One Identity Safeguard for Privileged Sessions 8.0 LTS

Administration Guide

Copyright 2025 One Identity LLC.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of One Identity LLC .

The information in this document is provided in connection with One Identity products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of One Identity LLC products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, ONE IDENTITY ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL ONE IDENTITY BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF ONE IDENTITY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. One Identity makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. One Identity does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

One Identity LLC.
Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (http://www.OneIdentity.com) for regional and international office information.

Patents

One Identity is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at http://www.OneIdentity.com/legal/patents.aspx.

Trademarks

One Identity and the One Identity logo are trademarks and registered trademarks of One Identity LLC. in the U.S.A. and other countries. For a complete list of One Identity trademarks, please visit our website at www.OneIdentity.com/legal/trademark-information.aspx. All other trademarks are the property of their respective owners.

Legend

WARNING: A WARNING icon highlights a potential risk of bodily injury or property damage, for which industry-standard safety precautions are advised. This icon is often associated with electrical hazards related to hardware.

CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

SPS Administration Guide
Updated - 05 November 2025, 13:45

For the most recent documents and product information, see Online product documentation.

Contents

Preface **19**

Introduction **20**

The major benefits of One Identity Safeguard for Privileged Sessions (SPS) 20

Application areas 21

The concepts of One Identity Safeguard for Privileged Sessions (SPS) **23**

The philosophy of One Identity Safeguard for Privileged Sessions (SPS) 23

Policies 25

Credential Stores 27

Plugin framework 28

Indexing 30

Supported protocols and client applications 31

- HTTP 32
- ICA 32
- MSSQL 32
- Remote Desktop Gateway Server Protocol (RDGSP) 33
- Remote Desktop Protocol (RDP) 33
- Secure Shell Protocol (SSH) 34
- Telnet 34
- VMware Horizon View 35
- Virtual Network Computing (VNC) 35

Modes of operation 35

- Transparent mode 36
- Single-interface transparent mode 36
- Non-transparent mode 38
- Inband destination selection 38

Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) 40

- Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using SSH 40
- Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using RDP 43
- Connecting to a server through One Identity Safeguard for Privileged Sessions 46

(SPS) using an RD Gateway
Archive and backup concepts 47
Configuration export 47
System backup 49
Connection backup 49
Connection archive 50
Support bundle 51
Debug logs 52
Connection logs 53
Core dump files 53
Maximizing the scope of auditing 54
IPv6 in One Identity Safeguard for Privileged Sessions (SPS) 57
SSH host keys 57
Authenticating clients using public-key authentication in SSH 58
The gateway authentication process 58
Four-eyes authorization 60
Network interfaces 61
High Availability support in One Identity Safeguard for Privileged Sessions (SPS) 61
Firmware and High Availability 62
Versions and releases of One Identity Safeguard for Privileged Sessions (SPS) 62
Accessing and configuring One Identity Safeguard for Privileged Sessions (SPS) 63
Cloud deployment considerations 65
AWS deployment 65
Azure deployment 67
Limitations 69
Prerequisites 71
High Availability and redundancy in Microsoft Azure 71
Redundancy 71
High Availability 72
The Welcome Wizard and the first login 73
The initial connection to One Identity Safeguard for Privileged Sessions (SPS) 73
Creating an alias IP address (Microsoft Windows) 74
Creating an alias IP address (Linux) 78
Modifying the IP address of One Identity Safeguard for Privileged Sessions (SPS) 79

Configuring One Identity Safeguard for Privileged Sessions (SPS) with the Welcome Wizard ... 80
Logging in to One Identity Safeguard for Privileged Sessions (SPS) and configuring the first connection ... 91
Basic settings ... 97
Supported web browsers ... 98
The structure of the web interface ... 100
Elements of the main workspace ... 104
Navigating on the SPS UI ... 106
Multiple users and locking ... 107
Preferences ... 107
Change password ... 109
Audit keystore ... 111
Adding the first private key to your audit keystore ... 111
Adding further private keys to your audit keystore ... 112
Unlocking your audit keystore ... 113
Deleting a private key from your audit keystore ... 114
Network settings ... 114
Routing table ... 114
IP forwarding ... 115
Naming ... 116
HTTPS proxy ... 116
Configuring user and administrator login addresses ... 117
Managing logical interfaces ... 119
Routing uncontrolled traffic between logical interfaces ... 122
Configuring the routing table ... 122
Configuring date and time ... 123
System logging, SNMP and e-mail alerts ... 125
Configuring system logging ... 125
Configuring e-mail alerts ... 127
Configuring SNMP alerts ... 129
Querying SPS status information using agents ... 132
Customize system logging in One Identity Safeguard for Privileged Sessions (SPS) 133
Configuring system monitoring on SPS ... 136
Configuring monitoring ... 136

Health monitoring ... 137
Preventing disk space fill-up ... 138
System related traps ... 139
Traffic related traps ... 142
Data and configuration backups ... 145
Creating a backup policy using Rsync over SSH ... 146
Creating a backup policy using SMB/CIFS ... 149
Creating a backup policy using NFS ... 153
Creating configuration backups ... 155
Creating data backups ... 156
Encrypting configuration backups with GPG ... 157
Archiving ... 158
Creating an archive policy using SMB/CIFS ... 159
Creating an archive policy using NFS ... 162
Archiving the collected data ... 165
Cleaning up audit data ... 166
Configuring cleanup policies ... 166
Running cleanup policies immediately ... 169
Using plugins ... 169
Uploading plugins ... 170
Verifying the integrity of a plugin ... 170
Forwarding data to third-party systems ... 172
Using the universal SIEM forwarder ... 173
Message types forwarded to SIEMs ... 175
Message format forwarded to SIEMs ... 176

User management and access control ... 324

Login settings ... 324
Protecting against brute-force attacks ... 325
Authentication banner ... 327
Web interface timeout ... 328
Managing One Identity Safeguard for Privileged Sessions (SPS) users locally ... 328
Creating local users in One Identity Safeguard for Privileged Sessions (SPS) ... 328
Deleting local users from One Identity Safeguard for Privileged Sessions (SPS) ... 330
Setting password policies for local users ... 330
Managing local user groups ... 333

Managing One Identity Safeguard for Privileged Sessions (SPS) users from an LDAP database 335
Adding a new Active Directory server 336
Adding a new POSIX LDAP server 340
Handling user names in User Principal Name (UPN) format 344
Authenticating users to a RADIUS server 347
Authenticating users with X.509 certificates 349
Authenticating users with SAML2 352
SAML2 login overview 352
SAML2 support in SPS 353
Identity Provider metadata 353
Service Provider metadata 354
User identifiers 355
Group membership 356
How to configure SAML2 login 356
Overview 356
Configure SPS as a SAML2 SP 356
Configure your IdP to trust SPS 357
Authenticating users with SAML2 login method 359
Managing user rights and user groups 361
Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface 363
Modifying group privileges 364
Finding specific user groups 365
Using user groups 367
Built-in user groups of One Identity Safeguard for Privileged Sessions (SPS) 368
ACL access and summary list 372
Creating rules for restricting access to search audit data 381
Displaying the privileges of users and user groups 385
Listing and searching configuration changes 388
Using the internal search interface 390
Filtering 392
Exporting the results 392
Customizing columns of the internal search interface 392
Managing One Identity Safeguard for Privileged Sessions (SPS) 394

Controlling One Identity Safeguard for Privileged Sessions (SPS): reboot, shutdown 394
Disabling controlled traffic 395
Disabling controlled traffic permanently 397
Managing One Identity Safeguard for Privileged Sessions (SPS) clusters 397
Cluster roles 398
Enabling cluster management 401
Creating a cluster 402
Joining to a cluster 404
Assigning roles to nodes in your cluster 408
Configuration synchronization across nodes in a cluster 411
Configuration synchronization and SSH keys 411
Using a configuration synchronization plugin 412
Monitoring the status of nodes in your cluster 415
Updating the IP address of a node in a cluster 417
Managing a cluster with configuration synchronization without central search 419
Managing a cluster with central search configuration and configuration synchronization 421
Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster 423
HA nodes configuration and management options 426
Adjusting the synchronization speed 427
Redundant heartbeat interfaces 428
Next-hop router monitoring 430
Upgrading One Identity Safeguard for Privileged Sessions (SPS) 432
Upgrade checklist 433
Upgrading One Identity Safeguard for Privileged Sessions (SPS) (single node) 434
Upgrading High Availability One Identity Safeguard for Privileged Sessions (SPS) nodes 435
Upgrading an SPS cluster 437
Troubleshooting 437
Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS) 438
Importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) 439
Managing the One Identity Safeguard for Privileged Sessions (SPS) license 440
Updating the SPS license 441

Accessing the One Identity Safeguard for Privileged Sessions (SPS) console 442
Using the console menu of One Identity Safeguard for Privileged Sessions (SPS) ... 442
Enabling SSH access to the One Identity Safeguard for Privileged Sessions (SPS) host 444
Changing the root password of One Identity Safeguard for Privileged Sessions (SPS) 446
Exporting and importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) using the console 447
Data migration from an SPS instance to another SPS instance 448
Sealed mode 449
Disabling sealed mode 449
Out-of-band management of One Identity Safeguard for Privileged Sessions (SPS) .. 450
Configuring the IPMI from the console 451
Configuring the IPMI from the BIOS 454
Managing the certificates used on One Identity Safeguard for Privileged Sessions (SPS) 459
Generating certificates for One Identity Safeguard for Privileged Sessions (SPS) 462
Uploading external certificates to One Identity Safeguard for Privileged Sessions (SPS) 463
Generating TSA certificate with Windows Certificate Authority on Windows Server 2016 or later 467

General connection settings 482

Configuring connections 482
Modifying the destination address 488
Configuring inband destination selection 490
Modifying the source address 493
Creating and editing channel policies 495
Real-time content monitoring with Content Policies 498
Creating a new content policy 499
Configuring time policies 503
Creating and editing user lists 504
Authenticating users to an LDAP server 505
Audit policies 512
Encrypting audit trails 512
Timestamping audit trails with built-in timestamping service 515
Timestamping audit trails with external timestamping service 517

Digitally signing audit trails ... 519
Verifying certificates with Certificate Authorities ... 521
Verifying certificates with Certificate Authorities using trust stores ... 523
Signing certificates on-the-fly ... 529
Creating an external Signing CA ... 531
Creating a Local User Database ... 538
Sharing SPS functions with SPP ... 539

HTTP-specific settings ... 541
Supported HTTP channel types ... 541
Limitations in handling HTTP connections ... 542
Authentication in HTTP and HTTPS ... 543
Creating a new HTTP authentication policy ... 543
Setting up HTTP connections ... 545
Setting up a transparent HTTP connection ... 545
Enabling One Identity Safeguard for Privileged Sessions (SPS) to act as an HTTP proxy ... 547
Enabling TLS encryption in HTTP ... 549
Configuring half-sided SSL encryption in HTTP ... 551
Session-handling in HTTP ... 552
Creating and editing protocol-level HTTP settings ... 553
Customizing HTTP error templates ... 557

ICA-specific settings ... 559
Setting up ICA connections ... 559
Supported ICA channel types ... 560
Creating and editing protocol-level ICA settings ... 562
One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment ... 563
Troubleshooting Citrix-related problems ... 567

MSSQL-specific settings ... 568
Setting up MSSQL connections ... 568
Limitations in handling MSSQL connections ... 569
Supported MSSQL channel types ... 569
Authentication in MSSQL ... 570
Creating a new MSSQL authentication policy ... 570
Creating and editing protocol-level MSSQL settings ... 572

Enabling TLS-encryption for MSSQL connections ... 574

RDP-specific settings ... 577

Kerberos ... 578

Supported RDP channel types ... 579

Creating and editing protocol-level RDP settings ... 582

Network Level Authentication (NLA) with One Identity Safeguard for Privileged Sessions (SPS) ... 588

Network Level Authentication (NLA) with domain membership ... 588

Using One Identity Safeguard for Privileged Sessions (SPS) across multiple domains ... 590

Verifying the certificate of the RDP server in encrypted connections ... 591

Enabling TLS-encryption for RDP connections ... 592

Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway ... 595

Configuring Remote Desktop clients for gateway authentication ... 598

Inband destination selection in RDP connections ... 602

Usernames in RDP connections ... 603

Saving login credentials for RDP on Windows ... 604

Configuring RemoteApps ... 605

Configuring the RemoteApp Launcher ... 607

Configuring SPS to enable exporting files from audit trails after RDP file transfer through clipboard or disk redirection ... 610

Configuring SPS to enable exporting sound from audit trails ... 611

Sharing RDP connection policies with SPP ... 612

Sharing RDP connection policies with SPS ... 613

Using credential injection in SPP-initiated RDP sessions ... 614

SSH-specific settings ... 616

Setting the SSH host keys of the connection ... 617

Setting the SSH host keys accepted on the server side ... 617

Setting the SSH host keys offered to the clients ... 618

Supported SSH channel types ... 619

Sharing SSH connection policies with SPP ... 623

Sharing SSH connection policies with SPS ... 624

Authentication Policies ... 625

Creating a new authentication policy ... 626

Client-side authentication settings ... 627

Local client-side authentication ... 629
Relayed authentication methods ... 630
Configuring your Kerberos environment ... 631
Kerberos authentication settings ... 632
Server host keys ... 634
Automatically adding the host keys of a server to One Identity Safeguard for Privileged Sessions (SPS) ... 635
Manually adding the host key of a server ... 635
Creating and editing protocol-level SSH settings ... 638
Supported encryption algorithms ... 641
Using Sudo with SPS ... 645
Setting up Sudo connections in SPS ... 645
Configuring Sudo ... 648
Telnet-specific settings ... 651
Enabling TLS-encryption for Telnet connections ... 652
Creating a new Telnet authentication policy ... 655
Extracting username from Telnet connections ... 658
Creating and editing protocol-level Telnet settings ... 658
Inband destination selection in Telnet connections ... 661
VMware Horizon View connections ... 662
One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a VMware environment ... 662
VNC-specific settings ... 664
Enabling TLS-encryption for VNC connections ... 665
Creating and editing protocol-level VNC settings ... 669
Indexing audit trails ... 672
Regenerate content stored in lucene indices ... 673
Configuring the internal indexer ... 675
Configuring external indexers ... 681
Prerequisites and limitations ... 682
Hardware requirements for the external indexer host ... 683
Configuring One Identity Safeguard for Privileged Sessions (SPS) to use external indexers ... 683
Installing the external indexer ... 684

Configuring the external indexer ... 685
Configuring a service pool ... 687
Uploading decryption keys to the external indexer ... 690
Configuring a hardware security module (HSM) or smart card to integrate with external indexer ... 691
Setting up and testing the environment ... 692
Encrypting a PKCS#11 PIN ... 693
Starting and restarting the external-indexer service when using a custom password for PKCS#11 PIN encryption ... 693
Configuring SoftHSM ... 694
Configuring AWS CloudHSM ... 696
Configuring a smart card ... 697
Customizing the indexing of HTTP traffic ... 698
Starting the external indexer ... 699
Disabling indexing on One Identity Safeguard for Privileged Sessions (SPS) ... 699
Managing the indexers ... 700
Upgrading the external indexer ... 700
Troubleshooting external indexers ... 701
Monitoring the status of the indexer services ... 702
Monitoring the status of the indexer services in classic view ... 704
HTTP indexer configuration format ... 706
HTTP indexer configuration options ... 706

Using the Sessions interface ... 710

Card view ... 711
Adding custom fields to the card view ... 711
Table view ... 713
Flow view ... 713
Assigning search privileges ... 717
Specifying time ranges ... 720
Using search queries ... 723
List of available search queries ... 724
Searching in the contents of audit trails ... 799
Audit trail downloads information on the Search interface ... 807
Displaying statistics on search results ... 808
Analyzing data using One Identity Safeguard for Privileged Analytics ... 810

The search and filter process ... 818
Viewing session details ... 821
Viewing session details for data recorded by SPS ... 822
Viewing session details for data recorded by SPP ... 827
Visualizing Frequent Item Sets on the FIS flow view ... 829
Viewing audited user details ... 835
Viewing session details associated with an audited user ... 837
Replaying audit trails in your browser ... 839
Using the browser to play video files ... 841
Streamable session recording playback with Safeguard Desktop Player started from the SPS UI ... 843
Viewing encrypted screenshots ... 846
Replaying encrypted audit trails in your browser ... 847
Following active sessions ... 848
Creating report subchapters ... 849
Creating search-based report subchapters from search results ... 849
Creating search-based report subchapters from scratch ... 850
Search interface changes between version 5.0 and 6.0 ... 852
Searching session data on a central node in a cluster ... 858

Advanced authentication and authorization techniques ... 860
Configuring usermapping policies ... 860
Configuring gateway authentication ... 862
Configuring out-of-band gateway authentication ... 863
Performing out-of-band gateway authentication on One Identity Safeguard for Privileged Sessions (SPS) ... 867
Performing inband gateway authentication in SSH and Telnet connections ... 868
Performing inband gateway authentication in RDP connections ... 868
Troubleshooting gateway authentication ... 869
Configuring four-eyes authorization ... 870
Configuring four-eyes authorization ... 870
Performing four-eyes authorization on One Identity Safeguard for Privileged Sessions (SPS) ... 872
Using credential stores for server-side authentication ... 873
Configuring local Credential Stores ... 874
Performing gateway authentication to RDP servers using local Credential Store and NLA ... 877

Configuring password-protected Credential Stores ... 878
Unlocking Credential Stores ... 882
Using a custom Credential Store plugin to authenticate on the target hosts ... 883
Integrating external authentication and authorization systems ... 884
How Authentication and Authorization plugins work ... 884
Using a custom Authentication and Authorization plugin to authenticate on the target hosts ... 886
Performing authentication with AA plugin in terminal connections ... 887
Performing authentication with AA plugin in Remote Desktop connections ... 888
Integrating ticketing systems ... 889
Performing authentication with ticketing integration in terminal connections ... 889
Performing authentication with ticketing integration in Remote Desktop connections ... 890
Creating a custom plugin ... 890
Plugin troubleshooting ... 891
Reports ... 892
Contents of the operational reports ... 893
Configuring custom reports ... 893
Creating report subchapters ... 896
Creating reports from audit trail content ... 896
Creating search-based report subchapters from search results ... 900
Creating search-based report subchapters from scratch ... 901
Creating PCI DSS reports ... 903
Contents of PCI DSS reports ... 904
Report output ... 908
Import and visualize audit data from SPS in the Power BI Desktop reporting application ... 909
The One Identity Safeguard for Privileged Sessions (SPS) REST API ... 910
One Identity Safeguard for Privileged Sessions (SPS) scenarios ... 911
Configuring public-key authentication on One Identity Safeguard for Privileged Sessions (SPS) ... 911
Configuring public-key authentication using local keys ... 911
Configuring public-key authentication using an LDAP server and a fixed key ... 912
Configuring public-key authentication using an LDAP server and generated keys ... 914
Organizing connections in non-transparent mode ... 915

Organizing connections based on port numbers ... 915
Organizing connections based on alias IP addresses ... 916
Using inband destination selection in SSH connections ... 916
Using inband destination selection with PuTTY ... 917
Using inband destination selection with OpenSSH ... 919
Using inband selection and nonstandard ports with PuTTY ... 919
Using inband selection and nonstandard ports with OpenSSH ... 921
Using inband destination selection and gateway authentication with PuTTY ... 922
Using inband destination selection and gateway authentication with OpenSSH ... 924
SSH usermapping and keymapping in AD with public key ... 925

Troubleshooting One Identity Safeguard for Privileged Sessions (SPS) ... 935
Network troubleshooting ... 935
Gathering data about system problems ... 936
Viewing logs on One Identity Safeguard for Privileged Sessions (SPS) ... 937
Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS) 938
Collecting logs and system information for error reporting ... 940
Collecting logs and system information of the boot process for error reporting ... 941
Support hotfixes ... 942
Status history and statistics ... 944
Connection statistics ... 946
Memory ... 947
Disk ... 948
CPU ... 949
Network connections ... 950
Interface ... 951
Load average ... 952
Number of processes ... 953
Displaying custom connection statistics ... 953
Troubleshooting a One Identity Safeguard for Privileged Sessions (SPS) cluster ... 954
Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses ... 954
Recovering One Identity Safeguard for Privileged Sessions (SPS) if both nodes broke down ... 957
Recovering from a split brain situation ... 957
Replacing a HA node in a One Identity Safeguard for Privileged Sessions (SPS) cluster ... 960

Resolving an IP conflict between cluster nodes ... 961
Understanding One Identity Safeguard for Privileged Sessions (SPS) RAID status ... 963
Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data ... 964
Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration ... 965
Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to a new SPS appliance ... 965
Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance ... 967
VNC is not working with TLS ... 968
Configuring the IPMI from the BIOS after losing IPMI password ... 968
Incomplete TSA response received ... 973

Using SPS with SPP ... 975
Configuring the Passwords-initiated workflow ... 978
Configuring SPP for Passwords-initiated workflow ... 980
Configuring the Sessions-initiated workflow ... 981
Configuring SPP for Sessions-initiated workflow ... 984
Configuring SPS for Sessions-initiated workflow ... 984
Configuring SPS for SRA-initiated workflow ... 986
Linking SPS to SPP ... 988
Switching seamlessly between SPS and SPP ... 991
Troubleshooting the SPS to SPP link ... 991
SPP to SPS link error resolution ... 991
SPP to SPS link issues ... 994

Configuring external devices ... 996
Configuring advanced routing on Linux ... 996
Configuring advanced routing on Cisco routers ... 998
Configuring advanced routing on Sophos UTM (formerly Astaro Security Gateway) firewalls ... 1002

Using SCP with agent-forwarding ... 1006

Security checklist for configuring One Identity Safeguard for Privileged Sessions (SPS) ... 1008
Encryption-related settings ... 1008
Connection policies ... 1009
Appliance access ... 1009

Networking considerations 1010

Jumplists for in-product help 1011

Basic Settings > Management 1011

Basic Settings > Local Services 1012

Basic Settings > System 1012

<Protocol name> Control > Global Options 1013

Configuring SPS to use an LDAP backend 1014

Overview 1015

Common to all backends 1016

POSIX LDAP backend 1016

Active Directory LDAP backend 1018

About us 1020

Contacting us 1020

Technical support resources 1020

Glossary 1021

1

Preface

Welcome to the One Identity Safeguard for Privileged Sessions 8.0.1 LTS Administrator Guide.

This document describes how to configure and manage the One Identity Safeguard for Privileged Sessions (SPS). Background information for the technology and concepts used by the product is also discussed.

2

Introduction

This section introduces One Identity Safeguard for Privileged Sessions (SPS) in a non-technical manner, discussing how and why is it useful, and what additional security it offers to an existing IT infrastructure.

The major benefits of One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) is part of the One Identity Safeguard solution, which in turn is part of One Identity's Privileged Access Management portfolio. Addressing large enterprise needs, SPS is a privileged session management solution which provides industry-leading access control, session recording and auditing to prevent privileged account misuse and accelerate forensics investigations.

SPS is a quickly deployable enterprise device, completely independent from clients and servers - integrating seamlessly into existing networks. It captures the activity data necessary for user profiling and enables full user session drill down for forensic investigations.

SPS has full control over the SSH, RDP, Telnet, TN3270, TN5250, Citrix ICA, and VNC connections, giving a framework (with solid boundaries) for the work of the administrators. The most notable features of SPS are the following:

Central policy enforcement

SPS acts as a centralized authentication and access-control point in your IT environment which protects against privileged identity theft and malicious insiders. The granular access management helps you to control who can access what and when on your critical IT assets.

Prevention of malicious activities

SPS monitors privileged user sessions in real-time and detects policy violations as they occur. In case of detecting a suspicious user activity (for example entering a destructive command, such as the "rm"), SPS can send you an alert or immediately terminate the connection.

Greater accountability (deterrance)

SPS audits "who did what", for example on your database- or SAP servers. Aware of this, your employees will do their work with a greater sense of responsibility leading to a reduction in human errors. By having an easily interpreted, tamper-proof record in encrypted, timestamped, and digitally signed audit trails, finger-pointing issues can be eliminated.

Faster, cost-effective compliance audits

SPS makes all user activity traceable by recording them in high quality, tamper-proof and easily searchable audit trails. All data is stored in encrypted, timestamped and signed files, preventing any modification or manipulation. The movie-like audit trails ensure that all the necessary information is accessible for ad-hoc analyses or audit reports.

Lower troubleshooting and forensics costs

When something wrong happens, everybody wants to know the real story. Analyzing thousands of text-based logs can be a nightmare and may require the participation of external experts. The ability to easily reconstruct user sessions allows you to shorten investigation time and avoid unexpected cost.

Application areas

Fastest return to value and extremely low TCO

One Identity Safeguard for Privileged Sessions (SPS) is a turnkey network appliance - its implementation and configuration is fast and simple. Compared to competitors, there is no need to purchase and install any additional software (for example, Windows or MS SQL servers) or hardware to have SPS fully functioning. Full implementation typically takes only 3-5 days. No need for long and costly professional services for implementation and customization. After deployment, SPS operates in the background like a black box of an airplane - there is no need for any extra workload to operate it.

Independent, agentless device

Compared to agent-based solutions, there is no need for installing and updating agents on clients or servers, eliminating unnecessary maintenance and potential security issues. As a host independent gateway, SPS can control and monitor access to any type of systems incl. all Windows/UNIX/Linux servers, mainframes, network devices, security devices, web-based applications or thin client environments, such as VMwareHorizon View (formerly known as VMware View), Citrix Virtual Apps (formerly known as Citrix XenApp) or Citrix Virtual Desktops (formerly known as Citrix XenDesktop).

Transparent, “router-like” operation

As a proxy gateway, SPS can operate as a router in the network – invisible to the user and to the server. As a transparent solution, SPS requires minimal changes to the existing network. Also, since it operates on the network level, users can keep using the client applications they are familiar with, and do not have to change their work processes, unlike jump host solutions.

Granular access control

Since SPS has full access to the inspected traffic, security managers can granularly control who can access what and when on the servers. For example, they can selectively permit or deny access to protocol channels: enable terminal sessions in SSH, but disable port-forwarding and file transfers, or enable desktop access for RDP, but disable file sharing. In addition, SPS supports real-time shadowing allowing an authorizer to follow the administrator's session in real-time and terminate his/her connection in case of detecting a policy violation.

Real-time prevention of malicious activities

SPS can monitor transferred content in real time and can send alerts or even block connections if a certain pattern is detected in the traffic. Predefined patterns can be a risky command in a text-oriented protocol or a suspicious application in a graphical connection. This command and application level policy can prevent malicious user activities as they happen instead of just recording or reporting them.

Industry-leading session recording and auditing

SPS is the leading session auditing solution on the market offering Optical Character Recognition (OCR) capabilities to log ALL data about privileged actions in graphical user interfaces as well as text-based protocols. SPS can support and audit file transfers, as well. All data is recorded into searchable movie-like audit trails, making it easy to find relevant information in forensics or troubleshooting situations. In case of any problems (server misconfiguration, database manipulation, unexpected shutdown), the circumstances of the event are readily available in the audit trails, thus the cause of the incident can be easily identified. Auditors can do free-text searches in the content of text-based and graphical sessions. They can search for EVERY events (for example, mouse clicks, pressing Enter) and texts seen by the user.

To protect the sensitive information included in the communication, the two directions of the traffic (client-server and server-client) can be separated and encrypted with different keys, thus sensitive information like passwords are displayed only when necessary.

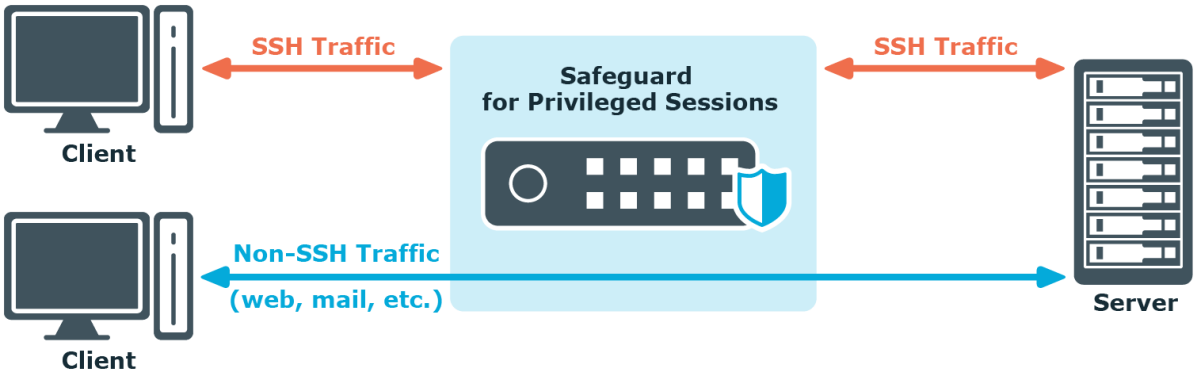
The concepts of One Identity Safeguard for Privileged Sessions (SPS)

This section discusses the technical concepts of One Identity Safeguard for Privileged Sessions (SPS).

The philosophy of One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) is a device that examines network traffic at the application level, that is, Layer 7 or the application layer of the OSI model. All communication must conform to the standards of the respective protocol. SPS examines Secure Shell (SSH, including forwarded X11 traffic), Secure Copy (SCP), SSH File Transfer Protocol (SFTP), Remote Desktop (RDP), HTTP, Independent Computing Architecture (Citrix ICA), Telnet, VMware Horizon View, and VNC connections, ignoring and simply forwarding all other types of traffic. SPS uses man-in-the-middle techniques to decrypt and terminate (when necessary) the inspected connections. It separates the connections into two parts (client — SPS, SPS — server) and inspects all traffic, so that no data can be directly transferred between the server and the client.

Figure 1: Inspecting SSH traffic with SPS



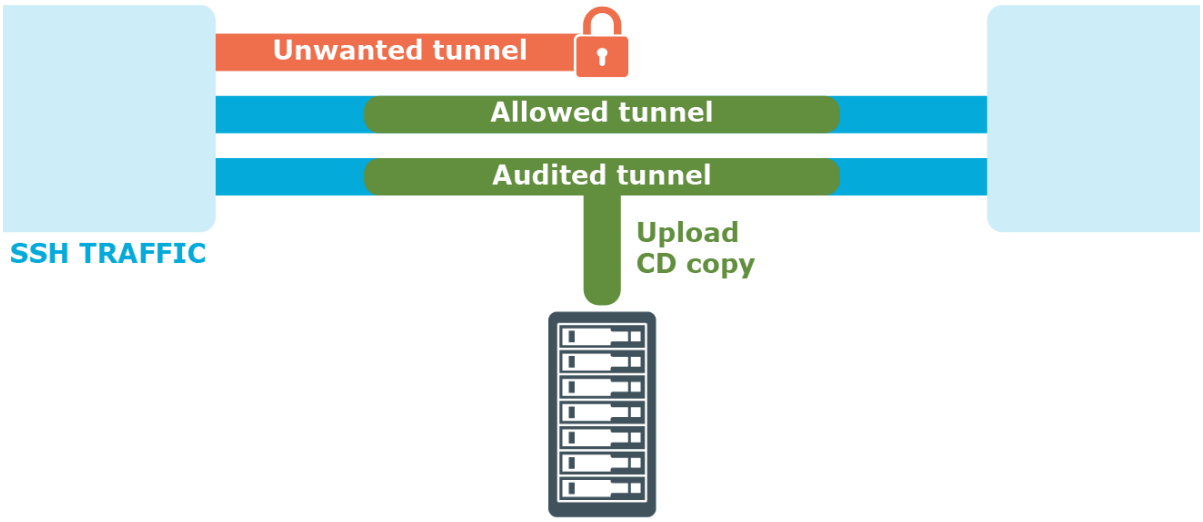
SPS has full control over the initial negotiation phase of the connection, when the client and the server decide the parameters of the encryption to be used in the communication. SPS can restrict the use of the various algorithms, forbidding the use of weak ones — an effective shield against downgrade attacks.

Since SPS isolates the client-server connection into two separate connections, the permitted algorithms can be different on the client and the server side.

SPS controls the connections right from the beginning — including user authentication. That way it is easy to mandate strong authentication for protocols where user information is available (for example, SSH), because SPS can limit the allowed authentication methods and also the users permitted to access the servers.

SPS uses various policies to restrict who, when, and how can access a connection or a specific channel of the protocol. These policies (based on username, authentication method used, and so on) can be applied to connections between particular clients and servers, or also to specific channels of a connection (for example, only to terminal-sessions in SSH, or desktop-sharing in RDP).

Figure 2: Controlling protocol channels

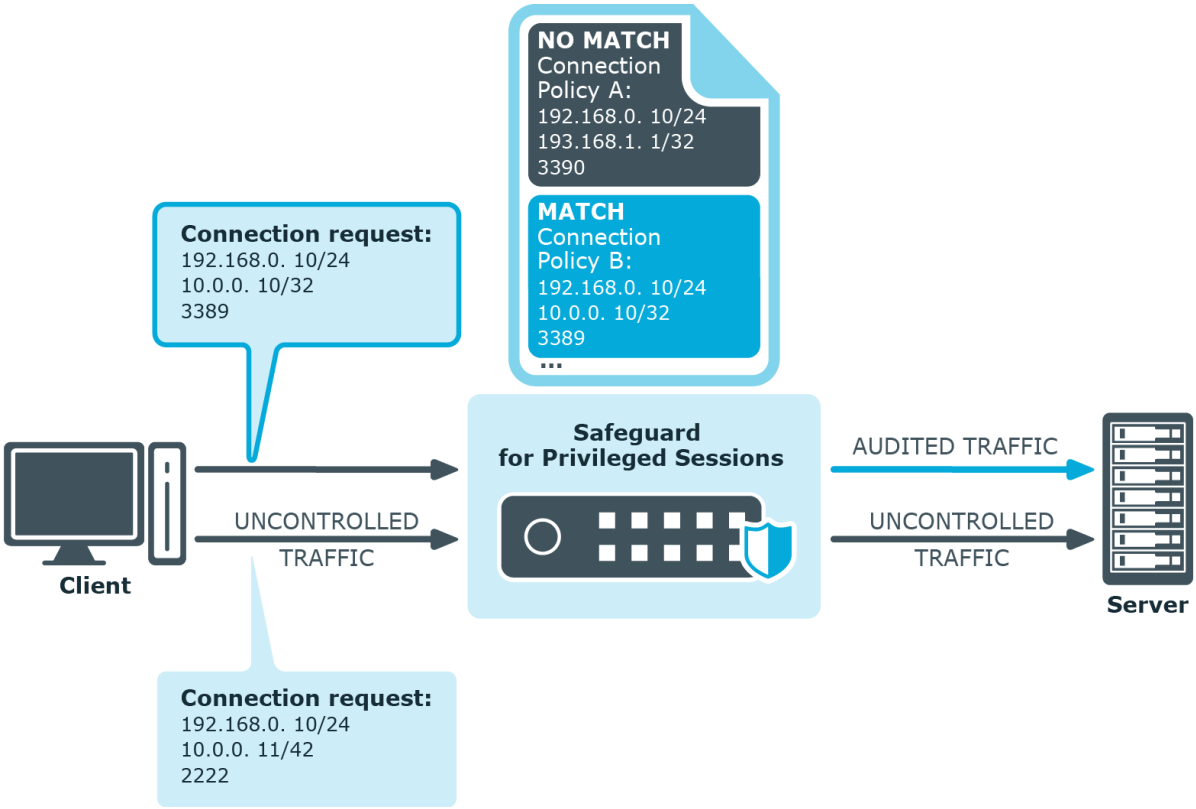


SPS is configured by an administrator or auditor using a web browser.

Policies

One Identity Safeguard for Privileged Sessions (SPS) controls access to connections through a set of policies. Policies let you specify various parameters of a connection, and so define the types of connections that SPS should monitor and restrict access to. When a connection request reaches SPS, SPS compares the connection policies to the parameters of the connection request one-by-one, starting with the first policy in the policy list. The first connection policy completely matching the connection request is applied to the connection.

Figure 3: Processing policies



This section provides a brief definition of each policy type and also explains the hierarchy between them.

A *connection policy* allows you to specify details of the connection between a particular client and server that you want to restrict in any way. In addition to setting basic details (such as the source and destination addresses), or more advanced ones (such as authentication to SPS or the server), the connection policy also references other policies that allow you to define further specifics of the connections you wish to control.

Depending on the protocol, the connection policy may also allow you to configure:

- A Credential Store that allows users auto logon to the target server.

  For information on Credential Stores, see Credential Stores on page 27.
- A plugin that allows integration with external systems, which users can be optionally authenticated to (before authenticating to the target server).

  For information on plugins, see Plugin framework on page 28.

For details on configuring connection policies, see Configuring connections on page 482.

A *channel policy* serves to control channel usage (for example, terminal session and Secure Copy in SSH, or drawing and clipboard in RDP) within a given connection. It lists channels that are allowed within a connection, and it also lets you specify restriction rules based on user lists, user groups, or the IP address of the client or server. You can also reference a content policy and a time policy within the channel policy, and it is also within the channel policy that you enable auditing for a specific channel.

For details on configuring channel policies, see Creating and editing channel policies on page 495.

A *content policy* lets you log an event, send an alert, or terminate a connection if a particular command or text (that you specify in the policy) appears in the command line or on the screen.

For details on creating a content policy, see Creating a new content policy on page 499.

A *time policy* specifies the timeframe when users are permitted to access a particular channel and so restricts the availability of that channel.

For details on configuring time policies, see Configuring time policies on page 503.

An *audit policy* enables you to prevent the manipulation of audit trails files that store the recorded activities of privileged users by providing you with options to encrypt, timestamp, and sign these files.

For details on creating audit policies, see Audit policies on page 512.

An *authentication policy* defines those client-side and server-side authentication methods that can be used in a connection.

For details on creating authentication policies, see Authentication Policies on page 625.

An *LDAP policy* lets you set details of the LDAP server to which you wish to authenticate users of the connections you are controlling.

For details on creating an LDAP policy, see Authenticating users to an LDAP server on page 505.

A *usermapping policy* specifies the usernames that are allowed access to the remote server and the user groups that are allowed to use the specified username.

For details on configuring usermapping policies, see Configuring usermapping policies on page 860.

An *archiving policy* lets you configure details of the archiving process that enables you to archive connection-related data and audit trails. You can configure, for example, the target server where archived files are to be stored, or the directory structure in which to organize your archived files.

For details on creating archiving policies, see Archiving on page 158.

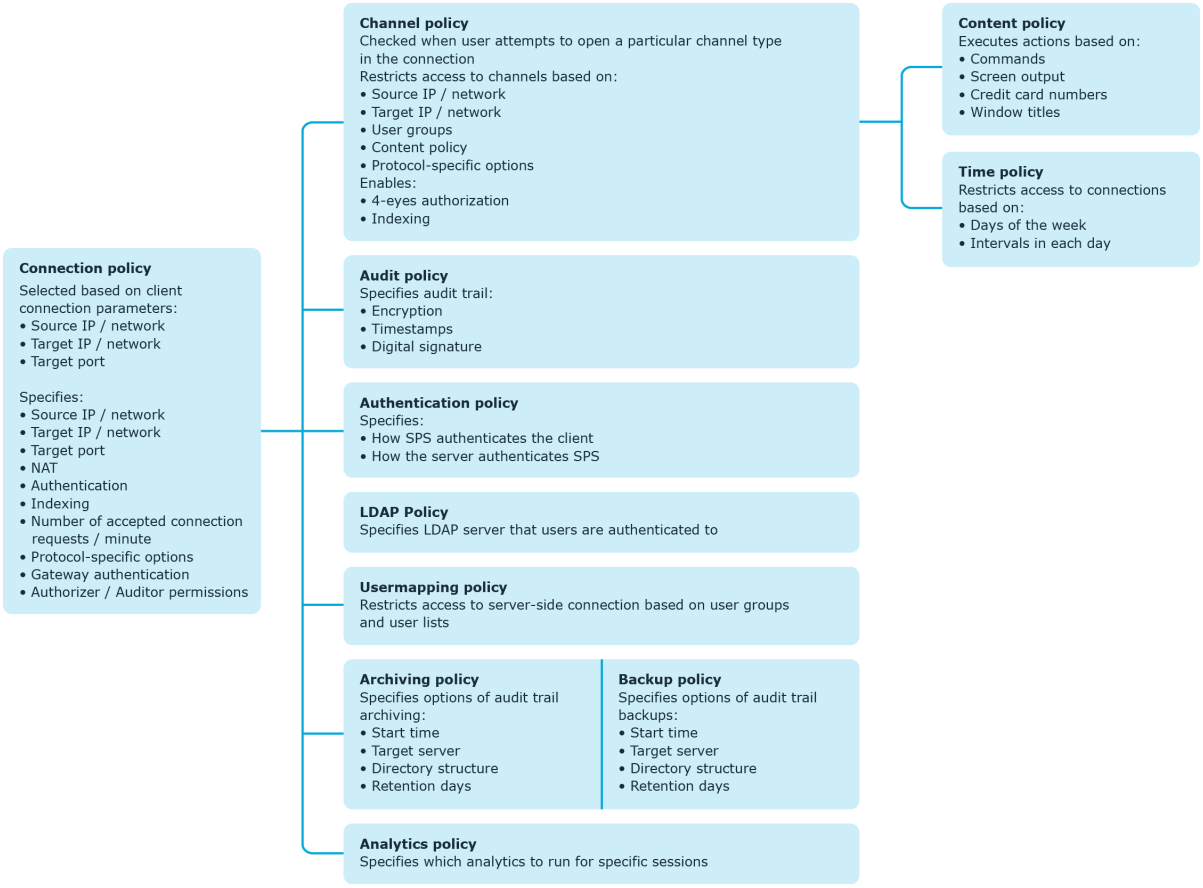
A *backup policy* defines the address of the backup server where you can back up connection data, the protocol to use to access it, details of authenticating to the backup server, and so on.

For details on creating backup policies, see Data and configuration backups on page 145.

An *analytics policy* lets you specify the analytics that you wish to run for specific sessions, and also determine the weight that scores given by the selected analytics should have in the final aggregated score.

For details on configuring analytics policies, see *Configure analytics* in the *Safeguard for Privileged Analytics Configuration Guide*.

Figure 4: Policies of SPS



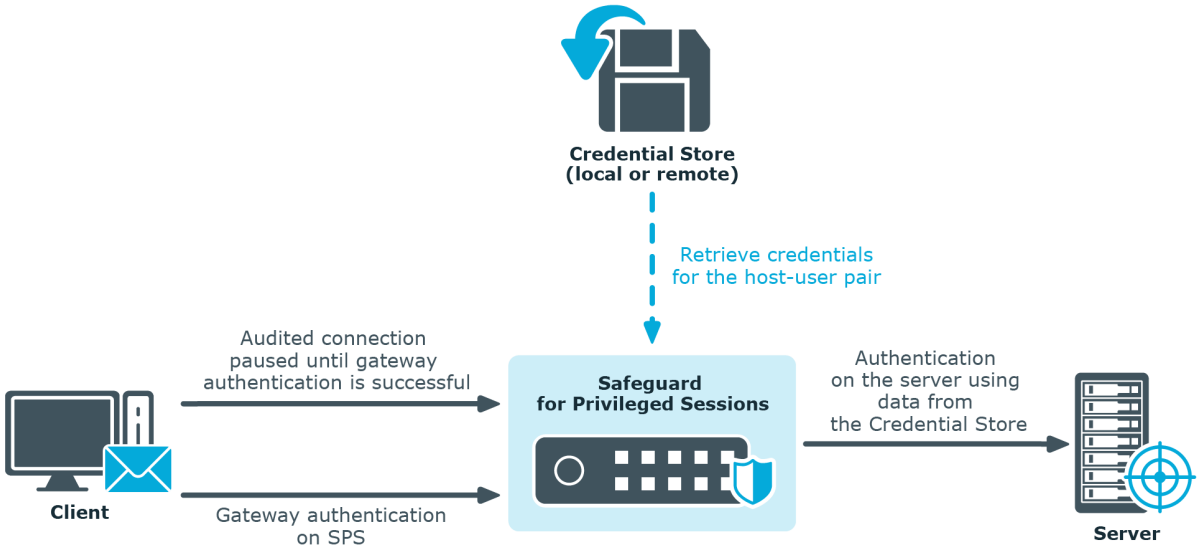
Credential Stores

Credential Stores are repositories of user credentials (for example, passwords, private keys, certificates). They are used for authenticating a user to the target server that the user wishes to access, without the user actually having access to those credentials. Credentials are retrieved transparently from One Identity Safeguard for Privileged Sessions's (SPS's) local Credential Store or an external, third-party password management system by SPS impersonating the authenticated user. This automatic password retrieval is

crucial, as this method protects the confidentiality of passwords since users can never access them.

Users accessing connections that use Credential Stores must authenticate on SPS using gateway authentication. They only have to use their gateway password to log in to SPS, and if they are allowed to access the target server, SPS automatically logs in using the Credential Store. For details on gateway authentication, see The gateway authentication process on page 58.

Figure 5: Authenticating using Credential Stores



Credential Stores can be stored locally on SPS, or on a remote device. For remote Credential Stores, SPS integrates with external authentication and authorization systems using plugins.

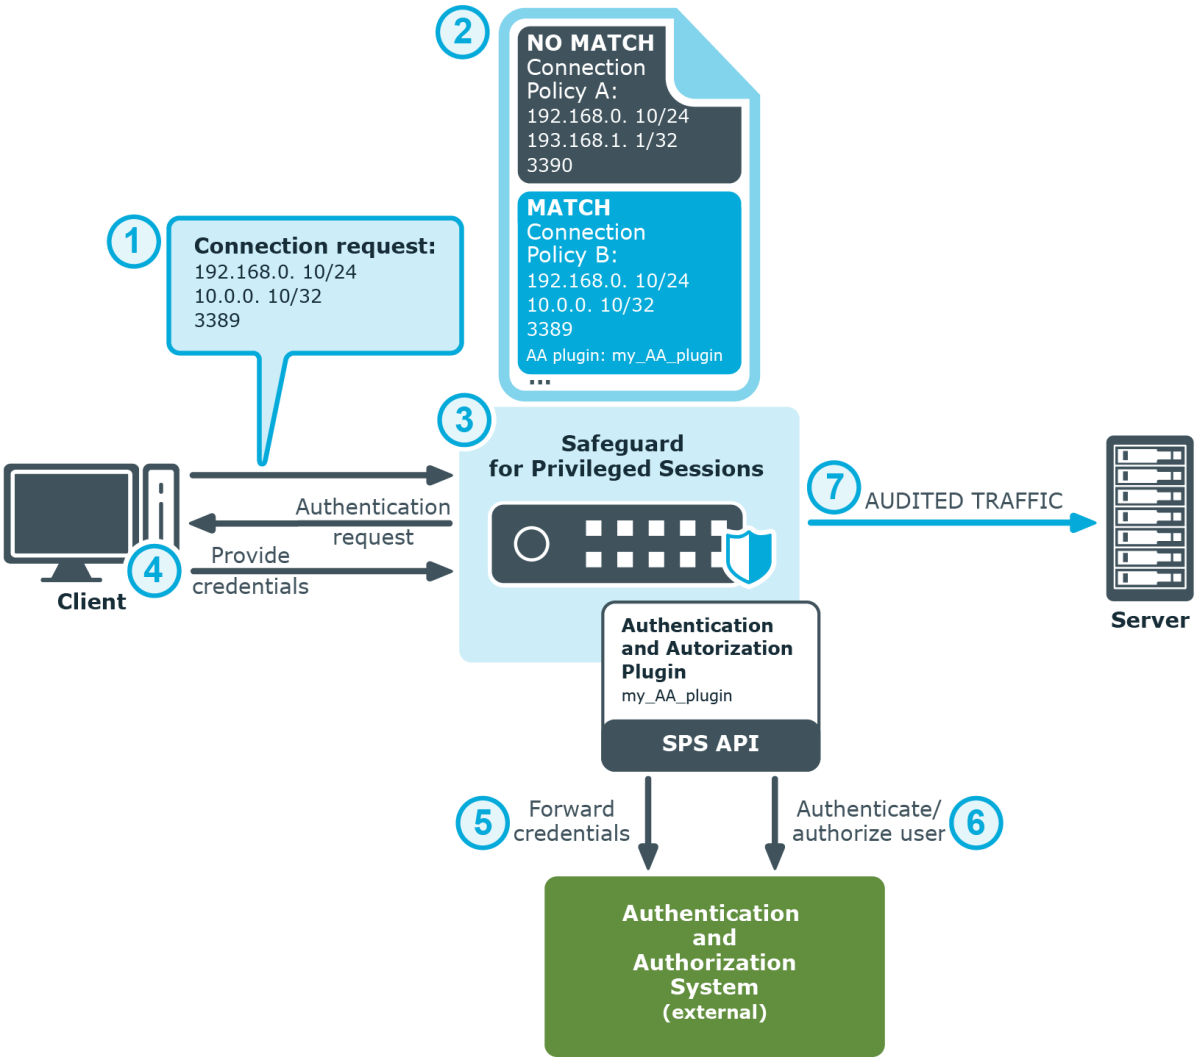
For further information on Credential Stores including configuration details, see Using credential stores for server-side authentication on page 873.

Plugin framework

One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS with external authentication and authorization systems, such as an external Credential Store, a ticketing system, or any third-party authentication or authorization solution.

Authenticating users to an external authentication and authorization system on page 29 and the process overview that follows describe how user authentication works at a high level when there is an external authentication and authorization system involved:

Figure 6: Authenticating users to an external authentication and authorization system



1. The client tries to establish a connection to the target server.
2. SPS notices that an AA plugin is configured in the connection policy matching the connection. This is treated as gateway authentication. For details on gateway authentication, see The gateway authentication process on page 58.
3. SPS prompts the client for credentials.
4. The client provides authentication details to SPS when prompted.
5. SPS forwards the client's details to the external authentication and authorization system using the SPS API.
6. The external authentication and authorization system verifies the data received and provides feedback to SPS about the result.
7. If the client is granted access by the external authentication and authorization

system, SPS authenticates the client to the target server, and establishes the connection.

For further information on plugins including configuration details, see Integrating ticketing systems on page 889 and Integrating external authentication and authorization systems on page 884.

Indexing

One Identity Safeguard for Privileged Sessions (SPS) can index the contents of audit trails, making the records of privileged users' activities easily searchable.

Audit trails contain user activity data recorded from terminal sessions (such as SSH and Telnet) and graphical protocols (such as RDP, Citrix ICA, and VNC). Examples of data recorded in audit trails are: mouse activity, keystrokes, and so on. Using its own indexer service or one or more external indexers, SPS determines elements of the content visible on the user's screen at a given point in time. Screen content elements include commands, window titles, IP addresses, user names, and so on.

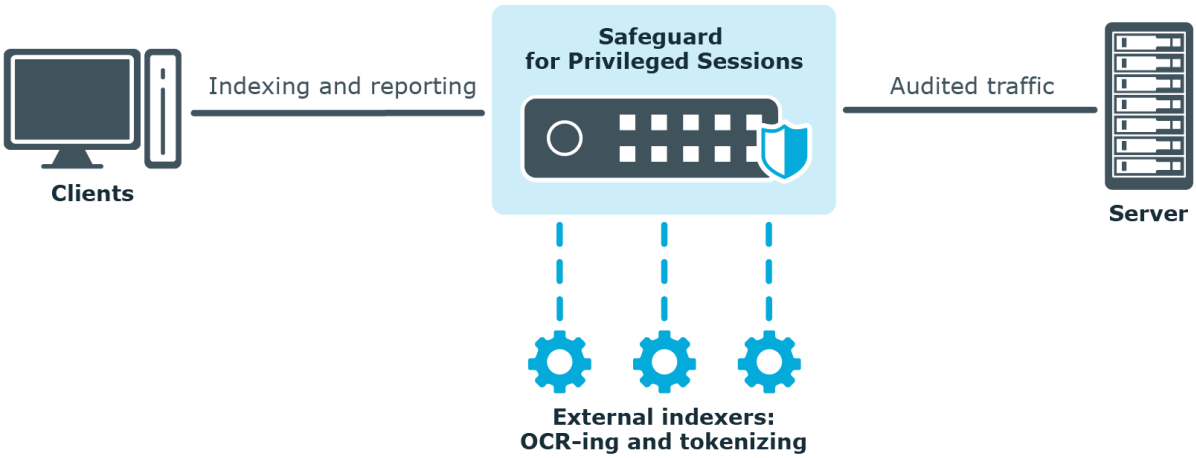
The indexer generates the following types of output as a result of processing the audit trail files:

- text
- screenshot files
- replayable video files

SPS then takes the output of indexing and breaks that down into searchable units.

Indexing and the process overview that follows describe how indexing works at a high level:

Figure 7: Indexing audit trail files



1. SPS monitors and records the protocol traffic in the audited connections passing through SPS. Protocol traffic data is recorded in audit trail files.
2. Once a connection has been closed, SPS sends the audit trail files to the indexer.
3. The indexer parses the contents of the audit trail files, and builds an "inventory" of the privileged user's activity data based on what appeared on their screen.

   In the case of a terminal session, screen content corresponds to the activity data that is captured in a terminal window. In the case of graphical protocols, screen content is whatever is visible in the graphical user interface of the applications the user is interacting with. In the latter case, the indexer's Optical Character Recognition (OCR) engine extracts text that appeared on the screen (for example, window titles).
4. The indexer returns the information extracted from the parsed audit trail files to SPS.
5. SPS processes the outcome of parsing and OCR-ing done in the previous phase and makes the data searchable.
6. Once indexed, the contents of the audit trails can be searched from SPS's web interface.

For details on how to configure SPS's internal indexer or one or more external indexers, see Indexing audit trails on page 672.

Supported protocols and client applications

One Identity Safeguard for Privileged Sessions (SPS) supports the following protocols and clients. As a general rule, client applications not specifically tested, but conforming to the relevant protocol standards, should work with SPS.

As a general rule, One Identity supports the listed client and server applications until their manufacturer provides mainstream support for them.

One Identity supports the listed client and server applications only on a best-effort basis after their vendor or manufacturer declares end-of-support or extended (or any other non-standard support) period for them. Best-effort basis means that without the vendor support we can only fix issues with our existing knowledge in the problematic area, and can only implement straightforward fixes.

Example

Microsoft provides mainstream and extended support periods for Windows Server 2019 Standard as described here. One Identity follows these periods and our best-effort support period starts at the same time when the mainstream period ends at Microsoft. The mainstream support for Windows Server 2019 will end on 09 January 2024 and after that, One Identity will support Windows Server 2019 on a best-effort basis.

HTTP

One Identity Safeguard for Privileged Sessions (SPS) supports the HTTP 1.0 and 1.1 standards.

ICA

One Identity Safeguard for Privileged Sessions (SPS) is certified for the following server versions:

- Citrix Virtual Apps (formerly known as Citrix XenApp) 6.5
- Citrix Virtual Apps 7.6
- Citrix Virtual Apps 7.15
- Citrix Virtual Apps 19.12
- Citrix Virtual Desktops (formerly known as Citrix XenDesktop) 6.5
- Citrix Virtual Desktops 7.6
- Citrix Virtual Desktops 7.15
- Citrix Virtual Desktops 19.12

For details on the deployment scenarios that support Citrix Virtual Desktops (formerly known as Citrix XenDesktop), see One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment on page 563.

The latest version of the Citrix Workspace app (formerly known as Citrix Receiver) for Windows, Linux and MacOS is supported.

SPS supports SecureICA using RC5 encryption. However, ICA with TLS basic encryption (non-RC5 algorithm) is not supported.

MSSQL

One Identity Safeguard for Privileged Sessions (SPS) supports the underlying TDS protocol from version 7.3 to 7.4. Due to the TDS protocol version requirement, Microsoft SQL Server 2008 or later is recommended.

Supported client and server applications

- MSSQL 2017 (client and server)
- MSSQL 2019 (client and server)
- Azure SQL

Remote Desktop Gateway Server Protocol (RDGSP)

One Identity Safeguard for Privileged Sessions (SPS) can act as a Remote Desktop Gateway (also called RD Gateway) and transfer the incoming connections to RDP connections.

Remote Desktop Protocol (RDP)

As a general rule, One Identity supports the listed client and server applications until their manufacturer provides mainstream support for them.

After the end date of the mainstream support, One Identity supports them on a best-effort basis.

Supported Windows client applications

As a general rule, One Identity supports the listed Microsoft client applications until Microsoft provides mainstream support for them.

After the end date of the mainstream support, One Identity supports them on a best-effort basis.

The built-in applications of the Windows Server 2016, Windows Server 2019, Windows Server 2022, Windows 10, and Windows 11 platforms are supported.

CAUTION: If you are using SHA1 (Secure Hash Algorithm 1) signed certificates, SPS does not allow Remote Desktop Protocol (RDP) connections to Windows Servers.

Use the Microsoft Management Console (MMC) to verify your certificate:

- If Remote Desktop Services (RDS) uses a self-signed certificate, make sure that you update your system to the latest patch level, then delete the certificate and restart the Remote Desktop Configuration service in order to re-generate the self-signed certificate.
- If RDS is using a certificate imported from a Public Key Infrastructure (PKI), contact your PKI admin for a new SHA256 certificate.

Supported Mac OS X client applications

- The latest released version of Royal TSX client application.
- The latest released version of Microsoft Remote Desktop client application.

NOTE: The Remote Desktop Connection Client for Mac application does not support RDP shadowing.

Other client applications

Other client applications are not explicitly supported, but may be compatible with One Identity Safeguard for Privileged Sessions (SPS).

Supported server (target) applications

As a general rule, One Identity supports the listed Microsoft server applications until Microsoft provides mainstream support for them.

The built-in applications of the Windows Server 2016, Windows Server 2019, Windows Server 2022, Windows 10, and Windows 11 platforms are supported.

After the end date of the mainstream support, One Identity supports them on a best-effort basis.

Accessing Remote Desktop Services (RemoteApp programs) is also supported.

Secure Shell Protocol (SSH)

One Identity Safeguard for Privileged Sessions (SPS) supports only the SSHv2 protocol. The older and insecure v1 version is not supported.

Supported client and server applications

- OpenSSH (client and server)

  Client and server tested with version OpenSSH_8.2 and OpenSSH_9.0.
- Dropbear (server)

  Tested with version v2019.78.
- SecureCRT (Windows, client)

  Tested with version 8.5.
- PUTTY (client)

  Tested with version 0.65.

Telnet

Telnet traffic must conform to RFC 854 and to various extensions described in the following RFCs: 856-861, 652-658, 698, 726-27, 732-736, 749, 779, 885, 927, 933, 1041, 1043, 1053, 1073, 1079, 1091, 1096-97, 1184, 1372, 1408, 1572, 2066, 2217, 2840, 2941, 2946.

TN3270

Telnet 3270 terminal protocol.

TN5250

Telnet 5250 terminal protocol, as described in RFC2877.

VMware Horizon View

VMware Horizon View Clients using the Remote Desktop (RDP) display protocol to access remote servers are supported. For details, see VMware Horizon View connections on page 662.

Virtual Network Computing (VNC)

One Identity Safeguard for Privileged Sessions (SPS) supports the Remote Framebuffer (RFB) protocol which is used in various open source VNC implementations. RFB versions 3.3-3.8 are supported.

The following client and server applications are supported if they are built on the open source RFB protocol:

- RealVNC
- UltraVNC
- TightVNC
- KVM
- Vino

Modes of operation

One Identity Safeguard for Privileged Sessions (SPS) can be configured to monitor both *transparent* and *non-transparent* connections.

- In *transparent* mode, SPS acts as a transparent router between two network segments. For details, see Transparent mode on page 36.
- You can also use policy-based routing to forward connections within the same network segment to SPS, in which case it acts like a *single interface transparent router*. For details, see Single-interface transparent mode on page 36.
- In *non-transparent* mode, users have to address SPS to initiate connections to protected servers. For details, see Non-transparent mode on page 38.
- When addressing SPS, you can also use *inband destination selection* to choose the server to connect to. For details, see Inband destination selection on page 38.

One Identity recommends that you design the network topology so that only management and server administration traffic passes SPS. This ensures that the services and applications running on the servers are accessible even in case SPS breaks down, so SPS cannot become a single point of failure.

Transparent mode

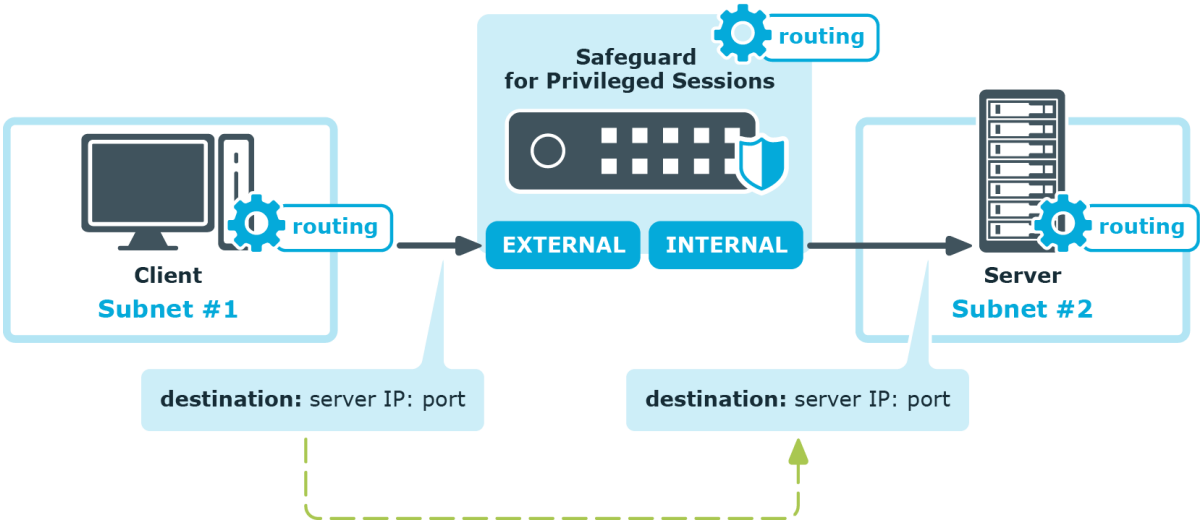
In *transparent* mode, One Identity Safeguard for Privileged Sessions (SPS) acts as a transparent router connecting the network segment of the administrators to the segment of the protected servers at the network layer (Layer 3 in the OSI model). All connections must pass through SPS to reach the servers — SPS is a proxy gateway, completely separating the protected servers from the rest of the network. Controlled connections and traffic are inspected on the application level, while other types of connections are simply forwarded on the packet level.

SPS can also be configured to act as a *single-interface transparent router*. For details, see Single-interface transparent mode on page 36.

CAUTION:

Transparent mode does not support multicast traffic.

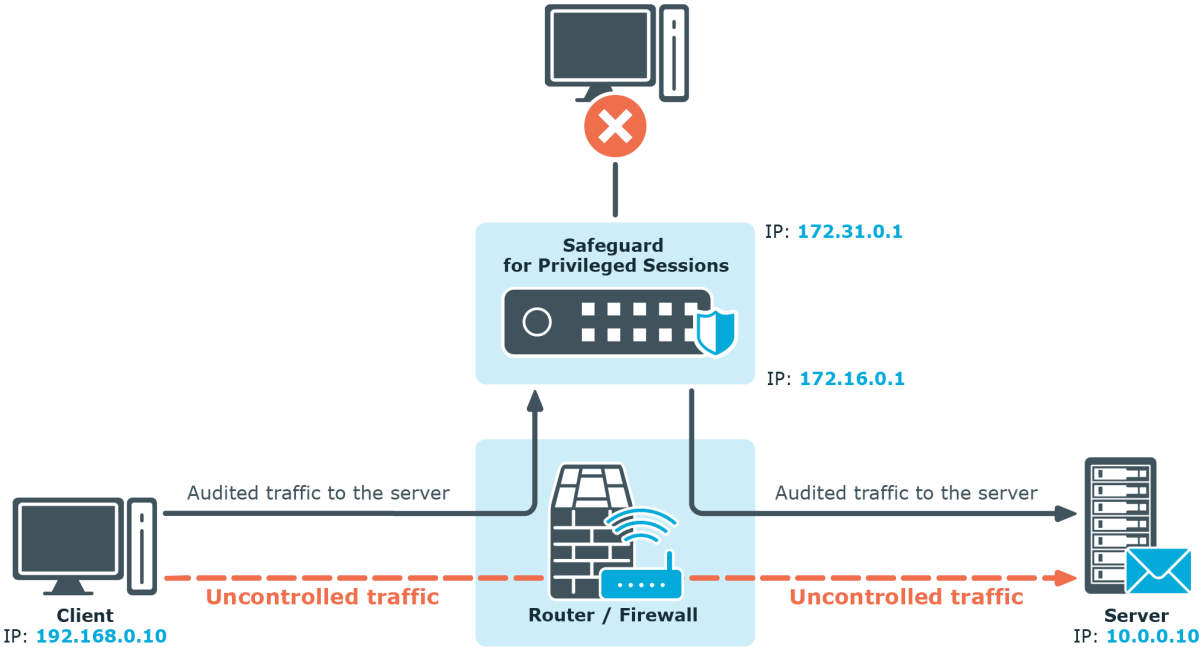
Figure 8: SPS in transparent mode



Single-interface transparent mode

Single-interface transparent mode is similar to transparent mode, but both client-side and server-side traffic use the same interface. An external device that actively redirects the audited traffic to One Identity Safeguard for Privileged Sessions (SPS) (typically, a firewall, a router, or a layer3 switch) is required . To accomplish this, the external device must support advanced routing (also called policy-based routing or PBR). For details on configuring an external device to work with SPS in single-interface transparent mode, see *Configuring external devices* in the *Administration Guide*.

Figure 9: SPS in single-interface transparent mode



Advantages

The advantages of using the single-interface transparent mode are:

- Totally transparent for the clients, no need to modify their configuration.
- The network topology is not changed.
- Only the audited traffic is routed to SPS, production traffic is not.

Disadvantages

The disadvantages of using the single-interface transparent mode are:

- SPS acts as a man-in-the-middle regarding the connection between the client and the target server. Instead of a single client-server connection, there are two separate connections: the first between the client and SPS, and a second between SPS and the server. Depending on how you configure SPS, the source IP in the SPS-server connection can be the IP address of SPS, or the IP address of the client. In the latter case — when operating in transparent mode (including single-interface transparent mode) — SPS performs IP spoofing. Consult the security policy of your organization to see if it permits IP spoofing on your network.
- Traffic must be actively routed to SPS using an external device. Consequently, a network administrator can disable SPS by changing routing rules.
- When adding a new port or subnet to the list of audited connections, the configuration of the external device must be modified as well.
- A network administrator can (intentionally or unintentionally) easily disable

monitoring of the servers, therefore additional measures have to be applied to detect such activities.

Non-transparent mode

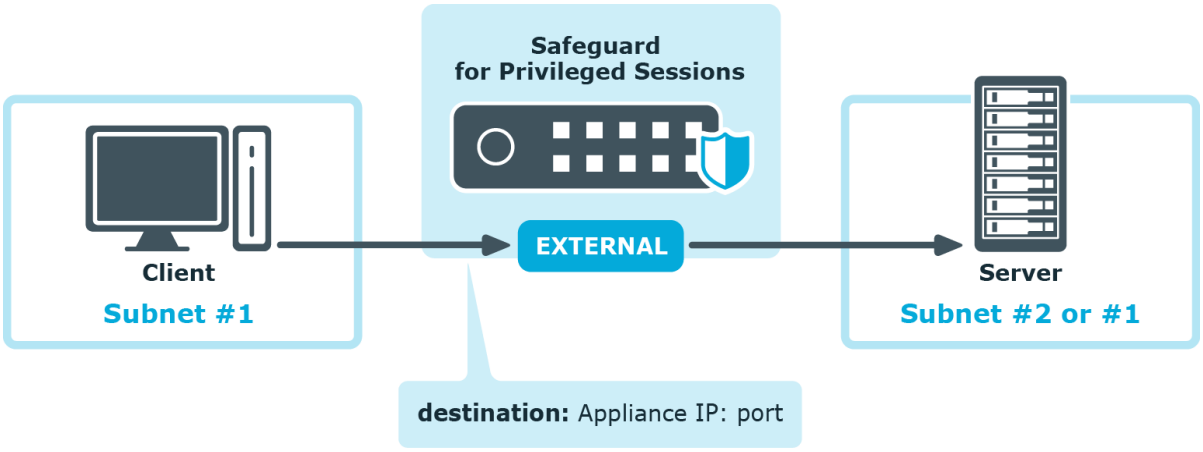
In *non-transparent* mode, One Identity Safeguard for Privileged Sessions (SPS) acts as a bastion host (that is, administrators can address only SPS, the administered servers cannot be targeted directly). The firewall of the network has to be configured to ensure that only connections originating from SPS can access the servers. SPS determines which server to connect to based on the parameters of the incoming connection (the IP address of the administrator and the target IP and port).

Non-transparent mode inherently ensures that only the controlled (management and server administration) traffic reaches SPS. Services and applications running on the servers are accessible even in case SPS breaks down, so SPS cannot become a single point of failure.

TIP: Non-transparent mode is useful if the general (that is, not inspected) traffic is very high and could not be forwarded by SPS.

NOTE: In case there is a high number of target devices, do not use fixed address rules in non-transparent mode, as configuration validation might fail. Consider using one of the dynamic configuration options, such as inband destination selection or transparent mode.

Figure 10: SPS in non-transparent mode



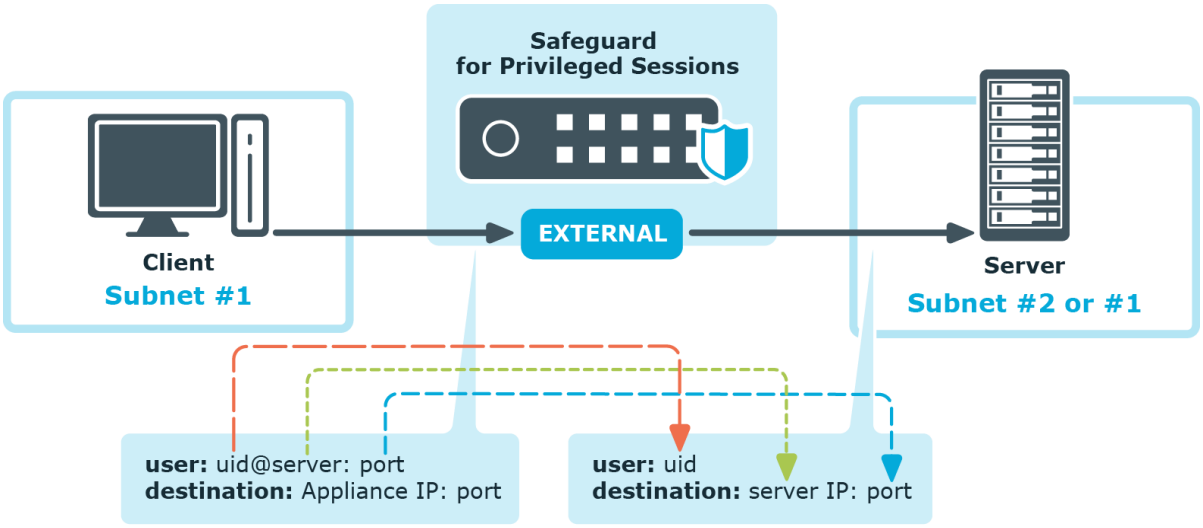
Non-transparent mode is often used together with inband destination selection. For details, see Inband destination selection on page 38).

Inband destination selection

Inband destination selection allows you to create a single connection policy and allow users to access any server by including the name of the target server in their username (for example, `ssh username@targetserver:port@scb_address`). One Identity Safeguard for

Privileged Sessions (SPS) can extract the address from the username and direct the connection to the target server.

Figure 11: Inband destination selection



Since some client applications do not permit the `@` and `:` characters in the username, alternative characters can be used as well:

- To separate the username and the target server, use the `@` or `%` characters, for example: **`username%targetserver@scb_address`**
- To separate the target server and the port number, use the `:`, `+`, or `/` characters, for example: `username%targetserver+port@scb_address`
- If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

In RDP, do not use the `@` character as an inband data separator but use alternative characters, for example, the `%` character.

You can use both IPv4 and IPv6 addresses with inband destination selection. For IPv6 addresses, add square brackets to separate the address and the port number:

```
username@[targetserver_ipv6]:port@[scb_address_ipv6]:port
```

When Network Level Authentication (NLA) is disabled, you can omit the username when starting an RDP connection (for example, use only `%targetserver`). The user can type the username later in the graphical login screen. However, the username must be specified if Network Level Authentication (NLA) is used in the connection.

For other details on inband destination selection in RDP connections, see Inband destination selection in RDP connections on page 602.

You can find examples of using inband destination selection in Using inband destination selection in SSH connections on page 916.

Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS)

When a client initiates a connection to a server, One Identity Safeguard for Privileged Sessions (SPS) performs a procedure similar to the ones detailed below. The exact procedure depends on the protocol used in the connection.

- For SSH connections, see Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using SSH on page 40.
- For RDP and other connections, see Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using RDP on page 43.

Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using SSH

The following describes what happens when a client connects to a server through One Identity Safeguard for Privileged Sessions (SPS) and how the different configuration options and policies of SPS affect this process. Note that this procedure does not cover the scenarios when inband destination selection is used.

1. **Client-side connection**

   The client tries to connect to the server. SPS receives the connection request and establishes the TCP connection with the client.

2. SPS examines the connection request: it checks the IP address of the client and the IP address and port number of the intended destination server. If these parameters of the request match a connection policy configured on SPS, SPS inspects the connection in detail. Other connections are ignored by SPS, and simply forwarded on the packet level.

   The selected connection policy determines all settings and parameters of the connection.

   > NOTE:
   >
   > SPS compares the connection policies to the parameters of the connection request one-by-one, starting with the first policy in the policy list. SPS applies to the connection the first connection policy that completely matches the connection request.
   >
   > For details, see Configuring connections on page 482.

3. SPS selects the destination server based on the **Target** parameter of the connection policy. Network address translation of the target address can be performed at this step. For details, see Modifying the destination address on page 488.
4.
5. SPS selects the source address used in the server-side connection based on the **SNAT** parameter of the connection policy. For details, see Modifying the source address on page 493.
6. The client authenticates itself using an authentication method permitted by the **Authentication policy** set in the Connection policy. Different connections can use different authentication policies, thus allow different authentication methods. The Authentication policy also restricts which users can connect to the server if public-key authentication is used. SPS can authenticate the user to a **Local User Database**, or to a remote LDAP (for example, Microsoft Active Directory) or RADIUS server. This is inband authentication, since it is performed in the same connection that the client originally established to communicate with the server.

   The username used in this authentication step is referred to as the **Gateway username** and is used to determine the **Gateway group** memberships of the user. For details, see Authentication Policies on page 625.

   If an **AA plugin** is configured in SPS, the client may be prompted to provide additional information when authenticating to the server. For details on the **AA plugin**, see Integrating external authentication and authorization systems on page 884. Note that if the plugin sets or overrides the username of the connection, a Usermapping policy needs to be configured and set in the Connection policy. For further information, see Configuring usermapping policies on page 860.
7. If the **Gateway authentication** option is set in the Connection policy, SPS pauses the connection until the user completes a gateway authentication on the SPS web interface. This is out-of-band authentication, since it is performed in an independent connection. For details, see The gateway authentication process on page 58.
8. If the **Usermapping policy** option is set in the Connection policy, SPS checks if the Usermapping policy permits the users of the gateway group to access the username used in the server-side connection (the remote username, for example, `root`). For details, see Configuring usermapping policies on page 860.
9. Before establishing the server-side connection, SPS can evaluate the channel policy to determine if the connection might be permitted at all, for example, it is not denied by a Time policy. SPS performs this check if the **Traffic Controls** > **SSH** > **Settings** > **Enable pre channel check** option is enabled. For details, see Creating and editing protocol-level SSH settings on page 638.

   For the SSH protocol, SPS checks the **From** (client address), **Gateway group**, and **Time policy** restrictions set in the Channel policy of the Connection policy. For details, see Creating and editing channel policies on page 495.

10. **Server-side connection**

    SPS sets up the server-side connection and does the following:

a. SPS establishes the TCP connection to the server.

b. SPS negotiates the protocol parameters of the connection (for example, SSH encryption parameters) according to the **Traffic Controls** > **SSH** > **Settings** of the connection policy.

c. SPS displays an SSH host key to the client. This host key is either generated on SPS, or it is the host key of the server (if it is available on SPS). The connection policy determines the host key shown to the client.

> IMPORTANT:
>
> If the SSH Settings of the Connection Policy enable only RSA keys, set the RSA key shown to the client in the Connection Policy.

d. SPS verifies the host key of the server according to the **Server side host key setting** option of the Connection policy (in general, you can manage the server host keys on the **Traffic Controls** > **SSH** > **Server Host Keys** page). If the server has not been contacted before, SPS can accept and store the host key of the server. Alternatively, the host key of the server can be manually uploaded to SPS. For details, see Server host keys on page 634.

11. SPS performs the authentication on the server, using the data received from:
    - the client during the client-side authentication, or
    - a local or external Credential Store (for details, see Using credential stores for server-side authentication on page 873).
12. SPS authorizes the connection based on the Channel policy. It checks:
    - If the Channel policy includes a **User List** restriction for the **Gateway group** or **Remote group**, SPS checks if the user can access the server. If needed, SPS connects to the LDAP servers set in the **LDAP Servers** policy to resolve the group memberships of the user. For details, see Creating and editing user lists on page 504.
    - SPS consults the **Time policy** assigned to the channel policy. Channels may be opened only within the allowed period.

      > TIP: Time policies are a good way to ensure that the server can be accessed only within the specified timeframe.
13. Both the server- and the client-side connections have been established. From this step, the client can try to open any type and any number of channels in the connection.
14.
15. If 4-eyes authorization is set in the Channel policy, the SSH session of the client is paused until the authorizer permits the client to connect to the server. Who can authorize the session depends on the **Access Control** settings of the Connection policy. For details, see Four-eyes authorization on page 60.
16. The client starts to work on the server. Information about the connection is now available on the **Sessions** page.

- SPS records the entire communication into digitally encrypted audit trails if auditing is enabled in the Channel policy, and encryption is configured in the Audit policy used in the Connection policy. For details, see Creating and editing channel policies on page 495 and Audit policies on page 512.
- If a Content policy is configured in the Channel policy, SPS monitors the connection in real time, and raises an alert or terminates the connection if the user performs an undesired action. For details, see Real-time content monitoring with Content Policies on page 498.

If the user opens another channel within the same connection, SPS consults the Channel policy of the connection to see if the channel is permitted, and processes it accordingly.

17. **Post-processing the connection**

    Once the connection has been closed, the following post-processing steps take place:

    a. After the client closes the connection, or it is terminated for some reason (for example, it times out, or a Content policy or a 4-eyes auditor terminates it), SPS indexes the contents of the audit trail (if the **Record audit trail** option of the Channel policy, and the **Enable indexing** option of the Connection policy are enabled).

    b. SPS creates a backup of the data and the audit trail of the connection, and archives it to a remote server, if a Backup policy and an Archive policy is set in the Connection policy. For more information, see Data and configuration backups on page 145 and Archiving on page 158.

    c. When the **Delete search metadata from SPS after** period expires, SPS deletes all data about the connection from its database.

Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using RDP

The following describes what happens when a client connects to a server through One Identity Safeguard for Privileged Sessions (SPS) and how the different configuration options and policies of SPS affect this process.

1. **Client-side connection**

   The client tries to connect to the server. SPS receives the connection request and establishes the TCP connection with the client.

2. SPS examines the connection request: it checks the IP address of the client and the IP address and port number of the intended destination server. If these parameters of the request match a connection policy configured on SPS, SPS inspects the connection in detail. Other connections are ignored by SPS, and simply forwarded on

the packet level.

The selected connection policy determines all settings and parameters of the connection.

> NOTE:
>
> SPS compares the connection policies to the parameters of the connection request one-by-one, starting with the first policy in the policy list. SPS applies to the connection the first connection policy that completely matches the connection request.
>
> For details, see Configuring connections on page 482.

3. SPS selects the destination server based on the **Target** parameter of the connection policy. Network address translation of the target address can be performed at this step. For details, see Modifying the destination address on page 488.
4.
5. SPS selects the source address used in the server-side connection based on the **SNAT** parameter of the connection policy. For details, see Modifying the source address on page 493.
6. If an **AA plugin** is configured in SPS, the client may be prompted to provide additional information when authenticating to the server. For details on the **AA plugin**, see Integrating external authentication and authorization systems on page 884. Note that if the plugin sets or overrides the username of the connection, a Usermapping policy needs to be configured and set in the Connection policy. For further information, see Configuring usermapping policies on page 860.
7. SPS checks if the client uses a version of the RDP protocol that is enabled in the **Protocol settings** of the Connection policy. Depending on the protocol version, different encryption is used in the connection, and different parameters are required in the Connection policy.
8. Before establishing the server-side connection, SPS can evaluate the channel policy to determine if the connection might be permitted at all, for example, it is not denied by a Time policy. SPS performs this check if the **Traffic Controls** > **RDP** > **Settings** > **Enable pre channel check** option is enabled. For details, see Creating and editing protocol-level RDP settings on page 582.
9. **Server-side connection**
    a. SPS establishes the TCP connection to the server.
    b. SPS checks the protocol parameters of the connection (for example, the version of the RDP protocol used) according to the **Protocol settings** of the Connection policy. The RDP handshake is performed simultaneously on the server- and the client-side.
10. The server opens a Drawing channel for the user to perform authentication.

11. SPS authorizes the connection based on the Channel policy. It checks:
    - If the Channel policy includes a **User List** restriction for the **Gateway group** or **Remote group**, SPS checks if the user can access the server. If needed, SPS connects to the LDAP servers set in the **LDAP Servers** policy to resolve the group memberships of the user. For details, see Creating and editing user lists on page 504.
    - SPS consults the **Time policy** assigned to the channel policy. Channels may be opened only within the allowed period.

      TIP: Time policies are a good way to ensure that the server can be accessed only within the specified timeframe.
12. If the **Gateway authentication** option is set in the Connection policy, SPS pauses the connection until the user completes a gateway authentication on the SPS web interface. This is out-of-band authentication, since it is performed in an independent connection. For details, see The gateway authentication process on page 58.

    It is also possible to perform gateway authentication inband, without having to access SPS's web interface. For details, see Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using an RD Gateway on page 46.
13. SPS performs the authentication on the server, using the data received from:
    - the client during the client-side authentication, or
    - a local or external Credential Store (for details, see Using credential stores for server-side authentication on page 873).
14. If the authentication fails for any reason, SPS terminates the client-side connection as well. This is required to verify the username of the client when it attempts to access the server again.
15. If 4-eyes authorization is set in the Channel policy, the RDP session of the client is paused until the authorizer permits the client to connect to the server. Who can authorize the session depends on the **Access Control** settings of the Connection policy. For details, see Four-eyes authorization on page 60.
16. The client starts to work on the server. Information about the connection is now available on the **Sessions** page.
    - SPS records the entire communication into digitally encrypted audit trails if auditing is enabled in the Channel policy, and encryption is configured in the Audit policy used in the Connection policy. For details, see Creating and editing channel policies on page 495 and Audit policies on page 512.
    - If a Content policy is configured in the Channel policy, SPS monitors the connection in real time, and raises an alert or terminates the connection if the user performs an undesired action. For details, see Real-time content monitoring with Content Policies on page 498.

    If the user opens another channel within the same connection, SPS consults the Channel policy of the connection to see if the channel is permitted, and processes it accordingly.

17. **Post-processing the connection**

    Once the connection has been closed, the following post-processing steps take place:

    a. After the client closes the connection, or it is terminated for some reason (for example, it times out, or a Content policy or a 4-eyes auditor terminates it), SPS indexes the contents of the audit trail (if the **Record audit trail** option of the Channel policy, and the **Enable indexing** option of the Connection policy are enabled).

    b. SPS creates a backup of the data and the audit trail of the connection, and archives it to a remote server, if a Backup policy and an Archive policy is set in the Connection policy. For more information, see Data and configuration backups on page 145 and Archiving on page 158.

    c. When the **Delete search metadata from SPS after** period expires, SPS deletes all data about the connection from its database.

Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using an RD Gateway

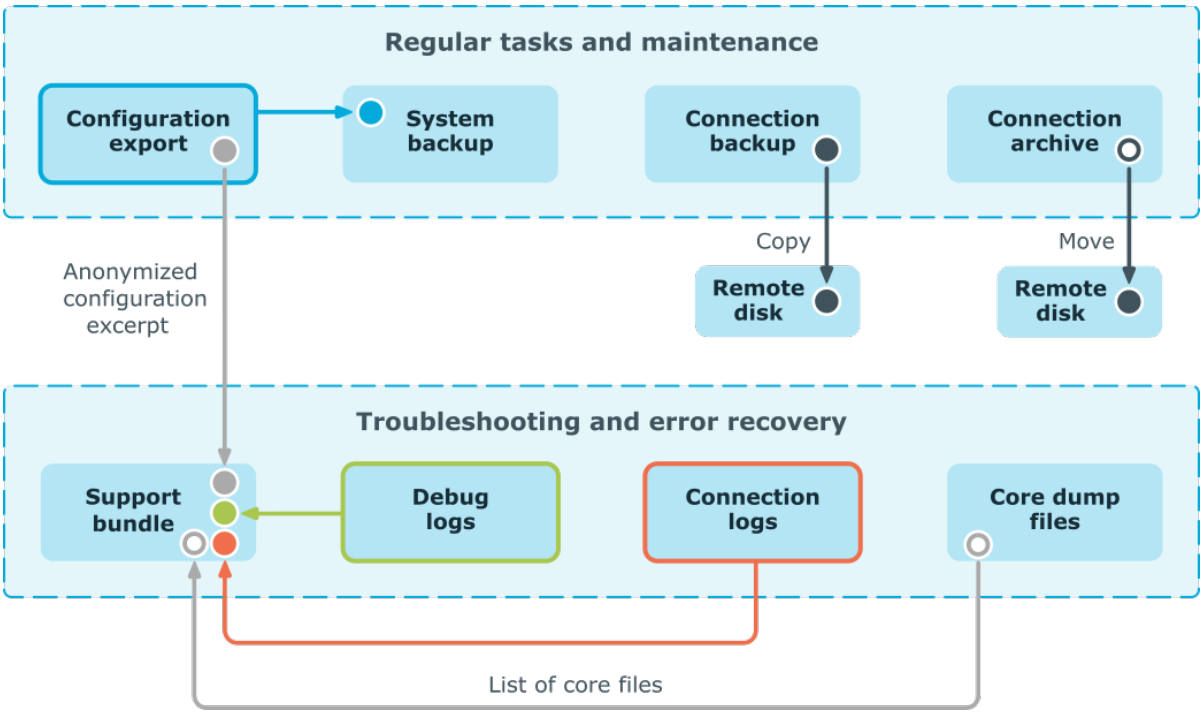
The following describes what happens when a client connects a server through One Identity Safeguard for Privileged Sessions (SPS) using a Remote Desktop Gateway (or RD Gateway), and how the different configuration options and policies of SPS affect this process. For details on the configuration process, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.

1. The client connects to port 443 of the Remote Desktop Gateway configured in the Remote Desktop software. The address of the Remote Desktop Gateway is an alias IP address of SPS. To process the connection request, SPS must have a Connection policy that is configured to handle RDP connection requests on the alias IP, and that has the **Act as a Remote Desktop Gateway** option enabled.
2. The client authenticates on Remote Desktop Gateway (that is, on SPS). Technically, this is an inband gateway authentication on the Domain Controller of SPS's domain (SPS must be the member of a domain, for details, see Network Level Authentication (NLA) with domain membership on page 588). The username used in this authentication step is referred to as the **Gateway username** and is used to determine the **Gateway group** memberships of the user.
3. The client tries to connect to the server. From this point on, this connection is processed as described in Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using RDP on page 43.

Archive and backup concepts

You can export, backup and save various types of data from One Identity Safeguard for Privileged Sessions (SPS), and it also creates log files, dumps and bundles to help the Support Team in troubleshooting errors.

Figure 12: Archive and backup concepts



The following sections describe these in detail:

- Configuration export
- System backup
- Connection backup
- Connection archive
- Support bundle
- Debug logs
- Connection logs
- Core dump files

Configuration export

The configuration of One Identity Safeguard for Privileged Sessions (SPS) can be exported to your local machine from the **Basic Settings** > **System** > **Export configuration** page.

The configuration export in itself is always a one-time action that cannot be configured in policies. However, the system backup (System backup on page 49), that contains the configuration export in addition to other items, can be configured as a scheduled policy and is saved to a backup server.

The exported file is a gzip-compressed archive. On Windows platforms, it can be decompressed with common archive managers such as the free 7-Zip tool.

The name of the exported file is `<hostname_of_SPS>-YYYYMMDDTHHMM.config`, the `-encrypted` or `-gpg` suffix is added for password-encrypted and GPG-encrypted files, respectively. Because the configuration export contains highly sensitive information, it is strongly suggested that you use encryption when generating the export.

For details on how to export the configuration of SPS, see: Exporting the configuration of SPS.

The configuration export is used for

- Manually archiving the configuration.
- Reinstalling a SPS machine and restoring its configuration.
- Migrating the configuration of an already installed SPS to a freshly installed SPS of the same version and therefore creating a machine with an identical configuration.

The configuration export contains the following

- Configuration XML file
- Every change of the configuration of SPS. You can also access these changes at **Users & Access Control** > **Configuration History** in a search interface.
- Certificates, for example:
    - CA certificates
    - TSA certificates
    - Signing CA
- Stored key files, for example:
    - Trusted keys
    - User keys
    - RDP5 RSA key
- User Preferences that are configured at **User Menu** > **Preferences**.
- Certificates and corresponding private keys in your private keystore that are configured at **User Menu** > **Private Keystore**. Only the content of the **Permanent** keystore is exported.
- Custom Report Logo configured at **Reporting** > **Create & Manage Reports**.
- Plugins and any data persisted by plugins.
- Local Credentials Store (the SQLite database) configured at **Policies** > **Credential Stores**.

System backup

The system backup contains the configuration export in addition to other items. It can be configured as a scheduled policy and is saved to a backup server.

Because the configuration export, which is part of the system backup contains highly sensitive information, it is strongly suggested that you use encryption when generating the export. For more information on encrypting the configuration export part, see: Encrypting configuration backups with GPG.

For more information on how to perform a system backup of One Identity Safeguard for Privileged Sessions (SPS), see: Creating configuration backups. It is a two-step process:

1. Create a backup policy at **Policies** > **Backup & Archive** > **Backup policies**.
2. Assign that policy to the system backup at **Basic Settings** > **Management** > **System backup** > **System backup policy**.

   Select **Encrypt configuration**.

For more information on how to restore the configuration and data of SPS from a complete backup, for example, after a hardware replacement, see: Restoring SPS configuration and data.

The system backup is used for

- Recovery in case of errors.

The system backup contains the following

- `config` directory:

  One configuration export file per scheduled backup.
- `db` directory:

  A database dump from SPS's connection metadata database, one `.sql` file overwritten with the actual dump on a daily basis.
- `reports` directory:

  The scheduled daily, weekly, monthly system reports that are accessible at **Reporting** > **Download reports** are saved in `.pdf` files.
- `rrd` directory:

  The output files of the internal system monitoring tool (Munin). These are the files that are used in generating graphs/charts on the **Basic Settings** > **Dashboard** page.
- `sql` directory:

  The internal SQLite databases, for example metadata about the reports.

Connection backup

The connection backup, also known as data backup contains the audit files and connection metadata of a connection. It can be configured as a scheduled policy and is saved to a backup server.

For more information on how to perform a connection backup of a connection, see: Creating data backups. It is a three-step process:

1. Configure a system backup. Restoring a data backup works only if a matching system configuration and metadata is available, that is, if a system backup is restored first.
2. Create a backup policy at **Policies** > **Backup & Archive** > **Backup policies**.
3. Navigate to **Traffic Controls** > **Protocol name** > **Connections**. Select the connection you want to back up. Select the previously created backup policy in the **Backup policy** field.

For more information on how to restore the configuration and data of One Identity Safeguard for Privileged Sessions (SPS) from a complete backup, for example, after a hardware replacement, see: Restoring SPS configuration and data.

The connection backup is used for

- Saving the created audit trail files and indexing metadata of a connection to a remote share. This is a copy operation in terms of data files.
- Recovery: In case of a hardware replacement, creating configuration export, system backup and connection backups is essential.
- Migration: Creating a machine identical to another SPS machine.

The connection backup contains the following

- The audit trails of the connection, that is, the `.zat` files storing the recorded activities of the administrators. For more information on audit trails, see Audit Policies.
- The index of the audit trail that makes the content of the audit trail searchable. For more information on indexing audit trails, see Indexing audit trails.

NOTE: Audit trails and index files are large. This means that backing up a connection requires a significant amount of free hardware space. Make sure you have enough free hardware space for those connections that you want to back up.

Connection archive

The connection archive, also known as data archive contains the audit files and connection metadata of a connection. In terms of contents, it is similar to a connection backup. It can be configured as a scheduled policy and is saved to an archive server. Archiving transfers data from One Identity Safeguard for Privileged Sessions (SPS) to an external storage solution, cleanup removes (deletes) old files. Archived data can be accessed and searched, but cannot be restored (moved back) to the SPS appliance.

For more information on how to perform a connection archive of a connection, see: Archiving or cleaning up the collected data. It is a two-step process:

1. Create an archive policy at **Policies** > **Backup & Archive** > **Archive policies**.
2. Navigate to **<Protocol name> Control** > **Connections**. Select the connection you want to archive. Select the previously created archive policy in the **Archive policy** field.

CAUTION:

Hazard of data loss *Never delete an Archive Policy if data has been archived with it. This will make the already archived data inaccessible.*

Do not "remake" an Archive Policy (that is, deleting an Archive Policy and then creating another one with the same name but different parameters). This will make data inaccessible, and identifying the root cause of the issue complicated.

If you want to change the connection parameters (that is when you perform a storage server migration), you must make sure that the share contents and file permissions are kept unmodified and there are no archiving or backup tasks running.

On the other hand, if you want to add a new network share to your archives, proceed with the following steps:

1. Create a new empty SMB/NFS network share.
2. Create a new Archive Policy that points to this network share.
3. Modify your Connection Policy(es) to archive using the newly defined Archive Policy.
4. Make sure to leave the existing Archive Policy unmodified.

It is also safe to extend the size of the network share on the server side.

The connection archive is used for

- Moving the created audit trail files and indexing metadata of a connection to a remote share. This is a move operation in terms of data files. Archived data can be accessed and searched, but cannot be restored (moved back) to the SPS appliance.
- Freeing up hardware space on SPS.

The connection archive contains the following

- The audit trails of the connection, that is, the `.zat` files storing the recorded activities of the administrators. For more information on audit trails, see Audit Policies.
- The index of the audit trail that makes the content of the audit trail searchable. For more information on indexing audit trails, see Indexing audit trails.

Support bundle

To track down support requests, the One Identity Support Team might request you to collect system-state and debugging information. This information is collected automatically, and contains log files, the anonymized excerpt of the configuration export file of One Identity Safeguard for Privileged Sessions (SPS), and various system-statistics.

To generate a support bundle, navigate to **Basic Settings** > **Troubleshooting** > **Create support bundle**.

The exported file is a zip-compressed archive.

The name of the exported file is `debug_info-<hostname>YYYYMMDDHHMM`. Sensitive data like key files and passwords are automatically removed from the configuration files.

For details on how to create a support bundle, see: Collecting logs and system information for error reporting.

The support bundle is used for

- Collecting a snapshot of the past week's system-state information for the One Identity Support Team for troubleshooting and debugging purposes.
- Collecting information about a specific error by generating data for a defined time interval where the event that causes the error is reproduced. This is also used by the One Identity Support Team for troubleshooting and debugging purposes.

The support bundle contains the following

- Debug logs, Connection logs and OS logs of the past week, one file per day. If there are too many events in a day, the log file in the support bundle only contains a truncated version of the connection logs. In this case, the complete log file is only accessible at `/var/log/messages-<day>`.
- An excerpt of the configuration export file:
    - The anonymized version of the configuration XML file
    - Plugins, which include the complete plugin .zip file and the plugin source code.
- System-state information, for example, version details, statistics, memory usage, system warnings, and so on.
- List of core files. This list might indicate previous system crashes.
- RAID controller information
- Upgrade logs
- Dashboard data

Debug logs

To increase the log level of the non-connection-related events, for example, to add the commands run by the One Identity Safeguard for Privileged Sessions (SPS) web interface to the logs, enable debug level logging at **Basic Settings** > **Management** > **Verbose system logs** > **Enable**.

These logs are accessible at `/var/log/scb-<day>`.

The debug logs are used for

- Our Support Team uses this to investigate the reasons behind a web user interface-related issue.

The debug logs contain the following

- Logs generated by the SPS web interface.
- System daemon logs.
- Logs of periodic cron jobs.

Connection logs

The connection logs contain all connection-related information of the past week, one file per day. A file contains all logs for all connections for a single day.

The logging level of One Identity Safeguard for Privileged Sessions (SPS) can be set separately for every protocol. To change the verbosity level of SPS, navigate to **Traffic Controls** > **Protocol name** > **Global Options**.

These logs are accessible at `/var/log/zorp-<protocol-name>-<day>`.

NOTE: The verbosity level ranges from 1 (no logging) to 10 (extremely detailed), with level 4 being the default normal level. To debug complex problems, you might have to increase the verbosity level to 7. Higher level is needed only in extreme cases.

CAUTION:

High verbosity levels generate very large amount of log messages and might result in a very high load on the machine.

For log levels 8-10, the logs contain highly sensitive data for all connections, as well as passwords and private keys in plain text format.

The connection logs are used for

- Our Support Team uses this to investigate the reasons behind a failed connection.

The connection logs contain the following

- Connection success/failure events
- Other connection-related events

Core dump files

One Identity Safeguard for Privileged Sessions (SPS) automatically generates core dump files if an important software component (for example, Zorp) of the system crashes for some reason. These core dump files can be of great help to the One Identity Support Team to identify problems. When a core dump file is generated, the SPS administrator receives

an alerting e-mail, and an SNMP trap is generated if alerting is properly configured (for details, see Configuring system monitoring on SPS on page 136 and System logging, SNMP and e-mail alerts on page 125).

To list and download the generated core dump files, navigate to **Basic Settings** > **Troubleshooting** > **Core files**.

For details on core dump files, see: Gathering data about system problems.

The core dump files are used for

- The One Identity Support Team uses this to investigate the reasons behind a system crash.

The core dump files contain the following

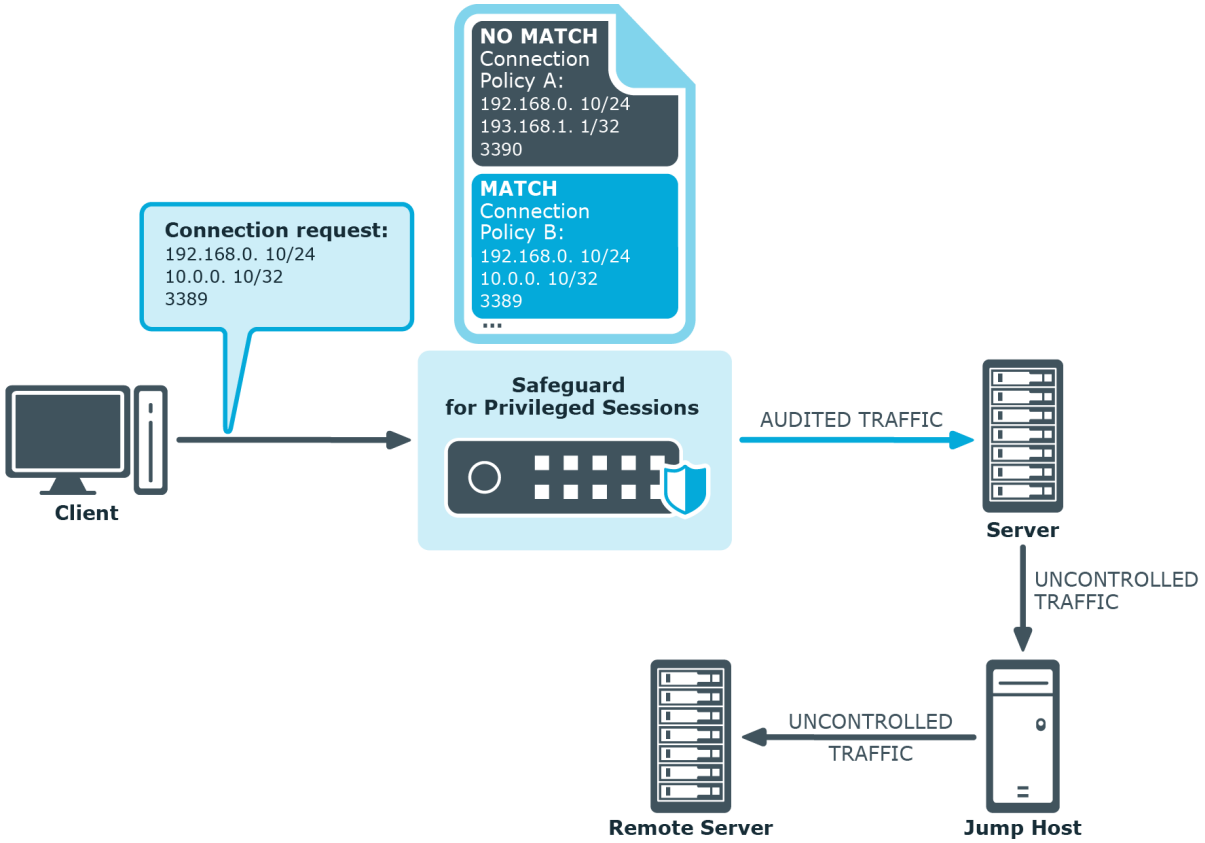
- The recorded state of the working memory of a computer program at a specific time, generally when the program has crashed or otherwise terminated abnormally.

Maximizing the scope of auditing

In certain special scenarios, One Identity Safeguard for Privileged Sessions (SPS) may examine and audit network traffic with some limitations, depending on the configuration.

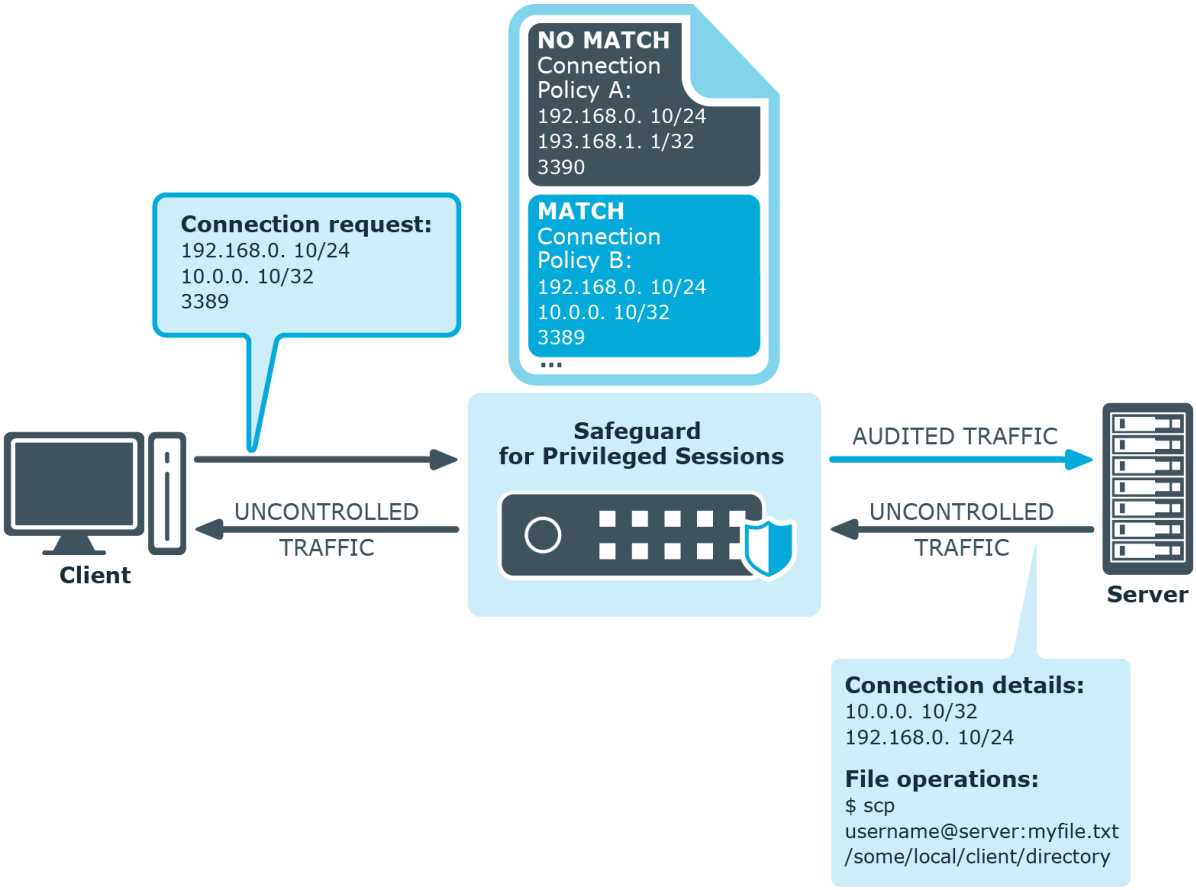
In the first scenario, your organization uses jump hosts to access remote servers or services. In this case, SPS ignores the connection between the target server and the remote server, as it does not go through SPS.

Figure 13: Connection to a remote server through a jump host



In the next scenario, a file operation is performed going from the target server to the client (for example, copying a file using SCP). In this case, the direction of the connection is switched, as compared to the initial client-to-server direction.

Figure 14: File operation in the "reverse" direction



In these scenarios, SPS may not:

- Restrict channels allowed in the connection.
- Audit file operations.

  When you wish to search for the audit files of these connections, there will be no results returned on the **Sessions** page.
- Allow authentication on the remote server if the user authenticates to the target server using a Credential Store.

If you want all connections in these scenarios to be audited, make sure that you add a connection policy for:

- The connection between the target server and any remote servers.
- The connection going from the target server to the client.

IPv6 in One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) supports IPv6 for monitoring connections only. You can define both IPv4 and IPv6 addresses for its logical network interfaces, and configure connections between IPv4 and IPv6 networks (for example, from a client with an IPv4 address to a target with an IPv6 address). You can also use IPv6 addresses with inband destination selection.

NOTE: IPv6 support in ICA connections is currently experimental only.

When configuring IPv6 addresses, SPS shortens the address to its canonical form (omitting leading zeroes, and replacing consecutive sections of zeroes with a double colon). Take the following address as an example:

`2001:0db8:0000:0000:0000:ff00:0042:8329`

SPS shortens the address to its canonical form:

`[2001:db8::ff00:42:8329]`

Additionally, where the IP address and the port is displayed together, IPv6 addresses are shown between brackets. For example, the same address with a port number of 443 is displayed as:

`[2001:db8::ff00:42:8329]:443`

You can search for both the initial (full) and the canonical form on the SPS Search page.

To provide the network range for IPv6 addresses, use network prefixes. Pay attention to the differences between IPv4 and IPv6 network ranges: for IPv4, you can limit the address range to a single address with a prefix of **/32**, but to achieve the same on an IPv6 network, you have to use set the prefix to **/128**.

SSH host keys

SSH communication authenticates the remote SSH server using public-key cryptography, using plain host keys.

The identity of the remote server can be verified by inspecting its host key. When trying to connect to a server through One Identity Safeguard for Privileged Sessions (SPS), the client sees a host key shown by SPS. This key is either the host key of SPS, or the original host key of the server, provided that the private key of the server has been uploaded to SPS. In the latter case, the client will not notice any difference and have no knowledge that it is not communicating directly with the server, but with SPS.

Supported SSH host keys

SPS allows you to use the following SSH host keys:

- RSA (ssh-rsa), which is the most widely used public-key algorithm for the SSH key. In SPS, uploading RSA SSH host keys are supported in PKCS #1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

  NOTE:

  TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
- Ed25519 (ssh-ed25519), which offers a better security and faster performance compared to RSA.

  In SPS, uploading Ed25519 SSH host keys are supported in PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.
- ECDSA NIST P-256 (ecdsa-sha2-nistp256), which is a variant of the Digital Signature Algorithm (DSA). In SPS, uploading ECDSA SSH host keys are supported in SEC1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

You can have multiple SSH server host keys on SPS for the same server, however, you cannot set more than one key for each type. For example, you can keep your old RSA SSH key and generate a new Ed25519 key but you cannot set two RSA keys.

Authenticating clients using public-key authentication in SSH

Public-key authentication requires a private and a public key (or an X.509 certificate) to be available on One Identity Safeguard for Privileged Sessions (SPS). First, the public key of the user is needed to verify the user's identity in the client-side SSH connection: the key presented by the client is compared to the one stored on SPS. SPS uses a private key to authenticate itself to the sever in the server-side connection. SPS can use the private key of the user if it is uploaded to SPS. Alternatively, SPS can generate a new keypair, and use its private key for the server-side authentication, or use agent-forwarding, and authenticate the client with its own key.

CAUTION:

If SPS generates the private key for the server-side authentication, then the public part of the keypair must be imported to the server, otherwise the authentication will fail. Alternatively, SPS can upload the public key (or a generated X.509 certificate) into an LDAP database.

The gateway authentication process

When gateway authentication is required for a connection, the user must authenticate on One Identity Safeguard for Privileged Sessions (SPS) as well.

This additional authentication can be performed:

- *Out-of-band*: in a protocol-independent way, on the web interface of SPS.

  That way the connections can be authenticated to the central authentication database (for example, LDAP or RADIUS), even if the protocol itself does not support authentication databases. Also, connections using general usernames (for example, **root**, **Administrator**, and so on) can be connected to real user accounts.

- *Inband*: when the protocol allows it, using the incoming connection itself for communication with the authentication database.

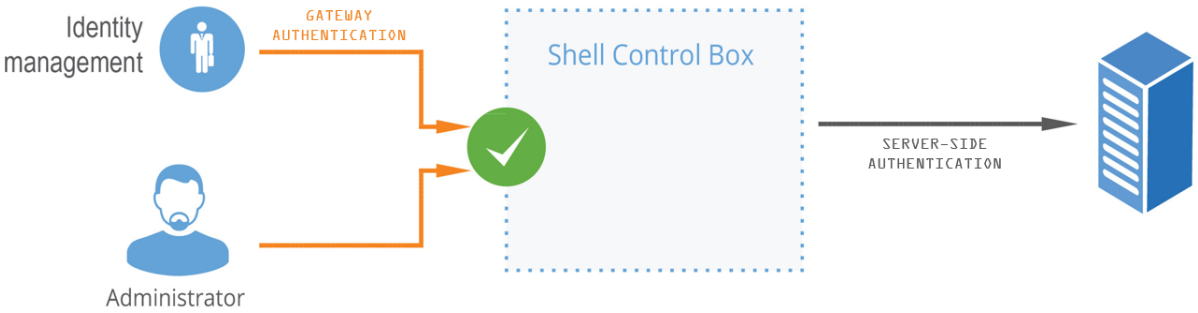
  It is the SSH, RDP, and Telnet protocols that allow gateway authentication to be performed also inband, without having to access the SPS web interface.

  For SSH and Telnet connections, inband gateway authentication must be performed when client-side authentication is configured. For details on configuring client-side authentication, see Client-side authentication settings on page 627.

  For RDP connections, inband gateway authentication must be performed when SPS is acting as a Remote Desktop Gateway (or RD Gateway). In this case, the client authenticates to the Domain Controller or a local user database. For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.

  In the case of RDP connections, inband gateway authentication can also be performed if an AA plugin is configured.

Figure 15: Gateway authentication



Technically, the process of gateway authentication is the following:

1. The user initiates a connection from a client.
2. If gateway authentication is required for the connection, SPS pauses the connection.
3. *Out-of-band authentication:*

   The user logs in to the SPS web interface, selects the connection from the list of paused connections, and enables it. It is possible to require that the authenticated session and the web session originate from the same client IP address.

   Inband authentication:

   SPS requests the username and optionally the credentials for gateway authentication. The user logs in to the SPS gateway.
4. The user performs the authentication on the server.

NOTE: Gateway authentication can be used together with other advanced authentication and authorization techniques like four-eyes authorization, client- and server-side authentication, and so on.

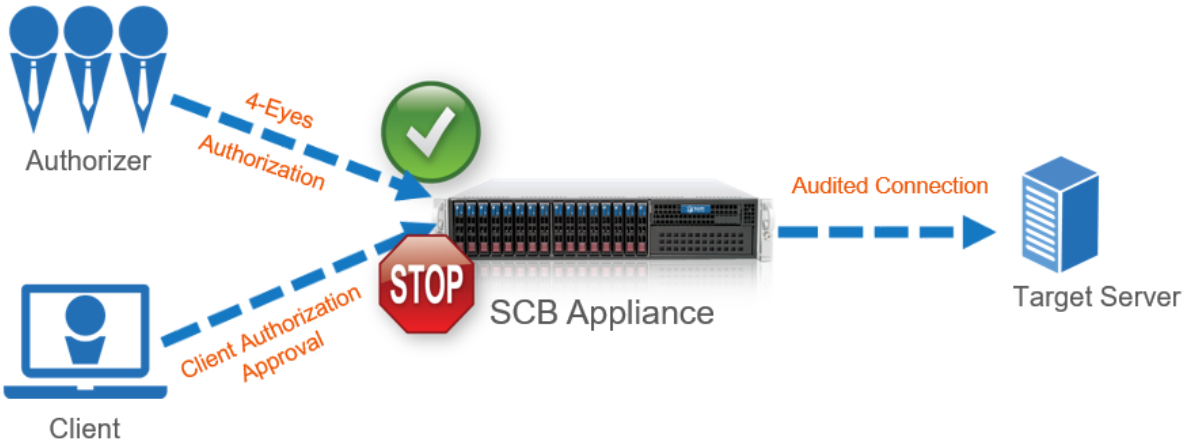
Four-eyes authorization

When four-eyes authorization is required for a connection, a user (called authorizer) must authorize the connection on One Identity Safeguard for Privileged Sessions (SPS) as well. This authorization is in addition to any authentication or group membership requirements needed for the user to access the remote server. Any connection can use four-eyes authorization, so it provides a protocol-independent, out-of-band authorization and monitoring method.

The authorizer has the possibility to terminate the connection any time, and also to monitor real-time the events of the authorized connections: SPS can stream the traffic to the Safeguard Desktop Player application, where the authorizer (or a separate auditor) can watch exactly what the user does on the server, just like watching a movie.

NOTE: The auditor can only see the events if the required decryption keys are available on the host running the Safeguard Desktop Player application.

Figure 16: Four-eyes authorization



Technically, the process of four-eyes authorization is the following:

NOTE: Four-eyes authorization can be used together with other advanced authentication and authorization techniques like gateway authentication , client- and server-side authentication, and so on.

1. The user initiates a connection from a client.
2. If four-eyes authorization is required for the connection, SPS pauses the connection.
3. The authorizer logs in to the SPS web interface, selects the connection from the list of paused connections, and enables it.
4. The user performs the authentication on the server.

5. The auditor (who can be the authorizer, but it is possible to separate the roles) watches the actions of the user real-time.

Network interfaces

The One Identity Safeguard for Privileged Sessions (SPS) hardware has five network interfaces: three physical interfaces for handling traffic, the HA interface for communicating with other nodes in a High Availability cluster, and the IPMI. The Appliance 3500 and 4000 hardware have two additional network interfaces available: the SFP+ interfaces labeled 5 and 6. For details on hardware installation, see *One Identity Safeguard for Privileged Sessions Hardware Installation Guide* in the *Installation Guide*.

You can assign any number of logical interfaces (alias IP addresses and netmasks) to a physical interface, and each logical interface can have its own VLAN ID. For more information on managing logical interfaces, see Managing logical interfaces on page 119.

The routing rules determine which interface is used for transferring remote backups and syslog messages of SPS.

> TIP: One Identity recommends that you direct backups, syslog and SNMP messages, and e-mail alerts to a dedicated interface. For details, see Configuring the routing table on page 122.

The HA interface is an interface reserved for communication between the nodes of SPS clusters. The HA interface uses the Ethernet connector labeled as `4` (or `HA`). For details on High Availability, see High Availability support in One Identity Safeguard for Privileged Sessions (SPS) on page 61.

In case of Appliance 3500 and 4000 hardware, the SFP+ interfaces are available for both proxy traffic and for local services. This means that these interfaces can be used for the same purposes as the other three physical interfaces.

The Intelligent Platform Management Interface (`IPMI`) allows system administrators to monitor system health and to manage SPS events remotely. IPMI operates independently of the operating system of SPS.

High Availability support in One Identity Safeguard for Privileged Sessions (SPS)

The goal of HA nodes is to support enterprise business continuity by providing failover.

In High Availability (HA) mode, two One Identity Safeguard for Privileged Sessions (SPS) units with identical configurations are operating simultaneously. These two units are the primary node and the secondary node (previously also referred to as the master node and the slave node). The primary node shares all data with the secondary node, and if the primary node stops functioning, the secondary node becomes active.

CAUTION: Depending on your environment, the shift between the two HA nodes might take 5 - 10 minutes. During this period, SPS connections are lost and they are not automatically restored. You might also encounter data loss during the shift.

You can find more information on managing a High Availability SPS cluster in Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster on page 423.

One Identity recommends using a High Availability SPS cluster instead of a standalone SPS appliance. A standalone SPS appliance can become a single point of failure (SPOF), and its failure can severely impact your business.

Firmware and High Availability

When powering on the One Identity Safeguard for Privileged Sessions (SPS) nodes in High Availability mode, both nodes boot and start the firmware. There is a difference, however, between the two nodes in the services that they start on booting. The secondary node will launch only a few services, those that are required for High Availability support (that is, for awareness of the primary node and data synchronization). The rest of the services (for example, managing connections) start only on the primary node.

Upgrading the SPS firmware via the web interface automatically upgrades the firmware on both nodes.

Versions and releases of One Identity Safeguard for Privileged Sessions (SPS)

The following release policy applies to One Identity Safeguard for Privileged Sessions (SPS):

One Identity Safeguard for Privileged Sessions customers choose between two paths for receiving SPS releases: Long Term Supported (LTS) release or feature release.

Releases

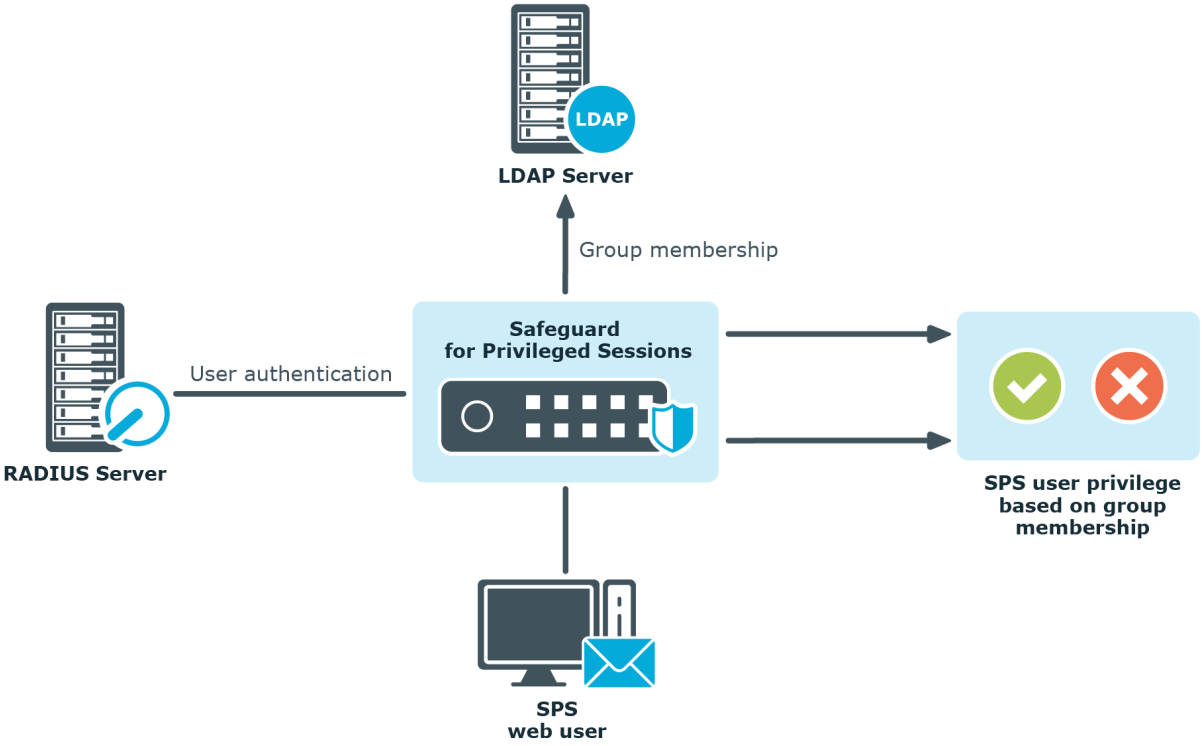
	LTS release	Feature release
Release frequency	**Frequency**: Typically, every 2 years **Scope**: Includes new features, resolved issues and security updates **Versioning**: First digit identifies the LTS and the second digit is a 0 (for example, 6.0, 7.0, and so	**Frequency**: Typically, every 2 months **Scope**: Includes the latest features, resolved issues, and other updates, such as security patches for the OS **Versioning**: First digit identifies the LTS and the second digit is a number identifying the feature release (for example, 6.1, 6.2, and so on)

	on)	
Maintenance release	**Frequency**:Typically, every 2 months during full support **Scope**: Includes important resolved issues and security updates **Versioning**: Third digit designates the LTS maintenance release (for example, 6.0.1)	**Frequency**:Only for highly critical issues **Scope**: Includes highly critical resolved issues **Versioning**: Third digit designates the feature maintenance release (for example, 6.1.1)
Support	For more information on the product support, see Product Support - One Identity Safeguard for Privileged Sessions. For a full description of long-term-supported and feature releases, see Product Life Cycle & Policies - One Identity Safeguard for Privileged Sessions.	

Accessing and configuring One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) has a web interface and is configured from a browser. The users of SPS can be authenticated using local, LDAP, or RADIUS databases. The privileges of users are determined by group memberships that can be managed either locally on SPS, or centrally in an LDAP database. Assigning privileges to groups is based on Access Control Lists (ACLs). It is also possible to match groups existing in the LDAP database to a set of SPS privileges. Access control in SPS is very granular, therefore it is possible to define exactly who can access which parts of the interface and of the stored data.

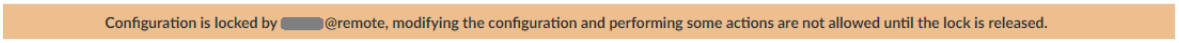
Figure 17: Authenticating the users of SPS



Configuration lock on the SPS web interface

When there is another user logged in to SPS, making configuration changes on legacy pages of the web interface, the configuration lock feature is automatically switched on. During this time, you cannot edit configuration parameters and commit changes on these legacy pages of the web interface. For example, you cannot edit the **Type** of your ACL objects, and certain UI elements will be grayed out or do not respond the same way as they do while the configuration lock is not activated. This state will stay active until the configuration lock is released by the user making configuration changes or until the user logs out of SPS.

TIP: If UI elements do not respond the way they usually do in SPS, check if there is a similar banner near the top of the web interface:



Cloud deployment considerations

One Identity Safeguard for Privileged Sessions (SPS) can be run from the cloud.

Before you start: platforms and resources

When setting up a virtual environment, carefully consider the configuration aspects such as CPU, memory availability, I/O subsystem, and network infrastructure to ensure the virtual layer has the necessary resources available. For more information on environment virtualization, see One Identity's Product Support Policies.

Platforms that have been tested with the cloud deployments are:

- AWS Virtual Machine (VM): AWS deployment
- Azure Virtual Machine (VM): Azure deployment

AWS deployment

One Identity Safeguard for Privileged Sessions (SPS) can be run in the cloud using Amazon Web Services (AWS).

To deploy the Amazon Machine Image (AMI) of SPS from AWS, visit the AWS marketplace listing for SPS (here) and follow the Deployment steps .

Limitations

- If High Availability (HA) operation mode is required in a virtual environment, use the HA function provided by the virtual environment.
- When running SPS in a virtual environment, use a single network interface.
- During AWS installation, connecting directly to the Internet using a public IP address is not supported. Instead, you must access the Internet via a Virtual Private Network or a jump host.
- Factory reset is not an option for virtual appliances. To factory reset a virtual appliance, just redeploy the appliance.

Disk size considerations

CAUTION: Before making any changes to the disk size, shut down the VM (stopped and deallocated).

SPS deploys with a minimal OS disk size. You should increase the size of the OS disk based on your estimated usage and budget.

To have your appliance set with the correct disk size, you must resize the OS disk right after creating the VM (after the Welcome Wizard finishes).

1. Log in via SSH or Amazon Elastic Compute Cloud (Amazon EC2) serial console. For more information, see Getting Started with Amazon EC2.
2. To resize the disk volume, select **Troubleshooting > Extend core partition**. For more information, see *Modifying the disk size of an SPS virtual appliance* in the *Installation Guide*.
3. Repeat steps 1 and 2 if you want to resize SPS again later.

AWS security considerations

Running SPS in AWS comes with some security considerations that do not apply to the hardware appliance. One Identity recommends:

- Do not give Safeguard a public IP address.
- Use the AWS key vault to encrypt the disk.
- Limit access within AWS to the Safeguard virtual machine. SPS in AWS cannot protect against rogue Administrators in the same way the hardware appliance can.

Static IP address required

Configure the SPS VM with a static IP address in AWS. In AWS, the IP address must not change after the VM is deployed. If you need to change the IP address, take a backup, deploy again, and restore the backup. You can script the VM deploy to pick up an existing virtual NIC with the IP address configuration. For details, see the Amazon Virtual Private Cloud (VPC) documentation.

Deployment steps

AWS automatically licenses the operating system during the deployment with an AWS KMS.

To deploy SPS in AWS

1. Open the AWS marketplace listing for SPS.
2. On the **One Identity Safeguard for Privileged Sessions** page, click **Continue to Subscribe**.
3. To configure your instance, advance through the resource creation screens. In addition to the Disk size considerations, AWS security considerations , and Static IP address required , One Identity recommends you select the instance type according to the intended usage:
    - **m5.xlarge** for standard usage.
    - **m5.2xlarge** or **m5.4xlarge** for heavy usage.
    - **c5.2xlarge**, **c5.4xlarge** or **c6.8xlarge** compute-optimized instances, if your traffic consists of heavy audit trails (for example, RDP or ICA).
    - **r5.large** or **r5.2xlarge** memory-optimized instances, if you need better performance with large Search databases, or when deploying search masters.

Alternatively, consider using instances with similar specifications from the **r5a**, **r5b**, **r5n**, **r6i** classes, and so on.

NOTE: SPS supports enhanced networking capabilities through the Elastic Network Adapter (ENA) on AWS.

For the list of instance types that support ENA, see table *Summary of networking and storage features* in chapter Instance types of the *Amazon EC2 documentation*.

For more information on ENA, see Enhanced networking: ENA in the *Amazon EC2 documentation*.

a. To enable enhanced networking through ENA for your SPS instance, select an instance type that supports ENA.

4. Once you have finished configuring the instance, select and launch it.

   NOTE: The instance launch process may take a few minutes to complete.

5. Once the instance has finished launching, complete the SPS Welcome Wizard. While advancing through the SPS Welcome Wizard, you can create the admin password. For more information, see The Welcome Wizard and the first login.

   NOTE: Ensure that you upgrade SPS to the latest available release. For more information, see the Upgrade Guide.

Azure deployment

The following describes how to have a One Identity Safeguard for Privileged Sessions running in Microsoft Azure.

To have a One Identity Safeguard for Privileged Sessions running in Microsoft Azure

1. **Deploy One Identity Safeguard for Privileged Sessions from the Microsoft Azure Marketplace**

   Create and configure a One Identity Safeguard for Privileged Sessions virtual machine (VM) in the Azure portal. For more information, see the Microsoft Azure documentation, here we just describe the SPS-specific settings.

   a. Login to the Azure portal, select **One Identity Safeguard for Privileged Sessions** from the Azure Marketplace, then click **Create**.

   b. Fill the required fields of the **Basics** blade. Note that you must fill the **User name** and **Authentication Password/SSH public key** fields, but SPS will not actually use these settings (SPS will use the parameters you configure in the SPS Welcome Wizard).

   c. Choose a size for the VM. If you want to use this machine in production and need help about sizing or architecture design, contact your One Identity sales

representative.

The number of interfaces you can use depends on the size of your Azure VM. If your VM allows you to use multiple interfaces, you can configure multiple interfaces in SPS. For details, see VM with multiple NICs.

d. On the **Settings** blade, disable monitoring.

e. When the deployment is finished, navigate to the network settings of the new VM in the Azure portal. Change the IP address of the SPS network interface to Static, and note down the IP address and the hostname (you will need it in the SPS Welcome Wizard).

f. If you want to backup or archive data from SPS into Azure, create an Azure File Share. Note down the following information of the file share, because you will need it to configure SPS backups and archiving: `URL`, `Username`, `Password`.

> **CAUTION:**
>
> If you have multiple SPS VMs, make sure to use a separate file share for each SPS.

2. Complete the SPS Welcome Wizard

Complete the SPS Welcome Wizard (for more information, see *Configuring One Identity Safeguard for Privileged Sessions (SPS) with the Welcome Wizard* in the *Administration Guide*). Note the following points specific for Azure deployments. When configuring the network settings of SPS note the following points.

> **CAUTION:**
>
> Do not export or import configuration between a physical SPS deployment and a virtual one. Because of the differences and limitations between physical and virtual appliances, configure the virtual appliance from scratch to ensure proper functionality. When you migrate a virtual SPS to another one, you can export and import the configuration.

a. Into the **Physical interface EXT or 1 — IP address** field, enter the static IP address of the SPS VM that you set on the Azure portal.

b. **Default GW**: The default gateway is usually the first address in a subnet (for example, if your subnet is `10.7.0.0/24`, then the gateway will be `10.7.0.1`).

c. **Hostname**: Use the hostname you have configured for the SPS VM on the Azure portal.

d. **DNS server**: You can use any DNS server that the SPS VM can access, even public ones.

3. Configure SPS

Login to SPS and configure it.

a. Configure backups for SPS. For backup and archiving purposes One Identity recommends the built-in file shares of Azure. For more information on configuring backups, see *Data and configuration backups* in the *Administration Guide*.

b. Configure archiving for SPS. For backup and archiving purposes One Identity recommends the built-in file shares of Azure. For more information on configuring backups, see *Archiving* in the *Administration Guide*. Configuring Archiving policy is highly recommended: because if the disk of the VM fills up, SPS stops working.

c. Configure a server: set up a host that is on the same subnet as SPS, and enable Remote Desktop (RDP) or Secure Shell (SSH) access to it.

d. Configure a connection on SPS to forward the incoming RDP or Secure Shell (SSH) connection to the host and establish a connection to the host. For more information, see *Logging in to One Identity Safeguard for Privileged Sessions (SPS) and configuring the first connection* in the *Administration Guide*.

e. Replay your session in the browser. For more information, see *Replaying audit trails in your browser* in the *Administration Guide*.

   In case you have questions about SPS, or need assistance, contact your One Identity representative.

Limitations

The following limitations apply to SPS when you deploy it from the Microsoft Azure Marketplace.

CAUTION:

Do not export or import configuration between a physical SPS deployment and a virtual one. Because of the differences and limitations between physical and virtual appliances, configure the virtual appliance from scratch to ensure proper functionality. When you migrate a virtual SPS to another one, you can export and import the configuration.

- Root login is not available on the console.
- SSH access is only available after you have completed the Welcome Wizard.
- Currently, the data that is entered during the provisioning phase (for example, the username and the IP address) of creating the virtual machine in Azure is not transferred to SPS. Therefore, only the data entered in the Welcome Wizard will be used.
- By default, you can only use Physical interface 1 (eth0) of SPS, with a single IP address. Aside from changing the IP address of SPS, do not modify other interface-related settings (additional logical interfaces, IP forwarding, and so on) on the **Basic Settings** > **Network** page of SPS.

The number of interfaces you can use depends on the size of your Azure VM. If your VM allows you to use multiple interfaces, you can configure multiple interfaces in SPS. For details, see VM with multiple NICs.

- The **Seal the box** functionality is not available.
- The High Availability support of SPS was designed to work between two physical SPS appliances. This feature is not available in Azure environments. For more information, see High Availability and redundancy in Microsoft Azure.
- Due to Azure requirements, an additional 5-minute delay has been added to the boot process. This ensures that the root device appears in the system.
- The size of the virtual disk in Azure is 100 GB, which you can increase. To increase the size of the virtual disk, see *Modifying the disk size of a SPS virtual appliance* in the *One Identity Safeguard for Privileged Sessions Installation Guide*.
- SPS currently cannot receive its IP address using DHCP. Make sure that:
    - The IP address you have configured in Azure and the IP address that you configure for SPS for the **Physical interface 1** on the Networking settings part of the Welcome Wizard are the same. Otherwise, you will not be able to access SPS.
    - You set the internal IP static on the Network Interfaces tab of the Virtual Machine.
    - Do not assign a public IP address to SPS, use SPS as a component of your internal infrastructure. If you absolutely must configure Welcome Wizard from a publicly accessible IP address, note that SPS will be publicly accessible. If you assign a public IP to the web management interface, consider the following:
        - Select a complex passphrase.
        - Limit access to the management interface based on the source IP address, and make sure that brute-force protection for the administrator web login is enabled (they are enabled by default). For more information, see *Configuring user and administrator login addresses* in the *Administration Guide*.
        - Configure an email alert or SNMP trap for administrator logon events. For more information, see *Configuring e-mail alerts* in the *Administration Guide*.
        - Forward the logs of SPS to a log server (for example, to a syslog-ng server, or an syslog-ng Store Box appliance) so that if the local logs are compromised, you still have an authentic copy of the original logs.
        - For security reasons, disable SSH access to SPS when it is not needed. Accessing the SPS host directly using SSH is not recommended or supported, except for troubleshooting purposes. If you enable SSH access, restrict the clients that can access SPS based on their source IP address, and make sure that brute-force protection is enabled (they are enabled by default). For more information, see Enabling SSH access to the *One Identity Safeguard for Privileged Sessions (SPS)* host in the

Administration Guide.

- To prevent unauthorized access to the audit trail files recorded on SPS, configure proper access control rules for the user groups and encrypt every audit trail. If you use encryption, store your keys in the personal or in the temporary key store. For more information, see *Encrypting audit trails* in the *Administration Guide*.

- Upgrading SPS in Azure is the same as upgrading a physical appliance: you have to upload the firmware on the SPS web interface. For more information, see the *One Identity Safeguard for Privileged Sessions Upgrade Guide*.
- Azure Accelerated Networking is not yet supported. Avoid deploying SPS Virtual Machines of types other than the recommended, especially those that have Accelerated Networking marked as Required in their specification. Using a VM other than the recommended prevents you from upgrading, and it may cause other unwanted behavior of the appliance. For more information on the specifications of Azure VM types, see Virtual machines in Azure.

Prerequisites

The following prerequisites must be met to deploy SPS in Microsoft Azure:

- You have a valid One Identity Safeguard for Privileged Sessions license. When deployed from the Microsoft Azure Marketplace, the One Identity Safeguard for Privileged Sessions uses the "Bring your own license" model. Note that to deploy two active SPS nodes as an availability set, you must purchase two standalone SPS licenses. To purchase a license, contact our Sales Team.
- Microsoft recommends to use the Azure Resource Manager (ARM) deployment model. When you install SPS from the Azure Marketplace, SPS supports only this deployment method. If you need to deploy SPS into and infrastructure that uses the Classic deployment model, contact your One Identity sales representative.
- You have a Microsoft Azure account.

High Availability and redundancy in Microsoft Azure

In a Microsoft Azure deployment, the high-availability and redundancy of the SPS appliance is provided by the Microsoft Azure infrastructure, according to the Azure Storage SLA.

Redundancy

The data in your Microsoft Azure storage account is always replicated to ensure durability and high availability, meeting the Azure Storage SLA. The exact type of replication depends

on your storage account settings, but every disk is stored in 3 copies.

For more information, see Locally redundant storage in the *Azure Storage replication* document and Service Healing - Auto-recovery of Virtual Machines.

High Availability

If a hardware failure occurs, Azure moves the Virtual Machine to another location and restarts it in 5-15 minutes. In case you require higher SLA, you are recommended to deploy two standalone SPS nodes into an availability set. Note that to deploy two active SPS nodes as an availability set, you must purchase two standalone SPS licenses.

For more information, see Locally redundant storage in the *Azure Storage replication* document and Service Healing - Auto-recovery of Virtual Machines.

4

The Welcome Wizard and the first login

This section describes the initial steps of configuring One Identity Safeguard for Privileged Sessions (SPS). Before completing the steps in The initial connection to One Identity Safeguard for Privileged Sessions (SPS), unpack, assemble, and power on the hardware. Connect interface 1 (labelled 1 or EXT) to the local network, or directly to the computer from which SPS will be configured.

NOTE: For details on unpacking and assembling the hardware, see *Hardware Installation Guide* in the *Installation Guide*. For details on how to create a High Availability SPS cluster, see *Installing two SPS units in HA mode* in the *Installation Guide*. For more information about the supported browsers, see Supported web browsers.

The initial connection to One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) can be connected from a client machine using web browsers and accessed from the local network.

NOTE: For details on supported browsers, see Supported web browsers.

Starting with version 3.1, SPS attempts to receive an IP address automatically via DHCP. If it fails to obtain an automatic IP address, it starts listening for HTTPS connections on the `192.168.1.1` IP address. Note that certain switch configurations and security settings can interfere with SPS receiving an IP address via DHCP. SPS accepts connections via its interface 1 (labelled *1* or *EXT*). For details on the network interfaces, see Network interfaces on page 61).

TIP: The SPS console displays the IP address on which interface 1 is listening.

If SPS is listening on the `192.168.1.1` address, note that the `192.168.1.0/24` subnet must be accessible from the client. If the client machine is in a different subnet (for example, its IP address is `192.168.10.X`), but in the same network segment, the easiest way is to assign an alias IP address to the client machine. Creating an alias IP on the client machine virtually puts both the client and SPS into the same subnet, so that they can communicate. To create an alias IP complete the following steps.

- For details on creating an alias IP on Microsoft Windows, see Creating an alias IP address (Microsoft Windows) on page 74.
- For details on creating an alias IP on Linux, see Creating an alias IP address (Linux) on page 78.
- If configuring an alias interface is not an option for some reason, you can modify the IP address of SPS. For details, see Modifying the IP address of One Identity Safeguard for Privileged Sessions (SPS) on page 79.

CAUTION:

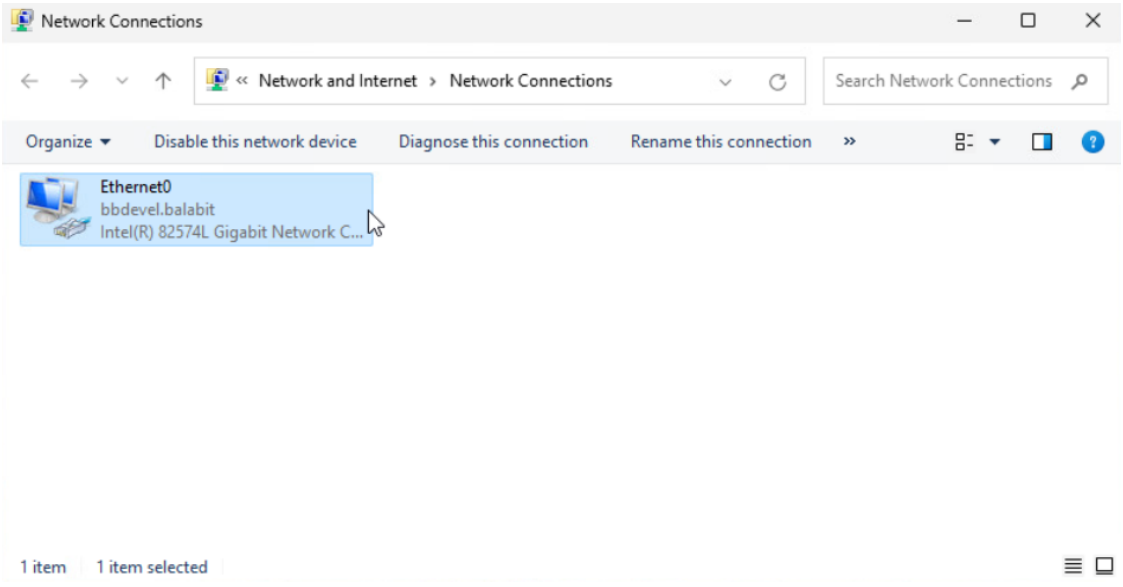
The Welcome Wizard can be accessed only using interface 1, as the other network interfaces are not configured yet.

Creating an alias IP address (Microsoft Windows)

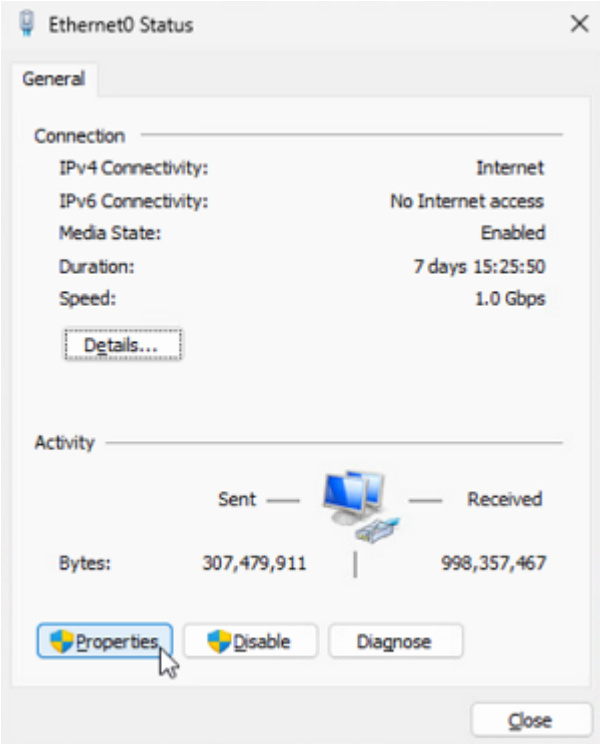
This procedure describes how to assign an alias IP address to a network interface on Microsoft Windows platforms.

To assign an alias IP address to a network interface on Microsoft Windows platforms

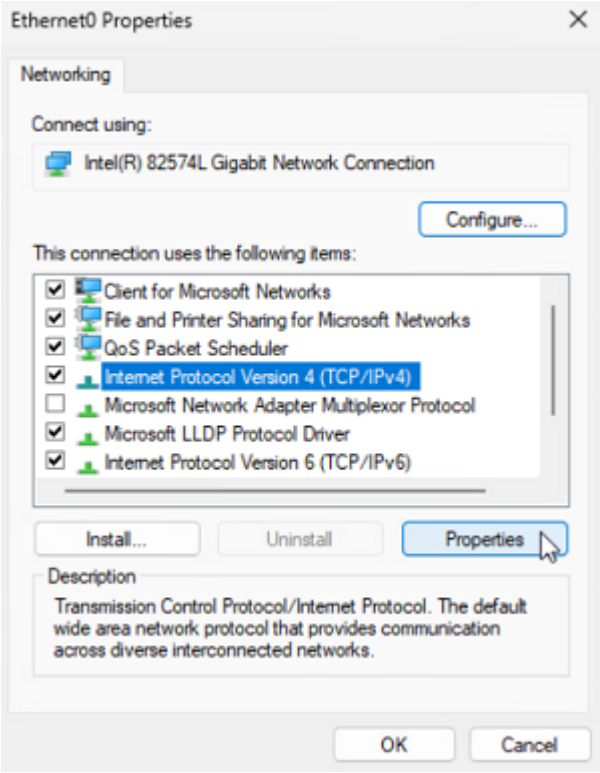
1. Open **Network Connections** via the **Run** dialog box (that is, press the **Windows key** + **R**, type `ncpa.cpl`, then press **Enter**.



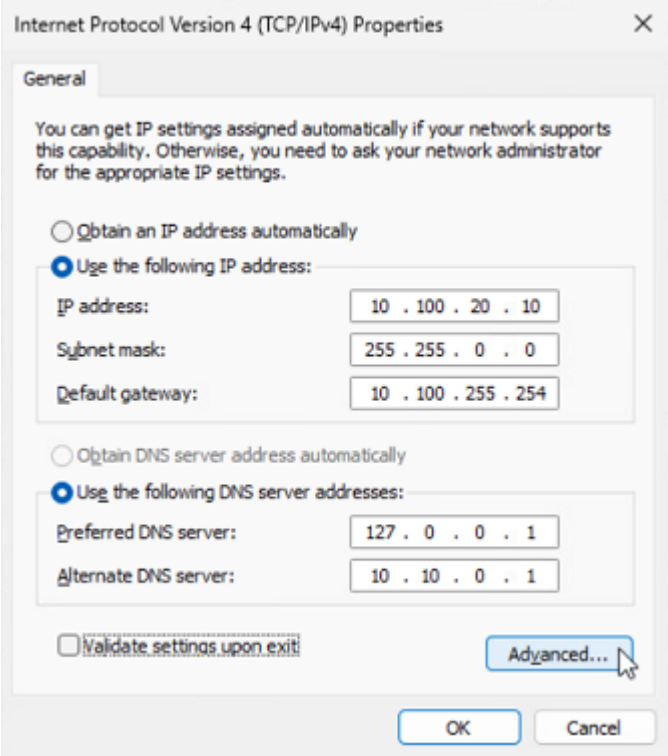
2. Double click the **Ethernet0** network interface, then click **Properties**.



3. Under the **Networking** tab, select **Internet Protocol Version 4 (TCP/IPv4)**.

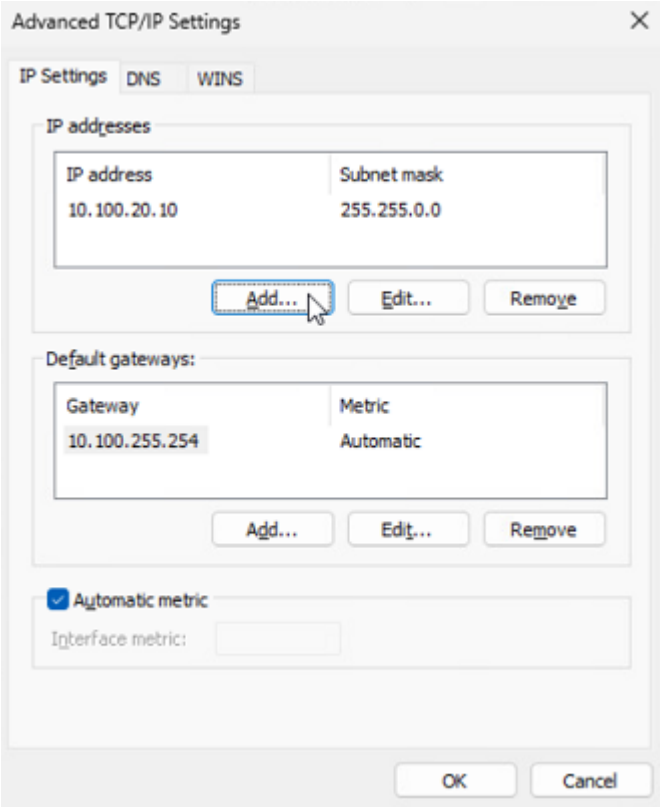


4. If the IP address is not configured automatically, under the **General** tab, select **Advanced**.

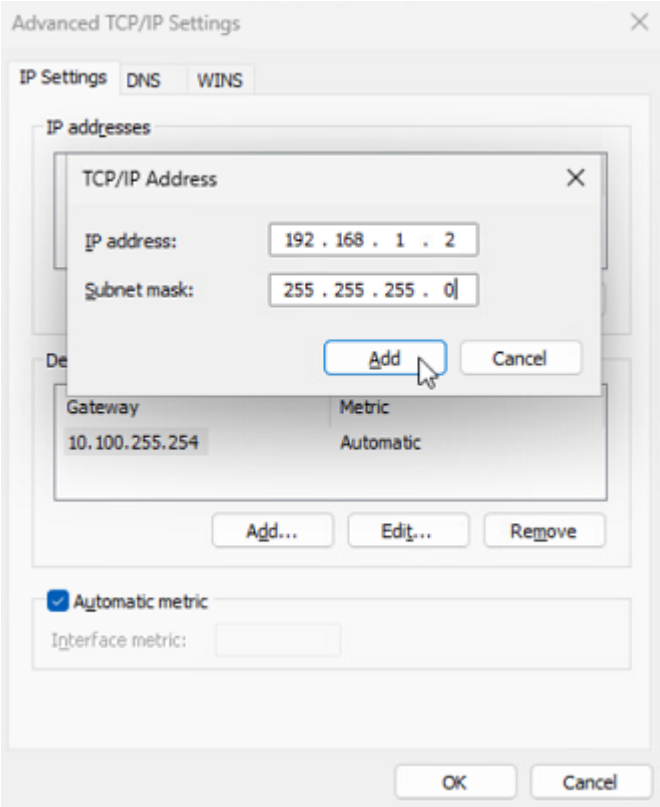


NOTE: You can only add a custom secondary IP address via manual configuration. If the IP address configuration is automatic, the **Advanced** button is inactive.

5. Under the **IP Settings** tab, select **Add** next to the **IP addresses** box.



6. Configure the **IP address** and **Subnet mask** fields.



7. Enter **192.168.1.2** into the **IP address** field, and **255.255.255.0** into the **Subnet mask** field.

   CAUTION:

   If your internal network uses the 192.168.1.0/24 IP range, it is possible that the 192.168.1.1 and 192.168.1.2 addresses are already in use. In this case, disconnect One Identity Safeguard for Privileged Sessions (SPS) from the network, and connect a computer directly to interface 1 (labelled 1 or EXT) using a standard cross-link cable.

8. To activate your settings, click **OK**.

Creating an alias IP address (Linux)

The following describes how to assign an alias IP address to a network interface on Linux platforms.

To assign an alias IP address to a network interface on Linux platforms

1. Start a terminal console (for example, `gnome-terminal`, `konsole`, `xterm`, and so on).
2. Issue the following command as root:

```
ifconfig <ethX>:0 192.168.1.2
```

where `<ethX>` is the ID of the network interface of the client, usually `eth0` or `eth1`.

3. Issue the `ifconfig` command. The `<ethX>:0` interface appears in the output, having `inet addr:192.168.1.2`.
4. Issue the `ping -c 3 192.168.1.1` command to verify that One Identity Safeguard for Privileged Sessions (SPS) is accessible. A similar result is displayed:

```
user@computer:~$ ping -c 3 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp-seq=1 ttl=63 time=0.357 ms
64 bytes from 192.168.1.1: icmp-seq=2 ttl=63 time=0.306 ms
64 bytes from 192.168.1.1: icmp-seq=3 ttl=63 time=0.314 ms

--- 192.168.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2013ms
rtt min/avg/max/mdev = 0.306/0.325/0.357/0.030 ms
```

5. Open the page `https://192.168.1.1` from your browser and accept the certificate shown. The Welcome Wizard of SPS appears.

Modifying the IP address of One Identity Safeguard for Privileged Sessions (SPS)

The following describes how to configure One Identity Safeguard for Privileged Sessions (SPS) to listen for connections on a custom IP address.

CAUTION:

Use this procedure only before the initial configuration of SPS, that is, before completing the Welcome Wizard. For details on changing the IP address or other network settings of a configured SPS system, see Network settings on page 114.

If you change the IP address of SPS, make sure that you use this address as the **Physical interface 1 — IP address** in the **Networking settings** section of the Welcome Wizard (see Configuring interface 1).

To configure SPS to listen for connections on a custom IP address

1. Access SPS from the local console, and log in with username `root` and password `default`.
2. Select **Shells** > **Core shell** in the Console Menu.
3. Change the IP address of SPS:

   ```
   ifconfig eth0 <IP-address> netmask 255.255.255.0
   ```

Replace `<IP-address>` with an IPv4 address suitable for your environment.

4. Set the default gateway using the following command:

   `route add default gw <IP-of-default-gateway>`

   Replace `<IP-of-default-gateway>` with the IP address of the default gateway.

5. Type `exit`, then select **Logout** from the Console Menu.

6. Open the page *https://<IP-address-you-set-for-SPS>* from your browser and accept the certificate shown. The Welcome Wizard of SPS appears.

Configuring One Identity Safeguard for Privileged Sessions (SPS) with the Welcome Wizard

The Welcome Wizard guides you through the basic configuration steps of One Identity Safeguard for Privileged Sessions (SPS). You can modify all parameters before the last step by using the **Back** button of the wizard, or later through the web interface of SPS.

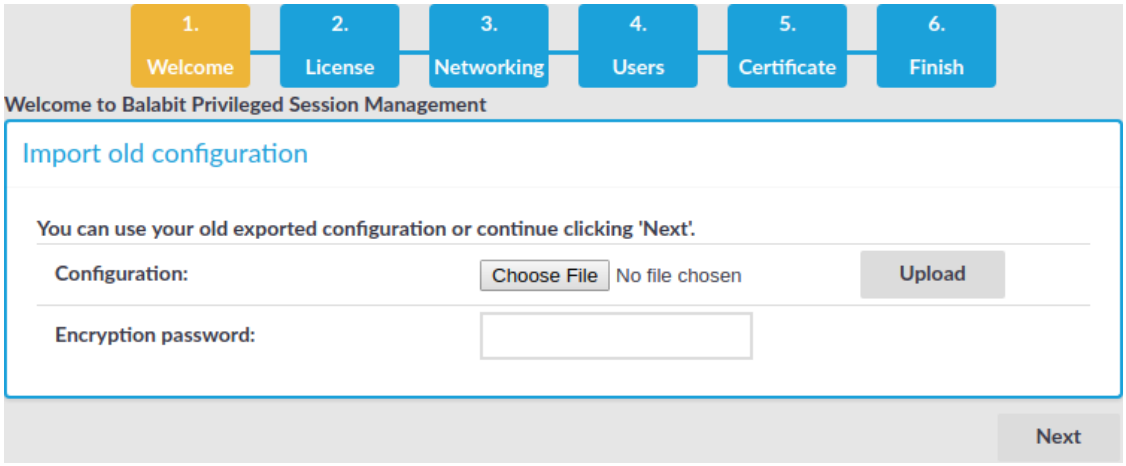
To configure SPS with the Welcome Wizard

1. Open the *https://<IP-address-of-SPS-interface>* page in your browser and accept the displayed certificate. The Welcome Wizard of SPS appears.

   TIP: The SPS console displays the IP address the interface is listening on. SPS either receives an IP address automatically via DHCP, or if a DHCP server is not available, listens on the `192.168.1.1` IP address.

2. When configuring SPS for the first time, click **Next**.

Figure 18: The Welcome Wizard



You can import an existing configuration from a backup file. Use this feature to restore a backup configuration after a recovery, or to migrate an existing SPS configuration to a new device.

CAUTION:

Do not export or import configuration between a physical SPS deployment and a virtual one. Because of the differences and limitations between physical and virtual appliances, configure the virtual appliance from scratch to ensure proper functionality. When you migrate a virtual SPS to another one, you can export and import the configuration.

a. Click **Browse** and select the configuration file to import.

   NOTE: It is not possible to directly import a GPG-encrypted configuration into SPS, it has to be decrypted locally first.

b. Enter the passphrase used when the configuration was exported into the **Encryption passphrase** field.

   For details on restoring configuration from a configuration backup, see Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance on page 967.

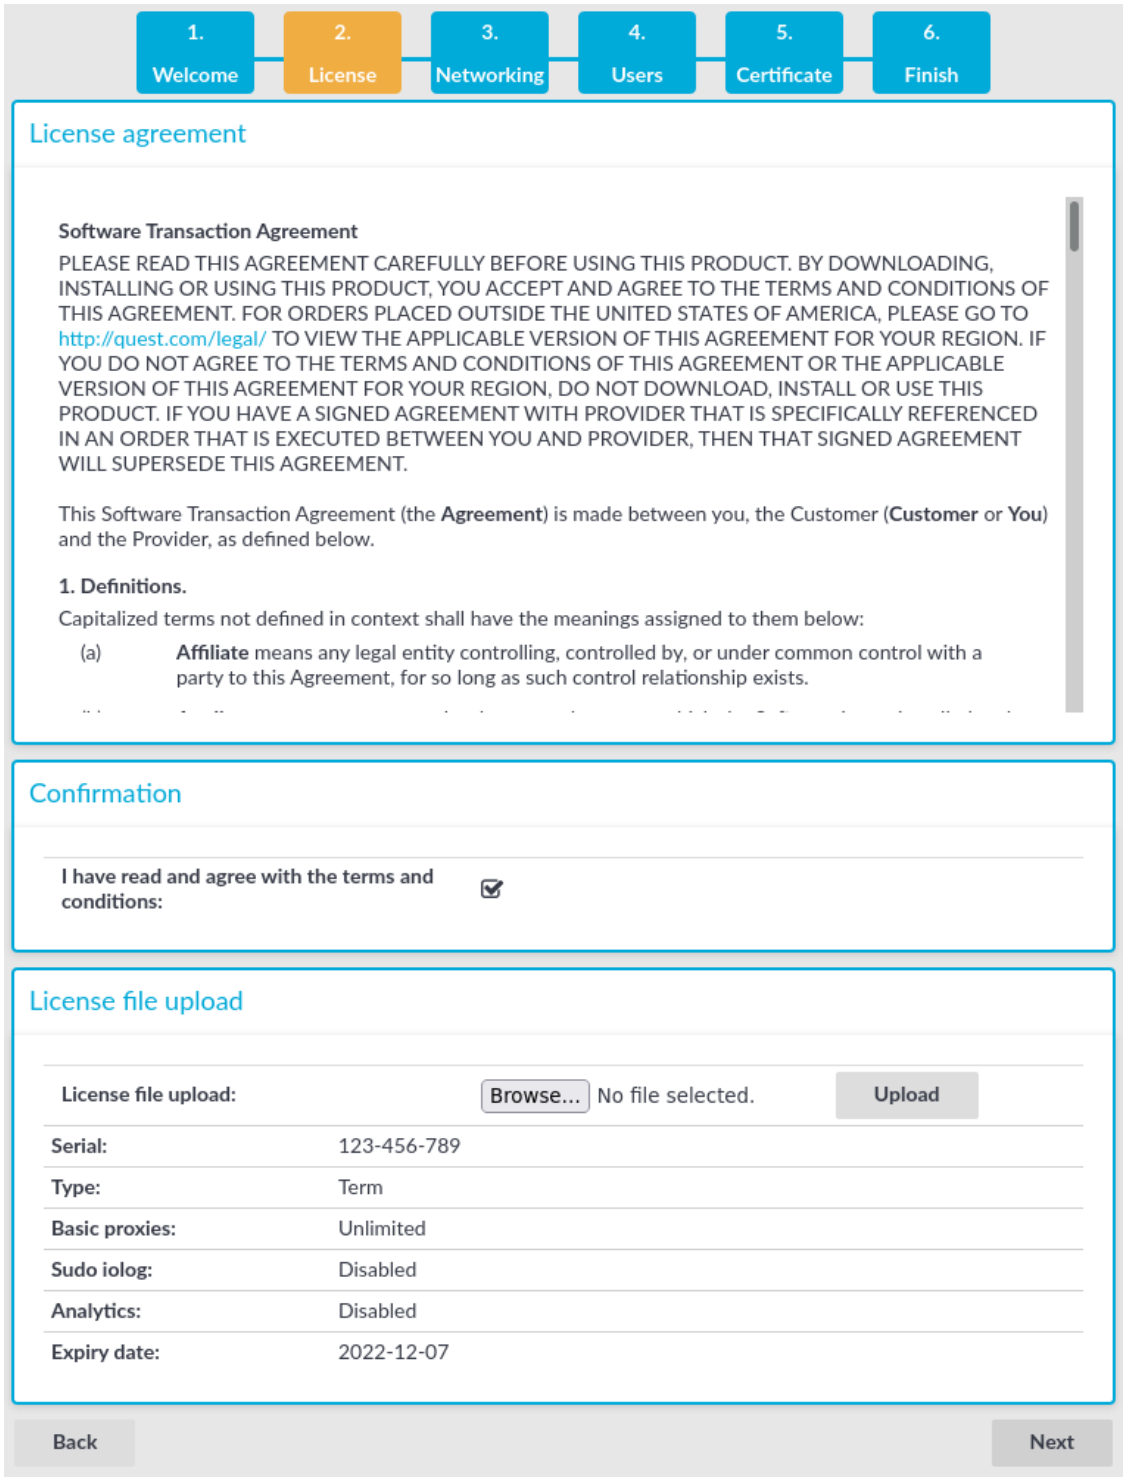
c. Click **Import**.

   CAUTION:

   If you use the Import function to copy a configuration from one SPS to another, do not forget to configure the IP addresses of the second SPS device. Having two devices with identical IP addresses on the same network leads to errors.

3. Accept the End User License Agreement and install the SPS license.

Figure 19: The EULA and the license key



a.

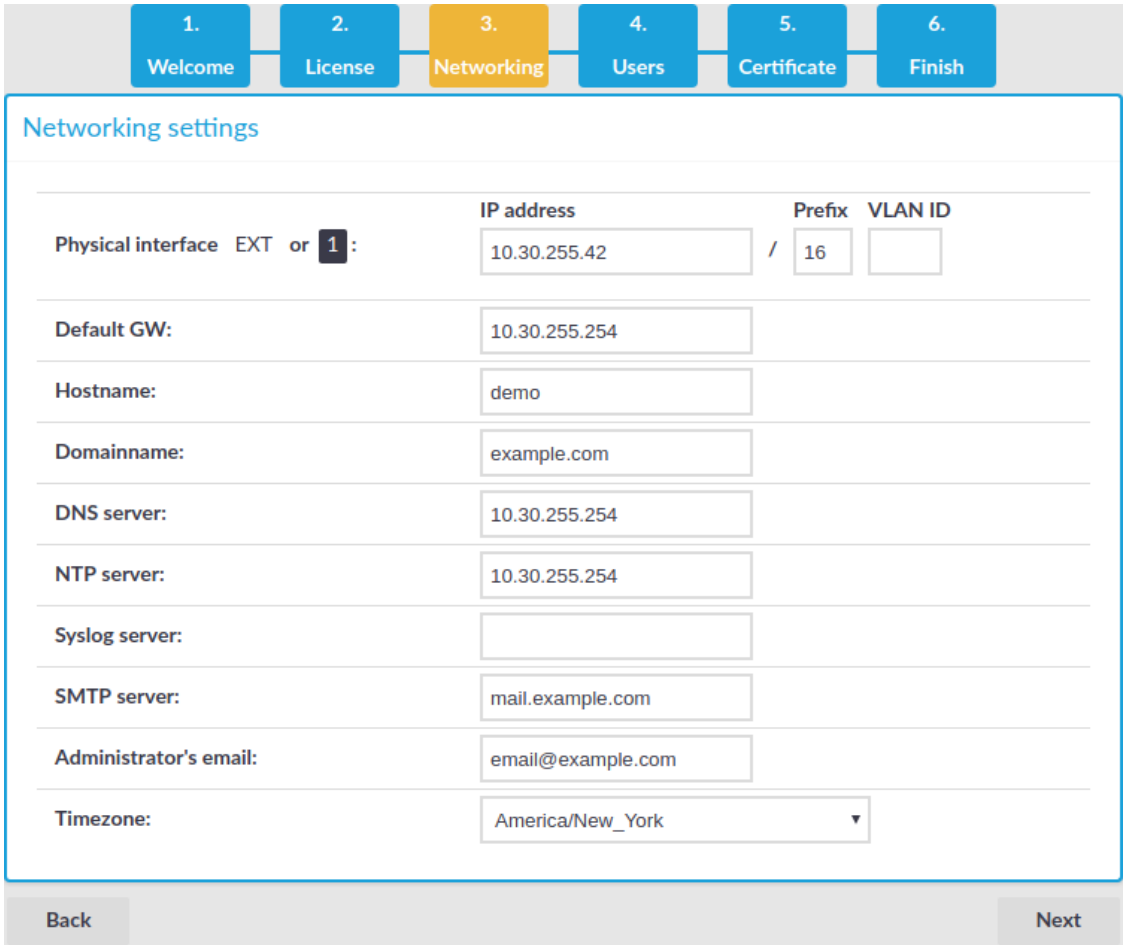
Read the End User License Agreement and select **I have read and agree with the terms and conditions**. The License Agreement covers both the traditional license, and subscription-based licensing as well. Clicking **I have read and agree with the terms and conditions** means that you accept the agreement that corresponds to the license you purchased. After the installation is complete, you can read the End User License Agreement at **Basic Settings** > **System** > **License**.

b. Click **Browse**, select the SPS license file received with SPS, then click **Upload**.

c. Click **Next**.

4. Configure networking. All settings can be modified later using the web interface of SPS.

Figure 20: Initial networking configuration



a. **Physical interface EXT or 1 — IP address**: The IP address of interface 1 (or EXT, for older hardware) of SPS (for example, `192.168.1.1`). You can choose the IP address from the range of the corresponding physical subnet. Clients will connect to this interface, so it must be accessible to them.

Use an IPv4 address.

NOTE: Do not use IP addresses that fall into the following ranges:

- `1.2.0.0/16` (reserved for communication between SPS cluster nodes)
- `127.0.0.0/8` (localhost IP addresses)

b. **Physical interface EXT or 1 — Prefix**: The IP prefix of the given range. For example, general class C networks have the /24 prefix.

c. **Physical interface EXT or 1 — VLAN ID**: (Optional) The VLAN ID of interface 1 (or EXT).

CAUTION:

Do not set the VLAN ID unless your network environment is already configured to use this VLAN. Otherwise, your SPS appliance will be unavailable using this interface.

d. **Default GW**: IP address of the default gateway.

Use an IPv4 address.

e. **Hostname**: Name of the machine running SPS (for example, `SPS`).

f. **Domainname**: Name of the domain used on the network.

g. **DNS server**: The IP address of the name server used for domain name resolution.

Use an IPv4 address.

h. **NTP server**: The IP address or the hostname of the NTP server.

Use an IPv4 address.

i. **Syslog server**: The IP address or the hostname of the syslog server.

Use an IPv4 address.

j. **SMTP server**: The IP address or the hostname of the SMTP server used to deliver e-mails.

Use an IPv4 address.

k. **Administrator's email**: E-mail address of the SPS administrator.

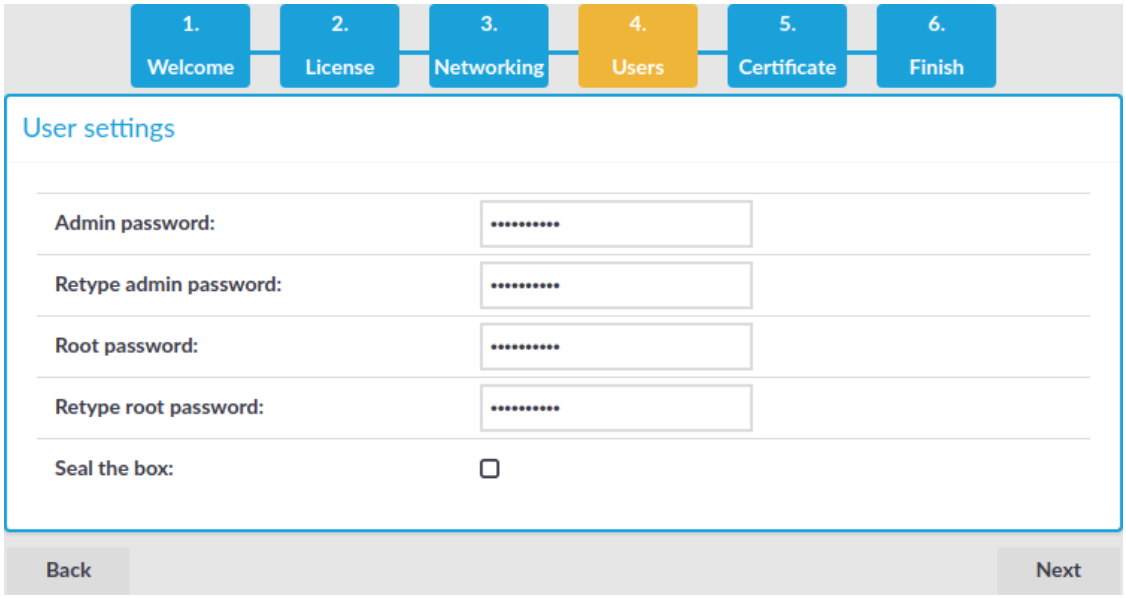
l. **Timezone**: The timezone where SPS is located.

m. **HA address**: The IP address of the High Availability (HA) interface. Leave this field on `auto` unless specifically requested by the support team.

n. Click **Next**.

5. Enter the passwords used to access SPS.

Figure 21: Passwords



NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

a. **Admin password**: The password of the `admin` user who can access the web interface of SPS.

b. **Root password**: The password of the `root` user, required to access SPS via SSH or from the local console.

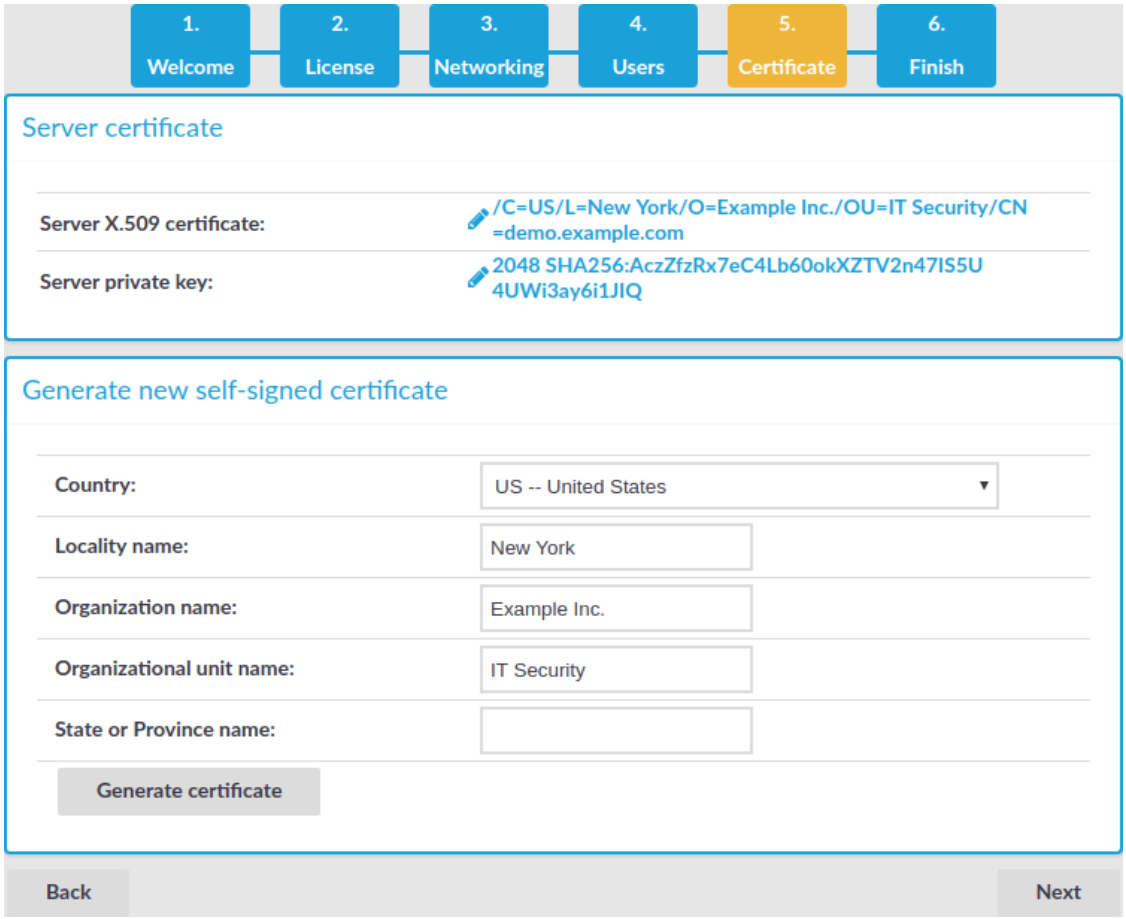
   NOTE: Accessing SPS using SSH is rarely needed, and One Identity recommends it only for advanced users for troubleshooting situations.

c. If you want to prevent users from accessing SPS remotely through SSH or changing the root password of SPS, select the **Seal the box** checkbox. You can activate sealed mode later from the web interface as well. For details, see *Sealed mode* in the *Administration Guide*.

d. Click **Next**.

6. Upload or create a certificate for the SPS web interface. This SSL certificate will be displayed by SPS to authenticate HTTPS connections to the web and the REST interface.

Figure 22: Creating a certificate for SPS



To create a self-signed certificate, fill the fields of the **Generate new self-signed certificate** section and click **Generate certificate**. The certificate will be self-signed by the SPS appliance. The hostname of SPS will be used as the issuer and common name.

a. **Country**: Select the country where SPS is located (for example, **HU-Hungary**).

b. **Locality name**: The city where SPS is located (for example, **Budapest**).

c. **Organization name**: The company that owns SPS (for example, **Example Inc.**).

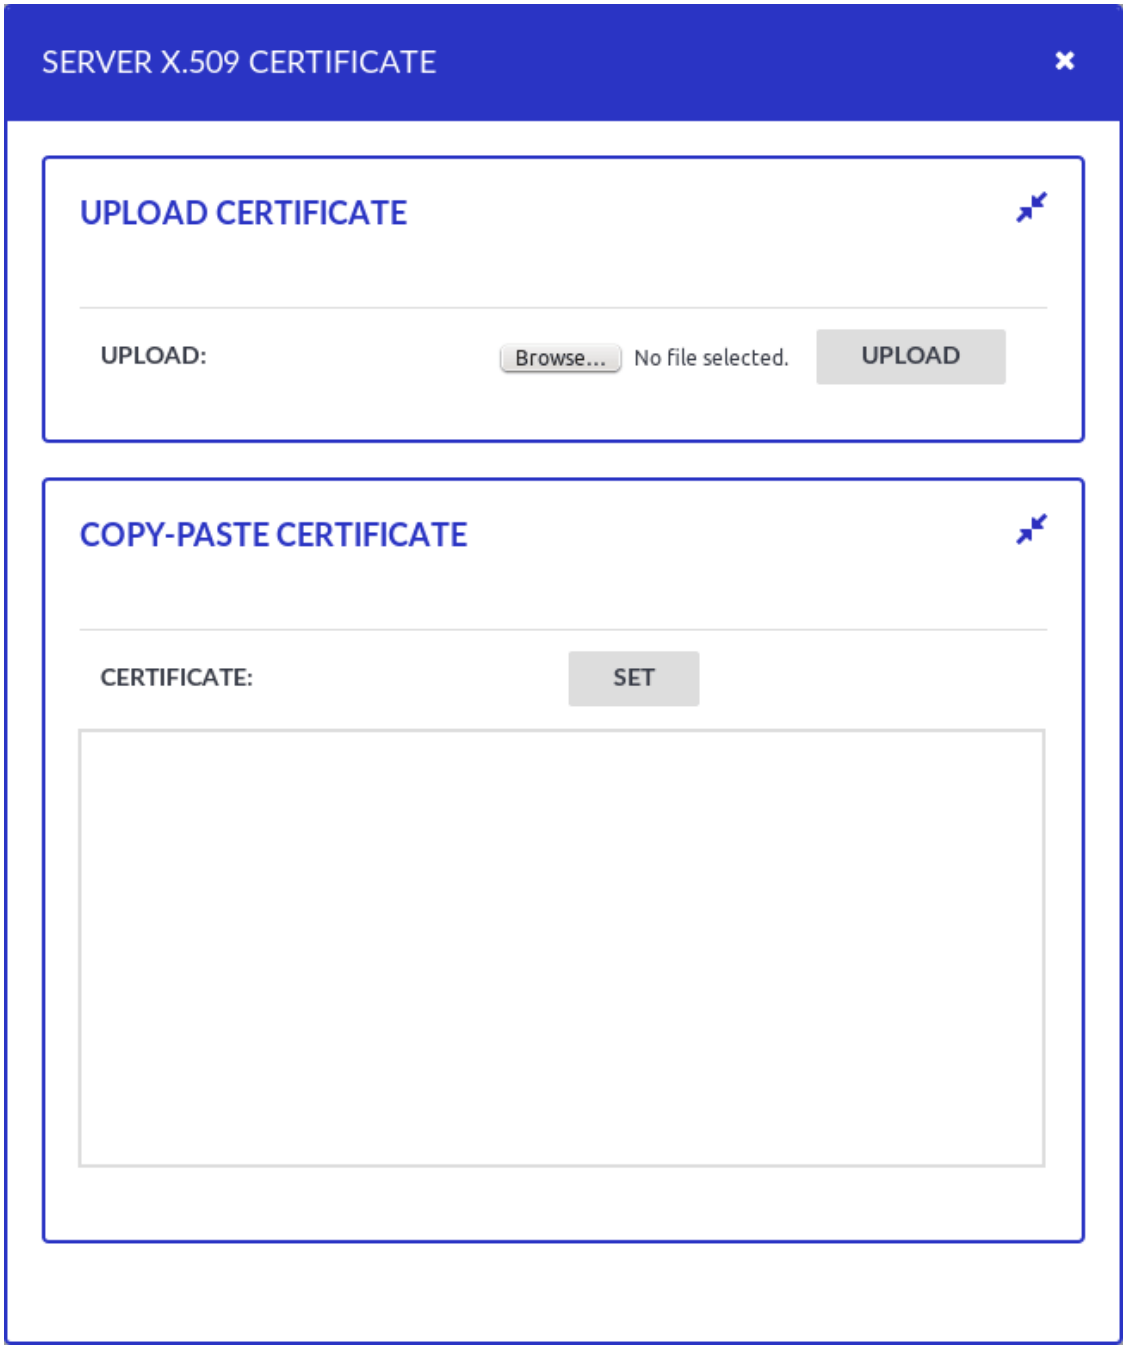
d. **Organizational unit name**: The division of the company that owns SPS (for example, **IT Security Department**).

e. **State or Province name**: The state or province where SPS is located.

f. Click **Generate certificate**.

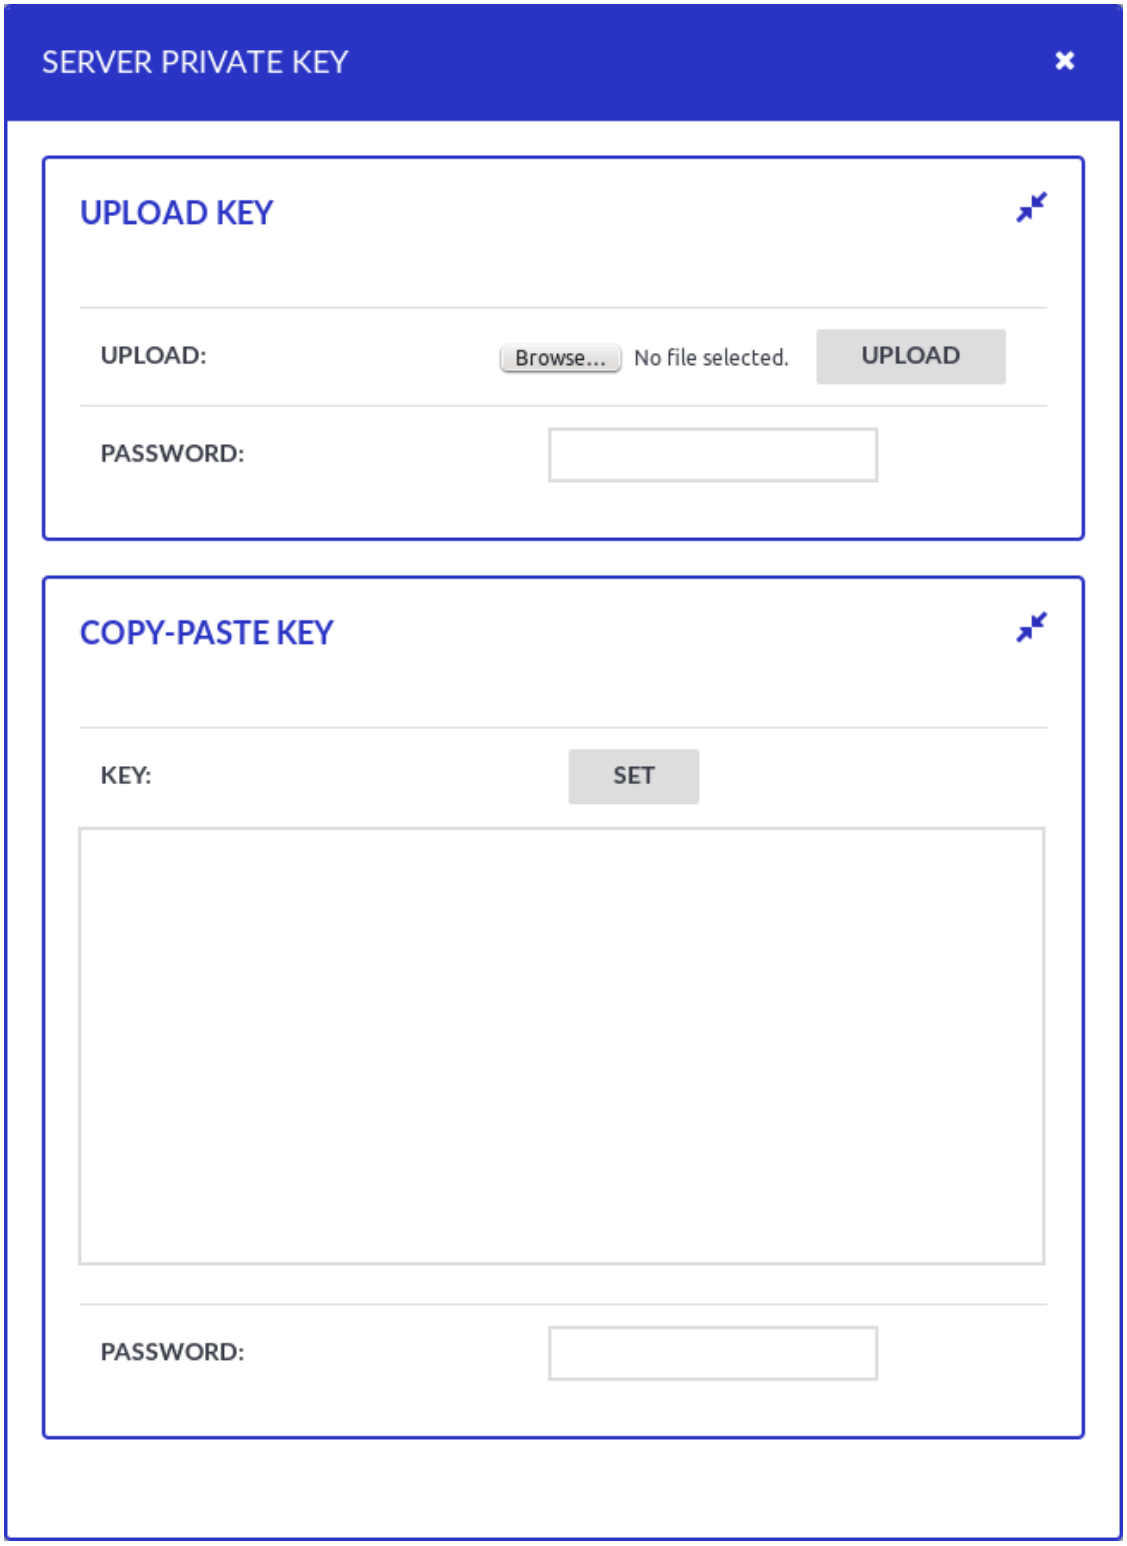
If you want to use a certificate that is signed by an external Certificate Authority, in the **Server X.509 certificate** field, click to upload the certificate.

Figure 23: Uploading a certificate for SPS



Then in the **Server private key** field click , upload the private key, and enter the password protecting the private key.

Figure 24: Uploading a private key



NOTE: SPS accepts private keys in PEM format, using RSA, DSA, and EC private key algorithms. Password-protected private keys are also supported.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

7. Review the data entered in the previous steps. This page also displays the certificate generated in the last step, the SSH RSA and Ed25519 keys of SPS, and information about the license file.

Figure 25: Review configuration data



If all information is correct, click **Finish**.

> **CAUTION:**
>
> The configuration takes effect immediately after clicking **Finish**. Incorrect network configuration data can render SPS unaccessible.

SPS is now accessible from the web interface through the IP address of interface 1 (or EXT).

8. Your browser is automatically redirected to the IP address set for interface 1 (or EXT) of SPS, where you can login to the web interface of SPS using the `admin` username and the password you set for this user in the Welcome Wizard.

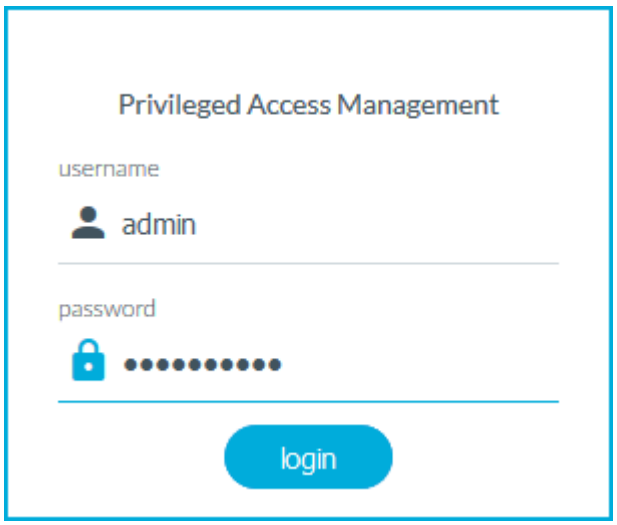
Logging in to One Identity Safeguard for Privileged Sessions (SPS) and configuring the first connection

After finishing the initial configuration of One Identity Safeguard for Privileged Sessions (SPS) using the Welcome Wizard, connections must be configured between the clients and the servers. SPS inspects only the connections that are configured from the web interface, all other connections are forwarded without any inspection.

To enable a simple SSH terminal or a Remote Desktop session over a transparent and a non-transparent connection

1. Login to SPS's web interface.

Figure 26: The first login



   a. Open the `https://IP-address-of-interface-1/` page from your browser to access the web interface of SPS. Replace the `IP-address-of-the-interface-1` string with the IP set for interface 1 in the **Networking settings** section of the Welcome Wizard (see Configuring interface 1) (for example, **`192.168.1.1`**).

   b. The certificate created in the **Certificate** section of the Welcome Wizard (see Creating the web interface certificate) is displayed. Accept it.

   c. Log in to the SPS web interface using the displayed login screen.

      - Enter `admin` into the **Login** field.
      - Enter the password set in the **Users** section of the Welcome Wizard (see Setting the administrator password) for the `admin` user into the **Password** field.

- Click **Login**. The main page of the SPS administration interface is displayed.

2. Configure a new transparent connection.

   a.
   - To configure an SSH connection, select **Traffic Controls** > **SSH** > **Connections** from the Main Menu. Only terminal sessions will be permitted.
   - To configure an RDP connection, click on the **Traffic Controls** > **RDP** > **Connections** from the Main Menu. Only basic Remote Desktop sessions will be permitted (no file-sharing).

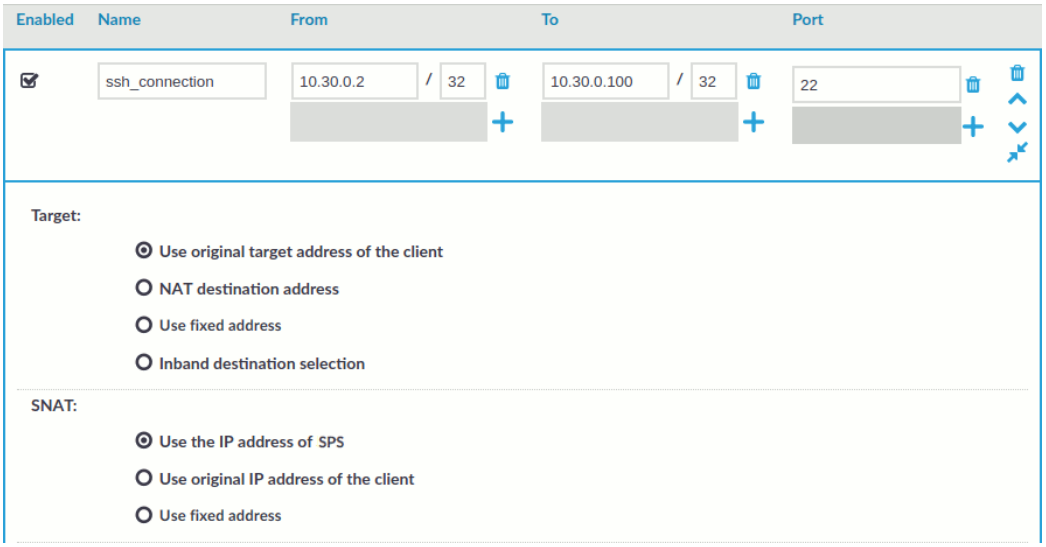
   b. Click the + icon on the right to create a new connection.

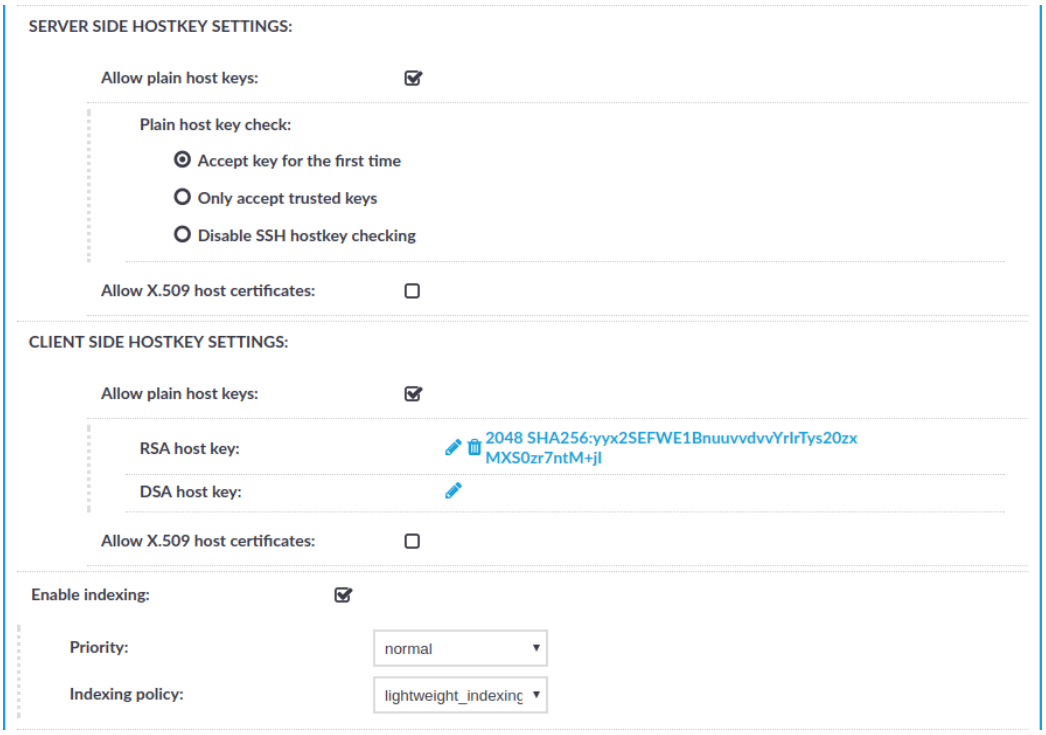
   c. Enter a name into the **Name** field that will identify the connection (for example, `admin-server-transparent`).

   TIP: One Identity recommends that you use descriptive names that give information about the connection (that is, they refer to the name of the accessible server, the allowed users, and so on).

   d. Enter the IP addresses defining the connection:

Figure 27: Traffic Controls > Protocol name > Connections — Configuring an SSH connection in transparent mode





- Enter the IP address of the client that will be permitted to access the server into the **From** field.

  You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).
- Enter the IP address of the server into the **To** field.

  You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).
- Enter the port number where the server is accepting connections into the **Port** field.

e. Select **Enable indexing**.

f. Click 

.

This connection allows any user from the client machine to connect to the specified server, but permits only terminal sessions — other SSH channels like TCP forwarding are disabled.

3. Configure a new non-transparent connection.

a.
- To configure an SSH connection, select **Traffic Controls** > **SSH** > **Connections** from the Main Menu. Only terminal sessions will be permitted.
- To configure an RDP connection, click on the **Traffic Controls** > **RDP** > **Connections** from the Main Menu. Only basic Remote Desktop sessions

will be permitted (that is, no clipboard or file-sharing).

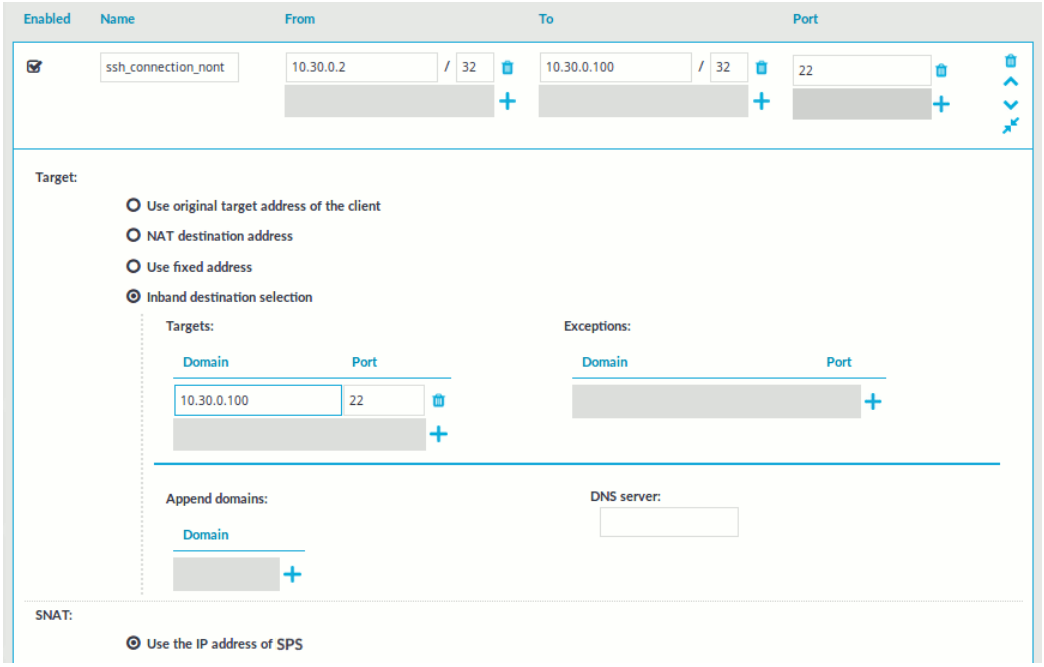
b. Click the + icon on the right to create a new connection.

c. Enter a name into the **Name** field that will identify the connection (for example, `admin-server-nontransparent`).

> TIP: One Identity recommends that you use descriptive names that give information about the connection (that is, they refer to the name of the accessible server, the allowed users, and so on).

d. Enter the IP addresses defining the connection:

Figure 28: Traffic Controls > Protocol name > Connections — Configuring an SSH connection in non-transparent mode



- Enter the IP address of the client that will be permitted to access the server into the **From** field.

  You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).

- Enter the IP address of SPS's physical interface 1 into the **To** field.

  You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).

- Enter a port number into the **Port** field.

- Enter the IP address of the server into the **Use fixed address** field of the **Target** section.

  You can use an IPv4 or an IPv6 address.

- Enter the port number where the server is accepting connections into the **Port** field of the **Target** section.

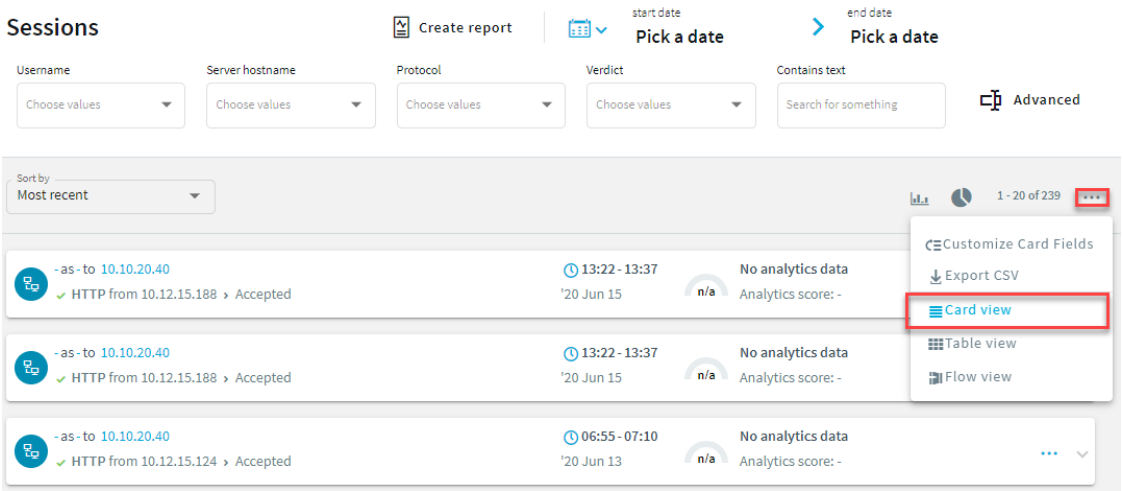
e. Select **Enable indexing**.

f. Click **Commit**.

This connection allows any user from the client machine to connect to the specified server, but permits only terminal sessions — other SSH channels like TCP forwarding are disabled.

4. Test the new configuration: try to initiate an SSH or and RDP connection from the client to the server.
   - For the transparent connection, use the IP address of the server (as configured in Configuring a connection in transparent mode).
   - For the non-transparent connection, use the IP address and port of SPS (as configured in Configuring a connection in non-transparent mode).
5. After successfully connecting to the server, do something in the connection, for example, execute a simple command in SSH (for example, `ls /tmp`), or launch an application in RDP (for example, the Windows Explorer), then disconnect from the server.
6. To access the Search interface, navigate to **Audit** > **Sessions**.

Figure 29: The Search interface



7. Find the session you want to replay on the **Sessions** page.

   For more information about search criteria and other search-related options, see Using the Sessions interface.

   For more information about the **session info** window and its contents, see Viewing session details.

8. Click ••• to display the details of the connection.

9. Click Start rendering to generate a video file from the audit trail you want to replay. Depending on the load of the indexer and the length and type of the audit trail, this can take several minutes.

   When the video is available, Start rendering changes to ▶ Play video and Delete video. You can use the Delete video button if you want to remove the generated video. After you remove the video file, the Start rendering button is available and you can use it to recreate the video file.

10. (Optional) If you have encrypted audit trails but the necessary certificates and private keys are not uploaded into your private keystore, you have to upload the keys first. After uploading them, click Unlock events. The Unlock events feature decrypts the encrypted upstream traffic elements. As a result, they will be displayed distributed in the generated video.

11. To replay the video, click ▶ Play video.

    The Player window opens.

12. Play the audit trail, and review your actions.

    For more information about audit trails, see sections Encrypting audit trails, Replaying audit trails in your browserand Replaying encrypted audit trails in your browser.

5

Basic settings

One Identity Safeguard for Privileged Sessions (SPS) is configured through the web interface. Configuration changes take effect automatically after clicking 

. Only the modifications of the current page or tab are activated — each page and tab must be committed separately.

- For details about the supported browsers, see Supported web browsers on page 98.
- For details on how to use the web interface of SPS, see The structure of the web interface on page 100.
- For details on how to configure the network interfaces, name resolution, and other networking-related settings, see Network settings on page 114.
- For details on how to control (for example reboot) SPS, upload a new firmware or license, export the current configuration, or stop the incoming syslog traffic, see Network settings on page 114.
- For details on how to set the system time and automatic time synchronization to an NTP server, see Configuring date and time on page 123.
- For details on how to configure where SNMP and e-mail alerts are sent, see System logging, SNMP and e-mail alerts on page 125.
- For details on how to configure system monitoring and alerts, see Configuring system monitoring on SPS on page 136.
- For details on how to configure data and configuration backups, see Data and configuration backups on page 145.
- For details on how to configure archiving, see Archiving on page 158.

Supported web browsers

Supported web browsers

Starting from version 6.13.0, SPS does not support Internet Explorer 11 (IE11). SPS version 6.12.0 and previous versions continue to support IE11.

Your browser must support:

- TLS-encrypted HTTPS connections with strong cipher algorithms
- JavaScript
- Cookies

SPS supports browsers as listed in the following table.

SPS version	IE11	Google Chrome	Safari	Mozzilla Firefox (latest version)	Microsoft EDGE	Microsoft EDGE Legacy
7.4.0	-	✓	✓	✓	✓	-
7.3.0	-	✓	✓	✓	✓	-
7.2.0	-	✓	✓	✓	✓	-
7.1.0	-	✓	✓	✓	✓	-
7.0 LTS	-	✓	✓	✓	✓	-
6.13.0	-	✓	✓	✓	✓	-
6.12.0	✓	✓	✓	✓	✓	-
6.11.0	✓	✓	✓	✓	✓	-
6.10.0	✓	✓	✓	✓	✓	-
6.9.0	✓	✓	✓	✓	✓	-
6.8.0	✓	✓	✓	✓	✓	-

Required applications and plugins

To use SPS, install and enable the following applications and plugins.

NOTE: To replay audit trails with SPS 6.9 or earlier versions with Internet Explorer 11 (IE11), install the Google WebM Video for Microsoft Internet Explorer plugin.

SPS version	JavaScipt	Google WebM Video for Microsoft Internet Explorer plugin
7.4.0	✓	-
7.3.0	✓	-
7.2.0	✓	-
7.1.0	✓	-
7.0 LTS	✓	-
6.13.0	✓	-
6.12.0	✓	-
6.11.0	✓	-
6.10.0	✓	-
6.9.0	✓	Required for IE11
6.8.0	✓	Required for IE11

Phased out browsers

SPS does not support anymore the browsers listed in the following table.

Browser	Phased out in
IE10	SPS 4 F3
IE9	SPS 4 F3

Starting from version 4 F3, SPS does not support Internet Explorer 9 and 10, because the official support for them ended in January, 2016.

SPS web UI

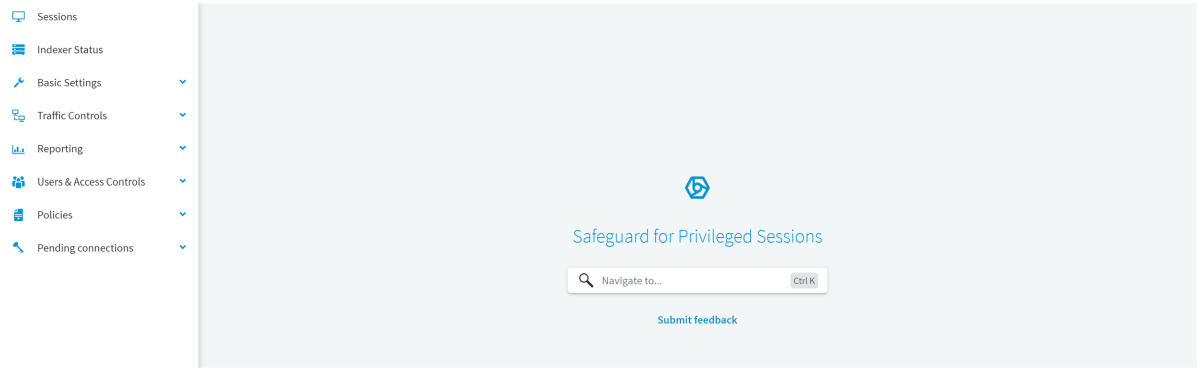
Opening the web UI of SPS in multiple browser windows or tabs is not supported.

NOTE: The minimum recommended screen resolution for viewing One Identity Safeguard for Privileged Sessions's (SPS's) web interface is 1366 x 768 pixels on a 14-inch widescreen (standard 16:9 ratio) laptop screen. Screen sizes and screen resolutions that are equal to or are above these values will guarantee an optimal display of the web interface.

The structure of the web interface

NOTE: The minimum recommended screen resolution for viewing One Identity Safeguard for Privileged Sessions's (SPS's) web interface is 1366 x 768 pixels on a 14-inch widescreen (standard 16:9 ratio) laptop screen. Screen sizes and screen resolutions that are equal to or are above these values will guarantee an optimal display of the web interface.

Figure 30: Structure of the web interface



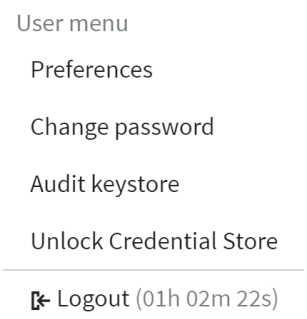
The web interface consists of the following main sections:

Main menu

Each menu item displays its options in the main workspace on one or more tabs. Click a **Main menu** item to display the list of tabs available under that particular menu item.

User menu

Figure 31: User menu

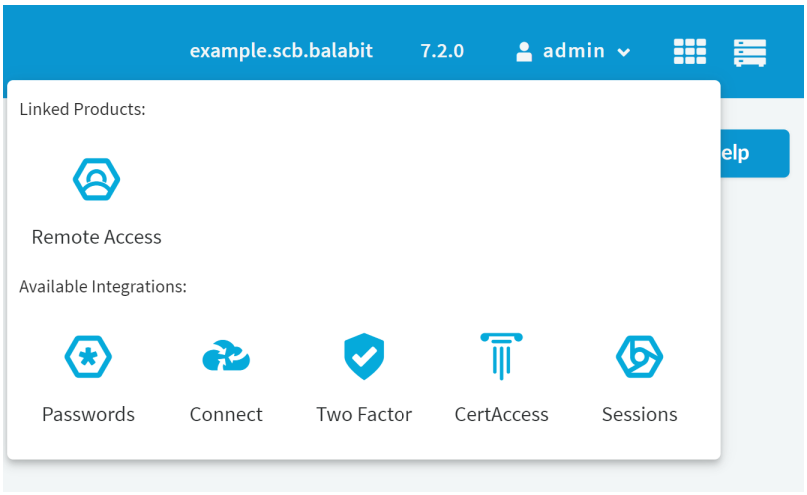


The **User menu** allows you to:

- Disable confirmation dialogs and tooltips using the **Preferences** option. For more information, see section Preferences.
- Change your SPS password using the **Change password** option. For more information, see section Change password.
- Upload your master password and permanent or temporary keys using the **Audit keystore** option. For more information, see Audit keystore, Viewing encrypted screenshots, and Replaying encrypted audit trails in your browser.
- Unlock a Credential Store.
- Log out.

App switcher

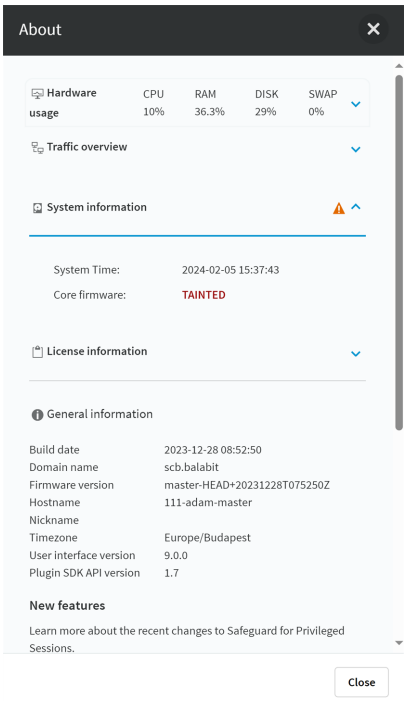
Figure 32: App switcher



After setting up a federated login method between SPS and a linked product, you can seamlessly switch between SPS and the linked product by selecting the app switcher icon.

About page

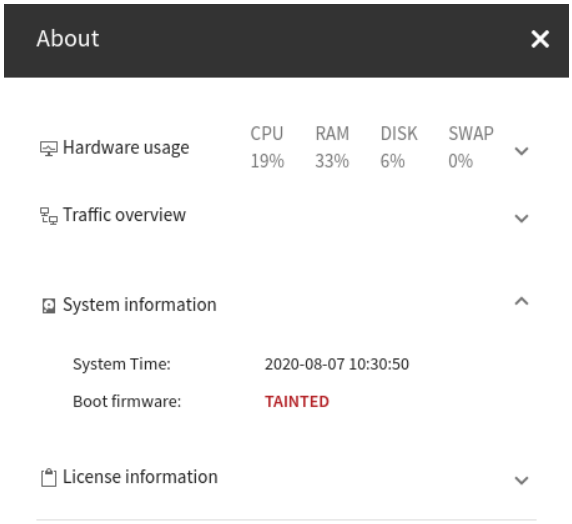
Figure 33: About page



Displays accessibility and system health information about SPS, including the following:

- **Hardware usage**: CPU, memory, hard disk, and swap use. Expand to see more details or navigate to **Basic Settings**> **Dashboard** for detailed reports.
- **Traffic overview**: The number of active **HTTP, ICA, MSSQL, RDP, SSH, TELNET**, and **VNC** connections. For **HTTP**, Traffic overview displays the number of active HTTP sessions.
- **System information**: Shows the system date, the system time, and the status of the core and boot firmware.
    - **Corrupted**: The firmware integrity check failed. If a firmware is shown as corrupted, contact our Support Team.
    - **Tainted**: It indicates that you have modified a file of the firmware locally. If you have modified a local file unintentionally, contact our Support Team.

Figure 34: Boot firmware - Tainted



- **General information** such as, current timezone or user interface version, and so on, as well as, links related to new features, help or feedback.

Context-sensitive help

SPS offers context-sensitive help, which is used to display information about the user interface relative to the task a user performs.

There are different levels of context sensitivity that have been implemented in SPS.

- **Screen-level help**

  When available, SPS opens the help topic for that screen. Instead of having the user browse through the help system to find the right topic, SPS can quickly and directly display the topic that corresponds to the screen.

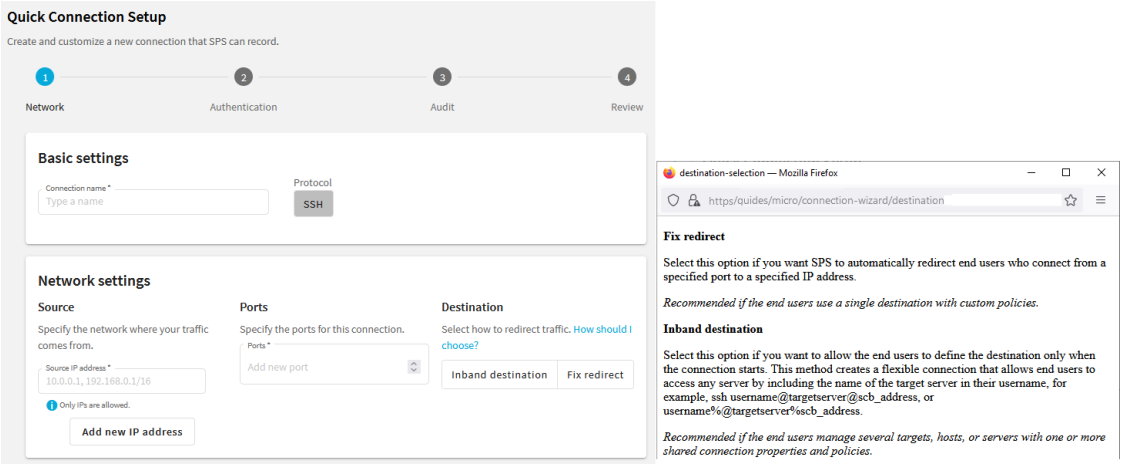
  To open a screen-level help, click **I need help**, when available.

- **Field-level help**

  When available, field-level help provides help text detailing the purpose and function of a field.

  The following is an example of a field-level help where you can click **How should I choose?** and the help opens with more details about making the relevant destination settings.

Figure 35: Quick Connection Setup — Example of a field-level help



Elements of the main workspace

The main workspace displays the configuration settings related to the selected **Main menu** item grouped into one or more submenus. Related parameters of a submenu are organized into labeled groups or sections, marked with blue outline

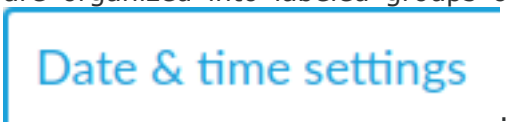
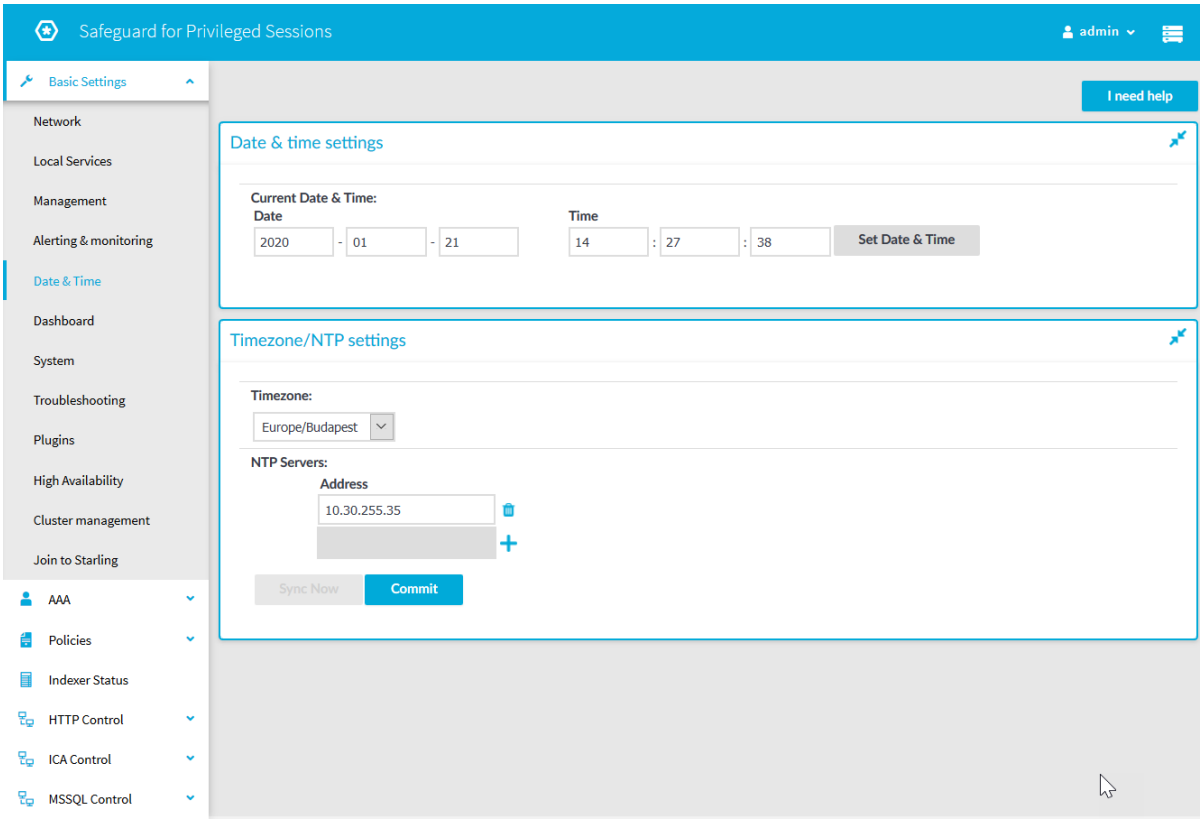


Figure 36: Main workspace

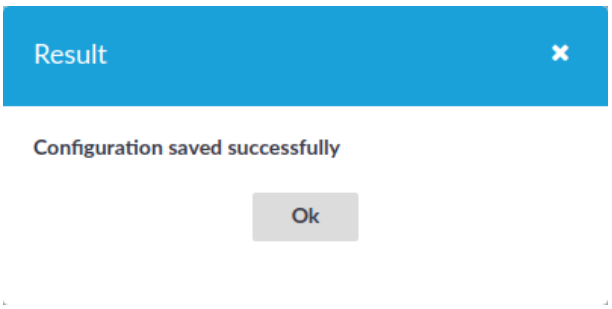


- Commit Each page includes one or more blue action buttons. The most common action button is Commit, which saves and activates the changes of the page.
- /*Show/Hide details*: Displays or hides additional configuration settings and options.
- *Create entry*: Create a new row or entry (for example an IP address or a policy).
- *Delete entry*: Delete a row or an entry (for example an IP address or a policy).
- *Modify entries or upload files*: Edit an entry (for example a host key, a list, and so on), or upload a file (for example a private key). These actions open a pop-up window where the actual modification can be performed.
- , *Position an item in a list*: Modify the order of items in a list. The order of items in a list (for example the order of connections, permitted channels in a channel policy, and so on) is important because when One Identity Safeguard for Privileged Sessions (SPS) is looking for a policy, it evaluates the list from top to down, and selects the first item completely matching the search criteria. For example, when a client initiates a connection to a protected server, SPS selects the first connection policy

matching the client's IP address, the server's IP address, and the target port (the From, To, and Port fields of the connection).

Message window: This pop-up window displays the responses of SPS to the user's actions, for example **Configuration saved successfully**. Error messages are also displayed here. All messages are included in the system log. For detailed system logs (including message history), see the **Troubleshooting** tab of the **Basic Settings**. To make the window appear only for failed actions, navigate to **User menu** > **Preferences** and enable the **Autoclose successful commit messages** option.

Figure 37: Message window



Navigating on the SPS UI

A quick navigation and search function has been introduced to improve the navigation experience and to provide a quick and efficient way to navigate to an SPS UI page.

This option has a keyboard-first design, aiming primarily at users who use the keyboard rather than the mouse cursor. However, you can use the mouse cursor as well.

To navigate on the SPS UI

1. To access the search field of the quick navigation function, select one of the following options:
    - Navigate to the SPS Home page by clicking on the logo.
    - In a Windows or Linux environment, press Ctrl+K.
    - On Mac OS, press Command+K.
2. (Optional) Insert question mark (?) in the search field to display the **Help** menu at the bottom of the search window.
3. Type a keyword in the search field. As you type, a drop-down dynamically provides you a list of search results.
4. (Optional) To use autocomplete to complete the keywords while typing, press Tab.
5. (Optional) To clear the content of the search field, press Esc.
6. (Optional) To close the search window, press Esc.

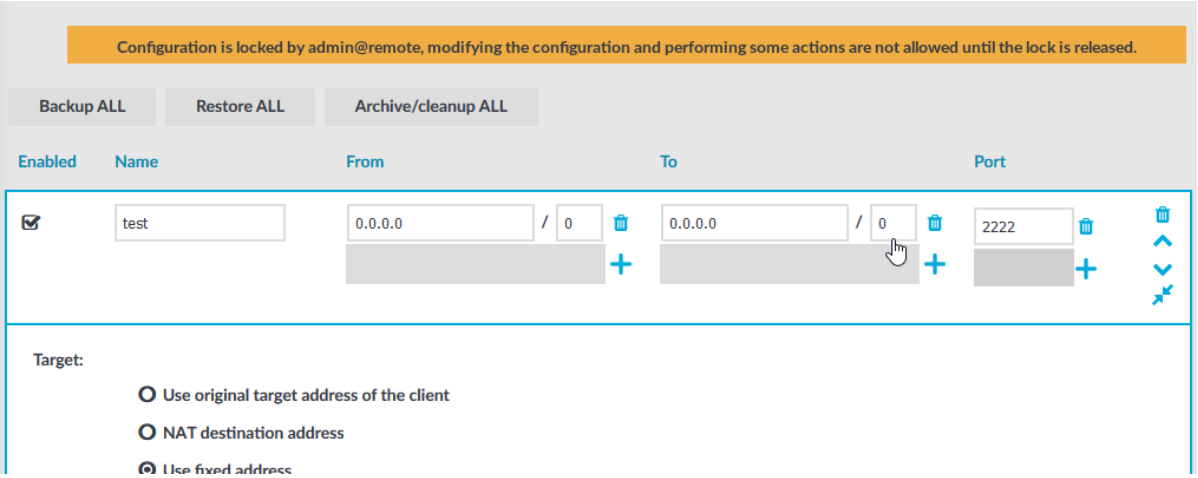
7. To navigate between the search results, use the up and down arrow buttons.
8. Press Enter to navigate to the page that corresponds to the highlighted search result.

Multiple users and locking

Multiple administrators can access the One Identity Safeguard for Privileged Sessions (SPS) web interface simultaneously, but only one of them can modify the configuration. This means that the configuration of SPS is automatically locked when the first administrator who can modify the configuration opens a configuration page (for example the **Basic Settings** or the **Users & Access Control** menu).

The warning message displays the username of the administrator locking the configuration as shown in the image below:

Figure 38: Configuration lock by remote administrator



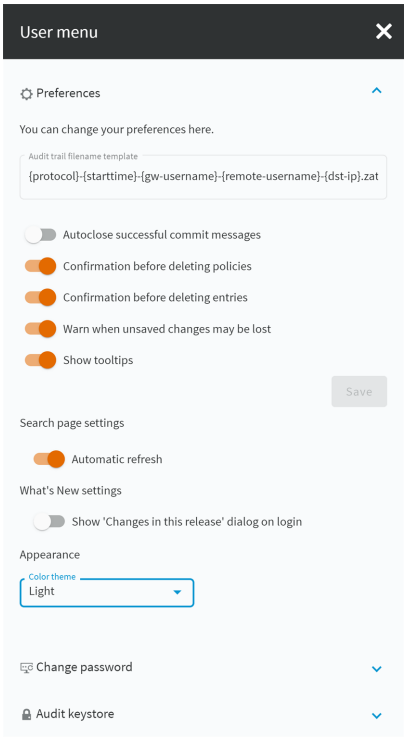
Other administrators can continue as read-only but must wait until the locking administrator navigates to an SPS page that does not require locking, the administrator logs out, or the session of the administrator times out. However, it is possible to access the **Sessions** and **Reporting** menus, and to perform gateway authentication and 4-eyes authorization or browse the configuration with only View rights (for details, see Managing user rights and user groups on page 361).

NOTE: If an administrator logs in to SPS using the local console or a remote SSH connection, the configuration is also locked. Inactive local and SSH connections timeout just like web connections. For details, see Accessing the One Identity Safeguard for Privileged Sessions (SPS) console on page 442.

Preferences

To configure your preferences about the web interface, navigate to **User Menu** > **Preferences**.

Figure 39: User Menu > Preferences



- **Audit trail filename template**:

  To change the filename of the audit trails, navigate to **User menu** > **Preferences** and change the **Audit trail filename template**. The default template is `{protocol}-{starttime}-{gw-username}-{remote-username}-{dst-ip}.zat`. The template can include anything, the keys (inside {} brackets) are replaced with their actual values. These keys are the following:

  - `connection-policy`: The connection policy
  - `dst-ip`: Destination IP address
  - `dst-port`: Destination port
  - `gw-username`: Gateway username
  - `protocol`: Protocol
  - `remote-username`: Remote username
  - `session-id`: Session ID
  - `src-ip`: Source IP address
  - `starttime`: Start time of the session
- **Autoclose successful commit messages**: General confirmation windows will not appear. (For example, **Configuration saved successfully** that appears after successfully committing a change). As a result, pop-up windows appear only for failed actions or errors.

- **Confirmation before deleting policies**: Display a pop-up window when you attempt to delete policies to prevent deleting policies accidentally.
- **Confirmation before deleting entries**: Display a pop-up window when you attempt to delete entries to prevent deleting entries accidentally.
- **Warn when unsaved changes may be lost**: Display a pop-up window to warn when you navigate to another window without committing your changes to prevent losing unsaved changes.
- **Show tooltips**: Display tooltips for user interface elements to help using the product.

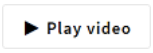
Search page settings

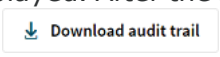
- **Automatic refresh**:If you select this option, the content of the details view on the **Sessions** page is refreshed automatically to provide up-to-date information about the sessions. On the **Sessions** page, you can open the details view of a session by clicking the ••• icon. **Automatic refresh** refreshes the content of all the tabs (**Overview**, **Details**, **Timeline**, **Analytics**) in the details view.

  If you select **Automatic refresh**, for example, if there is an active connection, the list of events is refreshed dynamically and the buttons are displayed according to the available options for a given session. For example, while a session is active, 

is displayed, and

and

are displayed. After the session is closed, the buttons dynamically change to 

and 

.

What's New settings

- **Show 'Changes in this release' dialog on login**: If you select this option, the **Changes in this release** dialog pops up every time you log in to SPS.

Appearance

- **Light**: Light mode is the default setting of the user interface. If this option is enabled, the user interface appears in light mode.
- **Dark (High contrast)**: If you select this option, the user interface appears in dark, high contrast mode.
- **Auto**: If you select this option, SPS detects the dark or light theme setting of your operating system, and displays the SPS login page according to the theme setting of your operating system.

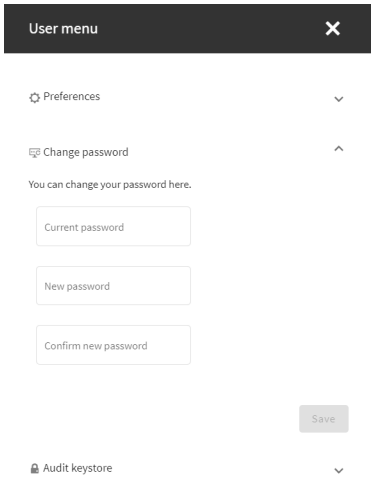
Change password

If you are a local user or an administrator, you can change your password for the web interface using the **User menu** > **Change password** option.

To change the password for SPS

1. Navigate to **User menu** > **Change password**.

Figure 40: User menu > Change password — Changing the password



2. Enter your current password in the **Current password** field.
3. Create a new, strong password and enter it in the **New password** field.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

   NOTE: If possible, use a random password generator.

   To create a strong password:

   - Use special characters
   - Use long passwords
   - Mix uppercase and lowercase letters

   For strong passwords, do not use:

   - Personal information in the passwords
   - Sequential letters or numbers
   - The word `password` as the password
   - Keyboard paths (for example, `qwerty`)

4. Enter your new password again in the **Confirm new password** field.
5. Click **Save**.

Audit keystore

To replay encrypted audit trails in your browser and to view encrypted screenshots, upload the necessary private keys to your audit keystore. In the audit keystore, only private keys are stored.

NOTE: Previously, the audit keystore was used to store certificates as well as private keys. From SPS version 6.10 and onwards, you must upload the certificates to **Basic settings** > **Local services** > **Indexer service**. For more information on how to add certificates, see Configuring the internal indexer.

Only RSA keys (in PEM-encoded X.509 certificates) can be uploaded to the private keystore.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

NOTE: Certificates are used as a container and delivery mechanism. For encryption and decryption, only the keys are used.

Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates (for example, do not use the certificate of the One Identity Safeguard for Privileged Sessions (SPS) webserver to encrypt audit trails, or the same keypair for signing and encrypting data).

To manage your audit keystore, see the following sections:

- Adding the first private key to your audit keystore
- Adding further private keys to your audit keystore
- Unlocking your audit keystore
- Deleting a private key from your audit keystore

Adding the first private key to your audit keystore

To replay encrypted audit trails in your browser and to view encrypted screenshots, upload the necessary private keys to your audit keystore. In the audit keystore, only private keys are stored.

For more information on the supported key format and the certificates that correspond to the private keys, see Audit keystore.

NOTE: The private keys are stored locally, in your browser.

To add the first key to your audit keystore

1. Navigate to **User menu** > **Audit keystore**.

If you want to open an encrypted audit trail or screenshot from the **Sessions** interface, but you have not added the corresponding private keys yet to your audit keystore, a dialog will take you to the **Audit Keystore** option.

2. Click **Add new key**.
    a. In the **Name** field, enter a name for the key.
    b. Click **Browse key**, select the file containing the key in PEM format, and click **Open**.

       Alternatively, you can also copy-paste or drag your key into the **Key content** field.
    c. (Optional) If you add a private key that is encrypted, an additional field, the **Key password** field is displayed. In the **Key password** field, enter the password for the encrypted key.
    d. Click **Add key**.
3. In the **Create master password** dialog, add a master password.

   The private key is stored in your audit keystore that is protected by the master password that you created.
4. (Optional) To lock your audit keystore, click **Lock keystore**.

   If you lock your audit keystore, you protect your private keys from unauthorized use. Your private keys can be used to decrypt content only if you unlock your audit keystore.

Adding further private keys to your audit keystore

This section describes how to add new private keys to your audit keystore.

NOTE: The private keys are stored locally, in your browser.

To add further keys to your audit keystore

1. Navigate to **User menu** > **Audit keystore**.
2. Click **Unlock private keystore**.

   Enter your master password and click **Unlock keystore**.

   The audit keystore is unlocked.

   If you forgot your master password, see section Unlocking your audit keystore.
3. Click **Add new key**.
    a. In the **Name** field, enter a name for the key.
    b. Click **Browse key**, select the file containing the key in PEM format, and click **Open**.

       Alternatively, you can also copy-paste or drag your key into the **Key content** field.

c. (Optional) If you add a private key that is encrypted, an additional field, the **Key password** field is displayed. In the **Key password** field, enter the password for the encrypted key.

d. Click **Add key**.

The private key is stored in your audit keystore that is protected by the master password that you created.

4. (Optional) To lock your audit keystore, click **Lock keystore**.

   If you lock your audit keystore, you protect your private keys from unauthorized use. Your private keys can be used to decrypt content only if you unlock your audit keystore.

Unlocking your audit keystore

This section provides information on:

- How to unlock your audit keystore
- How to reset your audit keystore if you forgot your master password

NOTE: The private keys are stored locally, in your browser.

In the audit keystore, the master password protects your private keys from unauthorized use, for example, if you share a computer with anyone.

To use the private keys that are stored in your audit keystore to decrypt audit items, you must unlock your audit keystore by providing your master password. After providing your master password, your audit keystore remains unlocked for the duration of your session, or until you click **Lock keystore** on **User menu** > **Audit keystore**.

Unlocking your audit keystore

To unlock your audit keystore

1. Navigate to **User menu** > **Audit keystore**.
2. Click **Unlock private keystore**.

   The **Unlock keystore** dialog is displayed.

   Enter your master password and click **Unlock keystore**.

   The audit keystore is unlocked.

   You can add new keys or manage your uploaded private keys.

Resetting your audit keystore

To reset your audit keystore if you forgot your master password

NOTE: The master password cannot be changed, but if you forget your master password, you can reset your audit keystore. If you reset your audit keystore, you must upload your private keys again.

1. Navigate to **User menu** > **Audit keystore**.
2. If you forgot your master password, click **Forget password?**.
3. In the **Forget password?** dialog, click **Reset keystore**.
4. Click **Add new key** and the **Create master password** dialog is displayed, where you can add a new master password.
5. Upload the necessary private keys again to your audit keystore.

Result

The audit keystore is unlocked and it remains open for the duration of your session or until you click **Lock keystore**.

Deleting a private key from your audit keystore

This section provides information on how to delete a key in your audit keystore.

NOTE: The private keys are stored locally, in your browser.

To delete a key from your audit keystore

1. Navigate to **User menu** > **Audit keystore**.
2. Click **Unlock private keystore**.

   The **Unlock keystore** dialog is displayed.

   Enter your master password and click **Unlock keystore**.

   The audit keystore is unlocked.
3. Click  next to the private key that you want to delete.

   NOTE: If you delete the last private key from your audit keystore, the audit keystore is reset and next time you add a private key to your audit keystore, you must define a master password again.

Result

The private key is deleted from your audit keystore.

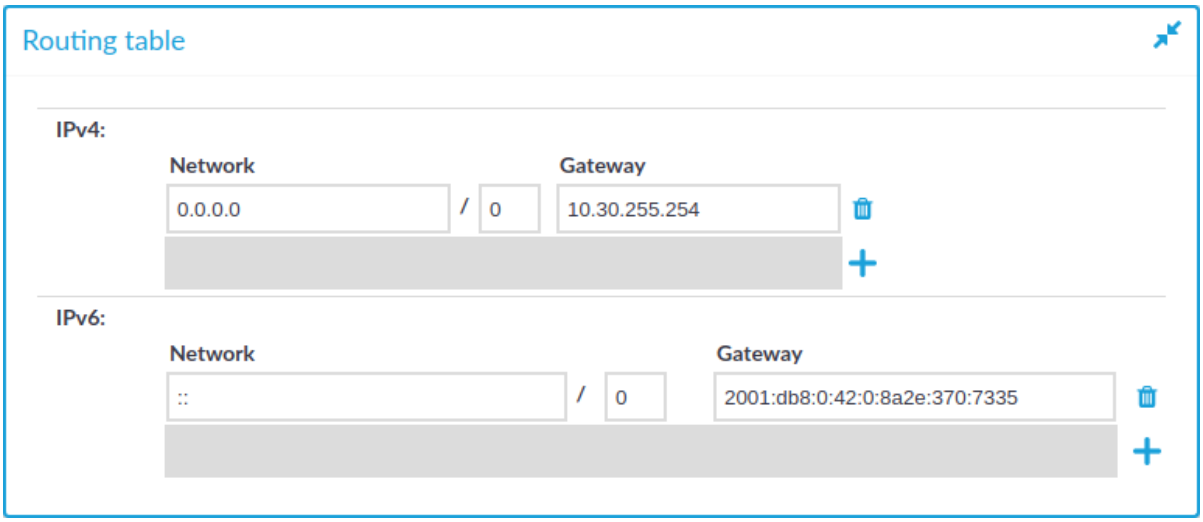
Network settings

The **Basic Settings** > **Network** tab contains the network interface and naming settings of One Identity Safeguard for Privileged Sessions (SPS).

Routing table

Routing table

Figure 41: Basic Settings > Network > Routing table



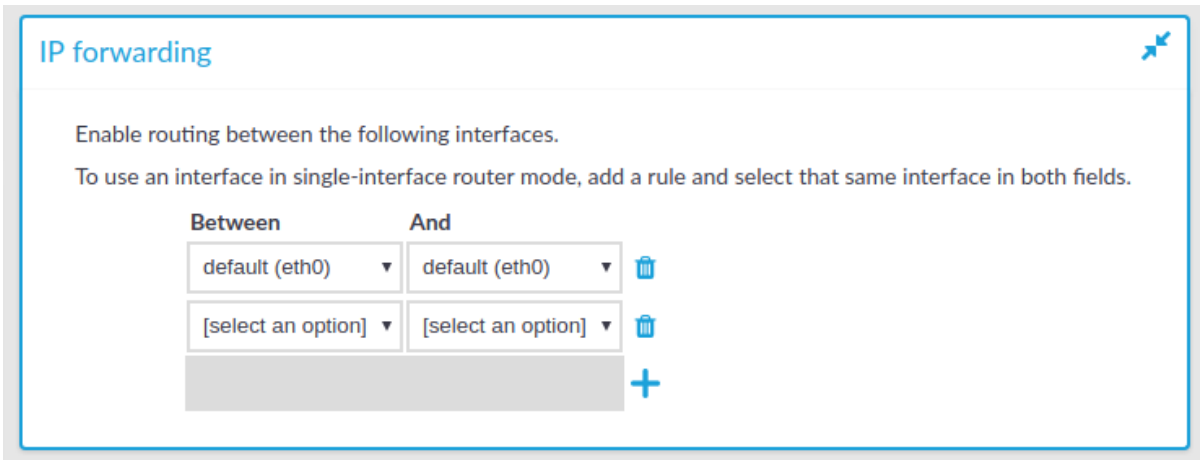
When sending a packet to a remote network, SPS consults the routing table to determine the path it should be sent. If there is no information in the routing table then the packet is sent to the default gateway. Use the routing table to define static routes to specific hosts or networks. You have to use the routing table if SPS interfaces are connected to multiple subnets.

Click the + and 🗑 icons to add new routes or delete existing ones. A route means that messages sent to the **Address**/**Netmask** network should be delivered to **Gateway**.

For more information, see Configuring the routing table.

IP forwarding

Figure 42: Basic Settings > Network > IP forwarding

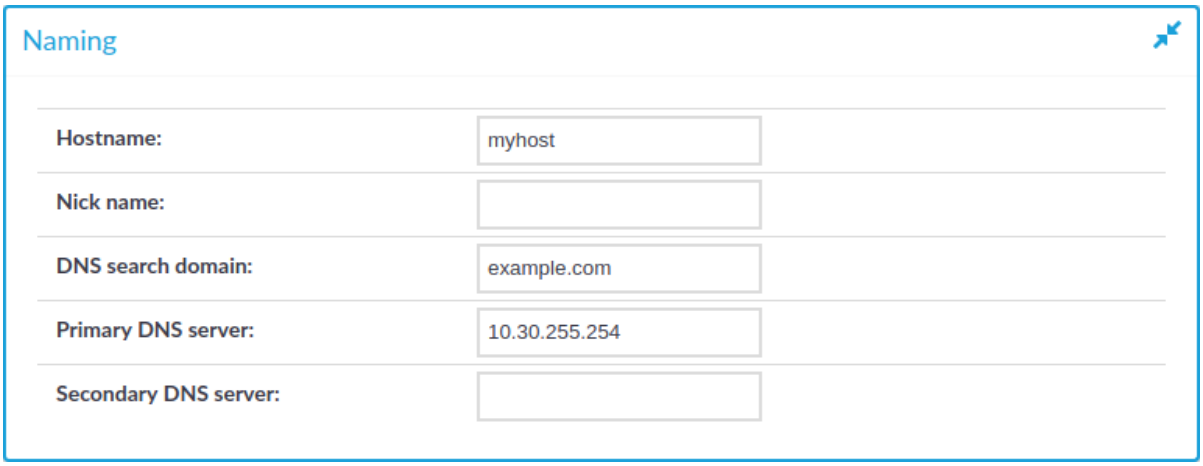


You can enable routing between logical interfaces, which allows you to direct uncontrolled traffic through SPS. For more information, see Routing uncontrolled traffic between logical interfaces.

To mimic the functionality of the deprecated Router mode, configure a logical interface for each physical interface you want to connect, and enable IP forwarding between them.

Naming

Figure 43: Basic Settings > Network > Naming

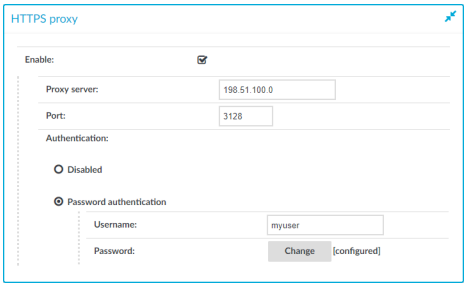


- **Hostname**: Name of the machine running SPS.
- **Nick name**: The nickname of SPS. Use it to distinguish the devices. It is displayed in the core and boot login shells.
- **DNS search domain**: Name of the domain used on the network. When resolving the domain names of the audited connections, SPS will use this domain to resolve the target hostname if the appended domain entry of a target address is empty.
- **Primary DNS server**: IP address of the name server used for domain name resolution.
- **Secondary DNS server**: IP address of the name server used for domain name resolution if the primary server is unaccessible.

HTTPS proxy

The **HTTPS proxy** settings must be configured if your company policies do not allow devices to connect directly to the web. Once configured, SPS uses the configured proxy server for outbound web requests to external integrated services, such as Join to Starling or SPS plugins.

Figure 44: Basic Settings > Network > HTTPS proxy



- **Proxy server**: The IP address or DNS name of the proxy server.
- **Port**: The IP address or DNS name of the proxy server.

  NOTE:

  If different ports are specified in the **Proxy server** and the **Port** field, the **Port** field takes precedence.
- **Username**: The user name used to connect to the proxy server.

  NOTE:

  The username and password are only required if your proxy server requires them to be specified.
- **Password**: The password required to connect to the proxy server.

  NOTE:

  The username and password are only required if your proxy server requires them to be specified.

Configuring user and administrator login addresses

You can configure two separate login addresses for accessing the web interface of One Identity Safeguard for Privileged Sessions (SPS):

- **Web login (admin)**: On this address, users can, depending on their access privileges, modify the configuration of SPS, and perform authentication-related activities (gateway authentication, 4-eyes authorization).
- **Web login (user)**: The configuration of SPS cannot be viewed or altered from this address. Users (even ones with administrator privileges) can only perform gateway authentication and 4-eyes authorization.

NOTE: For more information on gateway authentication and 4-eyes authorization, see Advanced authentication and authorization techniques on page 860.

Both login addresses can be configured to restrict connections to a configured set of IP addresses only.

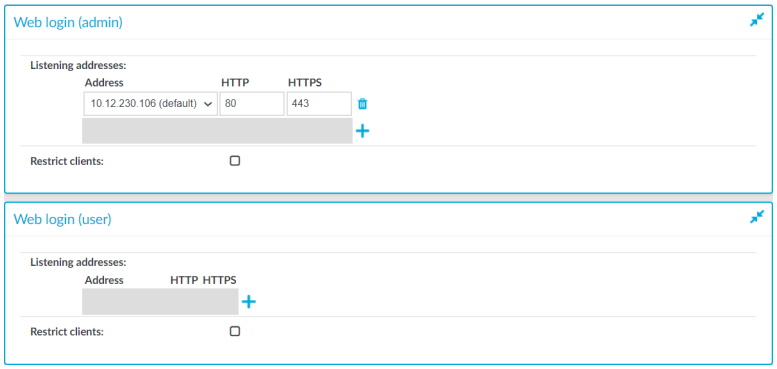
NOTE:

When configuring HTTP or SSH connections, avoid using the IP address configured for administrator or user login on SPS.

To configure two separate login addresses for accessing the web interface of SPS

1. Navigate to **Basic Settings** > **Local Services** > **Web login**.

   Figure 45: Basic Settings > Local Services > Web login — Configuring web login address

2. Choose + in the **Listening addresses** field.
3. Enter the IP address to use for connecting to SPS's user interface into the **Address** field.

   The available addresses correspond to the interface addresses configured in **Basic Settings** > **Network** > **Interfaces**. Only IPv4 addresses can be selected.
4. Enter the port number for HTTP connections into the **HTTP** field.
5. Enter the port number for HTTPS connections into the **HTTPS** field.
6. (Optional) To permit access to the SPS web interface only from selected subnets or IP addresses, select **Restrict clients**, click + and enter the IP address and netmask of the allowed clients. Note that these settings do not affect the SSH access to SPS.

   CAUTION:

   Permit administrative access to SPS only from trusted networks. If possible, monitored connections and administrative access to the SPS web interface should originate from separate networks.

   After comitting the changes, the web interface will be available only from the configured subnets or IP addresses.

   Use an IPv4 address.
7. Recommended: configure a separate login address for user connections in **Web login (user)**. The configuration settings of SPS cannot be viewed or modified from

this address.

8. Click **Commit**.

Managing logical interfaces

You can assign logical interfaces to a physical interface. Each logical interface must have its own VLAN ID, and can have its own set of (alias) IP addresses and prefixes. The configured name for each logical interface is visible on One Identity Safeguard for Privileged Sessions (SPS)'s user interface only.

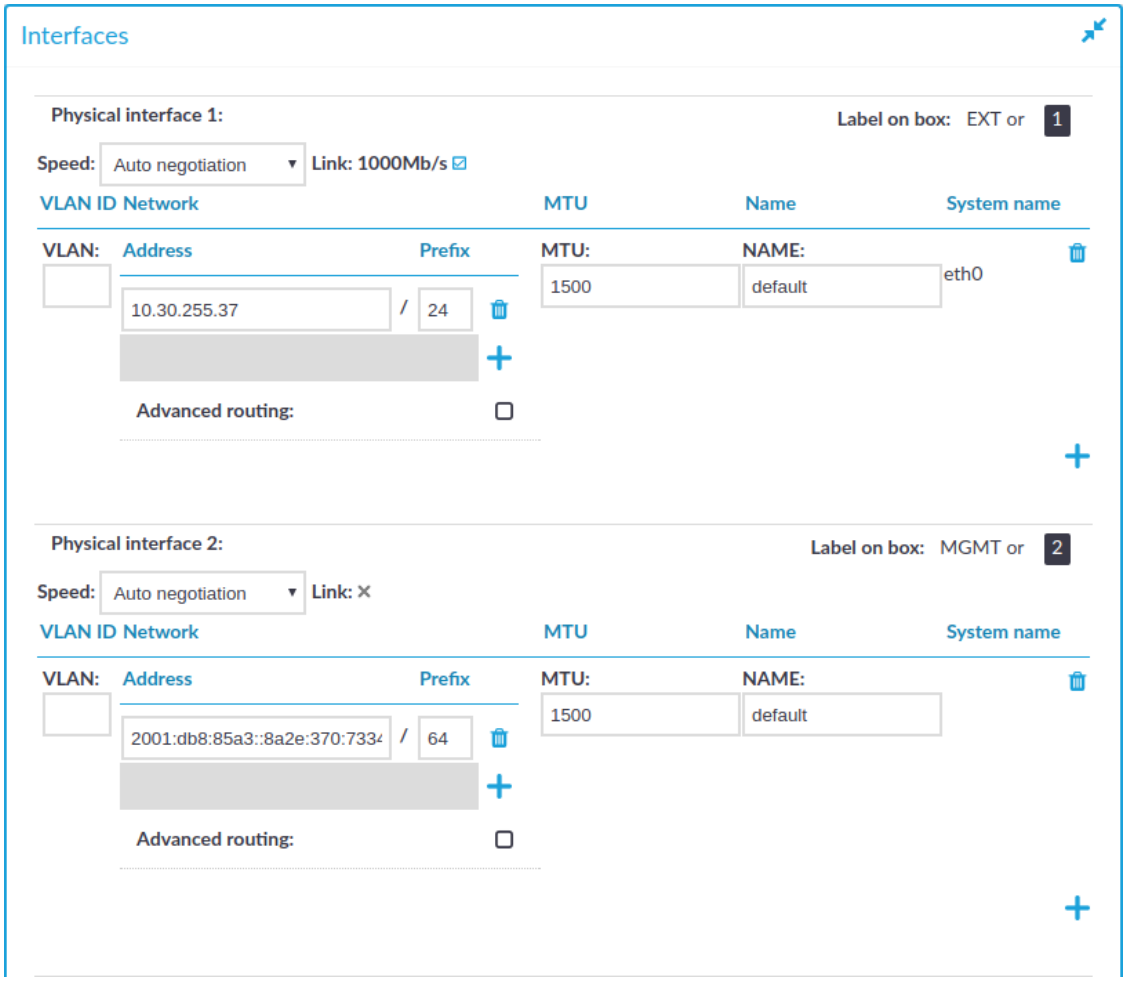
You can configure IPv4 and IPv6 addresses as well. IPv6 is intended for configuring monitored connections. Local services (including the web login) require IPv4 addresses. An interface can have multiple IP addresses, including a mix of IPv4 and IPv6 addresses.

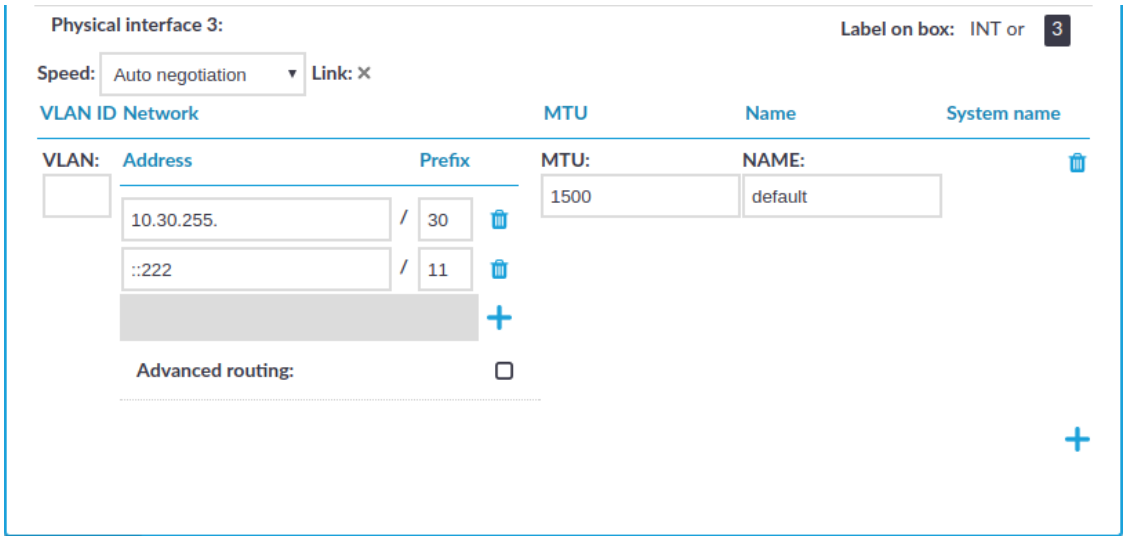
NOTE: SPS does not support scenarios with two hosts using the same IP address on different VLAN groups.

To manage logical interfaces

1. Navigate to **Basic Settings** > **Network** > **Interfaces**.

Figure 46: Basic Settings > Network > Interfaces — Managing the logical interfaces





2. If necessary, use the label on the SPS hardware to identify the physical interface to which you want to assign a logical interface.

3. Choose + to add a new logical interface. Provide the following:
   - **VLAN**: The VLAN ID of the logical interface. Optional.

     CAUTION:

     Do not set the VLAN ID unless your network environment is already configured to use this VLAN. Otherwise, your SPS appliance will be unavailable using this interface.
   - **Address**: The IP address of the logical interface.

     Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

     NOTE: Note the following limitations:
     - To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
     - If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.

     NOTE: Do not use IP addresses that fall into the following ranges:
     - `1.2.0.0/16` (reserved for communication between SPS cluster nodes)
     - `127.0.0.0/8` (localhost IP addresses)
   - **Prefix**: The IP range of the logical interface.
   - Optional: To add additional (alias) IP addresses and prefixes to a logical interface, click +. To remove an alias IP address, click the corresponding 🗑.

- **MTU**: Maximum Transmission Unit (MTU) to set per network interface (VLAN or network interface card). The default value is `1500`.
- **Name**: The name of the logical interface. This name is visible on SPS's user interface only.

To remove a logical interface, choose the on the right side.

4. Click Commit.

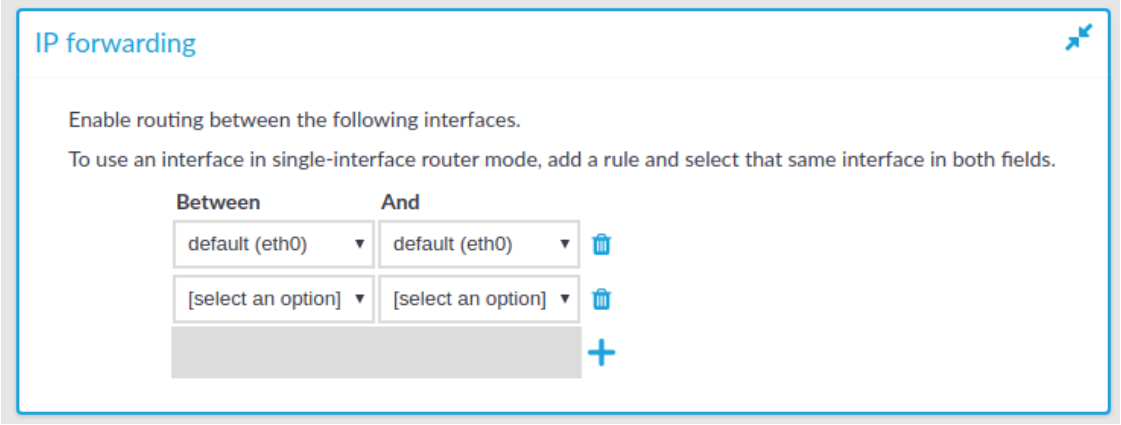
Routing uncontrolled traffic between logical interfaces

You can enable routing between logical interfaces, which allows you to direct uncontrolled traffic through SPS.

To enable routing between logical interfaces

1. Navigate to **Basic Settings** > **Network** > **IP forwarding**.

Figure 47: Basic Settings > Network > IP forwarding — IP forwarding between interfaces



2. To add a new forwarding rule, choose and select the two logical interfaces to connect. You can select the same interface in both fields to use that logical interface in single-interface router mode.

   To delete an existing rule, choose .

3. Click Commit.

Configuring the routing table

The routing table contains the network destinations SPS can reach. You have to make sure that both the monitored connections, and the local services of SPS (including connections made to the backup and archive servers, the syslog server, and the SMTP server) are routed properly.

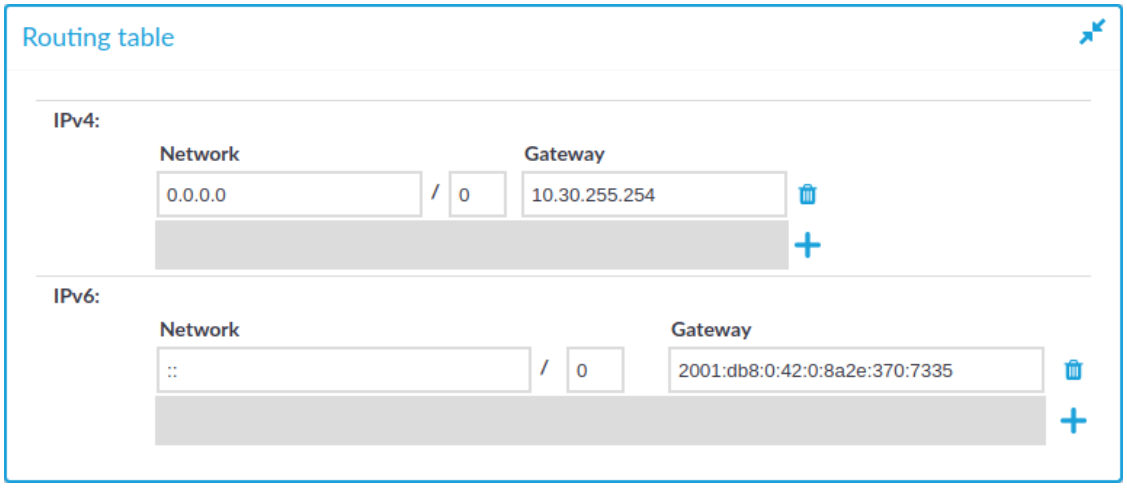
You can add multiple IPv4 and IPv6 addresses and address ranges along with their respective gateways.

To configure the routing table

1. To add a new routing entry, navigate to **Basic Settings** > **Network**.

   You can add interface-specific network routes using the **Advanced routing** option of each interface. Otherwise, use the **Routing table** option to manage networking routes.

   Figure 48: Basic Settings > Network > Routing table — Routing

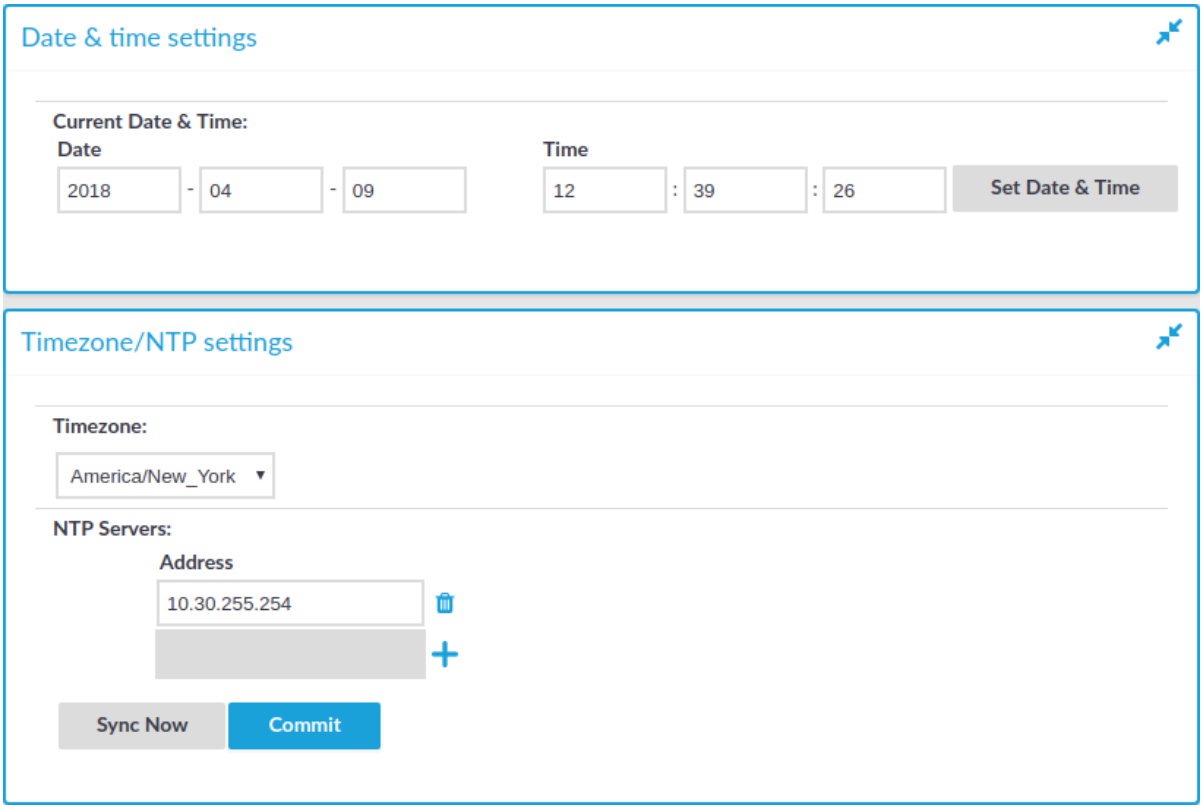
   

2. Click +, then enter the IP address and the network prefix into the **Network** field.
3. Enter the IP address of the gateway used on that subnetwork into the **Gateway** field.
4. Click Commit.

Configuring date and time

To configure the date and time-related settings of SPS, navigate to **Basic Settings** > **Date & Time**.

Figure 49: Basic Settings > Date & Time — Date and time management



CAUTION:

It is essential to set the date and time correctly on SPS, otherwise the date information of the logs and audit trails will be inaccurate.

SPS displays a warning on this page and sends an alert if the time becomes out of sync.

To explicitly set the date and time on SPS, enter the current date into respective fields of the **Date & Time settings** group and click **Set Date & Time**.

When two SPS units are operating in High Availability mode, the secondary node automatically synchronizes its time and date to the primary node. To manually synchronize the time between the nodes, click **Sync Master** (available only in High Availability mode).

To retrieve the date automatically from a time server, complete the following steps:

1. Select your timezone in the **Timezone** field.
2. Enter the IP address of an NTP time server into the **Address** field.

   Use an IPv4 address.
3. Click Commit.
4. Click the + and 🗑 icons to add new servers or delete existing ones.

5. Optional: If the time setting of SPS is very inaccurate (that is, the difference between the system time and the actual time is great), it might take a long time to retrieve the date from the NTP server. In this case, click **Sync Now** or **Sync Master** to sync the time immediately using SNTP.

System logging, SNMP and e-mail alerts

E-mail alerts and system logging can be configured on the **Basic Settings** > **Management** page.

Configuring system logging

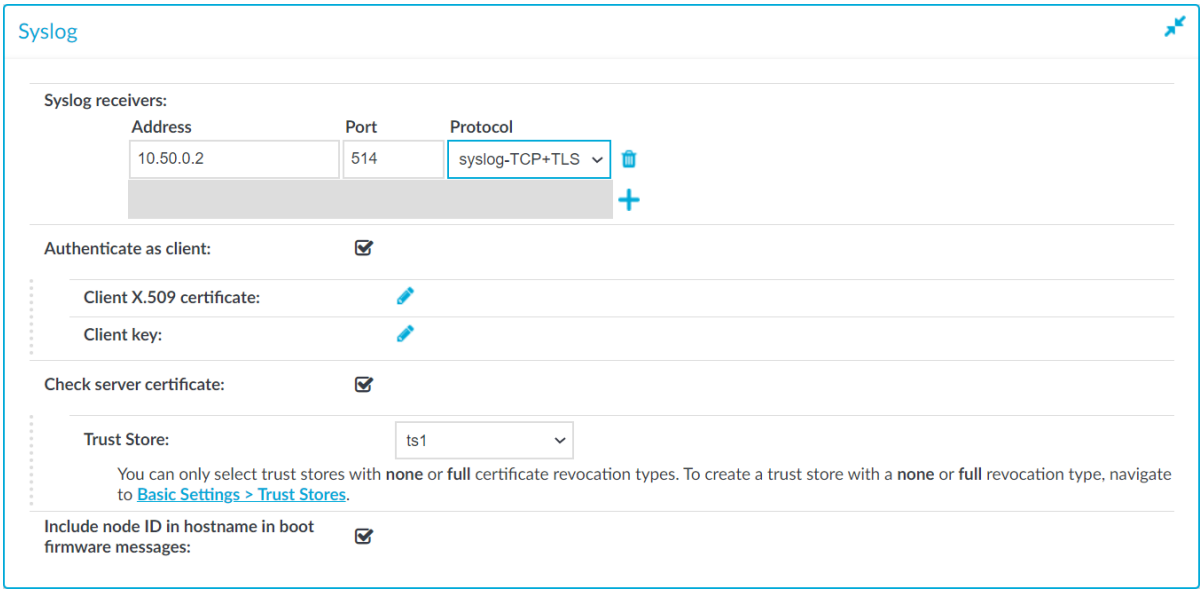
One Identity Safeguard for Privileged Sessions (SPS) can send its system log messages to remote syslog servers (for example, syslog-ng Premium Edition, syslog-ng Store Box, Splunk, or HPE ArcSight Data Platform).

NOTE: To send log messages in any custom format, contact our Support Team.

CAUTION:

The retention time for local logs of SPS is seven days. To retain them longer, forward them to a remote logserver.

Figure 50: Basic Settings > Management > Syslog — Configuring system logging



To configure logging to a remote server

1. Navigate to **Basic Settings** > **Management**.
2. Click  in the **Syslog** > **Syslog receivers** field to add a new syslog server.
3. Enter the IP address and port of the syslog server into the respective fields.

   Use an IPv4 address.
4. Select the network protocol used to transfer the messages in the **Protocol** field. The `legacy-` prefix corresponds to the legacy BSD-syslog protocol described in RFC3164, while the `syslog-` prefix corresponds to the new IETF-syslog protocol described in RFC5424. Note that not every syslog server supports the IETF protocol yet.

   Select `TCP+TLS` to send the log messages using a TLS-encrypted connection.

   > TIP: Transferring the syslog messages using TCP ensures that the server receives them.
   >
   > Transferring the syslog messages using TLS encryption ensures that third parties cannot read the messages. However, not every syslog server accepts encrypted connections. The syslog-ng Premium Edition and Open Source Edition applications, and the syslog-ng Store Box (which is a log-collector appliance similar to SPS) support both encrypted connections and the new IETF-syslog protocol as well. For details on these products, see syslog-ng Premium Edition and syslog-ng Store Box.
5. If the syslog server requires mutual authentication, that is, a certificate from SPS, check **Authenticate as client**.

   a. Generate and sign a certificate for SPS, then click the  icon in the **Client X.509 certificate** field to upload the certificate. After that, click the  icon in the **Client key** field and upload the private key corresponding to the certificate.
6. If you have selected the `TCP+TLS` protocol and you want SPS to verify the certificate of the syslog server, complete the following steps. Otherwise, click  .

   a. Select **Check server certificate**.

   b. Select a **Trust Store**.

      > NOTE: You can only select a trust store with **None** or **Full** revocation check type.

      SPS will use this trust store to verify the certificate of the server, and reject the connections if the verification fails. For more information on creating trust stores, see Verifying certificates with Certificate Authorities using trust stores.

   c. Click  .

7. To display separate hostnames for syslog messages sent by the nodes of a SPS HA nodes, select the **Include node ID in hostname in boot firmware messages** option. The node ID included in the hostname file of the syslog message is the MAC address of the node's HA interface. (Messages of the core firmware are always sent by the primary node.)

   The boot firmware boots up SPS, provides High Availability support, and starts the core firmware. The core firmware, in turn, handles everything else: provides the web interface, manages the connections, and so on.

8. Click the + and 🗑 icons to add new servers or delete existing ones.

   NOTE: To reduce the risk of the syslog server not receiving log messages from SPS because of a network outage or other problem with the syslog server, SPS buffers up to 10 Megabytes of log messages to its hard disk in case the syslog server becomes unaccessible.

Configuring e-mail alerts

The following describes how to configure e-mail alerts.

To configure e-mail alerts

1. Navigate to **Basic Settings** > **Management** > **Mail settings**.
2. If you want to encrypt the communication between SPS and the SMTP server, in **Encryption**, select the **STARTTLS** option and complete the following steps:
   - If you want SPS to verify the certificate of the server, select **Only accept certificates issued by the specified CA certificate** and click the ✎ icon in the **CA X.509 certificate** field. A pop-up window is displayed.

     Click **Browse**, select the certificate of the Certificate Authority (CA) that issued the certificate of the SMTP server, then click **Upload**. Alternatively, you can paste the certificate into the **Copy-paste** field and click **Set**.

     SPS will use this CA certificate to verify the certificate of the server, and reject the connections if the verification fails.
   - If the SMTP server requires mutual authentication, that is, it expects a certificate from SPS, enable **Authenticate as client**. Generate and sign a certificate for SPS, then click ✎ in the **Client X.509 certificate** field to upload the certificate. After that, click ✎ in the **Client key** field and upload the private key corresponding to the certificate.

   TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
3. If you want SPS to authenticate to the SMTP server, in **Authentication**, select the **Enabled** option. Enter the **Username** to authenticate with.

To configure or change the password to use to authenticate to the SMTP server, click **Change** and enter the password. Click **Update**. Click Commit.

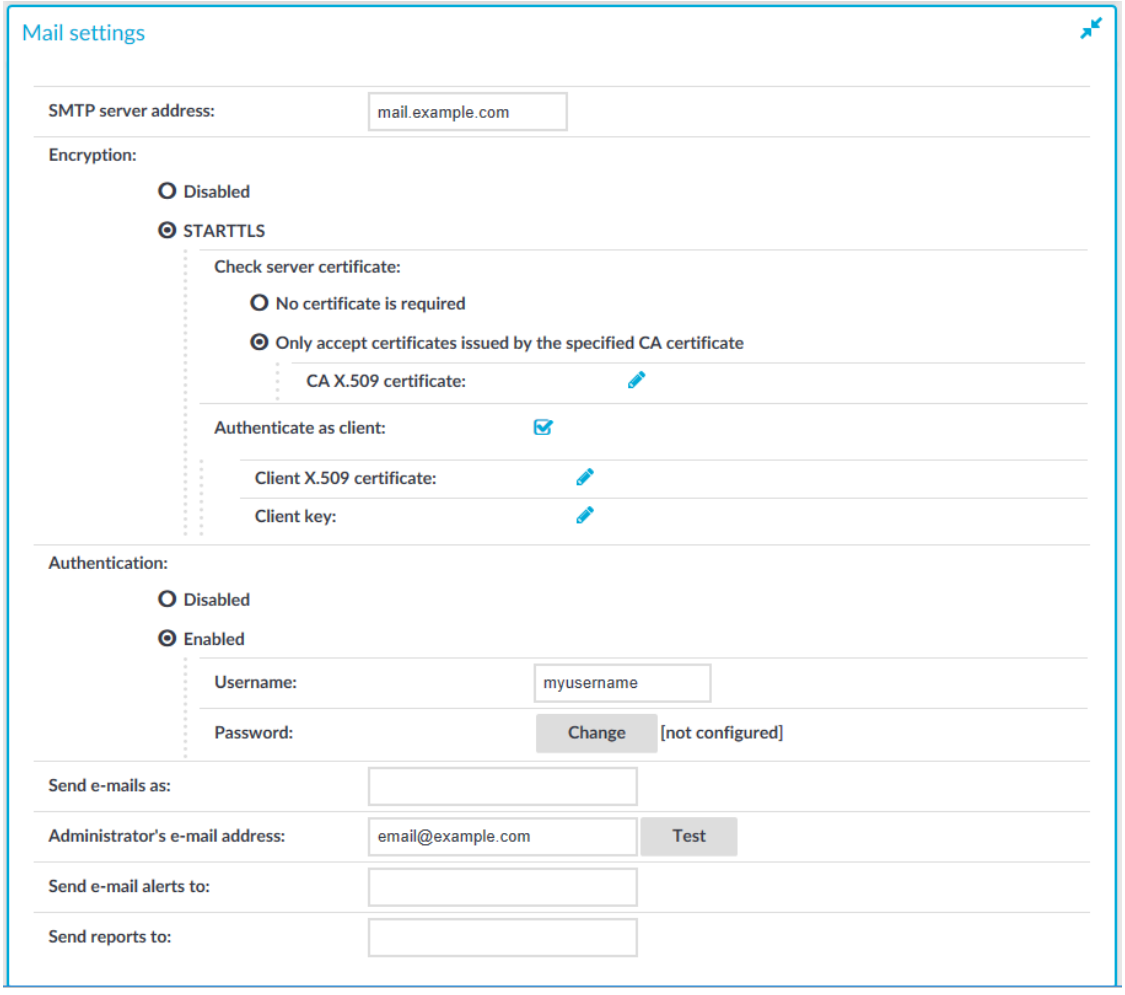
NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

4. Enter the IP address or the hostname of the mail server into the **SMTP server address** field.

   Use an IPv4 address.

Figure 51: Basic Settings > Management > Mail settings — Configuring e-mail sending



5. Enter the e-mail address where you want to receive e-mails from into the **Send e-mails as** field. This can be useful for e-mail filtering purposes. SPS sends e-mails from the address provided here. If no e-mail address is entered, e-mails will be sent from the default e-mail address.
6. Enter the e-mail address of the administrator into the **Administrator's e-mail address** field. SPS sends notifications related to system-events (but not alerts and reports) to this address.
7. Enter the e-mail address of the administrator into the **Send e-mail alerts to** field. SPS sends monitoring alerts to this address.
8. Enter the e-mail address the person who should receive traffic reports from SPS into the **Send reports to** field. For details on reports, see Reports on page 892.
9. Click 

.
10. Click **Test** to send a test message.

    If the test message does not arrive to the server, check if SPS can access the server. For details, see Troubleshooting One Identity Safeguard for Privileged Sessions (SPS) on page 935.
11. Navigate to **Basic Settings** > **Alerting & Monitoring** and select in which situations should SPS send an e-mail alert. For details, see Configuring system monitoring on SPS on page 136.
12. Click 

.

Configuring SNMP alerts

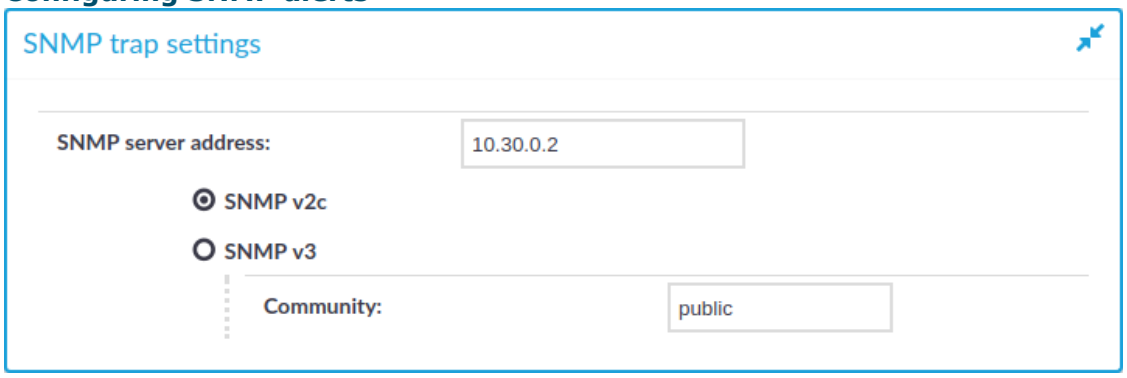
SPS can send alerts to a central monitoring server through SNMP (Simple Network Management Protocol).

To configure SNMP alerts

1. Navigate to **Basic Settings** > **Management** > **SNMP trap settings**.
2. Enter the IP address or the hostname of the SNMP server into the **SNMP server address** field.

   Use an IPv4 address.

Figure 52: Basic Settings > Management > SNMP trap settings — Configuring SNMP alerts



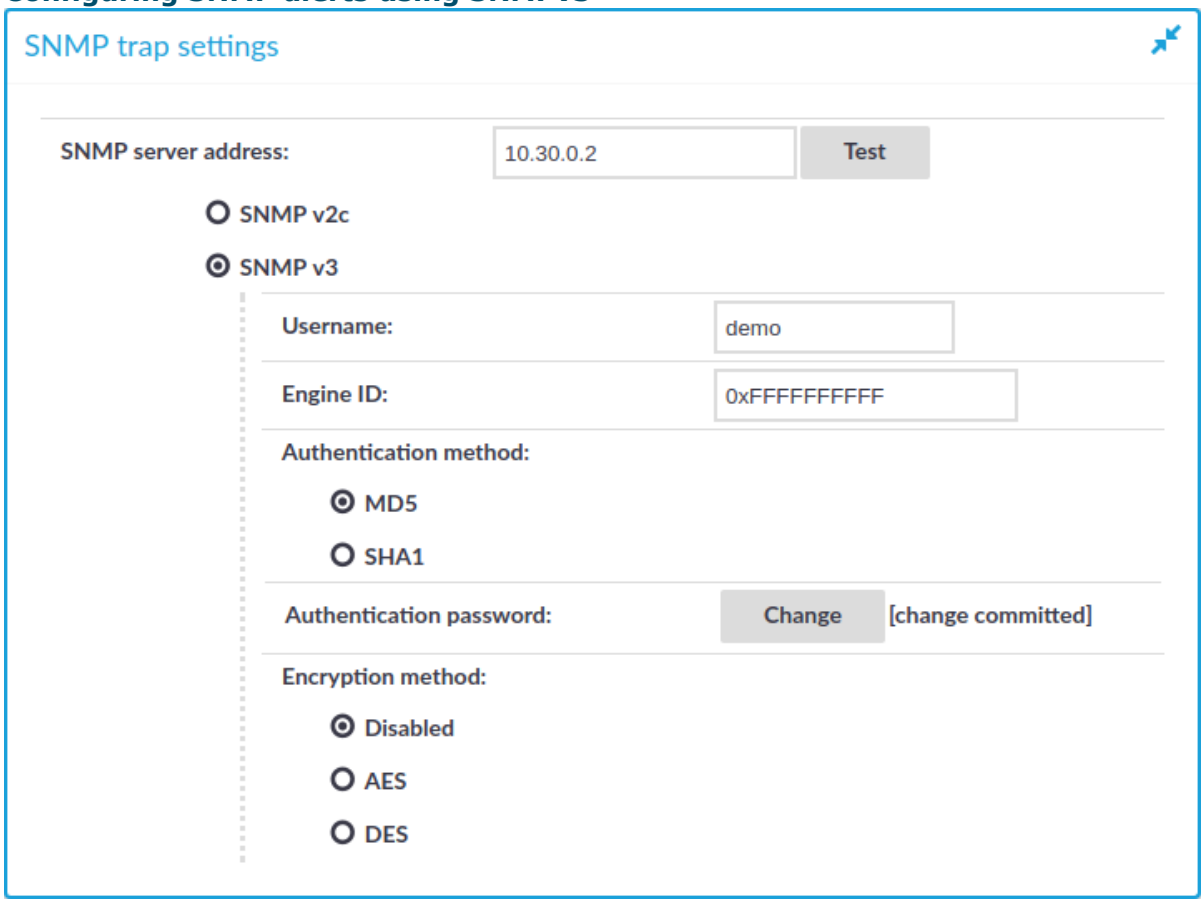
3. Select the SNMP protocol to use.

   (Optional) To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c**, and enter the community to use into the **Community** field. Otherwise, skip these steps.

   (Optional) To use the SNMP v3 protocol, select **SNMP v3** and complete the following steps. Otherwise, skip these steps.

   1. Enter the username to use into the **Username** field.
   2. Enter the engine ID to use into the **Engine ID** field. The engine ID is a hexadecimal number at least 10 digits long, starting with `0x`. For example, `0xABABABABAB`.
   3. Select the authentication method (**MD5 or SHA1**) to use from the options under the **Authentication method:** field.
   4. Enter the password to use into the **Authentication password** field.
   5. Select the encryption method (**Disabled, DES or AES**) to use from the options under the **Encryption method:** field.
   6. Enter the encryption password to use into the **Encryption password:** field.

Figure 53: Basic Settings > Management > SNMP trap settings — Configuring SNMP alerts using SNMPv3



NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

4. Click **Commit**.
5. Navigate to **Basic Settings** > **Alerting & Monitoring** and select in which situations SPS should send an SNMP alert. For details, see Configuring system monitoring on SPS on page 136.
6. Click Commit.

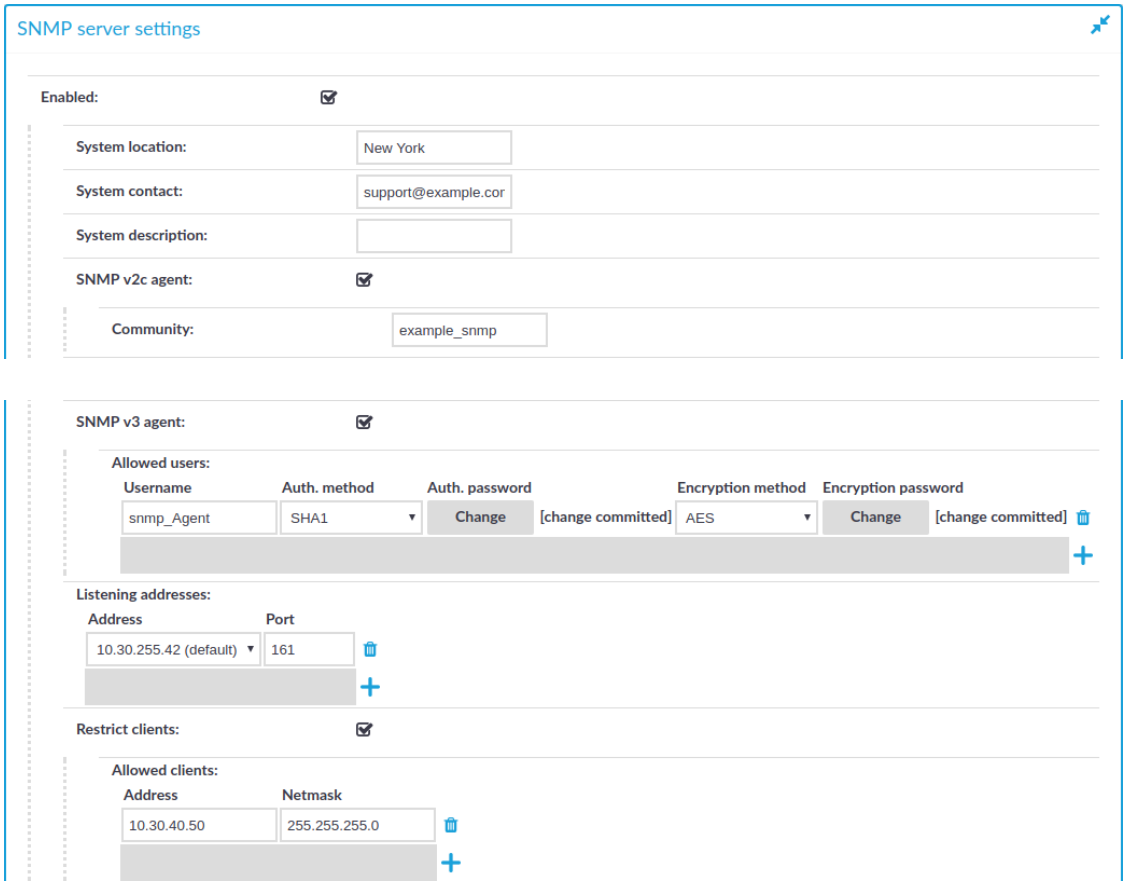
Querying SPS status information using agents

External SNMP agents can query the basic status information of SPS.

To configure which clients can query status information

1. Navigate to **Basic Settings** > **Local Services** > **SNMP server settings**.

Figure 54: Basic Settings > Local Services > SNMP server settings — Configuring SNMP agent access



2. Enable the SNMP server.
3. Optionally, you can enter the details of the SNMP server into the **System location**, **System contact**, and **System description** fields.
4. To use the SNMP v2c protocol for SNMP queries, enable **SNMP v2c agent**, and enter the community to use into the **Community** field.
5. To use the SNMP v3 protocol, select **SNMP v3 agent** and complete the following steps:

a. Click 
b. Enter the username used by the SNMP agent into the **Username** field.
c. Select the authentication method (**MD5 or SHA1**) to use from the options under the **Auth. method** field.
d. Enter the password used by the SNMP agent into the **Auth. password** field.
e. Select the encryption method (**Disabled, DES or AES**) to use from the options under the **Encryption method** field.
f. Enter the encryption password to use into the **Encryption password** field.
g. To add other agents, click  .

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

6. In the **Listening addresses** field, choose  and select the IP address and port for the SNMP server.

   The available addresses correspond to the interface addresses configured in **Basic Settings** > **Network** > **Interfaces**. Only IPv4 addresses can be selected.

   Repeat this step to add multiple addresses.
7. (Optional) To permit access to the SNMP server only from selected subnets or IP addresses, select **Restrict clients**, click  and enter the IP address and netmask of the allowed clients.

   Use an IPv4 address.

   Repeat this step to add multiple addresses.
8. Click  .

Customize system logging in One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) uses the syslog-ng Open Source Edition application (version 3.16) for system logging. Starting with SPS 5 LTS, you can customize its configuration to better integrate SPS into your logging infrastructure. If you are not familiar with syslog-ng Open Source Edition, read how syslog-ng OSE works. Customizing the configuration of syslog-ng Open Source Edition allows you to better integrate the log messages of SPS into your environment, for example, to:

- change the message format or rename message fields,
- send the messages to multiple logservers or SIEMs, or to
- select (filter) which messages to send to your logserver.

Limitations

Note that not every feature described in the syslog-ng Open Source Edition documentation is available on SPS. Typically, features that are only rarely used on logging clients are not available (for example, Java-based destinations, such as the hdfs( ) and elasticsearch2( ) destinations). For a detailed list of available modules, run the `syslog-ng --module-registry` command.

Customize the syslog-ng configuration

Do not change the syslog configuration of SPS unless you know exactly what you are doing. Incorrect changes can decrease the performance of SPS, deactivate system logging, or cause message loss.

While customizing the syslog-ng configuration, note the following points in particular:

1. Create a SPS configuration snippet in a file. Make sure that the filename ends with `.conf`. Note that syslog-ng OSE uses the configuration objects defined in these files only if they are used in a log path as well, so make sure to include a log path.

   Do not loop messages. That is, make sure that the destination does not send a message back to the original source of the message (doing so would cause an infinite loop).
2. Copy it to the `/etc/syslog-ng/conf.d/` directory of the core firmware. (If you are using a high-availability SPS cluster, SPS automatically copies the file to the secondary node as well.)

   Files located in this directory do not taint the SPS configuration and SPS automatically includes them in the configuration of syslog-ng Open Source Edition.

   Do not modify the original configuration files (for example, `/etc/syslog-ng/syslog-ng.conf` or `/etc/syslog-ng/conf.d/message-queue-client.conf`).
3. Verify that the resulting syslog-ng OSE configuration file is syntactically valid. The configuration is valid if running the following command does not show any syntax errors: `syslog-ng --syntax-only --no-caps`
4. Your changes will take effect only after you reload the configuration of syslog-ng Open Source Edition using the following command: `syslog-ng-ctl reload`

   If there are any errors in the configuration, SPS keeps on using the earlier configuration. In this case, correct the configuration, because if SPS reboots while the syslog-ng OSE configuration is invalid, SPS will not be able to log messages.

Available sources

You can use the following sources in your custom configuration. These sources are defined in the stock configuration file of syslog-ng OSE, and are in regular syslog message format (except for `s_message_queue_client`).

- `s_core_journal`: Logs of the SPS host, including log messages about the audited sessions.
- `s_message_queue_client`: Logs about the audited sessions in JSON format.
- `s_slave_boot`: Logs from the boot firmware of the secondary node in a high-availability SPS cluster.
- `src`: Logs messages of local SPS services.
- `src-internal`: The internal logs of syslog-ng OSE running on SPS.

Certificates and encrypted connections

If you are using a custom destination that requires a certificate (either to authenticate SPS, or to verify the identity of the logserver). In this case, copy the certificates to SPS into the `/etc/syslog-ng/conf.d/` directory. In the custom syslog configuration you cannot use the certificates uploaded to SPS using the web interface.

SIEM integration

Customizing the syslog configuration of SPS allows you to send log messages directly to your SIEM (for example, Splunk) in a format that your SIEM can interpret.

One Identity can provide you the configuration files needed to send the log messages of SPS to Splunk in the Splunk Common Information Model (CIM) format. If you are interested, contact our Support Team.

If you need assistance to use another SIEM format, contact professionalservices@balabit.com.

Examples

The following configuration snippet reads the messages from the built-in `s_message_queue_client` source, parses the JSON message, and sends the messages to a remote destination using the RFC5424 message format (the body of the message remains in JSON).

```
parser json {
    json-parser(
        prefix(".scb.")
        template("$MSG")
    );
};

destination d_custom_remote {
    syslog(
        "192.168.1.1"
```

```
        transport(tcp)
        port(6514)
        template("$(format-json --key .scb.*)\n")
    );
};

log {
    source(s_message_queue_client);
    parser(json);
    destination(d_custom_remote);
};
```

A sample log message using the above configuration is the following (line-breaks added for clarity):

```
{"_scb":
    {
       "username":"frobert",
             "timestamp":"1499673286.380022",
             "session_id":"svc/hm4soFrNYeAmkyzLVkzi1o/ssh_demo:4",
             "server_port":"22",
             "server_ip":"10.30.255.70",
             "protocol":"ssh",
             "event_type":"server_authentication_success",
             "connection_name":"ssh_demo",
             "client_port":"50830",
             "client_ip":"10.30.255.70"
       }
 }
```

To use this configuration snippet on your SPS, copy it to a file (make sure that the filename ends with `.conf`), change the IP address and port number to match your environment, copy it to the core firmware of your SPS into the `/etc/syslog-ng/conf.d` directory, then reload the syslog-ng configuration using `syslog-ng-ctl reload`.

Configuring system monitoring on SPS

SPS supports the SNMPv2c and SNMPv3 protocols. The SNMP server set on the **Management** tab can query status information from SPS.

NOTE: In order to have your central monitoring system to recognize the SNMP alerts sent by SPS, import the SPS-specific Management Information Base (MIB) into your monitoring system. Download all MIBs by navigating to **Basic Settings** > **Alerting & Monitoring** and clicking **Download MIBs** and import them into your monitoring system. For details, see the documentation of your monitoring system.

Configuring monitoring

The following describes how to configure monitoring.

To configure monitoring

1. Navigate to **Basic Settings** > **Alerting & Monitoring**.
2. The default threshold values of the parameters are suitable for most situations. Adjust the threshold values only if needed.
3. Click Commit.
4. Navigate to **Basic Settings** > **Management** and verify that the **SNMP settings** and **Mail settings** of SPS are correct. SPS sends alerts only to the alert e-mail address and to the SNMP server.

   CAUTION:

   Sending alerts fails if these settings are incorrect.

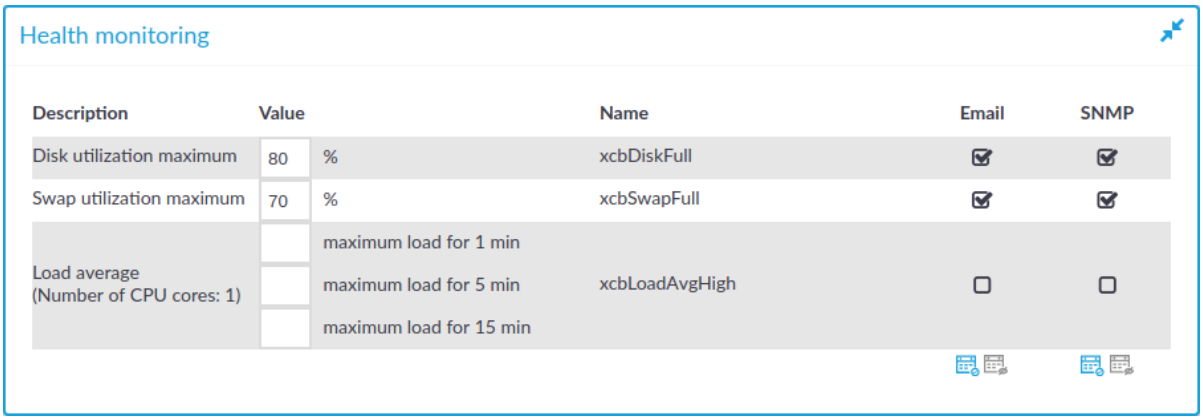
The following sections describe the parameters you can receive alerts on.

- For details on health-monitoring alerts, see Health monitoring on page 137.
- For details on system-monitoring alerts, see System related traps on page 139.
- For details on traffic-monitoring alerts, see Traffic related traps on page 142.

Health monitoring

SPS continuously monitors a number of parameters of the SPS hardware and its environment. If a parameter reaches a critical level (set in its respective **Maximum** field), SPS sends e-mail or SNMP messages to alert the administrator.

Figure 55: Basic Settings > Alerting & Monitoring — Health monitoring



- **Disk utilization maximum**: Ratio of free space available on the hard disk. SPS sends an alert if the audit trails use more space than the set value. Archive the audit trails to a backup server to free disk space. For details, see Archiving on page 158.

  NOTE: The alert message includes the actual disk usage, not the limit set on the web interface. For example, you set SPS to alert if the disk usage increases above `10` percent. If the disk usage of SPS increases above this limit (for example to `17` percent), you receive the following alert message: `less than 90% free (= 17%)`. This means that the amount of used disk space increased above 10% (what you set as a limit, so it is less than 90%), namely to 17%.
- **Load average**: The average load of SPS during the last one, five, or 15 minutes.
- **Swap utilization maximum**: Ratio of the swap space used by SPS. SPS sends an alert if it uses more swap space than the set value.

Preventing disk space fill-up

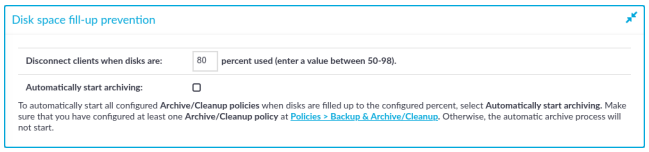
The following describes how to prevent disk space from filling up.

NOTE: One Identity highly recommends this if One Identity Safeguard for Privileged Sessions (SPS) is hosted in a virtual environment.

To prevent disk space from filling up

1. Navigate to **Basic Settings** > **Management** > **Disk space fill-up prevention**.

   Figure 56: Basic Settings > Management > Disk space fill-up prevention — Preventing disk space fill-up

2. Enter the limit of maximum disk utilization in percents in the **Disconnect clients when disks are: x percent used** field. Make sure to enter a value between `50-98` percent. When disk space is used above the configured limit, SPS disconnects all clients. The default value is `80`.
3. (Optional) To automatically start all configured archiving/cleanup jobs when disk usage goes over the limit, select the **Automatically start archiving** option.

   For more information on configuring an archiving policy, see Archiving on page 158.

   NOTE: If there is no archiving policy configured, selecting this option will not trigger automatic archiving.
4. Click Commit.

5. Navigate to **Basic Settings** > **Alerting & Monitoring** > **Health monitoring** and enable alert **Disk utilization maximum**.
6. Click 

.

System related traps

SPS can send the following system related alerts in e-mail or as SNMP trap. To configure these alerts, see Configuring e-mail alerts on page 127 and Configuring SNMP alerts on page 129.

NOTE:

Configure **Disk space fill-up prevention**, and configure SPS to send an alert if the free space on the disks of SPS is low. For details, see *Preventing disk space fill-up* in the *Administration Guide*.

Configure SPS to send an alert if a user fails to login to SPS. For details, see the **Login failed** alert in *System related traps* in the *Administration Guide*.

Figure 57: Basic Settings > Alerting & Monitoring — health monitoring

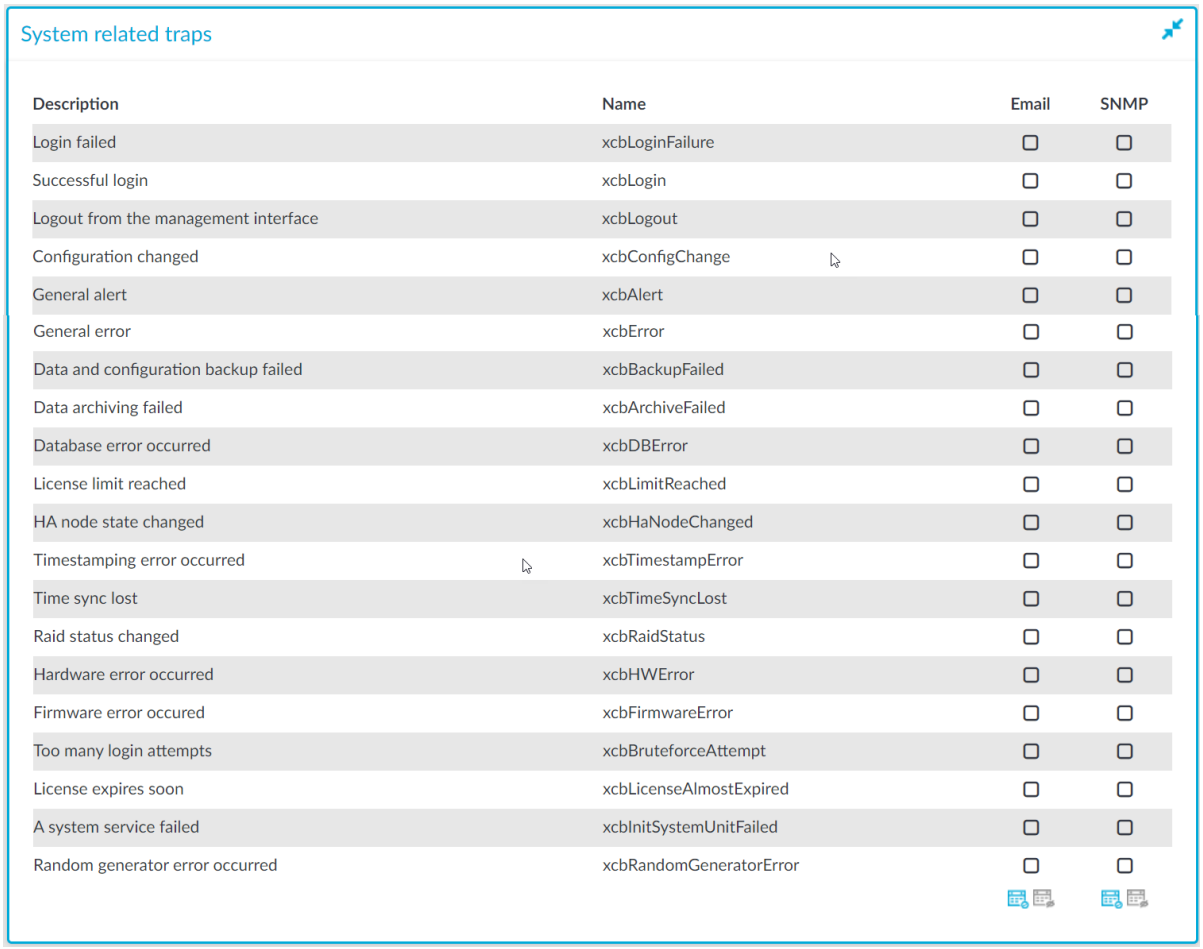


Table 1: System related traps

Name	SNMP alert ID	Description
Login failed	`xcbLoginFailure`	Failed login attempts from SPS web interface.
Successful login	`xcbLogin`	Successful login attempts into SPS web interface.
Logout from the management interface	`xcbLogout`	Logouts from SPS web interface.
Configuration changed	`xcbConfigChange`	Any modification of SPS's configuration.
General alert	`xcbAlert`	General alerts and error messages occurring on SPS.

Name	SNMP alert ID	Description
		Note that alerts on general alerts and errors are sent whenever there is an alert or error level message in the SPS system log. These messages are very verbose and mainly useful only for debugging purposes. Enabling these alerts may result in multiple e-mails or SNMP traps sent about the same event.
General error	`xcbError`	
Data and configuration backup failed	`xcbBackupFailed`	Alerts if the backup procedure is unsuccessful.
Data archiving failed	`xcbArchiveFailed`	Alerts if the archiving procedure is unsuccessful.
Database error occurred	`xcbDBError`	An error occurred in the database where SPS stores the connection metadata. For assistance, contact our Support Team.
License limit reached	`xcbLimitReached`	The number of protected servers (or concurrent sessions) reached the limit set in the SPS license. Clients cannot connect to new servers using SPS.
HA node state changed	`xcbHaNodeChanged`	A node of the SPS cluster changed its state (for example, a takeover occurred).
Timestamping error occurred	`xcbTimestampError`	An error occurred during the timestamping process (for example, the timestamping server did not respond).
Time sync lost	`xcbTimeSyncLost`	The system time became out of sync.
Raid status changed	`xcbRaidStatus`	The status of the node's RAID device changed its state.
Hardware error occurred	`xcbHWError`	SPS detected a hardware error.
Firmware error occured	`xcbFirmwareError`	SPS detected a firmware error, which can be as follows: • **Corrupted**: The firmware integrity check failed. If a firmware is shown as corrupted,

Name	SNMP alert ID	Description
		contact our Support Team. • **Tainted**: It indicates that you have modified a file of the firmware locally. If you have modified a local file unintentionally, contact our Support Team.
Too many login attempts	`xcbBruteforceAttempt`	SPS has detected a possible brute-force attack.
License expires soon	`xcbLicenseAlmostExpired`	Your SPS license will expire within 60 days.
A system service failed	`xcbInitSystemUnitFailed`	A system service has failed. Note that one alert is sent for each failed service.
Random generator error occurred	`xcbRandomGeneratorError`	The random generator repeatedly created the same byte sequence. To fix this issue, you can restart your SPS. If the error persists, contact our Support Team.

Traffic related traps

SPS can send the following traffic related alerts in e-mail or as SNMP trap. To configure these alerts, see Configuring e-mail alerts on page 127 and Configuring SNMP alerts on page 129.

Figure 58: Basic Settings > Alerting & Monitoring — health monitoring

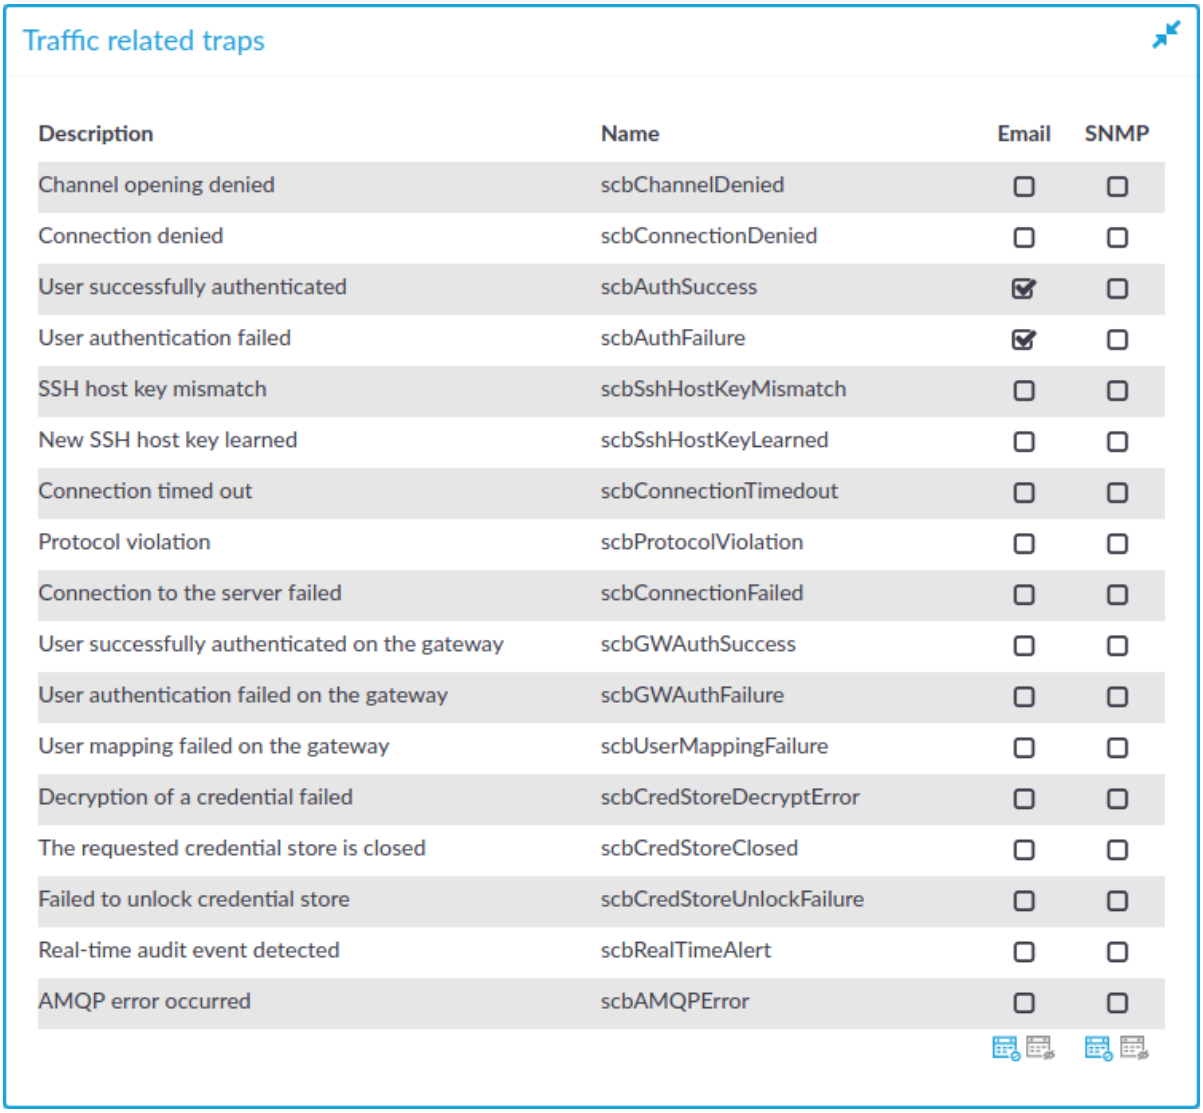


Table 2: Traffic related traps

Name	SNMP alert ID	Description
Channel opening denied	`scbChannelDenied`	A user attempted to open a channel not permitted by the channel policy.
Connection denied	`scbConnectionDenied`	A user attempted to connect a server not permitted in the connection policies.
User successfully authenticated	`scbAuthSuccess`	A user successfully authenticated on a protected server.

Name	SNMP alert ID	Description
User authentication failed	`scbAuthFailure`	A user failed to complete the authentication on a protected server.
SSH host key mismatch	`scbSshHostKeyMismatch`	The SSH host key of a server did not match the key stored on SPS.
New SSH host key learned	`scbHostKeyLearned`	SPS learned a new SSH host key.
Connection timed out	`scbConnectionTimedout`	A connection to a protected server timed out.
Protocol violation	`scbProtocolViolation`	A connection violated the protocol as specified in the RFC or protocol documentation. This may have been caused by an incompatible application or a deliberate attack.
Connection to the server failed	`scbConnectionFailed`	A connection to a protected server failed.
User successfully authenticated on the gateway	`scbGWAuthSuccess`	A user has successfully authenticated a connection on SPS as part of a gateway-authentication process.
User authentication failed on the gateway	`scbGWAuthFailure`	The gateway-authentication of a connection has failed.
User mapping failed on the gateway	`scbUserMappingFailure`	A usermapping policy did not find a suitable mapping for the connection.
Decryption of a credential store failed	`scbCredStoreDecrpytError`	SPS could not unlock a password-protected Credential Store. Navigate to **User menu** > **Unlock Credential Store** and enter the password(s) to open the Credential Store.
The requested credential store is closed	`scbCredStoreClosed`	A user attempted to access a connection policy that uses a password-protected Credential Store, and the Credential Store has not been unlocked. Navigate to **User menu** > **Unlock Credential Store** and enter the password(s) to open the Credential Store.

Name	SNMP alert ID	Description
Failed to unlock credential store	`scbCredStoreUnlockFailure`	A user attempted to unlock a password-protected Credential Store with an incorrect password. Navigate to **User menu** > **Unlock Credential Store** and enter the correct password (s) to open the Credential Store.
Real time audit event detected	`scbRealTimeAlert`	A real-time audit event has occurred.
AMQP error occurred	`scbAMQPError`	An error occurred in the event queue where SPS forwards session data. contact our Support Team.

Data and configuration backups

Backups create a snapshot of the configuration of One Identity Safeguard for Privileged Sessions (SPS) or the data which can be used for recovery in case of errors. SPS can create automatic backups of its configuration and the stored audit-trails to a remote server.

Configuring backups is a two-step process:

1. Create a backup policy.
2. Assign that policy to the system or a connection — depending on what it is that you wish to back up, SPS's configuration or a connection.

Creating a backup policy

Backup policies define the address of the backup server, which protocol to use to access it, and other parameters. SPS can be configured to use the Rsync, SMB/CIFS, and NFS protocols to access the backup server:

- To configure backups using Rsync over SSH, see Creating a backup policy using Rsync over SSH on page 146.
- To configure backups using SMB/CIFS, see Creating a backup policy using SMB/CIFS on page 149.
- To configure backups using NFS, see Creating a backup policy using NFS on page 153.

The different backup protocols assign different file ownerships to the files saved on the backup server. The owners of the backup files created using the different protocols are the following:

- *Rsync*: The user provided on the web interface.
- *SMB/CIFS*: The user provided on the web interface.
- *NFS*: `root` with `no-root-squash`, `nobody` otherwise.

CAUTION:

SPS cannot modify the ownership of a file that already exists on the remote server. If you change the backup protocol but you use the same directory of the remote server to store the backups, make sure to adjust the ownership of the existing files according to the new protocol. Otherwise SPS cannot overwrite the files and the backup procedure fails.

Assigning a backup policy

Once you have configured a backup policy, set it as a system backup policy (for configuration backups) or data backup policy (for connections backups):

- To configure a system backup policy, see Creating configuration backups on page 155.
- To configure a data backup policy, see Creating data backups on page 156.

NOTE: Backup deletes all other data from the target directory. Restoring a backup deletes all other data from SPS. For details on restoring configuration and data from backup, see Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance on page 967.

Creating a backup policy using Rsync over SSH

The **Rsync over SSH** backup method connects the target server with SSH and executes the `rsync` UNIX command to copy the data to the remote server. SPS authenticates itself with a public key — password-based authentication is not supported.

CAUTION:

The backup server must run rsync version 3.0 or newer.

To create a backup policy using Rsync over SSH

1. Navigate to **Policies** > **Backup & Archive** and click + in the **Backup policies** section to create a new backup policy.
2. Enter a name for the backup policy (for example, `config-backup`).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example, **`23:00`**).

   You can add the start time for additional backup processes.

   CAUTION:

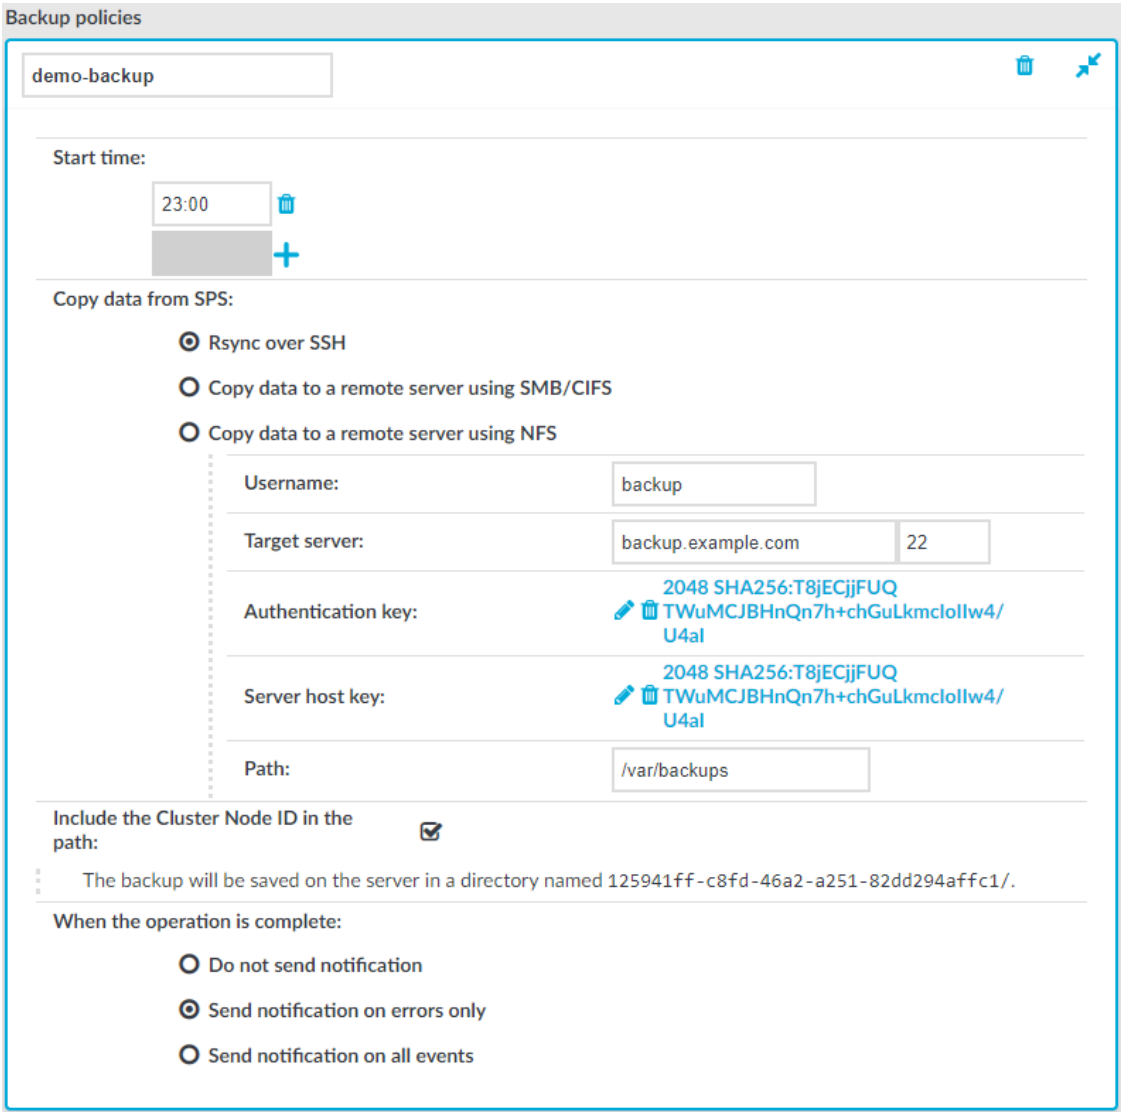
   When specifying an additional start time, ensure that the previous backup process finishes before the new backup process starts.

4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example, `backup.example.com`).

   Use an IPv4 address.

5. Select **Rsync over SSH** from the **Copy data from PSM** radio buttons.

Figure 59: Policies > Backup & Archive > Backup policies — Configuring backups using rsync

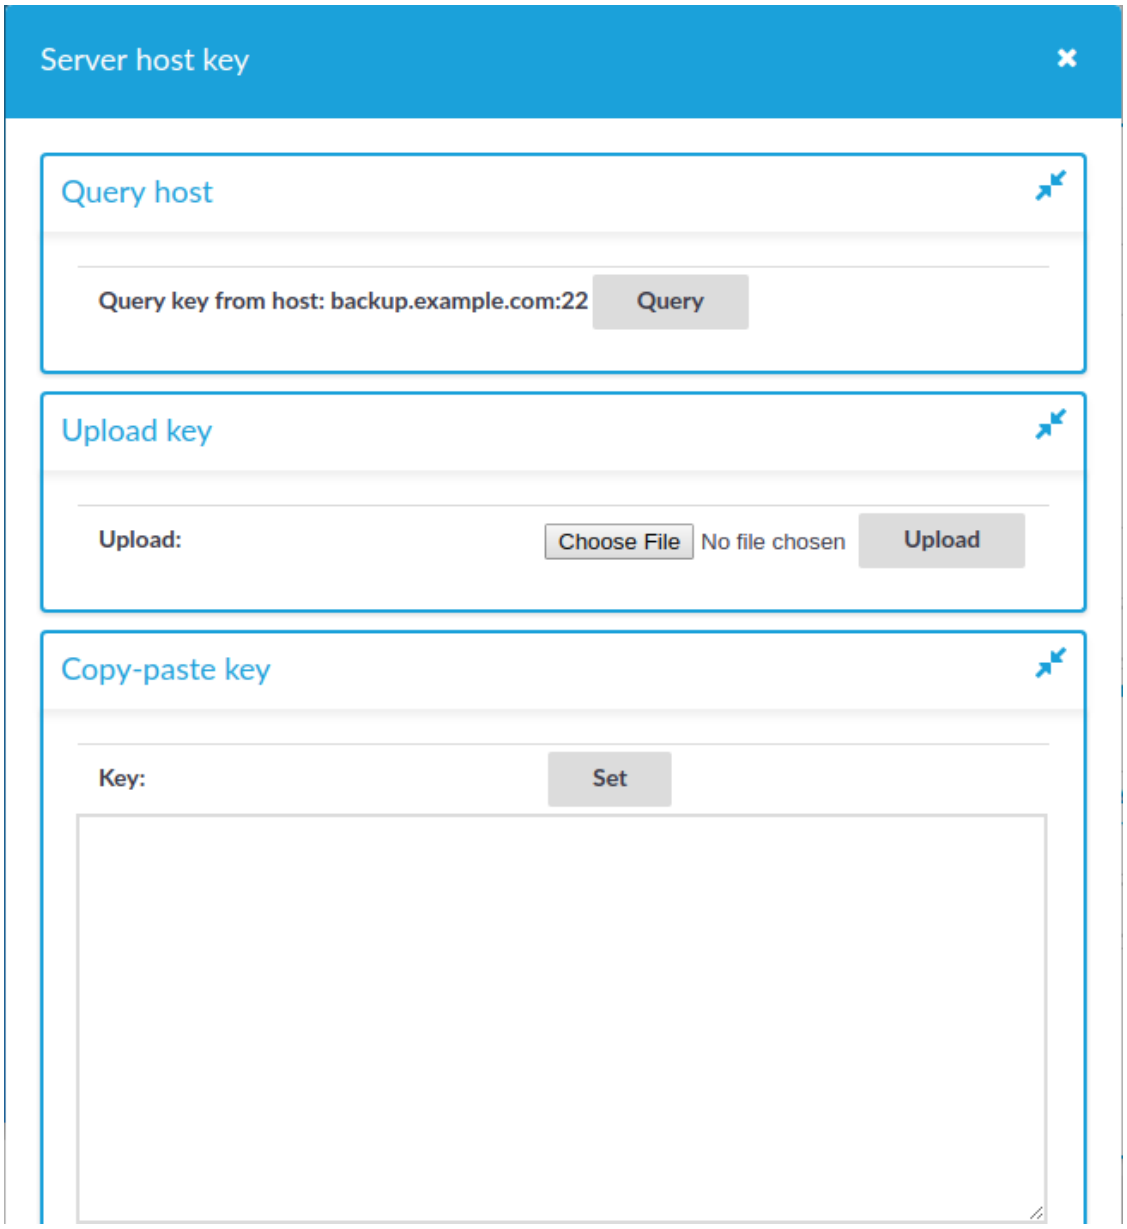


6. Enter the username used to log on to the remote server into the **Username** field.
7. Click  in the **Authentication key** field. A popup window is displayed.
8. Generate a new keypair by clicking **Generate** or upload or paste an existing one. This key will be used to authenticate SPS on the remote server. The public key of this

keypair must be imported to the remote server.

9. Click  in the **Server host key** field. A popup window is displayed.
10. Click **Query** to download the host key of the server, or upload or paste the host key manually. SPS will compare the host key shown by the server to this key, and connect only if the two keys are identical.

Figure 60: Policies > Backup & Archive > Backup policies — Configuring SSH keys



11. Enter the port number of the SSH server running on the remote machine into the **Port** field.
12. Enter the path to the backup directory on the target server into the **Path** field (for example `/backups`).

    SPS saves all data into this directory, automatically creating the subdirectories. Backups of audit-trails are stored in the `data`, configuration backups in the `config` subdirectory.
13. When your SPS instance is a node in a cluster, select **Include the Cluster Node ID in the path**. This ensures that the ID of the node is included in the path of the relevant directory, which is required to prevent cluster nodes from backing up data to the same location, and so overwriting each other's data and resulting in data loss. In addition, having the node's ID in the directory name also enables easy identification.

    CAUTION:

    Hazard of data loss

    Unchecking **Include the Cluster Node ID in the path** when your SPS is a node in a cluster can result in data loss if you have configured configuration synchronization across your cluster nodes.
14. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab.

    To include the list of files in the e-mail, select **Send notification on all events** and enable the **Include file list** option. However, note that if the list is very long (for example, SPS stores over 20000 audit trails), the SPS web interface might become unaccessible. In this case, set the **Maximum number of files in notification** lower. After this number has been reached, file names will be omitted from the notification.

    NOTE: This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see Configuring system monitoring on SPS on page 136).
15. Click 

.
16. To assign the backup policy to a connection, see Creating data backups on page 156.

Creating a backup policy using SMB/CIFS

The **Copy data to a remote server using SMB/CIFS** backup method connects to a share on the target server with Server Message Block protocol. SMB/CIFS is mainly used on Microsoft Windows Networks.

When deployed from the Azure Marketplace, you can use Azure File storage shares for Backup and Archive Policies. This is very useful as you can change the quota for the file storage dynamically, so the cumulative size of the audit trails is not limited to the OS disk

size. You can set up this share as normal SMB shares in your Backup and Archive policies. You can obtain the parameters for the policy from the Azure portal.

NOTE: Backup and archive policies only work with existing shares and subdirectories.

If a server has a share at, for example, `archive` and that directory is empty, when the user configures `archive/scb1` (or similar) as a backup/archive share, it will fail.

CAUTION:

When you try to create backups and archives from SPS to NetApp devices using the CIFS protocol, the operation may fail with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on.`

To overcome this problem, grant the SPS user "Full Control" access rights to the CIFS share on the NetApp device.

1. Navigate to **Policies** > **Backup & Archive** and click + in the **Backup policies** section to create a new backup policy.
2. Enter a name for the backup policy (for example, `config-backup`).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example, **`23:00`**).

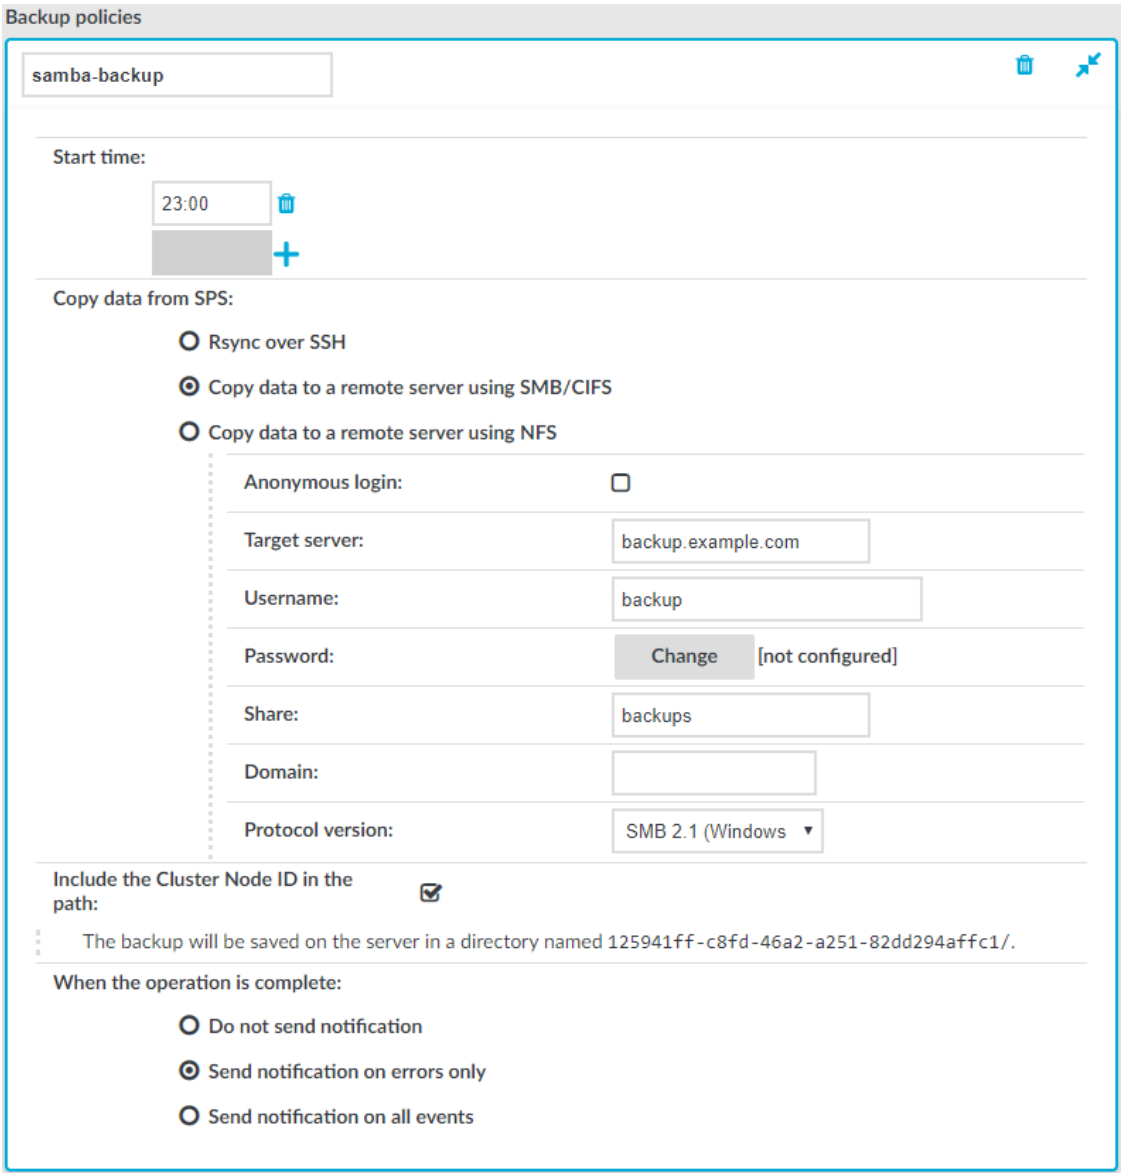
   You can add the start time for additional backup processes.

   CAUTION:

   When specifying an additional start time, ensure that the previous backup process finishes before the new backup process starts.
4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example, `backup.example.com`).

   Use an IPv4 address.
5. Select **Copy data to a remote server using SMB/CIFS** from the **Copy data from PSM** radio buttons.

Figure 61: Policies > Backup & Archive > Backup policies — Configuring backups through SMB/CIFS



6. Enter the username used to log on to the remote server into the **Username** field, or select the **Anonymous login** option.

   Usernames can contain space.

7. Enter the password corresponding to the username into the **Password** field.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9

- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

8. Enter the name and directory path of the share into the **Share** field. Use the following format:

   ```
   share_name/path/to/directory
   ```

   You can use backslashes and forward slashes as well.

   SPS saves all data into this directory, automatically creating the subdirectories. Backups of audit-trails are stored in the `data`, configuration backups in the `config` subdirectory.

9. Enter the domain name of the target server into the **Domain** field.
10. Select which SMB protocol to use when SPS connects to the server in the **Protocol version** field. Servers are usually backwards compatible with earlier protocol versions (for example, a server that supports version 2.1 supports versions 2.0 and 1.0 as well).
11. When your SPS instance is a node in a cluster, select **Include the Cluster Node ID in the path**. This ensures that the ID of the node is included in the path of the relevant directory, which is required to prevent cluster nodes from backing up data to the same location, and so overwriting each other's data and resulting in data loss. In addition, having the node's ID in the directory name also enables easy identification.

    CAUTION:

    Hazard of data loss

    Unchecking **Include the Cluster Node ID in the path** when your SPS is a node in a cluster can result in data loss if you have configured configuration synchronization across your cluster nodes.

12. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab.

    To include the list of files in the e-mail, select **Send notification on all events** and enable the **Include file list** option. However, note that if the list is very long (for example, SPS stores over 20000 audit trails), the SPS web interface might become unaccessible. In this case, set the **Maximum number of files in notification** lower. After this number has been reached, file names will be omitted from the notification.

    NOTE: This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see Configuring system monitoring on SPS on page 136).

13. Click Commit.
14. To assign the backup policy to a connection, see Creating data backups on page 156.

Creating a backup policy using NFS

The **Copy data to a remote server using NFS** backup method connects to a shared directory of the target server with the Network File Share protocol.

NOTE: Backup and archive policies only work with existing shares and subdirectories.

If a server has a share at, for example, `archive` and that directory is empty, when the user configures `archive/scb1` (or similar) as a backup/archive share, it will fail.

The version of NFS used is automatically detected. All versions of NFS, up to and including NFS version 4 protocol, are supported.

1. Navigate to **Policies** > **Backup & Archive** and click  in the **Backup policies** section to create a new backup policy.
2. Enter a name for the backup policy (for example, `config-backup`).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example, **`23:00`**).

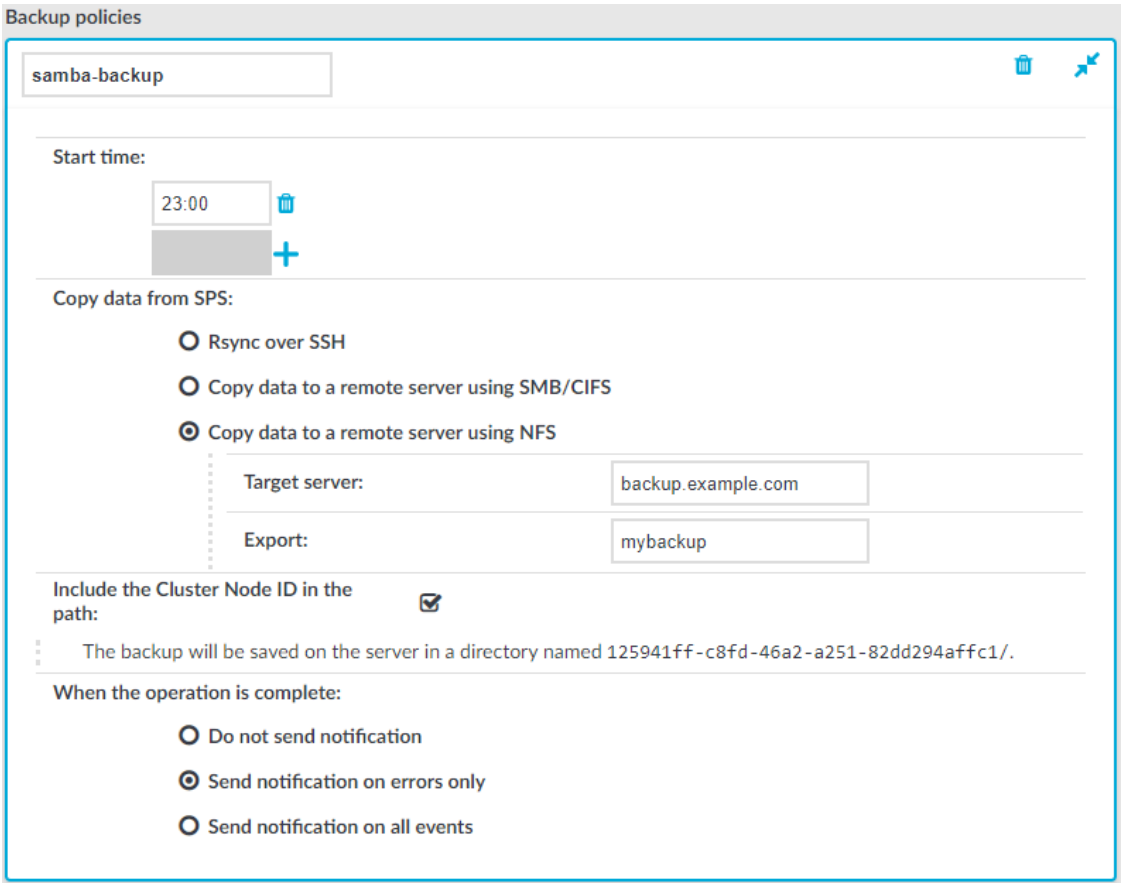
   You can add the start time for additional backup processes.

   CAUTION:

   When specifying an additional start time, ensure that the previous backup process finishes before the new backup process starts.
4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example, `backup.example.com`).

   Use an IPv4 address.
5. Select **Copy data to a remote server using NFS** from the **Copy data from PSM** radio buttons.

Figure 62: Policies > Backup & Archive > Backup policies — Configuring NFS backups



6. Enter the domain name of the remote server into the **Target server** field.
7. Enter the name of the NFS export into the **Export** field.

   SPS saves all data into this directory, automatically creating the subdirectories. Audit-trail backups are stored in the `data`, configuration backups in the `config` subdirectory.
8. The remote server must also be configured to accept backups from SPS.

   Add a line that corresponds to the settings of SPS to the `/etc/exports` file of the backup server. This line should contain the following parameters:

   - The path to the backup directory as set in the **Export** field of the SPS backup policy.
   - The IP address of the SPS interface that is used to access the remote server. For more information on the network interfaces of SPS, see Network settings on page 114.

     Use an IPv4 address.
   - The following parameters: `(rw,no_root_squash,sync)`.

Example: Configuring NFS on the remote server

For example, if SPS connects the remote server from the `192.168.1.15` IP address and the data is saved into the `/var/backups/SPS` directory, add the following line to the `/etc/exports` file:

```
/var/backups/SPS 192.168.1.15(rw,no_root_squash,sync)
```

9. On the remote server, run the following command:

   ```
   exportfs -a
   ```

   Verify that the `rpc portmapper` and `rpc.statd` applications are running.

10. When your SPS instance is a node in a cluster, select **Include the Cluster Node ID in the path**. This ensures that the ID of the node is included in the path of the relevant directory, which is required to prevent cluster nodes from backing up data to the same location, and so overwriting each other's data and resulting in data loss. In addition, having the node's ID in the directory name also enables easy identification.

    CAUTION:

    Hazard of data loss

    Unchecking **Include the Cluster Node ID in the path** when your SPS is a node in a cluster can result in data loss if you have configured configuration synchronization across your cluster nodes.

11. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab.

    To include the list of files in the e-mail, select **Send notification on all events** and enable the **Include file list** option. However, note that if the list is very long (for example, SPS stores over 20000 audit trails), the SPS web interface might become unaccessible. In this case, set the **Maximum number of files in notification** lower. After this number has been reached, file names will be omitted from the notification.

    NOTE: This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see Configuring system monitoring on SPS on page 136).

12. Click Commit.
13. To assign the backup policy to a connection, see Creating data backups on page 156.

Creating configuration backups

To create a configuration backup, assign a backup policy as the **System backup policy** of SPS.

TIP: To create an immediate backup of SPS's configuration to your machine (not to the backup server), select **Basic Settings** > **System** > **Export configuration**. Note that the configuration export contains only the system settings and configuration files (including changelogs). System backups includes additional information like reports and alerts, and also the connection database.

When exporting the configuration of SPS, or creating configuration backups, always use encryption. Handle the exported data with care, as it contains sensitive information, including credentials. For details on encrypting the configuration, see *Encrypting configuration backups with GPG* in the *Administration Guide*.

To encrypt your configuration backups, see Encrypting configuration backups with GPG on page 157.

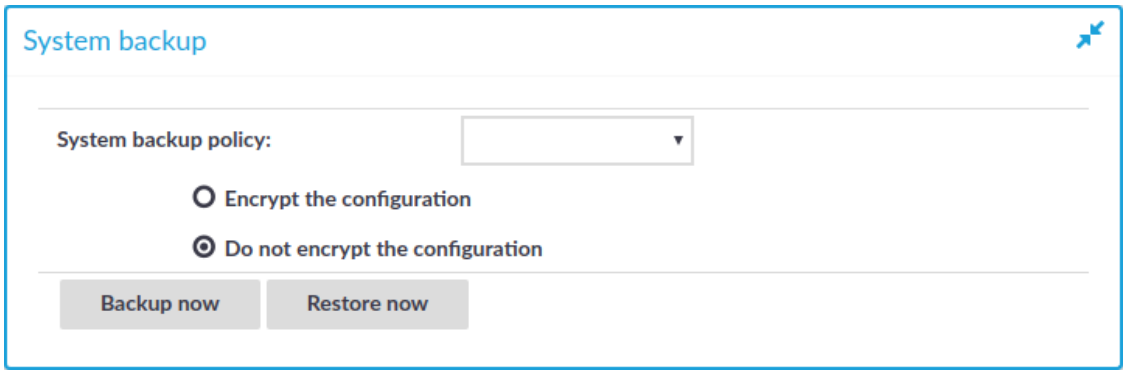
Prerequisites

You have to configure a backup policy before starting this procedure. For details, see Data and configuration backups.

To create a configuration backup

1. Navigate to **Basic Settings** > **Management** > **System backup**.

   Figure 63: Basic Settings > Management > System backup — Configuring system backups

2. Select the backup policy you want to use for backing up the configuration of SPS in the **System backup policy** field.
3. Click Commit.
4. *Optional:* To start the backup process immediately, click **Backup now**. The **Backup now** functionality works only after a backup policy has been selected and committed.

Creating data backups

To configure data backups, assign a backup policy to the connection.

NOTE:

When exporting the configuration of SPS, or creating configuration backups, always use encryption. Handle the exported data with care, as it contains sensitive information, including credentials. For details on encrypting the configuration, see *Encrypting configuration backups with GPG* in the *Administration Guide*.

Prerequisites

- Configure the system backup. Restoring a data backup works only if a matching system configuration and metadata is available, that is, if a system backup is restored first. For details, see Creating configuration backups on page 155.
- Configure a backup policy before starting this procedure. For details, see Data and configuration backups on page 145.

To configure data backups

1. Navigate to **[Your chosen protocol] Control** > **Connections**.
2. Select the connection you want to back up.
3. Select a backup policy in the **Backup policy** field.
4. Click 

.
5. *Optional:* To start the backup process immediately, click **Backup** or **Backup ALL**. The **Backup** and **Backup ALL** functionalities work only after a backup policy has been selected and committed.

Encrypting configuration backups with GPG

You can encrypt the configuration file of SPS during system backups using the public-part of a GPG key. The system backups of SPS contain other information as well (for example, databases), but only the configuration file is encrypted. Note that system backups do not contain audit-trail data.

When exporting the configuration of SPS, or creating configuration backups, always use encryption. Handle the exported data with care, as it contains sensitive information, including credentials. For details on encrypting the configuration, see *Encrypting configuration backups with GPG* in the *Administration Guide*.

For details on restoring configuration from a configuration backup, see Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance on page 967.

NOTE: It is not possible to directly import a GPG-encrypted configuration into SPS, it has to be decrypted locally first.

Prerequisites

You have to configure a backup policy before starting this procedure. For details, see Data and configuration backups on page 145.

You need a GPG key which must be permitted to encrypt data. Keys that can be used only for signing cannot be used to encrypt the configuration file.

To encrypt the configuration file of SPS during system backup

1. Navigate to **Basic Settings** > **Management** > **System backup**.
2. Select **Encrypt configuration**.
3. Click  .
    - To upload a key file, click **Browse**, select the file containing the public GPG key, and click **Upload**. SPS accepts both binary and ASCII-armored GPG keys.
    - To copy-paste the key from the clipboard, copy it, paste it into the **Key** field, then click **Set**.
4. Click 

.

Archiving

Archiving transfers data from SPS to an external storage solution. Archived data can be accessed and searched, but cannot be restored (moved back) to the SPS appliance. Only those closed audit-trail files are archived where the retention time has already elapsed.

To configure archiving, you first have to create an archive policy. Archive policies define the retention time, the address of the remote backup server, which protocol to use to access it, and other parameters. SPS can be configured to use the SMB/CIFS and NFS protocols to access the archive server:

- To configure archiving using SMB/CIFS, see Creating an archive policy using SMB/CIFS on page 159.
- To configure archiving using NFS, see Creating an archive policy using NFS on page 162.

CAUTION:

Hazard of data loss *Never delete an Archive Policy if data has been archived with it. This will make the already archived data inaccessible.*

Do not "remake" an Archive Policy (that is, deleting an Archive Policy and then creating another one with the same name but different parameters). This will make data inaccessible, and identifying the root cause of the issue complicated.

If you want to change the connection parameters (that is when you perform a storage server migration), you must make sure that the share contents and file permissions are kept unmodified and there are no archiving or backup tasks running.

On the other hand, if you want to add a new network share to your archives, proceed with the following steps:

1. Create a new empty SMB/NFS network share.
2. Create a new Archive Policy that points to this network share.
3. Modify your Connection Policy(es) to archive using the newly defined Archive Policy.
4. Make sure to leave the existing Archive Policy unmodified.

It is also safe to extend the size of the network share on the server side.

The different protocols assign different file ownerships to the files saved on the remote server. The owners of the archives created using the different protocols are the following:

- *SMB/CIFS*: The user provided on the web interface.
- *NFS*: `root` with `no-root-squash`, `nobody` otherwise.

CAUTION:

SPS cannot modify the ownership of a file that already exists on the remote server.

Once you have configured an archive policy, assign it to the connection you want to archive. For details, see Archiving the collected data on page 165.

Data about archived connections can be automatically deleted from the connection database. For details, see Configuring cleanup policies on page 166.

Creating an archive policy using SMB/CIFS

The **Move data to a remote server using SMB/CIFS** archive method connects to a share on the target server with Server Message Block protocol. SMB/CIFS is mainly used on Microsoft Windows Networks.

NOTE: Backup and archive policies only work with existing shares and subdirectories.

If a server has a share at, for example, `archive` and that directory is empty, when the user configures `archive/scb1` (or similar) as a backup/archive share, it will fail.

When deployed from the Azure Marketplace, you can use Azure File storage shares for Backup and Archive Policies. This is very useful as you can change the quota for the file storage dynamically, so the cumulative size of the audit trails is not limited to the OS disk

size. You can set up this share as normal SMB shares in your Backup and Archive policies. You can obtain the parameters for the policy from the Azure portal.

CAUTION:

When you try to create backups and archives from SPS to NetApp devices using the CIFS protocol, the operation may fail with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on.`

To overcome this problem, grant the SPS user "Full Control" access rights to the CIFS share on the NetApp device.

1. Navigate to **Policies** > **Backup & Archive** and click + in the **Archive policies** section to create a new archive policy.
2. Enter a name for the archive policy.
3. Enter the time when the archive process should start into the **Start time** field in HH:MM format (for example `23:00`).

   You can add the start time for additional archive processes.

   CAUTION:

   When specifying an additional start time, ensure that the previous archive process finishes before the new archive process starts.

4. To archive the data collected on + more than once a day, click +. You can schedule multiple archive times.

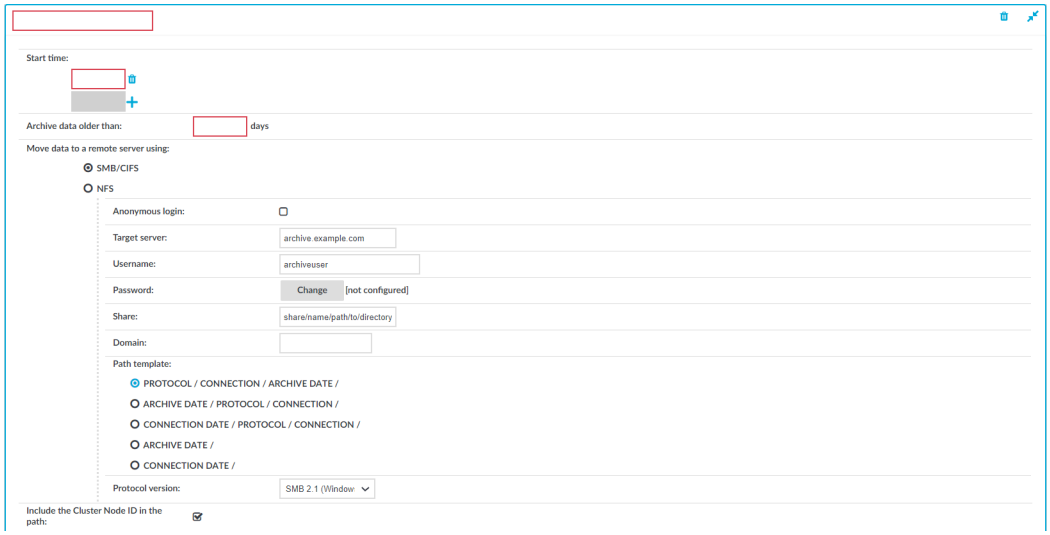
   NOTE: In case an archive process is not finished before the next one would start, the next archive process waits for the previous process to be completed.

5. Fill the **Delete data from SPS after** field. Data older than this value is archived to the external server.

   NOTE: The archived data is deleted from SPS.

6. Select **Move data to a remote server using SMB/CIFS** from the **Before deleting data from PSM** radio buttons.

Figure 64: Policies > Backup & Archive — Configuring archiving



7. Enter the username used to log on to the remote server into the **Username** field, or select the **Anonymous login** option.

   Usernames can contain space.

8. Enter the password corresponding to the username into the **Password** field.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

9. Enter the name and directory path of the share into the **Share** field. Use the following format:

   ```
   share_name/path/to/directory
   ```

   You can use backslashes and forward slashes as well.

   SPS saves all data into this directory, automatically creating the subdirectories. Archives of audit-trails are stored in the `data`, configuration backups in the `config` subdirectory.

10. Enter the domain name of the target server into the **Domain** field.

11. Select which SMB protocol to use when SPS connects to the server in the **Protocol version** field. Servers are usually backwards compatible with earlier protocol versions (for example, a server that supports version 2.1 supports versions 2.0 and 1.0 as well).

12. SPS organizes the audit trails into directories based on the date or the protocol. The subdirectories are created directly into the archive directory. Select one of the following directory structures:
    - **Protocol/Connection/Archive Date/**
    - **Archive Date/Connection/Protocol/**
    - **Connection Date/Protocol/Connection/**
    - **Archive Date/**
    - **Connection Date/**

    For example, the **Protocol/Connection/Archive Date** template will have create subdirectories for the audited protocols (that is, `ssh`, `rdp`, `telnet`, `vnc`), for the name of the connection policy, and finally, for the date (YEAR-MONTH-DAY in `YYYY-MM-DD` format).

    NOTE: **Connection Date** refers to the time the connection started, while **Archive Date** to the time it was archived. The difference between the two dates depends on the retention time set for the archiving policy.
13. When your SPS instance is a node in a cluster, select **Include the Cluster Node ID in the path**. This ensures that the node ID is included in the path of the relevant directory, which is required to prevent cluster nodes from archiving data to the same location. Archiving data to the same location would result in data loss. In addition, including the node ID in the directory name also enables easy identification.

    CAUTION:

    Hazard of data loss

    If you configured configuration synchronization across your cluster nodes, unchecking **Include the Cluster Node ID in the path** when your SPS is a node in a cluster can result in data loss.
14. To receive e-mail notifications, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab, and include the list of the files that were backed up.

    NOTE: This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see Configuring system monitoring on SPS on page 136).
15. Click **Commit**.
16. To assign the archive policy to the connection you want to archive, see Archiving the collected data on page 165.

Creating an archive policy using NFS

The **Move data to a remote server using NFS** archive method connects to a shared directory of the target server with the Network File Share protocol.

NOTE: Backup and archive policies only work with existing shares and subdirectories.

If a server has a share at, for example, `archive` and that directory is empty, when the user configures `archive/scb1` (or similar) as a backup/archive share, it will fail.

The version of NFS used is automatically detected. All versions of NFS, up to and including NFS version 4 protocol, are supported.

1. Navigate to **Policies** > **Backup & Archive** and click + in the **Archive policies** section to create a new archive policy.
2. Enter a name for the archive policy.
3. Enter the time when the archive process should start into the **Start time** field in HH:MM format (for example `23:00`).

   You can add the start time for additional archive processes.

   CAUTION:

   When specifying an additional start time, ensure that the previous archive process finishes before the new archive process starts.

4. To archive the data collected on + more than once a day, click +. You can schedule multiple archive times.

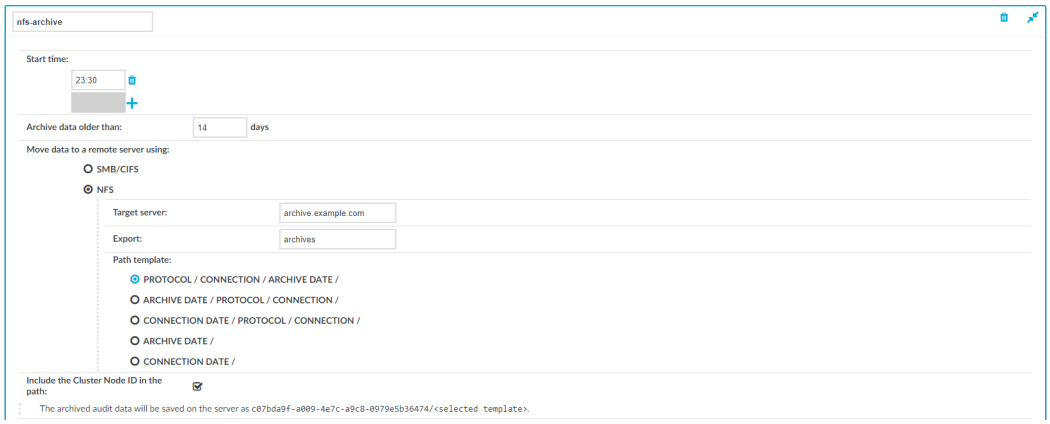
   NOTE: In case an archive process is not finished before the next one would start, the next archive process waits for the previous process to be completed.

5. Fill the **Delete data from SPS after** field. Data older than this value is archived to the external server.

   NOTE: The archived data is deleted from SPS.

6. Select **Move data to a remote server using NFS** from the **Before deleting data from PSM** radio buttons.

Figure 65: Policies > Backup & Archive —Configuring archiving



7. Enter the domain name of the remote server into the **Target server** field.
8. Enter the name of the NFS export into the **Export** field.

   SPS saves all data into this directory, automatically creating the subdirectories.

9. The remote server must also be configured to accept connections from SPS.

   Add a line that corresponds to the settings of SPS to the `/etc/exports` file of the remote server. This line should contain the following parameters:

   - The path to the archive directory as set in the **Export** field of the SPS archive policy.
   - The IP address of the SPS interface that is used to access the remote server. For more information on the network interfaces of SPS, see Network settings on page 114.

     Use an IPv4 address.
   - The following parameters: `(rw,no_root_squash,sync)`.

   Example: Configuring NFS on the remote server

   For example, if SPS connects the remote server from the `192.168.1.15` IP address and the data is saved into the `/var/backups/SPS` directory, add the following line to the `/etc/exports` file:

   ```
   /var/backups/SPS 192.168.1.15(rw,no_root_squash,sync)
   ```

   CAUTION:

   To enable non-root users to access the directories and subdirectories on the NFS server through SPS, assign them read permissions. Non-root users without a read permission cannot access the directories and subdirectories on the NFS server.

10. On the remote server, run the following command:

    ```
    exportfs -a
    ```

    Verify that the `rpc portmapper` and `rpc.statd` applications are running.

11. SPS organizes the audit trails into directories based on the date or the protocol. The subdirectories are created directly into the archive directory. Select one of the following directory structures:

    - **Protocol/Connection/Archive Date/**
    - **Archive Date/Connection/Protocol/**
    - **Connection Date/Protocol/Connection/**
    - **Archive Date/**
    - **Connection Date/**

    For example, the **Protocol/Connection/Archive Date** template will have create subdirectories for the audited protocols (that is, `ssh, rdp, telnet, vnc`), for the

name of the connection policy, and finally, for the date (YEAR-MONTH-DAY in YYYY-MM-DD format).

NOTE: **Connection Date** refers to the time the connection started, while **Archive Date** to the time it was archived. The difference between the two dates depends on the retention time set for the archiving policy.

12. When your SPS instance is a node in a cluster, select **Include the Cluster Node ID in the path**. This ensures that the node ID is included in the path of the relevant directory, which is required to prevent cluster nodes from archiving data to the same location. Archiving data to the same location would result in data loss. In addition, including the node ID in the directory name also enables easy identification.

    CAUTION:

    Hazard of data loss

    If you configured configuration synchronization across your cluster nodes, unchecking **Include the Cluster Node ID in the path** when your SPS is a node in a cluster can result in data loss.

13. To receive e-mail notifications, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab, and include the list of the files that were backed up.

    NOTE: This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see Configuring system monitoring on SPS on page 136).

14. Click **Commit**.
15. To assign the archive policy to the connection you want to archive, see Archiving the collected data on page 165.

Archiving the collected data

To configure data archiving, assign an archive policy to the connection.

Prerequisites

You have to configure an archive policy before starting this procedure. For details, see Archiving on page 158.

To assign an archive policy to the connection

1. Navigate to the connection (for example, to **Traffic Controls** > **SSH** > **Connections**).
2. Select the connection.
3. Select the archive policy you want to use in the **Archive policy** field.

4. Click 

.
5. *Optional:* To start the archiving process immediately, click **Archive now**. This functionality works only after a corresponding policy has been configured.

Cleaning up audit data

One Identity Safeguard for Privileged Sessions (SPS) can automatically archive audit trails older than a specified retention time. However, the `.zat` file and the metadata of the corresponding connections are not deleted from the SPS connection database. Deleting the stored data of old connections decreases the size of the database, making searches faster, and may also be required by certain policies or regulations.

In an audit data cleanup policy, you can specify the period after which the `zat` file and the metadata is deleted. You can also provide a lucene-like query, with which you can specify which sessions you want to delete. For example, using the query, you can create a filter for a specific protocol.

For more information, see Configuring cleanup policies.

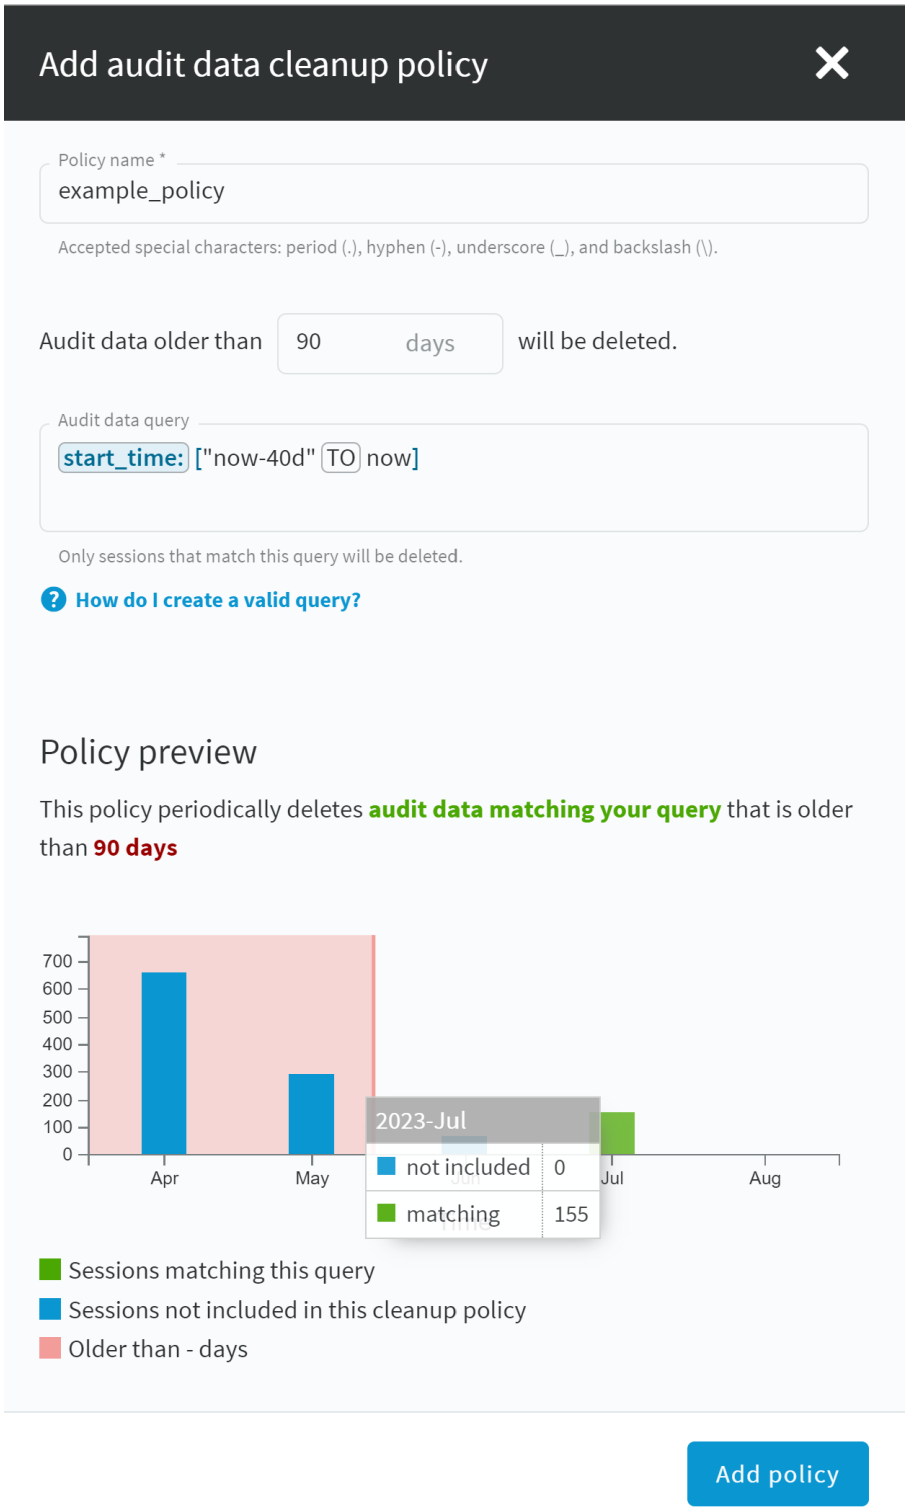
Configuring cleanup policies

In an audit data cleanup policy, you can specify the period after which the `zat` file and the metadata is deleted. You can also provide a lucene-like query, with which you can specify which sessions you want to delete. For example, using the query, you can create a filter for a specific protocol.

To add a new audit data cleanup policy

1. Navigate to **Policies** > **Audit Data Cleanup Policies**.
2. Select **Add policy**.

Figure 66: Policies > Audit Data Cleanup Policies — Configuring an audit data cleanup policy



3. In **Policy name**, specify a unique name for the audit data cleanup policy.
4. In the **Audit data older than** field, enter how long (in days) SPS must keep the `zat` file and the metadata. For example, if you specify `365`, SPS deletes the audit data of connections older than a year.

   The accepted value range is `30-100,000` days.

   NOTE: The database cleanup occurs once a day at 22:01 PM.

   NOTE: Since the database cleanup happens once a day at 22:01 PM, if you specify the same retention time for an archive policy, for example, `90` days in the **Audit data retention period** field, ensure that the archiving is set to start before 22:01 PM.
5. In **Audit data query**, which is a lucene-like query, specify to which audit data the cleanup policy applies.

   To fill this query, specify, for example, a field and the related term. Optionally, you may add a boolean operator and specify another field and related term. For example, to specify the audit data of the SSH protocol and the `ssh-connection-policy` connection policy to be cleaned up, in **Audit data query**, type `protocol:SSH AND recording.connection_policy:ssh-connection-policy`
6. To save your changes, click **Add policy**.
7. Optionally, repeat the steps to create new audit data cleanup policies for other protocols and connections.

Expected outcome

Every day, SPS deletes the `zat` file and the metadata of connections that are older than the given cleanup time from the connection database.

Preview the effect of the cleanup policy

The preview chart of a cleanup policy predicts how the respective cleanup policy will affect your audit data.

Preview charts are available in the following places:

- **Audit Data Cleanup Policies** page.

  You can preview one or more cleanup policies at the same time in one chart.
- **Add new audit data cleanup policy** and the **Edit cleanup policy** side sheets.

  You can preview the actual policy that you are creating or editing.
- By clicking the chart icons in the policy lists.

  You can preview the respective policy.

Reading the charts

- The charts display data in monthly increments.
- The vertical line or lines represent the end of the data retention period for the respective policy or policies.
- To the right of the vertical line, you can see the sessions which are not scheduled for deletion yet.
- To the left of the vertical line, you can see those sessions which are to be deleted by the next cleanup event.
- **Sessions matching this query** (green) represents the sessions which are affected by the respective cleanup policy.
- **Sessions not included in this cleanup policy** (blue) represents the sessions which are not affected by the respective cleanup policy.

Running cleanup policies immediately

You can run all audit data cleanup policies immediately. This procedure deletes all data matching the queries and older than the retention period.

Prerequisites

- You have configured audit data cleanup policies.

Running cleanup policies immediately

1. Navigate to **Policies** > **Audit Data Cleanup Policies**.
2. Select **Run all policies now**.
3. Select **Run policies** to confirm running all policies immediately in the pop-up dialog.

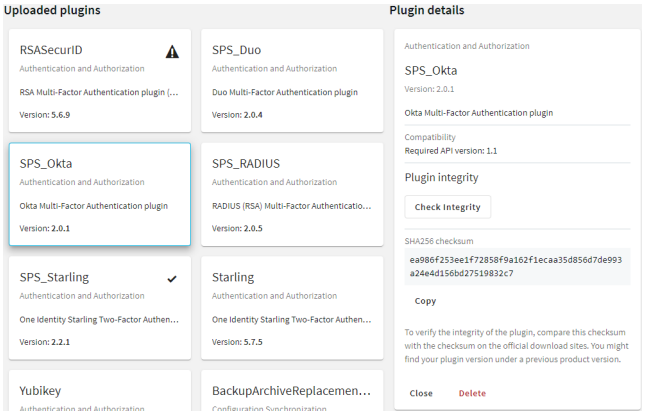
   After confirming the cleanup, you still have 10 seconds to undo the procedure in the pop-up dialog by selecting **Undo**. After 10 seconds or after selecting **Okay**, SPS completes the cleanup.

Using plugins

To download the official plugins for your product version, navigate to the product page on the Support Portal. The not officially supported plugins are also available on GitHub.

To write your own custom plugin, feel free to use our Plugin SDK.

Figure 67: Basic Settings > Plugins — Viewing the uploaded plugins



The following plugin types can be uploaded to SPS:

- Authentication and Authorization plugins

  For more information, see Integrating external authentication and authorization systems.
- Credential Store plugins

  For more information, see Using a custom Credential Store plugin to authenticate on the target hosts.
- Configuration Synchronization plugins

  For more information, see Using a configuration synchronization plugin.
- Signing CA plugins

  For more information, see Signing certificates on-the-fly.

  For more information about how to create an external Signing CA plugin, see Creating an external Signing CA.

Uploading plugins

To upload a plugin to SPS

1. Navigate to **Basic Settings** > **Plugins** and click **Upload plugin**.
2. Browse for the plugin `.zip` file and click **Open**.

   NOTE: It is not possible to upload or delete plugins if SPS is in sealed mode. For more information, see *Sealed mode* in the *Administration Guide*.

Verifying the integrity of a plugin

To verify the integrity of the plugin archive (that is, that the `.zip` file has officially been issued by One Identity and has not been tampered with before its extraction and uploading the plugin), complete the following steps. These also verify whether the plugin has been modified after upload or not.

This procedure only applies to plugins downloaded from the official repositories.

Prerequisites

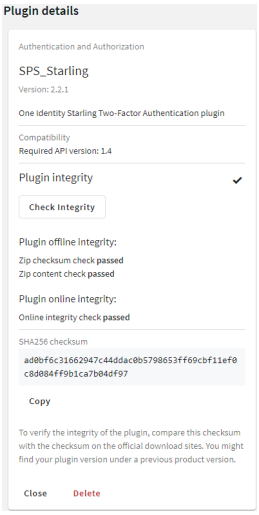
- Make sure that you have already uploaded a plugin.

  For more information on uploading plugins, see Uploading plugins.

To verify the integrity of a plugin

1. Navigate to **Basic Settings** > **Plugins**.
2. Select the plugin that you want to verify.
3. Under **Plugin details** > **Plugin integrity**, click **Check integrity**.

Figure 68: Basic Settings > Plugins > Plugin details — Verifying plugin integrity



There are three integrity checks:

- **Plugin offline integrity** > **Zip checksum check**

  This check verifies whether the recalculated checksum is the same as the checksum that has been stored in the configuration after uploading plugin.

- **Plugin offline integrity** > **Zip content check**

  This check verifies whether the plugin runtime files are the same since you have uploaded the plugin `.zip`.

- **Online integrity check**

  This check verifies whether the plugin `.zip` checksums match with the `.zip` checksums stored online.

To verify the integrity of a plugin manually

1. Under **Plugin details** > **Plugin integrity** > **SHA256 checksum**, click **Copy**.
2. To verify the integrity of the plugin, compare this checksum with the checksum on the official download sites. You might find your plugin version under a previous product version.

   On the support portal:

   a. Navigate to the product page on the Support Portal that you have downloaded the plugin from. Click on the name of the plugin.

   b. Next to the **sha256** section, you will see the checksum of the official One Identity plugin.

   On GitHub:

   a. Navigate to the GitHub plugin repository that you have downloaded the plugin from. Click on the name of the plugin.

   b. Navigate to the **releases** tab.

   c. Scroll to the specific release that you use.

   d. Under the **SHA256 checksum** section, you will see the checksum of the official One Identity plugin.
3. Compare the checksum of the official One Identity plugin with the one you have copied from **Plugin details** > **Plugin integrity** > **SHA256 checksum**.

Forwarding data to third-party systems

SPS can forward session data to Splunk, ArcSight, or other third-party systems that enable you to search, analyze, and visualize the forwarded data.

NOTE: Since SPS version 5.11, the universal SIEM forwarder supports Splunk easier than in previous versions. If you want to integrate your SPS with Splunk, One Identity recommends using the universal SIEM forwarder instead of the Splunk forwarder (which has been deprecated as of SPS version 6.4).

Using the universal SIEM forwarder

The universal SIEM forwarder can automatically send data about the audited sessions to Splunk, ArcSight, or other third-party systems. The messages are standard syslog messages in RFC3164 format (also called legacy-syslog or BSD-syslog format). The body of the syslog message (the MESSAGE part) can be formatted as JavaScript Object Notation (JSON), Common Event Format (CEF), or JSON-CIM format. For information about the details of the messages that the universal SIEM forwarder sends to the external SIEM network elements, see *Message format forwarded to SIEMs* in the *Administration Guide*.

One of the main advantages of the universal SIEM forwarder is that it has a lower impact on network and performance.

Each message contains the minimal information relevant to the event. Use the built-in correlation feature of the SIEM to combine events by session ID and view all information in one place.

Use the universal SIEM forwarder if you need a less resource-heavy solution. For more information, see Using the universal SIEM forwarder.

Using the universal SIEM forwarder

The universal SIEM forwarder can automatically send data about the audited sessions to Splunk, ArcSight, or other third-party systems. The messages are standard syslog messages in RFC3164 format (also called legacy-syslog or BSD-syslog format). The body of the syslog message (the MESSAGE part) can be formatted as JavaScript Object Notation (JSON), Common Event Format (CEF), or JSON-CIM format. For information about the details of the messages that the universal SIEM forwarder sends to the external SIEM network elements, see *Message format forwarded to SIEMs* in the *Administration Guide*.

One of the main advantages of the universal SIEM forwarder is that it has a lower impact on network and performance.

Each message contains the minimal information relevant to the event. Use the built-in correlation feature of the SIEM to combine events by session ID and view all information in one place.

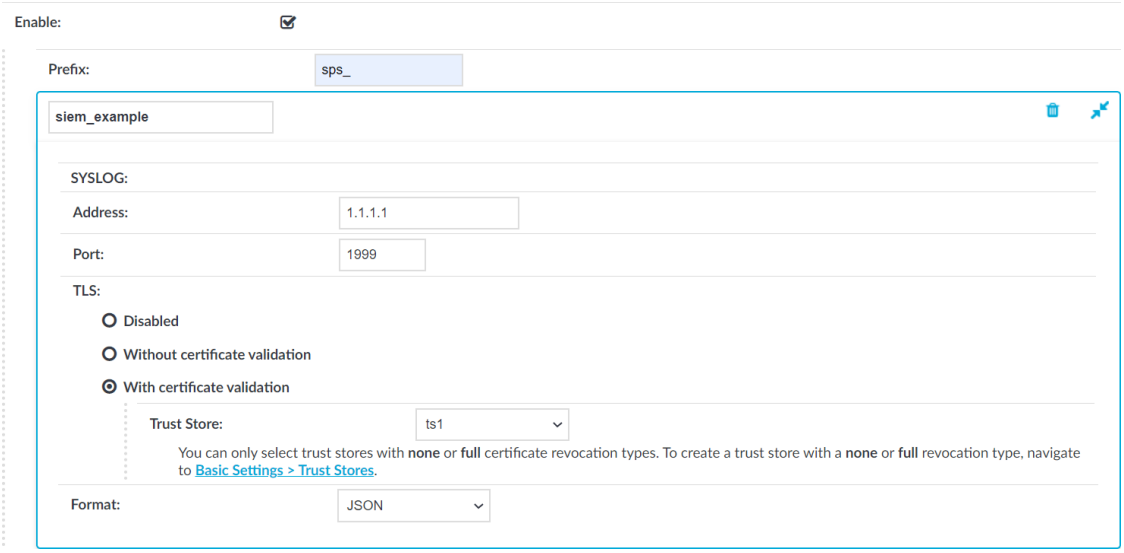
Prerequisites and restrictions

- SPS version 5 F9 or later
- Splunk version 6.5 or later
- The CEF format is supported on all currently supported versions of ArcSight ESM, IBM QRadar and Microsoft Azure Sentinel.
- SPS does not send historical data, only data from the sessions started after you complete this procedure.

To use the universal SIEM forwarder

1. Log in to SPS and navigate to **Basic Settings** > **Management** > **Universal SIEM forwarder**.

Figure 69: Basic Settings > Management > Universal SIEM forwarder — Sending session data to SIEM



2. Enter the IPv4 address or hostname of your third-party system, into the **Address** field.
3. Enter the port number where your third-party system is accepting connections into the **Port** field. For example, if you use Splunk, use port **1999**.
4. Select the appropriate **TLS** setting:
   - If your third-party system accepts unencrypted connections, select **TLS** > **Disabled**.

     Because the data forwarded contains sensitive information, One Identity recommends to use TLS encryption between SPS and your SIEM.
   - To use TLS encryption between SPS and your third-party system, select **TLS** > **Without certificate validation**.
   - To use TLS encryption between SPS and your third-party system and also verify the identity of your third-party system server, select **TLS** > **With certificate validation**, then select the trust store you want to use to validate the certificate of the third-party system in the **Trust Store** field.

     NOTE: You can only select a trust store with **None** or **Full** revocation check type.

     For more information on creating trust stores, see Verifying certificates with Certificate Authorities using trust stores.
5. Select the format of the message:
   - **JSON-CIM**: if using Splunk.
   - **CEF**: if using CEF-compatible SIEMs, for example, Microsoft Azure Sentinel.
   - **JSON**: for general use.

6. (Optional) You can specify a prefix to make the data more readable. Enter the prefix you want to use into the **Prefix** field.

   The prefix is added to each JSON key. For example, if you use **sps_** as a prefix, in the forwarded JSON message the {"protocol": "ssh"} key changes to {"**sps_**protocol": "ssh"}, which allows you to identify the forwarded data more easily.

   Other formats ignore the Prefix option.
7. Click **Commit**. From now on, SPS forwards session data to your third-party system.

Message types forwarded to SIEMs

There are three major categories of messages that One Identity Safeguard for Privileged Sessions (SPS): forwards to the SIEM: content, meta, and score.

- Content messages represents events when SPS detects interesting textual content in the session, such as a command execution or new window title.
- Meta messages represent events that change the session state and/or carry new information about a session.
- Score messages represent scoring events when SPS has calculated an initial score for the session, or updated the score for the session.

The following tables provide a summary of events for the different message types.

Content messages

Table 3: Summary of events for content messages

Event Id	Event Name	Description
127084214	CommandChannelEvent	Emitted when a command is detected in the session text.
911383355	WindowTitleChannelEvent	Emitted when a window title is detected in a graphical session.
1127618380	FileTransfer	Emitted when SCP file transfer is detected in the SSH protocol.

Meta messages

Table 4: Summary of events for meta messages

Event Id	Event Name	Description
1843867026	GatewayAuthenticationFailure	Emitted if gateway authentication is configured and the user failed to

Event Id	Event Name	Description
		authenticate through the gateway.
1865245228	ServerAuthenticationSuccess	Comes after the server authentication successfully happened.
1262825953	ServerAuthenticationFailure	Emitted if the server authentication failed.
107115592	ServerConnect	Comes after the server authentication successfully happened.
998298775	RdpEmbeddedInTsg	Emitted when the gateway user is acquired in a Terminal Service Gateway authentication scenario. This message will only contain the gateway_username optional field.
1639978560	ServerNameResolved	Emitted when the server_name field was successfully resolved to an ip address. This message will only contain the server_address optional field.
449510124	SessionClosed	Emitted when the session ends.

Score messages

Table 5: Summary of events for score messages

Event Id	Event Name	Description
1991765353	SessionScored	The message contains the aggregate score and one scoring algorithm name and score.

Message format forwarded to SIEMs

The messages are standard syslog messages in RFC3164 format (also called legacy-syslog or BSD-syslog format). The body of the syslog message (the MESSAGE part) can be formatted as one of:

- Common Event Format (CEF), based on the ArcSight CEF specification rev. 16, 22 July 2010
- JavaScript Object Notation (JSON)
- JSON-CIM format (available in SPS version 5.11 and later).

CEF

CEF (Common Event Format): the mapping to CEF will be described in terms of mapping from the JSON format to CEF. In CEF all relevant keys are present, but the value may be empty if it is not known.

Header

Here <...> is substituted with the actual values.

CEF:0|OneIdentity|SPS|<SPS_version>|<event_type_id>|<event_name>|<severity>|

Extensions

CEF extensions that are always present:

app: string, equal to Application protocol

cs1: string, equal to session_id

cs1Label: string, equal to literal "Session ID"

dst: string, equal to Destination address

duser: string, equal to Destination username

dvc: string, equal to Device address

src: equal to Source address

start: equal to timestamp

suser: equal to Source username

For details on the exact messages and the fields they contain, see CEF messages on page 178.

JSON

JSON (JavaScript Object Notation): the generated JSON structure is flat and the keys in the JSON depend on what kind of event is described. Some keys are always present in all messages. There are also keys that are message type specific, but may be missing if the related information is not available.

Keys that are always present and filled:

base_type_name: string, specifies the main category of the message, one of "meta", "content" or "score".

client_address: string, the IP address of the client.

client_name: string, the client hostname or IP address if hostname is not known.

client_port: integer, the port number of the client.

connection_policy: string, the name of the Connection Policy related to the session.

event_type_id: integer, a unique number specifying the message type (primarily for CEF).

event_name: string, the name of the event type.

gateway_username: string, the authenticated gateway username if there was a successful gateway authentication.

protocol: string, the application-level protocol.

session_id: string, the unique identifier of the session.

severity: integer, 0-10, the score of the session divided by 10 at the time of the message was created. The value is 0 if the score is not available.

timestamp: string, milliseconds since Unix epoch.

For details on the exact messages and the fields they contain, see JSON messages on page 234.

JSON-CIM

In One Identity Safeguard for Privileged Sessions (SPS) version 5.11 and later versions of SPS, the JSON-CIM external message format is also supported. The JSON-CIM format is a JSON format following Splunk's CIM field names. As a result, Splunk applications can interpret the JSON-CIM format.

Keys that are always present and filled:

dvc: string, equal to Device FQDN

event_name: string, the name of the event

product: string, the short name of the product and its version number

session_id: The unique ID of the session

_time: Timestamp when the event occurred

vendor: Contains the OneIdentity string

For details on the exact messages and the fields they contain, see JSON_CIM messages on page 283.

CEF messages

SessionClosed due to content policy violation

Description of the message: Emitted when content policy with termination action enabled is violated

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|449510124|SessionClosed|0|app=SSH cs1=svc-
9S9nqpGqdns6GAJxULWjHp-my_connection-52 cs1Label=Session ID cs2=TERMINATED
cs2Label=Verdict dhost=server.acme.com dpt=22 dst=10.170.255.206 duser=root
dvc=10.30.24.20 shost=client.acme.com spt=45928 src=10.30.0.24
start=1568639938032 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:
Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity
Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS
Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS
Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type
Example: `449510124`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message
Example: `SessionClosed`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
cs2	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `TERMINATED`

Field	Name	Scope	Present
cs2Label	Verdict label	message	always

Description: fixed to Verdict

Example: `Verdict`

ChannelAlert triggered

Description of the message: Emitted when channel alert triggered by content policy

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1244069864|ChannelAlert|0|app=SSH cs1=svc-
fPr7beYhfY11DuFUXa2628-my_connection-17 cs1Label=Session ID cs2=Commands
cs2Label=Event type cs3=sudo cs3Label=Matched regexp dst=10.170.255.206
duser=root dvc=10.30.24.20 reason=PatternMatcherRule src=10.30.0.24
start=1567600928995 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1244069864`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `ChannelAlert`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
cs2	Event type	message	sometimes

Description: the type of the event triggering the alert e.g. Command, Full screen content

Example: `Command`

Field	Name	Scope	Present
cs2Label	Event type label	message	sometimes

Description: fixed to Event type

Example: `Event type`

Field	Name	Scope	Present
cs3	Matched regexp	message	sometimes

Description: the regexp matching the content that triggered the alert

Example: `sudo`

Field	Name	Scope	Present
cs3Label	Matched regexp label	message	sometimes

Description: fixed to Matched regexp

Example: `Matched regexp`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
reason	Reason	message	sometimes

Description: the rule triggering alert

Example: `PatternMatcherRule`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

ServerConnect on initial contact

Description of the message: Emitted when SPS connects to the serverfor the first time in the session

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|107115592|ServerConnect|0|app=SSH cs1=svc-
hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com
dpt=22 dst=10.170.255.206 duser= dvc=10.30.24.20 shost=client.acme.com spt=38014
src=10.30.0.24 start=1554470650290 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `107115592`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

ServerConnect for secondary channels

Description of the message: Emitted when SPS connects to the serverfor opening further channels. The difference from initial connection is that the server user name is known and authenticated this time.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|107115592|ServerConnect|0|app=SSH cs1=svc-
hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com
dpt=22 dst=10.170.255.206 duser=root dvc=10.30.24.20 shost=client.acme.com
spt=38014 src=10.30.0.24 start=1554470650290 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `107115592`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

ServerAuthenticationSuccess

Description of the message: Emitted after the server authentication successfully happened

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1865245228|ServerAuthenticationSuccess|0|app=SSH
cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID
dhost=server.acme.com dpt=22 dst=10.170.255.206 duser=root dvc=10.30.24.20
shost=client.acme.com spt=38014 src=10.30.0.24 start=1554470652340
suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1865245228`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `ServerAuthenticationSuccess`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

ServerAuthenticationFailure

Description of the message: Emitted after the server authentication failed

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1262825953|ServerAuthenticationFailure|0|app=SSH
cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID
dhost=server.acme.com dpt=22 dst=10.170.255.206 duser=root dvc=10.30.24.20
shost=client.acme.com spt=38014 src=10.30.0.24 start=1554470652340
suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1262825953`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `ServerAuthenticationFailure`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: contains the non authenticated server username
Example: `root`

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

GatewayAuthenticationFailure

Description of the message: Emitted after a failed gateway authentication. Note that the gateway username here is not authenticated and will not be retained in further messages to avoid confusion with an authenticated gateway user.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1843867026|GatewayAuthenticationFailure|0|app=SSH
cs1=svc-mBbMWzauBWHQN9TpoZz8mD-my_connection-3 cs1Label=Session ID dhost= dpt=
dst= duser= dvc=10.30.24.20 shost=client.acme.com spt=46296 src=10.30.0.24
start=1557912667169 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1843867026`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `GatewayAuthenticationFailure`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
duser	Destination username	message	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	message	always

Description: the non authenticated gateway username

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

SessionClosed of successfully authenticated session

Description of the message: Emitted when the session ends and server authentication and any gateway authentication was successful. There may be further messages related to the session after this message due to post processing of session data!

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|449510124|SessionClosed|0|app=SSH cs2=ACCEPT
cs2Label=Verdict cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session
ID dhost=server.acme.com dpt=22 dst=10.170.255.206 duser=root dvc=10.30.24.20
shost=client.acme.com spt=38014 src=10.30.0.24 start=1554470652340
suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
cs2	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `ACCEPT`

Field	Name	Scope	Present
cs2Label	Verdict label	message	always

Description: fixed to Verdict

Example: `Verdict`

SessionClosed after a failed gateway authentication

Description of the message: Emitted when the session ends because gateway authentication failed.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|449510124|SessionClosed|0|app=SSH cs1=svc-
iiCfsG48oJG5smpuocBLAN-my_connection-25 cs1Label=Session ID dhost= dpt= dst=
duser= dvc=10.30.24.20 shost=client.acme.com spt=54632 src=10.30.0.24
start=1557913042048 suser= cs2=AUTH_FAIL cs2Label=Verdict
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
duser	Destination username	message	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
cs2	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

Field	Name	Scope	Present
cs2Label	Verdict label	message	always

Description: fixed to Verdict

Example: `Verdict`

SessionClosed after a failed server authentication

Description of the message: Emitted when the session ends because server authentication failed.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|449510124|SessionClosed|0|app=SSH cs1=svc-
iiCfsG48oJG5smpuocBLAN-my_connection-27 cs1Label=Session ID
dhost=server.acme.com dpt=22 dst=10.170.255.206 duser= dvc=10.30.24.20
shost=client.acme.com spt=55084 src=10.30.0.24 start=1557913066163
suser=gwtestauto cs2=AUTH_FAIL cs2Label=Verdict
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
duser	Destination username	message	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
cs2	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

Field	Name	Scope	Present
cs2Label	Verdict label	message	always

Description: fixed to Verdict

Example: `Verdict`

RdpEmbeddedInTsg

Description of the message: Emitted when the gateway user is acquired in a Terminal Service Gateway authentication scenario.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|998298775|RdpEmbeddedInTsg|0|app=RDP cs1=svc-
oUDm7arcL8zNb3t2CVwSQr-my_connection-44-1 cs1Label=Session ID dhost= dpt= dst=
duser= dvc=10.30.24.20 shost=client.acme.com spt=51083 src=10.30.0.24
start=1558006199668 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `998298775`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `RdpEmbeddedInTsg`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dhost	Destination host name	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
duser	Destination username	message	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dpt	Destination port	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
shost	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
suser	Source username	session	always

Description: the authenticated gateway username

Example: `gwtestauto`

Field	Name	Scope	Present
spt	Source port	session	always

Description: the port number on the client

Example: `38014`

SessionScored

Description of the message: Score messages represent scoring events when SPS has calculated an initial or changed score for the session.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1991765353|SessionScored|7|app=SSH cs1=svc-
822TNSfws1M6qixvRjQX8b-my_connection-4 cs1Label=Session ID cs2=70
cs2Label=Aggregated session score cs3=keystroke cs3Label=Scorer algorithm name
cs4=18 cs4Label=Score given by algorithm dst=10.170.255.206 duser=root
dvc=10.30.24.20 src=10.30.0.24 start=1558008998716 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1991765353`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `SessionScored`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
cs2	Aggregated score	message	always

Description: the average score from all enabled analytics algorithms

Example: `50`

Field	Name	Scope	Present
cs2Label	Aggregated score label	message	always

Description: fixed to Aggregated session score

Example: `Aggregated session score`

Field	Name	Scope	Present
cs3	Algorithm name	message	always

Description: the name of the algorithm that changed value

Example: `keystroke`

Field	Name	Scope	Present
cs3Label	Algorithm name label	message	always

Description: fixed to Scorer algorithm name

Example: `Scorer algorithm name`

Field	Name	Scope	Present
cs4	Algorithm score	message	always

Description: the new score value of the algorithm that changed value

Example: `60`

Field	Name	Scope	Present
cs4Label	Algorithm score label	message	always

Description: fixed to Score given by algorithm

Example: `Score given by algorithm`

CommandChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|127084214|CommandChannelEvent|0|app=SSH cs1=svc-
sZZoAcZZz9CbtCzTKWXgao-my_connection-0 cs1Label=Session ID cs2=exit
cs2Label=Command dst=10.170.255.206 duser=root dvc=10.30.24.20 src=10.30.0.24
start=1556287687858 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `127084214`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `CommandChannelEvent`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
cs2	Command	message	always

Description: the full command detected

Example: `exit`

Field	Name	Scope	Present
cs2Label	Command label	message	always

Description: fixed to Command

Example: `Command`

WindowTitleChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|911383355|WindowTitleChannelEvent|0|app=RDP
cs1=svc-oUDm7arcL8zNb3t2CVwSQr-my_connection-44-4 cs1Label=Session ID
cs2=Shortcut Tools Application Tools Administrative Tools cs2Label=Window title
dst=10.170.255.206 duser=Administrator dvc=10.30.24.20 src=10.30.0.24
start=1558006237095 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `911383355`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `WindowTitleChannelEvent`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
cs2	Window title	message	always

Description: the window title detected in graphical protocol

Example: `firefox`

Field	Name	Scope	Present
cs2Label	Window title label	message	always

Description: fixed to Window title

Example: `Window title`

FileTransfer

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
CEF:0|OneIdentity|SPS|5.11.0|1127618380|FileTransfer|0|act=UPLOAD app=SSH
cs1=svc-2L83Phh9J6GKLWTc881awk-my_connection-308 cs1Label=Session ID
dst=10.170.255.206 duser=root dvc=10.30.24.20 filePath=/cpuinfo fname=cpuinfo
src=10.30.0.24 start=1558023621127 suser=gwtestauto
```

The message contains the following fields.

Field	Name	Scope	Present
index 0	CEF version	product	always

Description:

Example: `CEF:0`

Field	Name	Scope	Present
index 1	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
index 2	Device product	product	always

Description: fixed to SPS

Example: `SPS`

Field	Name	Scope	Present
index 3	Device version	product version	always

Description: version of SPS

Example: `5.11.0`

Field	Name	Scope	Present
index 4	Signature ID	message	always

Description: numeric identifier of message type

Example: `1127618380`

Field	Name	Scope	Present
index 5	Name	message	always

Description: the type of the message

Example: `FileTransfer`

Field	Name	Scope	Present
index 6	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
index 7	CEF extensions	product	always

Description: contains the payload in key-value form

Example: `app=SSH cs1=svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0 cs1Label=Session ID dhost=server.acme.com dpt=22 ...`

Field	Name	Scope	Present
start	Start time	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
dvc	Device address	device	always

Description: IP address of SPS

Example: `10.30.24.20`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
cs1	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
cs1Label	Session ID label	product	always

Description: fixed to Session ID

Example: `Session ID`

Field	Name	Scope	Present
dst	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
duser	Destination username	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
src	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
suser	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
act	Operation	message	always

Description: the operation on the file such as UPLOAD/DOWNLOAD. It may contain the suffix 'WARNING', if the operation failed

Example: `UPLOAD`

Field	Name	Scope	Present
fname	Filename	message	always

Description: the file name

Example: `foobar.txt`

Field	Name	Scope	Present
filePath	Full file path	message	always

Description: the name of the file including its path on the server (in case of RDP protocol, this field is empty, in this case the full path of the file is in the filename field)

Example: `/tmp/foobar.txt`

JSON messages

SessionClosed due to content policy violation

Description of the message: Emitted when content policy with termination action enabled is violated

Example message:


```
{"verdict":"TERMINATED","timestamp":"1568640063579","severity":"0","session_
id":"svc-9S9nqpGqdns6GAJxULWjHp-my_connection-53","server_
username":"root","server_port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","event_type_id":"449510124","event_
name":"SessionClosed","connection_policy":"my_connection","client_
port":"45946","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta","auth_method":"password","gateway_
username":"gwtestauto"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol
Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
auth_method	Authentication method	session	always

Description: the type of authentication used in gateway authentication

Example: `password`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `TERMINATED`

ChannelAlert triggered

Description of the message: Emitted when channel alert triggered by content policy

Example message:


```
{"event_type_id":"1244069864","event_name":"ChannelAlert","session_id":"svc-
eyKp4M2pDBpbwHW4nCSe36-my_connection-
14","severity":"0","timestamp":"1567509110329","server_username":"root",
"gateway_username":"gwtestauto","server_name":"server.acme.com","server_
address":"10.170.255.206","server_port":"22","client_
name":"client.acme.com","client_address":"10.30.0.24","client_
port":"56988","protocol":"SSH","connection_policy":"my_connection","base_type_
name":"content_alert","alerting_type":"adp.event.command","matched_
regexp":"sudo","matched_content":"sudo","rule_name":"PatternMatcherRule"}
```


The message contains the following fields.

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1244069864`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ChannelAlert`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: content_alert

Example: `content_alert`

Field	Name	Scope	Present
alerting_type	Event type	message	sometimes

Description: the type of the event triggering the alert e.g. Command, Full screen content

Example: `Command`

Field	Name	Scope	Present
matched_regexp	Matched regexp	message	sometimes

Description: the regexp matching the content that triggered the alert

Example: `sudo`

Field	Name	Scope	Present
matched_content	Matched content	message	sometimes

Description: the screen content violating channel policy

Example: `$ sudo`

Field	Name	Scope	Present
rule_name	Reason	message	sometimes

Description: the rule triggering alert

Example: `PatternMatcherRule`

ServerConnect on initial contact

Description of the message: Emitted when SPS connects to the serverfor the first time in the session

Example message:


```
{"timestamp":"1557913242888","severity":"0","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-43","server_port":"22","server_
name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"107115592","event_
name":"ServerConnect","connection_policy":"my_connection","client_
port":"59190","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `107115592`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol
Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name
Example: `my_connection`

ServerConnect for secondary channels

Description of the message: Emitted when SPS connects to the serverfor opening further channels. The difference from initial connection is that the server user name is known and authenticated this time.

Example message:


```
{"timestamp":"1557913242888","severity":"0","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-43","server_port":"22","server_
name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"107115592","event_
name":"ServerConnect","connection_policy":"my_connection","server_
username":"root","client_port":"59190","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `107115592`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

ServerAuthenticationSuccess

Description of the message: Emitted after the server authentication successfully happened

Example message:

```
{"timestamp":"1557913243423","severity":"0","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-43","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"1865245228","event_
name":"ServerAuthenticationSuccess","connection_policy":"my_connection","client_
port":"59190","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```

The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1865245228`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerAuthenticationSuccess`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

ServerAuthenticationFailure

Description of the message: Emitted after the server authentication failed

Example message:


```
{"timestamp":"1557913134598","severity":"0","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-33","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"1262825953","event_
name":"ServerAuthenticationFailure","connection_policy":"my_connection","client_
port":"56692","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1262825953`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerAuthenticationFailure`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: contains the non authenticated server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the non authenticated server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

GatewayAuthenticationFailure

Description of the message: Emitted after a failed gateway authentication. Note that the gateway username here is not authenticated and will not be retained in further messages to avoid confusion with an authenticated gateway user.

Example message:

```
{"timestamp":"1557913110027","severity":"0","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-31","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"1843867026","event_
name":"GatewayAuthenticationFailure","connection_policy":"my_
connection","client_port":"56020","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1843867026`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `GatewayAuthenticationFailure`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred
Example: `1554470652340`

Field	Name	Scope	Present
gateway_username	Gateway username	message	always

Description: the non authenticated gateway username
Example: `gwtestauto`

Field	Name	Scope	Present
gateway_username	Gateway user domain	session	sometimes

Description: the non authenticated gateway user domain if known
Example: `acme.com`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

SessionClosed of successfully authenticated session

Description of the message: Emitted when the session ends and server authentication and any gateway authentication was successful. There may be further messages related to the session after this message due to post processing of session data!

Example message:


```
{"timestamp":"1557912701233","severity":"0","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-6","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"449510124","event_
name":"SessionClosed","connection_policy":"my_connection","client_
port":"46958","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta","auth_
method":"password","verdict":"ACCEPT"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
auth_method	Authentication method	session	always

Description: the type of authentication used in gateway authentication

Example: `password`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `ACCEPT`

SessionClosed after a failed gateway authentication

Description of the message: Emitted when the session ends because gateway authentication failed.

Example message:

```
{"timestamp":"1557912725391","severity":"0","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-9","protocol":"SSH","event_type_
id":"449510124","event_name":"SessionClosed","connection_policy":"my_
connection","client_port":"47444","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta","verdict":"AUTH_FAIL"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred
Example: `1554470652340`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol
Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name
Example: `my_connection`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

SessionClosed after a failed server authentication

Description of the message: Emitted when the session ends because server authentication failed.

Example message:


```
{"timestamp":"1557912748990","severity":"0","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-11","verdict":"AUTH_FAIL","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"449510124","event_
name":"SessionClosed","connection_policy":"my_connection","client_
port":"47840","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `449510124`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol
Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name
Example: `my_connection`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

RdpEmbeddedInTsg

Description of the message: Emitted when the gateway user is acquired in a Terminal Service Gateway authentication scenario.

Example message:

```
{"timestamp":"1558007294417","severity":"0","session_id":"svc-
oUDm7arcL8zNb3t2CVwSQr-my_connection-50-4","protocol":"RDP","gateway_
username":"gwtestauto","event_type_id":"998298775","event_
name":"RdpEmbeddedInTsg","connection_policy":"my_connection","client_
port":"51270","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"meta"}
```

The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: meta

Example: `meta`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `998298775`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `RdpEmbeddedInTsg`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
gateway_username	Gateway username	session	always

Description: the authenticated gateway username

Example: `gwtestauto`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

SessionScored

Description of the message: Score messages represent scoring events when SPS has calculated an initial or changed score for the session.

Example message:


```
{"timestamp":"1558009822701","severity":"7","session_id":"svc-
62a6XGcPzaFvLYDhVYDYXj-my_connection-0","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"1991765353","event_
name":"SessionScored","connection_policy":"my_connection","client_
port":"35620","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"score","algorithm_
score":"18","algorithm_name":"keystroke","aggregated_score":"70"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: score

Example: `score`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1991765353`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionScored`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
aggregated_score	Aggregated score	message	always

Description: the average score from all enabled analytics algorithms

Example: `50`

Field	Name	Scope	Present
algorithm_name	Algorithm name	message	always

Description: the name of the algorithm that changed value

Example: `keystroke`

Field	Name	Scope	Present
algorithm_score	Algorithm score	message	always

Description: the new score value of the algorithm that changed value

Example: `60`

CommandChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:


```
{"timestamp":"1557912701166","severity":"0","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-6","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","event_type_id":"127084214","event_
name":"CommandChannelEvent","connection_policy":"my_
connection","command":"exit","client_port":"46958","client_
name":"client.acme.com","client_address":"10.30.0.24","base_type_
name":"content"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: content

Example: `content`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `127084214`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `CommandChannelEvent`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
command	Command	message	always

Description: the full command detected

Example: `exit`

WindowTitleChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:


```
{"window_title":"Shortcut Tools Application Tools Administrative
Tools","timestamp":"1558007305516","severity":"0","session_id":"svc-
oUDm7arcL8zNb3t2CVwSQr-my_connection-50-4","server_
username":"Administrator","server_port":"3389","server_
name":"server.acme.com","server_
```

```
address":"10.170.255.206","protocol":"RDP","gateway_
username":"gwtestauto","event_type_id":"911383355","event_
name":"WindowTitleChannelEvent","connection_policy":"my_connection","client_
port":"51270","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"content"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: content

Example: `content`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `911383355`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `WindowTitleChannelEvent`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client
Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol
Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
window_title	Window title	message	always

Description: the window title detected in graphical protocol

Example: `firefox`

FileTransfer

Description of the message: Emitted when a command is detected in the session channel text.

Example message:


```
{"timestamp":"1558023671115","severity":"0","session_id":"svc-
2L83Phh9J6GKLWTc881awk-my_connection-316","server_username":"root","server_
port":"22","server_name":"server.acme.com","server_
address":"10.170.255.206","protocol":"SSH","gateway_
username":"gwtestauto","filepath":"","filename":"cpuinfo","file_
operation":"UPLOAD","event_type_id":"1127618380","event_
name":"FileTransfer","connection_policy":"my_connection","client_
port":"44292","client_name":"client.acme.com","client_
address":"10.30.0.24","base_type_name":"content"}
```


The message contains the following fields.

Field	Name	Scope	Present
base_type_name	Basic type	message	always

Description: basic message type: content

Example: `content`

Field	Name	Scope	Present
event_type_id	Signature ID	message	always

Description: numeric identifier of message type

Example: `1127618380`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `FileTransfer`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
severity	Severity	message	always

Description: number between 0-10 inclusive, equal to aggregated analytics score divided by 10 or 0 if analytics is disabled

Example: `0`

Field	Name	Scope	Present
timestamp	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
server_username	Server user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
server_domain	Server user domain if known	session	sometimes

Description: the server domain, if known

Example: `acme.com`

Field	Name	Scope	Present
gateway_username	Gateway username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
gateway_domain	Gateway user domain	session	sometimes

Description: the authenticated gateway user domain if there was a successful gateway authentication and known

Example: `acme.com`

Field	Name	Scope	Present
server_name	Server name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
server_address	Server address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
server_port	Server port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
client_name	Client name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
client_address	Client address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
client_port	Client port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
protocol	Application protocol	session	always

Description: SPS supported protocol

Example: `SSH`

Field	Name	Scope	Present
connection_policy	Connection policy name	session	always

Description: SPS connection policy name

Example: `my_connection`

Field	Name	Scope	Present
file_operation	Operation	message	always

Description: the operation on the file such as UPLOAD/DOWNLOAD. It may contain the suffix 'WARNING', if the operation failed

Example: `UPLOAD`

Field	Name	Scope	Present
filename	Filename	message	always

Description: the file name

Example: `foobar.txt`

Field	Name	Scope	Present
filepath	File path	message	always

Description: the path to the file on the server

Example: `/tmp`

JSON_CIM messages

SessionClosed due to content policy violation

Description of the message: Emitted when content policy with termination action enabled is violated

Example message:


```
{"verdict":"TERMINATED","vendor":"OneIdentity","user":"root","transport":"tcp","
src_user":"gwtestauto","src_port":"57542","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
w6rJcFNZ3c6Bqqu2pAoeoS-my_connection-1","product":"SPS-5.11.0","event_
name":"SessionClosed","dvc":"sps1.acme.com","dest_port":"22","dest_
ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","_
time":"1568984418014"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message
Example: `SessionClosed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred
Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: the server username
Example: `root`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `TERMINATED`

ChannelAlert triggered

Description of the message: Emitted when channel alert triggered by content policy

Example message:


```
{"vendor":"OneIdentity","user":"root","transport":"tcp","subject":"PatternMatcherRule","src_user":"gwtestauto","src_port":"57542","src_ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-w6rJcFNZ3c6Bqqu2pAoeoS-my_connection-1","product":"SPS-5.11.0","matched_regexp":"free","event_name":"ChannelAlert","dvc":"sps1.acme.com","dest_port":"22","dest_ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","alerting_type":"Full screen content","_time":"1568984413910"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `ssh`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client

Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
matched_regexp	Matched regexp	message	sometimes

Description: the regexp matching the content that triggered the alert

Example: `sudo`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ChannelAlert`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server

Example: `22`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server

Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known

Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: the server username

Example: `root`

Field	Name	Scope	Present
alerting_type	Event type	message	sometimes

Description: the type of the event triggering the alert e.g. Command, Full screen content

Example: `Command`

Field	Name	Scope	Present
subject	Reason	message	sometimes

Description: the rule triggering alert

Example: `PatternMatcherRule`

ServerConnect on initial contact

Description of the message: Emitted when SPS connects to the serverfor the first time in the session

Example message:

```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_
user":"gwtestauto","src_port":"58140","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-39","product":"SPS-5.11.0","event_
name":"ServerConnect","dvc":"sps1.acme.com","dest_port":"22","dest_
ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","action":"added","_
time":"1557913195000"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
action	Action	message	always

Description: the taken by the device according to CIM model

Example: `added`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

ServerConnect for secondary channels

Description of the message: Emitted when SPS connects to the serverfor opening further channels. The difference from initial connection is that the server user name is known and authenticated this time.

Example message:


```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_
user":"gwtestauto","src_port":"58140","src_
ip":"10.30.0.24","src":"client.acme.com","user":"root","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-39","product":"SPS-5.11.0","event_
name":"ServerConnect","dvc":"sps1.acme.com","dest_port":"22","dest_
ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","action":"added","_
time":"1557913195000"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerConnect`

Field	Name	Scope	Present
action	Action	message	always

Description: the taken by the device according to CIM model

Example: `added`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: the server username
Example: `root`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

ServerAuthenticationSuccess

Description of the message: Emitted after the server authentication successfully happened

Example message:


```
{"vendor":"OneIdentity","user":"root","transport":"tcp","src_
user":"gwtestauto","src_port":"57982","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-38","product":"SPS-5.11.0","event_
name":"ServerAuthenticationSuccess","dvc":"sps1.acme.com","dest_
port":"22","dest_
ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","action":"success","_
time":"1557913189329"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerAuthenticationSuccess`

Field	Name	Scope	Present
action	Action	message	always

Description: marks a successful authentication

Example: `success`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: the server username
Example: `root`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

ServerAuthenticationFailure

Description of the message: Emitted after the server authentication failed

Example message:


```
{"vendor":"OneIdentity","user":"root","transport":"tcp","src_
user":"gwtestauto","src_port":"58140","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
iiCfsG48oJG5smpuocBLAN-my_connection-39","product":"SPS-5.11.0","event_
name":"ServerAuthenticationFailure","dvc":"sps1.acme.com","dest_
port":"22","dest_
ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","action":"failure","_
time":"1557913197211"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `ServerAuthenticationFailure`

Field	Name	Scope	Present
action	Action	message	always

Description: marks a failed authentication

Example: `failure`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: contains the non authenticated server username
Example: `root`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

GatewayAuthenticationFailure

Description of the message: Emitted after a failed gateway authentication. Note that the gateway username here is not authenticated and will not be retained in further messages to avoid confusion with an authenticated gateway user.

Example message:


```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_
user":"gwtestauto","src_port":"49070","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-15","product":"SPS-5.11.0","event_
name":"GatewayAuthenticationFailure","dvc":"sps1.acme.com","dest_port":"","dest_
ip":"","dest":"","app":"ssh","action":"failure","_time":"1557912792360"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `GatewayAuthenticationFailure`

Field	Name	Scope	Present
action	Action	message	always

Description: marks a failed authentication

Example: `failure`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
user	Name of the user	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred
Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	message	always

Description: the non authenticated gateway username

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

SessionClosed of successfully authenticated session

Description of the message: Emitted when the session ends and server authentication and any gateway authentication was successful. There may be further messages related to the session after this message due to post processing of session data!

Example message:


```
{"vendor":"OneIdentity","user":"root","transport":"tcp","src_
user":"gwtestauto","src_port":"48302","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-12","product":"SPS-5.11.0","event_
name":"SessionClosed","dvc":"sps1.acme.com","verdict":"ACCEPT","dest_
port":"22","dest_ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","_
time":"1557912765545"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	session	always

Description: the server username
Example: `root`

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `ACCEPT`

SessionClosed after a failed gateway authentication

Description of the message: Emitted when the session ends because gateway authentication failed.

Example message:

```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_user":"","src_
port":"49070","src_ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-15","product":"SPS-5.11.0","event_
name":"SessionClosed","dvc":"sps1.acme.com","dest_port":"","dest_
ip":"","dest":"","app":"ssh","_time":"1557912792398","verdict":"AUTH_FAIL"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
user	Name of the user	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	always

Description: empty, not known in this message type

Example:

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

SessionClosed after a failed server authentication

Description of the message: Emitted when the session ends because server authentication failed.

Example message:

```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_
user":"gwtestauto","src_port":"49426","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
mBbMWzauBWHQN9TpoZz8mD-my_connection-17","product":"SPS-5.11.0","event_
name":"SessionClosed","dvc":"sps1.acme.com","dest_port":"22","verdict":"AUTH_
FAIL","dest_ip":"10.170.255.206","dest":"server.acme.com","app":"ssh","_
time":"1557912813792"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionClosed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol

Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: the IP address of the server
Example: `10.170.255.206`

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: the server hostname or IP address if hostname is not known
Example: `server.acme.com`

Field	Name	Scope	Present
user	Name of the user	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: the port number on the server
Example: `22`

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known
Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	sometimes

Description: the authenticated gateway username if there was a successful gateway authentication

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

Field	Name	Scope	Present
verdict	Verdict	session	always

Description: describes how the session ended, e.g. ACCEPT, AUTH_FAIL, DENY, FAIL, TERMINATED

Example: `AUTH_FAIL`

RdpEmbeddedInTsg

Description of the message: Emitted when the gateway user is acquired in a Terminal Service Gateway authentication scenario.

Example message:

```
{"vendor":"OneIdentity","user":"","transport":"tcp","src_
user":"gwtestauto","src_port":"51204","src_
ip":"10.30.0.24","src":"client.acme.com","session_id":"svc-
oUDm7arcL8zNb3t2CVwSQr-my_connection-47-4","product":"SPS-5.11.0","event_
name":"RdpEmbeddedInTsg","dvc":"sps1.acme.com","dest_port":"","dest_
ip":"","dest":"","app":"rdp","action":"allowed","_time":"1558006936608"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `RdpEmbeddedInTsg`

Field	Name	Scope	Present
action	Action	message	always

Description: the action taken by the device according to CIM model

Example: `allowed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
app	Application protocol	session	always

Description: SPS supported protocol
Example: `ssh`

Field	Name	Scope	Present
dest_ip	Destination address	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest	Destination host name	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
user	Name of the user	message	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
dest_port	Destination port	session	always

Description: empty, not known in this message type
Example:

Field	Name	Scope	Present
src_ip	Source address	session	always

Description: the IP address of the client
Example: `10.30.0.24`

Field	Name	Scope	Present
src	Source host name	session	always

Description: the client hostname or IP address if hostname is not known

Example: `client.acme.com`

Field	Name	Scope	Present
src_user	Source username	session	always

Description: the authenticated gateway username

Example: `gwtestauto`

Field	Name	Scope	Present
src_port	Source port	session	always

Description: the port number on the client

Example: `38014`

Field	Name	Scope	Present
transport	Transport	session	always

Description: the layer 3 protocol

Example: `tcp`

SessionScored

Description of the message: Score messages represent scoring events when SPS has calculated an initial or changed score for the session.

Example message:


```
{"vendor":"OneIdentity","signature":"keystroke","session_id":"svc-
416YVFZMy7rT8RA7T7yeAs-my_connection-0","product":"SPS-5.11.0","event_
name":"SessionScored","dvc":"sps1.acme.com","algorithm_score":"18","algorithm_
name":"keystroke","aggregated_score":"70","action":"allowed","_
time":"1558010880806"}
```


The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `SessionScored`

Field	Name	Scope	Present
action	Action	message	always

Description: the action taken by the device according to CIM model

Example: `allowed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
aggregated_score	Aggregated score	message	always

Description: the average score from all enabled analytics algorithms

Example: `50`

Field	Name	Scope	Present
algorithm_name	Algorithm name	message	always

Description: the name of the algorithm that changed value

Example: `keystroke`

Field	Name	Scope	Present
signature	Signature	message	always

Description: the algorithm name as CIM intrusion detection signature

Example: `hostlogin`

Field	Name	Scope	Present
algorithm_score	Algorithm score	message	always

Description: the new score value of the algorithm that changed value

Example: `60`

CommandChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
{"vendor":"OneIdentity","session_id":"svc-mBbMWzauBWHQN9TpoZz8mD-my_connection-12","product":"SPS-5.11.0","event_name":"CommandChannelEvent","dvc":"sps1.acme.com","command":"exit","action":"allowed","_time":"1557912765461"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `CommandChannelEvent`

Field	Name	Scope	Present
action	Action	message	always

Description: the action taken by the device according to CIM model

Example: `allowed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
command	Command	message	always

Description: the full command detected

Example: `exit`

WindowTitleChannelEvent

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
{"window_title":"Shortcut Tools Application Tools Administrative
Tools","vendor":"OneIdentity","session_id":"svc-oUDm7arcL8zNb3t2CVwSQr-my_
connection-47-4","product":"SPS-5.11.0","event_
name":"WindowTitleChannelEvent","dvc":"sps1.acme.com","action":"allowed","_
time":"1558007001482"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `WindowTitleChannelEvent`

Field	Name	Scope	Present
action	Action	message	always

Description: the action taken by the device according to CIM model

Example: `allowed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
window_title	Window title	message	always

Description: the window title detected in graphical protocol

Example: `firefox`

FileTransfer

Description of the message: Emitted when a command is detected in the session channel text.

Example message:

```
{"vendor":"OneIdentity","session_id":"svc-2L83Phh9J6GKLWTc881awk-my_connection-
324","product":"SPS-5.11.0","file_path":"/cpuinfo","file_
operation":"UPLOAD","file_name":"cpuinfo","event_
name":"FileTransfer","dvc":"sps1.acme.com","action":"allowed","_
time":"1558023721326"}
```

The message contains the following fields.

Field	Name	Scope	Present
vendor	Device vendor	product	always

Description: fixed to OneIdentity

Example: `OneIdentity`

Field	Name	Scope	Present
product	Product version	product	always

Description: short product name with version

Example: `SPS-5.11.0`

Field	Name	Scope	Present
dvc	Device fqdn	device	always

Description: the hostname of SPS

Example: `sps1.acme.com`

Field	Name	Scope	Present
session_id	Session ID	session	always

Description: the unique identifier of the session

Example: `svc-hjdBxA2UWkTadH3juDVwrT-my_connection-0`

Field	Name	Scope	Present
event_name	Event name	message	always

Description: the type of the message

Example: `FileTransfer`

Field	Name	Scope	Present
action	Action	message	always

Description: the action taken by the device according to CIM model

Example: `allowed`

Field	Name	Scope	Present
_time	Timestamp	message	always

Description: the UNIX time stamp when the event occurred

Example: `1554470652340`

Field	Name	Scope	Present
file_operation	Operation	message	always

Description: the operation on the file such as UPLOAD/DOWNLOAD. It may contain the suffix 'WARNING', if the operation failed

Example: `UPLOAD`

Field	Name	Scope	Present
file_name	Filename	message	always

Description: the file name

Example: `foobar.txt`

Field	Name	Scope	Present
file_path	Full file path	message	always

Description: the name of the file including its path on the server (in case of RDP protocol, this field is empty, in this case the full path of the file is in the filename field)

Example: `/tmp/foobar.txt`

User management and access control

The **Users & Access Control** menu (previously named **AAA** menu) allows you to configure multiple login options and to control the authentication, authorization, and accounting settings of users accessing One Identity Safeguard for Privileged Sessions (SPS). The following topics are detailed in the next sections:

- How to authenticate locally on SPS. For more information, see Managing One Identity Safeguard for Privileged Sessions (SPS) users locally on page 328.
- How to authenticate users using an external LDAP (for example Microsoft Active Directory) database. For more information, see Managing One Identity Safeguard for Privileged Sessions (SPS) users from an LDAP database on page 335.
- How to authenticate users using an external RADIUS server. For more information, see Authenticating users to a RADIUS server on page 347.
- How to authenticate users with X.509 certificates. For more information, see Authenticating users with X.509 certificates.
- How to control the privileges of users and usergroups. For more information, see Managing user rights and user groups on page 361.
- How to display the history of changes of SPS configuration. For more information, see Listing and searching configuration changes on page 388.

Login settings

You can configure the following login options under **Users & Access Control** > **Settings**:

- **Protect against brute-force attack**

  By default, the login addresses are protected against brute-force attacks. After the users reach the configured number of unsuccessful login attempts, SPS denies all following attempts for the configured time.

  For more information, see Protecting against brute-force attacks.
- **Authentication banner**

On the web and console login screen of SPS, you can display a banner that the users see every time they try to log in to SPS.

For more information, see Authentication banner.

- **Web interface timeout**

  You can configure the time after which SPS automatically logs the users out of the session.

  For more information, see Web interface timeout.

Protecting against brute-force attacks

This section describes the **Protect against brute-force attacks** option that you can configure on **Users & Access Control** > **Settings**.

How the Protect against brute-force attacks option works

By default, the web login addresses of administrators and users are protected against brute-force attacks. After the users reach the configured number of unsuccessful login attempts, SPS denies all following attempts for the configured time.

The **Protect against brute-force attacks** option blocks the user name or the IP address based on the following logic:

- If the number of unsuccessful login attempts from the same IP address with any user name exceeds the threshold, the IP address is blocked.
- If the number of unsuccessful login attempts with a user name from different IP addresses exceeds the configured threshold, the user name is blocked for all IP addresses.

The authentication attempts rejected by SPS during the blocking do not increase the lockout counters.

NOTE: The `admin` user is also subject to brute-force attack protection.

By default, **Protect against brute-force attacks** blocks the user name or the IP address for 10 minutes after 20 unsuccessful login attempts.

Accepted values:

- **Attempt limit**: `1-50` attempts
- **Lockout period**: `1-720` minutes

Blocked users receive the `Unable to authenticate` error message, regardless of whether they enter valid or invalid credentials.

NOTE: The `Unable to authenticate` error message does not provide details about the error and the possible solutions to prevent providing information for attackers.

Log messages about blocked user names and IP addresses

If a user name or an IP address is blocked, a log event is created, which provides the details about the blocking. The log event contains the following information:

- Cause of the blocking
- User name
- IP address
- Duration of the blocking

Example: log message about a blocked user name

The following example provides the details about the blocking of a user name. The blocked user name is `admin` and the IP address used is `1.2.3.4`. The reason for the blocking is that the user has exceeded the allowed number of unsuccessful authentication attempts. This user is blocked for `60` minutes.

```
Authentication denied, too many attempts, username is locked out;
username='admin', remote_address='1.2.3.4', lockout='60 min'
```

Example: log message about a blocked IP address

The following example provides the details about the blocking of an IP address. The user is `admin` and the blocked IP address is `1.2.3.4`. The reason for the blocking is that the allowed number of unsuccessful authentication attempts has been reached from this IP address. This IP address is blocked for `40` minutes.

```
Authentication denied, too many attempts, remote_address is locked out;
username='admin', remote_address='1.2.3.4', lockout='40 min'
```

Unblocking blocked user names and IP addresses

SPS resets the web lockout counter for a user name or IP address if:

- The lockout period is over.
- The server is rebooted.
- The secondary node becomes active after a High Availability (HA) failover.
- After the `root` user clears the list of blocked users/IP addresses on the **Troubleshooting** page of the text-based physical or SSH console.

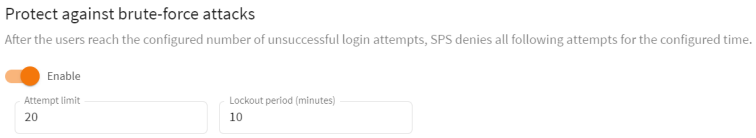
NOTE: If you are the root user, on the **Troubleshooting** page of the text-based physical or SSH console, you can clear the list of blocked user names and IP addresses using the **Clear list of blocked users/IPs** option. If you clear the list, users and IP addresses that previously were blocked due to exceeding the allowed number of web login attempts can attempt logging in again. Clearing the list does not disable the **Protect against brute-force attacks** option.

Configuring the Protect against brute-force attack option

To configure the Protect against brute-force attacks option

1. Navigate to **Users & Access Control** > **Settings**— **Protect against brute-force attacks**.

   Figure 70: Users & Access Control > Settings — Protect against brute-force attacks

2. (Optional) Modify the default values of **Attempt limit**, **Lockout period**, or both, according to your security requirements.
3. To save the modifications, click **Commit changes**.

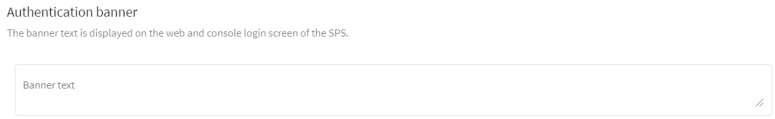
Authentication banner

You can display a banner with a configurable text on the web and console login screen of One Identity Safeguard for Privileged Sessions (SPS).

Users will see the banner every time they try to log in to SPS. The login screen displays the banner text as plain text, with whitespaces preserved.

To enter a banner message, navigate to **Users & Access Control** > **Settings** — **Authentication banner**.

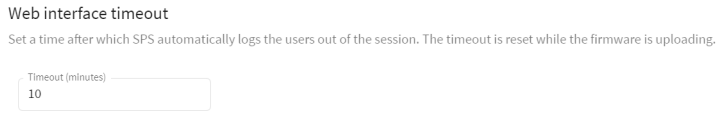
Figure 71: Users & Access Control > Settings> Authentication banner — Enter a banner message



Web interface timeout

By default, One Identity Safeguard for Privileged Sessions (SPS) terminates the web session of a user after 10 minutes of inactivity. To change the value of this timeout, adjust the **Users & Access Control** > **Settings** > **Web interface timeout** option.

Figure 72: Users & Access Control > Settings — Web interface timeout



Managing One Identity Safeguard for Privileged Sessions (SPS) users locally

By default, One Identity Safeguard for Privileged Sessions (SPS) users are managed locally on SPS. To add local users in SPS, complete all steps of the following procedure:

1. Create users.

   For detailed instructions on how to create local users, see Creating local users in One Identity Safeguard for Privileged Sessions (SPS) on page 328.

2. Assign users to groups.

   For details about how to add a usergroup, see Managing local user groups on page 333.

3. Assign privileges to groups.

   For information on how to control the privileges of usergroups, see Managing user rights and user groups on page 361.

Creating local users in One Identity Safeguard for Privileged Sessions (SPS)

This section describes how to create local users.

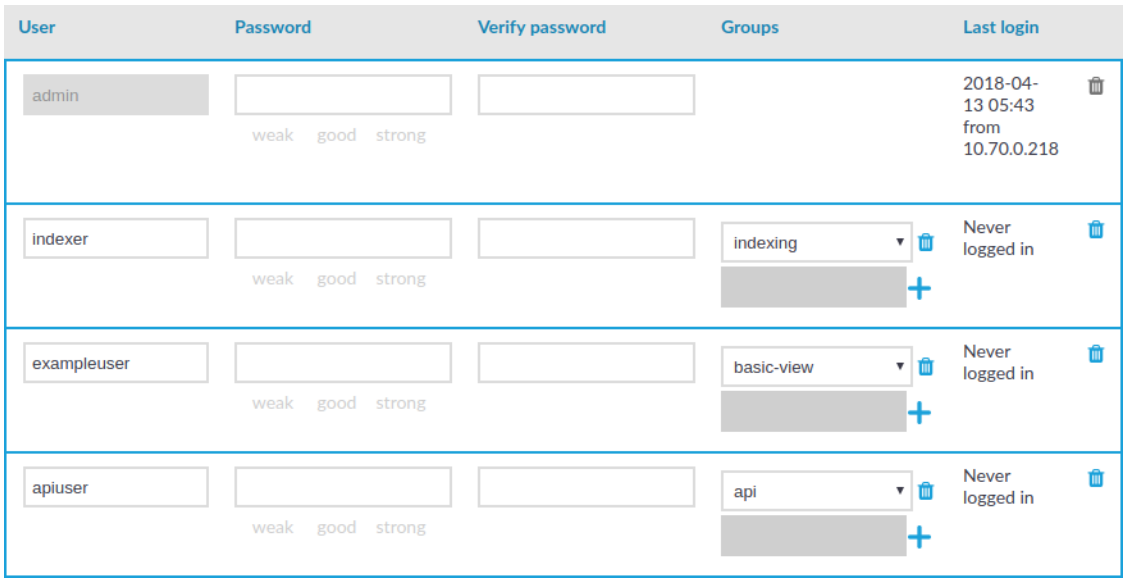
NOTE: The `admin` user is available by default, and has all possible privileges. It is not possible to delete this user.

When using RADIUS authentication together with local users, the users are authenticated to the RADIUS server, only their group memberships must be managed locally on One Identity Safeguard for Privileged Sessions (SPS). For details, see Authenticating users to a RADIUS server on page 347.

To create a local user

1. Navigate to **Users & Access Control** > **Local Users** and click  .

Figure 73: Users & Access Control > Local Users — Creating local users



2. Enter the username in the **User** field.

   NOTE: For the username of SSH users, only valid UTF-8 strings are allowed.

   The following characters cannot be used in usernames: `<>\/[]:;|=,+*?`

3. Enter a password for the user in the **Password** and **Verify password** fields.

   The strength of the password is indicated below the **Password** field, as you type. To set a policy for password strength, see Setting password policies for local users on page 330. The user can change the password later from the SPS web interface, and you can modify the password of the user here.

   Use strong passwords: at least 8 characters that include numbers, letters, special characters, and capital letters. For local One Identity Safeguard for Privileged Sessions (SPS) users, require the use of strong passwords (set **Users & Access Control** > **Login options** > **Minimal password strength** to `strong`). For more information, see *Setting password policies for local users* in the *Administration Guide*.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

NOTE: If possible, use a random password generator.

To create a strong password:

- Use special characters
- Use long passwords
- Mix uppercase and lowercase letters

For strong passwords, do not use:

- Personal information in the passwords
- Sequential letters or numbers
- The word `password` as the password
- Keyboard paths (for example, `qwerty`)

4. Click  in the **Groups** section and select a group that the user will be a member of. Repeat this step to add the user to multiple groups.

   You can also modify the group membership of local users here.

   To remove a user from a group, click  next to the group.
5. To save the modifications, click **Commit**.

Deleting local users from One Identity Safeguard for Privileged Sessions (SPS)

This section describes how to delete local users from One Identity Safeguard for Privileged Sessions (SPS).

To delete a local user from SPS

1. Navigate to **Users & Access Control** > **Local Users**.
2. Find the user you want to delete.
3. Click  next to the user, at the right edge of the screen.
4. To save your modifications, click **Commit**.

Setting password policies for local users

One Identity Safeguard for Privileged Sessions (SPS) can use password policies to enforce the use of:

- Password history
- Password strength

- Password length
- Password expiry
- Cracklib protection

Limitations

Consider the following limitations when configuring password policies:

- Password policies apply only to locally managed users.
- Password policies do not apply to users managed from an LDAP database, or authenticated to a RADIUS server.

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

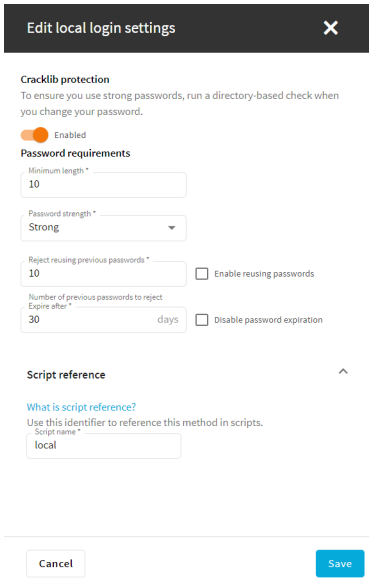
- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

To edit a password policy

1. Navigate to **Users & Access Control** > **Login options**.
2. Under **Login Options**, select the **Built-in Local login** method and click **Edit**.

   NOTE: Only the local users are affected by the password policy. The password rules of LDAP and RADIUS authentication are not managed by SPS.

Figure 74: Users & Access Control > Login options — Configuring password policies



3. (Optional) To do a basic dictionary check on passwords, enable **Cracklib protection**.
4. In the **Expire after** text box, configure the validity period of passwords (in days). Once the configured period expires, SPS users must change their passwords. The value range is `1-365`. You can also disable the password expiration.
5. **Reject reusing previous passwords**: use this option to prevent using the same password again for the configured number of password changes. For example, if the value is set to `10`, the users must use 10 different passwords consecutively, until the first password can be used again. The acceptable values are `1-30`. To disable this option, select **Enable reusing passwords**.
6. Set the required password complexity level in **Password strength**. The available values are `disabled`, `good`, and `strong`.

   NOTE: The strength of a password is determined by its length and complexity: the variety of numbers, letters, capital letters, and special characters used.

   To run simple dictionary-based attacks to find weak passwords, enable **Cracklib (eg. dictionary) protection**.
7. In **Password length**, set the minimum number of characters for the passwords. The acceptable values are `1-99`.
8. In the **Script reference** text box, specify a unique, human readable ID for referencing the configured settings in scripts (for example, to enable the REST API clients to select the login method).
9. To save your modifications, click **Commit**.

   NOTE: The changes you make in the password policy do not affect existing passwords. However, configuring password expiration will require every user to

change their password after the expiration date. Also, new passwords must comply with the current password strength settings.

Managing local user groups

You can use local groups to control the privileges of One Identity Safeguard for Privileged Sessions (SPS) local users — who can view and configure what.

For the description of built-in groups, see Built-in user groups of One Identity Safeguard for Privileged Sessions (SPS) on page 368.

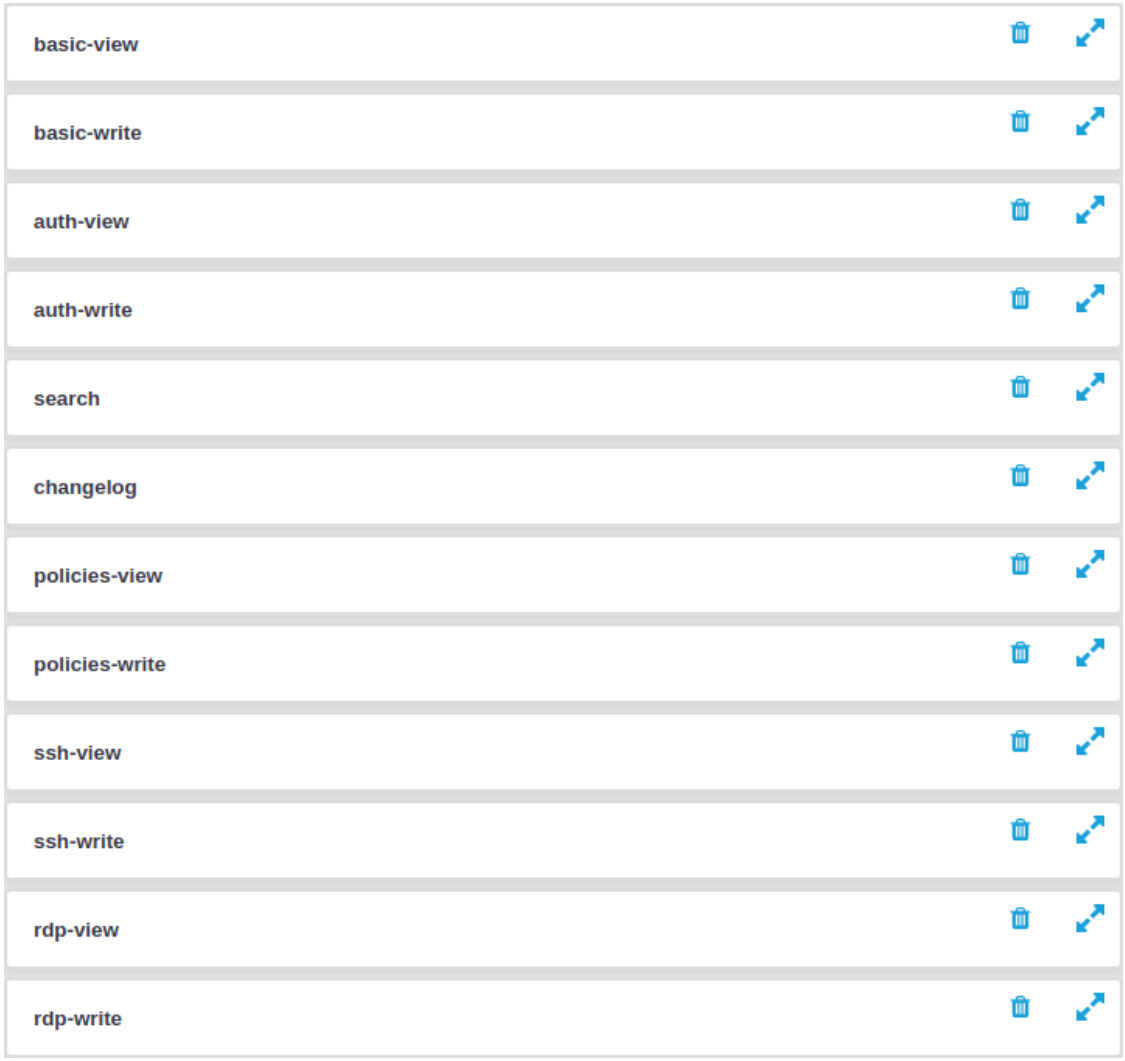
Use **Users & Access Control** > **Local User Groups** to:

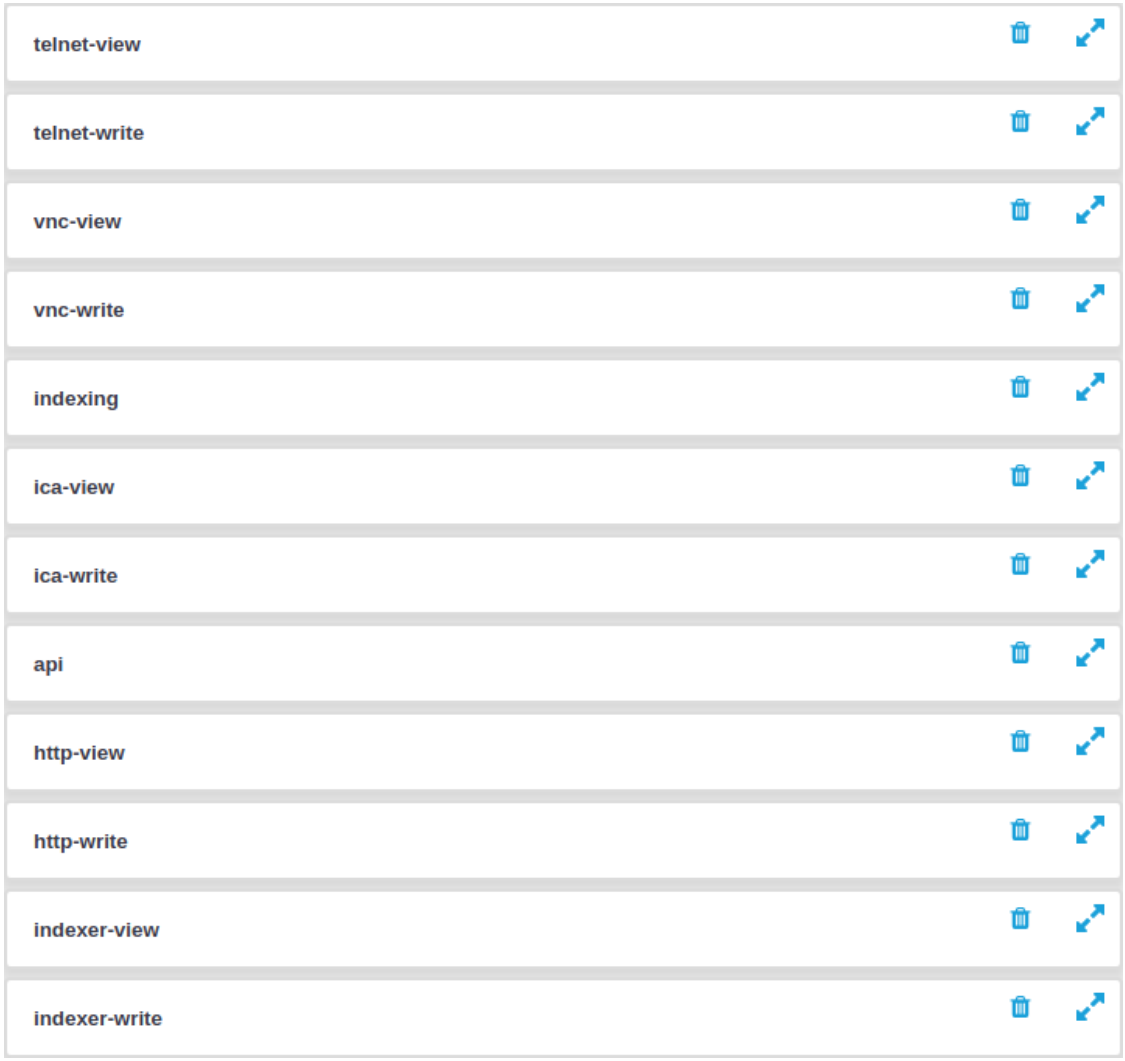
- Create a new user group.
- Display which users belong to a particular local user group.
- Edit group membership.

To create a new user group

1. Navigate to **Users & Access Control** > **Local User Groups** and click  .

Figure 75: Users & Access Control > Local User Groups — Group management





2. Enter a name for the group.
3. Enter the names of the users belonging to the group. Click + to add more users.
4. To save your modifications, click **Commit**.

   Once you have added your user groups, the next step is to start assigning privileges to them. For more information, see Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface on page 363.

Managing One Identity Safeguard for Privileged Sessions (SPS) users from an LDAP database

The One Identity Safeguard for Privileged Sessions (SPS) web interface can authenticate users to an external LDAP database to simplify the integration of SPS to your existing infrastructure. You can also specify multiple LDAP servers, so that if the first server becomes unavailable, SPS can try to connect to another server.

NOTE: Consider the following:

- Local users, including the `admin` user, are available by default.
- The `admin` user has every privilege and cannot be deleted.
- SPS accepts both pre-Windows 2000 style and Windows 2003 style User Principal Names (UPNs). UPNs consist of a username, the at (@) character, and a domain name, for example **`administrator@example.com`**.
- For SSH usernames, SPS supports only valid UTF-8 strings.
- The following characters cannot be used in:
  - user names: `/\[]:;|=+*?<>"`
  - group names: `/\[]:;|=+*?<>"@,`
- When using RADIUS authentication with LDAP users, the users are authenticated to the RADIUS server. However, their group memberships are managed in LDAP. For details, see *Authenticating users to a RADIUS server* in the *Administration Guide*.
- If the matching rule for an attribute is case insensitive in the LDAP database, SPS treats user names and group names in a case-insensitive manner.

Prerequisites

Make sure that the response timeout of the LDAP/Active Directory server is set to a minimum of 120 seconds.

To configure an LDAP server

1. Navigate to **Users & Access Control** > **Login options** > **Manage AD/LDAP Servers**.
2. Select the LDAP server from the list. Alternatively, if no LDAP server exists yet, click **Add new server** and select the server type you want to create:
   - **Active directory**

     For more information, see Adding a new Active Directory server.
   - **POSIX LDAP**

     For more information, see Adding a new POSIX LDAP server.

Adding a new Active Directory server

This section describes how to configure Active Directory (AD) servers.

Configuring AD servers requires the following steps:

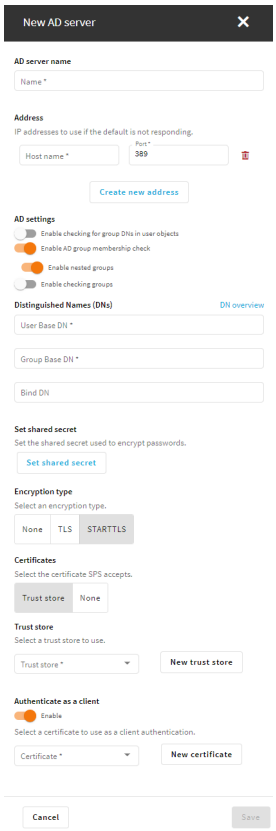
1. Creating a new AD server
2. Adding the new directory as a login method

Creating a new AD server

To create a new Microsoft Active Directory (AD) server

1. Navigate to **Users & Access Control** > **Login options** > **Manage AD/LDAP Servers**, click **Add new server**, and select **Active directory**.

   Figure 76: Users & Access Control > Login Options> Manage AD/LDAP Servers — Active Directory

2. In the **Name** field, enter the server name.
3. Enter the IP address/hostname and the port of the LDAP server in the respective text boxes.

   Consider the following when specifying the address information:

   - If you want to encrypt the communication between SPS and the LDAP server, use the following port numbers:

    - For **TLS**, specify 636 as the port number.
    - For **STARTTLS**, specify 389 as the port number.
  - Use an IPv4 adress or a hostname.
  - To add multiple servers, click  and enter the address of the next server. If a server is unreachable, SPS will try to connect to the next server in the list in failover mode.
  - When you configure the location of the LDAP server, that is, the IP address or hostname and the port number, you can use a Service record (SRV record), which is a type of information record in the DNS that maps the name of a service to the DNS name of the server. SRV records have the following format: `_ldap._tcp.<SITE_NAME>._sites.dc._msdcs.<DOMAIN.NAME>` in the **Address** field. SPS looks up the SRV record during committing the configuration change.

    For more information on SRV records, see the relevant Microsoft documentation.
  - **CAUTION:**

    If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.
4. Configure **AD settings**.
  - To also check group membership based on group Distinguished Names (DNs) in a user attribute, select **Enable checking for group DNs in user objects** and enter the name of the user attribute, for example, **`memberOf`** in the **User attribute of group DNs** field.

    CAUTION:

    If you have too many groups, using this option significantly slows down logging in to the SPS web interface.

    Use this option only if you have an LDAP schema where the user groups can only be determined from a user attribute that contains the group DNs.
  - To enable nested groups, select **Enable AD group membership check**, then **Enable nested groups**.

    CAUTION:

    Nested groups can slow down the query and cause the connection to timeout if the LDAP tree is very large. In this case, disable the **Enable nested groups** option.
  - To check for group membership based on user DNs in group attributes, use the **Check the user DN in these groups** option.

    For more information, see Active Directory LDAP backend.
5. Configure the options of the **Distinguished Names** field.

- In the **User Base DN** field, enter the name of the DN to be used as the base of queries regarding users (for example: `OU=People,DC=demodomain,DC=exampleinc`).

  NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

  To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

- In the **Group Base DN** field, enter the name of the DN to be used as the base of queries regarding groups (for example: `OU=Groups,DC=demodomain,DC=exampleinc`).

  NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

  To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

- In the **Bind DN** field, enter the Distinguished Name that SPS must use to bind to the LDAP directory (for example: `CN=Administrator,DC=demodomain,DC=exampleinc`).

  NOTE: SPS accepts both pre Windows 2000-style and Windows 2003-style account names, or User Principal Names (UPNs). For example, `administrator@example.com` is also accepted.

6. Configure the **Set shared secret** option.

   To configure or change the password to use when binding to the LDAP server, click **Set password**, enter the password, and click **Update**.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

7. If you want to encrypt the communication between SPS and the LDAP server, in **Encryption**, select the **TLS** or the **STARTTLS** option and verify the certificate of the server.

   - If you want SPS to verify the certificate of the server, under **Certificate requirements**, select **Trust Store**.
     - In the **Trust Store** field, select a trust store.

       SPS will use the selected trust store to verify the certificate of the server, and reject the connections if the verification fails.

     - To add a new trust store, click **New trust store**.

For more information, see Verifying certificates with Certificate Authorities using trust stores.

CAUTION:

SPS checks if the certificate revocation list (CRL) has expired and that the CRL has been signed by the same certificate authority (CA).

CAUTION:

If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.

- If the LDAP server requires mutual authentication, that is, it expects a certificate from SPS, enable **Authenticate as a client**. Generate and sign a certificate for SPS, upload the certificate and its private key, and click **Save**.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

8. To save your modifications, click **Commit**.

   NOTE: You must configure the usergroups in SPS, and possibly in your LDAP database. For details on using usergroups, see Using user groups.

Adding the new directory as a login method

After the directory has been configured and saved, you must add it as a login method.

To add the newly configured and saved directory as a login method

1. Navigate to **Users & Access Control** > **Login options** > **Create new login method**.
2. Select **Active Directory**.
3. Enter a name for the login method that will be displayed on the login page.
4. Select your previously created directory from the drop-down list under the **Authorized backed**.

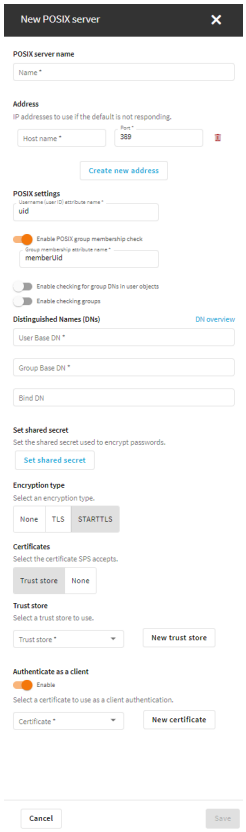
Adding a new POSIX LDAP server

This section describes how to configure POSIX LDAP servers.

To create a new POSIX LDAP server

1. To create a POSIX LDAP server, navigate to **Users & Access Control** > **Login options** > **Manage AD/LDAP Servers**, click **Add new server** and select **POSIX LDAP**.

Figure 77: Users & Access Control > Login Options> Manage AD/LDAP Servers — POSIX LDAP



2. In the **Name** field, enter the server name.
3. Enter the IP address/hostname and the port of the LDAP server in the respective text boxes.

   Consider the following when specifying the address information:

   - If you want to encrypt the communication between SPS and the LDAP server, use the following port numbers:
     - For **TLS**, specify 636 as the port number.
     - For **STARTTLS**, specify 389 as the port number.
   - Use an IPv4 adress or a hostname.
   - To add multiple servers, click + and enter the address of the next server. If a server is unreachable, SPS will try to connect to the next server in the list in failover mode.
   - When you configure the location of the LDAP server, that is, the IP address or hostname and the port number, you can use a Service record (SRV record), which is a type of information record in the DNS that maps the name of a service to the DNS name of the server. SRV records have the following format:

`_ldap._tcp.<SITE_NAME>._sites.dc._msdcs.<DOMAIN.NAME>` in the **Address** field. SPS looks up the SRV record during committing the configuration change.

For more information on SRV records, see the relevant Microsoft documentation.

- **CAUTION:**

  If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.

4. Configure **POSIX settings**.

   If your LDAP server uses a custom POSIX LDAP scheme, you might need to set which LDAP attributes store the username, or the attributes that set group memberships. For example, if your LDAP scheme does not use the `uid` attribute to store the usernames, set the **Username (user ID) attribute name** option.

   In addition to the POSIX LDAP backend membership checking, you can allow checking for supplementary group memberships by selecting **Enable POSIX group membership check** and specifying the **POSIX group membership attribute name** field.

   To also check group membership based on group Distinguished Names (DNs) in a user attribute, select **Enable checking for group DNs in user objects**. Then, enter the name of the user attribute (for example, `memberOf`) in the **User attribute of group DNs** field, and objectClass (for example, `groupOfNames`) in the **Group objectClass** field.

   CAUTION:

   If you have too many groups, using this option significantly slows down logging in to the SPS web interface.

   Use this option only if you have an LDAP schema where the user groups can only be determined from a user attribute that contains the group DNs.

   To check for group membership based on user DNs in group attributes, use the **Check the user DN in these groups** option.

5. Configure the options of the **Distinguished Names** field.

   - In the **User Base DN** field, enter the name of the DN to be used as the base of queries regarding users (for example: **`OU=People,DC=demodomain,DC=exampleinc`**).

     NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

     To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

   - In the **Group Base DN** field, enter the name of the DN to be used as the base of queries regarding groups (for example: **`OU=Groups,DC=demodomain,DC=exampleinc`**).

NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

- In the **Bind DN** field, enter the Distinguished Name that SPS must use to bind to the LDAP directory (for example: `CN=Administrator,DC=demodomain,DC=exampleinc`).

  NOTE: SPS accepts both pre Windows 2000-style and Windows 2003-style account names, or User Principal Names (UPNs). For example, `administrator@example.com` is also accepted.

6. If you want to encrypt the communication between SPS and the LDAP server, in **Encryption**, select the **TLS** or the **STARTTLS** option and verify the certificate of the server.
   - If you want SPS to verify the certificate of the server, under **Certificate requirements**, select **Trust Store**.
     - In the **Trust Store** field, select a trust store.

       SPS will use the selected trust store to verify the certificate of the server, and reject the connections if the verification fails.
     - To add a new trust store, click **New trust store**.

       For more information, see Verifying certificates with Certificate Authorities using trust stores.

     CAUTION:

     SPS checks if the certificate revocation list (CRL) has expired and that the CRL has been signed by the same certificate authority (CA).

     CAUTION:

     If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.
   - If the LDAP server requires mutual authentication, that is, it expects a certificate from SPS, enable **Authenticate as a client**. Generate and sign a certificate for SPS, upload the certificate and its private key, and click **Save**.

   TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
7. To save your modifications, click **Commit**.

   NOTE: You must configure the usergroups in SPS, and possibly in your LDAP database. For details on using usergroups, see Using user groups.

Handling user names in User Principal Name (UPN) format

This section and its subsections describe how you can handle user names in User Principal Name (UPN) format in One Identity Safeguard for Privileged Sessions (SPS).

- About the UPN format
- Multiple identifiers of a single user entity while using UPN user names
- Case-sensitivity while using UPN user names
- Differences in protocols while using UPN user names
- Using UPN user names in usermapping policies
- Using UPN user names in groups (user lists)
- UPN user names and AA plugins

About the UPN format

The UPN format is a user name in an email-like format, which consists of three parts:

- a local identifier
- a separator (the `@` character)
- a domain identifier (for example, `user@domain.com`).

UPN user names are common in many applications and protocols. In Active Directory, the domain identifier selects the Active Directory domain that the user belongs to.

> NOTE: Even though a UPN user name may look like an email address, it is not always a valid email address. Similarly, a person's valid email address is not necessarily a valid UPN user name, either.

Multiple identifiers of a single user entity while using UPN user names

In SPS, user identifiers are used in several places. For example, users can specify their user name as part of their login credentials, administrators can create groups of users for authorization purposes, and auditors can search for the activities of a certain person.

In Active Directory, several user name forms may refer to the same person. For example, these three can all belong to a single Active Directory user object:

- `EXAMPLE\exampleuser` (down-level user name)
- `exampleuser@example.local`
- `ExampleUser@example.com` (UPN with an alternative suffix).

When it comes to Windows authentication, the user is free to choose from any of the user name forms above. However, SPS components are only aware of the specified form of the user name.

Case-sensitivity while using UPN user names

User identifiers are case-insensitive in Windows and Active Directory, and in most POSIX LDAP schemas as well. On the other hand, other operating systems, and SPSin particular, handle user identifiers in a case-sensitive manner.

Consequently, the `username` strings, which are used for authentication at a remote Windows or AD/LDAP, are case-insensitive. For example, one can specify **`Administrator`**, **`administrator`**, or even **`aDMinIstRaTor`** during login, they are not different from the authentication point of view.

For convenience reasons, search also works in a case-insensitive manner.

However, for the following functions, SPS requires a case exact match:

- Credential stores
- Remote authentication (SAML2, Radius, X.509) with a local backend
- Usermapping policies
- Local user lists (remote or gateway groups)

For example, if there is a local credential store entry for **`administrator`** at a given target server, the user name in the connection must be exactly **`administrator`**, even though authentication would succeed with other forms of capitalization as well. Remote credential stores (such as One Identity Safeguard for Privileged Passwords), on the other hand, may accept user names in a case-insensitive manner.

For more information about how case-sensitivity and case matching applies to using the UPN format in usermapping policies, see Using UPN user names in usermapping policies.

Differences in protocols while using UPN user names

Unlike other protocols, RDP allows the explicit specification of the Windows domain during authentication. When the user specifies the target `username` string as a down-level user name in NetBIOS format, SPS splits the string into an actual `username` (that is, the part after the `\` character) and a `domain`. In this case, only the actual user name is used for usermapping policies, credential stores, AA plugins and user lists.

As an example, when the `username` string is specified as **DOMAINNAME\example**, the usermapping policy requires a record for the actual user name **`example`**, and not for the whole `username` string.

For protocols other than RDP, SPS does not split the `username` string to `username` and `domain`.

> NOTE: Historically, SPS versions before 8.0 LTS also split UPN user names to a local `username` and a `domain` for the RDP protocol. However, in SPS version 8.0 LTS and later versions, this operation is no longer performed.

For specific information about how differences in protocols affect AA plugins and the UPN format, see UPN user names and AA plugins.

Using UPN user names in usermapping policies

In SPS, you can use the User Principal Name (UPN) format in usermapping policies that let administrators control who can use specific user names at target servers.

For more information about usermapping policies, see Configuring usermapping policies.

When the **Ignore UPN suffix in user names** option is selected on the SPS web interface (it is selected by default), only the local identifier part (that is, the characters before the `@` character) will be used to match the remote user name. This means that if you specify `administrator` as the remote user name, it will match for `administrator@example.domain` and `administrator@other.domain` as well. The session is authorized on the target server side, depending on the target server's configuration.

If you want to be more specific in allowing which remote users can be used, you can disable the **Ignore UPN suffix in user names** option. In this case, you have to specify the exact user name in the remote user name. For example, if you have `example.domain` and `other.domain` specified as alternate UPN suffixes, and additionally, you want to enable the use of user names in NetBIOS format, you have to add the following three entries to the usermapping policy:

- `administrator@example.domain`
- `administrator@other.domain`
- `administrator`

Using UPN user names in groups (user lists)

You can use the UPN format in groups (that is, user lists) in SPS.

For more information about user lists, see Creating and editing user lists.

In SPS, local user lists do not allow specifying the members in UPN format because they can be used as exclusion lists (also known as deny lists), and the exclusion could be easily circumvented by using alternate user name forms. Therefore, when a user name is checked against a local user list, the UPN suffix is always ignored.

User groups managed in AD / LDAP can be used with UPN user names as well.

UPN user names and AA plugins

On the SPS web interface, the **ignore_upn_suffix_for_rdp** option is used for AA plugins. This option controls how AA plugins receive the `username` string from SPS, and it is set to `true` by default.

When the **ignore_upn_suffix_for_rdp** option is set and the RDP protocol is in use, UPN user names are split into a local `username` part and a `domain` part. Otherwise, the `domain` will be empty and the `username` string is unchanged.

When the AA plugin is invoked for an RDP session, `username` strings in the NetBIOS format (that is, in the `DOMAIN\username` format) are always split into a `username` and a `domain` part before being passed to the plugin.

For any other protocols, the `username` strings are never altered, and the `domain` part is never set.

Authenticating users to a RADIUS server

One Identity Safeguard for Privileged Sessions (SPS) can authenticate its users to an external RADIUS server. Group memberships of the users must be managed either locally on SPS or in an LDAP database.

CAUTION:

The challenge/response authentication method is currently not supported. Use other authentication methods (for example password, SecureID).

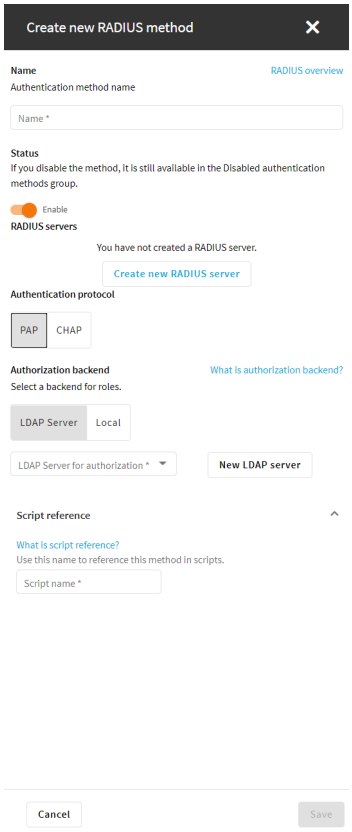
Authenticating SPS users to a RADIUS server

To authenticate SPS users to a RADIUS server, complete the following steps.

1. Navigate to **Users & Access Control** > **Login Options**.
2. To configure a RADIUS login method, select one of the following options:
   - Select an existing RADIUS login option and click **Edit**.
   - Click **Create new authentication method** and select **RADIUS**.

   The following figure shows the configuration options of the RADIUS login method.

Figure 78: Users & Access Control > Login options — Configuring RADIUS authentication



3. In the **Name** field, specify a name for the login option.
4. (Optional) Enable the RADIUS login method.
5. To add a new RADIUS server, click **Create new RADIUS server**.
    a. In the **Address** field, enter the IP address or domain name of the RADIUS server. Use an IPv4 address or hostname.
    b. In the **Server port** field, enter the port number.
    c. In the **Shared secret** field, enter the password that SPS can use to access the RADIUS server.

        NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

        - Letters A-Z, a-z
        - Numbers 0-9
        - The space character
        - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`
    d. Click **Save**.

6. (Optional) To add more RADIUS servers, click + and repeat the procedure for adding a new RADIUS server.

   If a server is unreachable, SPS tries to connect to the next server in the list in failover mode.

7. Select the authentication protocol.
   - To use the Password Authentication Protocol, select **PAP**.
   - To use the Challenge-Handshake Authentication Protocol, select **CHAP**.
8. Select **LDAP server** or **Local** as the **Authorization Backend**.
9. (Optional) To add a new LDAP server, click **New LDAP server** under **Authorization backend** and select one of the server types:
   - **Active Directory**

     For more information, see Adding a new Active Directory server.
   - **POSIX**

     For more information, see Adding a new POSIX LDAP server.
10. **Script reference** is filled out automatically when you specify the name for the login option. Special characters are automatically replaced with dashes ("-"). The **Script name** is a unique, human readable ID that is used by the REST API clients to select the login method.
11. To save your modifications, click **Commit**.

    CAUTION:

    After you commit this configuration, the SPS web interface will be available only after successfully authenticating to the RADIUS server. Note that the default `admin` account of SPS will be able to login normally, even if the RADIUS server is unaccessible.

Authenticating users with X.509 certificates

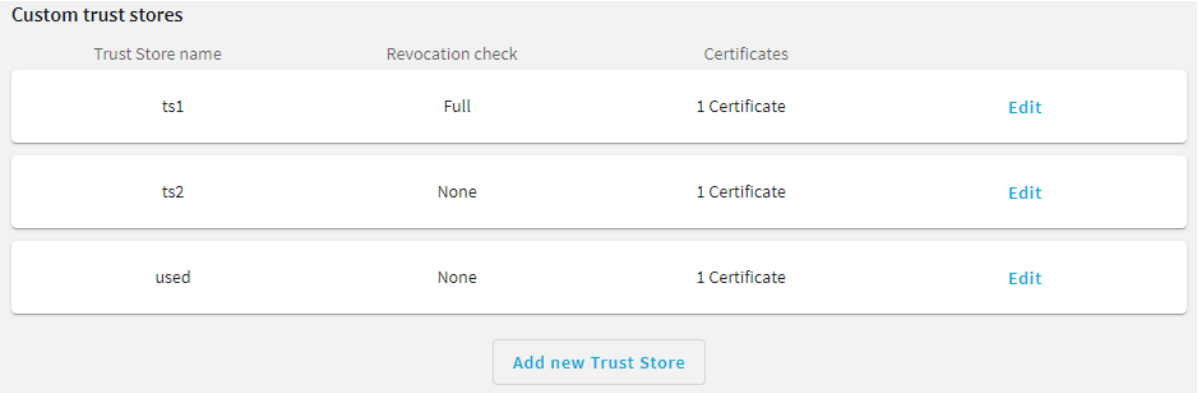
One Identity Safeguard for Privileged Sessions (SPS) provides a method to authenticate the users of the web interface with X.509 client certificates. The client certificate is validated against a trust store, and the username is exported from the client certificate for identification.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

Prerequisites

Navigate to **Basic Settings** > **Trust Stores** and create a custom trust store. For more information, see Verifying certificates with Certificate Authorities using trust stores.

Figure 79: Basic Settings > Trust Stores— Creating a custom trust store



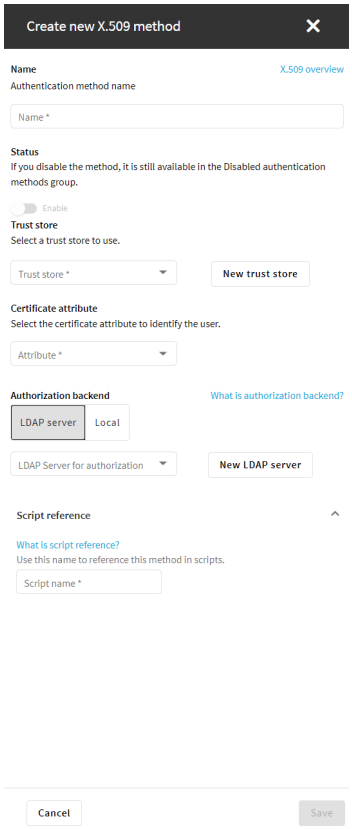
Authenticating SPS users on the SPS interface with X.509 certificates

To authenticate SPS users on the SPS web interface with X.509 client certificates, complete the following steps.

1. Navigate to **Users & Access Control** > **Login options**.
2. To configure an X.509 login method, select one of the following options:
   - Select an existing X.509 login option and click **Edit**.
   - Click **Create new authentication method** and select **X.509**.

   The following figure shows the configuration options of the X.509 login method.

Figure 80: Users & Access Control > Login options — Configuring X.509 authentication



3. In the **Name** field, specify a name for the login option.
4. (Optional) Enable the X.509 login method.

   NOTE: You can enable only one X.509 login method at a time.
5. Select the trust store you created as a prerequisite in **Basic Settings** > **Trust Stores**.

CAUTION: From version 6.8, SPS changes authenticating the users of the web interface with X.509 client certificates: certificates are validated against a trust store instead of a trusted CA list. During the upgrade, the trusted CA list formerly used for authentication is copied to a trust store that has revocation check disabled by default.

If you have previously enabled revocation check for your trusted CA list and already added the URLs of Certificate Revocation Lists (CRL), or you want to enable revocation check, you must edit the trust store settings manually.

- Navigate to **Basic Settings** > **Trust Stores**.
- Select the revocation check type **Leaf** or **Full** for the trust store.
- Add a **CRL URL** for each root and intermediate CA.

For more information about trust stores and how to configure them, see Verifying certificates with Certificate Authorities using trust stores.

6. In the **Certificate attribute** field, select the user certificate field that contains the username. In most cases, it is the `commonName` or `userid` field, but SPS supports the `emailAddress` and `userPrincipalName` fields as well.
7. Select **LDAP server** or **Local** as the **Authorization Backend**.
8. (Optional) To add a new LDAP server, click **New LDAP server** under **Authorization backend** and select one of the server types:
   - **Active Directory**

     For more information, see Adding a new Active Directory server.
   - **POSIX**

     For more information, see Adding a new POSIX LDAP server.
9. **Script reference** is filled out automatically when you specify the name for the login option. Special characters are automatically replaced with dashes ("-"). The **Script name** is a unique, human readable ID that is used by the REST API clients to select the login method.
10. To save your modifications, click **Commit**.

Authenticating users with SAML2

The topics below discuss how to authenticate users with SAML2.

SAML2 login overview

During SAML2 login, the Service Provider (SP) makes the authorization decision about a subject (user) based on an assertion, which is created by an Identity Provider (IdP). To

make this decision, the SP must trust the IdP and the IdP must provide enough information about the user to make this authorization decision.

When you configure SAML2 login in One Identity Safeguard for Privileged Sessions, SPS serves as the SP. The SAML2 authentication flow consists of several HTTP redirects, where the information exchange between the SP and the IdP is performed using the user's browser. This means that there is no direct network communication between SPS and the IdP.

The process of the SAML2 authentication flow is the following:

1. The user goes to the SPS login page and clicks on the login button associated with your SAML2 login method, see Authenticating users with SAML2 login method for the details.
2. SPS redirects the user to the login page of the IdP. The redirect contains the SAML2 authentication request.
3. The IdP authenticates the user. This step is entirely within the domain of the IdP. It may require a password and/or a second factor, or, if the user has an active session at a different application, the authentication may be performed without any user interaction, providing a single sign-on experience.
4. The IdP redirects the user back to SPS. This redirect conveys the SAML2 response, which contains information about the user.
5. SPS processes the response and based on the information within, it maps the user either to a local or to an AD/LDAP user entry. If the IdP provides the user's groups in the response, then SPS evaluates the permissions assigned to the groups given by the IdP. Otherwise, the permissions are evaluated based on the groups in which the mapped local or AD/LDAP user entry is a member.

SAML2 support in SPS

SAML 2.0 is a complex standard, and it requires that both the Identity Provider (IdP) and the Service Provider (SP) are configured in a way to interoperate correctly. This section is provided to help you integrate SPS with your IdP.

Identity Provider metadata

To authenticate users securely, SPS needs to know many technical details about the Identity Provider (IdP). The standard way of representing this information is SAML metadata, which is an XML file. You must obtain this file from your IdP and upload it to SPS.

The XML file must contain a single IdP entity. If you want to allow logins to SPS from multiple IdPs, you must create additional login methods with different metadata files, see Authenticating users with SAML2 login method. Optionally, the IdP entity element can be wrapped into an EntitiesDescriptor element.

Service Provider metadata

SPS provides Service Provider (SP) metadata at the following location:

```
https://<ADDRESS-OF-YOUR-SPS>/sts/saml2/sp-metadata.xml
```

This file is accessible also for unauthenticated users since it only contains public information about the SAML2 SP configuration.

NOTE: Many Identity Providers (IdPs) do not consume SP metadata directly, therefore you might need to configure your IdP manually. This guide uses the terms that are defined in the SAML 2.0 standard, which may differ from how the actual configuration parameters are called by your IdP implementation. For example, depending on your IdP implementation, the SP Entity ID may be referred to as **Audience** or **Audience URI**.

Entity ID

The entity ID of SPS is an opaque string, which is generated automatically, for example:

```
https://example.com/sts/saml2/1278910176319e703186a8
```

This is a technical identifier of your SP, which should be unique in your federation, but is not meant to be visible by users. When needed, you can customize this string using the SPS REST API, see Configuring SPS login methods in the *One Identity Safeguard for Privileged Sessions REST API Reference Guide*.

Signing and encryption keys

SPS SP has an automatically generated private key and a corresponding self-signed certificate, which are used for the following purposes:

- sign the SAML2 authentication request sent to the IdP,
- decrypt the assertion in the IdP's response when it is encrypted (encryption is not required),
- sign the metadata file, to prove the possession of the private key.

When needed, you can change the private key and the corresponding certificate using the SPS REST API.

Assertion Consumer Service URLs

Assertion Consumer Service (ACS) URL is a parameter in the SAML2 authentication request, to which the IdP should redirect the user back after completing the authentication. To prevent request forging, most IdPs validate that the ACS specified in the authentication request match any of the preconfigured ACS values for the SP.

Since SPS is accessible at multiple addresses, there can be several ACS URLs. The URL is a string made up from three parts:

- the fixed string `https://`
- the configured host name or the IP address of SPS with an optional port number. If the port number is specified, it must match the port number used in the URL bar of the browser.
- the fixed string `sts/saml2/acs/post`

All ACS URLs support the HTTP POST binding.

NOTE: When you configure your SP, you must include all host names in the list of ACS URLs where SAML2 login is allowed. If the users attempt to log in to SPS using an alternative host name or IP address in the URL bar, the login will fail. The host name might contain a port number, but in this case, you must enter the port number into the URL bar of the browser too.

Since a managed SPS cluster appears as a single SP entity for IdPs, you must enter all host names (or IP addresses) of all managed SPS hosts into the central configuration. For example, if a host is represented by three different hostnames, such as **`10.12.231.241`**, **`example.com`**, or **`user.example.com:8081`**, you must enter all three hostnames pertaining to that host to make SAML2 login method available for users. You must also enter the hostnames of the central configuration.

User identifiers

In SAML, the IdP can provide the user's long-term identifier in two ways:

- using a NameID, or
- using an attribute.

If the NameID format of the assertion is any of the following, then the user identifier will be obtained from the NameID value:

- `urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress`
- `urn:oasis:names:tc:SAML:2.0:nameid-format:persistent`
- `urn:oasis:names:tc:SAML:1.1:nameid-format:X509SubjectName`
- `urn:oasis:names:tc:SAML:1.1:nameid-format:WindowsDomainQualifiedName`

If the assertion has a different NameID format, then one of the following user attributes should contain the user identifier value:

- **`subjectId`**
- **`pairwiseId`**
- **`eduPersonUniqueId`**
- **`eduPersonPrincipalName`**
- **`upn`**
- **`primarySid`**
- **`uid`**

- **eduPersonOrcid**
- **email**
- **emailAdress**
- **mail**

The attribute names above can be expressed in various standard attribute name formats, such as `urn:oasis:names:tc:SAML:2.0:attrname-format:basic` or `urn:oasis:names:tc:SAML:2.0:attrname-format:uri`.

Group membership

If the IdP provides the user's groups in attribute values, then SPS evaluates the permissions assigned to these groups, therefore user authorization is performed based on the assertion only. SPS supports the following attributes for groups:

- **eduPersonEntitlement**
- **isMemberOf**
- **group**

How to configure SAML2 login

The topics below discuss how to configure SAML2 login.

Overview

Configuring SAML2 login contains the following steps:

- Configure local or AD/LDAP users and assign permissions to groups, see Managing One Identity Safeguard for Privileged Sessions (SPS) users locally and Managing One Identity Safeguard for Privileged Sessions (SPS) users from an LDAP database for details.
- Configure SPS as a SAML2 Service Provider (SP).
- Configure your Identity Provider (IdP) to trust SPS.
- Configure a SAML2 login method.

Configure SPS as a SAML2 SP

NOTE: Authentication configuration is shared between the SPS central configuration and the managed hosts, therefore you must configure the Service Provider (SP) settings and the SAML2 login methods on the central configuration node.

Navigate to **Users & Access Control** > **Login options**. The **SSO Configuration** menu contains information about your SP configuration.

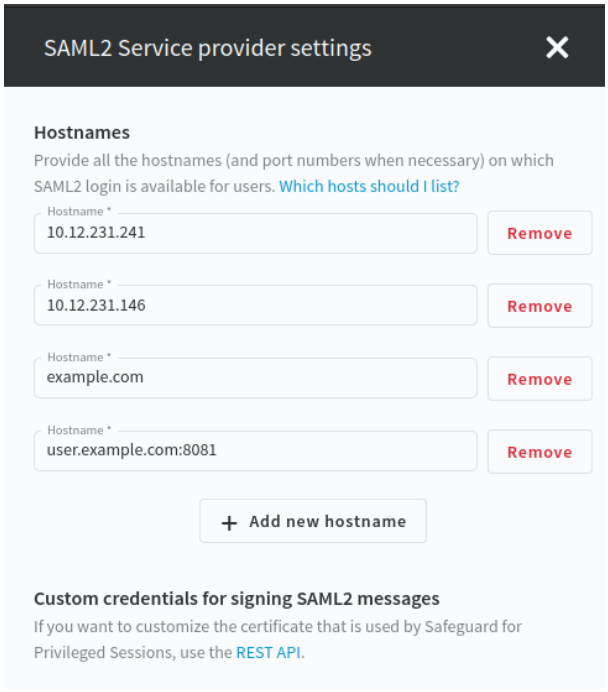
To configure SPS as a SAML2 Service Provider (SP), complete the following steps.

1. Navigate to **Users & Access Control** > **Login options**, and click **SSO Configuration**.

   The **SSO Configuration** menu contains information about your SP configuration.
2. Click **SAML2 Service Provider settings** from the drop down menu.
3. Click **Add new hostname** to provide all the host names on which the SAML2 login method is available for users. When necessary, provide the port number as well. For example, `10.12.231.241`, `example.com`, or `user.example.com:8081`.

   The web user interface allows you to configure the Assertion Consumer Service URLs. You can configure the entityID and the custom credentials on the REST API, if the defaults are not suitable.

   Figure 81: Users & Access Control > Login Options > SSO Configuration > SAML2 Service Provider settings – Configuring SP entity ID and host names

4. Click **Save**.

Configure your IdP to trust SPS

To configure your Identity Provider (IdP) to trust One Identity Safeguard for Privileged Sessions (SPS) as a Service Provider (SP), you must provide the SP metadata XML file of

your SPS to your IdP. If your IdP supports the import of SP metadata, then you can choose either of the following methods to download the SP metadata XML file:

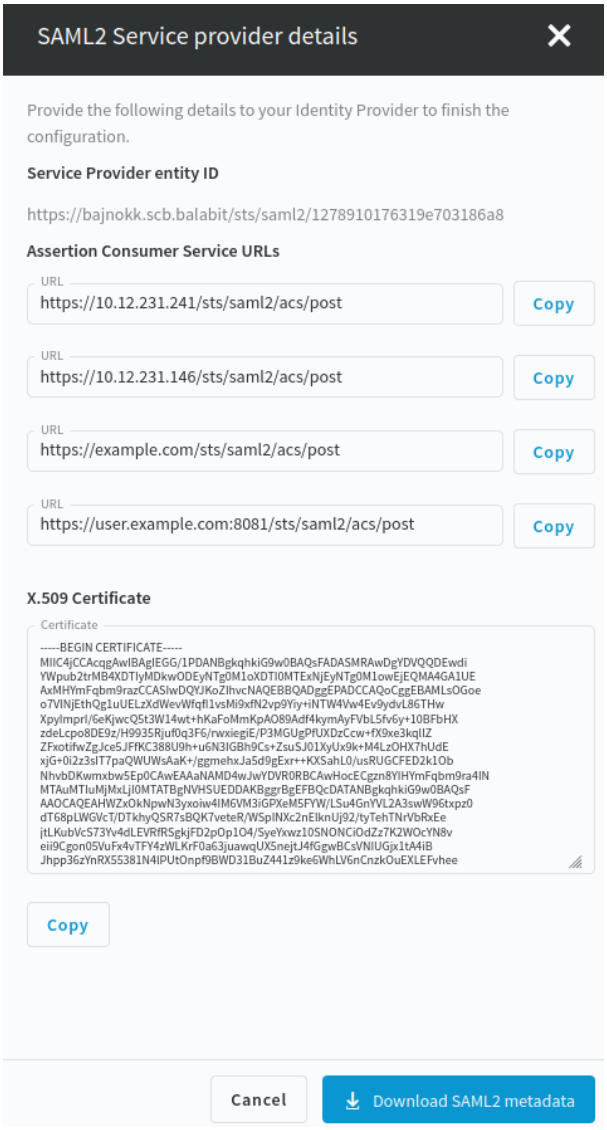
- Download the SP metadata XML file by clicking **Users & Access Control** > **Login Options** > **SSO Configuration** > **SAML2 Service Provider details** > **Download SAML2 metadata** on the SPS web interface.
- Download the SP metadata XML file from your SPS at the following location:

  ```
  https://<ADDRESS-OF-YOUR-SPS>/sts/saml2/sp-metadata.xml
  ```

  This file is accessible also for unauthenticated users since it only contains public information about the SAML2 SP configuration.

If your IdP does not support the import of SP metadata, then you must configure your IdP based on the summary of the **SAML2 Service Provider details** page.

Figure 82: Users & Access Control > Login Options > SSO Configuration > SAML2 Service Provider details



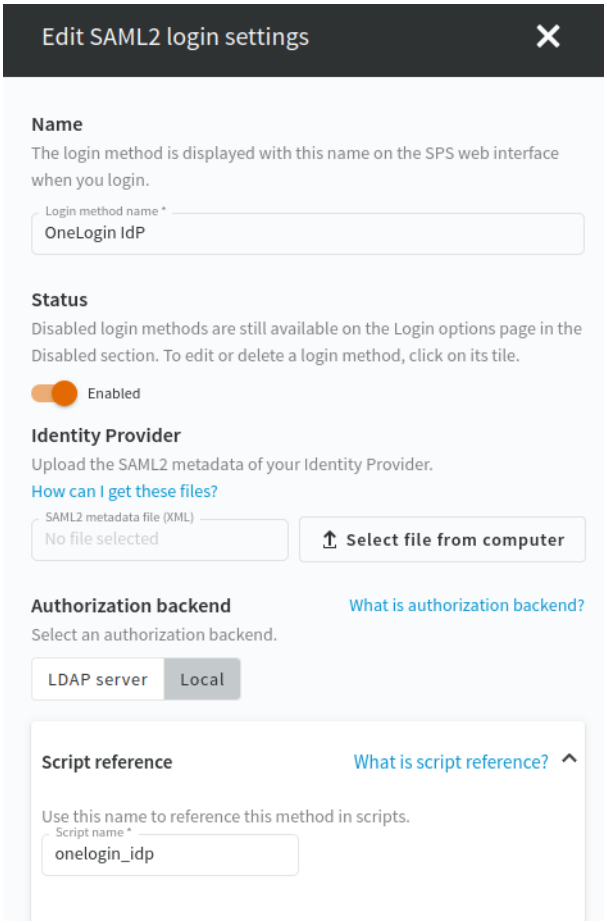
Authenticating users with SAML2 login method

To authenticate SPS users with a SAML2 login method, complete the following steps.

1. Navigate to **Users & Access Control** > **Login Options**.
2. To configure a SAML2 login method, select one of the following options:
    - Select an existing SAML2 login option and click **Edit**.
    - Click **Create new login method** and select **SAML2**.

The following figure shows the configuration options of the SAML2 login method.

Figure 83: Users & Access Control > Login options — Configuring SAML2 authentication



3. In the **Name** field, specify a name for the login option.
4. Upload the SAML2 metadata XML of your **Identity Provider**.

   NOTE: Your IdP metadata file should contain a single SAML2 IdP metadata entity. To support SAML2 login with multiple IdPs, you have to configure additional SAML2 login methods.
5. Select **LDAP server** or **Local** as the **Authorization backend**.
6. (Optional) To add a new LDAP server, click **Add new LDAP server** under **Authorization backend** and select one of the server types:
   - **Active Directory**

     For more information, see Adding a new Active Directory server.
   - **POSIX**

     For more information, see Adding a new POSIX LDAP server.

7. **Script reference** is filled out automatically when you specify the name for the login option. Special characters are automatically replaced with dashes ("-"). The **Script name** is a unique, human readable ID that is used by the REST API clients to select the login method.
8. To save your modifications, click **Commit**.

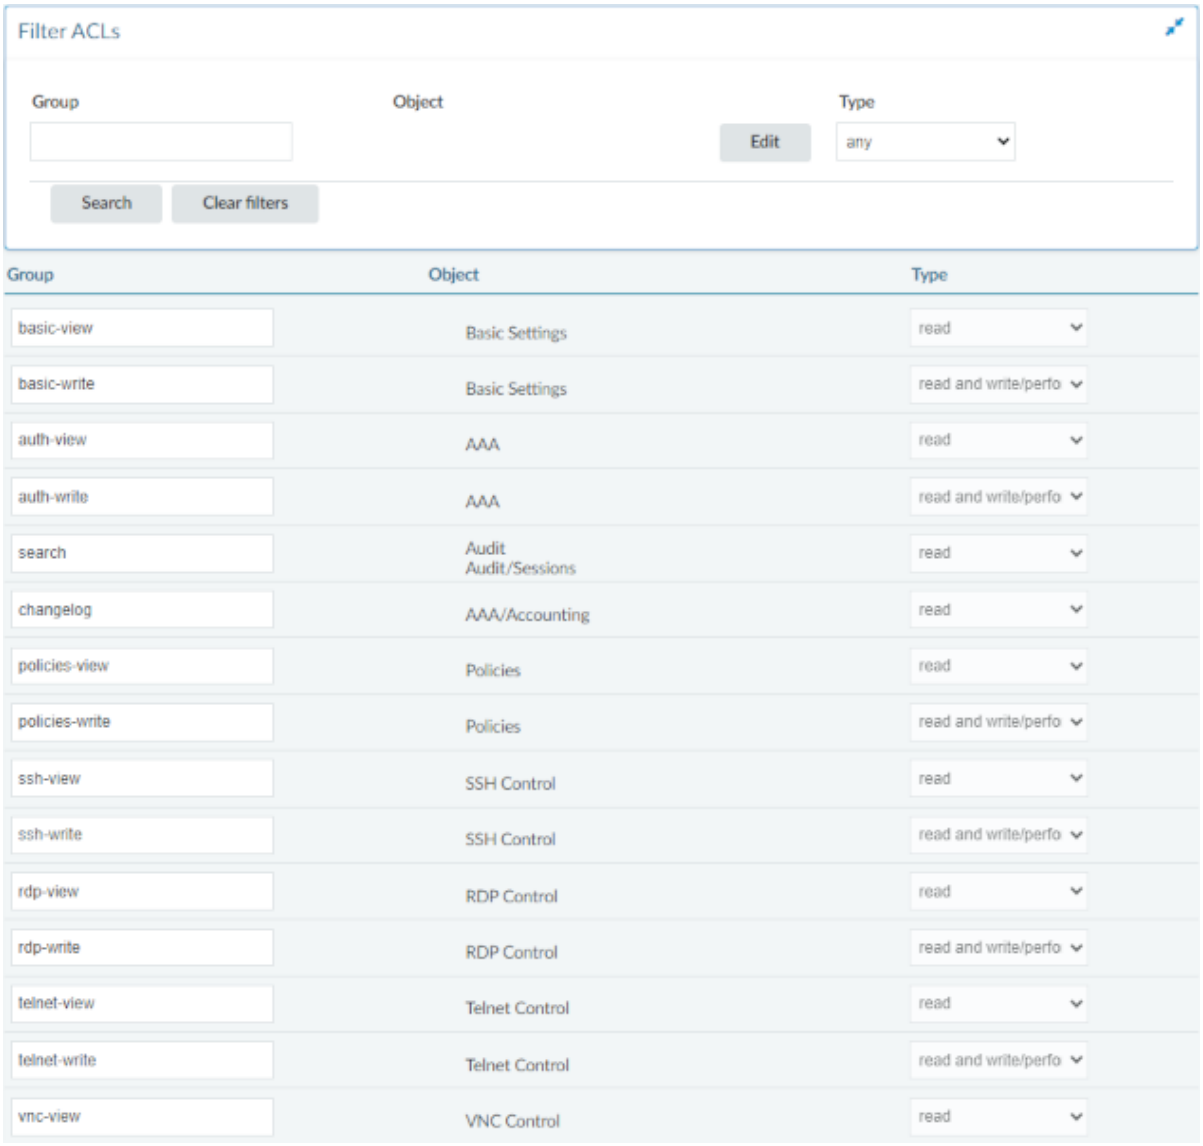
Managing user rights and user groups

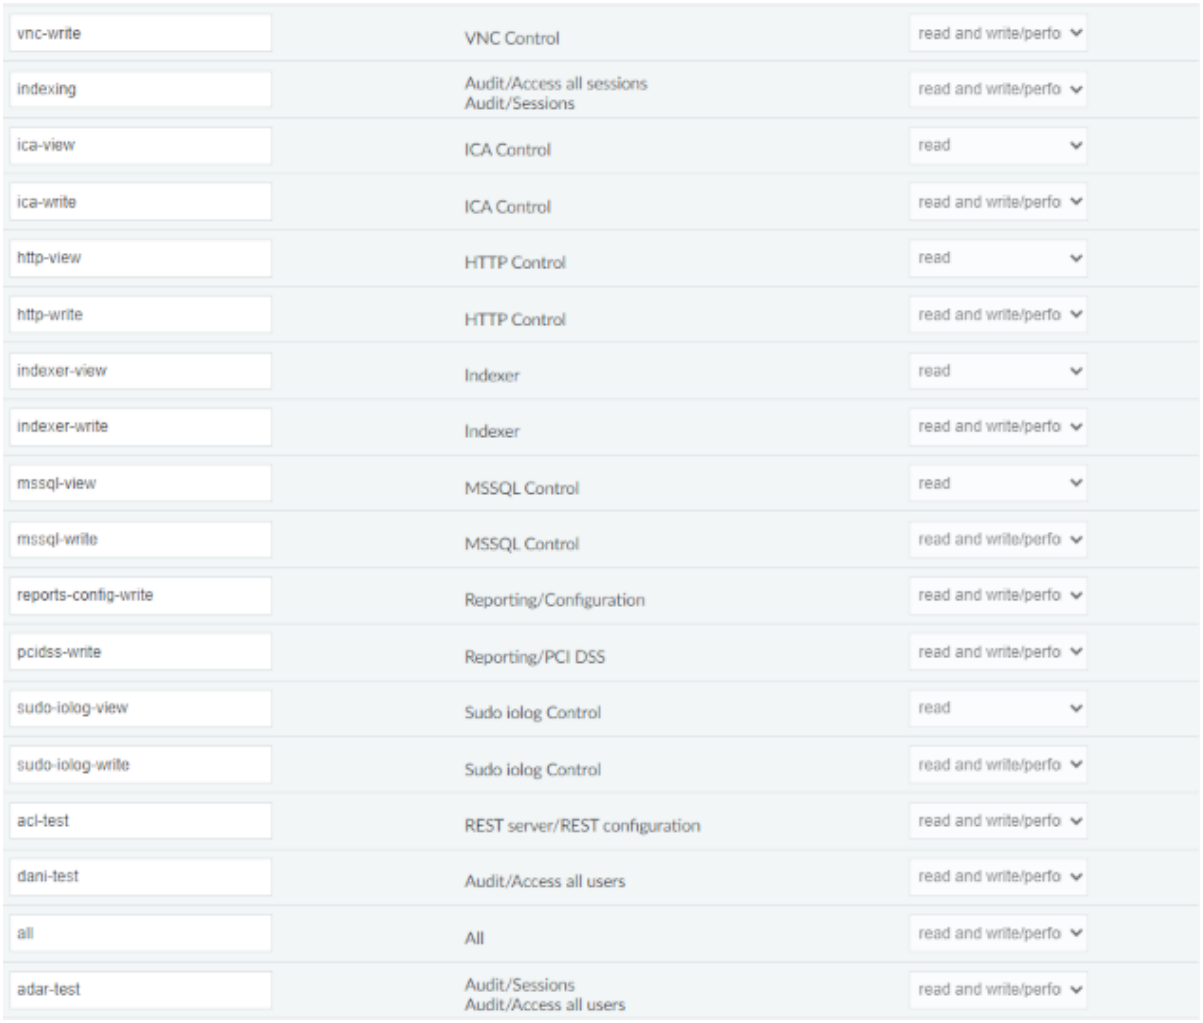
In One Identity Safeguard for Privileged Sessions (SPS), you can assign user rights to user groups. SPS has several user groups defined by default, but you can also define additional custom user groups. Every group has a set of privileges that affects which pages of the SPS web interface it can access, and whether it can only view (read) or also modify (read & configure) those pages or perform specific actions.

NOTE: Every group has either read or read & write/perform privileges to a set of pages.

- For more information about assigning privileges to a new user group, see Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface.
- For more information about modifying existing groups, see Modifying group privileges.
- For more information about finding user groups that have a specific privilege, see Finding specific user groups.
- For more information about using user groups, see Using user groups.
- For more information about the privileges of the built-in user groups, see Built-in user groups of One Identity Safeguard for Privileged Sessions (SPS).
- For more information about the access and summary list of Access Control Lists (ACLs) in SPS, see the ACL access and summary list.

Figure 84: Users & Access Control > Appliance Access — Managing SPS users





Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface

This section describes how you can assign privileges to a new group.

To assign privileges to a new group

1. Navigate to **Users & Access Control** > **Appliance Access** and click +.
2. Find your user group. If you start typing the name of the group you are looking for, the auto-complete function will make finding your group easier for you.
3. Click **Edit** located next to the name of the group. The list of available privileges is displayed.

4. Select the privileges (that is, the pages of the One Identity Safeguard for Privileged Sessions (SPS) interface) to which the group will have access and click **Save**.

   NOTE: To export the configuration of SPS, the **Export configuration** privilege is required.

   To import a configuration to SPS, the **Import configuration** privilege is required.

   To update the firmware and set the active firmware, the **Firmware** privilege is required.

5. Select the type of access (**read** or **read & write**) from the **Type** field.
6. Click 

.

Modifying group privileges

This section describes how you can modify the privileges of an existing group in One Identity Safeguard for Privileged Sessions (SPS).

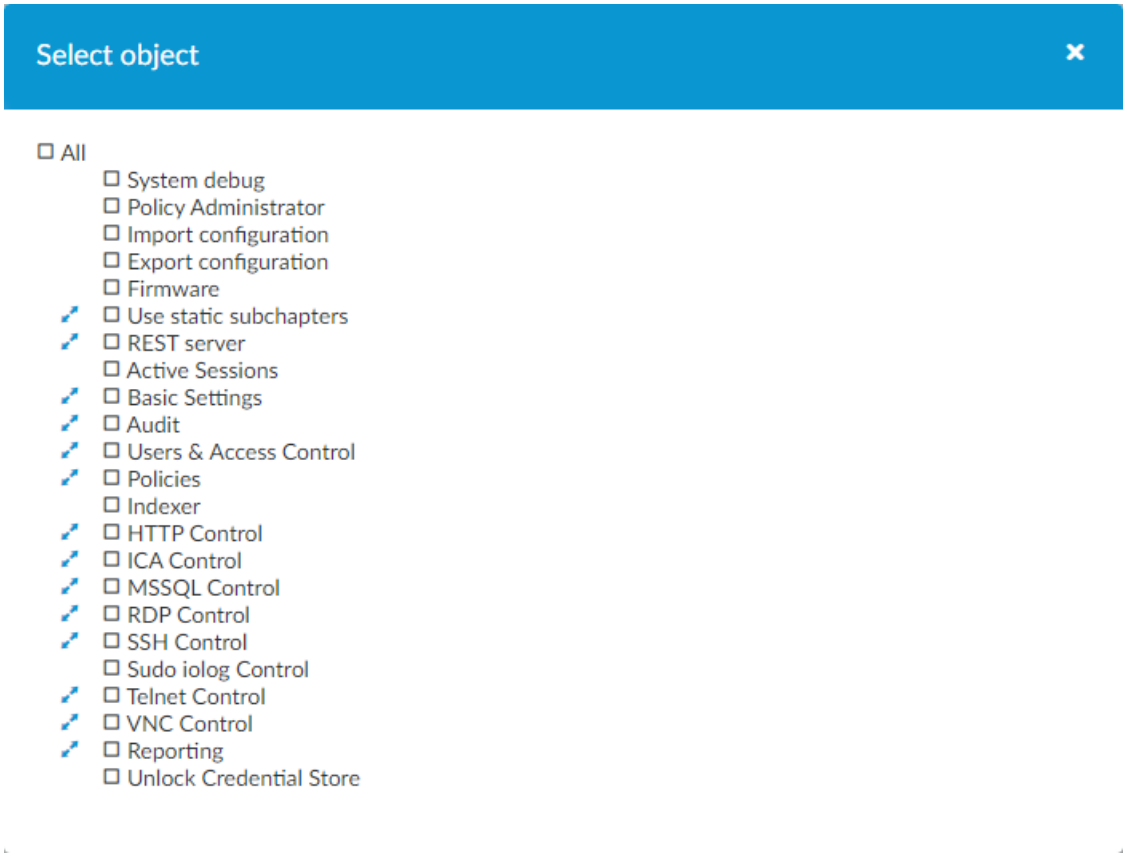
To modify the privileges of an existing group

1. Navigate to **Users & Access Control** > **Appliance Access**.
2. Find the group you want to modify and click **Edit**.

   The list of available privileges appears.

   For more information about the available privileges known as Access Control Lists (ACLs), see the ACL access and summary list.
3. Select the privileges (that is, the pages of the SPS interface) to which the group will have access and click **Save**.

Figure 85: Users & Access Control > Appliance Access > Edit — Modifying group privileges



CAUTION:

Assigning the **Audit** privilege to a user on the Users & Access Control page automatically enables the **Access all sessions** privilege, and grants the user access to every audit trail, even if the user is not a member of the groups listed in the **Access Control** option of the particular connection policy.

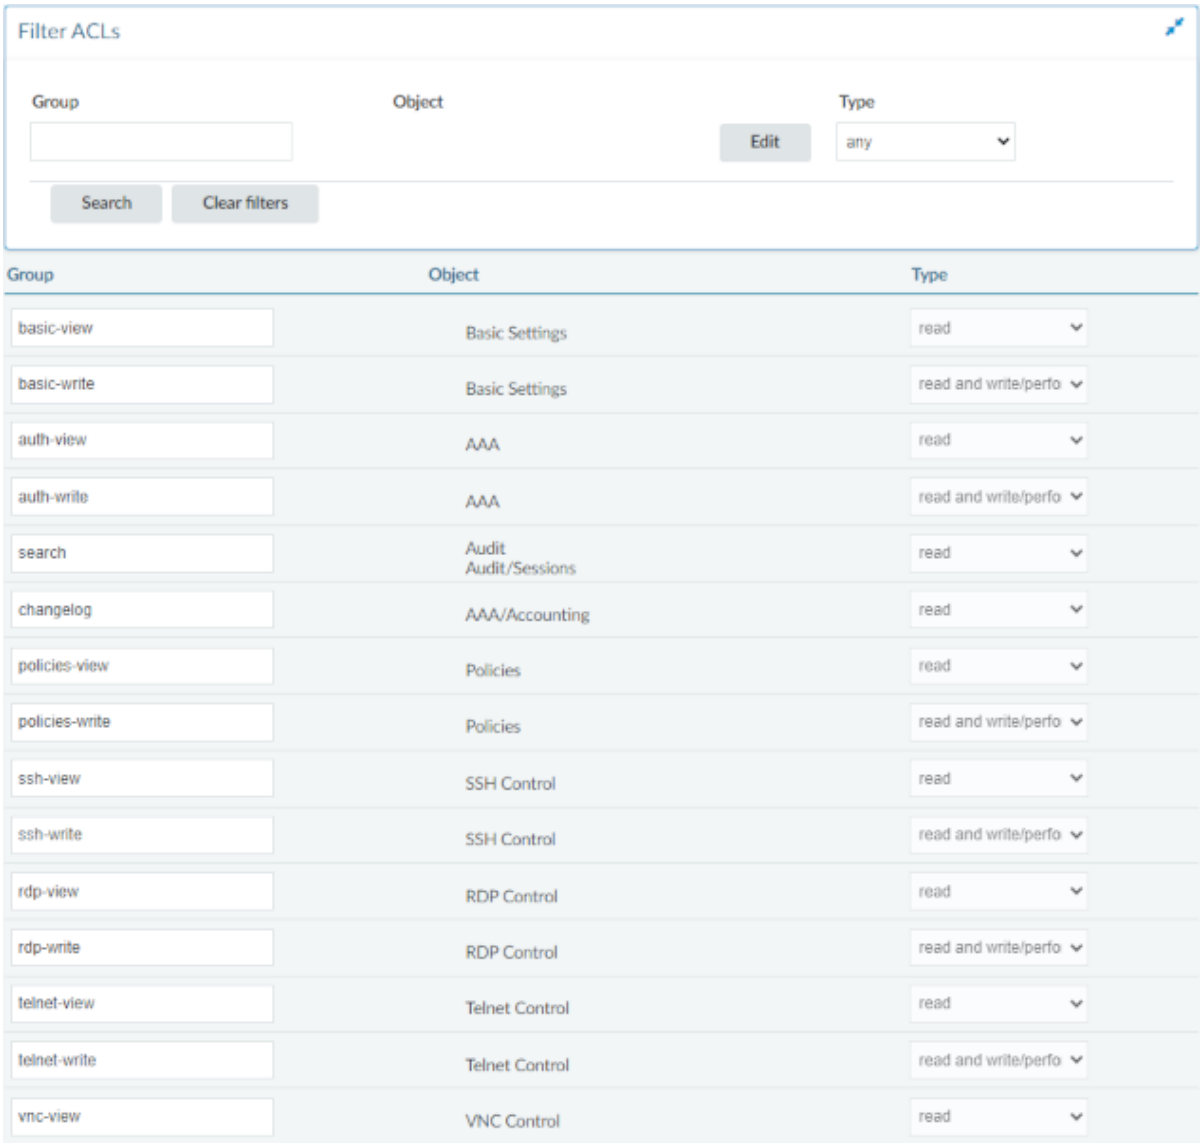
NOTE: The `admin` user is available by default and has all privileges. This user cannot be deleted.

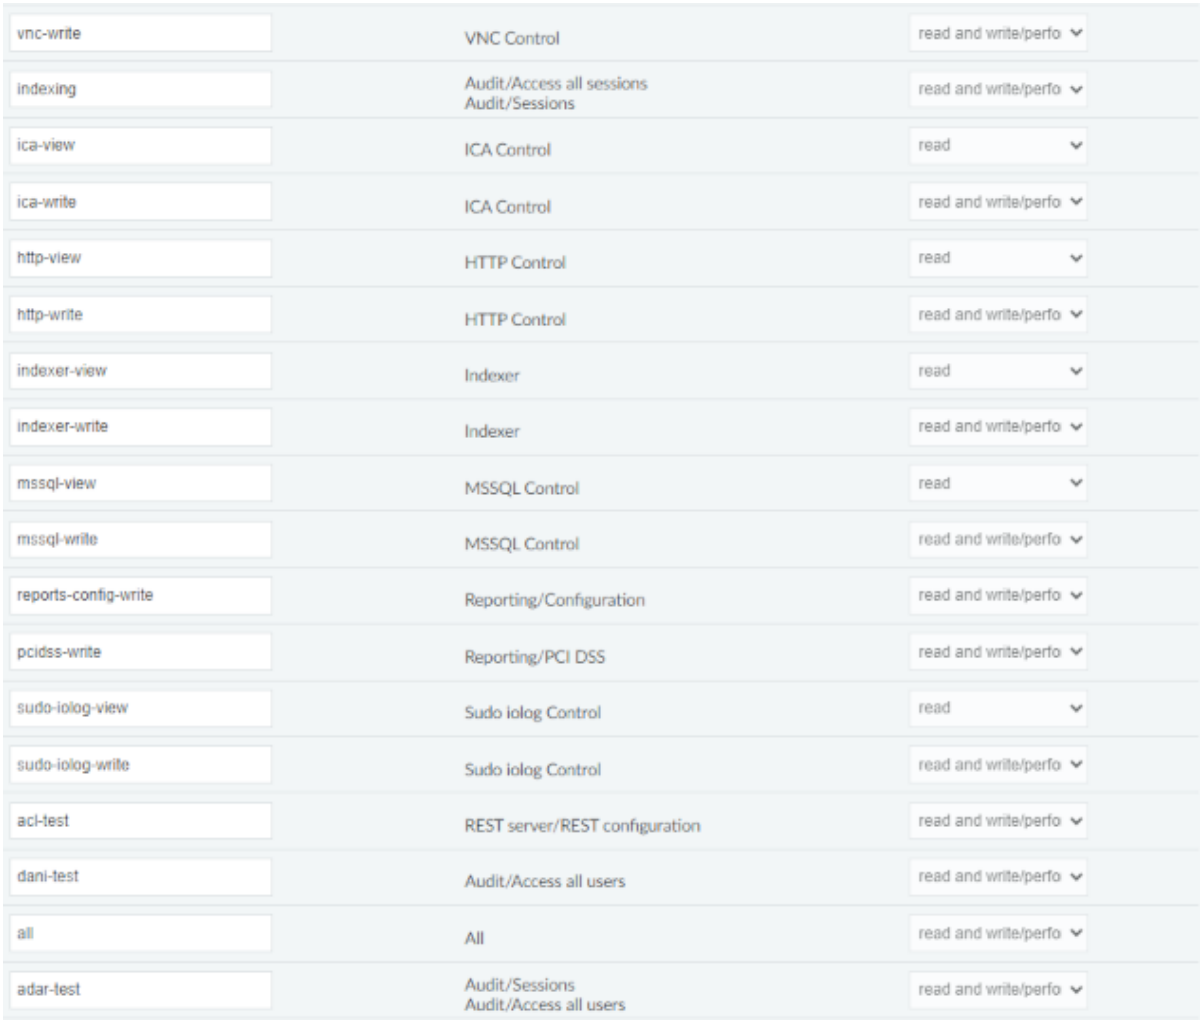
4. From the **Type** field, select the type of access (**read** or **read & write**) you want to assign to the user.
5. Click Commit.

Finding specific user groups

To search the names and privileges of user groups, use the search and filter interface of the **Filter ACLs** section of the **Users & Access Control** > **Appliance Access** page.

Figure 86: Users & Access Control > Appliance Access — Finding specific user groups





- To filter on a specific user group, enter the name of the group into the **Group** field and select **Search**.
- To select user groups who have a specific privilege, click **Edit**, select the privilege(s) you are looking for, and click **Search**.
- To filter for read or write access, use the **Type** option.

Using user groups

How you should name user groups depends on the way you manage your One Identity Safeguard for Privileged Sessions (SPS) users.

- *Local users*: If you use only local users, create or modify user groups on the **Users & Access Control** > **Local User Groups** page, assign or modify privileges on the **Users & Access Control** > **Appliance Access** page, and add users to the groups on the **Users & Access Control** > **Local Users** or the **Users & Access Control** >

Local User Groups page.

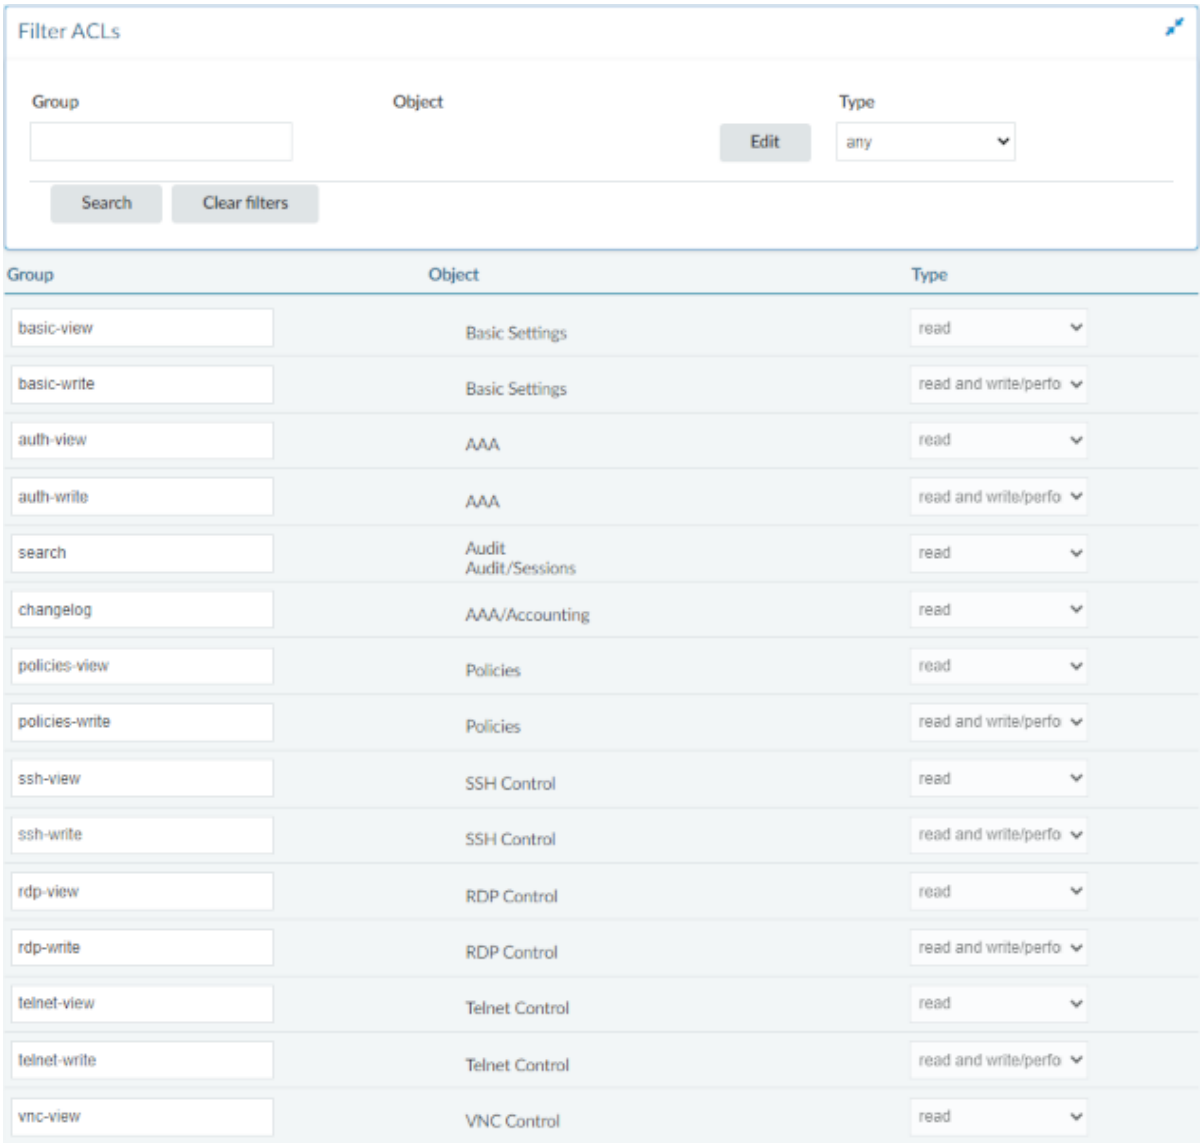
- *LDAP users and LDAP groups*: If you manage your users from LDAP and also have LDAP groups that match the way you want to group your SPS users, create or modify your user groups on the **Users & Access Control** > **Appliance Access** page and ensure that the name of your LDAP group and the SPS user group is the same. For example, to make members of the `admins` LDAP group be able to use SPS, create a user group called `admins` on the **Users & Access Control** > **Appliance Access** page and edit the privileges of the group as you need.
- *RADIUS users and local groups*: This is the case when you manage users from RADIUS, but you cannot or do not want to create groups in LDAP. Create your local groups on the **Users & Access Control** > **Appliance Access** page, and add your RADIUS users to these groups on the **Users & Access Control** > **Local User Groups** page.

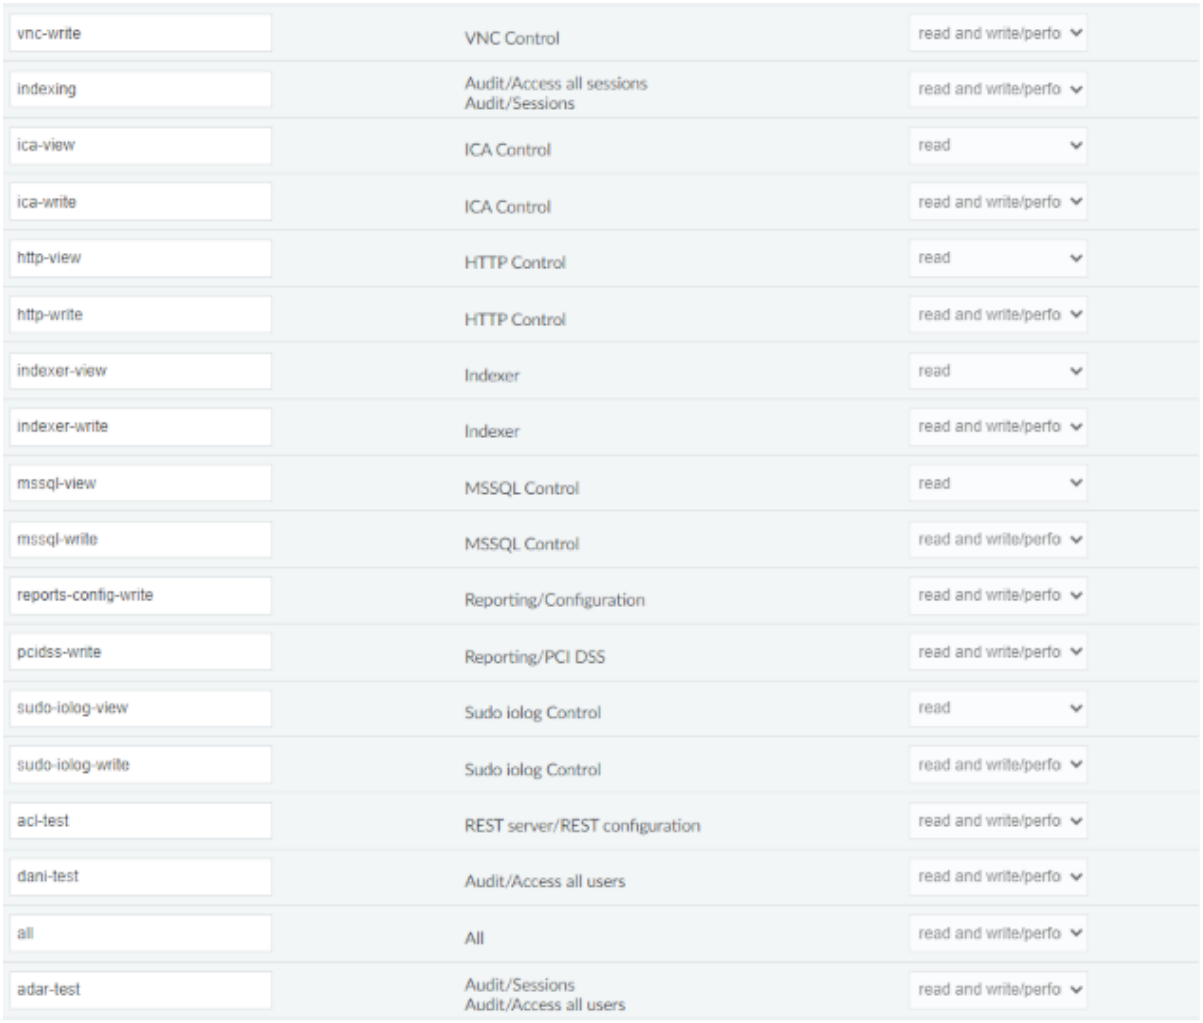
Built-in user groups of One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) has the following user groups built in by default.

NOTE: You can modify and delete these user groups as you see fit.

Figure 87: Users & Access Control > Appliance Access — Built-in user groups of SPS





CAUTION:

If you use LDAP authentication on the SPS web interface and want to use the default user groups, you have to create these groups in your LDAP database and assign users to them. For more information on using user groups, see Using user groups.

- **basic-view**: View the settings in the **Basic Settings** menu, including the system logs of SPS. Members of this group can also run commands on the **Troubleshooting** tab.
- **basic-write**: Edit the settings in the **Basic Settings** menu. Members of this group can manage SPS as a host.
- **auth-view**: View the names and privileges of the SPS administrators, the configured user groups, and the authentication settings in the **Users & Access Control** menu. Members of this group can also view the history of configuration changes.
- **auth-write**: Edit authentication settings and manage users and user groups.

CAUTION:

Members of the `auth-write` group, or any other group with write privileges to the **Users & Access Control** menu are essentially equivalent to SPS system administrators, as they can assign any privilege to themselves. Users with limited rights must never have such privileges.

If users with write privileges to the **Users & Access Control** menu assign new privileges to themselves (for example, they assign themselves group membership to a new group), then they can apply the new privilege by logging in to the SPS web interface again.

- **search**: Browse and download various logs and alerts in the **Sessions** menu. The members of this group have access to the audit trail files as well. Note that to open encrypted audit trail files, the proper encryption keys are required.
- **changelog**: View the history of SPS configuration changes in the **Users & Access Control** > **Configuration History** menu.
- **report**: Browse, create and manage reports, and add statistics-based chapters to the reports in the **Reports** menu. Users with this privilege can access every report. To give access to users only to specific reports, use the **Reports are accessible by the following groups** option of the report. For more information, see Configuring custom reports on page 893.

  NOTE: To control exactly which statistics-based chapters and reports can the user include in a report, use the `Use static subchapters` privileges.
- **policies-view**: View the policies and settings in the **Policies** menu.
- **policies-write**: Edit the policies and settings in the **Policies** menu.
- **ssh-view**: View all connection and policy settings in the **Traffic Controls** > **SSH** menu.
- **ssh-write**: Edit all connection and policy settings in the **Traffic Controls** > **SSH** menu.
- **rdp-view**: View all connection and policy settings in the **Traffic Controls** > **RDP** menu.
- **rdp-write**: Edit all connection and policy settings in the **Traffic Controls** > **RDP** menu.
- **telnet-view**: View all connection and policy settings in the **Traffic Controls** > **Telnet** menu.
- **telnet-write**: Edit all connection and policy settings in the **Traffic Controls** > **Telnet** menu.
- **vnc-view**: View all connection and policy settings in the **Traffic Controls** > **VNC** menu.
- **vnc-write**: Edit all connection and policy settings in the **Traffic Controls** > **VNC** menu.

- **indexing**: Allows hosts running external indexers to access and download audit trails for automatic indexing.

  NOTE: The members of this group can access the SPS web interface as well, and download any audit trail directly.
- **ica-view**: View all connection and policy settings in the **Traffic Controls** > **ICA** menu.
- **ica-write**: Edit all connection and policy settings in the **Traffic Controls** > **ICA** menu.
- **http-view**: View all connection and policy settings in the **Traffic Controls** > **HTTP** menu.
- **http-write**: Edit all connection and policy settings in the **Traffic Controls** > **HTTP** menu.
- **indexer-view**: View all connection and policy settings in the **Indexer** menu.
- **indexer-write**: Edit all connection and policy settings in the **Indexer** menu.

ACL access and summary list

This section contains the access and summary list of Access Control Lists (ACLs) available in One Identity Safeguard for Privileged Sessions (SPS).

For more information about managing user rights and user groups in SPS in general, see Managing user rights and user groups.

NOTE:

Starting from One Identity Safeguard for Privileged Sessions (SPS) 8.0 LTS, AA plugins are no longer required to connect SPS and One Identity Safeguard Remote Access.

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
All	Every menu point	Gives access to every setting.
System Debug	**Basic Settings** > **Troubleshooting** > **Create support bundle**	Enables support bundle generation.
Policy Administrator	**Traffic Control** > **Quick Connection Setup**	Enables configuring SSH connections and their connection-related policies.
	Traffic Control >	

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
	SSH	
	Policies > **Analytics policies**	
	Policies > **Audit policies**	
	Policies > **Content policies**	
	Policies > **Local User Databases**	
	Policies > **Time policies**	
	Policies > **User Lists**	
	Policies > **Usermapping policies**	
Import configuration	**Basic Settings** > **System** > **Import configuration**	Enables configuration import.
Export configuration	**Basic Settings** > **System** > **Export configuration**	Enables configuration export.
Firmware	**Basic Settings** > **System** > **Firmwares**	Enables uploading, installing and deleting firmwares on SPS.
Use static subchapters ACLs	Does not affect menu points.	Restricts users to add only specific subchapter types to reports.
REST server / **REST configuration**	**Traffic Controls** > **Quick Connection Setup**	Enables configuring connections via the Connection

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
		Setup Wizard and the REST API.
REST Server / **REST**	**Users & Access Control** > **Audit Data Access**	Enables configuring audit data access rules (ADARs) for user groups.
Active sessions	**Pending Connections** > **Active connections**	Gives access to view and terminate all active sessions monitored by SPS.
Basic Settings / **Network**	**Basic Settings** > **Network**	Enables configuring network-related settings.
Basic Settings / **Local services**	**Basic Settings** > **Local services**	Enables configuring various services of SPS.
Basic Settings / **Management**	**Basic Settings** > **Management**	Gives access to configuring system-level settings.
	Basic Settings > **Cluster Management**	Enables configuring multiple SPS appliances joining into cluster for scalability, easier monitoring and auditing. Additionally, it enables joining SPS to One Identity Safeguard for

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
		Privileged Passwords (SPP).
	Basic Settings > **Trust stores**	Enables configuring local certificate storage for storing trusted certificates that validate TLS connections. Additionally, it authorizes the user to manage cryptography settings.
Basic settings / **Alerting Monitoring**	**Basic Settings** > **Alerting and monitoring**	Gives access to alerting and monitoring configuration.
Basic settings / **Date and time**	**Basic Settings** > **Date & Time**	Gives access to date and time configuration.
Basic settings / **Dashboard**	**Basic Settings** > **Dashboard**	Gives access to various statistics and status history of system data and performance on the dashboard.
Basic settings / **System**	**Basic Settings** > **System**	Gives access to system-related configuration.
Basic settings / **Troubleshooting**	**Basic settings** > **Troubleshooting**	Gives access to tools that help troubleshooting network issues, viewing log files and generating support bundles.

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
Basic settings / **Plugins**	**Basic settings** > **Plugins**	Gives access to plugin management.
Basic settings / **High availability**	**Basic settings** > **High availability**	Gives access to high availability setup.
Basic settings / **Cluster management**	**Basic settings** > **Cluster management**	Gives access to cluster management.
Basic settings / **Join to Starling**	**Basic settings** > **Starling integration**	Gives access to the **Starling Integration** configuration and status page.
Audit / **Sessions**	**Audit** > **Sessions**	Enables querying sessions. However, without an audit data access rule (ADAR), no session data is visible. Gives access to the **Audit** > **Sessions** page.
Audit / **Access all sessions**	Does not affect menu points.	Authorizes to search and view all sessions, but does not give access to the **Audit** > **Sessions** page. In other words, it works like an unfiltered ADAR.
Audit / **Access all users**	**Audit** > **Users**	Enables viewing data about audited users.

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
		However, it does not give access to view user-related session data. To view user-related session data, you either need the combination of **Audit** / **Sessions ACL** with ADAR, or the combination of the **Audit** / **Sessions** ACL and the **Audit** / **Access all sessions** ACL.
Users & Access Control / **Settings**	**Users & Access Control** > **Settings**	Enables configuring the login settings of SPS.
Users & Access Control / **Local Users and Groups**	**Users & Access Control** > **Local Users and Groups**	Enables configuring local users and local user groups for accessing SPS.
Users & Access Control / **Access Control**	**Users & Access Control** > **Appliance Access**	Enables configuring access control for user groups.
Users & Access Control / **Accounting**	**Users & Access Control** > **Configuration History**	Enables access to view the configuration history of SPS.
Users & Access Control / **Permission Query**	**Users & Access Control** > **Access Rights Report**	Enables viewing access rights reports.
Policies / **AA Plugin Configurations**	**Policies** > **AA**	Enables

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
	Plugin Configurations	configuring authentication and authorization plugins to authenticate to target hosts.
Policies / **Analytics Policies**	**Policies** > **Analytics Policies**	Enables configuring analytics policies that specify algorithms to analyze monitored user behavior and build user baselines.
Policies / **Audit Policies**	**Policies** > **Audit Policies**	Enables configuring audit policies specifying encryption, signature and timestamping of the recorded audit trail files.
Policies / **Backup & Archive**	**Policies** > **Backup & Archive**	Enables configuring policies to backup and archive audit data.
Policies / **Audit Data Cleanup Policies**	**Policies** > **Audit Data Cleanup Policies**	Enables configuring policies to cleanup audit data.
Policies / **Content Policies**	**Policies** > **Content Policies**	Enables configuring policies to take action when a

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
		predefined content appears in a session.
Policies / **Credential Stores**	**Policies** > **Credential Stores**	Enables configuring user credential stores that store credentials for authenticating to target servers.
Policies / **Indexer Policies**	**Policies** > **Indexer Policies**	Enables configuring policies specifying what data to capture and index from sessions.
Policies / **LDAP Servers**	**Policies** > **LDAP Servers**	Enables configuring LDAP servers to query users and user groups for authentication purposes.
Policies / **Local User Databases**	**Policies** > **Local User Databases**	Enables configuring databases storing user credentials for authenticating clients on SPS in gateway authentication scenarios.
Policies / **Signing CAs**	**Policies** > **Signing CAs**	Enables configuring certificate authorities (CAs) for signing certificates.

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
Policies / **Time Policies**	**Policies** > **Time Policies**	Enables configuring time intervals to use in other policy settings.
Policies / **Trusted CA Lists**	**Policies** > **Trusted CA Lists**	Enables configuring trusted certificate authorities for validating certificates.
Policies / **User Lists**	**Policies** > **User Lists**	Enables configuring user lists specifying user access for sessions.
Policies / **Usermapping Policies**	**Policies** > **Usermapping Policies**	Enables configuring policies specifying who can use a specific user name to access the remote server.
Indexer	**Indexer Status**	Enables viewing information related to session indexing.
HTTP Control ACLs	Menu points under **Traffic Controls** > **HTTP**.	Enables configuring HTTP connections.
ICA Control ACLs	Menu points under **Traffic Controls** > **ICA**.	Enables configuring ICA connections.
MSSQL Control ACLs	Menu points under **Traffic Controls** > **MSSQL**.	Enables configuring MSSQL connections.

ACL element	Menu items to which access is granted on the SPS UI	ACL summary
RDP Control ACLs	Menu points under **Traffic Controls** > **RDP**.	Enables configuring RDP connections.
SSH Control ACLs	Menu points under **Traffic Controls** > **SSH**.	Enables configuring SSH connections.
Sudo iolog Control	**Traffic Controls** > **Sudo iolog**	Enables configuring Sudo iolog connections.
Telnet Control ACLs ACLs	Menu points under **Traffic Controls** > **Telnet**.	Enables configuring Telnet connections.
VNC Control ACLs	Menu points under **Traffic Controls** > **VNC**.	Enables configuring VNC connections.
Reporting ACLs	Menu points under **Reporting**.	Enables configuring, generating and downloading reports via SPS.
Unlock Credential Store	**User menu** > **Unlock Credential store**	Enables accessing user credential stores that store credentials for authenticating to target servers.

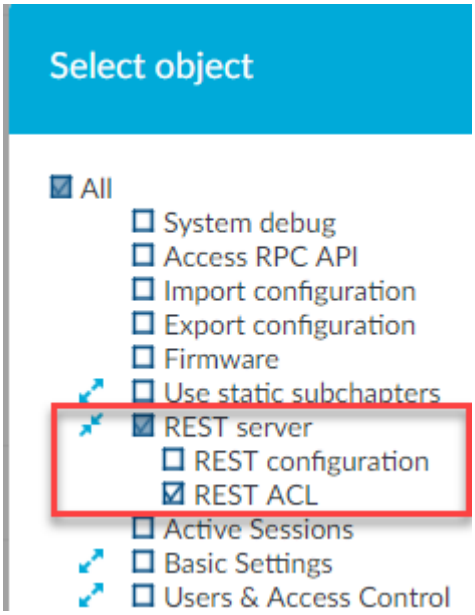
Creating rules for restricting access to search audit data

If you want users to access audit data on the Search interface only for sessions for which they are granted permission, complete the following steps.

To be able to see the **Audit Data Access** menu item and use this functionality, you must enable the **REST ACL** check box as described below.

1. Navigate to **Users & Access Control** > **Appliance Access**.
2. Click **Edit**.
3. Expand **REST server**.
4. Select the **REST ACL** check box.

Figure 88: Users & Access Control > Appliance Access — Select REST ACL



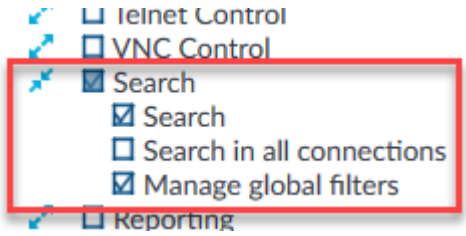
The following describes how you can create rules to restrict the search by providing access privileges for users to audit data.

Prerequisites

1. You have created a local user group as described in Managing local user groups.
2. You have added search access rights to your local user group as described in Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface.

   NOTE: Ensure that you clear the **Access all sessions** check box as shown below.

Figure 89: Users & Access Control > Appliance Access — Add search access rights to local user groups



3. You have created a local user, and added the user to the local user group. For more information, see Managing user rights and user groups.

To create search rules to access audit data

1. If you have multiple SPS appliances and they are organized into a cluster where one of the nodes is the Search Master (or Central Search) node, log in to that node.
2. Navigate to **Users & Access Control** > **Audit Data Access**.
3. Click **Create new**.
4. In the **Name** field, enter a name for your rule.
5. In the **Groups** field, enter an existing local user group for which you want to restrict access to audit data.
6. In the **Query** field, enter the correct query syntax to define which audit data this user group can access.
7. Optionally, for a quick visualization of the audit data that each group can access with this search query, click **Preview**.

   TIP: If required, modify the query syntax using the **Query** field, and the preview is updated accordingly.

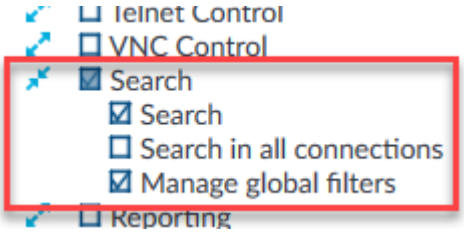
Example: Restrict access to search audit data for a user

You want a user to access audit data on the Search interface related to the SSH protocol only.

1. Create a local user group, for example, **`search-only-ssh`** as described in Managing local user groups.
2. Add search access rights to your **`search-only-ssh`** user group as described in Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface.

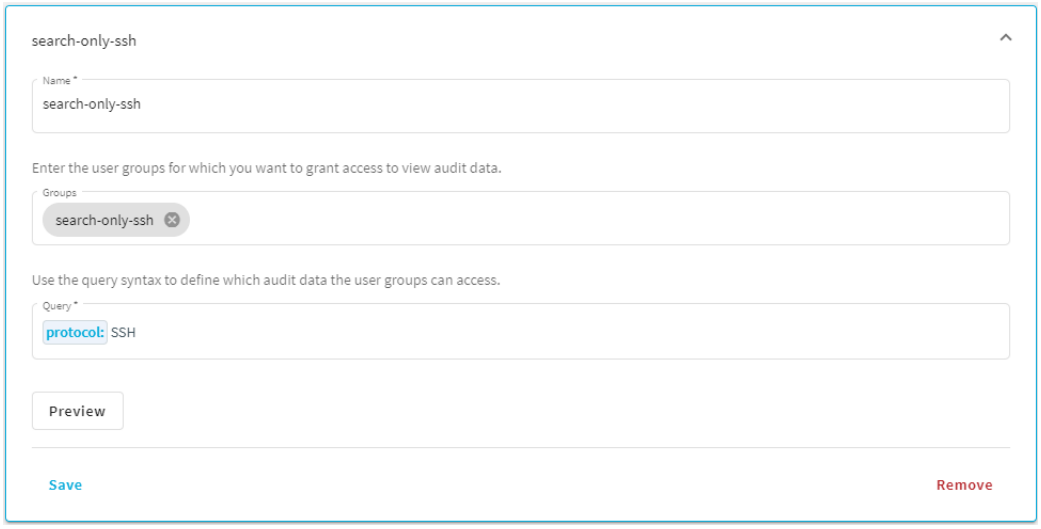
   NOTE: Ensure that you clear the **Access all sessions** check box.

Figure 90: Users & Access Control > Appliance Access — Add search access rights to local user groups



3. Create an audit data access rule (ADAR) for the **search-only-ssh** user group. Use the correct query syntax, for example, `protocol: SSH`.

Figure 91: Users & Access Control > Audit Data Access — Create new rule

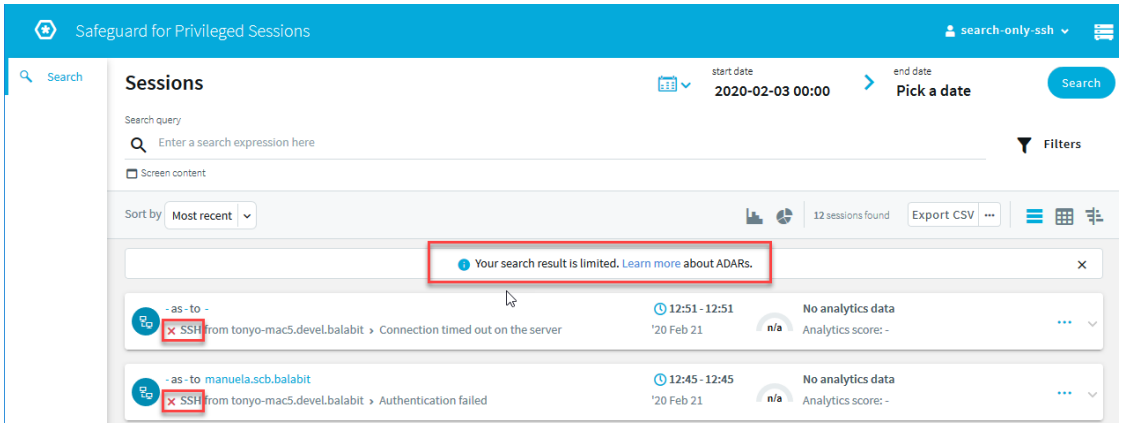


Optionally, for a quick visualization of the audit data that the **search-only-ssh** group can access with this search query, click **Preview**.

4. Create a local user, and add the user to the **search-only-ssh** group. For more information, see Managing user rights and user groups.

Result: Your local user will have access to audit data related to the SSH protocol only on the Search interface.

Figure 92: Audit > Sessions — Only SSH audit data is displayed for the user



Displaying the privileges of users and user groups

One Identity Safeguard for Privileged Sessions (SPS) version 3.2 and later provides an interface to query the user-rights and privileges of individual users and user groups. To display the privileges of a user or usergroup, navigate to **Users & Access Control** > **Access Rights Report**, enter the name of the user or group into the respective field, then click **Filter**. Note that:

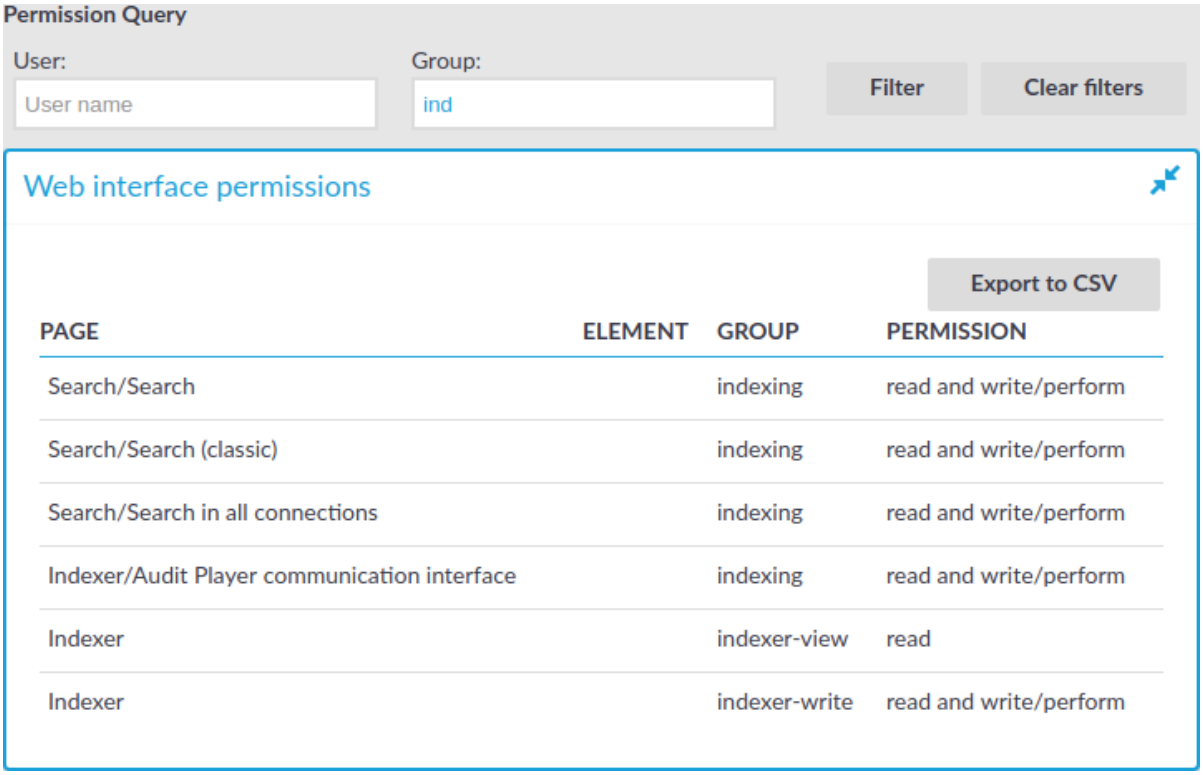
- It is not possible to filter on both the username and the group at the same time.
- Partial matches are also displayed.
- Usergroups can be local usergroups, userlists, or LDAP usergroups.

Web interface permissions

For usergroups accessing the SPS web interface, a table is displayed that lists the pages of the SPS web interface that the user or usergroup can access. The following information is displayed:

- **Page**: The name of the page or group of pages, for example, **Basic Settings**.
- **Element**: If a group has access only to a section of a page, the name of the element is listed here. For example, a particular Channel Policy.
- **Group**: The name of the usergroup.
- **Permission**: The type of access that the user or usergroup has to the page: `read` or `read and write/perform`.

Figure 93: Users & Access Control > Access Rights Report — Displaying web interface permissions

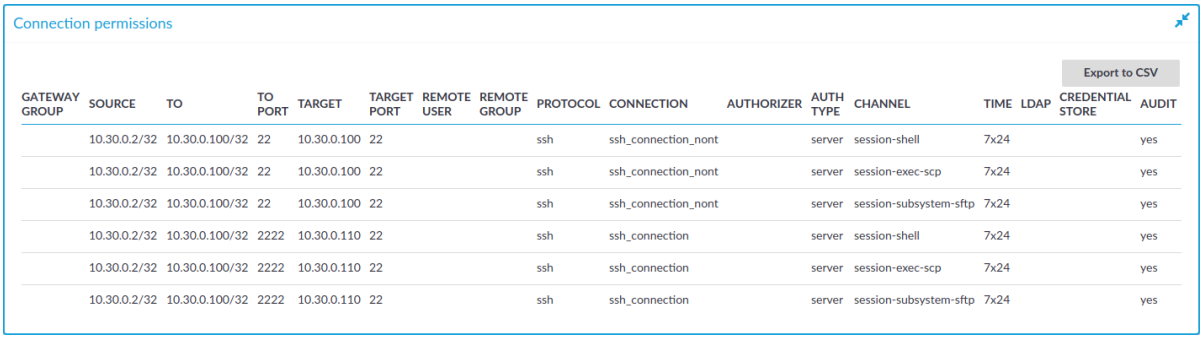


Connection permissions

To review which servers a user or usergroup can access, SPS collects the main information about the connections the user or group is permitted to use. The following information is displayed.

NOTE: To display the usergroups that can access a specific Connection Policy, open the Connection Policy, then on the **Connections** page, select **Show connection permissions** > **Show**.

Figure 94: Users & Access Control > Connection permissions — Displaying connection permissions



- **Gateway group**: Lists the group memberships required to access the connection. Group memberships can be restricted at the following places:
  - **Connection** > **Gateway authentication** > **Groups**
  - **Channel Policies** > **Gateway group**
  - **Policies** > **Usermapping Policies** > **Groups**
- **Source**: Refers to the following field from the session database:

  Source IP: The IP address of the client.
- **To**: Refers to the following field from the session database:

  Destination IP: The IP address of the server as requested by the client.
- **To port**: Refers to the following field from the session database:

  Destination port: The port number of the server as requested by the client.
- **Target**: Refers to the following field from the session database:

  Server IP: The IP address of the server connected by SPS.
- **Target port**: Refers to the following field from the session database:

  Server port: The port number of the server connected by SPS.
- **Remote user**: Refers to the following field from the session database:

  Username on server: The username used to log in to the remote server. This username can differ from the client-side username if usermapping is used in the connection. For details on usermapping, see Configuring usermapping policies on page 860.
- **Remote group**: The group that can access the destination server, as set in the Usermapping Policy (if any).
- **Protocol**: The protocol used in the connection (Citrix ICA, HTTP, RDP, SSH, Telnet, or VNC).
- **Connection**: Refers to the following field from the session database:

  Connection policy ID: The identifier of the connection policy.
- **Authorizer**: Refers to the following field from the session database:

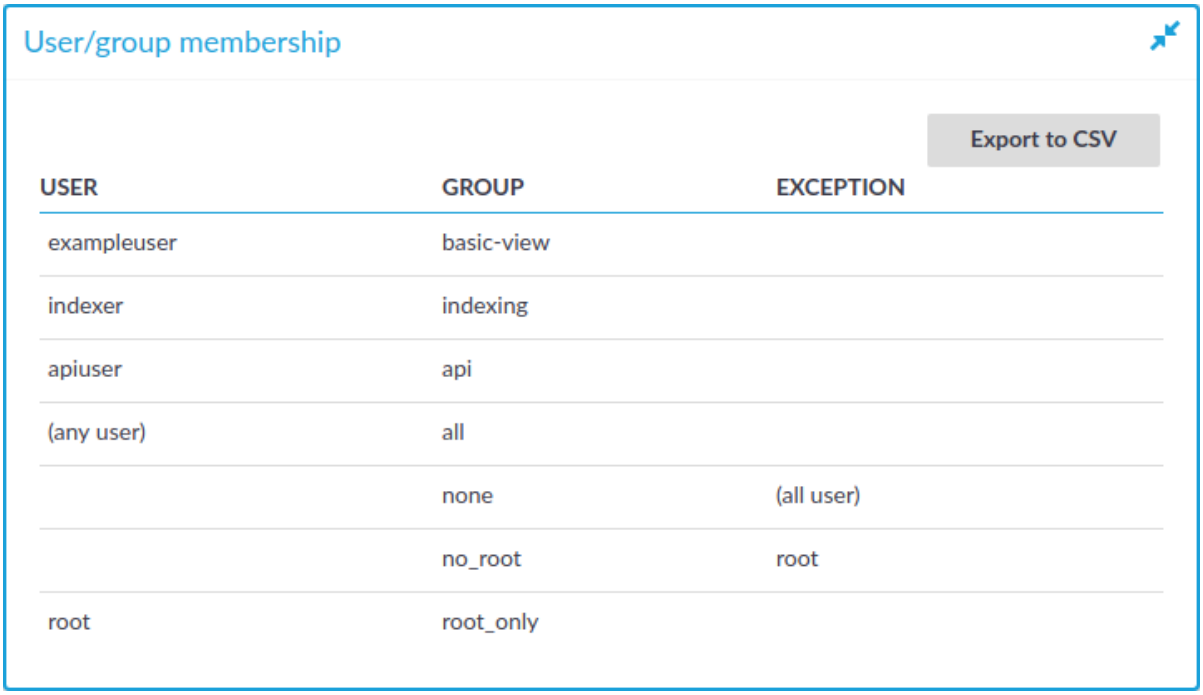
  Four-eyes authorizer: The username of the user who authorized the session. Available only if 4-eyes authorization is required for the channel. For details on 4-eyes authorization, see Configuring four-eyes authorization on page 870.
- **Auth type**: The authentication method used in the client-side connection during gateway authentication.
- **Channel**: The type of the channel, for example, `session-shell`.
- **Time**: The name of the Time Policy used in the connection.
- **LDAP**: The name of the LDAP Server used in the connection (if any).
- **Credential store**: The name of the Credential Store used in the connection (if any).
- **Audit**: Indicates if the connection is recorded into audit trails.

Usergroup memberships

When searching for users, the table displays the group memberships of the matching users. When searching for usergroups, the table displays the members of the matching groups. The following information is displayed:

- **User**: The username of the user.
- **Group**: The name of the usergroup or userlist.
- **Exception**: Usernames that are denied in case of default-deny userlists managed locally on SPS.

Figure 95: Users & Access Control > Connection permissions — Displaying usergroup and userlist memberships

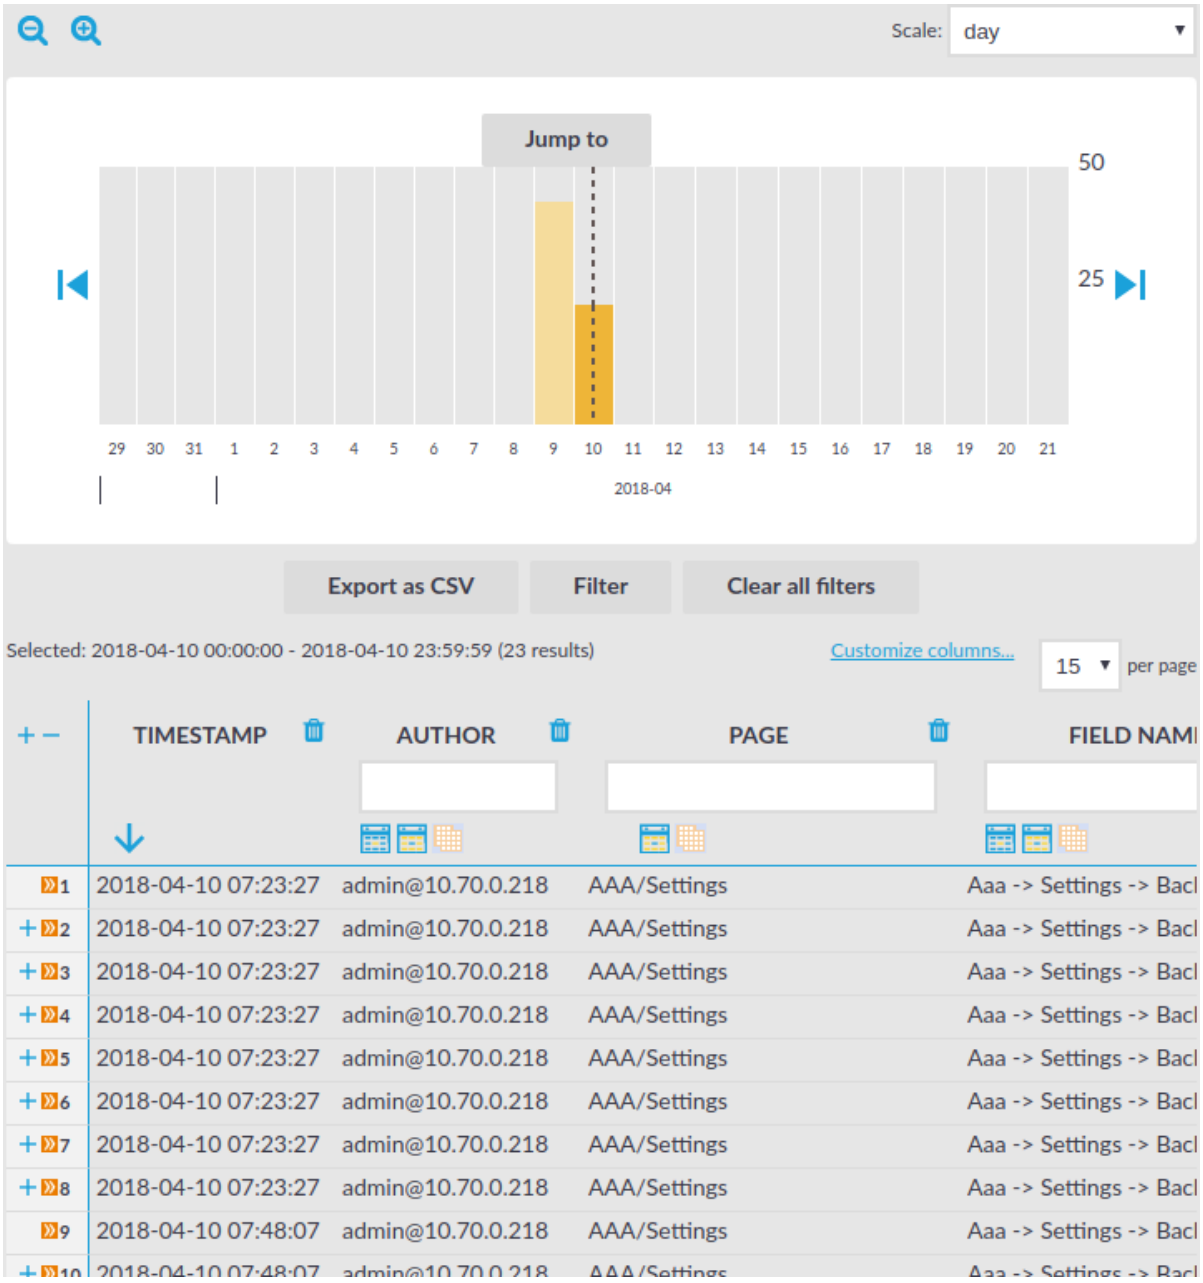


Listing and searching configuration changes

One Identity Safeguard for Privileged Sessions (SPS) automatically tracks every change of its configuration. To display the history of changes, select **Users & Access Control** > **Configuration History**. The changes are displayed on a search interface. For more information on using and customizing this interface, see Using the internal search interface on page 390.

The following information is displayed about each modification:

Figure 96: Users & Access Control > Configuration History — Browsing configuration changes



- **Timestamp**: The date of the modification.
- **Author**: Username of the administrator who modified the configuration of SPS.
- **Page**: The menu item that was modified.
- **Field name**: The name of the field or option that was modified.
- **New value**: The new value of the configuration parameter.

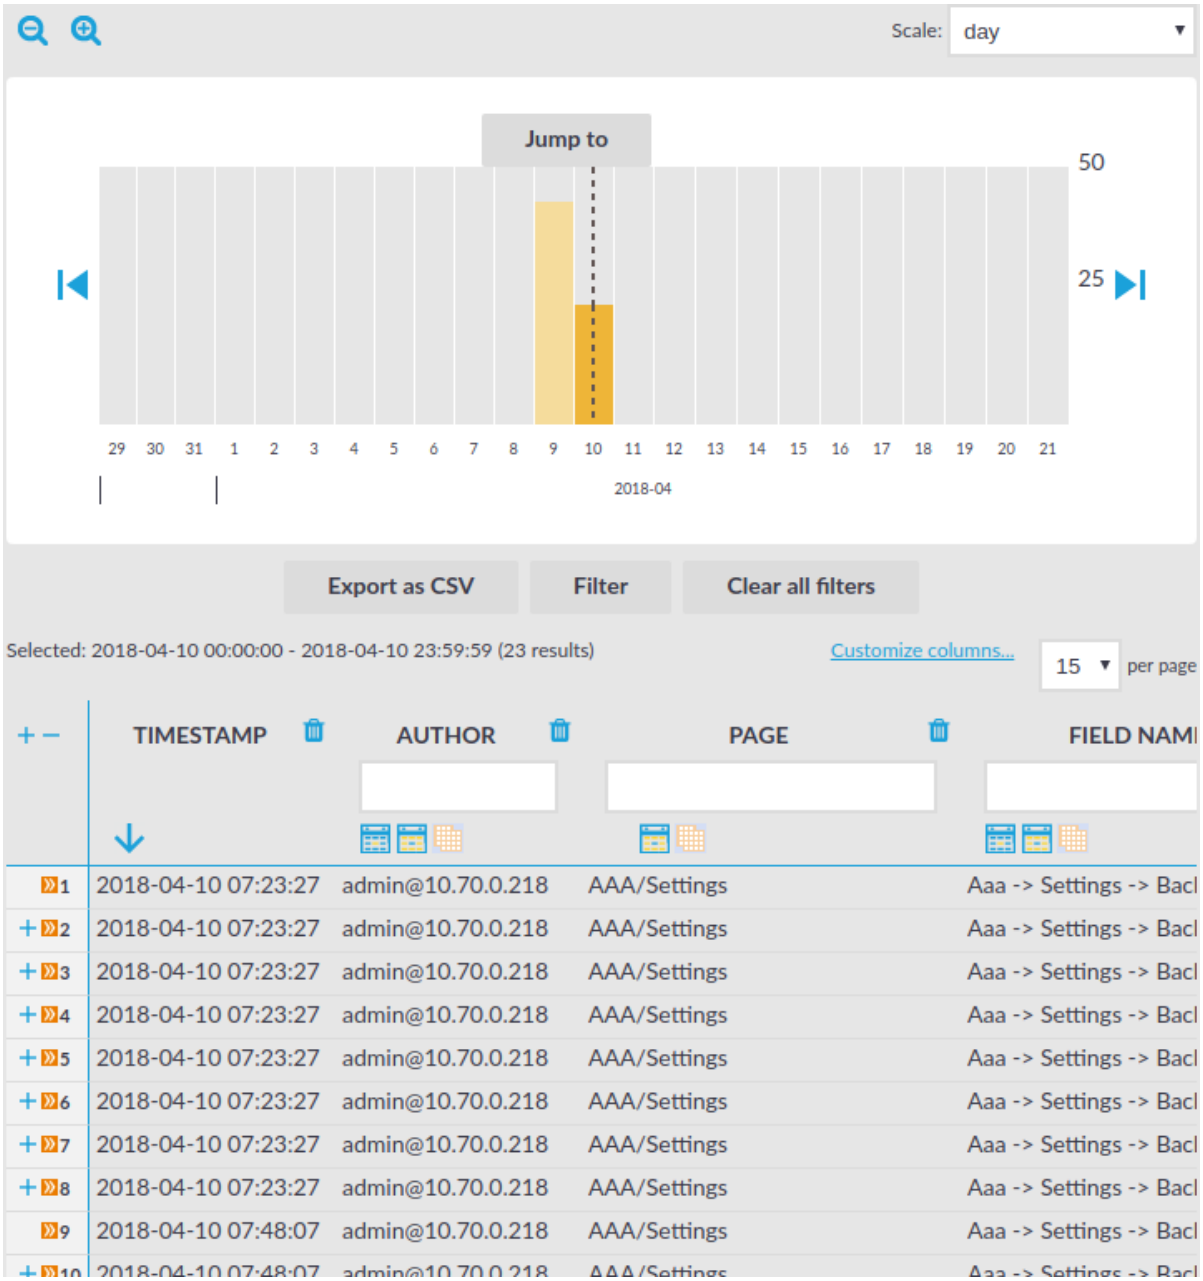
- **Message**: The changelog or commit log that the administrator submitted. This field is available only if the **Require commit log** option is enabled (see below).
- **Old value**: The old value of the configuration parameter.
- **Swap**: Signs if the order of objects was modified on the page (for example the order of two policies in the list).

To request the administrators to write an explanation to every configuration change, navigate to **Basic Settings** > **Management** > **Accounting settings** and select the **Require commit log** option.

Using the internal search interface

The internal search interface is for browsing and filtering the configuration changes and reports of One Identity Safeguard for Privileged Sessions (SPS).

Figure 97: Users & Access Control > Configuration History — The internal search interface



The bars display the number of results in the selected interval. Use the and icons to zoom, and the arrows to display the previous or the next intervals. To explicitly select a date, select **Jump to** and set the date in the calendar. You can change the length of the displayed interval with the **Scale** option.

Hovering the mouse above a bar displays the number of entries and the start and end date of the period that the bar represents. Click a bar to display the entries of that period in the table. Use Shift+Click to select multiple bars.

If data is too long to fit on one line, it is automatically wrapped and only the first line is displayed. To expand a row, click +. To shrink the row back to its original size, click −. To expand/shrink all rows, click the respective button on the header of the table. The rows can also be expanded/shrunk by double clicking on the respective row.

Filtering

The tables can be filtered for any parameter, or a combination of parameters. To filter the list, enter the filter expression in the input field of the appropriate column, and press `Enter`, or click on an entry in the table.

NOTE: When you use filters, the bars display the statistics of the filtered results.

Filtering displays also partial matches. For example, filtering the **Author** column on the **Users & Access Control** > **Configuration History** screen for `adm` displays all changes performed by users whose username contains the `adm` string.

You can use the icon to perform an exact search, and the icon for inverse filtering ("does not include"). To clear filters from a column, click .

To restore the original table, click **Clear all filters**.

Exporting the results

To save the table of search results as a file, click **Export as CSV**. This saves the table as a text file containing comma-separated values. Note that if an error occurs when exporting the data, the exported CSV file will include a line (usually as the last line of the file) starting with a zero and the details of the problem, for example `0;description_of_the_error`.

CAUTION:

Do not use **Export as CSV** to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load.

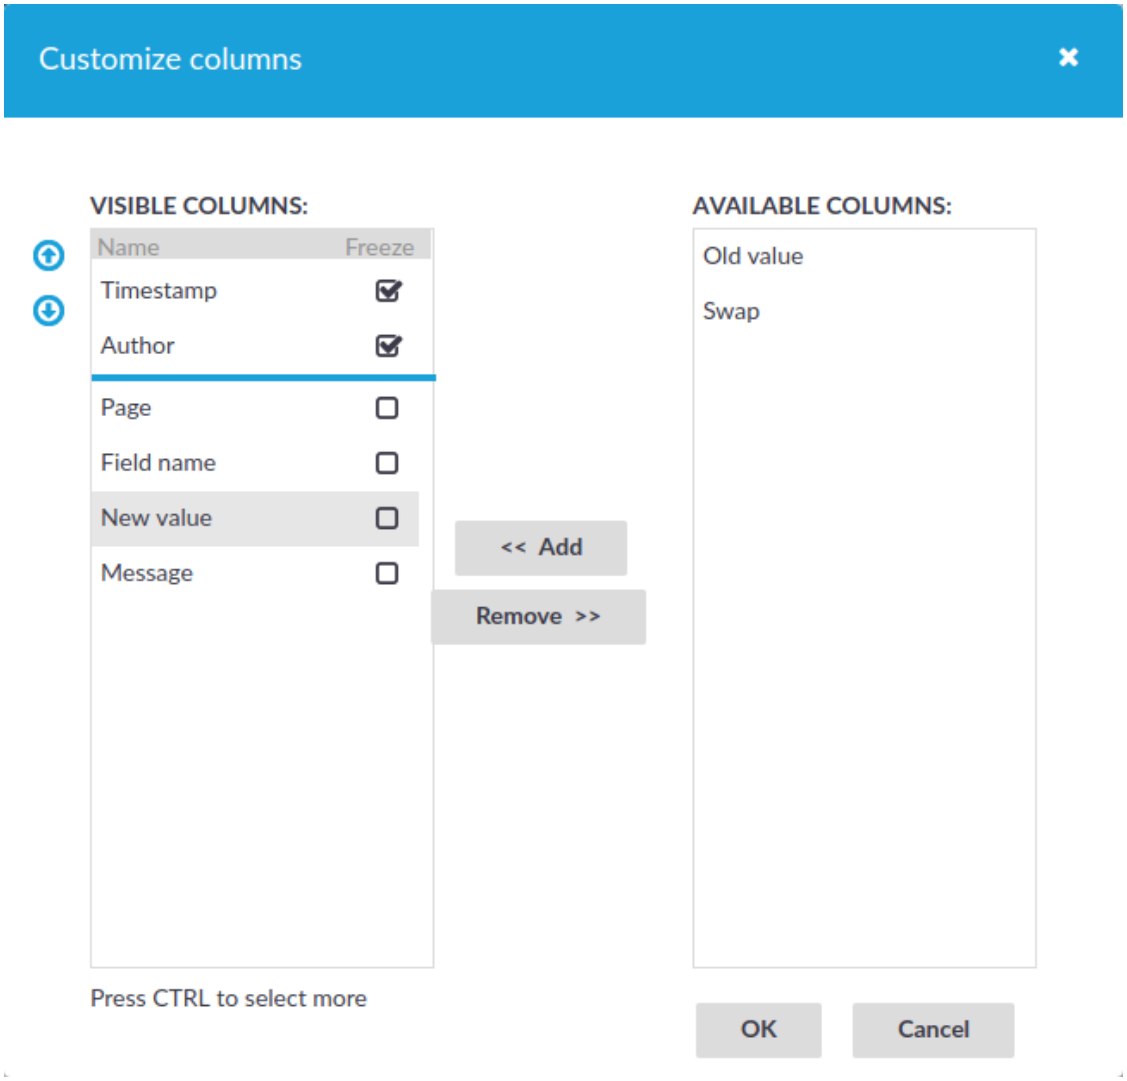
Customizing columns of the internal search interface

The following describes how to customize the data displayed on the interface.

To customize the data displayed on the interface

1. Navigate to the database you want to browse, for example **Users & Access Control** > **Configuration History**.
2. Click **Customize Columns**. A pop-up window containing the list of visible and available columns is displayed.

Figure 98: Users & Access Control > Configuration History — Customizing columns of the general search interfaces



3. The displayed parameters are enlisted in the **Visible columns** field. All other available parameters are enlisted in the **Available columns** field.
    - To add parameters to the **Visible columns** field, select the desired parameter (s) and click **Add**.
    - To remove parameters from the **Visible columns** field, select the desired parameter(s) and click **Remove**.
    - To freeze columns (to make them permanently visible, even when scrolling horizontally), enable the **Freeze** option next to the desired parameter.

    NOTE: To select multiple parameters, press`Ctrl`while clicking the items.
4. Click **OK**. The selected information is displayed.

Managing One Identity Safeguard for Privileged Sessions (SPS)

The following sections explain the basic management tasks of One Identity Safeguard for Privileged Sessions (SPS.

- For basic management tasks (reboot and shutdown, disabling traffic), see Controlling One Identity Safeguard for Privileged Sessions (SPS): reboot, shutdown on page 394.
- For information on managing a cluster of two or more SPS instances, see Managing One Identity Safeguard for Privileged Sessions (SPS) clusters on page 397.
- For managing a High Availability cluster, see Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster on page 423.
- For instructions on upgrading SPS, see Upgrading One Identity Safeguard for Privileged Sessions (SPS) on page 432.
- For instructions on accessing SPS through console and SSH, see Accessing the One Identity Safeguard for Privileged Sessions (SPS) console on page 442.
- For enabling sealed mode (which disables basic configuration changes from a remote host), see Sealed mode on page 449.
- For information on configuring the out-of-band (IPMI) interface, see Out-of-band management of One Identity Safeguard for Privileged Sessions (SPS) on page 450.
- For managing certificates used on SPS, see Managing the certificates used on One Identity Safeguard for Privileged Sessions (SPS) on page 459.

Controlling One Identity Safeguard for Privileged Sessions (SPS): reboot, shutdown

To reboot or shut down One Identity Safeguard for Privileged Sessions (SPS)

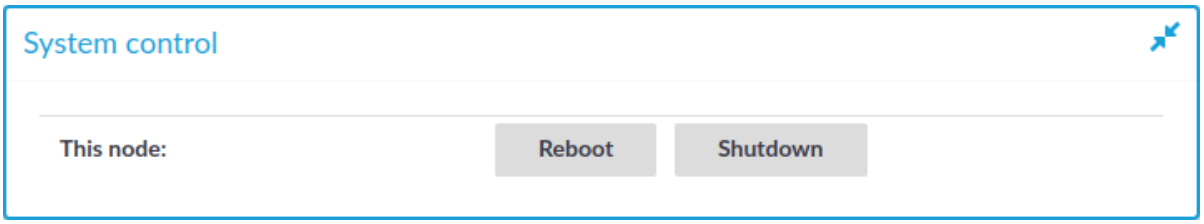
1. Navigate to **Basic Settings** > **System** > **System control** > **This node**.
2. Click the respective action button.

The **Other node** refers to the secondary node of a High Availability SPS cluster. For details on High Availability clusters, see Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster on page 423.

CAUTION:

- When rebooting the nodes of a cluster, reboot the other (secondary) node first to avoid unnecessary takeovers.
- When shutting down the nodes of a cluster, shut down the other (secondary) node first. When powering on the nodes, start the primary node first to avoid unnecessary takeovers.
- When both nodes are running, avoid interrupting the connection between the nodes: do not unplug the Ethernet cables, reboot the switch or router between the nodes (if any), or disable the HA interface of SPS.

Figure 99: Basic Settings > System > System Control — Performing basic management



NOTE: Web sessions to the SPS interface are persistent and remain open after rebooting SPS, so you do not have to relogin after a reboot.

During the reboot process, SPS displays information about the progress of the reboot and any possible problems in the following places:

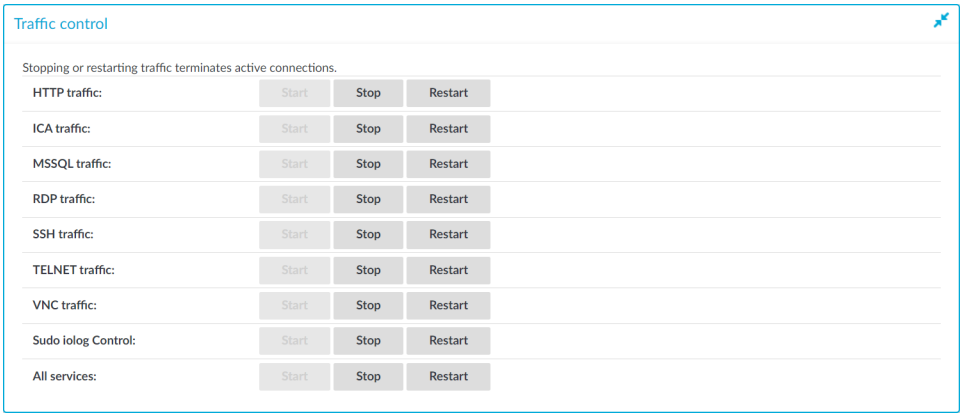
- On the web interface of SPS, at any of the **Listening addresses** configured at **Basic settings** > **Local Services** > **Web login (admin and user)**. (After booting, you are directed to the login screen of SPS.)
- On the console, which you can monitor with IPMI (ILOM) or console access.

The information displayed in the browser and on the console is the same.

Disabling controlled traffic

This section describes how to temporarily disable some or all of the controlled traffic to the protected servers.

Figure 100: Basic Settings > System > Traffic control — Disabling the controlled traffic



CAUTION:

Using the **Stop** option of the respective traffic types on **Basic Settings** > **System** > **Traffic control** disables the traffic only temporarily. Committing specific changes on the SPS web interface enables the connections again. For details on how to permanently disable a type of traffic, see Disabling controlled traffic permanently on page 397.

NOTE: Disabling the traffic affects only the traffic configured in the Connection policies. Other network traffic (such as web management, DNS, LDAP, SNMP, SMTP, and so on) can pass SPS even if all traffic is disabled. For details on configuring Connection policies, see General connection settings on page 482.

To temporarily disable some or all of the controlled traffic to the protected servers

1. Navigate to **Basic Settings** > **System** > **Traffic control**.
2. To disable any of the supported traffic types, click **Stop** in the respective traffic field. You can disable the following traffic types on this page:
   - HTTP
   - ICA
   - MSSQL
   - RDP
   - SSH
   - Telnet
   - VNC
   - Sudo iolog
   - All services

   NOTE: Disabling SSH traffic also disables all other traffic forwarded in SSH, for example X11.

Disabling Telnet also disables TN3270 traffic.

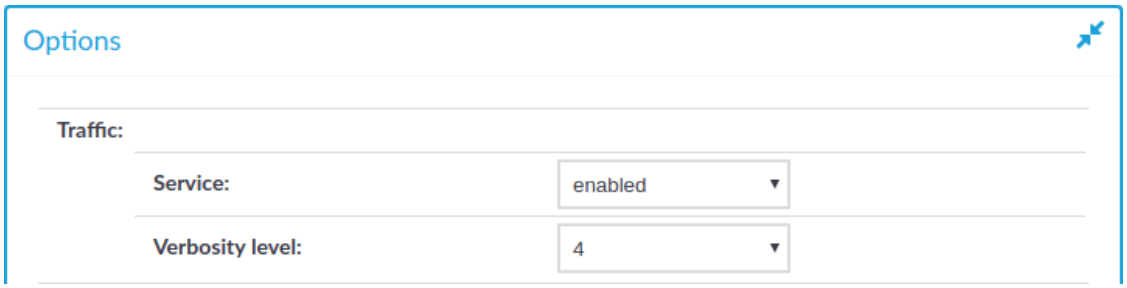
The **Traffic overview** under **About** displays the status of all traffic types.

Disabling controlled traffic permanently

NOTE: Disabling the traffic affects only the traffic configured in the Connection policies, other traffic can pass SPS even if the all traffic is disabled. For details on configuring Connection policies, see General connection settings on page 482.

To disable controlled traffic permanently

1. **Figure 101: Traffic Controls > Protocol name > Global Options — Disabling the controlled traffic persistently**



   Navigate to the **Global Options** page of the traffic type you want to disable, for example to **Traffic Controls** > **SSH** > **Global Options** to disable SSH traffic.
2. Set the **Traffic** > **Service** field to `disabled`.
3. Click Commit.

Managing One Identity Safeguard for Privileged Sessions (SPS) clusters

When you have a set of two or more One Identity Safeguard for Privileged Sessions (SPS) instances in your deployment, you can join them into a cluster. This has several advantages. You can:

- Manage the nodes from one central location.
- Monitor their status and update their configuration centrally.
- Search all session data recorded by all nodes in the cluster on a single node.
- Scale the performance of the cluster by adding new nodes and joining them to the cluster easily.
- Extend auditing to other networks by adding new nodes to the cluster and joining them to the cluster.

This is achieved by assigning roles to the individual nodes in your cluster: you can set one of your SPS nodes to be the Central management node and the rest of the nodes are managed from this central node.

NOTE: Consider the following:

- All nodes in a cluster must run the same version of SPS.
- One Identity recommends managing not more than a few tens of instances from the Central management node.
- Nodes in the cluster connect to each other using IPsec.

You can configure your One Identity Safeguard for Privileged Sessions (SPS) cluster in the following ways:

- **Configuration synchronization without a central search**: This method allows you to perform your configuration settings on your Central management node. Managed host nodes periodically fetch and merge the settings into their own: this is called "configuration synchronization". Central search is not configured in this method, so you can search for sessions on each node, including the Central management node.

  For more information on this method, see Configuration synchronization without a central search.
- **Central search with configuration synchronization**: This method allows you to use a Central management node with a Search master role to view session data recorded by the minion nodes of your cluster, as well as manage all the nodes in the cluster from one central location.

  For more information on this method, see Central search with configuration synchronization.

  IMPORTANT: One Identity does not recommend having a central search configuration without configuration synchronization.

Cluster roles

One Identity Safeguard for Privileged Sessions (SPS) offers several types of roles that you can assign to the nodes of an SPS cluster. These roles are listed below.

NOTE:

- Nodes keep their assigned role in the cluster even after a system restore.
- Depending on their type, you can assign certain roles only to a single node in the cluster, while others to multiple nodes.
- One node can also have multiple roles assigned to it, with some restrictions depending on the selected roles. For details, see the description of the roles.

Central management

NOTE: You can assign the **Central management** role only to a single node in the cluster.

The Central management node is used to set up a central configuration within the cluster, by having the configuration of this node synchronized to the other nodes of the cluster. Any changes that you make in the cluster configuration of this node (for example, role changes, host address changes, and so on) are fetched by the rest of the nodes in the cluster, and are merged into their configuration.

The **Central management** node also contains status information about all the other nodes in the cluster, so you can check the health of the whole cluster on this node. This status information contains:

- The time and result of the last attempt at fetching and updating the configuration.
- The errors and warnings that may have occurred during fetching and updating the configuration.

Managed host

NOTE: You can assign the **Managed host** role to multiple nodes in the cluster.

Managed host nodes in a cluster synchronize their entire configuration from the **Central management** node, not just the cluster-related elements of their configuration. Managed host nodes send their status information to the **Central management** node every 10 seconds.

Search master

NOTE:

- A One Identity Safeguard for Privileged Sessions (SPS) node with the **Search master** role cannot be used for monitoring network traffic, or for session recording and auditing purposes. Before assigning this role to a node, read carefully the limitations that apply to **Search master** nodes in Managing a cluster with central search configuration and configuration synchronization.
- You can assign the **Search master** role only to a single node in the cluster.

Search master: The Search master node allows you to search all the session data recorded by other nodes in the cluster, provided that those other nodes are set to the **Search minion** role.

This role can only be assigned to nodes that either have the **Managed host** or **Central management** role. This is required so that the configuration of **Search minion** nodes and the **Search master** node are always in sync.

If there is no configuration synchronization between the node acting as the **Search master** and the **Search minion** nodes, then session data may show up on the **Sessions** interface of the **Search master** that come from connections that do not match the connection policies set up on the **Search master** (because they come from session data recorded by the **Search minions**).

Search minion

NOTE: You can assign the **Search minion** role to multiple nodes in the cluster.

Search minion nodes in a cluster send session data that they recorded to the **Search master** for central search purposes. The session data recorded by a Search minion node is not searchable on the node itself, only on the **Search master**.

This role can only be assigned to nodes that either have the **Managed host** or **Central management** role. This is required so that the configuration of **Search minion** nodes and the **Search master** node are always in sync.

If there is no configuration synchronization between the node acting as the **Search master** and the **Search minion** nodes, then session data may show up on the **Sessions** interface of the **Search master** that come from connections that do not match the connection policies set up on the **Search master** (because they come from session data recorded by the **Search minions**).

Search local

NOTE: You can assign the **Search local** role to multiple nodes in the cluster.

Search local nodes keep their recorded session data for local searching. Therefore, the session data recorded by a Search local node is searchable on the node itself, but not on the **Search master** node (if there is one configured in the cluster).

SPP fetcher

NOTE:

- You can assign the **SPP fetcher** role only to a single node in the cluster.
- You can only assign the **SPP fetcher** role to a node that also has either the **Search local** or **Search minion** role assigned.
- Make sure that the One Identity Safeguard for Privileged Passwords (SPP) node is already connected to the SPS cluster before assigning the **SPP fetcher** role to a node. For details, see Linking SPS to SPP.

The SPP fetcher node fetches data from an SPP node linked to the cluster. The fetched data includes:

- Workflow data from the past 5 minutes first, and then new workflow data near real-time.
- Historical data from the past 1 year, starting with the past 1 week.

The time required for fetching depends on the amount of data to fetch. This data will be then available on the **Sessions** interface after the fetching process has finished. To track the fetch progress of historical data, check the status on the **About** page.

No role

Nodes in the cluster that have no roles assigned fetch only the cluster-related elements of their configuration from the **Central management** node.

Such nodes without any roles send their status information to the **Central management** node every 10 seconds.

Related information

- For more information about configuration synchronization, see Configuration synchronization across nodes in a cluster on page 411 and Managing a cluster with configuration synchronization without central search.
- For more information about central search, see Searching session data on a central node in a cluster on page 858 and Managing a cluster with central search configuration and configuration synchronization.

Enabling cluster management

To enable cluster management, enable the cluster interface on all nodes that you want to be part of your One Identity Safeguard for Privileged Sessions (SPS) cluster. Complete the following steps on each node of the cluster.

NOTE: All nodes in a cluster must run the same version of SPS.

Prerequisites

The nodes in the cluster must connect to each other using IPsec.

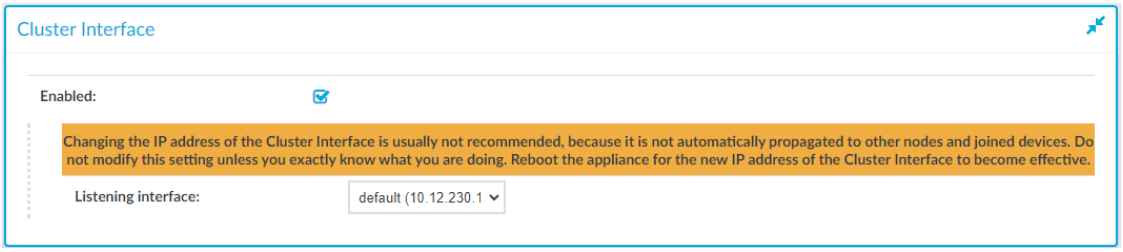
NOTE: IPsec requires UDP ports `500` and `4500` to be open bidirectionally in the firewalls between the nodes.

To enable cluster management

1. Navigate to **Basic Settings** > **Local Services** > **Cluster Interface**.
2. Select **Enabled**.

   The **Listening interface** field is displayed, showing the interfaces configured in the **Basic Settings** > **Network** page.
3. Select a cluster interface for the node to listen on.

Figure 102: Basic Settings > Local Services > Cluster Interface — Enabling cluster management



4. Click Commit.

Creating a cluster

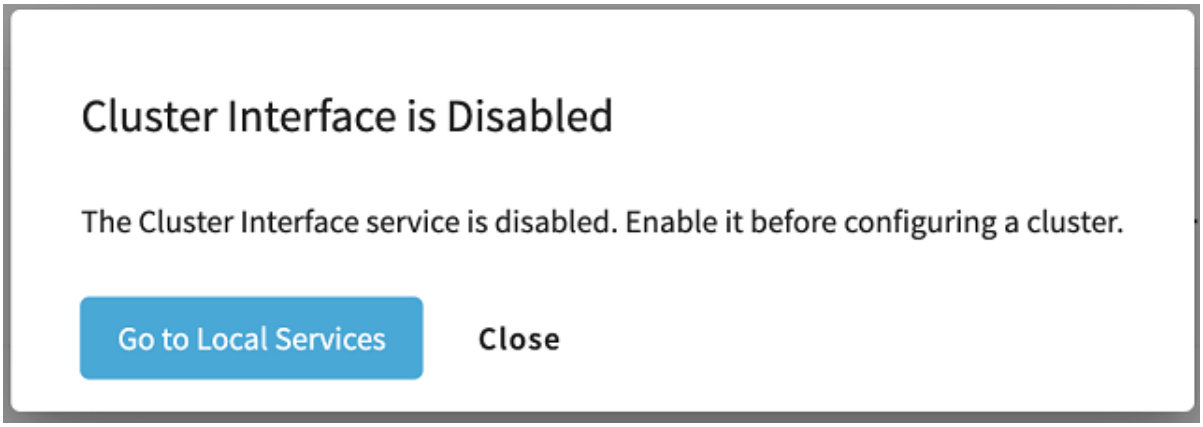
Create a cluster by promoting a One Identity Safeguard for Privileged Sessions (SPS) node to the role of the **Central Management** node. Next, join additional nodes to the cluster as described in Joining to a cluster.

NOTE: You can also promote a node through the REST API. For more information, see *Promote a SPS node to be the Central Management node in a new cluster* in the *REST API Reference Guide*.

Prerequisites

Enable the cluster interface on all nodes that you want to be part of your cluster. For more information, see Enabling cluster management on page 401. If no cluster interface is enabled, then clicking **Basic Settings** > **Cluster Management** results in the following pop-up dialog appearing:

Figure 103: Basic Settings > Cluster management — No cluster interface configured

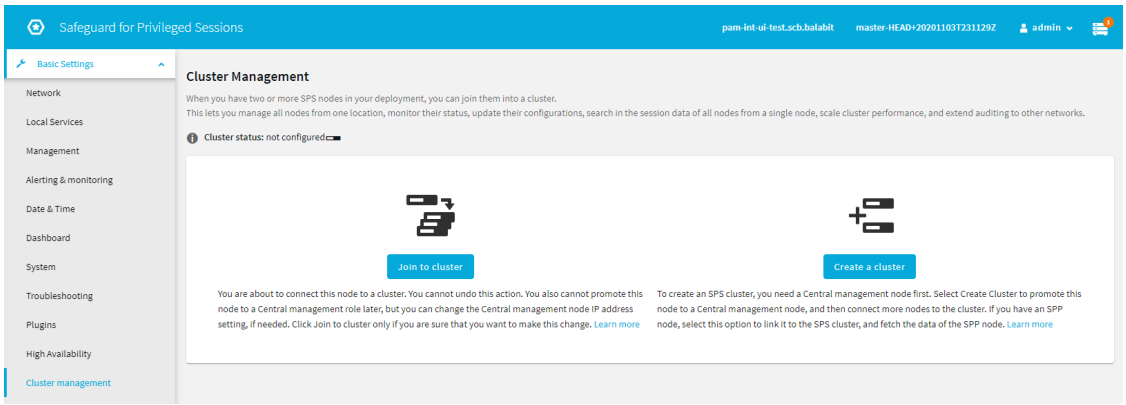


If this happens, click **Go to Local Services** to open **Basic Settings** > **Local Services** > **Cluster Interface**, and enable a cluster interface. After that, click **Basic Settings** > **Cluster management** again.

To create a cluster by promoting a node to Central management role

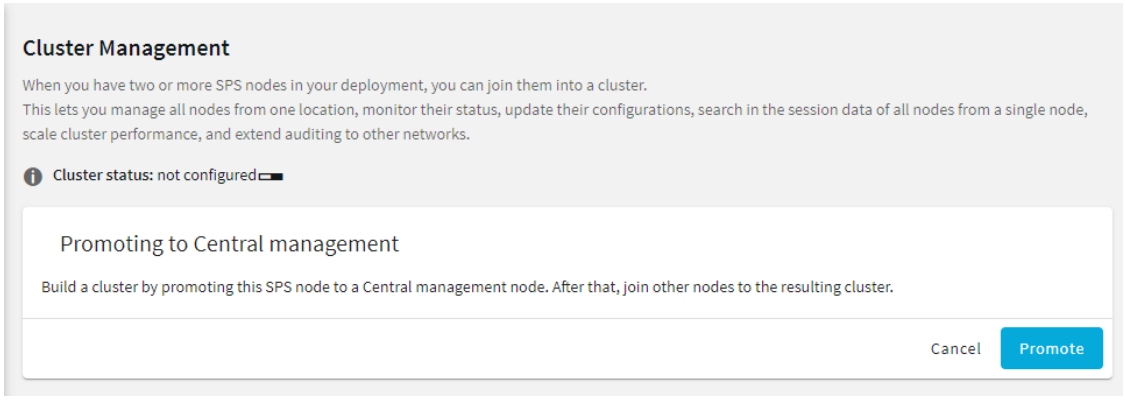
1. Navigate to **Basic Settings** > **Cluster management**. The main Cluster management window appears, allowing you to either join the node to an existing cluster, or create a new one.

Figure 104: Basic Settings > Cluster management — Join and create cluster options



2. Click **Create a cluster** to open the node promotion window.

Figure 105: Basic Settings > Cluster management — Promote node to Central management role to create cluster



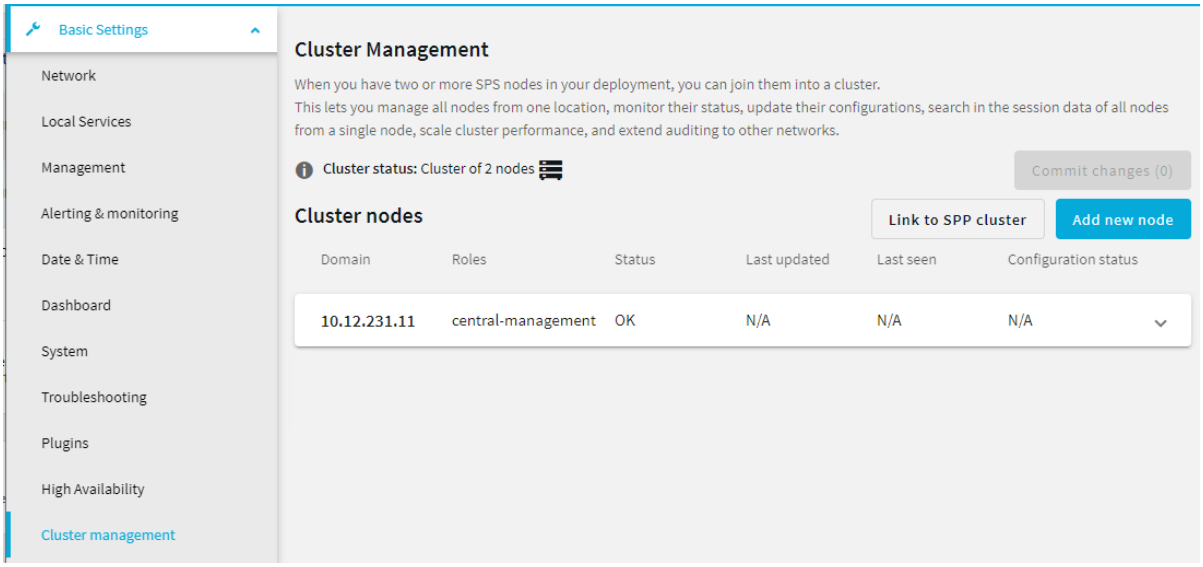
3. On the **Promoting to Central management role** dialog, click **Promote** to promote the node to Central management role, and create the cluster.

   NOTE: You cannot undo or modify this action later.

   The promotion status is indicated in the **Cluster status** row above the **Promoting to Central Management** role window.

Once the node is promoted, the **Cluster management** window is displayed, with the now-promoted Central management node listed in the screen.

Figure 106: Basic Settings > Cluster management — Cluster management dialog on the Central management node



Using this screen, you can join additional nodes to the cluster. For details on how to do that, see Joining to a cluster.

Joining to a cluster

To join additional One Identity Safeguard for Privileged Sessions (SPS) nodes to a cluster, generate a join token on the node that you want to add to the cluster, and then use that token on the Central management node to finish adding the node to the cluster.

CAUTION:

Configuration options that you set on a node before joining it to the cluster will be overwritten by the configuration of the Central management node. For example, policies and protocol-specific settings will be overwritten once you assign the Managed host role to the node. Managed host roles periodically fetch the configuration of the Central management node and merge it into their own. This is called configuration synchronization.

To avoid the loss of policies and settings that are specific to your Managed host node, use a configuration synchronization plugin. Such plugins enable you to limit the scope of configuration synchronization.

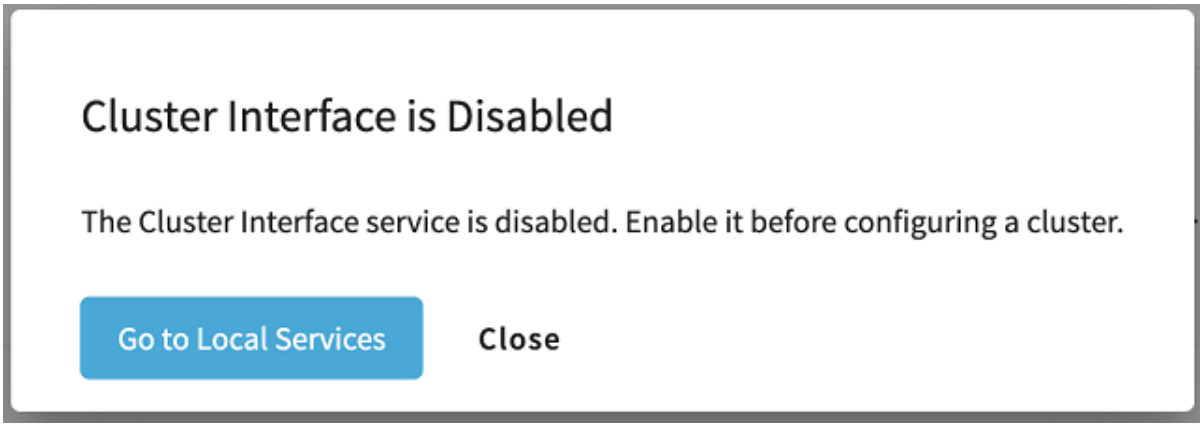
For more information, see Configuration synchronization across nodes in a cluster on page 411.

NOTE: You can also join additional nodes to your cluster through the REST API, too. For more information, see *Join node(s) to the cluster* in the *REST API Reference Guide*.

Prerequisites

Enable the cluster interface on all nodes that you want to be part of your cluster. For more information, see Enabling cluster management on page 401. If no cluster interface is enabled, then clicking **Basic Settings** > **Cluster Management** results in the following pop-up dialog appearing:

Figure 107: Basic Settings > Cluster management — No cluster interface configured

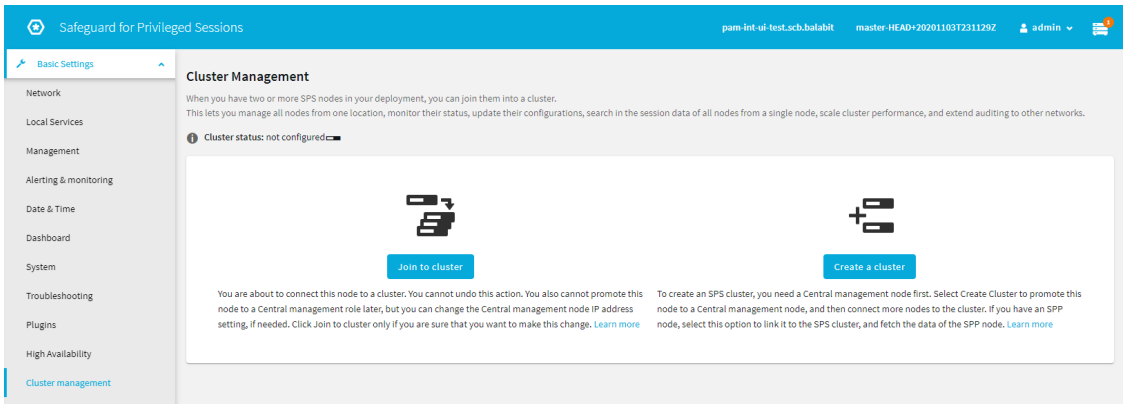


If this happens, click **Go to Local Services** to open **Basic Settings** > **Local Services** > **Cluster Interface**, and enable a cluster interface. After that, click **Basic Settings** > **Cluster management** again.

To join additional nodes to a cluster by generating and using join tokens

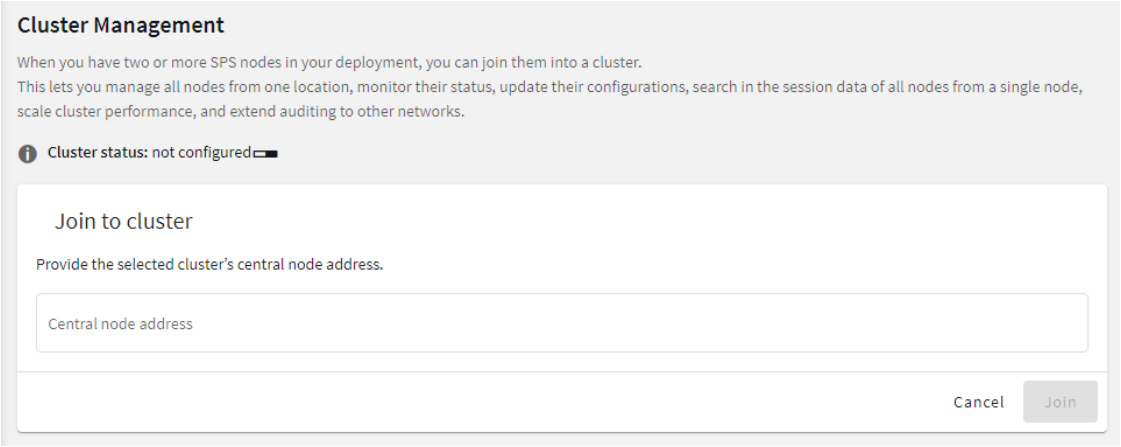
1. Navigate to **Basic Settings** > **Cluster management**. The main Cluster management window appears, allowing you to either join the node to an existing cluster, or create a new one.

   Figure 108: Basic Settings > Cluster management — Join and create cluster options

2. Click **Join to cluster** to open the cluster join window.

Figure 109: Basic Settings > Cluster management — Join to cluster window



3. On the **Join to cluster** dialog, in the **Central node address** field, enter the IP address of the Central management node of the cluster you want to join.

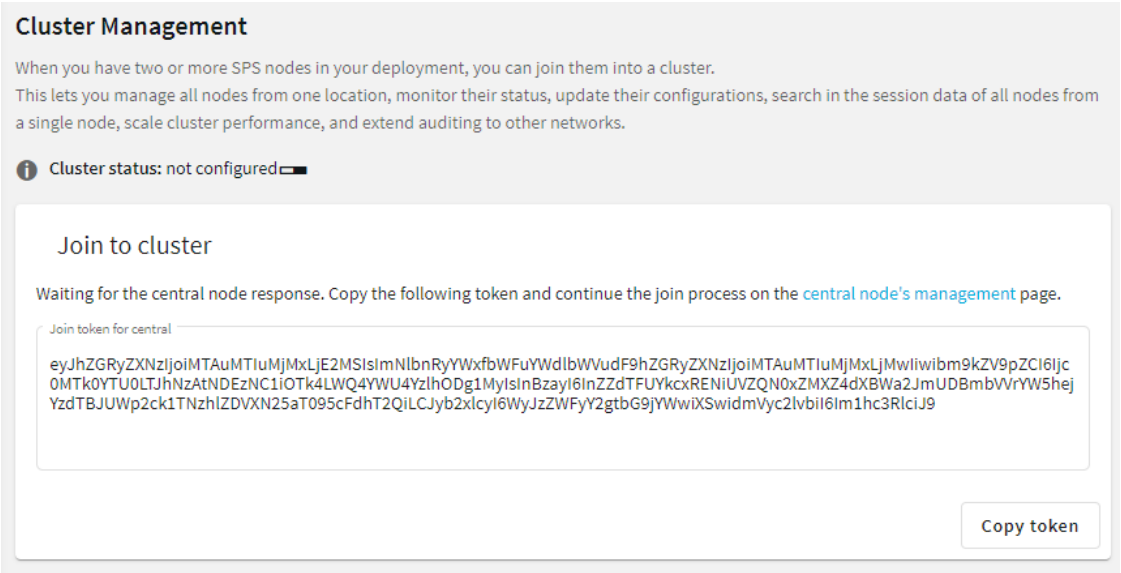
   TIP: You can check the IP address on the **Basic Settings** > **Cluster management** screen of the Central management node.

4. Click **Join**. A confirmation dialog appears. Click **Join to cluster** again in the dialog to proceed.

   NOTE: Once you click **Join to cluster**, you cannot undo the join process, and you will not be able to promote the node you are currently configuring to a Central management role later. However, you will still be able to change the IP address of the Central management node in the **Central Node Address** field, if needed.

5. A window with a join token appears.

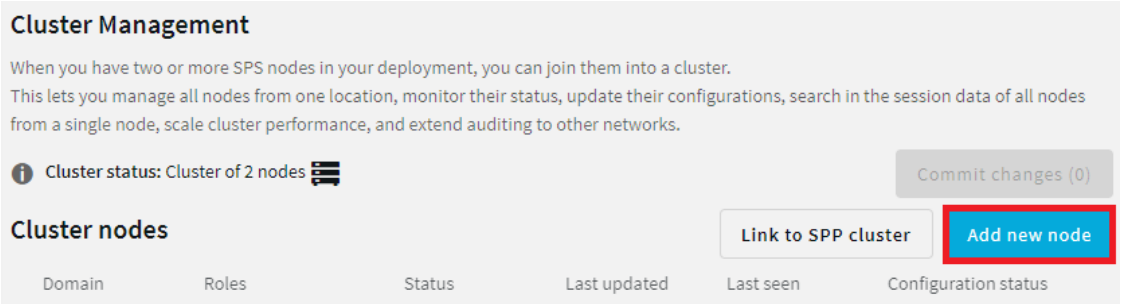
Figure 110: Basic Settings > Cluster management — Join token used to join a cluster



Click **Copy token** to copy the join token to the clipboard.

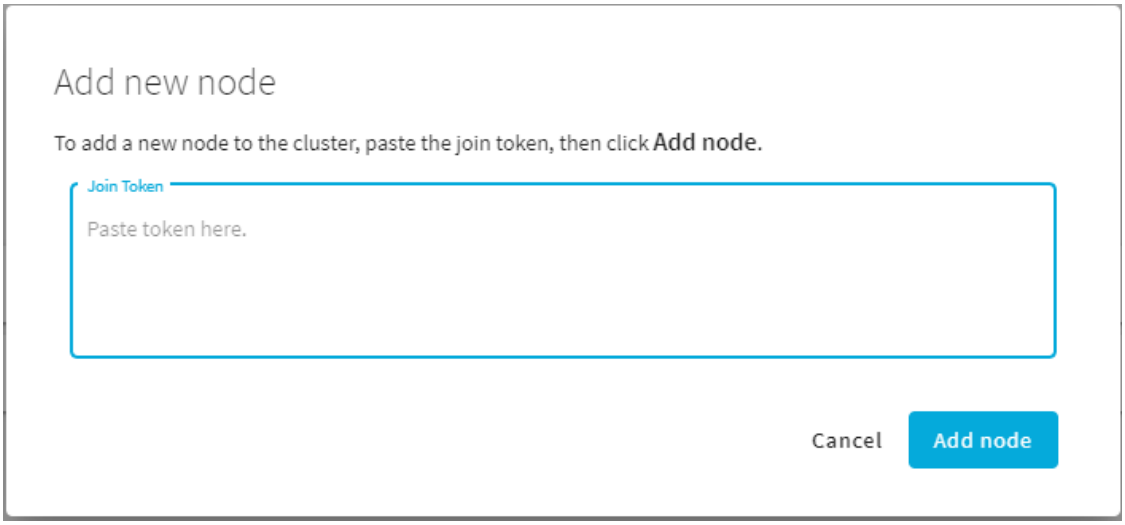
6. On the Central management node of the cluster you want to join, navigate to **Basic Settings** > **Cluster management**, and click **Add new node**.

Figure 111: Basic Settings > Cluster management — Cluster management dialog on the Central management node



7. The **Add new node** dialog appears.

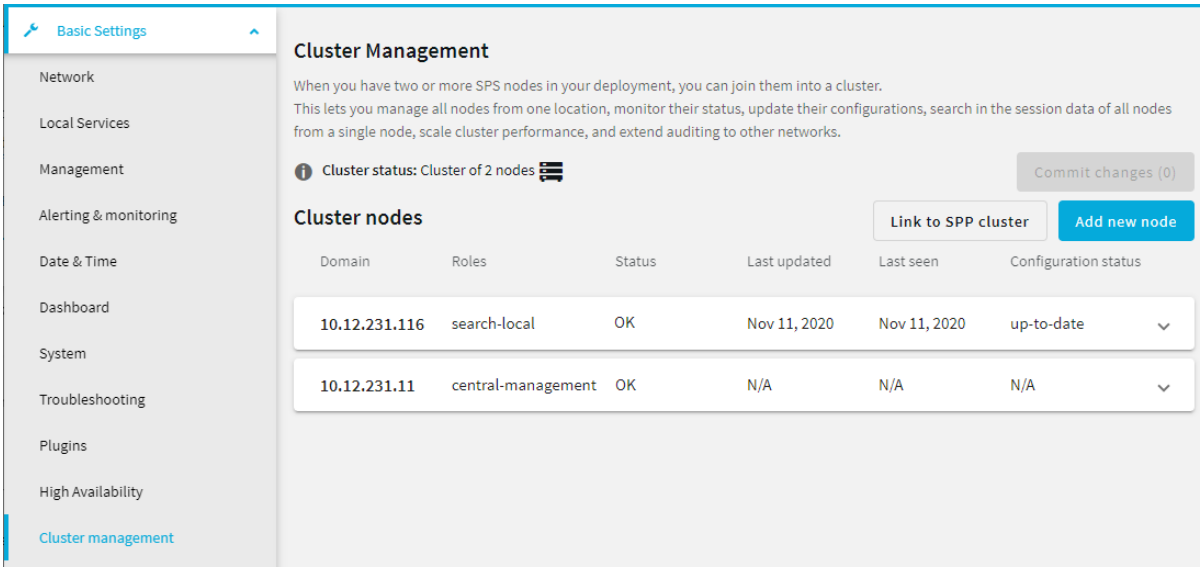
Figure 112: Basic Settings > Cluster management — Add new node dialog



Paste the token in the **Join Token** field, then click **Add node**.

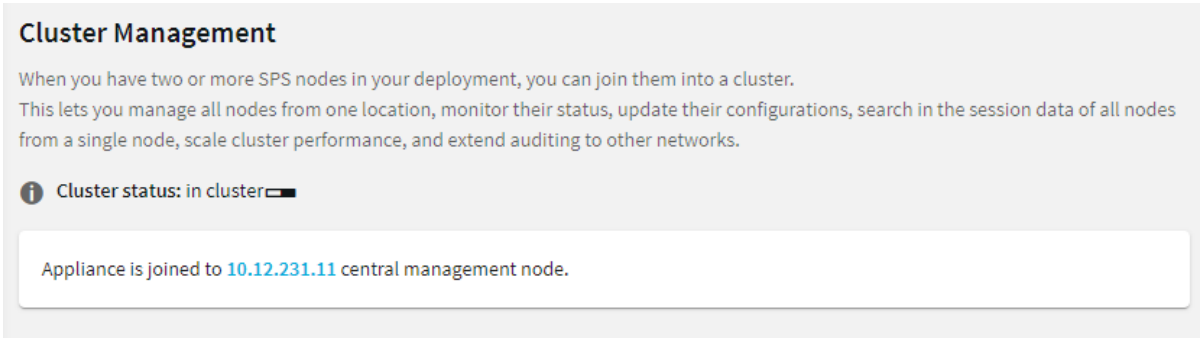
Once the node joined the cluster, it is displayed in the list of nodes on the **Basic Settings** > **Cluster management** window of the Central management node.

Figure 113: Basic Settings > Cluster management — New node indicated on the Cluster management node



At the same time, on the node that you joined to the cluster, the **Basic Settings** > **Cluster management** page shows the IP address of the Central management node of the cluster.

Figure 114: Basic Settings > Cluster management — Join status indicated on the node that newly joined to the cluster



If you want to centrally manage the configuration of the node(s) you have joined to the cluster, assign the **Managed host** role to them. For details, see Assigning roles to nodes in your cluster on page 408.

Assigning roles to nodes in your cluster

By default, nodes do not have any roles assigned to them. The only exception is the Central management node, which you specifically promoted to fulfill that role. To assign a role to a node in the cluster, complete the following steps.

To assign roles to nodes in your cluster

1. On the web interface of your Central management node, navigate to **Basic Settings** > **Cluster management**. This page displays all nodes in the cluster.
2. Click ⌄ at the right side of the row of the node that node that you want to update. The node row is expanded, showing the node address and the available roles.
3. Select the role that you want to assign to the node. For details on what each role means, see Cluster roles on page 398.

   CAUTION:

   Configuration options that you set on a node before joining it to the cluster will be overwritten by the configuration of the Central management node. For example, policies and protocol-specific settings will be overwritten once you assign the Managed host role to the node. Managed host roles periodically fetch the configuration of the Central management node and merge it into their own. This is called configuration synchronization.

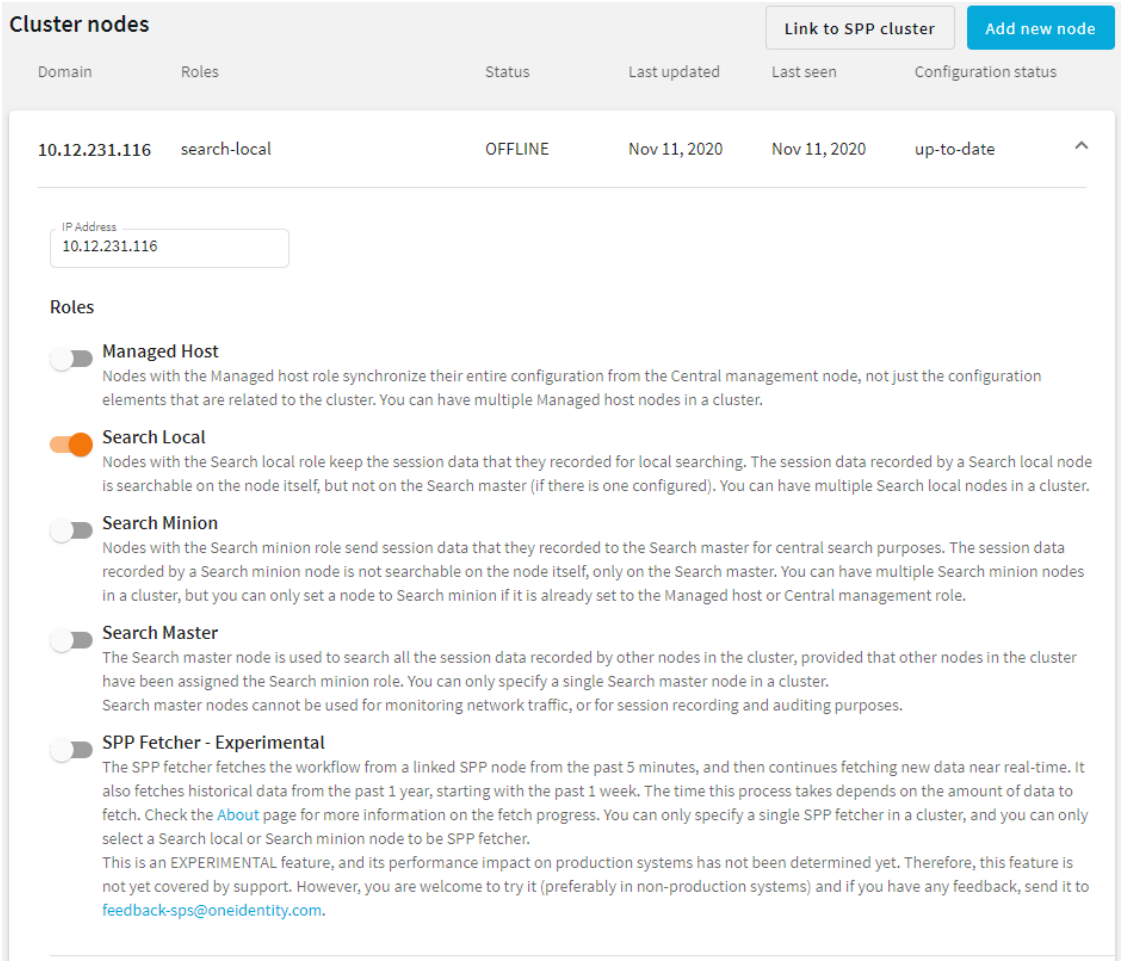
   To avoid the loss of policies and settings that are specific to your Managed host node, use a configuration synchronization plugin. Such plugins enable you to limit the scope of configuration synchronization.

   For more information, see Configuration synchronization across nodes in a cluster on page 411.

   NOTE: When assigning search roles, consider the following:

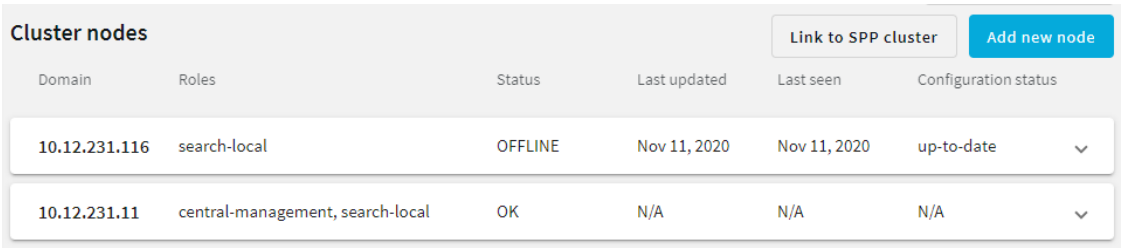
   - Ensure that each node has a search role.
   - Ensure that each node has only one search role.
   - You must assign the Search master role before you can assign Search minion roles.

Figure 115: Basic Settings > Cluster management — Assigning the Search local role to the selected node



4. Click **Update** to apply the selected roles. The role you assigned (in this case, the **Search local** role) is then displayed next to the node, under the **Roles** column.

Figure 116: Basic Settings > Cluster management — Search local role is assigned to node



You can assign roles to your nodes through the REST API, too. For more information, see *Assign a role to a node* in the *REST API Reference Guide*.

Configuration synchronization across nodes in a cluster

Nodes fetch their configuration from the Central management node, and merge it into their own configuration. Depending on their role, nodes may merge the whole configuration into their own (Managed host nodes), or only the cluster-specific parts (nodes with no roles assigned). Whenever a configuration change is made on the Central management node and the change is committed, it is synchronized to all nodes in the cluster as soon as the nodes fetch the latest configuration from the Central management node.

Configuration synchronization in One Identity Safeguard for Privileged Sessions (SPS) has some implications for the SSH keys (if any) that have been recorded on your nodes before they were joined to the cluster. For details, see Configuration synchronization and SSH keys on page 411.

In some cases, you may want to keep certain parts of the configuration on your nodes outside the scope of configuration synchronization. In that case, use a configuration synchronization plugin. For more information, see Using a configuration synchronization plugin on page 412.

The following configuration settings are never overwritten by configuration synchronization, even when not using a configuration synchronization plugin:

- Settings related to networking (**Basic Settings** > **Network**).
- Settings related to local services (**Basic Setting**s > **Local Services**).
- Settings related to the management of SPS (**Basic Settings** > **Management**).

For more information, see the following resources:

- For more information about configuration synchronization, see Configuration synchronization across nodes in a cluster on page 411 and Managing a cluster with configuration synchronization without central search.
- For more information about central search, see Searching session data on a central node in a cluster on page 858 and Managing a cluster with central search configuration and configuration synchronization.

Configuration synchronization and SSH keys

The only SSH keys present on Managed host nodes will always be the ones that have been recorded by the Central management node. This is because the SSH keys stored on the Central management node get synced to the Managed host nodes during configuration synchronization. This means that the SSH keys recorded on the Managed host nodes before they were joined to the cluster are overwritten by the keys stored on the Central management node.

The Central management node records new SSH keys in the following cases:

- The Central management node is configured to **Accept key for the first time**, and a new key is automatically recorded when the Central management node interacts with a server for the first time.
- A new key is recorded on the Central management node on the **Traffic Controls** > **SSH** > **Server Host Keys**page and this change is committed.

These are the keys that get synced to your Managed host nodes.

Using a configuration synchronization plugin

When synchronizing the central configuration across nodes, you may want to:

- Keep certain parts in the configuration of individual nodes unchanged.
- Customize certain parts of the central configuration to specific needs of individual nodes in the cluster (for example, your nodes may access external services through different network addresses).

You can achieve all of these by using a configuration synchronization plugin that contains transformations for the problematic parts. The plugin only runs on nodes that have the Managed host role.

Customizing certain parts or features of a node using a configuration synchronization plugin has the same limitations as configuring One Identity Safeguard for Privileged Sessions (SPS) through the REST API. In other words, whatever you can configure through the REST API, you can configure the exact same settings using the plugin. One notable difference between the REST API and the plugin is that using the REST API, you can only read certain types of data (such as keys and passwords), while using the configuration synchronization plugin, you can write these types of data as well.

For details on how to configure SPS using the REST API, see *REST API Reference Guide*.

Data structures in the plugin are represented as nested JSON objects. For object references, the plugin uses keys.

The plugin works with the following key parameters:

- `local_config`: The current configuration of a Managed Host node (those parts that can be configured through the REST API).
- `merged_config`: The configuration of the Central management node that is about to be synced to the Managed host node (those parts that can be configured through the REST API), with settings related to networking, local services, and management whitelisted. These settings are never overwritten by configuration synchronization.
- `node_id`: The unique ID of the Managed host node in the cluster (you can retrieve this identifier by querying the `/api/cluster/nodes` endpoint through the REST API).
- `plugin_config`: The configuration of the plugin provided as free-form text. Specifying the configuration of the plugin is optional. It enables you to run configuration synchronization on each cluster with different parameters if you have multiple clusters.

Example: Customizing an IP address in a connection policy

For example, an RDP connection policy on a Managed host node specifies the following client and target addresses:

```
$ curl ... https://<url-of-Central-Management-
node>/api/configuration/rdp/connections/<id-of-the-connection-policy>
```

```
{
       "body": {
              "network": {
                    "clients": [
                          "0.0.0.0/0"
                          ],
                    "ports": [
                          3389
                          ],
                "targets": [
                      "10.30.255.28/24"
                      ]
              },
       },
       ...
}
```

In the following example, an RDP connection policy is configured with the following details on the Central management node:

```
$ curl ... https://<url-of-Managed-
Node>/api/configuration/rdp/connections/<id-of-the-connection-policy>
```

```
{
       "body": {
              "network": {
                    "clients": [
                          "0.0.0.0/0"
                          ],
                    "ports": [
                          3389
                          ],
                "targets": [
                      "10.30.255.8/24"
                      ]
              },
       },
       ...
}
```

To ensure that the details of the connection policy on the Managed host node are kept as-is after configuration synchronization, add the following lines to the plugin `main.py` file:

```
$ cat main.py
def merge(local_config: dict, merged_config: dict, node_id: str, plugin_
config: str, **kwargs):
    merged_config['rdp']['connections'][<id-of-the-connection-policy>]
['network']['targets'][0] = "10.30.255.8/24"
    return merged_config
```

Due to possible new (as yet undefined) parameters, it is good practice to close the parameter list of the merge function with `**kwargs`.

If you need assistance with writing customized transformations, contact our Professional Services Team, and a One Identity Service Delivery Engineer will help you.

NOTE: Configuration settings related to networking (**Basic Settings** > **Network**) and local services (**Basic Settings** > **Local Services**), with the exception of Safeguard for Privileged Analytics, are not overwritten on the nodes by configuration synchronization even if you are not using a plugin.

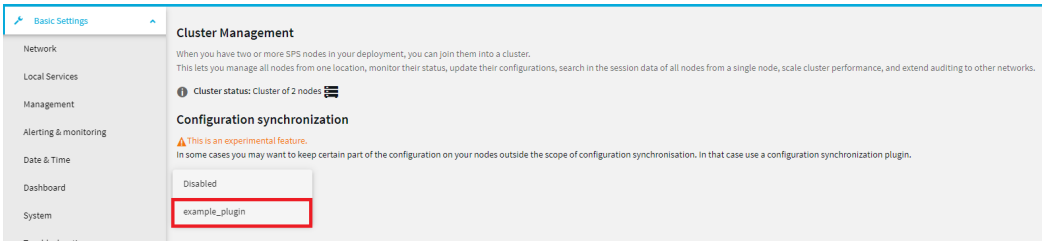
For the management of SPS (**Basic Settings** > **Management**), the following configuration settings are not overwritten:

- Email settings
- SSL certificates

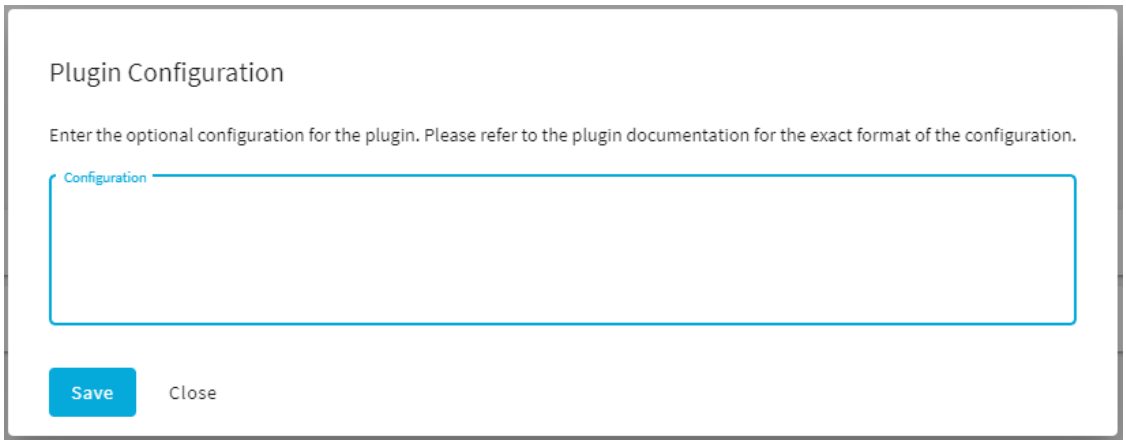
To use a configuration synchronization plugin

1. Upload a configuration synchronization plugin:
    a. Navigate to **Basic Settings** > **Plugins**.
    b. Browse for the file, and click **Upload plugin**.
    NOTE: It is not possible to upload or delete plugins if SPS is in Sealed mode.
2. Enable the plugin:
    a. Navigate to **Basic Settings** > **Cluster management** > **Configuration synchronization**.
    b. Select the plugin you have uploaded.

Figure 117: Basic Settings > Cluster management — Select configuration synchronization plugin



3. (Optional) Once you selected the plugin, to open the **Plugin Configuration** dialog, click **Configure** next to the plugin drop-down box.



   In the **Configuration** field, enter the configuration of the plugin. If you have multiple clusters, specifying the configuration of the plugin enables you to run configuration synchronization on each cluster with different parameters. To save your changes and return to the **Cluster management** window, click **Save**.

4. To save your changes, click **Commit changes**.

You can also upload and enable the configuration synchronization plugin through the REST API. For more information, see *Upload and enable a configuration synchronization plugin* in the *REST API Reference Guide*.

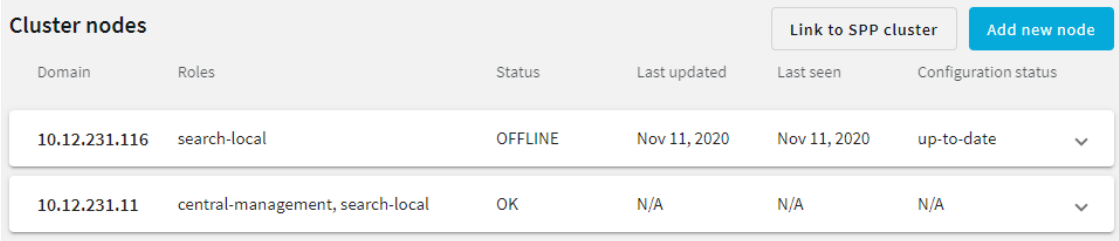
Monitoring the status of nodes in your cluster

This section describes how to monitor the status of nodes in your cluster.

To monitor the status of nodes in your cluster

1. On the web interface of your Central management node, navigate to **Basic Settings > Cluster management**. This page displays all nodes of the cluster.

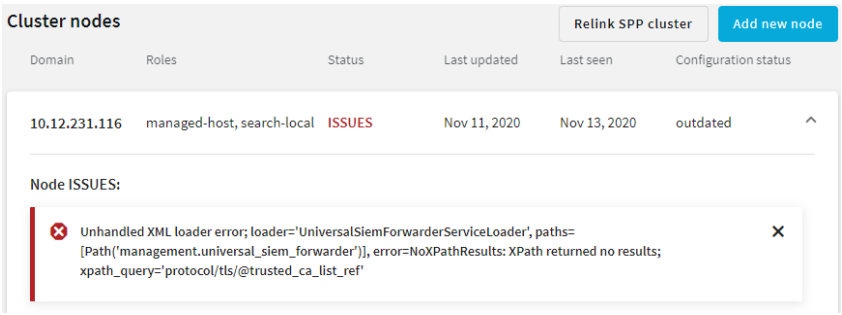
Figure 118: Basic Settings > Cluster management — Monitoring the status of the cluster nodes



The following status information is displayed for each node:

- **Status**: Indicates whether any issues occurred during configuration synchronization. It has the following values:

  Table 6: Basic Settings > Cluster management — Status values in the Cluster nodes screen

Value	Description
OK	Configuration synchronization was successful, no issues occurred.
ISSUES	While synchronizing the configuration, some issue(s) occurred. If this happens, click **ISSUES** in the **Status** column to reveal more information about the detected problem. 
OFFLINE	Status information was sent by the node longer than 60 seconds ago.

- **Last updated**: Indicates the last time the configuration of the node was synchronized, in ISO 8601 format.
- **Last seen**: Indicates the last time the node sent status information to the Central Management node, in ISO 8601 format.
- **Configuration status**: Indicates the status of configuration synchronization. It has the following values:

Table 7: Basic Settings > Cluster management — Configuration status values in the Cluster nodes screen

Value	Description
UP-TO-DATE	The node has fetched the latest configuration from the Central management node, and has applied it. It is in sync with the Central management node.
PENDING	There has been a configuration change on the Central management node, but the change has not yet been synchronized to the node.
OUTDATED	There has been some error on the node, therefore it is running an old configuration.
NOT FETCHED	The node has not fetched any configuration yet.
N/A	The node is the Central management node, so it is not fetching its configuration from any other node.

You can monitor the status of your nodes through the REST API, too. For more information, see *Query the status of all nodes in the cluster* in the *REST API Reference Guide* and *Query one particular node* in the *REST API Reference Guide*.

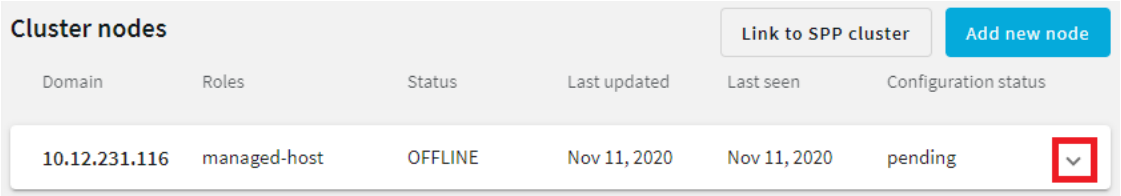
Updating the IP address of a node in a cluster

If the node that you joined to the cluster is a Managed host node, you can still change its IP address even after the join.

NOTE: This is not available for nodes that do not have the Managed host role assigned to them.

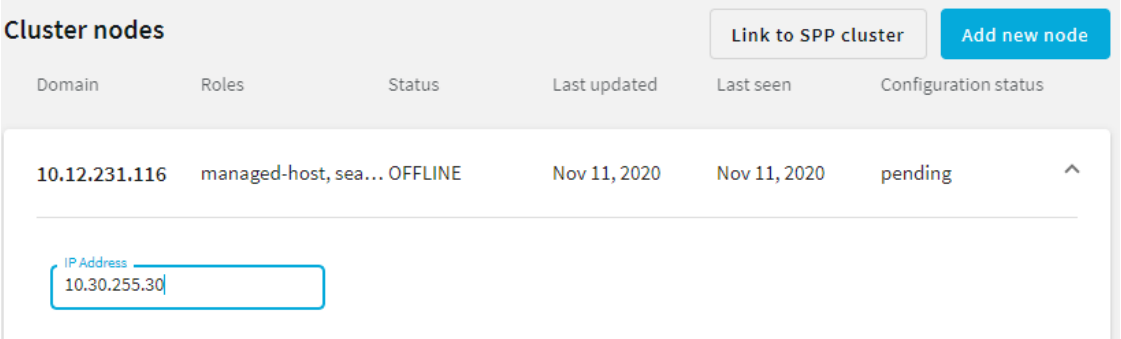
To update the IP address of a Managed host node that is already the member of a cluster

1. On the web interface of your Central management node, navigate to **Basic Settings > Cluster management**.
2. Click ⌄ on the row of the node that node that you want to update.



The node row is expanded, showing the node address and the available roles.

Figure 119: Basic Settings > Cluster management — Update IP address of node



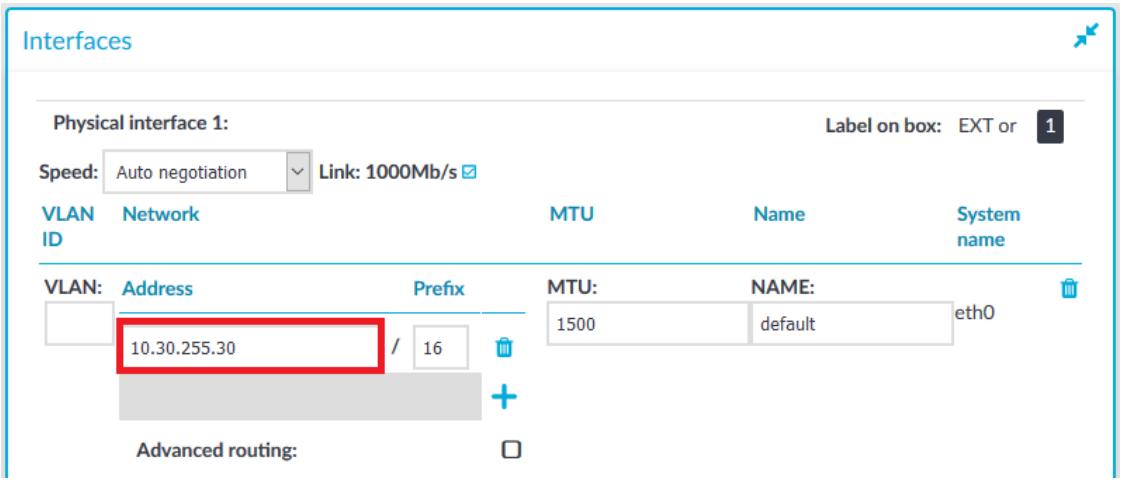
3. In the **IP Address** field, update the IP address of the node.

   CAUTION:

   Ensure that you are making the change for the Managed Host node. Do not change the IP address of the Central Management node.

4. Click **Update**.
5. On the web interface of the node with the IP address to update, navigate to **Basic Settings** > **Network** > **Interfaces**.
6. In the **Address** field, update the IP address of the node.

Figure 120: Basic Settings > Network > Interfaces — Updating the IP address of your node



7. Click Commit.

Managing a cluster with configuration synchronization without central search

You can configure your One Identity Safeguard for Privileged Sessions (SPS) cluster in the following ways:

- **Configuration synchronization without a central search**: This method allows you to perform your configuration settings on your Central management node. Managed host nodes periodically fetch and merge the settings into their own: this is called "configuration synchronization". Central search is not configured in this method, so you can search for sessions on each node, including the Central management node.

  For more information on this method, see Configuration synchronization without a central search.

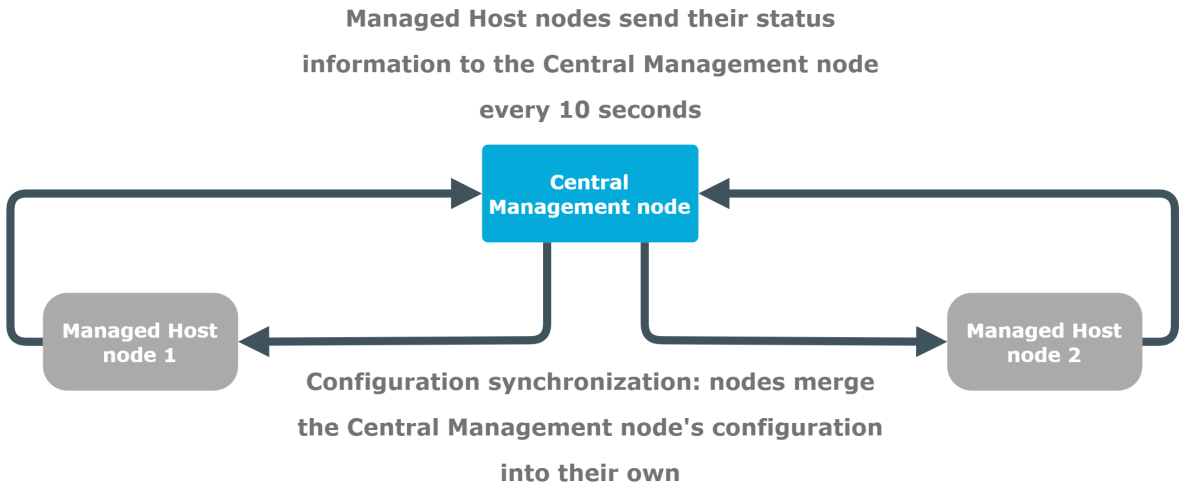
- **Central search with configuration synchronization**: This method allows you to use a Central management node with a Search master role to view session data recorded by the minion nodes of your cluster, as well as manage all the nodes in the cluster from one central location.

  For more information on this method, see Central search with configuration synchronization.

  > IMPORTANT: One Identity does not recommend having a central search configuration without configuration synchronization.

The following figure shows a cluster with configuration synchronization without central search.

Figure 121: Configuration synchronization without central search



The figure above is an example of an SPS cluster configured as follows:

- There is a Central management node.
- There are two Managed host nodes (Managed host node 1 and 2).
- The Central Management node is connected to the two Managed host nodes.
- The Managed host nodes fetch their configuration from the Central management node, and merge it into their own configuration.
- The Managed host nodes send their status information to the Central management node every 10 seconds.

The Central management node and the connected Managed host nodes require different configuration settings as described in the table below:

Table 8: Managing a configuration synchronization without a central search

Role	Use and configuration settings
Central management node	• Use it as a node with a central configuration, which is synchronized to the other nodes of the cluster. • Perform your configuration settings on this node. Managed host nodes periodically fetch and merge these configuration settings into their own (configuration synchronization). • For backup and archive, configure a backup and archive server on your minion node, as well as on your Central management node. • Ensure that you configure high availability (HA) for each node (for both your Central management node and the Managed host nodes). Also ensure that the Central management node has a system backup configured. • You can search for all the sessions recorded on this node.
Managed host node	• Use it to record sessions and send status information to the Central management node. • Do not perform configuration settings on the minion. These are overwritten during configuration synchronization. NOTE: All configuration settings that you make on the minions are overwritten during configuration synchronization except the node specific configuration. • Set external and internal indexers. • For backup and archive, configure a backup and archive server on your minion node, as well as on your Central management node. • Ensure that you configure high availability (HA) for each

Role	Use and configuration settings
	node (for both your Central management node and the Managed host nodes). Also ensure that the Central management node has a system backup configured. • You can search for all the sessions recorded on this node.

For more information on each role, see Cluster roles.

Managing a cluster with central search configuration and configuration synchronization

You can configure your One Identity Safeguard for Privileged Sessions (SPS) cluster in the following ways:

- **Configuration synchronization without a central search**: This method allows you to perform your configuration settings on your Central management node. Managed host nodes periodically fetch and merge the settings into their own: this is called "configuration synchronization". Central search is not configured in this method, so you can search for sessions on each node, including the Central management node.

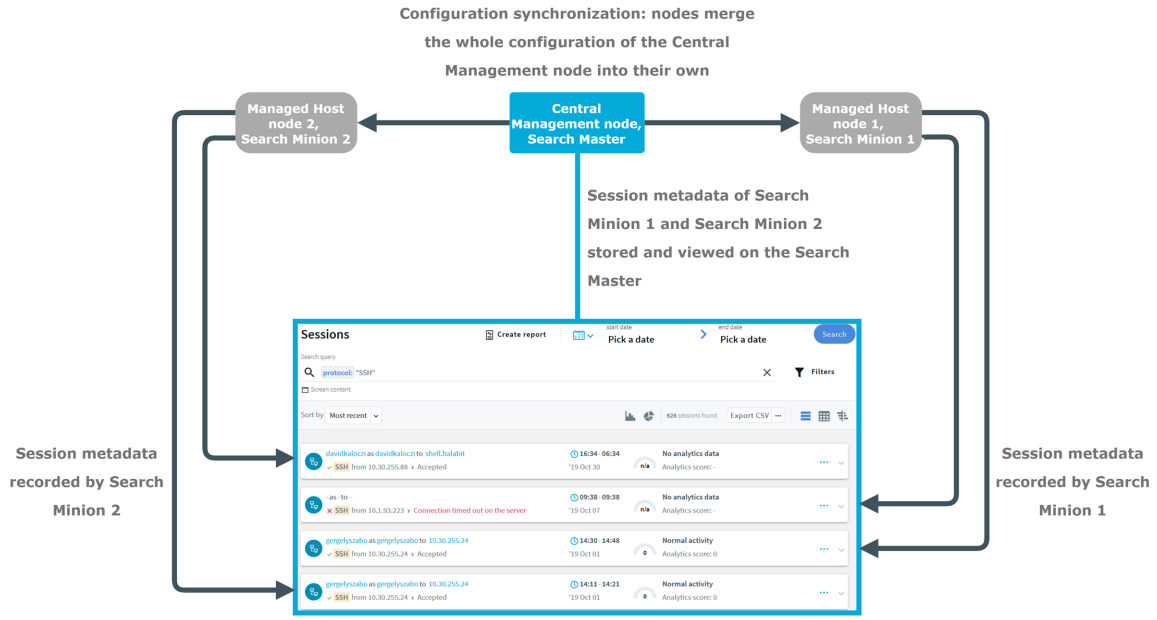
  For more information on this method, see Configuration synchronization without a central search.
- **Central search with configuration synchronization**: This method allows you to use a Central management node with a Search master role to view session data recorded by the minion nodes of your cluster, as well as manage all the nodes in the cluster from one central location.

  For more information on this method, see Central search with configuration synchronization.

  > IMPORTANT: One Identity does not recommend having a central search configuration without configuration synchronization.

The following figure shows a cluster configured for central search with configuration synchronization.

Figure 122: Central search with configuration synchronization



The figure above is an example of an SPS cluster configured as follows:

- There is a Central management node, which has a Search master role.
- There are two Managed host nodes (Managed host node 1 and 2), each configured with a Search minion role.
- The Central management node is connected to the two minion nodes.
- The minion nodes record sessions, which are displayed on the Search interface of the Central management node.
- The minion nodes fetch their configuration from the Central management node, and merge it into their own configuration.

The Central management node with a Search master role and the connected Managed host nodes with Search minion roles require different configuration settings as described in the table below:

Table 9: Managing a central search configuration

Role	Use and configuration settings
Central management node, Search master	• Use it for viewing session data recorded by minions as well as managing all the nodes in the cluster. • Perform your configuration settings on this node. Managed host nodes periodically fetch and merge these configuration settings into their own (configuration synchronization). • For backup and archive, configure a backup server on

Role	Use and configuration settings
	your Central management node and an archive server on your minion node. • Ensure that you configure high availability (HA) for each node (for both your Central management node and the Managed host nodes). Also ensure that the Central management node has a system backup configured. • This node cannot be used to record sessions.
Managed host node, Search minion	• Use it to record sessions and store audit trail files. • Do not perform configuration settings on the minion. These are overwritten during configuration synchronization. NOTE: All configuration settings that you make on the minions are overwritten during configuration synchronization except the node specific configuration. • Set external and internal indexers. • For backup and archive, configure an archive server on your minion node, and a backup server on your Central management node. • Ensure that you configure high availability (HA) for each node (for both your Central management node and the Managed host nodes). Also ensure that the Central management node has a system backup configured.

For more information on each role, see Cluster roles.

Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster

The goal of HA nodes is to support enterprise business continuity by providing failover.

In High Availability (HA) mode, two One Identity Safeguard for Privileged Sessions (SPS) units with identical configurations are operating simultaneously. These two units are the primary node and the secondary node (previously also referred to as the master node and the slave node). The primary node shares all data with the secondary node, and if the primary node stops functioning, the secondary node becomes active.

CAUTION: Depending on your environment, the shift between the two HA nodes might take 5 - 10 minutes. During this period, SPS connections are lost and they are not automatically restored. You might also encounter data loss during the shift.

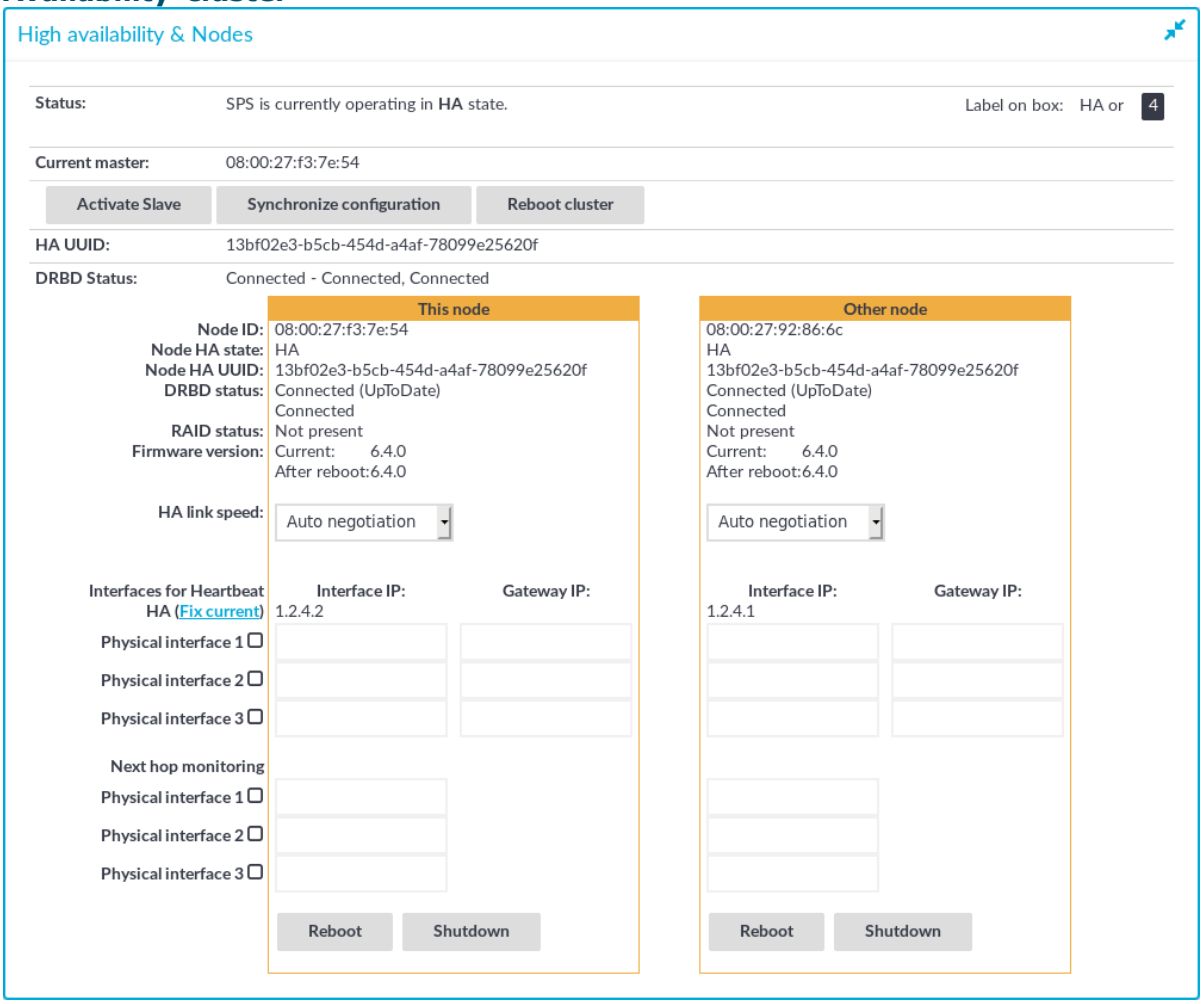
NOTE: To ensure the stability of the connection, One Identity recommends a direct physical connection between the nodes in the HA nodes. Gratuitous ARP requests are sent to inform hosts on the local network that the MAC addresses behind these IP addresses have changed.

The primary node shares all data with the secondary node using the HA network interface (labeled as 4 or HA on the SPS appliance). The disks of the primary and the secondary node must be synchronized for the HA support to operate correctly. Interrupting the connection between running nodes (unplugging the Ethernet cables, rebooting a switch or a router between the nodes, or disabling the HA interface) disables data synchronization and forces the secondary node to become active. This might result in data loss. You can find instructions to resolve such problems and recover a SPS cluster in Troubleshooting a One Identity Safeguard for Privileged Sessions (SPS) cluster on page 954.

NOTE: HA functionality was designed for physical SPS units. If SPS is used in a virtual environment, use the fallback functionalities provided by the virtualization service instead.

The **Basic Settings** > **High Availability** page provides information about the status of the HA nodes and its nodes.

Figure 123: Basic Settings > High Availability — Managing a High Availability cluster



The following information is available about the cluster:

- **Status**: Indicates whether the SPS nodes recognize each other properly and whether those are configured to operate in High Availability mode. For details, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.
- **Current master**: The MAC address of the High Availability interface (`4` or `HA`) of the primary node. This address is also printed on a label on the top cover of the SPS unit.
- **HA UUID**: A unique identifier of the HA nodes. Only available in High Availability mode.
- **DRBD status**: Indicates whether the SPS nodes recognize each other properly and whether those are configured to operate in High Availability mode. For details, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.
- **DRBD sync rate limit**: The maximum allowed synchronization speed between the primary and the secondary node. For details, see Adjusting the synchronization speed on page 427.

The active (that is, primary) SPS node is labeled as **This node**. This unit inspects the SSH traffic and provides the web interface. The SPS unit labeled as **Other node** is the secondary node that is activated if the primary node becomes unavailable.

The following information is available about each node:

- **Node ID**: The MAC address of the `HA` interface of the node. This address is also printed on a label on the top cover of the SPS unit.

  For SPS clusters, the IDs of both nodes are included in the internal log messages of SPS. Note that if the central log server is a syslog-ng server, the `keep-hostname` option should be enabled on the syslog-ng server.
- **Node HA state**: Indicates whether the SPS nodes recognize each other properly and whether those are configured to operate in High Availability mode. For details, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.
- **Node HA UUID**: A unique identifier of the cluster. Only available in High Availability mode.
- **DRBD status**: The status of data synchronization between the nodes. For details, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.
- **RAID status**: The status of the RAID device of the node. If it is not **Optimal**, there is a problem with the RAID device. For details, see Understanding One Identity Safeguard for Privileged Sessions (SPS) RAID status on page 963.
- **Firmware version**: Version number of the firmware.
- **HA link speed**: The maximum allowed speed between the primary node and the secondary node. The HA link's speed must exceed the **DRBD sync rate limit**, else the web UI might become unresponsive and data loss can occur.

- **Interfaces for Heartbeat**: Virtual interface used only to detect that the other node is still available. This interface is not used to synchronize data between the nodes (only heartbeat messages are transferred).

  You can find more information about configuring redundant heartbeat interfaces in Redundant heartbeat interfaces on page 428.
- **Next hop monitoring**: IP addresses (usually next hop routers) to continuously monitor both the primary node and the secondary node by using ICMP echo (ping) messages. If any of the monitored addresses becomes unreachable from the primary node while being reachable from the secondary node (in other words, more monitored addresses are accessible from the secondary node), then it is assumed that the primary node is unreachable and a forced takeover occurs – even if the primary node is otherwise functional. For details, see Next-hop router monitoring on page 430.

HA nodes configuration and management options

This section is about the available configuration and management options for HA nodess.

Setting up a High Availability cluster

For detailed instructions about setting up a HA nodes, see *Installing two SPS units in HA mode* in the *Installation Guide*.

Adjust the DRBD (primary-secondary) synchronization speed

You can change the limit of the DRBD synchronization rate. Note that this does not change the speed of normal data replication. For details, see Adjusting the synchronization speed on page 427.

Configure redundant heartbeat interfaces

You can configure virtual interfaces for each HA node to monitor the availability of the other node. For details, see Redundant heartbeat interfaces on page 428.

Configure next-hop monitoring

You can provide IP addresses (usually next hop routers) to continuously monitor both the primary node and the secondary node by using ICMP echo (ping) messages. If any of the monitored addresses becomes unreachable from the primary node while being reachable from the secondary node (in other words, more monitored addresses are accessible from the secondary node), then it is assumed that the primary node is unreachable and a forced takeover occurs – even if the primary node is otherwise functional. For details, see Next-hop router monitoring on page 430.

Reboot the HA nodes

To reboot both nodes, click **Reboot Cluster**. To prevent takeover, a token is placed on the secondary node. While this token persists, the secondary node halts its boot process to make sure that the primary node boots first. Following reboot, the primary node removes this token from the secondary node, allowing it to continue with the boot process.

If the token still persists on the secondary node following reboot, the **Unblock Slave Node** button is displayed. Clicking the button removes the token, and reboots the secondary node.

Reboot a node

This option reboots the selected node.

When rebooting the nodes of a cluster, reboot the other node (that is, the secondary node) first to avoid unnecessary takeovers.

Shutdown a node

This option forces the selected node to shut down.

When shutting down the nodes of a cluster, shut down the other node (that is, the secondary node) first. When powering on the nodes, start the primary node first to avoid unnecessary takeovers.

Manual takeover

To activate the other node (that is, the secondary node) and disable the currently active node, click **Activate slave**.

Activating the secondary node terminates all connections of One Identity Safeguard for Privileged Sessions (SPS) and might result in data loss. The secondary node becomes active after about 60 seconds, during which the protected servers cannot be accessed.

Adjusting the synchronization speed

One Identity Safeguard for Privileged Sessions (SPS) synchronizes the content of the hard disk of the primary node (previously also referred to as master node) and the secondary node (previously also referred to as slave node) in the following cases.

- When you configure two SPS units to operate in High Availability mode (converting a single node to a High Availability cluster),
- when you replace a node from a cluster, or
- when recovering from a split-brain situation.
- Normal data replication (copying incoming data, for example, audit trails from the primary node to the secondary node is not synchronization.

Since this synchronization can take up significant system-resources, the maximal speed of the synchronization is limited, by default, to `10 Mbps`. However, this means that

synchronizing large amount of data can take very long time, so it is useful to increase the synchronization speed in certain situations —.

To change the limit of the DRBD synchronization rate, navigate to **Basic Settings** > **High Availability** > **DRBD sync rate limit**, and select the desired value. Note the following points before changing the **DRBD sync rate limit** option.

- The **Basic Settings** > **High Availability** > **DRBD sync rate limit** option is visible only when synchronization is in progress, or when you have clicked **Convert to Cluster** but have not rebooted the cluster yet.
- Changing this option does not change the limit of the data replication speed.
- Set the sync rate carefully. A high value is not recommended if the load of SPS is very high, as increasing the resources used by the synchronization process may degrade the general performance of SPS. On the other hand, the HA link's speed must exceed the speed of the incoming data, else the web UI might become unresponsive and data loss can occur.

The **Basic Settings** > **High Availability** > **DRBD status** field indicates whether the latest data (including SPS configuration, audit trails, log files, and so on) is available on both SPS nodes. For a description of each possible status, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.

Redundant heartbeat interfaces

To avoid unnecessary takeovers and to minimize the chance of split-brain situations, you can configure additional heartbeat interfaces in One Identity Safeguard for Privileged Sessions (SPS). These interfaces are used only to detect that the other node is still available, they are not used to synchronize data between the nodes (only heartbeat messages are transferred). For example, if the main HA interface breaks down, or is accidentally unplugged and the nodes can still access each other on the redundant HA interface, no takeover occurs, but no data is synchronized to the secondary node until the main HA link is restored. Similarly, if connection on the redundant heartbeat interface is lost, but the main HA connection is available, no takeover occurs.

If a redundant heartbeat interface is configured, its status is displayed in the **Basic Settings** > **High Availability** > **Redundant Heartbeat status** field, and also in the **HA** > **Redundant** field of the System monitor. For a description of each possible status, see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.

The redundant heartbeat interface is a virtual interface with a virtual MAC address that uses an existing interface of SPS. The MAC address of the virtual redundant heartbeat interface is displayed as **HA MAC**. The MAC address of the redundant heartbeat interface is generated in a way that it cannot interfere with the MAC addresses of physical interfaces. Similarly, the HA traffic on the redundant heartbeat interface cannot interfere with any other traffic on the interface used.

If the nodes lose connection on the main HA interface, and after a time the connection is lost on the redundant heartbeat interfaces as well, the secondary node becomes active. However, as the primary node was active for a time when no data synchronization was possible between the nodes, this results in a split-brain situation, which must be resolved

before the HA functionality can be restored. For details, see Recovering from a split brain situation on page 957.

NOTE: Even if redundant HA links are configured, if the dedicated HA link fails, the secondary node will not be visible on the High Availability page anymore.

SPS nodes use UDP port 694 to send each other heartbeat signals.

To configure a redundant heartbeat interface

1. Navigate to **Basic Settings** > **High Availability** > **Interfaces for Heartbeat**.
2. Select the interface you want to use as redundant heartbeat interface (for example `Physical interface 1`). Using an interface as a redundant heartbeat interface does not affect the original traffic of the interface.

   Figure 124: Basic Settings > High Availability — Configuring redundant heartbeat interfaces

3. Enter an IP address into the **This node** > **Interface IP** field of the selected interface. Note the following:
   - The two nodes must have different **Interface IP**.
   - If you do not use next hop monitoring on the redundant interface, you can use any **Interface IP** (even if otherwise it does not exist on that network).
   - If you use next hop monitoring on the redundant interface, the **Interface IP** address must be a real IP address that is visible from the other node.
   - If you use next hop monitoring on the redundant interface, the **Interface IP** must be accessible from the next-hop address, and vice-versa. For details on next hop monitoring, see Next-hop router monitoring on page 430.

   Use an IPv4 address.
4. If the two nodes are in a different subnetwork, enter the IP address of the local gateway into the **This node** > **Gateway IP** field. The **Interface IP** address of the node must be accessible from the **Gateway IP** address.

   Use an IPv4 address.
5. Enter an IP address into the **Other node** > **Interface IP** field of the selected interface. Note the following:
   - The two nodes must have different **Interface IP**.
   - If you do not use next hop monitoring on the redundant interface, you can use any **Interface IP** (even if otherwise it does not exist on that network).

- If you use next hop monitoring on the redundant interface, the **Interface IP** address must be a real IP address that is visible from the other node.
- If you use next hop monitoring on the redundant interface, the **Interface IP** must be accessible from the next-hop address, and vice-versa. For details on next hop monitoring, see Next-hop router monitoring on page 430.

Use an IPv4 address.

6. If the two nodes are in a different subnetwork, enter the IP address of the local gateway into the **Other node** > **Gateway IP** field. The **Interface IP** address of the node must be accessible from the **Gateway IP** address.

   Use an IPv4 address.

7. Repeat the previous steps to add additional redundant heartbeat interfaces if needed.
8. Click 

.
9. Restart the nodes for the changes to take effect: click **Reboot Cluster**.

Next-hop router monitoring

By default, HA takeover occurs only if the primary node stops working or becomes unreachable from the secondary node. However, this does not cover the scenario when the primary node becomes unaccessible to the outside world while the secondary node would be still accessible (for example because it is connected to a different router).

To address such situations, you can specify IP addresses (usually next hop routers) to continuously monitor both the primary node and the secondary node by using ICMP echo (ping) messages. One such address can be set up for every interface.

When setting up next hop monitoring, you have to make sure that the primary and secondary nodes can ping the specified address directly. You can either:

- Choose the addresses of the redundant-HA One Identity Safeguard for Privileged Sessions (SPS) interfaces so that they are on the same subnet as the next-hop address
- Configure the next-hop device with an additional IP-address that is on the same subnet as the redundant-HA SPS interfaces facing it

If any of the monitored addresses becomes unreachable from the primary node while being reachable from the secondary node (in other words, more monitored addresses are accessible from the secondary node) then it is assumed that the primary node is unreachable and a forced takeover occurs — even if the primary node is otherwise functional.

Naturally, if the secondary node is not capable of taking over the primary node (for example, because there is data not yet synchronized from the current primary node), no takeover is performed.

Before configuring next-hop monitoring

To use next-hop router monitoring, you must configure the switch ports connecting to SPS. Follow one of the following methods to configure the switch ports:

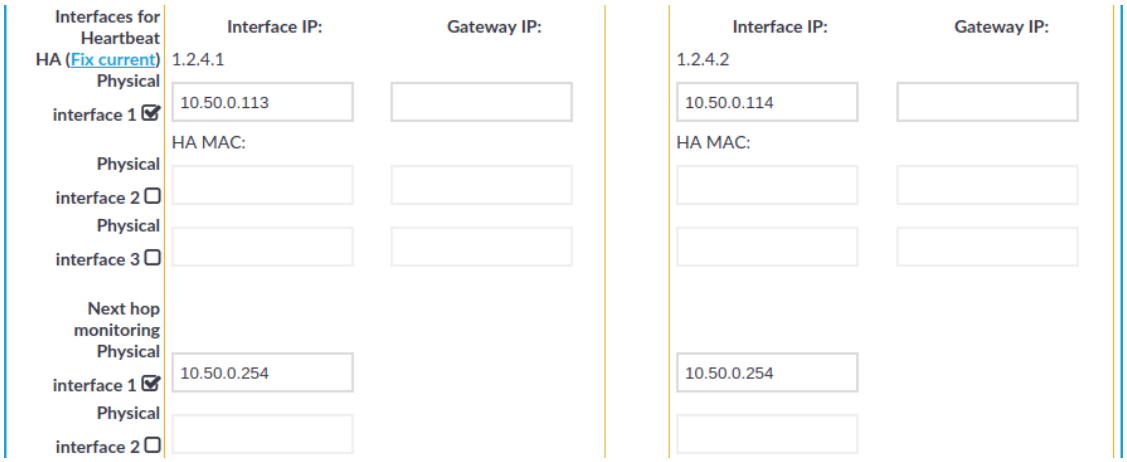
- Disable spanning tree protocol (STP).
- Enable portfast in case of a Cisco switch.
- Apply another switch vendor specific solution to mitigate any delay caused by STP.

CAUTION: Make sure that you always configure the switch ports before configuring next-hop monitoring. Using STP on SPS connected switch ports can cause an infinite HA failover loop which renders SPS unreachable.

To configure next-hop monitoring

1. Navigate to **Basic Settings** > **High Availability** > **Next hop monitoring**.
2. Select the interface to use for monitoring its next-hop router.

Figure 125: Basic Settings > High Availability — Configuring next hop monitoring



3. Enter the IP address to monitor from the current primary node (for example, the IP address of the router or the switch connected to the interface) into the **This node** > **Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.

   Use an IPv4 address.
4. Enter the IP address to monitor from the current secondary node (for example the IP address of the router or the switch connected to the interface) into the **Other node** > **Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.

   Use an IPv4 address.

5. Repeat the previous steps to add IP addresses to be monitored from the other interfaces if needed.
6. Click 

.

CAUTION:

For the changes to take effect, you have to restart both nodes. To restart both nodes, click **Reboot Cluster**.

Upgrading One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) appliances are preinstalled with a Long Term Support (LTS) release. One Identity recommends that you upgrade to the latest LTS maintenance release as soon as possible. Each LTS release is supported for 3 years after original publication date, and for 1 year after the succeeding LTS release is published (whichever date is later). You are encouraged to upgrade to succeeding LTS releases.

Feature Releases provide additional features which are not yet consolidated to an LTS release. To gain access to these features, you may install a supported Feature Release on the appliance, with the following conditions:

- You cannot roll back to an LTS release from a Feature Release.
- Feature Releases are released and supported in a timeline of 6 months. You have to keep upgrading SPS to the latest Feature Release to ensure that your appliance is supported.

For both LTS and Feature Releases, One Identity regularly incorporates security patches and bugfixes, and issues updated Revisions of the released product. One Identity strongly recommends always installing the latest Revision of the used software Release.

CAUTION:

Downgrading from the latest Feature Release, even to an LTS release, voids support for SPS.

The following sections describe how to keep SPS up to date, and how to install a new license:

- Prerequisites: Upgrade checklist on page 433.
- Upgrading a single node: Upgrading One Identity Safeguard for Privileged Sessions (SPS) (single node) on page 434.
- Upgrading a High Availability cluster: Upgrading High Availability One Identity Safeguard for Privileged Sessions (SPS) nodes on page 435.
- Troubleshooting: Troubleshooting on page 437.

- Renewing the SPS license: Updating the SPS license on page 441.
- Exporting the configuration of SPS: Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 438.
- Importing the configuration of SPS: Importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 439.

Upgrade checklist

The following list applies to all configurations:

- You have created a configuration backup of One Identity Safeguard for Privileged Sessions (SPS).

  For detailed instructions, refer to Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 438.
- You have a valid support portal account.

  To download the required firmware file and license, you need a valid support portal account. Consider that registration is not automatic, and might require up to two working days to process.
- You have downloaded the latest SPS firmware from the Downloads page.
- You have read the Release Notes of the firmware before updating. The Release Notes might include additional instructions specific to the firmware version.

  The Release Notes are available on the Downloads page.
- You have verified that SPS is in good condition (no issues are displayed on the System Monitor).
- Optional: You have exported core dump files, if necessary for debugging, from **Basic Settings** > **Troubleshooting** > **Core files**. These files are removed during upgrade.

If you have a High Availability cluster:

- You have IPMI access to the secondary node. You can find detailed information on using the IPMI in the following documents:

  For Safeguard Sessions Appliance 3000 and 3500, see the X9 SMT IPMI User's Guide.

  For Safeguard Sessions Appliance 4000, see the X12 H12 BMC User's Manual.
- You have verified on the **Basic Settings** > **High Availability** page that the HA status is not degraded.

If you are upgrading SPS in a virtual environment:

- You have created a snapshot of the virtual machine before starting the upgrade process.
- You have configured and enabled console redirection (if the virtual environment allows it).

During the upgrade, SPS displays information about the progress of the upgrade and any possible problems in the following places:

- On the web interface of SPS, at any of the **Listening addresses** configured at **Basic settings** > **Local Services** > **Web login (admin and user)**. (After booting, you are directed to the login screen of SPS.)

  NOTE: If you are upgrading to version 8.0.1 LTS from version 5.0.x, this feature is enabled after the first boot to version 8.0.1 LTS. So during the upgrade to version 8.0.1 LTS, you will not be able to see any upgrade logs on the web interface.

- On the console, which you can monitor with IPMI (ILOM) or console access.

The information displayed in the browser and on the console is the same.

One Identity strongly recommends that you test the upgrade process in a non-production (virtual, and so on) environment first.

Upgrading SPS requires a reboot. One Identity strongly recommends that you perform the upgrade on the production appliance during maintenance hours only, to avoid any potential data loss.

Upgrading One Identity Safeguard for Privileged Sessions (SPS) (single node)

This section describes how to upgrade One Identity Safeguard for Privileged Sessions (SPS) to a newer firmware version. One Identity recommends that you always use the latest maintenance release available.

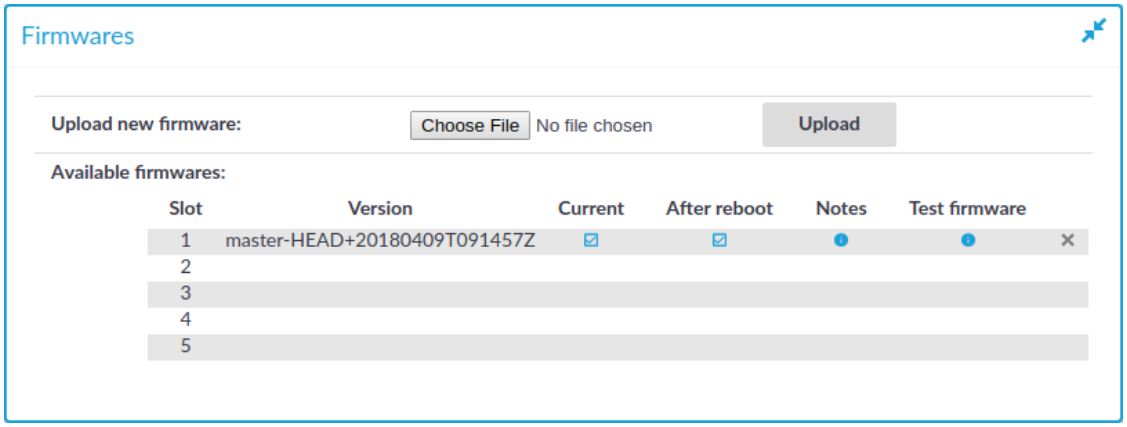
CAUTION:

When upgrading to a new major release (that is, to a new Feature Release or a new Long-Term Supported release), always follow the instructions of the *How to upgrade to One Identity Safeguard for Privileged Sessions* guide for that release, as it contains more detailed instructions (available at the Safeguard for Privileged Sessions Documentation page).

To upgrade SPS to a newer firmware version

1. Navigate to **Basic Settings** > **System** > **Firmwares**.

   Figure 126: Basic Settings > System > Firmwares — Managing the firmwares

2. Upload the new firmware: **Browse** for the firmware `.iso` file and then click **Upload**.

   To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.

3. Click **Test** for the new firmware to check if your configuration can be upgraded to version 8.0.1 LTS. If the test returns any errors, correct them before continuing the upgrade process. If you encounter any problems, contact our Support Team.

   CAUTION:

   Proceed only if the upgrade test is successful.

4. For the new firmware, select **After reboot**. This will activate the new firmware after reboot.

5. Navigate to **Basic Settings** > **System** > **System control** > **This node**, and choose **Reboot**.

   SPS attempts to boot with the new firmware. Wait for the process to complete.

6. Login to the SPS web interface to verify that the upgrade was successful.

   Navigate to **Basic Settings** > **System** > **Version details**, or check the system log for the version numbers SPS reports on boot. In case you encounter problems, you can find common troubleshooting steps in Troubleshooting on page 437.

Upgrading High Availability One Identity Safeguard for Privileged Sessions (SPS) nodes

The following describes how to upgrade a High Availability One Identity Safeguard for Privileged Sessions (SPS) nodes to a newer firmware version. One Identity recommends that you always use the latest maintenance release available.

CAUTION:

When upgrading to a new major release (that is, to a new Feature Release or a new Long-Term Supported release), always follow the instructions of the *How to upgrade to One Identity Safeguard for Privileged Sessions* guide for that release, as it contains more detailed instructions (available at the Safeguard for Privileged Sessions Documentation page).

To upgrade your High Availability SPS nodes

1. Navigate to **Basic Settings** > **System** > **Firmwares**.
2. Upload the new firmware: **Browse** for the firmware `.iso` file and then click **Upload**.

   To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.
3. Click **Test** for the new firmware to check if your configuration can be upgraded to version 8.0.1 LTS. If the test returns any errors, correct them before continuing the upgrade process. If you encounter any problems, contact our Support Team.

   CAUTION:

   Proceed only if the upgrade test is successful.
4. For the new firmware, select **After reboot**. This will activate the new firmware after reboot.
5. Navigate to **Basic Settings** > **High Availability**, and verify that the new firmware is active on the secondary node. This might take a few minutes.
6. In **Basic Settings** > **High Availability** > **Other node**, click **Shutdown**.
7. Restart the primary node: click **This node** > **Reboot**.

   SPS attempts to boot with the new firmware. Wait for the process to complete.
8. Login to the SPS web interface to verify that the primary node upgrade was successful.

   Navigate to **Basic Settings** > **System** > **Version details**, or check the system log for the version numbers SPS reports on boot. In case you encounter problems, you can find common troubleshooting steps in Troubleshooting on page 437.
9. Use the IPMI to power on the secondary node.

   The secondary node attempts to boot with the new firmware, and reconnects to the primary node to sync data. During the sync process, certain services (including Heartbeat) are not available. Wait for the process to finish, and the secondary node to boot fully.
10. Navigate to **Basic Settings** > **High Availability** and verify that the secondary node is connected, and has the same firmware versions as the primary node.
11. To finish the upgrade process and verify HA functionality, click **Activate Slave**.

Upgrading an SPS cluster

The following section describes how you can upgrade your One Identity Safeguard for Privileged Sessions (SPS). One Identity recommends that you always use the latest available release version.

Prerequisites

NOTE: One Identity strongly suggests that you reserve an adequate uninterrupted maintenance window to have sufficient time to upgrade every node of the cluster.

CAUTION: One Identity does not support having different SPS versions in the cluster.

For details on the different cluster roles, see *Cluster roles* in the *Administration Guide*.

Upgrading an SPS cluster

NOTE: If you are upgrading from SPS version 8.x, follow these steps.

If you are upgrading from a 7.x version, see the documentation for the appropriate 7.x version in the *One Identity Safeguard for Privileged Sessions documentation*.

To upgrade an SPS cluster

1. Start your upgrade with the **Search Master** node (if you have one).

   CAUTION: The period until the **Search Master** node gets fully upgraded is critical. While the **Search Master** node is rebooting, the **Search Minion** nodes cannot connect to the **Search Master** node. If the **Search Master** node cannot successfully start up, the **Search Minion** nodes might fill up and cannot accept new connections.

2. When your **Search Master** node is fully upgraded, upgrade your other nodes, one after the other.

Troubleshooting

If you experience any strange behavior of the web interface, first try to reload the page by holding the SHIFT key while clicking the Reload button of your browser to remove any cached version of the page.

In the unlikely case that One Identity Safeguard for Privileged Sessions (SPS) encounters a problem during the upgrade process and cannot revert to its original state, SPS performs the following actions:

- Initializes the network interfaces using the already configured IP addresses.
- Enables SSH-access to SPS, unless SPS is running in sealed mode. That way it is possible to access the logs of the upgrade process that helps the One Identity

Support Team to diagnose and solve the problem. Note that SSH access will be enabled on every active interface, even if management access has not been enabled for the interface.

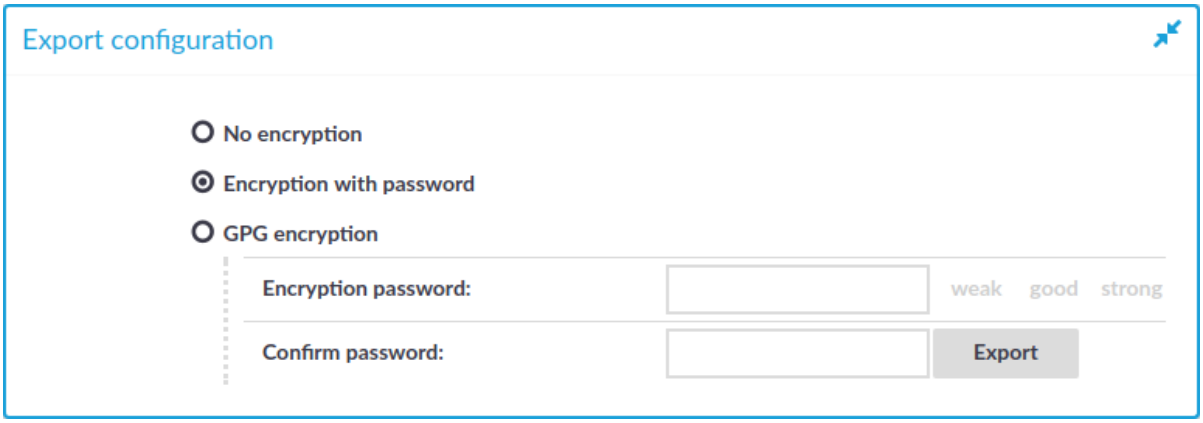
In case the web interface is not available within 30 minutes of rebooting SPS, check the information displayed on the local console and contact our Support Team.

Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS)

The configuration of One Identity Safeguard for Privileged Sessions (SPS) can be exported (for manual archiving, or to migrate it to another SPS unit) from the **Basic Settings** > **System** page. Use the respective action buttons to perform the desired operation.

You also have the option to export the configuration SPS into a local file using the console. For details, see Exporting and importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) using the console on page 447.

Figure 127: Basic Settings > System — Exporting the SPS configuration



To export the configuration of SPS

1. Navigate to **Basic Settings** > **System** > **Export configuration**.
2. Select how to encrypt the configuration:
   - To export the configuration file without encryption, select **No encryption**.

     CAUTION:

     One Identity does not recommend exporting the SPS configuration without encryption, as it contains sensitive information such as password hashes and private keys.
   - To encrypt the configuration file with a simple password, select **Encrypt with password** and enter the password into the **Encryption password** and

Confirm password fields.

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

- To encrypt the configuration file with GPG, select **GPG encryption**. Note that this option uses the same GPG key that is used to encrypt automatic system backups, and is only available if you have uploaded the public part of a GPG key to SPS at **Basic Settings** > **Management** > **System backup**. For details, see Encrypting configuration backups with GPG.

3. Click **Export**.

NOTE: The exported file is a gzip-compressed archive. On Windows platforms, it can be decompressed with common archive managers such as the free 7-Zip tool.

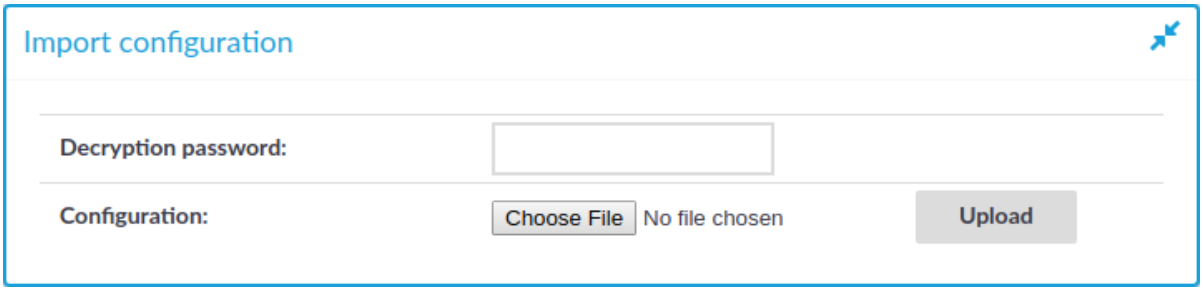
The name of the exported file is `<hostname_of_SPS>-YYYYMMDDTHHMM.config`, the `-encrypted` or `-gpg` suffix is added for password-encrypted and GPG-encrypted files, respectively.

Importing the configuration of One Identity Safeguard for Privileged Sessions (SPS)

The configuration of One Identity Safeguard for Privileged Sessions (SPS) can be imported from the **Basic Settings** > **System** page. Use the respective action buttons to perform the desired operation.

You also have the option to import configuration of SPS from a local file using the console. For details, see Exporting and importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) using the console on page 447.

Figure 128: Basic Settings > System — Importing the SPS configuration



CAUTION:

It is not possible to import the configuration of an older major release (for example, 1.0) into a newer release (for example, 2.0).

To import the configuration of SPS

1. **CAUTION:**

   Do not export or import configuration between a physical SPS deployment and a virtual one. Because of the differences and limitations between physical and virtual appliances, configure the virtual appliance from scratch to ensure proper functionality. When you migrate a virtual SPS to another one, you can export and import the configuration.

   Navigate to **Basic Settings** > **System** > **Import configuration**.
2. Click **Browse** and select the configuration file to import.
3. Enter the password into the **Encryption password** field and click **Upload**.

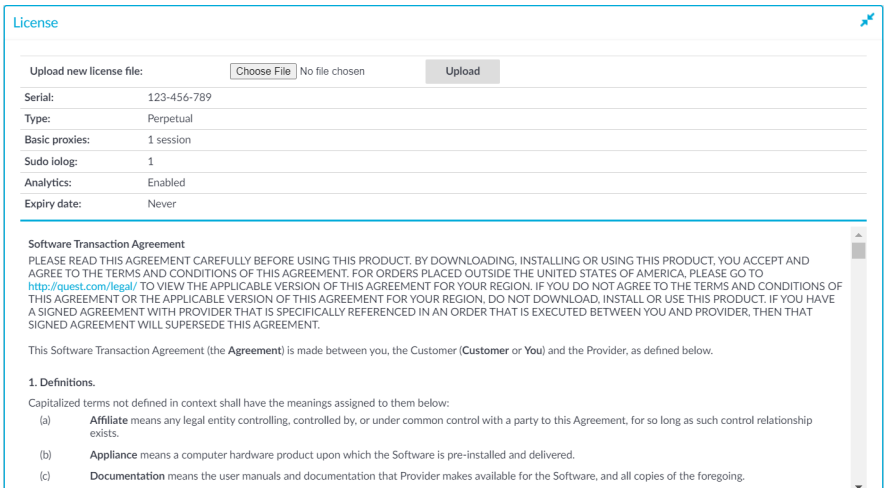
   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

Managing the One Identity Safeguard for Privileged Sessions (SPS) license

The **Basic Settings** > **System** > **License** page displays the following information of the current license:

Figure 129: Basic Settings > System > License — Updating the license



- **Serial**: The unique serial number of the license.
- **Type**: The license type, for example: `Perpetual`.
- **Basic proxies**: The actual value of the sessions or hosts. You can select the following sub-options:
    - **Host**: Limits the number of servers (individual IP addresses) that can be connected through SPS.
    - **Session**: Limits the number of concurrent sessions (parallel connections) that can pass through SPS at a time (for example `25`).
- **Sudo iolog**: Number of enabled Sudo iolog connections.

  For more information, see Using Sudo with SPS.
- **Analytics**: Enables the **Analytics** option.

  For more information, see Analyzing data using One Identity Safeguard for Privileged Analytics.
- **Expiry Date**: The license expiration date. The date is displayed in *YYYY-MM-DD* format.

SPS starts sending automatic alerts daily, `60` days before the license expires.

Updating the SPS license

One Identity recommends that you update the SPS license before the existing license expires or when you purchase a new license.

CAUTION:

From SPS 8.0, SPS requires a new license. To avoid possible downtimes due to certain features not being available, before starting the upgrade, ensure that you have a valid SPS license for 8.0.

Upgrade as follows

1. Perform the upgrade to 8.0 with your current license.
2. Update your SPS license to 8.0.

For a new SPS license for 8.0, contact our Licensing Team.

CAUTION:

Before uploading a new license, One Identity recommends that you back up the configuration of SPS. For more information, see Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 438.

To update the license

1. Navigate to **Basic Settings** > **System** > **License**.
2. Click **Browse** and select the new license file.
3. Click **Upload**, then **Commit**.

Accessing the One Identity Safeguard for Privileged Sessions (SPS) console

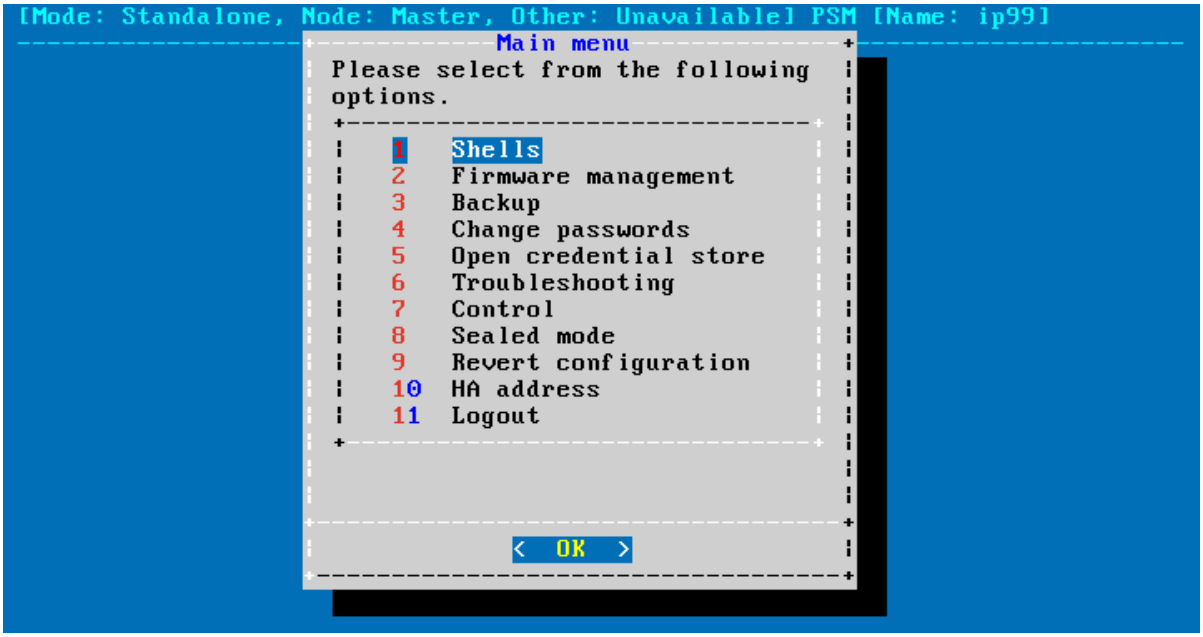
The following topics describe how to use the console menu of One Identity Safeguard for Privileged Sessions (SPS), how to enable remote SSH access to SPS, and how to change the root password from the web interface.

Using the console menu of One Identity Safeguard for Privileged Sessions (SPS)

Connecting to One Identity Safeguard for Privileged Sessions (SPS) locally or remotely using Secure Shell (SSH) allows you to access the console menu of SPS. The console menu provides access to the most basic configuration and management settings of SPS. It is mainly used for troubleshooting purposes, the primary interface of SPS is the web interface.

The console menu is accessible to the `root` user using the password set during completing the Welcome Wizard.

Figure 130: The console menu



The console menu provides allows you to perform the following actions.

Access the local core and boot shells.

This is usually not recommended and only required in certain troubleshooting situations. Select the boot/core shell's keyboard layout for the local console. This will not affect the

keyboard layout if you have accessed the shell via SSH.

The boot firmware boots up SPS, provides High Availability support, and starts the core firmware. The core firmware, in turn, handles everything else: provides the web interface, manages the connections, and so on.

Select the active firmware, and delete unneeded firmwares.

Accessing the firmware management is useful if after an update the new firmware does not operate properly and the web interface is not available to activate the previous firmware.

Start backup processes.

For more information about backup processes, see Data and configuration backups.

Change the passwords of the root and admin users.

For details, see Changing the root password of One Identity Safeguard for Privileged Sessions (SPS).

Access the network-troubleshooting functions and display the available log files.

If the web interface is inaccessible, it can be the result of an internal locking error. To resolve this issue, delete the lock files. After deletion, they are archived, and included in the support bundle if they are not older than 30 days. To create a support bundle, if the web interface is inaccessible, select **Create support bundle**.

NOTE: If deleting the lock files did not resolve the issue, contact our Support Team.

Reboot and shutdown the system.

For details, see Controlling One Identity Safeguard for Privileged Sessions (SPS): reboot, shutdown.

Enable and disable sealed mode.

For details, see Sealed mode on page 449.

Set the IP address of the HA interface.

For more information about assigning an IP address to the HA interface of a node, see Resolving an IP conflict between cluster nodes.

NOTE: Note that logging in to the console menu automatically locks the SPS interface, meaning that users cannot access the web interface while the console menu is used. The console menu can be accessed only if there are no users accessing the web interface. The connection of web-interface users can be terminated to force access to the console menu.

Enabling SSH access to the One Identity Safeguard for Privileged Sessions (SPS) host

Exclusively for troubleshooting purposes, you can access the One Identity Safeguard for Privileged Sessions (SPS) host using SSH.

Completing the Welcome Wizard automatically disables SSH access to SPS. Re-enabling it allows you to connect remotely to the SPS host and login using the `root` user. The password of the root user is the one you provided in the Welcome Wizard. For details on how to change the root password from the web interface, see Changing the root password of One Identity Safeguard for Privileged Sessions (SPS) on page 446.

CAUTION:

Accessing the One Identity Safeguard for Privileged Sessions (SPS) host directly using SSH is not recommended or supported, except for troubleshooting purposes. In such case, the One Identity Support Team will give you exact instructions on what to do to solve the problem.

For security reasons, disable SSH access to SPS when it is not needed. For details, see *Enabling SSH access to the One Identity Safeguard for Privileged Sessions (SPS) host* in the *Administration Guide*.

The following encryption algorithms are configured on the local SSH service of SPS:

- Key exchange (KEX) algorithms:

  ```
  diffie-hellman-group-exchange-sha256
  ```

- Ciphers:

  ```
  aes256-ctr,aes128-ctr
  ```

- Message authentication codes:

  ```
  hmac-sha2-512,hmac-sha2-256
  ```

NOTE: From SPS version 7.4, the local SSH console of SPS does not support signature algorithms based on SHA-1. The stronger RFC8332 RSA/SHA-256/512 algorithms are transparently used instead if client support is present (like in OpenSSH versions since 7.2 and in PuTTY versions since 0.75).

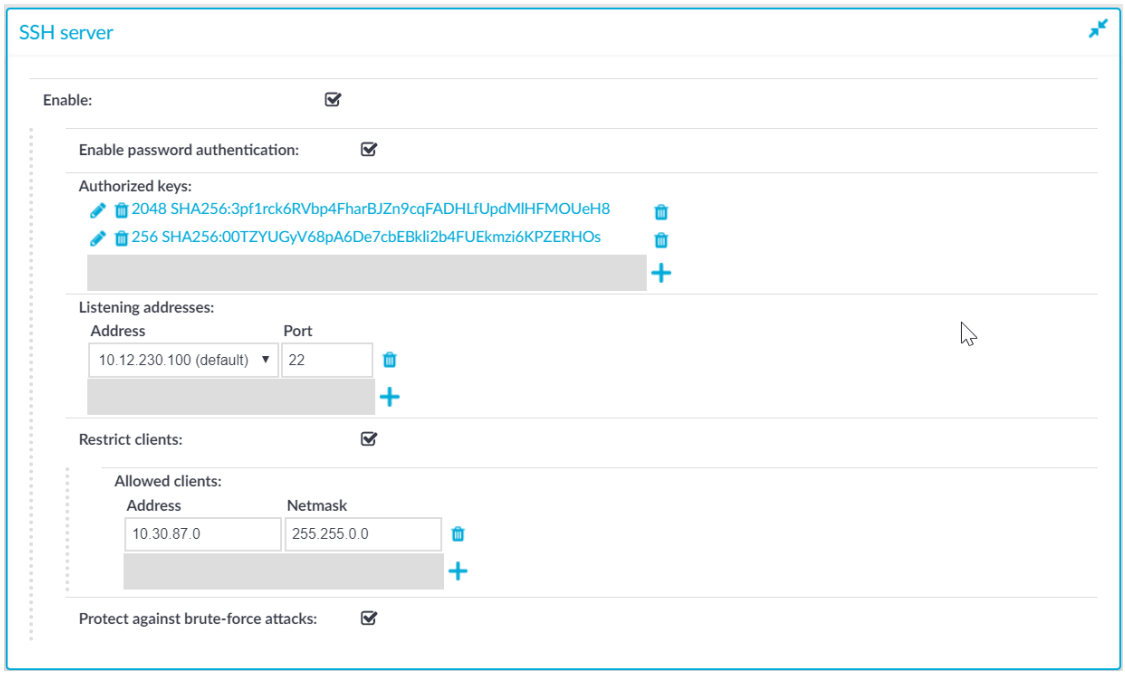
SSH access is, by default, protected against brute-force attacks: after 20 unsuccessful login attempts, the offending IP is blocked from accessing the SSH service for ten minutes.

You can turn off brute force protection by unselecting the **Protect against brute-force attacks** option for the SSH server.

To enable SSH access to the SPS host

1. Navigate to **Basic Settings** > **Local Services** > **SSH server**.

Figure 131: Basic Settings > Local Services > SSH server — Enabling remote SSH access to SPS



2. Select the **Enable** option.

   NOTE: Remote SSH access is automatically disabled if Sealed mode is enabled. For details, see Sealed mode on page 449.

3. Choose the authentication method for the remote SSH connections.
   - To enable password-based authentication, select the **Enable password authentication** option.
   - To enable public-key authentication, click + in the **Authorized keys** field, click ✎ and upload the public keys of the users who can access and manage SPS remotely via SSH.

     SPS allows you to use the following SSH host keys:
     - RSA (ssh-rsa), which is the most widely used public-key algorithm for the SSH key. In SPS, uploading RSA SSH host keys are supported in PKCS #1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

       TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

- Ed25519 (ssh-ed25519), which offers a better security and faster performance compared to RSA.

  In SPS, uploading Ed25519 SSH host keys are supported in PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

You can have multiple SSH server host keys on SPS for the same server, however, you cannot set more than one key for each type. For example, you can keep your old RSA SSH key and generate a new Ed25519 key but you cannot set two RSA keys.

4. Choose an address and a port for the SSH server in the **Listening addresses** section.

   The available addresses correspond to the interface addresses configured in **Basic Settings** > **Network** > **Interfaces**. Only IPv4 addresses can be selected.

   To add multiple addresses, click  .
5. (Optional) To permit SSH acces only from selected subnets or IP addresses, select **Restrict clients**, click  and enter the IP address and netmask of the allowed clients.

   Use an IPv4 address.

   To add multiple addresses, click  .
6. Click  .

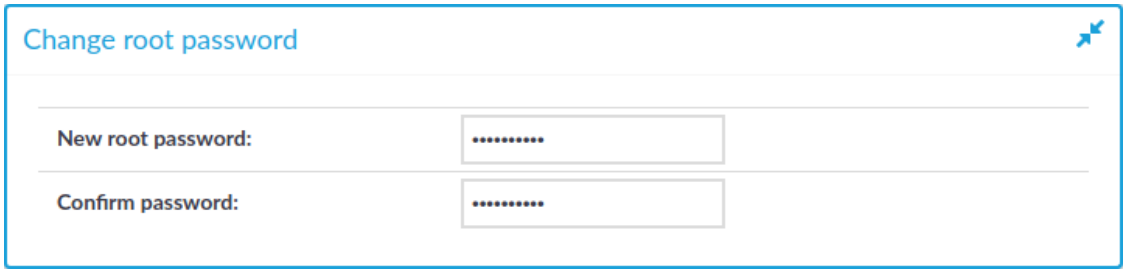
Changing the root password of One Identity Safeguard for Privileged Sessions (SPS)

The root password is required to access One Identity Safeguard for Privileged Sessions (SPS) locally, or remotely via an SSH connection. Note that the password of the `root` user can be changed from the console menu as well. For details, see Accessing the One Identity Safeguard for Privileged Sessions (SPS) console on page 442.

To change the root password of SPS

1. Navigate to **Basic Settings** > **Management** > **Change root password**.

   Figure 132: Basic Settings > Management > Change root password — Changing the root password of SPS

2. Enter the new password into the **New root password** and **Confirm password** fields.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

3. Click Commit.

Exporting and importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) using the console

For manual archiving, or to migrate it to another One Identity Safeguard for Privileged Sessions (SPS) unit, you can export/import the configuration of SPSfrom the console using the `/opt/scb/bin/configbundle.py` script.

NOTE: You must run the `/opt/scb/bin/configbundle.py` script using the `root` user.

NOTE: The configuration of your SPS may contain sensitive information. Make sure you delete any configuration export files that are not needed anymore.

To export/import the configuration of SPS from the console

1. Execute the following command to export the configuration of your SPS:

   ```
   /opt/scb/bin/configbundle.py create --bundle /<my destination
   folder>/bundle.tar.gz
   ```

   Where:

   - `/opt/scb/bin/configbundle.py`: The script you execute to export the configuration.
   - `create`: The option that lets you export a configuration.
   - `--bundle`: The option used to specify the bundle file.
   - `/<my destination folder>/bundle.tar.gz`: The path to the file where you wish to export the configuration.

     Replace `<my destination folder>` with the name of the folder where you wish to store the exported configuration.

2. Execute the following command to import the configuration of your SPS:

   ```
   /opt/scb/bin/configbundle.py import --bundle /<my destination
   folder>/bundle.tar.gz
   ```

   Where:

   - `/opt/scb/bin/configbundle.py`: The script you execute to import the configuration.
   - `import`: The option that lets you import a configuration.
   - `--bundle`: The option used to specify the bundle file.
   - `/<my destination folder>/bundle.tar.gz`: The path to the file from which you wish to import the configuration.

     Replace `<my destination folder>` with the name of the folder where your configuration export file is stored.

Data migration from an SPS instance to another SPS instance

If you need to switch from an One Identity Safeguard for Privileged Sessions (SPS) instance to another SPS instance, for example, your SPS appliance is old and you want to switch it to a new one, you can use the console menu to copy all data between SPS instances.

You can perform data migration between SPS instances having different versions. Data migration has the same version requirements as upgrade. For more information about upgrading, see Upgrading One Identity Safeguard for Privileged Sessions (SPS).

To copy all data and switch to the new SPS instance

1. From the console menu, select **Data migration between SPS instances**.
2. To initiate a preliminary copying of all the data from the source SPS to the target SPS without stopping the data traffic on the source SPS, select **Copy-only of all data to a new SPS instance (Optional)**.

   This step is optional but recommended to decrease production downtime caused by data and role migration. You can perform this step several times if required, for example, if the volume of your daily traffic is high.

   Skip this step only if a possible downtime due to having no preliminary copy is not an issue.
3. When you are ready to make the final switch to the new SPS instance, select **Copy and switch to the new SPS instance (Required)**.

   This process stops the data traffic on the source SPS, then copies all data from the source SPS to the target SPS. The target SPS also acquires the IP address of the source SPS.

Sealed mode

When sealed mode is enabled, the following settings are automatically applied:

- One Identity Safeguard for Privileged Sessions (SPS) cannot be accessed remotely via SSH for maintenance.
- The root password of SPS cannot be changed in sealed mode.
- It is not possible to upload or delete plugins in sealed mode.
- Sealed mode can be disabled only from the local console. For details, see Disabling sealed mode on page 449.

To enable sealed mode use one of the following methods:

- Select the **Sealed mode** option during the Welcome Wizard.
- Select **Basic Settings** > **System** > **Sealed mode** > **Activate sealed mode** on the SPS web interface.
- Log in to SPS as root using SSH or the local console, and select **Sealed mode** > **Enable** from the console menu.

Disabling sealed mode

The event of disabling sealed mode is logged. The following describes how to disable sealed mode.

To disable sealed mode

1. Go to the One Identity Safeguard for Privileged Sessions (SPS) appliance and access the local console.
2. Log in as `root`.
3. From the console menu, select **Sealed mode** > **Disable**
4. Select **Back to Main menu** > **Logout**.

Out-of-band management of One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) 8.0.1 LTS includes a dedicated out-of-band management interface conforming to the Intelligent Platform Management Interface (IPMI) v2.0 standards. The IPMI allows system administrators to monitor the system health of SPS and to manage the computer events remotely, independently of the operating system of SPS. SPS is accessible using the IPMI only if the IPMI is physically connected to the network.

NOTE: IPMI supports only 100 Mbps Full-Duplex speed.

Note that the IPMI supports only 100 Mbps Full-Duplex speed.

- For details on connecting the IPMI, see *Installing the SPS hardware* in the *Installation Guide*.
- For details on configuring the IPMI, see Configuring the IPMI from the console on page 451.
- For details on using the IPMI to remotely monitor and manage SPS, see the following document:

  For Safeguard Sessions Appliance 4000, see the X12 H12 BMC User's Manual

Basic information about the IPMI is available also on the SPS web interface on the **Basic Settings** > **High Availability** page. The following information is displayed:

Figure 133: Basic Settings > High Availability — Information about the IPMI SPS



- **Hardware serial number**: The unique serial number of the appliance.
- **IPMI IP address**: The IP address of the IPMI.
- **IPMI subnet mask**: The subnet mask of the IPMI.
- **IPMI default gateway**: The address of the default gateway configured for the IPMI.

- **IPMI IP address source**: Shows how the IPMI receives its IP address: dynamically from a DHCP server, or it uses a fixed static address.

Configuring the IPMI from the console

This section describes how you can modify the network configuration of IPMI from the console of One Identity Safeguard for Privileged Sessions (SPS).

Prerequisites

SPS is accessible using the IPMI only if the IPMI is physically connected to the network. For details on connecting the IPMI, see *Installing the SPS hardware* in the *Installation Guide*.

CAUTION:

IPMI searches for available network interfaces during boot. Make sure that IPMI is connected to the network through the dedicated Ethernet interface before SPS is powered on.

CAUTION: Security hazard!

The IPMI, like all out-of-band management interfaces, has known vulnerabilities that One Identity cannot fix or have an effect on. To avoid security hazards, One Identity recommends that you only connect the IPMI to well-protected, separated management networks with restricted accessibility. Failing to do so may result in an unauthorized access to all data stored on the SPS appliance. Data on the appliance can be unencrypted or encrypted, and can include sensitive information, for example, passwords, decryption keys, private keys, and so on.

For more information, see Best Practices for managing servers with IPMI features enabled in Datacenters.

NOTE: The administrator of SPS must be authorized and able to access the IPMI for support and troubleshooting purposes in case vendor support is needed.

The following ports are used by the IPMI:

- Port 22 (TCP): SSH (configurable)
- Port 80 (TCP): Web (configurable)
- Port 161 (UDP, TCP): SNMP (configurable)
- Port 443 (TCP): Web SSL (configurable)
- Port 623 (UDP): Virtual Media (configurable)
- Port 5900 (TCP): IKVM Server (configurable)
- Port 5985 (TCP): Wsman (configurable)

The SSH encrypted connection (port 22) works with the following properties:

Supported:	Safeguard Sessions Appliance 3000	Safeguard Sessions Appliance 3500
Ciphers	aes128-ctr, aes256-ctr	3des-cbc, aes128-ctr, aes128-cbc, aes256-ctr, aes256-cbc
KEX algorithm	curve25519-sha256, ecdh-sha2-nistp256, curve25519-sha256@libssh.org, ecdh-sha2-nistp384, diffie-hellman-group1-sha1, ecdh-sha2-nistp521, diffie-hellman-group14-sha1	curve25519-sha256, ecdh-sha2-nistp256, curve25519-sha256@libssh.org, ecdh-sha2-nistp384, diffie-hellman-group1-sha1, ecdh-sha2-nistp521, diffie-hellman-group14-sha1
MACs	hmac-sha1, hmac-sha2-256, hmac-sha1-96, hmac-sha2-512	hmac-md5, hmac-sha2-256, hmac-sha1, hmac-sha2-512, hmac-sha1-96
HostKey algorithms	ssh-rsa, ssh-dss	ssh-rsa, ssh-dss
Compression	enabled	enabled

SSL-encrypted connections work with the following properties:

Supported:	Safeguard Sessions Appliance 3000	Safeguard Sessions Appliance 3500
TLSv1.2	enabled	enabled
TLS Fallback SCSV	supported	supported
Heartbleed	not vulnerable	not vulnerable
Server Ciphers	Preferred TLSv1.2 256 bits ECDHE-RSA-AES256-GCM-SHA384 Curve P-256 DHE 256	Preferred TLSv1.2 256 bits ECDHE-RSA-AES256-GCM-SHA384 Curve P-256 DHE 256
	Accepted TLSv1.2 256 bits ECDHE-RSA-AES256-SHA384 Curve P-256 DHE 256	Accepted TLSv1.2 256 bits ECDHE-RSA-AES256-SHA384 Curve P-256 DHE 256
	Accepted TLSv1.2 256 bits DHE-RSA-AES256-GCM-SHA384 DHE 2048 bits	Accepted TLSv1.2 256 bits DHE-RSA-AES256-GCM-SHA384 DHE 2048 bits
	Accepted TLSv1.2 256 bits DHE-RSA-AES256-SHA256 DHE 2048 bits	Accepted TLSv1.2 256 bits DHE-RSA-AES256-SHA256 DHE 2048 bits
	Accepted TLSv1.2 256 bits AES256-GCM-SHA384	Accepted TLSv1.2 256 bits AES256-GCM-SHA384
	Accepted TLSv1.2 256 bits AES256-SHA256	Accepted TLSv1.2 256 bits AES256-SHA256

Supported:	Safeguard Sessions Appliance 3000	Safeguard Sessions Appliance 3500
	Accepted TLSv1.2 128 bits ECDHE-RSA-AES128-GCM-SHA256 Curve P-256 DHE 256	Accepted TLSv1.2 128 bits ECDHE-RSA-AES128-GCM-SHA256 Curve P-256 DHE 256
	Accepted TLSv1.2 128 bits ECDHE-RSA-AES128-SHA256 Curve P-256 DHE 256	Accepted TLSv1.2 128 bits ECDHE-RSA-AES128-SHA256 Curve P-256 DHE 256
	Accepted TLSv1.2 128 bits DHE-RSA-AES128-GCM-SHA256 DHE 2048 bits	Accepted TLSv1.2 128 bits DHE-RSA-AES128-GCM-SHA256 DHE 2048 bits
	Accepted TLSv1.2 128 bits DHE-RSA-AES128-SHA256 DHE 2048 bits	Accepted TLSv1.2 128 bits DHE-RSA-AES128-SHA256 DHE 2048 bits
	Accepted TLSv1.2 128 bits AES128-GCM-SHA256	Accepted TLSv1.2 128 bits AES128-GCM-SHA256
	Accepted TLSv1.2 128 bits AES128-SHA256	Accepted TLSv1.2 128 bits AES128-SHA256
Server Key Exchange Groups	TLSv1.2 128 bits secp256r1 (NIST P-256)	TLSv1.2 128 bits secp256r1 (NIST P-256)
Server Signature Algorithms	TLSv1.2 Server accepts all signature algorithms.	TLSv1.2 Server accepts all signature algorithms.

To modify the network configuration of IPMI from the console of SPS

1. Use the local console (or SSH) to log in to SPS as root.
2. Choose **Shells** > **Boot shell**.
3. Check the network configuration of the interface:

   ```
   # ipmitool lan print
   ```

   This guide assumes that `channel 1` is used for LAN. If your setup differs, adjust the following commands accordingly.
4. Configure the interface. You can use DHCP or configure a static IP address manually.

   Use an IPv4 address.

   - To use DHCP, enter the following command:

     ```
     # ipmitool lan set 1 ipsrc dhcp
     ```
   - To use static IP, enter the following command:

     ```
     # ipmitool lan set 1 ipsrc static
     ```

     Set the IP address:

```
# ipmitool lan set 1 ipaddr <IPMI-IP>
```

Set the netmask:

```
# ipmitool lan set 1 netmask <IPMI-netmask>
```

Set the IP address of the default gateway:

```
# ipmitool lan set 1 defgw ipaddr <gateway-IP>
```

5. Verify the network configuration of IPMI:

   ```
   # ipmitool lan print 1
   ```

   Use a browser to connect to the reported network address.

6. Change the default password:
   a. Log in to the IPMI web interface using the default login credentials (username: `ADMIN`, password: `ADMIN` or `changeme`, depending on your hardware).

      NOTE: The login credentials are case sensitive.

   b. Navigate to **Configure** > **Users**.
   c. Select **ADMIN**, and choose **Modify User**.
   d. Change the password, and save the changes with **Modify**.

Configuring the IPMI from the BIOS

To configure IPMI from the BIOS when configuring your One Identity Safeguard for Privileged Sessions (SPS) physical appliance for the first time, complete the following steps.

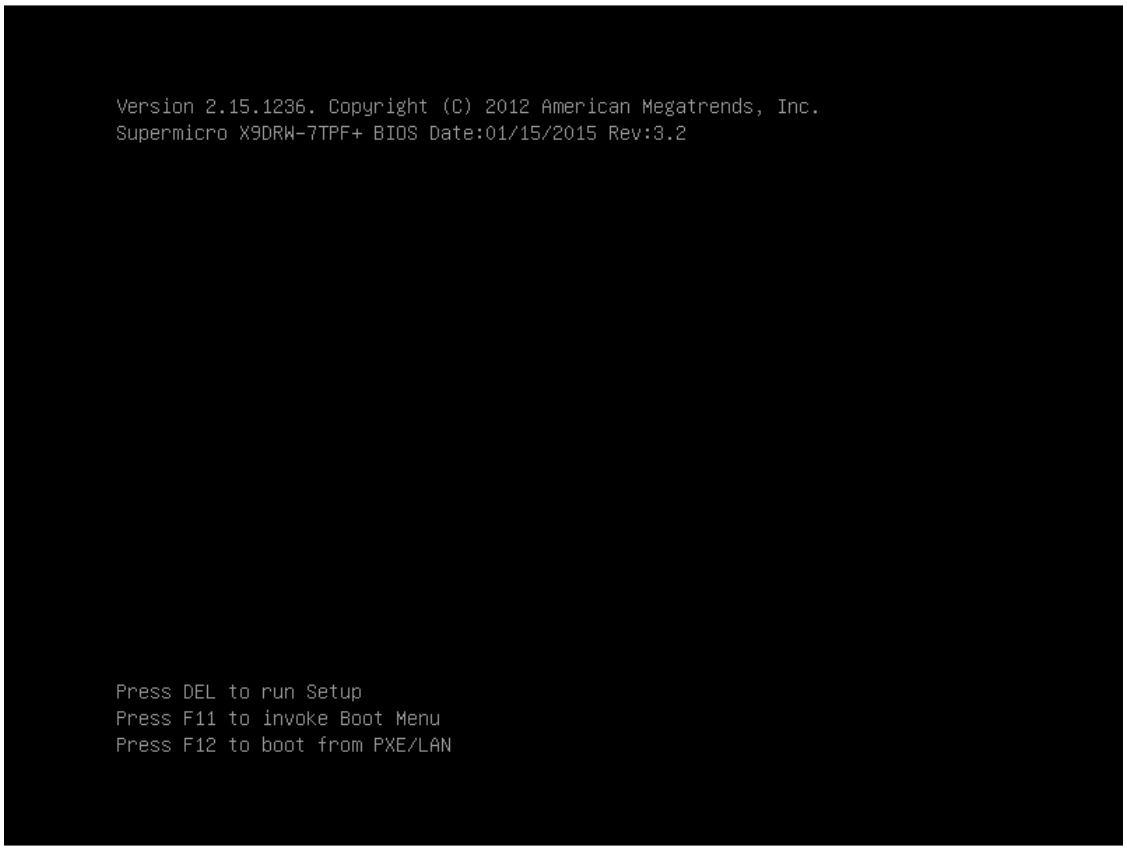
Prerequisites

To apply the procedure outlined here, you will need physical access to a monitor and keyboard.

To configure the IPMI from the BIOS

1. Press the DEL button when the POST screen comes up while the appliance is booting.

Figure 134: POST screen during booting



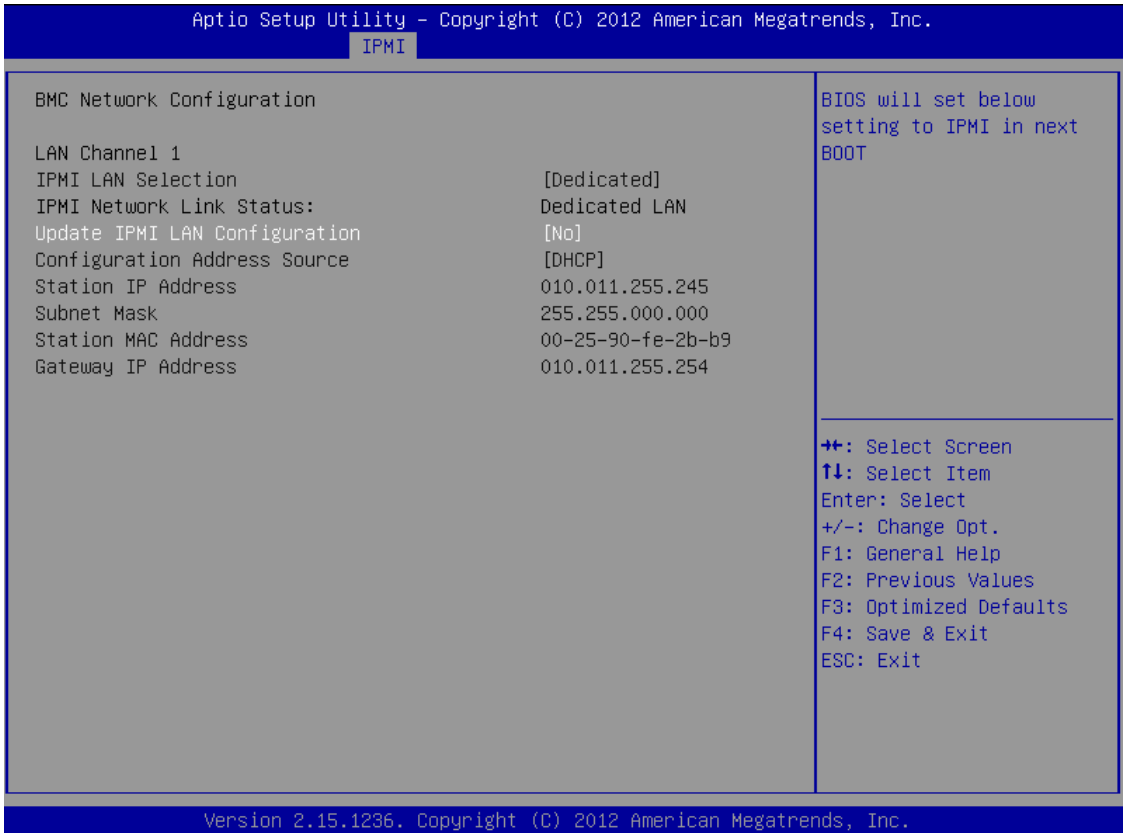
2. In the BIOS, navigate to the **IPMI** page.
3. On the **IPMI** page, select **BMC Network Configuration**, and press Enter.

Figure 135: IPMI page > BMC Network Configuration option



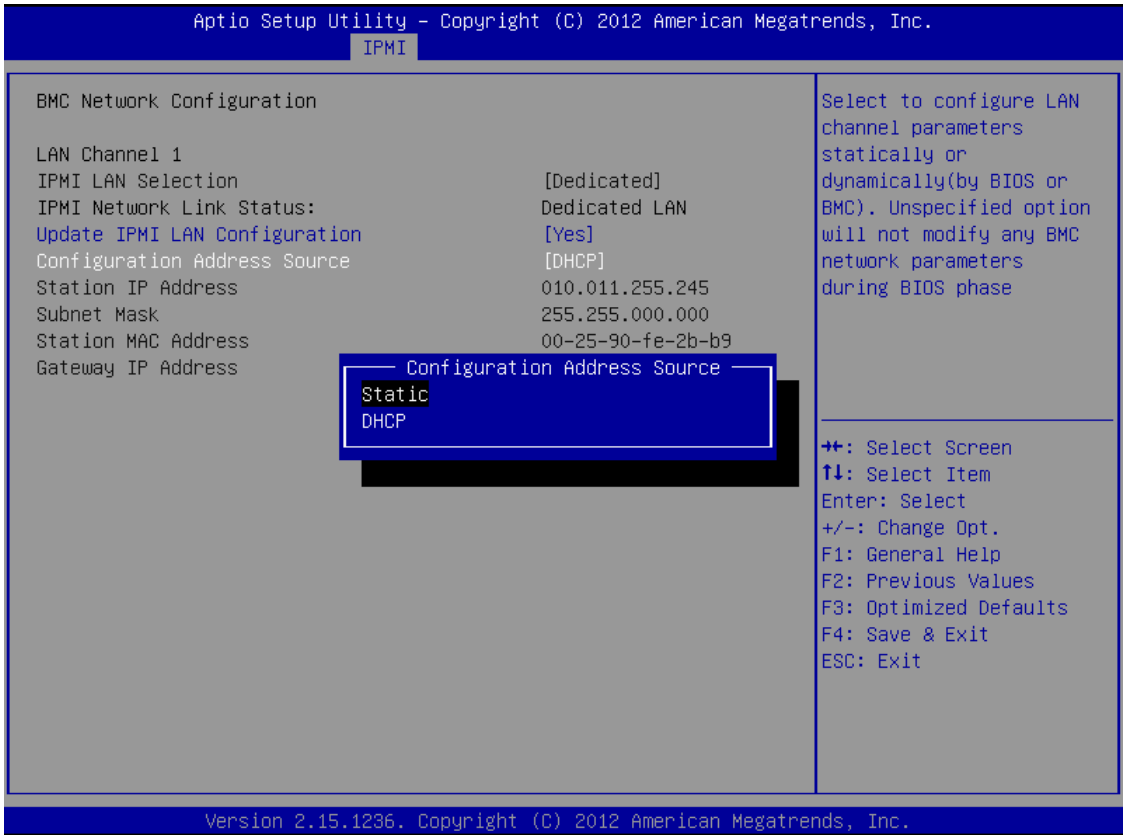
4. On the **BMC Network Configuration** page, select **Update IPMI LAN Configuration**, press `Enter`, and select **Yes**.

Figure 136: BMC Network Configuration page > Update IPMI LAN Configuration



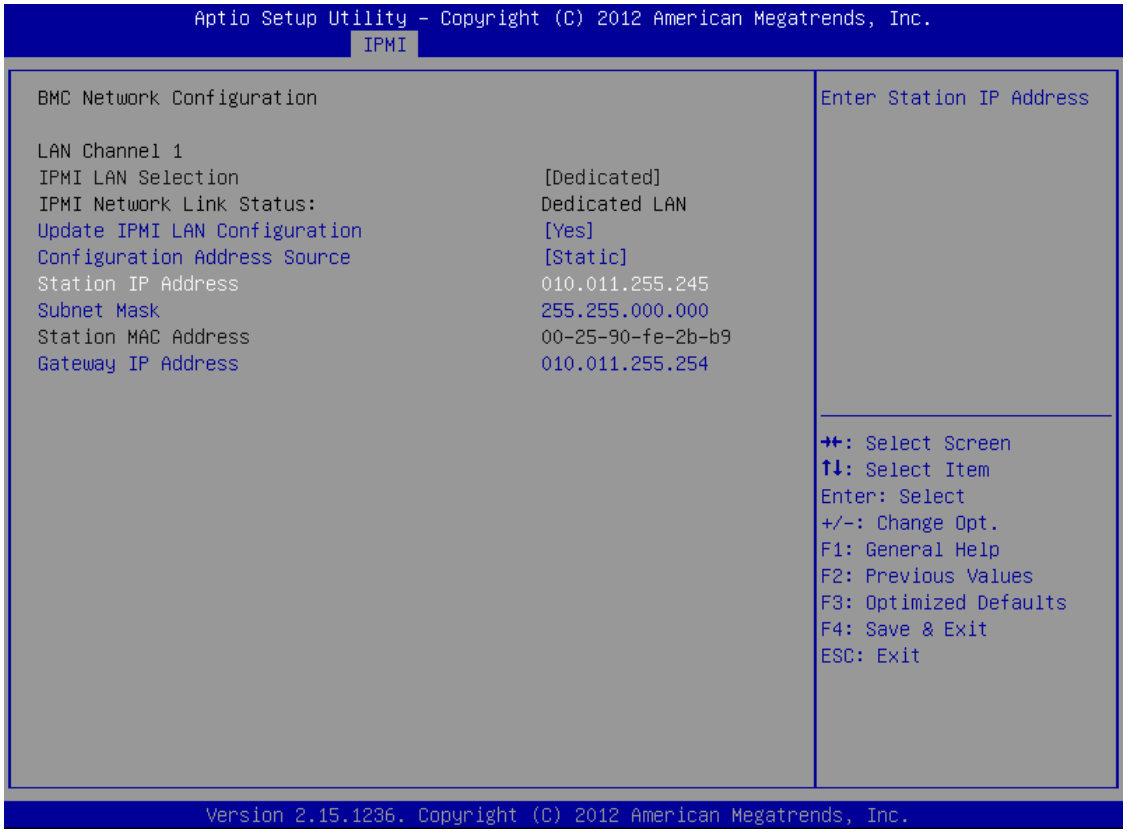
5. Stay on the **BMC Network Configuration** page, select **Configuration Address Source**, press **Enter**, and select **Static**.

Figure 137: BMC Network Configuration page > Configuration Address Source



6. Still on the **BMC Network Configuration** page, configure the **Station IP Address**, **Subnet Mask**, and **Gateway IP Address** individually.

Figure 138: BMC Network Configuration page > Station IP Address, Subnet Mask, Gateway IP Address



7. Press F4 to save the settings, and exit from the BIOS.

   About a minute later, you will be able to log in on the IPMI web interface.

Managing the certificates used on One Identity Safeguard for Privileged Sessions (SPS)

One Identity Safeguard for Privileged Sessions (SPS) uses a number of certificates for different tasks that can be managed from the **Basic Settings** > **Management** > **SSL certificates** menu.

CAUTION:

Starting from 6.10.0, SPS (SPS) has changed to hardened SSL settings. As a result, during TLS session establishment, the following items are not considered secure:

- Private keys and X.509 certificates having RSA or DSA keys shorter than 2048 bits, or ECC keys shorter than 224 bits.
- Certificates (other than Root CA certificates) with signatures that use the SHA-1 or the MD5 hashing algorithm.

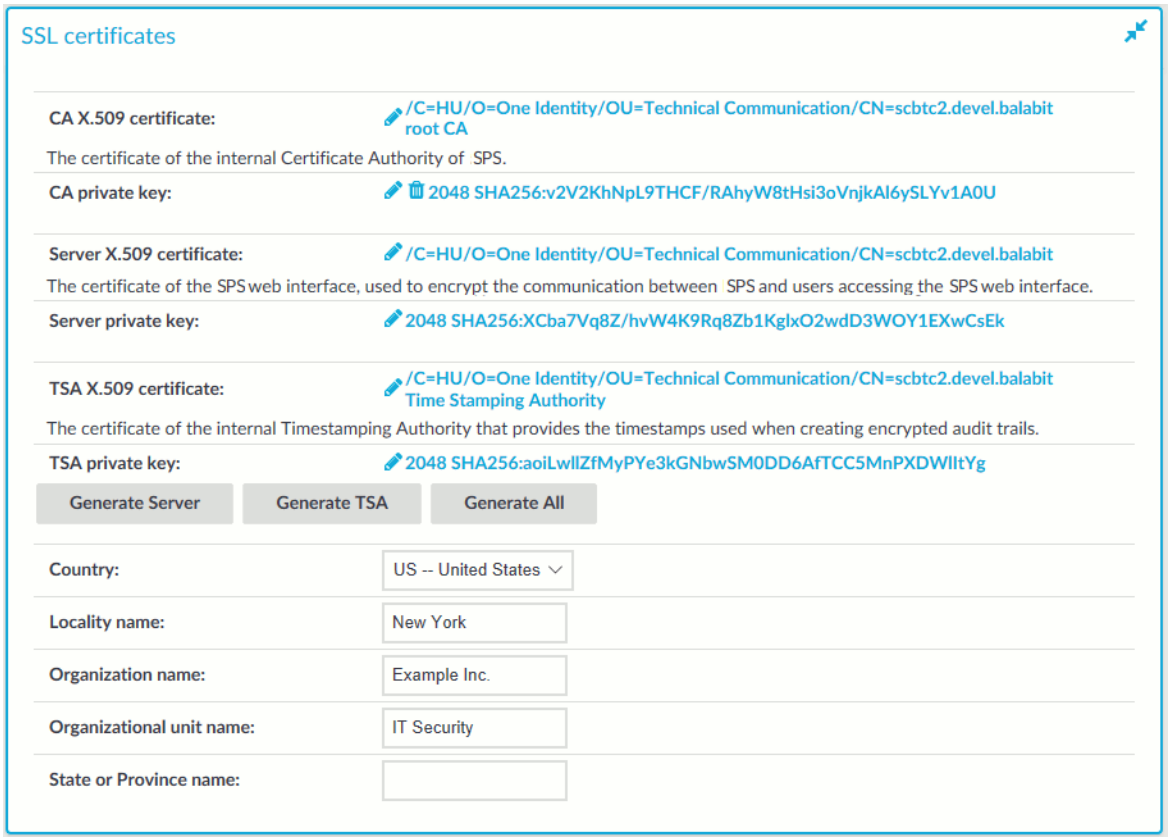
With the hardened SSL settings, SPS will not connect to remote systems that are protected with weak certificates.

You cannot upgrade SPS if your configuration contains insecure certificates, keys or certificate chains in any of the following sections:

- SPS web interface
- internal CA certificate
- connection policy TLS settings
- client X.509 credentials for external LDAP, SMTP or Syslog connections
- server X.509 certificates for external SMTP or Splunk servers
- external indexer credentials (only writable over the REST API)
- CA certificates in Trusted CA Lists and Trust Stores

Note that the certificates and keys that are used for signing, timestamping, encryption or decryption are not affected by this change.

Figure 139: Basic Settings > Management > SSL certificates — Changing the web certificate of SPS



The following certificates can be modified here:

- **CA certificate**: The certificate of the internal Certificate Authority of SPS.
- **Server certificate**: The certificate of the SPS web interface, used to encrypt the communication between SPS and the administrators.

  NOTE: If this certificate is changed, the browser of SPS users will display a warning stating that the certificate of the site has changed.
- **TSA certificate**: The certificate of the internal Timestamping Authority that provides the timestamps used when creating encrypted audit-trails.

NOTE: SPS uses other certificates for different purposes that are not managed here, for example, to encrypt data stored on SPS. For details, see Encrypting audit trails on page 512.

Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates (for example, do not use the certificate of the One Identity Safeguard for Privileged Sessions (SPS) webserver to encrypt audit trails, or the same keypair for signing and encrypting data).

For every certificate, the distinguished name (DN) of the X.509 certificate and the fingerprint of the private key is displayed. To display the entire certificate click on the DN. To display the public part of the private key, click on the fingerprint. It is not possible to download the private key itself from the SPS web interface, but the public part of the key

can be downloaded in different formats (for example PEM, DER, or OpenSSH). Also, the X.509 certificate can be downloaded in PEM and DER formats.

During the initial configuration, SPS creates a self-signed CA certificate, and uses this CA to issue the certificate of the web interface (see **Server certificate**) and the internal Timestamping Authority (**TSA certificate**).

There are two methods to manage certificates of SPS:

- **Recommended**: Generate certificates using your own PKI solution and upload them to SPS.

  Generate a CA certificate and two other certificates signed with this CA using your PKI solution and upload them to SPS. For the Server and TSA certificates, upload the private key as well. One Identity recommends using 2048-bit RSA keys (or stronger), and to use certificates that have the appropriate `keyUsage` or `extendedKeyUsage` fields set (for example, **`extendedKeyUsage=serverAuth`** for the SPS web server certificate).

  For details on uploading certificates and keys created with an external PKI, complete Uploading external certificates to One Identity Safeguard for Privileged Sessions (SPS) on page 463.

  CAUTION:

  The Server and the TSA certificates must be issued by the same Certificate Authority.

- Use the certificates generated on SPS. In case you want to generate new certificates and keys for SPS using its self-signed CA certificate, or generate a new self-signed CA certificate, complete Generating certificates for One Identity Safeguard for Privileged Sessions (SPS) on page 462.

  NOTE: Generate certificates using your own PKI solution and upload them to SPS whenever possible. Certificates generated on SPS cannot be revoked, and can become a security risk if they are somehow compromised.

Generating certificates for One Identity Safeguard for Privileged Sessions (SPS)

Create a new certificate for the One Identity Safeguard for Privileged Sessions (SPS) webserver or the Timestamping Authority using the internal CA of SPS, or create a new, self-signed CA certificate for the internal Certificate Authority of SPS.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To create a new certificate for the SPS webserver

1. Navigate to **Basic Settings** > **Management** > **SSL certificates**.
2. Fill the fields of the new certificate:

a. **Country**: Select the country where SPS is located (for example HU - Hungary).
b. **Locality name**: The city where SPS is located (for example Budapest).
c. **Organization name**: The company who owns SPS (for example Example Inc.).
d. **Organization unit name**: The division of the company who owns SPS (for example IT Security Department).
e. **State or Province name**: The state or province where SPS is located.

3. Select the certificate you want to generate.
   - To create a new certificate for the SPS web interface, select **Generate Server**.
   - To create a new certificate for the Timestamping Authority, select **Generate TSA**.
   - To create a new certificate for the internal Certificate Authority of SPS, select **Generate All**. Note that in this case new certificates are created automatically for the server and TSA certificates as well.

   NOTE: When generating new certificates, the server and TSA certificates are signed using the certificate of the CA. Consider the following:
   - If you have configured an external CA certificate along with its private key, it will be used to create the new server and TSA certificates.
   - If you have configured an external CA certificate without its private key, use your external PKI solution to generate certificates and upload them to SPS. In this case, only the **Generate All** button is visible on the UI, while **Generate TSA** and **Generate Server** are not available.

   CAUTION:
   Generating a new certificate automatically deletes the earlier certificate.

4. Click Commit.

Uploading external certificates to One Identity Safeguard for Privileged Sessions (SPS)

Upload a certificate generated by an external PKI system to One Identity Safeguard for Privileged Sessions (SPS).

Prerequisites

The certificate to upload. For the **TSA X.509 Certificate** and **Server X.509 Certificate**, the private key of the certificate is needed as well. The certificates must meet the following requirements:

- SPS accepts certificates in PEM format. The DER format is currently not supported.
- SPS accepts private keys in PEM format, using RSA, DSA, and EC private key algorithms. Password-protected private keys are also supported.

  NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

  - Letters A-Z, a-z
  - Numbers 0-9
  - The space character
  - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

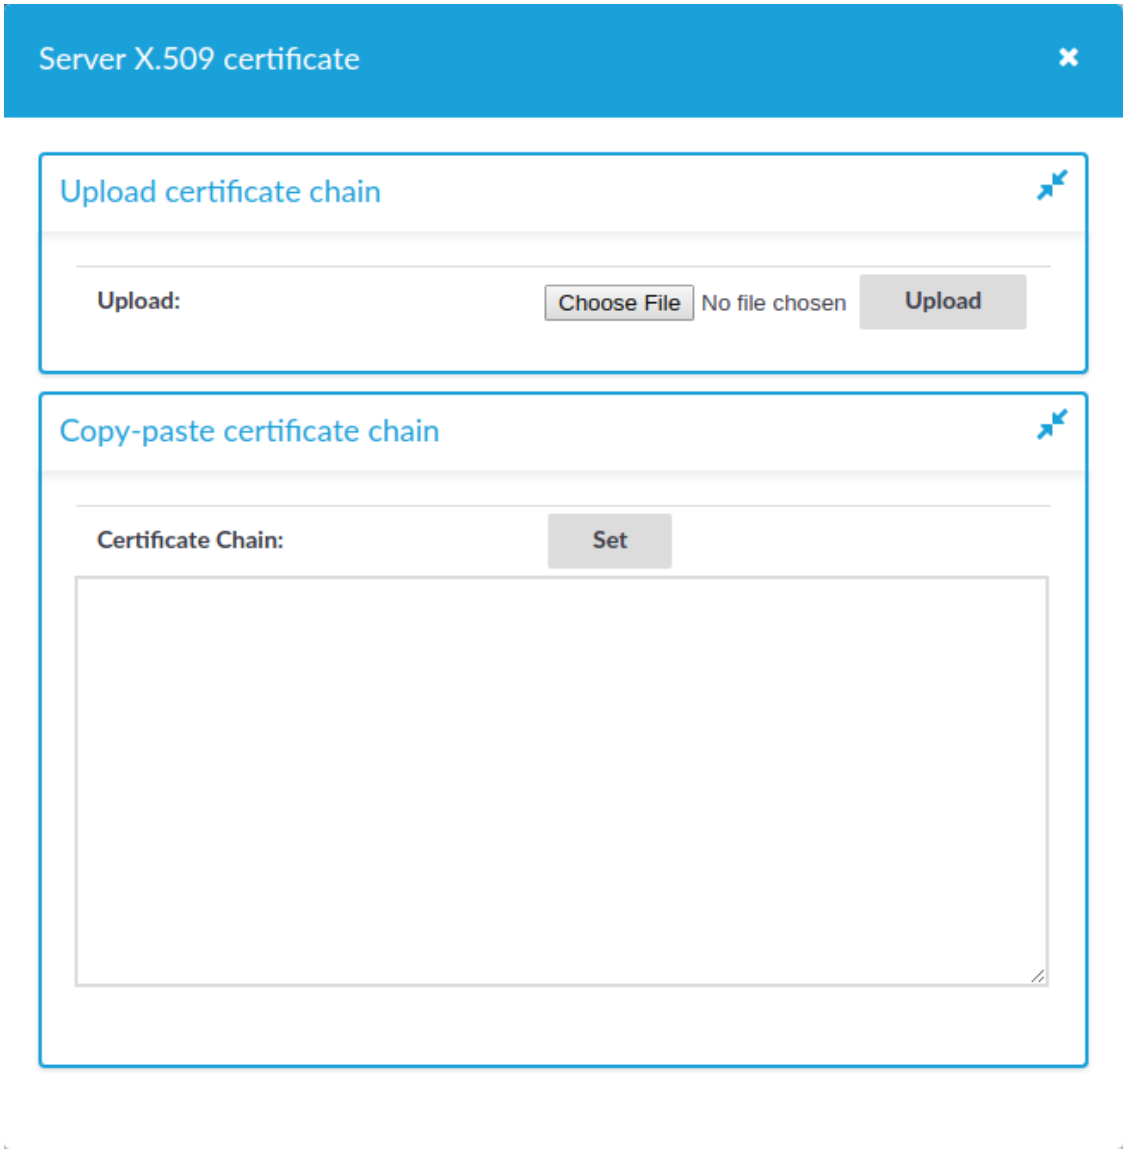
  For the internal CA certificate of SPS, uploading the private key is not required.
- For the TSA certificate, the `X509v3 Extended Key Usage` extension must be set to **`critical`** with the value **`Time Stamping`**. Also, the `Key Usage` extension must be `non repudiation` and `digital signature` (that is, without key encipherment or other key usage).
- For the Server certificate, the `X509v3 Extended Key Usage` extension must be set to **`TLS Web Server Authentication`**. Also, the `Common Name` of the certificate must contain the domain name or the IP address of the SPS host. If the web interface is accessible from multiple interfaces or IP addresses, list every IP address using the `Subject Alt Name` extension.
- For the certificate used to sign audit trails, the `X509v3 Extended Key Usage` extension must be set to **`Sign (downloadable) executable code`**.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To upload a certificate generated by an external PKI system to SPS

1. Navigate to **Basic Settings** > **Management** > **SSL certificates**.
2. Click  to upload the new certificate. A pop-up window is displayed.

Figure 140: Basic Settings > Management > SSL certificates — Uploading certificates



Select **Browse**, select the file containing the certificate, and click **Upload**.

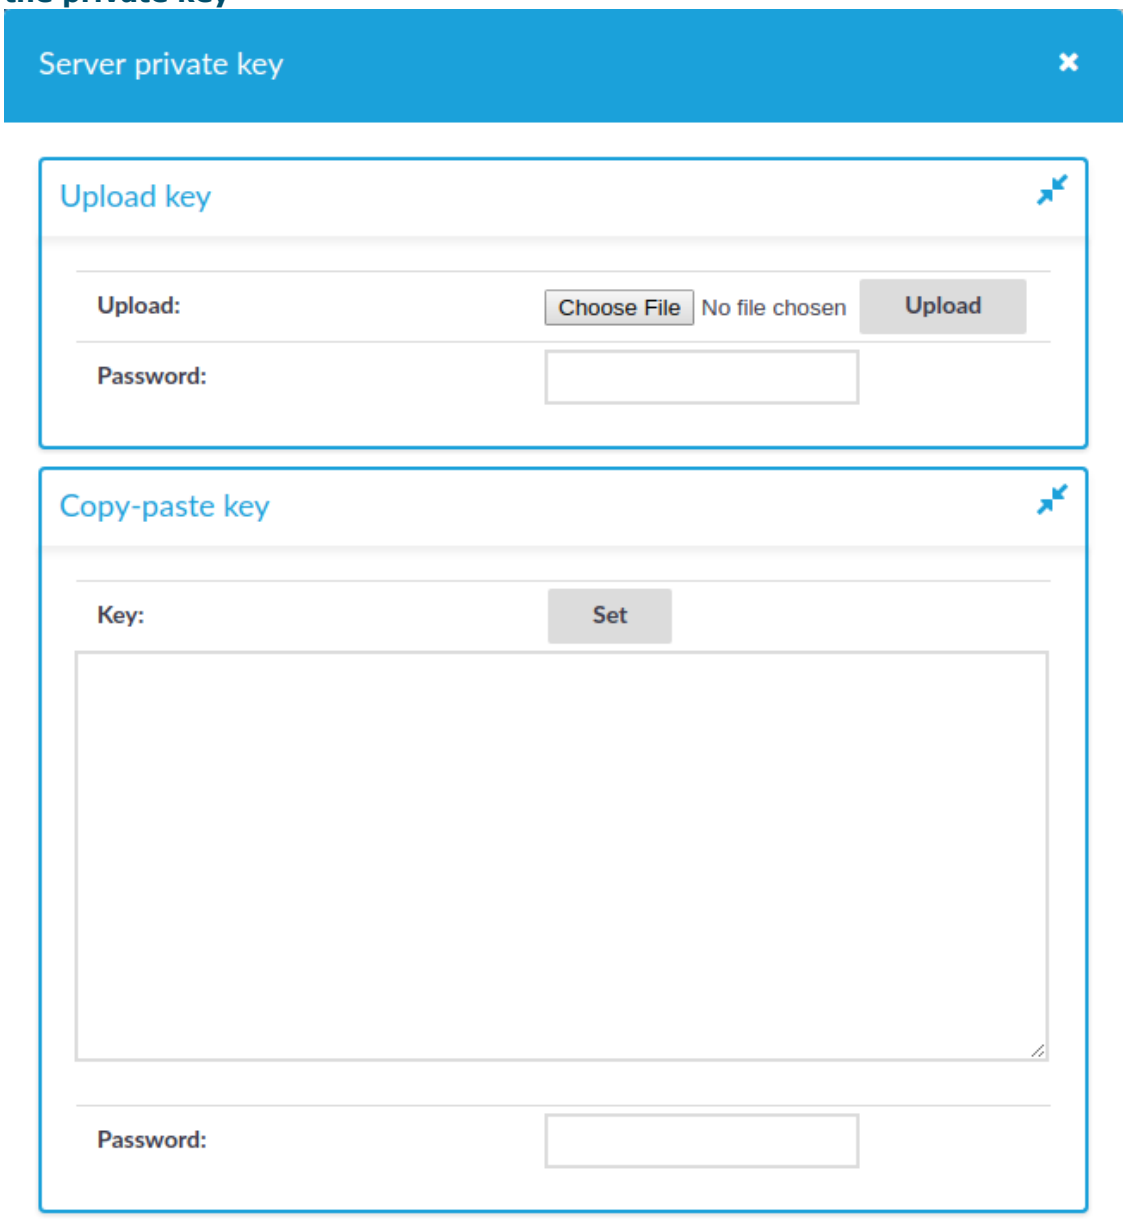
For the Server X.509 Certificate

For the **Server X.509 Certificate**, you can also upload a certificate chain. For that, copy the certificates after each other in a single file. Alternatively, you can copy and paste the certificates one by one after each other into the **Certificate** field and click **Set**. The certificates do not have to be in order, SPS will order them and validate the chain: if a member of the chain is missing, an error message is displayed.

NOTE: Certificate chains are supported only for the **Server X.509 Certificate**.

To upload the private key corresponding to the certificate, click icon. A pop-up window is displayed.

Figure 141: Basic Settings > Management > SSL certificates — Uploading the private key



3.

Select **Browse**, select the file containing the private key, provide the **Password** if the key is password-protected, and click **Upload**. Alternatively, you can also copy-paste the private key into the **Key** field, provide the **Password** there, and click **Set**.

In the case of a certificate chain, the private key has to be the same as the bottom level certificate.

Expected result

The new certificate is uploaded. If you receive the `Certificate issuer mismatch` error message after importing a certificate, you must import the CA certificate which signed the certificate as well (the private key of the CA certificate is not mandatory).

NOTE: To download previously uploaded certificates, click on the certificate and either download the certificate (or certificate chain) in one single PEM or DER file, or you can download single certificate files separately (if it is a certificate chain).

Generating TSA certificate with Windows Certificate Authority on Windows Server 2016 or later

To generate a TSA certificate with Windows Certificate Authority (CA) that works with One Identity Safeguard for Privileged Sessions (SPS), generate a CSR (certificate signing request) on a computer running OpenSSL and sign it with Windows CA, then import this certificate into SPS for timestamping.

Prerequisites

A valid configuration file for OpenSSL with the following extensions:

```
[ tsa_cert ]
extendedKeyUsage = critical,timeStamping
```

TIP: You can copy `/etc/xcb/openssl-ca.cnf` from SPS to the computer that will be used for signing. Rename the file to `openssl-temp.cnf`.

The TSA certificate is considered valid, in terms of compatibility with SPS, if the following conditions are met:

- The `X509v3 Extended Key Usage` extension of the TSA certificate is set to **`critical`** with the value `Time Stamping`.
- The `Key Usage` extension is `non repudiation` and `digital signature` (that is, without key encipherment or other key usage).

  CAUTION:

  In **Encryption**, do not select **Allow key exchange only with key encryption (key encipherment)**, because it will result in errors.

The following X509v3 extensions are supported:

- Hard requirement:

  X509v3 Extended Key Usage must be `critical`, and must only contain `Time Stamping`.

- Optional:

  X509v3 Key Usage, if present, must be `digitalSignature` and/or `nonRepudiation`.

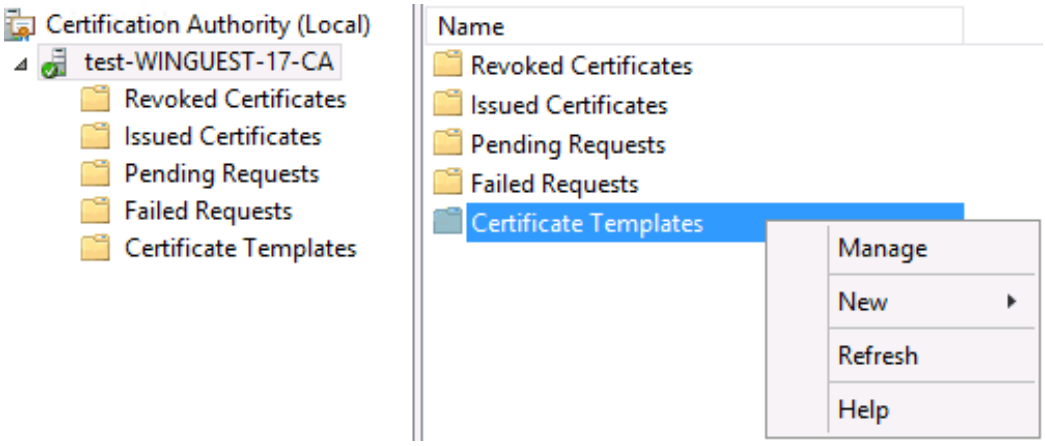
To generate TSA certificate with Windows Certificate Authority on Windows Server 2016 or later

1. Create CSR using the new configuration file: `openssl req -set_serial 0 -config openssl-temp.cnf -reqexts tsa_cert -new -newkey rsa:2048 -keyout timestamp.key -out timestamp.csr -nodes`
2. Complete the required fields according to your environment:

   ```
   Generating a 2048 bit RSA private key
   ......................+++
   ....................................+++
   writing new private key to 'timestamp.key'
   -----
   You are about to be asked to enter information that will be incorporated
   into your certificate request.
   What you are about to enter is what is called a Distinguished Name or a
   DN.
   There are quite a few fields but you can leave some blank
   For some fields there will be a default value,
   If you enter '.', the field will be left blank.
   -----
   Country Name (2 letter code) [AU]:US
   State or Province Name (full name) []:New York
   Locality Name (eg, city) []:New York
   Organization Name (eg, company) [Internet Widgits Pty Ltd]:Examplecompany
   IT Security
   Organizational Unit Name (eg, section) []:Service Delivery
   Common Name (eg, YOUR name) []:scb35-1-i1.tohuvabohu.examplecompany
   Email Address []:vlad@examplecompany.com
   ```

3. Create and configure a time stamping web server template in the Certificate Authority, and use that to generate the TSA certificate.
   a. Start the Certification Authority Microsoft Management Console, and select the CA server.
   b. Right-click on **Certificate Templates**, and choose **Manage**.

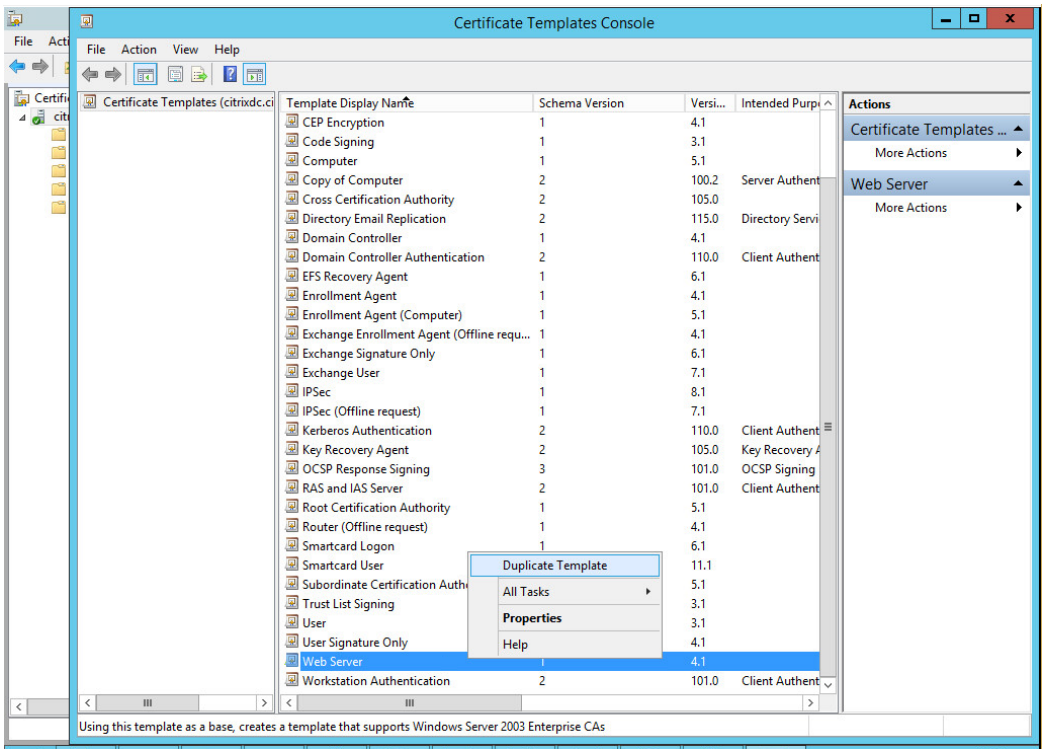
Figure 142: Managing certificate templates



The *Certificate Templates Console* opens.

c. Right-click the **Web Server** template, and choose **Duplicate Template**.

Figure 143: Duplicating a Template

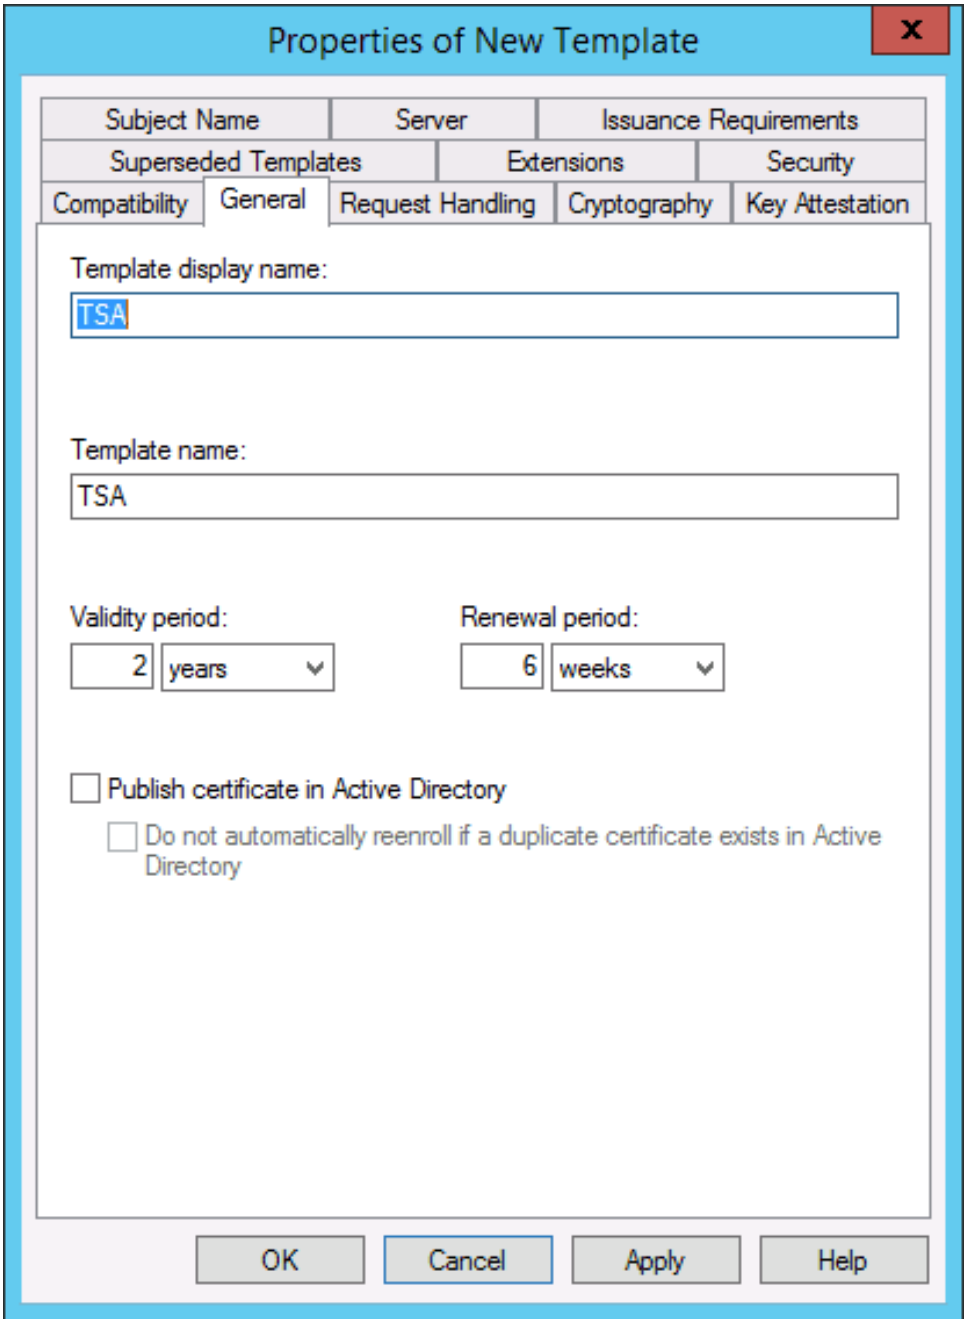


The *Properties of New Template* window is displayed.

d. Make the following changes to the new template:

- On the *General* tab, change the **Template display name** to TSA.

Figure 144: Creating the new template

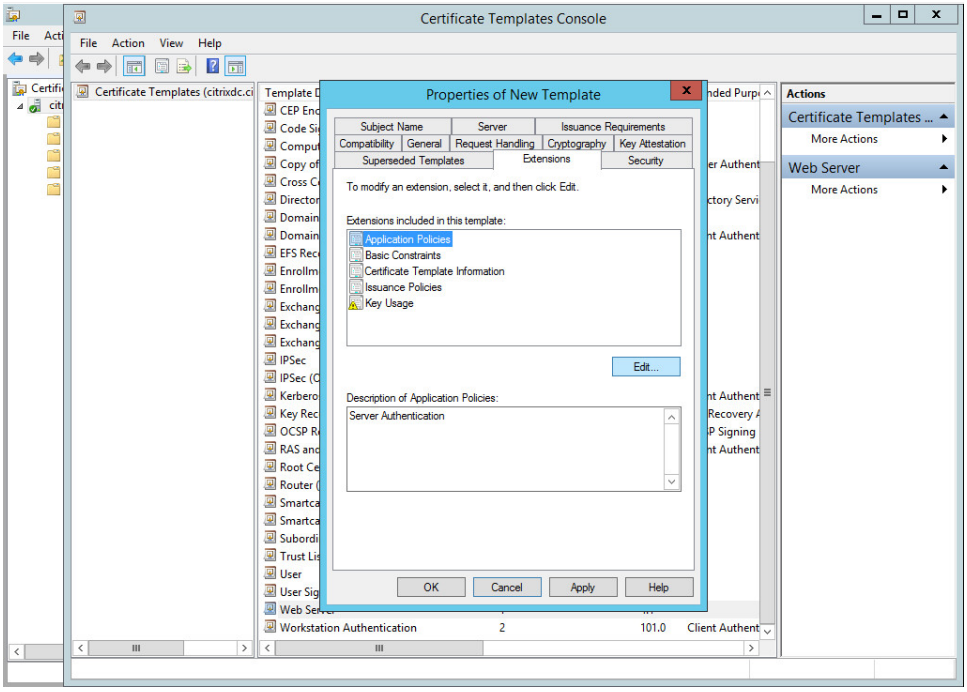


- On the *Request Handling* tab, enable the **Allow private key to be exported** option.
- On the *Extensions* tab, make the following changes:

Edit Application Policies

Select **Application Policies** and click **Edit** below the list of extensions.

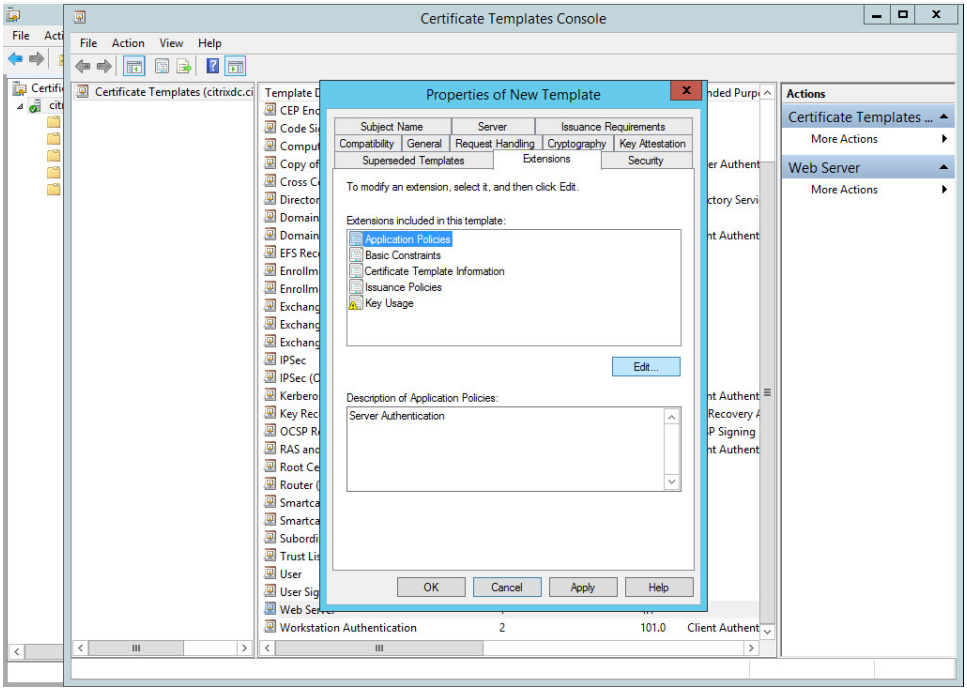
Figure 145: Editing Application Policies



Remove Server Authentication

Select **Server Authentication** and click **Remove**.

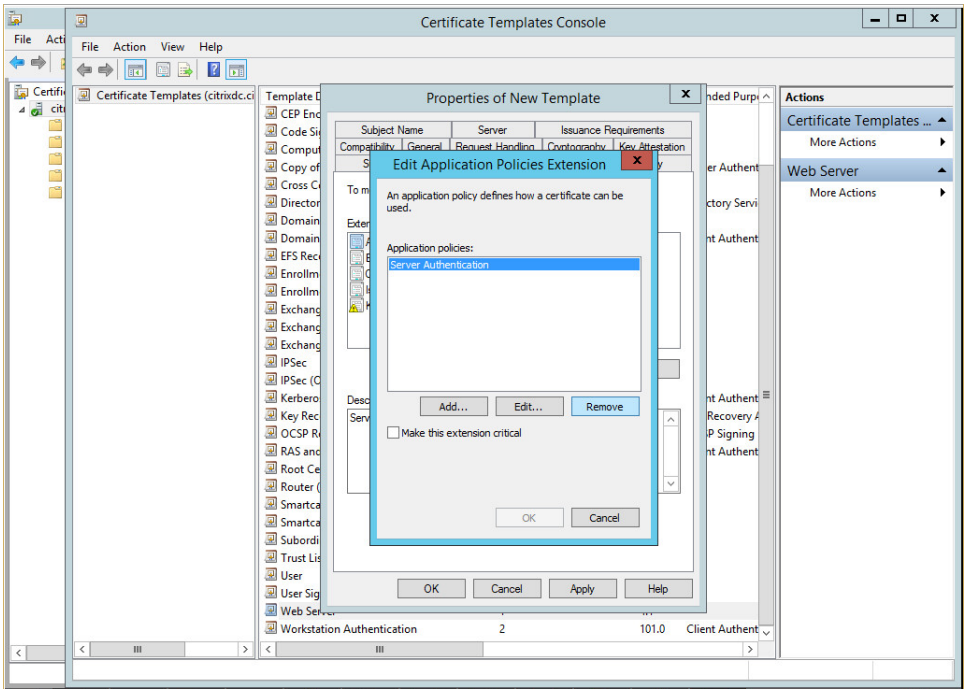
Figure 146: Removing Server Authentication



Add Time Stamping

Click **Add**, select **Time Stamping** and click **OK**.

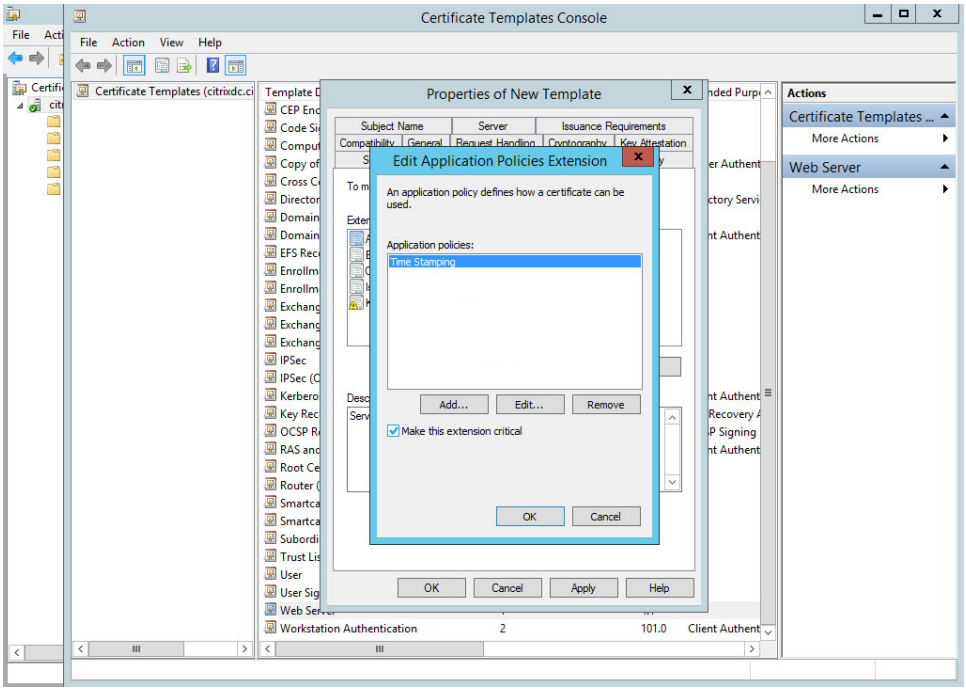
Figure 147: Adding Time Stamping



Make Time Stamping critical

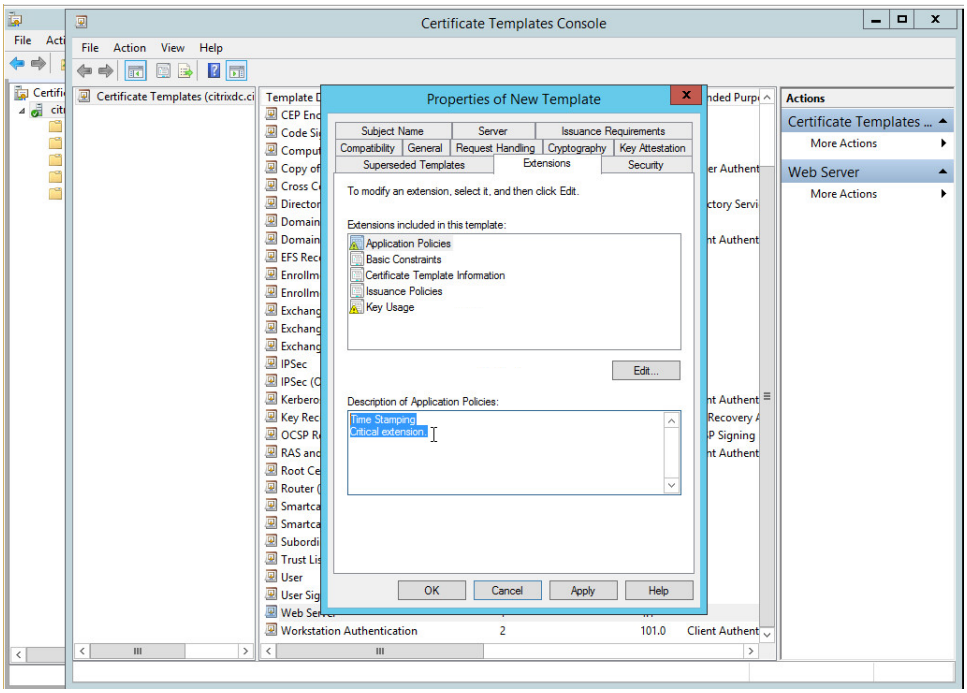
Select **Time Stamping** and enable the **Make this extension critical** option, then click **OK**.

Figure 148: Making Time Stamping critical



Time Stamping and **Critical extension** are listed in the **Description of Application Policies**.

Figure 149: Description of Application Policies



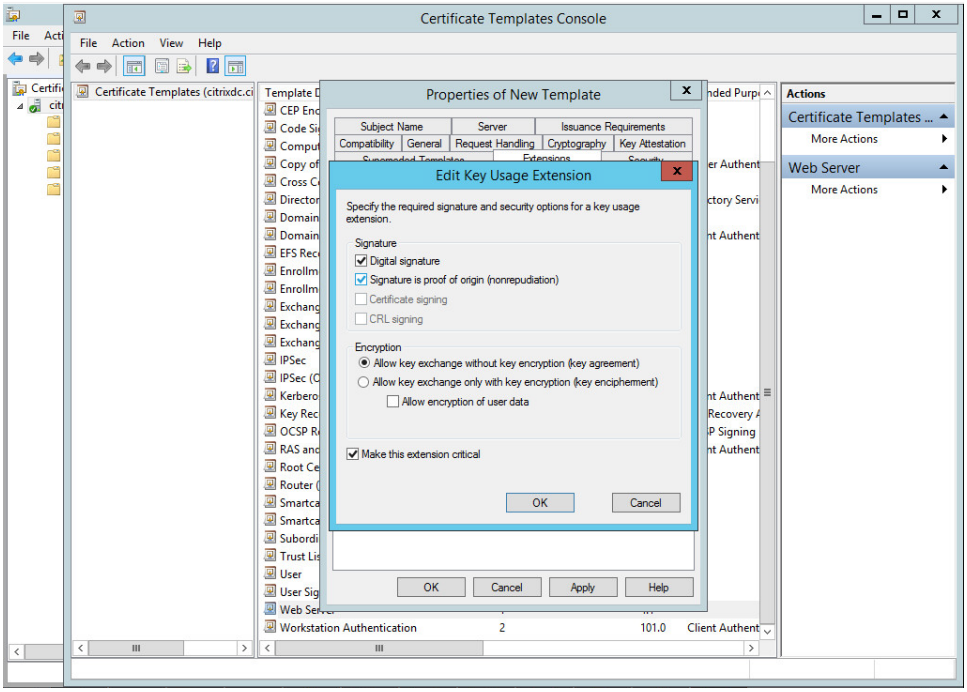
Edit Key Usage

Select **Key usage**, click **Edit**. Enable the **Signature is proof of origin (nonrepudiation)** option.

Select **Allow key exchange without key encryption (key agreement)**.

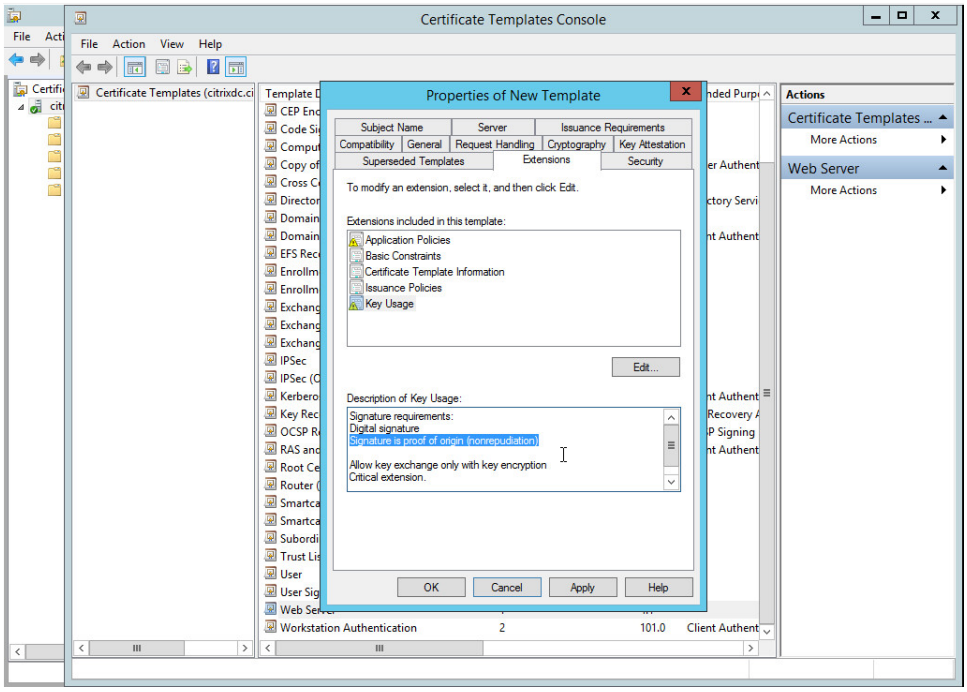
Click **OK**.

Figure 150: Editing Key Usage



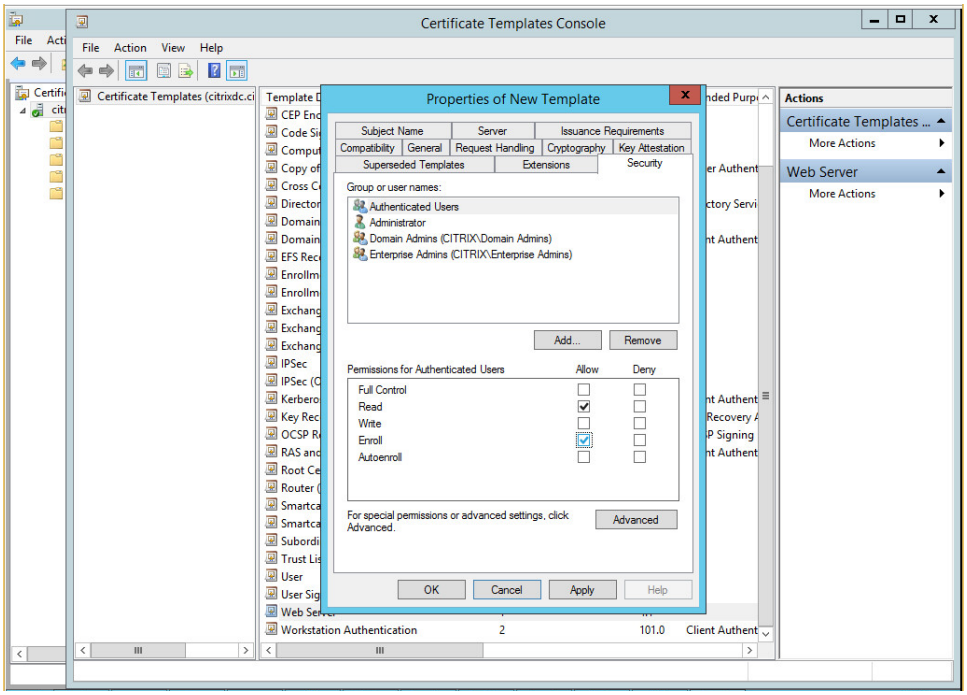
The following are listed in the **Description of Key Usage**.

Figure 151: Description of Key Usage



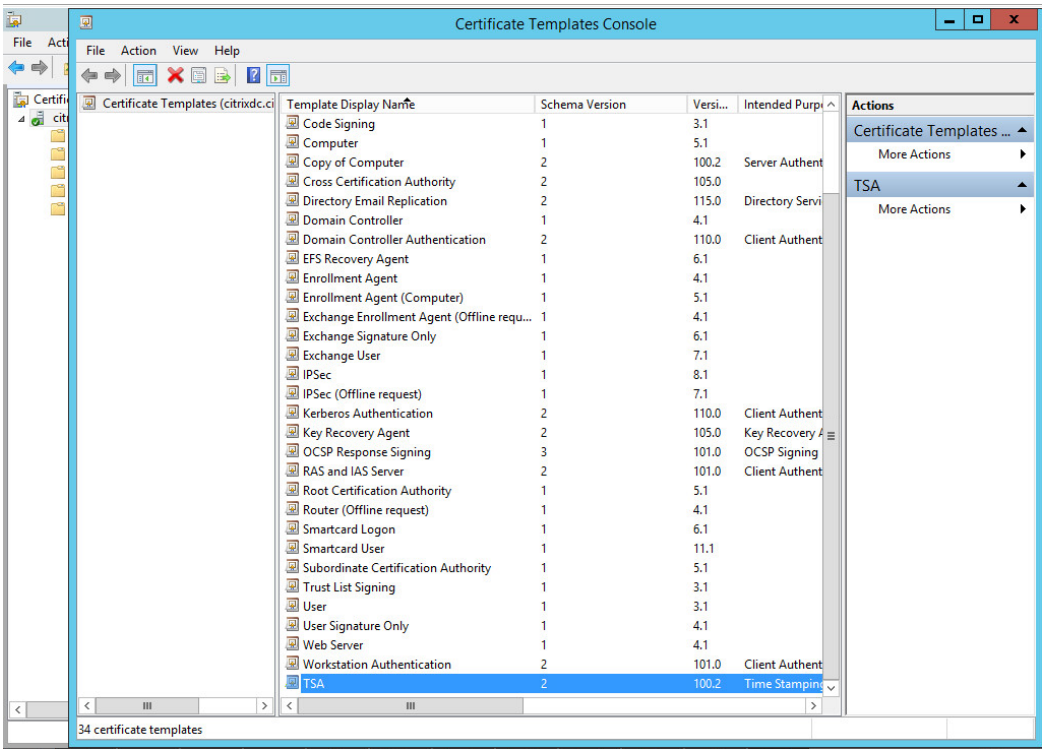
- On the *Security* tab, select **Authenticated Users**, and set **Enroll** to **Allow**.

Figure 152: Configuring permissions for the template



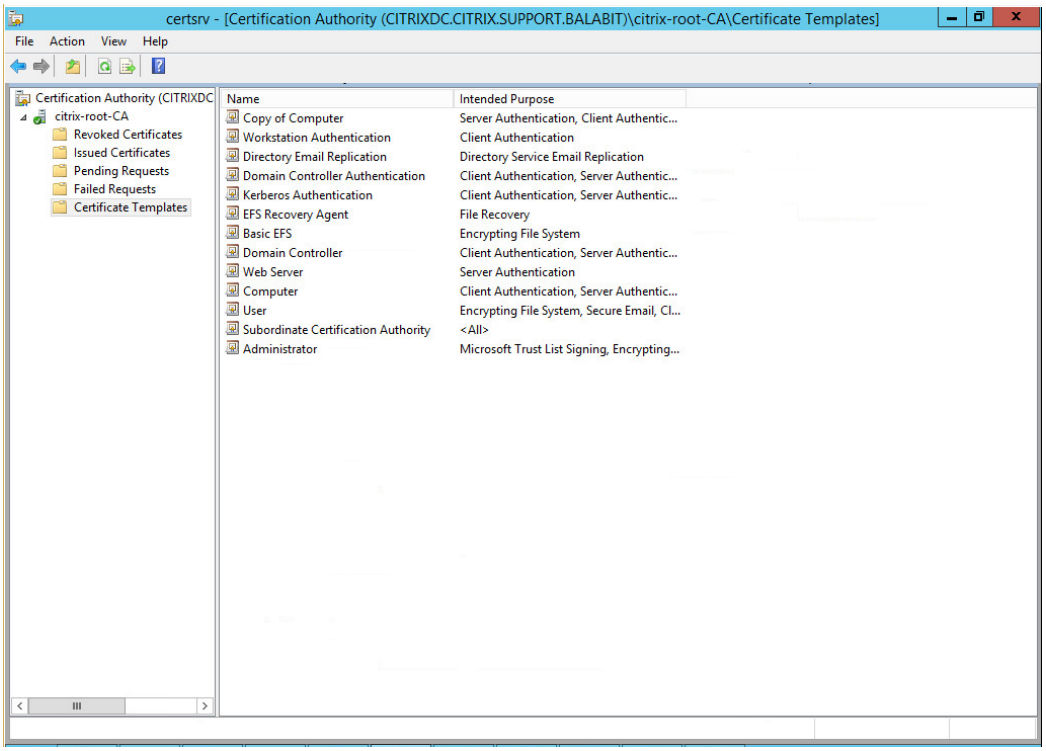
e. Click **Apply**. Click **OK**. The new TSA template is now displayed in the list of templates.

Figure 153: The new TSA template is now displayed in the list of templates



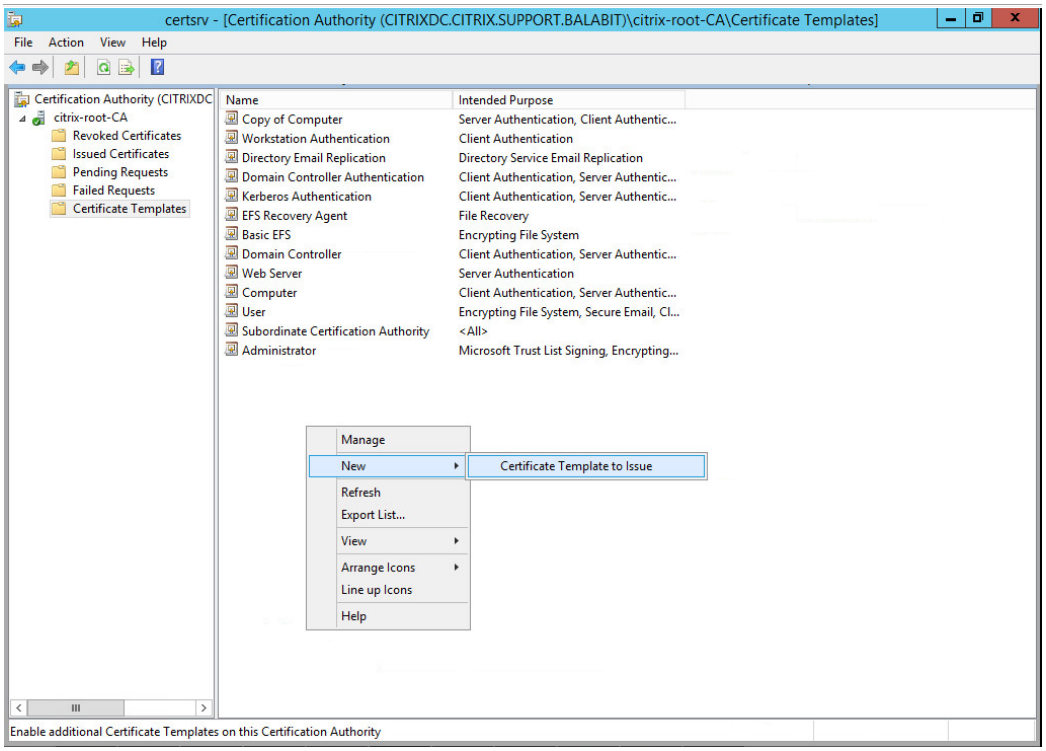
f. Close this window and return to the Certification Authority main screen, and select the **Certificate Templates** folder.

Figure 154: Certificate Templates



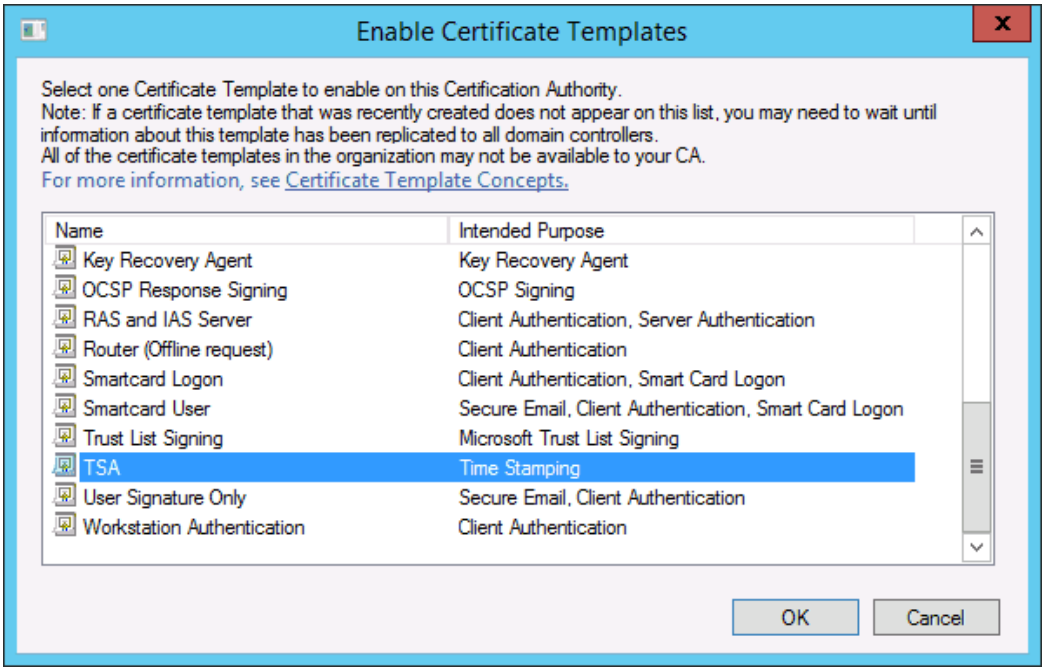
Right-click under the list, and choose **New** > **Certificate Template to Issue**.

Figure 155: Certificate Template to Issue



The *Enable Certificate Templates* window is displayed.

Figure 156: Enable the new template



g. Select the TSA certificate template, and choose **OK**. Close this window.

h. Open the command line, and issue the following command:

```
certreq -submit -attrib "CertificateTemplate:TSA" <CSR>
```

Replace <CSR> with the full path of the CSR created earlier (in the second step).

i. The *Certification Authority List* is displayed. Select the CA.

j. The *Save Certificate* window is displayed. Choose an output folder.

The certificate is generated to the specified folder.

4. In SPS, navigate to **Basic Settings** > **Management** > **SSL certificates**.

5. Click  next to **TSA X.509 certificate**, browse for the previously generated certificate, and click **Upload**.

6. Click  next to **TSA private key**, browse for the previously generated key, and click **Upload**.

NOTE: If the root CA (the **CA X.509 certificate** field under **Basic Settings** > **Management** > **SSL certificates**) that is used for other certificates on SPS is different from the CA that was used to sign the TSA certificate, a warning is displayed. In this scenario, ignore this warning.

General connection settings

Connections determine if a server can be accessed from a particular client.

- The policies used in the connection definition can restrict the availability of the *connection* based on the user name, time, authentication method, and so on. Channel policies (see Creating and editing channel policies on page 495) determine if a particular channel can be used within an already established connection.
- The policies used in the channel policy can restrict the availability of the *channel* based on the server and the client IP address, user name, and so on. The types of policies available in a connection depend on the protocol (SSH, RDP, and so on) enabled in the connection.

SPS compares the connection policies to the parameters of the connection request one-by-one, starting with the first policy in the policy list. SPS applies to the connection the first connection policy that completely matches the connection request.

This section describes how to configure connections, and details the general configuration options and policies that apply to every type of connection that SPS can control: HTTP, ICA, RDP, SSH, Telnet, and VNC. For a detailed list of supported protocol versions, see Supported protocols and client applications on page 31.

Protocol-specific configuration options are described in their respective sections: HTTP-specific settings on page 541, ICA-specific settings on page 559, RDP-specific settings on page 577, SSH-specific settings on page 616, Telnet-specific settings on page 651, and VNC-specific settings on page 664.

Configuring connections

This section describes how to configure connections.

NOTE:

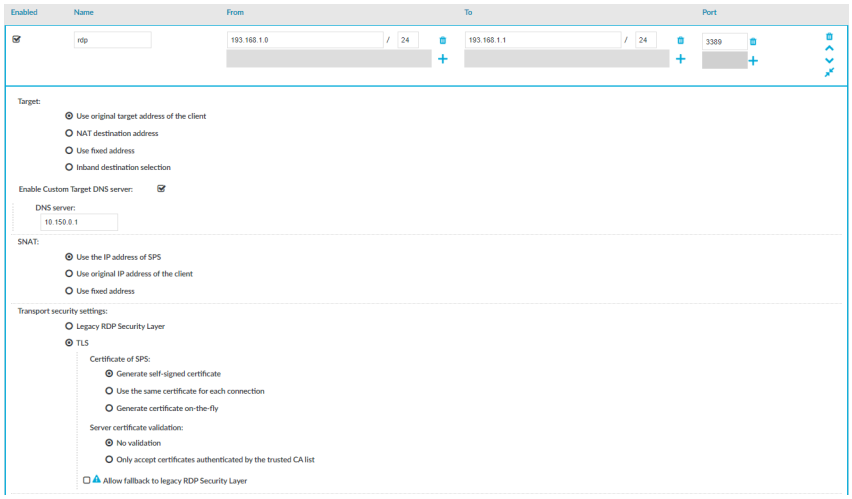
When configuring HTTP or SSH connections, avoid using the IP address configured for administrator or user login on SPS.

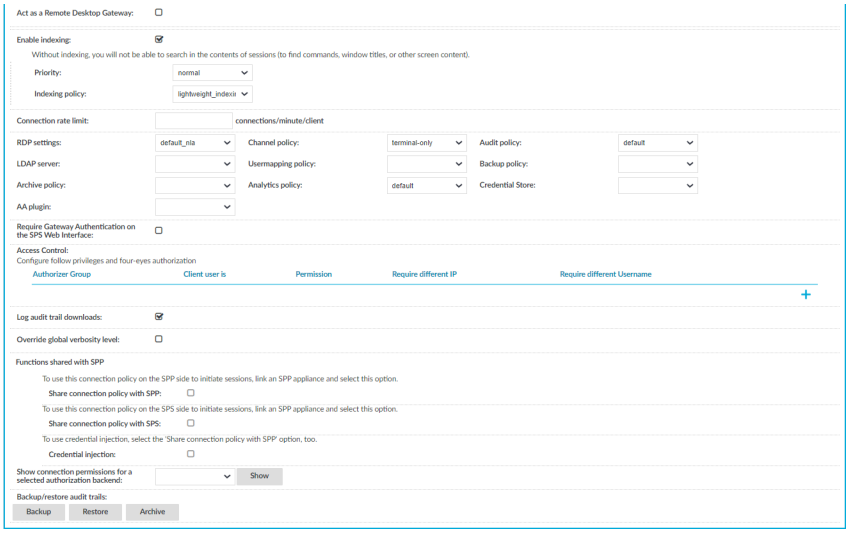
To configure connections

1. Select the type of connection from the main menu.
    - To configure an HTTP connection, select **HTTP Control** > **Connections**.
    - To configure an ICA connection, select **ICA Control** > **Connections**.
    - To configure a Remote Desktop connection, select **RDP Control** > **Connections**.
    - To configure a Secure Shell connection, select **SSH Control** > **Connections**.
    - To configure a Telnet connection, select **Telnet Control** > **Connections**.
    - To configure a VNC connection, select **VNC Control** > **Connections**.
2. Click  to define a new connection and enter a name that identifies the connection (for example, `admin_mainserver`).

    TIP: Use descriptive names that give information about the connection, for example, refer to the name of the accessible server, the allowed clients, and so on.

Figure 157: <Protocol name> Control> Connections — Configuring connections





3. In the **From** field, enter the IP address of the client that is permitted to access the server. To list additional clients, click + .

   You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).

   Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

   NOTE: Note the following limitations:

   - To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
   - If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.

4. In the **To** field, enter the IP address that the clients request.

   You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).

   Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

   NOTE: Note the following limitations:

   - To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
   - If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.

- In non-transparent mode, enter the IP address of an SPS logical interface.

  For more information on setting up logical network interfaces on SPS, see Managing logical interfaces on page 119.

  For more information, see Non-transparent mode.
- In transparent mode, enter the IP address of the protected server.

  For more information, see Transparent mode.

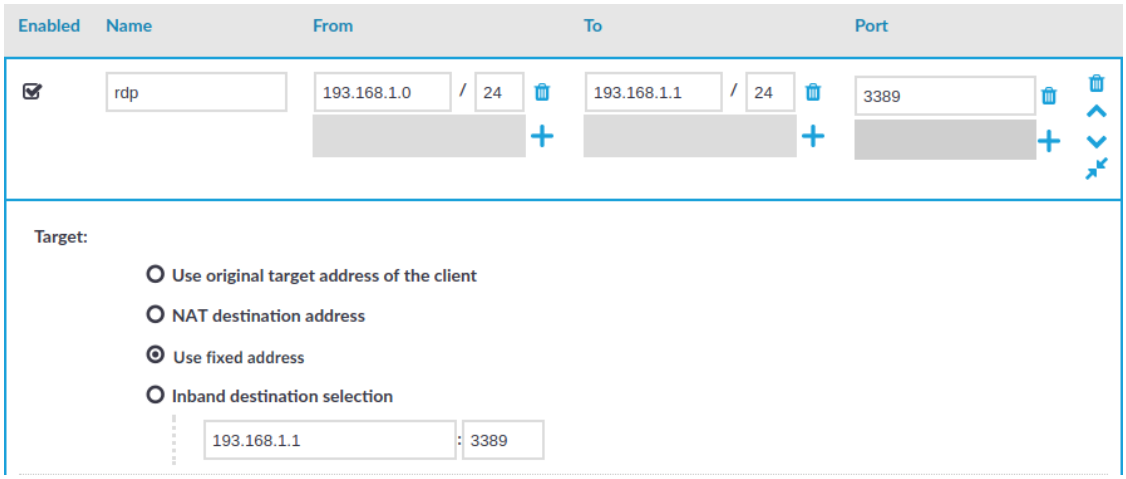
To add additional IP addresses, click + .

5. If the clients use a custom port to address the server instead of the default port of the protocol, in the **Port** field, enter the port number that the clients request. To list additional port numbers, click + .

   NOTE: SPS can handle a maximum of 15 unique ports per connection policy. If you want to add more than 15 custom ports, create additional connection policies.
6. *Non-transparent mode*: In the **Target** field, enter the IP address and port number of the target server. SPS connects all incoming client-side connections to this server.

   For details on organizing connections in non-transparent mode, see Organizing connections in non-transparent mode on page 915.

Figure 158: <Protocol name> Control> Connections — Configuring non-transparent connections



7. If needed, configure advanced settings (for example, network address translation, channel policy, gateway authentication, various policies, or other settings).
8. To save the connection, click **Commit**.

   TIP: To temporarily disable a connection, deselect the checkbox of the connection.
9. If needed, reorder the list of the connection policies. You can move connection policies by clicking the ^ and ˅ buttons.

SPS compares the connection policies to the parameters of the connection request one-by-one, starting with the first policy in the policy list. SPS applies to the connection the first connection policy that completely matches the connection request.

10. Depending on your needs and on your environment, you can configure the following settings for your connections:

    - Modify the destination or source addresses of the connections.

      For more information, see Modifying the destination address on page 488 and Modifying the source address on page 493.

    - Select a **Backup Policy** and an **Archiving Policy** for the audit trails and indexes of the connection.

      For more information on creating backup and archive policies, see Data and configuration backups on page 145 and Archiving on page 158.

      If you have indexed trails, the index is archived every 30 days.

      CAUTION:

      Hazard of data loss! Make sure you also back up your data besides archiving it.

      For more information, see Data and configuration backups on page 145.

      If a system crash occurs, you can lose up to 30 days of index, since the index is only archived every 30 days.

    NOTE: The backup and archive policies set for the connection apply only to the audit trails and indexes of the connection. General data about the connections that is displayed on the **Search** page is archived and backed up as part of the system-backup process of SPS.

    - To timestamp, encrypt, or sign the audit trails, configure an **Audit Policy** to suit your needs.

      For more information, see Audit policies on page 512.

      CAUTION:

      In RDP connections, if the client uses the Windows login screen to authenticate on the server, the password of the client is visible in the audit trail. To avoid displaying the password when replaying the audit trail, encrypt the upstream traffic in the audit trail using a separate certificate from the downstream traffic.

      For more information, see *Encrypting audit trails* in the *Administration Guide*.

    - Require the users to authenticate themselves not only on the target server, but on SPS as well.

      For more information, see Configuring gateway authentication on page 862.

- Require four-eyes authorization on the connections, with the possibility of an auditor monitoring the connection in real-time.

  For more information, see Configuring four-eyes authorization on page 870.

- In the case of certain connections and scenarios (for example SSH authentication, gateway authentication, Network Level Authentication (NLA) connections), SPS can authenticate you to an LDAP database, or retrieve your group memberships. To use these features, select an **LDAP Server**.

  For more information, see Authenticating users to an LDAP server on page 505.

  NOTE: To display the usergroups that can access a specific Connection Policy, open the Connection Policy, then on the **Connections** page, select **Show connection permissions** > **Show**.

- To limit the number of new connection requests accepted from a single client IP address per minute, in the **Connection rate limit** field, enter the maximum number of accepted connections.

- If you have joined an SPP appliance to SPS and want to share specific SPS functions with SPP, use the **Functions shared with SPP** option.

  For more information, see Sharing SPS functions with SPP.

  To share an RDP or an SSH connection policy with SPP to initiate sessions, select **Share connection policy with SPP**.

  For more information, see sections Sharing RDP connection policies with SPP and Sharing SSH connection policies with SPP.

NOTE: Protocol-specific configuration options are described in their respective sections:

- HTTP-specific settings on page 541
- ICA-specific settings on page 559
- RDP-specific settings on page 577
- SSH-specific settings on page 616
- Telnet-specific settings on page 651
- VNC-specific settings on page 664

11. If your clients and servers support it, configure the connection to use strong encryption.
    - For HTTP connections, see Enabling TLS encryption in HTTP on page 549.
    - For Citrix ICA connections, use the following scenario: Client - Broker - original secure gateway - Secure Ticket Authority (STA) - SPS - Server.
    - For RDP connections, see Enabling TLS-encryption for RDP connections on page 592.
    - For SSH connections, see Creating and editing protocol-level SSH settings on page 638.

- For Telnet connections, see Enabling TLS-encryption for Telnet connections on page 652.
- For VNC connections, see Enabling TLS-encryption for VNC connections on page 665.

12. For graphical connections, adjust the settings of your servers for optimal performance:

    - **CAUTION:**

      For optimal performance and text recognition in graphical protocols, disable antialiasing on your servers. Antialiased text in the audit trails of RDP, VNC, and X11 connections is not recognized by the OCR engine of the Audit Player. The indexer service recognizes antialiased text, but its accuracy depends on the exact antialiasing settings. To properly index the trails of these connections, disable antialiasing.

      Note that by default, antialiasing is enabled on Windows Vista and later versions. Antialiasing is also called font smoothing. To optimize performance, disable ClearType, which is an antialiasing technology used on Microsoft Windows.

    - When processing RDP connections, SPS attempts to extract the username from the connection.

      To ensure that your users can access the target servers only when their username is recorded, see Usernames in RDP connections on page 603.

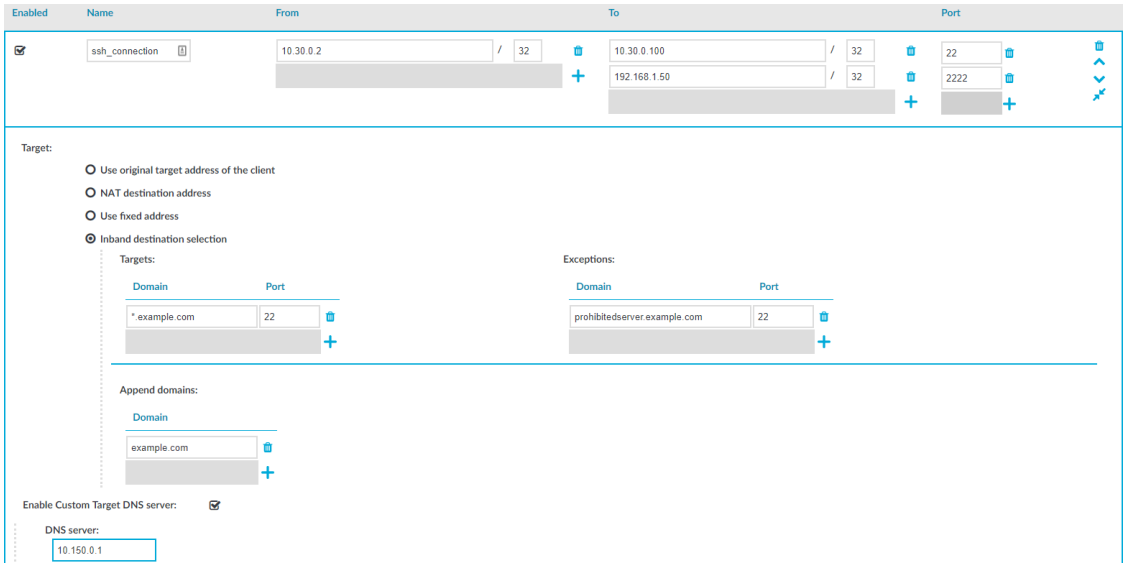
Modifying the destination address

The destination address is the address of the server where the clients finally connect to.

To modify the destination address of a connection

1. Navigate to the **Connections** tab storing the connection and click to display the details of the connection.

Figure 159: Traffic Controls > Protocol name > Connections — Configuring connections



2. The **Target** section allows you to configure Network Address Translation (NAT) on the server side of One Identity Safeguard for Privileged Sessions (SPS). Destination NAT determines the target IP address of the server-side connection. Set the destination address as required. The following options are available:

   NOTE: It is not possible to direct the traffic to the IP addresses belonging to SPS.

   - **Use the original target address of the client**: Connect to the IP address targeted by the client. This is the default behavior in transparent mode. This option is not available in non-transparent mode. For HTTP connections, you can use the **Use the original target address of the client** option only when the **Act as HTTP proxy** option is disabled.
   - **NAT destination address**: Perform a network address translation on the target address. Enter the target address in IP address/Prefix format.

     Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

     NOTE: Note the following limitations:

     - To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
     - If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.
   - **Use fixed address**: Enter the IP address and port number of the server. The connection will connect always to this address, redirecting the clients to the server.

Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

NOTE: Note the following limitations:

- To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
- If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.

- **Inband destination selection**: Extract the address of the server from the username. Note that for HTTP connections, you can use the **Inband destination selection** option only when the **Act as HTTP proxy** option is enabled. For details, see Configuring inband destination selection on page 490.

3. *Optional Step:* to enable a custom DNS server to be used for target selection in server-side Channel Policies, select **Enable Custom Target DNS server**, then enter the IP address of the custom DNS server to look up target addresses and resolve FQDN or wildcard FQDN addresses in the **Target** fields of your Channel Policies.
4. Click **Commit**.

Configuring inband destination selection

With inband destination selection, you can create a single connection policy and allow users to access any server by including the name of the target server in their username (for example, `ssh username@targetserver@scb_address`, or `username%@targetserver%scb_address`). If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays a terminal prompt where the user can enter the username and the server address.

Prerequisites

- Inband destination selection is not available for Virtual Networking (VNC) connections.

  NOTE:

  When using inband destination selection and TN3270 pattern sets in a connection, only destinations that are consistent with the specified pattern set will work.

- To use inband destination selection in HTTP connections, you must enable the **Act as HTTP proxy** option. For details, see Enabling One Identity Safeguard for Privileged Sessions (SPS) to act as an HTTP proxy on page 547.
- To use inband destination selection with RDP connections, it is recommended to use SPS as a Remote Desktop Gateway (or RD Gateway). For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on

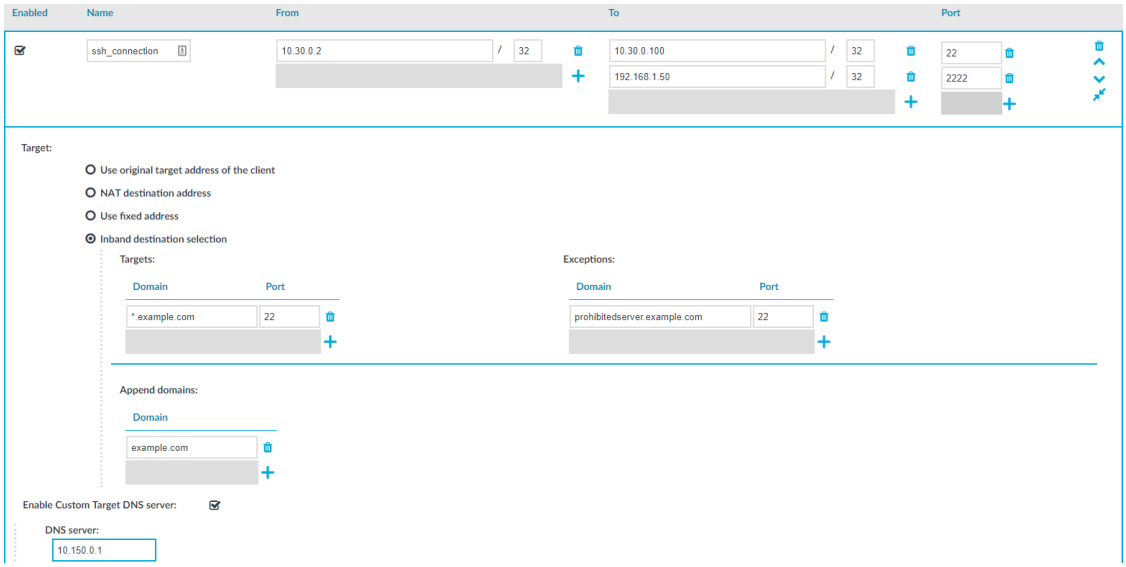
page 595.

- To use inband destination selection with RDP connections without using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway (or RD Gateway), you must use SSL-encrypted RDP connections (see Enabling TLS-encryption for RDP connections on page 592).
- For details on setting the clients to use inband destination selection in SSH connections, see Using inband destination selection in SSH connections on page 916.
- For details on setting the clients to use inband destination selection in Telnet connections, see Inband destination selection in Telnet connections on page 661.

To configure a Connection Policy to extract the address of the server from the username

1. Navigate to the Connection policy you want to modify, for example, to **Traffic Controls** > **SSH** > **Connections**.
2. Select **Inband destination selection**.

Figure 160: Traffic Controls > Protocol name > Connections — Configuring inband destination selection



3. *Optional Step*: Enter the IP address or the hostname of the domain name server used to resolve the address of the target server into the **DNS Server** field.

   If you do not set the **DNS Server** field, SPS will use the global DNS server (set on the **Basic Settings** > **Networking** page) to resolve the hostnames in this connection.
4. *Optional Step*: Configure domain names and CNAME records.

   If the clients do not include the domain name when addressing the server (for example they use `username@server` instead of `username@server.example.com`, or `username%server` for RDP connections), SPS can automatically add domain

information (for example, `example.com`). Enter the domain name to add into the **Append domain** field.

SPS can also resolve CNAME records.

To enter more domain names (for example, because connections extend through subnets), click  . In case of more domain names in the **Append domain** field, SPS appends the first domain name in the list that the target can be resolved with.

5. Enter the addresses of the servers that the users are permitted to access into the **Targets** field. Note the following points:
   - Use the IP address/prefix (for example **`192.168.2.16/32`**, or **`10.10.0.0/16`**) format. Alternatively, you can use the FQDN of the server. To permit access to any server, enter **`*`**.
   - For FQDN, you can use the **`*`** and **`?`** wildcard characters.

     CAUTION:

     If only the hostname of the server is listed and the client targets the server using its IP address, SPS refuses the connection.

     CAUTION:

     When the client uses hostname in inband destination selections, the hostname must comply with RFC5890: Internationalized Domain Names for Applications (IDNA). For example, from the ASCII characters only letters, digits, and the hyphen character is permitted.

     Starting with version 6.1.0, SPS rejects connection requests where the hostname does not comply with RFC5890.
   - If the clients target the server using its IP address, include the IP address of the server in the **Targets** > **Domain** list. This is required because SPS resolves the hostnames to IP addresses, but does not reverse-resolve IP addresses to hostnames.
   - If the clients target the server using its hostname, then the `hostname-from-the-client-request + the-value-of-the-Append-domain-option` must appear in the **Targets** > **Domain** list. Alternatively, you must include the IP address of the `hostname-from-the-client-request + the-value-of-the-Append-domain-option` host.

Example: Hostnames and inband destination selection

For example, you have set **Append domain** to **`example.com`**, and your clients use the `username%servername` request, then you must include either the `servername.example.com` host or its IP address in the **Targets** > **Domain** list.

6. If the clients can access only a specified port on the server, enter it into the **Port** field. If the **Port** is not set, the clients may access any port on the server.
7. If there are any servers that the users cannot target using inband destination selection, add them to the **Exceptions** field.
8. To use inband destination selection with RDP connections without using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway (or RD Gateway), you must use SSL-encrypted RDP connections (see Enabling TLS-encryption for RDP connections on page 592).
9. Click **Commit**.

Expected result

The connection policy will extract the address of the destination server from the protocol information.

> NOTE: For examples on using inband destination selection to establish an SSH connection, including scenarios where non-standard ports or gateway authentication is used, see Using inband destination selection in SSH connections on page 916.

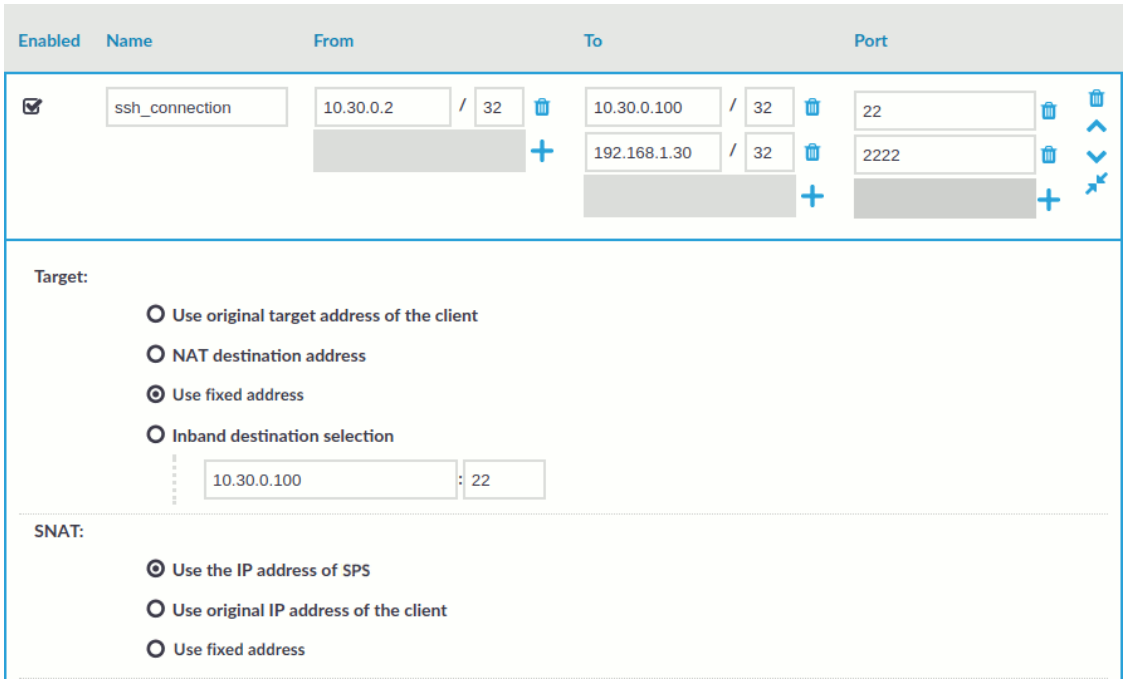
Modifying the source address

The source address is the address that One Identity Safeguard for Privileged Sessions (SPS) uses to connect the server. The server sees this address as the source of the connection.

To modify the source address of a connection

1. Navigate to the **Connections** tab storing the connection and click  to display the details of the connection.

Figure 161: Traffic Controls > Protocol name > Connections — Configuring connections



2. The **SNAT** section allows you to configure Source Network Address Translation (SNAT) on the server side of SPS. SNAT determines the IP address SPS uses in the server-side connection. The target server will see the connection coming from this address. The following options are available:
   - **Use the IP address of a SPS logical interface**: Server-side connections will originate from SPS's logical network interface. This is the default behavior of the connection.
   - **Use the original IP address of the client**: Server-side connections will originate from the client's IP address, as seen by SPS.
   - **Use fixed address**: Enter the IP address that will be used as the source address in server-side connections.

     Alternatively, you can enter a hostname instead. SPS automatically resolves the hostname to an IP address.

     NOTE: Note the following limitations:
     - To resolve the hostnames, SPS uses the Domain Name Servers set in the **Basic Settings**> **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields.
     - If the Domain Name Server returns multiple IP addresses, SPS randomly selects from the list.

CAUTION:

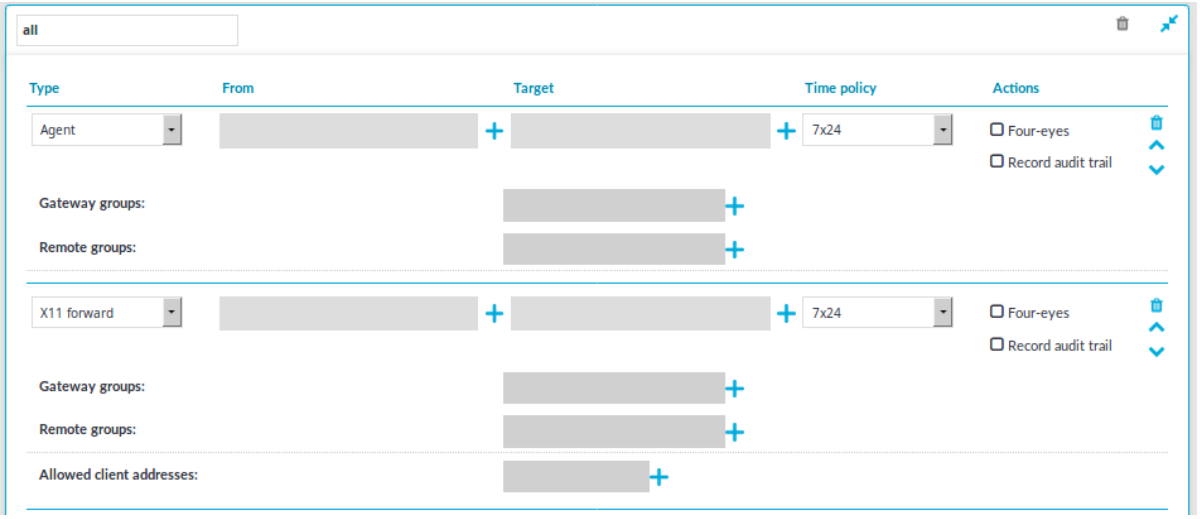
Do not forget to properly configure routers and other network devices when using the **Use fixed address** option: messages sent by the server to this address must reach SPS.

3. Click **Commit**.

Creating and editing channel policies

The Channel Policy lists the channels (for example, terminal session and SCP in SSH, or Drawing and Clipboard in RDP) that can be used in the connection, and also determines if the channel is audited or not. The Channel Policy can also restrict access to each channel based on the IP address of the client or the server, a user list, user group, or a time policy. For example, all clients may access the servers defined in a connection via SSH terminal, but the channel policy may restrict SCP access only to a single client. The policies set in the Channel Policy are checked when the user attempts to open a particular channel type in the connection.

Figure 162: Traffic Controls > Protocol name > Channel Policies — Configuring channel policies



To create a new channel policy or edit an existing one

1. Channel policies are configured individually for every protocol. Navigate to the **Channel Policies** tab of the respective protocol (for example, **Traffic Controls** > **SSH** > **Channel Policies**) and click + to create a new channel policy. Enter a name for the policy (for example, `shell_and_backup`).
2. Click + to add a new channel.

3. Select the channel to be enabled in the connection from the **Type** field. All restrictions set in the following steps will be effective on this channel type. The available channels are different for every protocol. For their descriptions, see the following sections:
    - Supported HTTP channel types on page 541 for the HTTP protocol.
    - Supported ICA channel types on page 560 for the Independent Computing Architecture protocol.
    - Supported RDP channel types on page 579 for the Remote Desktop protocol.
    - Supported SSH channel types on page 619 for the Secure Shell protocol.
    - The Telnet protocol has only one channel type with no special configuration options.
    - The VNC protocol has only one channel type with no special configuration options.
4. To restrict the availability of the channel only to certain clients, click + in the **From** field and enter the IP address of the client allowed to use this type of the channel. Repeat this step until all required client IP addresses are listed.

    Alternatively, you can also enter a hostname instead. One Identity Safeguard for Privileged Sessions (SPS) saves the hostname and resolves it when opening channels, therefore SPS can trace dynamic IP addresses.

    NOTE: Note the following limitations:
    - The Domain Name Servers you set must be able to resolve the hostnames you enter into the **From** and **Target** fields, otherwise this function (and, therefore, the sessions using this Channel Policy) will not work.
    - SPS Channel Policies support wildcard characters in the `*.example.com` format. If the channel opening request contains an IP address, SPS uses a reverse lookup method to resolve this IP address into a hostname for a match.
    - SPS uses the Domain Name Servers set in the **Basic Settings** > **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields to resolve the hostnames.
5. To restrict the availability of the channel only to certain servers, click + in the **Target** field and enter the IP address of the server allowed to use this type of the channel. Repeat this step until all required server IP addresses are listed.

    NOTE: Use the real IP address of the server, which may be different from the one addressed by the clients, specified in the **Target** field of the connection policy.

    Alternatively, you can also enter a hostname instead. One Identity Safeguard for Privileged Sessions (SPS) saves the hostname and resolves it when opening channels, therefore SPS can trace dynamic IP addresses.

    NOTE: Note the following limitations:

- The Domain Name Servers you set must be able to resolve the hostnames you enter into the **From** and **Target** fields, otherwise this function (and, therefore, the sessions using this Channel Policy) will not work.
- SPS Channel Policies support wildcard characters in the `*.example.com` format. If the channel opening request contains an IP address, SPS uses a reverse lookup method to resolve this IP address into a hostname for a match.
- SPS uses the Domain Name Servers set in the **Basic Settings** > **Network** > **Naming** > **Primary DNS server** and **Secondary DNS server** fields to resolve the hostnames.

Alternatively, you can configure a custom DNS server to be used for target selection. Select **Enable Custom DNS server** under the **Target** section of your connection policies (set under **Traffic Controls** > **Protocol name** > **Connections**) and enter the IP address of the custom DNS server.

6. To restrict the availability of the channel only to certain users, click + in the **Remote Group** field and enter the name of the user group allowed to use this type of the channel. Repeat this step until all permitted groups are listed.

   CAUTION:

   Adding more than approximately 1000 remote groups to a channel policy may cause configuration, performance, and authentication issues when connecting to LDAP servers.

   To restrict the availability of the channel when using gateway authentication, click + in the **Gateway Group** field and enter the name of the user group allowed to use this type of the channel. Repeat this step until all permitted groups are listed.

   You may list local user lists as defined in Creating and editing user lists on page 504, or LDAP groups (for details on accessing LDAP servers from SPS, see Authenticating users to an LDAP server on page 505). Note the following behavior of SPS:

   - If you list multiple groups, members of any of the groups can access the channel.

     NOTE: >When listing both a whitelist and blacklist in the **Remote Group** section and a username appears on both lists, the user will be able to access the channel.

   - If you do not list any groups, anyone can access the channel.

     NOTE: When the channel opens, there are certain cases when the remote group is not known yet. For example, in case of an RDP or ICA login screen, the drawing channel has to be opened first to properly display the logon screen. Only those channel rules will apply, where the **Remote group** field is empty. In case of network level authentication, all required information is present already so this limitation does not apply.

   - If a local user list and an LDAP group has the same name and the LDAP server is configured in the connection that uses this channel policy, both the members

of the LDAP group and the members of the local user list can access the channel.

NOTE: User lists and LDAP support is currently available only for the SSH and RDP protocols. For other protocols, see Configuring gateway authentication on page 862.

7. Select a time policy to narrow the availability of the channel. If the time policy of the channel policy is set to **7x24**, the channel is always available. For details, see Configuring time policies on page 503.
8. Some channel types require additional parameters, for example port forwarding in SSH needs the IP addresses and ports of the source and destination machines. Click + in the **Details** field and enter the required parameters. For a list of parameters used by the different channels, see Supported SSH channel types on page 619 and Supported RDP channel types on page 579.
9. Select the **Record audit trail** option to record the activities of the channel into audit trails. Typically large file-transfers (for example system backups, SFTP channels) are not audited because they result in very large audit trails. Check regularly the free hard disk space available on SPS if you do audit such channels. You can also receive alerts about disk space fill-up if you set these. For details, see Preventing disk space fill-up on page 138 and System related traps on page 139.
10. Select the **4 eyes** option to require four-eyes authorization to access the channel. For details, see Configuring four-eyes authorization on page 870.
11. Repeat Steps 2-10 to add other channels to the policy.

    NOTE: The order of the rules matters. The first matching rule will be applied to the connection. Also, note that you can add the same channel type more than once, to fine-tune the policy.
12. To save the list, click **Commit**.

Real-time content monitoring with Content Policies

You can monitor the traffic of certain connections in real time, and execute various actions if a certain pattern (for example, a particular command or text) appears in the command line or on the screen, or if a window with a particular title appears in a graphical protocol. Since content-monitoring is performed real-time, One Identity Safeguard for Privileged Sessions (SPS) can prevent harmful commands from being executed on your servers. SPS can also detect numbers that might be credit card numbers. The patterns to find can be defined as regular expressions. In case of ICA, RDP, and VNC connections, SPS can detect window title content.

The following actions can be performed:

- Log the event in the system logs.
- Immediately terminate the connection.

- Send an e-mail or SNMP alerts about the event.
- Store the event in the connection database of SPS.

SPS currently supports content monitoring in SSH session-shell connections, Telnet connections, RDP and Citrix ICA Drawing channels, and in VNC connections.

NOTE: Command, credit card and window detection algorithms use heuristics. In certain (rare) situations, they might not match the configured content. In such cases, contact our Support Team to help analyze the problem.

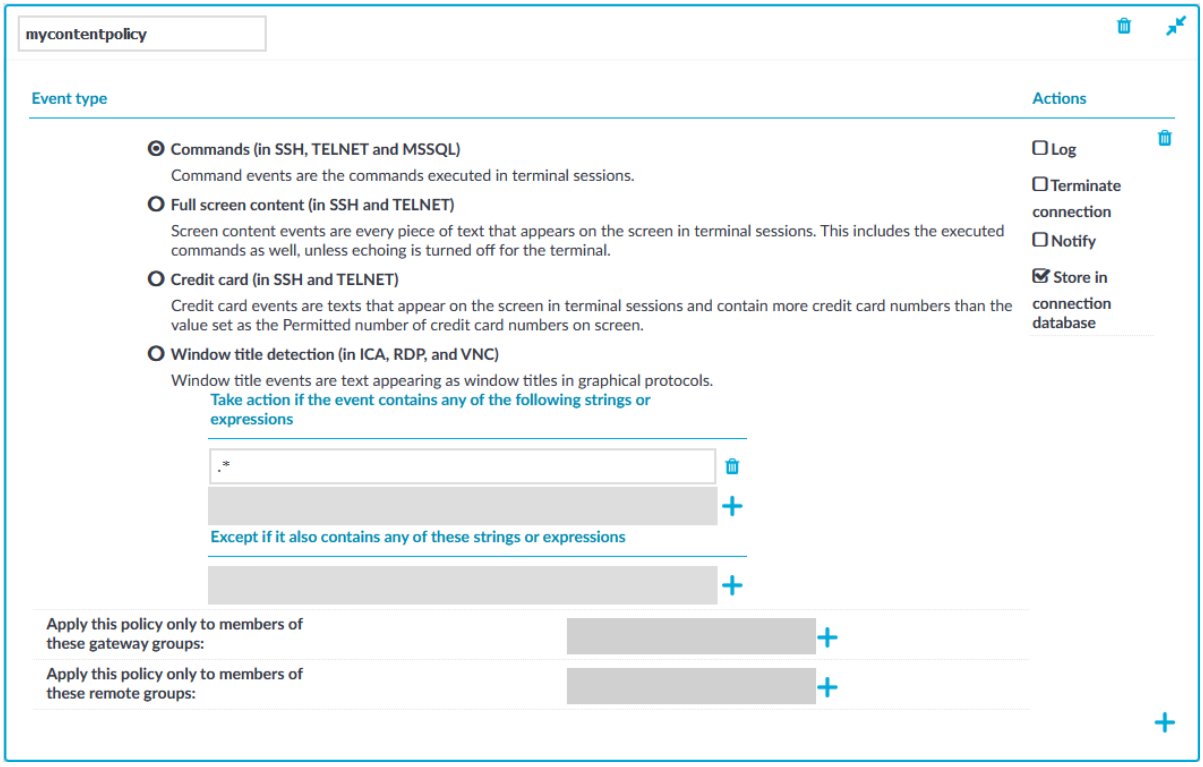
Real-time content monitoring in graphical protocols is not supported for Arabic and CJK languages.

Creating a new content policy

The following describes how to create a new content policy that performs an action if a predefined content appears in a connection.

NOTE: Using content policies significantly slows down connections (approximately 5 times slower), and can also cause performance problems when using the indexer service.

Figure 163: Policies > Content Policies — Content policies



To create a new content policy that performs an action if a predefined content appears in a connection

1. Navigate to **Policies** > **Content Policies**, click + and enter a name for the policy.
2. Select the type of event that you want to monitor:
   - **Commands**: The commands executed in the `session-shell` channel of SSH connections, or in Telnet connections.

     CAUTION:

     During indexing, if a separate certificate is used to encrypt the upstream traffic, command detection works only if the upstream key is accessible on the machine running the indexer.

     NOTE: Command detection is case-insensitive.
   - **Screen content**: Every text that appears on the screen. For example, every text that is displayed in the terminal of SSH or Telnet connections. This includes the executed commands as well, unless echoing is turned off for the terminal.
   - **Credit card**: Process every text that appears on the screen and attempt to detect credit card numbers in SSH or Telnet connections. One Identity Safeguard for Privileged Sessions (SPS) performs an action if the number of detected credit card numbers exceeds the value set as **Permitted number of credit card numbers**.

     Credit card number detection is based on the Luhn algorithm and lists of known credit card number prefixes.
   - **Window title detection**: Text appearing as window titles that can be detected on the screen in RDP, Citrix ICA, and VNC connections. Window title detection involves Optical Character Recognition (OCR) on parts of the screen, and can be slightly resource-intensive. SPS versions up till 6.2 only detected only the active window in the screen. From SPS version 6.3, multiple windows can be detected.

     ### Limitations

     - Default Windows themes are supported.
     - Windows that do not have an X (close window) button in the top-right corner (or it is not visible) are not detected.
     - Use window title detection for sessions that use a single monitor. The feature works in multi-monitor environments as well, but becomes very slow, therefore it is not recommended.
     - Window title detection is case-insensitive.
3. Select **Take action if the event contains any of the following strings or expressions**, click

+ and enter a string or regular expression. SPS will perform an action if this expression is found in the connection, unless it is listed in the **Except if it also contains any of these strings or expressions** list. For example, SPS can terminate the connection if the user issues the `rm -rf *` in an SSH connection. Repeat this step to add further expressions if needed.

- Use Perl Compatible Regular Expressions (PCRE).
- The following characters must be escaped using a backslash character: `'` (single-quote). For example, instead of `.*'` use `.*\'`
- SPS uses substring search to find the expression in the content. That is, SPS finds the expression even if there is more content before or after the matching part. For example, the `conf` pattern will match the following texts: `conf`, `configure`, `reconfigure`, `arcconf`, and so on.
- Using complicated regular expressions or using many regular expressions will affect the performance of SPS.
- If the multiple expressions are set, SPS processes them one after the other, and stops processing the content if the first match is found, even if other expressions would also match the content. Therefore, when using multiple expressions, start with the most specific one, and add general expressions afterward.

Example: Sample regular expressions for content policies

The following simple regular expressions are samples to demonstrate what kinds of events that can be detected using content policies.

- The `enable` command on Cisco devices: the user enters privileges mode.
- The `conf term` command on Cisco devices: the user configures the networking parameters of the device.
- The `sudo` and `su -` commands: the user enters privileged mode Linux and other UNIX platforms.

4. To add an exception to the **Take action if the event contains any of the following strings or expressions** rule, select **Except if it also contains any of these strings or expressions**, click + and enter a string or regular expression. SPS will not perform any action if this expression is found in the connection. For example, to permit the users to delete only the `/tmp` directory in an SSH connection, enter `rm -rf /tmp`. Repeat this step to add further expressions if needed.

Example: Sample content policies using Ignore rules

The following expressions can be used to perform an action if any SQL command is used in MySQL, except for the `select` and `help` commands:

- Into the **Take action if the event contains any of the following strings or expressions** expression, enter `mysql>.*`
- Add two **Except if it also contains any of these strings or Except if it also contains any of these strings or Except if it also contains any of these strings or expressions** expressions: `mysql> select.*` and `mysql> help.*`

5. Select the action to perform.
   - **Log**: Send a log message into the system logs. The log message includes the expression that matched the content. On log level 6, the message includes the matching content as well.
   - **Terminate connection**: Immediately terminate the connection. When using the **Terminate connection** action for the **Command** event type, and a command matches an expression, the connection is terminated before the command is executed. When using the **Terminate connection** action, note the following points.
     - Select the **Log** or **Notify** action as well so that it is easy to find out why a connection was terminated.
     - If the connection is terminated by a content policy, the **Verdict** of the connection becomes `ACCEPT-TERMINATED`.
   - **Notify**: Send an e-mail or SNMP alert about the event. To configure the alerts, navigate to **Basic Settings** > **Alerting & Monitoring** and set the required alerts for the **Real time audit event detected** (`scbAuditRealTime`) event.
   - **Store in connection database**: Add the event to the SPS connection database. These events are displayed in the **Alerts** column of the **Sessions** page. If the column is not visible, click **Customize columns...**.
6. To apply the content policy only for users belonging to specific groups, select **Apply this policy only to members of these gateway groups** or **Apply this policy only to members of these remote groups**, + and specify the usergroups as needed. If **Apply this policy only to members of these gateway groups** or **Apply this policy only to members of these remote groups** is set, the content policy is applied only to connections of these usergroups.
7. To add a new rule to the policy, click + and repeat Steps 2-6.

   Note that if you have more than one rules in a policy, SPS evaluates them as follows.

a. SPS evaluates the first (top) rule.

b. If the rule contains **Apply this policy only to members of these gateway groups** or **Apply this policy only to members of these remote groups** restrictions, SPS checks if the current user belongs to any of the specified groups. If the groups do not match, SPS skips the rule.

c. If the content matches any entry of the **Except if it also contains any of these strings or expressions** list, SPS skips the rule.

d. If the content matches any entry of the **Take action if the event contains any of the following strings or expressions** list, SPS performs the action configured for the rule. Otherwise, SPS skips the rule.

e. If the current rule did not match the content, SPS evaluates the next rule of the policy (if any).

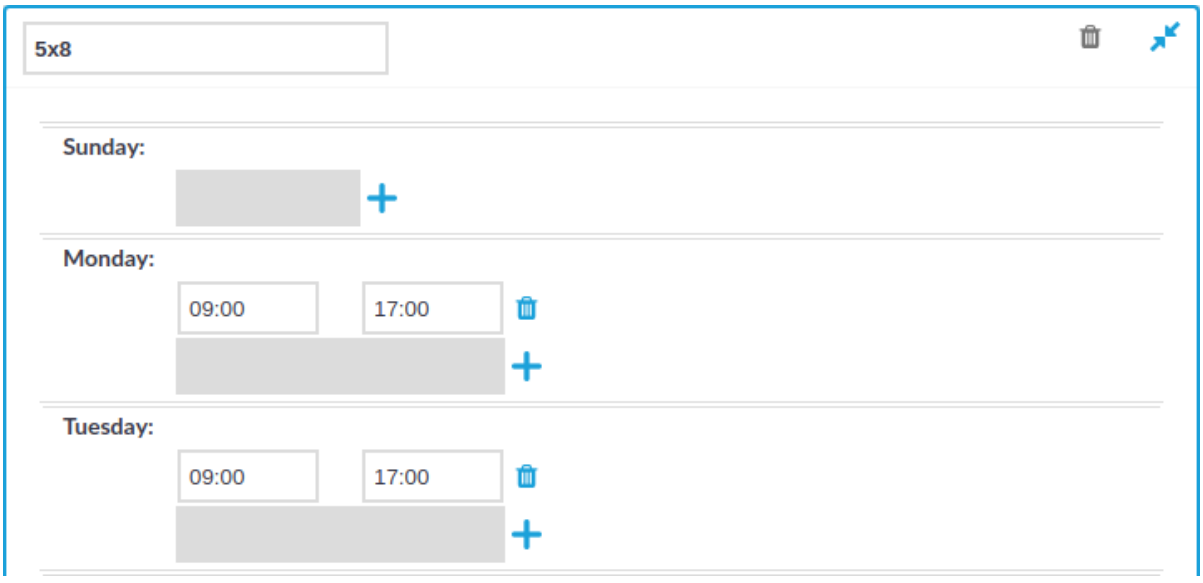
8. Click **Commit**. A new content policy is created.
9. To use the content policy created in the previous steps, select the policy in the channel policy that is used to control the connections.

   NOTE: It is not required to enable auditing to use content policies.

Configuring time policies

The time policy determines the timeframe when the users are permitted to access a particular channel. By default, there is no time-based restriction, all channels are available 7x24.

Figure 164: Policies > Time Policies — Configuring time policies



To create a time policy or edit an existing one

1. Navigate to the **Time Policies** tab of the **Policies** menu item and click + to create a new time policy. Enter a name for the policy (for example **workhoursonly**).
2. Click ⤢ to display the days of the week and the allowed intervals.
3. Enter the intervals for each day when the users are allowed to access the connection. Use the `hh:mm` format (for example from 08:00 to 16:00).
4. To add multiple intervals for a day, click +.
5. Click **Commit**.
6. To actually restrict access to a connection or a channel based on the policy created in the previous steps:
   - Select this policy in the **Time Policy** field of the channel policy.
   - Click **Commit**.

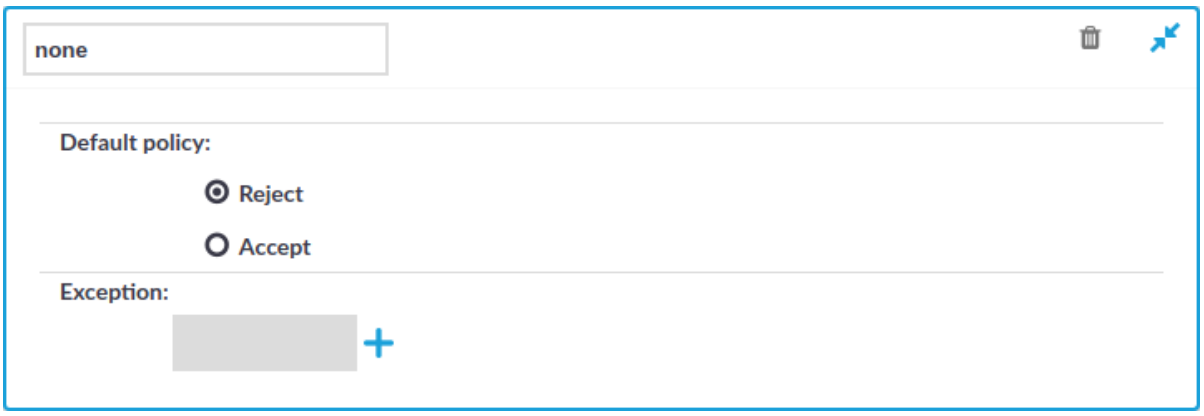
Creating and editing user lists

User lists are white- or blacklists of usernames that allow fine-control over who can access a connection or a channel.

CAUTION:

User Lists are white- or blacklists of usernames that determine who can access the server remotely. However, this cannot prevent a user from accessing the server from a local terminal.

Figure 165: Policies > User Lists — Configuring user lists



To create a new user list or edit an existing one

1. Navigate to the **User Lists** tab of the **Policies** menu and click  to create a new user list. Enter a name for the list **User List** field (for example **`serveradmins`**).

   CAUTION:

   Usernames, the names of user lists, and the names of usergroups are case sensitive.

2. Click  to display the list of users.
3. Select the default policy of the user list. Select **Reject** for a whitelist, that is, to allow access only to the members of the list. Select **Accept** for a blacklist, that is, to allow access to everyone except the members of the list.
4. Click  and enter a username into the displayed field. Repeat this step until all required usernames are listed.

   CAUTION:

   Usernames, the names of user lists, and the names of usergroups are case sensitive.

5. To save the list, click **Commit**.
6. To actually restrict access to a channel based on the user list created in the previous steps:
   - Navigate to the **Channel Policies** tab of the type of connection you want to control and click  to display the details of the policy.
   - Click  in the **Group** section to add a new group to the policy and enter the name of the group. Repeat this step to add other groups.

     CAUTION:

     Usernames, the names of user lists, and the names of usergroups are case sensitive.

     NOTE: When listing more groups, users of any of the listed groups can access the channel. For details, see Creating and editing channel policies on page 495.

     When listing both a whitelist and blacklist in the **Group** section and a username appears on both lists, the user will be able to access the channel.
   - Click **Commit**.

Authenticating users to an LDAP server

You can use the LDAP policy to set the details of the LDAP server you wish to use to:

- authenticate gateway users (available in SSH and Telnet as Authentication Policy)
- query gateway groups (available for RDP, Telnet, SSH, and ICA)
- query remote groups (available for RDP, Telnet, SSH, ICA, and HTTP)

NOTE: This feature is not available for Virtual Network Computing (VNC) connections.

Prerequisites

Make sure that the response timeout of the LDAP/Active Directory server is set to a minimum of 120 seconds.

NOTE: Consider the following:

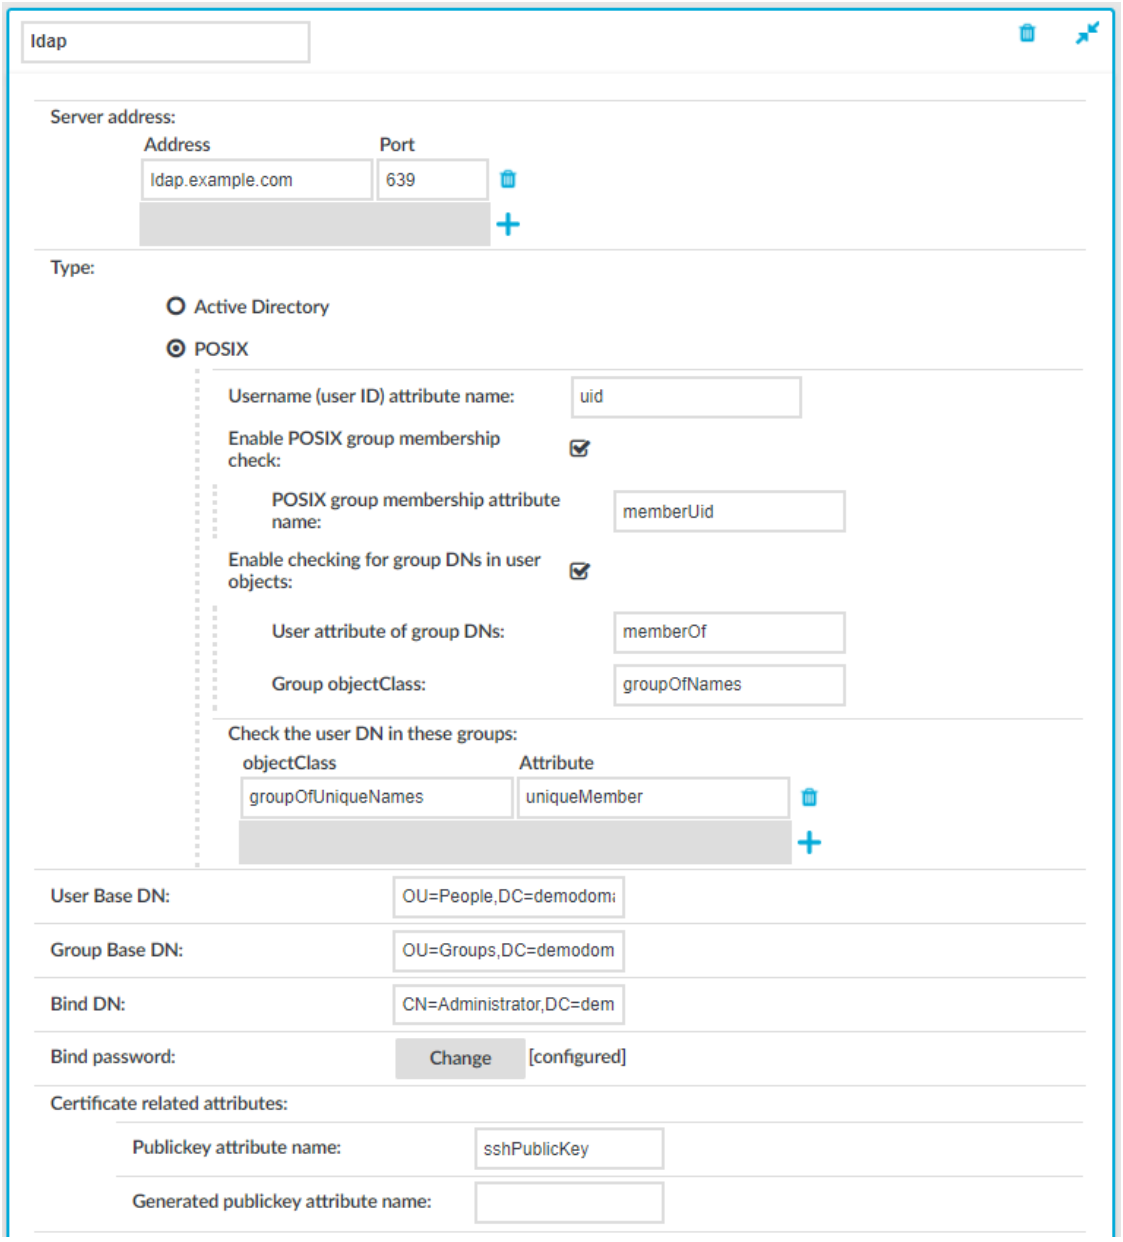
- In RDP (including RDG) connections, you can use the LDAP policy for group membership check only, you cannot use it as the authentication backend. However, you can use a trusted AD domain for authentication and LDAP for group membership check.

  In this case, LDAP will only use the username without the domain name to verify the group membership.
- One Identity Safeguard for Privileged Sessions (SPS) treats user and group names in a case insensitive manner if the matching rule for the attribute in question is case insensitive in the LDAP database.

To configure an LDAP policy for a connection

1. Navigate to **Policies** > **LDAP Servers** and click + to create a new LDAP policy.

Figure 166: Policies > LDAP Servers — Configuring LDAP Server policies



2. Enter a name for the policy (for example `ldapservers`).
3. Select the type of your LDAP server in the **Type** field. Select:

- **Active Directory** to connect to Microsoft Active Directory servers.

  You can enable nested groups. Select **Enable AD group membership check**, then **Enable nested groups**.

  > **CAUTION:**
  >
  > Nested groups can slow down the query and cause the connection to timeout if the LDAP tree is very large. In this case, disable the **Enable nested groups** option.

  To also check group membership based on group DNs in a user attribute, select **Enable checking for group DNs in user objects** and enter the name of the user attribute, for example, `memberOf` in the **User attribute of group DNs** field.

  > **CAUTION:**
  >
  > If you have too many groups, using this option significantly slows down logging in to the SPS web interface.
  >
  > Use this option only if you have an LDAP schema where the user groups can only be determined from a user attribute that contains the group DNs.

  To check for group membership based on user DNs in group attributes, use the **Check the user DN in these groups** options.

- **POSIX** to connect to servers that use the POSIX LDAP scheme.

  If your LDAP server uses a custom POSIX LDAP scheme, you might need to set which LDAP attributes store the username, or the attributes that set group memberships. For example, if your LDAP scheme does not use the `uid` attribute to store the usernames, set the **Username (user ID) attribute name** option.

  In addition to the POSIX LDAP backend membership checking, you can allow checking for supplementary group memberships by selecting **Enable POSIX group membership check** and specifying the **POSIX group membership attribute name** field.

  To also check group membership based on group Distinguished Names (DNs) in a user attribute, select **Enable checking for group DNs in user objects**. Then, enter the name of the user attribute (for example, `memberOf`) in the **User attribute of group DNs** field, and objectClass (for example, `groupOfNames`) in the **Group objectClass** field.

  > **CAUTION:**
  >
  > If you have too many groups, using this option significantly slows down logging in to the SPS web interface.
  >
  > Use this option only if you have an LDAP schema where the user groups can only be determined from a user attribute that contains the group DNs.

To check for group membership based on user DNs in group attributes, use the **Check the user DN in these groups** option.

4. Enter the IP address/hostname and the port of the LDAP server in the respective text boxes.

   Consider the following when specifying the address information:

   - If you want to encrypt the communication between SPS and the LDAP server, use the following port numbers:
     - For **TLS**, specify 636 as the port number.
     - For **STARTTLS**, specify 389 as the port number.
   - Use an IPv4 adress or a hostname.
   - To add multiple servers, click  and enter the address of the next server. If a server is unreachable, SPS will try to connect to the next server in the list in failover fashion.
   - When you configure the location of the LDAP server, that is, the IP address or hostname and the port number, you can use a Service record (SRV record), which is a type of information record in the DNS that maps the name of a service to the DNS name of the server. SRV records have the following format: `_ldap._tcp.<SITE_NAME>._sites.dc._msdcs.<DOMAIN.NAME>` in the **Address** field. SPS looks up of the SRV record during committing the configuration change.

     For more information on SRV records, see the relevant Microsoft documentation.
   - **CAUTION:**

     If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.

5. In the **User Base DN** field, enter the name of the DN to be used as the base of queries regarding users (for example: **`OU=People,DC=demodomain,DC=exampleinc`**).

   NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

   To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

6. In the **Group Base DN** field, enter the name of the DN to be used as the base of queries regarding groups (for example: **`OU=Groups,DC=demodomain,DC=exampleinc`**).

   NOTE: This field is mandatory. You can use the same value for the **User Base DN** and the **Group Base DN** settings.

   To speed up LDAP operations, specify a sufficiently narrow base for the LDAP subtrees where users and groups are stored.

7. In the **Bind DN** field, enter the Distinguished Name that SPS must use to bind to the LDAP directory (for example: **CN=Administrator,DC=demodomain,DC=exampleinc**).

   NOTE: SPS accepts both pre Windows 2000-style and Windows 2003-style account names, or User Principal Names (UPNs). For example, **administrator@example.com** is also accepted.

8. To configure or change the password to use when binding to the LDAP server, click **Change** and enter the password. Click **Update**. Click **Commit**.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

9. *Skip this step if you use passwords to authenticate the users.*

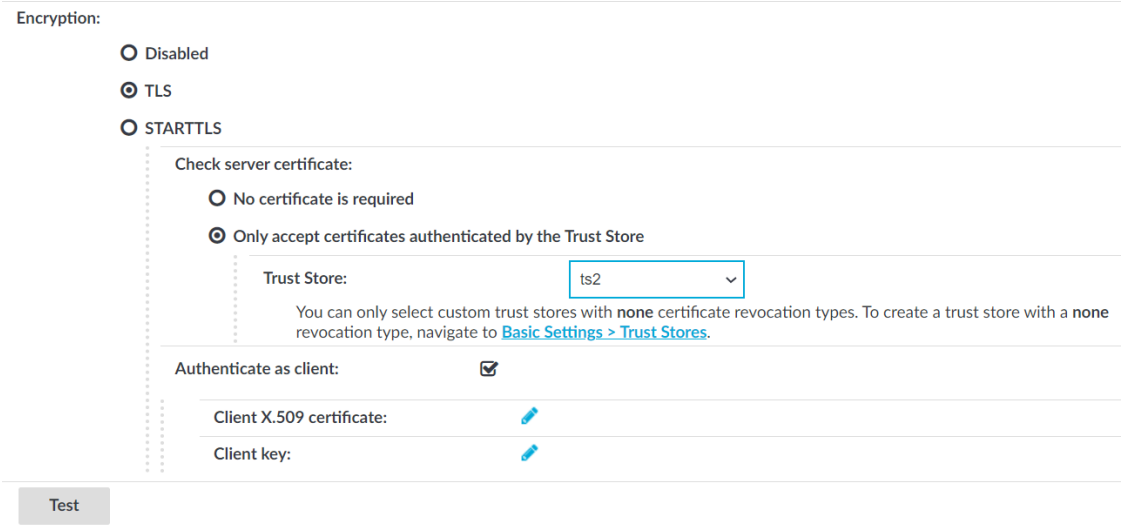
   - If you use public-key authentication and receive the public key of the users from the LDAP database, enter the name of the LDAP attribute that stores the public keys of the users into the **Publickey attribute name** field. For details on using public-key authentication with the LDAP database, see Configuring public-key authentication on One Identity Safeguard for Privileged Sessions (SPS) on page 911.
   - If you use X.509 certificate for authentication and receive the certificates of the users from the LDAP database, enter the name of the LDAP attribute that stores the certificates of the users into the **Certificate attribute name** field.

10. *Skip this step if you use passwords to authenticate the users.*

    - If you use public-key authentication and want SPS to generate server-side encryption keys on-the-fly and store them in a separate attribute on the LDAP server, enter the name of the attribute into the **Generated publickey attribute name** field.
    - If you use certificate authentication and want SPS to generate server-side certificates on-the-fly and store them in a separate attribute on the LDAP server, enter the name of the attribute into the **Generated certificate attribute name** field.

11. If you want to encrypt the communication between SPS and the LDAP server, in **Encryption**, select the **TLS** or the **STARTTLS** option and complete the following steps:

Figure 167: Policies > LDAP Servers — Configuring encryption



Verify the certificate of the server

- If you want SPS to verify the certificate of the server, select **Only accept certificates authenticated by the Trust Store** and select a trust store in the **Trust Store** field.

  SPS will use the selected trust store to verify the certificate of the server, and reject the connections if the verification fails.

  CAUTION:

  SPS checks if the certificate revocation list (CRL) has expired and that the CRL has been signed by the same certificate authority (CA).

  CAUTION:

  If you connect to the LDAP server over a TLS-encrypted connection with certificate verification, you must fill the `Address` field with a name or IP address, which must be present in the certificate.

  Authenticate as client

- If the LDAP server requires mutual authentication, that is, it expects a certificate from SPS, enable **Authenticate as client**. Generate and sign a certificate for SPS, then click in the **Client X.509 certificate** field to upload the certificate. After that, click in the **Client key** field and upload the private key corresponding to the certificate.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

12. To commit the changes, click **Commit**.
13. Click **Test** to test the connection.

NOTE: Testing TLS and STARTTLS-encrypted connections is not supported.

Audit policies

An audit trail is a file storing the recorded activities of the administrators. Audit trails are not created automatically for every connection: auditing must be enabled manually in the channel policy used in the connection. The available default channel policies enable auditing for the most common channels. Audit trails are automatically compressed, and can be encrypted, timestamped, and signed as well. Audit trails can be replayed using the Safeguard Desktop Player application (for details, see *Safeguard Desktop Player User Guide*), or directly in your browser (for details, see Replaying audit trails in your browser on page 839).

TIP: By default, every connection uses the built-in **default** audit policy. Unless you use a custom audit policy, modifying the **default** audit policy will affect every audited channel of the connections passing through One Identity Safeguard for Privileged Sessions (SPS).

CAUTION:

In RDP connections, if the client uses the Windows login screen to authenticate on the server, the password of the client is visible in the audit trail. To avoid displaying the password when replaying the audit trail, encrypt the upstream traffic in the audit trail using a separate certificate from the downstream traffic.

For more information, see *Encrypting audit trails* in the *Administration Guide*.

- For details on how to configure audit trail encryption, see Encrypting audit trails on page 512.
- For details on how to configure timestamping, see Timestamping audit trails with built-in timestamping service on page 515 and Timestamping audit trails with external timestamping service on page 517.
- For details on how to configure audit trail signing, see Digitally signing audit trails on page 519.

Encrypting audit trails

To prevent unauthorized access to the audit trail files, One Identity Safeguard for Privileged Sessions (SPS) can encrypt:

- The entire trail.
- The entire trail, and the upstream part with an additional (set of) certificate(s).
- Only the upstream part.

With upstream encryption, the passwords are visible only with the private key of the certificate used for encrypting the upstream traffic.

NOTE: Even if the upstream traffic is encrypted with a separate certificate, only one audit trail file is created for a session.

Encrypting the upstream part has the following limitations:

- During indexing, command detection does not work without the upstream encryption keys.

TIP: For more information on uploading certificates for indexing and replaying audit trails, see:

- Configuring the internal indexer on page 675 and Replaying encrypted audit trails in your browser on page 847 for uploading certificates for the indexer service.
- Replaying encrypted audit trails in your browser on page 847 for uploading certificates to a user's private keystore.
- *Replaying encrypted audit trails* in the *Safeguard Desktop Player User Guide* for uploading certificates to the Safeguard Desktop Player application.

Encrypting audit trails requires one or more X.509 certificate in PEM format that uses an RSA key, depending on the configuration.

NOTE: Certificates are used as a container and delivery mechanism. For encryption and decryption, only the keys are used.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates (for example, do not use the certificate of the One Identity Safeguard for Privileged Sessions (SPS) webserver to encrypt audit trails, or the same keypair for signing and encrypting data).

The following encryption options are available:

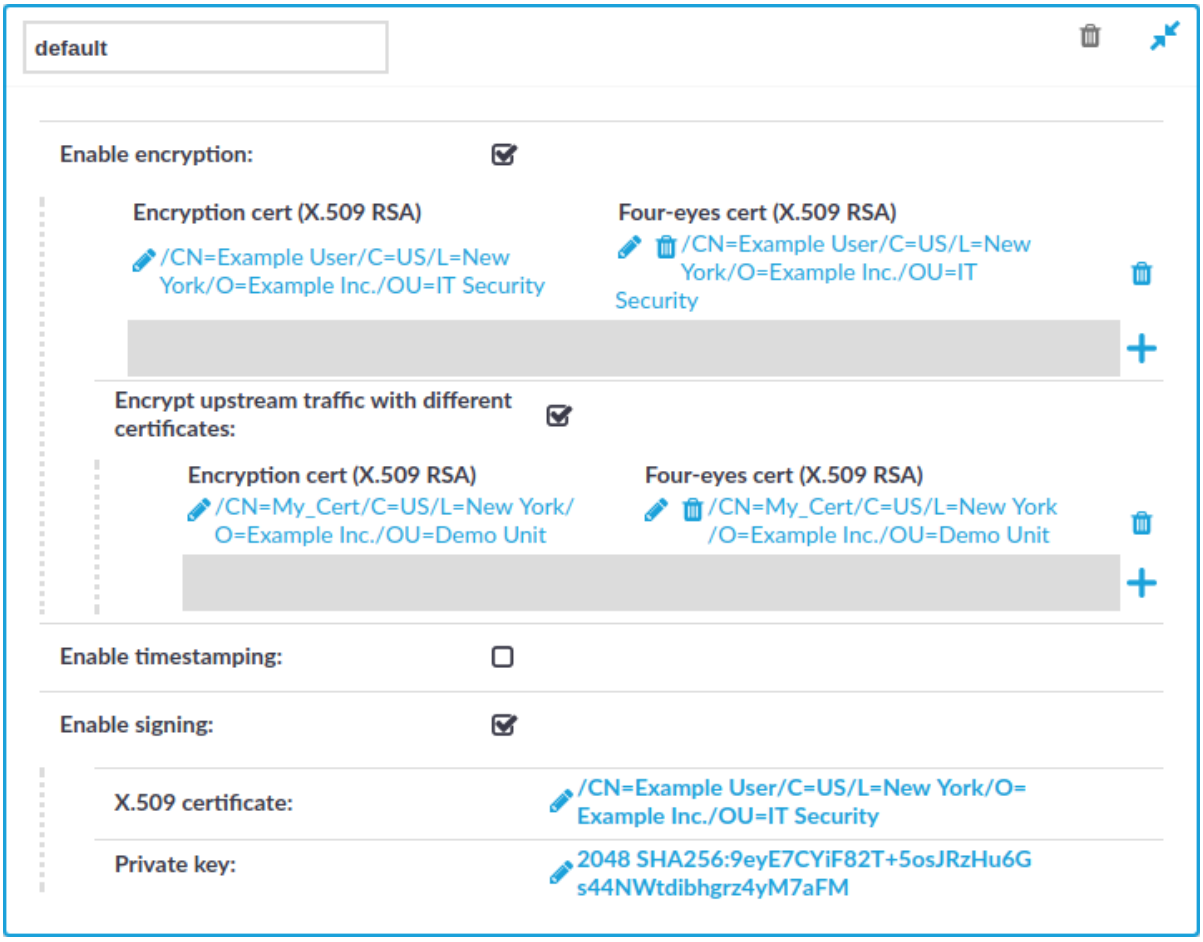
- *Encrypt with a single certificate.* This is the most simple approach: SPS uses one certificate to encrypt the audit trails, and anyone who has the private key of that certificate can replay the audit trails. If that key is lost, there is no way to open the audit trails.
- *Encrypt separately with multiple certificates.*SPS uses two or more certificates separately to encrypt the audit trails, and anyone who has the private key of one of the encryption certificates can replay the audit trails.
- *Encrypt jointly with two certificates.*SPS uses two certificates together (a certificate-pair) to encrypt the audit trails. The private keys of both encryption certificates are needed to replay the audit trails. This is a kind of "four-eyes in auditing".

You can combine the different encryption methods. For example, you can encrypt the audit trails with multiple certificate-pairs, and replay the trails only if the private keys of a certificate-pair are available. This is true for encrypting the upstream traffic as well. At the extreme, you will need four private keys to fully replay an audit trail: two to open the normal traffic, and two more to display the upstream traffic.

Note that SPS itself cannot create the certificates used to encrypt the audit trails.

TIP: If two certificates are displayed in a row, they are a certificate-pair and you need the private key of both to replay the audit trails. If two certificates are displayed in separate rows, you need the one of the private keys to replay the audit trails. If there are multiple rows containing two certificates, you need the private keys of the certificate(s) listed in any of the rows.

Figure 168: Policies > Audit Policies — Encrypting audit trails: joint encryption with a certificate pair



Each audit policy can have up to 8 lines of certificate pairs.

To encrypt audit trails

1. Navigate to **Policies** > **Audit Policies** and select the audit policy you will use in your connections.

   TIP: By default, every connection uses the built-in **default** audit policy. Unless you use a custom audit policy, modifying the **default** audit policy will affect every audited channel of the connections passing through One Identity Safeguard for Privileged Sessions (SPS).

2. Select the **Enable encryption** option.
3. To upload a certificate for encrypting the entire trail:

   a. Click the + icon under the **Encryption cert (X.509 RSA) 4-eyes cert (X.509 RSA)** row.

b. Click on the left  icon and upload a certificate to SPS. This certificate will be used to encrypt the audit trails, and it must not include the private key.

   NOTE: To replay the audit trails, you need the private key of the certificate on the computer running the Safeguard Desktop Player application.

c. (Optional) To encrypt the audit trails jointly with another certificate, click on the right  icon and upload a certificate to SPS. Note that the private key of both certificates will be required to replay the audit trails.

d. Repeat these steps to encrypt the audit trails separately with additional certificates.

4. To upload a certificate for encrypting the upstream traffic:

   a. Select **Encrypt upstream traffic with different certificates**.

   b. Click the + icon under the **Encryption cert (X.509 RSA) 4-eyes cert (X.509 RSA)** row.

   c. Click on the left  icon and upload a certificate to SPS. This certificate will be used to encrypt the audit trails, and it must not include the private key.

      NOTE: To replay the upstream part of the audit trails, you need the private key of the certificate on the computer running the Safeguard Desktop Player application.

   d. (Optional) To encrypt the audit trails jointly with another certificate, click on the right  icon and upload a certificate to SPS. Note that the private key of both certificates will be required to replay the audit trails.

   e. Repeat these steps to encrypt the upstream separately with additional certificates.

5. Click **Commit**.

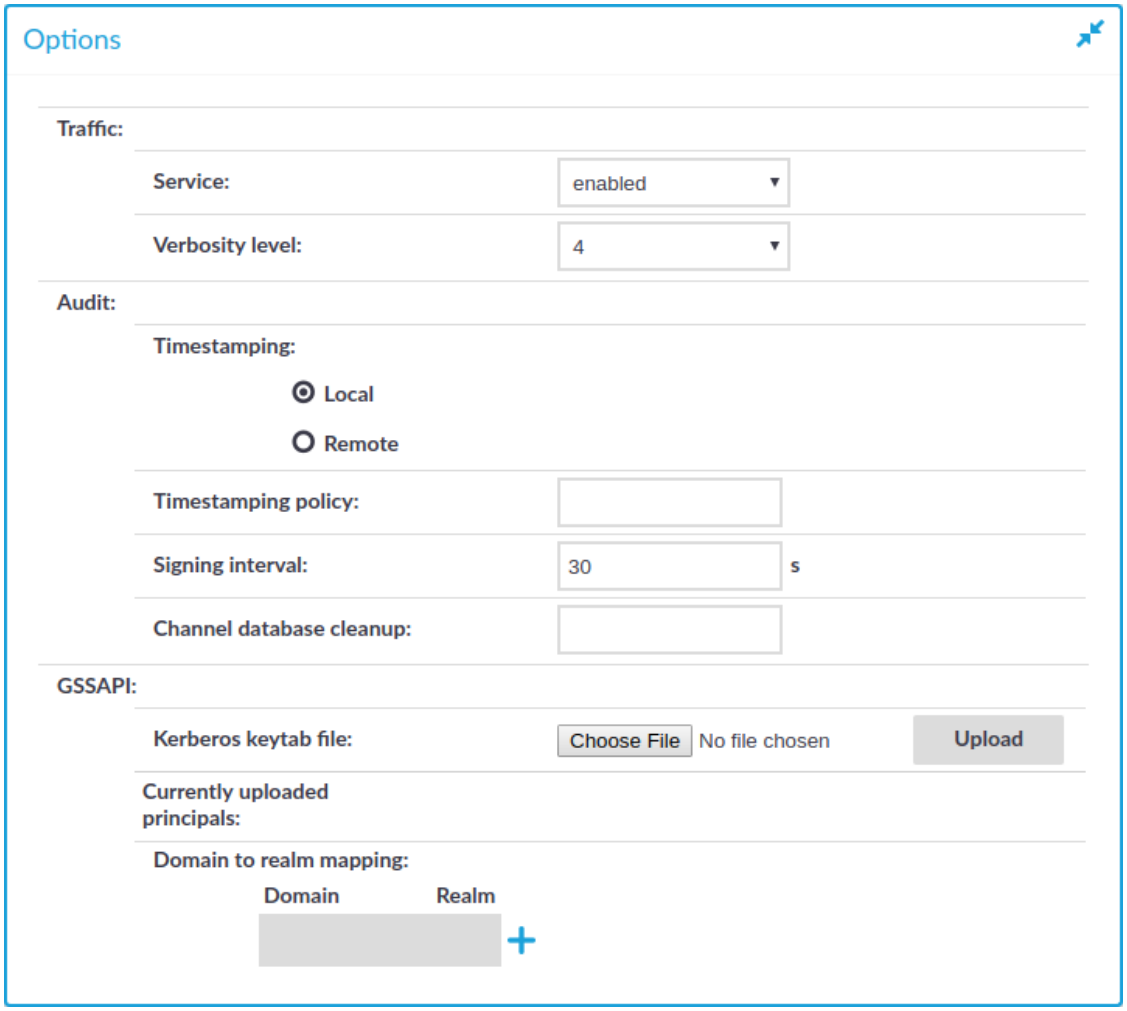
Timestamping audit trails with built-in timestamping service

The following describes how to add timestamps to the audit trails by using the built-in timestamping service of One Identity Safeguard for Privileged Sessions (SPS).

To add timestamps to the audit trails by using the built-in timestamping service of SPS

1. Configure the timestamping interval. You have to repeat these steps for each protocol (HTTP, ICA, RDP, SSH, Telnet, and VNC) you want to configure:

Figure 169: Traffic Controls > Protocol name > Global Options — Configuring local timestamping



a. In the protocol control settings, navigate to **Global Options** > **Timestamping** (for example, **Traffic Controls** > **SSH** > **Global Options** > **Timestamping**).

b. Select **Local**.

   NOTE: Make sure that you leave the **Timestamping policy** field empty. **Timestamping policy** has relevance only when **Timestamping** is set to **Remote**.

c. Set the **Signing interval**. You can choose any value between 10 and 100 000 seconds.

   NOTE:The same interval setting applies to timestamping and signing.

d. Click **Commit**.

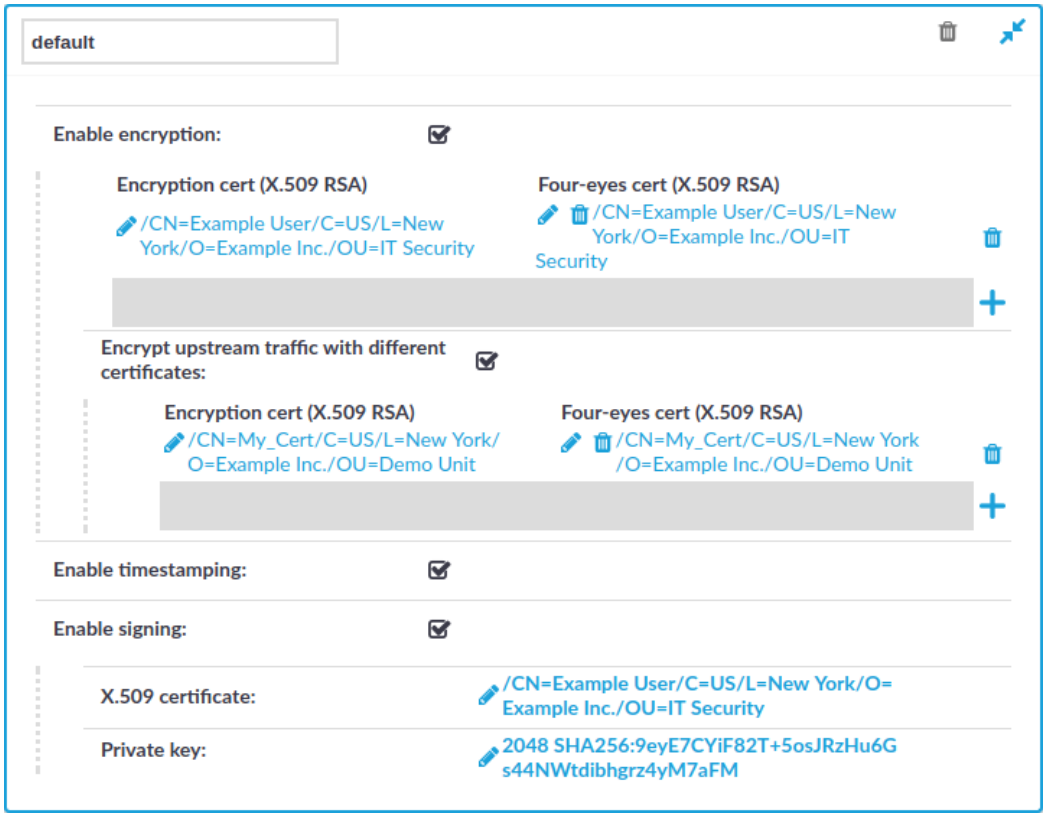
2. Configure audit policies to use timestamping. You have to repeat these steps for each audit policy you want to configure:

a. Navigate to **Policies** > **Audit Policies** and select the audit policy you will use in your connections.

   TIP: By default, every connection uses the built-in **default** audit policy. Unless you use a custom audit policy, modifying the **default** audit policy will affect every audited channel of the connections passing through One Identity Safeguard for Privileged Sessions (SPS).

b. Select the **Enable timestamping** option.

Figure 170: Policies > Audit Policies — Timestamping audit trails



c. Click **Commit**. SPS will automatically add timestamps to the audit trails of every connection that is audited and uses this audit policy.

   NOTE: For details on how to change the certificate used for timestamping, see Managing the certificates used on One Identity Safeguard for Privileged Sessions (SPS) on page 459.

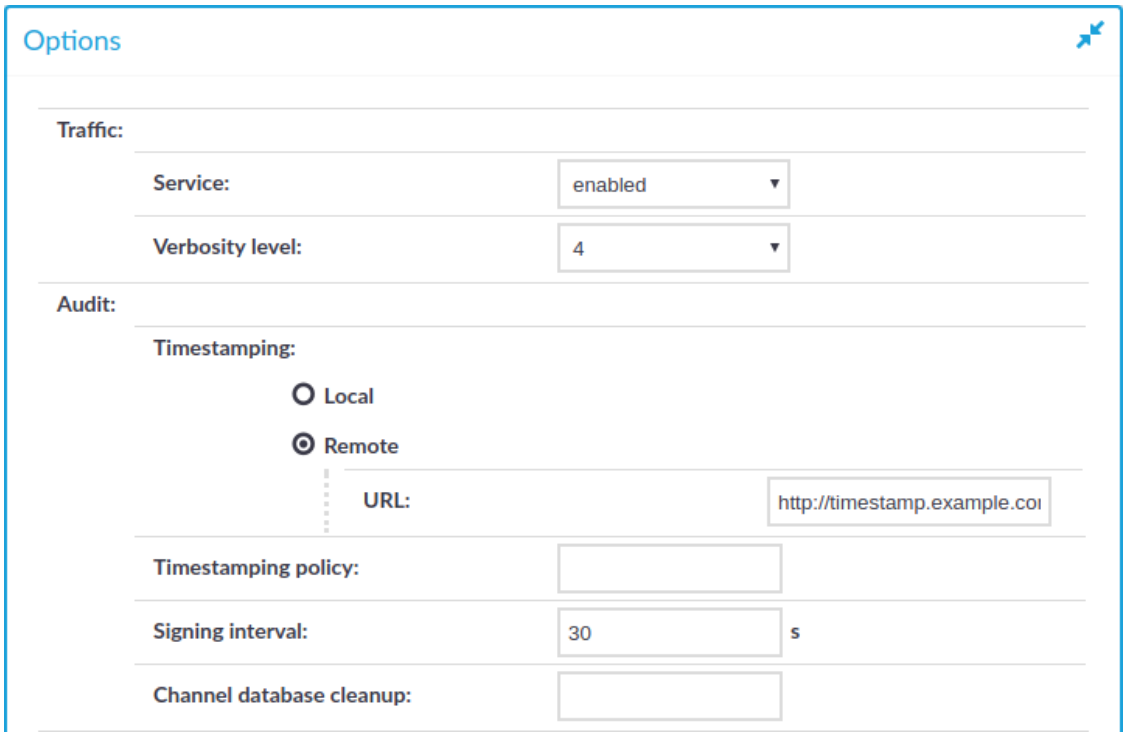
Timestamping audit trails with external timestamping service

The following describes how to request timestamps from a remote Timestamping Authority (TSA).

To request timestamps from a remote TSA

1. Configure the remote TSA, and the timestamping interval. You have to repeat these steps for each protocol (HTTP, ICA, RDP, SSH, Telnet, and VNC) you want to configure:

Figure 171: Traffic Controls > Protocol name > Global Options — Configuring a remote TSA



   a. In the protocol control settings, navigate to **Global Options** > **Timestamping** (for example, **Traffic Controls** > **SSH** > **Global Options** > **Timestamping**).
   b. Select **Remote**.
   c. Enter the address of the timestamping server into the **URL** field. Note that currently only plain HTTP services are supported, password-protected and HTTPS services are not supported.
   d. If the Timestamping Server has timestamping policies configured, enter the OID of the policy to use into the **Timestamping policy** field. One Identity Safeguard for Privileged Sessions (SPS) will include this ID in the timestamping requests sent to the TSA.
   e. Set the **Signing interval**. You can choose any value between 10 and 100 000 seconds.

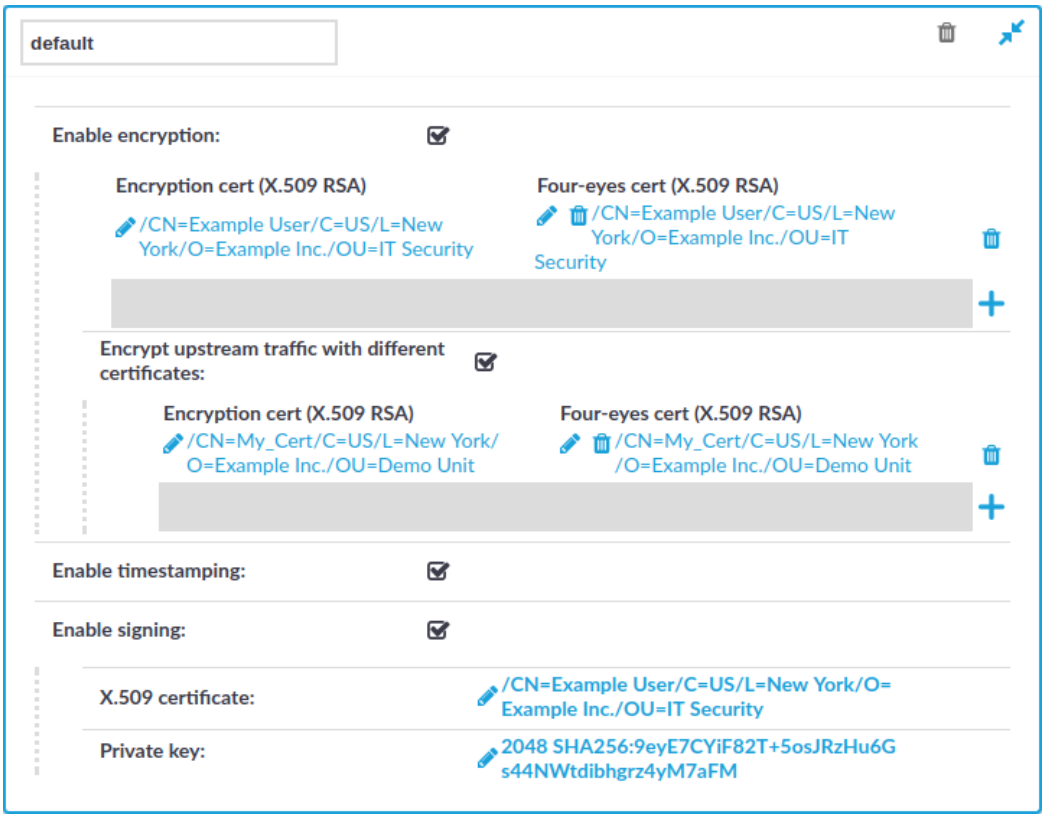
      NOTE: The same interval setting applies to timestamping and signing.
   f. Click **Commit**.

2. Configure audit policies to use timestamping. You have to repeat these steps for each audit policy you want to configure:
    a. Navigate to **Policies** > **Audit Policies** and select the audit policy you will use in your connections.

    TIP: By default, every connection uses the built-in **default** audit policy. Unless you use a custom audit policy, modifying the **default** audit policy will affect every audited channel of the connections passing through One Identity Safeguard for Privileged Sessions (SPS).

    b. Select the **Enable timestamping** option.

Figure 172: Policies > Audit Policies — Timestamping audit trails



    c. Click **Commit**. SPS will automatically add timestamps to the audit trails of every connection that is audited and uses this audit policy.

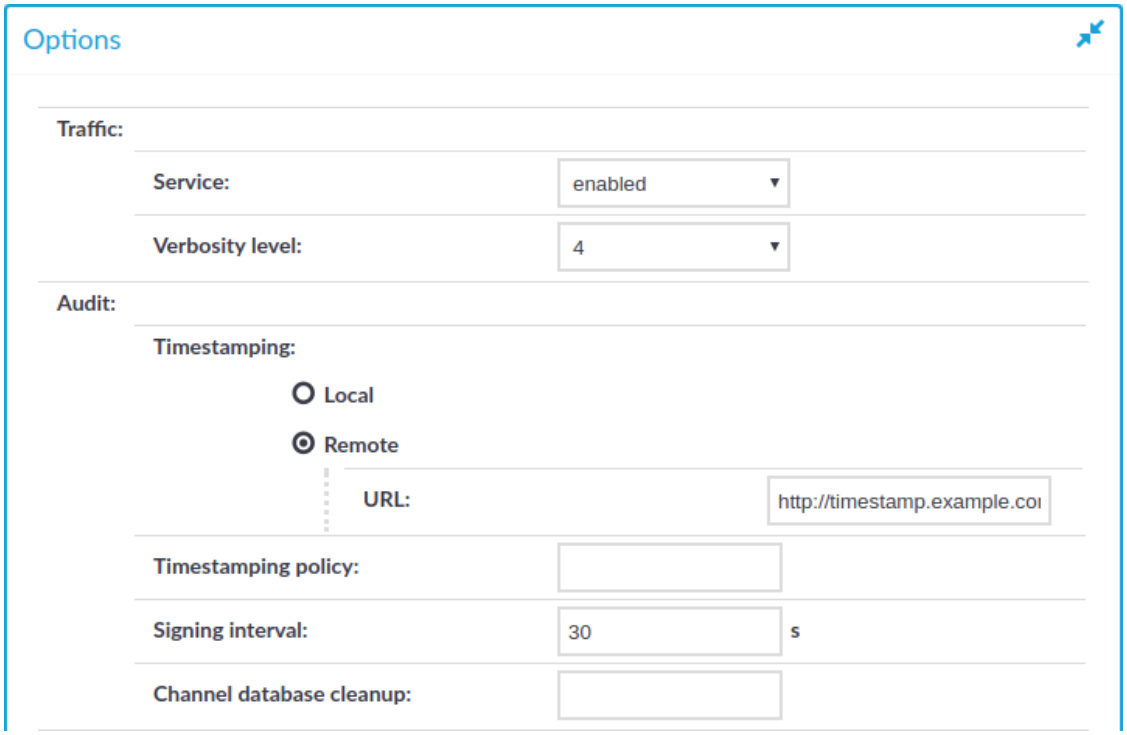
Digitally signing audit trails

One Identity Safeguard for Privileged Sessions (SPS) can digitally sign the audit trails to prevent the manipulation of the audit trail files. This requires an X.509 certificate and also the private key of the certificate. Note that SPS can generate a private key that you can use to create a certificate, but SPS itself cannot create the certificate used to sign the audit trails.

To enable the digital signing of the audit trails

1. Configure the signing interval. You have to repeat these steps for each protocol (HTTP, ICA, RDP, SSH, Telnet, and VNC) you want to configure:

   Figure 173: Traffic Controls > Protocol name > Global Options — Configuring the signing interval

   a. In the protocol control settings, navigate to **Global Options** > **Timestamping** (for example, **Traffic Controls** > **SSH** > **Global Options** > **Timestamping**).

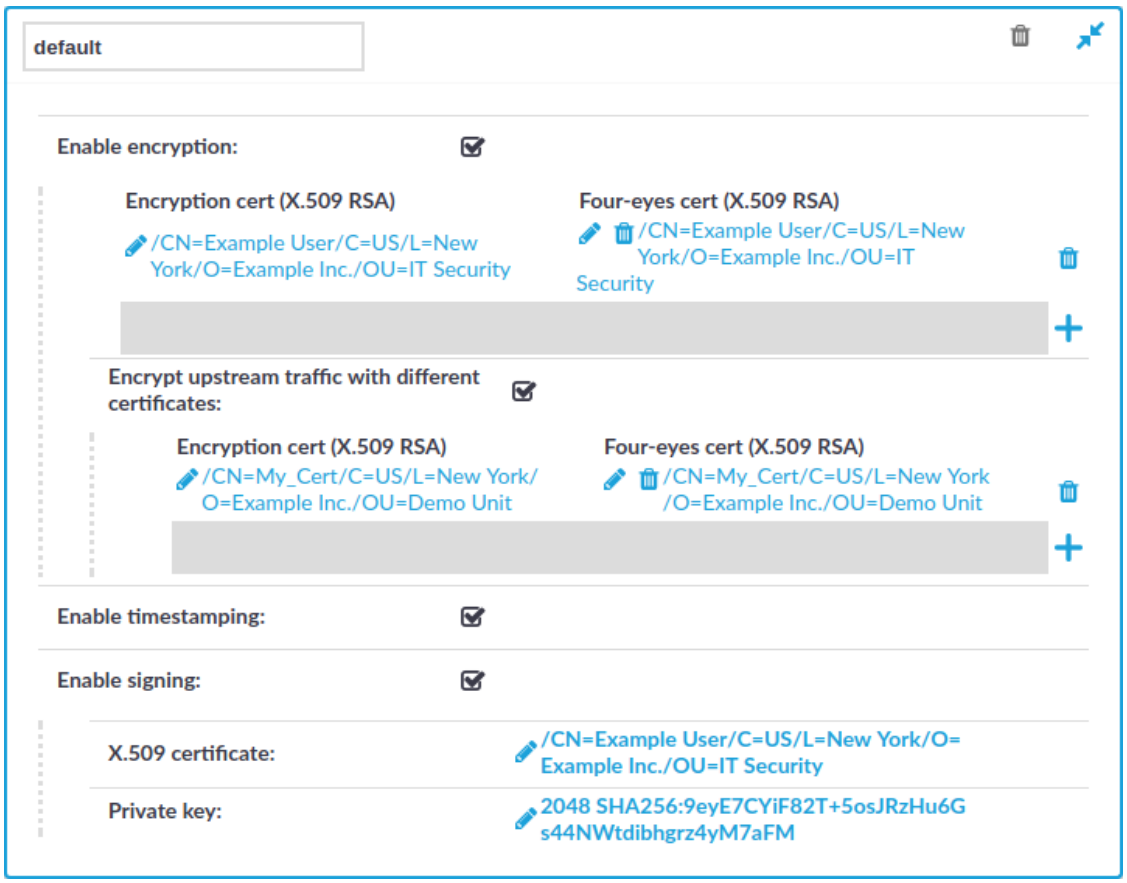
   b. Set the **Signing interval**. You can choose any value between 10 and 100 000 seconds.

      NOTE: The same interval setting applies to timestamping and signing.

   c. Click **Commit**.

2. Navigate to **Policies** > **Audit Policies** and select the audit policy you will use in your connections.

Figure 174: Policies > Audit Policies — Signing audit trails



TIP: By default, every connection uses the built-in **default** audit policy. Unless you use a custom audit policy, modifying the **default** audit policy will affect every audited channel of the connections passing through One Identity Safeguard for Privileged Sessions (SPS).

3. Select the **Enable signing** option.
4. Upload a certificate and the corresponding private key to SPS.
5. Click **Commit**. SPS will automatically sign the audit trails of every connection that is audited and uses this audit policy.
6. Repeat the above steps for other audit policies if needed.

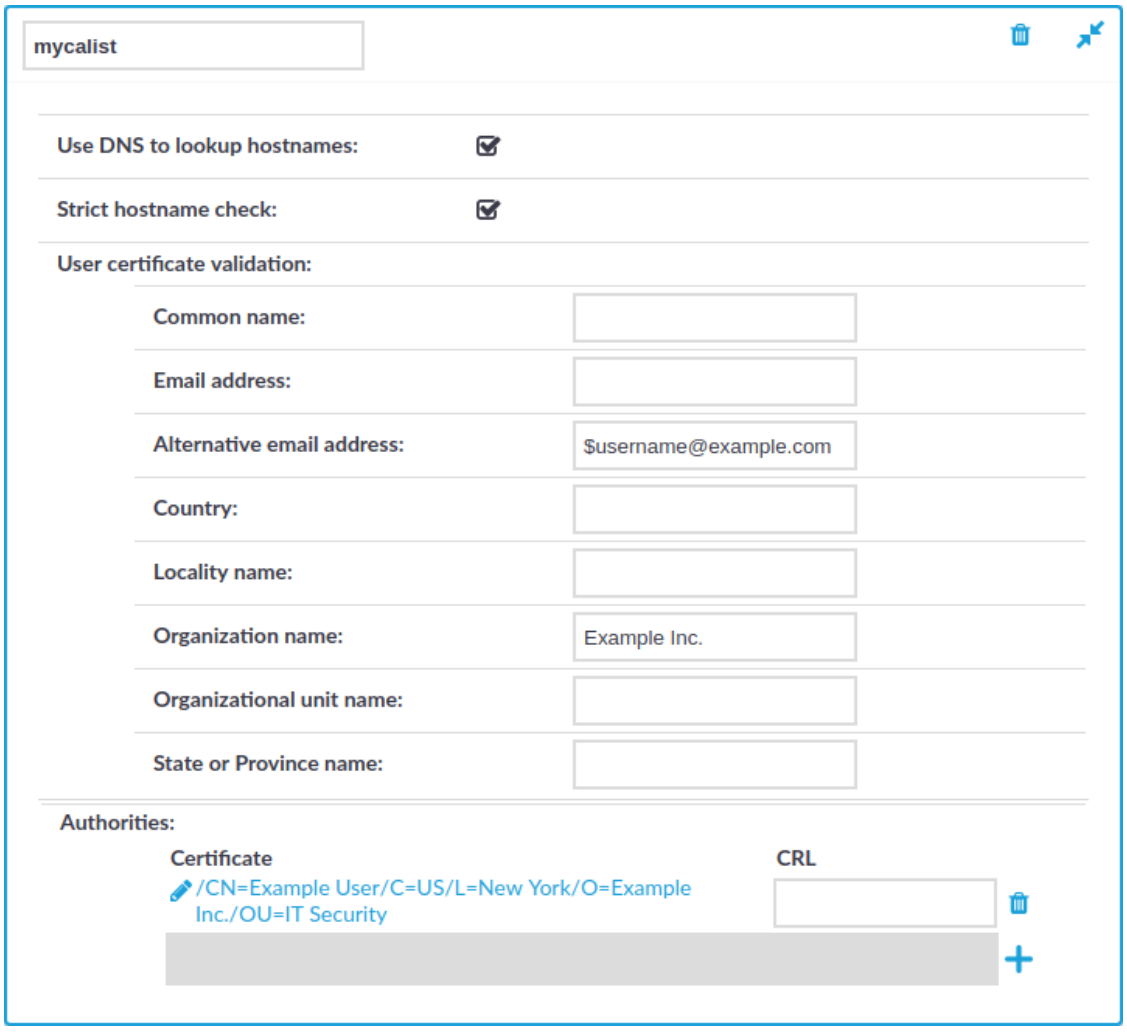
Verifying certificates with Certificate Authorities

One Identity Safeguard for Privileged Sessions (SPS) can check the validity of certificates using the certificates and certificate-revocation lists of the certificate authorities that issued the certificates. This can be used for example to verify the certificates of the servers in SSH/RDP connections.

To create a list of CA certificates to use during the certificate validation

1. Navigate to **Policies** > **Trusted CA Lists** and click + to create a new list.

Figure 175: Policies > Trusted CA Lists — Creating trusted CA lists



2. Enter a name for the CA list into the topmost field.
3. Click ✎ in the **Certificate** field, and upload the certificate of the Certificate Authority (CA) that will be used to validate the certificates.
4. Enter the URL of the Certificate Revocation List of the CA into the **CRL** field. Certificates appearing on the CRL list will be automatically rejected.

   NOTE: Note that only X.509 CRLs are accepted in either PEM and DER format. PKCS7 CRLs are not accepted.

CAUTION:

From SPS version 6.5.0, SPS verifies the signature and the expiration on the Certificate Revocation Lists (CRL) configured at **Policies** > **Trusted CA Lists**. The CRLs must be signed by the same Certificate Authority for which it was uploaded. If the validation of the remote CRL fails (or when the remote CRL is unavailable), an alert is generated. Despite the alert, SPS still updates the active CRL with the remote CRL unless the previously downloaded CRL local copy is still valid. Note that it is a security issue if it is not possible to validate the signature on the CRL. If the local copy of the CRL expires, connections that rely on the Trusted CA may fail.

In conjunction with this change, if the remote CRL is specified as a HTTPS URL, the web server certificate is no longer verified before the download.

CRL refresh is implemented as a background service running hourly. Therefore if the problem with the remote CRL persists, you may get alerts every hour.

5. To further limit which certificates are accepted, you may use the following options:
   - **Strict hostname check**: Select this option to accept only certificates when the Common Name of the certificate contains the hostname or the IP address of the host showing the certificate.
   - **Use DNS to lookup hostnames**: Select this option to use the domain name server set on **Basic Settings** > **Network** > **Naming** to resolve the hostnames and IP addresses for certificate validation. If you have enabled the **Strict hostname check** option, you probably want to enable this option as well.
   - To restrict the accepted certificates based on the content of the certificate, enter the required value into the appropriate field of the **User certificate validation** section. For example, to accept only certificates that contain `Example Inc.` in their Organization Name field, enter **`Example Inc.`** in to the **Organization Name** field. In the Common name, E-mail address, and Alternative e-mail address fields you can use the `$username` macro to refer to the username used in the connection. This macro makes it possible to check that the user is using his own certificate.
6. Click **Commit**.

Verifying certificates with Certificate Authorities using trust stores

One Identity Safeguard for Privileged Sessions (SPS) can check the validity of certificates using the certificates and certificate revocation lists of the certificate authorities (CA) that issued the certificates. This can be used to verify the certificates of the servers in TLS connections.

Trust stores serve as local certificate storages where you can store the certificate chains of trusted CAs. Create and configure custom trust stores to verify the certificates in TLS

connections. Note that you cannot modify the built-in trust store, which contains common CA certificates that the operational system of SPS uses.

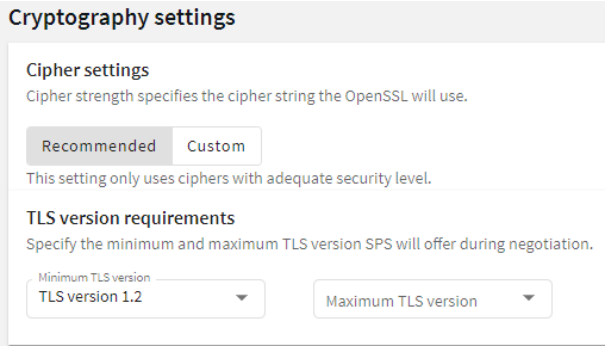
You can use custom trust stores in the following locations:

- **Basic Settings** > **Management** > **Syslog — Check server certificate**. For more information, see Configuring system logging.
- **Basic Settings** > **Management** > **Universal SIEM forwarder — TLS with certificate validation**. For more information, see Using the universal SIEM forwarder.
- **Users & Access Control** > **Settings** > **Authentication settings — Configuring X.509 authentication**. For more information, see Authenticating users with X.509 certificates.
- **Users & Access Control** > **Settings** > **Authentication settings — Configuring LDAP, TLS and STARTLS encryption method**. For more information, see Managing One Identity Safeguard for Privileged Sessions (SPS) users from an LDAP database.
- **Policies** > **LDAP Servers — Configuring encryption**. For more information, see Authenticating users to an LDAP server.

Prerequisites

To specify cipher settings and version requirements that SPS uses for establishing TLS connections in the **Cryptography Settings** section, navigate to **Basic Settings** > **Trust Stores**.

Figure 176: Basic Settings > Trust Stores — Cryptography settings



Cipher strength specifies the cipher string OpenSSL will use. The following settings options are possible:

- **Recommended**: this setting only uses ciphers with adequate security level.
- **Custom**: this setting allows you to specify the list of ciphers you want to permit SPS to use in the connection. This setting is only recommended to ensure compatibility with older systems. For more details on customizing this list, check the 'openssl-ciphers' manual page on your SPS appliance.

  For example: `ALL:!aNULL:@STRENGTH`

Minimum TLS version specifies the minimal TLS version SPS will offer during negotiation. The following settings options are possible:

- **TLS 1.2**: this setting only offers TLS version 1.2 during the negotiation. This is the recommended setting.
- **TLS 1.1**: this setting offers TLS version 1.1 and later versions during the negotiation.
- **TLS 1.0**: this setting offers TLS version 1.0 and later versions during the negotiation.

Maximum TLS version specifies the maximal TLS version SPS will offer during negotiation. The following settings options are possible:

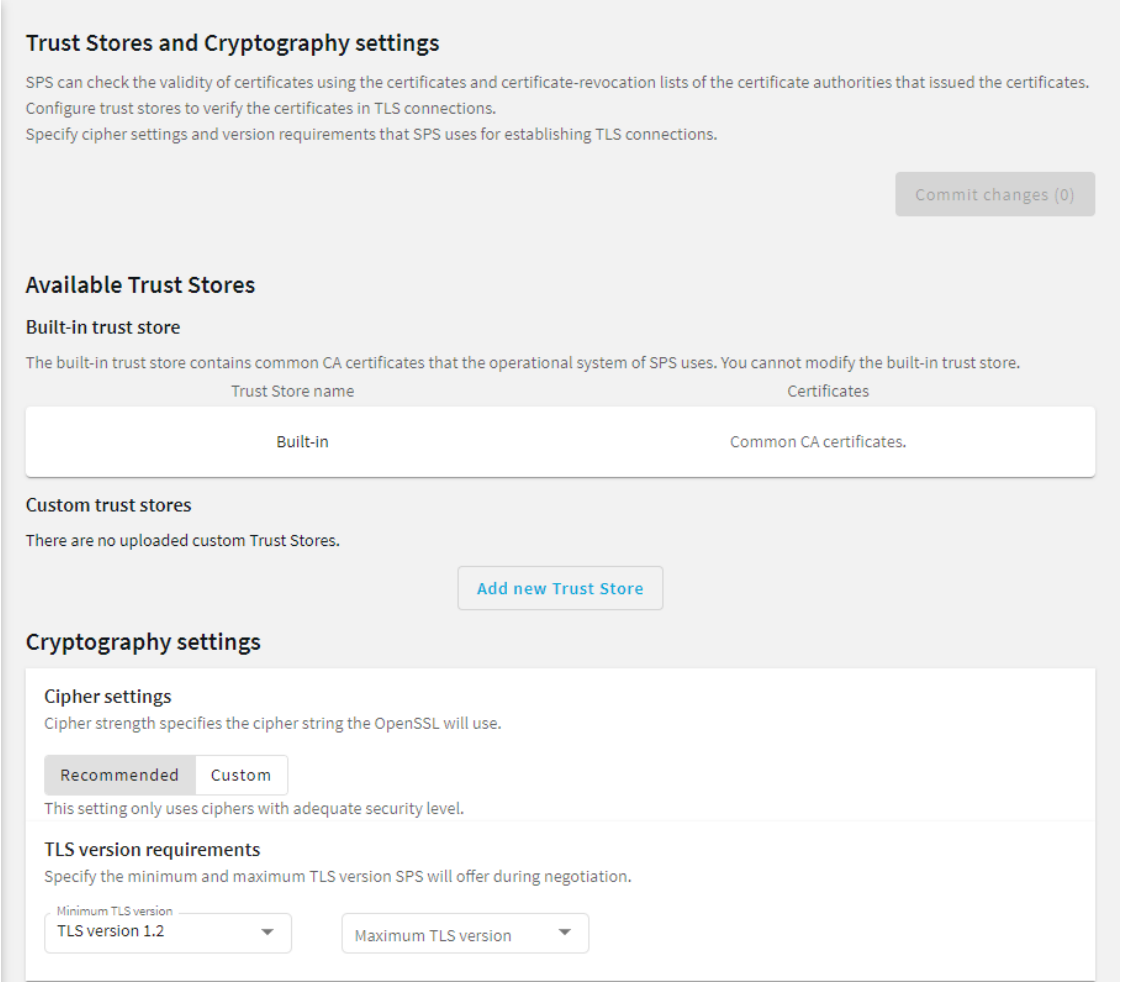
- **TLS 1.2**: this setting will only offer TLS version 1.2 during negotiation. This is the recommended setting.
- **TLS 1.1**: this setting will offer TLS version 1.1 and later versions during negotiation.
- **TLS 1.0**: this setting will offer TLS version 1.0 and later versions during negotiation.
- **Latest**: this setting will offer the latest TLS version during negotiation.

NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported.

To create a custom trust store or edit an existing custom trust store

1. Navigate to **Basic Settings** > **Trust Stores**.

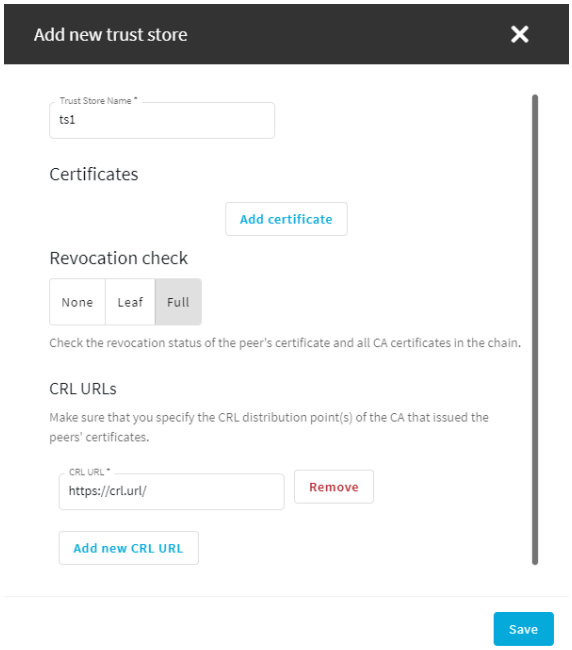
Figure 177: Basic Settings > Trust Stores



2. In the **Custom trust stores** section, click **Add new Trust Store**. To change the settings of an existing trust store, click **Edit**.

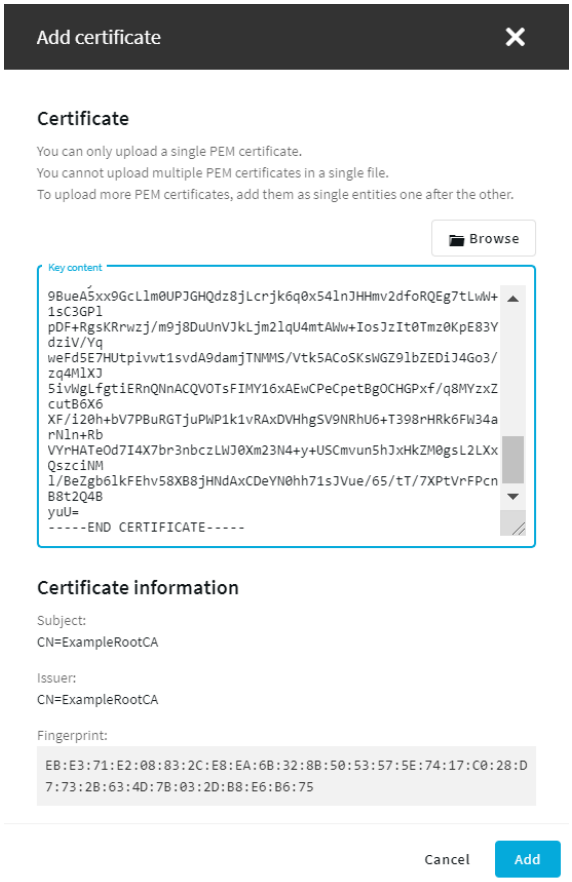
   The **Add new trust store** or **Edit trust store** window opens.

Figure 178: Basic Settings > Trust Stores — Add new trust store



3. In the **Trust Store Name** field, enter a name for the trust store. For example, **ts1**.
4. To add the certificate of the CA to the trust store that will be used to validate the certificates, click **Add certificate**. The **Add certificate** window opens.
5. Paste or drag the PEM certificate file to the **Key content** field. You can also click **Browse** to navigate to the PEM certificate file on your computer.

Figure 179: Basic Settings > Trust Stores — Add certificate



- You must upload all the CA certificates from the users' certificate chains (including the root CA). Therefore, these CA certificates must be on your computer in PEM format.
- The certificates of the users must contain the user name used to authenticate on SPS. You must know which certificate field will contain the user names (for example, `CN` or `UID`).
- The certificates must be imported into the browsers of the users. With SPS, you can authenticate with a certificate only if a personal certificate is available in the browser.

NOTE:

You can only upload a single PEM certificate.

You cannot upload multiple PEM certificates in a single file.

To upload more PEM certificates, add them as single entities one by one.

6. Click **Add**.
7. Set the revocation check type for the trust store.

- **None**: Do not check certificate revocation status.
- **Leaf**: Check the revocation status of the peer's certificate, but do not check the revocation status of the CA certificates in the chain.
- **Full**: Check the revocation status of the peer's certificate and all CA certificates in the chain.

8. If you set the revocation check type to **Leaf** or **Full**, click **Add new CRL URL** to specify the CRL distribution point(s) of the CA that issued the peers' certificates.
    a. Enter the URL of the Certificate Revocation List (CRL) of the CA into the **CRL URL** field. Certificates appearing on the CRL list will be automatically rejected.
    b. To remove a CRL distribution point from the trust store, click **Remove**.
9. Click **Save** and close the window.
10. To delete a certificate from a trust store, click **Edit** to open the **Edit trust store** window, click on the certificate, and click **Delete**.
11. To delete a trust store, click **Edit** to open the **Edit trust store** window, and click **Delete**.
12. Click **Commit**.

Signing certificates on-the-fly

At a number of places, One Identity Safeguard for Privileged Sessions (SPS) can generate the server certificates on the fly. This technique is used for example in SSL-encrypted RDP sessions, RDP sessions that use Network Level Authentication (CredSSP), or SSH connections that use X.509-based authentication.

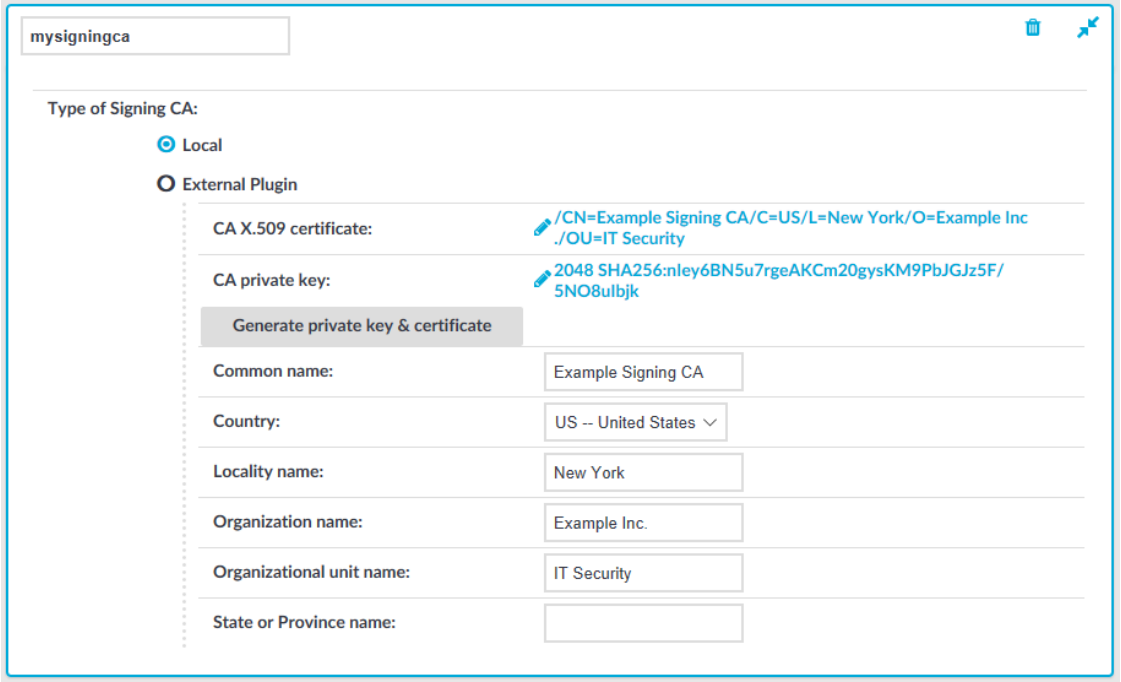
NOTE: Note the following points about using signing CAs:

- Signing CAs require a CA certificate permitted to sign certificates, and also the corresponding private key.
- These CAs cannot be used to sign audit trails. For details on how to configure the certificates used to sign audit trails, see Digitally signing audit trails on page 519.
- The version of the generated certificates will be the same as the version of the signing CA.
- SPS ignores the CRL (from the crlDistributionPoints extension) of the signing CA when generating certificates. If you want to include a CRL in the generated certificates, you must set it manually. See the following steps for details.

To create a signing CA using the built-in signing CA solution

1. Navigate to **Policies** > **Signing CAs** and click  .
2. Select **Local**.
3. Enter a name for the CA into the topmost field.

Figure 180: Policies > Signing CAs — Creating Signing CAs - Local



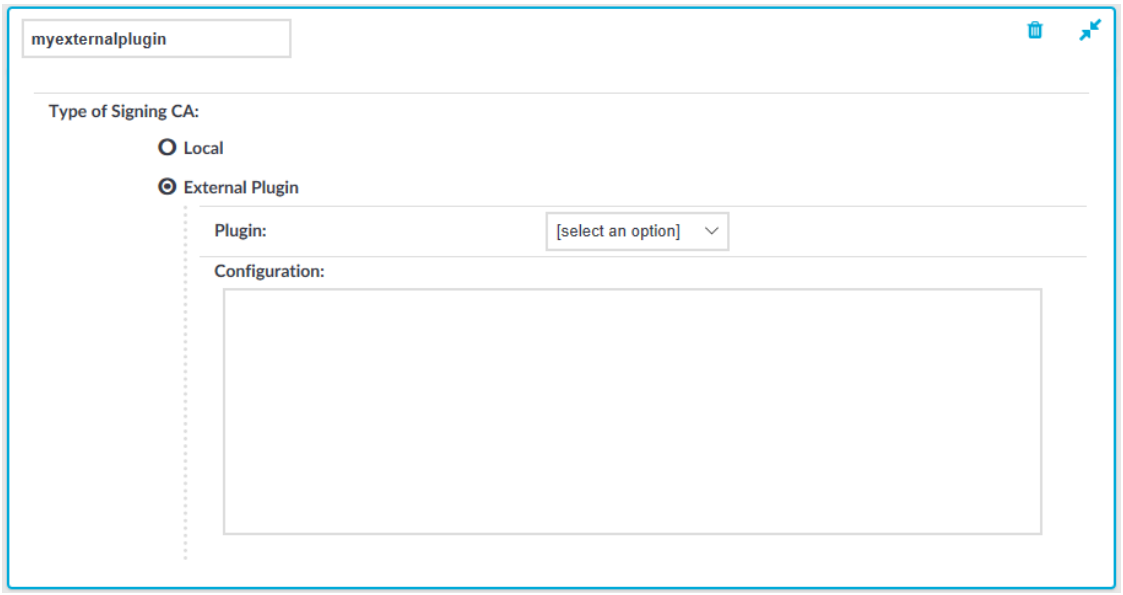
4. To upload a CA certificate and its private key, complete the following steps. Skip this step if you want to generate a CA on SPS.
    a. Click **Edit** in the **CA X.509 certificate** field and upload the certificate of the certificate authority. Alternatively, you can upload a certificate chain, where one member of the chain is the CA that will sign the certificates.
    b. Click **Edit** in the **CA private key** field and upload the private key of the certificate authority that will sign the certificates.
    c. (Optional) Enter the URL of the Certificate Revocation List (CRL) that you generated using your Certificate Authority in your Public Key Infrastructure (PKI) solution. The URL pointing to this CRL will be included in the certificate. This is the CRL information that will be shown to clients connecting to SPS.

        Note that the CRL list is not generated by the internal CA of SPS. The list must come from your own PKI solution.
    d. Click **Commit**.
5. To generate a CA certificate on SPS, complete the following steps:
    a. Enter the Common Name for the CA certificate into the **Common Name** field. This name will be visible in the `Issued By` field of the certificates signed by this CA.
    b. Fill the other fields as required, then click **Generate private key and certificate**.
    c. Click **Commit**.

To create a signing CA using an external signing CA plugin

1. Navigate to **Policies** > **Signing CAs** and click + .
2. Select **External Plugin**.
3. Enter a name for the CA into the topmost field.

Figure 181: Policies > Signing CAs — Creating Signing CAs - External Plugin



4. From the **Plugin** field, select an uploaded external plugin using the drop-down menu.

   To be able to select from the drop-down menu, you must have an external plugin uploaded in **Basic Settings** > **Plugins**.

   For more information about how to create an external Signing CA plugin, see Creating an external Signing CA.
5. Optionally, fill the **Configuration** field as required by the uploaded plugin.

   The input you enter in the **Configuration** field is passed down to the plugin.

Creating an external Signing CA

Overview

The External Signing CA plugin's purpose is to generate certificate and private key pairs signed by a Certificate Authority. By using this type of plugin the certificate signing can be tailored to fit any custom requirement.

Details

The plugin is a `.zip` file containing a `MANIFEST` and a `main.py` file.

The MANIFEST file

The MANIFEST file is a YAML file, and should conform to version 1.2 of the YAML specification. It should contain the following information about the plugin:

```
api: 1.0
type: signingca
name: MySigningCaPlugin
version: 1.0
description: My custom Signing CA
```

The type of the plugin must be `signingca`.

The main.py file

A `Plugin` class containing the following methods must be defined in the `main.py` file:

- `generate_for_addresses`: generates a key/certificate pair for the given addresses (IP/DNS)
- `generate_for_username`: generates a key/certificate pair for the given username
- `generate_for_subject`: generates a key/certificate pair for the given subject values

Method arguments

Each method must take the following arguments:

- `generate_for_addresses`:
  - `addresses`: `{list of str}` - contains either IP or DNS addresses for which the certificate shall be issued
  - `keytype`: `{str}` - contains a fixed value of 'RSA' or 'DSS' indicating the requested key type for the certificate
- `generate_for_username`:
  - `username` `{str}` - contains the username for which the certificate shall be issued
  - `keytype`: `{str}` - contains a fixed value of 'RSA' or 'DSS' indicating the requested key type for the certificate
- `generate_for_subject`:
  - `subject` `{list of (str, str)}` - contains the certificate subject as type-value pairs (tuples). Valid types are the following:
    - `'C'` country name
    - `'ST'` state or province name

- `'L'` locality name
- `'O'` orangization name
- `'OU'` organizational unit
- `'CN'` common name
- `'emailAddress'` email address

- `keytype: {str}` - contains a fixed value of 'RSA' or 'DSS' indicating the requested key type for the certificate

Method return values

Each method returns a `{dict}` that must have the following keys:

- `key: {str}` the generated key
- `chain: {list of str}` a PEM encoded certificate chain containing the generated certificate as the first element

Example

The code below demonstrates a simple plugin that can sign certificates with a built in CA. By default it uses a pre-generated CA certificate and key to complete signing requests. To use a custom certificate, provide a certificate and a key in a python dict format in the configuration field.

NOTE: If you wish to try this sample code, you will need to provide a MANIFEST file (see below) and the following package dependencies in the .zip file alongside the plugin:

- `asn1crypto`
- `certbuilder`
- `oscrypto`

`main.py`:

```
#!/usr/bin/env pluginwrapper3

#
# Copyright (c) One Identity
# All Rights Reserved.
#

from ast import literal_eval
from certbuilder import CertificateBuilder, pem_armor_certificate
from ipaddress import ip_address
```

```
from oscrypto import asymmetric

class Plugin(object):

    plugin_root_ca = """-----BEGIN CERTIFICATE-----
MIIDhjCCAm6gAwIBAgIIW+mOlk1Cu4swDQYJKoZIhvcNAQELBQAwYTELMAkGA1UE
BhMCSFUxETAPBgNVBAgMCEJ1ZGFwZXN0MREwDwYDVQQHDAhCdWRhcGVzdDEQMA4G
A1UECgwHQmFsYWJpdDEaMBgGA1UEAwwRQmFsYWJpdCBQbHVnaW4gQ0EwHhcNMTgx
MTEyMTQzMDQ2WhcNMTkxMTEyMTQzMDQ2WjBhMQswCQYDVQQGEwJIVTERMA8GA1UE
CAwIQnVkYXBlc3QxETAPBgNVBAcMCEJ1ZGFwZXN0MRAwDgYDVQQKDAdCYWxhYml0
MRowGAYDVQQDDBFCYWxhYml0IFBsdWdpbiBDQTCCASIwDQYJKoZIhvcNAQEBBQAD
ggEPADCCAQoCggEBAO73W0ONVwIaBJas+qUe0VBZ4rtk6PtzRNenZcBkTCkITuuF
DAQ3T1qLUsCyQ4uHMo+yKZUqR3HxbWGxS2l4IaHP6Hbna2kNEyYEsg16mGVUz6tc
D6bxFu3EpB7eU/OXh8RS8URIfZbLNrql1sKe7k1hpXUDS74Ra/avUIYKIpZ5sCjs
F6MBZWz5u3tNUa53xVmqgpnQ6pozN+OQ6k74DjK4xqWqJgTWcN6rxZ9k2voQYE3s
H66jl53q+Zl0D4w/AEW5W3OYNHJtx3tsc36sD2i0doqBCAAvflcSDEs7TXhfXSkC
qCBKyx8ics5EL9h49MDPGwDTehzwvXusz8LlxeMCAwEAAaNCMEAwDwYDVR0TAQH/
BAUwAwEB/zAdBgNVHQ4EFgQUyFWUMJli0q5CtJOp25IqK2M70oAwDgYDVR0PAQH/
BAQDAgEGMA0GCSqGSIb3DQEBCwUAA4IBAQCWoQCJPqfM4Sjg0R2O42yrE2GtQsXf
Qb3Dur+CefWLcvjI28t1xuj31khDgpNTwk4IVYrvarNX33C3tjYKgcimwWRMijbA
p8kZzFajOZSWC32CQtkWL79LLkJCTJB7b/4E41oNQPHtOoNCqFY+uQogP90qZ1w1
xlFX8ie/W3cuqhfzW6+/M3iCIwdjhBquvOo6mE3t2/lUcGXE20GayFsKnEmgpDJa
nxoG1+m+s5zCwDuukX8Lr1O7maTMwNVhm5P5QWeEPbGRN7yw+CfzcvPIbFYwnZ5x
XeC9Vtoj2Jbom8RV9uus8R5LfYBJ+HZh74wbGhIC2Kf9LrJTK7r92uVA
-----END CERTIFICATE-----"""

    plugin_root_ca_key = """-----BEGIN PRIVATE KEY-----
MIIEvwIBADANBgkqhkiG9w0BAQEFAASCBKkwggSlAgEAAoIBAQDu91tDjVcCGgSW
rPqlHtFQWeK7ZOj7c0TXp2XAZEwpCE7rhQwEN09ai1LAskOLhzKPsimVKkdx8W1h
sUtpeCGhz+h252tpDRMmBLINephlVM+rXA+m8RbtxKQe3lPzl4fEUvFESH2Wyza6
pdbCnu5NYaV1A0u+EWv2r1CGCiKWebAo7BejAWVs+bt7TVGud8VZqoKZ0OqaMzfj
kOpO+A4yuMalqiYE1nDeq8WfZNr6EGBN7B+uo5ed6vmZdA+MPwBFuVtzmDRybcd7
bHN+rA9otHaKgQgAL35XEgxLO014X10pAqggSssfInLORC/YePTAzxsA03oc8L17
rM/C5cXjAgMBAAECggEBAIucxpw76naW3tFtNG7eB2pLaZUUSq4F1VWtPlxd/MUI
Tpt5OuEHs3vx5CIixCWzkk2zyGmWrvEaHU6zN5ziC7wu7ODzKaTRd7uBiMkpM/oX
x9CU06w0NLIrbbt/J0ss36xKzRyYwY8lIM+Bbmx8UDuzbehkSY89PHd+S6xUJYsF
YmOVM00wx1N6yZGKHUV9GLRnysHb+DBbjGIcjDqmlsdyuAzlB7/DAeToLFNkZvzx
Qzza6whMILXS9Qp39dzn7nJuJywfoOAX2q5LgOrPise2QY1FuAy0GTfqvBDR1eGd
NwFW5YtH89l347AIDPgklKvKaii2iIw1ZEMf1AlX7mECgYEA+0JM9sToMgLhHJDj
cUsznVn3xzjDT/4O95LdAq1fVrn10wh/SwGCBBPMHQkV1nS70d//1aGctZLWk5+F
K3aPGV9Eas70mOGNcXdU1ITpFNuVfbKK8uH5NF/oEuoD4zRabunrj/zEk0Qu/D9v
pN4qEwJoV1SD/9HpfpaUG/xuBLECgYEA83mtE34zYTY2TLBr4HvoiWrFYoEpPldN
64oD+w1/D0Wd9hxCyzO3y2SmrBmmbzoawTckxD/VKndeRdV5dLlEnAV4F35bPsQl
dRJJEAAQPqqc1z4x6c2my27WPSbm4mIcvfTc65UFu4ovm/koywc96fwvpTX6JIN1
X8zHZ/tQaNMCgYEAl3yk3I9hk22K/ecZSiBWEUPCETpW/66kpX3FhKy085wQ61iP
```

```
LtDM69pn0QW+RduBtgsAu3PCAPN0LfManlbP9jMrE96NOHOdDNEusycjRHET04JH
JiM6VeqRCH5RM7ZH4+FjJh/3APc2AN3aWSOdaHKmKCkLoLyVs73jtG/ggTECgYBp
reCf22E1yrAa7WCFmYK/UqbGMMXUF1Ts7YT4zUzfNhpwHqgnRxV5pQBrJt8E3DWM
tACzZfmCazlyGkyTi27qQb10hRXZ0o1nmT45Qa3LZYaaLpa/otHI7xzyghYpIOjU
0pmpb4+DbWFo0+cO6N/I1ftgPGOMwbqKkHnk+kJWnQKBgQDQwITXna8OVjn86AQP
m36JHXi0RNVO/x4+b8T7nurU6XCPIzE0PxfVVSXXsbKTWlq48GIw9ZNpPKPSTCQy
fnYC+Pcu+4A+bfUwFk21khnN/fP5vyFlFhTrGneZeKWhUxv5iOqASEaizfOePmtj
et/4B9LUf9KQFstlhuIR4AP2OA==
-----END PRIVATE KEY-----"""

    def __init__(self, configuration=""):
        self.cert_x509name = {
            'country_name': 'HU',
            'state_or_province_name': 'Budapest',
            'locality_name': 'Budapest',
            'organization_name': 'One Identity',
            'common_name': '',
        }
        self.cert_alt_domains = []
        self.cert_alt_ips = []

        if configuration is not "":
            try:
                configdict = literal_eval(configuration)
                self.plugin_root_ca = configdict['ca']
                self.plugin_root_ca_key = configdict['key']
            except Exception:
                pass

        self.root_ca_certificate = asymmetric.load_certificate(bytes
(self.plugin_root_ca, 'utf-8'))
        self.root_ca_key = asymmetric.load_private_key(bytes(self.plugin_
root_ca_key, 'utf-8'))

    def _sign_certificate(self):
        end_entity_public_key, end_entity_private_key =
asymmetric.generate_pair(self.cert_keytype, bit_size=2048)
        end_entity_private_key = asymmetric.dump_private_key(end_entity_
private_key, None)
        end_entity_subject = self.cert_x509name

        builder = CertificateBuilder(
            end_entity_subject,
            end_entity_public_key
        )
```

```
        builder.subject_alt_domains = self.cert_alt_domains
        builder.subject_alt_ips = self.cert_alt_ips
        builder.issuer = self.root_ca_certificate
        end_entity_certificate = builder.build(self.root_ca_key)
        end_entity_certificate = pem_armor_certificate(end_entity_
certificate)
        return end_entity_certificate, end_entity_private_key

    _subject_type_mapping = {
        'C': 'country_name',
        'ST': 'state_or_province_name',
        'L': 'locality_name',
        'O': 'organization_name',
        'OU': 'organizational_unit_name',
        'CN': 'common_name',
        'emailAddress': 'email_address',
    }

    def _set_certificate_x509_name(self, x509name):
        for (t, v) in x509name:
            self.cert_x509name[self._subject_type_mapping[t]] = v

    def _set_certificate_keytype(self, keytype):
        # Certbuilder lists DSS key as DSA so we have to translate the
string here
        if keytype == "dss": keytype = "dsa"
        if keytype not in ['dsa', 'rsa']:
            raise ValueError('Certificate type should be either \'rsa\' or
\'dss\'')
        else:
            self.cert_keytype = keytype

    def _set_certificate_cn(self, cn):
        self.cert_x509name['common_name'] = cn

    def _set_certificate_addresses(self, addresses):
        for address in addresses:
            try:
                ip_address(address)
                self.cert_alt_ips.append(address)
            except ValueError:
                self.cert_alt_domains.append(address)

    def _build_response(self, cert, key):
```

```
        return {'key': key.decode('ascii'), 'chain': [cert.decode
('ascii'), self.plugin_root_ca]}

    def generate_for_addresses(self, addresses: list, keytype: str):
        self._set_certificate_keytype(keytype)
        self._set_certificate_addresses(addresses)
        self._set_certificate_cn(addresses[0])
        return self._build_response(
            *self._sign_certificate()
        )

    def generate_for_username(self, username: str, keytype: str):
        self._set_certificate_keytype(keytype)
        self._set_certificate_cn(username)
        return self._build_response(
            *self._sign_certificate()
        )

    def generate_for_subject(self, x509name: list, keytype: str):
        self._set_certificate_keytype(keytype)
        self._set_certificate_x509_name(x509name)
        return self._build_response(
            *self._sign_certificate()
        )
```

Use the following snippet as the `MANIFEST` file:

```
# Name of the plugin, may contain [a-zA-Z0-9]
name: HelloSigningCaPlugin

# Version of the plugin, only for display purposes
version: 0.1

# Type of the plugin - this is a signingca plugin
type: signingca

# API version of the SCB the plugin was written for, in major.minor format
api: 1.0

# Free form description.
description: This is an example plugin used for testing.

# Entry point for the plugin (also for running with python3)
entry_point: main.py
```

Creating a Local User Database

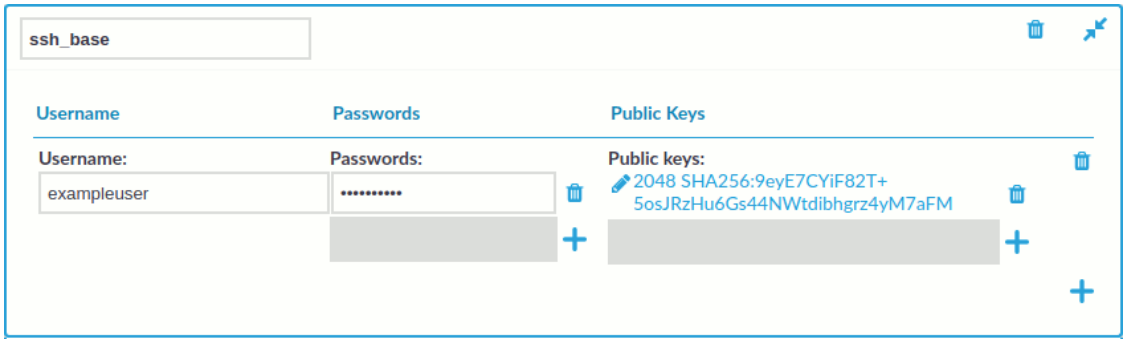
Local User Databases are available for HTTP, RDP, SSH and Telnet protocols, and can be used to authenticate the clients to credentials that are locally available on One Identity Safeguard for Privileged Sessions (SPS). Such credentials include passwords and public keys. **Local User Databases** are most commonly used in inband gateway authentication scenarios.

NOTE: To store credentials on SPS and use them to authenticate on the server, use a local Credential Store. For details, see Using credential stores for server-side authentication on page 873.

To create a Local User Database

1. Navigate to **Policies** > **Local User Databases** and click +.
2. Enter the name of the **Local User Database**.
3. Click + to add entries.

   Figure 182: Policies > Local User Databases — Mapping keys

4. Enter the name of the user into the **Username** field.

   NOTE: If you also use Usermapping policies, enter the username that the client will use on the server side. If you also use gateway authentication, the gateway username can be used as well.

5. If you use public-key based authentication on the client side, click the + icon in the **Public Keys** field, and upload the public key of the client.

   SPS will verify that a client trying to use the username set in Step 3 is authenticating itself with the private key that corresponds to the uploaded public key or certificate.

   TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

6. Repeat the above steps to add other users as required.
7. Click Commit.

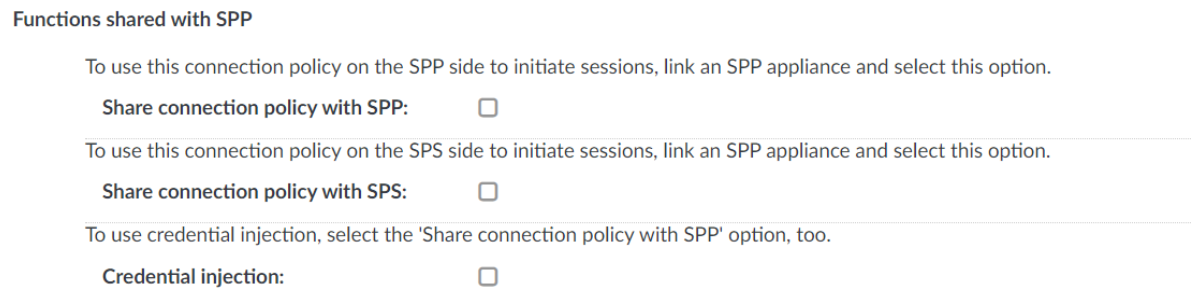
8. Navigate to the **Authentication Policies** tab of the respective protocol and select the **Local User Database** there.

Sharing SPS functions with SPP

If you have linked an SPP appliance to SPS and want to share specific functions with SPP, use the **Functions shared with SPP** option. This option lists the shared SPS functions that you can use in SPP without further configuration.

The following example displays the available shared functions under the **Traffic Controls** > **SSH** option.

Figure 183: Traffic Controls > RDP > Connections — Functions shared with SPP



Share connection policy with SPP

The **Share connection policy with SPP** option is supported currently with SSH and RDP connections. To make a specific SSH or RDP connection policy available in SPP to initiate sessions, select **Share connection policy with SPP** and set the configuration specified in the respective section:

- Sharing RDP connection policies with SPP
- Sharing SSH connection policies with SPP

Share connection policy with SPS

The **Share connection policy with SPS** option is supported currently with SSH and RDP connections. If you have joined an SPP to SPS, you can use the **Share connection policy with SPS** option.

To initiate sessions from SPS without using SPP directly, but using the credentials provided by and stored in SPP, select the Share connection policy with SPS option.

Set the configuration specified in the respective section:

- Sharing RDP connection policies with SPS
- Sharing SSH connection policies with SPS

Credential injection

The RDP Application session initiated on the SPP side provides the password automatically for the `RemoteApp Launcher`. To use credential injection, use a connection policy for the RDP Application session that has `Credential injection` selected.

NOTE: To use the credential injection function, you must select the **Share connection policy with SPP** option as well.

For more information, see Using credential injection in SPP-initiated RDP sessions.

HTTP-specific settings

HTTP request-response pairs do not form a well-defined, continuous connection. SPS will group request to the same session if the following points are true:

- The IP address of the client is the same.
- The hostname of the target server (not the IP address) is the same.
- The username is the same (if the user has performed inband authentication).
- The time elapsed since the last request-response pair between the same client and server is less than the session timeout value.

The following sections describe configuration settings available only for the HTTP protocol.

Use the following policies to control who, when, and how can access the HTTP connection. For details on configuring Channel Policies, see Creating and editing channel policies on page 495. For a list of supported client applications, see Supported protocols and client applications on page 31.

Auditing HTTP and HTTPS connections is possible in both transparent and non-transparent modes. SPS can also be used as an HTTP/HTTPS proxy to simplify client configuration and integration into your network environment, or it can forward HTTP traffic, behaving as a HTTP tunnel.

- *Channel Policy*: The channel policy determines which HTTP channels can be used in the connection, and whether they are audited or not. The different channels may be available only under certain restrictions, as set in the channel policy. For details on configuring these options, see Creating and editing channel policies on page 495.
- *HTTP connections*: For details, see Setting up HTTP connections on page 545.
- *HTTP sessions*: HTTP settings determine the parameters of the connection on the protocol level, including timeout value, and so on. For details, see Session-handling in HTTP on page 552.
- *HTTP settings*: HTTP settings determine the parameters of the connection on the protocol level, including timeout value, and so on. For details, see Creating and editing protocol-level HTTP settings on page 553.

For more information, see the latest One Identity Safeguard for Privileged Sessions Administration Guide.

Supported HTTP channel types

The available HTTP channel types and their functionalities are described below. For details on configuring Channel Policies, see Creating and editing channel policies on page 495. For a list of supported client applications, see Supported protocols and client applications on page 31.

- **HTTP**: Enables you to use the HTTP protocol. This channel must be enabled for HTTP to work.

  The available channel policy options are the following: **From**, **Target**, **Time policy**, **Record audit trail**, and **Remote groups**. Note that the **Remote groups** option is used only if the user performs inband authentication using one of the supported HTTP authentication methods (see Authentication in HTTP and HTTPS on page 543). To retrieve the groups of an authenticated user from an LDAP database, you must also set an **LDAP Server** in the Connection Policy (for HTTP/HTTPS connections, One Identity Safeguard for Privileged Sessions (SPS) uses this server only to retrieve the group membership of authenticated users, you cannot authenticate the users to LDAP from SPS). For details on configuring these options, see Creating and editing channel policies on page 495.

  When setting **Target**, note the following:

  - If the connection uses DNAT (NAT destination address), the target address of the original client will be compared to the **Target** parameter of the Channel policy, that is not necessarily equivalent with the server's address.
  - If the connection is redirected to a Fix address, the redirected address will be compared to the **Target** parameter of the Channel policy.
- **WebSocket**: Enables all WebSocket traffic. If the WebSocket channel type is not allowed, HTTP requests trying the WebSocket upgrade are rejected.

  WebSocket/VNC audit trails: You can replay audit trails of a WebSocket connection in your browser or using the Safeguard Desktop Player application only if it contains Virtual Network Computing (VNC) traffic. For all other WebSocket connections, export the audit trail as a PCAP file and replay it using the Safeguard Desktop Player application.

Limitations in handling HTTP connections

- When configuring HTTP or SSH connections, avoid using the IP address configured for administrator or user login on SPS.

  The current version of SPS does not support the following features that are available for other protocols:

  - Four-eyes authorization
- Forwarding HTTP connections to an HTTP proxy is not supported. If your clients use an HTTP proxy to access the target servers, place SPS behind the proxy: *Clients* -

HTTP Proxy - SPS.

> **CAUTION:**
>
> The *Clients - SPS - HTTP Proxy* scenario is not supported.

- HTTP (and thus WebSocket) sessions cannot be terminated. They neither have a **Terminate** button on the SPS side, nor would be a termination request successful initiated from Desktop Player.

  There is a related issue to remove the **Terminate** button from the SDP in case of HTTP / VNC over WebSocket sessions.

Authentication in HTTP and HTTPS

For the audited HTTP and HTTPS connections, One Identity Safeguard for Privileged Sessions (SPS) supports the following inband authentication methods for the HTTP protocol. These authentication methods are automatically supported for every Connection policy, without further configuration.

- Basic Access Authentication (according to RFC2617)
- The NTLM authentication method commonly used by Microsoft browsers, proxies, and servers

SPS records the username used in the authentication process into the **Username** and **Remote username** fields of the connection database.

For authenticated sessions, SPS can perform group-based user authorization that allows you to finetune access to your servers and services: you can set the required group membership in the Channel policy of the HTTP connection. Note that group-based authorization in HTTP works only for authenticated sessions (for HTTP/HTTPS connections, SPS uses this server only to retrieve the group membership of authenticated users, you cannot authenticate the users to LDAP from SPS). If a username is not available for the session, SPS will permit the connection even if the **Remote groups** field is set.

SPS does not store failed HTTP authentication attempts in the connection database. This means that the **Verdict** field of the Search page will never contain `CONN-AUTH-FAIL` values for HTTP connections.

Note that authentication also affects the way SPS handles HTTP sessions. For details, see Session-handling in HTTP on page 552.

Creating a new HTTP authentication policy

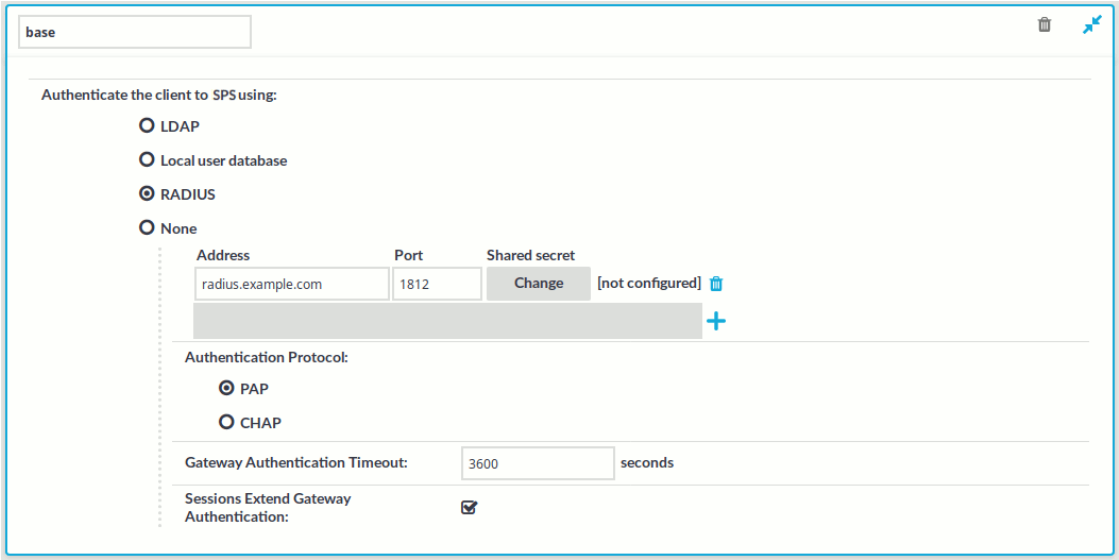
An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the

client can authenticate to One Identity Safeguard for Privileged Sessions (SPS).

To create a new authentication policy

1. Navigate to **Traffic Controls** > **HTTP** > **Authentication Policies**, and click + .

Figure 184: Traffic Controls > HTTP > Authentication Policies — Configuring HTTP authentication policies



2. Enter a name for the policy into the **Name** field.
3. Select the authentication method used on the client-side in the **Authenticate the client to SPS using** field. For the client-side connection, SPS can authenticate the client inband (within the HTTP protocol) using the following authentication methods:
   - **LDAP**: SPS will authenticate the client to the LDAP database set in the **LDAP Server** of the connection policy. To use LDAP authentication on the client side, select **Authenticate the client to SPS using** > **LDAP**.

     NOTE: SPS will authenticate the client-side connection to the LDAP server configured in the connection policy. This is not necessarily the same as the LDAP server used to authenticate the users accessing the SPS web interface.
   - **Local user database**: Authenticate the client locally on the SPS gateway using a **Local user database**. Select the database to use in the **Local user database** field. For details on creating a **Local User Database**, see Creating a Local User Database on page 538.
   - **RADIUS**: SPS will authenticate the client to the specified RADIUS server. Select **Authenticate the client to SPS using** > **RADIUS**, enter the IP address (use an IPv4 address) or hostname of the RADIUS server into the **Address** field, the port number of the RADIUS server into the **Port** field, and the shared secret of the RADIUS server into the **Shared secret** field. Only

password-authentication is supported (including one-time passwords), challenge-response based authentication is not.

To add more RADIUS servers, click  and fill in the respective fields.

- **None**: Do not perform client-side authentication, the client will authenticate only on the target server.

4. Specify the time remaining until a successful gateway authentication times out into the **Gateway Authentication Timeout** field.

   To avoid interruptions for active HTTP sessions, select the **Sessions Extend Gateway Authentication** checkbox. When enabled, active HTTP sessions can extend the gateway authentication beyond the configured timeout.

5. Click **Commit**.

   NOTE: The client-side authentication settings apply for authenticating the user inband (that is, within the HTTP protocol) to the SPS gateway.

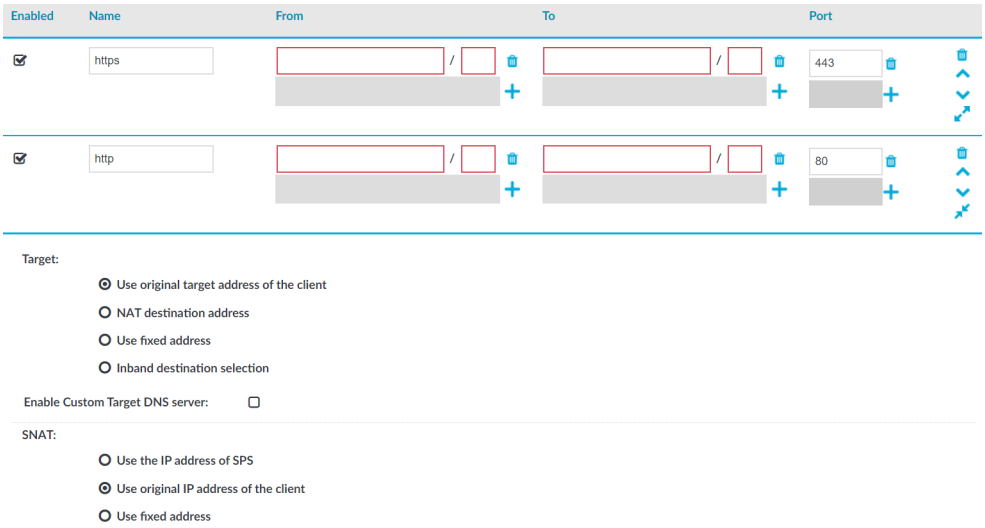
Setting up HTTP connections

This section focuses on describing the HTTP-specific details of connection configuration. For a detailed description on configuring connections, see General connection settings on page 482.

Setting up a transparent HTTP connection

This section describes how to set up a transparent HTTP connection. To audit HTTP connections in non-transparent mode, see Enabling One Identity Safeguard for Privileged Sessions (SPS) to act as an HTTP proxy on page 547.

Figure 185: Traffic Controls > HTTP > Connections — Transparent HTTP connection



To set up a transparent HTTP connection

1. In the **Name** field, enter the name of the connection that identifies the connection policy.
2. In the **From** field, enter the IP address and prefix of the client that you can use to access the server.

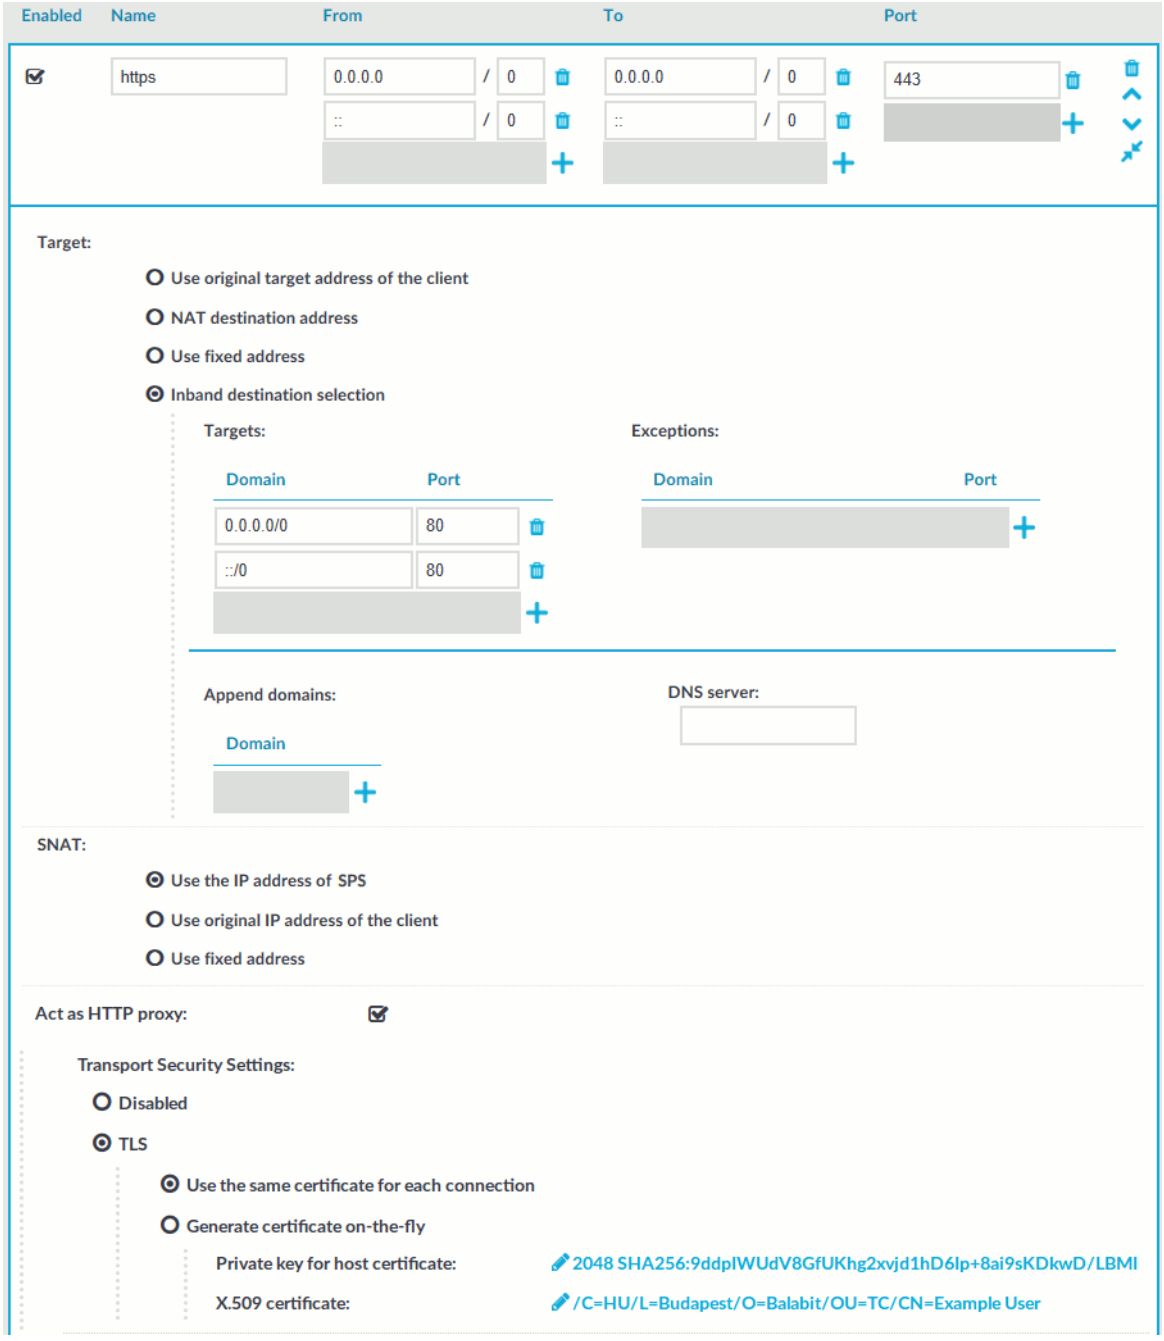
   You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).
3. In the **To** field, enter the IP address and prefix that the clients target.

   You can use an IPv4 or an IPv6 address. To limit the IP range to the specified address, set the prefix to **32** (IPv4) or **128** (IPv6).
4. In the **Target** section, select **Use original target address of the client**.
5. In the **SNAT** section, select **Use original IP address of the client**.
6. Since SPS cannot automatically decide whether the incoming sessions are encrypted or not, set up another identical connection policy for the same sessions, for HTTPS. As a result, HTTP and HTTPS sessions are saved into separate trails.
   a. Setup a new connection policy with the same settings as above.
   b. Set the **Port** to **443**.
   c. Enable TLS encryption. For more information, see Enabling TLS encryption in HTTP on page 549.

Enabling One Identity Safeguard for Privileged Sessions (SPS) to act as an HTTP proxy

This section describes how to enable One Identity Safeguard for Privileged Sessions (SPS) to act as an HTTP proxy.

Figure 186: Traffic Controls > HTTP > Connections — Act as HTTP proxy



To enable SPS to act as an HTTP proxy

1. Enable **Act as HTTP proxy** to configure the client to use SPS as an HTTP proxy.

   You can use SPS as an HTTP proxy through TLS. All traffic between the browser and SPS is tunneled through TLS.

   To use this feature, ensure that the client software can establish secure web proxy connections and supports client software configuration, such as proxy auto-configuration files.

   For more information about making browser specific settings for Chromium, see Secure Web Proxy.

2. Select whether you want encrypted web proxy connection between the HTTP client and SPS.

   Since there is now a secure channel between the web browser and SPS, you can also enable proxy authentication. This makes it possible for the web browser to do an inband gateway authentication to SPS before being able to issue HTTP requests through SPS.

   - To disable encryption between the HTTP client and SPS, select **Disabled**.

     NOTE: Since the forwarded data may contain sensitive information, One Identity recommends using encryption between the HTTP client and SPS.

   - To use encryption between the HTTP client and SPS, select one of the following options:

     - To use a fix certificate, select **Use the same certificate for each connection** and copy or upload the certificate.
     - To generate a certificate on-the-fly, signed by a provided Signing CA, select **Generate certificate on-the-fly**. It uses the parameters of the signing CA, excluding the CN field, which is filled with the name of the target host name.

       NOTE: When **Generate certificate on-the-fly** is selected and the connection is in transparent setup, the **CN** field is filled in using Server Name Indication (SNI). If the client does not support SNI, the **CN** field will contain the target IP, which may cause certificate verification warning on the client browser.

3. Select **Inband destination selection** as **Target**.
4. To permit access to any HTTP servers, enter `0.0.0.0/0` in the **Domain** field. Alternatively, enter the IP address or subnet of the HTTP address you want permit access to. For IPv6 addresses, add **::/0** as well.
5. To permit HTTP access to the destination servers on any port, leave the **Domain** > **Port** field empty. Otherwise, clients will be permitted only to access the specified port.
6. Enter the port where SPS should accept HTTP connections into the **To** > **Port** field. The default port number when using the **Act as HTTP proxy** setting is `3128`. This value should be the same as the proxy port setting on your clients.
7. Ensure that you have set SPS as proxy on the clients.

CAUTION:

To perform gateway authentication on SPS, the client browsers must be configured to use a Proxy Auto-Configuration (PAC) script.

To perform gateway authentication in a TLS-encrypted channel, the script must return an HTTPS address. Note that currently the Safari browsers do not support TLS-encryption in gateway authentication. For example:

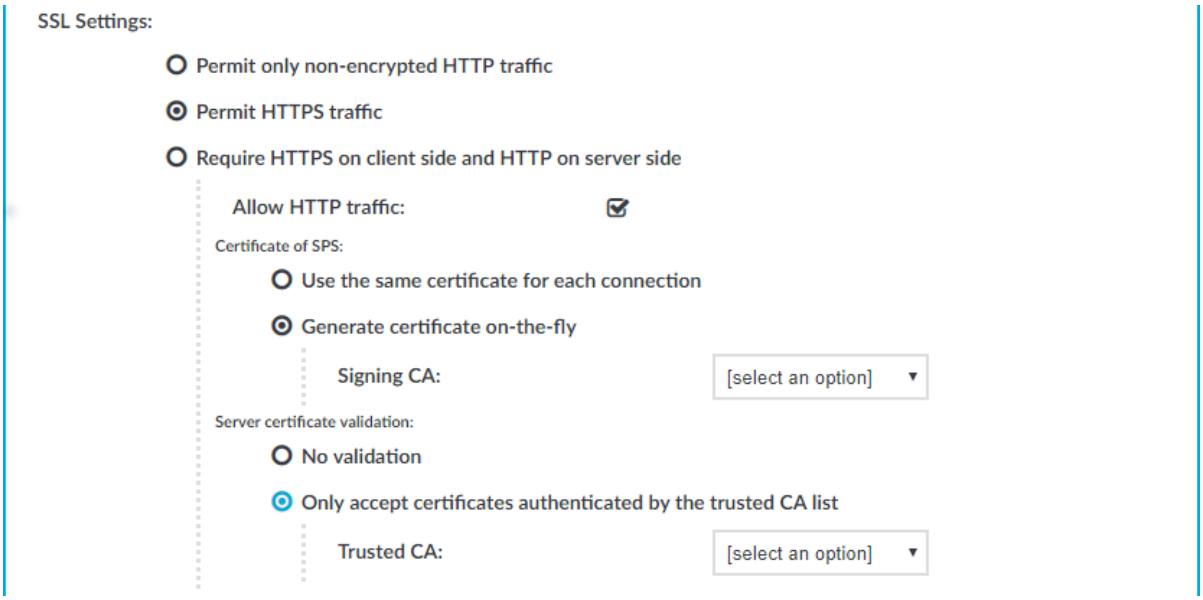
```
if (dnsDomainIs(host, "example-domain.com") || dnsDomainIs(host,
"www.example-domain.com"))
return "HTTPS 192.168.11.121:3128";
```

The client browsers might require the certificate of SPS to contain the **`Subject Alternate Name`** field. Certificates generated on SPS using the **Generate certificate on-the-fly** option automatically contain this field. If you **Use the same certificate for each connection**, make sure this field is present and properly set.

Enabling TLS encryption in HTTP

This setting either enforces TLS encryption or accepts both HTTP and HTTPS requests.

Figure 187: Traffic Controls > HTTP > Connections > SSL Settings — Enabling SSL encryption in HTTP



To enable SSL encryption

1. In **SSL Settings**, select **Permit HTTPS traffic**. To control plain HTTP traffic with the same connection policy, enable **Allow HTTP traffic**.
2. Select the certificate to show to the clients.
    - To use the same certificate for every peer, complete the following steps.
        1. Generate and sign a certificate for One Identity Safeguard for Privileged Sessions (SPS) in your PKI system, then export the certificate and its private key.
        2. Select **Use the same certificate for each connection**.
        3. Select **Private key for host certificate**, click  and upload the private key.
        4. Select **X.509 host certificate**, click  and upload the certificate.

    NOTE: When using the **Use the same certificate for each connection** option and the connection policy that allows access to multiple servers using HTTPS, the client browsers will display a warning because the certificate used in the connection will be invalid (namely, the Common Name of the certificate will not match the hostname or IP address of the server).

    - To use a separate certificate for every connection, complete the following steps.
        1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
        2. Select **Generate certificate on-the-fly**.
        3. Select the certificate authority to use in the **Signing CA** field.

Limitations

NOTE: When **Generate certificate on-the-fly** is selected and the connection is in transparent setup, the **CN** field is filled in using Server Name Indication (SNI). If the client does not support SNI, the **CN** field will contain the target IP, which may cause certificate verification warning on the client browser.

NOTE: Import the certificate of the signing Certificate Authority to your clients. Otherwise, the client browsers will display a warning due to the unknown Certificate Authority.

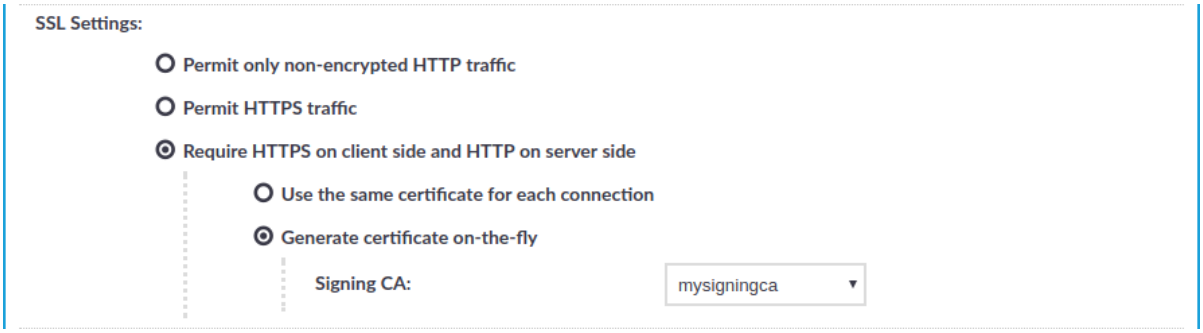
3. Select how SPS should authenticate the server.
    - To permit connections to servers without requesting a certificate, select **No validation**.
    - To permit connections only to servers with a valid certificate that was signed by a specific CA, complete the following steps.

1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the servers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
2. Select **Only accept certificates authenticated by the trusted CA list**.
3. In the **Trusted CA** field, select the certificate authority list to use.

Configuring half-sided SSL encryption in HTTP

The following steps describe how to enable half-sided SSL encryption (which requires HTTPS on client side, and HTTP on server side).

Figure 188: Traffic Controls > HTTP > Connections > SSL Settings — Enabling half-sided TLS encryption in HTTP



To enable half-sided TLS encryption, require HTTPS on client side, and HTTP on server side

1. In **SSL Settings**, select **Require HTTPS on client side and HTTP on server side**.

   NOTE: If the connection is configured at **Target** to **Use fixed address** and the port number is set to `443`, One Identity Safeguard for Privileged Sessions (SPS) will still automatically use port `80` to connect to the server, when **Require HTTPS on client side and HTTP on server side** is selected.
2. Select the certificate to show to the clients.
   - To use the same certificate for every peer, complete the following steps.
     1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
     2. Select **Use the same certificate for each connection**.

3. Select **Private key for host certificate**, click  and upload the private key.

4. Select **X.509 host certificate**, click  and upload the certificate.

NOTE: When using the **Use the same certificate for each connection** option and the connection policy that allows access to multiple servers using HTTPS, the client browsers will display a warning because the certificate used in the connection will be invalid (namely, the Common Name of the certificate will not match the hostname or IP address of the server).

- To use a separate certificate for every connection, complete the following steps.
  1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
  2. Select **Generate certificate on-the-fly**.
  3. Select the certificate authority to use in the **Signing CA** field.

Limitations

NOTE: When **Generate certificate on-the-fly** is selected and the connection is in transparent setup, the **CN** field is filled in using Server Name Indication (SNI). If the client does not support SNI, the **CN** field will contain the target IP, which may cause certificate verification warning on the client browser.

NOTE: Import the certificate of the signing Certificate Authority to your clients. Otherwise, the client browsers will display a warning due to the unknown Certificate Authority.

Session-handling in HTTP

Communication over HTTP consists of client requests and server responses (also called exchanges). Unlike in other protocols, for example SSH, these request-response pairs do not form a well-defined, continuous connection. Therefore, One Identity Safeguard for Privileged Sessions (SPS) assumes that an HTTP request-response pair belongs to a specific session if the following points are true:

- The IP address of the client is the same
- The hostname of the target server (not the IP address) is the same
- The username is the same (if the user has performed inband authentication)
- The time elapsed since the last request-response pair between the same client and server is less than the session timeout value (15 minutes by default).
- The first session cookie SPS finds within the request is the same. Note that the cookie must be listed in the **Session Cookie Settings** option. For more information, see Creating and editing protocol-level HTTP settings.

SPS creates a separate audit trail and records the accessed URLs for every session. These are displayed on the **Sessions** page. If any of the columns is not visible, click **Customize columns...**.

For technical reasons, in authenticated sessions the login page where the user provides the credentials is not part of the session associated with the username. This means that even if the login page is the first that the user visits, SPS will record two sessions: the first does not include a username, the second one does. These two sessions are visible on the **Pending Connections** > **Active Connections** page (until the unauthenticated session times out).

Creating and editing protocol-level HTTP settings

This section describes the HTTP settings that determine the parameters of the connection on the protocol level, including timeout value, session cookies, and TLS settings.

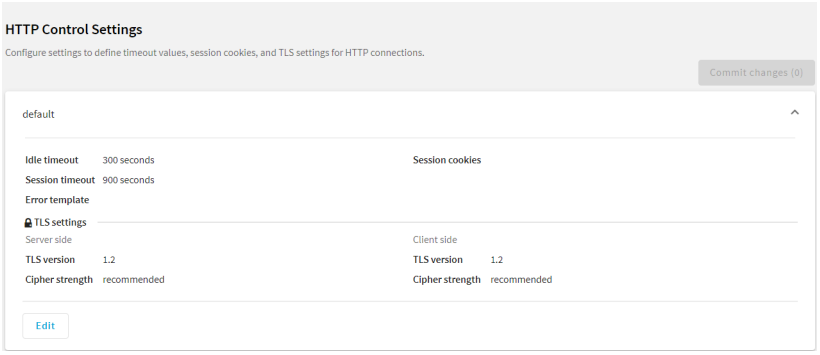
CAUTION:

Modifying timeout settings is recommended only for advanced users. Do not modify these settings unless you exactly know what you are doing.

To create a new HTTP setting profile or edit an existing one

1. Navigate to the **Settings** tab of the **Traffic Controls** > **HTTP** menu item.
2. Click to display the parameters of a profile.

Figure 189: Traffic Controls > HTTP > Settings — Creating and editing protocol-level HTTP settings



3. To create a new HTTP setting, click **Create new**.
4. On **Name and timeout**, enter a name for the new profile and configure the timeout values. The following parameters are available:

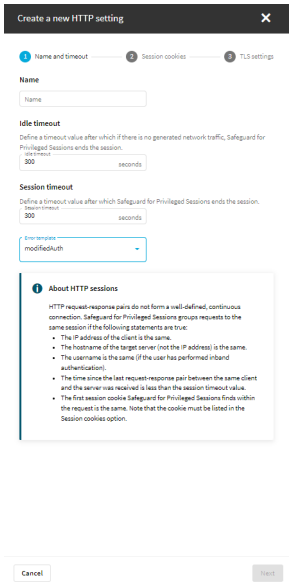
- **Name**: This is a required field. The name must be unique, and the accepted characters are the letters of the English alphabet (A-Z, and a-z) and the underscore (_) character.
- **Idle timeout**: This is a required field. Timeout value for the session in seconds. The accepted values are positive integers. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

  CAUTION:

  Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.
- **Session timeout**: This is a required field. The timeout value specifically for the HTTP session in seconds. HTTP is a request-based protocol. SPS uses this timeout value to group the requests into the same session if the client and target addresses are not different.
- **Error template**: Use the error templates to send customized HTTP error messages to the users. The error templates, created on **Traffic Controls** > **HTTP** > **Error templates**, contain the following data:
  - Customizable HTTP error messages
  - Brand name
  - Color
  - (Optional) Logo

  For more information, see Customizing HTTP error templates.

Figure 190: Traffic Controls > HTTP > Settings > Create a new HTTP setting - Name and timeout

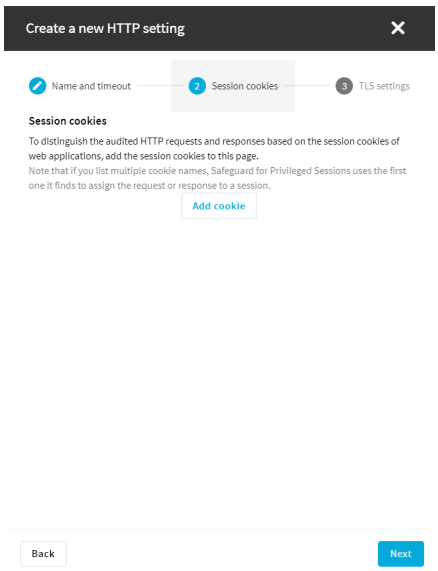


5. Proceed to **Session cookies**.

   To distinguish the audited HTTP requests and responses based on the session cookies of web applications, click **Add cookie**, and enter the name of the session cookie, for example, **`PHPSESSID`**, **`JSESSIONID`**, or **`ASP.NET_SessionId`**. Note that the names of session cookies are case sensitive.

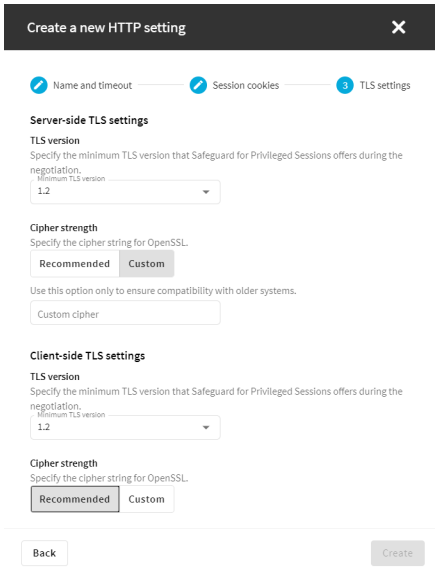
   Repeat this step to add multiple cookie names. Note that if you list multiple cookie names, SPS will use the first one it finds to assign the requests to a session.

Figure 191: Traffic Controls > HTTP > Settings > Create a new HTTP setting - Session cookies



6. To configure TLS security settings on both the **Client side** and the **Server side**, proceed to **TLS security settings**.

Figure 192: Traffic Controls > HTTP > Settings > Create a new HTTP setting - TLS settings



NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported.

- **Minimum TLS version** specifies the minimum TLS version SPS offers during the negotiation. The following options are available:
  - **TLS 1.3**: This setting only accepts TLS version 1.3 during negotiation.
  - **TLS 1.2**: This setting accepts TLS versions 1.2 and later during negotiation.
  - **TLS 1.1**: This setting accepts TLS versions 1.1 and later during negotiation.
  - **TLS 1.0**: This setting accepts TLS versions 1.0 and later during negotiation.
- **Cipher strength** specifies the cipher string for OpenSSL. The following options are available:
  - **Recommended**: this setting only uses ciphers with adequate security level.
  - **Custom**: this setting allows you to specify the list of ciphers you want to permit SPS to use in the connection. This setting is only recommended to ensure compatibility with older systems. For more details on customizing this list, check the 'openssl-ciphers' manual page on your SPS appliance.

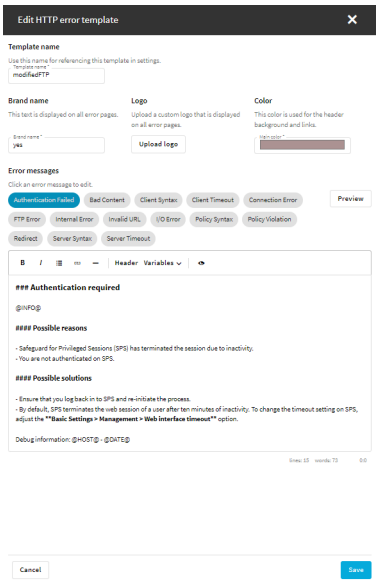
    For example: `ALL:!aNULL:@STRENGTH`

7. To create the new setting and to save it, click **Create** and after that, click **Commit changes**.
8. Select this profile in the **HTTP settings** field of your connections.

Customizing HTTP error templates

This section describes how to create and customize HTTP error templates that you can use to send customized HTTP error messages to the users.

Figure 193: Traffic Controls > HTTP — Error templates



To customize HTTP error templates

1. To create or edit an HTTP error template, navigate to **Traffic Controls** > **HTTP** > **Error templates**.
    a. To create a new error template, select **Create new template** and name your template.

       The new template contains all the error message templates listed on the **Error templates** page.
    b. To modify an error template, select the template and click **Edit**.

    The following error message templates are available:

    - Authentication Failed
    - Bad Content
    - Client Syntax

- Client Timeout
- Connection Error
- FTP Error
- Internal Error
- Invalid URL
- I/O Error
- Policy Syntax
- Policy Violation
- Redirect
- Server Syntax
- Server Timeout

2. Select or modify the brand name, the color, and optionally the logo that is displayed on the customized HTTP error pages.
3. (Optional) Using Markdown language, customize the default content of the error message templates.
4. To check the layout and the content of an error message template, click **Preview**.
5. Save the error template.
6. In **Traffic Controls** > **HTTP** > **Settings**, select a setting, and in the **Error template** option, select the error template that you created or customized in this procedure.

   For more information, see Creating and editing protocol-level HTTP settings.

ICA-specific settings

The following sections describe configuration settings available only for the ICA protocol. Use the following policies to control who, when, and how can access the ICA connection.

NOTE: As an experimental feature, IPv6 addresses can be configured for ICA connections.

- *ICA connections*: For details, see Setting up ICA connections on page 559.
- *Channel Policy*: The channel policy determines which ICA channels (for example clipboard, file-sharing, and so on) can be used in the connection, and whether they are audited or not. The different channels may be available only under certain restrictions, as set in the channel policy. For details, see Creating and editing channel policies on page 495.
- *ICA settings*: ICA settings determine the parameters of the connection on the protocol level, including timeout value and display parameters. For details, see Creating and editing protocol-level ICA settings on page 562.
- *Deployment scenarios*: These describe the available One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment. For details, see One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment on page 563

Setting up ICA connections

This section focuses on describing the ICA-specific details of connection configuration. For a detailed description on configuring connections, see General connection settings on page 482.

CAUTION:

If the clients are accessing a remote application or desktop that is shared for Anonymous users (that is, the **Users** properties of the application is set to **`Allow anonymous users`** in the Citrix Delivery Services Console), the actual remote session will be running under an Anonymous account name (for example, `Anon001`, `Anon002`, and so on), not under the username used to access the remote server. Therefore, you need to enable usermapping to the `Anon*` usernames.

To accomplish this, create a usermapping policy and set the **Username on the server** option to **`Anon*`**, and the **Groups** option to **`*`**, then use this usermapping policy in your ICA connections.

For more information on using usermapping policies, see Configuring usermapping policies on page 860.

Reliable connection is also known as Common Gateway Protocol (CGP). It attempts to reconnect to the server in case of a network failure. To use this feature, enable **Reliable** and enter the default port in the **Port** field in the upper right corner.

Enable **Act as SOCKS proxy** to configure the client to use One Identity Safeguard for Privileged Sessions (SPS) as a SOCKS proxy. If you have enabled this option, you can select **Inband destination selection** as **Target**. Enter the IP address or the IP address/Prefix of the brokers (Citrix XML Brokers) used by the client in this connection policy into the **Address** field. It is also recommended to enable access to the brokers on port 443, as the clients usually try to access the broker using this port first. Disabling port 443 will cause a denied connection to appear on the SPS Search interface for every connection attempt (but the clients will be able to connect the server).

CAUTION:

SPS does not audit or monitor the traffic between the brokers and the clients in any way, and are not listed on the SPS search interface. Only the connections between the clients and the actual servers are audited.

CAUTION:

If SPS is acting as a SOCKS proxy and a client attempts to access a server that it is not permitted to access according to the configuration of SPS, SPS will deny the connection. However, the Citrix client application will automatically attempt to connect the server directly without using a proxy and will succeed if the server is directly accessible from the client. Ensure that your firewalls are configured properly to prevent such connections, as these direct connections cannot be audited by SPS.

NOTE: When enabling **Reliable connection** or **Act as SOCKS proxy** the first time, a warning is displayed suggesting the default port to be used based on the specific settings. Also, read the tooltips on these options as they contain up-to-date information about the default port numbers.

Supported ICA channel types

The available ICA channel types and their functionalities are described below. For a list of supported client applications, see Supported protocols and client applications on page 31.

- **Drawing (Thinwire)**: Enables access to the server's desktop (screen). This channel is for remoting graphics and user input (keyboard, mouse). This channel must be enabled for ICA to work.
- **Audio Mapping**: Enable access to the sound device of the server.
- **Drive Mapping**: Enable access to the client's hard drives on the server.
- **Clipboard**: Enable access to the server's clipboard: the clipboard of the remote desktop can be pasted into local applications (and vice-versa). Note that One Identity Safeguard for Privileged Sessions (SPS) can audit the clipboard channel, but the Safeguard Desktop Player currently cannot search or display its contents.
- **Smartcard**: Enable using client side installed smartcards in server-side applications.
- **Printer (COM1)**: Enable access to the serial port COM1.
- **Printer (COM2)**: Enable access to the serial port COM2.
- **Printer (LPT1)**: Enable access to the parallel port LPT1.
- **Printer (LPT2)**: Enable access to the parallel port LPT2.
- **Printer Spooler**: Enable access to the client's printer from the remote desktops and applications.
- **HDX Mediastream**: Some user widgets (for example Flash player) will not run on the server but on the client. These widgets are controlled from the server side using this channel. This is not supported by Safeguard Desktop Player and it is disabled by default.
- **USB**: Enable using client side installed USB devices in server-side applications.
- **Seamless**: Enable seamless channels that run a single application on the ICA server, instead of accessing the entire desktop. When disabled, the application window will be accessed along with an empty desktop.
- **Speedbrowse**: Speeds up web browsing. Not currently supported by Safeguard Desktop Player, should be disabled by default.
- **Custom**: Applications can open custom channels to the clients connecting remotely to the server. Enabling the **Custom** channel allows the clients to access all of these custom channels. To permit only specific channels, enter the unique names of the channel into the **Details** field.

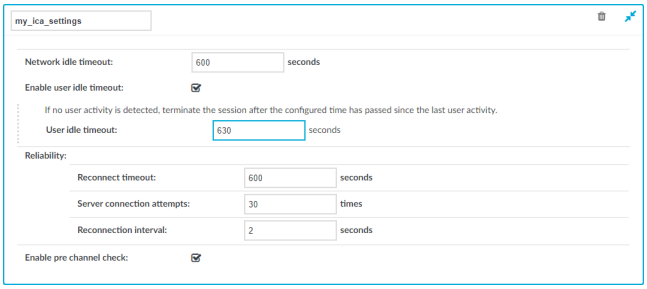
NOTE: When the channel opens, there are certain cases when the remote group is not known yet. For example, in case of an RDP or ICA login screen, the drawing channel has to be opened first to properly display the logon screen. Only those channel rules will apply, where the **Remote group** field is empty. In case of network level authentication, all required information is present already so this limitation does not apply.

NOTE: Multi-stream ICA is not supported in SPS 8.0.1 LTS.

Creating and editing protocol-level ICA settings

ICA settings determine the parameters of the connection on the protocol level, including timeout value, and so on.

Figure 194: Traffic Controls > ICA > Settings — ICA settings



CAUTION:

Modifying the ICA settings is recommended only to advanced users. Do not modify these settings unless you exactly know what you are doing.

To create a new ICA settings profile or edit an existing one

1. Navigate to the **Settings** tab of the **Traffic Controls** > **ICA** menu item and click + to create an ICA setting profile. Enter a name for the profile (for example `ica_special`).
2. Click ⤢ to display the parameters of the ICA connection.
3. Modify the parameters as needed. The following parameters are available:
    - **Network idle timeout**: Connection timeout value in seconds. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

      Even if the user is not active, the session can contain activity that must be audited (for example, the output of a script). The idle timeout period will start only after this activity has stopped.

      CAUTION:

      Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.

- **User idle timeout**: If no user activity is detected, terminate the session after the configured time has passed since the last user activity.

  This can be useful if only user-generated network traffic is important in a session. By using this option, situations described in the caution of **Network idle timeout** (such as a taskbar clock keeping the network traffic open indefinitely) can be avoided. To enable user idle timeout, select **Enable user idle timeout** and enter a value that is greater than or equal to the value of **Network idle timeout**.
- **Reconnect timeout**: How many seconds SPS waits for reconnections when reliable connections are used. Reliable connections use the Common Gateway Protocol (CGP).
- **Server connection attempts**: How many times SPS tries to connect to the target server.
- **Reconnection intervals**: How many seconds SPS waits between two connection attempts on the server side.
- **Enable pre channel check**: Select this option to evaluate the connection and channel policies before establishing the server-side connection. That way if the connection is not permitted at all, SPS does not establish the server-side connection.

NOTE: Reliability settings only apply if you have enabled **Reliable connection** in **Traffic Controls** > **ICA** > **Connections**.

4. Click **Commit**.
5. Select this settings profile in the **ICA settings** field of your connections.

One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment

This section enlists the available One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a Citrix environment. The text on the arrows are formatted in (<step number>) <target port> format. The target ports define the protocols used in the communication:

- *80*: Web service, HTTP: the list of available resources fetched in an XML format from the broker (v12 and v11 with Citrix Virtual Apps (formerly known as Citrix XenApp) only). The broker sends all the necessary information, including secure gateway and server addresses to the client.
- *8080*: XML service, HTTP+XML: application discovery, load balancing (v12 and v11 with Citrix Virtual Apps (formerly known as Citrix XenApp) only), used to fetch target to the application/desktop by the client from the broker (used for load balancing, and so on).

- *443*: XML service access or SOCKS/ICA or CGP/ICA wrapped in TLS. The client communicates with the secure gateway on this port for everything.
- *1080*: SOCKS. The client can be configured to access the target server and the broker using a SOCKS proxy.
- *1494*: Plain ICA.
- *2598*: CGP/ICA (reliable mode enabled).

CAUTION:

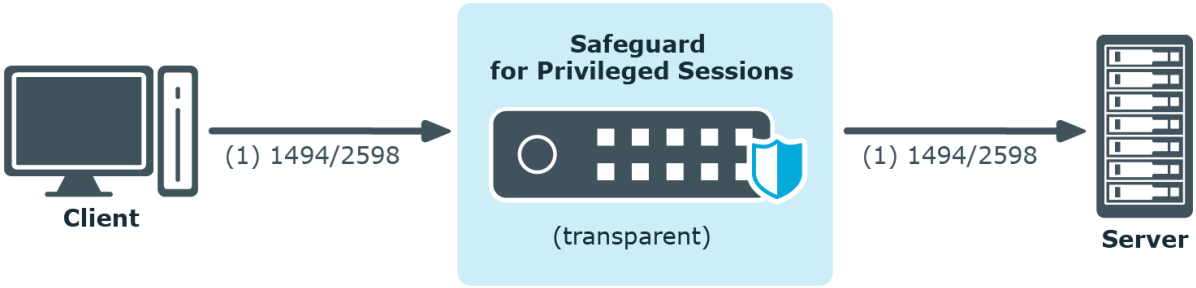
Accessing Citrix Virtual Desktops (formerly known as Citrix XenDesktop) is supported only in the following scenarios. Only reliable connections (CGP) are supported.

- Client - Broker - SPS - Server (Transparent mode) on page 565
- Client - Broker - SPS as socks proxy - Server on page 566

Client - SPS - Server (Transparent mode)

The SPS is deployed between the client and the server and the clients use predefined connection files or Program Neighbourhood, without a broker or secure gateway. The clients try to connect to their original ICA/CGP server.

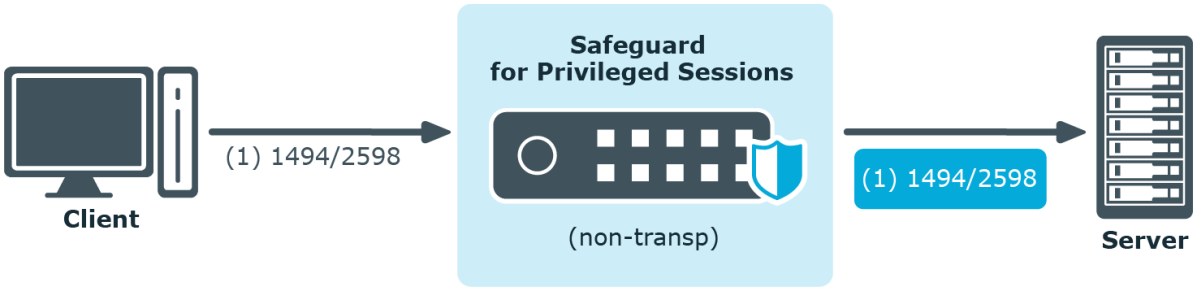
Figure 195: Client - SPS - Server (Transparent mode)



Client - SPS - Server (Non-transparent mode)

The SPS is deployed between the client and the server and the clients use predefined connection files or Program Neighbourhood, without a broker or secure gateway. The clients try to connect to SPS, which can distinguish between the potential targets for example by source IP, or by having multiple IP addresses itself.

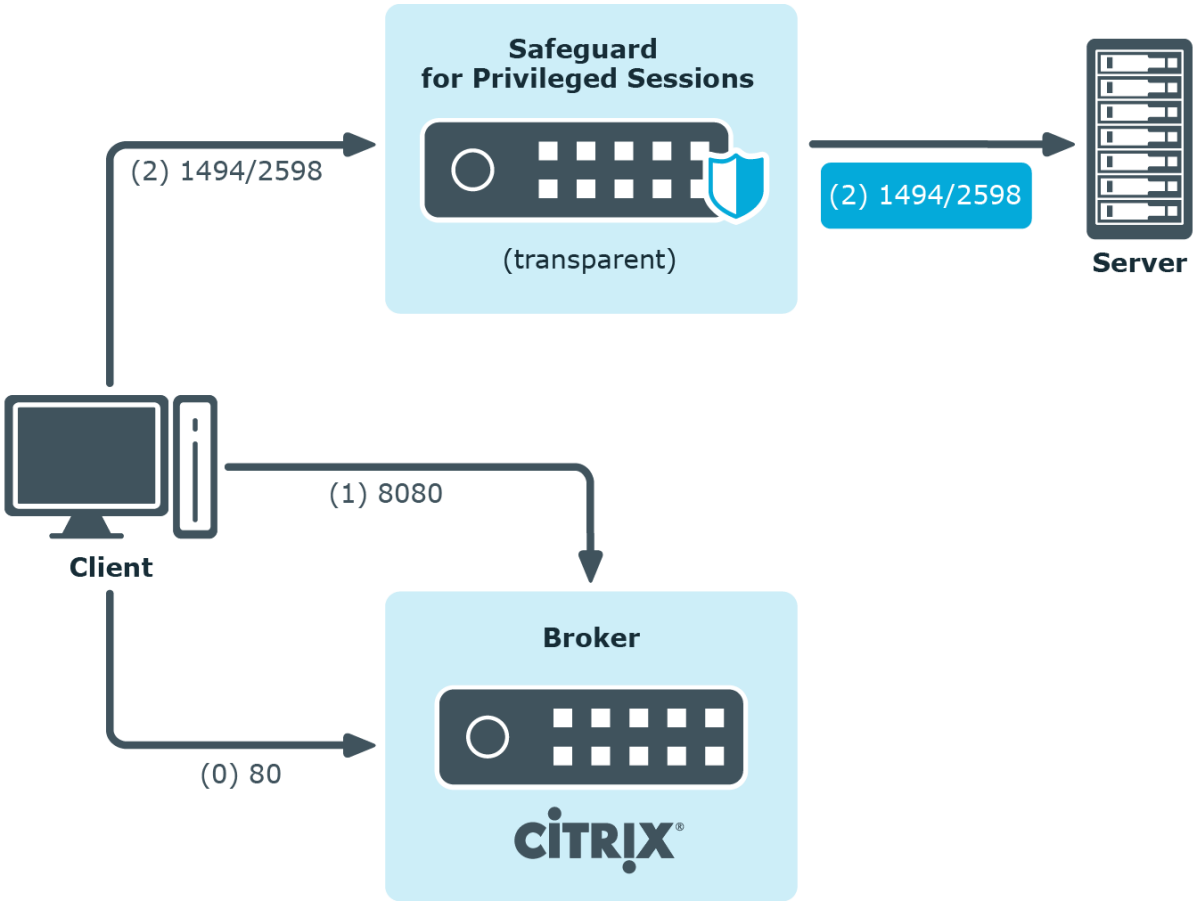
Figure 196: Client - SPS - Server (Non-transparent mode)



Client - Broker - SPS - Server (Transparent mode)

The clients are using a farm broker which gives them a list of the available applications and servers, but they do not use a secure gateway in the network. The SPS is placed between the clients and the servers in transparent mode, and it catches the connections when the clients try to connect to the server IP addresses they got from the broker.

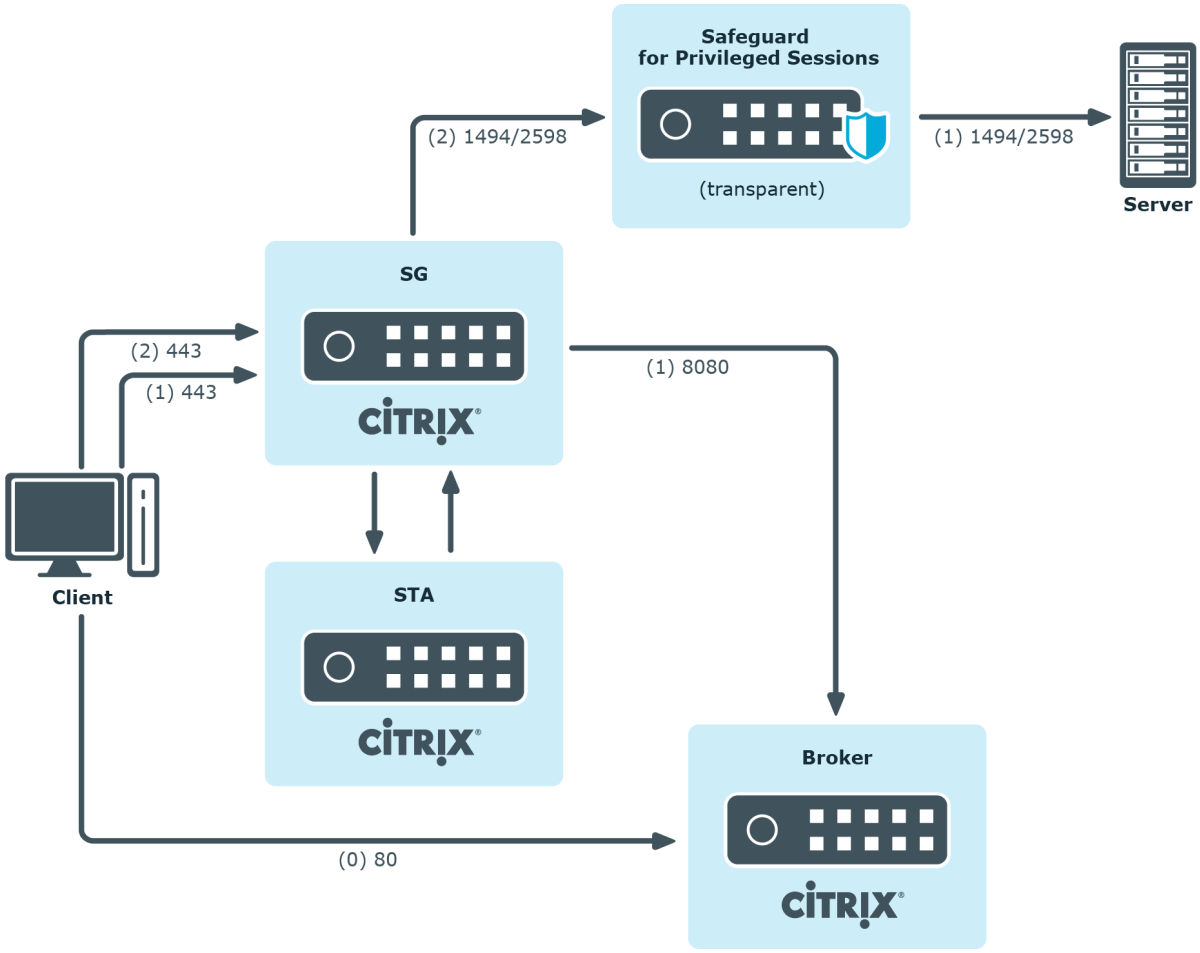
Figure 197: Client - Broker - SPS - Server (Transparent mode)



Client - Broker - original secure gateway - Secure Ticket Authority (STA) - SPS - Server

In this setup, a secure gateway is used in the network and the SPS is placed between this gateway and the servers in transparent mode. The clients connect to the broker for the list of available applications/servers and then make their further connections through the original secure gateway. That gateway forwards the connections either to the broker or to the CGP/ICA servers, which latter the SPS intercepts and audits/controls.

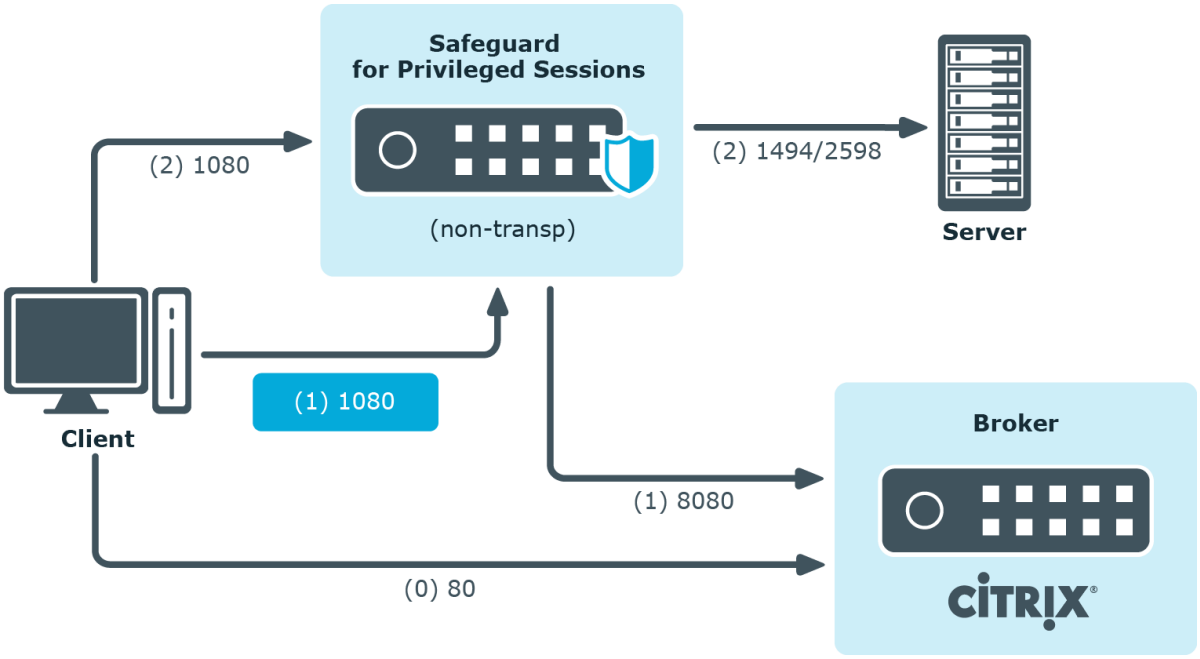
Figure 198: Client - Broker - original secure gateway - Secure Ticket Authority (STA) - SPS - Server



Client - Broker - SPS as socks proxy - Server

In this setup, the SPS acts as a SOCKS proxy for the client. It can be set either manually or specified by the broker. The client then makes all its connections to the broker or to the server using SPS as a proxy and hence it can audit/control these connections.

Figure 199: Client - Broker - SPS as socks proxy - Server



To configure such a scenario, you must set the ICA Connection Policy as follows:

- Enter the IP address of SPS into the **To** field. This must be the public IP address that the clients will target.
- Select **Inband destination selection**, and list the IP addresses or networks of target servers in the **Targets** field. (For details, see Configuring inband destination selection on page 490.)
- Select **Act as a SOCKS proxy**.
- Add the IP addresses of your brokers to the **Brokers** field.

Troubleshooting Citrix-related problems

Accessing Citrix servers using the Remote Desktop Protocol

Accessing Citrix servers using the Remote Desktop Protocol may fail in certain situations, and the connection is terminated with the `ERROR: error while decompressing packet` error message on the client, or with the `Event56, TermDD, The Terminal Server security layer detected an error in the protocol stream and has disconnected the client.` message on the server.

To overcome this problem, modify the settings of the network card of the server, and disable the **Large Send Offload** option.

NOTE: The problem is not related to using One Identity Safeguard for Privileged Sessions (SPS) in your environment.

MSSQL-specific settings

The following sections describe configuration settings available only for the MSSQL protocol. Use the following policies to control who, when, and how can access the MSSQL connection. For a list of supported client applications, see Supported protocols and client applications on page 31.

- *Channel Policy*: The MSSQL protocol has only one channel type with no special configuration options. The available channel policy options are the following: **Type**, **From**, **Target**, **Time policy**, **Four-eyes**, **Record audit trail**, **Gateway groups**, **Remote groups**, and **Content policy**. For details on configuring these options, see Creating and editing channel policies on page 495.
- *TLS support*: To enable TLS-encryption for your MSSQL connections, see Enabling TLS-encryption for MSSQL connections on page 574.
- *Authentication Policy*: Authentication policies describe the authentication methods allowed in a connection. Different methods can be used for the client and server-side connections. For details, see Creating a new MSSQL authentication policy on page 570.
- *MSSQL settings*: MSSQL settings determine the parameters of the connection on the protocol level, including timeout value, and so on. For details, see Creating and editing protocol-level MSSQL settings on page 572.
- *User lists in Channel Policies*: User lists affect MSSQL connections only when they are used together with Gateway Authentication. For details, see Configuring gateway authentication on page 862.
- *Content Policy*: Content policies allow you to inspect the content of the connections for various text patterns, and perform an action if the pattern is found. For example, One Identity Safeguard for Privileged Sessions (SPS) can send an e-mail alert if a specific command is used in a MSSQL terminal session. For details, see Creating a new content policy on page 499.
- *Authentication and Authorization plugin*:

  One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS to external systems to authenticate or authorize the user before authenticating on the target server. Such plugins can also be used to request additional information from the users, for example, to perform multi-factor authentication.

  For details, see Integrating external authentication and authorization systems on page 884.

Setting up MSSQL connections

This section focuses on describing the MSSQL-specific details of connection configuration. For a detailed description on configuring connections, see General connection settings on page 482.

Limitations in handling MSSQL connections

The current version of One Identity Safeguard for Privileged Sessions (SPS) has the following limitations:

- TDS protocol version 7.3 or later is required.
- Due to the TDS protocol version requirement, Microsoft® SQL Server® 2008, or later, is recommended.
- The **Require Gateway Authentication on the SPS Web Interface** option in **Traffic Controls** > **MSSQL** > **Connections** does not work in case of MSSQL connections.
- MSSQL server with TCP dynamic port settings is not supported.

  You must specify a static TCP port for every instance in the SQL Server Configuration Manager you want to audit. By doing so, you can configure the access to multiple MSSQL instances with multiple connection policies and specify the instances with inband or fixed targets and ports. You can also create and assign different Credential Store policies to check out SQL users' passwords of the instances.

  In the MSSQL client program, always specify the address with the port number of the SPS connection policy you want to connect to.

Supported MSSQL channel types

The available MSSQL channel types and their functionalities are described below. For details on configuring Channel Policies, see Creating and editing channel policies on page 495. For a list of supported client applications, see Supported protocols and client applications on page 31.

- **MSSQL**: Enables you to use the MSSQL protocol. This channel must be enabled for MSSQL to work.

  The available channel policy options are the following: **From**, **Target**, **Time policy**, **Four-eyes**, **Record audit trail**, **Gateway groups**, **Remote groups**, and **Content policy**. Note that the **Gateway groups** option is used only if the user performs inband authentication using one of the supported MSSQL authentication methods (see Authentication in MSSQL on page 570). To retrieve the groups of an authenticated user from an LDAP database, you must also set an **LDAP Server** in the Connection Policy (for MSSQL connections, One Identity Safeguard for Privileged Sessions (SPS) uses this server only to retrieve the group membership of authenticated users, you cannot authenticate the users to LDAP from SPS). For details on configuring these options, see Creating and editing channel policies on page 495.

  When setting **Target**, note the following:

  - If the connection uses DNAT (NAT destination address), the target address of the original client will be compared to the **Target** parameter of the Channel

policy, that is not necessarily equivalent with the server's address.

- If the connection is redirected to a Fix address, the redirected address will be compared to the **Target** parameter of the Channel policy.

Authentication in MSSQL

For the audited MSSQL connections, SPS(SPS) supports SQL Server Authentication.

SPS records the username used in the authentication process into the **Username** and **Remote username** fields of the connection database.

For authenticated sessions, SPS can perform group-based user authorization that allows you to finetune access to your servers and services: you can set the required group membership in the Channel policy of the MSSQL connection.

NOTE: Group-based authorization in MSSQL works only for authenticated sessions (for MSSQL connections, SPS uses this server only to retrieve the group membership of authenticated users, you cannot authenticate the users to LDAP from SPS). If a user name is not available for the session, SPS will permit the connection even if the **Remote groups** field is set.

Note that Group-based authorization in MSSQL works only for authenticated sessions (for MSSQL connections, SPS uses this server only to retrieve the group membership of authenticated users, you cannot authenticate the users to LDAP from SPS). If a username is not available for the session, SPS will permit the connection even if the **Remote groups** field is set.

SPS does not store failed MSSQL authentication attempts in the connection database. This means that the **Verdict** field of the Search page will never contain `CONN-AUTH-FAIL` values for MSSQL connections.

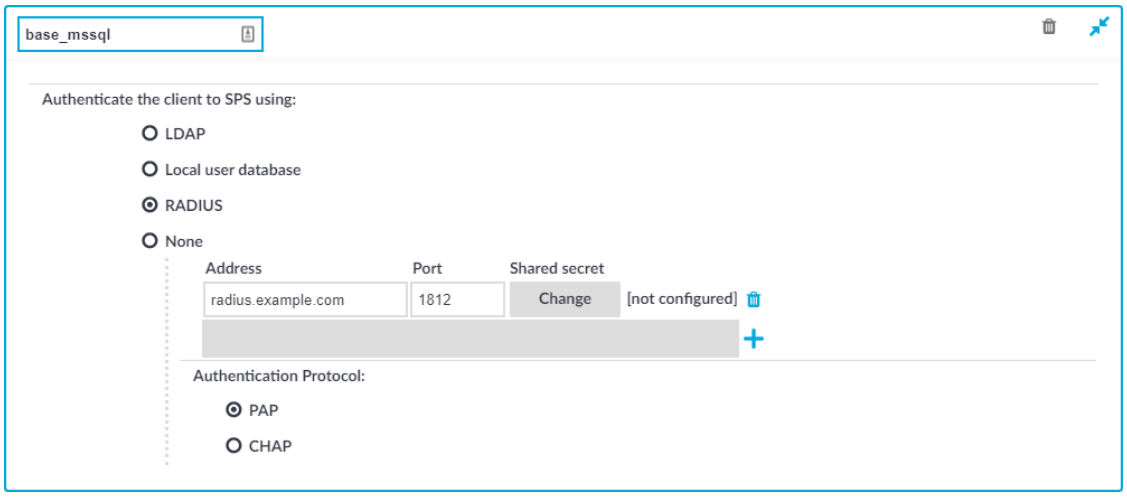
Creating a new MSSQL authentication policy

An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the client can authenticate to the target server. Separate authentication methods can be used on the client and the server-side of the connection.

To create a new authentication policy

1. Navigate to **Traffic Controls** > **MSSQL** > **Authentication Policies**, and click + .

Figure 200: Traffic Controls > MSSQL > Authentication Policies — Configuring MSSQL authentication policies



2. Enter a name for the policy into the **Name** field.
3. Select the authentication method used on the client-side in the **Authenticate the client to SPS using** field. For the client-side connection, SPS can authenticate the client inband (within the MSSQL protocol) using the following authentication methods:
   - **LDAP**: SPS will authenticate the client to the LDAP database set in the **LDAP Server** of the connection policy. To use LDAP authentication on the client side, select **Authenticate the client to SPS using** > **LDAP**.

     NOTE: SPS will authenticate the client-side connection to the LDAP server configured in the connection policy. This is not necessarily the same as the LDAP server used to authenticate the users accessing the SPS web interface.
   - **Local user database**: Authenticate the client locally on the SPS gateway using a **Local user database**. Select the database to use in the **Local user database** field. For details on creating a **Local User Database**, see Creating a Local User Database on page 538.
   - **RADIUS**: SPS will authenticate the client to the specified RADIUS server. Select **Authenticate the client to SPS using** > **RADIUS**, enter the IP address or hostname of the RADIUS server into the **Address** field, the port number of the RADIUS server into the **Port** field, and the shared secret of the RADIUS server into the **Shared secret** field. Only password-authentication is supported (including one-time passwords), challenge-response based authentication is not.

     Use an IPv4 address.

To add more RADIUS servers, click  and fill in the respective fields.

- **None**: Do not perform client-side authentication, the client will authenticate only on the target server.

  CAUTION:

  Hazard of security breach. If the **None** authentication option is selected on the client side and SPS is configured to use public-key or certificate based authentication on the server, the user will not be authenticated at all unless gateway authentication is required for the connection.

4. Click **Commit**.

Creating and editing protocol-level MSSQL settings

Procedure

MSSQL settings determine the parameters of the connection on the protocol level, including timeout value, and so on. Complete the following procedure to create a new MSSQL settings profile or edit an existing one:

CAUTION:

Modifying the MSSQL settings is recommended only to advanced users. Do not modify these settings unless you exactly know what you are doing.

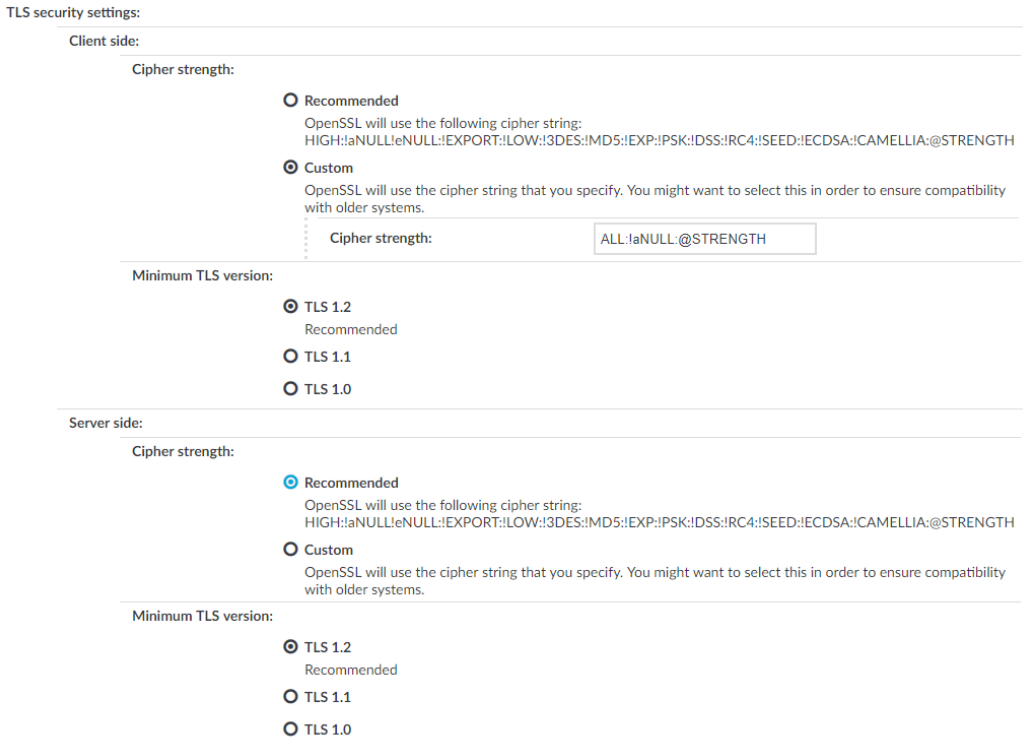
To create and edit protocol-level MSSQL settings

1. Navigate to the **Settings** tab of the **Traffic Controls** > **MSSQL** menu item and click  to create a MSSQL setting profile. Enter a name for the profile (for example, `mssql_special`).
2. Click  to display the parameters of the connection.
3. Modify the parameters as needed. The following parameters are available:
   - **Idle timeout**: Timeout value for the connection in seconds. To avoid early timeout, set it to a larger value, for example, a week (604800 seconds).

CAUTION:

Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.

- **Enable pre channel check**: Select this option to evaluate the connection and channel policies before establishing the server-side connection. That way if the connection is not permitted at all, SPS does not establish the server-side connection.
- To configure TLS security settings on both the **Client side** and the **Server side**, proceed to **TLS security settings**.**<Protocol> Control > Settings > TLS security settings** - configuring TLS security settings



Cipher strength specifies the cipher string OpenSSL will use. The following options are possible:**Minimum TLS version** specifies the minimal TLS version SPS will offer during negotiation. The following options are possible:NOTE: Setting up sessions to legacy systems that do not support at least TLS 1.2 is only possible when the security level of the connection is degraded to 0, which is possible by specifying the TLS ciphers manually and appending the string `:@SECLEVEL=0` to the cipher list. However, this setting also enables the use of known vulnerable algorithms and key sizes, therefore it is absolutely critical

to only use such connection settings when it is necessary and when you can fully trust your network between SPS and the legacy system. It is strongly recommended to use different security settings on the server and the client side of the connection, when degrading the security level of a connection is unavoidable. NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported..

4. Click 

.
5. Select this settings profile in the **MSSQL settings** field of your connections.

Enabling TLS-encryption for MSSQL connections

The following steps describe how to enable TLS-encryption in a MSSQL connection policy. Note that when using encryption, One Identity Safeguard for Privileged Sessions (SPS) automatically changes the port number of the connection policy to 992.

Prerequisites

Depending on your requirements, one or more of the following might be needed:

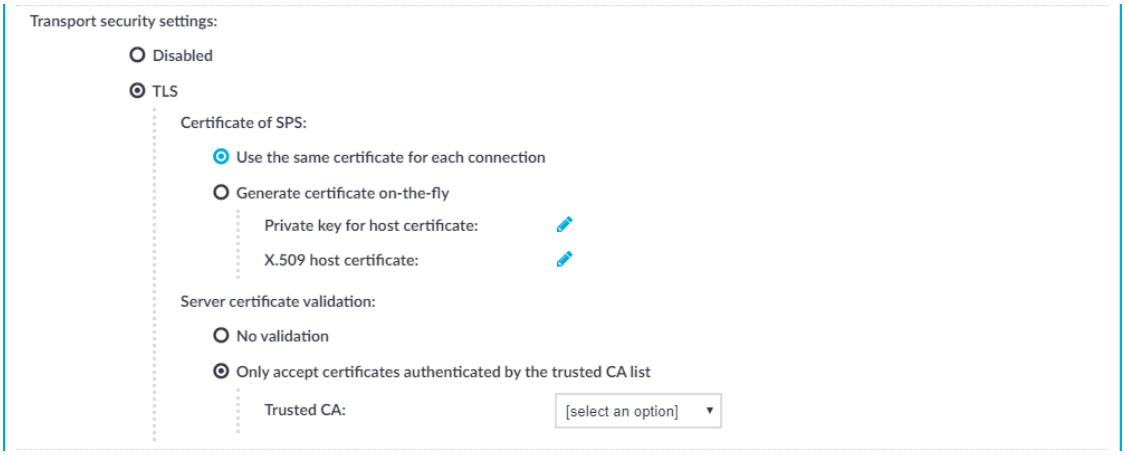
- An X.509 certificate and its private key. SPS can display the same certificate to the peers on both the client and the server side. You can also use different certificates for the client and server sides. Use your own PKI system to generate these certificates, as they cannot be created on SPS. Note that the Common Name of the certificate must contain the domain name or the IP address of SPS, otherwise the clients might reject the certificate.
- To generate certificates on-the-fly for a connection, a signing certificate authority is required. For details on creating a signing CA, see Signing certificates on-the-fly on page 529.
- To require the peers of SPS to have an X.509 certificate signed by a specific Certificate Authority, a list of the trusted certificate authorities is needed. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To enable TLS-encryption in a MSSQL connection policy

1. Navigate to **Traffic Controls** > **MSSQL** > **Connections** and select the connection policy in which you want to enable TLS.

Figure 201: Traffic Controls > MSSQL > Connections — Enabling TLS-encryption for MSSQL connections



2. Set the encryption settings in the **Transport security settings** section.
   - To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.
3. Select the certificate to show to the peers.
   - To use the same certificate for every peer, complete the following steps.
     1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
     2. Select **Use the same certificate for each connection**.
     3. Select **Private key for host certificate**, click and upload the private key.
     4. Select **X.509 host certificate**, click and upload the certificate.
   - To use a separate certificate for every connection, complete the following steps.
     1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
     2. Select **Generate certificate on-the-fly**.
     3. In the **Signing CA** field, select the certificate authority to use.
4. Select which certificatie validation method SPS should use in the **Server certificate validation** section.
   - To permit connections from peers without requesting a certificate, select **No validation**.

- To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
    1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
    2. Select **Only accept certificates authenticated by the trusted CA list**.
    3. Select the certificate authority list to use in the **Trusted CA** field.

5. Click **Commit**.

Expected result

The encryption settings are applied to the connection policy.

RDP-specific settings

The following sections describe configuration settings available only for the RDP protocol. Use the following policies to control who, when, and how can access the RDP connection.

- *Channel Policy*: The channel policy determines which RDP channels (for example clipboard, file-sharing, and so on) can be used in the connection, and whether they are audited or not. The different channels may be available only under certain restrictions, as set in the channel policy. For details, see Creating and editing channel policies on page 495.
- *RDP settings*: RDP settings determine the parameters of the connection on the protocol level, including timeout value, display parameters, and the version of RDP permitted. For details, see Creating and editing protocol-level RDP settings on page 582.
- *Domain membership*: When using Network Level Authentication (CredSSP) One Identity Safeguard for Privileged Sessions (SPS) must be a member of the domain. For details, see Network Level Authentication (NLA) with domain membership on page 588.
- *TLS-encrypted connections*: For details on how to setup TLS-encrypted RDP connections, see Enabling TLS-encryption for RDP connections on page 592 and Verifying the certificate of the RDP server in encrypted connections on page 591.
- *SPS as a Remote Desktop Gateway*: For details on how to configure SPS to accept connections using the Remote Desktop Gateway Server Protocol, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.
- *Content Policy*: Content policies allow you to inspect the content of the connections for various text patterns, and perform an action if the pattern is found. For example, SPS can send an e-mail alert if a specific window title appears in RDP and VNC connections. For details, see Creating a new content policy on page 499.
- *Authentication and Authorization plugin*:

  One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS to external systems to authenticate or authorize the user before authenticating on the target server. Such plugins can also be used to request additional information from the users, for example, to perform multi-factor authentication.

  For details, see Integrating external authentication and authorization systems on page 884.

Using multiple monitors (Multimon) is supported. To enable Multimon, use one of the following three methods:

- enable **Display** > **Use all my monitors for the remote session** option in the Remote Desktop Client (`mstsc.exe`) window of the client machine
- use the `/multimon` switch on the `mstsc.exe` command line
- add the `use multimon:i:1` row to the RDP file

NOTE: The **Maximum display width** and **Maximum display height** options should be high enough to cover the combined resolution of the client monitor setup. Connections that exceed these limits will automatically fail. Make sure to adjust these settings if your clients use multiple monitors. For example, if your clients use two monitors that have a resolution of 1920x1080 pixels each, set **Maximum display width** to 4000, and **Maximum display height** to 2200.

Limitations

The RDP connection fails due to the following Windows-side settings:

- If the target user is a member of the Protected Users security group.
- If Remote Credential Guard is enabled and used.
- If Restricted Admin mode is enabled and used. For more information on how to enable or disable the Restricted Admin mode on Windows, see Remote Desktop Services: Enable Restricted Admin mode.

NOTE: Due to the way RDP handles device redirection, these channels work only if the **Sound** channel type is also enabled. Make sure that you enable the **Sound** channel if you enable one of the specific redirection types, for example, **Serial**, **Parallel**, **Printer**,**Disk**, **SCard**, or **Custom** redirect.

Kerberos

This section describes the machine account password processes (including the creation of the machine account for SPS), the mandatory password change periods, the password rotation details, the details of the Service Principal Name for SPS and inband target selection.

Machine Account Password Process

When SPS joins a domain, a machine account is created for SPS. By default, Active Directory requires password change for machine accounts every 30 days. Since SPS requires an up-to-date machine account with a current password, SPS forces the machine account password change every 5 days.

If your machine account passwords are rotated more frequently than 5 days, users might experience Kerberos login failures. If your password policy requires more frequent machine account password rotation, contact our Support Team.

Service Principal Name (SPN) for SPS

When a user wants to authenticate with Kerberos, the server has to be addressed with FQDN so the client program can search for the SPN of the server (`TERMSRV/target-server.target-realm`). Therefore, during the process, SPS domain join registers `TERMSRV` SPN for its computer account (`TERMSRV/<SPS-hostname>.<realm-fqdn>`) address.

Make sure that `<SPS-hostname>.<realm-fqdn>` is resolvable from DNS.

Inband target selection

Since the user has to address the SPS during the Kerberos authentication, the connection policy that uses Kerberos authentication has to use inband target selection so that SPS can forward the incoming service ticket to the required target server.

Supported RDP channel types

The available RDP channel types and their functionalities are described below. For details on configuring Channel Policies, see Creating and editing channel policies on page 495. For a list of supported client applications, see Supported protocols and client applications on page 31.

- **Drawing**: Enables access to the server's graphical desktop (screen). This channel must be enabled for RDP to work.

  NOTE: In case the Drawing channel is disabled and the load of One Identity Safeguard for Privileged Sessions (SPS) is high, or the connection requires four-eyes authorization and the Authorizer is slow to accept the connection, the client might receive the following error message:

  ```
  The Remote Desktop Gateway server administrator has ended the connection.
  Try reconnecting later or contact your network administrator for assistance
  ```

- **Clipboard**: Enables access to the server's clipboard: the clipboard of the remote desktop can be pasted into local applications (and vice-versa). Note that SPS can audit the clipboard channel, and that files transferred via the clipboard can be audited Configuring SPS to enable exporting files from audit trails after RDP file transfer through clipboard or disk redirection.

  If the **Clipboard** channel is enabled, it implicitly enables copying files as well, as the user can simply copy-paste the file. Copy-pasted files will not be visible in the logs or the **File operations** column of the **Sessions** page. To ensure that SPS records file transfer events, you must disable the **Clipboard** channel.

- **Redirects**: Enables access to every device redirection available in RDP, like file-sharing, printer sharing, device (for example, CD-ROM) sharing, and so on.

- To make the list of file operations available in the **File operations** column of the **Sessions** page, navigate to the Channel Policies page of the protocol, and enable the **Log file transfers to database** option. This option is disabled by default.
- To send the file operations into the system log, enable the **Log file transfers to syslog** option. This option is disabled by default.

  NOTE: Turning logging on might result in a slight performance penalty. If traffic load slows processes down, disable the option.

To enable only specific types of redirections, use the following channels:

NOTE: Due to the way RDP handles device redirection, these channels work only if the **Sound** channel type is also enabled. Make sure that you enable the **Sound** channel if you enable one of the specific redirection types, for example, **Serial**, **Parallel**, **Printer**,**Disk**, **SCard**, or **Custom** redirect.

- **Serial redirect**: Enables access to serial-port redirections.
- **Parallel redirect**: Enables access to parallel-port redirections.
- **Printer redirect**: Enables access to shared printers.

  When enabling printer redirection, you may need to use TSVCTKT and XPSRD channels — these enable XPS printing.

  Note that these channels are dynamic virtual channels and you have to be enable them using the **Custom** channel type.

  For more information on TSVCTKT and XPSRD channels, see *section 2.1 Transport* in Microsoft Technical Document [MS-RDPEXPS].

  Before consulting the cited Microsoft Technical Document, it is recommended to start by reading [MS-RDSOD]: Remote Desktop Services Protocols Overview.
- **Disk redirect**: Enables access to shared disk drives.
  - To make the list of file operations available in the **File operations** column of the **Sessions** page, navigate to the Channel Policies page of the protocol, and enable the **Log file transfers to database** option. This option is disabled by default.
  - To send the file operations into the system log, enable the **Log file transfers to syslog** option. This option is disabled by default.

    NOTE: Turning logging on might result in a slight performance penalty. If traffic load slows processes down, disable the option.

  If the **Clipboard** channel is enabled, it implicitly enables copying files as well, as the user can simply copy-paste the file. Copy-pasted files will not be visible in the logs or the **File operations** column of the **Sessions** page. To ensure that SPS records file transfer events, you must disable the **Clipboard** channel.
- **SCard redirect**: Enables access to shared SCard devices.

To permit only specific redirections, enter the unique name of the redirection into the **Details** field. For example, if you want to enable access only to the shared disk drive

C:, enable the **Disk redirect** channel and enter `C:` into the **Permitted devices** field. Note that the name of the device comes from the device itself, so it is case sensitive, and may not always be reliable from a security point of view.

- **Sound**: Enables access to the sound that is played on the server.

  NOTE: Due to the way RDP handles device redirection, these channels work only if the **Sound** channel type is also enabled. Make sure that you enable the **Sound** channel if you enable one of the specific redirection types, for example, **Serial**, **Parallel**, **Printer**,**Disk**, **SCard**, or **Custom** redirect.

  To enable auditing the sound that is transferred between an RDP client and the server, make sure that the **Record audit trail** checkbox is selected for the **Sound** and the **Dynamic virtual channel**.

  NOTE: If the **Record audit trail** checkbox is selected only for the **Sound** channel, only the output sound (the one that is received by the audited user) is recorded in the audit trail.

  If the **Record audit trail** checkbox is selected for the **Dynamic virtual channel**, both the input sound (the one that comes from the audited user) and output sound (the one that is received by the audited user) are recorded.

- **Custom**: Applications can open custom channels to the clients connecting remotely to the server. Enabling the **Custom** channel allows the clients to access all of these custom channels. To permit only specific channels, enter the unique names of the channel into the **Permitted devices** field.

  For example, to monitor RemoteApp connections, you need to configure custom channels. For more information, see Configuring RemoteApps on page 605.

- **Seamless**: Enables seamless channels that run a single application on the RDP server, instead of accessing the entire desktop.

- **Dynamic virtual channel**: Enables the server to open channels back to the client dynamically. To restrict which dynamic channels are permitted, select **Channel details**, click + and enter the name of the permitted channel.

  Additionally, you may need to use one or more of the following:

  - PNPDR and FileRedirectorChannel channels: Enable Plug and Play devices.

    For more information, see *section 2.1 Transport* in Microsoft Technical Document [MS-RDPEPNP].

  - URBDRC channels: Enable USB redirection.

    For more information, see *section 2.1 Transport* in Microsoft Technical Document [MS-RDPEUSB].

  Before consulting any of the listed Microsoft Technical Documents, it is recommended to start by reading [MS-RDSOD]: Remote Desktop Services Protocols Overview.

  To enable auditing the sound that is transferred between an RDP client and the server, make sure that the **Record audit trail** checkbox is selected for the **Sound** and the **Dynamic virtual channel**.

  NOTE: If the **Record audit trail** checkbox is selected only for the **Sound** channel, only the output sound (the one that is received by the audited user) is recorded in

the audit trail.

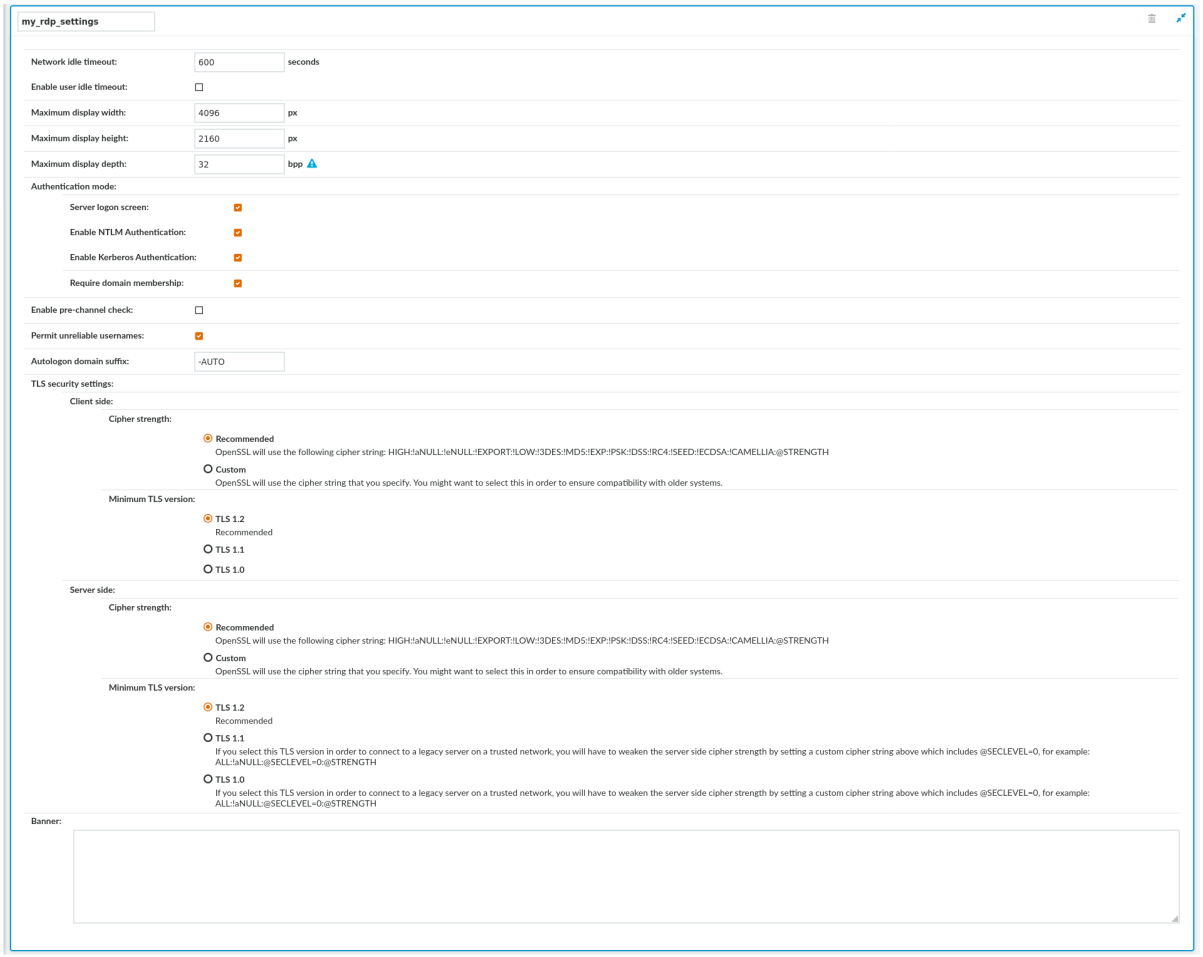
If the **Record audit trail** checkbox is selected for the **Dynamic virtual channel**, both the input sound (the one that comes from the audited user) and output sound (the one that is received by the audited user) are recorded.

NOTE: When the channel opens, there are certain cases when the remote group is not known yet. For example, in case of an RDP or ICA login screen, the drawing channel has to be opened first to properly display the logon screen. Only those channel rules will apply, where the **Remote group** field is empty. In case of network level authentication, all required information is present already so this limitation does not apply.

Creating and editing protocol-level RDP settings

RDP settings determine the parameters of the connection on the protocol level. For example, timeout value, the supported authentication modes, and display parameters.

Figure 202: Traffic Controls > RDP > Settings — RDP settings



CAUTION:

Modifying the RDP settings is recommended only to advanced users. Do not modify these settings unless you exactly know what you are doing.

To create a new RDP settings profile or edit an existing one

1. Navigate to **Traffic Controls** > **RDP** > **Settings** and click  to create an RDP setting profile. Enter a name for the profile (for example, **rdp5only**).
2. Click  to display the parameters of the RDP connection.
3. Modify the parameters as needed. The following parameters are available:
   - **Network idle timeout**: Connection timeout value in seconds. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

     Even if the user is not active, the session can contain activity that must be audited (for example, the output of a script). The idle timeout period will start only after this activity has stopped.

     CAUTION:

     Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.

     CAUTION:

     If the value is set below 31 seconds, MSTSC can fail and prevent new connections if **Act as a Remote Desktop Gateway** is enabled in **Traffic Controls** > **RDP** > **Connections**. To prevent this, set the **Idle timeout** value to at least **31** seconds.
   - **User idle timeout**: If no user activity is detected, terminate the session after the configured time has passed since the last user activity.

     This can be useful if only user-generated network traffic is important in a session. By using this option, situations described in the caution of **Network idle timeout** (such as a taskbar clock keeping the network traffic open indefinitely) can be avoided. To enable user idle timeout, select **Enable user idle timeout** and enter a value that is greater than or equal to the value of **Network idle timeout**.
   - **Maximum display width**: The maximum allowed width of the remote desktop in pixels (for example 1024).

     NOTE: The **Maximum display width** and **Maximum display height** options should be high enough to cover the combined resolution of the client

monitor setup. Connections that exceed these limits will automatically fail. Make sure to adjust these settings if your clients use multiple monitors. For example, if your clients use two monitors that have a resolution of 1920x1080 pixels each, set **Maximum display width** to 4000, and **Maximum display height** to 2200.

- **Maximum display height**: The maximum allowed height of the remote desktop in pixels (for example 768).

  NOTE: The **Maximum display width** and **Maximum display height** options should be high enough to cover the combined resolution of the client monitor setup. Connections that exceed these limits will automatically fail. Make sure to adjust these settings if your clients use multiple monitors. For example, if your clients use two monitors that have a resolution of 1920x1080 pixels each, set **Maximum display width** to 4000, and **Maximum display height** to 2200.

- **Maximum display depth**: The maximum allowed color depth of the remote desktop in bits (for example 24). The following values are valid: 8, 15, 16, 24.

  CAUTION:

  - Using 32-bit color depth is currently not supported: client connections requesting 32-bit color depth automatically revert to 24-bit.
  - Certain Windows versions do not support 24-bit color depth. In this case, those versions can only be displayed in 16-bit color depth. SPS automatically changes its settings to 16-bit.

- **Authentication mode**: Enable the authentication option that best matches the authentication method supported by your server. The available options are **Server logon screen** and **Enable NTLM Authentication**. Network Level Authentication (NLA) is also called Credential Security Service Provider (CredSSP).

  By default, both the **Server logon screen** and the **Enable NTLM Authentication** options are enabled. If you want to ensure a higher level of security, enable only the **Enable NTLM Authentication** option. In this case, the connection must use TLS as the transport security. For information on enabling the TLS encryption, see Enabling TLS-encryption for RDP connections.

  NOTE: Smartcard authentication cannot be used if Network Level Authentication is negotiated at the beginning of the connection.

  - **Server logon screen**: Use this option if the RDP server is not NLA-enabled. This option allows you to log in directly to the RDP server through its graphical login screen.
  - **Enable NTLM Authentication**: Use this option to connect to NLA-enabled RDP servers with RDP 6 or later versions. NTLM is the preferred option and considered more secure than the authentication provided by the **Server logon screen** option.

- **Enabling Kerberos Authentication**: Use this option to connect to RDP servers that support Kerberos authentication. Kerberos is the preferred option and considered more secure than the authentication provided by the **Server logon screen** option or **NTLM**.

  NOTE: When enabling Kerberos authentication, note that

  - Kerberos authentication can only be used with the **Require domain membership** sub-option enabled.
  - The AA plugin is not supported with Kerberos authentication.
  - SPS has to forward the incoming service ticket to the required target server. Therefore, the referrer connection policy has to use inband target selection. For more information, see Kerberos.
  - The **Act as Remote Desktop Gateway** option cannot be selected for the referrer connection policy.

- **Require domain membership**: This is a sub-option of the **Enable NTLM Authentication** option. By default, the **Require domain membership** option is not enabled. In the default operation, you can use SPS to monitor RDP access to servers that accept only NLA, even if the client, SPS, and the server are not in the same domain. You can also use this option if the RDP server is a standalone server and is not part of a domain, or if, for some reason, you cannot add SPS to the domain.

  If you enable the **Require domain membership** option, you can only authenticate to the RDP server successfully if SPS is a member of the domain to which the RDP server belongs. To configure SPS to join your domain, see Network Level Authentication (NLA) with domain membership on page 588.

- **Enable pre-channel check**: Select this option to evaluate the connection and channel policies before establishing the server-side connection. That way if the connection is not permitted at all, SPS does not establish the server-side connection.
- **Permit unreliable usernames**: SPS automatically terminates RDP connections if it cannot reliably extract the username from the RDP connection. Enable this option to permit connections with unreliable usernames. For details on ensuring that the usernames in RDP connections are reliable, see Usernames in RDP connections on page 603.

Known issue

When accessing a Windows Server 2003 R2 host, the **Permit unreliable usernames** option is disabled, and the username is unreliable, SPS terminates the connection, but only after the user logs in. As a result, the session is not closed on the server-side.

- **Autologon domain suffix**: Enter the suffix that the client will append to the domain when using autologon in conjunction with Network Level

Authentication (CredSSP).

4. To display a banner message to the clients before authentication, enter the message into the Banner field. For example, this banner can inform the users that the connection is audited. SPS displays this banner in a graphical window that has only an OK button. Note the following points:
   - You can write a plain-text or a basic HTML-formatted banner.

     CAUTION:

     If the banner is overly complex HTML using deeply embedded structures, displaying the banner will fail, causing the RDP connections to time out.

   - When using HTML markup, the entire banner must be a single HTML object (for example, a div).

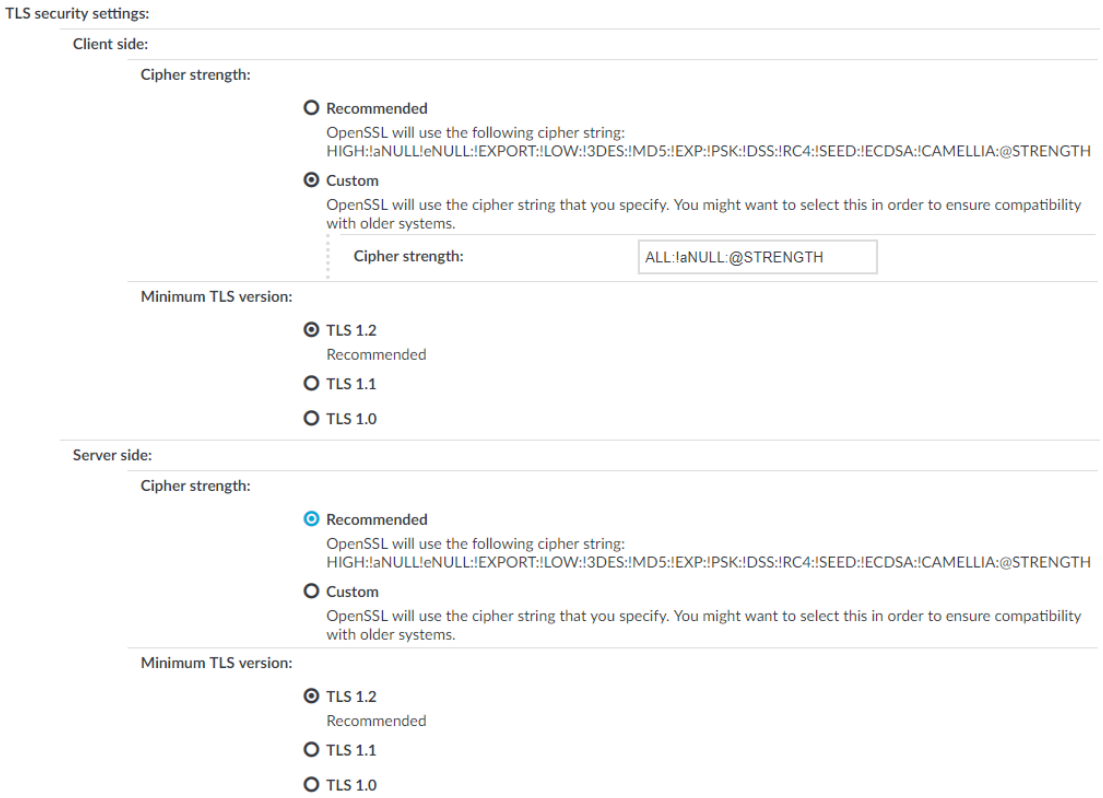
     ```
     <div align="center"><b>Your session is recorded using Privileged
     Session Monitoring</b></div>
     ```

   - In HTML, you can embed images (for example, a company logo) as data URLs in an img tag:

     ```
     To include a logo or other image, use a base64-encoded data url
     within an, like this: <img alt="Embedded Image"
     src="..." />.
     ```

   - Note that while you can include links in the text, your users cannot click or copy them.
5. To configure TLS security settings on both the **Client side** and the **Server side**, proceed to **TLS security settings**.

Figure 203: <Protocol> Control > Settings > TLS security settings - configuring TLS security settings



- **Cipher strength** specifies the cipher string OpenSSL will use. The following options are possible:
  - **Recommended**: this setting only uses ciphers with adequate security level.
  - **Custom**: this setting allows you to specify the list of ciphers you want to permit SPS to use in the connection. This setting is only recommended to ensure compatibility with older systems. For more details on customizing this list, check the 'openssl-ciphers' manual page on your SPS appliance.

    For example: `ALL:!aNULL:@STRENGTH`
- **Minimum TLS version** specifies the minimal TLS version SPS will offer during negotiation. The following options are possible:
  - **TLS 1.2**: this setting only offers TLS version 1.2 during the negotiation. This is the recommended setting.
  - **TLS 1.1**: this setting offers TLS version 1.1 and later versions during the negotiation.

- **TLS 1.0**: this setting offers TLS version 1.0 and later versions during the negotiation.

NOTE: Setting up sessions to legacy systems that do not support at least TLS 1.2 is only possible when the security level of the connection is degraded to 0, which is possible by specifying the TLS ciphers manually and appending the string `:@SECLEVEL=0` to the cipher list. However, this setting also enables the use of known vulnerable algorithms and key sizes, therefore it is absolutely critical to only use such connection settings when it is necessary and when you can fully trust your network between SPS and the legacy system. It is strongly recommended to use different security settings on the server and the client side of the connection, when degrading the security level of a connection is unavoidable.

NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported.

6. Click **Commit**.
7. Select this settings profile in the **RDP settings** field of your connections.

Network Level Authentication (NLA) with One Identity Safeguard for Privileged Sessions (SPS)

Network Level Authentication (NLA) with domain membership

You can use Credential Security Service Provider (CredSSP, also called Network Level Authentication or NLA) when One Identity Safeguard for Privileged Sessions (SPS) is member of the domain.

Prerequisites

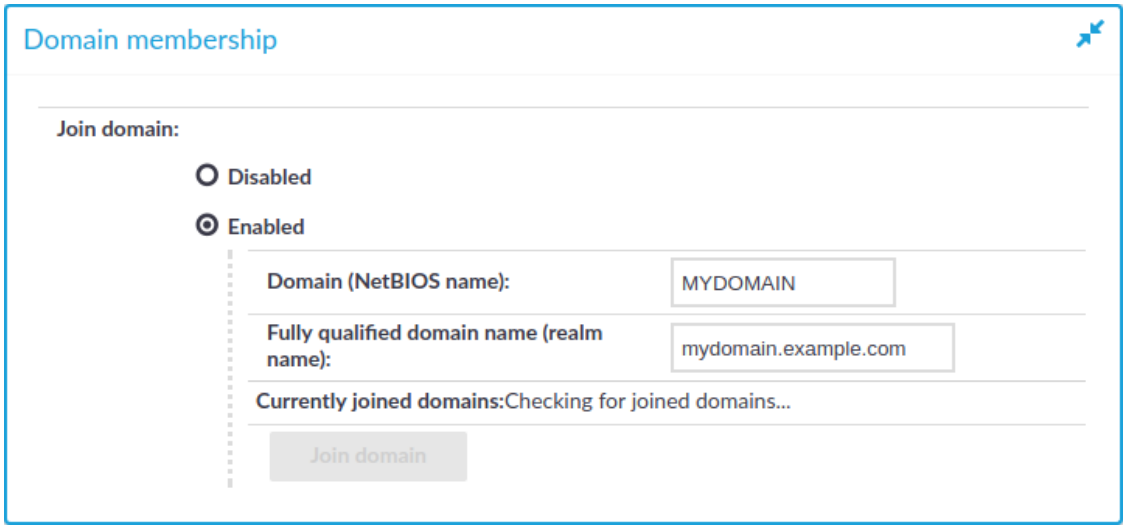
- The target servers and SPS must be in the same domain, or you must establish trust between the domains that contain the target servers and SPS. For details on the type of trust required, see *Using One Identity Safeguard for Privileged Sessions (SPS) across multiple domains* in the *Administration Guide*.

To use NLA with domain membership

1. Navigate to **Traffic Controls** > **RDP** > **Settings**, and select the RDP settings policy that you use in your connection policies.

2. Select the **Enable Network Level Authentication** > **Require domain membership** option.
3. Navigate to **Traffic Controls** > **RDP** > **Domain membership**.
4. Enter the name of the domain (for example `mydomain`) into the **Domain** field.

Figure 204: Traffic Controls > RDP > Domain membership — Joining a domain



5. Enter the name of the realm (for example `mydomain.example.com`) into the **Full domain name** field.

   NOTE: Ensure that your DNS settings are correct and that the full domain name can be resolved from SPS. To check this, navigate to **Basic Settings** > **Troubleshooting** > **Ping**, enter the full domain name into the **Hostname** field, and select **Ping host**.
6. Click **Commit**.
7. Click **Join domain**. A pop-up window is displayed.
8. SPS requires an account to your domain to be able to join the domain. Enter the following information:
   - The name of the user into the **Username** field.
   - The password into the **Password** field.

     NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:
     - Letters A-Z, a-z
     - Numbers 0-9
     - The space character
     - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

- The name of your domain controller into the **Domain controller** field. If you leave this field blank, SPS tries to find the domain controller automatically.

  NOTE: Ensure that your DNS settings are correct and that the hostname of the domain controller can be resolved from SPS. To check this, navigate to **Basic Settings** > **Troubleshooting** > **Ping**, enter the name of the domain controller into the **Hostname** field, and select **Ping host**.

- The organizational unit (OU) into the **Organization unit** field.

  The OU string reads from top to bottom without RDNs, and is delimited by a '/'. Note that '\' is used for escape by both the shell and ldap, so it may need to be doubled or quadrupled to pass through, and it is not used as a delimiter.

9. Click **Join domain**.
10. If successful, SPS displays the name of the domain it joined.

    NOTE: If you need SPS to leave the domain for some reason, click **Leave domain**.

Using One Identity Safeguard for Privileged Sessions (SPS) across multiple domains

If your users are in a domain (`EXAMPLE-DOMAIN`), One Identity Safeguard for Privileged Sessions (SPS) is also in that domain (`EXAMPLE-DOMAIN`), but your users need to access servers that are in a different domain (`OTHER-DOMAIN`), you must establish a level of trust between the domains. This is summarized in the following table.

Domain username of the client	Domain of the target server	Result
EXAMPLE-DOMAIN\my-username	EXAMPLE-DOMAIN	Connection is established
EXAMPLE-DOMAIN\my-username	OTHER-DOMAIN	If OTHER-DOMAIN trusts EXAMPLE-DOMAIN, the connection is established
OTHER-DOMAIN\my-username	OTHER-DOMAIN	If two-way trust is established between OTHER-DOMAIN and EXAMPLE-DOMAIN, the connection is established
OTHER-DOMAIN\my-username	EXAMPLE-DOMAIN	If two-way trust is established between OTHER-DOMAIN and EXAMPLE-DOMAIN, the connection is established

NOTE: If you use an LDAP database when using SPS accross multiple domains, LDAP will only use the username without the domain name to verify the group membership.

Verifying the certificate of the RDP server in encrypted connections

By default, One Identity Safeguard for Privileged Sessions (SPS) accepts any certificate shown by the server. This section describes how to accept only verified certificates.

To accept only verified certificates

1. Create a list of trusted CA certificates that you will use to verify the certificate of the server. For details, see Verifying certificates with Certificate Authorities on page 521.
2. Navigate to **Traffic Controls** > **RDP**> **Connections** and select the connection policy to modify.
3. Select **TLS**.
4. Select the required option under **Server certificate validation**.

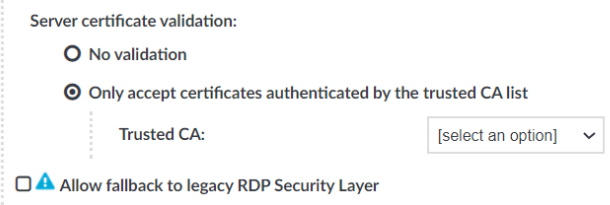
   You have the following options:

   - **No validation.**
   - **Only accept certificates authenticated by the trusted CA list.**

   NOTE:The **Only accept certificates authenticated by the trusted CA list.** option has no effect if the session uses Network Level Authentication, because in such cases SPS uses a different method to validate the server certificate.

   If you use Network Level Authentication (NLA, also called CredSSP), there is no verification performed in the TLS layer due to the TLS session-binding. For more information on TLS session-binding, see section *[MS-CSSP]: Credential Security Support Provider (CredSSP) Protocol* in the Microsoft documentation.

   Figure 205: Traffic Controls > RDP> Connections — Using SSL-encryption in RDP connections

5. (Optional) If you selected **Only accept certificates authenticated by the trusted CA list.**, in **Trusted CA list**, select the CA list to use for verifying the certificate of the server.
6. Click **Commit**.
7. (Optional) Configure your Windows servers to display a certificate signed with the above Certificate Authority for incoming RDP connections. To do this, complete the following steps:

a. Generate a certificate that contains the IP address or the hostname of the target server in its Common Name (CN) field and sign it with the Certificate Authority whose certificate you added to the **Trusted CA list** of SPS.

b. Convert the signed certificate of the target server to PKCS12 format that includes the private key.

c. Start the Microsoft Management Console (MMC) on the target server and select **Add Snap-in**> **Certificates** > **Computer Account**.

d. Right-click the **Personal** store, then select **All Tasks** > **Import**, and select the certificate created for the server.

e. Complete the Certificate Import Wizard, but do not select the **Extended certificate properties** option.

f. Select **Start** > **Administrative tools** > **Remote Desktop** > **Remote Desktop Session Host Configuration**.

g. Right-click the connection you want to configure and select **Properties** > **General**.

h. Set the **Security layer** to **SSL**.

i. Click **Certificate**> **Select** and select the imported certificate. The server uses this certificate to verify its identity for the incoming RDP connections.

Enabling TLS-encryption for RDP connections

To enable TLS-encryption in an RDP connection policy, you have two options:

- Enable Network Level Authentication (NLA, also called CredSSP). To enable NLA in RDP connections, see Network Level Authentication (NLA) with One Identity Safeguard for Privileged Sessions (SPS) on page 588.

  NOTE: Network Level Authentication uses SSL-encryption with self-signed certificates, so you do not have to configure a signing CA.

- Complete the following steps to configure TLS-encryption.

Prerequisites

Depending on your requirements, you might need one or more of the following:

- To use the same certificate for each session, an X.509 certificate and its private key are required. One Identity Safeguard for Privileged Sessions (SPS) can display this certificate to the peers on the client side. Use your own PKI system to generate these certificates, as they cannot be created on SPS.

  NOTE: The Common Name of the certificate must contain the domain name or the IP address of target machine, otherwise the clients might reject the certificate.

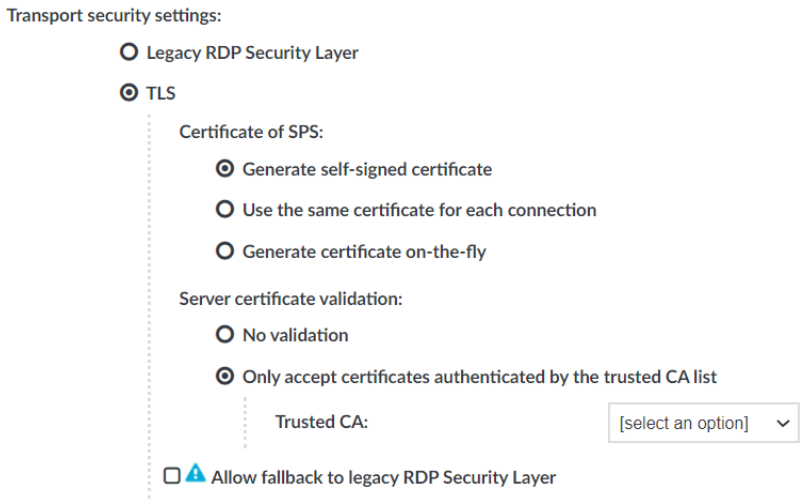
- To generate certificates on-the-fly for a connection, you need a signing certificate authority. For details on creating a signing CA, see Signing certificates on-the-fly on page 529.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To configure TLS-encryption

1. Navigate to **Traffic Controls** > **RDP**> **Connections** and select the connection policy in which you want to enable TLS.

Figure 206: Traffic Controls > RDP > Connections — Enabling TLS-encryption for RDP connections



2. Set the encryption settings used between the client/server and SPS in the **Transport security settings** section.

   To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.

3. Select the certificate to show to the peers.
   - If you want to enable TLS-encryption, but you do not have a certificate that is generated by an external CA, or a signing CA, select **Generate self-signed certificate**. By default, this option is selected.
   - To use the same certificate for every peer, complete the following steps.
     1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
     2. Select **Use the same certificate for each connection**.
     3. Select **Private key for host certificate**, click and upload the

private key.

4. Select **X.509 host certificate**, click  and upload the certificate.

- If you want to use your own Signing CA, complete the following steps.
  1. Create a certificate authority that is used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
  2. Select **Generate certificate on-the-fly**.
  3. In the **Signing CA** field, select the certificate authority to use.

  NOTE: Import the certificate of the signing Certificate Authority to your clients. Otherwise, the client applications will display a warning because of the unknown Certificate Authority.
- To disable TLS-encryption for RDP connections completely, select **Legacy RDP Security Layer** (also known as: Standard RDP Security). You might want to do this if you were using legacy RDP encryption, and you are experiencing compatibility issues. For example, you might experience a compatibility issue if you attempt to connect to a very old Windows machine (for example, Windows Server 2003 or older).

4. (Optional) Even if you choose TLS-encryption, you have the option to choose using legacy RDP encryption as well. If you experience compatibility issues (for example, if you attempt to connect to a very old Windows machine, such as Windows Server 2003 or older) and want to allow using legacy RDP encryption if TLS-encryption is not possible, select **Allow fallback to legacy RDP Security Layer** (also known as: Standard RDP Security).

   CAUTION: Security hazard!

   Selecting the Legacy RDP Security Layer or the Allow fallback to legacy RDP Security Layer options can significantly reduce the strength of the encryption used.

   Selecting these options is only recommended if you cannot overcome compatibility issues in any other way.

   To avoid security hazard, we recommend using TLS encryption.

5. Click **Commit**.

Expected result

The encryption settings are applied to the connection policy.

Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway

Remote Desktop Gateway is a Remote Desktop Services server role. This role allows authorized remote users to connect to resources located on an internal or private network from any Internet-connected device. The accessible resources can be, for example:

- Terminal servers
- Remote desktops
- Remote applications

The Remote Desktop Gateway Server Protocol is a remote procedure call (RPC) protocol using HTTPS as the transport mechanism, used primarily for tunneling client to server traffic across firewalls. One Identity Safeguard for Privileged Sessions (SPS) can act as a Remote Desktop Gateway, receiving connections using the Remote Desktop Gateway Server Protocol and transferring them to the target servers using the RDP protocol.

The Remote Desktop Gateway Server Protocol enables inband destination selection, meaning that SPS can extract the address of the target server from the client connections. This greatly simplifies managing connections on SPS without having to encode the name of the target server in the username, which was problematic as the length of the username is limited on many platforms — especially in non-transparent mode.

Prerequisites

- To access remote servers using a Remote Desktop Gateway, the clients must use version 6.1 or newer of the Remote Desktop application.
- SPS must be a member of a Windows Domain (for details on joining a domain, see Network Level Authentication (NLA) with domain membership on page 588), or you must use a **Local User Database** (for details, see Creating a Local User Database on page 538).
- Ensure that the following are synchronized:
    - System times of the Domain Controller
    - SPS
    - Clients
    - Target servers
- You cannot use the Gateway authentication on the SPS web interface for connection policies that use SPS as a Remote Desktop Gateway. However, you can configure the Remote Desktop applications of the clients to perform two separate authentications, one on the Remote Desktop Gateway (that is, on SPS), and one on the target server. For details on configuring the Remote Desktop applications of the clients to perform gateway authentications, see Configuring Remote Desktop clients for gateway authentication on page 598.

- The Remote Desktop Gateway Server Protocol supports various authentication methods. SPS acting as a Remote Desktop Gateway supports only NTLM authentication.
- You can use SPS as a Remote Desktop Gateway. You must configure the terminal service clients to use SPS as the Remote Desktop Gateway. SPS connects the server (selected inband) after authentication.
- Remote Desktop Gateway requires a certificate. Decide whether you want to use a fix certificate, or an on-the-fly generated certificate before performing the steps below and prepare the certificate.
- You may also need to adjust the port settings of the connections. The default port for RDP connections is 3389, but the Remote Desktop Gateway Server Protocol uses port 443. However, the SPS web interface uses port 443 as well, and other connection policies might already use port 443. Therefore, if administrator or user login is enabled on the interface that receives the Remote Desktop Services connections, add a new alias IP address to SPS interface. After that, use this alias in your connection policy and the client configurations. For details on creating IP aliases on SPS, see Managing logical interfaces on page 119.

CAUTION:

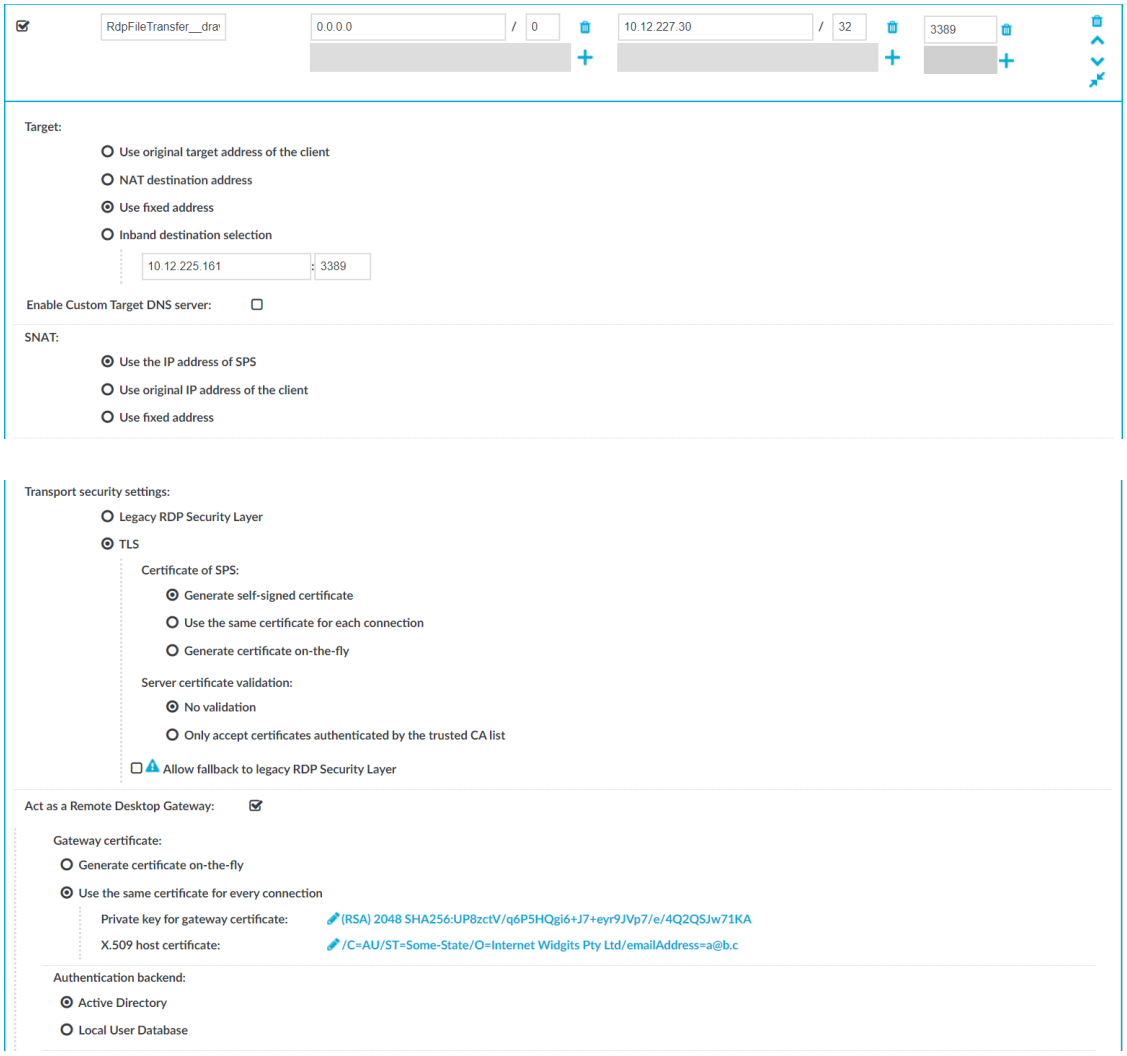
When the client uses hostname in inband destination selections, the hostname must comply with RFC5890: Internationalized Domain Names for Applications (IDNA). For example, from the ASCII characters only letters, digits, and the hyphen character is permitted.

Starting with version 6.1.0, SPS rejects connection requests where the hostname does not comply with RFC5890.

To use SPS as a Remote Desktop Gateway

1. Navigate to **Traffic Controls** > **RDP**> **Connections** and create a new connection policy that will handle the incoming client connections that use the Remote Desktop Gateway Server Protocol.
2. Enable the **Act as a Remote Desktop Gateway** option.

Figure 207: Traffic Controls > RDP > Connections — Configuring SPS as a Remote Desktop Gateway (or RD Gateway)



3. Set the target of the connections.
    - To direct every incoming connection to a single target server, select **Use fixed address** and specify the address of the target server.
    - To extract the destination address from the Remote Desktop Gateway Server Protocol, select **Inband destination selection** and set the address of the servers the clients are allowed to access in the **Target**> **Domain** fields. For details on using inband destination selection, see Modifying the destination address on page 488.

NOTE: In non-transparent mode, enter the IP address generated for the Remote Desktop Gateway service into the **To** field. Do not enter the IP address configured for administrator or user login.

4. To act as a Remote Desktop Gateway, SPS needs to display a certificate to the clients.
    - To display always the same certificate, select **Use the same certificate for every connection** and upload the X.509 certificate and the matching private key.

      TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

      CAUTION:

      The Common Name (CN) of the certificate must be the FQDN of SPS, which is the address of the Remote Desktop Gateway specified in the client applications. Otherwise the clients reject the connections.
    - To automatically create new certificates on SPS for every client, select **Generate certificate on-the-fly**, then select the Certificate Authority (CA) to sign the generated certificates with from the **Signing CA** field. For details on creating a signing CA, see Signing certificates on-the-fly on page 529.

      By default, the Common Name (CN) of the generated certificate is `<SPS-hostname.domainname>`. You can set a custom Common Name in the **Custom Common Name** field.

      NOTE: Save the CA certificate used to sign the certificate that SPS shows into DER format and import it to the clients into the **Local Computer**> **Trusted Root Certificate** store of the clients so that the clients can verify the identity of SPS.
5. In **Authentication backend**:
    - To use Active Directory for authentication, select **Active Directory**.
    - To use a Local User Database for authentication, select **Local User Database**, enter the **Domain**, and select the **Local User Database** from the list.
6. Configure other parameters of the connection policy as needed for your environment.
7. Click **Commit**.

Configuring Remote Desktop clients for gateway authentication

To configure the Remote Desktop applications of the clients to perform two separate authentications. One of these authentications is on the Remote Desktop Gateway, that is, on One Identity Safeguard for Privileged Sessions (SPS). The other authentication is on the target server. For details on configuring SPS to act as a Remote Desktop Gateway (or RD Gateway), see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.

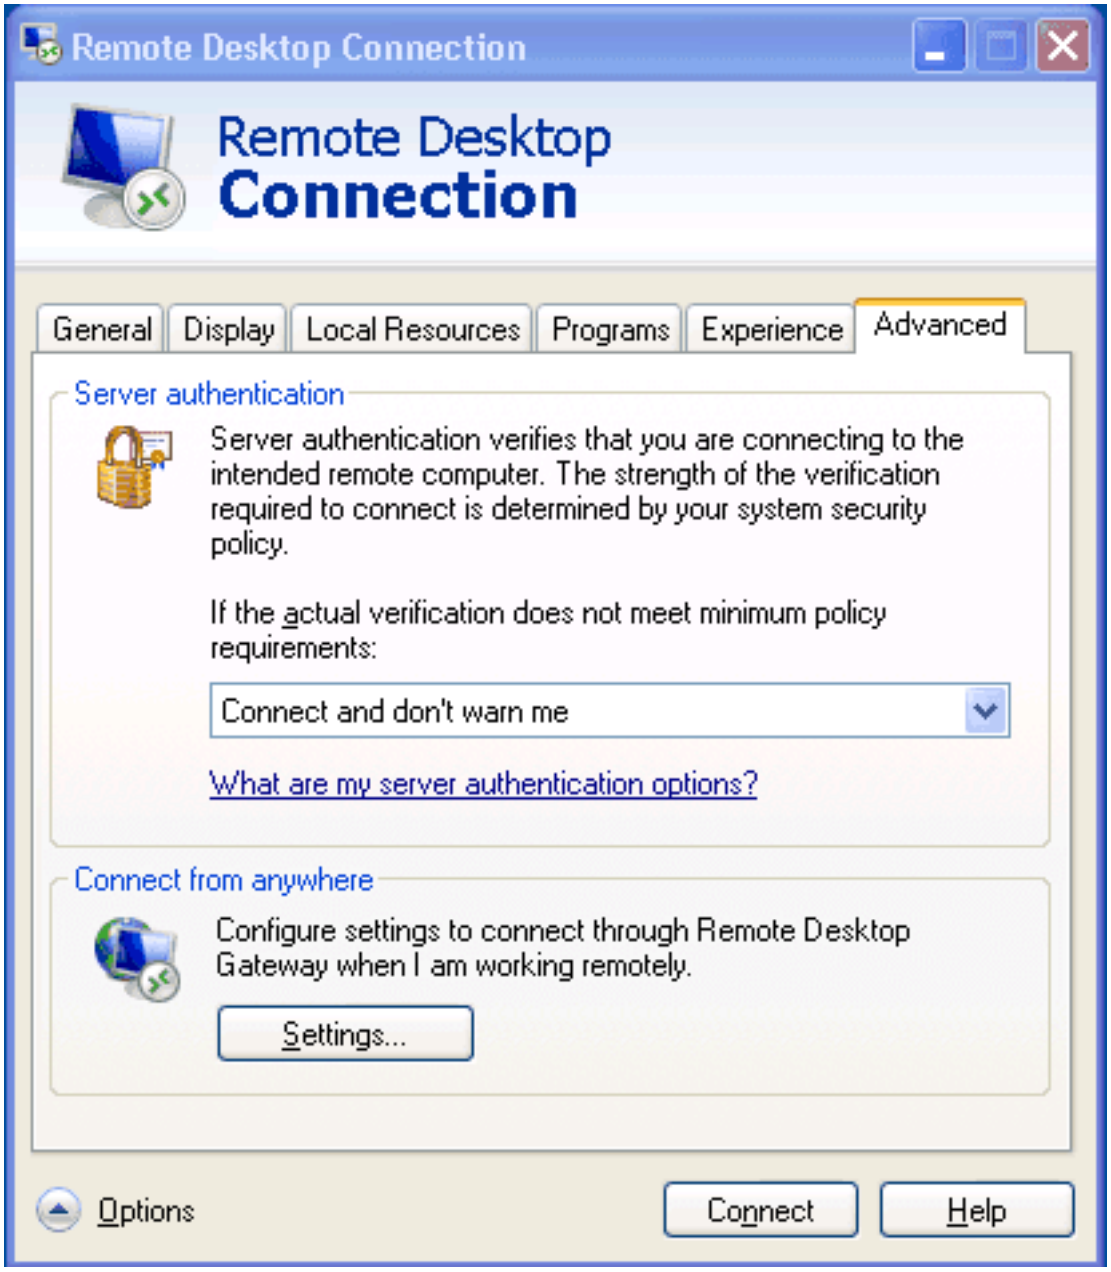
Prerequisites

- SPS must be configured to act as a Remote Desktop Gateway. For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.
- The client must use version 6.1 or newer of the Remote Desktop application.
- The target server must be member of a domain.
- The logical interface of SPS must be accessible from the client. You might have to add the address of the logical interface to the `Windows/System32/Drivers/etc/hosts` file to accomplish this.

To configure Remote Desktop clients for gateway authentication

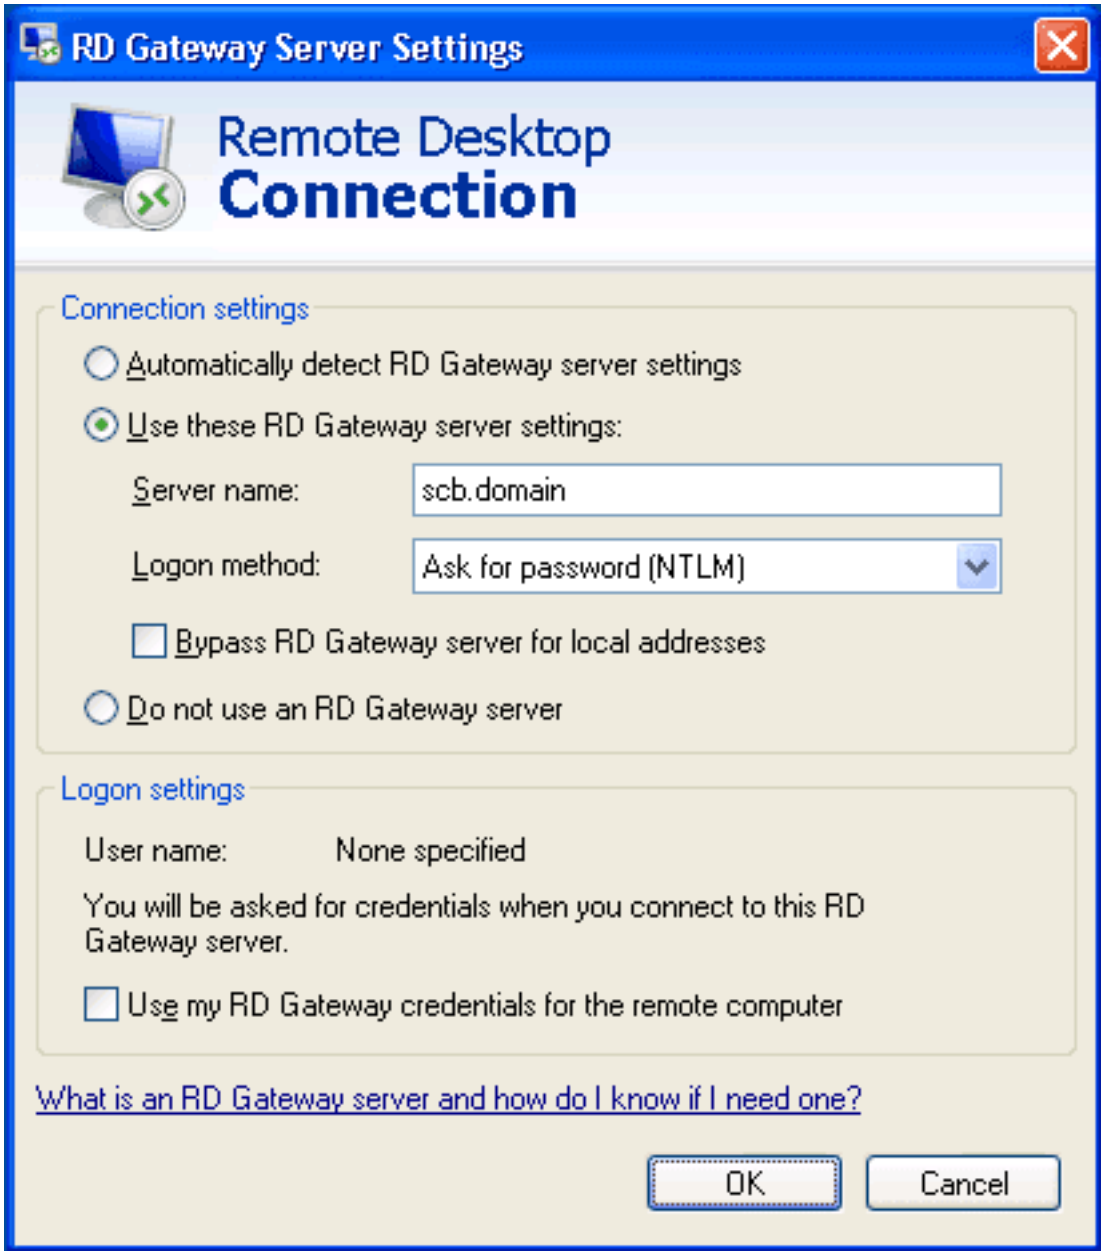
1. On your Windows client, start the **Remote Desktop Connection** application and select **Advanced** > **Settings**.

Figure 208: Configuring Remote Desktop clients to use SPS as a Remote Desktop Gateway (or RD Gateway)



2. Configure the client to use SPS as its Remote Desktop Gateway. Select **Connection settings** > **Use these RD Gateway settings**.

Figure 209: Configuring Remote Desktop clients to use SPS as a Remote Desktop Gateway (or RD Gateway)



3. Enter the address of SPS into the **Server name** field. Use the address of the SPS's logical interface that you have configured to accept RDP connections.
4. Select **Logon method** > **Ask for password (NTLM)**.
5. Uncheck the **Bypass RD Gateway server for local addresses** and **Use my RD Gateway credentials for the remote computer** options.

   NOTE:Technically, gateway authentication is performed even if the **Use my RD Gateway credentials for the remote computer** option is selected, but the

same credentials are used on the gateway and on the remote server.

6. Click **OK**.
7. Into the **Username** enter the domain username (for example, `exampledomain\exampleusername`).
8. Click **Connect**.

   NOTE: Depending on your network environment, it might take up to a minute until the connection is established.

Inband destination selection in RDP connections

To use inband destination selection with RDP connections, it is recommended to use One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway (or RD Gateway). For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.

To use inband destination selection with RDP connections without using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway (or RD Gateway), you must use SSL-encrypted RDP connections (see Enabling TLS-encryption for RDP connections on page 592).

Configure your RDP clients so SPS can record the username of client uses in the connection. If you do not configure these settings on the clients, SPS will automatically display a login screen for the users to enter their usernames and passwords.

Navigate to **Local Group Policy Editor** > **Computer Configuration** > **Administrative Templates** > **Windows Components** > **Remote Desktop Services** > **Remote Desktop Connection Client** and enable the **Prompt for credentials on the client computer** option in the clients. For details, see the Microsoft Documentation.

Also, your users have the option to encode the address of the destination server in their username, in the username field of their client application. Note that SPS automatically displays a login screen if it cannot determine the username used in the connection, or you have not encoded a destination server in the username field. You can specify the destination address in the login screen when prompted.

When encoding the address of the destination server in the username, there are a few points to keep in mind. Since most RDP client applications limit which special characters can be used in usernames, this is not always intuitive.

For the Microsoft Remote Desktop application (`mstsc`) and the login screen that SPS displays, note the following points:

- Use the `%` character to separate the fields, for example: `username%my-targetserver`
- Do not use the `@` character.

- To specify the port number of the server (if it does not use the default port), use the caret ^ character, for example: **username%my-targetserver^6464**
- To specify an IPv6 address, replace the colons with carets, and enclose the address in parentheses. For example, to target the ::1 IP address, use username%(^^1). To target port 6464 of the same server, use username%(^^1)^6464.

Usernames in RDP connections

When processing RDP connections, One Identity Safeguard for Privileged Sessions (SPS) attempts to extract the username from the connection. For example, you need the username to:

- Use gateway authentication for the connection. For details on gateway authentication, see *Configuring gateway authentication* in the *Administration Guide*.
- Use usermapping policies. In this case, SPS compares the username on the server with the username on the gateway. For more information on usermapping policies and gateway authentication, see *Configuring usermapping policies* in the *Administration Guide* and *Configuring gateway authentication* in the *Administration Guide*, respectively.

  NOTE: In certain cases, SPS receives an empty username from the server, and the connection will be denied by the usermapping policy unless a policy is set for the connection that allows every user for the given group. To add such a policy, specify * in the **Username on the server** field of the usermapping policy. For a list of cases when SPS receives empty username, see Windows settings that interfere with username extraction.
- Search or filter connections by the username on the SPS search interface, or create automatic statistics based on the username.
- Find the connection of the user on the **Pending Connections** > **Four Eyes** and **Pending Connections** > **Active Connections** pages.
- Usernames are also essential if you want to use One Identity Safeguard for Privileged Analytics. If you are interested in One Identity Safeguard for Privileged Analytics, contact our Sales Team, or your One Identity representative.

SPS can record the username automatically if the RDP connection is using Network Level Authentication (CredSSP), and usually in other scenarios as well. If SPS cannot automatically extract the username, it displays the following login screen, which allows you to paste text-based clipboard contents. Note that SPS can display this login screen only in TLS-encrypted connections.

The known scenarios that interfere with RDP usernames are listed in Windows settings that interfere with username extraction.

NOTE: SPS supports usernames both in UPN (such as username@domain) and down-level logon name (such as DOMAIN\username) formats.

To ensure that your users can access the target servers only when their username is recorded, you can configure SPS to terminate RDP connections if it cannot reliably extract

the username. To terminate such connections, clear the **Traffic Controls** > **RDP** > **Settings** > **Permit unreliable usernames** option.

Windows settings that interfere with username extraction

The following settings on the Windows client or server can prevent SPS from correctly extracting the username from the RDP connection. As a result, the username is not visible on the **Sessions** , **Pending Connections** > **Four Eyes** and **Pending Connections** > **Active Connections** pages.

- The **DontDisplayLastUserName** option is enabled on the server. The **DontDisplayLastUserName** security setting of Windows servers specifies whether the username from the last successful login is displayed on the login screen as a default for the next login. To disable the **DontDisplayLastUserName** security setting, do one of the following.
  - Disable the `HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System\DontDisplayLastUsername` registry setting. For more details, see the DontDisplayLastUserName TechNet article.

    NOTE: Registry settings can be overridden by Group Policy settings.
  - Disable this option in the `Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options` policy. For details, see Do not display last user name in logon screen TechNet article.
- There is no server-side authentication. To avoid this problem, ensure that your server requires authentication from the users.

Saving login credentials for RDP on Windows

You can use automatic RDP login on Windows, but the stored credentials are not trusted by default, and you have to enter the password for each connection. Create the following local policies on the client to allow delegating saved credentials:

1. Start the Group Policy Editor: run `gpedit.msc`.
2. Navigate to **Local Computer Policy** > **Computer Configuration** > **Administrative Templates** > **System** > **Credentials Delegation**.
3. Open the **Allow Delegating Saved Credentials with NTLM-only Server Authentication** policy.
4. Click **Show** and enter `TERMSRV/*`.
5. Click **Apply**.
6. Open the **Allow Delegating Saved Credentials** policy.

7. Click **Show** and enter `TERMSRV/*`.
8. Click **Apply**.
9. Open the **Allow Delegating Default Credentials with NTLM-only Server Authentication** policy.
10. Click **Show** and enter `TERMSRV/*`.
11. Click **Apply**.
12. Open the **Allow Delegating Default Credentials** policy.
13. Click **Show** and enter `TERMSRV/*`.
14. Click **Apply**.
15. Verify that the **Deny Delegating Saved Credentials** policy does not contain `TERMSRV/*` in the list.
16. Close the Group Policy Editor.
17. From the command line, issue the `gpupdate /force` command.

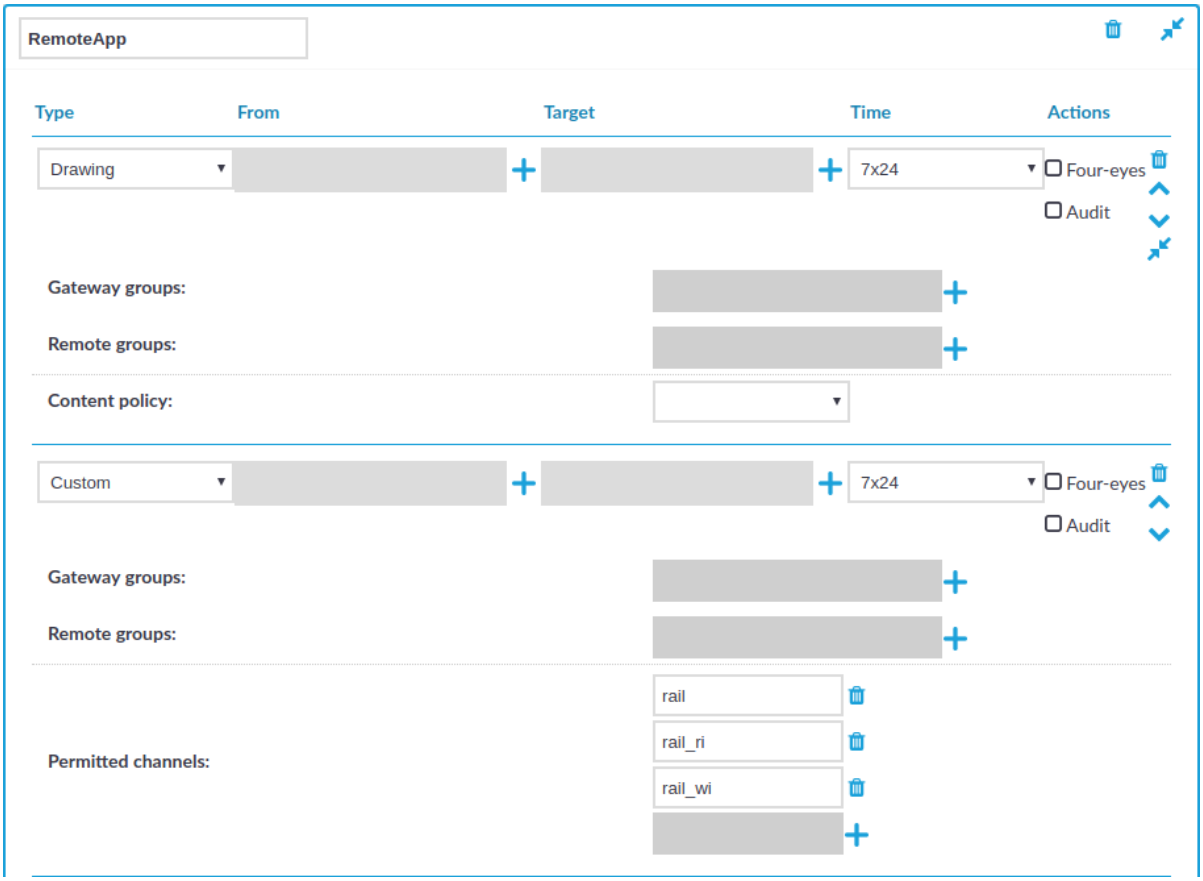
Configuring RemoteApps

Overview

RemoteApps use RDP channels that are denied by default. When configuring RDP connections for RemoteApps on One Identity Safeguard for Privileged Sessions (SPS), create a custom channel policy which enables the following channels:

- Drawing
- rail
- rail_ri
- rail_wi

Figure 210: Traffic Controls > RDP > Channel Policies — Configuring the required channels for RemoteApps



Prerequisites

- You must disable the **Use advanced RemoteFX graphics for RemoteApp** group policy on the RDP server.

  The policy is available at **Computer Configuration** > **Policies** > **Administrative Templates** > **Windows Components** > **Remote Desktop Services** > **Remote Desktop Session Host** > **Remote Session Environment** > **Use advanced RemoteFX graphics for RemoteApp**.
- You must have the Remote Desktop (RD) Licensing role installed.

To configure RemoteApps

1. Navigate to **Traffic Controls** > **RDP** > **Channel Policies**.
2. Click + to create a new channel policy.
3. Enter the name for the channel policy.
4. Choose **Drawing** as the channel type.

5. Click + to add an additional channel type.
6. Choose **Custom** as the second channel type.
7. In **Permitted channels**, click + to add the following channels:
   - rail
   - rail_ri
   - rail_wi

   (You have to click + for each channel.)
8. To save the channel policy, click **Commit**.
9. You have created a channel policy for RemoteApps.

   When you configure a connection that uses RemoteApps in **Traffic Controls** > **RDP** > **Connections**, select this channel policy as the **Channel policy** of the connection.

Configuring the RemoteApp Launcher

The RemoteApp Launcher enables users the access they need to an application without revealing credentials and passwords. By using the RemoteApp Launcher, you can protect shared credentials and limit an end user's access to an allowed or required application.

Prerequisites

- SPS version 6.11 or later.
- Windows Server 2016 or later.
- An RDP client.

Limitations

WARNING: Task managers and monitoring tools can expose the password as the part of a command line argument.

One Identity recommends hardening the remote host in use-cases when the remote app is not capable of using `stdin`. The remote app user should not be able to access task managers or monitoring tools. Alternatively, you can limit the access level of the monitoring agents.

To configure RemoteApp Launcher-specific settings on the SPS side

1. In SPS, configure a credential store plugin with the **get_remote_app_credentials** method implemented. This is required for negotiating application credentials with SPS.

Ensure that you specify the **asset** parameter. If your SPS is linked to SPP, you must also specify the **connection name** parameter.

Example: do_get_remote_app_credentials configuration

```
def do_get_remote_app_credentials(self):
    try:
        credential = self._get_credential_for_asset("password",
self.remote_app_asset, self.remote_app_account)
        return {"passwords": [credential]}
    except SafeguardException as exc:
        self.logger.error("Error checking out %s for %s@%s: '%s'",
"password", self.remote_app_account, self.remote_app_asset, exc)
        return None
```

For more information on how to configure a custom credential store plugin, see *Creating Custom Credential Store Plugins*.

2. Optionally, to enable automatic password handling by the RemoteApp Launcher, configure a custom AA plugin.

   By specifying the **asset** and **account name** parameters in an AA plugin, the RemoteApp Launcher will not prompt the user for these credentials and will start the configured application without the user having to know any credentials.

   For more information on how to configure a custom AA plugin, see *Creating Custom Authentication and Authorization Plugins*.

3. Configure RemoteApps in SPS.

   For more information, see Configuring RemoteApps.

4. Configure an RDP connection that uses RemoteApps in **Traffic Controls** > **RDP** > **Connections** and select the created RemoteApps channel policy as the **Channel policy** of the connection.

To configure RemoteApp Launcher-specific settings on the Windows side

1. Download the RemoteApp Launcher.
2. Install Remote Desktop Services (RDS).
3. Create a RemoteApp program and make the following settings:
   - For the **RemoteApp program location**, enter the path to your RemoteApp Launcher, for example, `C:\Program Files\OneIdentity\RemoteApp Launcher\OI-SG-RemoteApp-Launcher.exe`.
   - Specify to always use command-line parameters as shown in the example below:

```
"C:\Program Files\DBeaver\dbeaver.exe" --args "-con user=
{username}|password=
{password}|host=localhost|driver=PostgreSQL|database=
{asset}|name=work|connect=true" --enable-debug
```

Table 10: List of parameters

Parameter	Required?	Description
help	No	Create the help message.
cmd	Yes	Specifies the application to launch. Either the path to an executable or command alias.
args	Yes	Specifies the CLI args to pass to <cmd>. It must contain the {username}, {password} and {asset} template options, which will be expanded with the corresponding information.
use-path	No	Specifies whether to look up the cmd argument in the $PATH variable. The default setting is false. When this flag is set to false, you need to specify the full path to the executable.
`use-stdin`	No, but highly recommended if the remote app can use it.	Specifies the launcher to pass the password, username and asset by using StdIn.
`enable-debug`	No	Specifies whether to write debug logs. The default setting is false. The logs are available in the `C:\Users\<username>\AppData\Roaming\OneIdentity\OI-SG-RemoteApp-Launcher` folder. The logs are rotated daily and the file name format is `oi_sg_rci_<yyyymmdd>`.

The process and the arguments template is defined by the `--cmd` and `--args` command line arguments. In the arguments template, the RemoteApp Launcher replaces the {username} {password} and {asset} with data read from the dynamic virtual channel.

Example: RemoteApp Launcher configuration to start the DBeaver application

```
C:> C:\Program Files\OneIdentity\RemoteApp Launcher\OI-SG-
RemoteApp-Launcher.exe --cmd "C:\Program
Files\DBeaver\dbeaver.exe" --args "-con user={username}|password=
{password}|host=localhost|driver=PostgreSQL|database=
{asset}|name=work|connect=true" --enable-debug
```

4. Disable the **Use advanced RemoteFX graphics for RemoteApp** group policy on the RDP server.

   The policy is available at **Computer Configuration** > **Policies** > **Administrative Templates** > **Windows Components** > **Remote Desktop Services** > **Remote Desktop Session Host** > **Remote Session Environment** > **Use advanced RemoteFX graphics for RemoteApp**.

   You must have the Remote Desktop (RD) Licensing role installed.

Expected result

An end user can open a required application without any credentials by simply starting the RemoteApp Launcher.

Configuring SPS to enable exporting files from audit trails after RDP file transfer through clipboard or disk redirection

In SPS versions 6.2 and later, you can export files from audit trails after RDP file transfer through clipboard.

In SPS versions 6.5 and later, you can export files from audit trails after RDP file transfer through disk redirection too.

NOTE: You can export files from audit trails that were recorded before the SPS versions described above, if the Clipboard or Disk redirect channel was recorded.

To export files from these audit trails, make sure that you have the appropriate version of the Safeguard Desktop Player installed.

For more information on the process in the Safeguard Desktop Player application, see *Exporting files from an audit trail after RDP file transfer through clipboard or disk redirection* in the *Safeguard Desktop Player User Guide*.

To configure SPS to enable extracting files from audit trails after RDP file transfer through clipboard or disk redirection

1. Navigate to **Traffic Controls** > **RDP** > **Connections** and open an existing connection (or create and configure a new connection).
2. Expand the connection tab, scroll down to the **Channel policy** drop-down list, and select a channel policy of your choice from the drop-down list options.
3. Navigate to **Traffic Controls** > **RDP** > **Channel Policies** and open the channel policy that you selected from the **Channel policy** drop-down list under **Traffic Controls** > **RDP** > **Connections**.
4. Ensure that the **Clipboard** or **Redirect** drop-down list option under **Type** and the **Record audit trail** checkbox are both selected.
5. (Optional) To save your configuration, click **Commit**.

Configuring SPS to enable exporting sound from audit trails

From SPS version 6.11, you can enable auditing the sound that is transferred between an RDP client and the server.

Using the Safeguard Desktop Player application, you can export the input and output sound recorded on the side of the audited user to a **.wav** file.

NOTE: To export audio files from the audit trails, make sure that you have the appropriate version of Safeguard Desktop Player installed.

For more information on the sound export process in Safeguard Desktop Player, see *Exporting the sound from an audit trail* in the *Safeguard Desktop Player User Guide*.

To configure SPS to enable extracting sound from audit trails

1. Navigate to **Traffic Controls** > **RDP** > **Channel Policies** and open the policy that you want to use for sound auditing or create and configure a new connection for this purpose.

   For example, you can use the **all** policy for sound auditing as well.
2. From the **Type** drop-down, open **Sound** and **Dynamic virtual channel** and make sure that the **Record audit trail** checkbox is selected for both of these channel types.

   NOTE: If the **Record audit trail** checkbox is selected only for the **Sound** channel, only the output sound (the one that is received by the audited user) is recorded in the audit trail.

If the **Record audit trail** checkbox is selected for the **Dynamic virtual channel**, both the input sound (the one that comes from the audited user) and output sound (the one that is received by the audited user) are recorded.

TIP: To play back the video file of the recorded audit trails with sound, enable the **Record audit trail** checkbox for the **Drawing** channel as well.

3. To save your configuration, click **Commit**.

Sharing RDP connection policies with SPP

The **Share connection policy with SPP** option enables you to use RDP connection policies in SPP to initiate sessions.

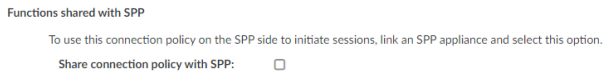
Prerequisites

Link an SPP appliance to SPS. For more information, see Linking SPS to SPP.

To share an RDP connection policy

1. Navigate to **Traffic Controls** > **RDP** > **Connections**.

Figure 211: Traffic Controls > RDP > Connections — Functions shared with SPP



2. Select **Share connection policy with SPP**.

NOTE: The **Share connection policy with SPP** checkbox is enabled only if you have linked an SPP appliance to SPS.

3. Set the following configuration for the RDP connection policy:
   - In **Target**, select **Inband destination selection**.
   - Unselect **Require Gateway Authentication on the SPS Web Interface**.
   - In **AA plugin**, consider the following, then select the AA plugin of your choice:
     - Select an **AA plugin** different from the **safeguard_default** plugin.
     - Alternatively, leave **AA plugin** unset.

WARNING: Do not delete or rename the **safeguard_default AA** plugin, otherwise, you are not able to use the **Share connection policy with SPS** option. If you have modified the **safeguard_default** plugin, to proceed, revert your changes.

- In **Credential Store**, select **SGCredStore**.

  NOTE: The **SGCredStore** option is available only if you have linked an SPP appliance to SPS.

4. To save your changes, click **Commit**.

Sharing RDP connection policies with SPS

If you have joined an SPP to SPS, the **Share connection policy with SPS** option enables you to use RDP connection policies in SPS to initiate sessions.

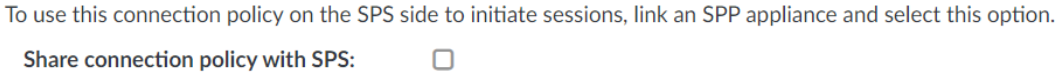
Prerequisites

Link an SPP appliance to SPS. For more information, see Linking SPS to SPP.

To share an RDP connection policy

1. Navigate to **Traffic Controls** > **RDP** > **Connections**.

   Figure 212: Traffic Controls > RDP > Connections — Functions shared with SPP

2. Select **Share connection policy with SPS**.

   NOTE: The **Share connection policy with SPS** checkbox is enabled only if you have linked an SPP appliance to SPS.

3. Set the following configuration for the RDP connection policy:
   - Unselect **Require Gateway Authentication on the SPS Web Interface**.
   - In **AA plugin**, consider the following:
     - Select an **AA plugin** different from the **safeguard_default** plugin.
     - Alternatively, leave **AA plugin** unset.

     WARNING: Do not delete or rename the **safeguard_defaultAA** plugin, otherwise, you are not able to use the **Share connection policy with SPS** option. If you have modified the **safeguard_default** plugin, to proceed, revert your changes.

4. To save your changes, click **Commit**.

Using credential injection in SPP-initiated RDP sessions

The **Credential injection** option enables you to use credential injection in SPP-initiated RDP sessions.

The RDP Application session initiated on the SPP side provides the password automatically for the `RemoteApp Launcher`. To use credential injection, use a connection policy for the RDP Application session that has `Credential injection` selected.

Prerequisites

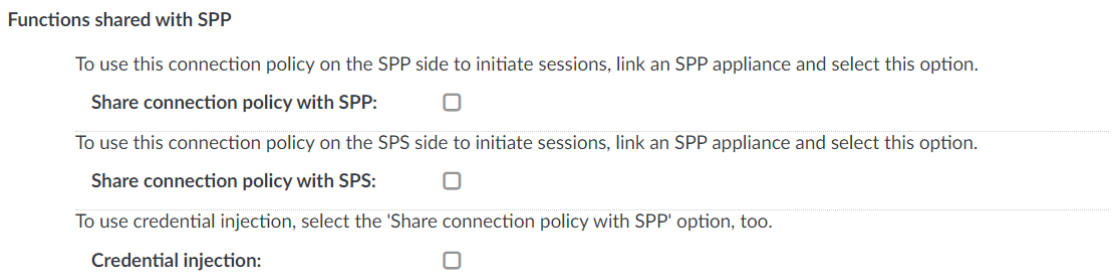
Link an SPP appliance to SPS. For more information, see Linking SPS to SPP.

To enable credential injection for SPP, select the **Share connection policy with SPP** option as well.

To use credential injection in SPP-initiated sessions

1. Navigate to **Traffic Controls** > **RDP** > **Connections**.

   Figure 213: Traffic Controls > RDP > Connections — Functions shared with SPP

2. Select **Share connection policy with SPP**.

   NOTE: The **Share connection policy with SPP** checkbox is enabled only if you have linked an SPP appliance to SPS.

   For more information on the required connection policy settings, see Sharing RDP connection policies with SPP.

3. Select **Credential injection** and set the following configuration in **RDP** > **Channel policies**:
   - In **Traffic Controls** > **RDP** > **Channel policies**, select **Drawing**.
   - In **Traffic Controls** > **RDP** > **Channel policies**, select **Dynamic virtual channel**.
   - In **Traffic Controls** > **RDP** > **Channel policies** > **Custom** > **Permitted channels**, set `rail`.

- In **Traffic Controls** > **RDP** > **Channel policies** > **Custom** > **Permitted channels**, set `rail_ri`.
- In **Traffic Controls** > **RDP** > **Channel policies** > **Custom** > **Permitted channels**, set `rail_wi`.

1. To save your changes, click **Commit**.

12

SSH-specific settings

The following sections describe configuration settings available only for the SSH protocol. Use the following policies to control who, when, and how can access the SSH connection.

- *Host keys and host certificates*: One Identity Safeguard for Privileged Sessions (SPS) allows you to set how the identity of the client hosts and servers is verified. For details, see Setting the SSH host keys of the connection on page 617.
- *Authentication Policy*: Authentication policies describe the authentication methods allowed in a connection. Different methods can be used for the client and server-side connections. For details, see Authentication Policies on page 625.
- *User List*: A user list is a list of usernames permitted to use — or forbidden from using — the connection. Essentially it is a blacklist or a whitelist. All users matching the other requirements of the connection are accepted by default. For details, see Creating and editing user lists on page 504.
- *Channel Policy*: The channel policy determines which SSH channels (for example terminal session, SCP, and so on) can be used in the connection, and whether they are audited or not. The different channels may be available only under certain restrictions, as set in the channel policy. For details, see Creating and editing channel policies on page 495.
- *SSH settings*: SSH settings determine the parameters of the connection on the protocol level, including timeout value and greeting message of the connection. The following parameters determine which algorithms are used in the connections, and can be set independently for the client and the server side: key exchange, host key, cipher, MAC, and compression algorithms. The default values include all possible algorithms. For details, see Creating and editing protocol-level SSH settings on page 638.
- *Content Policy*: Content policies allow you to inspect the content of the connections for various text patterns, and perform an action if the pattern is found. For example, SPS can send an e-mail alert if a specific command is used in an SSH terminal session. For details, see Creating a new content policy on page 499.
- *Authentication and Authorization plugin*:

  One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS to external systems to authenticate or authorize the user before authenticating on the target server. Such plugins can also be used to request additional information from the users, for example, to perform multi-factor authentication.

For details, see Integrating external authentication and authorization systems on page 884.

Setting the SSH host keys of the connection

Setting the SSH host keys accepted on the server side

You can specify how SPS verifies the identity of the servers based on their host keys.

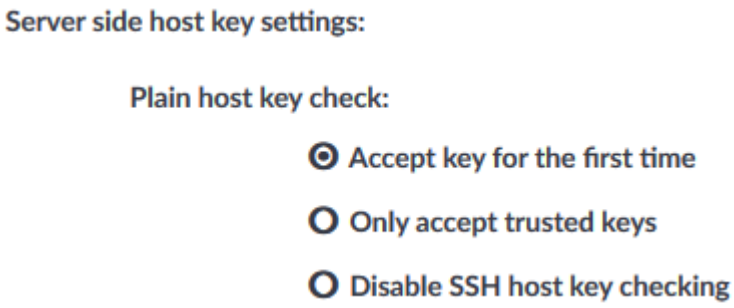
By default, SPS automatically records a host key shown by the server on the first connection. SPS will accept only this key from the server in later connections. If a host key is stored on SPS, SPS will accept only the stored key from the server.

The following describes how to set the SSH host keys accepted on the server side.

To set the SSH host keys accepted on the server side

1. Navigate to **Traffic Controls** > **SSH** > **Connections** and click to display the details of the connection.

   Figure 214: Traffic Controls > SSH > Connections — Server side host key settings

2. Verify the identity of the servers based on their host keys as follows:
   - To automatically record a host key shown by the server on the first connection, select **Accept key for the first time**.

     If no host key is stored on SPS for the target server at the time of a connection, SPS will accept and record any key shown by the server. Otherwise, if one or more host keys are already available on SPS for the target server, only the already recorded host keys are accepted. This is the default behavior of SPS.

NOTE: When your deployment consists of two or more instances of SPS organized into a cluster, the SSH keys recorded on the Managed Host nodes before they were joined to the cluster are overwritten by the keys on the Central Management node.

For more information, see Configuration synchronization and SSH keys on page 411.

- If the keys of the server are already available on SPS, select **Only accept trusted keys**. SPS will accept only the stored keys from the server.

  For more information on setting the host keys of the server, see Server host keys on page 634.

  NOTE: When your deployment consists of two or more instances of SPS organized into a cluster, the SSH keys recorded on the Managed Host nodes before they were joined to the cluster are overwritten by the keys on the Central Management node.

  For more information, see Configuration synchronization and SSH keys on page 411.

- To disable SSH host key verification, select **Disable SSH host key checking**.

  CAUTION:

  Disabling SSH host key verification makes it impossible for SPS to verify the identity of the server and prevent man-in-the-middle (MITM) attacks.

Setting the SSH host keys offered to the clients

By default, SPS automatically generates the private host keys for all supported host key algorithms when creating a new connection. You only have to make the settings described below if you would like to use your own host keys, or you would like to remove certain host keys.

The following describes how to upload or paste the private part of the SSH host key. SPS will offer the host keys to the clients.

To set the SSH host keys offered to the clients

1. Navigate to **Traffic Controls** > **SSH** > **Connections** and click to display the details of the connection.

Figure 215: Traffic Controls > SSH > Connections — Client side host key settings



2. Upload or paste the private part of the SSH host key.

   SPS allows you to use the following SSH host keys:

   - RSA (ssh-rsa), which is the most widely used public-key algorithm for the SSH key. In SPS, uploading RSA SSH host keys are supported in PKCS #1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

     NOTE:

     TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

   - Ed25519 (ssh-ed25519), which offers a better security and faster performance compared to RSA.

     In SPS, uploading Ed25519 SSH host keys are supported in PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

   - ECDSA NIST P-256 (ecdsa-sha2-nistp256), which is a variant of the Digital Signature Algorithm (DSA). In SPS, uploading ECDSA SSH host keys are supported in SEC1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

   You can have multiple SSH server host keys on SPS for the same server, however, you cannot set more than one key for each type. For example, you can keep your old RSA SSH key and generate a new Ed25519 key but you cannot set two RSA keys.

   TIP: Click on the fingerprint to display the public part of the key.

Supported SSH channel types

The available SSH channel types and their functionalities are described below. For details on configuring Channel Policies, see Creating and editing channel policies on page 495. For a list of supported client applications, see Supported protocols and client applications on page 31.

- **Agent**: Forwards the SSH authentication agent from the client to the server.

  NOTE: To perform agent-based authentication on the target server, it is not required to enable the `Agent-forwarding` channel in the Channel Policy used by the connection. The `Agent-forwarding` channel is needed only to establish connections from the target server to other devices and authenticate using the agent running

on the client.

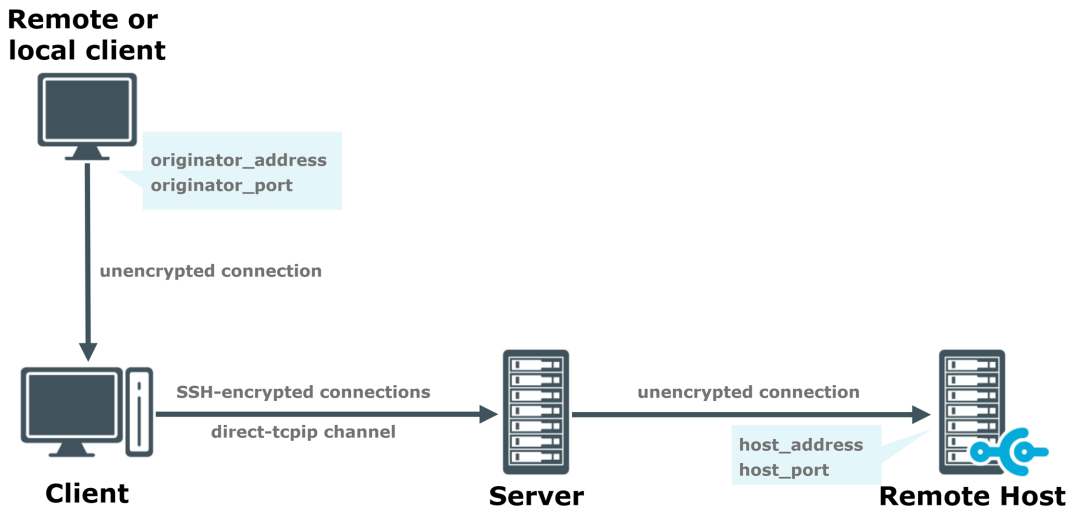
- **X11 Forward**: Forwards the graphical X-server session from the server to the client. Enter the address of the client into the **Allow client address** field to permit X11-forwarding only to the specified clients. Specify IP addresses or networks (in IP address/Prefix format).

  NOTE: Certain client applications send the Target address as a hostname, while others as an IP address. If you are using a mix of different client applications, you might have to duplicate the channel rules and create IP-address and hostname versions of the same rule.

- **Local Forward**: Forwards traffic arriving to a local port of the client to a remote host. To enable forwarding only between selected hosts, enter their IP addresses into the **Details** field. If the **Details** field is empty, local forwarding is enabled without restriction, the client may forward any traffic to the remote host. Enter the source of the forwarded traffic into the **Originator**, the target of the traffic into the **Target** field. Specify IP addresses or networks (in IP address/Prefix format). These parameters are the end-points of the forwarded traffic (that is, the local host that sends data to the remote host), and not the SSH server or the client.

  For example, to enable forwarding from the `192.168.20.20` host to the remote host `192.168.50.50`, enter **`192.168.20.20`** into the **Originator**, and **`192.168.50.50`** into the **Target** field.

Figure 216: Local TCP forwarding



NOTE: Certain client applications send the Originator and Target addresses as hostnames, while others as IP addresses. If you are using a mix of different client applications, you might have to duplicate the channel rules and create IP-address and hostname versions of the same rule.

CAUTION:

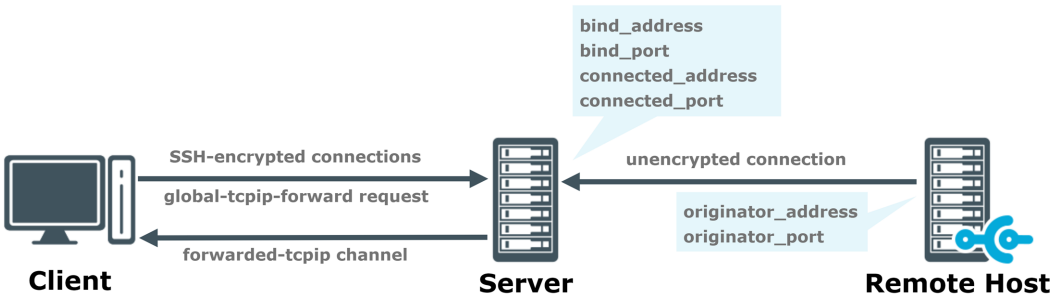
Port forwarding across One Identity Safeguard for Privileged Sessions (SPS) may fail for certain SSH client-server combinations. This happens if within the protocol, the address of the remote host is specified as a hostname during the port-forwarding request (`SSH_MSG_GLOBAL_REQUEST`), but the hostname is resolved to IP address in the channel opening request (`SSH_MSG_CHANNEL_OPEN`. By default, SPS rejects such connections.

To enable these connections, navigate to **Traffic Controls** > **SSH** > **Settings**, and disable the **Strict mode** option.

- **Remote Forward**: Forwards traffic arriving a remote port of the server to the client. To enable forwarding only between selected hosts, enter their IP addresses into the **Details** field. If the **Details** field is empty, remote forwarding is enabled without restriction, the SSH server may forward any traffic to the client. Enter the source of the forwarded traffic into the **Originator**, the target of the traffic into the **Target** field. Specify IP addresses or networks (in IP address/Prefix format). These parameters are the end-points of the forwarded traffic (that is, the remote host that sends data to the client), and not the SSH server.

  For example, to enable forwarding from the `192.168.20.20` remote host to the client `192.168.50.50`, enter **`192.168.20.20`** into the **Originator**, and **`192.168.50.50`** into the **Target** field.

Figure 217: Remote TCP forwarding



NOTE: Certain client applications send the Originator and Target addresses as hostnames, while others as IP addresses. If you are using a mix of different client applications, you might have to duplicate the channel rules and create IP-address and hostname versions of the same rule.

CAUTION:

Port forwarding across SPS may fail for certain SSH client-server combinations. This happens if within the protocol, the address of the remote host is specified as a hostname during the port-forwarding request (`SSH_MSG_GLOBAL_REQUEST`), but the hostname is resolved to IP address in the channel opening request (`SSH_MSG_CHANNEL_OPEN`. By default, SPS rejects such connections.

To enable these connections, navigate to **Traffic Controls** > **SSH** > **Settings**, and disable the **Strict mode** option.

- **Session Exec**: Execute a remote command (for example `rsync`) without opening a session shell. Enter the permitted command into the **Permitted commands** field. You can use regular expressions to specify the commands. This field can contain only letters (`a-z, A-Z`), numbers (`0-9`), and the following special characters (`{}()*?\\|[]`).

  CAUTION:

  Restricting the commands available in **Session Exec** channels does not guarantee that no other commands can be executed. Commands can be renamed, or executed from shell scripts to circumvent such restrictions.

- **Session Exec SCP**: Transfers files using the Secure Copy (SCP) protocol.
  - To make the list of file operations available in the **File operations** column of the **Sessions** page, navigate to the Channel Policies page of the protocol, and enable the **Log file transfers to database** option. This option is disabled by default.
  - To send the file operations into the system log, enable the **Log file transfers to syslog** option. This option is disabled by default.

    NOTE: Turning logging on might result in a slight performance penalty. If traffic load slows processes down, disable the option.

  CAUTION:

  The WinSCP application does not follow the RFC of the SCP protocol properly, but transfers files in a **Session Shell** channel instead of a **Session Exec SCP** channel. This has the following results:
  - If the **Session Shell** channel is enabled in a Channel Policy (this is needed for SSH terminal sessions as well), and your users use WinSCP using the **File protocol** > **SCP** option, they will be able to transfer files to and from the server. Also, these files will not be listed in the **File operations** field of the **Sessions** page.
  - To avoid these problems, you have to enforce that your clients use WinSCP with the **File protocol** > **SFTP** option. WinSCP uses SFTP by default, but can be changed manually to use SCP, and also falls back to using SCP if a server rejects SFTP.
  - To terminate the connection when a user transfers a file with WinSCP using the Session Shell channel, create a Content Policy that matches the `WinSCP: this is end-of-file` string in screen content, and use this policy in your Connection Policies. For details on Content Policies, see Real-time content monitoring with Content Policies on page 498. This solution has been tested with WinSCP version 5.1.5: if it does not work for your version, contact our Support Team.

- **Session Subsystem**: Use a subsystem. Enter the name of the permitted subsystem into the **Permitted subsystem** field.
- **Session SFTP**: Transfers files using the Secure File Transfer Protocol (SFTP).

- To make the list of file operations available in the **File operations** column of the **Sessions** page, navigate to the Channel Policies page of the protocol, and enable the **Log file transfers to database** option. This option is disabled by default.
- To send the file operations into the system log, enable the **Log file transfers to syslog** option. This option is disabled by default.

  NOTE: Turning logging on might result in a slight performance penalty. If traffic load slows processes down, disable the option.

- **Session Shell**: The traditional remote terminal session.

  CAUTION:

  The WinSCP application does not follow the RFC of the SCP protocol properly, but transfers files in a **Session Shell** channel instead of a **Session Exec SCP** channel. This has the following results:

  - If the **Session Shell** channel is enabled in a Channel Policy (this is needed for SSH terminal sessions as well), and your users use WinSCP using the **File protocol** > **SCP** option, they will be able to transfer files to and from the server. Also, these files will not be listed in the **File operations** field of the **Sessions** page.
  - To avoid these problems, you have to enforce that your clients use WinSCP with the **File protocol** > **SFTP** option. WinSCP uses SFTP by default, but can be changed manually to use SCP, and also falls back to using SCP if a server rejects SFTP.
  - To terminate the connection when a user transfers a file with WinSCP using the Session Shell channel, create a Content Policy that matches the `WinSCP: this is end-of-file` string in screen content, and use this policy in your Connection Policies. For details on Content Policies, see Real-time content monitoring with Content Policies on page 498. This solution has been tested with WinSCP version 5.1.5: if it does not work for your version, contact our Support Team.

Sharing SSH connection policies with SPP

The **Share connection policy with SPP** option enables you to use SSH connection policies in SPP to initiate sessions.

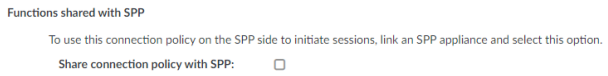
Prerequisites

Link an SPP appliance to SPS. For more information, see Linking SPS to SPP.

To share an SSH connection policy

1. Navigate to **Traffic Controls** > **SSH** > **Connections**.

   Figure 218: Traffic Controls > SSH > Connections — Functions shared with SPP

2. Select **Share connection policy with SPP**.

   NOTE: The **Share connection policy with SPP** checkbox is enabled only if you have linked an SPP appliance to SPS.

3. Set the following configuration for the SSH connection policy:
   - In **Target**, select **Inband destination selection**.
   - Unselect **Require Gateway Authentication on the SPS Web Interface**.
   - In **AA plugin**, consider the following, then select the AA plugin of your choice:
     - Select an **AA plugin** different from the **safeguard_default** plugin.
     - Alternatively, leave **AA plugin** unset.

     WARNING: Do not delete or rename the **safeguard_default AA** plugin, otherwise, you are not able to use the **Share connection policy with SPS** option. If you have modified the **safeguard_default** plugin, to proceed, revert your changes.

   - In **Credential Store**, select **SGCredStore**.

     NOTE: The **SGCredStore** option is available only if you have linked an SPP appliance to SPS.

   - In **Authentication policy**, select an authentication policy.

     After that, navigate to **Traffic Controls** > **SSH** > **Authentication Policies** and open the respective authentication policy. In **Relayed authentication methods**, select **Keyboard-interactive** and **Password**.

4. To save your changes, click **Commit**.

Sharing SSH connection policies with SPS

If you have joined an SPP to SPS, the **Share connection policy with SPS** option enables you to use SSH connection policies in SPS to initiate sessions.

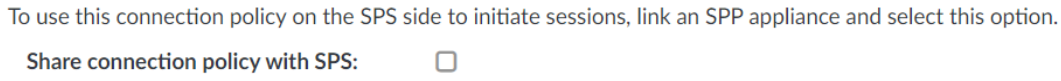
Prerequisites

Link an SPP appliance to SPS. For more information, see Linking SPS to SPP.

To share an SSH connection policy

1. Navigate to **Traffic Controls** > **SSH** > **Connections**.

   Figure 219: Traffic Controls > SSH > Connections — Functions shared with SPS

2. Select **Share connection policy with SPS**.

   NOTE: The **Share connection policy with SPS** checkbox is enabled only if you have linked an SPP appliance to SPS.

3. Set the following configuration for the SSH connection policy:
   - Unselect **Require Gateway Authentication on the SPS Web Interface**.
   - In **AA plugin**, consider the following:
     - Select an **AA plugin** different from the **safeguard_default** plugin.
     - Alternatively, leave **AA plugin** unset.

     WARNING: Do not delete or rename the **safeguard_default AA** plugin, otherwise, you are not able to use the **Share connection policy with SPS** option. If you have modified the **safeguard_default** plugin, to proceed, revert your changes.
   - In **Authentication policy**, select an authentication policy.

     After that, navigate to **Traffic Controls** > **SSH** > **Authentication Policies** and open the respective authentication policy. In **Gateway authentication method**, select at least one of the following options:
     - **Password**
     - **Public key**
     - **Kerberos**

4. To save your changes, click **Commit**.

Authentication Policies

An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the client can authenticate to the target server. Separate authentication methods can be used on the client and the server-side of the connection.

Figure 220: Authentication policies



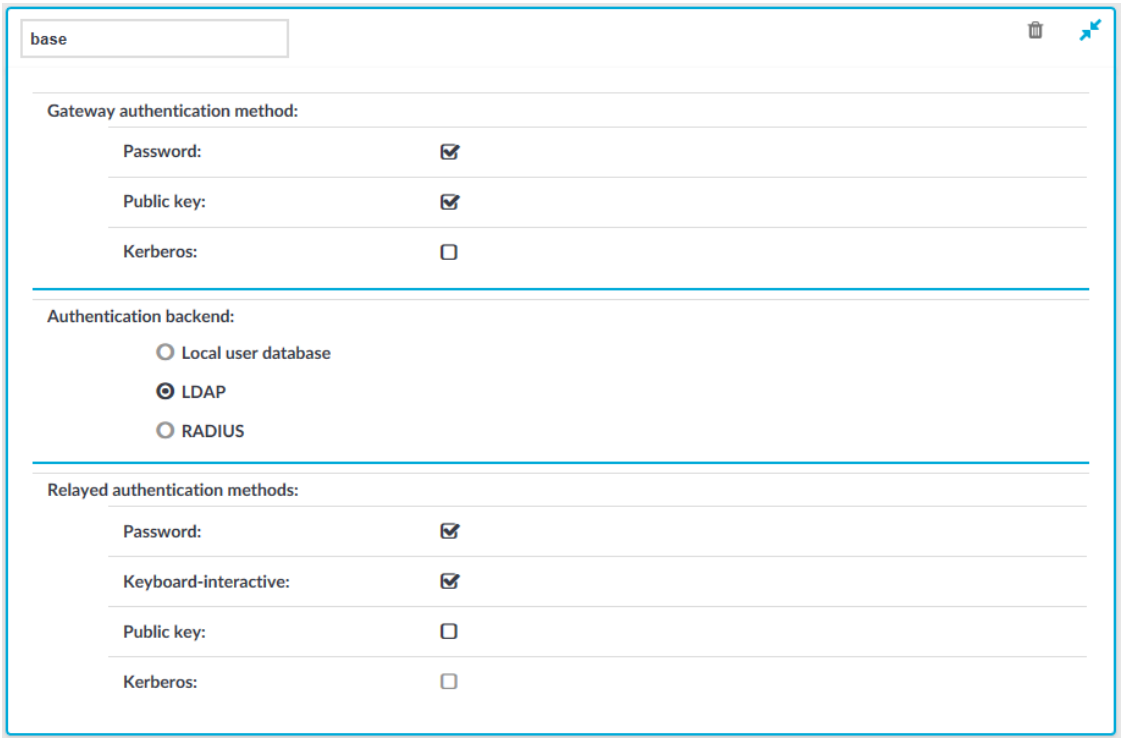
Creating a new authentication policy

The following describes how to create a new authentication policy.

To create a new authentication policy

1. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies**, and click .

   Figure 221: Traffic Controls > SSH > Authentication Policies — Configuring authentication policies

2. Enter a name for the policy.

3. Select the gateway authentication method to allow authentication on SPS. Note that this is an inband authentication that happens within the SSH protocol.

   If you selected Kerberos as the gateway authentication method, SPS disables all other authentication backends and authentication methods.

4. Select the authentication database used on the client-side in the **Authentication backend** field. For details on the client-side authentication settings, see Client-side authentication settings on page 627.

   If you selected Kerberos as the gateway authentication method, skip this step. For details, see Kerberos authentication settings.

5. Select the authentication method used on the server-side in the **Relayed authentication methods** field. For details on the relayed authentication settings, see Relayed authentication methods on page 630. Note the following:

   - If you selected Kerberos as the gateway authentication method, skip this step. For details, see Kerberos authentication settings.
   - If you selected **Public key** > **Agent** as the relayed authentication method:

     If this option is used, SPS requests the client to use its SSH agent to authenticate on the target server. Therefore, you must configure your clients to enable agent forwarding, otherwise authentication will fail. For details on enabling agent forwarding in your SSH application, see the documentation of the application.

   To be able to select Kerberos as the relayed authentication method, ensure that you also select Kerberos as the gateway authentication method.

6. Click **Commit**.

   NOTE: Consider the following:

   - The client-side authentication settings apply for authenticating the user inband (that is, within the SSH protocol) to the One Identity Safeguard for Privileged Sessions (SPS) gateway, and is independent from the gateway authentication performed on the SPS web interface. The web-based gateway authentication is an out-of-band gateway authentication method that can be required by the connection policy. For details on out-of-band gateway authentication, see Configuring out-of-band gateway authentication on page 863.

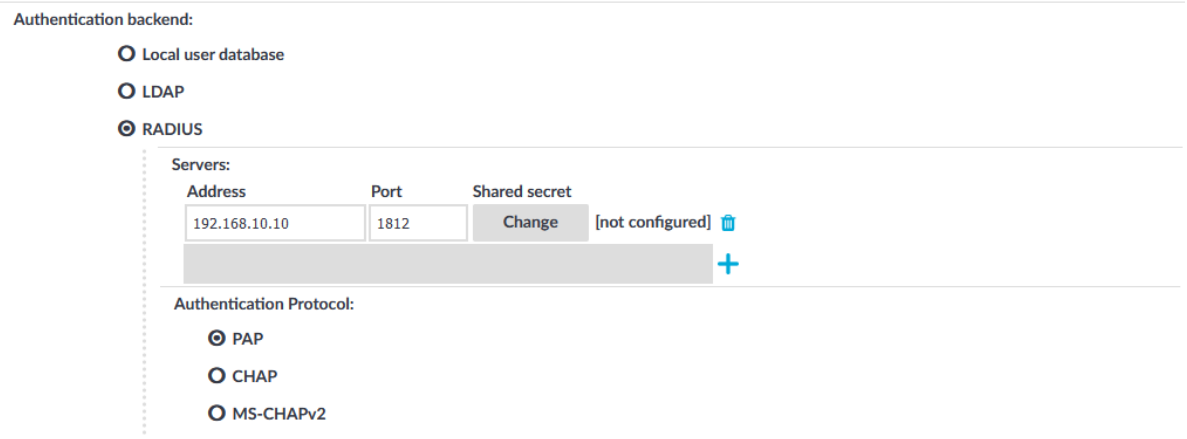
     Gateway authentication on the SPS web interface can be used together with authentication policies. In an extreme setting, this would mean that the user has to perform three authentications: a client-side gateway authentication within the SSH protocol to SPS, an out-of-band gateway authentication on the SPS web interface, and a final authentication on the target server.

   - The Connection Policy will ignore the settings for server-side authentication (set under **Relayed authentication methods** for SSH protocol) if a Credential Store is used in the Connection Policy.

Client-side authentication settings

For the client-side connection, One Identity Safeguard for Privileged Sessions (SPS) can authenticate the client inband (within the SSH protocol) using the following authentication methods:

Figure 222: Traffic Controls > SSH > Authentication Policies — Configuring client-side authentication methods



- **Local user database**: Authenticate the client locally on the SPS gateway. For details, see Local client-side authentication on page 629.
- **LDAP**: SPS will authenticate the client to the LDAP database set in the **LDAP Server** of the connection policy. To use LDAP authentication on the client side, select **LDAP**, and select the permitted authentication methods (**Password**, **Public key**). More than one method can be permitted.

  NOTE:

  - SPS will authenticate the client-side connection to the LDAP server configured in the connection policy. This is not necessarily the same as the LDAP server used to authenticate the users accessing the SPS web interface.
  - The public keys of the users stored in the LDAP database must be in OpenSSH format.
- **RADIUS**: SPS will authenticate the client to the specified RADIUS server. Select **RADIUS**, enter the IP address or hostname of the RADIUS server into the **Address** field, the port number of the RADIUS server into the **Port** field, and the shared secret of the RADIUS server into the **Shared secret** field. Only password-authentication is supported (including one-time passwords), challenge-response based authentication is not.

  To use the:

  - Password Authentication Protocol, select **PAP**.
  - Challenge-Handshake Authentication Protocol, select **CHAP**.
  - Microsoft version of the Challenge-Handshake Authentication Protocol, select **MS-CHAPv2**.

Use an IPv4 address.

To add more RADIUS servers, click + and fill in the respective fields.

To use certificates to authenticate the client, you can use the **LDAP** and the **Local user database** backends.

Figure 223: Client-side inband gateway authentication with different certificates

	Trusted CA list is set in the Authentication Policy	
	YES	NO
The certificate shown by the client is self-signed AND the user is in the Local User Database and has a self-signed certificate set in the database	successful	successful
The certificate shown by the client is CA-signed	successful	unsuccessful

Local client-side authentication

The following describes how to perform authentication locally on One Identity Safeguard for Privileged Sessions (SPS) for client-side connections.

NOTE: The users can be authenticated to their passwords or public-keys uploaded to SPS.

The accounts created to access the SPS web interface cannot be used to authenticate SSH connections.

Prerequisites

To perform authentication locally on SPS for client-side connections, an existing **Local User Database** is needed. To create a **Local User Database**, complete the following procedure: Creating a Local User Database on page 538.

To perform authentication locally on SPS for client-side connections

1. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies**, and select the authentication policy to modify.
2. Select the permitted authentication methods (**Password**, **Public key**).
3. Select **Local user database**.
4. Select the **Local user database** from the list that defines the users who can access the server.
5. Click **Commit**.

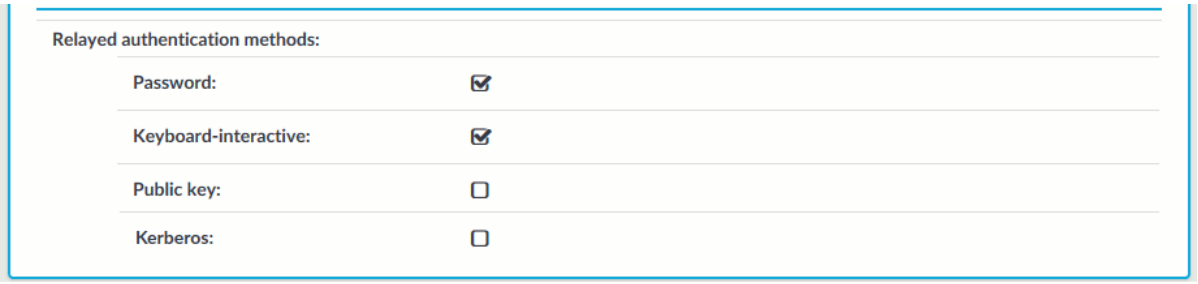
Relayed authentication methods

For the server-side connection (between One Identity Safeguard for Privileged Sessions (SPS) and the target server), the following authentication methods are available.

NOTE: Even though these settings refer to the server-side connection, the client must support the selected authentication method and have it enabled. For example, to use publickey authentication on the server side, the client must support publickey authentication as well as provide a fake publickey, even if a different authentication method is used on the client side.

The Connection Policy will ignore the settings for server-side authentication (set under **Relayed authentication methods** for SSH protocol) if a Credential Store is used in the Connection Policy.

Figure 224: Traffic Controls > SSH > Authentication Policies — Configuring relayed authentication methods



- **Password**: Authentication based on username and password. The server will request a password from the user, even if a password-based authentication was already successful on the client-side.

- **Keyboard-Interactive**: Authentication based on exchanging messages between the user and the server. This method includes authentication schemes like S/Key or TIS authentication. Note that depending on the configuration of the SSH server, password-based authentication can also require using the keyboard-interactive authentication method.
- **Public Key**: Authentication based on public-private encryption keypairs. SPS supports the following public-key authentication scenarios:
    - **Publish to LDAP**: SPS generates a keypair, and uses this keypair in the server-side connection. The public key of this keypair is also uploaded to the LDAP database set in the LDAP Server of the connection policy. That way the server can authenticate the client to the generated public key stored under the user's username in the LDAP database.
    - **Fix**: Uses the specified private key in the server-side connection.
    - **Agent**: Allow the client to use agent-forwarding, and use its own keypair on the server-side.

      During agent-forwarding, the following keys are accepted:

        - rsa
        - ed25519
        - ecdsa-sha2-nistp256
        - ecdsa-sha2-nistp384
        - ecdsa-sha2-nistp521

      If this option is used, SPS requests the client to use its SSH agent to authenticate on the target server. Therefore, you must configure your clients to enable agent forwarding, otherwise authentication will fail. For details on enabling agent forwarding in your SSH application, see the documentation of the application.

      TIP: Some clients may override agent forwarding requests for SFTP and SCP by default. For further information about ensuring access to the server in this case, see Using SCP with agent-forwarding.

  TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
- **Kerberos**: Authentication based on Kerberos. Only available if you selected Kerberos as the gateway authentication method. For more information, see Kerberos authentication settings.

Configuring your Kerberos environment

To integrate One Identity Safeguard for Privileged Sessions (SPS) with your Kerberos environment, so that your clients can authenticate on the target servers using Kerberos tickets, you have to configure your environment appropriately.

To configure your Kerberos environment

1. Configure your DNS server.
    a. On your Domain Name Server (DNS), add SRV records that describe which Key Distribution Center (KDC) belongs to the domain. Add both TCP and UDP entries for each domain. For example, if your domain is `example.com` and the hostname of your KDC server is `kdc.example.com`, this entry looks like:

    ```
    _kerberos._tcp.example.com 0 0 88 kdc.example.com
    _kerberos._udp.example.com 0 0 88 kdc.example.com
    ```

    b. If your environment uses multiple realms, repeat the previous step for every realm.
    c. Verify that the servers that your clients will connect to via SPS have proper reverse-dns entries. Otherwise, your clients cannot access the target servers if you use the **Inband destination selection** feature of SPS.
2. Create a keytab file for SPS.
    a. On your KDC server, create a principal for the SPS host, using the domain name of your SPS. For example:

    ```
    host/scb.example.com@EXAMPLE.COM
    ```

    b. If your environment uses multiple realms, repeat the previous step on the KDC of every realm.
    c. Export the key of the principal into a keytab file.
    d. If your environment uses multiple realms, merge the keytab files of the different realms into a single file, for example, using the `ktadd` or the `ktutil` utilities.
    e. If your environment uses multiple realms, repeat the previous step on the KDC of every realm.
3. Configure the SSH application of your client hosts to enable Kerberos (GSSAPI) ticket forwarding. (In most applications this is disabled by default.)

Expected result

You have configured your environment to use Kerberos authentication with SPS, and created a keytab file for your SPS host. For details on uploading the keytab file and configuring SPS see Kerberos authentication settings on page 632.

Kerberos authentication settings

The following describes how to perform authentication with Kerberos. One Identity Safeguard for Privileged Sessions (SPS) supports both end-to-end Kerberos authentication, when the client authenticates on SPS gateway and on the target server using Kerberos, and also the half-sided Kerberos scenario when Kerberos is used only on the SPS gateway.

Prerequisites

Before configuring Kerberos authentication on One Identity Safeguard for Privileged Sessions (SPS), make sure you have configured your Kerberos environment correctly and have retrieved the keytab file. For details, see Configuring your Kerberos environment on page 631.

To perform authentication with Kerberos

1. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies**.
2. Create a new Authentication Policy.
3. Select the authentication methods to use on the SPS gateway and on the target server.
   - To use Kerberos authentication on the target server, you must use Kerberos authentication both on the SPS gateway and on the target server. Select **Gateway authentication method** > **Kerberos** and **Relayed authentication methods** > **Kerberos**.
   - To use Kerberos authentication only on the SPS gateway (that is, in the client-side connection), select **Gateway authentication method** > **Kerberos**. If required, you can select other gateway authentication methods in addition to Kerberos, and also authentication backends and related to the selected gateway authentication methods.

     Select the authentication methods you want to use on the target server in the **Relayed authentication methods** field.
4. Click **Commit**.
5. Navigate to **Traffic Controls** > **SSH** > **Global Options** > G**SSAPI**.
6. **Browse** for the **Kerberos keytab file**, and click **Upload**. The uploaded principals are displayed in **Currently uploaded principals**.

   If a Connection Policy uses an SSH Authentication Policy with **Kerberos** authentication together with a Usermapping Policy, then SPS stores the user principal as the gateway user, and the target username as the server username in the session database. If you want to allow your users to use a username on the target server that is different from their principal, configure a Usermapping Policy for your SSH connections. For details, see Configuring usermapping policies.

   Click **Commit**.
7. (Optional) If more than one realm is deployed on your network, you have to specify the mapping from the server's DNS domain name to the name of its realm. To map hostnames onto Kerberos realms, click  .
8. Navigate to **Traffic Controls** > **SSH** > **Connections** and configure the SSH connection as follows. For details on configuring connections in general, see Configuring connections on page 482.
   a. Select **Use fixed address** or **Inband destination selection** as **Target**.

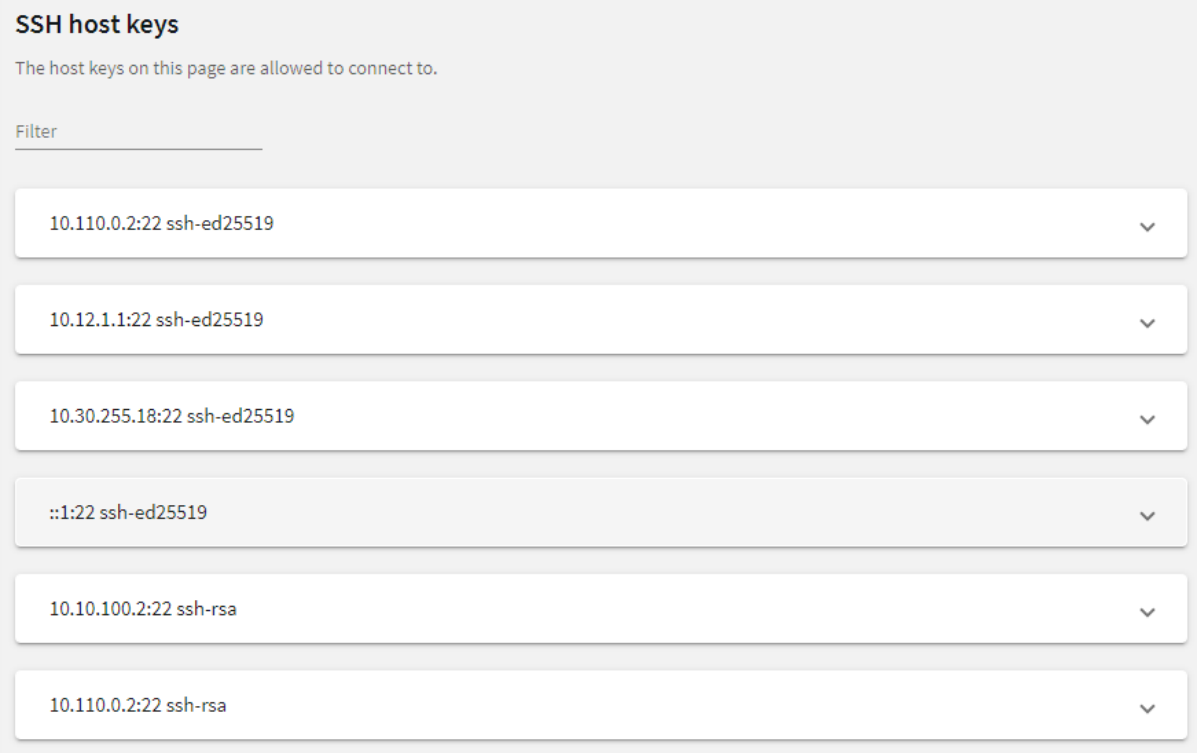
b. Select the Kerberos **Authentication policy**.

c. Click **Commit**.

Server host keys

The host keys of the trusted servers can be managed on the **Traffic Controls** > **SSH** > **Server Host Keys** page. When a client tries to connect to a server, One Identity Safeguard for Privileged Sessions (SPS) verifies the host key of the server. SPS allows connections only to the servers listed on this page, unless the **Accept key for the first time** or the **Accept certificate for the first time** option is enabled in the connection policy.

Figure 225: Traffic Controls > SSH > Server Host Keys — Server host keys



To filter, start typing into the **Filter** field. For example, enter **rsa** or an IP address and only relevant results are displayed.

Figure 226: Traffic Controls > SSH > Server Host Keys — Filtered RSA host keys



Automatically adding the host keys of a server to One Identity Safeguard for Privileged Sessions (SPS)

The host keys of the servers can be added either automatically or manually.

To add the host key automatically

1. Navigate to the **Traffic Controls** > **SSH** > **Connections**.
2. Configure a connection: fill the **From**, **To**, and **Port** fields.

   You can use IPv4 and IPv6 addresses as well.

   - To configure a transparent connection, enter the IP address of the server into the **To** field.
   - To configure a non-transparent connection, enter the IP address of SPS into the **To** field, and the address of the target server into the **Target** field.
3. Click to display the advanced settings and verify that the **Server side host key settings** > **Plain host key check** option is set to **Accept key for the first time**.

   Click **Commit**.
4. Initiate an SSH connection from the client to the server. SPS will automatically record the host key of the server — the server's IP address and the host key will be listed on the **Traffic Controls** > **SSH** > **Server Host Keys** page.

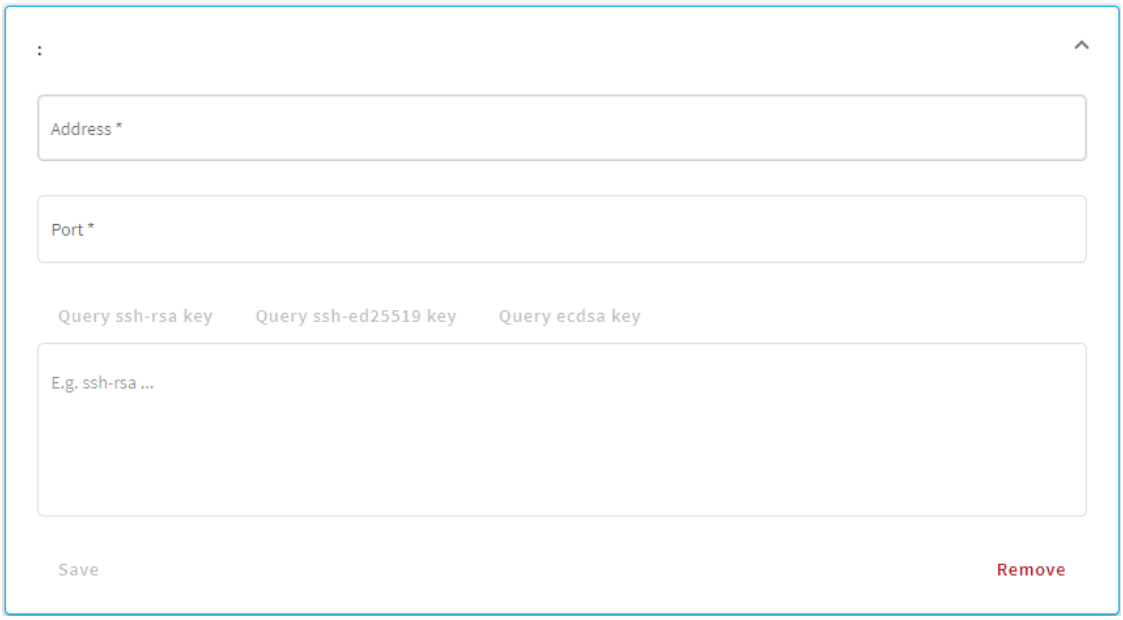
Manually adding the host key of a server

The following describes how to add the host key manually.

To add the host key manually

1. Navigate to the **Traffic Controls** > **SSH** > **Server Host Keys** and click **Create new**.

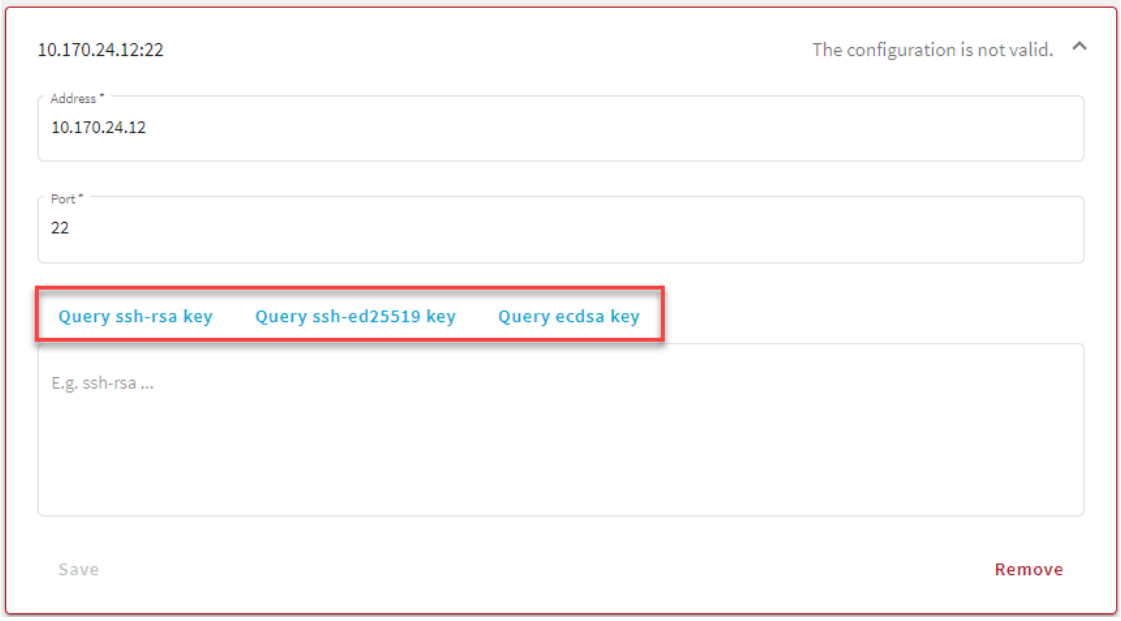
   Figure 227: Traffic Controls > SSH > Server Host Keys — Create new

2. Enter the IP address and port of the server into the **Address** and **Port** fields.

   You can use IPv4 and IPv6 addresses as well.

Figure 228: Traffic Controls > SSH > Server Host Keys — Example of active Query fields



3. Once you fill out the **Address** and **Port** fields, the **Query...** fields become active and you can query the public part of the host key of the server.

   SPS allows you to use the following SSH host keys:

   - RSA (ssh-rsa), which is the most widely used public-key algorithm for the SSH key. In SPS, uploading RSA SSH host keys are supported in PKCS #1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

     NOTE:

     TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

   - Ed25519 (ssh-ed25519), which offers a better security and faster performance compared to RSA.

     In SPS, uploading Ed25519 SSH host keys are supported in PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

   - ECDSA NIST P-256 (ecdsa-sha2-nistp256), which is a variant of the Digital Signature Algorithm (DSA). In SPS, uploading ECDSA SSH host keys are supported in SEC1 PEM, PKCS #8 PEM, OpenSSH (openssh-key-v1), and PuTTY private key formats.

   You can have multiple SSH server host keys on SPS for the same server, however, you cannot set more than one key for each type. For example, you can keep your old RSA SSH key and generate a new Ed25519 key but you cannot set two RSA keys.

   Alternatively, paste the public part of the host key of the server.

4. Click **Save**.

Creating and editing protocol-level SSH settings

SSH settings determine the parameters of the connection on the protocol level. For example, when the server-side connection is built, the timeout value, and greeting message of the connection. The following parameters determine which algorithms are used in the connections, and can be set independently for the client and the server side: key exchange, host key, cipher, MAC, and compression algorithms.

CAUTION:

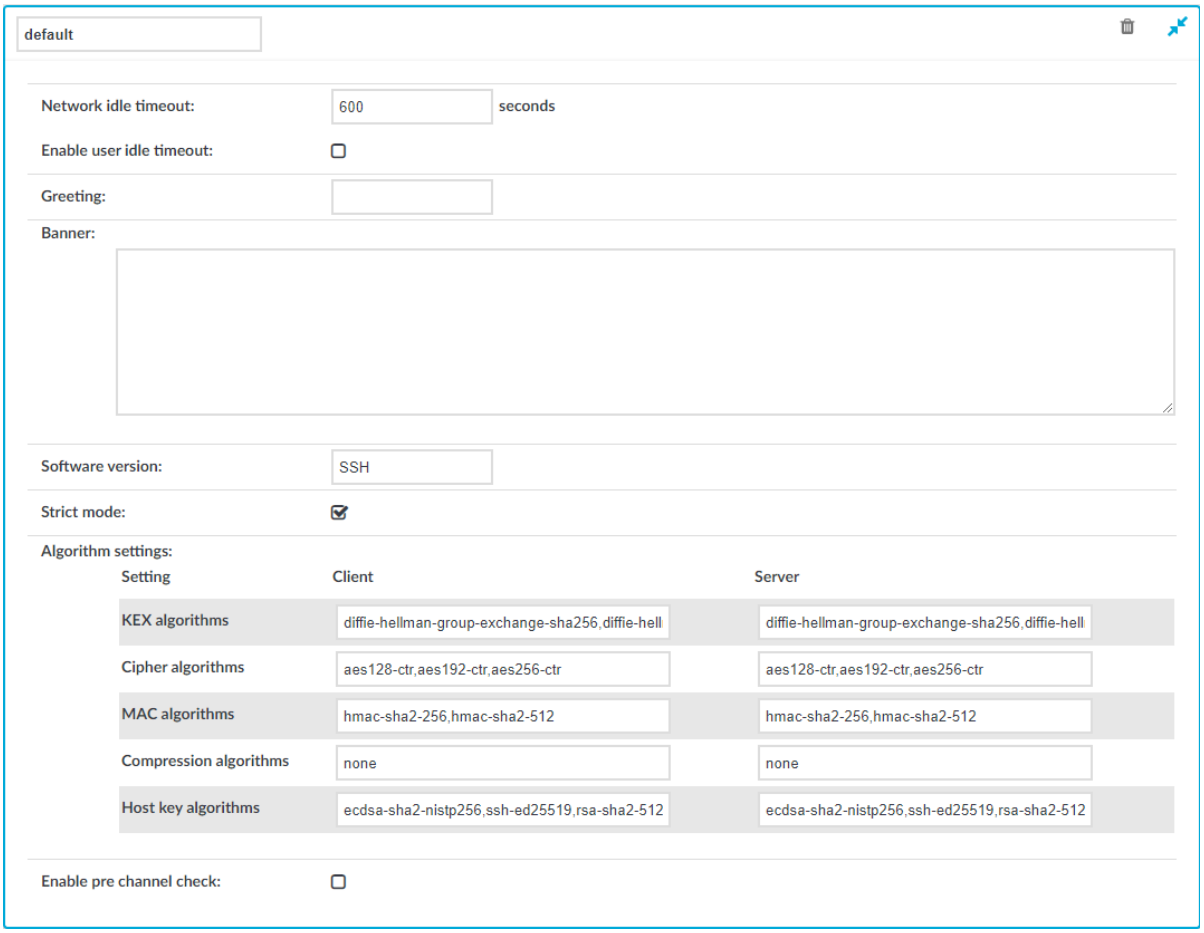
Before modifying any of the algorithm settings, check whether the default algorithms are supported by your SSH client and server.

If yes, then you can leave these settings untouched.

If not and you need to amend the default algorithm settings, ensure that the client and server sides are harmonized. You can either do that in One Identity Safeguard for Privileged Sessions (SPS) or on the client/server itself.

Note that modifying algorithm settings in SPS is recommended to advanced users only. If you are unsure about which settings to amend, then contact our Support Team for assistance.

Figure 229: Traffic Controls > SSH > Settings — SSH settings



To create a new SSH settings profile or edit an existing one

1. Navigate to the **Traffic Controls** > **SSH** > **Settings** and click to create an SSH setting profile. Enter a name for the profile (for example **`strongencryption`**).
2. Click to display the parameters of the SSH connection.
3. **Network idle timeout**: Connection timeout value in seconds. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

   Even if the user is not active, the session can contain activity that must be audited (for example, the output of a script). The idle timeout period will start only after this activity has stopped.

CAUTION:

Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.

4. **User idle timeout**: If no user activity is detected, terminate the session after the configured time has passed since the last user activity.

   This can be useful if only user-generated network traffic is important in a session. By using this option, situations described in the caution of **Network idle timeout** (such as a taskbar clock keeping the network traffic open indefinitely) can be avoided. To enable user idle timeout, select **Enable user idle timeout** and enter a value that is greater than or equal to the value of **Network idle timeout**.

5. To display a greeting message to the clients after connecting the server, enter the message into the **Greeting** field.

6. To display a banner message to the clients before authentication (as specified in RFC 4252 — The Secure Shell (SSH) Authentication Protocol), enter the message into the **Banner** field. For example, this banner can inform the users that the connection is audited.

7. Optional. You can specify additional text to append to the SSH protocol banner, for example to mask the OpenSSH version upon connection. Enter the text in the **Software version** field.

8. If needed, modify the encryption parameters. SPS enforces policies on the various elements of the encrypted SSH communication, such as the MAC, key-exchange, and cipher algorithms that are permitted to be used. The parameters can be set separately for the client and for the server side. The attributes are comma-separated strings listing the enabled methods/algorithms, in the order of preference.

   For a complete list of the available parameters, see Supported encryption algorithms on page 641.

   NOTE:

   Do not use the CBC block cipher mode, or the `diffie-hellman-group1-sha1` key exchange algorithm. For details, see *Supported encryption algorithms* in the *Administration Guide*.

9. To check the protocol-level parameters of the connections very strictly, select the **Strict mode** option. This option is enabled by default. When this option is enabled:

   SPS will reject connections that use unrealistic parameters, for example:

   - The number of columns and rows of the terminal is bigger or equal than 512.
   - The size of the screen is greater than 8192 pixels in either directions.

   SPS will reject port-forwarding connections where the address in the port-forwarding request and the channel-opening request does not match.

NOTE: Strict mode can interfere with certain client or server applications.

NOTE: Strict mode is not working with the Windows internal Bash/WSL feature, because it uses a very large terminal window size. Using Windows internal Bash/WSL is not supported.

10. Before establishing the server-side connection, SPS can evaluate the connection and channel policies to determine if the connection might be permitted at all, for example it is not denied by a Time Policy. To enable this function, select the **Enable pre channel check** option. That way SPS establishes the server-side connection only if the evaluated policies permit the client to access the server.
11. Click **Commit**.
12. Select this settings profile in the **SSH settings** field of your connections.

Supported encryption algorithms

The following tables contain all the encryption algorithms you can configure One Identity Safeguard for Privileged Sessions (SPS) to recognize. If you use a configuration that is only partially supported, SPS might ignore the connection without warning.

NOTE: Do not use the CBC block cipher mode, or any sha1-based KEX, MAC, or host key algorithm, which are considered weak.

Key exchange algorithms

The default SPS configuration for both the client and the server is the following:

```
ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,diffie-hellman-group14-sha256
```

The following key exchange (KEX) algorithms are recognized:

Figure 230: Key exchange (KEX) algorithms

Key exchange (KEX)	Default	Comment
ecdh-sha2-nistp256	✓	
ecdh-sha2-nistp384	✓	
ecdh-sha2-nistp521	✓	
diffie-hellman-group1-sha1	-	Not recommended
diffie-hellman-group14-sha1	-	Not recommended
diffie-hellman-group14-sha256	✓	

Key exchange (KEX)	Default	Comment
diffie-hellman-group15-sha512	-	
diffie-hellman-group16-sha512	✓	
diffie-hellman-group17-sha512	-	
diffie-hellman-group18-sha512	✓	
diffie-hellman-group-exchange-sha256	✓	
diffie-hellman-group-exchange-sha1	-	Not recommended

During an SSH session, SPS performs a key re-exchange after each gigabyte of transmitted data or after each hour of connection time, whichever comes sooner.

Cipher algorithms

The default SPS configuration for both the client and the server is the following:

```
aes128-ctr,aes192-ctr,aes256-ctr
```

The following cipher algorithms are recognized:

Figure 231: Cipher algorithms

Cipher algorithm	Default	Comment
3des-cbc	–	Not recommended
blowfish-cbc	–	Not recommended
twofish256-cbc	–	Not recommended
twofish-cbc	–	Not recommended
twofish192-cbc	–	Not recommended
twofish128-cbc	–	Not recommended
aes256-cbc	–	Not recommended
aes192-cbc	–	Not recommended
aes128-cbc	–	Not recommended
aes256-ctr	✓	
aes192-ctr	✓	
aes128-ctr	✓	
serpent256-cbc	–	Not recommended

Cipher algorithm	Default	Comment
serpent192-cbc	–	Not recommended
serpent128-cbc	–	Not recommended
arcfour	–	Not recommended
idea-cbc	–	Not recommended
cast128-cbc	–	Not recommended
none	–	Means no cipher algorithm; not recommended

Message authentication code (MAC) algorithms

The default SPS configuration for both the client and the server is the following:

```
hmac-sha2-256,hmac-sha2-512
```

The following MAC algorithms are recognized:

Figure 232: Message Authentication Code (MAC) algorithms

MAC	Default	Comment
hmac-sha1	–	Not recommended
hmac-sha1-96	–	Not recommended
hmac-md5	–	Not recommended
hmac-md5-96	–	Not recommended
hmac-sha2-256	✓	
hmac-sha2-512	✓	

SSH compression algorithms

The default SPS configuration for both the client and the server is the following:

```
none
```

The following SSH compression algorithms are recognized:

Figure 233: SSH compression algorithms

SSH compression algorithm	Default	Comment
zlib	–	Not recommended
none	✓	Means no compression

Host key algorithms

The default SPS configuration for both the client and the server is the following:

```
ecdsa-sha2-nistp256,ssh-ed25519,rsa-sha2-512,rsa-sha2-256,ssh-rsa
```

The following host key algorithms are recognized:

Figure 234: Host key algorithms

Host key algorithms	Default	Comment
ecdsa-sha2-nistp256	✓	
ssh-ed25519	✓	
rsa-sha2-512	✓	
rsa-sha2-256	✓	
ssh-rsa	✓	Not recommended NOTE: The ssh-rsa public key signature algorithm that depends on SHA-1 is not recommended and will be disabled in a future release.

Using Sudo with SPS

With the SPS and Sudo integration, you can collect and analyze Sudo session recordings, called iologs in Sudo terminology, in SPS.

By using SPS to collect and analyze Sudo session recordings, your Sudo recordings are stored and indexed by SPS, and you can use the Search interface, for example, to view the recordings, list commands run during a Sudo session, and so on.

To be able to use Sudo with SPS, perform the following settings:

1. In SPS, create a new Sudo connection as described in Setting up Sudo connections in SPS.
2. On the client side, configure Sudo as described in Configuring Sudo.

Setting up Sudo connections in SPS

To enable the SPS and Sudo integration, perform the following steps in SPS.

For a detailed description on configuring connections, see General connection settings on page 482.

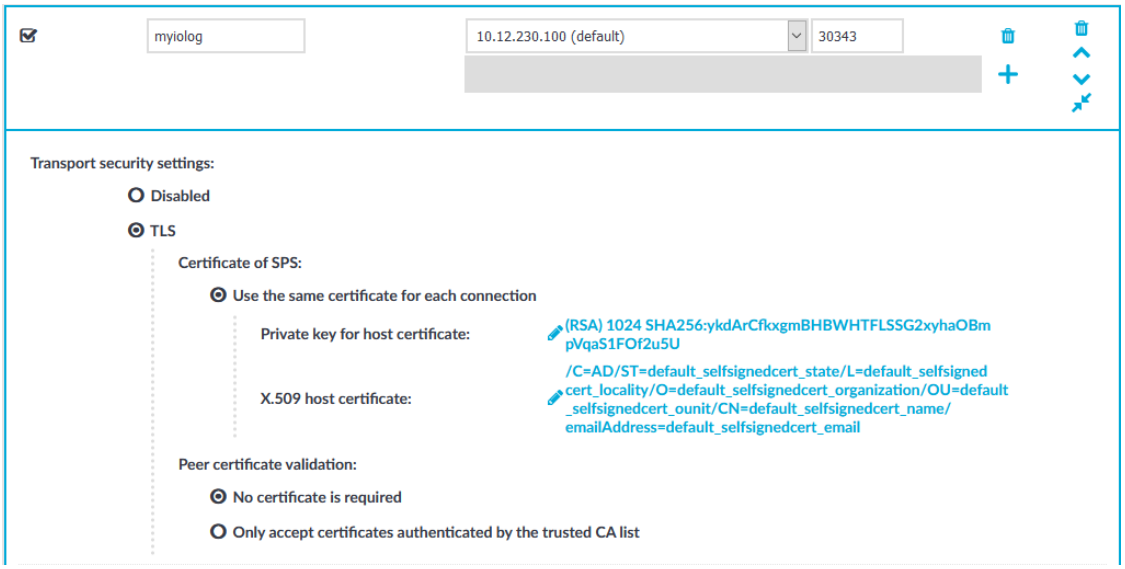
Prerequisites

- You have the SPS license required to integrate SPS and Sudo. For more information, contact our Sales Team.
- An X.509 certificate and its private key to encrypt the communication between the client and SPS. Use your own PKI system to generate these certificates, as they cannot be created on SPS.

To create a Sudo connection in SPS

1. Navigate to **Traffic Controls** > **Sudo iolog** > **Connections**.

Figure 235: Traffic Controls > Sudo iolog > Connections — Creating a Sudo connection in SPS



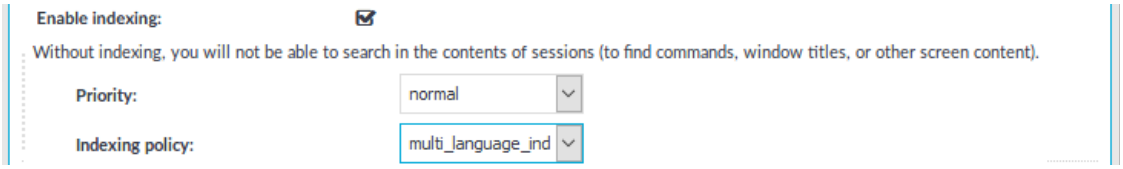
2. Since SPS can have multiple network interfaces, select an IP address where the Sudo clients can send the iologs. If required, you can change the port number.
3. TLS is disabled by default as you have to configure certificates manually. Make sure that you enable it as iologs carry highly sensitive information.

   To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.
4. Select the certificate to show to the peers.
   a. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
   b. Select **Use the same certificate for each connection**.
   c. Select **Private key for host certificate**, click and upload the private key.
   d. Select **X.509 host certificate**, click and upload the certificate.
5. Select how SPS should authenticate the peers.
   - To permit connections from peers without requesting a certificate, select **No certificate is required**.

- To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
    1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
    2. Select **Only accept certificates authenticated by the trusted CA list**.
    3. Select the certificate authority list to use in the **Trusted CA** field.

6. Select **Enable indexing**.

Figure 236: Traffic Controls > Sudo iolog > Enable indexing — Select indexing policy

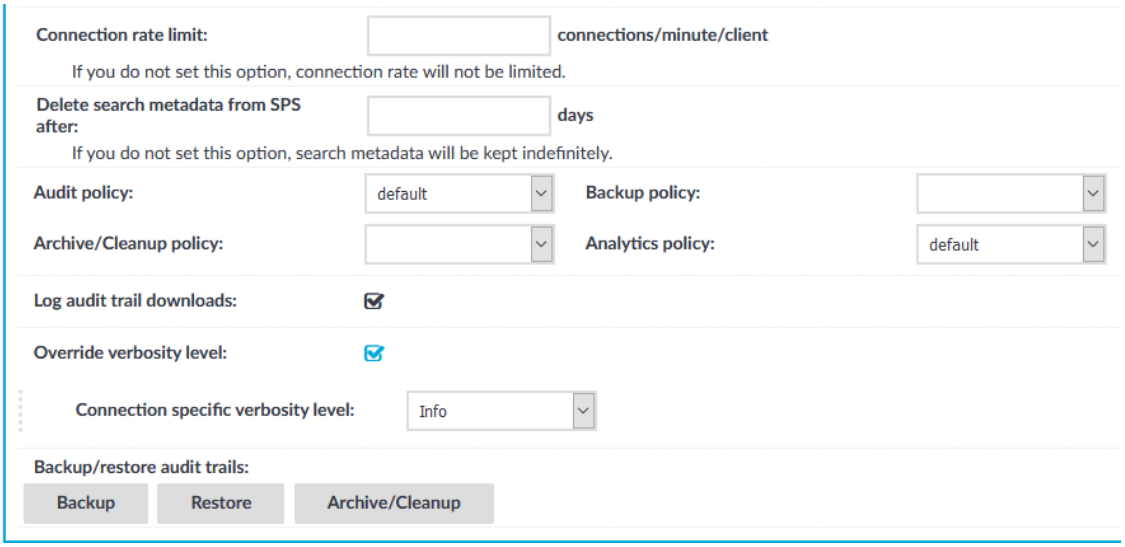


- To determine the priority level of indexing this connection, select the appropriate **Priority** level. Selecting a high priority level means that the trails of this connection are indexed first. Selecting a low priority level means that the trails of this connection are also indexed, but there might be a delay in indexing if there are a lot of high-priority connections waiting to be indexed. Selecting **near real-time** means that the indexing of sessions starts when sessions are still ongoing.
- Select the **Indexing policy** to be used.

  Both built-in indexer policies feature automatic language detection. To specify a particular language detection configuration, select the indexing policy if you have created before. For more information, see Configuring the internal indexer.

7. To limit the number of new connection requests accepted from a single client IP address per minute, enter the maximum number of accepted connections into the **Connection rate limit** field.

Figure 237: Traffic Controls > Sudo iolog > Connections — Additional settings



8. Enter how long SPS (in days) should keep the metadata into the **Delete search metadata from SPS after** field. For example, if you specify `365`, SPS deletes the data of connections older than a year. Enter zero (`0`) to keep the data indefinitely (this is also the default behavior of SPS).
9. If configured, select a policy to control various aspects of the connection. For more information, see Policies.
10. If you want to find out if the audit trail file of a relevant indexed session has already been downloaded, select **Log audit trail downloads**.
11. *Optional:* To set a verbosity level for this connection, select **Override verbosity level** and select the desired log level from the **Connection specific verbosity level** field.

    NOTE: The new verbosity level applies only to new sessions started after committing the change. The verbosity level of active sessions do not change.
12. Click **Commit**.

    Before you can use the SPS and Sudo integration, perform the Sudo-side configuration steps as described in Configuring Sudo.

Configuring Sudo

To enable the SPS and Sudo integration, perform the following steps on the client side.

Prerequisites

- For Sudo, you need a host with Sudo 1.9 or higher installed.
- You have the SPS license required to integrate SPS and Sudo. For more information, contact our Sales Team.
- For SPS, you have configured Sudo iolog connections as described in Setting up Sudo connections in SPS.

To enable the SPS and Sudo integration

1. In Sudo, open the sudoers file using **visudo**.
2. Add the following lines:

   ```
   Defaults ignore_iolog_errors
   Defaults log_servers = <IP of SPS>:<PORT of SPS>
   Defaults log_output
   Defaults log_input
   ```

   Where the options are as follows:

   Table 11: Options for the Sudo-side configuration

Option	Description
ignore_ iolog_ errors	Allow running commands even if `sudoers` cannot write to the I/O log.
log_ servers	Specify the IP address of your SPS. Additionally, you can specify the port number of SPS. If you use Transport Layer Security (TLS) for encryption, you must also specify it as described below.
log_ output	Enable recording any change, which appears on the screen.
log_input	Ensure recording of any user input, that is, anything the user types.

3. To require encryption (recommended), use TLS as follows:
   - Configure the `log_servers` option in the `sudoers` file:

     ```
     Defaults log_servers = <IP of SPS>:<PORT of SPS>(tls)
     ```

   - Configure the `log_server_cabundle`, `log_server_peer_cert`, or `log_server_peer_key` settings with the required TLS settings. For more information, see Securing server connections in the *Sudoers Manual*.

     For example, add the path to the certificate authority bundle file in `.pem` format

as shown in the example below:

```
Defaults log_server_cabundle = <path_to_PEM_file>/<file_name>.pem
```

Expected result

On the client side, start typing `sudo` and open any program.

In SPS, you can view the session using the Search interface.

Telnet-specific settings

The following sections describe configuration settings available only for the Telnet protocol. Use the following policies to control who, when, and how can access the Telnet connection. For a list of supported client applications, see Supported protocols and client applications on page 31.

- *Channel Policy*: The Telnet protocol has only one channel type with no special configuration options. The available channel policy options are the following: **Type**, **From**, **Target**, **Time policy**, **Four-eyes**, **Record audit trail**, **Gateway groups**, **Remote groups**, and **Content policy**. For details on configuring these options, see Creating and editing channel policies on page 495.
- *TLS support*: To enable TLS-encryption for your Telnet connections, see Enabling TLS-encryption for Telnet connections on page 652.
- *Authentication Policy*: Authentication policies describe the authentication methods allowed in a connection. Different methods can be used for the client and server-side connections. For details, see Creating a new Telnet authentication policy on page 655.
- *Telnet settings*: Telnet settings determine the parameters of the connection on the protocol level, including timeout value, and so on. For details, see Creating and editing protocol-level Telnet settings on page 658.
- *User lists in Channel Policies*: User lists affect Telnet connections only when they are used together with Gateway Authentication. For details, see Configuring gateway authentication on page 862.
- *Content Policy*: Content policies allow you to inspect the content of the connections for various text patterns, and perform an action if the pattern is found. For example, One Identity Safeguard for Privileged Sessions (SPS) can send an e-mail alert if a specific command is used in a Telnet terminal session. For details, see Creating a new content policy on page 499.
- *Authentication and Authorization plugin*:

  One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS to external systems to authenticate or authorize the user before authenticating on the target server. Such plugins can also be used to request additional information from the users, for example, to perform multi-factor authentication.

  For details, see Integrating external authentication and authorization systems on page 884.

Enabling TLS-encryption for Telnet connections

The following steps describe how to enable TLS-encryption in a Telnet connection policy. Note that when using encryption, One Identity Safeguard for Privileged Sessions (SPS) automatically changes the port number of the connection policy to 992.

Prerequisites

Depending on your requirements, one or more of the following might be needed:

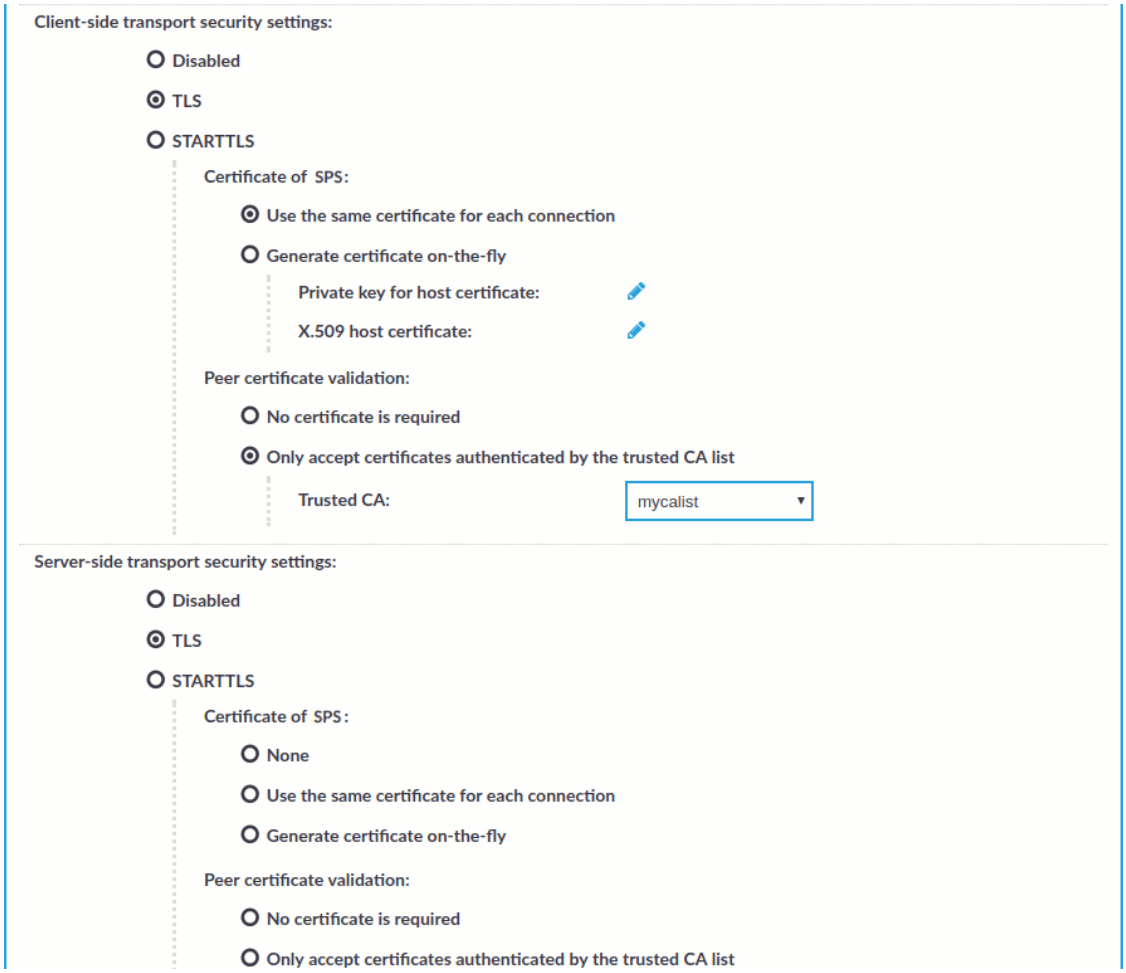
- An X.509 certificate and its private key. SPS can display the same certificate to the peers on both the client and the server side. You can also use different certificates for the client and server sides. Use your own PKI system to generate these certificates, as they cannot be created on SPS. Note that the Common Name of the certificate must contain the domain name or the IP address of SPS, otherwise the clients might reject the certificate.
- To generate certificates on-the-fly for a connection, a signing certificate authority is required. For details on creating a signing CA, see Signing certificates on-the-fly on page 529.
- To require the peers of SPS to have an X.509 certificate signed by a specific Certificate Authority, a list of the trusted certificate authorities is needed. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To enable TLS-encryption in a Telnet connection policy

1. Navigate to **Traffic Controls** > **Telnet** > **Connections** and select the connection policy in which you want to enable TLS.

Figure 238: Traffic Controls > Telnet > Connections — Enabling TLS-encryption for Telnet connections



2. Set the encryption settings used between the client and SPS in the **Client-side transport security settings** section.
    - To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.
    - To enable encrypted connections that use the STARTTLS method, select **STARTTLS**. Note that the peer must use the STARTTLS method. Unencrypted connections will be terminated after a brief period.
3. Select the certificate to show to the peers.
    - To use the same certificate for every peer, complete the following steps.
        1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
        2. Select **Use the same certificate for each connection**.

3. Select **Private key for host certificate**, click  and upload the private key.
4. Select **X.509 host certificate**, click  and upload the certificate.

- To use a separate certificate for every connection, complete the following steps.
  1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
  2. Select **Generate certificate on-the-fly**.
  3. In the **Signing CA** field, select the certificate authority to use.

4. Select how SPS should authenticate the peers.
   - To permit connections from peers without requesting a certificate, select **No certificate is required**.
   - To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
     1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
     2. Select **Only accept certificates authenticated by the trusted CA list**.
     3. Select the certificate authority list to use in the **Trusted CA** field.
5. Set the encryption settings used between SPS and the server in the **Server-side transport security settings** section.
   - To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.
   - To enable encrypted connections that use the STARTTLS method, select **STARTTLS**. Note that the peer must use the STARTTLS method. Unencrypted connections will be terminated after a brief period.
6. Select the certificate to show to the server.
   - If the server does not require a certificate from SPS, select **None**.
   - To use the same certificate for every peer, complete the following steps.
     1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
     2. Select **Use the same certificate for each connection**.
     3. Select **Private key for host certificate**, click  and upload the

private key.

4. Select **X.509 host certificate**, click  and upload the certificate.

- To use a separate certificate for every connection, complete the following steps.
  1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
  2. Select **Generate certificate on-the-fly**.
  3. Select the certificate authority to use in the **Signing CA** field.

Limitations

NOTE: When using the **Use the same certificate for each connection** option and the connection policy that allows access to multiple servers using HTTPS, the client applications will display a warning because the certificate used in the connection will be invalid (namely, the Common Name of the certificate will not match the hostname or IP address of the server).

NOTE: Import the certificate of the signing Certificate Authority to your clients. Otherwise, the client applications will display a warning due to the unknown Certificate Authority.

7. Select how SPS should authenticate the peers.
   - To permit connections from peers without requesting a certificate, select **No certificate is required**.
   - To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
     1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
     2. Select **Only accept certificates authenticated by the trusted CA list**.
     3. Select the certificate authority list to use in the **Trusted CA** field.
8. Click **Commit**.

Expected result

The encryption settings are applied to the connection policy.

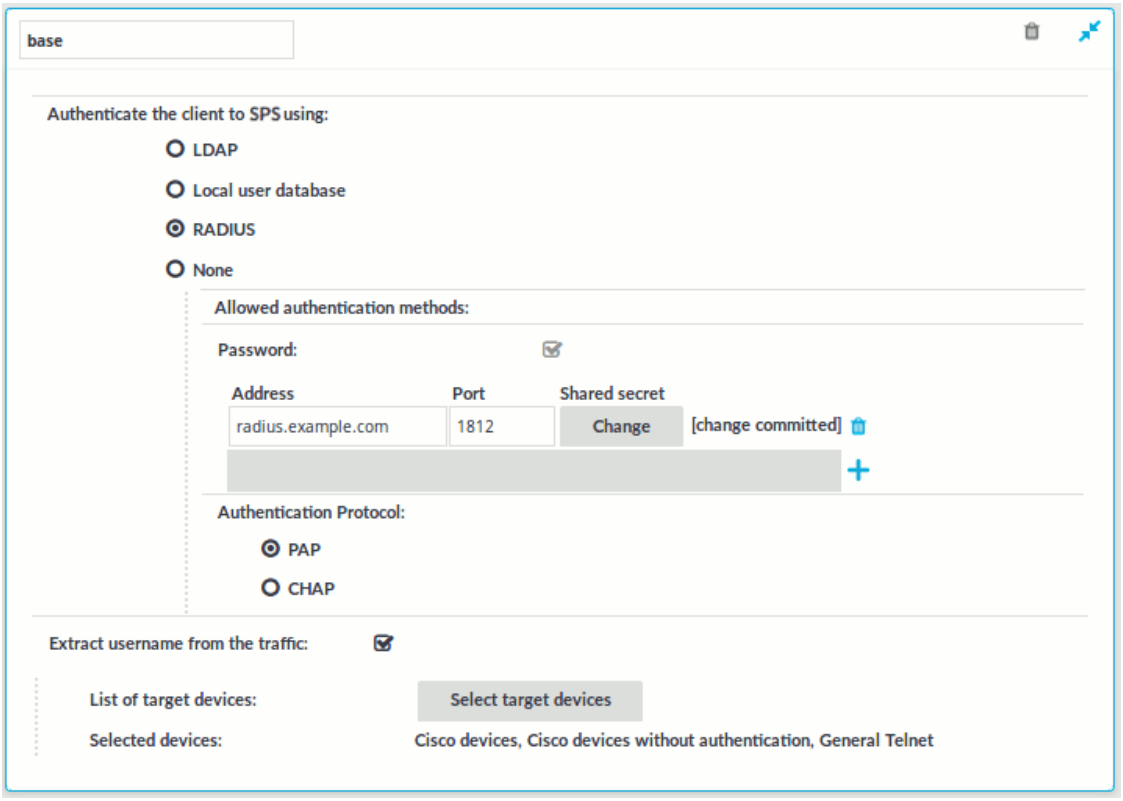
Creating a new Telnet authentication policy

An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the client can authenticate to the target server. Separate authentication methods can be used on the client and the server-side of the connection.

To create a new authentication policy

1. Navigate to **Traffic Controls** > **Telnet** > **Authentication Policies**, and click  .

Figure 239: Traffic Controls > Telnet > Authentication Policies — Configuring Telnet authentication policies



2. Enter a name for the policy into the **Name** field.
3. Select the authentication method used on the client-side in the **One Identity Safeguard for Privileged Sessions (SPS)Authenticate the client to SPS using** field. For the client-side connection, SPS can authenticate the client inband (within the Telnet protocol) using the following authentication methods:
   - **LDAP**: SPS will authenticate the client to the LDAP database set in the **LDAP Server** of the connection policy. To use LDAP authentication on the client side, select **Authenticate the client to SPS using** > **LDAP**.

     NOTE: SPS will authenticate the client-side connection to the LDAP server configured in the connection policy. This is not necessarily the same as the

LDAP server used to authenticate the users accessing the SPS web interface.

- **Local user database**: Authenticate the client locally on the SPS gateway using a **Local user database**. Select the database to use in the **Local user database** field. For details on creating a **Local User Database**, see Creating a Local User Database on page 538.
- **RADIUS**: SPS will authenticate the client to the specified RADIUS server. Select **Authenticate the client to SPS using** > **RADIUS**, enter the IP address or hostname of the RADIUS server into the **Address** field, the port number of the RADIUS server into the **Port** field, and the shared secret of the RADIUS server into the **Shared secret** field. Only password-authentication is supported (including one-time passwords), challenge-response based authentication is not.

  Use an IPv4 address.

  To add more RADIUS servers, click  and fill in the respective fields.
- **None**: Do not perform client-side authentication, the client will authenticate only on the target server.

  CAUTION:

  Hazard of security breach. If the **None** authentication option is selected on the client side and SPS is configured to use public-key or certificate based authentication on the server, the user will not be authenticated at all unless gateway authentication is required for the connection.

4. Click **Commit**.

NOTE: Consider the following:

- The client-side authentication settings apply for authenticating the user inband to the SPS gateway, and is independent from the gateway authentication performed on the SPS web interface. The web-based gateway authentication is an out-of-band gateway authentication method that can be required by the connection policy. For details on out-of-band gateway authentication, see Configuring out-of-band gateway authentication on page 863.

  Gateway authentication on the SPS web interface can be used together with authentication policies. In an extreme setting, this would mean that the user has to perform three authentications: a client-side gateway authentication within the SSH protocol to SPS, an out-of-band gateway authentication on the SPS web interface, and a final authentication on the target server.
- The Connection Policy will ignore the settings for server-side authentication (set under **Relayed authentication methods** for SSH protocol) if a Credential Store is used in the Connection Policy.

Extracting username from Telnet connections

For specific devices, it is now possible to extract the username from Telnet connections with the help of patterns (including TN3270 and TN5250 systems).

To select patterns or request a custom pattern

1. Navigate to **Traffic Controls** > **Telnet** > **Authentication Policies** and enable **Extract username from the traffic**.
2. Click **Select target devices** to display the list of available target devices. Select the respective device(s) in the **Available devices** column and click **Add**.

   NOTE: You can only add one TN3270 specific device to the authentication policy.

   To remove a device from the **Target devices** column, select it and click **Remove**.
3. Click **OK**. The target devices are listed after **Selected devices**.
4. If you cannot find your device in the list of available target devices, request a custom Pattern Set. To do this, contact our Support Team.
5. To upload the custom pattern set you received, navigate to **Traffic Controls** > **Telnet** > **Pattern Sets**, browse for the file and click **Upload**.
6. To delete a custom Pattern Set from One Identity Safeguard for Privileged Sessions (SPS), click in the respective row. Generic Pattern Sets cannot be deleted.

Creating and editing protocol-level Telnet settings

Procedure

Telnet settings determine the parameters of the connection on the protocol level, including timeout value, and so on. Complete the following procedure to create a new Telnet settings profile or edit an existing one:

CAUTION:

Modifying the Telnet settings is recommended only to advanced users. Do not modify these settings unless you exactly know what you are doing.

To create and edit protocol-level Telnet settings

1. Navigate to the **Settings** tab of the **Traffic Controls** > **Telnet** menu item and click + to create a Telnet setting profile. Enter a name for the profile (for example `telnet_special`).

2. Click  to display the parameters of the connection.
3. Modify the parameters as needed. The following parameters are available:
    - **Network idle timeout**: Connection timeout value in seconds. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

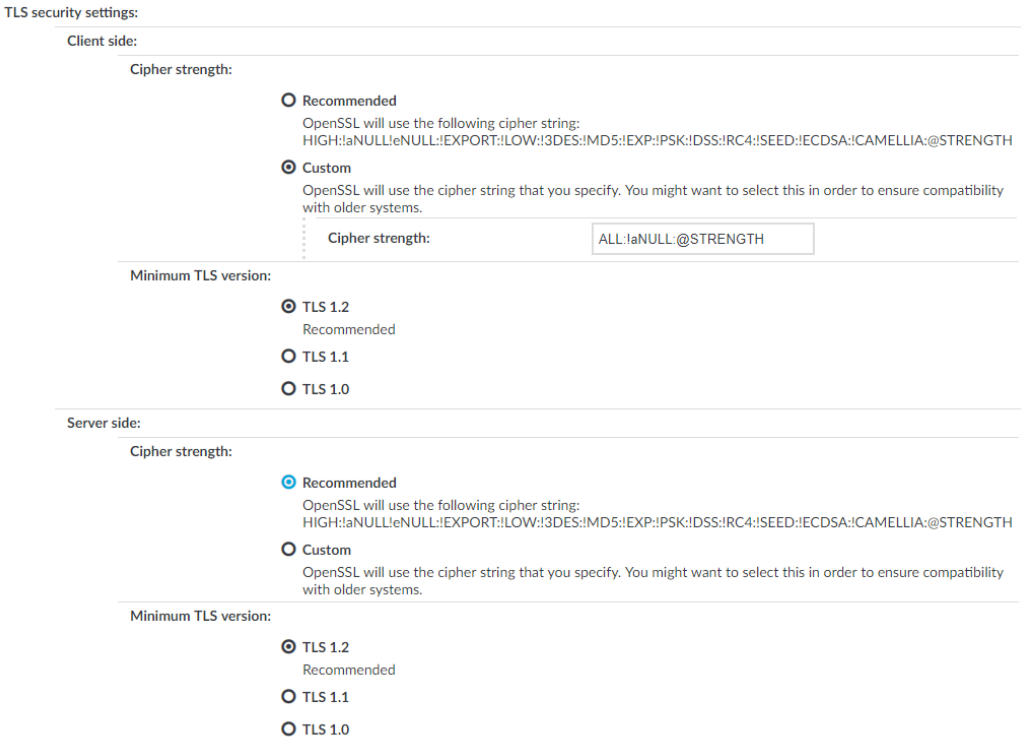
      Even if the user is not active, the session can contain activity that must be audited (for example, the output of a script). The idle timeout period will start only after this activity has stopped.

      CAUTION:

      Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.
    - **User idle timeout**: If no user activity is detected, terminate the session after the configured time has passed since the last user activity.

      This can be useful if only user-generated network traffic is important in a session. By using this option, situations described in the caution of **Network idle timeout** (such as a taskbar clock keeping the network traffic open indefinitely) can be avoided. To enable user idle timeout, select **Enable user idle timeout** and enter a value that is greater than or equal to the value of **Network idle timeout**.
    - **Enable pre channel check**: Select this option to evaluate the connection and channel policies before establishing the server-side connection. That way if the connection is not permitted at all, SPS does not establish the server-side connection.
    - To configure TLS security settings on both the **Client side** and the **Server side**, proceed to **TLS security settings**.

Figure 240: <Protocol> Control > Settings > TLS security settings - configuring TLS security settings



- **Cipher strength** specifies the cipher string OpenSSL will use. The following options are possible:
  - **Recommended**: this setting only uses ciphers with adequate security level.
  - **Custom**: this setting allows you to specify the list of ciphers you want to permit SPS to use in the connection. This setting is only recommended to ensure compatibility with older systems. For more details on customizing this list, check the 'openssl-ciphers' manual page on your SPS appliance.

    For example: `ALL:!aNULL:@STRENGTH`
- **Minimum TLS version** specifies the minimal TLS version SPS will offer during negotiation. The following options are possible:
  - **TLS 1.2**: this setting only offers TLS version 1.2 during the negotiation. This is the recommended setting.
  - **TLS 1.1**: this setting offers TLS version 1.1 and later versions during the negotiation.
  - **TLS 1.0**: this setting offers TLS version 1.0 and later versions during the negotiation.

NOTE: Setting up sessions to legacy systems that do not support at least TLS 1.2 is only possible when the security level of the connection is degraded to 0, which is possible by specifying the TLS ciphers manually and appending the string `:@SECLEVEL=0` to the cipher list. However, this setting also enables the use of known vulnerable algorithms and key sizes, therefore it is absolutely critical to only use such connection settings when it is necessary and when you can fully trust your network between SPS and the legacy system. It is strongly recommended to use different security settings on the server and the client side of the connection, when degrading the security level of a connection is unavoidable.

NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported.

4. Click **Commit**.
5. To display a banner message to the clients before authentication, enter the message into the **Banner** field. For example, this banner can inform the users that the connection is audited.
6. Select this settings profile in the **TELNET settings** field of your connections.

Inband destination selection in Telnet connections

When using inband destination selection in Telnet connections, the user can provide the server address and the username using the following methods:

- By setting the `TELNET ENVIRON` option using the `SERVER` environment variable in the `server:port` format.
- By setting the `TELNET ENVIRON` option using the `USER` environment variable in the `user@server:port` format.
- If neither the `SERVER` nor the `USER` environment variable, One Identity Safeguard for Privileged Sessions (SPS) displays a terminal prompt where the user can enter the username and the server address.

VMware Horizon View connections

The following sections describe how to use One Identity Safeguard for Privileged Sessions (SPS) to control and audit VMware Horizon View (formerly known as VMware View) connections. When using SPS to control and audit VMware Horizon View connections, the following requirements and restrictions apply:

- Only connections using the Remote Desktop (RDP) display protocol are supported. Connections using the PCoIP or HP Remote Graphics Software display protocols are not supported.
- Both direct connections and tunnel connections are supported.
- The VMware Horizon View connections must pass SPS directly. It is best if SPS is deployed directly before the Virtual Desktops accessed with VMware Horizon View, and connections are configured in transparent mode.

  Deploying SPS that way has the advantage of auditing connections even if the clients access the Virtual Desktops directly, without using a View Connection Server.

  NOTE: Using non-transparent mode is also possible if the VMware Horizon View traffic is routed to SPS with an external device (for example, a firewall).

SPS treats VMware Horizon View connections that satisfy these criteria as common RDP connections. All the features of SPS that are available for RDP connections can be used with VMware Horizon View connections as well, for example, four-eyes authorization, auditing and replaying, indexing the recorded audit trails, and so on. For details on RPD-specific settings, see RDP-specific settings on page 577.

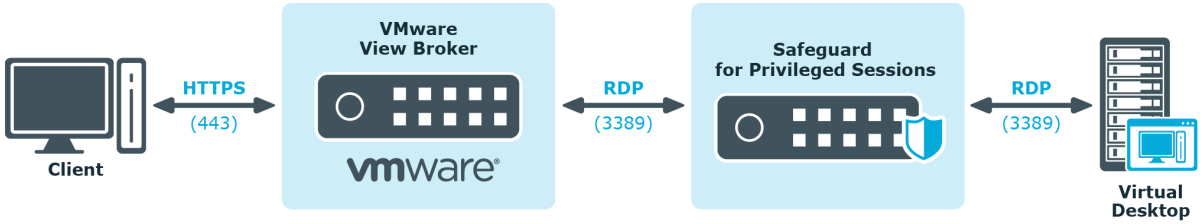
One Identity Safeguard for Privileged Sessions (SPS) deployment scenarios in a VMware environment

One Identity Safeguard for Privileged Sessions (SPS) supports a variety of deployment scenarios, which make it really flexible when it comes to deployment. The following network topologies illustrate typical SPS VMware Horizon View deployment scenarios.

Client - Broker - SPS - Server

SPS is deployed between the Broker and the virtual desktop, where the RDP traffic is embedded into a HTTPS tunnel between the Client and the Broker.

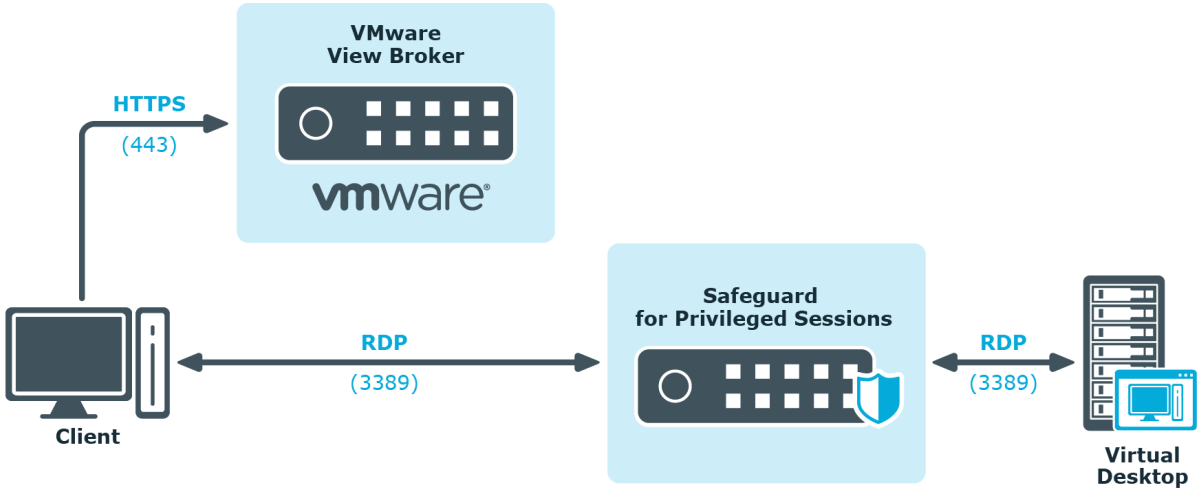
Figure 241: Client - Broker - SPS - Server



Client - SPS - Server

SPS is deployed between the Client and the virtual desktop, the client makes a direct RDP connection to the Server, without tunneling it through the Broker.

Figure 242: Client - SPS - Server



VNC-specific settings

The following sections describe configuration settings available only for the Virtual Networking (VNC) protocol. Use the following policies to control who, when, and how can access the VNC connections. For a list of supported client applications, see Supported protocols and client applications on page 31.

CAUTION:

To monitor VNC connections, enable user authentication on your VNC server. One Identity Safeguard for Privileged Sessions (SPS) automatically terminates unauthenticated connections.

- *Channel Policy*: The VNC protocol has only one channel type with no special configuration options. The available channel policy options are the following: **Type**, **From**, **Target**, **Time policy**, **Four-eyes**, **Record audit trail**, **Gateway groups**, **Remote groups**, and **Content policy**. For details on configuring these options, see Creating and editing channel policies on page 495.
- *TLS support*: To enable TLS-encryption for your VNC connections, see Enabling TLS-encryption for VNC connections on page 665.
- *VNC settings*: VNC settings determine the parameters of the connection on the protocol level, including timeout value, and so on. For details, see Creating and editing protocol-level VNC settings on page 669.
- *User lists in Channel Policies*: User lists affect VNC connections only when they are used together with Gateway Authentication. For details, see Configuring gateway authentication on page 862.
- *Content Policy*: Content policies allow you to inspect the content of the connections for various text patterns, and perform an action if the pattern is found. For example, SPS can send an e-mail alert if a specific window title appears in RDP and VNC connections. For details, see Creating a new content policy on page 499.
- *WebSocket/VNC audit trails*: You can replay audit trails of a WebSocket connection in your browser or using the Safeguard Desktop Player application only if it contains Virtual Network Computing (VNC) traffic. For all other WebSocket connections, export the audit trail as a PCAP file and replay it using the Safeguard Desktop Player application.

  For more information, see Supported HTTP channel types.

Enabling TLS-encryption for VNC connections

The following steps describe how to enable TLS-encryption in a VNC connection policy.

NOTE:

Some vendors may use custom protocol elements and TLS-encryption that do not have available documentation. As a result, these cannot be audited by One Identity Safeguard for Privileged Sessions (SPS). Regardless of vendors, only the custom features described in the RFC 6143 are supported. As for encryptions, only those completely TLS-encapsulated streams can be processed where the TLS encryption process was started before the VNC protocol handshake.

Prerequisites

Depending on your requirements, one or more of the following might be needed:

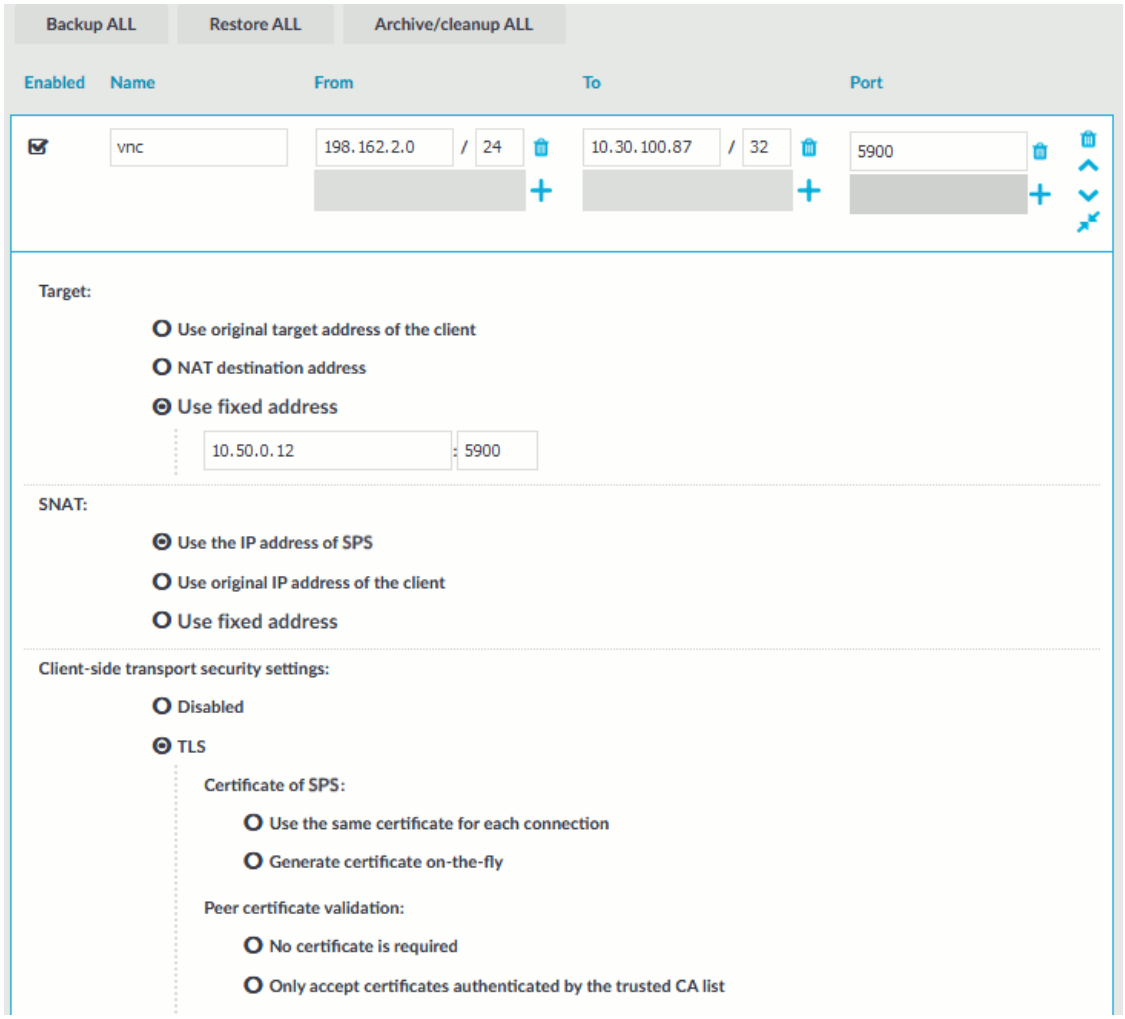
- An X.509 certificate and its private key. SPS can display the same certificate to the peers on both the client and the server side. You can also use different certificates for the client and server sides. Use your own PKI system to generate these certificates, as they cannot be created on SPS. Note that the Common Name of the certificate must contain the domain name or the IP address of SPS, otherwise the clients might reject the certificate.
- To generate certificates on-the-fly for a connection, a signing certificate authority is required. For details on creating a signing CA, see Signing certificates on-the-fly on page 529.
- To require the peers of SPS to have an X.509 certificate signed by a specific Certificate Authority, a list of the trusted certificate authorities is needed. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To enable TLS-encryption in a VNC connection policy

1. Navigate to **Traffic Controls** > **VNC** > **Connections** and select the connection policy in which you want to enable TLS.

Figure 243: Traffic Controls > VNC > Connections — Enabling TLS-encryption for VNC connections



2. Set the encryption settings used between the client and SPS in the **Client-side transport security settings** section.

   To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.

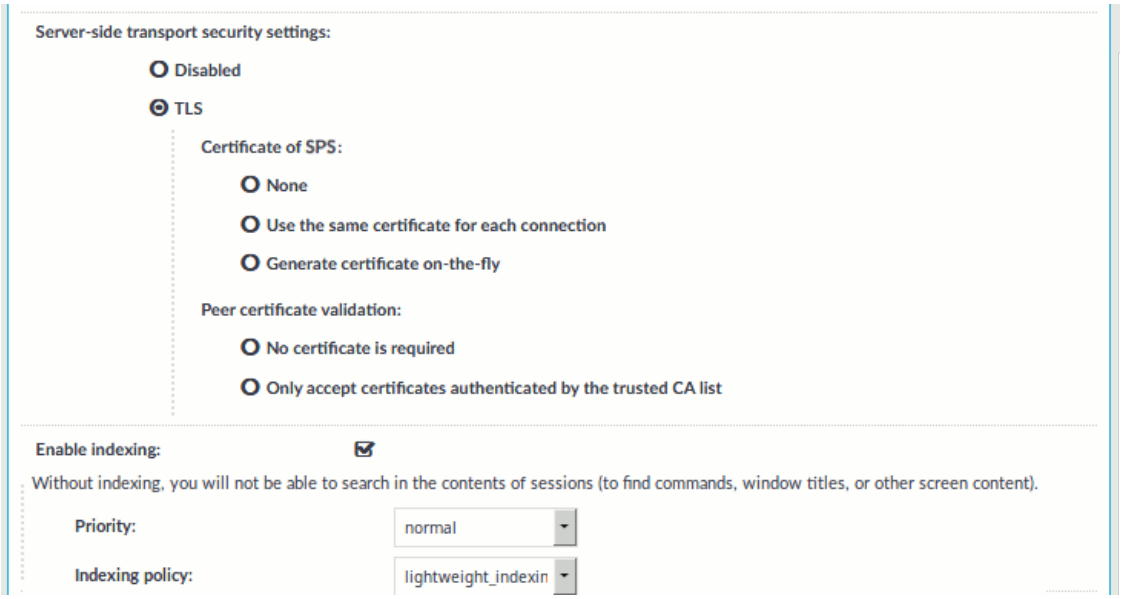
3.
   - To use the same certificate for every peer, complete the following steps.
     1. Generate and sign a certificate for One Identity Safeguard for Privileged Sessions (SPS) in your PKI system, then export the certificate and its private key.
     2. Select **Use the same certificate for each connection**.
     3. Select **Private key for host certificate**, click and upload the

private key.

4. Select **X.509 host certificate**, click  and upload the certificate.

- To use a separate certificate for every connection, complete the following steps.
  1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
  2. Select **Generate certificate on-the-fly**.
  3. Select the certificate authority to use in the **Signing CA** field.

4. Select how SPS should authenticate the peers.
   - To permit connections from peers without requesting a certificate, select **No certificate is required**.
   - To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
     1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.
     2. Select **Only accept certificates authenticated by the trusted CA list**.
     3. Select the certificate authority list to use in the **Trusted CA** field.
5. Set the encryption settings used between SPS and the server in the **Server-side transport security settings** section.



To require encryption, select **TLS**. When the connection is encrypted, SPS has to show a certificate to the peer.

6. Select the certificate to show to the server.
    - If the server does not require a certificate from SPS, select **None**.
    - To use the same certificate for every peer, complete the following steps.
        1. Generate and sign a certificate for SPS in your PKI system, and export the certificate and its private key.
        2. Select **Use the same certificate for each connection**.
        3. Select **Private key for host certificate**, click  and upload the private key.
        4. Select **X.509 host certificate**, click  and upload the certificate.
    - To use a separate certificate for every connection, complete the following steps.
        1. Create a certificate authority that will be used to sign the certificates that SPS shows to the peer. For details, see Signing certificates on-the-fly on page 529.
        2. Select **Generate certificate on-the-fly**.
        3. Select the certificate authority to use in the **Signing CA** field.

Limitations

NOTE: When using the **Use the same certificate for each connection** option and the connection policy that allows access to multiple servers using HTTPS, the client applications will display a warning because the certificate used in the connection will be invalid (namely, the Common Name of the certificate will not match the hostname or IP address of the server).

NOTE: Import the certificate of the signing Certificate Authority to your clients. Otherwise, the client applications will display a warning due to the unknown Certificate Authority.

7. Select how SPS should authenticate the peers.
    - To permit connections from peers without requesting a certificate, select **No certificate is required**.
    - To permit connections only from peers with a valid certificate that was signed by a specific CA, complete the following steps.
        1. Create a list of trusted Certificate Authorities that will be used to validate the certificates of the peers. For details on creating a trusted CA list, see Verifying certificates with Certificate Authorities on page 521.

2. Select **Only accept certificates authenticated by the trusted CA list**.
3. Select the certificate authority list to use in the **Trusted CA** field.

8. Click **Commit**.

Expected result

The encryption settings are applied to the connection policy.

Creating and editing protocol-level VNC settings

VNC settings determine the parameters of the connection on the protocol level, including timeout value, and so on.

CAUTION:

Modifying the VNC settings is recommended only to advanced users. Do not modify these settings unless you exactly know what you are doing.

To create a new VNC settings profile or edit an existing one

1. Navigate to **Traffic Controls** > **VNC** > **Settings** and click  to create a VNC setting profile. Enter a name for the profile (for example **`vnc_special`**).
2. Click  to display the parameters of the connection.
3. Modify the parameters as needed. The following parameters are available:
   - **Network idle timeout**: Connection timeout value in seconds. To avoid early timeout, set it to a larger value, for example a week (604800 seconds).

     Even if the user is not active, the session can contain activity that must be audited (for example, the output of a script). The idle timeout period will start only after this activity has stopped.

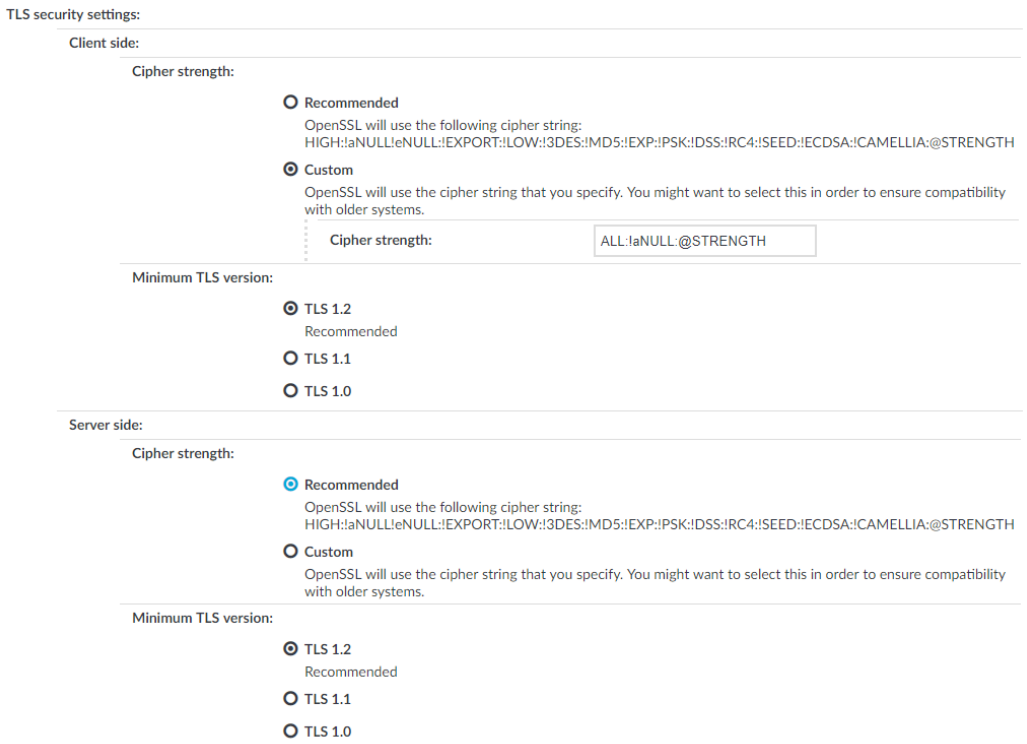
     CAUTION:

     Determining if a connection is idle is based on the network traffic generated by the connection, not the activity of the user. For example, if an application or the taskbar of a graphical desktop displays the time which is updated every minute, it generates network traffic every minute, negating the effects of timeout values greater than one minute and preventing One Identity Safeguard for Privileged Sessions (SPS) from closing the connection.
   - **User idle timeout**: If no user activity is detected, terminate the session after the configured time has passed since the last user activity.

This can be useful if only user-generated network traffic is important in a session. By using this option, situations described in the caution of **Network idle timeout** (such as a taskbar clock keeping the network traffic open indefinitely) can be avoided. To enable user idle timeout, select **Enable user idle timeout** and enter a value that is greater than or equal to the value of **Network idle timeout**.

- **Enable pre channel check**: Select this option to evaluate the connection and channel policies before establishing the server-side connection. That way if the connection is not permitted at all, SPS does not establish the server-side connection.
- To configure TLS security settings on both the **Client side** and the **Server side**, proceed to **TLS security settings**.

Figure 244: <Protocol> Control > Settings > TLS security settings - configuring TLS security settings



  - **Cipher strength** specifies the cipher string OpenSSL will use. The following options are possible:
    - **Recommended**: this setting only uses ciphers with adequate security level.
    - **Custom**: this setting allows you to specify the list of ciphers you want to permit SPS to use in the connection. This setting is only recommended to ensure compatibility with older systems. For

more details on customizing this list, check the 'openssl-ciphers' manual page on your SPS appliance.

For example: `ALL:!aNULL:@STRENGTH`

- **Minimum TLS version** specifies the minimal TLS version SPS will offer during negotiation. The following options are possible:
  - **TLS 1.2**: this setting only offers TLS version 1.2 during the negotiation. This is the recommended setting.
  - **TLS 1.1**: this setting offers TLS version 1.1 and later versions during the negotiation.
  - **TLS 1.0**: this setting offers TLS version 1.0 and later versions during the negotiation.

  NOTE: Setting up sessions to legacy systems that do not support at least TLS 1.2 is only possible when the security level of the connection is degraded to 0, which is possible by specifying the TLS ciphers manually and appending the string `:@SECLEVEL=0` to the cipher list. However, this setting also enables the use of known vulnerable algorithms and key sizes, therefore it is absolutely critical to only use such connection settings when it is necessary and when you can fully trust your network between SPS and the legacy system. It is strongly recommended to use different security settings on the server and the client side of the connection, when degrading the security level of a connection is unavoidable.

NOTE: Note that SPS only permits TLS-encrypted connections. SSLv3 is not supported.

4. Click **Commit**.
5. Select this settings profile in the **VNC settings** field of your connections.

Indexing audit trails

One Identity Safeguard for Privileged Sessions (SPS) can index the contents of audit trails using its own indexer service or external indexers. Indexing extracts the text from the audit trails and segments it to tokens. A token is a segment of the text that does not contain whitespace: for example words, dates (`2009-03-14`), MAC or IP addresses, and so on. The indexer returns the extracted tokens to SPS, which builds a comprehensive index from the tokens of the processed audit trails.

Once indexed, the contents of the audit trails can be searched from the web interface. SPS can extract the commands typed and the texts seen by the user in terminal sessions, and text from graphical protocols like RDP, Citrix ICA, and VNC. Window titles are also detected.

SPS has an internal indexer, which runs on the SPS appliance. In addition to the internal indexer, external indexers can run on Linux hosts.

Processing and indexing audit trails requires significant computing resources. If you have to audit lots of connections, or have a large number of custom reports configured, consider using an external indexer to decrease the load on SPS. For sizing recommendations, ask your One Identity partner or contact our Support Team.

- *The internal indexer service* runs on the SPS appliance. It supports languages based on the Latin-, Greek- and Cyrillic alphabets, as well as Chinese, Japanese and Korean languages, allowing it to recognize texts from graphical audit trails in 100+ languages. It can also generate screenshots for content search results.
- *The external indexer* runs on Linux hosts and instances. It uses the same engine as the indexer service of SPS, and has the same capabilities and limitations.

  SPS can work with multiple external indexers to process audit trails.

NOTE: The version of the external indexer must be equal to or greater than the version of One Identity Safeguard for Privileged Sessions (SPS). To make sure you meet this criterion, One Identity recommends that you always upgrade your external indexer when you upgrade SPS. You can check that SPS has established a connection to the external indexer on the **Indexer** > **Worker status** page of the SPS web interface.

NOTE: If a text is displayed for less than 1 second, it is not indexed.

If you have indexed trails, the index is archived every 30 days.

CAUTION:

Hazard of data loss! Make sure you also back up your data besides archiving it.

For more information, see Data and configuration backups on page 145.

If a system crash occurs, you can lose up to 30 days of index, since the index is only archived every 30 days.

- To configure SPS to index the entire content of the audited connections, complete Configuring the internal indexer on page 675.

  Indexing also needs to be enabled in the connection policy of the monitored connections.
- To configure external indexers, complete Configuring external indexers on page 681.
- To monitor the status of the servers indexing the audit trails, see Monitoring the status of the indexer services on page 702.
- To create custom reports from the contents of the audit trails, complete Creating reports from audit trail content on page 896.

Reindex audit trails

In certain cases, reindexing already indexed audit trails might be necessary, for example, if the audit trails were indexed without full screen content but you still need to search in the screen content. In this case, the audit trails can be reindexed with a different indexer configuration to perform screen content extraction. For more information, contact our Support Team.

Regenerate content stored in lucene indices

Reindexing lucene indices enables you to use reindexed indices with the officially supported search database of your choice.

NOTE: Reindexing can be very time consuming. To avoid unwanted reindexing, see the following examples below for setting time and initiating dry start.

Prerequisites

- Sessions and audit trails are still available (were not cleaned up).
- Make sure you have SPS version 6.0 (or later).

To reindex lucene indices

1. Create and run the following Python script from the core shell.

```
#!/usr/bin/env python
from datetime import datetime
import os.path
content_store_changed_filename="/opt/scb/var/upgrade/content-store-changed-time/-
content_store_changed_timestamp"
if not os.path.isfile(content_store_changed_filename):
 print("The "+content_store_changed_filename+" file does not exist. This file
contains the information when lucene storage was updated to search database.")
 exit()
f=open(content_store_changed_filename)
s=float(f.read())
f.close()
my_date =  datetime.fromtimestamp(s)
print("Run this command to reindex everything before lucene->search database
update:")
print("indexerctl reindex query \"*\" --end "+my_date.isoformat())
```

   You will receive a result similar to this:

   Example

   ```
   indexerctl reindex query "*" --end 2021-05-22T20:21:05.604341
   ```

   The example shows that all items will be reindexed until 2021-05-22 20:21:05.

2. You can set the start time parameter and check the list of items scheduled for reindexing:

   Example

   ```
   indexerctl reindex query "*" --end 2021-05-22T20:21:05.604341 --dry --
   start 2021-05-22T18:21:05.604341
   ```

   The example shows that all indices will be reindexed between 2021-05-22 18:21:05 and 2021-05-22 20:21:05.

3. If you are satisfied with the results, start reindexing:

Example

```
indexerctl reindex query "*" --end 2021-05-22T20:21:05.604341 --start
2021-05-22T18:21:05.604341
```

Configuring the internal indexer

This section describes how to configure One Identity Safeguard for Privileged Sessions (SPS) to index the audit trails.

Indexing is a resource intensive (CPU and hard disk) operation, and depending on the number of processed audit trails and parallel connections passing SPS, may affect the performance of SPS. Test it thoroughly before enabling it in a production environment that is under heavy load. If your SPS appliance cannot handle the connections and the indexing, consider using external indexers (see *Configuring external indexers* in the *Administration Guide*) to decrease the load on SPS. For sizing recommendations, ask your One Identity partner or contact our Support Team.

Note that the minimum value of **Backup & Archive** > **Archive policies** > **Delete data from SPS after** is `30` days when using the indexer service. If you previously had a setting lower than this, it will still archive the index after `30` days when the indexer service is used.

NOTE: Only those audit trails will be processed that were created after full-text indexing had been configured for the connection policy. It is not possible to process already existing audit trails.

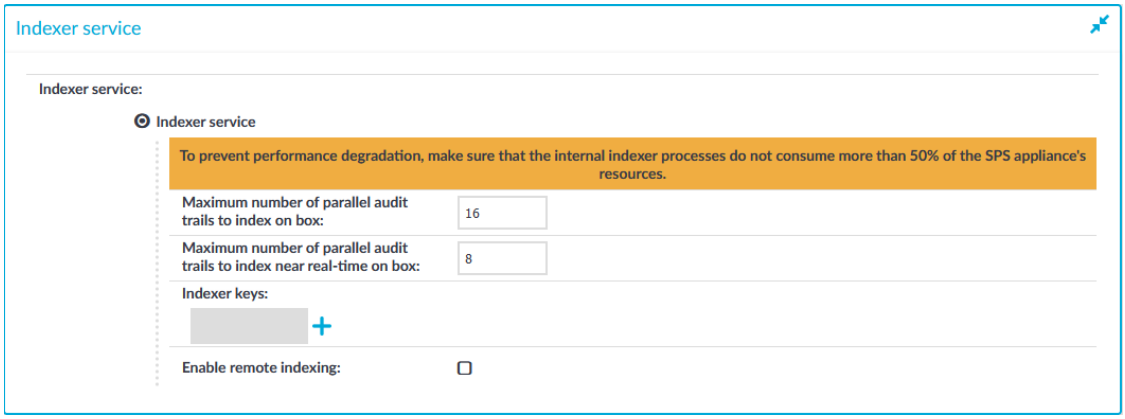
NOTE: Using content policies significantly slows down connections (approximately 5 times slower), and can also cause performance problems when using the indexer service.

NOTE: The version of the external indexer must be equal to or greater than the version of One Identity Safeguard for Privileged Sessions (SPS). To make sure you meet this criterion, One Identity recommends that you always upgrade your external indexer when you upgrade SPS. You can check that SPS has established a connection to the external indexer on the **Indexer** > **Worker status** page of the SPS web interface.

To configure SPS to index the audit trails

1. Navigate to **Basic Settings** > **Local Services** > **Indexer service**.

Figure 245: Basic Settings > Local Services > Indexer service > Configure the Indexer service of SPS



2. Define the **Maximum number of parallel audit trails to index on box**.

   This option determines the maximum number of parallel indexing tasks that the SPS appliance performs. The default value is set to the number of detected CPU cores. Note that indexing audit trails requires about 50-100 Mbytes of memory for terminal sessions (SSH, Telnet, TN3270), and 150-300 Mbytes for graphical sessions (RDP, ICA, VNC, X11). Consider the memory usage of your SPS host before modifying this value.

3. Define the **Maximum number of parallel audit trails to index near real-time on box**.

   This option determines the maximum number of parallel indexing tasks that the SPS appliance performs near real-time, meaning that indexing starts when sessions are still ongoing. The default value is set to 0.

   NOTE: A connection policy configured with near real-time priority (**Connection policy** > **Enable indexing** > **Priority**) requires that you set **Maximum number of parallel audit trails to index near real-time on box** to a value other than `0`.

4. (Optional) If you have encrypted audit trails and you want to index them, upload the necessary RSA private keys.

   Click  , and then click the  icon to upload a private key. A pop-up window is displayed.

   Select **Browse**, select the file containing the private key, provide the **Password** if the key is password-protected, and click **Upload**. Alternatively, you can also copy-paste the private key into the **Key** field, provide the **Password** there, and click **Set**.

   TIP: If you want to search in the trail content on the web interface: to view screen-shots generated from encrypted audit trails, you also have to upload the necessary

private encryption keys to your audit keystore. For more information, see Audit keystore.

5. Click 

.
6. Navigate to **Policies** > **Indexer Policies**.
7. Two Indexer Policies are available by default, both with automatic language detection:
   - `full_indexing`: Slower, indexes the complete content of the screen, including all events.
   - `lightweight_indexing`: Significantly faster, but it extracts only the executed commands (Command event) and the window titles (Window title event) that appear on the screen. It does not index any other screen content (for example, text that is displayed in a terminal or that appears in an RDP window).

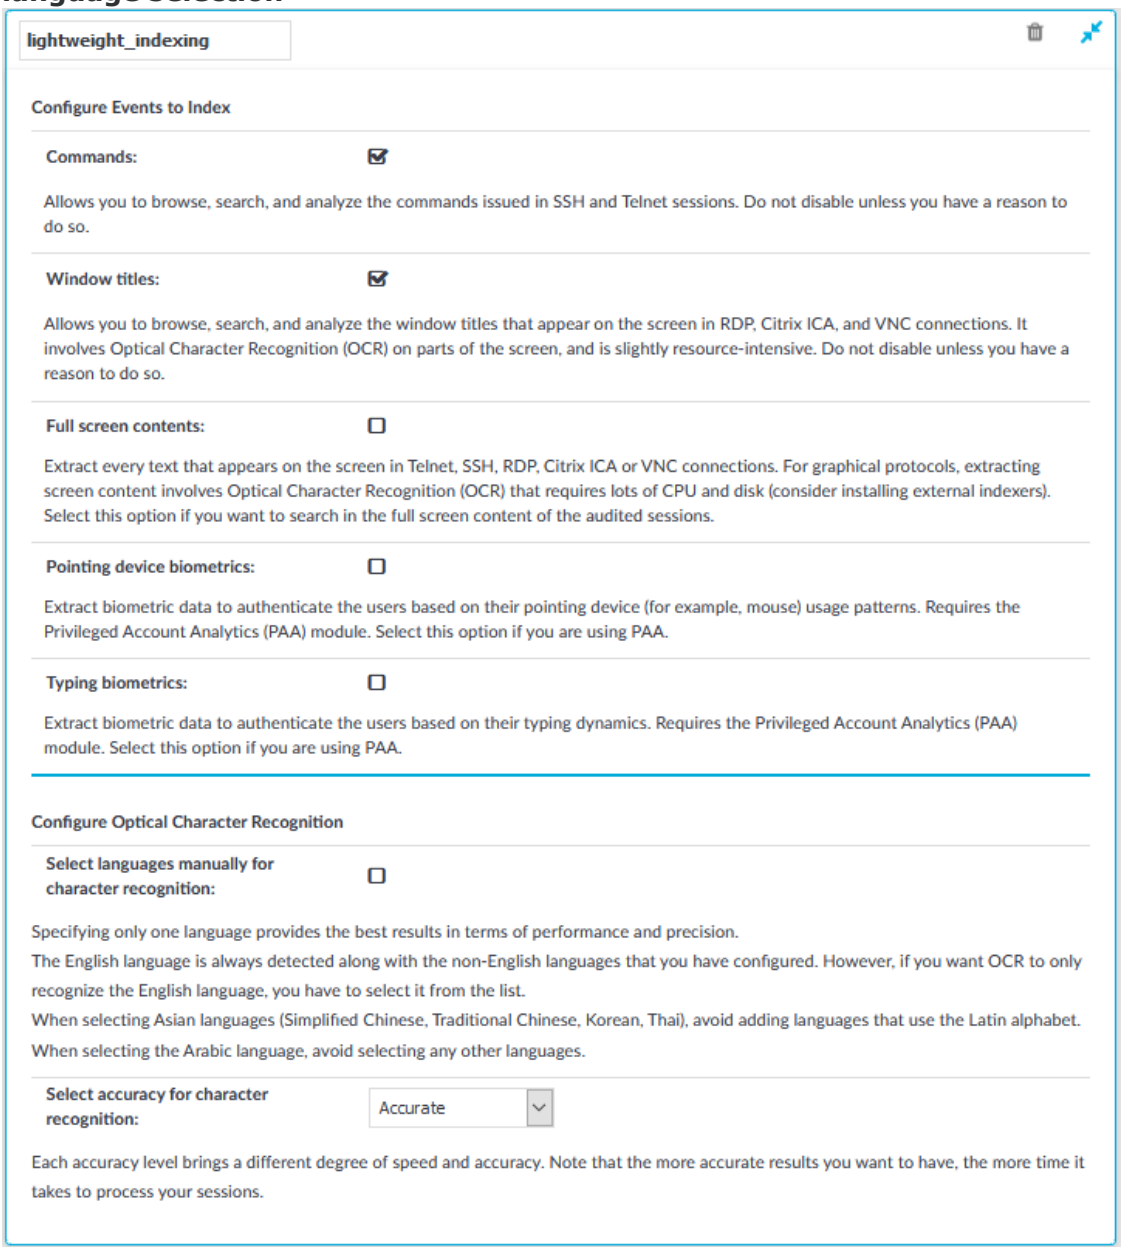
     For example, in the case of an SSH protocol, `lightweight_indexing` will index a command with parameters, such as `cat --help`, but will not index terminal printouts such as the help content itself.

     When you add a new Connection Policy, the `lightweight_indexing` Indexer Policy is assigned to it by default.

NOTE: In the case of graphical protocols, the default Optical Character Recognition (OCR) configuration is automatic language detection. This means that the OCR engine will attempt to detect the languages of the indexed audit trails automatically. However, if you know in advance what language(s) will be used, create a new Indexer Policy.

To create a new Indexer Policy, click 
.

Figure 246: Policies > Indexer Policies > Indexing options and manual language selection



8. Select from the indexing options as follows:
    - **Commands:** Allows you to browse, search, and analyze the commands issued in SSH and Telnet sessions.

      CAUTION:

      Do not disable unless you have a reason to do so.

- **Window titles:**Text appearing as window titles that can be detected on the screen in RDP, Citrix ICA, and VNC connections. Window title detection involves Optical Character Recognition (OCR) on parts of the screen, and can be slightly resource-intensive. SPS versions up till 6.2 only detected only the active window in the screen. From SPS version 6.3, multiple windows can be detected.

  ### Limitations

  - Default Windows themes are supported.
  - Windows that do not have an X (close window) button in the top-right corner (or it is not visible) are not detected.
  - Use window title detection for sessions that use a single monitor. The feature works in multi-monitor environments as well, but becomes very slow, therefore it is not recommended.
  - Window title detection is case-insensitive.

  CAUTION:

  Do not disable unless you have a reason to do so.

- **Full screen contents:** Select this option if you want to search in the full screen content of the audited sessions.

  Extract every text that appears on the screen in Telnet, SSH, RDP, Citrix ICA or VNC connections. For graphical protocols, extracting screen content involves Optical Character Recognition (OCR) that requires lots of CPU and disk (consider installing external indexers).

- **Pointing device biometrics:** Select this option only if you are using One Identity Safeguard for Privileged Analytics (SPA)).

  Extract biometric data to authenticate the users based on their pointing device (for example, mouse) usage patterns. SPA can analyze mouse movement patterns of your users as a biometric identity verification method to protect against account theft.

- **Typing biometrics:** Select this option only if you are using One Identity Safeguard for Privileged Analytics (SPA)).

  Extract biometric data to authenticate the users based on their typing dynamics. SPA can analyze the typing patterns of your users as a biometric identity verification method to protect against account theft.

9. To configure what languages to detect, select **Select languages manually for character recognition**. Select the language(s) to detect. Note the following:
   - Specifying only one language provides the best results in terms of performance and precision.
   - The English language is always detected along with the non-English languages that you have configured. However, if you want the OCR to only recognize the

English language, you have to select it from the list of languages.

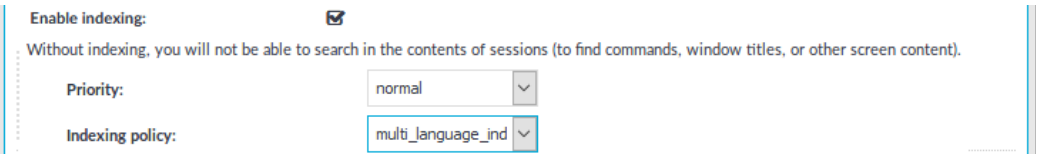
- There are certain limitations in the OCR engine when recognizing languages with very different character sets. For this reason, consider the following:
  - When selecting Asian languages (Simplified Chinese, Traditional Chinese, Korean), avoid adding languages that use the Latin alphabet.
  - When selecting the Arabic language, avoid selecting any other languages.
  - The Thai language is currently not supported. If you are interested in using SPS to index Thai texts, contact our Sales Team.

10. Specify an accuracy level for Optical Character Recognition (OCR). Each accuracy level brings a different degree of speed and accuracy:
    - **Fast**: The fastest option with potentially less accurate results. Select this option if speed is more important to you than getting the most accurate results possible.
    - **Balanced** (default setting): Fairly accurate option with less than optimum speed. Select this option if you want results to be fairly accurate but you have more than a few sessions to process and processing time is less of a concern.
    - **Accurate**: The most accurate option with less optimal speed. Select this option if you must have the most accurate results possible and speed is less important or you only have a few sessions to process.
11. Configure the Indexing policy for the Connection policy that you want to index:

    By default, the `lightweight_indexing` Indexing policy is enabled for every Connection policy with `normal` priority. If this is ideal for you, skip this step and continue with the next step. If you want to use a different policy, for example because you want to OCR the complete screen content, or because you have created a language-specific indexer policy, complete the following substeps.

    a. Navigate to the **Control** > **Connections** page of the traffic type (for example **Traffic Controls** > **SSH**), and select the connection policy to index.

    b. **Figure 247: Traffic Controls > Protocol name > Connections > Enable indexing — Select Indexing Policy**

    Select the **Indexing Policy** to be used. Both built-in Indexer Policies feature automatic language detection. To specify a particular language detection configuration, select the Indexing Policy you have created before (in Step 6).

    c. To determine the priority level of indexing this connection, select the appropriate **Priority** level. Selecting a high priority level means that the trails of this connection will be indexed first. Selecting a low priority level means that

the trails of this connection will be indexed also, but there might be a delay in indexing if there are a lot of high-priority connections waiting to be indexed. Selecting **near real-time** means that the indexing of sessions starts when sessions are still ongoing.

d. Click 

.

12. Check which channel policy is used in the connection, and navigate to the **Traffic Controls** > **Protocol name** > **Connections** page. Select the channel policy used in the connection to index.
13. On the **Traffic Controls** > **Protocol name** > **Channel Policies** page, verify that the **Record audit trail** option is selected for the channels you want to index (for example, the `Session shell` channel in SSH, or the `Drawing` channel in RDP).
14. Click 

.

TIP: To verify that indexing works as configured, start a session that uses this connection policy (connect from a client to a server).

When the session is finished, navigate to the **Indexer** > **Indexer status** page to verify that the indexer service is processing the audit trail.

If the audit trails are encrypted, ensure that the required decryption keys have been uploaded to **Basic Settings** > **Local Services** > **Indexer service** > **Indexer keys**.

Configuring external indexers

If One Identity Safeguard for Privileged Sessions (SPS) audits lots of connections, processing and indexing the created audit trails requires significant computing resources, which may not be available in the SPS appliance. To decrease the load on the SPS appliance, you can install the indexer service on external Linux hosts. These external indexer hosts run the same indexer service as the SPS appliance, and can index audit trails, or generate screenshots and replayable video files from the audit trails as needed. The external indexers register on SPS, wait for SPS to send an audit trail to process, process the audit trail, then return the processed data to SPS. The external indexer hosts do not store any data, thus any sensitive data is available on the host while it is being processed.

To use external indexers to process your audit trails, you have to complete the following steps.

- Read the conditions and limitations related to external indexers in Prerequisites and limitations on page 682.
- Install and configure the hosts (physical or virtual) that will run the external indexer service. For details on the hardware requirements, see Hardware requirements for the external indexer host on page 683.

- Configure SPS to use external indexers. For details, see Configuring One Identity Safeguard for Privileged Sessions (SPS) to use external indexers on page 683.
- Install and configure the indexer application on the external hosts. For details, see Installing the external indexerConfiguring the external indexer on page 685.
- If you enabled audit trail encrypting on SPS, you will also need to upload the necessary certificates to the external indexer to allow indexing the encrypted trails. For details, Uploading decryption keys to the external indexer on page 690.

Prerequisites and limitations

Before starting to use One Identity Safeguard for Privileged Sessions (SPS) with external indexers, consider the following:

- If there is a firewall between the host of the external indexer and SPS, enable connections from the external indexer to SPS.

  The default port is TCP/12345. To change the port number, you have to modify the indexer settings on SPS, and upload the new configuration to the external indexer(s).
- To protect the sensitive data in the audit trails, ensure that the audit trails are encrypted. For details on encrypting audit trails, see Encrypting audit trails on page 512.
- Make sure to permit indexer access only to the hosts that really run external indexers on the **Basic Settings** > **Local Services** > **Indexer service** page of the SPS web interface.
- NOTE: The current OCR engine cannot guarantee accurate character recognition for non-Latin characters smaller than 30 x 30 pixels. If you encounter problems with character recognition for non-Latin characters, increase resolution settings in your connection.
- The external indexer can be installed on the following 64-bit operating systems: Red Hat Enterprise Linux Server 8 and its derivatives, such as CentOS, Oracle Linux, AlmaLinux, Rocky Linux, etc.

  NOTE: Derivatives are supported only if an issue can be reproduced on an official RHEL distribution. Do not report issues specific to a derivative OS but not to RHEL.
- Update your system:

  ```
  yum update
  ```

  Download the External Indexer bundle from the SPS box itself:

  ```
  curl https://<SPS-IP>/external-indexer.rpm -o external-indexer.rpm
  ```

  Install the bundle:

  ```
  yum install external-indexer.rpm
  ```

If your security policy does not permit the above limitations, or your environment does not make it possible to fulfill them, do not use external indexers with SPS.

Hardware requirements for the external indexer host

NOTE: This is a data-driven part of the product. Hardware requirements and exact memory usage cannot be safely predicted as the actual memory usage depends on the contents of the sessions.

- CPU: You can configure the number of audit trails that an indexer host processes at the same time. For optimal performance, each indexer process should have a dedicated CPU core.
- Memory requirements: In addition to the memory requirements of the operating system of the host, the indexer requires about 300 MB memory for each worker process, depending on the protocol of the indexed audit trails. The audit trails of terminal connections require less memory.
- Disk: The indexer requests the data from One Identity Safeguard for Privileged Sessions (SPS) in small chunks, it does not store the entire audit trail nor any temporary files. You will need only disk space for the operating system, and a few GB to store logs.

For example, if you want to have a host that can process 6 audit trails at the same time, you need 6 CPU cores and 1.8 GB of memory for the indexer service. If you install only a minimal operating system and the external indexer on the host, 6 GB disk space should be enough.

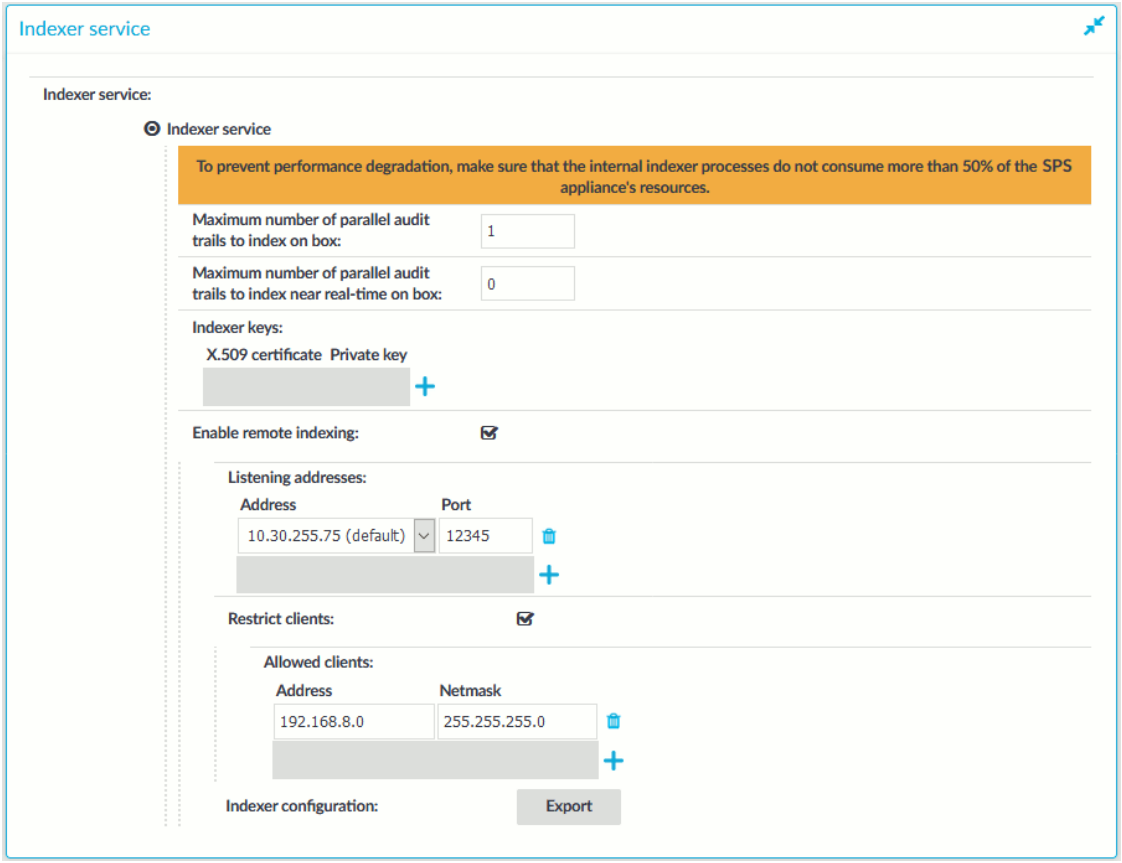
Configuring One Identity Safeguard for Privileged Sessions (SPS) to use external indexers

The following describes how to configure One Identity Safeguard for Privileged Sessions (SPS) to accept connections from external indexer services.

To configure SPS to accept connections from external indexer services

1. Log in to the SPS web interface, and navigate to **Basic Settings** > **Local Services** > **Indexer service**.
2. Select **Indexer service**.
3. Select **Enable remote indexing**.

Figure 248: Basic Settings > Local Services > Indexer service > Enable remote indexing — Configure external indexers



4. In the **Listening addresses** > **Address** field, select the network interface where SPS should accept external indexer connections. Repeat this step to add other interfaces if needed.

   The available addresses correspond to the interface addresses configured in **Basic Settings** > **Network** > **Interfaces**. Only IPv4 addresses can be selected.
5. Select **Restrict clients**, and list the IP address and netmask of your external indexer hosts.

   Use an IPv4 address.
6. Click **Commit**.

Installing the external indexer

Prerequisites

The external indexer can be installed on the following 64-bit operating systems: Red Hat Enterprise Linux Server 8 and its derivatives, such as CentOS, Oracle Linux, AlmaLinux,

Rocky Linux, etc.

NOTE: Derivatives are supported only if an issue can be reproduced on an official RHEL distribution. Do not report issues specific to a derivative OS but not to RHEL.

To install the external indexer

1. Log in as root to the host that you want to use to index your audit trails.
2. Copy the installer package to the host.

   NOTE: Due to legal reasons, installation packages of the external indexer application will be available only from the SPS web interface. After SPS versions 6.4 and 6.0.3 are released, the installation packages will be removed from our website.
3. Install the package using the package manager of the operating system. For example:

   ```
   yum install external-indexer-standalone-<version-number>.x86_64.rpm
   ```

   The installer performs the following steps automatically.

   - Unpacks the indexer files into the `/opt/external-indexer/` directory.
   - Installs the related init scripts (the `/etc/init.d/external-indexer` init script, and adds the init script configuration to `/etc/sysconfig/external-indexer`).
   - Creates the `indexer` user and usergroup. This is an unprivileged user that is used to run the indexer application.
   - Registers the service to start automatically on system boot. Note that the indexer init script uses bind mount points.
4. Configure the indexer. For details, see Configuring the external indexer on page 685.

Configuring the external indexer

To connect to One Identity Safeguard for Privileged Sessions (SPS) and index the audit trails, you must configure the external indexer.

CAUTION:

Unless you know exactly what you are doing, modify only the parameters you are instructed to.

To configure the external indexer

1. Log in to the SPS web interface, and navigate to **Basic Settings** > **Local Services** > **Indexer service**.
2. To export the configuration file for the external indexer, click **Export**.

   NOTE: The **Export** button is displayed only after the configuration that enables SPS to use remote indexers is committed.

The configuration file contains the listening address (IP and port) of SPS and the necessary keys for SSL authentication.

Upload the file to the external indexer host.

3. On the external indexer host, import the configuration file with the following command:

```
indexer-box-config <configuration-file>.config
```

4. To configure the external indexer service, open the `/etc/indexer/indexerworker.cfg` configuration file for editing.
5. To edit the number of worker groups assigned to a certain worker process type, find the `worker_groups` line.

   A worker group has the following parameters:

   - `name`: the name of the worker group
   - `count`: the number of workers to use for processing
   - `capabilities`: the types of jobs the process performs (`index`, `screenshot`, `video`, `near-realtime`)

     NOTE: Setting the `near-realtime` capability exclusively determines whether active or closed sessions are indexed. For example:

     - If you set `[near-realtime, index]` capabilities for a worker, that worker only indexes active and ongoing sessions.
     - If you set `[index, screenshot, video]` capabilities for a worker, that worker only indexes closed sessions.

     NOTE: If you configure a connection policy with near real-time priority (**Connection policy** > **Enable indexing** > **Priority**), you must configure indexer workers that are capable of near real-time indexing. To configure such indexer workers, set the `near-realtime` capability for the relevant workers.

     NOTE: Indexer workers with the `near-realtime` capability require fewer CPU cores but more memory than indexer workers that do not have this capability.

   Make sure that the sum value of the workers are equal to the number of CPU cores in the host (or the number of CPU cores minus one if you want to save resources for other tasks).

   One Identity recommends using dedicated hosts for external indexing. If the host is not dedicated exclusively to the external indexer, decrease the number of workers accordingly.
6. (Optional) To fine-tune performance, you can configure the number of OCR threads each worker can initiate using the `ocr_thread_count` key.

   The default setting is `3`. When configuring this setting, pay attention to the available CPU cores, as raising the number of possible threads too high can impact performance negatively.

7. (Optional) If instructed by One Identity Support, configure the OCR engine.

   Find the `engine` key, and change its value to one of the following options:

   - `omnipage-external` is the default setting. It provides the best performance and stability by allowing workers to initiate multiple OCR threads.

     This setting also allows you to search for images where OCR could not be performed. On the search UI of SPS, enter the `OOCCRRCCRRAASSHH` search string to list all such images. If possible, contact our Support Team, to help One Identity to further improve the engine.

     NOTE: Multiple OCR threads can only speed up processing graphical protocols (RDP, VNC, and ICA trails), and do not affect the processing speed for terminal-based protocols (telnet and SSH).

   - `omnipage` only supports one OCR thread per worker.

     If you have to use this option, make sure to also set the `ocr_thread_count` to **`0`**.

8. Save your changes.
9. (Optional) Continue with uploading decryption keys to index encrypted audit trails. For more information, see Uploading decryption keys to the external indexer on page 690.
10. (Optional) Start the indexer service. For more information, see Starting the external indexer on page 699.

Configuring a service pool

You can share your worker resources between multiple indexer services, for example behind a load-balancer, or in an SPS cluster if you want to share some of your indexer workers between multiple SPS instances.

To enable resource sharing

1. Use the `indexer-box-config` script with an additional `--service-pool` option to enumerate the service configurations.

   NOTE: Running this script can overwrite your custom modifications in your `indexerworker.cfg` file. Make a backup of `indexerworker.cfg` before running the script.

indexer-box-config

```
cp /etc/indxer/indexerworker.cfg /etc/indxer/indexerworker.cfg.bak
indexer-box-config $dedicated-server-config --service_pool $cfg1 $cfg2
$cfg3
```

If you run the script with `--service-pool` option, the script creates a `service-pool` field in your configuration file which can be used by any worker-group. The script also creates a default worker-group called `shared`, with an additional `balancing` field. These workers run in one-shot mode, and connect to an indexer-service randomly selected from the pool.

Example: output after running script with two configuration files

```
{
  "service": {
    "address": "169.254.1.1",
    "port": 12345,
    "ssl": {
      "ca_certificate": "/etc/indexer/cacert.pem",
      "certificate": "/etc/indexer/worker.pem",
      "enabled": false,
      "private_key": "/etc/indexer/worker.key"
    }
  },
  "service_pool": [
    {
      "address": "192.168.1.111",
      "port": 12345,
      "ssl": {
        "ca_certificate": "/etc/indexer/idx-external-indexer-
20220902T1208/ca.pem",
        "certificate": "/etc/indexer/idx-external-indexer-20220902T1208/-
worker.pem",
        "enabled": true,
        "private_key": "/etc/indexer/idx-external-indexer-20220902T1208/-
worker.key"
      }
    },
    {
      "address": "192.168.1.118",
      "port": 12345,
```

```
      "ssl": {
        "ca_certificate": "/etc/indexer/idx-external-indexer-
20220902T1207/ca.pem",
        "certificate": "/etc/indexer/idx-external-indexer-20220902T1207/-
worker.pem",
        "enabled": true,
        "private_key": "/etc/indexer/idx-external-indexer-20220902T1207/-
worker.key"
      }
    }
  ],
  "settings": {
    "log_level": 3,
    "ocr": {
      "engine": "omnipage-external",
      "minimal_time_distance": 1,
      "ocr_thread_count": 2
    },
    "pkcs11": {
      "custom_password": false,
      "slots": []
    },
    "terminal": {
      "extract_buffer": true
    },
    "worker_arguments": "--http-config /opt/external-index-
er/httpconfig.json"
  },
  "worker_groups": [
    {
      "balancing": false,
      "capabilities": [
        "index"
      ],
      "count": 4,
      "name": "workers"
    },
    {
      "balancing": false,
      "capabilities": [
        "screenshot",
        "video"
      ],
      "count": 1,
      "name": "screenshot-and-video"
    },
    {
      "balancing": false,
      "capabilities": [
        "video"
      ],
```

```
      "count": 1,
      "name": "video"
    },
    {
      "balancing": false,
      "capabilities": [
        "near-realtime"
      ],
      "count": 0,
      "name": "near-realtime"
    },
    {
      "balancing": true,
      "capabilities": [],
      "count": 0,
      "name": "shared"
    }
  ],
  "workercontroller": {
    "log_level": "info"
  }
}
```

These worker processes disconnect from the service when they finish processing a job, or when they do not receive a job within 60 seconds after connecting to the service.

2. Define additional worker-groups.

   You can define any number of worker-groups with different capabilities, but you can have only one dedicated service and one service pool.

3. Add `"balancing": true` to any worker-group to share those workers between the services configured in the `service-pool` field.

   Worker groups without balancing option, or balancing set to `false` connect to the dedicated service to fetch jobs.

Uploading decryption keys to the external indexer

If the audit trails you want to index are encrypted, complete the following steps to make the decryption keys available for the indexer.

To make the decryption keys available for the external indexer

1. Obtain the RSA private key and copy it to the external indexer's host.
2. Use the `indexer-keys-json` utility to transform the private key to the required JSON format. When executed, the script asks for the path to the private key, and the password of the private key. After the conversion, the password is removed.

   The utility automatically adds the private key to the `/etc/indexer/indexer-keys.cfg` keystore file. If you want to use a different keystore file, use the `--keystore` argument to specify another file. If the keystore already includes the private key you want to add, it will be ignored.

   a. In the `/opt/external-indexer/usr/bin/` folder, issue the following command: `indexer-keys-json`

   b. Provide the absolute path to the private key. Alternatively, you can include this information as a parameter: `indexer-keys-json --private-key <path-to-private-key>`

   c. If the key is password protected, enter the password to the private key.

   d. To add additional keys, re-run the `indexer-keys-json` command.
3. You can now start the indexer service. For more information, see Starting the external indexer on page 699.

Configuring a hardware security module (HSM) or smart card to integrate with external indexer

It is possible to use a hardware security module (HSM) or a smart card to store the decryption keys required for decrypting audit trails. An HSM or a smart card is a tamper-resistant physical, software, or cloud solution that can securely store digital keys used for authentication.

The main steps of configuring a hardware security module (HSM) or smart card to integrate with an external indexer are as follows:

1. Set up and test the environment.
2. Encrypt the PKCS#11 PIN.

To see examples of how to configure various HSM or smart card solutions that you wish to integrate with your external indexer(s), consult the following sections:

- Configuring SoftHSM on page 694
- Configuring AWS CloudHSM on page 696
- Configuring a smart card on page 697

Setting up and testing the environment

To access an HSM or smart card with the external indexer, a PKCS#11 shared library plugin must be used. In most cases, these libraries also need a background daemon or environment variables set. The PKCS#11 library must be accessible to the external indexer with a proper environment.

To set up the environment and test it, complete the following steps.

1. Load the environment for the indexer commands:

   ```
   source /etc/indexer/external-indexer.env
   ```

2. Test your environment.
   - *Option #1: Use the pkcs11-tool to test your environment:*
     1. List the available slots.

        ```
        pkcs11-tool --modul <path-to-pkcs11-library> -L
        ```

     2. List the objects in a slot.

        ```
        pkcs11-tool --modul <path-to-pkcs11-library> -l --slot <id> -O
        ```

   - *Option #2: Use the indexerworker with the log level set to dump to see the available keys:*

     ```
     indexerworker -l -v 7 --pkcs11-lib <path-to-pkcs11-library> --pkcs11-slot-id <id> --pkcs11-pin <pin>
     ```

3. Assuming that the environment is ready, the external indexer must be configured to use the PKCS#11 library. To do so, edit `/etc/indexer/indexerworker.cfg` as follows:

   ```
   ...
   "settings": {
      "pkcs11": {
                 "custom_password": false
                 "slots": [
                    {
                       "library": "<path-to-pkcs11-library>",
                       "slot_id": <slot-number>,
                       "pin": "<your-encrypted-PIN>"
                    }
                 ]
       }
   }
   ...
   ```

Encrypting a PKCS#11 PIN

The PKCS#11 PIN(s) must be protected by additional encryption. The indexerconfigcrypter tool must be used to encrypt the PIN(s).

To encrypt the PIN(s)

1. Encrypt the PIN.

   The PINs can be encrypted with a custom passphrase or a default one is used if no custom passphrase is provided. A custom passphrase is more secure, but interaction is needed to start or restart the external-indexer service. Using a custom passphrase is supported on hosts running CentOS 7 or later.

   Issue either of the following commands:

   - Using a default password (CentOS 6 or 7): `indexerconfigcrypter --input <your-PIN>`
   - Using a custom password (CentOS 7 or later): `indexerconfigcrypter --input <your-PIN> --password`

   It is possible to configure multiple slots. In that case, the PINs must be encrypted using the same passphrase.

2. Update the `"pkcs11"` object in the `indexerworker.cfg` file.

   The encrypted PINs must be stored in the `"pin"` field of the configuration file (in the example, a SoftHSM is used):

```
...
     "pkcs11": {
           "custom_password": true
           "slots": [
              {
                 "library": "/usr/lib/softhsm/libsofthsm.so",
                 "slot_id": 0,
                 "pin": "<your-encrypted-PIN>"
              }
           ]
     }
...
```

Starting and restarting the external-indexer service when using a custom password for PKCS#11 PIN encryption

When you choose to encrypt the PKCS#11 PIN(s) using a custom password, on starting or restarting the external-indexer service, you are asked to enter your password using a special tool.

To provide your password using the required tool

1. Start the external-indexer service:

   ```
   systemctl start external-indexer
   ```

2. The external-indexer service prompts you to provide a password using the systemd-ask-password tool. Issue:

   ```
   systemd-tty-ask-password-agent
   ```

3. Provide the password at the prompt. You can use multiple agents to enter the password.
4. Once the external indexer(s) have been started or restarted, make sure that all the indexers have started up successfully.

   For example, on CentOS 7, you can use:

   ```
   systemctl status external-indexer
   ```

Configuring SoftHSM

SoftHSM is the software implementation of an HSM. It can be installed from the EPEL repository. The configuration of SoftHSM can be found at `/etc/softhsm2.conf` (CentOS 7), or `/etc/softhsm.conf` (CentOS 6).

The following describes how to configure SoftHSM.

NOTE: Depending on the exact SoftHSM solution that you are using, the steps described here may slightly differ.

NOTE: The following steps assume that:

- You are on the host operating system.
- The external indexer has been installed.

Prerequisites

The indexer user/group has the rights to read the data directory of SoftHSM and its contents, which defaults to `/var/lib/softhsm`.

To configure SoftHSM

1. Initialize directories for SoftHSM.

   ```
   mkdir -p /var/lib/softhsm
   chgrp -R indexer /var/lib/softhsm
   ```

2. Configure slots for softhsm1 (CentOS 6). For softhsm2 (CentOS 7), you can skip this step.

```
cat /etc/softhsm.conf
0:/var/lib/softhsm/slot0.db
1:/var/lib/softhsm/slot1.db
```

3. Initialize slot 0 (softhsm1).

```
softhsm --init-token --slot 0 --label "<your-slot-label>" --<so-pin>
topsecret --pin <your-SoftHSM-PIN>
```

4. Initialize a new slot (softhsm2) and get the slot ID:

```
softhsm2-util --init-token --free --label "<your-slot-label>" --<so-pin>
topsecret --pin <your-SoftHSM-PIN>
SLOT_ID=$(softhsm2-util --show-slots | grep -B 15 "<your-slot-label>" |
grep "Slot [0-9]" | head -n 1 | cut -d ' ' -f 2)
```

5. Import your keys. Your keys must be in the `.der` format.

   For softhsm1, use:

```
pkcs11-tool --module /usr/lib/softhsm/libsofthsm.so -l -y privkey --slot 0
-w key.der -d 001 -a <your-key-label> --pin <your-SoftHSM-PIN>
```

   For softhsm2, use:

```
pkcs11-tool --module /usr/lib/softhsm/libsofthsm2.so -l -y privkey --slot
0 -w key.der -d 001 -a <your-key-label> --pin <your-SoftHSM-PIN>
```

6. Make sure that the indexer user/group has execute right to the token directory and read right to the token files below the `/var/lib/softhsm/tokens/` directory.
7. Test your SoftHSM configuration with the indexer.

```
source /etc/indexer/external-indexer.env
indexerworker -l -v 7 --pkcs11-lib "<your-SoftHSM-library>" --pkcs11-slot-
id 0 --pkcs11-pin "<your-SoftHSM-PIN>"
```

8. Encrypt the PKCS#11 PIN(s). For detailed instructions, see Encrypting a PKCS#11 PIN on page 693.
9. Update the `"pkcs11"` object in the `/etc/indexer/indexerworker.cfg` file.

```
...
   "pkcs11": {
       "slots": [
         {
           "library": "/usr/lib/softhsm/libsofthsm.so",
```

```
          "slot_id": 0,
          "pin": "<your-encrypted-PIN>"
        }
      ]
  }
...
```

Configuring AWS CloudHSM

Amazon Web Services (AWS) CloudHSM provides hardware security modules in the AWS Cloud.

The following describes how to configure CloudHSM.

NOTE: The following steps assume that:

- You have set up your AWS CloudHSM, that is, you have created a user for the indexer, imported/generated keys, and so on.

  For detailed information on AWS CloudHSM, see the AWS CloudHSM User Guide.
- The CloudHSM PKCS#11 library is installed.
- The external indexer has been installed.

To configure CloudHSM

1. Test your environment as described in Setting up and testing the environment on page 692.

   Note that you will need to provide your CloudHSM PIN in the following format:

   ```
   "<your-CloudHSM-username:your-CloudHSM-PIN>"
   ```

2. Encrypt the PKCS#11 PIN(s). For detailed instructions, see Encrypting a PKCS#11 PIN on page 693.
3. Update the `"pkcs11"` object in the `/etc/indexer/indexerworker.cfg` file.

   ```
   ...
        "pkcs11": {
              "custom_password": true
              "slots": [
                 {
                    "library": "/opt/cloudhsm/lib/libcloudhsm_pkcs11.so",
                    "slot_id": 1,
                    "pin": "<your-encrypted-PIN>"
                 }
              ]
        }
   ...
   ```

Configuring a smart card

NOTE: Using the external indexer with a smart card is currently an experimental feature only.

To configure a smart card

1. Install OpenSC, for example, from the EPEL repository of CentOS.
2. Ensure that the PC/SC Smart Card Daemon (pcscd) service is running:
   - On CentOS 6:

     ```
     service pcscd start
     ```

   - On CentOS 7:

     ```
     systemctl enable pcscd
     systemctl start pcscd
     ```

     Alternatively, you can use:

     ```
     systemctl enable pcscd.socket
     systemctl start pcscd.socket
     ```

     This ensures that the pcscd service will not start at system startup, it will only start when there is an attempt (for example, by the indexerworker) to connect to it.
3. Test your environment as described in Setting up and testing the environment on page 692.
4. Encrypt the PKCS#11 PIN(s). For detailed instructions, see the Encrypting a PKCS#11 PIN on page 693.
5. Update the `"pkcs11"` object in the `/etc/indexer/indexerworker.cfg` file, for example:

   ```
   ...
        "pkcs11":  {
              "slots":  [
                 {
                    "library":  "/usr/lib64/pkcs11/opensc-pkcs11.so",
                    "slot_id":  1,
                    "pin":  "encrypted_pin"
                 }
              ]
        }
   ...
   ```

Customizing the indexing of HTTP traffic

Use this section to customize how One Identity Safeguard for Privileged Sessions (SPS) indexes HTTP traffic.

Prerequisites

You can customize only the configuration of external indexers. The indexer running on the SPS host always uses the default HTTP configuration, which is the following:

```
{
    "General": {
       "Whitelist": ["text/.*", ".*json.*", "application/x-www-form-
urlencoded", "multipart/.*"],
       "Blacklist": ["text/css", "application/javascript", "text/xslt",
".*xml.*"]
    },
    "Form": {
         "Blacklist": ["password", "pass"]
    },
    "Html": {
       "Attributes": ["href", "name", "value", "title", "id", "src"],
       "StrippedTags": ["script", "object", "style", "noscript", "embed",
"video", "audio", "canvas", "svg"]
    }
 }
```

To customize how SPS indexes HTTP traffic

1. Create a configuration file for the HTTP indexer using a text editor. The configuration file uses the JSON format. For details on the configuration format, see HTTP indexer configuration format on page 706.

   NOTE: If you want to index HTTP POST messages, include the `"application/x-www-form-urlencoded"` Content-Type in the `General > WhiteList` list. The indexer will decode URL encoding (percentage encoding), and create `key=value` pairs from the form fields and their values. Note that in the values, the indexer will replace whitespace with the underscore (_) character. To avoid indexing sensitive information (for example, passwords from login forms), use the `Form > Blacklist` option.

2. Copy the configuration file to the external hosts, to the `/opt/external-indexer/usr/share/adp/httpconfig.json` file.
3. Reload the indexer service: `systemctl restart external-indexer.service`
4. Repeat the above steps for your other external indexer hosts. Otherwise, it is possible that certain audit trails will be indexed using different indexer configuration.
5. Disable the indexer that is running on the SPS host. Otherwise, it is possible that certain audit trails will be indexed using different indexer configuration.

Navigate to **Basic Settings** > **Local Services** > **Indexer service**, and set the **Maximum parallel audit trails to index on box** option to `0`.

Starting the external indexer

When you have configured the external indexer, and added all decryption keys, you can start running the service.

To start the external indexer

1. Start the indexer service using the following command.

   ```
   systemctl start external-indexer.service
   ```

2. Verify that the indexer service is running. Execute the `ps aux` command. In the output, you should see a `workercontroller` and one or more `indexerworker` processes. The number of the `indexerworker` processes should be the same number you set for the `number_of_workers` key of the `/etc/indexer/indexerworker.cfg` file.
3. Verify the `indexer-certs.cfg` configuration file.

   Check the system logs of the host of the external indexer. The "Error loading key store" log indicates that there was a problem with the `indexer-certs.cfg` configuration file.
4. Verify that the indexer host is displayed in the list of indexers on the **Indexer** > **Worker status** page of the One Identity Safeguard for Privileged Sessions (SPS) web interface.

Disabling indexing on One Identity Safeguard for Privileged Sessions (SPS)

To reduce load on One Identity Safeguard for Privileged Sessions (SPS), you can disable indexing audit trails on the box. Note that this introduces delays when generating on-demand screenshots for audit trail searches.

Prerequisites

Disabling indexing on the SPS box works only if an external indexer is available. If SPS cannot detect the presence of an external indexer (for example, because of a network outage), indexing is re-enabled on SPS automatically with one indexing process.

To disable indexing on SPS

1. On the SPS web interface, navigate to **Basic Settings** > **Local Services** > **Indexer service**.
2. Set the **Maximum parallel audit trails to index on box** to `0`.
3. Click 

.

Managing the indexers

The indexers that run on an external host send log messages into the standard syslog of the external host. These log messages are not visible on One Identity Safeguard for Privileged Sessions (SPS).

The indexers use the standard `init.d` framework of the host. You can restart the indexer processes using the `/etc/init.d/indexerworker restart` command, and the entire indexer service using the `/etc/init.d/external-indexer restart` command. Note that restarting the indexer service automatically restarts the worker processes as well.

The hosts that are running indexers should be visible in the list of indexers on the **Indexer** > **Worker status** page of the SPS web interface.

Upgrading the external indexer

This section describes how to upgrade the indexer application on your external indexer hosts.

CAUTION:

After SPS 6.5, CentOS 6 operating systems will not be supported for external indexers. This means that after upgrading to SPS 6.5, or the LTS maintanance release in that cadence, you will not be able to use your external indexers that are running on CentOS 6. Make sure that you prepare your affected systems for this change and upgrade to CentOS 7 or later.

NOTE: The version of the external indexer must be equal to or greater than the version of One Identity Safeguard for Privileged Sessions (SPS). To make sure you meet this criterion, One Identity recommends that you always upgrade your external indexer when you upgrade SPS. You can check that SPS has established a connection to the external indexer on the **Indexer** > **Worker status** page of the SPS web interface.

Prerequisites

Before you start, create a backup copy of the `/etc/indexer/indexerworker.cfg` and `/etc/indexer/indexer-certs.cfg` indexer configuration files. After SPS 6.13, the

`/etc/indexer/indexer-certs.cfg` indexer configuration file is automatically renamed to `/etc/indexer/indexer-keys.cfg`.

To upgrade the indexer application on your external indexer hosts

1. Download the latest indexer `.rpm` package from the **Basic Settings > Local Services > Indexer service** page of the SPS web interface.

   NOTE: Due to legal reasons, installation packages of the external indexer application will be available only from the SPS web interface. After SPS versions 6.4 and 6.0.3 are released, the installation packages will be removed from our website.
2. Copy the downloaded `.rpm` package to your external indexer hosts.
3. Stop the indexer by using the following command.
   - On Red Hat or CentOS 6.5:

     ```
     service external-indexer stop
     ```
   - On Red Hat or CentOS 7:

     ```
     systemctl stop external-indexer.service
     ```
4. Execute the following command: `yum upgrade -y indexer.rpm`
5. Resolve any warnings displayed during the upgrade process.
6. Restart the indexer by using the following command.
   - On Red Hat or CentOS 6.5:

     ```
     service external-indexer start
     ```
   - On Red Hat or CentOS 7:

     ```
     systemctl start external-indexer.service
     ```
7. Repeat this procedure on every indexer host.

Troubleshooting external indexers

The indexers that run on an external host send log messages into the standard syslog of the external host. These log messages are not visible on One Identity Safeguard for Privileged Sessions (SPS). If a problem occurs, check the logs of SPS and the external indexer to find out which component on which host causes the problem. If the problem is on the external indexer host, verify that the required decryption keys are available on the host, then restart the indexer service using the following command.

```
systemctl restart external-indexer.service
```

If the problem persists, contact our Support Team. You can increase the log level of the indexer processes from the configuration file.

Monitoring the status of the indexer services

You can monitor the status of your indexer services in a summarized view by navigating to the **Indexer Status** page of the **Main Menu**.

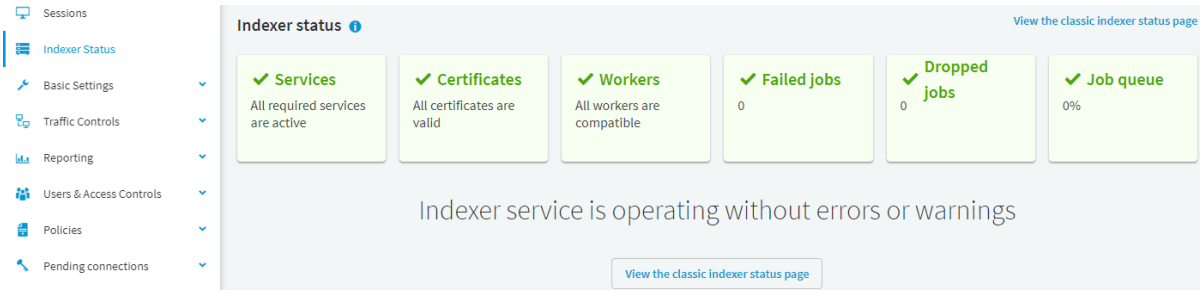
> TIP: To view the status of your indexer services in classic view, click the **View the classic indexer status page** link in the upper right corner, or the **View the classic indexer status page** button, if visible.
>
> For more information, see Monitoring the status of the indexer services in classic view.

The **Indexer Status** page displays the overall health of your indexer services, summarizing the current state of your:

- Services
- Certificates
- Workers
- Failed jobs
- Dropped jobs
- Job queue

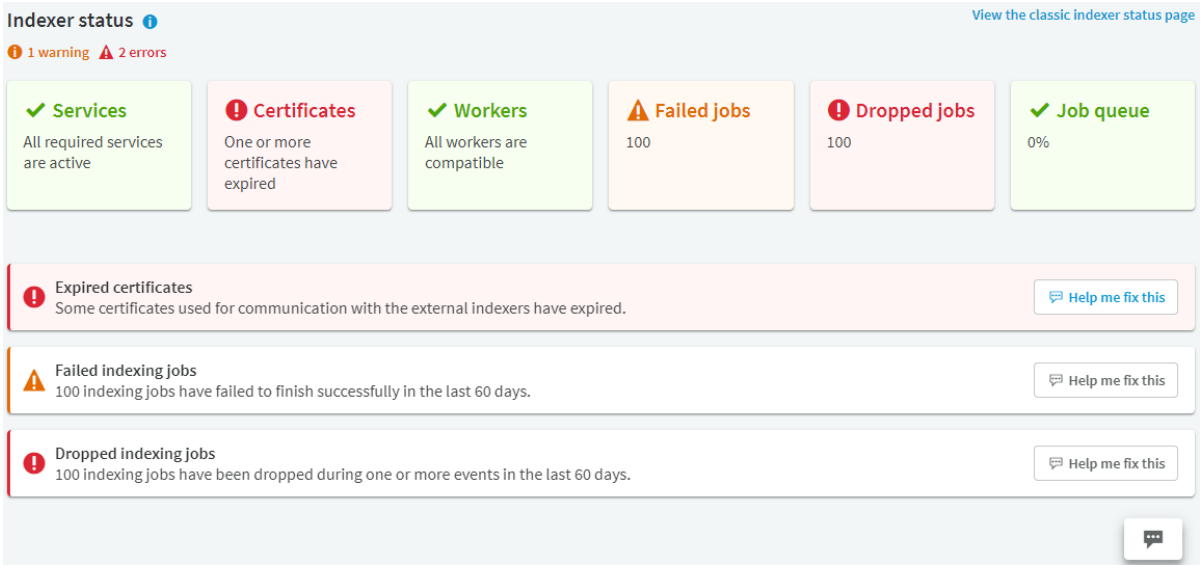
In an optimal state, or in case of a fairly new production environment, the page will display no errors or warnings:

Figure 249: Main Menu > Indexer Status — Indexer services working without errors or warnings



In a production environment that has been in use for some time, you will likely some have errors or warnings:

Figure 250: Main Menu > Indexer Status — Indexer services status with errors and warnings



NOTE: If you have errors and warnings listed on the Indexer Status page, you can use the help chat function of the SPS web UI by clicking the **Help me fix this** button on the error or warning of your choice.

The help chat function gives you a summary of the issue, explain how the issue affects your production environment, and guide you through the process of fixing the issue.

When the help chat function is still in use (for example, if you have not yet fixed the issue, but have clicked away from the chat window to use a terminal window), the minimized chat window is still accessible by clicking on the help chat ( ) icon, located at the bottom right corner of the web UI.

Services

The general workflow of indexing requires a few internal services to be active and running. Inactive internal services interrupt your workflow, causing delays or data loss. You can monitor the state of these internal services and see if any of them are inactive.

Certificates

Certificates and keys used to encrypt the communication between the indexer service and the external indexer ensure the processing of audit trails. You can monitor the validity of these credentials and see if any of them have expired or is about to expire soon.

Workers

The indexer service and the indexer service you use on your workers should always be compatible, otherwise external indexers will not process audit trails, and the queue load of the indexer service may be affected. You can monitor the state of the compatibility of your indexer service and the external indexers.

Failed jobs

When some of your indexing jobs are not finished successfully (for example, as a result of audit trail files moved or deleted during indexing, unsupported protocol versions used during remote sessions, or misconfigured indexer worker key stores), some of your recorded audit trails are not processed completely. As a result, the search function to the affected audit trail files is limited. You can monitor which indexing jobs may be affected and find the reason for the failure.

Dropped jobs

When your indexing jobs are dropped during an unknown even (for example, a mismanaged upgrade, an internal service shutdown, or an ill-timed system reboot), the affected audit trails are not processed completely, and some of the recorded contents are not indexed. As a result, the search function to the affected audit trail files is limited. You can monitor which indexing jobs may be affected and find the reason for the failure.

Job queue

When there are no free indexer workers to process your audit trails, indexing jobs will wait in a priority queue. Long queues may cause delays in using detailed search, and if your queue gets full, your most recent indexing jobs get dropped. By monitoring the status of your job queue, you can see if you may encounter delays, or in case of a full queue, dropped indexing jobs.

Monitoring the status of the indexer services in classic view

You can monitor the status of audit trail processing in detail by navigating to the **Indexer Status** page of the **Main Menu**, and clicking the **View the classic indexer status page** link in the upper right corner, or the **View the classic indexer status page** button, if visible.

TIP: To automatically refresh the **Indexer Status** page every 5 seconds, select **Auto refresh**. To refresh the page immediately, click **Refresh now**.

TIP: To view the status of your indexer services in a summarized view, click the **View the new indexer status page** button in the top right corner of the web UI.

Elements of the Indexer Status page in classic view

The following list describes the elements of the **Indexer** page and their functions.

- **Worker status**: displays information about the worker groups.
    - **Indexer IP address**: the IP address of the indexer running on One Identity Safeguard for Privileged Sessions (SPS) or an external indexer.

NOTE: **127.0.0.1** indicates the indexer running on SPS, while any IP address other than **127.0.0.1** indicates an external indexer.

- **Capabilities**: the type of job(s) this worker can perform. Possible job types are **index**, **near-realtime**, **screenshot**, and **video**.
- **Capacity (free / total)**: the available and total **Capacity** of the indexer. The value of the **total** capacity indicates the number of maximum parallel audit trails that the indexer can process.
- **Indexing**: the number of the active processes currently indexing an audit trail.
- **Generating video**: the number of the active processes currently generating a video.
- **Generating screenshot**: the number of the active processes currently generating a screenshot.

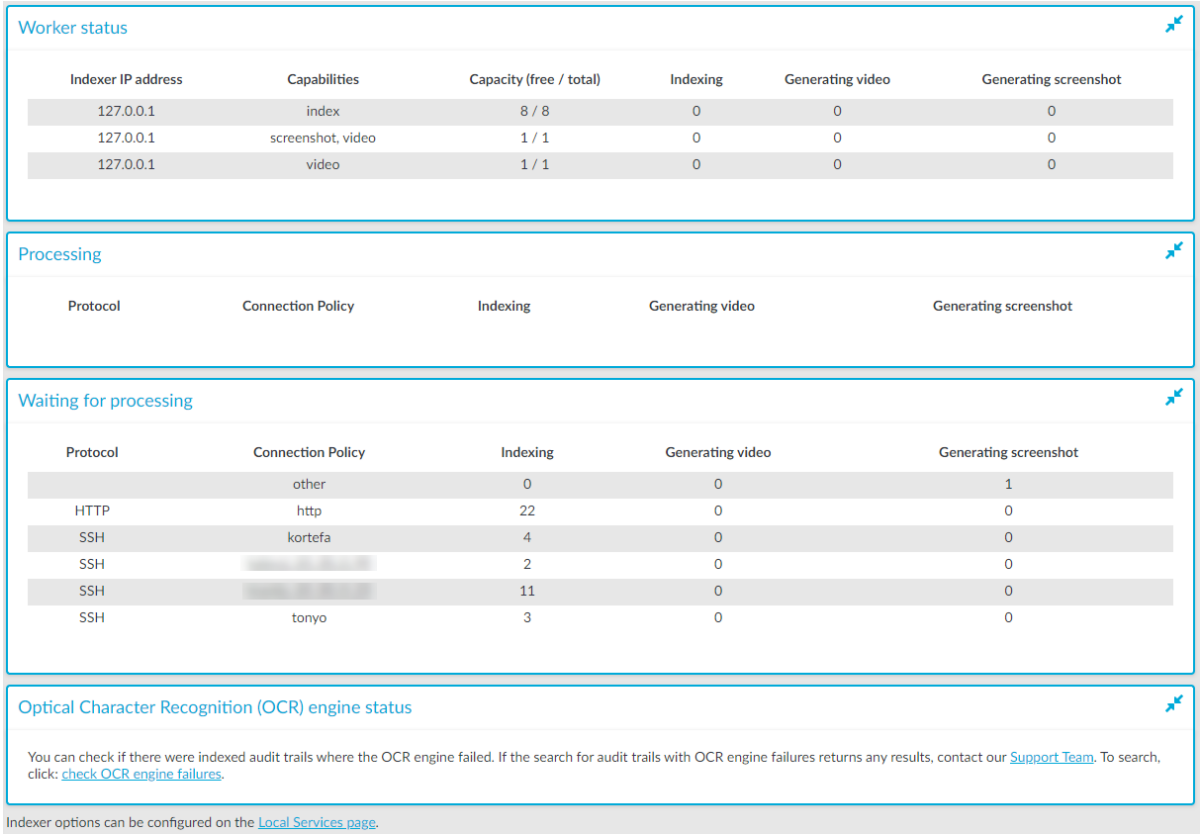
- **Processing**: audit trails currently being processed.
- **Waiting for processing**: audit trails waiting to be processed.

NOTE: Audit trails in the **Indexing** column may indicate any of the following:

- The maximal queue size is `1000`. If there are several trails waiting to be indexed, SPS will keep numerous trails in the queue.
- The worker with the appropriate key for decryption is not available at the moment, and there are no other workers with the required key to take over indexing.
- There are no workers with the required capacity available.

- **Optical Character Recognition (OCR) engine status**: It allows you to check and report indexed audit trails where the OCR engine failed. You can perform a search on the Search interface using the provided link and if the search returns any results, you can contact our Support Team to submit a report.

Figure 251: Indexer > Indexer status — Monitoring the status of the indexers



HTTP indexer configuration format

This section describes the configuration format and options of the HTTP indexer (that is, how and which fields of the HTTP audit trails are indexed). For details on how to customize HTTP indexing, see Customizing the indexing of HTTP traffic on page 698.

NOTE: If you want to index HTTP POST messages, include the `"application/x-www-form-urlencoded"` Content-Type in the `General > WhiteList` list. The indexer will decode URL encoding (percentage encoding), and create `key=value` pairs from the form fields and their values. Note that in the values, the indexer will replace whitespace with the under-score (_) character. To avoid indexing sensitive information (for example, passwords from login forms), use the `Form > Blacklist` option.

HTTP indexer configuration options

General

Type: Top level item

Description: Determines which HTTP Content-Types are indexed. An HTTP message is indexed only if its Content-Type is listed in `Whitelist` and is not listed in `Blacklist`.

For example:

```
"General": {
                  "Whitelist": ["text/.*", ".*json.*", "multipart/.*",
"application/x-www-form-urlencoded"],
                  "Blacklist": ["text/css", "application/javascript",
"text/xslt", ".*xml.*"]
},
```

General (Whitelist)

Type: list

Description: The list of HTTP Content-Types to index. Every entry of the list is treated as a regular expression.

For example:

```
"Whitelist": ["text/.*", ".*json.*", "multipart/.*", "application/x-www-form-
urlencoded"],
```

General (Blacklist)

Type: list

Description: The list of HTTP Content-Types that are not indexed. Every entry of the list is treated as a regular expression.

For example:

```
"Blacklist": ["text/css", "application/javascript", "text/xslt", ".*xml.*"]
```

Form

Type: Top level item

Description: Determines which fields are indexed in HTTP POST messages.

For example:

```
"Form": {
             "Blacklist": ["password", "pass"]
},
```

NOTE: If you want to index HTTP POST messages, include the `"application/x-www-form-urlencoded"` Content-Type in the `General > WhiteList` list. The indexer will decode URL encoding (percentage encoding), and create `key=value` pairs from the form fields and their values. Note that in the values, the indexer will replace whitespace with the under-score (_) character. To avoid indexing sensitive information (for example, passwords

from login forms), use the `Form > Blacklist` option.

Form (Blacklist)

Type: list

Description: The list of fields that are not indexed in HTTP POST messages (for example, when submitting forms, such as login forms). Every entry of the list is treated as a regular expression.

For example:

```
"Blacklist": ["password", "pass"]
```

Html

Type: Top level item

Description: Include this section in the configuration to process `text/html` messages. HTML tags are stripped from the text, and only their content is indexed (for example, `<html><title>Title</title></html>` becomes `Title`).

For example:

```
"Html": {
            "Attributes": ["href", "name", "value", "title", "id", "src"],
            "StrippedTags": ["script", "object", "style", "noscript",
"embed", "video", "audio", "canvas", "svg"]
}
```

Html (Attributes)

Type: list

Description: The list of HTML attributes that extracted as `key=value` pairs and indexed. Note that in the values, the indexer will replace whitespace with the underscore (_) character, and decode URL encoding. For example:

```
"Attributes": ["href", "name", "value", "title", "id", "src"],
```

Note that for the `content` attribute of the `meta name="description"`, `meta name="keywords"`, `meta name="author"` and `meta name="application-name"` is always indexed.

For example, if an audit trail contains the following HTML:

```
<head>
      <meta name="description" content="Web page description">
      <meta name="keywords" content="HTML,CSS,XML,JavaScript">
      <meta name="author" content="OI SA">
      <meta charset="UTF-8">
</head>
```

Then the index will contain the following text:

```
description=Web_page_description keywords=HTML,CSS,XML,JavaScript author=OI_SA
```

Html (StrippedTags)

Type: list

Description: The list of HTML tags that are not indexed.

For example:

```
"StrippedTags": ["script", "object", "style", "noscript", "embed", "video",
"audio", "canvas", "svg"]
```

17

Using the Sessions interface

This section provides an overview on how to use the Sessions interface. It describes how you can access the Sessions interface, lists the steps to take to search effectively, view the details of a connection, replay the audit trails, or export the search results as a comma-separated text file.

Prerequisites

Users need the **Search** privilege to access the Sessions interface.

NOTE: Assigning the **Search** privilege to a user on the **Users & Access Control** > **Appliance Access** page, automatically enables the **Access all sessions** privilege, and grants the user access to every audit trail. Access is granted even if the user is not a member of the groups listed in the **Access Control** option of the particular connection policy.

If you want users to access audit trails only for connections for which they are granted permission, see Assigning search privileges.

For information on configuring:

- Authorizers for a connection, see Configuring four-eyes authorization on page 870.
- User rights, see Managing user rights and user groups on page 361.

Accessing and using the Sessions interface

To access and use the Sessions interface

1. Navigate to **Audit** > **Sessions**.

   Sessions are displayed sorted by date. For ongoing sessions, the Sessions interface is updated in real-time to always show the most up-to-date information.

   You can view sessions in a card, table or flow view. Click ⋮ for more details and select from the list.
2. Specify a date and time range to restrict your search criteria as described in Specifying time ranges on page 720.
3. Filter connections as described in Using search queries on page 723.

4. Search the contents of audit trails as described in Searching in the contents of audit trails on page 799.
5. View connection details as described in Viewing session details on page 821.
6. Download and replay audit trails as described in Replaying audit trails in your browser on page 839.
7. To export the search results as a comma-separated text file, click ••• for more details and select **Export CSV**.

   NOTE: If your search returns more than 10,000 results, only the first 10,000 rows are exported. If you want to see all results, refine your search.

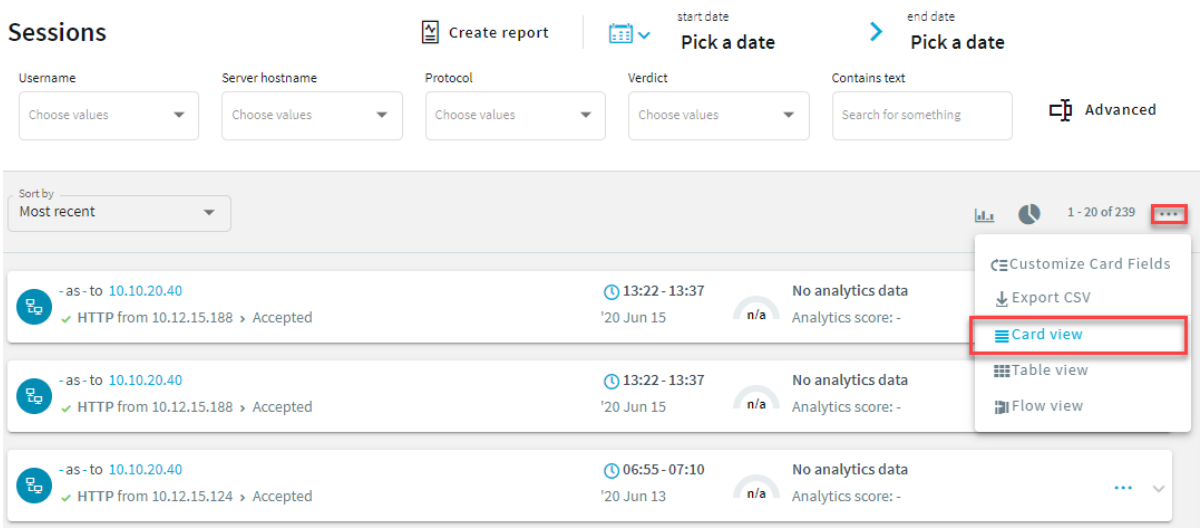
   To customize which fields are exported, click ⋮ for more details and select **Export CSV**.

Card view

You can view sessions in a card, table or flow view. Click ⋮ for more details and select from the list.

In card view, you can add additional search fields to the Search interface. For more information, see Adding custom fields to the card view.

Figure 252: Audit > Sessions — Card view



Adding custom fields to the card view

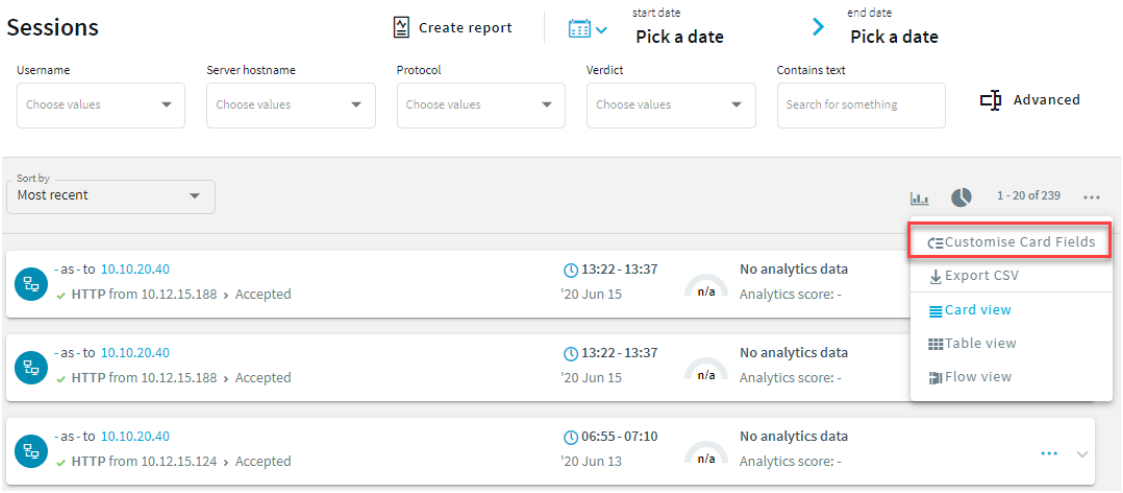
In card view, you can add additional search fields to the Search interface. This allows quick visualization of your preferred fields from the main page of the Search interface for each

session.

Adding custom fields to the card view

1. Click ••• for more details and select **Customize Card Fields** from the list.

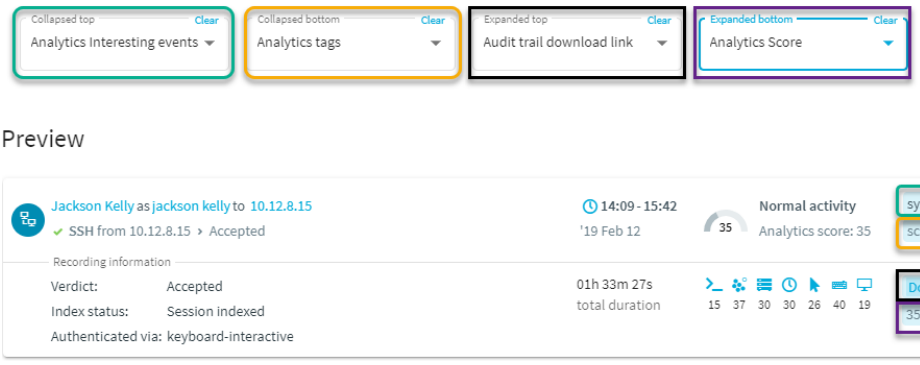
Figure 253: Audit > Sessions — Customize Card Fields



2. Select fields from **Collapsed top** and **Collapsed bottom** and the fields will be added as separate fields to the collapsed view of each session.

   The fields you select from **Expanded top** and **Expanded bottom** will be added to the expanded view of each session as shown in the figure below.

Figure 254: Audit > Sessions — Preview

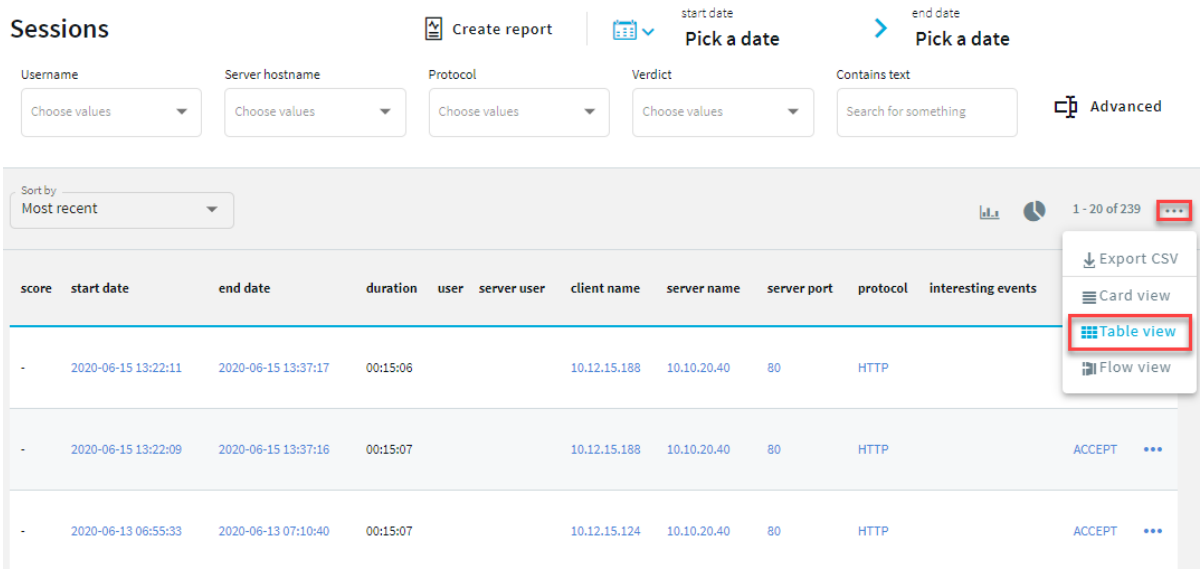


NOTE: Some fields may not be available for every session. When the field is not available, for example, if you have data recorded by SPP, the field will be empty.

Table view

You can view sessions in a card, table or flow view. Click ⋮ for more details and select from the list.

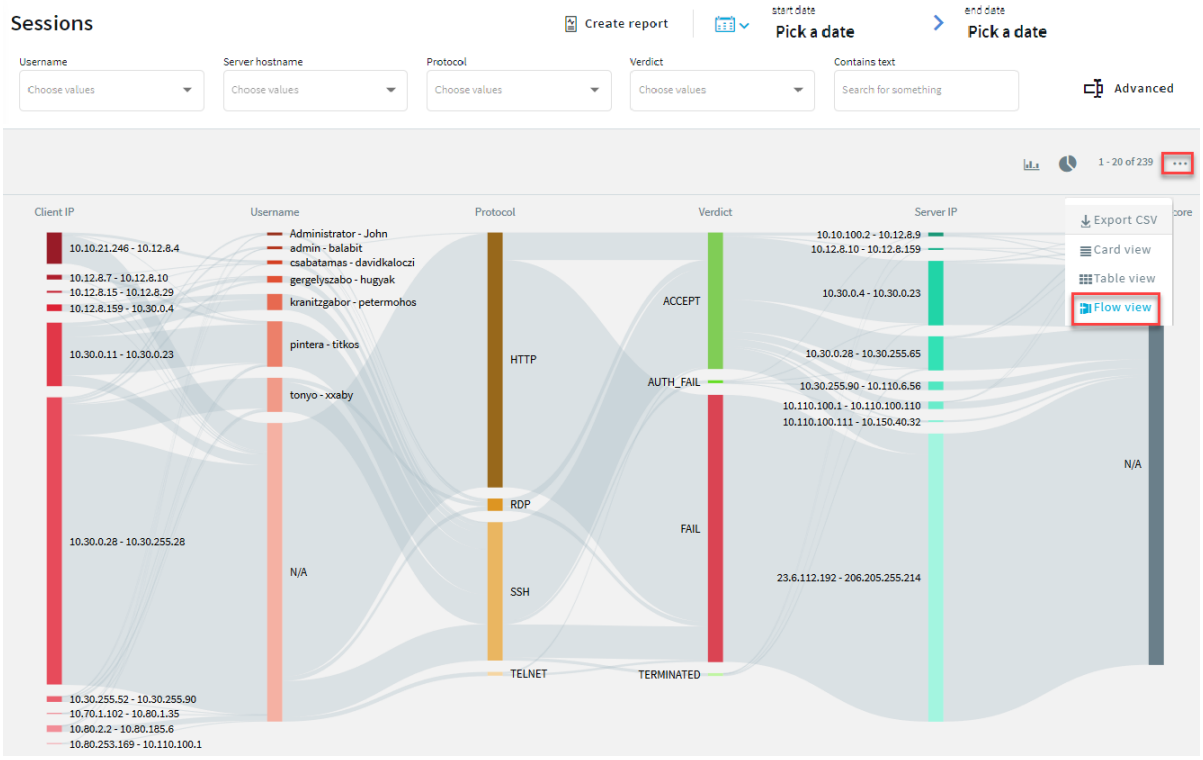
Figure 255: Audit > Sessions — Table view



Flow view

You can view sessions in a card, table or flow view. Click ⋮ for more details and select from the list.

Figure 256: Audit > Sessions — Flow view



The flow view allows you to:

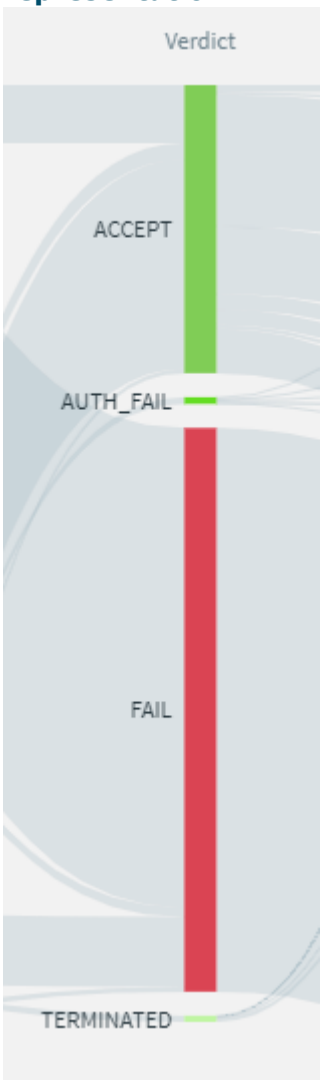
- Quickly visualize the distribution of the sessions based on their various metadata, such as, client address, username, protocol, verdict, server address, and One Identity Safeguard for Privileged Analytics (SPA) score.

  The metadata of the sessions are presented as vertical bars and each bar represents the proportional value of the data.

  Example: Proportional data representation

  The **Verdict** column shows that most of the sessions failed, a large number were accepted, and the rest of the sessions fall into the category of **AUTH_FAIL**, and **TERMINATED**.

Figure 257: Audit > Sessions > Flow view — proportional data representation



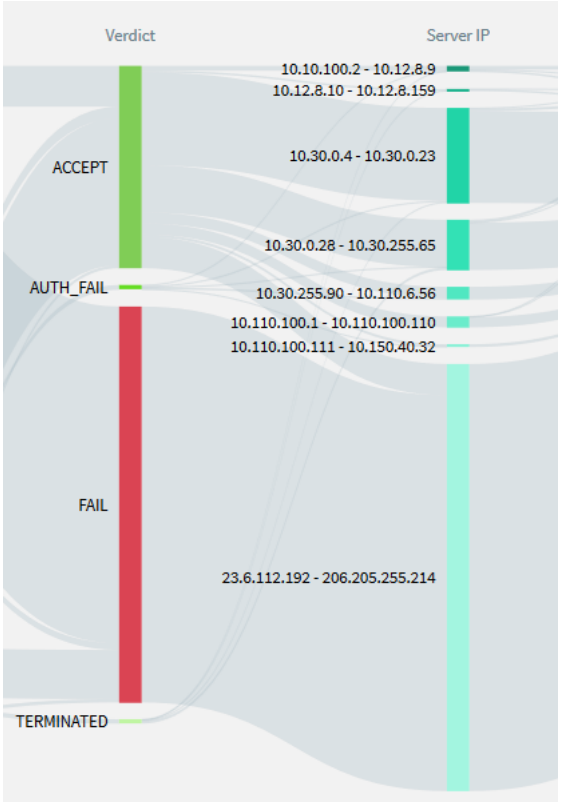
- See at a glance the relationship between various metadata and identify patterns in user behavior.

Example: Relationship between metadata

You want to have an overview of activities where access was denied.

A quick look at the **Verdict** column shows that there were several accesses where the authentication failed (**AUTH_FAIL**) and the lines from the **AUTH_FAIL** field point to several server addresses.

Figure 258: Audit > Sessions > Flow view — relationship between metadata



- Use it interactively to drill down further on information.

  To drill down on information, click on an item, then click **Search**.

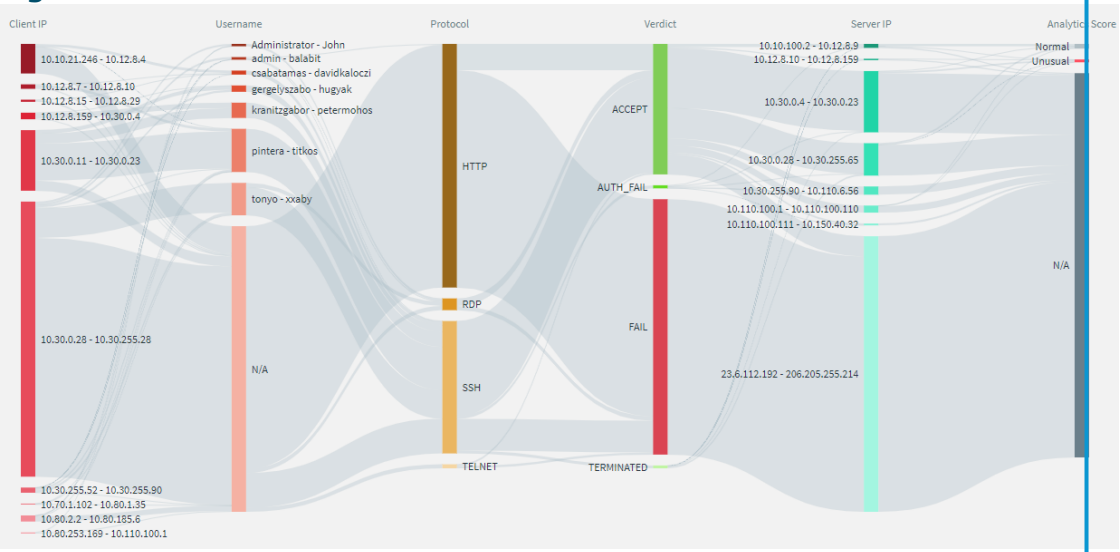
  TIP:To exclude an item, press `Ctrl` while clicking the item.

Example: Interactive drill down

You want to investigate if there were any unusual activities. To take a closer look, in the **Analytics Score** column, click **Unusual**, then click **Search**.

The flow view now only displays the unusual session activities. You can further narrow your search as required.

Figure 259: Audit > Sessions > Flow view — interactive drill down



Assigning search privileges

The following describes how to assign users to access sessions only for connections for which they are granted permission.

Users need the **Search** privilege to access the Search interface.

Assigning the **Search** privilege to a user on the **Users & Access Control** > **Appliance Access** page, automatically enables the **Access all sessions** privilege, and grants the user access to every session. Access is granted even if the user is not a member of the groups listed in the **Access Control** option of the particular connection policy.

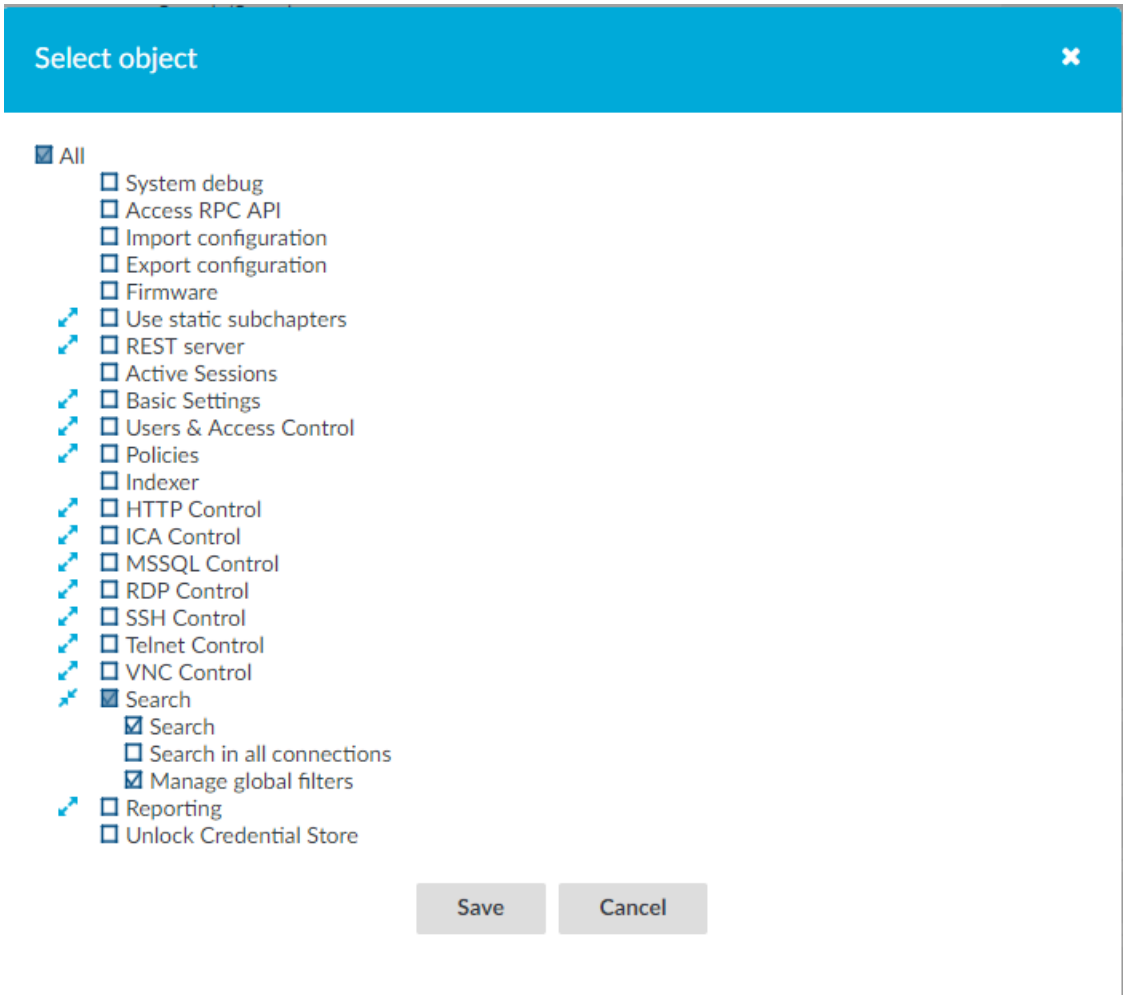
Prerequisites

- You have created a user for which you want to assign the search privilege. For more information, see Creating local users in One Identity Safeguard for Privileged Sessions (SPS).
- You have created a usergroup. For more information, see Managing local user groups.

To assign users to access sessions only for connections for which they are granted permission

1. Navigate to **Users & Access Control** > **Appliance Access**.

Figure 260: Users & Access Control > Appliance Access — Configuring search privileges



2. Assign the **Search** privilege to your usergroup as described in Assigning privileges to user groups for the One Identity Safeguard for Privileged Sessions (SPS) web interface.
3. Deselect the **Access all sessions** privilege so that users can access sessions only for connections for which they are granted permission.
4. To grant permission to a specific connection, navigate to the **Connections** page of the traffic (for example to **Traffic Controls** > **SSH** > **Connections**), and select the connection policy to modify.

5. Scroll down to the **Access Control** section, and click  .
6. Enter the name of the usergroup whose members are permitted to access the Search interface into the **Authorizer Group** field. This group must exist on the **Users & Access Control** > **Local User Groups** page.

   CAUTION:

   Usernames, the names of user lists, and the names of usergroups are case sensitive.

7. Set the permissions of the usergroup.
   - If the usergroup can authorize (that is, enable) and audit (that is, monitor in real-time and download the audit trails) the sessions, select **Permission** > **Follow&Authorize**.
   - If the usergroup can only audit the sessions but cannot authorize, select **Permission** > **Follow**.

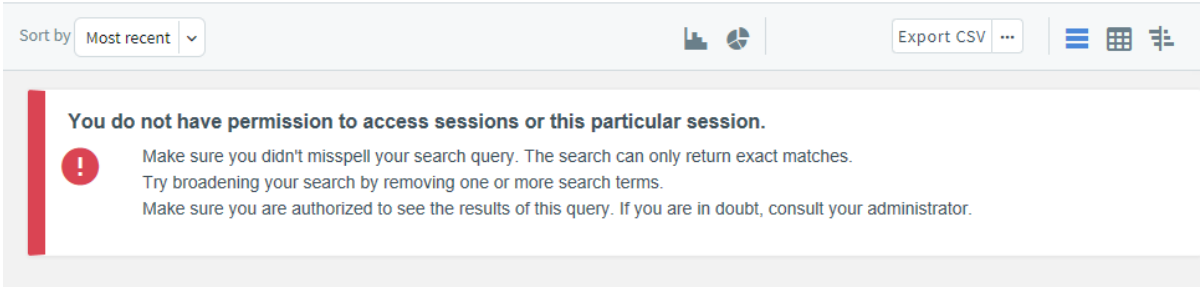
   NOTE: If the **Client user is** > **Member of** field is set, the auditor can only monitor the sessions of the specified usergroup. However, if **Client user is** > **Member of** field is set, the Auditor cannot access the **Sessions** page. To avoid this problem, add another Access Control rule for the **Authorizer Group** without setting the **Client user is**field.

   The `admin` user of One Identity Safeguard for Privileged Sessions (SPS) can audit and authorize every connection.

Result

Users with the relevant privileges can now access the sessions for which they are granted permission. If users do not have the required permission to access sessions, a warning message is displayed and no session is visible as shown below:

Figure 261: Audit > Sessions — Permission denied



Specifying time ranges

Specify a time range to restrict, or filter your search criteria by setting boundaries on your searches. You can restrict the search to one of the preset time ranges, or use a custom time range for a more specific search.

When you specify a time range, the search result includes:

- Connections started and finished anywhere between the start time and end time you specified.
- Connections started anywhere between the start time and end time you specified.
- Connections ended anywhere between the start time and end time you specified.
- Active connections if they were started anywhere between the start time and the end time you specified.

For example, at `17:00` PM you specify a start date of `10:00` AM and end date of `15:00` PM for your search. The search result includes:

- Connections started at `8:00` AM and ended at `14:00` PM.
- Connections started at `11:00` AM and ended at `14:00` PM.
- Connections started at `11:00` AM and ended at `16:00` PM.
- Active connections started at `11:00` AM.
- Active connections started at `10:00` AM.

To specify time ranges

1. To select the start date of your search, click **Pick a date**.

   Alternatively, use the (shortcuts) button to restrict the search to one of the preset time ranges. For example, to investigate an incident that occurred sometime in the last hour, you can select **Today**, but a better option is **Last 60 minutes**.

   Figure 262: Audit > Sessions — Pick a date

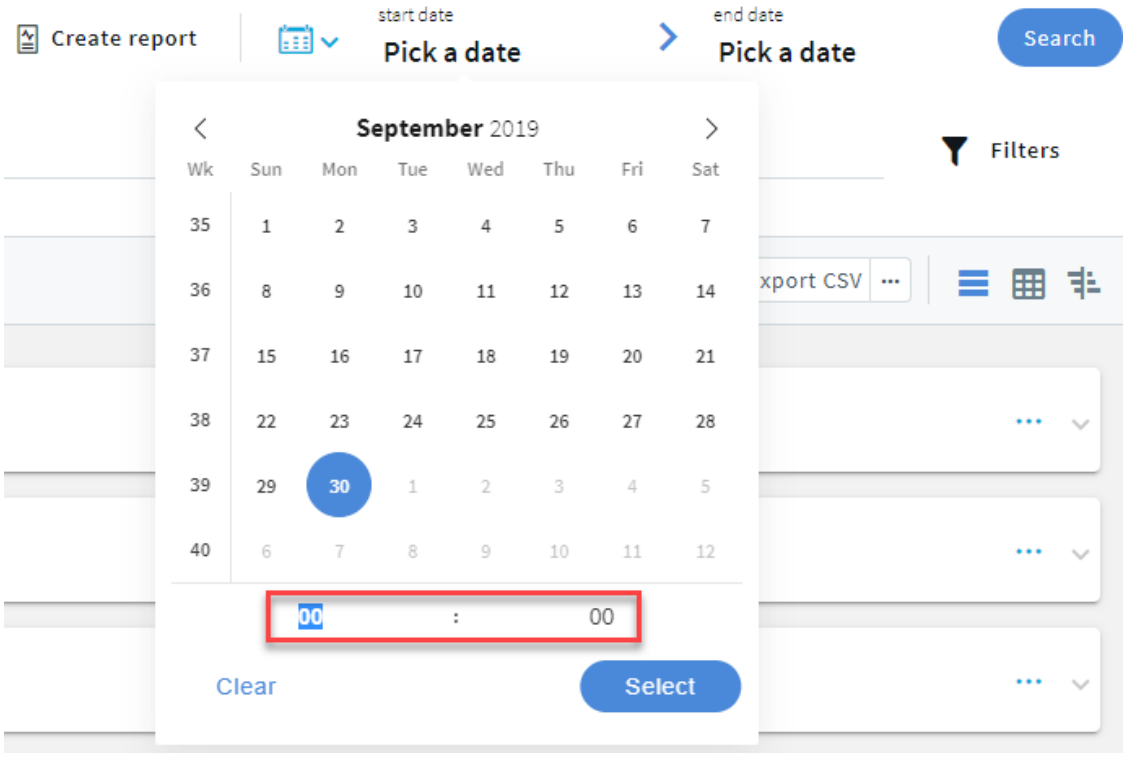
   

2. From the calendar, select the start date as required.

   NOTE: The date refers to the timezone configured on SPS.

3. For exact time ranges, specify to search by the hour and minute.

Figure 263: Audit > Sessions — Specify hour and minute



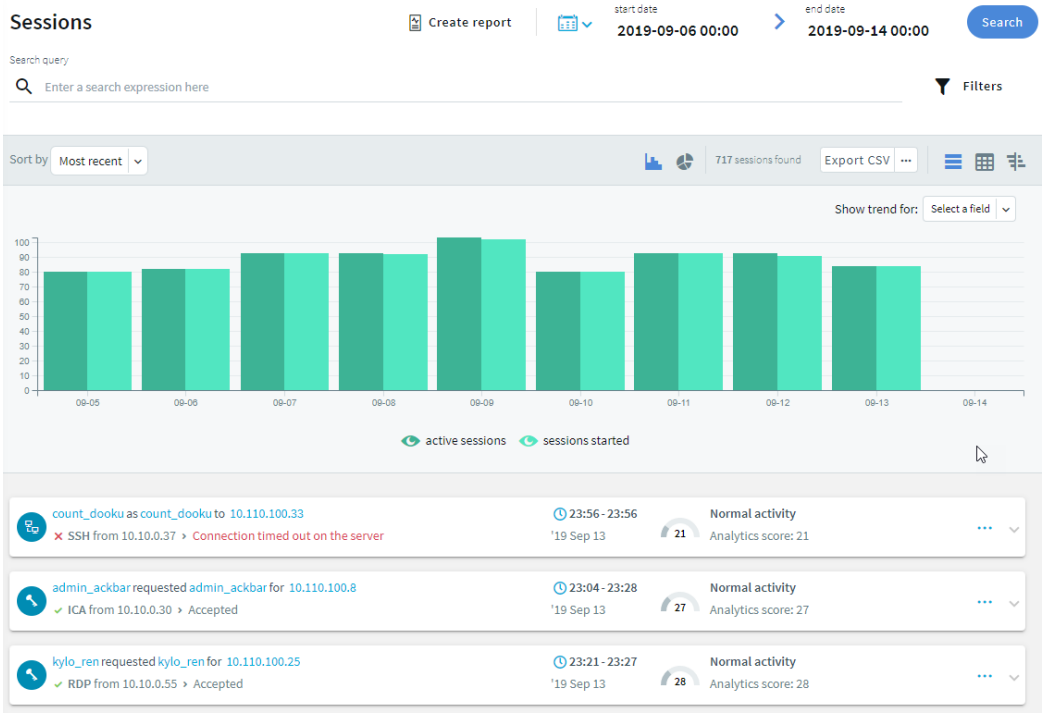
4. To select the end date of your search, click **Pick a date** and select a date as required.

   If you specify only the start date, the end date is set to the current time.

5. *Optional:* To clear the start and end date, click (shortcuts) > **All time**.
6. *Optional:* You can use the timeline for a quick time range selection and visual representation of sessions in the selected interval.

Click the  icon.

Figure 264: Audit > Sessions — Using the timeline



a.

The bars display the number of results in the selected interval.

The **active sessions** columns indicate all the sessions, which were active in the selected interval. The **sessions started** columns indicate all the sessions started during the selected interval. For example, if the selected interval is today between `8:00` AM and `9:00` AM, then a session started at `7:00` AM but lasting after `8:00` AM is displayed in the **active sessions** column. A session started at `8:30` AM is displayed in the **sessions started** column. Since the session was active during the selected time interval, the session started at `8:30` AM is also displayed in the **active sessions** column.

To disable the active sessions and view only the started sessions in the timeline, click active sessions. To disable the started sessions and view only the active sessions in the timeline, click sessions started.

Hovering the mouse above a bar displays the number of entries and the start and end date of the period that the bar represents.

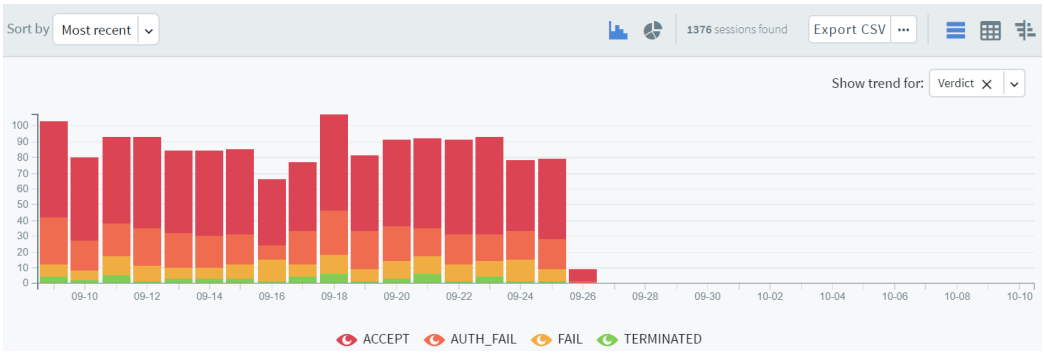
Trend analysis allows you to use the timeline to find changes over time. For example, to find the time range where terminated connections had a significant peak compared to other days, from the **Show trend for** drop-down menu, select **Verdict**. Note that you can only view trend analysis for Active, Analytics

Score, Client name, Protocol, Server hostname, Server port, Server username, Username and Verdict. All the other selections are grayed out.

The colors of the bars in the timeline allow you to quickly find the time range with a higher number of terminated sessions.

Optional: To clear the trend analysis view, from the **Show trend for** drop-down menu, select **X**.

Figure 265: Audit > Sessions — Using the timeline - trend analysis



b. To select a range, drag the mouse pointer across the timeline or use Shift+Click and select multiple bars.

Using search queries

This section describes how you can use search queries to perform a more specific search.

To search using search queries

1. Enter a search query in the **Search query** field, or click on an entry in the table.

   To search, enter a valid search field followed by a value in the **search field: VALUE** format. For example, if you enter **protocol: SSH**, the search returns all the SSH sessions.

   > TIP: Search is case insensitive. To make the search case sensitive, enclose the search keywords in double quotes.
   >
   > The search queries can include only alphanumerical characters. You can use complex expressions and boolean operators, for example, AND, OR, <, >, and so on.
   >
   > To search for sessions by fields containing special characters, you need to use the exact match by wrapping your search expression in double quotation marks. For example, to return sessions targeting a specific server domain, you must enter `server.name: "foo-bar.domain.com"`.

For the list of search fields that you can use, see List of available search queries on page 724.

For more information on how to use more complex keyphrases that are not covered in this guide, see the Apache Lucene documentation.

There are search fields that are not displayed but you can still use them to query the sessions. For example, you can search for active connections using the `active` search field, and search results are listed accordingly, but there is no **active** field displayed in the search table or in the **Overview**, **Details**, and **Timeline** tabs.

Alternatively, click  and set the filters you need from the appropriate columns. For example, to search for a specific username, select it using the drop-down menu of the **Username** column. For a more generic search, you can enter any text in the **Contains text** column.

2. After specifying the relevant query, click **Search** or press **Enter**.

   TIP:To save the queries for future use, simply save the URL or bookmark it in your browser.

Expected result

Session metadata is displayed in columns that you can query for any parameter, or a combination of parameters. You can view the metadata in the search columns and also displayed as fields in the **Overview**, **Details**, and **Timeline** tabs.

List of available search queries

This section lists the search fields that you can use to perform a more specific search. For information about how to use the search fields listed below, see Using search queries.

The following table provides an explanation to the search field tables listed in this section.

Name:	Specifies the meaningful and easily readable name of the search field.
Search field:	Specifies the filter expression that you can use to filter the audit trails. For example, to narrow your search to a specific server-side IP address, you can enter the **`server.address: 10.30.255.70`** search query in the **Search query** field. All search results that contain that specific server IP address are listed.
Displayed:	Specifies if the search query result is displayed as a field in the search columns or in the **Overview**, **Details**, and **Timeline** tabs. There are search fields that are not displayed but you can still use them to filter the audit trails. For example, you can search for active connections using the `active` search field, and search results are listed accordingly, but there is no **active** field displayed in the search table or in the **Overview**, **Details**, and **Timeline** tabs.

The following search fields are available:

alert

Alert type

Name:	Alert type
Search field:	alert_type
Type:	enum
Displayed:	True

The type of the alert.

Possible values:

- *adp.event.command*: A command entered in SSH or Telnet.
- *adp.event.screen.content*: Alert triggered by the screen content.
- *adp.event.screen.creditcard*: Credit card numbers detected. Displayed only as an alert, not visible in the events.
- *adp.event.screen.windowtitle*: The title of the window in graphic protocols.

Channel ID

Name:	Channel ID
Search field:	channel_id
Type:	string
Displayed:	True

The id of the channel the alert belongs to.

Matched regexp on action

Name:	Matched regexp on action
Search field:	matched_action
Type:	string
Displayed:	True

The regular expression that matched the command line without prompt

Matched content

Name:	Matched content
Search field:	matched_content
Type:	string
Displayed:	True

The content the alert matched.

Note that this value contains the context of the match as well. For example, if a Content Policy triggers an alert if a user types the `sudo` command, then the `psm.alerts.matched_content` value contains the entire command line, including the command prompt, for example, **`myuser@examplehost:~$ man sudo`**.

Matched regexp

Name:	Matched regexp
Search field:	matched_regexp
Type:	string
Displayed:	True

The regular expression that matched the content.

For details, see Real-time content monitoring with Content Policies on page 498.

Alert ID

Name:	Alert ID
Search field:	record_id
Type:	long
Displayed:	True

The identifier of the alert within the audit trail (.zat file).

Rule name

Name:	Rule name
Search field:	rule_name
Type:	string
Displayed:	True

The name of the content policy rule.
Note that this is not the name of the Content Policy.

Alert time

Name:	Alert time
Search field:	time
Type:	date
Displayed:	False

The timestamp of the alert.

channel

Channel is active

Name:	Channel is active
Search field:	active
Type:	boolean
Displayed:	True

True if the session has not ended yet.

Application

Name:	Application
Search field:	application
Type:	string
Displayed:	True

The name of the application accessed in a seamless Citrix ICA connection.

Audit stream ID

Name:	Audit stream ID
Search field:	audit_stream_id

Type:	string
Displayed:	True

The identifier of the channel's audit stream. If the session does not have an audit trail, this element is not used.

Channel ID

Name:	Channel ID
Search field:	channel_id
Type:	long
Displayed:	True

The unique ID of the channel.

Client X.509 Subject

Name:	Client X.509 Subject
Search field:	client_x509_subject
Type:	string
Displayed:	True

The client's certificate in TELNET or VNC sessions. Available only if the 'Client-side transport security settings > Peer certificate validation' option is enabled in One Identity Safeguard for Privileged Sessions.

Executed commands

Name:	Executed commands
Search field:	command
Type:	string
Displayed:	True

Lists the commands run in an SSH session.

Port-forward target IP

Name:	Port-forward target IP
Search field:	connected.ip
Type:	ip
Displayed:	True

The traffic was forwarded to this IP address in Remote Forward and Local Forward channels.

Port-forward target name

Name:	Port-forward target name
Search field:	connected.name
Type:	string
Displayed:	True

The traffic was forwarded to this host in Remote Forward and Local Forward channels. If the hostname is not available, this field contains the IP address of the host

Port-forward target port

Name:	Port-forward target port
Search field:	connected.port
Type:	port
Displayed:	True

The traffic was forwarded to this port in Remote Forward and Local Forward channels.

Device name

Name:	Device name
Search field:	device_name
Type:	string
Displayed:	True

The name or ID of the shared device (redirect) used in the RDP connection.

Description: Used with the `serial redirect`, `parallel redirect`, `printer redirect`, `disk redirect`, and `scard redirect` RDP channel types.

The name of the device.

Channel duration

Name:	Channel duration
Search field:	duration
Type:	long
Displayed:	True

The length of the channel (how long the channel lasted).

Dynamic channel

Name:	Dynamic channel
Search field:	dynamic_channel
Type:	string
Displayed:	True

The name or ID of the dynamic channel opened in the RDP session.

Channel end time

Name:	Channel end time
Search field:	end_time
Type:	date
Displayed:	True

Date when the channel was closed.

Environment

Name:	Environment
Search field:	environment
Type:	string
Displayed:	True

Date when the channel was closed.

Four-eyes authorizer

Name:	Four-eyes authorizer
Search field:	four_eyes_authorizer
Type:	string
Displayed:	True

The username of the user who authorized the session. Available only if four-eyes authorization is required for the channel.

Four-eyes description

Name:	Four-eyes description
Search field:	four_eyes_description
Type:	string
Displayed:	True

The description submitted by the authorizer of the session.

Channel originator IP address

Name:	Channel originator IP address
Search field:	originator.ip
Type:	ip
Displayed:	True

The IP address of the host initiating the channel in Remote Forward and Local Forward channels. Note that this host is not necessarily the client or the server of the SSH connection.

Channel originator name

Name:	Channel originator name
Search field:	originator.name

Type:	string
Displayed:	True

The hostname of the host initiating the channel in Remote Forward and Local Forward channels. Note that this host is not necessarily the client or the server of the SSH connection. If the hostname is not available, this field contains the IP address of the host.

Originator port

Name:	Originator port
Search field:	originator.port
Type:	port
Displayed:	True

The number of the forwarded port in Remote Forward and Local Forward SSH channels.

Rule number

Name:	Rule number
Search field:	rule_num
Type:	string
Displayed:	True

The number of the line in the Channel policy applied to the channel.

SCP path

Name:	SCP path
Search field:	scp_path
Type:	string
Displayed:	True

Name and path of the file copied via SCP. Available only for SCP sessions (Session exec SCP SSH channels) if the Log file transfers to database option isenabled in the Channel Policy of the connection.

Channel start time

Name:	Channel start time
Search field:	start_time
Type:	date
Displayed:	True

Date when the channel was started.

Subsystem name

Name:	Subsystem name
Search field:	subsystem_name
Type:	string
Displayed:	True

Name of the SSH subsystem used in the channel.

Channel type

Name:	Channel type
Search field:	type
Type:	enum
Displayed:	True

Type of the channel.
Possible values:

- *#drawing*: Drawing
- *CTXCAM*: Audio
- *CTXCDM*: Drive
- *CTXCLIP*: Clipboard
- *CTXCOM1*: Printer (COM1)
- *CTXCOM2*: Printer (COM2)
- *CTXCPM*: Printer Spooler
- *CTXFLSH*: HDX Mediastream
- *CTXLPT1*: Printer (LPT1)

- *CTXLPT2*: Printer (LPT2)
- *CTXSCRD*: Smartcard
- *CTXTW*: Drawing (Thinwire)
- *CTXTWI*: Seamless
- *CTXUSB*: USB
- *SPDBRS*: Speedbrowse
- *auth-agent*: Agent
- *cliprdr*: Clipboard
- *custom*: Custom
- *direct-tcpip*: Local forward
- *drawing*: Drawing
- *drdynvc*: Dynamic virtual channel
- *forwarded-tcpip*: Remote forward
- *http*: HTTP
- *rdpdr*: Redirects
- *rdpdr-disk*: Disk redirect
- *rdpdr-parallel*: Parallel redirect
- *rdpdr-printer*: Printer redirect
- *rdpdr-scard*: SCard redirect
- *rdpdr-serial*: Serial redirect
- *rdpsnd*: Sound
- *seamrdp*: Seamless
- *session-exec*: Session exec
- *session-exec-scp*: Session exec SCP
- *session-shell*: Session shell
- *session-subsystem*: Session subsystem
- *session-subsystem-sftp*: Session SFTP
- *telnet*: Telnet
- *vnc*: VNC
- *websocket*: WebSocket
- *x11*: X11 forward

Channel verdict

Name:	Channel verdict
Search field:	verdict
Type:	enum
Displayed:	True

Indicates what One Identity Safeguard for Privileged Sessions decided about the channel.

Possible values:

- *ACCEPT*: Accepted
- *DENY*: Denied
- *FOUR_EYES_DEFERRED*: Waiting for remote username
- *FOUR_EYES_ERROR*: Internal error during four-eyes authorization
- *FOUR_EYES_REJECT*: Four-eyes authorization rejected
- *FOUR_EYES_TIMEOUT*: Four-eyes authorization timed out

content

Window title

Name:	Window title
Search field:	title
Type:	string
Displayed:	True

The content of the title bar in the active window. The window title typically contains the name of the application, or the name of the dialogue box. Only available in graphical sessions (for example, RDP), if indexing is enabled.

Command

Name:	Command
Search field:	command
Type:	string
Displayed:	True

The commands that the user run in the session. Only available in terminal sessions (for example, SSH), if indexing is enabled.

event

Event Action

Name:	Event Action
Search field:	action
Type:	string
Displayed:	True

The command line without prompt in commands

Channel ID

Name:	Channel ID
Search field:	channel_id
Type:	string
Displayed:	True

The id of the channel the event belongs to.

Event content

Name:	Event content
Search field:	content
Type:	string
Displayed:	True

The command run, or the window title detected in the channel (for example, ls, exit, or Firefox).

Protocol details

Name:	Protocol details
Search field:	details
Type:	string
Displayed:	True

The details of the protocol used for the operation.

Event ID

Name:	Event ID
Search field:	event_id
Type:	string
Displayed:	True

The identifier of the vault event.

Operation

Name:	Operation
Search field:	operation
Type:	string
Displayed:	True

The type of the operation that occurred, for example, Create file (in the case of FTP) or GET (in the case of HTTP).

Path

Name:	Path
Search field:	path
Type:	string
Displayed:	True

The path (if any) used by the operation that occurred.

Event ID

Name:	Event ID
Search field:	record_id
Type:	long
Displayed:	True

The identifier of the event within the audit trail (.zat file).

Response code

Name:	Response code
Search field:	response_code
Type:	long
Displayed:	True

The status code of the protocol response (if any) returned.

Event date

Name:	Event date
Search field:	time
Type:	date
Displayed:	False

The date when the event happened.

Event type

Name:	Event type
Search field:	type
Type:	string
Displayed:	True

The type of the event, for example, command, screen_content, window_title.

indexer_info

Commands indexed

Name:	Commands indexed
Search field:	config.command.enabled

Type:	boolean
Displayed:	True

True if commands were extracted while indexing the session.

Keyboard buffering interval

Name:	Keyboard buffering interval
Search field:	config.keyboard.buffer_interval
Type:	double
Displayed:	True

The buffering interval in milliseconds used when extracting keyboard events while indexing the session.

Keyboard extracted

Name:	Keyboard extracted
Search field:	config.keyboard.enabled
Type:	boolean
Displayed:	True

True if keyboard events were extracted while indexing the session.

Mouse buffering interval

Name:	Mouse buffering interval
Search field:	config.mouse.buffer_interval
Type:	double
Displayed:	True

The buffering interval in milliseconds used when extracting mouse events while indexing the session.

Mouse extracted

Name:	Mouse extracted
Search field:	config.mouse.enabled
Type:	boolean
Displayed:	True

True if mouse events were extracted while indexing the session.

Near real-time indexing

Name:	Near real-time indexing
Search field:	config.near_realtime
Type:	boolean
Displayed:	True

True if indexing this session was done near real-time (when the session was still active).

OCR languages

Name:	OCR languages
Search field:	config.ocr_languages
Type:	string
Displayed:	True

The language configuration for optical character recognition used when indexing the session.

Screen content indexed

Name:	Screen content indexed
Search field:	config.screen.enabled
Type:	boolean
Displayed:	True

True if screen content was extracted while indexing the session.

OCR tradeoff

Name:	OCR tradeoff
Search field:	config.screen.omnipage_trade_off
Type:	string
Displayed:	True

The tradeoff used for optical character recognition when extracting screen content while indexing the session.

Titles indexed

Name:	Titles indexed
Search field:	config.title.enabled
Type:	boolean
Displayed:	True

True if window titles were extracted while indexing the session.

Indexing error

Name:	Indexing error
Search field:	error.message
Type:	string
Displayed:	True

The reason why indexing failed

Indexing cpu time

Name:	Indexing cpu time
Search field:	statistics.cpu_time
Type:	long
Displayed:	True

The CPU time that indexing this session took in milliseconds.

Indexing duration

Name:	Indexing duration
Search field:	statistics.duration
Type:	long
Displayed:	True

The duration of time that indexing this session took in milliseconds.

Indexing start time

Name:	Indexing start time
Search field:	statistics.start_time
Type:	date
Displayed:	True

The time and date when indexing this session started.

Indexing status

Name:	Indexing status
Search field:	status
Type:	string
Displayed:	True

Shows if the channel has been indexed successfully or not.

Indexer ADP version

Name:	Indexer ADP version
Search field:	version.adp
Type:	string
Displayed:	True

The version of the audit data processor used for indexing the session

Indexer version

Name:	Indexer version
Search field:	version.worker
Type:	string
Displayed:	False

The version of the indexer worker used for indexing the session

ZAC created

Name:	ZAC created
Search field:	config.zac.enabled
Type:	boolean
Displayed:	False

True if an Audit Content file was created while indexing the session.

screen

Screen content

Name:	Screen content
Search field:	content
Type:	string
Displayed:	False

Text that appeared on the screen in the session.

Channel id in trail

Name:	Channel id in trail
Search field:	channel_id_in_trail
Type:	long
Displayed:	False

The ID of the channel where this content appeared. To check the channel ID (channel_id), select a session and click details. Navigate to details > Channels and click the channel type.

Screen content creation time

Name:	Screen content creation time
Search field:	time
Type:	date
Displayed:	False

The creation time of the indexed screen content.

Screen content ID

Name:	Screen content ID
Search field:	id
Type:	string
Displayed:	False

The ID of a screen content event.

session

Active

Name:	Active
Search field:	active
Type:	boolean
Displayed:	True

The session is still open.

Analytics Interesting events

Name:	Analytics Interesting events
Search field:	analytics.interesting_events
Type:	string
Displayed:	True

Collection of interesting command(s) and window title(s) from the session.

Analytics Score

Name:	Analytics Score
Search field:	analytics.score.aggregated
Type:	long
Displayed:	True

The risk score that the Analytics Module assigned to the session.Ranges from 0 to 100, 100 is the highest risk score.

Score time

Name:	Score time
Search field:	analytics.score.time
Type:	date
Displayed:	False

The scoring time of the given analytics. The different analytics are scored at different times based on the type of the analytics and certain configuration settings.

Command score

Name:	Command score
Search field:	analytics.score.details.command.score
Type:	long
Displayed:	True

Score given by the Command algorithm.

FIS score

Name:	FIS score
Search field:	analytics.score.details.fis.score
Type:	long
Displayed:	True

Score given by the Frequent Item Set (FIS) algorithm

Host login score

Name:	Host login score
Search field:	analytics.score.details.hostlogin.score
Type:	long
Displayed:	True

Score given by the Host login algorithm.

Login time score

Name:	Login time score
Search field:	analytics.score.details.logintime.score
Type:	long
Displayed:	True

Score given by the Login time algorithm.

Keystroke score

Name:	Keystroke score
Search field:	analytics.score.details.keystroke.score
Type:	long
Displayed:	True

Score given by the Keystroke algorithm.

Mouse score

Name:	Mouse score
Search field:	analytics.score.details.mouse.score
Type:	long
Displayed:	True

Score given by the Mouse algorithm.

Windowtitle score

Name:	Windowtitle score
Search field:	analytics.score.details.windowtitle.score
Type:	long
Displayed:	True

Score given by the Window title algorithm.

Scripted

Name:	Scripted
Search field:	analytics.scripted
Type:	boolean
Displayed:	True

True if the One Identity Safeguard for Privileged Analytics module marked the session as scripted because of non-human activity

Similar Sessions

Name:	Similar Sessions
Search field:	analytics.similar_sessions
Type:	string
Displayed:	True

Collection of similar sessions from different sources.

Bucketed duration

Name:	Bucketed duration
Search field:	analytics.bucketed_duration
Type:	string
Displayed:	True

Categorized length of session

Bucketed starting hour

Name:	Bucketed starting hour
Search field:	analytics.bucketed_starting_hour
Type:	string
Displayed:	True

Session start time categorized by hours

Analytics tags

Name:	Analytics tags
Search field:	analytics.tags
Type:	string
Displayed:	True

The Analytics tags section in Search > details.

Client IP

Name:	Client IP
Search field:	client.ip
Type:	ip
Displayed:	True

The IP address of the client that initiated the session.

Client name

Name:	Client name
Search field:	client.name
Type:	string
Displayed:	True

The name of the client that initiated the session.

Client port

Name:	Client port
Search field:	client.port
Type:	port
Displayed:	True

The port number of the client that initiated the session.

Creation time

Name:	Creation time
Search field:	creation_time
Type:	date
Displayed:	True

The first time the pipeline created the session. It is different from start_time and can be later than start_time.

Duration

Name:	Duration
Search field:	duration
Type:	long
Displayed:	True

The length of the session (how long the session lasted).

End time

Name:	End time
Search field:	end_time
Type:	date
Displayed:	True

Date when the session was closed.

For ongoing connections, the value is `null`.

Starting with SPS 5 LTS, the timestamp is in ISO 8601 format, for example, `2018-10-11T09:23:38.000+02:00`. In earlier versions, it was in UNIX timestamp format.

Log adapter

Name:	Log adapter
Search field:	log.adapter_name
Type:	string
Displayed:	True

The name of the Log Adapter Plugin. This plugin can be uploaded at Basic Settings > Plugins.

Log auth method

Name:	Log auth method
Search field:	log.auth_method
Type:	string
Displayed:	True

SSH relayed authentication method. It is configured at Traffic Controls > SSH > Authentication Policies > Relayed authentication methods.

Log syslog time

Name:	Log syslog time
Search field:	log.syslog_time
Type:	date
Displayed:	True

Date of the message in the ISO 8601 compatible standard timestamp format.

Node ID

Name:	Node ID
Search field:	node_id

Type:	string
Displayed:	True

The node ID of the Safeguard for Privileged Sessions machine

Origin

Name:	Origin
Search field:	origin
Type:	string
Displayed:	True

The source from where One Identity Safeguard for Privileged Sessions (SPS) received this session: sessions recorded by SPS, sessions recorded by and fetched from One Identity Safeguard for Privileged Passwords, or logs for sessions built from log data.

Protocol

Name:	Protocol
Search field:	protocol
Type:	enum
Displayed:	True

The protocol used in the session: Citrix ICA, HTTP, RDP, SSH, Telnet (including TN3270 and TN5250), MSSQL or VNC.

Possible values:

- *HTTP*: HTTP
- *ICA*: ICA
- *RDP*: RDP
- *SSH*: SSH
- *TELNET*: TELNET
- *VNC*: VNC
- *MSSQL*: MSSQL

Appliance id

Name:	Appliance id
Search field:	vault.appliance_id
Type:	string
Displayed:	True

The appliance's id

Appliance name

Name:	Appliance name
Search field:	vault.appliance_name
Type:	string
Displayed:	True

The appliance's name

Access request type

Name:	Access request type
Search field:	vault.access_request_type
Type:	string
Displayed:	True

Access request type can be: SSH, RDP, Password

Asset partition id

Name:	Asset partition id
Search field:	vault.asset_partition_id
Type:	long
Displayed:	True

ID of asset partition which represents a collection of assets and accounts along with management configuration

Asset partition name

Name:	Asset partition name
Search field:	vault.asset_partition_name
Type:	string
Displayed:	True

Name of the asset partition which represents a collection of assets and accounts along with management configuration

Broker id

Name:	Broker id
Search field:	vault.broker_id
Type:	long
Displayed:	True

ID of the broker who made the access request

Broker name

Name:	Broker name
Search field:	vault.broker_name
Type:	string
Displayed:	True

The broker's name who made the access request

Account id

Name:	Account id
Search field:	vault.account_id
Type:	long
Displayed:	True

Database ID of the account being requested

Account name

Name:	Account name
Search field:	vault.account_name
Type:	string
Displayed:	True

Name of the account being requested

System id

Name:	System id
Search field:	vault.system_id
Type:	long
Displayed:	True

Database ID of the system that has been requested access to. Should be displayed as assetId

System name

Name:	System name
Search field:	vault.system_name
Type:	string
Displayed:	True

Name of the system that has been requested access to. Should be displayed as assetName

Ticket number

Name:	Ticket number
Search field:	vault.ticket_number
Type:	string
Displayed:	True

Number of the help desk ticket as required by policy

Reason name

Name:	Reason name
Search field:	vault.reason_name
Type:	string
Displayed:	True

Reason's name for why the access request is needed

Is emergency

Name:	Is emergency
Search field:	vault.is_emergency
Type:	boolean
Displayed:	True

True when the access request was submitted as being an emergency

Offline workflow

Name:	Offline workflow
Search field:	vault.offline_workflow
Type:	boolean
Displayed:	True

True when the access request is an offline workflow

Auto approved

Name:	Auto approved
Search field:	vault.auto_approved
Type:	date
Displayed:	True

Date when access request was auto-approved to see the workflow's life in a timeline

Emergency access granted

Name:	Emergency access granted
Search field:	vault.emergency_access_granted
Type:	date
Displayed:	True

Date when the emergency access request was granted to see the workflow's life in a timeline

Available

Name:	Available
Search field:	vault.available
Type:	date
Displayed:	True

Date when the request is available for access

Checked in

Name:	Checked in
Search field:	vault.checked_in
Type:	date
Displayed:	True

Date when the access request is checked-in to see the workflow's life in a timeline

Expired

Name:	Expired
Search field:	vault.expired
Type:	date
Displayed:	True

Date when the access request will expire

Created user id

Name:	Created user id
Search field:	vault.created.user.user_id
Type:	long
Displayed:	True

Database ID of the user who created the access request

Created user name

Name:	Created user name
Search field:	vault.created.user.user_name
Type:	string
Displayed:	True

Name of the user who made the access request

Created domain name

Name:	Created domain name
Search field:	vault.created.user.domain_name
Type:	string
Displayed:	True

Domain mame of the user who made the access request

Created user display name

Name:	Created user display name
Search field:	vault.created.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who made the access request

Created client ip address

Name:	Created client ip address
Search field:	vault.created.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who created the access request

Created comment

Name:	Created comment
Search field:	vault.created.comment
Type:	string
Displayed:	True

Comment for the created access request

Created timestamp

Name:	Created timestamp
Search field:	vault.created.timestamp
Type:	date
Displayed:	True

Date when the access request was created

Denied user id

Name:	Denied user id
Search field:	vault.denied.user.user_id
Type:	long
Displayed:	True

Database ID of the user who denied the access request

Denied user name

Name:	Denied user name
Search field:	vault.denied.user.user_name
Type:	string
Displayed:	True

Name of the user who denied the access request

Denied domain name

Name:	Denied domain name
Search field:	vault.denied.user.domain_name
Type:	string
Displayed:	True

The user's domain name who denied the access request

Denied user display name

Name:	Denied user display name
Search field:	vault.denied.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who denied the access request

Denied client ip address

Name:	Denied client ip address
Search field:	vault.denied.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who denied the access request

Denied comment

Name:	Denied comment
Search field:	vault.denied.comment
Type:	string
Displayed:	True

Comment made by approver to describe denial

Denied timestamp

Name:	Denied timestamp
Search field:	vault.denied.timestamp
Type:	date
Displayed:	True

Date when the access request was denied

Revoked user id

Name:	Revoked user id
Search field:	vault.revoked.user.user_id
Type:	long
Displayed:	True

Database ID of the user who revoked the access request

Revoked user name

Name:	Revoked user name
Search field:	vault.revoked.user.user_name
Type:	string
Displayed:	True

Username of the user who revoked the access request

Revoked domain name

Name:	Revoked domain name
Search field:	vault.revoked.user.domain_name
Type:	string
Displayed:	True

The user's domain name who revoked the access request

Revoked user display name

Name:	Revoked user display name
Search field:	vault.revoked.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who revoked the access request

Revoked client ip address

Name:	Revoked client ip address
Search field:	vault.revoked.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who revoked the access request

Revoked comment

Name:	Revoked comment
Search field:	vault.revoked.comment
Type:	string
Displayed:	True

Comment made by approver to describe the revoke

Revoked timestamp

Name:	Revoked timestamp
Search field:	vault.revoked.timestamp
Type:	date
Displayed:	True

Date when the access request was revoked

Closed user id

Name:	Closed user id
Search field:	vault.closed.user.user_id
Type:	long
Displayed:	True

User ID of user who closed the access request. Closing access request used by an admin when a review cannot be completed

Closed user name

Name:	Closed user name
Search field:	vault.closed.user.user_name
Type:	string
Displayed:	True

Username of the user who closed the access request. Closing access request used by an admin when a review cannot be completed

Closed domain name

Name:	Closed domain name
Search field:	vault.closed.user.domain_name
Type:	string
Displayed:	True

The user's domain name who closed the access request

Closed user display name

Name:	Closed user display name
Search field:	vault.closed.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who closed the access request

Closed client ip address

Name:	Closed client ip address
Search field:	vault.closed.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who closed the access request

Closed comment

Name:	Closed comment
Search field:	vault.closed.comment
Type:	string
Displayed:	True

Comment for the request

Closed timestamp

Name:	Closed timestamp
Search field:	vault.closed.timestamp
Type:	date
Displayed:	True

Date when the access request was closed

Reviewed user id

Name:	Reviewed user id
Search field:	vault.reviewed.user.user_id
Type:	long
Displayed:	True

Database ID of the user who reviewed the access request

Reviewed user name

Name:	Reviewed user name
Search field:	vault.reviewed.user.user_name
Type:	string
Displayed:	True

Username of the user who reviewed the access request

Reviewed domain name

Name:	Reviewed domain name
Search field:	vault.reviewed.user.domain_name
Type:	string
Displayed:	True

User's domain name who reviewed the access request

Reviewed user display name

Name:	Reviewed user display name
Search field:	vault.reviewed.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who reviewed the access request

Reviewed client ip address

Name:	Reviewed client ip address
Search field:	vault.reviewed.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who reviewed the access request

Reviewed comment

Name:	Reviewed comment
Search field:	vault.reviewed.comment
Type:	string
Displayed:	True

Comment made by reviewer to describe review

Reviewed timestamp

Name:	Reviewed timestamp
Search field:	vault.reviewed.timestamp
Type:	date
Displayed:	True

Date when the access request was reviewed

Approved user id

Name:	Approved user id
Search field:	vault.approved.user.user_id
Type:	long
Displayed:	True

Database ID of the user who approved the access request

Approved user name

Name:	Approved user name
Search field:	vault.approved.user.user_name
Type:	string
Displayed:	True

Username of the user who approved the access request

Approved domain name

Name:	Approved domain name
Search field:	vault.approved.user.domain_name
Type:	string
Displayed:	True

User's domain name who approved the access request

Approved user display name

Name:	Approved user display name
Search field:	vault.approved.user.user_display_name
Type:	string
Displayed:	True

Display name of the user who approved the access request

Approved client ip address

Name:	Approved client ip address
Search field:	vault.approved.user.client_ip_address
Type:	ip
Displayed:	True

IP address of the user who approved the access request

Approved comment

Name:	Approved comment
Search field:	vault.approved.comment
Type:	string
Displayed:	True

Comment made by approver to describe approval

Approved timestamp

Name:	Approved timestamp
Search field:	vault.approved.timestamp
Type:	date
Displayed:	True

Date when the access request was approved

Additional metadata

Name:	Additional metadata
Search field:	recording.additional_metadata
Type:	string
Displayed:	False

Data about the session recorded by the different plugins of One Identity Safeguard for Privileged Sessions, for example, when using an Authentication and Authorization plugin.

Recording Archive date

Name:	Recording Archive date
Search field:	recording.archive.date
Type:	date
Displayed:	True

The date when the connection was archived or cleaned up.

Recording Archive path

Name:	Recording Archive path
Search field:	recording.archive.path
Type:	string
Displayed:	True

The path where the audit trail was archived on the remote server.

Recording Archive policy

Name:	Recording Archive policy
Search field:	recording.archive.policy
Type:	string
Displayed:	True

The archive policy used to archive the audit trail.

Recording Archive server

Name:	Recording Archive server
Search field:	recording.archive.server
Type:	ip
Displayed:	True

The hostname or IP address of the remote server where the audit trail was archived.

Recording Archived

Name:	Recording Archived
Search field:	recording.archived
Type:	boolean
Displayed:	True

Shows if the data (metadata, audit trail) about the session was archived to a remote server.

Audit trail path

Name:	Audit trail path
Search field:	recording.audit_trail
Type:	string
Displayed:	False

The path to the audit trail file on One Identity Safeguard for Privileged Sessions. If One Identity Safeguard for Privileged Sessions has already archived the audit trail, see the Archive path field instead.

. If the session does not have an audit trail, this element is not used. To download the audit trail, see Replaying audit trails in your browser on page 839.

Audit trail download link

Name:	Audit trail download link
Search field:	trail_download_link
Type:	string
Displayed:	True

The download link to the audit trail file on One Identity Safeguard for Privileged Sessions.

Recording Authentication method

Name:	Recording Authentication method
Search field:	recording.auth_method
Type:	string
Displayed:	True

The authentication method used in the session.

Recording Channel policy

Name:	Recording Channel policy
Search field:	recording.channel_policy
Type:	string
Displayed:	True

The Channel policy applied to the session. Channel policy determines the channels permitted in the connection, and if the channel is audited or not. The Channel policy can restrict access based on IP address, user list, user group, or time policy.

You can find the list of channel policies for each protocol at the **<Protocol> Control > Channel Policies** page.

Commands available

Name:	Commands available
Search field:	recording.command_extracted
Type:	boolean
Displayed:	True

True if commands have been extracted from the session. The extracted commands are in the Events field.

Recording Connection policy

Name:	Recording Connection policy
Search field:	recording.connection_policy
Type:	string
Displayed:	True

The name of the Connection policy that handled the client's connection request.

This is the name displayed on the **<Protocol> Control > Connections** page of the SPS web interface, and in the `name` field of the Connection Policy object. You can find the list of connection policies for each protocol at the **<Protocol> Control > Connections** page.

Recording Connection policy ID

Name:	Recording Connection policy ID
Search field:	recording.connection_policy_id
Type:	string
Displayed:	True

The ID of the Connection policy that handled the client's connection request.

You can find the list of connection policies for each protocol at the **<Protocol> Control > Connections** page.

Recording Content reference ID

Name:	Recording Content reference ID
Search field:	recording.content_reference_id
Type:	long
Displayed:	True

The unique identifier for the session content search.

Deny Reason

Name:	Deny Reason
Search field:	recording.deny_reason
Type:	string
Displayed:	True

The failure reason in case of a DENY verdict sent by an AA plugin.

Recording Indexing status

Name:	Recording Indexing status
Search field:	recording.index_status
Type:	enum
Displayed:	True

Shows if the channel has been indexed.
Possible values:

- *CHANNEL_OPEN*: Session is active
- *INDEXED*: Session indexed
- *INDEXING_FAILED*: Session indexing failed
- *INDEXING_IN_PROGRESS*: Session indexing in progress
- *INDEXING_NOT_REQUIRED*: Session indexing not required
- *NOT_INDEXED*: Session is not indexed
- *NO_TRAIL*: Auditing not enabled
- *INDEXING_ABORTED*: Session indexing in progress was aborted

Has ZAC

Name:	Has ZAC
Search field:	recording.has_zac
Type:	boolean
Displayed:	False

Audit Content file is available for the session. This file allows the user to search the content of graphical sessions using the Safeguard Desktop Player.

Recording Network namespace

Name:	Recording Network namespace
Search field:	recording.network_id
Type:	string
Displayed:	True

The ID of the Linux network namespace where the session originated from.

Server local IP address

Name:	Server local IP address
Search field:	recording.server_local.ip
Type:	ip
Displayed:	True

The IP address of One Identity Safeguard for Privileged Sessions used in the server-side connection.

Server local name

Name:	Server local name
Search field:	recording.server_local.name
Type:	string
Displayed:	True

The hostname of One Identity Safeguard for Privileged Sessions used in the server-side connection. If the hostname is not available, this field contains the IP address of One Identity Safeguard for Privileged Sessions.

Recording Server local port

Name:	Recording Server local port
Search field:	recording.server_local.port
Type:	port
Displayed:	True

The port number of One Identity Safeguard for Privileged Sessions used in the server-side connection.

Recording Session ID

Name:	Recording Session ID
Search field:	recording.session_id
Type:	string
Displayed:	True

A globally unique string that identifies the session. Log messages related to the session contain this ID.

Target IP address

Name:	Target IP address
Search field:	recording.target.ip
Type:	ip
Displayed:	True

The client originally tried to access this IP address. This can differ from the destination address, for example, when One Identity Safeguard for Privileged Sessions is configured to redirect the connection. The address that the client actually connected to is in the Server address field.

Target name

Name:	Target name
Search field:	recording.target.name
Type:	string
Displayed:	True

The client originally tried to access this host. This can differ from the destination address, for example, when One Identity Safeguard for Privileged Sessions is configured to redirect the connection. The address that the client actually connected to is in the Server address field. If the hostname is not available, this field contains the IP address of the host.

Recording Target port

Name:	Recording Target port
Search field:	recording.target.port
Type:	port
Displayed:	True

The client originally tried to access this port. This can differ from the port of the destination server, for example, when One Identity Safeguard for Privileged Sessions is configured to redirect the connection. The port that the client actually connected to is in the Server port field.

Recording Verdict

Name:	Recording Verdict
Search field:	recording.verdict
Type:	enum
Displayed:	True

Indicates what One Identity Safeguard for Privileged Sessions decided about the session.

Possible values:

- *ACCEPT*: Accepted
- *ACCEPT_TERMINATED*: Terminated by a content policy
- *AUTH_FAIL*: Authentication failed
- *DENY*: Connection rejected
- *FAIL*: Connection timed out on the server

- *GW_AUTH_FAIL*: Gateway authentication failed
- *KEY_ERROR*: Hostkey mismatch
- *USER_MAPPING_FAIL*: Usermapping failed

Recording Window titles available

Name:	Recording Window titles available
Search field:	recording.window_title_extracted
Type:	boolean
Displayed:	True

True if window titles have been extracted from the session. The extracted window titles are in the Window title field.

Server IP

Name:	Server IP
Search field:	server.ip
Type:	ip
Displayed:	True

The IP address of the server that One Identity Safeguard for Privileged Sessions connected to. This address was the remote end of the server-side connection.

Server hostname

Name:	Server hostname
Search field:	server.name
Type:	string
Displayed:	True

The hostname of the server that One Identity Safeguard for Privileged Sessions connected to.

Server port

Name:	Server port
Search field:	server.port
Type:	port
Displayed:	True

The port number of the server that One Identity Safeguard for Privileged Sessions connected to.

Start time

Name:	Start time
Search field:	start_time
Type:	date
Displayed:	True

Date when the session was started.

Starting with SPS 5 LTS, the timestamp is in ISO 8601 format, for example, `2018-10-11T09:23:38.000+02:00`. In earlier versions, it was in UNIX timestamp format.

Gateway username

Name:	Gateway username
Search field:	user.gateway_username
Type:	string
Displayed:	True

The username used to authenticate on the One Identity Safeguard for Privileged Sessions gateway (that is, in the client-side connection). Sometimes it is also called client-side username.

Gateway username domain

Name:	Gateway username domain
Search field:	user.gateway_username_domain
Type:	string
Displayed:	True

The domain of the username used to authenticate on the One Identity Safeguard for Privileged Sessions gateway (that is, in the client-side connection).

User ID

Name:	User ID
Search field:	user.id
Type:	string
Displayed:	True

The ID of the user.

Username

Name:	Username
Search field:	user.name
Type:	string
Displayed:	True

This field contains the username, which was used by the user to authenticate to the remote server. Its value is the same as the gateway username when it is available, otherwise, it will be filled with the server username.

Name domain

Name:	Name domain
Search field:	user.name_domain
Type:	string
Displayed:	True

This field contains the domain of the username, which was used by the user to authenticate to the remote server. Its value is the same as the gateway domain when it is available, otherwise, it will be filled with the server domain.

Server username

Name:	Server username

Search field:	user.server_username
Type:	string
Displayed:	True

The username used to log in to the remote server. This username can differ from the client-side username if usermapping is used in the connection.

Server username domain

Name:	Server username domain
Search field:	user.server_username_domain
Type:	string
Displayed:	True

The domain of the username used to log in to the remote server.

Verdict

Name:	Verdict
Search field:	verdict
Type:	enum
Displayed:	True

Indicates what One Identity Safeguard for Privileged Sessions decided about the session. A session verdict that originates from log events or other external events.

Possible values:

- *ACCEPT*: Accepted
- *AUTH_FAIL*: Authentication failed
- *DENY*: Connection rejected
- *FAIL*: Connection timed out on the server
- *PENDING*: Connection is pending
- *TERMINATED*: Connection terminated

Channel is active

Name:	Channel is active
Search field:	channel.active
Type:	boolean
Displayed:	False

True if the session has not ended yet.

Application

Name:	Application
Search field:	channel.application
Type:	string
Displayed:	False

The name of the application accessed in a seamless Citrix ICA connection.

Audit stream ID

Name:	Audit stream ID
Search field:	channel.audit_stream_id
Type:	string
Displayed:	False

The identifier of the channel's audit stream. If the session does not have an audit trail, this element is not used.

Channel ID

Name:	Channel ID
Search field:	channel.channel_id
Type:	long
Displayed:	False

The unique ID of the channel.

Client X.509 Subject

Name:	Client X.509 Subject
Search field:	channel.client_x509_subject
Type:	string
Displayed:	False

The client's certificate in TELNET or VNC sessions. Available only if the 'Client-side transport security settings > Peer certificate validation' option is enabled in One Identity Safeguard for Privileged Sessions.

Executed commands

Name:	Executed commands
Search field:	channel.command
Type:	string
Displayed:	False

Lists the commands run in an SSH session.

Port-forward target IP

Name:	Port-forward target IP
Search field:	channel.connected.ip
Type:	ip
Displayed:	False

The traffic was forwarded to this IP address in Remote Forward and Local Forward channels.

Port-forward target name

Name:	Port-forward target name
Search field:	channel.connected.name
Type:	string
Displayed:	False

The traffic was forwarded to this host in Remote Forward and Local Forward channels. If the hostname is not available, this field contains the IP address of the host

Port-forward target port

Name:	Port-forward target port
Search field:	channel.connected.port
Type:	port
Displayed:	False

The traffic was forwarded to this port in Remote Forward and Local Forward channels.

Device name

Name:	Device name
Search field:	channel.device_name
Type:	string
Displayed:	False

The name or ID of the shared device (redirect) used in the RDP connection.

Channel duration

Name:	Channel duration
Search field:	channel.duration
Type:	long
Displayed:	False

The length of the channel (how long the channel lasted).

Dynamic channel

Name:	Dynamic channel
Search field:	channel.dynamic_channel
Type:	string
Displayed:	False

The name or ID of the dynamic channel opened in the RDP session.

Used with the `dynamic virtual` RDP channel type.

Channel end time

Name:	Channel end time
Search field:	channel.end_time
Type:	date
Displayed:	False

Date when the channel was closed.

Environment

Name:	Environment
Search field:	channel.environment
Type:	string
Displayed:	False

Date when the channel was closed.

Four-eyes authorizer

Name:	Four-eyes authorizer
Search field:	channel.four_eyes_authorizer
Type:	string
Displayed:	False

The username of the user who authorized the session. Available only if four-eyes authorization is required for the channel.

Four-eyes description

Name:	Four-eyes description
Search field:	channel.four_eyes_description
Type:	string
Displayed:	False

The description submitted by the authorizer of the session.

Channel originator IP address

Name:	Channel originator IP address
Search field:	channel.originator.ip
Type:	ip
Displayed:	False

The IP address of the host initiating the channel in Remote Forward and Local Forward channels. Note that this host is not necessarily the client or the server of the SSH connection.

Channel originator name

Name:	Channel originator name
Search field:	channel.originator.name
Type:	string
Displayed:	False

The hostname of the host initiating the channel in Remote Forward and Local Forward channels. Note that this host is not necessarily the client or the server of the SSH connection. If the hostname is not available, this field contains the IP address of the host.

Originator port

Name:	Originator port
Search field:	channel.originator.port
Type:	port
Displayed:	False

The number of the forwarded port in Remote Forward and Local Forward SSH channels.

Rule number

Name:	Rule number

Search field:	channel.rule_num
Type:	string
Displayed:	False

The number of the line in the Channel policy applied to the channel.

SCP path

Name:	SCP path
Search field:	channel.scp_path
Type:	string
Displayed:	False

Name and path of the file copied via SCP. Available only for SCP sessions (Session exec SCP SSH channels) if the Log file transfers to database option isenabled in the Channel Policy of the connection.

Channel start time

Name:	Channel start time
Search field:	channel.start_time
Type:	date
Displayed:	False

Date when the channel was started.

Subsystem name

Name:	Subsystem name
Search field:	channel.subsystem_name
Type:	string
Displayed:	False

Name of the SSH subsystem used in the channel.

Channel type

Name:	Channel type
Search field:	channel.type
Type:	enum
Displayed:	False

Type of the channel.

Possible values:

- *#drawing*: Drawing
- *CTXCAM*: Audio
- *CTXCDM*: Drive
- *CTXCLIP*: Clipboard
- *CTXCOM1*: Printer (COM1)
- *CTXCOM2*: Printer (COM2)
- *CTXCPM*: Printer Spooler
- *CTXFLSH*: HDX Mediastream
- *CTXLPT1*: Printer (LPT1)
- *CTXLPT2*: Printer (LPT2)
- *CTXSCRD*: Smartcard
- *CTXTW*: Drawing (Thinwire)
- *CTXTWI*: Seamless
- *CTXUSB*: USB
- *SPDBRS*: Speedbrowse
- *auth-agent*: Agent
- *cliprdr*: Clipboard
- *custom*: Custom
- *direct-tcpip*: Local forward
- *drawing*: Drawing
- *drdynvc*: Dynamic virtual channel
- *forwarded-tcpip*: Remote forward
- *http*: HTTP
- *rdpdr*: Redirects
- *rdpdr-disk*: Disk redirect
- *rdpdr-parallel*: Parallel redirect
- *rdpdr-printer*: Printer redirect
- *rdpdr-scard*: SCard redirect

- *rdpdr-serial*: Serial redirect
- *rdpsnd*: Sound
- *seamrdp*: Seamless
- *session-exec*: Session exec
- *session-exec-scp*: Session exec SCP
- *session-shell*: Session shell
- *session-subsystem*: Session subsystem
- *session-subsystem-sftp*: Session SFTP
- *telnet*: Telnet
- *vnc*: VNC
- *websocket*: WebSocket
- *x11*: X11 forward

Channel verdict

Name:	Channel verdict
Search field:	channel.verdict
Type:	enum
Displayed:	False

Indicates what One Identity Safeguard for Privileged Sessions decided about the channel. Possible values:

- *ACCEPT*: Accepted
- *DENY*: Denied
- *FOUR_EYES_DEFERRED*: Waiting for remote username
- *FOUR_EYES_ERROR*: Internal error during four-eyes authorization
- *FOUR_EYES_REJECT*: Four-eyes authorization rejected
- *FOUR_EYES_TIMEOUT*: Four-eyes authorization timed out

Event Action

Name:	Event Action
Search field:	event.action
Type:	string
Displayed:	False

The command line without prompt in commands

Channel ID

Name:	Channel ID
Search field:	event.channel_id
Type:	string
Displayed:	False

The id of the channel the event belongs to.

Event content

Name:	Event content
Search field:	event.content
Type:	string
Displayed:	False

The command run, or the window title detected in the channel (for example, ls, exit, or Firefox).

Protocol details

Name:	Protocol details
Search field:	event.details
Type:	string
Displayed:	False

The details of the protocol used for the operation.

Event ID

Name:	Event ID
Search field:	event.event_id
Type:	string
Displayed:	False

The identifier of the vault event.

Operation

Name:	Operation
Search field:	event.operation
Type:	string
Displayed:	False

The type of the operation that occurred, for example, Create file (in the case of FTP) or GET (in the case of HTTP).

Path

Name:	Path
Search field:	event.path
Type:	string
Displayed:	False

The path (if any) used by the operation that occurred.

Event ID

Name:	Event ID
Search field:	event.record_id
Type:	long
Displayed:	False

The identifier of the event within the audit trail (.zat file).

Response code

Name:	Response code
Search field:	event.response_code
Type:	long
Displayed:	False

The status code of the protocol response (if any) returned.

Event date

Name:	Event date
Search field:	event.time
Type:	date
Displayed:	False

The date when the event happened.

Event type

Name:	Event type
Search field:	event.type
Type:	string
Displayed:	False

The type of the event, for example, command, screen_content, window_title.

Alert type

Name:	Alert type
Search field:	alert.alert_type
Type:	enum
Displayed:	False

The type of the alert.

Possible values:

- *adp.event.command*: A command entered in SSH or Telnet.
- *adp.event.screen.content*: Alert triggered by the screen content.
- *adp.event.screen.creditcard*: Credit card numbers detected. Displayed only as an alert, not visible in the events.
- *adp.event.screen.windowtitle*: The title of the window in graphic protocols.

Channel ID

Name:	Channel ID
Search field:	alert.channel_id
Type:	string
Displayed:	False

The id of the channel the alert belongs to.

Matched regexp on action

Name:	Matched regexp on action
Search field:	alert.matched_action
Type:	string
Displayed:	False

The regular expression that matched the command line without prompt

Matched content

Name:	Matched content
Search field:	alert.matched_content
Type:	string
Displayed:	False

The content the alert matched.

Matched regexp

Name:	Matched regexp
Search field:	alert.matched_regexp
Type:	string
Displayed:	False

The regular expression that matched the content.

Alert ID

Name:	Alert ID
Search field:	alert.record_id
Type:	long
Displayed:	False

The identifier of the alert within the audit trail (.zat file).

Rule name

Name:	Rule name
Search field:	alert.rule_name
Type:	string
Displayed:	False

The name of the content policy rule.

Alert time

Name:	Alert time
Search field:	alert.time
Type:	date
Displayed:	False

The timestamp of the alert.

From API

Name:	From API
Search field:	trail_download.from_api
Type:	boolean
Displayed:	False

The audit trail downloaded via API or not.

Trail download ID

Name:	Trail download ID
Search field:	trail_download.id
Type:	string
Displayed:	False

The ID of an audit trail download event.

Download ip

Name:	Download ip
Search field:	trail_download.ip_address
Type:	ip
Displayed:	False

The ip address from where the download is requested.

Download time

Name:	Download time
Search field:	trail_download.time
Type:	date
Displayed:	False

The exact time when the user downloaded the audit trail file.

Downloader username

Name:	Downloader username
Search field:	trail_download.username
Type:	string
Displayed:	False

The name of user who downloaded the audit trail of the session.

Commands indexed

Name:	Commands indexed
Search field:	indexer_info.config.command.enabled
Type:	boolean
Displayed:	False

True if commands were extracted while indexing the session.

Keyboard buffering interval

Name:	Keyboard buffering interval
Search field:	indexer_info.config.keyboard.buffer_interval
Type:	double
Displayed:	False

The buffering interval in milliseconds used when extracting keyboard events while indexing the session.

Keyboard extracted

Name:	Keyboard extracted
Search field:	indexer_info.config.keyboard.enabled
Type:	boolean
Displayed:	False

True if keyboard events were extracted while indexing the session.

Mouse buffering interval

Name:	Mouse buffering interval
Search field:	indexer_info.config.mouse.buffer_interval
Type:	double
Displayed:	False

The buffering interval in milliseconds used when extracting mouse events while indexing the session.

Mouse extracted

Name:	Mouse extracted
Search field:	indexer_info.config.mouse.enabled
Type:	boolean
Displayed:	False

True if mouse events were extracted while indexing the session.

Near real-time indexing

Name:	Near real-time indexing
Search field:	indexer_info.config.near_realtime
Type:	boolean
Displayed:	False

True if indexing this session was done near real-time (when the session was still active).

OCR languages

Name:	OCR languages
Search field:	indexer_info.config.ocr_languages
Type:	string
Displayed:	False

The language configuration for optical character recognition used when indexing the session.

Screen content indexed

Name:	Screen content indexed
Search field:	indexer_info.config.screen.enabled
Type:	boolean
Displayed:	False

True if screen content was extracted while indexing the session.

OCR tradeoff

Name:	OCR tradeoff
Search field:	indexer_info.config.screen.omnipage_trade_off
Type:	string
Displayed:	False

The tradeoff used for optical character recognition when extracting screen content while indexing the session.

Titles indexed

Name:	Titles indexed
Search field:	indexer_info.config.title.enabled
Type:	boolean
Displayed:	False

True if window titles were extracted while indexing the session.

Indexing error

Name:	Indexing error
Search field:	indexer_info.error.message
Type:	string
Displayed:	False

The reason why indexing failed

Indexing cpu time

Name:	Indexing cpu time
Search field:	indexer_info.statistics.cpu_time
Type:	long
Displayed:	False

The CPU time that indexing this session took in milliseconds.

Indexing duration

Name:	Indexing duration
Search field:	indexer_info.statistics.duration
Type:	long
Displayed:	False

The duration of time that indexing this session took in milliseconds.

Indexing start time

Name:	Indexing start time
Search field:	indexer_info.statistics.start_time
Type:	date
Displayed:	False

The time and date when indexing this session started.

Indexing status

Name:	Indexing status
Search field:	indexer_info.status
Type:	string
Displayed:	False

Shows if the channel has been indexed successfully or not.

Indexer ADP version

Name:	Indexer ADP version
Search field:	indexer_info.version.adp
Type:	string
Displayed:	False

The version of the audit data processor used for indexing the session

Indexer version

Name:	Indexer version
Search field:	indexer_info.version.worker
Type:	string
Displayed:	False

The version of the indexer worker used for indexing the session

ZAC created

Name:	ZAC created
Search field:	indexer_info.config.zac.enabled
Type:	boolean
Displayed:	False

True if an Audit Content file was created while indexing the session.

Screen content

Name:	Screen content
Search field:	screen.content
Type:	string
Displayed:	False

Text that appeared on the screen in the session.

Channel id in trail

Name:	Channel id in trail
Search field:	screen.channel_id_in_trail
Type:	long
Displayed:	False

The ID of the channel where this content appeared. To check the channel ID (channel_id), select a session and click details. Navigate to details > Channels and click the channel type.

Screen content creation time

Name:	Screen content creation time
Search field:	screen.time
Type:	date
Displayed:	False

The creation time of the indexed screen content.

Screen content ID

Name:	Screen content ID
Search field:	screen.id
Type:	string
Displayed:	False

The ID of a screen content event.

trail_download

From API

Name:	From API
Search field:	from_api
Type:	boolean
Displayed:	True

The audit trail downloaded via API or not.

Trail download ID

Name:	Trail download ID
Search field:	id
Type:	string
Displayed:	True

The ID of an audit trail download event.

Download ip

Name:	Download ip
Search field:	ip_address
Type:	ip
Displayed:	True

The ip address from where the download is requested.

Download time

Name:	Download time
Search field:	time
Type:	date
Displayed:	False

The exact time when the user downloaded the audit trail file.

Downloader username

Name:	Downloader username
Search field:	username
Type:	string
Displayed:	True

The name of user who downloaded the audit trail of the session.

Searching in the contents of audit trails

NOTE: This feature is available only if auditing and content indexing was requested for the connection.

For more information, see Configuring the internal indexer on page 675.

You can search in the contents of the audit trails as follows:

- **From your browser**: Use this method to find all the sessions containing your search query.

  Enter the **screen.content: expression** search query in the **Search query** field. For example: **screen.content="exit"**. The search returns all the sessions where **exit** was on the screen.

- **From the Safeguard Desktop Player application**: Use this method to find the exact location of the search query within a specific audit trail.

  Download the relevant audit trail, open it in the Safeguard Desktop Player application, and use the Search feature. You can also search in the contents of the audit trails for trails of graphical sessions created and indexed with One Identity Safeguard for Privileged Sessions (SPS) 6.0.

There are various ways you can refine your content query, you can:

- use wildcards
- use boolean expressions
- search in the commands of terminal connections (for example, `command:"sudo su"`)
- search in the window titles of graphical connections (for example, `title:settings`)

Search query examples

The following sections provide examples for different search queries.

- For examples of exact matches, see Searching for exact matches on page 800.
- For examples of using boolean operators to combine search keywords, see Combining search keywords on page 801.
- For examples of wildcard searches, see Using wildcard searches on page 802.
- For examples of searching with special characters, see Searching for special characters on page 804.
- For examples of fuzzy search that finds words with similar spelling, see Searching for fuzzy matches on page 806.
- For examples of proximity search to find words that appear within a special distance, see Proximity search on page 806.
- For examples of adjusting the relevance of a search term, see Adjusting the relevance of search terms on page 806.

For details on how to use more complex keyphrases that are not covered in this guide, see the Apache Lucene documentation.

Searching for exact matches

By default, One Identity Safeguard for Privileged Sessions (SPS) searches for keywords as whole words and returns only exact matches. Note that if your search keywords include special characters, you must escape them with a backslash (\) character. For details on special characters, see Searching for special characters on page 804. The following characters are special characters: `+ - & | ! ( ) { } [ ] ^ " ~ * ? : \ /`

Example: Searching for exact matches

Search expression	`example`
Matches	example
Does not match	examples example.com query-by-example exam

To search for a string that includes a backslash characters, for example, a Windows path, use two backslashes (`\\`).

Search expression	`C\:\\Windows`
Matches	C:\Windows

Combining search keywords

You can use boolean operators – `AND`, `OR`, `NOT`, and `+` (required), – to combine search keywords. More complex search expressions can also be constructed with parentheses. If you enter multiple keywords,

Example: Combining keywords in search

Search expression	`keyword1 AND keyword2`
Matches	(returns hits that contain both keywords)

Search expression	`keyword1 OR keyword2`
Matches	(returns hits that contain at least one of the keywords)

Search expression	`keyword1 NOT keyword2`
Matches	(returns hits that contain the first phrase, but not the second)

Search expression	`+keyword1 keyword2`
Matches	(returns hits that contain keyword1, and may contain keyword2)

To search for expressions that can be interpreted as boolean operators (for example: **AND**), use the following format: **"AND"**.

Example: Using parentheses in search

Use parentheses to create more complex search expressions:

Search expression	`(keyword1 OR keyword2) AND keyword3`
Matches	(returns hits that contain either keyword1 and keyword3, or keyword2 and keyword3)

Using wildcard searches

You can use the ? and * wildcards in your search expressions.

Example: Using wildcard ? in search

The ? (question mark) wildcard means exactly one arbitrary character. Note that it does not work for finding non-UTF-8 or multibyte characters. If you want to search for these characters, the expression ?? might work, or you can use the * wildcard instead.

You cannot use a * or ? symbol as the first character of a search.

Search expression	`example?`
Matches	example1 examples example?
Does not match	example.com

	example12 query-by-example
Search expression	`example??`
Matches	example12
Does not match	example.com example1 query-by-example

Example: Using wildcard * in search

The * wildcard means 0 or more arbitrary characters. It finds non-UTF-8 and multibyte characters as well.

Search expression	`example*`
Matches	example examples example.com
Does not match	query-by-example example*

Example: Using combined wildcards in search

Wildcard characters can be combined.

Search expression	`ex?mple*`
Matches	example1 examples example.com exemple.com

	example12
Does not match	exmples
	query-by-example

Searching for special characters

To search for the special characters, for example, question mark (?), asterisk (*), backslash (\) or whitespace ( ) characters, you must prefix these characters with a backslash (\). Any character after a backslash is handled as character to be searched for. The following characters are special characters: `+ - & | ! ( ) { } [ ] ^ " ~ * ? : \ /`

Example: Searching for special characters

To search for a special character, use a backslash (\).

Search expression	`example\?`
Matches	example?
Does not match	examples
	example1

To search for a string that includes a backslash characters, for example, a Windows path, use two backslashes (\\).

Search expression	`C\:\\Windows`
Matches	C:\Windows

To search for a string that includes a slash character, for example, a UNIX path, you must escape the every slash with a backslash (\/).

Search expression	`\/var\/log\/messages`
Matches	/var/log/messages

Search expression	`\(1\+1\)\:2`
Matches	(1+1):2

Searching in commands and window titles

For terminal connections, use the `command:` prefix to search only in the commands (excluding screen content). For graphical connections, use the `title:` prefix to search only in the window titles (excluding screen content). To exclude search results that are commands or window titles, use the following format: `keyword AND NOT title:[* TO *]`.

You can also combine these search queries with other expressions and wildcards, for example, `title:properties AND gateway`.

Example: Searching in commands and window titles

Search expression	`command:sudo su`
Matches	`sudo su` as a terminal command
Does not match	`sudo su` in general screen content

Search expression	`title:settings`
Matches	`settings` appearing in the title of an active window
Does not match	`settings` in general screen content

To find an expression in the screen content and exclude search results from the commands or window titles, see the following example.

Search expression	`properties AND NOT title:[* TO *]`
Matches	`properties` appearing in the screen content, but not as a window title.
Does not match	`properties` in window titles.

You can also combine these search filters with other expressions and wildcards.

Search expression	`title:properties AND gateway`
Matches	A screen where `properties` appears in the window title, and `gateway` in the screen content (or as part of the window title).
Does not match	Screens where both `properties` and `gateway` appear, but `properties` is not in the window title.

Searching for fuzzy matches

Fuzzy search uses the tilde ~ symbol at the end of a single keyword to find hits that contain words with similar spelling to the keyword.

Example: Searching for fuzzy matches

Search expression	`roam~`
Matches	roams foam

Proximity search

Proximity search uses the tilde ~ symbol at the end of a phrase to find keywords from the phrase that are within the specified distance from each other.

Example: Proximity search

Search expression	`keyword1 keyword2 ~10`
Matches	(returns hits that contain keyword1 and keyword2 within 10 words from each other)

Adjusting the relevance of search terms

By default, every keyword or phrase of a search expression is treated as equal. Use the caret ^ symbol to make a keyword or expression more important than the others.

Example: Adjusting the relevance of search terms

Search expression	`keyword1^4 keyword2`
Matches	(returns hits that contain keyword1 and keyword2, but keyword1 is 4-times more relevant)

Search expression	`keyword1^5 keyword2`
Matches	(returns hits that contain keyword1 and keyword2, but keyword1 is 5-times more relevant)

Audit trail downloads information on the Search interface

If you want to find out if the audit trail file of a relevant indexed session has already been downloaded (or you are interested in the details of the session's audit trail downloads), the **Details** tab will provide information.

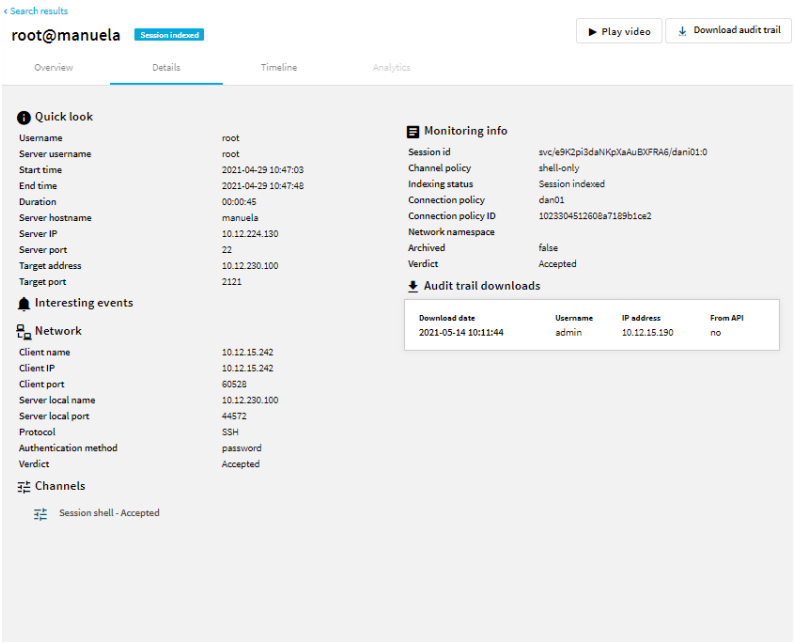
Prerequisites

To log audit trail downloads of a certain connection, enable **Log audit trail downloads** in the **Connections** menu of the connection.

To display audit trail downloads information about the relevant indexed session

1. Navigate to **Audit**, select the **Sessions** tab, and find the relevant indexed session.
2. Click **Session Details**.

Figure 266: Audit > Sessions > Session Details — The indexed session's available information displayed on the Details tab



If no audit trail file has been downloaded for the relevant indexed session yet, the **Details** tab displays: "The audit trail for this session has never been downloaded."

3. If you want to download an audit trail file for the session, click **Download audit trail**. In this case, the displayed information will contain information about your current session.

   If a downloaded audit trail file already exists for the relevant indexed session, the **Details** tab will display the following information:

   - **Download date**: The exact time when the user downloaded the audit trail file.
   - **Username**: The username that was used to download the audit trail file for the session.
   - **IP address**: The IP address from where the audit trail download was requested.
   - **From API**: Indicates if the audit trail file was downloaded through API or not.

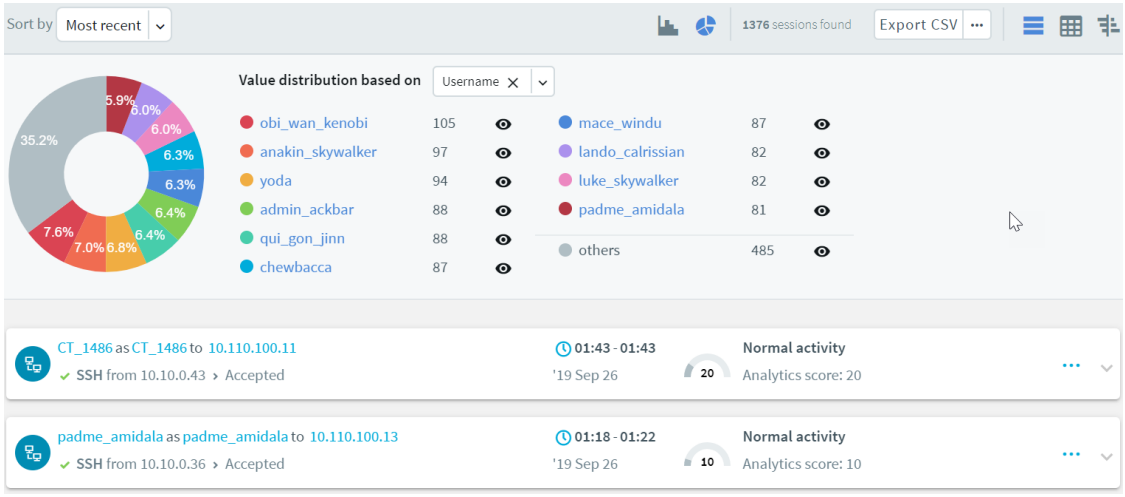
Displaying statistics on search results

You can quickly sort and visualize the distribution of the sessions based on their various metadata, for example, username, server address, and so on.

To display statistics on search results

1. Click the  icon.
2. Select the type of metadata you want to create statistics on from the **Value distribution based on** field, for example, select **Username** to display sessions based on username.

Figure 267: Audit > Sessions — Displaying statistics



3. To exclude items from the pie chart, click the  icon next to the metadata you want to exclude.

   For example, if you want to exclude results by a user called **testbot**, select the  icon next to the item.

Figure 268: Audit > Sessions — Excluding items from the pie chart



The pie chart now does not display results for the excluded item. The percentages always add up to 100%.

You can continue to restrict or refine your search results and view statistics as required.

Analyzing data using One Identity Safeguard for Privileged Analytics

One Identity Safeguard for Privileged Sessions (SPS) integrates data from SPS to use as the basis of user behavior analysis. SPA uses machine learning algorithms to scrutinize behavioral characteristics (using data from SPS), and generates user behavior profiles for each individual privileged user. SPA compares actual user activity to user profiles in real time, with profiles being continually adjusted using machine learning. When SPA detects unusual activity, this is indicated on the user interface of SPS in the form of high scores and visualized insight.

Prerequisites

Make sure that you have session data from network traffic that:

- contains real, unique user names linked to users other than root/administrator or a shared account

  To check this, navigate to **Sessions**, and check whether the **Username** column contains data. This is important, because session data will be linked to users.

  If you do not have unique user names in your session data, review your authentication settings and consult with the One Identity Professional Services team to learn about your options to tie accounts to users.
- has commands extracted (using lightweight or full indexing, or in real-time through content policies)

  For more information on how to configure indexing and include commands in the scope of indexing, see *Indexing audit trails* in the *Administration Guide*.

  For details on how to configure real-time command extraction using a content policy, see *Creating a new content policy* in the *Administration Guide*.
- has keystrokes extracted (using lightweight or full indexing, or in real-time through content policies)

  The minimum required amount of data for reliable insight is 5 sessions, with approximately 200 keystrokes each.

  For more information on how to configure indexing and include typing biometrics in the scope of indexing, see *Indexing audit trails* in the *Administration Guide*.

  For details on how to configure real-time extraction of keystroke-related data using a content policy, see *Creating a new content policy* in the *Administration Guide*.
- has pointing device (mouse) biometrics extracted (using lightweight or full indexing, or in real-time through content policies)

For more information on how to configure indexing and include pointing device biometrics in the scope of indexing, see *Indexing audit trails* in the *Administration Guide*.

For more information on how to configure real-time extraction of pointing device-related data using a content policy, see *Creating a new content policy* in the *Administration Guide*.

- has window titles extracted (using lightweight or full indexing, or in real-time through content policies)

  For more information on how to configure indexing and include window titles in the scope of indexing, see *Indexing audit trails* in the *Administration Guide*.

  For more information on how to configure real-time window title extraction using a content policy, see *Creating a new content policy* in the *Administration Guide*.

SPA limitations

SPS used in combination with SPA currently has the following limitations:

- SPA requires at least 12GB RAM to operate. If you are interested in upgrading your appliance, contact our Support Team.
- SPA requires a lot of computation, which can increase the load on SPS:
  - The keystroke algorithm requires significantly more resources than the other algorithms. As a result One Identity recommends that you start analyzing data using the algorithms that require less resources.
  - Before you start using SPA, make sure that at least half the capacity of SPS is available.
- SPA only analyzes audit trails and SPS metadata, it does not analyze log data.

Algorithm limitations

Generating more extracted data from sessions results in better analytics score calculations. However, you must meet certain algorithm-specific limitations. These limitations are necessary to ensure that the appliance stays operational even under high load.

Algorithms with limits:

- window title
- host login
- keystroke
- mouse
- command
- FIS

Window title and host login

SPS cannot build the baseline for these algorithms if the number of users observed by the appliance is over 4096.

Keystroke, mouse and command

These algorithms generate extracted data from every analyzed session.

SPS cannot build the baseline for a user for these algorithms if the collected extraction data amount is over 128MB in the last 90 days. Since these algorithms are very complex, it is not possible to estimate how much activity is needed to reach this threshold.

FIS

SPS cannot build the FIS baseline if there are more than 20000 recorded sessions in the last 90 days for a user.

The following describes how to analyze data using One Identity Safeguard for Privileged Analytics.

To start using SPA

1. **Start getting scores.**

   Scoring for sessions

   Scoring happens in real time, meaning that as soon as new data (even data from an ongoing session) is available, SPA immediately scores it.

   > TIP: When data is not immediately available to you and you are unable to wait until sufficient amount of data comes in from production traffic, you can resort to manually reindexing historical sessions. For more information, see *Reindex historical sessions* in the *Safeguard for Privileged Analytics Configuration Guide*.

   Scores represent an aggregated amount. Session data is scored by multiple algorithms independent from each other. Scores given by individual algorithms are aggregated to create a single score.

   For detailed instructions on how to configure SPA, see *Safeguard for Privileged Analytics Configuration Guide*.

   Scoring for users

   The goal of the algorithm is to create a score for the user to represent recent activities. The algorithm does this by averaging recent event scores and weighing the top 3 highest scores and taking in consideration the elapsed time. The user score is calculated hourly and weighs more recent activities with a bias.

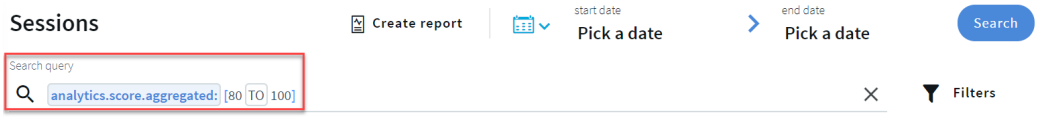
2. **Search for sessions with high scores.**
   a. Go to **Sessions**.

Sessions are displayed sorted by date. For ongoing sessions, the Search interface is updated in real time to always display the most up-to-date information.

b. In the **Search query** field, type `analytics.score.aggregated: [80 TO 100]`, and click **Search**.

   A score between `80` and `100` indicates unusual user behavior.

Figure 269: Searching for sessions with unusual user behavior using a search query



Results that show sessions with high scores are displayed.

Figure 270: Sessions with high scores — table view

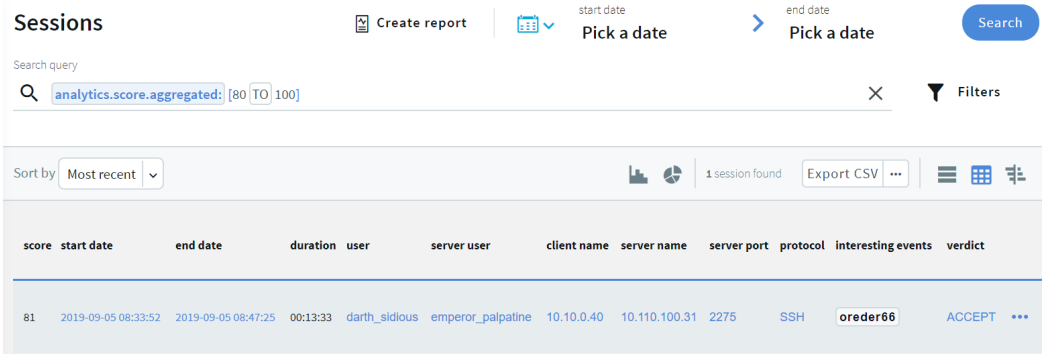
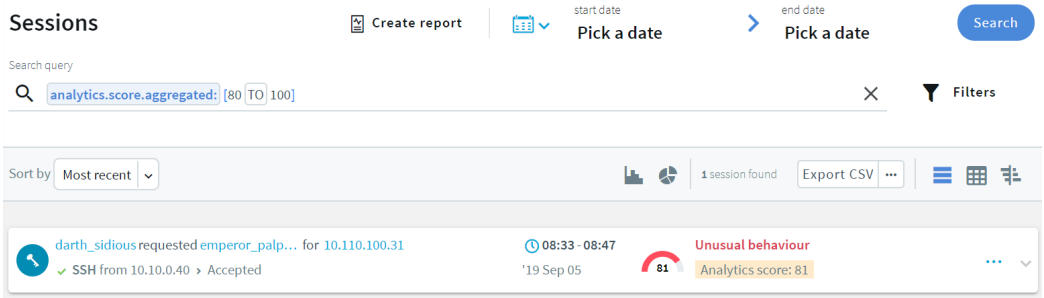


Figure 271: Sessions with high scores — card view



3. **Alternatively, search for scripted sessions.**

   In the **Search query** field, type `analytics.scripted:true`, and click **Search**.

4. **View the details of a session.**

   To view details of a session, click ••• .

5. **View session analytics.**

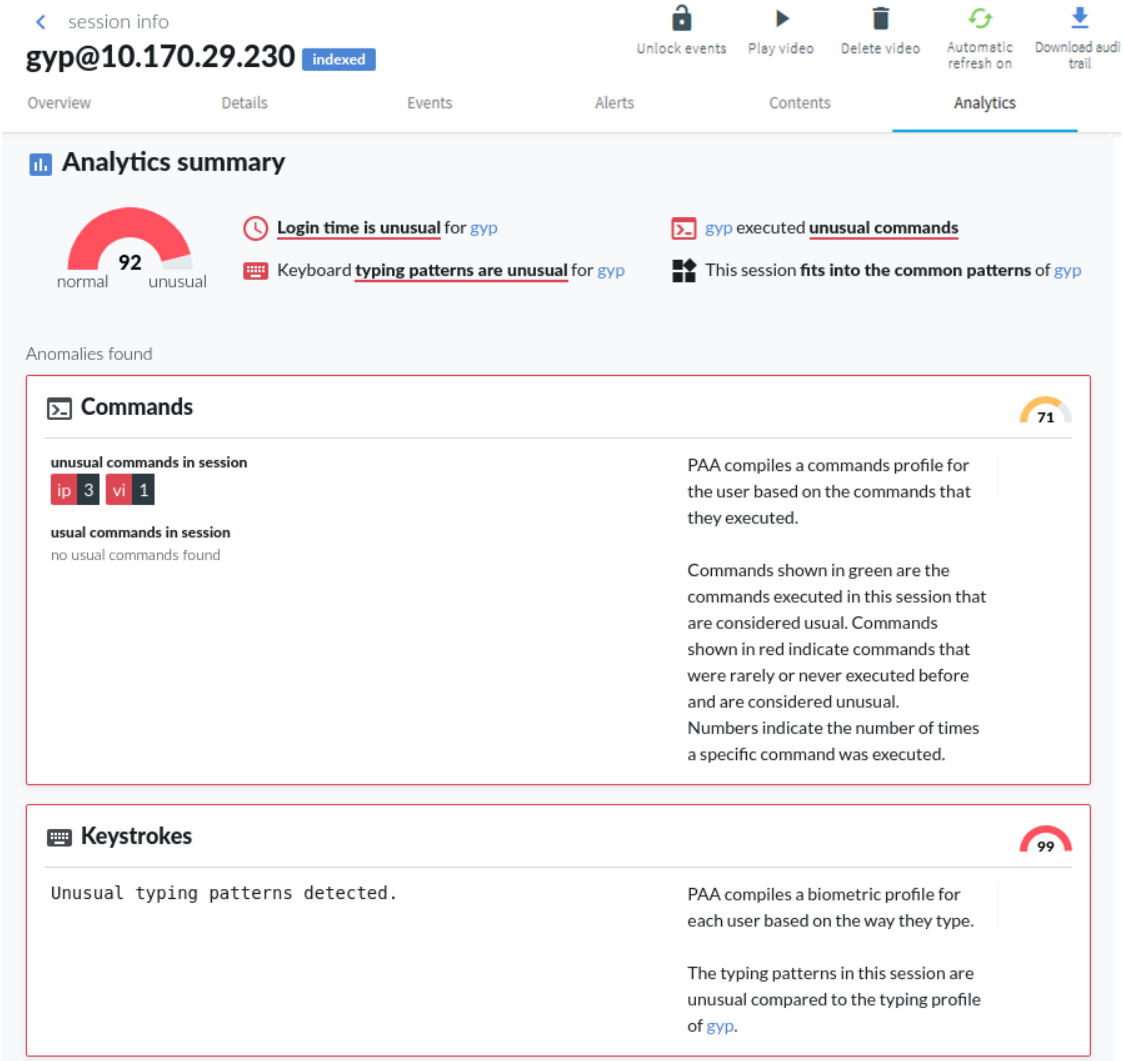
   Click the **Analytics** tab.

   The top of the page displays a summary of key insights about the session, such as:

   - The aggregated score (indicated by a gauge). The following color codes are used:
     - Scores between `80-100` indicate unusual behavior, their color code is red.
     - Scores between `70-79` indicate behavior that might require further analysis and attention, their color code is amber.
     - Scores between `0-69` indicate normal behavior, their color code is gray.
   - A one-sentence summary of each algorithm's verdict about the session and user behavior.

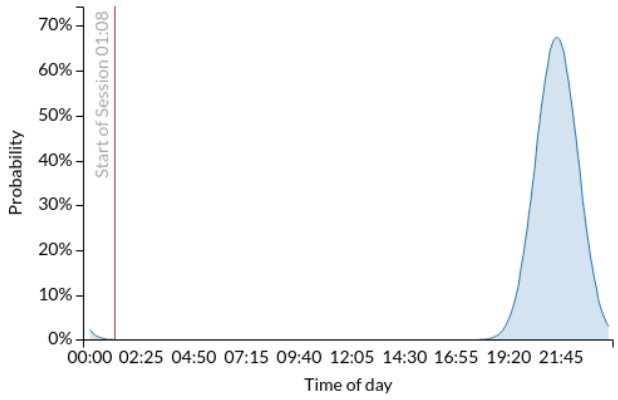
   The **Anomalies found** and **Normal behavior** sections of the page display detailed analyses provided by each of the configured algorithms. This includes short information on how a particular algorithm works and how to read the visualized insight, as well as scores given by the individual algorithms.

Figure 272: Audit > Sessions — Viewing details on the Analytics tab: Anomalies found



Login time

83



PAA compiles a login time profile for the user based on the exact time in a day when the user logs in.
The area shows the probability of login by this user at a certain time.

Login time is unusual for gyp , the probability of login around 01:02 is 2%.

Normal behavior

Frequent item set

39

This session **fits into the common patterns** of gyp

The FIS algorithm examines multiple attributes of sessions and attempts to find values that frequently appear together, forming a set. Using this information, the algorithm is able to discover patterns in user behavior.

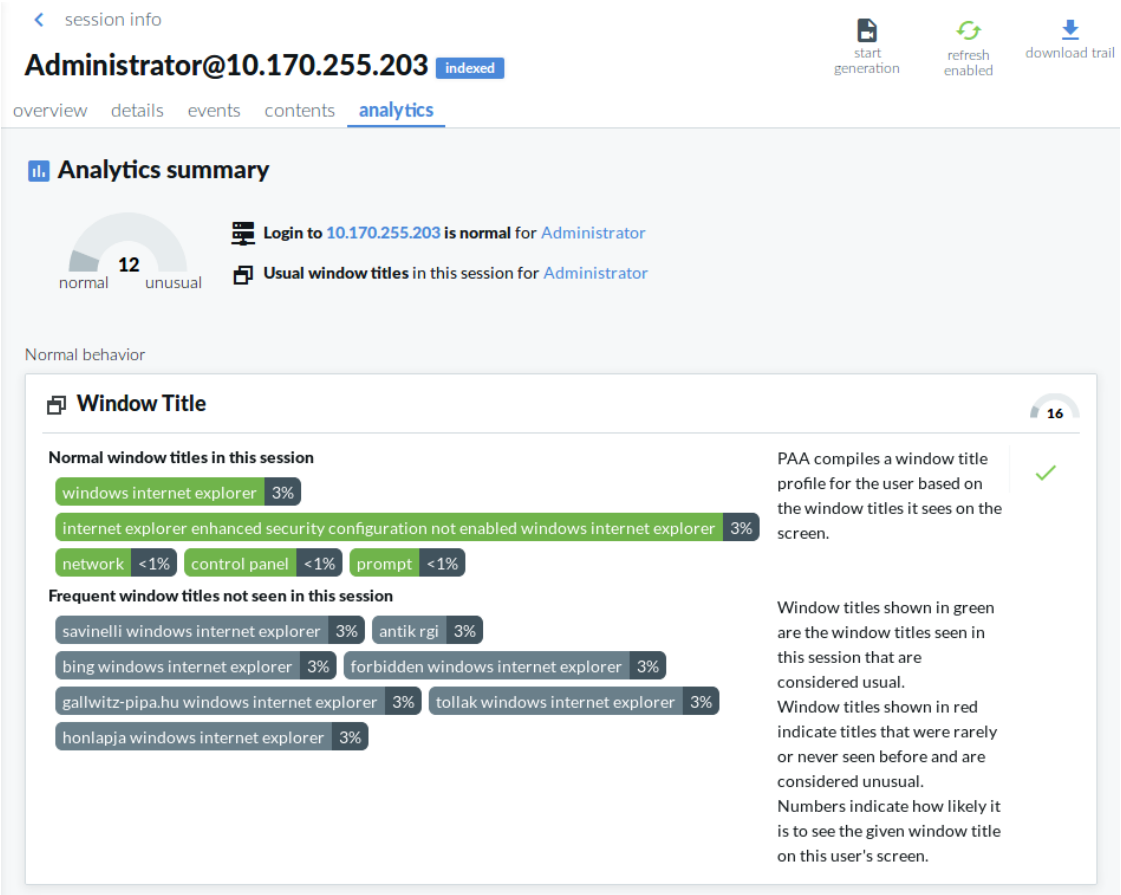
Not scored

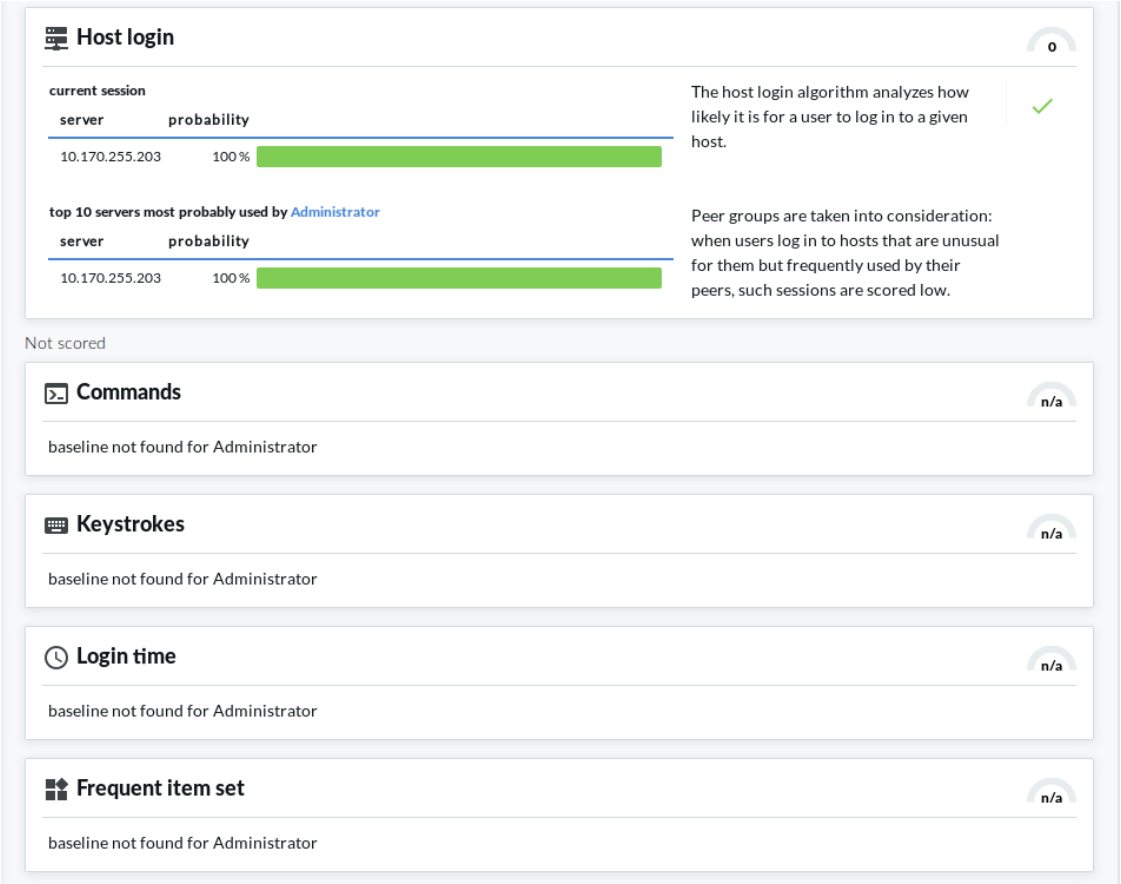
Host login

n/a

This user regularly logs in to the great majority of hosts. We cannot assign a score to this session.

Figure 273: Sessions — Viewing details on the Analytics tab: Normal behavior

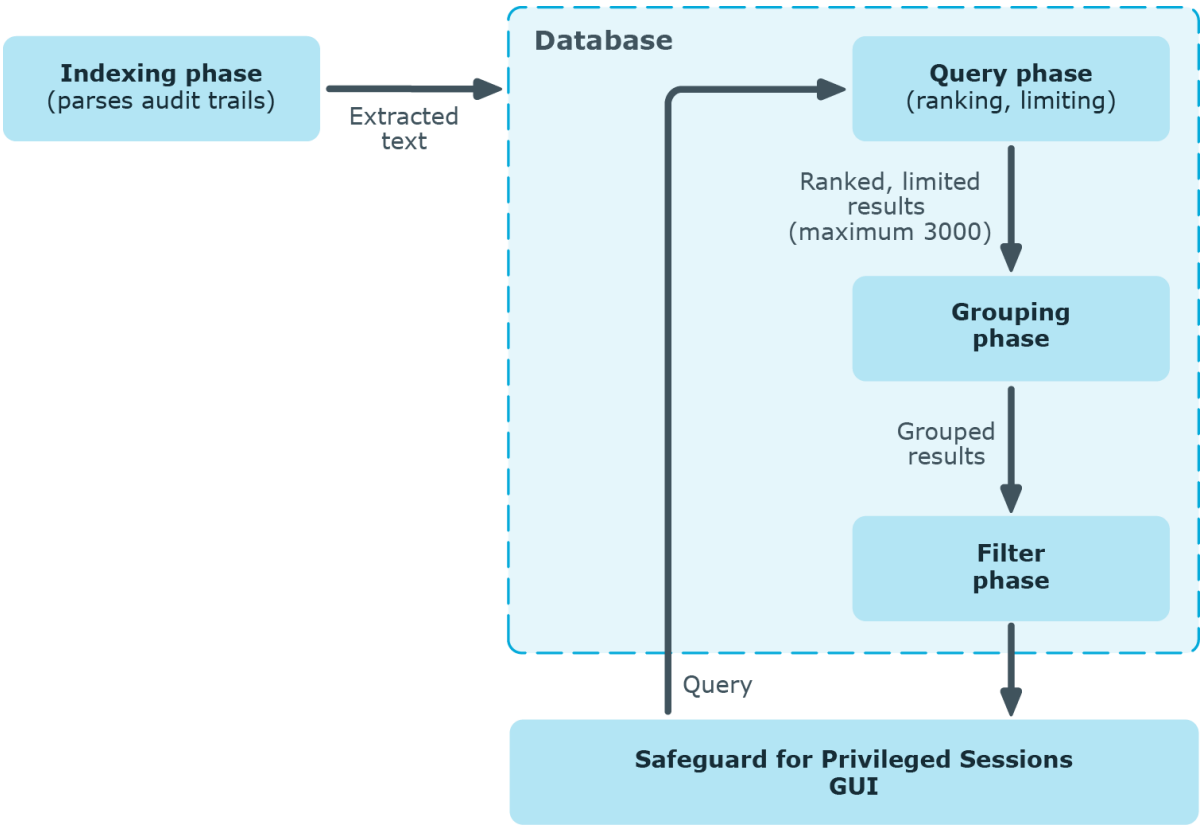




The search and filter process

The screen content is first indexed, then processed with the search backend, and finally, the filter expressions are applied. This process is described in detail in the following sections.

Figure 274: The search and filter process



Prerequisites - Indexing phase

First, as a prerequisite of the search process, screen content is indexed. The indexing phase generates a database that the search and filter processes will run on.

The indexer parses the audit trail files, and builds an "inventory" of the privileged user's activity data based on what appeared on their screen.

1. In the case of a terminal session, screen content corresponds to the activity data that is captured in a terminal window. In the case of graphical protocols, screen content is whatever is visible in the graphical user interface of the applications the user is interacting with. In the latter case, the indexer's Optical Character Recognition (OCR) engine extracts text that appeared on the screen (for example, window titles).

   NOTE: If a piece of text is displayed for less than 1 second, it is not extracted.

2. The indexer returns the information extracted from the parsed audit trail files to One Identity Safeguard for Privileged Sessions (SPS). In the case of a terminal session, the captured text is put in the backend database as one document per one second of screen content. Because of this, the content that you have searched for might only partially appear in the screenshot. In the case of graphical protocols, the captured text is put in the backend database as one document per screenshot.
3. The queries will be run on this database during the search process.

For details on indexing, see Indexing audit trails on page 672.

Search and filter process phases

The search and filter process consists of three major phases:

- Query phase
- Grouping phase
- Filter phase

Query phase

In the query phase, the backend ranks and then limits the number of results.

1. The result of one query is the top 3000 documents, ordered by the default ranking system of the backend.

   This means that if there are more than 3000 results, those of the lowest rank will not be passed to the next phase at all.

   The ranking system cannot be modified, so there is no way to "upvote" those results of lower ranks.

   If you want to ensure that all important results are passed to the grouping phase, use a smaller time range that you run the query on. If there are fewer than 3000 results, it is certain that the events you are interested in will be included in the grouping phase.
2. The grouping phase receives the results.

Grouping phase

The grouping phase groups the results that were passed on from the query phase.

1. First, the results with the same trail IDs are grouped together. A trail ID group contains all search hits that are in that trail.
2. The trail ID groups are then further grouped by seach expression and time range. This group is essentially the time range during which the expression is displayed on the screen (for example, if the text `root` is displayed from 00:00:12 to 00:01:45, this will be one group).
3. This grouped result is displayed in the search screen as one row.

Filter phase

The filter phase applies filter expressions to these grouped results.

NOTE: If there were screen content search results that were excluded during the query phase, the filter expressions will not be applied to them.

Example: Filtering for search results that were excluded in the query phase

For example, if you want to filter for Telnet connections where the text `root` was displayed, the following can happen:

You search for the **Screen content**: `root`. There are 3100 search results that consist of 3050 SSH connections and 50 Telnet connections. In this example, Telnet connections received the lowest ranks for some reason. 100 results that have received the lowest rank are excluded, and in this example it means all Telnet connections.

If you filter for protocol `Telnet` now, you will not see any results.

To solve the problem, try searching in a smaller time range to make sure that there are less than 3,000 search results.

Viewing session details

View the session details of each session for in-depth information on each of the indexed session stored in the connection database. You can use it to gain contextual insight about the indexed session and its events.

You can view session details for data recorded by:

- One Identity Safeguard for Privileged Sessions (SPS). For more information, see *Viewing session details for data recorded by SPS* in the *Administration Guide*.
- One Identity Safeguard for Privileged Passwords (SPP). For more information, see *Viewing session details for data recorded by SPP* in the *Administration Guide*.

Frequent Item Set (FIS) flow view visuals

Frequent Item Set (FIS) flow view visuals are also available on the **Analytics** tab. The FIS flow view is similar to the flow view analytics overview, except that the FIS flow view only displays data narrowed down to a single user's previous sessions in the analysis period (which is the previous 90 days by default). For more information, see Visualizing Frequent Item Sets on the FIS flow view.

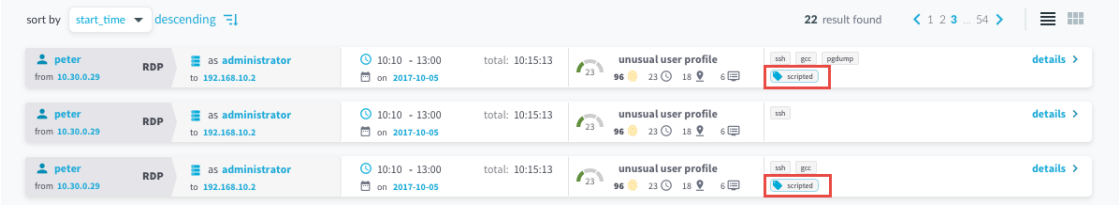
Session tags

Session tags allow you to get basic information about the session and its contents at a glance.

Scripted session tag: One Identity Safeguard for Privileged Sessions (SPS) currently supports the scripted session tag. SPS uses One Identity Safeguard for Privileged Analytics to detect if sessions are generated using human interaction or automation. If sessions are generated using automation, SPS displays the scripted tag in the search interface as shown below:

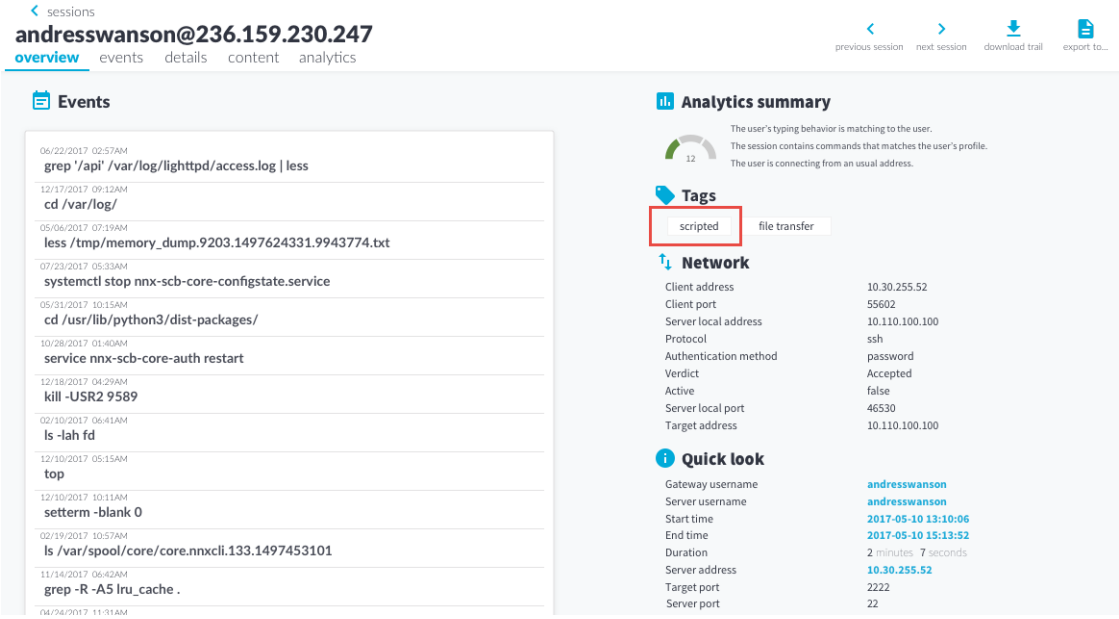
- Scripted sessions are shown on the main search screen.

Figure 275: Scripted sessions — cards view



- Scripted sessions are shown on the **Overview** tab.

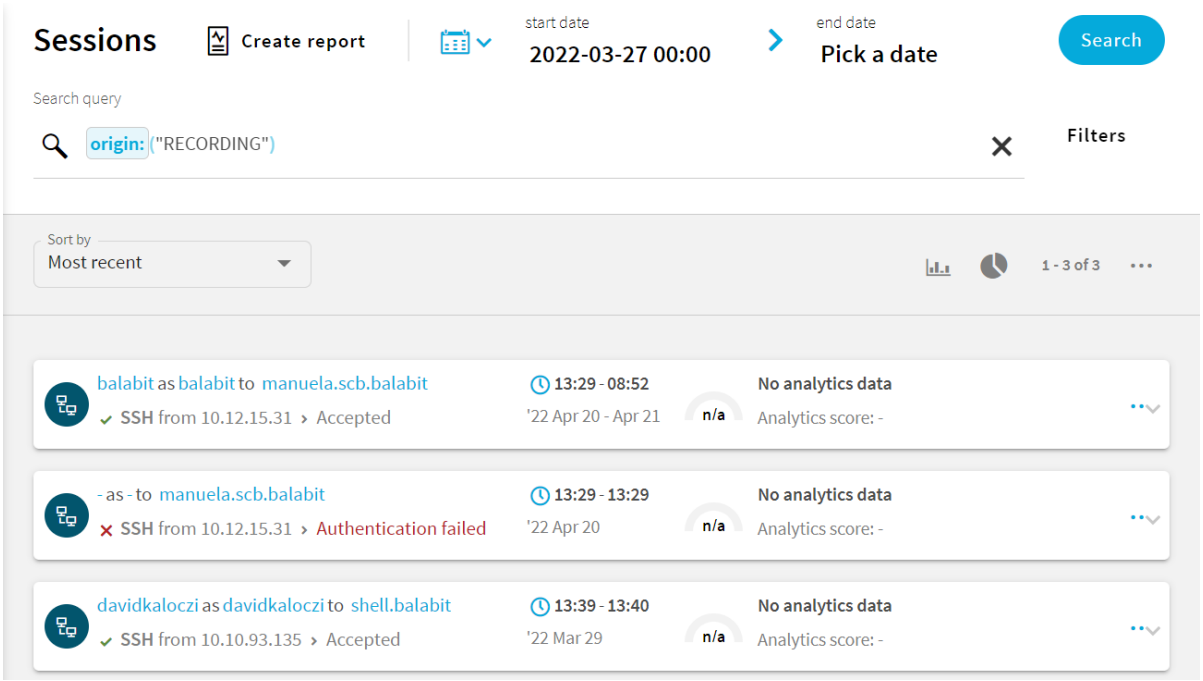
Figure 276: Scripted sessions — Overview tab



Viewing session details for data recorded by SPS

From the Search interface, you can view session details for data recorded by SPS. To view session details for data recorded by One Identity Safeguard for Privileged Passwords (SPP), see Viewing session details for data recorded by SPP.

The  icon in the Search interface indicates that data was recorded by SPS. To search only for data recorded by SPS, enter `origin:("RECORDING")` in the **Search query** field.



To view session details, click the ••• button in the last column of the relevant session.

Figure 277: Sessions — Accessing session details

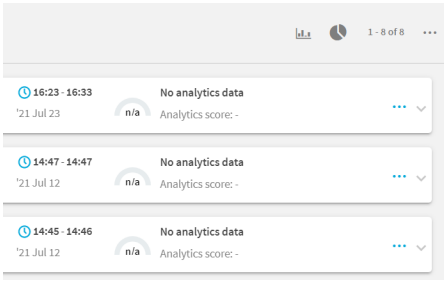
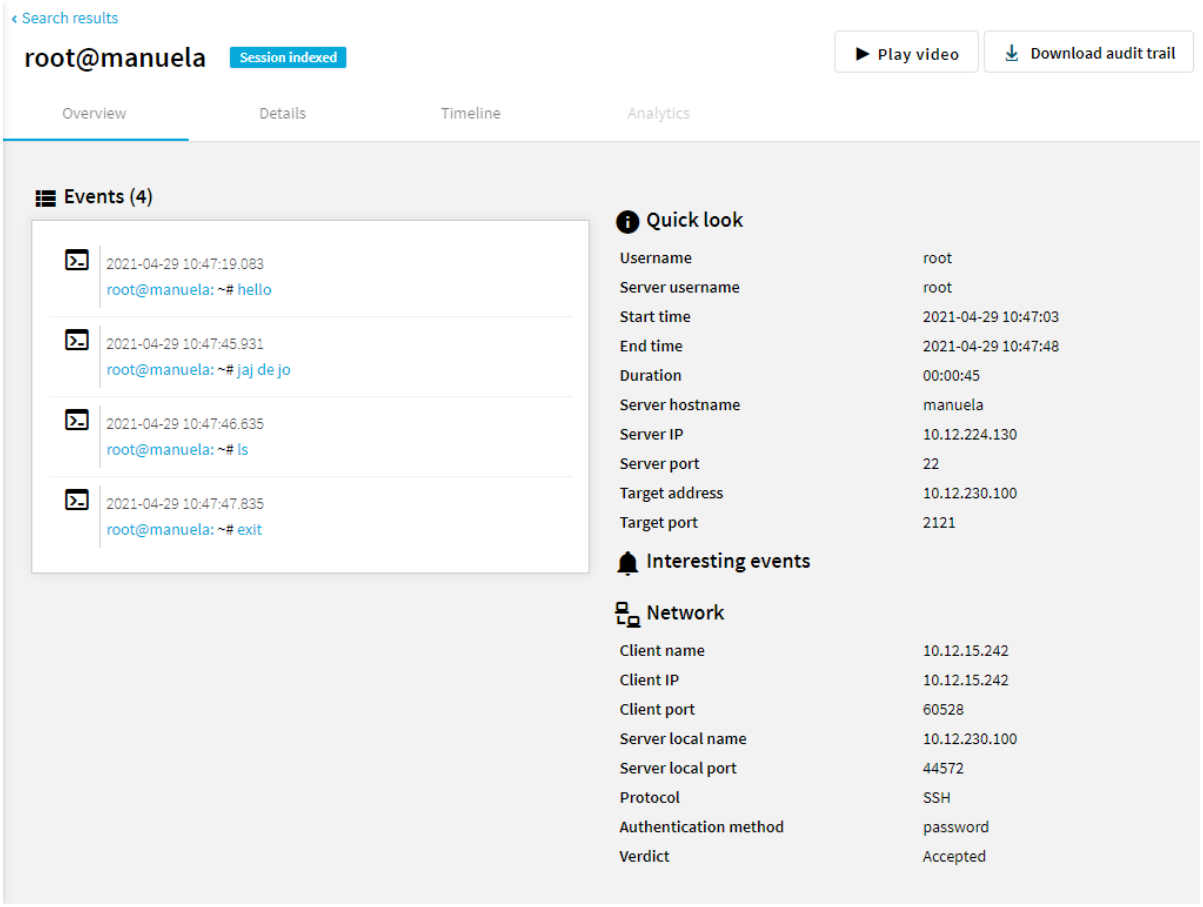


Figure 278: SPS session details



The session details window provides details about the sessions on tabs.

Overview tab

The **Overview** tab is divided into the following main areas:

- The **Events** area displays session events in chronological order.

  View the date and time of the event, the event type and event details. To filter events, use the **Timeline** tab.
- The **Score** area shows the risk score that the Analytics Module assigned to the session. Ranges from 0 to 100, 100 is the highest risk score.
- The **Quick look** area contains user information, for example:
  - Gateway and server username.
  - Start and end time of the session.

  The gateway username corresponds to the `Username` field of the connection metadata database. Note the following:

    - If the user performed inband gateway authentication in the connection, the field contains the username from the gateway authentication (gateway username).
    - Otherwise, the field contains the username used on the remote server.
- The **Interesting events** area displays events selected as interesting, for example, a list of commands and window titles from the session that could be interesting from security point of view.

  The list of interesting events is currently hard-coded and cannot be modified. For terminal sessions, it includes commands such as:

    - `chmod`
    - `ssh`
    - `shutdown`
    - `sudo`
    - `su`
    - `mount`
    - `adduser`
    - `addgroup`

      For graphical sessions, it contains window titles such as:
    - Management Console
    - Control Panel
    - Server Manager
    - PowerShell
    - Security Settings
    - Windows Security Center
- The **Network** area displays session information, for example:
    - Verdict
    - Protocol
    - Connection policy
    - Client and server address

Details tab

In addition to the **Quick look**, **Interesting events**, and **Network** screen areas, the **Details** tab also provides monitoring information, audit trail downloads information, and channels information.

If there is a gateway authentication or authorization failure due to an **AA plugin**, the reason of the failure is displayed in the **Deny reason** field.

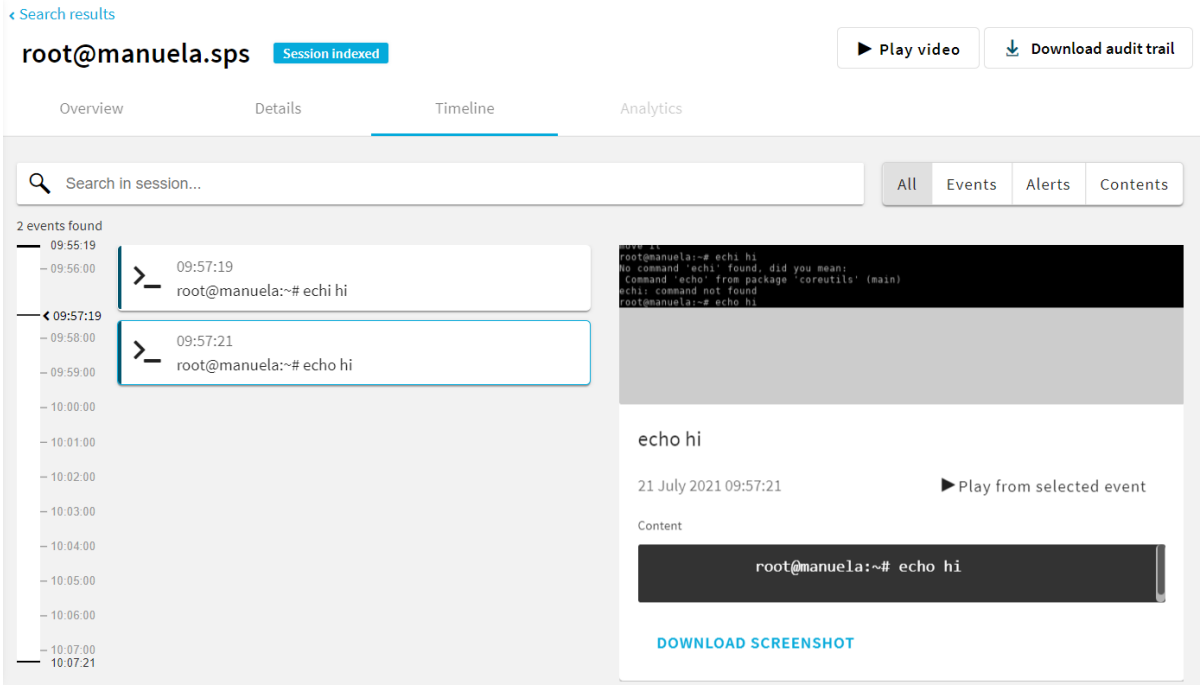
Timeline tab

On the **Timeline** of the Search interface, for data recorded by SPS, you can do the following:

- View session events and alerts.
- Search in the contents of the audit trail.
- Watch the video of an audit trail from a selected event.

The **Timeline** tab replaces the now deprecated **Events**, **Alerts**, and **Contents** tabs.

Figure 279: Timeline tab



Events:

- Session events in chronological order.

  You can narrow the event list by entering the event name in the **Search in session** field.
- Date and time of the event.
- Event type (command, screen content, and window title).
- Event details.
- Screenshots.

Alerts:

- Content policy alerts triggered in the session, in chronological order.

  You can narrow the alert list by entering the alert name in the **Search in session** field.
- Date and time of the alert.
- Alert details.
- Screenshots.

Contents:

You can search in the contents of the audit trail using plain-text search. Searching using complex keyphrases is not supported from the **Timeline** tab.

Screenshots:

If screenshots are available for the session, you can click each event or alert to view the corresponding screenshot.

Screenshots are not available for:

- Ongoing sessions.
- Unindexed sessions.
- Trails of HTTP sessions.
- Encrypted trails (without the necessary certificate).

  If screenshots are encrypted, you have to upload the necessary encryption key to your keystore. For more information, see Viewing encrypted screenshots.

Analytics tab

If you use the One Identity Safeguard for Privileged Analytics, you can view detailed analyses provided by the configured algorithms. For more information, see Analyzing data using One Identity Safeguard for Privileged Analytics on page 810.

Managing active sessions

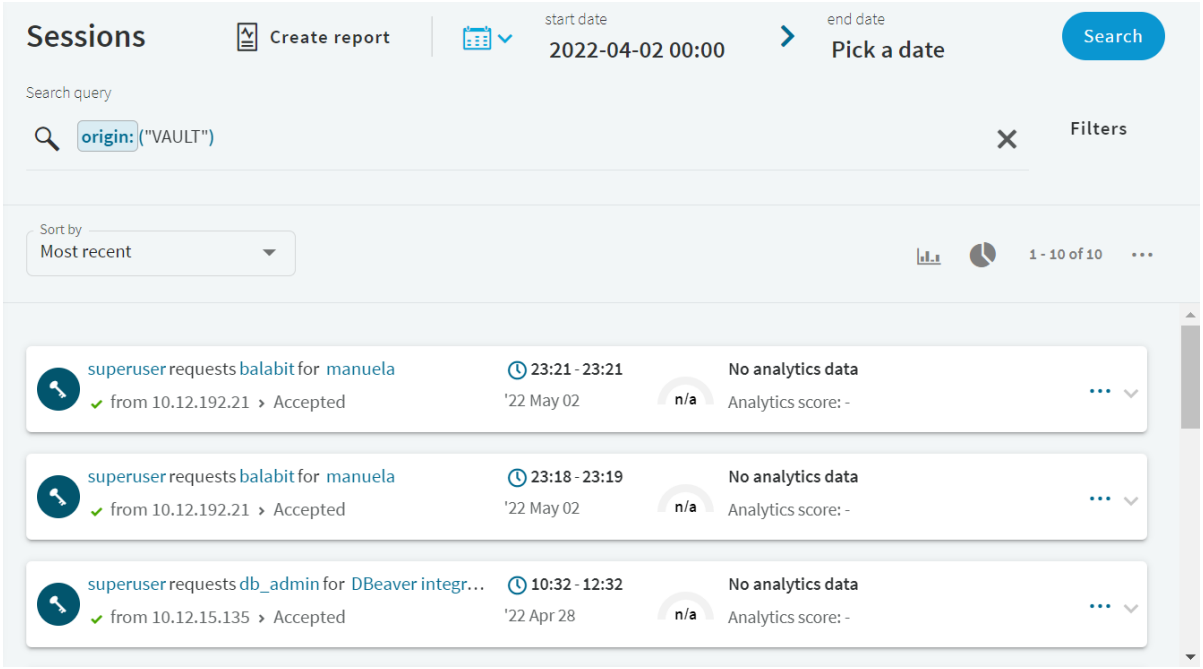
For information on how to follow and terminate active sessions, see section Following active sessions.

Viewing session details for data recorded by SPP

From the Search interface, you can view session details for data recorded by SPP. To view session details for data recorded by SPS, see Viewing session details for data recorded by SPS.

If you have linked your SPS to your SPP deployment, the icon in the Search interface indicates that data was recorded by SPP.

To search only for data recorded by SPP, enter `origin:("VAULT")` in the **Search query** field.



To view session details, click the button in the last column of the relevant session.

Figure 280: Sessions — Accessing session details

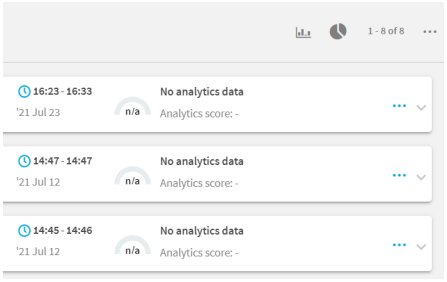
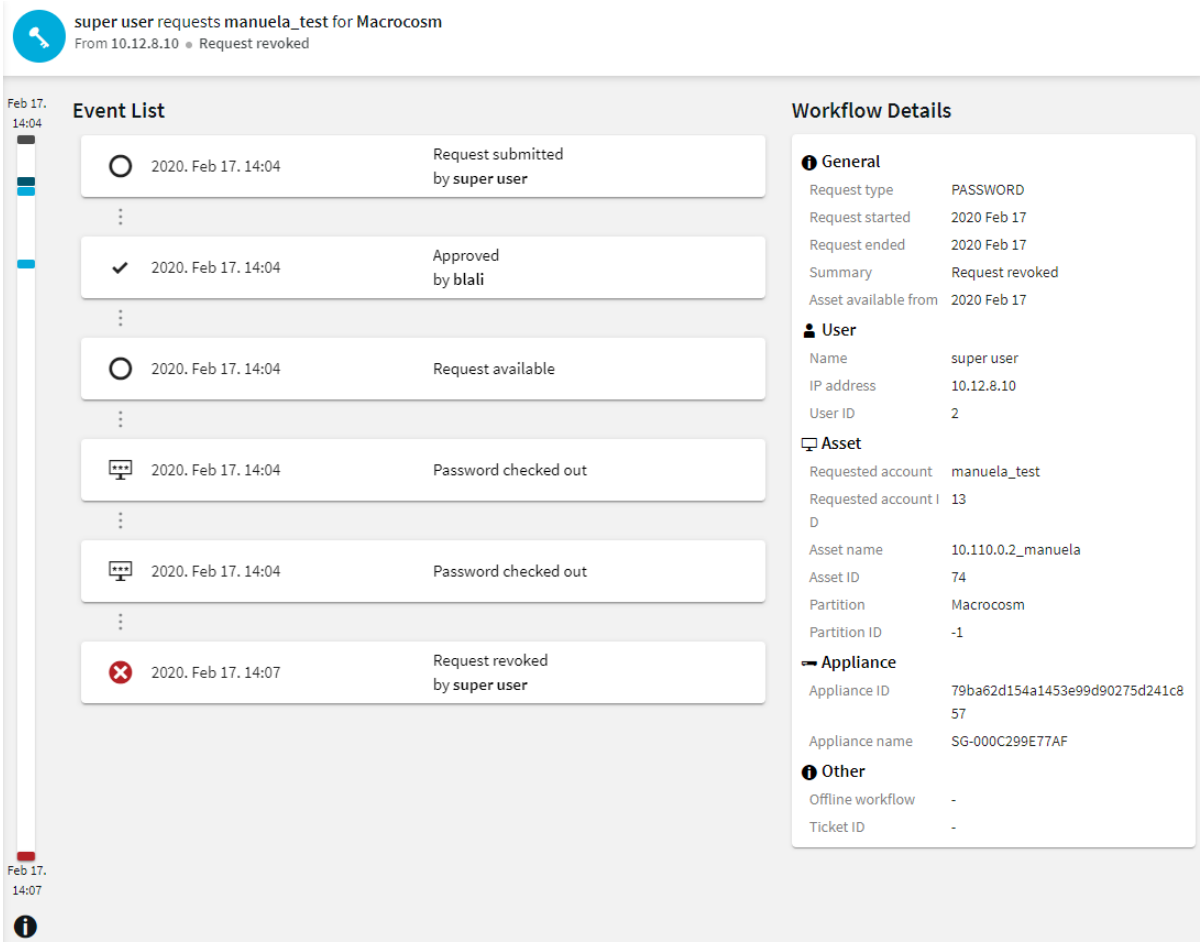


Figure 281: SPP session details



You can display the session details in a timeline. To access the color legend for the timeline, click (Timeline Color Legend). Hovering your mouse over an event shows the position of that event in the timeline sidebar. It also shows the event as active.

Visualizing Frequent Item Sets on the FIS flow view

This section provides an overview of the Frequent Item Sets (FIS) flow view feature on the Search interface. It describes the underlying component (that is, the FIS algorithm), the elements of the FIS flow view visual, and possible scenarios (depending on your session details and pattern option choice).

From One Identity Safeguard for Privileged Sessions (SPS) version 6.2., a visual overview of Frequent Item Set (FIS) analysis is available on the **Sessions** interface. The FIS flow view is essentially similar to the flow view analytics overview, except that the FIS flow view only displays data narrowed down to a single user's previous sessions in the analysis period (which is the previous 90 days by default).

The component behind the analysis is the FIS algorithm, which examines multiple attributes of sessions and attempts to find values that frequently appear together, forming a set. Using this information, the algorithm can discover patterns in user behavior.

NOTE: For the FIS algorithm to be able to score a user's sessions, the user needs at least 1 FIS baseline built. Algorithm baselines (including the FIS baselines) are built automatically every day (usually during hours with less heavy traffic).

Elements of the FIS flow view

To access the FIS flow view feature, click on **Details** of the session of your choice in **Audit** > **Sessions** > **Sessions Details**, then click **Analytics**.

Figure 282: Audit > Sessions — The FIS flow view on the Analytics tab

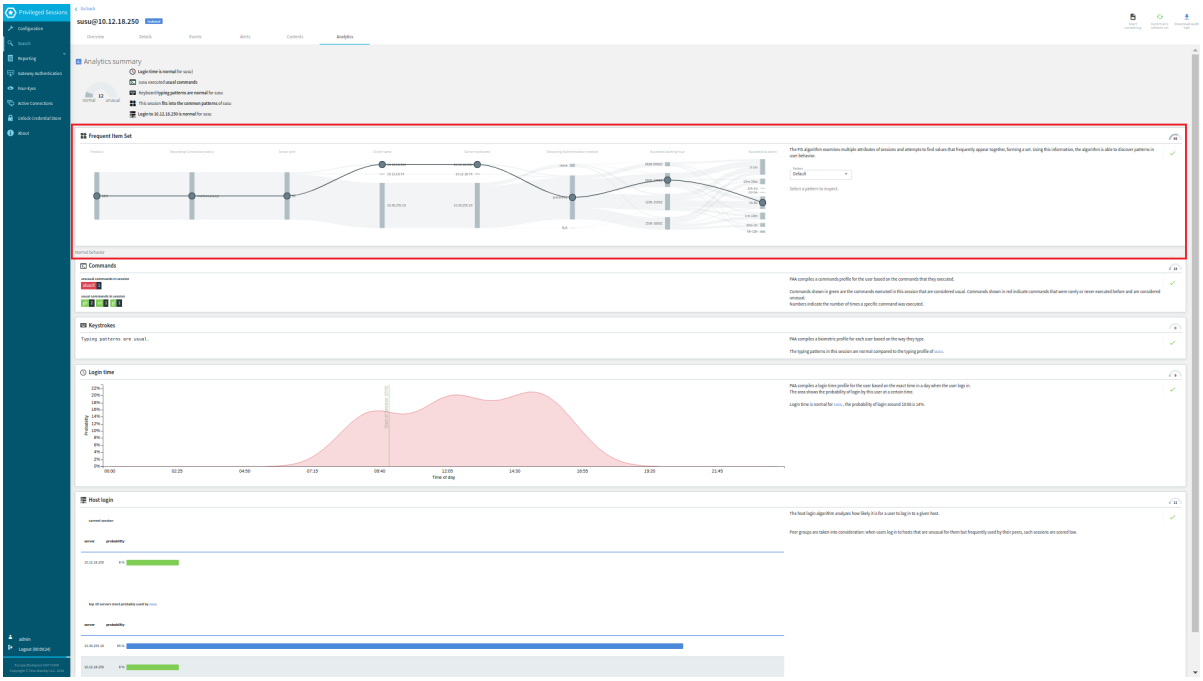
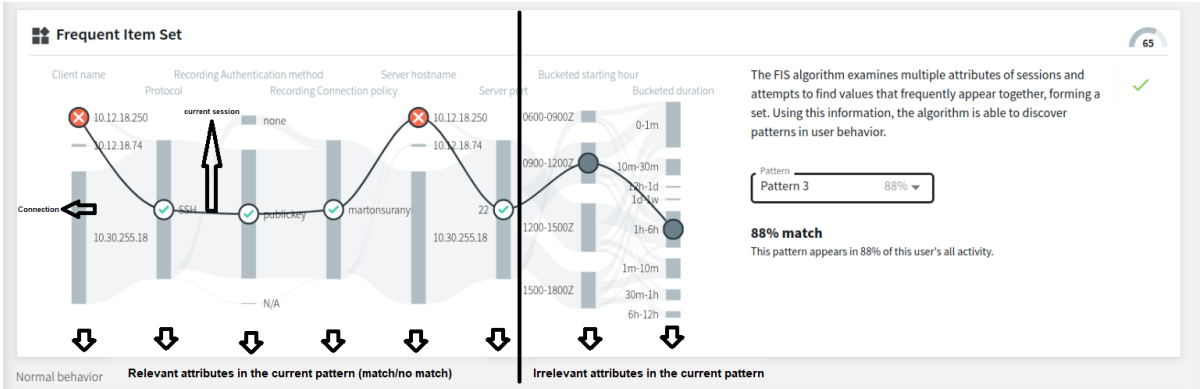


Figure 283: Audit > Sessions — The FIS flow view (close-up)



The layout of the FIS flow view is similar to the attributes-grouped flow view, with the sessions forming a flow line (from here on referred to as a Connection). The relevant examined attributes (including the attribute nodes) are visible on the left hand side, while the irrelevant attributes (including the attribute nodes) are on the right hand side of the flow line visual.

The continuous line, with circles representing attribute nodes (that can be gray (●) or contain the ✓ or the ✗ mark), displays the current session. Each attribute node represents that particular attribute's value in the current session. The FIS algorithm attempts to match these attribute values to the patterns in the user's baseline.

The different attribute node markings stand for:

● : session data for this attribute is irrelevant for pattern match comparison (or session data is part of the **Default** option)

✓ : session data for this attribute is a match for the selected pattern

✗ : session data for this attribute is not a match for the selected pattern

Attributes

Similarly to the flow view, session data in the FIS flow view is grouped according to attributes (such as Protocol, Client name, Server hostname, Server port, and so on) that come from session data. The two attributes in the FIS flow view that do not come from session data as-is, but are further grouped instead:

- **Bucketed starting hour**: the most frequently used session starting hours grouped into intervals of 3 hours each
- **Bucketed duration**: the most frequently occurring session duration values, grouped into intervals of various length

NOTE: As a rule, the relevant attributes and attribute nodes (marked with ✓ or ✗) are located on the left hand side of the flow line visual, while the irrelevant attribute nodes (marked with ●) are located on the right hand side of the flow line visual. The number of relevant attributes for the pattern match comparison, as well as the attribute groups' (and, as a result, the attribute nodes') relative position and relative order (going from left to right) changes from pattern to pattern on the flow line visual.

The FIS score gauge

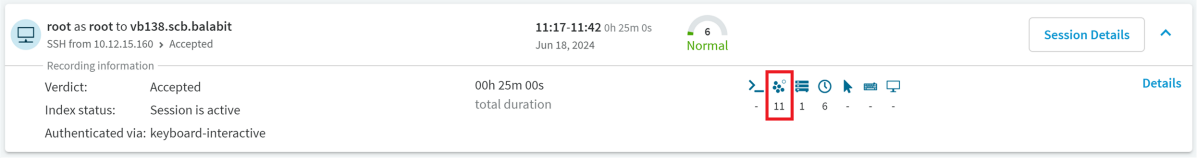
The FIS score gauge (located in the upper right corner of the FIS flow view visual) indicates the FIS score of the selected session.

The FIS score is assigned to the session after the FIS algorithm analyzes it in comparison with the latest available baseline. If the session has a high amount of matches to the user's baseline, the FIS score's value will be low (indicating normal user behavior).

NOTE: The closer the FIS score gets to 100, the more it indicates unusual user behavior. For more information about normal and unusual user behavior, see Analyzing data using One Identity Safeguard for Privileged Analytics, Viewing session details, and the *View session analytics* step in Analyzing data using One Identity Safeguard for Privileged Analytics.

The value of the FIS score (also visible on the **Sessions** > **Sessions list**, as well as in the aggregated score summary, usually above the FIS flow line visual) is one of the several components of the session's aggregated score.

Figure 284: Audit > Sessions — The FIS score visible on Sessions > Sessions list



For more information about analytics algorithms and scores, see this section.

Pattern selection drop-down list

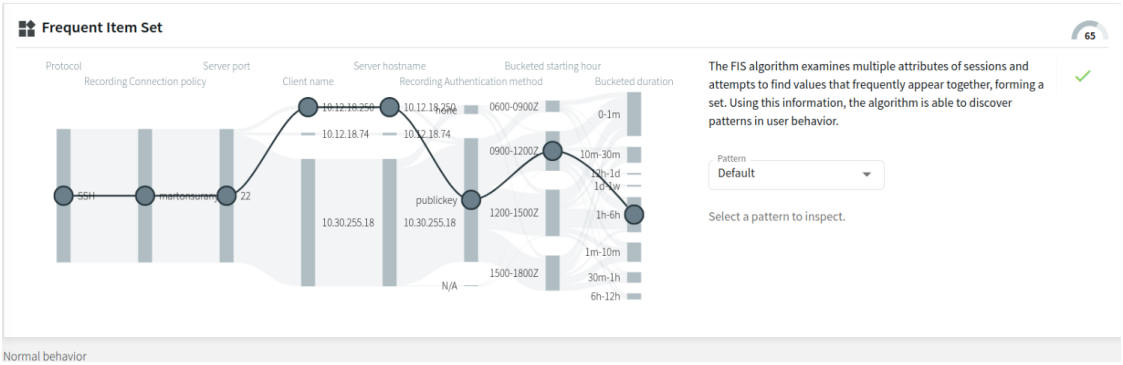
The pattern selection drop-down list (by default set to the **Default** option) offers the available patterns that the FIS algorithm generates after baseline analysis.

Depending on which available pattern option you select, the visual will display one of these cases:

- The **Default** option

  The **Default** option is the set of values that the FIS algorithm takes as a reference point. If you select the **Default** option, the flow line visual displays the user's previous sessions in the observed analysis period, with no patterns selected yet. The flow line visual below (set to the **Default** option) displays all possible attributes, marked with ● in all attributes.

  Figure 285: Audit > Sessions — FIS flow view - the Default option (no patterns displayed yet)

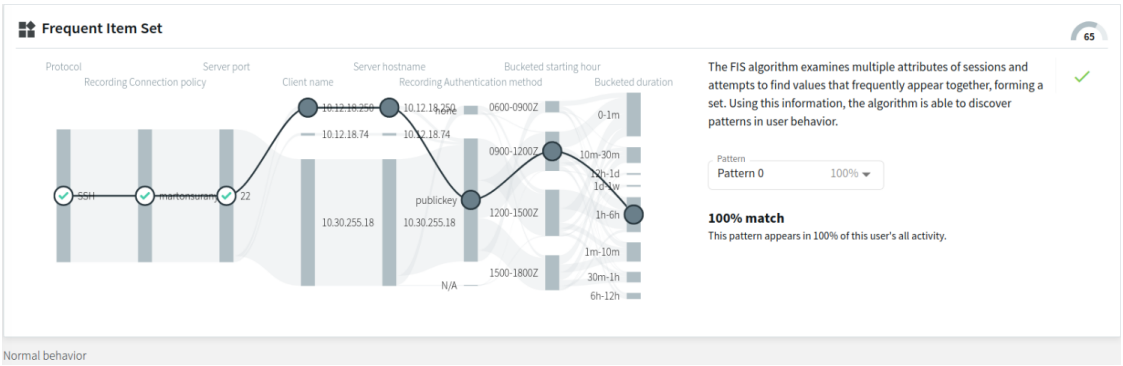
  

- **Pattern with a 100% match to this user's sessions during the observed analysis period**

  The example below displays a pattern that appears in 100% of the user's previous sessions during the analysis period. As mentioned before, the relevant attribute nodes (marked with the ✓ or the ✗ sign, depending on whether that attribute value matches or does not match the **Default** option in that particular session) are arranged on the left-hand side, while the irrelevant ones (marked with ●) are arranged on the right-hand side of the flow line visual. There are 3 attribute matches

(namely, Protocol, Recording Connection policy, and Server port) to this particular pattern in this particular session.

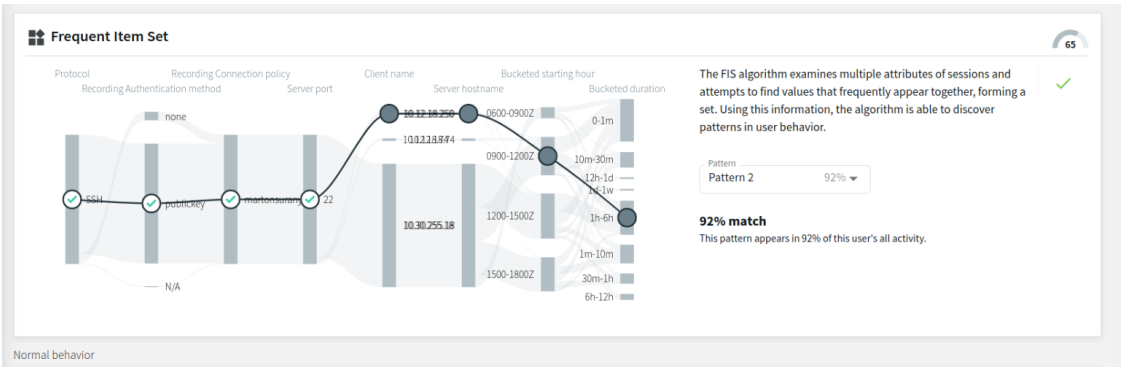
Figure 286: Audit > Sessions — FIS flow view - a pattern with a 100% match



- **Pattern with a high match percentage to this user's sessions during the observed analysis period**

  The example below displays a pattern that appears in 92% of the user's previous sessions during the observed analysis period. Similarly to the previous example, the relevant attribute nodes are arranged on the left-hand side, while the irrelevant ones are arranged on the right-hand side of the flow line visual. This example contains different relevant attribute nodes than the previous example, with 4 attribute matches to this particular pattern in this particular session.

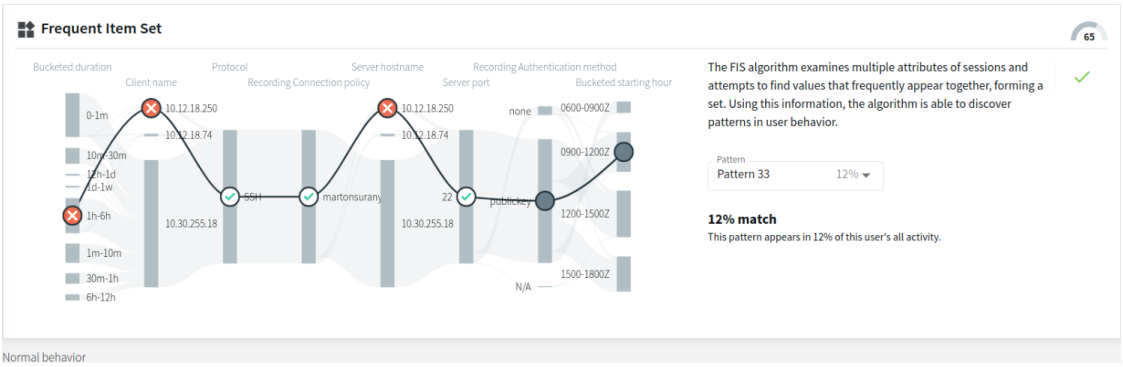
Figure 287: Audit > Sessions — FIS flow view - a pattern with a high match percentage



- **Pattern with a low match percentage to this user's sessions during the observed analysis period**

  The example below displays a pattern with a low match percentage, appearing in only 12% of this user's previous sessions during the observed analysis period. There are 3 attribute nodes that match (marked with ✓) and 3 attribute nodes that do not match (marked with ✗) this particular pattern in this particular session.

Figure 288: Audit > Sessions — FIS flow view - a pattern with a low match percentage



With the baseline generated every day on average (usually during hours with less heavy traffic), the baseline itself is continuously changing. As a result, the available patterns are also continuously changing over time.

Pattern match percentage

The pattern match percentage is a percentage value displayed under the pattern selection drop-down list, next to the pattern name (for example, **Pattern 0** | **Pattern 1**, and so on).

Figure 289: Audit > Sessions — FIS flow view - the pattern match percentage

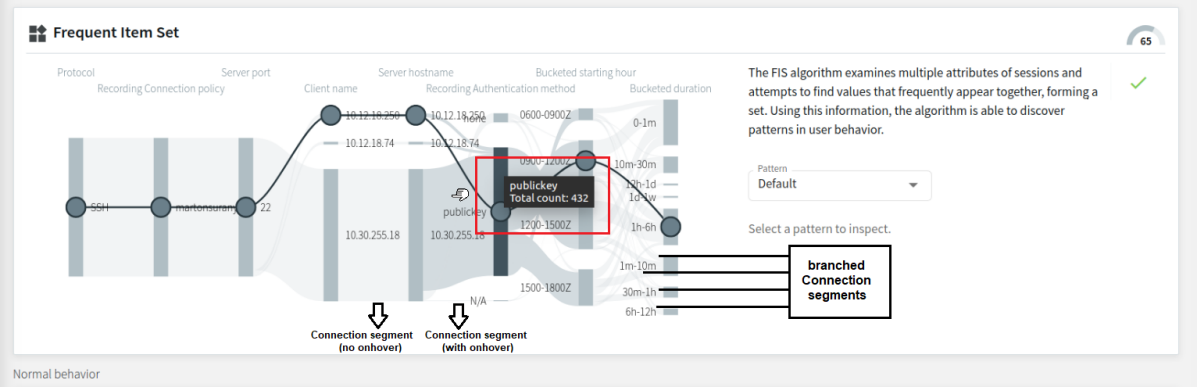
The percentage value always matches the value displayed next to the **Pattern** options (other than the **Default** option) in the drop-down list. The percentage indicates to what percent the selected pattern matches this user's sessions during the observed analysis period. Patterns that appear in less than 10% of this user's sessions during the observed analysis period are not displayed in the drop-down list as **Pattern** options (but may appear in the overall sum of **Total count** data, as mentioned in Displaying further details of individual Connection segments.

Displaying further details of individual Connection segments

Hovering on an individual Connection segment of the visual displays the exact number of occurrences of that particular attribute value during the observed analysis period. When the Connection splits to several branches (for example, in the example below, the Connection splits into separate branches according to **Bucketed starting hour**), hovering over the branched Connection segment of your choice will display that particular attribute's further details (namely, the attribute data (for example, IP addresses), and the **Total count**).

NOTE: The **Total count** values include session data originating from patterns that appear in less than 10% of the user's sessions during the observed analysis period.

Figure 290: Search — Hovering on particular Connection segments to display further details



Viewing audited user details

View the audited users details for in-depth information on their activity, sessions, aggregated user data, and score. You can use it to gain contextual insight about your audited users.

Prerequisites

You must have the proper Access Control Lists (ACLs) access set for viewing audited user details. For more information on ACLs, see ACL access and summary list.

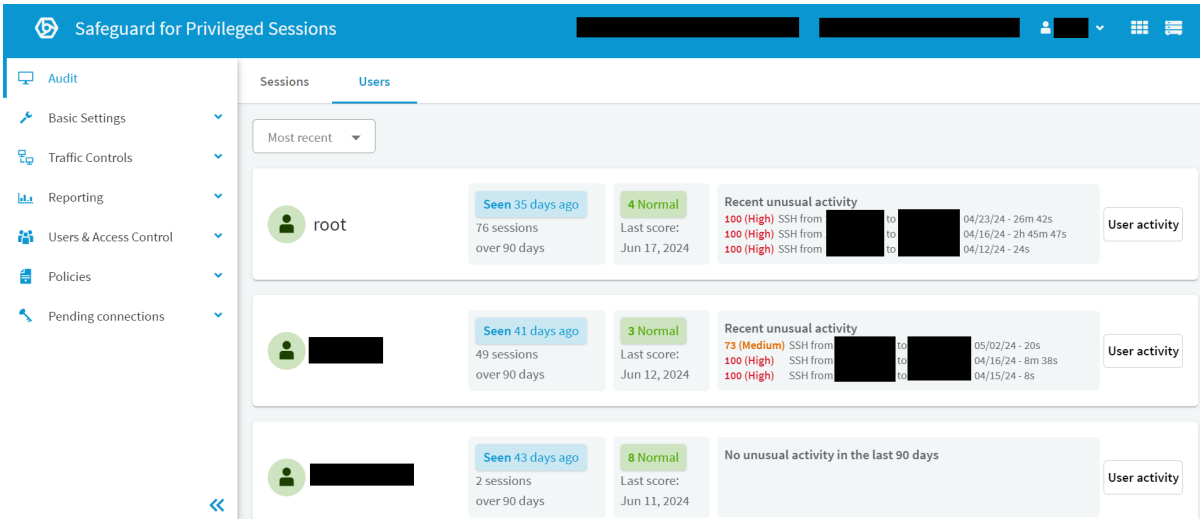
Limitations

If the auditor is restricted by Audit Data Access Rules (ADARs), the aggregated session details of the users are displayed accordingly. This means that the restricted auditor can only view the aggregated session details that the ADAR settings allow.

Audited users list

To view the audited users list, navigate to **Audit** > **Users**. On the main screen of the **Users** tab, you can view the audited users.

Figure 291: Audited users view



Each audited user is assigned three cards that are displayed on the main screen:

- Activity card: The Activity card displays information on the time of the user's last activity and the number of their sessions in the last 90 days.
- Score card: The Score card displays the score of the user calculated by numerous algorithms based on the aggregated user data and the date of the last score. Based on the score calculated for the user, their score can be evaluated as **Normal**, **Medium** or **High** by the algorithms. If the score is not available or it cannot be calculated, the score is evaluated as **Unavailable** score.
- Unusual activity card: The unusual activity card displays information on the most recent unusual activities conducted by the user. The card is limited to show the maximum of the three most recent unusual activity associated with the user.

Sorting audited users

On the main screen of the **Users** tab, the audited users are sorted by the most recent sorting filter by default. You can sort the audited users by the following sorting parameters as well:

- **Name**
- **Score**
- **Last seen**
- **Score timestamp**.

You can also use the **Search** field to search for sorting parameters.

Viewing session details associated with an audited user

View the session details associated with an audited user to gain contextual insight about the user's sessions and activity.

To view session details associated with an audited user

1. Navigate to **Audit** > **Users**.
2. Locate the user whose sessions you want to view.
3. Click **User activity**.

On the detailed individual user view screen, you can audit the user's score, last activity, session statistics and session details.

Specifying time ranges

Specify a time range for your audit. For more information, see Specifying time ranges.

For the activity heatmap, you can use the squares of the heatmap to select the start and end date of your audit. For more information, see Activity heatmap.

Session statistics and baselines

For the audited user, a **Session statistics** card and a **Baselines** card is displayed on the detailed individual user view screen.

The **Session statistics** card provides the following information on the audited user's sessions:

- **Failed login attempts**
- **Total target assets**
- **Terminated sessions**
- **Primary target**.

The **Baselines** card provides information on the build status of the baselines. For each user, the following baselines are built:

- **Login time**
- **Keystroke**
- **Mouse**
- **Frequent item set (FIS)**
- **Host login**
- **Command**
- **Window title**.

If the baseline build was successful, the (check) icon is displayed next to the baseline. If the baseline build was not successful the (warning) icon is displayed next to the baseline.

NOTE: Some baselines might not be built for the audited user if:

- Analytics is not enabled on SPS.
- One or more of your analytics, content, or indexing policies is not configured properly.
- There is no sufficient session data available for building the user baseline.

For more information on the algorithms that build the baselines, see *Algorithms* in the *Safeguard for Privileged Analytics Configuration Guide*.

Activity heatmap

In the activity heatmap view, you can see individual user activity in a heatmap format. Each square of the heatmap represents a specific day. The darker blue the square is, the more sessions the audited user conducted on a specific day.

NOTE: The number of sessions represented as shades in the activity heatmap:

- darkest shade: more than 5 sessions
- second darkest shade: 3-5 sessions
- second lightest shade: 1-2 sessions
- lightest shade: 0 sessions

You can change the audited time range by clicking a square that represents the start date and a square that represents the end date. When selecting a time range that is more than 77 days, you can scroll the heatmap to view more data. In this case, you need to scroll vertically, and the activity heatmap will scroll horizontally.

Sessions of the audited user

On the detailed individual user view screen, you can view, filter and search the audited user's sessions.

To view details about the audited user's single session, select the session and click **Session Details**. For more information on viewing session details, see Viewing session details for data recorded by SPS.

Click **Filters/Advanced** to filter or search in the audited user's sessions. You can use the following filtering parameters for filtering the sessions:

- **Origin**
- **Server hostname**
- **Protocol**
- **Verdict**
- **Contains text**.

Click **Timeline** to view the timeline of the audited user's sessions. You can view and filter the timeline of the active and started sessions.

Click **Quick statistics** to view quick statistics on the audited user's sessions. You can choose from several value distribution options.

Click **Create a reporting subchapter** to create a reporting subchapter based on your search parameters.

You can use several viewing options to view the details of your audited user's session. For more information on the view options, see Card view, Table view, Flow view and Adding custom fields to the card view.

Replaying audit trails in your browser

The following section describes how to replay an audit trail in your browser.

NOTE: You can replay audit trails in your browser, or using the Safeguard Desktop Player application. Note that there are differences between these solutions.

For details on the Safeguard Desktop Player application, see *Safeguard Desktop Player User Guide*.

The following table details the differences between the solutions provided by the browser and the Safeguard Desktop Player application when replaying audit trails.

	Browser	Safeguard Desktop Player
Works without installation	✓	-
Works on any operating system	✓	Windows, Linux, Mac
Replays audit trails recorded with SPS 5 F4 and newer	✓	✓
Replays TN5250 sessions	✓	✓
Extracts files from SCP, SFTP, HTTP and RDP sessions	-	✓
Replays HTTP sessions	-	Only exports raw files from the command line
Replays X11 sessions	✓	✓
Starts replay while rendering is in progress	✓	✓
Follows 4-eyes connections	-	✓
Replays live streams in follow mode	✓	✓

	Browser	Safeguard Desktop Player
Exports to PCAP	-	✓
Displays user input	✓	✓
Displays subtitles for video	✓	✓
Exports audit trail as video	-	✓
Exports screen content text	-	✓
Searches in the contents of the audit trails	-	✓

CAUTION:

From version 6.13.0, SPS does not support Internet Explorer 11 (IE11) anymore. SPS version 6.12.0 and previous versions continue to support IE11.

From SPS version 6.10, the Google WebM Video for Microsoft Internet Explorer plugin is not required for replaying audit trails in your browser. The supported browsers are:

- Google Chrome
- Firefox
- Safari
- Internet Explorer 11 (IE11) - supported until SPS version 6.12.0

For SPS version 6.9 and earlier versions, even though the One Identity Safeguard for Privileged Sessions (SPS) web interface supports Internet Explorer and Microsoft Edge in general, to replay audit trails, you need to use Internet Explorer 11, and install the Google WebM Video for Microsoft Internet Explorer plugin.

If you cannot install Internet Explorer 11 or another supported browser on your computer, use the the Safeguard Desktop Player application. For details, see *Replaying audit trails in your browser* in the *Administration Guide* and *Safeguard Desktop Player User Guide*.

NOTE: From SPS version 6.10 and onwards, you can play video files of audit trails both in central and in cluster environments. Previously, you could play the video files only on the Search Master node in a cluster environment. From SPS version 6.10 and onwards, you can play video files on Search Minion nodes as well.

To replay an audit trail in your browser

1. On the **Sessions** page, select the audit trail you want to replay.
2. To display the details of a connection, click ••• . This page is called the details view.
3. To play the video file, click ▶ Play video .
4. (Optional) For encrypted audit trails, upload any missing private keys and certificates. For more information on these procedures, see section Replaying

encrypted audit trails in your browser.

After uploading the private keys and the corresponding certificates, the encrypted upstream traffic elements are decrypted. The trail is decrypted and decompressed on the client-side. As a result, the encrypted elements will be displayed distributed in the video.

5. (Optional) After uploading any missing private keys or certificates, to play the video, click ▶ Play video.
    1. (Optional) If there are more channels that can be played, select the channel that you want to play in the **Multiple channels are available** dialog and click **Play video**.

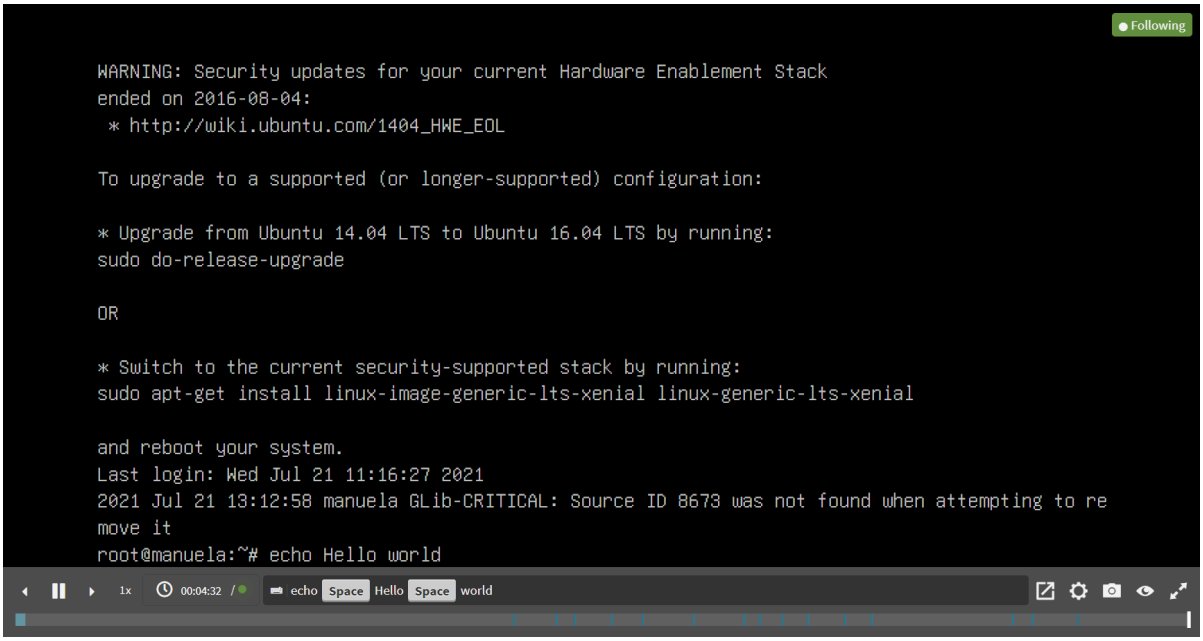
       The Player window opens.

For more information on how to use the browser to play video files, see Using the browser to play video files.

Using the browser to play video files

This section provides information on how to use the browser to play video files of recorded audit trails, or how to follow active sessions.

Figure 292: Replaying audit trails or following active sessions in your browser



Following active sessions

If you are following an active session, a green **Following** label is displayed in the top-right corner of the screen, and there is a green, pulsing dot on the menu bar. If you stop following the session, for example, by jumping back in the video, a green **Follow** button is

displayed on the menu bar, which allows you to continue to follow the active session. If a session is not active anymore, the **Following** label and the green dot are not displayed.

Player window controls

TIP: You can quickly zoom in or out by clicking anywhere in the Player window.

NOTE: From the video, you can copy the content of text-based protocols to the clipboard:

1. Select the required content.
2. Right click on the selected content.
3. From the list of available options, select **Copy**.

The Player window has the following controls:

- , : *Play*, *Pause*.
- , : *Jump to previous event*, *Jump to next event*.

  If you use the button, the video playing is stopped so that you can jump back to preceding events.
- 1x: *Adjust replay speed*. The possible increments are: `0,25`, `0,5`, `1`, `1,5`, `2`, `3`, and `5`.
- Time-related information and options of audit trails:
  - 00:00:45 / 00:13:13: *Time since the audit trail started / Length of the audit trail*.
  - To jump to a timestamp in the video, click on the time.

    00:00:14 MM:SS / 00:13:13: *Time since the audit trail started (top) / Jump to a timestamp (left) / Length of the audit trail (right)*.
  - 00:02:50 /: *Time since the audit trail started* when following an active session.
- Left Shift | Left Shift + Space | Left Shift + .: *List of keyboard events*. Special characters like SHIFT, ENTER, F1, and so on, and mouse usage are displayed as buttons. If the upstream traffic is encrypted, upload your permanent or temporary keys to **User menu** > **Audit keystore** to display the keyboard events. This will not be displayed if your upstream traffic is encrypted but not unlocked.
- : *Share*. You can copy the link of a session, or optionally, obtain the link of a session starting from a specific timestamp.
- : *Encoding settings*. This option enables you to set the encoding for terminal-based and graphical protocols.

  For terminal-based protocols, you can set the following:

- Terminal encoding layout
- Telnet codec
- Telnet alternate codec

For graphical protocols, you can set the keyboard layout.

You can save your video encoding settings, which are stored locally, in your browser.

- : *Create a screenshot*.
- : *Show / hide events*. Select the types of events to display. The available options are **Keystroke**, **Mouse activity**, and **On-screen changes**. Depending on the protocol used and how the audit trail was processed, SPS can display keyboard events, commands, mouse events, and window titles. Commands and window titles are displayed as subtitles at the bottom of the screen. This will not be displayed if your upstream traffic is encrypted but not unlocked.
- , : *Switching fullscreen mode on and off*
- : Progress bar and distribution of events. Light blue - screen change, blue - keyboard or mouse event, dark blue - command or title event. This will not be displayed if your upstream traffic is encrypted but not unlocked.
- : *Close the player*, and return to the **Connection details** page.

Streamable session recording playback with Safeguard Desktop Player started from the SPS UI

From SPS version 7.2, you can play back your session recordings with the Safeguard Desktop Player application started from the SPS UI. With this method, you do not have to start the Safeguard Desktop Player application outside SPS to start replaying your session recording. This method starts a streamable, on-demand playback. As a result, you do not have to wait for a potentially large file to completely download before you start replaying the recording.

NOTE: You can replay session recordings in your browser, or using the Safeguard Desktop Player application. Note that there are differences between these solutions.

For details on the Safeguard Desktop Player application, see *Safeguard Desktop Player User Guide*.

The following table details the differences between the solutions provided by the browser and the Safeguard Desktop Player application when replaying session recordings.

	Browser	Safeguard Desktop Player
Works without installation	✓	-
Works on any operating system	✓	Windows, Linux, Mac
Replays session recordings recorded with SPS 5 F4 and newer	✓	✓
Replays TN5250 sessions	✓	✓
Extracts files from SCP, SFTP, HTTP and RDP sessions	-	✓
Replays HTTP sessions	-	Only exports raw files from the command line
Replays X11 sessions	✓	✓
Starts replay while rendering is in progress	✓	✓
Follows 4-eyes connections	-	✓
Replays live streams in follow mode	✓	✓
Exports to PCAP	-	✓
Exports audit trail as video	-	✓
Exports screen content text	-	✓
Searches in the contents of the session recordings	-	✓

CAUTION:

From version 6.13.0, SPS does not support Internet Explorer 11 (IE11) anymore. SPS version 6.12.0 and previous versions continue to support IE11.

From SPS version 6.10, the Google WebM Video for Microsoft Internet Explorer plugin is not required for replaying audit trails in your browser. The supported browsers are:

- Google Chrome
- Firefox
- Safari
- Internet Explorer 11 (IE11) - supported until SPS version 6.12.0

For SPS version 6.9 and earlier versions, even though the One Identity Safeguard for Privileged Sessions (SPS) web interface supports Internet Explorer and Microsoft Edge in general, to replay audit trails, you need to use Internet Explorer 11, and install the Google WebM Video for Microsoft Internet Explorer plugin.

If you cannot install Internet Explorer 11 or another supported browser on your computer, use the the Safeguard Desktop Player application. For details, see *Replaying audit trails in your browser* in the *Administration Guide* and *Safeguard Desktop Player User Guide*.

NOTE: From SPS version 6.10 and onwards, you can play video files of audit trails both in central and in cluster environments. Previously, you could play the video files only on the Search Master node in a cluster environment. From SPS version 6.10 and onwards, you can play video files on Search Minion nodes as well.

Prerequisites

- The Safeguard Desktop Player application is installed and properly configured.

  NOTE: Trying to open your session recording in Safeguard Desktop Player while the application is not installed or not properly configured opens a snackbar notification, taking you directly to the *Installing Safeguard Desktop Player* section in the One Identity Safeguard for Privileged SessionsSafeguard Desktop Player User Guide.

- At least one recorded session is available on the **Sessions** page.

Replaying session recordings with Safeguard Desktop Player started from the SPS UI

With Safeguard Desktop Player installed, and at least one session recorded, you can start replaying started from the SPS UI.

To replay your session recordings with Safeguard Desktop Player started from the SPS UI

1. On the **Sessions** page, select the session recording you want to replay.
2. To display the details of a connection, click ••• . This page is called the details view.
3. Next to ▶ Play in browser, click the playback options (⋮) and choose Play in Safeguard Desktop Player.

   Depending on your browser, a confirmation window may appear. Verify that you want to open the Safeguard Desktop Player application.

   TIP: You can set your browser to always allow your `localhost:<port-number>` to open session recordings in Safeguard Desktop Player.
4. (Optional) For encrypted sessions, upload any missing private keys or certificates, then click Play in Safeguard Desktop Player.

   Safeguard Desktop Player opens in a separate window.

   NOTE: SPS saves your most recently selected replay method for recorded sessions. Depending on which method you used before logging out previously, the SPS UI will display ▶ Play in browser or Play in Safeguard Desktop Player when you next log in to SPS and navigate to the details view of your sessions. You can use replay options (⋮) to switch between replay methods.
5. Authenticate with your SPS username and password.

   The recorded session starts replaying.

   For more information about using the Safeguard Desktop Player, see the One Identity Safeguard for Privileged SessionsSafeguard Desktop Player User Guide.

Viewing encrypted screenshots

This section provides information on how to view encrypted screenshots and on the private keys that are necessary to view them.

Prerequisites

To view encrypted screenshots in the Search interface, you have to upload the necessary private encryption keys to your audit keystore. Only RSA keys (in PEM-encoded X.509 certificates) can be uploaded to the audit keystore.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

The private keys are stored locally, in your browser.

NOTE: If you clear your browser's data, your audit keystore could be deleted. If your audit keystore is deleted, upload your private keys again.

To view encrypted screenshots

1. On the **Sessions** page, select an encrypted audit trail.
2. To display the details of a connection, click •••. This page is called the details view.
3. To access the list of encrypted screenshots of the selected audit trail, click **Events** in the details view.
4. Click the screenshot that you want to view.
5. (Optional) Unlock your keystore.

   If you have private keys stored in your audit keystore, the **Unlock keystore** dialog is displayed.

   Enter your master password and click **Unlock keystore**.

   The keystore is unlocked and you can use your keys or you can add new keys.
6. (Optional) If there is a missing private key, the **Private keys missing** dialog opens.

   Click **Upload private keys**, which takes you to **User menu**> **Audit keystore**, where you can add the keys.

   - If you have not added any private keys yet to your audit keystore, see section Adding the first private key to your audit keystore.
   - If you have previously added private keys to your audit keystore, see section Adding further private keys to your audit keystore.
7. (Optional) If necessary, add any missing certificates on **Basic settings** > **Local services** > **Indexer service**.

   Previously, the audit keystore was used to store certificates as well as private keys. From SPS version 6.10 and onwards, you must upload the certificates to **Basic settings** > **Local services** > **Indexer service**.

   For more information on how to add certificates, see Configuring the internal indexer.

Result

Once you add the necessary private keys to your audit keystore, you can view the encrypted screenshots in the **Sessions** interface.

Replaying encrypted audit trails in your browser

This section provides information on the necessary private keys and certificates for playing encrypted audit trails in your browser and on how to open a video file of an audit trail.

Prerequisites

To replay encrypted audit trails in your browser, you have to upload the necessary private keys to your audit keystore and the corresponding certificates to **Basic settings** > **Local services** > **Indexer service**. Depending on the encryption, decrypting the upstream part of an audit trail may require an additional set of certificates and keys.

Only RSA keys (in PEM-encoded X.509 certificates) can be uploaded to the private keystore.

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

NOTE: Certificates are used as a container and delivery mechanism. For encryption and decryption, only the keys are used.

Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates (for example, do not use the certificate of the One Identity Safeguard for Privileged Sessions (SPS) webserver to encrypt audit trails, or the same keypair for signing and encrypting data).

For more information, see Encrypting audit trails on page 512 and Replaying audit trails in your browser.

NOTE: You can only upload certificates permanently. Uploading certificates temporarily is not supported anymore.

To replay encrypted audit trails in your browser

1. On the **Sessions** page, select the audit trail you want to replay.
2. To display the details of a connection, click ••• . This page is called the details view.

   NOTE: If you select the **Automatic refresh** option in **User menu** > **Preferences**, all the details of the sessions are refreshed automatically in the details view. For more information on the **Automatic refresh** option, see Preferences.
3. To play the video file, click **Play video**.
4. (Optional) Unlock your keystore.

   If you have private keys stored in your audit keystore, the **Unlock keystore** dialog is displayed.

Enter your master password and click **Unlock keystore**.

The keystore is unlocked and you can use your keys or you can add new keys.

5. (Optional) If there is a missing private key, the **Private keys missing** dialog opens.

   Click **Upload private keys**, which takes you to **User menu**> **Audit keystore**, where you can add the keys.

   - If you have not added any private keys yet to your audit keystore, see section Adding the first private key to your audit keystore.
   - If you have previously added private keys to your audit keystore, see section Adding further private keys to your audit keystore.

6. (Optional) If necessary, add any missing certificates on **Basic settings** > **Local services** > **Indexer service**.

   Previously, the audit keystore was used to store certificates as well as private keys. From SPS version 6.10 and onwards, you must upload the certificates to **Basic settings** > **Local services** > **Indexer service**.

   For more information on how to add certificates, see Configuring the internal indexer.

   After uploading the private keys and the corresponding certificates, the encrypted upstream traffic elements are decrypted. The trail is decrypted and decompressed on the client-side. As a result, the encrypted elements will be displayed distributed in the video (for more information, see Player window controls in Replaying audit trails in your browser).

7. (Optional) After uploading any missing private keys or certificates, to play the video, click **Play video**.

   1. (Optional) If there are more channels that can be played, select the channel that you want to play in the **Multiple channels are available** dialog and click **Play video**.

      The Player window opens.

8. (Optional) To download the video file, click **Download audit trail**.

For more information on how to use the browser to play video files, see Using the browser to play video files.

Following active sessions

If a session is not closed but it is still active, the active connection label is displayed in the session details window and you can follow or terminate active sessions.

Following active sessions

1. On the **Sessions** page, select the audit trail you want to replay.
2. To display the details of a connection, click ••• . This page is called the details view.

NOTE: If you select the **Automatic refresh** option in **User menu** > **Preferences**, all the details of the sessions are refreshed automatically in the details view. For more information on the **Automatic refresh** option, see Preferences.

3. Click Follow .
4. Select where you want to follow the active session.
    a. To follow the session in the browser, select **Follow in browser**.
    b. To follow the session on Safeguard Desktop Player, select **Follow in Safegurad Desktop Player**.

       The trail data is exported in `.srs` format, which you can download and open with Safeguard Desktop Player.

       For more information on the Safeguard Desktop Player, see *Safeguard Desktop Player User Guide*.
5. (Optional) To terminate the active session, click Terminate.

   This button is displayed only if you have the rights to perform this operation.

Creating report subchapters

Creating search-based report subchapters from search results

You can turn any search query or statistics into a subchapter to add to your reports. This is an easy and flexible way of creating reports to monitor traffic, track certain parameters, or get alerted about particular events.

To create a search-based report subchapter from search results

1. Navigate to **Audit** > **Sessions** and define a valid search query.
2. Click **Create report**. The **Create new subchapter** page is displayed, with the **Search query** field populated with your query.
3. In the **Subchapter title** field, add a title to your subchapter.
4. In **Subchapter type**, select the type that fits your query:
    - **Sessions list view**: Displays a list of sessions.

      Set **Number of sessions** and from **Subchapter columns**, select the session parameters to be displayed in a table in the report. You can add a maximum of 10 columns to the table.
    - **Timeline view**: Visualizes the timeline of sessions that meet the criteria of the **Search query** field, depending on the time range (day/month/week) selected

at the **Scheduling & Delivery** step of the report configuration.

- **Statistics view**: Visualizes the statistics data for the option you select in **Field**, for sessions that meet the criteria of the **Search query** field.

  Select a presentation option for your report, such as **List**, **Pie chart**, or **Bar chart**. In **Field**, select the data field to create your statistics on.

5. Select **Add to a report**, and select from the list of available reports.

   Alternatively, to configure a custom report from scratch and include this subchapter in it, select **Include in a new report**. For more information, see Configuring custom reports on page 893.

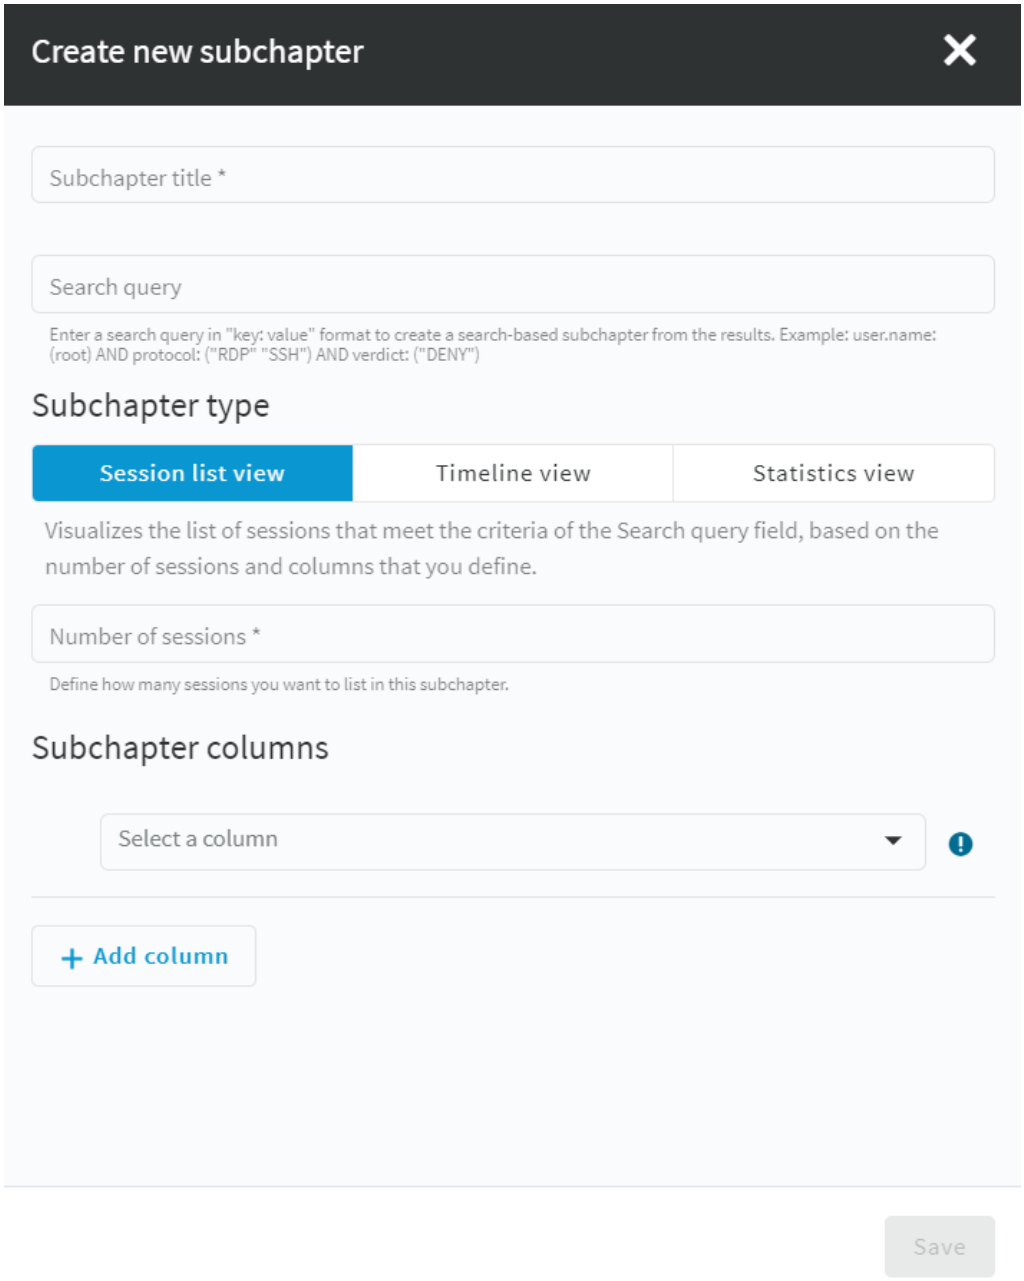
Creating search-based report subchapters from scratch

This section describes how to create a search-based subchapter from scratch to include the subchapter in a custom report.

To create a search-based report subchapter from scratch

1. If you have multiple SPS appliances organized into a cluster where one of the nodes is the Search Master (or Central Search) node, log in to that node.
2. Navigate to **Reporting**> **Create and Manage Reports** > **View & edit subchapters** > **Search-based**.
3. Select **Create new**. The **Create new subchapter** page is displayed.

Figure 293: Reporting> View & edit subchapters > Search-based — Create new subchapter



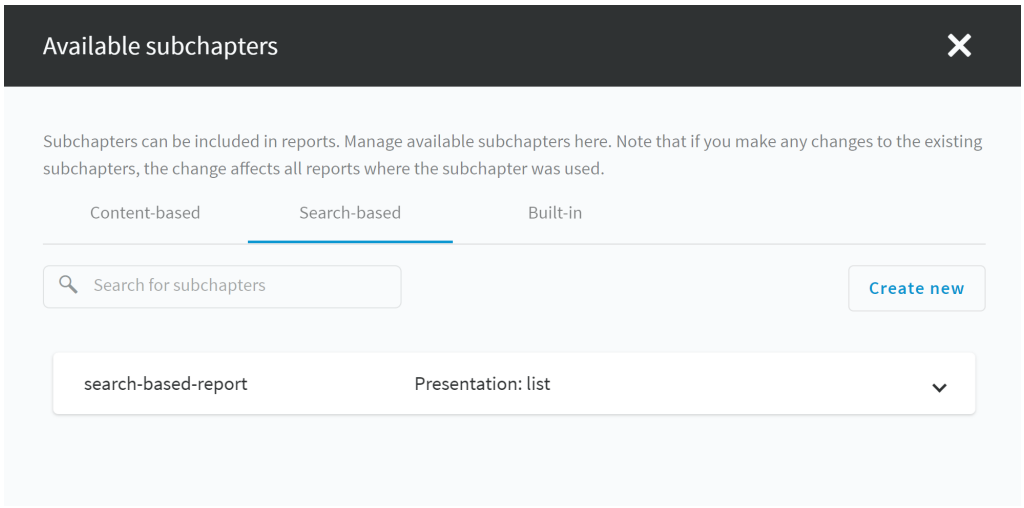
4. In **Subsection title**, add a title to your subchapter.
5. In **Search query**, enter a valid query.

6. In **Subchapter type**, select the type that fits your query:
   - **Sessions list view**: Displays a list of sessions.

     Set **Number of sessions** and from **Subchapter columns**, select the session parameters to be displayed in a table in the report. You can add a maximum of 10 columns to the table.
   - **Timeline view**: Visualizes the timeline of sessions that meet the criteria of the **Search query** field, depending on the time range (day/month/week) selected at the **Scheduling & Delivery** step of the report configuration.
   - **Statistics view**: Visualizes the statistics data for the option you select in **Field**, for sessions that meet the criteria of the **Search query** field.

     Select a presentation option for your report, such as **List**, **Pie chart**, or **Bar chart**. In **Field**, select the data field to create your statistics on.
7. Select **Save**.
8. To save your changes, navigate to **Create and Manage reports** and select **Commit**.
9. Add your subchapter to a new report or to an existing report. For more information, see Configuring custom reports.

   To find and add the subchapter you created to a report, navigate to **Reporting**> **Create and Manage Reports** > **View & edit subchapters** > **Search-based**.



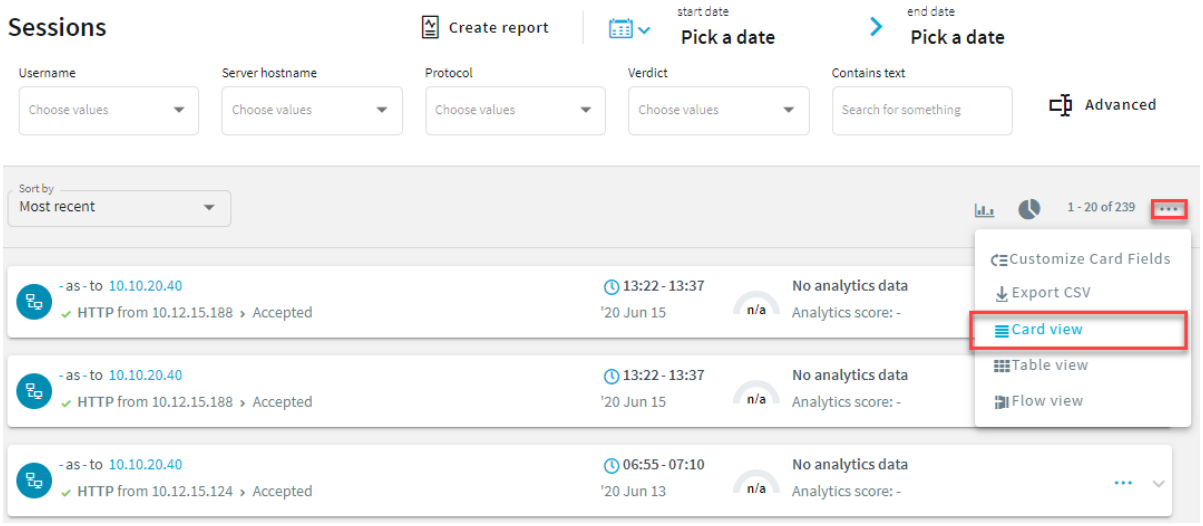
Search interface changes between version 5.0 and 6.0

Between versions 5 LTS (5.0) and 6.0 of One Identity Safeguard for Privileged Sessions (SPS), we have completely redesigned the **Sessions** page, and improved it with several new features. This section highlights the most important changes, and helps you find how to do the common search tasks on the new page. For the detailed documentation of the new **Sessions** page, see Using the Sessions interface on page 710.

Table view, card view

In addition to listing sessions and search results as a table, the new card view highlights the most important details of a session at a glance.

Figure 294: Search interface improvements

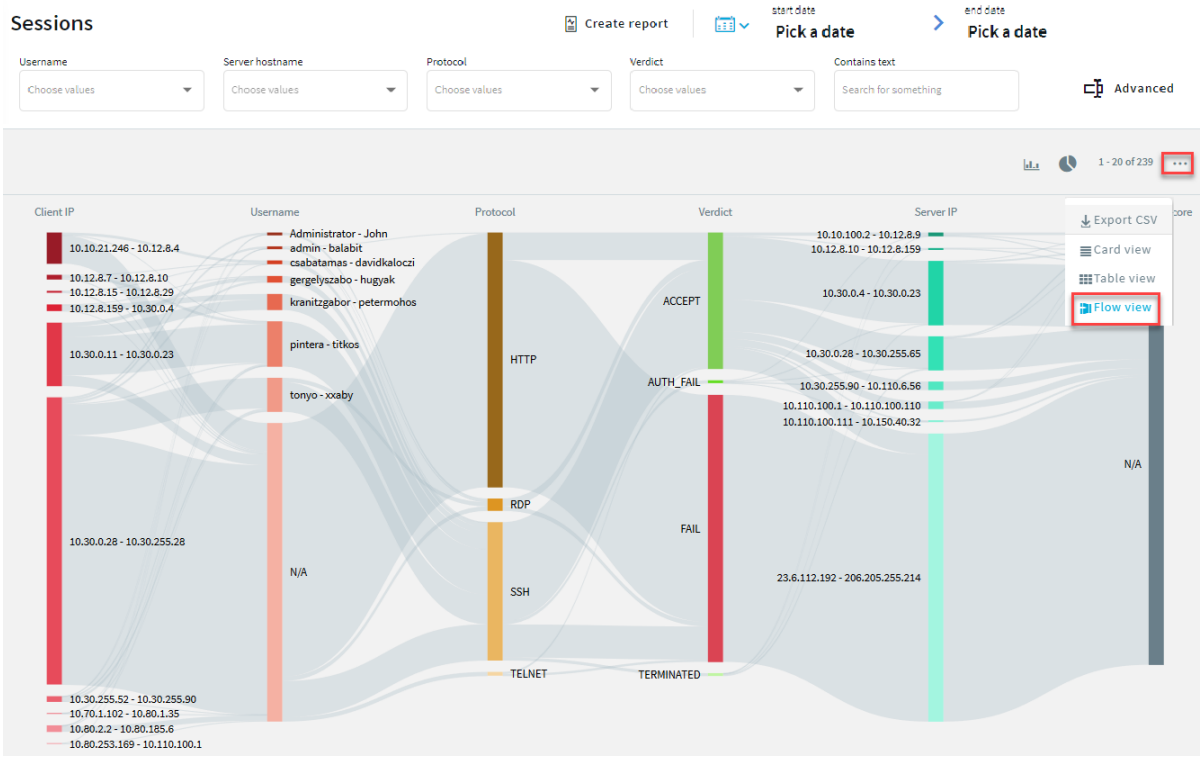


Note that in table view now the list of displayed columns is fixed and cannot be modified. However, if you search for specific values of fields that are not displayed, the values of these fields will be visible in card view.

Quick session analytics with the flow view

Display an interactive, visual overview of your search results to quickly visualize their distribution along multiple attributes, such as client and target IP addresses, protocol, or usernames. Helps to identify patterns in user behavior and to drill down fast to the most relevant sessions. For details, see Using the Sessions interface on page 710.

Figure 295: Audit > Sessions — Flow view



Timeline

The Search interface can now display a timeline showing the search results. Also, you can quickly sort and visualize the distribution of the sessions based on their various metadata, for example, username, server address, and so on.

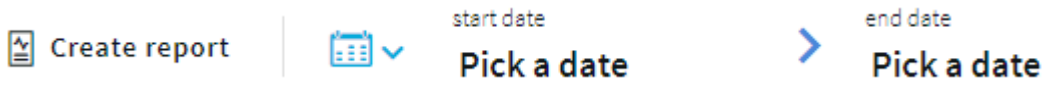
Figure 296: Audit > Sessions — Displaying statistics and timeline



Set a custom or preset date range

Specify a time range to restrict or filter your search criteria by setting boundaries on your searches. Use one of the preset time ranges, or use a custom time range for a more specific search.

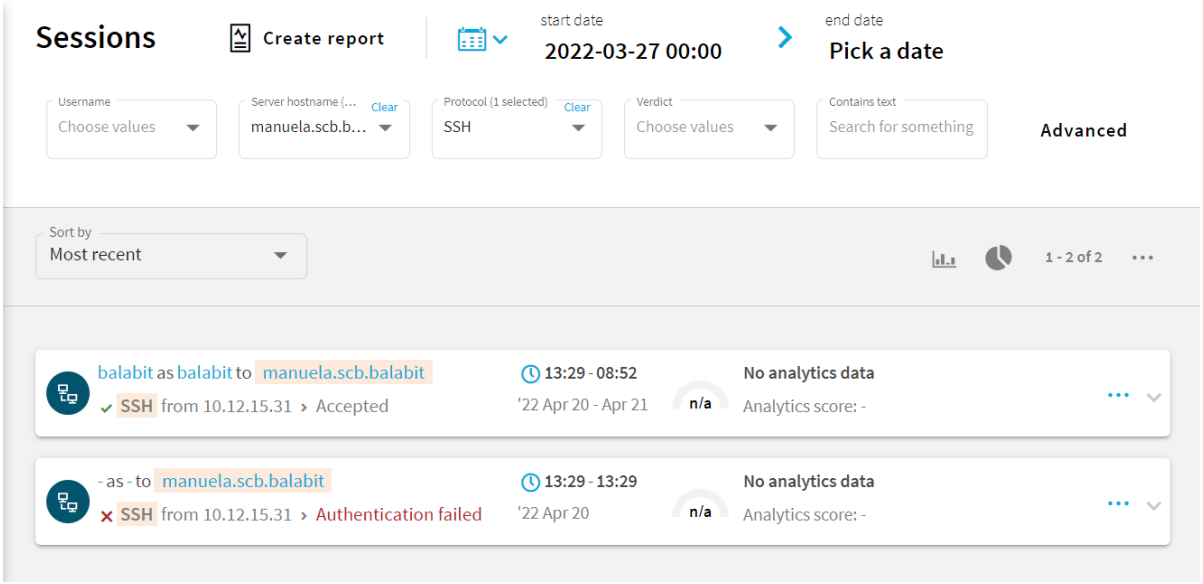
Figure 297: Audit > Sessions — Pick a date



Unified search field

Find everything from a single search box, filter search to a specific field, and combine searches in multiple fields using logical operators. You can also combine content search queries arbitrarily with other search queries. Flow view and quick statistics charts can handle content searches as well. For details, see Using search queries on page 723.

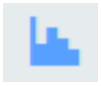
Figure 298: Audit > Sessions — Search filters



Further functions of the old search page (and where they are located on the new search page)

Some functions of the old search page are located somewhere else on the new page. Here is a list of the important ones.

Download the audit trail of a session: Click **details** or **...** to open the details of the session, then click **download trail**. For details, see Viewing session details on page 821.

Display the timeline: Click the  icon. To limit the date range of the search, click **Pick a date** or **Shortcuts**. For details, see Specifying time ranges on page 720.

Change the time interval of the search: To limit the date range of the search, click **Pick a date** or **Shortcuts**. For details, see Specifying time ranges on page 720. Alternatively, you can select a period on the timeline: click at the beginning of the interval, keep the button pressed, then move the pointer to the end of the interval. The timeline and the search results will be updated automatically.

Search in the screen content: To search in the content of the audited sessions, use the `screen.content` field in your search query. For example: `screen.content="exit"`. To search in the contents of a specific session, download the audit trail, open it in the Safeguard Desktop Player application, and use the Search feature of the Safeguard Desktop Player. For details, see Searching in the contents of audit trails on page 799.

Search or filter in a specific field: To search in a specific data field, start typing the name of the field into the search field. The possible field names and their description is automatically displayed. For example, to search for a specific username, use the `user.name:"my-username"` search query. For details, see Viewing session details on page 821.

Alternatively, click  and set the filters you need from the appropriate columns. For example, to search for a specific username, select it using the drop-down menu of the **Username** column. For a more generic search, you can enter any text in the **Contains text** column.

Save a filter or a search query: SPS does not store filters anymore, but you can bookmark the page.

Searching session data on a central node in a cluster

The central search functionality is available when your deployment consists of two or more instances of One Identity Safeguard for Privileged Sessions (SPS) organized into a cluster. When you have a cluster of nodes set up, you have the possibility to search all session data recorded by all nodes in the cluster on a single node. This is achieved by assigning roles to the individual nodes in your cluster: you can set up one of your SPS nodes to be the Search Master and the rest of the nodes to be Search Minions. Search Minions send session data that they record to the Search Master, and the Search Master acts as a central search node.

To set up your environment for central searching, complete the following steps:

1. Enable cluster management on the nodes that you want to be part of your cluster.
2. Build a cluster.
3. Assign roles to nodes in your cluster.

   Familiarize yourself with:

   - The available search roles before assigning them to nodes. For more information, see Cluster roles on page 398.
   - Managing a central search configuration. For more information, see Managing a cluster with central search configuration and configuration synchronization.

Once you have your cluster set up and the appropriate roles assigned, you can start searching session data using the Search interface.

NOTE: Central search is not available on the **Search (classic)** interface.

Limitations of the central search functionality

Currently, the central search functionality comes with the following limitations:

- Session data recorded by a node before it was joined to the cluster will not be searchable centrally. Only session data recorded after the node has been joined to the cluster is available for central search.
- You cannot to run the indexer process on unindexed sessions after assigning the Search Master role to a node. Make sure all important sessions are indexed before assigning the Search Master role to the node.

- The Search Master node cannot run internal indexer processes, nor does it receive connections from external indexers. Indexers work only with Search Minion nodes.
- It is not possible to replay audit trail files in your browser from the Search Master node.
- When near real-time indexing is configured on a Search Minion node, while session data from active connections is visible on the Search interface of the Search Master node, it is not possible to:
    - export the audit trail of an active connection,
    - follow an active connection, and
    - terminate an active connection.

  Note, however, that you can terminate the active, ongoing connection on the Search Minion node that is recording the connection in question.
- A reliable, high-bandwidth connection is required between the nodes. Small loss of connection is handled well but if the connection between the Search Minions and the Search Master is lost for a longer period of time, the Search Minions will stop accepting new connections until the connection is repaired. Data is automatically pushed to the Search Master after the connection is restored.

NOTE: Search Minion nodes do not send the files storing the audit trails to the Search Master node. When a user clicks **Download audit trail**, the Search Master node streams the trail files to the user from the original Search Minion node that recorded the sessions. If a Search Minion node does not have a backup policy set up and an error occurs that causes data loss, then session data recorded by that node will not be available.

Advanced authentication and authorization techniques

This section describes the advanced authentication and authorization techniques available in One Identity Safeguard for Privileged Sessions.

- For details on creating usermapping policies, see Configuring usermapping policies on page 860.
- For details on configuring gateway authentication, see Configuring gateway authentication on page 862.
- For details on configuring four-eyes authorization and real-time monitoring, see Configuring four-eyes authorization on page 870.
- For details on configuring Credential Stores, see Using credential stores for server-side authentication on page 873.

Configuring usermapping policies

For SSH, RDP, Telnet, and Citrix ICA connections, you can define usermapping policies. A usermapping policy describes who can use a specific user name to access the remote server. Only members of the specified local or LDAP user groups (for example, `administrators`) can use the specified user name (for example, `root`) on the server.

CAUTION:

In SSH connections, the users must use the following as their username: **`gu=username@remoteusername`**, where `username` is the username used in the LDAP directory, SPS will use this username to determine their group memberships, and `remoteusername` is the username they will use on the remote server. For example, to access the example.com server as root, use:

```
gu=yourldapusername@root@example.com
```

NOTE: For the user name of SSH users, only valid UTF-8 strings are allowed.

CAUTION:

In Telnet connections, usermapping policies work only if **Extract username from the traffic** is enabled.

For more information, see Extracting username from Telnet connections on page 658.

When configuring ICA connections, also consider the following:

CAUTION:

If the clients are accessing a remote application or desktop that is shared for Anonymous users (that is, the **Users** properties of the application is set to **`Allow anonymous users`** in the Citrix Delivery Services Console), the actual remote session will be running under an Anonymous account name (for example, `Anon001`, `Anon002`, and so on), not under the username used to access the remote server. Therefore, you need to enable usermapping to the `Anon*` usernames.

To accomplish this, create a usermapping policy and set the **Username on the server** option to **`Anon*`**, and the **Groups** option to **`*`**, then use this usermapping policy in your ICA connections.

For more information on using usermapping policies, see Configuring usermapping policies on page 860.

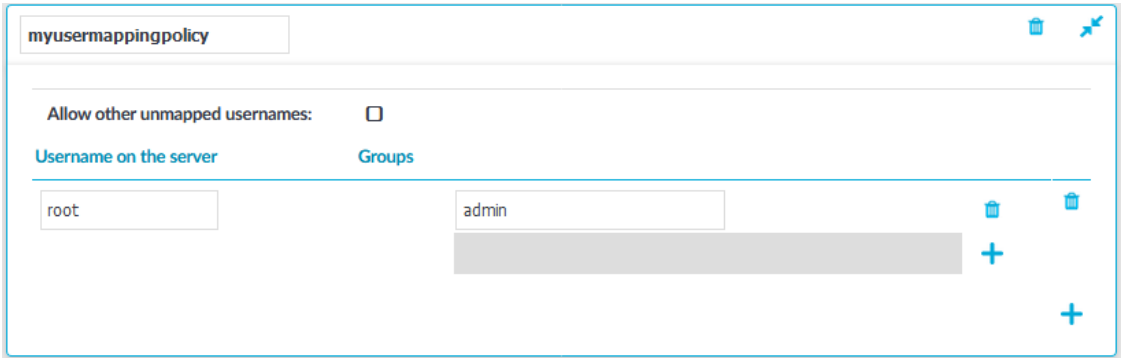
NOTE: Usermapping policies require gateway authentication as well.

For more information, see Configuring out-of-band gateway authentication on page 863.

To configure usermapping

1. Navigate to **Policies** > **Usermapping Policies**.

Figure 299: Policies > Usermapping Policies — Configuring usermapping policies



2. To create a new policy, click + (**Add row**), then enter a name for the policy.
3. Click + (**Add row**), then in the **Username on the server** field, enter the user name that can be used to access the remote server (for example, **`root`**). SPS will use this user name in the server-side connection. To permit any user name on the server side, enter an asterisk (**`*`**).
4. Select **Groups**, click + (**Add row**), then specify who is permitted to use the remote user name specified in the **Username on the server** field.

- (Optional) If you have an LDAP server set in the connection policy where you will use usermapping, enter the name of the local or LDAP user group (for example, **admins**) whose members will be permitted to use the remote user name.

  For more information on LDAP authentication, see Authenticating users to an LDAP server on page 505.

  NOTE: The LDAP server configured in the connection policy is not necessarily the same as the LDAP server used to authenticate the users accessing the SPS web interface.

- If you do not authenticate the connections to an LDAP server, enter the name of the userlist whose members will be permitted to use the remote username.

  For more information on using userlists, see Creating and editing user lists on page 504.

If needed, to add further groups, repeat this step.

5. Repeat steps 3-4 to add further usernames if needed.
6. To permit other users who are not explicitly listed in the Usermapping Policy access the remote servers, select the **Allow other unmapped usernames** option.

   NOTE: These users must use the same user name on the SPS gateway and the remote server.
7. Click **Commit**.
8. Navigate to the **Connections** page of the traffic (for example to **SSH Control** > **Connections**), and select the connection policy to modify.
9. Select the usermapping policy created in Step 2 from the **Usermapping policy** field.
10. Click **Commit**.

Configuring gateway authentication

When gateway authentication is required for a connection, the user must authenticate on One Identity Safeguard for Privileged Sessions (SPS) as well. This additional authentication can be performed:

- Out-of-band, on the SPS web interface, for every protocol.
- Inband, using the incoming connection, for the SSH, Telnet, and RDP protocols.

For details about the concepts of gateway authentication, see The gateway authentication process. You can use gateway authentication to authenticate the real person when the user is using a shared account to access the target server.

NOTE: For SSH, Telnet, and RDP connections, gateway authentication can be performed also inband, without having to access the SPS web interface.

- For SSH and Telnet connections, inband gateway authentication must be performed when client-side authentication is configured. For details on configuring client-side authentication, see Client-side authentication settings on page 627.

- For RDP connections, inband gateway authentication must be performed when SPS is acting as a Remote Desktop Gateway (or RD Gateway). In this case, the client authenticates to the Domain Controller or a local user database. For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway on page 595.

  In the case of RDP connections, inband gateway authentication can also be performed if an AA plugin is configured.

NOTE: Gateway authentication can be used together with other advanced authentication and authorization techniques like four-eyes authorization, client- and server-side authentication, and so on.

CAUTION:

If the username used within the protocol to access the remote server is different from the username used to perform gateway authentication (for example, because the user uses a shared account in the remote server, but a personal account for gateway authentication), usermapping must be configured for the connection. For details on usermapping, see Configuring usermapping policies on page 860.

NOTE: To configure a credential store for gateway authentication, see Using credential stores for server-side authentication on page 873.

Configuring out-of-band gateway authentication

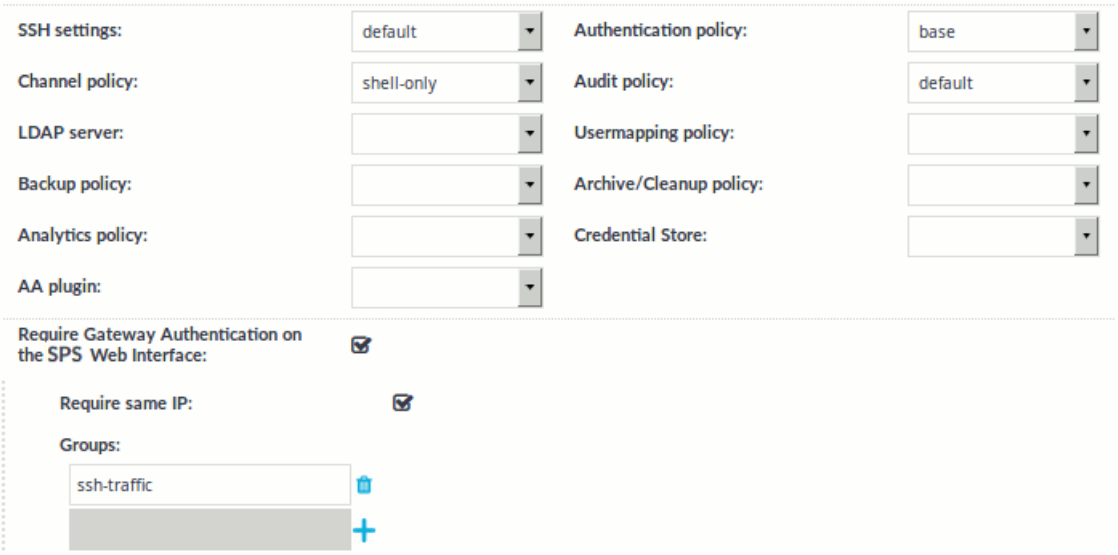
CAUTION:

- The `admin` user is a special One Identity Safeguard for Privileged Sessions (SPS) user and not a member of any user groups, nor can it belong to any group. Since usermapping policies are based on user groups, performing gateway authentication with the `admin` user is likely to result in usermapping errors.
- When using SSL-encrypted RDP connections, or connections that use the Credential Security Service Provider (CredSSP) authentication method, some Microsoft RDP clients restart the connection during the authentication process. This would require the user to perform gateway authentication on the SPS web interface twice. To avoid this situation, SPS temporarily caches the successful gateway authentication results if the client terminates the connection at a certain step while establishing the connection. The cache is used to automatically authenticate the restarted connection without user interaction.
- If the clients are behind a device that performs network address translation (NAT), it will seem to SPS as if every connection was initiated from the same IP address. Therefore, in such cases using out-of-band gateway authentication is not recommended for security reasons, especially for RDP connections. If possible, use inband gateway authentication instead.

To configure gateway authentication

1. Navigate to the **Connections** page of the traffic (for example to **Traffic Controls** > **SSH** > **Connections**), and select the connection policy to modify.
2. Select the **Require Gateway Authentication on the SPS Web Interface** option. This is the option to configure gateway authentication via the web interface of SPS.

Figure 300: <Protocol name> Connections > Require Gateway Authentication on the SPS Web Interface — Configuring gateway authentication



3. To accept the gateway authentication only from the host that initiated the connection, select **Require same IP**.

   NOTE: This option has no effect if the clients are behind a device that performs network address translation (NAT). In such cases, use inband gateway authentication instead.

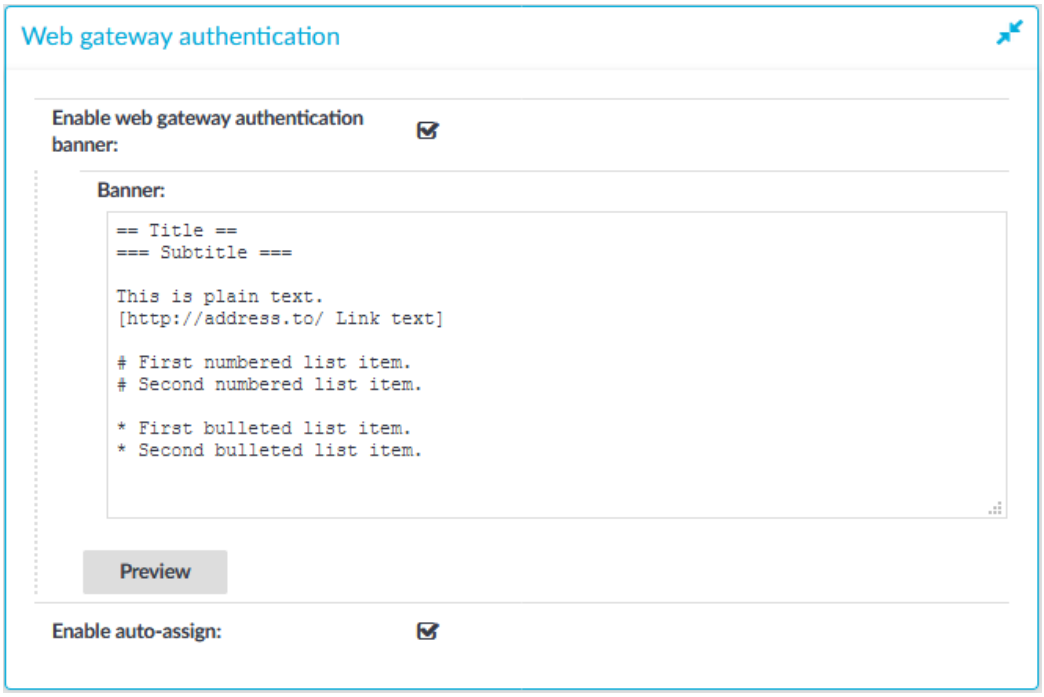
4. By default, any user can perform gateway authentication for the connections. To allow only members of a specific group authenticate the connections of this connection policy, select **Groups**, click + and enter the name of the group whose members can authenticate the connections. This group must exist on the **Users & Access Control** > **Local User Groups** page. For details on creating and managing usergroups, see Managing user rights and user groups on page 361. Repeat this step to add further groups if needed.
5. For SSH, RDP, Telnet and Citrix ICA connections, you can set a usermapping policy in the **Usermapping policy** field. For details on usermapping policies, see Configuring usermapping policies on page 860.
6. Click **Commit**. After that, users accessing these connections must perform gateway authentication as described in Performing out-of-band gateway authentication on One Identity Safeguard for Privileged Sessions (SPS) on page 867.

7. (Optional) To restrict the availability of selected channels of the connection based on the username used for gateway authentication, edit the channel policy used in the connection.
    a. Navigate to the channel policy used in the connection (for example, **Traffic Controls** > **SSH** > **Channel Policies**).
    b. Select **Gateway Group**, click + and enter the name of the user group allowed to use this type of the channel. The user group must correspond to the username used for the gateway authentication. Repeat this step until all permitted groups are listed.

       You may list local user lists as defined in Creating and editing user lists on page 504, or LDAP groups (for details on accessing LDAP servers from SPS, see Authenticating users to an LDAP server on page 505).

       NOTE: Consider the following behaviors of SPS:
       - If you list multiple groups, members of any of the groups can access the channel.
       - When listing both a whitelist and a blacklist in the **Gateway Group** section and a username appears on both lists, the user will be able to access the channel.
       - If a local user list and an LDAP group has the same name and the LDAP server is configured in the connection that uses this channel policy, both the members of the LDAP group and the members of the local user list can access the channel.
    c. Click **Commit**.
8. (Optional) If you want to provide a limited SPS web interface to your users that can be used only for gateway authentication and 4-eyes authorization, set up a dedicated user-only web login address. For details, see Configuring user and administrator login addresses on page 117.
9. (Optional) You can configure a message for users accessing SPS for out-of-band authentication. The message is displayed when they log in to SPS.
    a. Navigate to **Basic Settings** > **Management** > **Web gateway authentication**.
    b. Select **Enable web gateway authentication banner**.

Figure 301: Basic Settings > Management > Web gateway authentication — Configuring a banner



c. Enter the message in the **Banner** field. You can use the following text formatting options:

```
== Title ==
=== Subtitle ===

This is plain text.
[http://address.to/ Link text]

# First numbered list item.
# Second numbered list item.

* First bulleted list item.
* Second bulleted list item.
```

d. Click **Commit**.

10. (Optional) If your users have sessions to several remote servers, or they access a server several times a day, performing the gateway authentication for every session can be a nuisance. To permit your users to authenticate on the SPS web interface once, and open sessions without repeating the gateway authentication, select **Enable auto-assign** and click **Commit**.

    NOTE: For auto-assign to work, users must leave the browser window (or tab) of SPS open.

Performing out-of-band gateway authentication on One Identity Safeguard for Privileged Sessions (SPS)

The following describes how to perform out-of-band gateway authentication on One Identity Safeguard for Privileged Sessions (SPS).

To perform out-of-band gateway authentication on SPS

1. Initiate a connection from a client. If gateway authentication is required for the connection, SPS will pause the connection.

   NOTE: For SSH and Telnet connections, when initiating the connection, you can use the following as your username: `gu=gatewayusername@remoteusername`, where `gatewayusername` is the username you will use to login to the SPS web interface (also called gateway user), and `remoteusername` is the username you will use on the remote server.

   NOTE: After initiating the connection, the administrator with the appropriate authorization rights has 3 minutes to approve the request.

2. Open a browser, preferably on the same host you initiated the connection from, and navigate to the login page of SPS.

   CAUTION:

   If the username used within the protocol is different from the username used to access the SPS web interface to perform gateway authentication, usermapping must be configured for the connection. For details on usermapping, see Configuring usermapping policies on page 860.

3. Log in to SPS, and select **Pending Connections** > **Gateway Authentication** from the main menu. The list of connections waiting for gateway authentication will be displayed.

   NOTE: If users accessing the SPS web interface are authenticated to an LDAP server, the users must successfully authenticate to the LDAP server set on the **Users & Access Control** > **Settings** page.

   Figure 302: Pending Connections > Gateway Authentication — Performing gateway authentication

4. Select the connection that you started, and click **Assign**.

5. Continue to authenticate on the server.
6. To authenticate another session, you must either:
    - Repeat this procedure.
    - If your SPS administrator enabled the auto-assign feature, you must keep open the browser window or tab in which you authenticated to SPS.

Performing inband gateway authentication in SSH and Telnet connections

The following describes how to perform inband gateway authentication in SSH and Telnet connections.

To perform inband gateway authentication in SSH and Telnet connections

1. Initiate a connection from a client. If gateway authentication is required for the connection, One Identity Safeguard for Privileged Sessions (SPS) will pause the connection.
2. SPS requests the username used for gateway authentication. Enter your gateway username into the **Gateway username** prompt. If password authentication is used, provide the password for the gateway user as well.
3. The login prompt for the remote server is displayed. Enter your username used on the remote server into the **Username** prompt. If password authentication is used, provide the password for the username as well.

    CAUTION:

    If the username used within the protocol to access the remote server is different from the username used to perform gateway authentication, usermapping must be configured for the connection. For details on usermapping, see Configuring usermapping policies on page 860.

    NOTE: When initiating the connection, you can use the following as your username: `gu=gatewayusername@remoteusername`, where `gatewayusername` is the username you will use to authenticate on SPS and `remoteusername` is the username you will use on the remote server. That way you do not have to provide the usernames in the prompt, only the passwords if password authentication is used.

    If SPS is configured to require client-side authentication, the `gatewayusername` user must authenticate on the client side.

Performing inband gateway authentication in RDP connections

The following describes how to perform inband gateway authentication in RDP connections.

To perform inband gateway authentication in RDP connections

1. Initiate a connection from a client.
2. The graphical login window is displayed.
    - If the **Advanced** > **Remote Desktop Gateway** > **Logon Settings** > **Use my Remote Desktop Gateway credentials for the remote computer** option of your Remote Desktop application is enabled, login to the remote server using your usual credentials. One Identity Safeguard for Privileged Sessions (SPS) will use these credentials for the gateway authentication on the Domain Controller as well.
    - If the **Advanced** > **Remote Desktop Gateway** > **Logon Settings** > **Use my Remote Desktop Gateway credentials for the remote computer** option of your Remote Desktop application is disabled, first you have to authenticate on the SPS gateway. Enter your username and password for the Domain Controller.

      If the first authentication is successful, a second login window is displayed. Enter your username and password for the remote server you are trying to access.
    - If SPS is configured to use a Credential Store to login to the target server, enter the following:
        - In the **Username** field, enter the domain name, the `-AUTO` suffix, and your username. For example, `EXAMPLEDOMAIN-AUTO\Administrator`.

          NOTE: The `-AUTO` suffix is the default value of the **Traffic Controls** > **RDP** > **Settings** > **Autologon domain suffix** option of One Identity Safeguard for Privileged Sessions (SPS). If your SPS administrator has changed this option, use the appropriate suffix instead of -AUTO.
        - Enter your username (only the username, without the domain, for example, **`Administrator`**) into the **Password** field.
3. If the authentication is successful, the desktop of the remote server is displayed.

Troubleshooting gateway authentication

If a user initiates a connection and then logs in to the One Identity Safeguard for Privileged Sessions (SPS) web interface, it might happen that his connection is not shown on the **Pending Connections** > **Gateway Authentication** page. SPS checks the following points to determine if a pending connection is listed for a user:

CAUTION:

The `admin` user is a special One Identity Safeguard for Privileged Sessions (SPS) user and not a member of any user groups, nor can it belong to any group. Since usermapping policies are based on user groups, performing gateway authentication with the `admin` user is likely to result in usermapping errors.

- The username used to access the SPS web interface is a member of a group listed in the **Gateway authentication** > **Groups** field of the connection policy.
- If SPS knows from the protocol the username that will be used to access the SPS web interface to perform the gateway authentication, the connection is displayed only to this user.

  For SSH connections, SPS can determine the username if:

  - The client specifies the username for the web interface within the protocol using the `gu=webusername@server-side-username@server` format.
  - The client specifies the username within the protocol using an interactive prompt.
  - If the client does not use any of the above options, SPS uses the remote username. In this case, the username for the web interface must be the same as the remote username, otherwise the connection is not displayed.
- If the **Gateway authentication** > **Require same IP** option is enabled, the pending connection is displayed only if the user accesses the SPS web interface from the same IP address as the client in the pending connection.

NOTE: The `admin` user sees every pending connection.

Configuring four-eyes authorization

When four-eyes authorization is required for a connection, a user (called authorizer) must authorize the connection on One Identity Safeguard for Privileged Sessions (SPS) as well. This authorization is in addition to any authentication or group membership requirements needed for the user to access the remote server. For details about the concepts of four-eyes authorization, see Four-eyes authorization on page 60.

Configuring four-eyes authorization

The following describes how to configure four-eyes authorization.

To require four-eyes authorization for a connection

1. Navigate to the **Connections** page of the traffic (for example to **Traffic Controls** > **SSH** > **Connections**), and select the connection policy to modify.
2. Expand the view of the connection policy and scroll down to the **Access Control** section, then click  .
3. Enter the name of the usergroup whose members are permitted to authorize the sessions of the connection policy into the **Authorizer Group** field. This group must exist on the **Users & Access Control** > **Local User Groups** page. For details on creating and managing usergroups, see Managing user rights and user groups on

page 361.

CAUTION:

Usernames, the names of user lists, and the names of usergroups are case sensitive.

4. By default, the authorizer can authorize any session of the connection policy.

   If the authorizer is permitted to authorize only the sessions of a certain usergroup, select **Client user is** > **Member of**, and enter the name of the userlist whose sessions the authorizer can authorize. If you use four-eyes authorization without gateway authentication, you can specify an LDAP group instead of a userlist.

   CAUTION:

   Usernames, the names of user lists, and the names of usergroups are case sensitive.

   CAUTION:

   When using both gateway authentication and four-eyes authorization in a Connection Policy, specify the usergroup of the gateway username. The specified group must be a local or LDAP group.

5. Set the permissions of the usergroup set in the **Authorizer Group** field.

   - If the **Authorizer** group can authorize (that is, enable) and audit (that is, monitor in real-time and download the audit trails) the sessions, select **Permission** > **Follow&Authorize**.
   - If the **Authorizer** group can only authorize (that is, enable) the sessions, select **Permission** > **Authorize**.

     NOTE: This option is not valid for HTTP connections.
   - If the **Authorizer** group can only audit (that is, monitor in real-time and download the audit trails) the sessions, select **Permission** > **Follow**.

   NOTE: If the **Client user is** > **Member of** field is set, the auditor can only monitor the sessions of the specified usergroup. However, if **Client user is** > **Member of** field is set, the Auditor cannot access the **Sessions** page. To avoid this problem, add another Access Control rule for the **Authorizer Group** without setting the **Client user is**field.

   The `admin` user of One Identity Safeguard for Privileged Sessions (SPS) can audit and authorize every connection.

6. To ensure that the client and the authorizer use different IP addresses and thus prevent self-authorization, enable **Require different IP**. If this is enabled, and the client and the authorizer do not have different IP addresses, it disables all actions for the connection and the four-eyes authorization, until they have different IP addresses.
7. To ensure that the client and the authorizer use different usernames and thus prevent self-authorization, enable **Require different Username**. If this is enabled,

and the client and the authorizer do not have different usernames, it disables all actions for the connection and the four-eyes authorization, until they have different usernames.

8. Repeat steps 2-6 to add other authorizers or usergroups if needed.
9. Click **Commit**.
10. Navigate to the **Channel Policies** page of the traffic (for example, to **Traffic Controls** > **SSH** > **Channel Policies**), and select the channel policy used in the connection.
11. Enable the **Four-eyes** option for the channels which should be accessed only using four-eyes authorization.

    NOTE: If a connection uses secondary channels that require four-eyes authorization — for example, a Remote Desktop connection allows a Drawing channel but requires four-eyes authorization for a Disk redirection channel — the connection is locked until the authorizer accepts the channel on the **Four-Eyes** page of SPS, or the four-eyes request times out. During this time, the client application can become nonresponsive, for example, display the graphical desktop but not react to mouse clicks.

    NOTE: In Citrix ICA connections, four-eyes authorization is required before the user logs in to the destination server. To request four-eyes authorization only after the log in, when the server-side username is already known, select the **Perform 4 eyes after user login** option.
12. Click **Commit**. After that, users accessing connections using the modified channel policy must be authorized as described in Performing four-eyes authorization on One Identity Safeguard for Privileged Sessions (SPS) on page 872.
13. (Optional) If you want to provide a limited SPS web interface to your users that can be used only for gateway authentication and 4-eyes authorization, set up a dedicated user-only web login address. For details, see Configuring user and administrator login addresses on page 117.

Performing four-eyes authorization on One Identity Safeguard for Privileged Sessions (SPS)

The following describes how to perform four-eyes authorization on One Identity Safeguard for Privileged Sessions (SPS).

To perform four-eyes authorization on SPS

1. When a user initiates a connection from a client and four-eyes authorization is required for the connection, SPS will pause the connection.

NOTE: Four-eyes authorization can be set separately for every channel. However, if a client of an existing connection opens a new channel that requires four-eyes authorization, every channel is paused until the authorization is completed.

2. Login to SPS, and select **Four-Eyes** from the main menu. The list of connections waiting for authorization will be displayed.

   NOTE: Only those connections will be listed, where your usergroup has the `Authorize` or the `Follow&Authorize` permissions. No other SPS privilege is required to access this page.

3. Select the connection and click **Accept** to enable the connection, **Reject** to deny the connection, or **Accept&Follow** to enable it and monitor in real-time.

   NOTE: Following a session requires the following:

   - The **Record audit trail** option must be enabled for the specific channel in the Channel policy of the connection.
   - The Audit Player application must be installed on the computer of the auditor.
   - If the Audit policy of the connection uses encryption, the appropriate decryption keys must be available on the computer of the auditor.

   The Safeguard Desktop Player application replays the live streams in `live` mode. For more information on how to monitor a connection in real-time using the Safeguard Desktop Player, see *Replaying audit files in follow mode* in the *Safeguard Desktop Player User Guide*.

4. Enter a note why the connection was accepted/rejected into the appearing dialog box. This description will be stored in the connection database together with other metadata about the connection.

5. If you have to terminate an ongoing connection for some reason, select **Pending Connections** > **Active Connections** from the main menu. The list of ongoing connections will be displayed.

6. Select the connection to stop, and click **Terminate**.

   NOTE: When following a connection in the Safeguard Desktop Player application, the auditor can also terminate the connection from the Audit Player by clicking **Terminate**.

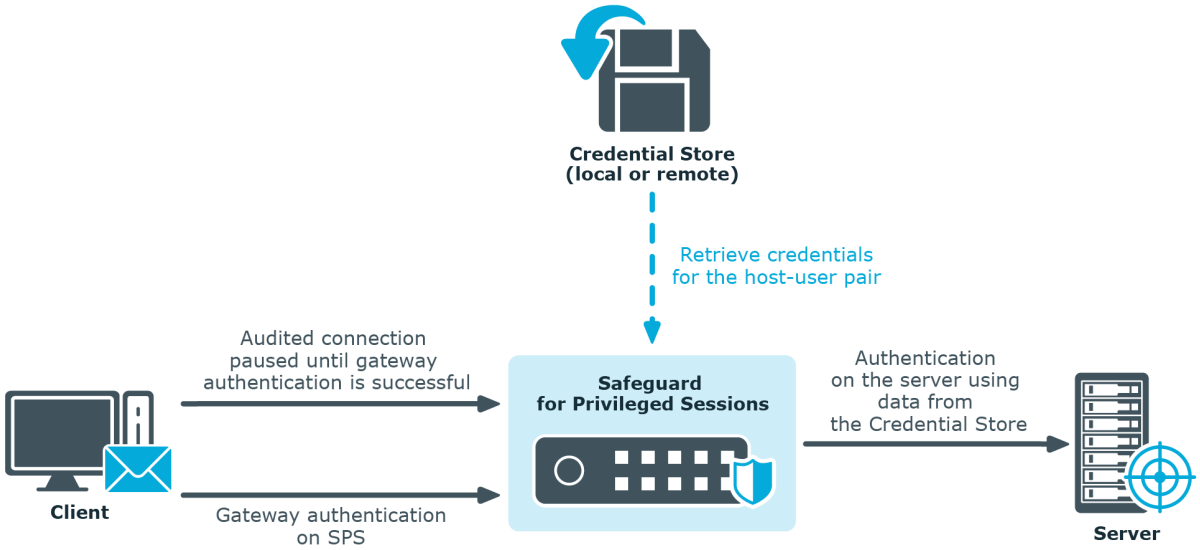
Using credential stores for server-side authentication

Credential Stores offer a way to store user credentials (for example, passwords, private keys, certificates) and use them to log in to the target server, without the user having access to the credentials. That way, the users only have to perform gateway authentication on SPS with their usual password (or to an LDAP database), and if the user is allowed to access the target server, SPS automatically logs in using the Credential Store.

For more information on gateway authentication, see Configuring gateway authentication on page 862.

NOTE: Keyboard-interactive authentication is not supported when using credential stores.

Figure 303: Authenticating using Credential Stores



Credential Stores can be stored locally on SPS, or on a remote device. For remote Credential Stores, SPS integrates with external authentication and authorization systems using plugins.

- To configure a local Credential Store, see Configuring local Credential Stores on page 874.
- To configure a local, password-protected Credential Store, see Configuring password-protected Credential Stores on page 878.
- To unlock a local, password-protected Credential Store, see Unlocking Credential Stores on page 882.
- To configure a custom Credential Store plugin, see Using a custom Credential Store plugin to authenticate on the target hosts on page 883.

NOTE: After performing a successful gateway authentication, if the credential store does not contain a password for the user, the user is prompted for the server-side password as a fallback.

In case of authenticating to RDP servers using Network Level Authentication (NLA), the server-side password is prompted at the start of the connection. If there is no password in the credential store for the user and the server-side password is incorrect, the connection is terminated.

Configuring local Credential Stores

The following describes how to configure a local Credential Store that stores the credentials used to login to the target host.

Prerequisites

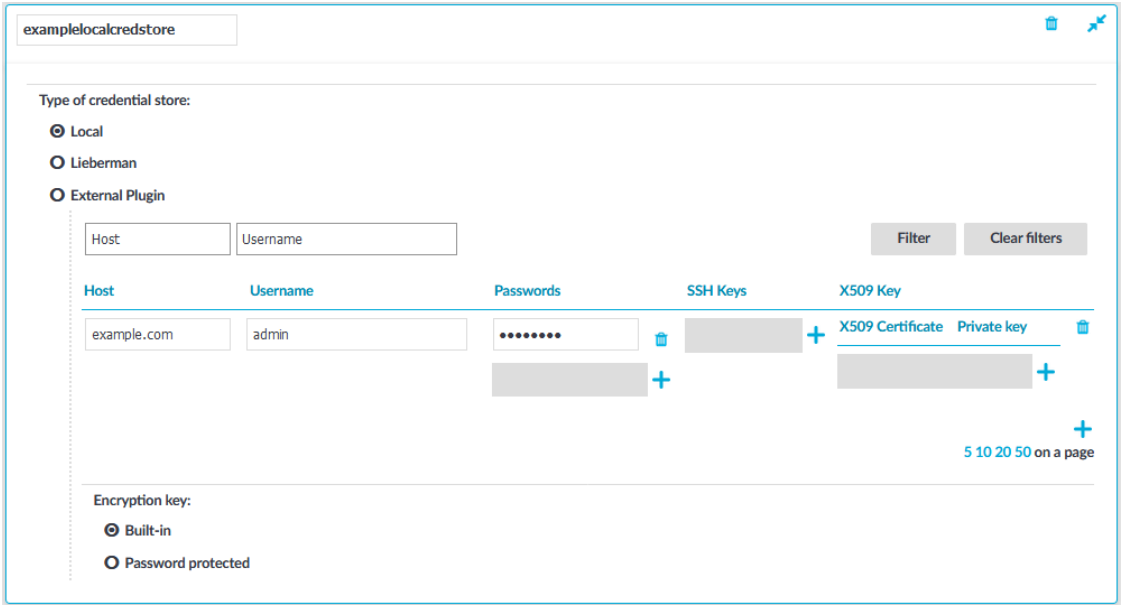
NOTE: Users accessing connections that use Credential Stores to authenticate on the target server must authenticate on SPS using gateway authentication or an AA plugin. Therefore gateway authentication or an AA plugin must be configured for these connections.

For more information, see *Configuring gateway authentication* in the *Administration Guide* and *Integrating external authentication and authorization systems* in the *Administration Guide*.

To configure a local Credential Store that stores the credentials used to login to the target host

1. Navigate to **Policies** > **Credential Stores**.
2. Click + and enter a name for the Credential Store.
3. Select **Local**.
4. Select **Encryption key** > **Built-in**. That way the credentials will be encrypted with a built-in password, and the Credential Store is automatically accessible when SPS boots up. To use custom passwords to encrypt the Credential Store, see Configuring password-protected Credential Stores on page 878.

Figure 304: Policies > Credential Stores > Local — Configuring local Credential Stores



5. Add credentials to the Credential Store.

   a. Click  and enter the destination host and the username. For the destination host, you can use hostname, IP address, or subnet as well. To use the same credentials for every destination host, enter the **`0.0.0.0/0`** subnet. To use the credentials only on the hosts of a specific domain, enter **`*.domain`**. Note that:

      - Usernames are case sensitive.
      - To authenticate users of a Windows domain, enter the name of the domain into the **Host** field.

      Use an IPv4 address.

   b. Set the credentials. SPS will use these credentials to login to the destination host if the credential store is selected in a Connection policy. If more than one credential is specified to a host-username pair, SPS will attempt to use the credentials as the destination host requests it.

      - To add a password, click **Passwords >**  , then enter the password corresponding to the username.
      - To upload a private key, click **SSH Keys >**  **>**  , then paste or upload a private key.

        NOTE: If the private key is protected by a passphrase, enter the passphrase. The passphrase is needed only once during the upload, it is not required for the later operation of the Credential Store.
      - To generate a keypair on SPS click **SSH Keys >**  **>**  , set the length and type of the key, then click **Generate**. After that, click the fingerprint of the key to download the public part of the keypair. There is no way to download the private key from the SPS web interface.

        NOTE:

        TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
      - To upload a certificate and the corresponding private key, click **X509 Keys >**  **>**  , then paste or upload a certificate and the private key.

        NOTE: If the private key is protected by a passphrase, enter the passphrase. The passphrase is needed only once during the upload, it is not required for the later operation of the Credential Store.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

c. Repeat the previous step to add further credentials to the username as necessary.

6. Repeat the previous step to add further hosts or usernames as necessary.

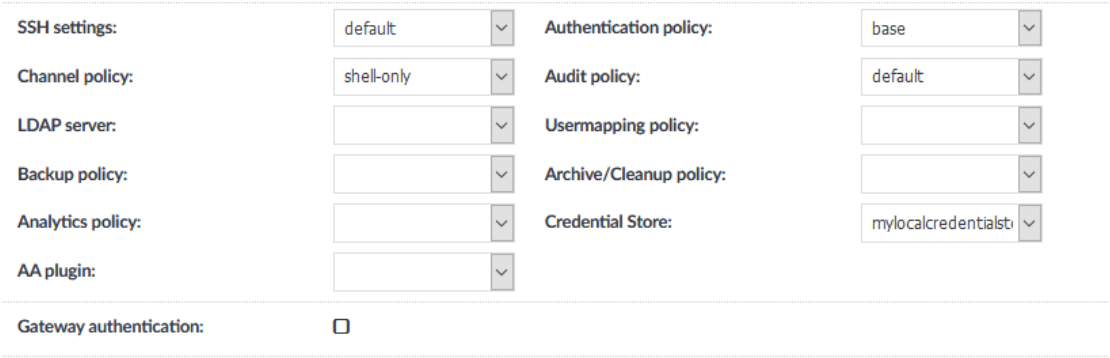
   NOTE: Credential Stores can be used together with usermapping policies to simplify the administration of users on the target hosts.

   For more information, see Configuring usermapping policies on page 860.

7. Click **Commit**.
8. Navigate to the Connection policy where you want to use the Credential Store (for example, to **Traffic Controls** > **SSH** > **Connections**), select the Credential Store to use in the **Credential Store** field, then click **Commit**.

   NOTE: The Connection Policy will ignore the settings for server-side authentication (set under **Relayed authentication methods**) if a Credential Store is used in the Connection Policy.

Figure 305: Traffic Controls > Protocol name > Connections — Select a Credential Store to use



Performing gateway authentication to RDP servers using local Credential Store and NLA

The following describes how to perform a gateway authentication to RDP servers using local Credential Store and Network Level Authentication (NLA).

To perform a gateway authentication to RDP servers using local Credential Store and NLA

1. Initiate the RDP connection.
2. Enter your gateway credentials during the gateway authentication. This can be web gateway authentication, or inband gateway authentication using RD Gateway.
3. Enter the following:
    - In the **Username** field, enter the domain name, the `-AUTO` suffix, and your username. For example, `EXAMPLEDOMAIN-AUTO\Administrator`.

      NOTE: The `-AUTO` suffix is the default value of the **Traffic Controls** > **RDP** > **Settings** > **Autologon domain suffix** option of One Identity Safeguard for Privileged Sessions (SPS). If your SPS administrator has changed this option, use the appropriate suffix instead of -AUTO.
    - Enter your username (only the username, without the domain, for example, **`Administrator`**) into the **Password** field.
4. If the authentication is successful, the desktop of the remote server is displayed.

Configuring password-protected Credential Stores

The following describes how to configure a local Credential Store that stores the credentials used to login to the target host. The Credential Store will be protected by custom passwords. This password must be entered every time One Identity Safeguard for Privileged Sessions (SPS) is rebooted to make the Credential Store available.

Prerequisites

NOTE: Users accessing connections that use Credential Stores to authenticate on the target server must authenticate on SPS using gateway authentication or an AA plugin. Therefore gateway authentication or an AA plugin must be configured for these connections.

For more information, see *Configuring gateway authentication* in the *Administration Guide* and *Integrating external authentication and authorization systems* in the *Administration Guide*.

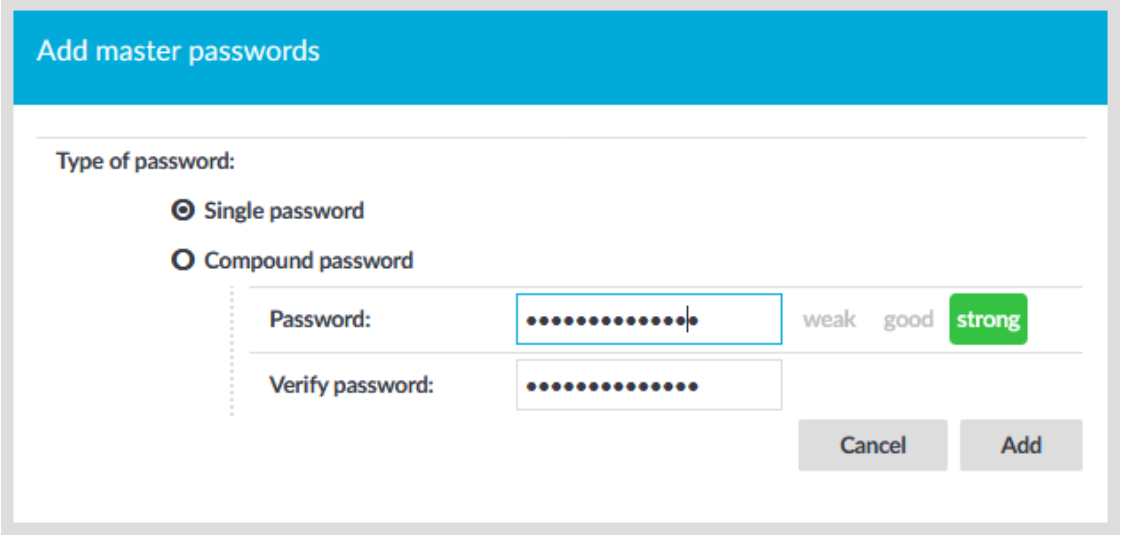
To configure a local Credential Store that stores the credentials used to login to the target host

1. Navigate to **Policies** > **Credential Stores**.
2. Click + and enter a name for the Credential Store.
3. Select **Local**.
4. Select **Encryption key** > **Password protected**.

NOTE: The contents of the Credential Store, as well as the passwords are included in the configuration backups of SPS. Make sure to encrypt the configuration backups.

5. Select **Master passwords** and click +.

Figure 306: Policies > Credential Stores > Local > Password protected — Set Master Password



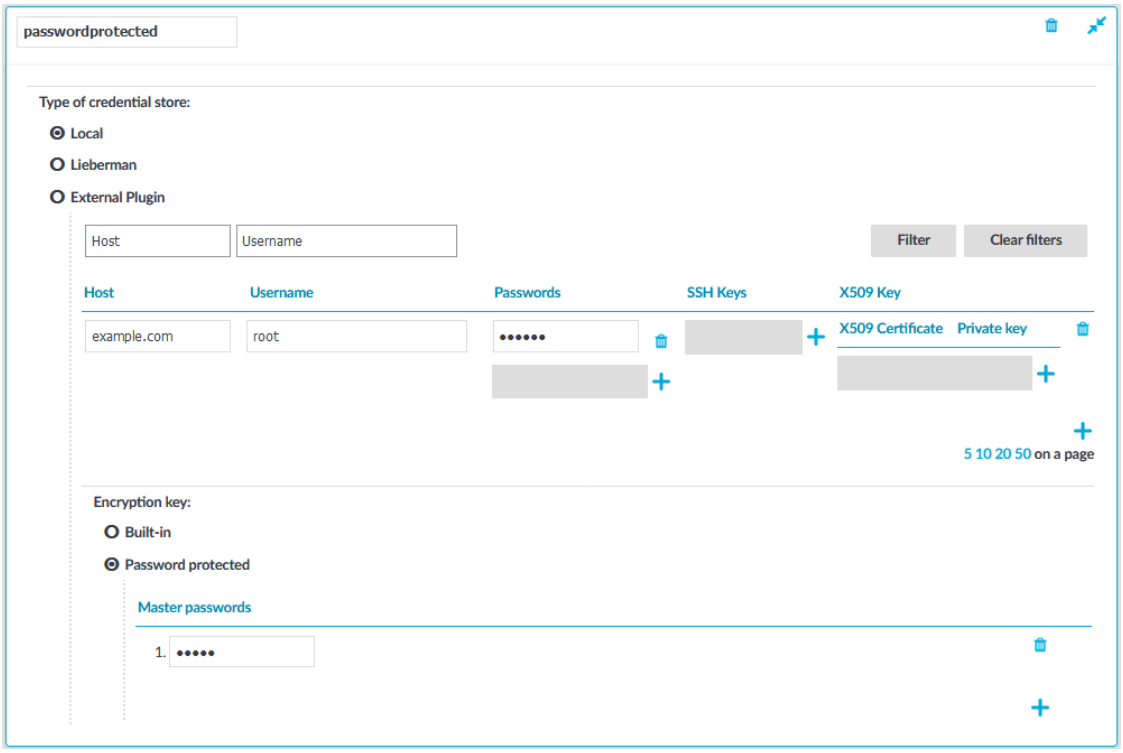
- To protect the Credential Store with a single password, select **Single password** and enter the password into the **Password** and **Verify password** fields. Anyone who knows this password and has the Unlock Credential Store privilege will be able to open the Credential Store. Password-protected Credential Stores must be unlocked on the SPS web interface or console after every SPS reboot.
- To protect the Credential Store with multiple passwords, select **Compound password**, click + and enter a password. Click + to add additional passwords. After finishing listing every password, click **Add**. All of these passwords will be needed to unlock the Credential Store.

SPS encrypts the master passwords using an `aes-256-cbc` cipher, and stores them in a local database.

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

Figure 307: Policies > Credential Stores > Local > Password protected — Configure Credential Store



6. Repeat the previous step to add another single or compound password. That way, different password sets can be defined for the Credential Store. For example, if a single and a compound password is configured, the chief administrator can unlock the Credential Store with a single password, and two of his subordinates can open the Credential Store together if they know one element each of the compound password.

   TIP: To change the password, just click to delete the old password. Then add new passwords as needed.

1. Add credentials to the Credential Store.

   a. Click and enter the destination host and the username. For the destination host, you can use hostname, IP address, or subnet as well. To use the same credentials for every destination host, enter the **`0.0.0.0/0`** subnet. To use the credentials only on the hosts of a specific domain, enter **`*.domain`**. Note that:

      - Usernames are case sensitive.
      - To authenticate users of a Windows domain, enter the name of the domain into the **Host** field.

      Use an IPv4 address.

   b. Set the credentials. SPS will use these credentials to login to the destination host if the credential store is selected in a Connection policy. If more than one

credential is specified to a host-username pair, SPS will attempt to use the credentials as the destination host requests it.

- To add a password, click **Passwords >**  , then enter the password corresponding to the username.
- To upload a private key, click **SSH Keys >**  **>**  , then paste or upload a private key.

  NOTE: If the private key is protected by a passphrase, enter the passphrase. The passphrase is needed only once during the upload, it is not required for the later operation of the Credential Store.

- To generate a keypair on SPS click **SSH Keys >**  **>**  , set the length and type of the key, then click **Generate**. After that, click the fingerprint of the key to download the public part of the keypair. There is no way to download the private key from the SPS web interface.

  NOTE:

  TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

- To upload a certificate and the corresponding private key, click **X509 Keys >**  **>**  , then paste or upload a certificate and the private key.

  NOTE: If the private key is protected by a passphrase, enter the passphrase. The passphrase is needed only once during the upload, it is not required for the later operation of the Credential Store.

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

c. Repeat the previous step to add further credentials to the username as necessary.

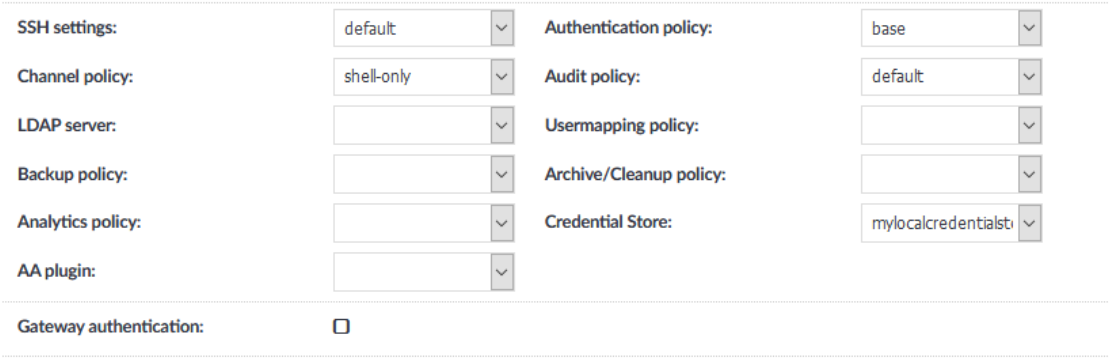
2. Repeat the previous step to add further hosts or usernames as necessary.

   NOTE: Credential Stores can be used together with usermapping policies to simplify the administration of users on the target hosts.

   For more information, see Configuring usermapping policies on page 860.

1. Click **Commit**.
2. Navigate to the Connection policy where you want to use the Credential Store (for example, to **Traffic Controls** > **SSH** > **Connections**), select the Credential Store to use in the **Credential Store** field, then click **Commit**.

Figure 308: Control > Connections — Select a Credential Store to use



3. Navigate to **Basic Settings** > **Alerting & Monitoring** > **Traffic related traps** and enable the **Decryption of a credential failed** (`scbCredStoreDecrpytError`) and **The requested credential store is closed** (`scbCredStoreClosed`) events. That way SPS sends automatic alerts if a Credential Store needs to be unlocked.

> **CAUTION:**
>
> Password-protected Credential Stores must be unlocked every time after SPS is rebooted. Connections using a password-protected Credential Store will automatically fail until the Credential Store is locked.
>
> To unlock a Credential Store, users must have the **User menu** > **Unlock Credential Store** privilege, or editing (read and write) privileges to the particular Credential Store.

Unlocking Credential Stores

To unlock a Credential Store and make it available for use, complete the following steps.

Prerequisites

To unlock a Credential Store, users must have the **User menu** > **Unlock Credential Store** privilege, or editing (read and write) privileges to the particular Credential Store.

1. Login to the One Identity Safeguard for Privileged Sessions (SPS) web interface.
2. Navigate to **User menu** > **Unlock Credential Store** and select the Credential Store to unlock.
3. Enter the password(s) for the Credential Store. For compound passwords, enter every element of the compound password in the correct order.

NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

- Letters A-Z, a-z
- Numbers 0-9
- The space character
- Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

4. Click **Unlock**.
5. Repeat the previous steps for other Credential Stores as needed.

   NOTE: Alternatively, Credential Stores can be unlocked also from the SPS Console Menu.

Using a custom Credential Store plugin to authenticate on the target hosts

The following describes how to configure One Identity Safeguard for Privileged Sessions (SPS) to retrieve the credentials used to login to the target host using a custom plugin.

Prerequisites

- To use a custom Credential Store plugin, you have to upload a working Credential Store plugin to SPS. This plugin is a script that can be used to access an external Credential Store or Password Manager. If you want to create such a custom Credential Store plugin, contact our Support Team or see *Introduction* in the *Creating Custom Credential Store Plugins*.

  For more information on uploading plugins, see Uploading plugins.

  NOTE: Users accessing connections that use Credential Stores to authenticate on the target server must authenticate on SPS using gateway authentication. Therefore, gateway authentication must be configured for these connections. For details, see Configuring gateway authentication.

- Verify the integrity of the plugin.

  For more information on verifying the integrity of plugins, see Verifying the integrity of a plugin.

To configure SPS to retrieve the credentials used to login to the target host using a custom plugin

1. Navigate to **Policies** > **Credential Stores**.
2. Click + and enter a name for the Credential Store.
3. Select **External Plugin**, then select the plugin to use from the **Plugin** list.

4. If your plugin supports configuration, then you can create multiple customized configuration instances of the plugin for your site. The **Configuration** textbox displays the example configuration of the plugin you selected. If you wish to create a customized configuration instance of the plugin for your site, then edit the configuration here.

   NOTE: Plugins created and issued before the release of SPS 5.1 do not support configuration. If you create a configuration for a plugin that does not support this, the affected connection will stop with an error message.

5. Click 

.
6. Navigate to the Connection policy where you want to use the Credential Store (for example, to **Traffic Controls** > **SSH** > **Connections**), select the Credential Store configuration instance to use in the **Credential Store** field, then click 

.

Integrating external authentication and authorization systems

One Identity Safeguard for Privileged Sessions (SPS) provides a plugin framework to integrate SPS to external systems to authenticate or authorize the user before authenticating on the target server. Such plugins can also be used to request additional information from the users, for example, to perform multi-factor authentication.

You can use an Authentication and Authorization plugin (AA plugin) in the following protocols:

- MSSQL
- Remote Desktop (RDP)
- Secure Shell (SSH)
- TELNET

To request a plugin that interoperates with your authentication or authorization system, contact our Support Team.

For details on configuring SPS to use a plugin, see Using a custom Authentication and Authorization plugin to authenticate on the target hosts.

How Authentication and Authorization plugins work

If a Connection Policy has an Authentication and Authorization plugin (**AA plugin**) configured, One Identity Safeguard for Privileged Sessions (SPS) runs the plugin as the last

step of the connection authorization phase. SPS can request the client to perform other types of authentication before running the plugin. Using an **AA plugin** in a Connection Policy is treated as gateway authentication if:

- the plugin authenticates the user
- authentication is successful
- the plugin returns the `gateway_user` and `gateway_groups` elements, identifying the user it has authenticated

Other types of gateway authentication will come before authentication by the **AA plugin**, so information from any other type of gateway authentication (for example, the username and usergroups of this authentication) will already be available and therefore can be used by the plugin. If the Authentication and Authorization plugin does perform gateway authentication, you can use a Credential Store as well.

However, for technical reasons, the web-based gateway authentication (that is, authenticating on the SPS web interface if the **Require Gatweay Authentication on the SPS Web Interface** option is selected in the Connection Policy) is performed after the **AA plugin**, so using **AA plugin** and ticking **Require Gateway Authentication on the SPS Web Interface** at the same time is not a valid configuration.

The plugin can interactively request additional information from the client in the SSH, Telnet, and RDP protocols.

NOTE: In SPS 5.8, a user's group membership is determined by querying only the relevant groups configured for the connection from the LDAP/AD server, instead of retrieving all groups of a given user.

This may cause problems when using AD/LDAP-based gateway authentication together with an AA plugin. The AA plugin `authorize()` hook may be called with only a subset of groups as group membership lookup does not consider groups referenced in the AA plugin code.

Consider that only groups queried by SPS are affected. Gateway groups returned by the AA plugin `authenticate()` hook are passed to the `authorize()` hook unchanged.

SPS runs the `authorize` method after the `authentication` method and any inband gateway authentication or inband destination selection steps. As a result, the `authorize` method already has access to the IP address of the target server and the remote username (the username used in the server-side connection).

Optionally, the plugin can return the `gateway_user` and `gateway_groups` values. SPS will only update the **gateway username** and **gateway groups** fields in the connection database if the plugin returns the `gateway_user` and `gateway_groups` values. The returned `gateway_user` and `gateway_groups` values override any such attributes already available on SPS about the connection (that means that channel policy evaluations will be affected), so make sure that the plugin uses the original values appropriately.

If the plugin returns the `gateway_user` and `gateway_groups` values, you may have to configure an appropriate **Usermapping policy** in the **Connection Policy**. If the plugin returns a `gateway_user` that is different from the remote user, the connection will fail without a usermapping policy. For details on usermapping policies, see *Configuring usermapping policies* in the *Administration Guide*.

Prerequisites

- SPS supports Authentication and Authorization plugins in the RDP, SSH, and TELNET protocols.
- In RDP, using an AA plugin together with Network Level Authentication in a Connection Policy has the same limitations as using Network Level Authentication without domain membership.
- In RDP, using an AA plugin requires TLS-encrypted RDP connections. For details, see *Enabling TLS-encryption for RDP connections* in the *Administration Guide*.

Using a custom Authentication and Authorization plugin to authenticate on the target hosts

The following describes how to configure One Identity Safeguard for Privileged Sessions (SPS) to use an Authentication and Authorization plugin (AA plugin) before accessing the target host.

Prerequisites

- To use a custom plugin, you need to upload a working **AA plugin** to SPS. This plugin is a script that uses the SPS API to access an external system. If you want to create such a plugin, contact our Support Team for details and instructions or see *Creating Custom Authentication and Authorization Plugins*.

  For more information on uploading plugins, see Uploading plugins.
- Verify the integrity of the plugin.

  For more information on verifying the integrity of plugins, see Verifying the integrity of a plugin.
- SPS supports AA plugins in the MSSQL, RDP, SSH, and Telnet protocols.
- In RDP, using an **AA plugin** together with Network Level Authentication in a Connection Policy has the same limitations as using Network Level Authentication without domain membership.
- In RDP, using an **AA plugin** requires TLS-encrypted RDP connections. For details, see *Enabling TLS-encryption for RDP connections* in the *Administration Guide*.

To configure SPS to use an Authentication and Authorization plugin before accessing the target host

1. If your plugin supports configuration, then you can create multiple customized configuration instances of the plugin for your site. Create an instance by completing the following steps:

a. Navigate to **Policies** > **AA Plugin Configurations**. Select the plugin to use from the **Plugin** list.

b. The **Configuration** textbox displays the example configuration of the plugin you selected. You can edit the configuration here if you wish to create a customized instance of the plugin.

NOTE: Plugins created and issued before the release of SPS 5.1 do not support configuration. If you create a configuration for a plugin that does not support this, the affected connection will stop with an error message.

Figure 309: Policies > AA Plugin Configurations — Creating a customized plugin configuration instance



2. Navigate to the Connection policy where you want to use the plugin (for example, to **Traffic Controls** > **RDP** > **Connections**), select the plugin configuration instance to use in the **AA plugin** field, then click Commit.
3. If the plugin sets or overrides the gateway username of the connection, configure a **Usermapping policy** and use it in the Connection policy. For details, see *Configuring usermapping policies* in the *Administration Guide*.
4. Verify that the configuration works properly: try to establish a test connection. For details, see *Performing authentication with AA plugin in Remote Desktop connections* in the *Administration Guide*. If the plugin is configured to store any metadata about the connection, these data will be available in the **Additional metadata** field of the SPS Search interface.

Performing authentication with AA plugin in terminal connections

The following describes how to establish a terminal connection (SSH, TELNET, or TN3270) to a server.

To establish a terminal connection (SSH, TELNET, or TN3270) to a server

1. Connect to the server.

   To encode additional data as part of the username, you can use the `@` as a field separator, for example:

   ```
   ssh token_id=id@user@server
   ```

   Replace `id` with your actual token ID.
2. If One Identity Safeguard for Privileged Sessions (SPS) prompts you for further information (for example, a one-time password), enter the requested information.
3. Authenticate on the server.
4. If authentication is successful, you can access the server.

Performing authentication with AA plugin in Remote Desktop connections

The following describes how to establish a Remote Desktop (RDP) connection to a server when the **AA plugin** is configured.

To establish an RDP connection to a server when the AA plugin is configured

1. Open your Remote Desktop client application.
2. If you have to provide additional information to authenticate on the server, you must enter this information in your Remote Desktop client application into the **User name** field, before the regular content (for example, your username) of the field.

   To encode additional data, you can use the following special characters:

   - `%` as a field separator
   - `~` as the equal sign
   - `^` as a colon (for example, to specify the port number or an IPv6 IP address)

   For example, to add a token ID before your username, use the following format:

   ```
   domain\token_id~12345%Administrator
   ```

   Note how domain information is provided. If your server is in a domain, make sure that you specify the domain in this format: putting it in front, followed by a backslash (`\`).
3. Connect to the server.
4. If One Identity Safeguard for Privileged Sessions (SPS) prompts you for further information (for example, a one-time password), enter the requested information.

5. Authenticate on the server.
6. If authentication is successful, you can access the server.

Integrating ticketing systems

The plugin framework provided by One Identity Safeguard for Privileged Sessions (SPS) can also be used to integrate SPS to external ticketing (or issue tracking) systems, allowing you to request a ticket ID from the user before authenticating on the target server. That way, SPS can verify that the user has a valid reason to access the server — and optionally terminate the connection if he does not. Requesting a ticket ID currently supports the following protocols:

- Remote Desktop (RDP)
- Secure Shell (SSH)
- TELNET
- TN3270
- To request a plugin that interoperates with your ticketing system, contact our Support Team.
- For details on configuring SPS to use a plugin, see *Using a custom Authentication and Authorization plugin to authenticate on the target hosts* in the *Administration Guide*.

Performing authentication with ticketing integration in terminal connections

The following describes how to establish a terminal connection (SSH, TELNET, or TN3270) to a server that requires you to enter a ticket ID.

To establish a terminal connection (SSH, TELNET, or TN3270) to a server that requires you to enter a ticket ID

1. Connect to the server.

   You have the option to use the ID of the ticket you are working on as part of the username (replace `id` with the ticket ID):

   ```
   ssh ticket_id=id@user@server
   ```

   NOTE: Your plugin may use a different name for the key `ticket_id` shown in the example. Plugins work with key-value pairs and the names of keys are entirely up to individual plugins.
2. If you did not provide a ticket ID, One Identity Safeguard for Privileged Sessions (SPS) now prompts you to enter it.

3. Authenticate on the server.
4. If the authentication is successful, you can access the server.

Performing authentication with ticketing integration in Remote Desktop connections

The following describes how to establish a Remote Desktop (RDP) connection to a server that requires you to enter a ticket ID.

To establish an RDP connection to a server that requires you to enter a ticket ID

1. Open your Remote Desktop client application.
2. Enter the ticket ID into your Remote Desktop client application into the *User name* field, before or after the regular content (for example, your username) of the field. You must provide the ticket ID in the following format:

   ```
   ticket_id~<your-ticket-id>%
   ```

   Replace `<your-ticket-id>` with your actual ticket number. For example:

   ```
   ticket_id~12345%Administrator
   ```

   NOTE: Your plugin may use a different name for the key `ticket_id` shown in the example. Plugins work with key-value pairs and the names of keys are entirely up to individual plugins.

   To encode additional data, you can use the following special characters:

   - `%` as a field separator
   - `~` as the equal sign
   - `^` as a colon (for example, to specify the port number or an IPv6 IP address)

   For example, to add a token ID before your username, use the following format:

   ```
   domain\token_id~12345%Administrator
   ```

   Note how domain information is provided. If your server is in a domain, make sure that you specify the domain in this format: putting it in front, followed by a backslash (`\`).
3. Connect to the server.
4. Authenticate on the server.
5. If the authentication is successful, you can access the server.

Creating a custom plugin

Creating a custom Authentication and Authorization plugin

For more information, see *Creating Custom Authentication and Authorization Plugins*.

Creating a custom Credential Store plugin

For more information, see *Creating Custom Credential Store Plugins*.

Plugin troubleshooting

On the default log level, One Identity Safeguard for Privileged Sessions (SPS) logs everything that the plugin writes to `stdout` and `stderr`. Log message lines are prefixed with the session ID of the proxy, which makes it easier to find correlating messages.

To transfer information between the methods of a plugin (for example, to include data in a log message when the session is closed), you can use a cookie.

If an error occurs while executing the plugin, SPS automatically terminates the session.

NOTE: This error is not visible in the verdict of the session. To find out why the session was terminated, you have to check the logs.

Reports

One Identity Safeguard for Privileged Sessions (SPS) periodically creates reports on the activity of the administrators, its system information, as well as the processed traffic. In addition, you can use the connection database for creating custom reports from connection statistics.

You can specify the following access rights on the **Reporting** > **Configuration** page:

- **Reporting** > **Configuration** allows you to configure new reports.
- **Reporting** > **Content subchapters** allows you to access and create content-based report subchapters.
- **Reporting** > **Search subchapters** allows you to access and create search-based report subchapters.
- **Reporting** > **Reports** allows you to create the available reports.

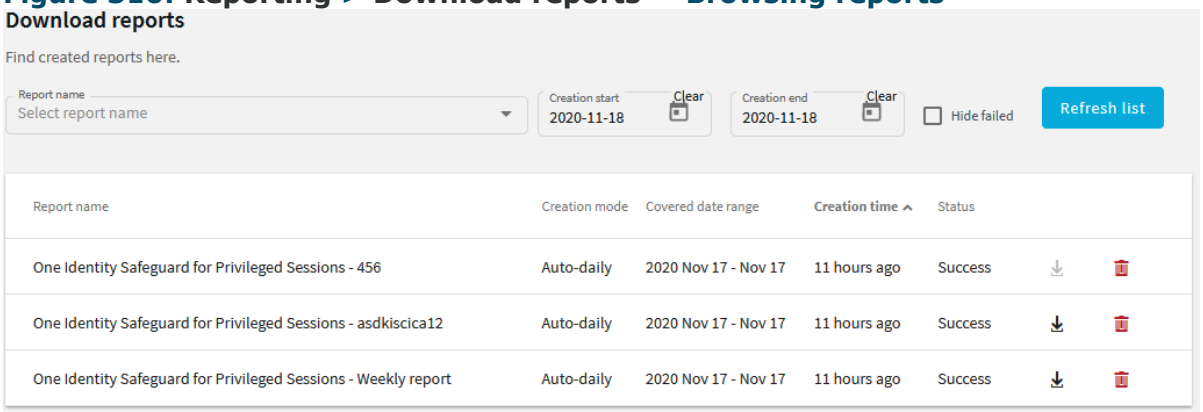
For more information on configuring user rights, see Managing user rights and user groups on page 361.

To access reports, select **Reporting** > **Download reports** from the Main Menu. The reports are displayed on a search interface. For more information on using and customizing this interface, see Using the internal search interface on page 390.

The reports are also sent to the email address set at **Basic Settings** > **Management** > **Mail settings** > **Send reports to**, unless specified otherwise in the configuration of the report.

NOTE: If the **Basic Settings** > **Management** > **Mail settings** > **Send reports to** address is not set, the system report is sent to the SPS administrator's email address.

Figure 310: Reporting > Download reports — Browsing reports



Reports can be generated for fixed periods:

- **Daily reports** are generated every day at 00:01.
- **Weekly reports** are generated every week on Monday at 00:01.
- **Monthly reports** are generated on the first day of every month at 00:01.

To access the reports from the SPS web interface, the user must have the appropriate privileges (for custom reports, the default requirement is membership in the `search` group). In addition, individual reports might have different access requirements configured. For more information on configuring user rights, see Managing user rights and user groups on page 361.

Contents of the operational reports

The operational reports of One Identity Safeguard for Privileged Sessions (SPS) are available in Adobe Portable Document Format (PDF), and contain the following information:

- **Configuration changes**: Lists the number of SPS configuration changes per page and per user. The frequency of the configuration changes is also displayed on a chart.
- **Main reports**: Contains statistics about the total traffic that passed SPS, including the number of sessions that passed for every connection policy, the used usernames, clients, and servers, and so on.

  NOTE: Connections that are still in progress when the report is generated are excluded from the report. Sessions that are being indexed and reporting jobs are listed in the *Sessions with in progress indexing or reporting jobs* section of the report.
- **Reports by connection**: Contains separate statistics about every connection policy configured on SPS.
- **System health information**: Displays information about the filesystem and network use of SPS, as well as the average load.

Configuring custom reports

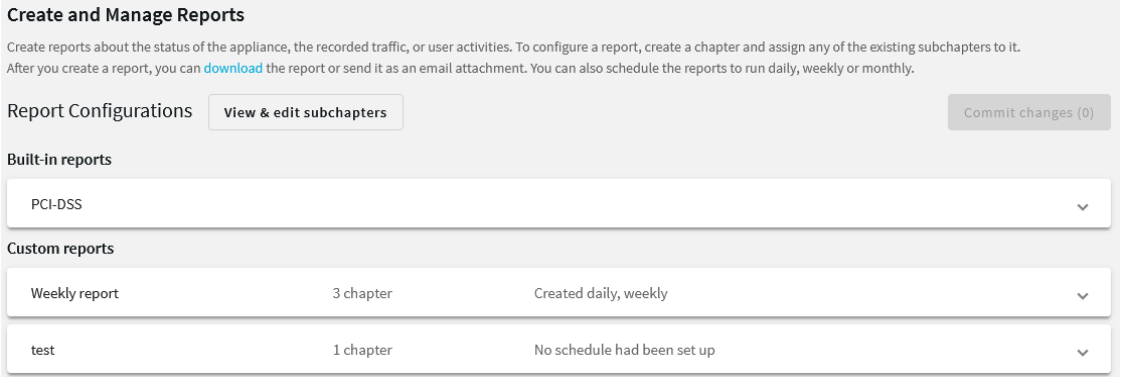
To configure a report, create a chapter and assign any of the existing subchapters to it. The following sources (statistics or other queries) are available as reporting subchapters:

- The indexed contents of audit trails, as described in Indexing audit trails on page 672.
- The statistics of an audit trail search, as described in Displaying statistics on search results on page 808.

To configure SPS to create custom reports

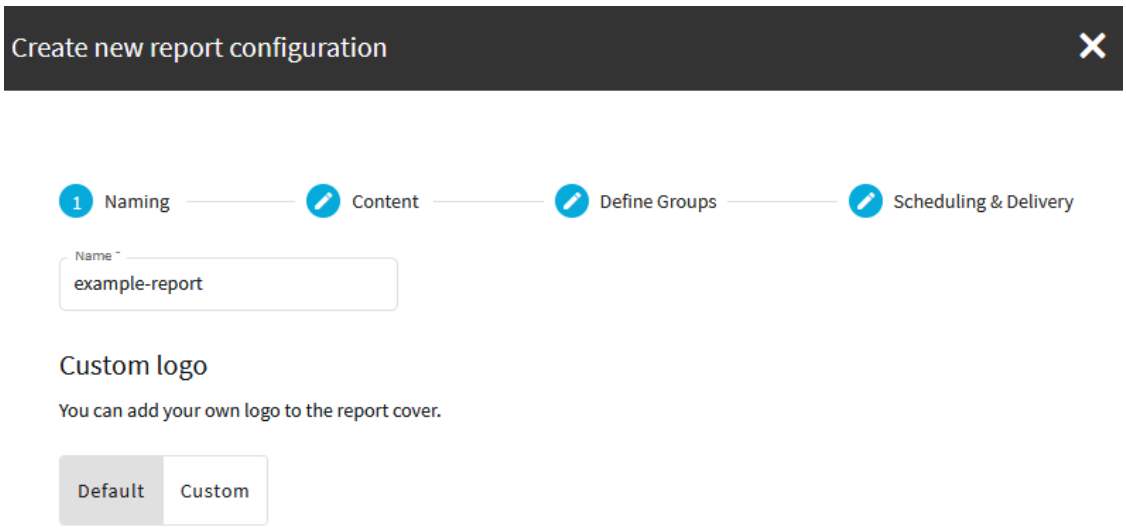
1. Login to the SPS web interface, and navigate to **Reporting** > **Create & Manage Reports**.

Figure 311: Reporting > Create & Manage Reports — Configuring custom reports



2. Click **Create new report configuration** and enter a name for the custom report.

Figure 312: Reporting > Create & Manage Reports — Configuring custom reports

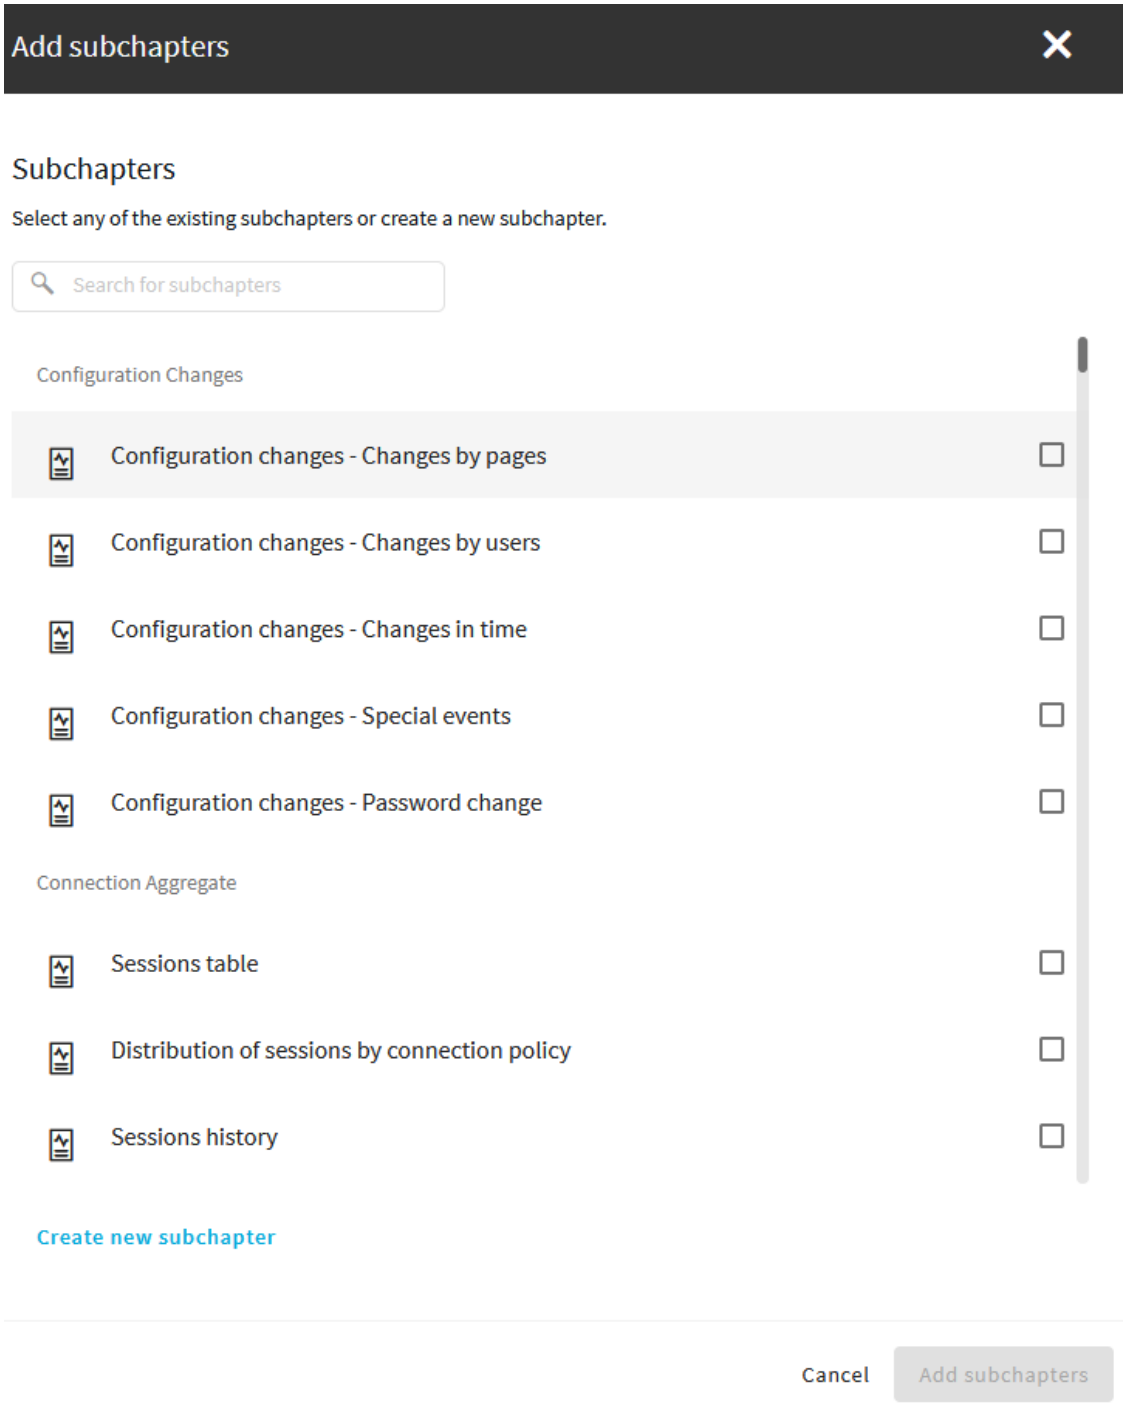


If you want to have the logo of your organization on the cover page of the report (instead of the One Identity logo), select **Custom**, select your logo file, then click **Upload**. You can upload GIF, JPEG, and PNG images. SPS will automatically resize the image to fit on the cover page.

3. Select **Create new chapter**, enter a name for the chapter. Repeat this step to create further chapters if needed.

4. Select **Add new subchapter** to add various reports and statistics to the chapter. Select any of the existing subchapters or create a new subchapter.

Figure 313: Reporting > Configuration > Add new subchapter — Adding subchapters to reports



NOTE: When creating a subchapter that searches for keywords in HTTP, only the timestamp of the results will be visible in the report, without data.

5. Once selected, use drag and drop to change the order of the subchapters if needed.
6. By default, members of the `report` group can access the custom reports through the SPS web interface. To change this, enter the name of a different group into the **Groups** field.

   NOTE: Members of the listed groups can access only these custom reports even if their groups do not have read access to the **Reporting** > **Download reports** page. However, only those reports will be listed, to which their group has access to.
7. Select how often SPS creates the report from the **Scheduling** field. Weekly reports are created on Mondays, while monthly reports on the first day of the month. If you want to create the report only manually, leave these fields empty.
8. By default, SPS sends out the reports in email to the address set in the **Basic Settings** > **Management** > **Mail settings** > **Send reports to** field.

   NOTE: If this address is not set, the report is sent to the SPS administrator's email address.

   - To disable email sending, clear the **Deliver in email** option.
   - To email the reports to a different address, select **Custom**, and enter the email address where the reports should be sent. Click **Add email** to list multiple email addresses if needed.
9. Click **Create report**.

Creating report subchapters

Creating reports from audit trail content

One Identity Safeguard for Privileged Sessions (SPS) can index the contents of audit trails using its own indexer service or external indexers. Indexing extracts the text from the audit trails and segments it to tokens. A token is a segment of the text that does not contain whitespace: for example words, dates (`2009-03-14`), MAC or IP addresses, and so on. The indexer returns the extracted tokens to SPS, which builds a comprehensive index from the tokens of the processed audit trails.

Once indexed, the contents of the audit trails can be searched from the web interface. SPS can extract the commands typed and the texts seen by the user in terminal sessions, and text from graphical protocols like RDP, Citrix ICA, and VNC. Window titles are also detected.

SPS has an internal indexer, which runs on the SPS appliance. In addition to the internal indexer, external indexers can run on Linux hosts.

Processing and indexing audit trails requires significant computing resources. If you have to audit lots of connections, or have a large number of custom reports configured, consider using an external indexer to decrease the load on SPS. For sizing recommendations, ask your One Identity partner or contact our Support Team.

SPS also creates statistics of the occurrences of the search keywords, as well as screenshots from the audit trail. These statistics and screenshots can be included in custom reports as subchapters.

NOTE: The screenshots of the sessions containing the search keywords are not included in the report output for security reasons, but you can access the screenshots by using the clickable QR codes.

For more information on accessing the screenshots, see section Report output.

NOTE: Consider the following:

- The screenshot generated from the search results contains the first occurrence of the search keywords. If your search keywords are visible in the audit trail for a longer period, it is possible that the first occurrence is not the most relevant.
- For technical reasons, trail data in terminal connections (SSH and Telnet) is aggregated for each second. The screenshot generated for the report reflects the terminal buffer, as it was visible at the end of that second. If data that contains the search keyword was pushed off-screen during this second, the search still finds it, but it will not be visible on the generated screenshot. Similarly, if you search for multiple keywords, it is possible that you will receive results that do not contain every keyword on the same screen (but they were separately visible within the one-second interval).

NOTE: Only audit trails created after the content subchapter has been configured will be processed. It is not possible to create reports from already existing audit trails.

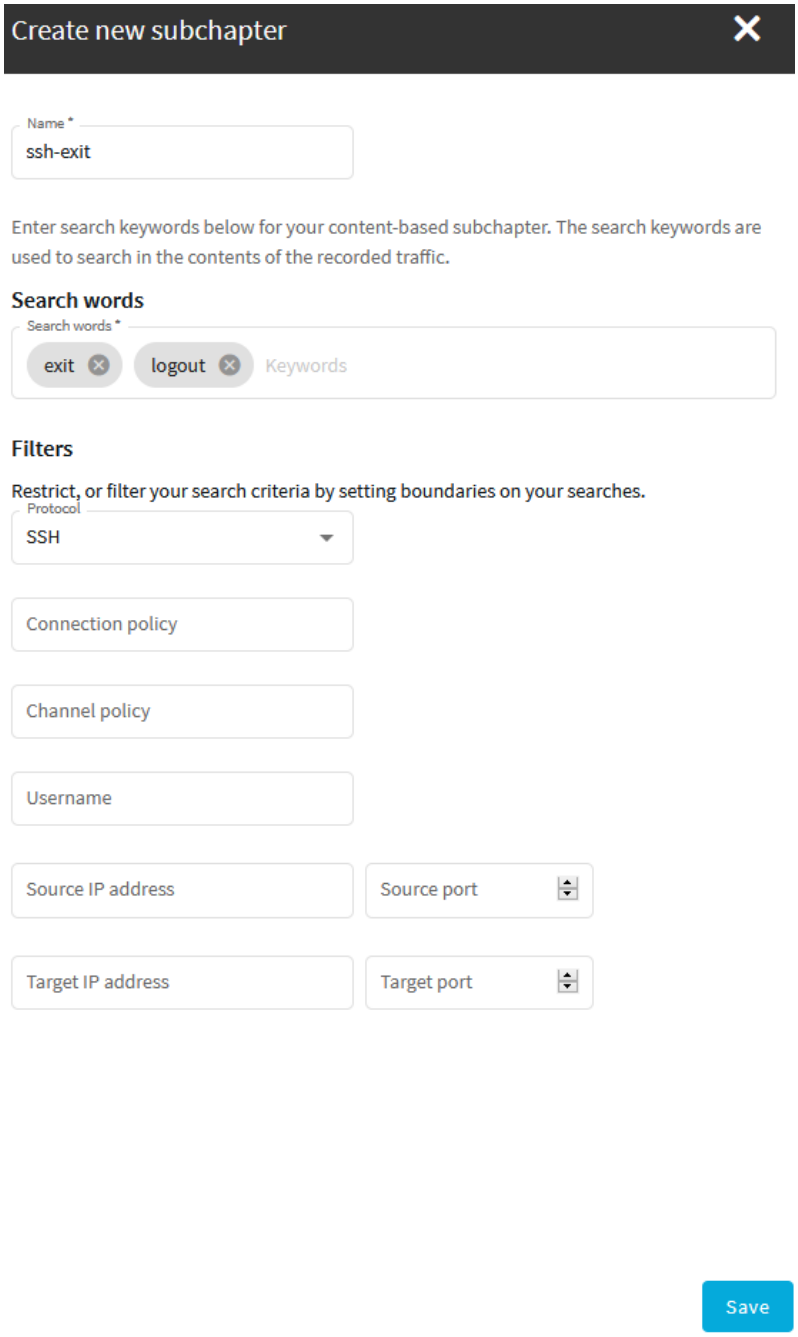
Prerequisites for the indexer service

If you are indexing the audit trails with the indexer service, make sure to enable indexing for the connections you want to include in the report. Otherwise, reporting on audit trail content will not work. For more information, see Configuring the internal indexer.

To configure SPS to create reports from the contents of audit trails

1. Login to the SPS web interface, and navigate to **Reporting** > **View & edit subchapters** > **Content-based**.

Figure 314: Reporting > View & edit subchapters > Content-based — Configuring audit trail reports



2. Click **Create new** and enter a name for the subchapter.
3. Enter the search keywords (or parts of the words) into the **Search words** field. The search keywords are used to search in the contents of the recorded traffic.

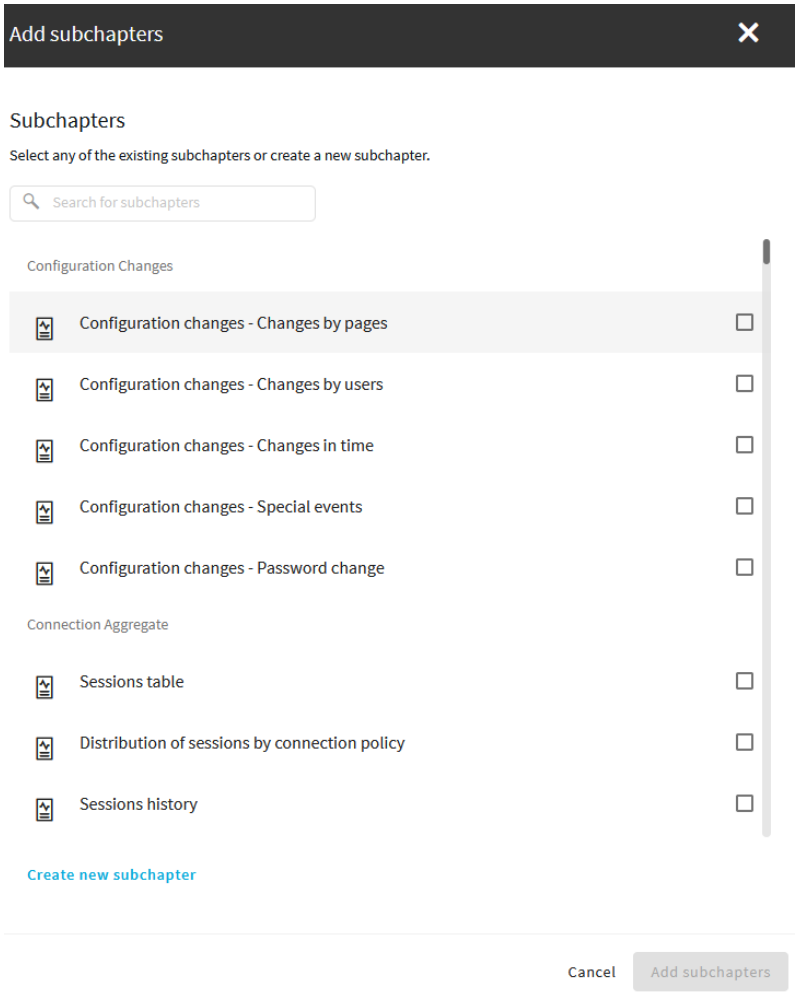
   Note the following points.

- Your search expression must be shorter than 255 characters.
- The search is not case sensitive.
- Wildcards and regular expressions are not supported.
- Searching for an exact phrase or expression by enclosing multiple keywords in double quotes (for example, **`"program files"`**) is not supported.

4. Configure filters to select the audit trails to index. The following filters are available:
   - **Protocol**: Process only audit trails of the specified traffic type (for example SSH).
   - **Connection policy**: Process only audit trails of the specified connection policy.
   - **Channel policy**: Process only audit trails of the specified channel policy.
   - **Username**: Process only audit trails where the specified username was used in the connection. Available only for protocols where the username is known (for example SSH).
   - **Source IP address**: Process only audit trails where the specified client IP address or port was used.
   - **Target IP address**: Process only audit trails where the specified server IP address or port was used.

   NOTE: If you do not configure any filters, every available audit trail will be processed. Audit trails are created only for channels where the **Record audit trail** option is enabled for the particular channel in the channel policy.
5. Click **Save**.
6. Navigate to **Reporting** > **Create & Manage Reports**, and add the new subchapter to an existing report, or create a new report. For more information, see Configuring custom reports.

Figure 315: Reporting > Configuration > Add Subchapter — Adding subchapters to reports



Creating search-based report subchapters from search results

You can turn any search query or statistics into a subchapter to add to your reports. This is an easy and flexible way of creating reports to monitor traffic, track certain parameters, or get alerted about particular events.

To create a search-based report subchapter from search results

1. Navigate to **Audit** > **Sessions** and define a valid search query.
2. Click **Create report**. The **Create new subchapter** page is displayed, with the **Search query** field populated with your query.
3. In the **Subchapter title** field, add a title to your subchapter.
4. In **Subchapter type**, select the type that fits your query:
   - **Sessions list view**: Displays a list of sessions.

     Set **Number of sessions** and from **Subchapter columns**, select the session parameters to be displayed in a table in the report. You can add a maximum of 10 columns to the table.
   - **Timeline view**: Visualizes the timeline of sessions that meet the criteria of the **Search query** field, depending on the time range (day/month/week) selected at the **Scheduling & Delivery** step of the report configuration.
   - **Statistics view**: Visualizes the statistics data for the option you select in **Field**, for sessions that meet the criteria of the **Search query** field.

     Select a presentation option for your report, such as **List**, **Pie chart**, or **Bar chart**. In **Field**, select the data field to create your statistics on.
5. Select **Add to a report**, and select from the list of available reports.

   Alternatively, to configure a custom report from scratch and include this subchapter in it, select **Include in a new report**. For more information, see Configuring custom reports on page 893.

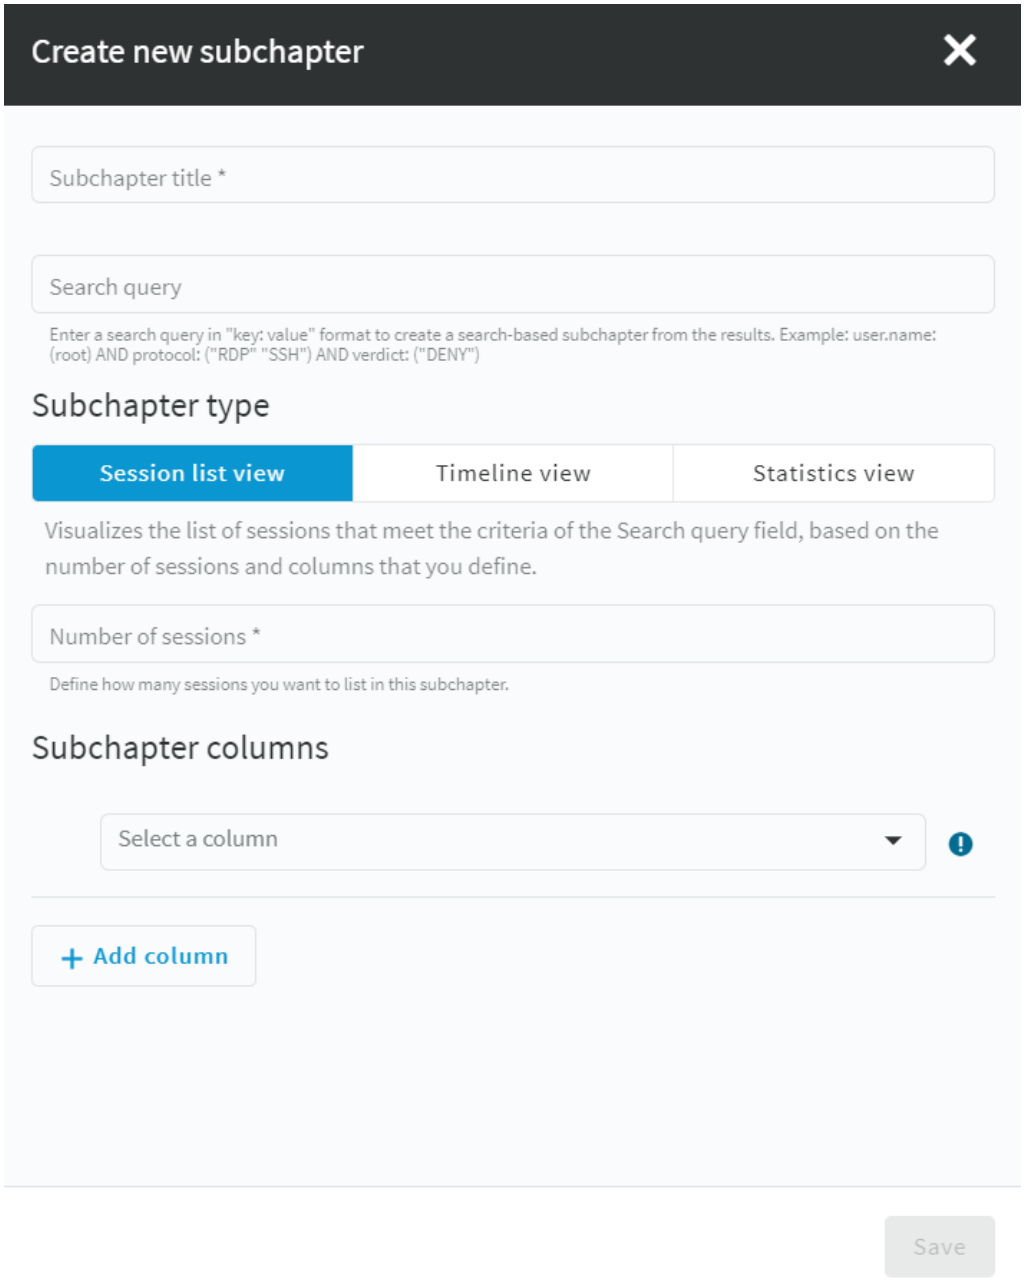
Creating search-based report subchapters from scratch

This section describes how to create a search-based subchapter from scratch to include the subchapter in a custom report.

To create a search-based report subchapter from scratch

1. If you have multiple SPS appliances organized into a cluster where one of the nodes is the Search Master (or Central Search) node, log in to that node.
2. Navigate to **Reporting**> **Create and Manage Reports** > **View & edit subchapters** > **Search-based**.
3. Select **Create new**. The **Create new subchapter** page is displayed.

Figure 316: Reporting> View & edit subchapters > Search-based — Create new subchapter



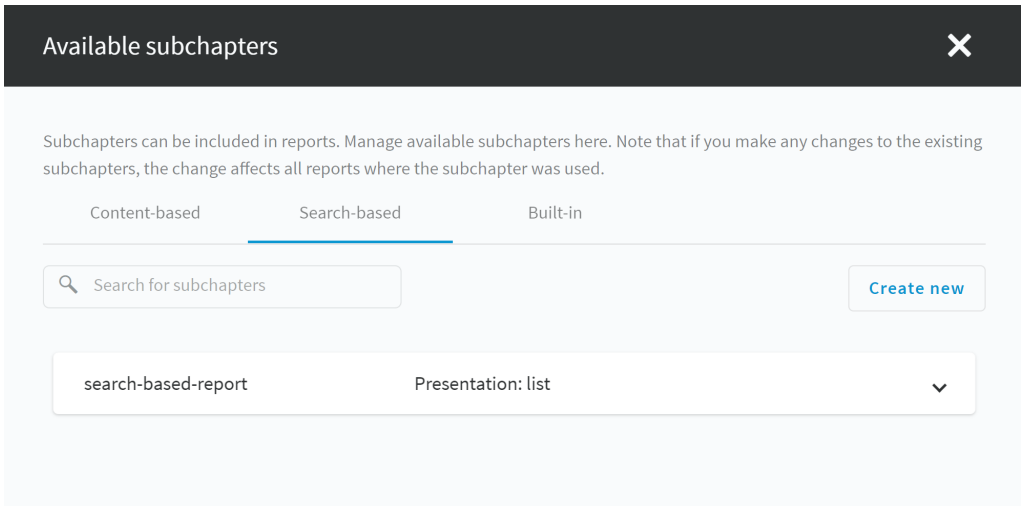
4. In **Subsection title**, add a title to your subchapter.
5. In **Search query**, enter a valid query.

6. In **Subchapter type**, select the type that fits your query:
   - **Sessions list view**: Displays a list of sessions.

     Set **Number of sessions** and from **Subchapter columns**, select the session parameters to be displayed in a table in the report. You can add a maximum of 10 columns to the table.
   - **Timeline view**: Visualizes the timeline of sessions that meet the criteria of the **Search query** field, depending on the time range (day/month/week) selected at the **Scheduling & Delivery** step of the report configuration.
   - **Statistics view**: Visualizes the statistics data for the option you select in **Field**, for sessions that meet the criteria of the **Search query** field.

     Select a presentation option for your report, such as **List**, **Pie chart**, or **Bar chart**. In **Field**, select the data field to create your statistics on.
7. Select **Save**.
8. To save your changes, navigate to **Create and Manage reports** and select **Commit**.
9. Add your subchapter to a new report or to an existing report. For more information, see Configuring custom reports.

   To find and add the subchapter you created to a report, navigate to **Reporting**> **Create and Manage Reports** > **View & edit subchapters** > **Search-based**.



Creating PCI DSS reports

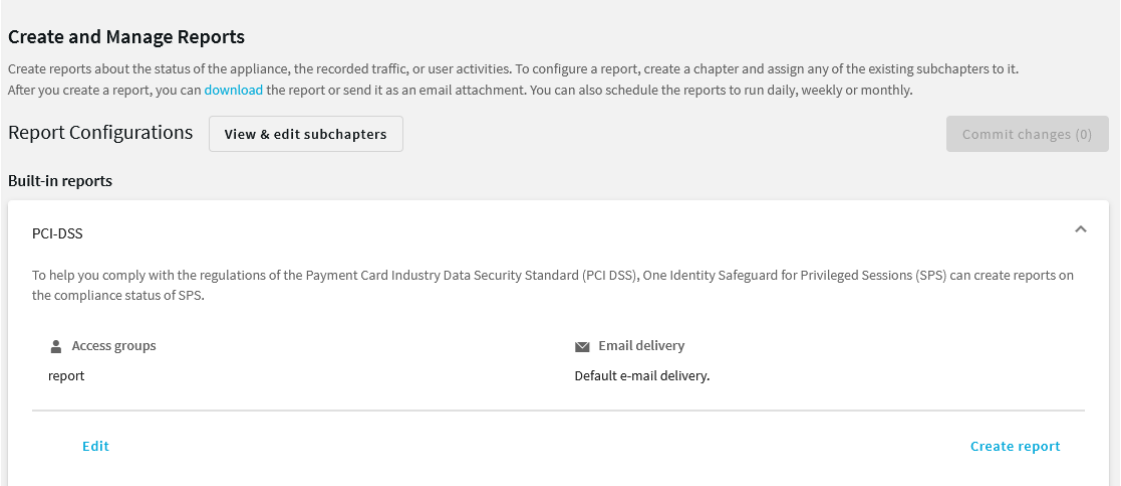
To help you comply with the regulations of the Payment Card Industry Data Security Standard (PCI DSS), One Identity Safeguard for Privileged Sessions (SPS) can generate reports on the compliance status of SPS. Note that this is not a fully-featured compliance report: it is a tool to enhance and complement your compliance report by providing information available in SPS. The report corresponds with the document *Payment Card Industry (PCI) Data Security Standard, Requirements and Security Assessment Procedures, Version 3.0*, published by the PCI Security Standards Council.

For details on the contents of the report, see Contents of PCI DSS reports on page 904.

To create PCI DSS reports

1. Log in to the SPS web interface, and navigate to **Reporting** > **Create & Manage Reports** > **Built-in reports** > **PCI DSS** > **Create report**.

   Figure 317: Reporting > Create & Manage Reports > Built-in reports > PCI DSS > Create report — Generating PCI DSS reports

2. By default, members of the `report` group can access the custom reports through the SPS web interface. To change this, click **Edit** and enter the name of a different group into the **Groups** field.

   NOTE: Members of the listed groups can access only these custom reports even if their groups do not have read access to the **Reporting** > **Download reports** page. However, only those reports will be listed, to which their group has access to.

3. By default, SPS sends out the reports in email to the address set in the **Basic Settings** > **Management** > **Mail settings** > **Send reports to** field.

   NOTE: If this address is not set, the report is sent to the SPS administrator's email address.

4. To disable email sending, clear the **Deliver in email** option.
5. To email the reports to a different address, select **Custom**, and enter the email address where the reports should be sent. Click **Add email** to list multiple email addresses if needed.
6. Click **Update report**.
7. Click **Create report**.

   The report will be automatically added in the list of reports (**Reporting** > **Download reports**), and also sent in an e-mail to the regular recipients of the report.

Contents of PCI DSS reports

To help you comply with the regulations of the Payment Card Industry Data Security Standard (PCI DSS), One Identity Safeguard for Privileged Sessions (SPS) can generate reports on the compliance status of SPS. Note that this is not a fully-featured compliance report: it is a tool to enhance and complement your compliance report by providing information available in SPS. The report corresponds with the document *Payment Card Industry (PCI) Data Security Standard, Requirements and Security Assessment Procedures, Version 3.0*, published by the PCI Security Standards Council.

For details on creating PCI DSS reports, see Creating PCI DSS reports on page 903. The following table details the information included in the SPS PCI DSS reports, and the relevant PCI compliance requirement.

Table 12: Contents of PCI DSS reports

PCI DSS Requirement	Compliance details
Requirement 1.1.6 Documentation and business justification for use of all services, protocols, and ports allowed, including documentation of security features implemented for those protocols considered to be insecure. Examples of insecure services, protocols, or ports include but are not limited to FTP, Telnet, POP3, IMAP, and SNMP v1 and v2.	The report lists the insecure connection policies configured in SPS, including SNMP server and agent settings, and the list of connection policies that permit unencrypted HTTP and Telnet. This list does not include any insecure connections that can be used to access SPS itself.
Requirement 2.1 Always change vendor-supplied defaults and remove or disable unnecessary default accounts before installing a system on the network. This applies to ALL default passwords, including but not limited to those used by operating systems, software that provides security services, application and system accounts, point-of-sale (POS) terminals, Simple Network Management Protocol (SNMP)community strings, etc.).	SPS can be accessed as "root" via the local management console, or - if explicitly enabled - remotely using a Secure Shell (SSH v2) connection. The report lists the local web user accounts that can access SPS. For details on configuring these accounts, see Managing One Identity Safeguard for Privileged Sessions (SPS) users locally on page 328.
Requirement 2.2.2 Enable only necessary services, protocols, daemons, etc., as required for the function of the system.	The report includes the list of services running on SPS.
Requirement 2.2.3 Implement additional security features for any required services, protocols, or daemons that are considered to be insecure. For example, use secured technologies such as	The report lists the connection policies enabling unencrypted HTTP and Telnet access, and any such session that was active when the report was generated.

PCI DSS Requirement	Compliance details
SSH, S-FTP, SSL, or IPSec VPN to protect insecure services such as NetBIOS, file-sharing, Telnet, FTP, etc.	
Requirement 2.3 Encrypt all non-console administrative access using strong cryptography. Use technologies such as SSH, VPN, or SSL/TLS for web-based management and other non-console administrative access.	Use HTTPS to connect to SPS. HTTP connections are forbidden.
Requirement 3.5.3 Store secret and private keys used to encrypt/decrypt cardholder data in one (or more) of the following forms at all times: • Encrypted with a key-encrypting key that is at least as strong as the dataencrypting key, and that is stored separately from the data-encrypting key • Within a secure cryptographic device (such as a host security module (HSM) or PTS-approved point-of-interaction device) • As at least two full-length key components or key shares, in accordance with an industry-accepted method Note: It is not required that public keys be stored in one of these forms.	Audit trails are encrypted with AES128-GCM (audit trails recorded with SPS 5 F3 and earlier are encrypted with AES128-CBC). The master key is encrypted with the key you provided.
Requirement 5.1.2 For systems considered to be not commonly affected by malicious software, perform periodic evaluations to identify and evaluate evolving malware threats in order to confirm whether such systems continue to not require anti-virus software.	SPS is an appliance running minimal services and using a hardened operating system. One Identity, the vendor of SPS, continuously monitors vulnerabilities and CVEs that might affect the components of SPS, and publishes security updates and announcements as needed. Using an up-to-date SPS version should keep the risk of SPS being affected by malicious software at a minimum level.
Requirement 6.2	The report includes the firmware version running on SPS. You can check which is

PCI DSS Requirement	Compliance details
Ensure that all system components and software are protected from known vulnerabilities by installing applicable vendor-supplied security patches. Install critical security patches within one month of release.	the latest version at the Downloads page.
Requirement 8.1.8 If a session has been idle for more than 15 minutes, require the user to reauthenticate to re-activate the terminal or session.	The report includes the timeout value to the SPS web interface (10 minutes by default). To change this value, see Web interface timeout on page 328.
Requirement 8.2.1 Using strong cryptography, render all authentication credentials (such as passwords/phrases) unreadable during transmission and storage on all system components.	The report lists where SPS stores passwords, and the hash used to secure them.
Requirement 8.2.4 Change user passwords/passphrases at least every 90 days.	The report lists the password expiry settings of local web users of SPS, and also the last time the password of each user was changed. For details on configuring these accounts, see Managing One Identity Safeguard for Privileged Sessions (SPS) users locally on page 328. For details on configuring password expiry for these accounts, see Setting password policies for local users on page 330.
Requirement 8.2.5 Do not allow an individual to submit a new password/phrase that is the same as any of the last four passwords/phrases he or she has used.	The report includes the password history settings of local web users of SPS. For details on configuring password history for these accounts, see Setting password policies for local users on page 330.
Requirement 10.5.3 Promptly back up audit trail files to a centralized log server or media that is difficult to alter.	The report lists the addresses of the logservers where SPS forwards its log messages. For details on forwarding log messages, see Configuring system logging on page 125.
Requirement 10.5.4 Write logs for external-facing technologies onto a secure, centralized, internal log	The report lists the security settings of the communication between SPS and the logservers where SPS forwards its log

PCI DSS Requirement	Compliance details
server or media device.	messages. For details on forwarding log messages, see Configuring system logging on page 125.
Requirement 10.7 Retain audit trail history for at least one year, with a minimum of three months immediately available for analysis.	The retention time for local logs of SPS is seven days. To retain them longer, forward them to a remote logserver.

Report output

The output of the reports are available in PDF format.

The reports enable you to check the sessions that contain the configured search expressions. For security reasons, the relevant sessions are available through links to avoid including the session details directly in the report. You can open the related sessions using one of the following methods:

- Clicking the link in the Sessions key column of the table
- Clicking the QR code
- Clicking the link provided with the QR code

Searching by sessions keys

From the report, you can access the relevant sessions on the Search interface by clicking the link in the **Sessions key** column.

NOTE: If clicking on the sessions key does not work, for example, because the IP address of your SPS has changed, you can still use the key to access the relevant session as follows:

1. Copy the key from the **Sessions key** column, for example, *svc-uXG6ciAkstSanfD8YhGHXVddavid-7* as shown in the following example.
2. Create a link, using the IP address of your SPSand the sessions key in the following format:

   https://<*your-SPS-IP*>/portal/#/audit/sessions/<*sessions-key*>

   For example, if using the following example and an SPS with an IP address of *10.10.10.10*, the link is *https://10.10.10.10/portal/#/audit/sessions/svc-uXG6ciAkstSanfD8YhGHXV-ddavid-7*

Figure 318: Reporting > Download reports — Accessing sessions from the PDF output

Nr.	Protocol	Username	Server IP	Source IP	Session end	Session start	Length	Sessions key
1	SSH	titkos	10.12.15.128	10.12.15.164	2020-06-05 11:30:38	2020-06-05 11:30:20	18.00 seconds (0:00:18)	svc-uXG6ciAkstSanfD8YhGHXV-ddavid-7

Accessing screenshots from the report output using QR codes

NOTE: The screenshots of the sessions containing the search keywords are not included in the report output for security reasons, but you can access the screenshots by using the clickable QR codes.

Alternatively, you can click the link provided with the QR codes, which makes it possible to display the link text if the report is printed on paper.

Figure 319: Example of a clickable QR code in the report output



Import and visualize audit data from SPS in the Power BI Desktop reporting application

From One Identity Safeguard for Privileged Sessions (SPS) version 7.3, you can use the One Identity Safeguard Power BI Connector (Power BI Connector) to import and visualize audit data from SPS in the Power BI Desktop reporting application.

NOTE: Power BI Desktop integration using the Power BI Connector is only supported for SPS 7.3 and later feature versions.

For more information, see the *One Identity Safeguard Power BI Connector Tutorial* on the One Identity Support Portal.

The One Identity Safeguard for Privileged Sessions (SPS) REST API

Starting with One Identity Safeguard for Privileged Sessions (SPS) version 4 F2, certain parts and features of SPS can be configured using a REST API (Representational State Transfer Application Programming Interface). The REST server conforms to the Hypermedia as the Engine of Application State (HATEOAS).

The SPS REST API uses JSON over HTTPS. The REST server has a single entry point and all resources are available at paths (URLs) returned in the response for a request sent to the entry point. The only path that is guaranteed not to change is `/api/authentication`. Every other path should be reached by navigating the links returned.

The SPS REST API allows you to create, read, update and delete (CRUD) the configuration resources of SPS.

The user accessing the SPS REST API must have the **REST server** privilege. For details, see *Modifying group privileges* in the *Administration Guide*. For details on using the REST API, see *REST API Reference Guide*.

22

One Identity Safeguard for Privileged Sessions (SPS) scenarios

This section discusses common scenarios for One Identity Safeguard for Privileged Sessions (SPS).

Configuring public-key authentication on One Identity Safeguard for Privileged Sessions (SPS)

If a protected server requires public-key authentication from the users, complete one of the following procedures.

- In Configuring public-key authentication using local keys on page 911, One Identity Safeguard for Privileged Sessions (SPS) stores the public keys of the users and the private-public keypair used in the server-side connection locally on SPS.
- In Configuring public-key authentication using an LDAP server and a fixed key on page 912, SPS receives the public keys of the users from an LDAP server and uses a locally-stored private-public keypair in the server-side connection.
- In Configuring public-key authentication using an LDAP server and generated keys on page 914, SPS receives the public keys of the users from an LDAP server. SPS generates a keypair that is used in the server-side connection on-the-fly, then uploads the public key of this pair to the LDAP database. That way the server can authenticate SPS to the (newly generated) public key of the user.

Configuring public-key authentication using local keys

The following describes how to store the public keys of the users and the private-public keypair used in the server-side connection locally on One Identity Safeguard for Privileged Sessions (SPS).

To configure public-key authentication using local keys

1. Navigate to **Policies** > **Local User Databases** and create a Local User Database. Add the users and their public keys to the database. SPS will authenticate the clients to this database. For details on creating and maintaining local user databases, see Creating a Local User Database on page 538.
2. Navigate to **Policies** > **Credential Stores** and create a Local Credential Store. Add hostnames and the users to the database. SPS will use these credentials to authenticate on the target server. For details on creating local credential stores, see Configuring local Credential Stores on page 874.
3. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies** and create a new Authentication Policy.
4. Select **Authenticate the client to SPS using** > **Local** > **Public key**, clear all other options.
5. Select the appropriate usergroup from the **Local User Database** field. SPS will authenticate the users to this local database.
6. Select **Relayed authentication methods** > **Public key** > **Fix**, clear all other options.
7. Click  > **Generate**. This will generate a private key that is needed only for the configuration, it will not be used in any connection.

   NOTE: The Connection Policy will ignore the settings for server-side authentication (set under **Relayed authentication methods**) if a Credential Store is used in the Connection Policy.
8. Click **Commit**.
9. Navigate to **Traffic Controls** > **SSH** > **Connections** and create a new Connection.
10. Enter the IP addresses of the clients and the servers into the **From** and **To** fields.
11. Select the authentication policy created in Step 1 in the **Authentication Policy** field.
12. Configure the other options of the connection as necessary.
13. Click **Commit**.
14. To test the above settings, initiate a connection from the client machine to the server.

Configuring public-key authentication using an LDAP server and a fixed key

The following describes how to fetch the public keys of the users from an LDAP server and use a locally-stored private-public keypair in the server-side connection.

NOTE:

TIP: One Identity recommends using 2048-bit RSA keys (or stronger).

To configure public-key authentication using an LDAP server and a fixed key

1. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies** and create a new Authentication Policy.
2. Select **Authenticate the client to SPS using** > **LDAP** > **Public key**, deselect all other options.
3. Select **Relayed authentication methods** > **Public key** > **Fix**, deselect all other options.
4. Select **Private key** and click  . A pop-up window is displayed.
5. Click **Browse** and select the private key of the user, or paste the key into the **Copy-paste** field. Enter the password for the private key into the **Password** field and click **Upload**.

   NOTE: SPS accepts passwords that are not longer than 150 characters and supports the following characters:

   - Letters A-Z, a-z
   - Numbers 0-9
   - The space character
   - Special characters: `!"#$%&'()*+,-./:;<>=?@[]\^-`{}_|`

   If the private key of the user is not available, click **Generate** to create a new private key. You can set the size of the key in the **Generate key** field. In this case, do not forget to export the public key from SPS and import it to the server. To export the key from SPS, just click on the key and save it to your local computer.
6. Click on the fingerprint of the key in the **Server side private and public key** > **Private key** field and save the public key. Do not forget to import this public key to the server: all connections that use this new authentication policy will use this keypair on the server side.
7. Click **Commit**.
8. Navigate to **Policies** > **LDAP Servers** and click  to create a new LDAP policy.
9. Enter the parameters of the LDAP server. For details, see Authenticating users to an LDAP server on page 505.
10. If different from `sshPublicKey`, enter the name of the LDAP attribute that stores the public keys of the users into the **Publickey attribute name** field.

    CAUTION:

    The public keys stored in the LDAP database must be in OpenSSH format.
11. Navigate to **Traffic Controls** > **SSH** > **Connections** and create a new Connection.
12. Enter the IP addresses of the clients and the servers into the **From** and **To** fields.
13. Select the authentication policy created in Step 1 from the **Authentication Policy** field.

14. Select the LDAP policy created in Step 7 from the **LDAP Server** field.
15. If the server accepts a user only from a specific IP address, select the **Use original IP address of the client** radiobutton from the **SNAT** field.
16. Configure the other options of the connection as necessary.
17. Click **Commit**.
18. To test the above settings, initiate a connection from the client machine to the server.

Configuring public-key authentication using an LDAP server and generated keys

The following describes how to fetch the public keys of the users from an LDAP server and have One Identity Safeguard for Privileged Sessions (SPS) generate a keypair that is used in the server-side connection on-the-fly, and upload the public key of this pair to the LDAP database.

To configure public-key authentication using an LDAP server and generated keys

1. Navigate to **Traffic Controls** > **SSH** > **Authentication Policies** and create a new Authentication Policy.
2. Select **Authenticate the client to SPS using** > **LDAP** > **Public key**, deselect all other options.
3. Select **Relayed authentication methods** > **Public key** > **Publish to LDAP**, deselect all other options.
4. Click **Commit**.
5. Navigate to **Policies** > **LDAP Servers** and click  to create a new LDAP policy.
6. Enter the parameters of the LDAP server. For details, see Authenticating users to an LDAP server on page 505.
7. If different from `sshPublicKey`, enter the name of the LDAP attribute that stores the public keys of the users into the **Publickey attribute name** field.

   CAUTION:

   The public keys stored in the LDAP database must be in OpenSSH format.
8. Enter the name of the LDAP attribute where SPS shall upload the generated keys into the **Generated publickey attribute name** field.
9. Click **Commit**.
10. Navigate to **Traffic Controls** > **SSH** > **Connections** and create a new Connection.
11. Enter the IP addresses of the clients and the servers into the **From** and **To** fields.

12. Select the authentication policy created in Step 1 from the **Authentication Policy** field.
13. Select the LDAP policy created in Step 7 from the **LDAP Server** field.
14. If the server accepts a user only from a specific IP address, select the **Use original IP address of the client** radiobutton from the **SNAT** field.
15. Configure the other options of the connection as necessary.
16. Click **Commit**.
17. To test the above settings, initiate a connection from the client machine to the server.

Organizing connections in non-transparent mode

When using One Identity Safeguard for Privileged Sessions (SPS) in non-transparent mode, the administrators must address SPS to access the protected servers. If an administrator has access to more than one protected server, SPS must be able to determine which server the administrator wants to access. For each protected server, the administrators must address either different ports of the configured interface, or different alias IP addresses.

Organizing connections based on port numbers

To allow the administrators to access protected servers by connecting to the IP address of One Identity Safeguard for Privileged Sessions (SPS), and use the port number to select which server they want to access. Organizing connections based on port numbers is advantageous if SPS has a public IP address and the protected servers must be administered from the Internet.

NOTE: Do not use the listening addresses configured for web login. For more details, see Configuring user and administrator login addresses on page 117.

For details on configuring alias IP addresses, see Managing logical interfaces on page 119.

To organize connections based on port numbers

1. Navigate to the **Connections** tab of the **Traffic Controls** > **SSH** menu.
2. Add a new connection. Enter the IP address of the administrators into the **From** fields, and the IP address and port number of the server into the **Target** field.
3. Enter the IP address of the logical interface of SPS into the **To** field, and enter a port number into the **Port** field.

4. Repeat Steps 2-3 for every protected server, but every time use a different port number in Step 3.
5. Click **Commit**.

Organizing connections based on alias IP addresses

To allow the administrators to access protected servers by connecting to an alias IP address of One Identity Safeguard for Privileged Sessions (SPS). The alias IP address determines which server they will access. Organizing connections based on alias IP addresses is advantageous if SPS is connected to a private network and many private IP addresses are available.

NOTE: Do not use the listening addresses configured for web login. For more details, see Configuring user and administrator login addresses on page 117.

To organize connections based on alias IP addresses

1. Navigate to **Basic Settings** > **Network**.
2. Set up a logical interface: click + and configure a new logical interface. Add alias IP addresses for every protected server. (Use a different IP address for each.)

   For more information on configuring logical interfaces and alias IP addresses, see Managing logical interfaces on page 119.
3. Navigate to **Traffic Controls** > **SSH** > **Connections**.
4. Add a new connection. Enter the IP address of the administrators into the **From** fields, and the IP address and port number of the target server into the **Target** field.
5. Enter an alias IP address of the configured logical interface of SPS into the **To** field.
6. Repeat Steps 4-5 for every protected server, but every time use a different alias IP address in Step 5.
7. Click **Commit**.

Using inband destination selection in SSH connections

The following sections provide examples for using inband destination selection to establish an SSH connection, including scenarios where nonstandard ports or gateway authentication is used.

Since some client applications do not permit the `@` and `:` characters in the username, alternative characters can be used as well:

- To separate the username and the target server, use the `@` or `%` characters, for example: **`username%targetserver@scb_address`**
- To separate the target server and the port number, use the `:`, `+`, or `/` characters, for example: `username%targetserver+port@scb_address`
- If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

In RDP, do not use the `@` character as an inband data separator but use alternative characters, for example, the `%` character.

For detailed instructions on configuring inband authentication, see Configuring inband destination selection on page 490.

Using inband destination selection with PuTTY

To establish an SSH connection through One Identity Safeguard for Privileged Sessions (SPS) with PuTTY, follow one of the methods:

Common method

To establish the SSH-connection using the most common method, enter the username, the target server's hostname (or IP address), and the hostname (or IP address) of SPS using the `<username>@<server>@<scb>` format in PuTTY.

If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

Example

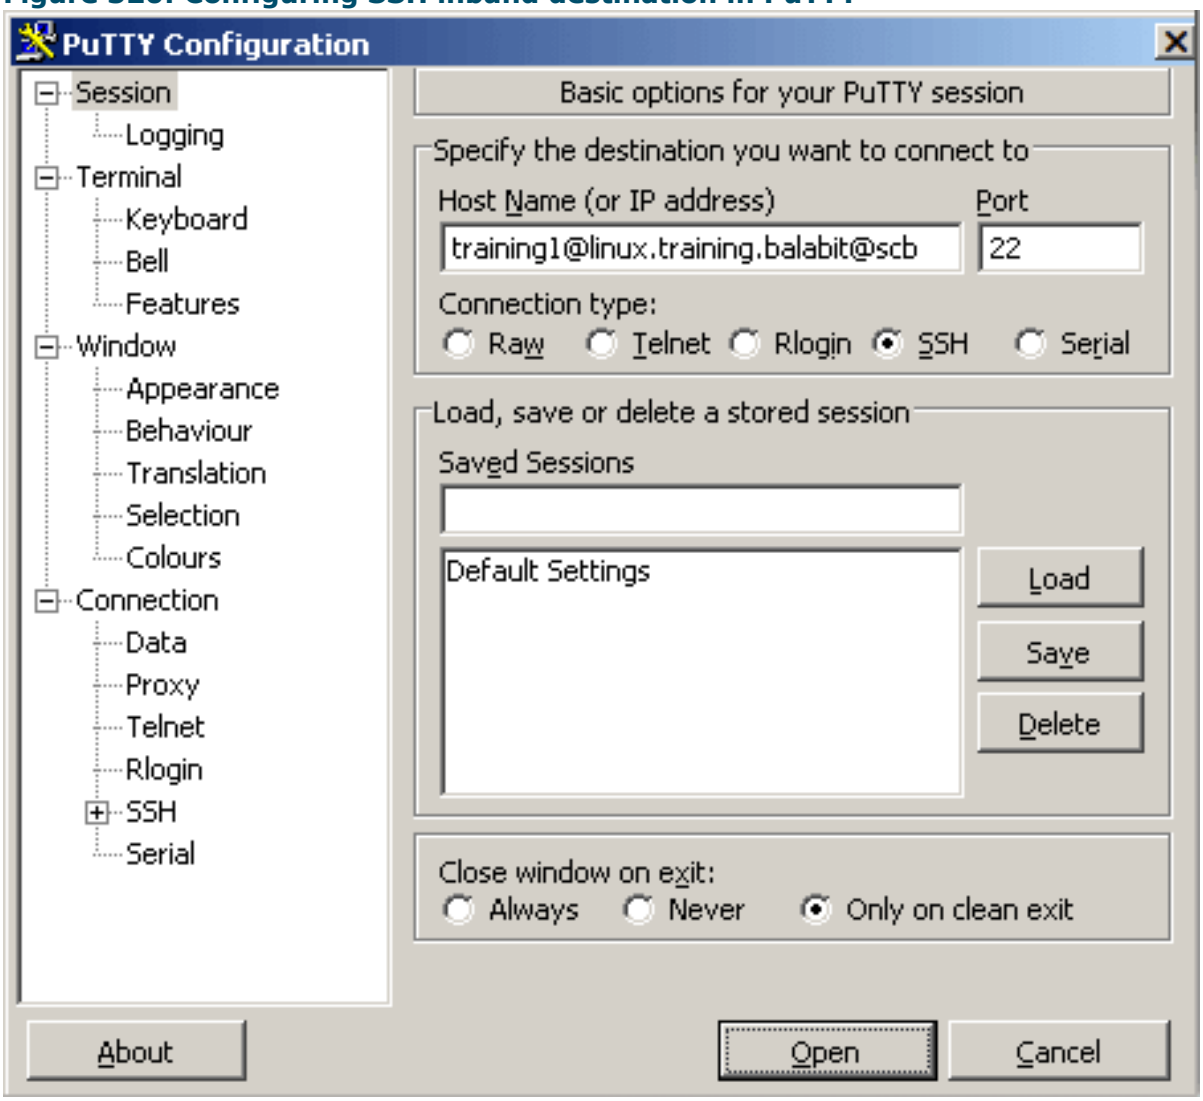
Assuming the following values:

- The username is `training1`
- The target server is `linux.training.example`
- The SPS server is `scb`

You can enter the following destination in PuTTY:

`training1@linux.training.example@scb`

Figure 320: Configuring SSH inband destination in PuTTY



Alternative method

To establish the SSH-connection using a different method,

1. Enter only the hostname (or IP address, depending on your configuration) of SPS in PuTTY.
2. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) using the `<username>@<server>` format.

Using inband destination selection with OpenSSH

To establish an SSH connection through One Identity Safeguard for Privileged Sessions (SPS), follow these steps:

1. Enter the following command:

   ```
   # ssh <username>@<server>@<scb>
   ```

   ...where `<username>` is the username, `<server>` is the target server's hostname (or IP address), and `<scb>` is the hostname (or IP address) of SPS.

   If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

   Example

   Assuming the following values:

   - The username is `training1`
   - The target server is `linux.training.example`
   - The SPS server is `scb`

   You can enter the following command:

   ```
   # ssh training1@linux.training.example@scb
   ```

2. Alternative approach:

   a. Enter only the hostname (or IP address, depending on your configuration) of SPS:

      ```
      # ssh <scb>
      ```

   b. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) using the `<username>@<server>` format

Using inband selection and nonstandard ports with PuTTY

The following steps provide instructions for establishing SSH connections with servers that are listening on a non-standard port (the **Inband destination selection** > **Targets** > **Port** option is not 22), and the port number targeted by the clients is also a non-standard port (the **To** > **Port** option of the Connection Policy).

1. Enter the following in PuTTY:
    a. In the **Host Name** field, enter the username on the target server, the target server's hostname (or IP address) and port number, and the hostname (or IP address) of One Identity Safeguard for Privileged Sessions (SPS) in the `<username>@<server>:<port>@<scb>` format

        If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

    b. In the **Port** field, enter the port number of the SPS server

Example

Assuming the following values:

- The username is `training1`
- The target server is `192.168.60.100`
- The target server is listening on port `2121`
- The SPS server is `scb`
- The SPS server is listening on port `4444`

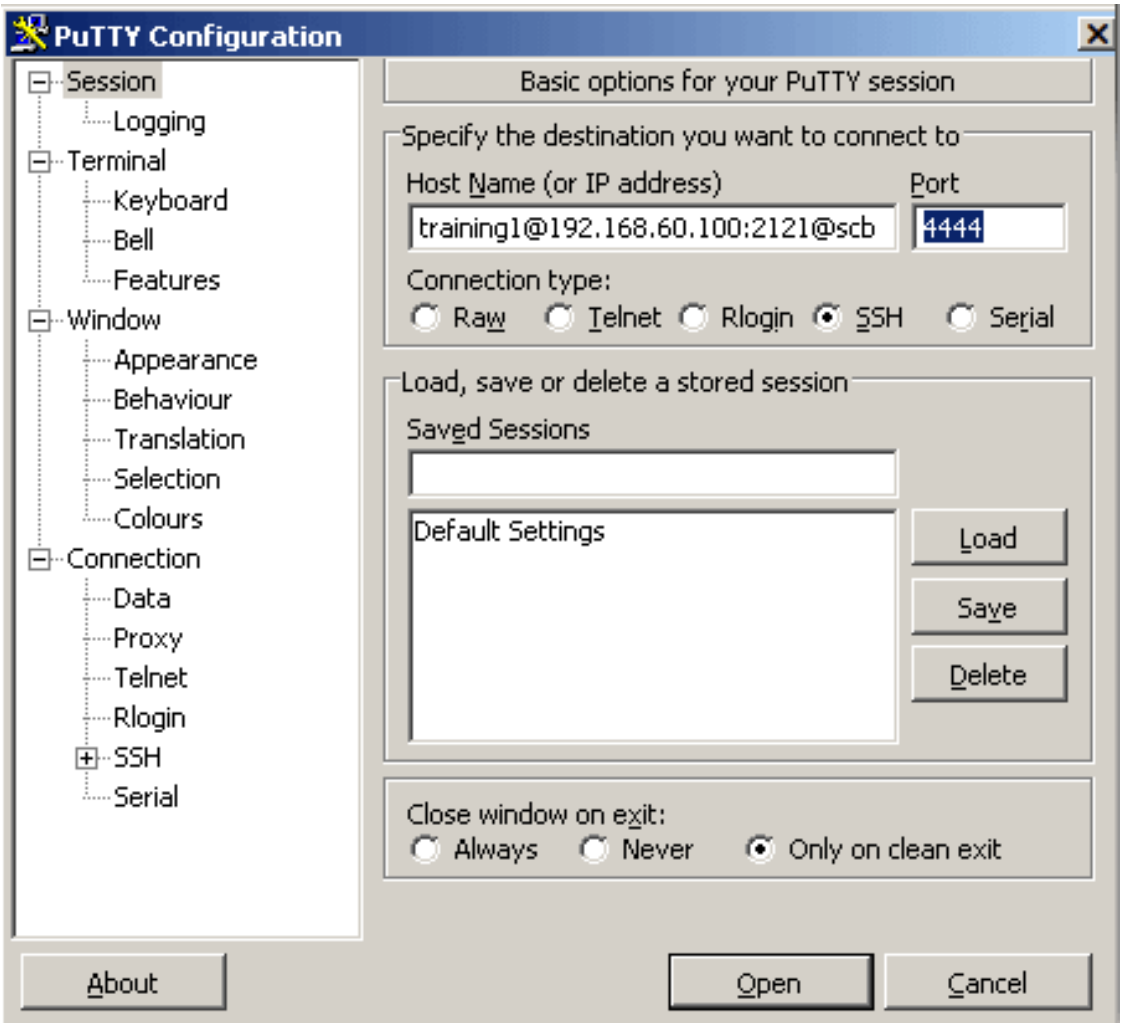
You can enter the following destination hostname in PuTTY:

`training1@192.168.60.100:2121@scb`

Also change the destination port to the SPS server's port number:

4444

Figure 321: Configuring SSH inband destination for nonstandard ports in PuTTY



2. Alternative approach:
    a. Enter only the hostname (or IP address, depending on your configuration) and port number of SPS in PuTTY.
    b. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) and port number using the `<username>@<server>:<port>` format.

Using inband selection and nonstandard ports with OpenSSH

The following steps provide instructions for establishing SSH connections with servers that are listening on a non-standard port (the **Inband destination selection** > **Targets** >

Port option is not 22), and the port number targeted by the clients is also a non-standard port (the **To** > **Port** option of the Connection Policy).

1. Enter the following command:

   ```
   # ssh -p <scb_port> <username>@<server>:<port>@<scb>
   ```

   ...where `<scb_port>` is the port number of One Identity Safeguard for Privileged Sessions (SPS), `<username>` is the username on the target server, `<server:port>` is the target server's hostname (or IP address), `<port>` is the target server's port number, and `<scb>` is the hostname (or IP address) of SPS.

   If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

   ### Example

   Assuming the following values:

   - The username is `training1`
   - The target server is `192.168.60.100`
   - The target server is listening on port `2121`
   - The SPS server is `scb`
   - The SPS server is listening on port `4444`

   You can enter the following command:

   ```
   # ssh -p 4444 training1@192.168.60.100:2121@scb
   ```

2. Alternative approach:

   a. Enter only the hostname (or IP address, depending on your configuration) and port number of SPS with the following command:

      ```
      # ssh -p <scb_port> <scb>
      ```

   b. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) and port number using the `<username>@<server>:<port>` format.

Using inband destination selection and gateway authentication with PuTTY

SPS can authenticate users attempting to establish an SSH connection against a gateway (see Configuring gateway authentication on page 862 for more details). You can provide the gateway login credentials in PuTTY:

1. Enter the gateway username, the username on the target server, the target server's hostname (or IP address), and the hostname (or IP address) of One Identity

Safeguard for Privileged Sessions (SPS) in the `gu=<gatewayusername>@<username>@<server>@<scb>` format in PuTTY

If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

Example

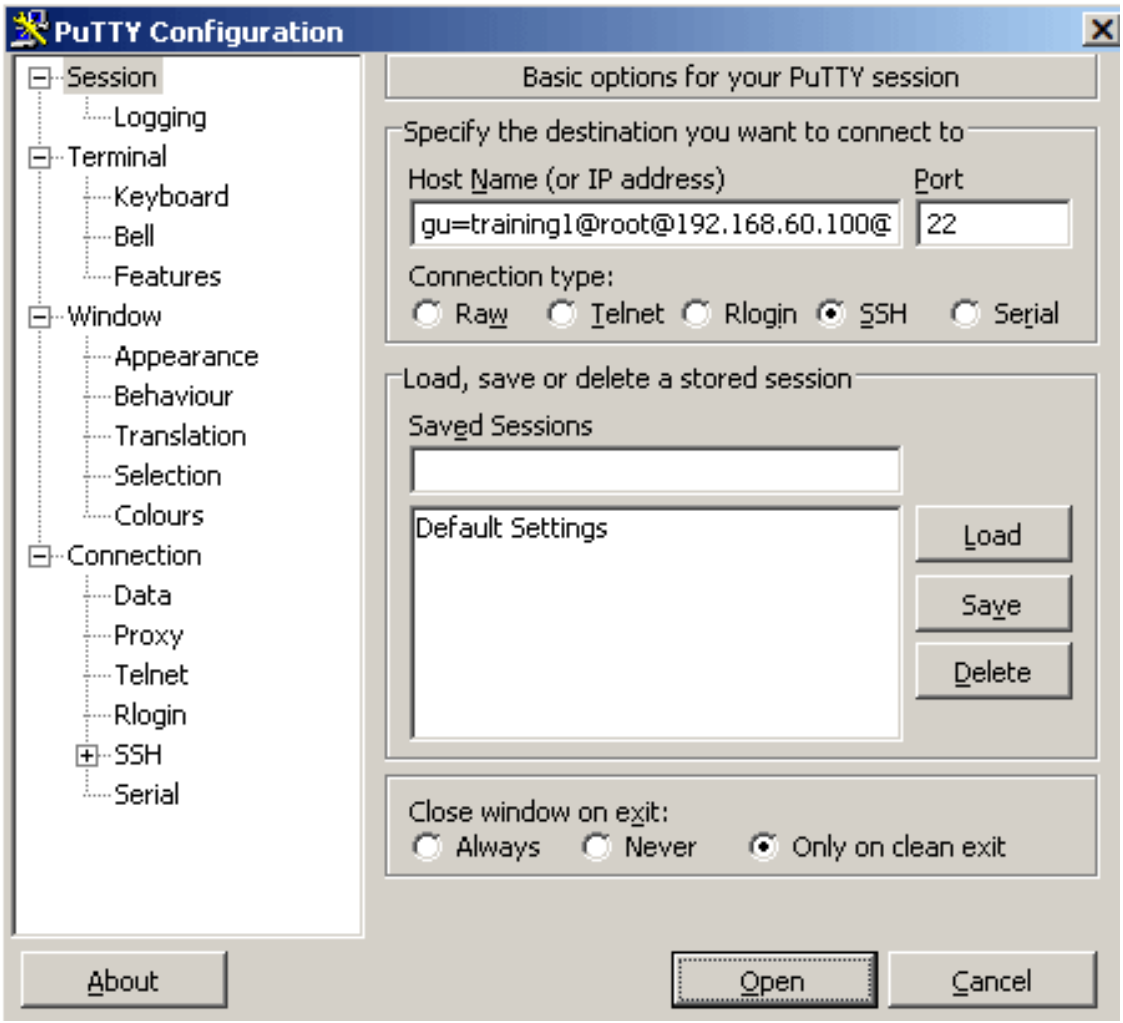
Assuming the following values:

- The gateway username is `training1`
- The username on the target server is `root`
- The target server is `192.168.60.100`
- The SPS server is `scb`

You can enter the following destination in PuTTY:

```
gu=training1@root@192.168.60.100@scb
```

Figure 322: Configuring SSH inband destination and gateway authentication in PuTTY



2. Alternative approach:
    a. Enter only the hostname (or IP address, depending on your configuration) of SPS in PuTTY.
    b. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) using the `<username>@<server>` format.
    c. When prompted, provide the gateway username.

Using inband destination selection and gateway authentication with OpenSSH

One Identity Safeguard for Privileged Sessions (SPS) can authenticate users attempting to establish an SSH connection against a gateway (see Configuring gateway authentication on

page 862 for more details). The following steps explain how you can provide the gateway login credentials:

1. Enter the following command:

   ```
   # ssh gu=<gatewayusername>@<username>@<server>@<scb>
   ```

   ...where `<gatewayusername>` is the gateway username, `<username>` is the username on the target server, `<server>` is the target server's hostname (or IP address), and `<scb>` is the hostname (or IP address) of SPS.

   If you do not specify the username or the address in nontransparent SSH and Telnet connections, One Identity Safeguard for Privileged Sessions (SPS) displays an interactive prompt where you can enter the username and the server address.

   Example

   Assuming the following values:

   - The gateway username is `training1`
   - The username on the target server is `root`
   - The target server is `192.168.60.100`
   - The SPS server is `scb`

   You can enter the following command:

   ```
   # ssh gu=training1@root@192.168.60.100@scb
   ```

2. Alternative approach:

   a. Enter only the hostname (or IP address, depending on your configuration) of SPS with the following command:

      ```
      # ssh <scb>
      ```

   b. At the login prompt, provide the username on the target server, and the target server's hostname (or IP address) using the `<username>@<server>` format.

   c. When prompted, provide the gateway username.

SSH usermapping and keymapping in AD with public key

A customer wants to be able to disable password authentication in SSH for `admin` users on the UNIX servers. However, the customer uses Active Directory, and would not like to enter the user name/password at gateway authentication for every login over and over again. Therefore, the customer needs a quasi SSO-like system, with only one group logging in as `root` and another group as `XY` user.

To perform SSH usermapping and keymapping in AD with public key

1. Create an LDAP authentication policy. For details on creating a new authentication policy, see Authentication Policies on page 625. In this scenario, only a few important details will be highlighted.
    a. In the **Authenticate the client to SPS using** field, set the authentication method used on the client-side to **LDAP**. This will be the Active Directory where the gateway will get the public key from, for authentication. Enable **Publickey** only from the **Allowed authentication methods** and disable all other methods.
    b. In the **Relayed authentication methods** field, enable **Public key** and select **Agent**. Disable all other methods.
2. Create a Credential Store that can return a private key for server-side authentication. It is local Credential Stores and external Credential Stores (with a suitable plugin) that can return a private key.

    For detailed step-by-step instructions, see Configuring local Credential Stores on page 874. In this scenario, only a few important details will be highlighted.

    a. Navigate to the bottom of the policy, and click  to add a new user.
    b. Enter the username in the **Username** field (for example: `root`). Generate a **Private key** and upload its public counterpart to the server.
3. Set an LDAP server policy where you set up the Active Directory. For details on authenticating users to an LDAP server, see Authenticating users to an LDAP server on page 505.

    Make sure that the **Publickey attribute name** field in this Active Directory LDAP policy is set to `sshPublicKey`.
4. By default, the Active Directory does not have any attribute that could store the SSH public key. To solve this, add an OpenSSH-LPK compatible schema to the Active Directory by doing any of the following:

    - Create an `sshPublicKey` attribute, and add that directly to one of the objectClasses of the user in question.
    - Create an `sshPublicKey` attribute and an `ldapPublicKey` auxiliary objectClass, and add the `ldapPublicKey` auxiliary objectClass to one of the objectClasses of the user in question.

    The `sshPublicKey` attribute must be compliant with the OpenSSH-LPK schema and have the following properties:

    - Name: sshPublicKey
    - Object ID: 1.3.6.1.4.1.24552.500.1.1.1.13
    - Syntax: Octet String
    - Multi-Valued

The `ldapPublicKey` auxiliary objectClass must be compliant with the OpenSSH-LPK schema and have the following properties:

- Name: ldapPublicKey
- OID: 1.3.6.1.4.1.24552.500.1.1.2.0

The OpenSSH-LPK schema is available on the openssh-lpk Google Code page.

The following steps describe how to create an `sshPublicKey` attribute and an `ldapPublicKey` auxiliary objectClass, and then add the `ldapPublicKey` auxiliary objectClass to one of the objectClasses of the user.

a. Enable Schema updates using the registry:

   i. Click **Start**, click **Run**, and then in the **Open** box, type: `regedit`. Press`Enter`.

   ii. Locate and click the following registry key: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters`.

   iii. On the **Edit** menu, click **New**, and then click **DWORD Value**.

   iv. Enter the value data when the following registry value is displayed:

   ```
   Value Name: Schema Update Allowed
   Data Type: REG_DWORD
   Base: Binary
   Value Data: 1
   ```

   NOTE:Type **`1`** to enable this feature, or **`0`** (zero) to disable it.

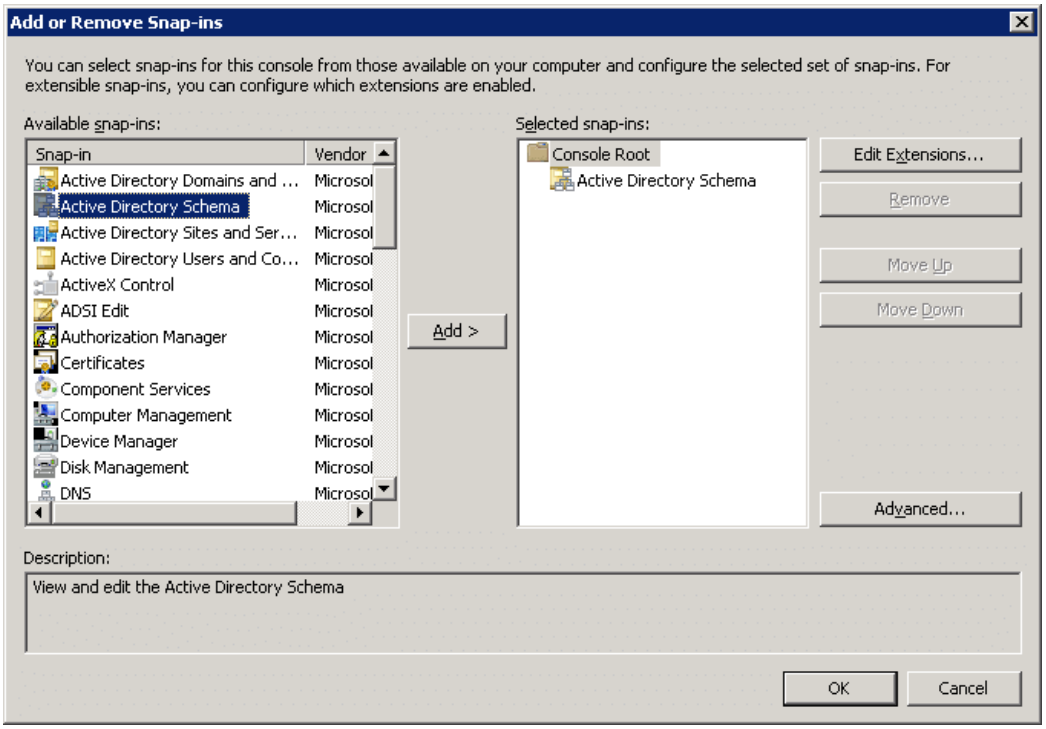
   v. Quit Registry Editor.

b. Install the Schema snap-in. For details, see the Microsoft Documentation. Note that you must have Administrator privileges to install the Schema snap-in.

c. Click **Start**, click **Run**, and then in the **Open** box, type: `MMC`. Press **Enter**.

d. Navigate to **File** > **Add or Remove Snap-in**, select **Active Directory Schema** and click **Add**. Note that you must have Schema Administrator privileges to complete the following steps.

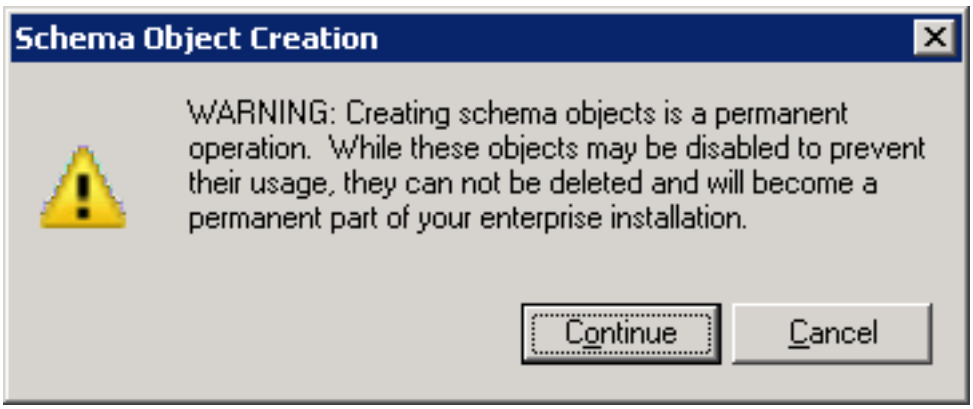
Figure 323: Add or Remove Snap-in



e. Expand the Active Directory schema and right-click **Attributes**.

f. Click **Create Attribute**. If a warning appears, click **Continue**.

Figure 324: Creating Attribute - Warning



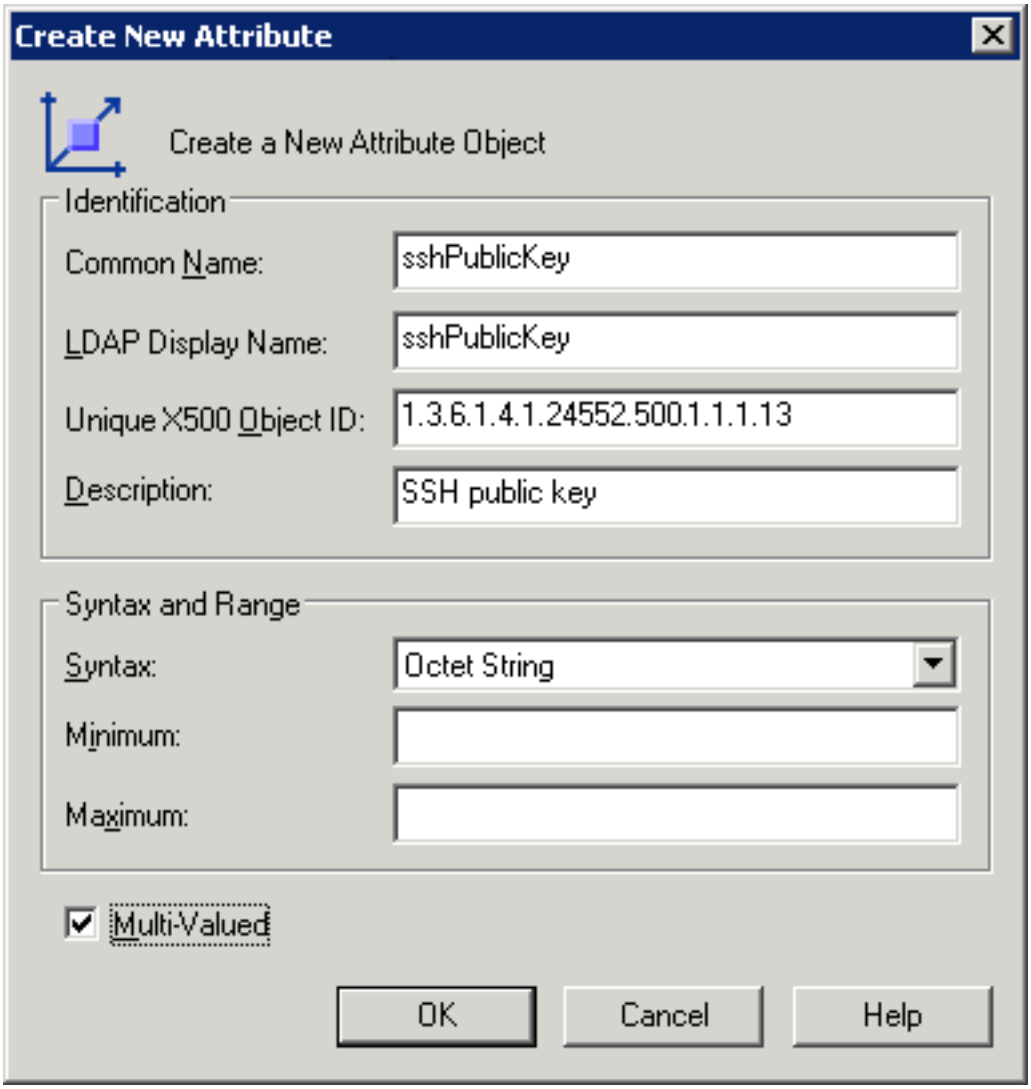
g. In **Common Name** and **LDAP Display Name**, enter `sshPublicKey`.

h. In **Unique X500 Object ID**, enter `1.3.6.1.4.1.24552.500.1.1.1.13`.

i. For **Syntax**, select `Octet String`.

j. Enable **Multi-Valued**. Click **OK**.

Figure 325: Create New Attribute



k. Right-click **Classes** and click **Create class**. If a warning appears, click **Continue**.

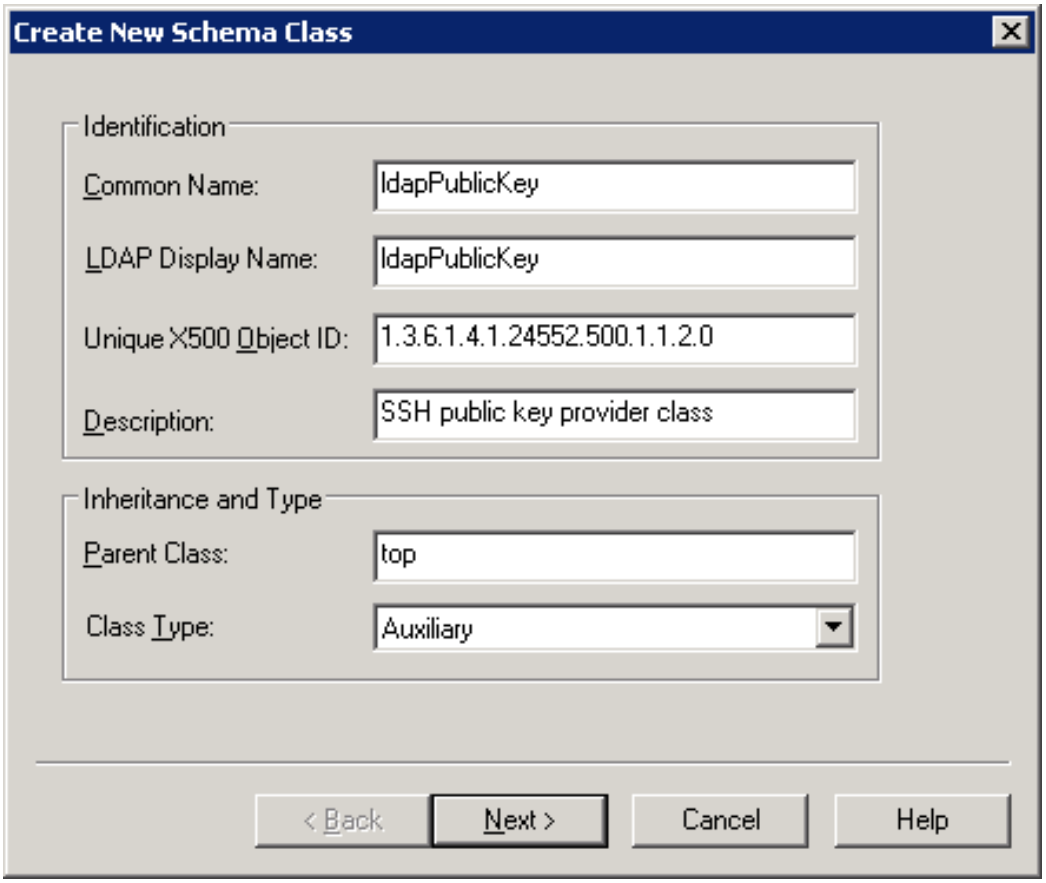
l. In **Common Name** and **LDAP Display Name**, enter `ldapPublicKey`.

m. In **Unique X500 Object ID**, enter `1.3.6.1.4.1.24552.500.1.1.2.0`

n. Create a new schema class.

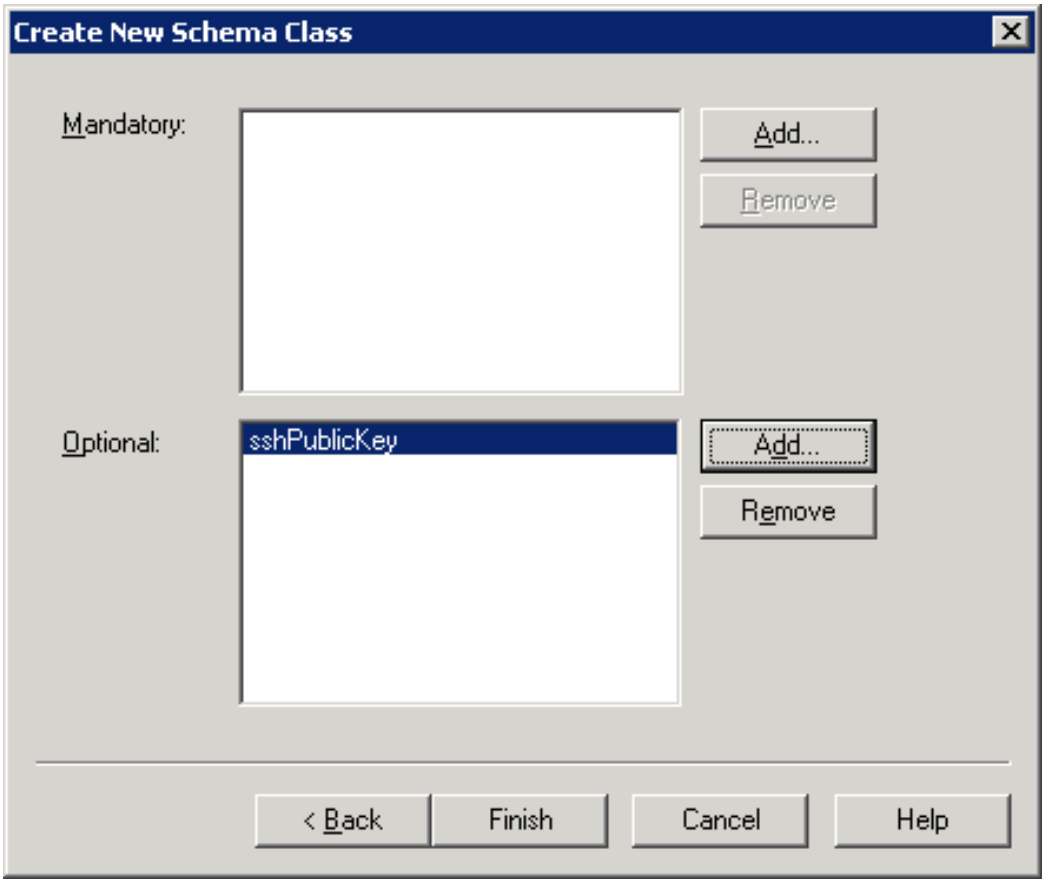
   In **Parent Class**, enter `top`, and in **Class Type**, select `Auxiliary`. Click **Next**.

Figure 326: Create New Schema Class — screen 2



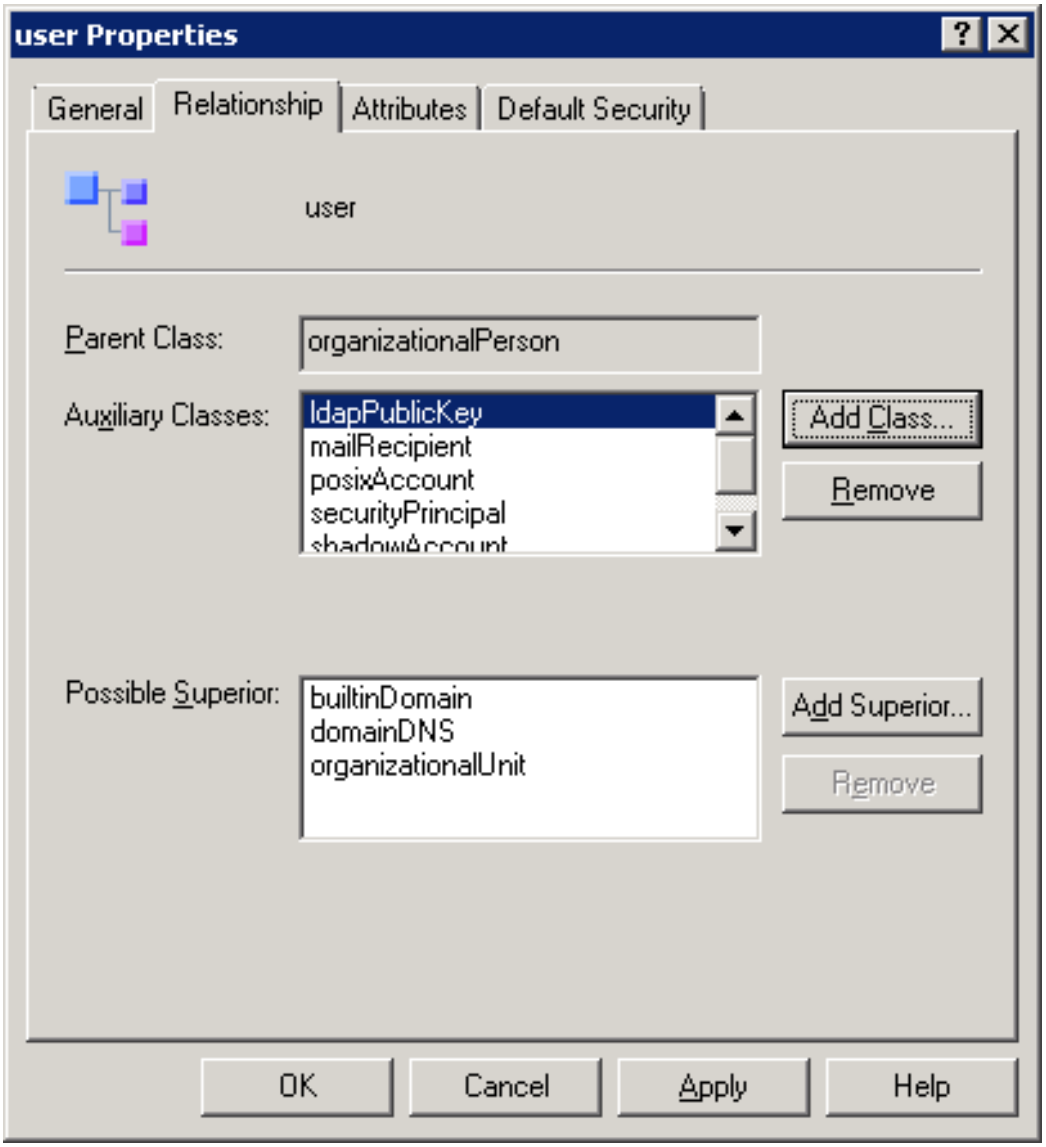
Add `sshPublicKey` to the **Optional** field. Click **Finish**.

Figure 327: Create New Schema Class — screen 1



o. Expand **Classes** and select **User**. Right-click **User** and select **Properties**. Navigate to **Relationship** > **Auxiliary Classes**, click **Add Class** and add the **ldapPublicKey** class. Click **Apply**.

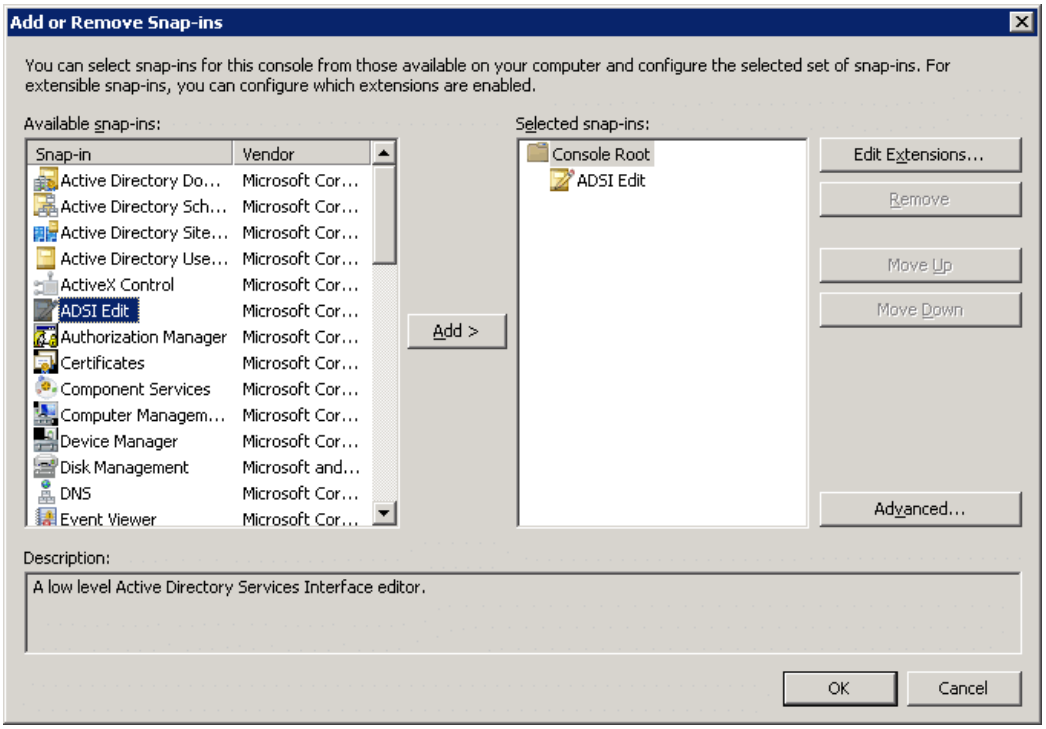
Figure 328: User Properties



5. The next step is to map the public keys to users. This is not possible in a user editor, use a low-level LDAP utility instead.

a. Add ADSI Edit as a snap-in to MMC.

Figure 329: Adding ADSI Edit



b. Right-click on the node and press **Enter**.

c. Search for the user in the tree, right-click on it and select *Properties.* All attributes can be edited there, so `sshPublicKey` too. Add the public keys to the Active Directory users.

> NOTE:It may happen that `sshPublicKey` is not visible in ADSI Edit. To make `sshPublicKey` visible, complete the procedure described in section Extending the Partial Attribute Set in https://blogs.technet.microsoft.com/scotts-it-blog/2015/02/28/ad-ds-global-catalogs-and-the-partial-attribute-set/.

6. Create a usermapping policy where you will set those groups from the Active Directory who can become root. For details on creating usermapping policies, see Configuring usermapping policies on page 860. In this scenario, only a few important details will be highlighted.

   a. Set **Username on the server** to `root` and select the group you intend to give these rights to.

   b. If you intend to allow other users in without usermapping, enable **Allow other unmapped usernames**.

7. Navigate to the relevant connection on the **Traffic Controls** > **SSH** > **Connections** page, and do the following:

a. In the **Authentication policy** field, add the LDAP authentication policy you created in Step 1.
b. In the **LDAP Server** field, add the LDAP server policy you created in Step 3.
c. In the **Credential Store** field, add the Credential Store you created in Step 2.
d. In the **Usermapping Policy** field, add the usermapping policy you created in Step 6.
e. To save the change, click **Commit**.

Troubleshooting One Identity Safeguard for Privileged Sessions (SPS)

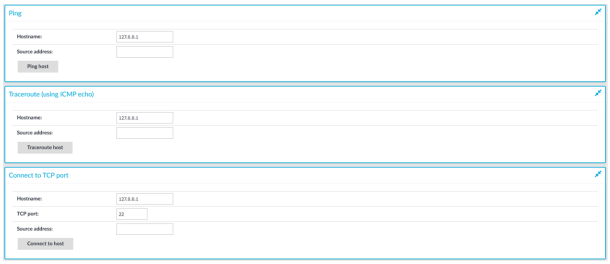
This section describes the tools to detect networking problems, and also how to collect core dump files and view the system logs of One Identity Safeguard for Privileged Sessions (SPS).

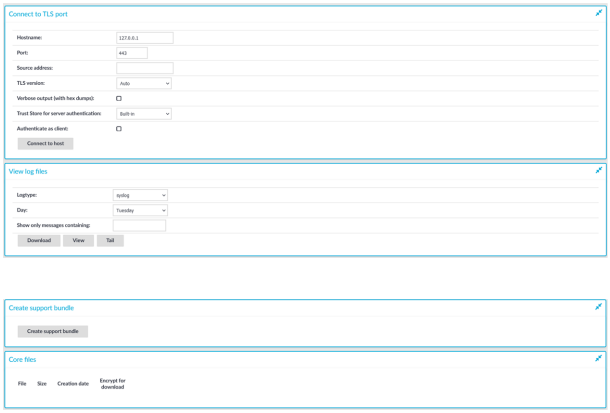
If you need to find the SPS appliance in the server room, you can use IPMI to control the front panel identify light. On One Identity Safeguard for Privileged Sessions N10000, navigate to **Basic Settings** > **System** > **Hardware information** > **Blink identification lights** and click **On** to blink the LEDs of hard disk trays on the front of the SPS appliance in red.

Network troubleshooting

The **Basic Settings** > **Troubleshooting** menu provides a number of diagnostic commands to resolve networking issues. Logfiles of One Identity Safeguard for Privileged Sessions (SPS) can also be displayed here — for details, see Viewing logs on One Identity Safeguard for Privileged Sessions (SPS) on page 937.

Figure 330: Basic Settings > Troubleshooting — Network troubleshooting with SPS





The following commands are available:

- `ping`: Sends a simple message to the specified host to test network connectivity.
- `traceroute`: Sends a simple message from SPS to the specified host and displays all hosts on the path of the message. It is used to trace the path the message travels between the hosts.
- `connect`: Attempts to connect the specified host using the specified port. It is used to test the availability or status of an application on the target host.

To run one of the above commands

1. Navigate to **Basic Settings** > **Troubleshooting**.
2. Enter the IP address or the hostname of the target host into the **Hostname** field of the respective command. For the `Connect` command, enter the target port into the **TCP port** field.

   Use an IPv4 address.
3. Click the respective action button to run the command.
4. Check the results in the pop-up window. Log files are displayed in a separate browser window.

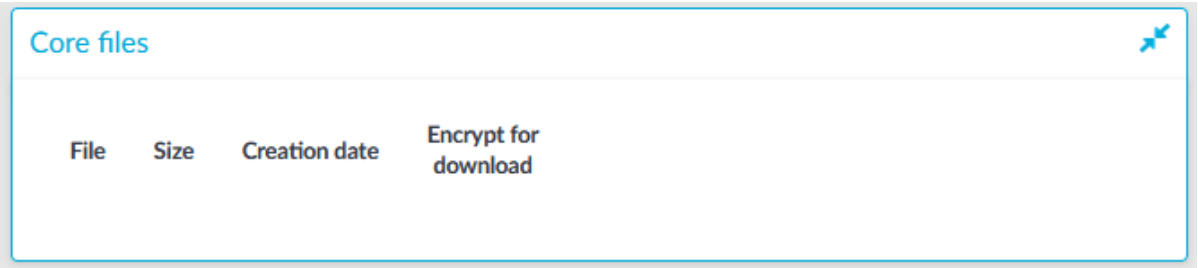
Gathering data about system problems

One Identity Safeguard for Privileged Sessions (SPS) automatically generates core dump files if an important software component (for example, Zorp) of the system crashes for some reason. These core dump files can be of great help to the One Identity Support Team to identify problems. When a core dump file is generated, the SPS administrator receives an alerting e-mail, and an SNMP trap is generated if alerting is properly configured (for details, see Configuring system monitoring on SPS on page 136 and System logging, SNMP and e-mail alerts on page 125).

To list and download the generated core dump files, navigate to **Basic Settings** > **Troubleshooting** > **Core files**.

By default, core dump files are deleted after 14 days. To change the deletion timeframe, navigate to **Basic Settings** > **Management** > **Core files**.

Figure 331: Basic Settings > Troubleshooting — System troubleshooting with SPS



Viewing logs on One Identity Safeguard for Privileged Sessions (SPS)

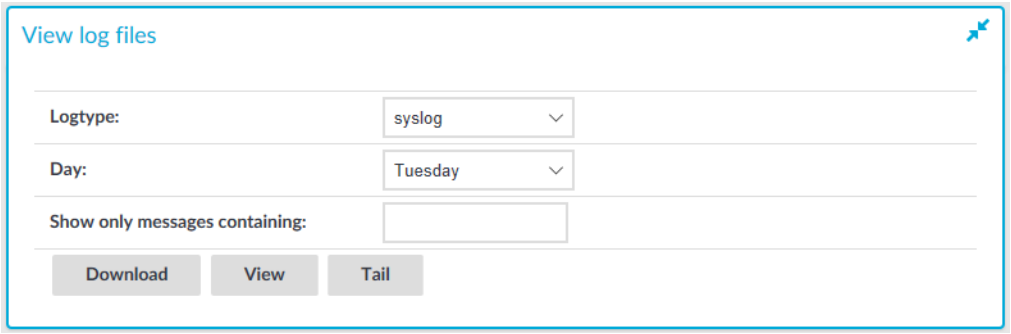
The **Troubleshooting** menu provides an interface to view the logs generated by the various components of One Identity Safeguard for Privileged Sessions (SPS).

NOTE: Because of performance reasons, log files larger than 2 Megabytes are not displayed in the web interface. To access these logs, download the file instead.

To view logs on SPS

1. Navigate to **Basic Settings** > **Troubleshooting** > **View log files**.

   Figure 332: Basic Settings > Troubleshooting — Viewing logs on SPS

2. Use the **Logtype** roll-down menu to select the message type.
   - `syslog`: All system logs of the SPS host.
   - `scb`: Logs of the SPS web interface.
   - `paa`: Logs related to the workings of the One Identity Safeguard for Privileged Analytics module.
   - `logadapter`: Logs of the log adapter plugin(s) and syslog instance(s) configured for ingesting logs from an external source.

- `http`: Logs of the HTTP connections passing through SPS.
- `ica`: Logs of the ICA connections passing through SPS.
- `rdp`: Logs of the RDP connections passing through SPS.
- `ssh`: Logs of the SSH connections passing through SPS.
- `telnet`: Logs of the Telnet connections passing through SPS.
- `vnc`: Logs of the VNC connections passing through SPS.

3. Use the buttons at the bottom of the dialog to perform the following tasks:
   - To download the log file, click **Download**.
   - To follow the current log messages real-time, click **Tail**.
   - To display the log messages, click **View**.
4. To display log messages of the last seven days, select the desired day from the **Day** field and click **View**.

   TIP: To display only the messages of a selected host or process, enter the name of the host or process into the **Show only messages containing** field.

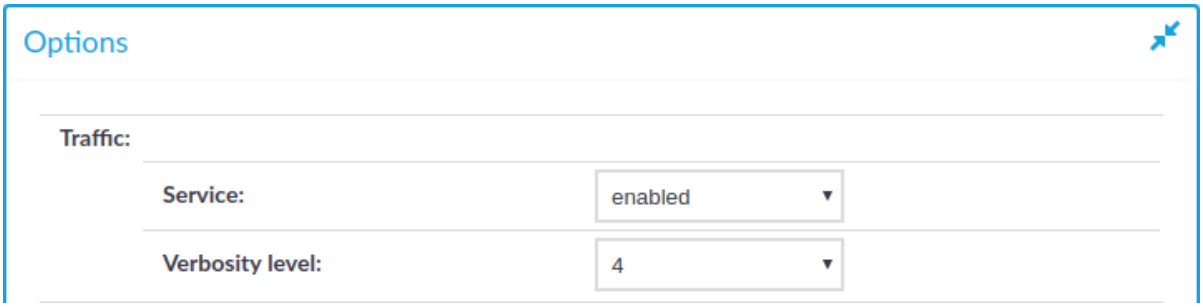
   The **Show only messages containing** field acts as a generic filter: enter a keyword or a regular expression to display only messages that contain the keyword or match the expression.

Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS)

The logging level of One Identity Safeguard for Privileged Sessions (SPS) can be set separately for every protocol.

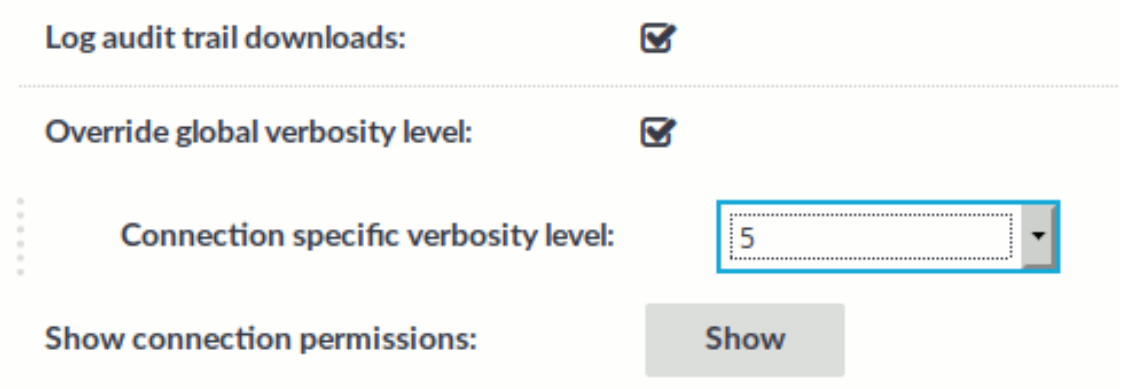
NOTE: The **Basic Settings** > **Management** > **Verbose system logs** > **Enable** option is not related to the verbosity of traffic logs: it increases the log level of the non-network-related events, for example adds the commands executed by the SPS web interface to the logs, and so on.

Figure 333: <Protocol name> Control > Global Options — Changing the verbosity level



To change the verbosity level of SPS

1. Navigate to the **Global Options** page of the traffic you want to change the log level of, for example, to **Traffic Controls** > **SSH** > **Global Options** to change the log level of SSH traffic, **Traffic Controls** > **RDP** > **Global Options** for remote desktop traffic, and so on.



2. Select the desired log level from the **Verbosity level** field. Note that the new verbosity level applies only to new sessions started after committing the change. The verbosity level of active sessions will not change.

   NOTE: The verbosity level ranges from 1 (no logging) to 10 (extremely detailed), with level 4 being the default normal level. To debug complex problems, you might have to increase the verbosity level to 7. Higher level is needed only in extreme cases.

   CAUTION:

   High verbosity levels generate very large amount of log messages and might result in a very high load on the machine.

   For log levels 8-10, the logs contain highly sensitive data for all connections, as well as passwords and private keys in plain text format.

3. Click Commit.
4. *Optional:* To set a different verbosity level for sessions that belong to a specific Connection Policy, complete the following steps:
   - Navigate to the Connection Policy you want to modify.
   - Select **Override global verbosity level**.
   - Select the desired log level from the **Connection specific verbosity level** field. Note that the new verbosity level applies only to new sessions started after committing the change. The verbosity level of active sessions will not change.
   - Click Commit.

Collecting logs and system information for error reporting

To track down support requests, the One Identity Support Team might request you to collect system-state and debugging information. This information is collected automatically, and contains log files, the configuration file of One Identity Safeguard for Privileged Sessions (SPS), and various system-statistics.

NOTE: Sensitive data like key files and passwords are automatically removed from the files, that is, configuration files do not contain passwords or keys. However, if you increase the proxy verbosity level to 8-10 in the Global Options, then for troubleshooting purposes, the logs can contain highly sensitive data, for example, passwords and keys in plain text format. If you are concerned about the presence of sensitive data, check the collected log files before submitting to the Support Portal.

The **Basic Settings** > **Management** > **Verbose system logs** > **Enable** option is not related to the verbosity of log messages: it adds the commands executed by the SPS web interface to the log.

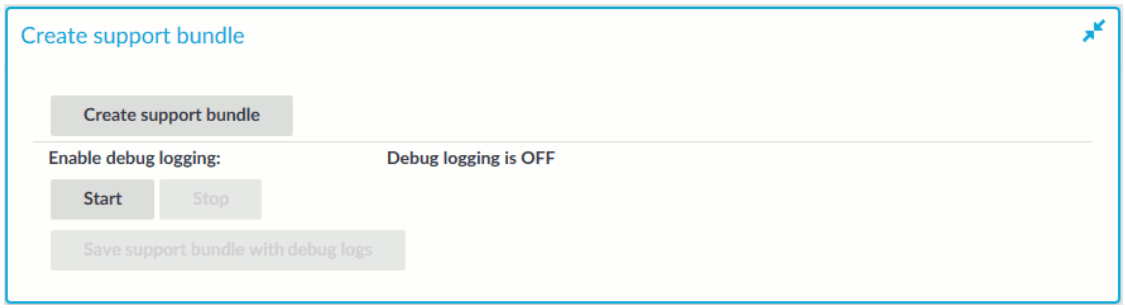
To collect system-state information (also known as a support bundle), navigate to **Basic Settings** > **Troubleshooting** > **Create support bundle** and click **Create support bundle**, then save the created zip file. The name of the file uses the `debug_info-<hostname>YYYYMMDDHHMM` format.

To collect information for a specific error

If the problem you want to reproduce requires a reboot, see Collecting logs and system information of the boot process for error reporting instead.

1. Navigate to **Basic Settings** > **Troubleshooting** > **Create support bundle**.

   Figure 334: Basic Settings > Troubleshooting > Create support bundle — Collecting debug information

2. Click **Start**.
3. (Optional) If the error you want to reproduce is related to the audited network traffic, for example, it occurs when a used connects to a protected SSH server, consider increasing the log level of the related traffic (for example, SSH). For details, see Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS).

Note that earlier versions of SPS increased the log level of the audited traffic automatically.

4. Reproduce the event that causes the error, for example connect to a server.
5. Click **Stop**.
6. Click **Save support bundle with debug logs** and save the created zip file. The name of the file uses the `debug_info-<hostname>YYYYMMDDHHMM` format.

   SPS includes the configuration files of any plugins installed. Note that depending on the plugin, these configuration files can contain sensitive information, such as passwords or API keys. In this case, edit the plugin-related files in the `plugins` directory of the support bundle and delete the sensitive information.
7. (Optional) If you have increased the verbosity level of the audited network traffic, decrease it to the default level.
8. Attach the file to your support ticket.

Collecting logs and system information of the boot process for error reporting

If you have a problem related to the boot process of the appliance, the One Identity Support Team might request you to collect system-state and debugging information about the boot process. This information is collected automatically, and contains log files, the configuration file of One Identity Safeguard for Privileged Sessions (SPS), and various system-statistics.

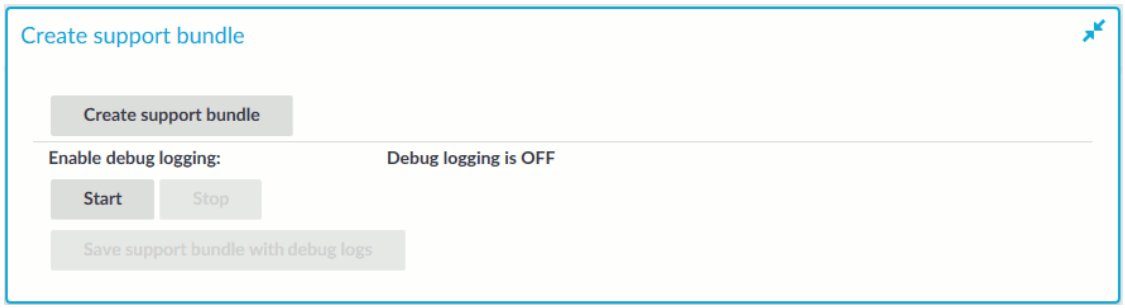
If the problem you want to reproduce does not require a reboot, see Collecting logs and system information for error reporting instead.

To collect information about the boot process

1. Navigate to **Basic Settings** > **Management** > **Verbose system logs** > **Enable**.
2. (Optional) If the error you want to reproduce is related to the audited network traffic, for example, it occurs when a used connects to a protected SSH server, consider increasing the log level of the related traffic (for example, SSH). For details, see Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS).

   Note that earlier versions of SPS increased the log level of the audited traffic automatically.
3. Reproduce the event that causes the error (for example, reboot the appliance).
4. Navigate to **Basic Settings** > **Troubleshooting** > **Create support bundle** and click **Create support bundle**.

Figure 335: Basic Settings > Troubleshooting > Create support bundle — Collecting debug information



5. Click **Save support bundle with debug logs** and save the created zip file. The name of the file uses the `debug_info-<hostname>YYYYMMDDHHMM` format.

   SPS includes the configuration files of any plugins installed. Note that depending on the plugin, these configuration files can contain sensitive information, such as passwords or API keys. In this case, edit the plugin-related files in the `plugins` directory of the support bundle and delete the sensitive information.
6. (Optional) If you have increased the verbosity level of the audited network traffic, decrease it to the default level.
7. Navigate to **Basic Settings** > **Management** > **Verbose system logs** > **Disable**.
8. Attach the file to your support ticket.

Support hotfixes

This section describes support hotfixes and their installation in One Identity Safeguard for Privileged Sessions (SPS).

Support hotfixes are official additions (signed .deb packages created by the Support Team) to a specific SPS release. By uploading a hotfix to an SPS appliance, it is possible to apply a modification (for example, a bugfix) quickly and without making the firmware Tainted. The hotfix files only work with the version of SPS they are created for.

You can upload the hotfix file you received from our Support Team in the SPS user interface.

CAUTION:

Consider the following:

- Clicking **Upload** immediately installs the hotfix to SPS.
- Installing multiple hotfix files to a single configuration of SPS is possible, but you cannot delete an individual hotfix file from SPS without the Support Team's assistance.
- Installing a new firmware will delete all hotfix files installed on the previous version of SPS.

If you have to delete an individual hotfix file from SPS without installing a new firmware first, contact our Support Team.

Installing support hotfixes

This section describes the most important requirements and information regarding the installation procedure of support hotfixes.

Prerequisites

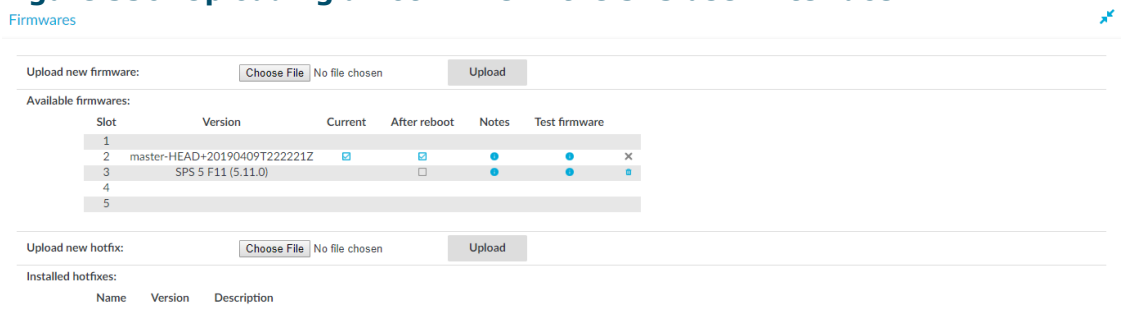
The hotfix files are normally not publicly accessible for download (unless attached to Knowledgebase Articles). As a result, if you want to install them to your SPS, you must first contact our Support Team for a hotfix file specifically created for your request. Consider that you cannot delete the installed hotfix file from SPS without the Support Team's assistance. In addition, rebooting the SPS appliance after deleting an installed hotfix is necessary. We strongly recommend that you only install hotfixes to SPS if you contact our Support Team for instructions beforehand.

NOTE: The hotfix files only work with the version of SPS they are created for. SPS automatically checks their version during upload.

To install the support hotfix file

1. Navigate to **Basic Settings** > **System** > **Firmwares**.

Figure 336: Uploading a hotfix file in the SPS user interface



2. Under the **Upload new hotfix:** section, click **Choose File** and select the hotfix file you want to upload.
3. Click **Upload**.

CAUTION:

Consider the following:

- Clicking **Upload** immediately installs the hotfix to SPS.
- Installing multiple hotfix files to a single configuration of SPS is possible, but you cannot delete an individual hotfix file from SPS without the Support Team's assistance.
- Installing a new firmware will delete all hotfix files installed on the previous version of SPS.

If you have to delete an individual hotfix file from SPS without installing a new firmware first, contact our Support Team.

4. If installation is successful, SPS will list information about the hotfix under **Installed hotfixes:**, such as **Name**, **Version** and **Description**.

   NOTE:

   Upload will fail in the following cases:

   - The hotfix file version does not pass the version check.
   - The hotfix file package is not properly signed by our Support Team.
   - The file you want to upload is not an appropriate .deb package or the file is corrupted.

   If upload fails, SPS will revert to its previous state automatically.

Status history and statistics

SPS displays various statistics and status history of system data and performance on the dashboard at **Basic Settings** > **Dashboard**. The dashboard is essentially an extension of the system monitor: the system monitor displays only the current values, while the dashboard creates graphs and statistics of the system parameters.

The dashboard consists of different modules. Every module displays the history of a system parameter for the current day. To display the graph for a longer period (last week, last month, or last year), select the **Week**, **Month**, or **Year** options, respectively. Hovering the mouse over a module enlarges the graph and displays the color code used on the graph.

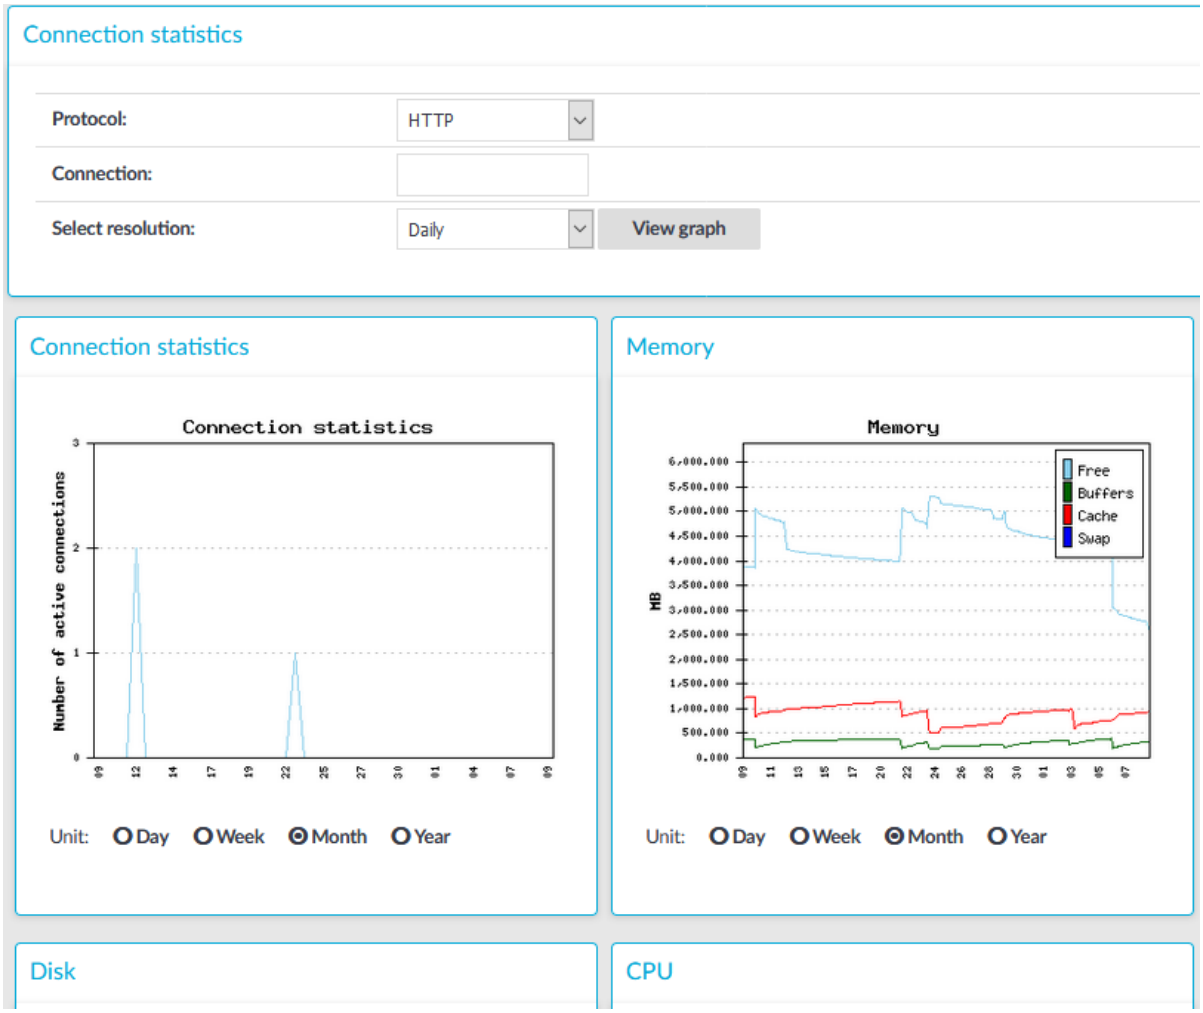
All types of data is collected every five minutes. This means that if changes are more frequent, it might not be represented in the graphs.

NOTE: If all parameters displayed are 0 at a certain point in time, it might mean that at that time One Identity Safeguard for Privileged Sessions (SPS) was not functional (for example, turned off or unresponsive). Or, in certain cases it might also mean that there was no information at that time.

NOTE: If you want to compare data displayed on the Dashboard to data displayed on the System Monitor, they might be different, because data on System Monitor is based on SNMP values, whereas data on the related Dashboard modules are based on the output of different commands.

To display statistics of a module as a table for the selected period, click on the graph.

Figure 337: Basic Settings > Dashboard — The dashboard

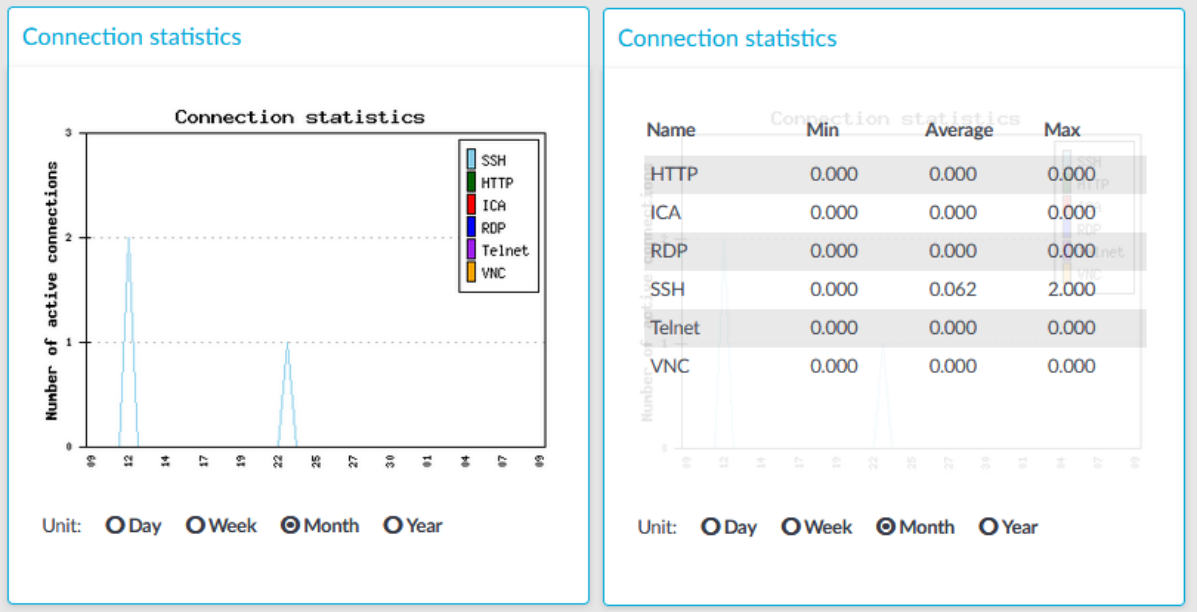


The following modules are displayed on the dashboard of SPS:

- **Connection statistics**: Number of active connections per protocol.
- **Memory**: The memory used by the system.
- **Disk**: Filesystem usage for the different partitions.
- **CPU**: CPU usage.
- **Network connections**: Number of network connections.
- **Physical interface 1 (eth0)**: Traffic on physical interface 1.
- **Physical interface 2 (eth1)**: Traffic on physical interface 2.
- **Physical interface 3 (eth2)**: Traffic on physical interface 3.
- **Load average**: Average load of the system.
- **Number of processes**: The number of running processes.

Connection statistics

Figure 338: Basic Settings > Dashboard > Connection statistics



The **Connection statistics** module on the **Dashboard** is based on statistics of high-level proxy-service protocols (SSH, RDP, VNC, ICA, and so on). These numbers display all active high-level proxy-service protocols, but these numbers are counted by all service connections too, which are connected to some protocols. Because of this, these numbers can differ from the numbers displayed on the **Pending Connections** > **Active Connections** page.

For example, if there are several active ICA connections in your system, it means that there are approximately the same number of CGP connections that are opened and counted in the **Connection statistics** module under the ICA label. If these CGP or ICA high-level proxy-service protocols are opening more than one TCP connections, these connections will be counted in the **Network connetion** module as different TCP connections, but these will count as only one connection on the **Pending Connections** > **Active Connections** page.

Statistics

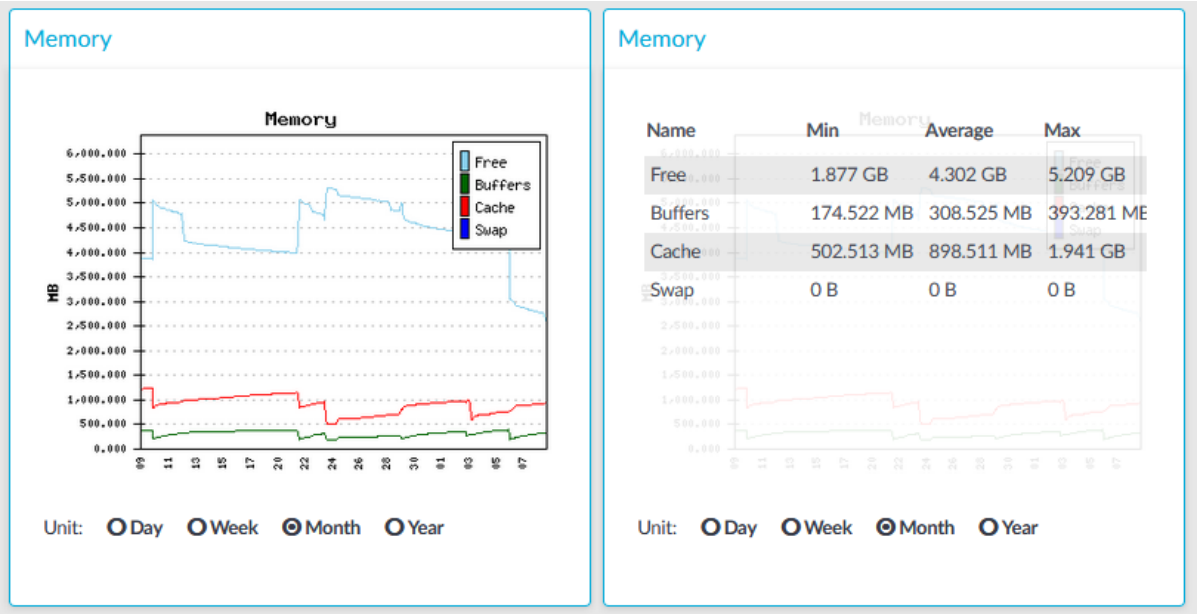
The connection types displayed can be the following:

- **RDP**: The number of RDP connections.
- **SSH**:The number of SSH connections.
- **HTTP**: The number of HTTP connections.
- **ICA**: The number of Citrix connections.
- **Telnet**: The number of Telnet connections.
- **VNC**: The number of VNC connections.

The **Min**, **Average** and **Max** values are displayed as a whole number if the value is constant for the statistics interval (the statistics are stored every 5 minutes). If minor changes occur in the actual values (for example, new connections are established), these changes can be displayed as fractions.

Memory

Figure 339: Basic Settings > Dashboard > Memory



The **Memory** module on the **Dashboard** is based on data provided by the Linux kernel (`/proc` and `/sys` directories). The standard Munin plugins query this information from these locations and they are displayed on the GUI.

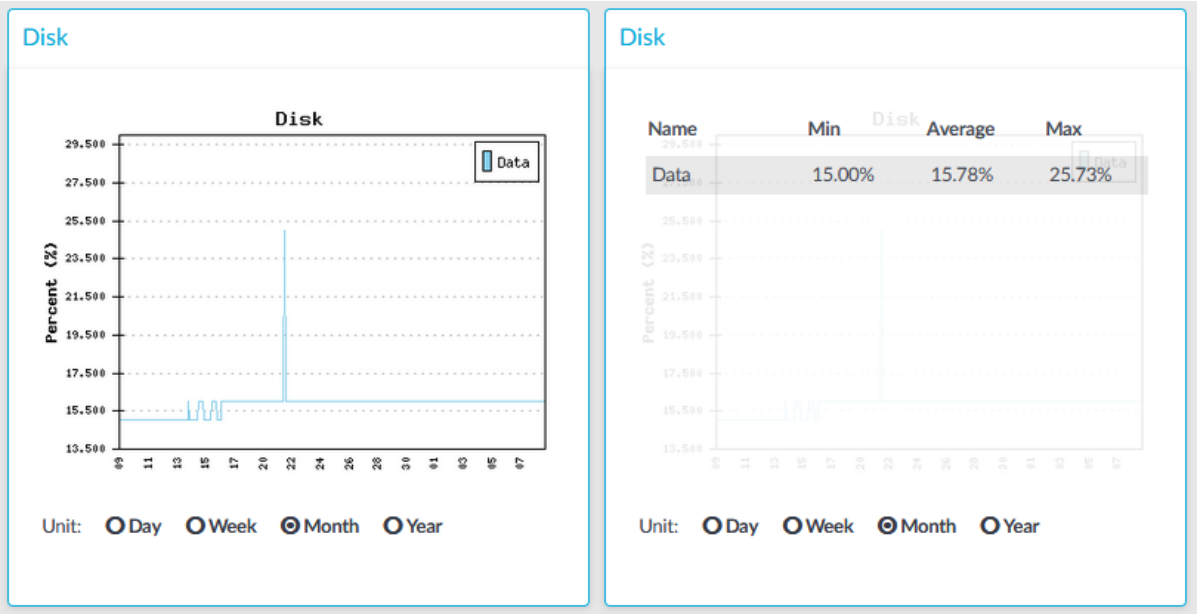
Statistics

The memory types displayed are the following:

- **Free**: Free memory
- **Buffers**: In-memory block I/O buffers.
- **Cache**: Memory used for disk caching. This does not count as "used" memory, because it is freed when it is required.
- **Swap**: Swap space usage (memory contents that have been temporarily moved to disk).This value might be high in case of lack of memory.

Disk

Figure 340: Basic Settings > Dashboard > Disk



The **Disk** module on the **Dashboard** is based on the output of the `df` command.

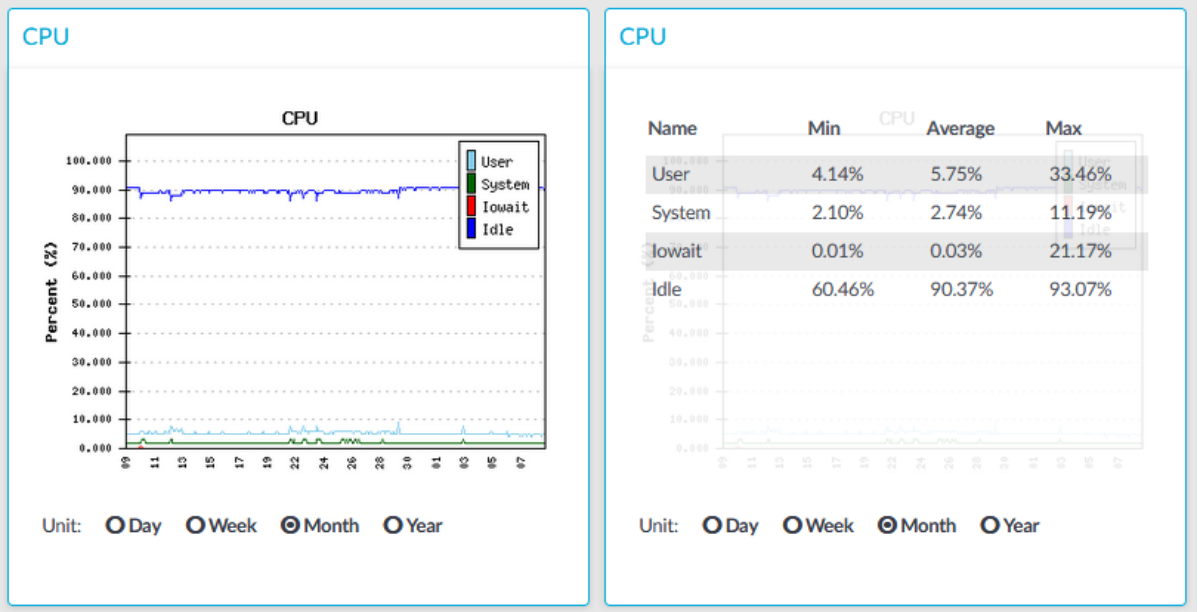
Statistics

The information displayed is the following:

- **Data**: The percent of disk that the core firmware uses.

CPU

Figure 341: Basic Settings > Dashboard > CPU



The **CPU** module on the **Dashboard** is based on data provided by the Linux kernel (`/proc` and `/sys` directories). The standard Munin plugins query this information from these locations and they are displayed on the GUI.

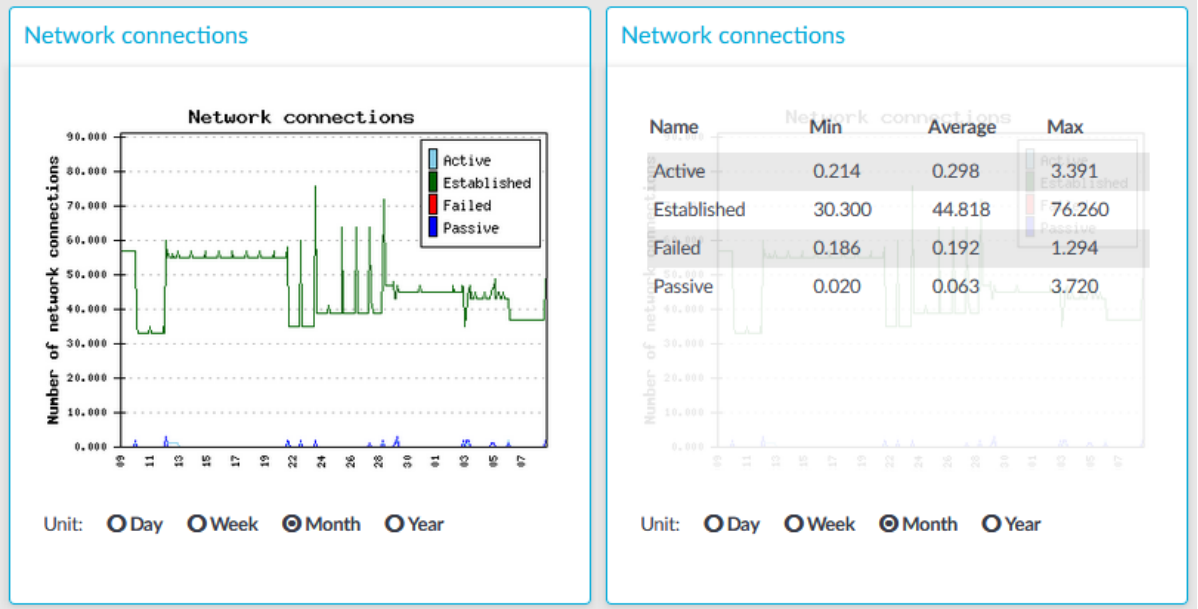
Statistics

The following details are displayed about CPU usage:

- **Idle**: Idle time of the processors. If there are more than one processors, they all add up to x100%, for example in case of 2 processors it adds up to 200% maximum.
- **Iowait**: Time spent receiving and handling hardware interrupts as a percentage of processor ticks. That is, waiting for IO.
- **System**: Kernel CPU usage.
- **User**: CPU usage of everything other than kernel.

Network connections

Figure 342: Basic Settings > Dashboard > Network connections



The **Network connetion** module on the **Dashboard** is based on the output of the `netstat -s` command. This command generates statistical information from all interfaces of all TCP connections. This means that in addition to the high-level proxy-service protocols (SSH, RDP, VNC, ICA, and so on), but all types of TCP connections are counted as well. The standard Munin plugins query this information and then it is displayed on the GUI. The graph itself displays the TCP activity of all network interfaces combined.

Statistics

The connection types displayed can be the following:

- **Active**: The number of active TCP openings per second.
- **Established**: The number of currently open connections.
- **Failed**: The number of failed TCP connection attempts per second.
- **Passive**: The number of passive TCP openings per second.
- **Resets**: The number of TCP connection resets.

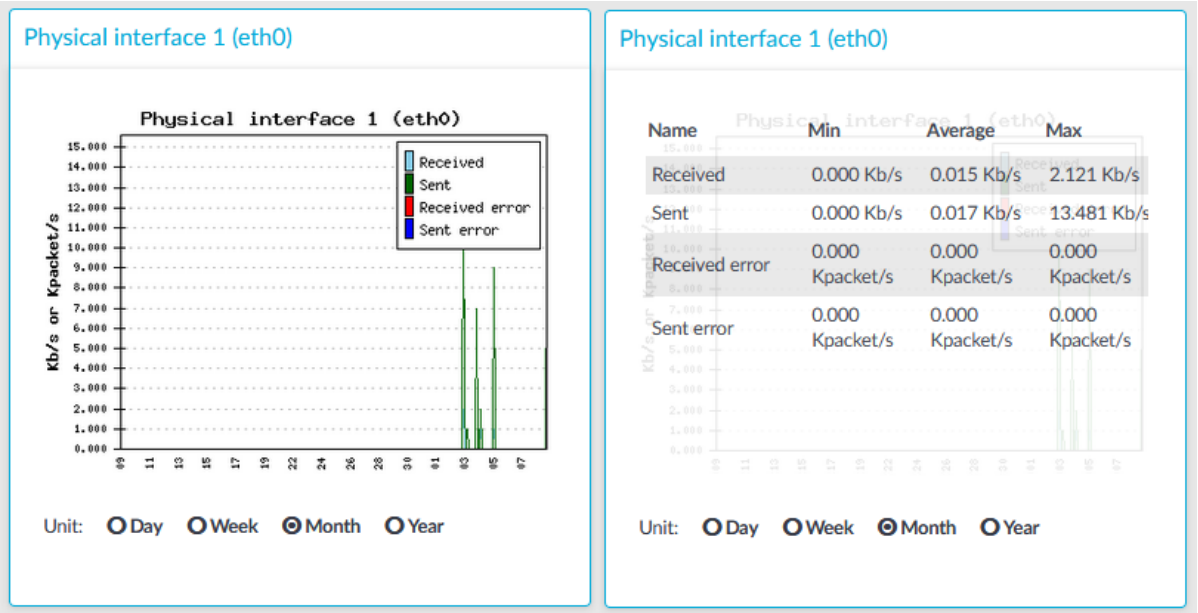
The **Min**, **Average** and **Max** values are displayed as a whole number if the value is constant for the statistics interval (the statistics are stored every 5 minutes). If minor changes occur in the actual values (for example, new connections are established), these changes can be displayed as fractions.

To determine the maximum values that the system can handle, consider the following:

- The type of machine that you run One Identity Safeguard for Privileged Sessions (SPS) on.
- The type of connections that are established and the content of these connections. For example:
    - If the users of RDP or ICA connections are watching videos, that can greatly reduce the amount of parallel connections that can pass through without experiencing speed reduction.
    - If the users mostly generate text-based content (for example, Excel, Word), then more connections can be used.
    - If a connection is not used actively, then it has minimal impact on SPS: only the memory allocation remains. In case of RDP, if the RDP client window is minimized, there is no network traffic at all.

Interface

Figure 343: Basic Settings > Dashboard > Interface



The **Interface** module on the **Dashboard** is based on data provided by the Linux kernel (`/proc` and `/sys` directories). The standard Munin plugins query this information from these locations and they are displayed on the GUI.

Statistics

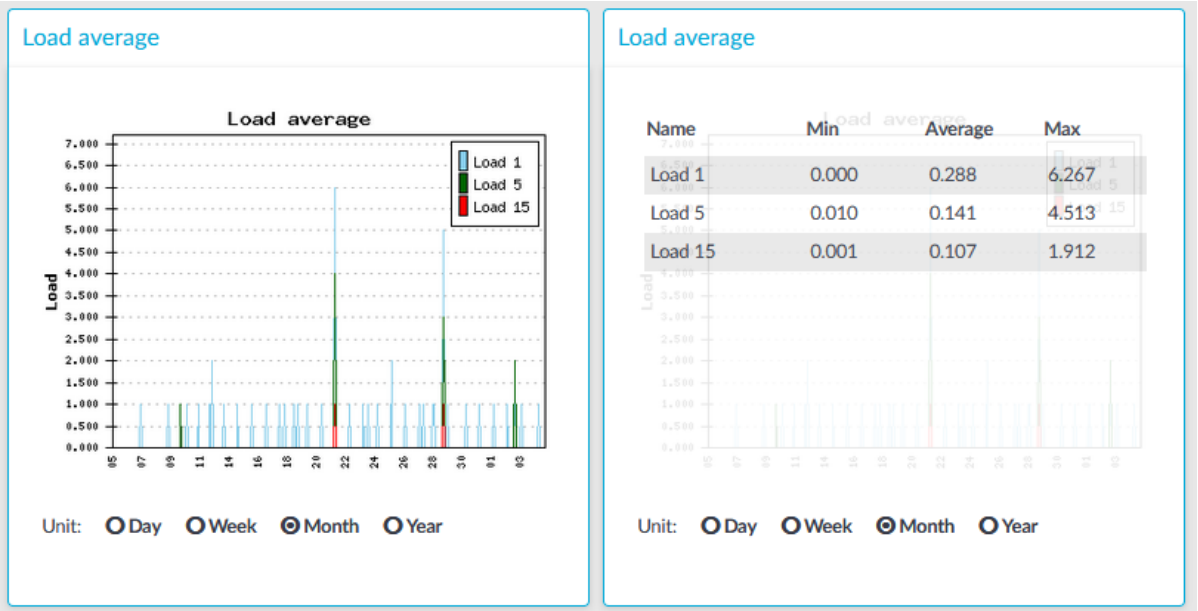
The memory types displayed are the following:

- **Received**: The network interface has received x Kilobytes per second.
- **Sent**: The network interface has sent x Kilobytes per second.

- **Received error**: The amount of errors, packet drops, and collisions on the network interface (in Kilopackets per second).
- **Sent error**: The amount of errors, packet drops, and collisions on the network interface (in Kilopackets per second).

Load average

Figure 344: Basic Settings > Dashboard > Load average



The **Load average** module on the **Dashboard** is based on data provided by the Linux kernel (`/proc` and `/sys` directories). The standard Munin plugins query this information from these locations and they are displayed on the GUI.

A measure of the amount of computational work that a computer system performs. The load average represents the average system load over a period of time. It conventionally appears in the form of three numbers which represent the system load during the last one-, five-, and fifteen-minute periods.

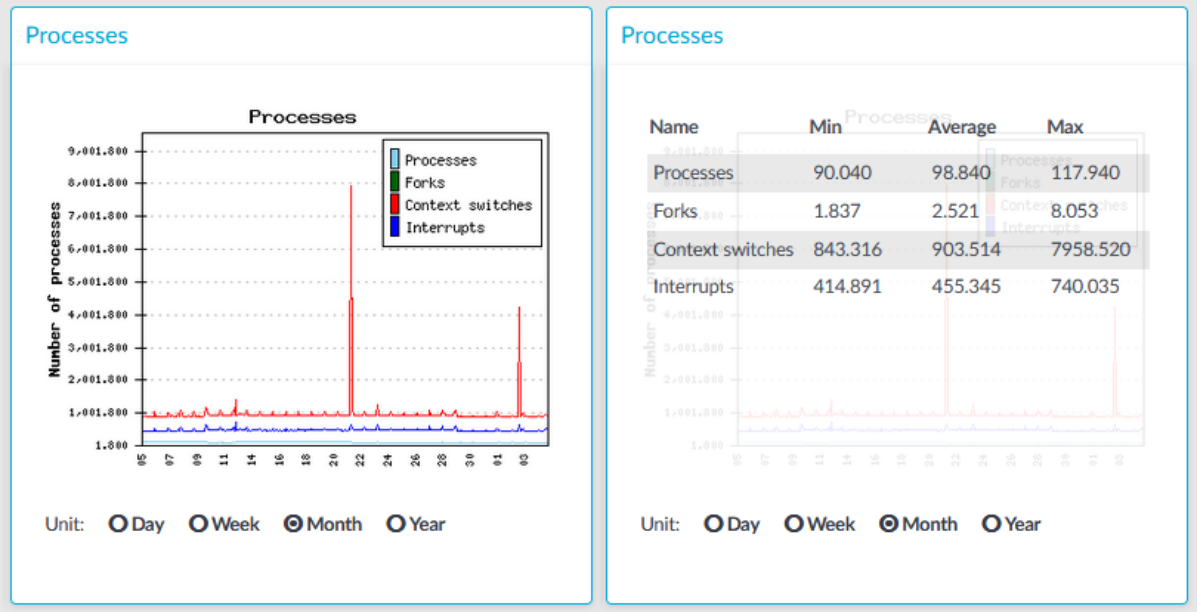
Statistics

The load average types displayed are the following:

- **Load 1**: Average load in 1 minute.
- **Load 5**: Average load in 5 minute.
- **Load 15**: Average load in 15 minute.

Number of processes

Figure 345: Basic Settings > Dashboard > Number of processes



The **Number of processes** module on the **Dashboard** is based on the output of the `ps` command.

Statistics

The process types displayed are the following:

- **Processes**: Number of running processes.
- **Forks**: Number of forks (system calls). It is an operation where a process creates a copy of itself.
- **Context switches**: Number of context switches. It is the switching of the CPU from one processor thread to another.
- **Interrupts**: Number of interrupts.

Displaying custom connection statistics

The following describes how to display statistics of a specific connection policy.

To display statistics of a specific connection policy

1. Navigate to **Basic Settings** > **Dashboard** > **Connection statistics**.
2. To display the statistics of a connection policy, enter the name of the policy into the **Connection**.
3. Select the time period to display from the **Select resolution** field.
4. Click **View graph**.

Troubleshooting a One Identity Safeguard for Privileged Sessions (SPS) cluster

The following sections help you to solve problems related to High Availability clusters.

- For a description of the possible statuses of the One Identity Safeguard for Privileged Sessions (SPS) cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured), see Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses on page 954.
- To recover a cluster that has broken down, see Recovering One Identity Safeguard for Privileged Sessions (SPS) if both nodes broke down on page 957.
- To resolve a split-bran situation when the nodes of the cluster were simultaneously active for a time, see Recovering from a split brain situation on page 957.
- To replace a broken node with a new appliance, see Replacing a HA node in a One Identity Safeguard for Privileged Sessions (SPS) cluster on page 960.

Understanding One Identity Safeguard for Privileged Sessions (SPS) cluster statuses

This section explains the possible statuses of the One Identity Safeguard for Privileged Sessions (SPS) cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured). SPS displays this information on the **Basic Settings** > **High Availability** page.

The **Status** field indicates whether the SPS nodes recognize each other properly and whether those are configured to operate in High Availability mode. The status of the individual SPS nodes is indicated in the **Node HA state** field of the each node. The following statuses can occur:

- **Standalone**: There is only one SPS unit running in `standalone` mode, or the units have not been converted to a cluster (the **Node HA state** of both nodes is `standalone`). Click **Convert to Cluster** to enable High Availability mode.

- **HA**: The two SPS nodes are running in High Availability mode. **Node HA state** is `HA` on both nodes, and the **Node HA UUID** is the same on both nodes.
- **Half**: High Availability mode is not configured properly, one node is in `standalone`, the other one in `HA` mode. Connect to the node in `HA` mode, and click **Join HA** to enable High Availability mode.
- **Broken**: The two SPS nodes are running in High Availability mode. **Node HA state** is `HA` on both nodes, but the **Node HA UUID** is different. For assistance, contact our Support Team.
- **Degraded**: SPS was running in High Availability mode, but one of the nodes has disappeared (for example broken down, or removed from the network). Power on, reconnect, or repair the missing node.
- **Degraded (Disk Failure)**: A hard disk of the secondary node is not functioning properly and must be replaced. To request a replacement hard disk and for details on replacing the hard disk, contact our Support Team.
- **Degraded Sync**: Two SPS units were joined to High Availability mode, and the first-time synchronization of the disks is currently in progress. Wait for the synchronization to complete. Note that in case of large disks with lots of stored data, synchronizing the disks can take several hours.
- **Split brain**: The two nodes lost the connection to each other, with the possibility of both nodes being active nodes (that is, primary nodes) for a time.

  CAUTION:

  Hazard of data loss In this case, valuable audit trails might be available on both SPS nodes, so special care must be taken to avoid data loss. For details on solving this problem, see Recovering from a split brain situation on page 957.

  Do NOT reboot or shut down the nodes.
- **Invalidated**: The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.
- **Converted**: After converting nodes to a cluster (clicking **Convert to Cluster**) or enabling High Availability mode (clicking **Join HA**) and before rebooting the node(s).

NOTE: If you experience problems because the nodes of the HA nodes do not find each other during system startup, navigate to **Basic Settings** > **High Availability** and select **HA (Fix current)**. That way the IP address of the HA interfaces of the nodes will be fix, which helps if the HA connection between the nodes is slow.

The **DRBD status** field indicates whether the latest data (including SPS configuration, audit trails, log files, and so on) is available on both SPS nodes. The primary node (this node) must always be in **consistent** status to prevent data loss. Inconsistent status means that the data on the node is not up-to-date, and should be synchronized from the node having the latest data.

The **DRBD status** field also indicates the connection between the disk system of the SPS nodes. The following statuses are possible:

- **Connected**: Both nodes are functioning properly.
- **Connected (Disk Failure)**: A hard disk of the secondary node is not functioning properly and must be replaced. To request a replacement hard disk and for details on replacing the hard disk, contact our Support Team.
- **Invalidated**: The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.
- **Sync source** or **Sync target**: One node (**Sync target**) is downloading data from the other node (**Sync source**).

  When synchronizing data, the progress and the remaining time is displayed in the **System monitor**.

  CAUTION:

  When the two nodes are synchronizing data, do not reboot or shutdown the primary node. If you absolutely must shutdown the primary node during synchronization, shutdown the secondary node first, and then the primary node.

- **Split brain**: The two nodes lost the connection to each other, with the possibility of both nodes being active nodes (that is, primary nodes) for a time.

  CAUTION:

  Hazard of data loss In this case, valuable audit trails might be available on both SPS nodes, so special care must be taken to avoid data loss. For details on solving this problem, see Recovering from a split brain situation on page 957.

- **WFConnection**: One node is waiting for the other node, the connection between the nodes has not been established yet.

If a redundant heartbeat interface is configured, its status is also displayed in the **Redundant Heartbeat status** field, and also in the **HA** > **Redundant** field of the System monitor. For a description of redundant heartbeat interfaces, see Redundant heartbeat interfaces on page 428.

The possible status messages are explained below.

- **NOT USED**: There are no redundant heartbeat interfaces configured.
- **OK**: Normal operation, every redundant heartbeat interface is working properly.
- **DEGRADED-WORKING**: Two or more redundant heartbeat interfaces are configured, and at least one of them is functioning properly. This status is displayed also when a new redundant heartbeat interface has been configured, but the nodes of the SPS cluster has not been restarted yet.
- **DEGRADED**: The connection between the redundant heartbeat interfaces has been lost. Investigate the problem to restore the connection.
- **INVALID**: An error occurred with the redundant heartbeat interfaces. Contact the One Identity Support Team for help. For assistance, contact our Support Team.

Recovering One Identity Safeguard for Privileged Sessions (SPS) if both nodes broke down

It can happen that both nodes break down simultaneously (for example because of a power failure), or the secondary node breaks down before the original primary node recovers.

NOTE: As of One Identity Safeguard for Privileged Sessions (SPS) version 2.0.2, when both nodes of a cluster boot up in parallel, the node with the `1.2.4.1` HA IP address will become the primary node.

To properly recover SPS

1. Power off both nodes by pressing and releasing the power button.

   CAUTION:

   Hazard of data loss If SPS does not shut off, press and hold the power button for approximately 4 seconds. This method terminates connections passing SPS and might result in data loss.

2. Power on the node that was the primary node before SPS broke down. Consult the system logs to find out which node was the primary node before the incident: when a node boots as primary node, or when a takeover occurs, SPS sends a log message identifying the primary node.

   TIP: Configure remote logging to send the log messages of SPS to a remote server where the messages are available even if the logs stored on SPS become unaccessible. For details on configuring remote logging, see System logging, SNMP and e-mail alerts on page 125.

3. Wait until this node finishes the boot process.
4. Power on the other node.

Recovering from a split brain situation

A split brain situation is caused by a temporary failure of the network link between the cluster nodes, resulting in both nodes switching to the active (that is, primary node) role while disconnected. This might cause new data (for example, audit trails) to be created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data have been created, which cannot be trivially merged.

CAUTION:

Hazard of data loss In a split brain situation, valuable audit trails might be available on both One Identity Safeguard for Privileged Sessions (SPS) nodes, so special care must be taken to avoid data loss.

The nodes of the SPS cluster automatically recognize the split brain situation once the connection between the nodes is reestablished, and do not perform any data synchronization to prevent data loss. When a split brain situation is detected, it is visible on

the SPS system monitor, in the system logs (`Split-Brain detected, dropping connection!`), on the **Basic Settings** > **High Availability** page, and SPS sends an alert as well.

Once the network connection between the nodes has been re-established, one of the nodes will become the active (that is, primary) node, while the other one will be the backup node (that is, the secondary node). This means that one node is providing services similar to normal operation, and the other one is kept passive (as a backup) to avoid network interferences. Note that there is no synchronization between the nodes at this stage.

To recover a SPS cluster from a split brain situation, complete the following steps.

CAUTION:

Do NOT shut down the nodes.

Data recovery

In the procedure described here, data will be saved from the host currently acting as the secondary node host. This is required because data on this host will later be overwritten by the data available on the current primary node.

NOTE: During data recovery, there will be no service provided by SPS.

To recover from a split brain situation

1. Log in to the primary node. If no Console menu is showing up after login, then this is the secondary node. In this case, try the other node.
2. Select **Shells** > **Boot Shell**.
3. Enter **`/usr/share/heartbeat/hb_standby`**. This will change the current secondary node to primary node and the current primary node to secondary node (HA failover).
4. Exit the console.
5. Wait a few seconds for the HA failover to complete.
6. Log in on the other host. If no Console menu is showing up, the HA failover has not completed yet. Wait a few seconds and try logging in again.
7. Select **Shells** > **Core Shell**.
8. Issue the `systemctl stop zorp-core.service` command to disable all traffic going through SPS.
9. Save the files from `/var/lib/zorp/audit` that you want to keep. Use `scp` or `rsync` to copy data to your remote host.

   TIP: To find the files modified in the last *n*24* hours, use `find . -mtime -n`.

   To find the files modified in the last *n* minutes, use `find . -mmin -n` .
10. Enter:

    ```
    pg_dump -U scb -f /root/database.sql
    ```

    Back up the `/root/database.sql` file.

11. Exit the console.
12. Log in again, and select **Shells** > **Boot Shell**.
13. Enter `/usr/share/heartbeat/hb_standby`. This will change the current secondary node to primary node and the current primary node to secondary node (HA failover).
14. Exit the console.
15. Wait a few minutes to let the failover happen, so the node you were using will become the secondary node and the other node will become the primary node.

    The nodes are still in a split-brain state but now you have all the data backed up from the secondary node, and you can synchronize the data from the primary node to the secondary node, which will turn the HA state from "Split-brain" to "HA". For details on how to do that, see HA state recovery on page 959.

HA state recovery

In the procedure described here, the "Split-brain" state will be turned to the "HA" state. Keep in mind that the data on the current primary node will be copied to the current secondary node and *data that is available only on the secondary node will be lost* (as that data will be overwritten).

Steps: Swapping the nodes (optional)

NOTE: If you completed the procedure described in Data recovery on page 958, you do not have to swap the nodes. You can proceed to the steps about data synchronization.

If you want to swap the two nodes to make the primary node the secondary node and the secondary node the primary node, perform the following steps:

1. Log in to the primary node. If no Console menu is showing up after login, then this is the secondary node. In this case, try the other node.
2. Select **Shells** > **Boot Shell**.
3. Enter `/usr/share/heartbeat/hb_standby`. This will output:

   ```
   Going standby [all]
   ```

4. Exit the console.
5. Wait a few minutes to let the failover happen, so the node you were using will become the secondary node and the other node will be the primary node.

Steps: Initializing data synchronization

To initialize data synchronization, complete the following steps:

1. Log in to the secondary node. If the Console menu is showing up, then this is the primary node. In this case, try logging in to the other node.
2. Enter the following commands. These commands will make the secondary node discard the data available only here, on this node.

```
drbdadm secondary r0
drbdadm connect --discard-my-data r0
```

3. Log out of the secondary node.
4. Log in to the primary node.
5. Select **Shells** > **Boot Shell**.
6. Enter:

```
drbdadm connect r0
```

7. Exit the console.
8. Check the High Availability state on the web interface of SPS, in the **Basic Settings** > **High Availability** > **Status** field. During synchronization, the status will say **Degraded Sync**, and after the synchronization completes, it will say **HA**.

Replacing a HA node in a One Identity Safeguard for Privileged Sessions (SPS) cluster

CAUTION:

Creating a High-availability (HA) node pair from different types of hardware is not possible. The primary and the secondary HA nodes have to run on the same type of hardware.

The following describes how to replace a unit in a One Identity Safeguard for Privileged Sessions (SPS) cluster with a new appliance.

To replace a unit in a SPS cluster with a new appliance

1. Verify the HA status on the working node. Select **Basic Settings** > **High Availability**. If one of the nodes has broken down or is missing, the **Status** field displays `DEGRADED`.
2. Note down the **Gateway IP** addresses, and the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces.
3. Perform a full system backup. Before replacing the node, create a complete system backup of the working node. For details, see Data and configuration backups on page 145.
4. Check which firmware version is running on the working node. Select **Basic Settings** > **System** > **Version details** and write down the exact version numbers.
5. Log in to your support portal and download the CD ISO for the same SPS version that is running on your working node.

6. Without connecting the replacement unit to the network, install the replacement unit from the ISO file. Use the IPMI if needed.
7. When the installation is finished, connect the two SPS units with an Ethernet cable via the Ethernet connectors labeled as 4 or HA.
8. Reboot the replacement unit and wait until it finishes booting.
9. Login to the working node and verify the HA state. Select **Basic Settings** > **High Availability**. The **Status** field should display HALF.
10. Reconfigure the **Gateway IP** addresses, and the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces. Click 

.
11. Click **Other node** > **Join HA**.
12. Click **Other node** > **Reboot**.
13. The replacement unit will reboot and start synchronizing data from the working node. The **Basic Settings** > **High Availability** > **Status** field will display DEGRADED SYNC until the synchronization finishes. Depending on the size of the hard disks and the amount of data stored, this can take several hours.
14. After the synchronization is finished, connect the other Ethernet cables to their respective interfaces (external to 1 or EXT, internal to 3 or INT, management to 2 or MGMT) as needed for your environment.

Expected result

A node of the SPS cluster is replaced with a new appliance.

Resolving an IP conflict between cluster nodes

The IP addresses of the HA interfaces connecting the two nodes are detected automatically, during boot. When a node comes online, it attempts to connect to the IP address 1.2.4.1. If no other node responds until timeout, then it sets the IP address of its HA interface to 1.2.4.1, otherwise (if there is a responding node on 1.2.4.1) it sets its own HA interface to 1.2.4.2.

Replaced nodes do not yet know the HA configuration (or any other HA settings), and will attempt to negotiate it automatically in the same way. If the network is, for any reason, too slow to connect the nodes on time, the replacement node boots with the IP address of 1.2.4.1, which can cause an IP conflict if the other node has also set its IP to that same address previously. In this case, the replacement node cannot join the HA nodes.

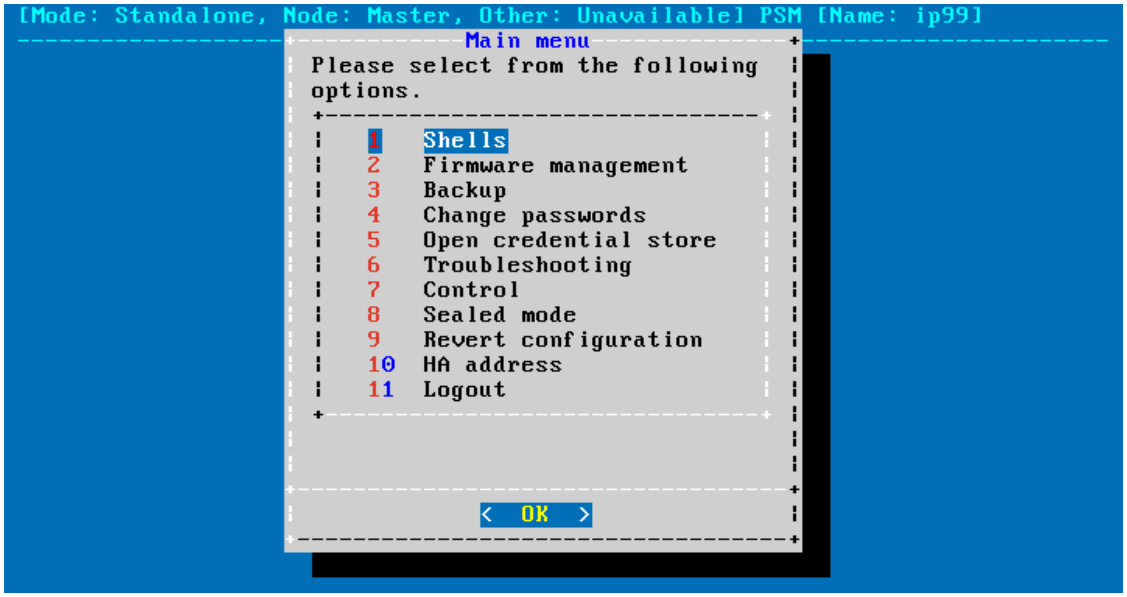
To manually assign the correct IP address to the HA interface of a node, perform the following steps:

1. Log in to the node using the IPMI or the physical console.

   Configuration changes have not been synced to the new (replacement) node, as it could not join the HA nodes. Use the default password of the root user of One Identity Safeguard for Privileged Sessions (SPS), see *Installing this hotfix* in the *Installation Guide*.

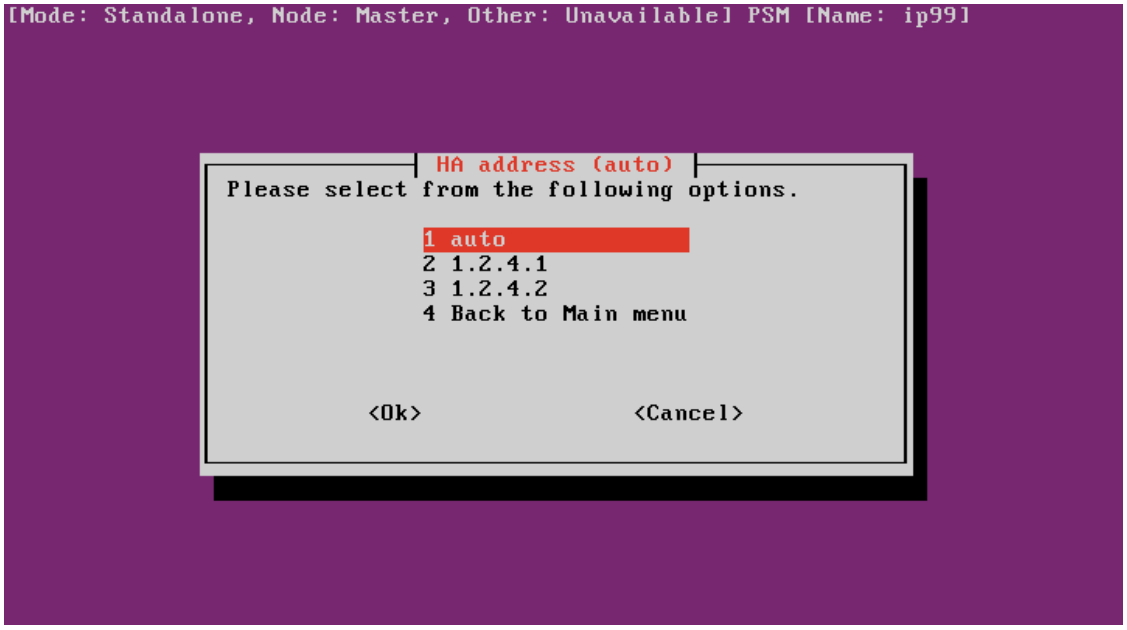
2. From the console menu, choose **10 HA address**.

Figure 346: The console menu



3. Choose the IP address of the node.

Figure 347: The console menu



4. Reboot the node.

Understanding One Identity Safeguard for Privileged Sessions (SPS) RAID status

This section explains the possible statuses of the One Identity Safeguard for Privileged Sessions (SPS) RAID device and the underlying hard disks. SPS displays this information on the **Basic Settings** > **High Availability** page. The following statuses can occur:

- **Optimal**: The hard disks are working as expected.
- **Degraded**: One or more hard disk has reported an error, and might have to be replaced. For assistance, contact our Support Team.
- **Failed stripes**: One or more stripes of data failed on the RAID device. It is possible that data loss occurred, but unfortunately there is no way to find out the extent of the data loss (if any).
    - *If you have a single SPS node*: You must reinstall SPS and restore the data from the latest backup. For details, see *One Identity Safeguard for Privileged Sessions Software Installation Guide* in the *Installation Guide* and Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance on page 967. If you do not have backup, contact our Support Team.

- *If you have a high-availability SPS cluster*: Shut the node down. Do NOT disconnect its HA interface. Reinstall the node (for details, see One Identity Safeguard for Privileged Sessions *Software Installation Guide*), power it on, then navigate to **Basic Settings** > **High Availability**, and click **Join HA**. For assistance, contact our Support Team.

- **Offline**: The RAID device is not functioning, probably because several disks have broken down. SPS cannot operate properly in this case. For assistance, contact our Support Team.

Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data

This section describes the options for restoring the configuration and the data of One Identity Safeguard for Privileged Sessions (SPS).

CAUTION: Do not use the `sps-<timestamp>.config` configuration file to clone SPS appliances. Doing so will result in errors and it might compromise the security of the appliances that are cloned this way.

You have the following options to restore SPS:

- Restore only the configuration.

  If the configuration file got corrupted, or it is not set properly and you want to use a previous version of the configuration that functioned correctly, restore only the SPS configuration.

  For more information, see Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration.

- Restore the configuration and the data.

  If you must completely restore the SPS appliance, restore both the configuration and the data.

  The possible use cases, when you must restore the configuration and the data are, for example, the following:

  - Disaster recovery
  - Planned hardware replacement

  The restore procedure differs if you perform the restore on the same, or on a new SPS appliance.

  For more information, see sections Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance and Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to a new SPS appliance.

Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration

The following procedure describes how to restore the configuration of One Identity Safeguard for Privileged Sessions (SPS).

CAUTION: Do not use the `sps-<timestamp>.config` configuration file to clone SPS appliances. Doing so will result in errors and it might compromise the security of the appliances that are cloned this way.

To restore the configuration of SPS

1. Select the configuration file you want to use to restore the SPS configuration. You can use either a configuration backup, or an exported configuration:
    - To use a configuration backup, connect to your backup server and locate the directory where SPS saves the backups. The configuration backups are stored in the `config` subdirectory, in timestamped files. Find the latest configuration file, named `sps-<timestamp>.config`.
    - To use an exported configuration, you must first export the configuration. To do so, navigate to **Basic Settings** > **System** > **Export configuration**.
2. Connect to SPS.
3. **CAUTION:**To minimize the amount of audit data that may be lost, disable the audit traffic by selecting **Stop** for **All services** on **Basic Settings** > **System** > **Traffic control**.
4. Navigate to **Basic Settings** > **System** > **Import configuration** > **Browse**, select the configuration file, and click **Import**.
5. To enable the audit traffic, navigate to **Basic Settings** > **System** > **Traffic control** and select **Start** for **All services**.

Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to a new SPS appliance

The following procedure describes how to restore the configuration and data of One Identity Safeguard for Privileged Sessions (SPS) from a complete backup to a new SPS appliance.

CAUTION:

Consider the following before starting the restore procedure:

- To minimize the amount of audit data that might be lost, perform the restore procedure as fast as possible.
- Do the restore procedure on the same SPS version. Restoring from an older version to a newer version, or the other way round, is not supported. For help, contact our Support Team.
- Ensure that you have enough free space to restore.
- During the restore procedure, the REST-based search might not function properly, since the data to search might still be incomplete.

To restore the configuration and data of SPS from a complete backup to a new SPS appliance

1. Connect to your backup server and locate the directory where SPS saves the backups. The configuration backups are stored in the `config` subdirectory, in timestamped files. Find the latest configuration file, named `sps-<timestamp>.config`.
2. Connect to SPS.

   The Welcome Wizard is displayed.

   CAUTION: To minimize the amount of audit data that might be lost, proceed to the following step to stop the traffic on **All services** as fast as possible.
3. **CAUTION:**To minimize the amount of audit data that may be lost, disable the audit traffic by selecting **Stop** for **All services** on **Basic Settings** > **System** > **Traffic control**.
4. Navigate to **Policies** > **Backup & Archive**. Verify that the settings of the target servers and the backup protocols are correct.
5. Navigate to **Basic Settings** > **Management** > **System backup**, click **Restore now** and wait for the process to finish.

   Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this might take a long time.
6. Enable the audit traffic: navigate to **Basic Settings** > **System** > **Traffic control** and select **Start** for **All services**.
7. For all protocols (or at least for the ones used in your system), navigate to **<Protocol-name> Control** > **Connections**, and click **Restore ALL**.

   Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

Restoring One Identity Safeguard for Privileged Sessions (SPS) configuration and data to the same SPS appliance

The following procedure describes how to restore the configuration and data of One Identity Safeguard for Privileged Sessions (SPS) from a complete backup to the same SPS appliance.

CAUTION:

- To minimize the amount of audit data that may be lost, perform the restore procedure as fast as possible.
- Do the restore procedure on the same SPS version. Restoring from an older version to a newer version, or the other way round, is not supported. For help, contact our Support Team.
- Ensure that you have enough free space to restore.
- During the restore procedure, the REST-based search might not function properly, since the data to search in might still be incomplete.

To restore the configuration and data of SPS from a complete backup to the same SPS appliance

1. Connect to your backup server and locate the directory where SPS saves the backups. The configuration backups are stored in the `config` subdirectory, in timestamped files. Find the latest configuration file, named `sps-<timestamp>.config`.
2. Connect to SPS.
3. The procedure differs depending on whether you have completed the Welcome Wizard or not.
    - If you have not yet completed the Welcome Wizard:
        1. On the Welcome Wizard, select the configuration file and import it.

            CAUTION: To minimize the amount of audit data that might be lost, proceed to the following step to stop the traffic on **All services** as fast as possible.

            For more information, see Configuring One Identity Safeguard for Privileged Sessions (SPS) with the Welcome Wizard.
        2. **CAUTION:**To minimize the amount of audit data that may be lost, disable the audit traffic by selecting **Stop** for **All services** on **Basic Settings** > **System** > **Traffic control**.
    - If you have previously completed the Welcome Wizard:
        1. **CAUTION:**To minimize the amount of audit data that may be lost,

disable the audit traffic by selecting **Stop** for **All services** on **Basic Settings** > **System** > **Traffic control**.

2. Navigate to **Basic Settings** > **System** > **Import configuration** > **Browse**, select the configuration file, and click **Import**.

4. Navigate to **Policies** > **Backup & Archive**. Verify that the settings of the target servers and the backup protocols are correct.
5. Navigate to **Basic Settings** > **Management** > **System backup**, click **Restore now** and wait for the process to finish.

   Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this might take a long time.
6. Enable the audit traffic: navigate to **Basic Settings** > **System** > **Traffic control** and select **Start** for **All services**.
7. For all protocols (or at least for the ones used in your system), navigate to **<Protocol-name> Control** > **Connections**, and click **Restore ALL**.

   Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

VNC is not working with TLS

Some vendors may use custom protocol elements and TLS-encryption that do not have available documentation. As a result, these cannot be audited by One Identity Safeguard for Privileged Sessions (SPS). Regardless of vendors, only the custom features described in the RFC 6143 are supported. As for encryptions, only those completely TLS-encapsulated streams can be processed where the TLS encryption process was started before the VNC protocol handshake.

Configuring the IPMI from the BIOS after losing IPMI password

It may happen that you inadvertently lose the IPMI password of your One Identity Safeguard for Privileged Sessions (SPS). The following procedure describes how you can re-configure your SPS if you lose your IPMI password.

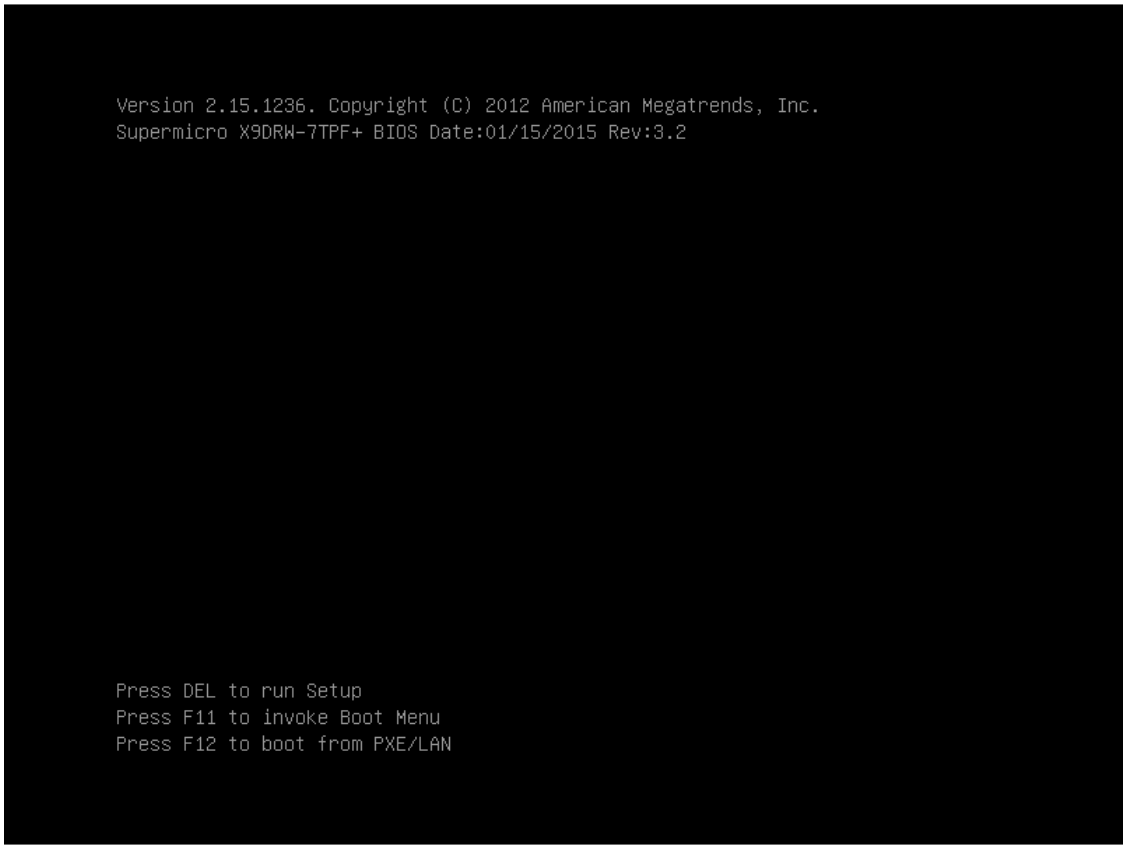
Prerequisites

To apply the procedure outlined here, you will need physical access to a monitor and keyboard.

To configure the IPMI from the BIOS after losing your IPMI password

1. Shut down SPS.
2. Unplug the SPS physical appliance's power cord.
3. Wait 30 seconds.
4. Replug the power cord.
5. Restart the appliance.
6. Press the DEL button when the POST screen comes up while the appliance is booting.

Figure 348: POST screen during booting



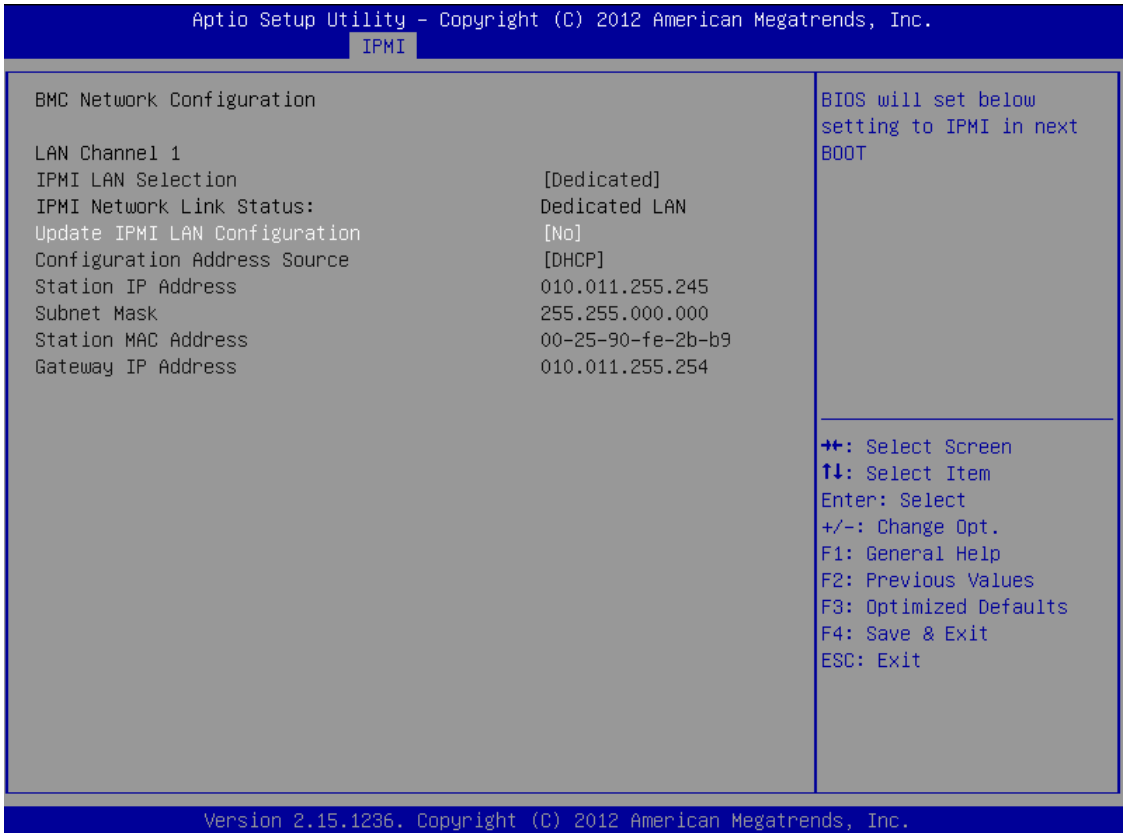
7. In the BIOS, navigate to the **IPMI** page.
8. On the **IPMI** page, select **BMC Network Configuration**, and press Enter.

Figure 349: IPMI page > BMC Network Configuration option



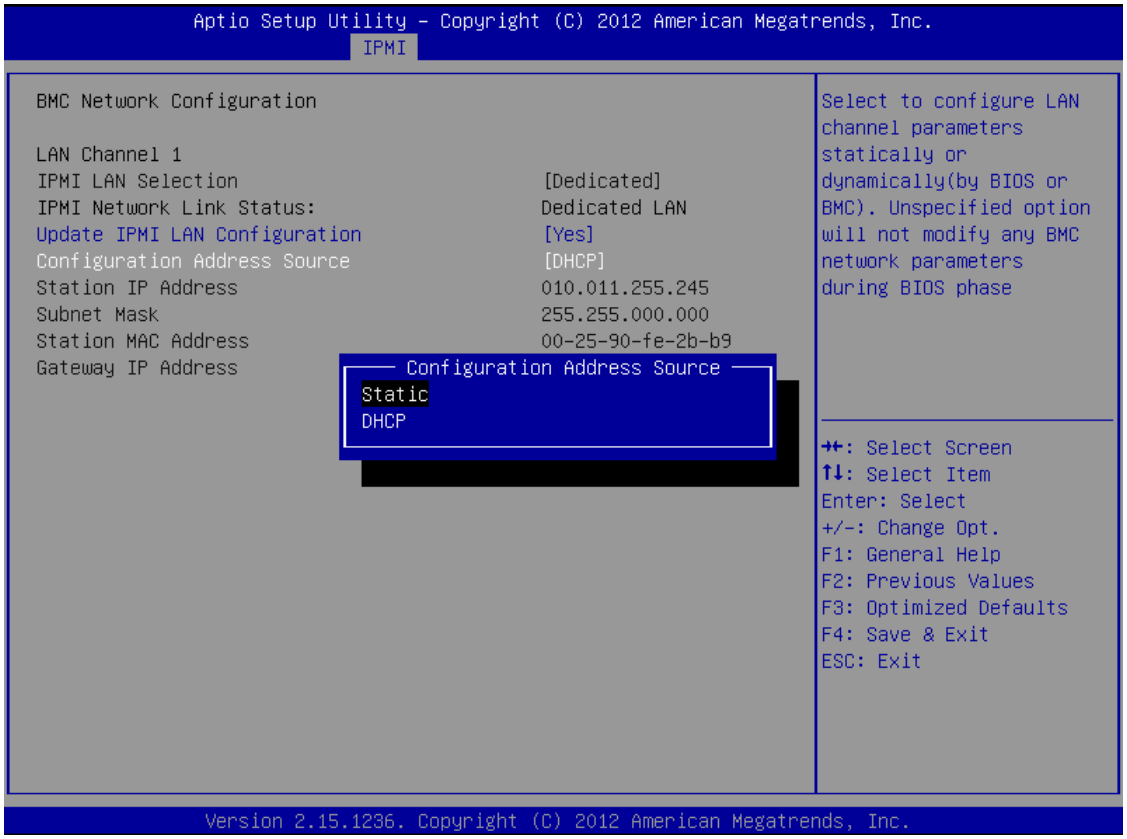
9. On the **BMC Network Configuration** page, select **Update IPMI LAN Configuration**, press `Enter`, and select **Yes**.

Figure 350: BMC Network Configuration page > Update IPMI LAN Configuration



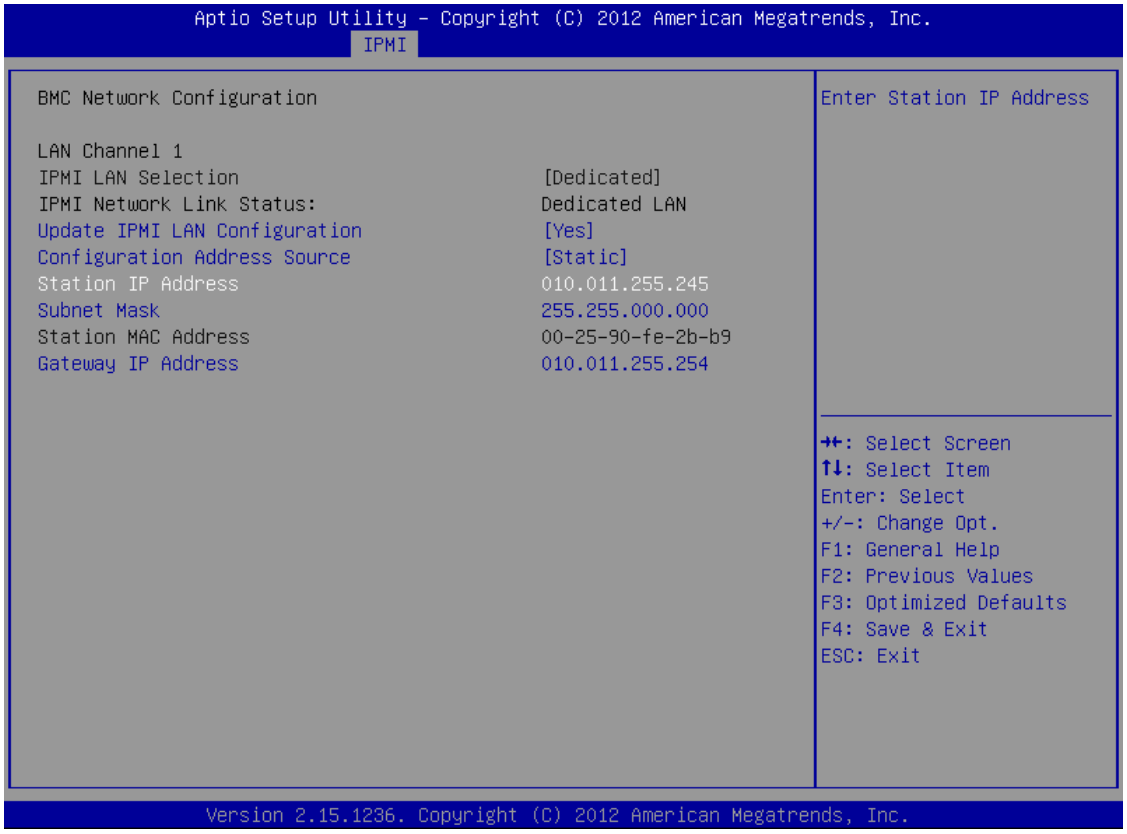
10. Stay on the **BMC Network Configuration** page, select **Configuration Address Source**, press **Enter**, and select **Static**.

Figure 351: BMC Network Configuration page > Configuration Address Source



11. Still on the **BMC Network Configuration** page, configure the **Station IP Address**, **Subnet Mask**, and **Gateway IP Address** individually.

Figure 352: BMC Network Configuration page > Station IP Address, Subnet Mask, Gateway IP Address



12. Press F4 to save the settings, and exit from the BIOS.

    About a minute later, you will be able to log in on the IPMI web interface.

Incomplete TSA response received

When using a TSA certificate generated with Windows Certificate Authority, you might see a similar error message:

```
Incomplete TSA response received, TSA HTTP server may be responding slowly;
errno='Success (0)', timeout_seconds='30'
```

When generating the certificate, make sure that you do the following:

Optional Key Usage: If **Key Usage** is present, it must be `digitalSignature` and/or `nonRepudiation`. Other values are not permitted. Make sure that in **Encryption**, **Allow key exchange without key encryption (key agreement)** is selected.

CAUTION:

In **Encryption**, do NOT select **Allow key exchange only with key encryption (key encipherment)**, because it will result in errors.

For more information, see Generating TSA certificate with Windows Certificate Authority on Windows Server 2016 or later.

24

Using SPS with SPP

You can link your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment. That way you can jointly use the features of the two deployments.

Both appliances provide different functionality. You can use them together or independently from each other.

SPP provides:

- Machine and account discovery
- Password rotation and management
- Advanced access request and approval workflows
- A user portal and desktop application to initiate connections

SPS provides:

- Transparent or non-transparent interception of remote admin protocols (SSH, RDP, Telnet, Citrix ICA, and VNC)
- Audit recording and video-like playback of sessions
- Inband authentication of the monitored users independently from the target servers
- Basic access control policy enforcement
- Advanced search and reporting capabilities in the audit records
- Built-in user behavior analytics for the recorded sessions (One Identity Safeguard for Privileged Analytics)

Prerequisites

Before you start, ensure that Network Level Authentication (NLA) is enabled in the RDP setting policies. Also ensure that the CVE-2018-0886 update of the Credential Security Support Provider protocol (CredSSP) from Microsoft has been installed. For more information, see Creating and editing protocol-level RDP settings on page 582.

CAUTION: When linking your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

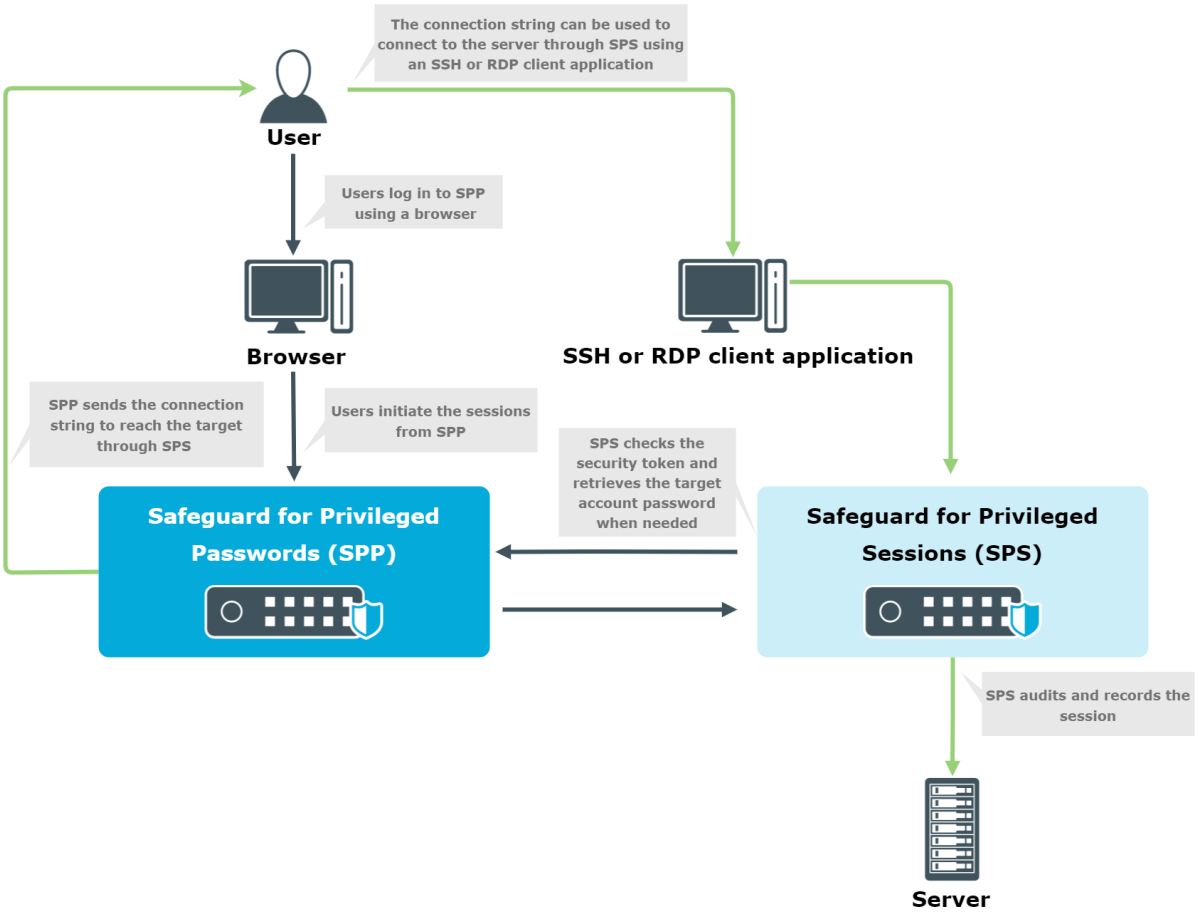
Passwords-initiated (SPP-initiated) workflow

In the Passwords-initiated workflow, the users initiate sessions from SPP. In this workflow, SPP uses SPS as a session-recording device.

You can use your browser to request access from SPP and initiate the connection to the target server through SPS. SPP creates an access string for the user's SSH or RDP client that allows these clients to connect to the target server through SPS, so that SPS can audit and record the session. In this sense, this workflow is nontransparent, the user must use a browser.

This is what all SPS users who bought the Sessions Module use before SPP version 2.7.

Figure 353: Passwords-initiated (SPP-initiated) workflow



For details on configuring this workflow, see Configuring SPP for Passwords-initiated workflow.

Sessions-initiated (SPS-initiated) workflow

In the Sessions-initiated workflow, the users initiate sessions from SPS. In this workflow SPS uses SPP as a credential store.

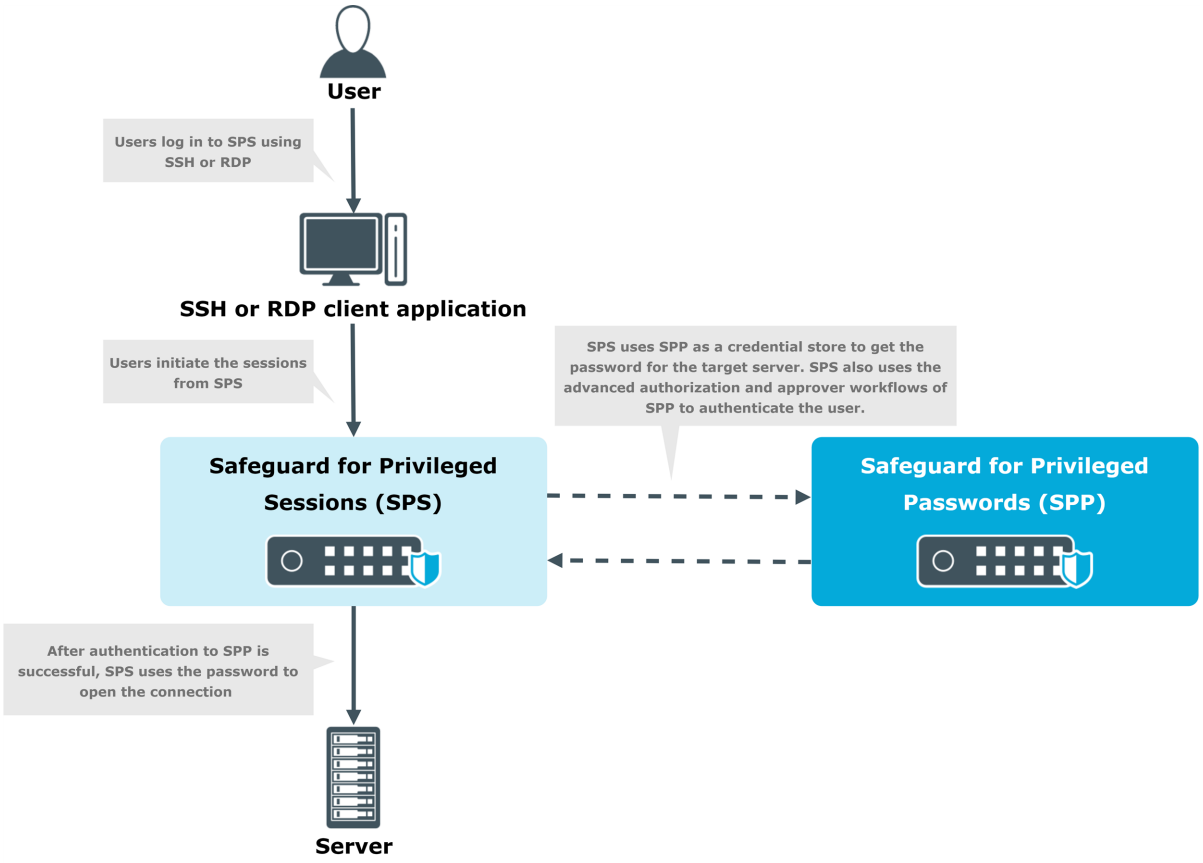
This workflow is transparent in the sense that you can connect to the target server or to SPS directly using your SSH or RDP client application. SPS authenticates these clients and communicates with SPP to get the password for the target server. It then uses that password to open the connection. Authentication happens on SPS, while authorization happens on SPP based on the user's entitlements.

This is what old and new users of standalone SPS are likely to prefer.

The usual SPP Access Requests workflows that SPP provides are supported:

- Auto-approved access request
- Approved/denied access request (similar to the four-eyes authorization feature of SPS)

Figure 354: Sessions-initiated (SPS-initiated) workflow



Configuring the Passwords-initiated workflow

Passwords-initiated (SPP-initiated) workflow

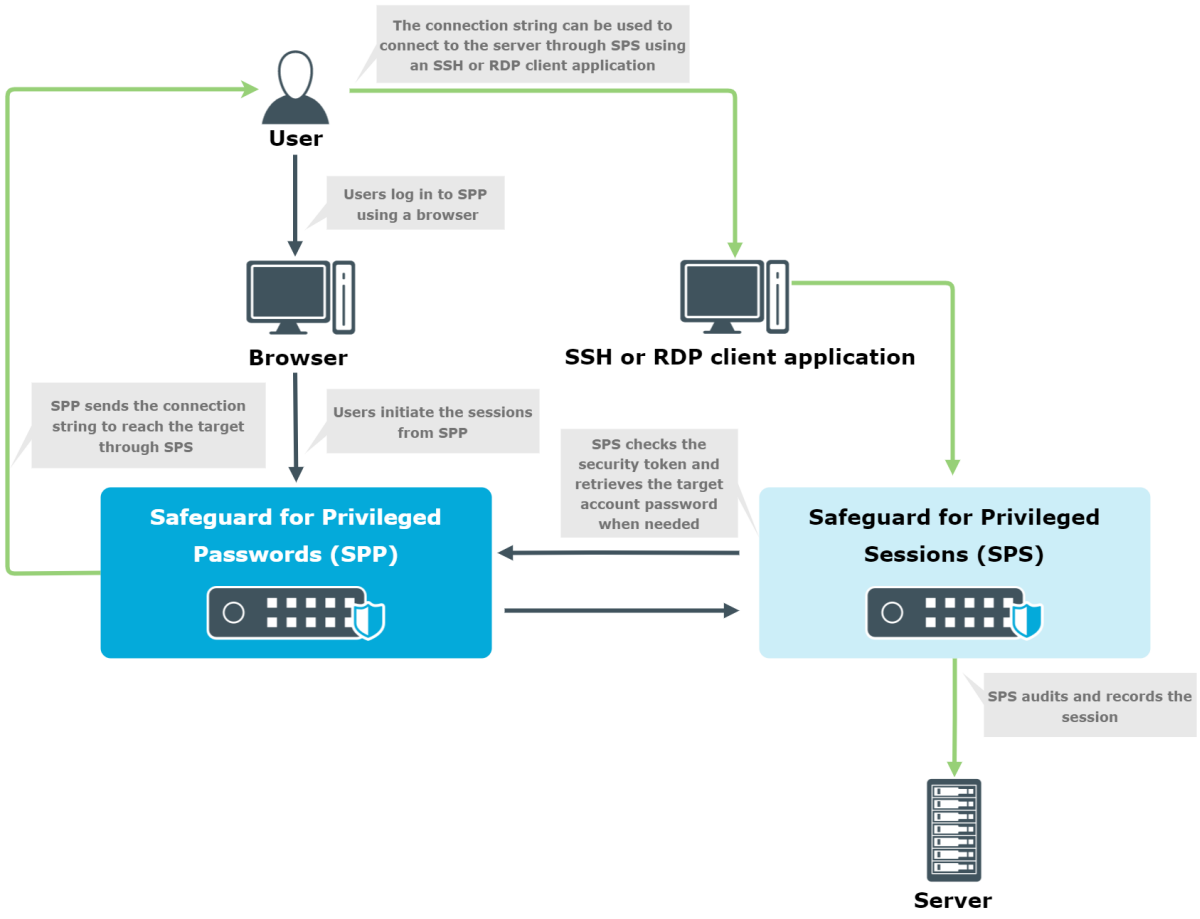
In the Passwords-initiated workflow, the users initiate sessions from SPP. In this workflow, SPP uses SPS as a session-recording device.

You can use your browser to request access from SPP and initiate the connection to the target server through SPS. SPP creates an access string for the user's SSH or RDP client that allows these clients to connect to the target server through SPS, so that SPS can audit

and record the session. In this sense, this workflow is nontransparent, the user must use a browser.

This is what all SPS users who bought the Sessions Module use before SPP version 2.7.

Figure 355: Passwords-initiated (SPP-initiated) workflow



For details on configuring this workflow, see Configuring SPP for Passwords-initiated workflow.

Prerequisites

- Minimum versions:
  - SPP version 2.7
  - SPS version 6.0
- You must have built an SPS cluster by promoting an SPS node to the role of Central Management node, even if it is a single node. For more information, see Creating a cluster.
- **CAUTION:** When linking your One Identity Safeguard for Privileged Sessions

(SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

To configure the Passwords-initiated (SPP-initiated) workflow

1. On SPS, link SPP and SPS as described in Linking SPS to SPP.
2. Configure SPP to use the linked SPS as described in Configuring SPP for Passwords-initiated workflow.
3. Optionally, customize monitoring settings as follows:
    - To make use of the more advanced features of SPS, you can change the **safeguard_default** Connection Policy or create a new Connection Policy and select that in SPP.
    - Follow the AA plugin settings listed in section Sharing RDP connection policies with SPS.

Configuring SPP for Passwords-initiated workflow

To configure SPP to use the linked SPS in Passwords-initiated (SPP-initiated) workflows, complete the following steps. For more information on the workflow, see Using SPS with SPP.

Prerequisites

- Minimum SPP version: 2.7.
- You have linked SPP and SPS as described in Linking SPS to SPP.

To configure SPP for Passwords-initiated workflow

1. On SPP, assign the managed networks for sessions management.
    a. Navigate to **Appliance Management** > **Cluster** > **Managed Networks**.
    b. Add the network you want to monitor with SPS and choose the SPS appliance for the **Sessions Managed By** field.
2. Select the SPS for the access request policy.
    a. Navigate to **Security Policy Management** and either select an existing **Entitlement**, or create one. After that, in **Entitlements**, on the **Access**

Request Policies tab, select the edit or the add icon.

b. On the **General** tab, select **Session** at **Choose Request Policy Type**, and select the appropriate session type at **Choose Session Type**.

c. On the **Security Tab**, select the SPS Connection Policy.

Configuring the Sessions-initiated workflow

Sessions-initiated (SPS-initiated) workflow

In the Sessions-initiated workflow, the users initiate sessions from SPS. In this workflow SPS uses SPP as a credential store.

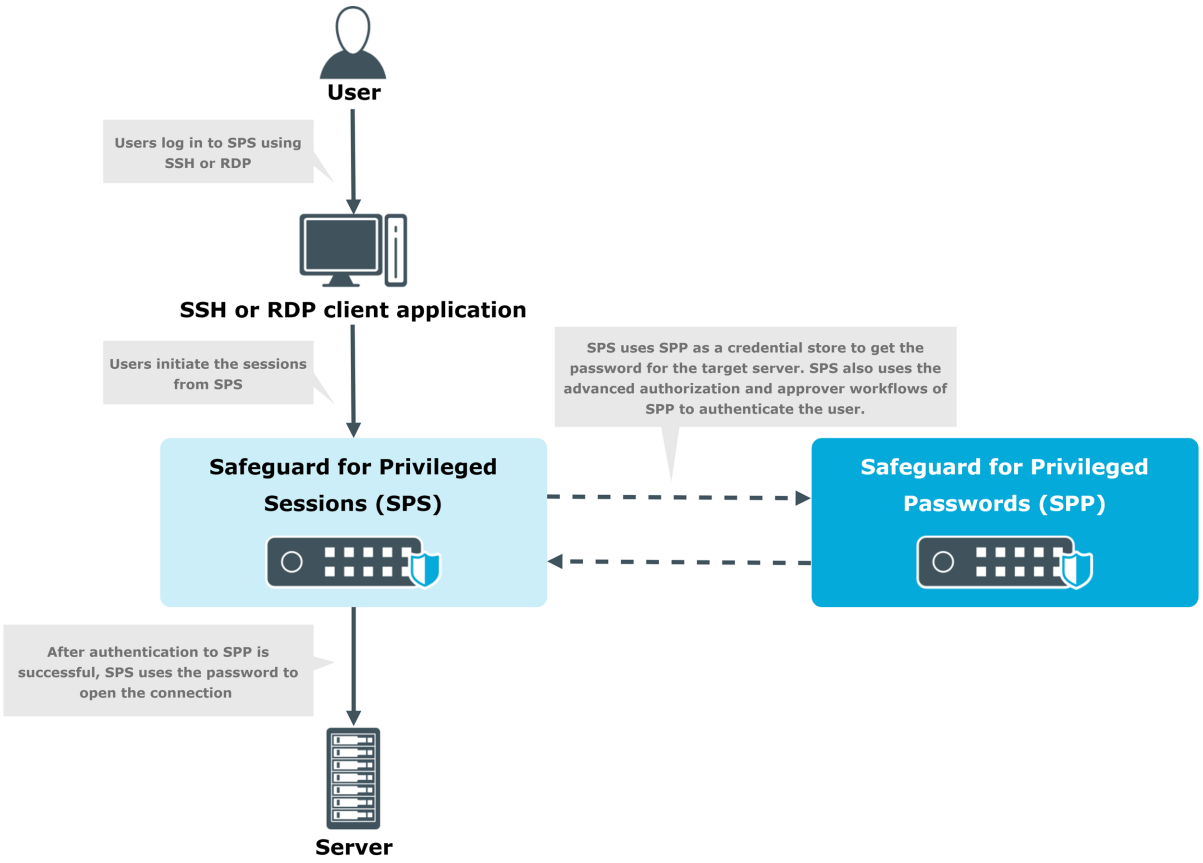
This workflow is transparent in the sense that you can connect to the target server or to SPS directly using your SSH or RDP client application. SPS authenticates these clients and communicates with SPP to get the password for the target server. It then uses that password to open the connection. Authentication happens on SPS, while authorization happens on SPP based on the user's entitlements.

This is what old and new users of standalone SPS are likely to prefer.

The usual SPP Access Requests workflows that SPP provides are supported:

- Auto-approved access request
- Approved/denied access request (similar to the four-eyes authorization feature of SPS)

Figure 356: Sessions-initiated (SPS-initiated) workflow



Prerequisites

- Minimum versions:
  - SPP version 2.8
  - SPS version 6.2 and newer, including 6.0.2 and newer versions of the 6.0.x branch, but excluding 6.1.x
- You must have built an SPS cluster by promoting an SPS node to the role of Central Management node, even if it is a single node. For more information, see Creating a cluster.

CAUTION: When linking your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

Limitations

- Only SSH and RDP sessions are supported.
- Users must perform gateway authentication on SPS with the same username they have Entitlements for in SPP.
    - For SSH sessions, the gateway authentication can use a Local User Database, an LDAP server, or an Active Directory server as authentication backend.

      Note that SPP does not support every type of LDAP and Active Directory settings that SPS does. Verify that you can configure both appliances to access and retrieve data from the LDAP or Active Directory server.
    - For RDP sessions, SPS must be configured to act as a Remote Desktop Gateway. For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway.

      The gateway authentication can use a Local User Database or an Active Directory server as authentication backend. When using an Active Directory server, note the following points.
        - Both SPS and SPP must use the same server, and be the member of the same domain as the Active Directory server.
        - SPP does not support every type of Active Directory settings that SPS does. Verify that you can configure both appliances to access and retrieve data from the Active Directory server.
        - SPS does not receive the domain of the authenticated user from the Domain Controller. SPS assumes that the user belongs to the same domain that SPS is joined into. (Configuring trust between the Domain Controller of SPS and the Domain Controller of the user does not solve this problem.)
- You must use a uniform address for the target server. Use either its IPv4 address or its hostname everywhere: when configuring the Assets in SPP, the Connection Policies and Channel Policies in SPS, and also when the user sets the target address in the SSH/RDP application. Otherwise, the authentication will fail.

To configure the Sessions-initiated (SPS-initiated) workflow

1. On SPS, link SPP and SPS as described in Linking SPS to SPP.
2. Configure SPP to allow SPS to request passwords from SPP as described in Configuring SPP for Sessions-initiated workflow.
3. Configure SPS to use the linked SPP as a Credential Store as described in Configuring SPS for Sessions-initiated workflow.
4. Optionally, customize monitoring settings as follows:

Configuring SPP for Sessions-initiated workflow

To configure SPP to use the linked SPS in Sessions-initiated (SPS-initiated) workflows, complete the following steps. For details on the workflow, see Using SPS with SPP.

Prerequisites

- Minimum versions:
    - SPP version 2.8
    - SPS version 6.2 and newer, including 6.0.2 and newer versions of the 6.0.x branch, but excluding 6.1.x
- You must have built an SPS cluster by promoting an SPS node to the role of Central Management node, even if it is a single node. For more information, see Creating a cluster.

CAUTION: When linking your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

To configure SPP for Sessions-initiated workflow

1. Ensure that SPP can access the usernames that authenticate on the SPS gateway either from a local user database, or from an Active Directory.
2. Ensure that these users have the proper entitlements to access the passwords to their target servers. Otherwise, SPS rejects their sessions.

Configuring SPS for Sessions-initiated workflow

To configure SPS to use the linked SPP in Sessions-initiated (SPS-initiated) workflows, complete the following steps. For details on the workflow, see Using SPS with SPP.

Prerequisites

- Minimum versions:
    - SPP version 2.8
    - SPS version 6.2 and newer, including 6.0.2 and newer versions of the 6.0.x branch, but excluding 6.1.x
- You must have built an SPS cluster by promoting an SPS node to the role of Central Management node, even if it is a single node. For more information, see Creating a cluster.

CAUTION: When linking your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

Limitations

- Only SSH and RDP sessions are supported.
- Users must perform gateway authentication on SPS with the same username they have Entitlements for in SPP.
    - For SSH sessions, the gateway authentication can use a Local User Database, an LDAP server, or an Active Directory server as authentication backend.

      Note that SPP does not support every type of LDAP and Active Directory settings that SPS does. Verify that you can configure both appliances to access and retrieve data from the LDAP or Active Directory server.
    - For RDP sessions, SPS must be configured to act as a Remote Desktop Gateway. For details, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway.

      The gateway authentication can use a Local User Database or an Active Directory server as authentication backend. When using an Active Directory server, note the following points.
        - Both SPS and SPP must use the same server, and be the member of the same domain as the Active Directory server.
        - SPP does not support every type of Active Directory settings that SPS does. Verify that you can configure both appliances to access and retrieve data from the Active Directory server.
        - SPS does not receive the domain of the authenticated user from the Domain Controller. SPS assumes that the user belongs to the same domain that SPS is joined into. (Configuring trust between the Domain

Controller of SPS and the Domain Controller of the user does not solve this problem.)

- You must use a uniform address for the target server. Use either its IPv4 address or its hostname everywhere: when configuring the Assets in SPP, the Connection Policies and Channel Policies in SPS, and also when the user sets the target address in the SSH/RDP application. Otherwise, the authentication will fail.

To configure SPS for Sessions-initiated workflow

Configure Connection Policies on SPS to audit your sessions. Note that you have to complete these steps for each Connection Policy that uses SPP as a Credential Store. For the general steps on configuring Connection Policies, see Configuring connections.

1. Select **Credential Store** > **safeguard_default**.
2. Select **Usermapping policy** > **safeguard_default**.
3. Configure gateway authentication. The users must perform gateway authentication on SPS with the same username they have Entitlements for in SPP. For details, see Configuring gateway authentication.
4. When you are using an Approve/deny workflow on SPP, increase the **Idle timeout** setting of the Connection Policy. SPS will wait for an approval from SPP until half the time set in **Idle timeout**. For example, if the authorizer on SPP has 2 minutes to approve the access request, set the **Idle timeout** option on SPS to more than double this value, for example, 5 minutes.

Configuring SPS for SRA-initiated workflow

To configure SPS to use SRA in Sessions-initiated (SPS-initiated) workflows, complete the following steps.

NOTE: Only SSH and RDP sessions are supported.

Prerequisites

- Minimum versions:
  - SPP version 2.8
  - SPS version 6.2 and newer, including 6.0.2 and newer versions of the 6.0.x branch, but excluding 6.1.x
- You must have built an SPS cluster by promoting an SPS node to the role of Central Management node, even if it is a single node. For more information, see Creating a cluster.

CAUTION: When linking your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment, ensure that the SPS and SPP versions match exactly, and you keep the versions synchronized during an upgrade. For example, you can only link SPS version 6.6 to SPP version 6.6, and if you upgrade SPS to version 6.7, you must also upgrade SPP to 6.7.

Make sure that you do not mix Long Term Supported (LTS) and feature releases. For example, do not link an SPS version 6.0 to an SPP version 6.1.

- You cannot use gateway authentication on SPS. For more information, see Configuring gateway authentication.
- For RDP sessions, you must not configure SPSto act as a Remote Desktop Gateway. For more information, see Using One Identity Safeguard for Privileged Sessions (SPS) as a Remote Desktop Gateway.
- You must use a uniform address for the target server. Use either its IPv4 address or its hostname everywhere: when configuring the Assets in SPP, the Connection Policies and Channel Policies in SPS, and also when the user sets the target address in the SSH/RDP application. Otherwise, the authentication will fail.
- Enable SRA on **Basic Settings** > **Starling Integration** > **Remote Access** > **Enable Remote Access**.

To configure SPS for SRA-initiated workflow

1. Download the SRA initiated plugin, then upload and configure the plugin as required.

   For more information, see Using a custom Authentication and Authorization plugin to authenticate on the target hosts.
2. Navigate to **Traffic Controls** > **Protocol name** > **Connections** and create a new connection.

   Configure Connection Policies on SPS to audit your sessions. Note that you have to complete these steps for each Connection Policy that uses SPP as a Credential Store. For the general steps on configuring Connection Policies, see Configuring connections.
3. Select **Credential Store** > **safeguard_default**.
4. For **AA plugin**, select the SRA plugin configuration instance you have uploaded.

   NOTE:

   Starting from One Identity Safeguard for Privileged Sessions (SPS) 8.0 LTS, AA plugins are no longer required to connect SPS and One Identity Safeguard Remote Access.
5. Select **Usermapping policy** > **safeguard_default**.
6. Ensure that **Act as a Remote Desktop Gateway** and **Require Gateway Authentication on the SPS Web Interface** are cleared.
7. Navigate to **Traffic Controls** > **Protocol name** > **Settings**.

When you are using an Approve/deny workflow on SPP, increase the **Idle timeout** setting of the Connection Policy. SPS will wait for an approval from SPP until half the time set in **Idle timeout**. For example, if the authorizer on SPP has 2 minutes to approve the access request, set the **Idle timeout** option on SPS to more than double this value, for example, 5 minutes.

Linking SPS to SPP

You can link your One Identity Safeguard for Privileged Sessions (SPS) deployment to your One Identity Safeguard for Privileged Passwords (SPP) deployment.

IMPORTANT: Once performed, you cannot unlink the SPS and SPP deployments.

If the primary IP address of your SPS deployment or SPP deployment changes, you must repeat this procedure to relink the clusters. To repeat the procedure to relink the clusters, navigate to **Basic Settings** > **Cluster management** and click **Relink SPP cluster**.

Prerequisites

Before you start linking your SPS deployment to your SPP deployment, consider the following:

- Your SPS deployment must be in an SPS cluster, set as a Central management node. Even if your SPS deployment consists of a single, standalone node, you must assign the Central management role to its own single-node cluster. For details, see Managing One Identity Safeguard for Privileged Sessions (SPS) clusters.

  Configuration synchronization must be enabled between the nodes of the SPS cluster. This is required so that SPP entitlements work properly for each SPS node.

  NOTE: If you have multiple standalone SPS appliances, consider joining them to a cluster before linking SPP. In general, One Identity recommends creating a cluster if the nodes can use a common configuration, or later you might want to centrally search the data of every node. Creating a cluster from the SPS nodes after linking SPP is problematic and should be avoided.
- You will need the primary IP address or the hostname of your SPP deployment that SPS can use to access SPP. Only IPv4 addresses are supported.
- You will need the username and password to an SPP account with "Appliance" and "Operations" permissions.
- Verify that your SPS policies do not contain the `safeguard_default` string in their names. During the linking process, SPS automatically creates and configures several policies and plugins. The name of these policies usually contains the string `safeguard_default`. Existing policies with such names will be overwritten.
- The SPP and SPS nodes must be able to communicate on the TCP `8649` port. If needed, update your firewall policies.

NOTE: When updating your firewall policies to enable the connection between SPS and SPP nodes, consider the following:

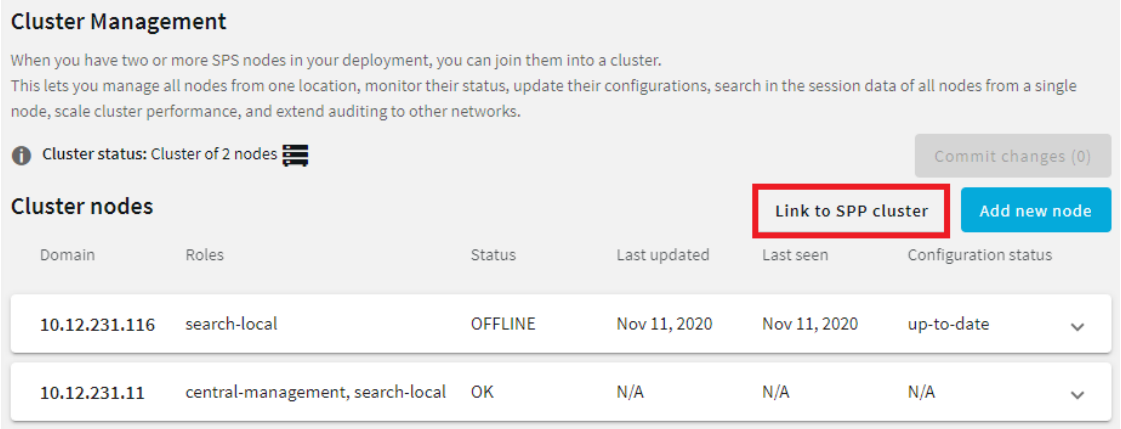
- **Connecting SPS nodes to SPS nodes**: Make sure that between all the SPS nodes, the `500` and `4500` UDP ports are opened bidirectionally. This is required so that every node can initiate and accept connections from every other node through the mentioned ports.
- **Connecting SPP nodes to SPP nodes**: For more information, see *Enrolling replicas into a cluster* in the One Identity Safeguard for Privileged Passwords Administration Guide.
- **Connecting SPS nodes to SPP nodes**: Make sure that between all the SPS and SPP nodes, the `8649` TCP port is opened bidirectionally. This is required so that every node can initiate and accept connections from every other node through the mentioned port. If there are `M` SPP and `N` SPS nodes in your setup, then create `M × N × 2` firewall rules to link the SPS nodes to the SPP nodes.

- During the linking process, SPS must be able to access SPP using HTTPS on the TCP `443` port. This is required only once during the linking process. If needed, update your firewall policies.

To link your SPS deployment to SPP

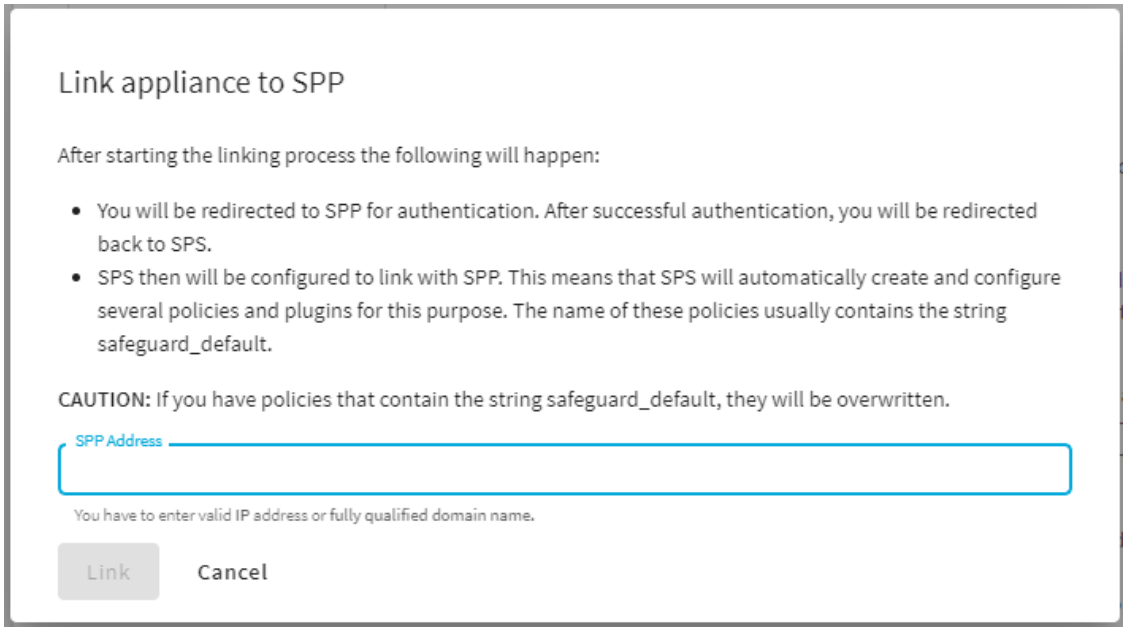
1. (Optional) Create a configuration backup of SPS. For details, see Creating configuration backups.
2. (Optional) Create a configuration backup of SPP. For details, see the Safeguard for Privileged Passwords Administration Guide, Backup and Retention settings.
3. Login to the Central management node of your SPS cluster. This node has **Central management** listed in the **Basic Settings** > **Cluster management** > **Roles** column.
4. Navigate to **Basic Settings** > **Cluster management**, and click **Link to SPP Cluster**.

Figure 357: Basic Settings > Cluster management — Linking SPS to SPP



The **Link Appliance to SPP** dialog then appears.

Figure 358: Basic Settings > Cluster management > Link to SPP Cluster — The Link Appliance to SPP dialog



5. Enter the primary IP address of SPP in the **SPP Address** field.

   NOTE: Only IPv4 addresses are supported.
6. Click **Link**. Wait until you are redirected to SPP.
7. Login to SPP. Wait until you are redirected to SPS.
8. Wait until SPS creates and configures the policies and plugins required for the joint operation of SPS and SPP. This step can usually take up to a minute.
9. You will receive a message:
   - If the linking is unsuccessful, this message displays: `Request failed`.

     If this happens, check the credentials and the IP address that you provided.

     For details on resolving errors, see SPP to SPS link issues on page 994 and SPP to SPS link error resolution on page 991.
   - If the linking is successful, this message displays: `SPS successfully linked to SPP`.

     SPP automatically closes any open access requests.
10. Log out from the SPS web interface.

CAUTION:

If the primary IP address of your SPS or SPP changes, you must repeat the linking procedure to relink the clusters. Use the Relink SPP cluster button to do so.

Switching seamlessly between SPS and SPP

You can seamlessly switch between the SPS and SPP web interfaces using the  icon located next to the user menu. Seamless switching uses a federated login method, which enables you to switch between SPS and SPP appliances without having to reenter your credentials.

Prerequisites

- SPP is joined to SPS, for more information, see Linking SPS to SPP.
- On SPP, when you create an external federation, unselect the **Require User to Always Authenticate** check box. If you select the **Require User to Always Authenticate** check box, you will always be required to enter your credentials on the external provider. For more information on SPP external federation settings, see *External Federation settings* in the *One Identity Safeguard for Privileged Passwords Administration Guide*.

To set up seamless switching between SPS and SPP

1. Complete the steps listed in Authenticating users with SAML2 login method.

   In **Script name**, enter the Login Provider ID of SPP.

   NOTE: The Login Provider ID is case sensitive.
2. If you are editing the same SAML2 login method that you are currently using, after editing the method, log out and log in again for the changes to take effect.
3. Select the  icon, and after that select the logo of SPP under **Linked Products**.

   You have now switched to the SPP appliance.

Troubleshooting the SPS to SPP link

SPP to SPS link error resolution

Common linking error resolutions follow which may occur when linking One Identity Safeguard for Privileged Passwords (SPP) to One Identity Safeguard for Privileged Sessions (SPS).

Typo in SPP's address, the address is not reachable

- Error: The browser reports errors when SPS redirects to SPP's login page, for example, `This site can’t be reached`. The exact error message depends on the browser.
- Resolution: Click the Back button of the browser and enter the correct address into the **SPP Address** field.

Typo in SPP's address, the address is alive, but not an SPP

- Error: After clicking the Join button, a web site other than SPP's Login interface is displayed.
- Resolution: Click the Back button of the browser and enter the correct address into the **SPP Address** field.

SPP's HTTPS certificate does not match its IP address or hostnam

- Error message: `Error joining to SPP: SPS has failed to join to SPP. For more information, see the error details. (JoinFailed)`
- Raw error:

```
{
    "response": "Error sending request: SSLError: HTTPSConnectionPool
(host='examplespp.company', port=443): Max retries exceeded with url:
/service/core/v3/Cluster/SessionModules (Caused by SSLError
(CertificateError(\"hostname 'examplespp.company' doesn't match
'192.0.2.123'\",),))",
    "status": null,
    "url":
"https://examplespp.company/service/core/v3/Cluster/SessionModules"
 }
```

- Resolution:
  - If SPP's certificate contains SPP's IPv4 address in the **Common Name** or **subjectAltName** field, then enter that IP address when linking SPS to SPP.
  - If SPP's certificate contains only its DNS name in the **Common Name** or **subjectAltName** field, then use that hostname when linking SPS to SPP.
  - Otherwise, set up an SSL server certificate for SPP which matches its IP address in the certificate's **Common Name** or **subjectAltName**fields (see SSL Certificates in the Safeguard Administration Guide) and retry linking. Wait about five minutes to let the timeout of the failed link request expire before starting a new link request after a failed incomplete one. (Alternatively, see Reversing the SPP to SPS join in the Safeguard Administration Guide.)

Typo in SPP credentials

- Error: Login to the SPP web interface fails.
- Raw error:

```
{
            "error": "invalid_request",
            "error_description": "Access denied.",
            "success": false
      }
```

- Resolution: Make sure that the correct username and password are entered in the **SPP username**: and **SPP password:** prompts.

SPP user has insufficient permission

- Error message: `Error joining to SPP: SPS has failed to join to SPP. For more information, see the error details. (JoinFailed)`

```
 Authorization is required for this request.

       Code: 60108
       URL: https://192.0.2.123/service/core/v3/Cluster/SessionModules
       Status: 403
```

- Resolution: When SPS redirects to SPP's Login interface, then login as an SPP user has "Appliance" and "Operations" permissions.

SPS is already linked to SPP

- Error message: `Error joining to SPP: SPS has failed to join to SPP. For more information, see the error details. (JoinFailed)`

```
The session connection has a missing, invalid, or non-unique
value [ NodeId ].

       Code: 60657
       URL: https://192.0.2.123/service/core/v3/Cluster/SessionModules
       Status: 400
```

- Resolution: See Reversing the SPP to SPS join in the Safeguard Administration Guide.

Linking takes too long (more than five minutes)

- `ERROR: Request to https://192.0.2.123/service/a2a/v2/PsmValidation failed, response (HTTP 403):`

```
{
      "Code": 60108,
      "Message": "Authorization is required for this request.",
      "InnerError": null
}
```

- Resolution: Make sure that SPS is not overloaded and try linking again.

SPP to SPS link issues

In addition to the monitoring tools in Safeguard for Privileged Passwords (SPP), you can use the SPP to SPS link issues on page 994 during the linking process. Several SPS tools are described below.

Link process fails and real-time monitoring

If the linking process fails for any reason, consult the system logs.

To view the One Identity Safeguard for Privileged Sessions logs, navigate to **Basic Settings** > **Troubleshooting** > **View log files**.

To show only the logs for the linking process:

1. Select a **Logtype** of **syslog**.
2. Select the **Day**; today is the default.
3. In the **Show only messages containing** text box, enter `SPP-join`.

Use the buttons at the bottom of the dialog to perform the following tasks:

- To download the log file, click **Download**.
- To follow the current log messages real-time, click **Tail**. The latest logs will update in a browser window while you interact with the linking process.
- To display the log messages, click **View**.

To increase the level of detail in the log, enable debug level logging:

1. Navigate to **Basic Settings** > **Troubleshooting** > **Create support bundle**.
2. Click **Start**.

Linking successful but connections do not work

When SPP and SPS report a successful linking, but the connections do not work, view the SPS connection logs.

In SPS, navigate to **Basic Settings** > **Troubleshooting** > **View log files**.

To show only the logs for the linking process:

1. Select a **Logtype** of **ssh** or **rdp**.
2. Select the **Day** (today is the default).
3. In the **Show only messages containing** text box, enter `SPP-join`.

To change the verbosity level of SPS, complete the following steps:

1. Navigate to the **Global Options** page of the traffic for which you want to change the log level. For example, go to **Traffic Controls** > **SSH** > **Global Options** to change the log level of SSH traffic, **Traffic Controls** > **RDP** > **Global Options** for remote desktop traffic, and so on.
2. Select the desired log level from the **Verbosity level** field. The verbosity level ranges from 1 (no logging) to 10 (extremely detailed), with level 4 being the default normal level.

CAUTION: High verbosity levels generate a very large amount of log messages and might result in a very high load on the machine. Log levels set around 9 to 10, may result in logs with highly sensitive data, for example, passwords in plain text format.

Testing network issues

You can use the Diagnostics tools of SPP and SPS to test network issues. The following commands are available:

- **ping**: Sends a simple message to the specified host to test network connectivity.
- **traceroute**: Sends a simple message from SPS to the specified host and displays all hosts on the path of the message. It is used to trace the path the message travels between the hosts.
- **connect**: Attempts to connect the specified host using the specified port. It is used to test the availability or status of an application on the target host.

To execute one of the above commands on SPS, see Network troubleshooting on page 935.

To execute one of the above commands on SPP, see Diagnostics tools of SPP.

Creating an SPS Support Bundle

If you have an issue which needs Support assistance, you may be asked to provide an SPS Support Bundle. To collect system-state information (also known as a debug bundle) in One Identity Safeguard for Privileged Sessions, see Collecting logs and system information for error reporting on page 940.

Configuring external devices

This section describes scenarios about configuring external devices to redirect selected traffic to One Identity Safeguard for Privileged Sessions (SPS).

Configuring advanced routing on Linux

The following describes how to configure a Linux-based router to redirect selected traffic to One Identity Safeguard for Privileged Sessions (SPS) instead of its original destination. This procedure should work on most modern Linux-based routers, including Check Point firewalls.

Prerequisites

The router must have the `iptables` and `ip` tools installed.

To configure a Linux-based router to redirect selected traffic to SPS instead of its original destination

1. Create the packet filter rules that will mark the connections to be sent to SPS using the CONNMARK feature of iptables. Mark only those connections that must be redirected to SPS.

```
# iptables -t mangle -I PREROUTING -i <interface-facing-the-clients> -p tcp -d <network-of-the-servers> --dport <port-to-access> -j CONNMARK --set-mark 1
```

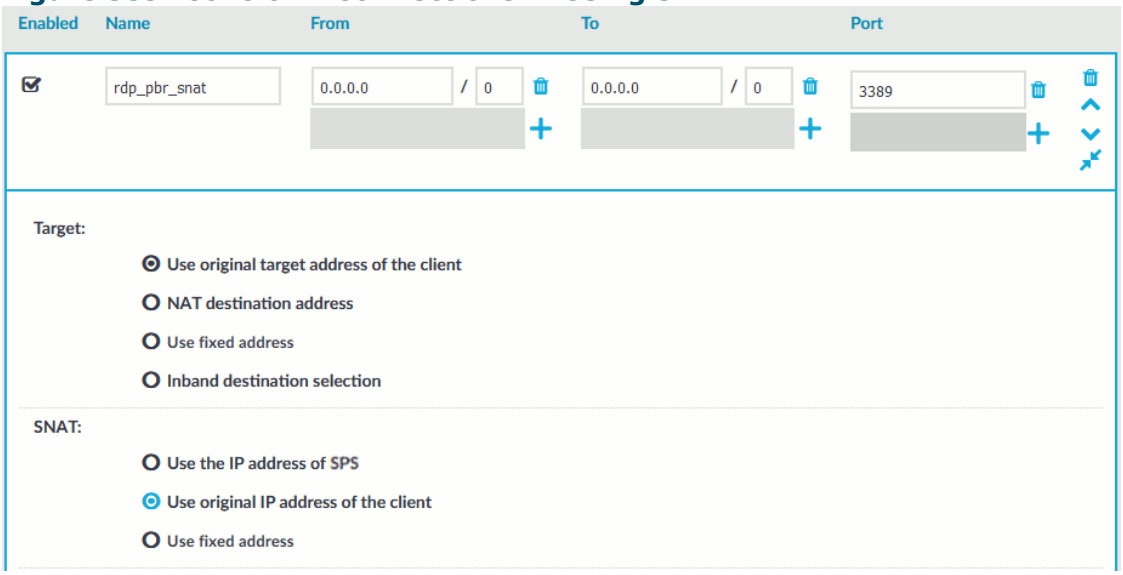
> **Example: Setting up a connection mark for Linux policy routing**
>
> For example, if the network interface of the router that faces the clients is called `eth0`, the servers are located in the `10.0.0.0/24` subnet, and the clients access the servers using port `3389` (the default port of the RDP protocol), then this command looks like:

```
# iptables -t mangle -I PREROUTING -i eth0 -p tcp -d 10.0.0.0/24 --
dport 3389 -j CONNMARK --set-mark 1
```

Create a rule that redirects the answers of the servers to SPS. That way both the client-to-server and the server-to-client traffic is routed to SPS.

NOTE: This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 359: Control > Connections — Using SNAT



2.

```
# iptables -t mangle -I PREROUTING -i <interface-facing-the-servers> -
p tcp -s <network-of-the-servers> --sport <port-to-access> -j CONNMARK
--set-mark 1
```

3. Convert the CONNMARK marks to MARK:

```
# iptables -t mangle -A PREROUTING ! -i <interface-facing-the-scb> -m
connmark --mark 1 -j MARK --set-mark 1
```

CAUTION:

This rule must be placed after the CONNMARK rules.

4. Add the table name to the `/etc/iproute2/rt_tables` of the router. Use the following format (for details on routing tables, see for example the Guide to IP Layer Network Administration with Linux):

```
103 scb
```

5. Create a routing table that has a single entry with a default route to SPS:

```
# /sbin/ip route add default via <ip-address-of-SPS> table scb
```

6. Create a routing rule that selects the routing table called `scb`, if the connection is marked.

```
# /sbin/ip rule add from all fwmark 1 table scb
```

7. If SPS is configured to spoof the IP address of the clients on the server side (that is, the **SNAT** > **Use original IP address of the client** option of the connection policies is selected), enable spoofing on the router for the interface connected to SPS.

```
# echo 0 > /proc/sys/net/ipv4/conf/<interface-facing-SPS>/rp_filter
# echo 0 > /proc/sys/net/ipv4/conf/all/rp_filter
```

Expected result

The traffic from the clients targeting the specified port of the servers is redirected to SPS. Therefore, SPS can be configured to control and audit this traffic.

Configuring advanced routing on Cisco routers

The following describes how to configure a Cisco router to redirect selected traffic to One Identity Safeguard for Privileged Sessions (SPS) instead of its original destination. This procedure should work on most modern Cisco IOS releases but was specifically tested on IOS version 12.3.

To configure a Cisco router to redirect selected traffic to SPS instead of its original destination

1. Create an ACL (Access Control List) entry that matches the client and server subnets and the to-be-audited port. Keep in mind that whatever is permitted by this ACL is what will be matched, so make sure that the scope of the ACL entry is narrowed down as much as possible.

```
#(config) ip access-list extended ssh-inbound
#(config-ext-nacl) permit tcp <src net> <src mask> <dst net> <dst mask>
eq <dst port>
```

Example: Configuring an ACL entry for Cisco policy routing

For example, if the clients are in the `192.168.0.0/24` subnet, the servers are located in the `10.0.0.0/24` subnet, and the clients access the servers using port `22` (the default port of the SSH protocol), then the permit clause should be:

```
#(config-ext-nacl) permit tcp 192.168.0.0 0.0.0.255 10.0.0.0
0.0.0.255 eq 22
```

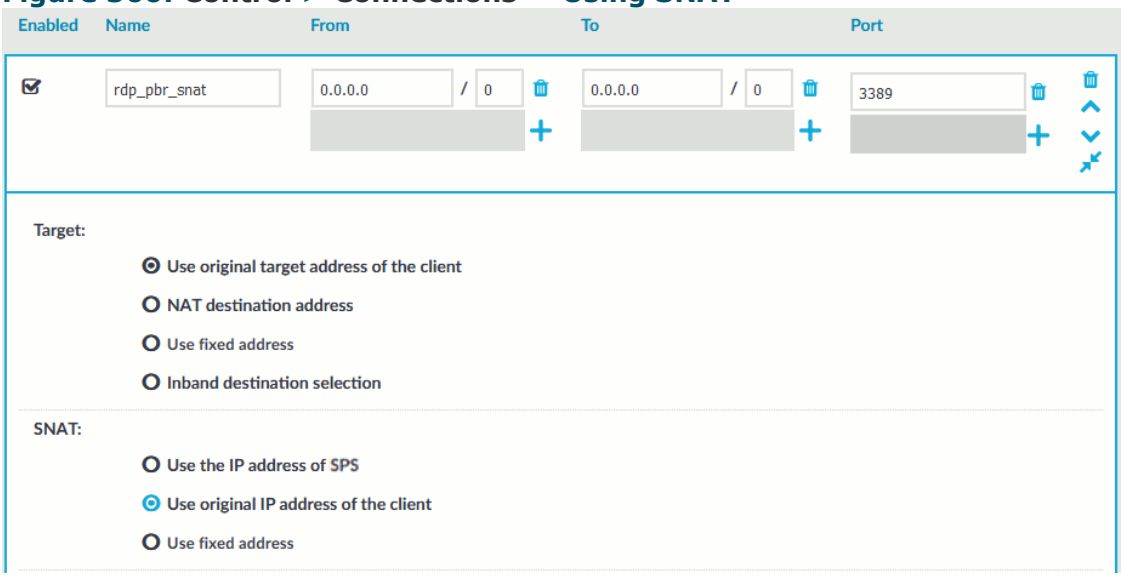
TIP: Cisco ACLs use inverse netmasks for defining network addresses. To calculate an inverse mask given a subnet mask, simply subtract each octet value from 255.

Create an ACL entry that matches the reply packets coming from the server zone and targeted at the client zone to make sure that replies are reaching the SPS.

```
#(config) ip access-list extended ssh-outbound
#(config-ext-nacl) permit tcp <dst net> <dst mask> eq <dst port> <src
net> <src mask>
```

NOTE: This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 360: Control > Connections — Using SNAT

2. 

Example: Configuring an ACL entry for reply packets with Cisco policy routing

In case of the example in step 1, the permit clause should be:

```
#(config-ext-nacl) permit tcp 10.0.0.0 0.0.0.255 eq 22
192.168.0.0 0.0.0.255
```

3. Create a route-map entry. It controls which packets are affected by policy routing and where they should be forwarded to. The `match` commands specify the conditions under which policy routing occurs. The `set` commands specify the routing actions to perform if the criteria enforced by the `match` commands are met. A new route-map can be defined as follows:

   ```
   #(config) route-map scb-inbound
   ```

   a. Set your route-map to match the traffic in ACL `ssh-inbound`:

      ```
      #(config-route-map) match ip address ssh-inbound
      ```

   b. Set an action on the matching traffic. Define a next-hop entry to redirect the traffic to the SPS.

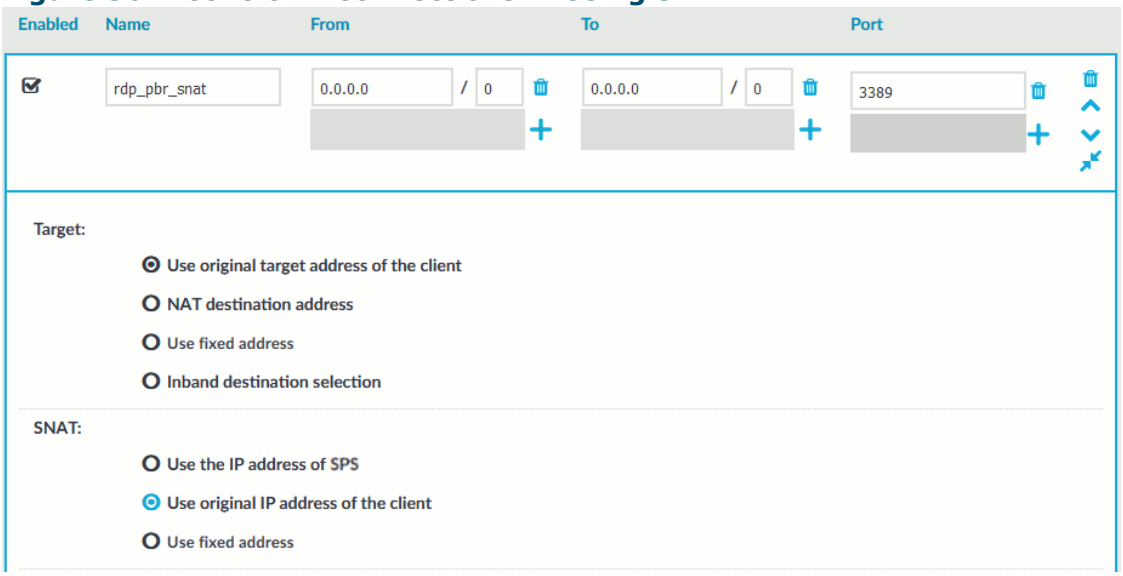
      ```
      #(config-route-map) set ip next-hop <SPS IP address>
      ```

4. Create another route-map that controls the reply packet flow.

   ```
   #(config) route-map scb-outbound
   #(config-route-map) match ip address ssh-outbound
   #(config-route-map) set ip next-hop <SPS IP address>
   ```

   NOTE: This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 361: Control > Connections — Using SNAT



5. Apply the route-map to the appropriate interfaces.
   a. First, add the ssh-inbound route-map entry to the interface facing the clients:

      ```
      #(config) interface <interface-facing-the-clients>
      #(config-if) ip policy route-map scb-inbound
      ```

   b. Then add the ssh-outbound route-map entry to the interface facing the servers:

      ```
      #(config) interface <interface-facing-the-servers>
      #(config-if) ip policy route-map scb-outbound
      ```

Expected result

The traffic from the clients targeting the specified port of the servers is redirected to SPS. Therefore, SPS can be configured to control and audit this traffic.

The full configuration for the above topology:

```
! interface facing the clients
interface FastEthernet0/0
 ip address 192.168.0.254 255.255.255.0
 ip policy route-map scb-inbound
 duplex full
 speed auto
 no mop enabled

! interface facing the SCB
interface FastEthernet0/1
 ip address 172.16.0.254 255.255.255.0
```

```
 duplex full
 speed auto
 no mop enabled

! interface facing the servers
interface FastEthernet1/0
 ip address 10.0.0.254 255.255.255.0
 ip policy route-map scb-outbound
 duplex full
 speed auto
 no mop enabled

! access lists matching the server and client subnets and the SSH port -
incoming packets
ip access-list extended ssh-inbound
 permit tcp 192.168.0.0 0.0.0.255 10.0.0.0 0.0.0.255 eq 22
! access lists matching the server and client subnets and the SSH port -
reply packets
ip access-list extended ssh-outbound
 permit tcp 10.0.0.0 0.0.0.255 eq 22 192.168.0.0 0.0.0.255

! policy routing entry matching on the incoming SSH connections and
! redirecting them to the SCB external interface
route-map scb-inbound permit 10
 match ip address ssh-inbound
 set ip next-hop 172.16.0.1

! the following part is only required for SNAT-based SCB configuration
! policy routing entry matching on the SSH reply packets and
! redirecting them to the SCB external interface
route-map scb-outbound permit 10
 match ip address ssh-outbound
 set ip next-hop 172.16.0.1
```

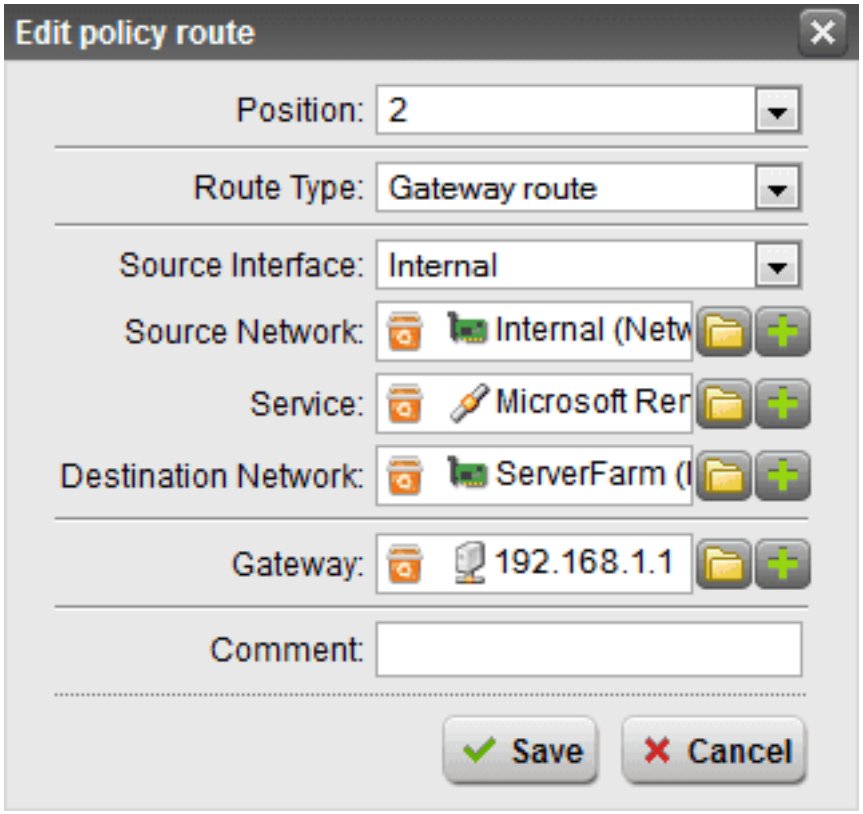
Configuring advanced routing on Sophos UTM (formerly Astaro Security Gateway) firewalls

The following describes how to configure a Sophos UTM firewall to redirect selected traffic to One Identity Safeguard for Privileged Sessions (SPS) instead of its original destination. Interface 1 will be referred to as 'Internal' and Interface 2 will be referred to as 'ServerFarm'.

To configure a Sophos UTM firewall to redirect selected traffic to SPS instead of its original destination

1. On the **Policy Routes** tab of the Sophos UTM firewall, click **New Policy Route**.

2. **Figure 362: New Policy Route**

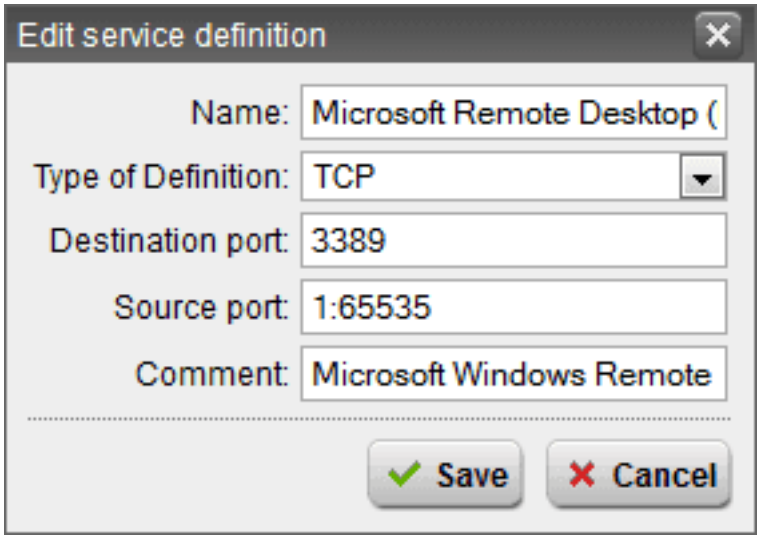


In the dialog box, enter the following settings:

- **Position**: Set the position number, defining the priority of the policy route. Lower numbers have higher priority. Routes are matched in ascending order. Once a route has been matched, routes with a higher number will not be evaluated anymore.
- **Route Type**: Select **`Gateway route`**. Packets will be sent to a particular host (gateway).
- **Source Interface**: Select **`Internal`**. This is the interface where the data packet to be routed arrives from.
- **Source Network**: Select **`Internal (Network)`**. This is the source network of the data packets to be routed.
- **Service**: Select **`Microsoft Remote Desktop Protocol`**. This is the service definition that matches the data packet to be routed.

- **Destination Network**: Select `ServerFarm (Network)`. This is the destination network of the data packets to be routed.
- **Gateway**: Select the IP address of SPS. This is the router where the gateway will forward data packets to.
- **Comment**: Optionally, enter a description or other information.

3. Click **Save**.
4. Click the status icon to activate the route.
5. Navigate to **Definitions & Users** > **Service Definitions** and click **New Service Definition**.
6. **Figure 363: New Service Definition**



In the dialog box, enter the following settings. It will ensure that the policy will apply to all TCP/3389:

- **Name**: Enter a descriptive name for the definition (for example Microsoft Remote Desktop Protocol).
- **Type of Definition**: Select `TCP`. This is the service type.

  NOTE: The definition type cannot be changed after saving the definition. To change the definition type, delete the service definition and create a new one with the desired settings.
- **Destination port**: Enter `3389`. This is the destination port that can either be entered as a single port number (for example `80`), or as a port range, using a colon as delimiter (for example 1024:64000).
- **Source port**: Enter `1:65535`. This is the source port that can either be entered as a single port number (for example `80`), or as a port range, using a colon as delimiter (for example 1024:64000).
- **Comment**: Optionally, enter a description or other information.

7. Click **Save**. The new definition appears in the service definition list.

   With this step, the client-server routing is configured.

8. To configure the server-client routing, create another policy route, and In the dialog box, enter the following settings:

   - **Position**: Set the position number, defining the priority of the policy route. Lower numbers have higher priority. Routes are matched in ascending order. Once a route has been matched, routes with a higher number will not be evaluated anymore.
   - **Route Type**: Select **`Gateway route`**. Packets will be sent to a particular host (gateway).
   - **Source Interface**: Select **`ServerFarm`**. This is the interface where the data packet to be routed arrives from.
   - **Source Network**: Select **`ServerFarm (Network)`**. This is the source network of the data packets to be routed.
   - **Service**: Select **`3389`**. This is the service definition that matches the data packet to be routed.
   - **Destination Network**: Select **`Internal (Network)`**. This is the destination network of the data packets to be routed.
   - **Gateway**: Select the IP address of SPS. This is the router where the gateway will forward data packets to.
   - **Comment**: Optionally, enter a description or other information.

Using SCP with agent-forwarding

When the client uses SSH to access a target server via One Identity Safeguard for Privileged Sessions (SPS) and authenticates with the public keys, the SPS Authentication Policy has **Public key** > **Agent** configured on the server-side. If the client supports agent-forwarding, this works well. However, scp does not: it always adds the `-a` option to the command-line to disable agent-forwarding. Explicitly allowing agent-forwarding with the `-A` or the `-oForwardAgent yes` command-line option, or writing `ForwardAgent yes` into the configuration has no effect, because the implicit `-a` at the end of the command-line takes precedence.

Solution 1: Use a wrapper script

The scp application can be started with the `-S` option to use an external application to create the encrypted connection. On Linux and UNIX platforms, this external application can be, for example, the following script that removes the unnecessary option from the scp command line.

```
#!/usr/bin/perl
exec '/usr/bin/ssh', '-A', map {$_ eq '-oForwardAgent=no' ? ( ) : $_} @ARGV
```

If you want your clients to use this script transparently, you can create an alias for it with the following command:

```
alias scp='scp -S <path-to-the-script-on-the-client>'
```

Solution 2: Use ssh master-channels

This solution relies on sending scp through an SSH master-control channel. In this case, scp does not need agent-forwarding, because it is already performed during the ControlMaster setup. The advantage of this solution is that the scp connection is setup quickly, because no authentication is needed, since the connection is already open. The disadvantage is that first a ControlMaster connection must be opened to the target host using the following command:

```
ssh -M -S /tmp/<address-of-the-target-server> <address-of-the-target-server>
```

When staring scp, reference the control path created with the previous command:

```
scp -oControlPath=/tmp/<address-of-the-target-server> [[user@]host1:]file1 ...
[[user@]host2:]file2
```

Solution 3: Patch the scp source

You can simply patch the scp source to overcome the problem, but then you need to recompile and re-install scp on every platform you use in your environment. The following is a sample patch for openssh-5.6p1:

```
--- scp-org.c    2010-07-02 05:37:33.000000000 +0200
+++ scp-new.c    2010-09-08 17:56:33.000000000 +0200
@@ -339,7 +339,6 @@
     args.list = NULL;
     addargs(&args, "%s", ssh_program);
     addargs(&args, "-x");
-    addargs(&args, "-oForwardAgent no");
     addargs(&args, "-oPermitLocalCommand no");
     addargs(&args, "-oClearAllForwardings yes");
```

Solution 4: Use fix or mapped keys on server-side

This is not agent-forwarding anymore, but scp still can use keys. Instead of passing the user-keys to the target server, SPS can authenticate on the server using a fix key, or a separate key for every user. Setting the server-side keys on SPS (or fetching them from LDAP), has the following advantages:

- The user cannot bypass SPS and directly connect to the target server
- Key-handling in the server environment becomes much simpler, because you do not have to import the user-keys to every host (if this is done locally, without a central identity management system)

For details on configuring server-side keys on SPS, see Relayed authentication methods on page 630.

Solution 5: WinSCP and agent-forwarding

WinSCP is a common tool for Windows to transfer files using SFTP/SCP. To use agent-forwarding in WinSCP, enable it in the **SSH** > **Authentication** options and load your keys.

Security checklist for configuring One Identity Safeguard for Privileged Sessions (SPS)

The following checklist is a set of recommendations and configuration best practices to ensure that your One Identity Safeguard for Privileged Sessions (SPS) is configured securely.

Encryption-related settings

- TIP: One Identity recommends using 2048-bit RSA keys (or stronger).
- Use strong passwords: at least 8 characters that include numbers, letters, special characters, and capital letters. For local One Identity Safeguard for Privileged Sessions (SPS) users, require the use of strong passwords (set **Users & Access Control** > **Login options** > **Minimal password strength** to `strong`). For more information, see *Setting password policies for local users* in the *Administration Guide*.
- When exporting the configuration of SPS, or creating configuration backups, always use encryption. Handle the exported data with care, as it contains sensitive information, including credentials. For details on encrypting the configuration, see *Encrypting configuration backups with GPG* in the *Administration Guide*.
- Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates (for example, do not use the certificate of the One Identity Safeguard for Privileged Sessions (SPS) webserver to encrypt audit trails, or the same keypair for signing and encrypting data).
- Do not use the CBC block cipher mode, or the `diffie-hellman-group1-sha1` key exchange algorithm. For details, see *Supported encryption algorithms* in the *Administration Guide*.
- Always encrypt your audit trails to protect sensitive data. For details, see *Encrypting audit trails* in the *Administration Guide*.

Connection policies

- When configuring connection policies, always limit the source of the connection to the client network that requires access to the connection.
- Always use gateway authentication to authenticate clients. Do not trust the source IP address of a connection, or the result of server authentication.
- To prevent Denial of Service (DoS) attacks against One Identity Safeguard for Privileged Sessions (SPS), set the **Connection rate limit** option of your connection policies. For details, see *Configuring connections* in the *Administration Guide*.
- Configure your RDP connection policies to use strong encryption. To enable SSL-encryption for the RDP protocol, see *Enabling TLS-encryption for RDP connections* in the *Administration Guide*.
- In RDP connections, if the client uses the Windows login screen to authenticate on the server, the password of the client is visible in the audit trail. To avoid displaying the password when replaying the audit trail, encrypt the upstream traffic in the audit trail using a separate certificate from the downstream traffic.

  For more information, see *Encrypting audit trails* in the *Administration Guide*.
- Ensure that host key verification is enabled in SSH connection policies. That is, the **Server side host key settings** > **Allow plain host keys** and **Server side host key settings** > **Allow X.509 host certificates** options do not have the **No check required** option selected. For more information, see *Setting the SSH host keys of the connection* in the *Administration Guide*.

Appliance access

- Accessing the One Identity Safeguard for Privileged Sessions (SPS) host directly using SSH is not recommended or supported, except for troubleshooting purposes. In such case, the One Identity Support Team will give you exact instructions on what to do to solve the problem.

  For security reasons, disable SSH access to SPS when it is not needed. For details, see *Enabling SSH access to the One Identity Safeguard for Privileged Sessions (SPS) host* in the *Administration Guide*.
- Permit administrative access to SPS only from trusted networks. If possible, monitored connections and administrative access to the SPS web interface should originate from separate networks.
- Configure SPS to send an alert if a user fails to login to SPS. For details, see the **Login failed** alert in *System related traps* in the *Administration Guide*.
- Configure **Disk space fill-up prevention**, and configure SPS to send an alert if the free space on the disks of SPS is low. For details, see *Preventing disk space fill-up* in the *Administration Guide*.

Networking considerations

- One Identity Safeguard for Privileged Sessions (SPS) stores sensitive data. Use a firewall and other appropriate controls to ensure that unauthorized connections cannot access it.
- If possible, enable management access to SPS only from trusted networks.
- Make sure that the HA interface of SPS is connected to a trusted network.

Jumplists for in-product help

To find the documentation for a specific UI element, browse the following sections.

Basic Settings > Management

- **Basic Settings** > **Management** > **Syslog**: For details, see Configuring system logging on page 125.
- **Basic Settings** > **Management** > **SNMP trap settings**: For details, see Configuring SNMP alerts on page 129.
- **Basic Settings** > **Management** > **Mail settings**: For details, see Configuring e-mail alerts on page 127.
- **Basic Settings** > **Management** > **Web interface timeout**: For details, see Web interface timeout on page 328.
- **Basic Settings** > **Management** > **Change root password**: For details, see Changing the root password of One Identity Safeguard for Privileged Sessions (SPS) on page 446.
- **Basic Settings** > **Management** > **System backup**: For details, see:
  - Creating configuration backups on page 155
  - Encrypting configuration backups with GPG on page 157
- **Basic Settings** > **Management** > **Verbose system logs**: For details, see Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS) on page 938.
- **Basic Settings** > **Management** > **SSL certificates**: For details, see Managing the certificates used on One Identity Safeguard for Privileged Sessions (SPS) on page 459.
- **Basic Settings** > **Management** > **Core files**: For details, see Gathering data about system problems on page 936.
- **Basic Settings** > **Management** > **Disk space fill-up prevention**: For details, see Preventing disk space fill-up on page 138.
- **Basic Settings** > **Management** > **Web gateway authentication**: For details, see Configuring out-of-band gateway authentication on page 863.

Basic Settings > Local Services

- **Basic Settings** > **Local Services** > **SSH server**: For details, see Enabling SSH access to the One Identity Safeguard for Privileged Sessions (SPS) host on page 444.
- **Basic Settings** > **Local Services** > **Web login (admin and user)**: For details, see Configuring user and administrator login addresses on page 117.
- **Basic Settings** > **Local Services** > **Web login (user only)**: For details, see Configuring user and administrator login addresses on page 117.
- **Basic Settings** > **Local Services** > **SNMP server settings**: For details, see Querying SPS status information using agents on page 132.
- **Basic Settings** > **Local Services** > **Indexer service**: For details, see Configuring the internal indexer on page 675.
- **Basic Settings** > **Local Services** > **Privileged Account Analytics**: Select this option only if you are also using One Identity Safeguard for Privileged Analytics.

  To enable One Identity Safeguard for Privileged Analytics and analyze the behavior of your users, One Identity Safeguard for Privileged Sessions (SPS) requires a special license. Also, depending on the number of your users and sessions, the performance and sizing of SPS must be considered. If you are interested in One Identity Safeguard for Privileged Analytics, contact our Sales Team, or your One Identity representative. For details on One Identity Safeguard for Privileged Analytics, see the One Identity One Identity Safeguard for Privileged Analytics website. For details on enabling One Identity Safeguard for Privileged Analytics, see *Safeguard for Privileged Analytics Configuration Guide*.
- **Basic Settings** > **Local Services** > **Cluster Interface**: This option is related to an experimental feature that will allow you to manage and synchronize the configuration of multiple SPS appliances from a central server. If you are interested in this feature, contact our Support Team.

Basic Settings > System

- **Basic Settings** > **System** > **System control**: For details, see Controlling One Identity Safeguard for Privileged Sessions (SPS): reboot, shutdown on page 394.
- **Basic Settings** > **System** > **Traffic control**: For details, see Disabling controlled traffic on page 395.
- **Basic Settings** > **System** > **Version details**: For details, see Upgrading One Identity Safeguard for Privileged Sessions (SPS) (single node) on page 434.
- **Basic Settings** > **System** > **Export configuration**: For details, see Exporting the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 438.
- **Basic Settings** > **System** > **Import configuration**: For details, see Importing the configuration of One Identity Safeguard for Privileged Sessions (SPS) on page 439.

- **Basic Settings** > **System** > **License**: For details, see Managing the One Identity Safeguard for Privileged Sessions (SPS) license on page 440.
- **Basic Settings** > **System** > **Sealed mode**: For details, see Sealed mode on page 449.
- **Basic Settings** > **System** > **Firmwares**: For details, see Upgrading One Identity Safeguard for Privileged Sessions (SPS) (single node) on page 434.

<Protocol name> Control > Global Options

- **Traffic Controls** > **Protocol name** > **Connections** > **Global Options** > **Traffic**: For details, see Changing log verbosity level of One Identity Safeguard for Privileged Sessions (SPS) on page 938.
- **Traffic Controls** > **Protocol name** > **Connections** > **Global Options** > **Audit** > **Timestamping**: For details, see:
  - Digitally signing audit trails on page 519
  - Timestamping audit trails with external timestamping service on page 517
  - Timestamping audit trails with built-in timestamping service on page 515

Configuring SPS to use an LDAP backend

To configure SPS to use an LDAP backend, read the following sections.

Overview

Access control in One Identity Safeguard for Privileged Sessions (SPS) is based on groups. Whenever a user needs to access a protected resource, like navigating to a configuration page on the SPS Web Interface, or opening a channel in a connection, SPS checks the access control list associated with the resource in question.

The access control lists grant access to groups. Therefore, SPS needs to determine which groups the user is a member of to evaluate the access rules.

When you configure SPS to use an LDAP backend, SPS will:

1. Identify the user. For more information, see *User identification* below.
2. Determine the relevant groups the user is a member of. For more information, see *Group membership resolution* below.

User identification

SPS works with plain usernames, for example, `administrator`. This must be unambiguously resolved to an LDAP user object in order to determine the user's groups. If a user identification returns multiple results, SPS treats this as an error, and access to the user in question is denied.

Only the user object returned in this phase is used for group membership checks, and *not* the original plain username.

User resolution depends on the type of the backend (POSIX or Active Directory).

For more information, see the backend-specific sections below.

Group membership resolution

SPS works with plain group names, for example, `superusers`. For group membership checks, SPS looks up a relevant group object in LDAP and checks if the user object returned during user identification is a member of that group. Since some of the group object's attributes are always used for group membership checks, the group object must also exist in LDAP.

Group membership resolution depends on the LDAP backend type.

For more information, see the backend-specific sections below.

Common to all backends

All backends have configurable parameters relevant for user identification and group membership:

- `bind_dn` and `bind_password`: *Bind DN* and *Bind password* are used for user identification and group membership check during authentication to the LDAP database. If you leave it empty, One Identity Safeguard for Privileged Sessions (SPS) will try to bind anonymously.
- `user_base_dn`: *User Base DN* is where SPS searches for users.
- `group_base_dn`: *Group Base DN* is where SPS searches for groups. Only groups under this base are considered for membership.
- `memberof_check`: the *Enable checking for group DNs in user objects* setting allows checking a configurable attribute in the user object. This attribute contains a list of group DNs the user is additionally a member of. This user attribute is usually `memberOf`. For more information, see the backend-specific sections below.
- `user_dn_in_groups`: *Check the user DN in these groups* is a list of additional group object classes and their respective attributes where SPS will look for member user DNs. For more information, see the backend-specific sections below.

All comparisons and searches are done by SPS in a way that plain user and group names are matched with attribute values by the LDAP server. As a result, user and group names are case insensitive if and only if the matching rule for the attribute in question is case insensitive in the LDAP database.

POSIX LDAP backend

In addition to the common parameters, the POSIX backend has the following configurable parameters:

- `username_attribute`: *Username (user ID) attribute name* is the name of the attribute in the user object, which contains the user's plain username.
- `membership_check`: *Enable POSIX group membership check* enables POSIX *primary* and *supplementary* group membership checking. When enabled, it has the following configurable parameter:
  - `member_uid_attribute`: the optional *POSIX group membership attribute name* is the name of the attribute in a **posixGroup** group object, which lists the plain usernames that are members of the group. These groups are usually referred to as *supplementary groups* of the referred user.

User identification in POSIX

To determine the user entry for a given plain username, One Identity Safeguard for Privileged Sessions (SPS) performs a search under `user_base_dn` for objects having the

`username_attribute` equal to the plain username of the user. The **objectClass** of the user object is not restricted.

The user object returned here is used for group membership checks.

Group membership resolution in POSIX

For all group membership checks, only the LDAP user object returned during user identification phase is used.

The plain group name is always compared to the **cn** attribute of the group object.

A user is treated as a member of a group given by its plain group name if the plain group name matches the **cn** attribute of the group object, and *any* of the following is true:

- The group is the user's primary group. That is, the group is a **posixGroup**, and the user's **gidNumber** attribute is equal to the group's **gidNumber** attribute.

  This check is performed only when the `membership_check` option is enabled for POSIX.

  > NOTE:
  >
  > It is OK for the user to have no **gidNumber** attribute, in which case this check will be skipped.

- The group lists the user's short username. That is, the group is a **posixGroup**, and it's `member_uid_attribute` contains the short username from the user object.

  This check is performed only when the `membership_check` option is enabled, and the `member_uid_attribute` is configured.

  > NOTE:
  >
  > For the purpose of this check, the user's short username is retrieved from the user object's `username_attribute`. Currently, this attribute should only contain a single username. A warning will appear in the logs if this is not the case, and the first value of the attribute will be used as returned by the server. This is a known limitation.

- The group lists the user's **dn** in any of the additional group objects configured in `user_dn_in_groups`.

  For example, if a row is added with `objectClass` set to **groupOfNames** and `attribute` set to **member**, SPS will treat the user as a member of all groups where the group is a **groupOfNames**, and the group's **member** attribute contains the user's **dn**.

- The user lists the group's **dn**. That is, the user's `memberof_user_attribute` contains the **dn** of the group, and the **objectClass** of the referred group is `memberof_group_objectclass`.

  This check is performed only when the `memberof_check` option is enabled for POSIX.

  > NOTE:
  >
  > SPS compares the **dn** stored in the `memberof_user_attribute` to the **dn** of the group object itself in a strict stringwise manner. Therefore, the user attribute must

contain the group DN exactly as it would be returned by the LDAP server. No case or accent differences are allowed.

Active Directory LDAP backend

In addition to the common parameters, the Active Directory (AD) backend has the following additional configurable parameters:

- `membership_check`: *Enable AD group membership check* enables AD specific non-primary group membership checking.

  NOTE:

  The AD user's primary group is always checked regardless of this setting.

  - `nested_groups`: *Enable nested groups* allows AD nested group support. See below for details.

Additionally, AD supports case and accent insensitive matching in many of the user and group name attributes. Since One Identity Safeguard for Privileged Sessions (SPS) relies on the server to perform comparisons, case and accent insensitive user and group name support depends solely on the server configuration.

User identification in AD

To determine the user entry for a given plain username, SPS performs a search under `user_base_dn` for objects having either the **sAMAccountName** or the **userPrincipalName** equal to the plain username of the user. The **objectClass** of the user object is not restricted.

NOTE:

Although **userPrincipalName** in AD is a Internet-style name like *user@example.com*, it matches simple names like *user*.

Only the user object returned here is used for group membership checks.

Group membership resolution in AD

For all group membership checks, only the LDAP user object returned during user identification phase is used.

The plain group name is always compared to the **cn** attribute of the group object.

A user is treated as a member of a group if both the group object's **objectClass** and **objectCategory** is **group**, and *any* of the following is true:

- The group is the user's primary group. That is, the **objectSID** attribute of the group matches the Security Identifier calculated from the user object's **objectSID** and **primaryGroupID** attributes, as described in the Microsoft Support article How to

use the PrimaryGroupID attribute to find the primary group for a user.

NOTE: When using the AD backend, this check is always performed, even if the `membership_check` option is disabled. However, it is OK for the user to have no primary group.

- The group lists the user's short username. That is, the group's **memberUid** attribute contains the short username from the user object.

  This check is performed only when the `membership_check` option is enabled for AD.

  NOTE: For the purpose of this check, the user's short username is retrieved from the user object's **sAMAccountName** attribute only, which is a single-valued attribute in AD. This is a known limitation.

  It is OK for the **sAMAccountName** attribute to be missing, in which case this check will be skipped.

- The group lists the user's **dn**. That is, the group object's **member** attribute contains the user's **dn**.

  This check is performed only when the `membership_check` option is enabled for AD.

  This is the only place where *nested groups* are supported. When the `nested_groups` setting is enabled in the configuration, SPS will also find groups which do not directly contain the user's **dn** in their **member** attribute, but do contain an intermediate group's **dn**, which in turn contains the user **dn** in its **member** attribute. This nesting can be arbitrarily deep, limited only by AD.

  NOTE: Due to the nature of the way AD resolves the nested group chain, intermediate groups might be outside the configured `group_base_dn`.

NOTE: Although an **objectCategory** in AD is a DN-valued attribute, it does match simple names like **group**.

Additionally, a user is treated as a member of a group if:

- The group lists the user's **dn** in any of the additional group objects configured in `user_dn_in_groups`.

  For example, if a row is added with `objectClass` set to **groupOfNames** and `attribute` set to **member**, SPS will treat the user as a member of all groups where the group is a **groupOfNames**, and the group's **member** attribute contains the user's **dn**.

  NOTE: There is no additional restriction on the group's **objectClass** in this case.

- The user lists the group's **dn**. That is, the user's `memberof_user_attribute` contains the **dn** of the group, and the **objectClass** of the referred group is **group**.

  This check is performed only when the `memberof_check` option is enabled for AD.

  NOTE: SPS compares the **dn** stored in the `memberof_user_attribute` to the **dn** of the group object itself in a strict stringwise manner. Therefore, this user attribute must contain the group DN exactly as it would be returned by the LDAP server. No case or accent differences are allowed.

About us

One Identity solutions eliminate the complexities and time-consuming processes often required to govern identities, manage privileged accounts and control access. Our solutions enhance business agility while addressing your IAM challenges with on-premises, cloud and hybrid environments.

Contacting us

For sales and other inquiries, such as licensing, support, and renewals, visit https://www.oneidentity.com/company/contact-us.aspx.

Technical support resources

Technical support is available to One Identity customers with a valid maintenance contract and customers who have trial versions. You can access the Support Portal at https://support.oneidentity.com/.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos at www.YouTube.com/OneIdentity
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product

#

4-eyes authorization
4-eyes authorization is an advanced authorization method where only two administrators logging in simultaneously are permitted to access the server. These administrators can monitor each other's work, reducing the chance of (accidental or intentional) human errors in the server administration process.

A

access policy
Collection of access policies. Access policies define who can authorize and audit a connection.

alias IP
An additional IP address assigned to an interface that already has an IP address. The normal and alias IP addresses both refer to the same physical interface.

Audit Player
Audit Player is a desktop application that can replay recorded audit trails like movie. The Audit Player is available for the Microsoft Windows and GNU/Linux platforms.

Audit trail
An audit trail is a file storing the recorded activities of the administrators in an encrypted format. Audit trails can be replayed using the Audit Player application.

auditing policy
The auditing policy determines which events are logged on host running Microsoft Windows operating systems.

authentication
The process of verifying the authenticity of a user or client before allowing access to a network system or service.

Authentication Policy
An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the client can authenticate to the target server.

B

BOM
The byte order mark (BOM) is a Unicode character used to signal the byte-order of the message text.

BSD-syslog protocol
The old syslog protocol standard described in RFC 3164. Sometimes also referred to as the legacy-syslog protocol.

C

CA
A Certificate Authority (CA) is an institute that issues certificates.

certificate
A certificate is a file that uniquely identifies its owner. Certificates contains information identifying the owner of the certificate, a public key itself, the expiration date of the certificate, the name of the CA that signed the certificate, and some other data.

Channel Policy
The channel policy lists the SSH channels (for example terminal session, SCP, and so on) that can be used in a connection. The channel policy can further restrict access to each channel based on the IP address of the client or the server, a user list, or a time policy.

client mode
In client mode, syslog-ng collects the local logs generated by the host and forwards them through a network connection to the central syslog-ng server or to a relay.

Common Gateway Protocol (CGP)
Reliable connection is also known as Common Gateway Protocol (CGP). It makes reconnection possible to the server in case of a network failure. Default port number is 2598.

Connection Policy
Connection policies determine if a server can be accessed from a particular client. Connection policies reference other resources (policies, usergroups, keys) that must be configured and available before creating a connection policy.

controlled traffic
SPS audits and controls only the traffic that is configured in the connection and channel policies, all other traffic is forwarded on the packet level without any inspection.

D

destination
A named collection of configured destination drivers.

destination driver
A communication method used to send log messages.

destination, local
A destination that transfers log messages within the host, for example writes them to a file, or passes them to a log analyzing application.

destination, network
A destination that sends log messages to a remote host (that is, a syslog-ng relay or server) using a network connection.

disk buffer
The Premium Edition of syslog-ng can store messages on the local hard disk if the central log server or the network connection to the server becomes unavailable.

disk queue
See disk buffer.

domain name
The name of a network, for example: balabit.com.

E

embedded log statement
A log statement that is included in another log statement to create a complex log path.

F

filter
An expression to select messages.

firmware
A firmware is a collection of the software components running on SPS. Individual software components cannot be upgraded on SPS, only the entire firmware. SPS contains two firmwares, an external (or boot) firmware and an internal (or core) firmware. These can be upgraded separately.

fully qualified domain name (FQDN)
A domain name that specifies its exact location in the tree hierarchy of the Domain Name System (DNS). For example, given a device with a local hostname myhost and a parent domain name example.com, the fully qualified domain name is myhost.example.com.

G

gateway
A device that connects two or more parts of the network, for example: your local intranet and the external network (the Internet). Gateways act as entrances into other networks.

H

HA network interface
The HA interface (labeled 4 or HA) is an interface reserved for communication between the nodes of SPS clusters.

High Availability
High Availability (HA) uses a second SPS unit (called secondary node) to ensure that the services are available even if the first unit (called primary node) breaks down.

host
A computer connected to the network.

hostname
A name that identifies a host on the network.

I

ICA
The base protocol of Citrix products (default port tcp/1494). It does desktop or application remoting through TCP or other network protocols. Independent Computing Architecture (ICA) is a proprietary protocol for an application server system, designed by Citrix Systems. The protocol lays down a specification for passing data between server and clients, but is not bound to any one platform. ICA is broadly similar in purpose to window servers such as the X Window System. It also provides for the feedback of user input from the client to the server, and a variety of means for the server to send graphical output, as well as other media such as audio, from the running application to the client.

IETF-syslog protocol
The syslog-protocol standard developed by the Internet Engineering Task Force (IETF), described in RFC 5424-5427.

K

key pair
A private key and its related public key. The private key is known only to the owner, while the public key can be freely distributed. Information encrypted with the private key can only be decrypted using the public key.

L

LDAP
The Lightweight Directory Access Protocol (LDAP), is an application protocol for querying and modifying data using directory services running over TCP/IP.

License
SPS's license determines the number of servers (IP addresses) that SPS protects. The license limits the number of IP addresses accessible.

log path
A combination of sources, filters, parsers, rewrite rules, and destinations: syslog-ng examines all messages arriving to the sources of the logpath and sends the messages matching all filters to the defined destinations.

log source host
A host or network device (including syslog-ng clients and relays) that sends logs to the syslog-ng server. Log source hosts can be servers, routers, desktop computers, or other devices capable of sending syslog messages or running syslog-ng.

log statement
See log path.

logstore
A binary logfile format that can encrypt, compress, and timestamp log messages.

Long Term Supported release
Long Term Supported releases are major releases of that are supported for three years after their original release.

LSH
See log source host.

N

name server
A network computer storing the IP addresses corresponding to domain names.

node
An SPS unit running in High Availability mode.

O

Oracle Instant Client
The Oracle Instant Client is a small set of libraries, which allow you to connect to an Oracle Database. A subset of the full Oracle Client, it requires minimal installation but has full functionality.

output buffer
A part of the memory of the host where syslog-ng stores outgoing log messages if the destination cannot accept the messages immediately.

output queue
Messages from the output queue are sent to the target syslog-ng server. The syslog-ng application puts the outgoing messages directly into the output queue,

unless the output queue is full. The output queue can hold 64 messages, this is a fixed value and cannot be modified.

overflow queue
See output buffer.

P

parser
A set of rules to segment messages into named fields or columns.

ping
A command that sends a message from a host to another host over a network to test connectivity and packet loss.

port
A number ranging from 1 to 65535 that identifies the destination application of the transmitted data. For example: SSH commonly uses port 22, web servers (HTTP) use port 80, and so on.

primary node
The active SPS unit that is inspecting the traffic when SPS is used in High Availability mode.

PSM
An old abbreviation of Safeguard for Privileged Sessions (SPS).

Public-key authentication
An authentication method that uses encryption key pairs to verify the identity of a user or a client.

R

redundant Heartbeat interface
A redundant Heartbeat interface is a virtual interface that uses an existing interface of the SPS device to detect that the other node of the SPS cluster is still available. The virtual interface is not used to synchronize data between the nodes, only Heartbeat messages are transferred.

regular expression
A regular expression is a string that describes or matches a set of strings.

relay mode
In relay mode, syslog-ng receives logs through the network from syslog-ng clients and forwards them to the central syslog-ng server using a network connection.

Remote Desktop Gateway
Remote Desktop Gateway (RD Gateway) is a role service in the Remote Desktop Services server role that allows authorized remote users to connect to resources located on an internal or private network from any Internet-connected device. The

accessible resources can be terminal servers, remote applications, remote desktops, and so on. This service is also called Remote Desktop Gateway or RD Gateway.

rewrite rule
A set of rules to modify selected elements of a log message.

S

SCB
An old abbreviation of Safeguard for Privileged Sessions (SPS).

secondary node
The passive SPS unit that replaces the active unit (the primary node) if the primary node becomes unavailable.

server mode
In server mode, syslog-ng acts as a central log-collecting server. It receives messages from syslog-ng clients and relays over the network, and stores them locally in files, or passes them to other applications, for example, log analyzers.

SNMP
Simple Network Management Protocol (SNMP) is an industry standard protocol used for network management. SPS can send SNMP alerts to a central SNMP server.

source
A named collection of configured source drivers.

source driver
A communication method used to receive log messages.

source, local
A source that receives log messages from within the host, for example, from a file.

source, network
A source that receives log messages from a remote host using a network connection, for example, network(), syslog().

split brain
A split brain situation occurs when for some reason (for example, the loss of connection between the nodes) both nodes of an SPS cluster become active (primary) nodes. This might cause that new data (for example, audit trails) is created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data are created, which cannot be trivially merged.

SPS
Safeguard for Privileged Sessions

SSH settings
SSH settings determine the parameters of the connection on the protocol level, including timeout value and greeting message of the connection, as well as the encryption algorithms used.

SSL
See TLS.

syslog-ng
The syslog-ng application is a flexible and highly scalable system logging application, typically used to manage log messages and implement centralized logging.

syslog-ng agent
The syslog-ng Agent for Windows is a commercial log collector and forwarder application for the Microsoft Windows platform. It collects the log messages of the Windows-based host and forwards them to a syslog-ng server using regular or SSL-encrypted TCP connections.

syslog-ng client
A host running syslog-ng in client mode.

syslog-ng Premium Edition
The syslog-ng Premium Edition is the commercial version of the open-source application. It offers additional features, like encrypted message transfer and an agent for Microsoft Windows platforms.

syslog-ng relay
A host running syslog-ng in relay mode.

syslog-ng server
A host running syslog-ng in server mode.

T

template
A user-defined structure that can be used to restructure log messages or automatically generate file names.

Time Policy
The time policy determines which hours of a day can the users access a connection or a channel.

TLS
Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which provide secure communications on the Internet. The application can encrypt the communication between the clients and the server using TLS to prevent unauthorized access to sensitive log messages.

traceroute
A command that shows all routing steps (the path of a message) between two hosts.

U

UNIX domain socket
A UNIX domain socket (UDS) or IPC socket (inter-procedure call socket) is a virtual socket, used for inter-process communication.

User List
User lists are white- or blacklists of usernames that allow fine-control over who can access a connection or a channel.