



syslog-ng Store Box 6.9.0

Administration Guide

Copyright 2021 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 Web site (<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. 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.**

Contents

Preface	12
Introduction	13
What SSB is	13
What SSB is not	14
Why is SSB needed	14
Who uses SSB	15
The concepts of SSB	17
The philosophy of SSB	17
Collecting logs with SSB	18
Managing incoming and outgoing messages with flow-control	20
Receiving logs from a secure channel	21
Advanced Log Transfer Protocol	22
Network interfaces	22
High Availability support in SSB	24
Firmware in SSB	24
Firmware and high availability	25
Versions and releases of SSB	25
Licensing model and modes of operation	26
Notes about counting the licensed hosts	27
Licensing benefits	27
License types	28
Perpetual license	28
Subscription-based license	28
Licensing examples	29
The structure of a log message	31
BSD-syslog or legacy-syslog messages	31
The PRI message part	32
The HEADER message part	33
The MSG message part	33
IETF-syslog messages	33
The PRI message part	34

The HEADER message part	36
The STRUCTURED-DATA message part	36
The MSG message part	37
The Welcome Wizard and the first login	38
The initial connection to SSB	38
Creating an alias IP address (Microsoft Windows)	39
Creating an alias IP address (Linux)	45
Modifying the IP address of SSB	45
Configuring SSB with the Welcome Wizard	46
Configuring SSB as a standalone unit, or as the primary node of a HA cluster	47
Configuring your SSB unit as the secondary node of a HA cluster	61
Basic settings	65
Supported web browsers	65
The structure of the web interface	66
Elements of the main workspace	69
Multiple web users and locking	71
Web interface and RPC API settings	71
Network settings	72
Configuring the external interface	76
Configuring the management interface	77
Configuring the routing table	78
Date and time configuration	79
Configuring a time (NTP) server	80
SNMP and e-mail alerts	81
Configuring e-mail alerts	81
Configuring SNMP alerts	82
Querying SSB status information using agents	84
View node ID and community	86
Configuring system monitoring on SSB	87
Configuring SNMP agent settings	87
Health monitoring	88
Preventing disk space fill up	90
Configuring message rate alerting	91
System related traps	94

Alerts related to syslog-ng	96
Data and configuration backups	98
Creating a backup policy using Rsync over SSH	99
Creating a backup policy using SMB/CIFS	103
Creating a backup policy using NFS	106
Creating configuration backups	108
Creating data backups	109
Encrypting configuration backups with GPG	110
Archiving and cleanup	110
Creating a cleanup policy	112
Creating an archive policy using SMB/CIFS	112
Creating an archive policy using NFS	114
Archiving or cleaning up the collected data	116
User management and access control	117
Managing SSB users locally	117
Creating local users in SSB	118
Deleting a local user from SSB	119
Setting password policies for local users	119
Managing local usergroups	121
Managing SSB users from an LDAP database	122
Authenticating users to a RADIUS server	125
Managing user rights and usergroups	127
Assigning privileges to usergroups for the SSB web interface	128
Modifying group privileges	128
Finding specific usergroups	129
How to use usergroups	130
Built-in usergroups of SSB	130
Listing and searching configuration changes	132
Managing SSB	134
Controlling SSB: restart, shutdown	134
Managing a high availability SSB cluster	135
Adjusting the synchronization speed	140
Asynchronous data replication	140
Redundant heartbeat interfaces	141

Next-hop router monitoring	143
Upgrading SSB	145
Upgrade checklist	145
Upgrading SSB	147
Upgrading an SSB cluster	147
Troubleshooting	148
Updating the SSB license	148
Exporting the configuration of SSB	150
Importing the configuration of SSB	151
Accessing the SSB console	152
Using the console menu of SSB	152
Enabling SSH access to the SSB host	154
Changing the root password of SSB	155
Sealed mode	155
Disabling sealed mode	156
Out-of-band management of SSB	156
Configuring the IPMI interface from the console	158
Configuring the IPMI interface from the BIOS	159
Managing the certificates used on SSB	164
Generating certificates for SSB	167
Uploading external certificates to SSB	167
Generating TSA certificate with Windows Certificate Authority on Windows Server 2008	171
Generating TSA certificate with Windows Certificate Authority on Windows Server 2012	176
Creating hostlist policies	189
Creating hostlists	190
Importing hostlists from files	191
Configuring message sources	193
Default message sources in SSB	193
Creating new message sources in SSB	194
Configuring your own, customized Syslog type message source	195
Configuring the Listening address and Listening port for your Syslog type message source	197
Configuring the Transport options for your Syslog type message source	198

Configuring the Hostname and timestamp-related settings for your Syslog type message source	200
Configuring the Monitoring settings for your Syslog type message source	202
Customizing encoding for your Syslog type message source	203
Configuring your own, customized SQL type message source	203
Setting up and testing the SQL database connection of your SQL type message source	206
Customizing fetching messages when using your SQL type message source	207
Configuring the fetching frequency settings for your SQL type message source	212
Configuring the Monitoring settings for your SQL type message source	213
Receiving SNMP messages	214
Storing messages on SSB	215
Using logstores	216
Creating logstores	217
Configuring the indexer service	220
Viewing encrypted logs with logcat	222
Creating text logspaces	222
Managing logspaces	226
Managing logspaces - Archive and backup logspaces	227
Assigning the SSB logspace of your choice to a custom cloud service provider data disk	229
Creating filtered logspaces	233
Creating remote logspaces	235
Creating multiple logspaces	236
Accessing log files across the network	239
Sharing log files in standalone mode	239
Sharing log files in Domain mode	241
Accessing shared files	243
Managing custom cloud service provider data disks for your logspaces in SSB	245
Possible use cases and scenarios for using custom cloud service provider data disks with SSB	246
Adding a new custom cloud service provider data disk to your SSB configuration ...	246
Adding a new custom cloud service provider data disk on the cloud service provider side	249
Adding a new custom cloud service provider data disk to your SSB configuration on the SSB side	253

Removing a custom cloud service provider data disk from your SSB configuration ..	261
Increasing the size of a custom cloud service provider data disk that you use in your SSB configuration ..	263
Forwarding messages from SSB ..	266
Forwarding log messages to SQL databases ..	266
SQL templates in SSB ..	270
The Legacy template ..	270
The Full template ..	271
The Custom template ..	271
Forwarding log messages to remote servers ..	271
Forwarding log messages to the Microsoft Azure Sentinel cloud ..	275
Prerequisites ..	275
Limitations ..	276
Configuring the Azure Sentinel destination: adding a new Azure Sentinel destination ..	276
Configuring the Azure Sentinel destination: Authentication and workspace settings ..	278
Configuring the Azure Sentinel destination: Advanced message parameters ..	279
Configuring the Azure Sentinel destination: Performance-related settings ..	280
Forwarding log messages to Google Pub/Sub ..	282
Prerequisites ..	282
Limitations ..	283
Configuring the Google Pub/Sub destination: adding a new Google Pub/Sub destination ..	284
Configuring the Google Pub/Sub destination: Authentication and workspace settings ..	285
Configuring the Google Pub/Sub destination: Advanced message parameters ..	287
Configuring the Google Pub/Sub destination: Performance-related settings ..	288
Forwarding log messages to Splunk ..	290
Prerequisites ..	291
Limitations ..	291
Transport settings for the Splunk destination ..	292
HTTP connection settings ..	294
HTTPS connection settings ..	295
JSON message body ..	297
Performance-related settings ..	298

Forwarding log messages to HDFS destinations	300
Configuring a Kerberos policy	301
Configuring the HDFS cluster	302
Configuring an HDFS destination	303
Log paths: routing and processing messages	306
Default logpaths in SSB	306
Creating new log paths	307
Filtering messages	309
Replace message parts or create new macros with rewrite rules	312
Find and replace the text of the log message	314
Parsing sudo log messages	316
Parsing key-value pairs	317
Configuring syslog-ng options	321
General syslog-ng settings	321
Time stamping configuration on SSB	323
Using name resolution on SSB	323
Setting the certificates used in TLS-encrypted log transport	325
Searching log messages	330
Using the search interface	330
Customizing columns of the log message search interface	337
Metadata collected about log messages	338
Using complex search queries	338
Browsing encrypted logspaces	345
Using persistent decryption keys	346
Using session-only decryption keys	348
Assigning decryption keys to a logstore	349
Creating custom statistics from log data	350
Displaying log statistics	351
Creating reports from custom statistics	353
Creating content-based alerts	354
Setting up alerts on the search interface	354
Setting up alerts on the Search > Content-Based Alerts page	357
Format of alert messages	358
Additional tools	359

Searching the internal messages of SSB	360
Using the internal search interfaces	361
Filtering	362
Exporting the results	363
Customizing columns of the internal search interfaces	363
Changelogs of SSB	364
Configuration changes of syslog-ng peers	366
Log message alerts	366
Notifications on archiving and backups	367
Status history and statistics	368
Displaying custom syslog-ng statistics	370
Statistics collection options	370
Reports	371
Contents of the default reports	372
Generating partial reports	373
Configuring custom reports	373
Classifying messages with pattern databases	376
The structure of the pattern database	377
How pattern matching works	379
Searching for rulesets	380
Creating new rulesets and rules	381
Exporting databases and rulesets	383
Importing pattern databases	383
Using pattern parsers	383
Using parser results in filters and templates	386
Using the values of pattern parsers in filters and templates	387
The SSB RPC API	389
Requirements for using the RPC API	389
RPC client requirements	389
Documentation of the RPC API	390
Monitoring SSB	391
Monitoring SSB's disk	391
Monitoring SSB's memory	394
Monitoring SSB's CPU	395

Monitoring CPU load	395
Monitoring CPU load averages	395
Monitoring CPU usage	396
Monitoring SSB's I/O	398
Monitoring SSB statistics	401
Monitoring the HA cluster	404
Monitoring hardware RAID	407
StorCLI's PD LIST (Physical Drive)	410
StorCLI's Drive State	413
StorCLI's VD LIST (Virtual Drive)	415
Monitoring software RAID	418
Monitoring web server, root CA and TSA certificates	422
Troubleshooting SSB	424
Locating the SSB appliance in the server room	424
Network troubleshooting	425
Gathering data about system problems	426
Viewing logs on SSB	426
Collecting logs and system information for error reporting	427
Troubleshooting an SSB cluster	429
Understanding SSB cluster statuses	429
Recovering SSB if both nodes broke down	431
Recovering from a split brain situation	432
Replacing a node in an SSB HA cluster	435
Resolving an IP conflict between cluster nodes	436
Restoring SSB configuration and data	438
Configuring the IPMI interface from the BIOS after losing IPMI password	439
Incomplete TSA response received	444
Correct Alerting & Monitoring and Management group privilege mismatch	445
Security checklist for configuring SSB	447
About us	450
Contacting us	450
Technical support resources	450
Glossary	451

Preface

Welcome to the syslog-ng Store Box (SSB) 6.9.0 Administration Guide.

This document describes how to configure and manage SSB. Background information for the technology and concepts used by the product is also discussed.

Introduction

This chapter introduces the syslog-ng Store Box (SSB), discussing how and why it is useful, and what benefits it offers to an existing IT infrastructure.

NOTE: Due to complexity of deployment, configuration, and design, you may require assistance from [One Identity Professional Services](#) while introducing new or additional:

- sources
- destinations
- log paths
- significant increases in log volume.

One Identity Professional Services is equipped and trained to evaluate the needs of any organization, and to provide configuration and architectural recommendations that help our users get the most out of any SSB version.

One Identity Professional Services offer assistance in planning and scoping for current needs, as well as recommendations for the future to ensure success.

What SSB is

syslog-ng Store Box (SSB) is a device that collects, processes, stores, monitors, and manages log messages. It is a central log server appliance that can receive system (syslog and eventlog) log messages and Simple Network Management Protocol (SNMP) messages from your network devices and computers, store them in a trusted and signed logstore, automatically archive and back up the messages, and also classify the messages using artificial intelligence.

The most notable features of SSB are as follows:

- Secure log collection using Transport Layer Security (TLS).
- Trusted, encrypted, and time stamped storage.
- Ability to collect log messages from a wide range of platforms, including Linux, Unix, BSD, Sun Solaris, HP-UX, IBM AIX, IBM System i, as well as Microsoft Windows.
- Forwards messages to log analyzing engines.

- Classifies messages using customizable pattern databases for real-time log monitoring, alerting, and artificial ignorance.
- High Availability (HA) support to ensure continuous log collection in business-critical environments.
- Real-time log monitoring and alerting.
- Retrieves group memberships of the administrators and users from a Lightweight Directory Access Protocol (LDAP) database.
- Strict, yet easily customizable access control to grant users access only to selected log messages.
- Ability to search log data in multiple logspaces, whether on the same SSB appliance or located on a different appliance, even in a remote location.

SSB is configured and managed from any modern web browser that supports HTTPS connections, JavaScript, and cookies.

Supported browsers:

Mozilla Firefox 52 ESR

We also test SSB on the following, unsupported browsers. The features of SSB are available and usable on these browsers as well, but the look and feel might be different from the supported browsers. Internet Explorer 11, Microsoft Edge, and the currently available version of Mozilla Firefox and Google Chrome.

What SSB is not

The syslog-ng Store Box (SSB) appliance is not a log analyzing engine, though it can classify individual log messages using artificial ignorance. SSB comes with a built-in feature to store log message patterns that are considered "normal". Messages matching these patterns are produced during the legitimate use of the applications (for example sendmail, Postfix, MySQL, and so on), and are unimportant from the log monitoring perspective, while the remaining messages may contain something "interesting". The administrators can define log patterns on the SSB interface, label matching messages (for example, security event, and so on), and request alerts if a specific pattern is encountered. For thorough log analysis, SSB can also forward the incoming log messages to external log analyzing engines.

Why is SSB needed

Log messages contain information about the events happening on the hosts. Monitoring system events is essential for security and system health monitoring reasons. A well-established log management solution offers several benefits to an organization. It ensures that computer security records are stored in sufficient detail, and provides a simple way to

monitor and review these logs. Routine log reviews and continuous log analysis help to identify security incidents, policy violations, or other operational problems.

Logs also often form the basis of auditing and forensic analysis, product troubleshooting and support. There are also several laws, regulations and industrial standards that explicitly require the central collection, periodic review, and long-time archiving of log messages. Examples of such regulations are the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS).

Built around the popular [syslog-ng Premium Edition](#) (syslog-ng PE) application used by thousands of organizations worldwide, the syslog-ng Store Box (SSB) brings you a powerful, easy-to-configure appliance to collect and store your logs. Using the features of the latest syslog-ng PE to their full power, SSB allows you to collect, process, and store log messages from a wide range of platforms and devices.

All data can be stored in encrypted and optionally time stamped files, preventing any modification or manipulation, satisfying the highest security standards and policy compliance requirements.

Who uses SSB

The syslog-ng Store Box (SSB) appliance is useful for everyone who has to collect, store, and review log messages. In particular, SSB is invaluable for:

- **Central log collection and archiving:** SSB offers a simple, reliable, and convenient way of collecting log messages centrally. It is essentially a high-capacity log server with high availability support. Being able to collect logs from several different platforms makes it easy to integrate into any environment.
- **Secure log transfer and storage:** Log messages often contain sensitive information and also form the basis of audit trails for several applications. Preventing eavesdropping during message transfer and unauthorized access once the messages reach the log server is essential for security and privacy reasons.
- **Policy compliance:** Many organization must comply with regulations like the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS). These regulations often have explicit or implicit requirements about log management, such as the central collection of log messages, the use of log analysis to prevent and detect security incidents, or guaranteeing the availability of log messages for an extended period of time — up to several years. SSB helps these organizations to comply with these regulations.
- **Automated log monitoring and log pre-processing:** Monitoring log messages is an essential part of system-health monitoring and security incident detection and prevention. SSB offers a powerful platform that can classify tens of thousands of messages real-time to detect messages that deviate from regular messages, and promptly raise alerts. Although this classification does not offer as complete an inspection as a log analyzing application, SSB can process many more messages than

a regular log analyzing engine, and also filter out unimportant messages to decrease the load on the log analyzing application.

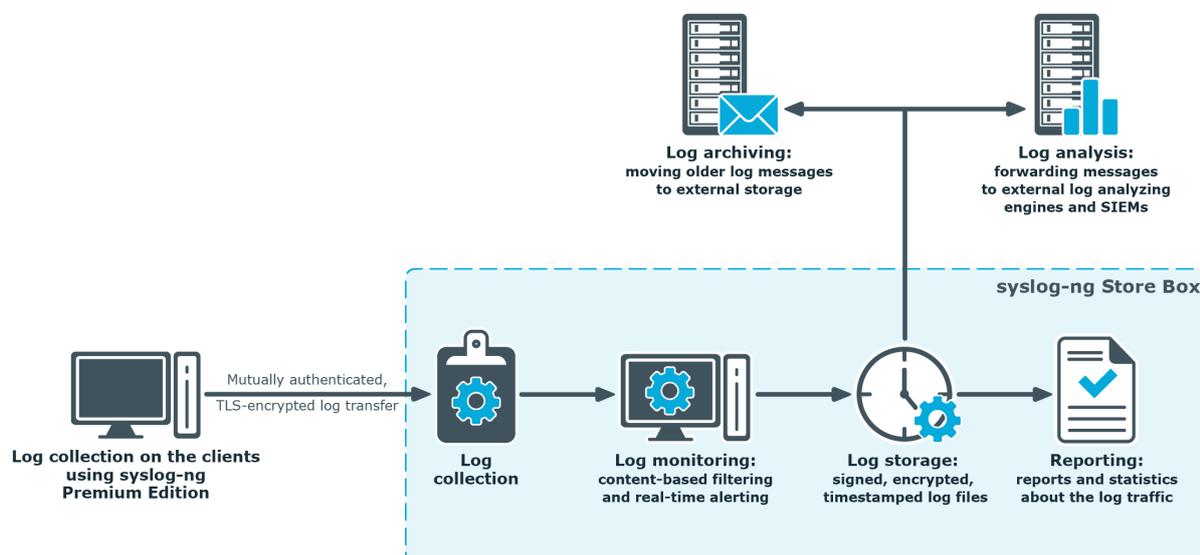
The concepts of SSB

This chapter discusses the technical concepts of syslog-ng Store Box (SSB).

The philosophy of SSB

The syslog-ng Store Box (SSB) appliance is a log server appliance that collects, stores and monitors log messages sent by network devices, applications and computers. SSB can receive traditional syslog messages, syslog messages that comply with the new Internet Engineering Task Force (IETF) standard ([RFC 5424-5428](#)), eventlog messages from Microsoft Windows hosts, as well as SNMP messages.

Figure 1: The philosophy of the syslog-ng Store Box



Clients can send messages to SSB using their own logging application if it supports the [BSD-syslog](#) (RFC 3164) or the [IETF-syslog](#) (RFC 5424-5428) protocol, or they can use the [syslog-ng Premium Edition](#) (syslog-ng PE) application to act as the log-forwarding agent of SSB.

The main purpose of SSB is to collect the logs from the clients and store them on its hard disk. The messages are stored in so-called *logspaces*. There are two types of logspaces: the first stores messages in traditional plain-text files, while the second one uses a binary format that can be compressed, encrypted, and time stamped.

You can also define multiple logspaces, remote logspaces, and configure filtered subsets of each logspace. A multiple logspace aggregates messages from multiple SSB appliances (located at different sites), allowing you to view and search the logs of several SSBs from a single web interface without having to log on to several different interfaces. Remote logspaces, on the other hand, enable you to access and search logspaces (including filtered logspaces) on other SSB appliances. Filtered logspaces allow the creation of a smaller, filtered subset of the logs contained in an existing local, remote or multiple logspace.

The syslog-ng PE application reads incoming messages and forwards them to the selected *destinations*. The syslog-ng PE application can receive messages from files, remote hosts, and other sources.

Log messages enter syslog-ng PE in one of the defined sources, and are sent to one or more *destinations*. In the case of the clients, one of the destinations is the syslog-ng Store Box. The destinations on the SSB can be logspaces or remote servers, such as database servers or log analyzing engines.

Sources and destinations are independent objects, *log paths* define what syslog-ng PE does with a message, connecting the sources to the destinations. A log path consists of one or more sources and one or more destinations: messages arriving to a source are sent to every destination listed in the log path. A log path defined in syslog-ng PE is called a *log statement*.

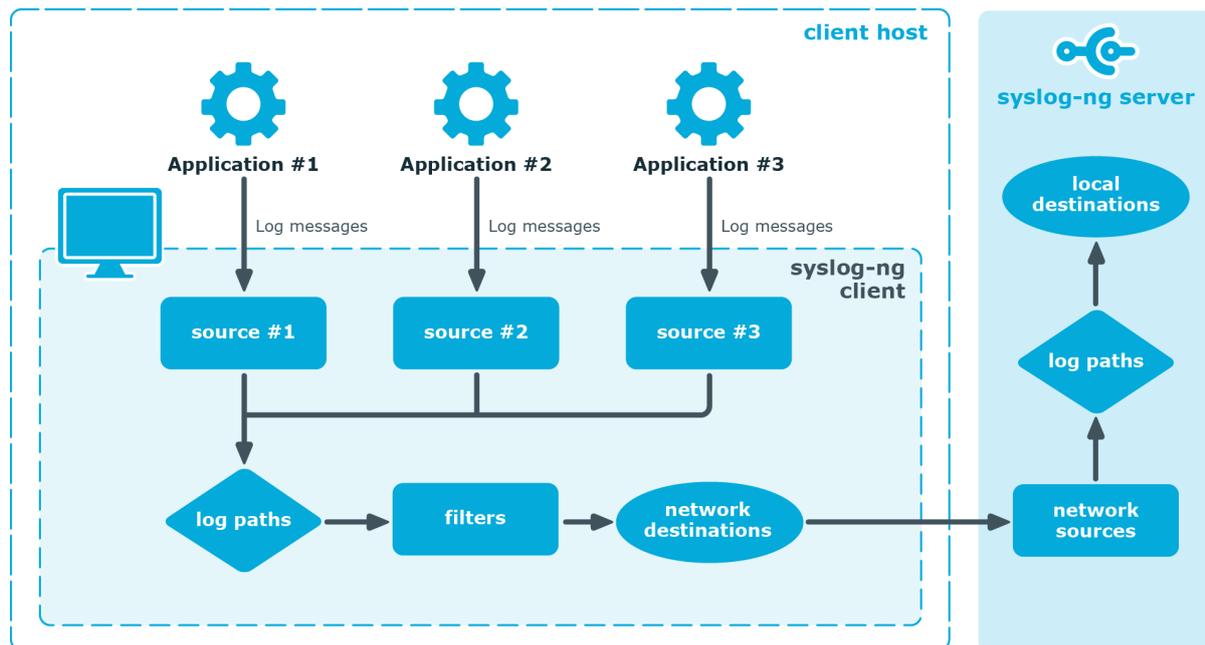
Optionally, log paths can include *filters*. Filters are rules that select only certain messages, for example, selecting only messages sent by a specific application. If a log path includes filters, syslog-ng PE sends only the messages satisfying the filter rules to the destinations set in the log path.

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

Collecting logs with SSB

The following procedure illustrates the route of a log message from its source on the [syslog-ng Premium Edition](#) (syslog-ng PE) client to the syslog-ng Store Box appliance.

Figure 2: The route of a log message



1. A device or application sends a log message to a source on the syslog-ng PE client. For example, an Apache web server running on Linux enters a message into the `/var/log/apache` file.
2. The syslog-ng PE client running on the web server reads the message from its `/var/log/apache` source.
3. The syslog-ng PE client processes the first log statement that includes the `/var/log/apache` source.
4. The syslog-ng PE client performs optional operations on the message, for example, it rewrites parts of the message or compares the message to the filters of the log statement (if any). If the message complies with all filter rules, syslog-ng PE sends the message to the destinations set in the log statement, for example, to the remote syslog-ng PE server.

After that, the syslog-ng PE client processes the next log statement that includes the `/var/log/apache` source, repeating Steps 3-4.

5. The message sent by the syslog-ng PE client arrives to a source set on the syslog-ng Store Box appliance.
6. The syslog-ng Store Box appliance reads the message from its source and processes the first log path that includes that source.
7. The syslog-ng Store Box appliance processes the message and performs the following operations. Note that most of these operations are optional, but the order of the processing steps is fixed.
 - a. Parse the message as a syslog message (unless message parsing is explicitly disabled for the source).

- b. [Classify the message](#) using a pattern database.
 - c. [Modify the message using rewrite rules](#) (before filtering).
 - d. [Filter the messages](#), for example, based on sender hostname or message content. If the message does not match the configured filter, syslog-ng Store Box(SSB) will not send it to the destination.
 - e. Parse the text of the message (that is, the `_${MESSAGE}` part) using a [key-value parser](#) or the [sudo parser](#).
 - f. [Modify the message using rewrite rules](#) (after filtering and other parsing).
 - g. SSB sends the message to the destinations set in the log path. The destinations are [local, optionally encrypted files on SSB](#), or [remote servers, such as a database server](#).
8. SSB processes the next log statement, repeating Steps 6-8.

NOTE: The syslog-ng PE application can stop reading messages from its sources if the destinations cannot process the sent messages. This feature is called flow-control and is detailed in [Managing incoming and outgoing messages with flow-control](#).

Managing incoming and outgoing messages with flow-control

This section describes the internal message-processing model of [syslog-ng Premium Edition](#) (syslog-ng PE), as well as the flow-control feature that can prevent message loss. To use flow-control, the **flow-control** option must be enabled for the particular log path.

The internal message-processing model of syslog-ng PE

1. The syslog-ng PE application checks the source for messages.
2. When a log message is found, syslog-ng PE reads the message.
3. The message is processed and put into the output buffer of the destination.
4. When the destination can accept the message, syslog-ng PE sends the message to the destination from the output buffer.

Flow-control

If the destination cannot send out messages, or not as fast as they arrive in the destination, the output buffer fills up. When the output buffer is full, the sources stop reading messages. This can prevent message loss.

If a message is successfully sent out from the destination, the source that sent that message starts reading logs again, until the destination buffer fills up.

Flow-control and multiple destinations

Using flow-control on a source has an important side-effect if the messages of the source are sent to multiple destinations. If flow-control is in use and one of the destinations cannot accept the messages, the other destinations do not receive any messages either, because syslog-ng PE stops reading the source. For example, if messages from a source are sent to a remote server and also stored locally in a file, and the network connection to the server becomes unavailable, neither the remote server nor the local file will receive any messages. This side-effect of the flow-control can be avoided by using the disk-based buffering feature of syslog-ng PE.

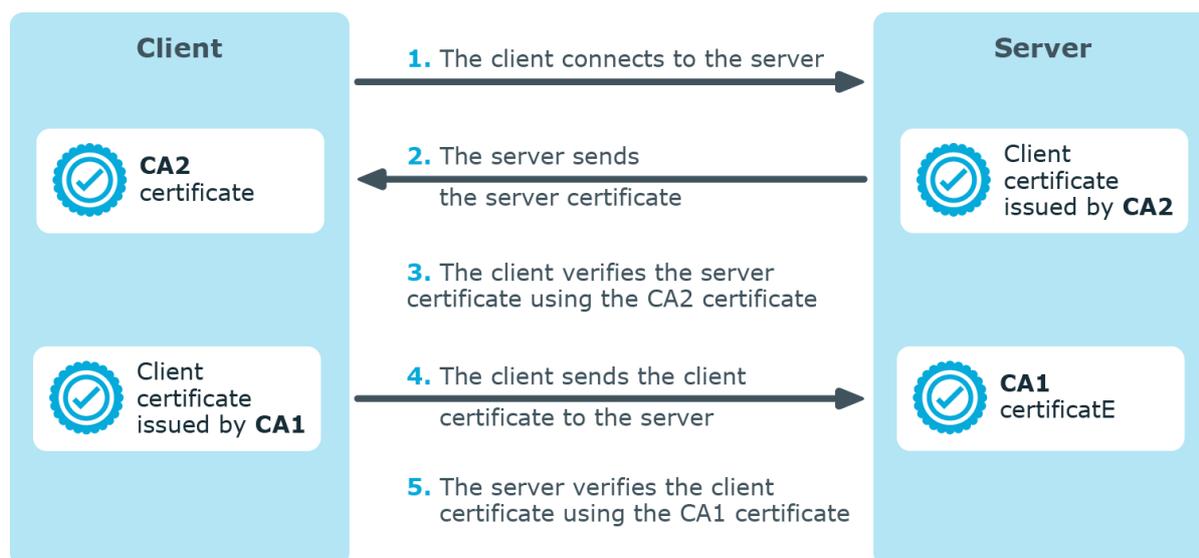
NOTE: Creating separate log paths for the destinations that use the same flow-controlled source does not help avoiding the problem.

Receiving logs from a secure channel

The syslog-ng Store Box (SSB) appliance receives log messages securely over the network using the Transport Layer Security (TLS) protocol (TLS is an encryption protocol over the TCP/IP network protocol).

TLS uses certificates to authenticate and encrypt communication, as illustrated in the following figure:

Figure 3: Certificate-based authentication



The client sending the logs authenticates SSB by requesting its certificate and public key. Optionally, SSB can also request a certificate from the client, thus mutual authentication is also possible.

In order to use TLS encryption in [syslog-ng Premium Edition](#) (syslog-ng PE), the following elements are required:

- A certificate on SSB that identifies SSB. This is available by default.
- The certificate of the Certificate Authority that issued the certificate of SSB must be available on the syslog-ng PE client.

When using mutual authentication to verify the identity of the clients, the following elements are required:

- A certificate must be available on the syslog-ng PE client. This certificate identifies the syslog-ng PE client.
- The certificate of the Certificate Authority that issued the certificate of the syslog-ng PE client must be available on SSB.

Mutual authentication ensures that SSB accepts log messages only from authorized clients.

For details on configuring TLS communication in syslog-ng PE, see [Configuring message sources](#) on page 193.

Advanced Log Transfer Protocol

The syslog-ng Store Box (SSB) appli can receive log messages in a reliable way over the TCP transport layer using the Advanced Log Transfer Protocol (ALTP). ALTP is a proprietary transport protocol that prevents message loss during connection breaks. The transport protocol is used between [syslog-ng Premium Edition](#) (syslog-ng PE) hosts and SSB (for example, a client and SSB, or a client-relay-SSB), and interoperates with the flow-control and reliable disk-buffer mechanisms of syslog-ng PE, thus providing the best way to prevent message loss.

The sender detects which messages the receiver has successfully received. If messages are lost during the transfer, the sender resends the missing messages, starting from the last successfully received message. Therefore, messages are not duplicated at the receiving end in case of a connection break (however, in failover mode this is not completely ensured). ALTP also allows for connections to be encrypted.

Network interfaces

The syslog-ng Store Box (SSB) hardware has five network interfaces: the external, the management, the internal (currently not used in SSB), the HA, and the IPMI interface. For details on hardware installation, see "[syslog-ng Store Box Hardware Installation Guide](#)" in the [Installation Guide](#).

External interface

The external interface is used for communication between SSB and the clients: clients send the syslog messages to the external interface of SSB. Also, the initial configuration of SSB is always performed using the external interface (for details on the initial configuration, see [Configuring SSB with the Welcome Wizard](#) on page 46). The external interface is used for

management purposes if the management interface is not configured. The external interface uses the Ethernet connector labeled as 1 (or EXT).

Using a 10Gbit interface as external interface

The SSB T-10 appliance is equipped with a dual-port 10Gbit interface. You can use the 10Gbit interface instead of the regular 1Gbit external (LAN 1) interface. That way, you can use SSB without any additional changes even if your network devices support only 10Gbit, and you must connect SSB to a 10Gbit-only network. This interface has SFP+ connectors (not RJ-45) labeled A and B, or labeled 5 and 6, depending on the hardware model, and can be found right of the Label 1 and 2 Ethernet interfaces.

NOTE: Only Intel-based SFP+ transceivers are compatible with the Intel 82599EB host chipset found in SSB.

The following Intel-based optical and Direct Attached Copper (DAC) SFP+ transceivers have been tested successfully with SSB, but in the future, their compatibility with SSB is not guaranteed:

- AOC-E10GSFPSR (optical)
- SFP-10GE-SR (DAC)
- FTLX8571D3BCVIT1 (DAC)

For a list of Intel-based connectors that may be compatible with the Intel 82599EB host chipset found in SSB, see the 82599-BASED ADAPTERS section on the Linux* Base Driver for the Intel(R) Ethernet 10 Gigabit PCI Express Adapters page at [82599-BASED ADAPTERS/Linux* Base Driver for the Intel\(R\) Ethernet 10 Gigabit PCI Express Adapters](#).

⚠ CAUTION:

Do not leave any unused SFP/SFP+ transceiver in the 10Gbit interface. It may cause network outage.

⚠ CAUTION:

Hazard of data loss One Identity recommends using a single interface (either 1, or A) and leaving the B interface unused.

If SSB detects a link on multiple interfaces, SSB will not switch to a different interface as long as the link is detected on the currently active interface, not even in case of packet loss or other network issues.

To ensure that your configuration is future-proof and to avoid having to reconfigure your appliance in the future, it is not recommended to use the B interface. In future releases of SSB, the B interface will be used exclusively in one particular type of scenario.

Management interface

The management interface is used exclusively for communication between SSB and the auditors or the administrators of SSB. Incoming connections are accepted only to access the SSB web interface, other connections targeting this interface are rejected. The management interface uses the Ethernet connector labeled as 2 (or MGMT).

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

TIP: It is recommended to direct backups, syslog and SNMP messages, and email alerts to the management interface. For details, see [Configuring the routing table](#) on page 78.

If the management interface is not configured, the external interface takes the role of the management interface.

NOTE: When deploying SSB in a virtual environment, it is sufficient to use only a single network interface. When only one network interface is defined, that interface will be the one used for management purposes, enabling access to SSB's web interface and the RPC API.

High availability interface

The high availability interface (HA) is an interface reserved for communication between the nodes of SSB clusters. The HA interface uses the Ethernet connector labeled as 4 (or HA). For details on high availability, see [High Availability support in SSB](#) on page 24.

IPMI interface

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

High Availability support in SSB

High availability clusters can stretch across long distances, such as nodes across buildings, cities or even continents. The goal of HA clusters is to support enterprise business continuity by providing location-independent load balancing and failover.

In high availability (HA) mode, two syslog-ng Store Box (SSB) units (called master and slave nodes) with identical configuration are operating simultaneously. The master shares all data with the slave node, and if the master node stops functioning, the other one becomes immediately active, so the servers are continuously accessible.

You can find more information on managing a high availability SSB cluster in [Managing a high availability SSB cluster](#) on page 135.

Firmware in SSB

The syslog-ng Store Box (SSB) firmware is separated into two parts: a boot and a core firmware.

- The boot firmware boots up SSB, provides the high availability support, and starts the core firmware.
- The core firmware handles everything else: provides the web interface, receives and processes log messages and so on.

When you upload a new .ISO file using the SSB web interface, it updates both firmware. For details, see [Upgrading SSB](#) on page 145.

Firmware and high availability

When powering on the syslog-ng Store Box (SSB) nodes in high availability mode, both nodes boot and start the boot firmware. The boot firmware then determines which unit is the master: the core firmware is started only on the master node.

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

Versions and releases of SSB

The following release policy applies to syslog-ng Store Box (SSB):

Long Term Support (LTS)

The initial release includes new features, bug fixes and security updates. After the initial release, only maintenance releases are published on this path, containing only bug fixes and security updates. The maintenance release frequency is typically six months.

Versioning: the first digit identifies the LTS main version (for example, 6.0.x), the second digit is always a 0, and the third digit designates the maintenance release (for example, 6.0.19). A long term support path is typically supported for three years after its original release.

For more information about versioning in SSB, see "[Version numbering in SSB](#)" in the [Upgrade Guide](#).

For more information about the upgrade paths between different version numbers in SSB, see "[Upgrade paths to SSB](#)" in the [Upgrade Guide](#).

Feature release

Feature releases include new features, bug fixes and security updates, building on the initially released LTS main version. Release frequency on this path is typically four months.

Versioning: the first digit identifies the LTS main version, the second digit designates the feature release (e.g. 6.4). Feature releases are typically supported for a year.

For more information about versioning in SSB, see "[Version numbering in SSB](#)" in the [Upgrade Guide](#).

For further information regarding the SSB LTS and Feature releases, see [the syslog-ng Store Box Product Life Cycle Table](#).

CAUTION:

Downgrading from a feature release to an earlier (and thus unsupported) feature release, or to the previous LTS release is officially not supported, but usually works as long as your SSB configuration file is appropriate for the old SSB version. However, persistent data like the position of the last processed message in a file source will be probably lost.

Logstore files created with a newer version of SSB might not be readable with an older version of SSB.

NOTE: Bug fixes and security updates are always issued in the latest & greatest releases, and never for previous releases. For example, in case of Long Term Support path, if a bug was reported by a customer for 6.0.1 LTS, the fix will be released in version 6.0.2 or in a later maintenance release. The same logic is true to rolling releases, for example, if a bug gets reported for 6.4, the fix will be issued in 6.5, or a later feature release.

NOTE: Consider the following information regarding hardware models:

- **Discontinued Support for Hardware Models N1000, N1000D, N10000:**

Last Day Sold: 01-Aug-2014

Last Day Renew: 01-Oct-2018

End of Support: 31-Mar-2019

- **Discontinued Support for T-Series hardware:**

Last Day Sold: 30-June-2019

Last Day Renew: 30-June-2021

End of Support: 30-June-2022

Licensing model and modes of operation

A Log Source Host (LSH) is any host, server, or device (including virtual machines, active or passive networking devices, syslog-ng clients and relays, and so on) that is capable of sending log messages. Log Source Hosts are identified by their IP addresses, so virtual machines and vhosts are separately counted.

The syslog-ng Store Box appliance as a central log-collecting server that receives messages through a network connection, and stores them locally, or forwards them to other destinations or external systems (for example, a SIEM or a database). The syslog-ng Store Box (SSB) appliance requires a license file, this license file determines the number of Log Source Hosts (LSHs) that can send log messages to the SSB server.

Note that the number of source hosts is important, not the number of hosts that directly sends messages to SSB: every host that send messages to the server (directly or using a relay) counts as a Log Source Host.

For technical reasons, the syslog-ng Store Box appliance itself counts as two LSHs in standalone mode, and three LSHs in high-availability (HA) mode. This is automatically adjusted when One Identity generates the license file.

Notes about counting the licensed hosts

⚠ CAUTION:

- If the actual IP address of the host differs from the IP address received by looking up its IP address from its hostname in the DNS, the syslog-ng server counts them as two different hosts.
- The `chain-hostnames()` option of syslog-ng can interfere with the way syslog-ng Store Box (SSB) counts the log source hosts, causing syslog-ng to think there are more hosts logging to the central server, especially if the clients sends a hostname in the message that is different from its real hostname (as resolved from DNS). Disable the `chain-hostnames()` option on your log source hosts to avoid any problems related to license counting.
- If the number of Log Source Hosts reaches the license limit, the SSB server will not accept connections from additional hosts. The messages sent by additional hosts will be dropped, even if the client uses a reliable transport method (for example, ALTP).
- If the `no-parse` flag is set in a message source on the SSB server, SSB assumes that the message arrived from the host (that is, from the last hop) that sent the message to SSB, and information about the original sender is lost.

Licensing benefits

Buying a syslog-ng Store Box (SSB) license permits you to perform the following:

- Deploy one instance of the SSB appliance as a central log collector server.
- The syslog-ng Store Box license also allows you to download the [syslog-ng Premium Edition](#) (syslog-ng PE) application (including the syslog-ng Agent for Windows application) and install it on hosts within your organization (on any supported platform) to use it as a log collector agent (client) for SSB. You cannot redistribute the application to third parties.

The syslog-ng Store Box license determines the number of individual hosts (also called log source hosts) that can send log messages to SSB.

License grants and legal restrictions are fully described in the [Software Transaction, License and End User License Agreements](#). Note that the [Software Transaction, License and End User License Agreements](#) and the syslog-ng Store Box [Product Guide](#) apply only to

scenarios where the Licensee (the organization who has purchased the product) is the end user of the product. In any other scenario — for example, if you want to offer services provided by SSB to your customers in an OEM or a Managed Service Provider (MSP) scenario — you have to negotiate the exact terms and conditions with One Identity.

License types

The following sections provide about the license types available for syslog-ng Store Box (SSB).

Perpetual license

Buying a license for a One Identity product allows you to use the product as described in the [Software Transaction, License and End User License Agreements](#).

You can download and use the latest Long Term Supported (LTS) Release of the product, and any subsequent Feature Release that is based on the Long Term Supported Release that was valid when you bought the license. To access the next Long Term Supported (LTS) Release, you must have a valid support package when the next Long Term Supported (LTS) Release is published.

Example: Accessing updates example

A customer's Support Service Agreement for syslog-ng Store Box (SSB) has expired and the customer did not renew it. At the time of expiration, the latest available versions were SSB 4 LTS and SSB 4 F3. In this case, the customer can access the current and future revisions of these versions, but they will not have access to future releases such as 4 F4 or 5 LTS when they are released.

Buying a subscription-based license automatically includes product support and access to the latest software versions.

You can download your licenses and the purchased software from the [support portal](#).

Subscription-based license

For virtual appliances, you can buy a subscription-based license that is valid for a fixed period of twelve (12) or thirty-six (36) months. The subscription-based license automatically includes product support and access to the latest software versions. For details, see the [Software Transaction, License and End User License Agreements](#).

Note that One Identity offers subscription-based licensing only in certain geographic regions and only for limited virtual appliance license options. For details, contact One Identity.

Licensing examples

Example: A simple example

Scenario:

- You want to deploy an syslog-ng Store Box (SSB) appliance as a log server.
- 45 servers with syslog-ng PE installed in client mode send logs to the SSB log server.
- 45 networks devices without syslog-ng PE installed send logs to the SSB log server.

License requirements: You need a syslog-ng Store Box license for at least 100 Log Source Host (LSH) as there are 90 LSHs ($45+45=90$) in this scenario.

Example: Using alternative log servers with syslog-ng PE clients

Scenario:

- You want to deploy an SSB appliance as a log server.
- 45 servers with syslog-ng PE installed in client mode send logs to the SSB log server.
- 45 networks devices without syslog-ng PE installed send logs to the SSB log server.
- 100 servers with syslog-ng PE installed send log messages to a log server without syslog-ng PE installed.

License requirements: You need a syslog-ng Store Box license for at least 200 LSHs as there are 190 LSHs ($45+45$ that send logs to a syslog-ng PE log server, and another 100 that run syslog-ng PE, $45+45+100=190$) in this scenario.

Example: Using syslog-ng PE relays

Scenario:

- You want to deploy an SSB appliance as a log server.
- 45 servers with syslog-ng PE installed in client mode send logs directly to the SSB log server.
- 5 servers with syslog-ng PE installed in relay mode send logs to the SSB log server.
- Every syslog-ng PE relay receives logs from 9 networks devices without syslog-ng PE installed (a total of 45 devices).
- 100 servers with syslog-ng PE installed send log messages to a log server without syslog-ng PE installed.

License requirements: You need a syslog-ng Store Box license for at least 200 LSH as there are 195 LSHs ($45+5+(5*9)+100=195$) in this scenario.

Example: Multiple facilities

You have two facilities (for example, data centers or server farms). Facility 1 has 75 AIX servers and 20 Microsoft Windows hosts, Facility 2 has 5 HP-UX servers and 40 Debian servers. That is 140 hosts altogether.

NOTE: If, for example, the 40 Debian servers at Facility 2 are each running 3 virtual hosts, then the total number of hosts at Facility 2 is 125, and the license sizes in the following examples should be calculated accordingly.

- **Scenario:** The log messages are collected to a single, central SSB log server.
License requirements: You need a syslog-ng Store Box license for 150 LSH as there are 140 LSHs ($75+20+5+40$) in this scenario.
- **Scenario:** Each facility has its own SSB log server, and there is no central log server.
License requirements: You need two separate licenses: a license for at least 95 LSHs ($75+20$) at Facility 1, and a license for at least 45 LSHs ($5+40$) at Facility 2. You need a license for 100 LSHs at Facility 1, and a license for 50 LSHs at Facility 2.
- **Scenario:** The log messages are collected to a single, central SSB log server. Facility 1 and 2 each have a syslog-ng PE relay that forwards the log messages to the central SSB log server.

License requirements: You need a syslog-ng Store Box license for 150 LSH as there are 142 LSHs (1+75+20+1+5+40) in this scenario (since the relays are also counted as an LSH).

- **Scenario:** Each facility has its own local SSB log server, and there is also a central SSB log server that collects every log message independently from the two local log servers.

License requirements: You need three separate licenses. A syslog-ng Store Box a license for at least 95 LSHs (75+20) at Facility 1, a license for at least 45 LSHs (5+40) at Facility 2, and also a license for at least 142 LSHs for the central syslog-ng Store Box log server (assuming that you want to collect the internal logs of the local log servers as well).

The structure of a log message

The following sections describe the structure of log messages. Currently there are two standard syslog message formats:

- The old standard described in RFC 3164 (also called the BSD-syslog or the legacy-syslog protocol): see [BSD-syslog or legacy-syslog messages](#) on page 31
- The new standard described in RFC 5424 (also called the IETF-syslog protocol): see [IETF-syslog messages](#) on page 33

BSD-syslog or legacy-syslog messages

This section describes the format of a syslog message, according to the legacy-syslog or BSD-syslog protocol (see [RFC 3164](#)). A syslog message consists of the following parts:

- [PRI](#)
- [HEADER](#)
- [MSG](#)

The total message must be shorter than 1024 bytes.

The following example is a sample syslog message:

```
<133>Feb 25 14:09:07 webserver syslogd: restart
```

The message corresponds to the following format:

```
<priority>timestamp hostname application: message
```

The different parts of the message are explained in the following sections.

NOTE: The syslog-ng application supports longer messages as well. For details, see the **Message size** option. However, it is not recommended to enable messages larger than the packet size when using UDP destinations.

The PRI message part

The PRI part of the syslog message (known as Priority value) represents the facility and severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the facility number by 8 and then adding the numerical value of the severity. The possible facility and severity values are presented below.

NOTE: Facility codes may slightly vary between different platforms.

The following table lists the facility values.

Table 1: syslog message facilities

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16-23	locally used facilities (local0-local7)

The following table lists the severity values.

Table 2: syslog Message severities

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

The HEADER message part

The HEADER part contains a time stamp and the hostname (without the domain name) or the IP address of the device. The time stamp field is the local time in the **Mmm dd hh:mm:ss** format, where:

- **Mmm** is the English abbreviation of the month: Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.
- **dd** is the day of the month in two digits. If the day of the month is less than 10, the first digit is replaced with a space. (For example Aug 7.)
- **hh:mm:ss** is the local time. The hour (hh) is represented in a 24-hour format. Valid entries are between 00 and 23, inclusive. The minute (mm) and second (ss) entries are between 00 and 59 inclusive.

The MSG message part

The MSG part contains the name of the program or process that generated the message, and the text of the message itself. The MSG part is usually in the following format:

```
program[pid]: message text
```

IETF-syslog messages

This section describes the format of a syslog message, according to the IETF-syslog protocol (see [RFC 5424-5428](#)). A syslog message consists of the following parts:

- [HEADER](#) (includes the [PRI](#) as well)
- [STRUCTURED-DATA](#)

- [MSG](#)

The following is a sample syslog message (source: <https://tools.ietf.org/html/rfc5424>):

```
<34>1 2003-10-11T22:14:15.003Z mymachine.example.com su - ID47 - BOM'su root'  
failed for lonvick on /dev/pts/8
```

The message corresponds to the following format:

```
<priority>VERSION ISOTIMESTAMP HOSTNAME APPLICATION PID MESSAGEID STRUCTURED-  
DATA MSG
```

- Facility is 4, severity is 2, so PRI is 34.
- The VERSION is 1.
- The message was created on 11 October 2003 at 10:14:15pm UTC, 3 milliseconds into the next second.
- The message originated from a host that identifies itself as "mymachine.example.com".
- The APP-NAME is "su" and the PROCID is unknown.
- The MSGID is "ID47".
- The MSG is "'su root' failed for lonvick...", encoded in UTF-8.
- In this example, the The encoding is defined by the BOM:
The byte order mark (BOM) is a Unicode character used to signal the byte-order of the message text.
- There is no STRUCTURED-DATA present in the message, this is indicated by "-" in the STRUCTURED-DATA field.

The HEADER part of the message must be in plain ASCII format, the parameter values of the STRUCTURED-DATA part must be in UTF-8, while the MSG part should be in UTF-8. The different parts of the message are explained in the following sections.

The PRI message part

The PRI part of the syslog message (known as Priority value) represents the facility and severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the facility number by 8 and then adding the numerical value of the severity. The possible facility and severity values are presented below.

| NOTE: Facility codes may slightly vary between different platforms.

The following table lists the facility values.

Table 3: syslog message facilities

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16-23	locally used facilities (local0-local7)

The following table lists the severity values.

Table 4: syslog message severities

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

The HEADER message part

The HEADER part contains the following elements:

- **VERSION:** The version number of the syslog protocol standard. Currently this can only be 1.
- **ISOTIMESTAMP:** The time when the message was generated in the ISO 8601 compatible standard time stamp format (yyyy-mm-ddThh:mm:ss+-ZONE), for example: 2006-06-13T15:58:00.123+01:00.
- **HOSTNAME:** The machine that originally sent the message.
- **APPLICATION:** The device or application that generated the message.
- **PID:** The process name or process ID of the syslog application that sent the message. It is not necessarily the process ID of the application that generated the message.
- **MESSAGEID:** The ID number of the message.

NOTE: The syslog-ng application supports other time stamp formats as well, like ISO, or the PIX extended format. The time stamp used in the IETF-syslog protocol is derived from RFC 3339, which is based on ISO 8601. For details, see the `ts_format()` option in .

The STRUCTURED-DATA message part

The STRUCTURED-DATA message part may contain meta-information about the syslog message, or application-specific information such as traffic counters or IP addresses. STRUCTURED-DATA consists of data elements enclosed in brackets ([]).

In the following example, you can see two STRUCTURED-DATA elements:

```
[exampleSDID@0 iut="3" eventSource="Application" eventID="1011"]  
[examplePriority@0 class="high"]
```

An element consists of an SD-ID (its identifier), and one or more parameters. Each parameter consists of a name and a value (for example, eventID="1011").

On SSB, the parameters (name-value pairs) parsed from these elements can be searched. From the example above, the following name-value pairs are parsed:

```
.sdata.exampleSDID@0.iut=3  
.sdata.exampleSDID@0.eventSource=Application  
.sdata.exampleSDID@0.eventID=1011  
.sdata.examplePriority@0.class=high
```

The syslog-ng application automatically parses the STRUCTURED-DATA part of syslog messages, which can be referenced in macros (see for details).

The MSG message part

The MSG part contains the text of the message itself. The encoding of the text must be UTF-8 if the BOM character is present in the message. If the message does not contain the BOM character, the encoding is treated as unknown. Usually messages arriving from legacy sources do not include the BOM character.

The Welcome Wizard and the first login

This chapter describes the initial steps of configuring syslog-ng Store Box (SSB). Before completing the steps below, unpack, assemble, and power on the hardware. Connect at least the external network interface to the local network, or directly to the computer from which SSB will be configured.

NOTE: Due to complexity of deployment, configuration, and design, you may require assistance from [One Identity Professional Services](#) while introducing new or additional:

- sources
- destinations
- log paths
- significant increases in log volume.

One Identity Professional Services is equipped and trained to evaluate the needs of any organization, and to provide configuration and architectural recommendations that help our users get the most out of any SSB version.

One Identity Professional Services offer assistance in planning and scoping for current needs, as well as recommendations for the future to ensure success.

NOTE: For details on unpacking and assembling the hardware, see "[syslog-ng Store Box Hardware Installation Guide](#)" in the [Installation Guide](#). For details on how to create a high availability SSB cluster, see "[Installing two SSB units in HA mode](#)" in the [Installation Guide](#).

The initial connection to SSB

The syslog-ng Store Box (SSB) appliance can be connected from a client machine using any modern web browser.

NOTE: For details on supported browsers, see [Supported web browsers](#) on page 65

SSB can be accessed from the local network. SSB 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 SSB receiving an IP address via DHCP. SSB accepts connections via its external interface (EXT, for details on the network interfaces, see [Network interfaces](#) on page 22).

TIP: The SSB console displays the IP address the external interface is listening on.

If SSB 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 SSB 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 39.
- For details on creating an alias IP on Linux, see [Creating an alias IP address \(Linux\)](#) on page 45.
- If configuring an alias interface is not an option for some reason, you can modify the IP address of SSB. For details, see [Modifying the IP address of SSB](#) on page 45.

CAUTION:

The Welcome Wizard can be accessed only using the external network interface of SSB, as the management interface is not configured yet.

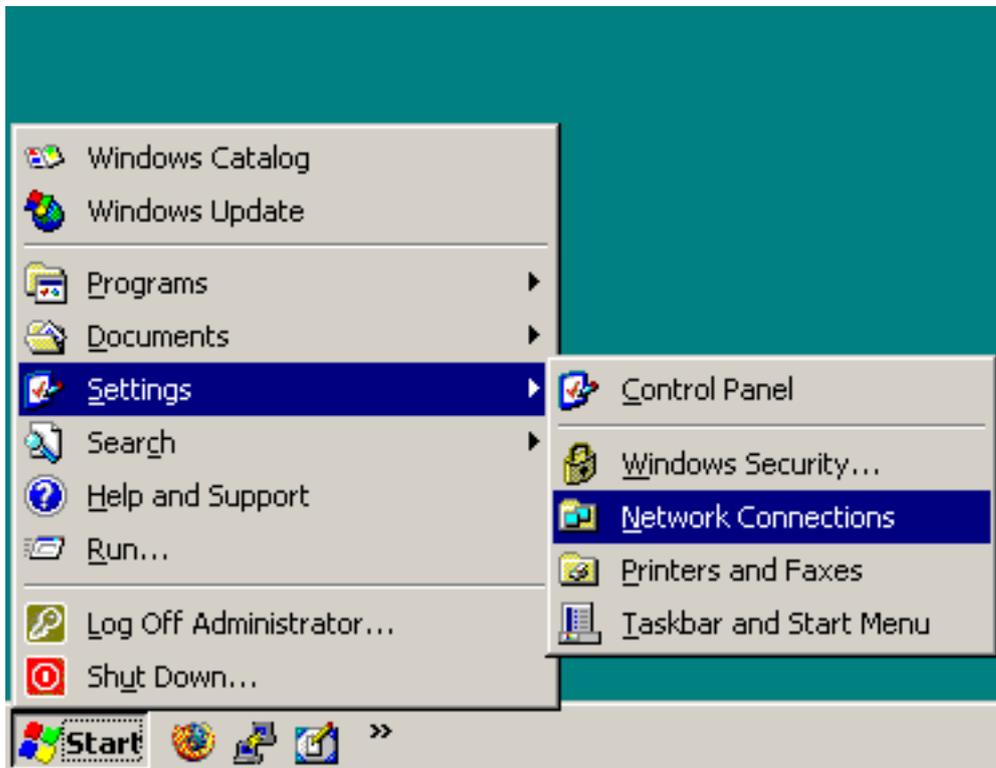
Open the page <https://192.168.1.1> from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

Creating an alias IP address (Microsoft Windows)

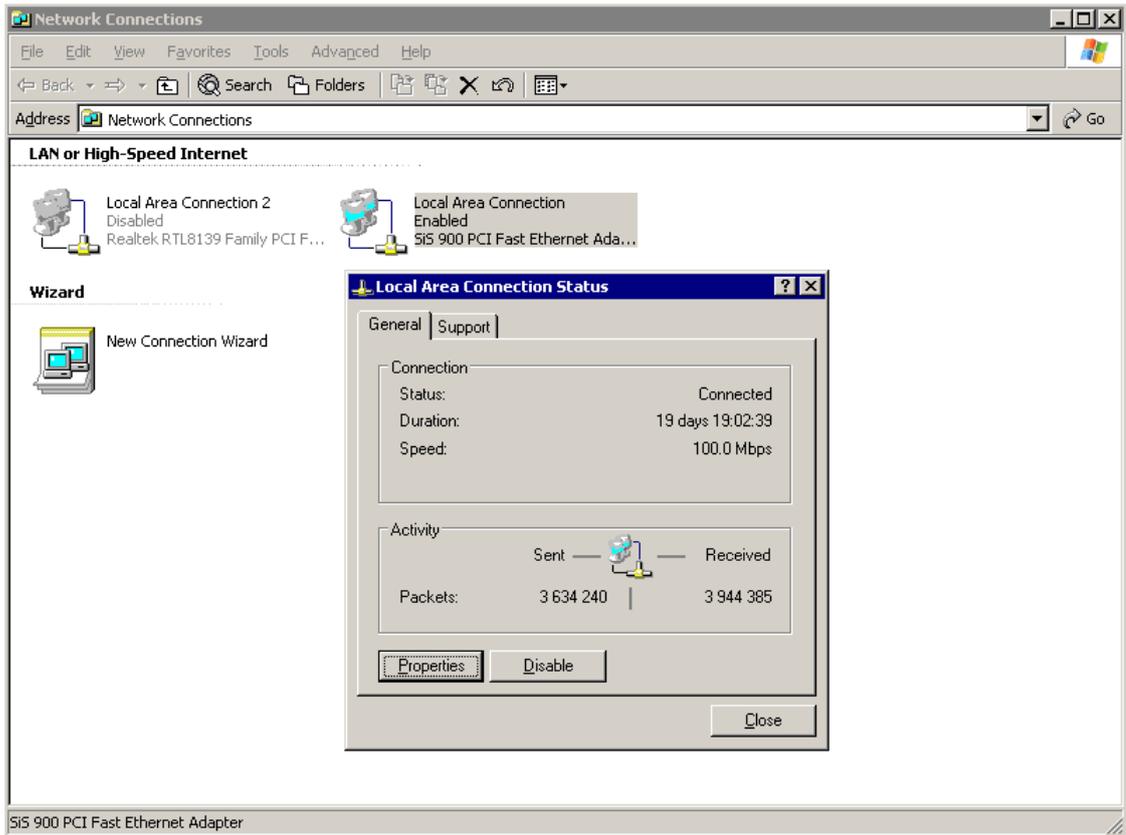
This procedure describes how to assign an alias IP address to a network interface on Microsoft Windows platforms.

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

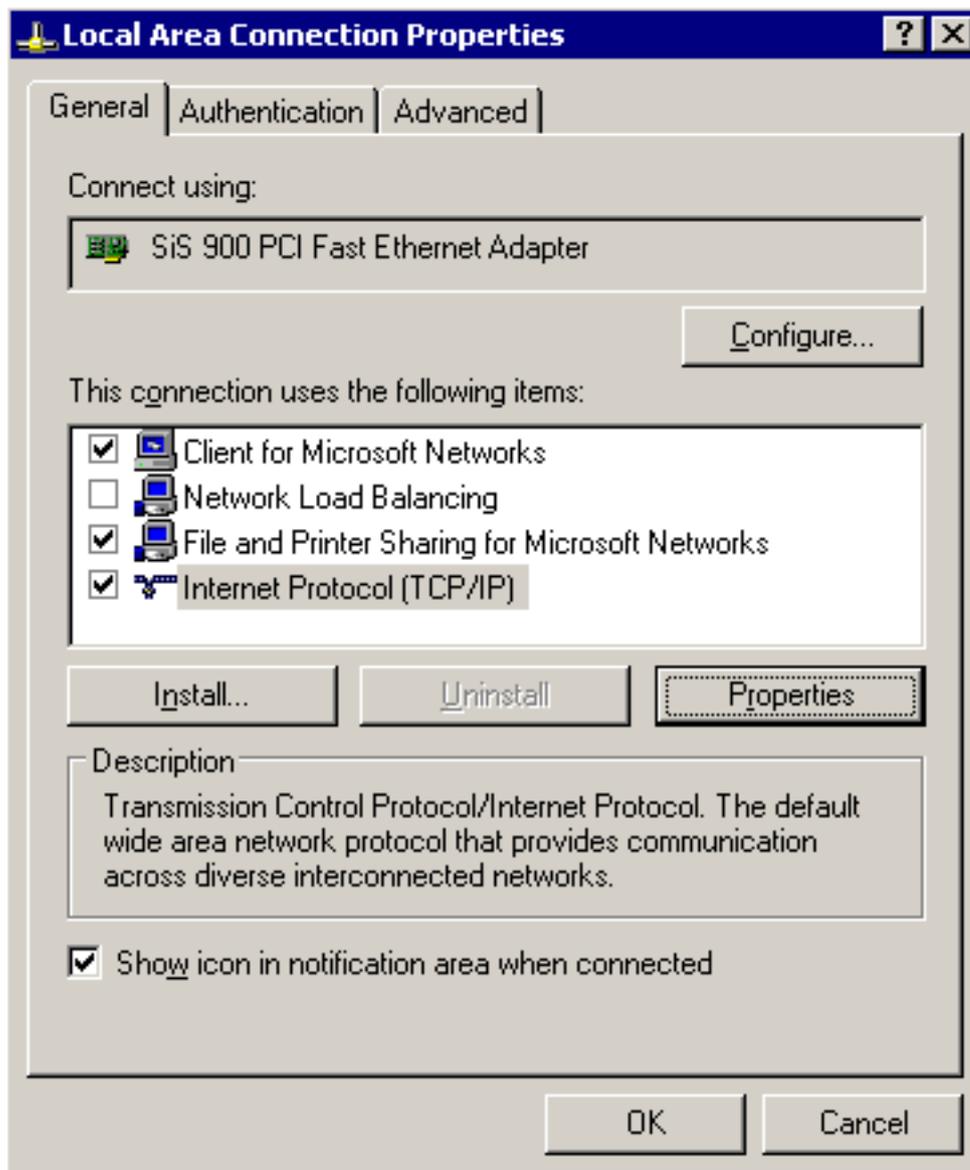
1. Navigate to **Start menu > Settings > Network Connections**.



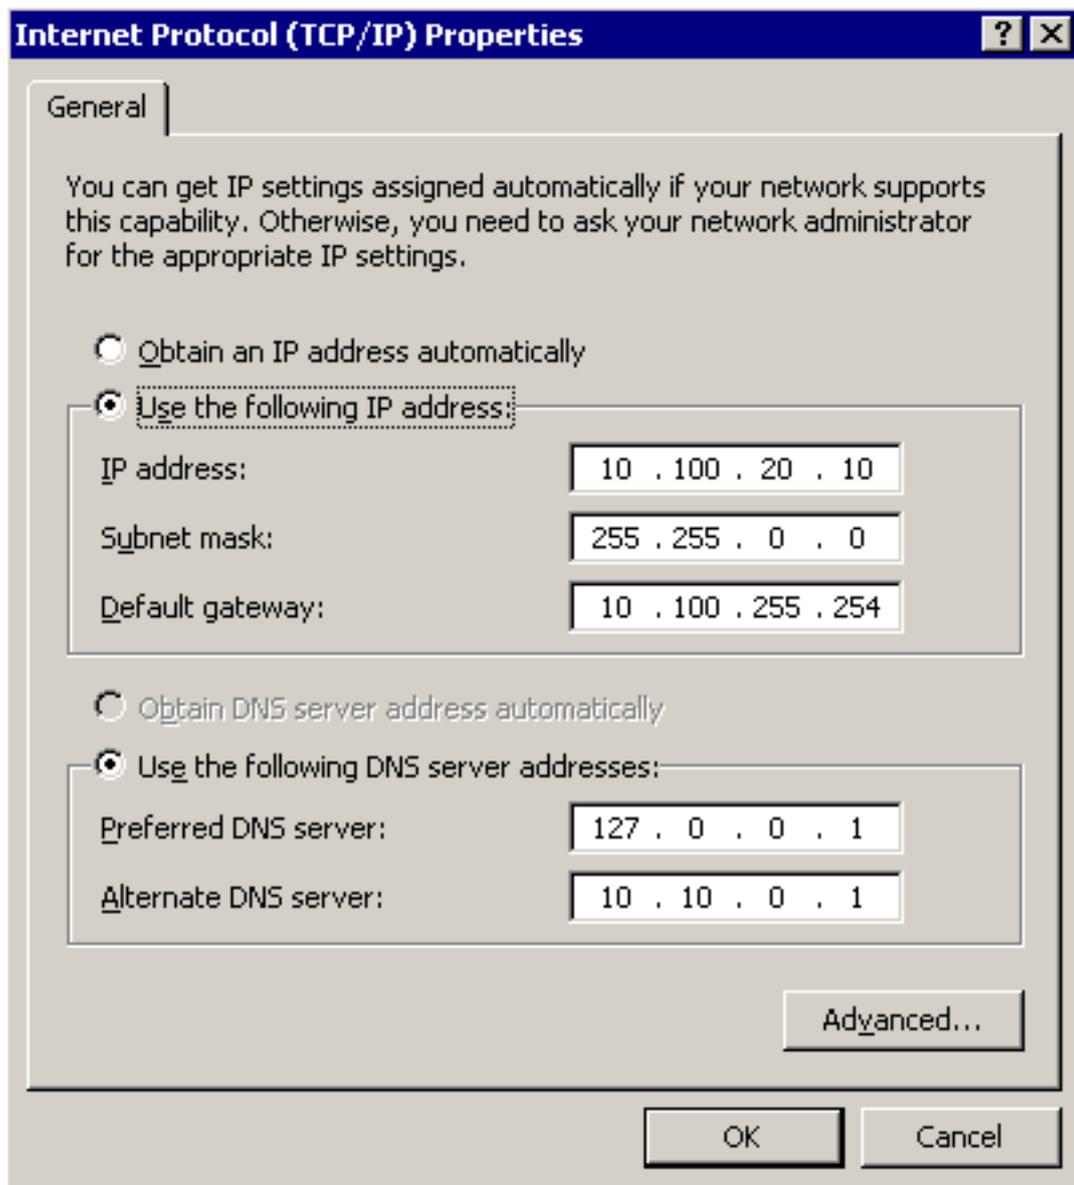
2. Double-click the **Local Area Connection** and then click **Properties**.



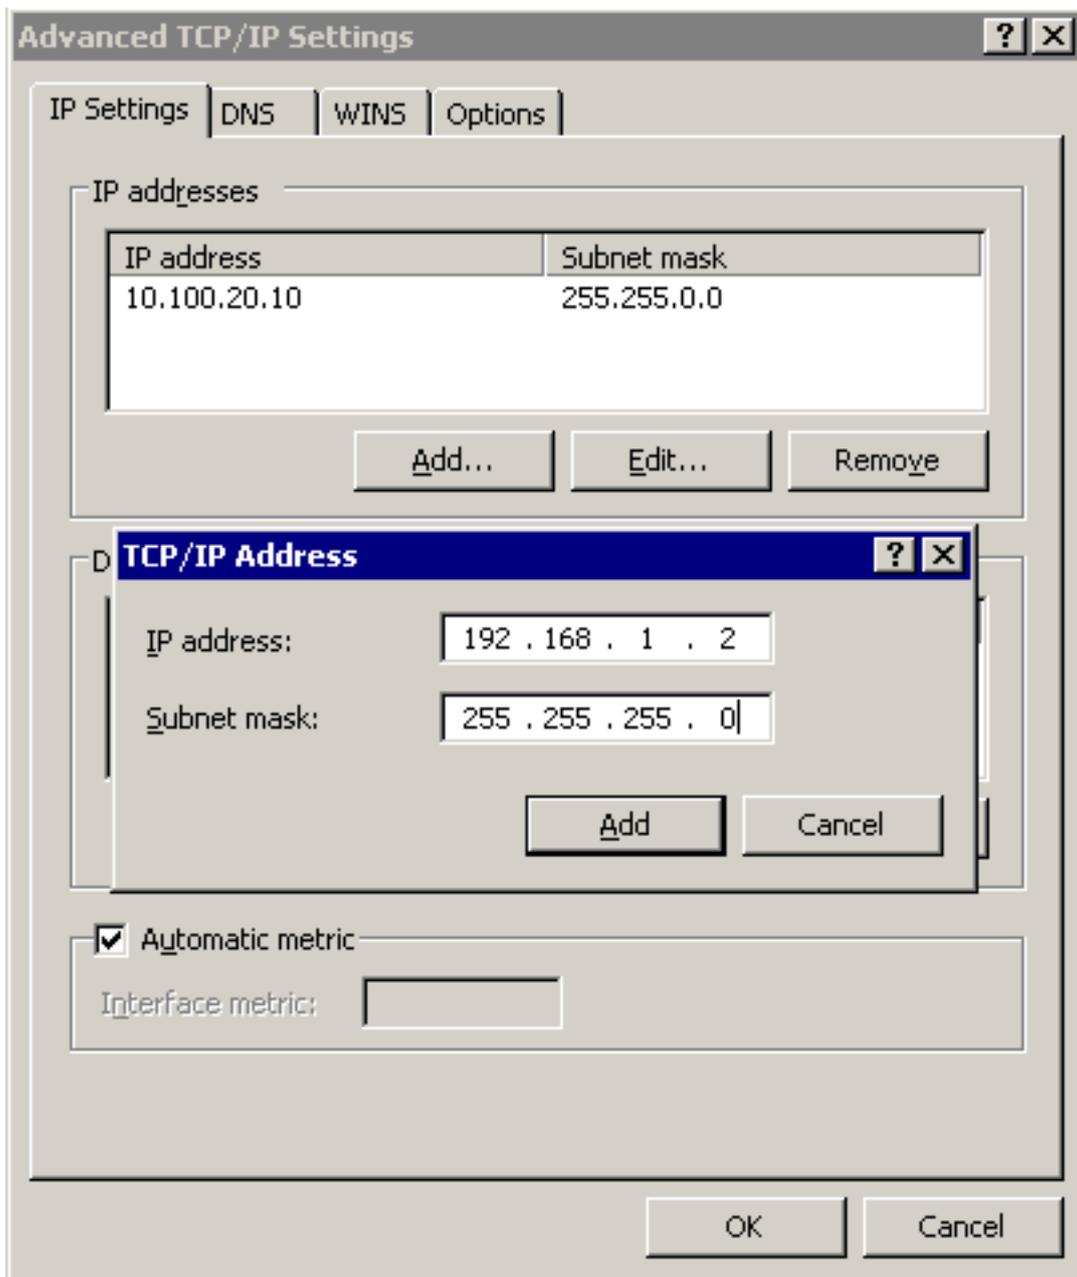
3. Select the **Internet Protocol (TCP/IP)** component in the list and click **Properties**.



4. To display the Advanced TCP/IP Settings window, click **Advanced**.



5. Select the **IP Settings** tab and in the **IP Addresses** section, click **Add**.



6. Into the **IP Address** field, enter 192.168.1.2. Into the **Netmask** field, enter 255.255.255.0.

⚠ CAUTION:

If your internal network uses the 192.168.1.0/24 IP range, the 192.168.1.1 and 192.168.1.2 addresses might already be in use. In this case, disconnect syslog-ng Store Box (SSB) from the network, and connect directly a computer to its external interface using a standard cross-link cable.

7. To complete the procedure, click **Add**.

Creating an alias IP address (Linux)

This procedure 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 syslog-ng Store Box (SSB) 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
```

Modifying the IP address of SSB

The following describes how to configure syslog-ng Store Box (SSB) to listen for connections on a custom IP address.

CAUTION:

Use this procedure only before the initial configuration of SSB, that is, before completing the Welcome Wizard. For details on changing the IP address or other network settings of a configured SSB system, see [Network settings on page 72](#).

If you change the IP address of SSB, make sure that you use this address as the External interface — IP address in [Configuring SSB with the Welcome Wizard](#).

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

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

```
ip addr add <IP-address>/24 dev eth0
```

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

4. Set the default gateway using the following command:

```
ip route add default via <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-SSB>> from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

Configuring SSB with the Welcome Wizard

The Welcome Wizard guides you through the basic configuration steps of syslog-ng Store Box (SSB). All parameters can be modified before the last step by using the **Back** button of the wizard, or later via the web interface of SSB.

The initial Welcome Wizard

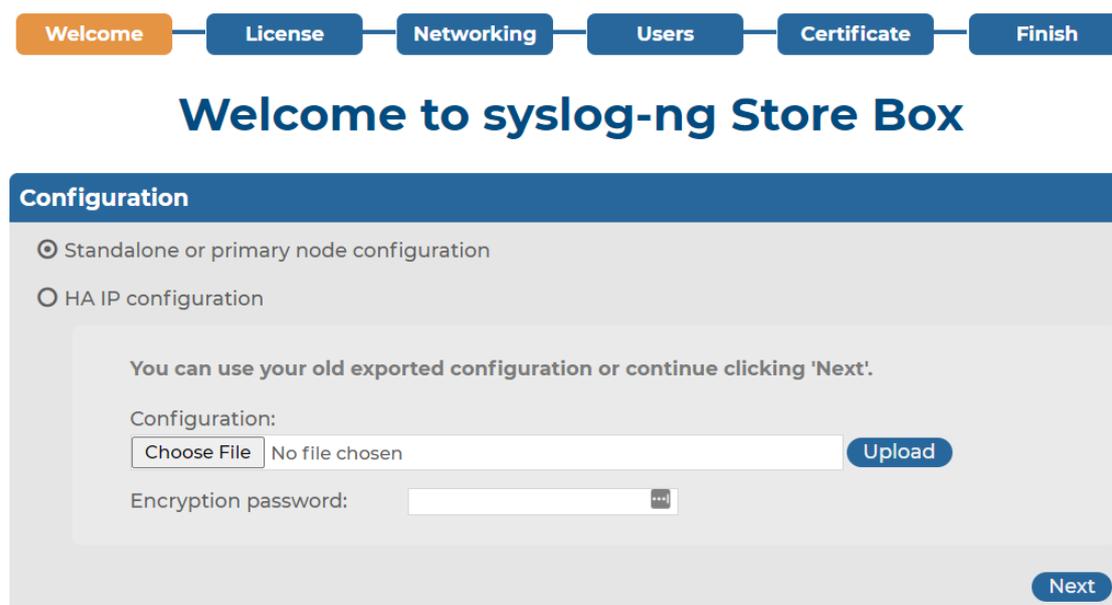
To start configuring SSB with the Welcome Wizard

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

TIP: The SSB console displays the IP address the external interface is listening on. SSB 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. Select your SSB configuration mode ([Standalone or primary node configuration](#), or [HA IP configuration](#)) for your SSB unit.

Figure 4: The Welcome Wizard



Configuring SSB as a standalone unit, or as the primary node of a HA cluster

This section describes how you can configure your unit as a standalone unit, or as the primary node of a HA cluster in the [syslog-ng Store Box \(SSB\) Welcome Wizard](#).

For details on how you can configure your unit as the secondary node of a HA cluster, see [Configuring your SSB unit as the secondary node of a HA cluster](#).

If you want to use your unit as a standalone unit, or as the primary node of a HA cluster, you can configure a new unit in the Welcome Wizard, or import an existing configuration from a backup file to restore a backup configuration after a recovery, or to migrate an existing SSB configuration to a new device.

On the initial screen, choose one of the following options:

- [Importing an existing SSB configuration](#)
- [Configuring a new SSB unit](#)

Importing an existing SSB configuration

To import an existing SSB configuration to be used as a standalone unit, or as the primary node of a HA cluster

1. On the initial **Configuration** screen, select **Standalone or primary node configuration**.

Figure 5: Standalone or primary node configuration

Welcome — License — Networking — Users — Certificate — Finish

Welcome to syslog-ng Store Box

Configuration

Standalone or primary node configuration

HA IP configuration

You can use your old exported configuration or continue clicking 'Next'.

Configuration:

No file chosen

Encryption password:

2. Click **Choose File** and select the configuration file to import.

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

3. Enter the password used when the configuration was exported into the **Encryption password** field.

For details on restoring configuration from a configuration backup, see [Restoring SSB configuration and data](#) on page 438.

4. Click **Import**.

CAUTION:

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

5. Accept the [Software Transaction, License and End User License Agreements](#) and install the SSB license.

Figure 6: The Software Transaction, License and End User License Agreements, and the license key

Welcome License Networking Users Certificate Finish

Welcome to syslog-ng Store Box

Software Transaction Agreement

PLEASE READ THIS AGREEMENT CAREFULLY BEFORE USING THIS PRODUCT. BY DOWNLOADING, INSTALLING OR USING THIS PRODUCT, YOU ACCEPT AND AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT. FOR ORDERS PLACED OUTSIDE THE UNITED STATES OF AMERICA, PLEASE GO TO <http://quest.com/legal/sta.aspx> TO VIEW THE APPLICABLE VERSION OF THIS AGREEMENT FOR YOUR REGION. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT OR THE APPLICABLE VERSION OF THIS AGREEMENT FOR YOUR REGION, DO NOT DOWNLOAD, INSTALL OR USE THIS PRODUCT. IF YOU HAVE A SIGNED AGREEMENT WITH PROVIDER THAT IS SPECIFICALLY REFERENCED IN AN ORDER THAT IS EXECUTED BETWEEN YOU AND PROVIDER, THEN THAT SIGNED AGREEMENT WILL SUPERSEDE THIS AGREEMENT.

This Software Transaction Agreement (the "Agreement") is made between you, the Customer ("Customer" or "You") and the Provider, as defined below.

Confirmation

I have read and agree with the terms and conditions:

By accepting the terms and conditions, you agree with the Software Transaction Agreement, in accordance with your purchase order.

License file upload

License file upload: No file chosen

Customer:
Serial:
Limit type:
Host limit:

- Read the [Software Transaction, License and End User License Agreements](#) and select **Accept**. The License Agreement covers both the traditional license, and subscription-based licensing as well. Clicking **Accept** means that you accept the agreement that corresponds to the license you purchased (for details on subscription-based licensing, see [License types](#) on page 28). After the installation is complete, you can read the [Software Transaction, License and End User License Agreements](#) at **Basic Settings > System > License**.
- Click **Browse**, select the SSB license file received with SSB, then click **Upload**. Without a license file, SSB will run in demo mode.

NOTE: It is not required to manually decompress the license file. Compressed licenses (for example .zip archives) can also be uploaded.

- c. Click **Next**.

Configuring a new SSB unit

To configure your SSB unit as a standalone unit, or as the primary node of a HA cluster

1. On the initial **Configuration** screen, select **Standalone or primary node configuration**.

Figure 7: Standalone or primary node configuration

The screenshot shows a progress bar at the top with steps: Welcome (highlighted), License, Networking, Users, Certificate, and Finish. Below the progress bar is the heading "Welcome to syslog-ng Store Box". The main content area is titled "Configuration" and contains two radio button options: "Standalone or primary node configuration" (selected) and "HA IP configuration". Below these options is a text box that says "You can use your old exported configuration or continue clicking 'Next'". Underneath is a "Configuration:" section with a "Choose File" button, a text input field containing "No file chosen", and an "Upload" button. Below that is an "Encryption password:" label and a password input field with a toggle icon. A "Next" button is located at the bottom right of the configuration area.

2. Click **Next**.
3. Accept the [Software Transaction, License and End User License Agreements](#) and install the SSB license.

Figure 8: The Software Transaction, License and End User License Agreements, and the license key

Welcome License Networking Users Certificate Finish

Welcome to syslog-ng Store Box

Software Transaction Agreement

PLEASE READ THIS AGREEMENT CAREFULLY BEFORE USING THIS PRODUCT. BY DOWNLOADING, INSTALLING OR USING THIS PRODUCT, YOU ACCEPT AND AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT. FOR ORDERS PLACED OUTSIDE THE UNITED STATES OF AMERICA, PLEASE GO TO <http://quest.com/legal/sta.aspx> TO VIEW THE APPLICABLE VERSION OF THIS AGREEMENT FOR YOUR REGION. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT OR THE APPLICABLE VERSION OF THIS AGREEMENT FOR YOUR REGION, DO NOT DOWNLOAD, INSTALL OR USE THIS PRODUCT. IF YOU HAVE A SIGNED AGREEMENT WITH PROVIDER THAT IS SPECIFICALLY REFERENCED IN AN ORDER THAT IS EXECUTED BETWEEN YOU AND PROVIDER, THEN THAT SIGNED AGREEMENT WILL SUPERSEDE THIS AGREEMENT.

This Software Transaction Agreement (the "Agreement") is made between you, the Customer ("Customer" or "You") and the Provider, as defined below.

Confirmation

I have read and agree with the terms and conditions:

By accepting the terms and conditions, you agree with the Software Transaction Agreement, in accordance with your purchase order.

License file upload

License file upload: No file chosen

Customer:
Serial:
Limit type:
Host limit:

- Read the [Software Transaction, License and End User License Agreements](#) and select **Accept**. The License Agreement covers both the traditional license, and subscription-based licensing as well. Clicking **Accept** means that you accept the agreement that corresponds to the license you purchased (for details on subscription-based licensing, see [License types](#) on page 28). After the installation is complete, you can read the [Software Transaction, License and End User License Agreements](#) at **Basic Settings > System > License**.
- Click **Browse**, select the SSB license file received with SSB, then click **Upload**. Without a license file, SSB will run in demo mode.

NOTE: It is not required to manually decompress the license file. Compressed licenses (for example .zip archives) can also be uploaded.

- c. Click **Next**.
- 4. Fill the fields to configure networking. The meaning of each field is described below. The background of unfilled required fields is red. All parameters can later be modified using the regular interface of SSB.

Figure 9: Initial networking configuration

Networking settings		
External interface:	IP address	Netmask
	<input type="text" value="10.30.255.76"/>	<input type="text" value="255.255.255.0"/>
Default GW:	<input type="text" value="10.30.255.254"/>	
Hostname:	<input type="text" value="demo"/>	
Domainname:	<input type="text" value="example.com"/>	
DNS server:	<input type="text" value="10.30.255.254"/>	
NTP server:	<input type="text" value="10.30.255.254"/>	
SMTP server:	<input type="text" value="mail.example.com"/>	
Administrator's email:	<input type="text" value="syslogadmin@example.com"/>	
Timezone:	<input type="text" value="Europe/Budapest"/>	

Make sure that you have entered the correct timezone. It is not recommended to change the timezone later, because logspace rotation is based on your local timezone. If you change the timezone later, you will not be able to properly search in your previously stored logs.

[Back](#) [Next](#)

- a. **External interface — IP address:** IP address of the external interface of SSB (for example, 192.168.1.1). The IP address can be chosen from the range of the corresponding physical subnet. Clients will connect the external interface, therefore it must be accessible to them.

If you have changed the IP address of SSB from the console before starting the Welcome Wizard, make sure that you use the same address here.

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

- IPv4 addresses
 - 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
 - 127.0.0.0/8 (localhost IP addresses)
- b. **External interface — Netmask:** The IP netmask of the given range in IP format. For example, general class C networks have the 255.255.255.0 netmask.
- c. **Default gateway:** IP address of the default gateway. When using several network cards, the default gateway is usually in the direction of the external interface.
- d. **Hostname:** Name of the machine running SSB (for example, SSB).
- e. **Domain name:** Name of the domain used on the network.
- f. **DNS server:** IP address of the name server used for domain name resolution.
- g. **NTP server:** The IP address or the hostname of the NTP server.
- h. **SMTP server:** The IP address or the hostname of the SMTP server used to deliver e-mails.
- i. **Administrator's e-mail:** E-mail address of the SSB administrator.
- j. **Timezone:** The timezone where the SSB is located.

 **CAUTION:**

Make sure that you have selected the correct timezone. It is not recommended to change the timezone later, because logspace rotation is based on your local timezone. If you change the timezone later, you will not be able to properly search in your previously stored logs.

- k. Click **Next**.
5. Enter the passwords used to access SSB.

Figure 10: Passwords

Welcome License Networking **Users** Certificate Finish

Welcome to syslog-ng Store Box

User settings

Admin password: [password field] weak good **strong**

Retype admin password: [password field]

Root password: [password field] weak good **strong**

Retype root password: [password field]

Seal the box:

Back Next

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

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

The default password policy on newly installed SSB appliances does not accept simple passwords for the admin and root users. As you type, SSB shows the strength of the password under the password field. Enter a password that gets at least a "good" rating.

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

The default password policy on newly installed SSB appliances does not accept simple passwords for the admin and root users. As you type, SSB shows the strength of the password under the password field. Enter a password that gets at least a "good" rating.

NOTE: Accessing SSB using SSH is rarely needed, and recommended only for advanced users for troubleshooting situations.

- c. If you want to prevent users from accessing SSB remotely via SSH or changing the root password of SSB, select the **Seal the box** checkbox. Sealed mode can be activated later from the web interface as well. For details, see [Sealed mode](#) on page 155.
- d. Click **Next**.

6. Upload or create a certificate for the SSB web interface. This SSL certificate will be displayed by SSB to authenticate administrative HTTPS connections to the web interface and RPC API.

Figure 11: Creating a certificate for SSB

Welcome License Networking Users **Certificate** Finish

Welcome to syslog-ng Store Box

Server certificate

Server X.509 certificate:
📄 /C=HU/L=Budapest/O=Example Inc/OU=IT Security/CN=demo.example.com

Server private key: 2048 SHA256:MC4o2KcZHg6ulN9EnEKrv7H1PMLaHUKy536XciNWBLM

Generate new self-signed certificate

Country: Hungary -- HU

Locality name: Budapest

Organization name: Example Inc

Organizational unit name: IT Security

State or Province name:

Generate certificate

Back Next

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

- Country:** Select the country where SSB is located (for example, HU-Hungary).
- Locality:** The city where SSB is located (for example, Budapest).
- Organization:** The company who owns SSB (for example, Example Inc.).
- Organization unit:** The division of the company who owns SSB (for example, IT Security Department).
- State or Province:** The state or province where SSB is located.
- Click **Generate**.

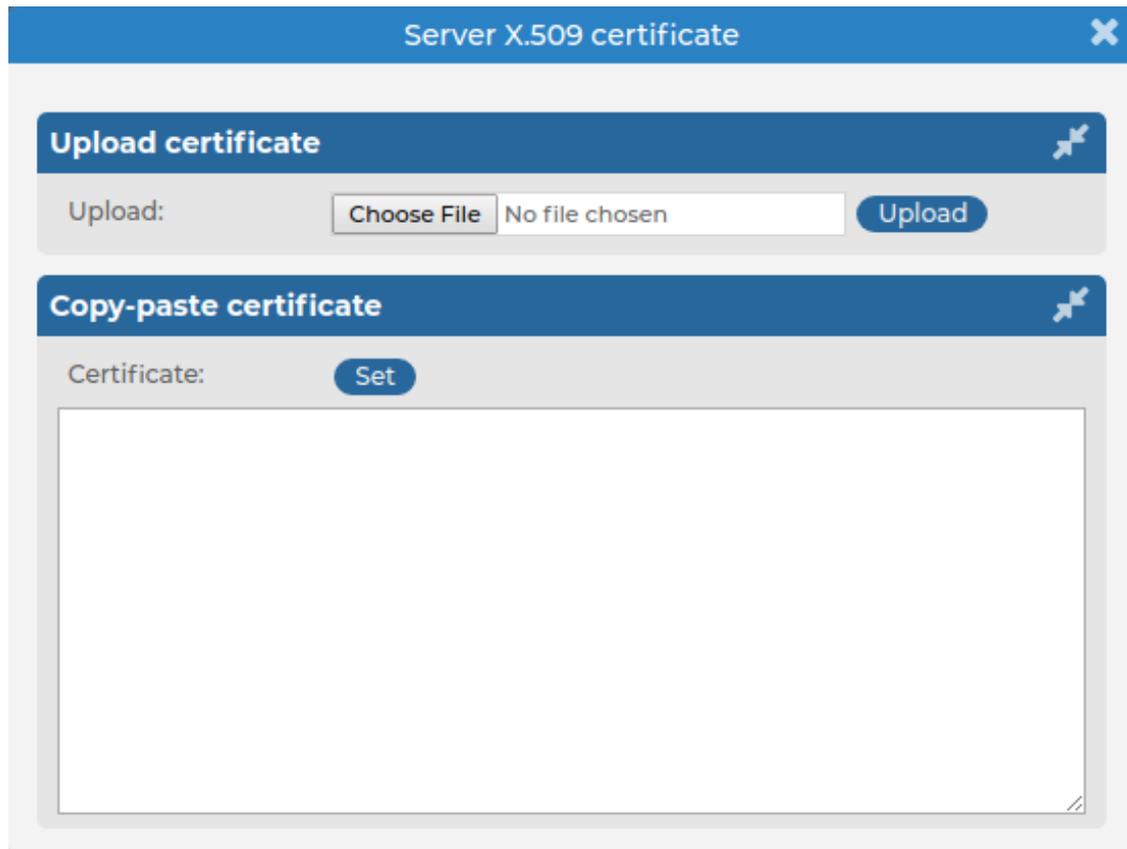
If you want to use a certificate that is signed by an external Certificate Authority, in the **Server X.509 certificate** field, click 📄 to upload the certificate.

NOTE: If you want to create a certificate with Windows Certificate Authority (CA) that works with SSB, generate a CSR (certificate signing request) on a computer

running OpenSSL (for example, using the **openssl req -set_serial 0 -new -newkey rsa:2048 -keyout ssbwin2k121.key -out ssbwin2k121.csr -nodes** command), sign it with Windows CA, then import this certificate into SSB.

- If you are using Windows Certificate Authority (CA) on Windows Server 2008, see [Generating TSA certificate with Windows Certificate Authority on Windows Server 2008](#) on page 171 for details.
- If you are using Windows Certificate Authority (CA) on Windows Server 2012, use the standard web server template to sign the certificate.

Figure 12: Uploading a certificate for SSB



You can choose to upload a single certificate or a certificate chain (that is, intermediate certificates and the end-entity certificate).

After uploading a certificate or certificate chain, you can review details by clicking the name of the certificate, and looking at the information displayed in the pop-up window that comes up.

Figure 13: Log > Options > TLS settings – X.509 certificate details



The pop-up window allows you to:

- Download the certificate or certificate chain.
 | **NOTE:** Certificate chains can only be downloaded in PEM format.
- View and copy the certificate or certificate chain.
- Check the names and the hierarchy of certificates (if it is a certificate chain and there is more than one certificate present).
 On hovering over a certificate name, the subject of the certificate is displayed, describing the entity certified.
- Check the validity dates of the certificate or certificates making up the chain.
 On hovering over a particular date, the exact time of validity is also displayed.

After uploading the certificate or certificate chain, the presence or absence of the string (**chain**) displayed after the name of the certificate will indicate whether the certificate is a certificate chain or a single certificate.

Then, back on the **Certificate** page of the Welcome Wizard, in the **Server private key** field, click , upload the private key, and enter the password protecting the private key.

Figure 14: Uploading a private key

The screenshot shows a web interface for configuring a server private key. The window title is "Server private key". It is divided into two sections: "Upload key" and "Copy-paste key".

- Upload key section:** Includes a "Choose File" button, a text box showing "No file chosen", and an "Upload" button. Below this is a "Password:" label and a text input field.
- Copy-paste key section:** Includes a "Key:" label, a "Set" button, a large text area for pasting the key, and a "Password:" label with a text input field at the bottom.

NOTE:

SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
- Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

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

Figure 15: Review configuration data

Welcome License Networking Users Certificate **Finish**

Welcome to syslog-ng Store Box

Configuration details

Hostname:	demo
Domainname:	example.com
External address:	10.30.255.76
Management address:	10.30.255.76
Default gateway:	10.30.255.254
DNS server:	10.30.255.254
Timezone:	Europe/Budapest
NTP server:	10.30.255.254
SSL certificate:	/C=HU/L=Budapest/O=Example Inc/OU=IT Security/CN=demo.example.com
SSH RSA key:	1024 SHA256:GMeBki+Xz0Vy8NrjmybzY53/eGAdUrVS350MdYeQ4wM
Licensed version:	5.0
Licensed customer:	BalaBit (Beta)
Licensed hosts:	Unlimited
License serial:	8080ca30-7eb0-ecb0-4130-9da0e800e840

Back Finish

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

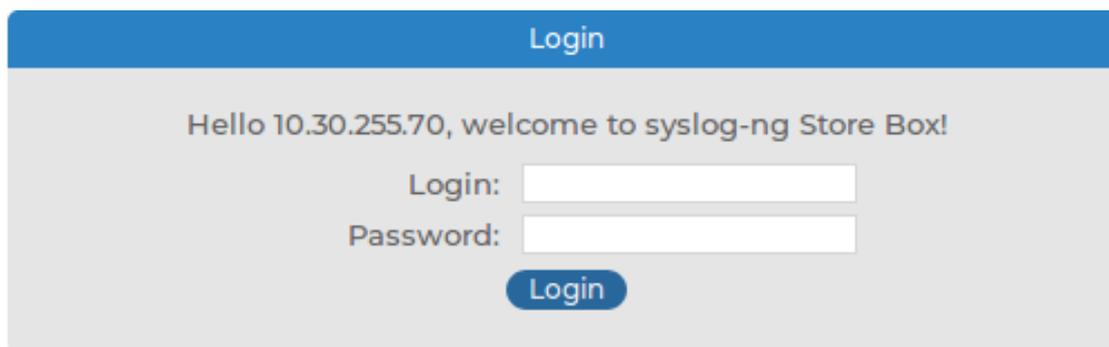
CAUTION:

The configuration takes effect immediately after clicking Finish. Incorrect network configuration data can render SSB inaccessible.

SSB is now accessible from the regular web interface via the IP address of its external interface.

After you finish configuring your SSB unit (which you can use as a standalone SSB unit, or as the primary node of a HA cluster), your browser is automatically redirected to the IP address set as the external interface of SSB, where you can log in to the web interface of SSB using the `admin` username and the password you set for this user in the Welcome Wizard.

Figure 16: Logging in to SSB



Preparing the nodes on the SSB web interface for establishing a HA cluster

If you want to use the newly configured SSB unit as the primary node in a future HA cluster, and you want to add an additional SSB unit as the secondary node in your future HA cluster, you have to configure the **IP addresses** that you want to use for your primary node (referred to as **This node** on the web interface, and occasionally as master node in error messages and warnings), and the secondary node (referred to as **Other node** on the web interface, and occasionally as slave node in error messages and warnings).

To prepare the nodes on your SSB web interface for establishing a HA cluster

1. Log in to the SSB unit configured as the primary node for your future HA cluster.
2. Navigate to **Basic Settings > High Availability**.

The newly configured standalone unit is displayed under **High availability & Nodes**, labeled as **This node**. The greyed out **Other node** is not yet configured, but in the **Interface IP** field, you can already set the **IP address** that you want to use on your secondary node later.

NOTE: Note that your **Cluster status** displays your primary SSB unit in a **STANDALONE** HA state.

3. In the **Interface IP** field on **This node**, set the **IP address** that you want to use for your primary node in your future HA cluster.
4. In the **Interface IP** field on **Other node**, set the **IP address** that you want to use for the secondary node in your future HA cluster.

NOTE: Make sure that the **IP address** you configure on **This node** is different from the **IP address** you configure on **Other node**.

5. **Commit** your changes.

NOTE: When your configuration changes are successfully saved, you will see a warning about the limitations of configuring your secondary node at this point. Click **OK**.

6. (Optional) Reboot your SSB unit. Alternatively, you can reboot your SSB unit later, after configuring a different unit as the secondary node of your future HA cluster.
7. [Configure a different SSB unit as the secondary node](#) of your future HA cluster.
8. [Convert your nodes into a HA cluster](#) on the SSB web interface.

Configuring your SSB unit as the secondary node of a HA cluster

This section describes how you can configure your syslog-ng Store Box (SSB) unit as the secondary node of a HA cluster in the in the [syslog-ng Store Box \(SSB\) Welcome Wizard](#).

Prerequisites

Before configuring your SSB unit as the secondary node of your future HA cluster, you must have a standalone SSB unit configured as the primary node of your HA cluster, and prepare the nodes on your SSB web interface to establish a HA cluster from your SSB units.

For details on how you can configure your SSB unit as a standalone unit, or as the primary node of a HA cluster, see [Configuring SSB as a standalone unit, or as the primary node of a HA cluster](#).

For details on how you can prepare the nodes on your SSB web interface to establish a HA cluster from your SSB units, see [Preparing the nodes on the SSB web interface for establishing a HA cluster](#).

HA IP configuration

If you want to use your SSB unit as the secondary node of a HA cluster, you can use the syslog-ng Store Box Welcome Wizard, but with fewer configuration steps than when you are configuring your primary node.

To configure your SSB unit as the secondary node of a HA cluster

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

TIP: The SSB console displays the IP address the external interface is listening on. SSB 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. On the initial **Configuration** screen, select **HA IP configuration**.

Figure 17: Select HA IP configuration

Welcome — License — **Networking** — Users — Certificate — Finish

Welcome to syslog-ng Store Box

Configuration

Standalone or primary node configuration

HA IP configuration

HA IP Address:

Next

3. Enter the HA IP Address that you [previously set](#) in the **Interface IP** field on **Other node**.

Figure 18: Enter the HA IP address for the secondary node

Welcome — License — **Networking** — Users — Certificate — Finish

Welcome to syslog-ng Store Box

Configuration

Standalone or primary node configuration

HA IP configuration

HA IP Address:

Next

4. Click **Next**.

The Welcome Wizard displays the confirmation that you have successfully configured your secondary node.

Figure 19: HA IP configuration successful



5. (Optional) To modify your configured HA IP Address, click **Back**.
6. If you do not want to change anything in your configuration, you can continue to [convert your nodes into a HA cluster](#).

Converting your primary node and secondary node to a HA cluster on the SSB web interface

1. (Optional) If you have not done it previously, reboot the SSB unit that you [previously configured as the primary node](#) for your HA cluster.
2. Log in to the SSB unit configured as the primary node, and navigate to **Basic Settings > High Availability**.
Under **High availability & Nodes**, both configured nodes are displayed, both in **STANDALONE** HA state.
3. Click **Convert to Cluster**.
Your **Cluster status** will display that you are in **CONVERTED** HA state.
4. Continue by either shutting down, and then powering up your HA cluster, or shutting down, and then powering up your nodes one by one:
 - **Rebooting the HA cluster:** If you do not want to closely monitor shutting down, and then rebooting your nodes separately, click **Reboot cluster**.
You will have to log in to the SSB web interface again.
 - **Shutting down, then powering up your nodes separately:**
 1. Click **Shutdown** on your secondary node (**Other node**).
While the node of your choice is shutting down, your **Cluster status** will display that you are in **DEGRADED** HA state.
 2. Click **Reboot** on the primary node (**This node**).

You will have to log in to the SSB web interface again.

3. Power up your secondary node (**Other node**).

5. Log in to the SSB unit configured as your primary node, and navigate to **Basic Settings > High Availability**.

While SSB is synchronizing the newly rebooted nodes, your **Cluster status** will display that you are in **DEGRADED SYNC** HA state. Depending on your configuration, synchronization may take a while.

When SSB successfully finishes synchronizing your nodes, your **Cluster status** displays that SSB is operating in **HA**.

6. (Optional) After your HA cluster is in **HA** state, you can change the configuration settings on your nodes if you want to.

Basic settings

syslog-ng Store Box (SSB) is configured via the web interface. Configuration changes take effect automatically after clicking **Commit**. Only the modifications of the current page or tab are activated — each page and tab must be committed separately.

- For the list of supported browsers, see [Supported web browsers](#).
- For a description of the web interface of SSB, see [The structure of the web interface](#).
- To configure network settings, see [Network settings](#).
- To configure date and time settings, see [Date and time configuration](#).
- To configure system logging and e-mail alerts, see [SNMP and e-mail alerts](#).
- To configure system monitoring, see [Configuring system monitoring on SSB](#).
- To configure data and configuration backups, see [Data and configuration backups](#).
- To configure archiving and clean-up, see [Archiving and cleanup](#).
- For a description of the backup and archiving protocols, see [Data and configuration backups](#).

Supported web browsers

The syslog-ng Store Box (SSB) web interface can be accessed only using TLS encryption and strong cipher algorithms. The browser must support HTTPS connections, JavaScript, and cookies. Make sure that both JavaScript and cookies are enabled.

NOTE: SSB displays a warning message if your browser is not supported or JavaScript is disabled.

If you have successfully accessed the SSB web interface using HTTPS at least once, your browser will remember this, and on any subsequent occasions, it will force you to access SSB using HTTPS, even if you try loading it through an HTTP connection. This is thanks to the HTTP Strict Transport Security (HSTS) policy, which enables web servers to enforce web browsers to restrict communication with the server over an encrypted SSL/TLS connection for a set period. Web servers declare the HSTS policy using a special Strict-Transport-Security response header field.

This might, however, cause issues in any of the following cases:

- When the SSL certificate of SSB's web interface has expired. In this case, any attempt to access the web interface using a secure connection will fail with an error message.
- When you switch the trusted CA-signed certificate to a self-signed certificate for SSB's web interface. As per HSTS design, a self-signed certificate is not taken to have been issued by a trusted CA, therefore any secure connections to the SSB web interface will fail with an error message.

The resolution to the above-mentioned issues is to:

- Remove the HSTS settings in your browser. This must be done locally, in a browser-specific way. For detailed instructions, consult the support site of the browser you are using.
- OR
- Upload a new certificate, using a different browser on a different machine. For detailed instructions on how to upload external certificates to SSB, see "[Uploading external certificates to SSB](#)" in the [Administration Guide](#).

Supported browsers:

Mozilla Firefox 52 ESR

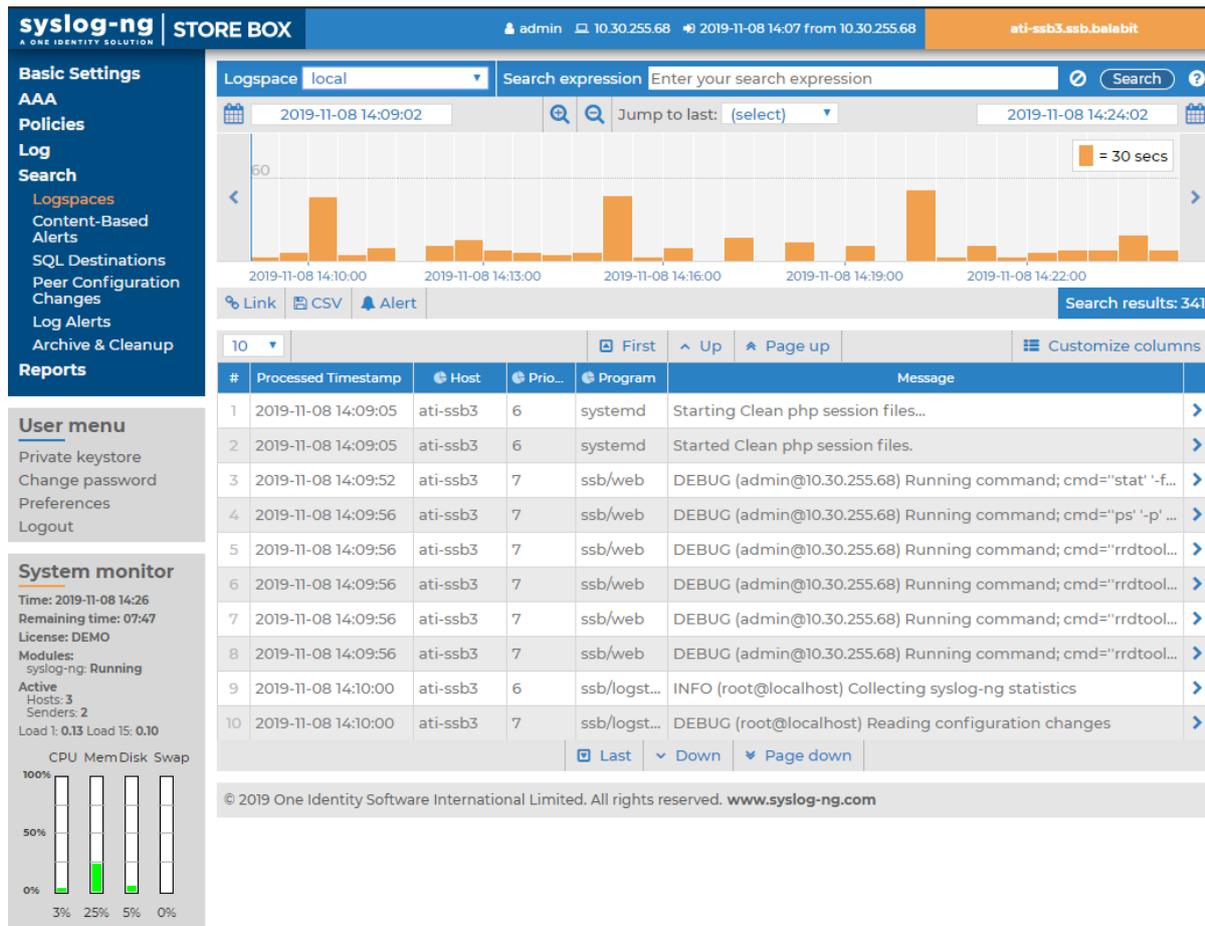
We also test SSB on the following, unsupported browsers. The features of SSB are available and usable on these browsers as well, but the look and feel might be different from the supported browsers. Internet Explorer 11, Microsoft Edge, and the currently available version of Mozilla Firefox and Google Chrome.

The structure of the web interface

The web interface consists of the following main sections:

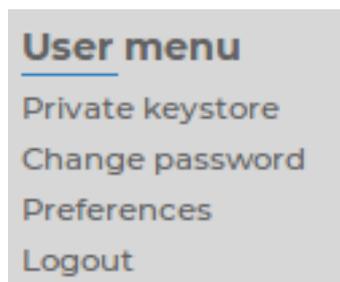
Main menu: Each menu item displays its options in the main workspace on one or more tabs. Click a menu item to display the list of available tabs.

Figure 20: Structure of the web interface



User menu: Provides possibilities to change your syslog-ng Store Box (SSB) password, to log out, and disable confirmation dialogs and tooltips using the **Preferences** option.

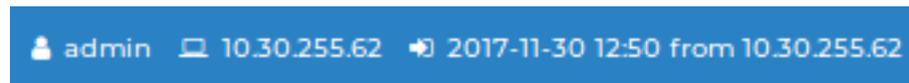
Figure 21: User menu



User info: Provides information about the user currently logged in:

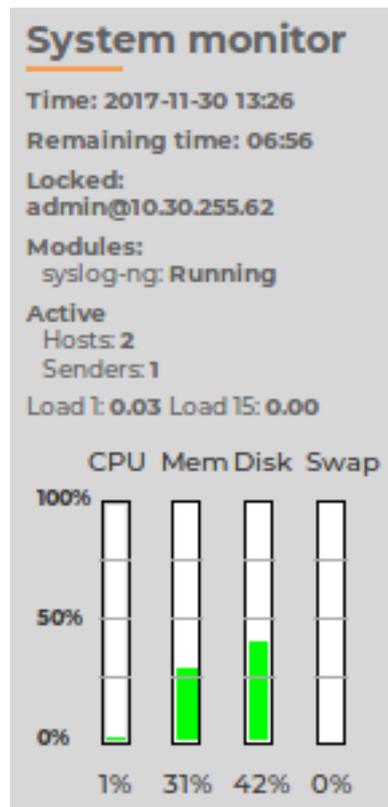
-  username
-  IP address of the user's computer
-  date and IP address of the user's last login

Figure 22: User info



System monitor: Displays accessibility and system health information about SSB, including the following:

Figure 23: System monitor



- **Time:** System date and time.
- **Remaining time:** The time remaining before the session to the web interface times out.
NOTE: To change timeout settings, navigate to **Basic Settings > Management > Web interface and RPC API settings > Session timeout** and enter the timeout value in minutes.
- **Locked:** Indicates that the interface is locked by another administrator (for details, see [Multiple web users and locking](#) on page 71).
- **Modules:** The status of syslog-ng running on SSB (ideally it is RUNNING).
- **License:** License information if the license is not valid, or an evaluation version license has expired.

- **Raid status:** The status of the RAID devices, if synchronization between the disks is in progress.
- **Active:**
 - **Hosts:** the number of clients (log source hosts) where the log messages originate from (for example computers)
 - **Senders:** the number of senders where the log messages directly come from (for example, relays)

Example: Number of hosts and senders

For example: if 300 clients all send log messages directly to SSB the Hosts and Senders are both 300.

If the 300 clients send the messages to 3 relays (assuming that the relays do not send messages themselves) and only the relays communicate directly with SSB then Hosts is 300, while Senders is 3 (the 3 relays).

If the relays also send messages, then Hosts is 303, while Senders is 3 (the 3 relays).

- **HA:** The HA status and the ID of the active node if two SSB units are running in a High Availability cluster. If there are redundant Heartbeat interfaces configured, their status is displayed as well. If the nodes of the cluster are synchronizing data between each other, the progress and the time remaining from the synchronization process is also displayed.
- Average system load during the
 - **Load 1:** last minute
 - **Load 15:** last fifteen minutes
- CPU, memory, hard disk, and swap use. Hover the mouse above the graphical bars to receive a more details in a tooltip, or navigate to **Basic Settings > Dashboard** for detailed reports.

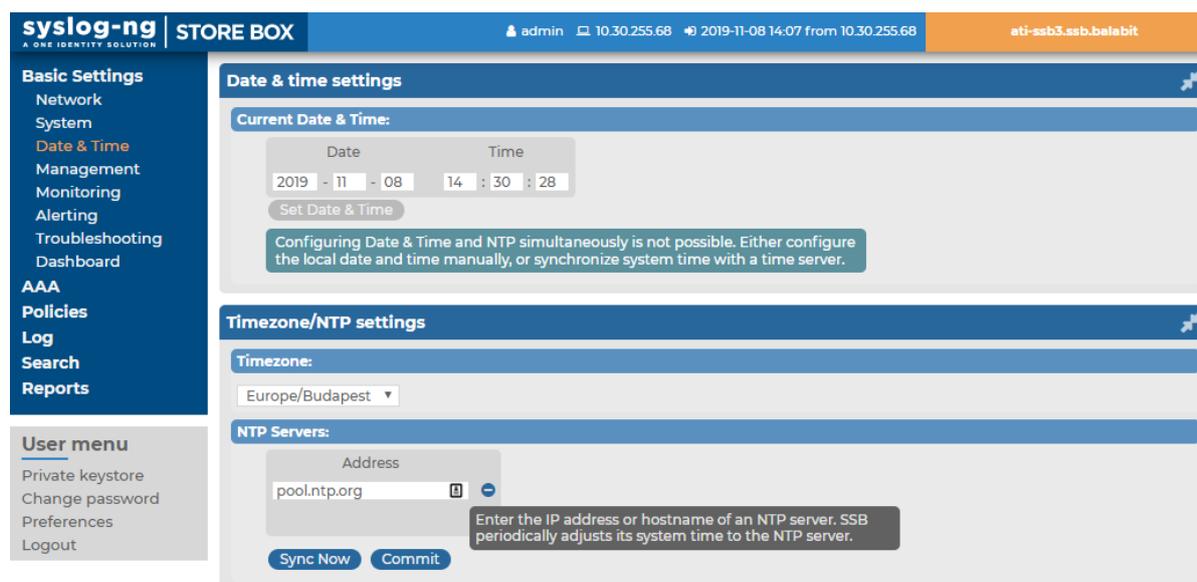
NOTE: If you have installed SSB from Azure, the swap column is not available, because in this case, swap memory is not used.

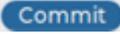
The System monitor displays current information about the state of SSB. To display a history of these parameters, go to **Basic Settings > Dashboard**. For details, see [Status history and statistics](#) on page 368.

Elements of the main workspace

The main workspace displays the configuration settings related to the selected main menu item.

Figure 24: Main workspace



-  **Commit** Each page includes one or more blue action buttons. The most common action button is the , 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).
-  **Open/collapse lists**: Open or close a list of options (for example, the list of available reports).
-  **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 popup 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 log paths) is important. For example, when syslog-ng Store Box (SSB) evaluates log paths, it looks at the log paths in descending order.

Message window: This popup window displays the responses of SSB 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 menu. To make the window appear only for failed actions, navigate to **User menu > Preferences** and enable the **Autoclose successful commit messages** option.

Figure 25: Message window

Result



Configuration saved successfully

OK

Multiple web users and locking

Multiple administrators can access the syslog-ng Store Box (SSB) web interface simultaneously, but only one of them can modify the configuration. This means that the configuration of SSB is automatically locked when the first administrator who can modify the configuration accesses a configuration page (for example, the **Basic Settings**, **AAA**, or **Logs** menu). The username and IP address of the administrator locking the configuration is displayed in the **System Monitor** field. Other administrators must wait until the locking administrator logs out, navigates to a page that is not concerned with modifying the configuration (for example, the **Search** page), or the session of the administrator times out. However, it is possible to access the **Search** and **Reporting** menus, or browse the configuration with only View rights (for details, see [Managing user rights and usergroups](#) on page 127).

NOTE: If an administrator logs in to SSB using the local console or a remote SSH connection, access via the web interface is completely blocked. Inactive local and SSH connections time out just like web connections. For details, see [Accessing the SSB console](#) on page 152.

Web interface and RPC API settings

syslog-ng Store Box (SSB) prevents brute force attacks when logging in. If you repeatedly try logging in to SSB using incorrect login details within a short period of time (10 times within 60 seconds), the source IP gets blocked on UI destination port 443 for 5 minutes. Your browser displays an **Unable to connect** page.

By default, SSB terminates the web session of a user after ten minutes of inactivity. To change this timeout value, adjust the **Basic Settings > Management > Web interface and RPC API settings > Session timeout** option.

In addition to controlling the web session timeout value, you can also specify the cipher suites to be permitted in the HTTPS connection.

The **Basic Settings > Management > Web interface and RPC API settings > Cipher suite** option allows you to choose the strength of the allowed cipher suites using one of the following options:

- **Compatible:** It is a large set of cipher suites determined by the following cipher string:

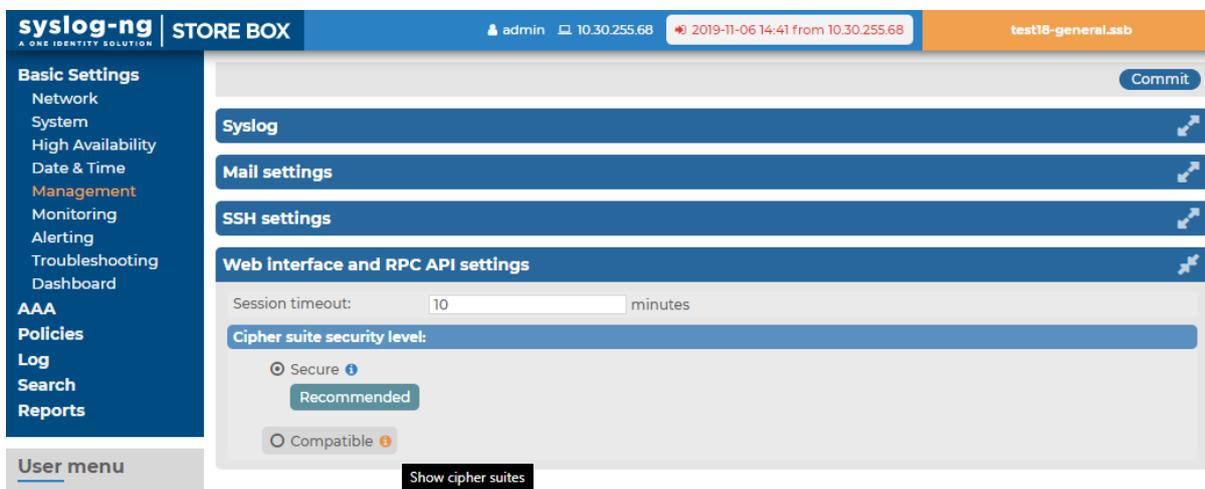
```
ALL:!aNULL:!eNULL
```

The Compatible setting may allow permitting (and hence not safe) cipher suites for the Transport Layer Security (TLS) negotiations.

- **Secure:** A smaller and more strict set of cipher suites where vulnerable cryptographic algorithms are eliminated. This cipher suite set is determined by the following cipher string:

```
HIGH:!COMPLEMENTOFDEFAULT:!aNULL:!eNULL:!DHE-RSA-AES128-SHA:!DHE-RSA-AES256-SHA:!ECDHE-RSA-AES128-SHA:!ECDHE-RSA-AES256-SHA:!AES128-SHA:!AES256-SHA
```

Figure 26: Basic Settings > Management > Web interface and RPC API settings – Set session timeout and Cipher suite



Network settings

The **Basic Settings > Network** tab contains the network interface and naming settings of syslog-ng Store Box (SSB).

Figure 27: Basic Settings > Network > Interfaces — Network settings

The screenshot displays the 'Network settings' page in the SSB web interface. On the left, there are three main navigation areas: 'Basic Settings' (Network, System, Date & Time, Management, Monitoring, Alerting, Troubleshooting, Dashboard), 'User menu' (Private keystore, Change password, Preferences, Logout), and 'System monitor' (Time: 2020-04-30 10:04, Remaining time: 06:18, Locked: admin@10.12.15.134, Modules: syslog-ng: Running, Active Hosts: 2, Senders: 1, Load 1: 0.20 Load 15: 0.02). The main content area is titled 'Interfaces' and contains several sections: 'External interface' with fields for IPv4 (10.12.165.27, 255.255.248.0) and IPv6 (fec0:b190::10:12:165:27, 32) addresses, a 'Management enabled' checkbox, 'Management interface' with an 'Enable management interface' checkbox, 'Routing table' with columns for Address, Netmask, and Gateway, 'Naming' with fields for Hostname (esx-ip27), Nick name, DNS search domain (ssb.balabit), Primary DNS server (10.12.192.1), and Secondary DNS server, and 'Ethernet links' showing status for Internal, External (10000Mb/s), Management, and HA interfaces.

- **Interfaces > External interface:** The address and netmask / prefix of the SSB network interface that receives client connections. Click the and icons to add new alias IP addresses (also called alias interfaces) or delete existing ones. You must configure at least one external interface. If you disable the management interface, you can also access the SSB web interface through the external interface. When multiple external interfaces are configured, the first one refers to the physical network interface, all others are alias interfaces. You can access the SSB web interface from all external interfaces (if no management interface is configured).

Optionally, you can enable access to the SSB web interface even if the management interface is configured by activating the **Management enabled** function.

⚠ CAUTION:

If you enable management access on an interface and configure alias IP address(es) on the same interface, SSB will accept management connections only on the original address of the interface.

NOTE: The IPv6 support on syslog-ng Store Box (SSB) only enables management and UI access, but SSB cannot send or accept logs through IPv6 addresses. As a result, the external interface must have at least one IPv4 address configured.

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

- IPv4 addresses
 - 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
 - 127.0.0.0/8 (localhost IP addresses)
- IPv6 addresses
 - ::1/128 (loopback address)
 - fe80::/10 (link-local addresses)

NOTE: The speed of the interface is displayed for every interface. In SSB version 4 F5 and later, you cannot manually change the speed of the interface.

On SSB T-10 appliances, if both the 1-Gb (label 1) and 10-Gb (label A) interfaces are plugged in, SSB displays the auto-detected speed of the interface where Ethernet link is detected (that is, the cable is plugged in, and the other side is powered on).

When SSB is deployed in a virtual environment and only a single network interface is configured, that single network interface starts to serve as the management interface. In such cases, the **Management enabled** function becomes redundant. Instead, a message displays that access to the web interface and the RPC API is enabled on every configured IP address.

Figure 28: Basic Settings > Network > Interfaces — Management enabled on every configured IP address

The screenshot displays three configuration sections:

- External interface:** A table with columns 'Address (IPv4 / IPv6)' and 'Netmask (IPv4) / Prefix (IPv6)'. It contains two rows: '10.12.165.27' with netmask '255.255.248.0' and 'fec0:b190::10:12:165:27' with netmask '32'. A 'Management enabled:' checkbox is checked.
- Management interface:** An 'Enable management interface:' checkbox is checked. Below it is a table with 'Address (IPv4 / IPv6)' and 'Netmask (IPv4) / Prefix (IPv6)' columns, containing 'fec0:a160::10:12:165:27' and '32'.
- Routing table:** A table with columns 'Address (IPv4 / IPv6)', 'Netmask (IPv4) / Prefix (IPv6)', and 'Gateway (IPv4 / IPv6)'. It contains three rows: '0.0.0.0' with gateway '10.12.167.254', '::' with gateway 'fec0:a160::1', and '::' with gateway 'fec1:b190::1'. A 'Reply on same interface:' checkbox is unchecked.

- **Interfaces > Management interface:** The address and netmask / prefix of the SSB network interface used to access the SSB web interface. If the management interface is configured, the web interface can be accessed only through this interface, unless:
 - Access from other interfaces is explicitly enabled.
 - Only one network interface has been defined, which then serves as the management interface.

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

- IPv4 addresses
 - 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
 - 127.0.0.0/8 (localhost IP addresses)
- IPv6 addresses
 - ::1/128 (loopback address)
 - fe80::/10 (link-local addresses)
- **Interfaces > Routing table:** When sending a packet to a remote network, SSB consults the routing table to determine the path it should be sent. If there is no information in the routing table, 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 the internal interface is connected to multiple subnets, because the default gateway is (usually) towards the external interface. Click the **+** and **x** icons to add new routes or delete existing ones. A route means that messages sent to the **Address (IPv4/IPv6) / Netmask (IPv4) / Prefix (IPv6)** network

should be delivered to **Gateway (IPv4 / IPv6)**. An option is also provided to override the default behavior of always routing outgoing packets based on the destination address and instead reply on the interface of the incoming packets.

For more information, see [Configuring the routing table](#) on page 78.

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

Configuring the external interface

This section describes how to activate the external interface.

To activate the external interface

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

Figure 29: Basic Settings > Network > Interfaces > External interface — Configuring the external interface

Address (IPv4 / IPv6)	Netmask (IPv4) / Prefix (IPv6)
10.12.165.27	255.255.248.0
fec0:b190::10:12:165:27	32

Management enabled:

2. Under **External interface**, select **Management enabled**.
3. **Address (IPv4 / IPv6):** enter the IP address of SSB's external interface.
4. **Netmask (IPv4) / Prefix (IPv6):** enter the netmask / prefix related to the IP address.

NOTE: The IPv6 support on syslog-ng Store Box (SSB) only enables management and UI access, but SSB cannot send or accept logs through IPv6 addresses. As a result, the external interface must have at least one IPv4 address configured.

NOTE: When entering IPv6 addresses, consider that the IPv6 address configured on the external interface and the IPv6 address configured on the management interface cannot be the same.

5. Click **Commit**.

Configuring the management interface

This section describes how to activate the management interface.

NOTE: When syslog-ng Store Box (SSB) is deployed in a virtual environment and only a single network interface is configured, that single network interface starts to serve as the management interface. In such cases, the **Enable management interface** function becomes redundant and is not displayed on the user interface.

To activate the management interface

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

Figure 30: Basic Settings > Network > Interfaces > Management interface – Configuring the management interface



Management interface:

Enable management interface:

Address (IPv4 / IPv6) Netmask (IPv4) / Prefix (IPv6)

fec0:a160::10:12:165:27 32

2. Under **Management interface**, select **Enable management interface**.
3. **Address (IPv4 / IPv6)**: enter the IP address of SSB's management interface.
4. **Netmask (IPv4) / Prefix (IPv6)**: enter the netmask / prefix related to the IP address.

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

- IPv4 addresses
 - 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
 - 127.0.0.0/8 (localhost IP addresses)
- IPv6 addresses
 - ::1/128 (loopback address)
 - fe80::/10 (link-local addresses)

NOTE: When entering IPv6 addresses, consider that the IPv6 address configured on the external interface and the IPv6 address configured on the management

interface cannot be the same.

5. **CAUTION:**

After clicking **Commit**, the web interface will be available only through the management interface. The web interface will not be accessible using the current (external) interface, unless the Management enabled option is selected for the external interface.

Make sure that the Ethernet cable is plugged in and that the management interface is connected to the network by checking the status of the Ethernet links in **Basic settings > Network > Ethernet links**. If the cable is plugged in and the interface is connected to the network, the  icon and the connection speed is displayed next to **Link:** in the respective interface field.

When using High Availability, ensure that the management interface of both SSB units is connected to the network.

The HA interface section indicates if a link is detected on the high availability interface.

Click **Commit**.

Configuring the routing table

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

You can add multiple addresses along with their respective gateways.

CAUTION:

Complete the following procedure only if the management interface is configured, otherwise the data sent by SSB will be lost. For details on configuring the management interface, see [Configuring the management interface on page 77](#).

To configure the routing table

1. To add a new routing entry, navigate to **Basic Settings > Network > Interfaces** and in the **Routing table** field, click .

Figure 31: Basic Settings > Network > Interfaces > Routing

Address (IPv4 / IPv6)	Netmask (IPv4) / Prefix (IPv6)	Gateway (IPv4 / IPv6)
0.0.0.0	0.0.0.0	10.12.167.254
::	0	fec0:a160::1
::	0	fec1:b190::1

Reply on same interface:

2. Enter the IP address of the remote server into the **Address (IPv4 / IPv6)** field.
3. Enter the related netmask /prefix into the **Netmask (IPv4) / Prefix (IPv6)** field.
4. Enter the IP address of the gateway used on that subnetwork into the **Gateway (IPv4 / IPv6)** field.

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

- IPv4 addresses
 - 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
 - 127.0.0.0/8 (localhost IP addresses)
- IPv6 addresses
 - ::1/128 (loopback address)
 - fe80::/10 (link-local addresses)

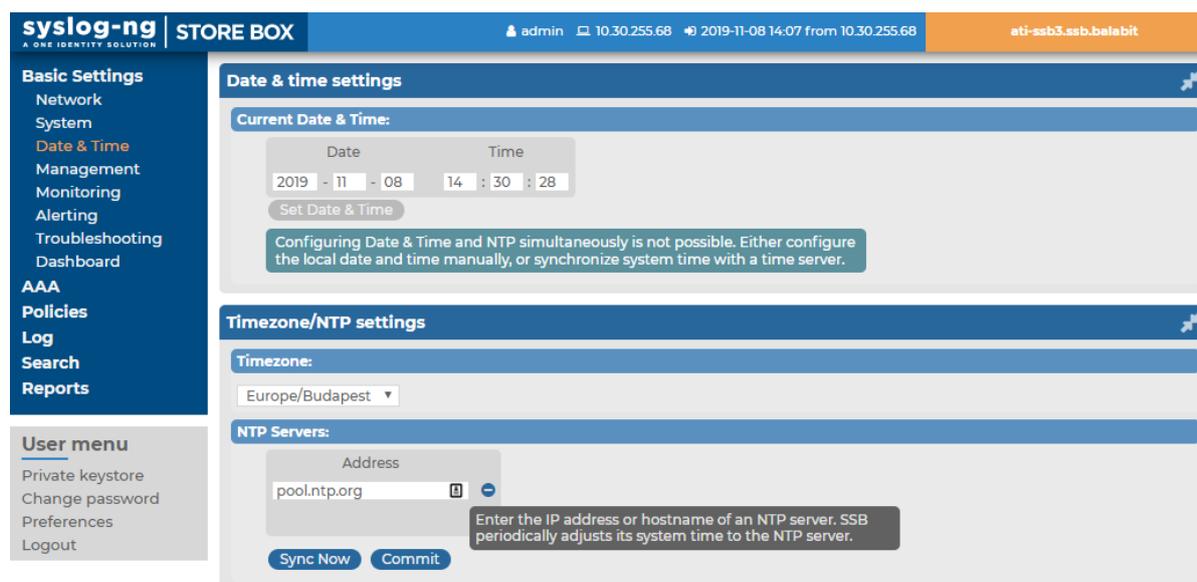
5. If you want to reply on the same interface where a packet came in, select **Reply on same interface**. This instructs SSB to disregard connected networks other than the network of the incoming packet's interface when routing reply packets.
6. Click **Commit**.

Date and time configuration

You can configure date and time-related settings of syslog-ng Store Box (SSB) on the **Date & Time** tab of the **Basic** page.

NOTE: Configuring Date & Time and NTP simultaneously is not possible. Either configure the local date and time manually, or synchronize system time with a time server.

Figure 32: Basic Settings > Date & Time — Set date and time



CAUTION:

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

SSB 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 SSB, enter the current date into respective fields of the **Date & Time Settings** group and click **Set Date & Time**.

NOTE:

If the time of SSB 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** to sync the time immediately using SNTP.

Configuring a time (NTP) server

This section describes how to retrieve the date automatically from a time server.

CAUTION:

One Identity recommends not changing the timezone, because logspace rotation is based on your currently configured local timezone. If you change the timezone, you will not be able to search in your previously stored logs. Before changing the timezone, [contact our Support Team](#).

To retrieve the date automatically from a time server

1. Select your timezone in the **Timezone** field.
2. Enter the IP address of an NTP time server into the **Address** field.
To add new servers or delete existing ones, click  and , respectively.

NOTE: Use an NTP server of high time accuracy. SSB needs high time accuracy for processing its logs with as exact timestamps as possible. Any inaccuracy will be detected by SSB, and the server will be rejected.

3. Click **Commit**.
4. To sync the time immediately using SNTP, click **Sync now**.

NOTE: If your local system time is different from the NTP server time, it can result in a time gap in the time stamp of the logs before synchronization and the logs after synchronization. The time gap may appear longer than the time that has actually passed. Therefore, logs collected before the synchronization might have time stamps that do not correspond to the NTP server time.

SNMP and e-mail alerts

The following sections describe how you can configure e-mail and SNMP alerts on syslog-ng Store Box (SSB).

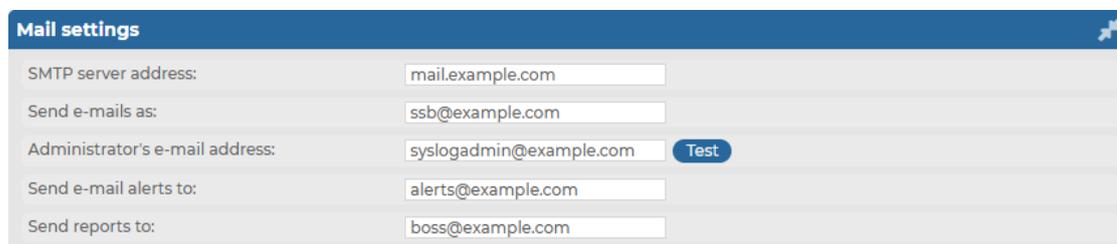
Configuring e-mail alerts

This section describes how to configure e-mail alerts.

To configure e-mail alerts

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

Figure 33: Basic Settings > Management > Mail settings — Configure e-mail sending



Mail settings 	
SMTP server address:	<input type="text" value="mail.example.com"/>
Send e-mails as:	<input type="text" value="ssb@example.com"/>
Administrator's e-mail address:	<input type="text" value="syslogadmin@example.com"/> Test
Send e-mail alerts to:	<input type="text" value="alerts@example.com"/>
Send reports to:	<input type="text" value="boss@example.com"/>

3. 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. syslog-ng Store Box (SSB) 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.
4. Enter the e-mail address of the administrator into the **Administrator's e-mail address** field. SSB sends notifications related to system-events (but not alerts and reports) to this address.
5. Enter the e-mail address of the administrator into the **Send e-mail alerts to** field. SSB sends monitoring alerts to this address.
6. Enter the e-mail address the person who should receive traffic reports from SSB into the **Send reports to** field. For details on reports, see [Reports](#) on page 371.

⚠ CAUTION:

To get alert e-mails, provide an e-mail address in this field. Sending alerts fails if these settings are incorrect, since the alerting e-mail address does not fall back to the administrator's e-mail address by default.

7. Click **Commit**.
8. Click **Test** to send a test message.
If the test message does not arrive to the server, check if SSB can access the server. For details, see [Troubleshooting SSB](#) on page 424.
9. Select in which situations SSB should send an e-mail alert. For details, see [Configuring system monitoring on SSB](#) on page 87.
10. Click **Commit**.

Configuring SNMP alerts

syslog-ng Store Box (SSB) can send alerts to a central monitoring server via SNMP (Simple Network Management Protocol). This section describes how you can configure SNMP trap settings and SNMP alerts.

To configure SNMP trap settings

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

Figure 34: Basic Settings > Alerting > SNMP trap settings — Configure SNMP alerts using SNMP v2c

The screenshot shows the 'SNMP trap settings' configuration page. At the top, the title is 'SNMP trap settings'. Below it, there is a field for 'SNMP server address' with the value '10.50.0.2' and a help icon. A tooltip below this field says 'Enter the IP address or hostname of the SNMP server.' There are two radio buttons: 'SNMP v2c' (which is selected) and 'SNMP v3'. Below these is a 'Community' field with the value 'public'.

3. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c**, and enter the community to use into the **Community** field.
- To use the SNMP v3 protocol, select **SNMP v3** and complete the following steps:

Figure 35: Basic Settings > Alerting > SNMP trap settings — Configure SNMP alerts using SNMP v3

The screenshot shows the 'SNMP trap settings' configuration page for SNMP v3. The title is 'SNMP trap settings'. The 'SNMP server address' field is '10.50.0.2' with a tooltip 'Enter the IP address or hostname of the SNMP server.' The 'SNMP v3' radio button is selected. Below are fields for 'Username' (value: 'demo') and 'Engine ID' (value: '0xFFFFFFFF'). There are two sections: 'Authentication method' with 'SHA1' selected and an 'Authentication password' field (value: '.....'); and 'Encryption method' with 'AES' selected and an 'Encryption password' field (value: '.....').

- a. Enter the username to use into the **Username** field.
- b. 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.
- c. Select the authentication method (**SHA1**) to use from the **Authentication method** field.
- d. Enter the password to use into the **Authentication password** field.

- e. Select the encryption method (**Disabled** or **AES**) to use from the **Encryption method** field.

The supported AES method is AES-128.

- f. In the case of AES, enter the encryption password to use into the **Encryption password** field.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#\$\$%&'()*+,-./:;<=>?@[]^_`{|}

4. Click **Commit**.
5. Select in which situations SSB should send an SNMP alert. For details, see [Configuring system monitoring on SSB](#) on page 87.
6. Click **Commit**.

Querying SSB status information using agents

External SNMP agents can query the basic status information of syslog-ng Store Box (SSB). This section describes how you can configure which clients can query SSB's basic status information.

To configure which clients can query this information

1. Navigate to **Basic Settings > Monitoring > SNMP agent settings**.

Figure 36: Basic Settings > Monitoring > SNMP agent settings — Configure SNMP agent access (only SNMP v2c agent enabled)

The screenshot shows the 'SNMP agent settings' configuration page. The 'Client address' is set to '10.50.0.2'. The 'SNMP v2c agent' checkbox is checked, while 'SNMP v3 agent' is unchecked. The 'Community' is set to 'public'. Below the settings is a table for 'Agent access' with the following data:

Node	Community (v2c)
active	public
087AC40C-7235-0706-0809-0CC47A083572	087AC40C7235070608090CC47A083572
017AC40C-AA59-0706-0809-0CC47A0159AA	017AC40CAA59070608090CC47A0159AA

Figure 37: Basic Settings > Monitoring > SNMP agent settings — Configure SNMP agent access (only SNMP v3 agent enabled)

SNMP agent settings

Client address: 10.50.0.2

System location:

System contact:

System description:

SNMP v2c agent:

SNMP v3 agent:

Allowed users

Username	Auth. method	Auth. password	Encryption method	Encryption password
snmpagent	SHA1	AES

Agent access:

Node	Context (v3)
active	...

Figure 38: Basic Settings > Monitoring > SNMP agent settings — Configure SNMP agent access (both SNMP v2c and SNMP v3 agent enabled)

SNMP agent settings

Client address: 10.50.0.2

System location:

System contact:

System description:

SNMP v2c agent:

Community: public

SNMP v3 agent:

Allowed users

Username	Auth. method	Auth. password	Encryption method	Encryption password
snmpagent	SHA1	AES

Agent access:

Node	Community (v2c)	Context (v3)
active	public	...

2. The status of SSB can be queried dynamically via SNMP. By default, the status can be queried from any host. To restrict access to these data to a single host, enter the IP address of the host into the **Client address** field.
3. Optionally, you can enter the details of the SNMP server into the **System location**, **System contact**, and **System description** fields.
4. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c agent**, and enter the community to use into the **Community** field.

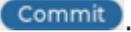
By default, information about SSB is available using the **public** community. If you are using a high-availability SSB cluster, then each node provides information about its own status using a specific community. This community is the Node ID of the node (as displayed in the **Basic Settings > High Availability > This node > Node ID** field).

Agent access:		
Node active	Community (v2c)	Context (v3)
087AC40C-7235-0706-0809-0CC47A083572	public	""
017AC40C-AA59-0706-0809-0CC47A0159AA	087AC40C7235070608090CC47A083572	087AC40C7235070608090CC47A083572
	017AC40CAA59070608090CC47A0159AA	017AC40CAA59070608090CC47A0159AA

- To use the SNMP v3 protocol, select **SNMP v3 agent** and complete the following steps:
 - Click .
 - Enter the username used by the SNMP agent into the **Username** field.
 - Select the authentication method (**MD5 or SHA1**) to use from the **Auth. method** field.
 - Enter the password used by the SNMP agent into the **Auth. password** field.
 - Select the encryption method (**Disabled, DES or AES**) to use from the **Encryption method** field.

The supported AES method is AES-128.
 - Enter the encryption password to use into the **Encryption password** field.
 - To add other agents, click .

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

- Click .

View node ID and community

This section provides information about monitoring primary and secondary nodes on syslog-ng Store Box (SSB) with the help of a node ID and community.

To monitor your primary and secondary node, you need the related community ID. Navigate to **Basic Settings > Monitoring > SNMP agent settings > Agent access**. This information is available if **SNMP v2c Agent** or **SNMP v3 Agent** is selected.

Figure 39: Basic Settings > Monitoring > SNMP agent settings > Agent access — Configure SNMP agent access

Agent access:		
Node	Community (v2c)	Context (v3)
active	public	""
087AC40C-7235-0706-0809-0CC47A083572	087AC40C7235070608090CC47A083572	087AC40C7235070608090CC47A083572
017AC40C-AA59-0706-0809-0CC47A0159AA	017AC40CAA59070608090CC47A0159AA	017AC40CAA59070608090CC47A0159AA

If either SNMP Agent is selected and your cluster is in HA state, both nodes appear under **Agent access**. In this case, the first row is the default community, the next rows are the nodes.

The community IDs are displayed under the respective **Community (v2c)** and **Context (v3)** columns.

You will use these communities in formulating SNMP queries. The **Node** names are used to indicate which SNMP objects can be queried from which community in [Monitoring SSB](#) on page 391.

To view the availability of the nodes, navigate to **Basic Settings > High Availability**. For details on what type of information you can see on that page, see [Managing a high availability SSB cluster](#) on page 135. The **Node ID** of the node is the same as the ID under **Agent access**

To gather more information your nodes, query the SSB-SNMP-MIB::ssbHAClusterStatus object. For details, see [The status of the HA cluster](#) on page 404.

If the query is not responding, it can mean that the target node is down or restarting. In this case, check the node status manually and attempt solving the issue, or [contact our Support Team](#).

NOTE: If SSB's core firmware is not accessible or down, SNMP queries will not work.

Configuring system monitoring on SSB

The following sections provide information about configuring system monitoring on syslog-ng Store Box (SSB).

Configuring SNMP agent settings

syslog-ng Store Box (SSB) continuously monitors a number of parameters of the SSB hardware and its environment. If a parameter reaches a critical level (set in its respective **Maximum** field), SSB sends e-mail and SNMP messages to alert the administrator.

SSB sends SNMP alerts using the external network interface by default, or using the management interface if it is enabled. SSB supports the SNMPv2c and SNMPv3 protocols. The SNMP server set on the **Alerting** tab can query status information from SSB.

TIP: To have your central monitoring system recognize the SNMP alerts sent by SSB, select **Basic Settings > Monitoring > Download MIBs** to download the SSB-specific Management Information Base (MIB), then import it into your monitoring system.

Figure 40: Basic Settings > Monitoring > SNMP agent settings

SNMP agent settings

Client address: 10.50.0.2

System location: Enter the IP address or network address of the hosts that can query SSB over SNMP.

System contact:

System description:

SNMP v2c agent:

Community: public

SNMP v3 agent:

Allowed users

Username	Auth. method	Auth. password	Encryption method	Encryption password
snmpagent	SHA1	AES

Agent access:

Node	Community (v2c)	Context (v3)
active	public	...
087AC40C-7235-0706-0809-0CC47A083572	087AC40C7235070608090CC47A083572	087AC40C7235070608090CC47A083572
017AC40C-AA59-0706-0809-0CC47A0159AA	017AC40CAA59070608090CC47A0159AA	017AC40CAA59070608090CC47A0159AA

The following sections describe the parameters you can receive alerts on.

- For details on health-monitoring alerts, see [Health monitoring](#) on page 88.
- For details on system-monitoring alerts, see [System related traps](#) on page 94.

For details on monitoring SSB with an external central monitoring system, see [Monitoring SSB](#) on page 391.

Health monitoring

This section provides information about health monitoring, SNMP trap settings, and health related traps.

Health monitoring

To configure health monitoring on syslog-ng Store Box (SSB), navigate to Basic Settings > Alerting.

Figure 41: Basic Settings > Alerting – Health monitoring

Health monitoring

Disk utilization maximum:	80	%
Load 1 maximum:	5	
Load 5 maximum:	4	
Load 15 maximum:	3	
Swap utilization maximum:	70	%

- **Disk utilization maximum:** Ratio of free space available on the hard disk. SSB sends an alert if the log files use more space than the set value. Archive the log files to a backup server to free disk space. For details, see [Archiving and cleanup](#) on page 110.

NOTE: The alert message includes the actual disk usage, not the limit set on the web interface. For example, you set SSB to alert if the disk usage increases above 10 percent. If the disk usage of SSB 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 1|5|15 maximum:** The average load of SSB during the last one, five, or 15 minutes.
- **Swap utilization maximum:** Ratio of the swap space used by SSB. SSB sends an alert if it uses more swap space than the set value.

Health related traps and SNMP trap settings

For health related alerts, SSB sends SNMP alerts using the management network interface by default, or using the external interface if the management interface is disabled. SSB supports the SNMPv2c and SNMPv3 protocols. The SNMP server set on the **Alerting** tab can query status information from SSB.

To configure SNMP trap settings

1. Navigate to **Basic Settings > Alerting > SNMP trap settings**.

Figure 42: Basic Settings > Alerting > SNMP trap settings.

2. If you want to configure SNMP v2c trap settings, select **SNMP v2c**, provide the SNMP server address you want to use, and enter the name of the SNMP community you want to use in the **Community:** field.

If you want to configure SNMP v3 trap settings, select **SNMP v3**, provide the SNMP server address, the **Username** and the **Engine ID** you want to use, then configure the **Authentication method** and **Encryption method** you want to use.

3. Click **Commit**.

To enable notifications about health related issues

1. Navigate to **Basic Settings > Alerting > Health related traps** and enable alert **Disk usage is above the defined ratio**.

Figure 43: Basic Settings > Alerting > Health related traps – Enable alert

Description	Name	Email	SNMP
Disk usage is above the defined ratio	xcbDiskFull	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>

2. Click **Commit**.

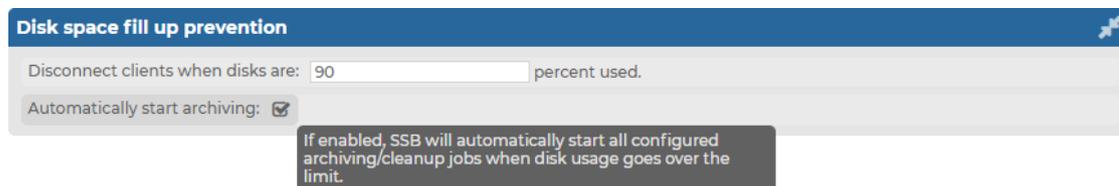
Preventing disk space fill up

The following section describes how to prevent disk space from filling up.

To prevent disk space from filling up

1. Navigate to **Basic Settings > Management > Disk space fill up prevention**.

Figure 44: Basic Settings > Management > Disk space fill up prevention



2. Set the limit of maximum disk utilization in percents in the respective field. When disk space is used above the set limit, syslog-ng Store Box (SSB) disconnects all clients. The default value is 90, and you can set values between 1-99.
3. *Optional step:* Enable the **Automatically start archiving** option to automatically start all configured archiving/cleanup jobs when disk usage goes over the limit.

NOTE: If there is no archiving policy set, enabling this option will not trigger automatic archiving.

4. Click **Commit**.
5. Navigate to **Basic Settings > Alerting > Health related traps** and enable alert **Disk usage is above the defined ratio**.

Figure 45: Basic Settings > Alerting > Health related traps – Enable alert

Health related traps			
Description	Name	Email	SNMP
Disk usage is above the defined ratio	xcbDiskFull	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>

6. Click **Commit**.

Configuring message rate alerting

With message rate alerting, you can detect the following abnormalities in syslog-ng Store Box(SSB):

- The syslog-ng inside SSB has stopped working.
- One of the clients/sites sending logs is not detectable.
- One of the clients/sites is sending too many logs, probably unnecessarily.

Message rate alerting can be set for sources, spaces and destinations (remote or local).

To configure message rate alerting

1. Navigate to **Log** and select **Sources**, **Spaces** or **Destinations**.
2. Enable **Message rate alerting**.
3. In case of **Sources**, select the counter to be measured:
 - Messages: Number of messages
 - Messages/sender: Number of messages per sender (the last hop)
 - Messages/hostname: Number of messages per host (based on the hostname in the message)

In case of **Spaces** or **Destinations**, the counter is the number of messages.

4. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.
5. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.
6. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.

Example: Creating an early time alert

In case you want an early time alert, can create a normal (non master) alert with a very low minimum number of messages and a low check interval.

Figure 46: Log > Sources > Message rate alerting — Create an early time alert

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages	24 hours	10000	1000000	Once	<input type="checkbox"/>
Messages	30 minutes	10	1000	Once	<input type="checkbox"/>

7. If you have set more than one message rate alerts, you can set a master alert where applicable. To set an alert to be a master alert, select the **Master alert** checkbox next to it.

When a master alert is triggered (and while it remains triggered), all other alerts for the given source/destination/space are suppressed. A master alert only blocks the other alerts that would be triggered at the given timeslot. A 24-hour alert does not block alerts that would be triggered at, for example 00:05.

Suggestions for setting the master alert:

- set the master alert to low check interval (5 minutes, if possible)
- set the master alert to a lower check interval than the alerts it suppresses
- set the master alert to have more lax limits than the alerts it suppresses

The following examples demonstrate a few common use cases of a **Master alert**.

Example: Using the master alert to indicate unexpected events

The user has 2 relays (sender) and 10 hosts per each relay (=20 hosts). Each host sends approximately 5-10 messages in 5 minutes. Two message rate alerts are set, and one master alert to signal extreme unexpected events. Such event can be that either a host is undetectable and probably has stopped working, or that it sends too many logs, probably due to an error. The following configuration helps detecting these errors without having to receive hundreds of alerts unnecessarily.

Figure 47: Log > Sources > Message rate alerting – Use a master alert to indicate unexpected events

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages/hostnar	5 minutes	50	100	Once	<input type="checkbox"/>
Messages/sender	5 minutes	500	1000	Once	<input type="checkbox"/>
Messages	5 minutes	0	10000	Once	<input checked="" type="checkbox"/>

8. *Optional step:* Global alerts count the number of all messages received by syslog-ng on all sources, including internal messages.
 - a. Navigate to **Log > Options > Message rate alerting statistics**. To add a global alert, click **+** at **Global alerts**.
 - b. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.
 - c. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.
 - d. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.
 - e. To set the alert as a system-wide master alert, select **Global master alert**. It will suppress all other log rate alerts on SSB when it is triggered.

NOTE: In the following cases, a so-called "always"-type super-master alert is triggered automatically.

If all or some of the statistics from syslog-ng cannot be fetched, an alert is sent out and all other errors are suppressed until the error is fixed.

If, for some reason, syslog-ng sends an unprocessable amount of statistics (for example because of some invalid input data), a similar super-master alert is triggered and stops processing the input.

9. *Optional step:* Navigate to **Log > Options > Message rate alerting statistics**. Set the maximum number of alerts you want to receive in **Limit of alerts sent out in a batch** to prevent alert flooding. SSB will send alerts up to the predefined value and then one single alert stating that too many message alerts were generated and the excess amount have not been sent.

CAUTION:

Hazard of data loss The alerts over the predefined limit will be unreachable.

System related traps

This section provides detailed information about system related SNMP traps. For information about configuring SNMP traps, see [Health related traps and SNMP trap settings](#).

To enable system related Email notifications and SNMP trap notifications, navigate to Basic Settings > Alerting > System related traps, select Email and/or SNMP for the traps you want to be notified about, then click Commit.

Figure 48: Basic Settings > Alerting > System related traps

System related traps			
Description	Name	Email	SNMP
Login failed	xcbLoginFailure	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Successful login	xcbLogin	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Logout from the management interface	xcbLogout	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Configuration changed	xcbConfigChange	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
General alert	xcbAlert	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
General error	xcbError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Data and configuration backup failed	xcbBackupFailed	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Data archiving failed	xcbArchiveFailed	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Database error occurred	xcbDBError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
License limit reached	xcbLimitReached	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
HA node state changed	xcbHaNodeChanged	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Timestamping error occurred	xcbTimestampError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Time sync lost	xcbTimeSyncLost	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Raid status changed	xcbRaidStatus	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Hardware error occurred	xcbHWEError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Firmware is tainted	xcbFirmwareTainted	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
		<input checked="" type="checkbox"/> <input type="checkbox"/>	<input checked="" type="checkbox"/> <input type="checkbox"/>

The following table provides detailed information about the system related SNMP traps .

Table 5: System related traps

Name	SNMP alert ID	Description
Login failed	xcbLoginFailure	Failed login attempts from syslog-ng Store Box (SSB) web interface.
Successful login	xcbLogin	Successful login attempts into SSB web interface.
Logout from the management interface	xcbLogout	Logouts from SSB web interface.
Configuration changed	xcbConfigChange	Any modification of SSB's configuration.
General alert	xcbAlert	<p>General alerts and error messages occurring on SSB.</p> <p>Note, that alerts on general alerts and errors are sent whenever there is an alert or error level message in the SSB system log. These messages are very verbose and mainly useful only for debugging purposes.</p> <p>Enabling these alerts may result in multiple e-mails or SNMP traps sent about the same event.</p>
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 SSB stores alerts and accounting information. Contact our support team (see About us on page 450 for contact information).
License limit reached	xcbLimitReached	Maximum number of clients has been reached.
HA node state changed	xcbHaNodeChanged	A node of the SSB cluster changed its state, for example, a takeover occurred.
Timestamping error occurred	xcbTimestampError	An error occurred during the time stamping process, for example the time stamping server did not respond.

Name	SNMP alert ID	Description
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	SSB detected a hardware error.
Firmware is tainted	xcbFirmwareTainted	A user has locally modified a file from the console.
Disk usage is above the defined ratio	xcbDiskFull	Disk space is used above the limit set in Disk space fill up prevention .

Alerts related to syslog-ng

This section provides detailed information about alerts related to syslog-ng.

To enable Email notifications and SNMP trap notifications related to syslog-ng, navigate to Basic Settings > Alerting > syslog-ng traps, select Email and/or SNMP for the traps you want to be notified about, then click [Commit](#).

Figure 49: Basic Settings > Alerting > syslog-ng traps

Description	Name	Email	SNMP
syslog-ng failure	syslogngFailureTrap	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Remote syslog-ng peer configuration changed	peerConfigChangeTrap	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Logspace exceeded warning size	spaceSizeLimit	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Message rate was outside the specified limits	ssbAbsoluteMessageRateAlert	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Too many message rate alerts were generated	ssbRateLimitTooManyAlerts	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Error during syslog-ng traffic statistics processing	ssbStatisticsError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Maximum number of connections has already been reached	syslogngConcurrentConnectionsReached	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
A destination path contains an invalid fragment	syslogngInvalidPathError	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>
Maximum number of dynamic clusters has been reached	syslogngDynamicClustersMaximumReached	Email <input checked="" type="checkbox"/>	SNMP <input checked="" type="checkbox"/>

The following table provides detailed information about the SNMP traps related to syslog-ng.

Table 6: Alerts related to syslog-ng

Name	SNMP alert ID	Description
syslog-ng failure	syslogngFailureTrap	The syslog-ng application did not start properly, shut

Name	SNMP alert ID	Description
Remote syslog-ng peer configuration changed	peerConfigChangeTrap	down unexpectedly, or encountered another problem. Depending on the error, syslog-ng Store Box (SSB) may not accept incoming messages or send them to the destinations.
Logspace exceeded warning size	spaceSizeLimit	The configuration of the syslog-ng application running on a remote host that sends its logs to SSB has been changed. Note that such changes are detected only if the remote peer uses at least version 3.0 of syslog-ng or version 3.0 of the syslog-ng Agent, and if messages from the internal source are sent to SSB.
Message rate was outside the specified limits	ssbAbsoluteMessageRateAlert	The size of a logspace has exceeded the size set as warning limit.
Too many message rate alerts were generated	ssbRateLimitTooManyAlerts	The message rate has exceeded the minimum or maximum value.
Error during syslog-ng traffic statistics processing	ssbStatisticsError	SSB is generating too many message rate alerts, probably due to unusual traffic that may need investigation and further user actions.
Maximum number of connections has already	syslogngConcurrentConnectionsReached	There was an error during querying and processing statistics of incoming, forwarded, stored, and dropped messages.
		There was an attempt to establish a new connection but this would have meant exceeding the log source's

Name	SNMP alert ID	Description
been reached		maximum number of allowed connections (set in Log > Sources > Maximum connections). The new connection was refused by syslog-ng.
A destination path contains an invalid fragment	syslogngInvalidPathError	syslog-ng was unable to open a specific logspace destination, because its path contains a prohibited fragment (such as a reference to a parent directory).
Maximum number of dynamic clusters has been reached	syslogngDynamicClustersMaximumReached	SSB collects various statistics about log messages received, processed, and dropped for objects (every source, destination, and individual application or program). To avoid performance issues, the maximal number of objects that SSB collects statistics for is 100000. This alert means that SSB has reached this limit.

Data and configuration backups

Backups create a snapshot of syslog-ng Store Box (SSB)'s configuration or the data which can be used for recovery in case of errors. SSB can create automatic backups of its configuration and the stored logs to a remote server.

To configure backups, you first have to create a backup policy. Backup policies define the address of the backup server, which protocol to use to access it, and other parameters. SSB 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 99.
- To configure backups using SMB/CIFS, see [Creating a backup policy using SMB/CIFS](#) on page 103.
- To configure backups using NFS, see [Creating a backup policy using NFS](#) on page 106.

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:

SSB 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 SSB cannot overwrite the files and the backup procedure fails.

Once you have configured a backup policy, set it as a system backup policy (for configuration backups) or data backup policy (for logspace backups):

- To configure a system backup policy, see [Creating configuration backups](#) on page 108.
- To configure a data backup policy, see [Creating data backups](#) on page 109.

NOTE: Backup deletes all other data from the target directory, restoring a backup deletes all other data from SSB. For details on restoring configuration and data from backup, see [Restoring SSB configuration and data](#) on page 438.

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. syslog-ng Store Box (SSB) 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/Cleanup** and click  in the **Backup policies** section to create a new backup policy.

Figure 50: Policies > Backup & Archive/Cleanup > Backup policies — Configure backup

The screenshot displays the configuration page for a backup policy. On the left, there is a navigation menu with sections: Basic Settings, AAA, Policies (selected), Hostlists, Shares, Backup & Archive/Cleanup (highlighted), Alert targets, Kerberos, HDFS Cluster, Log, Search, and Reports. Below this is a 'User menu' with options: Private keystore, Change password, Preferences, and Logout. At the bottom left is a 'System monitor' box showing system time (2018-10-04 13:46), remaining time (57:08), and lock status (admin@10.12.16.250). The main content area is titled 'Backup policies' and contains a form for the 'main-backup' policy. The 'Start time' is set to '23:30'. The 'Target settings' section has three radio buttons: 'Rsync over SSH', 'SMB/CIFS', and 'NFS' (selected). The 'Target server' field contains 'backup.example.com' and the 'Export' field contains 'backups'. The 'Send notification' section has three radio buttons: 'No notification e-mail', 'Send notification on errors only', and 'Send notification on all events' (selected). There is a checked checkbox for 'Include file list' and a text field for 'Maximum number of files in notification' set to '10240'. A 'Commit' button is located in the top right corner of the configuration area.

2. Enter a name for the backup policy (for example `main-backup`).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example `23:30`).
4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example `backup.example.com`).

5. Select **Rsync over SSH** from the **Target settings** radio buttons.

Figure 51: Policies > Backup & Archive/Cleanup > Backup policies – Configure backup using rsync

The screenshot shows the configuration page for a backup policy named 'main-backup'. The interface includes a left-hand navigation menu with sections like 'Basic Settings', 'Policies', and 'System monitor'. The main content area is titled 'Backup policies' and contains the following fields and options:

- Start time:** 23:30
- Target settings:**
 - Rsync over SSH
 - SMB/CIFS
 - NFS
- Username:** backup
- Target server:** 10.30.255.70
- Authentication key:** 2048 SHA256:9eyE7CYiF82T+5osJRzHu6Gs44NWtdibhgrz4yM7aFM
- Server host key:** 2048 SHA256:9eyE7CYiF82T+5osJRzHu6Gs44NWtdibhgrz4yM7aFM
- Path:** /backup/
- Send notification:**
 - No notification e-mail
 - Send notification on errors only
 - Send notification on all events
- Include file list:**
- Maximum number of files in notification:** 10240

6. Enter the username used to logon 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 SSB 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. SSB will compare the host key shown by the server to this key, and connect only if the two keys are identical.

Figure 52: Policies > Backup & Archive/Cleanup > Backup policies > Rsync over SSH > Server host key — Configure SSH keys

The screenshot shows a configuration window titled "Server host key". It is divided into three main sections:

- Query host:** A blue header with a close icon. Below it, a grey box contains the text "No host address or port is given".
- Upload key:** A blue header with a close icon. Below it, the text "Upload:" is followed by a "Browse..." button, a text input field containing "No file selected.", and an "Upload" button.
- Copy-paste key:** A blue header with a close icon. Below it, the text "Key:" is followed by a "Set" button and a large, empty text area for pasting a key.

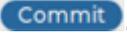
11. Enter the port number of the SSH server running on the remote machine into the **Port** field.
12. Enter the path to the backup directory on the target server into the **Path** field (for example /backups).

SSB saves all data into this directory, automatically creating subdirectories for logspaces. As a result of this, the same backup policy can be used for multiple logspaces. To ensure that a restore can be performed even if the logspace has been renamed, the subdirectories are created using a persistent internal ID of the logspace. To facilitate manual debugging, a text file is also saved in the directory with the name of the logspace, containing the internal ID for the logspace. This text file is only provided for troubleshooting purposes and is not used by SSB in any way.

13. 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 list is very long, the SSB 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring on SSB](#)).

14. Click .
15. To assign the backup policy to a logspace, see [Creating data backups](#).

Creating a backup policy using SMB/CIFS

The **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.

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/ssb1 (or similar) as a backup/archive share, it will fail.

CAUTION:

The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in syslog-ng Store Box (SSB), therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: /opt/ssb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

1. Navigate to **Policies > Backup & Archive/Cleanup** and click  in the **Backup policies** section to create a new backup policy.

Figure 53: Policies > Backup & Archive/Cleanup > Backup policies — Configure backup

The screenshot displays the configuration page for a backup policy named 'main-backup'. The left sidebar contains navigation menus for 'Basic Settings', 'AAA', 'Policies', 'Hostlists', 'Shares', 'Backup & Archive/Cleanup', 'Alert targets', 'Kerberos', 'HDFS Cluster', 'Log', 'Search', and 'Reports'. Below these are sections for 'User menu' (Private keystore, Change password, Preferences, Logout) and 'System monitor' (Time: 2018-10-04 13:46, Remaining time: 57:08, Locked: admin@10.12.16.250, Modules:). The main configuration area includes a 'Commit' button at the top right. The policy name 'main-backup' is entered in a text field. The 'Start time' is set to '23:30'. Under 'Target settings', 'NFS' is selected with a radio button. The 'Target server' is 'backup.example.com' and the 'Export' is 'backups'. Under 'Send notification', 'Send notification on all events' is selected, and 'Include file list' is checked. The 'Maximum number of files in notification' is set to '10240'.

2. Enter a name for the backup policy (for example `main-backup`).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example `23:30`).
4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example `backup.example.com`).
5. Select **Target settings > SMB/CIFS**.

NOTE: From SSB version 5.2.0, SSB only supports SMB 2.1 and later. Make sure that your operating system with the Samba share that you want to mount, supports SMB 2.1 or later. Otherwise, SSB cannot mount the remote share.

Figure 54: Policies > Backup & Archive/Cleanup > Backup policies — Configure backup via SMB/CIFS

6. Enter the username used to logon to the remote server into the **Username** field, and corresponding password into the **Password** field.

NOTE: NULL sessions (sessions without authentication) are not supported, authentication is required in all cases.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%\$%'()*+,-./:;<=>?@[]^_`{|}

7. Enter the name of the share into the **Share** field.
SSB saves all data into this directory, automatically creating the subdirectories. Backups of log files are stored in the data, configuration backups in the config subdirectory.
8. Enter the domain name of the target server into the **Domain** field.
9. 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 list is very long, the SSB 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring](#))

| on SSB)..

10. Click **Commit**.
11. To assign the backup policy to a logspace, see [Creating data backups](#).

Creating a backup policy using NFS

The **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/ssb1 (or similar) as a backup/archive share, it will fail.

1. Navigate to **Policies > Backup & Archive/Cleanup** and click **+** in the **Backup policies** section to create a new backup policy.

Figure 55: Policies > Backup & Archive/Cleanup > Backup policies — Configure backup

The screenshot shows the configuration page for a backup policy named 'main-backup'. The interface includes a left-hand navigation menu with sections for 'Basic Settings', 'Policies', 'Log', 'Search', 'Reports', 'User menu', and 'System monitor'. The main content area is titled 'Backup policies' and contains the following fields and options:

- Start time:** 23:30
- Target settings:**
 - Rsync over SSH
 - SMB/CIFS
 - NFS
- Target server:** backup.example.com
- Export:** backups
- Send notification:**
 - No notification e-mail
 - Send notification on errors only
 - Send notification on all events
- Include file list:**
- Maximum number of files in notification:** 10240

2. Enter a name for the backup policy (for example main-backup).
3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example 23:30).
4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example backup.example.com).

5. Select **NFS** from the **Target settings** radio buttons.

Figure 56: Policies > Backup & Archive/Cleanup > Backup policies – Configure NFS backups

The screenshot shows the 'Backup policies' configuration page for a policy named 'main-backup'. The interface includes a left-hand navigation menu with sections for 'Basic Settings', 'AAA', 'Policies', 'Hostlists', 'Shares', 'Backup & Archive/Cleanup', 'Alert targets', 'Kerberos', 'HDFS Cluster', 'Log', 'Search', and 'Reports'. Below this is a 'User menu' with options like 'Private keystore', 'Change password', 'Preferences', and 'Logout', and a 'System monitor' section showing system time and status. The main configuration area is titled 'Backup policies' and contains the following fields and options:

- Start time:** 23:30
- Target settings:**
 - Rsync over SSH
 - SMB/CIFS
 - NFS
- Target server:** backup.example.com
- Export:** backups
- Send notification:**
 - No notification e-mail
 - Send notification on errors only
 - Send notification on all events
- Include file list:**
- Maximum number of files in notification:** 10240

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.

syslog-ng Store Box (SSB) saves all data into this directory, automatically creating the subdirectories.

8. The remote server must also be configured to accept backups from SSB.

Add a line that corresponds to the settings of SSB 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 SSB backup policy.
- The IP address of the SSB interface that is used to access the remote server. For more information on the network interfaces of SSB, see [Network settings](#) on page 72.
- The following parameters: `(rw,no_root_squash, sync)`.

Example: Configuring NFS on the remote server

For example, if SSB connects the remote server from the 192.168.1.15 IP address and the data is saved into the `/var/backups/SSB` directory, add the following line to the `/etc/exports` file:

```
/var/backups/SSB 192.168.1.15(rw,no_root_squash, sync)
```

9. On the remote server, execute the following command:

```
exportfs -a
```

Verify that the `rpc portmapper` and `rpc.statd` applications are running.

10. 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 list is very long, the SSB 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring on SSB](#)).

11. Click **Commit**.
12. To assign the backup policy to a logspace, see [Creating data backups](#).

Creating configuration backups

To create a configuration backup, assign a backup policy as the **System backup policy** of syslog-ng Store Box (SSB).

TIP: To create an immediate backup of SSB'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.

To encrypt your configuration backups, see [Encrypting configuration backups with GPG](#) on page 110.

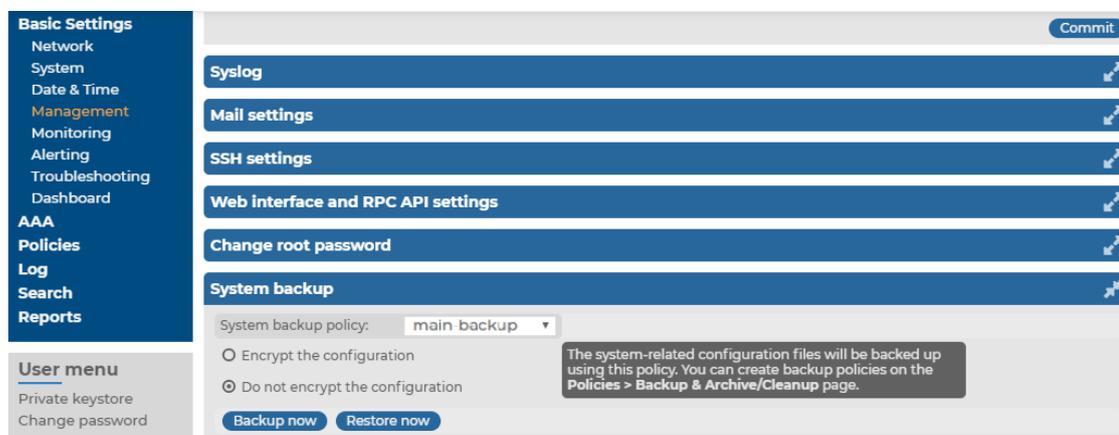
Prerequisites

You have to configure a backup policy before starting this procedure. For details, see [Data and configuration backups](#) on page 98.

To create a configuration backup

1. Navigate to **Basic Settings > Management > System backup**.

Figure 57: Basic Settings > Management > System backup — Configure system backup



2. Select the backup policy you want to use for backing up the configuration of SSB 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 logspace.

TIP: Data that is still in the memory of syslog-ng Store Box (SSB) is not copied to the remote server, only data that was already written to disk.

To make sure that all data is backed up (for example, before an upgrade), shut down syslog-ng before initiating the backup process.

CAUTION:

Statistics about syslog-ng and logspace sizes are not backed up. As a result, following a data restore, the Basic Settings > Dashboard page will not show any syslog-ng and logspace statistics about the period before the backup.

You have to configure a backup policy before starting this procedure. For details, see [Data and configuration backups](#) on page 98.

To configure data backups

1. Navigate to **Log > Logspaces**.
2. Select the logspace you want to back up.
3. Select a backup policy in the **Backup policy** field.
4. Click **Commit**.
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 syslog-ng Store Box (SSB) during system backups using the public-part of a GPG key. The system backups of SSB contain other information as well (for example, databases), but only the configuration file is encrypted. Note that system backups do not contain logspace data.

For details on restoring configuration from a configuration backup, see [Restoring SSB configuration and data](#) on page 438.

NOTE: It is not possible to directly import a GPG-encrypted configuration into SSB, 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 98.

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 configuration backups with GPG

1. Navigate to **Basic > System > Management > System backup**.
2. Select **Encrypt configuration**.
3. Select 
 - To upload a key file, click **Browse**, select the file containing the public GPG key, and click **Upload**. SSB accepts both binary and ASCII-armored GPG keys.
 - To copy-paste the key from the clipboard, paste it into the **Key** field and click **Set**.
4. Click **Commit**.

Archiving and cleanup

Archiving transfers data from syslog-ng Store Box (SSB) 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 SSB appliance.

To configure archiving and cleanup, you first have to create an archive/cleanup policy. Archive/cleanup policies define the retention time, the address of the remote backup server, which protocol to use to access it, and other parameters. SSB can be configured to use the SMB/CIFS and NFS protocols to access the backup server:

- To configure a cleanup policy that does not archive data to a remote server, see [Creating a cleanup policy](#) on page 112.
- To configure archiving using SMB/CIFS, see [Creating an archive policy using SMB/CIFS](#) on page 112.
- To configure archiving using NFS, see [Creating an archive policy using NFS](#) on page 114.

⚠ CAUTION:

Hazard of data loss Never delete an Archive Policy if data has been archived to 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 Logspace(s) 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:

SSB cannot modify the ownership of a file that already exists on the remote server.

Once you have configured an archive/cleanup policy, assign it to the logspace you want to archive. For details, see [Archiving or cleaning up the collected data](#) on page 116.

Creating a cleanup policy

Cleanup permanently deletes all log files and data that is older than **Retention time in days** without creating a backup copy or an archive. Such data is irrecoverably lost. Use this option with care.

NOTE: This policy does not delete existing archives from an external CIFS or NFS server.

1. Navigate to **Policies > Backup & Archive/Cleanup** and click  in the **Archive/Cleanup policies** section to create a new cleanup policy.
2. Enter a name for the cleanup policy.
3. Enter the time when the cleanup process should start into the **Start time** field in HH:MM format (for example 23:00).
4. Fill the **Retention time in days** field. Data older than this value is deleted from syslog-ng Store Box (SSB).
5. 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring on SSB](#) on page 87)

6. Click .
7. To assign the cleanup policy to the logspace you want to clean up, see [Archiving or cleaning up the collected data](#) on page 116.

Creating an archive policy using SMB/CIFS

The **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/ssb1 (or similar) as a backup/archive share, it will fail.

CAUTION:

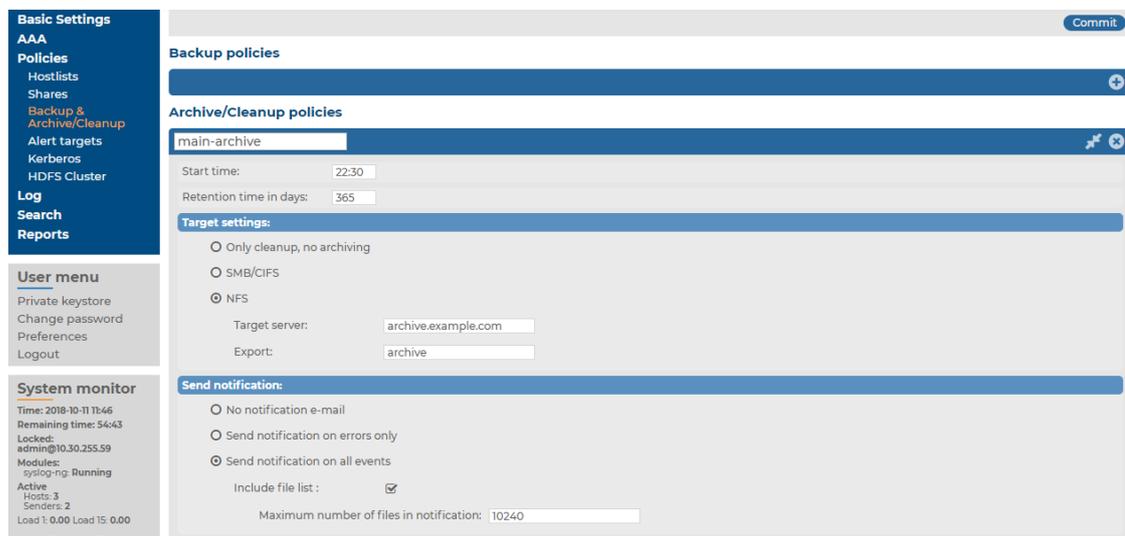
The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in syslog-ng Store Box (SSB), therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: /opt/ssb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

1. Navigate to **Policies > Backup & Archive/Cleanup** and click  in the **Archive/Cleanup policies** section to create a new archive policy.

Figure 58: Policies > Backup & Archive/Cleanup > Archive/Cleanup Policies — Configure cleanup and archiving



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).
4. Select **Target settings > SMB/CIFS**.
NOTE: From SSB version 5.2.0, SSB only supports SMB 2.1 and later. Make sure that your operating system with the Samba share that you want to mount, supports SMB 2.1 or later. Otherwise, SSB cannot mount the remote share.
5. Enter the username used to logon to the remote server into the **Username** field, and corresponding password into the **Password** field. For anonymous login, enter anonymous as username, and leave the **Password** field empty.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

6. Enter the name of the share into the **Share** field.

SSB saves all data into this directory, automatically creating the subdirectories. Archives of log files are stored in the `data`, configuration backups in the `config` subdirectory.

7. Enter the domain name of the target server into the **Domain** field.
8. Fill the **Retention time in days** field. Data older than this value is archived to the external server.

NOTE: The archived data is deleted from SSB.

9. 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring on SSB](#)).

10. Click **Commit**.
11. To assign the archive policy to the logspace you want to archive, see [Archiving or cleaning up the collected data](#).

Creating an archive policy using NFS

The **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/ssb1` (or similar) as a backup/archive share, it will fail.

1. Navigate to **Policies > Backup & Archive/Cleanup** and click  in the **Archive/Cleanup policies** section to create a new archive policy.

Figure 59: Policies > Backup & Archive/Cleanup > Archive/Cleanup Policies — Configure cleanup and archiving

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).
4. Select **NFS** from the **Target settings** radio buttons.
5. Enter the domain name of the remote server into the **Target server** field.
6. Enter the name of the NFS export into the **Export** field.

The syslog-ng Store Box (SSB) appliance saves all data into this directory, automatically creating the subdirectories.

7. The remote server must also be configured to accept connections from SSB.

Add a line that corresponds to the settings of SSB 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 SSB archive policy.
- The IP address of the SSB interface that is used to access the remote server. For more information on the network interfaces of SSB, see [Network settings](#) on page 72.
- The following parameters: `(rw,no_root_squash, sync)`.

Example: Configuring NFS on the remote server

For example, if SSB connects the remote server from the 192.168.1.15 IP address and the data is saved into the `/var/backups/SSB` directory, add the

following line to the `/etc/exports` file:

```
/var/backups/SSB 192.168.1.15(rw,no_root_squash, sync)
```

8. On the remote server, execute the following command:

```
exportfs -a
```

Verify that the `rpc portmapper` and `rpc.statd` applications are running.

9. Fill the **Retention time in days** field. Data older than this value is archived to the external server.
NOTE: The archived data is deleted from SSB.
10. 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 sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see [Configuring system monitoring on SSB](#)).
11. Click **Commit**.
12. To assign the archive policy to the logspace you want to archive, see [Archiving or cleaning up the collected data](#).

Archiving or cleaning up the collected data

To configure data archiving/cleanup, assign an archive/cleanup policy to the logspace.

Prerequisites:

You have to configure an archive/cleanup policy before starting this procedure. For details, see [Archiving and cleanup](#) on page 110.

To configure data archiving/cleanup

1. Navigate to **Log > Spaces**.
2. Select the logspace.
3. Select the archive/cleanup policy you want to use in the **Archive/Cleanup policy** field.
4. Click **Commit**.
5. *Optional:* To start the archiving or clean up process immediately, click **Archive now**. This functionality works only after a corresponding policy has been configured.

User management and access control

The **AAA** menu (Authentication, Authorization, and Accounting) allows you to control the authentication, authorization, and accounting settings of the users accessing syslog-ng Store Box (SSB). The following will be discussed in the next sections:

- For details on how to authenticate locally on SSB, see [Managing SSB users locally](#).
- For details on how to authenticate users using an external LDAP (for example Microsoft Active Directory) database, see [Managing SSB users from an LDAP database](#).
- For details on how to authenticate users using an external RADIUS server, see [Authenticating users to a RADIUS server](#).
- For details on how to control the privileges of users and usergroups, see [Managing user rights and usergroups](#).
- For details on how to display the history of changes of SSB configuration, see [Listing and searching configuration changes](#).

Managing SSB users locally

By default, syslog-ng Store Box (SSB) users are managed locally on SSB. In order to add local users in SSB, all steps of the following procedure need to be completed:

1. Create users.
For detailed instructions on how to create local users, see [Creating local users in SSB](#) on page 118.
2. Assign users to groups.
For details about how to add a usergroup, see [Managing local usergroups](#) on page 121.
3. Assign privileges to groups.
For information on how to control the privileges of usergroups, see [Managing user rights and usergroups](#) on page 127.

Creating local users in SSB

The following describes how to create a local user in syslog-ng Store Box (SSB).

NOTE: The admin user is available by default and has all possible privileges. It is not possible to delete this user.

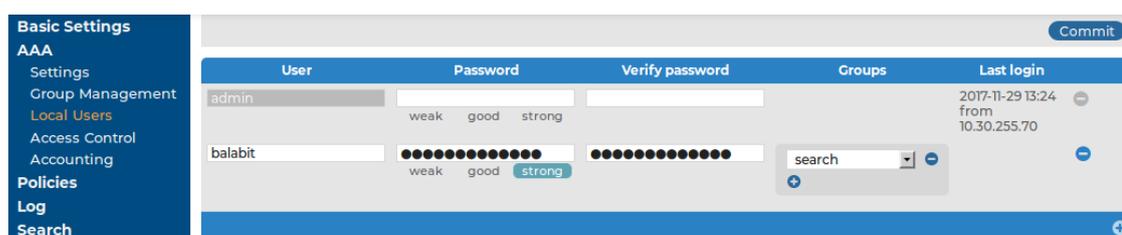
Local users cannot be managed when LDAP authentication is used (see [Managing SSB users from an LDAP database](#) on page 122). When LDAP authentication is enabled, the accounts of local users are disabled, they are not displayed on the **AAA > Local Users** page, but they are not deleted, either.

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 SSB. For details, see [Authenticating users to a RADIUS server](#) on page 125.

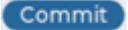
To create a local user in SSB

1. Navigate to **AAA > Local Users** and click .

Figure 60: AAA > Local Users — Create local user



User	Password	Verify password	Groups	Last login
admin	weak good strong			2017-11-29 13:24 from 10.30.255.70
balabit	●●●●●●●● weak good strong	●●●●●●●●	search +	

2. Enter the username into the **User** field.
NOTE: The following characters cannot be used in usernames: <>;\/[]:|=,+*?
3. Enter a password for the user into the **Password** and **Verify password** fields.
NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}~
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 119. The user can change the password later from the SSB web interface, and you can modify the password of the user here.
4. Click  in the **Groups** section and select a group that the user will be member of. Repeat this step to add the user to multiple groups.
If you wish to modify the group membership of a local user later on, you can do that here.
To remove a user from a group, click  next to the group.
5. Click .

Deleting a local user from SSB

This section describes how to delete a local user from syslog-ng Store Box (SSB).

To delete a local user from SSB

1. Navigate to **AAA > Local Users**.
2. Find the user you wish to delete.
3. Click next to the user, at the right edge of the screen.
4. Click **Commit**.

Setting password policies for local users

The syslog-ng Store Box (SSB) appliance can use password policies to enforce minimal password strength and password expiry. Password policies apply only to locally managed users, including the built-in admin and root users. They have no effect on users managed from an LDAP database, or if you authenticate your users to a RADIUS server.

The following describes how to create a password policy.

To create a password policy

1. Navigate to **AAA > Settings**.

Figure 61: AAA > Settings > User database — Configure password policies

The screenshot displays the configuration interface for AAA settings. On the left, there is a navigation menu with sections: Basic Settings (AAA, Settings, Group Management, Local Users, Access Control, Accounting), Policies, Log, Search, and Reports. Below this is a User menu (Private keystore, Change password, Preferences, Logout) and a System monitor section showing system time, remaining time, locked status, active modules, and system load. The main content area is titled 'Authentication settings' and includes: 'Authentication method' with radio buttons for 'Password provided by database' (selected) and 'RADIUS'; 'User database' with radio buttons for 'Local' (selected) and 'LDAP'; 'Password settings for root, admin and local users' with informational text and fields for 'Password expiration' (0 days), 'Number of passwords to remember' (10), and 'Minimal password strength' (good); and 'Cracklib (e.g. dictionary) check on password' with radio buttons for 'Enabled' and 'Disabled' (selected). At the bottom, there is an 'Accounting settings' section with a 'Require commit log' checkbox.

2. Verify that the **Authentication method** is set to **Password provided by database** and that the **User database** is set to **Local**.
NOTE: If the setting of these fields is different (for example LDAP or RADIUS), then SSB manages the passwords of the admin and root users locally.
3. Set how long the passwords are valid in the **Password expiration** field. After this period, SSB users will have to change their password. To disable password expiry, enter 0.
4. To prevent password-reuse (for example, when a user has two passwords and instead of changing to a new password only switches between the two), set how many different passwords the user must use before reusing an old password.
5. To enforce the use of strong passwords, select the level of password-complexity from the **Minimal password strength** field. As you type, SSB shows the strength of the password under the password field.

NOTE: The strength of the password is determined by its entropy: the variety of numbers, letters, capital letters, and special characters used, not only by its length. A strong password must have at least 12 characters, including lowercase and uppercase letters, numbers, and special characters.

The **Enable cracklib** option executes some simple dictionary-based attacks to find weak passwords.

6. Click **Commit**.

NOTE: If you increase the **Minimal password strength**, users whose existing password is weaker than required are forced to change their passwords immediately after their next login. The new passwords must comply with the strength requirements set in the password policy.

Managing local usergroups

You can use local groups to control the privileges of syslog-ng Store Box (SSB)'s local users (that is, who can view and configure what). Groups can be also used to control access to the logfiles available via a shared folder. For details, see [Accessing log files across the network](#) on page 239.

For the description of built-in groups, see [Built-in usergroups of SSB](#) on page 130.

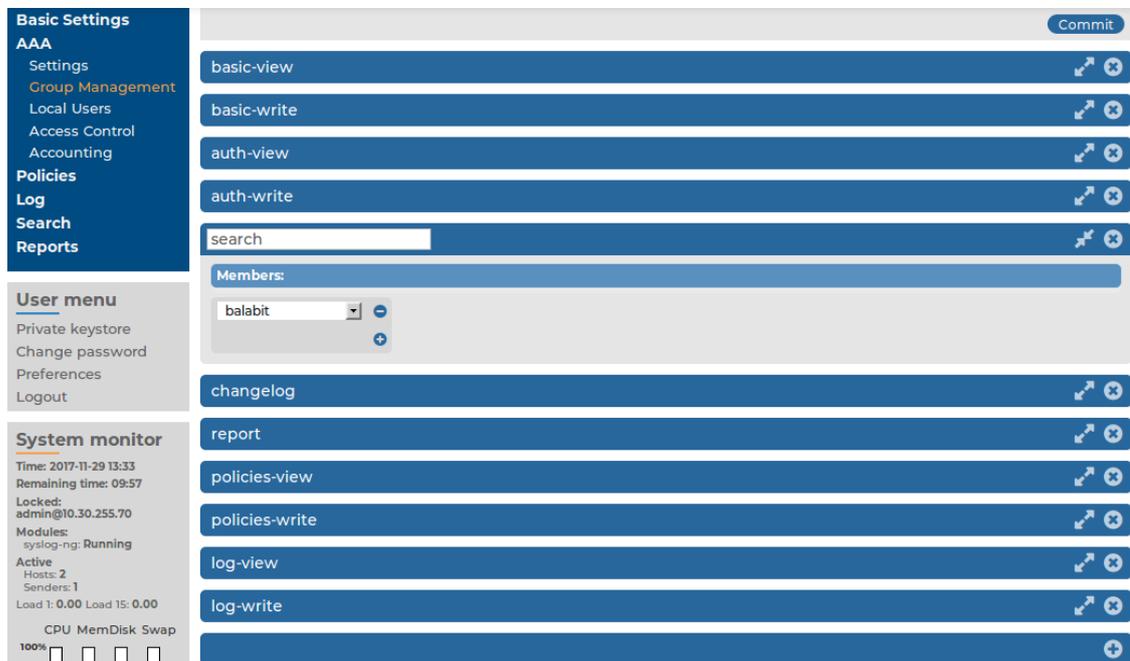
Use the **AAA > Group Management** page to:

- Create a new usergroup.
- Display which users belong to a particular local usergroup.
- Edit group memberships.

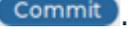
To create a new group

1. Navigate to **AAA > Group Management** and click .

Figure 62: AAA > Group Management — Manage local usergroups



2. Enter a name for the group.

3. Enter the names of the users belonging to the group. Click  to add more users.
4. Click .

Once you have added your usergroups, the next step is to start assigning privileges to them. For details on how to do that, see [Assigning privileges to usergroups for the SSB web interface](#) on page 128.

Managing SSB users from an LDAP database

The syslog-ng Store Box (SSB) web interface can authenticate users to an external LDAP database to simplify the integration of SSB to your existing infrastructure. You can also specify multiple LDAP servers, if the first server is unavailable, SSB will try to connect to the second server.

As in the case of locally managed users, use groups to control access to the logfiles available via a shared folder. For details, see [Accessing log files across the network](#) on page 239.

The following describes how to enable LDAP authentication.

NOTE: The admin user is available by default and has all privileges. It is not possible to delete this user.

The admin user can login to SSB even if LDAP authentication is used.

Enabling LDAP authentication automatically disables the access of every local user except for admin.

SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names). User Principal Names (UPNs) consist of a username, the at (@) character, and a domain name, for example administrator@example.com.

The following characters cannot be used in usernames and group names: <>\/[] ; | = , + *) ? @ "

When using RADIUS authentication together with LDAP users, the users are authenticated to the RADIUS server, only their group memberships must be managed in LDAP. For details, see [Authenticating users to a RADIUS server](#) on page 125.

 **CAUTION:**

A user can belong to a maximum of 10,000 groups, further groups are ignored.

 **CAUTION:**

By default, SSB uses nested groups when querying the LDAP server. Nested groups are mostly useful when authenticating the users to Microsoft Active Directory, but can slow down the query and cause the connection to time out if the LDAP tree is very large. In this case, disable the Enable nested groups option.

To enable LDAP authentication

1. Navigate to **AAA > Settings > Authentication settings**.
2. Select the **LDAP** option and enter the parameters of your LDAP server.

Figure 63: AAA > Settings > User database – Configure LDAP authentication

The screenshot shows the configuration interface for LDAP authentication. On the left, there is a navigation menu with sections: Basic Settings, AAA, User menu, and System monitor. The main content area is titled 'Authentication settings' and includes a 'Commit' button in the top right. The 'Authentication method' section has radio buttons for 'Password provided by database' (selected) and 'RADIUS'. The 'User database' section has radio buttons for 'Local' and 'LDAP' (selected). Below this, the 'Server address' is configured with a table for 'Address' and 'Port'. The first entry shows 'ldap.example.com' and '389'. There are plus and minus icons to add or remove servers. Below the table are fields for 'Base DN:', 'Bind DN:', and 'Bind password:'. The 'Type' section has radio buttons for 'Active Directory' and 'POSIX' (selected). It includes fields for 'Username (user ID) attribute name:' (uid), 'POSIX group membership attribute name:' (memberUid), and 'GroupOfUniqueNames membership attribute name:' (uniqueMember). The 'Encryption' section has radio buttons for 'Disabled', 'SSL/TLS', and 'STARTTLS' (selected). The 'Server certificate check' section has radio buttons for 'No certificate is required' and 'Only accept certificates authenticated by the specified CA certificate' (selected). There is a checkbox for 'Authenticate as client:' which is checked. Below it are fields for 'Client X.509 certificate:' (with a trash icon and the value '/C=HU/L=Budapest/O=Balabit/OU=TC/CN=Example User') and 'Client key:' (with a trash icon and the value '2048 SHA256:9ddpiWUdV8GfUKhg2xvjdthD6lp+8ai9sKDKwD/LBMI'). A 'Test' button is located at the bottom left of the configuration area.

- a. Enter the IP address or hostname and port of the LDAP server into the **Server Address** field. If you want to encrypt the communication between SSB and the LDAP server, in case of SSL/TLS, enter 636 as the port number, or in case of STARTTLS, enter 389 as the port number.

To add multiple servers, click  and enter the address of the next server. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.

CAUTION:

If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example ldap.example.com) in the Server Address field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the Common Name of the certificate.

- b. Enter the name of the DN to be used as the base of the queries into the **Base DN** field (for example DC=demodomain,DC=exampleinc).
- c. Enter the name of the DN where SSB should bind to before accessing the database into the **Bind DN** field.

For example: CN=Administrator,CN=Users,DC=demodomain,DC=exampleinc.

NOTE: SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names), for example administrator@example.com is also accepted.

NOTE: Do not use sAMAccountName, as the bind DN expects a CN.

- d. Enter the password to use when binding to the LDAP server into the **Bind Password** field.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used:

!"#\$%&'()*+,-./:;<=>?@[]^`{|}

- e. Select the type of your LDAP server in the **Type** field. Select **Active Directory** to connect to Microsoft Active Directory servers, or **Posix** to connect to servers that use the POSIX LDAP scheme.

3. If you want to encrypt the communication between SSB and the LDAP server, in **Encryption**, select the **SSL/TLS** or the **STARTTLS** option and complete the following steps:

NOTE:

TLS-encrypted connection to Microsoft Active Directory is supported only on Windows 2003 Server and newer platforms. Windows 2000 Server is not supported.

- If you want SSB to verify the certificate of the server, leave **Only accept certificates authenticated by the specified CA certificate** selected and click the  icon in the **CA X.509 certificate** field. A popup window is displayed.

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

SSB will use this CA certificate to verify the certificate of the server, and reject the connections if the verification fails.

CAUTION:

If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example `ldap.example.com`) in the **Server Address** field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the **Common Name** of the certificate.

- If the LDAP server requires mutual authentication, that is, it expects a certificate from SSB, enable **Authenticate as client**. Generate and sign a certificate for SSB, 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.

SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
 - Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.
4. (Optional) 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 (userid) attribute name** option. You can customize group-membership attributes using the **POSIX group membership attribute name** and **GroupOfUniqueNames membership attribute name** options.
 5. Click .
- NOTE:** You also have to configure the usergroups in SSB and possibly in your LDAP database. For details on using usergroups, see [How to use usergroups](#) on page 130.
6. Click **Test** to test the connection. Note that the testing of SSL-encrypted connections is currently not supported.

Authenticating users to a RADIUS server

The syslog-ng Store Box (SSB) appliance can authenticate its users to an external RADIUS server. Group memberships of the users must be managed either locally on SSB or in an LDAP database.

CAUTION:

The challenge/response authentication methods is currently not supported. Other authentication methods (for example, password, SecureID) should work.

To authenticate SSB users to a RADIUS server

1. Navigate to **AAA > Settings**.

Figure 64: AAA > Settings — Configuring RADIUS authentication

The screenshot shows the 'Authentication settings' page. On the left is a navigation menu with 'Basic Settings', 'AAA', 'Settings', 'Group Management', 'Local Users', 'Access Control', 'Accounting', 'Policies', 'Log', 'Search', 'Reports', and a 'User menu' with 'Private keystore', 'Change password', 'Preferences', and 'Logout'. The main content area has a 'Commit' button at the top right. The 'Authentication method' section has two radio buttons: 'Password provided by database' (unselected) and 'RADIUS' (selected). Below this is a 'Server address' section with a table:

Address	Port	Shared secret
radius.example.com	1812	●●●●●●
radius2.example.com	1812	●●●●●●

Below the table is a 'User database' section with two radio buttons: 'Local' (selected) and 'LDAP' (unselected).

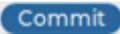
2. Set the **Authentication method** field to **RADIUS**.
3. Enter the IP address or domain name of the RADIUS server into the **Address** field.
4. Enter the password that SSB can use to access the server into the **Shared secret** field.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

5. To add more RADIUS servers, click  and repeat Steps 2-4.
Repeat this step to add multiple servers. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.
6. When configuring RADIUS authentication with a local user database, complete the following steps.
 - a. Set **Password expiration** to 0.
 - b. Set **Number of passwords to remember** to 0.
 - c. Set **Minimal password strength** to disabled.
 - d. Set **Cracklib check on password** to disabled.

7.

 **CAUTION:**

After clicking , the SSB web interface will be available only after successfully authenticating to the RADIUS server. Note that the default admin account of SSB will be able to login normally, even if the RADIUS server is inaccessible.

Click .

Managing user rights and usergroups

In syslog-ng Store Box (SSB), user rights can be assigned to usergroups. SSB has numerous usergroups defined by default, but custom user groups can be defined as well. Every group has a set of privileges: which pages of the SSB web interface it can access, and whether it can only view (read) or also modify (read & write/perform) those pages or perform certain actions.

Figure 65: AAA > Access Control – Managing SSB users

Group	Object	Type
basic-view	Basic Settings	read
basic-write	Basic Settings	read and write/perform
auth-view	AAA	read
auth-write	AAA	read and write/perform
search	Search	read
changelog	AAA/Accounting	read
report	Use static subchapters Reports	read and write/perform
policies-view	Policies	read
policies-write	Policies	read and write/perform
log-view	Log	read
log-write	Log	read and write/perform

NOTE: Every group has either read or read & write/perform privileges to a set of pages.

- For details on assigning privileges to a usergroup, see [Assigning privileges to usergroups for the SSB web interface](#) on page 128.
- For details on modifying existing groups, see [Modifying group privileges](#) on page 128.
- For details on finding usergroups that have a specific privilege, see [Finding specific usergroups](#) on page 129.
- For tips on using usergroups, see [How to use usergroups](#) on page 130.
- For a detailed description about the privileges of the built-in usergroups, see [Built-in usergroups of SSB](#) on page 130.

The admin user is available by default and has all privileges, except that it cannot remotely access the shared logspaces. It is not possible to delete this user.

Assigning privileges to usergroups for the SSB web interface

This section describes how to assign privileges to a new group.

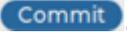
To assign privileges to a new group

1. Navigate to **AAA > Access Control** and click .
2. Find your usergroup. If you start typing the name of the group you are looking for, the autocomplete function will make finding your group easier for you.
3. Click  located next to the name of the group. The list of available privileges is displayed.
4. Select the privileges (pages of the syslog-ng Store Box (SSB) interface) to which the group will have access and click **Save**.

NOTE: To export the configuration of SSB, the **Export configuration** privilege is required.

To import a configuration to SSB, the **Import configuration** privilege is required.

To update the firmware and set the active firmware, the **Basic settings > System** privilege is required.

5. Select the type of access (read or read & write) from the **Type** field.
6. Click .

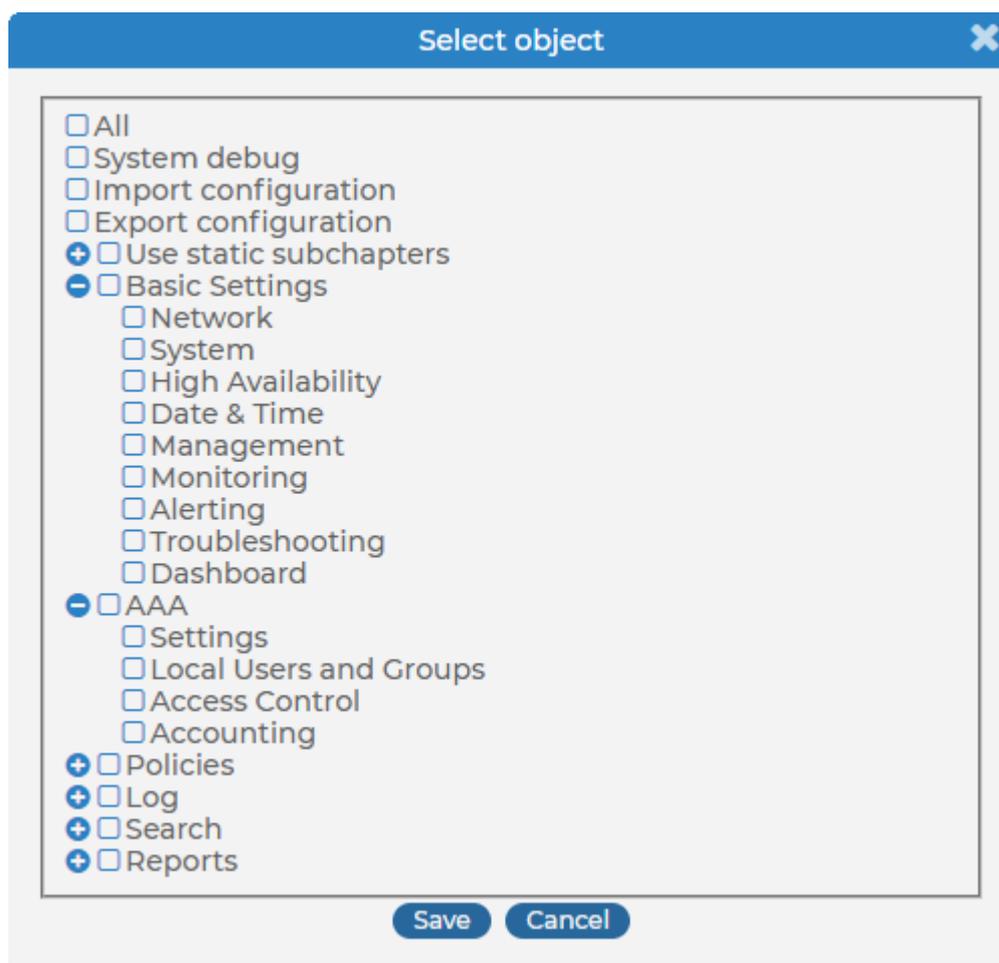
Modifying group privileges

This section describes how to modify the privileges of an existing group.

To modify the privileges of an existing group

1. Navigate to **AAA > Access Control**.
2. Find the group you want to modify and click . The list of available privileges is displayed.
3. Select the privileges (pages of the syslog-ng Store Box (SSB) interface) to which the group will have access and click **Save**.

Figure 66: AAA > Access Control — Modifying group privileges



CAUTION:

Assigning the Search privilege to a user on the AAA page grants the user search access to every logspace, even if the user is not a member of the groups listed in the Access Control option of the particular logspace.

4. Select the type of access (read or read & write) from the **Type** field.
5. Click **Commit**.

Finding specific usergroups

The **Filter ACLs** section of the **AAA > Access Control** page provides you with a simple searching and filtering interface to search the names and privileges of usergroups.

Figure 67: AAA > Access Control – Finding specific usergroups

The screenshot shows the 'Filter ACLs' section of the AAA > Access Control interface. It features a sidebar on the left with navigation options: Basic Settings, AAA, Settings, Group Management, Local Users, Access Control, and Accounting. The main content area has a 'Filter ACLs' header with a 'Commit' button in the top right. Below the header is a table with columns for 'Group', 'Object', and 'Type'. The 'Type' column has a dropdown menu currently set to 'any'. There are 'Search' and 'Clear filters' buttons below the table.

- To select usergroups starting with a specific string, enter the beginning of the name of the group into the **Group** field and select **Search**.
- To select usergroups who have a specific privilege, click , select the privilege or privileges you are looking for, and click **Search**.
- To filter for read or write access, use the **Type** option.

How to use usergroups

How you should name usergroups depends on the way you manage your syslog-ng Store Box (SSB) users.

- Local users: If you use only local users, create or modify usergroups on the **AAA > Group Management** page, assign or modify privileges on the **AAA > Access Control** page, and add users to the groups on the **AAA > Local Users** or the **AAA > Group Management** 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 SSB users, create or modify your usergroups on the **AAA > Access Control** page and ensure that the name of your LDAP group and the SSB usergroup is the same. For example, to make members of the admins LDAP group be able to use SSB, create a usergroup called admins on the **AAA > Access Control** page and edit the privileges of the group as needed.

CAUTION:

A user can belong to a maximum of 10,000 groups, further groups are ignored.

- 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 **AAA > Access Control** page, and add your RADIUS users to these groups on the **AAA > Group Management** page.

Built-in usergroups of SSB

The syslog-ng Store Box (SSB) appliance has the following usergroups by default. Note that you can modify and delete these usergroups as you see fit.

CAUTION:

If you use LDAP authentication on the SSB web interface and want to use the default usergroups, you have to create these groups in your LDAP database and assign users to them. For details on using usergroups, see [How to use usergroups on page 130](#).

- **basic-view:** View the settings in the **Basic Settings** menu, including the system logs of SSB. Members of this group can also execute commands on the **Troubleshooting** tab.
- **basic-write:** Edit the settings in the **Basic Settings** menu. Members of this group can manage SSB as a host.
- **auth-view:** View the names and privileges of the SSB administrators, the configured usergroups, and the authentication settings in the **AAA** menu. Members of this group can also view the history of configuration changes.
- **auth-write:** Edit authentication settings and manage users and usergroups.

CAUTION:

Members of the **auth-write** group, or any other group with write privileges to the **AAA** menu are essentially equivalent to system administrators of SSB, because they can give themselves any privilege. Users with limited rights should never have such privileges.

If a user with write privileges to the **AAA** menu gives himself new privileges (for example gives himself group membership to a new group), then he has to relogin to the SSB web interface to activate the new privilege.

- **search:** Browse and download various logs and alerts in the **Search** menu.
NOTE: The `admin` user is not a member of this group by default, so it cannot remotely access the shared logspaces.
- **changelog:** View the history of SSB configuration changes in the **AAA > Accounting** menu.
- **report:** Browse, create and manage reports, and add statistics-based chapters to the reports in the **Reports** menu.
NOTE: To control exactly which statistics-based chapters and reports can the user include in a report, use the `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.

CAUTION:

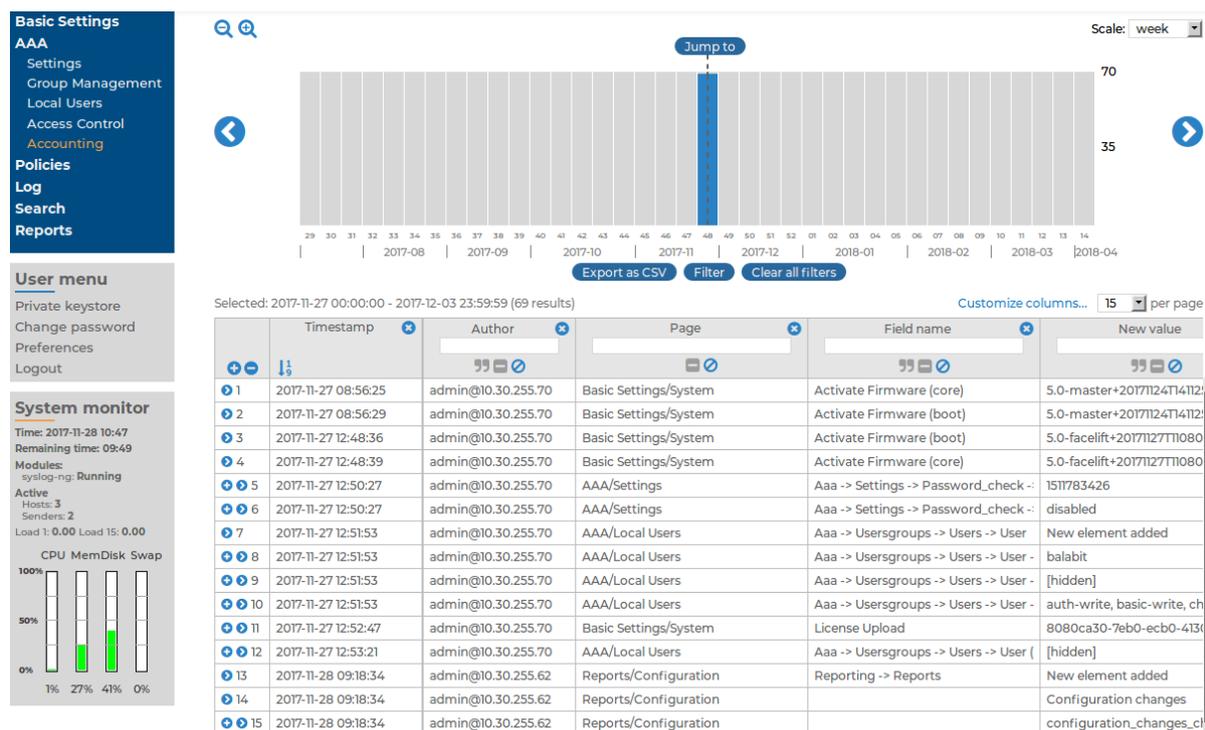
Members of this group can make the logs stored on SSB available as a shared network drive. In case of unencrypted logfiles, this may result in access to sensitive data.

- **log-view:** View the logging settings in the **Log** menu.
- **log-write:** Configure logging settings in the **Log** menu.

Listing and searching configuration changes

The syslog-ng Store Box (SSB) appliance automatically tracks every change of its configuration. To display the history of changes, select **AAA > Accounting**. The changes are organized as log messages, and can be browsed and searched using the regular SSB search interface (for details, see [Searching log messages](#) on page 330). The following information is displayed about each modification:

Figure 68: AAA > Accounting — Browsing configuration changes



- **Timestamp:** The date of the modification.
- **Author:** Username of the administrator who modified the configuration of SSB.
- **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:** Indicates 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 **AAA > Settings > Accounting settings** and select the **Require commit log** option.

Managing SSB

The following sections explain the basic management tasks of syslog-ng Store Box (SSB).

- For basic management tasks (reboot and shutdown, disabling traffic), see [Controlling SSB: restart, shutdown](#).
- For managing a high availability cluster, see [Managing a high availability SSB cluster](#).
- For instructions on upgrading SSB, see [Upgrading SSB](#).
- For instructions on accessing SSB through console and SSH, see [Accessing the SSB console](#).
- For enabling sealed mode (which disables basic configuration changes from a remote host), see [Sealed mode](#).
- For information on configuring the out-of-band (IPMI) interface, see [Out-of-band management of SSB](#).
- For managing certificates used on SSB, see [Managing the certificates used on SSB](#).
- For creating hostlist policies, see [Creating hostlist policies](#).

Controlling SSB: restart, shutdown

To restart or shut down syslog-ng Store Box (SSB), navigate to **Basic Settings > System > System control** and click the respective action button. The **Other node** refers to the slave node of a high availability SSB cluster. For details on high availability clusters, see [Managing a high availability SSB cluster](#) on page 135.

CAUTION:

- **When rebooting the nodes of a cluster, reboot the other (slave) node first to avoid unnecessary takeovers.**
- **When shutting down the nodes of a cluster, shut down the other (slave) node first. When powering on the nodes, start the master 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 SSB.**

Figure 69: Basic Settings > System > System control — Performing basic management



NOTE: Web sessions to the SSB interface are persistent and remain open after rebooting SSB, so you do not have to relogin after a reboot.

Managing a high availability SSB cluster

High availability (HA) clusters can stretch across long distances, such as nodes across buildings, cities or even continents. The goal of HA clusters is to support enterprise business continuity by providing location-independent failover and recovery.

To set up a high availability cluster, connect two syslog-ng Store Box (SSB) units with identical configurations in high availability mode. This creates a primary-secondary node pair.

NOTE: Primary and secondary nodes have the following functions, and are sometimes referred to, and displayed, as follows:

- **Primary node:** Functions as the active node in the HA cluster. Sometimes the primary node is also referred to as the master node. On the **Basic Settings > High Availability > High availability & Nodes** page, the primary node is displayed as **This node**.
- **Secondary node:** Functions as the backup node in the HA cluster. Sometimes the secondary node is also referred to as the slave node. On the **Basic Settings > High Availability > High availability & Nodes** page, the primary node is displayed as **Other node**

Should the primary node stop functioning, the secondary node takes over the functionality of the primary node. This way, the SSB servers are continuously accessible.

NOTE: To use the management interface and high availability mode together, connect the management interface of both SSB nodes to the network, otherwise you will not be able to access SSB remotely when a takeover occurs.

The primary node shares all data with the secondary node using the HA network interface (labeled as 4 or HA on the SSB 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 an SSB cluster in [Troubleshooting an SSB cluster](#) on page 429.

NOTE: HA functionality was designed for physical SSB units. If SSB is used in a virtual environment, use the fallback functionalities provided by the virtualization service instead.

On virtual SSB appliances, or if you have bought a physical SSB appliance without the high availability license option, the **Basic Settings > High Availability** menu item is not displayed anymore.

The **Basic Settings > High Availability** page provides information about the status of the HA cluster and its nodes.

Figure 70: Basic Settings > High Availability — Managing a high availability cluster

The screenshot shows the Syslog-ng web interface for managing a high availability cluster. The top navigation bar includes 'syslog-ng STORE BOX', user 'admin', IP '10.12.15.122', and session info. The left sidebar contains menu items like 'Basic Settings', 'Network', 'System', 'High Availability', 'Date & Time', 'Management', 'Monitoring', 'Alerting', 'Troubleshooting', 'Dashboard', 'AAA', 'Policies', 'Log', 'Search', and 'Reports'. Below the sidebar is a 'User menu' with options like 'Private keystore', 'Change password', 'Preferences', and 'Logout'. The 'System monitor' section shows system time, remaining time, and resource usage (CPU, Mem, Disk, Swap).

The main content area is titled 'High availability & Nodes' and shows the following information:

- Cluster status:** SSB is currently operating in HA.
- Redundant Heartbeat status:** OK (All redundant HA links are functioning properly)
- Current master:** 49434D53-0200-90C6-2500-C690250095E9
- Buttons:** Activate Slave, Reboot cluster
- HA UUID:** fd52d9df-3861-48c9-8eb1-486985437d25
- DRBD Status:** Connected - Connected, Connected
- DRBD sync rate limit:** 1 Gbit/s
- DRBD asynchronous mode:**

The configuration is split into two columns: 'This node' and 'Other node'.

	This node	Other node
Node ID:	49434D53-0200-90C6-2500-C690250095E9	49434D53-0200-90E9-2500-E990250003BA
Node HA status:	HA	HA
Node HA UUID:	fd52d9df-3861-48c9-8eb1-486985437d25	fd52d9df-3861-48c9-8eb1-486985437d25
DRBD status:	Connected (UpToDate) Connected	Connected (UpToDate) Connected
RAID status:	Boot partition: active Data partition: active	Boot partition: active Data partition: active
Boot firmware versions:	Current: 6.6.0 Active: 6.6.0	Current: 6.6.0 Active: 6.6.0
IPMI IP address:	10.12.193.67	10.12.193.66
IPMI subnet mask:	255.255.255.0	255.255.255.0
IPMI default gateway:	10.12.193.254	10.12.193.254
IPMI IP address source:	DHCP Address	DHCP Address
DRBD and heartbeat	Interface IP: 192.168.0.2	Interface IP: 192.168.0.1
HA interface	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Link speed:	Auto negotiation	Auto negotiation
Redundant heartbeat	Interface IP: 10.12.165.24	Interface IP: 10.12.165.23
External interface	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
HA MAC:	aa:7f:d8:71:f4:06	5a:0d:1b:fa:3c:4b
Management interface	<input type="checkbox"/>	<input type="checkbox"/>
Internal interface	<input type="checkbox"/>	<input type="checkbox"/>
Next hop monitoring	Interface IP: 10.12.165.21	Interface IP: 10.12.165.20
External interface	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Management interface	<input type="checkbox"/>	<input type="checkbox"/>
Internal interface	<input type="checkbox"/>	<input type="checkbox"/>
	Reboot Shutdown	Reboot Shutdown

Information about the HA cluster

The following information is available about the cluster:

- **Cluster status:** Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

You can find the description of each HA status in [Understanding SSB cluster statuses](#) on page 429.

- **Redundant Heartbeat status:** Indicates whether the redundant HA links in your HA cluster are functioning properly.
- **Current master:** The MAC address of the high availability interface (4 or HA) of the node.
- **HA UUID:** A unique identifier of the HA cluster. Only available in High Availability mode.

- **DRBD status:** Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

You can find the description of each DRBD status in [Understanding SSB cluster statuses](#) on page 429.

- **DRBD sync rate limit:** The maximum allowed synchronization speed between the master and the slave node.

You can find more information about configuring the DRBD sync rate limit in [Adjusting the synchronization speed](#) on page 140.

- **DRBD asynchronous mode:** Enable to compensate for high network latency and bursts of high activity by enabling asynchronous data replication between the primary and the secondary node.

You can find more information about configuring asynchronous data replication in [Asynchronous data replication](#) on page 140.

Information about the HA nodes

The active (primary) SSB node is labeled as **This node**. This unit receives the incoming log messages and provides the web interface. The SSB 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 universally unique identifier (UUID) of the physical or virtual machine.

NOTE: Due to backward compatibility, in the case of upgrades, the Node ID is the MAC address of the node's HA interface.

For SSB clusters, the IDs of both nodes are included in the internal log messages of SSB.

- **Node HA status:** Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

You can find the description of each HA status in [Understanding SSB cluster statuses](#) on page 429.

- **Node HA UUID:** A unique identifier of the cluster. It is a software-generated identifier. Only available in High Availability mode.
- **DRBD status:** The status of data synchronization between the nodes.
You can find the description of each DRBD status in [Understanding SSB cluster statuses](#) on page 429.
- **RAID status:** The status of the RAID device of the node.
- **Boot firmware versions:** The version numbers of the boot firmware.
You can find more information about the boot firmware in [Firmware in SSB](#) on page 24.
- **IPMI IP address:** The node's IPMI IP address.
- **IPMI subnet mask:** The node's IPMI subnet mask.
- **IPMI default gateway:** The node's IPMI default gateway.
- **IPMI IP address source:** The node's IPMI IP address source.
- **DRBD and heartbeat:** You can enable and configure the following options for DRBD and heartbeat monitoring:
 - **HA Interface:** Enable to customize the IP address of the high availability (HA) interface.
NOTE: You can only modify the **IP address** in the **Interface IP** field when your HA cluster is in **STANDALONE** or in **HA** state.
 - **Link speed:** The maximum allowed speed between the primary 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.
Leave this field on **Auto negotiation** unless specifically requested by the support team.
- **Redundant heartbeat:**
You can enable and configure the following:
 - **External interface:** Enable to set the **Interface IP address** that you want to use for redundant heartbeat monitoring on the node's external interface.
 - **Management interface:** Enable to set the **Interface IP address** that you want to use for redundant heartbeat monitoring on the node's management interface.
 - **Internal interface:** Enable to set the **Interface IP address** that you want to use for redundant heartbeat monitoring on the node's internal interface..
 You can find more information about configuring redundant heartbeat interfaces in [Redundant heartbeat interfaces](#) on page 141.
- **Next hop monitoring:** The interfaces you enable and configure here are virtual interfaces, used only to detect that the other node is still available. The interfaces you enable here are not used to synchronize data between the nodes, and only heartbeat messages are transferred.
You can enable and configure the following:

- **External interface:** Enable to set the **Interface IP address** that you want to use for next hop monitoring on the node's external interface.
- **Management interface:** Enable to set the **Interface IP address** that you want to use for next hop monitoring on the node's management interface.
- **Internal interface:** Enable to set the **Interface IP address** that you want to use for next hop monitoring on the node's internal interface.

You can find more information about configuring next hop monitoring in [Next-hop router monitoring](#).

Configuration and management options for HA clusters

The following configuration and management options are available for HA clusters:

- Set up a high availability cluster: You can find detailed instructions for setting up a HA cluster in "[Installing two SSB 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.

You can find more information about configuring the DRBD synchronization speed in [Adjusting the synchronization speed](#) on page 140.

- Enable asynchronous data replication: You can compensate for high network latency and bursts of high activity by enabling asynchronous data replication between the primary and the secondary node with the **DRBD asynchronous mode** option.

You can find more information about configuring asynchronous data replication in [Asynchronous data replication](#) on page 140.

- Configure redundant heartbeat interfaces: You can configure virtual interfaces for each HA node to monitor the availability of the other node.

You can find more information about configuring redundant heartbeat interfaces in [Redundant heartbeat interfaces](#) on page 141.

- Configure next-hop monitoring: You can provide IP addresses (usually next hop routers) to continuously monitor from both the primary and the secondary nodes 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.

You can find more information about configuring next-hop monitoring in [Next-hop router monitoring](#) on page 143.

- Reboot the HA cluster: 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: Reboots the selected node.

When rebooting the nodes of a cluster, reboot the other (secondary) node first to avoid unnecessary takeovers.

- Shutdown a node: Forces the selected node to shutdown.

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.

- Manual takeover: To activate the other node and disable the currently active node, click **Activate slave**.

Activating the secondary node terminates all connections of SSB and might result in data loss. The secondary node becomes active after about 60 seconds, during which SSB cannot accept incoming messages. Enable disk-buffering on your syslog-ng clients and relays to prevent data loss in such cases.

Adjusting the synchronization speed

When operating two syslog-ng Store Box (SSB) units in High Availability mode, every incoming data copied from the master (active) node to the slave (passive) node. Since synchronizing data 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 — for example, when synchronizing the disks after converting a single node to a high availability cluster.

The **Basic Settings > High Availability > DRBD status** field indicates whether the latest data (including SSB configuration, log files, and so on) is available on both SSB nodes. For a description of each possible status, see [Understanding SSB cluster statuses](#) on page 429.

To change the limit of the DRBD synchronization rate, navigate to **Basic Settings > High Availability**, select **DRBD sync rate limit**, and select the desired value.

Set the sync rate carefully. A high value is not recommended if the load of SSB is very high, as increasing the resources used by the synchronization process may degrade the general performance of SSB. On the other hand, the HA link's speed must exceed the speed of the incoming logs, else the web UI might become unresponsive and data loss can occur.

If you experience bursts of high activity, consider turning on asynchronous data replication.

Asynchronous data replication

When a high availability syslog-ng Store Box (SSB) cluster is operating in a high-latency environment or during brief periods of high load, there is a risk of slowness, latency or

package loss. To manage this, you can compensate latency with asynchronous data replication.

Asynchronous data replication is a method where local write operations on the primary node are considered complete when the local disk write is finished and the replication packet is placed in the local TCP send buffer. It does not impact application performance, and tolerates network latency, allowing the use of physically distant storage nodes. However, because data is replicated at some point after local acknowledgement, the remote storage nodes are slightly out of step: if the local node at the primary data center breaks down, data loss occurs.

To turn asynchronous data replication on, navigate to **Basic Settings > High Availability**, and enable **DRBD asynchronous mode**. You have to reboot the cluster (click **Reboot cluster**) for the change to take effect.

Under prolonged heavy load, asynchronous data replication might not be able to compensate for latency or for high packet loss ratio (over 1%). In this situation, stopping the slave machine is recommended to avoid data loss at the temporary expense of redundancy.

Redundant heartbeat interfaces

To avoid unnecessary takeovers and to minimize the chance of split-brain situations, you can configure additional heartbeat interfaces in syslog-ng Store Box (SSB). 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 SSB cluster statuses](#) on page 429.

The redundant heartbeat interface is a virtual interface with a virtual MAC address that uses an existing interface of SSB (for example, the external or the management interface). 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 432.

NOTE: Even if redundant HA links are configured, if the dedicated HA link fails, the secondary node will not be visible on the **Basic Settings > High Availability > High availability & Nodes** page anymore.

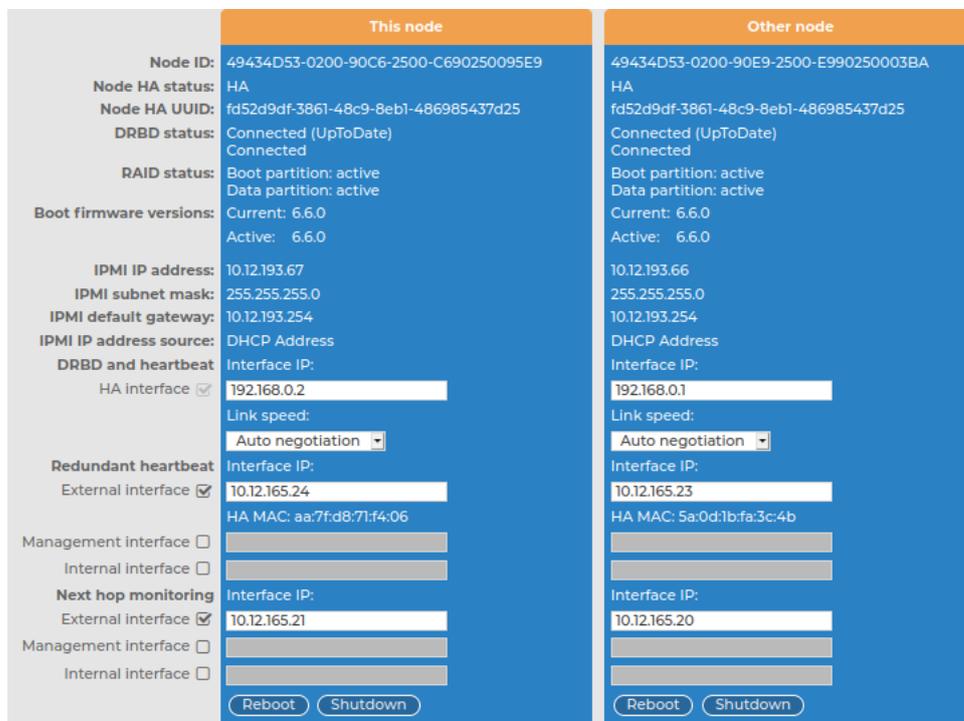
SSB nodes use UDP port **694** to send each other heartbeat signals.

The following describes how to configure a redundant heartbeat interface.

To configure a redundant heartbeat interface

1. Navigate to **Basic Settings > High Availability > Redundant heartbeat**.
2. Enable the interface you want to use as redundant heartbeat interface (for example, **External interface**). Using an interface as a redundant heartbeat interface does not affect the original traffic of the interface.

Figure 71: Basic Settings > High Availability > Redundant heartbeat — Configuring redundant heartbeat interfaces



3. Enter an IP address into the **This node > Redundant heartbeat > Interface IP** field of the selected interface.

NOTE: Consider 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](#).
4. Enter an IP address into the **Other node > Redundant heartbeat > Interface IP** field of the selected interface.

NOTE: Consider 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](#).
 5. Repeat the previous steps to add additional redundant heartbeat interfaces if needed.
 6. Click **Commit**.
 7. 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 (**This node**) stops working or becomes unreachable from the secondary node (**Other node**). However, this does not cover the scenario when the primary node becomes inaccessible to the outside world (for example its external interface or the router or switch connected to the external interface breaks down) 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 from both the primary and the secondary nodes 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 syslog-ng Store Box (SSB) 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 SSB 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.

The following describes how to configure next hop monitoring.

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 72: Basic Settings > High Availability > Next hop monitoring – Configuring next hop monitoring

	This node	Other node
Node ID:	49434D53-0200-90C6-2500-C690250095E9	49434D53-0200-90E9-2500-E990250003BA
Node HA status:	HA	HA
Node HA UUID:	fd52d9df-3861-48c9-8eb1-486985437d25	fd52d9df-3861-48c9-8eb1-486985437d25
DRBD status:	Connected (UpToDate) Connected	Connected (UpToDate) Connected
RAID status:	Boot partition: active Data partition: active	Boot partition: active Data partition: active
Boot firmware versions:	Current: 6.6.0 Active: 6.6.0	Current: 6.6.0 Active: 6.6.0
IPMI IP address:	10.12.193.67	10.12.193.66
IPMI subnet mask:	255.255.255.0	255.255.255.0
IPMI default gateway:	10.12.193.254	10.12.193.254
IPMI IP address source:	DHCP Address	DHCP Address
DRBD and heartbeat	Interface IP: 192.168.0.2	Interface IP: 192.168.0.1
HA interface <input checked="" type="checkbox"/>	Link speed: Auto negotiation	Link speed: Auto negotiation
Redundant heartbeat	Interface IP: 10.12.165.24	Interface IP: 10.12.165.23
External interface <input checked="" type="checkbox"/>	HA MAC: aa:7f:d8:71:f4:06	HA MAC: 5a:0d:1b:fa:3c:4b
Management interface <input type="checkbox"/>		
Internal interface <input type="checkbox"/>		
Next hop monitoring	Interface IP: 10.12.165.21	Interface IP: 10.12.165.20
External interface <input checked="" type="checkbox"/>		
Management interface <input type="checkbox"/>		
Internal interface <input type="checkbox"/>		
	Reboot Shutdown	Reboot Shutdown

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 monitoring > Interface 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.
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 monitoring > Interface 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.
5. Repeat the previous steps to add IP addresses to be monitored from the other interfaces if needed.
6. Click **Commit**.

⚠ CAUTION:

For the changes to take effect, you have to restart both nodes. To restart both nodes, click Reboot Cluster.

Upgrading SSB

The syslog-ng Store Box (SSB) appliances are preinstalled with the latest available Long Term Support (LTS) release.

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 condition:

Feature Releases are released and supported in a timeline of 6 (+2) months. You have to keep upgrading SSB 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. We strongly recommend 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 SSB.

The following sections describe how to keep SSB up to date, and how to install a new license:

- Prerequisites: [Upgrade checklist](#) on page 145.
- Upgrading a single node: [Upgrading SSB](#) on page 147.
- Upgrading a high availability cluster: [Upgrading an SSB cluster](#) on page 147.
- Troubleshooting: [Troubleshooting](#) on page 148.
- Renewing the SSB license: [Updating the SSB license](#) on page 148.
- Exporting the configuration of SSB: [Exporting the configuration of SSB](#) on page 150.
- Importing the configuration of SSB: [Importing the configuration of SSB](#) on page 151.

Upgrade checklist

The following list applies to all configurations:

- The firmware of your SSB appliance is not tainted (that is, none of its files were modified locally). If the firmware is tainted, a warning is displayed on the **Basic**

Settings > System > Version details page.

Upgrading is not supported if the firmware is tainted. If your firmware is tainted, [contact our Support Team](#).

- You have created a configuration backup of syslog-ng Store Box (SSB).
For detailed instructions, refer to [Exporting the configuration of SSB](#) on page 150.
- You have a valid [support portal](#) account.
To download the required .ISO file and the license, you need a valid [support portal](#) account. Note that the registration is not automatic, and might require up to two working days to process.
- You have downloaded the latest .ISO file from the [Downloads](#) page. For further information, see [Firmware in SSB](#) on page 24.
- You have read the Release Notes of the firmware(s) before updating. The Release Notes might include additional instructions specific to the firmware version.
The Release Notes are available here on the [Downloads](#) page.

If you have a high availability cluster:

- You have IPMI access to the slave node. You can find detailed information on using the IPMI interface in the following documents:
For syslog-ng Store Box (SSB) Appliance 3000 and 3500, see the [Latest IPMI firmware for Syslog-ng Store Box 3000/3500 and One Identity Safeguard for Privileged Sessions 3000/3500 appliances](#).
For syslog-ng Store Box (SSB) T1 and syslog-ng Store Box (SSB) T4, see the [IPMI User's Guide](#).
- You have verified on the **Basic Settings > High Availability** page that the HA status is not degraded.
- If you have a high availability cluster with geoclustering enabled:
Perform the firmware upload steps an hour before the actual upgrade. Geoclustering can introduce delays in master-slave synchronization, and the slave node might not be able to sync the new firmware from the master node on time.

If you are upgrading SSB 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, SSB displays information about the progress of the upgrade and any possible problems to the console, which you can monitor with IPMI (ILOM) or console access.

We recommend that you test the upgrade process in a non-productive (virtual, etc.) environment first.

Upgrading SSB requires a reboot. We strongly suggest that you perform the upgrade on the productive appliance during maintenance hours only, to avoid any potential data loss.

Upgrading SSB

To upgrade syslog-ng Store Box (SSB) (single node)

1. Navigate to **Basic Settings > System > Upgrade**.



2. Click **Choose File** and select the .ISO file you want to upload.
3. Click **Upload**.
When the upload is finished, read the **Upgrade notes** pop-up window.
4. Click **Upgrade and reboot node**. SSB will automatically upgrade and reboot the new version. Wait for the process to complete.
5. Navigate to **Basic Settings > System > Version details** and check the version numbers of SSB. In case you encounter problems, you can find common troubleshooting steps in [Troubleshooting](#) on page [148](#).

Upgrading an SSB cluster

To upgrade syslog-ng Store Box (SSB)

1. Navigate to **Basic Settings > System > Upgrade**.



2. Click **Choose File** and select the .ISO file you want to upload.
3. Click **Upload**. When the upload is finished, read the **Upgrade notes** pop-up window.
4. Click . SSB will automatically upgrade and reboot the new version. Wait for the process to complete.
NOTE: In High Availability mode, you have to start the slave node through the IPMI interface. Failing to start the slave node results in a **DEGRADED** HA status.
5. Navigate to **Basic Settings > System > Version details** and check the version numbers of SSB. In case you encounter problems, you can find common troubleshooting steps in [Troubleshooting](#) on page 148.

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 (or the F5 key on your keyboard) to remove any cached version of the page.

In the unlikely case that syslog-ng Store Box (SSB) encounters a problem during the upgrade process and cannot revert to its original state, SSB performs the following actions:

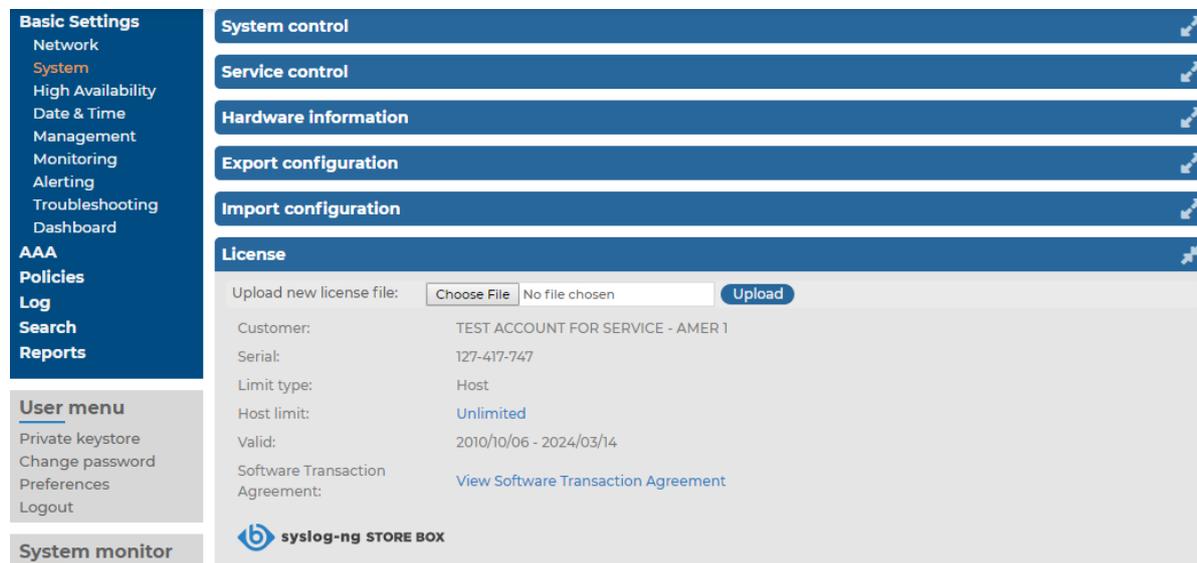
- Initializes the network interfaces using the already configured IP addresses.
- Enables SSH-access to SSB, unless SSB is running in sealed mode. That way it is possible to access the logs of the upgrade process that helps the our Support Team to diagnose and solve the problem. Note that SSH access 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 SSB, check the information displayed on the local console and [contact our Support Team](#).

Updating the SSB license

The syslog-ng Store Box (SSB) license must be updated before the existing license expires or when you purchase a new license. Information of the current license of SSB is displayed on the **Basic Settings > System > License** page. The following information is displayed:

Figure 73: Basic Settings > System > License – Updating the license



- **Customer:** The company permitted to use the license (for example Example Ltd.).
- **Serial:** The unique serial number of the license.
- **Host limit:** The number of peers SSB accepts log messages from.
- **Valid:** The period in which the license is valid. The dates are displayed in YYYY/MM/DD format.

SSB gives an automatic alert one week before the license expires. An alert is sent also when the number of peers exceeds 90% of the limit set in the license.

The following describes how to update the license.

⚠ CAUTION:

Before uploading a new license, you are recommended to backup the configuration of SSB. For details, see [Exporting the configuration of SSB](#) on page 150.

To update the license

1. Navigate to **Basic Settings > System > License**.
2. Click **Choose File** and select the new license file.

NOTE: It is not required to manually decompress the license file. Compressed licenses (for example .zip archives) can also be uploaded.

3. Click **Upload**, then **Commit**.

4. To activate the new license, navigate to **Service control > Syslog traffic, indexing & search:** and click **Restart syslog-ng.**

Exporting the configuration of SSB

The configuration of syslog-ng Store Box (SSB) can be exported (for manual archiving, or to migrate it to another SSB unit) from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

Figure 74: Basic Settings > System — Exporting the SSB configuration

The screenshot shows the 'Basic Settings > System' page. The left sidebar contains a navigation menu with items like 'Basic Settings', 'Network', 'System', 'High Availability', 'Date & Time Management', 'Monitoring', 'Alerting', 'Troubleshooting', 'Dashboard', 'AAA', 'Policies', 'Log', 'Search', and 'Reports'. The main content area is titled 'System control' and contains several sections: 'Service control', 'Hardware information', and 'Export configuration'. The 'Export configuration' section has three radio button options: 'No encryption', 'Encryption with password' (which is selected), and 'GPG encryption'. Below these options are two password input fields: 'Encryption password:' and 'Confirm password:'. The 'Encryption password' field has a strength indicator showing 'weak', 'good' (highlighted), and 'strong'. An 'Export' button is located at the bottom of the form.

To export the configuration of SSB

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:

Exporting the SSB configuration without encryption is not recommended, 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: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used:
! "\$ % & ' () * + , - . / : ; < = > ? @ [] ^ _ ` { | }

- 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 SSB at **Basic Settings > Management > System backup**. For details, see [Encrypting configuration backups with GPG](#) on page 110.

3. Click **Export**.

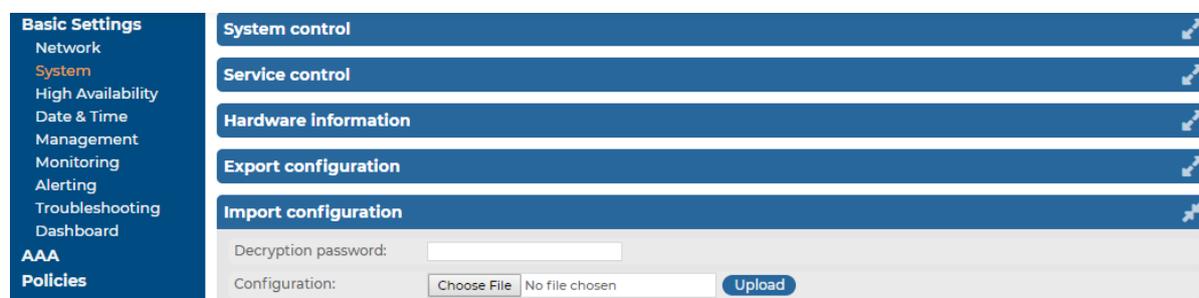
NOTE: The exported file is a gzip-compressed archive. On Windows platforms, it can be decompressed with common archive managers such as the [free 7-Zip tool](#).

The name of the exported file is <hostname_of_SSB>-YYYYMMDDTHHMM.config. The -encrypted or -gpg suffix is added for password-encrypted and GPG-encrypted files, respectively.

Importing the configuration of SSB

The configuration of syslog-ng Store Box (SSB) can be imported from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

Figure 75: Basic Settings > System — Importing the SSB configuration



CAUTION:

It is possible to import a configuration exported from SSB 2.0 or 3.0 into SSB 6.9.0, but it is not possible to restore an 1.1 or 1.0 backup into 6.9.0.

To import the configuration of SSB

1. Navigate to **Basic Settings > System > Import configuration**.
2. Click **Choose File** and select the configuration file to import.
3. Enter the password into the **Decryption password** field and click **Upload**.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

⚠ CAUTION:

When importing an older configuration, it is possible that there are logspaces on SSB that were created after the backing up of the old configuration. In such case, the new logspaces are not lost, but are deactivated and not configured. To make them accessible again, you have to:

1. Navigate to **Log > Logspaces** and configure the logspace. Filling the **Access Control** field is especially important, otherwise the messages stored in the logspace will not be available from the **Search > Logspaces** interface.
2. Adjust your log path settings on the **Log > Paths** page. Here you have to re-create the log path that was sending messages to the logspace.

Accessing the SSB console

This section describes how to use the console menu of syslog-ng Store Box (SSB), how to enable remote SSH access to SSB, and how to change the root password from the web interface.

Using the console menu of SSB

Connecting to the syslog-ng Store Box (SSB) locally or remotely using Secure Shell (SSH) allows you to access the console menu of SSB. The console menu provides access to the most basic configuration and management settings of SSB. It is mainly used for troubleshooting purposes, the primary interface of SSB is the web interface.

NOTE: Detailed host information is displayed in the shell prompt:

The format of the bash prompt is:

```
(firmware_type/HA_node/hostname)username@HA_node_name:current_working_directory#
```

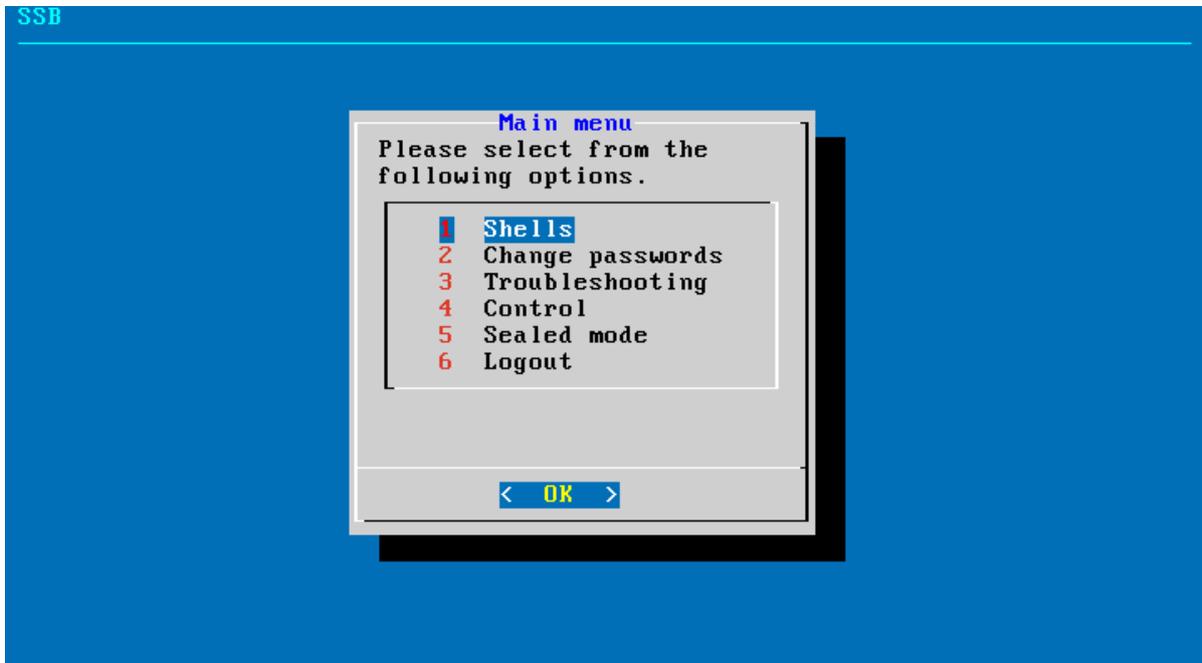
For example:

```
(core/master/documentation-ssb)root@ssb1:/etc#
```

- `firmware_type` is either `boot` or `core`
- `HA_node` is either `master` or `slave`
- `hostname` is the FQDN set on the GUI
- `username` is always `root`

The console menu is accessible to the root user using the password set during completing the Welcome Wizard.

Figure 76: The console menu



The console menu allows you to perform the following actions:

- Change the passwords of the root and admin users.
- Access the local core and boot shells. This is usually not recommended and only required in certain troubleshooting situations. In a high-availability cluster, you can also access the boot shell of the remote node.
- Access the network-troubleshooting functions and display the available log files.
- Reboot and shut down the system. In a high-availability cluster, you can also reboot or shutdown the remote node, or initiate a take over.
- Enable and disable sealed mode. For details, see [Sealed mode](#).
- Set the IP address of the HA interface.

This option is not available on virtual appliances, or if your SSB license does not include the HA option. If High Availability (HA) operation mode is required in a virtual environment, use the HA function provided by the virtual environment.

NOTE: Note that logging in to the console menu automatically locks the SSB 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.

On the secondary node of a high-availability cluster, the console menu is limited to troubleshooting and diagnostic functions (for example, you cannot change passwords or manage sealed mode).

Enabling SSH access to the SSB host

Exclusively for troubleshooting purposes, you can access the syslog-ng Store Box (SSB) host using SSH. Completing the Welcome Wizard automatically disables SSH access. To enable it again, complete the following steps:

⚠ CAUTION:

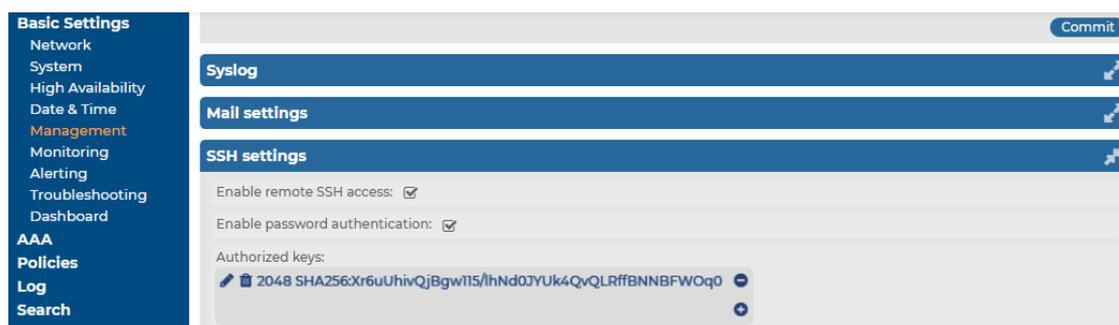
Accessing the SSB host directly using SSH is not recommended nor supported, except for troubleshooting purposes. In such case, the One Identity Support Team will give you exact instructions on what to do to solve the problem.

Enabling the SSH server allows you to connect remotely to the SSB host and login using the root user. The password of the root user is the one you had to provide in the Welcome wizard. For details on how to change the root password from the web interface, see [Changing the root password of SSB](#) on page 155.

To enable SSH access to the SSB host

1. Navigate to **Basic Settings > Management > SSH settings**.

Figure 77: Basic Settings > Management > SSH settings — Enabling remote SSH access to SSB



2. Select the **Enable remote SSH access** option.

NOTE: Remote SSH access is automatically disabled if Sealed mode is enabled. For details, see [Sealed mode](#) on page 155.

3. Set 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 SSB remotely via SSH.
4. Click **Commit**.

The SSH server of SSB accepts connections only on the management interface if the management interface is configured. If the management interface is not configured, the SSH server accepts connections on the external interface. If possible, avoid

enabling the SSH server of SSB when the management interface is not configured. For details on enabling the management connection, see [Configuring the management interface](#) on page 77.

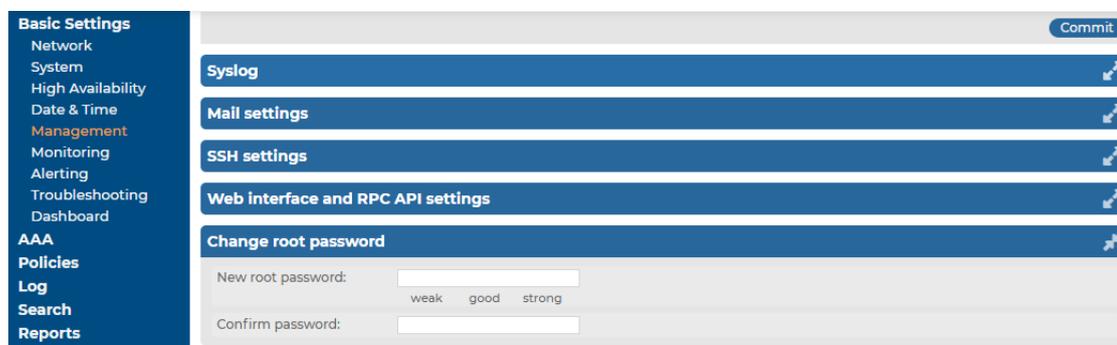
Changing the root password of SSB

The root password is required to access syslog-ng Store Box (SSB) 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 SSB console](#) on page 152.

To change the root password of SSB

1. Navigate to **Basic Settings > Management > Change root password**.

Figure 78: Basic Settings > Management > Change root password — Changing the root password of SSB



2. Enter the new password into the **New root password** and **Confirm password** fields. The password must meet the requirements of the **AAA > Settings > Password settings > Minimal password strength** option.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

3. Click **Commit**.

Sealed mode

When sealed mode is enabled, the following settings are automatically applied:

- The syslog-ng Store Box (SSB) appliance cannot be accessed remotely via SSH for maintenance. Also, configuration settings related to remote SSH access are deleted.
- The root password of SSB cannot be changed in sealed mode.
- Sealed mode can be disabled only from the local console. For details, see [Disabling sealed mode](#) on page 156.

To enable sealed mode use one of the following methods:

- Select the **Sealed mode** option during the Welcome Wizard.
- Navigate to **Basic Settings > System > Sealed mode > Activate sealed mode** on the SSB web interface and click **Enable**.
- Log in to SSB as root using SSH or the local console, and select **Sealed mode > Enable** from the console menu.

Disabling sealed mode

This section describes how to disable sealed mode.

To disable sealed mode

1. Go to the syslog-ng Store Box (SSB) appliance and access the local console.
2. Login as root.
3. From the console menu, select **Sealed mode > Disable**.
4. Select **Back to Main menu > Logout**.
5. If you want to access SSB remotely using SSH, configure SSH access. Disabling sealed mode does not restore any previous SSH configuration. For details, see [Enabling SSH access to the SSB host](#) on page 154.

Out-of-band management of SSB

Physical syslog-ng Store Box (SSB) appliances include a dedicated out-of-band management interface conforming to the Intelligent Platform Management Interface (IPMI) v2.0 standards. The IPMI interface allows system administrators to monitor the system health of SSB and to manage the computer events remotely, independently of the operating system of SSB. SSB is accessible using the IPMI interface only if the IPMI interface is physically connected to the network.

Note that the IPMI interface supports only 100Mbps Full-Duplex speed.

- For details on connecting the IPMI interface, see ["Installing the SSB hardware" in the Installation Guide](#).
- For details on configuring the IPMI interface, see [Configuring the IPMI interface from the console](#) on page 158.
- For details on using the IPMI interface to remotely monitor and manage SSB, see the following documents:

For syslog-ng Store Box (SSB) Appliance 3000 and 3500, see the [Latest IPMI firmware for Syslog-ng Store Box 3000/3500 and One Identity Safeguard for Privileged Sessions 3000/3500 appliances](#).

For syslog-ng Store Box (SSB) T1 and syslog-ng Store Box (SSB) T4, see the [IPMI User's Guide](#).

Basic information about the IPMI interface is available also on the SSB web interface on the **Basic Settings > High Availability** page. The following information is displayed:

Figure 79: Basic Settings > High Availability – Information about the IPMI interface SSB

The screenshot displays the 'High availability & Nodes' configuration page. On the left is a navigation sidebar with sections: Basic Settings (Network, System, High Availability, Date & Time, Management, Monitoring, Alerting, Troubleshooting, Dashboard), AAA, Policies, Log, Search, and Reports. Below the sidebar is a 'User menu' (Private keystore, Change password, Preferences, Logout) and a 'System monitor' (Time: 2019-11-07 13:30, Remaining time: 11:58:36, Locked: admin@10.30.255.68, HA status: HA active: 087AC40C-7235-0706-0809-0CC47A083572, Modules: syslog-ng: Running, Active Hosts: 3, Senders: 2, Load 1: 0.03 Load 15: 0.07, and a CPU/Mem/Disk/Swap bar chart). The main content area has a 'Commit' button at the top right. The 'High availability & Nodes' section shows: Cluster status: SSB is currently operating in HA. Current master: 087AC40C-7235-0706-0809-0CC47A083572. HA UUID: e05f8350-c212-408f-b0d1-b039c8f8983b. DRBD Status: Connected - Connected, Connected. DRBD sync rate limit: 1 Gbit/s. DRBD asynchronous mode: . Below are two columns: 'This node' and 'Other node'. 'This node' details: Node ID: 087AC40C-7235-0706-0809-0CC47A083572, Node HA status: HA, Node HA UUID: e05f8350-c212-408f-b0d1-b039c8f8983b, DRBD status: Connected (UpToDate) Connected, RAID status: All partition: active, Boot firmware versions: Current: 6.3.0-RC2, Active: 6.3.0-RC2, IPMI IP address: 10.12.193.68, IPMI subnet mask: 255.255.255.0, IPMI default gateway: 10.12.193.254, IPMI IP address source: DHCP Address, HA link speed: Auto negotiation, Interfaces for Heartbeat: HA (Fix current) 1.2.4.2 (FIX), External , Internal , Management 1.6.6.2, HA MAC: Next hop monitoring: External 10.120.75.151, Internal , Management . 'Other node' details: Node ID: 017AC40C-AA59-0706-0809-0CC47A0159AA, Node HA status: HA, Node HA UUID: e05f8350-c212-408f-b0d1-b039c8f8983b, DRBD status: Connected (UpToDate) Connected, RAID status: All partition: active, Boot firmware versions: Current: 6.3.0-RC2, Active: 6.3.0-RC2, IPMI IP address: 10.12.193.69, IPMI subnet mask: 255.255.255.0, IPMI default gateway: 10.12.193.254, IPMI IP address source: DHCP Address, HA link speed: Auto negotiation, Interfaces for Heartbeat: Interface IP: 1.2.4.1 (FIX), External , Internal , Management 1.6.6.1, HA MAC: Next hop monitoring: External 10.120.75.11, Internal , Management . Buttons for Reboot and Shutdown are at the bottom of each node configuration.

- **Hardware serial number:** The unique serial number of the appliance.
- **IPMI IP address:** The IP address of the IPMI interface.
- **IPMI subnet mask:** The subnet mask of the IPMI interface.
- **IPMI default gateway IP:** The address of the default gateway configured for the IPMI interface.
- **IPMI IP address source:** Shows how the IPMI interface receives its IP address: dynamically from a DHCP server, or it uses a fixed static address.

Configuring the IPMI interface from the console

The following describes how to modify the network configuration of IPMI from the console of syslog-ng Store Box(SSB).

Prerequisites:

SSB is accessible using the IPMI interface only if the IPMI interface is physically connected to the network. For details on connecting the IPMI interface, see ["Installing the SSB 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 SSB is powered on.

It is not necessary for the IPMI interface to be accessible from the Internet, but the administrator of SSB must be able to access it for support and troubleshooting purposes in case vendor support is needed. The following ports are used by the IMPI interface:

- Port 623 (UDP): IPMI (cannot be changed)
- Port 5123 (UDP): floppy (cannot be changed)
- Port 5901 (TCP): video display (configurable)
- Port 5900 (TCP): HID (configurable)
- Port 5120 (TCP): CD (configurable)
- Port 80 (TCP): HTTP (configurable)

To modify the network configuration of IPMI from the console of SSB

1. Use the local console (or SSH) to log in to SSB 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.
 - 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. Configure IPMI to use the dedicated Ethernet interface. Issue the following command, which works for all [appliances](#) (that is, for syslog-ng Store Box T1, syslog-ng Store Box T4, syslog-ng Store Box T10, syslog-ng Store Box Appliance 3000, and syslog-ng Store Box Appliance 3500):

```
ipmitool raw 0x30 0x70 0xc 1 0
```

6. Verify the network configuration of IPMI:

```
# ipmitool lan print 1
```

Use a browser to connect to the reported network address.

7. Change the default password:

- a. Log in to the IPMI web interface using the default login credentials (username: ADMIN, password: ADMIN).

| **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 interface from the BIOS

To configure IPMI from the BIOS when configuring your syslog-ng Store Box (SSB) physical appliance for the first time, complete the following steps.

Prerequisites:

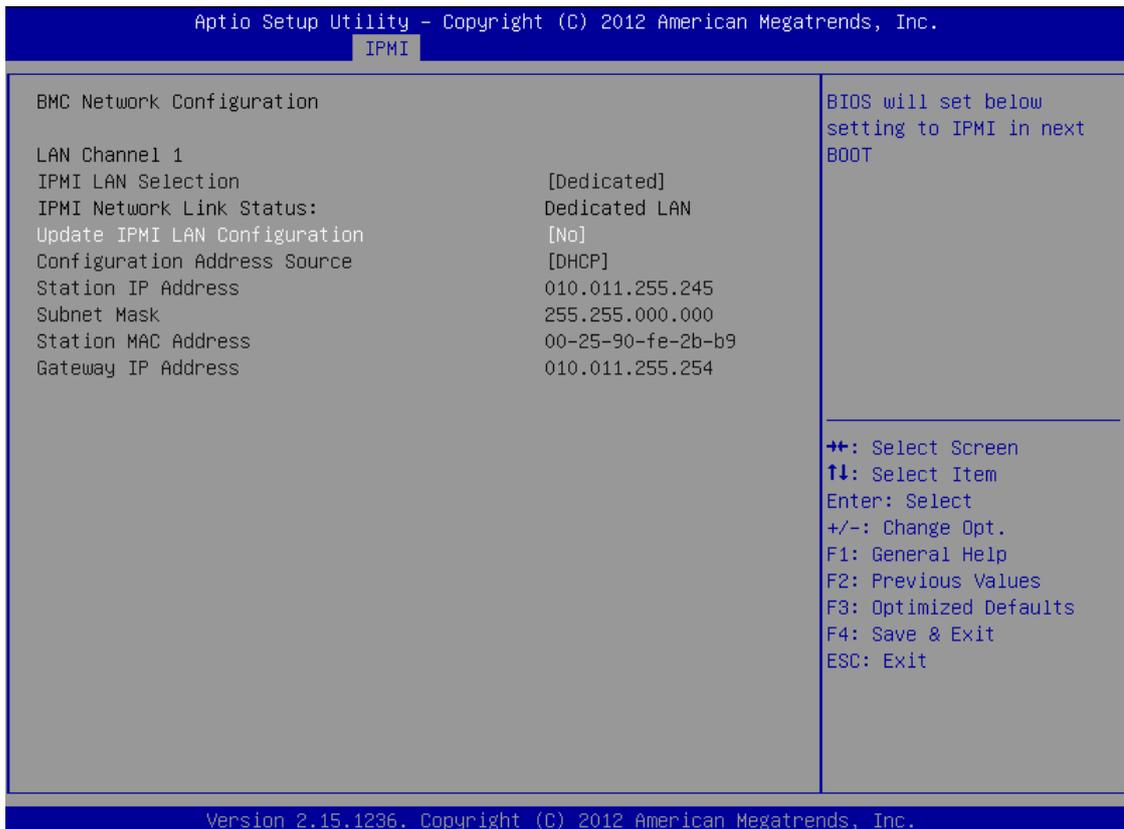
To apply the procedure outlined here, you will need physical access to a monitor and keyboard.

Figure 81: IMPI page > BMC Network Configuration option



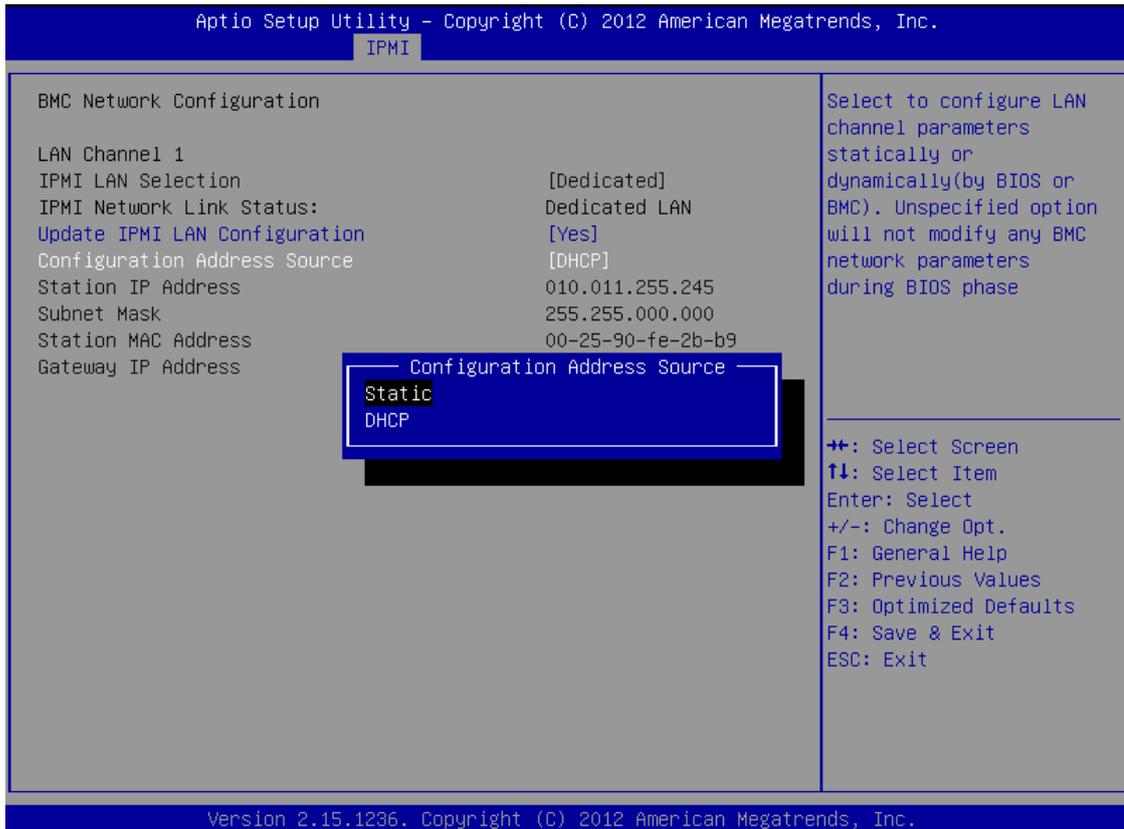
4. On the **BMC Network Configuration** page, select **Update IPMI LAN Configuration**, press Enter, and select **Yes**.

Figure 82: BMC Network Configuration page > Update IPMI LAN Configuration



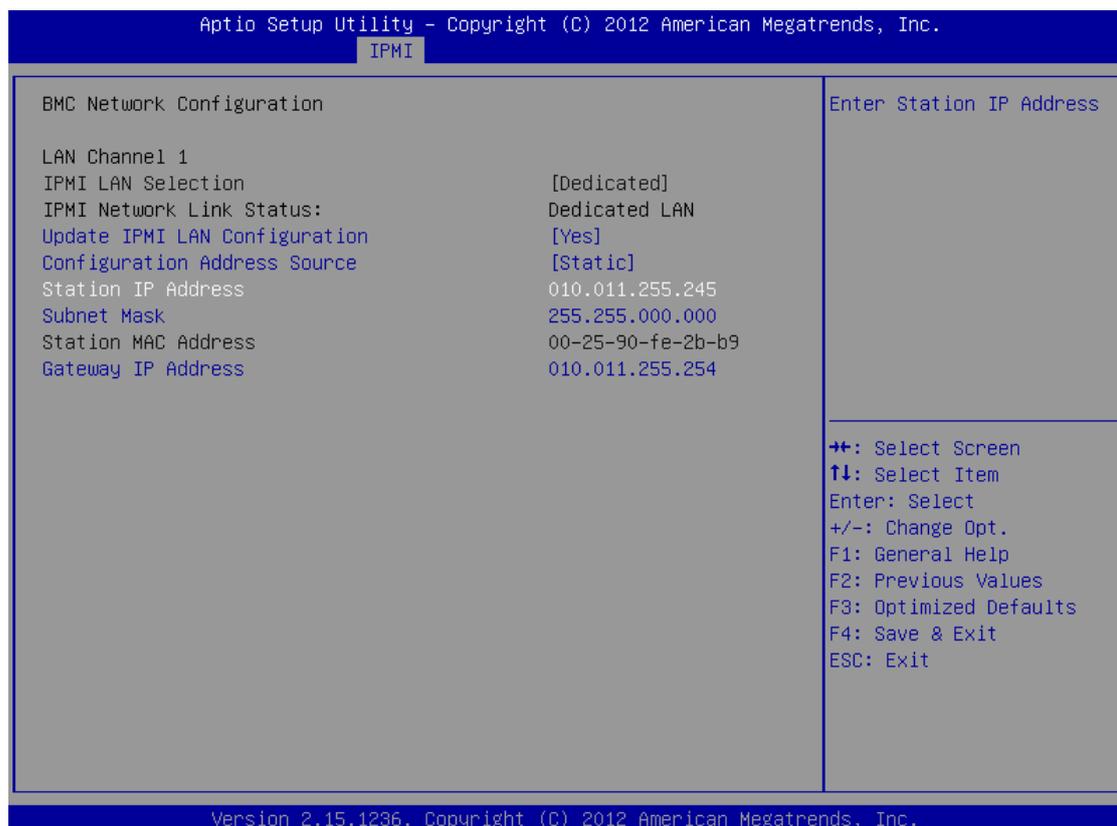
5. Stay on the **BMC Network Configuration** page, select **Configuration Address Source**, press **Enter**, and select **Static**.

Figure 83: 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 84: 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 SSB

SSB uses a number of certificates for different tasks that can be managed from the **Basic Settings > Management > SSL certificate** menu.

Figure 85: Basic Settings > Management > SSL certificate — Changing the web certificate of SSB

The following certificates can be modified here:

- **CA certificate:** The certificate of the internal Certificate Authority of SSB.
 - NOTE:** When you upload your own CA certificate, make sure that the certificate you upload is the issuer certificate of the Server and TSA certificates.
- **Server certificate:** The certificate of the SSB web interface, used to encrypt the communication between SSB and the administrators.
 - NOTE:** If this certificate is changed, the browser of SSB users will display a warning stating that the certificate of the site has changed.
 - NOTE:** When you have a certificate chain, you have to upload the entire chain in a single file, using PEM format. The uploaded file (or pasted text) must contain the following elements, concatenated in this order:
 1. the server certificate
 2. the issuer CA

- 3. the root CA certificates.
- **TSA certificate:** The certificate of the internal Time Stamping Authority that provides the time stamps used when creating encrypted logstores.

NOTE:

SSB uses other certificates for different purposes that are not managed here, for example, to encrypt data stored on SSB. For details, see [Creating logstores](#) on page 217.

Use every keypair or certificate only for one purpose. Do not reuse cryptographic keys or certificates, for example, do not use the same certificate for the SSB webserver and for encrypting logstores.

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 SSB web interface, but the public part of the key can be downloaded in different formats (for example, PEM, DER, OpenSSH, Tectia). Also, the X.509 certificate can be downloaded in PEM and DER formats, with the exception of certificate chains, which can only be downloaded in PEM format.

NOTE: Other parts of SSB may use additional certificates that are not managed here.

During the initial configuration, SSB creates a self-signed CA certificate, and uses this CA to issue the certificate of the web interface (see **Server certificate**) and the internal Time Stamping Authority (**TSA certificate**).

There are two methods to manage certificates of SSB:

- **Recommended:** Generate certificates using your own PKI solution and upload them to SSB.

Generate a CA certificate and two other certificates signed with this CA using your PKI solution and upload them to SSB. For the Server and TSA certificates, upload the private key as well. One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
- Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.

For details on uploading certificates and keys created with an external PKI, complete [Uploading external certificates to SSB](#) on page 167.

CAUTION:

The Server and the TSA certificates must be issued by the same Certificate Authority.

- Use the certificates generated on SSB. In case you want to generate new certificates and keys for SSB using its self-signed CA certificate, or generate a new self-signed CA certificate, complete [Generating certificates for SSB](#) on page 167.

NOTE: Generate certificates using your own PKI solution and upload them to SSB whenever possible. Certificates generated on SSB cannot be revoked, and can become a security risk if they are somehow compromised.

Generating certificates for SSB

Create a new certificate for the syslog-ng Store Box (SSB) webserver or the Time Stamping Authority using the internal CA of SSB, or create a new, self-signed CA certificate for the internal Certificate Authority of SSB.

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

To generate certificates for SSB

1. Navigate to **Basic Settings > Management > SSL certificate**.
2. Fill the fields of the new certificate:
 - a. **Country**: Select the country where SSB is located (for example HU - Hungary).
 - b. **Locality**: The city where SSB is located (for example Budapest).
 - c. **Organization**: The company who owns SSB (for example Example Inc.).
 - d. **Organization unit**: The division of the company who owns SSB (for example IT Security Department).
 - e. **State or Province**: The state or province where SSB is located.
3. Select the certificate you want to generate.
 - To create a new certificate for the SSB web interface, select **Generate Server certificate**.
 - To create a new certificate for the Time Stamping Authority, select **Generate TSA certificate**.
 - To create a new certificate for the internal Certificate Authority of SSB, 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. If you have uploaded an external CA certificate along with its private key, it will be used to create the new server and TSA certificates. If you have uploaded an external CA certificate without its private key, use your external PKI solution to generate certificates and upload them to SSB.

CAUTION:
Generating a new certificate automatically deletes the earlier certificate.

4. Click **Commit**.

Uploading external certificates to SSB

Upload a certificate generated by an external PKI system to syslog-ng Store Box (SSB).

The certificate to upload. For the TSA and Server certificate, the private key of the certificate is needed as well. The certificates must meet the following requirements:

- SSB accepts certificates in PEM format. The DER format is currently not supported.
- SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

For the internal CA certificate of SSB, uploading the private key is not required.

- One Identity recommends:
 - Using 2048-bit RSA keys (or stronger).
 - Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.
- For the TSA certificate, the X509v3 Extended Key Usage attribute must be enabled and set to `critical`. Also, its default value must be set to `Time Stamping`.
- For the Server certificate, the X509v3 Extended Key Usage attribute must be enabled and its default value set to `TLS Web Server Authentication`. Also, the Common Name of the certificate must contain the domain name or the IP address of the SSB host. If the web interface is accessible from multiple interfaces or IP addresses, list every IP address using the `Subject Alt Name` option.

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

To upload a certificate generated by an external PKI system to SSB

1. Navigate to **Basic Settings > Management > SSL certificate**.
2. To upload a new certificate, click  next to the certificate you want to modify. A pop-up window is displayed.

Figure 86: Basic Settings > Management > SSL certificate – Uploading certificates

The screenshot shows a pop-up window titled "Server X.509 certificate" with a close button in the top right corner. Inside the window, there are two main sections:

- Upload certificate:** This section has a sub-header "Upload:" followed by a "Choose File" button, a text input field containing "No file chosen", and an "Upload" button.
- Copy-paste certificate:** This section has a sub-header "Certificate:" followed by a "Set" button. Below this is a large, empty text area for pasting the certificate content.

Select **Browse**, select the file containing the certificate, and click **Upload**. Alternatively, you can also copy-paste the certificate into the **Certificate** field and click **Set**.

You can choose to upload a single certificate or a certificate chain (that is, intermediate certificates and the end-entity certificate).

After uploading a certificate or certificate chain, you can review details by clicking the name of the certificate, and looking at the information displayed in the pop-up window that comes up.

Figure 87: Log > Options > TLS settings — X.509 certificate details



The pop-up window allows you to:

- Download the certificate or certificate chain.
 | **NOTE:** Certificate chains can only be downloaded in PEM format.
- View and copy the certificate or certificate chain.
- Check the names and the hierarchy of certificates (if it is a certificate chain and there is more than one certificate present).
 On hovering over a certificate name, the subject of the certificate is displayed, describing the entity certified.
- Check the validity dates of the certificate or certificates making up the chain.
 On hovering over a particular date, the exact time of validity is also displayed.

After uploading the certificate or certificate chain, the presence or absence of the string (**chain**) displayed after the name of the certificate will indicate whether the certificate is a certificate chain or a single certificate.

3. To upload the private key corresponding to the certificate, click  icon. 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**.

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 download the certificate in one single PEM or DER file.

Note that certificate chains can only be downloaded in PEM format.

Generating TSA certificate with Windows Certificate Authority on Windows Server 2008

To generate a TSA certificate with Windows Certificate Authority (CA) that works with syslog-ng Store Box (SSB), generate a CSR (certificate signing request) on a computer running OpenSSL and sign it with Windows CA, then import this certificate into SSB for time stamping.

Prerequisites:

A valid configuration file for OpenSSL with the following extensions:

```
[ tsa_cert ]
extendedKeyUsage = critical,timeStamping
```

TIP: You can copy `/etc/ssb/openssl-ca.cnf` from SSB 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 SSB, if the following conditions are met:

- Must be a valid CA certificate (**CA** is true).
- **Key Usage:** Time Stamping is required. No other key usage is permitted.
- **Extended Key Usage:** Must be set to `critical`.
- **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.

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.

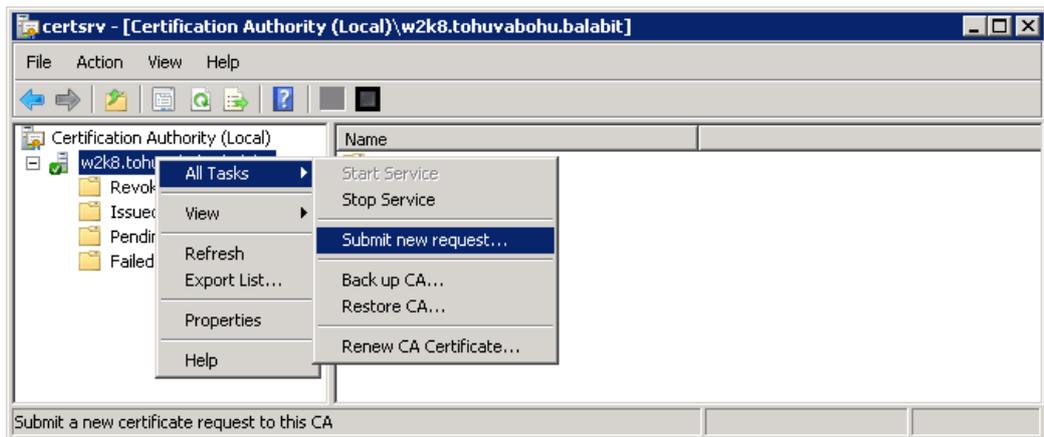
To generate TSA certificate with Windows Certificate Authority on Windows Server 2008

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]:HU
State or Province Name (full name) []:Budapest
Locality Name (eg, city) []:Budapest
Organization Name (eg, company) [Internet Widgits Pty Ltd]:BalaBit IT
Security
Organizational Unit Name (eg, section) []:Service Delivery
Common Name (eg, YOUR name) []:scb35-1-i1.tohuvabohu.balabit
Email Address []:vlad@balabit.com
```

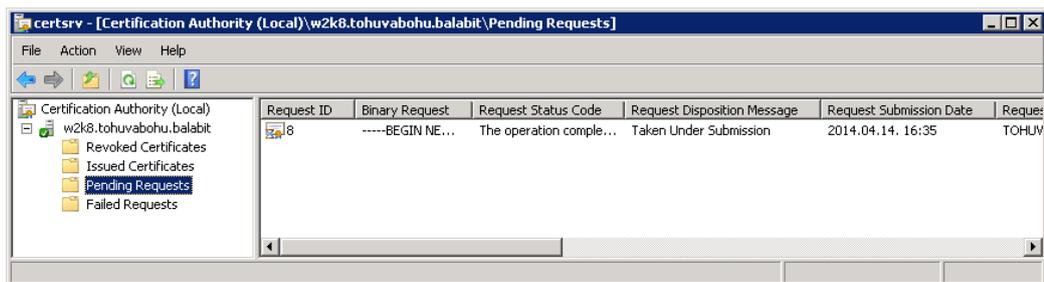
3. Sign the generated CSR with your Windows CA. Make sure that the CSR file is accessible from your Windows CA server.
 - a. To issue and sign the new certificate request, open the Microsoft Certification Authority Management Console: **Start > Run** and run `certsrv.msc`.
 - b. Right-click on the server name and navigate to **All Tasks > Submit new request**.

Figure 88: Submitting a new request



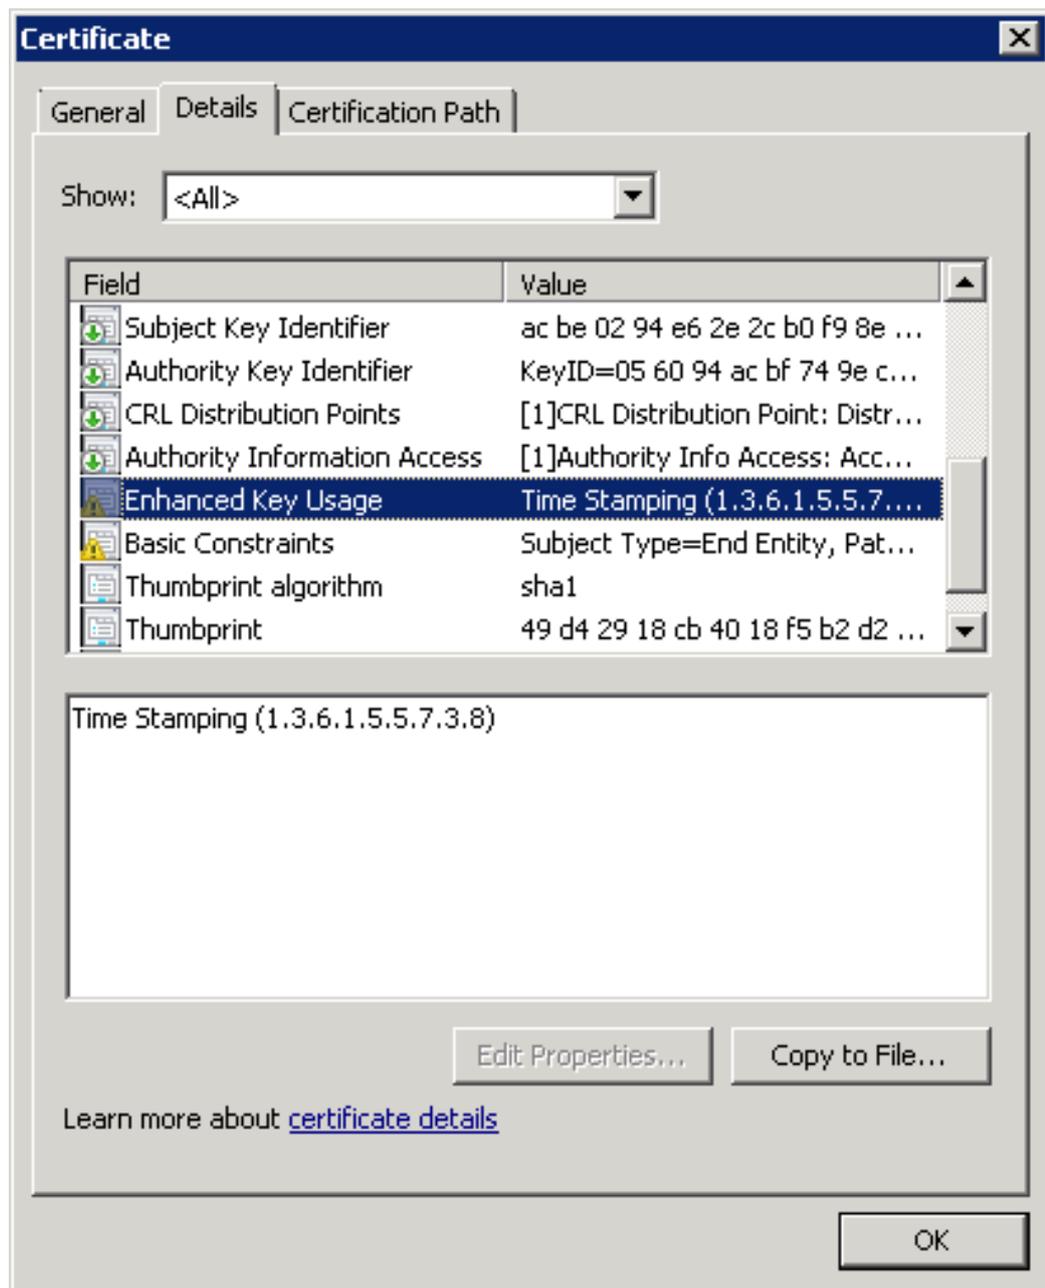
- c. Select the CSR created in the second step.
- d. On the left pane, click **Pending Requests**. The new certificate request is displayed in the right pane.

Figure 89: Issuing a new certificate



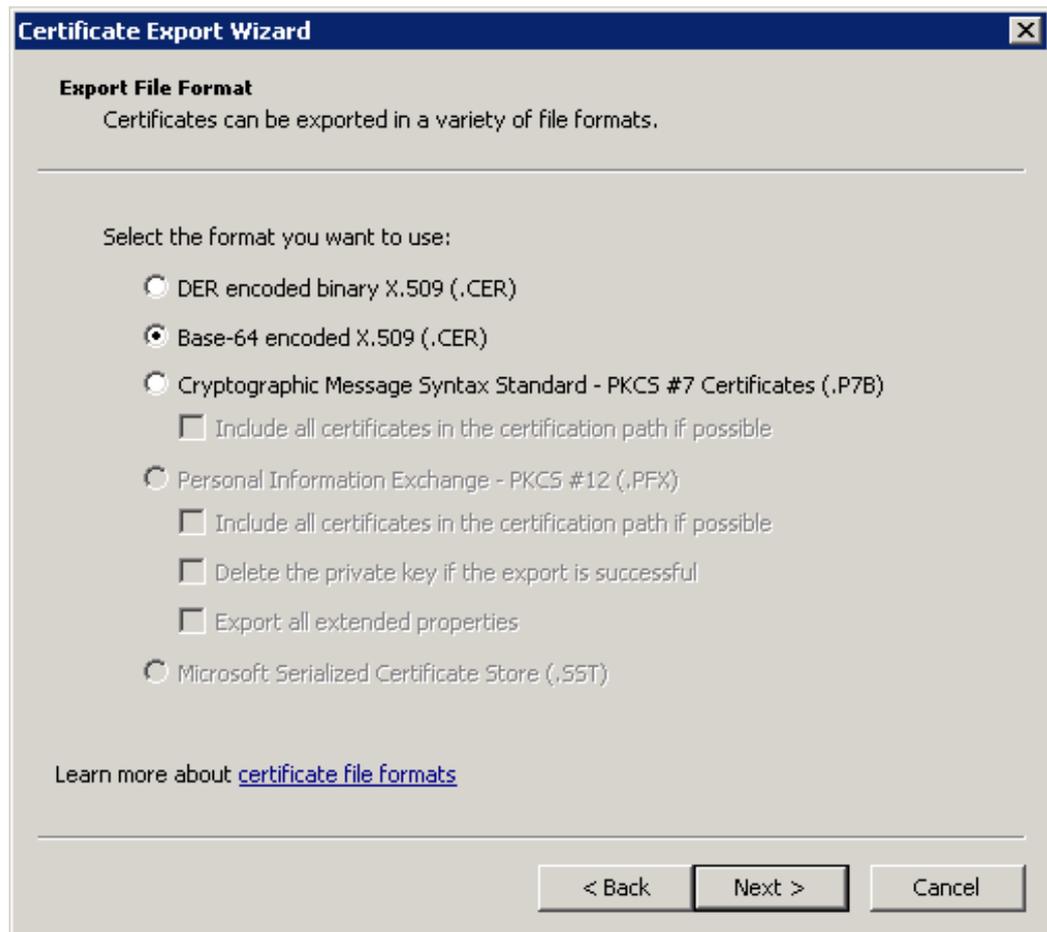
- e. To issue the new SSL certificate, right-click the pending certificate request, select **All Tasks** and click **Issue**.
- f. Select **Issued Certificates** and double-click on the certificate issued in the previous step.
- g. The CA Certificate window opens. Navigate to the **Details** tab. Ensure that the required **Enhanced Key Usage** field is visible and contains the Time Stamping value.

Figure 90: Verifying certificate details



- h. Click **Copy to File**. The Certificate Export Wizard launches. Click **Next**.
- i. Select the format of the certificate: **Base-64 encoded X.509 (.CER)**. Click **Next**.

Figure 91: Selecting certificate file format



- j. Select location to save the certificate, and save it.
 - k. The **Completing the Certificate Export Wizard** screen is displayed. Click **Finish**.
4. In SSB, navigate to **Basic Settings > Management > SSL certificate**.
 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 certificate**) that is used for other certificates on SSB is different from the CA that was used to sign the TSA certificate, a warning is displayed. In this scenario, ignore this warning.

Generating TSA certificate with Windows Certificate Authority on Windows Server 2012

To generate a TSA certificate with Windows Certificate Authority (CA) that works with syslog-ng Store Box (SSB), generate a CSR (certificate signing request) on a computer running OpenSSL and sign it with Windows CA, then import this certificate into SSB for time stamping.

Prerequisites:

A valid configuration file for OpenSSL with the following extensions:

```
[ tsa_cert ]
extendedKeyUsage = critical,timeStamping
```

TIP: You can copy `/etc/ssb/openssl-ca.cnf` from SSB 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 SSB, if the following conditions are met:

- Must be a valid CA certificate (**CA** is true).
- **Key Usage:** Time Stamping is required. No other key usage is permitted.
- **Extended Key Usage:** Must be set to `critical`.
- **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.

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 2012

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`**

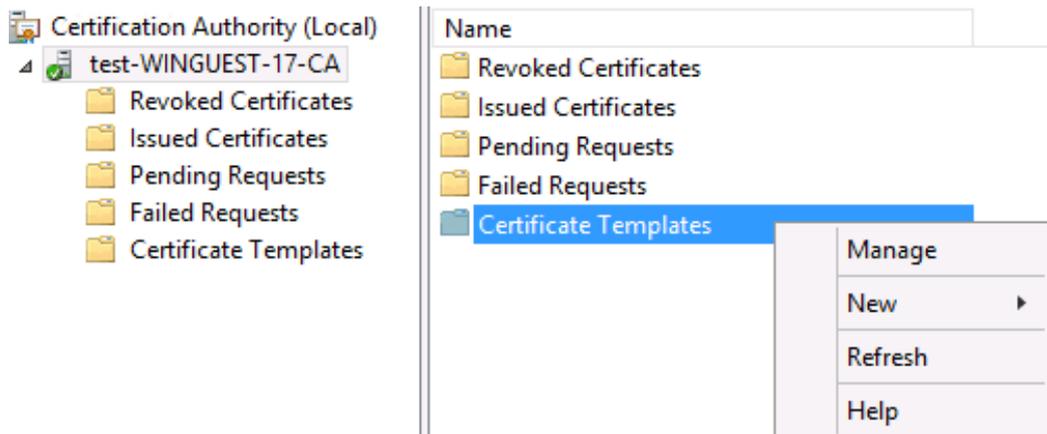
timestamp.key -out timestamp.csr -nodes

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]:HU
State or Province Name (full name) []:Budapest
Locality Name (eg, city) []:Budapest
Organization Name (eg, company) [Internet Widgits Pty Ltd]:BalaBit IT
Security
Organizational Unit Name (eg, section) []:Service Delivery
Common Name (eg, YOUR name) []:scb35-1-i1.tohuvabohu.balabit
Email Address []:vlad@balabit.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 **Certificate Templates**, and choose **Manage**.

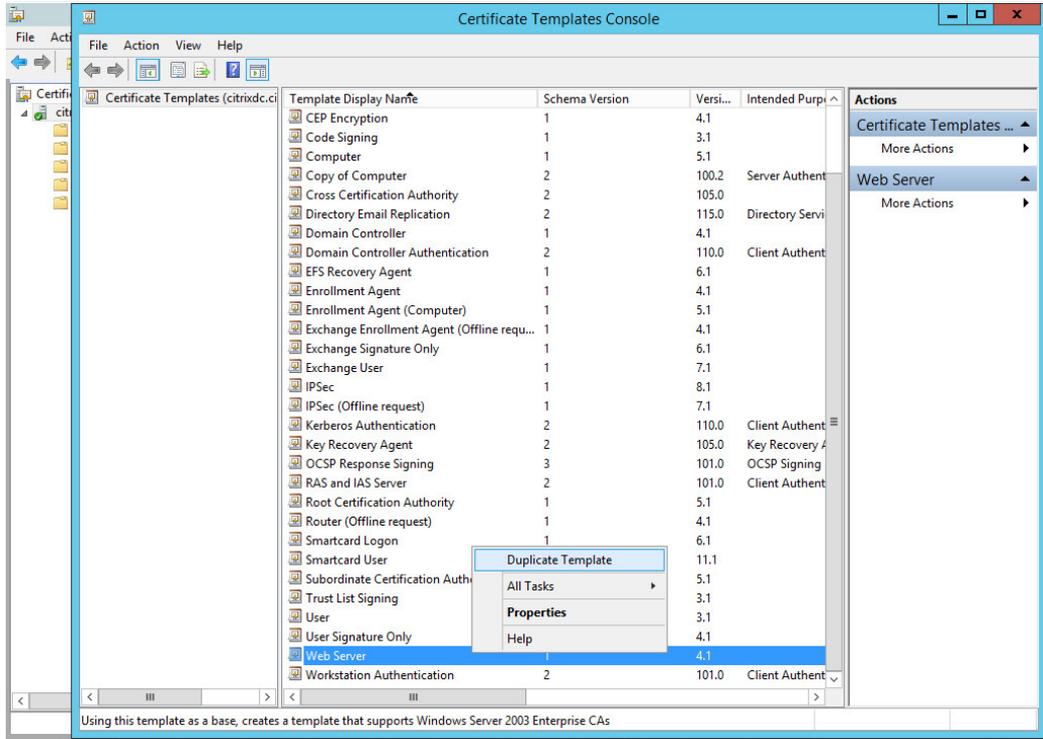
Figure 92: Managing certificate templates



The **Certificate Templates Console** opens.

- c. Right-click the **Web Server** template, and choose **Duplicate Template**.

Figure 93: 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 94: Creating the new template

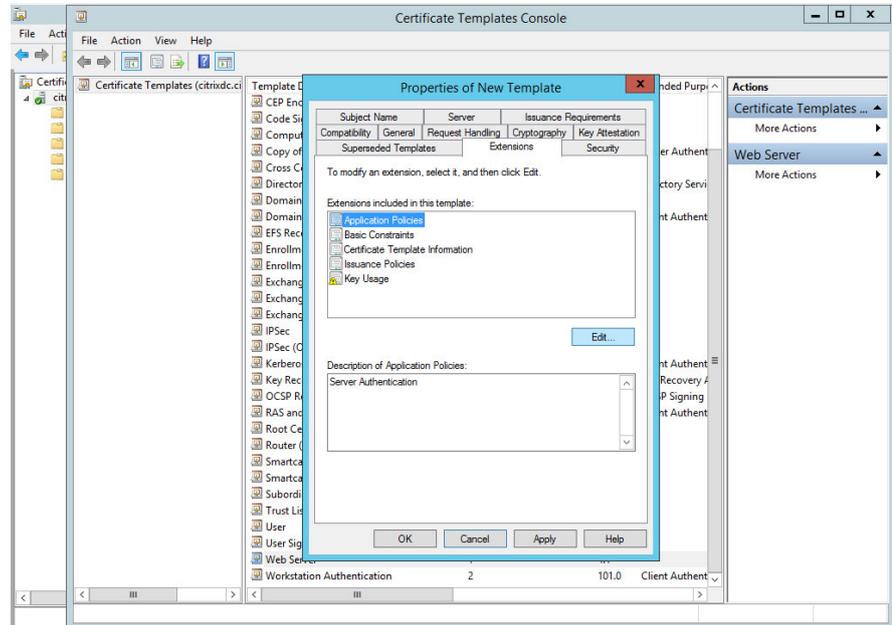
The screenshot shows a dialog box titled "Properties of New Template" with a close button (X) in the top right corner. The dialog has several tabs: "Subject Name", "Server", "Issuance Requirements", "Superseded Templates", "Extensions", "Security", "Compatibility", "General", "Request Handling", "Cryptography", and "Key Attestation". The "General" tab is selected. Inside the dialog, there are two text input fields: "Template display name:" and "Template name:", both containing the text "TSA". Below these are two sets of dropdown menus for "Validity period:" (set to 2 years) and "Renewal period:" (set to 6 weeks). There are two unchecked checkboxes: "Publish certificate in Active Directory" and "Do not automatically reenroll if a duplicate certificate exists in Active Directory". At the bottom of the dialog are four buttons: "OK", "Cancel", "Apply", and "Help".

- On the **Request Handling** tab, enable the **Allow private key to be exported** option.
- On the **Extensions** tab, make the following changes:

To edit Application Policies

- Select **Application Policies** and click **Edit** below the list of extensions.

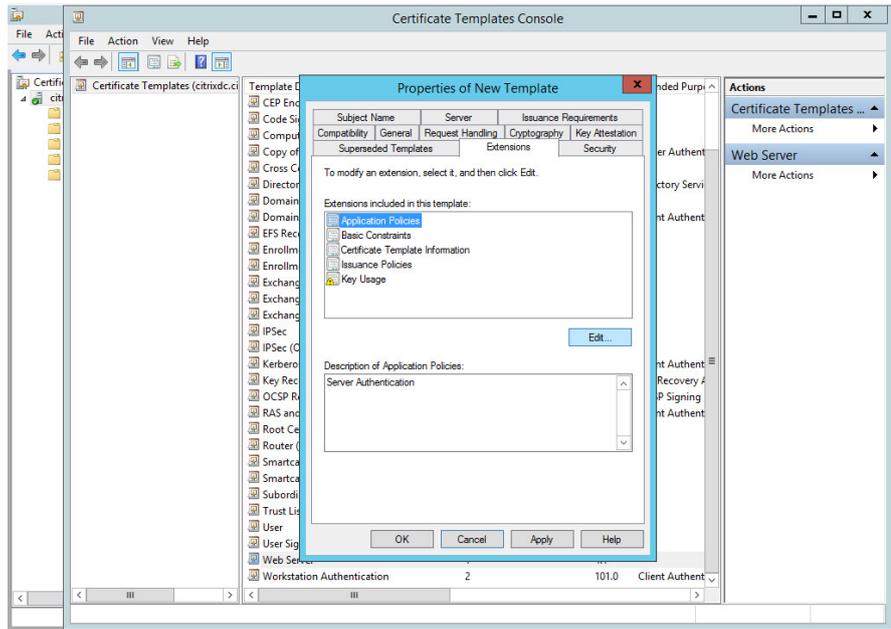
Figure 95: Editing Application Policies



To remove Server Authentication

-
- Select **Server Authentication** and click **Remove**.

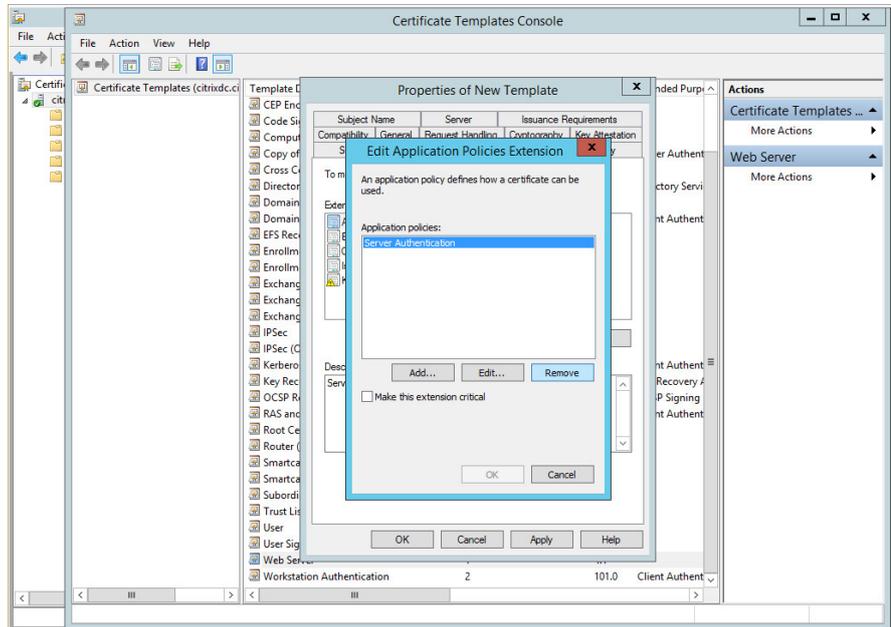
Figure 96: Removing Server Authentication



To add Time Stamping

- Click **Add**, select **Time Stamping** and click **OK**.

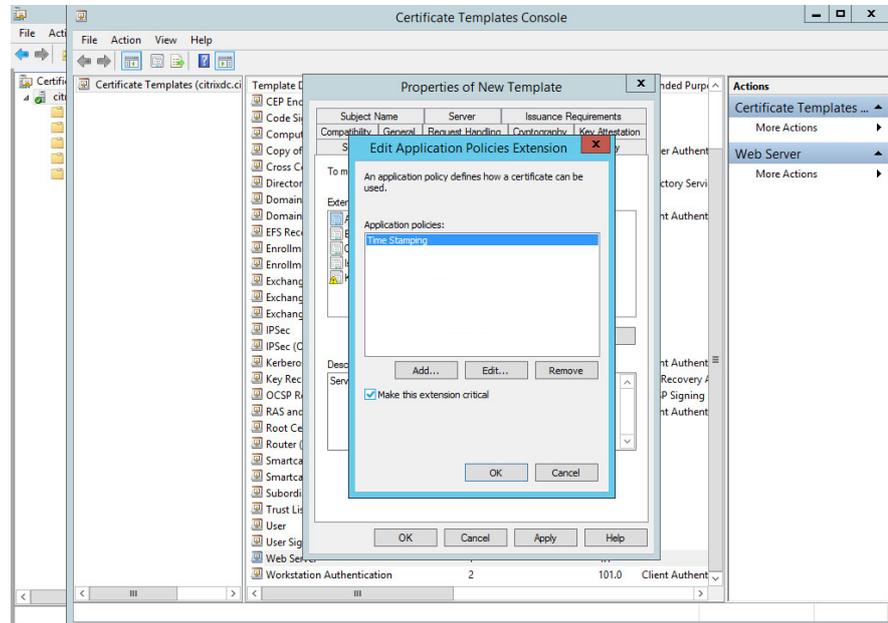
Figure 97: Adding Time Stamping



To make Time Stamping critical

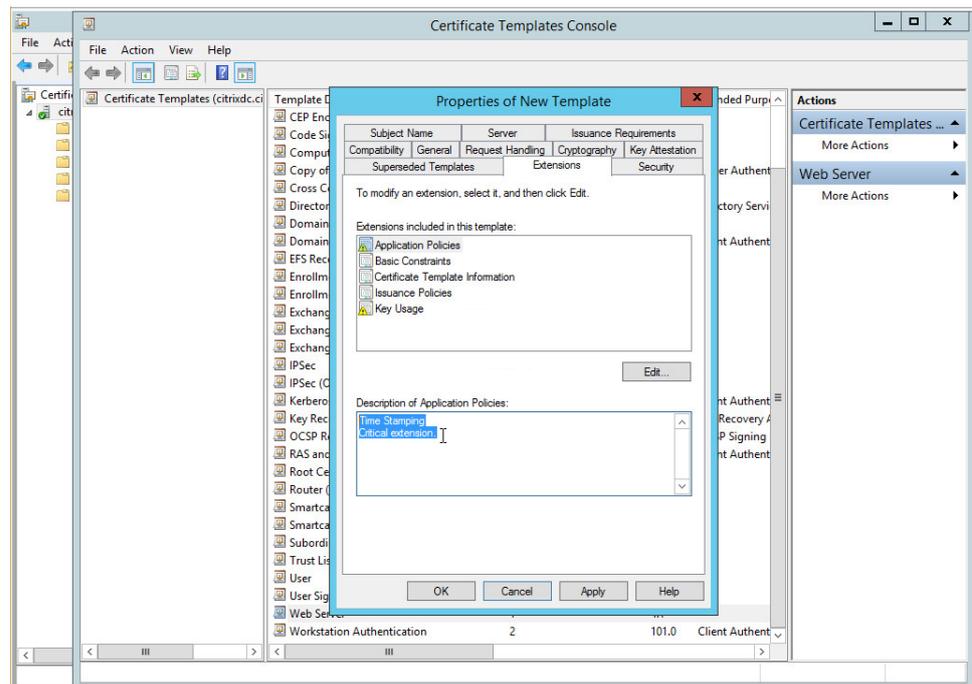
- Select **Time Stamping** and enable the **Make this extension critical** option, then click **OK**.

Figure 98: Making Time Stamping critical



Time Stamping and Critical extension are listed in the Description of Application Policies.

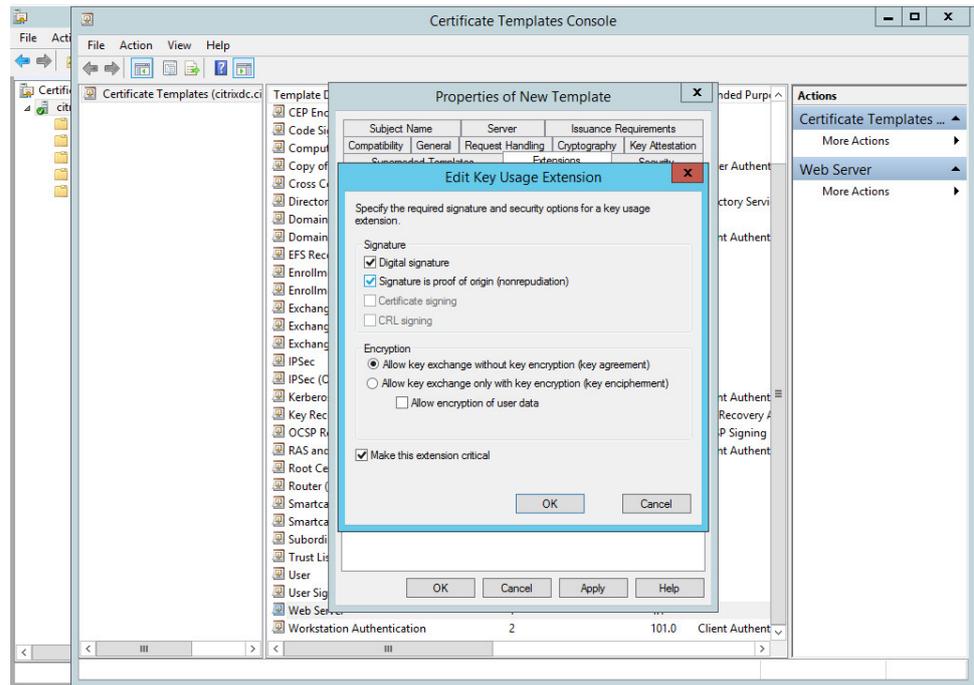
Figure 99: Description of Application Policies



To edit Key Usage

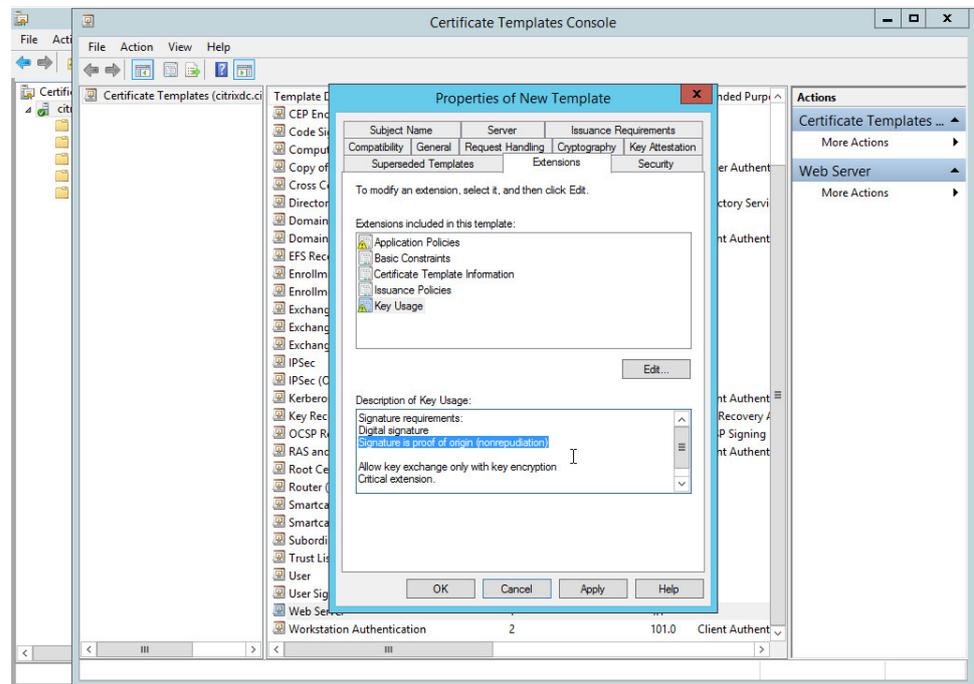
1. Select **Key usage**, click **Edit**. Enable the **Signature is proof of origin (nonrepudiation)** option.
2. Select **Allow key exchange without key encryption (key agreement)**.
3. Click **OK**.

Figure 100: Editing Key Usage



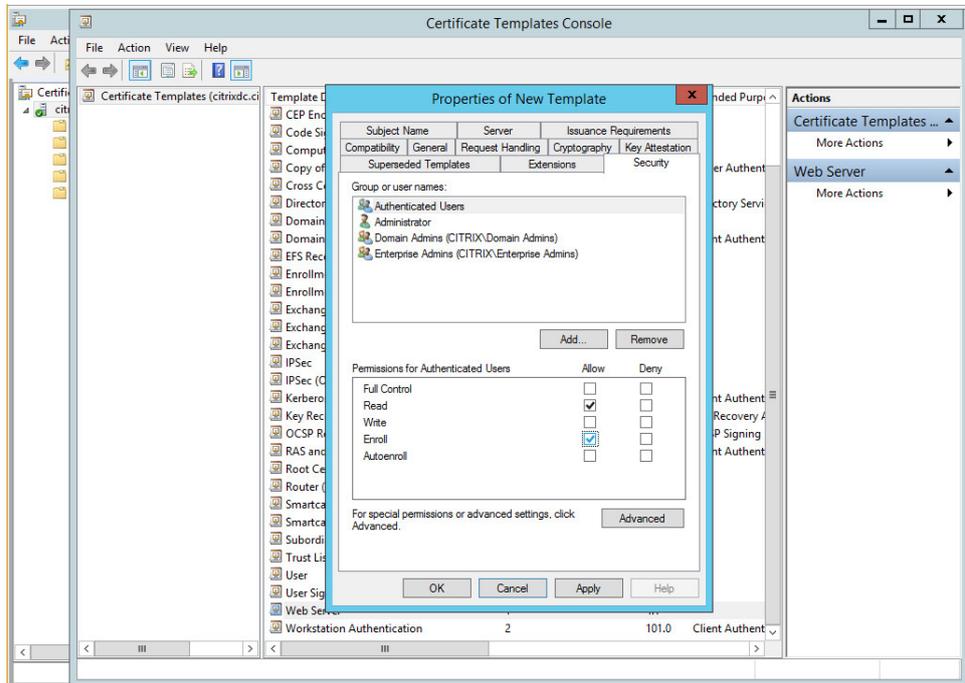
The following are listed in the **Description of Key Usage**.

Figure 101: Description of Key Usage



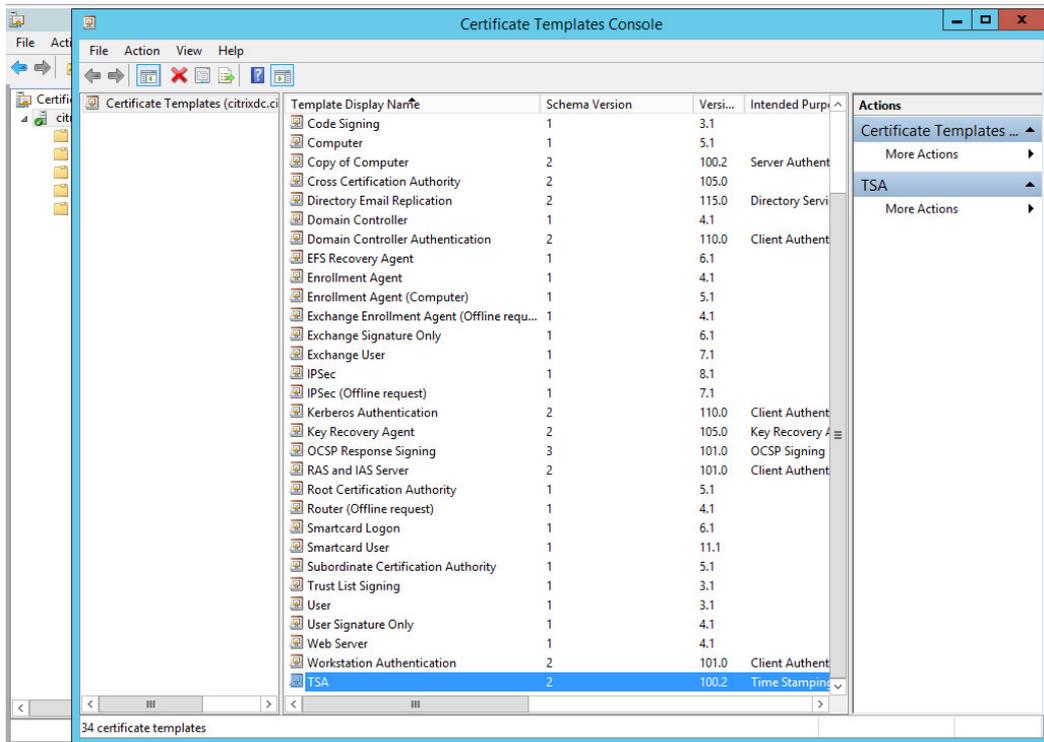
- On the **Security** tab, select **Authenticated Users**, and set **Enroll** to **Allowed**.

Figure 102: Configuring permissions for the template



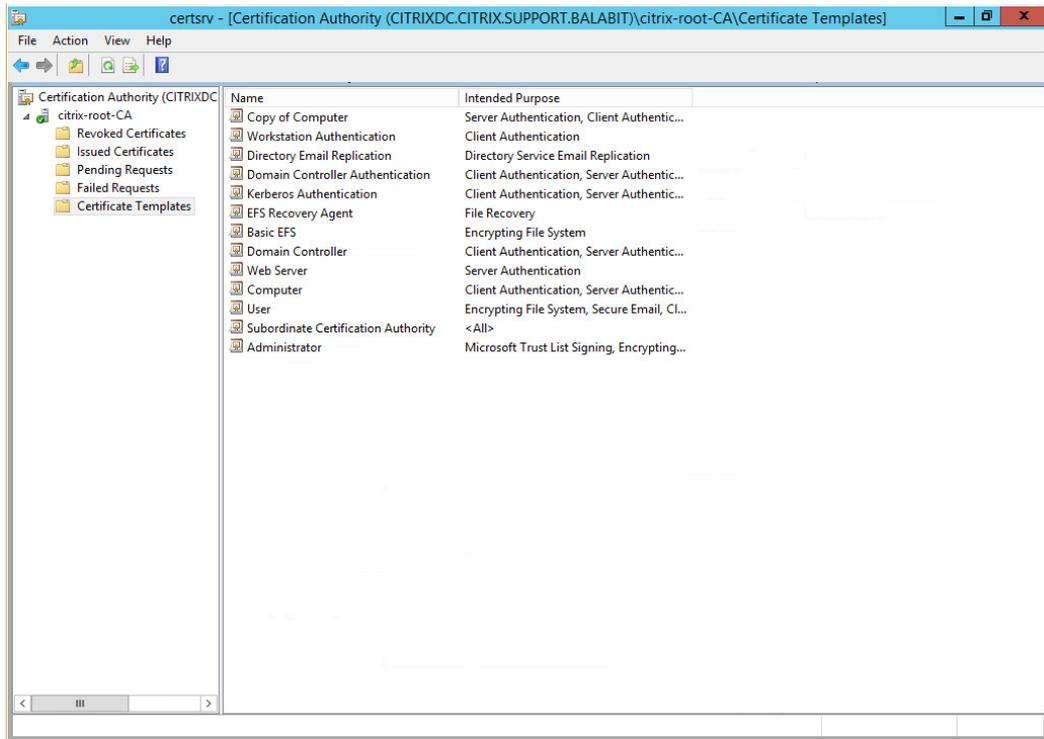
- e. Click **Apply**. Click **OK**. The new TSA template is now displayed in the list of templates.

Figure 103: The new TSA template is now displayed in the list of templates



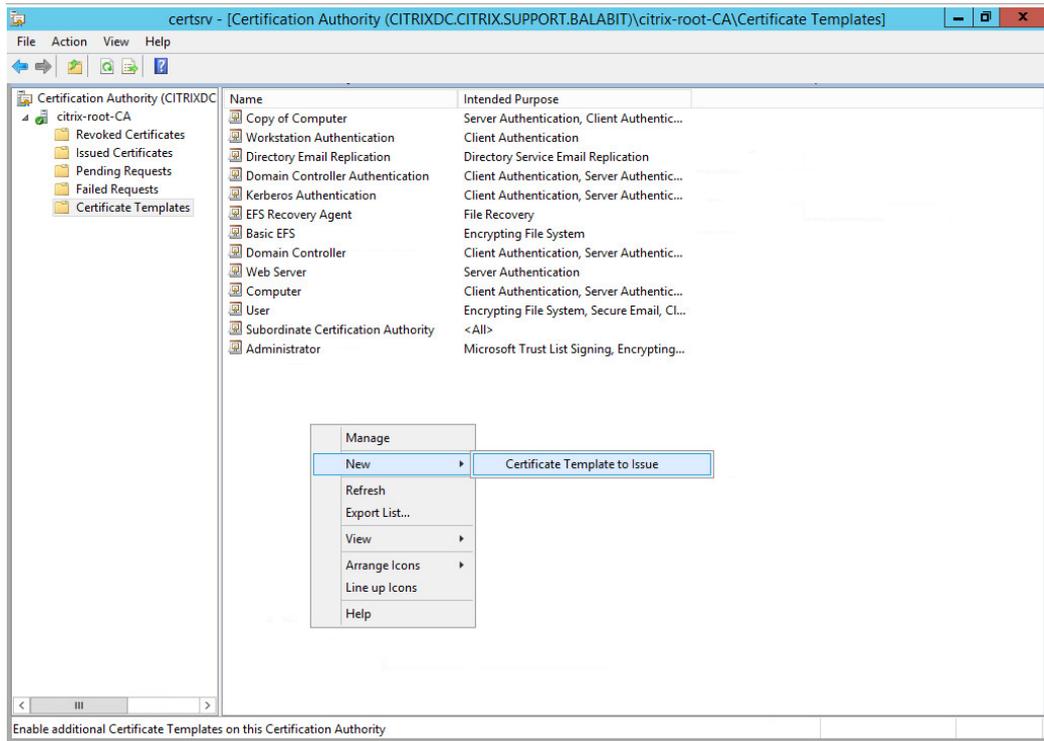
- f. Close this window and return to the Certification Authority main screen, and select the **Certificate Templates** folder.

Figure 104: Certificate Templates



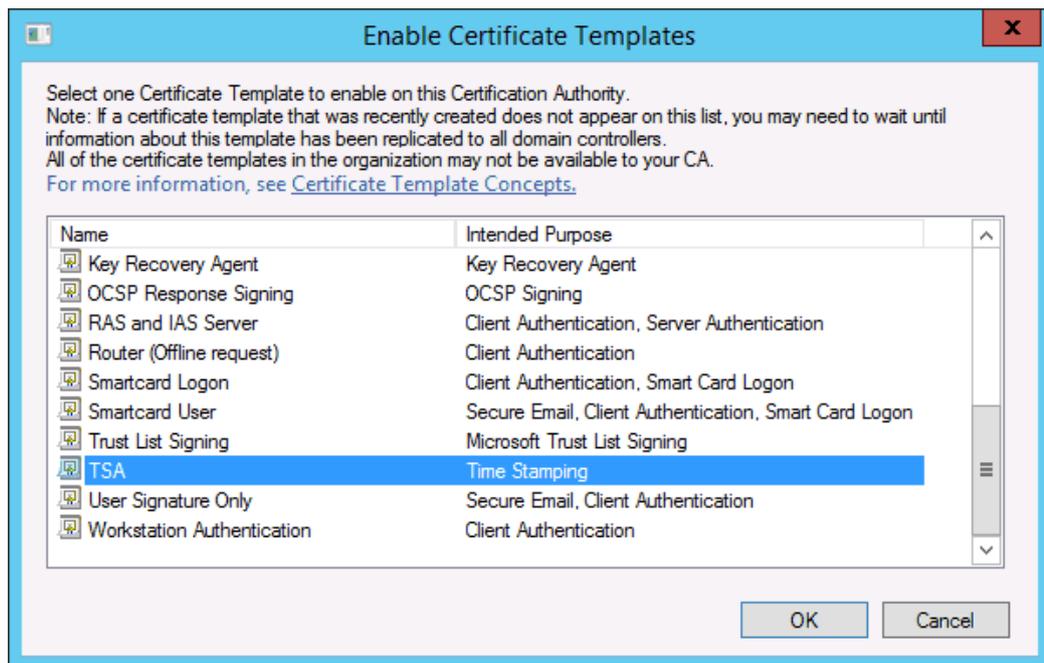
Right-click under the list, and choose **New > Certificate Template to Issue**.

Figure 105: Certificate Template to Issue



The **Enable Certificate Templates** window is displayed.

Figure 106: 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 SSB, navigate to **Basic Settings > Management > SSL certificate**.
 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 certificate**) that is used for other certificates on SSB is different from the CA that was used to sign the TSA certificate, a warning is displayed. In this scenario, ignore this warning.

Creating hostlist policies

The syslog-ng Store Box (SSB) appliance can use a list of host and network addresses at a number of places, for example for limiting the client that can send log messages to a log source, or the hosts that can access shared logspaces.

- For details on how to create a new hostlist, see [Creating hostlists](#) on page 190.
- For details on how to import hostlists from a file, see [Importing hostlists from files](#) on page 191.

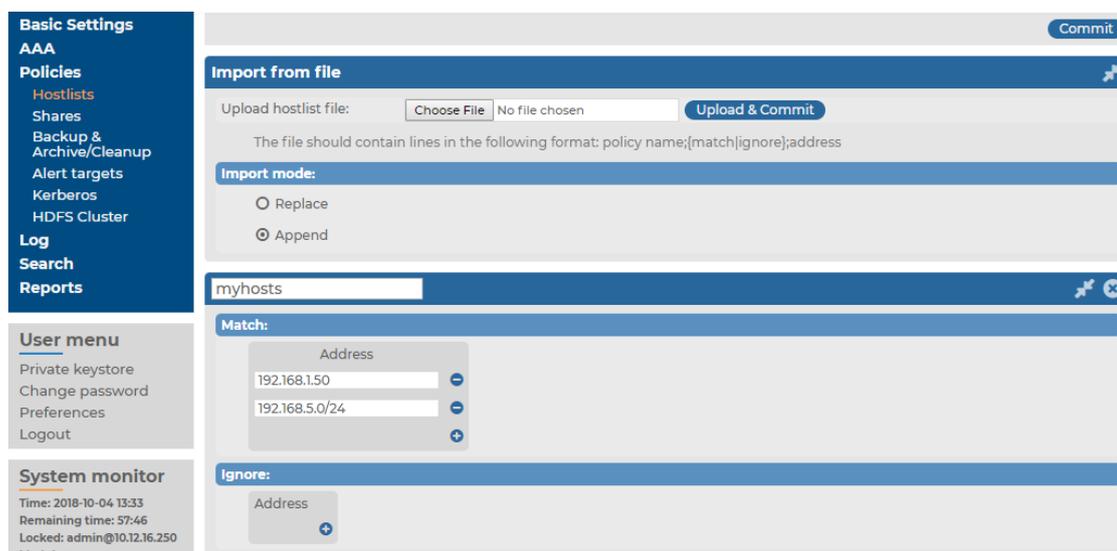
Creating hostlists

This section describes how to create a new hostlist.

To create a new hostlist

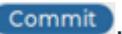
1. Navigate to **Policies > Hostlists** and select .
2. Enter a name for the hostlist (for example servers).

Figure 107: Policies > Hostlists — Creating hostlists



3. Enter the IP address of the permitted host into the **Match > Address** field. You can also enter a network address in the IP address/netmask format (for example 192.168.1.0/24). To add more addresses, click  and repeat this step.
4. To add hosts that are excluded from the list, enter the IP address of the denied host into the **Ignore > Address** field.

TIP: To add every address except for a few specific hosts or networks to the list, add the 0.0.0.0/0 network to the **Match** list, and the denied hosts or networks to the **Ignore** list.

5. Click .

NOTE: If you modify a hostlist, you only need to restart syslog-ng if a host, which is already connected, needs to be ignored with a hostlist. Navigate to **Basic Settings > System > Service control > Syslog traffic, indexing & search:** and select **Restart syslog-ng** for the changes to take effect.

Importing hostlists from files

This section describes how to import hostlists from a text file.

To import hostlists from a text file

1. Create a plain text file containing the hostlist policies and IP addresses to import. Every line of the file will add an IP address or network to a policy. Use the following format:

```
name_of_the_policy;match
```

or

```
ignore;IP address
```

For example, a policy that ignores the 192.168.5.5 IP address and another one that matches on the 10.70.0.0/24 subnet, use:

```
policy1;ignore;192.168.5.5  
policy2;match;10.70.0.0/24
```

To add multiple addresses or subnets to the same policy, list every address or subnet in a separate line, for example:

```
policy1;ignore;192.168.7.5  
policy1;ignore;192.168.5.5  
policy1;match;10.70.0.0/24
```

2. Navigate to **Policies > Hostlists > Import from file > Browse** and select the text file containing the hostlist policies to import.

Figure 108: Policies > Hostlists — Importing hostlists

3. If you are updating existing policies and want to add new addresses to them, select **Append**.

If you are updating existing policies and want to replace the existing addresses with the ones in the text file, select **Replace**.

4. Click **Upload**, then **Commit**.

NOTE: If you modify a hostlist, you only need to restart syslog-ng if a host, which is already connected, needs to be ignored with a hostlist. Navigate to **Basic Settings > System > Service control > Syslog traffic, indexing & search:** and select **Restart syslog-ng** for the changes to take effect.

Configuring message sources

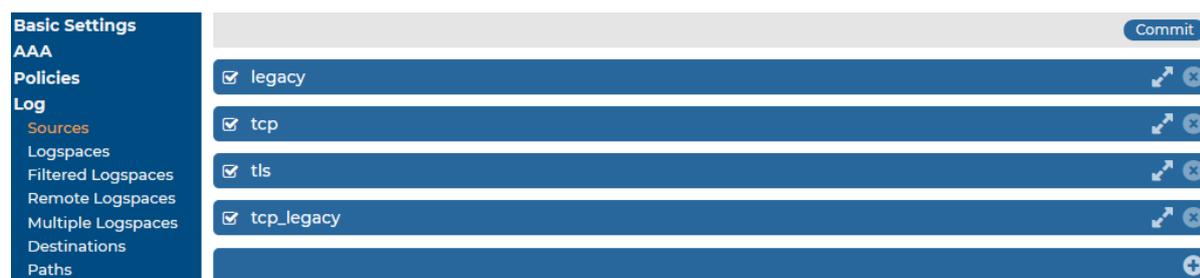
The syslog-ng Store Box (SSB) appliance receives log messages from remote hosts via sources. A number of sources are available by default, but you can also create your own customized message sources. In addition to creating your own, customized message sources based on the Syslog or SQL protocol, SSB can also receive messages via the SNMP protocol, and convert these messages to Syslog messages.

- For details on using the default message sources of SSB, see [Default message sources in SSB](#) on page 193.
- For details on how to create your new, customized message sources based on the Syslog or SQL protocol, see [Creating new message sources in SSB](#) on page 194.
- For details on receiving SNMP messages, see [Receiving SNMP messages](#) on page 214.

Default message sources in SSB

The syslog-ng Store Box (SSB) appliance automatically accepts messages from the following built-in sources:

Figure 109: Log > Sources — Default message sources in SSB



- legacy: Accepts UDP messages using the legacy BSD-syslog protocol on the port 514.
- tcp: Accepts TCP messages using the IETF-syslog protocol (RFC 5424) on port 601.
- tls: Accepts TLS-encrypted messages using the IETF-syslog protocol on port 6514. Mutual authentication is required: the client must show a (not necessarily valid) certificate, SSB sends the certificate created with the Welcome Wizard.

- tcp_legacy: Accepts TCP messages using the BSD-syslog protocol (RFC 3164) on port 514.

NOTE: All default sources have name resolution enabled.

In addition to the default message sources in the previous list, you can also create your own, customized message sources. For the details of the various settings, see [Creating new message sources in SSB](#) and its subsections.

Creating new message sources in SSB

The syslog-ng Store Box (SSB) appliance receives log messages from remote hosts via sources. A number of sources are available by default, but you can also create your own, customized message sources, based on the Syslog or SQL protocol.

For details on using the default message sources of SSB, see [Default message sources in SSB](#) on page 193.

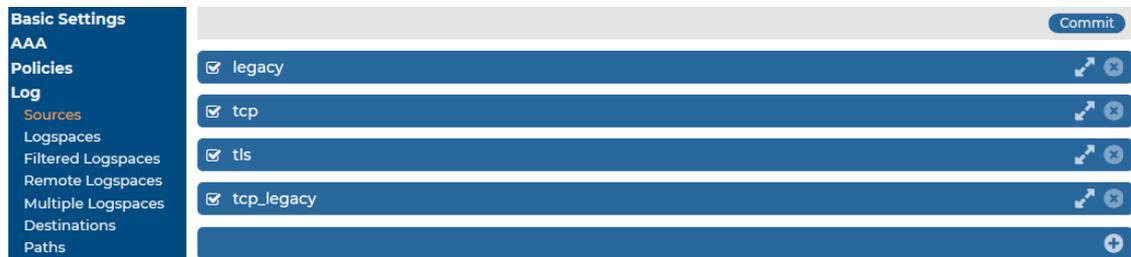
Creating your own, customized message source

If you do not want to use the default message sources available in SSB, you can create your own, customized message source.

To create your own, customized message source

1. Navigate to **Log > Sources** and click .

Figure 110: Log > Sources — Creating new message sources



2. Enter a name for the source into the top field. Use descriptive names (for example, sql_source, or syslog_source) that help you to identify the source easily.

NOTE: In these sections and subsections, some figures show a custom message source named your-new-source, but you can use any descriptive name to identify your message source.

3. In your new source, select your preferred **Source type**.

NOTE: When configuring new message sources in SSB, you can configure two source types: **Syslog**, or **SQL**.

Figure 111: Two available source types under Log > Sources > <your-new-source>



For further details about each source type, see the following subsections:

Configuring your own, customized Syslog type message source

When configuring your own, customized message sources, you can configure two source types: **Syslog**, or **SQL**.

For more information about configuring an SQL source type in your own, customized message source, see [Configuring your own, customized SQL type message source](#).

Figure 112: Log > Sources > <your-new-source> — Syslog source type in your own, customized message source

Source type:

Syslog
 SQL

Listening address: [select an option] ▾
 Listening port: 1433

Transport:

UDP
 TCP
 TLS
 ALTP
 ALTP TLS

Incoming log protocol and message format:

Syslog (IETF-syslog, RFC 5424)
 Legacy (BSD-syslog, RFC 3164)
 Do not parse

This option completely disables syslog message parsing and stores the complete log in the message part.
 It is useful if incoming messages do not comply to the syslog format.

Other source options:

Maximum connections: 1000

Hostname and timestamp-related settings:

Hostlist: None ▾
 Timezone: Europe/Budapest ▾
 Trusted:
 Use FQDN:

Use DNS:

Yes
 No
 Only from persistent configuration

Monitoring:

Message rate alerting:

Alerts: Global settings

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages ▾	30 minutes ▾	1	30	Once ▾	<input type="checkbox"/>

Other options:

Encoding: ▾

Prerequisites for configuring your own, customized Syslog type message source

- The **Listening address** of the Syslog source.
- The **Listening port** of the Syslog source.

Limitations to configuring your own, customized Syslog type message source

- **Listening address IP addresses**

NOTE: The current version (version 6.9.0) of SSB only supports configuring IPv4 addresses for **Source type > Syslog > Listening address** in message sources.

Configuration options for your own, customized Syslog type message source

While configuring your own, customized Syslog type message source, you can customize the following:

- The **Listening address** and **Listening port** of your **Syslog** source type.
- **Transport** protocol options, including **Incoming log protocol and message format** options, and the number of maximum connections under **Other source options**.
- **Hostname and timestamp-related settings**.
- **Monitoring** options.
- Setting the Syslog source type's **Encoding** under **Other options**.

For further details on the configuration options, see the following subsections:

Configuring the Listening address and Listening port for your Syslog type message source

Under **Log > Sources > <your-new-source> > Syslog**, you can first customize the **Listening address** and **Listening port** for your Syslog type message source.

To configure the Listening address and Listening port of your Syslog type message source

1. Navigate to **Log > Sources > <your-new-source> > Source type** and select **Syslog**.

Figure 113: Log > Sources > <your-new-source> > Syslog — Customizing the Listening address and the Listening port for your Syslog type message source

The screenshot shows a configuration window for a Syslog source. At the top, there's a header with a checkmark and the name 'your-new-source'. Below that, a section titled 'Source type:' contains two radio buttons: 'Syslog' (which is selected) and 'SQL'. Underneath, there are two fields: 'Listening address:' with a dropdown menu showing '[select an option]', and 'Listening port:' with a text input field containing the number '1433'.

2. Select the **Listening address** of your choice.

NOTE: Although from version 6.4, syslog-ng Store Box (SSB) supports IPv6 addresses on the external interface and the management interface, you can only select IPv4 addresses for your Syslog sources in the **Listening address** field.

3. Enter the **Listening port** on which you want your log source to listen.

Configuring the Transport options for your Syslog type message source

Under **Log > Sources > <your-new-source> > Syslog > Transport**, you can customize your **Transport** settings for your Syslog type message source.

To customize your Transport settings for your Syslog type message source

1. Navigate to **Log > Sources > <your-new-source> > Syslog > Transport**.

Figure 114: Log > Sources > <your-new-source> > Syslog > Transport — Configuring transport options for your Syslog type message source

The screenshot shows the 'Transport' configuration page. It starts with a 'Transport:' header and five radio buttons: 'UDP', 'TCP' (selected), 'TLS', 'ALTP', and 'ALTP TLS'. Below this is a section titled 'Incoming log protocol and message format:' with three radio buttons: 'Syslog (IETF-syslog, RFC 5424)' (selected), 'Legacy (BSD-syslog, RFC 3164)', and 'Do not parse'. A tooltip points to the 'Do not parse' option, stating: 'This option completely disables syslog message parsing and stores the complete log in the message part. It is useful if incoming messages do not comply to the syslog format.' At the bottom, there's a section 'Other source options:' with a 'Maximum connections:' label and a text input field containing '1000'.

2. In the **Transport** field, select the networking protocol (**UDP, TCP, TLS, ALTP** or **ALTP TLS**) that your clients use to transfer the messages to syslog-ng Store Box (SSB).
3. In case of UDP, TCP or TLS: select the Syslog protocol used by the clients from the **Incoming log protocol and message format** section. The **ALTP** and **ALTP TLS** sources only work with the IETF-syslog protocol.
 - If the clients use the legacy BSD-syslog protocol (RFC3164), select **Legacy (BSD-syslog, RFC3164)**. This protocol is supported by most devices and applications capable to send syslog messages.
 - If the clients use the new IETF-syslog protocol (for example the clients are syslog-ng 3.0 applications that use the `syslog` driver, or other drivers with the `flags(syslog-protocol)` option), select **Syslog (IETF-syslog, RFC 5452)**.

To disable Syslog message parsing and store the complete log in the message part, select **Do not parse**. It is useful if incoming messages do not comply with the Syslog format.

4. When using TLS, SSB displays a certificate to the client. This certificate can be set at **Log > Options > TLS settings** (for details, see [Setting the certificates used in TLS-encrypted log transport](#) on page 325). Optionally, SSB can perform mutual authentication and request and verify the certificate of the remote host (peer). Select the verification method to use from the **Peer verification** field.
 - None: Do not request a certificate from the remote host, and accept any certificate if the host sends one.
 - Optional trusted: If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.
 - Optional untrusted: Accept any certificate shown by the remote host. Note that the host must show a certificate.
 - Required trusted (default setting): Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See [Uploading external certificates to SSB](#) for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
 - Required untrusted: SSB requests a certificate from the remote host, and rejects the connection if no certificate is received. However, SSB accepts the connection if:
 - the certificate is not valid (expired), or
 - the Common Name of the certificate does not contain the domain name or the IP address of the host.

When using ALTP TLS, SSB only accepts Required-trusted peer verification.

| **NOTE:** For details on ALTP, see [Advanced Log Transfer Protocol](#) on page 22.

CAUTION:

UDP is a highly unreliable protocol, when using UDP, a large number of messages may be lost without any warning. Use TCP, TLS or ALTP whenever possible.

5. Configure other, source-related options in the **Other source options** section, depending on what transport you have selected.
 - When using TCP or TLS, you can set the maximum number of parallel connections in the **Maximum connections** field. This option corresponds to the `max_connections()` `syslog-ng` parameter.

In case of ALTP or ALTP TLS: enter the number of maximum connections. The default value is 1000 connections. Select **Allow compression** to allow compression on level 6. Compression level cannot be changed.
 - When using TLS or ALTP TLS, configure the strength of the allowed cipher suites using one of the following options:
 - **Compatible:** It is a large set of cipher suites determined by the following cipher string:

```
ALL: !aNULL: !eNULL
```

The Compatible setting may allow permitting (and hence not safe) cipher suites for the Transport Layer Security (TLS) negotiations.

- **Secure:** A smaller and more strict set of cipher suites where vulnerable cryptographic algorithms are eliminated. This cipher suite set is determined by the following cipher string:

```
HIGH: !COMPLEMENTOFDEFAULT: !aNULL: !eNULL: !DHE-RSA-AES128-SHA: !DHE-RSA-AES256-SHA: !ECDHE-RSA-AES128-SHA: !ECDHE-RSA-AES256-SHA: !AES128-SHA: !AES256-SHA
```

6. (Optional) Customize the number of your **Maximum connections**.

Configuring the Hostname and timestamp-related settings for your Syslog type message source

Under **Log > Sources > <your-new-source> > Syslog > Hostname and timestamp-related settings**, you can first customize your settings related to hostnames and timestamps for your Syslog type message source.

To customize your settings related to hostnames and timestamps for your Syslog type message source

1. Navigate to **Log > Sources > <your-new-source> > Syslog > Hostname and timestamp-related settings**.

Figure 115: Log > Sources > <your-new-source> > Syslog > Hostname and timestamp-related settings — Customizing your settings related to hostnames and timestamps for your Syslog type message source

Hostname and timestamp-related settings:

Hostlist:

Timezone:

Trusted:

Use FQDN:

Use DNS:

Yes

No

Only from persistent configuration

2. In the **Hostname and time stamp related settings** section, configure the following, based on your preferences:
 - To accept messages only from selected hosts, create a hostlist and select it in the **Hostlist** field. For details on creating hostlists, see [Creating hostlist policies](#) on page 189.
 - Set the **Timezone** option of the incoming messages if needed.
 - If the information sent by the hosts to this source can be trusted, enable the **Trusted** option. The syslog-ng Store Box (SSB) appliance keeps the time stamps and the hostname of the messages sent by trusted clients. This corresponds to enabling the `keep_timestamp()` and `keep_hostname()` syslog-ng Premium Edition (syslog-ng PE) options for the source.
 - Select the **Use FQDN** option if you wish to store the full domain name of the sender host.
3. Select the name resolving method to use from the **Use DNS** field.
 - To allow using DNS, select **Yes**.
 - To disable using DNS, select **No**.
 - To only allow using DNS when using a persistent configuration, select **Only from persistent configuration**.

Configuring the Monitoring settings for your Syslog type message source

Under **Log > Sources > <your-new-source> > Syslog > Monitoring**, you can customize your monitoring settings for your Syslog type message source.

To customize your settings related to monitoring for your Syslog type message source

Navigate to **Log > Sources > <your-new-source> > Syslog > Monitoring**.

Log > Sources > <your-new-source> > Syslog > Monitoring.

Figure 116: Log > Sources > <your-new-source> > Syslog > Monitoring — Customizing your monitoring settings your own, customized Syslog type message source

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages	30 minutes	1	30	Once	<input type="checkbox"/>

- 1.
2. (Optional) Enable **Message rate alerting**.
3. Select the basis of your alerts under **Counter**.
4. Select the frequency of alerts (in minutes) under **Period**.
5. Specify the amount of alerts you want to receive within the specified **Period** (ranging between the minimum and maximum numbers of your choice) under **Minimum** and **Maximum**.
6. Select the alerting frequency in the **Alert** field.
Once sends only one alert (and after the problem is fixed, a "Fixed" message).
Always sends an alert each time the result of the measurement falls outside the preset range.
7. (Optional) To set the configured alert settings as your default, enable **Master alert**.
8. (Optional) To leave the **Log > Sources > <your-new-source> > SQL > Monitoring** page and customize **Message rate alerting statistics** settings that apply to the entire syslog-ng Store Box (SSB) appliance, clicking **Global settings** takes you to **Log > Options > Message rate alerting statistics**.

For more information about the configurable settings you can customize under **Log > Options > Message rate alerting statistics**, see [Configuring message rate alerting](#).

NOTE: You can configure multiple alerts under **Monitoring** and pick the alert of your choice as your **Master alert**. To add a new alert under **Message rate alerting**, click .

To delete a redundant alert, click **X**.

1. To configure **Message rate alerting** for the source, see [Configuring message rate alerting](#) on page 91.
2. Set the character **Encoding** option of the incoming messages if needed.
3. Click **Commit**.

NOTE: In order to actually store the messages arriving to this source, you have to include this source in a log path. For details, see [Log paths: routing and processing messages](#) on page 306.

4. (Optional) If you want to receive messages using the ALTP or ALTP TLS protocol, make sure that you have configured your syslog-ng PE clients to transfer the messages to SSB using ALTP or ALTP TPS protocol. For details, see [Advanced Log Transfer Protocol](#).

Customizing encoding for your Syslog type message source

Under **Log > Sources > <your-new-source> > Syslog > Other options**, you can customize your encoding preferences for your Syslog type message source.

To customize your encoding preferences for your Syslog type message source

1. Navigate to **Log > Sources > <your-new-source> > Syslog > Other options**.

Figure 117: Log > Sources > <your-new-source> > Syslog > Other options — customizing your encoding preferences for your Syslog type message source



2. Select the **Encoding** type you want the Syslog source type to use.

Configuring your own, customized SQL type message source

When configuring your own, customized message sources in syslog-ng Store Box (SSB), you can configure two source types: **Syslog**, or **SQL**.

For more information about configuring a Syslog source type in your own, customized message source, see [Configuring your own, customized Syslog type message source](#).

There are many applications that natively store their log messages in SQL databases. The SSB appliance can pull messages from SQL database tables in real-time, similarly to receiving messages over the network.

Figure 118: Log > Sources > <your-new-source> – SQL source type in your own, customized message source

your-new-source
⌵ ⌵

Source type:

Syslog

SQL

Database connection:

Database type: Microsoft SQL Server

Address: 10.12.165.21

Port: 1433 Set Default Port (1433)

Username: test_user

Password: • 🔊

Database name: test_db

Test database connection

Fetching messages:

Create fetch query:

```
select __row_id as uid, date_time as datetime, ssb_sqldest_messages_2020_08.* from
ssb_sqldest_messages_2020_08
```

The query must return message parts with the following column names (note that depending on the SQL server, column names may be case-sensitive):

```
uid
datetime
message
host (optional)
program (optional)
```

Timezone: Europe/Budapest

Facility: Local0

Severity: Info (6)

Put all columns to SDATA:

Fetch history:

After the first successfully fetched message, the SQL source does not use these options anymore, even after restart or reload. It will fetch the messages from the last fetched record.

No history

Partial history

Start UID: 1

The fetch query results will only include records with a strictly larger UID than the Start UID.

Test message fetching

Fetching frequency:

Fetch data every: 10 seconds

Fast follow mode:

Monitoring:

Message rate alerting:

Prerequisites for configuring your own, customized SQL type message source

Currently, configuring your own, customized SQL type message source in SSB has the following prerequisites:

- The **Address** for the SQL type message source (that is, IP address of the remote database server to collect messages from).
- The **Port** of the database server that the SQL type message source will connect to.
- The **Username** of the database user.
- The **Password** of the database user.
- A previously configured SQL database that you can specify the SQL type message source to connect to, by entering the database name into **Database name**.

Limitations to configuring your own, customized SQL type message source

When configuring your message sources, consider that currently, configuring your own, customized SQL type message source in SSB has the following limitations:

- **Database servers**

SSB 6.9.0 was tested with the following database servers:

- MS SQL (with "select @@version")

```
Microsoft SQL Server 2017 (RTM-CU13) (KB4466404) - 14.0.3048.4 (X64)
Nov 30 2018 12:57:58
Copyright (C) 2017 Microsoft Corporation Developer Edition (64-bit)
on Linux (Ubuntu 16.04.5 LTS)
```

- Oracle (with "SELECT * FROM V\$VERSION;")

```
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 : 64bit
Production
PL/SQL Release 11.2.0.4.0 Production : "CORE 11.2.0.4.0 Production"
TNS for Linux: Version 11.2.0.4.0 Production
NLSRTL : Version 11.2.0.4.0 Production
```

- **Remote server IP addresses**

NOTE: The current version (version 6.9.0) of SSB only supports configuring IPv4 addresses for **Source type > SQL > Database connection > Address** in message sources.

Configuration options for your own, customized SQL type message source

While configuring your own, customized SQL type message source, you can customize the following:

- Configuring and testing the **Database connection** of your SQL source.
- Customizing your fetch queries and message fetching settings under **Fetching messages**.
- Customizing your fetch history and test message fetching under **Fetch history**.
- Customizing your **Fetching frequency**.
- Customizing your **Monitoring** options.

For further details on the configuration options, see the following subsections:

Setting up and testing the SQL database connection of your SQL type message source

The following section describes how to setup the SQL database connection of your SQL type message source, and how the connection is tested.

To setup and test the SQL database connection in your SQL type message source

1. Navigate to **Log > Sources > <your-new-source> SQL > Database connection**.

Figure 119: Log > Sources > <your-new-source> > SQL > Database connection – Setting up and testing the SQL database connection for your SQL type message source

The screenshot shows a configuration form titled "Database connection:". It contains the following fields and controls:

- Database type:** A dropdown menu set to "Microsoft SQL Server".
- Address:** A text input field containing "10.12.165.21".
- Port:** A text input field containing "1433", with a blue button labeled "Set Default Port (1433)" to its right.
- Username:** A text input field containing "test_user".
- Password:** A text input field with a masked password (dots) and a small eye icon to toggle visibility.
- Database name:** A text input field containing "test_db".
- Test database connection:** A blue button at the bottom of the form.

2. Select the **Database type** to collect log messages from.
3. Enter the name or the IP address of the remote database server to collect messages from.
4. Enter the port of the database server to connect to. To use the default port of the database, click **Set Default Port**.

NOTE: The default port depends on the database type, but you can configure it according to your preferences.

5. Enter the **Username** and the **Password** of the database user.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

6. Enter the **Database name** to connect to.

7. Click **Test database connection**. SSB reads the tables from the database.

NOTE: The syslog-ng Store Box (SSB) appliance can only read table names that contain numbers, uppercase and lowercase characters, hyphen (-), underscore (_), hashtag (#), at sign (@), or the dollar sign (\$). Tables with names that contain other characters, including full stop (.), cannot be monitored.

Customizing fetching messages when using your SQL type message source

For a more flexible SQL source configuration, you can customize your message fetching preferences (including customizing your fetch query and fetch history) for your SQL type message source.

Figure 120: Log > Sources > <your-new-source> > SQL > Fetching messages — Customizing fetching messages for your SQL type message source

Fetching messages:

Create fetch query:

```
select __row_id as uid, date_time as datetime, ssb_sqldest_messages_2020_08.* from ssb_sqldest_messages_2020_08
```

The query must return message parts with the following column names (note that depending on the SQL server, column names may be case-sensitive):

- uid
- datetime
- message
- host (optional)
- program (optional)

Timezone: Europe/Budapest

Facility: Local0

Severity: Info (6)

Put all columns to SDATA:

Fetch history:

After the first successfully fetched message, the SQL source does not use these options anymore, even after restart or reload. It will fetch the messages from the last fetched record.

No history

Partial history

Start UID: 1

The fetch query results will only include records with a strictly larger UID than the Start UID.

Test message fetching

For more information about creating your customized fetch query and customizing your fetch history, see the following subsections:

Customizing the fetch query for your SQL type message source

Under **Log > Sources > <your-new-source> > SQL > Fetching messages**, you can create a customized fetch query for your SQL type message source.

To create a customized fetch query for your SQL type message source

1. Navigate to **Log > Sources > <your-new-source> > SQL > Fetching messages**.

Figure 121: Log > Sources > <your-new-source> > SQL > Fetching messages — Customizing the fetch query for your SQL type message source

Fetching messages:

Create fetch query:

```
select __row_id as uid, date_time as datetime, message from ssb_sqldest_messages_2020_08
```

The query must return message parts with the following column names (note that depending on the SQL server, column names may be case-sensitive):

- uid
- datetime
- message
- host (optional)
- program (optional)

Timezone: Europe/Budapest

Facility: Local0

Severity: Info (6)

Put all columns to SDATA:

2. Enter your customized fetch query in the suitable format.
For more information, see [Limitations, considerations, and best practices](#).
3. Select the **Timezone** of your choice.
4. Select the **Facility** level of your choice.
5. Select the **Severity** level of your choice.
6. (Optional) Enable **Put all columns to SDATA**.

Limitations, considerations, and best practices

Before you configure your customized fetch query, consider the limitations, considerations, and best practices in the current version (version 6.9.0) of syslog-ng Store Box (SSB):

- **Required column names**

The custom fetch query has the following required column names it must return:

- **uid:**
The **uid** column must contain a unique number. This number must increase monotonously. SSB will store the last read **uid**.
| NOTE: SSB automatically substitutes the last read **uid** after the initial fetch.
- **datetime or date and time:**

SSB will use the content of the datetime column as the time stamp of the log message. The following column types are supported:

- *Oracle*: timestamp, int
- *MSSQL*: datetime, int

If the type is int, SSB will assume that it contains a UNIX time stamp.

When using separate date and time columns, the date column must be date type, the time column must be time type.

- message:

The message field must contain the message to be logged.

For example:

```
SELECT "uniq_id_number" AS "uid", "bsd_datetime" AS "datetime",  
"message" AS "message" FROM "test_table"
```

• Optional column names

The custom fetch query may contain the following optional column names:

- host (optional):
Optional column name. When in use, host overwrites the host field of the incoming message.
- program (optional):
Optional column name. When in use, program overwrites the program field of the incoming message.

• Considerations and best practices

NOTE: The query must not contain any comments.

NOTE: The host, program, and timezone parameters can be selected from columns or set as a fix value. The timezone must contain time-shifting value and not the name of the time zone. For example:

```
SELECT "myhost" AS "host", "myprogram" AS "program", "+01:00" AS  
"timezone", <further-parts-of-the-query>
```

⚠ CAUTION:

The syslog-ng Store Box (SSB) appliance does not validate or limit the contents of customized queries. Consequently, queries performed with a user with write-access can potentially modify or even harm the database. Use customized queries with care, and only for your own responsibility.

Examples: creating a customized fetch query

Example: SQL source fetch_query

The following custom fetch query queries records that are older than the last read record:

```
SELECT * FROM <table_name>
```

Or a more detailed example:

```
SELECT "uniq_id_number" AS "uid", "date_string" AS "datetime", "message"  
AS "message" FROM "test_mysql"
```

Customizing the fetch history settings for your SQL type message source

After [customizing your fetch query](#) settings for your SQL type message source, you can customize your fetch history settings as well.

NOTE: Customizing your fetch history has relevance when you are configuring your SQL type message source for the first time. After the initial message history fetching, the message source does not use these settings, but fetches the messages from the last fetched record instead.

To customize the fetch history settings for your SQL type message source

1. Navigate to **Log > Source > <your-new-source> > SQL > Fetch history**.

Figure 122: Log > Source > <your-new-source> > SQL > Fetch history – Customizing the fetch history settings for your SQL type message source

Fetch history:

After the first successfully fetched message, the SQL source does not use these options anymore, even after restart or reload. It will fetch the messages from the last fetched record.

No history

Partial history

Start UID:

The fetch query results will only include records with a strictly larger UID than the Start UID.

[Test message fetching](#)

2. Select how extensive you want your fetch history.

- To skip the initial message fetching from the database, select **No history**.
- To specify from which particular UID you want the message source to fetch messages from your database, select **Partial history**, and specify the **Start UID** of your preference.

⚠ CAUTION:

Consider that the message source only fetches messages with a UID that is strictly larger than the specified Start UID. As a result, the message with the UID specified as the Start UID will not be fetched from the database.

3. To quickly check your fetch history settings before committing your message source configuration, click **Test message fetching**. The message source will fetch a sample of the database, and you can check if your settings are right and if the fetched data match your database and your desired configuration and settings.

Configuring the fetching frequency settings for your SQL type message source

Under **Log > Sources > <your-new-source> > SQL > Fetching frequency**, you can customize your fetching frequency preferences for your SQL type message source.

To customize your fetching frequency preferences for your SQL type message source

1. Navigate to **Log > Sources > <your-new-source> > SQL > Fetching frequency**.

Figure 123: Log > Sources > <your-new-source> > SQL > Fetching frequency — Customizing your fetching frequency preferences for your SQL type message source

2. Specify how often you want syslog-ng Store Box (SSB) to periodically fetch data by entering the number of your choice (in seconds) in **Fetch data every:** <number> **seconds**.
3. (Optional) To allow SSB to read the database as fast as possible, enable **Fast follow mode**.

NOTE: SSB reads the database periodically, performing one query (that fetches multiple records) at a time. With **Fast follow mode** enabled, SSB continuously keeps querying the database until it fetches all records available at the time.

Configuring the Monitoring settings for your SQL type message source

Under **Log > Sources > <your-new-source> > SQL > Monitoring** you can customize your monitoring settings, including enabling **Message rate alerting** and customizing your **Alerts**, for your SQL type message source.

To customize the Monitoring settings for your SQL type message source

1. Navigate to **Log > Sources > <your-new-source> > SQL > Monitoring**.

Figure 124: Log > Sources > <your-new-source> > SQL > Monitoring – Configuring the Monitoring settings for your SQL type message source

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages	30 minutes	1	30	Once	<input type="checkbox"/>

2. (Optional) Enable **Message rate alerting**.
3. Select the basis of your alerts under **Counter**.
4. Select the frequency of alerts (in minutes) under **Period**.
5. Specify the amount of alerts you want to receive within the specified **Period** (ranging between the minimum and maximum numbers of your choice) under **Minimum** and **Maximum**.
6. Select the alerting frequency in the **Alert** field.
Once sends only one alert (and after the problem is fixed, a "Fixed" message).
Always sends an alert each time the result of the measurement falls outside the preset range.
7. (Optional) To set the configured alert settings as your default, enable **Master alert**.
8. (Optional) To leave the **Log > Sources > <your-new-source> > SQL > Monitoring** page and customize **Message rate alerting statistics** settings that apply to the entire syslog-ng Store Box (SSB) appliance, clicking **Global settings** takes you to **Log > Options > Message rate alerting statistics**.

For more information about the configurable settings you can customize under **Log > Options > Message rate alerting statistics**, see [Configuring message rate alerting](#).

NOTE: You can configure multiple alerts under **Monitoring** and pick the alert of your choice as your **Master alert**. To add a new alert under **Message rate alerting**, click . To delete a redundant alert, click .

Receiving SNMP messages

The syslog-ng Store Box (SSB) appliance can receive SNMP messages using the SNMPv2c protocol and convert these messages to syslog messages. SNMP messages are received using a special SNMP source that can be used in log paths like any other source. The following describes how to configure receiving SNMP messages.

To configure receiving SNMP messages

1. Navigate to **Log > Options > SNMP source**.
2. Ensure that the **SNMP source** option is enabled.

Figure 125: Log > Options > SNMP source — Receiving SNMP messages

The screenshot shows the configuration page for the SNMP source. The left sidebar contains a navigation menu with categories: Basic Settings, AAA, Policies, Log, Search, and Reports. Under 'Log', there are sub-items: Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options (highlighted), and Pattern Database. Under 'Search', there are: Private keystore, Change password, Preferences, and Logout. The main content area has a 'Commit' button at the top right. Below it are several configuration sections: 'Options', 'TLS settings', 'SNMP source' (checked), 'Alerting' (checked), 'Artificial ignorance' (checked), 'Name resolving', 'Dashboard statistics' (checked), and 'Message rate alerting statistics'. The 'SNMP source' section includes: 'Use DNS' (checked), 'Community' (text input with 'public'), 'Hostlist' (dropdown menu), and 'Rate limit' (text input with 'messages/sec').

3. The default community of the SNMP messages is `public`. Modify the **Community** field if your hosts use a different community.
NOTE: SSB can receive messages only from a single community.
4. To limit which hosts can send SNMP messages to SSB, create a hostlist policy, add the permitted hosts to the policy, and select the policy from the **Hostlist** field. For details on creating hostlists, see [Creating hostlist policies](#) on page 189.
5. To limit the rate of messages a host can send to SSB, enter the maximum number of packets (not messages) that SSB is allowed to accept from a single host into the **Rate limit** field. (This parameter sets the `hashlimit` parameter of the **iptables** packet filter that is applied to the source.)

CAUTION:

When rate limiting is enabled, and a host sends a large number of messages, SSB processes only the amount set in the Rate limit field. Any additional messages are dropped, and most probably lost.

6. To use name resolution for SNMP messages, enable the **Use DNS** option.
7. Click **Commit**.

Storing messages on SSB

The syslog-ng Store Box (SSB) appliance stores log messages in binary or plain text log files called logspaces. You can define multiple logspaces, remote logspaces, and configure filtered subsets of each logspace.

Binary log files (logstores) correspond to the encrypted `logstore()` destination of syslog-ng. Logstores can be compressed, encrypted, and time stamped by an external Time Stamping Authority (TSA). To make the contents of the logstore searchable, you can create a separate indexer configuration for each logstore.

A multiple logspace aggregates messages from multiple SSBs (located at different sites), allowing you to view and search the logs of several SSBs from a single web interface without having to log on to several different interfaces.

Remote logspaces enable you to access and search logspaces (including filtered logspaces) on other SSB appliances.

Filtered logspaces allow you to create a smaller, filtered subset of the logs contained in an existing local, remote or multiple logspace. Assigning a user group to a filtered logspace enables fine-grained access control by creating a group that sees only a subset of the logs from a logspace.

[Summary of multiple, remote, and filtered logspace types](#) on page 215 provides a summary and comparison of these three logspace types.

Table 7: Summary of multiple, remote, and filtered logspace types

Logspace type	Source	Main use case	Can be searched	Can be filtered
Multiple	Multiple SSBs located at different sites	Aggregate messages from multiple logspaces into a single logspace Pre-filter log messages and share with only select user groups	✓	✓
Remote	Remote SSB	Access a logspace on another SSB	✓	✓
Filtered	Local / multiple / remote SSB (s)	Control access to a logspace at a granular level by granting access only to a subset of a logspace	✓	N/A

By default, SSB has the following logspaces:

SSB

6.9.0

215

Figure 126: Log > Logspaces — Default logspaces in SSB



- local: An unencrypted, binary logspace for storing the log messages of SSB.
- center: An unencrypted, binary logspace for storing the log messages sent by the clients.

Logspaces are stored locally on the hard disk of SSB. To access a logspace remotely, you can configure another SSB to view and search the logspace as a remote logspace, or you can make the logspace accessible as a network drive.

- For information on using encrypted log files (logstores), see [Using logstores](#) on page 216.
- For details on creating plain-text logspaces, see [Creating text logspaces](#) on page 222.
- For details on managing logspaces, see [Managing logspaces](#) on page 226.
- For details on creating filtered logspaces, see [Creating filtered logspaces](#) on page 233.
- For details on creating remote logspaces, see [Creating remote logspaces](#) on page 235.
- For details on creating multiple logspaces, see [Creating multiple logspaces](#) on page 236.
- For details on making the log files accessible remotely as a network drive, see [Accessing log files across the network](#) on page 239.

Using logstores

Logstores are logspaces with binary log files for storing log messages sent by the clients. Logstores can be compressed, encrypted, and time stamped by an external Time Stamping Authority (TSA). To make the contents of the logstore searchable, you can create a separate indexer configuration for each logstore.

The following limitations apply to logstores:

- Indexing logstore files is currently limited: the indexer can handle only one file from a logstore for every day (syslog-ng Store Box (SSB) automatically starts a new log file for every day).
- Logstore files consist of chunks. In rare cases, if the syslog-ng application running on SSB crashes for some reason, it is possible that a chunk becomes broken: it contains log messages, but the chunk was not finished completely. However, starting with SSB version 2 F1 the syslog-ng application running on SSB processes log messages into a journal file before writing them to the logstore file, reducing message loss even

in the case of an unexpected crash.

Similarly, if the indexer application crashes for some reason, it may be possible that some parts of a logstore file are not indexed, and therefore the messages from this part of the file do not appear in search results. This does not mean that the messages are lost. Currently it is not possible to reindex a file.

These limitations will be addressed in future versions of SSB.

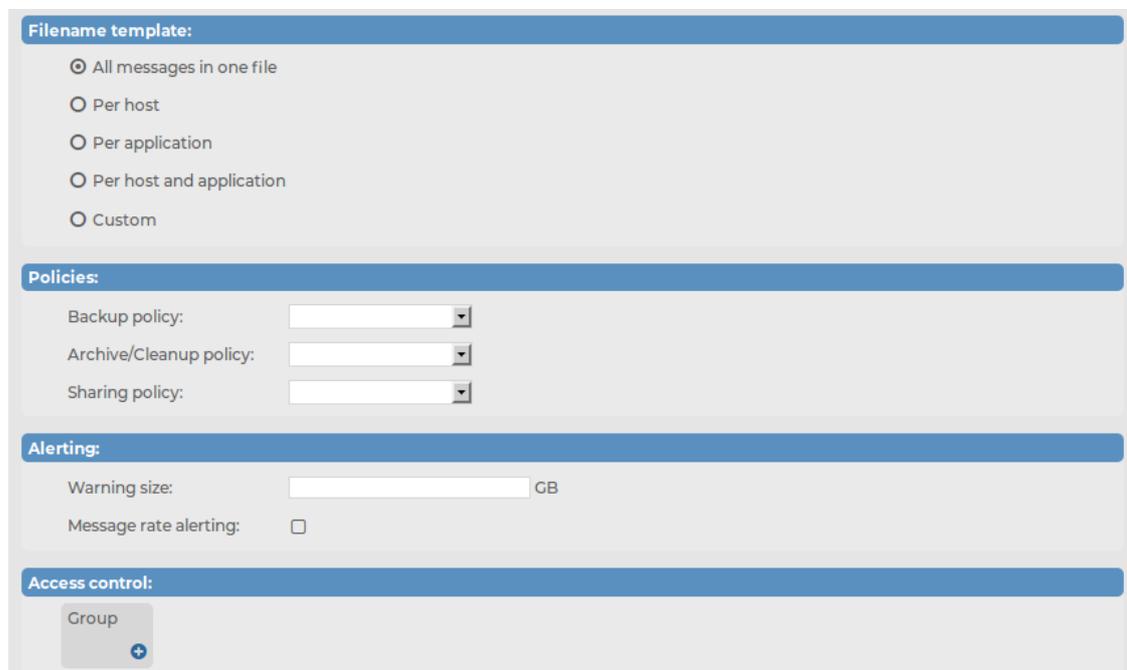
- For details on how to create logstores, see [Creating logstores](#) on page 217.
- For details on configuring indexing for logstores, see [Configuring the indexer service](#) on page 220.
- For details on displaying the contents of a logstore file, including encrypted logs, see [Viewing encrypted logs with logcat](#) on page 222.

Creating logstores

To create logstores

1. Navigate to **Log > Logspaces** and click .
2. Enter a name for the logspace into the top field. Use descriptive names that help you to identify the source easily. Note that the name of the logspace must begin with a number or a letter.

Figure 127: Log > Logspaces — Creating a new logstore



Filename template:

- All messages in one file
- Per host
- Per application
- Per host and application
- Custom

Policies:

Backup policy:

Archive/Cleanup policy:

Sharing policy:

Alerting:

Warning size: GB

Message rate alerting:

Access control:

Group

The screenshot displays the Syslog-ng Store Box (SSB) configuration page for a new logstore named 'newlogstore'. The interface is divided into several sections:

- Basic Settings:** Includes buttons for Backup ALL, Restore ALL, Archive/cleanup ALL, Empty ALL, and Commit.
- Logstore Name:** The name 'newlogstore' is entered in the input field.
- Type:** The 'LogStore' radio button is selected, with options for 'Text file'.
- Encryption:** The 'Compressed logstore' checkbox is checked. The 'Encryption certificate' field has a browse icon (🔑).
- Indexer:** The 'Enabled' radio button is selected. The 'Maximum number of search results' is set to 50000000, and the 'Memory limit' is 1024 MB.
- Indexed fields:** A list of fields with checkboxes: Facility (unchecked), Priority (unchecked), Program (checked), Pid (unchecked), Host (checked), Tags (unchecked), Name/value pairs (checked), Message (checked), and Delimiters (set to ';&~?![]=,;()').
- System monitor:** A sidebar widget showing system status: Time: 2018-02-18 19:44, Remaining time: 01:44, Locked: admin@10.30.255.62, Modules: syslog-ng: Running, Active: Hosts: 3, Senders: 2, Load 1: 0.03, Load 15: 0.00. Below this is a bar chart for CPU, Mem, Disk, and Swap usage.

3. Select **LogStore** from the **Type** field.
4. To encrypt the log files using public-key encryption, click  in the **Encryption certificate** field.

A pop-up window is displayed.

Click **Browse**, select the certificate you want to use to encrypt the log files, then click **Upload**. Alternatively, you can paste the certificate into the **Certificate** field and click **Upload**.

NOTE: To view encrypted log messages, you will need the private key of this certificate. For details on browsing encrypted logstores online on the syslog-ng Store Box (SSB) web interface, see [Browsing encrypted logspaces](#) on page 345. Encrypted log files can be displayed using the **logcat** command-line tool as well. The **logcat** application is currently available only for UNIX-based systems.

One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
- Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.

NOTE: Each certificate or encryption-related setting described above only takes effect from the next day.

However, if you use decryption private keys, you can search in the encrypted logstores immediately after the private keys are uploaded. For more information, see [Assigning decryption keys to a logstore](#) on page 349.

5. By default, SSB requests a time stamp every ten minutes from the internal Time Stamping Authority. Adjust the frequency of time stamping requests in the **Timestamping frequency** field if needed. For details on how to request time stamps from an external provider, see [Time stamping configuration on SSB](#) on page 323.
6. Indexing is enabled by default. For detailed instructions on configuring indexing, see [Configuring the indexer service](#) on page 220.
7. Logstore files are compressed by default. If you do not want to use compression, uncheck the **Compressed logstore** option.
8. Select how to organize the log files of this logspace from the **Filename template** field.
 - To save every message received during a day into a single file, select **All messages in one file**.
 - To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `${HOST}` macro of syslog-ng.
 - To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `${PROGRAM}` macro of syslog-ng.
 - To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `${HOST}-${PROGRAM}` macros of syslog-ng.
 - To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field.

NOTE:

For details on using filename templates, see .

Templates that generate an invalid path (for example, they use a filename longer than 246 characters or refer to a parent directory) will not work.

9. To create automatic daily backups of the logspace to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see [Data and configuration backups](#).
10. To archive the logspace automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies,

see [Archiving and cleanup](#).

⚠ CAUTION:

Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.

11. To make the log files of this logspace available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see [Accessing log files across the network](#).
12. Set a size for the logspace in the **Warning size** field: SSB will send an alert if the size of this logspace exceeds the limit.

⚠ CAUTION:

Make sure that the Logspace exceeded warning size alert is enabled in Basic Settings > Alerting > syslog-ng traps, and that the mail settings of Basic Settings > Management, and the SNMP settings of Basic Settings > Alerting are correct. Otherwise, you will not receive any alert when the logspace exceeds the size limit. For details on alerting and monitoring, see also [Configuring system monitoring on SSB](#).

13. By default, members of the search group can view the stored messages online. Use the **Access control** option to control which usergroups can access the logspace. For details, see also [Managing user rights and usergroups](#).
14. Click **Commit**.

Configuring the indexer service

The indexer service saves the indexes for the fields that are selected and makes them searchable. Indexing fields consumes disk space and processing power.

This section lists the limitations of the indexer service, and provides instructions for configuring indexing for logstores.

Limitations:

- Messages are tokenized based on the specified separator characters. Only the first 512 tokens are indexed in a message, the rest are ignored. This limitation does not affect other static fields (PROGRAM, HOST, and so on) or name-value pairs added by the pattern database or values coming from the SDATA part of incoming messages.
- Whitespace characters (space, tabulator and so on) are always treated as delimiters.
- Tokens that are shorter than 2 characters are not indexed.
- Tokens are truncated to 59 characters. Therefore, tokens with at least 59 characters long common prefix will be handled as identical ones.
- When indexing name-value pairs, the 59 characters limitation is applied to this format: "<name-of-nvpair>=<value-of-nvpair>". Do not use long name parts, in

order to avoid the premature truncation of the value part.

- The shortest timeframe for searching and creating statistics is 1 second. Smaller interval cannot be used.
- The order of the tokens in a message is not preserved. Therefore, if one message contains 'first_token second_token' and another message contains 'second_token first_token' search expressions such as 'first_token second_token' will find both messages.

To configure the indexer service

1. Navigate to **Log > Logspaces** and select the logstore to index.
2. To enable automatic indexing of the logstore files, select the **Enable** option of the **Indexer** field.
3. To limit the number of hits when searching in the logstore, enter the maximum number of search result hits in the **Maximum number of search results** field.
To disable the limit, enter 0.
4. Enter the maximum amount of memory the indexer can use for the current logspace in the **Memory limit** field.

⚠ CAUTION:

Hazard of data loss. Increasing the Memory limit option too high (1280 MB) can cause message loss and degraded performance. The exact values that can cause problems depend on your configuration and environment.

Make sure that the sum of the memory of all indexed logspaces is smaller than the available memory in your syslog-ng Store Box (SSB) appliance.

5. Configure the fields to be indexed in the **Indexed fields**.

| **NOTE:** At least one field must be selected.

The following fields can be indexed: **Facility, Priority, Program, Pid, Host, Tags, Name/value pairs, Message**.

For the **Name/value pairs** field, select **All** to index all Name/value fields or enter the names to be indexed in the **Only with the name** field as comma-separated names.

If the indexing of the **Message** field is enabled, the current **Delimiters** are displayed. By default, the indexer uses the following delimiter characters to separate the message into words (tokens): **& : ~ ? ! [] = , ; () ' "**

If your messages contain segments that include one of these delimiters, and you want to search for these segments as a whole, remove the delimiter from the list. For example, if your log messages contain MAC addresses, and you want to be able to search for messages that contain a particular MAC address, delete the colon (:) character from the list of delimiters. Otherwise, the indexer will separate the MAC address into several tokens.

NOTE: It is not possible to search for the whitespace () character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

Viewing encrypted logs with logcat

To access logstore files, you can:

- Access the logstores using a network share.
This is the recommended method. For details, see [Accessing log files across the network](#) on page 239.
- Log in to syslog-ng Store Box (SSB) locally, or remotely using SSH.

To display the contents of a logstore file, use the **logcat** command supplied with syslog-ng. For example:

```
logcat /var/log/messages.lgs
```

To display the contents of encrypted log files, specify the private key of the certificate used to encrypt the file. For example:

```
logcat -k private.key /var/log/messages.lgs
```

The contents of the file are sent to the standard output, so it is possible to use grep and other tools to find particular log messages. For example:

```
logcat /var/log/messages.lgs |grep 192.168.1.1
```

Every record that is stored in the logstore has a unique record ID. The **logcat** application can quickly jump to a specified record using the **-- seek** option.

For files that are in use by syslog-ng, the last chunk that is open cannot be read. Chunks are closed when their size reaches the limit set in the `chunk_size` parameter, or when the time limit set in the `chunk_time` parameter expires and no new message arrives.

When the logstore file is encrypted, a hash is also generated for every chunk to verify the integrity of the chunk. The hashes of the chunks are chained together to prevent injecting chunks into the logstore file. The encryption algorithm used is aes128 in CBC mode, the hashing (HMAC) algorithm is hmac-sha1.



CAUTION:

If the syslog-ng Premium Edition application or the computer crashes, an unclosed chunk remains at the end of the file. This chunk is marked as broken, its data stays there but is not shown by logcat.

Creating text logspaces

This section describes how to create a new logspace that stores messages in plain text files.

⚠ CAUTION:

Compared to binary logspaces (LogStore files), plain text logspaces have the following limitations.

- Plain text logspaces are not indexed, and you cannot browse or search them on the syslog-ng Store Box (SSB) search interface.
- You cannot create remote, filtered, or multiple logspaces using text logspaces.
- You cannot access text logspaces using the SSB RPC API.

Use text logspaces only if you want to access them as a shared file from an external application. For details, see [Accessing log files across the network on page 239](#).

You can also configure SSB to store the messages in a plain text logspace (so you can share it) and in a LogStore file at the same time, so you can access them from the SSB search interface. To accomplish this, configure a log path that has two destinations (one plain text, one LogStore), and disable the Log > Paths > Final option for the first path.

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this as member**. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

To create a new logspace that stores messages in plain text files

1. Navigate to **Log > Logspaces** and click .
2. Enter a name for the logspace into the top field. Use descriptive names that help you to identify the source easily.

Figure 128: Log > Logspaces — Creating a new text logspace



3. Select **Text file** from the **Type** field.
4. Select the template to use for parsing the log messages. The following templates are available:

- Legacy corresponds to the following syslog-ng template:

```
template("${DATE} ${HOST} ${MSGHDR}${MSG\n}")
```

- ISO date corresponds to the following syslog-ng template:

```
template("${ISODATE} ${HOST} ${MSGHDR}${MSG\n}")
```

- Extended is a deprecated option. Currently it duplicates the functionality of *ISO date*.
- Custom specifies a custom syslog-ng template in the appearing **Template** field.

For details on using syslog-ng templates, see .

5. Select how to organize the log files of this logspace from the **Filename template** field.
 - To save every message received during a day into a single file, select **All messages in one file**.
 - To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `${HOST}` macro of syslog-ng.
 - To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `${PROGRAM}` macro of syslog-ng.
 - To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `${HOST}-${PROGRAM}` macros of syslog-ng.
 - To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field.

NOTE:

For details on using filename templates, see .

Templates that generate an invalid path (for example, they use a filename longer than 246 characters or refer to a parent directory) will not work.

6. To create automatic daily backups of the logspace to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see [Data and configuration backups](#).
7. To archive the logspace automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies, see [Archiving and cleanup](#).

CAUTION:

Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.

8. To make the log files of this logspace available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see [Accessing log files across the network](#).
9. Set a size for the logspace in the **Warning size** field: SSB will send an alert if the size of this logspace exceeds the limit.

CAUTION:

Make sure that the Logspace exceeded warning size alert is enabled in **Basic Settings > Alerting > syslog-ng traps**, and that the mail settings of **Basic Settings > Management**, and the SNMP settings of **Basic Settings > Alerting** are correct. Otherwise, you will not receive any alert when the logspace exceeds the size limit. For details on alerting and monitoring, see also [Configuring system monitoring on SSB](#).

10. By default, members of the search group can view the stored messages online. Use the **Access control** option to control which usergroups can access the logspace. For details, see also [Managing user rights and usergroups](#).
11. Click **Commit**.

Managing logspaces

Logspaces are mostly managed automatically using backup and archiving policies, as described in [Data and configuration backups](#) on page 98 and [Archiving and cleanup](#) on page 110. However, backup and archiving can be started manually as well. To display the details of a logspace, click . A number of action buttons is shown in the top row.

NOTE: These options are not available for filtered and remote logspaces.

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this as member**. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

Figure 129: Log > Logspaces > Get current size – Managing logspaces



Archive, backup, cleanup, restore and emptying options

TIP: The size of the logspace is displayed in the **Size** row of the logspace details. To refresh the data, select **Get current size**.

- To start the backup process manually, click **Backup**.
- To restore the log files from the backup server to syslog-ng Store Box (SSB), click **Restore**.

⚠ CAUTION:

Restoring the backup replaces every log file of the logspace with the files from the backup. Any log message saved into the logspace since the backup is irrevocably lost.

- To start the archiving and the cleanup process manually, click **Archive/Cleanup**.

⚠ CAUTION:

If the archiving policy selected for the logspace is set to perform only cleanup, log messages older than the Retention Time are deleted and irrevocably lost. For details, see [Archiving and cleanup on page 110](#).

- To delete every log file in the logspace, click **Empty**. This option can be useful if you have to quickly free up space on SSB, or if you want to delete a logspace.

⚠ CAUTION:

This action deletes every file of the logspace. Any log message not archived or backed up is irrevocably lost.

You can still search archived logs of the logspace.

Similar action buttons are available at the top of the **Log > Logspaces** page to backup, archive, or delete the contents of every logspace. These actions are performed on every logspace with their respective settings, that is, clicking **Backup All** creates a backup of every logspace using the backup policy settings of the individual logspace.

Disk space

Type

Managing logspaces - Archive and backup logspaces

Logspaces are mostly managed automatically using backup and archiving policies, as described in [Data and configuration backups](#) on page 98 and [Archiving and cleanup](#) on page 110. However, backup and archiving can be started manually as well. To display the details of a logspace, click . A number of action buttons is shown in the top row.

NOTE: These options are not available for filtered and remote logspaces.

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this**

as member. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

Figure 130: Log > Logspaces — Managing logspaces



TIP: The size of the logspace is displayed in the **Size** row of the logspace details. To refresh the data, select **Get current size**.

- To start the backup process manually, click **Backup**.
- To restore the log files from the backup server to syslog-ng Store Box (SSB), click **Restore**.

⚠ CAUTION:

Restoring the backup replaces every log file of the logspace with the files from the backup. Any log message saved into the logspace since the backup is irrevocably lost.

- To start the archiving and the cleanup process manually, click **Archive/Cleanup**.

⚠ CAUTION:

If the archiving policy selected for the logspace is set to perform only cleanup, log messages older than the Retention Time are deleted and irrevocably lost. For details, see [Archiving and cleanup](#) on page 110.

- To delete every log file in the logspace, click **Empty**. This option can be useful if you have to quickly free up space on SSB, or if you want to delete a logspace.

⚠ CAUTION:

This action deletes every file of the logspace. Any log message not archived or backed up is irrevocably lost.

You can still search archived logs of the logspace.

Similar action buttons are available at the top of the **Log > Logspaces** page to backup, archive, or delete the contents of every logspace. These actions are performed on every logspace with their respective settings, that is, clicking **Backup All** creates a backup of every logspace using the backup policy settings of the individual logspace.

Assigning the SSB logspace of your choice to a custom cloud service provider data disk

From syslog-ng Store Box (SSB) version 6.9.0, you have more options for managing data disks for your logspaces.

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

This section describes how you can assign the SSB logspace of your choice to a custom cloud service provider data disk.

For more information about how you can add a custom cloud service provider data disk on the cloud service provider side, see [Adding a new custom cloud service provider data disk on the cloud service provider side](#).

For more information about adding a newly added custom cloud service provider data disk to your SSB configuration, see [Adding a new custom cloud service provider data disk to your SSB configuration](#).

For more information about the data disk information for your newly added custom cloud service provider data disks, see [Data disk information for your custom cloud service provider data disks](#).

Prerequisites

Assigning the SSB logspace of your choice to a custom cloud service provider data disk has the following prerequisites:

- A valid cloud provider subscription, and [a new custom cloud service provider data disk added on the cloud service provider side](#).
- Shutting your SSB device down before you begin adding a new managed disk on the cloud service provider side (for example, the Microsoft Azure portal).
- [Adding a new custom cloud service provider data disk to your SSB configuration](#).

Limitations

Assigning the SSB logspace of your choice to a custom cloud service provider data disk has the following limitations:

- **CAUTION:**
HAZARD OF DATA LOSS!
One Identity does not support decreasing the size of a Microsoft

Azure managed disk added to your SSB configuration as a custom cloud service provider data disk.

- **NOTE:** If you fail to assign the newly created SSB logspace of your choice to the new custom cloud service provider data disk, SSB will not automatically detect and use the new custom cloud service provider data disk to store messages on it.
- **NOTE:** Resizing your custom cloud service provider data disk on the Microsoft Azure side is quick, but your SSB configuration must resize the file system on your SSB side to match the resizing on the Microsoft Azure side. As a result, depending on the original size of the custom cloud service provider data disk and on the size you will expand it to, starting the virtual machine for your SSB device after you expand the size of your custom cloud service provider data disk may take a long time.
- In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.
- Transporting custom cloud service provider data disks between SSB devices is not supported.
- You can only add, remove, or resize a custom cloud service provider data disk in your configuration after shutting your SSB device down first. Hot swaps are not supported.
- You can only add an Azure managed disk as a custom cloud service provider data disk to your SSB on the Microsoft Azure portal side.
- SSB only accepts and detects empty disks as custom cloud service provider data disks. As a result, you can only add a newly created Azure managed disk as a custom cloud service provider data disk to your SSB configuration. Adding an already existing Azure managed disk (by clicking **Attach existing disks** on the Microsoft Azure portal) is not supported.
- One Identity does not support adding Azure managed disks to your SSB logspaces as a shared disk. You can only add one Azure managed disk to one SSB device.
- The number of managed disks you can add on the Microsoft Azure portal depends on the type of Microsoft Azure subscription you are using.

Assigning the SSB logspace of your choice to a custom cloud service provider data disk

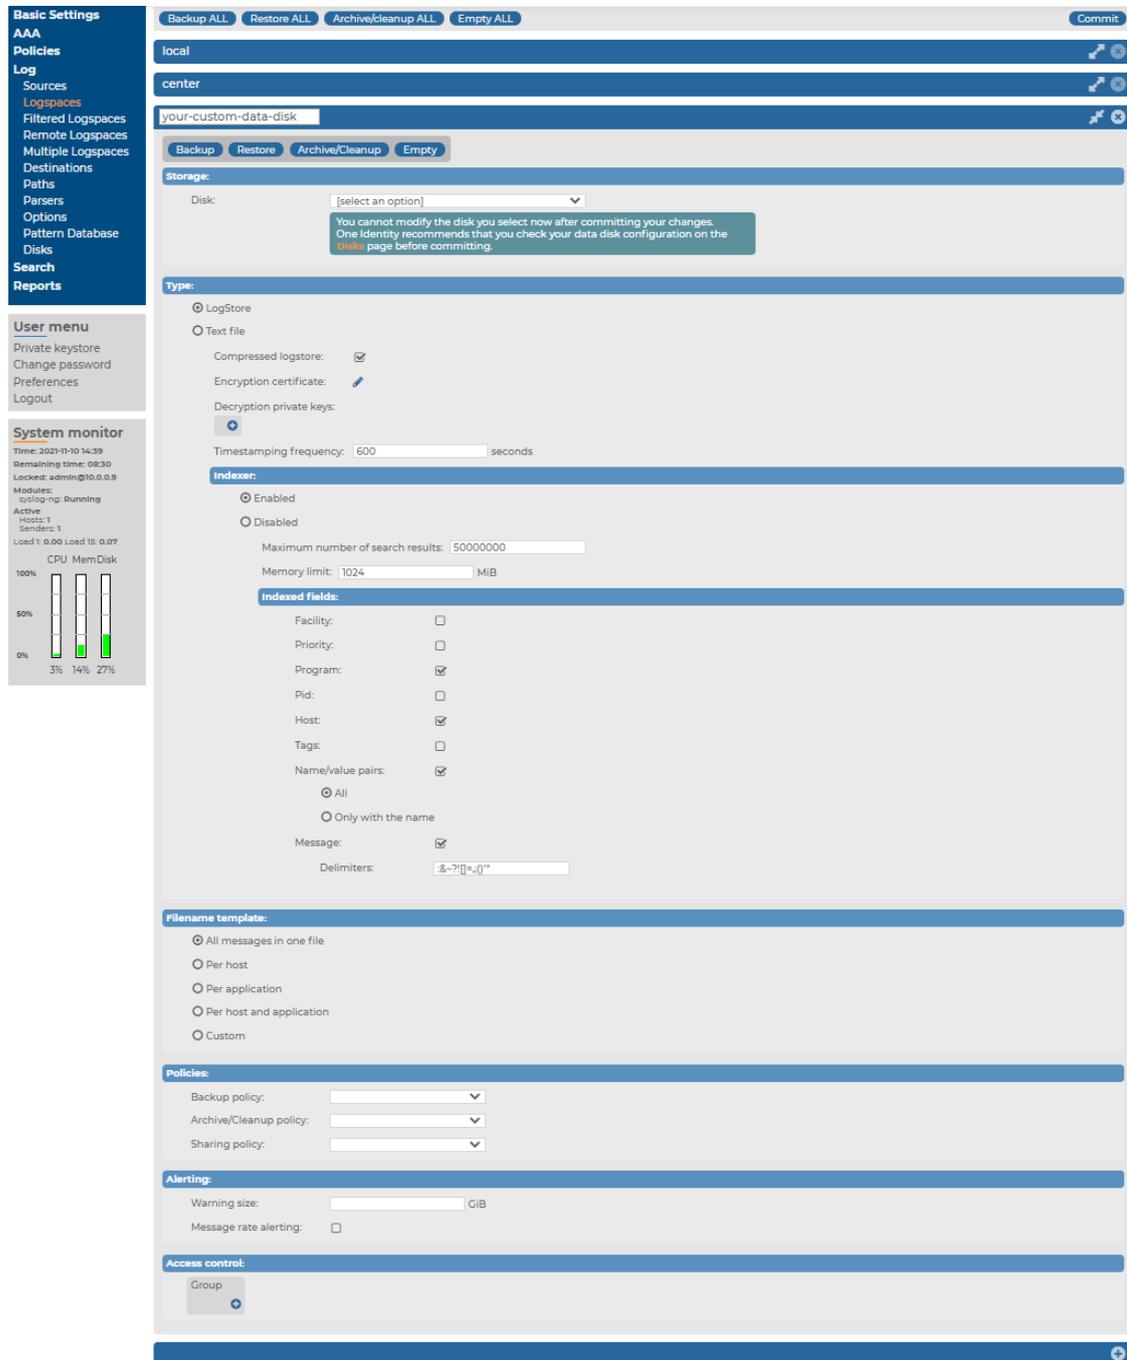
After adding a new managed disk on your cloud service provider side, and adding it as a custom cloud service provider data disk to your SSB configuration, you have to assign the SSB logspace of your choice to a custom cloud service provider data disk.

To assign the SSB logspace of your choice to a custom cloud service provider data disk,

1. [Add a new managed disk on your cloud service provider side.](#)
2. [Add the newly created managed disk as a custom cloud service provider data disk to your SSB configuration.](#)

- Navigate to **Log > Logspaces**, then click  to add a new logspace, and find a descriptive name for it (for example, the same name that you added for the new custom cloud service provider data disk, in this case, **your-custom-data-disk**).

Figure 131: Log > Logspaces — Adding a new logspace that you will assign to the new custom cloud service provider data disk



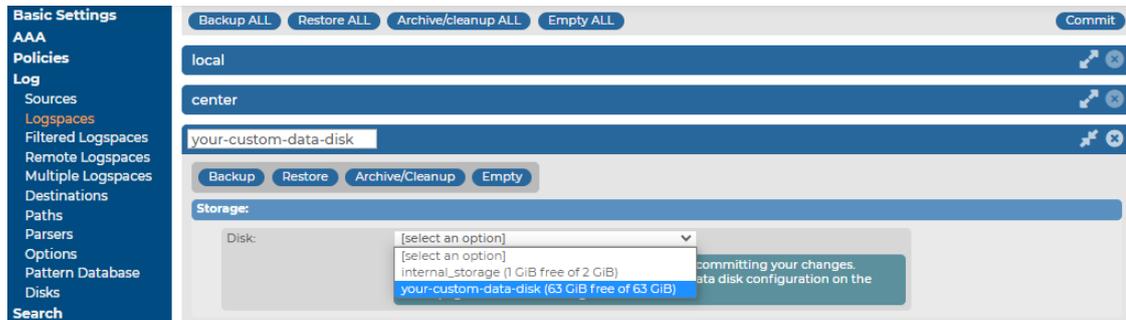
The screenshot displays the configuration interface for a new logspace. On the left, there is a navigation menu with sections for Basic Settings, AAA Policies, Log Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options, Pattern Database, Disks, Search, and Reports. Below the menu are sections for User menu and System monitor. The main configuration area includes:

- Storage:** A dropdown menu for 'Disk' with a warning message: "You cannot modify the disk you select now after committing your changes. One Identity recommends that you check your data disk configuration on the disks page before committing."
- Type:** Radio buttons for 'LogStore' (selected) and 'Text file'.
 - Compressed logstore:
 - Encryption certificate:
 - Decryption private keys:
 - Timestamping frequency: 600 seconds
- Indexer:** Radio buttons for 'Enabled' (selected) and 'Disabled'.
 - Maximum number of search results: 50000000
 - Memory limit: 1024 MIB
- Indexed fields:** Checkboxes for Facility, Priority, Program, Pid, Host, Tags, Name/value pairs, and Message.
 - Radio buttons for 'All' (selected) and 'Only with the name'
 - Delimiters: `;&-?[\+,-,()]*`
- Filename template:** Radio buttons for 'All messages in one file', 'Per host', 'Per application', 'Per host and application', and 'Custom'.
- Policies:** Dropdown menus for Backup policy, Archive/Cleanup policy, and Sharing policy.
- Alerting:** Text input for Warning size and a checkbox for Message rate alerting.
- Access control:** A dropdown menu for Group.

- Under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**, select the data disk you want to use for your new logspace.

In addition to the default SSB internal storage option, you will have more data disk options in the dropdown list, including any newly created custom cloud service provider data disks you recently added to your SSB configuration.

Figure 132: Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk — Selecting the data disk for your new logspace



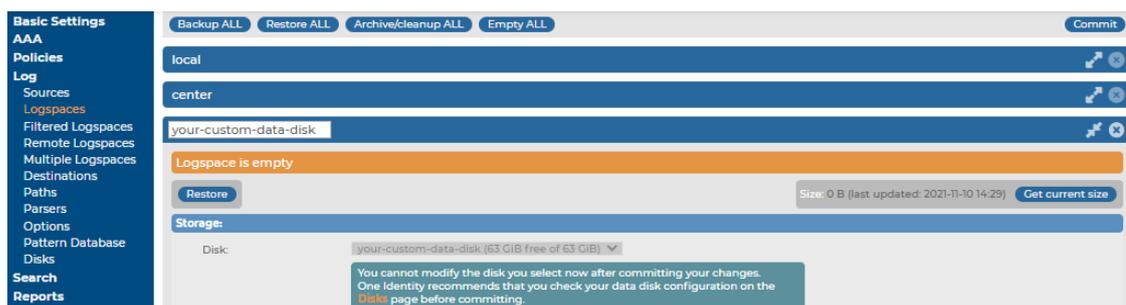
NOTE:

If you set a custom cloud service provider data disk under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**, you can not change it after committing your logspace configuration.

NOTE: For newly created logspaces, selecting a data disk for your logspace is required.

- After selecting the data disk of your choice (in this case, the custom cloud service provider data disk we added earlier, that is, **your-custom-data-disk**), click **Commit**.

Figure 133: Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk — New logspace assigned to your new custom cloud service provider data disk



NOTE: After adding a newly created custom cloud service provider data disk to your SSB configuration, and assigning the SSB logspace of your choice to it, removing the custom cloud service provider data disk from your SSB configuration has a strict order, especially if there are referenced objects tied to the custom cloud

service provider data disk. For more information about the process to remove such custom cloud service provider data disks from your SSB configuration, see [Removing a custom cloud service provider data disk from your SSB configuration](#).

TIP: The size of the logspace is displayed in the **Size** row of the logspace details. To refresh the data, select **Get current size**.

Creating filtered logspaces

Filtered logspaces allow you to create a smaller, filtered subset of the logs contained in an existing local, remote or multiple logspace. Assigning a user group to a filtered logspace enables fine grained access control by creating a group which sees only a subset of the logs from a logspace.

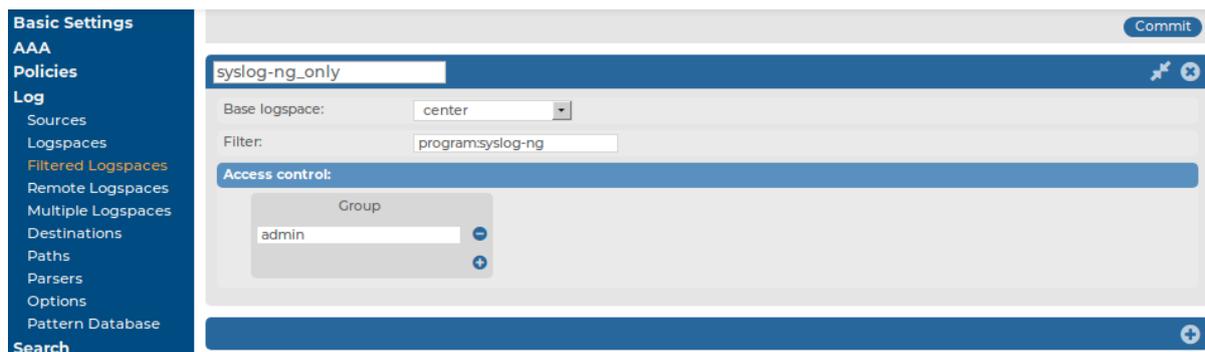
You can use the same search expressions and logic as on the Search interface to create a filtered logspace. In the following example, we have configured a filtered logspace that only contains messages from syslog-ng:

NOTE: The filtered logspace is only a view of the base logspace. The log messages are still stored in the base logspace (if the base logspace is a remote logspace, the log messages are stored on the remote syslog-ng Store Box (SSB) appliance). Therefore, you cannot alter any configuration parameters of the logspace directly. To do this, navigate to the base logspace itself.

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this as member**. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

Figure 134: Log > Filtered Logspaces — Filtered logspaces



To create filtered logspaces

1. Navigate to **Log > Filtered Logspaces** and click .
2. Enter a name for the logspace into the top field. Use descriptive names that help you to identify the source easily. Note that the name of the logspace must begin with a number or a letter.
3. Choose which logspace to filter in **Base logspace**.
4. Enter the search expression in the **Filter** field.

You can create complex searches using wildcards and boolean expressions. For more information and practical examples, see [Using complex search queries](#) on page 338.

NOTE: SSB only indexes the first 59 characters of every name-value pair (parameter). This has two consequences:

- If the parameter is longer than 59 characters, an exact search might deliver multiple, imprecise results.

Consider the following example. If the parameter is:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

SSB indexes it only as:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

This corresponds to the first 59 characters. As a result, searching for:

```
nvpair:.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

returns all log messages that contain:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

- Using wildcards might lead to the omission of certain messages from the search results.

Using the same example as above, searching for the value:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-12345
```

does not return any results (as the 12345 part was not indexed). Instead, you have to search for:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-*
```

This, as explained above, might find multiple results.

5. By default, members of the search group can view the stored messages online. Use the **Access control** option to control which usergroups can access the logspace. For

details, see also [Managing user rights and usergroups](#) on page 127.

6. Click **Commit**.

Creating remote logspaces

The syslog-ng Store Box (SSB) appliance can access and search logspaces (including filtered logspaces) on other SSB appliances. To configure SSB to access a logspace on another (remote) SSB, set up a remote logspace.

Once configured, remote logspaces can be searched like any other logspace on SSB. You can also create filtered logspaces that are based on the remote logspace.

NOTE: Note that you cannot alter the configuration, archive, back up, or empty the contents of the logspace on the remote SSB.

NOTE: If the remote logspace becomes inaccessible, you will not be able to view the contents of that logspace.

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this as member**. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

Figure 135: Log > Remote Logspaces — Remote logspaces

The screenshot shows the SSB web interface for configuring a remote logspace. The left sidebar contains a navigation menu with categories: Basic Settings, AAA, Policies, Log (Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options, Pattern Database), Search, and Reports. The main content area is titled 'Cassini' and includes a 'Commit' button in the top right. The configuration fields are: Host (10.40.22.31), Username (remote_log_viewer), Password (masked with dots), Remote logspace name (Huygens), and Remote certificate authority (C=HU/L=Budapest/O=Balabit/OU=TC/CN=Example User). Below these fields is an 'Access control' section with a 'Group' dropdown menu currently set to 'sec_ops'.

Prerequisites:

- You have verified that the version number of the remote SSB equals (or exceeds) the version number of the SSB where the remote logspace is created.
- You have configured a user on the remote SSB that can access the logspace you want to reach.

- If the logspace is encrypted, you have verified that the user has the necessary certificates.
- You have downloaded the CA X.509 certificate of the remote SSB.

To download the server certificate, navigate to **Basic Settings > Management > SSL certificate > CA X.509 certificate**, and click on the certificate.

To create remote logspaces

1. Navigate to **Log > Remote Logspaces** and click .
2. Enter a name for the logspace into the top field. Use descriptive names that help you to identify the source easily. Note that the name of the logspace must begin with a number or a letter.
3. Enter the IP address or hostname of the remote SSB in the **Host** field.
4. Enter the username of the user configured for accessing the logspace on the remote SSB in the **Username** field.
5. Enter the password of the same user in the **Password** field.
6. Enter the name of the logspace as it appears on the remote SSB in the **Remote logspace name** field.
7. In the **Remote certificate authority** section, click  to upload the server certificate of the remote SSB. A pop-up window is displayed.

NOTE:

It is not possible, nor required to upload a certificate chain to Remote certificate authority. The certificate chain is used by the server, not the remote logspace.

If you want to use a certificate chain when authenticating to a remote logspace, do the following:

1. Upload the root CA to **Log > Remote Logspaces > Remote certificate authority**.
2. Upload the intermediate CA and end-entity (server) certificate to **Basic Settings > Management > SSL certificate > Server X.509 certificate**.

Click **Browse**, select the certificate of the remote SSB, then click **Upload**.

8. By default, members of the search group can view the stored messages online. Use the **Access control** option to control which usergroups can access the logspace. For details, see also [Managing user rights and usergroups](#) on page 127.
9. Click .

Creating multiple logspaces

If you have several syslog-ng Store Box (SSB) appliances located at different sites, you can view and search the logs of these machines from the same web interface without having to log on to several different interfaces.

Creating multiple logspaces can also be useful if you want to pre-filter log messages based on different aspects and then share these filtered logs only with certain user groups.

Multiple logspaces aggregate the messages that arrive from the member logspaces within the group. The new log messages are listed below each other every second.

Once configured, multiple logspaces can be searched like any other logspace on SSB. You can also create filtered logspaces that are based on the multiple logspaces.

NOTE: Multiple logspaces are only a view of the member logspaces. The log messages are still stored in the member logspaces (if the member logspace is a remote logspace, the log messages are stored on the remote SSB). As a result, you cannot alter any configuration parameters of the logspace directly. To do this, navigate to the member logspace itself.

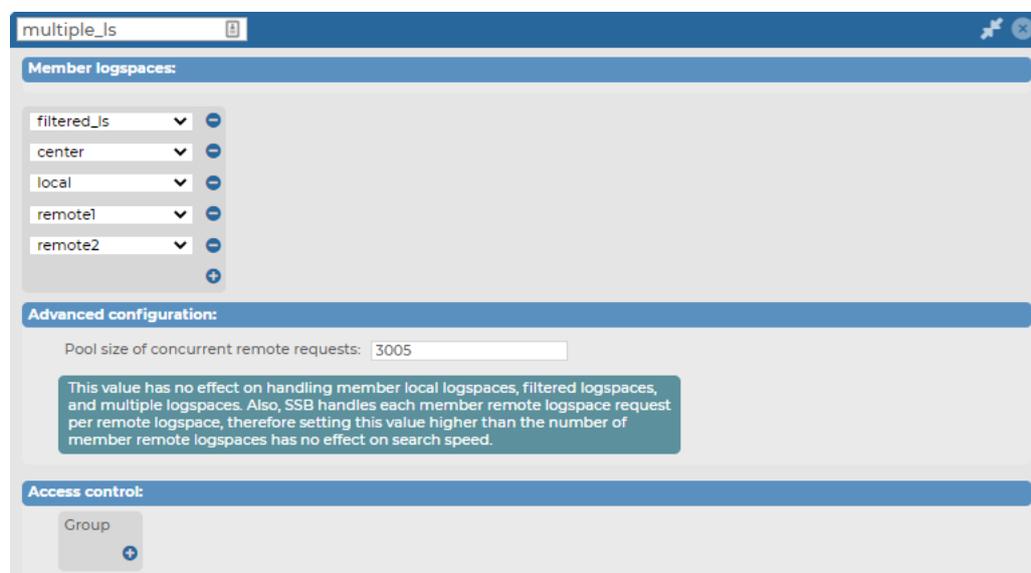
NOTE: If a remote member logspace becomes inaccessible, you will not be able to view the contents of that logspace.

NOTE: Using multiple logspaces may decrease the performance of the appliance. If possible, manage your logspaces without using multiple logspaces (for example, instead of including several filtered logspaces into a multiple logspace, use several search expressions in a filtered logspace).

NOTE: If there are any multiple logspaces using your logspace as a member logspace, the multiple logspaces in question will be listed under **Multiple logspaces using this as member**. The list items are clickable links that will take you directly to the logspaces on the SSB web interface.

This list is only visible on the SSB web interface for **Logspaces**, **Filtered Logspaces**, **Multiple Logspaces**, and **Remote Logspaces** if they are member logspaces in any multiple logspaces.

Figure 136: Log > Multiple Logspaces — Creating multiple logspaces



To create multiple logspaces

1. Navigate to **Log > Multiple Logspaces** and click .
2. Enter a name for the multiple logspace. Use descriptive names that help you to identify the source easily.

NOTE: When naming your multiple logspace, consider that the name of your multiple logspace must begin with either of the following:

- a letter of the English alphabet
- a number
- an underscore (_)
- a hyphen (-)

3. Select the **Member logspaces** from the list. To add a new member logspace, click  and select a different logspace.

NOTE: Consider the following:

- You can only select logspaces that you previously configured as a local logspace / filtered logspace / remote logspace / multiple logspace.
- You cannot add a multiple logspace to itself, only to a different multiple logspace. For example, you cannot add **multiple_Is** as a member of your **multiple_Is** multiple logspace.
- You can add several of the same logspace types (that is, local / filtered / remote / multiple logspaces), but you cannot add the same unique logspace more than once.

4. Under **Advanced configuration**, set the pool size of the concurrent remote requests for the configured multiple logspaces. The default value is 2, and the minimum configurable value is 1.

NOTE: Consider the following:

- Higher pool size settings may increase search speed in remote logspaces.
- The **Pool size of concurrent remote requests** value is set per multiple logspace.
- SSB only accesses remote logspaces parallelly. Parallel access is not effective for the following types of member logspaces:
 - filtered logspaces (even if the base of the filtered logspace is a remote logspace)
 - multiple logspaces
 - local logspaces
- If more than one member remote logspace within the multiple logspace is located on the same remote SSB, then one of the following scenarios are possible:

- The parallel or concurrent remote requests may result in performance issues on the remote SSB appliance.
 - Increasing the **Pool size of concurrent remote requests** value may affect search speed for remote logspaces located on the same remote host.
- Increasing the **Pool size of concurrent remote requests** value will increase your network usage.
 - SSB handles each member remote logspace request per remote logspace, as a single element of the available pool. As a result, setting the **Pool size of concurrent remote requests** value to the same value as the number of member remote logspaces may increase your search speed, but setting it any higher will not increase your search speed further.
5. By default, members of the search group can view the stored messages online. Use the **Access control** option to control which usergroups can access the logspace. For details, see also [Managing user rights and usergroups](#) on page 127.
 6. Click **Commit**.

Accessing log files across the network

The log files stored on syslog-ng Store Box (SSB) can be accessed as a network share if needed using the Samba (CIFS) or Network File System (NFS) protocols. Sharing is controlled using policies that specify the type of the share and the clients (hosts) and users who can access the log files. Sharing is possible also if SSB is part of a domain.

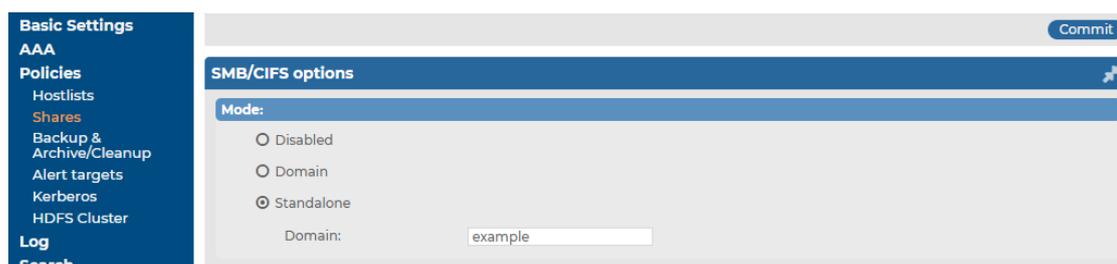
- If you manage SSB users locally, users who have SSB account can access the shared folders. Complete [Sharing log files in standalone mode](#) on page 239.
- If you manage SSB users from LDAP, you must join SSB to your domain. Complete [Sharing log files in Domain mode](#) on page 241.
- For details on how to access the shared files, see [Accessing shared files](#) on page 243.
- You can access logspaces (local, filtered, remote and multiple) through RPC API as well. For details on RPC API, see [The SSB RPC API](#) on page 389.

Sharing log files in standalone mode

To share log files in standalone mode

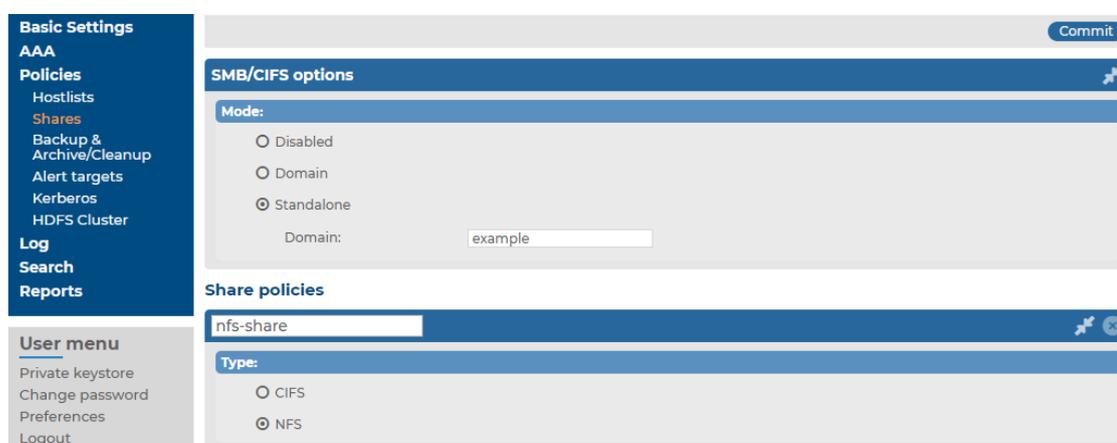
1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Standalone mode**.

Figure 137: Policies > Shares > SMB/CIFS options — Sharing logspaces



2. Select  to create a new share policy and enter a name for the policy.
3. Select the type of the network share from the **Type** field.

Figure 138: Policies > Shares > Share policies — Creating share policies



- To access the log files using NFS (Network File System), select **NFS**.
- To access the log files using Samba (Server Message Block protocol), select **CIFS**.

NOTE: From SSB version 5.2.0, SSB only supports SMB 2.1 and later. If you are using a Windows version earlier than Windows 2008R2, make sure that it supports SMB 2.1 or later. Otherwise, the Windows machine cannot connect to the SSB share.

4. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an syslog-ng Store Box (SSB) account has access to every shared log file.
 - To control which users can access the shared files, enter the name of the usergroup who can access the files into the **Allowed group** field. For details on local user groups, see [Managing local usergroups](#) on page 121.
 - To limit the hosts from where the shares can be accessed, create a hostlist and

select it from the **Hostlist** field. For details on creating hostlists, see [Creating hostlist policies](#) on page 189.

5. Click **Commit**.
6. To display the details of the logspace, navigate to **Log > Logspaces** and click .
7. Select the share policy to use from the **Sharing policy** field.

Figure 139: Log > Logspaces > Policies — Setting the share policy of a logspace



The screenshot shows a configuration panel titled "Policies:". It contains three dropdown menus: "Backup policy:", "Archive/Cleanup policy:", and "Sharing policy:". The "Sharing policy:" dropdown is currently set to "myshare".

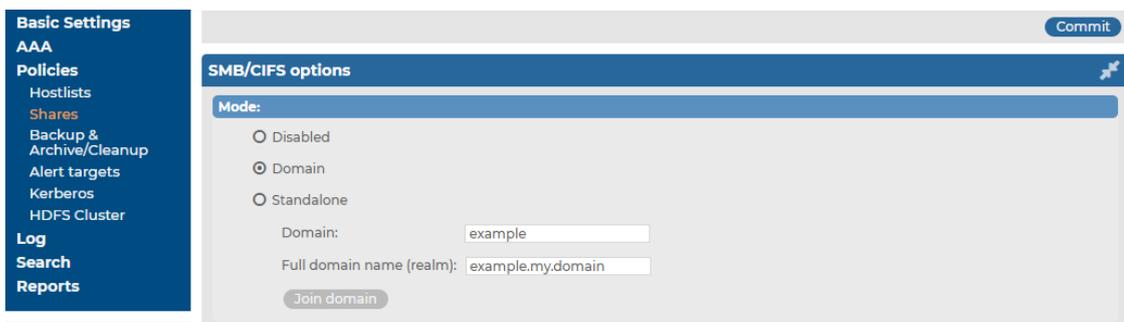
8. Click **Commit**.
9. [Mount the shared logspace](#) from your computer to access it.

Sharing log files in Domain mode

To share log files in Domain mode

1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Domain mode**.
2. Enter the name of the domain (for example mydomain) into the **Domain** field.

Figure 140: Policies > Shares > SMB/CIFS options — Joining a domain



The screenshot shows the "SMB/CIFS options" configuration page. The "Mode:" section has three radio buttons: "Disabled", "Domain" (which is selected), and "Standalone". Below this, there are two text input fields: "Domain:" with the value "example" and "Full domain name (realm):" with the value "example.my.domain". A "Join domain" button is located at the bottom of the form.

3. 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 syslog-ng Store Box (SSB). To check this, navigate to **Basic Settings > Troubleshooting > Ping**, enter the full domain name into the **Hostname** field, and select **Ping host**.

Click **Commit**.

4. Click **Join domain**. A pop-up window is displayed.
5. SSB requires an account to your domain to be able to join the domain. Enter the name of the user into the **Username** field, and the corresponding password into the **Password** field.

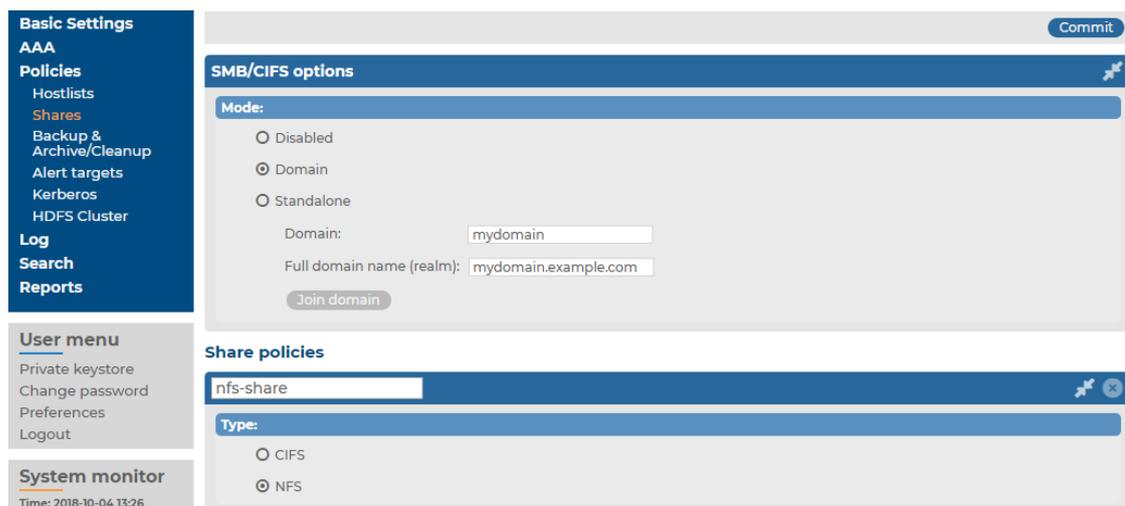
NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

Optionally, you can enter the name of your domain controller into the **Domain controller** field. If you leave this field blank, SSB will try 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 SSB. To check this, navigate to **Basic Settings > Troubleshooting > Ping**, enter the name of the domain controller into the **Hostname** field, and select **Ping host**.

6. Click **Join domain**.
7. Select  to create a new share policy and enter a name for the policy.

Figure 141: Policies > Shares > Share policies – Creating share policies



The screenshot shows the SSB web interface. On the left is a navigation menu with sections: Basic Settings, AAA, Policies (with sub-items: Hostlists, Shares, Backup & Archive/Cleanup, Alert targets, Kerberos, HDFS Cluster), Log, Search, and Reports. Below this is a 'User menu' with options: Private keystore, Change password, Preferences, and Logout. At the bottom left is a 'System monitor' section showing the time: 2018-10-04 13:26. The main content area is divided into two panels. The top panel is titled 'SMB/CIFS options' and has a 'Commit' button in the top right. It contains a 'Mode:' section with three radio buttons: 'Disabled', 'Domain' (which is selected), and 'Standalone'. Below these are two text input fields: 'Domain:' with the value 'mydomain' and 'Full domain name (realm):' with the value 'mydomain.example.com'. A 'Join domain' button is at the bottom of this panel. The bottom panel is titled 'Share policies' and has a search bar containing 'nfs-share'. It contains a 'Type:' section with two radio buttons: 'CIFS' and 'NFS' (which is selected).

8. Select the type of the network share from the **Type** field.
 - To access the log files using NFS (Network File System), select **NFS**.
 - To access the log files using Samba (Server Message Block protocol), select **CIFS**.

NOTE: From SSB version 5.2.0, SSB only supports SMB 2.1 and later. If you are using a Windows version earlier than Windows 2008R2, make sure that it

supports SMB 2.1 or later. Otherwise, the Windows machine cannot connect to the SSB share.

9. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an SSB account has access to every shared log file.
 - To control which users can access the shared files, enter the name of the domain that can access the files (specified in Step 2) into the **Allowed group** field. Note that the users and SSB must be members of the same domain.
 - To limit the hosts from where the shares can be accessed, create a hostlist and select it from the **Hostlist** field. For details on creating hostlists, see [Creating hostlist policies](#) on page 189.
10. Click **Commit**.
11. To display the details of the logspace, navigate to **Log > Logspaces** and click .
12. Select the share policy to use from the **Sharing policy** field.

Figure 142: Log > Logspaces > Policies — Setting the share policy of a logspace



Policies:	
Backup policy:	<input type="text"/>
Archive/Cleanup policy:	<input type="text"/>
Sharing policy:	myshare

13. Click **Commit**.
14. [Mount the shared logspace](#) from your computer to access it.

Accessing shared files

This section describes how to access log files that are shared using a share policy. For details on sharing log files, see [Accessing log files across the network](#) on page 239.

Every shared logspace is available as a separate shared folder, even if they all use a single share policy. The name of the shared folder is the name of the logspace.

Within the shared folder, the log files are organized into the following directory structure: YEAR/MM-DD/. The files are named according to the filename template set for the logspace. The extension of logstore files is .store, while the extension of text files is .log. Note that the root directory of the share may also contain various files related to the logspace, like index files for logstores. All files are read-only.

When using NFS for sharing the logspace, the name of the shared folder will be the following: /exports/{logspace_id}/....

Mount a shared logspace

The following examples show how to mount a shared logspace.

On Linux NFS:

```
mount -t nfs {ssb_ip}:/exports/{logspace_id} {where_to_mount}
```

On Linux SMB:

From SSB version 5.2.0, SSB only supports SMB 2.1 and later. If you are using a Linux version that uses SMB protocol version earlier than 2.1, add the option `-o vers=2.1` to ensure that SSB uses SMB 2.1. For example:

```
mount -t cifs //{ssb_ip}/{logspace_id} /path/to/mount/shared/logspace/ -o  
username={username},password={password},vers=2.1
```

On Windows NFS:

1. Make sure that you have the "Services for NFS" Windows component installed. If not, you can install the [NFS client](#) from the Windows interface.
2. Open **regedit**, and navigate to `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default`
3. Create two new DWORD keys called `AnonymousGID` and `AnonymousUID`. Set their values to 0.
4. Restart the NFS client service from an elevated privilege command prompt. Use the following commands: **`nfsadmin client stop`**, then **`nfsadmin client start`**
5. Mount the share from the command prompt. (Alternatively, you can also use the 'Map network drive...' function of the File Explorer.)

```
mount {ssb_ip}://exports/{logspace_id} {drive_letter}:
```

For example, the following command mounts the local logspace as drive G:

```
mount 192.168.1.1://exports/local G:
```

After mounting the shared logspace, it is visible in the file explorer. If it is not visible in the file explorer, you have probably used a different user to mount the share. To avoid this problem, you can mount the share again with the same user. Otherwise, you can access it from the command prompt using the **`{drive_letter}:`** command, even if it is not visible in the file explorer.

On Windows SMB:

Map the share from the command prompt. (Alternatively, you can also use the 'Map network drive...' function of the file explorer.)

```
net use {drive_letter}: \\{ssb_ip}\{logspace_name} /user:{user_name} "  
{password}"
```

For example, the following command maps the local logspace as drive G:

```
net use G: \\192.168.1.1:\local /user:myuser "mypassword"
```

After mapping the shared logspace, it is visible in the File Explorer. If it is not visible in the file explorer, you have probably used a different user to mount the share. To avoid this problem, you can mount the share again with the same user. Otherwise, you can access it from the command prompt using the **{drive_letter}**: command, even if it is not visible in the file explorer.

NOTE: NOTE: In case of accessing shared files in domain mode, also include the domain name in the command: `net use {drive_letter}: \\{ssb_ip}\{logspace_name} /user: {domain_name}\{user_name} "{password}"`

For example, the following command maps the local logspace as drive G:

```
net use G: \\192.168.1.1:\local /user:mydomain\myuser "mypassword"
```

For information on viewing encrypted logspace files, see [Viewing encrypted logs with logcat](#) on page 222.

Managing custom cloud service provider data disks for your logspaces in SSB

From syslog-ng Store Box (SSB) version 6.9.0, you have more options for managing data disks for your logspaces.

For more information about assigning the SSB logspace of your choice to the new custom cloud service provider data disk, see [Assigning the SSB logspace of your choice to a custom cloud service provider data disk](#).

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

NOTE:

In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.

Adding a custom cloud service provider data disk to your SSB configuration requires procedures that you must complete in a strict order, both on the SSB side, and on the cloud service provider side, too.

Possible use cases and scenarios for using custom cloud service provider data disks with SSB

From syslog-ng Store Box (SSB) version 6.9.0, you have more options for managing data disks for your logspaces.

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

This section lists the possible use cases and scenarios for using custom cloud service provider data disks with SSB.

From version 6.9.0, you can use Microsoft Azure managed disks added as SSB custom cloud service provider data disks in these scenarios and use cases:

- [Adding a new custom cloud service provider data disk to your SSB configuration](#)
- [Assigning the SSB logspace of your choice to a custom cloud service provider data disk](#)
- [Increasing the size of a custom cloud service provider data disk that you use in your SSB configuration](#)
- [Removing a custom cloud service provider data disk from your SSB configuration](#)

Adding a new custom cloud service provider data disk to your SSB configuration

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

This section and its subsections describe how you can add a new custom cloud service provider data disk to your SSB configuration.

For more information about how you can add a custom cloud service provider data disk on the cloud service provider side, see [Adding a new custom cloud service provider data disk on the cloud service provider side](#).

For more information about the data disk information for your custom cloud service provider data disks, see [Data disk information for your custom cloud service provider data disks](#).

For more information about assigning the SSB logspace of your choice to a custom cloud service provider data disk, see [Assigning the SSB logspace of your choice to a custom cloud service provider data disk](#).

Prerequisites

When adding a new custom cloud service provider data disk to your SSB configuration, consider the following prerequisites in advance:

- If you want to add a custom cloud service provider data disk to your SSB configuration, you must [add a new custom cloud service provider data disk on your cloud service provider side](#) before you can add it to your configuration on the SSB side.
- If you want to add custom cloud service provider data disks to your configuration in SSB, you must add the custom cloud service provider data disk under **Log > Disks > <your-custom-data-disk>**, and commit your changes. After that, the newly added custom cloud service provider data disk will be available for selection in the dropdown list under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**.
- You can only add, remove, or resize a custom cloud service provider data disk in your configuration after shutting your SSB device down first. Hot swaps are not supported.

Limitations

When adding a new custom cloud service provider data disk to your SSB configuration, consider the following limitations in advance:

CAUTION:

HAZARD OF DATA LOSS!

One Identity does not support decreasing the size of a Microsoft Azure managed disk added to your SSB configuration as a custom cloud service provider data disk.

- **NOTE:** After adding a newly created custom cloud service provider data disk to your SSB configuration, and assigning the SSB logspace of your choice to it, removing the custom cloud service provider data disk from your SSB configuration has a strict order, especially if there are referenced objects tied to the custom cloud service provider data disk. For more information about the process to remove such custom cloud service provider data disks from your SSB configuration, see [Removing a custom cloud service provider data disk from your SSB configuration](#).
- **NOTE:** If you fail to assign the newly created SSB logspace of your choice to the new custom cloud service provider data disk, SSB will not automatically detect and use the new custom cloud service provider data disk to store messages on it.
- **NOTE:** Resizing your custom cloud service provider data disk on the Microsoft Azure side is quick, but your SSB configuration must resize the file system on your SSB side to match the resizing on the Microsoft Azure side. As a result, depending on

the original size of the custom cloud service provider data disk and on the size you will expand it to, starting the virtual machine for your SSB device after you expand the size of your custom cloud service provider data disk may take a long time.

- In version 6.9.0, SSB supports adding Microsoft Azure managed disks as custom cloud service provider data disks.
- In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.
- If you set a custom cloud service provider data disk under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**, you can not change it after committing your logspace configuration.
- You can only configure custom cloud service provider data disks for logspaces that you newly create. One Identity does not support configuring custom cloud service provider data disks in logspaces that you configured and committed earlier.
- Transporting custom cloud service provider data disks between SSB devices is not supported.
- You can only add, remove, or resize a custom cloud service provider data disk in your configuration after shutting your SSB device down first. Hot swaps are not supported.
- You can only add an Azure managed disk as a custom cloud service provider data disk to your SSB on the Microsoft Azure portal side.
- SSB only accepts and detects empty disks as custom cloud service provider data disks. As a result, you can only add a newly created Azure managed disk as a custom cloud service provider data disk to your SSB configuration. Adding an already existing Azure managed disk (by clicking **Attach existing disks** on the Microsoft Azure portal) is not supported.
- One Identity does not support adding Azure managed disks to your SSB logspaces as a shared disk. You can only add one Azure managed disk to one SSB device.
- The number of managed disks you can add on the Microsoft Azure portal depends on the type of Microsoft Azure subscription you are using.

Adding a new custom cloud service provider data disk to your SSB configuration

Adding a new custom cloud service provider data disk to your SSB configuration requires procedures that you must complete in a strict order, both on the SSB side, and on the cloud service provider side, too.

To add a new custom cloud service provider data disk to your SSB configuration,

1. [Add a new custom cloud service provider data disk \(for example, a Microsoft Azure managed disk\) on your cloud service provider side.](#)
2. [Add the custom cloud service provider data disk to your SSB configuration.](#)

3. [Assign the SSB logspace of your choice to the newly added custom cloud service provider data disk.](#)

NOTE: After adding a newly created custom cloud service provider data disk to your SSB configuration, and assigning the SSB logspace of your choice to it, removing the custom cloud service provider data disk from your SSB configuration has a strict order, especially if there are referenced objects tied to the custom cloud service provider data disk. For more information about the process to remove such custom cloud service provider data disks from your SSB configuration, see [Removing a custom cloud service provider data disk from your SSB configuration.](#)

Adding a new custom cloud service provider data disk on the cloud service provider side

This section and its subsections describe the possible ways you can add custom cloud service provider data disks to your syslog-ng Store Box (SSB) configuration on the cloud service provider side.

For more information about how you can add custom cloud service provider disks to your SSB configuration on the SSB side, see [Adding a new custom cloud service provider data disk to your SSB configuration on the SSB side.](#)

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.

Adding a new Microsoft Azure managed disk as a custom cloud service provider data disk on the Microsoft Azure portal side

This section describes how you can add a new Microsoft Azure managed disk as a custom cloud service provider data disk to your syslog-ng Store Box (SSB) configuration on the Microsoft Azure portal side.

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.

Adding a new Microsoft Azure managed disk as a custom cloud service provider data disk to your SSB configuration requires procedures that you must complete in a strict order, both on the SSB side, and on the Microsoft Azure portal, too.

NOTE: After you configure your custom Microsoft Azure managed disk on the Microsoft Azure side, you must add the disk to your SSB configuration, and then assign the SSB logspace of your choice to it on the SSB side.

For more information about how you can add a Microsoft Azure managed disk to your SSB configuration on the SSB side, see [Adding a new custom cloud service provider data disk to your SSB configuration on the SSB side](#).

Prerequisites

When adding a new Microsoft Azure managed disk that you will later want to add to your SSB configuration, consider the following prerequisites in advance:

- You can only add an Azure managed disk as a custom cloud service provider data disk to your SSB on the Microsoft Azure portal side.
- The number of managed disks you can add on the Microsoft Azure portal depends on the type of Microsoft Azure subscription you are using.

Limitations

When adding a new Microsoft Azure managed disk that you will later want to add to your SSB configuration, consider the following limitations in advance:

CAUTION:

HAZARD OF DATA LOSS!

Although it is possible to decrease the size of Microsoft Azure managed disks, One Identity does not support decreasing the size of custom cloud service provider data disks already added to your SSB configuration. One Identity only supports increasing the size of such custom cloud service provider data disks after adding them to your SSB configuration.

- In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.
- You can only add, remove, or resize a custom cloud service provider data disk in your configuration after shutting your SSB device down first. Hot swaps are not supported.
- SSB only accepts and detects empty disks as custom cloud service provider data disks. As a result, you can only add a newly created Azure managed disk as a custom cloud service provider data disk to your SSB configuration. Adding an already existing Azure managed disk (by clicking **Attach existing disks** on the Microsoft Azure portal) is not supported.

- Transporting custom cloud service provider data disks between SSB devices is not supported.
- One Identity does not support adding Azure managed disks to your SSB logspaces as a shared disk. You can only add one Azure managed disk to one SSB device.

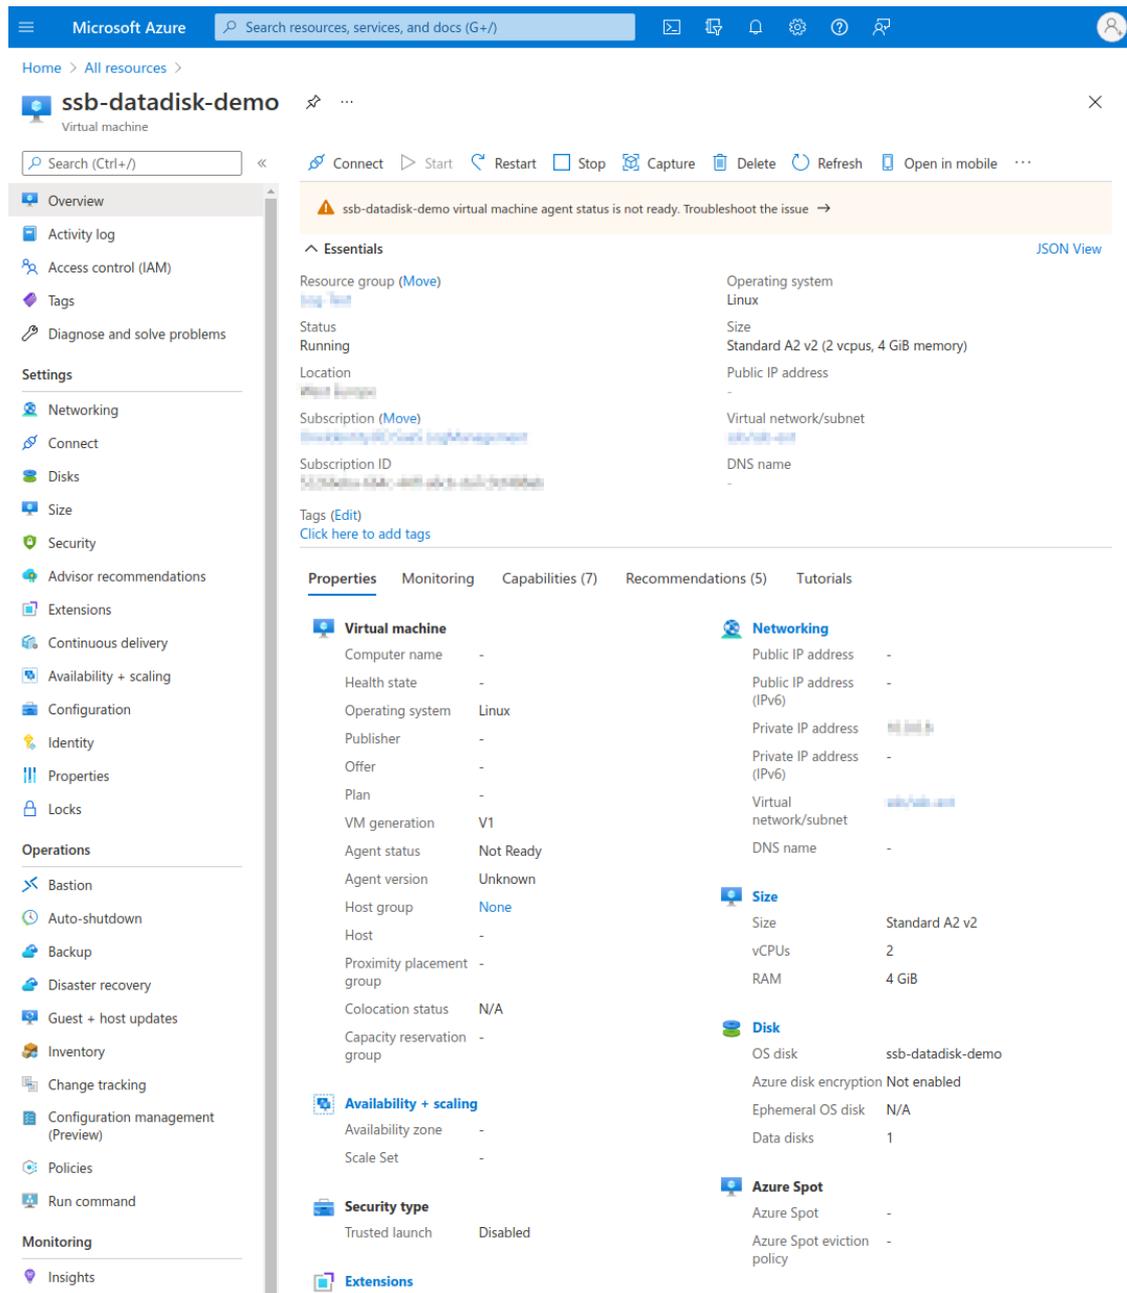
Adding a new Microsoft Azure managed disk as a custom cloud service provider data disk on the Microsoft Azure portal side

Adding a new Microsoft Azure managed disk as a custom cloud service provider data disk to your SSB configuration requires procedures that you must complete in a strict order.

To add a new Microsoft Azure managed disk as a custom cloud service provider data disk on the Microsoft Azure portal side,

1. [Shut down your SSB device](#) completely.
2. Log in to the [Microsoft Azure portal](#).
3. Navigate to the virtual machine of your SSB device.

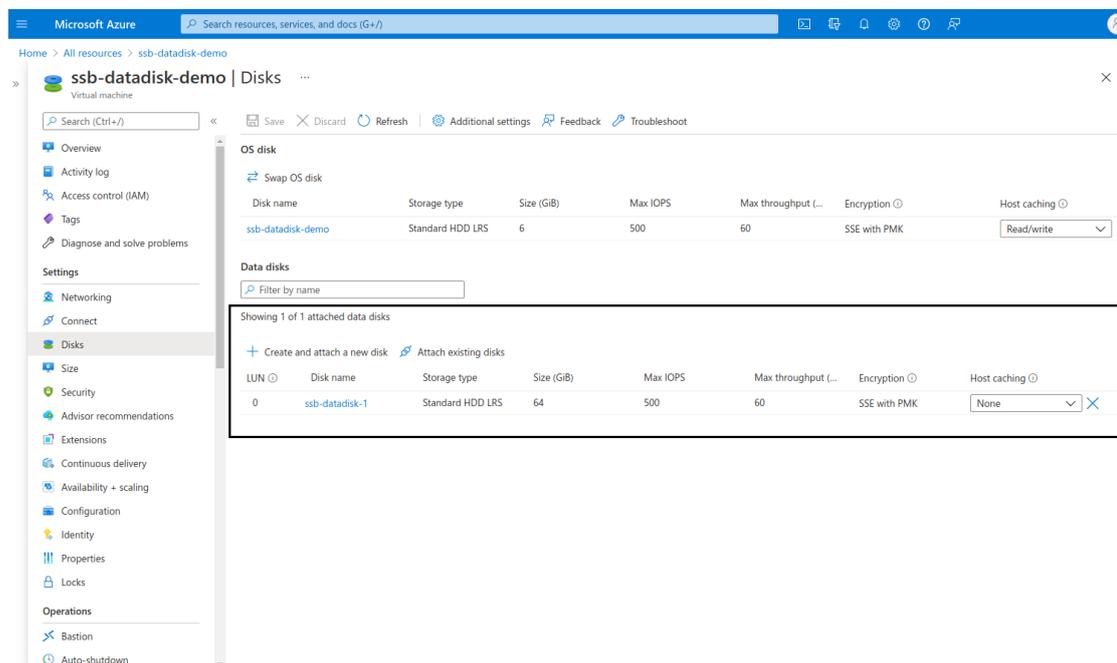
Figure 143: Home > All resources > <your-virtual-machine-in-microsoft-azure>



4. Follow the instructions of the *Microsoft Azure online tutorial documentation* to [attach a new data disk](#) with the size of your choice to the virtual machine of your SSB device.

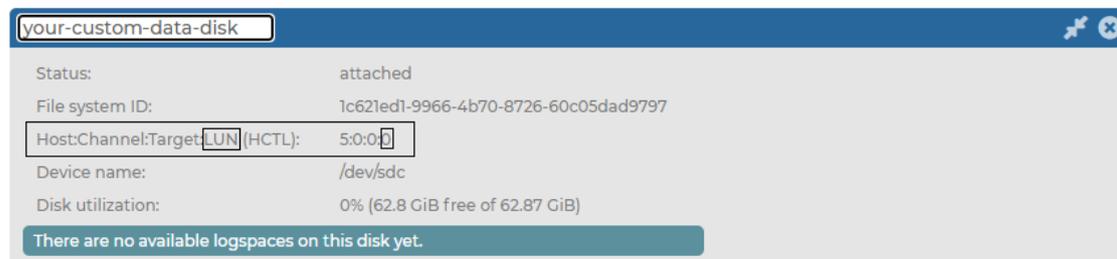
After successfully attaching the data disk, it appears under the **Disks** of the virtual machine of your SSB device.

Figure 144: Home > All resources> <your-virtual-machine-in-microsoft-azure> > Settings > Disk - The new data disk attached to the virtual machine of your SSB device



NOTE: The **LUN** (Logical Unit Number) value of your attached data disk corresponds to the last digit of the **Host:Channel:Target:LUN (HCTL)** of your custom cloud service provider data disk on the SSB side.

Figure 145: Log > Disks> <your-custom-data-disk> - The LUN value corresponds to the last digit of the HCTL on the SSB side



- Restart the virtual machine for your SSB device, add the newly created Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and then assign the SSB logspace of your choice to the newly added custom cloud service provider data disk.

Adding a new custom cloud service provider data disk to your SSB configuration on the SSB side

If you need more disk space than the disk size of your default data disk available in syslog-ng Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

This section describes how you can add a new custom cloud service provider data disk to your syslog-ng Store Box (SSB) configuration on the SSB side.

For more information about how you can add a custom cloud service provider data disk on the cloud service provider side, see [Adding a new custom cloud service provider data disk on the cloud service provider side](#).

For more information about assigning the SSB logspace of your choice to a custom cloud service provider data disk, see [Assigning the SSB logspace of your choice to a custom cloud service provider data disk](#).

For more information about the data disk information for your custom cloud service provider data disks, see [Data disk information for your custom cloud service provider data disks](#).

Prerequisites

When adding a new custom cloud service provider data disk to your SSB configuration on the SSB side, consider the following prerequisites in advance:

- If you want to add a custom cloud service provider data disk to your SSB configuration, you must [add a new custom cloud service provider data disk on your cloud service provider side](#) before you can add it to your configuration on the SSB side.
- If you want to add custom cloud service provider data disks to your configuration in SSB, you must add the custom cloud service provider data disk under **Log > Disks > <your-custom-data-disk>**, and commit your changes. After that, the newly added custom cloud service provider data disk will be available for selection in the dropdown list under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**.

Limitations

When adding a new custom cloud service provider data disk to your SSB configuration on the SSB side, consider the following limitations in advance:

CAUTION:

HAZARD OF DATA LOSS!

Although it is possible to decrease the size of Microsoft Azure managed disks, One Identity does not support decreasing the size of custom cloud service provider data disks already added to your SSB configuration. One Identity only supports increasing the size of such

custom cloud service provider data disks after adding them to your SSB configuration.

- **NOTE:** Resizing your custom cloud service provider data disk on the Microsoft Azure side is quick, but your SSB configuration must resize the file system on your SSB side to match the resizing on the Microsoft Azure side. As a result, depending on the original size of the custom cloud service provider data disk and on the size you will expand it to, starting the virtual machine for your SSB device after you expand the size of your custom cloud service provider data disk may take a long time.
- **NOTE:** If you fail to assign the newly created SSB logspace of your choice to the new custom cloud service provider data disk, SSB will not automatically detect and use the new custom cloud service provider data disk to store messages on it.
- You can only configure custom cloud service provider data disks for logspaces that you newly create. One Identity does not support configuring custom cloud service provider data disks in logspaces that you configured and committed earlier.
- In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.
- If you want to add custom cloud service provider data disks to your configuration in SSB, you must add the custom cloud service provider data disk under **Log > Disks > <your-custom-data-disk>**, and commit your changes. After that, the newly added custom cloud service provider data disk will be available for selection in the dropdown list under **Log > Logspaces > <your-new-logspace-with-custom-data-disk> > Storage > Disk**.
- Transporting custom cloud service provider data disks between SSB devices is not supported.
- You can only add, remove, or resize a custom cloud service provider data disk in your configuration after shutting your SSB device down first. Hot swaps are not supported.
- You can only add an Azure managed disk as a custom cloud service provider data disk to your SSB on the Microsoft Azure portal side.
- SSB only accepts and detects empty disks as custom cloud service provider data disks. As a result, you can only add a newly created Azure managed disk as a custom cloud service provider data disk to your SSB configuration. Adding an already existing Azure managed disk (by clicking **Attach existing disks** on the Microsoft Azure portal) is not supported.
- One Identity does not support adding Azure managed disks to your SSB logspaces as a shared disk. You can only add one Azure managed disk to one SSB device.
- The number of managed disks you can add on the Microsoft Azure portal depends on the type of Microsoft Azure subscription you are using.

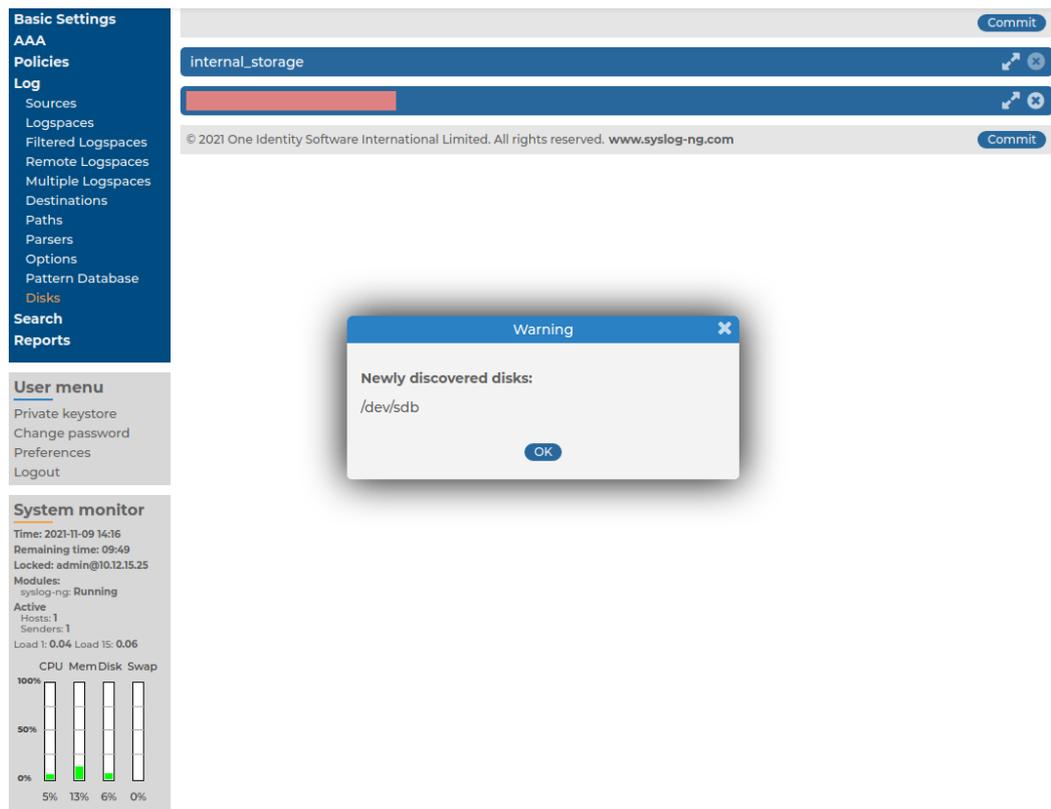
Adding a new custom cloud service provider data disk to your SSB configuration on the SSB side

After [adding a new custom cloud service provider data disk on your cloud service provider side](#), you can add the new custom cloud service provider data disk on the SSB side.

To add a new custom cloud service provider data disk to your SSB configuration,

1. **On the SSB side, shut down your SSB device**
 - a. Navigate to **Basic Settings > System**.
 - b. Under **System control**, click **Shutdown**.
2. **On the cloud service provider side, complete the following steps:**
 - a. Log in to your cloud service provider's portal (in this case, the [Microsoft Azure portal](#)).
 - b. Wait until the power state of your SSB device's virtual machine is **Stopped**.
For more information about power states in Microsoft Azure, see [VM power states](#) in the *Microsoft Azure online tutorial documentation*.
 - c. Deallocate the virtual machine for your SSB device, then wait until its power state is **Stopped (deallocated)**.
 - d. To [attach a new data disk](#) with the size of your choice to the virtual machine of your SSB device, follow the instructions of the *Microsoft Azure online tutorial documentation*.
 - e. Start the virtual machine for your SSB device.
3. **In the virtual machine for your SSB device, complete the following steps on the SSB web interface:**
 - a. Navigate to **Log > Disks**.
SSB will automatically detect newly added custom cloud service provider data disks and display them:

Figure 146: Log > Disks - New disks detected



- b. Enter a name for your new custom cloud service provider data disk, and click **Commit**.

Figure 147: Log > Disks > <your-custom-data-disk>



- c. Navigate to **Log > Logspaces**, and assign the SSB logspace of your choice to the custom cloud service provider data disk.

NOTE: If you fail to assign the newly created SSB logspace of your choice to the new custom cloud service provider data disk, SSB will not automatically detect and use the new custom cloud service provider data disk to store messages on it.

NOTE: After adding a newly created custom cloud service provider data disk to your SSB configuration, and assigning the SSB logspace of your choice to it, removing the custom cloud service provider data disk from your SSB configuration has a strict order, especially if there are referenced objects tied to the custom cloud service provider data disk. For more information about the process to remove such custom cloud service provider data disks from your SSB configuration, see [Removing a custom cloud service provider data disk from your SSB configuration](#).

NOTE: For newly created logspaces, selecting a data disk for your logspace is required.

Data disk information for your custom cloud service provider data disks

If you need more disk space than the disk size of your default data disk available in syslog-ing Store Box (SSB), from version 6.9.0, you can also add a custom cloud service provider data disk to your configuration. From 6.9.0, SSB supports adding a Microsoft Azure managed disk to your SSB configuration as a custom cloud service provider data disk, and assigning the SSB logspace of your choice to it.

In version 6.9.0, SSB supports adding Microsoft Azure managed disks of up to 32,767 GiB storage space to your SSB device. For more information about managed data disks in Microsoft Azure, see [Disk type comparison](#) in the *Microsoft Azure online documentation*.

This section describes the data disk information available for custom cloud service provider data disks on the web interface.

For more information about custom cloud service provider data disks on the cloud service provider side, see [Adding a new custom cloud service provider data disk on the cloud service provider side](#).

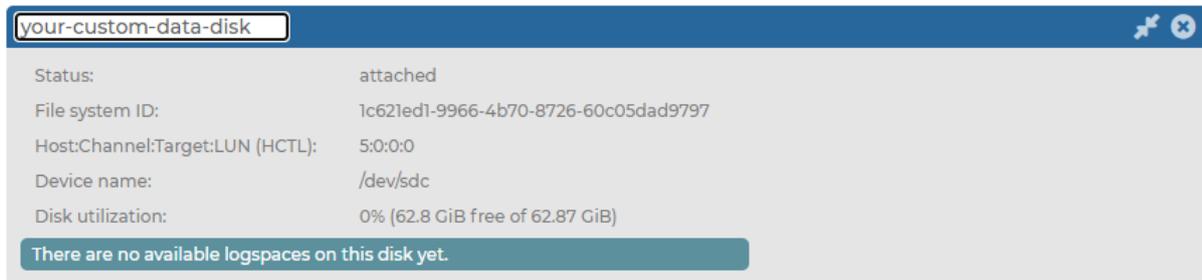
For more information about custom cloud service provider data disks on the side, see [Adding a new custom cloud service provider data disk to your SSB configuration](#).

For more information about assigning the SSB logspace of your choice to a custom cloud service provider data disk, see [Assigning the SSB logspace of your choice to a custom cloud service provider data disk](#).

Data disk information for your custom cloud service provider data disks

After adding the new custom cloud service provider data disk to your SSB configuration, the following data disk information will be displayed under **Log > Disks > <your-custom-data-disk>**:

Figure 148: Log > Disks > <your-custom-data-disk> – Data disk information for your custom cloud service provider data disks



- **Status:** The connectivity status of the custom cloud service provider data disk between the SSB side and the cloud service provider side.

The possible status states are the following:

- **attached:** The custom cloud service provider data disk is successfully attached to the virtual machine for your SSB device on the cloud service provider (in this case, Microsoft Azure) side.
- **missing:** The custom cloud service provider data disk is committed to your SSB configuration, but the SSB device can not detect or access it on the cloud service provider side for some reason (for example, it was removed on the Microsoft Azure side earlier, or there is a connectivity issue with the Microsoft Azure portal).
- **File system ID:** The unique identifier for your file system on the SSB side.
- **Host:Channel:Target:LUN (HCTL):** The locally unique identifier for your custom cloud service provider data disk, which consists of the colon-separated data of the **Host**, the **Channel**, the **Target**, and the **LUN** (Logical Unit Number) value.

The last digit of the custom cloud service provider data disk's **Host:Channel:Target:LUN (HCTL)** corresponds to the attached custom cloud service provider data disk's **LUN** value on the Microsoft Azure portal:

Figure 149: Log > Disks > <your-custom-data-disk> – Data disk information for your newly added custom cloud service provider data disks

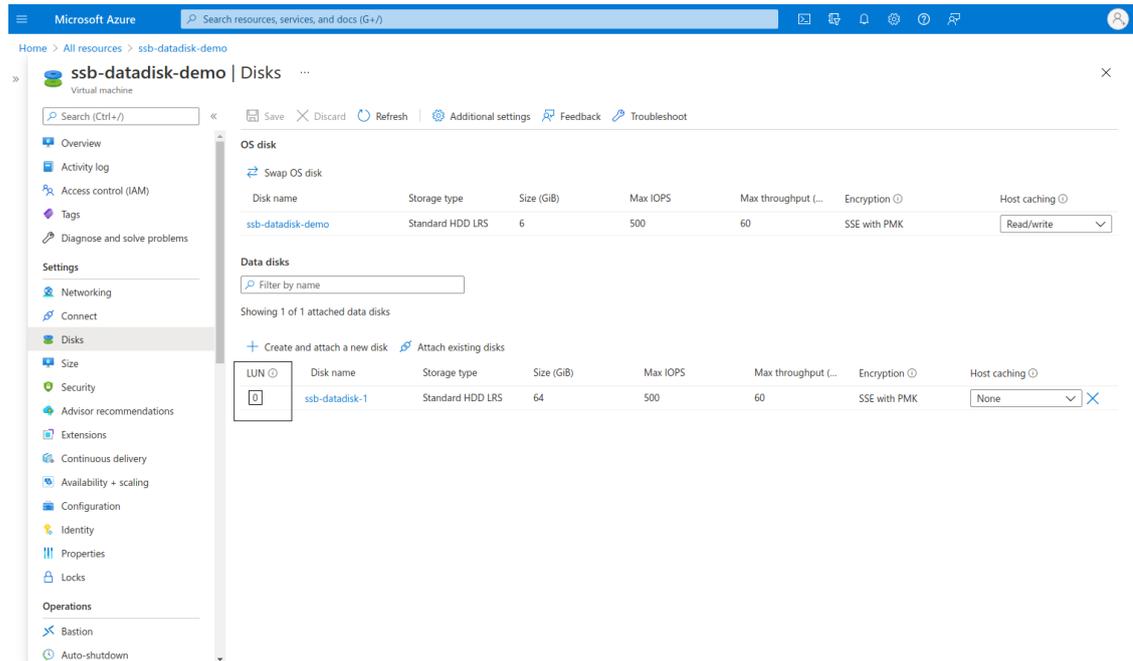
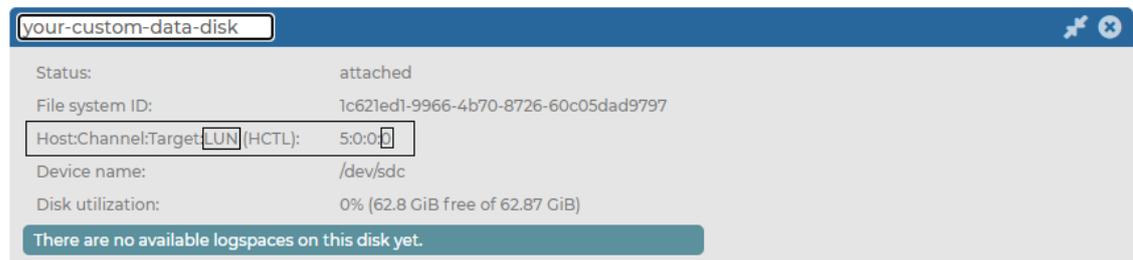
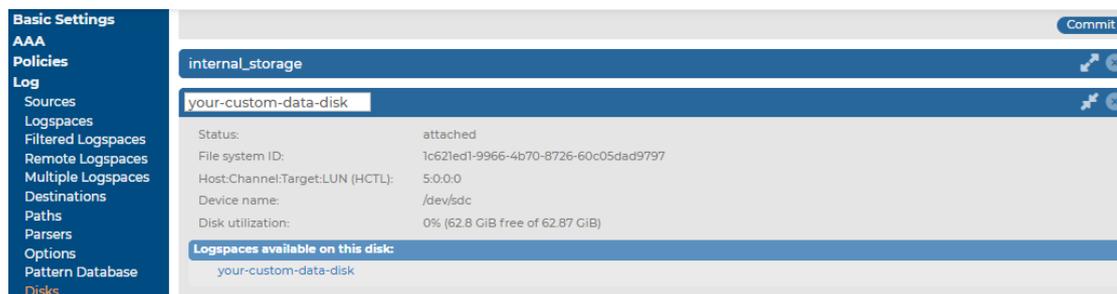


Figure 150: Log > Disks > <your-custom-data-disk> – Data disk information for your custom cloud service provider data disks, with no SSB logspace assigned to them yet



- **Device name:** The navigation path to your device on the SSB side.
- **Disk utilization:** The utilization ratio of your custom data disk (in percentage), followed by the exact amount of free space, and the full size of the disk, in brackets.
- **(Optional) Onscreen information about logspaces assigned to the data disk:** When you have not [assigned the SSB logspace of your choice to the custom cloud service provider data disk](#) yet, you will see the **There are no available logspaces on this disk yet.** onscreen information displayed on the SSB web interface.
- **(Optional) Logspaces available on this disk:** When you have already assigned the SSB logspace of your choice to the custom cloud service provider data disk, the assigned logspaces will be listed here.

Figure 151: Log > Disks > <your-custom-data-disk> – Data disk information for your custom cloud service provider data disks, with the SSB logspace of your choice assigned to them



Removing a custom cloud service provider data disk from your SSB configuration

This section describes the required procedure if you want to remove a previously added custom cloud service provider data disk (for example, a Microsoft Azure managed disk) from your syslog-ng Store Box (SSB) configuration.

For more information about adding a custom cloud service provider data disk to your SSB configuration, see [Adding a new custom cloud service provider data disk to your SSB configuration](#).

Prerequisites

When removing a custom cloud service provider data disk from your SSB configuration, consider the following prerequisites in advance:

- One Identity recommends that you archive or backup all data from the logspaces that you want to delete from your custom cloud service provider data disk.
- If there are any referenced objects (for example, logspaces, filtered logspaces, multiple logspaces, remote logspaces, and logpaths) connected to your custom cloud service provider data disk, you must disconnect them from your custom cloud service provider data disk before removing it from your SSB configuration.

Limitations

When removing a custom cloud service provider data disk from your SSB configuration, consider the following limitations in advance:

- **NOTE:** After adding a newly created custom cloud service provider data disk to your SSB configuration, and attaching it to a logspace, removing the custom data disk from your SSB configuration has a strict order, especially if there are referenced objects connected to the custom cloud service provider data disk. One Identity

recommends that you always complete the following steps in this particular order when removing a custom cloud service provider data disk from your SSB configuration.

- You can not remove a custom cloud service provider data disk from your SSB configuration if there are any referenced objects connected to it. In these cases, you must disconnect them from your custom cloud service provider data disk before removing it from your SSB configuration.

Removing a custom cloud service provider data disk from your SSB configuration

If you have carefully considered the prerequisites and limitations, complete the following steps in this exact same order.

To remove a custom cloud service provider data disk from your SSB configuration,

- 1. On the SSB side, archive or backup your logspace data, remove the logspaces from your configuration, then shut down your SSB device.**
 - a. Navigate to **Log > Logspaces**.
 - b. (Optional) Archive or backup all data from the logspaces that you want to delete from your custom cloud service provider data disk.
 - c. Delete all logspaces on your custom cloud service provider data disk and click **Commit**.
 - d. Navigate to **Log > Disks**, delete the custom cloud service provider data disk of your choice, and click **Commit**.
 - e. Navigate to **Basic Settings > System**.
 - f. Under **System control**, click **Shutdown**.
- 2. On the Microsoft Azure side, stop, deallocate, and detach your custom cloud service provider data disk, then restart the virtual machine for your SSB device.**
 - a. Log in to your cloud service provider's portal (in this case, the [Microsoft Azure portal](#)).
 - b. Wait until the power state of your SSB device's virtual machine is **Stopped**.
For more information about power states in Microsoft Azure, see [VM power states](#) in the *Microsoft Azure online tutorial documentation*.
 - c. Deallocate the virtual machine for your SSB device, then wait until its power state is **Stopped (deallocated)**.
 - d. Follow the instructions of the *Microsoft Azure online tutorial documentation* to [detach your custom data disk](#) from the virtual machine for your SSB device.
 - e. Restart the virtual machine for your SSB device.

Increasing the size of a custom cloud service provider data disk that you use in your SSB configuration

This section describes how you can increase the size of a custom cloud service provider data disk that you use in your syslog-ng Store Box (SSB) configuration.

Prerequisites

When increasing the size of a custom cloud service provider data disk that you use in your SSB configuration, consider the following prerequisites in advance:

- A custom cloud service provider data disk, previously [added on your cloud service provider side](#), then [added to your SSB configuration](#), and [the SSB logspace of your choice assigned to the custom cloud service provider data disk](#).
- A cloud service provider subscription that allows increasing the size of your custom cloud service provider data disk on the cloud service provider side.

Limitations

When increasing the size of a custom cloud service provider data disk that you use in your SSB configuration, consider the following limitations in advance:

⚠ CAUTION:

HAZARD OF DATA LOSS!

Although it is possible to decrease the size of Microsoft Azure managed disks, One Identity does not support decreasing the size of custom cloud service provider data disks already added to your SSB configuration. One Identity only supports increasing the size of such custom cloud service provider data disks after adding them to your SSB configuration.

Increasing the size of a custom cloud service provider data disk that you use in your SSB configuration

If the custom cloud service provider data disk that you recently added to your SSB configuration does not have enough disk space, you can increase its size on the cloud service provider side.

To increase the size of a custom cloud service provider data disk that you use in your SSB configuration,

1. **On the SSB side, shut down your SSB device**
 - a. Navigate to **Basic Settings > System**.
 - b. Under **System control**, click **Shutdown**.
2. **On the cloud service provider side, complete the following steps:**
 - a. Log in to your cloud service provider's portal (in this case, the [Microsoft Azure portal](#)).
 - b. Wait until the virtual machine of your SSB device is in the **Stopped** power state.

For more information about the power states of virtual machines in Microsoft Azure, see [VM power states](#) in the *Microsoft Azure online tutorial documentation*.
 - c. Deallocate the virtual machine for your SSB device, then wait until it is in the **Stopped (deallocated)** power state.
 - d. Follow the instructions of [Resize a managed disk in the Azure portal](#) in the *Microsoft Azure online tutorial documentation* to increase the size of your custom cloud service provider data disk.

Figure 152: Home > All resources > <your-virtual-machine-in-microsoft-azure> > <your-data-disk-in-microsoft-azure> - The Size + performance page for your managed disk in Microsoft Azure

The screenshot shows the Microsoft Azure portal interface for a managed disk. The breadcrumb navigation is: Home > All resources > ssb-datadisk-demo > ssb-datadisk-1 > ssb-datadisk-1 | Size + performance. A notification states: "changes to the disk size can be made only when the disk is unattached or the managing virtual machine(s) are deallocated." The current Disk SKU is "Standard HDD (locally-redundant storage)".

Size	Disk tier	Provisioned IOPS	Provisioned through...	Max Shares	Max burst IOPS	Max burst throughput
32 GiB	S4	500	60	-	-	-
64 GiB	S6	500	60	-	-	-
128 GiB	S10	500	60	-	-	-
256 GiB	S15	500	60	-	-	-
512 GiB	S20	500	60	2	-	-
1024 GiB	S30	500	60	5	-	-
2048 GiB	S40	500	60	5	-	-
4096 GiB	S50	500	60	5	-	-
8192 GiB	S60	1300	300	5	-	-
16384 GiB	S70	2000	500	5	-	-
32767 GiB	S80	2000	500	5	-	-

At the bottom, the "Custom disk size (GiB)" is set to 64. Buttons for "Resize" and "Discard" are visible.

- e. Start the virtual machine for your SSB device.

NOTE: Resizing your custom cloud service provider data disk on the Microsoft Azure side is quick, but your SSB configuration must resize the file system on

your SSB side to match the resizing on the Microsoft Azure side. As a result, depending on the original size of the custom cloud service provider data disk and on the size you will expand it to, starting the virtual machine for your SSB device after you expand the size of your custom cloud service provider data disk may take a long time.

Forwarding messages from SSB

The syslog-ng Store Box (SSB) appliance can forward log messages to remote destinations. The remote destination can be an SQL database running on a remote server, a syslog or log analyzing application running on a remote server, or a Hadoop Distributed File System (HDFS) destination.

- To forward messages to a remote SQL database, complete the procedure in [Forwarding log messages to SQL databases](#) on page 266. Currently Oracle, Microsoft SQL (MSSQL), MySQL, and PostgreSQL databases are supported.
- To forward messages to a remote server, complete the procedure in [Forwarding log messages to remote servers](#) on page 271.

To forward messages to an HDFS destination, complete the procedure in [Forwarding log messages to HDFS destinations](#) on page 300.

Forwarding log messages to SQL databases

This section describes how to forward log messages from syslog-ng Store Box (SSB) to a remote SQL database server.

Tested SQL destinations:

SSB6.9.0 was tested with the following database servers:

- MS SQL (with "select @@version")

```
Microsoft SQL Server 2005 - 9.00.5057.00 (Intel X86) Mar 25 2011
13:50:04 Copyright (c) 1988-2005 Microsoft Corporation Standard Edition
on Windows NT 5.2 (Build 3790: Service Pack 2)
```

- PostgreSQL (with "select version()")

```
PostgreSQL 8.3.15 on i486-pc-linux-gnu, compiled by GCC cc (GCC) 4.2.4
(Ubuntu 4.2.4-1ubuntu4)
```

- MySQL (with "select version()")

```
5.0.51a-3ubuntu5.8-log
```

- Oracle (with "SELECT * FROM V\$VERSION;")

```
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit  
Production  
PL/SQL Release 11.2.0.4.0 - Production  
"CORE 11.2.0.4.0      Production"  
TNS for Linux: Version 11.2.0.4.0 - Production  
NLSRTL Version 11.2.0.4.0 - Production
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit  
Production  
PL/SQL Release 12.1.0.2.0 - Production  
"CORE 12.1.0.2.0      Production"  
TNS for Linux: Version 12.1.0.2.0 - Production  
NLSRTL Version 12.1.0.2.0 - Production
```

To forward log messages from SSB to a remote SQL database server

1. To create a new remote destination, navigate to **Log > Destinations** and select .
2. Enter a name for the destination.
NOTE: This name will be used in the name of the database tables created by SSB. For compatibility reasons, it can contain only numbers, lowercase characters, and the underscore (_) character, for example example_database_destination.
3. Select **Database Server**.

Figure 153: Log > Destinations — Creating database destinations

The screenshot shows the configuration page for a new database destination. The left sidebar contains navigation menus for Basic Settings, AAA, Policies, Log, Paths, Parsers, Options, Pattern Database, Search, and Reports. The main content area is titled 'example_db_destination' and includes a 'Commit' button in the top right. The configuration is organized into several sections:

- Destination type:** Radio buttons for 'Database server' (selected), 'Remote host', and 'HDFS destination'.
- Database access:** Fields for 'Database type' (PostgreSQL), 'Address' (db.example.com), 'Port' (5432), 'Username' (ssb), and 'Password' (masked). A 'Test connection' button is present.
- Table schema:** Radio buttons for 'Legacy', 'Full' (selected), and 'Custom columns'.
- Database handling:** Fields for 'Flush lines' (1000) and 'Retention time' (31 days).
- Table rotation:** Radio buttons for 'Daily', 'Monthly' (selected), and 'Custom'.
- Access control:** A 'Group' field.

At the bottom left, there is a 'System monitor' section showing system status: Time: 2018-09-25 18:01, Remaining time: 11:57:28, Locked: admin@10.12.16.250, Modules: syslog-ng: Running, Active Hosts: 3, Senders: 2, Load 1: 0.00 Load 15: 0.00. Below this are three bar charts for CPU, MemDisk, and Swap usage.

4. Select the type of the remote database from the **Database type** field.
5. Enter the IP address or hostname of the database server into the **Address** field. If the database is running on a non-standard port, adjust the **Port** setting.
6. Enter the name and password of the database user account used to access the database into the **Username** and **Password** fields, respectively. This user needs to have the appropriate privileges for creating new tables.

NOTE: The syslog-ng Store Box (SSB) appliance accepts passwords that are not longer than 150 characters. The following special characters can be used: !"#%&'()*+,-./:;<=>?@[]^_`{|}

7. Enter the name of the database that will store the log messages into the **Database name** field.
8. *Optional step:* Enter the number of log message lines into the **Flush lines** field that SSB should wait before sending them off in a single batch. Setting this number high increases throughput as fully filled frames are sent to the network. However, it also increases message latency.

NOTE: **Flush lines** is in connection with the **Output memory buffer** value. (To set the **Output memory buffer** value, navigate to **Log > Destinations**). The

value of **Output memory buffer** has to be greater than or equal to the value of **Flush lines**.

9. SSB will automatically start a new table for every day or every month. Optionally, you can also create custom tables. Select the table naming template from the **Table rotation** field.
10. Select which columns should SSB insert into the database. You can use one of the predefined templates, or select **Custom columns** to create a custom template. The available templates are described in [SQL templates in SSB](#) on page 270.
11. SSB can automatically delete older messages and tables from the database. By default, messages are deleted after one month. Adjust the **Retention time** as needed for your environment.
12. The logs stored in the database can be accessed using the search interface of SSB. Enter the name of the usergroup who can access the logs into the **Access control > Group** field. To add more groups (if needed), click .
13. The time stamps of most log messages is accurate only to the second. The syslog-ng Store Box(SSB) appliance can include more accurate time stamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the `frac_digits()` parameter of syslog-ng.
14. If the server and SSB are located in a different timezone and you use the Legacy message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
15. Set the size of the disk buffer (in Megabytes) in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the `log_disk_fifo_size()` parameter of syslog-ng.

Note that SSB does not pre-allocate the hard disk required for the disk buffer, so make sure that the required disk space is available on SSB. For details on creating archiving policies and adjusting the disk-fillup prevention, see [Archiving and cleanup](#) and [Preventing disk space fill up](#).

Example: Calculating disk buffer size

The size of the disk buffer you need depends on the rate of the incoming messages, the size of the messages, and the length of the network outage that you want to cover. For example:

- SSB is receiving 15000 messages per second
- On the average, one message is 250 bytes long
- You estimate that the longest time the destination will be unavailable is 4 hours

In this case, you need a disk buffer for $250 \text{ [bytes]} * 15000 \text{ [messages per second]} * 4*60*60 \text{ [seconds]} = 54000000000 \text{ [bytes]}$, which is 54000 Megabytes (in other words, a bit over 50 GB).

16. Click **Commit**.
17. To start sending messages to the destination, include the new destination in a logpath. For details, see [Log paths: routing and processing messages](#).
18. To test if the database is accessible, select **Test connection**.

SQL templates in SSB

The following sections describe the SQL templates available in syslog-ng Store Box(SSB):

- [Legacy](#)
- [Full](#)
- [Custom](#)

The Legacy template

The **Legacy** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- *insert_time*: The date when syslog-ng Store Box(SSB) received the message in Unixtime format.
- *rule_id*: ID of the pattern database rule that matched the message.
- *__row_id*: Identifier of the row.
- *date_time*: The date the message was sent in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- *facility*: The facility that sent the message.
- *priority*: The priority level of the message.
- *host*: The IP address or hostname of the host where the message was generated.
- *program*: The name of the application that generated the message.
- *pid*: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- *message*: The text of the log message.

The `insert_time`, `rule_id`, `date_time`, `facility`, `host`, and `program` columns are indexed.

The Full template

The **Full** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- *insert_time*: The date when syslog-ng Store Box(SSB) received the message in Unixtime format.
- *rule_id*: ID of the pattern database rule that matched the message.
- *__row_id*: Identifier of the row.
- *date_time*: The date the message was sent in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- *facility*: The facility that sent the message.
- *priority*: The priority level of the message.
- *sourceip*: The IP address of the host that sent the message.
- *host*: The IP address or hostname of the host where the message was generated.
- *program*: The name of the application that generated the message.
- *pid*: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- *message*: The text of the log message.

The `insert_time`, `rule_id`, `date_time`, `facility`, `host`, `sourceip`, and `program` columns are indexed.

The Custom template

The **Custom** template allows you to specify the columns to use. Enter a name for the column, select its type, and specify its content using macros. For details on using macros, see [Macros of syslog-ng PE](#) in the *syslog-ng PE Administration Guide* and [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*.

Select the **Indexed** option if you want the database to index the column.

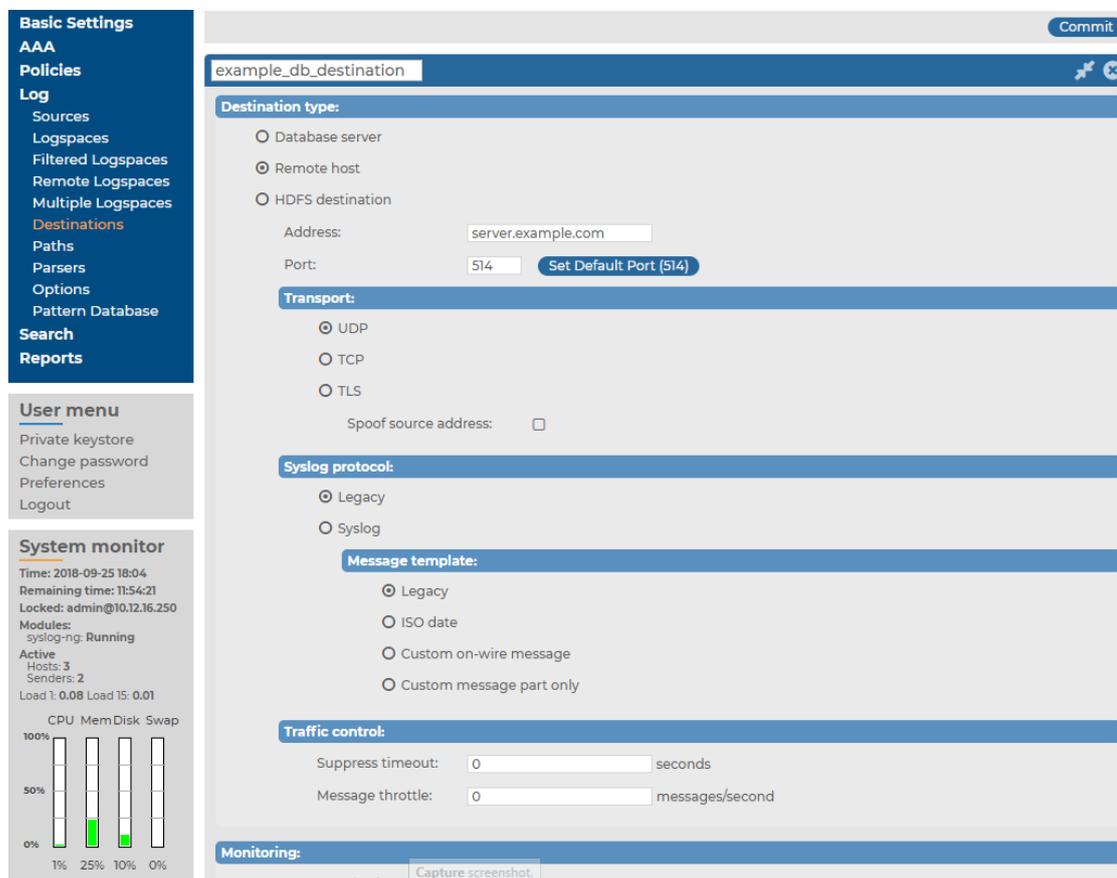
Forwarding log messages to remote servers

This section describes how to forward messages from syslog-ng Store Box(SSB) to a remote server.

To forward messages from SSB to a remote server

1. Navigate to **Log > Destinations** and select  to create a new remote destination.
2. Select **Remote host**.

Figure 154: Log > Destinations – Creating server destinations



The screenshot shows the SSB configuration interface for creating a new remote destination. The left sidebar contains navigation menus for Basic Settings, AAA, Policies, Log, User menu, and System monitor. The main content area is titled "example_db_destination" and includes a "Commit" button. The configuration is organized into several sections:

- Destination type:** Radio buttons for Database server, Remote host (selected), and HDFS destination.
- Address:** Text input field containing "server.example.com".
- Port:** Text input field containing "514" and a "Set Default Port (514)" button.
- Transport:** Radio buttons for UDP (selected), TCP, and TLS. A "Spoof source address" checkbox is also present.
- Syslog protocol:** Radio buttons for Legacy (selected) and Syslog.
- Message template:** Radio buttons for Legacy (selected), ISO date, Custom on-wire message, and Custom message part only.
- Traffic control:** Input fields for "Suppress timeout: 0 seconds" and "Message throttle: 0 messages/second".
- Monitoring:** A "Capture screenshot" button.

3. Enter the IP address or hostname of the remote server into the **Address** field. Enter the port where the server is accepting syslog messages into the **Port** field.

Note that the **Address** and **Port** pair must be unique for each remote destination.

4. Select the network protocol used to transfer the log messages from the **Transport** field. The UDP, TCP, and the encrypted TLS protocols are available. The UDP and TLS protocols have additional parameters.

When forwarding messages using UDP, the remote host will see the messages as if they originated from SSB. Select the **Spoof source address** option to make them seem to originate from their original sender.

CAUTION:

When using the **Spoof source address** option, SSB automatically truncates long messages to 1024 bytes, regardless of the **Log > Options > Message size setting**.

For TLS, select a method to verify the identity of the remote host. The following options are available:

- **None:** Do not request a certificate from the remote host, and accept any certificate if the host sends one.
- **Optional trusted:** If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.
- **Optional untrusted:** Accept any certificate shown by the remote host. Note that the host must show a certificate.
- **Required trusted (default setting):** Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See [Uploading external certificates to SSB](#) for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
- **Required untrusted:** SSB requests a certificate from the remote host, and rejects the connection if no certificate is received. However, SSB accepts the connection if:
 - the certificate is not valid (expired), or
 - the Common Name of the certificate does not contain the domain name or the IP address of the host.

NOTE: Consult the documentation of the remote server application to determine which protocols are supported.

UDP is a highly unreliable protocol and a high amount of messages may be lost without notice during the transfer. Use TCP or TLS instead whenever possible.

5. Select the syslog protocol to use from the **Syslog protocol** field.
 - To use the legacy BSD-syslog protocol described in RFC 3164, select **Legacy** and specify the message template to use. Select **Legacy** to use the message format described in the RFC, or **ISO date** to replace the original time stamp with an ISO8061 compliant time stamp that includes year and timezone information. To customize the format of the message contents using macros, select **Custom message part only**, or **Custom on-wire message** to completely reformat the message (including the headers). For details on using macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*. If you have no special requirements, use the **ISO date** template.
 - Use the new IETF-syslog protocol. Note that most syslog applications and devices currently support only the legacy protocol. Consult the documentation of the remote server application to determine which protocols are supported. If you need, you can customize the contents of the message using macros. Note that for the IETF-syslog protocol, the header cannot be customized. For details on using macros, see .

6. If SSB would send several messages with identical content to the destination, it can send only a single message and a line `Last message repeated n times..` Enter the number of seconds to wait for identical messages into the **Suppress timeout** field. This option corresponds to the `suppress()` parameter of `syslog-ng`.
7. To limit the maximum number of messages sent to the destination per second, enter the maximum number of messages into the **Message throttle** field. Use this output-rate-limiting functionality only when using `disk-buffer` as well to avoid the risk of losing messages. Specifying 0 or a lower value sets the output limit to unlimited. This option corresponds to the `throttle()` parameter of `syslog-ng`.
8. The time stamps of most log messages is accurate only to the second. The `syslog-ng` Store Box(SSB) appliance can include more accurate time stamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the `frac_digits()` parameter of `syslog-ng`.
9. If the server and SSB are located in a different timezone and you use the Legacy message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
10. Set the size of the disk buffer (in Megabytes) in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the `log_disk_fifo_size()` parameter of `syslog-ng`.

Note that SSB does not pre-allocate the hard disk required for the disk buffer, so make sure that the required disk space is available on SSB. For details on creating archiving policies and adjusting the disk-fillup prevention, see [Archiving and cleanup](#) and [Preventing disk space fill up](#).

Example: Calculating disk buffer size

The size of the disk buffer you need depends on the rate of the incoming messages, the size of the messages, and the length of the network outage that you want to cover. For example:

- SSB is receiving 15000 messages per second
- On the average, one message is 250 bytes long
- You estimate that the longest time the destination will be unavailable is 4 hours

In this case, you need a disk buffer for $250 \text{ [bytes]} * 15000 \text{ [messages per second]} * 4 * 60 * 60 \text{ [seconds]} = 54000000000 \text{ [bytes]}$, which is 54000 Megabytes (in other words, a bit over 50 GB).

11. Click **Commit**.
12. To start sending messages to the destination, include the new destination in a `logpath`. For details, see [Log paths: routing and processing messages](#).

Forwarding log messages to the Microsoft Azure Sentinel cloud

This section describes how to forward messages from syslog-ng Store Box (SSB) to [the Microsoft Azure Sentinel cloud](#) (Azure Sentinel).

Azure Sentinel is Microsoft's native cloud-based SIEM solution. Beside Microsoft's own cloud services, it can accept log messages from external sources. Microsoft provides 26 predefined Data connectors and a HTTP Data Collector API for further integrations.

From version 6.7.0, SSB uses the syslog-ng Premium Edition (syslog-ng PE) application's support to post messages to Azure Sentinel over HTTP and Secure HTTP (HTTPS).

For further details about how the application forwards messages to Azure Sentinel, see [sentinel: Sending logs to the Microsoft Azure Sentinel cloud](#) in *the syslog-ng PE Administration Guide*.

For more information about Azure Sentinel, see [Microsoft Azure: Azure Sentinel documentation](#) in *the Microsoft Azure Documentation*.

For more information about Data connectors used in Azure Sentinel, see [Microsoft Azure: Connect data sources](#) in *the Microsoft Azure Documentation*.

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

Prerequisites

This section describes the prerequisites for forwarding messages from syslog-ng Store Box (SSB) to [the Microsoft Azure Sentinel cloud](#) (Azure Sentinel).

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

Prerequisites to using the Azure Sentinel destination

- **WORKSPACE ID** (which will function as the **Workspace id** on the SSB web interface)
- **PRIMARY KEY** (which will function as the **Auth secret** on the SSB web interface)

NOTE: For more information about the **WORKSPACE ID** and the **PRIMARY KEY** on the Azure Sentinel side, see [Getting the required credentials to configure syslog-ng PE as a](#)

Limitations

This section describes the limitations for forwarding messages from syslog-ng Store Box (SSB) to [the Microsoft Azure Sentinel cloud](#) (Azure Sentinel).

- **Messages with HTTP 400 response code will be dropped**

If the message sent to Azure Sentinel is invalid, the Azure Sentinel messaging service will reply with an HTTP 400 response code.

The message can be invalid for either of these reasons:

- A required argument is missing from the message.
- The message size exceeds limits.
- The message itself has an invalid format.

In these cases, SSB cannot successfully send the messages to Azure Sentinel. These messages would prevent SSB from sending further messages to the messaging service, therefore SSB must drop them.

- **Proxy limitations**

If you use a proxy, consider that only HTTP proxies are supported.

Configuring the Azure Sentinel destination: adding a new Azure Sentinel destination

This section describes the first steps of configuring the Azure Sentinel destination of syslog-ng Store Box (SSB), that is, adding the new Azure Sentinel destination on the SSB web interface.

For information about configuring the authentication and workspace settings of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Authentication and workspace settings](#).

For information about configuring the advanced message parameters of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Advanced message parameters](#).

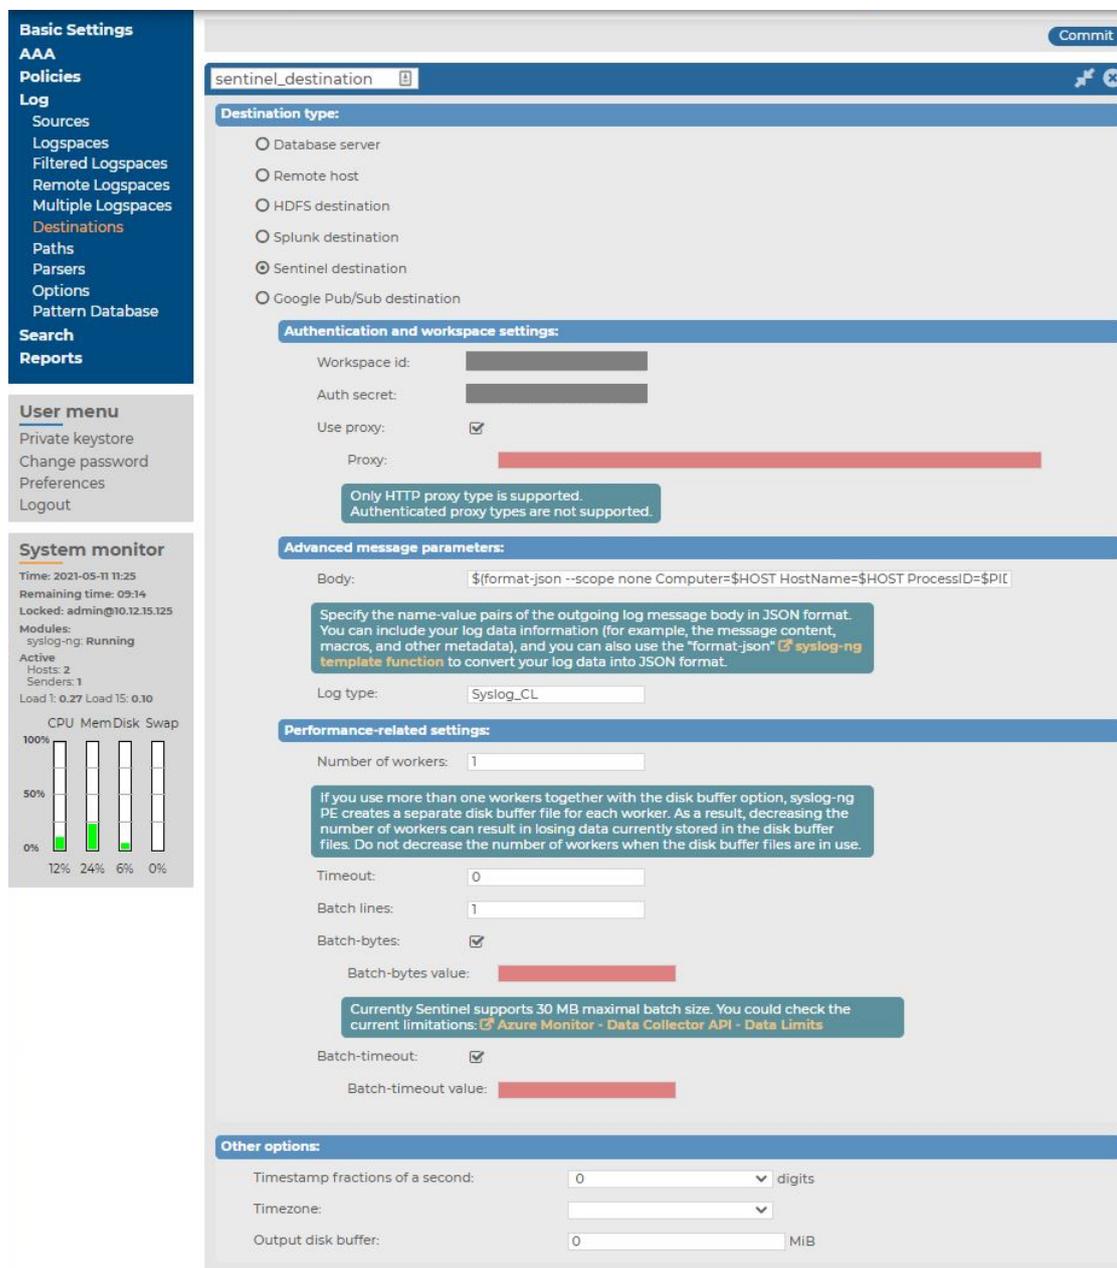
For information about configuring the performance-related settings of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Performance-related settings](#).

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To create your custom Azure Sentinel destination on the SSB web interface

1. Navigate to **Log > Destinations**, and select  to create a new destination.

Figure 155: Log > Destinations > <your-sentinel-destination> – Adding a new Sentinel destination



The screenshot displays the configuration page for a new Sentinel destination in the SSB web interface. The page is titled "sentinel_destination" and includes a "Commit" button in the top right corner. The left sidebar contains navigation menus for "Basic Settings", "User menu", and "System monitor". The main content area is divided into several sections:

- Destination type:** A list of radio buttons with "Sentinel destination" selected.
- Authentication and workspace settings:** Fields for "Workspace id" and "Auth secret", a checked "Use proxy" checkbox, and a "Proxy" field. A note states: "Only HTTP proxy type is supported. Authenticated proxy types are not supported."
- Advanced message parameters:** A "Body" field containing a JSON format string: `${format-json --scope none Computer=$HOST HostName=$HOST ProcessID=$PID}`. A note explains the "format-json" template function. The "Log type" is set to "Syslog_CL".
- Performance-related settings:** A "Number of workers" field set to "1". A note explains that decreasing the number of workers can result in losing data. Other fields include "Timeout" (0), "Batch lines" (1), a checked "Batch-bytes" checkbox, a "Batch-bytes value" field, a note about the 30 MB batch size limit, a checked "Batch-timeout" checkbox, and a "Batch-timeout value" field.
- Other options:** Fields for "Timestamp fractions of a second" (0 digits), "Timezone", and "Output disk buffer" (0 MiB).

2. Under **Destination type**, select **Sentinel destination**.
3. After creating your Azure Sentinel destination, continue customizing it by configuring the following:

- [Authentication and workspace settings](#)
- [Advanced message parameters](#)
- [Performance-related settings](#)

Configuring the Azure Sentinel destination: Authentication and workspace settings

This section describes configuring the authentication and workspace settings after [Configuring the Azure Sentinel destination: adding a new Azure Sentinel destination](#).

For information about configuring the advanced message parameters of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Advanced message parameters](#).

For information about configuring the performance-related settings of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Performance-related settings](#).

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the authentication and workspace settings of your Azure Sentinel destination

1. Navigate to **Log > Destination > <your-sentinel-destination> > Authentication and workspace settings**.

Figure 156: Log > Destinations > <your-sentinel-destination> – Configuring the authentication and workspace settings



2. In the **Workspace id** field, copy the **WORKSPACE ID** from your Azure Sentinel side.

NOTE: The workspace ID is a unique hexadecimal number provided by Microsoft, with the purpose of identifying your Sentinel instance.

3. In the **Auth secret** field, copy the **PRIMARY KEY** from your Azure Sentinel side.

NOTE: The primary key is a Base64-encoded secret provided by Microsoft, with the purpose of identifying your application.

- (Optional) Enable **Use proxy**, and in the **Proxy** field, enter the HTTP proxy address that you want to use.

| **NOTE:** If you have to use a proxy, consider that only HTTP proxies are supported.

Configuring the Azure Sentinel destination: Advanced message parameters

This section describes configuring the advanced message parameters of the Azure Sentinel destination after [configuring the authentication and workspace settings](#).

For information about adding a new Azure Sentinel destination, see [Configuring the Azure Sentinel destination: adding a new Azure Sentinel destination](#).

For information about configuring the performance-related settings of your Azure Sentinel destination, see [Configuring the Azure Sentinel destination: Performance-related settings](#).

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the advanced message parameters of your Azure Sentinel destination

- Navigate to **Log > Destination > <your-sentinel-destination> > Advanced message parameters**.

Figure 157: Log > Destinations > <your-sentinel-destination> – Configuring the advanced message parameters

Advanced message parameters:

Body:

Specify the name-value pairs of the outgoing log message body in JSON format. You can include your log data information (for example, the message content, macros, and other metadata), and you can also use the "format-json" [syslog-ng template function](#) to convert your log data into JSON format.

Log type:

- To customize the message format sent to Azure Sentinel, specify the name-value pairs of the outgoing message in JSON format in the **Body** field.

⚠ CAUTION: Hazard of data loss!

Make sure that the customized message format is accepted by Azure Sentinel. For invalid messages, SSB will receive an HTTP 400 response code and messages with such a response code will be dropped.

For more information on dropping messages with HTTP 400 response code, see [Limitations](#).

For more information on the `$(format-json)` template function, see

3. In the **Log type** field, enter `Syslog_CL`.

NOTE: For more information about log types on the Azure Sentinel side, and how they connect to the functionalities of the syslog-ng Premium Edition (syslog-ng PE) application (and, as a result, to SSB), see [Log types](#) in *the syslog-ng PE Administration Guide*.

Configuring the Azure Sentinel destination: Performance-related settings

This section describes configuring the performance-related settings of the Azure Sentinel destination after configuring [the authentication and workspace settings](#) and [the advanced message parameters](#).

For more information about adding a new Azure Sentinel destination, see [Configuring the Azure Sentinel destination: adding a new Azure Sentinel destination](#).

NOTE: This section and the other Azure Sentinel-related sections in this documentation are based on Azure Sentinel messaging service concepts and terminology. If you do not use the Azure Sentinel messaging service on a regular basis, One Identity recommends that you read [the Azure Sentinel quick-start documentation](#) to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the performance-related settings of your Azure Sentinel destination

1. Navigate to **Log > Destination > <your-sentinel-destination> > Performance-related settings**.

Figure 158: Log > Destinations > <your-sentinel-destination> – Configuring the performance-related settings

Performance-related settings:

Number of workers:

If you use more than one workers together with the disk buffer option, syslog-ng PE creates a separate disk buffer file for each worker. As a result, decreasing the number of workers can result in losing data currently stored in the disk buffer files. Do not decrease the number of workers when the disk buffer files are in use.

Timeout:

Batch lines:

Batch-bytes:

Batch-bytes value:

Currently Sentinel supports 30 MB maximal batch size. You could check the current limitations: [Azure Monitor - Data Collector API - Data Limits](#)

Batch-timeout:

Batch-timeout value:

2. In the **Number of workers** field, set the number of worker threads that you want SSB to use when sending messages to the server.

⚠ CAUTION: Hazard of data loss!

When you use more than one worker threads together with disk-buffering enabled, consider that the syslog-ng PE application behind SSB creates a separate disk-buffer file for each worker thread. This means that decreasing the number of workers can result in losing data currently stored in the disk-buffer files.

To avoid data loss, One Identity recommends that you do not decrease the number of workers when the disk-buffer files are in use.

NOTE: Increasing the number of worker threads can drastically improve the performance of the destination.

3. In the **Timeout** field, specify the timeout (in **seconds**) that you want SSB to wait for an operation to complete, and attempt to reconnect to the server if the configured timeout limit is exceeded.

The default value of the **Timeout** setting is **0**, which means that it is disabled by default.

4. In the **Batch lines** field, specify how many lines you want SSB to flush to a destination in one batch.

NOTE: SSB waits for the configured number of lines to accumulate, and when this number is reached, SSB sends the message lines to the destination in a single batch. For example, if you set **Batch lines** to **100**, SSB waits for 100 message lines before sending them in one batch.

NOTE: Consider the following when configuring the number of batch lines:

- Increasing the number of batch lines increases throughput (because more messages are sent in a single batch), but also increases message latency.
- If the **Batch-timeout** option is disabled, the syslog-ng PE application behind SSB flushes the messages if it has sent the number of messages specified in **Batch lines**, or the queue became empty. If you stop or reload the syslog-ng PE application behind SSB, or if in case of network sources, the connection with the client is closed, the syslog-ng PE behind SSB automatically sends the unsent messages to the destination.
- If the **Batch-timeout** option is enabled and the queue becomes empty, SSB flushes the messages only if **Batch timeout** expires, or if the batch reaches the limit set in **Batch lines**.

NOTE: Depending on your source configuration settings, your batch may not reach the **Batch lines** limit before your queue becomes empty, and SSB forwards your messages.

5. In the **Batch-bytes** field, set the maximum size of payload in a batch (in **bytes**).

NOTE: When configuring **Batch-bytes**, consider the following:

- If the size of the messages reaches this value, the syslog-ng PE application behind SSB sends the batch to Azure Sentinel even if the number of messages is less than the value you configure in the **Batch-bytes** field.
 - If **Batch-timeout** is enabled and the queue becomes empty, SSB flushes the messages only if the value specified in **Batch-timeout** expires, or if the message batch reaches the limit set in the **Batch-bytes** field.
6. (Optional) Select **Batch-timeout**, and in the **Batch-timeout** value field, specify the time SSB waits for **Batch lines** to accumulate in the output buffer.

SSB sends batches to the destinations evenly. The timer starts when the first message arrives to the buffer, so if only few messages arrive, SSB sends messages to the destination once every **Batch timeout** milliseconds at most.

Forwarding log messages to Google Pub/Sub

This section describes how to forward messages from syslog-ng Store Box (SSB) to the [Google Pub/Sub messaging service](#) (Google Pub/Sub).

From version 6.7.0, SSB uses the syslog-ng Premium Edition (syslog-ng PE) application's support to generate your own messaging Google Pub/Sub infrastructure with syslog-ng PE as a "Publisher" entity, utilizing the HTTP REST interface of the service. The Google Pub/Sub destination is an asynchronous messaging service connected to Google's infrastructure.

For more information about Google Pub/Sub's messaging service, see [What Is Pub/Sub?](#) in the *Google Pub/Sub online documentation*.

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

Prerequisites

This section describes the prerequisites to forwarding messages from syslog-ng Store Box (SSB) to the [Google Pub/Sub messaging service](#) (Google Pub/Sub).

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the Google Pub/Sub destination, you must have:

- Appropriate credentials through a Google Cloud Pub/Sub service account.
For more information, see [Building a Pub/Sub messaging system > Create service account credentials](#) or [Managing topics and subscriptions > Resource names](#) in the *Google Pub/Sub online documentation*.
- A Google Pub/Sub topic, and a Google Pub/Sub subscription.
For more information, see [Building a Pub/Sub messaging system > Tutorial setup > Set up your Google Cloud project and Pub/Sub topic and subscriptions](#) in the *Google Pub/Sub online documentation*.

Limitations

This section describes the limitations to using syslog-ng Store Box (SSB) to forward messages to the [Google Pub/Sub messaging service](#) (Google Pub/Sub).

The current implementation of the Google Pub/Sub destination has the following limitations:

- **No message-based acknowledgement**

While Google Pub/Sub acknowledges the batch of received messages, it also sends individual acknowledgement IDs to each message. However, SSB currently does not track individual messages inside Google Pub/Sub. Under normal operational circumstances, the lack of tracking individual messages has no effect on message delivery, and even allows flow control to work properly. However, in case of an error, the only solution is to repeat the entire batch, which can lead to message duplication in case Google Pub/Sub acknowledged part of the previous batch in spite of indicating an overall error.

NOTE: This behavior, called [At-Least-Once delivery](#), means that if an error occurs, it is more acceptable to duplicate messages than to lose any of them.

NOTE: The Google Pub/Sub destination can not fetch logs, only serve as a "Publisher" entity in the Google Pub/Sub service.

- **Messages with HTTP 400 response code will be dropped**

If the message sent to Google Pub/Sub is invalid, the Google Pub/Sub messaging service will reply with an HTTP 400 response code.

The message can be invalid for either of these reasons:

- A required argument is missing from the message.
- The message size exceeds limits.
- The message itself has an invalid format.

In these cases, SSB cannot successfully send the messages to Google Pub/Sub. These messages would prevent SSB from sending further messages to the messaging service, therefore SSB must drop them.

For further details on HTTP error codes of the Google Pub/Sub messaging service, see [Cloud Pub/Sub > Documentation > Reference > Error Codes](#) in the *Google Pub/Sub Reference Documentation*.

- **Proxy limitations**

If you use a proxy, consider that only HTTP proxies are supported.

Configuring the Google Pub/Sub destination: adding a new Google Pub/Sub destination

This section describes the first steps of configuring the Google Pub/Sub destination, that is, adding a new Google Pub/Sub destination on the syslog-ng Store Box (SSB) web interface.

For more information about configuring the authentication and workspace settings of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Authentication and workspace settings](#).

For more information about configuring the advanced message parameters of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Advanced message parameters](#).

For more information about configuring the performance-related settings of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Performance-related settings](#).

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To create your custom Google Pub/Sub destination on the SSB web interface

1. Navigate to **Log > Destinations**, and select  to create a new destination.
2. Under **Destination type**, select **Google Pub/Sub destination**.

Figure 159: Log > Destination — Adding a new Google Pub/Sub destination

The screenshot displays the configuration interface for a new Google Pub/Sub destination. On the left, a navigation menu includes 'Basic Settings', 'AAA', 'Policies', 'Log' (with sub-items like Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options, Pattern Database), 'Search', and 'Reports'. Below the menu is a 'User menu' with options like 'Private keystore', 'Change password', 'Preferences', and 'Logout', and a 'System monitor' showing system status and resource usage (CPU, Mem, Disk, Swap).

The main configuration area is titled 'Destination type:' and has 'Google Pub/Sub destination' selected. Below this are three sections:

- Authentication and workspace settings:** Includes fields for Project, Topic, Credentials (with a 'Choose File' button and 'Upload' button), OAuth 2.0 Scope (set to 'https://www.googleapis.com/auth/cloud-platform'), and Use proxy (checked). A note states: 'Only HTTP proxy type is supported. Authenticated proxy types are not supported.'
- Advanced message parameters:** Includes a 'Body' field with a JSON format template. A help text box explains: 'Specify the name-value pairs of the outgoing log message body in JSON format. You can include your log data information (for example, the message content, macros, and other metadata), and you can also use the "format-json" syslog-ng template function to convert your log data into JSON format.'
- Performance-related settings:** Includes 'Number of workers' (set to 1), 'Timeout' (set to 0), 'Batch lines' (set to 1000), 'Batch-bytes value' (set to 10000000), and 'Batch-timeout' (checked). A note states: 'Currently Google Pub/Sub supports 10 MB (10000000 bytes) maximal batch size. You can check the current limitations at: Quotas and limits | Cloud Pub/Sub Documentation | Google Cloud.'

- After creating your Google Pub/Sub destination, continue customizing it by configuring the following:
 - Authentication and workspace settings
 - Advanced message parameters
 - Performance-related settings

Configuring the Google Pub/Sub destination: Authentication and workspace settings

This section describes how you can configure the authentication and workspace settings after adding a new Google Pub/Sub destination.

For information about configuring the advanced message parameters of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Advanced message parameters](#).

For information about configuring the performance-related settings of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Performance-related settings](#).

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the authentication and workspace settings for your Google Pub/Sub destination

1. Navigate to **Log > Destination > <your-pubsub-destination> > Authentication and workspace settings**.

Figure 160: Log > Destination > <your-pubsub-destination> — Configuring the authentication and workspace settings

The screenshot shows the 'Authentication and workspace settings' configuration page. It features the following elements:

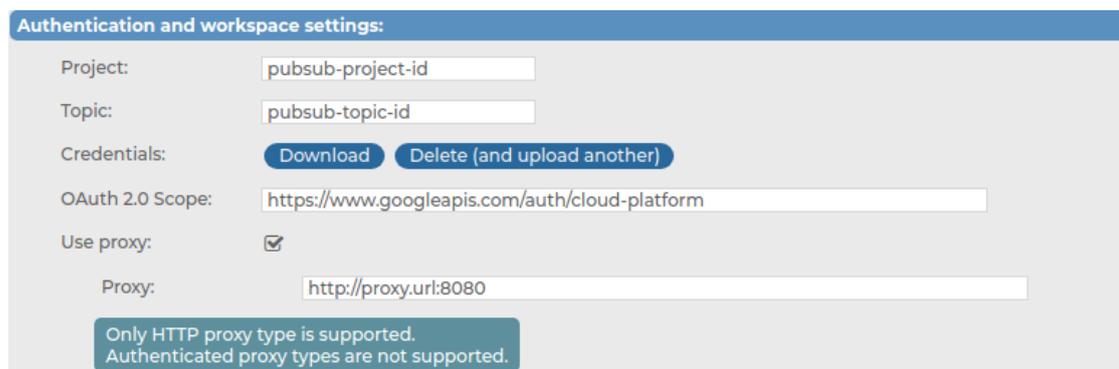
- Project:** A text input field containing 'pubsub-project-id'.
- Topic:** A text input field containing 'pubsub-topic-id'.
- Credentials:** A file selection interface with a 'Choose File' button, a text input field containing 'credentials.json', and an 'Upload' button. Below this is a blue callout box with the text: 'Upload the credential JSON file that you can get from [Google Pub/Sub UI](#)'.
- OAuth 2.0 Scope:** A text input field containing 'https://www.googleapis.com/auth/cloud-platform'.
- Use proxy:** A checked checkbox.
- Proxy:** A text input field containing 'http://proxy.url:8080'.
- A blue callout box at the bottom states: 'Only HTTP proxy type is supported. Authenticated proxy types are not supported.'

2. In the **Project** field, enter the Google Pub/Sub project ID of your choice.
3. In the **Topic** field, enter the Google Pub/Sub topic of your choice.

For more information about setting up Google Pub/Sub projects and topics, see [Building a Pub/Sub messaging system > Tutorial setup > Set up your Google Cloud project and Pub/Sub topic and subscriptions](#) in the *Google Pub/Sub online documentation*.

4. In the **Credentials** field, upload your Google Pub/Sub service account credentials. After you upload your Google Pub/Sub service account credentials JSON file, the web interface will look like this:

Figure 161: Log > Destination > <your-pubsub-destination> — Credentials uploaded



Authentication and workspace settings:

Project: pubsub-project-id

Topic: pubsub-topic-id

Credentials: Download Delete (and upload another)

OAuth 2.0 Scope: https://www.googleapis.com/auth/cloud-platform

Use proxy:

Proxy: http://proxy.url:8080

Only HTTP proxy type is supported.
Authenticated proxy types are not supported.

For more information about creating Google Pub/Sub credentials, see [Building a Pub/Sub messaging system > Tutorial setup > Create service account credentials](#) in the *Google Pub/Sub online documentation*.

5. In the **OAuth 2.0 Scope** field, enter the OAuth scope to be able to publish into the Google Pub/Sub messaging service.

For more information about Google Pub/Sub OAuth scopes, see [OAuth 2.0 Scopes for Google APIs > Cloud Pub/Sub API](#) or [OAuth Scopes](#) in the *Google Identity online documentation*

6. (Optional) Select **Use proxy**, and in the **Proxy** field, enter the HTTP proxy URL.

Configuring the Google Pub/Sub destination: Advanced message parameters

This section describes how you can configure the advanced message parameters for your Google Pub/Sub destination after you finish configuring [the authentication and workspace settings](#).

For information about adding a new Google Pub/Sub destination on the SSB web interface, see [Configuring the Google Pub/Sub destination: adding a new Google Pub/Sub destination](#).

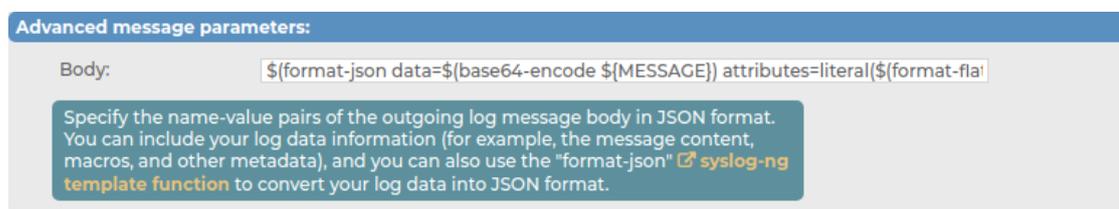
For information about configuring the performance-related settings of your Google Pub/Sub destination, see [Configuring the Google Pub/Sub destination: Performance-related settings](#).

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the advanced message parameters for your Google Pub/Sub destination

1. Navigate to **Log > Destination > <your-pubsub-destination> > Advanced message parameters**.

Figure 162: Log > Destination > <your-pubsub-destination> — Configuring the advanced message parameters for your Google Pub/Sub destination



2. In the **Body** field, customize the message format sent to the Google Pub/Sub messaging service by entering the message format you want.

CAUTION: Hazard of data loss!

Make sure that the customized message format is accepted by the Google Pub/Sub messaging service. For invalid messages, SSB will receive an HTTP 400 response code and messages with such a response code will be dropped.

For more information on dropping messages with HTTP 400 response code, see [Limitations](#).

For more information on the `$(format-json)` and the `$(format-flat-json)` template functions, see [Manipulating messages > Customizing message format using macros and templates > Template functions of syslog-ng PE](#) in the *syslog-ng PE Administration Guide*.

For more information on Google Pub/Sub message format, see [Publishing messages to topics > Message format](#) or [Pubsub Message](#) in the *Google Pub/Sub Reference Documentation*.

Configuring the Google Pub/Sub destination: Performance-related settings

This section describes configuring the performance-related settings of your Google Pub/Sub destination after you finish configuring [the authentication and workspace settings](#) and [the advanced message parameters](#) for the Google Pub/Sub destination.

NOTE: This section and the other Google Pub/Sub-related sections in this documentation are based on Google Pub/Sub messaging service concepts and terminology. If you do not use the Google Pub/Sub messaging service on a regular basis, One Identity recommends that you read the [Google Pub/Sub overview documentation](#) in the *Google Pub/Sub online documentation* to familiarize yourself with the messaging service's concepts and terminology before you continue reading these sections.

To configure the performance-related settings of your Google Pub/Sub destination

1. Navigate to **Log > Destination > <your-pubsub-destination> > Performance-related settings**.

Figure 163: Log > Destination > <your-pubsub-destination> — Configuring the authentication and workspace settings

Performance-related settings:

Number of workers:

If you use more than one workers together with the disk buffer option, syslog-ng PE creates a separate disk buffer file for each worker. As a result, decreasing the number of workers can result in losing data currently stored in the disk buffer files. Do not decrease the number of workers when the disk buffer files are in use.

Timeout:

Batch lines:

Batch-bytes value:

Currently Google Pub/Sub supports 10 MB (10000000 bytes) maximal batch size. You can check the current limitations at: [Quotas and limits | Cloud Pub/Sub Documentation | Google Cloud](#).

Batch-timeout:

Batch-timeout value:

2. In the **Number of workers** field, set the number of worker threads that you want SSB to use when sending messages to the server.

⚠ CAUTION: Hazard of data loss!

When you use more than one worker threads together with disk-buffering enabled, consider that the syslog-ng PE application behind SSB creates a separate disk-buffer file for each worker thread. This means that decreasing the number of workers can result in losing data currently stored in the disk-buffer files. To avoid data loss, One Identity recommends that you do not decrease the number of workers when the disk-buffer files are in use.

NOTE: Increasing the number of worker threads can drastically improve the performance of the destination.

3. In the **Timeout** field, specify the timeout (in **seconds**) that you want SSB to wait for an operation to complete, and then attempt to reconnect the server if the configured timeout limit is exceeded.

The default value of the **Timeout** setting is **0**, which means that it is disabled by default.

4. In the **Batch lines** field, specify how many lines you want SSB to flush to a destination in one batch.

For further information on the maximum number of **Batch lines**, see the *Publish request* limit at [Cloud Pub/Sub > Documentation > Resources > Quotas and limits > Resource limits](#).

NOTE: SSB waits for the configured number of lines to accumulate, and when this number is reached, SSB sends the message lines to the destination in a single batch. For example, if you set **Batch lines** to **100**, SSB waits for 100 message lines before sending them in one batch.

NOTE: Consider the following when configuring the number of batch lines:

- Increasing the number of batch lines increases throughput (because more messages are sent in a single batch), but also increases message latency.
- If the **Batch-timeout** option is disabled, the syslog-ng PE application behind SSB flushes the messages if it has sent the number of messages specified in **Batch lines**, or the queue became empty. If you stop or reload the syslog-ng PE application behind SSB, or if in case of network sources, the connection with the client is closed, the syslog-ng PE behind SSB automatically sends the unsent messages to the destination.
- If the **Batch-timeout** option is enabled and the queue becomes empty, SSB flushes the messages only if **Batch timeout** expires, or if the batch reaches the limit set in **Batch lines**.

NOTE: Depending on your source configuration settings, your batch may not reach the **Batch lines** limit before your queue becomes empty, and SSB forwards your messages.

5. In the **Batch-bytes** field, set the maximum size of payload in a batch (in **bytes**).

For the maximum value of **Batch-bytes**, see the *Publish request* limit at [Cloud Pub/Sub > Documentation > Resources > Quotas and limits > Resource limits](#).

NOTE: When configuring **Batch-bytes**, consider the following:

- If the size of the messages reaches this value, the syslog-ng PE application behind SSB sends the batch the Google Pub/Sub messaging service even if the number of messages is less than the value you configure in the **Batch-bytes** field.
- If **Batch-timeout** is enabled and the queue becomes empty, SSB flushes the messages only if **Batch-timeout** value expires, or if the message batch reaches the limit set in the **Batch-bytes** field.

6. (Optional) Select **Batch-timeout**, and in the **Batch-timeout** value field, specify the time SSB waits for **Batch lines** to accumulate in the output buffer.

SSB sends batches to the destinations evenly. The timer starts when the first message arrives to the buffer, so if only few messages arrive, SSB sends messages to the destination once every **Batch timeout** milliseconds at most.

Forwarding log messages to Splunk

This section describes how to forward messages from syslog-ng Store Box (SSB) to [Splunk](#).

From version 6.6.0, SSB uses the syslog-ng Premium Edition (syslog-ng PE) application's support to post messages to a Splunk deployment in JSON format, using the HTTP Event Collector (HEC) over HTTP and Secure HTTP (HTTPS).

For further details about how the application forwards messages to Splunk, see [splunk-hec: Sending messages to Splunk HTTP Event Collector](#) in the syslog-ng PE Administration Guide.

Prerequisites

This section describes the prerequisites for forwarding messages from syslog-ng Store Box (SSB) to [Splunk](#).

- On your Splunk deployment
 - You must enable HTTP Event Collector (HEC).
 - You must create a token for SSB.

NOTE: One Identity recommends that you use the default `syslog` source type for the token.

For details about enabling HEC and creating a token on your Splunk deployment, see [Set up and use HTTP Event Collector in Splunk Web](#).

- On the SSB appliance
 - If you want to use **Peer verification** for your Splunk destination, consider that the CA certificate must be added under **Log > Options > TLS settings** before you enable **Peer verification** under **Log > Destination > <your-splunk-destination> > Transport > HTTPS connection settings > Verification**.

Limitations

This section describes the limitations for forwarding messages from syslog-ng Store Box (SSB) to [Splunk](#).

- **Messages with HTTP 400 response code will be dropped**

If the message sent to Splunk is invalid, Splunk will reply with an HTTP 400 response code.

The message can be invalid for either of these reasons:

- A required argument is missing from the message.
- The message size exceeds limits.
- The message itself has an invalid format.

In these cases, SSB cannot successfully send the messages to Splunk. These messages would prevent SSB from sending further messages to the messaging service, therefore SSB must drop them.

- **IP address limitation**

IPv6 addresses are not supported as HTTP Event Collector (HEC) URLs.

- **Proxy type limitations**

Only HTTP proxy types are supported.

- Authenticated proxy types are not supported.

- **Batch-bytes unit of measurement limitation**

When configuring performance-related settings, you must enter the value of **Batch-bytes** in bytes.

Transport settings for the Splunk destination

This section describes how you can configure the transport settings (HTTP or HTTPS transport) for your Splunk destination.

To configure the transport settings for your Splunk destination

1. Navigate to **Log > Destinations** and select  to create a new destination.
2. Select **Splunk destination**.

Figure 164: Log > Destinations > <your-splunk-destination> - Creating your new Splunk destination

Destination type:

- Database server
- Remote host
- Splunk destination
- HDFS destination

Transport:

- HTTP
- HTTPS

HTTP connection settings:

Splunk HTTP Event Collector (HEC) URL

Splunk authentication and index settings:

Token:

Index:

Use proxy:

Proxy:

Only HTTP proxy type is supported.
Authenticated proxy types are not supported.

JSON message body:

Event:

Time:

Source type:

Host:

Source:

Performance-related settings:

Number of workers:

If you use more than one workers together with the disk buffer option, syslog-ng PE creates a separate disk buffer file for each worker. As a result, decreasing the number of workers can result in losing data currently stored in the disk buffer files. Do not decrease the number of workers when the disk buffer files are in use.

Timeout:

Batch lines:

Batch-bytes:

Batch-bytes value:

Batch-timeout:

Batch-timeout value:

3. Select the transport type ([HTTP connection settings](#) or [HTTPS connection settings](#)) that you want to use for your Splunk destination, then continue configuring the

respective connection type settings.

HTTP connection settings

This section describes the HTTP connection settings for your Splunk destination.

After you have [set the transport settings for the Splunk destination](#), and selected the HTTP transport type, you can configure the following:

- [The Splunk HTTP Event Collector \(HEC\) URL](#)
- [Splunk authentication and index settings](#)

Figure 165: Log > Destinations > <your-splunk-destination> Transport > HTTP > HTTP connection settings - Configuring the HTTP connection settings for your new Splunk destination

The screenshot shows the configuration page for an HTTP destination in Splunk. At the top, the 'Transport' section is set to 'HTTP'. Below this, the 'HTTP connection settings' section contains a text input field for the 'Splunk HTTP Event Collector (HEC) URL' with the value 'http://example.com:8088/services/collector/event'. The 'Splunk authentication and index settings' section includes fields for 'Token' and 'Index', both of which are redacted. The 'Use proxy' checkbox is checked, and a 'Proxy' field is present but redacted. A warning box at the bottom states: 'Only HTTP proxy type is supported. Authenticated proxy types are not supported.'

The Splunk HTTP Event Collector (HEC) URL

When configuring the HTTP or HTTPS connection settings for your Splunk destination for syslog-ng Store Box (SSB), you have to configure the HTTP Event Collector (HEC) URL first.

To configure the Splunk HEC URL for your Splunk destination

1. Navigate to **Log > Destinations > <your-splunk-destination>**.
2. Select the **Transport** type you want to use (**HTTP** or **HTTPS**).
3. Under **HTTP connection settings** or **HTTPS connection settings** (depending on your **Transport** type), enter the **Splunk HTTP Event Collector (HEC) URL** of your choice.

NOTE: IPv6 addresses are not supported.

Splunk authentication and index settings

After setting the transport settings, and then configuring the Splunk HEC URL for your Splunk destination, you have to configure your Splunk authentication and index settings.

To configure your Splunk authentication and index settings

1. Copy the Splunk HTTP Event Collector token provided by your Splunk deployment into the **Token** field.

Because the Splunk HTTP Event Collector token permits SSB to send messages to Splunk, configuring the **Token** field is required for the Splunk destination.

For details about setting up and using the Splunk HTTP Event Collector token on your Splunk deployment, see [Set up and use HTTP Event Collector in Splunk Web](#).

2. Enter the name of the Splunk index where Splunk will store the messages received from SSB into the **Index** field.

Configuring the **Index** field is required for the Splunk destination.

3. (Optional) Enable **Use proxy** to use a proxy address while forwarding messages to your Splunk deployment.

NOTE: When configuring your proxy address, consider the following:

- Only HTTP proxy type is supported.
- Authenticated proxy types are not supported.

HTTPS connection settings

This section describes the HTTPS connection settings for your Splunk destination.

After you have [set the transport settings for the Splunk destination](#), and selected the HTTPS transport type, you can configure the following:

- [The Splunk HTTP Event Collector \(HEC\) URL](#)
- [Splunk authentication and index settings](#)
- [Verification](#)

Figure 166: Log > Destinations > <your-splunk-destination> Transport > HTTPS > HTTPS connection settings - Configuring the HTTPS connection settings for your new Splunk destination

The screenshot shows the configuration page for a Splunk destination. At the top, under the 'Transport' section, 'HTTPS' is selected. Below this is the 'HTTPS connection settings' section, which includes a text input field for the 'Splunk HTTP Event Collector (HEC) URL' containing the value 'https://example.com:8088/services/collector/event'. The 'Splunk authentication and index settings' section contains fields for 'Token' and 'Index', a checked 'Use proxy' checkbox, and a 'Proxy' field. A warning message states: 'Only HTTP proxy type is supported. Authenticated proxy types are not supported.' The 'Verification' section has a checked 'Peer verification' checkbox and a message: 'The CA certificate should be added under the Log Options TLS settings menu point'. At the bottom, 'Client authentication' is checked, with fields for 'Client X.509' and 'Client key'.

The Splunk HTTP Event Collector (HEC) URL

When configuring the HTTP or HTTPS connection settings for your Splunk destination for syslog-ng Store Box (SSB), you have to configure the HTTP Event Collector (HEC) URL first.

To configure the Splunk HEC URL for your Splunk destination

1. Navigate to **Log > Destinations > <your-splunk-destination>**.
2. Select the **Transport** type you want to use (**HTTP** or **HTTPS**).
3. Under **HTTP connection settings** or **HTTPS connection settings** (depending on your **Transport** type), enter the **Splunk HTTP Event Collector (HEC) URL** of your choice.

NOTE: IPv6 addresses are not supported.

Splunk authentication and index settings

After setting the transport settings, and then configuring the Splunk HEC URL for your Splunk destination, you have to configure your Splunk authentication and index settings.

To configure your Splunk authentication and index settings

1. Copy the Splunk HTTP Event Collector token provided by your Splunk deployment into the **Token** field.

Because the Splunk HTTP Event Collector token permits SSB to send messages to Splunk, configuring the **Token** field is required for the Splunk destination.

For details about setting up and using the Splunk HTTP Event Collector token on your Splunk deployment, see Set up and use HTTP Event Collector in Splunk Web.

2. Enter the name of the Splunk index where Splunk will store the messages received from SSB into the **Index** field.

Configuring the **Index** field is required for the Splunk destination.

3. (Optional) Enable **Use proxy** to use a proxy address while forwarding messages to your Splunk deployment.

NOTE: When configuring your proxy address, consider the following:

- Only HTTP proxy type is supported.
- Authenticated proxy types are not supported.

Verification

After setting the transport settings, configuring the Splunk HEC, then configuring your Splunk authentication and index settings, you can configure your verification settings for your Splunk destination.

NOTE: Both **Peer verification** and **Client authentication** are optional. If you want to use **Peer verification** for your Splunk destination, consider that the CA certificate must be added under **Log > Options > TLS settings** before you enable the **Peer verification** option.

To configure the verification settings for your Splunk destination

1. Navigate to **Log > Destination > <your-splunk-destination> > Transport > HTTPS connection settings > Verification**.
2. (Optional) To use the CA certificate that you previously added under **Log > Options > TLS settings**, enable **Peer verification**.
3. (Optional) To use the client authentication method for verification, enable **Client authentication**, and add your **Client X.509** certificate and your **Client key**.

JSON message body

When configuring your Splunk destination for syslog-ng Store Box (SSB), you have to configure JSON message body settings after you finish [configuring the Transport settings for the Splunk destination](#).

Figure 167: Log > Destinations > <your-splunk-destination> - Configuring the JSON message body for your new Splunk destination

JSON message body:	
Event:	<input type="text" value="\${S_ISODATE} \${HOST} \$"/>
Time:	<input type="text" value="\${S_UNIXTIME}.\${S_MSEC}"/>
Source type:	<input type="text" value=".app.name:-syslog"/>
Host:	<input type="text" value="\${HOST}"/>
Source:	<input type="text" value="syslog-ng"/>

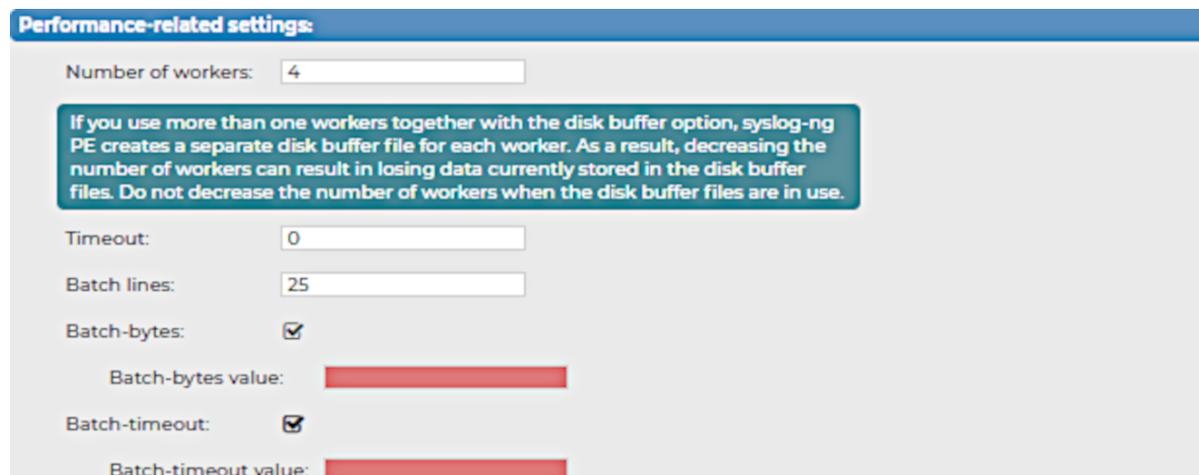
To configure the JSON message body settings for your Splunk destination

1. In the **Event** field, specify the body of the message that SSB sends to your Splunk deployment.
NOTE: SSB sends the contents of the message body (which will be identical to the value of the `${MSG}` macro in a syslog-ng PE configuration) to your Splunk deployment in JSON format.
Default syntax: `${S_ISODATE} ${HOST} ${MSGHDR}${MSG}\n`
2. In the **Time** field, specify the timestamp for the message that SSB sends to your Splunk deployment.
Default syntax: `${S_UNIXTIME}.${S_MSEC}`
3. In the **Source type** field, specify the same source type that you configured on your Splunk deployment for the Splunk token assigned to this SSB destination.
Default syntax: `.app.name:-syslog`
For details about source types, see [Why source types matter](#) in the Splunk online documentation.
4. In the **Host** field, specify the hostname for the message that SSB sends to your Splunk deployment.
NOTE: On your Splunk deployment, the message will appear as it is sent by the host specified in this field.
Default syntax: `${HOST}`
5. In the **Source** field, enter the value of the source field configured on your Splunk deployment into the **Source** field.
Default value on Splunk deployments: `syslog-ng`.

Performance-related settings

When configuring your Splunk destination for syslog-ng Store Box (SSB), you have to configure the performance-related settings after you finish [configuring the JSON message body settings for the Splunk destination](#).

Figure 168: Log > Destinations > <your-splunk-destination> - Configuring the Performance-related settings for your new Splunk destination



Performance-related settings:

Number of workers:

If you use more than one workers together with the disk buffer option, syslog-ng PE creates a separate disk buffer file for each worker. As a result, decreasing the number of workers can result in losing data currently stored in the disk buffer files. Do not decrease the number of workers when the disk buffer files are in use.

Timeout:

Batch lines:

Batch-bytes:

Batch-bytes value:

Batch-timeout:

Batch-timeout value:

To configure the performance-related settings for your Splunk destination

1. Specify the **Number of workers** that you want SSB to use when sending messages to the server.

⚠ CAUTION: Hazard of data loss!

When you use more than one worker threads together with disk-buffering enabled, consider that the syslog-ng PE application behind SSB creates a separate disk-buffer file for each worker thread. This means that decreasing the number of workers can result in losing data currently stored in the disk-buffer files. To avoid data loss, One Identity recommends that you do not decrease the number of workers when the disk-buffer files are in use.

NOTE: Increasing the number of worker threads can drastically improve the performance of the destination.

2. Specify the **Timeout** (in seconds) that you want SSB to wait for an operation to complete, and attempt to reconnect the server if the configured timeout limit is exceeded.
3. In the **Batch lines** field, specify how many lines you want SSB to flush to a destination in one batch.

NOTE: SSB waits for the configured number of lines to accumulate, and when this number is reached, SSB sends the message lines to the destination in a single batch. For example, if you set **Batch lines** to **100**, SSB waits for 100 message lines before sending them in one batch.

Consider the following when configuring the number of batch lines:

- Increasing the number of batch lines increases throughput (because more messages are sent in a single batch), but also increases message latency.
- If **Batch-timeout** option is disabled, the syslog-ng PE application behind SSB flushes the messages if it has sent the number of messages specified in **Batch lines**, or the queue became empty. If you stop or reload the syslog-ng PE application behind SSB, or if in case of network sources, the connection with the client is closed, the syslog-ng PE application behind SSB automatically sends the unsent messages to the destination.
- If the **Batch-timeout** option is enabled and the queue becomes empty, SSB flushes the messages only if **Batch timeout** expires, or the batch reaches the limit set in **Batch lines**.

NOTE: Depending on your source configuration settings, your batch may not reach the **Batch lines** limit before your queue becomes empty, and SSB forwards your messages.

4. (Optional) Select **Batch-bytes**, and in the **Batch-bytes value** field, set the maximum size of payload in a batch (in bytes).

NOTE: When configuring **Batch-bytes**, consider the following:

- If the size of the messages reaches this value, the syslog-ng PE application behind SSB sends the batch to your Splunk deployment even if the number of messages is less than the value you configure in the **Batch-bytes** field.
- Consider that if **Batch-timeout** is enabled and the queue becomes empty, SSB flushes the messages only if **Batch-timeout value** expires, or if the message batch reaches the limit set in the **Batch-bytes** field.

5. (Optional) Select **Batch-timeout**, and in the **Batch-timeout value** field, specify the time SSB waits for **Batch lines** to accumulate in the output buffer.

SSB sends batches to the destinations evenly. The timer starts when the first message arrives to the buffer, so if only few messages arrive, SSB sends messages to the destination once every **Batch timeout** milliseconds at most.

Forwarding log messages to HDFS destinations

You can forward log messages from syslog-ng Store Box(SSB) to [Hadoop Distributed File System \(HDFS\)](#) servers, allowing you to store your log data on a distributed, scalable file system. This is especially useful if you have huge amounts of log messages that would be difficult to store otherwise, or if you want to process your messages using Hadoop tools.

Forwarding log messages from SSB to an HDFS destination comprises the following steps:

1. [Configure a Kerberos policy.](#)
2. [Configure the HDFS cluster.](#)
3. [Configure an HDFS destination.](#)
4. [Create a log path.](#)

Configuring a Kerberos policy

The syslog-ng Store Box(SSB) appliance authenticates to the HDFS cluster through a trusted third party, a Kerberos server. Once SSB has been granted a ticket by Kerberos, it is then able to write data to the HDFS servers.

To configure a Kerberos policy

1. Navigate to **Policies > Kerberos** and select  to create a new policy.
2. In the **Default realm** field, enter the name of the Kerberos realm where your SSB resides.
3. If you have to specify the address of the Key Distribution Center (KDC) server, click  first, and then enter the FQDN or IP address of the KDC server that is issuing tickets within your Kerberos realm.

If your DNS server is configured to map Kerberos realms to KDC hostnames, you do not need to specify KDC servers here.

4. Add a Kerberos principal policy. First, select  under **Kerberos principals**.
5. Enter a name for your policy. This name will be used later, on the **Policies > HDFS Cluster** page of SSB, to identify the Kerberos principal policy you want to use. For more information, see [Configuring the HDFS cluster](#) on page 302.
6. Upload the keytab file that contains keys for your principal.

This is the principal that has write access to the HDFS cluster.

The keytab file was provided to you by the Kerberos administrator, and it contains the encrypted key required to authenticate your principal to Kerberos.

7. Select your principal from the **Principal** list.
The keytab file you have uploaded may contain keys for several principals. This list displays all the principals with keys in the uploaded keytab file.
8. Test whether or not your principal is able to authenticate to Kerberos. Click **Test authentication**.

9. When all works fine, click **Commit**.

Figure 169: Policies > Kerberos — Configuring a Kerberos policy

The screenshot shows the configuration page for a Kerberos policy. On the left is a navigation menu with sections: Basic Settings, AAA, Policies (with sub-items: Hostlists, Shares, Backup & Archive/Cleanup, Alert targets, Kerberos, HDFS Cluster), Log, Search, and Reports. Below this is a 'User menu' with options: Private keystore, Change password, Preferences, and Logout. At the bottom left is a 'System monitor' section showing 'Time: 2018-08-03 12:10' and 'Remaining time: 09:51'. The main content area is titled 'Kerberos' and has a 'Commit' button at the top right. It is divided into two sections: 'Default realm' and 'Kerberos principals'. The 'Default realm' section has a 'Default realm:' field with the value 'SSB.BALABIT'. Below it is a 'KDC servers:' section with a message: 'If your DNS server is configured to map Kerberos realms to KDC hostnames, you do not need to specify KDC servers here.' There is an 'Address' field with a dropdown menu showing 'kerberos-1.ssb.balabit'. The 'Kerberos principals' section has a 'hdfs-principal' field. Below it are 'Upload keytab file:' buttons for 'Download' and 'Delete (and upload another)'. The 'Principal:' field has a dropdown menu showing 'testuser@SSB.BAL'. At the bottom of this section is a 'Test authentication' button. A plus sign icon is at the bottom right of the main content area.

Configuring the HDFS cluster

The following describes how to configure settings related to the HDFS cluster where you want to forward logs.

Prerequisites:

- You have write/perform permission for the **Basic Settings > System** page. For details on how to assign user rights, see [Managing user rights and usergroups](#) on page 127.
- You have configured a Kerberos policy. For more information, see [Configuring a Kerberos policy](#) on page 301.

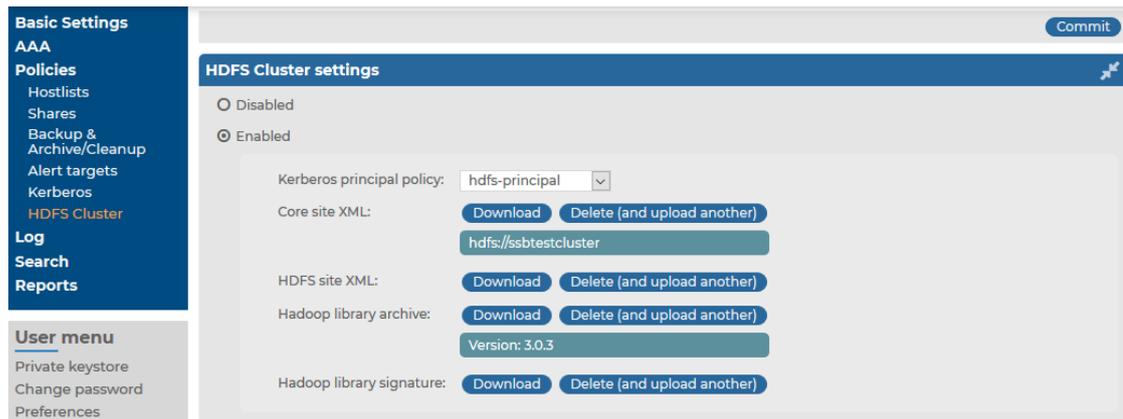
To configure settings related to the HDFS cluster where you want to forward logs

1. Navigate to **Policies > HDFS Cluster** and select **Enabled**.
2. Select the **Kerberos principal policy** configured previously (for details, see [Configuring a Kerberos policy](#) on page 301).
3. Upload the **Core site XML** file of your HDFS cluster. You may have to ask for this file from your HDFS cluster administrator.
4. Upload the **HDFS site XML** file of your HDFS cluster. You may have to ask for this file from your HDFS cluster administrator.

5. In the **Hadoop library archive** field, upload the Hadoop binary tarball matching the version of your HDFS cluster infrastructure. Binary tarballs are distributed on the [official Apache site](#).
6. In the **Hadoop library signature** field, upload the signature GPG file matching the used binary version. Signature GPG files are distributed on the [official Apache site](#).
7. Click **Commit**.

The version number of the Hadoop library archive is displayed.

Figure 170: Policies > HDFS Cluster — Configuring the HDFS cluster



Configuring an HDFS destination

The following describes how to configure the HDFS destination where you want to forward logs.

Prerequisites:

- You have configured a Kerberos policy. For more information, see [Configuring a Kerberos policy](#) on page 301.
- You have enabled the HDFS cluster. For more information, see [Configuring the HDFS cluster](#) on page 302.

To configure the HDFS destination

1. Navigate to **Log > Destinations** and select  to add a new destination.
2. Enter a name for the destination.
3. Select **HDFS destination**.

- In **File path**, specify the absolute path for the destination file on the HDFS server.

Figure 171: Log > Destinations – Configuring an HDFS destination

- Select the syslog protocol to use from the **Syslog protocol** field.
 - To use the legacy BSD-syslog protocol described in RFC 3164, select **Legacy** and specify the message template to use. Select **Legacy** to use the message format described in the RFC, or **ISO date** to replace the original time stamp with an ISO8061 compliant time stamp that includes year and timezone information. To customize the format of the message contents using macros, select **Custom message part only**, or **Custom on-wire message** to completely reformat the message (including the headers). For details on using macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide* and [Templates and macros](#) in the *syslog-ng PE Administration Guide*. If you have no special requirements, use the **ISO date** template.
 - To use the new IETF-syslog protocol, select **Syslog**. Note that most syslog applications and devices currently support only the legacy protocol. If you need, you can customize the contents of the message using macros. Note that for the IETF-syslog protocol, the header cannot be customized. For details on using macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide* and [Templates and macros](#) in the *syslog-ng PE Administration Guide*.
- Configure message rate alerting. For detailed instructions, see [Configuring message rate alerting](#) on page 91.

7. The time stamps of most log messages is accurate only to the second. The syslog-ng Store Box(SSB) appliance can include more accurate time stamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the `frac_digits()` parameter of syslog-ng.
8. If the server and SSB are located in a different timezone and you use the Legacy message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
9. Set the size of the disk buffer (in Megabytes) in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the `log_disk_fifo_size()` parameter of syslog-ng.

Note that SSB does not pre-allocate the hard disk required for the disk buffer, so make sure that the required disk space is available on SSB. For details on creating archiving policies and adjusting the disk-fillup prevention, see [Archiving and cleanup](#) and [Preventing disk space fill up](#).

Example: Calculating disk buffer size

The size of the disk buffer you need depends on the rate of the incoming messages, the size of the messages, and the length of the network outage that you want to cover. For example:

- SSB is receiving 15000 messages per second
- On the average, one message is 250 bytes long
- You estimate that the longest time the destination will be unavailable is 4 hours

In this case, you need a disk buffer for $250 \text{ [bytes]} * 15000 \text{ [messages per second]} * 4 * 60 * 60 \text{ [seconds]} = 54000000000 \text{ [bytes]}$, which is 54000 Megabytes (in other words, a bit over 50 GB).

10. Click **Commit**.
11. To start sending messages to the destination, include the new destination in a logpath. For details, see [Log paths: routing and processing messages](#).

On the **Log > Paths** page, the HDFS destination will be displayed in the **remote** category.

Log paths: routing and processing messages

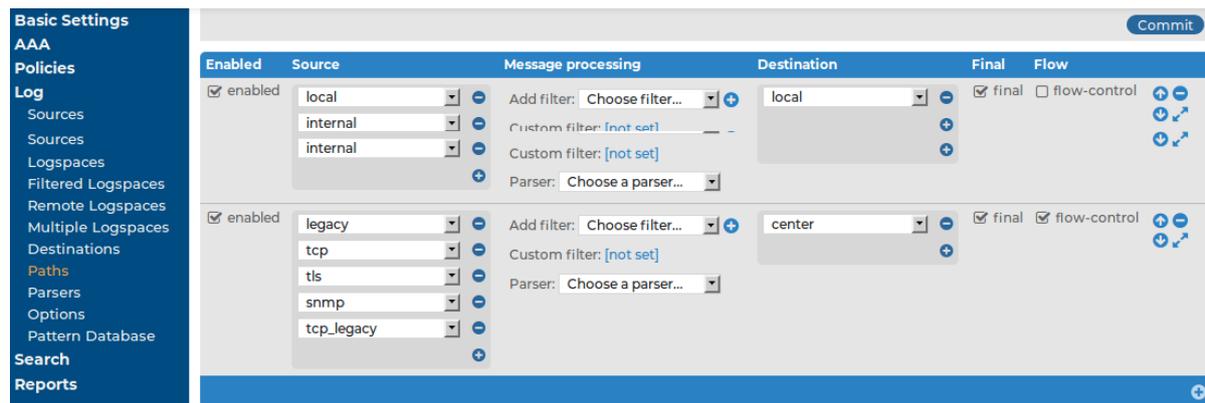
This section describes how to create and configure log paths in syslog-ng Store Box(SSB). Log paths and filters allow you to select and route messages to specific destinations. You can also parse and modify the log messages in log path using message parsers and rewriter rules. The log path processes the incoming messages as follows.

1. Parse the message as a syslog message (unless message parsing is explicitly disabled for the source).
2. [Classify the message](#) using a pattern database.
3. [Modify the message using rewrite rules](#) (before filtering).
4. [Filter the messages](#), for example, based on sender hostname or message content. If the message does not match the configured filter, syslog-ng Store Box(SSB) will not send it to the destination.
5. Parse the text of the message (that is, the `${MESSAGE}` part) using a [key-value parser](#) or the [sudo parser](#).
6. [Modify the message using rewrite rules](#) (after filtering and other parsing).
7. SSB sends the message to the destinations set in the log path. The destinations are [local](#), [optionally encrypted files on SSB](#), or [remote servers, such as a database server](#).
 - For a list of default log paths, see [Default logpaths in SSB](#) on page 306.
 - For details on how to create a new log path, see [Creating new log paths](#) on page 307.
 - For details on how to send only selected messages to a destination, see [Filtering messages](#) on page 309.
 - To modify parts of a message, see [Replace message parts or create new macros with rewrite rules](#) on page 312.

Default logpaths in SSB

Two log paths are available by default in the syslog-ng Store Box(SSB) appliance (see [Log > Paths](#)):

Figure 172: Log > Paths – Default logpaths of SSB



- The first log path collects the local messages of SSB. It sends every message of the web interface, the built-in syslog-ng server, and other internal components to the **local** logspace.
- The second log path collects messages sent to SSB using the default syslog sources (for details, see [Default message sources in SSB](#) on page 193) or via SNMP (for details, see [Receiving SNMP messages](#) on page 214). These messages are stored in the **center** logspace.

NOTE: Note that both default log paths are marked as **Final**: if you create a new log path that collects logs from the default sources, make sure to adjust the order of the log paths, or disable the **Final** option for the default log path.

Creating new log paths

This section describes how to create a new log path.

To create a new log path

1. Navigate to **Log > Paths** and select . A new log path is added to the list of log paths.
2. Select a source for the log path from the **Source** field. Messages arriving to this source will be processed by this log path. To add more sources to the log path, select in the source field and repeat this step.

Figure 173: Log > Paths – Creating a new logpath

Remote sources receive messages from the network, while built-in sources are messages that originate on syslog-ng Store Box(SSB). However, note that the SNMP source (for details, see [Receiving SNMP messages](#) on page 214) is listed in the built-in section.

TIP: To process every message of every source, leave the source option on all. This is equivalent to using the catchall flag of syslog-ng.

3. Select a destination for the log path from the **Destination** field. Messages arriving to this source will be forwarded to this destination. To add more destinations to the log path, select in the destination field and repeat this step.

NOTE: Remote destinations forward the messages to external servers or databases and are configured on the **Log > Destinations** page (for details, see [Forwarding messages from SSB](#) on page 266).

Local destinations store the messages locally on SSB and are configured on the **Log > Logspaces** page (for details, see [Storing messages on SSB](#) on page 215).

If you do not want to store the messages arriving to this log path, leave the **Destination** field on none.

CAUTION:

The none destination discards messages — messages sent only to this destination will be lost irrevocably.

- If you do not want other log paths to process the messages sent to a destination by this log path, select the **Final** option.

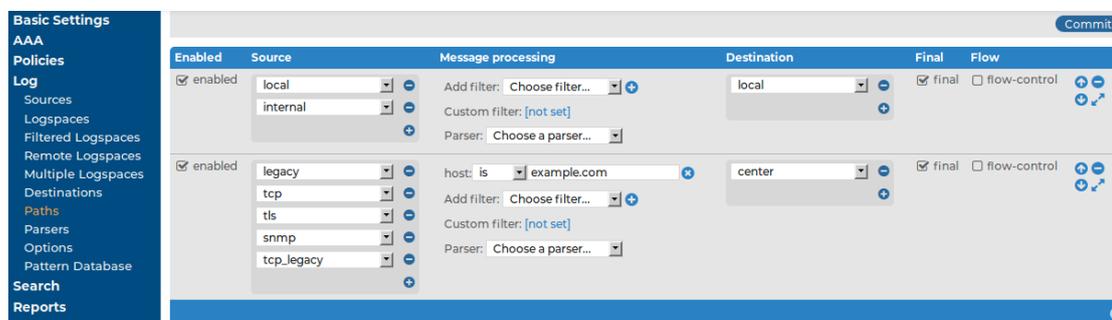
The order of the log paths is important, especially if you use the **Final** option in one or more destinations, because SSB evaluates log paths in descending order. Use the ,  buttons to position the log path if needed.

- To enable flow-control for this log path, select the **flow-control** option. For details on how flow-control works, see [Managing incoming and outgoing messages with flow-control](#) on page 20.

NOTE: As a result of toggling the flow-control status of the logpath, the output buffer size of the logpath's destination(s) will change. For the changes to take effect, navigate to **Basic Settings > System > Service control** and click **Restart syslog-ng**.

- If you do not want to send every message from the sources to the destinations, use filters. Select the filter to use from the **Filter** field, click , and configure the filter as needed. To apply more filters, click  and select a new filter. Note that SSB sends only those messages to the destinations that pass every listed filter of the log path. The available filters are described in [Filtering messages](#) on page 309.

Figure 174: Log > Paths — Filtering log messages



- Click **Commit**. After that, the new log path will start to collect log messages.

TIP: If you do not want to activate the log path immediately, deselect the **Enable** option.

Filtering messages

This section describes the filters that can be used in log paths. Every filter can be used to select (for example, `priority is`) or exclude (for example, `priority is not`) messages. The following filters are available:

- **facility:** Select messages sent by a specific facility (for example, `kernel`).
- **host:** Select messages sent by a specific host. Enter the a hostname, IP address, or a POSIX (extended) regular expression.
- **message:** Select messages containing a specific keyword or POSIX (extended) regular expression in the text of the log message (excluding the headers).
- **priority:** Select messages of a specific priority.
- **program:** Select messages sent by a specific application. Enter the name of the application or a POSIX (extended) regular expression.
- **sender:** Filter on the address of the host that sent the message to syslog-ng Store Box(SSB).

NOTE: To be able to use this filter, as a prerequisite, you must have a hostlist defined. For more information, see [Creating hostlist policies](#) on page 189.

NOTE: When using the host, message, and program filters, remember to escape special characters. The characters `()[]{}.*?+^$|\` are treated as special symbols and have to be escaped with a backslash (`\`) in order to be interpreted as literal characters.

NOTE: The effect of the sender and the host filters is the same if every client sends the logs directly to SSB. But if SSB receives messages from relays, then the host filter applies to the address of the clients, while the sender applies to the address of the relays.

If multiple filters are set for a log path, only messages complying to every filter are sent to the destinations. (In other words, filters are added using the logical AND operation.)

Figure 175: Log > Paths – Using custom filters

The screenshot displays the 'Log Paths' configuration page. On the left is a navigation menu with sections like 'Basic Settings', 'Policies', 'Log', 'User menu', and 'System monitor'. The main area shows a table of log paths:

Enabled	Source	Message processing	Destination	Final	Flow
<input checked="" type="checkbox"/>	local internal	Add filter: Choose filter... Custom filter: [not set] Parser: Choose a parser...	local	<input checked="" type="checkbox"/>	<input type="checkbox"/> flow-control
<input checked="" type="checkbox"/>	legacy tcp tls snmp tcp_legacy	Add filter: Choose filter... Custom filter: [not set] Parser: Choose a parser...	center	<input checked="" type="checkbox"/>	<input type="checkbox"/> flow-control
<input checked="" type="checkbox"/>	[all]	Add filter: Choose filter... Custom filter: "\$SOURCEIP..." Parser: Choose a parser...	center	<input checked="" type="checkbox"/>	<input type="checkbox"/> flow-control

The selected path (third row) has a 'Custom filter' field containing: `"$SOURCEIP" == "192.168.1.1";`. Below this, there are sections for 'Before message processing' and 'After message processing', each with a search and replace interface.

If you need more complex filtering in your log path, select the  of the log path and enter a custom filter into the appearing field. The contents of the **Custom filter** field are pasted into the `filter()` parameter of the `syslog-ng` log path definition.

When defining custom filters, you can use regular expressions. By default, custom filters use POSIX-style (extended) regular expressions.

NOTE: When using POSIX regular expressions, the characters `()[]{}.*?+^$|\` are used as special symbols. Depending on how you want to use these characters and which quotation mark you use, these characters must be used differently, as summarized below:

- When enclosing strings between double-quotes ("string"), the string is interpreted and you have to escape special characters, that is, prefix them with a backslash (`\`) if they are meant literally.
- Strings between single quotes ('string') are treated as literals and are not interpreted at all, so you do not have to escape special characters.

To use other expression types, add the `type()` option after the regular expression. For example:

```
message("([0-9]+)=\\1" type("pcre"))
```

In this example, a PCRE regular expression with backreference is used and a match is returned if the message contains identical numbers separated by the equal sign (=). For example:

```
123=123
```

Replace message parts or create new macros with rewrite rules

The syslog-ng Store Box(SSB) appliance can rewrite parts of the messages using rewrite rules. Almost all part (macro) of the message can be rewritten. The rules use a key-value pair format.

The **Replace with** value completely replaces the old value of the message part. If the message part does not already exist, SSB automatically creates it. If you want to perform search and replace in the text of the log message, see [Find and replace the text of the log message](#) on page 314 instead.

NOTE: You cannot change the values of hard macros in rewrite rules. For the list of hard macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*.

To replace message parts or create new macros with rewrite rules

1. Navigate to **Log > Paths**.
2. Select the path(s) where you want to use rewrite rules.
3. In the **Rewrites** section, click  to add a new rewrite rule. Rewrite rules can be applied before filtering, or after filtering.

The sequence of filtering and rewrite rules depends on how it was specified in the log path. The sequence of the process is the following:

- a. Parse the message as a syslog message (unless message parsing is explicitly disabled for the source).
- b. [Classify the message](#) using a pattern database.
- c. [Modify the message using rewrite rules](#) (before filtering).
- d. [Filter the messages](#), for example, based on sender hostname or message content. If the message does not match the configured filter, syslog-ng Store Box(SSB) will not send it to the destination.
- e. Parse the text of the message (that is, the `_${MESSAGE}` part) using a [key-value parser](#) or the [sudo parser](#).
- f. [Modify the message using rewrite rules](#) (after filtering and other parsing).
- g. SSB sends the message to the destinations set in the log path. The destinations

are local, optionally encrypted files on SSB, or remote servers, such as a database server.

Figure 176: Log > Paths – Modifying messages using rewrite

4. Enter the part of the message to rewrite into the **In Message part** field. For example, MESSAGE, HOST, .SDATA.meta.custom. If the specified field does not exist, it is automatically created and set to the **Replace with** field.
5. Enter the value of the message part after rewriting into the **Replace with** field. To use macros, begin with a \$ sign and enclose the name of the macro between braces, for example `${MSG}`, `${.SDATA.meta.custom}`.

NOTE:

- The replacement value completely replaces the old value of the message part.
- You cannot change the values of hard macros in rewrite rules. For the list of hard macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*.

6. Click **Commit**.

Find and replace the text of the log message

You can perform search and replace operations on the log messages to rewrite the messages. Almost all part (macro) of the message can be rewritten. You can use PCRE regular expressions.

If you want to completely replace a message part, or create a new one that does not already exist, see [Replace message parts or create new macros with rewrite rules](#) on page 312 instead.

NOTE: You cannot change the values of hard macros in rewrite rules. For the list of hard macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*.

To find and replace the text of the log message

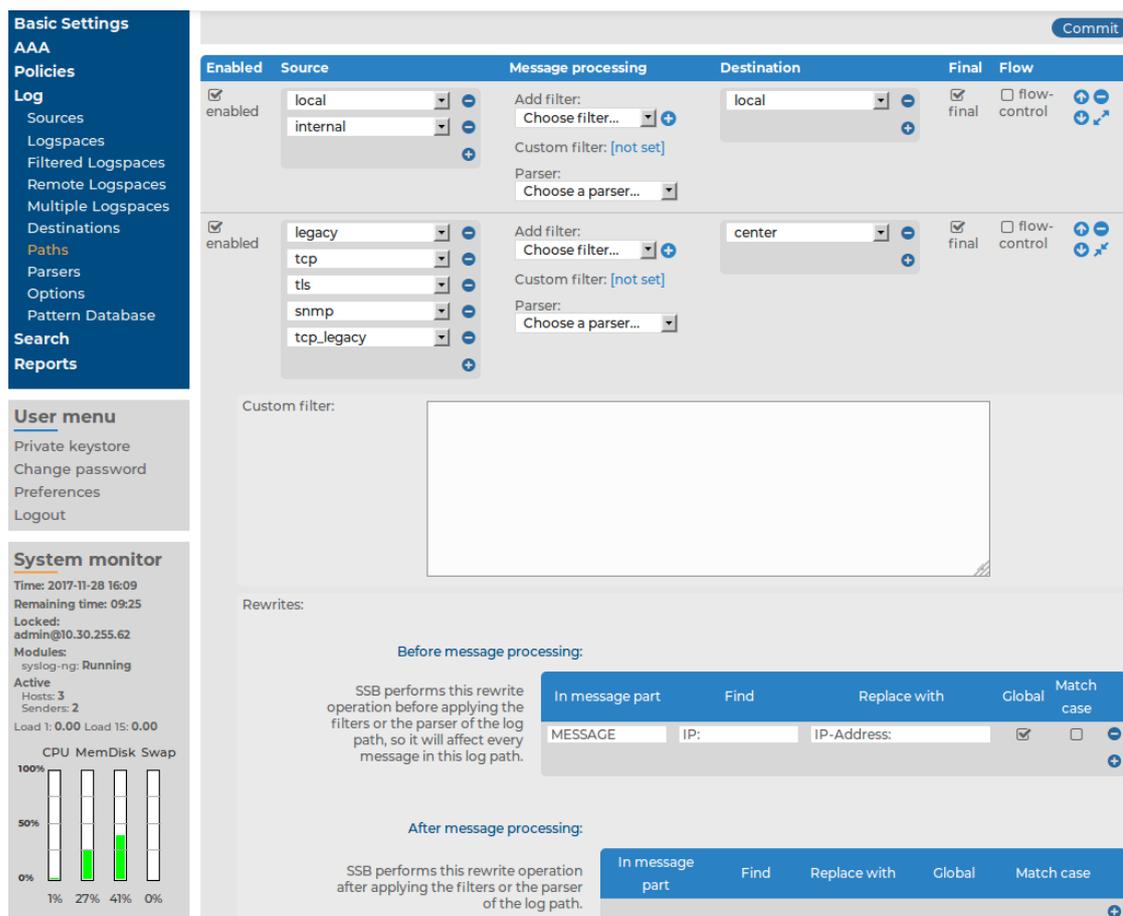
1. Navigate to **Log > Paths**.
2. Select the path(s) where you want to use rewrite rules.
3. In the **Rewrites** section, click  to add a new rewrite rule. Rewrite rules can be applied before filtering, or after filtering.

The sequence of filtering and rewrite rules depends on how it was specified in the log path. The sequence of the process is the following:

- a. Parse the message as a syslog message (unless message parsing is explicitly disabled for the source).
- b. [Classify the message](#) using a pattern database.
- c. [Modify the message using rewrite rules](#) (before filtering).
- d. [Filter the messages](#), for example, based on sender hostname or message content. If the message does not match the configured filter, syslog-ng Store Box(SSB) will not send it to the destination.
- e. Parse the text of the message (that is, the `${MESSAGE}` part) using a [key-value parser](#) or the [sudo parser](#).
- f. [Modify the message using rewrite rules](#) (after filtering and other parsing).
- g. SSB sends the message to the destinations set in the log path. The destinations are [local, optionally encrypted files on SSB](#), or [remote servers, such as a database server](#).

The message part you want to modify must already exist — if you want to modify a macro that a parser creates, you must add the rewrite rule into the **After filtering** section.

Figure 177: Log > Paths — Find and replace in the text of log messages



4. Enter the part of the message to modify into the **In Message part** field. For example, MESSAGE, HOST, .SDATA.meta.custom.
You cannot change the values of hard macros in rewrite rules. For the list of hard macros, see [Hard versus soft macros](#) in the *syslog-ng PE Administration Guide*.
5. Enter the expression you want to find into the **Find** field. You can use PCRE regular expressions.
6. Enter the expression that will replace the **Find** expression into the **Replace with** field. By default, SSB replaces the first occurrence of the expression. To use macros, begin with a \$ sign and enclose the name of the macro between braces, for example `${MSG}`, `${.SDATA.meta.custom}`.
You can use matches of the **Find** expression as well: `${0}` stores the entire match, `${1}` is the first group of the match (parentheses), and so on. If you use named patterns in the **Find** expression (`?<name>pattern`), you can use `${name}` as well.
7. To replace every occurrence of the **Find** expression, select the **Global** option.
8. To make the **Find** expression case sensitive, select the **Match case** option.
9. Click **Commit**.

Parsing sudo log messages

The sudo parser separates **sudo** log messages into name-value pairs.

Use this parser to enrich your log message data with details of privilege escalation events, such as who initiated the event, what command was issued, and so on. The parsed values are automatically assigned metadata, which you can then display on the syslog-ng Store Box(SSB) search interface as dynamic columns.

The aim is to enrich log data with semantic value, and consistently apply the same metadata to the same type of log message data. For example, any information about the client where **sudo** was executed will always be displayed in the **dest** and **src** dynamic columns.

The sudo parser maps the contents of log messages to the dynamic columns listed in [Mapping sudo log message contents to dynamic columns](#) on page 316.

Example log message:

```
2016-08-12T06:57:12+02:00 mail sudo: johndoe : TTY=tty ; PWD=pwd ; USER=usr ; GROUP=grp ; TSID=000001 ; ENV=PATH=/usr/local/bin ; COMMAND=cmd -w 20 -c 40
```

Table 8: Mapping sudo log message contents to dynamic columns

Dynamic column	Parsed value	Description
action	success	The action performed on the resource. Possible values: success
app	sudo	The application where the command was issued. Currently, the value of this column is always sudo .
dest	mail	The IP address or hostname of the entity that validates the authentication request.
src	mail	The IP address or hostname of the entity that sends the authentication request.
src_user	johndoe	The user identifier showing which user executed sudo .
tty	tty	The terminal device name where sudo was executed.
pwd	pwd	The working directory where sudo was issued.
user	usr	The user identifier showing who the new user is after executing sudo .
group	grp	The sudo group target (if present).
tsid	000001	The sudo terminal session (log) identifier (if

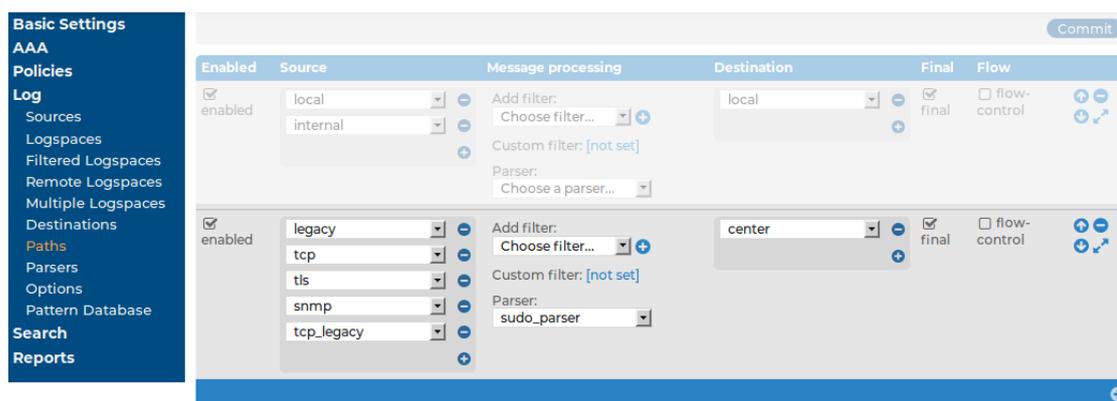
Dynamic column	Parsed value	Description
		present).
env	PATH=/usr/local/bin	The sudo environment variable (if present).
command	cmd -w 20 -c 40	The command that was issued by the src_user as a superuser.
tags	authentication privileged	The metadata to flag the message as a sudo log message.

You can also use the enriched metadata generated from the parsed values in statistics and custom reports.

To use the sudo parser in a specific log path

1. Navigate to **Log > Paths**.
2. Select the path where you want to use the parser.
3. In the **Parser** field, **Predefined** group, select **sudo_parser** from the drop-down list.

Figure 178: Log > Paths – Using the sudo_parser in the log path



4. Click **Commit**.

Parsing key-value pairs

The syslog-ng Store Box(SSB) appliance can separate a message consisting of key-value pairs (for example, Postfix log messages) into name-value pairs. The parsed values are automatically added to the metadata about the message, and you can display them on the SSB search interface as dynamic columns. You can specify the separator character to parse different log messages, for example, colon (:) to parse MySQL log messages, or the equal sign (=) for firewall logs. For details on when the key-value parser is executed related to other message processing operations, see the following list.

CAUTION:

If the names of keys in the message is the same as the names of SSB soft macros, the value from the parsed message will overwrite the value of the macro. For example, the `PROGRAM=value1, MESSAGE=value2` content will overwrite the `${PROGRAM}` and `${MESSAGE}` macros. To avoid overwriting such macros, use the `prefix()` option.

Hard macros cannot be modified, so they will not be overwritten. For details on the macro types, see .

The parser discards message sections that are not key=value pairs, even if they appear between key=value pairs that can be parsed.

1. Parse the message as a syslog message (unless message parsing is explicitly disabled for the source).
2. [Classify the message](#) using a pattern database.
3. [Modify the message using rewrite rules](#) (before filtering).
4. [Filter the messages](#), for example, based on sender hostname or message content. If the message does not match the configured filter, syslog-ng Store Box(SSB) will not send it to the destination.
5. Parse the text of the message (that is, the `${MESSAGE}` part) using a [key-value parser](#) or the [sudo parser](#).
6. [Modify the message using rewrite rules](#) (after filtering and other parsing).
7. SSB sends the message to the destinations set in the log path. The destinations are [local](#), [optionally encrypted files on SSB](#), or [remote servers, such as a database server](#).

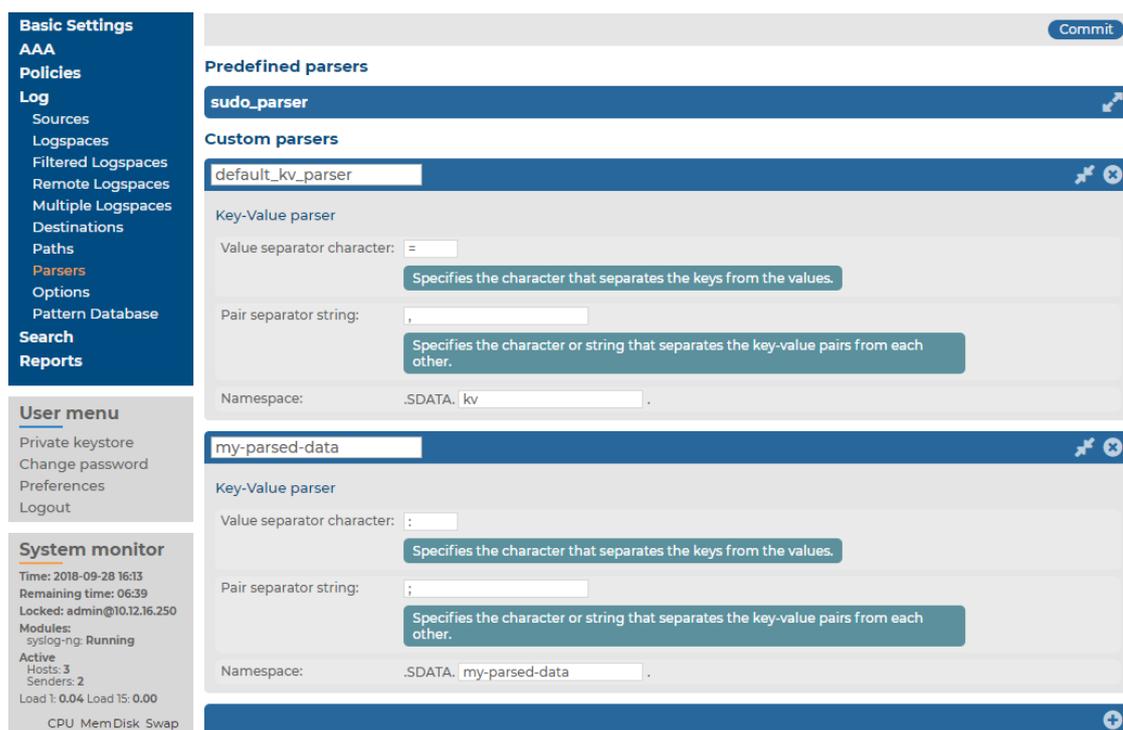
NOTE: If a log message contains the same key multiple times (for example, `key1=value1, key2=value2, key1=value3, key3=value4, key1=value5`), then SSB stores only the last (rightmost) value for the key. Using the previous example, SSB will store the following pairs: `key1=value5, key2=value2, key3=value4`.

NOTE: The names of the keys can contain only the following characters: numbers (0-9), letters (a-z,A-Z), underscore (_), dot (.), hyphen (-). Other special characters are not permitted.

To configure parsing key-value pairs

1. Navigate to **Log > Parsers** and select . A new parser is added to the list of parsers.

Figure 179: Log > Parsers — Creating a key=value parser

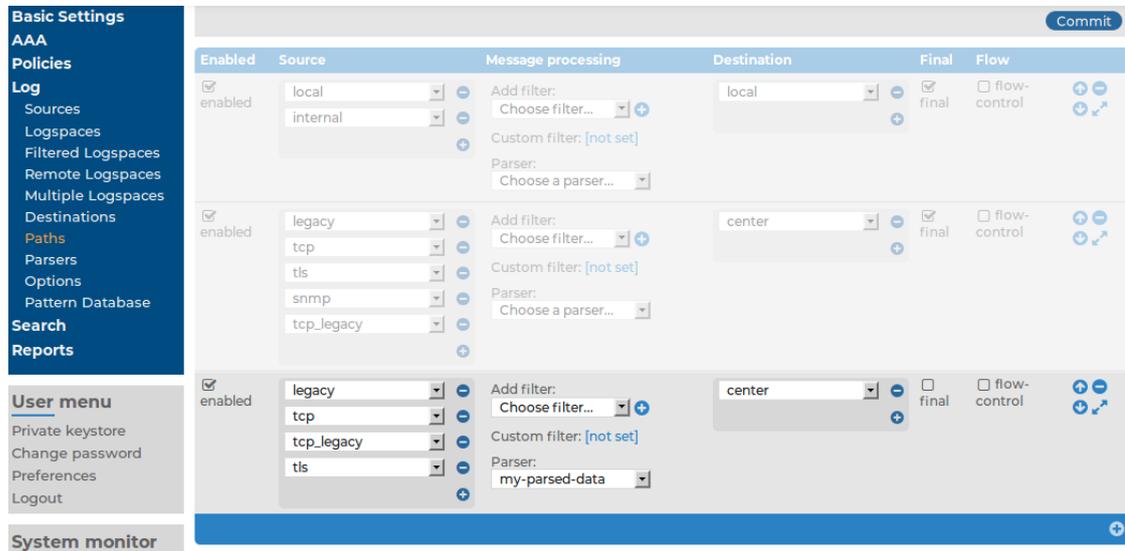


2. Enter a name for the parser.
3. Enter the character that separates the keys from the values in the incoming messages into the **Value separator character** field. For example, if your messages look like key1:value1, key2:value2, key3:value3, enter :.
4. Enter the character or string that separates the key-value pairs from each other into the **Pair separator string** field. For example, if your messages look like key1:value1, key2:value2, key3:value3, enter ;. If you use a string, it can be a maximum of 20 characters long.
5. Enter a prefix before the key part of the parsed key-value pairs to help further processing into the **Namespace** field. For example, to insert the my-parsed-data prefix, enter my-parsed-data. Note the following points:
 - SSB automatically adds the .SDATA. prefix before the value you enter into the **Namespace** field. That way these values are automatically included in the structured data (SDATA) part of the log message if you forward the message using the IETF-syslog protocol.
 - SSB automatically adds a dot (.) character as a separator between the namespace and the key parsed from the message.

For example, if you entered my-parsed-data as the namespace, and the keys in the message are key1, key2, and so on, then the full name of the macro that contains the parsed values is \${.SDATA.my-parsed-data.key1}, \${.SDATA.my-parsed-data.key2}, and so on. The parsed values are also automatically available as dynamic columns in the SSB search interface (the name of the column is the name of the macro).

6. Click **Commit**.
7. Navigate to **Log > Paths**.
8. Select the path where you want to use the parser.
9. In the **Parser** field, **Custom** group, select the parser you want to use in this log path.

Figure 180: Log > Paths – Using a key=value parser in the log path



10. Click **Commit**.

Configuring syslog-ng options

There are several options of the syslog-ng server running on syslog-ng Store Box(SSB) that can be configured. These include:

- For details on general syslog-ng settings — see [General syslog-ng settings](#) on page 321.
- For details on time stamping-related options — see [Time stamping configuration on SSB](#) on page 323.
- For details on certificate management for receiving and sending log messages in TLS-encrypted channels —, see [Setting the certificates used in TLS-encrypted log transport](#) on page 325.
- For details on managing domain name resolution for log messages — see [Using name resolution on SSB](#) on page 323.

General syslog-ng settings

To configure the general options of the syslog-ng server running on the syslog-ng Store Box(SSB) appliance, navigate to **Log > Options**. The following options are available (note that options related to name resolution are discussed in [Using name resolution on SSB](#) on page 323):

Figure 181: Log > Options — Configuring syslog-ng options

- Message size: Specifies the maximum length of incoming log messages in bytes. This option corresponds to the `log-msg-size()` parameter of `syslog-ng`. The maximum value of this parameter is `1000000` (1 MB).

NOTE: To be able to edit the **Message size**, you must have write/perform permission for the **Basic Settings > System** page. For details on how to assign user rights, see [Managing user rights and usergroups](#) on page 127.

- Wait time between polls: The time to wait in milliseconds before checking if new messages have arrived to a source. This option corresponds to the `time-sleep()` parameter of `syslog-ng`.
- Idle time before destination is closed: The time to wait in seconds before an idle destination file is closed. This option corresponds to the `time-reap()` parameter of `syslog-ng`.
- Cipher: Select the cipher method used to encrypt the logstore. The following cipher methods are available: `aes-128-cbc`, `aes-128-cfb`, `aes-128-cfb1`, `aes-128-cfb8`, `aes-128-ecb`, `aes-128-ofb`, `aes-192-cbc`, `aes-192-cfb`, `aes-192-cfb1`, `aes-192-cfb8`, `aes-192-ecb`, `aes-192-ofb`, `aes-256-cbc`, `aes-256-cfb`, `aes-256-cfb1`, `aes-256-cfb8`, `aes-256-ecb`, `aes-256-ofb`, `aes128`, `aes192`, `aes256`, `bf`, `bf-cbc`, `bf-cfb`, `bf-ecb`, `bf-ofb`, `cast5-cbc`, `cast5-cfb`, `des-cbc`, `des-cfb`, `des-cfb1`, `des-cfb8`, `des-ecb`, `des-edede`, `des-edede-cbc`, `des-edede-cfb`, `des-edede-ofb`, `des-edede3`, `des-edede3-cbc`, `des-edede3-cfb`, `des-edede3-ofb`, `des-ofb`, `desx-cbc`, `rc2-40-cbc`, `rc2-64-cbc`, `rc2-cbc`, `rc2-cfb`, `rc2-ecb`, `rc2-ofb`, `rc4`, and `rc4-40`.

By default, SSB uses the aes-256-cbc method.

- Digest: Select the digest method to use. The following digest methods are available: MD4, MD5, SHA-1, RIPEMD-160, SHA-224, SHA-256, SHA-384, and SHA-512.

By default, SSB uses the SHA-256 method.

⚠ CAUTION:

The size of the digest hash must be equal to or larger than the key size of the cipher method. For example, to use the aes-256-cbc cipher method, the digest method must be at least SHA-256.

Time stamping configuration on SSB

To configure the time stamping options of syslog-ng Store Box(SSB), navigate to **Log > Options**. The following options are available:

- Time stamp server: Select the time stamping server to use for signing encrypted logspaces. To use the built-in time stamp server of SSB, select **Local**.

To use an external time stamping server, select **Remote** and enter the address of the server into the **Server URL** field in the following format:

```
http://<IP address>:<port number>/
```

For example:

```
http://10.50.50.50:8080/
```

Note that currently only plain HTTP services are supported, password-protected and HTTPS services are not supported.

⚠ CAUTION:

SSB currently supports only time stamping servers that use the Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP) described in RFC 3161.

- Time stamp policy OID: If the Time Stamping Server has time stamping policies configured, enter the OID of the policy to use into the Timestamping policy field. SSB will include this ID in the time stamping requests sent to the TSA.

NOTE: The time stamp requests are handled by a separate process in syslog-ng, message processing is not affected if the time stamping server is slow or cannot be accessed.

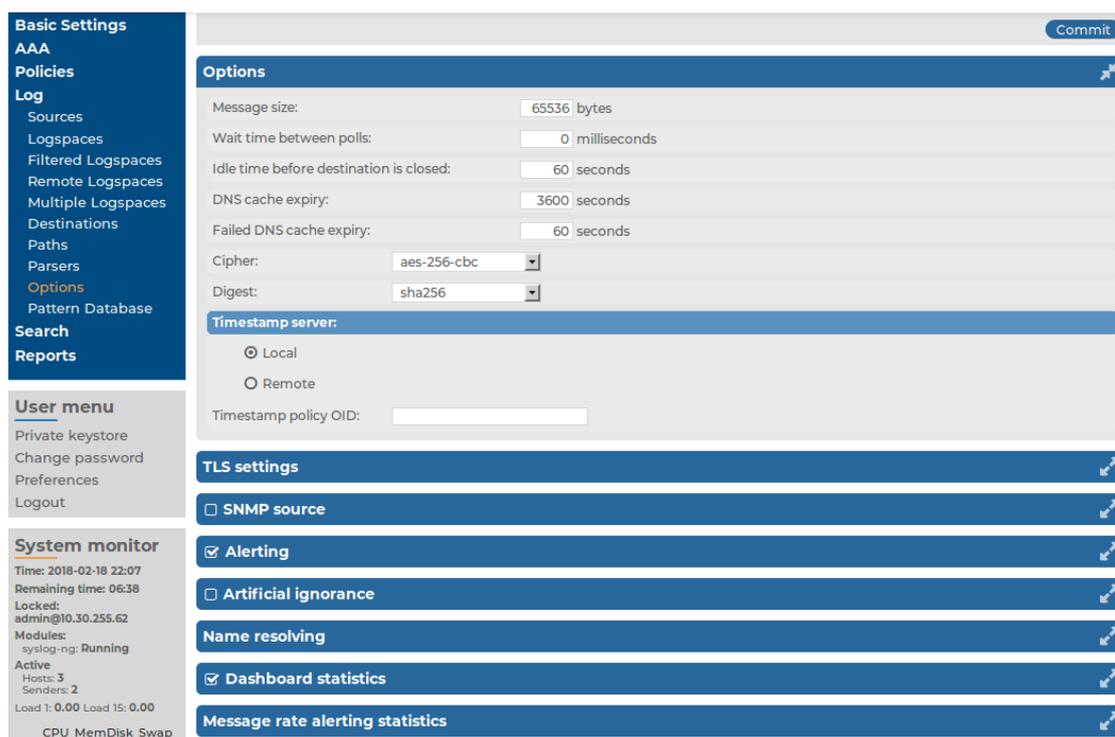
Using name resolution on SSB

The syslog-ng Store Box(SSB) appliance can resolve the hostnames of the clients and include them in the log messages. However, the performance of SSB can be severely

degraded if the domain name server is unaccessible or slow. Therefore, SSB automatically caches the results of name resolution. If you experience performance problems under high load, it is recommended to disable name resolution. If you must use name resolution, consider the following:

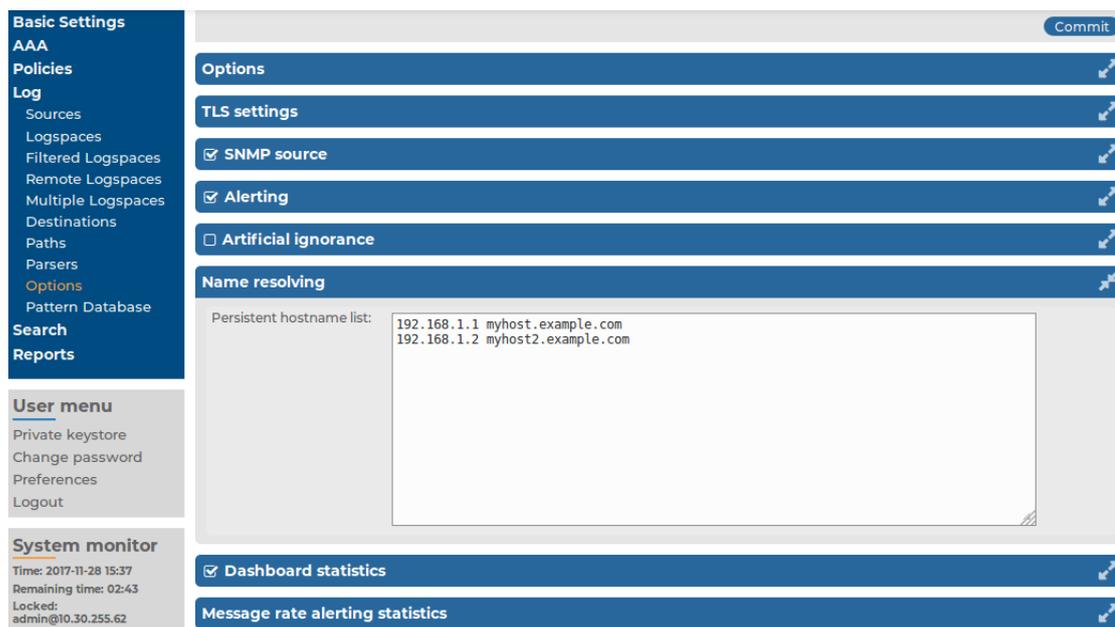
- If the IP addresses of the clients change only rarely, set the expiry of the DNS cache to a large value. By default, SSB caches successful DNS lookups for an hour, and failed lookups for one minute. These parameters can be adjusted under **Log > Options > Options > DNS Cache expiry** and **Failed DNS cache expiry**.

Figure 182: Log > Options > Options > DNS Cache expiry – Configuring DNS options



- Resolve the hostnames locally. Resolving hostnames locally enables you to display hostnames in the log files for frequently used hosts, without having to rely on a DNS server. The known IP address – hostname pairs are stored locally in a file. In the log messages, syslog-ng will replace the IP addresses of known hosts with their hostnames. To configure local name resolution, select **Log > Options > Name resolving**, and enter the IP Address - hostname pairs in (for example 192.168.1.1 myhost.example.com) into the **Persistent hostname list** field. Then navigate to **Log > Sources**, and set the **Use DNS** option of your sources to **Only from persistent configuration**.

Figure 183: Log > Options > Name resolving — Configuring persistent name resolution



Setting the certificates used in TLS-encrypted log transport

This section describes how to set a custom certificate and a CA certificate for encrypting the transfer of log messages.

NOTE: If you do not upload a certificate to encrypt the TLS-communication (that is, the **TLS certificate** and **TLS private key** options are not set), syslog-ng Store Box(SSB) uses the certificate and CA certificate set for the web interface (set under **Basic Settings > Management > SSL certificates**) for this purpose as well.

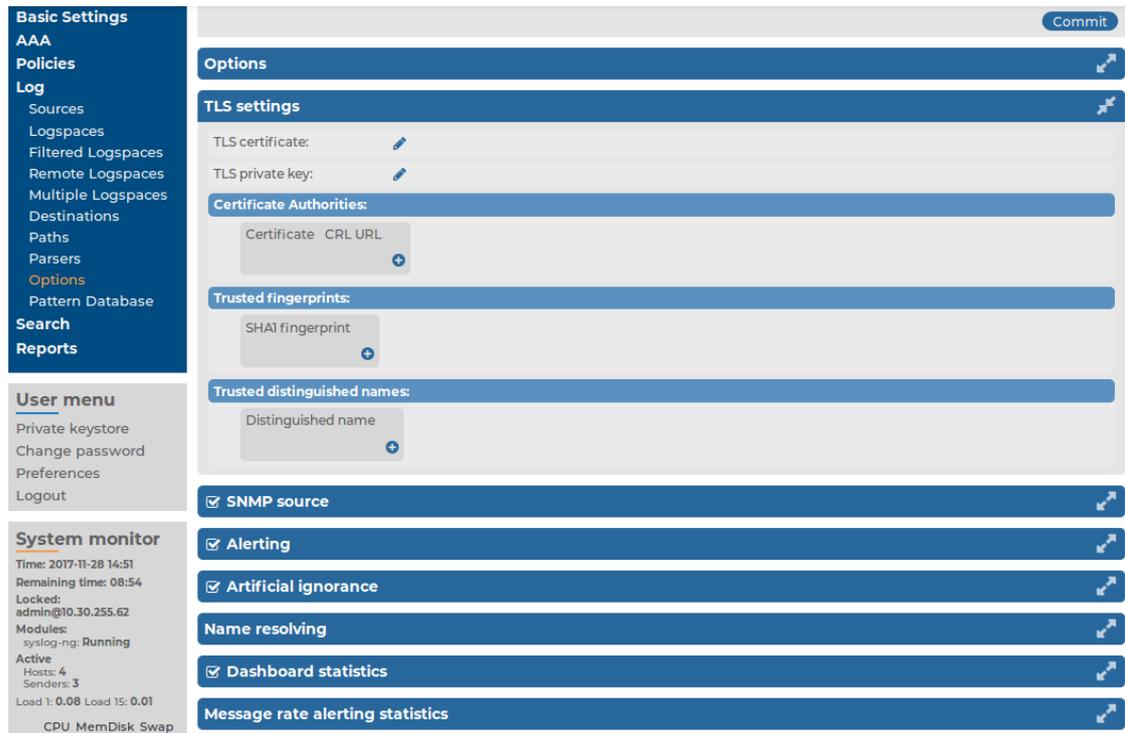
One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
- Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.

To set a custom certificate and a CA certificate for encrypting the transfer of log messages

1. In your PKI system, generate and sign a certificate for SSB, then navigate to **Log > Options > TLS settings**.
2. Click the  icon in the **TLS certificate** field to upload the certificate.

Figure 184: Log > Options > TLS settings — Configuring TLS settings for syslog-ng



To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**. Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

You can choose to upload a single certificate or a certificate chain (that is, intermediate certificates and the end-entity certificate).

After uploading a certificate or certificate chain, you can review details by clicking the name of the certificate, and looking at the information displayed in the pop-up window that comes up.

Figure 185: Log > Options > TLS settings — X.509 certificate details



The pop-up window allows you to:

- Download the certificate or certificate chain.
 | **NOTE:** Certificate chains can only be downloaded in PEM format.
- View and copy the certificate or certificate chain.
- Check the names and the hierarchy of certificates (if it is a certificate chain and there is more than one certificate present).
 On hovering over a certificate name, the subject of the certificate is displayed, describing the entity certified.
- Check the validity dates of the certificate or certificates making up the chain.
 On hovering over a particular date, the exact time of validity is also displayed.

After uploading the certificate or certificate chain, the presence or absence of the string (**chain**) displayed after the name of the certificate will indicate whether the certificate is a certificate chain or a single certificate.

3. Click the  icon in the **TLS private key** field and upload the private key corresponding to the certificate.
4. To set the certificate of the Certificate Authority (CA) used to verify the identity of the peers, click  in the **Certificate Authorities** field, then click .

Figure 186: Log > Options > TLS settings > Certificate Authorities — Uploading certificates

The screenshot shows a configuration window titled "Server X.509 certificate". It contains two main sections for adding certificates:

- Upload certificate:** This section includes a label "Upload:", a "Choose File" button, a text input field with the placeholder text "No file chosen", and an "Upload" button.
- Copy-paste certificate:** This section includes a label "Certificate:", a "Set" button, and a large, empty text area for pasting the certificate data.

To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**.

Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

Repeat this step to add more CA certificates if needed.

5. If the CA issues a Certificate Revocation List (CRL), enter its URL into the **CRL URL** field. SSB periodically downloads the list and refuses certificates that appear on the list.

NOTE: Note that only .pem format CRLs are accepted. CRLs that are in PKCS7 format (.crl) are not accepted.

6. If you want to accept connections only from hosts using certain certificates signed by the CA, click  in the **Trusted distinguished names** field and enter the distinguished name (DN) of the accepted certificates into the **Distinguished name** field. This option corresponds to the trusted-dn() parameter of syslog-ng.

Example: *, O=Example Inc, ST=Some-State, C=* accepts only certificates issued for the Example Inc organization in Some-State state.

7. If you want to accept connections only from hosts using certain certificates that have specific SHA-1 fingerprints, click  in the **Trusted fingerprints** field and enter the SHA-1 fingerprint of the accepted certificates into the **SHA-1 fingerprint** field. This option corresponds to the `trusted-keys()` parameter of `syslog-ng`.

Example: `00:EF:ED:A4:CE:00:D1:14:A4:AB:43:00:EF:00:91:85:FF:89:28:8F,0C:42:00:3E:B2:60:36:64:00:E2:83:F0:80:46:AD:00:A8:9D:00:15` adds these specific SHA-1 fingerprints:

`00:EF:ED:A4:CE:00:D1:14:A4:AB:43:00:EF:00:91:85:FF:89:28:8F` and
`0C:42:00:3E:B2:60:36:64:00:E2:83:F0:80:46:AD:00:A8:9D:00:15`.

NOTE: When using the `trusted-keys()` and `trusted-dn()` parameters at the same time, note the following:

- If the fingerprint of the peer is listed in the `trusted-keys()` parameter and the DN of the peer is listed in the `trusted-dn()` parameter, then the certificate validation is performed.
- If either the fingerprint of the peer is not listed in the `trusted-keys()` parameter or the DN of the peer is not listed in the `trusted-dn()` parameter, then the authentication of the peer fails and the connection is closed.

Searching log messages

This section describes how to browse the log messages collected on syslog-ng Store Box(SSB).

- [Using the search interface](#) on page 330 explains how to use and customize the search interface, describes the log message data that is available on SSB, and provides examples of the the wildcard and boolean search operators you can use.
- [Browsing encrypted logspaces](#) on page 345 describes how to decrypt and browse encrypted logspaces.
- [Creating custom statistics from log data](#) on page 350 explains how to create custom statistics from the available log data, and how to save them for reports.
- [Creating content-based alerts](#) describes how to create content-based alerts.
- [Additional tools](#) on page 359 provides information about functionalities that allow you to obtain further data about log messages from pattern database alerts and reports.

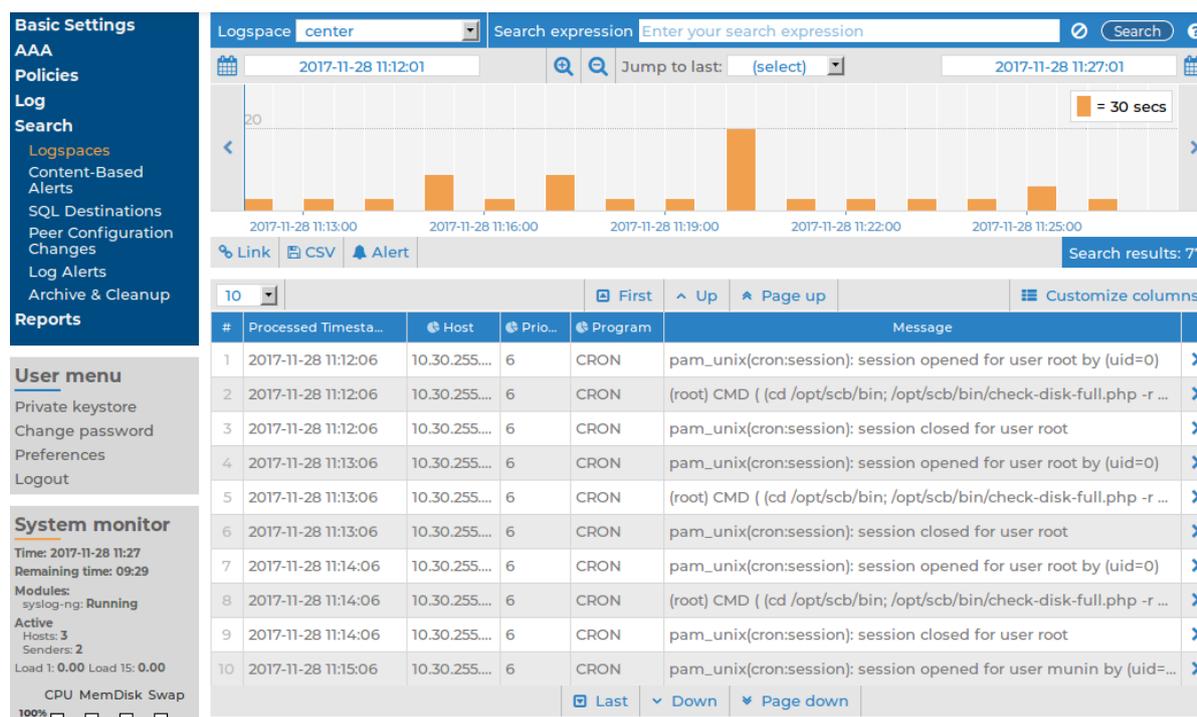
Using the search interface

The syslog-ng Store Box(SSB) appliance has a search interface for browsing the collected log messages. You can choose the logspace, enter a search expression, specify the timeframe, and browse the results here.

This section walks you through the main parts of the search interface.

To access the search interface, navigate to **Search > Logspaces**.

Figure 187: Search > Logspaces — The log message search interface



Logspaces:

To choose the appropriate logspace, use the **Logspace name** menu. Note that you cannot access plain text logspaces on the SSB search interface.

For more information on the available logspaces, and how to configure them, see ["Storing messages on SSB" in the Administration Guide](#).

Search:

On the log message search interface, you can use the **Search expression** field to search the full list of log messages. Search expressions are case insensitive, with the exception of operators (like AND, OR, etc.), which must always be capitalized. Click the  icon, or see [Using complex search queries](#) for more details.

When searching log messages, the capabilities of the search engine depend on the delimiters used to index the particular logspace. For details on how to configure the delimiters used for indexing, see ["Creating logstores" in the Administration Guide](#).

NOTE: You can search in indexed logspaces even if log traffic is disabled.

You can create complex searches using wildcards and boolean expressions. For more information and practical examples, see [Using complex search queries](#).

NOTE: SSB only indexes the first 59 characters of every name-value pair (parameter). This has two consequences:

- If the parameter is longer than 59 characters, an exact search might deliver multiple, imprecise results.

Consider the following example. If the parameter is:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

SSB indexes it only as:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

This corresponds to the first 59 characters. As a result, searching for:

```
nvpair:.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

returns all log messages that contain:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

- Using wildcards might lead to the omission of certain messages from the search results.

Using the same example as above, searching for the value:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-12345
```

does not return any results (as the 12345 part was not indexed). Instead, you have to search for:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-*
```

This, as explained above, might find multiple results.

Overview:

Displays the number of log messages in the selected time interval.

Figure 188: Search > Logspaces — Log message overview



Use the  and  icons to zoom, and the arrows to display the previous or the next intervals. To change the timeframe, you can:

- Change the beginning and the end date.
- Click and drag the pointer across a period on the calendar bars to select a specific interval and zoom in.
- Use the **Jump to last** option to select the last 15 minutes, hour, 6 hours, day, or week.

Hovering the mouse above a bar displays the number of results, and the start and end date of the period that the bar represents. Click a bar to display the results of that period in the table. Use **Shift+Click** to select multiple bars.

Action bar:

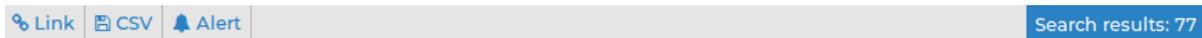
The search interface provides an action bar that allows you to:

- Fetch a [link to a search query](#).
- Export search results into a [csv](#) file.
- Create a [content-based alert](#).

It also displays the following information:

- [Error and warning messages](#).
- The number of [search results](#) returned by a search query.

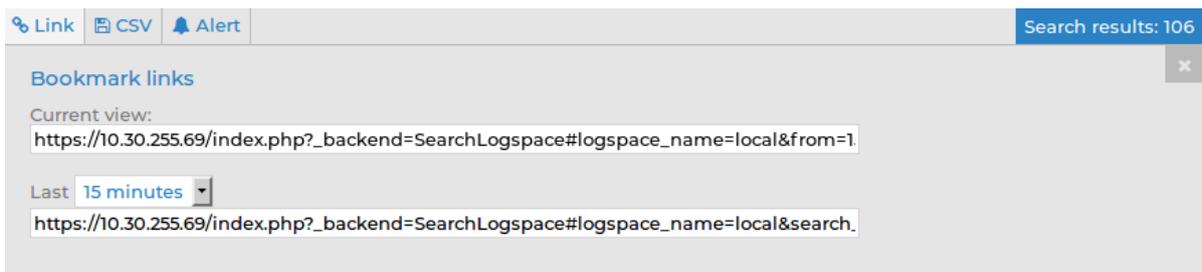
Figure 189: Search > Logspaces: Action bar



Link to a search query:

On clicking , the **Bookmark links** panel is displayed:

Figure 190: Search > Logspaces – Bookmark links panel



Bookmark links allow you to fetch a link to a search query so that you can:

- Share your search queries with colleagues, who can then access the relevant search results in one click.
- Save frequently used search queries as bookmark links.

The link in the **Current view** field provides a direct link to your search query and its results currently displayed on your screen. Whenever you open the bookmarked link from your browser, it will always return the same, fixed set of results. The start and end date that you set when executing the search query and fetching the link from the **Bookmark links** panel remain fixed.

The **Last** menu, on the other hand, allows you to specify an interval of time, for example, the last 15 minutes or the last hour, and fetch search results generated within that period. The search results that you access using this link may differ on two different occasions as the start point of the specified interval is always the moment you open the bookmarked link from your browser.

CSV export:

On clicking , the **CSV export** panel is displayed:

Figure 191: Search > Logspaces — CSV export panel



Clicking  exports your search results into a CSV file. 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 Download CSV export to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load. If you regularly need a large portion of your data in plain text format, consider using the SSB RPC API (for details, see ["The SSB RPC API" in the Administration Guide](#)), or sharing the log files on the network and processing them with external tools (for details, see ["Accessing log files across the network" in the Administration Guide](#)).

Alert:

The alert functionality enables you to set up content-based alerts for search expressions of your choice. You will receive an alert when a match is found between the search expression and the contents of a log message. Note that the alerts are generated for only those log messages that are stored in the logspace(s) for which you set up the alert.

For detailed information on content-based alerts, see ["Creating content-based alerts" in the Administration Guide](#).

Errors and warnings:

When any user action results in an error condition (for example, if you enter an invalid search expression, display statistics for a column that has not been indexed), an error or warning notification will be displayed on the action bar. Errors are shown in red letters, warnings are displayed in amber.

If there is more than one notification, the latest will be displayed and the number of notifications triggered will also be indicated. Clicking the notification will open an **Errors and warnings** panel:

Figure 192: Search > Logspaces — Errors and warnings panel



The **Errors and warnings** panel displays a list of errors/warnings with their time stamp and details of their cause.

You can clear notifications one by one by clicking  next to the them, or clear all of them by clicking .

Search results:

After running a search query, the action bar displays the number of search results returned by the query. This is useful information when you are trying to find out how often a certain element appears in the logs.

List of log messages:

Use the arrow keys and the Page Up and Page Down keys to navigate the listed log messages, or use the mouse wheel to scroll. You can disable mouse wheel scrolling in your **User menu > Preferences**. If data is too long to fit on one line, it is automatically wrapped and only the first line is displayed.

Figure 193: Search > Logspaces — List of log messages

#	Processed Timesta...	Host	Prio...	Progra...	Message
1	2017-11-28 13:27:48	ssbdemo	6	index-loc...	Indexer performance statistics; tokenizer_queue_length='1', toke...
2	2017-11-28 13:27:48	ssbdemo	6	index-ce...	Indexer performance statistics; tokenizer_queue_length='1', toke...
3	2017-11-28 13:28:01	ssbdemo	6	CRON	pam_unix(cron:session): session opened for user root by (uid=0)
4	2017-11-28 13:28:01	ssbdemo	6	CRON	(root) CMD ([-x /opt/ssb/bin/check-disk-full.php] && (cd /opt/ssb/...
5	2017-11-28 13:28:01	ssbdemo	6	CRON	pam_unix(cron:session): session closed for user root
6	2017-11-28 13:29:01	ssbdemo	6	CRON	pam_unix(cron:session): session opened for user root by (uid=0)
7	2017-11-28 13:29:01	ssbdemo	6	CRON	(root) CMD ([-x /opt/ssb/bin/check-disk-full.php] && (cd /opt/ssb/...
8	2017-11-28 13:29:01	ssbdemo	6	CRON	pam_unix(cron:session): session closed for user root
9	2017-11-28 13:29:01	ssbdemo	6	index-ce...	Scan directory for potential days of logs (YYYY/MM-DD); directory='...
10	2017-11-28 13:29:01	ssbdemo	6	index-loc...	Scan directory for potential days of logs (YYYY/MM-DD); directory='...

Details of a log message:

To expand a row in the list of log messages, click [>](#). The complete log message is displayed:

Figure 194: Search > Logspaces — Viewing a single log message

Message 6 of 2407	
Processed timestamp:	2018-02-23 12:37:01
Host:	real
Facility:	10
Unique ID:	810062992740732938
Tags:	
Message:	pam_unix(cron:session): session closed for user root
Dynamic columns:	mytag=ALMA sdata.timequality.issynced=0

Use the arrow keys to jump to the previous or the next log message.

Use the Page Up and Page Down to jump to the 10th log message before or after the currently displayed log message. You can also jump to the previous or the next log message with the mouse wheel.

If the displayed log message consists of several pages of data, you can configure the mouse wheel to be able to use it for scrolling the message vertically. To do this, navigate to **User menu > Preferences**, deselect **Mousewheel scrolling of search results** and click **Set options**. This will disable jumping between log messages with the mouse wheel.

You can perform the following actions:

- Click any word in the message to copy it to the Search field.
- Click any of the dynamic columns (name-value pairs) to add it as a column to the list

of log messages.

- Click any of the  icons to view the statistics of the selected category.

To return to the list of all log messages, click .

Customizing columns of the log message search interface

The following describes how to customize the data displayed on the log message search interface.

To customize the data displayed on the log message search interface

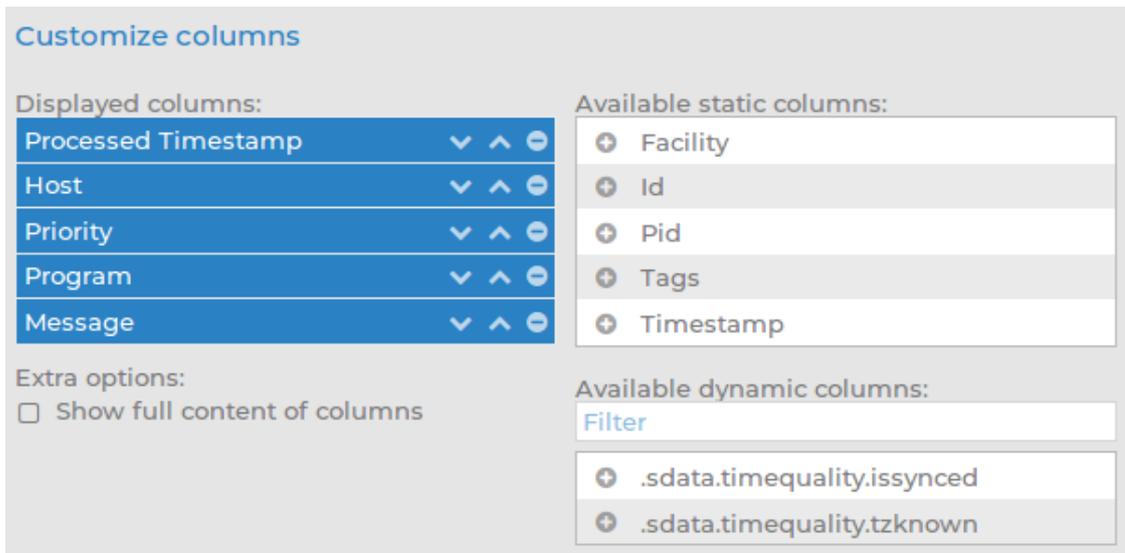
1. Click **Customize columns**.

The parameters used for the columns when displaying log messages are listed under **Displayed columns**. All other available parameters are listed under **Available static columns** and **Available dynamic columns**.

Dynamic columns are created from structured data parameters (name-value pairs) in log messages stored on syslog-ng Store Box(SSB). Structured data parameters are detected and added to the list of customizable columns automatically. (For more information on the structured data part of log messages, see "[The STRUCTURED-DATA message part](#)" in the [Administration Guide](#).)

NOTE: To export the search results into a CSV file, click  on the action bar. Note that the CSV file includes all the static columns and the displayed dynamic columns.

Figure 195: Search > Logspaces > Customize columns — Customizing columns of the log message search interface



Customize columns

Displayed columns:

Processed Timestamp	▼ ▲ -
Host	▼ ▲ -
Priority	▼ ▲ -
Program	▼ ▲ -
Message	▼ ▲ -

Available static columns:

+ Facility
+ Id
+ Pid
+ Tags
+ Timestamp

Extra options:

Show full content of columns

Available dynamic columns:

Filter

+ .sdata.timequality.issynced
+ .sdata.timequality.tzknown

2. To add a static column to the **Displayed columns**, click  .
3. To add a dynamic column to the **Displayed columns**, choose a name-value pair from **Available dynamic columns** and click  .

The selected name generates a new, separate dynamic column with a **<name>** heading (where **<name>** is the name of the key). The relevant values are displayed in the cells of the respective column.

4. To remove parameters from the **Visible columns**, click  .
5. To display the full content of each column (including the log messages), enable **Show full content of columns**.

Metadata collected about log messages

The following information is available about the log messages:

- **Processed Timestamp:** The date when syslog-ng Store Box(SSB) received the log message in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Timestamp:** The time stamp received in the message — the time when the log message was created in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Facility:** The facility that sent the message.
- **Priority:** The priority value of the message.
- **Program:** The application that created the message.
- **Pid:** The program identifier of the application that created the message.
- **Host:** The IP address or hostname of the client that sent the message to SSB.
- **Message:** The text of the log message.
- **Tag:** Tags assigned to the message matching certain pattern database rules.
- **Id:** Unique ID of the message.
- **classifier.rule_id:** ID of the pattern database rule that matched the message.
- **classifier.class:** Description of the pattern database rule that matched the message.
- Dynamic columns, created from additional name-value pairs, might also be available.

Using complex search queries

You can use wildcards and boolean expressions, and search specific parts of the log messages collected on syslog-ng Store Box(SSB).

NOTE: When searching log messages, the capabilities of the search engine depend on the delimiters used to index the particular logspace. By default, the indexer uses the following delimiter characters to separate the message into words (tokens): `& : ~ ? ! [] = , ; () ' "`. For details on how to configure the delimiters used for indexing, see ["Creating logstores" in the Administration Guide](#).

NOTE: It is not possible to search for the whitespace () character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

The following sections provide examples for different search queries:

- For examples of exact matches, see [Searching for exact matches and using complex queries](#) on page 339.
- For examples of using boolean operators to combine search keywords, see [Combining search keywords](#) on page 339.
- For examples of wildcard searches, see [Using wildcard searches](#) on page 340.
- For examples of searching for special characters, see [Searching for special characters](#) on page 342.
- For examples of searching in a specific part of the message, see [Searching in a specific part of the message](#) on page 343.
- For examples of searching name-value pairs, see [Searching the name-value pairs of the message](#) on page 343.

Searching for exact matches and using complex queries

By default, SSB searches for keywords as whole words in the MESSAGE part of the log message and returns only exact matches.

Example: Searching for exact matches

Search expression	example
Matches	example Example EXAMPLE
Does not match	examples example.com query-by-example exam

Combining search keywords

You can use boolean operators - AND, OR, and NOT - to combine search keywords. Note that the boolean operators are case sensitive, and must be in all caps. More complex search expressions can also be constructed with parentheses.

Example: Combining keywords in search

Search expression	keyword1 AND keyword2
-------------------	-----------------------

Matches	(returns log messages that contain both keywords)
---------	---

Search expression	keyword1 OR keyword2
-------------------	----------------------

Matches	(returns log messages that contain at least one of the keywords)
---------	--

Search expression	keyword1 AND NOT keyword2
-------------------	---------------------------

Matches	(returns log messages that contain only keyword1)
---------	---

To search for expressions that can be interpreted as boolean operators (for example: AND), use the following format: `message:AND`.

Example: Using parentheses in search

Use parentheses to create more complex search expressions:

Search expression	(keyword1 OR keyword2) AND keyword3
-------------------	-------------------------------------

Matches	(returns log messages 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 when trying to find 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.

Search expression	example?
Matches	example1 examples
Does not match	example.com example12 query-by-example example?
Search expression	?example?
Matches	1example2
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. Wildcard characters also work in any message part, for example, `program:postfix*`.

Search expression	example*
Matches	example examples example.com

Does not match	query-by-example example*
----------------	------------------------------

Search expression	*example
-------------------	----------

Matches	example query-by-example
---------	-----------------------------

Does not match	example.com example12
----------------	--------------------------

Search expression	*example*
-------------------	-----------

Matches	example query-by-example example.com example12
---------	---

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 question mark (?), asterisk (*), backslash (\) or whitespace () characters, you must prefix these characters with a backslash (\). Any character after a backslash is handled as a character to be searched for.

NOTE: Delimiter characters are an exception to the rule. It is not possible to search for delimiter characters, even when they are prefixed.

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 the backslash character, use two backslashes (\\).

Search expression	C:\\Windows
Matches	C:\Windows
Search expression	nvpair:path=C:\\Program\ Files
Matches	C:\Program Files

Searching in a specific part of the message

You can search in a specific part of the message using the <type>: prefix. The message: (or msg:) prefix means the message part and can be omitted. For example, use the program: prefix to search for the name of an application, or use the host: prefix to search for a host name, and so on.

Example: Searching specific parts of messages

Search expression	program:syslog-ng
Matches	All log messages from the syslog-ng application.

Searching the name-value pairs of the message

You can search the structured data part of log messages using the nvpair: prefix. Use the = delimiter to separate the name and the value of structured data parameters, and remove the quote marks from the values.

Example: Searching the structured data part of messages

Search expression `nvpair:.sdata.win@18372.4.event_type=Alert`

Matches All log messages where there is a win@18372.4 element with the event_type="Alert" parameter. For example:
`[win@18372.4 EVENT_TYPE="Alert"]`

Example: Using wildcard * to search the structured data

You can use the asterisk (*) wildcard to broaden the search to all structured data elements.

Search expression `nvpair:*event_type=Alert*`

Matches All log messages where the "event_type" name has the "Alert" value.

Example: Searching for parameter names

To search for a specific name, add the "=" character after the name.

Search expression `nvpair:*event_type=*`

Matches All log messages where an "event_type" name exists.

Example: Searching for parameter values

To search for a specific value, add the "=" character before the value.

Search expression `nvpair:*=Alert`

Matches All log messages where a name has the "Alert" value.

NOTE: SSB only indexes the first 59 characters of every name-value pair (parameter). This has two consequences:

- If the parameter is longer than 59 characters, an exact search might deliver multiple, imprecise results.

Consider the following example. If the parameter is:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

SSB indexes it only as:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

This corresponds to the first 59 characters. As a result, searching for:

```
nvpair:.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345
```

returns all log messages that contain:

```
.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-
```

- Using wildcards might lead to the omission of certain messages from the search results.

Using the same example as above, searching for the value:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-12345
```

does not return any results (as the 12345 part was not indexed). Instead, you have to search for:

```
nvpair:*=2011-12-08T12:32:25.024+01:00-hostname-*
```

This, as explained above, might find multiple results.

Search performance tips

To decrease the load on SSB when searching and receive your search results faster, note the following points.

- Use as small a time range as possible
- Prefer AND instead of OR
- Avoid unneeded wildcard characters, such as * and ?
- Use wildcard characters at the end of the tokens if possible

Browsing encrypted logspaces

By default, you cannot browse encrypted logstores from the syslog-ng Store Box(SSB) web interface, because the required decryption keys are not available on SSB. To make browsing and searching encrypted logstores possible, SSB provides the following options:

- Use persistent decryption key(s) for a single user.
For details, see [Using persistent decryption keys](#) on page 346.
- Use decryption keys for the duration of the user session only.
For details, see [Using session-only decryption keys](#) on page 348.
- Assign decryption keys to a logstore (making them available to every SSB user). This option might raise security concerns.
For details, see [Assigning decryption keys to a logstore](#) on page 349.

NOTE: Do not use SSB's own keys and certificates for encrypting or decrypting.

One Identity recommends:

- Using 2048-bit RSA keys (or stronger).
- Using the SHA-256 hash algorithm (or stronger) when creating the public key fingerprint.

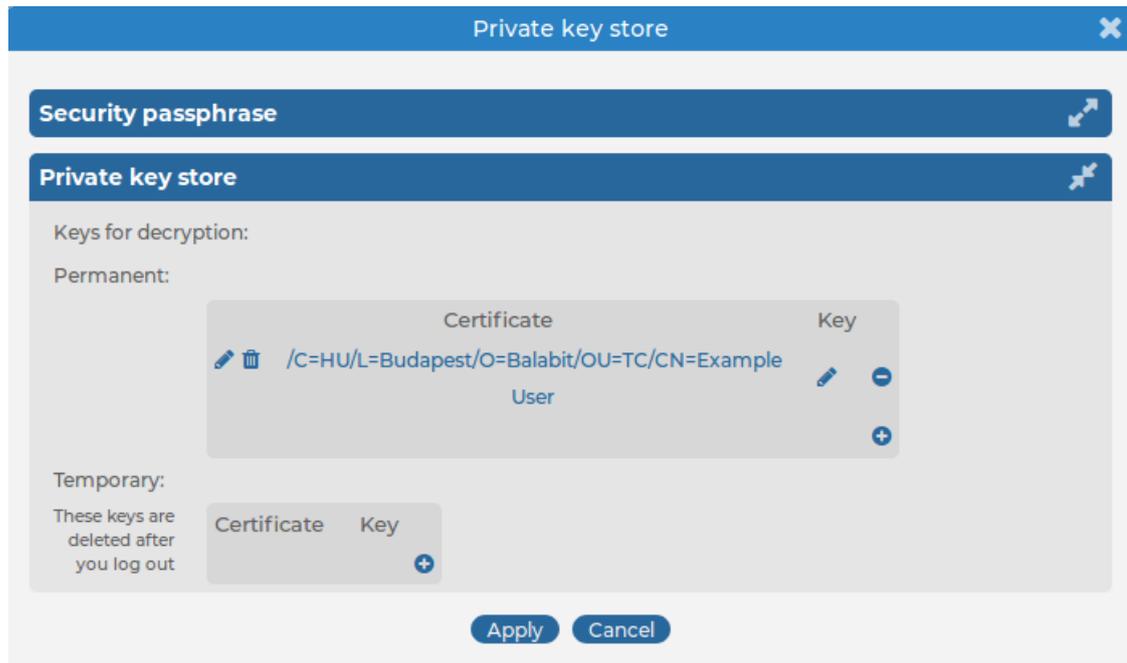
Using persistent decryption keys

You can upload decryption keys and bind them to your account. The decryption keys are stored on syslog-ng Store Box(SSB), but they are only made available for this user account, and can also be protected (encrypted) with a passphrase.

To use persistent decryption keys

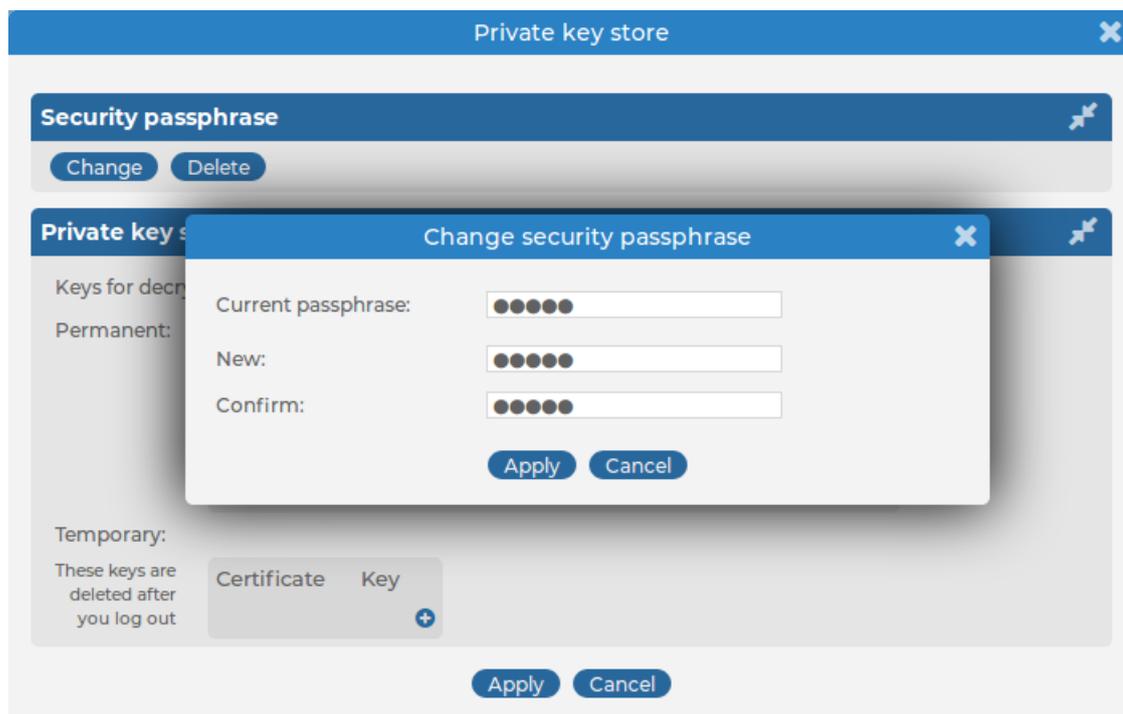
1. Select **User menu** > **Private keystore**. A pop-up window is displayed.
2. Select **Permanent** > , then select **Certificate** > . A pop-up window is displayed.

Figure 196: User menu > Private keystore — Adding decryption keys to the private keystore



3. Paste or upload the certificate used to encrypt the logstore.
4. Select **Key** > . A pop-up window is displayed.
5. Paste or upload the private key of the certificate used to encrypt the logstore.
6. Repeat Steps 2-5 to upload additional keys if needed.
7. Select **Security passphrase** > **Change**, and enter a passphrase to protect the private keys.

Figure 197: User menu > Private keystore — Securing the private keystore with a passphrase



8. Click **Apply**.

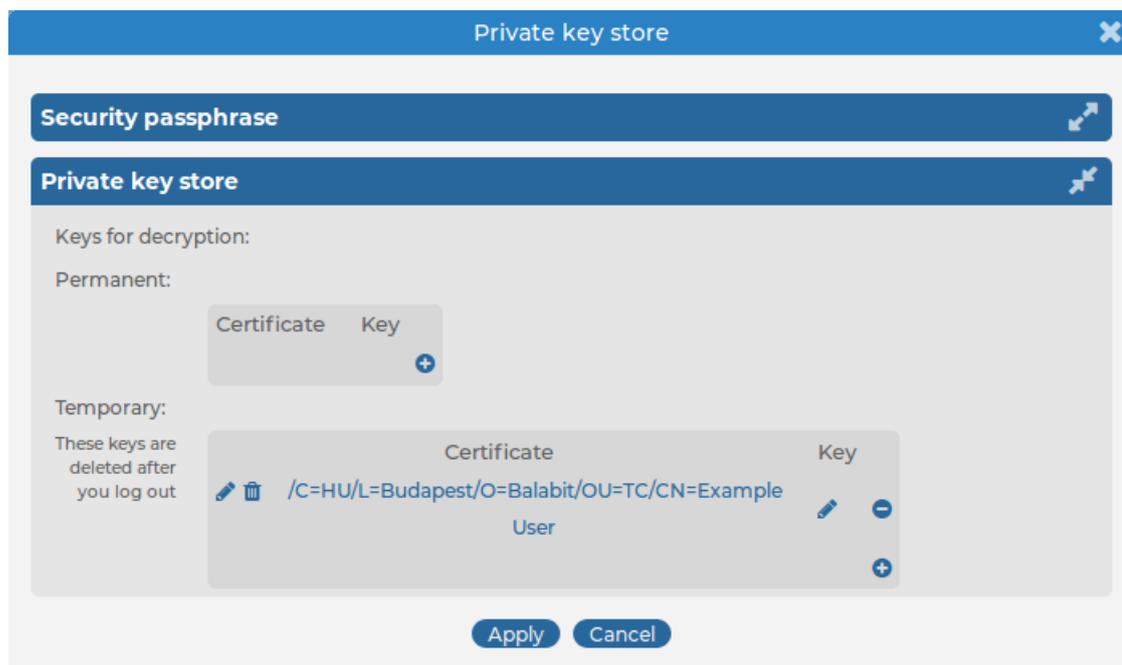
Using session-only decryption keys

You can upload decryption keys to browse encrypted logspaces for the duration of the session only. These keys are automatically deleted when you log out from syslog-ng Store Box(SSB).

To use session-only decryption keys

1. Select **User menu > Private keystore**. A pop-up window is displayed.
2. Select **Temporary** > , then select **Certificate** > . A pop-up window is displayed.

Figure 198: User menu > Private keystore — Adding decryption keys to the private keystore



3. Paste or upload the certificate used to encrypt the logstore.
4. Select **Key** > . A pop-up window is displayed.
5. Paste or upload the private key of the certificate used to encrypt the logstore.
6. Repeat Steps 2-5 to upload additional keys if needed.
7. Click **Apply**.

Assigning decryption keys to a logstore

You can add a private key (or set of keys) to a logstore, and use these keys to decrypt the logstore files. This way, anyone who has the right to search a particular logspace can search the messages. These decryption keys are stored unencrypted in the syslog-ng Store Box(SSB) configuration file.

As this may raise security concerns, avoid this solution unless absolutely necessary.

To assign decryption keys to a logstore

1. Navigate to **Log > Logspaces** and select the encrypted logspace you want to make searchable for every user via the SSB web interface.
2. Select **Decryption private keys** > . A pop-up window is displayed.

Figure 199: Log > Logspaces — Adding decryption keys to a logstore



3. Paste or upload the private key of the certificate used to encrypt the logstore.
4. Repeat Steps 2-3 to upload additional keys if needed.

An additional key is needed when the certificate used to encrypt a logstore expires. When this happens, you have to upload a new certificate. However, to be able to read the logstore encrypted with the old (expired) certificate(s), you need to keep the old encryption key(s) with the new one.

5. Click **Commit**.

Creating custom statistics from log data

The syslog-ng Store Box(SSB) appliance can create statistics from the Facility, Priority, Program, Pid, Host, Tags, and `.classifier.class` columns. Use **Customize columns** to add the required column, if necessary.

NOTE: The `.classifier.class` data is the class assigned to the message when pattern database is used. For details, see "[Classifying messages with pattern databases](#)" in the

[Administration Guide](#). The pattern databases provided by One Identity currently use the following message classes by default: system, security, violation, or unknown.

You can display statistics on the web interface, export the related data as CSV, and also save the statistics to include in a report.

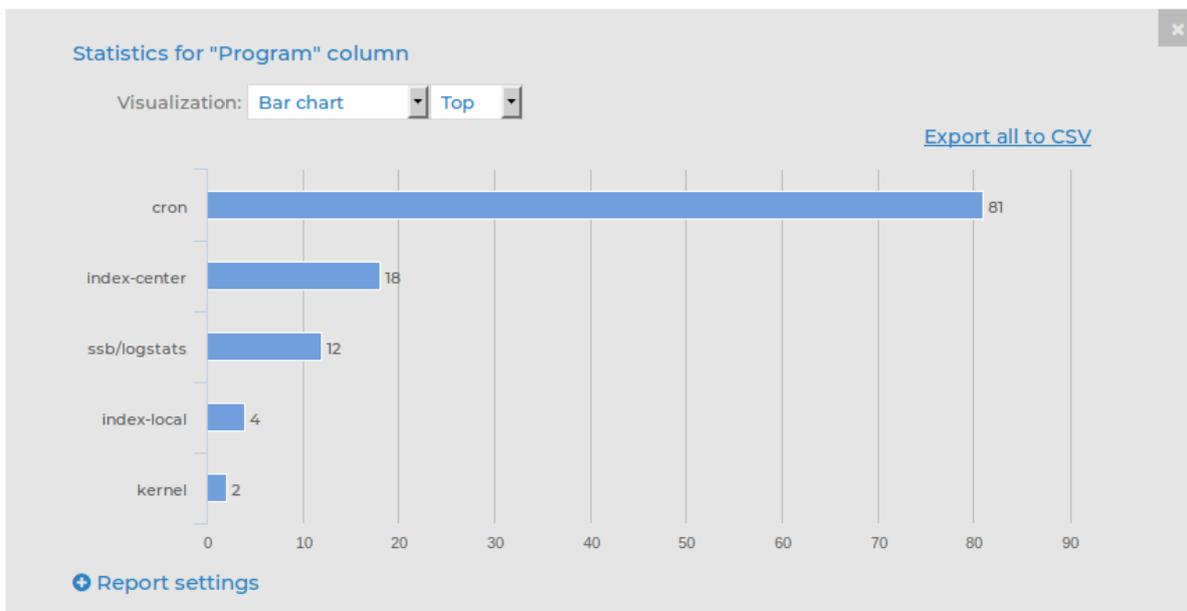
Displaying log statistics

To display statistics about the log messages, click the  icon in the appropriate header of the table.

You can choose from **Bar chart** or **Pie chart & List**.

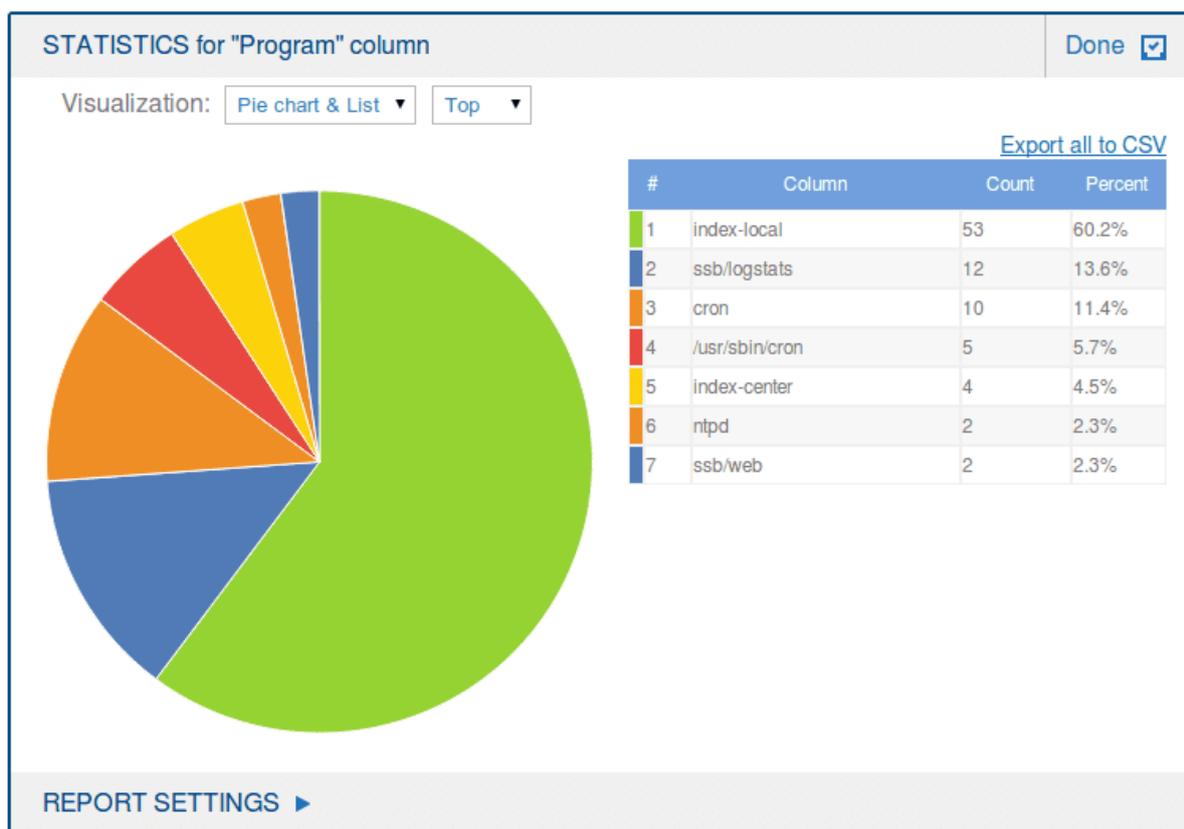
NOTE: For performance reasons, when creating statistics for a **Multiple Logspace** (see ["Creating multiple logspaces" in the Administration Guide](#)), syslog-ng Store Box(SSB) does not create statistics if the data upon which the statistics is based (for example, the hostname) has over 1000 entries in any of the member logspaces. In this case, SSB displays the Number of member statistics has too many entries error message.

Figure 200: Search > Logspaces — Displaying log statistics as Bar chart



In **Pie chart & List** view, percentages add up to 100%. The only exception to this is when statistics are based on **Tags**. Since statistics are provided for tags rather than messages, when messages have multiple tags, the percentages may add up to more than 100%.

Figure 201: Search > Logspaces — Displaying log statistics as Pie chart & List



Statistics will show the item with the largest number of entries first. To display the item with the least number of entries first, select **Least**.

NOTE: When navigating to the "future" in the search bar, it is possible that the number of logs displayed in the **Search results** differs from the number of logs displayed in the **Count** part of the **Host** pie chart.

To avoid this, do not navigate to the "future".

If this has already happened, save the search expression that you have used somewhere, and then refresh the page by clicking **Log > Search** again. Note that it will display the original state of the Search page, meaning that for example it will remove all search expressions that you have entered before.

You can export these statistics in CSV format using the **Export all to CSV** option, or you can include them in reports as a subchapter.

CAUTION:

Do not use Export all to CSV to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load. If you regularly need a large portion of your data in plain text format, consider using the syslog-ng Store Box(SSB) RPC API (for details, see "The SSB RPC API" in the Administration Guide), or sharing the log files on the network and processing them with external tools (for details, see "Accessing log files across the network" in the Administration Guide).

Creating reports from custom statistics

You can save log statistics to include them in reports as a subchapter.

Figure 202: Search > Logspaces — Creating reports from custom log statistics

Statistics for "Program" column

Visualization: **Pie chart & List** **Top**

[Export all to CSV](#)

#	Column	Count	Percent
1	cron	81	69.2%
2	index-center	18	15.4%
3	ssb/logstats	12	10.3%
4	index-local	4	3.4%
5	kernel	2	1.7%

Report settings

Report subchapter name:

Visualization: **List** **Top**

Number of entries: **15**

Grant access for the following user groups:

- **+**

Save as Report subchapter

1. In the **Statistics** view, click **Report settings**.
2. Add a name for the statistics in the **Report subchapter name** field.
3. Select the **Visualization** for the report: List, Pie chart, or Bar chart.
4. Choose how the entries are sorted: descending (**Top**) or ascending (**Least**).
5. Choose the **Number of entries** to include.

NOTE: Selecting **All** includes only the first 1000 results. The remaining results are aggregated as 'others'.

NOTE: For performance reasons, when creating statistics for a **Multiple Logspace** (see "[Creating multiple logspaces](#)" in the [Administration Guide](#)), syslog-ng Store Box(SSB) does not create statistics if the data upon which the statistics is based (for example, the hostname) has over 1000 entries in any of the member logspaces. In this case, SSB displays the Number of member statistics has too many entries error message.

6. Select the user group that can access the subchapter in the **Grant access for the following user groups** field.
7. Click **Save as Report subchapter**.
8. To add the saved subchapter to a report, follow the instructions provided in [Configuring custom reports](#).

Creating content-based alerts

The syslog-ng Store Box(SSB) appliance can create content-based alerts about log messages based on specific search expressions. Search queries are run every few seconds and an alert is triggered whenever a match between the contents of a log message and a search expression is found. Alerts are collected and sent to a pre-defined email address (or email addresses).

Some log messages might have particular significance and therefore getting notifications about those can often be more efficient than searching for them manually.

You can set up or modify alerts for local logspaces or those logspaces to which you have the relevant privileges, meaning that:

- Either the relevant user group has been assigned read and write/perform access to the **Search > Logs** object on the **AAA > Access Control** page.
- Or the user group has been added under the **Access control** option of the relevant logspace on the **Log > Logspaces** page.

There are two ways to create alerts, using the search interface or the **Search > Content-Based Alerts** page:

- For details on how to set up alerts on the search interface, see [Setting up alerts on the search interface](#).
- For details on how to set up alerts on the **Search > Content-Based Alerts** page, see [Setting up alerts on the Search > Content-Based Alerts page](#).

NOTE: Content-based alerting is currently not available for filtered, multiple, and remote logspaces.

NOTE: In the case of encrypted logspaces, no decryption key is required for content-based alerting to work. SSB has access to the log messages while processing them, and the indexer and content-based alerting services run before encryption happens.

Setting up alerts on the search interface

This section describes how to set up alerts using the search interface.

To set up alerts using the search interface

1. Configure a target where you wish to send your content-based alerts.

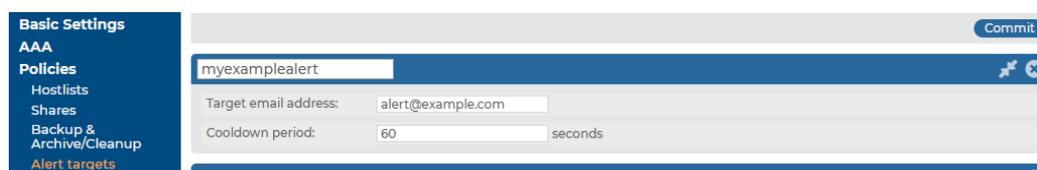
Alert targets are set up and modified by superusers or user groups that have been assigned read and write/perform access to the Policies object on the **AAA > Access Control** page.

To specify an alert target:

- a. Go to **Policies > Alert targets**.
- b. Click .

The new tab that opens allows you to record an alert target.

Figure 203: Policies > Alert targets – Alert targets page



- c. Enter a name for your alert target.
- d. In the **Target email address** field, enter the email address where you wish to send alerts.

NOTE: Alert target names must be unique.

NOTE: You can specify only one email address per target. However, you can add multiple targets per alert, which allows you to send a specific alert to more than one email addresses (if required).

- e. In the **Cooldown period** field, enter the minimum amount of time (in seconds) that should pass between the sending of two alert messages to this target.

The minimum value is 60 seconds, and the maximum value is 999999 seconds.

NOTE: An alert message is sent only when a match is found between the contents of log messages and a search expression. This means that if no match is found, more time may pass between two alert messages than the interval specified as the cooldown period.

- f. Click **Commit** to save your details.

Expected result:

You have successfully configured a target for your alert where alert messages will be sent.

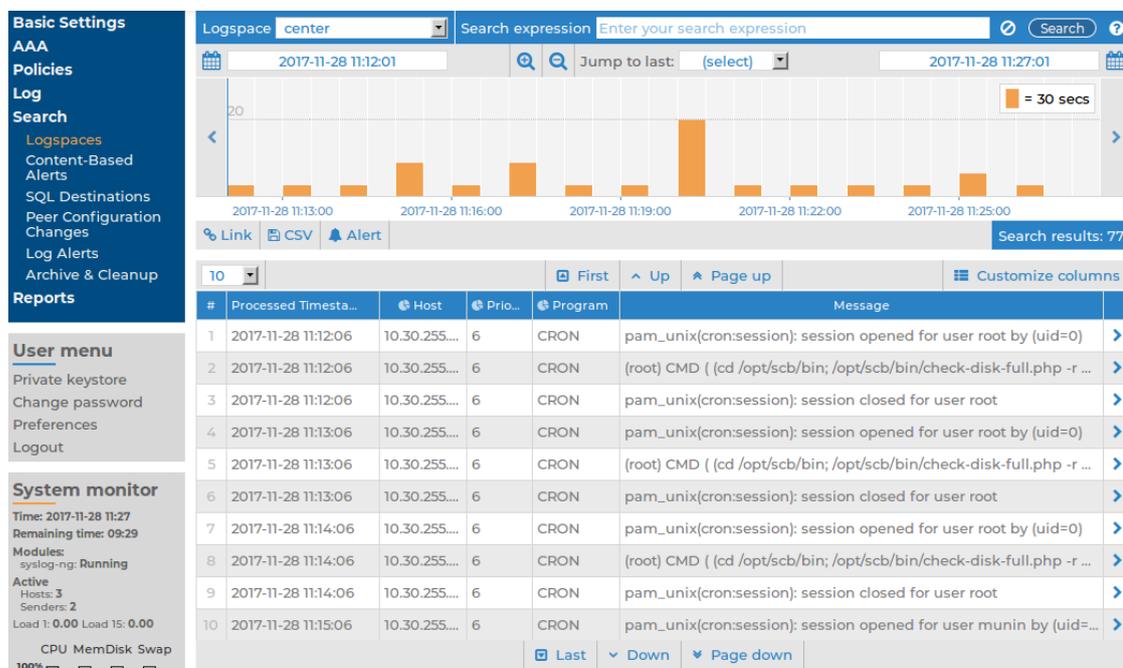
2. *Optional step:* You can also specify the email address from which the alerts are sent to your targets. Configuring an email address from where you wish to receive emails can be useful for filtering purposes. If you do not specify such an email address, a

default one will be used.

For detailed instructions, see the steps describing how to specify a **Send e-mails as** email address in "[Configuring e-mail alerts](#)" in the [Administration Guide](#).

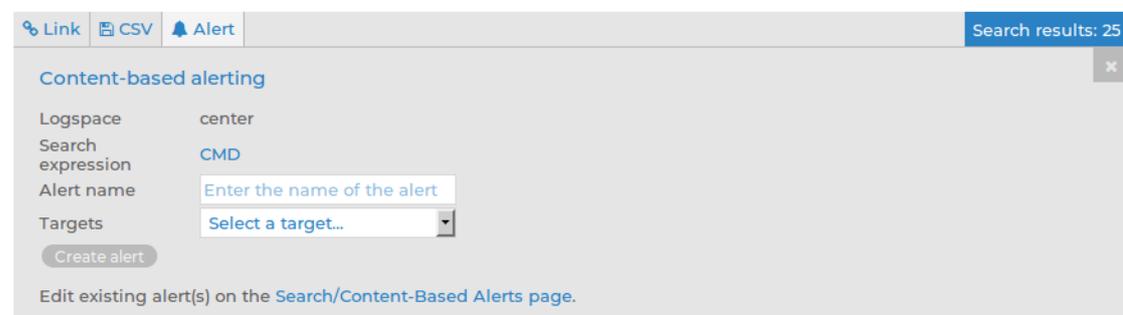
3. Once you have set up a target or targets, navigate to the search interface by going to **Search > Logspaces**.

Figure 204: Search > Logspaces — Setting up alerts on the search interface



4. In the **Logspace name** menu, select the relevant logspace.
5. In the **Search expression** field, enter the search expression that you wish to receive alerts about and click **Search**.
6. To configure additional details for the alert, click **Alert**. The **Content-based alerting** panel is displayed.

Figure 205: Search > Logspaces — Content-based alerting panel



The **Logspace** field displays the name of the logspace that you have selected from the **Logspace name** menu. The **Search expression** field displays the search expression that you entered in the **Search expression** field.

7. Enter a name for your alert in the **Alert name** field.

NOTE: Alert names must be globally unique. Using a prefix before alert names can help avoid specifying a name that is already in use.

8. Select a target from **Targets**. You can select multiple targets if you wish to distribute the alert to multiple email addresses.

You can remove targets you have already added by clicking  in front of the target's name.

9. To save your details, click .

NOTE: If you wish to modify your alert later on, you can make changes via **Search > Content-Based Alerts**. For details, see [Setting up alerts on the Search > Content-Based Alerts page](#).

Setting up alerts on the Search > Content-Based Alerts page

This section describes how to set up alerts on the **Search > Content-Based Alerts** page.

To set up alerts on the Search > Content-Based Alerts page

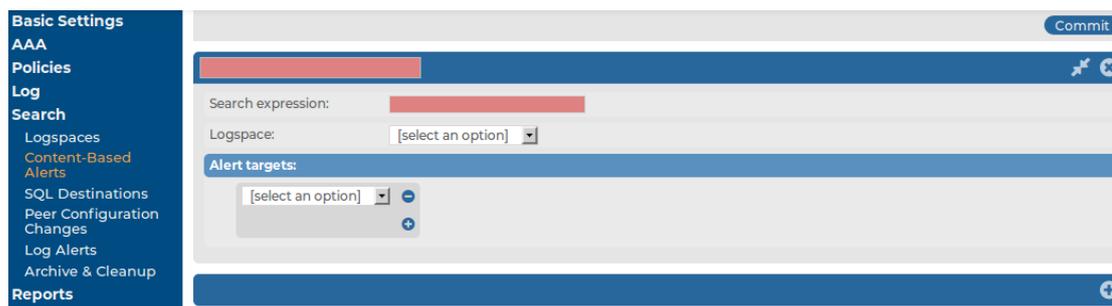
1. Configure a target where you wish to send content-based alerts. For details on how to do this, see Step 1 in [Setting up alerts on the search interface](#).
2. *Optional step:* You can also specify the email address from which alerts are sent. Configuring an email address from where you wish to receive emails can be useful for filtering purposes. If you do not specify such an email address, a default one will be used.

For detailed instructions, see the steps describing how to specify a **Send e-mails as** email address in "[Configuring e-mail alerts](#)" in the [Administration Guide](#).

3. Once you have set up a target or targets, navigate to **Search > Content-Based Alerts**.
4. Click .

The new tab that opens allows you to specify a content-based alert.

Figure 206: Search > Content-Based Alerts — Setting up content-based alerts on the Search



5. Enter a name for your alert.

NOTE: Alert names must be globally unique. Using a prefix before alert names can help avoid specifying a name that is already in use.

6. In the **Search expression** field, enter the search expression that you wish to receive alerts about.
7. Select the appropriate logspace from the **Logspace** menu.
8. Select a target or targets from the **Alert targets** menu. You can select multiple targets if you wish to distribute the alert to multiple email addresses.

You can remove targets you have already added by clicking .

9. To save your details, click **Commit**.

NOTE: If you wish to modify your alert later on, you can make changes by revisiting the relevant steps on the **Search > Content-Based Alerts** page.

Format of alert messages

Once content-based alerts have been created, syslog-ng Store Box(SSB) will send alert messages to the configured targets.

The alert email's subject line will follow this format:

```
Alert: [myalert][mylogspace]
```

Alert messages will be presented in the following format:

```
Alert: There were at least 10000 matches between Mon 18 Apr 2016 10:45:38 CEST
and Mon 18 Apr 2016 10:45:43 CEST on
* logspace: "<mylogspace>"
* alert: "<myalert>"
* search expression: "<mysearchexpression>"
```

To review these matches on your SSB appliance, see:

```
https://<IP_address_of_SSB>:<port_number>/index.php?_
backend=SearchLogspace#logspace_name=mylogspace&
from=1460976338&to=1460976343&search_expression=mysearchexpression
```

Note: You will not receive a new alert message for a cooldown period of 1 minute for this alert.

Note that the contents of the log messages are not shared in the alert message. A URL is provided to direct users to their SSB appliance.

Additional tools

The syslog-ng Store Box(SSB) appliance provides additional tools to obtain information about log messages that can come from external sources. They are as follows:

- Pattern database: You can use the pattern database of SSB to alert on certain log messages. If you are using the pattern database for such purposes and you wish to check the history of the alerts raised by SSB, then refer to [Log message alerts](#) on page 366.
- Reports: SSB periodically creates reports on processed traffic. If you wish to retrieve information available in such reports, see [Reports](#) on page 371.

Searching the internal messages of SSB

The syslog-ng Store Box(SSB) appliance allows you to search, filter, and export internal messages. These internal messages contain the logs created by SSB itself (not the messages collected from external sources), including log messages of the SSB appliance, configuration changes, notifications, alerts, and dashboard statistics.

Log messages of the SSB appliance:

- All available log messages are listed in the **local** logspace in **Search > Logspaces**. For detailed instructions on using the log search interface, see [Using the search interface](#).
- Recent log messages are also available in **Basic settings > Troubleshooting**. For detailed instructions on using the troubleshooting tools, see [Troubleshooting SSB](#) on page [424](#).

Configuration changes:

- The configuration-related activity of SSB users and administrators is available at **AAA > Accounting**. The configuration changes performed on the SSB web interface are all listed here. For the list of displayed parameters, see [Changelogs of SSB](#) on page [364](#).
- Peers (client computers) that use syslog-ng Premium Edition 3.0 or newer send a special log message to SSB when their configuration is modified. These changes are listed at **Search > Peer configuration change**. For the list of displayed parameters, see [Configuration changes of syslog-ng peers](#) on page [366](#).

Alerts and notifications:

- If you use the pattern database of SSB to alert on certain log messages, then a history of the alerts is available at **Search > Alerts**.

For the list of displayed parameters, see [Log message alerts](#) on page 366.

- Backup and archive notifications, including errors encountered during backup or archiving, are stored at **Search > Archive & Cleanup**.

For the list of displayed parameters, see [Notifications on archiving and backups](#) on page 367.

Dashboard statistics and reports:

- The statistics of SSB are available at **Basic settings > Dashboard**.

For detailed information and the list of available options, see [Status history and statistics](#) on page 368.

- PDF reports about the configuration changes, system health parameters, and other activities of SSB are available at **Reporting > Reports**.

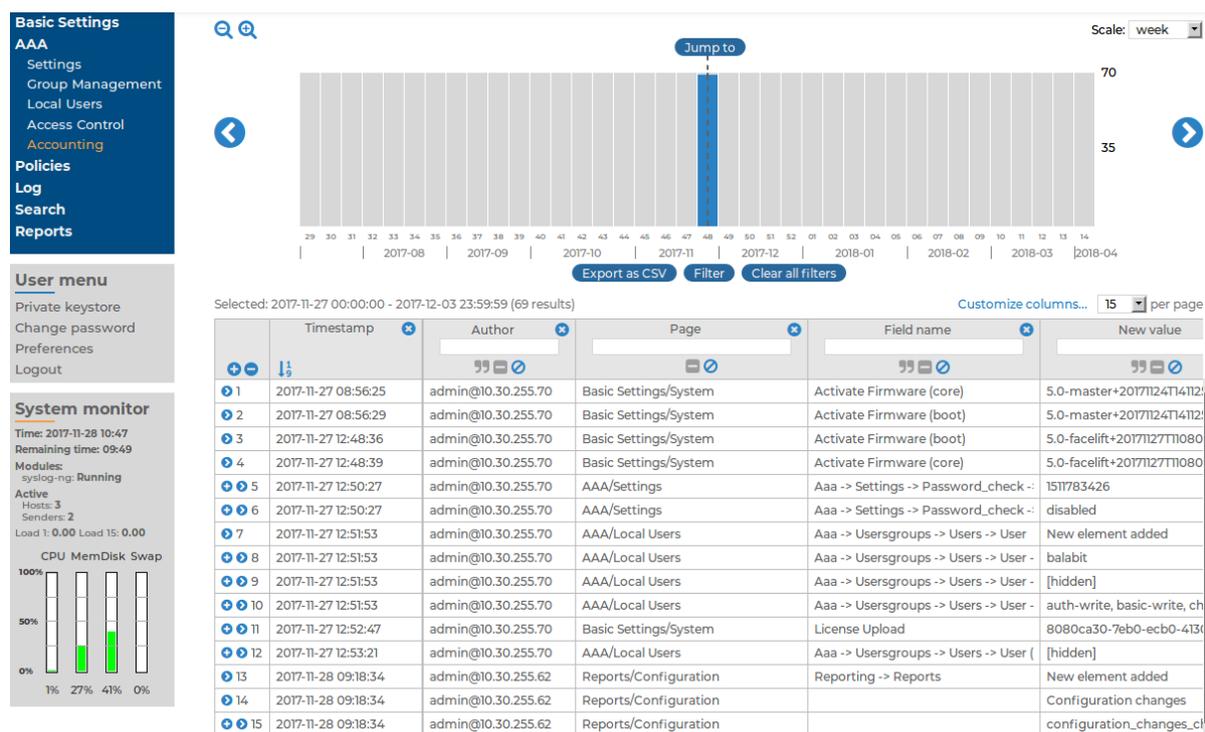
For the list of displayed parameters, see [Reports](#) on page 371.

Using the internal search interfaces

The internal search interfaces that allow you to browse and filter the configuration changes, alerts, notifications, and reports of syslog-ng Store Box(SSB) are located across various pages. The way the user interface works, however, is uniform across all these pages. This section walks you through the main functionalities that are available to you when browsing internal messages.

The example in [AAA > Accounting — An example of an internal search interface](#) on page 362 shows the **AAA > Accounting** page but all the search interfaces listed under [Configuration changes:](#), [Alerts and notifications:](#), and [Dashboard statistics and reports:](#) have similar features and look and feel.

Figure 207: AAA > Accounting — An example of an internal search interface



The bars display the number of log messages 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 also displays partial matches. For example, filtering the **Author** column on the **AAA > Accounting** page for and 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 all to CSV** to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load. If you regularly need a large portion of your data in plain text format, consider using the syslog-ng Store Box(SSB) RPC API (for details, see ["The SSB RPC API" in the Administration Guide](#)), or sharing the log files on the network and processing them with external tools (for details, see ["Accessing log files across the network" in the Administration Guide](#)).

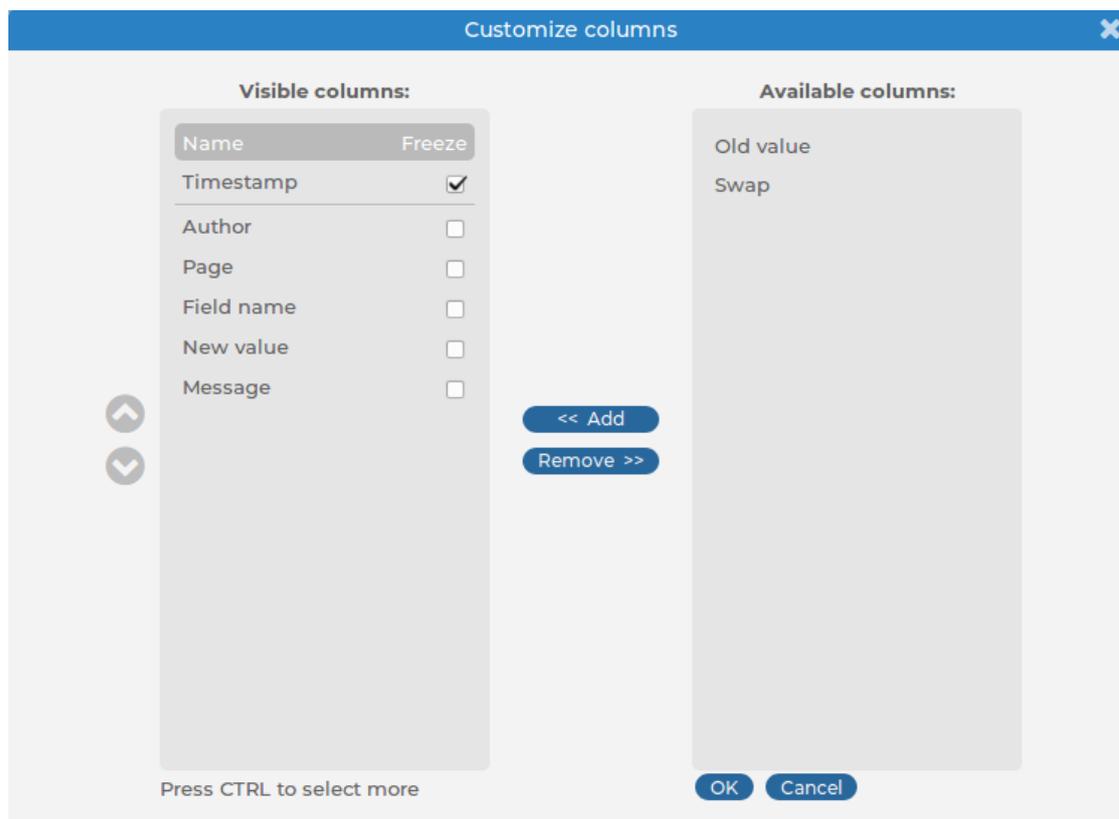
Customizing columns of the internal search interfaces

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, **AAA > Accounting**.
2. Click **Customize Columns**. A pop-up window containing the list of visible and available columns is displayed.

Figure 208: AAA > Accounting > Customize Columns — Customizing columns of the search interfaces



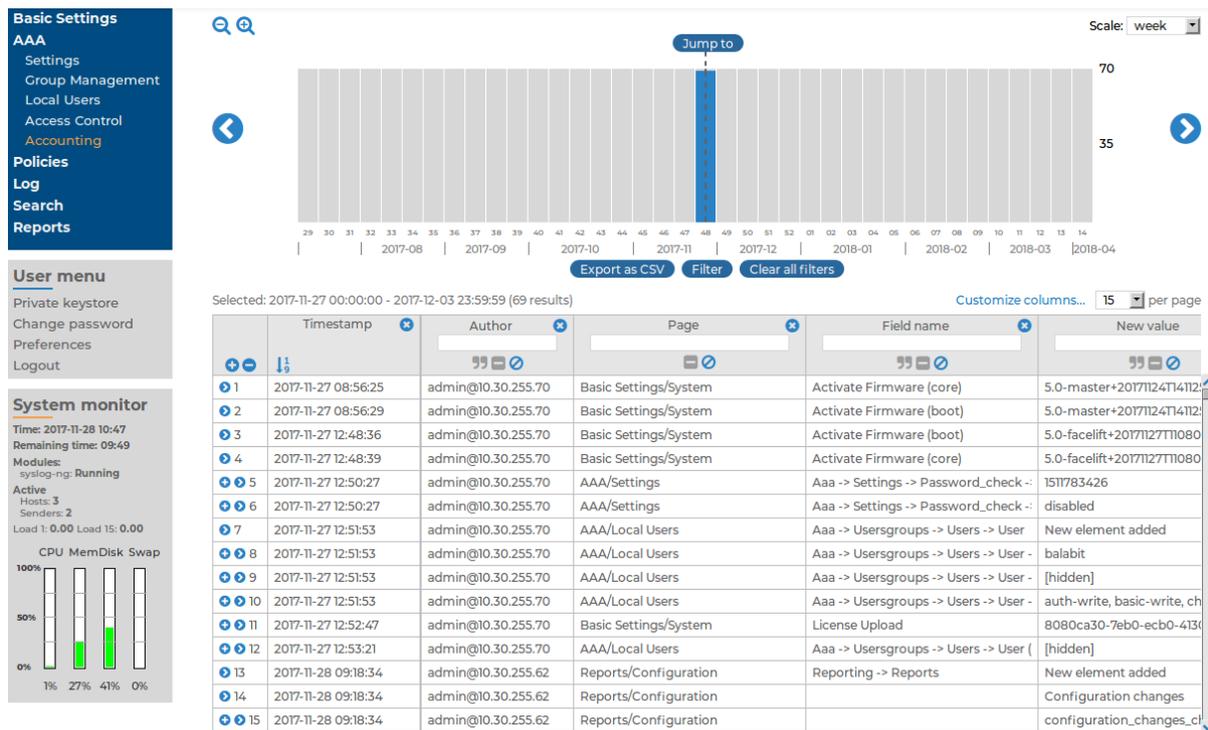
3. The displayed parameters are listed in the **Visible columns** field. All other available parameters are listed in the **Available columns** field.
 - To add parameters to the **Visible columns** field, select the desired parameter (s) and click **Add**.
 - To remove parameters from the **Visible 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.

Changelogs of SSB

The syslog-ng Store Box(SSB) appliance automatically records the activity of its users and administrators. These activities are displayed at **AAA > Accounting**. The following information is available:

Figure 209: AAA > Accounting — Displaying configuration changes



- **Timestamp:** The date when the modification was committed in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Author:** The SSB user who performed the modification.
- **Page:** The main menu item that was modified (for example, Basic Settings > Management).
- **Field name:** The name of the field on the page that was modified.
- **New value:** The new value of the field after the modification.
- **Description:** The changelog entered by the SSB administrator. Changelogs are available only if the **AAA > Settings > Require commit log** option was enabled at the time of the change.
- **Old value:** The original value of the field.
- **Swap:** Indicates if the order of objects was modified on the page (for example the order of two policies in the list).

For details on how to navigate around the user interface and interact with features such as filtering and exporting results, and customizing what data is displayed, see [Using the internal search interfaces](#) on page 361.

Configuration changes of syslog-ng peers

Peers running syslog-ng Premium Edition 3.0-6.0.x automatically send a notification to syslog-ng Store Box(SSB) when their configuration has changed since the last configuration reload or restart. Note that peers running syslog-ng Premium Edition version 7.0.x do not send such notifications. These log messages are available at **Search > Peer Configuration Change**. Note that the log messages do not contain the actual modification, only indicate that the configuration was modified. The following information is available:

- **Timestamp:** The time stamp received in the message — the time when the log message was created in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Hostname:** The hostname or IP address of the client whose configuration has been changed.
- **Version:** The version number of the syslog-ng application that sent the message.
- **Sender address:** The IP address of the client or relay that sent the message directly to SSB.
- **Signature:** The signature of the syslog-ng client.
- **Fingerprint:** The SHA-1 hash of the new configuration file.

For details on how to navigate around the user interface and interact with features such as filtering and exporting results, and customizing what data is displayed, see [Using the internal search interfaces](#) on page 361.

Log message alerts

When using the pattern database, syslog-ng Store Box(SSB) raises alerts for messages that are classified as Violation. The history of these alerts is available at **Search > Alerts**. The following information is available about the alerts:

Figure 210: Search > Log Alerts — Displaying alert messages



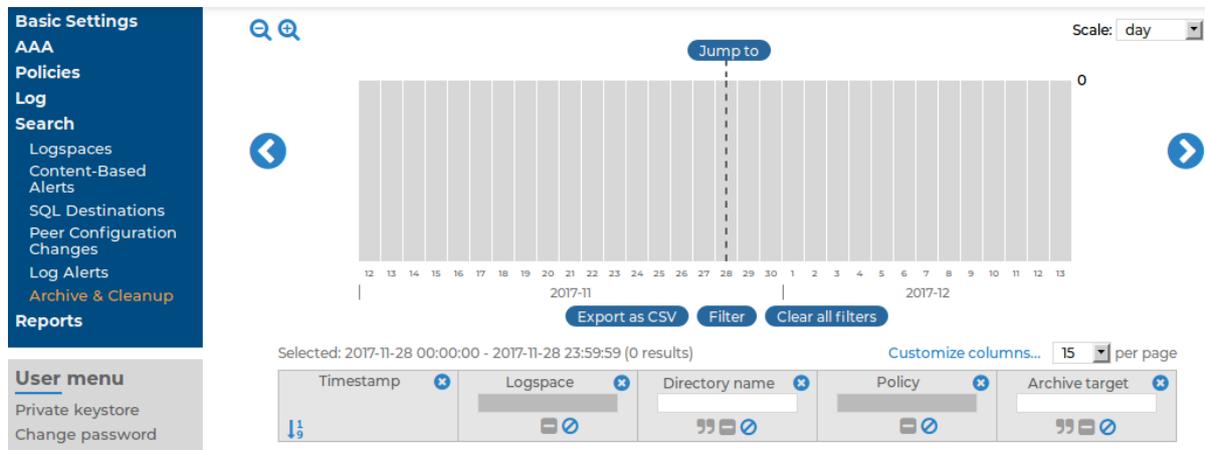
- **Timestamp:** The date of the alert in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Sender address:** The IP address of the client or relay that sent the message directly to SSB.
- **Hostname:** The hostname or IP address of the client that sent the message.
- **Program:** The application that generated the message.
- **Message:** The content of the message.
- **Rule ID:** The ID of the classification rule in the pattern database that matched the message. For details, see [Classifying messages with pattern databases](#) on page 376.
- **Rule description:** The description of the classification rule that matched the message. For details, see [Classifying messages with pattern databases](#) on page 376.

For details on how to navigate around the user interface and interact with features such as filtering and exporting results, and customizing what data is displayed, see [Using the internal search interfaces](#) on page 361.

Notifications on archiving and backups

Notifications and error messages of the archiving, cleanup and backup procedures are available at **Search > Archive & Cleanup**. The following information is available:

Figure 211: Search > Archive & Cleanup — Displaying archiving and backup notifications



- **Timestamp:** The date of the message in YEAR-MONTH-DAY HOUR:MINUTE:SECOND format.
- **Logspace:** The name of the archived or backed up logspace.
- **Directory name:** The name of the folder where the archives and backups are located. A new folder is created each day, using the current date as the folder name.
- **Policy:** The name of the archive or backup policy used.
- **Archive target:** The address of the remote server used in the policy.

- **Manual archiving:** Indicates if the archiving or backup process was started manually.

For details on how to navigate around the user interface and interact with features such as filtering and exporting results, and customizing what data is displayed, see [Using the internal search interfaces](#) on page 361.

Status history and statistics

The syslog-ng Store Box(SSB) appliance displays various statistics and the 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.

To display the statistics of a module as a table for the selected period, click on the graph.

Figure 212: Basic Settings > Dashboard — The dashboard



The following modules are displayed on the dashboard of SSB:

⚠ CAUTION:

Statistics about syslog-ng and logspace sizes are not backed up. As a result, following a data restore, the Basic Settings > Dashboard page will not show any syslog-ng and logspace statistics about the period before the backup.

- **syslog-ng:** syslog-ng statistics about the received, processed, and dropped messages. See also [Displaying custom syslog-ng statistics](#) on page 370.
- **Connected syslog peers:** A list of hosts that actively send messages to SSB. Note that these values are updated periodically based on the **Sampling interval** set on page **Log > Options > Dashboard Statistics**. For details, see [Displaying custom syslog-ng statistics](#) on page 370.

- **syslog-ng statistics:** The rate of incoming messages in messages/second. Note that the values displayed are average values calculated for the last 15 minutes.
- **Logspaces:** The size of the logspaces. Note that these values are updated only every ten minutes.
- **Memory:** The memory used by the system.
- **Disk:** Filesystem usage for the different partitions.
- **CPU:** CPU usage.
- **Network connections:** The number of network connections.
- **External interface:** Traffic on the external interface.
- **Management interface:** Traffic on the management interface.
- **Load average:** Average load of the system.
- **Processes:** The number of running processes.

For details about setting the statistics collection options, see [Statistics collection options](#) on page 370.

Displaying custom syslog-ng statistics

The following describes how to display statistics of a specific source, destination, or host.

To display statistics of a specific source, destination, or host

1. Navigate to **Basic Settings > Dashboard > syslog-ng statistics**.
2.
 - To display the statistics of a particular source, select source from the **Search in** field, and enter the name of the source into the **Search** field. Source names all start with the s character.
 - To display the statistics of a particular destination, select destination from the **Search in** field, and enter the name of the destination into the **Search** field. Destination names all start with the d character.
 - To display the statistics of a particular host, select src.host from the **Search in** field, and enter the hostname or IP address of the host into the **Search** field.
3. Select the time period to display from the **Select resolution** field.
4. Click **View graph**.

Statistics collection options

To control the quantity and quality of the statistics collected to the **Dashboard**, set the statistics collection options.

Navigate to **Log > Options > Dashboard statistics**.

Time-based statistics: The default setting is **Enabled**.

- **Cleanup if unchanged for:** Statistics unchanged (not present in syslog-ng statistics output anymore) for this number of days will be cleaned up from the system. Enter 0 here to keep them forever. To start the cleanup process immediately, click **Cleanup now**.
- **Enable statistics for:** The default setting is that all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Time-based statistics** and select the desired option.

NOTE: When disabling an option, the data will only be deleted after the first cleanup. Until then, the data already collected is still accessible on the dashboard.

Top/Least statistics: the default setting is **Enabled** and all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Top/Least statistics** and select the desired option.

Maximum number of statistics to process: Enter the number of statistics files to keep on the system. Enter 0 here to store unlimited number of statistics files. Statistics over this limit will be dropped, and SSB sends an error message containing the number of entries dropped and the first dropped entry. This setting needs to be increased only if you have more than 10000 hosts.

Sampling interval: Select the sampling interval for the statistics here. A more frequent sampling interval results in more precise graphs at the cost of heavier system load. The default setting is 5 minutes. The possible parameters are 5 minutes, 10 minutes, 30 minutes, 60 minutes, 2 hours, 4 hours, 8 hours, 1 day.

⚠ CAUTION:

Hazard of data loss When changing the Sampling interval, the already existing statistics are not converted to the new sampling rate, but are deleted.

To clear all statistics, click **Clear all statistics**. It is advised to clear statistics if you have changed the number of the statistics files to keep, or if you have disabled the time-based statistics collection.

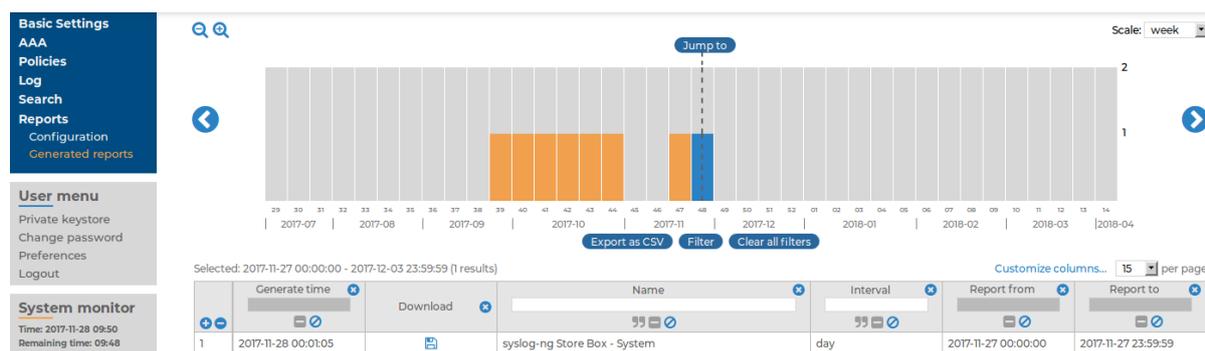
Reports

The syslog-ng Store Box(SSB) appliance periodically creates reports on the activity of the administrators, the system-health information of SSB, as well as the processed traffic. These reports are available in Portable Document (PDF) format by selecting **Reports > Generated reports** from the Main menu. 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.

To access the reports from the SSB web interface, the user must have the appropriate privileges.

NOTE: If the **Basic Settings > Management > Mail settings > Send reports to** address is not set, the report is sent to the SSB administrator's email address.

Figure 213: Reports > Generated reports — Browsing reports



Reports are generated as follows:

- **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.

TIP: Use the time bar to find reports that apply a particular period. If you select a period (for example, click on a bar), only those reports will be displayed that contain information about the selected period.

The following information is available about the reports:

- **Download:** A link to download the report.
- **Name:** The name of the report.
- **Interval:** The length of the reported period, for example, week, month, and so on.
- **Report from:** The start of the reported interval.
- **Report to:** The end of the reported interval.
- **Generate time:** The date when the report was created.

TIP: To create a report for the current day, select **Generate reports for today**. The report will contain data for the 00:00 - current time interval. If artificial ignorance (for details, see [Classifying messages with pattern databases](#) on page 376) is enabled, an artificial ignorance report is created as well.

For details on how to navigate around the user interface and interact with features such as filtering and exporting results, and customizing what data is displayed, see [Using the internal search interfaces](#) on page 361.

Contents of the default reports

The default report of the syslog-ng Store Box(SSB) appliance (called System) is available in Adobe Portable Document Format (PDF), and contains the following information for the given period:

- **Configuration changes:** Lists the number of SSB configuration changes per page and per user. The frequency of the configuration changes is also displayed on a chart.
- **Peer configuration:** Lists the number of times the configuration of a syslog-ng client was changed per client, as well as the version number of the syslog-ng application running on the client (if this information is available).
- **Alerts:** Various statistics about the alerts received from classifying messages using the pattern database (if pattern databases have been uploaded to SSB).
- **syslog-ng traffic statistics:** Displays the rate of incoming, forwarded, stored, and dropped messages in messages/second.
- **System health information:** Displays information about the filesystem and network use of SSB, as well as the average load.

Generating partial reports

The following describes how to generate a report manually for a period that has not been already covered in an automatic report.

To generate a report manually for a period that has not been already covered in an automatic report

1. Log in to the syslog-ng Store Box(SSB) web interface, and navigate to **Reports > Configuration**.
2. Select the report you want to generate.
3.
 - To create a report from the last daily report till now, click **Generate partial daily report**. For example, if you click this button at 11:30 AM, the report will include the period from 00:01 to 11:30.
 - To create a report from the last weekly report till now, click **Generate partial weekly report**. For example, if you click this button on Wednesday at 11:30 AM, the report will include the period from Monday 00:01 to Wednesday 11:30.
 - To create a report from the last monthly report till now, click **Generate partial monthly report**. For example, if you click this button at 11:30 AM, December 13, the report will include the period from December 1, 00:01 to December 13, 11:30.

The report will be automatically added in the list of reports (**Reports > Generated reports**), and also sent in an email to the regular recipients of the report.

4. Click **Commit**.

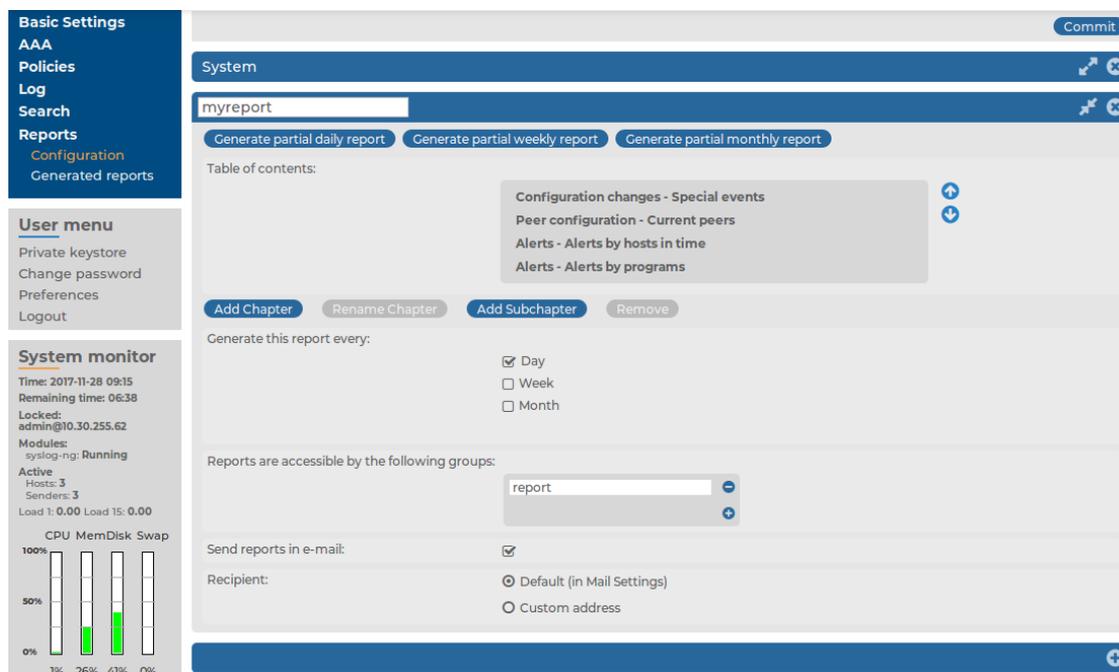
Configuring custom reports

The following describes how to configure syslog-ng Store Box(SSB) to create custom reports. Make sure that the user account has read & write/perform access to the **use static subchapters** privilege.

To configure SSB to create custom reports

1. Log in to the SSB web interface, and navigate to **Reports > Configuration**.

Figure 214: Reports > Configuration — Configuring custom reports



2. Click  and enter a name for the custom report.
3. Reports are organized into chapters and subchapters. To add a new chapter, go to **Table of contents**, click **Add Chapter**, enter a name for the chapter, then click **OK**. Repeat this step to create further chapters if needed.
4. Click **Add Subchapter** to add various reports and statistics to the chapter. The available reports will be displayed in a pop-up window. The reports created from custom statistics are listed at the end.
5. Use the arrows to change the order of the subchapters if needed.
6. To specify how often SSB should create the report, select the relevant **Generate this report every (Day, Week, Month)** option. Weekly reports are created on Mondays, while monthly reports on the first day of the month. You can select multiple options simultaneously.

If you want to generate the report only manually, leave this field empty.

7. By default, members of the search group can access the custom reports via the SSB web interface. To change this, enter the name of a different group into the **Reports are accessible by the following groups** field, or click  to grant access to other groups.

NOTE: Members of the listed groups will be able to access only these custom reports even if their groups do not have read access to the **Reporting > Reports**

page. However, only those reports will be listed, to which their group has access.

8. By default, SSB 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 SSB administrator's email address.

- To disable email sending, unselect the **Send reports in e-mail** option.
- To email the reports to a different address, select **Recipient > Custom address**, and enter the email address where the reports should be sent. Click  to list multiple email addresses if needed.

9. Click .

Classifying messages with pattern databases

Using the pattern database allows you to classify messages into various categories, receive alerts on certain messages, and to collect unknown messages using artificial ignorance.

Figure 215: Log > Pattern Database – Pattern database

The screenshot displays the 'Pattern Database' configuration interface. On the left is a navigation sidebar with the following items: Basic Settings, AAA, Policies, Log, Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options, Pattern Database (highlighted), Search, and Reports. The main content area is split into two panels. The top panel, titled 'Database', shows 'Publication date: 2009-11-26 (version 3)' and 'Rulesets/rules: 0/0'. It contains 'Empty' and 'Export' buttons, and an 'Upload database:' section with a 'Browse...' button, the text 'No file selected.', and an 'Upload' button. The bottom panel, titled 'Search', has a 'Ruleset name:' input field with a 'Show' button. Below it is a 'Match for log message:' section with 'Program:' and 'Message:' input fields, a 'Search' button, and a 'Create new ruleset' button.

Note that the classification of messages is always performed, but its results are used only if you specifically enable the relevant options on the **Log > Options** page.

Figure 216: Log > Options — Enabling artificial ignorance and pattern-matching alerts

The screenshot displays the 'Log > Options' configuration page. On the left is a navigation menu with categories: Basic Settings, AAA, Policies, Log (with sub-items: Sources, Logspaces, Filtered Logspaces, Remote Logspaces, Multiple Logspaces, Destinations, Paths, Parsers, Options, Pattern Database), Search, and Reports. Below this is a 'User menu' with options: Private keystore, Change password, Preferences, and Logout. The main content area has a 'Commit' button at the top right. Below it are several configuration sections: 'Options', 'TLS settings', 'SNMP source' (checked), 'Alerting' (checked) with 'Rate limit' set to 10 alerts/minute and an 'Alert email' field; 'Artificial ignorance' (unchecked) with 'Database retention' set to 7 days and a 'Report email' field; 'Name resolving'; 'Dashboard statistics' (checked); and 'Message rate alerting statistics'.

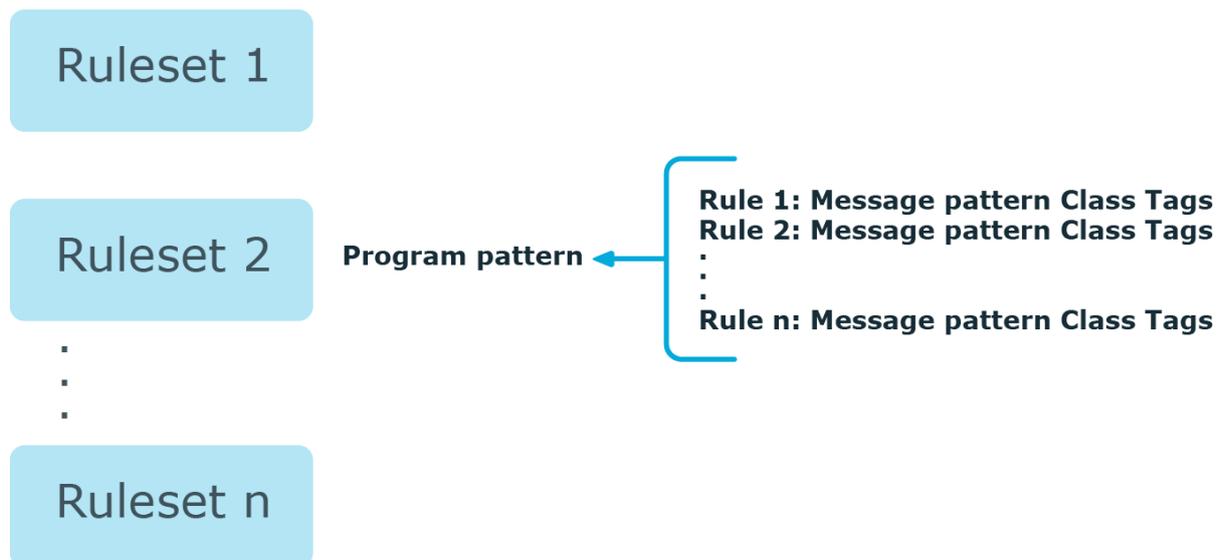
- To receive alerts on messages classified as Violation, navigate to **Log > Options** and enable the **Alerts** option.
- To receive reports on messages not included in the pattern database, navigate to **Log > Options** and enable the **Artificial ignorance** option.

The structure of the pattern database

The pattern database is organized as follows:

Figure 217: The structure of the pattern database

Pattern database



- The pattern database consists of rulesets. A ruleset consists of a Program Pattern and a set of rules: the rules of a ruleset are applied to log messages if the name of the application that sent the message matches the Program Pattern of the ruleset. The name of the application (the content of the `${PROGRAM}` macro) is compared to the Program Patterns of the available rulesets, and then the rules of the matching rulesets are applied to the message.
- The Program Pattern can be a string that specifies the name of the application or the beginning of its name (for example, to match for sendmail, the program pattern can be sendmail, or just send), and the Program Pattern can contain pattern parsers. Note that pattern parsers are completely independent from the syslog-ng parsers used to segment messages. Additionally, every rule has a unique identifier: if a message matches a rule, the identifier of the rule is stored together with the message.
- Rules consist of a message pattern and a class. The Message Pattern is similar to the Program Pattern, but is applied to the message part of the log message (the content of the `${MESSAGE}` macro). If a message pattern matches the message, the class of the rule is assigned to the message (for example, Security, Violation, and so on).
- Rules can also contain additional information about the matching messages, such as the description of the rule, an URL, name-value pairs, or free-form tags. This information is displayed by the syslog-ng Store Box(SSB) appliance in the e-mail alerts (if alerts are requested for the rule), and are also displayed on the search interface.
- Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers.

NOTE: If the `${PROGRAM}` part of a message is empty, rules with an empty Program Pattern are used to classify the message.

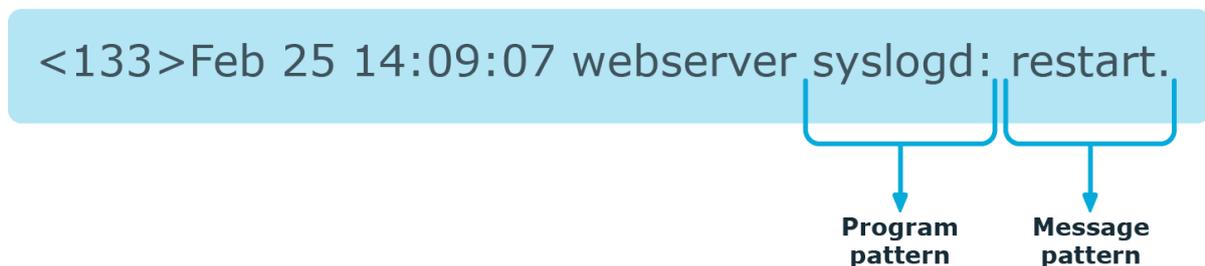
If the same Program Pattern is used in multiple rulesets, the rules of these rulesets are merged, and every rule is used to classify the message. Note that message patterns must be unique within the merged rulesets, but the currently only one ruleset is checked for uniqueness.

How pattern matching works

This section describes how patterns work. This information applies to program patterns and message patterns alike, even though message patterns are used to illustrate the procedure.

Figure 218: Applying patterns

A sample log message:



Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers. Pattern parsers attempt to parse a sequence of characters according to certain rules.

NOTE: Wildcards and regular expressions cannot be used in patterns. The @ character must be escaped, that is, to match for this character, you have to write @@ in your pattern. This is required because pattern parsers of syslog-ng are enclosed between @ characters.

When a new message arrives, syslog-ng attempts to classify it using the pattern database. The available patterns are organized alphabetically into a tree, and syslog-ng inspects the message character-by-character, starting from the beginning. This approach ensures that only a small subset of the rules must be evaluated at any given step, resulting in high processing speed. Note that the speed of classifying messages is practically independent from the total number of rules.

For example, if the message begins with the Apple string, only patterns beginning with the character A are considered. In the next step, syslog-ng selects the patterns that start with Ap, and so on, until there is no more specific pattern left.

Note that literal matches take precedence over pattern parser matches: if at a step there is a pattern that matches the next character with a literal, and another pattern that would match it with a parser, the pattern with the literal match is selected. Using the previous example, if at the third step there is the literal pattern Apport and a pattern parser Ap@STRING@, the Apport pattern is matched. If the literal does not match the incoming string (for example, Apple), syslog-ng attempts to match the pattern with the parser. However, if there are two or more parsers on the same level, only the first one will be applied, even if it does not perfectly match the message.

If there are two parsers at the same level (for example, Ap@STRING@ and Ap@QSTRING@), it is random which pattern is applied (technically, the one that is loaded first). However, if the selected parser cannot parse at least one character of the message, the other parser is used. But having two different parsers at the same level is extremely rare, so the impact of this limitation is much less than it appears.

Searching for rulesets

To display the rules of a ruleset, enter the name of the ruleset into the **Search > Ruleset name** field, and click **Show**. If you do not know the name of the ruleset, type the beginning letter(s) of the name, and the names of the matching rulesets will be displayed. If you are looking for a specific rule, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

| NOTE: Rulesets containing large number of rules may not display correctly.

Figure 219: Log > Pattern Database > Search > Ruleset name — Searching rules

Search

Ruleset name: **Show**

Match for log message:

Program:

Message:

Search

Create new ruleset

hddtemp

Name:

URL:

Program pattern:

Description:

Rules:

Pattern	Class
<input type="text" value="/dev@QSTRING::/@QSTRING:: :@ @NUMBER::@"/>	System
<input type="text" value="/dev/sg@NUMBER::/@QSTRING:: :@ drive is sleeping"/>	System
<input type="text" value="/dev@QSTRING::/@QSTRING:: :@ drive is sleeping"/>	System
<input type="text" value="/dev/sg@NUMBER::/@QSTRING:: :@ @NUMBER::@"/>	System

Creating new rulesets and rules

The following describes how to create a new ruleset and new rules.

To create a new ruleset and new rules

1. Select **Log > Pattern Database > Create new ruleset**.

TIP: If you search for a ruleset that does not exist, syslog-ng Store Box(SSB) offers you to create a new ruleset with the name you were searching for.

2. Enter a name for the ruleset into the **Name** field.

Figure 220: Log > Pattern Database > Create new ruleset — Creating pattern database rulesets

New Ruleset

Name:

URL:

Program pattern:

Description:

Rules:

Pattern **Class**

System

Description:

URL:

Tags:

Rule ID:

Values:

Name	Value
<input type="text" value="username"/>	<input type="text" value="\$username"/>
<input type="text" value="eventtype"/>	<input type="text" value="login"/>

3. Enter the name of the application or a pattern that matches the applications into the **Program pattern** field. For details, see [Using pattern parsers](#) on page 383.
4. Optionally, add a description to the ruleset.
5. Add rules to the class.

- a. Click  in the **Rules** section.
 - b. Enter the beginning of the log message or a pattern that matches the log message into the **Pattern** field. For details, see [Using pattern parsers](#) on page 383. Note that only messages sent by applications matching the **Program pattern** will be affected by this pattern.
 - c. Select the type of the message from the **Class** field. This class will be assigned to messages matching the pattern of this rule. The following classes are available: Violation, Security, and System.
If alerting is enabled at **Log > Options > Alerting**, SSB automatically sends an alert if a message is classified as Violation.
 - d. Optionally, you can add a description, custom tags, and name-value pairs to the rule. Note that the values of name-value pairs can contain macros in the `#{macroname}` format. For details on pattern databases and macros, see [The syslog-ng Premium Edition Administration Guide](#).
6. Repeat the previous step to add more rules.
 7. Click .

Exporting databases and rulesets

To export the entire pattern database, navigate to **Log > Pattern Database** and select **Export**.

To export a ruleset, enter the name of the ruleset into the **Search > Ruleset name** field, click **Show**, and select **Export ruleset**. If you do not know the name of the ruleset, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

Importing pattern databases

You can upload official databases distributed by One Identity or pattern databases that you have exported from syslog-ng Store Box(SSB).

To import a ruleset, navigate to Log > Pattern Database, select Browse, then locate the database file to upload, and click Upload.

NOTE: Imported rules are effective immediately after the upload is finished.

If you have modified a rule that was originally part of an official database, then the update will not modify this rule.

Using pattern parsers

Pattern parsers attempt to parse a part of the message using rules specific to the type of the parser. Parsers are enclosed between @ characters. The syntax of parsers is the following:

- a beginning @ character,
- the type of the parser written in capitals,
- optionally a name,
- parameters of the parser, if any, and
- a closing @ character.

Example: Pattern parser syntax

A simple parser:

```
@STRING@
```

A named parser:

```
@STRING:myparser_name@
```

A named parser with a parameter:

```
@STRING:myparser_name:*@
```

A parser with a parameter, but without a name:

```
@STRING::*@
```

The following parsers are available:

- **@ANYSTRING@**: Parses everything to the end of the message. You can use it to collect everything that is not parsed specifically to a single macro. In that sense its behavior is similar to the greedy() option of the CSV parser.
- **@DOUBLE@**: An obsolete alias of the **@FLOAT@** parser.
- **@ESTRING@**: This parser has a required parameter that acts as the stopcharacter: the parser parses everything until it find the stopcharacter. For example to stop by the next " (double quote) character, use **@ESTRING::"@**. As of syslog-ng 3.1, it is possible to specify a stopstring instead of a single character, for example **@ESTRING::stop_here.@**.
- **@FLOAT@**: A floating-point number that may contain a dot (.) character. (Up to syslog-ng 3.1, the name of this parser was **@DOUBLE@**.)
- **@IPv4@**: Parses an IPv4 IP address (numbers separated with a maximum of 3 dots).
- **@IPv6@**: Parses any valid IPv6 IP address.

- @IPvANY@: Parses any IP address.
- @NUMBER@: A sequence of decimal (0-9) numbers (for example 1, 0687, and so on). Note that if the number starts with the 0x characters, it is parsed as a hexadecimal number, but only if at least one valid character follows 0x.
- @QSTRING@: Parse a string between the quote characters specified as parameter. Note that the quote character can be different at the beginning and the end of the quote, for example: @QSTRING: : "@ parses everything between two quotation marks ("), while @QSTRING: <>@ parses from an opening bracket to the closing bracket.
- @STRING@: A sequence of alphanumeric characters (0-9, A-z), not including any whitespace. Optionally, other accepted characters can be listed as parameters (for example to parse a complete sentence, add the whitespace as parameter, like: @STRING: : @). Note that the @ character cannot be a parameter, nor can line-breaks or tabs.

Patterns and literals can be mixed together. For example, to parse a message that begins with the Host: string followed by an IP address (for example Host: 192.168.1.1), the following pattern can be used: Host:@IPv4@.

NOTE: Note that using parsers is a CPU-intensive operation. Use the ESTRING and QSTRING parsers whenever possible, as these can be processed much faster than the other parsers.

Example: Using the STRING and ESTRING parsers

For example, if the message is user=joe96 group=somegroup, @STRING:mytext:@ parses only to the first non-alphanumeric character (=), parsing only user. @STRING:mytext:=@ parses the equation mark as well, and proceeds to the next non-alphanumeric character (the whitespace), resulting in user=joe96 being parsed. @STRING:mytext:= @ will parse the whitespace as well, and proceed to the next non-alphanumeric non-equation mark non-whitespace character, resulting in user=joe96 group=somegroup.

Of course, usually it is better to parse the different values separately, like this: "user=@STRING:user@ group=@STRING:group@".

If the username or the group may contain non-alphanumeric characters, you can either include these in the second parameter of the parser (as shown at the beginning of this example), or use an ESTRING parser to parse the message till the next whitespace: "user=@ESTRING:user: @group=@ESTRING:group: @".

Example: Patterns for multiline messages

Patterns can be created for multiline log messages. For example, the following pattern will find the multiline message where a line ends with `first` and the next line starts with `second`:

```
first
second
```

Using parser results in filters and templates

The results of message classification and parsing can be used in custom filters and file and database templates as well. There are two built-in macros in syslog-ng Store Box (SSB) that allow you to use the results of the classification: the `.classifier.class` macro contains the class assigned to the message (for example `violation`, `security`, or `unknown`), while the `.classifier.rule_id` macro contains the identifier of the message pattern that matched the message.

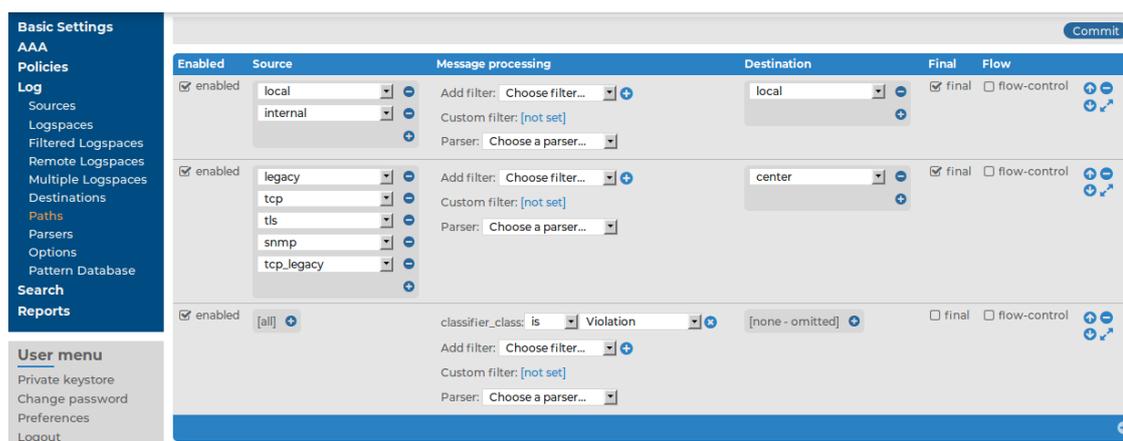
NOTE: ID of the message pattern is automatically inserted into the template if the messages are forwarded to an SQL database.

The following describes how to use these macros as filters in a log path.

To use macros as filters in a log path

1. Navigate to **Log > Paths** and select the log path to use.
2. To filter on a specific message class, select **Add filter > classifier_class**, select , then select the class to match (for example `violation`) from the **classifier_class** field.

Figure 221: Log > Paths — Filtering messages based on the classification



3. To filter on messages matching a specific classification rule, **Add filter > classifier_rule_id**, select **+**, then enter the unique identifier of the rule (for example e1e9c0d8-13bb-11de-8293-000c2922ed0a) into the **classifier_rule_id** field.

NOTE: To filter messages based on other classification data like tags, you have to use Custom filters. For details, see [Filtering messages](#) on page 309.

4. Click **Commit**.

Using the values of pattern parsers in filters and templates

Similarly to [Using parser results in filters and templates](#) on page 386, the results of pattern parsers can be used as well. To accomplish this, you have to add a name to the parser, and then you can use this name as a macro that refers to the parsed value of the message.

For example, you want to parse messages of an application that look like "Transaction: <type>.", where <type> is a string that has different values (for example refused, accepted, incomplete, and so on). To parse these messages, you can use the following pattern:

```
'Transaction: @ESTRING:..@'
```

Here the @ESTRING@ parser parses the message until the next full stop character. To use the results in a filter or a filename template, include a name in the parser of the pattern, for example:

```
'Transaction:
    @ESTRING:TRANSACTIONTYPE:..@'
```

After that, add a custom template to the logpath that uses this template. For example, to select every accepted transaction, use the following custom filter in the log path:

```
match("accepted" value("TRANSACTIONTYPE"));
```

NOTE: The above macros can be used in database columns and filename templates as well, if you create custom templates for the destination or logspace.

The SSB RPC API

The syslog-ng Store Box(SSB) appliance can be accessed using a Remote-Procedure Call Application Programming Interface (RPC API).

The SSB RPC API allows you to access and query SSB logspaces from remote applications. You can access the API using a RESTful protocol over HTTPS, meaning that you can use any programming language that has access to a RESTful HTTPS client to integrate SSB to your environment. Sample shell code snippets are provided in the API documentation.

Accessing SSB with the RPC API offers several advantages:

- Integration into custom applications and environments
- Flexible, dynamic search queries

SSB prevents brute force attacks when logging in. If you repeatedly try logging in to SSB using incorrect login details within a short period of time (10 times within 60 seconds), the source IP gets blocked for 5 minutes.

Requirements for using the RPC API

To access syslog-ng Store Box(SSB) using the RPC API, the following requirements must be met:

- The appliance can be accessed using a RESTful protocol over authenticated HTTPS connections.
- The user account used to access SSB via RPC must have **Search** privilege (which provides access to all logspaces), or must be a member of the groups listed in the **Access Control** option of the particular logspace. For details on managing user privileges, see [Modifying group privileges](#).

RPC client requirements

The client application used to access syslog-ng Store Box(SSB) must meet the following criteria:

- Support RESTful web APIs over HTTPS
- Properly handle complex object types
- Include a JSON decoder for interpreting the results of search operations

Documentation of the RPC API

The documentation of the syslog-ng Store Box(SSB) RPC API is available online from the following URL: <https://<ip-address-of-SSB>/api/4/documentation>. This documentation contains the detailed description of public calls, with examples. For a quickstart guide, see [RPC API Quickstart Guide](#).

Monitoring SSB

This section describes monitoring syslog-ng Store Box(SSB).

If you intend to monitor SSB, the general rules of monitoring an Ubuntu system apply.

Depending on your monitoring tool, check out the available SNMP templates related to your system, and use those to monitor the general parts of SSB.

⚠ CAUTION:

Do not install any monitoring agent on the SSB appliance. Any alteration to the system is unsupported.

For details on receiving e-mail alerts or SNMP traps on health- and system-related issues, see [Configuring system monitoring on SSB](#) on page 87.

For details on SNMP communities, see [View node ID and community](#) on page 86.

Monitoring SSB's disk

SNMP object:	HOST-RESOURCES-MIB::hrStorageTable
---------------------	------------------------------------

Community (v2c) /	Data and system
--------------------------	-----------------

Context (v3)	
---------------------	--

This section describes monitoring syslog-ng Store Box(SSB)'s disk usage. Disk usage is measured per mountpoint (that is, partition).

Data partition

Mountpoint	/mnt/firmware
------------	---------------

Partition	data partition
-----------	----------------

Community	<community-of-extended-node> (Basic Settings > Monitoring > SNMP agent settings > Community)
-----------	--

Context	"" (empty string)
---------	-------------------

Free disk space that can be acquired by your main production system (for example, logspaces).

To make sure that you have the free disk space you are comfortable with, monitor the free disk space. This partition can fill up with for example the following:

- received logs
- generated reports
- collected statistics
- core dumps

If you cannot keep your free disk space in your comfortable interval (with scheduled cleanups) you probably need to purchase more SSB appliances. For assistance, [contact our Support Team](#).

NOTE: If you have configured **Basic Settings > Management > Disk space fill up prevention**, be aware that SSB will stop receiving log after reaching the configured threshold. By default, clients are disconnected when disks are 90 percent full. For details, see [Preventing disk space fill up](#) on page 90.

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	The file system's available space is continuously decreasing, depending on logspaces, reports, statistics, and core files.
Related issues and issue indicators	The monitored values may indicate that the used storage space has reached the level configured under Basic Settings > Management > Disk space fill up prevention . In this case, archiving starts (if archiving is configured), and incoming logs are rejected. Note that if the disk fill up prevention level is reached too often and too soon, it indicates a serious issue.

⚠ CAUTION:

Hazard of data loss If the used storage space of the file system under `/mnt/firmware` exceeds the allowed rate configured under **Basic Settings > Management > Disk space fill up prevention**, SSB will stop receiving log messages. To avoid data loss, archive your logs or store them in some other way (for example, forward your log messages to different logstores). Alternatively, reconsider your configuration settings, upgrading the capacity of your SSB appliance, or purchasing more SSB appliances.

Solution:

- Consider upgrading the capacity of your SSB appliance or purchasing more SSB appliances (for more information, see "[Hardware specifications](#)" in the [Installation Guide](#) or [contact our Sales Team](#)).
- Reconsider your configuration settings.
- Archive your logs from your logspaces to clean up logstores on your SSB appliance.

System partition

Mountpoint	/initrd/mnt
Partition	system partition (if SSB is in HA mode, the system firmware partition of both nodes)
Community	<id-of-the-node>
Context	<id-of-the-node>

The space on this mountpoint is required only by the system. Generally, this is independent from how you use SSB. The only important thing here to have some free space on the mountpoint.

Make sure that you have some free space on this mountpoint. As a recommended threshold, set a trigger to 80% in your monitoring system. If there is only about 20% free space left on this mountpoint, [contact our Support Team](#).

NOTE: Monitoring the size of specific logspaces is not possible this way. If you are interested in the size of a specific logspace, you can configure a size limit alert for that logspace on the SSB web interface. For details on configuring a disk size alert for a specific logspace, see [Creating logstores](#) on page 217.

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	The available storage space is not supposed to decrease significantly.
Related issues and issue indicators	The monitored values may indicate that the system is out of storage space, or that the available storage space is running very low.

CAUTION:

Hazard of data loss If the used storage space of the file system under /mnt/firmware exceeds the allowed rate configured under **Basic Settings > Management > Disk space fill up prevention**, SSB will stop receiving log messages. To avoid data loss, archive your logs or store them in some other way (for example, forward your log messages to different logstores). Alternatively, reconsider your configuration settings, upgrading the capacity of your SSB appliance, or purchasing more SSB appliances.

Solution:

If you encounter an issue related to the file system under /initrd/mnt, [contact our Support Team](#).

Monitoring SSB's memory

Swap usage

SNMP object:	<ul style="list-style-type: none">• UCD-SNMP-MIB::memTotalSwap• UCD-SNMP-MIB::memAvailSwap
---------------------	---

Community (v2c) / Context (v3)	Data and system
---------------------------------------	-----------------

The syslog-ng Store Box(SSB) appliance is using swap (except on Azure) as part of its normal operation.

You can receive swap usage alerts by configuring an alert on **Basic Settings > Alerting**. For details, see [Configuring system monitoring on SSB](#) on page 87.

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing.
Related issues and issue indicators	If the computed returned value is close to 0, memory usage is too high.

Solution:

- Decrease load.
- Purchase a new SSB appliance.

Memory usage

SNMP object:	<ul style="list-style-type: none">• UCD-SNMP-MIB::memTotalReal• UCD-SNMP-MIB::memAvailReal
---------------------	---

Community (v2c) / Context (v3)	Data and system
---------------------------------------	-----------------

If SSB's memory and swap usage are both above 90%, fine-tune your configuration or purchase more SSB appliances to balance the load. For assistance, [contact our Support Team](#).

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing.
Related issues and issue indicators	If the computed returned value is close to 0, memory usage is too high.

Solution:

- Decrease load.
- Purchase a new SSB appliance.

Monitoring SSB's CPU

This section describes monitoring syslog-ng Store Box(SSB)'s CPU.

NOTE: This document uses the concept of 'logical CPU' used by the Linux kernel. In this section, 'logical CPU' will be abbreviated as 'CPU'. To determine the number of CPUs in your SSB machine, enter the `lscpu` command in your console, or send an SNMP request.

Monitoring CPU load

SNMP object:	HOST-RESOURCES-MIB::hrProcessorLoad
Community (v2c) / Context (v3)	Data and system

CPU load shows the availability of a CPU. Its value is a rational number ranging from 0.0 or 0% to 1.0 or 100%. If the CPU load is 0.0 or 0%, then the measured CPU is idle. If the value is 1.0 or 100%, then the CPU is fully loaded.

If the load of one CPU is above 90% and other threads are significantly less loaded for a longer time period, fine-tune your configuration or purchase more syslog-ng Store Box (SSB) appliances. For assistance, [contact our Support Team](#).

Monitoring CPU load averages

SNMP object:	UCD-SNMP-MIB::laLoad
---------------------	----------------------

Community (v2c) / Data and system
Context (v3)

CPU load averages (or system load averages) is the average load of syslog-ng Store Box (SSB)'s CPUs and the size of the task queue, during the past 1, 5, and 15 minutes, respectively.

If the load is constantly equal to or higher than the number of CPUs in your appliance, fine-tune your configuration or purchase more SSB appliances. For assistance, [contact our Support Team](#).

If you query `UCD-SNMP-MIB::1aTable`, the table of returned values will contain the returned values you would get when querying `UCD-SNMP-MIB::1aLoad`, but in a more structured, easy-to-read format.

For which systems and configurations is it applicable? Applicable for all configurations and systems.

Value change frequency Its value is continuously changing.

Related issues and issue indicators Load too high (see [Monitoring SSB's CPU](#)).

Solution:

- Decrease load.
- Purchase a new SSB appliance.

Monitoring CPU usage

User CPU time:

SNMP object: UCD-SNMP-MIB::ssCpuUser

Community (v2c) / Data and system
Context (v3)

If processor is not in idle (for example, there is live log traffic or report generation), it is quite normal that the majority of the CPU time is spent on running user space processes.

For which systems Applicable for all configurations and systems.

and configurations is it applicable?

Value change frequency Its value is continuously changing.

Related issues and issue indicators If the returned value is too high, it indicates a highly loaded CPU on user space.*Solution:*

- Decrease load.
- Purchase a new syslog-ng Store Box(SSB) appliance
- Reconsider your configuration settings.

System CPU time:

SNMP object: UCD-SNMP-MIB::ssCpuSystem

Community (v2c) / Context (v3) Data and system

The amount of time spent in the kernel should be as low as possible. Ideally, around 0.5% of the time given to the different processes is spent in the kernel. This number can peak much higher, especially when there are a lot of I/O activities.

For which systems and configurations is it applicable? Applicable for all configurations and systems.

Value change frequency Its value is continuously changing.

Related issues and issue indicators If the returned value is too high, it indicates high kernel-intensive operations running on the CPU.*Solution:*

- Decrease load.
- Purchase a new SSB appliance
- Reconsider your configuration settings.

Idle CPU time:

SNMP object:	UCD-SNMP-MIB::ssCpuIdle
Community (v2c) / Context (v3)	Data and system

The total of the user CPU time percentage and the idle CPU percentage should be close to 100%. If the CPU spends a lot more time in other states, it is worth investigating the root cause, because it can indicate issues.

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing.

Monitoring the details of your CPU usage:

To fine-tune monitoring your CPU usage, you can use the following values. It is also possible to generate a chart from these values.

Raw CPU user time	UCD-SNMP-MIB::ssCpuRawUser
Raw CPU nice time	UCD-SNMP-MIB::ssCpuRawNice
Raw CPU system time	UCD-SNMP-MIB::ssCpuRawSystem
Raw CPU idle time	UCD-SNMP-MIB::ssCpuRawIdle
Raw CPU wait time	UCD-SNMP-MIB::ssCpuRawWait
Raw CPU kernel time	UCD-SNMP-MIB::ssCpuRawKernel
Raw CPU interrupt time	UCD-SNMP-MIB::ssCpuRawInterrupt
Raw CPU Soft IRQ time	UCD-SNMP-MIB::ssCpuRawSoftIRQ
Raw CPU steal time	UCD-SNMP-MIB::ssCpuRawSteal
Raw CPU guest time	UCD-SNMP-MIB::ssCpuRawGuest
Raw CPU guest nice time	UCD-SNMP-MIB::ssCpuRawGuestNice
Community (v2c) / Context (v3)	Data and system for all of the above

Monitoring SSB's I/O

Disk I/O per partition

SNMP object:	UCD-DISKIO-MIB::diskIOTable
Community (v2c) / Context (v3)	Data and system

- sda

If the 15-minute load (for details, see [Monitoring CPU load averages](#) on page 395) is getting close to 90%, your system does not have enough resources and you probably need to purchase more syslog-ng Store Box(SSB) appliances. For assistance, [contact our Support Team](#).

- sdb

| **NOTE:** This is only available on SSB T1 appliances.

If the 15-minute load (for details, see [Monitoring CPU load averages](#) on page 395) is getting close to 90%, your system does not have enough resources and you probably need to purchase more SSB appliances. For assistance, [contact our Support Team](#).

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value changes quite often, depending on the I/O load and its type.
Related issues and issue indicators	When I/O load is too high, the system slows down.

Solution:

- Reconsider your configuration settings.
- Purchase a new SSB appliance.
- For technical assistance, [contact our Support Team](#).

Interfaces I/O by interface name

SNMP object:	RFC1213-MIB::ifTable
Community (v2c) / Context (v3)	Data and system

The following interfaces can be monitored (the type of traffic that can affect the load):

- eth0 - external (network, redundant HA, next hop monitoring)
- eth1 - management (network, redundant HA, next hop monitoring)
- eth2 - internal (redundant, next hop monitoring)
- eth3 - HA

If the load on an interface seems to be too high, check whether you have configured SSB in a way that affects that node. For example, if you do not use a management interface, the load on the external interface can be higher. Or, configuring next hop monitoring can also increase the load on an interface.

RFC1213-MIB:ifTable

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing, depending on incoming logs and DRBD sync.
Related issues and issue indicators	I/O load may become too high on network interfaces, which may result in log loss, slow sync, and HA in degraded mode.

Solution:

- Reconsider your configuration settings.
- Purchase a new SSB appliance.
- For technical assistance, [contact our Support Team](#).

RFC1213-MIB:ifTable - ETH 0, ETH3

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing, depending on the number of incoming logs.
Related issues and issue indicators	If the I/O load is too high, the network will not handle it, which may result in log loss.

CAUTION:

Hazard of data loss If the I/O load becomes too high for the network to handle, it may cause log loss. To avoid log loss, reconsider your configuration settings, Alternatively, reconsider your configuration settings, upgrading the capacity of your SSB appliance, purchasing more SSB appliances, or [contact our Support Team](#).

Solution:

- Reconsider your configuration settings.
- Consider upgrading the capacity of your SSB appliance or purchasing more SSB appliances (for more information, see ["Hardware specifications" in the Installation Guide](#) or [contact our Sales Team](#)).
- For technical assistance, [contact our Support Team](#).

RFC1213-MIB:ifTable - ETH3

For which systems and configurations is it applicable?	Only applicable for HA clusters.
Value change frequency	Its value is continuously changing, depending on the number of incoming logs.
Related issues and issue indicators	Network traffic load too high for the NIC to handle. (It rarely ever happens.)

Solution:

- Reconsider your configuration settings.
- Consider upgrading the capacity of your SSB appliance or purchasing more SSB appliances (for more information, see ["Hardware specifications" in the Installation Guide](#) or [contact our Sales Team](#)).
- For technical assistance, [contact our Support Team](#).

Monitoring SSB statistics

SSB's version number

SNMP object:	SSB-SNMP-MIB::ssbFirmwareVersion
Type:	String
Community (v2c) / Context (v3)	Data
Short description:	Current version of the syslog-ng Store Box(SSB).

Description: The current version number of SSB. This always changes after a successful upgrade.

Number of session files on SSB

SNMP object:	SSB-SNMP-MIB::ssbHTTPSessions
Type:	Integer32
Community (v2c) / Context (v3)	Data
Short description:	Number of recently active HTTP-based connections to SSB.

Description: The number of session files on SSB. These are generated as a result of the following events:

- Accessing the web user interface of SSB.
- Accessing a remote logspace.
- Performing an RPC API call.

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously changing, depending on the number of active connections.
Related issues and issue indicators	If the returned value changes too often within a short period of time, it can indicate a brute force attack.

Solution:

- Cooperate with your network administrator to fend off the external brute force attack.

Number of core files on SSB

SNMP object:	SSB-SNMP-MIB::ssbCoreFiles
Type:	Integer (number of)
Community (v2c) / Context (v3)	Data
Short description:	The number of core files in SSB's core firmware.

Description: If the value of this parameter is larger than 0, [contact our Support Team](#).

For which systems and configurations is it applicable?	Applicable for all configurations and systems, but unless a core file is generated, its returned value is 0.
---	--

Value change frequency	Its value does not change often, only when a core file is generated.
Related issues and issue indicators	Even a single core file indicates an issue. When more than one of them appear, it indicates a more serious issue.

Solution:

- Check the state of syslog-ng/indexer.
- Restart your syslog-ng application.
- For technical assistance, [contact our Support Team](#).

Available free space on SSB

SNMP object:	SSB-SNMP-MIB::ssbUnusedLogStorageCapacity
Type:	Integer (% percent)
Community (v2c) / Context (v3)	Data
Short description:	Ratio of free space on SSB compared to the Disk space fill up prevention limit.

Description: The available free space on SSB.

⚠ CAUTION:

Hazard of data loss If the value of this parameter is constantly close to 0%, fine-tune your configuration or purchase more SSB appliances. For assistance, [contact our Support Team](#).

If the value of this parameter reaches 0%, SSB will stop receiving logs.

If you have an Archive Policy configured, archiving will start after the value of this parameter reaches 0%. Therefore, SSB might start receiving logs again after some time has passed.

Make sure that you always have enough free space.

The definition of "enough" varies based on your specific configuration settings, for example:

- The disk size of your SSB appliance.
- The size, number and frequency of your incoming logs.
- Your **Policies > Backup & Archive/Cleanup settings** configuration.
- Your **Basic Settings > Management > Disk space fill up prevention limit** configuration. For details, see [Preventing disk space fill up](#) on page 90
- and so on

Example: Available free space on SSB

To calculate the available free space on SSB, the following formula is used:

[Disk capacity of the core partition] - [The free space above the **Basic Settings > Management > Disk space fill up prevention limit**] - [The space that is already in use].

For example:

- *Disk capacity of the core partition:* This is always 100%
- *The free space above the **Basic Settings > Management > Disk space fill up prevention limit**:* If SSB is configured to **Disconnect clients when disks are 90 percent used**, this value is $100\% - 90\% = 10\%$
- *The space that is already in use:* 35%

Available free space on SSB = $100\% - 10\% - 35\% = 55\%$

For which systems and configurations is it applicable?	Applicable for all configurations and systems.
Value change frequency	Its value is continuously decreasing, depending on available log storage capacity.
Related issues and issue indicators	As the returned value approaches 0, the available log storage capacity is continuously decreasing.

Solution:

- Archive your logs or store them in some other way (for example, forward your log messages to different logstores).
- Consider upgrading the capacity of your SSB appliance or purchasing more SSB appliances (for more information, see ["Hardware specifications" in the Installation Guide](#) or [contact our Sales Team](#)).

Monitoring the HA cluster

The status of the HA cluster

SNMP object: SSB-SNMP-MIB::ssbHAClusterStatus

Type:	String
Community (v2c) / Context (v3)	Data
Short description:	Status of the HA cluster.

Description: The status of the syslog-ng Store Box(SSB) cluster. For details, see [Status](#) on page [429](#).

For which systems and configurations is it applicable?	Only applicable for HA clusters.
Value change frequency	When the HA cluster functions properly, this SNMP object should be in ha status in the majority of the cases. The rest of the status returned values (for example, degraded) may also occur occasionally, but the ha status should be dominant as a rule.
Related issues and issue indicators	If while in a HA cluster, the status returned value is not ha or sync, the HA cluster is in degraded mode.

Solution:

- Check your HA network.
- Reboot the Secondary node.
- Reboot the HA cluster.
- For technical assistance, [contact our Support Team](#).

The status of the Redundant Heartbeat interface

SNMP object:	SSB-SNMP-MIB: :ssbHARedundantHeartbeatStatus
Type:	String
Community (v2c) / Context (v3)	Data
Short description:	Status of the Redundant Heartbeat interface.

Description: The status of the Redundant Heartbeat interface. For details, see [Redundant Heartbeat status](#) on page [431](#).

For which systems and configurations is it applicable?	Only applicable for HA clusters, but it only has a returned value if Redundant HA is configured.
---	--

Value change frequency	When the cluster functions properly, it should be in ok returned value status in the majority of the cases. The rest of the status returned values (for example, degraded) may also occur occasionally, but the ok status should be dominant as a rule.
Related issues and issue indicators	Sometimes this SNMP objects has an ok status, but the HAclusterStatus is not ok. The HA cluster will function properly in this case, too.

The synchronization progress of HA nodes

SNMP object:	SSB-SNMP-MIB::ssbHASynchronizationProgress
Type:	Integer32 (0..100 %)
Community (v2c) / Context (v3)	Data
Short description:	HA cluster synchronization progress (in percent). 100%, if the cluster is fully synchronized.

Description: This value can be important in the following cases:

- When enabling HA mode the first time, after navigating to **Basic Settings > High Availability** and clicking **Convert to Cluster**, the synchronization process starts. This value will start at 0% and will gradually increase to 100%. When it reaches 100%, it means that the conversion has been finished and the nodes are now in HA status.
- If one of your nodes becomes unavailable and you decide to reinstall SSB, you will have to rejoin your cluster again by navigating to **Basic Settings > High Availability** and clicking **Join HA**. This will start the synchronization progress from 0% again and will gradually increase to 100%. When it reaches 100%, it means that the join progress has been finished and the nodes are now in HA status again.
- If a node becomes unavailable for a longer period and then gets joined again, it can be possible that the configuration of the two nodes become different. In this case, the two nodes start the synchronization process again so that the new changes are transferred to the previously unavailable node. This does not necessarily mean that the synchronization value will start at 0%, it is possible that it starts from a number somewhere between 0% and 100%.

For which systems and configurations is it applicable?	Only applicable for HA clusters.
Value change frequency	Following a conversion to a HA cluster, its returned value continuously increases till reaching 100%. After reaching 100%, its returned value rarely changes - or does not change at all.

Related issues and issue indicators	When 100% has not yet been reached, but the process still does not change for a long time.
--	--

Solution:

- Check your HA network.
- Reboot the Secondary node.
- Reboot the HA cluster.
- For technical assistance, [contact our Support Team](#).

Determining whether the HA node is the primary node

SNMP object:	SSB-SNMP-MIB::ssbHAIsPrimary
Type:	TruthValue (SNMP boolean value)
Community (v2c) / Context (v3)	System
Short description:	The current HA node is the primary node

Description: This information is only supplied on HA-cluster nodes and it is available on the SNMP community provided by the boot-firmwares (the ID-based communities on the **Basic Settings > Monitoring > SNMP agent settings** page).

You can monitor which node is the primary HA node, that is, which node is responsible for SSB's business logic. For example, HTTP configuration, log management (syslog-ng, archive, backup), and so on.

Monitoring hardware RAID

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaid
Type:	This is a grouping node
Community (v2c) / Context (v3)	System
Short description:	Detailed information about hardware Raid devices

Description: Monitor syslog-ng Store Box(SSB)'s hardware RAID, which is responsible for providing disk availability (<https://en.wikipedia.org/wiki/RAID#Hardware-based>), for example, if a disk fails. It is used to monitor the status of the disks in an SSB appliance.

Available on SSB appliances (except T1), on the SNMP community provided by the boot-firmwares, that is, the ID-based communities on the **Basic Settings > Monitoring > SNMP agent settings** page.

RAID controller battery state

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidBatteryState
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The battery state of the raid controller. The value of State from Cachevault_Info, or BBU_Info table of StorCLI.
For which systems and configurations is it applicable?	When this particular hardware supports this particular (T4, T10, S, M) RAID type.
Value change frequency	Its returned value is not supposed to change. When it does, it indicates an issue.
Related issues and issue indicators	The issue generally occurs due to power outage or hardware error (for example, natural battery amortization).

Solution:

- For technical assistance, [contact our Support Team](#).
- Check power supply (for example, check if the power cord is damaged, or if the machine is running from battery, and so on).

RAID controller firmware version

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidControllerFirmwareVersion
Type:	string
Community (v2c) / Context (v3)	System
Short description:	Version of the controller's firmware. This value is reported by StorCLI.
For which systems and configurations	When this particular hardware does not support this particular (T4, T10, S, M) RAID type.

is it applicable?

Value change frequency Not too often, only in case of RAID firmware update.

Hardware RAID status

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidStatus

Type: string

Community (v2c) / Context (v3) System

Short description: Status of the hardware raid.

For which systems and configurations is it applicable? When this particular hardware supports this particular (T4, T10, S, M) RAID type.

Value change frequency Its value does not change too often.

Related issues and issue indicators A status returned value other than the optimal active indicates that the RAID is in degraded mode.

Hardware RAID synchronization progress

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidSyncRate

Type: Integer32

Community (v2c) / Context (v3) System

Short description: Progress of hardware raid synchronization (in percent).

For which systems and configurations is it applicable? When this particular hardware supports this particular (T4, T10, S, M) RAID type.

Value change frequency When not syncing, it has no returned value. Otherwise, its value is continuously changing and may drop when resyncing.

Related issues and issue indicators When the progress does not change for a longer period of time, it indicates an issue.

Solution:

- Reboot your SSB appliance.
- For technical assistance, [contact our Support Team](#).

StorCLI's PD LIST (Physical Drive)

Device ID (DID)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdDeviceID
Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'DID' output in 'PD LIST'.

Drive Group (DG)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdDriveGroup
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'DG' output in 'PD LIST'.

Enclosure Device ID and Slot Number (EID:SlT)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdEnclosureDeviceIDAndSlotNumber
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'EID:SlT' output in 'PD LIST'.

Interface (Intf)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdInterface
---------------------	--

Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Intf' output in 'PD LIST'.

Media Type (Med)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdMediaType
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Med' output in 'PD LIST'.

Model

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdModel
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Model' output in 'PD LIST'.

Protection Info (PI)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdProtectionInfo
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'PI' output in 'PD LIST'.

Sector Size (SeSz)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdSectorSize
Type:	string

Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'SeSz' output in 'PD LIST'.

Self Encrypting Drive (SED)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdSelfEncryptingDrive
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'SED' output in 'PD LIST'.

Size

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdSize
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Size' output in 'PD LIST'. The size of the physical drive.

Spun (Sp)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdSpun
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Sp' output in 'PD LIST'.

State

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPdState
Type:	string

Community (v2c) / Context (v3) System

Short description: The value of StorCLI's 'State' output in 'PD LIST'.

SSB-SNMP-MIB::ssbHardwareRaidPDTable

For which systems and configurations is it applicable? When this particular hardware supports this particular (T4, T10, S, M) RAID type.

Value change frequency The table's contents are not supposed to change (except for when one or more disks within the RAID fail, or in case of disk error).

Related issues and issue indicators In case of hardware error or disk error, one or more disks within the RAID may fail. In this case, the HotSpare will take the faulty disks' place and it is advisable to find out what caused disk failure.

Solution:

- Replace faulty disk (for technical assistance, [contact our Support Team](#), or an official supplier).

StorCLI's Drive State

Drive Temperature

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidDriveTemperature

Type: string

Community (v2c) / Context (v3) System

Short description: The value of StorCLI's 'Drive Temperature' output from the 'Drive State' information list.

Media Error Count

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidMediaErrorCount

Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Media Error Count' output from the 'Drive State' information list.

Other Error Count

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidOtherErrorCount
Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Other Error Count' output from the 'Drive State' information list.

Predictive Failure Count

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidPredictiveFailureCount
Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Predictive Failure Count' output from the 'Drive State' information list.

Shield Counter

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidShieldCounter
Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Shield Counter' output from the 'Drive State' information list.

S.M.A.R.T alert flagged by drive

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidSmartAlertFlaggedByDrive
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'S.M.A.R.T alert flagged by drive' output from the 'Drive State' information list.

S.M.A.R.T BBM Error Count

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidSmartBBMErrorCount
Type:	Integer32
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'BBM Error Count' output from the 'Drive State' information list. This information only present on machines T4, and T10.

S.M.A.R.T Enclosure Device ID and Slot Number (EID:Slit)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidSmartEnclosureDeviceIDAndSlotNumber
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'EID:Slit' output from the 'Drive' information list.

StorCLI's VD LIST (Virtual Drive)

Access

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidVdAccess
Type:	string

Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Access' output in 'VD LIST'.

Cache

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidVdCache
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Cache' output in 'VD LIST'.

Consist

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidVdConsist
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'Consist' output in 'VD LIST'.

Drive Group and Virtual Drive (DG/VD)

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidVdDriveGroupAndVirtualDrive
Type:	string
Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'DG/VD' output in 'VD LIST'.

Name

SNMP object:	SSB-SNMP-MIB::ssbHardwareRaidVdName
Type:	string
Community (v2c) / Context (v3)	System

Context (v3)

Short description: The value of StorCLI's 'Name' output in 'VD LIST'. The name of the virtual drive.

Scheduled Check Consistency (sCC)

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidVdScheduledCheckConsistency

Type: string

**Community (v2c) /
Context (v3)** System

Short description: The value of StorCLI's 'sCC' output in 'VD LIST'.

Size

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidVdSize

Type: string

**Community (v2c) /
Context (v3)** System

Short description: The value of StorCLI's 'Size' output in 'VD LIST'. The size of the virtual drive.

State

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidVdState

Type: string

**Community (v2c) /
Context (v3)** System

Short description: The value of StorCLI's 'State' output in 'VD LIST'. The type of the virtual drive.

Type

SNMP object: SSB-SNMP-MIB::ssbHardwareRaidVdType

Type: string

Community (v2c) / Context (v3)	System
Short description:	The value of StorCLI's 'TYPE' output in 'VD LIST'. The type of the virtual drive.

SSB-SNMP-MIB::ssbHardwareRaidVDTable

For which systems and configurations is it applicable?	When this particular hardware supports this particular (T4, T10, S, M) RAID type.
Value change frequency	The values in this table are not supposed to change.
Related issues and issue indicators	If the value of table element <code>ssbHardwareRaidVdState</code> is other than <code>Opt1</code> and there is no RAID synchronization (which is indicated by object value <code>SSB-SNMP-MIB::ssbHardwareRaidSyncRate</code>), it indicates a RAID status error.

Solution:

- The statuses indicating that the RAID is in degraded mode probably indicate physical disk error. For technical assistance, [contact our Support Team](#).

Monitoring software RAID

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaid
Type:	This is a grouping node
Community (v2c) / Context (v3)	System
Short description:	Detailed information about software Raid devices

Description: Monitor syslog-ng Store Box(SSB)'s software RAID, which is responsible for providing disk availability (<https://en.wikipedia.org/wiki/RAID#Software-based>) for example, if a disk fails. It is used to monitor the status of the disks in an SSB T1 appliance.

Available on SSB appliances (except T1), on the SNMP community provided by the boot-firmwares, that is, the ID-based communities on the **Basic Settings > Monitoring > SNMP agent settings** page.

Software RAID device name

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidDeviceName
Type:	string
Community (v2c) / Context (v3)	System
Short description:	Name of the software Raid device

Software RAID device status

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidDeviceStatus
Type:	string
Community (v2c) / Context (v3)	System
Short description:	Status of the software Raid device (active / degraded / resync / resync-delayed)

Software RAID device RAID level

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidDeviceRaidLevel
Type:	string
Community (v2c) / Context (v3)	System
Short description:	Current Raid level of the software Raid device

RAID synchronization progress

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidDeviceSyncProgress
Type:	Integer32 (%)
Community (v2c) / Context (v3)	System
Short description:	Raid synchronization progress

RAID synchronization speed

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidDeviceSyncSpeed
Type:	Integer32 (KiB/sec)
Community (v2c) / Context (v3)	System
Short description:	Raid sync speed

SSB-SNMP-MIB::ssbSoftwareRaidDeviceTable

For which systems and configurations is it applicable?	When this particular hardware has a software RAID (T1).
Value change frequency	The values of these table elements rarely change. In particular, the table columns <code>ssbSoftwareRaidDeviceSyncProgress</code> and <code>ssbSoftwareRaidDeviceSyncSpeed</code> only have values in case of RAID synchronization.
Related issues and issue indicators	<ol style="list-style-type: none"> 1. If <code>ssbSoftwareRaidDeviceSyncProgress</code> does not reach 100% for a long time, and the percentage or progress value does not change, either. 2. If the returned value for <code>ssbSoftwareRaidDeviceSyncSpeed</code> is too low or stops for a long time. 3. If <code>ssbSoftwareRaidDeviceStatus</code> returns degraded.

Solution:

- Reboot your SSB appliance.
- For technical assistance, [contact our Support Team](#).

Block device name

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidBlockDeviceName
Type:	string
Community (v2c) / Context (v3)	System
Short description:	Name of the block device

Block device S.M.A.R.T health check status

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidBlockDeviceSMARTHealthCheckStatus
Type:	string
Community (v2c) / Context (v3)	System
Short description:	S.M.A.R.T health check status for the block device

Block device temperature

SNMP object:	SSB-SNMP-MIB::ssbSoftwareRaidBlockDeviceTemperature
Type:	Integer32 (Celsius)
Community (v2c) / Context (v3)	System
Short description:	Current temperature of the block device

SSB-SNMP-MIB::ssbSoftwareRaidBlockDeviceTable

For which systems and configurations is it applicable?	When this particular hardware has a software RAID (T1).
Value change frequency	<ol style="list-style-type: none"> BlockDeviceTemperature may change depending on disk load. It is an ambient analog value, and thus continuously changing, but not problematic within a given value range. Note that manufacturing standards allow for a maximum of 60-70 °C, with a normal temperature range of 35-45 °C. RAIDBlockDeviceSMARTHealthCheckStatus is not supposed to change. Ideally, it is supposed to return Passed. If it returns Failed, it indicates an issue.
Related issues and issue indicators	<ol style="list-style-type: none"> If BlockDeviceTemperature is too high for too long. Note that manufacturing standards allow for a maximum of 60-70 °C, with a normal temperature range of 35-45 °C. If RAIDBlockDeviceSMARTHealthCheckStatus returns Failed.

Solution:

- For technical assistance, [contact our Support Team](#).

Monitoring web server, root CA and TSA certificates

Certain properties of the SSL certificates on **Basic Settings > Management > SSL certificate** can be monitored in a table format.

SNMP object:	SSB-SNMP-MIB::ssbServerCertificateTable
Type:	This is a grouping node
Community (v2c) / Context (v3)	Data
Short description:	Detailed information about syslog-ng Store Box(SSB) server certificates

Certificate table index

SNMP object:	SSB-SNMP-MIB::ssbCertificateIndex
Type:	index
Short description:	The index of the certificate or certificate chain in this table.

Certificate Title field

SNMP object:	SSB-SNMP-MIB::ssbCertificateTitle
Type:	string
Community (v2c) / Context (v3)	Data
Short description:	The purpose SSB uses the certificate or certificate chain for.

Certificate "Subject" field

SNMP object:	SSB-SNMP-MIB::ssbCertificateSubject
Type:	string
Community (v2c) / Context (v3)	Data

Short description: The 'Subject' field of the certificate. In case of certificate chains, the 'Subject' field of the end-entity certificate (also called the server certificate).

Certificate "Not before" field

SNMP object: SSB-SNMP-MIB::ssbCertificateNotBefore

Type: date-time string

**Community (v2c) /
Context (v3)** Data

Short description: The 'Not before' field of the certificate in generalized ASN.1 format (YYYYMMDDhhmmssZ or YYYYMMDDhhmmss+hhmm or YYYYMMDDhhmmss-hhmm). In case of certificate chains, this is the latest 'Not before' date of the certificates in the chain.

Certificate "Not after" field

SNMP object: SSB-SNMP-MIB::ssbCertificateNotAfter

Type: date-time string

**Community (v2c) /
Context (v3)** Data

Short description: The 'Not after' field of the certificate in generalized ASN.1 format (YYYYMMDDhhmmssZ or YYYYMMDDhhmmss+hhmm or YYYYMMDDhhmmss-hhmm). In case of certificate chains, this is the earliest 'Not after' date of the certificates in the chain.

Troubleshooting SSB

This section describes the tools to detect networking problems, and also how to collect core files and view the system logs of syslog-ng Store Box (SSB).

Locating the SSB appliance in the server room

SSB T4 and SSB T10

To find syslog-ng Store Box (SSB) appliance T4 and SSB T10 in the server room, you can use IPMI to control the front panel and back panel identify lights.

1. On SSB T4 and SSB T10, navigate to **Basic Settings > System > Hardware information > Blink system identification lights**.
2. To blink the blue LEDs on the front and back of the SSB appliance, click **On**.

| **NOTE:** SSB T1 does not support identify lights.

Alternatively, use the command line as follows:

1. To start the blinking of the blue LEDs on the front and back of the SSB appliance, enter:
ipmitool chassis identify force
2. To stop the blinking of the blue LEDs on the front and back of the SSB appliance, enter:
ipmitool chassis identify 0

syslog-ng Store Box Appliance 3000 and syslog-ng Store Box Appliance 3500

To locate syslog-ng Store Box Appliance 3000 and syslog-ng Store Box Appliance 3500, look for an appliance that has a white front panel with a light blue One Identity logo on it.

Network troubleshooting

The **Troubleshooting** menu provides a number of diagnostic commands to resolve networking issues. Logfiles of syslog-ng Store Box(SSB) can also be displayed here — for details, see [Viewing logs on SSB](#) on page 426.

Figure 222: Basic Settings > Troubleshooting > System debug — Network troubleshooting with SSB

The screenshot displays the SSB Network Troubleshooting interface. On the left, there is a navigation menu with categories like Basic Settings, User menu, and System monitor. The main content area is divided into several functional sections: Ping, Traceroute, Connect to TCP port, View log files, System debug, and Core files. The System monitor section provides system status information, including time, remaining time, and resource usage (CPU, Mem, Disk, Swap). The System debug section allows users to toggle debug mode and collect system state information.

The following commands are available:

- **ping**: Sends a simple message to the specified host to test network connectivity.
- **traceroute**: Sends a simple message from SSB 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 commands above

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 **Port** field.
3. Click the respective action button to execute the command.
4. Check the results in the popup window. Log files are displayed in a separate browser window.

Gathering data about system problems

syslog-ng Store Box (SSB) automatically generates core files if an important software component (for example, syslog-ng, or the indexer) of the system crashes for some reason. These core files can be of great help to the One Identity Support Team to identify problems. To display a list of alerts, navigate to **Search > Log Alerts**.

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

By default, core files are deleted after 14 days. To change the deletion timeframe, navigate to **Basic Settings > Management > Core files**.

Viewing logs on SSB

The **Troubleshooting** menu provides an interface to view the logs generated by the various components of syslog-ng Store Box (SSB). For details on how to browse the log messages received by SSB from its peers, see [Searching log messages](#) on page 330.

NOTE: For performance considerations, 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 SSB

1. Navigate to **Basic Settings > Troubleshooting > View log files**.
2. Use the **Logtype** roll-down menu to select the message type.
 - SSB: Logs of the SSB web interface.
 - syslog: All system logs of the SSB host.
 - syslog-ng: Internal log messages of the built-in syslog-ng server. These logs do not contain messages received from the peers.
3.
 - 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 **Message:** field.

The **Message:** field acts as a generic filter: enter a keyword or a POSIX (basic) regular expression to display only messages that contain the keyword or match the expression.

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 syslog-ng Store Box(SSB), and various system-statistics.

NOTE: Sensitive data like key files and passwords are automatically removed from the files.

The **Basic Settings > Management > Debug logging > Enable debug logs** option is not related to the verbosity of log messages: it adds the commands executed by the SSB web interface to the log.

To collect system-state information, navigate to **Basic Settings > Troubleshooting > System debug** and click **Collect and save current system state info**, 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

1. Navigate to **Basic Settings > Troubleshooting > System debug**.

Figure 223: Basic Settings > Troubleshooting > System debug – Collecting debug information

The screenshot shows the 'System debug' configuration page. On the left is a navigation menu with categories: Basic Settings (Network, System, High Availability, Date & Time, Management, Monitoring, Alerting, Troubleshooting, Dashboard), AAA, Policies, Log, Search, and Reports. Below this is a 'User menu' with options: Private keystore, Change password, Preferences, and Logout. The 'System monitor' section displays system status: Time: 2019-11-08 12:04, Remaining time: 11:59:22, Locked: admin@10.12.15.104, HA status: HA active, Modules: syslog-ng: Running, Active Hosts: 3, Senders: 2, Load 1: 0.10, Load 15: 0.09. A bar chart shows CPU (4%), Mem (9%), Disk (0%), and Swap (0%) usage. The main content area includes: 'Ping' (Hostname: 127.0.0.1, Ping host button), 'Traceroute' (Hostname: 127.0.0.1, Traceroute host button), 'Connect to TCP port' (Hostname: 127.0.0.1, TCP port: 22, Connect to host button), 'View log files' (Logtype: syslog, Day: Friday, Search patterns: Message: [input], Download, View, Tail buttons), 'System debug' (Debug mode is OFF, Collect and save current system state info button, Start/Stop buttons, Save collected debug info button), and 'Core files' table with columns: File, Size, Creation date, Encrypt for download, Download, Delete.

2. Click **Start**.

NOTE: Starting debug mode increases the log level of SSB, and might cause performance problems if the system is under a high load.

3. Reproduce the event that causes the error, for example send a log message from a client.
4. Click **Stop**.
5. Click **Save the collected debug info** and save the created zip file. The name of the file uses the debug_info-<hostname>YYYYMMDDHHMM format.
6. Attach the file to your support ticket.

Troubleshooting an SSB cluster

The following sections help you to solve problems related to high availability clusters.

- For a description of the possible statuses of the syslog-ng Store Box (SSB) cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured), see [Understanding SSB cluster statuses](#) on page 429.
- To recover a cluster that has broken down, see [Recovering SSB if both nodes broke down](#) on page 431.
- To resolve a split-brain situation when the nodes of the cluster were simultaneously active for a time, see [Recovering from a split brain situation](#) on page 432.
- To replace a broken node with a new appliance, see [Replacing a node in an SSB HA cluster](#) on page 435.

Understanding SSB cluster statuses

This section explains the possible statuses of the syslog-ng Store Box (SSB) cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured). SSB displays this information on the **Basic Settings > High Availability** page.

Status

The **Status** field indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode. The status of the individual SSB nodes is indicated in the **Node HA status** field of the each node. The following statuses can occur:

- **Standalone:** There is only one SSB unit running in standalone mode, or the units have not been converted to a cluster (the **Node HA status** of both nodes is standalone). Click **Convert to Cluster** to enable High Availability mode.
- **HA:** The two SSB nodes are running in High Availability mode. **Node HA status** 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 SSB nodes are running in High Availability mode. **Node HA status** is HA on both nodes, but the **Node HA UUID** is different. Contact the One Identity Support Team for help. For contact details, see [About us](#) on page 450.
- **Degraded:** SSB 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 slave 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 SSB 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 (master) for a time.

⚠ CAUTION:

Hazard of data loss In this case, valuable log messages might be available on both SSB 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 432](#).

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 cluster 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 fixed, which helps if the HA connection between the nodes is slow.

DRBD status

The **DRBD status** field indicates whether the latest data (including SSB configuration, log files, and so on) is available on both SSB nodes. The master 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 SSB nodes. The following statuses are possible:

- **Connected:** Both nodes are functioning properly.
- **Connected (Disk Failure):** A hard disk of the slave 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 master node. If you absolutely must shutdown the master node during synchronization, shutdown the slave node first, and then the master node.

- **Split brain:** The two nodes lost the connection to each other, with the possibility of both nodes being active (master) for a time.

⚠ CAUTION:

Hazard of data loss In this case, valuable log messages might be available on both SSB 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 432.

- **WFConnection:** One node is waiting for the other node. The connection between the nodes has not been established yet.

Redundant Heartbeat status

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 141.

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 SSB 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 contact details, see [About us](#) on page 450.

Recovering SSB if both nodes broke down

It can happen that both nodes break down simultaneously (for example because of a power failure), or the slave node breaks down before the original master node recovers. The following describes how to properly recover syslog-ng Store Box (SSB).

NOTE: When both nodes of a cluster boot up in parallel, the node with the 1.2.4.1 HA IP address will become the master node.

To properly recover SSB

1. Power off both nodes by pressing and releasing the power button.

CAUTION:

Hazard of data loss If SSB does not shut off, press and hold the power button for approximately 4 seconds. This method terminates connections passing SSB and might result in data loss.

2. Power on the node that was the master before SSB broke down. Consult the system logs to find out which node was the master before the incident: when a node boots as master, or when a takeover occurs, SSB sends a log message identifying the master node.

TIP: Configure remote logging to send the log messages of SSB to a remote server where the messages are available even if the logs stored on SSB become unaccessible. For details on configuring remote logging, see [SNMP and e-mail alerts](#) on page 81.

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 (master) role while disconnected. This might cause new data (for example, log messages) 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 easily merged.

CAUTION:

Hazard of data loss In a split brain situation, valuable log messages might be available on both syslog-ng Store Box (SSB) nodes, so special care must be taken to avoid data loss.

The nodes of the SSB cluster automatically recognize the split brain situation once the connection between the nodes is re-established, and do not perform any data synchronization to prevent data loss. When a split brain situation is detected, it is visible on the SSB system monitor, in the system logs (Split-Brain detected, dropping connection!), on the **Basic Settings > High Availability** page, and SSB sends an alert as well.

NOTE: After the connection between the nodes has been restored, the split brain situation will persist. The core firmware will be active on one of the nodes, while it will not start on the other node.

Once the network connection between the nodes has been re-established, one of the nodes will become the master node, while the other one will be the slave node. Find out which node is the master node. There are two ways to identify the master node:

- Locally: Log in to each SSB locally, and wait for the console menu to come up. The console menu only appears on the master node.
- Remotely: Try connecting to each SSB using SSH. It is only the master node that you can directly connect to via SSH. The slave node cannot be accessed externally, only via SSH from the master.

To recover an SSB cluster from a split brain situation, complete the procedures described in [Data recovery](#) and [HA state recovery](#).

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 slave host. This is required because data on this host will later be overwritten by the data available on the current master.

NOTE: During data recovery, there will be no service provided by SSB.

To configure recovering from a split brain situation

1. Log in to the master node as root locally (or remotely using SSH) to access the Console menu. If no menu is showing up after login, then this is the slave node. Try the other node.
2. Select **Shells > Boot Shell**.
3. Enter `/usr/share/heartbeat/hb_standby`. This will change the current slave node to master and the current master node to slave (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 syslog-ng.service** command to disable all traffic going through SSB.
9. Save the files from `/opt/ssb/var/logspace/` 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. Exit the console.
11. Log in again, and select **Shells > Boot Shell**.
12. Enter `/usr/share/heartbeat/hb_standby`. This will change the current slave node to master and the current master node to slave (HA failover).
13. Exit the console.

14. Wait a few minutes to let the failover happen, so the node you were using will become the slave node and the other node will become the master node.

The nodes are still in a split-brain state but now you have all the data backed up from the slave node, and you can synchronize the data from the master node to the slave node, which will turn the HA state from "Split-brain" to "HA". For details on how to do that, see [HA state recovery](#).

HA state recovery

In the procedure described here, the "Split-brain" state will be turned to the "HA" state.

⚠ CAUTION:

Keep in mind that the data on the current master node will be copied to the current slave node and *data that is available only on the slave node will be lost* (as that data will be overwritten).

To swap the nodes (optional)

NOTE: If you completed the procedure described in [Data recovery](#), 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 master node the slave node and the slave node the master node, perform the following steps.

1. Log in to the master node as root locally (or remotely using SSH) to access the Console menu. If no menu is showing up after login, then this is the slave node. 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 slave node and the other node will be the master node.

To initialize data synchronization

1. Log in to the slave node as root locally (or remotely using SSH) to access the Console menu. If the menu is showing up, then this is the master node. Try logging in to the other node.

Note that you are in the boot shell now as on the slave node, only the boot shell is available.

2. Invalidate the DRBD. Issue the following commands:

```
drbdadm secondary r0
```

```
drbdadm connect --discard-my-data r0
```

```
ssh sb-other
```

drbdadm connect r0

3. Reboot the slave node.

Following this step, the master will be in **Standalone** state, while the slave's DRBD status will be **WFCnection**.

The console will display an **Inconsistent (10)** message. This is normal behavior, and it is safe to ignore this message.

4. Reboot the master node. The SSB cluster will now be functional, accepting traffic as before.
5. After both nodes reboot, the cluster should be in **Degraded Sync** state, the master being **SyncSource** and the slave being **SyncTarget**. The master node should start synchronizing its data to the slave node. Depending on the amount of data, this can take a long time. To adjust the speed of synchronization, see [Adjusting the synchronization speed](#) on page 140.
6. Enable all incoming traffic on the master node. Navigate to **Basic Settings > System > Service control > Syslog traffic, indexing & search:** and click **Enable**.

If the web interface is not accessible or unstable, complete the following steps on the active SSB:

- a. Log in to SSB as root locally (or remotely using SSH) to access the console menu.
- b. Select **Shells > Core Shell**, and issue the **systemctl start syslog-ng.service** command.
- c. Issue the **date**, and check the system date and time. If it is incorrect (for example, it displays 2000 January), replace the system battery. For details, see the hardware manual of the appliance.

Replacing a node in an SSB HA cluster

This section describes how to replace a unit in an syslog-ng Store Box (SSB) cluster with a new appliance.

To replace a unit in an SSB 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 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 98.
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](#) account and download the CD ISO for the same SSB 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 interface if needed.
7. When the installation is finished, connect the two SSB 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. Log in to the working node and verify the HA state. Select **Basic Settings > High Availability**. The **Status** field should display HALF.
10. Reconfigure the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces. Click **Commit**.
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, management to 2 or MGMT) as needed for your environment.

Expected result:

A node of the SSB cluster is replaced with a new appliance.

Resolving an IP conflict between cluster nodes

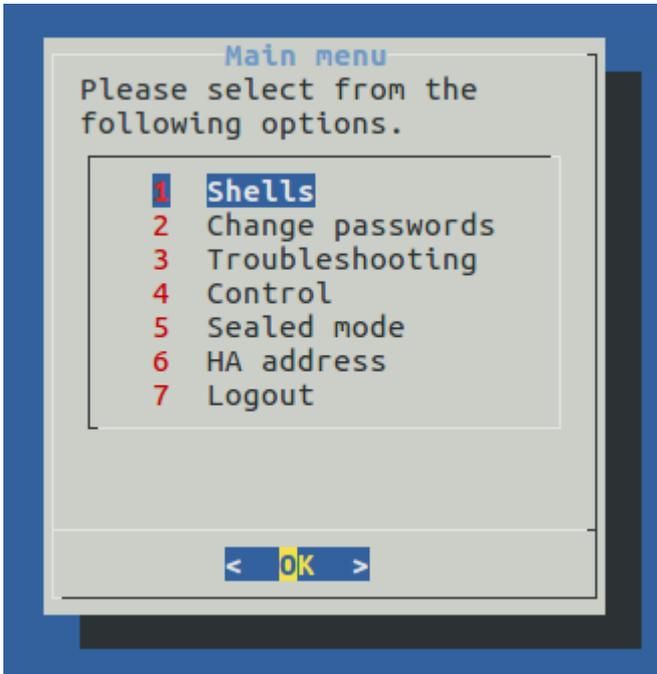
The IP addresses of the HA interfaces connecting the two nodes are detected automatically, during boot. When a node comes online, it attempts to connect to the IP address 1.2.4.1. If no other node responds until timeout, then it sets the IP address of its HA interface to 1.2.4.1, otherwise (if there 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 cluster.

To manually assign the correct IP address to the HA interface of a node, perform the following steps:

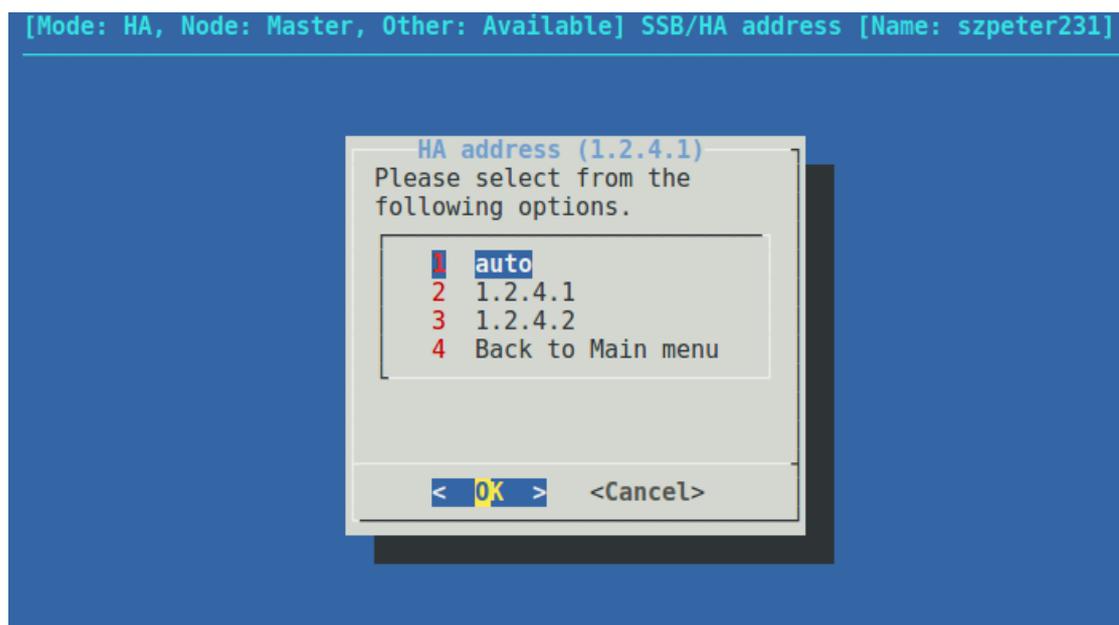
1. Log in to the node using the IPMI interface or the physical console.
Configuration changes have not been synced to the new (replacement) node, as it could not join the HA cluster. Use the default password of the root user of syslog-ng Store Box (SSB), see ["Installing the SSB hardware" in the Installation Guide](#).
2. From the console menu, choose **6 HA address**.

Figure 224: The console menu



3. Choose the IP address of the node.

Figure 225: The console menu



4. Reboot the node.

Restoring SSB configuration and data

The following procedure describes how to restore the configuration and data of syslog-ng Store Box (SSB) from a complete backup, for example, after a hardware replacement.

NOTE: It is possible to receive indexer errors following data restore. Data that was still in the memory of SSB during backup might have been indexed, but as it was not on the hard drive, it was not copied to the remote server.

To make sure that all data is backed up (for example, before an upgrade), shut down syslog-ng before initiating the backup process.

⚠ CAUTION:

Statistics about syslog-ng and logspace sizes are not backed up. As a result, following a data restore, the Basic Settings > Dashboard page will not show any syslog-ng and logspace statistics about the period before the backup.

To restore the configuration and data of SSB from a complete backup

1. Connect to your backup server and locate the directory where SSB saves the backups. The configuration backups are stored in the config subdirectory in time stamped files. Find the latest configuration file (the configuration files are called SSB-timestamp.config).
2. Connect to SSB.

If you have not yet completed the Welcome Wizard, click **Browse**, select the configuration file, and click **Import**.

If you have already completed the Welcome Wizard, navigate to **Basic Settings > System > Import configuration > Browse**, select the configuration file, and click **Import**.

3. Navigate to **Policies > Backup & Archive/Cleanup**. Verify that the settings of the target servers and the backup protocols are correct.
4. 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 may take a long time.
5. Navigate to **Log > Logspaces**, 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.

Configuring the IPMI interface from the BIOS after losing IPMI password

It may happen that you inadvertently lose the IPMI password of your syslog-ng Store Box (SSB). In that case, you will be required to:

1. Shut down SSB.
2. Unplug the SSB physical appliance's power cord.
3. Wait 30 seconds.
4. Replug the power cord.
5. Restart the appliance.
6. Re-configure the IPMI interface from the BIOS.

To configure IPMI from the BIOS, complete the following steps.

Prerequisites:

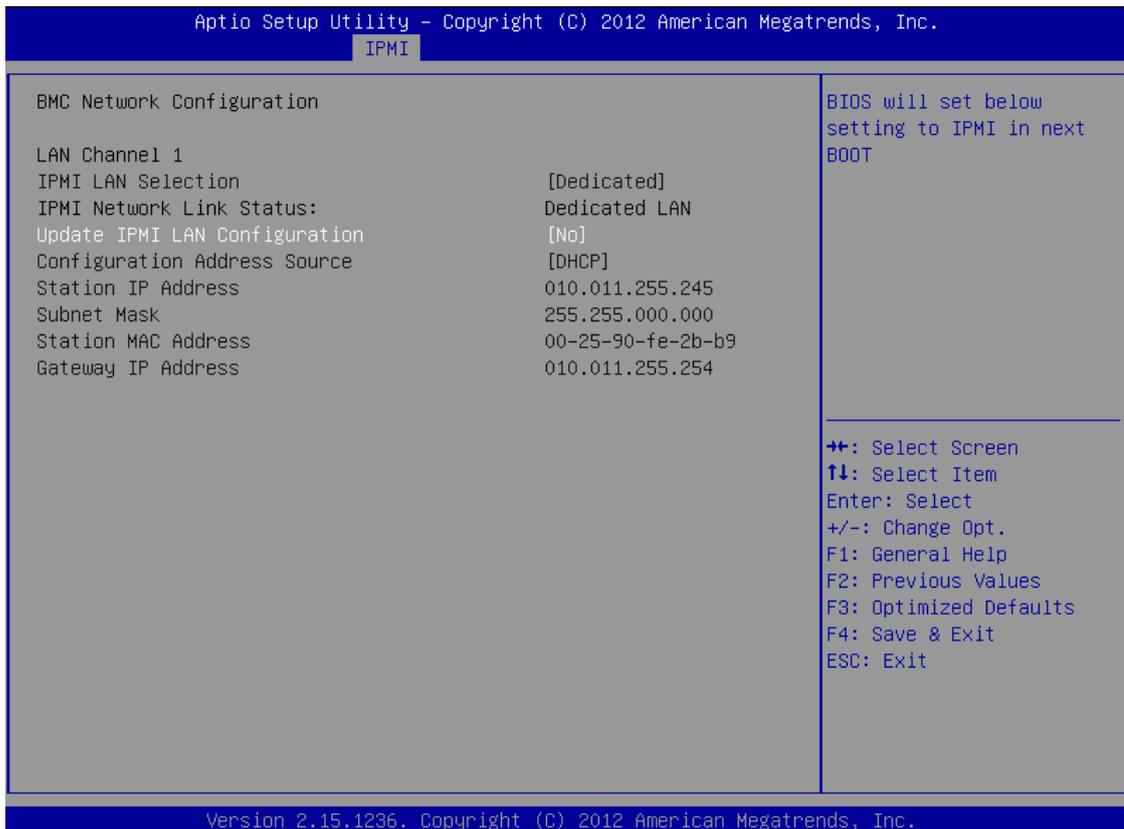
To apply the procedure outlined here, you will need physical access to a monitor and keyboard.

Figure 227: IMPI page > BMC Network Configuration option



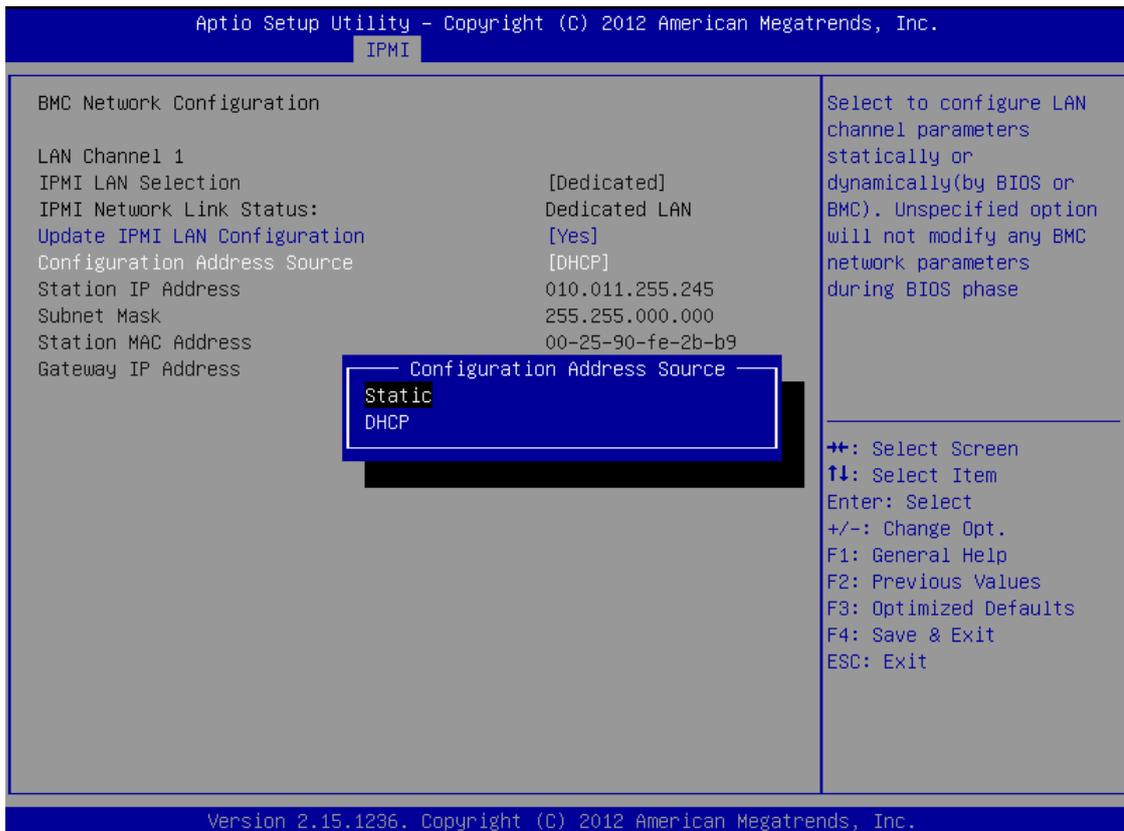
4. On the **BMC Network Configuration** page, select **Update IPMI LAN Configuration**, press Enter, and select **Yes**.

Figure 228: BMC Network Configuration page > Update IPMI LAN Configuration



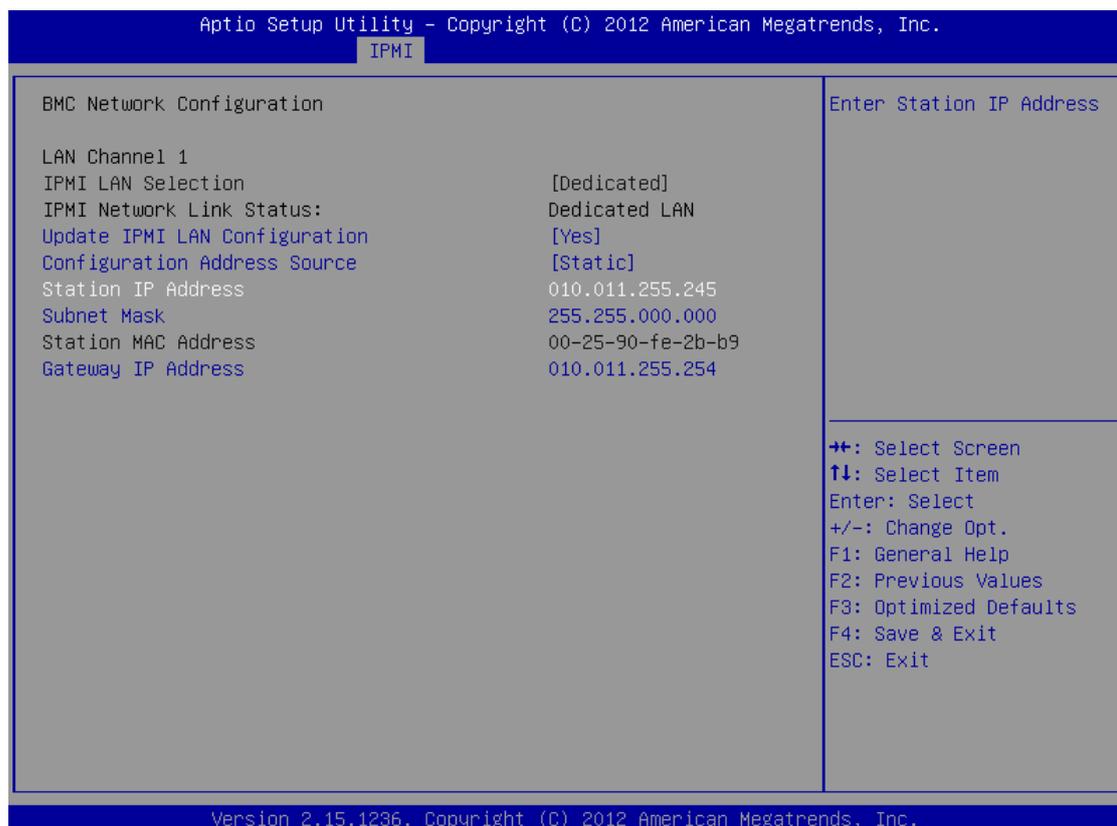
5. Stay on the **BMC Network Configuration** page, select **Configuration Address Source**, press **Enter**, and select **Static**.

Figure 229: 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 230: 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.

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 details, see [Generating TSA certificate with Windows Certificate Authority on Windows Server 2008](#) or [Generating TSA certificate with Windows Certificate Authority on Windows Server 2012](#).

Correct Alerting & Monitoring and Management group privilege mismatch

When attempting to upgrade to syslog-ng Store Box(SSB) version 6.3 (or newer), or upload an older configuration, you might see an error message similar to the following:

```
Pre-check failed
Error upgrading XML database - One or more user groups have group privilege
mismatch for the Basic Settings > Alerting & Monitoring and Basic Settings >
Management pages.
To continue the upgrade or import process, first, complete the steps in section
'Correct Alerting & Monitoring and Management group privilege mismatch' in
https://docadmin.quest.com/syslog-ng-store-box/6.3.0/administration-
guide/troubleshooting-ssb
```

This error is probably due to the fact that one or more of your usergroups currently have

- **Basic Settings > Alerting & Monitoring** group privilege, but no **Basic Settings > Management** group privilege, or have
- **Basic Settings > Management** group privilege, but no **Basic Settings > Alerting & Monitoring** group privilege.

The content of the **Alerting & Monitoring** and **Management** group privileges changed in SSB version 6.3 the following way:

- Instead of a single menu item, the configuration of **Basic Settings > Alerting & Monitoring** will be split to **Basic Settings > Alerting** and **Basic Settings > Monitoring**.
- The configuration of **Basic Settings > Management > SNMP trap settings** and **Basic Settings > Management > SNMP agent settings** will be moved to **Basic Settings > Alerting > SNMP trap settings** and **Basic Settings > Monitoring > SNMP agent settings** in the menu, respectively.

These changes to the group privilege contents will NOT be mapped during the upgrade process.

As a result, the usergroups that currently have **Basic Settings > Alerting & Monitoring** group privilege would gain additional access rights to **SNMP trap settings** and **SNMP agent settings**.

Before you start the upgrade process, you must change the group privileges of these usergroups the following way:

1. Navigate to **AAA > Access Control**, and **Edit** the group privileges of the usergroup that
 - only has **Basic Settings > Alerting & Monitoring** group privilege, but no **Basic Settings > Management** group privilege, or
 - only has **Basic Settings > Management** group privilege, but no **Basic Settings > Alerting & Monitoring** group privilege.
2. Change your configuration so that both group privileges are the same. Either grant both **Basic Settings > Management** and **Basic Settings > Alerting & Monitoring** group privileges to the usergroup, or revoke both group privileges.
3. Repeat the previous steps for all affected usergroups.
4. **Commit** your changes and retry upgrading your machine to SSB version 6.3.

Security checklist for configuring SSB

The following checklist is a set of recommendations and configuration best practices to ensure that your syslog-ng Store Box(SSB) is configured securely.

General security recommendations

- As a general recommendation, use 2048-bit RSA keys (or stronger), AES-256-CBC cipher (or stronger), and SHA-256 hash algorithm (or stronger). For more specific information, see the relevant sections of the [Administration Guide](#).
- Use mutual authentication whenever possible, as detailed below, when configuring log sources, log destinations or LDAP user database.
- One Identity recommends that you generate certificates using your own public key infrastructure (PKI) solution and then upload them to SSB. Certificates generated by SSB cannot be revoked, therefore, they can become a security risk if compromised.
- When exporting the configuration of SSB, or creating configuration backups, always use encryption. Handle the exported data with care, as it contains sensitive information, including credentials. For more information 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 same certificate for the SSB webserver and for encrypting logstores.
- For backward compatibility reasons, SSB does not enforce strict security configuration for backup, archive, and share - using SMB/CIFS and NFS - therefore, any security expectations need to be ensured by the joining peers and the underlying network architecture. For more information on backups and archiving, see "[Data and configuration backups](#)" in the [Administration Guide](#) and "[Archiving and cleanup](#)" in the [Administration Guide](#).

Log traffic and storage specific security recommendations

- When creating logspaces on **Log > Logspaces**, use **LogStore** type rather than plain text files and apply encryption.
- When encrypting log files, One Identity recommends:
 - Using 2048-bit RSA keys (or stronger). For more information, see ["Creating logstores" in the Administration Guide](#).
 - Using AES-256-CBC cipher (or stronger) and SHA-256 hash algorithm (or stronger). For more information, see [Administration Guide](#).
- One Identity recommends using User Temporary private key store for decrypting and viewing encrypted logs on the **Search > Logspaces** interface. Avoid using User Permanent private key store or shared decryption private key uploaded on the **Log > Logspaces** interface. For more information, see ["Browsing encrypted logspaces" in the Administration Guide](#).
- For the Server certificate and the Time Stamping Authority (TSA) certificate, upload the private key as well. One Identity recommends using 2048-bit RSA keys (or stronger). These two certificates must be issued by the same Certificate Authority. For more information on uploading certificates and keys created with an external PKI, see ["Uploading external certificates to SSB" in the Administration Guide](#).
- When granting user privileges, make sure that only the intended users can access logspaces.

By default, members of the search group can view the stored messages online. Use the *Access control* option to control which usergroups can access a logspace. For more information, see ["Managing user rights and usergroups" in the Administration Guide](#).

- Configure each logsource in SSB at **Log > Sources** as follows:
 1. For **Transport**, select **TLS**.
 2. For **Incoming log protocol and message format**, select **Syslog (IETF-syslog, RFC 5452)**.
 3. For **Peer verification**, select **Required-trusted**.
 4. For **Cipher suite**, select **Secure**.

By applying the **Secure** cipher suite, SSB will not allow permissive cipher suites to be used for remote connections.
- If log messages must be forwarded outside the box, configure log destinations at **Log > Destinations** in a similar way as the logsources described above (Steps 1-4). Note that you cannot set cipher suites since the TLS server is the remote side (Step 5). For more information, see [Administration Guide](#).
- Consider that connections for log source or destination types UDP, TCP, SQL, and SNMP are not encrypted. Even though ALTP is encrypted, it can still be compromised. For more information, see [Administration Guide](#).
- Enable flow-control to prevent message loss. For more information, see [Administration Guide](#).

Accessing SSB

- Disallow permissive cipher suites for HTTPS connections towards the SSB webserver. When configuring the cipher suite capability for HTTPS connections, use the **Secure** cipher suite set under **Basic Settings > Management > Web interface and RPC API settings > Cipher suite**. For more information, see "[Web interface and RPC API settings](#)" in the *Administration Guide*.
- Use strong passwords, which have at least 12 characters including lower case letters, upper case letters, numbers, and special characters. For local SSB users, set the password policy strength to strong on **AAA > Settings > Minimal password strength**. For more information, see "[Setting password policies for local users](#)" in the *Administration Guide*.
- Accessing the SSB 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 SSB when it is not needed. For more information, see "[Enabling SSH access to the SSB host](#)" in the *Administration Guide*.
- Permit administrative access to SSB only from trusted networks. If possible, log messages from clients and administrative access to the SSB web interface should be originated from separate networks.
- Configure SSB to send an alert if a user fails to login to SSB. For more information, see the **Login failed** alert in "[System related traps](#)" in the *Administration Guide*.
- Configure **Disk space fill up prevention**, and configure SSB to send an alert if the free space on the disks of SSB is low. For more information, see "[Preventing disk space fill up](#)" in the *Administration Guide*.
- Prefer configuring SSB to use the local user database. If LDAP is needed, make sure to configure mutual authentication. For more information on local user management, see "[Managing SSB users locally](#)" in the *Administration Guide*.

Networking considerations

- SSB stores sensitive data. Use a firewall and other appropriate controls to ensure that unauthorized connections cannot access it.
- If possible, enable management access to SSB only from trusted networks.
- Make sure that the HA interface of SSB is connected to a trusted network.
- Make sure that for the communication between the peer nodes, for example, log sending, log receiving, or webserver interface communication, you have the properly secure configuration as described above.

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

A

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.

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.

B

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.

Cadence

[[[Undefined variable TemplateGuideVariables.OneIdentityNameShort]]] font that contains standard icons used in the user interfaces for various [[[Undefined variable TemplateGuideVariables.OneIdentityNameShort]]] products.

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.

certificate chain

An ordered list of certificates, containing an end-user subscriber (or server) certificate and intermediate certificates (that represent the intermediate CAs). A certificate chain enables the receiver to verify that the sender and all intermediate certificates are trustworthy.

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.

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 to a logspace.

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.

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.

Drop-down

Flare default style, that can be used to group content within a topic. It is a resource to structure and collapse content especially in non-print outputs.

E

External network interface

The external interface (labeled 1 or EXT) is used for general communication between the clients and the servers. If the management interface is not configured, the external interface is used for management purposes as well.

F

filter

An expression to select messages.

filtered logspace

The filtered subset of logs contained in an existing local, remote, or multiple logspace. A filtered logspace is created by using the same search expressions and logic as on the Search interface. See also multiple logspace and remote logspace.

firmware

A firmware is a collection of the software components running on SSB. Individual software components cannot be upgraded on SSB, only the entire firmware. SSB

contains two firmwares (an external (or boot) firmware, and an internal (or core) firmware). These are bundled into a single ISO file.

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.

Glossary

List of short definitions of product specific terms.

H

HA network interface

The HA interface (labeled 4 or HA) is an interface reserved for communication between the nodes of SSB clusters.

High Availability

High Availability (HA) uses a second SSB unit (called slave node) to ensure that the services are available even if the first unit (called master 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.

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.

logspace

The virtual container on SSB of log messages collected from clients and from SSB itself. Can be of the type: logstore or plain text logspace. See also logstore and plain text logspace.

logstore

A binary logfile format that can encrypt, compress, and time stamp 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.

M

Management network interface

The management interface (labeled 2 or MGMT) is used exclusively for communication between SSB and the auditor or the administrator of the syslog-ng Store Box.

master node

The active SSB unit that is inspecting the traffic when SSB is used in High Availability mode.

multiple logspace

A logspace that aggregates log messages from several logspaces. A multiple logspace can be searched like any other logspace on SSB, and you can also create filtered logspaces that are based on a multiple logspace. See also filtered logspace.

N**name server**

A network computer storing the IP addresses corresponding to domain names.

node

An SSB unit running in High Availability mode.

Note

Circumstance, that needs special attention.

O**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.

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 SSB device to detect that the other node of the SSB 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**SaaS**

Software-as-a-Service.

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.

Skin

Used to design the online output window.

slave node

The passive SSB unit that replaces the active unit (the master node) if the master becomes unavailable.

Snippet

Flare file type that can be used to reuse content. The One Identity SSB contains various default snippets.

SNMP

Simple Network Management Protocol (SNMP) is an industry standard protocol used for network management. SSB 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 SSB cluster become active (master). 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.

SSB

syslog-ng Store Box

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.

Tip

Additional, usefull information.

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.