



One Identity Safeguard for Privileged Sessions 5.9

Deployment in Single-Interface Router Mode

Copyright 2018 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 indicates a potential for property damage, personal injury, or death.
-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.
-  **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Overview	4
Introduction	5
Inline routing scenario in SPS	6
Single-interface transparent mode	7
Advanced or Policy-based routing	9
Example scenarios	10
Configuring advanced routing on Linux	10
Configuring advanced routing on Cisco routers	12
Configuring advanced routing on Cisco ASA firewalls	15
About us	18
Contacting us	18
Technical support resources	18

Overview

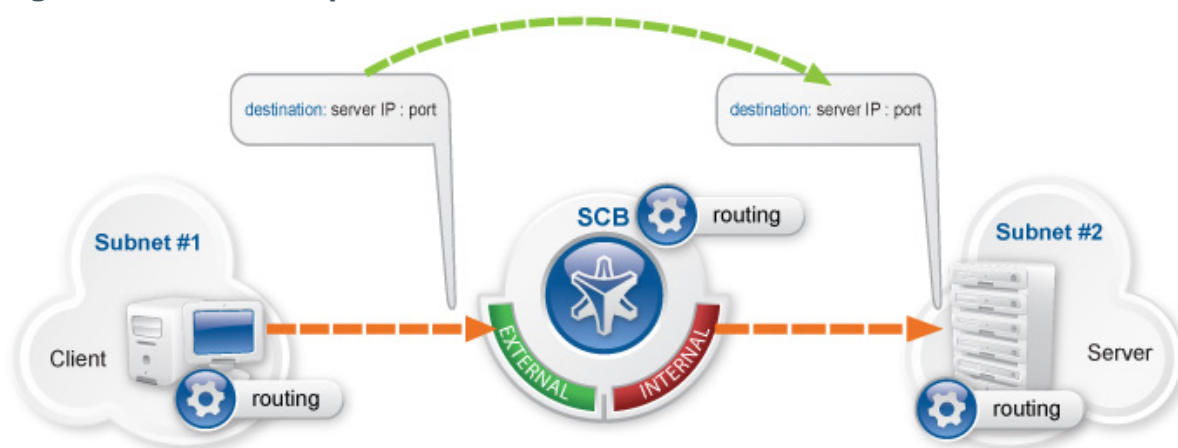
This short document is about a special implementation of the One Identity Safeguard for Privileged Sessions. This makes it possible to deploy the device without changing the network topology, but keeping all the advantages of the transparent mode of the SPS.

Introduction

The One Identity Safeguard for Privileged Sessions connection policies can work in different network models to make it easy to integrate it into an existing network. These two modes are transparent, and non-transparent modes (for details on modes of operation, see ["Modes of operation" in the Administration Guide](#)). The aim is usually the transparent implementation. Although the non-transparent mode can provide some transparency, it is not the best to be used for that purpose.

For the easy-to-deploy and totally transparent solution the transparent mode would be the best. This mode requires integrating SPS in the network level, so all the administrative traffic could pass the box to make it controllable and auditable (for details and illustrations on transparent mode, see ["Transparent mode" in the Administration Guide](#)).

Figure 1: SPS in transparent mode



In most cases it is not possible, or not optimal to integrate SPS into the network as in the abovementioned example, because it would require significant changes to the network topology, and SPS could act as a single point of failure. However, it is possible to use SPS in transparent mode transparently without changing the network layout, with a few additional configuration steps in some of the active network devices (firewalls or routers) and the SPS itself. The following sections will describe this in detail.

Inline routing scenario in SPS


SPS has two physical network interfaces that are normally used for production (monitored) traffic: "External" and "Internal". The function of these interfaces is interchangeable, the names are only used in this document for easier identification. SPS is implemented as it is visible on Figure *SPS in transparent mode*. In this case the connections are coming from the client's network (define it as 192.168.0.0/24) and heading towards the server's network (define it as 10.0.0.0/24). The routing on the client is configured so that it uses SPS as a gateway, when the server network is accessed. The servers are configured so that they send the answers into the client network through SPS. Also, all the networks and gateways are defined in the routing table of SPS, to send the traffic out on the appropriate interface.

SPS does not check whether the client is coming from the "External" interface or if the connections are going out on the "Internal" interface. Because of this, it is possible to create a topology, where both the clients and the servers are located on the "External" side.

Single-interface transparent mode

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

Figure 2: SPS in single-interface transparent mode



Advantages:

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

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

Disadvantages:

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

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

Advanced or Policy-based routing

Usually there is a central network device somewhere close to the location where SPS is planned to be implemented. This central network device (a router or firewall) can facilitate improving the previous layout into a real working scenario, as it is visible on Figure *SPS in single-interface transparent mode*.

There is a router (or firewall, or layer3 switch) between the zones, and SPS is installed into a new, separated network. Here, all the devices (including SPS) are configured to use the central router as their default gateway. So, if a client is trying to reach a server, the connection is going through the router. However, SPS is not able to audit the remote administration with this configuration.

Here comes the router into the play. For example if we have to audit RDP connections, the router can be configured to route all the connections to SPS (connections are coming from the client network, going to the server network, and the destination port is 3389). With this configuration, all RDP connections are "redirected" to SPS. It sends the traffic back to the router, which then sends the connection to its original destination, to the server zone. SPS receives a connection on its "External" interface, and it routes it back on the same interface. It creates a "hook" in the network traffic, but it also makes the connection totally transparent: the client IP can be the same (optional), and the client does not have to know anything about SPS. Non-administrative traffic can also be unaffected, as we can selectively route the necessary traffic to SPS.

The configuration of the router can vary depending on the type of the router. The following two procedures describe configuration scenarios for a Linux and a Cisco router.

Example scenarios

On the SPS side, no special configuration is required. SPS has to be in transparent mode. The "External" interface has to be configured correctly, and has to be connected to the router. On the "Internal" interface any IP address can be configured that is not used on the network, as we will not use this interface at all. The default gateway should be the router.

Configuring advanced routing on Linux

Purpose:

To configure a Linux-based router to redirect selected traffic to SPS instead of its original destination, complete the following steps. This procedure should work on most modern Linux-based routers, including Check Point firewalls.

Prerequisites:

The router must have the `iptables` and `ip` tools installed.

Steps:

1. Create the packet filter rules that will mark the connections to be sent to SPS using the CONNMARK feature of iptables. Mark only those connections that must be redirected to SPS.

```
# iptables -t mangle -I PREROUTING -i <interface-facing-the-clients> -p tcp -d <network-of-the-servers> --dport <port-to-access> -j CONNMARK --set-mark 1
```

Example: Setting up a connection mark for Linux policy routing

For example, if the network interface of the router that faces the clients is called `eth0`, the servers are located in the `10.0.0.0/24` subnet, and the clients

access the servers using port 3389 (the default port of the RDP protocol), then this command looks like:

```
# iptables -t mangle -I PREROUTING -i eth0 -p tcp -d 10.0.0.0/24 --dport 3389 -j CONNMARK --set-mark 1
```

2. Create a rule that redirects the answers of the servers to SPS. That way both the client-to-server and the server-to-client traffic is routed to SPS.

NOTE:

This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 3: Control > Connections — Using SNAT

```
# iptables -t mangle -I PREROUTING -i <interface-facing-the-servers> -p tcp -s <network-of-the-servers> --sport <port-to-access> -j CONNMARK --set-mark 1
```

3. Convert the CONNMARK marks to MARK:

```
# iptables -t mangle -A PREROUTING ! -i <interface-facing-the-scb> -m connmark -mark 1 -j MARK --set-mark 1
```

CAUTION:

This rule must be placed after the CONNMARK rules.

4. Add the table name to the `/etc/iproute2/rt_tables` of the router. Use the following format (for details on routing tables, see for example the [Guide to IP Layer Network Administration with Linux](#)):

```
103 scb
```

5. Create a routing table that has a single entry with a default route to SPS:

```
# /sbin/ip route add default via <ip-address-of-SPS> table scb
```

6. Create a routing rule that selects the routing table called `scb`, if the connection is marked.

```
# /sbin/ip rule add from all fwmark 1 table scb
```

7. If SPS is configured to spoof the IP address of the clients on the server side (that is, the **SNAT > Use original IP address of the client** option of the connection policies is selected), enable spoofing on the router for the interface connected to SPS.

```
# echo 0 > /proc/sys/net/ipv4/conf/<interface-facing-SPS>/rp_filter
# echo 0 > /proc/sys/net/ipv4/conf/all/rp_filter
```

Expected result:

The traffic from the clients targeting the specified port of the servers is redirected to SPS. Therefore, SPS can be configured to control and audit this traffic.

Configuring advanced routing on Cisco routers

Purpose:

To configure a Cisco router to redirect selected traffic to SPS instead of its original destination, complete the following steps. This procedure should work on most modern Cisco IOS releases but was specifically tested on IOS version 12.3.

Steps:

1. Create an ACL (Access Control List) entry that matches the client and server subnets and the to-be-audited port. Keep in mind that whatever is permitted by this ACL is what will be matched, so make sure that the scope of the ACL entry is narrowed down as much as possible.

```
!(config) ip access-list extended ssh-inbound
!(config-ext-nacl) permit tcp <src net> <src mask> <dst net> <dst mask>
eq <dst port>
```

Example: Configuring an ACL entry for Cisco policy routing

For example, if the clients are in the 192.168.0.0/24 subnet, the servers are located in the 10.0.0.0/24 subnet, and the clients access the servers using port 22 (the default port of the SSH protocol), then the permit clause should be:

```
!(config-ext-nacl) permit tcp 192.168.0.0 0.0.0.255 10.0.0.0 0.0.0.255
eq 22
```

TIP:

Cisco ACLs use inverse netmasks for defining network addresses. To calculate an inverse mask given a subnet mask, simply subtract each octet value from 255.

2. Create an ACL entry that matches the reply packets coming from the server zone and targeted at the client zone to make sure that replies are reaching the SPS.

```
 #(config) ip access-list extended ssh-outbound
 #(config-ext-nacl) permit tcp <dst net> <dst mask> eq <dst port> <src
 net> <src mask>
```

NOTE:

This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 4: Control > Connections — Using SNAT

Example: Configuring an ACL entry for reply packets with Cisco policy routing

In case of the example in step 1, the permit clause should be:

```
 #(config-ext-nacl) permit tcp 10.0.0.0 0.0.0.255 eq 22 192.168.0.0
 0.0.0.255
```

3. Create a route-map entry. It controls which packets are affected by policy routing and where they should be forwarded to. The `match` commands specify the conditions under which policy routing occurs. The `set` commands specify the routing actions to perform if the criteria enforced by the `match` commands are met. A new route-map can be defined as follows:

```
 #(config) route-map scb-inbound
```

- a. Set your route-map to match the traffic in ACL `ssh-inbound`:

```
 #(config-route-map) match ip address ssh-inbound
```

- b. Set an action on the matching traffic. Define a next-hop entry to redirect the traffic to the SPS.

```
 #(config-route-map) set ip next-hop <SPS IP address>
```

4. Create another route-map that controls the reply packet flow.

```
 #(config) route-map scb-outbound
```

```
 #(config-route-map) match ip address ssh-outbound
```

```
 #(config-route-map) set ip next-hop <SPS IP address>
```

NOTE:

This step is only required if you want to use Source NAT (IP Spoofing) instead of SPS's address towards the monitored servers.

Figure 5: Control > Connections — Using SNAT

5. Apply the route-map to the appropriate interfaces.
 - a. First, add the ssh-inbound route-map entry to the interface facing the clients:

```
 #(config) interface <interface-facing-the-clients>
 #(config-if) ip policy route-map scb-inbound
```
 - b. Then add the ssh-outbound route-map entry to the interface facing the servers:

```
 #(config) interface <interface-facing-the-servers>
 #(config-if) ip policy route-map scb-outbound
```

Expected result:

The traffic from the clients targeting the specified port of the servers is redirected to SPS. Therefore, SPS can be configured to control and audit this traffic.

The full configuration for the above topology:

```
! interface facing the clients
interface FastEthernet0/0
 ip address 192.168.0.254 255.255.255.0
 ip policy route-map scb-inbound
 duplex full
 speed auto
 no mop enabled

! interface facing the SCB
interface FastEthernet0/1
 ip address 172.16.0.254 255.255.255.0
 duplex full
 speed auto
 no mop enabled

! interface facing the servers
interface FastEthernet1/0
 ip address 10.0.0.254 255.255.255.0
 ip policy route-map scb-outbound
 duplex full
 speed auto
 no mop enabled

! access lists matching the server and client subnets and the SSH port -
incoming packets
ip access-list extended ssh-inbound
 permit tcp 192.168.0.0 0.0.0.255 10.0.0.0 0.0.0.255 eq 22
! access lists matching the server and client subnets and the SSH port - reply
packets
```

```

ip access-list extended ssh-outbound
 permit tcp 10.0.0.0 0.0.0.255 eq 22 192.168.0.0 0.0.0.255

! policy routing entry matching on the incoming SSH connections and
! redirecting them to the SCB external interface
route-map scb-inbound permit 10
 match ip address ssh-inbound
 set ip next-hop 172.16.0.1

! the following part is only required for SNAT-based SCB configuration
! policy routing entry matching on the SSH reply packets and
! redirecting them to the SCB external interface
route-map scb-outbound permit 10
 match ip address ssh-outbound
 set ip next-hop 172.16.0.1

```

Configuring advanced routing on Cisco ASA firewalls

The configuration of Cisco ASA firewalls follows the same rules as the Cisco router configuration, however the commands are slightly different.

⚠ CAUTION:

Source NAT (IP spoofing) is not supported in case of Cisco ASA firewalls. This means that with Cisco ASA, you cannot spoof the source IP towards the destination servers, therefore the source of the connections will be SCB's IP address.

Purpose:

To configure a Cisco ASA Firewall to redirect selected traffic to SCB instead of its original destination, complete the following steps. This procedure should work on most modern Cisco ASA software releases, but was specifically tested on Cisco Adaptive Security Appliance Software Version 9.6(2)3

Steps:

1. Define network objects that match the subnets or hosts that you want to monitor:

```

!Define SSH and RDP hosts/subnets as desired below
object network SSHHosts
 subnet <SSHHosts Subnet IP> <SSHHosts Subnet Netmask>
object-group network SSHtoSCB
network-object object SSHHosts
object network RDPHost

```

```
host <RDPHost IP>
object-group network RDPtoSCB
network-object object RDPHost
```

2. Create an ACL (Access Control List) entry that matches the objects above

```
!Allow RDP and SSH and their reply packets to SCB
access-list acl_pbr_ToSCB extended permit object rdp3389 any object-group
RDPtoSCB
access-list acl_pbr_ToSCB extended permit object rdp3389-response object-group
RDPtoSCB any
access-list acl_pbr_ToSCB extended permit object ssh22 any object-group SSHtoSCB
access-list acl_pbr_ToSCB extended permit object ssh22-response object-group
SSHtoSCB any
```

Keep in mind that whatever is permitted by this ACL is what will be matched, so make sure that the scope of the ACL entry is narrowed down as much as possible.

TIP:

Cisco ACLs use inverse netmasks for defining network addresses. To calculate an inverse mask given a subnet mask, simply subtract each octet value from 255.

3. Create a route-map entry. It controls which packets are affected by policy routing and where they should be forwarded to. The **match** commands specify the conditions under which policy routing occurs. The **set** commands specify the routing actions to perform if the criteria enforced by the **match** commands are met. A new route-map can be defined as follows:

```
!Define routing to SCB
route-map ToSCB permit
match ip address acl_pbr_ToSCB
set ip next-hop <SCB IP>
```

Apply the route-map to the appropriate interfaces.

```
!Set it on interface as needed
interface <interface-facing-to-the-servers>
ip policy route-map ToSCB
```

Expected result:

The traffic from the clients targeting the specified port of the servers is redirected to SPS. Therefore, SPS can be configured to control and audit this traffic.

The full configuration for the above topology:

```
!
!Define SSH and RDP hosts/subnets as desired below
object network SSHHosts
subnet <SSHHosts Subnet IP> <SSHHosts Subnet Netmask>
object-group network SSHtoSCB
network-object object SSHHosts
```



```

object network RDPHost
  host <RDPHost IP>
object-group network RDPtoSCB
  network-object object RDPHost
!
!Allow RDP and SSH and their reply packets to SCB
access-list acl_pbr_ToSCB extended permit object rdp3389 any object-group
RDPtoSCB
access-list acl_pbr_ToSCB extended permit object rdp3389-response object-group
RDPtoSCB any
access-list acl_pbr_ToSCB extended permit object ssh22 any object-group SSHtoSCB
access-list acl_pbr_ToSCB extended permit object ssh22-response object-group
SSHtoSCB any
!
!Define routing to SCB
route-map ToSCB permit
  match ip address acl_pbr_ToSCB
  set ip next-hop <SCB IP>
!
!Set it on interface as needed
interface <interface-facing-to-the-servers>
  ip policy route-map ToSCB

```

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 or other inquiries, visit <https://www.oneidentity.com/company/contact-us.aspx> or call +1-800-306-9329.

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