ONE IDENTITY

by Quest

One Identity Safeguard for Privileged Sessions 7.0.5 LTS

Getting Started with Safeguard for Privileged Sessions as a Virtual Appliance

Contents

Getting started with One Identity Safeguard for Privileged Sessions

Before you start:

This guide will help you get started with One Identity Safeguard for Privileged Sessions (SPS). It will explain the basic concepts and features of the product and walk you through a simple control and monitor setup for SSH and RDP connections. It will give you pointers to the relevant sections of the One Identity Safeguard for Privileged Sessions - Administration Guide to help you learn more about the more advanced options.

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

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

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

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

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

Central policy enforcement

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

Prevention of malicious activities

SPS monitors privileged user sessions in real-time and detects policy violations as they occur. In case of detecting a suspicious user activity (for example entering a destructive

command, such as the "rm"), SPS can send you an alert or immediately terminate the connection.

Greater accountability (deterrance)

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

Faster, cost-effective compliance audits

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

Lower troubleshooting and forensics costs

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

Supported virtual environments for quick starting One Identity Safeguard for Privileged Sessions

To start using One Identity Safeguard for Privileged Sessions as a virtual appliance, you can download and install the latest SPS ISO file into a virtual machine. The following virtual environments are supported for evaluation:

- Kernel-based Virtual Machine (KVM)
- Microsoft Hyper-V
- VMware
- vSphere (VMware ESX)
- Azure Marketplace
- Amazon Web Services (AWS)

SPS may work in other virtual environments like VirtualBox as well, although these are officially not supported. You can obtain an quick start license and the ISO file using your support portal account.

Setting up SPS and the virtual environment

To start using SPS, first install it in a virtual machine.

vSphere

Follow the instructions provided in *One Identity Safeguard for Privileged Sessions VMware Installation Guide* in the *Installation Guide*.

VirtualBox

Follow the instructions provided in *One Identity Safeguard for Privileged Sessions VMware Installation Guide* in the *Installation Guide*.

Hyper-V

Follow the instructions provided in *One Identity Safeguard for Privileged Sessions Hyper-V Installation Guide* in the *Installation Guide*.

Kernel-based Virtual Machine (KVM)

Follow the instructions provided in *Installing One Identity Safeguard for Privileged Sessions as a Kernel-based Virtual Machine* in the *Installation Guide*.

Azure Marketplace

Follow the instructions provided in *Azure deployment* in the *Administration Guide*.

Amazon Web Services (AWS)

Follow the instructions provided in *AWS deployment* in the *Administration Guide*.

Setting up a basic monitoring and controlling configuration for SSH and RDP

The rest of the guide will walk you through the steps required after the installation to get you started. You will learn how to do the following in the next sections:

1. Review and customize your network settings.

2. Configure a basic SSH Connection Policy.

3. Understand and customize the most important options around authentication and destination selection.

4. Understand and test real-time activity monitoring and Four Eyes Authorization.

5. Configure a basic RDP Connection Policy.

6. Understand and configure screen content indexing.

7. Understand and test the main functions of the search interface for audit recordings.

Initial setup

To start the setup of SPS

1. Connect to SPS.

   The SPS virtual machine acquires an IP address from your DHCP server accessible in the virtual environment. After SPS has booted up, the console displays the IP address of the SPS web interface at login prompt. To connect to SPS, use this IP address. For details, or tips if SPS cannot receive an IP address, see *The initial connection to One Identity Safeguard for Privileged Sessions (SPS)* in the *Administration Guide*.

2. Complete the Welcome Wizard as described in *Configuring One Identity Safeguard for Privileged Sessions (SPS) with the Welcome Wizard* in the *Administration Guide*. Upload the evaluation license file you have downloaded with your support portal account.

General connection settings

SPS supports transparent and non-transparent proxy operation modes to make deployments in existing network infrastructures as easy as possible. SPS will automatically handle non-transparent and transparent connections simultaneously.

Modes of operation

The following operation modes are possible:

- *Non-transparent proxy operation*: This guide will focus on this operation mode.
- *Transparent mode*: If you configure SPS proxies in transparent mode, the client usually addresses the target server directly. Therefore, you have to configure the connection policies in SPS accordingly. For more information, see *Transparent mode* in the *Administration Guide*.
- *Single-interface transparent mode*: For more information, see *Single-interface transparent mode* in the *Administration Guide*.

Non-transparent proxy operation

This guide focuses on non-transparent proxy operation, which is the easiest to implement. In this configuration, clients connect to a server through SPS. That is, end-users address SPS explicitly, which then forwards connections to target systems based on various parameters depending on what destination selection method you select.

Figure 1: SPS in non-transparent mode
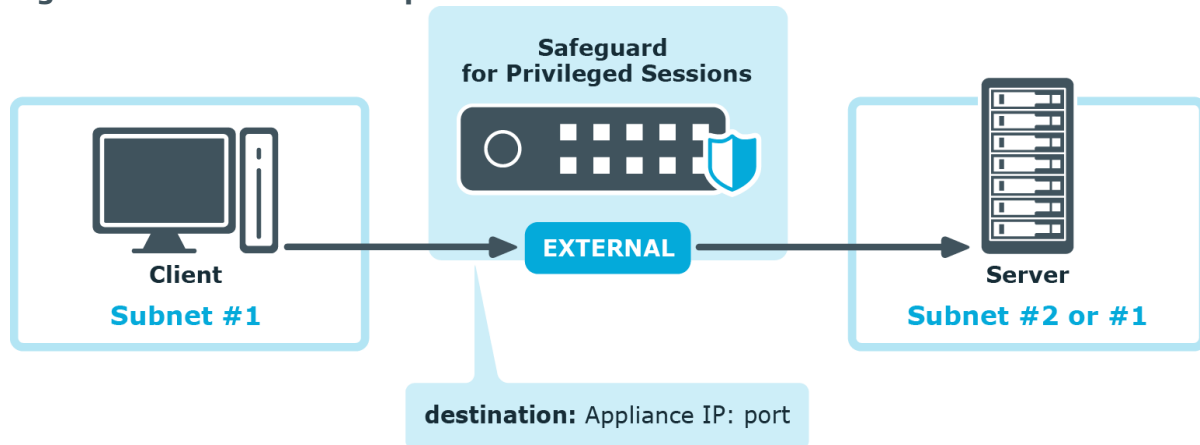
For an illustration of what happens when a client connects a server through SPS and how the different configuration options and policies of SPS affect this process, see:

- *Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using SSH* in the *Administration Guide*

- *Connecting to a server through One Identity Safeguard for Privileged Sessions (SPS) using RDP* in the *Administration Guide*

Network settings

Assigning logical interfaces to physical interfaces

To assign logical interfaces to the three physical interfaces of SPS, navigate to **Basic Settings > Network > Interfaces**.

Each logical interface must have its own VLAN ID, and can have its own set of (alias) IP addresses and prefixes. The configured name for each logical interface is visible on SPS's user interface only.

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

For details, see *Network settings* in the *Administration Guide*.

Routing uncontrolled traffic

To control how SPS routes uncontrolled traffic (that is, traffic that passes SPS but is not inspected or audited) between its network interfaces, navigate to **Basic Settings > Network > IP forwarding**.

You can connect interface pairs to each other, and SPS will route all uncontrolled traffic

between these. To add a new forwarding rule, choose ✚ and select the two logical

interfaces to connect. You can select the same interface in both fields to use that logical interface in single-interface router mode.

For details, see *Routing uncontrolled traffic between logical interfaces* in the *Administration Guide*.

Configuring connections: SSH

The following procedures provide a skeleton to configure SSH connections in SPS. If you want to have a deeper understanding, read the in-depth detailed procedure in *Configuring connections* in the *Administration Guide*.

Configure an SSH connection with fixed destination IP

The following describes how to configure a basic Secure Shell (SSH) connection in SPS. This Connection Policy uses a fixed destination IP, that is, it receives connections on an IP address of SPS, and forwards them to a server explicitly set in the policy.

The destination address is the address of the server where the clients finally connect to. To modify the destination address of a connection, complete the following steps.

Prerequisites:

- A SPS appliance where you have already completed the Welcome Wizard.

- An SSH server that is running on a host that you can access from SPS. That is, SPS must be able to access the network of the SSH server (adjust any routing and firewall settings in your network to permit this connection). If you only want to do a quick test, you can install an SSH server on the host you are configuring SPS from.

To configure a basic SSH connection in SPS

1. Navigate to **SSH Control > Connections**.

2. Click ➕ to define a new connection and enter a name that will identify the connection (for example `admin_mainserver`).

   TIP: It is recommended to use descriptive names that give information about the connection, for example refer to the name of the accessible server, the allowed clients, and so on.

3. Enter the IP address of the client that will be permitted to access the server into the **From** field. Click ✚ to list additional clients.

   Enter the IP address that the clients will request into the **To** field. To test SPS the easiest is to use the IP address of SPS, meaning that the connection will be non-transparent. (To test transparent connections, you must place SPS into the network between the client and the server, or route the traffic that way.)

   Figure 2: Configuring fix destination selection

4.

5. The **Target** section allows you to configure Network Address Translation (NAT) on the server side of SPS. Destination NAT determines the target IP address of the server-side connection. You can set the destination address as required for your environment. For this example non-transparent connection, select **Use fixed address**.

6. Enter the IP address and port number of the server. SPS will connect all incoming client-side connections to this server. For example, to redirect the connections to your computer (if it is running an SSH server), enter the IP address of your computer.

   You can also enter a hostname instead of the IP address, and SPS automatically resolves the hostname to IP address. Note the following limitations:

   - SPS uses the Domain Name Servers set **Basic Settings > Network > Naming > Primary DNS server** and **Secondary DNS server** fields to resolve the hostnames.

   - Only IPv4 addresses are supported.

- If the Domain Name Server returns multiple IP addresses, SPS selects randomly from the list.

7. If the clients use a custom port to address the server instead of the default port used by the protocol, enter the port number that the clients will request into the **Port** field. Click ➕ to list additional port numbers. For details on organizing connections in non-transparent mode, see *Organizing connections in non-transparent mode* in the *Administration Guide*.

8. Click [Commit] to save the connection.

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

   TIP: To temporarily disable a connection, deselect the checkbox before the name of the connection.

9. Test the new configuration: try to initiate an SSH connection from the client (your computer) to the server.

10. After successfully connecting to the server, do something in the connection, for example, execute a simple command in SSH (for example, ls /tmp), then disconnect from the server.

11. Navigate to **Search** on the SPS web interface. Your sessions are displayed in the list of connections. Note that for the transparent connection, the client addresses the target server, while the non-transparent connection addresses SPS.

12. Click the [i] icon. A summary will be displayed about the connection.

Server-side (only) password authentication

The default authentication method for SSH connection policies is to let the target system check credentials as it would happen when users access the server directly without SPS in place.

If you want to configure a different authentication method, create an authentication policy.

Figure 3: Authentication policy

An authentication policy is a list of authentication methods that can be used in a connection. Connection definitions refer to an authentication policy to determine how the client can authenticate to the target server. Separate authentication methods can be used on the client and the server-side of the connection.

To create a new authentication policy, navigate to **SSH Control > Authentication Policies**.

For details, see *Authentication Policies* in the *Administration Guide*.

Permitting or denying access to SSH channels

For certain protocols, multiple channels are defined each of which is responsible for a specific functionality supported by the protocol. For example, the **Session Shell** channel is the traditional remote terminal session, while the **Session Exec** channel allows to execute a remote command (for example rsync without opening a session shell.

For details on the supported SSH channel types, see *Supported SSH channel types* in the *Administration Guide*.

SPS can permit/deny access to these functionalities based on various parameters of a connection (for example time of the day, username, and so on) to provide an additional level of access control and protection.

Figure 4: Controlling protocol channels

Access to sub-channels is controlled by channel policies. The default SSH channel policy allows session shell access only.

For details, see *Creating and editing channel policies* in the *Administration Guide*.

Configuring SCP and SFTP access in SSH

***To configure SCP and SFTP access in SSH***

1. Navigate to **SSH Control > Channel Policies** and click ✚ to create a new channel policy. Enter a name for the policy into the **Channel Policy** field (for example, `shell_and_backup`).

2. Click ✚ to add a new channel.

3. Select **Session Exec SCP** from the **Type** field.

4. Restrict the availability of the channel based on your preferences.

   For details, see *Creating and editing channel policies* in the *Administration Guide*.

5. To be able to extract the original file from the corresponding audit trail for further inspection, select the **Record audit trail** option to record the activities of the channel into audit trails.

6. (Optional) To also configure SFTP channel access, add a new channel and repeat the steps above, but this time, select **Session SFTP** from the **Type** field.

Authorizing and monitoring a connection personally in real-time

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

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

Figure 5: Four-eyes authorization

For details on four-eyes authorization, see *Four-eyes authorization* in the *Administration Guide*.

Configuring four-eyes authorization

To configure four-eyes authorization

1. To enforce four-eyes authorization, navigate to **SSH Control > Connections**.

2. Select the connection policy to modify. Navigate to **Access Control** and click ✚ .

3. Enter the name of the usergroup whose members are permitted to authorize the sessions of the connection policy into the **Authorizer** field.

4. Configure the parameters of four-eyes authorization. For details, see *Configuring four-eyes authorization* in the *Administration Guide*.

5. Navigate to **SSH Control > Channel Policies**, and select the channel policy used in the connection.

6. Enable the **4 eyes** option for the channels which should be accessed only using four-eyes authorization.

7. Click **Commit** .

Inband destination selection

Using fix destination selection has the disadvantage of requiring one connection policy per protected server, because policies are mapped to servers based on IP addresses or port numbers.

Inband destination selection allows you to create a single connection policy and allow end-users to access any server. by including the name of the target server in their username (for example `ssh username@targetserver:port@scb_address`). SPS can extract the address from the username and direct the connection to the target server.

Figure 6: Inband destination selection

The process looks like the following:

1. End-users specify the destination server as part of the username, for example in the format of `<username>@<server address>:<port>@<SPS address>`, where the server and SPS address can be either a hostname or an IP address.

2. SPS tokenizes the username and the server address to forward the connection to.

3. SPS forwards the connection to the server.

For details, see *Using inband destination selection in SSH connections* in the *Administration Guide*.

Configuring inband destination selection

The following describes how to configure a Connection Policy to extract the address of the server from the username.

To configure a Connection Policy to extract the address of the server from the username

1. Navigate to the Connection policy you want to modify, for example, to **SSH Control > Connections**.

2. Select **Inband destination selection**.

Figure 7: Configuring inband destination selection

3. Enter the addresses of the servers that the users are permitted to access into the **Targets** field.

4. If the clients can access only a specified port on the server, enter it into the **Port** field. If the **Port** is not set, the clients may access any port on the server.

5. If there are any servers that the users cannot target using inband destination selection, add them to the **Exceptions** field.

6. Click **Commit** .

Example: Initiating a connection

Once the connection policy is configured correctly, a sample connection initiation would look like the following:

```
$ ssh root@192.168.56.10@192.168.56.200
```

- `root` = server user
- `192.168.56.10` = target server
- `192.168.56.200` = IP address of SPS

Gateway authentication

When gateway authentication is required for a connection, the user must authenticate on SPS as well. This additional authentication can be performed on the SPS web interface, so it provides a protocol-independent, outband authentication method. That way the connections can be authenticated to the central authentication database (for example LDAP or RADIUS), even if the protocol itself does not support authentication databases. Also, connections using general usernames (for example `root`, `Administrator`, and so on) can be connected to real user accounts.

Figure 8: Gateway authentication

For details on gateway authentication, see *The gateway authentication process* in the *Administration Guide*.

Password-based gateway (local) + password-based server-side authentication from credential store

The goal of this scenario is to demonstrate an SSH connection in which end-users must authenticate themselves successfully with their own passwords against a local user database maintained on SPS and have a session opened to the requested destination with a different account without any further interaction (that is, have SPS complete the password-based login process).

To configure password-based gateway (local) + password-based server-side authentication from credential store

1. **Create a local user database:**

   Navigate to **Policies > Local User Databases** and create a local user database.

   For details, see *Creating a Local User Database* in the *Administration Guide*.

2. **Connect the local user database with a client-side gateway authentication policy:**

   Navigate to **SSH Control > Authentication Policies**. Create an authentication policy. Select **Authenticate the client to SPS using > Local**. Select **Password**. Configure the required settings.

   For details, see *Local client-side authentication* in the *Administration Guide*.

3. **Create a user list:**

   Navigate to **Policies > User lists** and create a user list.

   For details, see *Creating and editing user lists* in the *Administration Guide*.

4. **Create a usermapping policy:**

   Navigate to **Policies > Usermapping policies** and create a usermapping policy.

   For details, see *Configuring usermapping policies* in the *Administration Guide*.

5. Create a local or remote credential store with the server user and its password. SPS provides a plugin framework to integrate with other remote credential stores/password management systems. For more information, see *Configuring local Credential Stores* in the *Administration Guide*.

   For details, see *Using a custom Credential Store plugin to authenticate on the target hosts* in the *Administration Guide*.

6. Expected outcome:

   If all prerequisites are met, SPS is ready to perform inband gateway authentication in an SSH session, which together with inband destination selection could be performed with the following connection string by an end-user:

   Example: Inband gateway authentication and destination selection

   ```
   $ ssh gu=myusername@root@192.168.56.10@192.168.56.200
   ```

   - `gu=myusername` = gateway user (myusername)
   - `root` = server user
   - `192.168.56.10` = target server
   - `192.168.56.200` = IP address of SPS

Public key-based gateway + password-based server-side authentication from credential store

This scenario differs from the previous one only in the client-side authentication method. In this case, the end-user is authenticated with the public key method, and if all permissions are granted by SPS (for example usermapping is allowed), they get logged in automatically to the requested server with the requested server account without having to enter a password.

To configure this, upload a public key for the user in the applied local user database, and make sure that the private key is accessible for the client application (openSSH, PuTTY, and so on).

To configure public key-based gateway + password-based server-side authentication from credential store

1. **Create a local user database:**

   Navigate to **Policies > Local User Databases** and create a local user database.

   For details, see *Creating a Local User Database* in the *Administration Guide*.

2. **Connect the local user database with a client-side gateway authentication policy:**

   Navigate to **SSH Control > Authentication Policies**. Create an authentication policy. Select **Authenticate the client to SPS using > Local**. Select **Public key**. Configure the required settings.

   For details, see *Local client-side authentication* in the *Administration Guide*.

3. **Create a user list:**

   Navigate to **Policies > User lists** and create a user list.

   For details, see *Creating and editing user lists* in the *Administration Guide*.

4. **Create a usermapping policy:**

   Navigate to **Policies > Usermapping policies** and create a usermapping policy.

   For details, see *Configuring usermapping policies* in the *Administration Guide*.

5. Create a local or remote credential store with the server user and its password. SPS provides a plugin framework to integrate with other remote credential stores/password management systems. For more information, see *Configuring local Credential Stores* in the *Administration Guide*.

For details, see *Using a custom Credential Store plugin to authenticate on the target hosts* in the *Administration Guide*.

6. Expected outcome:

If all prerequisites are met, SPS is ready to perform inband gateway authentication in an SSH session, which together with inband destination selection could be performed with the following connection string by an end-user:

Example: Inband gateway authentication and destination selection

```
$ ssh gu=balabit@root@192.168.56.10@192.168.56.200
```

- gu=balabit = gateway user (balabit)
- root = server user
- 192.168.56.10 = target server
- 192.168.56.200 = IP address of SPS

Configuring connections: RDP

The following procedures will provide a skeleton of configuring RDP connections in SPS. If you want to have a deeper understanding, see the in-depth detailed procedure in *Configuring connections* in the *Administration Guide*.

Configure an RDP connection with fixed destination IP

The following describes how to configure a basic Remote Desktop (RDP) connection in SPS. This Connection Policy uses a fixed destination IP, that is, it receives connections on an IP address of SPS (on the default RDP port `3389`), and forwards them to a server explicitly set in the policy.

The destination address is the address of the server where the clients finally connect to. To modify the destination address of a connection, complete the following steps.

Prerequisites:

- A SPS appliance where you have already completed the Welcome Wizard.
- A computer that accepts Remote Desktop connections (and RDP server). SPS must be able to access the network of the RDP server (adjust any routing and firewall settings in your network to permit this connection).

To configure a basic RDP connection in SPS

1. Navigate to **RDP Control > Connections**.

2. Click ![+] to define a new connection and enter a name that will identify the connection (for example `admin_mainserver`).

   TIP: It is recommended to use descriptive names that give information about the connection, for example refer to the name of the accessible server, the allowed clients, and so on.

3. Enter the IP address of the client that will be permitted to access the server into the **From** field. Click ✚ to list additional clients.

   Enter the IP address that the clients will request into the **To** field. To test SPS the easiest is to use the IP address of SPS, meaning that the connection will be non-transparent. (To test transparent connections, you must place SPS into the network between the client and the server, or route the traffic that way.)

Figure 9: Configuring fixed IP destination selection for RDP

4.
5. The **Target** section allows you to configure Network Address Translation (NAT) on the server side of SPS. Destination NAT determines the target IP address of the server-side connection. You can set the destination address as required for your environment. For this example non-transparent connection, select **Use fixed address**.

6. Enter the IP address and port number of the server. SPS will connect all incoming client-side connections to this server.

   You can also enter a hostname instead of the IP address, and SPS automatically resolves the hostname to IP address. Note the following limitations:

   - SPS uses the Domain Name Servers set **Basic Settings > Network > Naming > Primary DNS server** and **Secondary DNS server** fields to resolve the hostnames.

   - Only IPv4 addresses are supported.

   - If the Domain Name Server returns multiple IP addresses, SPS selects randomly from the list.

7. Click [**Commit**] to save the connection.

   This connection allows any user from the client machine to connect to the specified server, but permits only Desktop sessions — other RDP channels like disk redirection are disabled.

8. Test the new configuration: try to initiate an RDP connection from the client (your computer) to the server.

9. After successfully connecting to the server, do something in the connection, then disconnect from the server.

10. Navigate to **Search** on the SPS web interface. Your sessions are displayed in the list of connections. Note that for the transparent connection, the client addresses the target server, while the non-transparent connection addresses SPS.

11. Click the icon. A summary will be displayed about the connection.

Inband destination selection with Remote Desktop Gateway

Non-transparent operation with inband destination selection in RDP is supported with the implementation of the Remote Desktop Gateway protocol. When it is enabled, end-users configure their MSTSC client to use SPS as an RDP proxy/gateway and keep specifying target server addresses on the **General** tab the way they are used to.

Figure 10: Inband destination selection via RDP with RD Gateway
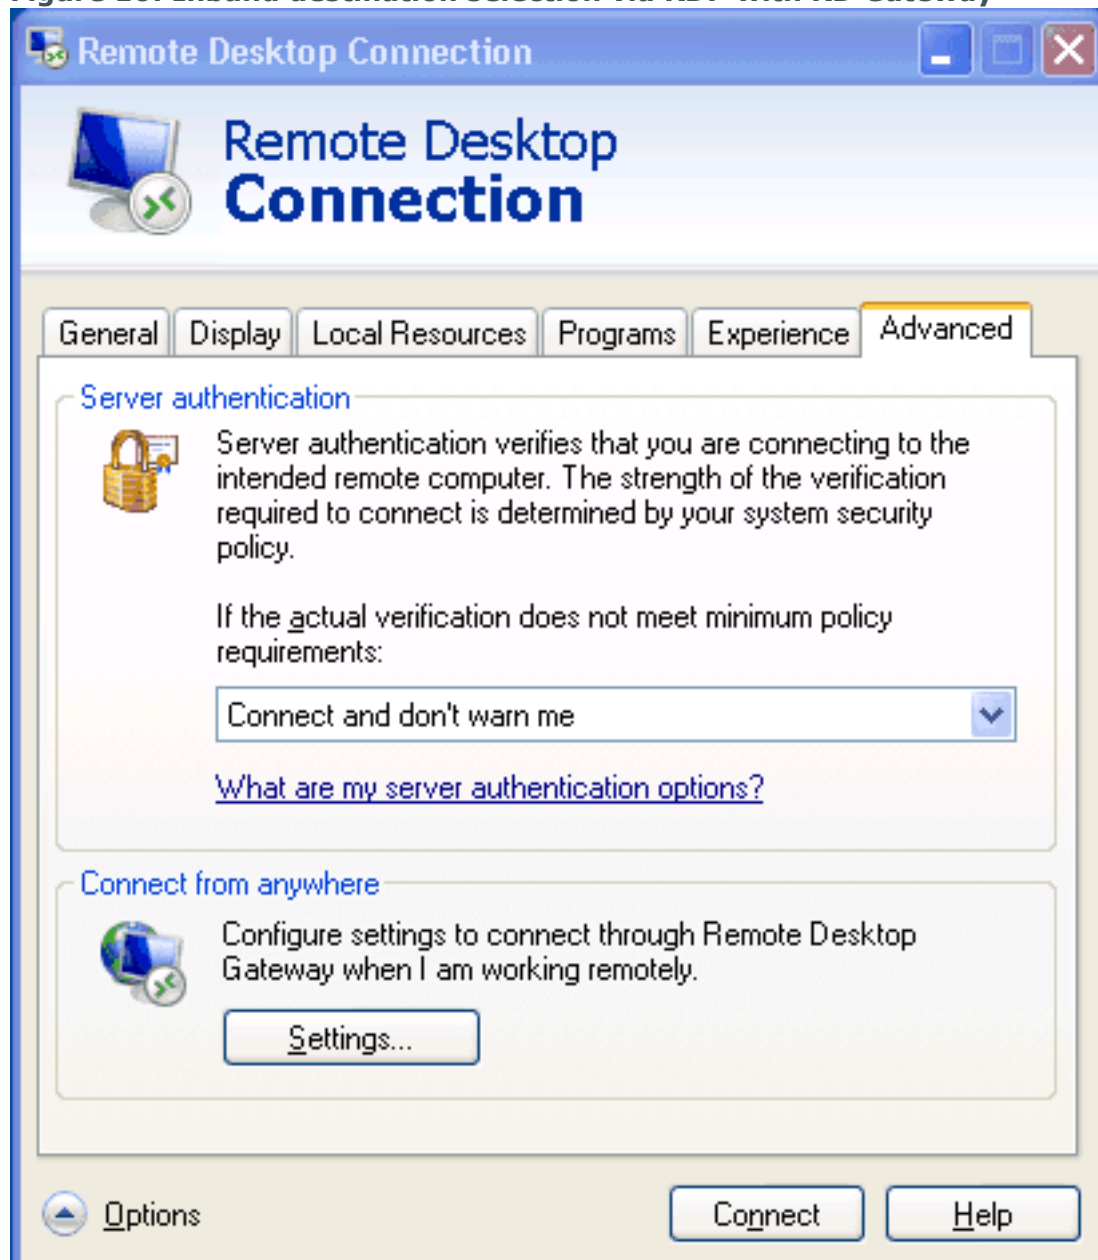
Configuring inband destination selection without RD Gateway

The following describes how to configure a Connection Policy to extract the address of the server from the username.

To configure a Connection Policy to extract the address of the server from the username

1. Navigate to the Connection policy you want to modify, for example, to **RDP Control > Connections**.

    Select **Inband destination selection**.

    Figure 11: Configuring inband destination selection for Windows connections

2.

3. Enter the addresses of the servers that the users are permitted to access into the **Targets** field.

4. If the clients can access only a specified port on the server, enter it into the **Port** field. If the **Port** is not set, the clients may access any port on the server.

5. If there are any servers that the users cannot target using inband destination selection, add them to the **Exceptions** field.

6. To use inband destination selection with RDP connections without using SPS as a Remote Desktop Gateway, you must use SSL-encrypted RDP connections.

   For details, see *Enabling TLS-encryption for RDP connections* in the *Administration Guide*.

7. Click Commit .

8. Start an RDP session from a Windows machine to SPS.

   Also, your users have the option to encode the address of the destination server in their username, in the username field of their client application. Note that SPS automatically displays a login screen if it cannot determine the username used in the connection, or you have not encoded a destination server in the username field. You can specify the destination address in the login screen when prompted.
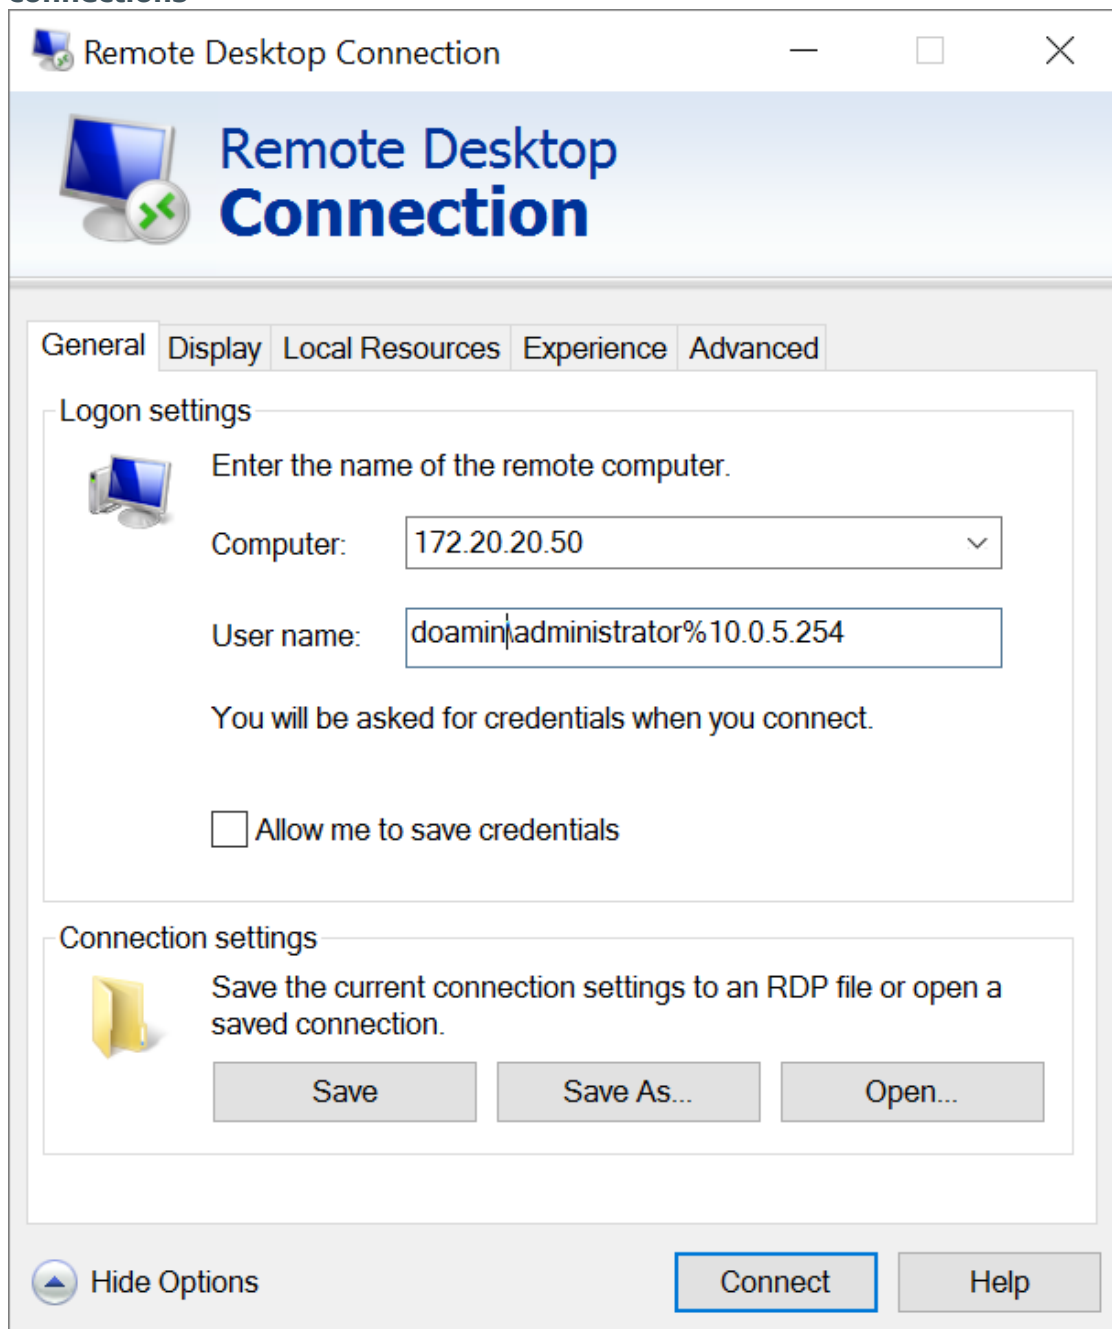
   When encoding the address of the destination server in the username, there are a few points to keep in mind. Since most RDP client applications limit which special characters can be used in usernames, this is not always intuitive.

   For the Microsoft Remote Desktop application (`mstsc`) and the login screen that SPS displays, note the following points:

   - Use the `%` character to separate the fields, for example: `username%my-targetserver`

   - Do not use the `@` character.

   - To specify the port number of the server (if it does not use the default port), use the caret `^` character, for example: `username%my-targetserver^6464`

   - To specify an IPv6 address, replace the colons with carets, and enclose the address in parentheses. For example, to target the `::1` IP address, use `username%(^^1)`. To target port 6464 of the same server, use `username%(^^1)^6464`.

   In the following example, a % symbol is passing the destination IP address to SPS, which redirects the connection to the proper client.

Figure 12: Configuring inband destination selection for Windows connections

Real-time content monitoring with Content Policies

You can monitor the traffic of certain connections in real time, and execute various actions if a certain pattern (for example, a particular command or text) appears in the command line or on the screen, or if a window with a particular title appears in a graphical protocol. Since content-monitoring is performed real-time, SPS can prevent harmful commands from being executed on your servers. SPS can also detect numbers that might be credit card numbers. The patterns to find can be defined as regular expressions. In case of ICA, RDP, and VNC connections, SPS can detect window title content.

The following channels support content policies:

- SSH Session shell (event type: Commands/Screen Content/Credit card)
- Telnet (event type: Commands/Screen Content/Credit card)
- RDP Drawing (event type: Window title detection)
- VNC (event type: Window title detection)
- ICA Drawing (event type: Window title detection)

For details, see *Real-time content monitoring with Content Policies* in the *Administration Guide*.

NOTE: Using content policies significantly slows down connections (approximately 5 times slower), and can also cause performance problems when using the indexer service.

The following describes how to create a new content policy that performs an action if a predefined content appears in a connection.

For details, see *Creating a new content policy* in the *Administration Guide*.

To create a new content policy that performs an action if a predefined content appears in a connection

1. Navigate to **Policies > Content Policies**, click ➕ and enter a name for the policy.
2. Select the **Event type** that you want to monitor.

3. Select **Match**, click and enter a string or regular expression. SPS will perform an action if this expression is found in the connection, unless it is listed in the **Ignore** list.

4. To add an exception to the **Match** rule, select **Ignore**, click and enter a string or regular expression.

5. Select the action to perform.

6. Click **Commit** .

7. To use the content policy created in the previous steps, select the policy in the channel policy that is used to control the connections.

Indexing service

One Identity Safeguard for Privileged Sessions (SPS) can index the contents of audit trails using its own indexer service or external indexers. Indexing extracts the text from the audit trails and segments it to tokens. A token is a segment of the text that does not contain whitespace: for example words, dates (2009-03-14), MAC or IP addresses, and so on. The indexer returns the extracted tokens to SPS, which builds a comprehensive index from the tokens of the processed audit trails.

Once indexed, the contents of the audit trails can be searched from the web interface. SPS can extract the commands typed and the texts seen by the user in terminal sessions, and text from graphical protocols like RDP, Citrix ICA, and VNC. Window titles are also detected.

SPS has an internal indexer, which runs on the SPS appliance. In addition to the internal indexer, external indexers can run on Linux hosts.

Processing and indexing audit trails requires significant computing resources. If you have to audit lots of connections, or have a large number of custom reports configured, consider using an external indexer to decrease the load on SPS. For sizing recommendations, ask your One Identity partner or contact our Support Team.

For details, see *Indexing audit trails* in the *Administration Guide*.

Using the content search

To most effectively search in the contents of the audit trails, make sure that the following prerequisites are met:

- Indexing was enabled in the connection policy related to the audit trail during the session, and
- the audit trail has already been indexed.

For details, see *Indexing audit trails* in the *Administration Guide*.

Configuring the internal indexer

Indexing audit trails allows you to search in the content of the audit trails, for example, to search for specific texts that the user has seen or typed in the session. The following describes how to configure SPS to index the audit trails. For details, see *Configuring the internal indexer* in the *Administration Guide*.

To configure SPS to index the audit trails

1. Navigate to **Basic Settings > Local Services > Indexer service**, and select **Indexer service**.

2. Define the **Maximum parallel audit trails to index on box**. The default value is set to the number of detected CPU cores.

3. (Optional) If you have encrypted audit trails and you want to index them, upload the necessary RSA keys (in PEM-encoded X.509 certificates).

4. Click **Commit**.

5. Navigate to **Policies > Indexer Policies**.

6. To create a new Indexer Policy, click **+**.

7. To configure what languages to detect, select **Manual language selection**. Select the language(s) to detect.

8. Navigate to the Control page of the traffic type (for example **SSH Control**), and select the connection policy to index.

9. Select **Enable indexing**.

10. To determine the priority level of indexing this connection, select the appropriate **Priority** level.

11. Select the **Indexing Policy** to be used.

12. Click **Commit**.

13. Check which channel policy is used in the connection, and navigate to the **Connection policies** page.

14. Select the channel policy used in the connection to index, and verify that the **Record audit trail** option is selected for the channels you want to index (for example, the Session shell channel in SSH, or the Drawing channel in RDP).

15. Click **Commit**.

16. Test the new configuration: try to initiate a connection from the client (your computer) to the server.

17. After successfully connecting to the server, do something in the connection, for example, execute a simple command in SSH (for example, `ls /tmp`), or launch an

application in RDP (for example, the Windows Explorer), then disconnect from the server.

18. Navigate to **Search** on the SPS web interface. Your sessions are displayed in the list of connections. Note that for the transparent connection, the client addresses the target server, while the non-transparent connection addresses SPS.

19. Click the icon. A summary will be displayed about the connection. Enter a text that was displayed in the connection into the search box, for example, the command you executed in SSH, or a menu item or other text you have seen in RDP (for example, Start). SPS will automatically generate a screenshot showing when the text was displayed in the connection.

Using the Search interface

This section provides an overview on how to use the Search interface. It describes how you can access the Search interface, lists the steps to take to search effectively, view the details of a connection, replay the audit trails, or export the search results as a comma-separated text file.

Prerequisites

Users need the **Search** privilege to access the Search interface.

NOTE: Assigning the **Search** privilege to a user on the **Users & Access Control > Appliance Access** page, automatically enables the **Search in all connections** privilege, and grants the user access to every audit trail, even if the user is not a member of the groups listed in the **Access Control** option of the particular connection policy.

If you want users to access audit trails only for connections for which they are granted permission, see *Assigning search privileges* in the *Administration Guide*.

For information on configuring:

- Authorizers for a connection, see *Configuring four-eyes authorization* in the *Administration Guide*.

- User rights, see *Managing user rights and usergroups* in the *Administration Guide*.

1. To access the Search interface, navigate to **Search**.

   Sessions are displayed sorted by date. For ongoing sessions, the Search interface is updated in real-time to always show the most up-to-date information.

   You can view sessions in a card, table or flow view. Click ••• for more details and select from the list.

2. Specify a date and time range to restrict your search criteria as described in Specifying time ranges on page 37.

3. Filter connections as described in Using search queries on page 41.

4. Search the contents of audit trails as described in *Searching in the contents of audit trails* in the *Administration Guide*.

5. View connection details as described in Viewing session details on page 44.

6. Download and replay audit trails as described in *Replaying audit trails in your browser* in the *Administration Guide*.

To export the search results as a comma-separated text file, click ••• for more details and select **Export CSV**. Note that if your search returns more than 10.000 results, only the first 10.000 rows are exported. If you want to see all results, refine your search.

To customize which fields are exported, click ••• for more details and select **Export CSV ...**.

7.

Specifying time ranges

Specify a time range to restrict, or filter your search criteria by setting boundaries on your searches. You can restrict the search to one of the preset time ranges, or use a custom time range for a more specific search.

When you specify a time range, the search result includes:

- Connections started and finished anywhere between the start time and end time you specified.
- Connections started anywhere between the start time and end time you specified.
- Connections ended anywhere between the start time and end time you specified.
- Active connections if they were started anywhere between the start time and the end time you specified.

For example, at 17:00 PM you specify a start date of 10:00 AM and end date of 15:00 PM for your search. The search result includes:

- Connections started at `8:00` AM and ended at `14:00` PM.
- Connections started at `11:00` AM and ended at `14:00` PM.
- Connections started at `11:00` AM and ended at `16:00` PM.
- Active connections started at `11:00` AM.
- Active connections started at `10:00` AM.

To specify time ranges

1. To select the start date of your search, click **Pick a date**.

   Alternatively, use the 🗓 ∨ (shortcuts) button to restrict the search to one of the preset time ranges. For example, to investigate an incident that occurred sometime in the last hour, you can select **Today**, but a better option is **Last 60 minutes**.

   Figure 13: Search — Pick a date

2. From the calendar, select the start date as required.

   | NOTE: The date refers to the timezone configured on SPS.

3. For exact time ranges, specify to search by the hour and minute.

Figure 14: Search — Specify hour and minute

4. To select the end date of your search, click **Pick a date** and select a date as required.

   If you specify only the start date, the end date is set to the current time.

5. *Optional:* To clear the start and end date, click 📅 ⌄ (shortcuts) > **All time**.

6. *Optional:* You can use the timeline for a quick time range selection and visual representation of sessions in the selected interval.

Click the icon.

Figure 15: Search — Using the timeline
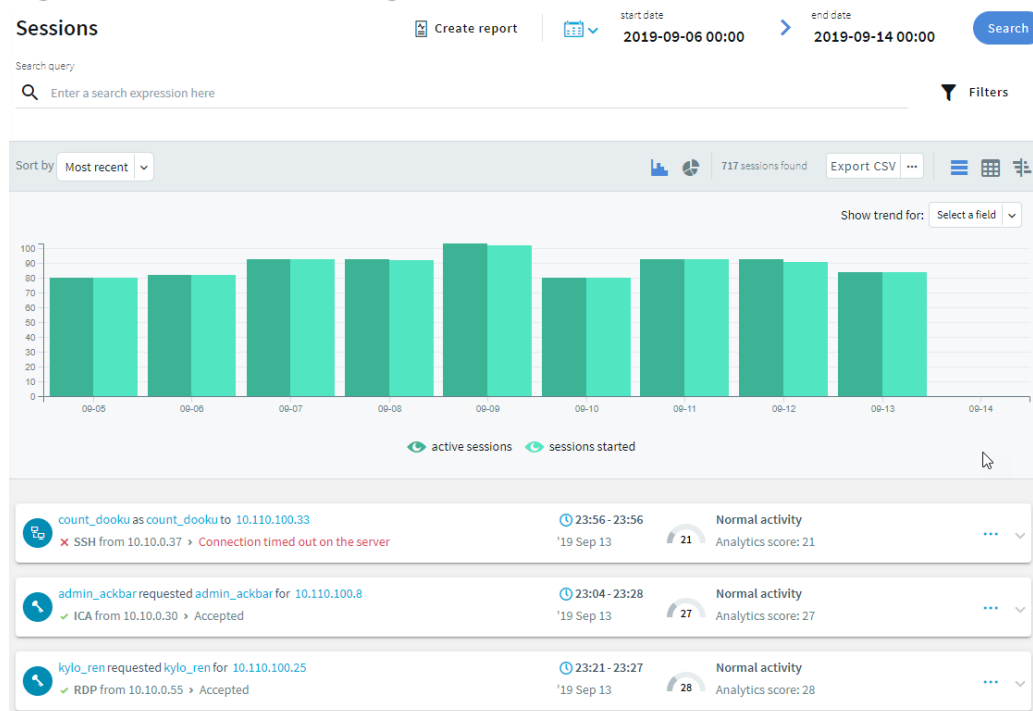
a.

The bars display the number of results in the selected interval.

The **active sessions** columns indicate all the sessions, which were active in the selected interval. The **sessions started** columns indicate all the sessions started during the selected interval. For example, if the selected interval is today between 8:00 AM and 9:00 AM, then a session started at 7:00 AM but lasting after 8:00 AM is displayed in the **active sessions** column. A session started at 8:30 AM is displayed in the **sessions started** column. Since the session was active during the selected time interval, the session started at 8:30 AM is also displayed in the **active sessions** column.

To disable the active sessions and view only the started sessions in the timeline, click . To disable the started sessions and view only the active sessions in the timeline, click .

Hovering the mouse above a bar displays the number of entries and the start and end date of the period that the bar represents.

Trend analysis allows you to use the timeline to find changes over time. For example, to find the time range where terminated connections had a significant peak compared to other days, from the **Show trend for** drop-down menu, select **Verdict**. Note that you can only view trend analysis for Active, Analytics

Score, Client name, Protocol, Server hostname, Server port, Server username, Username and Verdict. All the other selections are grayed out.

The colors of the bars in the timeline allow you to quickly find the time range with a higher number of terminated sessions.

Optional: To clear the trend analysis view, from the **Show trend for** drop-down menu, select **X**.

Figure 16: Search — Using the timeline - trend analysis

b. To select a range, drag the mouse pointer across the timeline or use Shift+Click and select multiple bars.

Using search queries

This section describes how you can use search queries to perform a more specific search.

To search using search queries

1. Enter a search query in the **Search query** field, or click on an entry in the table.

   To search, enter a valid search field followed by a value in the **search field: VALUE** format. For example, if you enter **protocol: SSH**, the search returns all the SSH sessions.

   TIP: Search is case insensitive. To make the search case sensitive, enclose the search keywords in double quotes.

   The search queries can include only alphanumerical characters. You can use complex expressions and boolean operators, for example, AND, OR, <,>, and so on.

   For the list of search fields that you can use, see *List of available search queries* in the *Administration Guide*.

   For more information on how to use more complex keyphrases that are not covered in this guide, see the Apache Lucene documentation.

There are search fields that are not displayed but you can still use them to query the sessions. For example, you can search for active connections using the `active` search field, and search results are listed accordingly, but there is no **active** field displayed in the search table or in the **Overview**, **Details**, and **Timeline** tabs.

Figure 17: Search — Search queries

Alternatively, click and set the filters you need from the appropriate columns. For example, to search for a specific username, select it using the drop-down menu of the **Username** column. For a more generic search, you can enter any text in the **Contains text** column.

Figure 18: Search — Search filters - Basic view

2. After specifying the relevant query, click **Search** or press **Enter**.

   TIP: To save the queries for future use, simply save the URL or bookmark it in your browser.

Expected result

Session metadata is displayed in columns that you can query for any parameter, or a combination of parameters. You can view the metadata in the search columns and also displayed as fields in the **Overview**, **Details**, and **Timeline** tabs.

Displaying statistics on search results

You can quickly sort and visualize the distribution of the sessions based on their various metadata, for example, username, server address, and so on.

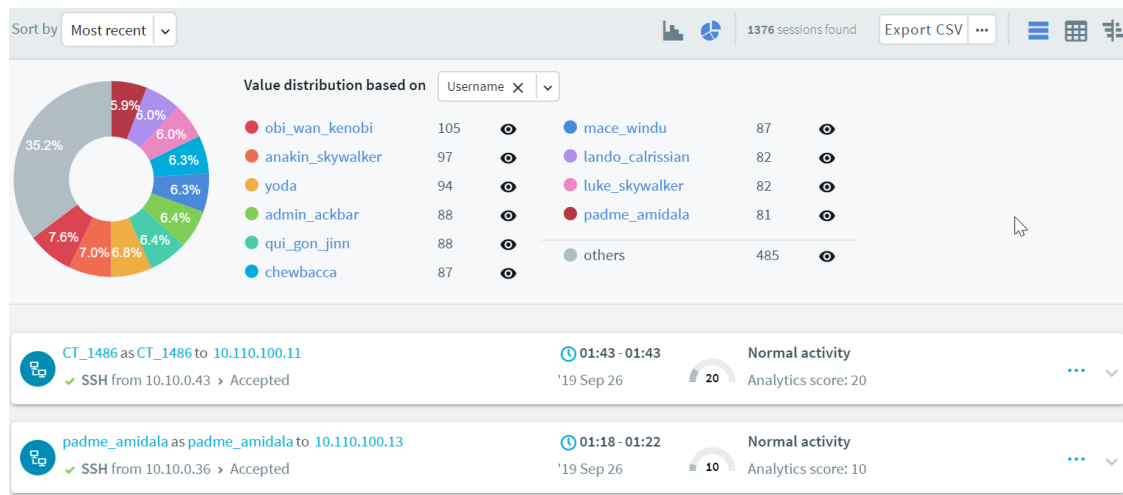***To display statistics on search results***

1. Click the icon.

2. Select the type of metadata you want to create statistics on from the **Value distribution based on** field, for example, select **Username** to display sessions based on username.

Figure 19: Search — Displaying statistics

3. To exclude items from the pie chart, click the ⦿ icon next to the metadata you want to exclude.
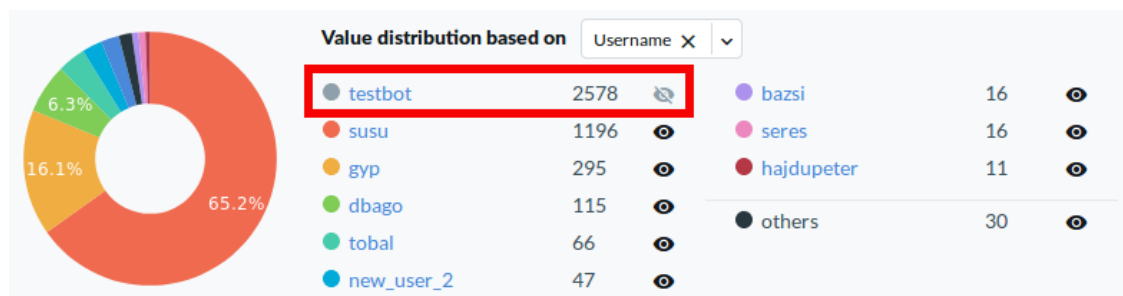
   For example, if you want to exclude results by a user called **testbot**, select the ⦿ icon next to the item.

Figure 20: Search — Excluding items from the pie chart

The pie chart now does not display results for the excluded item. The percentages always add up to 100%.

You can continue to restrict or refine your search results and view statistics as required.

Viewing session details

View the session details of each session for in-depth information on each of the indexed session stored in the connection database. You can use it to gain contextual insight about the indexed session and its events.

You can view session details for data recorded by:

- One Identity Safeguard for Privileged Sessions (SPS). For more information, see *Viewing session details for data recorded by SPS* in the *Administration Guide*.

- One Identity Safeguard for Privileged Passwords (SPP). For more information, see *Viewing session details for data recorded by SPP* in the *Administration Guide*.

Frequent Item Set (FIS) flow view visuals

Frequent Item Set (FIS) flow view visuals are also available on the **Analytics** tab. The FIS flow view is similar to the flow view analytics overview, except that the FIS flow view only displays data narrowed down to a single user's previous sessions in the analysis period (which is the previous 90 days by default). For more information, see *Visualizing Frequent Item Sets on the FIS flow view* in the *Administration Guide*.

Session tags

Session tags allow you to get basic information about the session and its contents at a glance.

Scripted session tag: One Identity Safeguard for Privileged Sessions (SPS) currently supports the scripted session tag. SPS uses One Identity Safeguard for Privileged Analytics to detect if sessions are generated using human interaction or automation. If sessions are generated using automation, SPS displays the scripted tag in the search interface as shown below:

- Scripted sessions are shown on the main search screen.

  Figure 21: Scripted sessions — cards view

- Scripted sessions are shown on the **Overview** tab.

Figure 22: Scripted sessions — Overview tab

SPS 7.0.5 LTS Getting Started with Safeguard for Privileged
Sessions as a Virtual Appliance
Using the Search interface

45

About us

One Identity solutions eliminate the complexities and time-consuming processes often required to govern identities, manage privileged accounts and control access. Our solutions enhance business agility while addressing your IAM challenges with on-premises, cloud and hybrid environments.

Contacting us

For sales and other inquiries, such as licensing, support, and renewals, visit https://www.oneidentity.com/company/contact-us.aspx.

Technical support resources

Technical support is available to One Identity customers with a valid maintenance contract and customers who have trial versions. You can access the Support Portal at https://support.oneidentity.com/.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos at www.YouTube.com/OneIdentity
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product