Engineering
Simplicity

# Juniper Secure Analytics Administration Guide

Published
2023-04-20

RELEASE
7.5.0

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Juniper Secure Analytics Administration Guide*
7.5.0

The information in this document is current as of the date on the title page.

## YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

## END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

# About This Guide

Use this guide to set up and manage JSA administrative functionality. Also learn about the new features and capabilities that make it easier for you to configure and administer your JSA deployment.

1

CHAPTER

# What's New for Administrators

# What's New for Administrators

Learn about the new features and capabilities that make it easier for you to configure and administer your JSA deployment.

RELATED DOCUMENTATION

# New Features and Enhancements in JSA 7.4.2

IN THIS SECTION

The following new features and enhancements make it easier for administrators to manage their JSA 7.4.2 deployment.

To view a list of all new features in this release, see What’s New Guide.

## Adjusting the Number of MAC Addresses Allowed for an Asset

In JSA 7.4.2, you can adjust the number of MAC addresses that are allowed for a single asset. In previous releases of JSA, administrators were not able to adjust this number, which resulted in an error message that stated that there were too many MAC addresses for the asset. Enter the number in the **Number of MAC Addresses Allowed for a Single Asset** field in the **Asset Profiler Configuration** window.

If you have users who log in from multiple wireless access points, or multiple users who log in remotely through a VPN, you can set the number of MAC addresses that are allowed for the asset in the same way that you can for IP addresses.

**Figure 1: Asset Profiler Configuration Window**

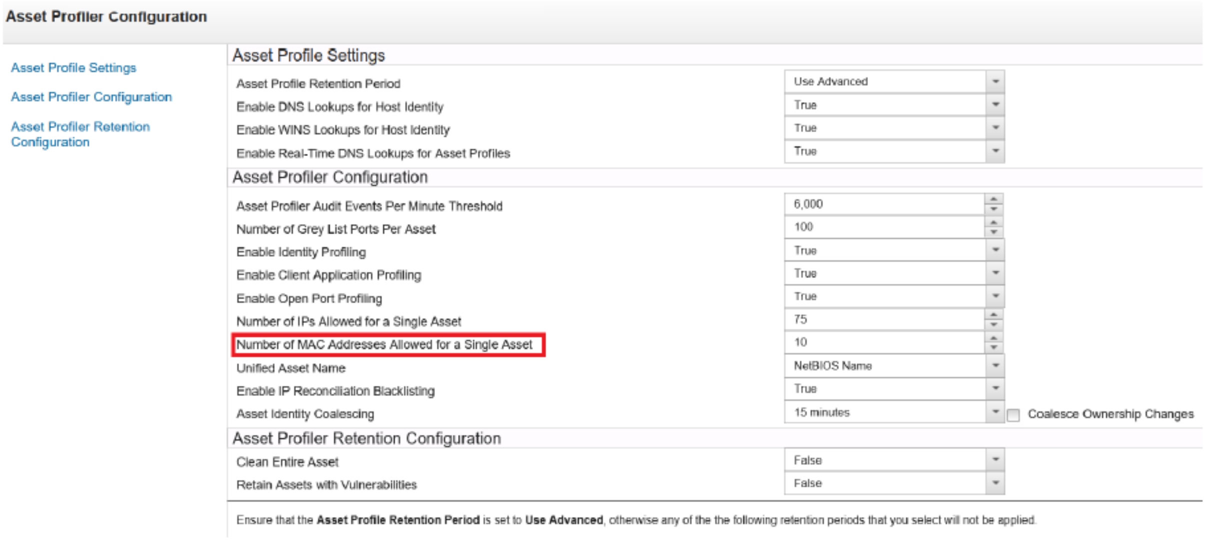

Asset Profiler Configuration
Asset Profile Settings
Asset Profiler Configuration
Asset Profiler Retention Configuration

Asset Profile Settings

| | |
|---|---|
| Asset Profile Retention Period | Use Advanced |
| Enable DNS Lookups for Host Identity | True |
| Enable WINS Lookups for Host Identity | True |
| Enable Real-Time DNS Lookups for Asset Profiles | True |

Asset Profiler Configuration

| | | |
|---|---|---|
| Asset Profiler Audit Events Per Minute Threshold | 6,000 | |
| Number of Grey List Ports Per Asset | 100 | |
| Enable Identity Profiling | True | |
| Enable Client Application Profiling | True | |
| Enable Open Port Profiling | True | |
| Number of IPs Allowed for a Single Asset | 75 | |
| Number of MAC Addresses Allowed for a Single Asset | 10 | |
| Unified Asset Name | NetBIOS Name | |
| Enable IP Reconciliation Blacklisting | True | |
| Asset Identity Coalescing | 15 minutes | □ Coalesce Ownership Changes |

Asset Profiler Retention Configuration

| | |
|---|---|
| Clean Entire Asset | False |
| Retain Assets with Vulnerabilities | False |

Ensure that the **Asset Profile Retention Period** is set to **Use Advanced**, otherwise any of the the following retention periods that you select will not be applied.

## Generating Regex for Parsing Event Properties

JSA 7.4.2 can suggest regular expressions (regex) when you enter event data in the **Workspace**. If you are not familiar with creating regex expressions, use this feature to generate your regex.

Highlight the payload text that you want to capture and in the **Properties** tab, click **Suggest Regex**. The suggested expression appears in the **Expression** field. Alternatively, you can click the **Regex** button in the **Workspace** and select the property that you want to write an expression for. If JSA is unable to generate a suitable regex for your data sample, a system message appears.

> **TIP**: The regex generator works best for fields in well-structured event payloads. If your payload consists of complex data from natural language or unstructured events, the regex generator might not be able to parse it and does not return a result.

The following figure shows how you can generate your regex with the **Suggest Regex** button in the **Properties** tab, or with the **Regex** button in the **Workspace**.

Figure 2: Suggest Regex Button

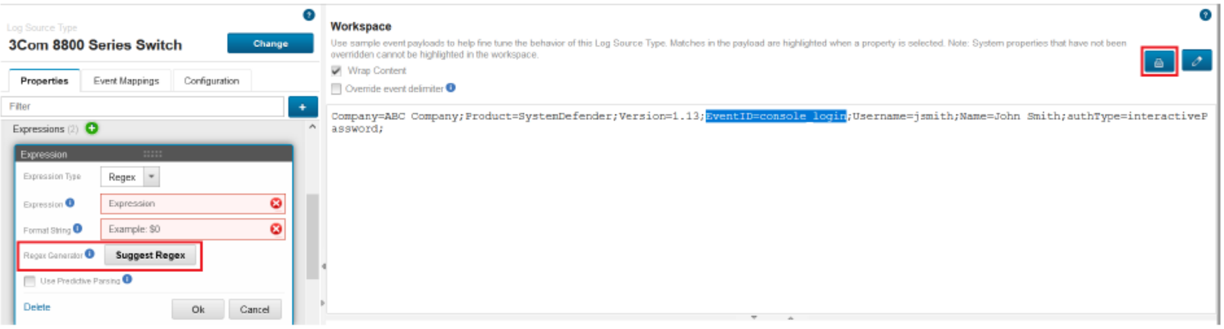

Expression Type: Regex

Expression: Expression

Format String: Example: $0

Regex Generator: Suggest Regex

□ Use Predictive Parsing

Use sample event payloads to help fine tune the behavior of this Log Source Type. Matches in the payload are highlighted when a property is selected. Note: System properties that have not been overridden cannot be highlighted in the workspace.

☑ Wrap Content

□ Override event delimiter

```
Company=ABC Company;Product=SystemDefender;Version=1.13;EventID=console login;Username=jsmith;Name=John Smith;authType=interactiveP
assword;
```

RELATED DOCUMENTATION

What's New for Administrators | 2

New Features and Enhancements in JSA 7.4.1 | 4

New Features and Enhancements in JSA 7.4.0 | 5

# New Features and Enhancements in JSA 7.4.1

The following new features and enhancements make it easier for administrators to manage their JSA 7.4.1 deployment.

To view a list of all new features in this release, see What’s New Guide.

RELATED DOCUMENTATION

What's New for Administrators | 2

New Features and Enhancements in JSA 7.4.2 | 2

New Features and Enhancements in JSA 7.4.0 | 5

# New Features and Enhancements in JSA 7.4.0

IN THIS SECTION

The following new features and enhancements make it easier for administrators to manage their JSA 7.4.0 deployment.

To view a list of all new features in this release, see What’s New Guide.

## Global System Notifications configuration

Global System Notifications are now local, making them host specific and more useful. As a result, the thresholds are now set automatically by JSA and the Global System Notification section of the Admin tab was removed.

## Secure email server

Send email to distribute alerts, reports, notifications, and event messages to mail servers that require authentication.

You can configure an email server for your entire JSA deployment, or multiple email servers.

## DSM Parameter support in the DSM Editor

In JSA 7.4.0, if your log source type has DSM parameters, you can use the DSM Editor to configure the DSM parameters. Enable the **Display DSM Parameters Configuration** option to view and edit the DSM parameters.

Figure 3: DSM Parameters Configuration

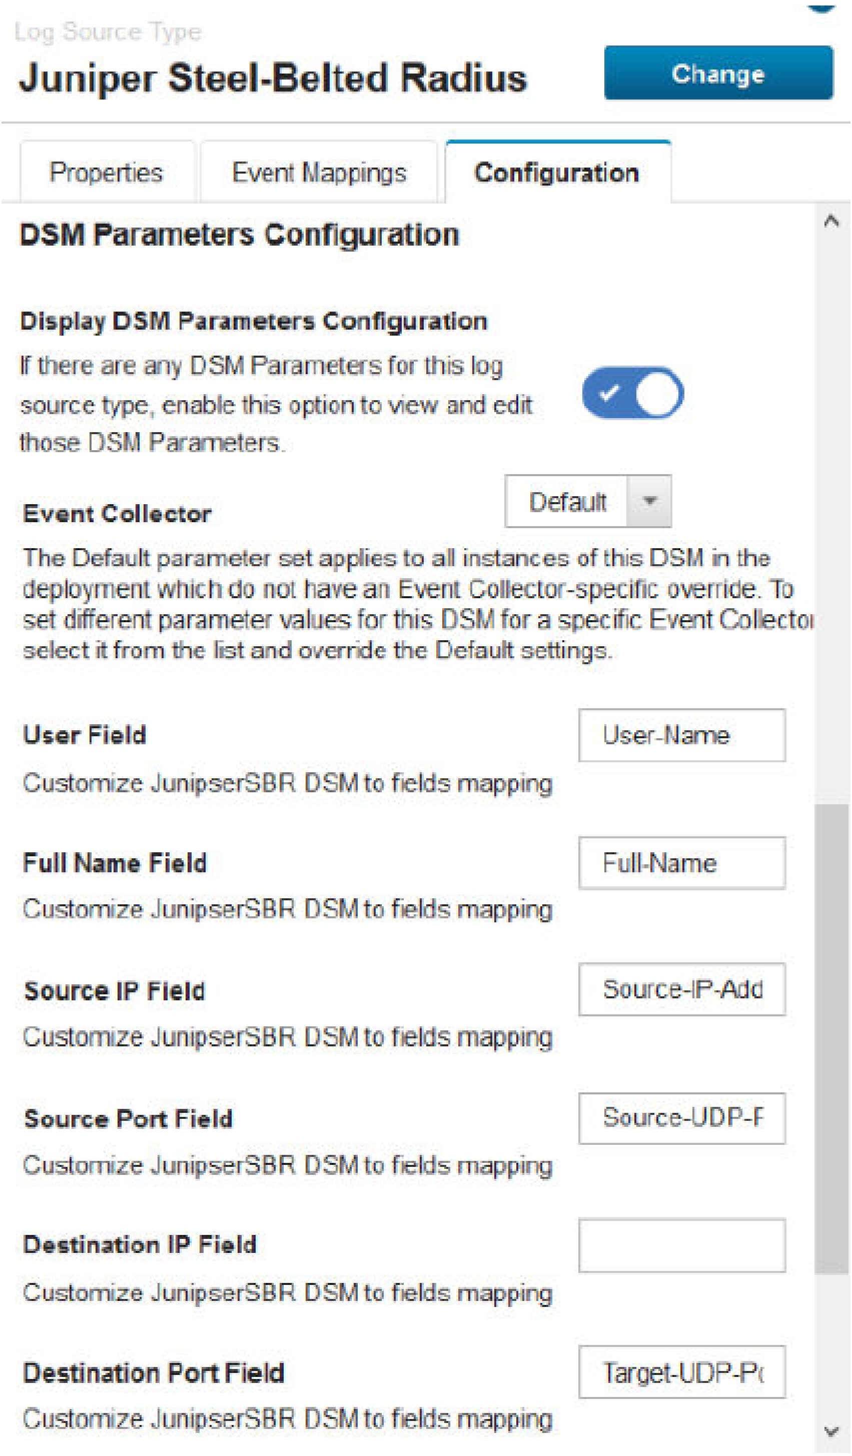

Log Source Type

Juniper Steel-Belted Radius

DSM Parameters Configuration

Display DSM Parameters Configuration

If there are any DSM Parameters for this log source type, enable this option to view and edit those DSM Parameters. ✓

Event Collector: Default

The Default parameter set applies to all instances of this DSM in the deployment which do not have an Event Collector-specific override. To set different parameter values for this DSM for a specific Event Collectoı select it from the list and override the Default settings.

User Field: User-Name

Customize JunipserSBR DSM to fields mapping

Full Name Field: Full-Name

Customize JunipserSBR DSM to fields mapping

Source IP Field: Source-IP-Add

Customize JunipserSBR DSM to fields mapping

Source Port Field: Source-UDP-F

Customize JunipserSBR DSM to fields mapping

Destination IP Field:

Customize JunipserSBR DSM to fields mapping

Destination Port Field: Target-UDP-P(

Customize JunipserSBR DSM to fields mapping

## Reverse tunnel initiation

The SSH tunnel between two managed hosts can now be initiated from the remote host instead of the local host. For example, you have a connection from an Event Processor in a secure environment to an Event Collector that is outside of the secure environment. You also have a firewall rule that prevents you from having a host outside the secure environment connect to a host in the secure environment. In JSA 7.4.0, you can switch which host creates the tunnel so that the connection is established from the Event Processor by selecting the Remote Tunnel Initiation checkbox for the Event Collector.

## Improved flow timestamp handling

Two new configuration settings provide more control over the way that flow timestamps are handled when Netflow V9 begins sending records with overflowed system uptime values. The new settings eliminate the need to reset the first and last switched times.

The new configuration options and the default values are shown here:

- NORMALISE_OVERFLOWED_UPTIMES=YES
- UPTIME_OVERFLOW_THRESHOLD_MSEC=86400000

The timestamps are corrected when the system uptime value is less than the first and last switched packet times by more than the value that is specified in the UPTIME_OVERFLOW_THRESHOLD_MSEC configuration. The timestamps are corrected based on the assumption that the system uptime wrapped around the maximum 32-bit value.

RELATED DOCUMENTATION

What's New for Administrators | 2

New Features and Enhancements in JSA 7.4.1 | 4

2

CHAPTER

# Overview of JSA Administration

# JSA Administration

As a JSA administrator, you have a variety of tools available to help you configure and manage your JSA deployment.

For example, using the tools on the **Admin** tab, you can perform the following tasks:

- Deploy and manage JSA hosts and licenses.
- Configure user accounts and authentication.
- Build a network hierarchy.
- Configure domains and set up a multi-tenant environment.
- Define and manage log and flow data sources.
- Manage JSA data retention.
- Manage assets and reference data.
- Schedule regular backups of JSA configuration and data.
- Monitor the system health of managed hosts.

RELATED DOCUMENTATION

Capabilities in Your JSA Product | 10

Supported Web Browsers | 13

# Capabilities in Your JSA Product

IN THIS SECTION

JSA product documentation describes functionality such as offenses, flows, assets, and historical correlation, that might not be available in all JSA products. Depending on the product that you are using, some documented features might not be available in your deployment.

- **Log Manager--** Log Manager is a basic, high-performance, and scalable solution for collecting, analyzing, storing, and reporting on large volumes of network and security event logs.
- **JSA--** JSA is an advanced offering that includes the full range of security intelligence capabilities for on-premises deployments. It consolidates log source and network flow data from thousands of assets, devices, endpoints, and applications that are distributed throughout your network, and performs immediate normalization and correlation activities on the raw data to distinguish real threats from false positives.

## JSA Product Capabilities

Review the following table to compare the capabilities in each JSA product.

**Table 1: Comparison Of JSA Capabilities**

| Capability | JSA | Log Manager |
|---|---|---|
| Full administrative capabilities | Yes | Yes |
| Supports hosted deployments | No | No |
| Customizable dashboards | Yes | Yes |
| Custom rules engine | Yes | Yes |
| Manage network and security events | Yes | Yes |
| Manage host and application logs | Yes | Yes |
| Threshold-based alerts | Yes | Yes |
| Compliance templates | Yes | Yes |

**Table 1: Comparison Of JSA Capabilities *(Continued)***

| Capability | JSA | Log Manager |
|---|---|---|
| Data archiving | Yes | Yes |
| Juniper X-Force Threat Intelligence IP reputation feed integration | Yes | Yes |
| WinCollect stand-alone deployments | Yes | Yes |
| WinCollect managed deployments | Yes | Yes |
| JSA Vulnerability Manager integration | Yes | Yes |
| Network activity monitoring | Yes | Yes |
| Asset profiling | Yes | No [1] |
| Offenses management | Yes | No |
| Network flow capture and analysis | Yes | Yes |
| Historical correlation | Yes | No |
| JSA Risk Manager integration | Yes | No |
| Vulnerability assessment scanners | Yes | Yes |

[1] Log Manager tracks asset data only if JSA Vulnerability Manager is installed.

Some documentation, such as the *Juniper Secure Analytics Administration Guide* and the *Juniper Secure Analytics Users Guide*, is common across multiple products and might describe capabilities that are not available in your deployment.

RELATED DOCUMENTATION

Supported Web Browsers | 13

# Supported Web Browsers

For the features in JSA products to work properly, you must use a supported web browser.

The following table lists the supported versions of web browsers.

**Table 2: Supported Web Browsers for JSA Products**

| Web browser | Supported versions |
|---|---|
| 64-bit Mozilla Firefox | 60 Extended Support Release and later |
| 64-bit Microsoft Edge | 38.14393 and later |
| 64-bit Google Chrome | Latest |

The Microsoft Internet Explorer web browser is no longer supported as of JSA 7.4.0 or later.

### Security Exceptions and Certificates

If you are using the Mozilla Firefox web browser, you must add an exception to Mozilla Firefox to log in to JSA. For more information, see your Mozilla Firefox web browser documentation.

### Navigate the Web-Based Application

When you use JSA, use the navigation options available in the JSA Console instead of your web browser **Back** button.

RELATED DOCUMENTATION

Capabilities in Your JSA Product | 10

3

CHAPTER

# User Management

# User Management

You define user roles, security profiles, and user accounts to control who has access to JSA, which tasks they can perform, and which data they have access to.

When you initially configure JSA, use the **User Management** feature on the **Admin** tab to configure and manage user accounts for all users that require access to JSA.

# User Roles

IN THIS SECTION

A user role defines the functions that a user can access in JSA.

During the installation, four default user roles are defined: **Admin**, All, WinCollect, and Disabled.

Before you add user accounts, you must create the user roles to meet the permission requirements of your users.

## Creating a User Role

Create user roles to manage the functions that a user can access in JSA.

By default, your system provides a default administrative user role, which provides access to all areas of JSA. Users who are assigned an administrative user role cannot edit their own account. This restriction applies to the default Admin user role. Another administrative user must make any account changes.

1. On the **Admin** tab, click **User Roles**.
2. On the toolbar, click **New**.

3. In the **User Role Name** field, type a unique name for this user role.

4. Select the permissions that you want to assign to the user role.

   The permissions that are visible on the **User Role Management** window depend on which JSA components are installed.

**Table 3: User Role Management window permissions**

| Permission | Description |
|---|---|
| **Admin** | Grants administrative access to the user interface. You can grant specific Admin permissions.<br>Users with **System Administrator** permission can access all areas of the user interface. Users who have this access cannot edit other administrator accounts.<br>**Administrator Manager**<br>Grants users permission to create and edit other administrative user accounts.<br>**Remote Networks and Services Configuration**<br>Grants users access to the Remote Networks and Services icon on the **Admin** tab.<br>**System Administrator**<br>Grants users permission to access all areas of user interface. Users with this access are not able to edit other administrator accounts. |
| **Delegated Administration** | Grant users permissions to perform limited administrative functions. In a multi-tenant environment, tenant users with **Delegated Administration** permissions can see only data for their own tenant environment. If you assign other administrative permissions that are not part of **Delegated Administration**, tenant users can see data for all tenants. |
| **Offenses** | Grants administrative access to all functions on the **Offenses** tab.<br>Users must have administrative access to create or edit a search group on the **Offenses** tab.<br>User roles must have the **Maintain Custom Rules** permission to create and edit custom rules. |

Table 3: User Role Management window permissions *(Continued)*

| Permission | Description |
|---|---|
| **Log Activity** | Grants access to functions in the **Log Activity** tab. You can also grant specific permissions:<br>**Maintain Custom Rules**<br>Grants permission to create or edit rules that are displayed on the **Log Activity** tab.<br>**Manage Time Series**<br>Grants permission to configure and view time series data charts.<br>**User Defined Event Properties**<br>Grants permission to create custom event properties.<br>**View Custom Rules**<br>Grants permission to view custom rules. If granted to a user role that does not also have the **Maintain Custom Rules** permission, the user role cannot create or edit custom rules. |
| **Network Activity** | Grants access to all the functions in the **Network Activity** tab. You can grant specific access to the following permissions:<br>**Maintain Custom Rules**<br>Grants permission to create or edit rules that are displayed on the **Network Activity** tab.<br>**Manage Time Series**<br>Grants permission to configure and view time series data charts.<br>**User Defined Flow Properties**<br>Grants permission to create custom flow properties.<br>**View Custom Rules**<br>Grants permission to view custom rules. If the user role does not also have the **Maintain Custom Rules** permission, the user role cannot create or edit custom rules.<br>**View Flow Content**<br>Grants permission to view source payload and destination payload in the flow data details. |

**Table 3: User Role Management window permissions *(Continued)***

| Permission | Description |
|---|---|
| Assets | This permission is displayed only if JSA Vulnerability Manager is installed on your system.<br>Grants access to the function in the **Assets** tab. You can grant specific permissions:<br>**Perform VA Scans**<br>Grants permission to complete vulnerability assessment scans. For more information about vulnerability assessment, see the *Managing Vulnerability Assessment Guide*.<br>**Remove Vulnerabilities**<br>Grants permission to remove vulnerabilities from assets.<br>**Server Discovery**<br>Grants permission to discover servers.<br>**View VA Data**<br>Grants permission to vulnerability assessment data. For more information about vulnerability assessment, see the *Managing Vulnerability Assessment guide*. |
| Reports | Grants permission to access all of the functions on the **Reports** tab.<br>**Distribute Reports via Email**<br>Grants permission to distribute reports through email.<br>**Maintain Templates**<br>Grants permission to edit report templates. |
| Risk Manager | Grants users permission to access JSA Risk Manager functions. JSA Risk Manager must be activated. |
| Vulnerability Manager | Grants permission to QRadar Vulnerability Manager function. QRadar Vulnerability Manager must be activated.<br>For more information, see the *Juniper Secure Analytics Vulnerability Manager User Guide*. |

Table 3: User Role Management window permissions *(Continued)*

| Permission | Description |
|---|---|
| IP Right Click Menu Extensions | Grants permission to options added to the right-click menu. |
| Platform Configuration | Grants permission to **Platform Configuration** services.<br>**Dismiss System Notifications**<br>Grants permission to hide system notifications from the **Messages** tab.<br>**View Reference Data**<br>Grants permission to view reference data when it is available in search results.<br>**View System Notifications**<br>Grants permission to view system notifications from the **Messages** tab. |
| JSA Log Source Management | Grants permission to the JSA Log Source Management app. |
| Pulse - Dashboard | Grants permission to dashboards in the QRadar Pulse app. |
| Pulse - Threat Globe | Grants permission to Threat Globe dashboard in the QRadar Pulse app. |
| QRadar Assistant | Grants permission to the IBM QRadar Assistant app. |
| QRadar Use Case Manager | Grants permission to the QRadar Use Case Manager app. |

5. In the **Dashboards** area, select the dashboards that you want the user role to access, and click **Add**.

> **NOTE**: A dashboard displays no information when the user role does not have permission to view dashboard data. If a user modifies the displayed dashboards, the defined dashboards for the user role appear at the next login.

6. Click **Save** and close the **User Role Management** window.

7. On the **Admin** tab menu, click **Deploy Changes**.

## Editing a User Role

You can edit an existing role to change the permissions that are assigned to the role.

To quickly locate the user role you want to edit on the **User Role Management** window, you can type a role name in the **Type to filter** text box.

1. On the **Admin** tab, click **User Roles**.
2. In the left pane of the **User Role Management** window, select the user role that you want to edit.
3. In the right pane, update the permissions, as necessary.
4. Modify the **Dashboards** options for the user role as necessary.
5. Click **Save**.
6. Close the **User Role Management** window.
7. On the **Admin** tab menu, click **Deploy Changes**.

## Deleting a User Role

If a user role is no longer required, you can delete the user role.

If user accounts are assigned to the user role you want to delete, you must reassign the user accounts to another user role. The system automatically detects this condition and prompts you to update the user accounts.

You can quickly locate the user role that you want to delete on the **User Role Management** window. Type a role name in the **Type to filter** text box, which is located above the left pane.

1. On the **Admin** tab, click **User Roles**.
2. In the left pane of the **User Role Management** window, select the role that you want to delete.
3. On the toolbar, click **Delete**.
4. Click **OK**.
    - If user accounts are assigned to this user role, the **Users are Assigned to this User Role** window opens. Go to Step "6" on page 21.

- If no user accounts are assigned to this role, the user role is successfully deleted. Go to Step "7" on page 21.

5. Reassign the listed user accounts to another user role:

   a. From the **User Role to assign** list box, select a user role.

   b. Click **Confirm**.

6. Close the **User Role Management** window.

7. On the **Admin** tab menu, click **Deploy Changes**.

RELATED DOCUMENTATION

Security Profiles | 21
User Accounts | 27
User Authentication | 34

# Security Profiles

IN THIS SECTION

Security profiles define which networks, log sources, and domains that a user can access.

JSA includes one default security profile for administrative users. The **Admin** security profile includes access to all networks, log sources, and domains.

Before you add user accounts, you must create more security profiles to meet the specific access requirements of your users.

## Domains

Security profiles must be updated with an associated domain. You must define domains on the **Domain Management** window before the **Domains** tab is shown on the **Security Profile Management** window. Domain-level restrictions are not applied until the security profiles are updated, and the changes are deployed.

Domain assignments take precedence over all settings on the **Permission Precedence**, **Networks**, and **Log Sources** tabs.

If the domain is assigned to a tenant, the tenant name appears in brackets beside the domain name in the **Assigned Domains** window.

## Permission Precedence

Permission precedence determines which security profile components to consider when the system displays events in the **Log Activity** tab and flows in the **Network Activity** tab.

Choose from the following restrictions when you create a security profile:

- **No Restrictions -**This option does not place restrictions on which events are displayed in the **Log Activity** tab, and which flows are displayed in the **Network Activity** tab.
- **Network Only -** This option restricts the user to view only events and flows that are associated with the networks that are specified in this security profile.
- **Log Sources Only -**This option restricts the user to view only events that are associated with the log sources that are specified in this security profile.
- **Networks AND Log Sources -**This option allows the user to view only events and flows that are associated with the log sources and networks that are specified in this security profile.

  For example, if the security profile allows access to events from a log source but the destination network is restricted, the event is not displayed in the **Log Activity** tab. The event must match both requirements.
- **Networks OR Log Sources -** This option allows the user to view events and flows that are associated with either the log sources or networks that are specified in this security profile.

For example, if a security profile allows access to events from a log source but the destination network is restricted, the event is displayed on the **Log Activity** tab if the permission precedence is set to **Networks OR Log Sources**. If the permission precedence is set to **Networks AND Log Sources**, the event is not displayed on the **Log Activity** tab.

### Permission Precedence for Offense Data

Security profiles automatically use the **Networks OR Log Sources** permission when offense data is shown. For example, if an offense has a destination IP address that your security profile permits you to see, but the security profile does not grant permissions to the source IP address, the **Offense Summary** window shows both the destination and source IP addresses.

## Creating a Security Profile

To add user accounts, you must first create security profiles to meet the specific access requirements of your users.

JSA includes one default security profile for administrative users. The Admin security profile includes access to all networks, log sources, and domains.

To select multiple items on the **Security Profile Management** window, hold the Control key while you select each network or network group that you want to add.

If after you add networks, log sources or domains you want to remove one or more before you save the configuration, you can select the item and click the **Remove (<)** icon. To remove all items, click **Remove All**.

1. On the **Admin** tab, click **Security Profiles**.

2. On the **Security Profile Management** window toolbar, click **New**.

3. Configure the following parameters:

    a. In the **Security Profile Name** field, type a unique name for the security profile. The security profile name must meet the following requirements: minimum of 3 characters and maximum of 30 characters.

    b. Optional: Type a description of the security profile. The maximum number of characters is 255.

4. Click the **Permission Precedence** tab.

5. In the Permission Precedence Setting pane, select a permission precedence option. See "Permission Precedence" on page 22.

6. Configure the networks that you want to assign to the security profile:

   a. Click the **Networks** tab.

   b. From the navigation tree in the left pane of the **Networks** tab, select the network that you want this security profile to have access to.

   c. Click the **Add (>)** icon to add the network to the Assigned Networks pane.

   d. Repeat for each network you want to add.

7. Configure the log sources that you want to assign to the security profile:

   a. Click the **Log Sources** tab.

   b. From the navigation tree in the left pane, select the log source group or log source you want this security profile to have access to.

   c. Click the **Add (>)** icon to add the log source to the Assigned Log Sources pane.

   d. Repeat for each log source you want to add.

8. Configure the domains that you want to assign to the security profile:

   Domains must be configured before the **Domains** tab appears.

   a. Click the **Domains** tab.

   b. From the navigation tree in the left pane, select the domain that you want this security profile to have access to.

   c. Click the **Add (>)** icon to add the domain to the Assigned Domains pane.

   d. Repeat for each domain that you want to add.

9. Click **Save**.

   > **NOTE**: The log sources and domains that are assigned to the security profile must match. If the log sources and domains do not match, you cannot save the security profile.

10. Close the **Security Profile Management** window.

11. On the **Admin** tab menu, click **Deploy Changes**.

## Editing a Security Profile

You can edit an existing security profile to update which networks and log sources a user can access and the permission precedence.

To quickly locate the security profile you want to edit on the **Security Profile Management** window, type the security profile name in the **Type to filter** text box.

1. On the **Admin** tab, click **Security Profiles**.
2. In the left pane, select the security profile that you want to edit.
3. On the toolbar, click **Edit**.
4. Update the parameters as necessary.
5. Click **Save**.
6. If the **Security Profile Has Time Series Data** window opens, select one of the following options:

| Option | Description |
|---|---|
| Keep Old Data and Save | Select this option to keep previously accumulated time series data. If you choose this option, users with this security profile might see previous data that they no longer have permission to see when they view time series charts. |
| Hide Old Data and Save | Select this option to hide the timeseries data. If you choose this option, time series data accumulation restarts after you deploy your configuration changes. |

7. Close the **Security Profile Management** window.
8. On the **Admin** tab menu, click **Deploy Changes**.

## Duplicating a Security Profile

If you want to create a new security profile that closely matches an existing security profile, you can duplicate the existing security profile and then modify the parameters.

To quickly locate the security profile you want to duplicate on the **Security Profile Management** window, type the security profile name in the **Type to filter** text box.

1. On the **Admin** tab, click **Security Profiles**.
2. In the left pane, select the security profile that you want to duplicate.
3. On the toolbar, click **Duplicate**.
4. In the **Confirmation** window, type a unique name for the duplicated security profile.
5. Click **OK**.
6. Update the parameters as necessary.
7. Close the **Security Profile Management** window.
8. On the **Admin** tab menu, click **Deploy Changes**.

## Deleting a Security Profile

If a security profile is no longer required, you can delete the security profile.

If user accounts are assigned to the security profiles you want to delete, you must reassign the user accounts to another security profile. JSA automatically detects this condition and prompts you to update the user accounts.

To quickly locate the security profile you want to delete on the **Security Profile Management** window, you can type the security profile name in the **Type to filter** text box.

1. On the **Admin** tab, click **Security Profiles**.
2. In the left pane, select the security profile that you want to delete.
3. On the toolbar, click **Delete**.
4. Click **OK**.
   - If user accounts are assigned to this security profile, the **Users are Assigned to this Security Profile** window opens. Go to "Deleting a User Role" on page 20.
   - If no user accounts are assigned to this security profile, the security profile is successfully deleted. Go to "Deleting a User Role" on page 20.
5. Reassign the listed user accounts to another security profile:
   a. From the **User Security Profile to assign** list box, select a security profile.

   b. Click **Confirm**.

6. Close the **Security Profile Management** window.

7. On the **Admin** tab menu, click **Deploy Changes**.

RELATED DOCUMENTATION

User Accounts | 27

User Authentication | 34

LDAP Authentication | 41

# User Accounts

IN THIS SECTION

The user account defines the unique user name that is used to log in to JSA, and specifies which user role, security profile, and tenant assignments the user is assigned to.

When you initially configure your system, you must create user accounts for each user who requires access to JSA.

## Viewing and editing Information About the Current User

You can view and edit account information for the current user through the main product interface.

1. Click the user icon in the upper right of the main product interface.

2. Click **User Preferences**.

3. Update the configurable user details.

| Parameter | Description |
|---|---|
| **Email** | Enter an email address to be associated with this user. The address cannot contain more than 255 characters, and cannot contain spaces. |
| **Current Password** | Enter your current password. |
| **New Password** | Enter a new password for the user to gain access. The password must meet the minimum length and complexity requirements that are enforced by the password policy |
| **Confirm New Password** | Enter the new password again. |
| **Locale** | Select a preferred language from the list. |
| **Enable Popup Notifications** | When enabled, system notification messages are displayed. To disable system notifications, set to off. |

4. Click **Save**.

## Viewing User Login History

You can view the login history of users to determine if there has been unauthorized access to their account. You can enable and disable the tracking of login attempts, and specify the retention period for tracking login attempts.

If you enable the login history display, a **Login History** window displays the date, time and IP address of the last successful login, and the number of unsuccessful login attempts of a user since the last successful login.

If you specify a retention period for tracking login attempts, JSA retains login history for the selected number of days.

When you change the login retention period, it takes effect for a user the next time they log in. For example if you change the login retention from 14 days to 7 days, any administrator continues to see 14 days of login history for any user that has not logged in since the change was made.

1. On the **Admin** tab, click **Authentication**.
2. Click **General Authentication Settings**.
3. Enable **Display Login History**.
4. Set the **Login History Retention (in days)** field to the number of days to retain the history of login attempts of a user.

   **NOTE**: The default is no value, which retains all login history.

5. Click **Save Settings**.
6. Close the **Authentication** window.

## Creating a User Account

When you create a new user account, you must assign access credentials, a user role, and a security profile to the user. User roles define what actions the user has permission to perform. Security profiles define what data the user has permission to access.

Before you can create a user account, you must ensure that the required user role and security profile are created.

You can create multiple user accounts that include administrative privileges; however, any user role with Administrator Manager privileges can create other administrative user accounts.

1. On the **Admin** tab, click **Users**.

   The **User Management** window opens.

2. Click **Add**.

3. Enter values for the following parameters:

| Parameter | Description |
|---|---|
| **User Name** | Enter a unique username for the new user. The username must contain 1 - 60 characters. |
| **User Description** | Enter a description for the user. The description cannot contain more than 2048 characters. |
| **Email** | Enter an email address to be associated with this user. The address cannot contain more than 255 characters, and cannot contain spaces. |
| **New Password** | Enter a new password for the user to gain access. The password must meet the minimum length and complexity requirements that are enforced by the password policy |
| **Confirm New Password** | Enter the new password again. |
| **User Role** | Select a role for this user from the list. |
| **Security Profile** | Select a security profile for this user from the list. |
| Override System Inactivity Timeout | Enable this setting to configure the inactivity timeout threshold for the user account. |

4. Click **Save**.

5. Close the **User Details** window.

6. On the **Admin** tab, click **Deploy Changes**.

## Editing a User Account

You can edit account information for the current user through the main product interface. To quickly locate the user account you want to edit on the **User Management** window, type the user name in the **Search User** text box on the toolbar.

1. On the **Admin** tab, click **Users**.
2. In the **User Management** window, select the user that you want to edit.

   You can use the **Advanced Filter** to search by User Role or Security Profile.
3. In the **User Details** window, click **Edit**.
4. Edit the account information for the user.
5. Click **Save**.
6. Close the **User Management** window.
7. On the **Admin** tab, click **Deploy Changes**.

## Disabling a User Account

You can disable a user account to restrict a user from accessing JSA. The option to disable a user account temporarily revokes a user's access without deleting the account.

If the user with the disabled account attempts to log in, a message is displayed to inform the user that the user name and password are no longer valid. Items that the user created, such as saved searches and reports, remain associated with the user.

1. On the **Admin** tab, click **Users**.
2. In the **User Management** window , click the user account that you want to disable.

   You can use the **Advanced Filter** to search by User Role or Security Profile.
3. Click **Edit**.
4. From the **User Details** window, select **Disabled** from the **User Role** list.
5. Click **Save**.
6. Close the **User Management** window.
7. On the **Admin** tab menu, click **Deploy Changes**.

## Deleting a User Account

If a user account is no longer necessary, you can delete the user account. After you delete a user, the user no longer has access to the user interface. If the user attempts to log in, a message is displayed to inform the user that the username and password is no longer valid.

To quickly locate the user account you want to delete on the **User Management** window, type the username in the **Search User** text box.

1. On the **Admin** tab, click **Users**.
2. In the **User Management** window, click the user account that you want to delete.

   You can use the **Advanced Filter** to search by User Role or Security Profile.
3. In the **User Details** window, click **Delete**. A search for dependents begins.
4. In the Found Dependents window, click **Delete** or **Re-Assign** dependents.
5. When the user has no dependents, click **Delete User**.
6. In the Confirm **Delete** window, click **Delete > OK**.
7. Click **Delete**.
8. Close the **User Management** window.
9. On the **Admin** tab, click **Deploy Changes**.

## Deleting Saved Searches of a Deleted User

If the saved searches of a deleted user are no longer necessary, you can delete the searches.

Saved searches that were created by a deleted user remain associated with the user until you delete the searches.

1. On the **Log Activity** or **Network Activity** tab, click **Search** > **Manage Search Results**.
2. Click the **Status** column to sort the saved searches.
3. Select the saved searches with a status of "ERROR!", then click **Delete**.

## Unlocking Locked User Accounts

New in 7.4.1 A user with root access can unlock user accounts that are locked out of JSA.

A user account can be locked out of JSA if there are too many failed login attempts for that account.

1. Using SSH, log in to your system as the root user.
2. Unlock specific user accounts or all user accounts.
   - Unlock specific user accounts by typing the following command:

     ```
     /opt/qradar/bin/runjava.sh com.ibm.si.security_model.authentication.AuthenticationLockoutCommandLineTool --
     removeaccount <user_account1> <user_account2> <user_account3>
     ```

   - Unlock all user accounts by typing the following command:

     ```
     /opt/qradar/bin/runjava.sh com.ibm.si.security_model.authentication.AuthenticationLockoutCommandLineTool --
     removeall- accounts
     ```

## Unlocking Locked Hosts

New in 7.4.1 A user with root access can unlock hosts that are locked out of JSA.

A host can be locked out of JSA if there are too many failed login attempts from that host.

1. Using SSH, log in to your system as the root user.
2. Unlock specific hosts or all user hosts.
   - Unlock specific hosts by typing the following command:

     ```
     /opt/qradar/bin/runjava.sh com.ibm.si.security_model.authentication.AuthenticationLockoutCommandLineTool --
     remove-ip <host_IP_address1> <host_IP_address2> <host_IP_address3>
     ```

   - Unlock all hosts by typing the following command:

     ```
     /opt/qradar/bin/runjava.sh com.ibm.si.security_model.authentication.AuthenticationLockoutCommandLineTool --
     removeall- ips
     ```

RELATED DOCUMENTATION

User Authentication | 34

LDAP Authentication | 41

Security Profiles | 21

# User Authentication

IN THIS SECTION

When authentication is configured and a user enters an invalid username and password combination, a message is displayed to indicate that the login was invalid.

If the user attempts to access the system multiple times with invalid information, the user must wait the configured amount of time before they can attempt to access the system again. You can configure console settings to determine the maximum number of failed logins, and other related settings. For more information about configuring console settings for authentication, see "JSA System Time" on page 89.

JSA supports the following authentication types:

- **System authentication** - Users are authenticated locally. System authentication is the default authentication type.
- **RADIUS authentication** - Users are authenticated by a Remote Authentication Dial-in User Service (RADIUS) server. When a user attempts to log in, JSA encrypts the password only, and forwards the username and password to the RADIUS server for authentication.
- **TACACS authentication** - Users are authenticated by a Terminal Access Controller Access Control System (TACACS) server. When a user attempts to log in, JSA encrypts the username and password, and forwards this information to the TACACS server for authentication. TACACS Authentication uses Cisco Secure ACS Express as a TACACS server. JSA supports up to Cisco Secure ACS Express 4.3.

- **Microsoft Active Directory -** Users are authenticated by a Lightweight Directory Access Protocol (LDAP) server that uses Kerberos.
- **LDAP -** Users are authenticated by an LDAP server.
- **SAML single sign-on authentication** – Users can easily integrate JSA with your corporate identity server to provide single sign-on, and eliminate the need to maintain JSA local users. Users who are authenticated to your identity server can automatically authenticate to JSA. They don't need to remember separate passwords or type in credentials every time they access JSA.

## Prerequisite Checklist for External Authentication Providers

Before you can configure an authentication type, you must complete the following tasks:

- Configure the authentication server before you configure authentication in JSA. For more information, see your server documentation.
- Ensure that the server has the appropriate user accounts and privilege levels to communicate with JSA. For more information, see your server documentation.
- Ensure that the time of the authentication server is synchronized with the time of the JSA server.
- Ensure that all users have appropriate user accounts and roles to allow authentication with the vendor servers.

## Configuring Inactivity Timeout for a User

If you have users who require longer periods of inactivity before they are logged out of the system, you can configure their inactivity timeout threshold individually.

1. On the **Admin** tab, click **Users**.
2. Select a user from the list and click **Edit**.
3. In the **User Details** pane, enable the **Override System Inactivity Timeout** setting.
4. Enter the number of minutes of inactivity before the user is logged out, and click **Save**.

## External Authentication Guidelines

You can configure an external authentication provider to allow JSA to authenticate users without JSA storing passwords locally for those users.

**NOTE**: You cannot configure more than one external authentication provider for JSA at a time. If you have set up one external authentication provider and you set up another external authentication provider, the configuration for the first external authentication provider is deleted.

When you choose to use an external authentication provider, consider these points:

- Ensure that your external provider is trustworthy because you are delegating an important security decision to this external provider. A compromised provider might allow access to your JSA to unintended parties.
- Ensure that the connection to the external provider is secure. Choose only secure communications protocols, by using LDAPS instead of LDAP, for example.
- Consider whether you want to enable a local authentication fallback in case the external provider is unavailable. If the external provider is compromised, it might be used as a denial of access attack.
- The decision to configure an external authentication provider applies to all administrator and nonadministrator JSA users. There is no such thing as a "local-only user" in JSA.
- If you enable auto-provisioning of JSA accounts, a compromised provider might be used to force the creation of a rogue JSA account, so use caution when you are combining these features.
- JSA users that do not have an entry in the external provider are relying on the fallback feature to check the local password. A compromised external authentication provider might be used to create a "shadow" for an existing JSA account, providing an alternative password for authentication.

### Local authentication fallback

Each non-administrator user can be configured for local authentication fallback. Local authentication fallback is turned off by default. If enabled, a non-administrator JSA user can access the system by using the locally stored password even if the external provider is unavailable, or if the password in the external provider is locked out or is unknown to the user. This also means that a rogue JSA administrator might change the locally stored password and log in as that user, so ensure that your JSA administrators are trustworthy. This is also the case if an external authentication provider is not configured.

The default administrator account, named admin, is always configured for local authentication fallback by default. This prevents the administrative user from being locked out of the system, but also means you must ensure that the configured external authentication provider has the correct entry for the admin user, and that the password is known only to the authorized JSA administrator. If you cannot maintain control of the admin entry in the external authentication provider, disable the admin account within JSA to prevent unauthorized users from logging in to JSA as admin. When you enable autoprovisioning, such as when you use LDAP group authentication, any user account that matches the LDAP query are created or reactivated with the appropriate roles as mapped. To prevent this from happening, disable auto-provisioning by using LDAP local.

For other privileged JSA users (users with the admin role), you can choose on a user-by-user basis whether to enable local authentication fallback. The ENABLE_FALLBACK_ALL_ADMINS setting (disabled by default) can be used to force all privileged users to use local authentication fallback. If local authentication fallback is configured, the same considerations apply as for the admin account.

When you configure an external authentication provider and create a new user, that user doesn't automatically have a local password set for JSA. If a user needs a local password, then you must configure local authentication fallback for that user. Local authentication fallback allows a user to authenticate locally if external authentication fails for any reason, including invalid passwords. Fallback users can then access JSA regardless of the state of the external authentication.

Even if local authentication fallback is enabled for a user account, JSA first attempts to authenticate the user to the external authentication module before it attempts local authentication. When external authentication fails, JSA automatically attempts to authenticate locally if local authentication fallback is enabled for that user. User accounts cannot be configured only to authenticate locally when an external authentication provider is configured. For this reason, it is important that all JSA user accounts correspond to an external authentication provider account of the same name associated with the same authorized user.

Ensure that the external authentication provider is trustworthy, as this configuration outsources a security decision and a rogue authentication admin can allow unauthorized access to your JSA. Make this connection in a secure way, by using the secure version of protocols (for example by using LDAPS rather than LDAP).

Local authentication fallback is not available with SAML authentication. No users are able to authenticate locally when you use SAML authentication.

When offboarding users, disable local authentication fallback for that user before you remove their authentication access from the external authentication provider.

## Configuring System Authentication

You can configure local authentication on your JSA system. You can specify length, complexity, and expiry requirements for local passwords.

The local authentication password policy applies to local passwords for administrative users. The policy also applies to non-administrative users if no external authentication is configured.

When the local authentication password policy is updated, users are prompted to change their password if they log in with a password that does not meet the new requirements.

1. On the **Admin** tab, click **Authentication**.
2. Click **Authentication Module Settings**.
3. Optional: From the **Authentication Module** list, select **System Authentication**.

   System authentication is the default authentication module. If you change from another authentication module, then you must deploy JSA before you do the next steps.
4. Click **Save Authentication Module**.
5. Click **Home**.
6. Click **Local Password Policy Configuration**.
7. Select the password complexity settings for local authentication.

## Configuring RADIUS authentication

You can configure RADIUS authentication on your JSA system.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. Click **System Configuration >User Management > Authentication**.
3. From the **Authentication Module** list box, select **RADIUS Authentication**.
4. Configure the parameters:
   a. In the **RADIUS Server** field, type the host name or IP address of the RADIUS server.

   b. In the **RADIUS Port** field, type the port of the RADIUS server.

   c. From the **Authentication Type** list box, select the type of authentication you want to perform.

      Choose from the following options:

| Option | Description |
|---|---|
| CHAP | Challenge Handshake Authentication Protocol (CHAP) establishes a Point-to-Point Protocol (PPP) connection between the user and the server. |
| MSCHAP | Microsoft Challenge Handshake Authentication Protocol (MSCHAP) authenticates remote Windows workstations. |
| ARAP | Apple Remote Access Protocol (ARAP) establishes authentication for AppleTalk network traffic. |
| PAP | Password Authentication Protocol (PAP) sends clear text between the user and the server. |

d. In the **Shared Secret** field, type the shared secret that JSA uses to encrypt RADIUS passwords for transmission to the RADIUS server.

5. Click **Save Authentication Module**.

## Configuring TACACS authentication

You can configure TACACS authentication on your JSA system.

1. On the navigation menu ( ≡ ), click **Admin**.
2. Click **System Configuration >User Management > Authentication.**
3. From the **Authentication Module** list box, select **TACACS Authentication**.
4. Configure the parameters:

    a. In the **TACACS Server** field, type the host name or IP address of the TACACS server.

    b. In the **TACACS Port** field, type the port of the TACACS server.

    c. From the **Authentication Type** list box, select the type of authentication you want to perform.

    Choose from the following options:

| Option | Description |
|---|---|
| ASCII | American Standard Code for Information Interchange (ASCII) sends the user name and password in clear text. |
| PAP | Password Authentication Protocol (PAP) sends clear text between the user and the server. PAP is the default authentication type. |
| CHAP | Challenge Handshake Authentication Protocol (CHAP) establishes a Point-to-Point Protocol (PPP) connection between the user and the server. |
| MSCHAP | Microsoft Challenge Handshake Authentication Protocol (MSCHAP) authenticates remote Windows workstations. |
| MSCHAP2 | Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAP2) authenticates remote Windows workstations by using mutual authentication. |
| EAPMD5 | Extensible Authentication Protocol using MD5 Protocol (EAPMD5) uses MD5 to establish a PPP connection. |

d. In the **Shared Secret** field, type the shared secret that JSA uses to encrypt TACACS passwords for transmission to the TACACS server.

5. Click **Save**.

## Configuring Active Directory authentication

You can configure Microsoft Active Directory authentication on your JSA system.

1. On the navigation menu
(
≡
), click **Admin**.

2. Click **System Configuration** >**User Management** > **Authentication**.

3. From the **Authentication Module** list box, select **Active Directory**.

   Configure the parameters:

a. In the **RADIUS Server** field, type the host name or IP address of the RADIUS server.

b. In the **RADIUS Port** field, type the port of the RADIUS server.

c. From the **Authentication Type** list, select **Active Directory** and configure the following parameters.

Configure the following parameters:

| Parameter | Description |
|---|---|
| Server URL | Type the URL used to connect to the LDAP server, for example, ldaps://host:port. |
| LDAP Context | Type the LDAP context you want to use. For example, DC=JSA,DC=INC. |
| LDAP Domain | Type the domain that you want to use. For example, jsa.inc. |

4. Click **Authentication Module**.

RELATED DOCUMENTATION

LDAP Authentication | 41

Security Profiles | 21

User Accounts | 27

# LDAP Authentication

IN THIS SECTION

You can configure JSA to use supported Lightweight Directory Access Protocol (LDAP) providers for user authentication and authorization.

JSA reads the user and role information from the LDAP server, based on the authorization criteria that you defined.

In geographically dispersed environments, performance can be negatively impacted if the LDAP server and the JSA Console are not geographically close to each other. For example, user attributes can take a long time to populate if the JSA Console is in North America and the LDAP server is in Europe.

You can use the LDAP authentication with an Active Directory server.

## Configuring LDAP Authentication

You can configure LDAP authentication on your JSA system.

If you plan to use SSL encryption or use TLS authentication with your LDAP server, you must import the SSL or TLS certificate from the LDAP server to the **/opt/qradar/conf/trusted_certificates** directory on your JSA Console. For more information about configuring the certificates, see "Configuring SSL or TLS Certificates" on page 49.

If you are using group authorization, you must configure a JSA user role or security profile on the JSA console for each LDAP group that is used by JSA. Every JSA user role or security profile must have at least one **Accept** group. The mapping of group names to user roles and security profiles is case-sensitive.

*Authentication* establishes proof of identity for any user who attempts to log in to the JSA server. When a user logs in, the username and password are sent to the LDAP directory to verify whether the credentials are correct. To send this information securely, configure the LDAP server connection to use Secure Socket Layer (SSL) or Transport Layer Security (TLS) encryption.

*Authorization* is the process of determining what access permissions a user has. Users are authorized to perform tasks based on their role assignments. You must have a valid bind connection to the LDAP server before you can select authorization settings.

User attribute values are case-sensitive. The mapping of group names to user roles and security profiles is also case-sensitive.

The user base DN is where JSA queries and finds users. Enable query permissions to allow your users to query against the user base DN.

1. On the **Admin** tab, click **Authentication**.

2. Click **Authentication Module Settings**.

3. From the **Authentication Module** list, select **LDAP**.

4. Click **Add** and complete the basic configuration parameters.

   There are three configuration types and each has specific requirements for the Server URL, SSL Connection, and TLS Authentication parameters:

   **Secure LDAP (LDAPS)**

   The **Server URL** parameter must use ldaps:// as the protocol, and specify an LDAP over SSL encrypted port (typically 636). For example ldaps://ldap1.example.com:636

   If you are using Global Catalog because you're using multiple domains, use port 3269. For example ldaps://ldap1.example.com:3269

   The **SSL Connection** parameter must be set to "True" and the **TLS Authentication** parameter must be set to "False".

   **LDAP with StartTLS**

   The **Server URL** parameter must use ldap:// as the protocol, and specify an LDAP unencrypted port that supports the StartTLS option (typically 389). For example ldap:// ldap1.example.com:389

   The **SSL Connection** parameter must be set to "False" and the **TLS Authentication** must be set to "True".

   TLS 1.2 using StartTLS is not the same as the LDAP SSL port.

   TLS Authentication does not support referrals, so referrals must be set to "ignore", and the LDAP server must include a complete structure to search.

   **Unencrypted**

   An unencrypted LDAP configuration is not recommended.

   The **Server URL** parameter must use the ldap:// protocol and specify an unencrypted port (typically 389). For example ldap://ldap1.example.com:389

   The **SSL Connection** parameter and the **TLS Authentication** parameter must both be set to "False".

**Table 4: LDAP Basic Configuration parameters**

| Parameter | Description |
|---|---|
| Search entire base | Select **True** to search all subdirectories of the specified Directory Name (DN).<br>Select **False** to search only the immediate contents of the Base DN. The subdirectories are not searched. This search is faster than one which searches all directories. |
| LDAP User Field | The user field identifier that you want to search on.<br>You can specify multiple user fields in a comma-separated list to allow users to authenticate against multiple fields. For example, if you specify **uid,mailid**, a user can be authenticated by providing either their user ID or their mail ID. |
| User Base DN | The Distinguished Name (DN) of the node where the search for a user would start. The **User Base DN** becomes the start location for loading users. For performance reasons, ensure that the **User Base DN** is as specific as possible.<br>For example, if all of your user accounts are on the directory server in the Users folder, and your domain name is ibm.com, the User Base DN value would be cn=Users,dc=ibm,dc=com. |
| Referral | Select **Ignore** or **Follow** to specify how referrals are handled. |

5. Under **Connection Settings**, select the type of bind connection.

**Table 5: LDAP bind connections**

| Bind connection type | Description |
|---|---|
| Anonymous bind | Use anonymous bind to create a session with the LDAP directory server that doesn't require that you provide authentication information. |

Table 5: LDAP bind connections *(Continued)*

| Bind connection type | Description |
|---|---|
| Authenticated bind | Use authenticated bind when you want the session to require a valid user name and password combination. A successful authenticated bind authorizes the authenticated user to read the list of users and roles from the LDAP directory during the session. For increased security, ensure that the user ID that is used for the bind connection does not have permissions to do anything other than reading the LDAP directory.<br>Provide the **Login DN** and **Password**. For example, if the login name is admin and the domain is juniper.com, the **Login DN** would be cn=admin,dc=juniper,dc=com. |

6. Click **Test connection** to test the connection information.

   You must provide user information to authenticate against the user attributes that you specified in the **LDAP User Field**. If you specified multiple values in **LDAP User Field**, you must provide user information to authenticate against the first attribute that is specified.

   > **NOTE**: The **Test connection** function tests the ability of JSA to read the LDAP directory, not whether you can log in to the directory.

7. Select the authorization method to use.

   Table 6: LDAP authorization methods

| Authorization method parameter | Description |
|---|---|
| Local | The user name and password combination is verified for each user that logs in, but no authorization information is exchanged between the LDAP server and JSA server. If you chose **Local** authorization, you must create each user on the JSA console. |

Table 6: LDAP authorization methods *(Continued)*

| Authorization method parameter | Description |
|---|---|
| User attributes | Choose **User Attributes** when you want to specify which user role and security profile attributes can be used to determine authorization levels.<br>You must specify both a user role attribute and a security profile attribute. The attributes that you can use are retrieved from the LDAP server, based on your connection settings. User attribute values are case-sensitive. |

Table 6: LDAP authorization methods *(Continued)*

| Authorization method parameter | Description |
|---|---|
| Group based | Choose **Group Based** when you want users to inherit role-based access permissions after they authenticate with the LDAP server. The mapping of group names to user roles and security profiles is case-sensitive.<br>**Group base DN**<br>Specifies the start node in the LDAP directory for loading groups.<br>For example, if all of your groups are on the directory server in the Groups folder, and your domain name is juniper.com, the **Group Base DN** value might be cn=Groups,dc=juniper,dc=com.<br>**Query limit enabled**<br>Sets a limit on the number of groups that are returned.<br>**Query result limit**<br>The maximum number of groups that are returned by the query. By default, the query results are limited to show only the first 1000 query results.<br>**By member**<br>Select **By Member** to search for groups based on the group members. In the **Group Member Field** box, specify the LDAP attribute that is used to define the users group membership.<br>For example, if the group uses the memberUid attribute to determine group membership, type memberUid in the **Group Member Field** box.<br>**By query**<br>Select **By Query** to search for groups by running a query. You provide the query information in the **Group Member Field** and **Group Query Field** text boxes.<br>For example, to search for all groups that have at least one **memberUid** attribute and that have a **cn** value that starts with the letter 's', type memberUid in **Group Member Field** and type cn=s* in **Group Query Field**. |

8. If you specified Group Based authorization, click **Load Groups** and click the plus (+) or minus (-) icon to add or remove privilege groups.

The user role privilege options control which JSA components the user has access to. The security profile privilege options control the JSA data that each user has access to.

> **NOTE**: Query limits can be set by selecting the **Query Limit Enabled** checkbox or the limits can be set on the LDAP server. If query limits are set on the LDAP server, you might receive a message that indicates that the query limit is enabled even if you did not select the **Query Limit Enabled** checkbox.

9. Click **Save**.

10. Click **Manage synchronization** to exchange authentication and authorization information between the LDAP server and the JSA console.

    a. If you are configuring the LDAP connection for the first time, click **Run Synchronization Now** to synchronize the data.

    b. Specify the frequency for automatic synchronization.

    c. Click **Close**.

11. Repeat the steps to add more LDAP servers, and click **Save** when complete.

## Synchronizing Data with an LDAP Server

You can manually synchronize data between the JSA server and the LDAP authentication server.

If you use authorization that is based on user attributes or groups, user information is automatically imported from the LDAP server to the JSA console.

Each group that is configured on the LDAP server must have a matching user role or security profile that is configured on the JSA console. For each group that matches, the users are imported and assigned permissions that are based on that user role or security profile.

> **NOTE**: If you manually run the synchronization, new data is not imported. LDAP users are imported only when you first log in to JSA.

By default, synchronization happens every 24 hours. The timing for synchronization is based on the last run time. For example, if you manually run the synchronization at 11:45 pm, and set the synchronization interval to 8 hours, the next synchronization will happen at 7:45 am. If the access permissions change for a user that is logged in when the synchronization occurs, the session becomes invalid. The user is redirected back to the login screen with the next request.

1. On the **Admin** tab, click **Authentication**.

2. Click **Authentication Module Settings**.

3. From the **Authentication Module** list, select **LDAP**.

4. Click **Manage Synchronization** >**Run Synchronization Now**.

## Configuring SSL or TLS Certificates

If you use an LDAP directory server for user authentication and you want to enable SSL encryption or TLS authentication, you must configure your SSL or TLS certificate.

1. Using SSH, log in to your system as the root user.

2. Type the following command to create the **/opt/qradar/conf/trusted_certificates/** directory:

   **mkdir -p /opt/qradar/conf/trusted_certificates**

3. Copy the SSL or TLS certificate from the LDAP server to the **/opt/qradar/conf/trusted_certificates** directory on your system.

4. Verify that the certificate file name extension is **.cert**, which indicates that the certificate is trusted.

   The JSA system loads only **.cert** files.

## Displaying Hover Text for LDAP Information

You create an LDAP properties configuration file to display LDAP user information as hover text. This configuration file queries the LDAP database for LDAP user information that is associated with events, offenses, or assets (if available).

The web server must be restarted after the LDAP properties is created. Consider scheduling this task during a maintenance window when no active users are logged in to the system.

The following example lists properties that you can add to an **ldap.properties** configuration file.

```
ldap.url=ldap://LDAPserver.example.com:389
ldap.authentication=simple
ldap.userName=user.name
ldap.password=your.encrypted.password
ldap.basedn=O=IBM,C=US
ldap.filterString=(&(objectclass=user)(samaccountname=%USER%))
ldap.attributes.displayName=Name
ldap.attributes.email=Email
ldap.attributes.employeeID=EmployeeID
ldap.attributes.department=Department
```

1. Use SSH to log in to JSA as a root user.

2. To obtain an encrypted LDAP user password, run the following **perl** script:

   **perl -I /opt/qradar/lib/Q1/ -e "use auCrypto; print Q1::auCrypto::encrypt *('<password>');*"**

3. Use a text editor to create the **/opt/qradar/conf/ldap.properties** configuration file.

4. Specify the location and authentication information to access the remote LDAP server.

   a. Specify the URL of the LDAP server and the port number.

      Use **ldaps://** or **ldap://** to connect to the remote server, for example, **ldap.url=ldaps://LDAPserver.example.com:389**.

   b. Type the authentication method that is used to access the LDAP server.

      Administrators can use the simple authentication method, for example, `ldap.authentication=simple`.

   c. Type the user name that has permissions to access the LDAP server.

      For example, `ldap.userName=`*`user.name`* .

   d. To authenticate to the remote LDAP server, type the encrypted LDAP user password for the user.

      For example, `ldap.password=`*`password`* .

   e. Type the base DN used to search the LDAP server for users.

      For example, `ldap.basedn=`*`BaseDN`* .

   f. Type a value to use for the search parameter filter in LDAP.

      For example, in JSA, when you hover over `ldap.filterString=(&(objectclass=user)(samaccountname=%USER%))`, the `%USER%` value is replaced by the user name.

5. Type one or more attributes to display in the hover text.

   You must include at least one LDAP attribute. Each value must use this format: `ldap.attributes.`*`AttributeName`*`=`*`Descriptive text to show in UI`*.

6. Verify that there is read-level permission for the **ldap.properties** configuration file.

7. Log in to JSA as an administrator.

8. On the **Admin** tab, select **Advanced** >**Restart Web Server**.

Administrators can hover over the **Username** field on the **Log Activity** tab and **Offenses** tab, or hover over the **Last User** field on the **Assets** tab (if available) to display more information about the LDAP user.

## Multiple LDAP Repositories

You can configure JSA to map entries from multiple LDAP repositories into a single virtual repository.

**NOTE**: If you configure the same user account in multiple LDAP servers, regardless of the User Base DN that is configured, a user can authenticate to either LDAP server. When they authenticate, the user is granted access to the same JSA account.

If multiple repositories are configured, when a user logs in, they must specify which repository to use for authentication. They must specify the full path to the repository and the domain name in the user name field. For example, if Repository_1 is configured to use domain `example.com` and Repository_2 is configured to use domain `example.ca.com`, the login information might look like these examples:

- `OU=User Accounts,OU=PHX,DC=qcorpaa,DC=aa,DC=example.com\`*username*

- `OU=Office,OU=User Accounts,DC=qcorpaa,DC=aa,DC=example.ca.com\`*username*

For an example using repository IDs, if the repository ID of Repository_1 is UsersJSA and the repository ID of Repository_2 is UsersJSAca, the login information might look like these examples:

- UsersJSA*<username>*

- UsersJSA*<username>*

User information is automatically imported from the LDAP server for repositories that use user attributes or group authorization. For repositories that use local authorization, you must create users directly on the JSA system.

## Example: Least Privileged Access Configuration and Set Up

Grant users only the minimum amount of access that they require to do their day-to-day tasks.

You can assign different privileges for JSA data and JSA capabilities. You can do this assignment by specifying different accept and deny groups for security profiles and user roles. Accept groups assign privileges and deny groups restrict privileges.

Let's look at an example. Your company hired a group of student interns. John is in his final year of a specialized cyber security program at the local university. He was asked to monitor and review known network vulnerabilities and prepare a remediation plan based on the findings. Information about the company's network vulnerabilities is confidential.

As the JSA administrator, you must ensure that the student interns have limited access to data and systems. Most student interns must be denied access to JSA Vulnerability Manager, but John's special assignment requires that he has this access. Your organization's policy is that student interns never have access to the JSA API.

The following table shows that John must be a member of the **company.interns** and **qvm.interns** groups to have access to JSA Risk Manager and JSA Vulnerability Manager.

**Table 7: User Role Privilege Groups**

| User Role | Accept | Deny |
|---|---|---|
| Admin | **jsa.admin** | **company.firedemployees** |
| QVM | **jsa.qvm**<br>**qvm.interns** | **company.firedemployees**<br>**jsa.qrm**<br>**company.interns** |
| QRM | **jsa.qrm**<br>**company.interns** | **company.firedemployees** |

The following table shows that the security profile for **qvm.interns** restricts John from accessing the JSA API.

Table 8: Security Profile Privilege Groups

| Security profile | Accept | Deny |
|---|---|---|
| QVM | qradar.secprofile.qvm | company.firedemployees |
| API | qradar.secprofile.qvm.api | company.firedemployees<br>qradar.secprofile.qvm.interns |

RELATED DOCUMENTATION

Security Profiles | 21
User Accounts | 27
User Authentication | 34

# SAML Single Sign-on Authentication

IN THIS SECTION

Security Assertion Markup Language (SAML) is a framework for authentication and authorization between a service provider (SP) and an identity provider (IDP) where authentication is exchanged using digitally signed XML documents. The service provider agrees to trust the identity provider to authenticate users. In return, the identity provider generates an authentication assertion, which indicates that a user has been authenticated.

By using the SAML authentication feature, you can easily integrate JSA with your corporate identity server to provide single sign-on, and eliminate the need to maintain JSA local users. Users who are authenticated to your identity server can automatically authenticate to JSA. They don't need to remember separate passwords or type in credentials every time they access JSA.

JSA is fully compatible with SAML 2.0 web SSO profile as a service provider. It supports both SP and IDP initiated single sign-on and single logout.

## Configuring SAML Authentication

You can configure JSA to use the Security Assertion Markup Language (SAML) 2.0 single signon framework for user authentication and authorization.

To complete SAML configuration in JSA, you must generate an XML metadata file on your Identity Provider (SAML) server.

Follow these steps to configure SAML authentication on your JSA host. After you complete this task, you must configure the Identity Provider to work with JSA.

1. On the **Admin** tab, click **Authentication**.

2. Click **Authentication Module Settings**.

3. From the **Authentication Module** list, select **SAML 2.0**.

4. In the **Identity Provider Configuration** section, click **Select Metadata File**, browse to the XML metadata file that was created by your Identity Provider, and then click **Open**.

5. In the **Service Provider Configuration** section, type the **Entity ID** URL.

6. Select a **NameID format**:

   - Unspecified (default)

   - Persistent

   - Email Address

   - X509 Certificate Subject Name

   - Windows Domain Name

   - Kerberos

**NOTE**: Use **Unspecified** unless your Identity Provider does not support it.

7. Select the **Request Binding Protocol**:

    - HTTP-POST
    - HTTP-Redirect

8. Select **Yes** for **Request Signed Assertion**, unless the device you are connecting to does not support signed assertions.

    ⚠ **CAUTION**: Selecting **No** leads to unauthenticated communication with the SAML device and is not recommended because it allows an unauthenticated network-based attacker to access protected resources.

9. If you want the assertion that is returned by the Identity Provider to be encrypted using a JSA certificate, select Yes for **Request Encrypted Assertion**.

    **NOTE**: Enabling encryption requires "Installing Unrestricted SDK JCE Policy Files" on page 57.

10. If you want to sign the authentication request by using a JSA certificate, select Yes for **Sign Authentication Request**.

11. If you want to automatically log users out of the Identity Provider when they log out of JSA, select Yes for **Enable Service Provider Initiated Single Logout**.

    **NOTE**: This option is available only if supported by your Identity Provider.

12. Use one of the following methods to configure a certificate for signing and decrypting:

**Table 9: Configure a Certificate for Signing and Decrypting**

| Option | Description |
|---|---|
| **Use the provided QRadar_SAML certificate** | Use the links in the tooltip to download the Root CA, Root CA CRL, Intermediate CA, and Intermediate CA CRL files of the certificate, which should be uploaded to the trusted certificate store of the Identity Provider server. |
| **Add a new certificate** | Click **Add** and follow the instructions in "Importing a New Certificate for Signing and Decrypting" on page 57 to add a custom certificate. |
| **Renew or update an existing certificate** | Click **Renew** to renew the `QRadar_SAML` certificate if it has expired or expires soon. Click **Update** to update a custom certificate that has expired or expires soon. These options appear based on which certificate you are using |

13. Select one of the following methods to authorize users:

**Table 10: Configure a Certificate for Signing and Decrypting**

| Option | Description |
|---|---|
| **Local** | You must create local JSA users and configure their roles and security profiles in User Manager. |
| **User Attributes** | JSA uses the attributes provided in SAML assertions to create local users automatically upon authentication requests. Roles and security profiles are assigned according to the value of the role attribute and the security profile attribute. These attributes must be provided in the assertions, and the roles and security profiles must already exist in JSA. Usernames, user roles and security profiles are case sensitive.<br>**NOTE**: When using a role with Admin capabilities, the value of the security profile attribute must be *Admin*.<br>**NOTE**: In a multi-tenancy environment, you must configure the *Tenant* attribute as well to assign users to tenants. If the tenant attribute is not provided, the user that is created is not assigned to any tenant. |

14. Click **Save Authentication Module**.

    The JSA SAML metadata file is automatically downloaded.

15. On the **Admin** tab, click **Deploy Changes**.

    After you configure JSA, you must configure your Identity Provider using the saved XML metadata file.

    If you selected Local authorization, go to "User Management" on page 15 to create local users. If you selected User Attributes, create roles, security profiles, and tenants as needed, then deploy.

## Installing Unrestricted SDK JCE Policy Files

Use of encryption technology is controlled by United States law. JSA Solution Developer Kits (SDKs) include strong but limited jurisdiction policy files. To support encrypted SAML assertions, with JSA, you must first obtain the unlimited jurisdiction Java Cryptography Extension (JCE) policy files.

1. Download the unrestricted Java Cryptography Extension (JCE) policy files.
2. Unpack the compressed file.

   Select the following JAR files from the unrestricted folder:

   - `local_policy.jar`
   - `US_export_policy.jar`
3. Place the files in the following directory on your JSA Console:

   **/opt/ibm/java-x86_64-80/jre/lib/security/**
4. On the Admin tab, click **Deploy Changes**.
5. Click **Advanced Settings** > **Restart Web Server**.

## Importing a New Certificate for Signing and Decrypting

The JSA SAML 2.0 feature has options to use an x509 certificate other than the provided `QRadar_SAML certificate` for signing and encryption.

1. For **Certificate for signing and encryption**, click **Add**.
2. In the **Import New Certificate** window, type a **Friendly Name** for the certificate.
3. Click **Browse** to select a **Private Key File**, and then click **Open**.
4. Click **Browse** to select a **Certificate File**, and then click **Open**.

5. If the certificate to upload has an Intermediate CA, click Browse to select the **Intermediate CA File**, and then click **Open**.

6. If the certificate's Root CA is not a common Root CA that is preinstalled with the operating system, click **Browse** to select the **Root CA File**, and then click **Open**.

7. Click **Upload** to upload the certificate.

## Setting up SAML with Microsoft Active Directory Federation Services

After you configure SAML in JSA, you can configure your Identity Provider by using the XML metadata file that you created during that process. This example includes instructions for configuring Microsoft Active Directory Federation Services (AD FS) to communicate with JSA using the SAML 2.0 single sign-on framework.

To configure the AD FS server, you must first configure SAML in JSA. Then copy the JSA SAML XML metadata file you created during that process to a location accessible to the AD FS server.

1. On the AD FS Management console, select the **Relying Party Trusts** folder.

2. On the Actions sidebar, click Standard **Relying Party Trust**, and click **Start**. This opens the **Add Relying Party Trust** wizard.

3. On the **Select Data Source** window, select **Import data about the relying party from a file**, browse to the JSA SAML XML metadata file, and click **Open**.

4. Click **Next**.

5. Type a **Display name** and add any relevant **Notes**, then click **Next**.

6. Select an access control policy, and click **Next**.

7. Configure any additional options you require, and click **Next**.

8. Click **Close**.

9. In the **Relying Party Trusts** folder, select the new trust you created, then click **Edit Claim Issuance Policy**.

10. Click **Add Rule**.

11. Select **Send LDAP Attributes as Claims** from the **Claim rule template** menu, then click **Next**.

12. Type a **Claim rule name**, and select the Attribute store.

13. Select the attributes to be sent in the assertion, map to the appropriate **Outgoing Claim Type**, and click **Finish**.

14. Click **Add Rule**.

15. Select **Transform an Incoming Claim** from the **Claim rule template** menu, then click **Next**.

16. Configure the following options:

    - Claim rule name

    - Incoming claim type - use value UPN

    - Outgoing claim type as NameID

    - Outgoing NameID format

17. Select **Pass through all claim values**, then click **Finish**.

18. If you configured JSA to use the provided QRadar_SAML certificate for SAML, copy the previously downloaded Root CA, intermediate CA, and CRL files to a directory on the Windows server. Then open a command line window as administrator on Windows OS and type the following commands:

    ```
    certutil -addstore -f ROOT <local_path>root-qradar-ca_ca
    certutil -addstore -f CA <local_path>QRadarSAML_ca.crt
    certutil -addstore -f ROOT <local_path>QRadarSAML_ca.crl
    certutil -addstore -f Root <local_path>root-qradar-ca_ca.crl
    ```

    The files are located in **/opt/qradar/ca/www**.

## Troubleshooting SAML Authentication

Use the following information to troubleshoot errors and issues when using SAML 2.0 with JSA.

### Sign on or logout failure

When single sign on or single logout fails, make sure that the JSA SAML metadata that you uploaded to the Identity Provider matches the latest deployed metadata at `https://`*<yourjsaserverhostname>*`/console/SAMLMetadata`. Also, make sure that you uploaded the root CA, root CA CRL, intermediate CA, intermediate

CA CRL files of your selected certificate to the right location of the IDP server's certificate stores. When the provided `QRadar_SAML` certificate is used, you can download these files at:

```
http://<yourjsaserverhostname>:9381/root-qradar-ca_ca
http://<yourjsaserverhostname>:9381/QRadarSAML_ca.crt
http://<yourjsaserverhostname>:9381/root-qradar-ca_ca.crl
http://<yourjsaserverhostname>:9381/QRadarSAML_ca.crl
```

**NOTE**: If you are using the provided QRadar_SAML certificate, the above steps are required after you restore JSA from a backup.

### Account not authorized

Certain configuration issues can produce this error:

```
This account is not authorized to access JSA. Logout from your SAML identity provider and use an authorized
account to login.
```

If you are using Local authorization, ensure that the **NameID** in the SAML assertion matches an existing JSA user name and that the user is deployed.

If you are using **User Attribute** authorization, ensure that the SAML assertion contains the configured role attribute and security profile attribute with values that match an existing deployed role and security profile in JSA. When using a role with *Admin* capabilities, the value of the security profile attribute must be Admin. If the assertion contains a tenant attribute in a multi-tenancy environment, ensure that the value of the attribute matches an existing tenant in JSA.

### Log files

You can diagnose many other issues by using the Identity Provider server logs and the **/var/log/qradar.error** log.

### Restore system login for investigation

To investigate issues with SAML 2.0, you can restore JSA to use the default system login.

Copy the content of the **/opt/qradar/conf/templates/login.conf into /opt/qradar/conf/ login.conf**

Alternatively, edit the /opt/qradar/conf/login.conf file and change

```
ModuleClass=com.q1labs.uiframeworks.auth.configuration.SAMLLoginModule
```

to

```
ModuleClass=com.q1labs.uiframeworks.auth.configuration.LocalPasswordLoginConfiguration
```

Clear the browser cache and login as an Admin user. After you complete your investigation, change the attribute back to `SAMLLoginModule` and clear the browser cache again.

### Unable to reach the JSA console after logging in with an identity provider

Ensure that the host name for the JSA console can be resolved by the local DNS server. Also, ensure that your computer can reach the JSA console by using the host name.

### Login or logout failures on the IDP server

Check the IDP server logs to determine if the failures are caused by errors in the CRL revocation checks. If so, import the QRadar_SAML certificate CRLs to the IDP server, or make sure that the IDP server can reach the JSA console by using an HTTP connection.

### Identity provider certificate is expired

When the certificate in the identity providers metadata file is expired, you cannot log in to JSA, and the following error appears in the **/var/log/qradar.error** file:

```
com.q1labs.uiframeworks.auth.saml.metadata.DefaultMetadataServiceImpl: [ERROR] NotAfter: <date>
java.security.cert.CertificateExpiredException: NotAfter:
```

To resolve this issue, ask your identity provider to update the certificate in the metadata file, and then reconfigure SAML in JSA to use the new IDP metadata file.

### QRadar_SAML certificate is expired

A JSA system notification is shown when the QRadar_SAML certificate is about to expire.

Before the certificate expires, you must renew it.

1. On the **Admin** tab, click **Authentication**.
2. Click **Authentication Module Settings**.
3. From the **Authentication Module** list, select **SAML 2.0**.
4. Click **Renew** to renew the QRadar_SAML certificate.
5. Click **Save Authentication Module**.

   The JSA SAML metadata file is automatically downloaded.

6. Click the links in the tooltip to download the JSA root CA and intermediate CA certificate, as well as the CRL files.

7. On the **Admin** tab, click **Deploy Changes**.

8. Send the following files to your IDP server to deploy the changes.

   - JSA metadata file
   - JSA root CA certificate
   - JSA intermediate CA certificate
   - `CRL` files

### Third-party certificate is expired

You do not have to use the QRadar_SAML certificate that is provided with JSA; you can use your own thirdparty certificate. When the certificate is about to expire, a JSAsystem notification is shown.

Before the third-party certificate expires, you must update the existing certificate or add a new certificate.

1. Click **Authentication Module Settings**.

2. From the **Authentication Module** list, select **SAML 2.0**.

3. Click **Add** or **Update**.

4. Click **Save Authentication Module**.

   The JSA SAML metadata file is automatically downloaded.

5. Click the links in the tooltip to download the JSA root CA and intermediate CA certificate, as well as the CRL files.

6. On the **Admin** tab, click **Deploy Changes**.

7. Send the following files to your IDP server to deploy the changes.

   - JSA metadata file
   - JSA root CA certificate
   - JSA intermediate CA certificate
   - `CRL` files

## RELATED DOCUMENTATION

Security Profiles | 21

User Accounts | 27

User Authentication | 34

4

CHAPTER

# License Management

# License Management

IN THIS SECTION

License keys entitle you to specific JSA products, and control the event and flow capacity for your JSA deployment. You can add licenses to your deployment to activate other JSA products, such as JSA Vulnerability Manager.

When you install JSA, the default license key is temporary and gives you access to the system for 35 days from the installation date. The email that you received from Juniper Networks when you purchased JSA contains your permanent license keys. These license keys extend the capabilities of your appliance, and you must apply them before the default license expires.

To apply a license key to the system, follow these steps:

1. Obtain the license key. For new or updated license keys, contact your local sales representative.
2. "Uploading a License Key" on page 71.
3. "Allocating a License Key to a Host" on page 72.
4. "Deploying Changes" on page 114.

After you apply the license keys to JSA, "Distributing Event and Flow Capacity" on page 73 to ensure that each of the managed hosts is allocated enough capacity to handle the average volume of network traffic, and still have enough EPS and FPM available to efficiently handle a data spike. You do not need to deploy the changes after you redistribute the EPS and FPM capacity.

## License Expiry

The processing capacity of the system is measured by the volume of events and flows that JSA can process in real time. The capacity can be limited by either the appliance hardware or the license keys. The temporary license key allows for 5,000 events per second (EPS) on the JSA console, and 10,000 EPS on each managed host. The FPM rate for the temporary license is 200,000 on both the JSA console and the managed hosts.

When a license expires, JSA continues to process events and flows up to the licensed capacity limits. If the EPS and FPM capacity of the expired license was allocated to a host, the shared license pool might go into a deficit, and cause JSA to block capabilities on the **Network Activity** and **Log Activity** tabs.

When JSA is not licensed to handle the volume of incoming network data, you can add a license that has more event or flow capacity.

# Event and Flow Processing Capacity

**IN THIS SECTION**

The capacity of a deployment is measured by the number of events per second (EPS) and flows per minute (FPM) that JSA can collect, normalize, and correlate in real time. The event and flow capacity is set by the licenses that are uploaded to the system.

Each host in your JSA deployment must have enough event and flow capacity to ensure that JSA can handle incoming data spikes. Most incoming data spikes are temporary, but if you continually receive system notifications that indicate that the system exceeded the license capacity, you can replace an existing license with a license that has more EPS or FPM capacity.

## Shared License Pool

The EPS and FPM rate that is set by each license is combined into a shared license pool. From the shared license pool, you can distribute the processing capacity to any host within a specific deployment or that is managed by a single console, regardless of which host the original license is allocated to.

By adjusting the allocation of the shared license pool, you ensure that the event and flow capacity is distributed according to the network workload, and that each JSA host has enough EPS and FPM to effectively manage periods of peak traffic.

In deployments that have separate event collector and event processor appliances, the event collector inherits the EPS rate from the event processor that it is attached to. To increase the capacity of the event collector, allocate more EPS from the shared license pool to the parent event processor.

### Contributions to the License Pool

A license that includes both event and flow capacity might not contribute both the EPS and FPM to the shared license pool. The license pool contributions are dependent on the type of appliance that the license is allocated to. For example, when you apply a license to a 16xx Event Processor, only the EPS is added to the license pool. The same license, when applied to a 17xx Flow Processor, contributes only the FPM to the license pool. Applying the license to an 18xx Event/Flow Processor contributes both EPS and FPM to the pool. With exception of software licenses for event or flow processors, all software licenses contribute both the EPS and FPM to the shared license pool, regardless of which type of appliance the license is allocated to.

As of JSA 7.3.2, you can now acquire stackable EPS/Flow increments instead of replacing existing console or other managed hosts license when you need to increase the overall event or flow thresholds of your deployment. After the licenses are uploaded and deployed, the event/flow capacity can then be reallocated through the License Pool Management.

### Exceeding Your Licensed Processing Capacity Limits

The license pool becomes overallocated when the combined EPS and FPM that is allocated to the managed hosts exceeds the EPS and FPM that is in the shared license pool. When the license pool is overallocated, the **License Pool Management** window shows a negative value for the EPS and FPM, and the allocation chart turns red. JSA blocks functionality on the **Network Activity** and **Log Activity** tabs, including the ability to view events and flows from the **Messages** list on the main JSA toolbar.

To enable the blocked functionality, reduce the EPS and FPM that you allocated to the managed hosts in your deployment. If the existing licenses do not have enough event and flow capacity to handle the volume of network data, upload a new license that includes enough EPS or FPM to resolve the deficit in the shared license pool.

### Expired Licenses

When a license expires, JSA continues to process events and flows at the allocated rate.

If the EPS and FPM capacity of the expired license was allocated to a host, the shared resources in the license pool might go into a deficit, and cause JSA to block functionality on the **Network Activity** and **Log Activity** tabs.

## Capacity Sizing

The best way to deal with spikes in data is to ensure that your deployment has enough events per second (EPS) and flows per minute (FPM) to balance peak periods of incoming data. The goal is to allocate EPS and FPM so that the host has enough capacity to process data spikes efficiently, but does not have large amounts of idle EPS and FPM.

When the EPS or FPM that is allocated from the license pool is very close to the average EPS or FPM for the appliance, the system is likely to accumulate data in a temporary queue to be processed later. The more data that accumulates in the temporary queue, also known as the burst-handling queue, the longer it takes JSA to process the backlog. For example, a JSA host with an allocated rate of 10,000 EPS takes longer to empty the burst handling queue when the average EPS rate for the host is 9,500, compared to a system where the average EPS rate is 7,000.

Offenses are not generated until the data is processed by the appliance, so it is important to minimize how frequently JSA adds data to the burst handling queue. By ensuring that each managed host has enough capacity to process short bursts of data, you minimize the time that it takes for JSA to process the queue, ensuring that offenses are created when an event occurs.

When the system continuously exceeds the allocated processing capacity, you cannot resolve the problem by increasing the queue size. The excess data is added to the end of the burst handling queue where it must wait to be processed. The larger the queue, the longer it takes for the queued events to be processed by the appliance.

## Internal Events

JSA appliances generate a small number of internal events when they communicate with each other as they process data.

To ensure that the internal events are not counted against the allocated capacity, the system automatically returns all internal events to the license pool immediately after they are generated.

RELATED DOCUMENTATION

Burst Handling | 69

Uploading a License Key | 71

Allocating a License Key to a Host | 72

# Burst Handling

IN THIS SECTION

JSA uses burst handling to ensure that no data is lost when the system exceeds the allocated events per second (EPS) or flows per minute (FPM) license limits.

When JSA receives a data spike that causes it to exceed the allocated EPS and FPM limits, the extra events and flows are moved to a temporary queue to be processed when the incoming data rate slows. When burst handling is triggered, a system notification alerts you that the appliance exceeded the EPS or FPM license limit.

The backlog in the temporary queue is processed in the order that the events or flows were received. The older data at the start of the queue is processed before the most recent data at the end of the queue. The rate at which the queue empties or fills is impacted by several factors, including the volume and duration of the data spike, the capacity of the appliance, and the payload size.

Hardware appliances normally can handle burst rates at least 50% greater than the appliance's stated EPS and FPM capability, and can store up to 5GB in the temporary queue. The actual burst rate capability depends upon the system load. VM appliances can achieve similar results if the VM is adequately sized and meets the performance requirements.

The burst recovery rate is the difference between the allocated rate and the incoming rate. When the volume of incoming data slows, the system processes the backlog of events or flows in the queue as fast as the recovery rate allows. The smaller the recovery rate, the longer it takes to empty the queue.

## Example: Incoming Data Spike

Every morning, between 8am and 9am, a company's network experiences a data spike as employees log in and begin to use the network resources.

The company's deployment includes a JSA Event and Flow Processor combo appliance that is allocated 5,000 events per second (EPS) and 100,000 flows per minute (FPM). The average capacity for this appliance is 4,000 EPS and 70,000 FPM.

During the data spike, which peaks around 9am, the appliance routinely receives up to 6,000 EPS and 120,000 FPM. JSA automatically moves the extra events and flows (1,000 EPS and 20,000 FPM) to the burst handling queue, and generates a system notification to alert the administrator that the appliance exceeded the allocated capacity.

The following images show a two-hour window when the incoming event and flow data exceeds the licensed capacity, which triggers a system notification, and a recovery period after the data volume returns to normal.

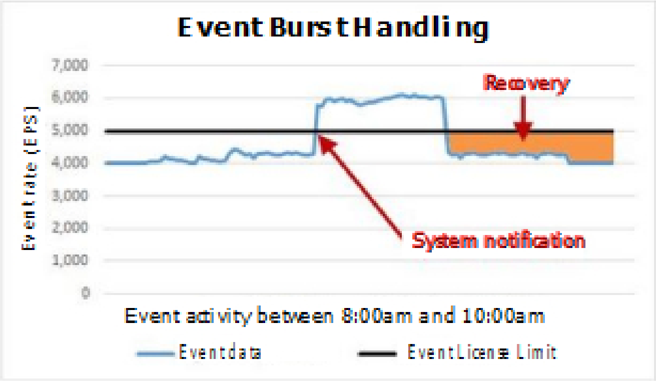

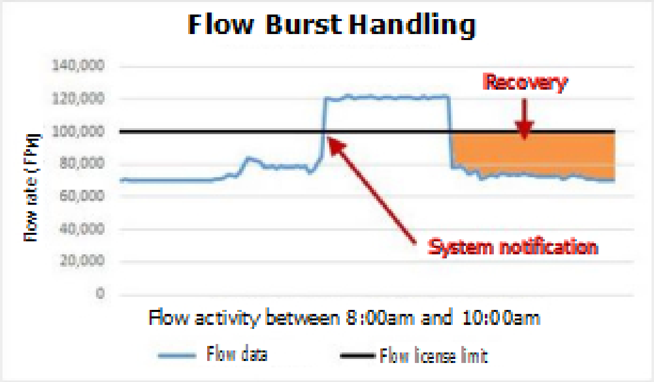

The recovery rate is the difference between the allocated EPS or FPM amount and the current incoming data rate. In this example, when the event and flow rates return to normal, the recovery rate is 1,000 EPS and 30,000 FPM.

**5,000 licensed events - 4,000 incoming events = 1,000 EPS recovery rate 100,000 licensed flows - 70,000 incoming flows = 30,000 FPM recovery rate**

Offenses are not generated until the data is processed by the appliance, so it is important to allocate enough EPS and FPM to the appliance to ensure that it can recover from a data spike quickly.

RELATED DOCUMENTATION

Uploading a License Key | 71
Allocating a License Key to a Host | 72
Distributing Event and Flow Capacity | 73

# Uploading a License Key

If you need assistance to obtain a new or updated license key, contact your local sales representative.

License keys determine your entitlement to JSA products and features, and the system capacity for handling events and flows.

You must upload a license key when you are doing these tasks:

- Updating an expired JSA console license
- Increasing the events per minute (EPS) or flows per minute (FPM) limits
- Adding a JSA product to your deployment, such as JSA Vulnerability Manager, to your deployment

As of JSA 7.3.0, you do not need to upload a new license when you add an Event Processor or Flow Processor to your deployment. Event and flow processors are automatically assigned a perpetual, or permanent, appliance license, and you can allocate EPS or FPM from the license pool to the appliance.

As of JSA 7.3.2, you can now acquire stackable EPS/Flow increments instead of replacing existing console or other managed hosts license when you need to increase the overall event or flow thresholds of your deployment. After the licenses are uploaded and deployed, the event/flow capacity can then be reallocated through the License Pool Management.

If the license key for your JSA console expires, you are automatically directed to the **System and License Management** window when you log on. You must upload a license key before you can continue.

If a managed host system has an expired license key, a message is displayed when you log in that indicates that a managed host requires a new license key. You use the **System and License Management** window to update the license key. If the license pool is not over allocated, delete the expired key and allocate EPS or FPM from the license pool to the managed host.

1. On the navigation menu (☰), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. On the toolbar, click **Upload License**.
4. In the dialog box, click **Select File**.
5. Select the license key, and click **Open**.
6. Click **Upload**, and then click **Confirm**.

The license is uploaded to your JSA console and is displayed in the **System and License Management** window.

By default, most licenses are not immediately allocated to a JSA host. However, the system automatically allocates all JSA Vulnerability Manager, and JSA Risk Manager keys to the JSA console.

"Allocating a License Key to a Host" on page 72.

RELATED DOCUMENTATION

Allocating a License Key to a Host | 72

Distributing Event and Flow Capacity | 73

Viewing License Details | 75

# Allocating a License Key to a Host

You must "Uploading a License Key" on page 71.

Allocate a license key to a JSA host when you want to replace an existing license, add new JSA products, or increase the event or flow capacity in the shared license pool.

You can allocate multiple licenses to a JSA console. For example, you can allocate license keys that add JSA Risk Manager and JSA Vulnerability Manager to your JSA console.

You cannot revert a license key after you add it to a JSA host. If you mistakenly allocate a license to the wrong host, you must deploy the change, and then delete the license from the system. After the license is deleted, you can upload the license again, and then reallocate it. After the license is allocated to the correct host, you must deploy the changes again.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. From the **Display** list, select **Licenses**.
4. Select the license, and click **Allocate System to License**.

   **TIP**: When you select **System** from the **Display** list, the label changes to **Allocate License to a System**.

5. To filter the list of licenses, type a keyword in the search box.
6. On the **Allocate a System to a License** window, select the host that you want to allocate the license to, and click **Allocate System to License**.

**NOTE**: You must purchase a separate HA license to set up high availability. However, the license should not be allocated to the secondary console or the secondary managed host.

RELATED DOCUMENTATION

Distributing Event and Flow Capacity | **73**

Viewing License Details | **75**

Deleting Expired Licenses | **76**

# Distributing Event and Flow Capacity

Use the **License Pool Management** window to ensure that the deployment is configured to use all of the events per second (EPS) and flows per minute (FPM) that you are entitled to. Also, ensure that JSA is configured to handle periodic bursts of data without dropping events or flows, or without having excessive amounts of unused EPS and FPM.

Ensure that the license pool has enough unallocated EPS or FPM. If the EPS or FPM in the license pool is fully allocated, redistribute the allocations.

Proper allocation of EPS and FPM capacity is important to ensure that JSA processes all events and flows in a timely manner. The goal is to allocate EPS and FPM so that the host has enough capacity to process data spikes efficiently, without having excessive idle EPS and FPM capacity.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **System and License Management.**
3. From the **Display** list, select **Licenses**.
4. Click **License Pool Management** and hover the mouse over the circle charts to see the total capacity for the deployment.
5. In the **License Allocations** table, review the data to determine whether the appliance has enough event and flow capacity to cover the average EPS and FPM, and still have enough left to cover the peak volumes.

   Learn more about reviewing the event and flow capacity data:

   - The **EPS Allocation** and **FPM Allocation** columns show the capacity that is assigned to each JSA processor or JSA console.
   - The **Average EPS** and **Average FPM** columns show the average number of events and flows that were processed by the JSA host over the last 30 days. The calculations use the **Event Rate (EPS)** and **Flow Rate (FPS)** saved searches. On deployments where the saved searches were deleted, the average event and flow rates appear as **N/A**.
   - Click the host name to view the details about the peak EPS and FPM rates for the past 30 days.
6. To change the allocated EPS or FPM rate for the JSA host, click the edit icon.
7. Update the **Allocated EPS** or **Allocated FPM** field, and click **Save**.

   The revised EPS and FPM allocations are validated against these criteria:

   - The EPS allocation must be a multiple of 500, and the FPM allocation must be a multiple of 5,000.
   - The allocated EPS or FPM does not cause the license pool to be over-allocated.
   - The allocated EPS or FPM does not exceed the hardware limits for the appliance type.

If your changes are not allocated correctly, click **Admin > Advanced > Restart Event Collection Services**.

If the issue persists, click **Admin > Advanced > Deploy Full Configuration**. If there are SourceMonitor Warning messages in the JSA logs, click **Admin > Advanced > Restart Event Collection Services**. A full deployment causes event collection to stop for several minutes.

RELATED DOCUMENTATION

Viewing License Details | 75
Deleting Expired Licenses | 76
Exporting License Information | 77

# Viewing License Details

View the license details to see information such as the status, expiration, and event and flow rate limits for each license that is uploaded to the system.

Licenses that are not yet allocated to a host appear at the top of the **License** table. Each host in the deployment has a summary row, which is shown in bold. The **Event Rate Limit** and the **Flow Rate Limit** fields on the summary row shows the EPS and FPM that is allocated to the host. If the host does not have any allocated EPS or FPM, **N/A** is shown in the **Event Rate Limit** and the **Flow Rate Limit** columns.

Licenses that are allocated to a JSA host appear as a child row, nested beneath the JSA host summary row. For the JSA console and Event and Flow Processor appliances, the child row shows the capacity and expiration dates for the EPS and FPM portion of the license. Before you manage the licenses, select the row that corresponds to the individual license.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. From the **Display** list, select **Licenses**.
4. To view detailed information about a specific host or license, select the nested row, and then click **Actions >View License**.

**Table 11: Status Of JSA Licenses**

| State | Description |
|---|---|
| **Unallocated** | The license is uploaded but not allocated to a JSA host. The EPS and FPM of the license don't contribute to the license pool. |
| **Undeployed** | The license is allocated to a JSA host, but is not deployed. The license is not yet active in your deployment. The EPS and FPM are included in the license pool. |

Table 11: Status Of JSA Licenses *(Continued)*

| State | Description |
|---|---|
| **Deployed** | The license is allocated and active in your deployment. The EPS and FPM are included in the license pool. |

RELATED DOCUMENTATION

Deleting Expired Licenses | 76

Exporting License Information | 77

Distributing Event and Flow Capacity | 73

# Deleting Expired Licenses

Delete a license if you mistakenly allocated it to the wrong JSA host. Also, delete an expired license to stop JSA from generating daily system notifications about the expired license.

You cannot delete a license if it causes the license pool to be over-allocated. JSA validates that the license pool has enough unallocated EPS and FPM capacity to cover the loss in capacity when the license is deleted. For example, if you want to delete a license that has 2,500 EPS associated with it, the license pool must have at least 2,500 EPS that has not been allocated to a JSA host.

If the license pool does not have enough unallocated EPS and FPM to cover the deficit, you must adjust the EPS and FPM allocations to ensure that the pool is not over-allocated when you delete the license.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. From the **Display** list, select **Licenses**.
4. In the host table, select the nested child row that contains the license that you want to delete.
5. Click **Actions** >**Delete License**.

   The **License Expiration Date** shows **Perpetual** with an **Event Rate Limit** and **Flow Rate Limit** of 0.

RELATED DOCUMENTATION

Exporting License Information | 77

Distributing Event and Flow Capacity | 73

Viewing License Details | 75

# Exporting License Information

For auditing, export information about the license keys that are installed on your system to an external **.xml** file.

You can't use the **.xml** file to move licenses to another system. Use it only for viewing detailed information about the individual license keys.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. From the **Display** list, select **Licenses**.
4. From the **Actions** menu, select **Export Licenses**.
5. Save the file locally and click **OK**.

RELATED DOCUMENTATION

Distributing Event and Flow Capacity | 73

Viewing License Details | 75

Deleting Expired Licenses | 76

5

CHAPTER

# System Management

# System Management

JSA has a modular architecture that supports deployments of varying sizes and topologies.

In a single-host deployment, all the software components run on a single appliance, and the JSA console provides the user interface, the real-time event and flow views, reports, offenses, asset information, and administrative functions.

To scale JSA, you can add non-console managed hosts to the deployment. You can configure a specific component type, such as collectors, processors, and data nodes, for each managed host, providing greater flexibility to manage data collection and processing in a distributed environment.

# System Health Information

The QRadar Deployment Intelligence app is a powerful monitoring application that consolidates historical health data for each managed host in your deployment. Use the app to monitor the health of your JSA deployment.

The Host status overview on the QRadar Deployment Intelligence dashboard shows the state of each appliance (active, standby, offline, or unknown), and the number of notifications for each host, the host name and appliance type, disk usage, status, and time changed. From the Host status overview, you can drill down to see more visual information about the status of the managed host, including the event and flow rates, system notifications, and disk information.

To assist with troubleshooting issues in your deployment, use the Get Logs capability to collect log files from the JSA Console and any other managed hosts in your deployment.

The QRadar Deployment Intelligence is available on the IBM Security App Exchange. You must install the app and then create an authorized service token to allow the app to use the JSA API to request data from the managed hosts.

The QRadar Deployment Intelligence uses JSA health metrics to monitor your deployment. Health metrics are essential, lightweight system events that do not count against your license.

RELATED DOCUMENTATION

JSA Component Types | 80

Data Nodes | 83

Network Interface Management | 87

# JSA Component Types

IN THIS SECTION

Each JSA appliance that is added to the deployment has configurable components that specify the way that the managed host behaves in JSA.

**Figure 4: JSA Event and Flow Components**

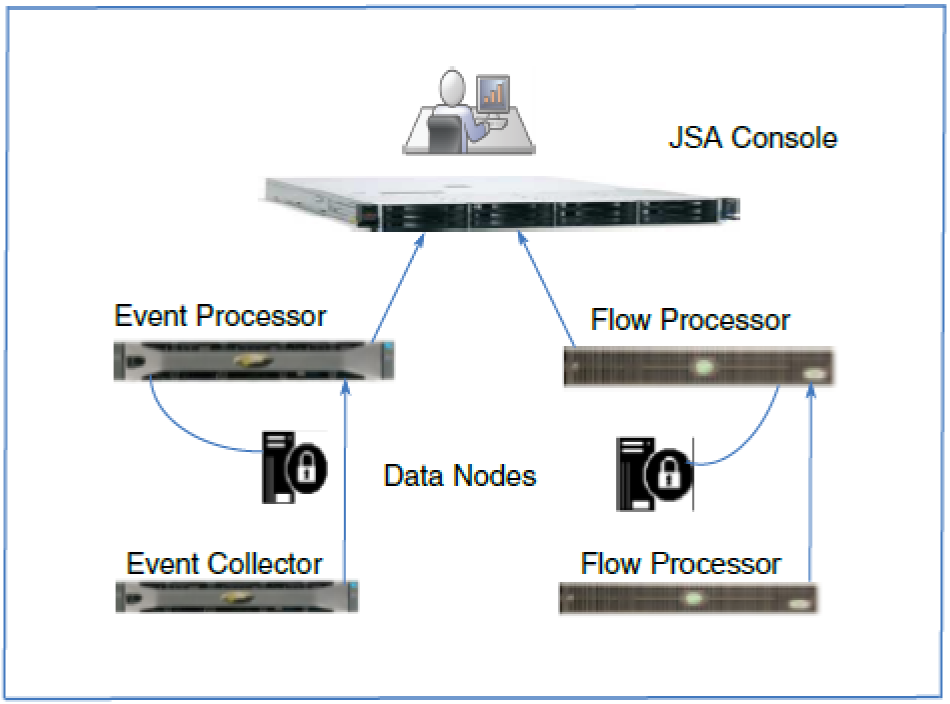

## JSA Console

The JSA console provides the JSA product interface, real-time event and flow views, reports, offenses, asset information, and administrative functions. In distributed environments, the JSA console is used to manage the other components in the deployment.

## Event Collector

The Event Collector collects events from local and remote log sources, and normalizes the raw event data so that it can be used by JSA. To conserve system resources, the Event Collector bundles identical events together and sends the data to the Event Processor.

## Event Processor

The Event Processor processes events that are collected from one or more Event Collector components. If events are matched to the custom rules that are defined on the Console, the Event Processor follows the action that is defined in the rule response.

Each Event Processor has local storage. Event data is stored on the processor, or it can be stored on a Data Node.

## JSA Flow Processor

JSA flow processor collects network flows from devices on your network. Live and recorded feeds are included, such as network taps, span ports, NetFlow, and JSA flow logs.

**NOTE**: Log Manager doesn't support flow collection.

## Flow Processor

The Flow Processor processes flows from one or more JSA flow processor appliances. The Flow Processor appliance can also collect external network flows such as NetFlow, J-Flow, and sFlow directly from routers in your network.

Flow Processors include an on-board processor and internal storage for flow data.

## Data Node

The Data Node receives security events and flows from event and flow processors, and stores the data to disk.

The Data Node is always connected to either an Event Processor or a Flow Processor.

## Off-site Source and Target Appliances

An off-site appliance is a JSA appliance that is not part of the deployment that is monitored by the JSA console.

An off-site source appliance forwards normalized data to an Event Collector. You can configure an off-site source to encrypt the data before forwarding.

An off-site target appliance receives normalized event or flow data from any Event Collector, or any processor in your deployment.

Later versions of JSA systems can receive data from earlier versions of JSA systems, but earlier versions can't receive data from later versions. To avoid problems, upgrade all receivers before you upgrade senders.

RELATED DOCUMENTATION

Data Nodes | 83
Network Interface Management | 87
JSA System Time | 89

# Data Nodes

IN THIS SECTION

A data node is an appliance that you can add to your event and flow processors to increase storage capacity and improve search performance. You can add an unlimited number of data nodes to your JSA deployment, and they can be added at any time. Each data node can be connected to only one processor, but a processor can support multiple data nodes.

For more information about planning your deployment, see the *Juniper Secure Analytics Architecture and Deployment Guide*.

## Data Rebalancing After a Data Node is Added

When you add a data node, JSA rebalances the data to improve search and overall system performance.

Data rebalancing includes decompressing older data, and moving data that was on the original storage device to evenly distribute it across all connected devices.

For example, your deployment has an event processor that receives 20,000 events per second (EPS). When you add data nodes, JSA automatically distributes the events across the event processor and all data nodes that are available to it. If you add three data nodes, the event processor stores 5,000 EPS and sends 5,000 EPS to each of the attached data nodes. The event processor is still processing all of the events, but the data nodes provide more storage, indexing, and search capabilities to improve the overall performance.

**How does rebalancing work?**

Cluster members consist of one event processor and one or more data nodes. Data can move between any members of the cluster in any direction. Data moves between members of the cluster transactionally by hourly folders. One hour of data is the smallest block of data that moves. If any file from an hourly folder is not copied, the entire transaction is rolled back.

Rebalancing does not merge hourly folders. For example, if an hourly folder exists on the destination, rebalancing does not move data from the same hourly folder from other members of the cluster. Before rebalancing starts, the cluster determines its target. The target is the percentage of free space that rebalancing tries to achieve on all members of the cluster. The target doesn't account for absolute free space in gigabytes, it accounts only for the percentage.

Members that have a higher percentage of free space are targets. After the cluster determines its target, the members that have a smaller percentage of free space than the target become sources. Each source connects, and pushes data, to each destination. Some components in your JSA deployment might restart and cause the rebalancing process to fail. Rebalancing restarts itself and continues from where it failed to completion. When rebalancing restarts, it does so with a progressively increasing timeout period (5 minutes, 10 minutes, 30 minutes, and so on) to avoid too many failed attempts during full deployment or maintenance. Whole rebalancing concludes between Ariel processes on members of the cluster.

**How does scattering work?**

Scattering distributes incoming data from the event processor among all members of the cluster. Scattering works with events and flows and is not bound to the smallest hourly block. For example, one hour of events is scattered across all clusters into the same hourly folder.

Scattering distributes events and flows proportionally to the amount of free space in percentage on the member of the cluster. Scattering moves data sequentially to the cluster hosts in round-robin fashion according to the free space percentage.

If any errors or connectivity issues occur, scattering tries to move the data to the next member of the cluster. If it is unsuccessful, it stores data locally on the event processor so that no data is lost. Data is scattered between the ecs-ep process (source) and multiple data node processes (destinations) on the data node.

**How is existing data moved between the event processor (source) and the data node (target)?**

When you add a data node, JSA calculates a target space. The target space is the amount of free space on the event processor, plus the amount of free space on the data nodes, divided by the total amount of event processors and data nodes. For example, you have one event processor and two data nodes. If the event processor has 60% free space and both data nodes have 100% free space, the target space is 86.6% (60 + 100 + 100 / 3). When the target is defined, the data is moved in one hour blocks at a time until the target space is reached (86.6% in this example) on any cluster hosts.

**How is new data moved between the event processor (source) and the data node (target)?**

When the initial balancing is complete, JSA scatters new data across the event processors and data nodes, according to the amount of free space available. For example, if an event processor has 25% free space and a data node has 40% free space, the data node receives 40 events, while the EP receives 25 events until both appliances have approximately the same amount of free space.

**When is balancing complete?**

The balancing process is complete when all source data is processed, or when the target space constraints are reached.

## Viewing the Progress Of Data Rebalancing

When you add a data node, JSA automatically redistributes the data to balance it across the storage volumes in your deployment.

Search performance improvements are realized only after the data rebalancing is complete. You can view the progress of the data rebalancing, and also see data such as the percentage of disk space that is used.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. In the host table, select the managed host that you want to view more information about.

   - To view information about the cluster of managed hosts, select the top-level host.

   - To view information about a specific data node, select the data node.

5. On the **Actions** menu, click **View and manage system**.

6. Click the **Security Data Distribution** tab to view the progress of data rebalancing and the capacity of the Data Node appliance.

   **NOTE**: You can also view information about the progress of data node rebalancing in the deployment status bar on the **Admin** tab.

## Saving All Event Data to a Data Node Appliance

To improve the performance of an event processor, configure JSA to save all event data on a Data Node appliance. With this configuration, the event processor only processes events; it doesn't store event data locally.

An event processor that is configured to only process events still saves event data locally when no active Data Node appliances are available. When a Data Node appliance becomes available, JSA transfers as much data as possible from the event processor to the Data Node.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.
4. Select the Event Processor from the host table, and on the **Deployment Actions** menu, click **Edit Host**.
5. Click the **Component Management** settings icon ( ⚙ ).
6. Under **Event Processor**, in the **Event Processor Mode** field, select **Processing-Only**.
7. Click **Save**, and then click **Save** again.
8. On the **Admin** tab, click **Deploy Changes**.

## Archiving Data Node Content

Configure a Data Node appliance to use **Archive** mode when you want the Data Node to provide online access to historical data without impacting storage for incoming data.

In **Archive** mode, the appliance does not receive new data, but existing data is saved.

> **NOTE**: No event retention policies are applied on the Data Node appliance in Archive mode.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.
4. Select the Data Node appliance in the host table, and on the **Deployment Actions** menu, click **Edit Host**.
5. Click the **Component Management** settings icon ( ⚙ ).
6. In the **Data Node Mode** field, select **Archive**, and then click **Save**.
7. On the **Admin** tab, click **Deploy Changes**.

To resume storing data on the Data Node appliance, set the mode back to **Active**.

RELATED DOCUMENTATION

Network Interface Management | 87
JSA System Time | 89
NAT-Enabled Networks | 92

# Network Interface Management

IN THIS SECTION

In addition to the default management interface, you can add extra network interfaces to your JSA appliances to provide alternative network connectivity.

Use extra network interfaces for the following purposes:

- Provide a dedicated crossover connection between high-availability (HA) peers. You configure a crossover connection during HA setup.

- Provide a dedicated data collection interface for inbound events or external flow sources. TCP-based data sources must be in the same subnet as the data collection interface.

Use a regular network interface card for:

- Data collection (logs/flows(NetFlow/s Flow))
- Web UI
- Backup/restore (not limited to iSCSI but can be NFS)

> **NOTE**: WinCollect configurations that are connected to a non-managed port are not supported.

## Configuring Network Interfaces

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. From the **Display** menu, click **Systems**.
4. Select the host for which you want to configure network interfaces.
5. Click **Actions** >**View and Manage System**, and click the **Network Interfaces** tab.
6. To edit a network interface, follow these steps:
   a. Select the device that you want to edit, and click **Edit**.
   b. In the **Role** list, select the role for the device:
      - Choose **Regular** when the device is used for:
        - Data collection (logs/flows(NetFlow/s Flow))
        - Web UI
        - Backup/restore (not limited to iSCSI but can be NFS)

        This interface must have an IP address. The subnet of the interface cannot be the same subnet used by the management interface.

- Choose **Monitor** when the device is a JSA Flow Processor that is used for packet collection. This interface does not require an IP address.
- Choose **Disabled** to prevent the device from being used for any network connectivity.

c. To apply the configuration to the active HA node, click **Apply this interface configuration and IP address to the active HA node**.

   This option is used only when HA is added to this system.

d. Click **Save**.

RELATED DOCUMENTATION

JSA System Time | 89

NAT-Enabled Networks | 92

Off-site Hosts Management | 95

# JSA System Time

IN THIS SECTION

When your deployment spans multiple time zones, configure all appliances to use the same time zone as the JSA Console. Alternatively, you can configure all appliances to use Greenwich Mean Time (GMT).

Configure the JSA system time from the JSA user interface. You can configure the time manually, or by configuring Network Time Protocol (NTP) servers to maintain the system time.

The time is automatically synchronized between the JSA console and the managed hosts.

## Problems That Are Caused by Mismatched Time Zones

To ensure that searches and data-related functions work properly, all appliances must synchronize time settings with the JSA console appliance. When the time zone settings are mismatched, you might see inconsistent results between JSA searches and report data.

The Accumulator service runs on all appliances with local storage to create minute by minute accumulations, and hourly and daily roll ups. JSA uses the accumulated data in reports and time series graphs. When the time zones are mismatched in a distributed deployment, report and time series graphs might show inconsistent results when compared to AQL query results due to the way that the accumulated data is aggregated.

JSA searches run against data that is stored in the Ariel databases, which use a date structure (YYYY/MM/DD/HH/MM) to store files to disk. Changing the time zone after the data is written to disk disrupts the file naming sequence in the Ariel databases and might cause data integrity problems.

## Configuring System Time

Configure system time on your JSA console by setting the time manually, or by using NTP servers to maintain the time. JSA synchronizes the JSA console time with the managed hosts in your deployment.

1. On the navigation menu (☰), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. Select the host for which you want to configure the system time settings.

5. From the **Actions** menu, click **View and Manage System**.

6. Click the **System Time** tab.

7. To configure time on the JSA console, follow these steps:

   a. In the **Time Zone** list, select the time zone that applies to the JSA console.

   b. To manually configure the time, click **Set time manually:**, and then set the date and time for the console.

> **NOTE**: If you set the system time to a future date that is affected by Daylight Saving Time (DST) changes, the time you set is adjusted by 1 hour. For example, on 4 July 2016 in the U.S.A, you set the date to December 16, 2016 and the time to 8:00 PM. The time that you set ignores the DST change and is adjusted to 7:00 PM.

c. To manage time by using NTP servers, follow these steps:

   i. Click **Specify NTP Servers**, and click **Add More**.

   ii. In the **Server 1 Address** field, type an IP address or a host name for the NTP server. Host names are resolved by a DNS server.

8. To configure time on a managed host, in the **Time Zone** list, select the time zone that applies to the host.

   On a managed host, you can configure only the time zone. The system time is synchronized with the JSA console but if the managed host is in a different time zone, then you can change to that time zone.

9. Click **Save**.

10. Click **OK** to accept that services are restarted, or **Cancel** to cancel the changes.

    Data collection for events and flows stops until the hostcontext and tomcat services are restarted.

When you set the system time on VMware systems and then restart the system, the changes might be lost. To prevent the time changes from being lost, you can disable time synchronization on the virtual device by editing the virtual machine's configuration file and adding these lines to the synchronization properties:

```
tools.syncTime = "FALSE"
time.synchronize.continue = "FALSE"
time.synchronize.restore = "FALSE"
time.synchronize.resume.disk = "FALSE"
time.synchronize.shrink = "FALSE"
time.synchronize.tools.startup = "FALSE"
```

The **.vmx** file is typically located in the directory where you created the virtual machine. For more information, see the vendor-specific documentation for your operating system.

RELATED DOCUMENTATION

NAT-Enabled Networks | 92
Off-site Hosts Management | 95
Managed Hosts | 102

# NAT-Enabled Networks

IN THIS SECTION

Network address translation (NAT) converts an IP address in one network to a different IP address in another network. NAT provides increased security for your JSA deployment because requests are managed through the conversion process and internal IP addresses are hidden. With NAT, computers that are located on a private, internal network are converted through a network device, typically a firewall, and can communicate to the public Internet through that network. Use NAT to map individual internal IP addresses to individual external IP addresses.

JSA NAT configuration requires static NAT translation and allows only one public IP address per managed host.

Any JSA host that is not in the same NAT group with its peer, or is in a different NAT group, is configured to use the public IP address of that host to reach it. For example, when you configure a public IP address on the JSA console, any host that is in the same NAT group uses the private IP address of the JSA console to communicate. Any managed host that is in a different NAT group uses the public IP address of the JSA console to communicate.

If you have a host in one of these NAT group locations that does not require external conversion, enter the private IP address in both the **Private IP** and **Public IP** fields. Systems in remote locations with a different NAT group than the console still require an external IP address and NAT, because they need to be able to establish connections to the console. Only hosts that are located in the same NAT group as the console can use the same public and private IP addresses.

## Configuring a NAT Group

Configure a Network Address Translation (NAT) group to limit the number of public IP addresses that are required for your JSA managed hosts to communicate with the Internet.

Ensure that the NAT-enabled network is using static NAT translation.

It is important to complete the NAT configuration for each managed host in your deployment before you deploy the changes. After deployment, managed hosts that aren't NAT-enabled might not be able to communicate with the JSA Console.

JSA can support multiple NAT networks when the public IP address for the JSA Console is the same in each network.

To configure a NAT group:

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. To configure a NAT group for the JSA Console, follow these steps:

    a. Select the JSA Console appliance in the host table.

    b. On the **Deployment Actions** menu, click **Edit Host**.

    c. Select the **Network Address Translation** check box.

    d. In the **NAT Group** list, select the NAT group that the console belongs to, or click the settings icon ( ⚙ ) to create a new NAT group.

    e. In the **Public IP** field, type the public IP address for the console, and then click **Save**.

5. Configure each managed host in the same network to use the same NAT group as the JSA Console.

    a. Select the managed host appliance in the host table.

    b. On the **Deployment Actions** menu, click **Edit Host**.

    c. Select the **Network Address Translation** check box.

   d. In the **NAT Group** list, select the NAT group that the JSA Console belongs to.

   e. In the **Public IP** field, type the public IP address for the managed host.

   > **NOTE**: Unless an event collector is connecting to a managed host that uses NAT, configure the managed host to use the same the public IP address and the private IP address.

   f. Click **Save**.

6. On the **Admin** tab menu, click **Advanced** > **Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

To fix communication issues between the JSA Console and hosts that are not NAT-enabled after deployment, edit the iptables rules for the managed host to configure the local firewall to allow the JSA Console to access the managed host.

## Changing the NAT Status for a Managed Host

Configure a managed host to use network address translation (NAT) to ensure that it can communicate with the JSA Console and other managed hosts in the same network.

Ensure that the NAT-enabled network is using static NAT translation.

The JSA Console and all managed hosts in the same network must be members of the same NAT group.

To change the NAT status for a managed host, make sure that you update the managed host configuration within JSA before you update the device. Updating the configuration first prevents the host from becoming unreachable, and ensures that you can continue to deploy changes to that host.

To change the NAT status for a managed host:

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.
4. Select the host in the host table, and on the **Deployment Actions** menu, click **Edit Host**.
5. To disable NAT, clear the **Network Address Translation** check box.
6. To enable NAT, follow these steps:
    a. Select the **Network Address Translation** check box.
    b. From the **NAT Group** list, select the group that the managed host belongs to.
    c. In the **Public IP** field, type the public IP address that the managed host uses to communicate with other hosts in a different NAT group.
7. Click **Save**.
8. On the **Admin** tab menu, click **Advanced** > **Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

If you enabled NAT, you might have to update the firewall configuration for the managed host that you want to communicate with.

RELATED DOCUMENTATION

Off-site Hosts Management | 95
Managed Hosts | 102
Deploying Changes | 114

# Off-site Hosts Management

IN THIS SECTION

An off-site host is a JSA appliance that can't be accessed through the JSA console in your current deployment. You can configure an off-site host to transfer data to or to receive data from your JSA deployment.

## Configuring an Off-site Source

To forward event and flow data to an Event Collector in another deployment, configure the target deployment to include an off-site source so that it knows which computer is sending the data.

To prevent connection errors, when you configure off-site source and target components, deploy the JSA Console with the off-site source first. Then, deploy the JSA console with the off-site target.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. On the **Deployment Actions** menu, click **Manage Off-site Sources**.

5. Click **Add** and configure the parameters.

   The name can be up to 20 characters in length and can include underscores or hyphens.

6. Click **Save**.

7. Click **Manage Connections** to specify which JSA hosts you want to receive the data.

   The host must have an Event Collector to receive the data.

8. Repeat the steps to configure all off-site sources that you want to configure.

9. Deploy the changes and restart the event collection service.

## Configuring an Off-site Target

To forward event and flow data to an Event Collector in another deployment, configure the source deployment to include an off-site target so that it knows which computer to send the data to.

You must know the listening ports for the off-site target appliance. By default, the listening port for events is 32004, and 32000 for flows.

To find the listening port on the target appliance, follow these steps:

1. In the target deployment, click the **System and License Management** icon.

2. Select the host and click **Deployment Actions >Edit Host**.

3. Click the **Component Management** settings icon ( ⚙ ), and find the ports in the **Event Forwarding Listening Port** and **Flow Forwarding Listening Port** fields.

To prevent connection errors, when you configure off-site source and target components, deploy the JSA Console with the off-site source first. Then, deploy the JSA console with the off-site target.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. On the **Deployment Actions** menu, click **Manage Off-site Targets**.

5. Click **Add** and configure the parameters.

   The name can be up to 20 characters in length and can include underscores or hyphens.

   The default port to listen for events is 32004, and 32000 for flows.

   > **NOTE**: If the off-site target is a managed host with encrypted host connections to its console, port 22 for SSH opens no matter which port is selected in the user interface.

6. Click **Save**.

7. Click **Manage Connections** to specify which JSA hosts you want to receive the data.

Only hosts that have an Event Collector are shown in the list.

8. Repeat the steps to configure all off-site targets that you want to configure.

9. On the **Admin** tab, click **Deploy changes**.

## Generating Public Keys for JSA Products

To forward normalized events in JSA, you must copy the public key file, **/root/.ssh/id_rsa.pub,** from the off-site source to the off-site target.

If the off-site source and off-site target are on separate systems, the public key is automatically generated. If the off-site source and target are both on an all-in-one system, the public key is not automatically generated. You must manually generate the public key.

To manually generate the public key, follow these steps:

1. Use SSH to log in to your system as the root user.

2. To generate the public key, type the following command:

   **opt/qradar/bin/ssh-key-generating**

3. Press Enter.

   The public and private key pair is generated and saved in the **/root/.ssh/id_rsa** folder.

## Forwarding Filtered Flows

You can set up forwarding of filtered flows. You can use filtered flows to split flow forwarding across multiple boxes, and to forward specific flows for specific investigations.

1. On the target system, set up the source system as an off-site source.

   a. On the navigation menu
   (
   ≡
   ), click **Admin**.

   b. Click **System and License Management** > **Deployment Actions** > **Manage Off-Site Sources.**

   c. Add the source system IP address, and select **Receive Events** and/or **Receive Flows.**

   d. Select **Manage Connections** and select which host is expecting to receive the off-site connection.

   e. Click **Save**.

   f. Select **Deploy Full Configuration** from the **Advanced** menu for the changes to take effect.

2. On the source system, set up the forwarding destination, IP address, and port number.

   a. Click **Main menu** > **Admin**.

   b. Click **Forwarding Destinations** > **Add**.

   c. Set the IP address of the target system and the destination port.

   d. Enter 32000 for the port number on the source system. Port 32000 is used for flow forwarding.

   e. Select **Normalized** from the **Event Format** list.

3. Set up routing rules.

   a. Click **Main menu** > **Admin**.

   b. Click **Routing Rules** > **Add**.

   c. Select the rules that you want to add.

   > **NOTE**: Rules forward flows that are based on offenses, or based on CRE information when **Offline Forwarding** is selected on the Routing Rules page.

   The flows that are filtered on the **Routing Rules** screen are forwarded.

## Example: Forwarding Normalized Events and Flows

To forward normalized events and flows, configure the target deployment to include an off-site source so that it knows which computer is sending the data. Configure the source deployment to include an off-site target so that it knows which computer to send the data to.

The following diagram shows forwarding event and flow data between deployments.

**Figure 5: Forwarding Data Between Deployments by Using SSH**

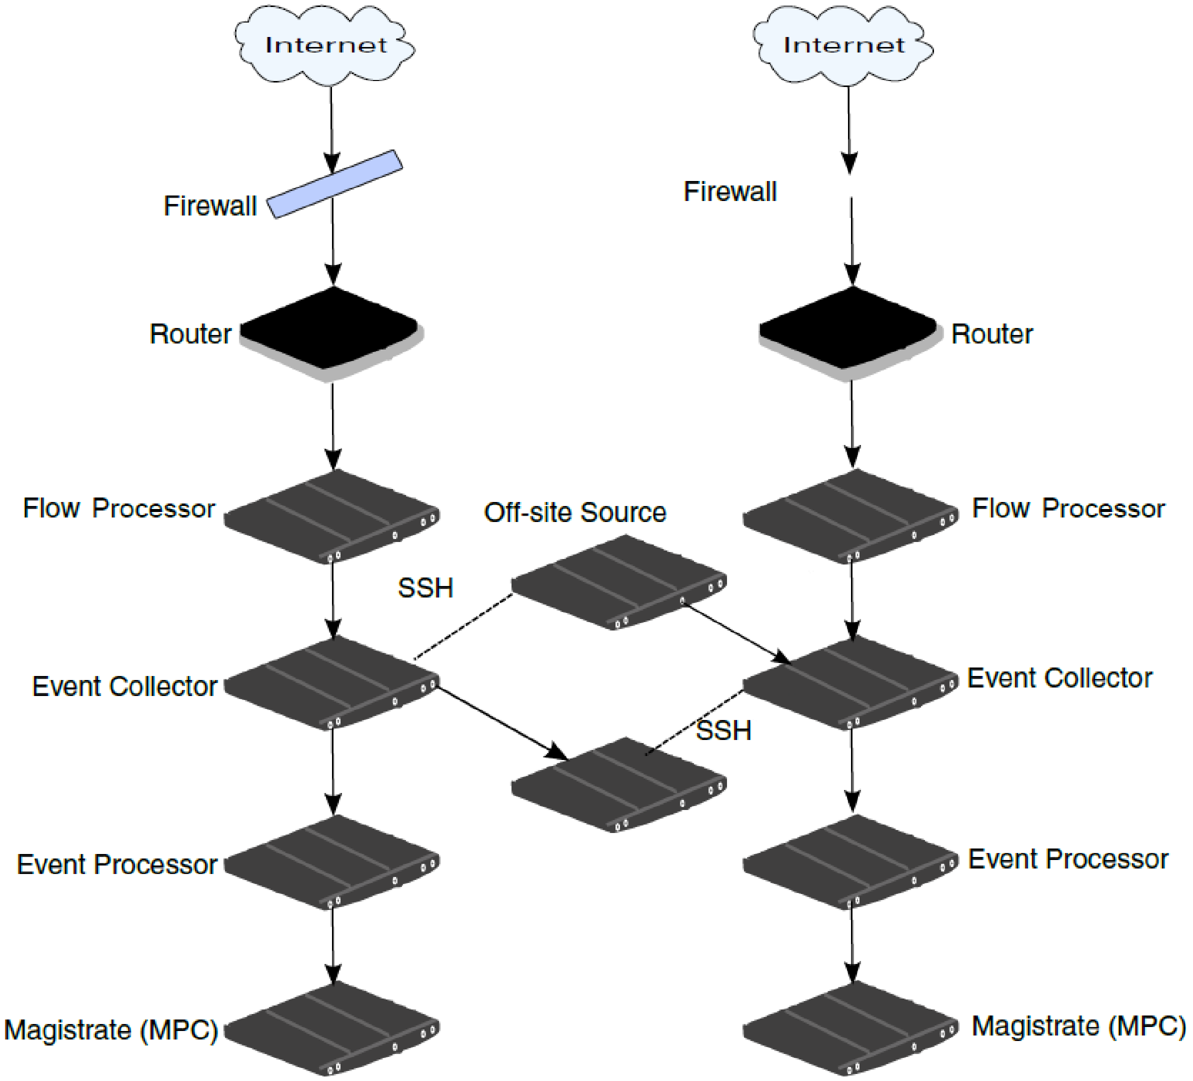

If the off-site source or target is an all-in-one system, the public key is not automatically generated; therefore, you must manually generate the public key. For more information, see "Generating public keys for JSA products" on page 98.

To forward normalized events and flows from Deployment A to Deployment B:

1. Configure an off-site target in Deployment A.

   The off-site target configuration includes the IP address of the Event Collector in Deployment B that receives the data.

2. Configure an off-site source in Deployment B.

The off-site source configuration includes the IP address and the port number of the Event Collector in Deployment A that is sending the data.

3. To transfer encrypted data, you must enable encryption on both the off-site source and the off-site target.

   To ensure appropriate access, the SSH public key for the source system (Deployment A) must be available to the target system (Deployment B). For example, to enable encryption between Deployment A and Deployment B, follow these steps:

4. Create ssh keys by using the **ssh-keygen -1 -t rsa** command, and press enter when prompted about the directory and passphrase.

   By default, the **id_rsa.pub** file is stored in the **/root/.ssh** directory.

5. Copy the **id_rsa.pub** file to the **/root/.ssh** directory on the Event Collector and on the JSA console in the source system (Deployment A).

6. Rename the file to **authorized_keys**.

   Ensure that the source system is configured with the appropriate permissions to send event and flow data to the target system.

7. If you didn't use the **chmod 600 authorized_keys** command to assign **rw** owner privileges to the file and the parent directory, use the **ssh-copy-id** command with the **-i** parameter to specify that the identity file **/root/.ssh/id_rsa.pub** be used.

   For example, type the following command to append entries or create a new **authorized_keys** file on the target console with the right privileges. This command does not check for duplicate entries.

   ```
   ssh-copy-id -i root@10.100.133.80
   ```

8. Configure the source system to ensure that forwarding of events and flows is not interrupted by other configuration activities, such as adding a managed host to one of the consoles.

   For example, if a managed host is added to a console that is forwarding events, then an **authorized_keys** file must exist in the **/root/.ssh** directory on the managed host. If not, adding a managed host fails. This file is required regardless of whether encryption is used between the managed host and the console.

9. On the JSA console in the source system (Deployment A), create a **ssh_keys_created** file under **/opt/qradar/conf**.

10. Change the owner and group to **nobody** and the permission to **775** to make sure that the file can be backed up and restored properly.

    ```
    chown nobody:nobody /opt/qradar/conf/ssh_keys_created chmod 775 /opt/qradar/conf/
    ssh_keys_created
    ```

11. To prevent connection errors, deploy the changes in the target system (Deployment B) before you deploy the changes in the source system (Deployment A).

If you update the Event Collector configuration or the monitoring ports, you must manually update the configuration for the off-site source and off-site target to maintain the connection between the two deployments.

If you want to disconnect the source system (Deployment A), you must remove the connections from both deployments. Remove the off-site target from the source system (Deployment A), and then remove the off-site source from the target system (Deployment B).

RELATED DOCUMENTATION

Managed Hosts | 102
Deploying Changes | 114
Shutting Down a System | 116

# Managed Hosts

IN THIS SECTION

For greater flexibility over data collection and event and flow processing, build a distributed JSA deployment by adding non-console managed hosts, such as gateways, processors, and data nodes.

For more information about planning and building your JSA environment, see the *Juniper Secure Analytics Architecture and Deployment Guide*.

## Software Compatibility Requirements

Software versions for all JSA appliances in your deployment must be at the same version and update package level. Deployments that use different versions of software are not supported because mixed software environments can prevent rules from firing, prevent offenses from being created or updated, or cause errors in search results.

When a managed host uses a software version that is different than the JSA Console, you might be able to view components that were already assigned to the host, but you cannot configure the component or add or assign new components.

## Internet Protocol (IP) Requirements

The following table describes the various combinations of IP protocols that are supported when you add non-console managed hosts:

**Table 12: Supported Combinations of IP protocols on Non-console Managed Hosts**

| Managed Hosts | JSA Console (IPv6, single) | JSA Console (IPv6, HA) |
|---|---|---|
| IPv4, single | No | No |
| IPv4, HA | No | No |
| IPv6, single | Yes | Yes |
| IPv6, HA | Yes | Yes |

## Bandwidth Considerations for Managed Hosts

To replicate state and configuration data, ensure that you have a minimum bandwidth of 100 Mbps between the JSA console and all managed hosts. Higher bandwidth is necessary when you search log and network activity, and you have over 10,000 events per second (EPS).

An Event Collector that is configured to store and forward data to an Event Processor forwards the data according to the schedule that you set. Ensure that you have sufficient bandwidth to cover the amount of data that is collected, otherwise the forwarding appliance cannot maintain the scheduled pace.

Use the following methods to mitigate bandwidth limitations between data centers:

- **Process and send data to hosts at the primary data center --** Design your deployment to process and send data as it's collected to hosts at the primary data center where the console resides. In this design, all user-based searches query the data from the local data center rather than waiting for remote sites to send back data.

  You can deploy a store and forward event collector, such as a JSA 15XX physical or virtual appliance, in the remote locations to control bursts of data across the network. Bandwidth is used in the remote locations, and searches for data occur at the primary data center, rather than at a remote location.

- **Don't run data-intensive searches over limited bandwidth connections --** Ensure that users don't run data-intensive searches over links that have limited bandwidth. Specifying precise filters on the search limits the amount of data that is retrieved from the remote locations, and reduces the bandwidth that is required to send the query result back.

## Encryption

To provide secure data transfer between each of the appliances in your environment, JSA has integrated encryption support that uses OpenSSH. Encryption occurs between managed hosts and is enabled by default when you add a managed host.

When encryption is enabled, a secure tunnel is created on the client that initiates the connection, by using an SSH protocol connection. When encryption is enabled on a managed host, an SSH tunnel is created for all client applications on the managed host. When encryption is enabled on a non-Console managed host, encryption tunnels are automatically created for databases and other support service connections to the Console. Encryption ensures that all data between managed hosts is encrypted.

The SSH tunnel between two managed hosts can be initiated from the remote host instead of the local host. For example, if you have a connection from an Event Processor in a secure environment to an Event Collector that is outside of the secure environment, and you have a firewall rule that would prevent you from having a host outside the secure environment connect to a host in the secure environment, you can switch which host creates the tunnel so that the connection is established from the Event Processor by selecting the **Remote Tunnel Initiation** checkbox for the Event Collector.

You cannot reverse the tunnels from your Console to managed hosts.

For example, with encryption enabled on an Event Processor, the connection between the Event Processor and Event Collector is encrypted, and the connection between the Event Processor and Magistrate is encrypted.

## Adding a Managed Host

Add managed hosts, such as event and flow processors and data nodes to distribute data collection and processing activities across your JSA deployment.

Ensure that the managed host has the same JSA version and update package as the JSA Console that you are using to manage it.

If you want to enable Network Address Translation (NAT) for a managed host, the network must use static NAT translation.

The following table describes the components that you can connect:

**Table 13: Supported Component Connections**

| Source Connection | Target Connection | Description |
|---|---|---|
| Flow Processor | Event Collector | You can connect a Flow Processor only to an Event Collector. The number of connections is not restricted.<br>You can't connect a Flow Processor to the Event Collector on a 15xx appliance. |
| Event Collector | Event Processor | You can connect an Event Collector to only one Event Processor.<br>You can connect a non-console Event Collector to an Event Processor on the same system.<br>A console Event Collector can be connected only to a console Event Processor. You can't remove this connection. |

**Table 13: Supported Component Connections *(Continued)***

| Source Connection | Target Connection | Description |
|---|---|---|
| Event Processor | Event Processor | You can't connect a console Event Processor to a non-console Event Processor.<br>You can connect a non-console Event Processor to another console or non-console Event Processor, but not both at the same time.<br>When a non-console managed host is added, the non-console Event Processor is connected to the console Event Processor. |
| Data Node | Event Processor | You can connect a data node to an event or flow processor only. You can connect multiple Data Nodes to the same processor to create a storage cluster. |
| Event Collector | Off-site target | The number of connections is not restricted. |
| Off-site source | Event Collector | The number of connections is not restricted.<br>An Event Collector that is connected to an event-only appliance can't receive an off-site connection from system hardware that has the Receive Flows feature enabled.<br>An Event Collector that is connected to a Flow-only appliance can't receive an off-site connection from a remote system that has the Receive Flows feature enabled. |

If you configured JSA Vulnerability Manager in your deployment, you can add vulnerability scanners and a vulnerability processor. For more information, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

If you configured JSA Risk Manager in your deployment, you can add a managed host. For more information, see the *Juniper Secure Analytics Risk Manager Installation Guide*.

To add a managed host:

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. On the **Deployment Actions** menu, click **Add Host**.

5. Configure the settings for the managed host by providing the fixed IP address, and the root password to access the operating system shell on the appliance.

6. Click **Add**.

7. Optional: Use the **Deployment actions** > **View Deployment** menu to see visualizations of your deployment. You can download a PNG image or a Microsoft Visio (2010) VDX file of your deployment visualization.

8. On the **Admin** tab menu, click **Advanced** > **Deploy Full Configuration**. When you deploy the full configuration, JSA restarts all services. Data collection for events and flows stops until the deployment completes.

**NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

## Configuring a Managed Host

Configure a managed host to specify which role the managed host fulfills in your deployment. For example, you can configure the managed host as a collector, processor, or a data node. You can also change the encryption settings, and assign the host to a network address translation (NAT) group.

To make network configuration changes, such as an IP address change to your JSA Console and managed host systems after you install your JSA deployment, use the **qchange_netsetup** utility. If you use `qchange_netsetup`, verify all external storage which is not **/store/ariel or /store** is not mounted.

For more information about network settings, see the *Installation Guide* for your product.

Ensure that the managed host has the same JSA version and update package as the JSA Console that is used to manage it. You can't edit or remove a managed host that uses a different version of JSA.

If you want to enable Network Address Translation (NAT) for a managed host, the network must use static NAT translation.

To configure a managed host:

1. On the navigation menu
(

≡ ), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.
4. Select the host in the host table, and on the **Deployment Actions** menu, click **Edit Host**.
   a. Optional: To initiate the tunnel between managed hosts from the remote host, select the **Remote Tunnel Initiation** checkbox.
   b. To configure the managed host to use a NAT-enabled network, select the **Network Address Translation** checkbox, and then configure the **NAT Group** and **Public IP address**.
   c. To configure the components on the managed host, click the **Component Management** settings icon ( ⚙ ) and configure the options.
   d. Click **Save**.
5. On the **Admin** tab menu, click **Advanced** > **Deploy Full Configuration**. When you deploy the full configuration, JSA restarts all services. Data collection for events and flows stops until the deployment completes.

## Removing a Managed Host

You can remove non-Console managed hosts from your deployment. You can't remove a managed host that hosts the JSA Console.

Ensure that the managed host has the same JSA version and update package as the JSA Console that is used to manage it. You can't remove a host that is running a different version of JSA.

To remove a managed host:

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.

4. On the **Deployment Actions** menu, click **Remove host** and click OK. You can't remove a JSA Console host.

5. On the **Admin** tab menu, click **Advanced** > **Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

## Configuring Your Local Firewall

Use the local firewall to manage access to the JSA managed host from specific devices that are outside the network. When the firewall list is empty, access to the managed host is disabled, except through the ports that are opened by default.

1. On the navigation menu (≡), click **Admin**.

2. In the **System Configuration** section, click **System and License Management**.

3. In the **Display** list, select **Systems**.

4. Select the host for which you want to configure firewall access settings.

5. From the **Actions** menu, click **View and Manage System**.

6. Click the **Firewall** tab and type the information for the device that needs to connect to the host.

   a. Configure access for devices that are outside of your deployment and need to connect to this host.

   b. Add this access rule.

7. Click **Save**.

   If you change the **External Flow Source Monitoring Port** parameter in the Flow configuration, you must also update your firewall access configuration.

## Adding an Email Server

JSA uses an email server to distribute alerts, reports, notifications, and event messages.

You can configure an email server for your entire JSA deployment, or multiple email servers.

> **NOTE**: JSA only supports encryption for the email server using STARTTLS.
>
> If you configure the mail server setting for a host as `localhost`, then the mail messages don't leave that host.

1. On the **Admin** tab, click **Email Server Management**.
2. Click **Add**, and configure the parameters for your email server.
3. Click **Save**.

   > **TIP**: Keep the **TLS** option set to **On** to send encrypted email. Sending encrypted email requires an external TLS certificate.

4. To edit an email server, click the **Other Settings** icon for the server, make your edits, and then click **Save**.
5. To delete an email server, click the **Other Settings** icon for the server, and then click **Delete**.
6. After you configure an email server, you can assign it to one or more hosts.

   a. On the **System and License Management** page, select a host.

   b. Change the **Display** list to show **Systems**.

   c. Click **Actions** > **View and Manage System**.

   d. On the **Email Server** tab, select an email server and click **Save**.

   e. Test the connection to the email server by clicking the **Test Connection** button.

   f. Click **Save**.

RELATED DOCUMENTATION

Deploying Changes | 114

Shutting Down a System | 116

# Configuration Changes in your JSA Environment

IN THIS SECTION

When you make configuration changes to JSA, the changes are saved to a staging area, and the deployment banner on the Admin tab is updated indicating that changes need to be deployed. Deploying the changes might require JSA services to restart.

JSA has two methods of deploying changes: standard and full configuration. The type of deployment that is required depends on the type of changes that were made.

## Standard Deployment

This deployment method restarts only those services that are directly affected by the changes that were made. You begin a standard deployment by clicking Deploy changes on the banner on the Admin tab.

The following list shows examples of changes that require a standard deployment:

- Adding or editing a new user or user role.
- Setting a password for another user.
- Changing a users' role or security profile.

## Full Configuration Deployment

Changes that affect the entire JSA deployment must be deployed by using the full configuration deployment method. You begin a full configuration deployment by clicking Deploy full configuration from the Advanced menu on the Admin tab.

This method rebuilds all configuration files on each of the managed hosts. To ensure that the new configuration is loaded properly, all services on the managed hosts are automatically restarted, except for the event collection service. While the other services restart, JSA continues collecting events and stores them in a buffer until the managed hosts come back online.

The following list shows examples of changes that require a full configuration deployment:

- Adding a managed host.
- Changing the configuration for a managed host.
- Configuring offsite hosts for sending or receiving data from the JSA Console.
- Restoring a configuration backup.

## Changes that Impact Event Collection

Events come into JSA through the ecs-ec-ingress event collection service. Starting in JSA 7.3.1, the service is managed separately from other JSA services. To minimize interruptions in collecting event data, the service does not automatically restart when the hostcontext service restarts.

The following situations can cause an interruption in event collection:

- Rebooting an appliance that collects events.
- Adding an HA managed host.
- During HA failover.
- Restoring a configuration backup.
- Adding or removing an off-site source connection.
- Whenever a partition's disk usage exceeds the maximum threshold.

When you deploy changes after you restore a configuration backup, you can restart the event collection service now or later. When you choose to restart the service later, JSA deploys all changes that don't depend on the event collection service, and continues to collect events while the other services restart.

The deployment banner continues to show undeployed changes, and the Event collection service must be restarted message is shown when you view the details.

## Configuring an Event Collector

Add a JSA Event Collector when you want to expand your deployment, either to collect more events locally or collect events from a remote location.

1. On the navigation menu, click Admin.
2. Click **System Configuration** > **System and License Management**.
3. Select the managed host that you want to configure.
4. Click **Deployment Actions** > **Edit Host**.
5. Click **Component Management**.
6. Enter values for the following parameters:

**Table 14: Event Collector Parameters**

| Parameter | Description |
|---|---|
| **Event Forwarding Listen Port** | The Event Collector event forwarding port. |
| **Flow Forwarding Listen Port** | The Event Collector flow forwarding port |
| **Autodetection Enabled** | **True** enables the Event Collector to automatically analyze and accept traffic from previously unknown log sources. The appropriate firewall ports are opened to enable Autodetection to receive events. This option is the default.<br>**False** prevents the Event Collector from automatically analyzing and accepting traffic from previously unknown log sources |
| **Autodetection - Use Global settings** | **True** specifies that the Event Collector uses global settings for Log Source Autodetection.<br>**False** specifies that the Event Collector uses individual, local settings (XML configuration file) for Log Source Autodetection. |

**Table 14: Event Collector Parameters *(Continued)***

| Parameter | Description |
|---|---|
| Flow Deduplication Enabled | |
| Flow Deduplication Filter Time | The amount of time in seconds that flows are buffered before they are forwarded. |
| Asymmetric Flow Filter Time | The amount of time in seconds that asymmetric flow is buffered before they are forwarded. |
| Forward Events Already Seen | **True** enables the Event Collector to forward events that were detected on the system.<br>**False** prevents the Event Collector from forwarding events that were detected on the system. This option prevents event-looping on your system. |
| Compress Event Processor Traffic | |

7. Click **Save**.

8. Repeat for all JSA Event Collectors in your deployment that you want to configure.

RELATED DOCUMENTATION

Deploying Changes | 114

Shutting Down a System | 116

Restarting a System | 117

# Deploying Changes

Changes that are made to the JSA deployment must be pushed from the staging area to the production area.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. Check the deployment banner to determine whether changes must be deployed.
3. Click **View Details** to view information about the undeployed configuration changes.
4. Choose the deployment method:

   a. To deploy changes from only the current session, click **Deploy Changes** on the deployment banner.

   b. To deploy all configuration changes that were made since the last deployment, click **Advanced** > **Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

RELATED DOCUMENTATION

Shutting Down a System | 116

Restarting a System | 117

Collecting Log Files | 117

# Restarting the Event Collection Service

There might be situations when you want to restart only the event collection service across all managed hosts in your JSA environment. For example, when a new version of the ecs-ec-ingress service is available for upgrade, or when you deferred restarting the service during an earlier deployment.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. On the **Advanced** menu, click **Restart Event Collection Services**. Event collection is briefly interrupted while the service restarts.

**NOTE**: You can also restart the event collection service at the command line by typing the following command:

```
systemctl restart ecs-ec-ingress
```

RELATED DOCUMENTATION

Shutting Down a System | **116**

Restarting a System | **117**

Collecting Log Files | **117**

# Shutting Down a System

When you shut down a system, the appliance is powered off. The JSA interface is unavailable and data collection stops while the system is shut down.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.
4. Select the system that you want to shut down.
5. From the **Actions** menu, select **Shutdown System**.

RELATED DOCUMENTATION

Restarting a System | **117**

Collecting Log Files | **117**

Changing the Root Password on Your JSA Console | **118**

# Restarting a System

When you restart a system, the JSA interface is unavailable and data collection stops while the system restarts.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.
4. Select the system that you want to restart.
5. From the **Actions** menu, select **Restart System**.

RELATED DOCUMENTATION

- Collecting Log Files | **117**
- Changing the Root Password on Your JSA Console | **118**
- Resetting SIM | **119**

# Collecting Log Files

JSA log files contain detailed information about your deployment, such as hostnames, IP addresses, and email addresses. If you need help with troubleshooting, you can collect the log files and send them to Juniper Networks Support.

You can collect the log files for one or more host systems at the same time. Depending on the size of your deployment and the number of managed hosts, collecting the log files might take a while. The JSA console log files are automatically included in each log file collection.

You can continue to use the JSA console while the log file collection is running. If the system is actively collecting log files, you can't begin a new collection request. Cancel the active collection process and start another collection.

When the log file collection process completes, a system notification appears on the **System Monitoring** dashboard.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System and License Management**.
3. In the **Display** list, select **Systems**.
4. Select the hosts in the host table.
5. Click **Actions** >**Collect Log Files**.
6. Click **Advanced Options** and choose the options for the log file collection.

   > **NOTE**: If you choose the **Encrypt compressed file** option, you must enter a password for the log file. If you are sending encrypted log files to Juniper Customer Support, you must also provide the password so that the log files can be decrypted.

   In previous releases, you could not specify a password and encrypted log files could only be decrypted by Juniper Customer Support.
7. Click **Collect Log Files**.

   Check the status of the collection process in the System Support Activities Messages section.
8. To download the log file collection, wait for the `Log file collection completed successfully` notification, and then click the **Click here to download file** link.

RELATED DOCUMENTATION

Changing the Root Password on Your JSA Console | **118**

Restarting a System | **117**

Resetting SIM | **119**

# Changing the Root Password on Your JSA Console

As a good security practice, change the root password on your JSA console at regular intervals.

1. Log in to your JSA console as the root user.
2. Use the **passwd** command to change your password.

RELATED DOCUMENTATION

Resetting SIM | 119

Restarting a System | 117

Collecting Log Files | 117

# Resetting SIM

After you tune your deployment, avoid receiving any additional false positive information by resetting SIM to remove all offense, and source and destination IP addresses from the database and the disk.

The SIM reset process can take several minutes, depending on the amount of data in your system. If you attempt to move to other areas of the JSA user interface during the SIM reset process, an error message is displayed.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. From the **Advanced** menu, select **Clean SIM Model**.
3. Read the information on the **Reset SIM Data Model** window.
4. Select one of the following options.

| Option | Description |
|---|---|
| **Soft Clean** | Closes all offenses in the database. If you select the **Soft Clean** option, you can also select the **Deactivate all offenses** check box. |
| **Hard Clean** | Purges all current and historical SIM data, which includes offenses, source IP addresses, and destination IP addresses. |

5. If you want to continue, select the **Are you sure you want to reset the data model?** check box.
6. Click **Proceed**.
7. When the SIM reset process is complete, click **Close**.
8. Refresh your web browser.

## RELATED DOCUMENTATION

Restarting a System | **117**

Collecting Log Files | **117**

Changing the Root Password on Your JSA Console | **118**

6

CHAPTER

# JSA Set Up Tasks

# JSA Set Up Tasks

Use the settings on the Admin tab to configure your JSA deployment, including your network hierarchy, automatic updates, system settings, event retention buckets, system notifications, console settings, and index management.

# Network Hierarchy

IN THIS SECTION

JSA uses the network hierarchy objects and groups to view network activity and monitor groups or services in your network.

When you develop your network hierarchy, consider the most effective method for viewing network activity. The network hierarchy does not need to resemble the physical deployment of your network. JSA supports any network hierarchy that can be defined by a range of IP addresses. You can base your network on many different variables, including geographical or business units.

## Guidelines for Defining Your Network Hierarchy

Building a network hierarchy in JSA is an essential first step in configuring your deployment. Without a well configured network hierarchy, JSA cannot determine flow directions, build a reliable asset database, or benefit from useful building blocks in rules.

Consider the following guidelines when you define your network hierarchy:

- Organize your systems and networks by role or similar traffic patterns.

  For example, you might organize your network to include groups for mail servers, departmental users, labs, or development teams. Using this organization, you can differentiate network behavior and enforce behaviour-based network management security policies. However, do not group a server that has unique behavior with other servers on your network. Placing a unique server alone provides the server greater visibility in JSA, and makes it easier to create specific security policies for the server.

- Place servers with high volumes of traffic, such as mail servers, at the top of the group. This hierarchy provides you with a visual representation when a discrepancy occurs.
- Avoid having too many elements at the root level.

  Large numbers of root level elements can cause the **Network hierarchy** page to take a long time to load.
- Do not configure a network group with more than 15 objects.

  Large network groups can cause difficulty when you view detailed information for each object. If your deployment processes more than 600,000 flows, consider creating multiple top-level groups.
- Conserve disk space by combining multiple Classless Inter-Domain Routings (CIDRs) or subnets into a single network group.

  For example, add key servers as individual objects, and group other major but related servers into multi-CIDR objects.

**Table 15: Example Of Multiple CIDRs and Subnets in a Single Network Group**

| Group | Description | IP addresses |
|---|---|---|
| 1 | Marketing | 10.10.5.0/24 |
| 2 | Sales | 10.10.8.0/21 |
| 3 | Database Cluster | 10.10.1.3/32<br>10.10.1.4/32<br>10.10.1.5/32 |

- Define an all-encompassing group so that when you define new networks, the appropriate policies and behavior monitors are applied.

  In the following example, if you add an HR department network, such as 10.10.50.0/24, to the Cleveland group, the traffic displays as Cleveland-based and any rules you apply to the Cleveland group are applied by default.

**Table 16: Example Of an All-encompassing Group**

| Group | Subgroup | IP address |
|---|---|---|
| Cleveland | Cleveland miscellaneous | 10.10.0.0/16 |
| Cleveland | Cleveland Sales | 10.10.8.0/21 |
| Cleveland | Cleveland Marketing | 10.10.1.0/24 |

- In a domain-enabled environment, ensure that each IP address is assigned to the appropriate domain.

## Acceptable CIDR Values

JSA accepts specific CIDR values.

The following table provides a list of the CIDR values that JSA accepts:

**Table 17: Acceptable CIDR Values**

| CIDR Length | Mask | Number of Networks | Hosts |
|---|---|---|---|
| /1 | 128.0.0.0 | 128 A | 2,147,483,392 |
| /2 | 192.0.0.0 | 64 A | 1,073,741,696 |
| /3 | 224.0.0.0 | 32 A | 536,870,848 |
| /4 | 240.0.0.0 | 16 A | 268,435,424 |
| /5 | 248.0.0.0 | 8 A | 134,217,712 |
| /6 | 252.0.0.0 | 4 A | 67,108,856 |

Table 17: Acceptable CIDR Values *(Continued)*

| CIDR Length | Mask | Number of Networks | Hosts |
|---|---|---|---|
| /7 | 254.0.0.0 | 2 A | 33,554,428 |
| /8 | 255.0.0.0 | 1 A | 16,777,214 |
| /9 | 255.128.0.0 | 128 B | 8,388,352 |
| /10 | 255.192.0.0 | 64 B | 4,194,176 |
| /11 | 255.224.0.0 | 32 B | 2,097,088 |
| /12 | 255.240.0.0 | 16 B | 1,048,544 |
| /13 | 255.248.0.0 | 8 B | 524,272 |
| /14 | 255.252.0.0 | 4 B | 262,136 |
| /15 | 255.254.0.0 | 2 B | 131,068 |
| /16 | 255.255.0.0 | 1 B | 65,534 |
| /17 | 255.255.128.0 | 128 C | 32,512 |
| /18 | 255.255.192.0 | 64 C | 16,256 |
| /19 | 255.255.224.0 | 32 C | 8,128 |
| /20 | 255.255.240.0 | 16 C | 4,064 |
| /21 | 255.255.248.0 | 8 C | 2,032 |

Table 17: Acceptable CIDR Values *(Continued)*

| CIDR Length | Mask | Number of Networks | Hosts |
|---|---|---|---|
| /22 | 255.255.252.0 | 4 C | 1,016 |
| /23 | 255.255.254.0 | 2 C | 508 |
| /24 | 255.255.255.0 | 1 C | 254 |
| /25 | 255.255.255.128 | 2 subnets | 124 |
| /26 | 255.255.255.192 | 4 subnets | 62 |
| /27 | 255.255.255.224 | 8 subnets | 30 |
| /28 | 255.255.255.240 | 16 subnets | 14 |
| /29 | 255.255.255.248 | 32 subnets | 6 |
| /30 | 255.255.255.252 | 64 subnets | 2 |
| /31 | 255.255.255.254 | none | none |
| /32 | 255.255.255.255 | 1/256 C | 1 |

For example, a network is called a supernet when the prefix boundary contains fewer bits than the natural (or classful) mask of the network. A network is called a subnet when the prefix boundary contains more bits than the natural mask of the network:

- 209.60.128.0 is a class C network address with a mask of /24.
- 209.60.128.0 /22 is a supernet that yields:
  - 209.60.128.0 /24
  - 209.60.129.0 /24
  - 209.60.130.0 /24

  - 209.60.131.0 /24

- 192.0.0.0 /25

  | Subnet | Host Range |
  |---|---|
  | 0 | 192.0.0.1-192.0.0.126 |
  | 1 | 192.0.0.129-192.0.0.254 |

- 192.0.0.0 /26

  | Subnet | Host Range |
  |---|---|
  | 0 | 192.0.0.1 - 192.0.0.62 |
  | 1 | 192.0.0.65 - 192.0.0.126 |
  | 2 | 192.0.0.129 - 192.0.0.190 |
  | 3 | 192.0.0.193 - 192.0.0.254 |

- 192.0.0.0 /27

  | Subnet | Host Range |
  |---|---|
  | 0 | 192.0.0.1 - 192.0.0.30 |
  | 1 | 192.0.0.33 - 192.0.0.62 |
  | 2 | 192.0.0.65 - 192.0.0.94 |
  | 3 | 192.0.0.97 - 192.0.0.126 |
  | 4 | 192.0.0.129 - 192.0.0.158 |
  | 5 | 192.0.0.161 - 192.0.0.190 |
  | 6 | 192.0.0.193 - 192.0.0.222 |
  | 7 | 192.0.0.225 - 192.0.0.254 |

## Defining Your Network Hierarchy

A default network hierarchy that contains pre-defined network groups is included in JSA. You can edit the pre-defined network hierarchy objects, or you can create new network groups or objects.

Network objects are containers for Classless Inter-Domain Routing (CIDR) addresses. Any IP address that is defined by a CIDR range in the network hierarchy is considered to be a local address. Any IP address that is not defined in a CIDR range in the network hierarchy is considered to be in a remote address. A CIDR can belong only to one network object, but subsets of a CIDR range can belong to another network object. Network traffic matches the most exact CIDR. A network object can have multiple CIDR ranges assigned to it.

Some of the default building blocks and rules in JSA use the default network hierarchy objects. Before you change a default network hierarchy object, search the rules and building blocks to understand how the object is used and which rules and building blocks might need adjustments after you modify the object. It is important to keep the network hierarchy, rules, and building blocks up to date to prevent false offenses.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.

2. In the **System Configuration** section, click **Network Hierarchy**.

3. From the menu tree on the **Network Views** window, select the area of the network in which you want to work.

4. To add network objects, click **Add** and complete the following fields:

   **Table 18: Add Network Objects**

| Option | Description |
|---|---|
| **Name** | The unique name of the network object.<br>**NOTE**: You can use periods in network object names to define network object hierarchies. For example, if you enter the object name D.E.F, you create a three-tier hierarchy with E as a subnode of D, and F as a subnode of E. |
| **Group** | The network group in which to add the network object. Select from the Group list, or click Add a New Group.<br>**NOTE**: When you add a network group, you can use periods in network group names to define network group hierarchies. For example, if you enter the group name A.B.C, you create a three-tier hierarchy with B as a subnode of A, and C as a subnode of B. |
| **IP/CIDR(s)** | Type an IP address or CIDR range for the network object, and click **Add**. You can add multiple IP addresses and CIDR ranges. |

Table 18: Add Network Objects *(Continued)*

| Option | Description |
|---|---|
| **Description** | A description of the network object. |
| **Country / Region** | The country or region in which the network object is located. |
| **Longitude and Latitude** | The geographic location (longitude and latitude) of the network object. These fields are co-dependent. |

5. Click **Create**.

6. Repeat the steps to add more network objects, or click **Edit** or **Delete** to work with existing network objects.

RELATED DOCUMENTATION

Automatic Updates | 130

Manual Updates | 140

Configuring System settings | 144

# Automatic Updates

IN THIS SECTION

You can automatically or manually update your configuration files to ensure that your configuration files contain the latest network security information.

Updated configuration files help to eliminate false positives and to protect your system from the latest malicious sites, botnets, and other suspicious Internet activity.

## Automatic Update Requirements

The JSA Console must be connected to the Internet to receive the updates. If your Console is not connected to the Internet, you must configure an internal update server for your Console to download the files from.

To maintain the integrity of your current configuration and information, either replace your existing configuration files or integrate the updated files with your existing files.

After you install updates on your Console and deploy your changes, the Console updates its managed hosts.

## Description Of Updates

Update files can include the following updates:

- Configuration updates that are based on content, including configuration file changes, vulnerabilities, QID maps, supportability scripts, and security threat information updates.

- DSM, scanner, and protocol updates that include corrections to parsing issues, scanner changes, and protocol updates.
- Major updates, such as updated JAR files or large patches, that require restarting the user interface service.
- Minor updates, such as daily automatic update logs or QID map scripts, that do not restart the user interface service.

## Automatic Updates for High Availability Deployments

When you update your configuration files on a primary host and deploy your changes, the updates are automatically made on the secondary host. If you do not deploy your changes, the updates are made on the secondary host through an automated process that runs hourly.

## Frequency Of Automatic Updates for New Installations and Upgrades

The default frequency of the automatic update is determined by the installation type and the JSA version.

- If you upgrade from JSA versions earlier than 2014.1, the value to which the update frequency is set remains the same after the upgrade. By default, the update is set to weekly, but you can manually change the frequency.
- If you install a new installation of JSA 2014.1 or later, the default frequency of the update is daily. You can manually change the frequency.

## Viewing Pending Updates

Your system is preconfigured for weekly automatic updates. You can view the pending updates in the **Updates** window.

Your system needs to be operational long enough to retrieve the weekly updates. If no updates are displayed in the **Updates** window, either your system has not been in operation long enough to retrieve the weekly updates or no updates have been issued. If this occurs, you can manually check for new updates. For more information about checking for new updates, see "Checking for New Updates" on page 137.

The **Check for Updates** toolbar provides the following functions:

**Table 19: Check for Updates Toolbar Functions**

| Function | Description |
|---|---|
| **Hide** | Select one or more updates, and then click **Hide** to remove the selected updates from the Check for Updates page. You can view and restore the hidden updates on the **Restore Hidden Updates** page. For more information, see "Restoring Hidden Updates" on page 139. |
| **Install** | You can manually install updates. When you manually install updates, the installation process starts within a minute. For more information, see "Manually Installing Automatic Updates" on page 138. |
| **Schedule** | You can configure a specific date and time to manually install selected updates on your Console. Scheduling is useful when you want to schedule the update installation during off-peak hours. For more information, see "Scheduling an Update" on page 136. |
| **Unschedule** | You can remove preconfigured schedules for manually installing updates on your Console. For more information, see "Scheduling an Update" on page 136. |
| **Search By Name** | You can locate a specific update by name. |
| **Next Refresh** | This counter displays the amount of time until the next automatic refresh. The list of updates on the **Check for Updates** page automatically refreshes every 60 seconds. The timer is automatically paused when you select one or more updates. |
| **Pause** | Pauses the automatic refresh process. To resume automatic refresh, click **Play**. |
| **Refresh** | Refreshes the list of updates. |

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Auto Update**.
3. To view details on an update, select the update.

## Configuring Automatic Update Settings

You customize the automatic update settings to change the frequency, update type, server configuration, and backup settings.

You can select the **Auto Deploy** to automatically deploy updates. If **Auto Deploy** is not selected, then you must manually deploy changes, from the **Dashboard** tab, after updates are installed.

> **NOTE**: In high-availability (HA) environment, automatic updates aren't installed when a secondary host is active. The updates are installed only after the primary host become the active node.

You can select **Auto Restart Service** to allow automatic updates that require the user interface to restart. A user interface disruption occurs when the service restarts. Alternatively, you can manually install the updated from the **Check for Updates** window.

1. On the navigation menu (≡), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Click **Change Settings**.

4. On the **Basic** tab, select the schedule for updates.

    a. In the **Configuration Updates** section, select the method that you want to use for updating your configuration files.

    - To merge your existing configuration files with the server updates without affecting your custom signatures, custom entries, and remote network configurations, select **Auto Integrate**.

    - To override your customizations with server settings, select **Auto Update**.

    b. In the **DSM, Scanner, Protocol Updates** section, select an option to install updates.

    c. In the **Major Updates** section, select an option for receiving major updates for new releases.

    d. In the **Minor Updates** section, select an option for receiving patches for minor system issues.

    e. If you want to deploy update changes automatically after updates are installed, select the **Auto Deploy** check box.

f. If you want to restart the user interface service automatically after updates are installed, select the **Auto Restart Service** check box.

5. Click the **Advanced** tab to configure the update server and backup settings.

   a. In **Web Server** field, type the web server from which you want to obtain the updates.

      The default web server is https://download.juniper.net/ .

   b. In the **Directory field**, type the directory location on which the web server stores the updates.

      The default directory is **software/strm/autoupdate/**.

   c. Optional: Configure the settings for proxy server.

      If the application server uses a proxy server to connect to the Internet, you must configure the proxy server. If you are using an authenticated proxy, you must provide the username and password for the proxy server.

   d. In the **Backup Retention Period** list, type or select the number of days that you want to store files that are replaced during the update process.

      The files are stored in the location that is specified in the **Backup Location**. The minimum is one day and the maximum is 65535 years.

   e. In the **Backup Location** field, type the location where you want to store backup files.

   f. In the **Download Path** field, type the directory path location to which you want to store DSM, minor, and major updates.

      The default directory path is **/store/configservices/staging/updates**.

## Configuring Updates Behind a Proxy Server That Uses SSL or TLS Interception

To configure JSA updates behind a proxy server, add your proxy server's CA certificate to the **ca-bundle.crt** file.

1. Create a backup copy of the **ca-bundle.crt** file in JSA. For example, use the copy command to create a .bak file: **cp /etc/ssl/certs/ca-bundle.crt{,bak}**.

2. Get the root CA certificate from your proxy server. For more information, see the proxy server documentation.

**NOTE**: You must use only the root CA certificate from your proxy server.

3. Add the CA certificate to the **ca-bundle.crt** file by typing the following command as one line:

   ```
   cp proxycert.pem/etc/pki/ca-trust/source/anchors
   ```

4. Extract the certificate by typing the following command:

   **update-ca-trust extract**

5. Type the following command to run the auto update:

   **/opt/qradar/bin/UpdateConfs.pl -ds runnow 1**

6. Verify that auto updates work by tailing the log in **/var/log/autoupdates/**.

## Scheduling an Update

Automatic updates occur on a recurring schedule according to the settings on the **Update Configuration** page. You can also schedule an update or a set of updates to run at a specific time.

To reduce performance impacts on your system, schedule a large update to run during off-peak hours.

For detailed information on each update, you can select the update. A description and any error messages are displayed in the right pane of the window.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Optional: If you want to schedule specific updates, select the updates that you want to schedule.

4. From the **Schedule** list box, select the type of update you want to schedule.

5. Using the calendar, select the start date and time of when you want to start your scheduled updates.

## Clearing Scheduled Updates

You can cancel any scheduled update.

Scheduled updates display a status of **Scheduled** in the **Status** field. After the schedule is cleared, the status of the update displays as **New**.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Click **Check for Updates**.

4. If you want to clear specific scheduled updates, select the updates that you want to clear.

5. From the **Unschedule** list box, select the type of scheduled update that you want to clear.

## Checking for New Updates

Juniper Networks provides updates on a regular basis. By default, the Auto Update feature is scheduled to automatically download and install updates. If you require an update at a time other than the preconfigured schedule, you can download new updates.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Click **Check for Updates**.

4. Click **Get new updates**.

## Manually Installing Automatic Updates

Juniper Networks provides updates regularly. By default, updates are automatically downloaded and installed on your system. However, you can install an update at a time other than the preconfigured schedule.

The system retrieves the new updates from https://download.juniper.net/. This might take an extended period. When complete, new updates are listed on the Updates window.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Click **Check for Updates**.

4. Optional: If you want to install specific updates, select the updates that you want to schedule.

5. From the **Install** list box, select the type of update you want to install.

## Viewing Your Update History

After an update was successfully installed or failed to install, the update is displayed on the **View Update History** page.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Auto Update**.

3. Click **View Update History**.

4. Optional: Using the **Search by Name** field, you can type a keyword and then press Enter to locate a specific update by name.

5. To investigate a specific update, select the update.

A description of the update and any installation error messages are displayed in the right pane of the View Update History page.

## Restoring Hidden Updates

You can remove updates from the **Check for Updates** page. You can view and restore the hidden updates on the **Restore Hidden Updates** page.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Auto Update**.
3. Click **Restore Hidden Updates**.
4. Optional: To locate an update by name, type a keyword in the **Search by Name** text box and press Enter.
5. Select the hidden update that you want to restore.
6. Click **Restore**.

## Viewing the Autoupdate Log

The autoupdate log contains the most recent automatic update that was run on your system.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Auto Update**.
3. On the navigation menu, click **View Log**.

### RELATED DOCUMENTATION

Manual Updates | **140**

Configuring System settings | **144**

IF-MAP Server Certificates | **153**

# Manual Updates

IN THIS SECTION

If your deployment includes a JSA Console that is unable to access the Internet, or you want to manually manage updates to your system, you can manage the update process manually by setting up a JSA update server.

The autoupdate package includes all files necessary to manually set up an update server in addition to the necessary system configuration files for each update. After the initial setup, you only need to download and uncompress the most current autoupdate package to manually update your configuration.

## Configuring an Update Server

Configure an Apache server as the update server for your JSA deployment.

Download the autoupdate package from Juniper Customer Support.

1. Access your Apache server and create an update directory named **software/jsa/autoupdate/**.

   By default, the update directory is in the web root directory of the Apache server. You can place the directory in another location if you configure JSA accordingly.

2. Optional: Create an Apache user account and password to be used by the update process.

3. Save the autoupdate package file on your Apache server in the **software/jsa/autoupdate/** directory that you created.

4. On the Apache server, type the following command to uncompress the autoupdate package.

   ```
   tar -zxf updatepackage-[timestamp].tgz
   ```

5. On the navigation menu
   (

≡
), click **Admin**.

6. In the **System Configuration** section, click **Auto Update**.

7. Click **Change Settings**, and click the **Advanced** tab.

8. In the **Server Config** pane parameters, configure the settings for the Apache server.

   a. In the **Web Server** field, type the address or directory path of your Apache server.

      If the Apache server runs on non-standard ports, add the port number to the end of the address. For example, type **https://<server name/ip>:<server port>**.

   b. In the **Directory field**, type the directory location where the web server stores the updates.

      The default directory is **software/jsa/autoupdate/**.

   c. Optional: If the application server uses a proxy server to connect to the Internet, type the URL in the **Proxy Server** field.

   d. Optional: If you are using an authenticated proxy, type the credentials in the **Proxy Username** and **Proxy Password** fields.

9. Click **Save**.

10. On the **Admin** tab, click **Deploy changes**.

11. Using SSH, log in to JSA as the root user.

12. Type the following command to configure the user name that you set for your Apache server:

    ```
    /opt/qradar/bin/UpdateConfs.pl -change_username <username>
    ```

13. Type the following command to configure the password that you set for your Apache server:

    ```
    /opt/qradar/bin/UpdateConfs.pl -change_password <password>
    ```

14. To test the update server, type the following command as a single line of text in the command line interface.

    ```
    interface.
    wget -q -O- --no-check-certificate
    https://<your update server>/<directory path to updates>/manifest_list
    ```

15. Type the user name and password.

## Configuring the JSA Console As the Update Server

To streamline your maintenance process, you can configure your JSA Console to be your update server so that JSA updates are automatically downloaded to the Console.

1. Download the auto update package from Juniper Customer Support.

2. Save the auto update package file in the **/tmp/** directory on your JSA Console.

   The size of the auto update file is approximately 2 - 5 GB.

3. Log in to JSA as the root user.

4. Type the following command to create the auto update directory:

   **mkdir -p /opt/qradar/www/software/jsa/autoupdate/**

5. To verify that the Console has enough space for the auto update file, type the following command:

   **df -h /opt/qradar/**

   If you do not have enough space in the current directory, you can create another directory structure in the **/store** directory, such as **/store/downloads**

   a. Create a symbolic link to **/opt/qradar/www/software/jsa/autoupdate/** by typing the following command:

      **ln -s /store/downloads/autoupdates /opt/qradar/www/software/jsa/autoupdate/**

   b. To verify that the symbolic link was created properly, type the following command:

      **touch /store/downloads/testfile**

   c. Confirm that the test file value is created in the **/opt/qradar/www/software/jsa/autoupdate/** directory by typing the following command:

      **ls /opt/qradar/www/software/jsa/autoupdate/**

6. Copy the **autoupdates-<version>.tqz>** file from the **/tmp/** directory to the JSA Console, and place it in the **/opt/qradar/www/software/jsa/autoupdate/** directory or the symbolic link directory that you created in "5 (a)" on page 142.

7. On the JSA Console, type the following commands to extract the auto update package:

   **cd /opt/qradar/www/software/jsa/autoupdate/**

   **tar -zxf /tmp/<name_of_autoupdate_file>**

8. Log in to JSA.

9. On the navigation menu ( ≡ ), click **Admin**.

10. In the **System Configuration** section, click **Auto Update**.

11. Click **Change Settings**, and select **Advanced** tab.

12. In the Directory field, type **software/jsa/autoupdate/**

13. In the **Web Server** field, type **https://<console_IP_address_or_hostname>**

14. Click **Save**.

Your JSA autoupdate directory is created, the autoupdate package downloaded, and the configuration for autoupdates is complete.

## Downloading Updates to the Update Server

You can download updates from Juniper Customer Support to your update server.

You must configure your update server and set up JSA to receive updates from the update server.

1. Download the autoupdate package from Juniper Customer Support.

2. Save the autoupdate package file on your update server in the **software/jsa/autoupdate/** directory that you created.

3. Type the following command to uncompress the autoupdate package:

   ```
   tar -zxf autoupdate-[timestamp].tgz
   ```

4. Log in to JSA as the root user.

5. Type the following command to test your update server:

   ```
   wget https://<your_update_server>/<directory_path_to_updates>/manifest_list
   ```

6. Type the user name and password of your update server.

### RELATED DOCUMENTATION

Configuring System settings | 144

IF-MAP Server Certificates | 153

# Configuring System settings

IN THIS SECTION

System settings specify how your JSA system components are configured for normal operation.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **System Settings**.
3. Configure the system settings. Click the **Help** button to see setting descriptions.
4. Click **Save**.
5. On the **Admin** tab menu, select **Advanced** >**Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

## Enhancing the Right-click Menu for Event and Flow Columns

You can add more actions to the right-click options that are available on the columns in the **Log Activity** table or the **Network Activity** table. For example, you can add an option to view more information about the source IP or destination IP.

**NOTE**: The right-click feature is not available on fields in the **Event Information** window.

You can pass any data that is in the event or flow to the URL or script.

1. Using SSH, log in to the JSA console appliance as the root user.

2. Go to the **/opt/qradar/conf** directory and create a file that is named **arielRightClick.properties**.

3. Edit the **/opt/qradar/conf/arielRightClick.properties** file. Use the following table to specify the parameters that determine the options for the right-click menu.

**Table 20: ArielRightClick.properties File Parameters**

| Parameter | Requirement | Description | Example |
|---|---|---|---|
| **pluginActions** | Required | Indicates either a URL or script action. | |
| **arielProperty** | Required | Specifies the column, or Ariel field name, for which the right-click menu is enabled. | **sourceIP**<br>**sourcePort**<br>**destinationIP**<br>**qid** |
| **text** | Required | Specifies the text that is displayed on the right-click menu. | Google search |
| **useFormattedValue** | Optional | Specifies whether formatted values are passed to the script.<br>Set to true to ensure that the formatted value for attributes, such as `username` and `payload`, are passed. Formatted values are easier for administrators to read than unformatted values. | If the parameter is set to true for the event name (QID) property, the event name of the QID is passed to the script.<br>If the parameter is set to false, the raw, unformatted QID value is passed to the script. |

Table 20: ArielRightClick.properties File Parameters *(Continued)*

| Parameter | Requirement | Description | Example |
|---|---|---|---|
| **url** | Required to access a URL | Specifies the URL, which opens in a new window, and the parameters to pass to the URL.<br>Use the format: $*Ariel_Field Name*$ | **sourceIPwebUrlAction.url= http://www.mywebsite.com? q= $sourceIP$** |
| **command** | Required if the action is a command | Specifies the absolute path of the command or script file. | **destinationPortScriptAction.command=/bin/echo** |
| **arguments** | Required if the action is a command | Specifies the data to pass to the script.<br>Use the following format: $*Ariel_Field Name*$ | **destinationPortScriptAction.arguments=$qid$** |

For each of the key names that are specified in the *pluginActions* list, define the action by using a key with the format *key name*, *property*.

4. Save and close the file.

5. Log in to the JSA user interface.

6. On the navigation menu ( ≡ ), click **Admin**.

7. Select **Advanced** >**Restart Web Server**.

The following example shows how to add *Test URL* as a right-click option for source IP addresses.

**pluginActions=sourceIPwebUrlAction**

**sourceIPwebUrlAction.arielProperty=sourceIP sourceIPwebUrlAction.text=Test URL sourceIPwebUrlAction.url=http://www.mywebsite.com?q=$sourceIP$**

The following example shows how to enable script action for destination ports.

**pluginActions=destinationPortScriptAction**

**destinationPortScriptAction.arielProperty=destination Port destinationPortScriptAction.text=Test Unformatted Command destinationPortScriptAction.useFormattedValue=false destinationPortScriptAction.command=/bin/echo destinationPortScriptAction.arguments=$qid$**

The following example shows adding several parameters to a URL or a scripting action.

**pluginActions=qidwebUrlAction,sourcePortScriptAction**

**qidwebUrlAction.arielProperty=qid,device,eventCount qidwebUrlAction.text=Search on Google qidwebUrlAction.url=http://www.google.com?q=$qid$-$device$-$eventCount$**

**sourcePortScriptAction.arielProperty=sourcePort sourcePortScriptAction.text=Port Unformatted Command sourcePortScriptAction.useFormattedValue=true sourcePortScriptAction.command=/bin/echo sourcePortScriptAction.arguments=$qid$-$sourcePort$-$device$-$CONTEXT$**

## Asset Retention Values Overview

Additional information for the period, in days, that you want to store the asset profile information.

- Assets are tested against the retention thresholds at regular intervals. By default, the cleanup interval is 12 hours
- All specified retention periods are relative to the last seen date of the information, regardless of whether the information was last seen by a scanner or passively observed by the system.
- Asset information is deleted as it expires, meaning that following a cleanup interval, all asset information within its retention threshold remains.
- By default, assets that are associated with un-remediated vulnerabilities (as detected by JSA Vulnerability Manager or other scanner) are retained.
- Assets can always be deleted manually through the user interface.

**Table 21: Asset Components**

| Asset component | Default retention (in days) | Notes |
|---|---|---|
| IP Address | 120 days | By default, user-supplied IP Addresses are retained until they are deleted manually. |

Table 21: Asset Components *(Continued)*

| Asset component | Default retention (in days) | Notes |
|---|---|---|
| MAC Addresses (Interfaces) | 120 days | By default, user-supplied interfaces are retained until they are deleted manually. |
| DNS and NetBIOS Hostnames | 120 days | by default, user-supplied hostnames are retained until they are deleted manually. |
| Asset Properties | 120 days | By default, user-supplied IP Addresses are retained until they are deleted manually.<br>The asset properties this value can affect are **Given Name**, **Unified Name**, **Weight**, **Description**, **Business Owner**, **Business Contact**, **Technical Owner**, **Technical Contact**, **Location**, **Detection Confidence**, **Wireless AP**, **Wireless SSID**, **Switch ID**, **Switch Port ID**, **CVSS Confidentiality Requirement**, **CVSS Integrity Requirement**, **CVSS Availability Requirement**, **CVSS Collateral Damage Potential**, **Technical User**, **User Supplied OS**, **OS Override Type**, **OS Override Id**, **Extended**, **Legacy (Pre-2014.x) Cvss Risk**, **VLAN**, and **Asset Type**. |
| Asset Products | 120 days | By default, user-supplied products are retained until they are deleted manually.<br>Asset products include Asset OS, Asset Installed Applications, and products that are associated with open asset ports |
| Asset "Open" Ports | 120 days | |

**Table 21: Asset Components *(Continued)***

| Asset component | Default retention (in days) | Notes |
|---|---|---|
| Asset netBIOS Groups | 120 days | NetBIOS groups are seldom used, and more customers may not be aware of their existence. In the case where they are used, they are deleted after 120 days. |
| Asset Client Application | 120 days | Client Applications are not yet leveraged in the user interface. This value can be ignored. |
| Asset Users | 30 days | |

## Adding or Editing a JSA Login Message

Create a new login message or edit an existing login message on your JSAConsole.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **User Management**.
3. Click **Authentication**, and then click **General Authentication Settings**.
4. To edit the login message, click **Login Page** and then set **Login Message** to **On**.
   a. Type your message in the Edit Login Message window.
   b. To force users to consent to the login message before they can log in, set **Require explicit consent of this message for login** to **On**.
   c. Click **Save Settings**.

      The login message is saved in the **opt/qradar/conf/LoginMessage.txt** file.

**NOTE**: You can also upload the **LoginMessage.txt** file to the **opt/qradar/conf/** directory.

5. On the **Admin** tab, click **Deploy Changes**.

6. To see your changes, log out of JSA.

## Turning on and Configuring Rule Performance Visualization

Use the Custom Rule Settings feature to turn on and configure metrics for rule performance analysis. Rule performance visualization extends the current logging around performance degradation and the expensive custom rules in the JSA pipeline. With rule performance visualization, you can determine the efficiency of rules in the JSA pipeline directly from the Rules page.

After you turn on rule performance visualization, the metrics remain blank unless an event or flow performance issue occurs.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **System Settings**.

3. On the **System Settings** page, click **Advanced**.

4. Configure the **Custom Rule Settings**.

   **Table 22: Custom Rule Settings**

| Setting | Description |
|---|---|
| **Enable Performance Analysis** | Enable cost performance analysis tracking for custom rules. The default is False. |
| **Reset Metrics on Rule Change** | Enable the reset of the rule performance analysis metrics when a rule is modified. The default is True.<br>**NOTE**: To reset metrics on a rule, edit the rule, and then save it. The metrics are cleared for the rule that you modified. |

Table 22: Custom Rule Settings *(Continued)*

| Setting | Description |
|---|---|
| **Performance Analysis Upper Limit** | The upper threshold (in EPS or FPS) that is used to determine the performance bar value for a rule.<br>• If the throughput for a rule drops below this limit and is above the Performance Analysis Lower Limit, the performance is displayed as two orange bars.<br>• If the throughput for a rule is above this limit, the performance is displayed as three green bars.<br>The default is 50,000. |
| **Performance Analysis Lower Limit** | The lower threshold (in EPS or FPS) used to determine the performance bar value for a rule. If the throughput for a rule drops below this limit, the performance is displayed as one red bar.<br>The default is 12,500. |

5. Click **Save**.

6. On the navigation menu (≡), click **Admin**.

7. Click **Deploy Changes.**

When rule performance visualization is turned on, the **Performance** column is added to the **Rules** page. The **Performance** column on the **Rules** page is blank until a performance issue occurs in the custom rule engine.

For more information about Rule performance visualization, see the *Juniper Secure Analytics User Guide.*

## Troubleshooting Rule Performance Visualization

This reference provides troubleshooting information for rule performance visualization.

**Why am I not seeing metrics for a rule?**

Table 23: Rule Metrics Issues

| Issue | Solution |
|---|---|
| Performance Analysis is not enabled. | Deploy the changes |
| Metrics do not display for rules that are not enabled. | Works as designed. Metrics display only for enabled rules. |
| Metrics do not display for offense rules. | Works as designed. Metrics are collected only for all event, common, and flow rules. |
| Metrics do not display for a rule. | The rule might be recently modified, which resets the metrics. The metrics are cleared for the rule that you modified. If you don't want the metric to be reset when a rule is resaved, disable **Reset Metrics on Rule Change**. |

### Why would I want to change the upper and lower thresholds?

Whether you would want to change the upper and lower threshold limits, depends on what you deem to be an acceptable event per second (EPS) or flows per second (FPS) throughput for your rules. You might want to start with your general system EPS or FPS throughput. Increase your upper threshold limit by a few thousand, and decrease your lower threshold limit by a few thousand. When you change these settings, keep in mind your license and hardware throughput limitations. Your upper limit doesn't need to go above your license or hardware capacity. Typically, as you use this feature to tune your rules, you might want to update the lower limit with a slightly higher value so that you can focus on the under-performing rules.

Example:

- General EPS load for system: 5,000 EPS
- Upper Limit: 8,000 EPS
- Lower Limit: 2,000 EPS

Rules that can process 8,001 EPS or more display three green bars. Rules that can process only 1,999 EPS or lower display 1 red bar. All rules between these ranges are marked with two orange bars. After you tune all of your rules that display red bars and only the orange and green bars display, you can increase the lower limit to 3,000 EPS.

### Why does a disabled rule show as expensive?

When rule performance is turned on, previous values might display for disabled rules, which might cause the rule to show as expensive.

If you selected **Reset Metrics on Rule Change** when you enabled rule performance, reset the metrics for the rule by editing the rule, and then saving it. The metrics are cleared for the rule that you modified.

You can view the metrics for a rule from the **Rules** page when you move the mouse pointer over the colored bars in the **Performance** column, and in the Performance Analysis textbox, which is in the lower-right corner of the **Rules** page. You can also view the metrics for a rule in the **Rule Wizard** when you edit a rule. The timestamp in the **Performance Analysis** textbox shows when the metrics for the rule were updated.

For more information about editing rules, see the *Juniper Secure Analytics User Guide.*

# IF-MAP Server Certificates

**IN THIS SECTION**

The Interface For Metadata Access Points (IF-MAP) rule response enables the JSA console to publish alert and offense data that is derived from events, flows, and offenses to an IF-MAP server.

Before you can configure IF-MAP authentication on the **System Settings** window, you must configure your IF-MAP server certificate.

## Configuring IF-MAP Server Certificate for Basic Authentication

This task provides instruction for how to configure your IF-MAP certificate for basic authentication.

Contact your IF-MAP server administrator to obtain a copy of the IF-MAP server public certificate. The certificate must have the **.cert** file extension.

1. Using SSH, log in to JSA as the root user.
2. Copy the certificate to the **/opt/qradar/conf/trusted_certificates** directory.

# Configuring IF-MAP Server Certificate for Mutual Authentication

Mutual authentication requires certificate configuration on your JSA console and on your IF-MAP server.

This task provides steps to configure the certificate on your JSA console. For assistance configuring the certificate on your IF-MAP server, contact your IF-MAP server administrator.

Contact your IF-MAP server administrator to obtain a copy of the IF-MAP server public certificate. The certificate must have the **.cert** file extension.

1. Using SSH, log in to JSA as the root user.

2. Access the certificate to the **/opt/qradar/conf/trusted_certificates** directory

3. Copy the SSL intermediate certificate and SSL Verisign root certificate to your IF-MAP server as CA certificates. For assistance, contact your IF-MAP server administrator.

4. Type the following command to create the Public-Key Cryptography Standards file with the .pkcs12 file extension:

   **openssl pkcs12 -export -inkey *<private_key>* -in *<certificate>* -out *<pkcs12_filename.pkcs12>* -name "IFMAP Client"**

5. Type the following command to copy the **pkcs12** file to the **/opt/qradar/conf/key_certificates** directory:

   **cp *<pkcs12_filename.pkcs12>* /opt/qradar/conf/key_certificates**

6. Create a client on the IF-MAP server with the certificate authentication and upload the SSL certificate. For assistance, contact your IF-MAP server administrator.

7. Type the following command to change the permissions of the directory:

   **chmod 755 /opt/qradar/conf/trusted_certificates chmod 644 /opt/qradar/conf/trusted_certificates/*.cert**

8. Type the following command to restart the Tomcat service:

   **systemctl restart tomcat**

RELATED DOCUMENTATION

IPv6 Addressing in JSA Deployments | **162**

Advanced Iptables Rules Examples | **164**

# SSL Certificates

IN THIS SECTION

Secure Sockets Layer (SSL) is an industry standard security protocol is used by websites to protect online transactions. It provides communication privacy so that client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery. To generate an SSL link, a web server requires an SSL certificate. SSL certificates are issued by internal or trusted third-party certifying authorities.

Browsers and operating systems include a preinstalled list of trusted certificates, which are installed in the Trusted Root Certification authorities store.

**Self-signed certificates** - A self-signed certificate provides basic security, enabling data encryption between the user and the application. Because self-signed certificates cannot be authenticated by any existing known root certificate authorities, users are warned about this unknown certificate and must accept it to proceed.

**Internal CA signed certificates** - Organizations that have their own internal root certificate authority (CA) can create a certificate by using that internal CA. This certificate is supported by JSA, and the internal root CA is also imported into the JSA environment.

**Public CA / Intermediate CA signed** - Certificates that are signed by known public CAs and intermediate certificates are supported by JSA.

Public signed certificates can be used directly in JSA, and certificates that are signed with Intermediate CA are installed by using both the signed certificate and the intermediate certificate to provide valid certificate functions.

**NOTE**: An intermediate certificate is commonly used by organizations that create multiple SSL keys in their environment, and want to have them signed by a known commercial certificate vendor. When they use the intermediate key, they can then create sub-keys from this intermediate key. When this configuration is used,JSA must be configured with both the intermediate certificate and the host SSL certificate so that connections to the host can verify the full certificate path.

## SSL Connections Between JSA Components

To establish all internal SSL connections between components, JSA uses the web server certificate that is preinstalled on theJSA Console.

All trusted certificates for JSA must meet the following requirements:

- The certificate must be an X.509 certificate and have PEM base64 encoding.
- The certificate must have a `.cert`, `.crt`, `.pem`, or .der file extension.
- Keystore files that contain certificates must have the `.truststore` file extension.
- The certificate file must be stored in the **/opt/qradar/conf/trusted_certificates** directory.

## Creating an SSL Certificate Signing Request with 2048-bit RSA Keys

1. Use SSH to log in to the JSA Console.
2. Generate a private key file by using the following command:

   ```
   openssl genrsa -out qradar.key 2048
   ```

   **NOTE**: Do not use the private encryption options, because they can cause compatibility issues.

   The qradar.key file is created in the current directory. Keep this file to use when you install the certificate.
3. Generate the certificate signing request (CSR) file.

The `qradar.csr` file is used to create the SSL Certificate, with an internal CA or commercial certificate authorities. Run the following command, and provide necessary information as prompted:

`openssl req -new -key qradar.key -out qradar.csr`

Example output:

```
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) [XX]:US
State or Province Name (full name) []:MyState
Locality Name (eg, city) [Default City]:MyCity
Organization Name (eg, company) [Default Company Ltd]:MyCompany
Organizational Unit Name (eg, section) []:MyCompanyOrg
Common Name (eg, your name or your server's hostname)
[]:qradar.mycompany.com
Email Address []:username@example.com
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

4. If you want to verify the information in the CSR before you send it, type the following command:

   ```
   openssl req -noout -text -in qradar.csr
   ```

   If incorrect information was entered, run the OpenSSL command again to re-create the CSR file.

5. Use the Secure File Transfer Protocol or another program to securely copy the CSR file to your computer.

6. Submit the CSR to your internal or commercial certificate authority for signing according to their instructions.

   **NOTE**: The CSR is identified as a certificate in Apache format.

## Creating a Multi-Domain (SAN) SSL Certificate Signing Request

1. Use SSH to log in to the JSA Console.

2. Create and save a `sancert.conf` configuration file containing the following information:

   ```
   [ req ] default_bits = 2048 # RSA key size encrypt_key = no # Protect private key default_md = sha256 # MD to
   use utf8 = yes # Input is UTF-8 string_mask = utf8only # Emit UTF-8 strings prompt = no # Prompt for DN
   distinguished_name = server_dn # DN template req_extensions = server_reqext # Desired extensions [ server_dn ]
   countryName = <country_or_region_code> # ISO 3166 stateOrProvinceName = <state_or_province> localityName =
   <city_or_locality> organizationName = <organization_name> organizationalUnitName = <organizational_unit_name
   commonName = <common_name> # Should match a SAN under alt_names [ server_reqext ] basicConstraints = CA:FALSE
   keyUsage = critical,digitalSignature,keyEncipherment extendedKeyUsage = serverAuth subjectKeyIdentifier = hash
   subjectAltName = @alt_names [alt_names] DNS.1 = qradar.example.com #Example DNS.2 = console.example.com
   #Example IP.3 = 192.0.2.0 #Example
   ```

3. Generate a private key and public certificate signing request (CSR) pair by using the following command:

   ```
   openssl req -new -nodes -sha256 -out <csr_filename>.csr -config sancert.conf -keyout <privatekey_filename>.key
   ```

   The CSR file is used to create the SSL certificate, with either an internal CA or commercial certificate authorities The key file is created in the current directory. Keep this file to use when you install the certificate.

4. If you want to verify the information in the CSR before you send it, type the following command:

   ```
   openssl req -noout -text -in <csr_filename>.csr
   ```

   If incorrect information was entered, update the `sancert.conf` configuration file and do "3" on page 158 again.

5. Use the Secure File Transfer Protocol or another program to securely copy the CSR file to your computer.

6. Submit the CSR to your internal or commercial certificate authority for signing, according to their instructions.

   > **NOTE**: The CSR is identified as a certificate in Apache format.

## Using Certificates That are Signed by an Internal Certificate Authority

If the certificate is issued by an internal certificate authority and not a commercial certificate provider, JSA must be updated to include the internal root certificate into the local certificate store for proper certificate validation. Root verification certificates are automatically included with the operating system.

Follow these steps to update the trust anchors root certificate store in RHEL on the JSA console and all JSA hosts.

1. Copy the CA's root certificate to **/etc/pki/ca-trust/source/anchors/** on the JSA console.
2. Run the following commands at the SSH command line on the console:

```
/opt/qradar/support/all_servers.sh -p /etc/pki/ca-trust/source/anchors/<root_certificate> - r /etc/pki/ca-trust/source/anchors
```

```
/opt/qradar/support/all_servers.sh -C update-ca-trust
```

## Installing a New SSL Certificate

By default, JSA is configured with a Security Sockets Layer (SSL) certificate that is signed by an internal CA. When you log in to the Console for the first time, you are prompted with a warning message that your connection is not secure or is not private. You can replace the SSL certificate with your own selfsigned certificate, a private certificate authority (CA) signed certificate, or a public CA signed certificate.

You must have the following information:

- The newly signed SSL CertificateFilefrom either your internal CA, or a public one.
- The `qradar.key` private key to generate the CSR file.
- An intermediate certificate, if used by your certificate provider.

> **NOTE**: If an intermediate certificate is used, run the install-ssl-cert.sh command with the -i flag to install both the new certificate and the intermediate certificate. When used, it prompts for three file paths:
>
> - `SSLCertifficateFile`
> - `SSLIntermediateCertificateFile`

- `SSLCertificateKeyFile`

If you use a DER certificate, you must convert it to a PEM certificate by typing the following command line:

```
openssl x509 -in <cert>.der -inform der -outform pem -out <cert>.pem
```

1. Use SSH to log in to the JSA Console as the root user. Install the certificate by entering the following command:

   ```
   /opt/qradar/bin/install-ssl-cert.sh
   ```

   a. At the Path to Public Key File (SSLCertificateFile) prompt, enter the path to the Public Key File. For example:

   ```
   /root/new.certs/cert.cert
   ```

   b. At the Path to Private Key File (SSLCertificateKeyFile) prompt, enter the path to the Private Key File. For example:

   ```
   /root/new.certs/qradar.key
   ```

   Example output:

   ```
   You have specified the following:
   SSLCertificateFile of /root/updated.certs/cert.cert
   SSLCertificateKeyFile of /root/updated.certs/qradar.key
   Re-configure Apache now (includes restart of httpd) (Y/[N])? y
   Backing up current SSL configuration ... (OK)
   Installing user SSL certificate ... (OK)
   Reloading httpd configuration:
   - Restarting httpd service ... (OK)
   Restarting services:
   - Stopping hostcontext ... (OK)
   - Restarting Tomcat ... (OK)
   - Starting hostcontext ... (OK)

   Updating deployment:
   - Copying certificate to managed hosts
   * 192.0.2.0 ...... (OK)
   - Restarting hostcontext on managed hosts
   ```

```
* 192.0.2.0 ...... (OK)
The event collection service must be restarted if WinCollect is used in your environment.
Restart the event collection service now (y/[n])? y
- Restarting ecs-ec-ingress on managed hosts
* 192.0.2.0 ...... (OK)
- Restarting ecs-ec-ingress on console ... (OK)
Fri Jan 17 10:33:42 EST 2020 [install-ssl-cert.sh] OK: Install SSL Cert
Completed
[root@qavm215 ~]#
```

**NOTE**: Data collection for events and flows stops while services are restarted.

2. If you are installing a certificate that was not generated by JSA or reinstalling an overwritten certificate that was not generated by JSA, do the following:

   a. Disable the CA framework from automatically replacing the certificate. Run the following command to disable the certificate check so that JSA does not overwrite the certificate:

   ```
   sed -e "s@\"CertSkip\":[ \t]*\".*\"@\"CertSkip\": \"true\"@" -i /opt/qradar/ca/conf.d/httpd.json
   ```

   b. Disable the certificate monitoring through modifying the **/opt/qradar/ca/conf.d/httpd.json** with the line

   ```
   "CertMonitorThreshold": 0
   ```

If the `install-ssl-cert.sh` script finished with the OK: Install SSL Cert Completed message, then the certificate was installed successfully. If you answered y (yes) to the prompt to reconfigure Apache, then you're done. Otherwise, you must deploy the full configuration. On the navigation menu, click **Admin**, then click **Advanced** > **Deploy Full Configuration**.

## Reverting to Certificates that are Generated by the JSA Local CA

If you have issues with your certificate, such as an incorrect name or IP address, the expiration date passed, or the IP or hostname on your console changed, follow these steps to generate certificates that are signed by the JSA local certificate authority.

1. Back up the certificates that were installed previously that are not working.

   Existing certificates are detected and reported when you run certificate generation, which can cause the generation process to stop.

   ```
   mkdir /root/backup.certs/cp /etc/httpd/conf/certs/cert.* /root/backup.certs/
   ```

2. Update the following items in the **/opt/qradar/ca/conf.d/httpd.json** file:

   - Set **CertMonitorThreshold** back to its original value. If the original value is not known, remove from the file so that the defaults are used.

   - Set **CertSkip** to false.

3. Run the **/opt/qradar/ca/bin/install_qradar_ssl_cert.sh** command to generate new certificates.

RELATED DOCUMENTATION

Advanced Iptables Rules Examples | 164
Data Retention | 169

# IPv6 Addressing in JSA Deployments

IN THIS SECTION

IPv4 and IPv6 addressing is supported for network connectivity and management of JSA software and appliances. When you install JSA, you are prompted to specify whether your Internet Protocol is IPv4 or IPv6.

## JSA Components That Support IPv6 Addressing

The following JSA components support IPv6: addressing.

- **Network Activity tab** -- Because **IPv6 Source Address** and **IPv6 Destination Address** are not default columns, they are not automatically displayed. To display these columns, you must select them when you configure your search parameters (column definition).

To save space and indexing in an IPv4 or IPv6 source environment, extra IP address fields are not stored or displayed. In a mixed IPv4 and IPv6 environment, a flow record contains both IPv4 and IPv6 addresses.

IPv6 addresses are supported for both packet data, including sFlow, and NetFlow V9 data. However, older versions of NetFlow might not support IPv6.

- **Log Activity tab** -- Because **IPv6 Source Address** and **IPv6 Destination Address** are not default columns, they are not automatically displayed. To display these columns, you must select them when you configure your search parameters (column definition).

  DSMs can parse IPv6 addresses from the event payload. If any DSM cannot parse IPv6 addresses, a log source extension can parse the addresses. For more information about log source extensions, see the *Juniper Secure Analytics Log Sources User Guide*.

- **Searching, grouping, and reporting on IPv6 fields** -- You can search events and flows by using IPv6 parameters in the search criteria.

  You can also group and sort event and flow records that are based on IPv6 parameters.

  You can create reports that are based on data from IPv6-based searches.

- **Custom rules** --The following custom rule to support IPv6 addressing was added: **SRC/DST IP = IPv6 Address**

  IPv6-based building blocks are available in other rules.

- **Device support modules (DSMs)** -- DSMs can parse IPv6 source and destination address from event payloads.

## Deploying JSA in IPv6 Environments

To log in to JSA in an IPv6 environment, wrap the IP address in square brackets:

**https://**[ ***<IP Address>*** ]

Both IPv4 and IPv6 environments can use a hosts file for address translation. In an IPv6 environment, the client resolves the Console address by its host name. You must add the IP address of the IPv6 console to the **/etc/hosts** file on the client.

Flow sources, such as NetFlow and sFlow, are accepted from IPv4 and IPv6 addresses. Event sources, such as syslog and SNMP, are accepted from IPv4 and IPv6 addresses. You can disable superflows and flow bundling in an IPv6 environment.

## IPv6 Addressing Limitations

When JSA is deployed in an IPv6 environment, the following limitations are known:

- The network hierarchy is not updated to support IPv6.

  Some parts of the JSA deployment, including surveillance, searching, and analysis, do not take advantage of the network hierarchy. For example, within the Log Activity tab, you cannot search or aggregate events By Network

- No IPv6-based asset profiles.
- Asset profiles are created only if JSA receives events, flows, and vulnerability data for IPv4 hosts.
- No host profile test in custom rules for IPv6 addresses.
- No specialized indexing or optimization of IPv6 addresses.
- No IPv6-based sources and destinations for offenses

RELATED DOCUMENTATION

Advanced Iptables Rules Examples | 164
Data Retention | 169

# Advanced Iptables Rules Examples

IN THIS SECTION

You can configure your iptables rules to better control access to JSA, restrict inbound data sources, and redirect traffic. The following examples can help you to gain better insight to your network, by manually adjusting your iptables.

## Blocking Access to SSH with Iptables

Consoles and unmanaged hosts allow SSH from any inbound request. When a host is added to the deployment, the managed hosts allow SSH access from the JSA console, and the console keeps port 22 open for inbound connections. You can limit the inbound connections on port 22 by modifying a host's iptables rules.

You can block SSH access from other managed hosts on your console, which can break encrypted connections.

```
-A INPUT -i eth0 -m state --state NEW -m tcp -p tcp --dport 22 -s 10.100.50.41 -j ACCEPT
-A INPUT -i eth0 -m state --state NEW -m tcp -p tcp --dport 22 -s 10.100.50.59 -j ACCEPT
-A INPUT -i eth0 -m state --state NEW -m tcp -p tcp --dport 22 -j DROP
```

## Enabling ICMP to JSA Systems

You can enable ping responses from your JSA system by adding the following rule to the **/opt/qradar/conf/iptables.pre** file.

```
-A INPUT -p icmp -j ACCEPT
```

Run the following script to create an entry in the `/etc/sysconfig/iptables` file.

**NOTE**: You can limit this rule to a specific host by adding the `-s source.ip.address` field.

## Blocking Unwanted Data Sources

You can block out a data source such as a log source or a netflow data source, for a short time, rather than disabling the original device. To block a particular host, you can add an entry similar to the following to **/opt/qradar/conf/iptables.pre**.

Block a netflow from the router:

-A INPUT -p udp -s <*IP Address*> --dport 2055 -j REJECT

Block a syslog from another source:

-A INPUT -p tcp -s <*IP Address*> --dport 514 -j REJECT -A INPUT -p udp -s <*IP Address*> --dport 514 -j REJECT

Block a syslog from a specific subnet:

-A INPUT -p tcp -s <*IP Address*> --dport 514 -j REJECT -A INPUT -p udp -s <*IP Address*> --dport 514 -j REJECT

## Redirecting Iptables to Syslog Ports

You can redirect syslog traffic on non-standard ports into port 514 on a JSA Event Collector.You can use the following steps to enable an iptables rule to redirect the alternative port back into 514 on the Event Collector.

1. Enable the NAT option in the Linux kernel by adding or updating the following line in the **/etc/sysctl.conf** file.

```
net.ipv4.ip_forward = 1
```

**NOTE**: For changes to take effect to the NAT rule, you might need to restart your service.

2. Enable ipforwarding in the current active kernel.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

3. Add the following lines to the **/opt/qradar/conf/iptables-nat.post**. Enter the port number that you want to redirect as the <*portnumber*>.

   ```
   -A PREROUTING -p udp --dport <portnumber> -j REDIRECT --to-ports 514 -A PREROUTING -p tcp
   --dport <portnumber> -j REDIRECT --to-ports 514
   ```

4. Enter the following command to rebuild your iptables.

   ```
   /opt/qradar/bin/iptables_update.pl
   ```

5. Verify the redirection by typing the following command.

   ```
   iptables -nvL -t nat
   ```

   The following code is an example of what the output might look like.

```
Chain PREROUTING (policy ACCEPT 140 packets, 8794 bytes) pkts bytes target prot opt in out source destination
0 0 REDIRECT udp -- * * 0.0.0.0/0 0.0.0.0/0 udp dpt:10529 redir ports 514 0 0 REDIRECT tcp -- * * 0.0.0.0/0
0.0.0.0/0 tcp dpt:10529 redir ports 514 Chain POSTROUTING (policy ACCEPT 207 packets, 25772 bytes) pkts bytes
target prot opt in out source destination Chain OUTPUT (policy ACCEPT 207 packets, 25772 bytes) pkts bytes
target prot opt in out source destination
```

## Redirecting Inbound Syslog Traffic

You can use your JSA console as a syslog message gateway to redirect inbound events, by configuring rules in iptables.

1. Enable the forwarding rule for a log source on your Event Collector.

2. Set the forwarding destination for the TCP syslog to be the console IP address on port 7780.

3. From the command line of the console, add the following iptables rule to redirect to another host.

   ```
   iptables -I OUTPUT --src 0/0 --dst 153.2.200.80 -p tcp --dport 7780 -j REDIRECT --to-ports 514
   ```

## Configuring Iptables Rules

Access to the JSA network services is controlled first on hosts with iptables. The iptables rules are adjusted and configured based on the requirements of the deployment. Ports for Ariel searching, streaming, and times when you are using encryption (tunneling) can update various iptables rules.

You can configure and check iptables rules for IPv4 and IPv6. The following procedure indicates how you can tune your iptables manually.

1. Log in to JSA as the root user by using SSH.

   Login: <**root**>

   Password: <**password**>

2. Type the following command to edit the pre rules iptables file:

   IPv4:

   **vi /opt/qradar/conf/iptables.pre**

   IPv6:

   **vi /opt/qradar/conf/ip6tables.pre**

   The iptables.pre configuration file is displayed.

3. Type the following command to edit the post rules iptables file:

   IPv4:

   **vi /opt/qradar/conf/iptables.post**

   IPv6:

   **vi /opt/qradar/conf/ip6tables.post**

   The iptables.post configuration file is displayed.

4. Add the following rule for JSA to access a specific port number, where *portnumber* is the port number:

   To accept UDP traffic for a specific port input:

   **-A INPUT -m udp -p udp --dport <*portnumber*> -j ACCEPT**

   To accept TCP traffic for a specific port input:

   **-A INPUT -m state --state NEW -m tcp -p tcp --dport <*portnumber*> -j ACCEPT**

5. Save your iptables configuration.

6. Run the following script to propagate the changes:

   **/opt/qradar/bin/iptables_update.pl**

7. Type the following commands to check for existing iptables:

   IPv4:

   **iptables -L -n -v**

   IPv6:

   **ip6tables -L -n -v**

RELATED DOCUMENTATION

Data Retention | 169

Custom Offense Close Reasons | 175

# Data Retention

IN THIS SECTION

Retention buckets define how long event and flow data is retained in JSA.

As JSA receives events and flows, each one is compared against the retention bucket filter criteria. When an event or flow matches a retention bucket filter, it is stored in that retention bucket until the deletion policy time period is reached. The default retention period is 30 days; then, the data is immediately deleted.

Retention buckets are sequenced in priority order from the top row to the bottom row. A record is stored in the bucket that matches the filter criteria with highest priority. If the record does not match any of your configured retention buckets, the record is stored in the default retention bucket, which is always located below the list of configurable retention buckets.

## Tenant Data

You can configure up to 10 retention buckets for shared data, and up to 10 retention buckets for each tenant.

When data comes into the system, the data is assessed to determine whether it is shared data or whether the data belongs to a tenant. Tenant-specific data is compared to the retention bucket filters that are defined for that tenant. When the data matches a retention bucket filter, the data is stored in that retention bucket until the retention policy time period is reached.

If you don't configure retention buckets for the tenant, the data is automatically placed in the default retention bucket for the tenant. The default retention period is 30 days, unless you configure a tenant-specific retention bucket.

For more information about tenant data retention, see "Retention Policies for Tenants" on page 384.

## Configuring Retention Buckets

Configure retention policies to define how long JSA is required to keep event and flow data, and what to do when that data reaches a certain age.

Changes to the retention bucket filters are applied immediately to incoming data only. For example, if you configured a retention bucket to retain all data from source IP address 10.0.0.0/8 for 1 day, and you later edit the filter to retain data from source IP 192.168.0.1, the change is not retroactive. Immediately upon changing the filter, the retention bucket has 24 hours of 10.0.0.0/8 data, and all data that is collected after the filter change is 192.168.0.1 data.

The retention policy on the bucket is applied to all data in the bucket, regardless of the filters criteria. Using the previous example, if you changed the retention policy from 1 day to 7 days, both the 10.0.0.0/8 data and the 192.168.0.1 data in the bucket is retained for 7 days.

The **Distribution** of a retention bucket indicates the retention bucket usage as a percentage of total data retention in all your retention buckets. The distribution is calculated on a per-tenant basis.

1. On the navigation menu
   (

☰
), click **Admin**.

2. In the **Data sources** section, click **Event Retention** or **Flow Retention**.

3. If you configured tenants, in the **Tenant** list, select the tenant that you want the retention bucket to apply to.

   **NOTE**: To manage retention policies for shared data in a multi-tenant configuration, choose **N/A** in the **Tenant** list.

4. To configure a new retention bucket, follow these steps:

   a. Double-click the first empty row in the table to open the **Retention Properties** window.

   b. Configure the retention bucket parameters.

   Learn more about retention bucket parameters:

| Properties | Description |
|---|---|
| **Name** | Type a unique name for the retention bucket. |
| **Keep data placed in this bucket for** | The retention period that specifies how long the data is to be kept. When the retention period is reached, data is deleted according to the **Delete data in this bucket** parameter. JSA does not delete data within the retention period. |
| **Delete data in this bucket** | Select **Immediately after the retention period has expired** to delete data immediately on matching the **Keep data placed in this bucket for** parameter. The data is deleted at the next scheduled disk maintenance process, regardless of disk storage requirements.<br>Deletions that are based on storage space begin when the free disk space drops to 15% or less, and the deletions continue until the free disk space is 18% or the policy time frame that is set in the **Keep data placed in this bucket for** field runs out. For example, if the used disk space reaches 85% for records, data is deleted until the used percentage drops to 82%. When storage is required, only data that matches the **Keep data placed in this bucket for** field is deleted.<br>If the bucket is set to **Delete data in this bucket: immediately after the retention period has expired**, no disk space checks are done and the deletion task immediately removes any data that is past the retention. |

*(Continued)*

| Properties | Description |
|---|---|
| **Description** | Type a description for the retention bucket. |
| **Current Filters** | Configure the filter criteria that each piece of data is to be compared against. |

c. Click **Add Filter** after you specify each set of filter criteria.

d. Click **Save**.

5. To edit an existing retention bucket, select the row from the table and click **Edit**.

   Refer to Step "4" on page 171 for information about the retention policy properties.

6. To delete a retention bucket, select the row from the table and click **Delete**.

7. Click **Save**.

   Incoming data that matches the retention policy properties is immediately stored in the retention bucket.

## Managing Retention Bucket Sequence

You can change the order of the retention buckets to ensure that data is being matched against the retention buckets in the order that matches your requirements.

Retention buckets are sequenced in priority order from the top row to the bottom row on the **Event Retention** and **Flow Retention** windows. A record is stored in the first retention bucket that matches the record parameters.

You cannot move the default retention bucket. It always resides at the bottom of the list.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **Data sources** section, click **Event Retention** or **Flow Retention**.

3. If you configured tenants, in the **Tenant** list, select the tenant for the retention buckets that you want to reorder.

> **NOTE**: To manage retention policies for shared data in a multi-tenant configuration, choose **N/A** in the **Tenant** list.

4. Select the row that corresponds to the retention bucket that you want to move, and click **Up** or **Down** to move it to the correct location.

5. Click **Save**.

## Enabling and Disabling a Retention Bucket

When you configure and save a retention bucket, it is enabled by default. You can disable a bucket to tune your event or flow retention.

When you disable a bucket, any new events or flows that match the requirements for the disabled bucket are stored in the next bucket that matches the event or flow properties.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **Data sources** section, click **Event Retention** or **Flow Retention**.

3. If you configured tenants, in the **Tenant** list, select the tenant for the retention bucket that you want to change.

   > **NOTE**: To manage retention policies for shared data in a multi-tenant configuration, choose **N/A** in the **Tenant** list.

4. Select the retention bucket you want to disable, and then click **Enable/Disable**.

## Deleting a Retention Bucket

When you delete a retention bucket, only the criteria that defines the bucket is deleted. The events or flows that were stored in the bucket are collected by the default retention bucket. The default retention period is 30 days; then, the data is immediately deleted.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. In the **Data sources** section, click **Event Retention** or **Flow Retention**.
3. If you configured tenants, in the **Tenant** list, select the tenant for the retention bucket that you want to delete.

   > **NOTE**: To manage retention policies for shared data in a multi-tenant configuration, choose N/A in the **Tenant** list.

4. Select the retention bucket you want to delete, and then click **Delete**.

RELATED DOCUMENTATION

Custom Offense Close Reasons | 175

Configuring a Custom Asset Property | 178

# System Notifications

JSA continuously monitors all appliances and delivers information, warning, and error notifications to the JSA Console, making it easier for you to monitor the status and health of your deployment.

Global System Notifications are host specific and the threshold for each notification is set automatically by JSA.

To show system notifications on your screen, you must configure your browser to allow pop-up windows and ensure that the Enable Popup Notifications check box is selected in your user preferences. If you disable desktop notifications for JSA, you can still view the system notifications under the notifications menu.

During installation, JSA automatically determines and configures the thresholds for all system notifications.

For information about system notifications, see the *Juniper Secure Analytics Troubleshooting Guide*.

**NOTE**: Browser notifications are supported for Mozilla Firefox, Google Chrome, and Microsoft Edge 10. Microsoft Internet Explorer does not support browser-based notifications. Notifications in Internet Explorer appear in a JSA notification box. The way that the notifications appear and how long the messages stay on the screen might vary between browsers.

RELATED DOCUMENTATION

Custom Offense Close Reasons | 175
Configuring a Custom Asset Property | 178

# Custom Offense Close Reasons

IN THIS SECTION

You can manage the options listed in the **Reason for Closing** list box on the **Offenses** tab.

When a user closes an offense on the **Offenses** tab, the Close Offense window is displayed. The user is prompted to select a reason from the **Reason for Closing** list box. Three default options are listed:

- False-positive, tuned
- Non-issue
- Policy violation

Administrators can add, edit, and delete custom offense close reasons from the **Admin** tab.

## Adding a Custom Offense Close Reason

When you add a custom offense close reason, the new reason is listed on the Custom Close Reasons window and in the **Reason for Closing** list box on the Close Offense window of the **Offenses** tab.

The Custom Offense Close Reasons window provides the following parameters.

**Table 24: Custom Close Reasons Window Parameters**

| Parameter | Description |
|---|---|
| Reason | The reason that is displayed in the **Reason for Closing** list box on the Close Offense window of the **Offenses** tab. |
| Created by | The user that created this custom offense close reason. |
| Date Created | The date and time of when the user created this custom offense close reason. |

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Custom Offense Close Reasons**.

3. Click **Add**.

4. Type a unique reason for closing offenses. Reasons must be between 5 and 60 characters in length.

5. Click **OK**.

   Your new custom offense close reason is now listed in the Custom Close Reasons window. The **Reason for Closing** list box on the Close Offense window of the **Offenses** tab also displays the custom reason you added.

## Editing Custom Offense Close Reason

Editing a custom offense close reason updates the reason in the Custom Close Reasons window and the **Reason for Closing** list box on the Close Offense window of the **Offenses** tab.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Custom Offense Close Reasons**.
3. Select the reason you want to edit.
4. Click **Edit**.
5. Type a new unique reason for closing offenses. Reasons must be between 5 and 60 characters in length.
6. Click **OK**.

## Deleting a Custom Offense Close Reason

Deleting a custom offense close reason removes the reason from the **Custom Close Reasons** window and the **Reason for Closing** list box on the **Close Offense** window of the **Offenses** tab.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Custom Offense Close Reasons**.
3. Select the reason you want to delete.
4. Click **Delete**.
5. Click **OK**.

### RELATED DOCUMENTATION

Configuring a Custom Asset Property | **178**

Index Management | **179**

Restrictions to Prevent Resource-intensive Searches | **183**

# Configuring a Custom Asset Property

Define asset properties to facilitate asset queries. Custom properties provide more query options.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Custom Asset Properties**.
3. In the **Name** field, enter a descriptor for the custom asset property.

> **NOTE**: The name must contain only alphanumeric characters, spaces, or underscores. No special characters are allowed.

4. In the **Type** drop-down menu, select **Numeric** or **Text** to define the information type for the custom asset property.
5. Click **OK**.
6. Click the **Assets** tab.
7. Click **Edit Asset >Custom Asset Properties**.
8. Enter the required information in the value field.
9. Click **OK**.

RELATED DOCUMENTATION

Index Management | **179**

Restrictions to Prevent Resource-intensive Searches | **183**

App Hosts | **187**

# Index Management

IN THIS SECTION

Use Index Management to control database indexing on event and flow properties. To improve the speed of searches in JSA, narrow the overall data by adding an indexed field in your search query.

An *index* is a set of items that specify information about data in a file and its location in the file system. Data indexes are built in real-time as data is streamed or are built upon request after data is collected. Searching is more efficient because systems that use indexes don't have to read through every piece of data to locate matches. The index contains references to unique terms in the data and their locations. Because indexes use disk space, storage space might be used to decrease search time.

Use indexing event and flow properties first to optimize your searches. You can enable indexing on any property that is listed in the Index Management window and you can enable indexing on more than one property. When a search starts in JSA, the search engine first filters the data set by indexed properties. The indexed filter eliminates portions of the data set and reduces the overall data volume and number of event or flow logs that must be searched. Without any filters, JSA takes more time to return the results for large data sets.

For example, you might want to find all the logs in the past six months that match the text: `The operation is not allowed`. By default, JSA stores full text indexing for the past 30 days. Therefore, to complete a search from the last 6 months, the system must reread every payload value from every event or flow in that time frame to find matches. Your results display faster when you search with an indexed value filter such as a **Log Source Type**, **Event Name**, or **Source IP**.

The Index Management feature also provides statistics, such as:

- The percentage of saved searches running in your deployment that include the indexed property
- The volume of data that is written to the disk by the index during the selected time frame

To enable payload indexing, you must enable indexing on the Quick Filter property.

## Enabling Indexes

The Index Management window lists all event and flow event properties that can be indexed and provides statistics for the properties. Toolbar options allow you to enable and disable indexing on selected event and flow event properties.

Modifying database indexing might decrease system performance. Ensure that you monitor the statistics after you enable indexing on multiple properties.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Index Management**.

3. Select one or more properties from the Index Management list.

4. Choose one of the following options:

| Situation | Time frame | Action | Reason |
|---|---|---|---|
| The index is disabled and **% of Searches Using Property** is above 30% and **% of Searches Missing Index** is above 30%. | 24 hours, 7 days, or 30 days | Click **Enable Index**. | This search property is used often. Enabling an index can improve performance. |
| The index is enabled and **% of Searches Using Property** is zero. | 30 days | Click **Disable Index**. | The enabled index is not used in the searches. Disable the indexed property to preserve disk space. |

5. Click **Save**.

6. Click **OK**.

In lists that include event and flow event properties, indexed property names are appended with the following text: `[Indexed]`. Examples of such lists include the search parameters on the **Log Activity** and **Network Activity Log Activity** tab search criteria pages and the **Add Filter** window.

## Enabling Payload Indexing to Optimize Search Times

To optimize event and flow search times, enable payload indexing on the **Quick Filter** property.

> **NOTE**: Use the **Quick Filter** feature in the **Log Activity** and **Network Activity** tab to search event and flow payloads by using a text string.
>
> Payload indexing increases disk storage requirements and might affect system performance. Enable payload indexing if your deployment meets the following conditions:
>
> - The event and flow processors are at less than 70% disk usage.
> - The event and flow processors are less than 70% of the maximum events per second (EPS) or flows per interface (FPI) rating.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Index Management**.
3. In the **Quick Search** field, type **Quick Filter**.

   The **Quick Filter** property is displayed.
4. Select the **Quick Filter** property that you want to index.

   In the results table, use the value in the **Database** column to identify the flows or events **Quick Filter** property.
5. On the toolbar, click **Enable Index**.

   A green dot indicates that the payload index is enabled.

   If a list includes event or flow properties that are indexed, the property names are appended with the following text: `[Indexed]`.
6. Click **Save**.

To manage payload indexes, see "Configuring the Retention Period for Payload Indexes" on page 182.

## Configuring the Retention Period for Payload Indexes

By default, JSA sets 30 days for the data retention period of the payload index. You can search for specific values in quick filter indexes beyond 30 days by changing the default retention in JSA.

Your virtual and physical appliances require a minimum of 24 GB of RAM to enable full payload indexing. However, 48 GB of RAM is suggested.

The minimum and suggested RAM values applies to all JSA systems that are processing events or flows.

The retention values reflect the time spans that you are typically searching. The minimum retention period is 1 day and the maximum is 2 years.

> **NOTE**: Quick Filter searches that use a time frame outside of the Payload Index Retention setting can trigger slow and resource-intensive system responses. For example, if the payload index retention is set for 1 day, and you use a time frame for the last 30 hours in the search.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **System Settings**.
3. In the **Database Settings** section, select a retention time period from the **Payload Index Retention** list.
4. Click **Save**.
5. Close the **System Settings** window.
6. On the **Admin** tab menu, click **Deploy Changes**.

If you retain payload indexes longer than the default value, extra disk space is used. After you select a greater value in the **Payload Index Retention** field, monitor system notifications to ensure that you do not fill disk space.

### RELATED DOCUMENTATION

- Restrictions to Prevent Resource-intensive Searches | **183**
- App Hosts | **187**
- Checking the Integrity Of Event and Flow Logs | **193**

# Restrictions to Prevent Resource-intensive Searches

IN THIS SECTION

You can balance the usage of your JSA infrastructure by setting resource restrictions on JSA event and flow searches.

Before you set resource restrictions, carefully consider the normal operational procedures in your environment. Try to set restrictions that ensure that all users have access to the data that they require, yet prevent them from inadvertently running large queries that negatively impact the system availability and performance for other users.

## Types Of Resource Restrictions

You can set limitations on searches by configuring either time or data set restrictions based on user, role, or tenant.

Resource restrictions are applied in the following order: user, user role, and tenant. For example, restrictions that are set for a user take precedence over restrictions that are set for the user role or tenant that the user is assigned to.

You can set the following types of restrictions on event and flow searches:

- The length of time that a search runs before data is returned
- The time span of the data to be searched
- The number of records that are processed by the Ariel query server

> **NOTE**: Ariel search stops when the record limit is reached, but all in-progress search results are returned to the search manager and are not truncated. Therefore, the search result set is often larger than the specified record limit.

### User-based Restrictions

User-based restrictions define limits for an individual user, and they take precedence over role and tenant restrictions.

For example, your organization hires university students to work with the junior analysts in your SOC. The students have the same user role as the other junior analysts, but you apply more restrictive user-based restrictions until the students are properly trained in building JSA queries.

### Role-based Restrictions

Role-based restrictions allow you to define groups of users who require different levels of access to your JSA deployment. By setting role-based restrictions, you can balance the needs of different types of users.

For example, a junior security analyst might focus on security incidents that happened recently, while a senior security analyst might be more involved in forensic investigations that review data over a longer period of time. By setting role-based restrictions, you can limit a junior analyst to accessing only the last 7 days of data, while a senior analyst has access to a much larger time span of data.

### Tenant-based Restrictions

In a Managed Security Service Provider (MSSP) or a multi-divisional organization, tenant-based restrictions can help you ensure quality of service by preventing resource contention and degradation of services. You can prevent a tenant from querying terabytes of data that can negatively impact the system performance for all other tenants.

As an MSSP, you might define standard resource restrictions based on a set of criteria that each tenant is compared to. For example, the standard configuration for a medium-sized tenant might include resource restrictions that limit searches to accessing only the last 14 days of data and a maximum of 10,000 records returned.

## Resource Restrictions in Distributed Environments

In a distributed environment, the timing of the data transfer between the JSA console and managed hosts can impact the search results.

When you run a search in JSA, the search runs on all nodes at the same time. Each managed host runs the search, and sends the aggregated results to the JSA console when the search is complete or when it reaches the predefined number of rows.

It is important to understand how the resource restrictions that you set might impact the search results that are returned to a user:

- **Canceled searches --** Each managed host periodically checks the state of the resource restriction limit. If a limit is reached, the search is automatically canceled to prevent the incomplete results from being cached and reused.

  Results that were collected before the search was canceled by the system can be viewed by clicking **Search >Manage Search Results** on the **Log Activity** or **Network Activity** tab.

- **Empty search results --** When you set time-limit or record-limit restrictions, the remote aggregation might cause the JSA console to reach the resource restriction limit before the managed host sends the partial aggregate to the console. In this situation, the search results might appear to be empty even though some data was collected.

- **Inconsistent search results --** JSA monitors the load on each managed host, and manages the search to ensure optimized performance throughout the entire deployment. Depending on the system load, searches that are run repeatedly might show slightly different results due to the managed hosts that return the data in a different order.

  For example, in a deployment that has six event processors, EP1, EP3, and EP5 might be the first processors to return data on the initial run. In subsequent runs, EP2, EP4, and EP6 might return data first, which accounts for the inconsistent search results.

- **Limited search results** — You can set a limit restriction on search results for JSA that limits the number of records that are read from the disk in a search query. A limit ensures that the query stops after any managed host that is participating in the search reads the restricted number of entries from the disk. The query does not retrieve all the data and gives you only the restricted number of rows. Setting this restriction can prevent a system crash in the instance of a large amount of data.

  For example, if you set the restriction at 10,000 rows, the query stops running after the managed host processes 10,000 records.

Depending on the frequency that users reach the resource restrictions, you can tune the limits to avoid restricting users from running reasonable searches to meet their job requirements. Users who consistently run searches that strain the system might benefit from more training in building JSA queries. For more information, see the *Juniper Secure Analytics Ariel Query Language Guide*.

## Configuring Resource Restrictions

Set resource restrictions to apply time or data limitations to event and flow searches.

You can set the following types of resource restrictions:

- **Execution time** restrictions specify the maximum amount of time that a query runs before data is returned.
- Time span restrictions specify the time span of the data to be searched.
- **Record limit** restrictions specify the number of data records that are returned by a search query.

Users who run searches that are limited by resource restrictions see the resource restriction icon
(

) next to the search criteria.

> **NOTE**: The search result set is often larger than the specified record limit. When the record limit is reached, the search manager signals all search participants to stop (there are multiple search participants even on a single system), but results continue to accumulate until the search fully stops on all participants. All search results are added to the result set.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Resource Restrictions**.
3. If your deployment has tenants that are configured, click **Role** or **Tenant** to specify the type of restrictions to set.
4. Double-click the role or tenant that you want to set restrictions for.
5. To set restrictions for all users who are assigned to the user role or tenant, follow these steps:
   a. Click the summary row at the top to open the **Edit Restriction** dialog box.
   b. Click **Enabled** for the type of restriction that you want to set, and specify the restriction values.
   c. Click **Save**.
6. To set restrictions for a specific user, follow these steps:
   a. Double-click the user that you want to set the restrictions for.

      To search for a user, type the user name in the filter field.

   b. Click **Enabled** for the type of restriction that you want to set, and specify the restriction values.

c. Click **Save**.

RELATED DOCUMENTATION

App Hosts | 187

Checking the Integrity Of Event and Flow Logs | 193

Adding Custom Actions | 197

# App Hosts

IN THIS SECTION

An App Host is a managed host that is dedicated to running apps. App Hosts provide extra storage, memory, and CPU resources for your apps without impacting the processing capacity of your JSA Console. Apps such as User Behavior Analytics with Machine Learning Analytics require more resources than are currently available on the Console.

The App Host replaces the App Node. Unlike the App Node, JSA manages all updates to the App Host. App Host supports high availability and you can include the App Host in your high availability deployments.

**NOTE:**

- The appliance ID for the App Host is 4000.
- You can have only one App Host in your deployment.
- Port 5000 must be open on your console.

- Port 443 must be open on your console.
- If your App Host is not encrypted, open port 9000 and 14433 for unidirectional communication from your console to the App Host.
- If your App Host is encrypted, open port 26000 and 26001 for unidirectional communication from your console to the App Host.

### App Host specifications

The following table shows the minimum requirements and suggested specifications for an App Host.

**NOTE**: *The suggested specifications for medium and large sized deployments haven't been tested. If you are using some of the larger apps, such as the Pulse Dashboard and User Behavior Analytics with Machine Learning, the minimum requirements are probably insufficient. Consider upgrading your deployment environment.

**Table 25: App Host Secifications**

| | CPU cores | RAM | Disk Space | Description |
|---|---|---|---|---|
| Small | 4 | 12 GB | 256 GB | Minimum requirements for an App Host. You can run most apps with the minimum requirements, but not larger apps such as JSA DNS Analyzer and User Behavior Analytics with Machine Learning. |
| Medium | 12 or more | 64 GB or more | 500 GB or more | *You can run all apps that exist today, but this specification does not give you room for future apps. |
| Large | 24 or more | 128 GB or more | 1 TB or more | *You can run all apps that exist today and you would have room for future apps. |

### Installation Scenarios

If you are installing an App Host and you do not have an App Node in your deployment, see "Installing an App Host" .

## Installing an App Host

You can run apps on an App Host instead of your Console to lessen the processing load on the Console. Install an App Host the same way you would any other managed host for JSA. You can install an App Host on hardware or in a VM, and as either an appliance install or a software install.

- This procedure assumes that you are doing an appliance install. For more information about appliance and software installations, see *Juniper Secure Analytics Administration Guide*.
- Ensure that all apps on your system are updated.
- Resolve any issues with applications in an error state or not displaying properly
- Schedule a maintenance window for this task and ensure that users do not do any of the following during the migration.
  - Do not install or uninstall apps.
  - Do not do a full deploy.
  - Do not do a restore.
  - Do not delete the App Host.
  - Do not re-IP the Console.

1. Type `root` at the login prompt to start the installation wizard. Type password if you are prompted for a password.
2. Accept the **End User License Agreement**.
3. Select **App Host Appliance** for the appliance type.
4. For the type of setup, select **Normal Setup (default)**, and set up the time.
5. Select the Internet Protocol version:
   - Select **ipv4** or **ipv6**.
6. If you selected **ipv6**, select **manual** or **auto** for the **Configuration type**.
7. Select the bonded interface setup, if required.
8. Select the management interface.

9. In the wizard, enter a fully qualified domain name in the **Hostname** field.

10. In the **IP address** field, enter a static IP address, or use the assigned IP address.

11. If you do not have an email server, enter `localhost` in the **Email server name** field.

12. Enter a `root` password that meets the following criteria:

    - Contains at least 5 characters

    - Contains no spaces

    - Can include the following special characters: @, #, ^, and *.

13. Click **Finish**.

14. Follow the instructions in the installation wizard to complete the installation.

    The installation process might take several minutes.

15. Add this managed host to your deployment and deploy changes.

## Changing Where Apps are Run

Change where apps are run when you add an App Host to your deployment.

Apps are disabled during the transfer process. Apps are enabled when the transfer is complete.

If you don't have enough disk space or memory available on the Console, moving your apps from the App Host to the Console moves only the apps themselves, but not app data. Any app data remains on your App Host, and apps do not start on the Console when the transfer is complete. Apps start again when they are transferred back to the App Host.

1. Sign in to the JSA user interface.

2. Click **Admin**.

3. On the **System and License Management** screen, click the **Click to change where apps are run** link.

4. 4. Choose where apps are run.

   - Click **App Host** to transfer apps to the App Host.

   - Click **Console** to transfer apps to the Console.

**NOTE**: The more apps and app data you have, the longer the transfer takes.

## Troubleshooting an App Node to App Host Migration

You can troubleshoot the following issues if you come across them with your App Node to App Host migration.

**App Node backup file MD5 checksum is d41d8cd98f00b204e9800998ecf8427e**

An MD5 checksum of d41d8cd98f00b204e9800998ecf8427e indicates that the App Node backup file is a zero-byte file. Insufficient disk space available for the backup file might be the cause. If you receive this value for the checksum:

1. Use `ssh` to log in to your App Node as the root user.
2. Type the following command and note the space available in **/store**.

   **df -h /store**

3. Type the following command and note the used space in **/store/backup/marathon**.

   **du -hs /store/backup/marathon/**

4. Type the following command and note the used space in **/store/docker/volumes**.

   **du -hs /store/docker/volumes/**

5. Compare the total space that is used by **/store/docker/volumes** to the total space available in **/ store**. You need at least 1 to 1.5 times as much free space in **/ store** as space used by /store/ docker/ volumes. If you don't have enough free space in **/ store**, check to see whether you have enough used space in **/store/backup/marathon** to make up the difference.

   For example, if the space used by **/store/docker/volumes** is 100 GB, and the space available in **/ store** is 90 GB, you don't have enough free space for the backup file. If **/store/backup/marathon** is using 10 GB or more of space, you can free space in **/store/backup/marathon**.

6. Back up and remove files from **/store/backup/marathon** to free up space by following these steps:

   a. On your Console, create an App Node back up directory under /store by typing the following command.

      **mkdir /store/app_node_backup**

   b. Copy marathon backup files from the App Node to the Console by typing the following command.

scp root@<appnode_IP_address>:/store/backup/marathon/backup.marathon-volumes.qapp.*.tgz / store/app_node_backup/

c. Check the MD5 checksum of the marathon backup files on the Console by typing the following command.

**ls /store/backup/marathon/backup.marathon-volumes.qapp.*.tgz | xargs md5sum**

d. On the App Node, check the MD5 checksum of the marathon backup files by typing the following command.

**ls /store/backup/marathon/backup.marathon-volumes.qapp.*.tgz | xargs md5sum**

e. Verify that both checksum values are the same. If they are, remove the marathon backup files from the App Node by typing the following command on the App Node.

> **NOTE**: The `rm -rf` command removes a directory and all files in it. Ensure that you enter the command exactly as shown here.

**rm -rf /store/backup/marathon/***

## Removing an App Host

You can't remove an App Host if you have any apps that are running on it.

1. Move your apps back to the Console. See "Changing Where Apps are Run".

> **NOTE**: If you don't have enough disk space or memory available on the Console, moving your apps to the Console moves only the apps themselves, but not app data. Any app data remains on your App Host.

2. Click **Admin**.
3. On the **System Configuration** screen, click **System and License Management**..
4. In the **Display** list, select **Systems**.
5. Select your App Host.
6. On the **Deployment Actions** menu, click **Remove Host**.

RELATED DOCUMENTATION

Checking the Integrity Of Event and Flow Logs | 193

Adding Custom Actions | 197

Managing Aggregated Data Views | 202

# Checking the Integrity Of Event and Flow Logs

**IN THIS SECTION**

When log hashing is enabled, any system that writes event and flow data creates hash files. Use these hash files to verify that the event and flow logs were not modified since they were originally written to disk.

The hash files are generated in memory before the files are written to disk, so the event and flow logs cannot be tampered with before the hash files are generated.

Ensure that log hashing is enabled for your JSA system. For more information about enabling log hashing, see "Enabling Log Hashing" on page 195.

You must log in to the system that has the data storage for events and flows, and run a utility to check the logs. You cannot check the log integrity in the event and flow viewer interface.

1. Use SSH to log in to JSA as the root user.
2. To run the utility, type the following command:

```
/opt/qradar/bin/check_ariel_integrity.sh -d <duration> -n <database name>
[-t <endtime>] [-a <hash algorithm>] [-r <hash root directory>] [-k <hmac key>]
```

This table describes the parameters that are used with the **check_ariel_integrity.sh** utility.

Table 26: Parameters for the check_ariel_integrity.sh Utility

| Parameter | |
|---|---|
| -d | Duration of time, in minutes, of the log file data to scan. The time period immediately precedes the end time that is specified using the **-t** parameter. For example, if **-d 5** is entered, all log data that was collected five minutes before the **-t** end time is scanned. |
| -n | The JSA database to scan. Valid options are events and flows. |
| -t | The end time for the scan. The format for the end time is "yyyy/mm/dd hh:mm" where hh is specified in 24-hour format. If no end time is entered, the current time is used. |
| -a | Hashing algorithm to use. This algorithm must be the same one that was used to create the hash keys. If no algorithm is entered, SHA-1 is used. |
| -r | The location of the log hashing. This argument is required only when the log hashing is not in the location that is specified in the configuration file, **/opt/qradar/conf/arielConfig.xml**. |
| -k | The key that is used for Hash-based Message Authentication Code (HMAC) encryption. If you do not specify an HMAC key and your system is enabled for HMAC encryption, the **check_ariel_integrity.sh** script defaults to the key specified in the system settings. |
| -h | Shows the help message for the **check_ariel_integrity.sh** utility. |

For example, to validate the last 10 minutes of event data, type the following command:

```
/opt/qradar/bin/check_ariel_integrity.sh -n events -d 10
```

If an `ERROR` or `FAILED` message is returned, the hash key that is generated from the current data on the disk does not match the hash key that was created when the data was written to the disk. Either the key or the data was modified.

## Enabling Log Hashing

Enable log hashing to have any system that writes event and flow data creates hash files. Use these hash files to verify that the event and flow logs were not modified since they were originally written to disk. The hash files are generated in memory before the files are written to disk, so the event and flow logs cannot be tampered with before the hash files are generated.

The system uses the following hashing algorithm types:

**Message-Digest Hash Algorithm**

Transforms digital signatures into shorter values called Message-Digests (MD).

**Secure Hash Algorithm (SHA) Hash Algorithm**

Standard algorithm that creates a larger (60 bit) MD.

1. On the **Admin** tab, click **System Settings**.
2. In the Ariel Database Settings section, select **Yes** in the **Flow Log Hashing** field and the **Event Log Hashing** field.
3. Select a hashing algorithm for database integrity.
   - If the **HMAC Encryption** parameter is disabled, the following hashing algorithm options are available:

     **MD2**

     Algorithm that is defined by RFC 1319.

     **MD5**

     Algorithm that is defined by RFC 1321.

     **SHA-1**

     Algorithm that is defined by Secure Hash Standard (SHS), NIST FIPS 180-1. This setting is the default.

     **SHA-256**

Algorithm that is defined by the draft Federal Information Processing Standard 180-2, SHS. SHA-256 is a 255-bit hash algorithm that is intended for 128 bits of security against security attacks.

SHA-384

Algorithm that is defined by the draft Federal Information Processing Standard 180-2, SHS. SHA-384 is a bit hash algorithm, which is created by truncating the SHA-512 output.

SHA-512

Algorithm that is defined by the draft Federal Information Processing Standard 180-2, SHS. SHA-512 is a bit hash algorithm that is intended to provide 256 bits of security.

- If the **HMAC Encryption** parameter is enabled, the following hashing algorithm options are available:

  HMAC-MD5

  An encryption method that is based on the MD5 hashing algorithm.

  HMAC-SHA-1

  An encryption method that is based on the SHA-1 hashing algorithm.

  HMAC-SHA-256

  An encryption method that is based on the SHA-256 hashing algorithm.

  HMAC-SHA-384

  An encryption method that is based on the SHA-384 hashing algorithm.

  HMAC-SHA-512

  An encryption method that is based on the SHA-512 hashing algorithm.

  If the **HMAC Encryption** parameter is enabled, you must specify an HMAC key in the **HMAC Key** and **Verify HMAC Key** fields.

4. Click **Save**.

## RELATED DOCUMENTATION

Adding Custom Actions | 197

Managing Aggregated Data Views | 202

Accessing a GLOBALVIEW Database | 204

# Adding Custom Actions

IN THIS SECTION

Attach scripts to custom rules to do specific actions in response to network events. Use the **Custom Action** window to manage custom action scripts.

Use custom actions to select or define the value that is passed to the script and the resulting action.

The following examples are custom actions that are the outcomes of passing values to a script:

- Block users and domains.
- Initiate work flows and updates in external systems.
- Update TAXI servers with a STIX representation of a threat.

Custom actions work best with low volume custom rule events and with custom rules that have a low response limiter value.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **Custom Action** section,, click **Define Custom Action**.

3. To upload scripts, click **Add**. Programming language versions that the product supports are listed in the **Interpreter** list.

   For the security of your deployment, JSA does not support the full range of scripting functionality that is provided by the Python, Perl, or Bash languages.

4. Specify the parameters to pass to the script that you uploaded.

Table 27: Custom Action Parameters

| Parameter | Description |
|---|---|
| **Fixed property** | Values that are passed to the custom action script.<br>These properties are not based on the events or flow themselves, but cover other defined values that you can use the script to act on.<br>For example, pass the fixed properties **username** and **password** for a third-party system to a script to send an SMS alert.<br>Encrypt fixed properties by selecting the **Encrypt value** check box. |
| **Network event property** | Dynamic Ariel properties that are generated by events. Select from the **Property** list.<br>For example, the network event property **sourceip** provides a parameter that matches the source IP address of the triggered event.<br>For more information about Ariel properties, see the *Juniper Secure Analytics Ariel Query Language Guide*. |

Parameters are passed into your script in the order in which you added them in the **Define Custom Action** dialog box.

When custom action scripts are run, a **chroot jail** is set up in the **/opt/qradar/bin/ca_jail/** directory. Any content in the **/opt/qradar/bin/ca_jail/** directory can be modified and written to by scripts. The custom action user's home directory (**/home/customactionuser**) can be modified.

A script can run only from inside the jail environment so that it does not interfere with the JSA run environment.

All file access during custom action execution is relative to the **/opt/qradar/bin/ca_jail/** directory.

The custom action user account might not have permission to run follow-up commands, such as logging into a firewall and blocking an IP address. Test whether your script runs successfully before you associate it with a rule.

**NOTE**: The type of custom action that you implement depends on your network infrastructure and its components. For example, you can configure REST APIs on Cisco devices to block suspect IP addresses. Other third-party vendors might not provide a REST interface, so you might need to develop your own web services solution to run custom actions.

You must run the dos2unix utility on scripts that originate from a Windows or DOS system. Windows or DOS systems typically add control characters. To successfully test custom action scripts by using the script Test Execution function in JSA, you must remove the control characters.

## Testing Your Custom Action

Test whether your script runs successfully and has the intended result before you associate it with a rule.

Custom action scripts run inside a testing environment on your JSA managed hosts that is isolated from your production environment. Custom action scripts typically run on the managed host that runs the event processor. However, if you have an All-In-One appliance, custom actions run on the JSA console.

**Test Execution** is supported only on the JSA console and is not supported on managed hosts.

If you must write to disk from a custom action script, you must use the following directory: **/home/customactionuser**.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **Custom actions** section, click **Define Actions**.

3. Select a custom action from the list and click **Test Execution >Execute** to test your script. The result of the test and any output that is produced by the script is returned.

4. After you configure and test your custom action, use the **Rule Wizard** to create a new event rule and associate the custom action with it.

   For more information about event rules, see the *Juniper Secure Analytics Users Guide*.

## Passing Parameters to a Custom Action Script

Sample scripts in Bash, Python, and Perl show how to pass parameters to custom action scripts.

The following simple sample scripts show how to query the asset model API for an asset with the supplied offense source IP address. For the sake of this example, the scripts output the JSON that is returned by the endpoint.

The scripts require three parameters:

- Console IP address
- API token
- Offense source IP address

These parameters are configured in the Define Custom Action window **Script Parameters** area:

**Figure 6: Custom Action Script Parameters**

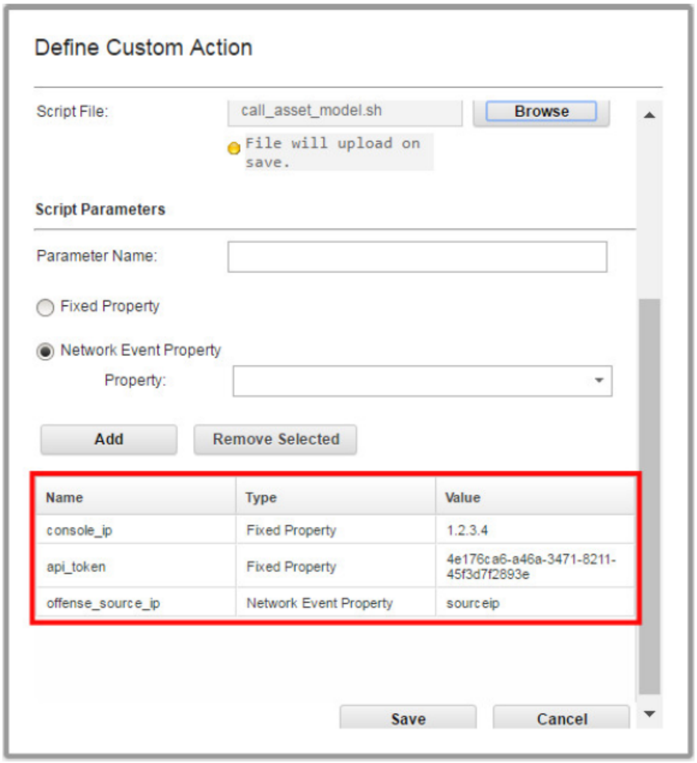

Define Custom Action

Script File: call_asset_model.sh

File will upload on save.

Script Parameters

Parameter Name:

○ Fixed Property

● Network Event Property

Property:

| Name | Type | Value |
|---|---|---|
| console_ip | Fixed Property | 1.2.3.4 |
| api_token | Fixed Property | 4e176ca6-a46a-3471-8211-45f3d7f2893e |
| offense_source_ip | Network Event Property | sourceip |

Each parameter is passed to the script in the order in which it was added in the **Define Custom Action** window. In this case:

1. console_ip
2. api_token
3. offense_source_ip

The variables that are defined at the beginning of each of the sample scripts use the sample parameter names that were added in the **Define Custom Action** window.

**Figure 7: Call_asset_model.sh**

```
#!/bin/bash
console_ip=$1
api_token=$2
offense_source_ip=$3
auth_header="SEC:$api_token"
output=$(curl -k -H $auth_header https://$console_ip/console/restapi/api/
asset_model/assets?filter=interfaces%20contains%20%20ip_addresses
%20contains%20%20value%20%3D%20%22$offense_source_ip%22)
# Basic print out of the output of the command
echo $output
```

**Figure 8: Call_asset_model.py**

```
#!/usr/bin/python
import sys
import requests
console_ip = sys.argv[1]
api_token = sys.argv[2]
offense_source_ip = sys.argv[3]
auth_header = {’SEC’ : api_token }
endpoint = "https://{0}/console/restapi/api/asset_model/
assets?filter=interfaces%20contains%20%20ip_addresses
%20contains%20%20value%20%3D%20%22{1}%22"
.format(console_ip, offense_source_ip)
response = requests.get(endpoint, headers=auth_header, verify=False)
# Basic print out of the output of the command
print(response.json())
```

**Figure 9: Call_asset_model.pl**

```
#!/usr/bin/perl
use strict;
use warnings;
use LWP::UserAgent;
my $console_ip = $ARGV[0];
my $api_token = $ARGV[1];
my $offense_source_ip = $ARGV[2];
my $endpoint = "https://$console_ip/console/restapi/api/asset_model/
assets?filter=interfaces%20contains%20%20ip_addresses
%20contains%20%20value%20%3D%20%22$offense_source_ip%22";
my $client = LWP::UserAgent -> new(ssl_opts => { verify_hostname => 0 });
my $response = $client -> get($endpoint, "SEC" => $api_token);
# Basic print out of the output of the command
print $response -> decoded_content;
```

### RELATED DOCUMENTATION

Managing Aggregated Data Views | **202**

Accessing a GLOBALVIEW Database | **204**

Checking the Integrity Of Event and Flow Logs | **193**

# Managing Aggregated Data Views

A large volume of data aggregation can decrease your system performance. The Ariel function uses a separate database for aggregated data in order to improve system performance and to make the data more readily available. You can disable, enable, or delete aggregated data views.

The items that appear in the **Display** list sort the data.

The Aggregated Data View is required to generate data for ADE rules, time series graphs, and reports.

Disable or delete views if the maximum number of views is reached.

Duplicate views can appear in the **Aggregated Data ID** column because an aggregated data view can include multiple searches.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Aggregated Data Management**.
3. To filter the list of aggregated data views, perform one the following options:
    - Select an option from the **View**, **Database**, **Show**, or **Display** list.

- Type an aggregated data ID, report name, chart name, or saved search name in the search field.

4. To manage an aggregated data view, select the view, and then click the appropriate action on the toolbar:

   - If you select **Disable View** or **Delete View**, content dependencies are displayed for the aggregated data view. After you disable or delete the view, the dependent components no longer use aggregated data.

   - Enable a previously disabled aggregated data view to restore the view.

**Table 28: Aggregated Data Management View Column Descriptions**

| Column | Description |
|---|---|
| Aggregated Data ID | Identifier for the aggregated data |
| Saved Search Name | Defined name for the saved search |
| Column Name | Column identifier |
| Times Searches | Search count |
| Data Written | The size of the written data |
| Database Name | Database where the file was written |
| Last Modified Time | Timestamp of the last data modification |
| Unique Count Enabled | True or False: Search the results to display unique event and flowcounts instead of average counts over time. |

## RELATED DOCUMENTATION

Accessing a GLOBALVIEW Database | 204

Checking the Integrity Of Event and Flow Logs | 193

Adding Custom Actions | 197

# Accessing a GLOBALVIEW Database

Use the JSA REST API documentation interface to get the GLOBALVIEW database results for a given saved search name and time range. The type of data contained in the database results corresponds to the type of saved search queried.

1. Find a saved search.

   a. On the navigation menu
   (
   ☰
   ), click **Admin**.

   b. In the **System Configuration** section, click **Aggregated Data Management**.

   c. Under the **Saved Search Name** column, record a saved search name from the list.

2. Query the JSA REST API to find a search ID.

   a. Log in to the JSA API, **https://<Console IP>/api_doc**, as an administrator.

   b. Click the most recent version of the JSA API.

   c. Click the **/ariel/searches** endpoint.

   d. Click **POST**.

   e. In the **query_expression** parameter field, type the following command: `select * from GLOBALVIEW('savedsearch','timerange')`

      Use one of the following values for the *timerange* variable:

      *NORMAL HOURLY DAILY*

      The following example shows query for Top Log Sources with a time range of the last two days:

      ```
      select * from GLOBALVIEW('Top Log Sources','DAILY') last 2 days
      ```

   f. Click **Try It Out!**

   g. Copy the search ID from the response body.

3. Get the search results.

   a. From the **/ariel/searches/search{id}/results** endpoint, click **GET**.

   b. In the **search_id** parameter field, enter the search ID.

   c. Click **Try It Out!**.

d. Ensure that the search successfully completes.

e. Get the database results from the response body.

## RELATED DOCUMENTATION

Checking the Integrity Of Event and Flow Logs | 193

Adding Custom Actions | 197

Managing Aggregated Data Views | 202

7

CHAPTER

# Event Data Processing in JSA

# Event Data Processing in JSA

In JSA, use the DSM Editor to solve parsing problems and to add custom parsing.

The DSM Editor provides real-time feedback so that you know whether your customization works the way that you expect it to.

RELATED DOCUMENTATION

# DSM Editor Overview

IN THIS SECTION

Instead of manually creating a log source extension to fix parsing issues or extend support for new log source types, use the DSM Editor. The DSM Editor provides different views of your data. You use the DSM Editor to extract fields, define custom properties, categorize events, and define new QID definition.

The DSM Editor provides the following views:

## Workspace

The **Workspace** shows you raw event data. Use sample event payloads to test the behavior of the log source type, and then the **Workspace** area shows you the data that you capture in real time.

All sample events are sent from the workspace to the DSM simulator, where properties are parsed and QID maps are looked up. The results are displayed in the **Log Activity Preview** section. Click the **pencil** icon to open in edit mode.

In the edit mode, you paste up to 100,000 characters of event data into the workspace or edit data directly. When you edit properties on the **Properties** tab, matches in the payload are highlighted in the workspace. Custom properties and overridden system properties are also highlighted in the **Workspace**.

New in JSA 7.4.1, you can specify a custom delimiter that makes it easier for JSA to ingest multiline events. To ensure that your event is kept intact as a single multiline event, select the **Override event delimiter** checkbox to separate the individual events based on another character or sequence of characters. For example, if your configuration is ingesting multiline events, you can add a special character to the end of each distinct event in the **Workspace**, and then identify this special character as the event delimiter.

New in JSA 7.4.2, JSA can suggest regular expressions (regex) when you enter event data in the **Workspace**. If you are not familiar with creating regex expressions, use this feature to generate your regex. Highlight the payload text that you want to capture and in the **Properties** tab, click **Suggest Regex**. The suggested expression appears in the **Expression** field. Alternatively, you can click the **Regex** button in the **Workspace** and select the property that you want to write an expression for. If JSA cannot generate a suitable regex for your data sample, a system message appears.

**TIP**: The regex generator works best for fields in well-structured event payloads. If your payload consists of complex data from natural language or unstructured events, the regex generator might not be able to parse it and does not return a result.

## Log Activity Preview

New in JSA 7.4.1 The **Parsing Status** column was added to the Log Activity Preview.

The **Log Activity Preview** simulates how the payloads in the workspace appear in the **Log Activity** viewer. The **Parsing Status** column indicates whether your event properties are successfully parsing and mapping to a QID record. Every standard property that is supported is displayed. The fields that are marked with an asterisk (*), for example, **Event name**, **Severity**, **Low-level category**, and **QID,** are populated from the QID map. Fields that are populated from the QID map cannot be parsed verbatim from the raw events data in the workspace, so they cannot be defined or edited. However, you can adjust their values by selecting the corresponding event ID and category combination from the **Event Mappings** tab.

**NOTE**: You must set an Event ID for any system properties to be parsed correctly.

Click the **wrench** icon to select which columns to show or to hide in the **Log Activity Preview** window, and to reorder the columns.

## Properties

The **properties** tab contains the combined set of system and custom properties that constitute a DSM configuration. Configuring a system property differs from configuring a custom property. You can override a property, by selecting the **Override system behaviour** check box and entering the regex or JSON expression.

**NOTE**: If you override the **Event Category** property, you must also override the **Event ID** property.

Matches in the payload are highlighted in the event data in the workspace. The highlighting color is two-toned, depending on what you capture. For example, the orange highlighting represents the capture group value while the bright yellow highlighting represents the rest of the regex that you specified. The feedback in the workspace shows whether you have the correct regex. If an expression is in focus, the highlighting in the workspace reflects only what that expression can match. If the overall property is in focus, then the highlighting turns green and shows what the aggregate set of expressions can match, taking into account the order of precedence.

In the format strings field, capture groups are represented by using *$<number>* notation. For example, $1 represents the first capture group from the regex, $2 is the second capture group, and so on.

You can add multiple expressions to the same property, and you can assign precedence by dragging and dropping the expressions to the top of the list.

A warning tool tip beside any of the properties indicates that no expression was added.

## Event Mappings Tab

New in JSA 7.4.1 Support for copying Event ID and Event Category fields was added to the **Event Mapping** tab.

The **Event Mappings** tab displays all the event ID and category combinations that exist in the system. If a new event mapping is created, it is added to the list of event ID and category combination that is displayed in the **Event Mappings** tab. In general, the **Event Mappings** tab displays all event ID and category combinations and the QID records that they are mapped to.

## Configuration Tab

You can configure Auto Property Discovery for structured data that are in JSON format. By default, log source types have Auto Property Discovery turned off.You can configure Auto Property Discovery for structured data that are in JSON format. By default, log source types have Auto Property Discovery turned off.

When you enable **Auto Property Discovery** on the **Configuration** tab, the property discovery engine automatically generates new properties to capture all fields that are present in the events that are received by a log source type. You can configure the number of consecutive events to be inspected for new properties in the **Discovery Completion Threshold** field. Newly discovered properties appear in the **Properties** tab, and are made available for use in the rules and search indexes. However, if no new properties are discovered before the threshold, the discovery process is considered complete and Auto Property Discovery for that log source type is disabled. You can manually enable the **Auto Property Discovery** on the Configuration tab at any time.

> **NOTE**: To continuously inspect events for a log source type, you must make sure that you set the **Discovery Completion Threshold** value to 0.

RELATED DOCUMENTATION

Properties in the DSM Editor | 211

Property Configuration in the DSM Editor | 212

Opening the DSM Editor | 225

# Properties in the DSM Editor

IN THIS SECTION

In the DSM Editor, normalized system properties are combined with custom properties and are sorted alphabetically.

A DSM cannot have multiple properties with the same name.

The configuration of a system property differs from a custom property.

## System Properties

System properties cannot be deleted but you can override the default behavior. There are two types of system properties:

- **Predefined system property** -- Displays the default JSA behavior that is used for the DSM.
- **Override system property** -- System properties with override configured (log source extension) show **Override** in the status line. When a system property has an override, a log source extension for that DSM uses the regular expressions that you entered for the configuration.

## Custom Properties

Custom properties show **Custom** in the status line.

Custom properties differ from system properties in these ways:

- Custom properties display **Custom** below their name.
- Custom properties have no **Override system behavior** check box.

- To make a custom property available for rules and search indexing, select the **Enable this Property for use in Rules and Search Indexing** check box when you create a custom property.

> **NOTE**: When you select this option, JSA attempts to extract the property from events as soon as they enter the pipeline. Extracted property information and the remainder of the event record are persisted. The property does not need to be extracted again when it is used in a search, or report. The process enhances performance when the property is retrieved, but the process can have a negative impact on performance during event collection and storage.

- Custom properties must have one or more expressions to be valid.

RELATED DOCUMENTATION

Property Configuration in the DSM Editor | **212**

Opening the DSM Editor | **225**

Configuring a Log Source Type | **226**

# Property Configuration in the DSM Editor

IN THIS SECTION

Configure properties in the DSM Editor to change the behavior of an overridden system property or the custom property of a DSM.

When you override the behavior of a system property, you must provide a valid expression on the **Property Configuration** tab. The **Format String** field is a combination of regex capture groups and literal characters. The string is used to populate system properties by one or more values that are captured from events, and with more formatting characters or injected information. For example, you might want to parse an IP address and a port to combine them both into a string. If your regular expression (regex) has two capture groups, you can combine them by using this format string: `$1:$2`.

**NOTE**: The DSM Editor allows capture group references of 1 through 9 in any given match. If you reference any capture group above 9, the log source extension might not work correctly.

You must configure each custom property that you create. You must provide a valid expression and capture group for a custom property on the **Property Configuration** tab. You can also define selectivity and enable or disable your expression.

## Referencing Capture Strings by using Format String Fields

Use the **Format String** field on the **Property Configuration** tab to reference capture groups that you defined in the regex. Capture groups are referenced in their order of precedence.

A capture group is any regex that is enclosed within parenthesis. A capture group is referenced with an $n notation, where n is a group number that contains a regular expression (regex). You can define multiple capture groups.

For example, you have a payload with company and host name variables.

**"company":"companyname", "hostname":"localhost.com" "company":"companyname", "hostname":"username.com"**

You can customize the host name from the payload to display *example.hostname.com* by using capture groups:

1. In the **regex** field, enter the following regular expression:

   **"company":"(.*?)".*"hostname":"(.*?)"**

2. In the **Format String** field, enter the capture group $1.$2 where $1 is the value for the company variable and $2 is the value for the host name in the payload.

   The following output is given:

```
companyname.localhost.com

companyname.username.com
```

## Regex for Well-structured Logs

Well-structured logs are a style of event formatting that is composed of a set of properties and are presented in the following way:

*<name_of_property_1> <assignment_character> <value_of_property_1> <delimiter_character> <name_of_property_2> <assignment_character> <value_of_property_2> <delimiter_character> <name_of_property_3> <assignment_character> <value_of_property_3> <delimiter_character>...*

Use the following general guidelines:

- The *<assignment_character>* either '=' or ':' or a multi-character sequence such as '->'.
- The *<delimiter_character>* either a white space character (space or tab) or a list delimiter, such as a comma or semi-colon.
- The *<value_of_property>* and sometimes *<name_of_property>* are encapsulated in quotation marks or other wrapping characters.

For example, consider a simple login event that is generated by a device or an application. The device might report on the account of a user who logged in, the time the login occurred, and the IP address of the computer from which the user logged in. A name/value pair-style event might look like this snippet:

**<13>Sep 09 22:40:40 192.0.2.12 action=login accountname=JohnDoe clientIP= 192.0.2.24 timestamp=01/09/2016 22:40:39 UTC**

**NOTE**: The string "<13>Sep 09 22:40:40 192.0.2.12" is a syslog header. The string is not part of the event body.

The following table shows how the properties of the well-structured log example above, can be captured:

**Table 29: Regex for Capturing Properties Of a Well-structured Log**

| Property | Regex |
|---|---|
| action | action=(.*?)\t |

**Table 29: Regex for Capturing Properties Of a Well-structured Log *(Continued)***

| Property | Regex |
|---|---|
| accountname | accountname=(.*?)\t |
| clientIP | clientIP=(.*?)\t |
| timestamp | timestamp=(.*?)\t |

The patterns that are enclosed within the brackets denote the capture group. Each regex in the table captures everything after the equal sign (=) and before the next tab character.

## Regex for Natural Language Logs

Natural language logs are presented in a sentence-like form and each event type might look different.

For example, a simple login event can be presented in the following form:

```
<13>Sep 09 22:40:40 192.0.2.12 Account JohnDoe initiated a login action from 192.0.2.24 at
01/09/2016 22:40:39 UTC
```

The following table shows how the properties of the natural language log in the example above, can be captured:

**Table 30: Regex for Capturing Properties Of a Natural Language Log**

| Property | Regex |
|---|---|
| action | initiated a (.*?) action |
| accountname | Account (.*?) initiated |
| clientIP | from (.*?) at |
| timestamp | at (.*?) |

**NOTE**: Writing regex for natural language logs requires you to look at the static information that surrounds the value you want to capture before you create the capture group.

## Expressions in JSON Format for Structured Data

Structured data in JSON format contains one or more properties, which are represented as a key-value pair.

You can extract properties from an event data presented in JSON format by writing a JSON expression that matches the property. The JSON expression must be a path in the format of /"<name of top-level field>".

For example, you have an event data formatted in JSON:

```
{ "action": "login", "user": "John Doe" }
```

or an event that has a nested JSON format, such as:

```
{ "action": "login", "user": { "first_name": "John", "last_name": "Doe" } }
```

To extract properties from event data, choose one of the following methods:

- To extract the 'user' property for event data that is formatted in JSON, type the expression /"user" in the **Expression** field.
- To extract the 'last_name' of the user for an event that has a nested JSON format, type the expression /"user"/"last_name" in the **Expression** field.

## JSON Keypath Expressions

To uniquely identify the fields that you want to extract from a JSON object, your JSON expression must follow specific JSON keypath conventions.

Use the following guidelines for your JSON keypath expressions:

- A forward slash (/) must be at the start of all JSON keypaths. All paths must start at the beginning of the root JSON object. Subsequent slashes in the keypath indicate access to fields that are nested in the JSON object.

- Field names must be enclosed in double quotation marks.

  A valid path might look like the following example:

  ```
  /"object"/"nestedObject"/"furtherNestedObject"/"desiredPropertyName"
  ```

- Square brackets indicate the handling of JSON arrays.

  If you do not supply an index in the square brackets, the entire body of the array is extracted. If you supply an index in the square bracket, that index in the array is extracted or nested. Arrays begin at a zero index, where 0 is the first index in the array, 1 is the second index in the array, and so on.

  In the following keypath example, the JSON parser looks into the second index of the "object" JSON array, and then within that array index, looks for a field called "desiredPropertyName".

  ```
  /"object"[1]/"desiredPropertyName"
  ```

- Within log source extensions, you can supply and combine together multiple JSON keypaths to give a single result; this convention excludes custom properties. You can also choose to include literal text. Each of the JSON keypaths must be enclosed in curly braces.

  Consider the following example:

  ```
  {/"object"/"nestedObject"/"desiredPropertyName1"}
  {/"object"/"nestedObject"/"desiredPropertyName2"}
  ```

  You get a parsed value from the first JSON keypath, a literal text space, and then a parsed value from the second JSON keypath.

  **Example:** The following two examples show how to extract data from a JSON object:

  The following table shows the values that are extractable from the keypaths in that sample object:

- Simple case of a JSON object:

  ```
  [{"name":"object1","field1":"value1"},
  {"name":"object2","field2":"value2"},
  {"name":"object3","field3":"value3"}]
  ```

The following table shows the values that are extractable from the keypaths in that sample object:

**Table 31: Keypaths from the Simple JSON Object**

| Keypath | Description | Value |
|---|---|---|
| /[] | Extracts the entire JSON array from the root of the JSON object | `[{"name":"object1","field1":"value1"}, {"name":"object2","field2":"value2"}, {"name":"object3","field3":"value3"}]` |
| /[1]/"name" | Extracts the value for the attribute called "name" from the JSON object at index 1 in the root JSON array. | object2 |

- Complex case of a JSON object:

```
<13>May 22 10:15:41 log.test.com {"module":"CPHalo","version":"1.0","user_name":"user123", "event_type":"File integrity scan request created", "event_category":"File Integrity Scanning Management","srcName":"domain-lab-123", "timestamp":"2018-12-02T15:36:17.486","user": {"email":"user123@example.com","first_name":"fname", "last_name":"lname","alias":["alias name","alias1","name"]},"client_ip":"10.12.12.12", "server_id":"12317412471421274","server_reported_fqdn":"None","actor_country":"USA", "server_group_name":"Example Server","server_platform":"Linux", "message":"A file integrity monitoring scan was requested for Linux server domain-lab-123 (10.13.13.13) by Halo user user123@example.com from IP address 10.12.12.12 (USA).", "type":"fim_scan_request_created","id":"c2e8bf72- b74f-11e2-9055-870a490fcfb6"}
```

The following table shows the values that are extractable from the keypaths in that sample object:

**Table 32: Keypaths from the Complex JSON Object**

| Keypath | Description | Value |
|---|---|---|
| /"user_name" | Extracts value of the "user_name" attribute from the root of the JSON object. | user123 |

Table 32: Keypaths from the Complex JSON Object *(Continued)*

| Keypath | Description | Value |
|---|---|---|
| /"user"/"alias"[] | Extracts the entire JSON array called "alias" that is nested under the "user" JSON object. | ["alias name","alias1","name"] |
| /"user"/"alias"[0] | Extracts the value at index 0 within the "alias" JSON array that is nested under the "user" JSON Object. | alias name |
| /"user"/'first_name" | Extracts the value of the property called "first_name" that is nested under the "user" JSON Object. | fname |
| {/"user"/"first_name"}.{/"user"/ "last_name"} | Extracts the value of the property called "first_name" that is nested under the "user" JSON object, then inserts a literal '.' character, and then extracts the value of the property called "second_name" that is nested under the "user" JSON object.<br>Pertains only to log source extensions and non-custom properties within the DSM Editor. This operation is not possible in custom properties. | fname.lname |

Table 32: Keypaths from the Complex JSON Object *(Continued)*

| Keypath | Description | Value |
|---|---|---|
| {/"user"/"alias"[1]}@{/"client_ip"} | Extracts the value at index 1 of the "alias" JSON array that is nested under the "user' JSON object, inserts a literal '@' character, and then extracts the value of the property called "client_ip" under the root JSON object.<br>Pertains only to log source extensions and non-custom properties within the DSM Editor. This operation is not possible in custom properties. | alias1@10.12.12.12 |

## Expressions in LEEF Format for Structured Data

Structured data in LEEF format contains one or more properties, which are represented as key-value pairs.

You can extract properties from an event that is presented in LEEF format by writing a LEEF expression that matches the property. Valid LEEF expressions are in the form of either a single key reference, or a special LEEF header field reference.

For example, you have an event that is formatted in LEEF V1.0, such as:

```
LEEF:1.0|ABC Company|SystemDefender|1.13|console_login|devTimeFormat=yyyy-
MM-dd'T'HH:mm:ss.SSSZ
devTime=2017-10-18T11:26:03.060+0200 usrName=flastname name=Firstname
Lastname
authType=interactivePassword src=192.168.0.1
```

or an event that is formatted in LEEF V2.0 with the caret (^) separator character, such as:

```
LEEF:2.0|ABC Company|SystemDefender|1.13|console_login|^|devTimeFormat=yyyy-
MMdd'T'HH:mm:ss.SSSZ^
```

```
devTime=2017-10-18T11:26:03.060+0200^usrName=flastname^name=Firstname
Lastname
^authType=interactivePassword^src=192.168.0.1
```

You can extract a property or a header key property from the event by choosing one of the following methods:

1. To extract the 'usrName' property, enter **usrName** in the **LEEF Key** field.

   The possible keys that can be extracted are:

   - devTimeFormat
   - devTime
   - usrName
   - name
   - authType
   - src

2. To extract a header key property, type the key in the following format in the **LEEF Key** field:

   **$eventid$**
   The LEEF header values can be extracted by using the following expressions:

   - $leefversion$
   - $vendor$
   - $product$
   - $version$
   - $eventid$

## Expressions in CEF Format for Structured Data

Structured data in CEF format contains one or more properties, which are represented as key-value pairs.

You can extract properties from an event that is presented in CEF format by writing a CEF expression that matches the property. Valid CEF expressions are in the form of either a single key reference, or a special CEF header field reference.

For example, you have an event that is formatted in CEF:

```
CEF:0|ABC Company|SystemDefender|1.13|console_login|Console Login|1|
start=Oct 18 2017 11:26:03
duser=jsmith cs1=John Smith cs1Label=Person Name cs2=interactivePassword
cs2Label=authType src=10.1.1.1
```

You can extract a property or a header key property from the event by choosing one of the following methods:

1. To extract the 'cs1' property, type cs1 in the **CEF Key** field.

   The possible keys that can be extracted are:

   - start
   - duser
   - cs1
   - cs1Label
   - cs2
   - cs2Label
   - src

2. To extract a header key property, type the key in the following format in the **CEF Key** field:

   **$id$**
   The CEF header values can be extracted by using the following expressions:

   - $cefversion$
   - $vendor$
   - $product$
   - $version$
   - $id$
   - $name$
   - $severity$

## Expressions in Name Value Pair Format for Structured Data

Structured data in Name Value Pair format contains one or more properties, which are represented as key-value pairs.

You can extract properties from an event that is in Name Value Pair format by writing an expression that matches the property. Valid Name Value Pair expressions are in the form of a single key reference.

The following example shows an event that is in Name Value Pair format:

```
Company=ABC
Company;Product=SystemDefender;Version=1.13;EventID=console_login;Username=jsmith;Name=John
Smith;authType=interactivePassword;
```

To extract the Username property, type `Username` in the **Expression** field.

In the **Value Delimiter** field, enter the key-value delimiter that is specific for your payload. In this example, the keyvalue delimiter is an equal sign (=).

In the **Delimiter** field, enter the delimiter between key-value pairs that is specific for your payload. In this example, the delimiter between key-value pairs is a semicolon (;).

Matches in the payload are highlighted in the event data in the **Workspace** of the DSM Editor.

## Expressions in Generic List Format for Structured Data

Structured data in Generic List format contains one or more properties, which are represented as list items.

You can extract properties from an event that is in Generic List format by writing an expression that matches the property. Valid Generic List expressions are in the form of a $<number> notation. For example, $0 represents the first property in the list, $1 is the second property, and so on.

The following example shows an event that is in Generic List format:

```
ABC Company;1.13;console_login;jsmith;John Smith;interactivePassword;
```

To extract the first property in the list, type $0 in the **Expression** field.

In the **Delimiter** field, enter the delimiter between list items that is specific for your payload. In this example, the delimiter between list items is a semicolon (;).

Matches in the payload are highlighted in the event data in the **Workspace** of the DSM Editor.

## Expressions in XML Format for structured data

Structured data in XML format contains one or more properties, which are represented as key-value pairs.

You can extract properties from an event that is in XML format by writing an expression that matches the property. Valid XML expressions are in the form of a single key reference.

Enter the path to the XML field that you want to use to populate the property's value. An XML key path must begin with a forward slash (/) to indicate the root of the XML object, and be followed by one or more XML field names within double quotation marks.

The following example shows an event that is in XML format:

```
<EPOEvent><MachineInfo><MachineName>NEPTUNE<
/MachineName><MachineName>VALUE23</MachineName>
<AgentGUID>9B-B5-A6-A8-37-B</AgentGUID>
<IPAddress someattrib="someattribvalue">192.0.2.0</IPAddress><OSName>Windows 7</
OSName><UserName>I am a test user</UserName></MachineInfo></EPOEvent>
```

To capture the value nested in the top-level OSName object, type **/"EPOEvent"/"MachineInfo"/"OSName"** in the **Expression** field.

To capture the attribute value, use a period (.) after the key path. For example, to capture some attribute value, type **/"EPOEvent"/"MachineInfo"/"IPAddress".someattrib** in the **Expression** field.

To combine multiple fields together with multiple paths, use set brackets to enclose each. For example,

```
{/"EPOEvent"/"MachineInfo"/"OSName"} {/"EPOEvent"/"MachineInfo"/"MachineName"[1]}
```

To capture the value that is nested within multiple tags with the same name, use [0], [1], and so on, after the key path. For example, to capture VALUE23, type, **/"EPOEvent"/"MachineInfo"/"MachineName"[1]** in the **Expression** field.

Matches in the payload are highlighted in the event data in the **Workspace** of the DSM Editor.

RELATED DOCUMENTATION

Opening the DSM Editor | 225

Configuring a Log Source Type | 226

Custom Property Definitions in the DSM Editor | 232

# Opening the DSM Editor

YYou can open the DSM Editor from the **Log activity** tab, or if you are an administrator, you can open it from the Admin tab. For example, if events that are sent to the system are not handled properly, you can select the event data from the **Log Activity** tab and send it to the DSM Editor. For events that are not yet sent to the system, you must be an administrator and access the DSM Editor from the **Admin** tab.

1. To open the DSM Editor from the Admin tab, follow these steps:

   a. On the navigation menu
   (
   ≡
   ), click **Admin**.

   b. In the **Data Sources** section, click **DSM Editor**.

2. To open the DSM Editor from the **Log Activity** tab, follow these steps:

   a. On the navigation menu
   (
   ≡
   ), click **Admin**.

   b. Pause the incoming results and then highlight one or more events.

   **NOTE**: If more than one event from two or more log sources are selected, you are prompted to select which log source type you want to operate on. You can select only a single log source type, and only the events from log activity that match the selected log source type are automatically added to the workspace.

   c. On the navigation menu, select **Actions** > **DSM Editor**

## RELATED DOCUMENTATION

Configuring a Log Source Type | 226

Custom Property Definitions in the DSM Editor | 232

Event Mapping | 237

# Configuring a Log Source Type

With the DSM Editor, you can configure a new log source type or use an existing one in JSA.

1. On the navigation menu (☰), click **Admin**.
2. In the **Data Sources** section, click **DSM Editor**.
3. Create a log source type or select an existing log source type:
   - To create a new log source type, click **Create New** and follow the prompts.
   - To locate an existing log source type, use the **filter** field and then click **Select**.

### RELATED DOCUMENTATION

Custom Property Definitions in the DSM Editor | **232**

Event Mapping | **237**

Exporting Contents from the DSM Editor | **240**

# Configuring Property Autodetection for Log Source Types

When you enable **Property Autodetection**, new properties are automatically generated to capture all fields that are in the events that the selected log source type receives. Configure property autodetection of new properties for a log source type so that you do not need to manually create a custom property for each instance.

By default, **Property Autodetection** for a log source type is disabled.

1. In the DSM Editor, select a log source type or create a new one from the **Select Log Source** page.
2. Click the **Configuration** tab.
3. Click **Enable Auto Property Discovery**.

**NOTE**: Property autodetection works only for structured data that is in JSON, CEF, LEEF, or Name Value Pair format.

4. Select the structured data format for the log source type from the **Property Discovery Format** list.

   If you choose **Name Value Pair**, in the **Delimiter In Name Value Pairs** section, enter the delimiter used to separate each name and value, and the delimiter used to separate each Name Value Pair. Delimiters for each pair are automatically created.

5. To enable new properties to use in rules and searches, click **Enable Properties for use in Rules and Search Indexing**.

6. In the **Autodetection Completion Threshold** field, set the number of consecutive events to inspect for new properties.

   If no new properties are discovered when the number of consecutive events are inspected, the discovery process is considered complete and Property Autodetection is disabled. You can manually re-enable Property Autodetection at any time. A threshold value of 0 means that the discovery process perpetually inspects events for the selected log source type.

7. Click **Save**.

   The newly discovered properties appear in the Properties tab of the DSM Editor.

RELATED DOCUMENTATION

Custom Property Definitions in the DSM Editor | 232

Event Mapping | 237

Exporting Contents from the DSM Editor | 240

# Configuring Log Source Autodetection for Log Source Types

Configure Log Source Autodetection for a log source type so that you don't need to manually create a log source for each instance. Log source autodetection configuration also helps to improve the accuracy of detecting devices that share a common format, and can improve pipeline performance by avoiding the creation of incorrectly detected devices.

In JSA 7.3.2, upgrades from previous versions enable global configuration settings, which are stored in the JSA database. The global settings are initially set based on the contents of the TrafficAnalysisConfig.xml file in **/opt/qradar/conf/** directory on the JSA Console. If this file was customized before you upgrade to 7.3.2, the customizations are preserved. If different customizations exist on other managed hosts in the deployment, these customizations aren't carried over to the global settings. You can still enable per-event processor autodetection settings by using the configuration file method. Disable global autodetection settings in **Admin > System & License Management > Edit Managed Host > Component Management**.

When Log Source Autodetection is enabled, if you create a custom log source type that has many instances in your network, you don't need to manually create a log source for each instance.

You can also use the JSA REST API or a command line script to enable and disable which log source types are autodetected. If you use a smaller number of log source types, you can configure which log sources are autodetected to improve the speed of detection

If you choose to revert to the file-based (non-global) settings, you can only configure autodetection by using the config file. The DSM Editor and REST API work only with global settings. Move any custom autodetection configurations to global settings and to the DSM Editor.

Tune the autodetection engine so that log sources aren’t incorrectly identified as the wrong type. Incorrect detection happens when a DSM incorrectly recognizes events as its own even though they don't originate from the type of system that the DSM corresponds to. For example, if the events are formatted similarly to the events the DSM supports, or they contain the same keywords that the DSM is looking for. It can also happen even if a DSM exists for the system that is generating the events, if the events are so similar that the incorrect DSM is successful at parsing the events like the correct DSM. That DSM incorrectly recognizes the events as its own, and the autodetection engine creates a log source that isn't of the correct type.

For example, if you have both Linux and AIX systems in your JSA deployment, and most of them are Linux. You can reduce the **Minimum Successful Events for Autodetection** parameter or the **Minimum Successful Events for Autodetection** for Linux. Alternatively, increase the **Minimum Successful Events for Autodetection** parameter or the **Minimum Successful Events for Autodetection** parameter for AIX.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **Data Sources** section, click **DSM Editor**.
3. Select a log source type or create a new one from the **Select Log Source Type** window.
4. Click the **Configuration** tab.
5. Click **Enable Auto Property Autodetection**.
6. Configure the following parameters:

Table 33: Global System Notifications Window Parameters

| Parameter | Description |
|---|---|
| **Log Source Name Template** | Enter the template for setting the name of autodetected log sources.<br>Two variables can be used:<br>• **$$DEVICE_TYPE$$** corresponds to the log source type name.<br>• **$$SOURCE_ADDRESS$$** corresponds to the source address the events originate from. |
| **Log Source Description Template** | Enter the template for setting the description of autodetected log sources<br>Two variables can be used:<br>• **$$DEVICE_TYPE$$** corresponds to the log source type name.<br>• **$$SOURCE_ADDRESS$$** corresponds to the source address the events originate from. |
| **Minimum Successful Events for Autodetection** | The minimum number of events from an unknown source that must be successfully parsed for autodetection to occur. |
| **Minimum Success Rate for Autodetection** | The minimum parsing success percentage for events from an unknown source for autodetection to occur |
| **Attempted Parse Limit** | The maximum number of events from an unknown source to attempt before abandoning autodetection |
| **Consecutive Failed Parse Limit** | The number of consecutive events from an unknown source to abandon autodetection. |

7. Click **Save**.

RELATED DOCUMENTATION

Custom Property Definitions in the DSM Editor | 232

Event Mapping | 237

Exporting Contents from the DSM Editor | 240

# Configuring DSM Parameters for Log Source Types

Use the DSM Editor to configure the DSM parameters for your log source type.

If your log source type has DSM parameters, you can use the DSM Editor to edit the parameters.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **Data Sources** section, click **Admin**.
3. Select a log source type or create a new one from the **Select Log Source Type** window.
4. Click the Configuration tab, and then click **Display DSM Parameters Configuration**.
5. Configure the parameters.

   The Default parameters apply to all instances of this DSM in your deployment which do not have an Event Collector-specific override. To set different parameter values for this DSM for a specific Event Collector, select it from the Event Collector list to override the Default settings.
6. Click **Save**.

RELATED DOCUMENTATION

Custom Log Source Types | 231

Custom Property Definitions in the DSM Editor | 232

Event Mapping | 237

Exporting Contents from the DSM Editor | 240

# Custom Log Source Types

IN THIS SECTION

Use the DSM Editor to create and configure a custom log source type to parse your events. If you create a log source type for your custom applications and systems that don't have a supported DSM, JSA analyzes the data in the same way that it does for supported DSMs.

You can select events from the **Log Activity** tab and send them directly to the DSM Editor to be parsed. Or you can open the DSM Editor from the **Admin** tab to create and configure a new log source type.

Complete the fields in the DSM Editor with the correct structured data to parse relevant information from the events. JSA uses the **Event Category and Event ID** fields to map a meaning to the event. The Event ID is a mandatory field that defines the event, and the category breaks down the event further. You can set the **Event Category** to the Device Type name, or you can leave it as unknown. If you leave the **Event Category** as unknown, you must set it to unknown for any event mappings that you create for this log source type.

Use the DSM Editor to map your Event ID/Event Category combinations that you are parsing from your events. Enter the Event ID/Event Category combination into the new entry in the **Event Mapping** tab. You can choose a categorization of the previously created QID map entry that is relevant to your event, or click **Choose QID** to create a new map entry.

## Creating a Custom Log Source Type to Parse Events

If you have events that are imported into JSA, you can select the events on which you want to base your custom log source type and send them directly to the DSM Editor.

1. Click the **Log Activity** tab.

2. Pause the incoming results and then highlight one or more events.

> **NOTE**: You can select only a single log source type, and only the events from log activity that match the selected log source type are automatically added to the workspace.

3. On the navigation menu, select **Actions** > **DSM Editor**, and choose one of the following options:

   - If you are parsing known events, select your log source type from the list.

   - If you are parsing stored events, click **Create New**. Enter a name for your log source type in the **Log Source Type Name** field and click **Save**.

4. In the **Properties** tab, select the **Override system properties** checkbox for the properties that you want to edit.

RELATED DOCUMENTATION

Event Mapping | 237

Exporting Contents from the DSM Editor | 240

Configuring a Log Source Type | 226

# Custom Property Definitions in the DSM Editor

IN THIS SECTION

You can define a custom property and reuse the same property in a separate DSM. Use these properties in searches, rules, and to allow specific user-defined behavior for parsing values into those fields.

Where relevant, each custom property has a set of configuration options that includes selectivity and data parsing. Each custom property definition within a DSM configuration is an ordered group of expressions that consists of regular expressions, a capture group, an optional selectivity configuration, and an enabled or disabled toggle button. You can't modify the **Name**, **Field type**, **Description**, **optimize** fields, or any advanced options for a custom property on the **Properties** tab in the DSM Editor.

A custom property is shared across all DSMs, while specific implementations for reading values from payloads are at the DSM level.

Selectivity is specified when you configure an expression to run only when certain conditions are met.

**NOTE**: The **Capture Group** field of a custom property cannot be assigned a value greater than the number of capture groups in the regex.

## Selectivity

In the DSM Editor, you can restrict running a custom property to certain criteria for better performance.

The following are the types of restrictions:

- **By high-level category and low-level category** -- A property is evaluated only when the high-level and low-level categories match a specific combination. For example, a property is evaluated only when the event is known to have a high-level category of **Authentication** and a low-level category of **Admin Logout**.
- **By specific QID** -- A property is evaluated only when the event that is seen maps to a specific QID. For example, when the event maps to a QID of **Login Failed**, the property is evaluated.

## Creating a Custom Property

In the DSM Editor, you can define a custom property for one or more log sources, whose events do not fit into the JSA normalized event model. For example, a system property might fail to capture data from some applications, operation systems, databases, and other systems.

You can create custom property for data that does not fit into JSA system properties. Use the custom properties in searches and test against them in rules.

1. On the **Properties** tab in the DSM Editor, click the **Add** (**+**).
2. To create a new custom property definition, use the following steps:
   a. On the **Choose a Custom Property Definition to Express** page, select **Create New**.
   b. On the **Create a new Custom Property Definition** page, configure the parameters in the following table.

Table 34: Custom Property Parameters

| Parameter | Description |
|---|---|
| Name | A descriptive name for the custom property that you create. |
| Field Type | The default is **Text**.<br>**NOTE**: When you select **Number** or **Date** from the **Field Type** list, extra fields are displayed. |
| Enable this Property for use in Rules and Search Indexing | When this option is enabled, during the parsing stage of the event pipeline, JSA attempts to extract the property from events immediately as they enter the system. Other components downstream in the pipeline such as rules, forwarding profiles and indexing can use the extracted values. Property information is persisted along with the rest of the event record and doesn't need to be extracted again when it is retrieved as part of a search or report. This option enhances performance when the property is retrieved, but can have a negative impact on performance during the event parsing process, and impacts storage.<br>When this option is not enabled, JSA extracts the property from the events only when they are retrieved or viewed.<br>**NOTE**: To use Custom Properties in rule tests, forwarding profiles, or for search indexing, make sure that this checkbox is selected. Rule evaluation, event forwarding, and indexing occur before events are written to disk, so the values must be extracted at the parsing stage. |
| Use number format from a Locale | This field displays when you select **Number** from the **Field Type** list. If you select the **Use number format from a Locale** check box, you must select an **Extracted Number Format** from the list. |
| Extracted Date/ Time Format | This field displays when you select **Date** from the **Field Type** list. You must provide a datetime pattern that matches how the datetime appears in the original event.<br>For example, 'MMM dd YYYY HH:mm:ss' is a valid datetime pattern for a time stamp like 'Apr 17 2017 11:29:00'. |

Table 34: Custom Property Parameters *(Continued)*

| Parameter | Description |
|---|---|
| Locale | This field displays when you select **Date** from the **Field Type** list. You must select the locale of the event.<br>For example, if the locale is **English**, it will recognizes 'Apr' as a short form of the month 'April'. But if the event is presented in French and the month token is 'Avr' (for Avril), then set the locale to a **French** one, or the code does not recognize it as a valid date. |

c. If you want to extract the property from events as they enter the system, select the **Enable this property for use in Rules and Search indexing** check box.

d. Click **Save**.

3. To use an existing custom property, use the following steps:

a. On the **Choose a Custom Property Definition to Express** page, search for an existing custom property from the Filter Definitions field.

b. Click **Select** to add the custom property.

## Expressions

You can define expressions for custom properties in the DSM Editor. Expressions are the mechanism that defines the behavior of a property. The main component of an expression is a valid regex or JSON. The data that makes up an expression depends on the property type.

For a custom property, you can choose only one capture group from the regex.

## Configuring a custom property expression

You can use different expressions to capture various custom properties for the same event. You can also use a combination of expression types to capture the same custom property if that property can be captured from multiple event formats.

JSA supports the following custom property expression types:

- Regex
- JSON
- LEEF
- CEF
- Name Value Pair
- Generic List
- XML

1. On the **Properties** tab, locate and select the custom property. Custom properties display the word **Custom** next to them to differentiate them from system properties.

2. Select an expression type from the **Expression Type** list and define a valid expression for it.

   **NOTE**:

   - For Regex, the expression must be a valid java-compatible regular expression. Case-insensitive matching is supported only by using the (?i) token at the beginning of the expression. The (?i) token is saved in the log source extension .xml file. To use other expressions, such as (?s), manually edit the log source extension .xml file.

   - For JSON, the expression must be a path in the format of */"<name of top-level field>"* with additional */"<name of sub-field>"* subobjects to capture subfields if any.

   - To capture the value of a key-value pair for LEEF and CEF, set the expression to the key.

   - To capture the value of a header field, set the expression to the corresponding reserved word for that header field.

3. If the expression type is Regex, select a capture group.

4. To limit an expression to run against a specific category, click **Edit** to add selectivity to the custom property, and select a **High Level Category** and a **Low Level Category**.

5. To limit an expression to run against a specific event or QID, click **Choose Event** to search for a specific QID.

6. In the **Expression** window, click **Ok**.

7. To add multiple expressions and reorder them, follow these steps:

   - Click Add (+) in the expressions list.

- Drag expressions in the order that you want them to run.

## Deleting a custom property expression

You can delete a custom property expression in the DSM Editor. If you delete a custom property expression, only the expression is deleted. The custom property is not deleted.

1. On the Admin tab, click **DSM Editor**.
2. In the **Select Log Source Type** window, choose a log source type and click **Select**.
3. In the Log Source Type pane, select the custom property with the expression that you want to delete.
4. In the Property Configuration section, select the expression that you want to delete and click the delete icon

   .
5. Click **Delete**.

RELATED DOCUMENTATION

Event Mapping | 237

Exporting Contents from the DSM Editor | 240

Configuring a Log Source Type | 226

# Event Mapping

IN THIS SECTION

In the DSM Editor, the event mapping shows all the event ID and category combinations that are in the system.

An event mapping represents an association between an event ID and category combination and a QID record (referred to as event categorization). Event ID and category values are extracted by DSMs from events and are then used to look up the mapped event categorization or QID. Event categorizations store extra metadata for the event that might not exist verbatim in the raw event data, such as a human-readable name and description, a severity value, or a low level category assignment. Low-level categorization and severity are useful for search and rule definitions.

**NOTE**: For multi-tenant environments, any user-defined mapping or event categorization information that is defined in the DSM Editor becomes visible across all tenants. You must ensure that no tenant-specific data is put in any event categorization names or descriptions.

## Identity Properties for Event Mappings

Identity data is a special set of system properties that includes **Identity Username**, **Identity IP**, **Identity NetBIOS Name**, **Identity Extended Field**, **Identity Host Name**, **Identity MAC**, **Identity Group Name**.

When identity properties are populated by a DSM, the identity data is forwarded to the asset profiler service that runs on the JSA console. The asset profiler is used to update the asset model, either by adding new assets or by updating the information on existing assets, including the **Last User** and **User Last Seen** asset fields when an **Identity Username** is provided.

JSA DSMs can populate identity data for certain events, such as those that establish an association or disassociation between identity properties. This association or disassociation is for performance and also for certain events that provide new or useful information that is needed for asset updates. For example, a login event establishes a new association between a user name and an asset (an IP address, a MAC address, or a host name, or a combination of them). The DSM generates identity data for any login events that it parses, but subsequent events of different types that involve the same user, provide no new association information. Therefore, the DSM does not generate identity for other event types.

Also, the DSMs for DHCP services can generate identity data for DHCP assigned events because these events establish an association between an IP address and a MAC address. DSMs for DNS services generate identity information for events that represents DNS lookups because these events establish an association between an IP address and a host name or DNS name.

You can configure the DSM Editor to override the behavior of the identity properties. However, unlike other system properties, overridden identity property has no effect unless it is linked to specific Event ID or Event Category combinations (event mappings). When identity property overrides are configured, you can go to the **Event Mappings** tab and select an event mapping to configure specific identity properties for that event. Only identity properties that are available and captured by the configured property regex are populated for an event.

**NOTE**: The **Identity Username** property is unique and cannot be independently configured. If any identity properties are enabled for a particular event mapping, then the **Identity Username** property is automatically populated for the event from the available **Username** property value.

## Creating an Event Map and Categorization

An event mapping is an event ID and category combination you use to map an event to a QID. With the DSM Editor, you can create a new event mapping to map all unknown events to an entry in the QID map. Also, you can remap existing ones to either a newly created event categorization (QIDs) or to an existing one in the system.

1. To add an event mapping, click the Add (**+**) icon on the **Event Mapping** tab of DSM Editor.

2. Ensure that values are entered for the **Event ID** and **Event Category** fields.

3. To create a new event categorization, use the following steps:

   a. From the **Create a new Event Mapping** window, click **Choose QID**.

   b. On the **QID Records** window, click **Create New QID Record**.

   c. Enter values for the **Name**, **Description** fields, and select a **Log Source Type**, a **High Level Category**, a **Low Level Category**, and a **Severity**.

   d. Click **Save** to create the new event categorization.

4. To use an existing event categorization, use the following steps:

   a. From the **Create a new Event Mapping** window, click **Choose Event**.

   b. Search for an existing event categorization on the **Event Categorizations** window.

   c. Select a **High Level category**, **Low Level category**, **Log Source Type** or **QID**. Results are shown in the **Search Results** pane.

   d. Click **Ok** to add the event category.

RELATED DOCUMENTATION

Exporting Contents from the DSM Editor | 240

Configuring a Log Source Type | 226

Custom Property Definitions in the DSM Editor | 232

# Exporting Contents from the DSM Editor

IN THIS SECTION

You can use a content management tool script to export custom content that is created in the DSM Editor. Contents can be exported from one JSA deployment and imported into another JSA deployment. You can also export custom content to external media.

The DSM Editor produces the following content types:

**Table 35: DSM Editor Content Types**

| Custom content type | String | ID |
|---|---|---|
| Custom properties | customproperty | 6 |
| Log source type | sensordevicetype | 24 |
| Log source extensions | deviceextension | 16 |
| Custom QidMap entries | qidmap | 27 |

The contentManagement.pl script is in the **/opt/qradar/bin** directory

## Exporting Contents As a Package

You can use the content management tool script to search for specific content that is created in the DSM Editor. These contents are exported as a package.

1. Use SSH to log in to JSA as the root user.
2. To search for specific content items to export, type the following command:

   **./contentManagement.pl -a search -c [content_type] -r [regex]**

   For example, to search for the content items of a log source type, type the following command:

   **/opt/qradar/bin/contentManagement.pl -a search -c 24 -r "*<search_name>*"**
3. Create a text file that lists the content that you want to export.

   Each line must include the custom content type followed by a comma-separated list of unique IDs for that type.

   For example, to export three log source types with ID 24, ID 26, and ID 95, all custom properties, create a text file with the following entries:

   **sensordevicetype, 24,26,95**
4. Export the content items as a package by using the following command:

   **/opt/qradar/bin/contentManagement.pl -a export -c package -f *<source_file>***

## Exporting Content for Single Custom Property

You can use the content management tool script to export content for each custom property that is created from the **Properties** tab in the DSM Editor.

When you use the DSM Editor to create custom properties, a **customproperty** entity is produced for each custom property that is created.

1. Use SSH to log in to JSA as the root user.
2. To search for specific content to export, type the following command:

   **./contentManagement.pl -a search -c [content_type] -r [regex]**

   For example, to search for the content of a custom property, type the following command:

   **/opt/qradar/bin/contentManagement.pl -a search -c 6 -r "*<name_of_custom_property>*"**

3. To export a custom property content, type the following command:

   ```
   /opt/qradar/bin/contentManagement.pl -a export -c [content_type] -i [content_identifier]
   ```

## RELATED DOCUMENTATION

Configuring a Log Source Type | 226

Custom Property Definitions in the DSM Editor | 232

Event Mapping | 237

8

CHAPTER

# Using Reference Data in JSA

# Using Reference Data in JSA

IN THIS SECTION

Use reference data collections to store and manage business data that you want to correlate against the events and flows in your JSA environment. You can add business data or data from external sources into a reference data collection, and then use the data in JSA searches, filters, rule test conditions, and rule responses.

Reference data collections are stored on the JSA console, but the collections are regularly copied to each managed host. For best performance on data lookups, the managed host caches the most frequently referenced data values.

## External Threat Intelligence Data

You can use reference data collections to integrate indicator of compromise (IOC) data from third-party vendors into JSA. JSA uses IOC data to detect suspicious behavior faster, which helps security analysts investigate threats and respond to incidents more quickly.

For example, you can import IOC data, such as IP addresses, DNS names, URLs, and MD5s, from open source or subscription-based threat data providers, and correlate it with events and incidents on your network.

## Business Data

Reference data collections can contain business data that is specific to your organization, such as a list of users with privileged system access. Use the business data to create blocklists and allowlists.

For example, use a reference set that contains the user IDs of terminated employees to prevent them from logging in to the network. Or, you can use business data to build an allowlist that allows only a limited set of IP addresses to do specific functions.

# Types Of Reference Data Collections

JSA has different types of reference data collections that can handle different levels of data complexity. The most common types are reference sets and reference maps.

If you want to use the same reference data in both JSA and JSA Risk Manager, use a reference set. You can't use other types of reference data collections with JSA Risk Manager.

**Table 36: Types of Reference Data Collections**

| Type of collection | Description | How to use | Examples |
|---|---|---|---|
| Reference set | A collection of unique values. | Use a reference set to compare a property value against a list, such as IP addresses or user names. | To verify whether a login ID that was used to log in to JSA is assigned to a user, create a reference set with the **LoginID** parameter. |
| Reference map | A collection of data that maps a unique key to a value. | Use a reference map to verify a unique combination of two property values. | To correlate user activity on your network, create a reference map that uses the **LoginID** parameter as a key, and the **Username** as a value. |
| Reference map of sets | A collection of data that maps a key to multiple values. Every key is unique and maps to one reference set. | Use a reference map of sets to verify a combination of two property values against a list. | To test for authorized access to a patent, create a map of sets that uses a custom event property for **Patent ID** as the key, and the **Username** parameter as the value. Use the map of sets to populate a list of authorized users. |

Table 36: Types of Reference Data Collections *(Continued)*

| Type of collection | Description | How to use | Examples |
|---|---|---|---|
| Reference map of maps | A collection of data that maps one key to another key, which is then mapped to a single value. Every key is unique and maps to one reference map. | Use a reference map of maps to verify a combination of three property values. | To test for network bandwidth violations, create a map of maps that uses the **Source IP** parameter as the first key, the **Application** parameter as the second key, and the **Total Bytes** parameter as the value. |
| Reference table | A collection of data that maps one key to another key, which is then mapped to a single value. The second key is assigned a data type. | Use a reference table to verify a combination of three property values when one of the properties is a specific data type. | Create a reference table that stores **Username** as the first key, **Source IP** as the second key with an assigned **cidr** data type, and **Source Port** as the value. |

RELATED DOCUMENTATION

Reference Sets Overview | 246

Creating Reference Data Collections by Using the Command Line | 254

Creating Reference Data Collections with the APIs | 259

# Reference Sets Overview

IN THIS SECTION

Use reference sets in JSA to store data in a simple list format.

You can populate the reference set with external data, such as indicators of compromise (IOCs), or you can use it to store business data, such as IP addresses and user names, that is collected from events and flows that occur on your network.

A reference set contains unique values that you can use in searches, filters, rule test conditions, and rule responses. Use rules to test whether a reference set contains a data element, or configure the rule response to add data to a reference set. For example, you can create a rule that detects when an employee accesses a prohibited website, and configure the rule response to add the employee's IP address or user name to a reference set.

For more information about configuring rule responses to add data to a reference set, see the *Juniper Secure Analytics Users Guide*.

Reference sets are the only type of reference data collection that you can manage in JSA. You can also use the "Creating Reference Data Collections by Using the Command Line" on page 254 and the "Creating Reference Data Collections with the APIs" on page 259 to manage reference sets.

## Adding, Editing, and Deleting Reference Sets

Use a reference set to compare a property value, such as an IP address or user name, against a list. You can use reference sets with rules to keep watch lists. For example, you can create a rule to detect when an employee accesses a prohibited website and then add that employee's IP address to a reference set.

After you add data to the reference set, the **Number of Elements** and **Associated Rules** parameters are automatically updated.

When you edit a reference set, you can change the data values, but you cannot change the type of data that the reference set contains.

Before a reference set is deleted, JSA runs a dependency check to see whether the reference set has rules that are associated with it.

**NOTE**: If you use techniques to obfuscate data on the event properties that you want to compare to the reference set data, use an alphanumeric reference set and add the obfuscated data values.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Reference Set Management**.

3. To add a reference set:

   a. Click **Add** and configure the parameters.

      Learn more about reference set parameters:

      The following table describes each of the parameters that are used to configure a reference set.

      **Table 37: Reference Set Parameters**

| Parameter | Description |
|---|---|
| **Name** | The maximum length of the reference set name is 255 characters. |
| **Type** | Select the data types for the reference elements. You can't edit the **Type** parameter after you create a reference set.<br>The **IP** type stores IPv4 addresses. **Alphanumeric (Ignore Case)** automatically changes any alphanumeric value to lowercase.<br>To compare obfuscated event and flow properties to the reference data, you must use an alphanumeric reference set. |
| **Time to Live of elements** | Specifies when reference elements expire. If you select the Lives Forever default setting, the reference elements don't expire.<br>If you specify an amount of time, indicate whether the time-to-live interval is based on when the data was first seen, or was last seen.<br>JSA removes expired elements from the reference set periodically (by default, every 5 minutes). |

**Table 37: Reference Set Parameters *(Continued)***

| Parameter | Description |
|---|---|
| **When elements expire** | Specifies how expired reference elements are logged in the `qradar.log` file when they are removed from the reference set.<br>The **Log each element in a separate log entry** option triggers an **Expired ReferenceData element** log event for each reference element that is removed. The event contains the reference set name and the element value.<br>The **Log elements in one log entry** option triggers one **Expired ReferenceData element** log event for all reference elements that are removed at the same time. The event contains the reference set name and the element values.<br>The **Do not log elements** option does not trigger a log event for removed reference elements. |

b. Click **Create**.

4. Click **Edit** or **Delete** to work with existing reference sets.

> **TIP**: To delete multiple reference sets, use the **Quick Search** text box to search for the reference sets that you want to delete, and then click **Delete Listed**.

## Viewing the Contents Of a Reference Set

View information about the data elements in the reference set, such as the domain assignment, the expiry on the data, and when the element was last seen in your network.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Reference Set Management**.
3. Select a reference set and click **View Contents**.
4. Click the **Content** tab to view information about each data element.

**TIP**: Use the search field to filter for all elements that match a keyword. You can't search for data in the **Time To Live** column.

Learn more about the data elements:

The following table describes the information that is shown for each data element in the reference set.

**Table 38: Information About the Reference Set Data Elements**

| | Description |
|---|---|
| **Domain** | Domain-specific reference data can be viewed by tenant users who have access to the domain, MSSP Administrators, and users who do not have a tenant assignment. Users in all tenants can view shared reference data. |
| **Value** | The data element that is stored in the reference set. For example, the value might show user names or IP addresses. |
| **Origin** | Shows the user name when the data element is added manually, and the file name when the data was added by importing it from an external file. Shows the rule name when the data element is added in response to a rule. |
| **Time to Live** | The time that is remaining until this element is removed from the reference set. |
| **Date Last Seen** | The date and time that this element was last detected on your network. |

5. Click the **References** tab to view the rules that use the reference set in a rule test or in a rule response.

**Table 39: Content Tab Parameters**

| Parameter | Description |
|---|---|
| **Rule Name** | Name of the rule that is configured to use the reference set. |
| **Group** | The group that the rule belongs to. |

Table 39: Content Tab Parameters *(Continued)*

| Parameter | Description |
|---|---|
| **Category** | Shows if the rule is a custom rule rule. |
| **Type** | Shows **event**, **flow**, **common**, or **offense** to indicate the type of data that the rule is tested against. |
| **Enabled** | A rule must be enabled for the custom rule engine to evaluate it. |
| **Response** | The responses that are configured for this rule. |
| **Origin** | **System** indicates a default rule.<br>**Modified** indicates that a default rule was customized.<br>**User** indicates a user-created rule. |

6. To view or edit an associated rule, double-click the rule in the **References** list and complete the rule wizard.

## Adding Elements to a Reference Set

Add elements to a reference set when you want JSA to compare a property to the element value. Use JSA to manually add elements to a reference set, or to import elements from a **.csv** file.

To import elements, make sure that the **.csv** file is stored locally.

Domain-specific reference data can be viewed by tenant users who have access to the domain, MSSP Administrators, and users who do not have a tenant assignment. Users in all tenants can view shared reference data.

You can assign reference data to a specific domain. Domain-specific reference data can be viewed by tenant users who have access to the domain, MSSP Administrators, and users who do not have a tenant assignment. Users in all tenants can view shared reference data. For example, MSSP users who are not administrators can view reference data that is assigned to a domain.

1. On the navigation menu
   (

≡
), click **Admin**.

2. In the **System Configuration** section, click **Reference Set Management**.
3. Select the reference set that you want to add the elements to, and click **View Contents**.
4. Click the **Content** tab.
5. To add data elements manually, follow these steps:

   a. Click **Add** and configure the parameters.

      Valid port values are 0 - 65535. Valid IP addresses are between 0 and 255.255.255.255.

      > **NOTE**: If you use data obfuscation techniques on the event properties that you want to compare to the reference set data, you must use an alphanumeric reference set that contains the obfuscated data values.

   b. Click **Add**.
6. To add elements from a **.csv** file, follow these steps:

   a. Click **Import**.

   b. Click **Select File** and browse to select the **.csv** file that you want to import.

      The **.csv** file must be formatted with all items comma-separated on a single line, or with each item on a separate line. A delimiter is not required when each item is on a separate line.

   c. Select the **Domain** that you want to add the reference set data to.

   d. Click **Import**.

      The import adds the content of the text file to the reference set.

## Exporting Elements from a Reference Set

Export reference set elements to a **.csv** file when you want to include the information in reports, or share the information with people who don't use JSA.

1. On the navigation menu
   (

☰

), click **Admin**.

2. In the **System Configuration** section, click **Reference Set Management**.
3. Select the reference set that you want to export, and click **View Contents**.
4. Click the **Content** tab, and click **Export**.
5. Choose whether to open the file immediately, or save the file, and then click **OK**.

## Deleting Elements from a Reference Set

You might need to delete elements from a reference set when an element is added to the reference set in error, or when you no longer need to compare the element with other JSA properties. For example, you might need to remove an asset that was mistakenly added to an asset exclusion blocklist.

1. On the navigation menu

   (

   ☰

   ), click **Admin**.
2. In the **System Configuration** section, click **Reference Set Management**.
3. Select the reference set that contains the elements that you want to delete, and click **View Contents**.
4. Click the **Content** tab and choose one of the following options:
   - To delete a single element, select the element from the list, and click **Delete**.
   - To delete multiple elements, use the search box to filter the list to show only the elements that you want to delete, and then click **Delete Listed**.

RELATED DOCUMENTATION

Creating Reference Data Collections by Using the Command Line | **254**

Creating Reference Data Collections with the APIs | **259**

Examples for Using Reference Data Collections | **263**

# Creating Reference Data Collections by Using the Command Line

IN THIS SECTION

Use the command line to manage reference data collections that cannot be managed in JSA, such as reference maps, map of sets, map of maps, and tables. Although it's easier to manage reference sets using JSA, use the command line when you want to schedule management tasks.

Use the `ReferenceDataUtil.sh` script to manage reference sets and other types of reference data collections.

When you use an external file to populate the reference data collection, the first non-comment line in the file identifies the column names in the reference data collection. Each line after that is a data record that gets added to the collection. While the data type for the reference collection values is specified when the collection is created, each key is an alphanumeric string.

The following table shows examples of how to format data in an external file that is to be used for populating reference maps.

**Table 40: Formatting Data in an External File to Be Used for Populating Reference Data Collections**

| Type of reference collection | Data formatting examples |
|---|---|
| Reference map | `key1,data`<br>`key1,value1`<br>`key2,value2` |
| Reference map of sets | `key1,data`<br>`key1,value1`<br>`key1,value2` |

Table 40: Formatting Data in an External File to Be Used for Populating Reference Data Collections *(Continued)*

| Type of reference collection | Data formatting examples |
|---|---|
| Reference map of maps | `key1,key2,data`<br>`map1,key1,value1`<br>`map1,key2,value2` |

You can also create reference data collections by using the **/reference_data** endpoint in the JSA RESTful API.

1. Using SSH, log in to JSA as the root user.

2. Go to the **/opt/qradar/bin** directory.

3. To create the reference data collection, type the following command:

   **./ReferenceDataUtil.sh create *name* [SET | MAP | MAPOFSETS | MAPOFMAPS | REFTABLE] [ALN | NUM | IP | PORT | ALNIC | DATE] [-timeoutType=[FIRST_SEEN | LAST_SEEN]] [-timeToLive=]**

4. To populate the map with data from an external file, type the following command:

   **./ReferenceDataUtil.sh load *name filename* [-encoding=...] [-sdf=" ... "]**

Here are some examples of how to use the command line to create different types of reference data collections:

- Create an alphanumeric map:

  **./ReferenceDataUtil.sh create testALN MAP ALN**

- Create a map of sets that contains port values that will age out 3 hours after they were last seen:

  **./ReferenceDataUtil.sh create testPORT MAPOFSETS PORT -timeoutType=LAST_SEEN -timeToLive='3 hours'**

- Create a map of maps that contains numeric values that will age out 3 hours 15 minutes after they were first seen:

  **./ReferenceDataUtil.sh create testNUM MAPOFMAPS NUM -timeoutType=FIRST_SEEN -timeToLive='3 hours 15 minutes'**

- Create a reference table where the default format is alphanumeric:

```
./ReferenceDataUtil.sh create testTable REFTABLE ALN -
keyType=ipKey:IP,portKey:PORT,numKey:NUM,dateKey:DATE
```

Log in to JSA to create rules that add data to your reference data collections. You can also create rule tests that detect activity from elements that are in your reference data collection.

## Command Reference for Reference Data Utilities

You can manage your reference data collections by using the `ReferenceDataUtil.sh` utility on the command line. The following commands are available to use with the script.

### Create

Creates a reference data collection.

**name**

The name of the reference data collection.

**[SET | MAP | MAPOFSETS | MAPOFMAPS | REFTABLE]**

The type of reference data collection.

**[ALN | ALNIC | NUM | IP | PORT | DATE]**

The type of data in the reference set.

- **ALN** specifies alphanumeric values. This data type supports IPv4 and IPv6 addresses.
- **ALNIC** specifies alphanumeric values, but rule tests ignore the case. This data type supports IPv4 and IPv6 addresses.
- **NUM** specifies numeric values.
- **IP** specifies IP addresses. This data type supports only IPv4 address.
- **PORT** specifies port addresses.
- **DATE** specifies date values.

**[-timeoutType=[FIRST_SEEN | LAST_SEEN]]**

Specifies whether the amount of time the data elements remain in the reference data collection is from the time the element was first seen or last seen.

[-TimeToLive="]

The amount of time the data elements remain in the reference data collection.

[-keyType=name:elementType,name:elementType,...]

A mandatory **REFTABLE** parameter of consisting of key name to **ELEMENTTYPE** pairs.

[-key1Label="]

An optional label for key1, or the primary key. A key is a type of information, such as an IP address.

[-valueLabel="]

An optional label for the values of the collection.

### Update

Updates a reference data collection.

name

The name of the reference data collection.

[-timeoutType=[FIRST_SEEN | LAST_SEEN]]

Specifies whether the amount of time the data elements remain in the reference data collection is from the time the element was first seen or last seen.

[-timeToLive="]

The amount of time the data elements remain in the reference data collection.

[-keyType=name:elementType,name:elementType,...]

A mandatory **REFTABLE** parameter of consisting of key name to **elementType** pairs.

[-key1Label="]

An optional label for key1.

[-valueLabel="]

An optional label for the values of the collection.

### Add

Adds a data element to a reference data collection.

name

The name of the reference data collection.

<value> <key1> <key2>

The key value pair that you want to add. The keys are alphanumeric strings.

- MAP and MAPOFSETS require Key 1.
- MAPOFMAPS and REFTABLE require Key 1, and the second-level Key 2.

[-sdf=" ... "]

The Simple Date Format string that is used to parse the date data.

### Delete

Deletes an element from a reference data collection.

name

The name of the reference data collection.

<value> <key1> <key2>

The key value pair that you want to add. The keys are alphanumeric strings.

- MAP and MAPOFSETS require Key 1.
- MAPOFMAPS and REFTABLE require Key 1, and the second-level Key 2.

[-sdf=" ... "]

The Simple Date Format string that is used to parse the date data.

### Remove

Removes a reference data collection.

name

The name of the reference data collection.

### Purge

Purges all elements from a reference data collection.

name

The name of the reference data collection.

### List

Lists elements in a reference data collection.

name

The name of the reference data collection.

[displayContents]

Lists all elements in the specified reference data collection.

### Listall

Lists all elements in all reference data collection.

| | |
|---|---|
| [displayContents] | Lists all elements in all reference data collections. |

### Load

Populates a reference data collection with data from an external **.csv** file.

| | |
|---|---|
| name | The name of the reference data collection. |
| filename | The fully qualified file name to be loaded. Each line in the file represents a record to be added to the reference data collection. |
| [-encoding=...] | Encoding that is used to read the file. |
| [-sdf=" ... "] | The Simple Date Format string that is used to parse the date data. |

RELATED DOCUMENTATION

Creating Reference Data Collections with the APIs | 259
Examples for Using Reference Data Collections | 263
Reference Sets Overview | 246

# Creating Reference Data Collections with the APIs

You can use the application program interface (API) to manage JSA reference data collections.

1. Use a web browser to access **https:// *<Console IP>*/api_doc** and log in as the administrator.
2. Select the latest iteration of the JSA API.
3. Select the **/reference_data** directory.
4. To create a new reference set, follow these steps:

   a. Select **/sets**.

   b. Click **POST** and enter the relevant information in the **Value** fields.

      Learn more about the parameters to create a reference set:

The following table provides information about the parameters that are required to create a reference set:

**Table 41: Parameters - Reference Set**

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| element_type | query | (required) | String | text/plain | String <one of: ALN, NUM, IP, PORT, ALNIC, DATE> |
| name | query | (required) | String | text/plain | String |
| fields | query | (required) | String | text/plain | field_one (field_two, field_three), field_four |
| time_to_live | query | (optional) | String | text/plain | String |
| timeout_type | query | (optional) | String | text/plain | String <one of: UNKNOWN, FIRST_SEEN, LAST_SEEN> |

c. Click **Try It Out!** to finish creating the reference data collection and to view the results.

5. To create a new reference map, follow these steps:

a. Click **/maps**.

b. Click **POST** and enter the relevant information in the **Value** fields.

Learn more about the parameters to create a reference map:

The following table provides information about the parameters that are required to create a reference map:

**Table 42: Parameters - Reference Map**

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| element_type | query | (required) | String | text/plain | String <one of: ALN, NUM, IP, PORT, ALNIC, DATE> |

Table 42: Parameters - Reference Map *(Continued)*

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| name | query | (required) | String | text/plain | String |
| fields | query | (optional) | String | text/plain | field_one (field_two, field_three), field_four |
| key_label | query | (optional) | String | text/plain | String |
| time_to_live | query | (optional) | String | text/plain | String |
| timeout_type | query | (optional) | String | text/plain | String <one of: UNKNOWN, FIRST_SEEN, LAST_SEEN> |
| value_label | query | (optional) | String | text/plain | String |

c. Click **Try It Out!** to finish creating the reference data collection and to view the results.

6. To create a new reference map of sets, follow these steps:

a. Select **/map_of_sets**.

b. Click **POST** and enter the relevant information in the **Value** fields.

Learn more about the parameters to create a reference map of sets:

The following table provides information about the parameters that are required to create a reference map of sets:

Table 43: Parameters - Reference Map of Sets

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| element_type | query | (required) | String | text/plain | String <one of: ALN, NUM, IP, PORT, ALNIC, DATE> |
| name | query | (required) | String | text/plain | String |

Table 43: Parameters - Reference Map of Sets *(Continued)*

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| fields | query | (optional) | String | text/plain | field_one (field_two, field_three), field_four |
| key_label | query | (optional) | String | text/plain | String |
| time_to_live | query | (optional) | String | text/plain | String |
| timeout_type | query | (optional) | String | text/plain | String <one of: UNKNOWN, FIRST_SEEN, LAST_SEEN> |
| value_label | query | (optional) | String | text/plain | String |

c. Click **Try It Out!** to finish creating the reference data collection and to view the results.

7. To create a new reference table or map of maps, follow these steps:

a. Click **/tables**.

b. Click **POST** and enter the relevant information in the **Value** fields.

Learn more about the parameters to create a reference table or a map of maps:

The following table provides information about the parameters that are required to create a reference table or a map of maps:

Table 44: Parameters - Reference Table

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| element_type | query | (required) | String | text/plain | String <one of: ALN, NUM, IP, PORT, ALNIC, DATE> |
| name | query | (required) | String | text/plain | String |
| fields | query | (optional) | String | text/plain | field_one (field_two, field_three), field_four |

Table 44: Parameters - Reference Table *(Continued)*

| Parameter | Type | Value | Data Type | MIME Type | Sample |
|---|---|---|---|---|---|
| key_name_types | query | (optional) | Array | application/ json | [ { "element_type": "String <one of: ALN, NUM, IP, PORT, ALNIC, DATE>", "key_name": "String" }] |
| outer_key_label | query | (optional) | String | text/plain | String |
| time_to_live | query | (optional) | String | text/plain | String |
| timeout_type | query | (optional) | String | text/plain | String <one of: UNKNOWN, FIRST_SEEN, LAST_SEEN> |

c. Click **Try It Out!** to finish creating the reference data collection and to view the results.

RELATED DOCUMENTATION

Examples for Using Reference Data Collections | 263
Reference Sets Overview | 246
Creating Reference Data Collections by Using the Command Line | 254

# Examples for Using Reference Data Collections

IN THIS SECTION

These examples show how you can use reference data collections to track and store data that you want to use in JSA searches, filters, rule test conditions, and rule responses.

## Tracking Expired User Accounts

Use reference data collections to identify stale data, such as expired user accounts, in your JSA environment.

By default, reference data remains in JSA until it is removed. However, when you create a reference data collection, you can configure JSA to remove the data after a specified period of time.

When the data element expires, JSA automatically deletes the value from the reference data collection and triggers an event to track the expiry.

1. Create a reference set to keep track of the time since a user last logged in.

    a. Set the **Time to Live of elements** to represent the period of time after which an unused user account is considered expired.

    b. Select the **Since last seen** button.

2. Create a custom event rule to add login data, such as the **username**, to the reference set.

    > **NOTE**: JSA tracks the **Date Last Seen** for each data element. If no data is added for a particular user within the time-to-live period, the reference set element expires, and a **Reference Data Expiry** event is triggered. The event contains the reference set name and the username that is expired.

3. Use the **Log Activity** tab to track the **Reference Data Expiry** events.

Use the reference set data in searches, filters, rule test conditions, and rule responses.

## Integrate Dynamic Data from External Sources

Large enterprise organizations can use reference data collections to share information about their IT assets with the security teams that manage the JSA deployment.

For example, the Information Technology (IT) team maintains an asset management database that includes information about all the network assets. Some of the information, such as the IP addresses for the web servers, changes frequently.

Once a week, the IT team exports the list of IP addresses for all of the web servers that are deployed in the network and provides the list to the security team. The security team imports the list into a reference set, which can then be used in rules, searches, and reports to provide more context to the events and flows that are processed by JSA.

## RELATED DOCUMENTATION

Reference Sets Overview | 246

Creating Reference Data Collections by Using the Command Line | 254

Creating Reference Data Collections with the APIs | 259

9

CHAPTER

# User Information Source Configuration

# User Information Source Configuration

Configure your JSA system to collect user and group information from Identity and Access Management endpoints.

JSA uses the information that is collected from the endpoints to enrich the user information that is associated with the traffic and events that occur on your network.

## User Information Source Overview

IN THIS SECTION

You can configure a user information source to enable user information collection from an Identity and Access Management endpoint.

An Identity and Access Management endpoint is a product that collects and manages electronic user identities, group memberships, and access permissions. These endpoints are called user information sources.

Use the following utilities to configure and manage user information sources:

- **Tivoli Directory Integrator -** You must install and configure a Tivoli Directory Integrator on a non-JSA host.
- **UISConfigUtil.sh -** Use this utility to create, retrieve, update, or delete user information sources. You can use user information sources to integrate JSA using a Tivoli Directory Integrator server.
- **GetUserInfo.sh -** Use this utility to collect user information from a user information source and store the information in a reference data collection. You can use this utility to collect user information on demand or on a schedule.

## User Information Sources

A user information source is a configurable component that enables communication with an endpoint to retrieve user and group information.

JSA systems support the following user information sources:

**Table 45: Supported Information Sources**

| Information Source | Information that is collected |
|---|---|
| MicrosoftWindows Active Directory (AD), version 2008 - MicrosoftWindows AD is a directory service that authenticates and authorizes all users and computers that use your Windows network. | • full_name<br>• user_name<br>• user_principal_name<br>• family_name<br>• given_name<br>• account_is_disabled<br>• account_is_locked<br>• password_is_expired<br>• password_can_not_be_changed<br>• no_password_expired<br>• password_does_not_expire |

## Reference Data Collections for User Information

This topic provides information about how reference data collections store data collected from user information sources.

When JSA collects information from a user information source, it automatically creates a reference data collection to store the information. The name of the reference data collection is derived from the user information source group name. For example, a reference data collection that is collected from MicrosoftWindows AD might be named Domain Admins.

The reference data collection type is a Map of Maps. In a Reference Map of Maps, data is stored in records that map one key to another key, which is then mapped to a single value.

For example:

- `#`
- `# Domain Admins`
- `# key1,key2,data`
- `smith_j,Full Name,John Smith`
- `smith_j,account_is_disabled,0`
- `smith_j,account_is_locked,0`
- `smith_j,account_is_locked,1`
- `smith_j,password_does_not_expire,1`

## Integration Workflow Example

After user and group information is collected and stored in a reference data collection, there are many ways in which you can use the data in JSA.

You can create meaningful reports and alerts that characterize user adherence to your company's security policies.

**NOTE**: If you want to collect application security logs, you must create a Device Support Module (DSM). For more information, see the *Configuring DSMs Guide*.

## User Information Source Configuration and Management Task Overview

To initially integrate user information sources, you must perform the following tasks:

1. Configure a Tivoli Directory Integrator server. See "Configuring the Tivoli Directory Integrator Server" on page 270.
2. Create and manage user information sources. See "Creating and Managing User Information Source" on page 273.

3. Collect user information. See "Collecting User Information" on page 277.

RELATED DOCUMENTATION

Configuring the Tivoli Directory Integrator Server | 270

Creating and Managing User Information Source | 273

Collecting User Information | 277

# Configuring the Tivoli Directory Integrator Server

For JSA to integrate with user information sources, you must install and configure a Tivoli Directory Integrator on a non-JSA host.

No configuration is required on your JSA system; however, you must access your Console to obtain the **QRadarIAM_TDI.zip** file. Then, install and configure a Tivoli Directory Integrator server on a separate host. Create and import a self-signed certificate.

When you extract the **QRadarIAM_TDI.zip** file on the Tivoli Directory Integrator server, the TDI directory is automatically created. The TDI directory includes the following files:

- **QradarIAM.sh**, which is the TDI start up script for Linux
- **QradarIAM.bat**, which is the TDI start up script for Microsoft Windows
- **QradarIAM.xml**, which is the TDI xml script and must be stored in the same location as the **QradarIAM.properties** file
- **QradarIAM.properties**, which is the properties file for TDI xml script

When you install Tivoli Directory Integrator, you must configure a name for the Solutions directory. This task requires you to access the Solutions directory. Therefore, in the task steps, *<solution_directory>* refers to the name that you gave to the directory.

The following parameters are used to create and import certificates:

**Table 46: Certification Configuration Parameters**

| Parameter | Description |
|---|---|
| *<server_ip_address>* | Defines the IP address of the Tivoli Directory Integrator server. |

**Table 46: Certification Configuration Parameters *(Continued)***

| Parameter | Description |
|---|---|
| *<days_valid>* | Defines the number of days that the certificate is valid. |
| *<keystore_file>* | Defines the name of the keystore file. |
| -storepass *<password>* | Defines the password for keystore. |
| - keypass *<password>* | Defines the password for the private/public key pair. |
| *<alias>* | Defines the alias for an exported certificate. |
| *<certificate_file>* | Defines the file name of the certificate. |

1. Install Tivoli Directory Integrator on a non-JSA host. For more information on how to install and configure Tivoli Directory Integrator, see your Tivoli Directory Integrator (TDI) documentation.
2. Using SSH, log in to your JSA Console as the root user.
   a. User name: **root**

   b. Password: ***<password>***
3. Copy the **QRadarIAM_TDI.zip** file to the Tivoli Directory Integrator server.
4. On the Tivoli Directory Integrator server, extract the **QRadarIAM_TDI.zip** file in the Solutions directory.
5. Configure your Tivoli Directory Integrator server to integrate with JSA.
   a. Open the Tivoli Directory Integrator ***<solution_directory>*/solution.properties** file.

   b. Uncomment the **com.ibm.di.server.autoload** property. If this property is already uncommented, note the value of the property.

   c. Choose one of the following options:

   - Change directories to the **autoload.tdi** directory, which contains the **com.ibm.di.server.autoload** property by default.

   - Create an **autoload.tdi** directory in the ***<solution_directory>*** to store the **com.ibm.di.server.autoload** property.

   d. Move the **TDI/QRadarIAM.xml** and **TDI/QRadarIAM.property** files from the Tivoli Directory Integrator directory to ***<solution_directory>*/autoload.tdi** directory or the directory you created in the previous step.

   e. Move the **QradarIAM.bat** and **QradarIAM.sh** scripts from the Tivoli Directory Integrator directory to the location from which you want to start the Tivoli Directory Integrator.

6. Create and import the self-signed certificate into the Tivoli Directory Integrator truststore.

   a. To generate a keystore and a private/public key pair, type the following command:

      - **keytool -genkey -dname cn=*<server_ip_address>* -validity *<days_valid>* -keystore *<keystore_file>* -storepass *<password>* - keypass *<password>***

      - For example, **keytool -genkey -dname cn=192.168.1.1 -validity 365 -keystore server.jks -storepass secret -keypass secret**

   b. To export the certificate from the keystore, type the following command:

      - **keytool -export -alias *<alias>* -file *<certificate_file>* - keystore *<keystore_file>* - storepass *<password>***

      - For example, **keytool -export -alias mykey -file server.cert -keystore server.jks -storepass secret**

   c. To import the primary certificate back into the keystore as the self-signed CA certificate, type the following command:

      - **keytool -import -trustcacerts -file *<certificate_file>* -keystore *<keystore_file>* -storepass *<password>* -alias *<alias>***

      - For example, **keytool -import -trustcacerts -file server.cert -keystore server.jks -storepass secret -alias mytrustedkey**

   d. Copy the certificate file to the **/opt/qradar/conf/trusted_certificates** on the JSA console.

7. Import the CA certificate into the Tivoli Directory Integrator truststore.

   a. To import the CA certificate into the keystore as the self-signed CA certificate, type the following command:

      - **keytool -import -trustcacerts -file *<certificate_file>* -keystore *<keystore_file>* -storepass *<password>* -alias *<alias>***

      - For example, **keytool -import -trustcacerts -file server.cert -keystore server.jks -storepass secret -alias mytrustedkey**

   b. Copy the CA certificate file to the **/opt/qradar/conf/trusted_certificates** on the JSA console.

8. Edit the **<solution_directory>/solution.properties** file to uncomment and configure the following properties:

- `javax.net.ssl.trustStore=`*`<keystore_file>`*
- `{protect}-javax.net.ssl.trustStorePassword=`*`<password>`*
- `javax.net.ssl.keyStore=`*`<keystore_file>`*
- `{protect}-javax.net.ssl.keyStorePassword=`*`<password>`*

> **NOTE**: The default unmodified password might be displayed in the following format: `{encr}EyHbak`. Enter the password as plain text. The password encrypts the first time that you start Tivoli Directory Integrator.

9. Start Tivoli Directory Integrator.

RELATED DOCUMENTATION

Creating and Managing User Information Source | 273

Collecting User Information | 277

User Information Source Overview | 267

# Creating and Managing User Information Source

IN THIS SECTION

Use the UISConfigUtil utility to create, retrieve, update, or delete user information sources.

## Creating a User Information Source

Use the UISConfigUtil utility to create a user information source.

Before you create a user information source, you must install and configure your Tivoli Directory Integrator server. For more information, see "Configuring the Tivoli Directory Integrator Server" on page 270.

When you create a user information source, you must identify the property values required to configure the user information source. The following table describes the supported property values:

**Table 47: Supported User Interface Property Values**

| Property | Description |
|---|---|
| tdiserver | Defines the host name of the Tivoli Directory Integrator server. |
| tdiport | Defines the listening port for the HTTP connector on the Tivoli Directory Integrator server. |
| hostname | Defines the host name of the user information source host. |
| port | Defines the listening port for the Identity and Access Management registry on the user information host. |
| username | Defines the user name that JSA and Log Manager use to authenticate to the Identity and Access Management registry. |
| password | Defines the password that is required to authenticate to the Identity and Access Management registry. |
| searchbase | Defines the base DN.<br>**NOTE**: All users that are referenced in all groups must be found in a search from the searchbase. |
| search filter | Defines the search filter that is required to filter the groups that are retrieved from the Identity and Access Management registry. |

1. Using SSH, log in to your JSA Console as the root user.

    a. User name: **root**

b. Password: **<password>**

2. To add a user information source, type the following command: **UISConfigUtil.sh add <name> -t <AD|ISAM|ISIM|ISFIM> [-d description] [-p prop1=value1,prop2=value2...,propn=valuen]**

   Where:

   - **<name>** Is the name of the user information source you want to add.
   - **<AD|ISAM|ISIM|ISFIM>** Indicates the user information source type.
   - **[-d description]** Is a description of the user information source. This parameter is optional.
   - **[-p prop1=value1,prop2=value2,...,propn=valuen]** Identifies the property values required for the user information source. For more information about the supported parameters, see "Creating a User Information Source" on page 274.

   For example:

   - **/UISConfigUtil.sh add "UIS_ISIM" -t ISIM -d "UIS for ISIM" -p "tdiserver=nc9053113023.tivlab.austin.ibm.com,tdiport=8080, hostname=vmibm7094.ottawa.ibm.com,port=389, username=cn=root,password=password,\"searchbase=ou=org,DC=COM\",\ "searchfilter=(| (objectClass=erPersonItem)(objectClass=erBPPersonItem) (objectClass=erSystemUser))\""**

## Retrieving User Information Sources

Use the UISConfigUtil utility to retrieve user information sources.

1. Using SSH, log in to your JSA Console as the root user.

   a. User name: **root**

   b. Password: **<password>**

2. Choose one of the following options:

   a. Type the following command to retrieve all user information sources: **UISConfigUtil.sh get <name>**

   b. Type the following command to retrieve a specific user information source: **UISConfigUtil.sh get <name>**

   Where **<name>** is the name of the user information source you want to retrieve.

   For example:

**[root@vmibm7089 bin]# .UISConfigUtil.sh get "UIS_AD"**

## Editing a User Information Source

Use the UISConfigUtil utility to edit a user information source.

1. Using SSH, log in to your JSA Console as the root user.

   a. User name: **root**

   b. Password: **<password>**

2. Type the following command to edit a user information source: **UISConfigUtil.sh update <name> -t <AD|ISAM|ISIM|ISFIM> [-d description] [-p prop1=value1,prop2=value2,...,propn=valuen]**

   Where:

   - **<name>** Is the name of the user information source you want to edit.

   - **<AD|ISAM|ISIM|ISFIM>** Indicates the user information source type. To update this parameter, type a new value.

   - **[-d description]** Is a description of the user information source. This parameter is optional. To update this parameter, type a new description.

   - **[-p prop1=value1,prop2=value2,...,propn=valuen]** Identifies the property values required for the user information source. To update this parameter, type **new properties**. For more information about the supported parameters, see "Creating a User Information Source" on page 274.

   For example:

   **./UISConfigUtil.sh update "UIS_AD_update" -t AD -d "UIS for AD" -p "searchbase=DC=local"**

## Deleting a User Information Source

Use the UISConfigUtil utility to delete a user information source.

1. Using SSH, log in to your JSA Console as the root user.

   a. User name: **root**

   b. Password: **<password>**

2. Type the following command to delete a user information source:

**UISConfigUtil.sh delete <name>**

Where **<name>** is the name of the user information source you want to delete.

The collected user information is stored in a reference data collection in the JSA database. If no reference data collection exists, a new reference data collection is created. If a reference data collection was previously created for this user information source, the reference map is purged of previous data and the new user information is stored. For more information about reference data collections, see "Reference Data Collections for User Information" on page 268.

RELATED DOCUMENTATION

Collecting User Information | **277**

User Information Source Overview | **267**

Configuring the Tivoli Directory Integrator Server | **270**

# Collecting User Information

Use the GetUserInfo utility to collect user information from the user information sources and store the data in a reference data collection.

Use this task to collect user information on demand. If you want to create automatic user information collection on a schedule, create a cron job entry. For more information about cron jobs, see your Linux documentation.

1. Using SSH, log in to your JSA Console as the root user.
   a. User name: **root**

   b. **<password>**
2. Type the following command to collect user information on demand:

   **GetUserInfo.sh <UISName>**

   Where **<UISName>** is the name of the user information source you want to collect information from.

The collected user information is stored in a reference data collection on the database. If no reference date collection exists, a new reference data collection is created. If a reference data collection was previously created for this user information source, the reference map is purged of previous data and the new user information is stored. For more information about reference data collections, see "Reference Data Collections for User Information" on page 268.

## RELATED DOCUMENTATION

User Information Source Overview | **267**

Configuring the Tivoli Directory Integrator Server | **270**

Creating and Managing User Information Source | **273**

10

CHAPTER

# Juniper Networks X-Force Integration

# Juniper Networks X-Force Integration

IN THIS SECTION

Juniper Networks X-Force security experts use a series of international data centers to collect tens of thousands of malware samples, analyze web pages and URLs, and run analysis to categorize potentially malicious IP addresses and URLs. X-Force Exchange is the platform for sharing this data, which can be used in JSA.

## X-Force Threat Intelligence Feed

You can integrate X-Force Exchange data into JSA to help your organization stay ahead of emerging threats by identifying and remediating undesirable activity in your environment before it threatens the stability of your network.

For example, you can identify and prioritize these types of incidents:

- A series of attempted logins for a dynamic range of IP addresses
- An anonymous proxy connection to a Business Partner portal
- A connection between an internal endpoint and a known botnet command and control
- Communication between an endpoint and a known malware distribution site

**NOTE**: X-Force integration allows you to use the X-Force Threat Intelligence data in JSA correlation rules and AQL queries. Access to the X-Force Exchange REST API is not included.

RELATED DOCUMENTATION

IBM QRadar Security Threat Monitoring Content Extension | 283
Juniper X-Force Exchange Plug-in for JSA | 285

# Enabling the X-Force Threat Intelligence Feed

You must enable the X-Force Threat Intelligence feed before you can use the enhanced content that is installed with the IBM QRadar Security Threat Monitoring Content Extension application.

JSA downloads approximately 30 MB of IP reputation data per day when you enable the X-Force Threat Intelligence feed.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **System Settings**.
3. Select **Yes** in the **Enable X-Force Threat Intelligence Feed** field.

Deploy the system setting changes to receive the data from the X-Force servers. For more information, see "Deploying Changes" on page 114.

RELATED DOCUMENTATION

Updating X-Force Data in a Proxy Server | **281**

Preventing X-Force Data from Downloading Locally | **282**

Juniper X-Force Exchange Plug-in for JSA | **285**

# Updating X-Force Data in a Proxy Server

JSA uses a reverse proxy lookup through an Apache server to collect data directly from Juniper X-Force Threat Intelligence servers on the Internet.

All JSA appliances in a deployment contact the Apache server to send cached requests. After the data is received by the JSA Console, the result is cached and replayed for all other managed hosts that make a request for new IP reputation data.

If a proxy is configured in your network, you must update the configuration to receive the X-Force data.

> **NOTE**: NTLM authentication is not supported.

1. Use SSH to log in to the JSA console.
2. Open the **/etc/httpd/conf.d/ssl.conf** file in a text editor.
3. Add the following lines before `</VirtualHost>`:

   **ProxyRemote https://license.xforce-security.com/ http:// *PROXY_IP:PROXY_PORT***

   **ProxyRemote https://update.xforce-security.com/ http:// *PROXY_IP:PROXY_PORT***
4. Update the IP address and port of the corporate proxy server to allow an anonymous connection to the X-Force security servers.
5. Save the changes to the **ssl.conf** file.
6. Restart the Apache server by typing the following command:

   **apachectl restart**

   Restarting the Apache server on the JSA console logs out all users and the managed hosts might produce error messages. Restart the Apache server during scheduled maintenance windows.

## RELATED DOCUMENTATION

Preventing X-Force Data from Downloading Locally | 282

Juniper X-Force Exchange Plug-in for JSA | 285

Enabling the X-Force Threat Intelligence Feed | 281

# Preventing X-Force Data from Downloading Locally

JSA downloads approximately 30 MB of IP reputation data per day. To stop JSA from downloading the X-Force data to your local system, disable the X-Force Threat Intelligence feed.

Before you disable the X-Force feed, ensure that the X-Force rules are disabled, and that you are not using X-Force functions in saved searches.

After the X-Force Threat Intelligence feed is disabled, the X-Force content is still visible in JSA, but you cannot use the X-Force rules or add X-Force functions to AQL searches.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. In the **System Configuration** section, click **System Settings**.

3. Select **No** in the **Enable X-Force Threat Intelligence Feed** field.

Deploy the system setting changes to receive the data from the X-Force servers.

RELATED DOCUMENTATION

Juniper X-Force Exchange Plug-in for JSA | 285
Enabling the X-Force Threat Intelligence Feed | 281
Updating X-Force Data in a Proxy Server | 281

# IBM QRadar Security Threat Monitoring Content Extension

**IN THIS SECTION**

The IBM QRadar Security Threat Monitoring Content Extension on the IBM Security App Exchange contains rules, building blocks, and custom properties that are intended for use with X-Force feed data.

The X-Force data includes a list of potentially malicious IP addresses and URLs with a corresponding threat score. You use the X-Force rules to automatically flag any security event or network activity data that involves the addresses, and to prioritize the incidents before you begin to investigate them.

The following list shows examples of the types of incidents that you can identify using the X-Force rules:

- when the *[source IP|destinationIP|anyIP]* is part of any of the following *[remote network locations]*
- when *[this host property]* is categorized by X-Force as *[Anonymization Servers|Botnet C&C|DynamicIPs|Malware|ScanningIPs|Spam]* with confidence value *[equal to] [this amount]*
- when *[this URL property]* is categorized by X-Force as *[Gambling|Auctions|Job Search|Alcohol|Social Networking|Dating]*

JSA downloads approximately 30 MB of IP reputation data per day when you enable the X-Force Threat Intelligence feed for use with the IBM QRadar Security Threat Monitoring Content Extension.

## Installing the IBM QRadar Security Threat Monitoring Content Extension Application

The IBM QRadar Security Threat Monitoring Content Extension application contains JSA content, such as rules, building blocks, and custom properties, that are designed specifically for use with X-Force data. The enhanced content can help you to identify and to remediate undesirable activity in your environment before it threatens the stability of your network.

Download the IBM QRadar Security Threat Monitoring Content Extension application from the IBM Security App Exchange.

To use X-Force data in JSA rules, offenses, and events, you must configure JSA to automatically load data from the X-Force servers to your JSA appliance.

To load X-Force data locally, enable the X-Force Threat Intelligence feed in the system settings. If new information is available when X-Force starts, the IP address reputation or URL database is updated. These updates are merged into their own databases and the content is replicated from the JSA console to all managed hosts in the deployment.

The X-Force rules are visible in the product even if the application is later uninstalled.

1. On the navigation menu
   (
   ≡
   ), click **Admin**

2. In the **System Configuration** section, click **Extensions Management**.

3. Upload the application to the JSA console by following these steps:

   a. Click **Add**.

   b. Click **Browse** and browse to find the extension.

   c. Click **Install immediately** to install the extension without viewing the contents.

   d. Click **Add**.

4. To view the contents of the extension, select it from the extensions list and click **More Details**.

5. To install the extension, follow these steps:

   a. Select the extension from the list and click **Install**.

   b. If the extension does not include a digital signature, or it is signed but the signature is not associated with the JSA Security certificate authority (CA), you must confirm that you still want to install it. Click **Install** to proceed with the installation.

c. Review the changes that the installation makes to the system.

d. Select **Overwrite** or **Keep existing data** to specify how to handle existing content items.

e. Click **Install**.

f. Review the installation summary and click **OK**.

The rules appear under the **Threats** group in the **Rules List** window. They must be enabled before they are used.

Enable the X-Force Threat Intelligence feed so that you can use the X-Force rules or add X-Force functions to AQL searches. For more information, see "Enabling the X-Force Threat Intelligence Feed" on page 281.

RELATED DOCUMENTATION

Juniper X-Force Exchange Plug-in for JSA | 285
Enabling the X-Force Threat Intelligence Feed | 281
Updating X-Force Data in a Proxy Server | 281

# Juniper X-Force Exchange Plug-in for JSA

IN THIS SECTION

The Juniper X-Force Exchange (XFE) is a sharing platform for threat intelligence that is used by security analysts, network security specialists, and security operations center teams.

The X-Force Exchange (XFE) plug-in provides the option to search the information on the X-Force Exchange website for IP addresses, URLs, CVEs, and web applications that are found in JSA.

For example, you can right-click a URL from a JSA event to see what data the X-Force Exchange contains about the URL.

You can also use the right-click lookup option to submit IP addresses or URL data from JSA searches, offenses, and rules to a public or private collection. The collection stores the information in one place as you use the data for more research.

Collections also contain a section that serves as a wiki-style notepad, where you can add comments or any free text that is relevant. You can use the collection to save X-Force reports, text comments, or any other content. An X-Force report has both a version of the report from the time that it was saved and a link to the current version of the report.

## Juniper X-Force Exchange Right-click Plug-in Installation

Install the X-Force Exchange plug-in on your JSA Console so that you have right-click functionality to access data in X-Force Exchange.

This procedure requires a web server restart from the **Admin** tab to load the plug-in after the RPM is installed. Restarting the web server logs out all JSA users, so it is advised that you install this plug-in during scheduled maintenance.

If your JSA system is version 2014.4 or later, the plug-in is already installed. Administrators can verify that the plug-in is installed by right-clicking on any IP address in JSA, and selecting **More Options >Plugin options**. If the Juniper X-Force Exchange lookup is displayed, then the plug-in is installed.

1. Download the X-Force Exchange right-click plug-in from https://ibm.biz/BdX4BW.

    a. Copy the RPM file to the JSA console.

    b. Type the following command to install the plug-in: **rpm -Uvh RightClick-XFE-7.2.<version>.x86_64.rpm**

2. Log in to the JSA console as an admin user.

3. On the navigation menu
   (
   ≡
   ), click **Admin**.

4. Select **Advanced >Restart Web Server**.

   After the web server restarts, the X-Force right-click plug-in is enabled for IP addresses in JSA for URL fields in the **Log Activity** tab.

5. Log in to the pop-up window for the X-Force Exchange website by using your IBM id, or continue as a guest.

   Guest users are not able to use all features on the X-Force Exchange website.

6. Close the browser window after the initial login to the Juniper X-Force Exchange website.

## RELATED DOCUMENTATION

Enabling the X-Force Threat Intelligence Feed | 281

Updating X-Force Data in a Proxy Server | 281

Preventing X-Force Data from Downloading Locally | 282

11

CHAPTER

# Managing Authorized Services

# Managing Authorized Services

You can configure authorized services on the **Admin** tab to authenticate an API call for your JSA deployment.

The JSA RESTful API uses authorized services to authenticate API calls to the JSA console. You can add or revoke an authorized service at any time. For more information about the RESTful API, see the *Juniper Secure Analytics API Guide*.

The **Manage Authorized Services** window provides the following information:

**Table 48: Parameters for Authorized Services**

| Parameter | Description |
|---|---|
| **Service Name** | The name of the authorized service. |
| **Authorized By** | The name of the user or administrator that authorized the addition of the service. |
| **Authentication Token** | The token that is associated with this authorized service. |
| **User Role** | The user role that is associated with this authorized service. |
| **Security Profile** | The security profile that is associated with this authorized service. |
| **Created** | The date that this authorized service was created. |
| **Expires** | The date and time that the authorized service expires. By default, the authorized service is valid for 30 days. |

# Viewing Authorized Services

The **Authorized Services** window displays a list of authorized services, from which you can copy the token for the service.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Authorized Services**.
3. From the **Manage Authorized Services** window, select the appropriate authorized service.

> **NOTE**: When you create an authorized service token, the token is displayed in the Authorized Service Created Successfully dialogue. From JSA 7.5.0, the authorized service token is no longer visible after you close the Authorised Service Created Successfully dialogue. You must copy the token to a secure location before you close the dialogue.

RELATED DOCUMENTATION

Adding an Authorized Service | **290**

Revoking Authorized Services | **291**

# Adding an Authorized Service

Use the **Add Authorized Service** window to add a new authorized service.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Authorized Services**.
3. Click **Add Authorized Service**.
4. In the **Service Name** field, type a name for this authorized service. The name can be up to 255 characters in length.
5. From the **User Role** list, select the user role that you want to assign to this authorized service. The user roles that are assigned to an authorized service determine the functions that this service can access on the JSA user interface.
6. From the **Security Profile** list, select the security profile that you want to assign to this authorized service. The security profile determines the networks and log sources that this service can access on the JSA user interface.

7. In the **Expiry Date** list, type or select a date that you want this service to expire. If an expiry date is not required, select **No Expiry**
8. Click **Create Service**.

   The confirmation message contains a token field that you must copy into your vendor software to authenticate with JSA.

RELATED DOCUMENTATION

Revoking Authorized Services | **291**

Viewing Authorized Services | **289**

# Revoking Authorized Services

Use the **Add Authorized Service** window to revoke an authorized service.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Authorized Services**.
3. From the **Manage Authorized Services** window, select the service that you want to revoke.
4. Click **Revoke Authorization**.

RELATED DOCUMENTATION

Viewing Authorized Services | **289**

Adding an Authorized Service | **290**

12

CHAPTER

# Backup and Recovery

# Backup and Recovery

You can back up and recover JSA configuration information and data.

You can use the backup and recovery feature to back up your event and flow data; however, you must restore event and flow data manually. For more information, see "Restoring Data" on page 309.

By default, JSA creates a backup archive of your configuration information daily at midnight. The backup archive includes configuration information, data, or both from the previous day.

You can use two types of backups: configuration backups and data backups.

Configuration backups include the following components:

- Application configuration
- Assets
- Custom logos
- Custom rules
- Device Support Modules (DSMs)
- Event categories
- Flow sources
- Flow and event searches
- Groups
- Index management information
- License key information
- Log sources
- Offenses
- Reference set elements
- Store and Forward schedules
- User and user roles information
- Vulnerability data (if JSA Vulnerability Manager is installed)

Data backups include the following information:

- Audit log information
- Event data
- Flow data
- Report data
- Indexes

The data backup does not included application data. To configure and manage backups for application data, see "Backing Up and Restoring App Data" on page 315.

RELATED DOCUMENTATION

Backup JSA Configurations and Data | 294
Manage Existing Backup Archives | 299
Restore JSA Configurations and Data | 301

# Backup JSA Configurations and Data

IN THIS SECTION

By default, JSA creates a backup archive of your configuration information daily at midnight. The backup archive includes your configuration information, data, or both from the previous day. You can customize this nightly backup and create an on-demand configuration backup, as required.

## Scheduling Nightly Backup

Use the **Backup Recovery Configuration** window to configure a night scheduled backup process.

By default, the nightly backup process includes only your configuration files. You can customize your nightly backup process to include data from your JSA Console and selected managed hosts. You can also customize your backup retention period, backup archive location, the time limit for a backup to process before timing out, and the backup priority in relation to other JSA processes.

**NOTE**: The nightly backup starts running at midnight in the timezone where the JSA Console is installed. If JSA automatic updates are scheduled to run at the same time, the performance of JSA might be impacted.

The Backup Recovery Configuration window provides the following parameters:

**Table 49: Backup Recovery Configuration Parameters**

| Parameter | Description |
|---|---|
| General Backup Configuration | |
| Backup Repository Path | Type the location where you want to store your backup file. The default location is **/store/backup**. This path must exist before the backup process is initiated. If this path does not exist, the backup process aborts.<br>If you modify this path, make sure the new path is valid on every system in your deployment.<br>• Active data is stored on the **/store** directory. If you have both active data and backup archives stored in the same directory, data storage capacity might easily be reached and your scheduled backups might fail. We recommend you specify a storage location on another system or copy your backup archives to another system after the backup process is complete. You can use a Network File System (NFS) storage solution in your JSA deployment. For more information on using NFS, see the *Configuring Offboard Storage Guide*. |
| Backup Retention Period (days) | Type or select the length of time, in days, that you want to store backup files. The default is 2 days.<br>This period of time only affects backup files generated as a result of a scheduled process. On-demand backups or imported backup files are not affected by this value. |
| Nightly Backup Schedule | Select a backup option. |

**Table 49: Backup Recovery Configuration Parameters *(Continued)***

| Parameter | Description |
|---|---|
| Select the managed hosts you would like to run data backups: | This option is only displayed if you select the **Configuration and Data Backups** option.<br>All hosts in your deployment are listed. The first host in the list is your Console; it is enabled for data backup by default, therefore no check box is displayed. If you have managed hosts in your deployment, the managed hosts are listed below the Console and each managed host includes a check box.<br>Select the check box for the managed hosts you want to run data backups on.<br>For each host (Console or managed hosts), you can optionally clear the data items you want to exclude from the backup archive. |
| Configuration Only Backup | |
| Backup Time Limit (min) | Type or select the length of time, in minutes, that you want to allow the backup to run. The default is 180 minutes. If the backup process exceeds the configured time limit, the backup process is automatically canceled. |
| Backup Priority | From this list box, select the level of importance that you want the system to place on the configuration backup process compared to other processes.<br>A priority of medium or high have a greater impact on system performance. |
| *Data Backup* | |
| Backup Time Limit (min) | Type or select the length of time, in minutes, that you want to allow the backup to run. The default is 1020 minutes. If the backup process exceeds the configured time limit, the backup is automatically canceled. |
| Backup Priority | From the list, select the level of importance you want the system to place on the data backup process compared to other processes.<br>A priority of medium or high have a greater impact on system performance. |

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.
3. On the toolbar, click **Configure**.
4. On the **Backup Recovery Configuration** window, customize your nightly backup.
5. Click **Save**.
6. Close the **Backup Archives** window.
7. On the **Admin** tab menu, click **Deploy Changes**.

## Creating an On-demand Configuration Backup Archive

If you must back up your configuration files at a time other than your nightly scheduled backup, you can create an on-demand backup archive. On-demand backup archives include only configuration information.

You initiate an on-demand backup archive during a period when JSA has low processing load, such as after normal office hours. During the backup process, system performance is affected.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Backup and Recovery**.
3. From the toolbar, click **On Demand Backup**.
4. Enter values for the following parameters:

| Option | Description |
|---|---|
| Name | Type a unique name that you want to assign to this backup archive. The name can be up to 100 alphanumeric characters in length. The name can contain following characters: underscore (_), dash (-), or period (.). |

*(Continued)*

| Option | Description |
|---|---|
| Description | Type a description for this configuration backup archive. The description can be up to 255 characters in length. |

5. Click **Run Backup**.

   You can start a new backup or restore processes only after the on-demand backup is complete. You can monitor the backup archive process in the **Backup Archives** window.

## Creating an Email Notification for a Failed Backup

To receive a notification by email about a backup failure on the JSA Console or a JSA Event Processor, create a rule that is based on the system notification message.

You must configure an email server to distribute system notifications in JSA. For more information, see "Configuring Your Local Firewall" on page 109.

If a backup fails, you see one of the following backup failure system notifications:

- `Backup: requires more disk space`
- `Backup: last Backup exceeded execution threshold`
- `Backup: unable to execute request`

1. Click the **Offenses** tab.
2. In the **Offenses** pane, click **Rules**.
3. Click **Actions** >**New Event Rule**.
4. In the **Rule Wizard**, check the **Skip this page when running this rules wizard** box and click **Next**.
5. In the filter box, type the following search query:

   **when the event QID is one of the following QIDs**
6. Click the green add **(+)** icon.
7. In the **Rule** pane, click the **QIDs** link.

8. In the **QID/Name** field, type **Backup:**

9. Select the following QIDs and click **Add +**:

   - **Backup requires more disk space**
   - **Backup: last backup exceeded execution threshold**
   - **Backup unable to execute request**

10. Click **Submit**.

11. In the **Rule** pane, type the following name for your rule test and click **Next**:

    **Backup Failure**

12. In the **Rule Response** section, check the **Email** box and type the email addresses you want to notify.

RELATED DOCUMENTATION

Manage Existing Backup Archives | **299**

Restore JSA Configurations and Data | **301**

Backup and Restore Applications | **314**

# Manage Existing Backup Archives

IN THIS SECTION

Use the **Backup Management Archive** window to view and manage all successful backup archives.

## Importing a Backup Archive

Importing a backup archive is useful if you want to restore a backup archive that was created on another JSA host.

If you place a JSA backup archive file in the **/store/backupHost/inbound** directory on the Console server, the backup archive file is automatically imported.

1. On the navigation menu (☰), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.

3. In the **Upload Archive** field, click **Browse**.

4. Locate and select the archive file that you want to upload. The archive file must include a **.tgz** extension.

5. Click **Open**.

6. Click **Upload**.

## Deleting a Backup Archive

To delete a backup archive file, the backup archive file and the Host Context component must be located on the same system. The system must also be in communication with the JSA Console and no other backup can be in progress.

If a backup file is deleted, it is removed from the disk and from the database. Also, the entry is removed from this list and an audit event is generated to indicate the removal.

1. On the navigation menu (☰), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.

3. In the **Existing Backups** section, select the archive that you want to delete.

4. Click **Delete**.

RELATED DOCUMENTATION

# Restore JSA Configurations and Data

IN THIS SECTION

Restoring a backup archive is useful if you want to restore previously archived configuration files, offense data, and asset data on your JSA system.

Before you restore a backup archive, note the following considerations:

- You can restore only a backup archive that is created within the same release of software and its software update level. For example, if you are running JSA 7.5.0 update package, make sure that , the backup archive is cretaed on the JSA 7.5.0 update package Console.

- The restore process restores only your configuration information, offense data, and asset data. For more information, see "Restoring Data".

- If the backup archive originated on a NATed Console system, you can restore only that backup archive on a NATed system.

- You cannot complete a configuration restore on a console in which the IP address matches the IP address of a managed host in the backup.

If possible, before you restore a configuration backup, run an on demand backup to preserve the current environment. The following description is a high-level view of the configuration restore process:

- Tomcat is shut down
- All system processes are shut down.
- All files are extracted from the backup archive and restored to disk.
- Database tables are restored.
- All system processes are restored.
- Tomcat is restarted.

For more information about how to backup or restore an archive, see the following topics:

- Restoring a Backup Archive
- Restoring a Backup Archive Created on a Different JSA System
- Restoring Data
- Verifying Restored Data
- Retrieving Backup Files Missing from the Disk
- WinCollect Files are not Restored During a Configuration Restore

**NOTE**: If you are restoring WinCollect data, you must install the WinCollect SFS that matches the version of WinCollect in your backup before you restore the configuration. For more information, see WinCollect Files are not Restored During a Configuration Restore" .

## Restoring a Backup Archive

You can restore a backup archive. Restoring a backup archive is useful if you have a system hardware failure or you want to restore a backup archive on a replacement appliance.

You can restart the Console only after the restore process is complete.

The restore process can take up to several hours; the process time depends on the size of the backup archive that must be restored. When complete, a confirmation message is displayed.

A window provides the status of the restore process. This window provides any errors for each host and instructions for resolving the errors.

The following parameters are available in the **Restore a Backup** window:

**Table 50: Restore a Backup Parameters**

| Parameter | Description |
|---|---|
| **Name** | The name of the backup archive. |
| **Description** | The description, if any, of the backup archive. |
| **Type** | The type of backup. Only configuration backups can be restored, therefore, this parameter displays **config**. |
| **Select All Configuration Items** | When selected, this option indicates that all configuration items are included in the restoration of the backup archive. |
| **Restore Configuration** | Lists the configuration items to include in the restoration of the backup archive. To remove items, you can clear the check boxes for each item you want to remove or clear the **Select All Configuration Items** check box. |
| **Select All Data Items** | When selected, this option indicates that all data items are included in the restoration of the backup archive. |
| **Restore Data** | Lists the configuration items to include in the restoration of the backup archive. All items are cleared by default. To restore data items, you can select the check boxes for each item you want to restore. |

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.

3. Select the archive that you want to restore.

4. Click **Restore**.

5. On the **Restore a Backup** window, configure the parameters.

   Select the **Custom Rules Configuration** check box to restore the rules and reference data that is used by apps. Select the **Users Configuratio**n check box to restore authorized tokens that are used by apps.

The following table lists the restore configurations and what is included in each:

**NOTE**: The content included in each configuration is not limited to the content that is listed.

| Restore Configuration | Content Included |
|---|---|
| **Custom Rules Configuration** | • Rules<br>• Reference Sets<br>• Reference Data<br>• Saved Searches<br>• Forwarding Destinations<br>• Routing Rules<br>• Custom Properties<br>• Historical Searches<br>• Historical Rules<br>• Retention Bucket Configuration |
| **Deployment Configuration** | All content.<br>If you select this option, it is recommended that you select all other configuration options. |

*(Continued)*

| Restore Configuration | Content Included |
|---|---|
| **Users Configuration** | • Users<br>• User Roles<br>• Security Profiles<br>• Authorized Services<br>• Dashboards<br>• User Settings<br>• User Quick Searches |
| **License** | • License keys<br>• License Pool Allocations<br>• License history |
| **Report Templates** | Report templates<br>This does not include generated report content. |
| **System Settings** | • System Settings<br>• Asset Profiler Configuration |
| **QVM Scan profiles and results** | QVM Scan profiles and results |

*(Continued)*

| Restore Configuration | Content Included |
|---|---|
| **Installed Applications Configuration** | App configurations<br>This does not include app data.<br>Apps depending on authorized services might not work as expected if **Users Configuration** is not selected.<br>When **Installed Applications Configuration** is selected, the **Deployment Configuration group** is auto-selected. |
| **Assets** | Asset model<br>When **Assets** is selected, the **Deployment Configuration group** is auto-selected. |
| **Offenses** | • Offense data<br>• Offense associations (for example, QID links, rule links, or asset links)<br>• Offense searches<br>**NOTE**: When Offenses is selected, the Deployment Configuration group is auto-selected. |

6. Click **Restore**.

7. Click **OK**.

8. Click **OK**.

9. Choose one of the following options:

   - If the user interface was closed during the restore process, open a web browser and log in to JSA.

   - If the user interface was not closed, the login window is displayed. Log in to JSA.

10. Follow the instructions on the status window.

After you verify that your data is restored to your system, ensure that your DSMs, vulnerability assessment (VA) scanners, and log source protocols are also restored.

If the backup archive originated on an HA cluster, you must click **Deploy Changes** to restore the HA cluster configuration after the restore is complete. If disk replication is enabled, the secondary host immediately synchronizes data after the system is restored. If the secondary host was removed from the deployment after a backup, the secondary host displays a failed status on the **System and License Management** window.

## Restoring a Backup Archive Created on a Different JSA System

Each backup archive includes the IP address information of the system where it was created. When you restore a backup archive from a different JSA system, the IP address of the backup archive and the system that you are restoring are mismatched. You can correct the mismatched IP addresses.

You can restart the Console only after the restore process is complete. The restore process can take up to several hours; the process time depends on the size of the backup archive that must be restored. When complete, a confirmation message is displayed.

A window provides the status of the restore process, and provides any errors for each host and instructions for resolving the errors.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.

3. Select the archive that you want to restore, and click **Restore**.

4. On the Restore a Backup window, configure the following parameters and then click **Restore**.

   **Table 51: Restore a Backup Parameters**

   | Parameter | Description |
   |---|---|
   | **Select All Configuration Items** | Indicates that all configuration items are included in the restoration of the backup archive. This check box is selected by default. |
   | **Restore Configuration** | Lists the configuration items to include in the restoration of the backup archive. All items are selected by default. |

Table 51: Restore a Backup Parameters *(Continued)*

| Parameter | Description |
|---|---|
| **Select All Data Items** | Indicates that all data items are included in the restoration of the backup archive. This check box is selected by default. |
| **Restore Data** | Lists the configuration items to include in the restoration of the backup archive. All items are cleared by default. |

5. Stop the table service on each managed host in your deployment. The IP tables is a Linux based firewall.

    a. Using SSH, log in to the managed host as the root user.

    b. For App Host, type the following commands:

    - `systemctl stop docker_iptables_monitor.timer`
    - `systemctl stop iptables`

    c. For all other managed hosts, type the following command:

    ```
    service iptables stop
    ```

    d. Repeat for all managed hosts in your deployment.

6. On the **Restore a Backup** window, click **Test Hosts Access**.

7. After testing is complete for all managed hosts, verify that the status in the **Access Status** column indicates a status of **OK**.

8. If the **Access Status** column indicates a status of **No Access** for a host, stop iptables again, and then click **Test Host Access** again to attempt a connection.

9. On the **Restore a Backup** window, configure the parameters.

    > **NOTE**: By selecting the Installed Applications Configuration checkbox, you restore the install app configurations only. Extension configurations are not restored. Select the Deployment Configuration checkbox if you want to restore extension configurations.

10. Click **Restore**.

11. Click **OK**.

12. Click **OK** to log in.

13. Choose one of the following options:

    - If the user interface was closed during the user restore process, open a web browser and log in to JSA.

    - If the interface was not closed, the login window is displayed. Log in to JSA.

14. View the results of the restore process and follow the instructions to resolve any errors.

15. Refresh your web browser window.

16. From the **Admin** tab, select **Advanced >Deploy Full Configuration**.

    **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

17. To enable the IP tables for an App Host, type the following command:

    ```
    systemctl start docker_iptables_monitor.timer
    ```

After you verify that your data is restored to your system, you must reapply RPMs for any DSMs, vulnerability assessment (VA) scanners, or log source protocols.

If the backup archive originated on an HA cluster, you must click **Deploy Changes** to restore the HA cluster configuration after the restore is complete. If disk replication is enabled, the secondary host immediately synchronizes data after the system is restored. If the secondary host was removed from the deployment after a backup, the secondary host displays a failed status on the **System and License Management** window.

## Restoring Data

You can restore the data on your JSA Console and managed hosts from backup files. The data portion of the backup files includes information such as source and destination IP address information, asset data, event category information, vulnerability data, flow data, and event data.

Each managed host in your deployment, including the JSA console, creates all backup files in the **/store/backup/** directory. Your system might include a **/store/backup** mount from an external SAN or NAS service. External services provide long term, offline retention of data, which is commonly required for compliancy regulations, such as PCI.

**NOTE**: If you are restoring data on a new JSA Console, the configuration backup must be restored before you restore the data backup.

Ensure that the following conditions are met:

- You know the location of the managed host where the data is backed up.
- If your deployment includes a separate mount point for that volume, the **/store** or **/store/ariel** directory has sufficient space for the data that you want to recover.
- You know the date and time for the data that you want to recover.
- If your configuration has been changed, before you restore the data backup, you must restore the configuration backup.

1. Use SSH to log in to JSA as the root user.
2. Go to the **/store/backup** directory.
3. To list the backup files, type the following command:

   ```
   ls -l
   ```

4. If backup files are listed, go to the root directory by typing the following command:

   ```
   cd /
   ```

   **NOTE**: The restored files must be in the **/store** directory. If you type **cd** instead of **cd /**, the files are restored to the **/root/store** directory.

5. To extract the backup files to their original directory, type the following command:

   ```
   tar -zxpvPf /store/backup/backup.name.hostname_hostID .target date.backup type.timestamp.tgz
   ```

   **Table 52: Description Of File Name Variables**

   | File name variable | Description |
   |---|---|
   | *name* | The name of the backup. |

Table 52: Description Of File Name Variables *(Continued)*

| File name variable | Description |
|---|---|
| *hostname_hostID* | The name of the JSA system that hosts the backup file followed by the identifier for the JSA system. |
| *target date* | The date that the backup file was created. The format of the target date is *day_month_year*. |
| *backup type* | The options are **data** or **config**. |
| *timestamp* | The time that the backup file was created. |

Daily backup of data captures all data on each host. If you want to restore data on a managed host that contains only event or flow data, only that data is restored to that host. If you want to maintain the restored data, increase your data retention settings to prevent the nightly disk maintenance routines from deleting your restored data.

## Verifying Restored Data

Verify that your data is restored correctly in JSA.

1. To verify that the files are restored, review the contents of one of the restored directories by typing the following command:

   **cd /store/ariel/flows/payloads/<*yyyy/mm/dd*>**

   **cd /store/ariel/events/payloads/<*yyyy/mm/dd*>**

   You can view the restored directories that are created for each hour of the day. If directories are missing, data might not be captured for that time period.

2. Verify that the restored data is available.

   a. Log in to the JSA interface.

   b. Click the **Log Activity** or **Network Activity** tab.

   c. Select **Edit Search** from the **Search** list on the toolbar.

d. In the **Time Range** pane of the **Search** window, select **Specific Interval**.

e. Select the time range of the data you restored and then click **Filter**.

f. View the results to verify the restored data.

g. If your restored data is not available in the JSA interface, verify that data is restored in the correct location and file permissions are correctly configured.

   Restored files must be in the **/store directory**. If you typed **cd** instead of **cd /** when you extracted the restored files, check the **/root/store** directory for the restored files. If you did not change directories before you extracted the restored files, check the **/store/backup/store** directory for the restored files.

   Typically, files are restored with the original permissions. However, if the files are not owned by the root user account, issues might occur. The correct ownership of directories and files in **/store/ariel/events/payloads** and **/store/ariel/flows/payloads** is `root:root`. If the files and folders do not have the correct ownership, change the ownership by using the **chown** command.

   The correct permissions of directories and files in **/store/ariel/events/payloads** and **/ store/ariel/ flows/payloads** is 755 for folders, and 644 for files. If the files and folders do not have the correct permissions, change the permissions by using the **chmod** command.

After you verified that your data is restored, you must complete an auto update in JSA. The auto update ensures DSMs, vulnerability assessment (VA) scanners, and log source protocols are at the latest version.

## Retrieving Backup Files Missing from the Disk

When the backup files are missing from the disk, the respective backup table entry on the Backup and Recovery page is marked with an exclamation icon to show that the file is not retrievable. Files that are missing cannot be downloaded or restored. This issue can occur when you are using external storage that is no longer available, or is offline.

1. On the **Admin** tab, click **Backup and Recovery**.

2. If the external storage is offline or no longer available, **delete** the table entry by using the Delete option at the top of the **Backup and Recovery** page.

> **NOTE**: If you are not expecting this behavior and are using external storage for your backup archive location, investigate whether the storage system is still accessible. If it is offline, and you are able to restore the directory, the indicator icons are automatically updated and removed when the system detects the restored files.

3. On the **Backup and Recovery** page, click **Configure** and take note of the **Backup Repository Path**.

4. Log out of JSA and log back in to ensure that the files are again accessible by fixing the external mount or restoring missing files to the appropriate backup location.

5. Refresh the **Backup and Recovery** page to synchronize the backups.

## WinCollect Files are not Restored During a Configuration Restore

When you complete a configuration restore and some WinCollect files are not restored, it might be because the installation ISO contains a previous version of WinCollect.

The JSA ISO contains a built-in version of WinCollect. When you restore by using that ISO, it deploys the WinCollect files that are stored in that ISO, rather than the files from your backup.

To remedy this issue, you must install the WinCollect SFS that matches the version of WinCollect in your backup before you restore the configuration. Perform the following tasks in this order:

- Perform JSA backup.
- Bring new hardware online and deploy the ISO.
- Install the WinCollect SFS that matches the version of WinCollect in your backup on the Console.
- Restore the configuration backup.

The appropriate WinCollect files are deployed with the configuration restore.

RELATED DOCUMENTATION

Backup and Restore Applications | 314
Backup JSA Configurations and Data | 294
Manage Existing Backup Archives | 299

# Backup and Restore Applications

IN THIS SECTION

JSA provides a way to backup and restore application configurations separate from the application data.

Application configurations are backed up as part of the nightly configuration backup. The configuration backup includes apps that are installed on the JSA console and on an App Host. You can restore the application configuration by selecting the **Installed Applications Configuration** option when you restore a backup.

Application data is backed up separate from the application configuration by using an easy-to-use script that runs nightly. You can also use the script to restore the app data, and to configure backup times and data retention periods for app data.

## Backing Up and Restoring Apps

Use the JSA **Backup and Recovery** window on the **Admin** tab to back up and restore apps.

You can back up your apps by creating a configuration back up. A configuration backup does not backup your app's data.

If an App Host is attached to your JSA console, the App Host’s configuration is backed up as part of the console's Deployment Configuration. You cannot restore an App Host on a JSA console with a different IP address than the App Host was initially configured with.

By default, apps are restored to console unless an App Host is present. If JSA cannot restore apps to your App Host , it attempts to back restore them to the JSA console. The number of App Host apps that can be restored onto the console is constrained by the amount of memory that is available on the JSA console. Apps that are defined as **node_only** in their application manifest file cannot be restored to the JSA console.

1. On the navigation menu
   (

≡
), click **Admin**.

2. In the **System Configuration** section, click **Backup and Recovery**.

3. Select an existing backup in the **Backup and Recovery** window and click **Restore**.

4. Ensure that the **Installed Applications Configuration** check box is selected, and click **Restore**.

> **NOTE**: By selecting the **Installed Applications Configuration** check box, you restore the install app configurations only. Extension configurations are not restored. Select the **Deployment Configuration** check box if you want to restore extension configurations.

## Backing Up and Restoring App Data

Use the **marathon-volume-backup.py** script to back up and recover app data.

A configuration backup that you do on the **backup and Recovery** window does not back up your apps' data. The **/usr/local/bin/marathon-volume-backup.py** script runs nightly at 2:30 AM, and backs up each installed application's **/store** mounted volume. By default, data is retained for 7 days.

Use the script to do the following tasks:

- Back up data manually for installed apps.

- List all installed app data backups on the system.

- Restore data for installed apps.

- Run the retention process and set the retention period for backups.

This script is on both the JSA console and App Host if one is installed. The script backs up app data only if apps are on the current host.

1. Use SSH to log in to your Console or your App Host as the root user.

2. Go to the **/usr/local/bin/** directory.

    - Use the following command to back up app data:

      The **marathon-volume-backup.py** script runs nightly at 2:30 AM local time to back up all installed apps. Backup archives are stored in the**/store/backup/marathon** folder. You can change the backup archives location by editing the APP_VOLUME_BACKUP_DIR variable in **/store/**

**configservices/staging/globalconfig/ nva.conf.**. You must deploy changes after you edit this variable.

- To view all data backups for installed apps, enter the following command:

  **./marathon-volume-backup.py ls**

  This command outputs all backup archives that are stored in the backup archives folder.

- To restore a backup archive, enter the following command:

  **./marathon-volume-backup.py restore -i <backup name> - <backup_name>**

  Use the **ls** command to find the name of a backup archive

- To restore data for a specific application instance, rather than restoring all instances, enter

  the following command:

  ```
  ./app-volume-backup.py restore-interactive -i <backup name>
  ```

  **NOTE**: This function was added in JSA 7.5.0, and works only with backups that were created after updating to version 7.5.0.

- By default, all backup archives are retained for one week. The retention process runs nightly at 2:30 AM local time with the backup.

  - To perform retention manually, and use the default retention period, enter the following command:

    **./marathon-volume-backup.py retention**

  - You can also set the retention period manually by adding **-t** (time - defaults to 1) and **-p** (period - defaults to 0) switches.

    The **-p** switch accepts three values: 0 for a week, 1 for a day, and 2 for an hour.

    For example, to set the retention period for a back up to 3 weeks, enter the following command:

    **./marathon-volume-backup.py retention -t 3 -p 0**

- If you want to change the retention time that is used by the nightly timer, add flags to the retention command found in the following systemd service file:

  **/etc/systemd/system/framework-apps-data-backup.service**

For example, to change the retention period that is used by the nightly retention process to 5 days, locate the following line:

```
ExecStart=/opt/qradar/bin/app-volume-backup.py retention
```

Replace it with:

**ExecStart=/usr/local/bin/marathon-volume-backup.py retention -t 5 -p 1**

Save your changes, and run the **systemctl daemon-reload** command for systemd to apply the changes.

App containers are restarted automatically after the restore is complete.

### RELATED DOCUMENTATION

Backup JSA Configurations and Data | 294
Manage Existing Backup Archives | 299
Restore JSA Configurations and Data | 301

# Data Redundancy and Recovery in JSA Deployments

IN THIS SECTION

To safeguard from data loss, configure your deployments to include data redundancy and recovery functionality. Data Synchronization is possible when you have two identical JSA systems in separate geographic environments that are a mirror of each other, and data is synchronized at both sites. Forwarding data uses off-site forwarding, which is set up on both the primary and secondary deployments. You can set up data synchronization with deployments that are in different geographical locations.

### Data Synchronization App

Implement the Data Synchronization app to safeguard your JSA configurations and data by mirroring your data to another identical JSA. Recovery from a data loss is possible when you have two identical JSA systems in separate geographic environments that are a mirror of each other, and data is collected at both sites.

If you do not meet the requirements for the Data Synchronization app, the following are some alternative solutions. Recovery from data loss is possible when you forward live data, for example, flows and events from a primary JSA, to a parallel system at another site.

### Primary JSA Console and backup console

A hardware failure solution, where the backup console is a copy of the primary server, with the same configuration but stays powered off. Only one console is operational at any one time. If the primary console fails, you manually turn the power on the backup console, apply the primary configuration backup, and use the IP address from the primary console. After you restore the primary server and before you turn it on, you manually turn off the backup server. If the system is down for a long time, apply the backup console configuration backup to the primary server.

### Event and flow forwarding

Events and flows are forwarded from a primary site to a secondary site. Identical architectures in two separate data centers are required.

### Distributing the same events and flows to the primary and secondary sites

Distribute the same event and flow data to two live sites by using a load balancer or other method to deliver the same data to mirrored appliances. Each site has a record of the log data that is sent.

## Primary JSA Console and backup JSA Console

When the primary JSA Console fails and you want the backup JSA Console to take up the role of the primary, you manually turn the power on the backup console, apply the configuration backup and the IP address from the primary. Use a similar switchover method for other appliances such as a JSA Flow Processor or an Event Collector, where each appliance has a cold backup or spare that is an identical appliance.

The backup console takes over the primary JSA Console role from the time of activation, and does not store past events, flow, or offenses from the original primary JSA Console. Use this type of deployment for your appliances, to minimize downtime, when there is a hardware failure.

- A backup console requires its own dedicated license key (matching the EPS and FPM values of the primary console).

- The license configuration of the backup console needs to match the values of the primary JSA Console; this includes the EPS and FPS values of the primary JSA Console.

  **Example:** If the primary JSA Event Processor was licensed for 15K EPS, the redundant backup console should also be licensed for 15K EPS.

- There are special failover upgrade parts that need to be purchased for the backup console.

- From a technical perspective, the license for both primary and backup consoles are identical, however for compliance reasons the backup console (and associated license) cannot not be processing live data unless a failure has occurred with the primary JSA Console.

- Data collected by the backup console will need to be copied back to the Primary console when the Primary console once again becomes functional.

If the primary fails, take the following steps to set up the backup console as the primary JSA Console:

1. Power on the backup console.

2. Add the IP address from the primary console.

3. Restore configuration backup data from the primary console to the backup console.

The backup console functions as the primary console until the primary console is brought back online. Ensure that both servers are not online at the same time.

### Configuring the IP Address on the Backup Console

When the primary JSA Console fails, you configure the secondary backup console to take on the primary console role. Add the IP address of the failed JSA Console to the backup console so that your JSA system continues to function.

Power on the backup console.

1. Use SSH to log in to as the root user.

2. To configure the IP address on the backup console, follow these steps:

   a. Type the following command:

      ```
      qchange_netsetup
      ```

      **NOTE**: Verify all external storage which is not **/store/ariel or /store** is not mounted.

   b. Follow the instructions in the wizard to enter the configuration parameters.

      After the requested changes are processed, the JSA automatically shuts down and restarts.

### Backup and Recovery

Back up your JSA configuration information and data so that you can recover from a system failure or data loss.

Use the backup and recovery that is built-in to JSA to back up your data. However, you must restore the data manually. By default, JSA creates a daily backup archive of your configuration information at midnight. The backup archive includes configuration information, generated data, or both from the previous day.

You can create the following types of backup:

- Configuration backups, which include system configuration data, for example, assets and log sources in your JSA deployment.
- Data backups, which include information that is generated by a working JSA deployment such as log information or event dates.

## Event and Flow Forwarding from a Primary Data Center to Another Data Center

To ensure that there is a redundant data store for events, flows, offenses, and that there is an identical architecture in two separate data centers, forward event and flow data from site 1 to site 2.

The following information is provided only for general guidance and is not intended or designed as a how-to guide.

This scenario is dependent upon site 1 remaining active. If site 1 fails, data is not transmitted to Site 2, but the data is current up to the time of failure. In the case of failure at site 1, you implement recovery of your data, by manually changing IP addresses and use a backup and restore to fail over from site 1 to site 2, and to switch to site 2 for all JSA hosts.

The following list describes the setup for event and flow forwarding from the primary site to the secondary site:

- There is an identical distributed architecture in two separate data centers, which includes a primary data center and a secondary data center.
- The primary JSA Console is active and collecting all events and flows from log sources and is generating correlated offenses.
- You configure off-site targets on the primary JSA Console to enable forwarding of event and flow data from the primary data center to the event and flow processors in another data center.

**Fast path:** Use routing rules instead of off-site targets because the setup is easier.

- Periodically, use the content management tool to update content from the primary JSA Console to the secondary JSA Console.

In the case of a failure at site 1, you can use a high-availability (HA) deployment to trigger an automatic failover to site 2. The secondary HA host on site 2 takes over the role of the primary HA host on site 1. Site 2 continues to collect, store, and process event and flow data. Secondary HA hosts that are in a standby state don't have services that are running but data is synchronized if disk replication is enabled.

**NOTE**: You can use a load balancer to divide events, and split flows such as NetFlow, J-Flow, and sFlow but you can't use a load balancer to split Flow Processors. Use external technologies such as a regenerative tap to divide Flow Processor and send to the backup site.

The following diagram shows how site 2 is used as a redundant data store for site 1. Event and flow data are forwarded from site 1 to site 2.

**Figure 10: Event and Flow Forwarding from Site 1 to Site 2 for Disaster Recovery**

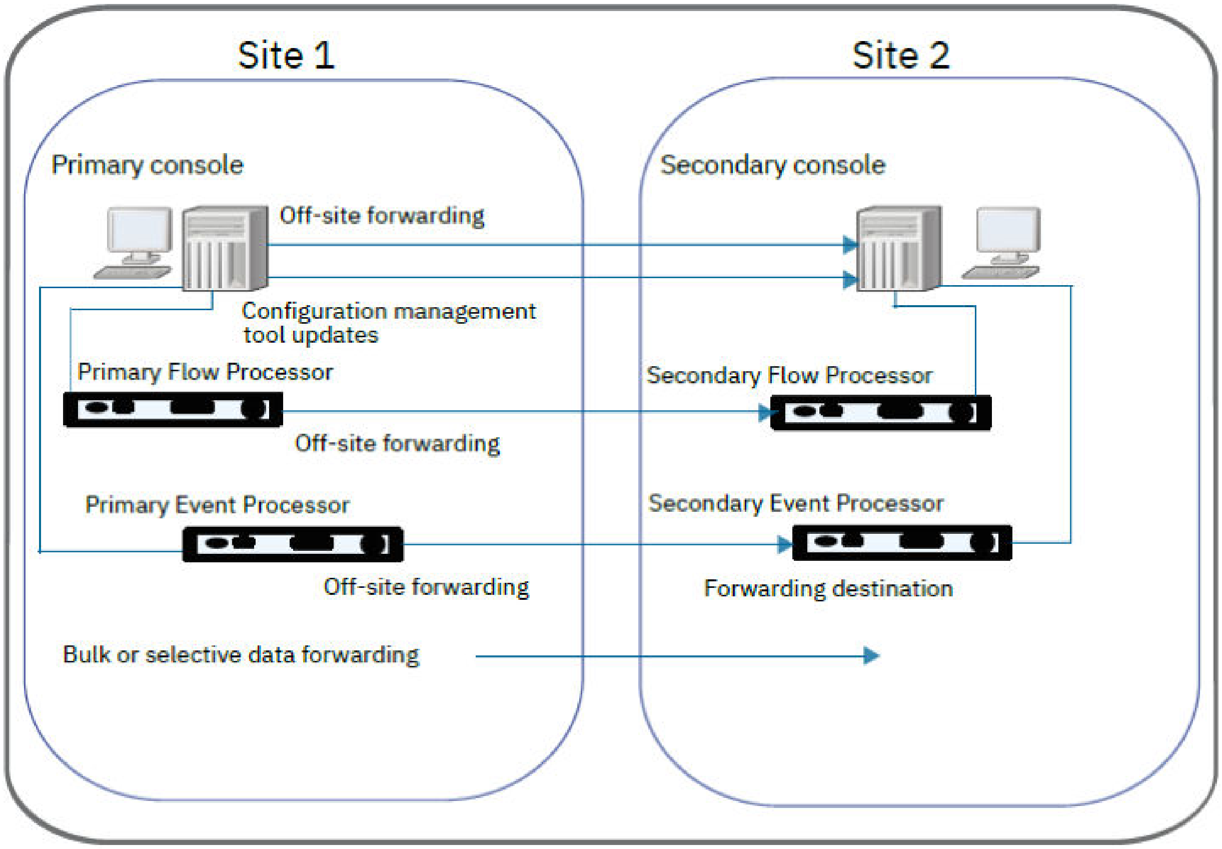

### Event and Flow Forwarding Configuration

For data redundancy, configure JSA to forward data from one site to a backup site.

The target system that receives the data from JSA is known as a forwarding destination. JSA ensure that all forwarded data is unaltered. Newer releases of JSA can receive data from earlier releases of JSA. However, earlier releases cannot receive data from later releases. To avoid compatibility issues, upgrade all receivers before you upgrade JSA systems that send data. Follow these steps to set up forwarding:

1. Configure one or more forwarding destinations.

   A forwarding destination is the target system that receives the event and flow data from the JSA primary console. You must add forwarding destinations before you can configure bulk or selective data forwarding.

2. Configure routing rules, custom rules, or both.

   After you add one or more forwarding destinations for your event and flow data, you can create filterbased routing rules to forward large quantities of data.

3. Configure data exports, imports, and updates.

   You use the content management tool to move data from your primary JSA Console to the JSA secondary console. Export security and configuration content from JSA into an external, portable format.

## Load Balancing of Events and Flows Between Two Sites

When you are running two live JSA deployments at both a primary and secondary site, you send event and flow data to both sites. Each site has a record of the log data that is sent. Use the content management tool to keep the data synchronized between the deployments

The following diagram shows two live sites, where data from each site is replicated to the other site.

**Figure 11: Load Balancing of Events and Flows Between Two Sites**

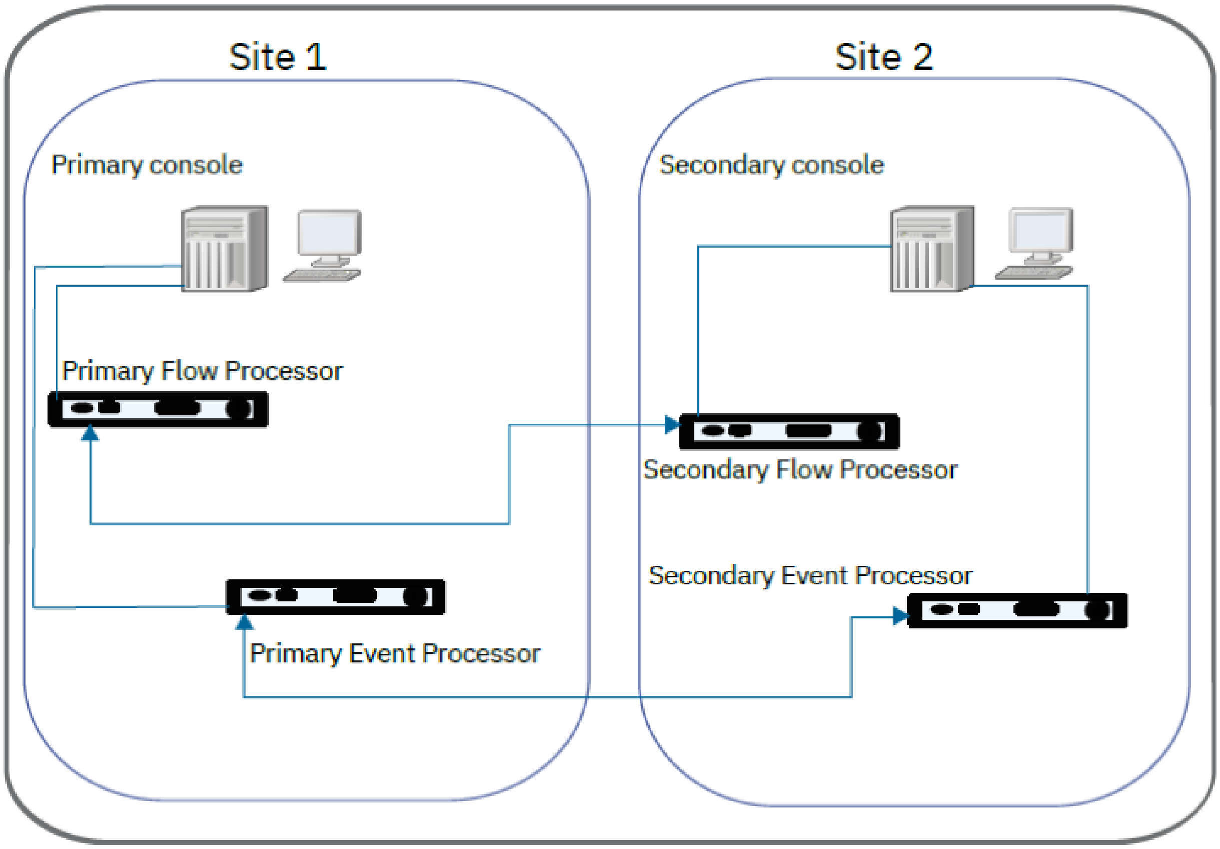

Send the same events and flows to separate data centers or geographically separate sites and enable data redundancy by using a load balancer or other method to deliver the same data to mirrored appliances.

## Restoring Configuration Data from the Primary to the Secondary JSA Console

After you set up the secondary JSA Console as the destination for the logs, you either add or import a backup archive from the primary JSA Console. You can restore a backup archive that is created on another JSA host. Log in to the secondary JSA Console and do a full restore of the primary console backup archive to the secondary JSA Console.

You must have a data backup from your primary console to complete this task.

1. On the navigation menu, click **Admin**.

2. On the navigation menu, click **System Configuration**.

3. Click the **Backup and Recovery** icon.
4. In the **Upload Archive** field, click **Browse**.
5. Locate and select the archive file that you want to upload.

   > **TIP**: If the JSA backup archive file is in the **/store/backupHost/inbound** directory on the console server, the backup archive file is automatically imported.

   The archive file must have a .tgz extension.
6. Click **Open**.
7. Click **Upload**.
8. Select the archive that you uploaded and click **Restore**.

   When the restore is finished, the secondary JSA Console becomes the primary console.

## Event and Flow Data Redundancy

Send the same events and flows to separate data centers or geographically separate sites and enable data redundancy by using a load balancer or other method to deliver the same data to mirrored appliances.

Configure the distribution of log and flow sources for data redundancy:

- Send log source data to the Event Processor on the second site.
- Send flow source data to the Flow Processor on the second site.

For more information about configuring log sources, see the *Juniper Secure Analytics Configuring DSMs Guide*.

Figure 12: Sending Events and Flows to Two Sites

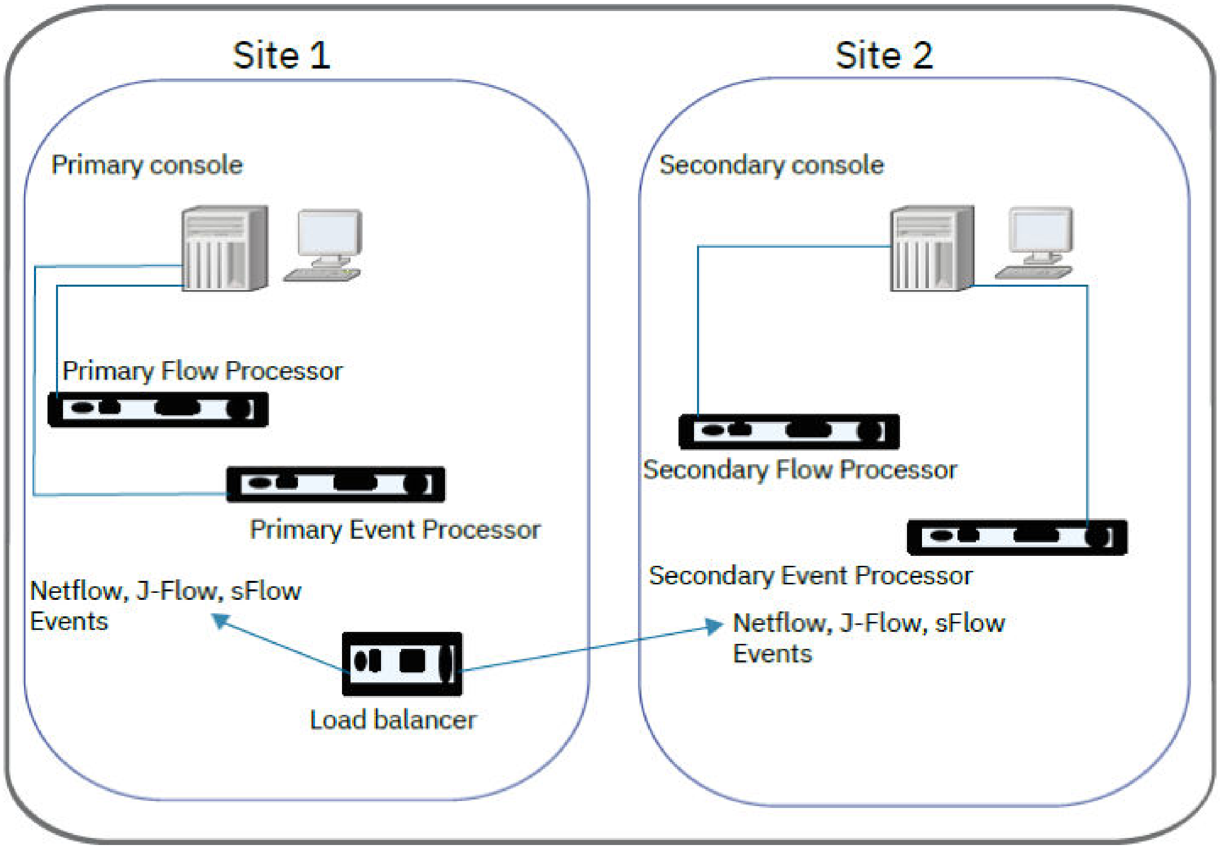

## Configure JSA to receive events

JSA automatically discovers many log sources that send syslog messages in your deployment. Log sources that are automatically discovered by JSA appear in the Log Sources window.

You configure the automatic discovery of log sources for each Event Collector by using the **Autodetection Enabled** setting in the Event Collector configuration. If you want to keep the log source event IDs synchronized with the primary Event Collector, you disable the **Autodetection** setting. In this situation, use the content management tool to synchronize the log source configuration or restore a configuration backup to the site.

For more information about auto discovered log sources and configurations specific to your device or appliance, see the *Juniper Secure Analytics Configuring DSMs Guide*.

## Configure JSA to Receive Flows

To enable data redundancy for flows, you need to send NetFlow, J-Flow, and sFlow to both sites for Flow Processor collection.

You can collect flows from a SPAN or tap and then send packets to your backup location, or you mirror the SPAN or tap in the backup location by using external technologies. A load balancer splits flows such as NetFlow, J-Flow, and sFlow but it can't split Flow Processor.

### Use the Content Management Tool (CMT)

If you want to ensure that the primary JSA Console from site 1 and the secondary JSA Console from site 2 have identical configurations, use the content management tool to update site 2 with the configurations from site 1.

RELATED DOCUMENTATION

Backup JSA Configurations and Data | 294

Manage Existing Backup Archives | 299

Restore JSA Configurations and Data | 301

# Backup and Restore the QRadar Analyst Workflow

If you need to restore QRadar Analyst Workflow to a different JSA console, you must reinstall QRadar Analyst Workflow after the JSA restore.

QRadar Analyst Workflow is located at https://exchange.xforce.ibmcloud.com/hub, where you can download the *QRadarAnalystWorkflow<x.x.x>.* zip file.

The JSA backup and recovery feature backs up and restores all the data for QRadar Analyst Workflow, and it can be restored to a different host. However, if you restore to a different host, QRadar Analyst Workflow docker images are not included in the recovery.

After you restore JSA, you must copy the *QRadarAnalystWorkflow<x.x.x>*..zip file to the new host, decompress the file, and install on the new host.

For instructions on installing the *QRadarAnalystWorkflow<x.x.x>*.zip file, see Installing the QRadar Analyst Workflow.

RELATED DOCUMENTATION

Backup JSA Configurations and Data | 294

Manage Existing Backup Archives | 299

Restore JSA Configurations and Data | 301

13

CHAPTER

# Flow Sources Management

# Flow Sources

For JSA appliances, JSA automatically adds default flow sources for the physical ports on the appliance. JSA also includes a default NetFlow flow source.

With JSA you can integrate flow sources.

RELATED DOCUMENTATION

# Types of Flow Sources

IN THIS SECTION

JSA Flow Processor can process flows from multiple sources, which are categorized as either internal or external sources.

## Internal Flow Sources

Sources that include packet data by connecting to a SPAN port or a network TAP are considered internal sources. These sources provide raw packet data to a monitoring port on the Flow Processor, which converts the packet details into flow records.

JSA does not keep the entire packet payload. Instead, it captures a snapshot of the flow, referred to as the payload or content capture, which includes packets from the beginning of the communication.

Flow collection from internal sources normally requires a dedicated Flow Processor.

## External Flow Sources

JSA supports the following external flow sources:

- NetFlow
- IPFIX
- sFlow
- J-Flow
- Packeteer
- Network interface

For more information about the fields that are supported for each flow source type, see the *Juniper Secure Analytics Users Guide*.

External sources do not require as much CPU utilization to process so you can send the flows directly to a Flow Processor. In this configuration, you may have a dedicated flow processor, receiving and creating flow data.

If your Flow Processor collects flows from multiple sources, you can assign each flow source a distinct name. A distinct name helps to distinguish the external flow data from other sources.

JSA can forward external flow source data by using the spoofing or non-spoofing method:

### Spoofing

Resends the inbound data that is received from a flow source to a secondary destination.

To configure the spoofing method, configure the flow source so that the **Monitoring Interface** is set to the management port on which the data is received.

When you use a specific interface, the Flow Processor uses a promiscuous mode capture to collect the flow data, rather than the default UDP listening port on port 2055. This way, the Flow Processor can capture and forward the data.

### Non-Spoofing

For the non-spoofing method, configure the **Monitoring Interface** parameter in the flow source configuration as Any.

The Flow Processor opens the listening port, which is the port that is configured as the **Monitoring Port**, to accept the flow data. The data is processed and forwarded to another flow source destination.

When the data is forwarded, the source IP address of the flow becomes the IP address of the JSA system, not the original router that sent the data.

RELATED DOCUMENTATION

Enabling and Disabling a Flow Source | **332**

Deleting a Flow Source | **333**

Flow Source Aliases Management | **333**

# Adding or Editing a Flow Source

Use the **Flow Source** window on the Admin tab to add or edit a flow source.

1. On the navigation menu (≡), click **Admin**.
2. In the **Data Sources** section, under **Flows**, click **Flow Sources**.
3. Do one of the following actions:
   - To add a flow source, click **Add**.
   - To edit a flow source, select the flow source and click **Edit**.
4. To create this flow source from an existing flow source, select the **Build from existing flow source** check box, and select a flow source from the **Use as Template** list.
5. Enter the name for the **Flow Source Name**.

   > **TIP**: If the external flow source is also a physical device, use the device name as the flow source name. If the flow source is not a physical device, use a recognizable name.

   For example, if you want to use IPFIX traffic, enter **ipf1**. If you want to use NetFlow traffic, enter **nf1**.
6. Select a flow source from the **Flow Source Type** list and configure the properties.
   - If you select the **Flowlog File** option, ensure that you configure the location of the Flowlog file for the **Source File Path** parameter.

- If you select the **JFlow**, **Netflow**, **Packeteer FDR**, or **sFlow** options in the **Flow Source Type** parameter, ensure that you configure an available port for the **Monitoring Port** parameter.

  The default port for the first NetFlow flow source that is configured in your network is 2055. For each additional NetFlow flow source, the default port number increments by 1. For example, the default NetFlow flow source for the second NetFlow flow source is 2056.

- If you select the **Network Interface** option, for the **Flow Interface**, configure only one log source for each Ethernet interface.

  > **NOTE**: You cannot send different flow types to the same port.

7. If traffic on your network is configured to take alternate paths for inbound and outbound traffic, select the **Enable Asymmetric Flows** check box.
8. Click **Save**.
9. On the **Admin** tab menu, click **Deploy Changes**.

RELATED DOCUMENTATION

Enabling and Disabling a Flow Source | 332

Deleting a Flow Source | 333

Flow Source Aliases Management | 333

# Enabling and Disabling a Flow Source

Using the **Flow Source** window, you can enable or disable a flow source.

1. On the navigation menu (≡), click **Admin**.
2. In the **Data Sources** section, under **Flows**, click **Flow Sources**.
3. Select the flow source that you want to enable or disable, and click **Enable/Disable**.
4. On the **Admin** tab, click **Deploy Changes**.

RELATED DOCUMENTATION

Deleting a Flow Source | 333

Flow Source Aliases Management | 333

Adding or Editing a Flow Source | 331

# Deleting a Flow Source

Use the **Flow Source** window to delete a flow source.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **Data Sources** section, under **Flows**, click **Flow Sources**.
3. Select the flow source that you want to delete, and click **Delete**.
4. On the **Admin** tab, click **Deploy Changes**.

RELATED DOCUMENTATION

Flow Source Aliases Management | 333

Enabling and Disabling a Flow Source | 332

Adding or Editing a Flow Source | 331

Correcting Flow Time Stamps | 335

# Flow Source Aliases Management

IN THIS SECTION

A flow source alias uses a virtual name to identify external flows that are sent to the same port on a flow processor. For example, the JSA Flow Processor can have a single NetFlow flow source that is listening on port 2055, and can have multiple NetFlow sources sending to the same JSA Flow Processor. By using flow source aliases, you can identify the different NetFlow sources based by their IP addresses.

When JSA Flow Processor receives traffic from a device that has an IP address but does not have a current alias, the JSA Flow Processor attempts a reverse DNS lookup. The lookup is used to determine the host name of the device.

You can configure the JSA Flow Processor to automatically create flow source aliases. When the JSA Flow Processor receives traffic from a device that has an IP address but does not have a current alias, it does a reverse DNS lookup to determine the host name of the device.

If the lookup is successful, the JSA Flow Processor adds this information to the database and reports the information to all JSA Flow Processor components in your deployment. If the lookup fails, JSA creates a default alias for the flow source based on the flow source name and the source IP address. For example, the default alias might appear as **default_NetFlow_172.16.10.139**.

## Adding or a Flow Source Alias

Use the **Flow Source Alias** window to add a flow source alias.

1. On the navigation menu (≡), click **Admin**.

2. In the **Data Sources** section, under **Flows**, click **Flow Source Aliases**.

3. Do one of the following actions:

   - To add a flow source alias, click **Add** and enter the values for the parameters.

   - To edit an existing flow source alias, select the flow source alias, click **Edit**, and update the parameters.

4. Click **Save**.

5. On the **Admin** tab, click **Deploy Changes**.

> **NOTE**: If you rename a flow source alias, you must use the original name to perform a historical search.

## Deleting a Flow Source Alias

Use the **Flow Source Alias** window to delete a flow source alias.

1. On the navigation menu (☰), click **Admin**.
2. In the **Data Sources** section, under **Flows**, click **Flow Source Aliases**.
3. Select the flow source alias that you want to delete, and then click **Delete**.
4. On the **Admin** tab, click **Deploy Changes**.

RELATED DOCUMENTATION

Enabling and Disabling a Flow Source | **332**

Deleting a Flow Source | **333**

Adding or Editing a Flow Source | **331**

Correcting Flow Time Stamps | **335**

# Correcting Flow Time Stamps

You can specify the way that you want flow time stamps to be handled when Netflow V9 begins sending records with overflowed system uptime values.

Two new configuration settings provide more control over the way that flow time stamps are handled when Netflow V9 begins sending records with overflowed system uptime values. The new settings eliminate the need to reset the first and last switched times.

The new configuration options and the default values are shown here:

- `NORMALISE_OVERFLOWED_UPTIMES=YES`
- `UPTIME_OVERFLOW_THRESHOLD_MSEC=86400000`

The time stamps are corrected when the system uptime value is less than the first and last switched packet times by more than the value that is specified in the `UPTIME_OVERFLOW_THRESHOLD_MSEC` configuration. The time stamps are corrected based on the assumption that the system uptime wrapped around the maximum 32-bit value.

1. To change these settings, add the settings to the **/store/configservices/staging/ globalconfig/ nva.conf** file.
2. To fine-tune the settings, specify a different time interval for the `UPTIME_OVERFLOW_THRESHOLD_MSEC` setting.
3. To disable this feature, set the `NORMALISE_OVERFLOWED_UPTIMES` to `NO`.

   When this feature is disabled, JSA does not modify the NetFlow v9 time stamps that meet this condition.
4. After you change the configuration settings, you must deploy the system.

## RELATED DOCUMENTATION

Flow Source Aliases Management | 333

Enabling and Disabling a Flow Source | 332

Adding or Editing a Flow Source | 331

14

CHAPTER

# Remote Networks and Services Configuration

# Remote Networks and Services Configuration

Use remote network and service groups to represent traffic activity on your network for a specific profile. Remote networks groups display user traffic that originates from named remote networks.

All remote network and service groups have group levels and leaf object levels. You can edit remote network and service groups by adding objects to existing groups or changing preexisting properties to suit your environment.

If you move an existing object to another group, the object name moves from the existing group to the newly selected group. However, when the configuration changes are deployed, the object data that is stored in the database is lost and the object ceases to function. To resolve this issue, create a new view and re-create the object that exists with another group.

You can group remote networks and services for use in the custom rules engine, flow, and event searches. You can also group networks and services in JSA Risk Manager, if it is available.

# Default Remote Network Groups

JSA includes default remote network groups.

The following table describes the default remote network groups.

**Table 53: Default Remote Network Groups**

| Group | Description |
|---|---|
| BOT | Specifies traffic that originates from BOT applications.<br>For more information, see Botnet Command and Control drop rules on the Emerging Threats website (http://rules.emergingthreats.net/blockrules/emerging-botcc.rules) |
| Bogon | Specifies traffic that originates from unassigned IP addresses.<br>For more information, see bogon reference on the Team CYMRU website (http://www.team-cymru.org/Services/Bogons/bogon-bn-nonagg.txt). |

**Table 53: Default Remote Network Groups *(Continued)***

| Group | Description |
|---|---|
| HostileNets | Specifies traffic that originates from known hostile networks.<br>HostileNets has a set of 20 (rank 1 - 20 inclusive) configurable CIDR ranges.<br>For more information, see HostileNets reference on the DShield website (http://www.dshield.org/ipsascii.html?limit=20) |
| Neighbours | Specifies traffic that originates from nearby networks that your organization has network peering agreements with.<br>This group is blank by default. You must configure this group to classify traffic that originates from neighboring networks. |
| Smurfs | Specifies traffic that originates from smurf attacks.<br>A smurf attack is a type of denial-of-service attack that floods a destination system with spoofed broadcast ping messages. |
| Superflows | This group is non-configurable.<br>A superflow is a flow that is an aggregate of a number of flows that have a similar predetermined set of elements. |
| TrustedNetworks | Specifies traffic from trusted networks, including business partners that have remote access to your critical applications and services.<br>This group is blank by default.<br>You must configure this group to classify traffic that originates from trusted networks. |
| Watchlists | Classifies traffic that originates from networks that you want to monitor.<br>This group is blank by default. |

Groups and objects that include superflows are only for informational purposes and cannot be edited. Groups and objects that include bogons are configured by the automatic update function.

**NOTE**: You can use reference sets instead of remote networks to provide some of this functionality. Although you can assign a confidence level to an IP value in a reference table, reference sets are used only with single IPs and cannot be used with CIDR ranges. You can use a CIDR value after a remote network update, but not with weight or confidence levels.

RELATED DOCUMENTATION

Default Remote Service Groups | 340
Guidelines for Network Resources | 341
Managing Remote Networks Objects | 342

# Default Remote Service Groups

JSA includes the default remote service groups.

The following table describes the default remote service groups.

**Table 54: Default Remote Network Groups**

| Parameter | Description |
|---|---|
| IRC_Servers | Specifies traffic that originates from addresses commonly known as chat servers. |
| Online_Services | Specifies traffic that originates from addresses commonly known online services that might involve data loss. |
| Porn | Specifies traffic that originates from addresses commonly known to contain explicit pornographic material. |
| Proxies | Specifies traffic that originates from commonly known open proxy servers. |
| Reserved_IP_ Ranges | Specifies traffic that originates from reserved IP address ranges. |

Table 54: Default Remote Network Groups *(Continued)*

| Parameter | Description |
|---|---|
| Spam | Specifies traffic that originates from addresses commonly known to produce SPAM or unwanted email. |
| Spy_Adware | Specifies traffic that originates from addresses commonly known to contain spyware or adware. |
| Superflows | Specifies traffic that originates from addresses commonly known to produce superflows. |
| Warez | Specifies traffic that originates from addresses commonly known to contain pirated software. |

RELATED DOCUMENTATION

Guidelines for Network Resources | 341

Managing Remote Networks Objects | 342

Managing Remote Services Objects | 343

# Guidelines for Network Resources

Given the complexities and network resources that are required for JSA in large structured networks, follow the suggested guidelines.

The following list describes some of the suggested practices that you can follow:

- Bundle objects and use the **Network Activity** and **Log Activity** tabs to analyze your network data.

  Fewer objects create less input and output to your disk.

- Typically, for standard system requirements, do not exceed more than 200 objects per group.

  More objects might impact your processing power when you investigate your traffic.

RELATED DOCUMENTATION

Managing Remote Networks Objects | 342

Managing Remote Services Objects | 343

QID Map Overview | 343

# Managing Remote Networks Objects

After you create remote network groups, you can aggregate flow and event search results on remote network groups. You can also create rules that test for activity on remote network groups.

Use the **Remote Networks** window, you can add or edit a remote networks object.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **Remote Networks and Services Configuration** section, click **Remote Networks and Services Configuration**.
3. To add a remote networks object, click **Add** and enter values for the parameters.
4. To edit a remote networks object, follow these steps:.
   a. Double-click the group name.
   b. Select the profile and click the edit icon to edit the remote profile.
5. Click **Save**.
6. Click the previous icon to go back to the Remote Networks and Services window.
7. On the **Admin** tab, click **Deploy Changes**.

RELATED DOCUMENTATION

Managing Remote Services Objects | 343

QID Map Overview | 343

Guidelines for Network Resources | 341

# Managing Remote Services Objects

Remote services groups organize traffic that originates from user-defined network ranges or the Juniper Networks automatic update server. After you create remote service groups, you can aggregate flow and event search results, and create rules that test for activity on remote service groups.

Use the **Remote Services** window to add or edit a remote services object.

1. On the navigation menu (≡), click **Admin**.
2. In the **Remote Networks and Services Configuration** section, click **Remote Networks and Services Configuration**.
3. To add a remote services object, click **Add** and enter the parameter values.
4. To edit a remote services object, click the group that you want displayed, click the **Edit** icon and change the values.
5. Click **Save**.
6. Click **Return**.
7. Close the **Remote Services** window.
8. On the **Admin** tab menu, click **Deploy Changes**.

## RELATED DOCUMENTATION

QID Map Overview | **343**

Guidelines for Network Resources | **341**

Managing Remote Networks Objects | **342**

# QID Map Overview

### IN THIS SECTION

Use the JSA Identifier (QID) map utility to create, export, import, or modify user-defined QID map entries.

The QID map associates an event on an external device to a (QID).

See the following tasks for QID management:

- "Creating a QID Map Entry" on page 345
- "Modifying a QID Map Entry" on page 346
- "Importing Qid Map Entries" on page 347
- "Exporting QID Map Entries" on page 349

To run the utility, use the following syntax:

```
qidmap_cli.sh [-l|-c|-m|-i[-f <filename>]|-e[-f <filename>]|-d]
```

The following table describes the command-line options for the QID map utility.

**Table 55: QID Map Utility Options**

| Options | Description |
|---|---|
| -l | Lists the low-level category. |
| -c | Creates a QID map entry |
| -m | Modifies an existing user-defined QID map entry. |
| -i | Imports QID map entries. |
| -e | Exports existing user-defined QID map entries. |

Table 55: QID Map Utility Options *(Continued)*

| Options | Description |
|---|---|
| -f <filename> | If you include the -i or -e option, specifies a file name to import or export QID map entries. |
| -d | If you include the -i or -e option, specifies a delimiter for the import or export file. The default is a comma. |
| -h | Displays the help options. |

## Creating a QID Map Entry

Create a JSA Identifier (QID) Map Entry to map an event of an external device to QID.

1. Using SSH, log in to JSA as the root user.

2. To locate the low-level category for the QID map entry that you want to create, type the following command:

   ```
   /opt/qradar/bin/qidmap_cli.sh -l
   ```

   If you want to search for a particular low-level category, you can use the `grep` command to filter the results:

   ```
   /opt/qradar/bin/qidmap_cli.sh -l | grep <text>
   ```

3. Type the following command:

   **qidmap_cli.sh -c --qname <name> --qdescription <description> --severity <severity> --lowlevelcategoryid <ID>**

   The following table describes the command-line options for the QID map utility:

   | Options | Description |
   |---|---|
   | -c | Creates a QID map entry. |

*(Continued)*

| Options | Description |
|---|---|
| --qname *<name>* | The name that you want to associate with this QID map entry. The name can be up to 255 characters in length.<br>If you include spaces in the name, enclose the name value in double quotation marks. |
| --qdescription *<description>* | The description for this QID map entry. The description can be up to 2048 characters in length.<br>If you include spaces in the description, enclose the description value in double quotation marks. |
| --severity *<severity>* | The severity level that you want to assign to this QID map entry. The valid range is 1 - 10. |
| --lowlevelcategoryid *<ID>* | The low-level category ID you want to assign to this QID map entry. For more information, see the *Juniper Secure Analytics Administration Guide*. |

## Modifying a QID Map Entry

Modify an existing user-defined JSA Identifier (QID) map entry.

> **NOTE**: The `qidmap_cli` script cannot interact with QID entries that are associated with a specific custom Log Source Type. JSA has public APIs that can interact with QIDs in this range. The API is used as the supported mechanism for the operation. The QID map API is at **/data_classification/qid_records**. The API supports the GET, CREATE, and UPDATE functions. It does not support the DELETE function.

1. Using SSH, log in to JSA as the root user.

2. Type the following command:

```
qidmap_cli.sh -m --qid<QID> --qname <name> --qdescription <description> --severity <severity>
```

The following table describes the command-line options for the QID map utility:

| Options | Description |
|---|---|
| -m | Modifies an existing user-defined QID map entry. |
| --qid *<QID>* | The QID that you want to modify. |
| --qname *<name>* | The name that you want to associate with this QID map entry. The name can be up to 255 characters in length with no spaces. |
| --qdescription *<description>* | The description for this QID map entry. The description can be up to 2048 characters in length with no spaces. |
| --severity *<severity>* | The severity level that you want to assign to this QID map entry. The valid range is 0 - 10. |

## Importing Qid Map Entries

Using the JSA Identifier (QID) map utility, you can import QID map entries from a .txt file.

1. Create a **.txt** file that includes the user-defined QID map entries that you want to import. Ensure that each entry in the file is separated with a comma. Choose one of the following options:

   - If you want to import a new list of user-defined QID map entries, create the file with the following format for each entry:

     ```
     ,<name>,<description>,<severity>,<category>
     ```

     **,buffer,buffer_QID,7,18401 ,malware,malware_misc,8,18403**

- If you want to import an existing list of user-defined QID map entries, create the file with the following format for each entry:

  ```
  <qid>,<name>,<description>,<severity>
  ```

  **2000002,buffer,buffer_QID,7 2000001,malware,malware_misc**

The following table describes the command-line options of the QID utility.

| Options | Description |
|---|---|
| *<qid>* | The existing QID for the entry. This option is required if you want to import an existing exported list of QID entries.<br>To import new QID entries, do not use this option. The QID map utility assigns an identifier (QID) for each entry in the file. |
| --qname *<name>* | The name that you want to associate with this QID map entry. The name can be up to 255 characters in length with no spaces. |
| --qdescription *<description>* | The description for this QID map entry. The description can be up to 2048 characters in length with no spaces. |
| --severity *<severity>* | The severity level that you want to assign to this QID map entry. The valid range is 0 - 10. |
| --lowlevelcategoryid *<ID>* | The low-level category ID that you want to assign to this QID map entry.<br>This option is only necessary if you want to import a new list of QID entries. |

2. Save and close the file.

3. Using SSH, log in to JSA as the root user:

4. To import the QID map file, type the following command:

```
/opt/qradar/bin/qidmap_cli.sh -i -f <filename.txt>
```

The *<filename.txt>* option is the directory path and name of the file that contains the QID map entries. If any of the entries in the file cause an error, no entries in the file are enforced.

## Exporting QID Map Entries

View the mappings between the events of external devices and their unique identifiers by exporting your QID entries.

For QID map entries that you created, use the QID map utility to export the entries to a **.txt** file.

For an entire QID map that includes the default system QID entries, use the **idlist.sh** command.

1. Using SSH, log in to JSA as the root user.

2. To export the QID map file of user-defined entries, type the following command:

   ```
   /opt/qradar/bin/qidmap_cli.sh -e -f <filename.txt>
   ```

   The *<filename.txt>* option is the directory path and name of the file that you want to contain your QID map entries.

3. To export the entire QID map, type the following command:

   **/opt/qradar/bin/idlist.sh -e qid > *<filename.txt>***

4. To determine the last modified date of your QID map, run an SQL query.

   For example, if a QID has the identification number 64250088, type the following SQL query to retrieve its last modified date:

   **psql -U qradar -c "select qid,to_timestamp(serial/1000) as date from qidmap_serial where qid = 64250088;"**

RELATED DOCUMENTATION

Guidelines for Network Resources | 341

Managing Remote Networks Objects | 342

Managing Remote Services Objects | 343

15

CHAPTER

# Server Discovery

# Server Discovery

The **Server Discovery** function uses the Asset Profile database to discover different server types that are based on port definitions. Then, you can select the servers to add to a server-type building block for rules.

The **Server Discovery** function is based on server-type building blocks. Ports are used to define the server type. Thus, the server-type building block works as a port-based filter when you search the Asset Profile database.

For more information about building blocks, see the *Juniper Secure Analytics Users Guide*.

Use the **Server Discovery** function with JSA Vulnerability Manager to create exception rules for benign vulnerabilities. Reduce the number of vulnerabilities that you see for the following **Server Types**:

**Table 56: Server Type Vulnerabilities**

| Server Type | Vulnerability |
|---|---|
| FTP Servers | **FTP Server Present** |
| DNS Servers | **DNS Server is Running** |
| Mail Servers | **SMTP Server Detected** |
| Web Servers | **Web Service is Running** |

For more information about false positive vulnerabilities, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

## RELATED DOCUMENTATION

Discovering Servers | 352

# Discovering Servers

Use the **Assets** tab to discover servers on your network.

1. On the navigation menu (≡), click **Assets** to open the **Assets** tab.
2. On the **Assets** navigation menu, click **Server Discovery**.
3. From the **Server Type** list, select the server type that you want to discover.
4. Select one of the following options to determine the servers you want to discover:
   - To use the currently selected **Server Type** to search all servers in your deployment, select **All**.
   - To search servers in your deployment that were assigned to the currently selected **Server Type**, select **Assigned**.
   - To search servers in your deployment that are not assigned, select **Unassigned**.
5. To edit the standard server port list, click **Edit ports**.
6. From the **Network** list, select the network that you want to search.
7. Click **Discover Servers**.
8. In the **Matching Servers** table, select the check boxes of all servers you want to assign to the server role.
9. Click **Approve Selected Servers**.

## RELATED DOCUMENTATION

Server Discovery | **351**

16

CHAPTER

# Domain Segmentation

# Domain Segmentation

Segmenting your network into different domains helps to ensure that relevant information is available only to those users that need it.

You can create security profiles to limit the information that is available to a group of users within that domain. Security profiles provide authorized users access to only the information that is required to complete their daily tasks. You modify only the security profile of the affected users, and not each user individually.

You can also use domains to manage overlapping IP address ranges. This method is helpful when you are using a shared JSA infrastructure to collect data from multiple networks. By creating domains that represent a particular address space on the network, multiple devices that are in separate domains can have the same IP address and still be treated as separate devices.

**Figure 13: Domain Segmentation**

# Overlapping IP Addresses

An overlapping IP address is an IP address that is assigned to more than one device or logical unit, such as an event source type, on a network. Overlapping IP address ranges can cause significant problems for companies that merge networks after corporate acquisitions, or for Managed Security Service Providers (MSSPs) who are bringing on new clients.

JSA must be able to differentiate events and flows that come from different devices and that have the same IP address. If the same IP address is assigned to more than one event source, you can create domains to distinguish them.

For example, let's look at a situation where Company A acquires Company B and wants to use a shared instance of JSA to monitor the new company's assets. The acquisition has a similar network structure that results in the same IP address being used for different log sources in each company. Log sources that have the same IP address cause problems with correlation, reporting, searching, and asset profiling.

To distinguish the origin of the events and flows that come in to JSA from the log source, you can create two domains and assign each log source to a different domain. If required, you can also assign each event collector and flow processor to the same domain as the log source that sends events to them.

To distinguish the origin of the events and flows that come in to JSA from the log source, you can create two domains and assign each log source to a different domain. If required, you can also assign each

event collector, flow processor, or data gateway to the same domain as the log source that sends events to them.

To view the incoming events by domain, create a search and include the domain information in the search results.

RELATED DOCUMENTATION

Domain Definition and Tagging | 355
Creating Domains | 360
Domain Privileges That Are Derived from Security Profiles | 364

# Domain Definition and Tagging

IN THIS SECTION

Domains are defined based on JSA input sources. When events and flows come into JSA, the domain definitions are evaluated and the events and flows are tagged with the domain information.

## Specifying Domains for Events

The following diagram shows the precedence order for evaluating domain criteria for events.

Figure 14: Precedence Order for Events

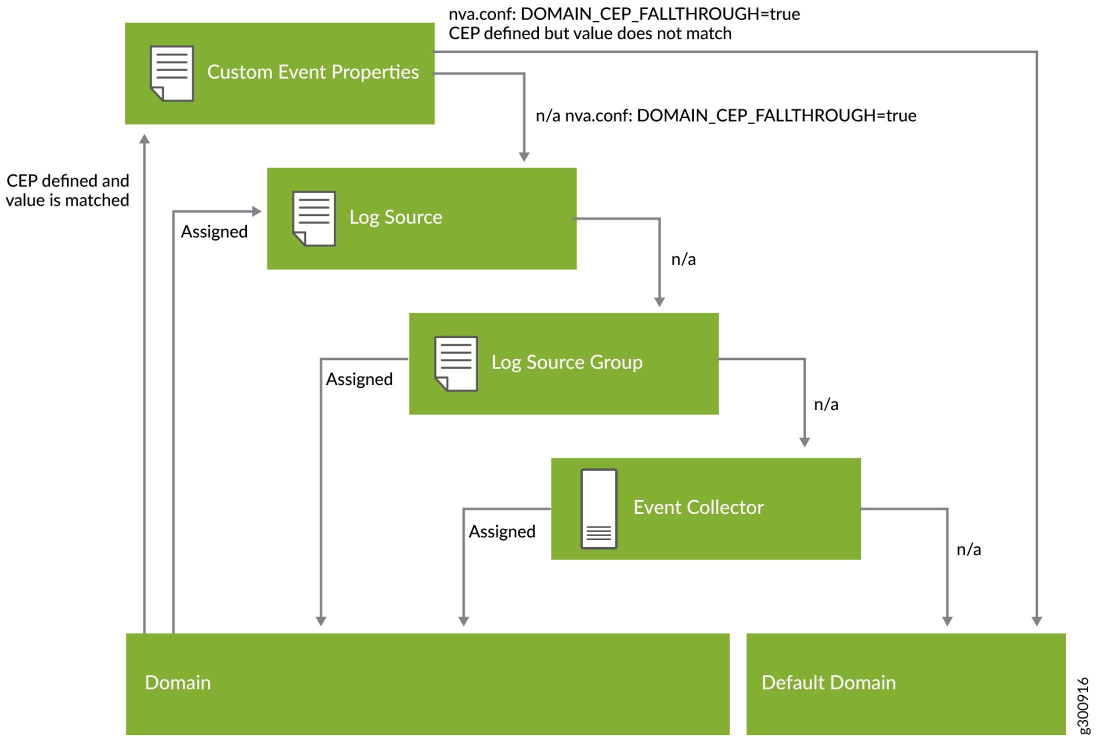
nva.conf: DOMAIN_CEP_FALLTHROUGH=true
CEP defined but value does not match
Custom Event Properties
n/a nva.conf: DOMAIN_CEP_FALLTHROUGH=true
CEP defined and value is matched
Assigned
Log Source
n/a
Assigned
Log Source Group
n/a
Assigned
Event Collector
n/a
Domain
Default Domain
g300916

These are the ways to specify domains for events:

- **Custom properties**--You can apply custom properties to the log messages that come from a log source.

  To determine which domain that specific log messages belong to, the value of the custom property is looked up against a mapping that is defined in the Domain Management editor.

  This option is used for multi-address-range or multi-tenant log sources, such as file servers and document repositories.

- **Disconnected Log Collector**--

  You can use a Disconnected Log Collector (DLC) for domain mapping. DLCs append their universally unique identifiers (UUIDs) to the Log Source Identifier value of the events they collect. Appending the UUID to the Log Source Identifier value ensures that the Log Source Identifier is unique.

- **Log sources**--You can configure specific log sources to belong to a domain.

  This method of tagging domains is an option for deployments in which an Event Collector can receive events from multiple domains.

- **Log source groups**--

You can assign log source groups to a specific domain. This option allows broader control over the log source configuration.

Any new log sources that are added to the log source group automatically get the domain tagging thatis associated with the log source group.

- **Event collectors--**If an event collector is dedicated to a specific network segment or IP address range, you can flag that entire event collector as part of that domain.

  All log sources that arrive at that event collector belong to the domain; therefore, any new auto-detected log sources are automatically added to the domain.

  > **NOTE**: If an event source is redirected from one event collector to another in a different domain, you must update its log source in one of the following ways:
  >
  > - Edit the log source to update the event collector information.
  > - Delete the log source and deploy the full configuration so that the event source is auto-detected on the new event collector.

  Unless the log source is updated, non-admin users with domain restrictions might not see offenses that are associated with the log source.

## Specifying Domains for Flows

The following diagram shows the precedence order for evaluating domain criteria for flows.

**Figure 15: Precedence Order for Flows**

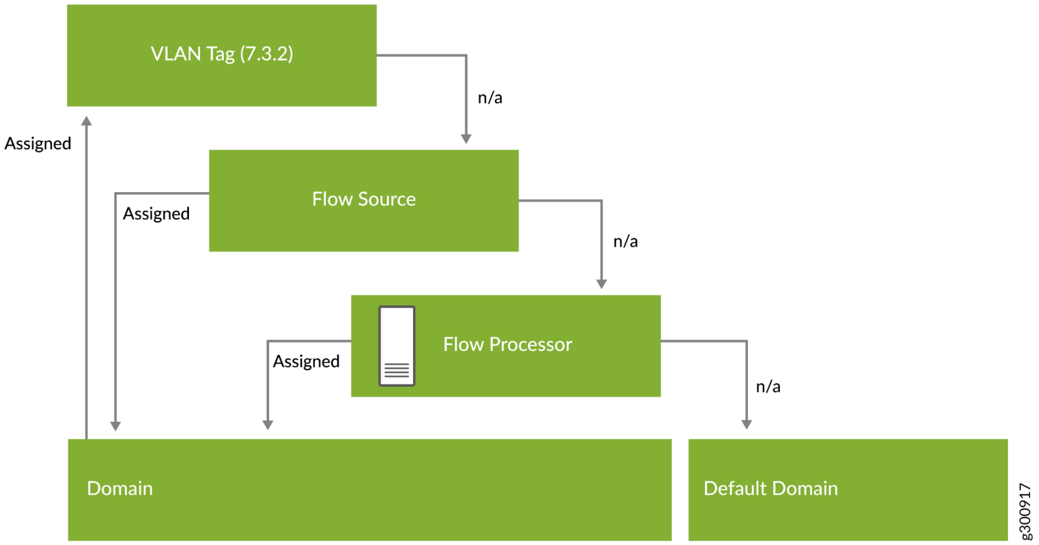

These are the ways to specify domains for flows:

- **Flow processors**-- You can assign specific Flow processors to a domain.

  All flow sources that arrive at that flow processor belong to the domain; therefore, any new auto-detected flow sources are automatically added to the domain.

- **Flow processors and data gateways**-- You can assign specific data gateways to a domain.

  All flow sources that arrive at that flow processor or data gateway belong to the domain; therefore, any new autodetected flow sources are automatically added to the domain.

- **Flow sources**-- You can designate specific flow sources to a domain.

  This option is useful when a single Flow processor is collecting flows from multiple network segments or routers that contain overlapping IP address ranges.

  This option is useful when a single flow processor or data gateway is collecting flows from multiple network segments or routers that contain overlapping IP address ranges.

- **Flow VLAN ID** —You can designate specific VLANs to a domain.

  This option is useful when you collect traffic from multiple network segments, often with overlapping IP ranges. This VLAN definition is based on the Enterprise and Customer VLAN IDs.

  The following information elements are sent from Flow Processor when flows that contain VLAN information are analyzed.

## Specifying Domains for Scan Results

You can also assign vulnerability scanners to a specific domain so that scan results are properly flagged as belonging to that domain. A domain definition can consist of all JSA input sources.

For information about assigning your network to preconfigured domains, see "Network Hierarchy" on page 123.

## Precedence Order for Evaluating Domain Criteria

When events and flows come into the JSA system, the domain criteria is evaluated based on the granularity of the domain definition.

If the domain definition is based on an event, the incoming event is first checked for any custom properties that are mapped to the domain definition. If the result of a regular expression that is defined in a custom property does not match a domain mapping, the event is automatically assigned to the default domain.

If the event does not match the domain definition for custom properties, the following order of precedence is applied:

1. DLC

2. Log source

3. Log source group

4. Event Collector

5. Event collector or data gateway

If the domain is defined based on a flow, the following order of precedence is applied:

1. Flow source

2. Flow Processor or data gateway

If a scanner has an associated domain, all assets that are discovered by the scanner are automatically assigned to the same domain as the scanner.

## Forwarding Data to Another JSA System

Domain information is removed when data is forwarded to another JSA system. Events and flows that contain domain information are automatically assigned to the default domain on the receiving JSA system. To identify which events and flows are assigned to the default domain, you can create a custom search on the receiving system. You might want to reassign these events and flows to a user-defined domain.

RELATED DOCUMENTATION

Creating Domains | 360

Domain Privileges That Are Derived from Security Profiles | 364

Domain-specific Rules and Offenses | 366

# Creating Domains

Use the **Domain Management** window to create domains based on JSA input sources.

Use the following guidelines when you create domains:

- Everything that is not assigned to a user-defined domain is automatically assigned to the default domain. Users who have limited domain access should not have administrative privileges because this privilege grants unlimited access to all domains.

- You can map the same custom property to two different domains, however the capture result must be different for each one.

- You cannot assign a log source, log source group, or event collector to two different domains. When a log source group is assigned to a domain, each of the mapped attributes is visible in the **Domain Management** window.

- You cannot assign a log source, log source group, event collector, or data gateway to two different domains. When a log source group is assigned to a domain, each of the mapped attributes is visible in the **Domain Management** window.

Security profiles must be updated with an associated domain. Domain-level restrictions are not applied until the security profiles are updated, and the changes deployed.

1. On the navigation menu
   (

☰
), click **Admin**.

2. In the **System Configuration** section, click **Domain Management**.
3. To add a domain, click **Add** and type a unique name and description for the domain.

> **TIP**: You can check for unique names by typing the name in the **Input domain name** search box.

4. Depending on the domain criteria to be defined, click the appropriate tab.
   - To define the domain based on a custom property, log source group, log source, or event collector, click the **Events** tab.
   - To define the domain based on a custom property, log source group, log source, event collector, or data gateway, click the Events tab.
   - To define the domain based on a flow source, flow processor or data gateway, click the **Flows** tab.
   - To define the domain based on a flow source, click the **Flows** tab.
   - To define the domain based on a scanner, including JSA Vulnerability Manager scanners, click the **Scanners** tab.
5. To assign a custom property to a domain, in the **Capture Result** box, type the text that matches the result of the regular expression (regex) filter.

> **NOTE**: You must select the **Optimize parsing for rules, reports, and searches** check box in the **Custom Event Properties** window to parse and store the custom event property. Domain segmentation will not occur if this option is not checked.

6. From the list, select the domain criteria and click **Add**.
7. After you add the source items to the domain, click **Create**.

"Create security profiles" on page 23 to define which users have access to the domains. After you create the first domain in your environment, you must update the security profiles for all non-administrative users to specify the domain assignment. In domain-aware environments, non-administrative users whose security profile does not specify a domain assignment will not see any log activity or network activity.

Review the hierarchy configuration for your network, and assign existing IP addresses to the proper domains. For more information, see "Network Hierarchy" on page 123.

RELATED DOCUMENTATION

Domain Privileges That Are Derived from Security Profiles | 364

Domain-specific Rules and Offenses | 366

Example: Domain Privilege Assignments Based on Custom Properties | 370

# Creating Domains for VLAN Flows

Use the **Domain Management** window to create domains based on JSA VLAN flow sources.

In JSA, you can assign domains to incoming flows based on the VLAN information that is contained in the flow. The incoming flows are mapped to domains that contain the same VLAN definition.

1. On the navigation menu
   (
   ☰
   ), click **Admin**.
2. In the **System Configuration** section, click **Domain Management**.
3. Click **Add** and type a unique name and description for the domain.

   **TIP**: You can check for unique names by typing the name in the **Input domain name** search box.

   **Figure 16: Input Domain Name**

   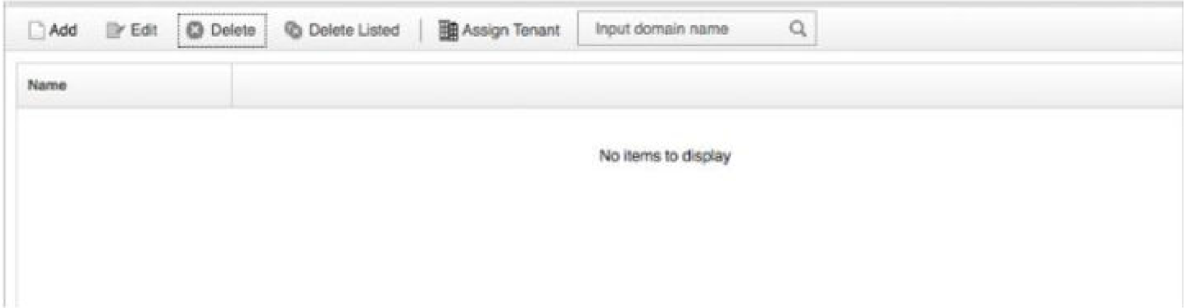

4. Click the **Flows** tab, and then select **Flow VLAN IDs**.
5. Select the enterprise VLAN ID and Customer VLAN ID values that match the values on the incoming flows, and then click **Add**.

**NOTE**:

- The Enterprise VLAN ID (IE): 82 is specified by Private Enterprise Number (PEN): 2, Information Element (IE) on incoming flows.

- The Customer VLAN ID is specified by PEN: 2 and IE: 83 on incoming flows.

**Figure 17: New Domain**

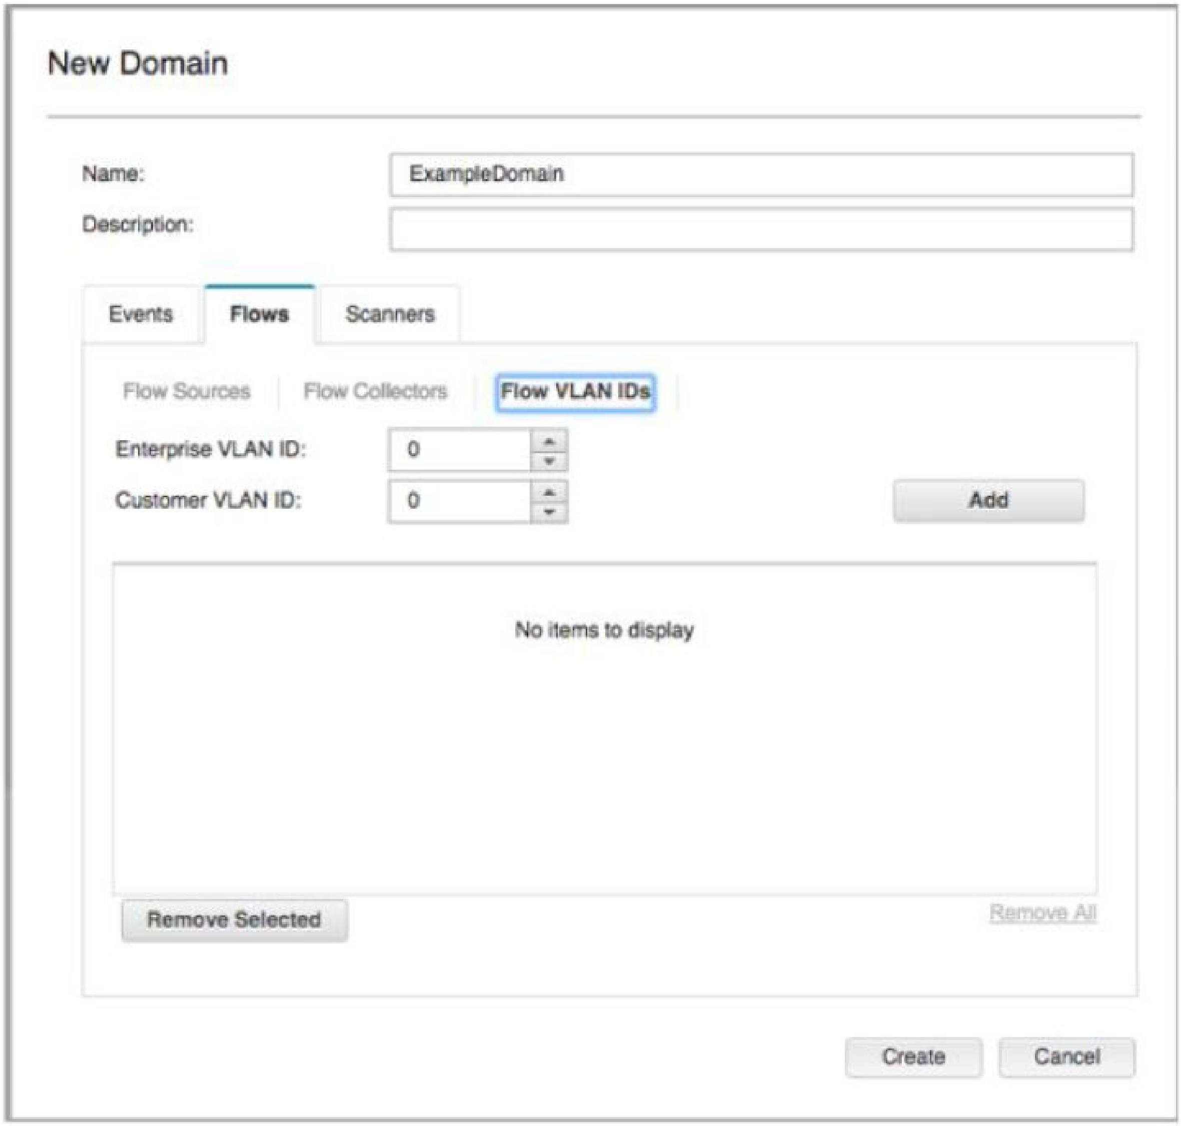

New Domain

Name: ExampleDomain

Description:

Enterprise VLAN ID: 0

Customer VLAN ID: 0

No items to display

6. In the Name field, type a unique name for the domain and then click **Create**.

The domain definition is created and incoming flows are mapped. Tenant assignment to a domain occurs as normal.

RELATED DOCUMENTATION

# Domain Privileges That Are Derived from Security Profiles

IN THIS SECTION

You can use security profiles to grant domain privileges and ensure that domain restrictions are respected throughout the entire JSA system. Security profiles also make it easier to manage privileges for a large group of users when your business requirements suddenly change.

Users can see only data within the domain boundaries that are set up for the security profiles that are assigned to them. Security profiles include domains as one of the first criteria that is evaluated to restrict access to the system. When a domain is assigned to a security profile, it takes priority over other security permissions. After domain restrictions are evaluated, individual security profiles are assessed to determine network and log permissions for that particular profile.

For example, a user is given privileges to Domain_2 and access to network 10.0.0.0/8. That user can see only events, offenses, assets, and flows that come from Domain_2 and contain an address from the 10.0.0.0/8 network.

As a JSA administrator, you can see all domains and you can assign domains to non-administrative users. Do not assign administrative privileges to users whom you want to limit to a particular domain.

Security profiles must be updated with an associated domain. Domain-level restrictions are not applied until the security profiles are updated, and the changes are deployed.

When you assign domains to a security profile, you can grant access to the following types of domains:

- **User-defined domains**--You can create domains that are based on input sources by using the Domain Management tool. For more information, see "Creating Domains" on page 360.
- **Default domain**--Everything that is not assigned to a user-defined domain is automatically assigned to the default domain. The default domain contains system-wide events.

  **NOTE**: Users who have access to the default domain can see system-wide events without restriction. Ensure that this access is acceptable before you assign default domain access to users. All administrators have access to the default domain.

  Any log source that gets auto-discovered on a shared event collector (one that is not explicitly assigned to a domain), is auto-discovered on the default domain. These log sources require manual intervention. To identify these log sources, you must periodically run a search in the default domain that is grouped by log source.

  Any log source that gets auto-discovered on a shared event collector or data gateway (one that is not explicitly assigned to a domain), is auto-discovered on the default domain. These log sources require manual intervention. To identify these log sources, you must periodically run a search in the default domain that is grouped by log source.
- **All domains**--Users who are assigned to a security profile that has access to **All Domains** can see all active domains within the system, the default domain, and any domains that were previously deleted across the entire system. They can also see all domains that are created in the future.

If you delete a domain, it cannot be assigned to a security profile. If the user has the **All domains** assignment, or if the domain was assigned to the user before it was deleted, the deleted domain is returned in historical search results for events, flows, assets, and offenses. You can't filter by deleted domains when you run a search.

Administrative users can see which domains are assigned to the security profiles on the **Summary** tab in the **Domain Management** window.

## Rule Modifications in Domain-aware Environments

Rules can be viewed, modified, or disabled by any user who has both the **Maintain Custom Rules** and **View Custom Rules** permissions, regardless of which domain that user belongs to.

**NOTE**: When you add the **Log Activity** capability to a user role, the **Maintain Custom Rules** and **View Custom Rules** permissions are automatically granted. Users who have these permissions have access to all log data for all domains, and they can edit rules in all domains, even if their security profile settings have domain-level restrictions. To prevent domain users from being able to access log data and modify rules in other domains, edit the user role and remove the **Maintain Custom Rules** and **View Custom Rules** permissions.

## Domain-aware Searches

You can use domains as search criteria in custom searches. Your security profile controls which domains you can search against.

System-wide events and events that are not assigned to a user-defined domain are automatically assigned to the default domain. Administrators, or users who have a security profile that provides access to the default domain, can create a custom search to see all events that are not assigned to a user-defined domain.

The default domain administrator can share a saved search with other domain users. When the domain user runs that saved search, the results are limited to their domain.

### RELATED DOCUMENTATION

Domain-specific Rules and Offenses | 366

Example: Domain Privilege Assignments Based on Custom Properties | 370

Creating Domains | 360

# Domain-specific Rules and Offenses

A rule can work in the context of a single domain or in the context of all domains. Domain-aware rules provide the option of including the **And Domain Is** test.

The following diagram shows an example using multiple domains.

**Figure 18: Domain Aware Rules**

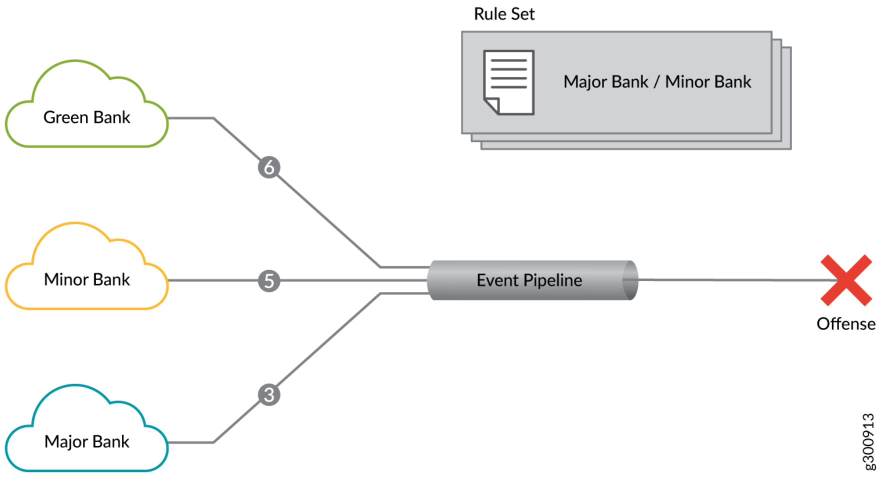

You can restrict a rule so that it is applied only to events that are happening within a specified domain. An event that has a domain tag that is different from the domain that is set on the rule does not trigger an event response.

In an JSA system that does not have user-defined domains, a rule creates an offense and keeps contributing to it each time the rule fires. In a domain-aware environment, a rule creates a new offense each time the rule is triggered in the context of a different domain.

Rules that work in the context of all domains are referred to as system-wide rules. To create a system-wide rule that tests conditions across the entire system, select **Any Domain** in the domain list for the **And Domain Is** test. An **Any Domain** rule creates an **Any Domain** offense.

- **Single-domain rule**--If the rule is a stateful rule, the states are maintained separately for each domain. The rule is triggered separately for each domain. When the rule is triggered, offenses are created separately for each domain that is involved and the offenses are tagged with those domains.

- **Single-domain offense**--The offense is tagged with the corresponding domain name. It can contain only events that are tagged with that domain.

- **System-wide rule**--If the rule is a stateful rule, a single state is maintained for the whole system and domain tags are ignored. When the rule runs, it creates or contributes to a single system-wide offense.

- **System-wide offense**--The offense is tagged with **Any Domain**. It contains only events that are tagged with all domains.

The following table provides examples of domain-aware rules. The examples use a system that has three domains that are defined: Domain_A, Domain_B, and Domain_C.

The rule examples in the following table may not be applicable in your JSA environment. For example, rules that use flows and offenses are not applicable in Log Manager.

**Table 57: Domain-aware Rules**

| Domain text | Explanation | Rule response |
|---|---|---|
| **domain is one of: Domain_A** | Looks only at events that are tagged with `Domain_A` and ignores rules that are tagged with other domains. | Creates or contributes to an offense that is tagged with `Domain_A`. |
| **domain is one of: Domain_A** and a stateful test that is defined as **when HTTP flow is detected 10 times within 1 minute** | Looks only at events that are tagged with `Domain_A` and ignores rules that are tagged with other domains. | Creates or contributes to an offense that is tagged with `Domain_A`. A single state, an HTTP flow counter, gets maintained for Domain_A. |
| **domain is one of: Domain_A, Domain_B** | Looks only at events that are tagged with `Domain_A` and `Domain_B` and ignores events that are tagged with `Domain_C`.<br>This rule behaves as two independent instances of a single domain rule, and creates separate offenses for different domains. | For data that is tagged with `Domain_A`, it creates or contributes to a single domain offense that is tagged with `Domain_A`.<br>For data that is tagged with `Domain_B`, it creates or contributes to a single domain offense that is tagged with `Domain_B`. |
| **domain is one of: Domain_A, Domain_B** and a stateful test that is defined as **when HTTP flow is detected 10 times within 1 minute** | Looks only at events that are tagged with `Domain_A` and `Domain_B` and ignores events that are tagged with `Domain_C`.<br>This rule behaves as two independent instances of a single domain rule, and maintains two separate states (HTTP flow counters) for two different domains. | When the rule detects 10 HTTP flows that are tagged with `Domain_A` within a minute, it creates or contributes to an offense that is tagged with `Domain_A`.<br>When the rule detects 10 HTTP flows that are tagged with `Domain_B` within a minute, it creates or contributes to an offense that is tagged with `Domain_B`. |

Table 57: Domain-aware Rules *(Continued)*

| Domain text | Explanation | Rule response |
|---|---|---|
| No domain test defined | Looks at events that are tagged with all domains and creates or contributes to offenses on a per-domain basis. | Each independent domain has offenses that are generated for it, but offenses do not contain contributions from other domains. |
| A rule has a stateful test that is defined as **when HTTP flow is detected 10 times within 1 minute** and no domain test is defined | Looks at events that are tagged with `Domain_A`, `Domain_B`, or `Domain_C`. | Maintains separate states and creates separate offenses for each domain. |
| **domain is one of: Any Domain** | Looks at all events, regardless of which domain it is tagged with. | Creates or contributes to a single system-wide offense that is tagged with `Any Domain`. |
| **domain is one of: Any Domain** and a stateful test that is defined as **when HTTP flow is detected 10 times within 1 minute** | Looks at all events, regardless of which domain it is tagged with, and it maintains a single state for all domains. | Creates or contributes to a single system-wide offense that is tagged with `Any Domain`.<br>For example, if it detects 3 events that are tagged with `Domain_A`, 3 events that are tagged with `Domain_B`, and 4 events that are tagged with `Domain_C` within 1 minute, it creates an offense because it detected 10 events in total. |
| **domain is one of: Any Domain, Domain_A** | Works the same as a rule that has **domain is one of: Any Domain.** | When the domain test includes `Any Domain`, any single domains that are listed are ignored. |

When you view the offense table, you can sort the offenses by clicking the **Domain** column. The **Default Domain** is not included in the sort function so it does not appear in alphabetical order. However, it appears at the top or bottom of the **Domain** list, depending on whether the column is sorted in ascending or descending order. **Any Domain** does not appear in the list of offenses.

RELATED DOCUMENTATION

Example: Domain Privilege Assignments Based on Custom Properties | 370

Creating Domains | 360

Domain Privileges That Are Derived from Security Profiles | 364

# Example: Domain Privilege Assignments Based on Custom Properties

If your log files contain information that you want to use in a domain definition, you can expose the information as a custom event property.

You assign a custom property to a domain based on the capture result. You can assign the same custom property to multiple domains, but the capture results must be different.

For example, a custom event property, such as `userID`, might evaluate to a single user or a list of users. Each user can belong to only one domain.

In the following diagram, the log sources contain user identification information that is exposed as a custom property, `userID`. The event collector returns two user files, and each user is assigned to only one domain. In this case, one user is assigned to Domain: 9 and the other user is assigned to Domain: 12.

**Figure 19: Assigning Domains by Using Custom Event Property**

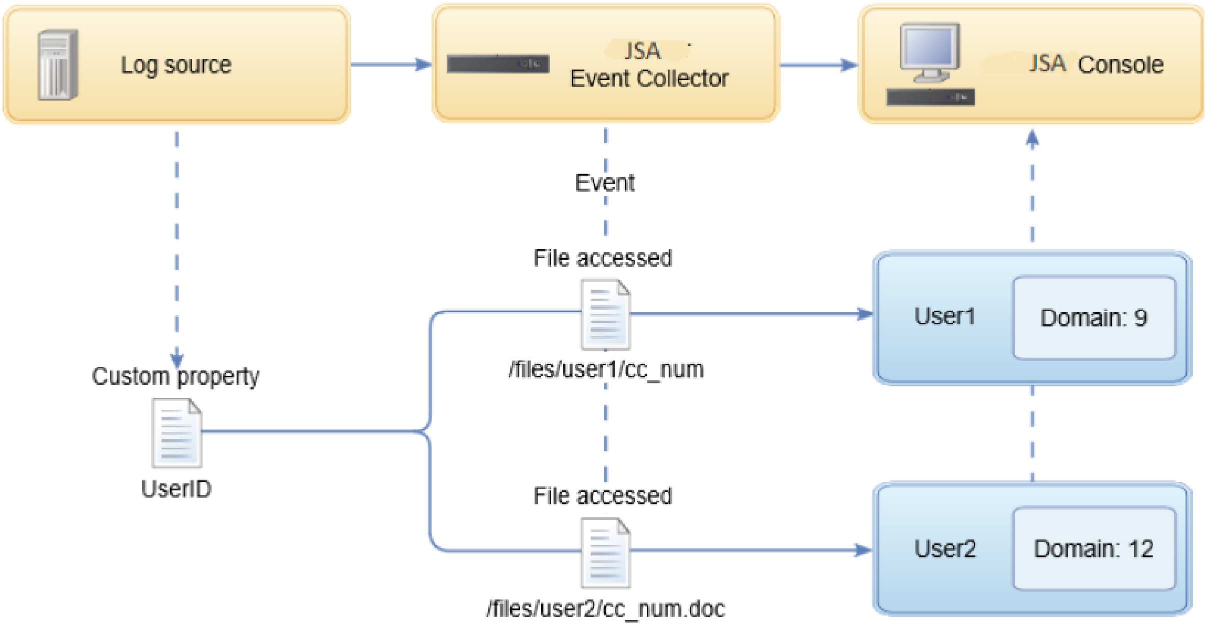

If the capture results return a user that is not assigned to a specific user-defined domain, that user is automatically assigned to the default domain. Default domain assignments require manual intervention. Perform periodic searches to ensure that all entities in the default domain are correctly assigned.

**NOTE**: Before you use a custom property in a domain definition, ensure that **Optimize parsing for rules, reports, and searches** is checked on the **Custom Event Properties** window. This option ensures that the custom event property is parsed and stored when JSA receives the event for the first time. Domain segmentation doesn't occur if this option is not checked.

## RELATED DOCUMENTATION

Creating Domains | 360

Domain Privileges That Are Derived from Security Profiles | 364

17

CHAPTER

# Multitenant Management

# Multitenant Management

Multitenant environments allow Managed Security Service Providers (MSSPs) and multi-divisional organizations to provide security services to multiple client organizations from a single, shared JSA deployment. You don't have to deploy a unique JSA instance for each customer.

**Figure 20: Multitenant Environments**

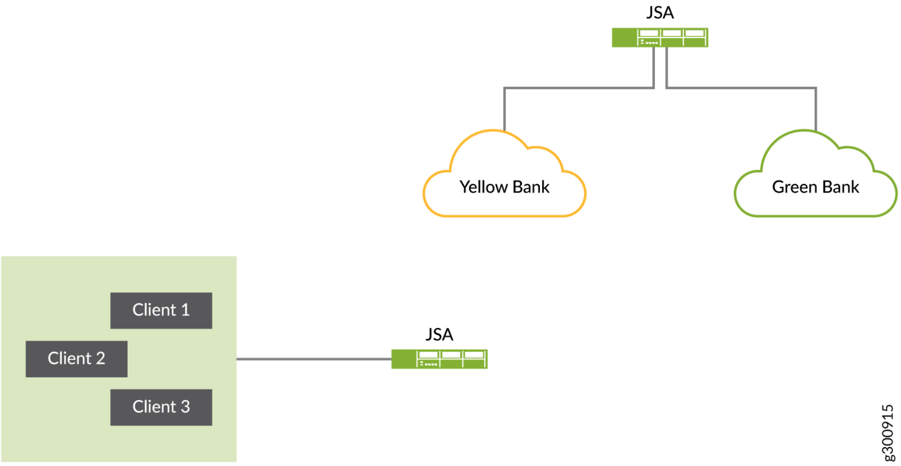

In a multitenant deployment, you ensure that customers see only their data by creating domains that are based on their JSA input sources. Then, use security profiles and user roles to manage privileges for large groups of users within the domain. Security profiles and user roles ensure that users have access to only the information that they are authorized to see.

# User Roles in a Multitenant Environment

IN THIS SECTION

Multitenant environments include a service provider and multiple tenants. Each role has distinct responsibilities and associated activities.

## Service Provider

The service provider owns the system and manages its use by multiple tenants. The service provider can see data across all tenants. The Managed Security Service Provider (MSSP) administrator is typically responsible for the following activities:

- Administers and monitors the system health of the JSA deployment.
- Provisions new tenants.
- Creates roles and security profiles for tenant administrators and users.
- Secures the system against unauthorized access.
- Creates domains to isolate tenant data.
- Deploys changes that the tenant administrator made in the tenant environment.
- Monitors JSA licenses.
- Collaborates with the tenant administrator.

## Tenants

Each tenancy includes a tenant administrator and tenant users. The tenant administrator can be an employee of the tenant organization, or the service provider can administer the tenant on behalf of the customer.

The tenant administrator is responsible for the following activities:

- Configures "Network Hierarchy" on page 123 definitions within their own tenancy.
- Configures and manages tenant data.
- Views log sources.
- Collaborates with the MSSP administrator.

The tenant administrator can configure tenant-specific deployments, but they can't access or change the configuration for another tenant. They must contact the MSSP administrator to deploy changes in the JSA environment, including network hierarchy changes within their own tenant.

Tenant users have no administrative privileges and can see only the data that they have access to. For example, a user can have privileges to view data from only 1 log source within a domain that has multiple log sources.

RELATED DOCUMENTATION

Domains and Log Sources in Multitenant Environments | 375
Provisioning a New Tenant | 377
Monitoring License Usage in Multitenant Deployments | 378

# Domains and Log Sources in Multitenant Environments

IN THIS SECTION

Use domains to separate overlapping IP addresses, and to assign sources of data, such as events and flows, into tenant-specific data sets.

When events or flows come into JSA, JSA evaluates the domain definitions that are configured, and the events and flows are assigned to a domain. A tenant can have more than one domain. If no domains are configured, the events and flows are assigned to the default domain.

## Domain Segmentation

Domains are virtual buckets that you use to segregate data based on the source of the data. They are the building blocks for multitenant environments. You configure domains from the following input sources:

- Event and flow processors
- Flow sources
- Log sources and log source groups
- Custom properties
- Scanners

A multitenant deployment might consist of a basic hardware configuration that includes one JSA Console, one centralized event processor, and then one event collector for each customer. In this configuration, you define domains at the collector level, which then automatically assigns the data that is received by JSA to a domain.

To consolidate the hardware configuration even further, you can use one collector for multiple customers. If log or flow sources are aggregated by the same collector but belong to different tenants, you can assign the sources to different domains. When you use domain definitions at the log source level, each log source name must be unique across the entire JSA deployment.

If you need to separate data from a single log source and assign it to different domains, you can configure domains from custom properties. JSA looks for the custom property in the payload, and assigns it to the correct domain. For example, if you configured JSA to integrate with a Check Point Provider-1 device, you can use custom properties to assign the data from that log source to different domains.

## Automatic Log Source Detection

When domains are defined at the collector level and the dedicated event collector is assigned to a single domain, new log sources that are automatically detected are assigned to that domain. For example, all log sources that are detected on `Event_Collector_1` are assigned to `Domain_A`. All log sources that are automatically collected on `Event_Collector_2` are assigned to `Domain_B`.

When domains are defined at the log source or custom property level, log sources that are automatically detected and are not already assigned to a domain are automatically assigned to the default domain. The MSSP administrator must review the log sources in the default domain and allocate them to the correct client domains. In a multitenant environment, assigning log sources to a specific domain prevents data leakage and enforces data separation across domains.

RELATED DOCUMENTATION

Provisioning a New Tenant | 377

Monitoring License Usage in Multitenant Deployments | 378

Rules Management in Multitenant Deployments | 382

# Provisioning a New Tenant

As a Managed Security Services Provider (MSSP) administrator, you are using a single instance of JSA to provide multiple customers with a unified architecture for threat detection and prioritization.

In this scenario, you are onboarding a new client. You provision a new tenant and create a tenant administrator account that does limited administrative duties within their own tenant. You limit the access of the tenant administrator so that they can't see or edit information in other tenants.

Before you provision a new tenant, you must create the data sources, such as log sources or flow processors, for the customer and assign them to a domain.

Complete the following tasks by using the tools on the **Admin** tab to provision the new tenant in JSA:

1. To create the tenant, click **Tenant Management**.

   For information about setting events per second (EPS) and flows per minute (FPM) limits for each tenant, see "Monitoring License Usage in Multitenant Deployments" on page 378.

2. To assign domains to the tenant, click **Domain Management**.

3. To create the tenant administrator role and grant the **Delegated Administration** permissions, click **User Roles**.

   In a multitenant environment, tenant users with **Delegated administration** permissions can see only data for their own tenant environment. If you assign other administrative permissions that are not part of **Delegated Administration**, access is no longer restricted to that domain.

4. To create the tenant security profiles and restrict data access by specifying the tenant domains, click **Security Profiles**.

5. To create the tenant users and assign the user role, security profile, and tenant, click **Users**.

RELATED DOCUMENTATION

Monitoring License Usage in Multitenant Deployments | 378
Rules Management in Multitenant Deployments | 382
Network Hierarchy Updates in a Multitenant Deployment | 384

# Monitoring License Usage in Multitenant Deployments

IN THIS SECTION

As the Managed Security Service Provider (MSSP) administrator, you monitor the event and flow rates across the entire JSA deployment.

When you create a tenant, you can set limits for both events per second (EPS) and flows per minute (FPM). By setting EPS and FPM limits for each tenant, you can better manage license capacities across multiple clients. If you have a processor that is collecting events or flows for a single customer, you do not need to assign tenant EPS and FPM limits. If you have a single processor that collects events or flows for multiple customers, you can set EPS and FPM limits for each tenant.

If you set the EPS and FPM limits to values that exceed the limits of either your software licenses or the appliance hardware, the system automatically throttles the events and flows for that tenant to ensure that the limits are not exceeded. If you do not set EPS and FPM limits for tenants, each tenant receives events and flows until either the license limits or the appliance limits are reached. The licensing limits are applied to the managed host. If you regularly exceed the license limitations, you can get a different license that is more suitable for your deployment.

## Viewing the Cumulative License Limits in Your Deployment

The EPS and FPM rates that you set for each tenant are not automatically validated against your license entitlements. To see the cumulative limits for the software licenses that are applied to the system as compared to the appliance hardware limits, do these steps:

1. On the navigation menu ( ≡ ), click **Admin** to open the admin tab.

2. In the **System Configuration** section, click **System and License Management**.

3. Expand **Deployment Details** and hover your mouse pointer over **Event Limit** or **Flow Limit**.

## Viewing EPS Rates Per Log Source

Use the **Advanced Search** field to enter an Ariel Query Language (AQL) query to view the EPS rates for log sources.

1. On the **Log Activity** tab, select **Advanced Search** from the list on the **Search** toolbar.

2. To view the EPS per log source, type the following AQL query in the **Advanced Search** field:

   **select logsourcename(logsourceid) as LogSource, sum(eventcount) / ( ( max(endTime) - min(startTime) ) / 1000 ) as EPS from events group by logsourceid order by EPS desc last 24 hours**

   The date values for (`endTime` and (`startTime`) must be represented in milliseconds since the UNIX Epoch January 1st 1970.

## Viewing EPS Rates Per Domain

Use the **Advanced Search** field to enter an Ariel Query Language (AQL) query to view the EPS rates for domains.

1. On the **Log Activity** tab, select **Advanced Search** from the drop-down list box on the **Search** toolbar.

2. To view the EPS per domain, type the following AQL query in the **Advanced Search** field:

   **select DOMAINNAME(domainid) as LogSource, sum(eventcount) / ( ( max(endTime) - min(startTime)) / 1000 ) as EPS from events group by domainid order by EPS desc last 24 hours**

The date values for `(endTime` and `(startTime)` must be represented in milliseconds since the UNIX Epoch January 1st 1970.

If you want to view average EPS rates for log sources only, click **Log Sources** in the **Data Sources** pane on the **Admin** tab. You can use this to quickly identify configuration issues with log sources that are failing to report.

## Viewing Individual License Limits in Your Deployment

The EPS and FPM rates that you set for each tenant are not automatically validated against your license entitlements. To see the individual limits for the software licenses that are applied to the system as compared to the appliance hardware limits, do these steps:

1. On the navigation menu
   (
   ≡
   ), click **Admin** to open the admin tab.

2. In the **System Configuration** section, click **System and License Management**.

3. Expand **Deployment Details** and hover your mouse over **Event Limit** or **Flow Limit**.

## Viewing the EPS Rate for an Individual Log Source

Use the **Advanced Search** field to enter an Ariel Query Language (AQL) query to view the EPS rate for an individual log source.

1. On the **Log Activity** tab, select **Advanced Search** from the list on the **Search** toolbar.

2. To get a log source ID, type the following AQL query in the Advanced Search field: Select `domainid,logsourceid,LOGSOURCENAME(logsourceid) from events GROUP BY domainid,logsourceid order by domainid ASC last 1 HOURS`.

3. To view the EPS rate for your selected log source, type the following AQL query in the **Advanced Search** field:

   **select logsourcename(logsourceid) as LogSource, sum(eventcount) / ( ( max(endTime) - min(startTime) ) / 1000 ) as EPS from events where logsourceid=*logsourceid* group by logsourceid order by EPS desc last 24 hours**

## Viewing the EPS Rate for an Individual Domain

Use the **Advanced Search** field to enter an Ariel Query Language (AQL) query to view the EPS rate for an individual domain.

1. On the **Log Activity** tab, select **Advanced Search** from the list on the **Search** toolbar.

2. To get a domain ID, type the following AQL query in the **Advanced Search** field: Select `select domainid, DOMAINNAME(domainid) from events GROUP BY domainid last 1 HOURS`.

3. To view the EPS rate for your selected domain, type the following AQL query in the **Advanced Search** field:

   **select DOMAINNAME(domainid) as LogSource, sum(eventcount) / ( ( max(endTime) - min(startTime)) / 1000 ) as EPS from events where domainid=*domainid* group by domainid order by EPS desc last 24 hours**

## Detecting Dropped Events and Flows

Events and flows are dropped when the JSA processing pipeline can't handle the volume of incoming events and flows, or when the number of events and flows exceeds the license limits for your deployment. You can look at the JSA log file messages when these situations occur.

1. Use SSH to log in to JSA as the root user.

2. View the **/var/log/qradar.error** log file and look for these messages:

   These messages indicate that events or flows were dropped:

   ```
   [Tenant:[tenantID]:[tenantName] Event dropped while attempting to add to Tenant Event Throttle queue. The Tenant Event Throttle queue is full.
   ```

   ```
   [Tenant:[tenantID]:[tenantName] Flow dropped while attempting to add to Tenant Flow Throttle queue. The Tenant Flow Throttle queue is full.
   ```

   These messages indicate that the processing pipeline was near capacity:

   ```
   Throttle processor cannot keep up with events. TENANT_QUEUE_THREAD_INTERVAL_IN_MILLISEC is probably too short.
   ```

   ```
   Throttle processor cannot keep up with flows. TENANT_QUEUE_THREAD_INTERVAL_IN_MILLISEC is probably too short.
   ```

   If this warning persists, JSA might drop events or flows.

If your system is dropping events and flows, you can expand your license to handle more data or you can set more restrictive EPS and FPM limits for each tenant.

RELATED DOCUMENTATION

# Rules Management in Multitenant Deployments

IN THIS SECTION

In a multitenant environment, you must customize rules to make them tenant-aware. Tenant-aware rules use the **when the domain is one of the following** rule test, but the domain modifier determines the scope of the rule.

The following table shows how you can use the domain modifier to change the scope of rules in a multitenant deployment.

**Table 58: Scope Of Rules in a Multitenant Environment**

| Rule scope | Description | Rule test example |
|---|---|---|
| Single domain rules | These rules include only 1 domain modifier. | **and when the domain is one of the following:** *manufacturing* |
| Single tenant rules | These rules include all the domains that are assigned to the tenant. Use single tenant rules to correlate events across multiple domains within a single tenant. | **and when the domain is one of the following:** *manufacturing, finance, legal* |
| Global rules | These rules use the **Any domain** modifier and run across all tenants. | **and when the domain is one of the following:** *Any domain* |

By being domain-aware, the custom rules engine (CRE) automatically isolates event correlations from different tenants by using their respective domains. For more information about working with rules in a domain-segmented network, see "Domain Segmentation" on page 376.

## Restricting Log Activity Capabilities for Tenant Users

To ensure that the tenant administrator and users can view the log data for only their tenant, you must restrict the permissions for the **Log Activity** capability.

When you add the **Log Activity** capability to a user role, the **Maintain Custom Rules** and **View Custom Rules** permissions are automatically granted. Users who have these permissions have access to all log data for all domains. They can edit rules in all domains, even if their security profile settings have domain-level restrictions.

To prevent users from being able to access log data and modify rules in other domains or tenants, edit the user role and remove the **Maintain Custom Rules** and **View Custom Rules** permissions. Without these permissions, the tenant administrator and users cannot change rules, including those rules in their own domain.

1. On the navigation menu (≡), click **Admin**.

2. In the **System Configuration** section, click **User Roles** and select the user role that you want to edit.

3. Under **Log Activity**, clear the **Maintain Custom Rules** and **View Custom Rules** check boxes.

4. Click **Save**.

### RELATED DOCUMENTATION

Network Hierarchy Updates in a Multitenant Deployment | **384**

Retention Policies for Tenants | **384**

Monitoring License Usage in Multitenant Deployments | **378**

# Network Hierarchy Updates in a Multitenant Deployment

JSA uses the network hierarchy to understand and analyze the network traffic in your environment.

Tenant administrators who have the **Define network hierarchy** permission can change the network hierarchy within their own tenant.

Network hierarchy changes require a full configuration deployment to apply the updates in the JSA environment. Full configuration deployments restart all JSA services, and data collection for events and flows stops until the deployment completes. Tenant administrators must contact the Managed Security Service Provider (MSSP) administrator to deploy the changes. MSSP administrators can plan the deployment during a scheduled outage, and notify all tenant administrators in advance.

In a multitenant environment, the network object name must be unique across the entire deployment. You cannot use network objects that have the same name, even if they are assigned to different domains.

RELATED DOCUMENTATION

Retention Policies for Tenants | 384

Monitoring License Usage in Multitenant Deployments | 378

Rules Management in Multitenant Deployments | 382

# Retention Policies for Tenants

You can configure up to 10 retention buckets for shared data, and up to 10 retention buckets for each tenant. The default retention period is 30 days; then, the tenant data is automatically deleted. To keep tenant data for longer than 30 days, you must configure a retention bucket. Until you configure a retention bucket, all events or flows are stored in the default retention bucket for each tenant.

If your JSA deployment has more than 10 tenants, you can configure a shared data retention policy and use the domain filter to create a domain-based retention policy for each of the domains within the tenant. Adding the domains specifies that the policy applies only to the data for that tenant.

## RELATED DOCUMENTATION

Monitoring License Usage in Multitenant Deployments | 378

Rules Management in Multitenant Deployments | 382

Network Hierarchy Updates in a Multitenant Deployment | 384

18

CHAPTER

# Asset Management

# Asset Management

Assets and asset profiles that are created for servers and hosts in your network provide important information to assist you in resolving security issues. Using the asset data, you can connect offenses that are triggered in your system to physical or virtual assets to provide a starting point in a security investigation.

The **Assets** tab in JSA provides a unified view of the known information about the assets in your network. As JSA discovers more information, the system updates the asset profile and incrementally builds a complete picture about the asset.

Asset profiles are built dynamically from identity information that is passively absorbed from event or flow data, or from data that JSA actively looks for during a vulnerability scan. You can also import asset data or edit the asset profile manually. For more information, see the topics *Importing Asset Profiles* and *Adding or editing an asset profile* in the *Juniper Secure Analytics Users Guide*.

**NOTE**: Log Manager tracks only asset data if JSA Vulnerability Manager is installed. For more information about the differences between JSA and Log Manager, see "Capabilities in Your JSA Product" on page 10.

RELATED DOCUMENTATION

Sources Of Asset Data | 387
Incoming Asset Data Workflow | 389
Updates to Asset Data | 391

# Sources Of Asset Data

IN THIS SECTION

Asset data is received from several different sources in your JSA deployment.

Asset data is written to the asset database incrementally, usually 2 or 3 pieces of data at a time. With exception of updates from network vulnerability scanners, each asset update contains information about only one asset at a time.

Asset data usually comes from one of the following asset data sources:

- **Events** -- Event payloads, such as those created by DHCP or authentication servers, often contain user logins, IP addresses, host names, MAC addresses, and other asset information. This data is immediately provided to the asset database to help determine which asset the asset update applies to.

  Events are the primary cause for asset growth deviations.

- **Flows** -- Flow payloads contain communication information such as IP address, port, and protocol that is collected over regular, configurable intervals. At the end of each interval, the data is provided to the asset database, one IP address at a time.

  Because asset data from flows is paired with an asset based on a single identifier, the IP address, flow data is never the cause of asset growth deviations.

  **NOTE**: Asset generation from IPv6 flows is not supported.

- **Vulnerability scanners** -- JSA integrates with both Juniper Networks and third-party vulnerability scanners that can provide asset data such as operating system, installed software, and patch information. The type of data varies from scanner to scanner and can vary from scan to scan. As new assets, port information, and vulnerabilities are discovered, data is brought into the asset profile based on the CIDR ranges that are defined in the scan.

  It is possible for scanners to introduce asset growth deviations but it is rare.

- **User interface** -- Users who have the Assets role can import or provide asset information directly to the asset database. Asset updates that are provided directly by a user are for a specific asset. Therefore the asset reconciliation stage is bypassed.

  Asset updates that are provided by users do not introduce asset growth deviations.

## Domain-aware Asset Data

When an asset data source is configured with domain information, all asset data that comes from that data source is automatically tagged with the same domain. Because the data in the asset model is domain-aware, the domain information is applied to all JSA components, including identities, offenses, asset profiles, and server discovery.

When you view the asset profile, some fields might be blank. Blank fields exist when the system did not receive this information in an asset update, or the information exceeded the asset retention period. The default retention period is 120 days. An IP address that appears as 0.0.0.0 indicates that the asset does not contain IP address information.

RELATED DOCUMENTATION

Incoming Asset Data Workflow | 389

Updates to Asset Data | 391

Identification Of Asset Growth Deviations | 395

# Incoming Asset Data Workflow

JSA uses identity information in an event payload to determine whether to create a new asset or update an existing asset.

**NOTE**: Asset generation from IPv6 flows is not supported.

**Figure 21: Asset Data Workflow Diagram**

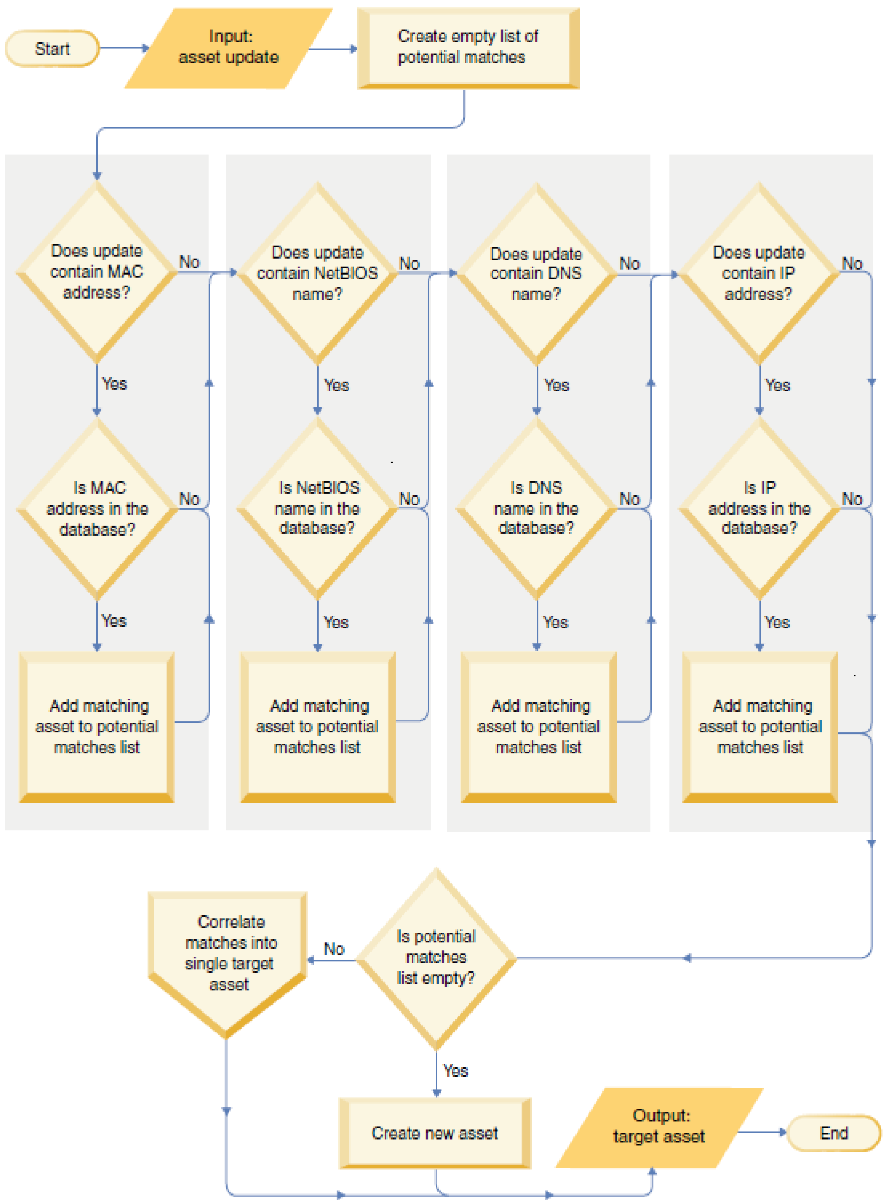

1. JSA receives the event. The asset profiler examines the event payload for identity information.

2. If the identity information includes a MAC address, a NetBIOS host name, or a DNS host name that are already associated with an asset in the asset database, then that asset is updated with any new information.

3. If the only available identity information is an IP address, the system reconciles the update to the existing asset that has the same IP address.

4. If an asset update has an IP address that matches an existing asset but the other identity information does not match, the system uses other information to rule out a false-positive match before the existing asset is updated.

5. If the identity information does not match an existing asset in the database, then a new asset is created based on the information in the event payload.

RELATED DOCUMENTATION

Updates to Asset Data | 391
Identification Of Asset Growth Deviations | 395
Prevention Of Asset Growth Deviations | 399

# Updates to Asset Data

IN THIS SECTION

JSA uses identity information in an event payload to determine whether to create a new asset or update an existing asset.

Each asset update must contain trusted information about a single asset. When JSA receives an asset update, the system determines which asset to which the update applies.

*Asset reconciliation* is the process of determining the relationship between asset updates and the related asset in the asset database. Asset reconciliation occurs after JSA receives the update but before the information is written to the asset database.

## Identity Information

Every asset must contain at least one piece of identity data. Subsequent updates that contain one or more pieces of that same identity data are reconciled with the asset that owns that data. Updates that are based on IP addresses are handled carefully to avoid false-positive asset matches. False positive asset matches occur when one physical asset is assigned ownership of an IP address that was previously owned by another asset in the system.

When multiple pieces of identity data are provided, the asset profiler prioritizes the information from the most deterministic to the least in the following order:

- MAC address
- NetBIOS host name
- DNS host name
- IP address

MAC addresses, NetBIOS host names, and DNS host names are unique and therefore are considered as definitive identity data. Incoming updates that match an existing asset only by the IP address are handled differently than updates that match more definitive identity data.

## Asset Reconciliation Exclusion Rules

With each asset update that enters JSA, the asset reconciliation exclusion rules apply tests to the MAC address, NetBIOS host name, DNS host name, and IP address in the asset update.

By default, each piece of asset data is tracked over a two-hour period. If any one piece of identity data in the asset update exhibits suspicious behavior two or more times within 2 hours, that piece of data is added to the asset blacklists. Each type of identity asset data that is tested results in a new blacklist.

**TIP**: JSA excludes events based on data that is received in the event, not on any data that is later inferred or linked to the event.

In domain-aware environments, the asset reconciliation exclusion rules track the behavior of asset data separately for each domain.

The asset reconciliation exclusion rules test the following scenarios:

**Table 59: Rule Tests and Responses**

| Scenario | Rule response |
|---|---|
| When a MAC address is associated to three or more different IP addresses in 2 hours or less | Add the MAC address to the Asset Reconciliation Domain MAC blacklist |
| When a DNS host name is associated to three or more different IP addresses in 2 hours or less | Add the DNS host name to the Asset Reconciliation Domain DNS blacklist |
| When a NetBIOS host name is associated to three or more different IP addresses in 2 hours or less | Add the NetBIOS host name to the Asset Reconciliation Domain NetBIOS blacklist |
| When an IPv4 address is associated to three or more different MAC addresses in 2 hours or less | Add the IP address to the Asset Reconciliation Domain IPv4 blacklist |
| When a NetBIOS host name is associated to three or more different MAC addresses in 2 hours or less | Add the NetBIOS host name to the Asset Reconciliation Domain NetBIOS blacklist |
| When a DNS host name is associated to three or more different MAC addresses in 2 hours or less | Add the DNS host name to the Asset Reconciliation Domain DNS blacklist |
| When an IPv4 address is associated to three or more different DNS host names in 2 hours or less | Add the IP address to the Asset Reconciliation Domain IPv4 blacklist |
| When a NetBIOS host name is associated to three or more different DNS host names in 2 hours or less | Add the NetBIOS host name to the Asset Reconciliation Domain NetBIOS blacklist |
| When a MAC address is associated to three or more different DNS host names in 2 hours or less | Add the MAC address to the Asset Reconciliation Domain MAC blacklist |
| When an IPv4 address is associated to three or more different NetBIOS host names in 2 hours or less | Add the IP address to the Asset Reconciliation Domain IPv4 blacklist |
| When a DNS host name is associated to three or more different NetBIOS host names in 2 hours or less | Add the DNS host name to the Asset Reconciliation Domain DNS blacklist |

**Table 59: Rule Tests and Responses *(Continued)***

| Scenario | Rule response |
|---|---|
| When a MAC address is associated to three or more different NetBIOS host names in 2 hours or less | Add the MAC address to the Asset Reconciliation Domain MAC blacklist |

You can view these rules on the **Offenses** tab by clicking **Rules** and then selecting the **asset reconciliation exclusion** group in the drop-down list.

## Asset Merging

*Asset merging* is the process where the information for one asset is combined with the information for another asset under the premise that they are actually the same physical asset.

Asset merging occurs when an asset update contains identity data that matches two different asset profiles. For example, a single update that contains a NetBIOS host name that matches one asset profile and a MAC address that matches a different asset profile might trigger an asset merge.

Some systems can cause high volumes of asset merging because they have asset data sources that inadvertently combine identity information from two different physical assets into a single asset update. Some examples of these systems include the following environments:

- Central syslog servers that act as an event proxy
- Virtual machines
- Automated installation environments
- Non-unique host names, common with assets like iPads and iPhones.
- Virtual private networks that have shared MAC addresses
- Log source extensions where the identity field is `OverrideAndAlwaysSend=true`

Assets that have many IP addresses, MAC addresses, or host names show deviations in asset growth and can trigger system notifications.

### RELATED DOCUMENTATION

Identification Of Asset Growth Deviations | 395

Prevention Of Asset Growth Deviations | 399

# Identification Of Asset Growth Deviations

IN THIS SECTION

Sometimes, asset data sources produce updates that JSA cannot handle properly without manual remediation. Depending on the cause of the abnormal asset growth, you can either fix the asset data source that is causing the problem or you can block asset updates that come from that data source.

*Asset growth deviations* occur when the number of asset updates for a single device grows beyond the limit that is set by the retention threshold for a specific type of the identity information. Proper handling of asset growth deviations is critical to maintaining an accurate asset model.

At the root of every asset growth deviation is an asset data source whose data is untrustworthy for updating the asset model. When a potential asset growth deviation is identified, you must look at the source of the information to determine whether there is a reasonable explanation for the asset to accumulate large amounts of identity data. The cause of an asset growth deviation is specific to an environment.

## DHCP Server Example Of Unnatural Asset Growth in an Asset Profile

Consider a virtual private network (VPN) server in a Dynamic Host Configuration Protocol (DHCP) network. The VPN server is configured to assign IP addresses to incoming VPN clients by proxying DHCP requests on behalf of the client to the network's DHCP server.

From the perspective of the DHCP server, the same MAC address repeatedly requests many IP address assignments. In the context of network operations, the VPN server is delegating the IP addresses to the clients, but the DHCP server can't distinguish when a request is made by one asset on behalf of another.

The DHCP server log, which is configured as a JSA log source, generates a DHCP acknowledgment (DHCP ACK) event that associates the MAC address of the VPN server with the IP address that it assigned to the VPN client. When asset reconciliation occurs, the system reconciles this event by MAC address, which results in a single existing asset that grows by one IP address for every DHCP ACK event that is parsed.

Eventually, one asset profile contains every IP address that was allocated to the VPN server. This asset growth deviation is caused by asset updates that contain information about more than one asset.

## Threshold Settings

When an asset in the database reaches a specific number of properties, such as multiple IP addresses or MAC addresses, JSA blocks that asset from receiving more updates.

The Asset Profiler threshold settings specify the conditions under which an asset is blocked from updates. The asset is updated normally up to the threshold value. When the system collects enough data to exceed the threshold, the asset shows an asset growth deviation. Future updates to the asset are blocked until the growth deviation is rectified.

## System Notifications That Indicate Asset Growth Deviations

JSA generates system notifications to help you identify and manage the asset growth deviations in your environment.

The following system messages indicate that JSA identified potential asset growth deviations:

- `The system detected asset profiles that exceed the normal size threshold`
- `The asset blacklist rules have added new asset data to the asset blacklists`

The system notification messages include links to reports to help you identify the assets that have growth deviations.

### Asset Data That Changes Frequently

Asset growth can be caused by large volumes of asset data that changes legitimately, such as in these situations:

- A mobile device that travels from office-to-office frequently and is assigned a new IP address whenever it logs in.
- A device that connects to a public wifi with short IP addresses leases, such as at a university campus, might collect large volumes of asset data over a semester.

## Example: How Configuration Errors for Log Source Extensions Can Cause Asset Growth Deviations

Customized log source extensions that are improperly configured can cause asset growth deviations.

You configure a customized log source extension to provide asset updates to JSA by parsing user names from the event payload that is on a central log server. You configure the log source extension to override the event host name property so that the asset updates that are generated by the custom log source always specify the DNS host name of the central log server.

Instead of JSA receiving an update that has the host name of the asset that the user logged in to, the log source generates many asset updates that all have the same host name.

In this situation, the asset growth deviation is caused by one asset profile that contains many IP addresses and user names.

## Troubleshooting Asset Profiles That Exceed the Normal Size Threshold

JSA generates the following system notification when the accumulation of data under a single asset exceeds the configured threshold limits for identity data.

**The system detected asset profiles that exceed the normal size threshold**

### Explanation

The payload of the notification shows a list of the top five most frequently deviating assets and why the system marked each asset as a growth deviation. As shown in the following example, the payload also shows the number of times that the asset attempted to grow beyond the asset size threshold.

```
Feb 13 20:13:23 127.0.0.1 [AssetProfilerLogTimer]
com.q1labs.assetprofile.updateresolution.UpdateResolutionManager:
[INFO] [NOT:0010006101][192.0.2.83/- -] [-/- -]
The top five most frequently deviating asset profiles between
```

```
Feb 13, 2015 8:10:23 PM AST and Feb 13, 2015 8:13:23 PM AST:
[ASSET ID:1003, REASON:Too Many IPs, COUNT:508],
[ASSET ID:1002, REASON:Too many DNS Names, COUNT:93],
[ASSET ID:1001, REASON:Too many MAC Addresses, COUNT:62]
```

When the asset data exceeds the configured threshold, JSA blocks the asset from future updates. This intervention prevents the system from receiving more corrupted data and mitigates the performance impacts that might occur if the system attempts to reconcile incoming updates against an abnormally large asset profile.

### Required User Action

Use the information in the notification payload to identify the assets that are contributing to the asset growth deviation and determine what is causing the abnormal growth. The notification provides a link to a report of all assets that experienced deviating asset growth over the past 24 hours.

After you resolve the asset growth deviation in your environment, you can run the report again.

1. Click the **Log Activity** tab and click **Search >New Search**.
2. Select the **Deviating Asset Growth: Asset Report** saved search.
3. Use the report to identify and repair inaccurate asset data that was created during the deviation.

## New Asset Data is Added to the Asset Blocklists

JSA generates the following system notification when a piece of asset data exhibits behavior that is consistent with deviating asset growth.

**The asset blacklist rules have added new asset data to the asset blacklists**

### Explanation

Asset exclusion rules monitor asset data for consistency and integrity. The rules track specific pieces of asset data over time to ensure that they are consistently being observed with the same subset of data within a reasonable time.

For example, if an asset update includes both a MAC address and a DNS host name, the MAC address is associated with that DNS host name for a sustained period. Subsequent asset updates that contain that MAC address also contain that same DNS host name when one is included in the asset update. If the MAC address suddenly is associated with a different DNS host name for a short period, the change is monitored. If the MAC address changes again within a short period, the MAC address is flagged as contributing to an instance of deviating or abnormal asset growth.

### Required User Action

Use the information in the notification payload to identify the rules that are used to monitor asset data. Click the **Asset deviations by log source** link in the notification to see the asset deviations that occurred in the last 24 hours.

If the asset data is valid, JSA administrators can configure JSA to resolve the problem.

- If your blocklists are populating too aggressively, you can tune the asset reconciliation exclusion rules that populate them.
- If you want to add the data to the asset database, you can remove the asset data from the blocklist and add it to the corresponding asset allowlist. Adding asset data to the whitelist prevents it from inadvertently reappearing on the blocklist.

RELATED DOCUMENTATION

Prevention Of Asset Growth Deviations | 399
Clean Up Asset Data After Growth Deviations | 414
Updates to Asset Data | 391

# Prevention Of Asset Growth Deviations

IN THIS SECTION

After you confirm that the reported asset growth is legitimate, there are several ways to prevent JSA from triggering growth deviation messages for that asset.

Use the following list to help you decide how to prevent asset growth deviations:

- "Stale Asset Data" on page 400
- "Tuning the Asset Profiler Retention Settings" on page 406
- "Tuning the Number Of IP Addresses Allowed for a Single Asset" on page 407
- "Identity Exclusion Searches" on page 409
- "Advanced Tuning Of Asset Reconciliation Exclusion Rules" on page 411
- Create asset allowlists to prevent data from reappearing on the asset blocklists.
- Modify the entries on the asset blocklists and asset allowlists.
- Ensure that your DSMs are up to date. JSA provides a weekly automatic update that might contain DSM updates and corrections to parsing issues.

## Stale Asset Data

Stale asset data can be problematic when the rate at which new asset records are created exceeds the rate at which stale asset data is removed. Controlling and managing asset retention thresholds is the key to addressing asset growth deviations that are caused by stale asset data.

*Stale asset data* is historical asset data that is not actively or passively observed within a specific time. Stale asset data is deleted when it exceeds the configured retention period.

The historical records become active again if they are observed by JSA passively, through events and flows, or actively, through port and vulnerability scanners.

Preventing asset growth deviations requires finding the right balance between the number of IP addresses allowed for a single asset and the length of time that JSA retains the asset data. You must consider the performance and manageability trade-offs before you configure JSA to accommodate high levels of asset data retention. While longer retention periods and higher per-asset thresholds might appear desirable all the time, a better approach is to determine a baseline configuration that is acceptable for your environment and test that configuration. Then, you can increase the retention thresholds in small increments until the right balance is achieved.

# Asset Blocklists and Allowlists

JSA uses a group of asset reconciliation rules to determine if asset data is trustworthy. When asset data is questionable, JSA uses asset blocklists and allowlists to determine whether to update the asset profiles with the asset data.

An *asset blocklist* is a collection of data that JSA considers untrustworthy. Data in the asset blocklist is likely to contribute to asset growth deviations and JSA prevents the data from being added to the asset database.

An *asset allowlist* is a collection of asset data that overrides the asset reconciliation engine logic about which data is added to an asset blocklist. When the system identifies a blocklist match, it checks the allowlist to see whether the value exists. If the asset update matches data that is on the allowlist, the change is reconciled and the asset is updated. Allowlisted asset data is applied globally for all domains.

The asset blocklists and allowlists are reference sets. You can view and modify the asset blocklist and allowlist data using the **Reference Set Management** tool in the JSA console. For more information about working with reference sets, see "Reference Sets Overview" on page 246.

Alternatively, you can use the command line interface (CLI) or the RestFUL API endpoint to update the content of the asset blocklists and allowlists.

## Asset Blocklists

An *asset blocklist* is a collection of data that JSA considers untrustworthy based on the asset reconciliation exclusion rules. Data in the asset blocklist is likely to contribute to asset growth deviations and JSA prevents the data from being added to the asset database.

Every asset update in JSA is compared to the asset blocklists. Blocklisted asset data is applied globally for all domains. If the asset update contains identity information (MAC address, NetBIOS host name, DNS host name, or IP address) that is found on a blocklist, the incoming update is discarded and the asset database is not updated.

The following table shows the reference collection name and type for each type of identity asset data.

Table 60: Reference Collection Names for Asset Blocklist Data

| Type of identity data | Reference collection name | Reference collection type |
|---|---|---|
| IP addresses (v4) | Asset Reconciliation IPv4 Blacklist | Reference Set [Set Type: IP] |
| DNS host names | Asset Reconciliation DNS Blacklist | Reference Set [Set Type: ALNIC*] |
| NetBIOS host names | Asset Reconciliation NetBIOS Blacklist | Reference Set [Set Type: ALNIC*] |
| MAC Addresses | Asset Reconciliation MAC Blacklist | Reference Set [Set Type: ALNIC*] |

ALNIC is an alphanumeric type that can accommodate both host name and MAC address values.

You can use the **Reference Set Management** tool to edit the blocklist entries. For information about working with reference sets, see *Juniper Secure Analytics Administration Guide*.

Your JSA administrator can modify the blocklist entries to ensure that new asset data is handled correctly.

## Asset Allowlists

You can use asset allowlists to keep JSA asset data from inadvertently reappearing in the asset blocklists.

An *asset allowlists* is a collection of asset data that overrides the asset reconciliation engine logic about which data is added to an asset blacklist. When the system identifies a blocklist match, it checks the allowlists to see whether the value exists. If the asset update matches data that is on the allowlists, the change is reconciled and the asset is updated. allowlisted asset data is applied globally for all domains.

You can use the **Reference Set Management** tool to edit the allowlist entries. For information about working with reference sets, see *Juniper Secure Analytics Administration Guide*.

Your JSA administrator can modify the allowlist entries to ensure that new asset data is handled correctly.

## Example of an Allowlist Use Case

The allowlist is helpful if you have asset data that continues to show up in the blocklists when it is a valid asset update. For example, you might have a round robin DNS load balancer that is configured to rotate across a set of five IP addresses. The Asset Reconciliation Exclusion rules might determine that the multiple IP addresses associated with the same DNS host name are indicative of an asset growth deviation, and the system might add the DNS load balancer to the blocklist. To resolve this problem, you can add the DNS host name to the **Asset Reconciliation DNS Whitelist**.

### Mass Entries to the Asset Allowlist

An accurate asset database makes it easier to connect offenses that are triggered in your system to physical or virtual assets in your network. Ignoring asset deviations by adding mass entries to the asset allowlist is not helpful in building an accurate asset database. Instead of adding mass allowlist entries, review the asset blocklist to determine what is contributing to the deviating asset growth and then determine how to fix it.

### Types Of Asset Allowlists

Each type of identity data is kept in a separate allowlist. The following table shows the reference collection name and type for each type of identity asset data.

**Table 61: Reference Collection Name for Asset Allowlist Data**

| Type of data | Reference collection name | Reference collection type |
|---|---|---|
| IP addresses | **Asset Reconciliation IPv4 Whitelist** | Reference Set [Set Type: IP] |
| DNS host names | **Asset Reconciliation DNS Whitelist** | Reference Set [Set Type: ALNIC*] |
| NetBIOS host names | **Asset Reconciliation NetBIOS Whitelist** | Reference Set [Set Type: ALNIC*] |
| MAC addresses | **Asset Reconciliation MAC Whitelist** | Reference Set [Set Type: ALNIC*] |

* ALNIC is an alphanumeric type that can accommodate host name and MAC address values.

## Updating the Asset Blocklists and Allowlists by Using Reference Set Utility

You can use the JSA reference set utility to add or modify the entries that are on the asset blocklists or allowlists.

To manage your reference sets, run the `ReferenceDataUtil.sh` utility from **/opt/qradar/bin** on the JSA console.

The commands to add new values to each list are described in the following table. The parameter values must exactly match the asset update values that are provided by the originating asset data source.

**Table 62: Command Syntax to Modify Asset Blocklist and Allowlist Data**

| Name | Command syntax |
|---|---|
| Asset Reconciliation IPv4 Blacklist | `ReferenceDataUtil.sh add "Asset Reconciliation IPv4 Blacklist" IP`<br>For example, this command adds IP address 192.168.3.56 to the blocklist:<br>`ReferenceDataUtil.sh add "Asset Reconciliation IPv4 Blacklist" 192.168.3.56` |
| Asset Reconciliation DNS Blacklist | `ReferenceDataUtil.sh add "Asset Reconciliation DNS Blacklist" DNS`<br>For example, this command adds domain name 'misbehaving.asset.company.com' to the blocklist:<br>`ReferenceDataUtil.sh add "Asset Reconciliation DNS Blacklist" "misbehaving.asset.company.com"` |
| Asset Reconciliation NetBIOS Blacklist | `ReferenceDataUtil.sh add "Asset Reconciliation NetBIOS Blacklist" NETBIOS`<br>For example, this command removes NetBIOS host name 'deviantGrowthAsset-156384' from the blocklist:<br>`ReferenceDataUtil.sh delete "Asset Reconciliation NetBIOS Blacklist" "deviantGrowthAsset-156384"` |
| Asset Reconciliation MAC Blacklist | `ReferenceDataUtil.sh add "Asset Reconciliation MAC Blacklist" MACADDR`<br>For example, this command adds MAC address '00:a0:1a:2b:3c:4d' to the blocklist:<br>`ReferenceDataUtil.sh add "Asset Reconciliation MAC Blacklist" "00:a0:1a:2b:3c:4d"` |
| Asset Reconciliation IPv4 Whitelist | `ReferenceDataUtil.sh add "Asset Reconciliation IPv4 Whitelist" IP`<br>For example, this command deletes IP address 192.0.2.0 from the allowlist:<br>`ReferenceDataUtil.sh delete "Asset Reconciliation IPv4 Whitelist" 192.0.2.0` |

Table 62: Command Syntax to Modify Asset Blocklist and Allowlist Data *(Continued)*

| Name | Command syntax |
|---|---|
| Asset Reconciliation DNS Whitelist | `ReferenceDataUtil.sh add "Asset Reconciliation DNS Whitelist" DNS`<br>For example, this command adds domain name 'loadbalancer.company.com' to the allowlist:<br>`ReferenceDataUtil.sh add "Asset Reconciliation DNS Whitelist" "loadbalancer.company.com"` |
| Asset Reconciliation NetBIOS Whitelist | `ReferenceDataUtil.sh add "Asset Reconciliation NetBIOS Whitelist" NETBIOS`<br>For example, this command adds NetBIOS name 'assetName-156384' to the allowlist:<br>`ReferenceDataUtil.sh add "Asset Reconciliation NetBIOS Whitelist" "assetName-156384"` |
| Asset Reconciliation MAC Whitelist | `ReferenceDataUtil.sh add "Asset Reconciliation MAC Whitelist" MACADDR`<br>For example, this command adds MAC address '00:a0:1a:2b:3c:4d' to the allowlist :<br>`ReferenceDataUtil.sh add "Asset Reconciliation MAC Whitelist" "00:a0:1a:2b:3c:4d"` |

## Updating the Blocklists and Allowlists Using the RESTful API

You can use the JSA RESTful API to customize the content of the asset blocklists and allowlists.

You must specify the exact name of the reference set that you want to view or update.

- **Asset Reconciliation IPv4 Blacklist**
- **Asset Reconciliation DNS Blacklist**
- **Asset Reconciliation NetBIOS Blacklist**
- **Asset Reconciliation MAC Blacklist**
- **Asset Reconciliation IPv4 Whitelist**
- **Asset Reconciliation DNS Whitelist**
- **Asset Reconciliation NetBIOS Whitelist**
- **Asset Reconciliation MAC Whitelist**

1. Type the following URL in your web browser to access the RESTful API interface:

https://*ConsoleIPaddress*/api_doc

2. In the navigation pane on the left, find **4.0>/reference_data >/sets > /{name}**.

3. To view the contents of an asset blocklist or allowlist, follow these steps:

   a. Click the **GET** tab and scroll down to the **Parameters** section.

   b. In the **Value** field for the **Name** parameter, type the name of the asset blocklist or allowlist that you want to view.

   c. Click **Try It Out** and view the results at the bottom of the screen.

4. To add a value to an asset blocklist or allowlist, follow these steps:

   a. Click the **POST** tab and scroll down to the **Parameters** section.

   b. Type in the values for the following parameters:

   **Table 63: Parameters That Are Required to Add New Asset Data**

   | Parameter name | Parameter description |
   |---|---|
   | name | Represents the name of the reference collection that you want to update. |
   | value | Represents the data item that you want to add to the asset blocklist or allowlist. Must exactly match the asset update values that are provided by the originating asset data source. |

   c. Click **Try It Out** to add the new value to the asset allowlist or asset blocklist.

For more information about using the RESTful API to change the reference sets, see the *Juniper Secure Analytics API Guide*.

## Tuning the Asset Profiler Retention Settings

JSA uses the asset retention settings to manage the size of the asset profiles.

The default retention period for most asset data is 120 days after the last time it was either passively or actively observed in JSA. User names are retained for 30 days.

Asset data that is added manually by JSA users does not usually contribute to asset growth deviations. By default, this data is retained forever. For all other types of asset data, the **Retain Forever** flag is suggested only for static environments.

You can adjust the retention time based on the type of asset identity data that is in the event. For example, if multiple IP addresses are merging under one asset, you can change the Asset IP Retention period from 120 days to a lower value.

When you change the asset retention period for a specific type of asset data, the new retention period is applied to all asset data in JSA. Existing asset data that already exceeds the new threshold is removed when the deployment is complete. To ensure that you can always identify named hosts even when the asset data is beyond the retention period, the asset retention cleanup process does not remove the last known host name value for an asset.

Before you determine how many days that you want to retain the asset data, understand the following characteristics about longer retention periods:

- provides a better historical view of your assets.
- creates larger data volumes per asset in the asset database.
- increases the probability that stale data will contribute to asset growth deviation messages.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Asset Profiler Configuration**.
3. Click **Asset Profiler Retention Configuration**.
4. Adjust the retention values and click **Save**.
5. Deploy the changes into your environment for the updates to take effect.

## Tuning the Number Of IP Addresses Allowed for a Single Asset

JSA monitors the number of IP addresses that a single asset accumulates over time.

By default, JSA generates a system message when a single asset accumulates more than 75 IP addresses. If you expect assets to accumulate more than 75 IP addresses, you can tune the **Number of IPs Allowed for a Single Asset** value to avoid future system messages.

Setting the limit for the number of IP addresses too high prevents JSA from detecting asset growth deviations before they have a negative impact on the rest of the deployment. Setting the limit too low increases the number of asset growth deviations that are reported.

You can use the following guideline when you tune the **Number of IPs Allowed for a Single Asset** setting for the first time.

Number of IP addresses that are allowed for a single asset = (*<retention time (days)>* x *<estimated IP addresses per day>*) + *<buffer number of IP addresses>*

Where

- *<estimated IP addresses per day>* is the number of IP addresses that a single asset might accumulate in one day under normal conditions

- *<retention time (days)>* is the preferred amount of time to retain the asset IP addresses

1. On the navigation menu
   (
   ≡
   ), click **Admin**.

2. In the **Assets** section, click **Asset Profiler Retention Configuration**.

3. Click **Asset Profiler Retention Configuration**.

4. Adjust the configuration values and click **Save**.

5. Deploy the changes into your environment for the updates to take effect.

## Tuning the Number of MAC Addresses Allowed for a Single Asset

JSA monitors the number of MAC addresses that a single asset accumulates over time.

By default, JSA generates a system message when a single asset accumulates more than ten IP addresses. If you expect assets to accumulate more than ten MAC addresses, you can tune the **Number of MAC Addresses Allowed for a Single Asset** value to avoid future system messages.

Setting the limit for the number of MAC addresses too high prevents JSA from detecting asset growth deviations before they have a negative impact on the rest of the deployment. Setting the limit too low increases the number of asset growth deviations that are reported.

You can use the following guideline when you tune the **Number of MAC Addresses Allowed for a Single Asset** setting for the first time.

Number of MAC addresses that are allowed for a single asset = *(<retention time (days)> x <estimated MAC addresses per day>) + <buffer number of MAC addresses>*

Where

- *<estimated MAC addresses per day>* is the number of MAC addresses that a single asset might accumulate in one day under normal conditions
- *<retention time (days)>is* the preferred amount of time to retain the asset MAC addresses

1. On the navigation menu, click **Admin**.
2. In the **Assets** section, click **Asset Profiler Configuration**.
3. Click **Asset Profiler Configuration**.
4. Adjust the **Number of MAC Addresses Allowed for a Single Asset** value and click **Save**.
5. Deploy the changes into your environment for the updates to take effect.

## Identity Exclusion Searches

Identity exclusion searches can be used to manage single assets that accumulate large volumes of similar identity information for known, valid reasons.

For example, log sources can provide large volumes of asset identity information to the asset database. They provide JSA with near real-time changes to asset information and they can keep your asset database current. But log sources are most often the source of asset growth deviations and other asset-related anomalies.

When a log source sends incorrect asset data to JSA, try to fix the log source so that the data it sends is usable by the asset database. If the log source cannot be fixed, you can build an identity exclusion search that blocks the asset information from entering the asset database.

You can also use an identity exclusion search where `Identity_Username+Is Any Of + Anonymous Logon` to ensure that you are not updating assets that are related to service accounts or automated services.

### Differences Between Identity Exclusion Searches and Blacklists

While identity exclusion searches appear to have similar functionality to asset blacklists, there are significant differences.

Blacklists can specify only raw asset data, such as MAC addresses and host names, that is to be excluded. Identity exclusion searches filter out asset data based on search fields like log source, category, and event name.

Blacklists do not account for the type of data source that is providing the data, whereas identity exclusion searches can be applied to events only. Identity exclusion searches can block asset updates based on common event search fields, such as event type, event name, category, and log source.

## Creating Identity Exclusion Searches

To exclude certain events from providing asset data to the asset database, you can create a JSA identity exclusion search.

The filters that you create for the search must match events that you want to exclude, not the events that you want to keep.

You might find it helpful to run the search against events that are already in the system. However, when you save the search, you must select **Real Time (streaming)** in the **Timespan** options. If you do not choose this setting, the search will not match any results when it runs against the live stream of events that are coming into JSA.

When you update the saved identity exclusion search without changing the name, the identity exclusion list that is used by the Asset Profiler is updated. For example, you might edit the search to add more filtering of the asset data that you want to exclude. The new values are included and the asset exclusion starts immediately after the search is saved.

1. Create a search to identify the events that do not provide asset data to the asset database.

   a. On the **Log Activity** tab, click **Search >New Search**.

   b. Create the search by adding search criteria and filters to match the events that you want to exclude from asset updates.

   c. In the **Time Range** box, select **Real Time (streaming)** and then click **Filter** to run the search.

   d. On the search results screen, click **Save Criteria** and provide the information for the saved search.

   > **NOTE**: You can assign the saved search to a search group. An Identity Exclusion search group exists in the **Authentication, Identity and User Activity** folder.

   e. Click **OK** to save the search.

2. Identify the search that you created as an identity exclusion search.

   a. On the navigation menu
   (
   ≡
   ), click **Admin**.

b. In the **System Configuration** section, click **Asset Profiler Configuration**.

c. Click **Manage Identity Exclusion** at the bottom of the screen.

d. Select the identity exclusion search that you created from the list of searches on the left and click the add icon (>).

> **TIP**: If you can't find the search, type the first few letters into the filter at the top of the list.

e. Click **Save**.

3. On the **Admin** tab, click **Deploy changes** for the updates to take effect.

## Advanced Tuning Of Asset Reconciliation Exclusion Rules

You can tune the Asset Reconciliation Exclusion rules to refine the definition of deviating asset growth in one or more of the rules.

For example, consider this normalized template from an Asset Reconciliation Exclusion rule.

**Apply *AssetExclusion: Exclude DNS Name By IP* on events which are detected by the *Local* system *and NOT* when any of *Identity Host Name* are contained in any of *Asset Reconciliation DNS Whitelist - AlphaNumeric (Ignore Case), Asset Reconciliation DNS Blacklist - AlphaNumeric (Ignore Case)* and when at least *N1* events are seen with the same *Identity Host Name* and different *Identity IP* in *N2***

This table lists the variables in the rule template that can be tuned and the result of the change. Avoid changing other variables in the template.

**Table 64: Options for Tuning the Asset Reconciliation Rules**

| Variable | Default value | Tuning result |
|---|---|---|
| N1 | 3 | Tuning this variable to a lower value results in more data being added to the blacklist because fewer events with conflicting data are needed for the rule to fire.<br>Tuning this variable to a higher value results in less data being added to the blacklist because more events with conflicting data are needed for the rule to fire. |

**Table 64: Options for Tuning the Asset Reconciliation Rules *(Continued)***

| Variable | Default value | Tuning result |
|---|---|---|
| N2 | 2 hours | Tuning this variable to a lower value reduces the window of time in which N1 events must be seen for the rule to fire. The time required to observe matching data is decreased, which results in less data being added to the blacklist.<br>Tuning this variable to a higher value increases the time in which N1 events must be seen for the rule to fire. The time to observe matching data is increased, which results in more data being added to the blacklist.<br>Increasing the time period might impact system memory resources as data is tracked over longer periods of time. |

The Asset Reconciliation Exclusion rules are system-wide rules. Changes to the rules affect the way that the rule behaves throughout the entire system.

## Applying Different Tuning for Rules

It might be necessary to apply different tuning for rules in different parts of the system. To apply different tuning for rules, you must duplicate the Asset Reconciliation Exclusion rules that you want to tune and add one or more tests to constrain the rules so that you test only certain parts of the system. For example, you might want to create rules that test only networks, log sources, or event types.

Always be cautious when you are adding new rules to the system because as some tasks and CRE rules might impact system performance. It might be beneficial to add the new rules to the top of each test stack to allow the system to bypass the remainder of the test logic whenever an asset update matches the criteria for the new rule.

1. Duplicate the rule.

   a. On the **Offenses** tab, click **Rules** and select the rule that you want to copy.

   b. Click **Actions** >**Duplicate**.

      It can be helpful if the name of the new rule is indicative of the reason for duplicating it.

2. Add a test to the rule.

   Determine a filter that you want to use to apply the rule only to a subset of system data. For example, you can add a test that matches only events from a specific log source.

3. Tune the variables of the rule to achieve the wanted behavior.

4. Update the original rule.

a. Add the same test that you added to the duplicate rule to the original rule, but this time invert the rules `AND` and `AND NOT` operators.

   Inverting the operators prevents events from being triggered in both rules.

## Example: Asset Exclusion Rules That Are Tuned to Exclude IP Addresses from the Blacklist

You can exclude IP addresses from being blacklisted by tuning the asset exclusion rules.

As the Network security administrator, you manage a corporate network that includes a public wifi network segment where IP address leases are typically short and frequent. The assets on this segment of the network tend to be transient, primarily notebooks and hand-held devices that log in and out of the public wifi frequently. Commonly, a single IP address is used multiple times by different devices over a short time.

In the rest of your deployment, you have a carefully managed network that consists only of inventoried, well-named company devices. IP address leases are much longer in this part of the network, and IP addresses are accessed by authentication only. On this network segment, you want to know immediately when there are any asset growth deviations and you want to keep the default settings for the asset reconciliation exclusion rules.

### Blacklisting IP Addresses

In this environment, the default asset reconciliation exclusion rules inadvertently blacklist the entire network in a short time.

Your security team finds the asset-related notifications that are generated by the wifi segment are a nuisance. You want to prevent the wifi from triggering any more deviating asset growth notifications.

### Tuning Asset Reconciliation Rules to Ignore Some Asset Updates

You review the **Asset deviation by log source** report in the last system notification. You determine that the blacklisted data is coming from the DHCP server on your wifi.

The values in the **Event Count** column, **Flow Count** column and the **Offenses** column for the row corresponding to the **AssetExclusion: Exclude IP By MAC Address** rule indicate that your wifi DHCP server is triggering this rule.

You add a test to the existing asset reconciliation exclusion rules to stop rules from adding wifi data to the blacklist.

**Apply AssetExclusion:Exclude IP by MAC address on events which are detected by the Local system and NOT when the event(s) were detected by one or more of MicrosoftDHCP @ microsoft.dhcp.test.com and NOT when any of Domain is the key and any of Identity IP is the value in any of Asset Reconciliation Domain IPv4 Whitelist - IP Asset Reconciliation Domain IPv4 Blacklist - IP and when at least 3 events are seen with the same Identity IP and different Identity MAC in 2 hours.**

The updated rule tests only the events from the log sources that are not on your wifi DHCP server. To prevent wifi DHCP events from undergoing more expensive reference set and behavior analysis tests, you also moved this test to the top of the test stack.

RELATED DOCUMENTATION

Clean Up Asset Data After Growth Deviations | **414**

Updates to Asset Data | **391**

Identification Of Asset Growth Deviations | **395**

# Clean Up Asset Data After Growth Deviations

**IN THIS SECTION**

JSA uses the asset model to connect offenses in your deployment to physical or virtual assets in your network. The ability to collect and view relevant data on how assets are used is an important step in resolving security issues. It is important to maintain the asset database to ensure that the data is current and accurate.

Whether you fix the source of the problem or block the asset updates, you must clean up the asset database by removing the invalid asset data and removing the asset blocklist entries.

## Deleting Invalid Assets

After you fix the assets that contributed to the asset growth deviation, clean up your asset artifacts by using selective clean up or rebuilding the asset database.

- **Selective clean up**--This method is for asset growth deviations of limited scope. Selectively removing the affected assets is the least invasive way to clean up asset artifacts, but if many assets were affected, it can also be the most tedious.

- **Rebuild the asset database**--Rebuilding the asset database from scratch is the most efficient and precise method of deleting assets when asset growth deviations are pervasive.

  This method passively regenerates assets in your database based on the new tuning that you configured to resolve the asset growth issues. With this approach, all scan results and residual asset data are lost, but the data can be reclaimed by rerunning a scan or re-importing scan results.

1. To selectively remove invalid artifacts in the asset database, perform these steps:

   a. On the **Log Activity** tab, run the **Deviating Asset Growth: Asset Report** event search.

      This search returns a report of assets that are affected by deviating asset growth and must be deleted.

   b. On the **Assets** tab, click **Actions >Delete Asset**

      There might be a delay before the asset no longer appears in JSA.

2. To rebuild the asset database from scratch, perform these steps:

   a. Use SSH to log in to the JSA console as an administrator.

   b. Run the **/opt/qradar/support/cleanAssetModel.sh** script from the console command line and select **Option 1** when prompted.

   Rebuilding the asset database restarts the asset reconciliation engine.

Purging a blocklist removes all blocklist entries, including those entries that were added manually. Blocklist entries that were manually added must be added again.

## Deleting Blacklist Entries

After you fixed the cause of the blacklist entries, you must clean up the remnant entries. You can remove the individual blacklist entries, however it is better to purge all blacklist entries and allow the blacklist values that are unrelated to the asset growth deviation to regenerate.

1. To purge a blacklist by using the JSA Console:

    a. On the navigation menu (☰), click **Admin**.

    b. In the **System Configuration** section, click **Reference Set Management**.

    c. Select a reference set and then click **Delete**.

    d. Use the quick search text box to search for the reference sets that you want to delete, and then click **Delete Listed**.

2. To purge a blacklist by using the JSA console command-line interface:

    a. Change directory to **/opt/qradar/bin**.

    b. Run the following command.

    **./ReferenceDataUtil.sh purge "***Reference Collection Name***"**

    where *Reference Collection Name* is one of the following lists:

    - Asset Reconciliation NetBIOS Blacklist
    - Asset Reconciliation DNS Blacklist
    - Asset Reconciliation IPv4 Blacklist
    - Asset Reconciliation MAC Blacklist

Purging a blacklist removes all blacklist entries, including those entries that were added manually. Blacklist entries that were manually added must be added again.

## RELATED DOCUMENTATION

Updates to Asset Data | **391**

Identification Of Asset Growth Deviations | **395**

Prevention Of Asset Growth Deviations | **399**

19

CHAPTER

# Configuring JSA to Forward Data to Other Systems

# Forward Data to Other Systems

Configure JSA to forward data to one or more vendor systems, such as ticketing or alerting systems.

You can also forward normalized data to other JSA deployments. The target system that receives the data from JSA is known as a *forwarding destination*. JSA ensures that all forwarded data is unaltered.

**NOTE**: Forwarded normalized data must match or exist in both JSA deployments. Otherwise, the event might have an incorrect associated QID or remain unparsed. This data includes QIDS, custom log source types, custom properties, event ID, and event category expressions. To prevent synchronization issues, forward the events by using raw format.

To avoid compatibility problems when you send event and flow data, ensure that the deployment that receives the data is the same version or higher than the deployment that sends the data by using the following workflow.

1. Configure one or more forwarding destinations.
2. To determine what data you want to forward, configure routing rules, custom rules, or both.
3. Configure the routing options to apply to the data.

For example, you can configure all data from a specific event collector to forward to a specific ticketing system. You can also bypass correlation by removing the data that matches a routing rule.

# Adding Forwarding Destinations

Before you can configure bulk or selective data forwarding, you must add a forwarding destination. Normalized events that you forward can be interpreted only by other JSA systems.

**NOTE**: You cannot forward data to systems that use dynamic IP addresses. The connection is established when the service starts, and changes to the IP address are not detected until the service restarts. The forwarding destination must have a static IP address.

1. On the navigation menu
   (

☰
), click **Admin**.

2. In the **System Configuration** section, click **Forwarding Destinations**.
3. On the toolbar, click **Add**.
4. In the **Forwarding Destinations** window, enter values for the parameters.

   The following table describes some of the **Forwarding Destinations** parameters.

**Table 65: Forwarding Destinations Parameters**

| Parameter | Description |
|---|---|
| Event Format | • **Payload** is the data in the format that the log source or flow source sent.<br>If you select this option, ensure that port 514 is open.<br>• **Normalized** is raw data that is parsed and prepared as readable information for the user interface. If you select this option, ensure that ports 32000 and 32004 are open.<br>• **JSON** (Javascript Object Notation) is a data-interchange format.<br>If you select this option, ensure that port 5141 is open. |
| Destination Address | The IP address or host name of the vendor system that you want to forward data to. |
| Protocol | Use the **TCP** protocol to send normalized data by using the TCP protocol. You must create an off-site source at the destination address on port 32004 for events, or on port 32000 for flows.<br>Use the TCP over SSL protocol to send normalized data securely by using the TCP protocol with an SSL certificate.<br>**NOTE**: You cannot transmit normalized and JSON data by using the UDP protocol. If you select the **Normalized or JSON** options, the **UDP** option in the Protocol list is disabled. |

Table 65: Forwarding Destinations Parameters *(Continued)*

| Parameter | Description |
|---|---|
| Prefix a syslog header if it is missing or invalid | Applicable only when the event format is **Payload**.<br>When JSA forwards syslog messages, the outbound message is verified to ensure that it has a valid syslog header.<br>If a valid syslog header is not detected on the original syslog message and this checkbox is selected, the prefixed syslog header includes the originating IP address from the packet that JSA received in the **Hostname** field of the syslog header. If this checkbox is not selected, the data is sent unmodified. |

5. Optional: If you are using the TCP over SSL protocol, do the following:

   a. From the command line of the event collector or processor that uses the routing rule to forward data, change the directory to **/tmp**.

   b. Run the following command: **/opt/qradar/bin/getcert.sh tlssyslog_server_iptlssyslog_port**

      A copy of the client certificate is downloaded from the target system and is titled with the IP and port you downloaded it from.

   c. Move the certificate to **/opt/qradar/conf/trusted_certificates/**

   d. If the certificate was signed by a commercial or private certificate authority (CA), copy the root CA and intermediate certificates to **/etc/pki/ca-trust/source/anchors**

   e. Run the following command: **update-ca-trust**

6. Click **Save**.

Setting up a forwarding destination does not automatically send data to that destination. You must configure either a routing rule or a custom rule to forward data to the destination.

## RELATED DOCUMENTATION

Configuring Forwarding Profiles | **421**

Configuring Routing Rules to Forward Data | **422**

Using Custom Rules and Rule Responses to Forward Data | **425**

# Configuring Forwarding Profiles

If you want to specify which properties to forward to the forwarding destination, configure a forwarding profile.

You must re-create JSON forwarding profiles that you created in JSA 2014.3 or earlier.

You can use forwarding profiles only when the event data is sent in JSON format.

You can select specific event or flow properties, including custom properties, to forward to an external destination. You can enhance the readability of the event data by specifying an alias name and default value for the attribute. Alias names and default values are specific to the profile they are defined in. If the attributes are used in other profiles, the alias names and default values must be redefined.

You can use a single profile that has multiple forwarding destinations. When you edit a profile, ensure that the changes are appropriate for all forwarding destinations that the profile is associated with.

When you delete a profile, all forwarding destinations that used the profile automatically revert to using the default profile.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Forwarding Destinations**.
3. On the toolbar, click **Profile Manager**.
4. To create a new profile, click **New**.
5. Type a name for the profile and select the check box beside the attributes that you want to include in the event data set.
6. To change an existing profile, select the profile and click **Edit** or **Delete**.
7. Click **Save**.

## RELATED DOCUMENTATION

Configuring Routing Rules to Forward Data | 422

Using Custom Rules and Rule Responses to Forward Data | 425

Viewing Forwarding Destinations | 427

# Configuring Routing Rules to Forward Data

Forward data by configuring filter-based routing rules.

You can configure routing rules to forward data in either online or offline mode:

- In **Online** mode, your data remains current because forwarding is done in real time. If the forwarding destination becomes unreachable, any data that is sent to that destination is not delivered, resulting in missing data on that remote system. To ensure that delivery is successful, use offline mode.
- In **Offline** mode, all data is first stored in the database and then sent to the forwarding destination. This mode ensures that no data is lost; however, delays in data forwarding can occur.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Routing Rules**.
3. On the toolbar, click **Add**.
4. In the Routing Rule window, type a name and description for your routing rule.
5. In the **Mode** field, select one of the following options: **Online** or **Offline**.
6. In the **Forwarding Event Collector** or **Forwarding Event Processor** list, select the event collector from which you want to forward data.

   **Learn more about the forwarding appliance:**

   - **Forwarding Event Collector -** Specifies the Event Collector that you want this routing rule to process data from. This option displays when you select the **Online** option.

     > **NOTE**: Online/Realtime forwarding is not impacted by any Rate Limit or Scheduling configurations that might be configured on a Store and Forward event collectors.

   - **Forwarding Event Processor -** Specifies the Event Processor that you want this routing rule to process data from. This option is displayed when you select the **Offline** option.

     > **NOTE**: This option is not available if **Drop** is selected from the **Routing Options** pane.

7. In the Data Source field, select which data source you want to route: **Events** or **Flows**.

   The labels for the next section change based on which data source you select.
8. Specify which events or flows to forward by applying filters:

a. To forward all incoming data, select the **Match All Incoming Events** or **Match All Incoming Flows** checkbox.

> **NOTE**: If you select this checkbox, you cannot add a filter.

b. To forward only some events or flows, specify the filter criteria, and then click **Add Filter**.

9. Specify the routing options to apply to the forwarded data:

   a. Optional: If you want to edit, add, or delete a forwarding destination, click the **Manage Destinations** link.

   b. To forward log data that matches the specified filters, select the **Forward** checkbox and then select the checkbox for each forwarding destination.

   > **NOTE**: If you select the **Forward** check box, you can select only one of these check boxes: **Drop**, **Bypass Correlation**, or **Log Only**.

10. Click **Save**.

### Routing options for rules

You can choose from four rule routing options: Forward, Drop, Bypass correlation, and Log Only. The following table describes the different options and how to use them.

**Table 66: Rule Routing Options**

| Routing type | Description |
|---|---|
| **Forward** | Data is forwarded to the specified forwarding destination. Data is also stored in the database and processed by the Custom Rules Engine (CRE). |
| **Drop** | Data is dropped. The data is not stored in the database and is not processed by the CRE. This option is not available if you select the **Offline** option. Any events that are dropped are credited back 100% to the license. |

Table 66: Rule Routing Options *(Continued)*

| Routing type | Description |
|---|---|
| **Bypass Correlation** | Data bypasses CRE, but it is stored in the database. This option is not available if you select the **Offline** option.<br>The **Bypass correlation** option does not require an entitlement for JSA Data Store. Bypass correlation allows events that are received in batches to bypass real-time rules. You can use the events in analytic apps and for historical correlation runs. For historical correlation runs, the events can be replayed as though they were received in real time. |
| **Log Only (Exclude Analytics)** | Events are stored and flagged in the database as **Log Only** and bypass CRE. These events are not available for historical correlation, and are credited back 100% to the license. This option is not available for flows or if you select the **Offline** option.<br>The **Log Only** option requires an entitlement for JSA Data Store. After the entitlement is purchased and the **Log Only** option is selected, events that match the routing rule are stored to disk and are available to view and for searches. The events bypass the custom rule engine and no real-time correlation or analytics occur. The events can't contribute to offenses and are ignored when historical correlation runs. |

The following table describes different routing option combinations that you can use. These options are not available in offline mode.

Table 67: Rule Routing Combination Options

| Routing combination | Description |
|---|---|
| **Forward** and **Drop** | Data is forwarded to the specified forwarding destination. Data is not stored in the database and is not processed by the CRE. Any events that are dropped are credited back 100% to the license. |
| **Forward** and **Bypass Correlation** | Data is forwarded to the specified forwarding destination. Data is stored in the database, but it is not processed by the CRE. |
| **Forward** and **Log Only (Exclude Analytics)** | Events are forwarded to the specified forwarding destination. Events are stored and flagged in the database as Log Only and bypass CRE. These events are not available for historical correlation, and are credited back 100% to the license. |

If data matches multiple rules, the safest routing option is applied. For example, if data that matches a rule that is configured to drop and a rule to bypass CRE processing, the data is not dropped. Instead, the data bypasses the CRE and is stored in the database.

RELATED DOCUMENTATION

Using Custom Rules and Rule Responses to Forward Data | 425

Viewing Forwarding Destinations | 427

Viewing and Managing Forwarding Destinations | 428

# Using Custom Rules and Rule Responses to Forward Data

Use the **Custom Rule** wizard to configure forwarding of event data that matches rules in your system. Configure the rule response to forward event data to one or more forwarding destinations.

The criteria that determines the event data that is sent to a forwarding destination is based on the tests and building blocks that are included in the rule.

When the rule is configured and enabled, all event data that matches the rule tests are automatically sent to the specified forwarding destinations. For more information about how to edit or add a rule, see the *Juniper Secure Analytics Users Guide* for your product.

1. Click the **Offenses Log Activity** tab.
2. On the **Rules** menu, select **Rules**.
3. In the Rules List window, select the rule to edit, or click **Actions** to create a new rule.
4. On the Rule Response page in the Rule wizard, ensure that you select the **Send to Forwarding Destinations** option.

RELATED DOCUMENTATION

Viewing Forwarding Destinations | 427

Viewing and Managing Forwarding Destinations | 428

Viewing and Managing Routing Rules | 429

# Configuring Routing Rules to Use the JSA Data Store

A new offering, JSA Data Store, normalizes and stores both security and operational log data for future analysis and review. The offering supports the storage of an unlimited number of logs without counting against your organization's Events Per Second JSA license, and enables your organization to build custom apps and reports based on this stored data to gain deeper insights into your environments.

Using the Log Only (Exclude Analytics) option requires entitlement for JSA Data Store, but is not currently enforced. In the future, when entitlement is enforced, access to the collected event data will be restricted to properly licensed systems. When the license is applied and the Log Only (Exclude Analytics) option is selected, events that match the routing rule will be stored to disk and will be available to view and for searches. The events bypass the custom rule engine and no real-time correlation or analytics occur. The events can't contribute to offenses and are ignored when historical correlation runs.

The following apps also ignore Log Only events:

- JSA User Behavior Analytics
- JSA Advisor with Watson

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Routing Rules**.
3. On the toolbar, click **Add**.
4. In the **Routing Rule** window, type a name and description for your routing rule.
5. In the **Mode** field, select **Online**.
6. In the **Forwarding Event Collector** list, select the event collector on which you want to apply the **Log Only (Exclude Analytics)** option.
7. In the **Data Source** field, select **Events**.
8. Specify which events to apply the **Log Only (Exclude Analytics)** option to by applying filters:

   a. To apply the **Log Only (Exclude Analytics)** option to all incoming data, select the **Match All Incoming Events** check box.

   > **NOTE**: If you select this check box, you cannot add a filter.

   b. To apply the **Log Only (Exclude Analytics)** option to only some events, specify the filter criteria, and then click **Add Filter**.

9. To apply the **Log Only (Exclude Analytics)** option to log data that matches the specified filters, select **Log Only (Exclude Analytics)**.

   > **NOTE**: The **Log Only (Exclude Analytics)** option specifies that events are stored and flagged in the database as Log Only and bypass CRE. These events are not available for historical correlation, and are credited back 100% to the license. This option is not available for flows.
   >
   > You can combine the **Forward** and **Log Only (Exclude Analytics)** options. Events are forwarded to the specified forwarding destination in online mode. Events are stored and flagged in the database as Log Only and bypass CRE. These events are not available for historical correlation, and are credited back 100% to the license. This option is not available in offline mode.
   >
   > If data matches multiple rules, the safest routing option is applied. For example, if data that matches a rule that is configured to drop and a rule to bypass CRE processing, the data is not dropped. Instead, the data bypasses the CRE and is stored in the database.

10. Click **Save**.

RELATED DOCUMENTATION

Using Custom Rules and Rule Responses to Forward Data | 425

Viewing Forwarding Destinations | 427

Viewing and Managing Forwarding Destinations | 428

# Viewing Forwarding Destinations

The Forwarding Destinations window provides valuable information about your forwarding destinations. Statistics for the data sent to each forwarding destination is displayed.

For example, you can see the following information:

- The total number events and flows that were seen for this forwarding destination.
- The number of events or flows that were sent to this forwarding destination.
- The number of events or flows that were dropped before the forwarding destination was reached.

1. On the navigation menu
   (

≡
), click **Admin**.

2. In the **System Configuration** section, click **Forwarding Destinations**.

   Statistics for the data sent to each forwarding destination is displayed. For example, you can see the following information:

   - The total number events and flows that were seen for this forwarding destination.
   - The number of events or flows that were sent to this forwarding destination.
   - The number of events or flows that were dropped before the forwarding destination was reached.
3. View the statistics for your forwarding destinations.

RELATED DOCUMENTATION

Viewing and Managing Forwarding Destinations | **428**

Viewing and Managing Routing Rules | **429**

Using Custom Rules and Rule Responses to Forward Data | **425**

# Viewing and Managing Forwarding Destinations

Use the Forwarding Destination window to view, edit, and delete forwarding destinations.

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Forwarding Destinations**.

   Statistics for the data sent to each forwarding destination is displayed. For example, you can see the following information:

   - The total number events and flows that were seen for this forwarding destination.
   - The number of events or flows that were sent to this forwarding destination.
   - The number of events or flows that were dropped before the forwarding destination was reached.
3. On the toolbar, click an action, as described in the following table.

**Table 68: Description of the Forwarding Destination Toolbar Actions**

| Action | Description |
|---|---|
| Reset Counters | Resets the counters for the **Seen**, **Sent**, and **Dropped** parameters to zero, and the counters start accumulating again.<br>**TIP**: You can reset the counters to provide a more targeted view of the performance of your forwarding destinations. |
| Edit | Changes the configured name, format, IP address, port, or protocol. |
| Delete | Deletes a forwarding destination<br>If the forwarding destination is associated with any active rules, you must confirm that you want to delete the forwarding destination. |

## RELATED DOCUMENTATION

Viewing and Managing Routing Rules | **429**

Using Custom Rules and Rule Responses to Forward Data | **425**

Viewing Forwarding Destinations | **427**

# Viewing and Managing Routing Rules

Use the Event Routing Rules window to enable or disable the rules, or to edit a rule to change the configured name, Event Collector, filters, or routing options.

1. On the navigation menu (≡), click **Admin**.
2. In the **System Configuration** section, click **Routing Rules**.
3. Select the routing rule you want to manage.
4. To edit the routing rule, on the toolbar, click **Edit** and update the parameters.
5. To remove the routing rule, on the toolbar, click **Delete**.

6. To enable or disable the routing rule, on the toolbar, click **Enable/Disable**.

   If you enable a routing rule that is configured to drop events, a confirmation message is displayed.

## RELATED DOCUMENTATION

Using Custom Rules and Rule Responses to Forward Data | **425**

Viewing Forwarding Destinations | **427**

Viewing and Managing Forwarding Destinations | **428**

20

CHAPTER

# Event Store and Forward

# Event Store and Forward

Use the Store and Forward feature to manage schedules for forwarding events from your dedicated Event Collector appliances to Event Processor components in your deployment.

The Store and Forward feature is supported on the Event Collector 1501 and Event Collector 1590. For more information about these appliances, see the *Juniper Secure Analytics Hardware Guide*.

A dedicated Event Collector does not process events and it does not include an on-board Event Processor. By default, a dedicated Event Collector continuously forwards events to an Event Processor that is connected to JSA.

You can schedule a time range for when you want the Event Collector to forward events to the Event Processor. By forwarding the events during non-business hours, you can ensure that the transmission does not negatively affect your network bandwidth. When event forwarding is scheduled, the events are stored locally on the Event Collector until the forwarding schedule kicks in. During this time, you cannot view the events in the JSA Console.

# Store and Forward Overview

The Store and Forward feature is supported on the Event Collector 1501 and Event Collector 1590. For more information about these appliances, see the *Juniper Secure Analytics Hardware Guide*.

A dedicated Event Collector does not process events and it does not include an on-board Event Processor. By default, a dedicated Event Collector continuously forwards events to an Event Processor that you must connect using the Deployment Editor. The Store and Forward feature allows you to schedule a time range for when you want the Event Collector to forward events. During the period of time when events are not forwarding, the events are stored locally on the appliance and are not accessible using the Console user interface.

This scheduling feature allows you to store events during your business hours and then forward the events to an Event Processor during periods of time when the transmission does not negatively affect your network bandwidth. For example, you can configure an Event Collector to only forward events to an Event Processor during non-business hours, such as midnight until 6 AM.

### RELATED DOCUMENTATION

Viewing the Store and Forward Schedule List | 433

Creating a Store and Forward Schedule | 436

Editing a Store and Forward Schedule | 437

# Viewing the Store and Forward Schedule List

You must create a schedule. By default, the first time that you access the **Store and Forward** window, no schedules are listed.

Use the **Store and Forward** window to see a list of schedules. The schedules include statistics that help you evaluate the status, performance, and progress of your schedules.

You can use options on the toolbar and the **Display** list box to change your view of the schedule list. Change your view of the list to focus on the statistics from various points of view. For example, if you want to view the statistics for a particular Event Collector, you can select **Event Collectors** from the **Display** list. The list then groups by the **Event Collector** column and makes it easier for you to locate the Event Collector that you want to investigate.

By default, the Store and Forward list is configured to display the list that is organized by the schedule (**Display** >**Schedules**).

1. On the navigation menu
   (
   ≡
   ), click **Admin**.
2. In the **System Configuration** section, click **Store and Forward**.
3. In the **Store and Forward** window, view the parameters for each schedule.

   The following table describes some of the parameters for the schedule.

   **Table 69: Store and Forward Window Parameters**

   | Parameter | Description |
   |---|---|
   | Display | The **Schedules** option shows a hierarchy of the parent-child relationship between the schedules, event processors, and the associated event collectors.<br>The **Event Collectors** option shows the lowest level in the hierarchy, which is a list of event collectors.<br>**Event Processors** option shows a hierarchy of the parent-child relationship between the event processors and the associated event collectors. |

**Table 69: Store and Forward Window Parameters** ***(Continued)***

| Parameter | Description |
|---|---|
| Name | For the **Schedules** option, the **Name** column is displayed the following format.<br>• **First Level** represents the name of the schedule.<br>• **Second Level** represents the name of the event processor.<br>• **Third Level** represents the name of the event collector.<br>For the **Event Processors** option, the column is displayed in the following format<br>• **First Level** represents the name of the event processor.<br>• **Second Level** represents the name of the event collector.<br>**TIP**: You can use the plus symbol (+) and minus symbol (-) beside the name or options on the toolbar to expand and collapse the hierarchy tree. You can also expand and collapse the hierarchy tree by using options on the toolbar. |
| Schedule Name | Displays the name of the schedule for the **Event Collectors** or **Event Processors** options.<br>If an event processor is associated with more than one schedule, the **Schedule Name** shows <Multiple>*n*, where *n* is the number of schedules.<br>**TIP**: Click the plus symbol (+) to view the associated schedules. |
| Last Status | Displays the status of the Store and Forward process:<br>• **Forwarding** indicates that event forwarding is in progress.<br>• **Forward Complete** indicates that event forwarding is successfully completed and events are stored locally on the event collector. The stored events are forwarded when the schedule indicates that forwarding can start again.<br>• **Warn** indicates that the percentage of events that are remaining in storage exceeds the percentage of time that is remaining in the Store and Forward schedule.<br>• **Error** indicates that event forwarding was stopped before all stored events were forwarded.<br>• **Inactive** indicates that no event collectors are assigned to the schedule, or the assigned event collectors are not receiving any events.<br>**TIP**: Move your mouse pointer over the **Last Status** column to view a summary of the status. |

Table 69: Store and Forward Window Parameters *(Continued)*

| Parameter | Description |
|---|---|
| Forwarded Events | Displays the number of events (in K, M, or G) forwarded in the current session.<br>**TIP**: Move your mouse pointer over the value in the **Forwarded Events** column to view the number of events. |
| Remaining Events | Displays the number of events (in K, M, or G) remaining to be forwarded in the current session.<br>**TIP**: Move your mouse pointer over the value in the **Remaining Events** column to view the number of events. |
| Average Event Rate | Displays the average rate at which events are forwarding from the event collector to the event processor.<br>**TIP**: Move your mouse pointer over the value in the **Average Event Rate** column to view the average events per second (EPS). |
| Current Event Rate | Displays the rate at which events are forwarding from the event collector to the event processor.<br>**TIP**: Move your mouse pointer over the value in the **Current Event Rate** column to view the current events per second (EPS) |
| Transfer Rate Limit | The transfer rate limit is configurable.<br>The transfer rate limit can be configured to display in kilobytes per second (KBs), megabytes per second (MBs), or gigabytes per second (GBs). |

## RELATED DOCUMENTATION

Creating a Store and Forward Schedule | 436

Editing a Store and Forward Schedule | 437

Deleting a Store and Forward Schedule | 437

# Creating a Store and Forward Schedule

Use the Store and Forward Schedule wizard to create a schedule that controls when your event collector starts and stops forwarding data to an event processor.

You can create and manage multiple schedules to control event forwarding from multiple JSA event collectors in a geographically distributed deployment.

Ensure that your dedicated event collector is added to your deployment and connected to an event processor. Use the **System and License Management** window to configure the connection between an event collector and an event processor.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Store and Forward**.
3. Click **Actions** >**Create**.
   a. Click **Next** to move to the **Select Collectors** page.

   b. On the **Select Collectors** page, configure the parameters.

      If the event collector that you want to configure is not listed, ensure that the event collector was added to the JSA deployment.

   c. On the **Schedule Options** page, configure the parameters.

      To configure the forward transfer rate, the minimum transfer rate is 0. The maximum transfer rate is 9,999,999. A value of 0 means that the transfer rate is unlimited.

   d. Finish the configuration.

      You can now view the schedule in the **Store and Forward** window. After you create a new schedule, it might take up to 10 minutes for statistics to start displaying in the **Store and Forward** window.

## RELATED DOCUMENTATION

Editing a Store and Forward Schedule | 437

Deleting a Store and Forward Schedule | 437

Viewing the Store and Forward Schedule List | 433

# Editing a Store and Forward Schedule

You can edit a **Store and Forward** schedule to add or remove JSA event collectors and change the schedule parameters. After you edit a **Store and Forward** schedule, the statistics that are displayed in the **Store and Forward** list are reset.

1. On the navigation menu ( ≡ ), click **Admin**.
2. In the **System Configuration** section, click **Store and Forward**.
3. Select the schedule that you want to edit.
4. Click **Actions** >**Edit**.

   You can also double-click a schedule for editing.
5. Click **Next** to move to the **Select Collectors** page.
6. On the **Select Collectors page**, edit the parameters.
7. Click **Next** to move to the **Schedule Options** page.
8. On the **Schedule Options** page, edit the scheduling parameters.
9. Click **Next** to move to the **Summary** page.
10. On the **Summary** page, confirm the options that you edited for this schedule.

    After you edit a schedule, it might take up to 10 minutes for statistics to update in the **Store and Forward** window.

RELATED DOCUMENTATION

Deleting a Store and Forward Schedule | **437**

Viewing the Store and Forward Schedule List | **433**

Creating a Store and Forward Schedule | **436**

# Deleting a Store and Forward Schedule

You can delete a **Store and Forward** schedule.

1. On the navigation menu, click **System Configuration** .
2. Click the **Store and Forward** icon.

3. Select the schedule that you want to delete.
4. Click **Actions** >**Delete**.

   After the schedule is deleted, the associated JSA event collectors resume continuous forwarding of events to their assigned event processor.

## RELATED DOCUMENTATION

Viewing the Store and Forward Schedule List | **433**

Creating a Store and Forward Schedule | **436**

Editing a Store and Forward Schedule | **437**

21

CHAPTER

# Security Content

# Security Content

You use the content management tools in JSA to import security content such as rules, reports, dashboards and applications into JSA. Security content can come from other JSA systems, or it can be developed independently to extend existing JSA capabilities.

# Types Of Security Content

IN THIS SECTION

JSA content is bundled into two types: content packs and extensions.

- **Content packs--** Security *content packs* contain enhancements to specific types of security content. Often, they include content for third-party integrations or operating systems. For example, a security content pack for a third-party integration might contain new custom event properties that make information in the event payload searchable for the log source and available for reporting.
- **Extensions--** Juniper and other vendors write security *extensions* that enhance or extend JSA capabilities. An extension can contain apps, content items, such as custom rules, report templates, saved searches, or contain updates to existing content items. For example, an extension might include an app to add a tab in JSA that provides visualizations for an offense.

  On IBM Security App Exchange, extensions are known as apps. You can download JSA apps from IBM Security App Exchange and use the **Extensions Management** tool to install them. Apps are not available as part of an auto-update.

## Sources Of Security Content

**JSA deployments--** You export custom content from a JSA deployment as an extension and then import it into another system when you want to reuse the content. For example, you can export content from your development environment to your production environment. You can use the content management script to export all content, or you can choose to export only some custom content.

RELATED DOCUMENTATION

Methods Of Importing and Exporting Content | 441
Content Type Identifiers for Exporting Custom Content | 458
Content Management Script Parameters | 460

# Methods Of Importing and Exporting Content

IN THIS SECTION

You can use the following tools to import and export content in your JSA deployment.

## Extensions Management Tool

Use the **Extensions Management** tool to add extensions to your JSA deployment. When you import content using the **Extensions Management** tool, you can view the content before it is installed. If the content items exist in your system, you can specify whether to replace the content item or skip the update.

You cannot use the **Extensions Management** tool to export content.

## Content Management Script

Use the content management script to export custom content from your JSA deployment into an external, portable format. You can then use the script to import the custom content into another JSA deployment. The script is useful when you want to automate moving content between your JSA deployments.

The **contentManagement.pl** script is in the **/opt/qradar/bin** directory.

You must use the content management script to export content from the JSA source deployment. You can use either the content management script or the **Extensions Management** tool to import the content to the target deployment.

DSM Editor

In JSA 7.3.3 and later, you can export your custom content that you create in the DSM Editor. Click the Export button in the DSM Editor to export your content from one JSA deployment to another, or to external media.

**NOTE**: You can export content from an earlier release of JSA and import into a later release. However, you cannot import content from a later release into an earlier release.

**NOTE**: If you move overridden rules from one JSA deployment to another, use the **Replace Existing Content Items** option to ensure that the rules are imported correctly.

## Exporting All Custom Content

You use the **contentManagement.pl** script to export all custom content in your JSA deployment.

1. Use SSH to log in to JSA as the root user.

2. Go to `/opt/qradar/bin` directory, and type the command to export all of the custom content:

   **./contentManagement.pl -a export -c all**

   **Examples:**

- To include accumulated data in the export, type the following command:

```
./contentManagement.pl --action export --content-type all -g
```

- To specify the directory for the exported file and change the compression format, type the following command:

```
./contentManagement.pl -a export -c all -o [filepath] -t [compression_type]
```

The content is exported to a compressed file, for example, **all-ContentExport-20151022101803.zip**. You can manually change the file name to a name that is more descriptive. The exported file might contain more content items than expected because all dependencies are exported with the specified content items. For example, if you export a report, the saved search that the report uses is also exported.

## Exporting All Custom Content Of a Specific Type

You can export all custom content of a specific type in one action.

The content management script uses text identifiers or numeric identifiers to specify the type of content that you want to export.

**Table 70: Content Type Identifiers for Exporting Custom Content**

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| Dashboards | **dashboard** | 4 |
| Reports | **report** | 10 |
| Saved searches | **search** | 1 |
| FGroups [1] | **fgroup** | 12 |
| FGroup types | **fgrouptype** | 13 |

**Table 70: Content Type Identifiers for Exporting Custom Content *(Continued)***

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| Custom rules | **customrule** | 3 |
| Custom properties | **customproperty** | 6 |
| Log sources | **sensordevice** | 17 |
| Log source types | **sensordevicetype** | 24 |
| Log source categories | **sensordevicecategory** | 18 |
| Log source extensions | **deviceextension** | 16 |
| Reference data collections | **referencedata** | 28 |
| Custom QID map entries | **qidmap** | 27 |
| Historical correlation profiles | **historicalsearch** | 25 |
| Custom functions | **custom_function** | 77 |
| Custom actions | **custom_action** | 78 |
| Applications | **installed_application** | 100 |
| DSM event mapping | **dsmevent** | 41 |
| [1] An FGroup represents a group of content, such as a log source group, reporting group, or search group. | | |

1. Use SSH to log in to JSA as the root user.

2. Go to the `/opt/qradar/bin` directory and type the command to export all content of the specified type:

./contentManagement.pl -a export --content-type *[content_type]* --id all

Parameters:

**Table 71: contentManagement.pl Script Parameters for exporting Custom Content of a Specific Type**

| Parameter | Description |
|---|---|
| **-c** [content_type] or **--content-type** [content_type] | Specifies the type of content.<br>You can type the corresponding text or numeric identifier to specify the content type.<br>**NOTE**: If you choose to export data of a specific content type, additional data from related content of any content type might be exported. |
| **-e** or **--include-reference-data-elements** | Set this flag to include reference data keys and elements in the export.<br>Reference data keys and reference data elements are applicable to the referencedata content type. This parameter is applicable only when you export reference data, or content items that are dependent on reference data. |
| **-g** or **--global-view** | Includes accumulated data in the export. |
| **-i** [content_identifier] or **--id** [content_identifier] | Specifies the identifier of a specific instance of custom content such as a single report or a single reference set.<br>You can specify all to export all content of the specified type. |
| **-o** [filepath] or **--output-directory** [filepath] | Specifies the full path to the directory where the export file is written.<br>If no output directory is specified, the content is exported to the current directory. If the specified output directory does not exist, it is created. |
| **-t** [compression_type] or **--compression-type** [compression_type] | Specifies the compression type of the export file.<br>Valid options are ZIP and TARGZ (case sensitive). If you do not specify a compression type, the default compression type is ZIP. |

Examples:

- To export all custom searches, type the following command:

```
./contentManagement.pl --action export --content-type search --id all
```

- To export all reports and include accumulated data, type the following command:

```
./contentManagement.pl -a export -c 10 --id all --global-view
```

The content is exported to a compressed file, for example, **reports-ContentExport-20151022101803.zip**. You can manually change the file name to a name that is more descriptive. The exported file might contain more content items than expected because all dependencies are exported with the specified content items. For example, if you export a report, the saved search that the report uses is also exported.

## Searching for Specific Content Items to Export

You use the content management script to search for specific content in your JSA deployment. After you find the content, you can use the unique identifier to export the content item.

The following table lists the identifiers to use when you want to search for specific types of content.

**Table 72: Content Type Identifiers for Searching Custom Content**

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| Dashboards | **dashboard** | 4 |
| Reports | **report** | 10 |
| Saved searches | **search** | 1 |
| FGroups [1] | **fgroup** | 12 |
| FGroup types | **fgrouptype** | 13 |

Table 72: Content Type Identifiers for Searching Custom Content *(Continued)*

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| Custom rules | **customrule** | 3 |
| Custom properties | **customproperty** | 6 |
| Log sources | **sensordevice** | 17 |
| Log source types | **sensordevicetype** | 24 |
| Log source categories | **sensordevicecategory** | 18 |
| Log source extensions | **deviceextension** | 16 |
| Reference data collections | **referencedata** | 28 |
| Custom QID map entries | **qidmap** | 27 |
| Historical correlation profiles | **historicalsearch** | 25 |
| Custom functions | **custom_function** | 77 |
| Custom actions | **custom_action** | 78 |
| Applications | **installed_application** | 100 |
| [1]An FGroup represents a group of content, such as a log source group, reporting group, or search group. | | |

1. Use SSH to log in to JSA as the root user.

2. Go to the `/opt/qradar/bin` directory and type the following command to search for custom content that matches a regular expression:

   **./contentManagement.pl -a search -c** ***[content_type]*** **-r** ***[regex]***

**Parameters:**

**Table 73: contentManagement.pl Script Parameters for Searching Content Items**

| Parameter | Description |
|---|---|
| **-c** [content_type] or **--content-type** [content_type] | Specifies the type of content to search for.<br>You must specify the type of content to search for. You cannot use -c package or -c all with the search action. |
| **-r** [regex] or **--regex** [regex] | Specifies the content to search for.<br>All content that matches the expression is displayed. |

**Examples:**

- To search for all reports that includes Overview in the description, type the following command:

```
/opt/qradar/bin/contentManagement.pl --action search
--content-type report --regex "Overview"
```

- To list all log sources, type the following command:

```
/opt/qradar/bin/contentManagement.pl -a search -c 17 -r "\w"
```

The search results list details, including the unique ID, for the content items that are found.

**[INFO] Search results: [INFO] - [ID] - [Name] - [Description] [INFO] - [67] - [Asset Profiler-2 :: hostname] - [Asset Profiler] [INFO] - [62] - [SIM Generic Log DSM-7 :: hostname] - [SIM Generic Log DSM] [INFO] - [63] - [Custom Rule Engine-8 :: hostname] - [Custom Rule Engine] [INFO] - [71] - [Pix @ apophis] - [Pix device] [INFO] - [70] - [Snort @ wolverine] - [Snort device] [INFO] - [64] - [SIM Audit-2 :: hostname] - [SIM Audit] [INFO] - [69] - [Health Metrics-2 :: hostname] - [Health Metrics]**

Use the unique identifier to export specific content items from JSA. For more information, see "Content Management Script Parameters" on page 460.

## Exporting a Single Custom Content Item

Export a single custom content item, such as a custom rule or a saved search, from JSA.

You must know the unique identifier for the custom content item that you want to export.

1. Us SSH to log in to JSA as the root user.

2. Go to the `/opt/qradar/bin` directory and type the command to export the content:

   **./contentManagement.pl -a export -c *[content_type]* -i *[content_identifier]***

   **Parameters:**

   **Table 74: contentManagement.pl Script Parameters for Exporting a Single Content Item**

| Parameter | Description |
|---|---|
| **-c** [content_type] or **--content-type** [content_type] | Specifies the type of content to export.<br>Type the corresponding text identifier or numeric identifier for specific content types. |
| **-e** or **--include-reference-data-elements** | Set this flag to include reference data keys and elements in the export.<br>Reference data keys and reference data elements are applicable to the referencedata content type. This parameter is applicable only when you export reference data, or content items that are dependent on reference data. |
| **-g** or **--global-view** | Includes accumulated data in the export. |
| **-i** [content_identifier] or **--id** [content_identifier] | Specifies the identifier of a specific instance of custom content such as a single report or a single reference set. |
| **-o** [filepath] or **--output-directory** [filepath] | Specifies the full path to the directory where the export file is written.<br>If no output directory is specified, the content is exported to the current directory. If the specified output directory does not exist, it is created. |

Table 74: contentManagement.pl Script Parameters for Exporting a Single Content Item *(Continued)*

| Parameter | Description |
|---|---|
| **-t** [compression_type] or **--compression-type** [compression_type] | Used with the export action.<br>Specifies the compression type of the export file. Valid options are ZIP and TARGZ (case sensitive). If you do not specify a compression type, the default compression type is ZIP. |

**Examples:**

- To export the dashboard that has ID 7 into the current directory, type the following command:

```
./contentManagement.pl -a export -c dashboard -i 7
```

- To export the log source that has ID 70, including accumulated data, into the **/store/cmt/exports** directory, type the following command:

```
./contentManagement.pl -a export -c sensordevice -i 70 -o /store/cmt/exports -g
```

The content is exported to a compressed **.zip** file. The exported file might contain more content items than expected because all dependencies are exported with the specified content items. For example, if you export a report, the saved search that the report uses is also exported. You can manually change the file name to a name that is more descriptive.

## Exporting Custom Content Items Of Different Types

Export multiple custom content items from JSA, such as custom rules, or dashboards and reports, by using the content management script.

You must know the unique identifiers for each custom content item that you want to export.

1. Use SSH to log in to JSA as the root user.

2. Create a text file that lists the content that you want to export.

   Each line must include the custom content type followed by a comma-separated list of unique IDs for that type.

**Example:** To export two dashboards that have ID 5 and ID 7, all custom rules, and a group, create a text file that has the following entries:

```
dashboard, 5,7
customrule, all
fgroup, 77
```

3. Go to `/opt/qradar/bin` and type the command to export the content:

   **./contentManagement.pl -a export -c package -f** ***[source_file]***

   **Parameters:**

   **Table 75: contentManagement.pl Script Parameters for Exporting Different Types of Content Item**

| Parameter | Description |
|---|---|
| **-c** [content_type] or **--content-type** [content_type] | Specifies the type of content.<br>Specifies the type of content. You can specify -c package, or you can type the corresponding text or numeric identifier for specific content types. When you use -c package, you must specify the --file or --name parameters. |
| **-e** or **--include-reference-data-elements** | Set this flag to include reference data keys and elements in the export.<br>Reference data keys and reference data elements are applicable to the referencedata content type. This parameter is applicable only when you export reference data, or content items that are dependent on reference data. |
| **-f** [source_file] or **--file** [source_file] | Specifies the path and file name of the text file that contains the list of custom content items that you want to export.<br>The first time you use the --file parameter, a package template file is written to the **/store/cmt/packages** directory so that you can reuse it.<br>The filename and path are case-sensitive. |
| **-g** or **--global-view** | Includes accumulated data in the export. |

**Table 75: contentManagement.pl Script Parameters for Exporting Different Types of Content Item *(Continued)***

| Parameter | Description |
|---|---|
| **-n** [name] or **--name** [name] | Specifies the name of the package template file that contains the list of custom content to export.<br>The package template file is created the first time that you use the --file parameter. By default, the --name parameter assumes that the text file is in the **/store/cmt/packages** directory.<br>You must specify the --file or --name parameter when --content-type package is used. |
| **-o** [filepath] or **--output-directory** [filepath] | Specifies the full path to the directory where the export file is written.<br>If no output directory is specified, the content is exported to the current directory. If the specified output directory does not exist, it is created. |
| **-t** [compression_type] or **--compression-type** [compression_type] | Specifies the compression type of the export file.<br>Valid compression types are ZIP and TARGZ (case sensitive). If you do not specify a compression type, the default compression type is ZIP. |

**Examples:**

- To export all items in the **exportlist.txt** file in the jsa directory, and save the exported file in the current directory, type the following command:

```
./contentManagement.pl -a export -c package -f /qradar/exportlist.txt
```

- To export all items in the exportlist.txt file in the jsa directory, including accumulated data, and save the output in the /store/cmt/exports directory, type the following command:

```
./contentManagement.pl -a export -c package
--file /qradar/exportlist.txt -o /store/cmt/exports -
```

When you use the **--file** parameter, a package template file is automatically generated in **/store/cmt/packages**. To use the package template file, specify the filename as the value for the **--name** parameter.

The content is exported to a compressed **.zip** file. The exported file might contain more content items than expected because all dependencies are exported with the specified content items. For example, if you export a report, the saved search that the report uses is also exported. You can manually change the file name to a name that is more descriptive.

## Installing Extensions by Using Extensions Management

Use the **Extensions Management** tool to add security extensions to JSA. The **Extensions Management** tool allows you to view the content items in the extension and specify the method of handling content updates before you install the extension.

Extensions must be on your local computer before you install them in JSA.

An extension is a bundle of JSA functions. An extension can include content such as rules, reports, searches, reference sets, and dashboards. It can also include applications that enhance JSA functions.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Extensions Management**.

3. To upload a new extension to the JSA console, follow these steps:

    a. Click **Add**.

    b. Click **Browse** and navigate to find the extension.

    c. Click **Install immediately** to install the extension without viewing the contents. See "5" on page 453.b.

    d. Click **Add**.

4. To view the contents of the extension, select it from the extensions list and click **More Details**.

5. To install the extension, follow these steps:

    a. Select the extension from the list and click **Install**.

    b. To assign a user to the app, select the **User Selection** menu, and select a user. For example, you might want to associate the app with a specified user that is listed in the **User Selection** menu who has the defined permissions.

**NOTE**: This screen appears only if any of the apps in the extension that you are installing are configured to request authentication for background processes.

c. If the extension does not include a digital signature, or it is signed but the signature is not associated with the JSA Security Certificate Authority (CA), you must confirm that you still want to install it. Click **Install** to proceed with the installation.

d. Review the changes that the installation makes to the system.

e. Select **Overwrite** or **Keep existing data** to specify how to handle existing content items.

**NOTE**: If the extension contains overridden system rules, select Replace Existing Items to ensure that the rules are imported correctly.

f. Click **Install**.

g. Review the installation summary and click **OK**.

## Uninstalling a Content Extension

Remove a content extension that isn't useful anymore or that adversely impacts the system. You can remove rules, custom properties, reference data, and saved searches. You might not be able to remove some content if another content item depends on it.

When you uninstall a content extension, any rules, custom properties, and reference data that were installed by the content extension are removed or reverted to their previous state. Saved searches can't be reverted. They can only be removed.

For example, if you've edited custom rules in an app that you now want to uninstall, you can preserve the changes you made for each customized rule. If the custom rule previously existed on the system, you can revert the rule to its previous state. If the custom rule didn't previously exist, you can remove it.

**NOTE**: If you have introduced an outside dependency on a content extension that is installed by the app, JSA doesn't remove that piece of content when you uninstall the app. For example, if you create a custom rule that uses one of the app's custom properties, that custom property isn't removed when you uninstall the app.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **System Configuration** section, click **Extensions Management**.

3. Select the extension that you want to uninstall and click **Uninstall**.

   JSA checks for any applications, rules, custom properties, reference data, and saved searches that are installed by the content extension that can be removed.

4. If you have manually altered any rules, custom properties, or reference data after you installed the app, choose whether to **Preserve** or **Remove/Revert** that content extension.

5. Click **Uninstall**, and then click **OK**.

## Importing Content by Using the Content Management Script

You can import custom content that you exported from another JSA system.

If you want to import content from another JSA system, you must first export the content and copy it to the target system. For more information about exporting content, see "Content Type Identifiers for Exporting Custom Content" on page 458.

When you import content that has log sources, confirm that DSM and protocol RPMs are installed and current on the target system.

> **NOTE**: If the content contains overridden system rules, use the update action instead of the import action to ensure that the rules are imported correctly.

You can export content from an earlier release of JSA and import into a later release. However, you cannot import content from a later release into an earlier release.

You do not have to export content in a specific order. However, do not start multiple imports on the same system at the same time. The imports fail due to conflicts with shared resources.

1. Use SSH to log in to JSA as the root user.

2. Go to the directory where the export content file is located.

3. Type this command to import the content:

   ```
   /opt/qradar/bin/contentManagement.pl -a import -f [source_file] -u [user]
   ```

**Parameters:**

**Table 76: contentManagement.pl Script Parameters for Importing Custom Content**

| Parameter | Description |
|---|---|
| **-f** [source_file] or **--file** [source_file] | Specifies the file that contains the content items to import.<br>Valid file types are zip, targz, and xml.<br>The file name and path are case-sensitive. |
| **-u** [user] or **--user** [user] | Specifies the user that replaces the current owner when you import user-specific data. The user must exist on the target system before you import the content. |

**Examples:**

- To import content from the **fgroup-ContentExport-20120418163707.tar.gz** file in the current directory, type the following command:

```
/opt/qradar/bin/contentManagement.pl --action import
-f fgroup-ContentExport-20120418163707.tar.gz
```

- To import content from the **fgroup-ContentExport-20120418163707.tar.gz** file in the current directory and make the admin user the owner of all sensitive data in the import, type the following command:

```
/opt/qradar/bin/contentManagement.pl --action import
--file fgroup-ContentExport-20120418163707.tar.gz --user admin
```

The import script displays the following message when reference data is actively collected while it is being exported: `Foreign key constraint violation`. To avoid this issue, run the export process when no reference data is being collected.

## Updating Content by Using the Content Management Script

Use the update action to update existing JSA content or add new content to the system.

If you want to update content with content that was exported from another JSA system, ensure that the exported file is on the target system. For more information about exporting content, see "Content Type Identifiers for Exporting Custom Content" on page 458.

When you import content that has log sources, confirm that DSM and protocol RPMs are installed and current on the target system.

You can export content from an earlier release of JSA and import into a later release. However, you cannot import content from a later release into an earlier release.

You do not have to export content in a specific order. However, do not start multiple imports on the same system at the same time. The imports will fail due to conflicts with shared resources.

1. Use SSH to log in to JSA as the root user.

2. To update content, type the following command:

   **/opt/qradar/bin/contentManagement.pl -a update -f *[source_file]***

   **Parameters:**

   **Table 77: contentManagement.pl Script Parameters for Updating Custom Content**

| Parameter | Description |
|---|---|
| -f [source_file] or --file [source_file] | Specifies the file that contains the content items to update.<br>Valid file types are zip, targz, and xml.<br>The filename and path are case-sensitive. |
| -u [user] or --user [user] | Specifies the user that replaces the current owner when you import user-specific data.<br>The user must exist on the target system before you import the content. |

   **Example:**

   - To update based on the content in the **fgroup-ContentExport- 20120418163707.zip** file, type the following command:

     ```
     /opt/qradar/bin/contentManagement.pl --action update
     -f fgroup-ContentExport-20120418163707.zip
     ```

RELATED DOCUMENTATION

Content Type Identifiers for Exporting Custom Content | 458

Content Management Script Parameters | 460

Types Of Security Content | 440

# Content Type Identifiers for Exporting Custom Content

When you export a specific type of custom content from JSA, you must specify the content type. You must use either the text identifier or the numeric identifier for the content type.

When you export content from a JSA appliance, the content management script checks content dependencies, and then includes associated content in the export.

For example, when the content management script detects that a saved search is associated with a report that you want to export, the saved search is also exported. You can't export offense, asset, or vulnerability saved searches.

You use the content type identifier when you want to export all custom content of a specific type. If you want to export a specific content item from your JSA deployment, you must know the unique identifier for that specific content item.

The following table describes the content type identifiers that are passed into the `contentManagement.pl` script for the **-c** parameter.

**Table 78: Content Type Identifiers for Exporting Custom Content**

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| All custom content | **all** | Not applicable |
| Custom list of content | **package** | Not applicable |
| Dashboards | **dashboard** | 4 |
| Reports | **report** | 10 |

Table 78: Content Type Identifiers for Exporting Custom Content *(Continued)*

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| Saved searches | **search** | 1 |
| FGroups [1] | **fgroup** | 12 |
| FGroup types | **fgrouptype** | 13 |
| Custom rules | **customrule** | 3 |
| Custom properties | **customproperty** | 6 |
| Log sources | **sensordevice** | 17 |
| Log source types | **sensordevicetype** | 24 |
| Log source categories | **sensordevicecategory** | 18 |
| Log source extensions | **deviceextension** | 16 |
| Reference data collections | **referencedata** | 28 |
| Custom QID map entries | **qidmap** | 27 |
| Historical correlation profiles | **historicalsearch** | 25 |
| Custom functions | **custom_function** | 77 |
| Custom actions | **custom_action** | 78 |
| Applications | **installed_application** | 100 |

**Table 78: Content Type Identifiers for Exporting Custom Content *(Continued)***

| Custom content type | Text identifier | Numeric identifier |
|---|---|---|
| [1]An FGroup is a group of content such as a log source group, reporting group, or search group. | | |

RELATED DOCUMENTATION

Content Management Script Parameters | 460

Types Of Security Content | 440

Methods Of Importing and Exporting Content | 441

# Content Management Script Parameters

Use the **contentManagement.pl** script to export content from one JSA deployment and import it to another deployment.

The following table describes the parameters for the **contentManagement.pl** script and the actions to which each parameter applies.

/opt/qradar/bin/contentManagement.pl --action *[action_type] [script_parameters]*

**Table 79: ContentManagement.pl Script Parameters**

| Parameter | Description |
|---|---|
| -a*[action_type]*<br>or<br>--action*[action_type]* | Required. Specifies the action.<br>Valid action types are `export`, `search`, `import`, and `update`.<br>The `import` action adds only content that does not exist in the deployment. |

Table 79: ContentManagement.pl Script Parameters *(Continued)*

| Parameter | Description |
|---|---|
| -c*[content_type]*<br>or<br>--content-type*[content_type]* | Used with the `export` and `search` actions. Specifies the type of content.<br>When used with the `export` action, you can specify `-c all` or `-c package`, or you can type the "Content Type Identifiers for Exporting Custom Content" on page 458. When you use `-c package`, you must specify the `--file` or `--name` parameters.<br>When used with the `search` action, you must specify the "Content Type Identifiers for Exporting Custom Content" on page 458. You cannot use `-c package` or `-c all` with the `search` action. |
| -d<br>or<br>--debug | Used with all actions.<br>Use debug level logging when you run the `contentManagement.pl` script to see more detailed information, such as logs for customer support. |
| -e<br>or<br>--include-reference-data-elements | Used with the `export` action.<br>Set this flag to include reference data keys and elements in the export.<br>Reference data keys and reference data elements are applicable to the `referencedata` content type. This parameter is applicable only when you export reference data, or content items that are dependent on reference data. |

Table 79: ContentManagement.pl Script Parameters *(Continued)*

| Parameter | Description |
|---|---|
| -f*[file_path]*<br>or<br>--file*[file_path]* | Used with `export`, `import`, and `update` actions.<br>When used with the `export` action, specifies the path and file name of the text file that contains the list of custom content items that you want to export. The first time you use the `--file` parameter, a package template file is written to the `/store/cmt/packages` directory so that you can reuse it.<br>When used with the `import` or `update` action, specifies the file that contains the content items to import. Valid file types are **zip**, **targz**, and **xml**.<br>The filename and path are case-sensitive. |
| -g<br>or<br>--global-view | Used with the `export` action.<br>Includes accumulated data in the export. |
| -h*[action_type]*<br>or<br>--help*[action_type]* | Used with all actions.<br>Displays help that is specific to the `action_type`. When no `action_type` is specified, displays a general help message. |
| -i*[content_identifier]*<br>or<br>--id*[content_identifier]* | Used with the `export` action.<br>. You can specify *all* to "Exporting All Custom Content Of a Specific Type" on page 443. |

**Table 79: ContentManagement.pl Script Parameters *(Continued)***

| Parameter | Description |
|---|---|
| **-n***[name]*<br>or<br>**--name***[name]* | Used with the `export` action.<br>Specifies the name of the package template file that contains the list of custom content to export.<br>The package template file is created the first time that you use the `--file` parameter. The `--name` parameter assumes that the package template file is in the `/store/cmt/packages` directory.<br>You must specify the `--file` or `--name` parameter when `--content-type package` is used. |
| **-o***[filepath]*<br>or<br>**--output-directory***[filepath]* | Used with the `export` action.<br>Specifies the full path to the directory where the export file is written.<br>If no output directory is specified, the content is exported to the current directory. If the specified output directory does not exist, it is created. |
| **-q**<br>or<br>**--quiet** | Used with all actions. No output appears on the screen. |
| **-r***[regex]*<br>or<br>**--regex***[regex]* | Used with the `search` action.<br>When searching, you must use the `--regex` parameter to specify the content to search for. All content that matches the expression is displayed. |
| **-t***[compression_type]*<br>or<br>**--compression-type***[compression_type]* | Used with the `export` action.<br>Specifies the compression type of the export file. Valid compression types are **ZIP** and **TARGZ** (case sensitive). If you do not specify a compression type, the default compression type is **ZIP**. |

**Table 79: ContentManagement.pl Script Parameters *(Continued)***

| Parameter | Description |
|---|---|
| -u*[user]*<br>or<br>--user*[user]* | Used with the `import` action.<br>Specifies the user that replaces the current owner when you import user-specific data. The user must exist on the target system before you import the content. |
| -v<br>or<br>--verbose | Used with all actions.<br>Use when you log in to view default-level information for the content management tool. |

## RELATED DOCUMENTATION

Types Of Security Content | **440**

Methods Of Importing and Exporting Content | **441**

Content Type Identifiers for Exporting Custom Content | **458**

22

CHAPTER

# SNMP Trap Configuration

# SNMP Trap Configuration

JSA uses the Net-SNMP agent, which supports various system resource monitoring MIBs. They can be polled by Network Management solutions for the monitoring and alerting of system resources. For more information about Net-SNMP, see Net-SNMP documentation.

In JSA, you can configure a rule to generate a rule response that sends an SNMP trap when configured conditions are met. JSA acts as an agent to send the SNMP traps to another system.

A Simple Network Management Protocol (SNMP) trap is an event or offense notification that JSA sends to a configured SNMP host for additional processing.

Customize the SNMP configuration parameters in the custom rules wizard and modify the SNMP traps that the custom rule engine sends to other software for management. JSA provides two default traps. However, you can add custom traps or modify the existing traps to use new parameters.

For more information on SNMP, go to the The Internet Engineering Task Force (http://www.ietf.org/) website and type **RFC 1157** in the search field.

**NOTE**: SNMPv3 rule responses are sent out as SNMP informs and not traps.

# Adding a Custom SNMP Trap to JSA

In JSA products, you can create a new option for the SNMP trap selection in the custom rules wizard. The trap names that are specified in the list box are configured in the **snmp-master.xml**configuration file.

**NOTE**: SNMPv3 rule responses are sent out as SNMP informs and not traps.

1. Use SSH to log in to JSA as the root user.
2. Go to the **/opt/qradar/conf** directory.
3. Create an SNMP settings file for the new trap.

   **TIP**: Copy, rename, and modify one of the existing SNMP settings files.

4. Make a backup copy of the **snmp-master.xml** file.

5. Open the **snmp-master.xml** file for editing.
6. Add a new <include> element.

   The <include> element has the following attributes:

   **Table 80: Attributes for the <include> Element**

   | Attribute | Description |
   |---|---|
   | `name` | Displayed in the list box |
   | `uri` | The name of the custom SNMP settings file |

   **Example:**

   ```
   <include name="Custom_Event_Name" uri="customSNMPdef01.xml"/>
   ```

   The traps are displayed in the menu in the same order in which they are listed in the **snmp-master.xml** file.
7. Save and close the file.
8. Copy the `snmp-master.xml` file and the `customSNMPdef01.xml file` from the **/opt/qradar/conf** directory to the **/store/configservices/staging/globalconfig** directory.
9. Log in to the JSA interface.
10. Log in to the JSA as an administrator.
11. On the navigation menu
    (
    ☰
    ), click **Admin**.
12. Select **Advanced** >**Deploy Full Configuration**.

    > **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

RELATED DOCUMENTATION

Sending SNMP Traps to a Specific Host | 468

# Sending SNMP Traps to a Specific Host

By default, in JSA products, SNMP traps are sent to the host that is identified in your **host.conf** file. You can customize the **snmp.xml** file to send SNMP traps to a different host.

> **NOTE**: SNMPv3 rule responses are sent out as SNMP informs and not traps.

1. Use SSH to log in to JSA as the root user.
2. Go to the **/opt/qradar/conf** directory and make backup copies of the following files:
   - **eventCRE.snmp.xml**
   - **offenseCRE.snmp.xml**
3. Open the configuration file for editing.
   - To edit the SNMP parameters for event rules, open the **eventCRE.snmp.xml** file.
   - To edit the SNMP parameters for offense rules, open the **offenseCRE.snmp.xml** file.
4. Add no more than one *<trapConfig>* element inside the *<snmp>* element inside the *<creSNMPTrap>* element and before any other child elements.

```
<trapConfig>
        <!-- All attribute values are default -->
            <snmpHost snmpVersion="3" port="162" retries="2" timeout="500">HOST
            </snmpHost>
        <!-- Community String for Version 2 -->
            <communityString>COMMUNITY_STRING</communityString>
        <!-- authenticationProtocol (MD5 or SHA)securityLevel (AUTH_PRIV, AUTH_NOPRIV
        or NOAUTH_PRIV) -->
            <authentication authenticationProtocol="MD5"securityLevel="AUTH_PRIV">
                AUTH_PASSWORD
            </authentication>
        <!-- decryptionProtocol (DES, AES128, AES192 or AES256) -->
            <decryption decryptionProtocol="AES256">
                    DECRYPTIONPASSWORD
            </decryption>
```

```
          <!-- SNMP USER-->
              <user> SNMP_USER </user>
      </trapConfig>
```

5. Use the following table to help you update the attributes.

**Table 81: Attribute Values to Update in the <trapConfig> Element**

| Element | Description |
|---|---|
| `</snmpHost>` | The new host to which you want to send SNMP traps.<br>The value for the`snmpVersion` attribute for `<snmpHost>` element must be 2 or 3. |
| `<communityString>` | The community string for the host. Do not use special characters. |
| `<authentication>` | An authentication protocol, security level, and password for the host. |
| `<decryption>` | The decryption protocol and password for the host. |
| `<user>` | SNMP user |

6. Save and close the file.
7. Copy the file from the **/opt/qradar/conf** directory to the **/store/ configservices/staging/ globalconfig** directory.
8. Log in to the JSA as an administrator.
9. On the navigation menu ( ≡ ), click **Admin**.
10. Select **Advanced** >**Deploy Full Configuration**.

> **NOTE**: JSA continues to collect events when you deploy the full configuration. When the event collection service must restart, JSA does not restart it automatically. A message displays that gives you the option to cancel the deployment and restart the service at a more convenient time.

RELATED DOCUMENTATION

Adding a Custom SNMP Trap to JSA | **466**

23

CHAPTER

# Protect Sensitive Data

# Sensitive Data Protection

Configure a data obfuscation profile to prevent unauthorized access to sensitive or personal identifiable information in JSA.

*Data obfuscation* is the process of strategically hiding data from JSA users. You can hide custom properties, normalized properties such as user names, or you can hide the content of a payload, such as credit card or social security numbers.

The expressions in the data obfuscation profile are evaluated against the payload and normalized properties. If the data matches the obfuscation expression, the data is hidden in JSA. The data might be hidden to all users, or only to users belonging to particular domains or tenants. Affected users who try to query the database directly cannot see the sensitive data. The data must be reverted to the original form, by uploading the private key that was generated when the data obfuscation profile was created.

To ensure that JSA can still correlate the hidden data values, the obfuscation process is deterministic. It displays the same set of characters each time the data value is found.

# How Does Data Obfuscation Work?

IN THIS SECTION

Before you configure data obfuscation in your JSA deployment, you must understand how it works for new and existing offenses, assets, rules, and log source extensions.

## Existing Event Data

When a data obfuscation profile is enabled, the system masks the data for each event as it is received by JSA. Events that are received by the appliance before data obfuscation is configured remain in the original unobfuscated state. The older event data is not masked and users can see the information.

## Assets

When data obfuscation is configured, the asset model accumulates data that is masked while the pre-existing asset model data remains unmasked.

To prevent someone from using unmasked data to trace the obfuscated information, purge the asset model data to remove the unmasked data. JSA will repopulate the asset database with obfuscated values.

## Offenses

To ensure that offenses do not display data that was previously unmasked, close all existing offenses by resetting the SIM model. For more information, see "Resetting SIM" on page 119.

## Rules

You must update rules that depend on data that was previously unmasked. For example, rules that are based on a specific user name do not fire when the user name is obfuscated.

## Log Source Extensions

Log source extensions that change the format of the event payload can cause issues with data obfuscation.

RELATED DOCUMENTATION

Data Obfuscation Profiles | 474

# Data Obfuscation Profiles

The data obfuscation profile contains information about which data to mask. It also tracks the keystore that is required to decrypt the data.

- **Enabled profiles**--Enable a profile only when you are sure that the expressions correctly target the data that you want to obfuscate. If you want to test the regular expression before you enable the data obfuscation profile, you can create a regex-based custom property.

  A profile that is enabled immediately begins obfuscating data as defined by the enabled expressions in the profile. The enabled profile is automatically locked. Only the user who has the private key can disable or change the profile after it is enabled.

  To ensure that obfuscated data can be traced back to an obfuscation profile, you cannot delete a profile that was enabled, even after you disable it.

- **Locked profiles**-- A profile is automatically locked when you enable it, or you can lock it manually.

  A locked profile has the following restrictions:

  - You cannot edit it.
  - You cannot enable or disable it. You must provide the keystore and unlock the profile before you can change it.
  - You cannot delete it, even after it is unlocked.
  - If a keystore is used with a profile that is locked, all other profiles that use that keystore are automatically locked.

  The following table shows examples of profiles that are locked or unlocked:

  **Table 82: Locked Profile Examples**

| Scenario | Result |
|---|---|
| Profile A is locked. It was created by using keystore A.<br>Profile B is also created by using keystore A. | Profile B is automatically locked. |

**Table 82: Locked Profile Examples *(Continued)***

| Scenario | Result |
|---|---|
| Profile A is created and enabled. | Profile A is automatically locked. |
| Profile A, Profile B, and Profile C are currently locked. All were created by using keystore A.<br>Profile B is selected and **Lock/Unlock** is clicked. | Profile A, Profile B, and Profile C are all unlocked. |

RELATED DOCUMENTATION

Data Obfuscation Expressions | 475

Scenario: Obfuscating User Names | 477

How Does Data Obfuscation Work? | 472

# Data Obfuscation Expressions

IN THIS SECTION

Data obfuscation expressions identify the data to hide. You can create data obfuscation expressions that are based on field-based properties or you can use regular expressions.

## Field-based Properties

Use a field-based property to hide user names, group names, host names, and NetBIOS names. Expressions that use field-based properties obfuscate all instances of the data string. The data is hidden regardless of its log source, log source type, event name, or event category.

If the same data value exists in more than one of the fields, the data is obfuscated in all fields that contain the data even if you configured the profile to obfuscate only one of the four fields. For example, if you have a host name that is called `JSAHost` and a group name that is called `JSAHost`, the value `JSAHost` is obfuscated in both the host name field and the group name field even if the data obfuscation profile is configured to obfuscate only host names.

## Regular Expressions

Use a regular expression to obfuscate one data string in the payload. The data is hidden only if it matches the log source, log source type, event name, or category that is defined in the expression.

You can use high-level and low-level categories to create a regular expression that is more specific than a field-based property. For example, you can use the following regex patterns to parse user names:

**Table 83: Regex User Name Parsing**

| Example regex patterns | Matches |
|---|---|
| usrName=([0-9a-zA-Z]([-.\w]*[0-9a-zA-Z])*@([0-9 a-zA-Z][-\w]*[0-9a-zA-Z]\.)+[a-zA-Z]{2,20})$ | `username@companyname.com` |
| usrName=(^([\w]+[^\W])([^\W]\.?)([\w]+[^\W]$)) | `username` |
| usrName=^([a-zA-Z])[a-zA-Z_-]*[\w_-]*[\S]$\|^([a -zA-Z])[0-9_-]*[\S]$\|^[a-zA-Z]*[\S]$ | `username` |
| usrName=(/S+) | Matches any non-white space after the equal, =, sign. This regular expression is non-specific and can lead to system performance issues. |

Table 83: Regex User Name Parsing *(Continued)*

| Example regex patterns | Matches |
|---|---|
| msg=([0-9a-zA-Z]([-.\w]*[0-9a-zA-Z]))*@\b(([01] ?\d?\d\|2[0-4]\d\|25[0-5])\.){3}([01]?\d?\d\|2[0-4 ]\d\|25[0-5])\b | Matches users with IP address. For example, `username@10.1.1.1` |
| src=\b(([01]?\d?\d\|2[0-4]\d\|25[0-5])\.){3}([01] ?\d?\d\|2[0-4]\d\|25[0-5])\b | Matches IP address formats. |
| host=^(([a-zA-Z0-9]\|[a-zA-Z0-9][a-zA-Z0-9\-]*[a -zA-Z0-9])\.)*([A-Za-z0-9]\|[A-Za-z0-9][A-Za-z0- 9\-]*[A-Za-z0-9])$ | `hostname.companyname.com` |

RELATED DOCUMENTATION

Scenario: Obfuscating User Names | **477**

How Does Data Obfuscation Work? | **472**

Data Obfuscation Profiles | **474**

# Scenario: Obfuscating User Names

**IN THIS SECTION**

You are an JSA administrator. Your organization has an agreement with the workers union that all personal identifiable information must be hidden from JSA users. You want to configure JSA to hide all user names.

Use the **Data Obfuscation Management** feature on the **Admin** tab to configure JSA to hide the data:

1. Create a data obfuscation profile and download the system-generated private key. Save the key in a secure location.

2. Create the data obfuscation expressions to target the data that you want to hide.

3. Enable the profile so that the system begins to obfuscate the data.

4. To read the data in JSA, upload the private key to deobfuscate the data.

## Creating a Data Obfuscation Profile

JSA uses data obfuscation profiles to determine which data to mask, and to ensure that the correct keystore is used to unmask the data.

You can create a profile that creates a new keystore or you can use an existing keystore. If you create a keystore, it must be downloaded and stored in a secure location. Remove the keystore from the local system and store it in a location that can be accessed only by users who are authorized to view the unmasked data.

Configuring profiles that use different keystores is useful when you want to limit data access to different groups of users. For example, create two profiles that use different keystores when you want one group of users to see user names and another group of users to see host names.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **Data Sources** section, click **Data Obfuscation Management**.

3. To create a new profile, click **Add** and type a unique name and description for the profile.

4. To create a new keystore for the profile, complete these steps:

    a. Click **System generate keystore**.

    b. In the **Provider** list box, select **IBMJCE**.

    c. In the **Algorithm** list box, select **JCE** and select whether to generate 512-bit or 1024-bit encryption keys.

       In the **Keystore Certificate CN** box, the fully qualified domain name for the JSA server is auto-populated.

    d. In the **Keystore password** box, enter the keystore password.

The keystore password is required to protect the integrity of the keystore. The password must be at least 8 characters in length.

e. In the **Verify keystore password**, retype the password.

5. To use an existing keystore with the profile, complete these steps:

    a. Click **Upload keystore**.

    b. Click **Browse** and select the keystore file.

    c. In the **Keystore password** box, type the password for the keystore.

6. Click **Submit**.

7. Download the keystore. Remove the keystore from your system and store it in a secure location.

"Creating Data Obfuscation Expressions" on page 479 that target the data that you want to hide.

## Creating Data Obfuscation Expressions

The data obfuscation profile uses expressions to specify which data to hide from JSA users. The expressions can use either field-based properties or regular expressions.

After an expression is created, you cannot change the type. For example, you cannot create a property-based expression and then later change it to a regular expression.

You cannot hide a normalized numeric field, such as port number or an IP address.

Multiple expressions that hide the same data cause data to be hidden twice. To decrypt data that is obfuscated multiple times, each keystore that is used in the obfuscation process must be applied in the order that the obfuscation occurred.

1. On the navigation menu ( ≡ ), click **Admin**.

2. In the **Data Sources** section, click **Data Obfuscation Management**.

3. Click the profile that you want to configure, and click **View Contents**.

    You cannot configure profiles that are locked.

4. To create a new data obfuscation expression, click **Add** and type a unique name and description for the profile.

5. Select the **Enabled** check box to enable the profile.

6. Optional: To apply the obfuscation expression to specific domains or tenants, select them from the Domain field. Or select **All Domains** to apply the obfuscation expression to all domains and tenants.

7. To create a field-based expression, click **Field Based** and select the field type to obfuscate.

8. To create a regular expression, click **RegEx** and configure the regex properties.

9. Click **Save**.

## Deobfuscating Data So That It Can Be Viewed in the Console

When data obfuscation is configured on an JSA system, the masked version of the data is shown throughout the application. You must have both the corresponding keystore and the password to deobfuscate the data so that it can be viewed.

You must be an administrator and have the private key and the password for the key before you can deobfuscate data. The private key must be on your local computer.

Before you can see the obfuscated data, you must upload the private key. After the key is uploaded, it remains available on the system for the duration of the current session. The session ends when you log out of JSA, when the cache is cleared on the JSA console, or when there is an extended period of inactivity. When the session ends, the private keys that were uploaded in the previous session are no longer visible.

JSA can use the keys available in the current session to automatically deobfuscate data. With auto-deobfuscation enabled, you do not have to repeatedly select the private key on the **Obfuscation Session Key** window each time that you want to view the data. Auto-deobfuscate is automatically disabled when the current session ends.

1. On the **Event Details** page, find the data that you want to deobfuscate.

2. To deobfuscate identity-based data:

   a. Click the lock icon next to the data that you want to deobfuscate.

   b. In the **Upload Key** section, click **Select File** and select the keystore to upload.

   c. In the **Password** box, type the password that matches the keystore.

   d. Click **Upload**.

      The **Deobfuscation** window shows the event payload, the profile names that are associated with the keystore, the obfuscated text, and the deobfuscated text.

   e. Optional:Click **Toggle Auto Deobfuscate** to enable auto-deobfuscation.

      After you toggle the auto-deobfuscation setting, you must refresh the browser window and reload the event details page for the changes to appear.

3. To deobfuscate payload data that is not identity-based:

   a. On the toolbar on the **Event Details** page, click **Obfuscation >Deobfuscation keys**.

   b. In the **Upload Key** section, click **Select File** and select the private key to upload.

   c. In the **Password** box, type the password that matches the private key and click **Upload**.

   d. In the **Payload information** box, select and copy the obfuscated text to the clipboard.

   e. On the toolbar on the **Event Details** page, click **Obfuscation >Deobfuscation**.

   f. Paste the obfuscated text in to dialog box.

   g. Select the obfuscation profile from the drop-down list and click **Deobfuscate**.

## Editing or Disabling Obfuscation Expressions Created in Previous Releases

When you upgrade to JSA 2014.6, data obfuscation expressions that were created in previous releases are automatically carried forward and continue to obfuscate data. These expressions appear in a single data obfuscation profile, named **AutoGeneratedProperty**.

Although you can see the expressions, you cannot edit or disable data obfuscation expressions that were created in earlier versions. You must manually disable them and create a data obfuscation profile that contains the revised expressions.

To disable an old expression, you must edit the `xml` configuration file that defines the attributes for the expression. You can then run the `obfuscation_updater.sh` script to disable it.

Ensure that you disable old expressions before you create new expressions that obfuscate the same data. Multiple expressions that obfuscate the same data cause the data to be obfuscated twice. To decrypt data that is obfuscated multiple times, each keystore that is used in the obfuscation process must be applied in the order that the obfuscation occurred.

1. Use SSH to log in to your JSA Console as the root user.

2. Edit the obfuscation expressions `.xml` configuration file that you created when you configured the expressions.

3. For each expression that you want to disable, change the **Enabled** attribute to **false**.

4. To disable the expressions, run the **obfuscation_updater.sh** script by typing the following command:

   **obfuscation_updater.sh [-p *<path_to_private_key>*] [-e *<path_to_obfuscation_xml_config_file>*]**

   The **obfuscation_updater.sh** script is in the **/opt/qradar/bin** directory, but you can run the script from any directory on your JSA Console.

"Creating Data Obfuscation Expressions" on page 479 to obfuscate data and manage obfuscation expressions directly in JSA.

## RELATED DOCUMENTATION

How Does Data Obfuscation Work? | **472**

Data Obfuscation Profiles | **474**

Data Obfuscation Expressions | **475**

24

CHAPTER

# Log Files

# Log Files

Operations performed in JSA are recorded in log files for tracking purposes. Log files can help you troubleshoot problems by recording the activities that take place when you work with a product.

The following log files can help you identify and resolve problems when they occur:

- **/var/log/qradar.log**
- **/var/log/qradar.error**
- **/var/log/qradar-sql.log**
- **/opt/tomcat6/logs/catalina.out**
- **/var/log/qflow.debug**

If you want to collect the JSA log files and review them later, see "Collecting Log Files" on page 117.

# Audit Logs

**IN THIS SECTION**

Changes that are made by JSA users are recorded in the audit logs.

All audit logs are stored in plain text and are archived and compressed when the audit log file reaches 50 MB. The current log file is named **audit.log**. When the file reaches 50 MB, the file is compressed and renamed to **audit.1.gz**. The file number increments each time that a log file is archived. JSA stores up to 25 archived log files.

Audit log data is also stored in the `SIM Audit-2` log source, which can be used for filtering and reporting to track how users interact with JSA. The data retention is determined by your event retention configuration.

## Viewing the Audit Log File

Use Secure Shell (SSH) to log in to your JSA system and monitor changes to your system.

You can use **Log Activity** tab to view normalized audit log events.

The maximum size of any audit message, excluding date, time, and host name, is 1024 characters.

Each entry in the log file displays by using the following format:

***<date_time> <host name> <user>@<IP address> (thread ID) [<category>] [<sub-category>] [<action>] <payload>***

The following table describes the log file format options.

**Table 84: Description Of the Parts Of the Log File Format**

| File format part | Description |
|---|---|
| *date_time* | The date and time of the activity in the format: Month Date HH:MM:SS |
| *host name* | The host name of the Console where this activity was logged. |
| *user* | The name of the user who changed the settings. |
| *IP address* | The IP address of the user who changed the settings. |
| *thread ID)* | The identifier of the Java thread that logged this activity. |
| *category* | The high-level category of this activity. |
| *sub-categor* | The low-level category of this activity. |
| *action* | The activity that occurred. |
| *payload* | The complete record, which might include the user record or event rule, that changed. |

1. Using SSH, log in to JSA as the root user:
2. **User Name: root**

3. Password: ***password***

4. Go to the following directory:

   **/var/log/audit**

5. Open and view the audit log file.

## Creating Reports from Audit Log Searches in JSA

To help you track how users interact with JSA, create reports that are based on your search results.

1. Click **Log Activity** > **Add Filter**.

2. In the **Add Filter** window, configure the following settings:

   **Table 85: Settings to Configure**

   | Settings to configure | Value |
   |---|---|
   | **Parameter** | Log Source [Indexed] |
   | **Operator** | Equals |
   | **Log Source** | SIM Audit-2 |

3. Click **Add Filter**.

4. If events are streaming into the **Log Activity** tab, click **Pause**.

5. From the **View** list, select a time interval.

6. To save the search, click **Save Criteria**, provide a name for the search, and then click **OK**.

7. To generate a report from your search result, follow these steps:

   a. From the **Reports** tab, click **Actions** > **Create**.

   b. Follow the report wizard.

   c. In the **Saved Searches** field, type the name of the search that you created for the SIM audit log source.

d. Click **Save Container Details**.

e. Finish the report wizard pages.

## Logged Actions

The JSAr audit logs are in the **/var/log/audit** directory.

The following list describes the categories of actions that are in the audit log file:

- **Administrator Authentication--**
  - Log in to the Administration Console.
  - Log out of the Administration Console.
- **Assets--**
  - Delete an asset.
  - Delete all assets.
- **Audit Log Access--**A search that includes events that have a high-level event category of Audit.
- **Backup and Recovery--**
  - Edit the configuration.
  - Initiate the backup.
  - Complete the backup.
  - Fail the backup.
  - Delete the backup.
  - Synchronize the backup.
  - Cancel the backup.
  - Initiate the restore.
  - Upload a backup.
  - Upload an invalid backup.
  - Initiate the restore.

  - Purge the backup.

- **Chart Configuration--**Save flow or event chart configuration.

- **Content Management--**

  - Content export initiated.

  - Content export complete.

  - Content import initiated.

  - Content import complete.

  - Content update initiated.

  - Content update complete.

  - Content search initiated.

  - Applications added.

  - Applications modified.

  - Custom actions added.

  - Custom actions modified.

  - Ariel property added.

  - Ariel property modified.

  - Ariel property expression added.

  - Ariel property expression modified.

  - CRE rule added.

  - CRE rule modified.

  - Dashboard added.

  - Dashboard modified.

  - Device extension added.

  - Device extension modified.

  - Device extension association modified.

  - Grouping added.

  - Grouping modified.
  - Historical correlation profile added.
  - Historical correlation profile modified.
  - QID map entry added.
  - QID map entry modified.
  - Reference data created.
  - Reference data updated.
  - Security profile added.
  - Security profile modified.
  - Sensor device added.
  - Sensor device modified.
- **Custom Properties--**
  - Add a custom event property.
  - Edit a custom event property.
  - Delete a custom event property.
  - Edit a custom flow property.
  - Delete a custom flow property.
- **Custom Property Expressions--**
  - Add a custom event property expression.
  - Edit a custom event property expression.
  - Delete a custom event property expression.
  - Add a custom flow property expression.
  - Edit a custom flow property expression.
  - Delete a custom flow property expression.
- **Flow Sources--**
  - Add a flow source.

  - Edit a flow source.
  - Delete a flow source.
- **Groups--**
  - Add a group.
  - Delete a group.
  - Edit a group.
- **Historical Correlation--**
  - Add a historical correlation profile.
  - Delete a historical correlation profile.
  - Modify a historical correlation profile.
  - Enable a historical correlation profile.
  - Disable a historical correlation profile.
  - Historical correlation profile is running.
  - Historical correlation profile is canceled.
- **Licensing--**
  - Add a license key.
  - Delete a license key.
  - Delete license pool allocation.
  - Update license pool allocation.
- **Log Source Extension--**
  - Add an log source extension.
  - Edit the log source extension.
  - Delete a log source extension.
  - Upload a log source extension.
  - Upload a log source extension successfully.
  - Upload an invalid log source extension.

  - Download a log source extension.
  - Report a log source extension.
  - Modify a log sources association to a device or device type.
- Offenses--
  - Create an offense.
  - Hide an offense.
  - Close an offense.
  - Close all offenses.
  - Add a destination note.
  - Add a source note.
  - Add a network note.
  - Add an offense note.
  - Add a reason for closing offenses.
  - Edit a reason for closing offenses.
- Protocol Configuration--
  - Add a protocol configuration.
  - Delete a protocol configuration.
  - Edit a protocol configuration.
- QIDmap--
  - Add a QID map entry.
  - Edit a QID map entry.
- JSA Vulnerability Manager --
  - Create a scanner schedule.
  - Update a scanner schedule.
  - Delete a scanner schedule.
  - Start a scanner schedule.

  - Pause a scanner schedule.
  - Resume a scanner schedule.
- **Reference Sets--**
  - Create a reference set.
  - Edit a reference set.
  - Purge elements in a reference set.
  - Delete a reference set.
  - Add reference set elements.
  - Delete reference set elements.
  - Delete all reference set elements.
  - Import reference set elements.
  - Export reference set elements.
- **Reports--**
  - Add a template.
  - Delete a template.
  - Edit a template.
  - Generate a report.
  - Delete a report.
  - Delete generated content.
  - View a generated report.
  - Email a generated report.
- **Retention Buckets--**
  - Add a bucket.
  - Delete a bucket.
  - Edit a bucket.
  - Enable or disable a bucket.

- **Root Login--**
  - Log in to JSA, as root user.
  - Log out of JSA, as root user.
- **Rules--**
  - Add a rule.
  - Delete a rule.
  - Edit a rule.
- **Scanner--**
  - Add a scanner.
  - Delete a scanner.
  - Edit a scanner.
- **Scanner Schedule--**
  - Add a schedule.
  - Edit a schedule.
  - Delete a schedule.
- **Session Authentication--**
  - Create an administration session.
  - Terminate an administration session.
  - Deny an invalid authentication session.
  - Expire a session authentication.
  - Create an authentication session.
  - Terminate an authentication session.
- **SIM--**Clean a SIM model.
- **Store and Forward--**
  - Add a Store and Forward schedule.
  - Edit a Store and Forward schedule.

  - Delete a Store and Forward schedule.
- Syslog Forwarding--
  - Add a syslog forwarding.
  - Delete a syslog forwarding.
  - Edit a syslog forwarding.
- System Management--
  - Shut down a system.
  - Restart a system.
- User Accounts--
  - Add an account.
  - Edit an account.
  - Delete an account.
- User Authentication--
  - Log in to the user interface.
  - Log out of the user interface.
- User Authentication Ariel --
  - Deny a login attempt.
  - Add an Ariel property.
  - Delete an Ariel property.
  - Edit an Ariel property.
  - Add an Ariel property extension.
  - Delete an Ariel property extension.
  - Edit an Ariel property extension.
- User Roles--
  - Add a role.
  - Edit a role.

  - Delete a role.

- VIS--

  - Discover a new host.

  - Discover a new operating system.

  - Discover a new port.

  - Discover a new vulnerability.

## RELATED DOCUMENTATION

Log Files | 484

25

CHAPTER

# Event Categories

# Event Categories

Event categories are used to group incoming events for processing by JSA. The event categories are searchable and help you monitor your network.

Events that occur on your network are aggregated into high-level and low-level categories. Each high-level category contains low-level categories and an associated severity level and ID number.

You can review the severity levels that are assigned to events and adjust them to suit your corporate policy needs.

You can run an AQL query by using high-level and low-level event category IDs. The category IDs for the associated category names can be retrieved from the event category tables.

For example, if you are developing applications on JSA, you can run an AQL search similar to the following query from the command line, to gather data from Ariel:

```
select qidname(qid) as 'Event', username as 'Username', devicetime as 'Time' from events where '<high-level category ID>' and '<Low-level category ID>' and LOGSOURCENAME(logsourceid) like "%Low-level category name%" last 3 days
```

# High-level Event Categories

Events in JSA log sources are grouped into high-level categories. Each event is assigned to a specific high-level category.

Categorizing the incoming events ensures that you can easily search the data.

The following table describes the high-level event categories.

**Table 86: High-level Event Categories**

| Category | Category ID | Description |
|---|---|---|
| "Recon" on page 501 | 1000 | Events that are related to scanning and other techniques that are used to identify network resources, for example, network or host port scans. |

Table 86: High-level Event Categories *(Continued)*

| Category | Category ID | Description |
|---|---|---|
| "DoS" on page 503 | 2000 | Events that are related to denial-of-service (DoS) or distributed denial-of-service (DDoS) attacks against services or hosts, for example, brute force network DoS attacks. |
| "Authentication" on page 510 | 3000 | Events that are related to authentication controls, group, or privilege change, for example, log in or log out. |
| "Access" on page 523 | 4000 | Events resulting from an attempt to access network resources, for example, firewall accept or deny. |
| "Exploit" on page 528 | 5000 | Events that are related to application exploits and buffer overflow attempts, for example, buffer overflow or web application exploits. |
| "Malware" on page 531 | 6000 | Events that are related to viruses, trojans, back door attacks, or other forms of hostile software. Malware events might include a virus, trojan, malicious software, or spyware. |
| "Suspicious Activity" on page 534 | 7000 | The nature of the threat is unknown but behavior is suspicious. The threat might include protocol anomalies that potentially indicate evasive techniques, for example, packet fragmentation or known intrusion detection system (IDS) evasion techniques. |
| "System" on page 542 | 8000 | Events that are related to system changes, software installation, or status messages. |
| "Policy" on page 551 | 9000 | Events regarding corporate policy violations or misuse. |
| "Unknown" on page 553 | 10000 | Events that are related to unknown activity on your system. |
| "CRE" on page 555 | 12000 | Events that are generated from an offense or event rule. |

Table 86: High-level Event Categories *(Continued)*

| Category | Category ID | Description |
|---|---|---|
| "Potential Exploit" on page 556 | 13000 | Events relate to potential application exploits and buffer overflow attempts. |
| "Flow" on page 558 | 14000 | Events that are related to flow actions. |
| "User Defined" on page 561 | 15000 | Events that are related to user-defined objects. |
| "SIM Audit" on page 566 | 16000 | Events that are related to user interaction with the Console and administrative functions. |
| "VIS Host Discovery" on page 568 | 17000 | Events that are related to the host, ports, or vulnerabilities that the VIS component discovers. |
| "Application" on page 569 | 18000 | Events that are related to application activity. |
| "Audit" on page 612 | 19000 | Events that are related to audit activity. |
| "Risk" on page 619 | 20000 | Events that are related to risk activity in JSA Risk Manager. |
| "Risk Manager Audit" on page 621 | 21000 | Events that are related to audit activity in JSA Risk Manager. |
| "Control" on page 622 | 22000 | Events that are related to your hardware system. |
| "Asset Profiler" on page 626 | 23000 | Events that are related to asset profiles. |
| "Sense" on page 635 | 24000 | Events that are related to UBA. |

RELATED DOCUMENTATION

Recon | 501

DoS | 503

Authentication | 510

# Recon

The Recon category contains events that are related to scanning and other techniques that are used to identify network resources.

The following table describes the low-level event categories and associated severity levels for the Recon category.

**Table 87: Low-level Categories and Severity Levels for the Recon Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Form of Recon | 1001 | An unknown form of reconnaissance. | 2 |
| Application Query | 1002 | Reconnaissance to applications on your system. | 3 |
| Host Query | 1003 | Reconnaissance to a host in your network. | 3 |
| Network Sweep | 1004 | Reconnaissance on your network. | 4 |
| Mail Reconnaissance | 1005 | Reconnaissance on your mail system. | 3 |
| Windows Reconnaissance | 1006 | Reconnaissance for Windows operating system. | 3 |

**Table 87: Low-level Categories and Severity Levels for the Recon Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Portmap / RPC r\Request | 1007 | Reconnaissance on your portmap or RPC request. | 3 |
| Host Port Scan | 1008 | Indicates that a scan occurred on the host ports. | 4 |
| RPC Dump | 1009 | Indicates that Remote Procedure Call (RPC) information is removed. | 3 |
| DNS Reconnaissance | 1010 | Reconnaissance on the DNS server. | 3 |
| Misc Reconnaissance Event | 1011 | Miscellaneous reconnaissance event. | 2 |
| Web Reconnaissance | 1012 | Web reconnaissance on your network. | 3 |
| Database Reconnaissance | 1013 | Database reconnaissance on your network. | 3 |
| ICMP Reconnaissance | 1014 | Reconnaissance on ICMP traffic. | 3 |
| UDP Reconnaissance | 1015 | Reconnaissance on UDP traffic. | 3 |
| SNMP Reconnaissance | 1016 | Reconnaissance on SNMP traffic. | 3 |
| ICMP Host Query | 1017 | Indicates an ICMP host query. | 3 |

Table 87: Low-level Categories and Severity Levels for the Recon Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| UDP Host Query | 1018 | Indicates a UDP host query. | 3 |
| NMAP Reconnaissance | 1019 | Indicates NMAP reconnaissance. | 3 |
| TCP Reconnaissance | 1020 | Indicates TCP reconnaissance on your network. | 3 |
| UNIX Reconnaissance | 1021 | Reconnaissance on your UNIX network. | 3 |
| FTP Reconnaissance | 1022 | Indicates FTP reconnaissance. | 3 |

RELATED DOCUMENTATION

DoS | 503

Authentication | 510

Access | 523

# DoS

The DoS category contains events that are related to denial-of-service (DoS) attacks against services or hosts.

The following table describes the low-level event categories and associated severity levels for the DoS category.

**Table 88: Low-level Categories and Severity Levels for the DoS Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown DoS Attack | 2001 | Indicates an unknown DoS attack. | 8 |
| ICMP DoS | 2002 | Indicates an ICMP DoS attack. | 9 |
| TCP DoS | 2003 | Indicates a TCP DoS attack. | 9 |
| UDP DoS | 2004 | Indicates a UDP DoS attack. | 9 |
| DNS Service DoS | 2005 | Indicates a DNS service DoS attack. | 8 |
| Web Service DoS | 2006 | Indicates a web service DoS attack. | 8 |
| Mail Service DoS | 2007 | Indicates a mail server DoS attack. | 8 |
| Distributed DoS | 2008 | Indicates a distributed DoS attack. | 9 |
| Misc DoS | 2009 | Indicates a miscellaneous DoS attack. | 8 |
| UNIX DoS | 2010 | Indicates a UNIX DoS attack. | 8 |
| Windows DoS | 2011 | Indicates a Windows DoS attack. | 8 |

Table 88: Low-level Categories and Severity Levels for the DoS Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Database DoS | 2012 | Indicates a database DoS attack. | 8 |
| FTP DoS | 2013 | Indicates an FTP DoS attack. | 8 |
| Infrastructure DoS | 2014 | Indicates a DoS attack on the infrastructure. | 8 |
| Telnet DoS | 2015 | Indicates a Telnet DoS attack. | 8 |
| Brute Force Login | 2016 | Indicates access to your system through unauthorized methods. | 8 |
| High Rate TCP DoS | 2017 | Indicates a high rate TCP DoS attack. | 8 |
| High Rate UDP DoS | 2018 | Indicates a high rate UDP DoS attack. | 8 |
| High Rate ICMP DoS | 2019 | Indicates a high rate ICMP DoS attack. | 8 |
| High Rate DoS | 2020 | Indicates a high rate DoS attack. | 8 |
| Medium Rate TCP DoS | 2021 | Indicates a medium rate TCP attack. | 8 |
| Medium Rate UDP DoS | 2022 | Indicates a medium rate UDP attack. | 8 |

Table 88: Low-level Categories and Severity Levels for the DoS Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Medium Rate ICMP DoS | 2023 | Indicates a medium rate ICMP attack. | 8 |
| Medium Rate DoS | 2024 | Indicates a medium rate DoS attack. | 8 |
| Low Rate TCP DoS | 2025 | Indicates a low rate TCP DoS attack. | 8 |
| Low Rate UDP DoS | 2026 | Indicates a low rate UDP DoS attack. | 8 |
| Low Rate ICMP DoS | 2027 | Indicates a low rate ICMP DoS attack. | 8 |
| Low Rate DoS | 2028 | Indicates a low rate DoS attack. | 8 |
| Distributed High Rate TCP DoS | 2029 | Indicates a distributed high rate TCP DoS attack. | 8 |
| Distributed High Rate UDP DoS | 2030 | Indicates a distributed high rate UDP DoS attack. | 8 |
| Distributed High Rate ICMP DoS | 2031 | Indicates a distributed high rate ICMP DoS attack. | 8 |
| Distributed High Rate DoS | 2032 | Indicates a distributed high rate DoS attack. | 8 |

Table 88: Low-level Categories and Severity Levels for the DoS Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Distributed Medium Rate TCP DoS | 2033 | Indicates a distributed medium rate TCP DoS attack. | 8 |
| Distributed Medium Rate UDP DoS | 2034 | Indicates a distributed medium rate UDP DoS attack. | 8 |
| Distributed Medium Rate ICMP DoS | 2035 | Indicates a distributed medium rate ICMP DoS attack. | 8 |
| Distributed Medium Rate DoS | 2036 | Indicates a distributed medium rate DoS attack. | 8 |
| Distributed Low Rate TCP DoS | 2037 | Indicates a distributed low rate TCP DoS attack. | 8 |
| Distributed Low Rate UDP DoS | 2038 | Indicates a distributed low rate UDP DoS attack. | 8 |
| Distributed Low Rate ICMP DoS | 2039 | Indicates a distributed low rate ICMP DoS attack. | 8 |
| Distributed Low Rate DoS | 2040 | Indicates a distributed low rate DoS attack. | 8 |
| High Rate TCP Scan | 2041 | Indicates a high rate TCP scan. | 8 |
| High Rate UDP Scan | 2042 | Indicates a high rate UDP scan. | 8 |

**Table 88: Low-level Categories and Severity Levels for the DoS Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| High Rate ICMP Scan | 2043 | Indicates a high rate ICMP scan. | 8 |
| High Rate Scan | 2044 | Indicates a high rate scan. | 8 |
| Medium Rate TCP Scan | 2045 | Indicates a medium rate TCP scan. | 8 |
| Medium Rate UDP Scan | 2046 | Indicates a medium rate UDP scan. | 8 |
| Medium Rate ICMP Scan | 2047 | Indicates a medium rate ICMP scan. | 8 |
| Medium Rate Scan | 2048 | Indicates a medium rate scan. | 8 |
| Low Rate TCP Scan | 2049 | Indicates a low rate TCP scan. | 8 |
| Low Rate UDP Scan | 2050 | Indicates a low rate UDP scan. | 8 |
| Low Rate ICMP Scan | 2051 | Indicates a low rate ICMP scan. | 8 |
| Low Rate Scan | 2052 | Indicates a low rate scan. | 8 |
| VoIP DoS | 2053 | Indicates a VoIP DoS attack. | 8 |
| Flood | 2054 | Indicates a Flood attack. | 8 |

Table 88: Low-level Categories and Severity Levels for the DoS Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| TCP Flood | 2055 | Indicates a TCP flood attack. | 8 |
| UDP Flood | 2056 | Indicates a UDP flood attack. | 8 |
| ICMP Flood | 2057 | Indicates an ICMP flood attack. | 8 |
| SYN Flood | 2058 | Indicates a SYN flood attack. | 8 |
| URG Flood | 2059 | Indicates a flood attack with the urgent (URG) flag on. | 8 |
| SYN URG Flood | 2060 | Indicates a SYN flood attack with the urgent (URG) flag on. | 8 |
| SYN FIN Flood | 2061 | Indicates a SYN FIN flood attack. | 8 |
| SYN ACK Flood | 2062 | Indicates a SYN ACK flood attack. | 8 |

## RELATED DOCUMENTATION

Authentication | 510

Access | 523

Exploit | 528

# Authentication

The authentication category contains events that are related to authentication, sessions, and access controls that monitor users on the network.

The following table describes the low-level event categories and associated severity levels for the authentication category.

**Table 89: Low-level Categories and Severity Levels for the Authentication Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Authentication | 3001 | Indicates unknown authentication. | 1 |
| Host Login Succeeded | 3002 | Indicates a successful host login. | 1 |
| Host Login Failed | 3003 | Indicates that the host login failed. | 3 |
| Misc Login Succeeded | 3004 | Indicates that the login sequence succeeded. | 1 |
| Misc Login Failed | 3005 | Indicates that login sequence failed. | 3 |
| Privilege Escalation Failed | 3006 | Indicates that the privileged escalation failed. | 3 |
| Privilege Escalation Succeeded | 3007 | Indicates that the privilege escalation succeeded. | 1 |
| Mail Service Login Succeeded | 3008 | Indicates that the mail service login succeeded. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Mail Service Login Failed | 3009 | Indicates that the mail service login failed. | 3 |
| Auth Server Login Failed | 3010 | Indicates that the authentication server login failed. | 3 |
| Auth Server Login Succeeded | 3011 | Indicates that the authentication server login succeeded. | 1 |
| Web Service Login Succeeded | 3012 | Indicates that the web service login succeeded. | 1 |
| Web Service Login Failed | 3013 | Indicates that the web service login failed. | 3 |
| Admin Login Successful | 3014 | Indicates that an administrative login was successful. | 1 |
| Admin Login Failure | 3015 | Indicates the administrative login failed. | 3 |
| Suspicious Username | 3016 | Indicates that a user attempted to access the network by using an incorrect user name. | 4 |
| Login with username/ password defaults successful | 3017 | Indicates that a user accessed the network by using the default user name and password. | 4 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Login with username/ password defaults failed | 3018 | Indicates that a user was unsuccessful accessing the network by using the default user name and password. | 4 |
| FTP Login Succeeded | 3019 | Indicates that the FTP login was successful. | 1 |
| FTP Login Failed | 3020 | Indicates that the FTP login failed. | 3 |
| SSH Login Succeeded | 3021 | Indicates that the SSH login was successful. | 1 |
| SSH Login Failed | 3022 | Indicates that the SSH login failed. | 2 |
| User Right Assigned | 3023 | Indicates that user access to network resources was successfully granted. | 1 |
| User Right Removed | 3024 | Indicates that user access to network resources was successfully removed. | 1 |
| Trusted Domain Added | 3025 | Indicates that a trusted domain was successfully added to your deployment. | 1 |
| Trusted Domain Removed | 3026 | Indicates that a trusted domain was removed from your deployment. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| System Security Access Granted | 3027 | Indicates that system security access was successfully granted. | 1 |
| System Security Access Removed | 3028 | Indicates that system security access was successfully removed. | 1 |
| Policy Added | 3029 | Indicates that a policy was successfully added. | 1 |
| Policy Change | 3030 | Indicates that a policy was successfully changed. | 1 |
| User Account Added | 3031 | Indicates that a user account was successfully added. | 1 |
| User Account Changed | 3032 | Indicates a change to an existing user account. | 1 |
| Password Change Failed | 3033 | Indicates that an attempt to change an existing password failed. | 3 |
| Password Change Succeeded | 3034 | Indicates that a password change was successful. | 1 |
| User Account Removed | 3035 | Indicates that a user account was successfully removed. | 1 |
| Group Member Added | 3036 | Indicates that a group member was successfully added. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Group Member Removed | 3037 | Indicates that a group member was removed. | 1 |
| Group Added | 3038 | Indicates that a group was successfully added. | 1 |
| Group Changed | 3039 | Indicates a change to an existing group. | 1 |
| Group Removed | 3040 | Indicates that a group was removed. | 1 |
| Computer Account Added | 3041 | Indicates that a computer account was successfully added. | 1 |
| Computer Account Changed | 3042 | Indicates a change to an existing computer account. | 1 |
| Computer Account Removed | 3043 | Indicates that a computer account was successfully removed. | 1 |
| Remote Access Login Succeeded | 3044 | Indicates that access to the network by using a remote login was successful. | 1 |
| Remote Access Login Failed | 3045 | Indicates that an attempt to access the network by using a remote login failed. | 3 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| General Authentication Successful | 3046 | Indicates that the authentication processes was successful. | 1 |
| General Authentication Failed | 3047 | Indicates that the authentication process failed. | 3 |
| Telnet Login Succeeded | 3048 | Indicates that the telnet login was successful. | 1 |
| Telnet Login Failed | 3049 | Indicates that the telnet login failed. | 3 |
| Suspicious Password | 3050 | Indicates that a user attempted to log in by using a suspicious password. | 4 |
| Samba Login Successful | 3051 | Indicates that a user successfully logged in by using Samba. | 1 |
| Samba Login Failed | 3052 | Indicates a user failed to log in by using Samba. | 3 |
| Auth Server Session Opened | 3053 | Indicates that a communication session with the authentication server was started. | 1 |
| Auth Server Session Closed | 3054 | Indicates that a communication session with the authentication server was closed. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Firewall Session Closed | 3055 | Indicates that a firewall session was closed. | 1 |
| Host Logout | 3056 | Indicates that a host successfully logged out. | 1 |
| Misc Logout | 3057 | Indicates that a user successfully logged out. | 1 |
| Auth Server Logout | 3058 | Indicates that the process to log out of the authentication server was successful. | 1 |
| Web Service Logout | 3059 | Indicates that the process to log out of the web service was successful. | 1 |
| Admin Logout | 3060 | Indicates that the administrative user successfully logged out. | 1 |
| FTP Logout | 3061 | Indicates that the process to log out of the FTP service was successful. | 1 |
| SSH Logout | 3062 | Indicates that the process to log out of the SSH session was successful. | 1 |
| Remote Access Logout | 3063 | Indicates that the process to log out using remote access was successful. | 1 |

**Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Telnet Logout | 3064 | Indicates that the process to log out of the Telnet session was successful. | 1 |
| Samba Logout | 3065 | Indicates that the process to log out of Samba was successful. | 1 |
| SSH Session Started | 3066 | Indicates that the SSH login session was initiated on a host. | 1 |
| SSH Session Finished | 3067 | Indicates the termination of an SSH login session on a host. | 1 |
| Admin Session Started | 3068 | Indicates that a login session was initiated on a host by an administrative or privileged user. | 1 |
| Admin Session Finished | 3069 | Indicates the termination of an administrator or privileged users login session on a host. | 1 |
| VoIP Login Succeeded | 3070 | Indicates a successful VoIP service login | 1 |
| VoIP Login Failed | 3071 | Indicates an unsuccessful attempt to access VoIP service. | 1 |
| VoIP Logout | 3072 | Indicates a user logout, | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| VoIP Session Initiated | 3073 | Indicates the beginning of a VoIP session. | 1 |
| VoIP Session Terminated | 3074 | Indicates the end of a VoIP session. | 1 |
| Database Login Succeeded | 3075 | Indicates a successful database login. | 1 |
| Database Login Failure | 3076 | Indicates a database login attempt failed. | 3 |
| IKE Authentication Failed | 3077 | Indicates a failed Internet Key Exchange (IKE) authentication was detected. | 3 |
| IKE Authentication Succeeded | 3078 | Indicates that a successful IKE authentication was detected. | 1 |
| IKE Session Started | 3079 | Indicates that an IKE session started. | 1 |
| IKE Session Ended | 3080 | Indicates that an IKE session ended. | 1 |
| IKE Error | 3081 | Indicates an IKE error message. | 1 |
| IKE Status | 3082 | Indicates IKE status message. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| RADIUS Session Started | 3083 | Indicates that a RADIUS session started. | 1 |
| RADIUS Session Ended | 3084 | Indicates a RADIUS session ended. | 1 |
| RADIUS Session Denied | 3085 | Indicates that a RADIUS session was denied. | 1 |
| RADIUS Session Status | 3086 | Indicates a RADIUS session status message. | 1 |
| RADIUS Authentication Failed | 3087 | Indicates a RADIUS authentication failure. | 3 |
| RADIUS Authentication Successful | 3088 | Indicates a RADIUS authentication succeeded. | 1 |
| TACACS Session Started | 3089 | Indicates a TACACS session started. | 1 |
| TACACS Session Ended | 3090 | Indicates a TACACS session ended. | 1 |
| TACACS Session Denied | 3091 | Indicates that a TACACS session was denied. | 1 |
| TACACS Session Status | 3092 | Indicates a TACACS session status message. | 1 |
| TACACS Authentication Successful | 3093 | Indicates a TACACS authentication succeeded. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| TACACS Authentication Failed | 3094 | Indicates a TACACS authentication failure. | 1 |
| Deauthenticating Host Succeeded | 3095 | Indicates that the deauthentication of a host was successful. | 1 |
| Deauthenticating Host Failed | 3096 | Indicates that the deauthentication of a host failed. | 3 |
| Station Authentication Succeeded | 3097 | Indicates that the station authentication was successful. | 1 |
| Station Authentication Failed | 3098 | Indicates that the station authentication of a host failed. | 3 |
| Station Association Succeeded | 3099 | Indicates that the station association was successful. | 1 |
| Station Association Failed | 3100 | Indicates that the station association failed. | 3 |
| Station Reassociation Succeeded | 3101 | Indicates that the station reassociation was successful. | 1 |
| Station Reassociation Failed | 3102 | Indicates that the station association failed. | 3 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Disassociating Host Succeeded | 3103 | Indicates that the disassociating a host was successful. | 1 |
| Disassociating Host Failed | 3104 | Indicates that the disassociating a host failed. | 3 |
| SA Error | 3105 | Indicates a Security Association (SA) error message. | 5 |
| SA Creation Failure | 3106 | Indicates a Security Association (SA) creation failure. | 3 |
| SA Established | 3107 | Indicates that a Security Association (SA) connection established. | 1 |
| SA Rejected | 3108 | Indicates that a Security Association (SA) connection rejected. | 3 |
| Deleting SA | 3109 | Indicates the deletion of a Security Association (SA). | 1 |
| Creating SA | 3110 | Indicates the creation of a Security Association (SA). | 1 |
| Certificate Mismatch | 3111 | Indicates a certificate mismatch. | 3 |
| Credentials Mismatch | 3112 | Indicates a credentials mismatch. | 3 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Admin Login Attempt | 3113 | Indicates an admin login attempt. | 2 |
| User Login Attempt | 3114 | Indicates a user login attempt. | 2 |
| User Login Successful | 3115 | Indicates a successful user login. | 1 |
| User Login Failure | 3116 | Indicates a failed user login. | 3 |
| SFTP Login Succeeded | 3117 | Indicates a successful SSH File Transfer Protocol (SFTP) login. | 1 |
| SFTP Login Failed | 3118 | Indicates a failed SSH File Transfer Protocol (SFTP) login. | 3 |
| SFTP Logout | 3119 | Indicates an SSH File Transfer Protocol (SFTP) logout. | 1 |
| Identity Granted | 3120 | Indicates that an identity was granted. | 1 |
| Identity Removed | 3121 | Indicates that an identity was removed. | 1 |
| Identity Revoked | 3122 | Indicates that an identity was revoked. | 1 |

Table 89: Low-level Categories and Severity Levels for the Authentication Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Policy Removed | 3123 | Indicates that a policy was removed. | 1 |
| User Account Lock | 3124 | Indicates that a user account was locked. | 1 |
| User Account Unlock | 3125 | Indicates that a user account was unlocked | 1 |
| User Account Expired | 3126 | Indicates that a user account is expired | 1 |

RELATED DOCUMENTATION

Access | 523
Exploit | 528
Malware | 531

# Access

The access category contains authentication and access controls that are used for monitoring network events.

The following table describes the low-level event categories and associated severity levels for the access category.

**Table 90: Low-level Categories and Severity Levels for the Access Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Network Communication Event | 4001 | Indicates an unknown network communication event. | 3 |
| Firewall Permit | 4002 | Indicates that access to the firewall was allowed. | 0 |
| Firewall Deny | 4003 | Indicates that access to the firewall was denied. | 4 |
| Flow Context Response (JSA only) | 4004 | Indicates events from the Classification Engine in response to a SIM request. | 5 |
| Misc Network Communication Event | 4005 | Indicates a miscellaneous communications event. | 3 |
| IPS Deny | 4006 | Indicates Intrusion Prevention Systems (IPS) denied traffic. | 4 |
| Firewall Session Opened | 4007 | Indicates that the firewall session was opened. | 0 |
| Firewall Session Closed | 4008 | Indicates that the firewall session was closed. | 0 |
| Dynamic Address Translation Successful | 4009 | Indicates that dynamic address translation was successful. | 0 |
| No Translation Group Found | 4010 | Indicates that no translation group was found. | 2 |

Table 90: Low-level Categories and Severity Levels for the Access Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Misc Authorization | 4011 | Indicates that access was granted to a miscellaneous authentication server. | 2 |
| ACL Permit | 4012 | Indicates that an Access Control List (ACL) allowed access. | 0 |
| ACL Deny | 4013 | Indicates that an Access Control List (ACL) denied access. | 4 |
| Access Permitted | 4014 | Indicates that access was allowed. | 0 |
| Access Denied | 4015 | Indicates that access was denied. | 4 |
| Session Opened | 4016 | Indicates that a session was opened. | 1 |
| Session Closed | 4017 | Indicates that a session was closed. | 1 |
| Session Reset | 4018 | Indicates that a session was reset. | 3 |
| Session Terminated | 4019 | Indicates that a session was allowed. | 4 |
| Session Denied | 4020 | Indicates that a session was denied. | 5 |

Table 90: Low-level Categories and Severity Levels for the Access Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Session in Progress | 4021 | Indicates that a session is in progress. | 1 |
| Session Delayed | 4022 | Indicates that a session was delayed. | 3 |
| Session Queued | 4023 | Indicates that a session was queued. | 1 |
| Session Inbound | 4024 | Indicates that a session is inbound. | 1 |
| Session Outbound | 4025 | Indicates that a session is outbound. | 1 |
| Unauthorized Access Attempt | 4026 | Indicates that an unauthorized access attempt was detected. | 6 |
| Misc Application Action Allowed | 4027 | Indicates that an application action was allowed. | 1 |
| Misc Application Action Denied | 4028 | Indicates that an application action was denied. | 3 |
| Database Action Allowed | 4029 | Indicates that a database action was allowed. | 1 |
| Database Action Denied | 4030 | Indicates that a database action was denied. | 3 |

Table 90: Low-level Categories and Severity Levels for the Access Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| FTP Action Allowed | 4031 | Indicates that an FTP action was allowed. | 1 |
| FTP Action Denied | 4032 | Indicates that an FTP action was denied. | 3 |
| Object Cached | 4033 | Indicates that an object was cached. | 1 |
| Object Not Cached | 4034 | Indicates that an object was not cached. | 1 |
| Rate Limiting | 4035 | Indicates that the network rate-limits traffic. | 4 |
| No Rate Limiting | 4036 | Indicates that the network does not rate-limit traffic. | 0 |
| P11 Access Permitted | 4037 | Indicates that P11 access is permitted. | 8 |
| P11 Access Denied | 4038 | Indicates that P11 access was attempted and denied. | 8 |
| IPS Permit | 4039 | Indicates an IPS permit. | 0 |

## RELATED DOCUMENTATION

Exploit | 528

Malware | 531

Suspicious Activity | 534

# Exploit

The exploit category contains events where a communication or an access exploit occurred.

The following table describes the low-level event categories and associated severity levels for the exploit category.

**Table 91: Low-level Categories and Severity Levels for the Exploit Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Exploit Attack | 5001 | Indicates an unknown exploit attack. | 9 |
| Buffer Overflow | 5002 | Indicates a buffer overflow. | 9 |
| DNS Exploit | 5003 | Indicates a DNS exploit. | 9 |
| Telnet Exploit | 5004 | Indicates a Telnet exploit. | 9 |
| Linux Exploit | 5005 | Indicates a Linux exploit. | 9 |
| UNIX Exploit | 5006 | Indicates a UNIX exploit. | 9 |
| Windows Exploit | 5007 | Indicates a MicrosoftWindows exploit. | 9 |
| Mail Exploit | 5008 | Indicates a mail server exploit. | 9 |
| Infrastructure Exploit | 5009 | Indicates an infrastructure exploit. | 9 |
| Misc Exploit | 5010 | Indicates a miscellaneous exploit. | 9 |

Table 91: Low-level Categories and Severity Levels for the Exploit Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Web Exploit | 5011 | Indicates a web exploit. | 9 |
| Session Hijack | 5012 | Indicates that a session in your network was interceded. | 9 |
| Worm Active | 5013 | Indicates an active worm. | 10 |
| Password Guess/Retrieve | 5014 | Indicates that a user requested access to their password information from the database. | 9 |
| FTP Exploit | 5015 | Indicates an FTP exploit. | 9 |
| RPC Exploit | 5016 | Indicates an RPC exploit. | 9 |
| SNMP Exploit | 5017 | Indicates an SNMP exploit. | 9 |
| NOOP Exploit | 5018 | Indicates an NOOP exploit. | 9 |
| Samba Exploit | 5019 | Indicates a Samba exploit. | 9 |
| SSH Exploit | 5020 | Indicates an SSH exploit. | 9 |
| Database Exploit | 5021 | Indicates a database exploit. | 9 |
| ICMP Exploit | 5022 | Indicates an ICMP exploit. | 9 |
| UDP Exploit | 5023 | Indicates a UDP exploit. | 9 |

**Table 91: Low-level Categories and Severity Levels for the Exploit Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Browser Exploit | 5024 | Indicates an exploit on your browser. | 9 |
| DHCP Exploit | 5025 | Indicates a DHCP exploit | 9 |
| Remote Access Exploit | 5026 | Indicates a remote access exploit | 9 |
| ActiveX Exploit | 5027 | Indicates an exploit through an ActiveX application. | 9 |
| SQL Injection | 5028 | Indicates that an SQL injection occurred. | 9 |
| Cross-Site Scripting | 5029 | Indicates a cross-site scripting vulnerability. | 9 |
| Format String Vulnerability | 5030 | Indicates a format string vulnerability. | 9 |
| Input Validation Exploit | 5031 | Indicates that an input validation exploit attempt was detected. | 9 |
| Remote Code Execution | 5032 | Indicates that a remote code execution attempt was detected. | 9 |
| Memory Corruption | 5033 | Indicates that a memory corruption exploit was detected. | 9 |

**Table 91: Low-level Categories and Severity Levels for the Exploit Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Command Execution | 5034 | Indicates that a remote command execution attempt was detected. | 9 |
| Code Injection | 5035 | Indicates that a code injection was detected. | 9 |
| Replay Attack | 5036 | Indicates that a replay attack was detected. | 9 |

RELATED DOCUMENTATION

Malware | 531

Suspicious Activity | 534

System | 542

# Malware

The malicious software (malware) category contains events that are related to application exploits and buffer overflow attempts.

The following table describes the low-level event categories and associated severity levels for the malware category.

**Table 92: Low-level Categories and Severity Levels for the Malware Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Malware | 6001 | Indicates an unknown virus. | 4 |

Table 92: Low-level Categories and Severity Levels for the Malware Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Backdoor Detected | 6002 | Indicates that a back door to the system was detected. | 9 |
| Hostile Mail Attachment | 6003 | Indicates a hostile mail attachment. | 6 |
| Malicious Software | 6004 | Indicates a virus. | 6 |
| Hostile Software Download | 6005 | Indicates a hostile software download to your network. | 6 |
| Virus Detected | 6006 | Indicates that a virus was detected. | 8 |
| Misc Malware | 6007 | Indicates miscellaneous malicious software | 4 |
| Trojan Detected | 6008 | Indicates that a trojan was detected. | 7 |
| Spyware Detected | 6009 | Indicates that spyware was detected on your system. | 6 |
| Content Scan | 6010 | Indicates that an attempted scan of your content was detected. | 3 |
| Content Scan Failed | 6011 | Indicates that a scan of your content failed. | 8 |

Table 92: Low-level Categories and Severity Levels for the Malware Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Content Scan Successful | 6012 | Indicates that a scan of your content was successful. | 3 |
| Content Scan in Progress | 6013 | Indicates that a scan of your content is in progress. | 3 |
| Keylogger | 6014 | Indicates that a key logger was detected. | 7 |
| Adware Detected | 6015 | Indicates that Ad-Ware was detected. | 4 |
| Quarantine Successful | 6016 | Indicates that a quarantine action successfully completed. | 3 |
| Quarantine Failed | 6017 | Indicates that a quarantine action failed. | 8 |
| Malware Infection | 6018 | Indicates that a malware infection was detected. | 10 |
| Remove Successful | 6019 | Indicates that the removal was successful. | 3 |
| Remove Failed | 6020 | Indicates that the removal failed. | 8 |

## RELATED DOCUMENTATION

Suspicious Activity | 534

# Suspicious Activity

The suspicious category contains events that are related to viruses, trojans, back door attacks, and other forms of hostile software.

The following table describes the low-level event categories and associated severity levels for the suspicious activity category.

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Suspicious Event | 7001 | Indicates an unknown suspicious event. | 3 |
| Suspicious Pattern Detected | 7002 | Indicates that a suspicious pattern was detected. | 3 |
| Content Modified By Firewall | 7003 | Indicates that content was modified by the firewall. | 3 |
| Invalid Command or Data | 7004 | Indicates an invalid command or data. | 3 |
| Suspicious Packet | 7005 | Indicates a suspicious packet. | 3 |
| Suspicious Activity | 7006 | Indicates suspicious activity. | 3 |
| Suspicious File Name | 7007 | Indicates a suspicious file name. | 3 |

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Suspicious Port Activity | 7008 | Indicates suspicious port activity. | 3 |
| Suspicious Routing | 7009 | Indicates suspicious routing. | 3 |
| Potential Web Vulnerability | 7010 | Indicates potential web vulnerability. | 3 |
| Unknown Evasion Event | 7011 | Indicates an unknown evasion event. | 5 |
| IP Spoof | 7012 | Indicates an IP spoof. | 5 |
| IP Fragmentation | 7013 | Indicates IP fragmentation. | 3 |
| Overlapping IP Fragments | 7014 | Indicates overlapping IP fragments. | 5 |
| IDS Evasion | 7015 | Indicates an IDS evasion. | 5 |
| DNS Protocol Anomaly | 7016 | Indicates a DNS protocol anomaly. | 3 |
| FTP Protocol Anomaly | 7017 | Indicates an FTP protocol anomaly. | 3 |
| Mail Protocol Anomaly | 7018 | Indicates a mail protocol anomaly. | 3 |

Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Routing Protocol Anomaly | 7019 | Indicates a routing protocol anomaly. | 3 |
| Web Protocol Anomaly | 7020 | Indicates a web protocol anomaly. | 3 |
| SQL Protocol Anomaly | 7021 | Indicates an SQL protocol anomaly. | 3 |
| Executable Code Detected | 7022 | Indicates that an executable code was detected. | 5 |
| Misc Suspicious Event | 7023 | Indicates a miscellaneous suspicious event. | 3 |
| Information Leak | 7024 | Indicates an information leak. | 1 |
| Potential Mail Vulnerability | 7025 | Indicates a potential vulnerability in the mail server. | 4 |
| Potential Version Vulnerability | 7026 | Indicates a potential vulnerability in the JSA version. | 4 |
| Potential FTP Vulnerability | 7027 | Indicates a potential FTP vulnerability. | 4 |
| Potential SSH Vulnerability | 7028 | Indicates a potential SSH vulnerability. | 4 |

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Potential DNS Vulnerability | 7029 | Indicates a potential vulnerability in the DNS server. | 4 |
| Potential SMB Vulnerability | 7030 | Indicates a potential SMB (Samba) vulnerability. | 4 |
| Potential Database Vulnerability | 7031 | Indicates a potential vulnerability in the database. | 4 |
| IP Protocol Anomaly | 7032 | Indicates a potential IP protocol anomaly | 3 |
| Suspicious IP Address | 7033 | Indicates that a suspicious IP address was detected. | 2 |
| Invalid IP Protocol Usage | 7034 | Indicates an invalid IP protocol. | 2 |
| Invalid Protocol | 7035 | Indicates an invalid protocol. | 4 |
| Suspicious Window Events | 7036 | Indicates a suspicious event with a screen on your desktop. | 2 |
| Suspicious ICMP Activity | 7037 | Indicates suspicious ICMP activity. | 2 |
| Potential NFS Vulnerability | 7038 | Indicates a potential network file system (NFS) vulnerability. | 4 |

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Potential NNTP Vulnerability | 7039 | Indicates a potential Network News Transfer Protocol (NNTP) vulnerability. | 4 |
| Potential RPC Vulnerability | 7040 | Indicates a potential RPC vulnerability. | 4 |
| Potential Telnet Vulnerability | 7041 | Indicates a potential Telnet vulnerability on your system. | 4 |
| Potential SNMP Vulnerability | 7042 | Indicates a potential SNMP vulnerability. | 4 |
| Illegal TCP Flag Combination | 7043 | Indicates that an invalid TCP flag combination was detected. | 5 |
| Suspicious TCP Flag Combination | 7044 | Indicates that a potentially invalid TCP flag combination was detected. | 4 |
| Illegal ICMP Protocol Usage | 7045 | Indicates that an invalid use of the ICMP protocol was detected. | 5 |
| Suspicious ICMP Protocol Usage | 7046 | Indicates that a potentially invalid use of the ICMP protocol was detected. | 4 |

Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Illegal ICMP Type | 7047 | Indicates that an invalid ICMP type was detected. | 5 |
| Illegal ICMP Code | 7048 | Indicates that an invalid ICMP code was detected. | 5 |
| Suspicious ICMP Type | 7049 | Indicates that a potentially invalid ICMP type was detected. | 4 |
| Suspicious ICMP Code | 7050 | Indicates that a potentially invalid ICMP code was detected. | 4 |
| TCP port 0 | 7051 | Indicates a TCP packet uses a reserved port (0) for source or destination. | 4 |
| UDP port 0 | 7052 | Indicates a UDP packet uses a reserved port (0) for source or destination. | 4 |
| Hostile IP | 7053 | Indicates the use of a known hostile IP address. | 4 |
| Watch list IP | 7054 | Indicates the use of an IP address from a watch list of IP addresses. | 4 |
| Known offender IP | 7055 | Indicates the use of an IP address of a known offender. | 4 |

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| RFC 1918 (private) IP | 7056 | Indicates the use of an IP address from a private IP address range. | 4 |
| Potential VoIP Vulnerability | 7057 | Indicates a potential VoIP vulnerability. | 4 |
| Blacklist Address | 7058 | Indicates that an IP address is on the blocklist. | 8 |
| Watchlist Address | 7059 | Indicates that the IP address is on the list of IP addresses being monitored. | 7 |
| Darknet Address | 7060 | Indicates that the IP address is part of a darknet. | 5 |
| Botnet Address | 7061 | Indicates that the address is part of a botnet. | 7 |
| Suspicious Address | 7062 | Indicates that the IP address must be monitored. | 5 |
| Bad Content | 7063 | Indicates that bad content was detected. | 7 |
| Invalid Cert | 7064 | Indicates that an invalid certificate was detected. | 7 |

**Table 93: Low-level Categories and Severity Levels for the Suspicious Activity Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| User Activity | 7065 | Indicates that user activity was detected. | 7 |
| Suspicious Protocol Usage | 7066 | Indicates that suspicious protocol usage was detected. | 5 |
| Suspicious BGP Activity | 7067 | Indicates that suspicious Border Gateway Protocol (BGP) usage was detected. | 5 |
| Route Poisoning | 7068 | Indicates that route corruption was detected. | 5 |
| ARP Poisoning | 7069 | Indicates that ARP-cache poisoning was detected. | 5 |
| Rogue Device Detected | 7070 | Indicates that a rogue device was detected. | 5 |
| Government Agency Address | 7071 | Indicates that a government agency address was detected. | 3 |

## RELATED DOCUMENTATION

System | 542

Policy | 551

Unknown | 553

# System

The system category contains events that are related to system changes, software installation, or status messages.

The following table describes the low-level event categories and associated severity levels for the system category.

**Table 94: Low-level Categories and Severity Levels for the System Events Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown System Event | 8001 | Indicates an unknown system event. | 1 |
| System Boot | 8002 | Indicates a system restart. | 1 |
| System Configuration | 8003 | Indicates a change in the system configuration. | 1 |
| System Halt | 8004 | Indicates that the system was halted. | 1 |
| System Failure | 8005 | Indicates a system failure. | 6 |
| System Status | 8006 | Indicates any information event. | 1 |
| System Error | 8007 | Indicates a system error. | 3 |
| Misc System Event | 8008 | Indicates a miscellaneous system event. | 1 |
| Service Started | 8009 | Indicates that system services started. | 1 |

Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Service Stopped | 8010 | Indicates that system services stopped. | 1 |
| Service Failure | 8011 | Indicates a system failure. | 6 |
| Successful Registry Modification | 8012 | Indicates that a modification to the registry was successful. | 1 |
| Successful Host-Policy Modification | 8013 | Indicates that a modification to the host policy was successful. | 1 |
| Successful File Modification | 8014 | Indicates that a modification to a file was successful. | 1 |
| Successful Stack Modification | 8015 | Indicates that a modification to the stack was successful. | 1 |
| Successful Application Modification | 8016 | Indicates that a modification to the application was successful. | 1 |
| Successful Configuration Modification | 8017 | Indicates that a modification to the configuration was successful. | 1 |
| Successful Service Modification | 8018 | Indicates that a modification to a service was successful. | 1 |

Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Failed Registry Modification | 8019 | Indicates that a modification to the registry failed. | 1 |
| Failed Host-Policy Modification | 8020 | Indicates that a modification to the host policy failed. | 1 |
| Failed File Modification | 8021 | Indicates that a modification to a file failed. | 1 |
| Failed Stack Modification | 8022 | Indicates that a modification to the stack failed. | 1 |
| Failed Application Modification | 8023 | Indicates that a modification to an application failed. | 1 |
| Failed Configuration Modification | 8024 | Indicates that a modification to the configuration failed. | 1 |
| Failed Service Modification | 8025 | Indicates that a modification to the service failed. | 1 |
| Registry Addition | 8026 | Indicates that a new item was added to the registry. | 1 |
| Host-Policy Created | 8027 | Indicates that a new entry was added to the registry. | 1 |

**Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| File Created | 8028 | Indicates that a new was created in the system. | 1 |
| Application Installed | 8029 | Indicates that a new application was installed on the system. | 1 |
| Service Installed | 8030 | Indicates that a new service was installed on the system. | 1 |
| Registry Deletion | 8031 | Indicates that a registry entry was deleted. | 1 |
| Host-Policy Deleted | 8032 | Indicates that a host policy entry was deleted. | 1 |
| File Deleted | 8033 | Indicates that a file was deleted. | 1 |
| Application Uninstalled | 8034 | Indicates that an application was uninstalled. | 1 |
| Service Uninstalled | 8035 | Indicates that a service was uninstalled. | 1 |
| System Informational | 8036 | Indicates system information. | 3 |
| System Action Allow | 8037 | Indicates that an attempted action on the system was authorized. | 3 |

Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| System Action Deny | 8038 | Indicates that an attempted action on the system was denied. | 4 |
| Cron | 8039 | Indicates a crontab message. | 1 |
| Cron Status | 8040 | Indicates a crontab status message. | 1 |
| Cron Failed | 8041 | Indicates a crontab failure message. | 4 |
| Cron Successful | 8042 | Indicates a crontab success message. | 1 |
| Daemon | 8043 | Indicates a daemon message. | 1 |
| Daemon Status | 8044 | Indicates a daemon status message. | 1 |
| Daemon Failed | 8045 | Indicates a daemon failure message. | 4 |
| Daemon Successful | 8046 | Indicates a daemon success message. | 1 |
| Kernel | 8047 | Indicates a kernel message. | 1 |
| Kernel Status | 8048 | Indicates a kernel status message. | 1 |

Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Kernel Failed | 8049 | Indicates a kernel failure message. | |
| Kernel Successful | 8050 | Indicates a kernel successful message. | 1 |
| Authentication | 8051 | Indicates an authentication message. | 1 |
| Information | 8052 | Indicates an informational message. | 2 |
| Notice | 8053 | Indicates a notice message. | 3 |
| Warning | 8054 | Indicates a warning message. | 5 |
| Error | 8055 | Indicates an error message. | 7 |
| Critical | 8056 | Indicates a critical message. | 9 |
| Debug | 8057 | Indicates a debug message. | 1 |
| Messages | 8058 | Indicates a generic message. | 1 |
| Privilege Access | 8059 | Indicates that privilege access was attempted. | 3 |

**Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Alert | 8060 | Indicates an alert message. | 9 |
| Emergency | 8061 | Indicates an emergency message. | 9 |
| SNMP Status | 8062 | Indicates an SNMP status message. | 1 |
| FTP Status | 8063 | Indicates an FTP status message. | 1 |
| NTP Status | 8064 | Indicates an NTP status message. | 1 |
| Access Point Radio Failure | 8065 | Indicates an access point radio failure. | 3 |
| Encryption Protocol Configuration Mismatch | 8066 | Indicates an encryption protocol configuration mismatch. | 3 |
| Client Device or Authentication Server Misconfigured | 8067 | Indicates that a client device or authentication server was not configured properly. | 5 |
| Hot Standby Enable Failed | 8068 | Indicates a hot standby enable failure. | 5 |
| Hot Standby Disable Failed | 8069 | Indicates a hot standby disable failure. | 5 |

Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Hot Standby Enabled Successfully | 8070 | Indicates that hot standby was enabled successfully. | 1 |
| Hot Standby Association Lost | 8071 | Indicates that a hot standby association was lost. | 5 |
| MainMode Initiation Failure | 8072 | Indicates MainMode initiation failure. | 5 |
| MainMode Initiation Succeeded | 8073 | Indicates that the MainMode initiation was successful. | 1 |
| MainMode Status | 8074 | Indicates a MainMode status message was reported. | 1 |
| QuickMode Initiation Failure | 8075 | Indicates that the QuickMode initiation failed. | 5 |
| Quickmode Initiation Succeeded | 8076 | Indicates that the QuickMode initiation was successful. | 1 |
| Quickmode Status | 8077 | Indicates a QuickMode status message was reported. | 1 |
| Invalid License | 8078 | Indicates an invalid license. | 3 |
| License Expired | 8079 | Indicates an expired license. | 3 |

**Table 94: Low-level Categories and Severity Levels for the System Events Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| New License Applied | 8080 | Indicates a new license applied. | 1 |
| License Error | 8081 | Indicates a license error. | 5 |
| License Status | 8082 | Indicates a license status message. | 1 |
| Configuration Error | 8083 | Indicates that a configuration error was detected. | 5 |
| Service Disruption | 8084 | Indicates that a service disruption was detected. | 5 |
| EPS or FPM allocation exceeded | 8085 | Indicates that the license pool allocations for EPS or FPM were exceeded. | 3 |
| Performance Status | 8086 | Indicates that the performance status was reported. | 1 |
| Performance Degradation | 8087 | Indicates that the performance is being degraded. | 4 |
| Misconfiguration | 8088 | Indicates that an incorrect configuration was detected. | 5 |

RELATED DOCUMENTATION

Policy | 551

# Policy

The policy category contains events that are related to administration of network policy and the monitoring network resources for policy violations.

The following table describes the low-level event categories and associated severity levels for the policy category.

**Table 95: Low-level Categories and Severity Levels for the Policy Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Policy Violation | 9001 | Indicates an unknown policy violation. | 2 |
| Web Policy Violation | 9002 | Indicates a web policy violation. | 2 |
| Remote Access Policy Violation | 9003 | Indicates a remote access policy violation. | 2 |
| IRC/IM Policy Violation | 9004 | Indicates an instant messenger policy violation. | 2 |
| P2P Policy Violation | 9005 | Indicates a Peer-to-Peer (P2P) policy violation. | 2 |
| IP Access Policy Violation | 9006 | Indicates an IP access policy violation. | 2 |
| Application Policy Violation | 9007 | Indicates an application policy violation. | 2 |

Table 95: Low-level Categories and Severity Levels for the Policy Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Database Policy Violation | 9008 | Indicates a database policy violation. | 2 |
| Network Threshold Policy Violation | 9009 | Indicates a network threshold policy violation. | 2 |
| Porn Policy Violation | 9010 | Indicates a porn policy violation. | 2 |
| Games Policy Violation | 9011 | Indicates a games policy violation. | 2 |
| Misc Policy Violation | 9012 | Indicates a miscellaneous policy violation. | 2 |
| Compliance Policy Violation | 9013 | Indicates a compliance policy violation. | 2 |
| Mail Policy Violation | 9014 | Indicates a mail policy violation. | 2 |
| IRC Policy Violation | 9015 | Indicates an IRC policy violation | 2 |
| IM Policy Violation | 9016 | Indicates a policy violation that is related to instant message (IM) activities. | 2 |
| VoIP Policy Violation | 9017 | Indicates a VoIP policy violation | 2 |
| Succeeded | 9018 | Indicates a policy successful message. | 1 |

Table 95: Low-level Categories and Severity Levels for the Policy Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Failed | 9019 | Indicates a policy failure message. | 4 |
| Data Loss Prevention Policy Violation | 9020 | Indicates a data loss prevention policy violation. | 2 |
| Watchlist Object | 9021 | Indicates a watchlist object. | 2 |
| Web Policy Allow | 9022 | Indicates a new web policy allowance. | 1 |

RELATED DOCUMENTATION

Unknown | 553
CRE | 555
Potential Exploit | 556

# Unknown

The Unknown category contains events that are not parsed and therefore cannot be categorized.

The following table describes the low-level event categories and associated severity levels for the Unknown category.

**Table 96: Low-level Categories and Severity Levels for the Unknown Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown | 10001 | Indicates an unknown event. | 3 |
| Unknown Snort Event | 10002 | Indicates an unknown Snort event. | 3 |
| Unknown Dragon Event | 10003 | Indicates an unknown Dragon event. | 3 |
| Unknown Pix Firewall Event | 10004 | Indicates an unknown Cisco Private Internet Exchange (PIX) Firewall event. | 3 |
| Unknown Tipping Point Event | 10005 | Indicates an unknown HP TippingPoint event. | 3 |
| Unknown Windows Auth Server Event | 10006 | Indicates an unknown Windows Auth Server event. | 3 |
| Unknown Nortel Event | 10007 | Indicates an unknown Nortel event. | 3 |
| Stored | 10009 | Indicates an unknown stored event. | 3 |
| Behavioral | 11001 | Indicates an unknown behavioral event. | 3 |
| Threshold | 11002 | Indicates an unknown threshold event. | 3 |

RELATED DOCUMENTATION

CRE | 555

Potential Exploit | 556

Flow | 558

# CRE

The custom rule event (CRE) category contains events that are generated from a custom offense, flow, or eventan event rule.

The following table describes the low-level event categories and associated severity levels for the CRE category.

**Table 97: Low-level Categories and Severity Levels for the CRE Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown CRE Event | 12001 | Indicates an unknown custom rules engine event. | 5 |
| Single Event Rule Match | 12002 | Indicates a single event rule match. | 5 |
| Event Sequence Rule Match | 12003 | Indicates an event sequence rule match. | 5 |
| Cross-Offense Event Sequence Rule Match | 12004 | Indicates a cross-offense event sequence rule match. | 5 |
| Offense Rule Match | 12005 | Indicates an offense rule match. | 5 |

RELATED DOCUMENTATION

Potential Exploit | 556
Flow | 558
User Defined | 561

# Potential Exploit

The potential exploit category contains events that are related to potential application exploits and buffer overflow attempts.

The following table describes the low-level event categories and associated severity levels for the potential exploit category.

**Table 98: Low-level Categories and Severity Levels for the Potential Exploit Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Potential Exploit Attack | 13001 | Indicates that a potential exploitative attack was detected. | 7 |
| Potential Buffer Overflow | 13002 | Indicates that a potential buffer overflow was detected. | 7 |
| Potential DNS Exploit | 13003 | Indicates that a potentially exploitative attack through the DNS server was detected. | 7 |
| Potential Telnet Exploit | 13004 | Indicates that a potentially exploitative attack through Telnet was detected. | 7 |

Table 98: Low-level Categories and Severity Levels for the Potential Exploit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Potential Linux Exploit | 13005 | Indicates that a potentially exploitative attack through Linux was detected. | 7 |
| Potential UNIX Exploit | 13006 | Indicates that a potentially exploitative attack through UNIX was detected. | 7 |
| Potential Windows Exploit | 13007 | Indicates that a potentially exploitative attack through Windows was detected. | 7 |
| Potential Mail Exploit | 13008 | Indicates that a potentially exploitative attack through mail was detected. | 7 |
| Potential Infrastructure Exploit | 13009 | Indicates that a potential exploitative attack on the system infrastructure was detected. | 7 |
| Potential Misc Exploit | 13010 | Indicates that a potentially exploitative attack was detected. | 7 |
| Potential Web Exploit | 13011 | Indicates that a potentially exploitative attack through the web was detected. | 7 |

**Table 98: Low-level Categories and Severity Levels for the Potential Exploit Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Potential Botnet Connection | 13012 | Indicates a potentially exploitative attack that uses botnet was detected. | 6 |
| Potential Worm Activity | 13013 | Indicates a potential attack that uses worm activity was detected. | 6 |

## RELATED DOCUMENTATION

Flow | **558**

User Defined | **561**

SIM Audit | **566**

# Flow

The flow category includes events that are related to flow actions.

The following table describes the low-level event categories and associated severity levels for the flow category.

**Table 99: Low-level Categories and Severity Levels for the Flow Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unidirectional Flow | 14001 | Indicates a unidirectional flow of events. | 5 |
| Low number of Unidirectional Flows | 14002 | Indicates a low number of unidirectional flows of events. | 5 |

Table 99: Low-level Categories and Severity Levels for the Flow Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Medium number of Unidirectional Flows | 14003 | Indicates a medium number of unidirectional flows of events. | 5 |
| High number of Unidirectional Flows | 14004 | Indicates a high number of unidirectional flows of events. | 5 |
| Unidirectional TCP Flow | 14005 | Indicates a unidirectional TCP flow. | 5 |
| Low number of Unidirectional TCP Flows | 14006 | Indicates a low number of unidirectional TCP flows. | 5 |
| Medium number of Unidirectional TCP Flows | 14007 | Indicates a medium number of unidirectional TCP flows. | 5 |
| High number of Unidirectional TCP Flows | 14008 | Indicates a high number of unidirectional TCP flows. | 5 |
| Unidirectional ICMP Flow | 14009 | Indicates a unidirectional ICMP flow. | 5 |
| Low number of Unidirectional ICMP Flows | 14010 | Indicates a low number of unidirectional ICMP flows. | 5 |
| Medium number of Unidirectional ICMP Flows | 14011 | Indicates a medium number of unidirectional ICMP flows. | 5 |
| High number if Unidirectional ICMP Flows | 14012 | Indicates a high number of unidirectional ICMP flows. | 5 |
| Suspicious ICMP Flow | 14013 | Indicates a suspicious ICMP flow. | 5 |
| Suspicious UDP Flow | 14014 | Indicates a suspicious UDP flow. | 5 |

Table 99: Low-level Categories and Severity Levels for the Flow Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Suspicious TCP Flow | 14015 | Indicates a suspicious TCP flow. | 5 |
| Suspicious Flow | 14016 | Indicates a suspicious flow. | 5 |
| Empty Packet Flows | 14017 | Indicates empty packet flows. | 5 |
| Low number of Empty Packet Flows | 14018 | Indicates a low number of empty packet flows. | 5 |
| Medium number of Empty Packet Flows | 14019 | Indicates a medium number of empty packet flows. | 5 |
| High number of Empty Packet Flows | 14020 | Indicates a high number of empty packet flows. | 5 |
| Large Payload Flows | 14021 | Indicates a large payload of flows. | 5 |
| Low number of Large Payload Flows | 14022 | Indicates a low number of large payload flows. | 5 |
| Medium number of Large Payload Flows | 14023 | Indicates a medium number of large payload flows. | 5 |
| High number of Large Payload Flows | 14024 | Indicates a high number of large payload flows. | 5 |
| One Attacker to Many Target Flows | 14025 | Indicates that one attacker is targeting many flows. | 5 |
| Many Attacker to one Target Flow | 14026 | Indicates that many attackers are targeting one flow. | 5 |

Table 99: Low-level Categories and Severity Levels for the Flow Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Unknown Flow | 14027 | Indicates an unknown flow. | 5 |
| Netflow Record | 14028 | Indicates a Netflow record. | 5 |
| Flow Record | 14029 | Indicates a Flow record. | 5 |
| SFlow Record | 14030 | Indicates an SFlow record. | 5 |
| Packeteer Record | 14031 | Indicates a Packeteer record. | 5 |
| Misc Flow | 14032 | Indicates a misc flow. | 5 |
| Large Data Transfer | 14033 | Indicates a large transfer of data. | 5 |
| Large Data Transfer Outbound | 14034 | Indicates a large transfer of outbound data. | 5 |
| VoIP Flows | 14035 | Indicates VoIP Flows. | 5 |

RELATED DOCUMENTATION

User Defined | 561

SIM Audit | 566

VIS Host Discovery | 568

# User Defined

The User Defined category contains events that are related to user-defined objects

The following table describes the low-level event categories and associated severity levels for the User Defined category.

**Table 100: Low-level Categories and Severity Levels for the User Defined Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Custom Sentry Low | 15001 | Indicates a low severity custom anomaly event. | 3 |
| Custom Sentry Medium | 15002 | Indicates a medium severity custom anomaly event. | 5 |
| Custom Sentry High | 15003 | Indicates a high severity custom anomaly event. | 7 |
| Custom Sentry 1 | 15004 | Indicates a custom anomaly event with a severity level of 1. | 1 |
| Custom Sentry 2 | 15005 | Indicates a custom anomaly event with a severity level of 2. | 2 |
| Custom Sentry 3 | 15006 | Indicates a custom anomaly event with a severity level of 3. | 3 |
| Custom Sentry 4 | 15007 | Indicates a custom anomaly event with a severity level of 4. | 4 |
| Custom Sentry 5 | 15008 | Indicates a custom anomaly event with a severity level of 5. | 5 |

**Table 100: Low-level Categories and Severity Levels for the User Defined Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Custom Sentry 6 | 15009 | Indicates a custom anomaly event with a severity level of 6. | 6 |
| Custom Sentry 7 | 15010 | Indicates a custom anomaly event with a severity level of 7. | 7 |
| Custom Sentry 8 | 15011 | Indicates a custom anomaly event with a severity level of 8. | 8 |
| Custom Sentry 9 | 15012 | Indicates a custom anomaly event with a severity level of 9. | 9 |
| Custom Policy Low | 15013 | Indicates a custom policy event with a low severity level. | 3 |
| Custom Policy Medium | 15014 | Indicates a custom policy event with a medium severity level. | 5 |
| Custom Policy High | 15015 | Indicates a custom policy event with a high severity level. | 7 |
| Custom Policy 1 | 15016 | Indicates a custom policy event with a severity level of 1. | 1 |
| Custom Policy 2 | 15017 | Indicates a custom policy event with a severity level of 2. | 2 |

**Table 100: Low-level Categories and Severity Levels for the User Defined Category** ***(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Custom Policy 3 | 15018 | Indicates a custom policy event with a severity level of 3. | 3 |
| Custom Policy 4 | 15019 | Indicates a custom policy event with a severity level of 4. | 4 |
| Custom Policy 5 | 15020 | Indicates a custom policy event with a severity level of 5. | 5 |
| Custom Policy 6 | 15021 | Indicates a custom policy event with a severity level of 6. | 6 |
| Custom Policy 7 | 15022 | Indicates a custom policy event with a severity level of 7. | 7 |
| Custom Policy 8 | 15023 | Indicates a custom policy event with a severity level of 8. | 8 |
| Custom Policy 9 | 15024 | Indicates a custom policy event with a severity level of 9. | 9 |
| Custom User Low | 15025 | Indicates a custom user event with a low severity level. | 3 |
| Custom User Medium | 15026 | Indicates a custom user event with a medium severity level. | 5 |

**Table 100: Low-level Categories and Severity Levels for the User Defined Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Custom User High | 15027 | Indicates a custom user event with a high severity level. | 7 |
| Custom User 1 | 15028 | Indicates a custom user event with a severity level of 1. | 1 |
| Custom User 2 | 15029 | Indicates a custom user event with a severity level of 2. | 2 |
| Custom User 3 | 15030 | Indicates a custom user event with a severity level of 3. | 3 |
| Custom User 4 | 15031 | Indicates a custom user event with a severity level of 4. | 4 |
| Custom User 5 | 15032 | Indicates a custom user event with a severity level of 5. | 5 |
| Custom User 6 | 15033 | Indicates a custom user event with a severity level of 6. | 6 |
| Custom User 7 | 15034 | Indicates a custom user event with a severity level of 7. | 7 |
| Custom User 8 | 15035 | Indicates a custom user event with a severity level of 8. | 8 |

**Table 100: Low-level Categories and Severity Levels for the User Defined Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Custom User 9 | 15036 | Indicates a custom user event with a severity level of 9. | 9 |

RELATED DOCUMENTATION

SIM Audit | 566

VIS Host Discovery | 568

Application | 569

# SIM Audit

The SIM Audit category contains events that are related to user interaction with the JSA Console and administrative features.

The following table describes the low-level event categories and associated severity levels for the SIM Audit category.

**Table 101: Low-level Categories and Severity Levels for the SIM Audit Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| SIM User Authentication | 16001 | Indicates a user login or logout on the Console. | 5 |
| SIM Configuration Change | 16002 | Indicates that a user changed the SIM configuration or deployment. | 3 |

Table 101: Low-level Categories and Severity Levels for the SIM Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| SIM User Action | 16003 | Indicates that a user initiated a process, such as starting a backup or generating a report, in the SIM module. | 3 |
| Session Created | 16004 | Indicates that a user session was created. | 3 |
| Session Destroyed | 16005 | Indicates that a user session was destroyed. | 3 |
| Admin Session Created | 16006 | Indicates that an admin session was created. | |
| Admin Session Destroyed | 16007 | Indicates that an admin session was destroyed. | 3 |
| Session Authentication Invalid | 16008 | Indicates an invalid session authentication. | 5 |
| Session Authentication Expired | 16009 | Indicates that a session authentication expired. | 3 |
| Risk Manager Configuration | 16010 | Indicates that a user changed the JSA Risk Manager configuration. | 3 |

## RELATED DOCUMENTATION

VIS Host Discovery | 568

Application | 569

Audit | 612

# VIS Host Discovery

When the VIS component discovers and stores new hosts, ports, or vulnerabilities that are detected on the network, the VIS component generates events. These events are sent to the Event Collector to be correlated with other security events.

The following table describes the low-level event categories and associated severity levels for the VIS host discovery category.

**Table 102: Low-level Categories and Severity Levels for the VIS Host Discovery Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| New Host Discovered | 17001 | Indicates that the VIS component detected a new host. | 3 |
| New Port Discovered | 17002 | Indicates that the VIS component detected a new open port. | 3 |
| New Vuln Discovered | 17003 | Indicates that the VIS component detected a new vulnerability. | 3 |
| New OS Discovered | 17004 | Indicates that the VIS component detected a new operating system on a host. | 3 |
| Bulk Host Discovered | 17005 | Indicates that the VIS component detected many new hosts in a short period. | 3 |

RELATED DOCUMENTATION

Application | 569

# Application

The application category contains events that are related to application activity, such as email or FTP activity.

The following table describes the low-level event categories and associated severity levels for the application category.

**Table 103: Low-level Categories and Severity Levels for the Application Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Mail Opened | 18001 | Indicates that an email connection was established. | 1 |
| Mail Closed | 18002 | Indicates that an email connection was closed. | 1 |
| Mail Reset | 18003 | Indicates that an email connection was reset. | 3 |
| Mail Terminated | 18004 | Indicates that an email connection was terminated. | 4 |
| Mail Denied | 18005 | Indicates that an email connection was denied. | 4 |
| Mail in Progress | 18006 | Indicates that an email connection is being attempted. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Mail Delayed | 18007 | Indicates that an email connection was delayed. | 4 |
| Mail Queued | 18008 | Indicates that an email connection was queued. | 3 |
| Mail Redirected | 18009 | Indicates that an email connection was redirected. | 1 |
| FTP Opened | 18010 | Indicates that an FTP connection was opened. | 1 |
| FTP Closed | 18011 | Indicates that an FTP connection was closed. | 1 |
| FTP Reset | 18012 | Indicates that an FTP connection was reset. | 3 |
| FTP Terminated | 18013 | Indicates that an FTP connection was terminated. | 4 |
| FTP Denied | 18014 | Indicates that an FTP connection was denied. | 4 |
| FTP In Progress | 18015 | Indicates that an FTP connection is in progress. | 1 |
| FTP Redirected | 18016 | Indicates that an FTP connection was redirected. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| HTTP Opened | 18017 | Indicates that an HTTP connection was established. | 1 |
| HTTP Closed | 18018 | Indicates that an HTTP connection was closed. | 1 |
| HTTP Reset | 18019 | Indicates that an HTTP connection was reset. | 3 |
| HTTP Terminated | 18020 | Indicates that an HTTP connection was terminated. | 4 |
| HTTP Denied | 18021 | Indicates that an HTTP connection was denied. | 4 |
| HTTP In Progress | 18022 | Indicates that an HTTP connection is in progress. | 1 |
| HTTP Delayed | 18023 | Indicates that an HTTP connection was delayed. | 3 |
| HTTP Queued | 18024 | Indicates that an HTTP connection was queued. | 1 |
| HTTP Redirected | 18025 | Indicates that an HTTP connection was redirected. | 1 |
| HTTP Proxy | 18026 | Indicates that an HTTP connection is being proxied. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| HTTPS Opened | 18027 | Indicates that an HTTPS connection was established. | 1 |
| HTTPS Closed | 18028 | Indicates that an HTTPS connection was closed. | 1 |
| HTTPS Reset | 18029 | Indicates that an HTTPS connection was reset. | 3 |
| HTTPS Terminated | 18030 | Indicates that an HTTPS connection was terminated. | 4 |
| HTTPS Denied | 18031 | Indicates that an HTTPS connection was denied. | 4 |
| HTTPS In Progress | 18032 | Indicates that an HTTPS connection is in progress. | 1 |
| HTTPS Delayed | 18033 | Indicates that an HTTPS connection was delayed. | 3 |
| HTTPS Queued | 18034 | Indicates that an HTTPS connection was queued. | 3 |
| HTTPS Redirected | 18035 | Indicates that an HTTPS connection was redirected. | 3 |
| HTTPS Proxy | 18036 | Indicates that an HTTPS connection is proxied. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| SSH Opened | 18037 | Indicates that an SSH connection was established. | 1 |
| SSH Closed | 18038 | Indicates that an SSH connection was closed. | 1 |
| SSH Reset | 18039 | Indicates that an SSH connection was reset. | 3 |
| SSH Terminated | 18040 | Indicates that an SSH connection was terminated. | 4 |
| SSH Denied | 18041 | Indicates that an SSH session was denied. | 4 |
| SSH In Progress | 18042 | Indicates that an SSH session is in progress. | 1 |
| RemoteAccess Opened | 18043 | Indicates that a remote access connection was established. | 1 |
| RemoteAccess Closed | 18044 | Indicates that a remote access connection was closed. | 1 |
| RemoteAccess Reset | 18045 | Indicates that a remote access connection was reset. | 3 |
| RemoteAccess Terminated | 18046 | Indicates that a remote access connection was terminated. | 4 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| RemoteAccess Denied | 18047 | Indicates that a remote access connection was denied. | 4 |
| RemoteAccess In Progress | 18048 | Indicates that a remote access connection is in progress. | 1 |
| RemoteAccess Delayed | 18049 | Indicates that a remote access connection was delayed. | 3 |
| RemoteAccess Redirected | 18050 | Indicates that a remote access connection was redirected. | 3 |
| VPN Opened | 18051 | Indicates that a VPN connection was opened. | 1 |
| VPN Closed | 18052 | Indicates that a VPN connection was closed. | 1 |
| VPN Reset | 18053 | Indicates that a VPN connection was reset. | 3 |
| VPN Terminated | 18054 | Indicates that a VPN connection was terminated. | 4 |
| VPN Denied | 18055 | Indicates that a VPN connection was denied. | 4 |
| VPN In Progress | 18056 | Indicates that a VPN connection is in progress. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| VPN Delayed | 18057 | Indicates that a VPN connection was delayed | 3 |
| VPN Queued | 18058 | Indicates that a VPN connection was queued. | 3 |
| VPN Redirected | 18059 | Indicates that a VPN connection was redirected. | 3 |
| RDP Opened | 18060 | Indicates that an RDP connection was established. | 1 |
| RDP Closed | 18061 | Indicates that an RDP connection was closed. | 1 |
| RDP Reset | 18062 | Indicates that an RDP connection was reset. | 3 |
| RDP Terminated | 18063 | Indicates that an RDP connection was terminated. | 4 |
| RDP Denied | 18064 | Indicates that an RDP connection was denied. | 4 |
| RDP In Progress | 18065 | Indicates that an RDP connection is in progress. | 1 |
| RDP Redirected | 18066 | Indicates that an RDP connection was redirected. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| FileTransfer Opened | 18067 | Indicates that a file transfer connection was established. | 1 |
| FileTransfer Closed | 18068 | Indicates that a file transfer connection was closed. | 1 |
| FileTransfer Reset | 18069 | Indicates that a file transfer connection was reset. | 3 |
| FileTransfer Terminated | 18070 | Indicates that a file transfer connection was terminated. | 4 |
| FileTransfer Denied | 18071 | Indicates that a file transfer connection was denied. | 4 |
| FileTransfer In Progress | 18072 | Indicates that a file transfer connection is in progress. | 1 |
| FileTransfer Delayed | 18073 | Indicates that a file transfer connection was delayed. | 3 |
| FileTransfer Queued | 18074 | Indicates that a file transfer connection was queued. | 3 |
| FileTransfer Redirected | 18075 | Indicates that a file transfer connection was redirected. | 3 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| DNS Opened | 18076 | Indicates that a DNS connection was established. | 1 |
| DNS Closed | 18077 | Indicates that a DNS connection was closed. | 1 |
| DNS Reset | 18078 | Indicates that a DNS connection was reset. | 5 |
| DNS Terminated | 18079 | Indicates that a DNS connection was terminated. | 5 |
| DNS Denied | 18080 | Indicates that a DNS connection was denied. | 5 |
| DNS In Progress | 18081 | Indicates that a DNS connection is in progress. | 1 |
| DNS Delayed | 18082 | Indicates that a DNS connection was delayed. | 5 |
| DNS Redirected | 18083 | Indicates that a DNS connection was redirected. | 4 |
| Chat Opened | 18084 | Indicates that a chat connection was opened. | 1 |
| Chat Closed | 18085 | Indicates that a chat connection was closed. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Chat Reset | 18086 | Indicates that a chat connection was reset. | 3 |
| Chat Terminated | 18087 | Indicates that a chat connection was terminated. | 3 |
| Chat Denied | 18088 | Indicates that a chat connection was denied. | 3 |
| Chat In Progress | 18089 | Indicates that a chat connection is in progress. | 1 |
| Chat Redirected | 18090 | Indicates that a chat connection was redirected. | 1 |
| Database Opened | 18091 | Indicates that a database connection was established. | 1 |
| Database Closed | 18092 | Indicates that a database connection was closed. | 1 |
| Database Reset | 18093 | Indicates that a database connection was reset. | 5 |
| Database Terminated | 18094 | Indicates that a database connection was terminated. | 5 |
| Database Denied | 18095 | Indicates that a database connection was denied. | 5 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Database In Progress | 18096 | Indicates that a database connection is in progress. | 1 |
| Database Redirected | 18097 | Indicates that a database connection was redirected. | 3 |
| SMTP Opened | 18098 | Indicates that an SMTP connection was established. | 1 |
| SMTP Closed | 18099 | Indicates that an SMTP connection was closed. | 1 |
| SMTP Reset | 18100 | Indicates that an SMTP connection was reset. | 3 |
| SMTP Terminated | 18101 | Indicates that an SMTP connection was terminated. | 5 |
| SMTP Denied | 18102 | Indicates that an SMTP connection was denied. | 5 |
| SMTP In Progress | 18103 | Indicates that an SMTP connection is in progress. | 1 |
| SMTP Delayed | 18104 | Indicates that an SMTP connection was delayed. | 3 |
| SMTP Queued | 18105 | Indicates that an SMTP connection was queued. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| SMTP Redirected | 18106 | Indicates that an SMTP connection was redirected. | 3 |
| Auth Opened | 18107 | Indicates that an authorization server connection was established. | 1 |
| Auth Closed | 18108 | Indicates that an authorization server connection was closed. | 1 |
| Auth Reset | 18109 | Indicates that an authorization server connection was reset. | 3 |
| Auth Terminated | 18110 | Indicates that an authorization server connection was terminated. | 4 |
| Auth Denied | 18111 | Indicates that an authorization server connection was denied. | 4 |
| Auth In Progress | 18112 | Indicates that an authorization server connection is in progress. | 1 |
| Auth Delayed | 18113 | Indicates that an authorization server connection was delayed. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Auth Queued | 18114 | Indicates that an authorization server connection was queued. | 3 |
| Auth Redirected | 18115 | Indicates that an authorization server connection was redirected. | 2 |
| P2P Opened | 18116 | Indicates that a Peer-to-Peer (P2P) connection was established. | 1 |
| P2P Closed | 18117 | Indicates that a P2P connection was closed. | 1 |
| P2P Reset | 18118 | Indicates that a P2P connection was reset. | 4 |
| P2P Terminated | 18119 | Indicates that a P2P connection was terminated. | 4 |
| P2P Denied | 18120 | Indicates that a P2P connection was denied. | 3 |
| P2P In Progress | 18121 | Indicates that a P2P connection is in progress. | 1 |
| Web Opened | 18122 | Indicates that a web connection was established. | 1 |
| Web Closed | 18123 | Indicates that a web connection was closed. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Web Reset | 18124 | Indicates that a web connection was reset. | 4 |
| Web Terminated | 18125 | Indicates that a web connection was terminated. | 4 |
| Web Denied | 18126 | Indicates that a web connection was denied. | 4 |
| Web In Progress | 18127 | Indicates that a web connection is in progress. | 1 |
| Web Delayed | 18128 | Indicates that a web connection was delayed. | 3 |
| Web Queued | 18129 | Indicates that a web connection was queued. | 1 |
| Web Redirected | 18130 | Indicates that a web connection was redirected. | 1 |
| Web Proxy | 18131 | Indicates that a web connection was proxied. | 1 |
| VoIP Opened | 18132 | Indicates that a Voice Over IP (VoIP) connection was established. | 1 |
| VoIP Closed | 18133 | Indicates that a VoIP connection was closed. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category** ***(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| VoIP Reset | 18134 | Indicates that a VoIP connection was reset. | 3 |
| VoIP Terminated | 18135 | Indicates that a VoIP connection was terminated. | 3 |
| VoIP Denied | 18136 | Indicates that a VoIP connection was denied. | 3 |
| VoIP In Progress | 18137 | Indicates that a VoIP connection is in progress. | 1 |
| VoIP Delayed | 18138 | Indicates that a VoIP connection was delayed. | 3 |
| VoIP Redirected | 18139 | Indicates that a VoIP connection was redirected. | 3 |
| LDAP Session Started | 18140 | Indicates an LDAP session started. | 1 |
| LDAP Session Ended | 18141 | Indicates an LDAP session ended. | 1 |
| LDAP Session Denied | 18142 | Indicates that an LDAP session was denied. | 3 |
| LDAP Session Status | 18143 | Indicates that an LDAP session status message was reported. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| LDAP Authentication Failed | 18144 | Indicates that an LDAP authentication failed. | 4 |
| LDAP Authentication Succeeded | 18145 | Indicates that an LDAP authentication was successful. | 1 |
| AAA Session Started | 18146 | Indicates that an Authentication, Authorization, and Accounting (AAA) session started. | 1 |
| AAA Session Ended | 18147 | Indicates that an AAA session ended. | 1 |
| AAA Session Denied | 18148 | Indicates that an AAA session was denied. | 3 |
| AAA Session Status | 18149 | Indicates that an AAA session status message was reported. | 1 |
| AAA Authentication Failed | 18150 | Indicates that an AAA authentication failed. | 4 |
| AAA Authentication Succeeded | 18151 | Indicates that an AAA authentication was successful. | 1 |
| IPSEC Authentication Failed | 18152 | Indicates that an Internet Protocol Security (IPSEC) authentication failed. | 4 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| IPSEC Authentication Succeeded | 18153 | Indicates that an IPSEC authentication was successful. | 1 |
| IPSEC Session Started | 18154 | Indicates that an IPSEC session started. | 1 |
| IPSEC Session Ended | 18155 | Indicates that an IPSEC session ended. | 1 |
| IPSEC Error | 18156 | Indicates that an IPSEC error message was reported. | 5 |
| IPSEC Status | 18157 | Indicates that an IPSEC session status message was reported. | 1 |
| IM Session Opened | 18158 | Indicates that an Instant Messenger (IM) session was established. | 1 |
| IM Session Closed | 18159 | Indicates that an IM session was closed. | 1 |
| IM Session Reset | 18160 | Indicates that an IM session was reset. | 3 |
| IM Session Terminated | 18161 | Indicates that an IM session was terminated. | 3 |
| IM Session Denied | 18162 | Indicates that an IM session was denied. | 3 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| IM Session In Progress | 18163 | Indicates that an IM session is in progress. | 1 |
| IM Session Delayed | 18164 | Indicates that an IM session was delayed | 3 |
| IM Session Redirected | 18165 | Indicates that an IM session was redirected. | 3 |
| WHOIS Session Opened | 18166 | Indicates that a WHOIS session was established. | 1 |
| WHOIS Session Closed | 18167 | Indicates that a WHOIS session was closed. | 1 |
| WHOIS Session Reset | 18168 | Indicates that a WHOIS session was reset. | 3 |
| WHOIS Session Terminated | 18169 | Indicates that a WHOIS session was terminated. | 3 |
| WHOIS Session Denied | 18170 | Indicates that a WHOIS session was denied. | 3 |
| WHOIS Session In Progress | 18171 | Indicates that a WHOIS session is in progress. | 1 |
| WHOIS Session Redirected | 18172 | Indicates that a WHOIS session was redirected. | 3 |
| Traceroute Session Opened | 18173 | Indicates that a Traceroute session was established. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Traceroute Session Closed | 18174 | Indicates that a Traceroute session was closed. | 1 |
| Traceroute Session Denied | 18175 | Indicates that a Traceroute session was denied. | 3 |
| Traceroute Session In Progress | 18176 | Indicates that a Traceroute session is in progress. | 1 |
| TN3270 Session Opened | 18177 | This category indicates that a TN3270 session was established. | 1 |
| TN3270 Session Closed | 18178 | Indicates that a TN3270 session was closed. | 1 |
| TN3270 Session Reset | 18179 | Indicates that a TN3270 session was reset. | 3 |
| TN3270 Session Terminated | 18180 | Indicates that a TN3270 session was terminated. | 3 |
| TN3270 Session Denied | 18181 | Indicates that a TN3270 session was denied. | 3 |
| TN3270 Session In Progress | 18182 | Indicates that a TN3270 session is in progress. | 1 |
| TFTP Session Opened | 18183 | Indicates that a TFTP session was established. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| TFTP Session Closed | 18184 | Indicates that a TFTP session was closed. | 1 |
| TFTP Session Reset | 18185 | Indicates that a TFTP session was reset. | 3 |
| TFTP Session Terminated | 18186 | Indicates that a TFTP session was terminated. | 3 |
| TFTP Session Denied | 18187 | Indicates that a TFTP session was denied. | 3 |
| TFTP Session In Progress | 18188 | Indicates that a TFTP session is in progress. | 1 |
| Telnet Session Opened | 18189 | Indicates that a Telnet session was established. | 1 |
| Telnet Session Closed | 18190 | Indicates that a Telnet session was closed. | 1 |
| Telnet Session Reset | 18191 | Indicates that a Telnet session was reset. | 3 |
| Telnet Session Terminated | 18192 | Indicates that a Telnet session was terminated. | 3 |
| Telnet Session Denied | 18193 | Indicates that a Telnet session was denied. | 3 |
| Telnet Session In Progress | 18194 | Indicates that a Telnet session is in progress. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category** ***(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Syslog Session Opened | 18201 | Indicates that a syslog session was established. | 1 |
| Syslog Session Closed | 18202 | Indicates that a syslog session was closed. | 1 |
| Syslog Session Denied | 18203 | Indicates that a syslog session was denied. | 3 |
| Syslog Session In Progress | 18204 | Indicates that a syslog session is in progress. | 1 |
| SSL Session Opened | 18205 | Indicates that a Secure Socket Layer (SSL) session was established. | 1 |
| SSL Session Closed | 18206 | Indicates that an SSL session was closed. | 1 |
| SSL Session Reset | 18207 | Indicates that an SSL session was reset. | 3 |
| SSL Session Terminated | 18208 | Indicates that an SSL session was terminated. | 3 |
| SSL Session Denied | 18209 | Indicates that an SSL session was denied. | 3 |
| SSL Session In Progress | 18210 | Indicates that an SSL session is in progress. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| SNMP Session Opened | 18211 | Indicates that a Simple Network Management Protocol (SNMP) session was established. | 1 |
| SNMP Session Closed | 18212 | Indicates that an SNMP session was closed. | 1 |
| SNMP Session Denied | 18213 | Indicates that an SNMP session was denied. | 3 |
| SNMP Session In Progress | 18214 | Indicates that an SNMP session is in progress. | 1 |
| SMB Session Opened | 18215 | Indicates that a Server Message Block (SMB) session was established. | 1 |
| SMB Session Closed | 18216 | Indicates that an SMB session was closed. | 1 |
| SMB Session Reset | 18217 | Indicates that an SMB session was reset. | 3 |
| SMB Session Terminated | 18218 | Indicates that an SMB session was terminated. | 3 |
| SMB Session Denied | 18219 | Indicates that an SMB session was denied. | 3 |
| SMB Session In Progress | 18220 | Indicates that an SMB session is in progress. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Streaming Media Session Opened | 18221 | Indicates that a Streaming Media session was established. | 1 |
| Streaming Media Session Closed | 18222 | Indicates that a Streaming Media session was closed. | 1 |
| Streaming Media Session Reset | 18223 | Indicates that a Streaming Media session was reset. | 3 |
| Streaming Media Session Terminated | 18224 | Indicates that a Streaming Media session was terminated. | 3 |
| Streaming Media Session Denied | 18225 | Indicates that a Streaming Media session was denied. | 3 |
| Streaming Media Session In Progress | 18226 | Indicates that a Streaming Media session is in progress. | 1 |
| RUSERS Session Opened | 18227 | Indicates that a (Remote Users) RUSERS session was established. | 1 |
| RUSERS Session Closed | 18228 | Indicates that a RUSERS session was closed. | 1 |
| RUSERS Session Denied | 18229 | Indicates that a RUSERS session was denied. | 3 |
| RUSERS Session In Progress | 18230 | Indicates that a RUSERS session is in progress. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Rsh Session Opened | 18231 | Indicates that a remote shell (rsh) session was established. | 1 |
| Rsh Session Closed | 18232 | Indicates that an rsh session was closed. | 1 |
| Rsh Session Reset | 18233 | Indicates that an rsh session was reset. | 3 |
| Rsh Session Terminated | 18234 | Indicates that an rsh session was terminated. | 3 |
| Rsh Session Denied | 18235 | Indicates that an rsh session was denied. | 3 |
| Rsh Session In Progress | 18236 | Indicates that an rsh session is in progress. | 1 |
| RLOGIN Session Opened | 18237 | Indicates that a Remote Login (RLOGIN) session was established. | 1 |
| RLOGIN Session Closed | 18238 | Indicates that an RLOGIN session was closed. | 1 |
| RLOGIN Session Reset | 18239 | Indicates that an RLOGIN session was reset. | 3 |
| RLOGIN Session Terminated | 18240 | Indicates that an RLOGIN session was terminated. | 3 |
| RLOGIN Session Denied | 18241 | Indicates that an RLOGIN session was denied. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| RLOGIN Session In Progress | 18242 | Indicates that an RLOGIN session is in progress. | 1 |
| REXEC Session Opened | 18243 | Indicates that a (Remote Execution) REXEC session was established. | 1 |
| REXEC Session Closed | 18244 | Indicates that an REXEC session was closed. | 1 |
| REXEC Session Reset | 18245 | Indicates that an REXEC session was reset. | 3 |
| REXEC Session Terminated | 18246 | Indicates that an REXEC session was terminated. | 3 |
| REXEC Session Denied | 18247 | Indicates that an REXEC session was denied. | 3 |
| REXEC Session In Progress | 18248 | Indicates that an REXEC session is in progress. | 1 |
| RPC Session Opened | 18249 | Indicates that a Remote Procedure Call (RPC) session was established. | 1 |
| RPC Session Closed | 18250 | Indicates that an RPC session was closed. | 1 |
| RPC Session Reset | 18251 | Indicates that an RPC session was reset. | 3 |
| RPC Session Terminated | 18252 | Indicates that an RPC session was terminated. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| RPC Session Denied | 18253 | Indicates that an RPC session was denied. | 3 |
| RPC Session In Progress | 18254 | Indicates that an RPC session is in progress. | 1 |
| NTP Session Opened | 18255 | Indicates that a Network Time Protocol (NTP) session was established. | 1 |
| NTP Session Closed | 18256 | Indicates that an NTP session was closed. | 1 |
| NTP Session Reset | 18257 | Indicates that an NTP session was reset. | 3 |
| NTP Session Terminated | 18258 | Indicates that an NTP session was terminated. | 3 |
| NTP Session Denied | 18259 | Indicates that an NTP session was denied. | 3 |
| NTP Session In Progress | 18260 | Indicates that an NTP session is in progress. | 1 |
| NNTP Session Opened | 18261 | Indicates that a Network News Transfer Protocol (NNTP) session was established. | 1 |
| NNTP Session Closed | 18262 | Indicates that an NNTP session was closed. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| NNTP Session Reset | 18263 | Indicates that an NNTP session was reset. | 3 |
| NNTP Session Terminated | 18264 | Indicates that an NNTP session was terminated. | 3 |
| NNTP Session Denied | 18265 | Indicates that an NNTP session was denied. | 3 |
| NNTP Session In Progress | 18266 | Indicates that an NNTP session is in progress. | 1 |
| NFS Session Opened | 18267 | Indicates that a Network File System (NFS) session was established. | 1 |
| NFS Session Closed | 18268 | Indicates that an NFS session was closed. | 1 |
| NFS Session Reset | 18269 | Indicates that an NFS session was reset. | 3 |
| NFS Session Terminated | 18270 | Indicates that an NFS session was terminated. | 3 |
| NFS Session Denied | 18271 | Indicates that an NFS session was denied. | 3 |
| NFS Session In Progress | 18272 | Indicates that an NFS session is in progress. | 1 |
| NCP Session Opened | 18273 | Indicates that a Network Control Program (NCP) session was established. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| NCP Session Closed | 18274 | Indicates that an NCP session was closed. | 1 |
| NCP Session Reset | 18275 | Indicates that an NCP session was reset. | 3 |
| NCP Session Terminated | 18276 | Indicates that an NCP session was terminated. | 3 |
| NCP Session Denied | 18277 | Indicates that an NCP session was denied. | 3 |
| NCP Session In Progress | 18278 | Indicates that an NCP session is in progress. | 1 |
| NetBIOS Session Opened | 18279 | Indicates that a NetBIOS session was established. | 1 |
| NetBIOS Session Closed | 18280 | Indicates that a NetBIOS session was closed. | 1 |
| NetBIOS Session Reset | 18281 | Indicates that a NetBIOS session was reset. | 3 |
| NetBIOS Session Terminated | 18282 | Indicates that a NetBIOS session was terminated. | 3 |
| NetBIOS Session Denied | 18283 | Indicates that a NetBIOS session was denied. | 3 |
| NetBIOS Session In Progress | 18284 | Indicates that a NetBIOS session is in progress. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| MODBUS Session Opened | 18285 | Indicates that a MODBUS session was established. | 1 |
| MODBUS Session Closed | 18286 | Indicates that a MODBUS session was closed. | 1 |
| MODBUS Session Reset | 18287 | Indicates that a MODBUS session was reset. | 3 |
| MODBUS Session Terminated | 18288 | Indicates that a MODBUS session was terminated. | 3 |
| MODBUS Session Denied | 18289 | Indicates that a MODBUS session was denied. | 3 |
| MODBUS Session In Progress | 18290 | Indicates that a MODBUS session is in progress. | 1 |
| LPD Session Opened | 18291 | Indicates that a Line Printer Daemon (LPD) session was established. | 1 |
| LPD Session Closed | 18292 | Indicates that an LPD session was closed. | 1 |
| LPD Session Reset | 18293 | Indicates that an LPD session was reset. | 3 |
| LPD Session Terminated | 18294 | Indicates that an LPD session was terminated. | 3 |
| LPD Session Denied | 18295 | Indicates that an LPD session was denied. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| LPD Session In Progress | 18296 | Indicates that an LPD session is in progress. | 1 |
| Lotus Notes Session Opened | 18297 | Indicates that a Lotus Notes session was established. | 1 |
| Lotus Notes Session Closed | 18298 | Indicates that a Lotus Notes session was closed. | 1 |
| Lotus Notes Session Reset | 18299 | Indicates that a Lotus Notes session was reset. | 3 |
| Lotus Notes Session Terminated | 18300 | Indicates that a Lotus Notes session was terminated. | 3 |
| Lotus Notes Session Denied | 18301 | Indicates that a Lotus Notes session was denied. | 3 |
| Lotus Notes Session In Progress | 18302 | Indicates that a Lotus Notes session is in progress. | 1 |
| Kerberos Session Opened | 18303 | Indicates that a Kerberos session was established. | 1 |
| Kerberos Session Closed | 18304 | Indicates that a Kerberos session was closed. | 1 |
| Kerberos Session Reset | 18305 | Indicates that a Kerberos session was reset. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Kerberos Session Terminated | 18306 | Indicates that a Kerberos session was terminated. | 3 |
| Kerberos Session Denied | 18307 | Indicates that a Kerberos session was denied. | 3 |
| Kerberos Session In Progress | 18308 | Indicates that a Kerberos session is in progress. | 1 |
| IRC Session Opened | 18309 | Indicates that an Internet Relay Chat (IRC) session was established. | 1 |
| IRC Session Closed | 18310 | Indicates that an IRC session was closed. | 1 |
| IRC Session Reset | 18311 | Indicates that an IRC session was reset. | 3 |
| IRC Session Terminated | 18312 | Indicates that an IRC session was terminated. | 3 |
| IRC Session Denied | 18313 | Indicates that an IRC session was denied. | 3 |
| IRC Session In Progress | 18314 | Indicates that an IRC session is in progress. | 1 |
| IEC 104 Session Opened | 18315 | Indicates that an IEC 104 session was established. | 1 |
| IEC 104 Session Closed | 18316 | Indicates that an IEC 104 session was closed. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| IEC 104 Session Reset | 18317 | Indicates that an IEC 104 session was reset. | 3 |
| IEC 104 Session Terminated | 18318 | Indicates that an IEC 104 session was terminated. | 3 |
| IEC 104 Session Denied | 18319 | Indicates that an IEC 104 session was denied. | 3 |
| IEC 104 Session In Progress | 18320 | Indicates that an IEC 104 session is in progress. | 1 |
| Ident Session Opened | 18321 | Indicates that a TCP Client Identity Protocol (Ident) session was established. | 1 |
| Ident Session Closed | 18322 | Indicates that an Ident session was closed. | 1 |
| Ident Session Reset | 18323 | Indicates that an Ident session was reset. | 3 |
| Ident Session Terminated | 18324 | Indicates that an Ident session was terminated. | 3 |
| Ident Session Denied | 18325 | Indicates that an Ident session was denied. | 3 |
| Ident Session In Progress | 18326 | Indicates that an Ident session is in progress. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| ICCP Session Opened | 18327 | Indicates that an Inter-Control Center Communications Protocol (ICCP) session was established. | 1 |
| ICCP Session Closed | 18328 | Indicates that an ICCP session was closed. | 1 |
| ICCP Session Reset | 18329 | Indicates that an ICCP session was reset. | 3 |
| ICCP Session Terminated | 18330 | Indicates that an ICCP session was terminated. | 3 |
| ICCP Session Denied | 18331 | Indicates that an ICCP session was denied. | 3 |
| ICCP Session In Progress | 18332 | Indicates that an ICCP session is in progress. | 1 |
| GroupWiseSession Opened | 18333 | Indicates that a GroupWisesession was established. | 1 |
| GroupWiseSession Closed | 18334 | Indicates that a GroupWise session was closed. | 1 |
| GroupWiseSession Reset | 18335 | Indicates that a GroupWisesession was reset. | 3 |

**Table 103: Low-level Categories and Severity Levels for the Application Category** ***(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| GroupWiseSession Terminated | 18336 | Indicates that a GroupWisesession was terminated. | 3 |
| GroupWiseSession Denied | 18337 | Indicates that a GroupWise session was denied. | 3 |
| GroupWiseSession In Progress | 18338 | Indicates that a GroupWise session is in progress. | 1 |
| Gopher Session Opened | 183398 | Indicates that a Gopher session was established. | 1 |
| Gopher Session Closed | 18340 | Indicates that a Gopher session was closed. | 1 |
| Gopher Session Reset | 18341 | Indicates that a Gopher session was reset. | 3 |
| Gopher Session Terminated | 18342 | Indicates that a Gopher session was terminated. | 3 |
| Gopher Session Denied | 18343 | Indicates that a Gopher session was denied. | 3 |
| Gopher Session In Progress | 18344 | Indicates that a Gopher session is in progress. | 1 |
| GIOP Session Opened | 18345 | Indicates that a General Inter-ORB Protocol (GIOP) session was established. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| GIOP Session Closed | 18346 | Indicates that a GIOP session was closed. | 1 |
| GIOP Session Reset | 18347 | Indicates that a GIOP session was reset. | 3 |
| GIOP Session Terminated | 18348 | Indicates that a GIOP session was terminated. | 3 |
| GIOP Session Denied | 18349 | Indicates that a GIOP session was denied. | 3 |
| GIOP Session In Progress | 18350 | Indicates that a GIOP session is in progress. | 1 |
| Finger Session Opened | 18351 | Indicates that a Finger session was established. | 1 |
| Finger Session Closed | 18352 | Indicates that a Finger session was closed. | 1 |
| Finger Session Reset | 18353 | Indicates that a Finger session was reset. | 3 |
| Finger Session Terminated | 18354 | Indicates that a Finger session was terminated. | 3 |
| Finger Session Denied | 18355 | Indicates that a Finger session was denied. | 3 |
| Finger Session In Progress | 18356 | Indicates that a Finger session is in progress. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Echo Session Opened | 18357 | Indicates that an Echo session was established. | 1 |
| Echo Session Closed | 18358 | Indicates that an Echo session was closed. | 1 |
| Echo Session Denied | 18359 | Indicates that an Echo session was denied. | 3 |
| Echo Session In Progress | 18360 | Indicates that an Echo session is in progress. | 1 |
| Remote .NET Session Opened | 18361 | Indicates that a Remote .NET session was established. | 1 |
| Remote .NET Session Closed | 18362 | Indicates that a Remote .NET session was closed. | 1 |
| Remote .NET Session Reset | 18363 | Indicates that a Remote .NET session was reset. | 3 |
| Remote .NET Session Terminated | 18364 | Indicates that a Remote .NET session was terminated. | 3 |
| Remote .NET Session Denied | 18365 | Indicates that a Remote .NET session was denied. | 3 |
| Remote .NET Session In Progress | 18366 | Indicates that a Remote .NET session is in progress. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| DNP3 Session Opened | 18367 | Indicates that a Distributed Network Proctologic (DNP3) session was established. | 1 |
| DNP3 Session Closed | 18368 | Indicates that a DNP3 session was closed. | 1 |
| DNP3 Session Reset | 18369 | Indicates that a DNP3 session was reset. | 3 |
| DNP3 Session Terminated | 18370 | Indicates that a DNP3 session was terminated. | 3 |
| DNP3 Session Denied | 18371 | Indicates that a DNP3 session was denied. | 3 |
| DNP3 Session In Progress | 18372 | Indicates that a DNP3 session is in progress. | 1 |
| Discard Session Opened | 18373 | Indicates that a Discard session was established. | 1 |
| Discard Session Closed | 18374 | Indicates that a Discard session was closed. | 1 |
| Discard Session Reset | 18375 | Indicates that a Discard session was reset. | 3 |
| Discard Session Terminated | 18376 | Indicates that a Discard session was terminated. | 3 |
| Discard Session Denied | 18377 | Indicates that a Discard session was denied. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Discard Session In Progress | 18378 | Indicates that a Discard session is in progress. | 1 |
| DHCP Session Opened | 18379 | Indicates that a Dynamic Host Configuration Protocol (DHCP) session was established. | 1 |
| DHCP Session Closed | 18380 | Indicates that a DHCP session was closed. | 1 |
| DHCP Session Denied | 18381 | Indicates that a DHCP session was denied. | 3 |
| DHCP Session In Progress | 18382 | Indicates that a DHCP session is in progress. | 1 |
| DHCP Success | 18383 | Indicates that a DHCP lease was successfully obtained | 1 |
| DHCP Failure | 18384 | Indicates that a DHCP lease cannot be obtained. | 3 |
| CVS Session Opened | 18385 | Indicates that a Concurrent Versions System (CVS) session was established. | 1 |
| CVS Session Closed | 18386 | Indicates that a CVS session was closed. | 1 |
| CVS Session Reset | 18387 | Indicates that a CVS session was reset. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| CVS Session Terminated | 18388 | Indicates that a CVS session was terminated. | 3 |
| CVS Session Denied | 18389 | Indicates that a CVS session was denied. | 3 |
| CVS Session In Progress | 18390 | Indicates that a CVS session is in progress. | 1 |
| CUPS Session Opened | 18391 | Indicates that a Common UNIX Printing System (CUPS) session was established. | 1 |
| CUPS Session Closed | 18392 | Indicates that a CUPS session was closed. | 1 |
| CUPS Session Reset | 18393 | Indicates that a CUPS session was reset. | 3 |
| CUPS Session Terminated | 18394 | Indicates that a CUPS session was terminated. | 3 |
| CUPS Session Denied | 18395 | Indicates that a CUPS session was denied. | 3 |
| CUPS Session In Progress | 18396 | Indicates that a CUPS session is in progress. | 1 |
| Chargen Session Started | 18397 | Indicates that a Character Generator (Chargen) session was started. | 1 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Chargen Session Closed | 18398 | Indicates that a Chargen session was closed. | 1 |
| Chargen Session Reset | 18399 | Indicates that a Chargen session was reset. | 3 |
| Chargen Session Terminated | 18400 | Indicates that a Chargen session was terminated. | 3 |
| Chargen Session Denied | 18401 | Indicates that a Chargen session was denied. | 3 |
| Chargen Session In Progress | 18402 | Indicates that a Chargen session is in progress. | 1 |
| Misc VPN | 18403 | Indicates that a miscellaneous VPN session was detected | 1 |
| DAP Session Started | 18404 | Indicates that a DAP session was established. | 1 |
| DAP Session Ended | 18405 | Indicates that a DAP session ended. | 1 |
| DAP Session Denied | 18406 | Indicates that a DAP session was denied. | 3 |
| DAP Session Status | 18407 | Indicates that a DAP session status request was made. | 1 |
| DAP Session in Progress | 18408 | Indicates that a DAP session is in progress. | 1 |

**Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| DAP Authentication Failed | 18409 | Indicates that a DAP authentication failed. | 4 |
| DAP Authentication Succeeded | 18410 | Indicates that DAP authentication succeeded. | 1 |
| TOR Session Started | 18411 | Indicates that a TOR session was established. | 1 |
| TOR Session Closed | 18412 | Indicates that a TOR session was closed. | 1 |
| TOR Session Reset | 18413 | Indicates that a TOR session was reset. | 3 |
| TOR Session Terminated | 18414 | Indicates that a TOR session was terminated. | 3 |
| TOR Session Denied | 18415 | Indicates that a TOR session was denied. | 3 |
| TOR Session In Progress | 18416 | Indicates that a TOR session is in progress. | 1 |
| Game Session Started | 18417 | Indicates that a game session was started. | 1 |
| Game Session Closed | 18418 | Indicates that a game session was closed. | 1 |
| Game Session Reset | 18419 | Indicates that a game session was reset. | 3 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Game Session Terminated | 18420 | Indicates that a game session was terminated. | 3 |
| Game Session Denied | 18421 | Indicates that a game session was denied. | 3 |
| Game Session In Progress | 18422 | Indicates that a game session is in progress. | 1 |
| Admin Login Attempt | 18423 | Indicates that an attempt to log in as an administrative user was detected. | 2 |
| User Login Attempt | 18424 | Indicates that an attempt to log in as a non-administrative user was detected. | 2 |
| Client Server | 18425 | Indicates client/server activity. | 1 |
| Content Delivery | 18426 | Indicates content delivery activity. | 1 |
| Data Transfer | 18427 | Indicates a data transfer. | 3 |
| Data Warehousing | 18428 | Indicates data warehousing activity. | 3 |
| Directory Services | 18429 | Indicates directory service activity. | 2 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| File Print | 18430 | Indicates file print activity. | 1 |
| File Transfer | 18431 | Indicates file transfer. | 2 |
| Games | 18432 | Indicates game activity. | 4 |
| Healthcare | 18433 | Indicates healthcare activity. | 1 |
| Inner System | 18434 | Indicates inner system activity. | 1 |
| Internet Protocol | 18435 | Indicates Internet Protocol activity. | 1 |
| Legacy | 18436 | Indicates legacy activity. | 1 |
| Mail | 18437 | Indicates mail activity. | 1 |
| Misc | 18438 | Indicates miscellaneous activity. | 2 |
| Multimedia | 18439 | Indicates multimedia activity. | 2 |
| Network Management | 18440 | Indicates network management activity. | |
| P2P | 18441 | Indicates Peer-to-Peer (P2P) activity. | 4 |

Table 103: Low-level Categories and Severity Levels for the Application Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Remote Access | 18442 | Indicates Remote Access activity. | 3 |
| Routing Protocols | 18443 | Indicates routing protocol activity. | 1 |
| Security Protocols | 18444 | Indicates security protocol activity. | 2 |
| Streaming | 18445 | Indicates streaming activity. | 2 |
| Uncommon Protocol | 18446 | Indicates uncommon protocol activity. | 3 |
| VoIP | 18447 | Indicates VoIP activity. | 1 |
| Web | 18448 | Indicates web activity. | 1 |
| ICMP | 18449 | Indicates ICMP activity | 1 |

RELATED DOCUMENTATION

Audit | 612

Risk | 619

Risk Manager Audit | 621

# Audit

The audit category contains events that are related to audit activity, such as email or FTP activity.

The following table describes the low-level event categories and associated severity levels for the audit category.

**Table 104: Low-level Categories and Severity Levels for the Audit Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| General Audit Event | 19001 | Indicates that a general audit event was started. | 1 |
| Built-in Execution | 19002 | Indicates that a built-in audit task was run. | 1 |
| Bulk Copy | 19003 | Indicates that a bulk copy of data was detected. | 1 |
| Data Dump | 19004 | Indicates that a data dump was detected. | 1 |
| Data Import | 19005 | Indicates that a data import was detected. | 1 |
| Data Selection | 19006 | Indicates that a data selection process was detected. | 1 |
| Data Truncation | 19007 | Indicates that the data truncation process was detected. | 1 |
| Data Update | 19008 | Indicates that the data update process was detected. | 1 |
| Procedure/Trigger Execution | 19009 | Indicates that the database procedure or trigger execution was detected. | 1 |

Table 104: Low-level Categories and Severity Levels for the Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Schema Change | 19010 | Indicates that the schema for a procedure or trigger execution was altered. | 1 |
| Create Activity Attempted | 19011 | Indicates that creating activity was attempted. | 1 |
| Create Activity Succeeded | 19012 | Indicates that creating activity was successful. | 1 |
| Create Activity Failed | 19013 | Indicates that creating activity failed. | 3 |
| Read Activity Attempted | 19014 | Indicates that a reading activity was attempted. | 1 |
| Read Activity Succeeded | 19015 | Indicates that a reading activity was successful. | 1 |
| Read Activity Failed | 19016 | Indicates that reading activity failed. | 3 |
| Update Activity Attempted | 19017 | Indicates that updating activity was attempted. | 1 |
| Update Activity Succeeded | 19018 | Indicates that updating activity was successful. | 1 |
| Update Activity Failed | 19019 | Indicates that updating activity failed. | 3 |
| Delete Activity Attempted | 19020 | Indicates that deleting activity was attempted. | 1 |

Table 104: Low-level Categories and Severity Levels for the Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Delete Activity Succeeded | 19021 | Indicates that deleting activity was successful. | 1 |
| Delete Activity Failed | 19022 | Indicates that deleting activity failed. | 3 |
| Backup Activity Attempted | 19023 | Indicates that backup activity was attempted. | 1 |
| Backup Activity Succeeded | 19024 | Indicates that backup activity was successful. | 1 |
| Backup Activity Failed | 19025 | Indicates that backup activity failed. | 3 |
| Capture Activity Attempted | 19026 | Indicates that capturing activity was attempted. | 1 |
| Capture Activity Succeeded | 19027 | Indicates that capturing activity was successful. | 1 |
| Capture Activity Failed | 19028 | Indicates that capturing activity failed. | 3 |
| Configure Activity Attempted | 19029 | Indicates that configuration activity was attempted. | 1 |
| Configure Activity Succeeded | 19030 | Indicates that configuration activity was successful. | 1 |

Table 104: Low-level Categories and Severity Levels for the Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Configure Activity Failed | 19031 | Indicates that configuration activity failed. | 3 |
| Deploy Activity Attempted | 19032 | Indicates that deployment activity was attempted. | 1 |
| Deploy Activity Succeeded | 19033 | Indicates that deployment activity was successful. | 1 |
| Deploy Activity Failed | 19034 | Indicates that deployment activity failed. | 3 |
| Disable Activity Attempted | 19035 | Indicates that disabling activity was attempted. | 1 |
| Disable Activity Succeeded | 19036 | Indicates that disabling activity was successful. | 1 |
| Disable Activity Failed | 19037 | Indicates that disabling activity failed. | 3 |
| Enable Activity Attempted | 19038 | Indicates that enabling activity was attempted. | 1 |
| Enable Activity Succeeded | 19039 | Indicates that enabling activity was successful. | 1 |
| Enable Activity Failed | 19040 | Indicates that enabling activity failed. | 3 |
| Monitor Activity Attempted | 19041 | Indicates that monitoring activity was attempted. | 1 |

Table 104: Low-level Categories and Severity Levels for the Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Monitor Activity Succeeded | 19042 | Indicates that monitoring activity was successful. | 1 |
| Monitor Activity Failed | 19043 | Indicates that monitoring activity failed. | 3 |
| Restore Activity Attempted | 19044 | Indicates that restoring activity was attempted. | 1 |
| Restore Activity Succeeded | 19045 | Indicates that restoring activity was successful. | 1 |
| Restore Activity Failed | 19046 | Indicates that restoring activity failed. | 3 |
| Start Activity Attempted | 19047 | Indicates that starting activity was attempted. | 1 |
| Start Activity Succeeded | 19048 | Indicates that starting activity was successful. | 1 |
| Start Activity Failed | 19049 | Indicates that starting activity failed. | 3 |
| Stop Activity Attempted | 19050 | Indicates that stopping activity was attempted. | 1 |
| Stop Activity Succeeded | 19051 | Indicates that stopping activity was successful. | 1 |
| Stop Activity Failed | 19052 | Indicates that stopping activity failed. | 3 |

Table 104: Low-level Categories and Severity Levels for the Audit Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Undeploy Activity Attempted | 19053 | Indicates that undeploy activity was attempted. | 1 |
| Undeploy Activity Succeeded | 19054 | Indicates that undeploy activity was successful. | 1 |
| Undeploy Activity Failed | 19055 | Indicates that undeploy activity failed. | 3 |
| Receive Activity Attempted | 19056 | Indicates that receiving activity was attempted. | 1 |
| Receive Activity Succeeded | 19057 | Indicates that receiving activity was successful. | 1 |
| Receive Activity Failed | 19058 | Indicates that receiving activity failed | 3 |
| Send Activity Attempted | 19059 | Indicates that sending activity was attempted. | 1 |
| Send Activity Succeeded | 19060 | Indicates that sending activity was successful. | 1 |
| Send Activity Failed | 19061 | Indicates that sending activity failed. | 3 |

## RELATED DOCUMENTATION

Risk | 619

Risk Manager Audit | 621

Control | 622

# Risk

The risk category contains events that are related to JSA Risk Manager.

The following table describes the low-level event categories and associated severity levels for the risk category.

**Table 105: Low-level Categories and Severity Levels for the Risk Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Policy Exposure | 20001 | Indicates that a policy exposure was detected. | 5 |
| Compliance Violation | 20002 | Indicates that a compliance violation was detected. | 5 |
| Exposed Vulnerability | 20003 | Indicates that the network or device has an exposed vulnerability. | 9 |
| Remote Access Vulnerability | 20004 | Indicates that the network or device has a remote access vulnerability. | 9 |
| Local Access Vulnerability | 20005 | Indicates that the network or device has local access vulnerability. | 7 |
| Open Wireless Access | 20006 | Indicates that the network or device has open wireless access. | 5 |
| Weak Encryption | 20007 | Indicates that the host or device has weak encryption. | 5 |

Table 105: Low-level Categories and Severity Levels for the Risk Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Un-Encrypted Data Transfer | 20008 | Indicates that a host or device is transmitting data that is not encrypted. | 3 |
| Un-Encrypted Data Store | 20009 | Indicates that the data store is not encrypted. | 3 |
| Mis-Configured Rule | 20010 | Indicates that a rule is not configured properly. | 3 |
| Mis-Configured Device | 20011 | Indicates that a device on the network is not configured properly. | 3 |
| Mis-Configured Host | 20012 | Indicates that a network host is not configured properly. | 3 |
| Data Loss Possible | 20013 | Indicates that the possibility of data loss was detected. | 5 |
| Weak Authentication | 20014 | Indicates that a host or device is susceptible to fraud. | 5 |
| No Password | 20015 | Indicates that no password exists. | 7 |
| Fraud | 20016 | Indicates that a host or device is susceptible to fraud. | 7 |
| Possible DoS Target | 20017 | Indicates a host or device is a possible DoS target. | 3 |

Table 105: Low-level Categories and Severity Levels for the Risk Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Possible DoS Weakness | 20018 | Indicates a host or device has a possible DoS weakness. | 3 |
| Loss of Confidentiality | 20019 | Indicates that a loss of confidentially was detected. | 5 |
| Policy Monitor Risk Score Accumulation | 20020 | Indicates that a policy monitor risk score accumulation was detected. | 1 |

RELATED DOCUMENTATION

Risk Manager Audit | 621
Control | 622
Asset Profiler | 626

# Risk Manager Audit

The risk category contains events that are related to JSA Risk Manager audit events.

The following table describes the low-level event categories and associated severity levels for the Risk Manager audit category.

**Table 106: Low-level Categories and Severity Levels for the Risk Manager Audit Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Policy Monitor | 21001 | Indicates that a policy monitor was modified. | 3 |
| Topology | 21002 | Indicates that a topology was modified. | 3 |
| Simulations | 21003 | Indicates that a simulation was modified. | 3 |
| Administration | 21004 | Indicates that administrative changes were made. | 3 |

## RELATED DOCUMENTATION

Control | 622

Asset Profiler | 626

Sense | 635

# Control

The control category contains events that are related to your hardware system.

The following table describes the low-level event categories and associated severity levels for the control category.

Table 107: Low-level Categories and Severity Levels for the Control Category

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Device Read | 22001 | Indicates that a device was read. | 1 |
| Device Communication | 22002 | Indicates communication with a device. | 1 |
| Device Audit | 22003 | Indicates that a device audit occurred. | 1 |
| Device Event | 22004 | Indicates that a device event occurred. | 1 |
| Device Ping | 22005 | Indicates that a ping action to a device occurred. | 1 |
| Device Configuration | 22006 | Indicates that a device was configured. | 1 |
| Device Registration | 22007 | Indicates that a device was registered. | 1 |
| Device Route | 22008 | Indicates that a device route action occurred. | 1 |
| Device Import | 22009 | Indicates that a device import occurred. | 1 |
| Device Information | 22010 | Indicates that a device information action occurred. | 1 |

Table 107: Low-level Categories and Severity Levels for the Control Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Device Warning | 22011 | Indicates that a warning was generated on a device. | 1 |
| Device Error | 22012 | Indicates that an error was generated on a device. | 1 |
| Relay Event | 22013 | Indicates a relay event. | 1 |
| NIC Event | 22014 | Indicates a Network Interface Card (NIC) event. | 1 |
| UIQ Event | 22015 | Indicates an event on a mobile device. | 1 |
| IMU Event | 22016 | Indicates an event on an Integrated Management Unit (IMU). | 1 |
| Billing Event | 22017 | Indicates a billing event. | 1 |
| DBMS Event | 22018 | Indicates an event on the Database Management System (DBMS). | 1 |
| Import Event | 22019 | Indicates that an import occurred. | 1 |
| Location Import | 22020 | Indicates that a location import occurred. | 1 |

Table 107: Low-level Categories and Severity Levels for the Control Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Route Import | 22021 | Indicates that a route import occurred. | 1 |
| Export Event | 22022 | Indicates that an export occurred. | 1 |
| Remote Signaling | 22023 | Indicates remote signaling. | 1 |
| Gateway Status | 22024 | Indicates gateway status. | 1 |
| Job Event | 22025 | Indicates that a job occurred. | 1 |
| Security Event | 22026 | Indicates that a security event occurred. | 1 |
| Device Tamper Detection | 22027 | Indicates that the system detected a tamper action. | 1 |
| Time Event | 22028 | Indicates that a time event occurred. | 1 |
| Suspicious Behavior | 22029 | Indicates that suspicious behavior occurred. | 1 |
| Power Outage | 22030 | Indicates that a power outage occurred. | 1 |
| Power Restoration | 22031 | Indicates that power was restored. | 1 |

**Table 107: Low-level Categories and Severity Levels for the Control Category *(Continued)***

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Heartbeat | 22032 | Indicates that a heartbeat ping occurred. | 1 |
| Remote Connection Event | 22033 | Indicates a remote connection to the system. | 1 |

RELATED DOCUMENTATION

Asset Profiler | 626

Sense | 635

Risk Manager Audit | 621

# Asset Profiler

The asset profiler category contains events that are related to asset profiles.

The following table describes the low-level event categories and associated severity levels for the asset profiler category.

**Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Created | 23001 | Indicates that an asset was created. | 1 |
| Asset Updated | 23002 | Indicates that an asset was updated. | 1 |
| Asset Observed | 23003 | Indicates that an asset was observed. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Moved | 23004 | Indicates that an asset was moved. | 1 |
| Asset Deleted | 23005 | Indicates that an asset was deleted. | 1 |
| Asset Hostname Cleaned | 23006 | Indicates that a host name was cleaned. | 1 |
| Asset Hostname Created | 23007 | Indicates that a host name was created. | 1 |
| Asset Hostname Updated | 23008 | Indicates that a host name was updated. | 1 |
| Asset Hostname Observed | 23009 | Indicates that a host name was observed. | 1 |
| Asset Hostname Moved | 23010 | Indicates that a host name was moved. | 1 |
| Asset Hostname Deleted | 23011 | Indicates that a host name was deleted. | 1 |
| Asset Port Cleaned | 23012 | Indicates that a port was cleaned. | 1 |
| Asset Port Created | 23013 | Indicates that a port was created. | 1 |
| Asset Port Updated | 23014 | Indicates that a port was updated. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Port Observed | 23015 | Indicates that a port was observed. | 1 |
| Asset Port Moved | 23016 | Indicates that a port was moved. | 1 |
| Asset Port Deleted | 23017 | Indicates that a port was deleted. | 1 |
| Asset Vuln Instance Cleaned | 23018 | Indicates that a vulnerability instance was cleaned. | 1 |
| Asset Vuln Instance Created | 23019 | Indicates that a vulnerability instance was created. | 1 |
| Asset Vuln Instance Updated | 23020 | Indicates that a vulnerability instance was updated. | 1 |
| Asset Vuln Instance Observed | 23021 | Indicates that a vulnerability instance was observed. | 1 |
| Asset Vuln Instance Moved | 23022 | Indicates that a vulnerability instance was moved. | 1 |
| Asset Vuln Instance Deleted | 23023 | Indicates that a vulnerability instance was deleted. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset OS Cleaned | 23024 | Indicates that an operating system was cleaned. | 1 |
| Asset OS Created | 23025 | Indicates that an operating system was created. | 1 |
| Asset OS Updated | 23026 | Indicates that an operating system was updated. | 1 |
| Asset OS Observed | 23027 | Indicates that an operating system was observed. | 1 |
| Asset OS Moved | 23028 | Indicates that an operating system was moved. | 1 |
| Asset OS Deleted | 23029 | Indicates that an operating system was deleted. | 1 |
| Asset Property Cleaned | 23030 | Indicates that a property was cleaned. | 1 |
| Asset Property Created | 23031 | Indicates that a property was created. | 1 |
| Asset Property Updated | 23032 | Indicates that a property was updated. | 1 |
| Asset Property Observed | 23033 | Indicates that a property was observed. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Property Moved | 23034 | Indicates that a property was moved. | 1 |
| Asset Property Deleted | 23035 | Indicates that a property was moved. | 1 |
| Asset IP Address Cleaned | 23036 | Indicates that an IP address was cleaned. | 1 |
| Asset IP Address Created | 23037 | Indicates that an IP address was created. | 1 |
| Asset IP Address Updated | 23038 | Indicates that an IP address was updated. | 1 |
| Asset IP Address Observed | 23039 | Indicates that an IP address was observed. | 1 |
| Asset IP Address Moved | 23040 | Indicates that an IP address was moved. | 1 |
| Asset IP Address Deleted | 23041 | Indicates that an IP address was deleted. | 1 |
| Asset Interface Cleaned | 23042 | Indicates that an interface was cleaned. | 1 |
| Asset Interface Created | 23043 | Indicates that an interface was created. | 1 |
| Asset Interface Updated | 23044 | Indicates that an interface was updated. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Interface Observed | 23045 | Indicates that an interface was observed. | 1 |
| Asset Interface Moved | 23046 | Indicates that an interface was moved. | 1 |
| Asset Interface Merged | 23047 | Indicates that an interface was merged. | 1 |
| Asset Interface Deleted | 23048 | Indicates that an interface was deleted. | 1 |
| Asset User Cleaned | 23049 | Indicates that a user was cleaned. | 1 |
| Asset User Observed | 23050 | Indicates that a user was observed. | 1 |
| Asset User Moved | 23051 | Indicates that a user was moved. | 1 |
| Asset User Deleted | 23052 | Indicates that a user was deleted. | 1 |
| Asset Scanned Policy Cleaned | 23053 | Indicates that a scanned policy was cleaned. | 1 |
| Asset Scanned Policy Observed | 23054 | Indicates that a scanned policy was observed. | 1 |
| Asset Scanned Policy Moved | 23055 | Indicates that a scanned policy was moved. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Scanned Policy Deleted | 23056 | Indicates that a scanned policy was deleted. | 1 |
| Asset Windows Application Cleaned | 23057 | Indicates that a Windows application was cleaned. | 1 |
| Asset Windows Application Observed | 23058 | Indicates that a Windows application was observed. | 1 |
| Asset Windows Application Moved | 23059 | Indicates that a Windows application was moved. | 1 |
| Asset Windows Application Deleted | 23060 | Indicates that a Windows application was deleted. | 1 |
| Asset Scanned Service Cleaned | 23061 | Indicates that a scanned service was cleaned. | 1 |
| Asset Scanned Service Observed | 23062 | Indicates that a scanned service was observed. | 1 |
| Asset Scanned Service Moved | 23063 | Indicates that a scanned service was moved. | 1 |
| Asset Scanned Service Deleted | 23064 | Indicates that a scanned service was deleted. | 1 |
| Asset Windows Patch Cleaned | 23065 | Indicates that a Windows patch was cleaned. | 1 |
| Asset Windows Patch Observed | 23066 | Indicates that a Windows patch was observed. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Windows Patch Moved | 23067 | Indicates that a Windows patch was moved. | 1 |
| Asset Windows Patch Deleted | 23068 | Indicates that a Windows patch was deleted. | 1 |
| Asset UNIX Patch Cleaned | 23069 | Indicates that a UNIX patch was cleaned. | 1 |
| Asset UNIX Patch Observed | 23070 | Indicates that a UNIX patch was observed. | 1 |
| Asset UNIX Patch Moved | 23071 | Indicates that a UNIX patch was moved. | 1 |
| Asset UNIX Patch Deleted | 23072 | Indicates that a UNIX patch was deleted. | 1 |
| Asset Patch Scan Cleaned | 23073 | Indicates that a patch scan was cleaned. | 1 |
| Asset Patch Scan Created | 23074 | Indicates that a patch scan was created. | 1 |
| Asset Patch Scan Moved | 23075 | Indicates that a patch scan was moved. | 1 |
| Asset Patch Scan Deleted | 23076 | Indicates that a patch scan was deleted. | 1 |
| Asset Port Scan Cleaned | 23077 | Indicates that a port scan was cleaned. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| Asset Port Scan Created | 23078 | Indicates that a port scan was cleaned. | 1 |
| Asset Port Scan Moved | 23079 | Indicates that a patch scan was moved. | 1 |
| Asset Port Scan Deleted | 23080 | Indicates that a patch scan was deleted. | 1 |
| Asset Client Application Cleaned | 23081 | Indicates that a client application was cleaned. | 1 |
| Asset Client Application Observed | 23082 | Indicates that a client application was observed. | 1 |
| Asset Client Application Moved | 23083 | Indicates that a client application was moved. | 1 |
| Asset Client Application Deleted | 23084 | Indicates that a client application was deleted. | 1 |
| Asset Patch Scan Observed | 23085 | Indicates that a patch scan was observed. | 1 |
| Asset Port Scan Observed | 23086 | Indicates that a port scan was observed. | 1 |
| NetBIOS Group Created | 23087 | Indicates that a NetBIOS group was created. | 1 |
| NetBIOS Group Updated | 23088 | Indicates that a NetBIOS group was updated. | 1 |

Table 108: Low-level Categories and Severity Levels for the Asset Profiler Category *(Continued)*

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| NetBIOS Group Observed | 23089 | Indicates that a NetBIOS group was observed. | 1 |
| NetBIOS Group Deleted | 23090 | Indicates that a NetBIOS group was deleted. | 1 |
| NetBIOS Group Cleaned | 23091 | Indicates that a NetBIOS group was cleaned. | 1 |
| NetBIOS Group Moved | 23092 | Indicates that a NetBIOS group was moved. | 1 |

## RELATED DOCUMENTATION

Sense | **635**

Risk Manager Audit | **621**

Control | **622**

# Sense

The sense category contains events that are related to sense user behavior analytics.

The following table describes the low-level event categories and associated severity levels for the sense category.

**Table 109:**

| Low-level event category | Category ID | Description | Severity level (0 - 10) |
|---|---|---|---|
| User Behavior | 24001 | Indicates the user's behavior. | 5 |
| User Geography | 24002 | Indicates the user's geography. | 5 |
| User Time | 24003 | Indicates the user's time. | 5 |
| User Access | 24004 | Indicates the user's access. | 5 |
| User Privilege | 24005 | Indicates the user's privilege. | 5 |
| User Risk | 24006 | Indicates the user's risk. | 5 |
| Sense Offense | 24007 | Indicates that a sense offense occurred. | 5 |
| Resource Risk | 24008 | Indicates the resources that are at risk. | 5 |

## RELATED DOCUMENTATION

Risk Manager Audit | 621

Control | 622

Asset Profiler | 626

26

CHAPTER

# Common Ports and Servers Used by JSA

# Common Ports and Servers Used by JSA

IN THIS SECTION

JSA requires that certain ports are ready to receive information from JSA components and external infrastructure. To ensure that JSA is using the most recent security information, it also requires access to public servers and RSS feeds.

## SSH Communication on Port 22

All the ports that are used by the JSA console to communicate with managed hosts can be tunneled, by encryption, through port 22 over SSH.

The console connects to the managed hosts using an encrypted SSH session to communicate securely. These SSH sessions are initiated from the console to provide data to the managed host. For example, the JSA console can initiate multiple SSH sessions to the Event Processor appliances for secure communication. This communication can include tunneled ports over SSH, such as HTTPS data for port 443 and Ariel query data for port 32006. JSA Flow Processor that use encryption can initiate SSH sessions to Flow Processor appliances that require data.

## Open Ports that are not Required by JSA

You might find additional open ports in the following situations:

- When you mount or export a network file share, you might see dynamically assigned ports that are required for RPC services, such as `rpc.mountd` and `rpc.rquotad`.

# JSA Port Usage

IN THIS SECTION

Review the list of common ports that JSA services and components use to communicate across the network. You can use the port list to determine which ports must be open in your network. For example, you can determine which ports must be open for the JSA console to communicate with remote event processors.

## WinCollect Remote Polling

WinCollect agents that remotely poll other MicrosoftWindows operating systems might require additional port assignments.

For more information, see the *Juniper Secure Analytics WinCollect User Guide*.

## JSA Listening Ports

The following table shows the JSA ports that are open in a `LISTEN` state. The `LISTEN` ports are valid only when iptables is enabled on your system. Unless otherwise noted, information about the assigned port number applies to all JSA products.

**Table 110: Listening Ports That Are Used by JSA Services and Components**

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 22 | SSH | TCP | Bidirectional from the JSA console to all other components. | Remote management access.<br>Adding a remote system as a managed host.<br>Log source protocols to retrieve files from external devices, for example the log file protocol.<br>Users who use the command-line interface to communicate from desktops to the Console.<br>High-availability (HA). |
| 25 | SMTP | TCP | From all managed hosts to the SMTP gateway. | Emails from JSA to an SMTP gateway.<br>Delivery of error and warning email messages to an administrative email contact. |
| 111 and random generated port | Port mapper | TCP/UDP | Managed hosts (MH) that communicate with the JSA console.<br>Users that connect to the JSA console. | Remote Procedure Calls (RPC) for required services, such as Network File System (NFS). |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 123 | Network Time Protocol (NTP) | UDP | Outbound from the JSA Console to the NTP Server<br>Outbound from the MH to the JSA Console | Time synchronization via Chrony between:<br>• JSA Console and NTP server<br>• JSA Managed Hosts and JSA Console |
| 135 and dynamically allocated ports above 1024 for RPC calls. | DCOM | TCP | Bidirectional traffic between WinCollect agents and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between JSA console components or JSA event collectors that use either Microsoft Security Event Log Protocol or Adaptive Log Exporter agents and Windows operating systems that are remotely polled for events. | This traffic is generated by WinCollect, Microsoft Security Event Log Protocol, or Adaptive Log Exporter.<br>**NOTE**: DCOM typically allocates a random port range for communication. You can configure Microsoft Windows products to use a specific port. For more information, see your Microsoft Windows documentation. |

**Table 110: Listening Ports That Are Used by JSA Services and Components** ***(Continued)***

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 137 | Windows NetBIOS name service | UDP | Bidirectional traffic between WinCollect agents and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between JSA console components or JSA Event Collectors that use either Microsoft Security Event Log Protocol or Adaptive Log Exporter agents and Windows operating systems that are remotely polled for events. | This traffic is generated by WinCollect, Microsoft Security Event Log Protocol, or Adaptive Log Exporter. |
| 138 | Windows NetBIOS datagram service | UDP | Bidirectional traffic between WinCollect agents and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between JSA console components or JSA Event Collectors that use either Microsoft Security Event Log Protocol or Adaptive Log Exporter agents and Windows operating systems that are remotely polled for events. | This traffic is generated by WinCollect, Microsoft Security Event Log Protocol, or Adaptive Log Exporter. |

**Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)***

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 139 | Windows NetBIOS session service | TCP | Bidirectional traffic between WinCollect agents and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between JSA console components or JSA Event Collectors that use either Microsoft Security Event Log Protocol or Adaptive Log Exporter agents and Windows operating systems that are remotely polled for events. | This traffic is generated by WinCollect, Microsoft Security Event Log Protocol, or Adaptive Log Exporter. |
| 162 | NetSNMP | UDP | JSA managed hosts that connect to the JSA console.<br>External log sources to JSA Event Collectors. | UDP port for the NetSNMP daemon that listens for communications (v1, v2c, and v3) from external log sources. The port is open only when the SNMP agent is enabled. |
| 199 | NetSNMP | TCP | JSA managed hosts that connect to the JSA console.<br>External log sources to JSA Event Collectors. | TCP port for the NetSNMP daemon that listens for communications (v1, v2c, and v3) from external log sources. The port is open only when the SNMP agent is enabled. |

**Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)***

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 443 | Apache/HTTPS | TCP | Bidirectional traffic for secure communications from all products to the JSA console.<br>Unidirectional traffic from the App Host to the JSA Console. | Configuration downloads to managed hosts from the JSA console.<br>JSA managed hosts that connect to the JSA console.<br>Users to have log in access to JSA.<br>JSA console that manage and provide configuration updates for WinCollect agents.<br>Apps that require access to the JSA API. |
| 445 | Microsoft Directory Service | TCP | Bidirectional traffic between WinCollect agents and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between JSA console components or JSA Event Collectors that use the Microsoft Security Event Log Protocol and Windows operating systems that are remotely polled for events.<br>Bidirectional traffic between Adaptive Log Exporter agents and Windows operating systems that are remotely polled for events. | This traffic is generated by WinCollect, Microsoft Security Event Log Protocol, or Adaptive Log Exporter. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 514 | Syslog | UDP/TCP | External network appliances that provide TCP syslog events use bidirectional traffic.<br>External network appliances that provide UDP syslog events use uni-directional traffic.<br>Internal syslog traffic from JSA hosts to the JSA console. | External log sources to send event data to JSA components.<br>Syslog traffic includes WinCollect agents, event collectors, and Adaptive Log Exporter agents capable of sending either UDP or TCP events to JSA. |
| 762 | Network File System (NFS) mount daemon (mountd) | TCP/UDP | Connections between the JSA console and NFS server. | The Network File System (NFS) mount daemon, which processes requests to mount a file system at a specified location. |
| 1514 | Syslog-ng | TCP/UDP | Connection between the local Event Collector component and local Event Processor component to the syslog-ng daemon for logging. | Internal logging port for syslog-ng. |
| 2049 | NFS | TCP | Connections between the JSA console and NFS server. | The Network File System (NFS) protocol to share files or data between components. |
| 2055 | NetFlow data | UDP | From the management interface on the flow source (typically a router) to the JSA Flow Processor. | NetFlow datagram from components, such as routers. |
| 2376 | Docker command port | TCP | Internal communications. This port is not available externally. | Used to manage JSA application framework resources. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 3389 | Remote Desktop Protocol (RDP) and Ethernet over USB is enabled | TCP/UDP | | If the MicrosoftWindows operating system is configured to support RDP and Ethernet over USB, a user can initiate a session to the server over the management network. This means the default port for RDP, 3389 must be open. |
| 4333 | Redirect port | TCP | | This port is assigned as a redirect port for Address Resolution Protocol (ARP) requests in JSA offense resolution. |
| 5000 | Used to allow communication to the docker si-registry running on the Console. This allows all managed hosts to pull images from the Console that will be used to create local containers. | TCP | Unidirectional from the JSA managed host to the JSA Console. The port is only opened on the Console. Managed hosts must pull from the Console. | Required for apps running on an App Host. |
| 5432 | Postgres | TCP | Communication for the managed host that is used to access the local database instance. | Required for provisioning managed hosts from the **Admin** tab. |
| 6514 | Syslog | TCP | External network appliances that provide encrypted TCP syslog events use bidirectional traffic. | External log sources to send encrypted event data to JSA components. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 7676, 7677, and four randomly bound ports above 32000. | Messaging connections (IMQ) | TCP | Message queue communications between components on a managed host. | Message queue broker for communications between components on a managed host.<br>**NOTE**: You must permit access to these ports from the JSA console to unencrypted hosts.<br>Ports 7676 and 7677 are static TCP ports, and four extra connections are created on random ports.<br>For more information about finding randomly bound ports, see "Viewing IMQ Port Associations" on page 656. |

**Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)***

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 5791, 7700, 7777, 7778, 7779, 7780, 7781, 7782, 7783, 7787, 7788, 7790, 7791, 7792, 7793, 7794, 7795, 7799, 8989, and 8990. | JMX server ports | TCP | Internal communications. These ports are not available externally. | JMX server (Java Management Beans) monitoring for all internal JSA processes to expose supportability metrics.<br>These ports are used by JSA support. |
| 7789 | HA Distributed Replicated Block Device (DRBD) | TCP/UDP | Bidirectional between the secondary host and primary host in an HA cluster. | Distributed Replicated Block Device (DRBD) used to keep drives synchronized between the primary and secondary hosts in HA configurations. |
| 7800 | Apache Tomcat | TCP | From the Event Collector to the JSA console. | Real-time (streaming) for events. |
| 7801 | Apache Tomcat | TCP | From the Event Collector to the JSA console. | Real-time (streaming) for flows. |
| 7803 | Anomaly Detection Engine | TCP | From the Event Collector to the JSA console. | Anomaly detection engine port. |

**Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)***

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 7804 | QRM Arc builder | TCP | Internal control communications between JSA processes and ARC builder. | This port is used for JSA Risk Manager only. It is not available externally. |
| 7805 | Syslog tunnel communication | TCP | Bidirectional between the JSA Console and managed hosts | Used for encrypted communication between the console and managed hosts. |
| 8000 | Event Collection service (ECS) | TCP | From the Event Collector to the JSA console. | Listening port for specific Event Collection Service (ECS). |
| 8001 | SNMP daemon port | TCP | External SNMP systems that request SNMP trap information from the JSA console. | Listening port for external SNMP data requests. |
| 8005 | Apache Tomcat | TCP | Internal communications. Not available externally. | Open to control tomcat.<br>This port is bound and only accepts connections from the local host. |
| 8009 | Apache Tomcat | TCP | From the HTTP daemon (HTTPd) process to Tomcat. | Tomcat connector, where the request is used and proxied for the web service. |
| 8080 | Apache Tomcat | TCP | From the HTTP daemon (HTTPd) process to Tomcat. | Tomcat connector, where the request is used and proxied for the web service. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 8082 | Secure tunnel for JSA Risk Manager | TCP | Bidirectional traffic between the JSA Console and JSA Risk Manager | Required when encryption is used between JSARisk Manager and the JSA Console. |
| 8413 | WinCollect agents | TCP | Bidirectional traffic between WinCollect agent and JSA console. | This traffic is generated by the WinCollect agent and communication is encrypted. It is required to provide configuration updates to the WinCollect agent and to use WinCollect in connected mode. |
| 8844 | Apache Tomcat | TCP | Unidirectional from the JSA console to the appliance that is running the JSA Vulnerability Manager processor. | Used by Apache Tomcat to read information from the host that is running the JSA Vulnerability Manager processor. |
| 9000 | Conman | TCP | Unidirectional from the JSA Console to a JSA App Host. | Used with an App Host. It allows the Console to deploy apps to an App Host and to manage those apps. |
| 9090 | XForce IP Reputation database and server | TCP | Internal communications. Not available externally. | Communications between JSA processes and the XForce Reputation IP database. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 9381 | Certificate files download | TCP | Unidirectional from JSA managed host or external network to JSA Console | Downloading JSA CA certificate and CRL files, which can be used to validate JSA generated certificates. |
| 9381 | localca-server | TCP | Bidirectional between JSA components. | Used to hold JSA local root and intermediate certificates, as well as associated CRLs. |
| 9393, 9394 | vault-qrd | TCP | Internal communications. Not available externally. | Used to hold secrets and allow secure access to them to services. |
| 9913 plus one dynamically assigned port | Web application container | TCP | Bidirectional Java Remote Method Invocation (RMI) communication between Java Virtual Machines | When the web application is registered, one additional port is dynamically assigned. |
| 9995 | NetFlow data | UDP | From the management interface on the flow source (typically a router) to the JSA flow processor. | NetFlow datagram from components, such as routers. |
| 9999 | JSA Vulnerability Manager processor | TCP | Unidirectional from the scanner to the appliance running the JSA Vulnerability Manager processor | Used for JSA Vulnerability Manager (QVM) command information. The JSA console connects to this port on the host that is running the JSA Vulnerability Manager processor. This port is only used when QVM is enabled. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 10000 | JSA web-based, system administration interface | TCP/UDP | User desktop systems to all JSA hosts. | In JSA 2014.5 and earlier, this port is used for server changes, such as the hosts root password and firewall access.<br>Port 10000 is disabled in 2014.6. |
| 10101, 10102 | Heartbeat command | TCP | Bidirectional traffic between the primary and secondary HA nodes. | Required to ensure that the HA nodes are still active. |
| 12500 | Socat binary | TCP | Outbound from MH to the JSA Console | Port used for tunneling chrony udp requests over tcp when JSA Console or MH is encrypted |
| 14433 | traefik | TCP | Bidirectional between JSA components. | Required for app services discovery. |
| 15432 | | | | Required to be open for internal communication between JSA Risk Manager and JSA. |
| 15433 | Postgres | TCP | Communication for the managed host that is used to access the local database instance. | Used for JSA Vulnerability Manager (QVM) configuration and storage. This port is only used when QVM is enabled. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 20000-23000 | SSH Tunnel | TCP | Bidirectional from the JSA Console to all other encrypted managed hosts. | Local listening point for SSH tunnels used for Java Message Service (JMS) communication with encrypted managed hosts. Used to perform long-running asynchronous tasks, such as updating networking configuration via System and License Management. |
| 23111 | SOAP web server | TCP | | SOAP web server port for the Event Collection Service (ECS). |
| 26000 | traefik | TCP | Bidirectional between JSA components. | Used with an App Host that is encrypted. Required for app services discovery. |
| 26001 | Conman | TCP | Unidirectional from the JSA Console to a JSA App Host. | Used with an App Host that is encrypted. It allows the Consoleto deploy apps to an App Host and to manage those apps. |
| 32000 | Normalized flow forwarding | TCP | Bidirectional between JSA components. | Normalized flow data that is communicated from an off-site source or between JSA Processors. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 32004 | Normalized event forwarding | TCP | Bidirectional between JSA components. | Normalized event data that is communicated from an off-site source or between JSA Event Collectors. |
| 32005 | Data flow | TCP | Bidirectional between JSA components. | Data flow communication port between JSA Event Collectors when on separate managed hosts. |
| 32006 | Ariel queries | TCP | Bidirectional between JSA components. | Communication port between the Ariel proxy server and the Ariel query server. |
| 32007 | Offense data | TCP | Bidirectional between JSA components. | Events and flows contributing to an offense or involved in global correlation. |
| 32009 | Identity data | TCP | Bidirectional between JSA components. | Identity data that is communicated between the passive Vulnerability Information Service (VIS) and the Event Collection Service (ECS). |
| 32010 | Flow listening source port | TCP | Bidirectional between JSA components. | Flow listening port to collect data from JSA Flow Processor. |

Table 110: Listening Ports That Are Used by JSA Services and Components *(Continued)*

| Port | Description | Protocol | Direction | Requirement |
|---|---|---|---|---|
| 32011 | Ariel listening port | TCP | Bidirectional between JSA components. | Ariel listening port for database searches, progress information, and other associated commands. |
| 32000-33999 | Data flow (flows, events, flow context) | TCP | Bidirectional between JSA components. | Data flows, such as events, flows, flow context, and event search queries. |
| 40799 | PCAP data | UDP | From Juniper Networks SRX Series appliances to JSA. | Collecting incoming packet capture (PCAP) data from Juniper Networks SRX Series appliances.<br>**NOTE**: The packet capture on your device can use a different port. For more information about configuring packet capture, see your Juniper Networks SRX Series appliance documentation. |
| ICMP | ICMP | | Bidirectional traffic between the secondary host and primary host in an HA cluster. | Testing the network connection between the secondary host and primary host in an HA cluster by using Internet Control Message Protocol (ICMP). |

RELATED DOCUMENTATION

Viewing IMQ Port Associations | 656

Searching for Ports in Use by JSA | 657

JSA Public Servers | 657

# Viewing IMQ Port Associations

Several ports that are used by JSA allocate extra random port numbers. For example, Message Queues (IMQ) open random ports for communication between components on a managed host. You can view the random port assignments for IMQ by using telnet to connect to the local host and doing a lookup on the port number.

Random port associations are not static port numbers. If a service is restarted, the ports that are generated for the service are reallocated and the service is provided with a new set of port numbers.

1. Using SSH, log in to the JSA console as the root user.
2. To display a list of associated ports for the IMQ messaging connection, type the following command:

   **telnet localhost 7676**

   The results from the telnet command might look similar to this output:

   ```
   [root@domain ~]# telnet localhost 7676 Trying 127.0.0.1... Connected to localhost. Escape character
   is '^]'. 101 imqbroker 4.4 Update 1 portmapper tcp PORTMAPPER 7676 [imqvarhome=/opt/
   openmq/mq/var,imqhome=/opt/openmq/mq,sessionid=<session_id>] cluster_discovery tcp
   CLUSTER_DISCOVERY 44913 jmxrmi rmi JMX 0 [url=service:jmx:rmi://domain.ibm.com/stub/
   <urlpath>] admin tcp ADMIN 43691 jms tcp NORMAL 7677 cluster tcp CLUSTER 36615
   ```

   The telnet output shows 3 of the 4 random high-numbered TCP ports for IMQ. The fourth port, which is not shown, is a JMX Remote Method Invocation (RMI) port that is available over the JMX URL that is shown in the output.

   If the telnet connection is refused, it means that IMQ is not currently running. It is probable that the system is either starting up or shutting down, or that services were shut down manually.

RELATED DOCUMENTATION

Searching for Ports in Use by JSA | 657

JSA Public Servers | 657

JSA Port Usage | 639

# Searching for Ports in Use by JSA

Use the **netstat** command to determine which ports are in use on the JSA Console or managed host. Use the **netstat** command to view all listening and established ports on the system.

1. Using SSH, log in to your JSA console, as the root user.
2. To display all active connections and the TCP and UDP ports on which the computer is listening, type the following command:

   **netstat -nap**
3. To search for specific information from the netstat port list, type the following command:

   **netstat -nap | grep *port***

   **Examples:**

   - To display all ports that match 199, type the following command:

     **netstat -nap | grep 199**

   - To display information on all listening ports, type the following command:

     **netstat -nap | grep LISTEN**

RELATED DOCUMENTATION

JSA Public Servers | 657

JSA Port Usage | 639

Viewing IMQ Port Associations | 656

# JSA Public Servers

To provide you with the most current security information, JSA requires access to a number of public servers.

## Public Servers

This table lists descriptions for the IP addresses or hostnames that JSA accesses. https://www.ibm.com/support/pages/node/6244622

**Table 111: Public Servers That JSA Must Access.**

| IP address or hostname | Description |
|---|---|
| 194.153.113.31 | JSA Vulnerability Manager DMZ scanner |
| 194.153.113.32 | JSA Vulnerability Manager DMZ scanner |
| https://download.juniper.net/ | JSA auto-update servers. |
| | |
| `update.xforce-security.com` | X-Force Threat Feed update server |
| `license.xforce-security.com` | X-Force Threat Feed licensing server |

### RELATED DOCUMENTATION

- JSA Port Usage | 639
- Viewing IMQ Port Associations | 656
- Searching for Ports in Use by JSA | 657

# Docker Containers and Network Interfaces

IN THIS SECTION

A Docker network defines a communication trust zone where communication is unrestricted between containers in that network.

Each network is associated with a bridge interface on the host, and firewall rules are defined to filter traffic between these interfaces. Typically, containers within a zone that share the same Docker network and host bridge interface can communicate with each other. An exception to this general rule is that apps run on the same dockerApps network, but are isolated from each other by the firewall.

## Docker Interfaces

To view a list of Docker interfaces, type the following command:

```
docker network ls
```

Here's an example of the output:

```
[root@q1dk00 ~]# docker network ls
NETWORK ID NAME DRIVER SCOPE
943dd35a4747 appProxy bridge local
9e2ba36111d1 dockerApps bridge local
514471d98b42 dockerInfra bridge local
```

The **dockerApps** interface is used to apply rules for communication between apps.

The **appProxy** interface displays the **nginx_framework_apps_proxy** container.

The **dockerInfra** interface is used to host service launcher and **qoauth**. Apps are isolated from most infrastructure components but they must be able to connect to service launcher and **qoauth**to manage secrets and authorization.

## Information about Docker Interfaces

Type the following command to get information about Docker interfaces:

```
docker inspect <docker_container_ID> | grep NetworkMode
```

Here's an example of the output:

```
"NetworkMode": "appProxy"
```

This example shows how you use the docker inspect *<docker_container_ID>* command and pipe it to less to view more network details:

```
docker inspect d9b3e58649de | less
```

Here's an example of the output:

```
"Networks": {
"dockerApps": {
"IPAMConfig": null,
"Links": null,
"Aliases": [
"d9b3e58649de"
], "NetworkID":
"79bc4716da5139a89cfa5360a3b72824e67701523768822d11b53caeaa5e349e",
"EndpointID":
"9dba9d9a174b037f72333945b72cdf60c3719fdb9a3a10a14a8ee3cc0e92a856",
"Gateway": "172.18.0.1",
"IPAddress": "172.18.0.2",
"IPPrefixLen": 16,
"IPv6Gateway": "2003:db8:1::1",
"GlobalIPv6Address": "2003:db8:1::2",
"GlobalIPv6PrefixLen": 64,
"MacAddress": "02:42:ac:12:00:02"
}
```

The output in this example shows the configuration of the network that is used by the specified container (d9b3e58649de), and shows the Docker network interface name (dockerApps) and the IP address of the network that is assigned to the Docker container.

## RELATED DOCUMENTATION

Viewing IMQ Port Associations | **656**

Searching for Ports in Use by JSA | **657**

JSA Public Servers | **657**

27

CHAPTER

# RESTful API

# RESTful API

IN THIS SECTION

The representational state transfer (REST) application programming interface (API) is useful when you want to integrate JSA with other solutions. You can perform actions on the JSA console by sending HTTPS requests to specific endpoints (URLs) on the JSA console.

Each endpoint contains the URL of the resource that you want to access and the action that you want to complete on that resource. The action is indicated by the HTTP method of the request: GET, POST, PUT, or DELETE. For more information about the parameters and responses for each endpoint, see the *Juniper Secure Analytics API Guide*.

## JSA API Forum and Code Samples

The API forum provides more information about the REST API, including the answers to frequently asked questions and annotated code samples that you can use in a test environment.

RELATED DOCUMENTATION

Accessing the Interactive API Documentation Page | 663

# Accessing the Interactive API Documentation Page

Use the interactive API documentation page to access technical details for the RESTful APIs and experiment with making API requests to your server.

The API documentation user interface provides descriptions and the ability to use the following REST API interfaces:

**Table 112: REST API Interfaces**

| REST API | Description |
|---|---|
| `/api/analytics` | Create, update, and remove custom actions for rules. |
| `/api/ariel` | View event and flow properties, create event and flow searches, and manage searches. |
| `/api/asset_model` | Returns a list of all assets in the model. You can also list all available asset property types and saved searches, and update an asset. |
| `/api/auth` | Log out and invalidate the current session. |
| `/api/config` | View and manage tenants, domains, and JSA extensions. |
| `/api/data_classification` | View all high and low-level categories, QRadar Identifier (QID) records, and event mappings. You can also create or edit QID records and mappings. |
| `/api/gui_app_framework` | Install and manage applications that are created by using the GUI Application Framework Software Development Kit. |
| `/api/help` | Returns a list of API capabilities. |
| `/api/qvm` | Retrieves assets, vulnerabilities, networks, open services, networks, and filters. You can also create or update remediation tickets. |
| `/api/reference_data` | View and manage reference data collections. |
| `/api/scanner` | View, create, or start a remote scan that is related to a scan profile. |

Table 112: REST API Interfaces *(Continued)*

| REST API | Description |
|---|---|
| `/api/services` | Perform tasks such as WHOIS lookups, port scan lookups, DNS lookups, and DIG lookups. You can also retrieve geolocation data for an IP or set of IP addresses. |
| `/api/siem` | View, update, and close offenses. You can also add notes and manage offense closing reasons. |
| `/api/system` | Manage server hosts, network interfaces, and firewall rules. |

1. To access the interactive API documentation interface, enter the following URL in your web browser: **https://*ConsoleIPaddress*/api_doc/**.
2. Select the API version that you want to use from the list.
3. Go to the endpoint that you want to access.
4. Read the endpoint documentation and complete the request parameters.
5. Click **Try it out** to send the API request to your console and receive a properly formatted HTTPS response.

   **NOTE**: When you click **Try it out**, the action is performed on the JSA system. Not all actions can be reversed, for example, you cannot reopen an offense after you close it.

6. Review and gather the information that you need to integrate with JSA.

## RELATED DOCUMENTATION

RESTful API | 663