

Juniper Secure Analytics Users Guide

Published
2022-07-05

RELEASE
7.5.0

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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 Users Guide

7.5.0

Copyright © 2022 Juniper Networks, Inc. All rights reserved.

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 | ix

1

What's New for JSA Users

New Features and Enhancements in JSA 7.5.0 | 2

2

Capabilities in your JSA product

Capabilities in Your JSA Product | 7

Supported Web Browsers | 9

JSA Login | 10

RESTful API | 11

JSA Common Procedures | 12

3

Dashboard Management

Dashboard Management | 20

Default Dashboards | 20

Custom Dashboards | 24

Creating a Custom Dashboard | 34

Using the Dashboard to Investigate Log or Network Activity | 35

Configuring Dashboard Chart Types | 36

Removing Dashboard Items | 37

Detaching a Dashboard Item | 37

Renaming a Dashboard | 38

Deleting a Dashboard | 38

Managing System Notifications | 39

Adding Search-based Dashboard Items to the Add Items List | 39

4

QRadar Analyst Workflow

QRadar Analyst Workflow | 42

What's New in QRadat Analyst Workflow | 43

Known Issues | 44

Installing the Stand-Alone Version | 44

Removing the Stand-Alone Version | 45

Installing the UBI App Version | 46

Offenses | 47

Querying Event and Flow Data to Find Specific Offenses | 52

Searching with the Visual Query Builder | 53

Events | 54

5

Offense Management

Offense Management | 58

Offense Prioritization | 58

Offense Chaining | 59

Offense Indexing | 60

Offense Retention | 62

Offense Investigations | 65

Offense Actions | 79

6

Log Activity Investigation

Log Activity Investigation | 88

Log Activity Tab Overview | 88

Log Activity Monitoring | 93

Viewing Associated Offenses | 111

Modifying Event Mapping | 111

[Tuning False Positive Events From Creating Offenses | 112](#)

[PCAP Data | 113](#)

[Exporting Events | 116](#)

7

Network Activity Monitoring

[Network Activity Monitoring | 119](#)

[Flow Pipeline | 120](#)

[Flow Sources | 121](#)

[Flow Aggregation | 132](#)

[Flow Direction | 134](#)

[Application Identification | 139](#)

[Superflows | 143](#)

[Deduplication | 146](#)

[View Flow Data on the Network Activity Tab | 147](#)

[VLAN Fields | 155](#)

[Configuring a Flow Processor | 158](#)

[Tuning False Positive Events from Creating Offenses | 162](#)

8

Asset Management

[Asset Management | 164](#)

[Sources Of Asset Data | 166](#)

[Incoming Asset Data Workflow | 168](#)

[Updates to Asset Data | 170](#)

[Identification Of Asset Growth Deviations | 175](#)

[Asset Blacklists and Whitelists | 180](#)

[Asset Profiles | 183](#)

9

Chart Management

Chart Management | 205

Time Series Chart Overview | 206

Chart Legends | 208

Configuring Charts | 208

10

Event and Flow Searches

Event and Flow Searches | 212

Offense Searches | 235

Finding IOCs Quickly with Lazy Search | 250

Deleting Search Criteria | 251

Using a Subsearch to Refine Search Results | 252

Managing Search Results | 253

Managing Search Groups | 254

Search Example: Daily Employee Reports | 259

11

Custom Event and Flow Properties

Custom Event and Flow Properties | 262

Creating a Custom Property | 263

Modifying or Deleting a Custom Property | 265

Defining Custom Properties by Using Custom Property Expressions | 265

Use case: Create a report that uses Event Data that is not Normalized | 272

12

Rules

Custom Rules in JSA | 275

Custom Rules | 277

Creating a Custom Rule | 279

Configuring an Event or Flow As False Positive | 284

Anomaly Detection Rules | 284

[Configuring a Rule Response to Add Data to a Reference Data Collection](#) | 294

[Editing Building Blocks](#) | 295

[Rule Performance Visualization](#) | 296

13

Historical Correlation

[Historical Correlation](#) | 300

[Historical Correlation Overview](#) | 301

[Creating a Historical Correlation Profile](#) | 302

[Viewing Information About Historical Correlation Runs](#) | 304

14

Juniper Networks X-Force Integration

[X-Force Integration in JSA](#) | 307

[X-Force Data on the Dashboard](#) | 307

[IP Address and URL Categories](#) | 308

[Confidence Factor and IP Address Reputation](#) | 311

[Searching Data from Juniper X-Force Exchange with Advanced Search Criteria](#) | 313

15

Report Management

[Report Management](#) | 316

[Report Layout](#) | 317

[Chart Types](#) | 317

[Report Tab Toolbar](#) | 320

[Graph Types](#) | 322

[Creating Custom Reports](#) | 325

[Editing a Report](#) | 330

[Viewing Generated Reports](#) | 330

[Deleting Generated Content](#) | 331

[Manually Generating a Report](#) | 332

Duplicating a Report | 333

Sharing a Report | 333

Branding Reports | 334

Report Groups | 335

About This Guide

Use this guide to manage JSA including the dashboard, offenses, reports, assets, log activity, and network activity interfaces.

1

CHAPTER

What's New for JSA Users

[New Features and Enhancements in JSA 7.5.0 | 2](#)

New Features and Enhancements in JSA 7.5.0

IN THIS SECTION

- [Multi-threaded processing for external flow sources | 2](#)
- [Sequence number verification | 3](#)
- [Support for Network Address Translation fields | 3](#)
- [New application determination algorithms | 3](#)
- [Support for more fields from AWS VPC flow logs | 4](#)
- [More improvements | 4](#)

For JSA users, JSA 7.5.0 introduces greater visibility into flow traffic, as well as performance and stability improvements.

Multi-threaded processing for external flow sources

Changed in 7.5.0 Update Pack 1

As part of ongoing improvements to the flow pipeline, the Flow Processor flow processing service now supports multi-threaded processing for external flow sources, such as IPFIX, NetFlow V9, and QRadar Network Insights flow sources.

JSA 7.5.0 introduced multi-threaded processing in the receiving, parsing, and normalization phases when processing external flow sources.

Building on those improvements, JSA 7.5.0 Update Pack 1 introduces multi-threaded processing for the analysis, sending, and garbage collection of flows within the Flow Processor flow processing service.

Now, the entire end-to-end Flow Processor flow processing service uses multi-threaded processing. Performance of the Flow Processor process is improved, which enables JSA to process more flows.

Multi-threaded processing is turned on by default, and the number of threads is automatically determined based on the capabilities of the appliance.

Sequence number verification

New in 7.5.0 Update Pack 1

Now that all stages of the Flow Processor flow processing service uses multi-threaded processing, JSA can use sequence number verification to detect when messages are dropped. Dropped messages might indicate that something is wrong in your network, such as a faulty flow exporter, a lossy network, or packet injection into the network by an attacker.

In JSA 7.5.0 Update Pack 1, missing sequence numbers are reported only once per minute to ensure that packets have time to fill in gaps in the sequence ranges. Missed sequence numbers are reported in the `/var/log/qradar.log` file.

Support for Network Address Translation fields

JSA can now receive network address translation (NAT) information from IPFIX and NetFlow V9 flow records.

The following NAT fields are supported in JSA 7.5.0:

- `postNATSourceIPv4Address` (IANA Element ID 225)
- `postNATDestinationIPv4Address` (IANA Element ID 226)
- `postNAPTSourceTransportPort` (IANA Element ID 227)
- `postNAPTDestinationTransportPort` (IANA Element ID 228)

The new fields are categorized under **Flow Data** on the **Flow Information** window. You can use them in filters, searches, and rules.

For more information about the supported fields for JSA flow sources, see the *Network activity monitoring* chapter in the *Juniper Secure Analytics Users Guide*.

New application determination algorithms

Now you can see more information about the application identification algorithm that is used for QRadar Network Insights flows.

The **QNI Inspectors (9)** algorithm is removed in this release. It is replaced by the following new algorithms:

- **QNI port heuristics (11)**

This algorithm is used when QRadar Network Insights identifies the application based on port heuristics. It represents the least degree of confidence in the application determination.

- **QNI initial data (12)**

This algorithm is used when QRadar Network Insights identifies the application based on the analysis of initial data in the flow session. It represents a medium degree of confidence.

- **QNI parsers (13)**

This algorithm is used when QRadar Network Insights is confident in determining the application based on the data that is available.

You can see the information in the **Application Determination Algorithm** field on the **Flow Information** window.

For more information about identifying flow application, see the *Juniper Secure Analytics Users Guide*.

Support for more fields from AWS VPC flow logs

JSA can now show more information from Amazon Web Services (AWS) Virtual Private Cloud (VPC) Version 3 flow logs.

JSA 7.5.0 introduces support for the following fields:

- VPC ID
- Subnet ID
- Instance ID

When an IPFIX flow record includes these fields, JSA shows the information on the **Flow Details** page under the **Cloud** category.

For more information about viewing AWS flow data, see the *Juniper Secure Analytics Users Guide*.

More improvements

JSA 7.5.0 also includes the following enhancements:

- On the **Component Management** window (**Admin > System and License Management > Edit Host**), the **Alias Autodetection** field is renamed to **DNS lookup for Alias Autodetection**.
- The flow direction algorithm is now applied at the beginning of the flow parsing process.

This change ensures that the destination port is determined before the payload content capture occurs so that the amount of captured payload always matches the setting in the common destination port configuration.

- Only the relevant IPFIX fields are encoded into the payload.

The default encoding method for some IPFIX fields changed, and they are no longer appended to the payload. Now, they are added to the flow as type-value-length (TLV) elements.

- You cannot delete the **Uncategorized** category for tagged flow fields from your system.

2

CHAPTER

Capabilities in your JSA product

[Capabilities in Your JSA Product](#) | 7

[Supported Web Browsers](#) | 9

[JSA Login](#) | 10

[RESTful API](#) | 11

[JSA Common Procedures](#) | 12

Capabilities in Your JSA Product

IN THIS SECTION

- [JSA Product Capabilities | 7](#)

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

JSA product documentation describes capabilities, such as offenses, flows, assets, and historical correlation, that might not be available in all JSA products. Review the following table to compare the capabilities in each 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

Table 1: Comparison Of JSA Capabilities (Continued)

Capability	JSA	Log Manager
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
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	No
Asset profiling	Yes	No ¹
Offenses management	Yes	No
Network flow capture and analysis	Yes	No
Historical correlation	Yes	No

Table 1: Comparison Of JSA Capabilities (Continued)

Capability	JSA	Log Manager
JSA Vulnerability Manager integration	Yes	Yes
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 | 9](#)

[JSA Login | 10](#)

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

Table 2: Supported Web Browsers for JSA Products (Continued)

Web browser	Supported versions
64-bit Google Chrome	Latest

The Microsoft Internet Explorer web browser is no longer supported on 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

[JSA Login | 10](#)

[RESTful API | 11](#)

JSA Login

JSA is a web-based application. JSA uses default login information for the URL, user name, and password.

Use the information in the following table when you log in to your JSA console.

Table 3: Default Login Information for JSA

Login information	Default
URL	<p>https://<IP Address>, where <IP Address> is the IP address of the JSA console.</p> <p>To log in to JSA in an IPv6, wrap the IP address in square brackets:</p> <p>https://[<IP Address>]</p>

Table 3: Default Login Information for JSA (Continued)

Login information	Default
User name	admin
Password	The password that is assigned to JSA during the installation process.
License key	A default license key provides you access to the system for 5 weeks.

RELATED DOCUMENTATION

[RESTful API | 11](#)

[User Interface Tabs](#)

[JSA Common Procedures | 12](#)

RESTful API

IN THIS SECTION

- [JSA API Forum and Code Samples | 12](#)

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

User Interface Tabs

[JSA Common Procedures | 12](#)

[JSA Login | 10](#)

JSA Common Procedures

IN THIS SECTION

- [Viewing Notifications | 12](#)
- [Refreshing and Pausing JSA | 14](#)
- [Investigating IP Addresses | 14](#)
- [System Time | 16](#)
- [Updating User Preferences | 17](#)

Various controls on the JSA user interface are common to most user interface tabs.

Information about these common procedures is described in the following sections.

Viewing Notifications

The **Notifications** menu, which is on the upper right corner of the user interface, provides access to a window in which you can read and manage your system notifications.

For system notifications to show on the **Notifications** window, the administrator must create a rule that is based on each notification message type and select the **Notify** check box in the **Custom Rules Wizard**.

The **Messages** menu indicates how many unread system notifications you have in your system. This indicator increments the number until you close system notifications. For each system notification, the **Messages** window provides a summary and the date stamp for when the system notification was created. You can hover your mouse pointer over a notification to view more detail. Using the functions on the **Messages** window, you can manage the system notifications.

System notifications are also available on the **Dashboard** tab and on an optional pop-up window that can be displayed on the lower left corner of the user interface. Actions that you perform in the **Messages** window are propagated to the **Dashboard** tab and the pop-up window. For example, if you close a system notification from the **Messages** window, the system notification is removed from all system notification displays.

For more information about Dashboard system notifications, see ["Managing System Notifications" on page 39](#).

1. Log in to JSA .
2. On the upper right corner of the user interface, click **Notifications**.
3. On the **Messages** window, view the system notification details.
4. Optional. To refine the list of system notifications, click one of the following options:
 - **Errors**
 - **Warnings**
 - **Information**
5. Optional. To close system notifications, choose of the following options:

Option	Description
Dismiss All	Click to close all system notifications.
Dismiss	Click the Dismiss icon next to the system notification that you want to close.

6. Optional. To view the system notification details, hover your mouse pointer over the system notification.

Refreshing and Pausing JSA

You can manually refresh, pause, and play the data that is displayed on tabs.

- **Dashboard tab**--The **Dashboard** tab automatically refreshes every 60 seconds. The timer, which is on the upper right corner of the interface, indicates the amount of time that remains until the tab is automatically refreshed.

Click the title bar of any dashboard item to automatically pause the refresh time. The timer flashes red to indicate that the current display is paused.

- **Log Activity and Network Activity tabs**--The **Log Activity** and **Network Activity** tabs automatically refresh every 60 seconds if you are viewing the tab in Last Interval (auto refresh) mode.

When you view the **Log Activity** or **Network Activity** tab in Real Time (streaming) or Last Minute (auto refresh) mode, you can use the **Pause** icon to pause the current display.

- **Offenses tab**--The **Offenses** tab must be refreshed manually. The timer, which is on the upper right corner of the interface, indicates the amount of time since the data was last refreshed. The timer flashes red when the timer is paused.

Investigating IP Addresses

You can use several methods to investigate information about IP addresses on the Dashboard, Log Activity, and Network Activity tabs.

1. Log in to JSA.
2. Click the tab that you want to view.
3. Move your mouse pointer over an IP address to view the location of the IP address.
4. Right-click the IP address or asset name and select one of the following options:

Table 4: IP Addresses Information

Option	Description
Navigate >View by Network	Displays the networks that are associated with the selected IP address.

Table 4: IP Addresses Information (Continued)

Option	Description
Navigate >View Source Summary	Displays the offenses that are associated with the selected source IP address.
Navigate >View Destination Summary	Displays the offenses that are associated with the selected destination IP address.
Information >DNS Lookup	Searches for DNS entries that are based on the IP address.
Information >WHOIS Lookup	Searches for the registered owner of a remote IP address. The default WHOIS server is whois.arin.net.
Information >Port Scan	Performs a Network Mapper (NMAP) scan of the selected IP address. This option is only available if NMAP is installed on your system. For more information about installing NMAP, see your vendor documentation.
Information >Asset Profile	<p>Displays asset profile information.</p> <p>This option is displayed if JSA Vulnerability Manager is purchased and licensed. For more information, see the <i>Juniper Secure Analytics Vulnerability Manager User Guide</i>.</p> <p>This menu option is available if JSA acquired profile data either actively through a scan or passively through flow sources.</p> <p>This menu option is available if JSA acquired profile data actively through a scan. For information, see the <i>Juniper Secure Analytics Administration Guide</i>.</p>
Information >Search Events	Searches for events that are associated with this IP address.
Information >Search Flows	Searches for flows that are associated with this IP address.

Table 4: IP Addresses Information (Continued)

Option	Description
Information >Search Connections	Searches for connections that are associated with this IP address. This option is only displayed if you purchased JSA Risk Manager and connected JSA and the JSA Risk Manager appliance. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i> .
Information >Switch Port Lookup	Determines the switch port on a Cisco IOS device for this IP address. This option applies only to switches that are discovered by using the Discover Devices option on the Risks tab. NOTE: This menu option isn't available in Log Manager
Information >View Topology	Displays the Risks tab, which depicts the layer 3 topology of your network. This option is available if you purchased JSA Risk Manager and connected JSA and the JSA Risk Manager appliance.
Run Vulnerability Scan	Select the Run Vulnerability Scan option to scan an JSA Vulnerability Manager scan on this IP address. This option is only displayed when JSA Vulnerability Manager has been purchased and licensed. For more information, see the <i>Juniper Secure Analytics Vulnerability Manager User Guide</i> .

System Time

The right corner of the JSA user interface displays system time, which is the time on the console.

The console time synchronizes JSA systems within the JSA deployment. The console time is used to determine what time events were received from other devices for correct time synchronization correlation.

In a distributed deployment, the console might be in a different time zone from your desktop computer.

When you apply time-based filters and searches on the **Log Activity** and **Network Activity** tabs, you must use the console system time to specify a time range.

When you apply time-based filters and searches on the **Log Activity** tab, you must use the console system time to specify a time range.

Updating User Preferences

You can set your user preference, such as locale, in the main JSA user interface.

1. To access your user information, click **Preferences**.
2. Update your preferences.

Option	Description
Username	Displays your user name. You cannot edit this field.
Password	<p>JSA user passwords are stored as a salted SHA-256 string.</p> <p>The password must meet the following criteria:</p> <ul style="list-style-type: none"> • Minimum of 6 characters • Maximum of 255 characters • Contain at least 1 special character • Contain 1 uppercase character
Password (Confirm)	Password confirmation
Email Address	<p>The email address must meet the following requirements:</p> <ul style="list-style-type: none"> • Minimum of 10 characters • Maximum of 255 characters
Locale	<p>JSA is available in the following languages: English, Simplified Chinese, Traditional Chinese, Japanese, Korean, French, German, Italian, Spanish, Russian, and Portuguese (Brazil).</p> <p>If you choose a different language, the user interface displays in English. Other associated cultural conventions, such as, character type, collation, format of date and time, currency unit are used.</p>

(Continued)

Option	Description
Enable Popup Notifications	Select this check box if you want to enable pop-up system notifications to be displayed on your user interface.

RELATED DOCUMENTATION

[JSA Login | 10](#)

[RESTful API | 11](#)

User Interface Tabs

3

CHAPTER

Dashboard Management

- [Dashboard Management | 20](#)
 - [Default Dashboards | 20](#)
 - [Custom Dashboards | 24](#)
 - [Creating a Custom Dashboard | 34](#)
 - [Using the Dashboard to Investigate Log or Network Activity | 35](#)
 - [Configuring Dashboard Chart Types | 36](#)
 - [Removing Dashboard Items | 37](#)
 - [Detaching a Dashboard Item | 37](#)
 - [Renaming a Dashboard | 38](#)
 - [Deleting a Dashboard | 38](#)
 - [Managing System Notifications | 39](#)
 - [Adding Search-based Dashboard Items to the Add Items List | 39](#)
-

Dashboard Management

Use the **Dashboard** tab, which is the default view when you log into JSA, to focus on specific areas of your network security. The workspace supports multiple dashboards on which you can display your views of network security, activity, or data that is collected.

TIP: Use the QRadar Pulse app for an enhanced dashboard experience. The Pulse app is included with JSA 7.4.0 and later. For more information about the Pulse app, see *Pulse App Guide*

You can customize your dashboard. The content that is displayed on the **Dashboard** tab is user-specific. Changes that are made within a session affect only your system. For example, you can make these customizations:

- Add and remove dashboard items from your dashboards.
- Move and position items to meet your requirements.

When you position items, each item is automatically resized in proportion to the dashboard.

- Add custom dashboard items that are based on any data.

For example, you can add a dashboard item that provides a time series graph or a bar chart that represents top 10 network activity.

To create custom items, you can create saved searches on the **Log Activity** tab and choose how you want the results that are represented in your dashboard. Each dashboard chart displays real-time upto-the-minute data. Time series graphs on the dashboard refresh every 5 minutes.

Default Dashboards

Use the default dashboard to customize your items into functional views. These functional views focus on specific areas of your network.

The **Dashboard** tab provides five default dashboards that are focused on security, network activity, application activity, system monitoring, and compliance.

Each dashboard displays a default that is set of dashboard items. The dashboard items act as starting point to navigate to more detailed data. The following table defines the default dashboards.

Table 5: Default Dashboards

Default dashboard	Items
Application Overview	<p>The Application Overview dashboard includes the following default items:</p> <ul style="list-style-type: none"> • Inbound Traffic by Country (Total Bytes) • Outbound Traffic by Country (Total Bytes) • Top Applications (Total Bytes) • Top Applications Inbound from Internet (Total Bytes) • Top Applications Outbound to the Internet (Total Bytes) • Top Services Denied through Firewalls (Event Count) • DSCP - Precedence (Total Bytes)
Compliance Overview	<p>The Compliance Overview dashboard includes the following default items:</p> <ul style="list-style-type: none"> • Top Authentications by User (Time Series) • Top Authentication Failures by User (Event Count) • Login Failures by User (real-time) • Compliance: Username Involved in Compliance Rules (time series) • Compliance: Source IPs Involved in Compliance Rules (time series) • Most Recent Reports •

Table 5: Default Dashboards (*Continued*)

Default dashboard	Items
Network Overview	<p>The Network Overview dashboard includes the following default items:</p> <ul style="list-style-type: none"> • Top Talkers (real time) • ICMP Type/Code (Total Packets) • Top Networks by Traffic Volume (Total Bytes) • Firewall Deny by DST Port (Event Count) • Firewall Deny by DST IP (Event Count) • Firewall Deny by SRC IP (Event Count) • Top Applications (Total Bytes) • Link Utilization (real-time) • DSCP - Precedence (Total Bytes)
System Monitoring	<p>The System Monitoring dashboard includes the following default items:</p> <ul style="list-style-type: none"> • Top Log Sources (Event Count) • Link Utilization (real-time) • System Notifications • Event Processor Distribution (Event Count) • Event Rate (Events per Second Coalesced - Average 1 Min) • Flow Rate (Flows per Second - Peak 1 Min)

Table 5: Default Dashboards (*Continued*)

Default dashboard	Items
Threat and Security Monitoring	<p>The Threat and Security Monitoring dashboard includes the following default items:</p> <ul style="list-style-type: none"> • Default-IDS/IPS-All: Top Alarm Signatures (real-time) • Top Systems Attacked (Event Count) • Top Systems Sourcing Attacks (Event Count) • My Offenses • Most Severe Offenses • Most Recent Offenses • Top Services Denied through Firewalls (Event Count) • Internet Threat Information Center • Flow Bias (Total Bytes) • Top Category Types • Top Sources • Top Local Destinations

RELATED DOCUMENTATION

[Custom Dashboards | 24](#)

[Creating a Custom Dashboard | 34](#)

Custom Dashboards

IN THIS SECTION

- [Flow Search | 25](#)
- [Offenses | 25](#)
- [Log Activity | 26](#)
- [System Summary | 28](#)
- [Risk Monitoring Dashboard | 28](#)
- [Monitoring Policy Compliance | 29](#)
- [Monitoring Risk Change | 31](#)
- [Vulnerability Management Items | 32](#)
- [System Notification | 32](#)

You can customize your dashboards. The content that is displayed on the **Dashboard** tab is user-specific. Changes that are made within a JSA session affect only your system.

To customize your **Dashboard** tab, you can perform the following tasks:

- Create custom dashboards that are relevant to your responsibilities. 255 dashboards per user is the maximum; however, performance issues might occur if you create more than 10 dashboards.
- Add and remove dashboard items from default or custom dashboards.
- Move and position items to meet your requirements. When you position items, each item automatically resizes in proportion to the dashboard.
- Add custom dashboard items that are based on any data.

For example, you can add a dashboard item that provides a time series graph or a bar chart that represents top 10 network activity.

To create custom items, you can create saved searches on the **Network Activity** or **Log Activity** tabs and choose how you want the results that are represented in your dashboard. Each dashboard chart displays real-time up-to-the-minute data. Time series graphs on the dashboard refresh every 5 minutes.

Flow Search

You can display a custom dashboard item that is based on saved search criteria from the **Network Activity** tab.

Flow search items are listed in the **Add Item >Network Activity >Flow Searches** menu. The name of the flow search item matches the name of the saved search criteria the item is based on.

Default saved search criteria is available and is preconfigured to display flow search items on your **Dashboard** tab menu. You can add more flow search dashboard items to your **Dashboard** tab menu. For more information, see ["Adding Search-based Dashboard Items to the Add Items List" on page 39](#).

On a flow search dashboard item, search results display real-time last-minute data on a chart. The supported chart types are time series, table, pie, and bar. The default chart type is bar. These charts are configurable. For more information about configuring charts, see ["Configuring Charts" on page 208](#).

Time series charts are interactive. Using the time series charts, you can magnify and scan through a timeline to investigate network activity.

Offenses

You can add several offense-related items to your dashboard.

NOTE: Hidden or closed offenses are not included in the values that are displayed in the **Dashboard** tab. For more information about hidden or closed events, see ["Offense Management" on page 58](#).

The following table describes the Offense items:

Table 6: Offense Items

Dashboard items	Description
Most Recent Offenses	The five most recent offenses are identified with a magnitude bar to inform you of the importance of the offense. Point to the offense name to view detailed information for the IP address.

Table 6: Offense Items (Continued)

Dashboard items	Description
Most Severe Offenses	The five most severe offenses are identified with a magnitude bar to inform you of the importance of the offense. Point to the offense name to view detailed information for the IP address.
My Offenses	The My Offenses item displays 5 of the most recent offenses that are assigned to you. The offenses are identified with a magnitude bar to inform you of the importance of the offense. Point to the IP address to view detailed information for the IP address.
Top Sources	The Top Sources item displays the top offense sources. Each source is identified with a magnitude bar to inform you of the importance of the source. Point to the IP address to view detailed information for the IP address.
Top Local Destinations	The Top Local Destinations item displays the top local destinations. Each destination is identified with a magnitude bar to inform you of the importance of the destination. Point to the IP address to view detailed information for the IP
Categories	The Top Categories Types item displays the top 5 categories that are associated with the highest number of offenses.

Log Activity

The **Log Activity** dashboard items will allow you to monitor and investigate events in real time.

NOTE: Hidden or closed events are not included in the values that are displayed in the **Dashboard** tab.

Table 7: Log Activity Items

Dashboard item	Description
Event Searches	<p>You can display a custom dashboard item that is based on saved search criteria from the Log Activity tab. Event search items are listed in the Add Item >Network Activity >Event Searches menu. The name of the event search item matches the name of the saved search criteria the item is based on.</p> <p>You can display a custom dashboard item that is based on saved search criteria from the Log Activity tab. Event search items are listed in the Add Item >Log Activity >Event Searches menu. The name of the event search item matches the name of the saved search criteria the item is based on.</p> <p>JSA includes default saved search criteria that is preconfigured to display event search items on your Dashboard tab menu. You can add more event search dashboard items to your Dashboard tab menu. For more information, see Adding search-based dashboard items to the Add Items list.</p> <p>On a Log Activity dashboard item, search results display real time last-minute data on a chart. The supported chart types are time series, table, pie, and bar. The default chart type is bar. These charts are configurable.</p> <p>Time series charts are interactive. You can magnify and scan through a timeline to investigate log activity.</p>
Events By Severity	<p>The Events By Severity dashboard item displays the number of active events that are grouped by severity. This item will allow you to see the number of events that are received by the level of severity assigned. Severity indicates the amount of threat an offense source poses in relation to how prepared the destination is for the attack. The range of severity is 0 (low) to 10 (high). The supported chart types are Table, Pie, and Bar.</p>
Top Log Sources	<p>The Top Log Sources dashboard item displays the top 5 log sources that sent events to JSA within the last 5 minutes.</p> <p>The Top Log Sources dashboard item displays the top 5 log sources that sent events to JSA Log Manager within the last 5 minutes.</p> <p>The number of events that are sent from the specified log source is indicated in the pie chart. This item will allow you to view potential changes in behavior, for example, if a firewall log source that is typically not in the top 10 list now contributes to a large percentage of the overall message count, you should investigate this occurrence. The supported chart types are Table, Pie, and Bar.</p>

System Summary

The **System Summary** dashboard item provides a high-level summary of activity within the past 24 hours.

Within the summary item, you can view the following information:

- **Current Flows Per Second** Displays the flow rate per second.
- **Flows (Past 24 Hours)** Displays the total number of active flows that are seen within the last 24 hours.
- **Current Events Per Second** Displays the event rate per second.
- **New Events (Past 24 Hours)** Displays the total number of new events that are received within the last 24 hours.
- **Updated Offenses (Past 24 Hours)** Displays the total number of offenses that have been either created or modified with new evidence within the last 24 hours.
- **Data Reduction Ratio** Displays the ratio of data reduced based on the total events that are detected within the last 24 hours and the number of modified offenses within the last 24 hours.

Risk Monitoring Dashboard

You use the **Risk Monitoring** dashboard to monitor policy risk and policy risk change for assets, policies and policy groups.

By default, the **Risk Monitoring** dashboard displays **Risk** and **Risk Change** items that monitor the policy risk score for assets in the High Vulnerabilities, Medium Vulnerabilities, and Low Vulnerabilities policy groups, as well as compliance pass rates and historical changes in policy risk score in the CIS policy group.

The Risk Monitoring dashboard items do not display any results unless JSA Risk Manager is licensed. For more information, see JSA Risk Manager Users Guide.

To view the default **Risk Monitoring** dashboard, select **Show Dashboard >Risk Monitoring** on the **Dashboard** tab.

Monitoring Policy Compliance

Create a dashboard item that shows policy compliance pass rates and policy risk score for selected assets, policies, and policies groups.

1. Click the **Dashboard** tab.
2. On the toolbar, click **New Dashboard**.
3. Type a name and description for your policy compliance dashboard.
4. Click **OK**.
5. On the toolbar, select **Add Item >Risk Manager >Risk**.

Risk Manager dashboard items are displayed only when JSA Risk Manager is licensed.

6. On the header of the new dashboard item, click the yellow **Settings** icon.
7. Use the **Chart Type**, **Display Top**, and **Sort** lists to configure the chart.
8. From the **Group** list, select the group that you want to monitor. For more information, see the table in step 9.

When you select the **Asset** option, a link to the **Risks >Policy Management >By Asset** page appears at the bottom of the **Risk** dashboard item. The **By Asset** page displays more detailed information about all results that are returned for the selected **Policy Group**. For more information on a specific asset, select **Table** from **Chart Type** list and click the link in the **Asset** column to view details about the asset in the **By Asset** page.

When you select the **Policy** option, a link to the **Risks >Policy Management >By Policy** page appears at the bottom of the **Risk** dashboard item. The **By Policy** page displays more detailed information about all results that are returned for the selected **Policy Group**. For more information on a specific policy, select **Table** from **Chart Type** list and click the link in the **Policy** column to view details about the policy in the **By Policy** page.

9. From the **Graph** list, select the graph type that you want to use. For more information, see the following table:

Group	Asset Passed Percentage	Policy Checks Passed Percentage	Policy Group Passed Percentage	Policy Risk Score
All	Returns the average asset percentage pass rate across assets, policies, and the policy group.	Returns the average policy check percentage pass rate across assets, policies, and the policy group.	Returns the average policy group pass rate across all assets, policies, and the policy group.	Returns the average policy risk score across all assets, policies, and the policy group.
Asset	Returns whether an asset passes asset compliance (100%=passed, 0%=failed). Use this setting to show which assets associated with a Policy Group pass compliance.	Returns percentage of policy checks that an asset passes. Use this setting to show the percentage of policy checks that passed for each asset that is associated with the Policy Group.	Returns the percentage of policy subgroups that are associated with the asset that pass compliance.	Returns the sum of all importance factor values for policy questions that are associated with each asset. Use this setting to view the policy risk for each asset that is associated with a selected policy group.
Policy	Returns whether all the assets associated with each policy in a Policy group pass compliance. Use this setting to monitor whether all the assets associated with each policy in a Policy Group pass or not.	Returns percentage of policy checks that pass per policy in the policy group. Use this setting to monitor how many policy checks are failing per policy.	Returns the percentage of policy subgroups of which the policy is a part that pass compliance.	Returns the importance factor values for each policy question in the Policy group. Use this setting to view the importance factor for each policy in a policy group.
Policy Group	Returns the percentage of assets that pass compliance for the selected Policy Group as a whole.	Returns the percentage of policy checks that pass per policy for the policy group as a whole.	Returns the percentage of policy subgroups within the Policy Group that pass compliance.	Returns the sum of all importance factor values for all policy questions in the Policy group.

10. From the **Policy Group** list, select the policy groups that you want to monitor.
11. Click **Save**.

Monitoring Risk Change

Create a dashboard item that shows policy risk change for selected assets, policies, and policies groups on a daily, weekly, and monthly basis.

Use this dashboard item to compare changes in the Policy Risk Score, Policies Checks, and Policies values for a policy group over time.

The **Risk Change** dashboard item uses arrows to indicate where policy risk for selected values that increased, decreased, or stayed the same over a chosen time period:

- The number beneath the red arrow indicates the values that show an increased risk.
- The number beneath the gray arrows indicates the values where there is no change in risk.
- The number beneath the green arrow indicates the values that show a decreased risk.

1. Click the **Dashboard** tab.
2. On the toolbar, click **New Dashboard**.
3. Type a name and description for your historical policy compliance dashboard.
4. Click **OK**.
5. On the toolbar, select **Add Item >Risk Manager >Risk Change**.

Risk Manager Dashboard items are displayed only when JSA Risk Manager is licensed.

6. On the header of the new dashboard item, click the yellow **Settings** icon.
7. From the **Policy Group** list, select the policy groups that you want to monitor.
8. Select an option from the **Value To Compare** list:
 - If you want to see the cumulative changes in importance factor for all policy questions within the selected policy groups, select **Policy Risk Score**.
 - If you want to see how many policy checks changed within the selected policy groups, select **Policies Checks**.
 - If you want to see how many policies changed within the selected policy groups, select **Policies**.

9. Select the risk change period that you want to monitor from the **Delta Time** list:
 - If you want to compare risk changes from 12:00 a.m. today with yesterday's risk changes, select **Day**.
 - If you want to compare risk changes from Monday 12:00 a.m. this week with last week's risk changes, select **Week**.
 - If you want to compare risk changes from the 12:00 a.m. on the first day of the current month with last month's risk changes, select **Month**.
10. Click **Save**.

Vulnerability Management Items

Vulnerability Management dashboard items are only displayed when JSA Vulnerability Manager is purchased and licensed.

For more information, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

You can display a custom dashboard item that is based on saved search criteria from the **Vulnerabilities** tab. Search items are listed in the **Add Item >Vulnerability Management >Vulnerability Searches** menu. The name of the search item matches the name of the saved search criteria the item is based on.

JSA includes default saved search criteria that is preconfigured to display search items on your **Dashboard tab** menu. You can add more search dashboard items to your **Dashboard tab** menu.

The supported chart types are table, pie, and bar. The default chart type is bar. These charts are configurable.

System Notification

The Systems Notification dashboard item displays event notifications that are received by your system.

For notifications to show in the **System Notification** dashboard item, the Administrator must create a rule that is based on each notification message type and select the **Notify** check box in the Custom Rules Wizard.

For more information about how to configure event notifications and create event rules, see the *Juniper Secure Analytics Administration Guide*.

On the **System Notifications** dashboard item, you can view the following information:

- **Flag** Displays a symbol to indicate severity level of the notification. Point to the symbol to view more detail about the severity level.
 - **Health** icon
 - **Information** icon (?)
 - **Error** icon (X)
 - **Warning** icon (!)
- **Created** Displays the amount of time elapsed since the notification was created.
- **Description** Displays information about the notification.
- **Dismiss icon (x)** Will allow you to close a system notification.

You can point your mouse over a notification to view more details:

- **Host IP** Displays the host IP address of the host that originated the notification.
- **Severity** Displays the severity level of the incident that created this notification.
- **Low Level Category** Displays the low-level category that is associated with the incident that generated this notification. For example: Service Disruption.
- **Payload** Displays the payload content that is associated with the incident that generated this notification.
- **Created** Displays the amount of time elapsed since the notification was created.

When you add the **System Notifications** dashboard item, system notifications can also display as pop-up notifications in the JSA user interface. These pop-up notifications are displayed in the lower right corner of the user interface, regardless of the selected tab.

Pop-up notifications are only available for users with administrative permissions and are enabled by default. To disable pop-up notifications, select **User Preferences** and clear the **Enable Pop-up Notifications** check box.

In the **System Notifications** pop-up window, the number of notifications in the queue is highlighted. For example, if (1 - 12) is displayed in the header, the current notification is 1 of 12 notifications to be displayed.

The **system notification** pop-up window provides the following options:

- **Next icon (>)** Displays the next notification message. For example, if the current notification message is 3 of 6, click the icon to view 4 of 6.
- **Close icon (X)** - Closes this notification pop-up window.

- **(details)** Displays more information about this system notification.

RELATED DOCUMENTATION

[Creating a Custom Dashboard | 34](#)

[Using the Dashboard to Investigate Log or Network Activity | 35](#)

Creating a Custom Dashboard

You can create a custom dashboard to view a group of dashboard items that meet a particular requirement.

After you create a custom dashboard, the new dashboard is displayed in the **Dashboard** tab and is listed in the **Show Dashboard** list box. A new custom dashboard is empty by default; therefore, you must add items to the dashboard.

1. Click the **Dashboard** tab.
2. Click the **New Dashboard** icon.
3. In the **Name** field, type a unique name for the dashboard. The maximum length is 65 characters.
4. In the **Description** field, type a description of the dashboard. The maximum length is 255 characters. This description is displayed in the tooltip for the dashboard name in the **Show Dashboard** list box.
5. Click **OK**.

RELATED DOCUMENTATION

[Using the Dashboard to Investigate Log or Network Activity | 35](#)

[Configuring Dashboard Chart Types | 36](#)

[Removing Dashboard Items | 37](#)

Using the Dashboard to Investigate Log or Network Activity

Search-based dashboard items provide a link to the **Log Activity** or **Network Activity** tabs, allowing you to further investigate log or network activity.

To investigate flows from a **Log Activity** dashboard item:

1. Click the **View in Log Activity** link. The **Log Activity** tab is displayed, displaying results and two charts that match the parameters of your dashboard item.

To investigate flows from a **Network Activity** dashboard item:

1. Click the **View in Network Activity** link. The **Network Activity** tab is displayed, displaying results and two charts that match the parameters of your dashboard item.

The chart types that are displayed on the **Log activity** or **Network Activity** tab depend on which chart is configured in the dashboard item:

Chart type	Description
Bar, Pie, and Table	The Log Activity or Network Activity tab displays a bar chart, pie chart, and table of flow details.
Time Series	<p>The Log Activity or Network Activity tab displays charts according to the following criteria:</p> <ol style="list-style-type: none"> 1. If your time range is less than or equal to 1 hour, a time series chart, a bar chart, and a table of event or flow details are displayed. 2. If your time range is more than 1 hour, a time series chart is displayed and you are prompted to click Update Details. This action starts the search that populates the event or flow details and generates the bar chart. When the search completes, the bar chart and table of event or flow details are displayed.

RELATED DOCUMENTATION

[Configuring Dashboard Chart Types | 36](#)

[Removing Dashboard Items | 37](#)

[Detaching a Dashboard Item | 37](#)

Configuring Dashboard Chart Types

You can configure different dashboard chart types for presenting your organization's data in meaningful ways.

Alternatively, use the JSA Pulse dashboard app to communicate insights and analysis about your network. Visualize offenses, network data, threats, and malicious user behavior from around the world in geographical maps, a built-in 3D threat globe, and auto updating charts.

1. Click the **Dashboard** tab.
2. From the **Show Dashboard** list box, select the dashboard that contains the item you want to customize.
3. On the header of the dashboard item you want to configure, click the **Settings** icon.
4. Configure the chart parameters.
 - a. From the **Value to Graph** list box, select the object type that you want to graph on the chart. Options include all normalized and custom event or flow parameters that are included in your search parameters.
 - b. Select a chart type:
 - Bar, pie, and table charts are only available for grouped events or flows.
 - Data accumulates so that when you run a time series saved search, a cache of event or flows data is available to display the data for the previous time period. Accumulated parameters are indicated by an asterisk (*) in the **Value to Graph** list box. If you select a value to graph that is not accumulated (no asterisk), time series data is not available.
 - Select the **Capture Time Series** Data checkbox to enable time series capture. When you select this checkbox, the chart feature accumulates data for time series charts. By default, this option is disabled.

Your custom chart configurations are retained so that they are displayed as configured each time that you access the **Dashboard** tab.

RELATED DOCUMENTATION

[Removing Dashboard Items | 37](#)

[Detaching a Dashboard Item | 37](#)

[Renaming a Dashboard | 38](#)

Removing Dashboard Items

You can remove items from a dashboard and add the item again at any time.

When you remove an item from the dashboard, the item is not removed completely.

1. Click the **Dashboard** tab.
2. From the **Show Dashboard** list box, select the dashboard from which you want to remove an item.
3. On the dashboard item header, click the red [x] icon to remove the item from the dashboard.

RELATED DOCUMENTATION

[Detaching a Dashboard Item | 37](#)

[Renaming a Dashboard | 38](#)

[Deleting a Dashboard | 38](#)

Detaching a Dashboard Item

You can detach an item from your dashboard and display the item in a new window on your desktop system.

When you detach a dashboard item, the original dashboard item remains on the **Dashboard** tab, while a detached window with a duplicate dashboard item remains open and refreshes during scheduled intervals. If you close the JSA application, the detached window remains open for monitoring and continues to refresh until you manually close the window or shut down your computer system.

1. Click the **Dashboard** tab.
2. From the **Show Dashboard** list box, select the dashboard from which you want to detach an item.
3. On the dashboard item header, click the green icon to detach the dashboard item and open it in separate window.

RELATED DOCUMENTATION

[Renaming a Dashboard | 38](#)

[Deleting a Dashboard | 38](#)

[Managing System Notifications | 39](#)

Renaming a Dashboard

You can rename a dashboard and update the description.

1. Click the **Dashboard** tab.
2. From the **Show Dashboard** list box, select the dashboard that you want to edit.
3. On the toolbar, click the **Rename Dashboard** icon.
4. In the **Name** field, type a new name for the dashboard. The maximum length is 65 characters.
5. In the **Description** field, type a new description of the dashboard. The maximum length is 255 characters
6. Click **OK**.

RELATED DOCUMENTATION

[Deleting a Dashboard | 38](#)

[Managing System Notifications | 39](#)

[Adding Search-based Dashboard Items to the Add Items List | 39](#)

Deleting a Dashboard

You can delete a dashboard.

After you delete a dashboard, the **Dashboard** tab refreshes and the first dashboard that is listed in the **Show Dashboard** list box is displayed. The dashboard that you deleted is no longer displayed in the **Show Dashboard** list box.

1. Click the **Dashboard** tab.
2. From the **Show Dashboard** list box, select the dashboard that you want to delete.
3. On the toolbar, click **Delete Dashboard**.

4. Click **Yes**.

RELATED DOCUMENTATION

[Managing System Notifications | 39](#)

[Adding Search-based Dashboard Items to the Add Items List | 39](#)

[Renaming a Dashboard | 38](#)

Managing System Notifications

Ensure the **System Notification** dashboard item is added to your dashboard.

You can specify the number of notifications that you want to display on your **System Notification** dashboard item and close system notifications after you read them.

1. On the System Notification dashboard item header, click the **Settings** icon.
2. From the **Display** list box, select the number of system notifications you want to view.
 - The options are **5**, **10** (default), **20**, **50**, and **All**.
 - To view all system notifications that are logged in the past 24 hours, click **All**.
3. To close a system notification, click the **Delete** icon.

RELATED DOCUMENTATION

[Adding Search-based Dashboard Items to the Add Items List | 39](#)

[Renaming a Dashboard | 38](#)

[Deleting a Dashboard | 38](#)

Adding Search-based Dashboard Items to the Add Items List

To add an event and flow search dashboard item to the **Add Item** menu on the **Dashboard** tab, you must access the **Log Activity** or **Network Activity** tab to create search criteria that specifies that the search

results can be displayed on the **Dashboard** tab. The search criteria must also specify that the results are grouped on a parameter.

You can add search-based dashboard items to your **Add Items** menu.

1. Choose:
 - To add a flow search dashboard item, click the **Network Activity** tab.
 - To add an event search dashboard item, click the **Log Activity** tab.
2. From the **Search** list box, choose one of the following options:
 - To create a search, select **New Search**.
 - To edit a saved search, select **Edit Search**.
3. Configure or edit your search parameters, as required.
 - On the Edit Search pane, select the **Include in my Dashboard** option.
 - On the Column Definition pane, select a column and click the **Add Column** icon to move the column to the **Group By** list.
4. Click **Filter**. The search results are displayed.
5. Click **Save Criteria**. See Saving search criteria on the Offense tab.
6. Click **OK**.
7. Verify that your saved search criteria successfully added the event or flow search dashboard item to the **Add Items** list
 - a. Click the **Dashboard** tab.
 - b. Choose one of the following options:
 - a. To verify an event search item, select **Add Item > Log Activity >Event Searches >Add Item**.
 - b. To verify a flow search item, select **Add Item >Network Activity >Flow Searches**.

The dashboard item is displayed on the list with the same name as your saved search criteria.

RELATED DOCUMENTATION

[Renaming a Dashboard | 38](#)

[Deleting a Dashboard | 38](#)

[Managing System Notifications | 39](#)

4

CHAPTER

QRadar Analyst Workflow

QRadar Analyst Workflow | 42

What's New in QRadar Analyst Workflow | 43

Known Issues | 44

Installing the Stand-Alone Version | 44

Removing the Stand-Alone Version | 45

Installing the UBI App Version | 46

Offenses | 47

Querying Event and Flow Data to Find Specific Offenses | 52

Searching with the Visual Query Builder | 53

Events | 54

QRadar Analyst Workflow

QRadar Analyst Workflow provides new methods for filtering offenses and events, and graphical representations of offenses, by magnitude, assignee, and type. The improved offenses workflow provides a more intuitive method to investigate offense to determine the root cause of an issue and work to resolve it. Use the built-in query builder to create AQL queries by using examples and saved or shared searches, or by typing plain text into the search field.

Offenses

The Offenses page displays a table of the offenses in your JSA environment that you can filter in many different ways. It also includes graphical representations of offenses, by magnitude, assignee, and type. From this page, you can investigate an offense to determine the root cause of an issue and work to resolve it.

Search

The Search page includes a Query Builder that you can use to build an Ariel Query Language (AQL) search to find specific offenses. Create a search using examples, saved or shared searches, or typing directly into the Query Builder. The Search page also includes links to many resources to learn about creating AQL queries.

Apps

The Apps list includes JSA apps that are compatible with the new Analyst Workflow. The first release of the workflow includes the Dashboards (Pulse) app. QRadar Pulse is a dashboard app that you can use to communicate insights and analysis about your network. For more information, see the *Pulse App Guide*.

RELATED DOCUMENTATION

[What's New in QRadar Analyst Workflow | 43](#)

[Known Issues | 44](#)

[Installing the Stand-Alone Version | 44](#)

[Offenses | 47](#)

[Querying Event and Flow Data to Find Specific Offenses | 52](#)

[Events | 54](#)

What's New in QRadar Analyst Workflow

Version 2.6.5

NOTE: This version of QRadar Analyst Workflow is only supported on JSA 7.4.3 Fix Pack 1 and later. Unlike previous versions of QRadar Analyst Workflow, which was installed through the command line, this extension can be installed through the Extensions Management tool, just like any other QRadar app. If you upgrade to JSA 7.4.3 Fix Pack 1 from an earlier version of JSA, you must do a fresh installation of this version of QRadar Analyst Workflow.

This release resolves the following new features:

- Added a toggle switch to the Event and Flow panels to hide or show empty custom properties.
- UI design improvements.
- Accessibility improvements.
- Performance improvements.

This release resolves the following known issues:

- Fixed an issue that produced an error when a query was started before the Events table was fully loaded.
- Fixed an issue that displayed a filter as "undefined" when the user applied a filter on a table item.

RELATED DOCUMENTATION

[Known Issues | 44](#)

[Installing the Stand-Alone Version | 44](#)

[Offenses | 47](#)

[Querying Event and Flow Data to Find Specific Offenses | 52](#)

[Events | 54](#)

Known Issues

QRadar Analyst Workflow contains required information for known issues.

QRadar Analyst Workflow 2.6.5 contains the following known issue:

- If you view event data by using the custom rules redirect, the query does not display.

RELATED DOCUMENTATION

[Installing the Stand-Alone Version | 44](#)

[Offenses | 47](#)

[Querying Event and Flow Data to Find Specific Offenses | 52](#)

[Events | 54](#)

Installing the Stand-Alone Version

You can install QRadar Analyst Workflow on JSA 7.4.0 or later.

The following instructions describe the installation process for JSA versions 7.4.0 to 7.4.3 GA only. For installations with JSA version 7.4.3 Fix Pack 1 and later, QRadar Analyst Workflow is installed as a standard application by using extensions management. For more information, see "[Installing the UBI App Version](#)" on page 46.

The stand-alone version of QRadar Analyst Workflow requires root access to install. If you are using the command line to enable root user privileges, you must use the following command:

```
sudo su -
```

If you use `sudo su` (without `-`), full root access is not granted.

1. If you have custom certificates, run the following commands on your JSA Console, in any directory:
 - `update-ca-trust`
 - `systemctl restart docker`
2. Download the latest version of the QRadar Analyst Workflow<x.x.x>.zip file from [Juniper Downloads](#). See instructions on the IBM Security App Exchange.

3. Copy the file onto your JSA host by using the Linux "secure copy" (scp) command or an FTP client.

Secure copy example: `scp QRadarAnalystWorkflow<x.x.x>.zip <QRadar host>:/<directory>`

4. Type the following command to create a new directory on the JSA host: `mkdir qradar-ui`

NOTE: If the directory exists from a previous installation, you must delete it before you extract the .zip file.

5. To extract the latest version of the `QRadarAnalystWorkflow<x.x.x>.zip` file on your JSA host, type the following command:

```
unzip QRadarAnalystWorkflow<x.x.x>.zip -d qradar-ui
```

6. Run `./qradar-ui/start.sh`, then wait for the logs to run.
7. Access the QRadar Analyst Workflow by using one of the following methods:
 - In the navigation menu, click **Try the New UI**.
 - Access the new UI in your browser at `https://<QRadar IP address>/console/ui`.

RELATED DOCUMENTATION

[Offenses | 47](#)

[Querying Event and Flow Data to Find Specific Offenses | 52](#)

[Events | 54](#)

[QRadar Analyst Workflow | 42](#)

[What's New in QRadar Analyst Workflow | 43](#)

Removing the Stand-Alone Version

If you upgrade to JSA version 7.4.3 Fix Pack 1 or later, you must uninstall the QRadar Analyst Workflow stand-alone version.

To remove the QRadar Analyst Workflow stand-alone version, run the following commands:

```
/opt/ibm/si/conman/bin/conman-api-cli.sh stop -n ui -f
```

```
/opt/ibm/si/conman/bin/conman-api-cli.sh stop -n graphql -f
```

After you uninstall the stand-alone version, you can ["Install the UBI app version" on page 46](#).

RELATED DOCUMENTATION

| [Installing the Stand-Alone Version | 44](#)

Installing the UBI App Version

After you download the Red Hat Universal Base Images (UBI) app version of the QRadar Analyst Workflow from the IBM Security App Exchange, install it by using the **Extensions Management** page on your JSA computer.

You must have an IBM ID to access the App Exchange.

NOTE: The UBI version is supported on JSA 7.4.3 Fix Pack 1 or later.

NOTE: If you upgrade to JSA version 7.4.3 Fix Pack 1 or later, you must uninstall the QRadar Analyst Workflow stand-alone version before you install the UBI app version. For instructions, see ["Removing the Stand-Alone Version" on page 45](#).

1. Download the QRadar Analyst Workflow UBI app archive from the IBM Security App Exchange.
2. On the JSA Console, click **Admin > Extensions Management**.
3. In the **Extension Management** window, click **Add** and select the app archive that you want to upload to the console.
4. Select the **Install immediately** checkbox.

NOTE: You might have to wait several minutes before your app becomes active.

5. To preview the contents of an app after it is added and before it is installed, select it from the list of extensions, and click **More Details**. Expand the folders to view the individual content items in each group.
6. Access the QRadar Analyst Workflow by using one of the following methods:

- In the navigation menu, click **Try the New UI**.
- Access the new UI in your browser at `https://<QRadar IP address>/console/ui`.

When the installation is complete, clear your browser cache and refresh the browser window before you use the app.

Offenses

IN THIS SECTION

- [Visualization of Offenses | 47](#)
- [Offense Investigation | 49](#)
- [Offense Actions | 50](#)

The QRadar Analyst Workflow Offenses overview page displays a table of the offenses in your JSA environment that you can filter in many different ways. It also includes graphical representations of offenses, by magnitude, assignee, and type.

On the offenses page, you can investigate an offense to determine the root cause of an issue and work to resolve it.

Visualization of Offenses

Filter the Offenses table in the QRadar Analyst Workflow to display the specific offenses you want to investigate.

As you apply filters, the offenses table displays only the offenses that meet your filter criteria. The graphs displayed on the page also change to reflect only the offenses in your filtered list.

TIP: You can copy and paste the URL from your browser to share the offenses page, including all filters and configuration options.

1. To apply a filter, click any of the following categories to see filtering options for that category:

- Magnitude
- Severity
- Assigned To
- Status
- Start Time
- Offense Type
- Log Source Name
- Log Source Type
- Destination Network
- Local Destination Addresses
- Source Addresses
- Rules
- Follow Up
- Protected

2. To include only offenses with specific attributes, select that attribute in the filters list. To exclude offenses with specific attributes, click the icon next to the attribute, and click **Apply IS NOT Filter**.

TIP: You can right-click on a Status, Type, Source IP, or Destination IP in the offenses table and quickly apply an IS or IS NOT filter to the offenses.

3. To sort the offenses table in ascending or descending order by an attribute, click the appropriate table heading.
4. To clear individual filters, click the **X** on the filter indicator. To clear all filters, click **Clear filters**.
5. To configure the number of offenses displayed in the table, click the **Items per page** drop-down at the bottom of the table.
6. To sort the offenses table in ascending or descending order by an attribute, click the appropriate table heading.

Offense Investigation

Begin your offense investigation in the QRadar Analyst Workflow by clicking an offense in the offense table. The offense details provide context to help you understand what happened and determine how to isolate and resolve the problem.

In addition to the basic information included in the offense table, the offense details page includes the following detailed information:

Feature	Description
Insights	The Insights section displays rules that triggered the event. Click a rule to see details about specific rules.
Events graph	The Events graph displays the number of events that occurred at a given time within the last 7 active days. Use the scrubber bar at the top of the graph to zoom in on specific times and event spikes. Click View Events to see a list of events that contributed to the offense and investigate event details.
Source and Destination IPs	If offenses include multiple source or destination IPs, you can click the IP lists to scroll through the entire list of IPs. Click a specific IP address to see details about that IP.
Magnitude	The Magnitude graph provides a visual representation of how the magnitude was calculated, based on relevance, credibility, and severity. Click the graph to see a detailed description of how the magnitude is calculated.
Notes	In the Notes section, you can click on a long note to see the entire text. Click Add note to add your own note to the offense details.

TIP: If an offense has a long title, click on the title to see the entire offense title.

Offense Actions

Use the QRadar Analyst Workflow to keep track of offenses throughout your investigation.

Knowing that an offense occurred is only the first step; identifying how it happened, where it happened, and who did it requires some investigation.

Marking an offense for follow-up

In the QRadar Analyst Workflow, you can mark an offense for follow-up when you want to flag it for further investigation.

1. From the Offenses table, do one of the following actions:
 - Select the offenses that you want to flag.
 - Click on a single offense listing to open the offense details.
2. From the **Actions** list, select **Follow up**.

TIP: To remove the flag, select **Unfollow** from the **Actions** list.

Protecting Offenses

You might have offenses that you want to retain regardless of the retention period. In the QRadar Analyst Workflow, you can protect offenses to prevent them from being removed from JSA after the retention period has elapsed.

By default, offenses are retained for thirty days. For more information about customizing the offense retention period, see the *Juniper Secure Analytics Administration Guide*.

1. From the Offenses table, do one of the following:
 - Select any offenses you want to protect.
 - Click on a single offense listing to open the offense details.
2. From the **Actions** list, select **Protect**.

TIP: To remove the protection from the offense, select **Unprotect** from the **Actions** list.

Hiding Offenses

Hide an offense to prevent it from being displayed in the QRadar Analyst Workflow offenses table. After you hide an offense, the offense is only displayed if you apply an IS filter for **Status = Hidden**.

1. From the Offenses table, do one of the following:
 - Select any offenses you want to hide.
 - Click on a single offense listing to open the offense details.
2. From the **Actions** list, select **Hide**.

TIP: To unhide the offense, filter to see hidden offenses, and select **Open** from the **Actions** list.

Closing Offenses

Close an offense in the QRadar Analyst Workflow to remove it completely from your system.

The default offense retention period is 30 days. After the offense retention period expires, closed offenses are deleted from the system. You can protect an offense to prevent it from being deleted when the retention period expires.

After you close an offense, the offense is only displayed if you apply an IS filter for **Status = Closed**. If more events occur for an offense that is closed, a new offense is created.

When you close offenses, you must select a reason for closing the offense. If you have the **Manage Offense Closing** permission, you can add custom closing reasons. For more information about user role permissions, see the *Juniper Secure Analytics Administration Guide*.

1. From the Offenses table, do one of the following:
 - Select any offenses you want to close.
 - Click on a single offense listing to open the offense details.
2. From the **Actions** list, select **Close**.
3. Specify a closing reason from the **Choose a resolution option** list.
4. In the text field, type a note to provide more information.

Notes must not exceed 1,984 characters.
5. Click **OK**.

RELATED DOCUMENTATION

[Querying Event and Flow Data to Find Specific Offenses | 52](#)

[Events | 54](#)

[Known Issues | 44](#)

[Installing the Stand-Alone Version | 44](#)

Querying Event and Flow Data to Find Specific Offenses

Search for specific event and flow data by creating Ariel Query Language (AQL) searches in the QRadar Analyst Workflow Query Builder.

Create searches by using search history or entering keywords directly into the Query Builder. This information populates a query template that you can further customize to suit your needs, or manually create your own searches.

NOTE: To build a query without using AQL, try the "[Visual Query Builder](#)" on page 53.

1. From the navigation menu, click **Search** and select the **Advanced builder** tab.
2. Type one of the following keywords in the **Query Builder** to start a query:
 - IP address
 - URL
 - MD5/SHA-1/SHA-256 hash
3. Select one of the predefined searches from the list that appears as you enter a keyword.
4. Review and edit the query template to refine your search, and then click **Run query**.

TIP:

- Syntax tokens are color-coded based on token class.

- For a syntactically correct AQL string, paired parentheses are underscored when the cursor is placed between them.

```
(startTime, 'MMM dd hh:mm a')
```

5. Click **Filter** to further refine your search results and then select an offense to view more details.
6. To run an existing search result, select the query in the **Last Search** field to add it to the Query Builder, and then click **Run query**.
7. Optional: Expand the **Training and resources** section to learn more about AQL queries.

The following is an example of an AQL query.

```
SELECT sourceip, destinationip, username
FROM events
WHERE username = 'test name'
GROUP by sourceip, destinationip
ORDER BY sourceip DESC
LIMIT 10
LAST 2 DAYS
```

RELATED DOCUMENTATION

[Events | 54](#)

[Known Issues | 44](#)

[Installing the Stand-Alone Version | 44](#)

[Offenses | 47](#)

Searching with the Visual Query Builder

You can use the visual builder to search for specific event and flow data without any knowledge of Ariel Query Language (AQL).

1. From the navigation menu, click **Search**, and select the **Visual builder** tab.
2. If you want to enable quick filtering of simple terms from your results, toggle the **Quick Filter** switch, type a value to filter, and press Enter.

3. Select a boolean **IF** operator for your condition set.
 - **AND** includes only data that meets all of the conditions in your set.
 - **OR** includes data that meets any of the conditions in your set.
4. Select a **Condition** to search for.

Examples include source or destination IP addresses, geographic locations, event names, log sources, and many more.
5. Select an Operator, such as **Equals**, **Does not equal**, or a value specific to that condition.
6. Type or select a **Value**, or use the toggle icon to display child values.
7. To add more conditions to your set, click **Add another condition**, then repeat Steps "4" on page 53 to "6" on page 53 to define each condition.
8. To add another condition set to your query, click **Add condition set**, select the **AND** or **OR** boolean operator, then repeat Steps "4" on page 53 to "6" on page 53 to define each condition in the new set.
9. After you finish defining conditions, click **Run Search**.

Events

IN THIS SECTION

- [Investigating Events | 55](#)
- [Filtering Events | 55](#)

Use the Events page to further investigate specific events to determine the root cause of an issue and work to resolve it.

The Events page displays a table of the events that contributed to a specific offense and a Time Series chart that shows the number of events sorted by date. You can filter these events to suit your needs.

Investigating Events

The Events graph on the offense details page displays the number of events that occurred at a given time within the last 7 active days.

1. From the offenses page, click on an offense in the offense table to open the details page.

TIP: Use the scrubber bar at the top of the Events graph to zoom in on specific times and event spikes.

2. Click **View Events** to see a list of events that contributed to the offense and investigate event details.
3. To configure the number of events returned in your filter results, click the arrows in the Result Limit indicator.
4. To configure the number of events displayed in the table, click the Items per page drop-down at the bottom of the table.
5. To sort the events table in ascending or descending order by an attribute, click the appropriate table heading.
6. Click on an event to see more details about that event. You can also click on a log source, source IP, or destination IP for specific information on that source or destination.
7. Click **Update events** to refresh the events results.

TIP: You can copy and paste the URL from your browser to share the events page, including all filters and configuration options.

Filtering Events

Filter the Events page to display only the specific events you want to investigate.

As you apply filters, the events table displays only the events that meet your filter criteria.

TIP: You can copy and paste the URL from your browser to share the events page, including all filters and configuration options.

1. To apply a filter, click any of the following categories to see filtering options for that category:
 - Event Time
 - Magnitude
 - Log Source Name
 - Category
 - Source IP
 - Source Port
 - Destination IP
 - Destination Port
 - Event Name
 - User
2. To include only events with specific attributes, select that attribute in the filters list. To exclude events with specific attributes, click the vertical ellipsis icon next to the attribute, and click **Apply IS NOT Filter**.

TIP: You can right-click on a Log Source, Source IP, Destination IP, Category, or Username in the events table and quickly apply an IS or IS NOT filter to the events.

3. To sort the events table in ascending or descending order by an attribute, click the appropriate table heading.
4. To clear individual filters, click the close icon [x] on the filter indicator. To clear all filters, click **Clear filters**.
5. Click **Update events** to refresh the events results.

RELATED DOCUMENTATION

[Known Issues | 44](#)

[Installing the Stand-Alone Version | 44](#)

[Offenses | 47](#)

5

CHAPTER

Offense Management

Offense Management | 58

Offense Prioritization | 58

Offense Chaining | 59

Offense Indexing | 60

Offense Retention | 62

Offense Investigations | 65

Offense Actions | 79

Offense Management

JSA reduces billions of events and flows into a manageable number of actionable offenses that are prioritized by their impact on your business operations. Use the **Offenses** tab to access all of the data that you need to understand even the most complex threats.

By providing immediate context for the offense, JSA helps you to quickly identify which offenses are the most important, and to begin an investigation to find the source of the suspected security attack or policy breach.

TIP: You can also manage offenses using the QRadar Analyst Workflow interface. For more information, see ["QRadar Analyst Workflow" on page 42](#).

NOTE: You cannot manage offenses in Log Manager. For more information about the differences between JSA and Log Manager, see ["Capabilities in Your JSA Product" on page 7](#).

Offense Prioritization

The magnitude rating of an offense is a measure of the importance of the offense in your environment. JSA uses the magnitude rating to prioritize offenses and help you to determine which offenses to investigate first.

The *magnitude rating* of an offense is calculated based on relevance, severity, and credibility.

- *Relevance* determines the impact of the offense on your network. For example, if a port is open, the relevance is high.
- *Credibility* indicates the integrity of the offense as determined by the credibility rating that is configured in the log source. Credibility increases as multiple sources report the same event.
- *Severity* indicates the level of threat that a source poses in relation to how prepared the destination is for the attack.

JSA uses complex algorithms to calculate the offense magnitude rating, and the rating is re-evaluated when new events are added to the offense and also at scheduled intervals. The following information is considered when the offense magnitude is calculated:

- the number of events and flows that are associated with the offense
- the number of log sources
- the age of the offense
- the weight of the network object associated with the offense
- the categories, severity, relevance, and credibility of the events and flows that contribute to the offense
- the vulnerabilities and threat assessment of the hosts that are involved in the offense

The magnitude rating of an offense is different from the magnitude rating for an event. You can influence the magnitude of an offense by setting the event magnitude in the rule actions, but you cannot bypass the JSA algorithms to set the offense magnitude yourself.

RELATED DOCUMENTATION

[Offense Chaining](#) | 59

[Offense Indexing](#) | 60

[Offense Retention](#) | 62

Offense Chaining

JSA chains offenses together to reduce the number of offenses that you need to review, which reduces the time to investigate and remediate the threat.

Offense chaining helps you find the root cause of a problem by connecting multiple symptoms together and showing them in a single offense. By understanding how an offense changed over time, you can see things that might be overlooked during your analysis. Some events that would not be worth investigating on their own might suddenly be of interest when they are correlated with other events to show a pattern.

Offense chaining is based on the offense index field that is specified on the rule. For example, if your rule is configured to use the source IP address as the offense index field, there is only one offense that has that source IP address for while the offense is active.

You can identify a chained offense by looking for preceded by in the **Description** field on the **Offense Summary** page. In the following example, JSA combined all of the events that fired for each of the three rules into one offense, and appended the rule names to the **Description** field:

```
Exploit Followed By Suspicious Host Activity - Chained  
preceded by Local UDP Scanner Detected  
preceded by XForce Communication to a known Bot Command and Control
```

RELATED DOCUMENTATION

[Offense Indexing | 60](#)

[Offense Retention | 62](#)

[Offense Investigations | 65](#)

Offense Indexing

IN THIS SECTION

- [Offense Indexing Considerations | 61](#)
- [Example: Detecting Malware Outbreaks Based on the MD5 Signature | 62](#)

Offense indexing provides the capability to group events or flows from different rules indexed on the same property together in a single offense.

JSA uses the offense index parameter to determine which offenses to chain together. For example, an offense that has only one source IP address and multiple destination IP addresses indicates that the threat has a single attacker and multiple victims. If you index this type of offense by the source IP address, all events and flows that originate from the same IP address are added to the same offense.

You can configure rules to index an offense based on any piece of information. JSA includes a set of predefined, normalized fields that you can use to index your offenses. If the field that you want to index on is not included in the normalized fields, create a custom event or a custom flow property to extract the data from the payload and use it as the offense indexing field in your rule. The custom property that you index on can be based on a regular expression, a calculation, or an AQL-based expression.

Offense Indexing Considerations

It is important to understand how offense indexing impacts your JSA deployment.

System Performance

Ensure that you optimize and enable all custom properties that are used for offense indexing. Using properties that are not optimized can have a negative impact on performance.

When you create a rule, you cannot select non-optimized properties in the **Index offense based on** field. However, if an existing rule is indexed on a custom property, and then the custom property is de-optimized, the property is still available in the offense index list. Do not de-optimize custom properties that are used in rules.

Rule Action and Response

When the indexed property value is null, an offense is not created, even when you select the **Ensure the detected event is part of an offense** check box in the rule action. For example, if a rule is configured to create an offense that is indexed by host name, but the host name in the event is empty, an offense is not created even though all of the conditions in the rule tests are met.

When the response limiter uses a custom property, and the custom property value is null, the limit is applied to the null value. For example, if the response is **Email**, and the limiter says **Respond no more than 1 time per 1 hour per custom property**, if the rule fires a second time with a null property within 1 hour, an email will not be sent.

When you index using a custom property, the properties that you can use in the rule index and response limiter field depends on the type of rule that you are creating. An event rule accepts custom event properties in the rule index and response limiter fields, while a flow rule accepts only custom flow properties. A common rule accepts either custom event or custom flow properties in the rule index and response limiter fields.

You cannot use custom properties to index an offense that is created by a dispatched event.

Payload Contents

Offenses that are indexed by the Ariel Query Language (AQL), a regular expression (regex), or by a calculated property include the same payload as the initial event that generated the offense.

Offenses that are indexed by a normalized event field, such as Source IP or Destination IP, include the event name and description as the custom rules engine (CRE) payload.

Example: Detecting Malware Outbreaks Based on the MD5 Signature

As a network security analyst for a large organization, you use JSA to detect when a malware outbreak occurs. You set the criteria for an outbreak as a threat that occurs across 10 hosts within 4 hours. You want to use the MD5 signature as the basis for this threat detection.

You configure JSA to evaluate the incoming logs to determine whether a threat exists, and then you group all of the fired rules that contain the same MD5 signature into a single offense.

1. [Creating a Regex-based Custom Property](#) to extract the MD5 signature from the logs. Ensure that the custom property is optimized and enabled.
2. ["Creating a Custom Rule" on page 279](#) and configure the rule to create an offense that uses the MD5 signature custom property as the offense index field. When the rule fires, an offense is created. All fired rules that have the same MD5 signature are grouped into one offense.
3. You can ["Searching for Offenses That Are Indexed on a Custom Property" on page 249](#) to find the offenses that are indexed by the MD5 signature custom property.

RELATED DOCUMENTATION

[Offense Retention | 62](#)

[Offense Investigations | 65](#)

[Offense Actions | 79](#)

Offense Retention

IN THIS SECTION

- [Protecting Offenses | 63](#)
- [Unprotecting Offenses | 64](#)

The state of an offense determines how long JSA keeps the offense in the system. The offense retention period determines how long inactive and closed offenses are kept before they are removed from the JSA console.

- **Active offenses**--When a rule triggers an offense, the offense is active. In this state, JSA is waiting to evaluate new events or flows against the offense rule test. When new events are evaluated, the offense clock is reset to keep the offense active for another 30 minutes.
- **Dormant offenses**--An offense becomes dormant if new events or flows are not added to the offense within 30 minutes, or if JSA did not process any events within 4 hours. An offense remains in a dormant state for 5 days. If an event is added while an offense is dormant, the five-day counter is reset.
- **Inactive offenses**--An offense becomes inactive after 5 days in a dormant state. In the inactive state, new events that trigger the offense rule test do not contribute to the inactive offense. They are added to a new offense.

Inactive offenses are removed after the offense retention period elapses.

- **Closed offenses**--Closed offenses are removed after the offense retention period elapses. If more events occur for an offense that is closed, a new offense is created.

If you include closed offenses in a search, and the offense wasn't removed from the JSA console, the offense is displayed in the search results.

The default offense retention period is 30 days. After the offense retention period expires, closed and inactive offenses are removed from the system. Offenses that are not inactive or closed are retained indefinitely.

NOTE: System performance is negatively impacted when the system retains many inactive and closed offenses. For optimum performance, set the retention period for the least amount of time possible. The suggested retention period is 3 days.

To prevent an offense from being removed from the system, you can protect it. Before you protect offenses, consider the performance impact that it might have. Some offenses impact system performance more than others. For example, offenses with large numbers of events and flows have a greater impact on performance. Offenses that have many targets and destinations impact performance more than an offense that has only a single target or destination.

If you need to re-create an offense after it is removed from the system, run a historical correlation job to analyze the historical data. For more information, see "[Historical Correlation](#)" on page 300.

Protecting Offenses

You might have offenses that you want to retain regardless of the retention period. You can protect offenses to prevent them from being removed from JSA after the retention period has elapsed.

By default, offenses are retained for thirty days. For more information about customizing the offense retention period, see the *Juniper Secure Analytics Administration Guide*.

1. Click the **Offenses** tab, and click **All Offenses**.
2. Choose one of the following options:
 - Select the offense that you want to protect, and then select **Protect** from the **Actions** list.
 - From the **Actions** list box, select **Protect Listed**.
3. Click **OK**.

The offense is protected and will not be removed from JSA. In the **Offense** window, the protected offense is indicated by a **Protected** icon in the Flag column.

Unprotecting Offenses

You can unprotect offenses that were previously protected from removal after the offense retention period has elapsed.

To list only protected offenses, you can perform a search that filters for only protected offenses. If you clear the **Protected** check box and ensure that all other options are selected under the **Excludes option** list on the Search Parameters pane, only protected offenses are displayed.

1. Click the **Offenses** tab.
2. Click **All Offenses**.
3. Optional: Perform a search that displays only protected offenses.
4. Choose one of the following options:
 - Select the offense that you no longer want to protect, and then select **Unprotect** from the Actions list box.
 - From the **Actions** list box, select **Unprotect Listed**.
5. Click **OK**.

RELATED DOCUMENTATION

[Offense Investigations | 65](#)

[Offense Actions | 79](#)

Offense Investigations

IN THIS SECTION

- [Selecting an Offense to Investigate | 67](#)
- [Investigating an Offense by Using the Summary Information | 68](#)
- [Investigating Events | 76](#)
- [Investigating Flows | 77](#)

JSA uses rules to monitor the events and flows in your network to detect security threats. When the events and flows meet the test criteria that is defined in the rules, an offense is created to show that a security attack or policy breach is suspected. But knowing that an offense occurred is only the first step; identifying how it happened, where it happened, and who did it requires some investigation.

The **Offense Summary** window helps you begin your offense investigation by providing context to help you understand what happened and determine how to isolate and resolve the problem.

Figure 1: Offense Summary View

The screenshot displays the 'Offense 31' summary window. It includes several sections with callout boxes asking specific questions:

- Offense Details:** Magnitude (highlighted), Status, Relevance (5), Severity (0), Credibility (3). Description: 'Large Outbound Transfer Slow Rate of Transfer preceded by Large Outbound Transfer High Rate of Transfer containing unknown'. Source IP(s): 10.105.99.148 (10.105.99.148). Destination IP(s): 192.168.100.147. Network(s): pfprd. Offense Type: Source IP. EventFlow count: 58 events and 3. Start: Apr 13, 2016, 4:00. Duration: 4d 18h 15m. Assigned to: Unassigned. Callout: 'What was the attack?' and 'Was it successful?'.
- Offense Source Summary:** IP: 10.105.99.148. Location: 192.168.100.148. Magnitude: (highlighted). Vulnerabilities: 0. Username: Unknown. Root Name: Unknown. Asset Name: 10.105.99.148. Offenses: 1. Events/Flows: 3,528. Callout: 'Who was responsible?' and 'Where can I find them?'.
- Top 5 Source IPs:** Table with columns: Source IP, Magnitude, Offenses, Destinat..., Last EventFlow, EventsF. Callout: 'How many targets are involved?' and 'Are the targets vulnerable?'.
- Top 5 Destination IPs:** Table with columns: Destination IP, Magnitude, Location, Vulnerability, Chained, User, MAC, Weight, Offenses, Source(s), Last EventFlow, EventsF. Callout: 'How valuable are the targets to the business?'.
- Last 10 Events:** Table with columns: Event Name, Magnitude, Log Source, Category, Destination, Time. Events include 'Authentication Fail' and 'Root Login Failed'. Callout: 'Where is the evidence?'.
- Top 5 Annotations:** Callout: 'Why does JSA consider the event threatening?'.

JSA does not use device level user permissions to determine which offenses each user is able to view. All users who have access to the network can view all offenses regardless of which log source or flow source is associated with the offense. For more information about restricting network access, see the security profiles documentation in the *Juniper Secure Analytics Administration Guide*.

Selecting an Offense to Investigate

Use the navigation options on the left to view the offenses from different perspectives. For example, select **By Source IP** or **By Destination IP** to view information about repeat offenders, IP addresses that generate many attacks, or systems that are continually under attack. You can further refine the offenses in the list by selecting a time period for the offenses that you want to view or by changing the search parameters.

You can also search for offenses that are based on various criteria. For more information about searching offenses, see "[Offense Searches](#)" on page 235.

1. Click the **Offenses** tab.
2. On the navigation menu, select the category of offenses that you want to view.
3. Depending on the category that you selected, you may be able to select the following filtering options:
 - a. From the **View Offenses** list, select an option to filter the list of offenses for a specific time frame.
 - b. In the **Current Search Parameters** pane, click **Clear Filter** links to refine the list of offenses.
4. To view all global offenses that are occurring on the network, click **All Offenses**.
5. To view all offenses that are assigned to you, click **My Offenses**.
6. To view offenses grouped on the high-level category, click **By Category**.
 - a. To view low-level category groups for a particular high-level category, click the arrow icon next to the high-level category name.
 - b. To view a list of offenses for a low-level category, double-click the low-level category.

Count fields, such as **Event/Flow Count** and **Source Count** do not consider the network permissions of the user.
7. To view offenses grouped by source IP address, click **By Source IP**.

The list of offenses displays only source IP addresses with active offenses.

- a. Double-click the **Source IP** group that you want to view.
- b. To view a list of local destination IP addresses for the source IP address, click **Destinations** on the **Source** page toolbar.
- c. To view a list of offenses that are associated with this source IP address, click **Offenses** on the **Source** page toolbar.

8. To view offenses grouped by destination IP address, click **By Destination IP**.
 - a. Double-click the **Source IP** address group that you want to view.
 - b. To view a list of offenses that are associated with the destination IP address, click **Offenses** on the **Destination** page toolbar.
 - c. To view a list of source IP addresses associated with the destination IP address, click **Sources** on the **Destination** page toolbar.
9. To view offenses grouped by network, click **By Network**.
 - a. Double-click the **Network** that you want to view.
 - b. To view a list of source IP addresses associated with this network, click **Sources** on the **Network** page toolbar.
 - c. To view a list of destination IP addresses associated with this network, click **Destinations** on the **Network** page toolbar.
 - d. To view a list of offenses that are associated with this network, click **Offenses** on the **Network** page toolbar.
10. Double-click the offense to see additional information.

Use the information in the offense summary and details to investigate the offense and take necessary actions.

Investigating an Offense by Using the Summary Information

The **Offense Summary** window provides the information that you need to investigate an offense in JSA. The information that is most important to you during your investigation might be different, depending on the type of offense that you are investigating.

To make it easier for you to investigate an offense, the bottom of the **Offense Summary** page groups information about top contributors to the offense. These fields show only the most recent or most important pieces of information in that category. Many fields show more information when you hover the mouse over them. Some fields have right-click menu options.

1. Click the **Offenses** tab and double-click the offense that you want to investigate.

The **Offense Summary** window opens.

2. Review the first row of data to learn about the level of importance that JSA assigned to the offense.

Learn more about the magnitude rating:

Table 8: Magnitude Rating Details

Parameter	Description
Magnitude	Indicates the relative importance of the offense. This value is calculated based on the relevance, severity, and credibility ratings.
Status	<p>Hover your mouse over the status icon to see the status.</p> <p>JSA does not display a status icon when an offense is active.</p>
Relevance	<p>Indicates the importance of the destination.</p> <p>JSA determines the relevance by the weight that the administrator assigned to the networks and assets.</p>
Severity	Indicates the threat that an attack poses in relation to how prepared the destination is for the attack.
Credibility	Indicates the integrity of the offense as determined by the credibility rating that is configured in the log source. Credibility increases as multiple sources report the same event. JSA administrators configure the credibility rating of log sources.

- Review the information in the top portion of the **Offense Summary** window to learn more about the type of attack and the timeframe when it occurred.

Learn more about the offense information:

Table 9: Offense Information

Parameter	Description
Description	<p>Shows the cause of the offense.</p> <p>Chained offenses show Preceded by, indicating that the offense changed over time as new events and flows were added to offense.</p>
Offense Type	<p>The offense type is determined by the rule that created the offense. The offense type determines what type of information is displayed in the Offense Source Summary pane.</p>
Event/Flow count	<p>To see the list of events and flows that contributed to the offense, click the Event or Flow links.</p>
Source IP(s)	<p>Specifies the device that attempts to breach the security of a component on your network. The device can have an IPv4 or IPv6 address.</p> <p>Offenses of type Source IP always originate from only one source IP address. Offenses of other types can have more than one source IP address. You can see more information about the source IP address by hovering the mouse over the address, or by using right-click and left-click mouse actions.</p>
Destination IP(s)	<p>Specifies the network device that the source IP address attempted to access. The network device can have an IPv4 or IPv6 address.</p> <p>If the offense has only one target, the IP address is displayed. If the offense has multiple targets, the number of local or remote IP addresses that were targeted. You can see more information by hovering the mouse over the address, or by using right-click and left-click mouse actions.</p>
Start	<p>Specifies the date and time when the first event or flow occurred for the offense.</p>

Table 9: Offense Information (*Continued*)

Parameter	Description
Duration	Specifies the amount of time that elapsed since the first event or flow associated with the offense was created.
Network(s)	Specifies the local networks of the local destination IP addresses that were targeted. JSA considers all networks that are specified in the network hierarchy as local. The system does not associate remote networks to an offense, even if they are specified as a remote network or a remote service on the Admin tab.

4. In the **Offense Source Summary** window, review the information about the source of the offense.

The information that is shown in the **Offense Source Summary** window depends on the **Offense Type** field.

Learn more about the source summary information:

Table 10: Source Summary Information:

Parameter	Description
Chained	<p>Specifies whether the destination IP address is chained.</p> <p>A chained IP address is associated with other offenses. For example, a destination IP address might become the source IP address for another offense. If the destination IP address is chained, click Yes to view the chained offenses.</p>

Table 10: Source Summary Information: *(Continued)*

Parameter	Description
Destination IP(s)	<p>Specifies the network device that the source IP address attempted to access. The network device can have an IPv4 or IPv6 address.</p> <p>If the offense has only one target, the IP address is displayed. If the offense has multiple targets, this field shows the number of local or remote IP addresses that were targeted. You can see more information by hovering the mouse over the address, or by using right-click and left-click mouse actions.</p>
Location	<p>Specifies the network location of the source or destination IP address. If the location is local, click the link to view the networks.</p>
Magnitude	<p>Specifies the relative importance of the source or destination IP address.</p> <p>The magnitude bar provides a visual representation of the CVSS risk value of the asset that is associated with the IP address. Hover your mouse over the magnitude bar to display the calculated magnitude.</p>
Severity	<p>Specifies the severity of the event or offense.</p> <p>Severity specifies the level of threat that an offense poses in relation to how prepared the destination IP address is for the attack. This value is directly mapped to the event category that correlates to the offense. For example, a Denial of Service (DoS) attack has a severity of 10, which specifies a severe occurrence.</p>

Table 10: Source Summary Information: *(Continued)*

Parameter	Description
Source IP(s)	<p>Specifies the device that attempted to breach the security of a component on your network. The device can have an IPv4 or IPv6 address.</p> <p>Offenses of type Source IP always originate from only one source IP address. Offenses of other types can have more than one source IP address. You can see more information about the source IP address by hovering the mouse over the address, or by using right-click and left-click mouse actions.</p>
Username	<p>Specifies the user name that is associated with the event or flow that created the offense.</p> <p>Hover your mouse over the user name to see the most recent information in the asset model database for the user.</p> <p>Events that do not include a user name in the payload, or system-generated events that belong to a local computer or a system account, show Unknown.</p> <p>To access more information that is associated with a selected user name, right-click the user name for View Assets and View Events menu options.</p>
Vulnerabilities	<p>Specifies the number of identified vulnerabilities that are associated with the source or destination IP address. This value also includes the number of active and passive vulnerabilities.</p>

When you view the summary information for historical offenses, the **Last Known** data fields are not populated.

- In the bottom portion of the **Offense Summary** window, review additional information about the offense top contributors, including notes and annotations that are collected about the offense.

To see all the information that JSA collected in a category, click the links on the right side of the category heading.

Learn more about the information presented in the offense details:

Table 11: Offense Details

Offense details category	Description
Last 5 Notes	Use notes to track important information that is gathered during the offense investigation. You can add a note to an offense, but you cannot edit or delete notes.
Top 5 Source IPs	<p>Shows the top 5 IP addresses with the highest magnitude, which is where the suspected attack or policy breach originated.</p> <p>Offenses that have only one source IP address show only one entry in the table.</p>
Top 5 Destination IPs	<p>Shows the top 5 local IP addresses with the highest magnitude, which might indicate the target of the attack. Offenses that target less than 5 local IP addresses show fewer entries in the table.</p> <p>The Chained column indicates whether the destination IP address is the source IP address of another offense. A Yes in this column indicates that an attacker has control over the system with this IP address and is using it to attack other systems.</p> <p>The Magnitude column shows the aggregate Common Vulnerability Scoring System (CVSS) score when it exists. When no CVSS score is available, the column shows the highest magnitude of all the offenses that the IP address is a part of.</p> <p>When you hover the mouse over the destination IP address, the Destination Magnitude shows the CVSS score. When no CVSS score is available, a zero is displayed.</p>

Table 11: Offense Details (Continued)

Offense details category	Description
Top 5 Log Sources	<p>Shows the log sources that contribute the most events to the offense.</p> <p>The Custom Rule Engine (CRE) creates an event and adds it to the offense when the test criteria that is specified in the custom rule matches the incoming event. A log source that displays Custom Rule Engine in the Description field indicates that JSA created the events from that log source.</p> <p>Total Events shows the sum of all the events received from this log source while the offense was active.</p>
Top 5 Users	<p>Events must include user information in order for JSA to populate this table.</p>
Top 5 Categories	<p>Shows the low-level categories that have the most events that contributed to the offense.</p> <p>Local Destination Count shows the number of local destination IP addresses affected by offenses with events in the category. When all destination IP addresses are remote, this field shows 0.</p>
Last 10 Events	<p>Shows information about the last 10 events that contributed to the offense.</p>
Last 10 Flows	<p>Shows information about the last 10 flows that contributed to the offense.</p> <p>The Total Bytes column shows the sum of the bytes transferred in both directions.</p>

Table 11: Offense Details (Continued)

Offense details category	Description
Annotations	<p>Annotations provide insight into why JSA considers the event or observed traffic to be threatening.</p> <p>JSA can add annotations when it adds events or flows to an offense. The oldest annotation shows information that JSA added when the offense was created. Users cannot add, edit, or delete annotations.</p>
Last 5 Search Results	Shows information about the results from the last five scheduled searches.

6. If you installed JSA Risk Manager, click **View Attack Path** to see which assets in your network are communicating to allow an offense to travel through the network.

Investigating Events

An event is a record from a log source, such as a firewall or router device, that describes an action on a network or host. Events that are associated with an offense provide evidence that suspicious activity is happening on your network. By examining the event data, you can understand what caused the offense and determine how best to isolate and mitigate the threat.

Some events are created based on an incoming raw event, while others are created by the JSA Custom Rule Engine (CRE). Events that are created by JSA do not have a payload because they are not based on raw events.

1. In the **Offense Summary** window, click **Events**.

The **List of Events** window shows all events that are associated with the offense.

2. Specify the **Start Time**, **End Time**, and **View** options to view events that occurred within a specific time frame.
3. Click the event column header to sort the event list.
4. In the list of events, right-click the event name to apply quick filter options to reduce the number of events to review.

You can apply quick filters to other columns in the event list as well.

5. Double-click an event to view the event details.

The **Event Information** and the **Source and Destination Information** window show only the information that is known about the event. Depending on the type of event, some fields might be empty.

Learn more about the time fields on the Event Information:

Table 12: Time Fields on the Event Information:

Field	Description
Start Time	The time that JSA received the raw event from the log source.
Storage Time	The time that JSA stored the normalized event.
Log Source Time	The time that is recorded in the raw event from the log source.

6. In the **Payload Information** box, review the raw event for information that JSA did not normalize.

Information that is not normalized does not appear in the JSA interface, but it may be valuable to your investigation.

For more information about how to use JSA to review event data, see "[Log Activity Monitoring](#)" on page 93 and [Searches](#).

Investigating Flows

JSA correlates flows into an offense when it identifies suspicious activity in network communications. The flow analysis provides visibility into layer 7, or the application layer, for applications such as web browsers, NFS, SNMP, Telnet, and FTP. A flow can include information such as IP addresses, ports, applications, traffic statistics, and packet payload from unencrypted traffic.

By default, JSA tries to extract normalized fields and custom flow properties from the first 64 bytes of flow data, but administrators can increase the content capture length to collect more data. For more information, see the *Juniper Secure Analytics Administration Guide*.

1. In the **Offense Summary** window, click **Flows** in the upper right menu.

The **Flow List** window shows all flows that are associated with the offense.

2. Specify the **Start Time**, **End Time**, and **View** options to view flows that occurred within a specific time frame.
3. Click the flow column header to sort the flow list.
4. In the list of flows, right-click the flow name to apply quick filter options to reduce the number of flows to review.

You can apply quick filters to other columns in the flow list as well.

5. Double-click a flow to review the flow details.

Learn more about the Flow Details:

Table 13: Flow Details

Field	Description
Event Description	When the application is not identified in the payload, JSA uses built-in decoding to determine the application, and shows Application detected with state-based decoding in Event Description .
Source Payload and Destination Payload	Shows the size of the payload. When the size exceeds 64 bytes, the payload might contain additional information that is not shown in the JSA interface.
Custom Rules Partially Matched	Shows rules for which the threshold value was not met, but otherwise the rule matched.
Flow Direction	Specifies the flow direction, where L indicates local network, and R indicates remote network.

For more information about how to use JSA to review flow data, see [Network Activity Investigation](#) and ["Event and Flow Searches"](#) on page 212.

RELATED DOCUMENTATION

[Offense Actions | 79](#)

[Offense Indexing | 60](#)

[Offense Retention | 62](#)

Offense Actions

IN THIS SECTION

- [Adding Notes | 79](#)
- [Hiding Offenses | 80](#)
- [Showing Hidden Offenses | 80](#)
- [Closing Offenses | 81](#)
- [Exporting Offenses | 82](#)
- [Assigning Offenses to Users | 83](#)
- [Sending Email Notifications | 83](#)
- [Marking an Offense for Follow-up | 85](#)

JSA provides the capability to act on the offenses as you investigate them. To help you track offenses that were acted upon, JSA adds an icon to the **Flag** column when you assign an offense to a user, protect or hide an offense, add notes, or mark the offense for follow-up.

To perform the same action on multiple offenses, hold the Control key while you select each offense you want to act on. To view offense details on a new page, press the Ctrl key while you double-click an offense.

Adding Notes

Add notes to an offense to track information that is collected during an investigation. Notes can include up to 2000 characters.

1. Click the **Offenses** tab.

2. Select the offense to which you want to add the note.

To add the same note to multiple offenses, press the Ctrl key while you select each offense.

3. From the **Actions** list, select **Add Note**.
4. Type the note that you want to include for this offense.
5. Click **Add Note**.

The note is displayed in the **Last 5 Notes** pane on the **Offense Summary** window. A **Notes** icon is displayed in the flag column of the offense list.

Hover your mouse over the notes indicator in the **Flag** column of the **Offenses** list to view the note.

Hiding Offenses

Hide an offense to prevent it from being displayed in the offense list. After you hide an offense, the offense is no longer displayed in any list on the **Offenses** tab, including the **All Offenses** list. However, if you perform a search that includes hidden offenses, the offense is displayed in the search results.

1. Click the **Offenses** tab.
2. Select the offense that you want to hide.

To hide multiple offenses, hold the Control key while you select each offense.

3. From the **Actions** list box, select **Hide**.
4. Click **OK**.

Showing Hidden Offenses

By default, the offense list on the **Offenses** tab filters to exclude hidden offenses. To view hidden offenses, clear the filter on the **Offenses** tab or perform a search that includes hidden offenses. When you include hidden offenses in the offense list, the offenses show the **Hidden** icon in the **Flag** column.

1. Click the **Offenses** tab.
2. To clear the filter on the offense list, click **Clear Filter** next to the **Exclude Hidden Offenses** search parameter.
3. To create a new search that includes hidden offenses, follow these steps:

- a. From the **Search** list box, select **New Search**.
 - b. In the **Search Parameters** window, clear the **Hidden Offenses** check box in the **Exclude** options list.
 - c. Click **Search**.
4. To remove the hidden flag from an offense, follow these steps:
- a. Select the offense for which you want to remove the hidden flag.
To select multiple offenses, hold the Control key while you click each offense.
 - b. From the **Actions** list box, select **Show**.

The hidden flag is removed and the offense appears in the offense list without having to clear the **Exclude Hidden Offenses** filter.

Closing Offenses

Close an offense to remove it completely from your system.

The default offense retention period is 30 days. After the offense retention period expires, closed offenses are deleted from the system. You can protect an offense to prevent it from being deleted when the retention period expires.

Closed offenses are no longer displayed in any list on the **Offenses** tab, including the **All Offenses** list. If you include closed offenses in a search, and the offense is still within the retention period, the offense is displayed in the search results. If more events occur for an offense that is closed, a new offense is created.

When you close offenses, you must select a reason for closing the offense. If you have the **Manage Offense Closing** permission, you can add custom closing reasons. For more information about user role permissions, see the *Juniper Secure Analytics Administration Guide*.

1. Click the **Offenses** tab.
2. Select the offense that you want to close.
To close multiple offenses, hold the Control key while you select each offense.
3. From the **Actions** list, select **Close**.
4. In the **Reason for Closing** list, specify a closing reason.

To add a close reason, click the icon beside **Reason for Closing** to open the **Custom Offense Close Reasons** dialog box.

5. In the **Notes** field, type a note to provide more information.

The **Notes** field displays the note that was entered for the previous offense closing. Notes must not exceed 2,000 characters.

6. Click **OK**.

After you close offenses, the counts that are displayed on the **By Category** window of the **Offenses** tab can take several minutes to reflect the closed offenses.

Exporting Offenses

Export offenses when you want to reuse the data or when you want to store the data externally. For example, you can use the offense data to create reports in a third-party application. You can also export offenses as a secondary long-term retention strategy. Customer Support might require you to export offenses for troubleshooting purposes.

You can export offenses in Extensible Markup Language (XML) or comma-separated values (CSV) format. The resulting XML or CSV file includes the parameters that are specified in the **Column Definition** pane of the search parameters. The length of time that is required to export the data depends on the number of parameters specified.

1. Click the **Offenses** tab.
2. Select the offenses that you want to export.

To select multiple offenses, hold the Control key while you select each offense.

3. Choose one of the following options:

- To export the offenses in XML format, select **Actions >Export to XML**.
- To export the offenses in CSV format, select **Actions >Export to CSV**.

NOTE: If you use Microsoft Excel to import the CSV file, you must select the correct locale to ensure that the data displays correctly.

4. Choose one of the following options:

- To open the file for immediate viewing, select **Open with** and select an application from the list.

- To save the file, select **Save File**.

5. Click **OK**.

The file, `<date>-data_export.xml.zip`, is saved in the default download folder on your computer.

Assigning Offenses to Users

By default, all new offenses are unassigned. You can assign an offense to an JSA user for investigation.

When you assign an offense to a user, the offense is displayed on the **My Offenses** page for that user. You must have the **Assign Offenses to Users** permission to assign offenses to users. For more information about user role permissions, see the *Juniper Secure Analytics Administration Guide*.

You can assign offenses to users from either the **Offenses** tab or **Offense Summary** pages. This procedure provides instruction on how to assign offenses from the **Offenses** tab.

1. Click the **Offenses** tab.
2. Select the offense that you want to assign.

To assign multiple offenses, hold the Control key while you select each offense.
3. From the **Actions** list, select **Assign**.
4. In the **Assign To User** list, select the user that you want to assign this offense to.

NOTE: The **Assign To User** list displays only those users who have privileges to view the **Offenses** tab. The security profile settings for the user are followed as well.

5. Click **Save**.

The offense is assigned to the selected user. The **User** icon is displayed in the Flag column of the **Offenses** tab to indicate that the offense is assigned. The designated user can see this offense on the **My Offenses** page.

Sending Email Notifications

Share the offense summary information with another person by sending an email.

The body of the email message includes the following information, if available:

- Source IP address
 - Source user name, host name, or asset name
 - Total number of sources
 - Top five sources by magnitude
 - Source networks
 - Destination IP address
 - Destination user name, host name, or asset name
 - Total number of destinations
 - Top five destinations by magnitude
 - Destination networks
 - Total number of events
 - Rules that caused the offense or event rule to fire
 - Full description of the offense or event rule
 - Offense ID
 - Top five categories
 - Start time of the offense or the time the event was generated
 - Top five annotations
 - Link to the offense user interface
 - Contributing CRE rules
1. Click the **Offenses** tab.
 2. Select the offense for which you want to send an email notification.
 3. From the **Actions** list box, select **Email**.
 4. Configure the following parameters:

Option	Description
Parameter	Description
To	Type the email address of the user you want to notify when a change occurs to the selected offense. Separate multiple email addresses with a comma.
From	Type the originating email address. The default is root@localhost.com.
Email Subject	Type the subject for the email. The default is Offense ID .
Email Message	Type the standard message that you want to accompany the notification email.

5. Click **Send**.

Marking an Offense for Follow-up

Mark an offense for follow-up when you want to flag it for further investigation.

1. Click the **Offenses** tab.
2. Find the offense that you want to mark for follow-up.
3. Double-click the offense.
4. From the **Actions** list, select **Follow up**.

The offense now displays the follow-up icon in the **Flag** column. To sort the offense list to show flagged offenses at the top, click the **Flags** column header.

RELATED DOCUMENTATION

[Offense Indexing | 60](#)

[Offense Retention | 62](#)



Log Activity Investigation

[Log Activity Investigation | 88](#)

[Log Activity Tab Overview | 88](#)

[Log Activity Monitoring | 93](#)

[Viewing Associated Offenses | 111](#)

[Modifying Event Mapping | 111](#)

[Tuning False Positive Events From Creating Offenses | 112](#)

[PCAP Data | 113](#)

[Exporting Events | 116](#)

Log Activity Investigation

You can monitor and investigate events in real time or perform advanced searches.

Using the **Log Activity** tab, you can monitor and investigate log activity (events) in real time or perform advanced searches.

TIP: You can also investigate log activity using the QRadar Analyst Workflow interface. For more information, see "[QRadar Analyst Workflow](#)" on page 42.

RELATED DOCUMENTATION

[Log Activity Tab Overview](#) | 88

[Log Activity Monitoring](#) | 93

[Viewing Associated Offenses](#) | 111

Log Activity Tab Overview

IN THIS SECTION

- [Log Activity Tab Toolbar](#) | 89
- [Right-click Menu Options](#) | 92
- [Status Bar](#) | 92

An event is a record from a log source, such as a firewall or router device, that describes an action on a network or host.

The **Log Activity** tab specifies which events are associated with offenses.

You must have permission to view the **Log Activity** tab.

Log Activity Tab Toolbar

You can access several options from the Log Activity toolbar

Using the toolbar, you can access the following options:

Table 14: Log Activity Toolbar Options

Option	Description
Search	<p>Click Search to perform advanced searches on events. Options include:</p> <ul style="list-style-type: none"> • New Search Select this option to create a new event search. • Edit Search Select this option to select and edit an event search. • Manage Search Results Select this option to view and manage search results.
Quick Searches	<p>From this list box, you can run previously saved searches. Options are displayed in the Quick Searches list box only when you have saved search criteria that specifies the Include in my Quick Searches option.</p>
Add Filter	<p>Click Add Filter to add a filter to the current search results.</p>
Save Criteria	<p>Click Save Criteria to save the current search criteria.</p>
Save Results	<p>Click Save Results to save the current search results. This option is only displayed after a search is complete. This option is disabled in streaming mode.</p>
Cancel	<p>Click Cancel to cancel a search in progress. This option is disabled in streaming mode.</p>
False Positive	<p>Click False Positive to open the False Positive Tuning window, which will allow you to tune out events that are known to be false positives from creating offenses.</p> <p>This option is disabled in streaming mode. For more information about tuning false positives, see "Tuning False Positives" on page 112.</p>

Table 14: Log Activity Toolbar Options (Continued)

Option	Description
Rules	<p>The Rules option is only visible if you have permission to view rules.</p> <p>Click Rules to configure custom event rules. Options include:</p> <ul style="list-style-type: none"> • Rules Select this option to view or create a rule. If you only have the permission to view rules, the summary page of the Rules wizard is displayed. If you have the permission to maintain custom rules, the Rules wizard is displayed and you can edit the rule. • Add Threshold Rule Select this option to create a threshold rule. A threshold rule tests event traffic for activity that exceeds a configured threshold. Thresholds can be based on any data that is collected JSA. For example, if you create a threshold rule indicating that no more than 220 clients can log in to the server between 8 am and 5 pm, the rules generate an alert when the 221st client attempts to log in. <p>When you select the Add Threshold Rule option, the Rules wizard is displayed, prepopulated with the appropriate options for creating a threshold rule.</p>
Rules (continued)	<ul style="list-style-type: none"> • Add Behavioral Rule Select this option to create a behavioral rule. A behavioral rule tests event traffic for abnormal activity, such as the existence of new or unknown traffic, which is traffic that suddenly ceases or a percentage change in the amount of time an object is active. For example, you can create a behavioral rule to compare the average volume of traffic for the last 5 minutes with the average volume of traffic over the last hour. If there is more than a 40% change, the rule generates a response. <p>When you select the Add Behavioral Rule option, the Rules wizard is displayed, prepopulated with the appropriate options for creating a behavioral rule.</p>

Table 14: Log Activity Toolbar Options (Continued)

Option	Description
Actions	<p>Click Actions to perform the following actions:</p> <ul style="list-style-type: none"> • Show All Select this option to remove all filters on search criteria and display all unfiltered events. • Print Select this option to print the events that are displayed on the page. • Export to XML > Visible Columns Select this option to export only the columns that are visible on the Log Activity tab. This is the recommended option. See Exporting events. • Export to XML > Full Export (All Columns) - Select this option to export all event parameters. A full export can take an extended period of time to complete. See "Exporting Events" on page 116. • Export to CSV >Visible Columns Select this option to export only the columns that are visible on the Log Activity tab. This is the recommended option. See "Exporting Events" on page 116. • Export to CSV >Full Export (All Columns) - Select this option to export all event parameters. A full export can take an extended period of time to complete. See "Exporting Events" on page 116. • Delete Select this option to delete a search result. See "Managing Search Results" on page 253 • Notify Select this option to specify that you want a notification emailed to you on completion of the selected searches. This option is only enabled for searches in progress. <p>NOTE: The Print, Export to XML, and Export to CSV options are disabled in streaming mode and when viewing partial search results.</p>
Search toolbar	<ul style="list-style-type: none"> • Advanced Search--Select Advanced Search from the list box to enter an Ariel Query Language (AQL) search string to specify the fields that you want returned. • Quick Filter--Select Quick Filter from the list box to search payloads by using simple words or phrases.
View	<p>The default view on the Log Activity tab is a stream of real-time events. The View list contains options to also view events from specified time periods. After you choose a specified time period from the View list, you can then modify the displayed time period by changing the date and time values in the Start Time and End Time fields.</p>

Right-click Menu Options

On the **Log Activity** tab, you can right-click an event to access more event filter information.

The right-click menu options are:

Table 15: Right-click Menu Options

Option	Description
Filter on	Select this option to filter on the selected event, depending on the selected parameter in the event.
False Positive	Select this option to open the False Positive window, which will allow you to tune out events that are known to be false positives from creating offenses. This option is disabled in streaming mode. See "Tuning False Positives" on page 112.
More options:	Select this option to investigate an IP address or a user name. For more information about investigating an IP address, see Investigating IP addresses . NOTE: This option is not displayed in streaming mode.
Quick Filter	Filter items that match, or do not match the selection.

Status Bar

When streaming events, the status bar displays the average number of results that are received per second.

This is the number of results the Console successfully received from the Event processors. If this number is greater than 40 results per second, only 40 results are displayed. The remainder is accumulated in the result buffer. To view more status information, move your mouse pointer over the status bar.

When events are not being streamed, the status bar displays the number of search results that are currently displayed on the tab and the amount of time that is required to process the search results.

RELATED DOCUMENTATION

[Log Activity Monitoring | 93](#)

[Viewing Associated Offenses | 111](#)

[Modifying Event Mapping | 111](#)

Log Activity Monitoring

IN THIS SECTION

- [Viewing Streaming Events | 93](#)
- [Viewing Normalized Events | 94](#)
- [Viewing Raw Events | 97](#)
- [Viewing Grouped Events | 99](#)
- [Viewing Event Details | 104](#)
- [Event Details Toolbar | 109](#)

By default, the **Log Activity** tab displays events in streaming mode, allowing you to view events in real time.

For more information about streaming mode, see "[Viewing Streaming Events](#)" on page 93. You can specify a different time range to filter events by using the **View** list box.

If you previously configured saved search criteria as the default, the results of that search are automatically displayed when you access the **Log Activity** tab. For more information about saving search criteria, see "[Saving Search Criteria](#)" on page 220.

Viewing Streaming Events

Streaming mode will enable you to view event data that enters your system. This mode provides you with a real-time view of your current event activity by displaying the last 50 events.

If you apply any filters on the **Log Activity** tab or in your search criteria before enabling streaming mode, the filters are maintained in streaming mode. However, streaming mode does not support searches that

include grouped events. If you enable streaming mode on grouped events or grouped search criteria, the **Log Activity** tab displays the normalized events. See "[Viewing normalized events](#)" on page 94

When you want to select an event to view details or perform an action, you must pause streaming before you double-click an event. When the streaming is paused, the last 1,000 events are displayed.

1. Click the **Log Activity** tab.
2. From the **View** list box, select **Real Time (streaming)**.

For information about the toolbar options, see Table 4-1. For more information about the parameters that are displayed in streaming mode, see Table 4-7.

3. Optional. Pause or play the streaming events. Choose one of the following options:
 - To select an event record, click the **Pause** icon to pause streaming.
 - To restart streaming mode, click the **Play** icon.

Viewing Normalized Events

Events are collected in raw format, and then normalized for display on the **Log Activity** tab.

Normalization involves parsing raw event data and preparing the data to display readable information about the tab. When events are normalized, the system normalizes the names as well. Therefore, the name that is displayed on the **Log Activity** tab might not match the name that is displayed in the event.

NOTE: If you selected a time frame to display, a time series chart is displayed. For more information about using time series charts, see "[Time Series Chart Overview](#)" on page 206.

By default, the **Log Activity** tab displays the following parameters when you view normalized events:

Table 16: Log Activity Tab - Default (Normalized) Parameters

Parameter	Description
Current Filters	<p>The top of the table displays the details of the filters that are applied to the search results. To clear these filter values, click Clear Filter.</p> <p>NOTE: This parameter is only displayed after you apply a filter.</p>

Table 16: Log Activity Tab - Default (Normalized) Parameters (Continued)

Parameter	Description
View	From this list box, you can select the time range that you want to filter for.
Current Statistics	<p>When not in Real Time (streaming) or Last Minute (auto refresh) mode, current statistics are displayed, including:</p> <p>NOTE: Click the arrow next to Current Statistics to display or hide the statistics</p> <ul style="list-style-type: none"> • Total Results Specifies the total number of results that matched your search criteria. • Data Files Searched Specifies the total number of data files searched during the specified time span. • Compressed Data Files Searched Specifies the total number of compressed data files searched within the specified time span. • Index File Count Specifies the total number of index files searched during the specified time span. • Duration Specifies the duration of the search. <p>NOTE: Current statistics are useful for troubleshooting. When you contact Juniper Customer Support to troubleshoot events, you might be asked to supply current statistical information.</p>
Charts	<p>Displays configurable charts that represent the records that are matched by the time interval and grouping option. Click Hide Charts if you want to remove the charts from your display. The charts are only displayed after you select a time frame of Last Interval (auto refresh) or above, and a grouping option to display. For more information about configuring charts, see "Chart Management" on page 205.</p> <p>NOTE: If you use Mozilla Firefox as your browser and an ad blocker browser extension is installed, charts do not display. To display charts, you must remove the ad blocker browser extension. For more information, see your browser documentation.</p>
Offenses icon	<p>Click this icon to view details of the offense that is associated with this event. For more information, see "Chart Management" on page 205.</p> <p>NOTE: Depending on your product, this icon is might not be available. You must have JSA.</p>
Start Time	Specifies the time of the first event, as reported to JSA by the log source.

Table 16: Log Activity Tab - Default (Normalized) Parameters (Continued)

Parameter	Description
Event Name	Specifies the normalized name of the event.
Log Source	Specifies the log source that originated the event. If there are multiple log sources that are associated with this event, this field specifies the term Multiple and the number of log sources.
Event Count	Specifies the total number of events that are bundled in this normalized event. Events are bundled when many of the same type of event for the same source and destination IP address are detected within a short time.
Time	Specifies the date and time when JSA received the event.
Low Level Category	Specifies the low-level category that is associated with this event. For more information about event categories, see the <i>Juniper Secure Analytics Administration Guide</i> .
Source IP	Specifies the source IP address of the event. NOTE: If you select the Normalized (With IPv6 Columns) display, refer to the Source IPv6 parameter for IPv6 events.
Source Port	Specifies the source port of the event.
Destination IP	Specifies the destination IP address of the event. NOTE: If you select the Normalized (With IPv6 Columns) display, refer to the Destination IPv6 parameter for IPv6 events.
Destination Port	Specifies the destination port of the event.
Username	Specifies the user name that is associated with this event. User names are often available in authentication-related events. For all other types of events where the user name is not available, this field specifies N/A.

Table 16: Log Activity Tab - Default (Normalized) Parameters (Continued)

Parameter	Description
Magnitude	Specifies the magnitude of this event. Variables include credibility, relevance, and severity. Point your mouse over the magnitude bar to display values and the calculated magnitude.

If you select the **Normalized (With IPv6 Columns)** display, then the **Log Activity** tab displays the following extra parameters:

Table 17: Table 14. Log Activity tab - Normalized (With IPv6 Columns) Parameters

Parameter	Description
Source IPv6	Specifies the source IP address of the event. NOTE: IPv4 events display 0.0.0.0.0.0.0 in the Source IPv6 and Destination IPv6 columns.
Destination IPv6	Specifies the destination IP address of the event. NOTE: IPv4 events display 0.0.0.0.0.0.0 in the Source IPv6 and Destination IPv6 columns.

1. Click the **Log Activity** tab.
2. Optional: From the **Display** list box, select **Normalized (With IPv6 Columns)**. The **Normalized (With IPv6 Columns)** display shows source and destination IPv6 addresses for IPv6 events.
3. From the **View** list box, select the time frame that you want to display.
4. Click the **Pause** icon to pause streaming.
5. Double-click the event that you want to view in greater detail. For more information, see "[Event details](#)" on page 104.

Viewing Raw Events

You can view raw event data, which is the unparsed event data from the log source.

When you view raw event data, the **Log Activity** tab provides the following parameters for each event.

Table 18: Raw Event Parameters

Parameter	Description
Current Filters	<p>The top of the table displays the details of the filters that are applied to the search results. To clear these filter values, click Clear Filter.</p> <p>NOTE: This parameter is only displayed after you apply a filter.</p>
View	<p>From this list box, you can select the time range that you want to filter for.</p>
Current Statistics	<p>When not in Real Time (streaming) or Last Minute (auto refresh) mode, current statistics are displayed, including:</p> <p>NOTE: Click the arrow next to Current Statistics to display or hide the statistics</p> <ul style="list-style-type: none"> • Total Results Specifies the total number of results that matched your search criteria. • Data Files Searched Specifies the total number of data files searched during the specified time span. • Compressed Data Files Searched Specifies the total number of compressed data files searched within the specified time span. • Index File Count Specifies the total number of index files searched during the specified time span. • Duration Specifies the duration of the search. <p>NOTE: Current statistics are useful for troubleshooting. When you contact Juniper Customer Support to troubleshoot events, you might be asked to supply current statistical information.</p>
Charts	<p>Displays configurable charts that represent the records that are matched by the time interval and grouping option. Click Hide Charts if you want to remove the charts from your display. The charts are only displayed after you select a time frame of Last Interval (auto refresh) or above, and a grouping option to display.</p> <p>NOTE: If you use Mozilla Firefox as your browser and an ad blocker browser extension is installed, charts do not display. To display charts, you must remove the ad blocker browser extension. For more information, see your browser documentation.</p>
Offenses icon	<p>Click this icon to view details of the offense that is associated with this event.</p>

Table 18: Raw Event Parameters (Continued)

Parameter	Description
Start Time	Specifies the time of the first event, as reported to JSA by the log source.
Log Source	Specifies the log source that originated the event. If there are multiple log sources that are associated with this event, this field specifies the term Multiple and the number of log sources.
Payload	Specifies the original event payload information in UTF-8 format.

1. Click the **Log Activity** tab.
2. From the **Display** list box, select **Raw Events**.
3. From the **View** list box, select the time frame that you want to display.
4. Double-click the event that you want to view in greater detail. See ["Event details" on page 104](#).

Viewing Grouped Events

Using the **Log Activity** tab, you can view events that are grouped by various options. From the **Display** list box, you can select the parameter by which you want to group events.

The Display list box is not displayed in streaming mode because streaming mode does not support grouped events. If you entered streaming mode by using non-grouped search criteria, this option is displayed.

The Display list box provides the following options:

Table 19: Grouped Events Options

Group option	Description
Low Level Category	<p>Displays a summarized list of events that are grouped by the low-level category of the event.</p> <p>For more information about categories, see the <i>Juniper Secure Analytics Administration Guide</i>.</p>

Table 19: Grouped Events Options (Continued)

Group option	Description
Event Name	Displays a summarized list of events that are grouped by the normalized name of the event.
Destination IP	Displays a summarized list of events that are grouped by the destination IP address of the event.
Destination Port	Displays a summarized list of events that are grouped by the destination port address of the event.
Source IP	Displays a summarized list of events that are grouped by the source IP address of the event.
Custom Rule	Displays a summarized list of events that are grouped by the associated custom rule.
Username	Displays a summarized list of events that are grouped by the user name that is associated with the events.
Log Source	Displays a summarized list of events that are grouped by the log sources that sent the event to JSA.
High Level Category	Displays a summarized list of events that are grouped by the high-level category of the event.
Network	Displays a summarized list of events that are grouped by the network that is associated with the event.
Source Port	Displays a summarized list of events that are grouped by the source port address of the event.

After you select an option from the **Display** list box, the column layout of the data depends on the chosen group option. Each row in the events table represents an event group. The **Log Activity** tab provides the following information for each event group.

Table 20: Grouped Event Parameters

Parameter	Description
Grouping By	Specifies the parameter that the search is grouped on.
Current Filters	The top of the table displays the details of the filter that is applied to the search results. To clear these filter values, click Clear Filter .
View	From the list box, select the time range that you want to filter for.
Current Statistics	<p>When not in Real Time (streaming) or Last Minute (auto refresh) mode, current statistics are displayed, including:</p> <p>NOTE: Click the arrow next to Current Statistics to display or hide the statistics.</p> <ul style="list-style-type: none"> • Total Results Specifies the total number of results that matched your search criteria. • Data Files Searched Specifies the total number of data files searched during the specified time span. • Compressed Data Files Searched Specifies the total number of compressed data files searched within the specified time span. • Index File Count Specifies the total number of index files searched during the specified time span. • Duration Specifies the duration of the search. <p>NOTE: Current statistics are useful for troubleshooting. When you contact Juniper Customer Support to troubleshoot events, you might be asked to supply current statistic information.</p>

Table 20: Grouped Event Parameters (Continued)

Parameter	Description
Charts	<p>Displays configurable charts that represent the records that are matched by the time interval and grouping option. Click Hide Charts if you want to remove the chart from your display.</p> <p>Each chart provides a legend, which is a visual reference to help you associate the chart objects to the parameters they represent. Using the legend feature, you can perform the following actions:</p> <ul style="list-style-type: none"> • Move your mouse pointer over a legend item to view more information about the parameters it represents. • Right-click the legend item to further investigate the item. • Click a legend item to hide the item in the chart. Click the legend item again to show the hidden item. You can also click the corresponding graph item to hide and show the item. • Click Legend if you want to remove the legend from your chart display. <p>NOTE: The charts are only displayed after you select a time frame of Last Interval (auto refresh) or above, and a grouping option to display.</p> <p>NOTE: If you use Mozilla Firefox as your browser and an ad blocker browser extension is installed, charts do not display. To display charts, you must remove the ad blocker browser extension. For more information, see your browser documentation.</p>
Source IP (Unique Count)	Specifies the source IP address that is associated with this event. If there are multiple IP addresses that are associated with this event, this field specifies the term Multiple and the number of IP addresses.
Destination IP (Unique Count)	Specifies the destination IP address that is associated with this event. If there are multiple IP addresses that are associated with this event, this field specifies the term Multiple and the number of IP addresses.
Destination Port (Unique Count)	Specifies the destination ports that are associated with this event. If there are multiple ports that are associated with this event, this field specifies the term Multiple and the number of ports.
Event Name	Specifies the normalized name of the event.

Table 20: Grouped Event Parameters (Continued)

Parameter	Description
Log Source (Unique Count)	Specifies the log sources that sent the event to JSA. If there are multiple log sources that are associated with this event, this field specifies the term Multiple and the number of log sources.
High Level Category (Unique Count)	Specifies the high-level category of this event. If there are multiple categories that are associated with this event, this field specifies the term Multiple and the number of categories. For more information about categories, see the <i>Juniper Secure Analytics Administration Guide</i> .
Low Level Category (Unique Count)	Specifies the low-level category of this event. If there are multiple categories that are associated with this event, this field specifies the term Multiple and the number of categories.
Protocol (Unique Count)	Specifies the protocol ID associated with this event. If there are multiple protocols that are associated with this event, this field specifies the term Multiple and the number of protocol IDs.
Username (Unique Count)	Specifies the user name that is associated with this event, if available. If there are multiple user names that are associated with this event, this field specifies the term Multiple and the number of user names.
Magnitude (Maximum)	Specifies the maximum calculated magnitude for grouped events. Variables that are used to calculate magnitude include credibility, relevance, and severity.
Event Count (Sum)	Specifies the total number of events that are bundled in this normalized event. Events are bundled when many of the same type of event for the same source and destination IP address are seen within a short time.
Count	Specifies the total number of normalized events in this event group.

1. Click the **Log Activity** tab.
2. From the **View** list box, select the time frame that you want to display.
3. From the Display list box, choose which parameter you want to group events on. See Table 2.

The events groups are listed. For more information about the event group details, see Table 1.

4. To view the **List of Events** page for a group, double-click the event group that you want to investigate.

The **List of Events** page does not retain chart configurations that you might have defined on the **Log Activity** tab. For more information about the **List of Events** page parameters, see Table 1.

5. To view the details of an event, double-click the event that you want to investigate. For more information about event details, see Table 2.

Viewing Event Details

You can view a list of events in various modes, including streaming mode or in event groups. In, whichever mode you choose to view events, you can locate and view the details of a single event.

The event details page provides the following information:

Table 21: Event Details

Parameter	Description
Event Name	Specifies the normalized name of the event.
Low Level Category	Specifies the low-level category of this event. For more information about categories, see the <i>Juniper Secure Analytics Administration Guide</i> .
Event Description	Specifies a description of the event, if available.
Magnitude	Specifies the relative importance of a particular offense. Magnitude is a weighted value calculated from relevance, severity, and credibility.
Relevance	Specifies the relative impact of an event, category, or offense on the network.
Severity	Specifies the relative threat that a source poses on a destination.

Table 21: Event Details (*Continued*)

Parameter	Description
Credibility	Specifies the integrity of an event or an offense. Credibility increases as multiple sources report the same event or offense.
Username	Specifies the user name that is associated with this event, if available. To access more information that is associated with a selected user name, right-click the user name for View Assets and View Events menu options.
Start Time	Specifies the time of the event was received from the log source.
Storage Time	Specifies the time that the event was stored in the JSA database.
Log Source Time	Specifies the system time as reported by the log source in the event payload.

Source and Destination information

Source IP	Specifies the source IP address of the event.
Destination IP	Specifies the destination IP address of the event.
Source Asset Name	Specifies the user-defined asset name of the event source. For more information about assets, see Asset management.
Destination Asset Name	Specifies the user-defined asset name of the event destination. For more information about assets, see Asset management
Source Port	Specifies the source port of this event.
Destination Port	Specifies the destination port of this event.
Pre NAT Source IP	For a firewall or another device capable of Network Address Translation (NAT), this parameter specifies the source IP address before the NAT values were applied. NAT translates an IP address in one network to a different IP address in another network.

Table 21: Event Details (Continued)

Parameter	Description
Pre NAT Destination IP	For a firewall or another device capable of NAT, this parameter specifies the destination IP address before the NAT values were applied.
Pre NAT Source Port	For a firewall or another device capable of NAT, this parameter specifies the source port before the NAT values were applied.
Pre NAT Destination Port	For a firewall or another device capable of NAT, this parameter specifies the destination port before the NAT values were applied.
Post NAT Source IP	For a firewall or another device capable of NAT, this parameter specifies the source IP address after the NAT values were applied.
Post NAT Destination IP	For a firewall or another device capable of NAT, this parameter specifies the destination IP address after the NAT values were applied.
Post NAT Source Port	For a firewall or another device capable of NAT, this parameter specifies the source port after the NAT values were applied.
Post NAT Destination Port	For a firewall or another device capable of NAT, this parameter specifies the destination port after the NAT values were applied.
Post NAT Source Port	For a firewall or another device capable of NAT, this parameter specifies the source port after the NAT values were applied.
Post NAT Destination Port	For a firewall or another device capable of NAT, this parameter specifies the destination port after the NAT values were applied.
IPv6 Source	Specifies the source IPv6 address of the event.
IPv6 Destination	Specifies the destination IPv6 address of the event.
Source MAC	Specifies the source MAC address of the event.

Table 21: Event Details (Continued)

Parameter	Description
Destination MAC	Specifies the destination MAC address of the event.
Payload information	
Payload	<p>Specifies the payload content from the event. This field offers 3 tabs to view the payload:</p> <ul style="list-style-type: none"> • Universal Transformation Format (UTF) - Click UTF. • Hexadecimal - Click HEX. • Base64 - Click Base64.
Additional information	
Protocol	Specifies the protocol that is associated with this event.
QID	Specifies the QID for this event. Each event has a unique QID. For more information about mapping a QID, see "Modifying Event Mapping" on page 111
Log Source	Specifies the log source that sent the event to JSA. If there are multiple log sources that are associated with this event, this field specifies the term Multiple and the number of log sources.
Event Count	Specifies the total number of events that are bundled in this normalized event. Events are bundled when many of the same type of event for the same source and destination IP address are seen within a short time.
Custom Rules	Specifies custom rules that match this event. .
Custom Rules Partially Matched	Specifies custom rules that partially match this event.

Table 21: Event Details (*Continued*)

Parameter	Description
Annotations	Specifies the annotation for this event. Annotations are text descriptions that rules can automatically add to events as part of the rule response.
Event Collector	Specifies the ID of the Event Collector component that parsed the event
QID Event ID	The primary value set by a DSM to identify an event. JSA uses this field together with the Event Category to map to a QID record for the event.
QID Event Category	The secondary value set by a DSM to identify an event. JSA uses this field together with the Event ID to map to a QID record for the event.
Log Source Identifier	Specifies the Log Source Identifier of the log source that received the event. If the event is routed to a SIM Generic Log type log source, set this value as the Log Source Identifier value when you create a log source to collect this event.
Truncated	Specifies whether the event payload was truncated because it exceeded the maximum allowable size of 32 KB for JSA. The parameter is only set to True if the payload was truncated before storage because it exceeded the maximum allowable size for JSA. The parameter is set to False if the payload was not truncated at all. It is also set to False if the payload was truncated by the log source protocol that collected it based on the maximum payload size parameter that was set in the log source configuration.
Stored for Performance	Set to True if an event was routed directly to storage due to performance problems. If the parameter is set to False, and the event has a Low Level Category of Stored, JSA attempted to parse it but the event was unrecognized by all available log sources that have a matching Log Source Identifier. In both cases, the event was stored without any parsing or normalization.
<p>Identity information JSA collects identity information, if available, from log source messages. Identity information provides extra details about assets on your network. Log sources only generate identity information if the log message sent to JSA contains an IP address and least one of the following items: User name or MAC address. Not all log sources generate identity information.</p>	
Identity Username	Specifies the user name of the asset that is associated with this event.

Table 21: Event Details (Continued)

Parameter	Description
Identity IP	Specifies the IP address of the asset that is associated with this event.
Identity Net Bios Name	Specifies the Network Base Input/Output System (Net Bios) name of the asset that is associated with this event.
Identity Extended field	Specifies more information about the asset that is associated with this event. The content of this field is user-defined text and depends on the devices on your network that are available to provide identity information. Examples include: physical location of devices, relevant policies, network switch, and port names.
Has Identity (Flag)	Specifies True if JSA has collected identify information for the asset that is associated with this event. For more information about which devices send identity information, see the <i>Juniper Secure Analytics Configuring DSMs</i> .
Identity Host Name	Specifies the host name of the asset that is associated with this event.
Identity MAC	Specifies the MAC address of the asset that is associated with this event.
Identity Group Name	Specifies the group name of the asset that is associated with this event.

Event Details Toolbar

The events details toolbar provides several functions for viewing events detail.

The **event details** toolbar provides the following functions:

Table 22: Event Details Toolbar

Parameter	Description
Return to Events List	Click Return to Events List to return to the list of events.
Offense	Click Offense to display the offenses that are associated with the event.
Anomaly	Click Anomaly to display the saved search results that caused the anomaly detection rule to generate this event. NOTE: This icon is only displayed if this event was generated by an anomaly detection rule.
Map Event	Click Map Event to edit the event mapping. For more information, see " Modifying Event Mapping " on page 111.
False Positive	Click False Positive to tune JSA to prevent false positive events from generating into offenses.
Extract Property	Click Extract Property to create a custom event property from the selected event.
Previous	Click Previous to view the previous event in the event list.
Next	Click Next to view the next event in the event list.
Print	Click Print to print the event details.

RELATED DOCUMENTATION

[Viewing Associated Offenses | 111](#)

[Modifying Event Mapping | 111](#)

[Tuning False Positive Events From Creating Offenses | 112](#)

Viewing Associated Offenses

From the Log Activity tab, you can view the offense that is associated with the event.

If an event matches a rule, an offense can be generated on the **Offenses** tab.

For more information about rules, see the *Juniper Secure Analytics Administration Guide*.

When you view an offense from the **Log Activity** tab, the offense might not display if the Magistrate has not yet saved the offense that is associated with the selected event to disk or the offense has been purged from the database. If this occurs, the system notifies you.

1. Click the **Log Activity** tab.
2. Optional. If you are viewing events in streaming mode, click the **Pause** icon to pause streaming.
3. Click the **Offense** icon beside the event you want to investigate.
4. View the associated offense.

RELATED DOCUMENTATION

[Modifying Event Mapping | 111](#)

[Tuning False Positive Events From Creating Offenses | 112](#)

Modifying Event Mapping

This manual action is used to map unknown log source events to known JSA events so that they can be categorized and processed appropriately.

You can manually map a normalized or raw event to a high-level and low-level category (or QID).

For normalization purposes, JSA automatically maps events from log sources to high- and low-level categories.

For more information about event categories, see the *Juniper Secure Analytics Administration Guide*.

If events are received from log sources that the system is unable to categorize, then the events are categorized as unknown. These events occur for several reasons, including:

- **User-defined Events**— Some log sources, such as Snort, allows you to create user-defined events.

- **New Events or Older Events**— Vendor log sources might update their software with maintenance releases to support new events that JSA might not support.

NOTE: The **Map Event** icon is disabled for events when the high-level category is SIM Audit or the log source type is Simple Object Access Protocol (SOAP).

1. Click the **Log Activity** tab.
2. Optional. If you are viewing events in streaming mode, click the **Pause** icon to pause streaming.
3. Double-click the event that you want to map.
4. Click **Map Event**.
5. If you know the QID that you want to map to this event, type the QID in the **Enter QID** field.
6. If you do not know the QID you want to map to this event, you can search for a particular QID:
 - a. Choose one of the following options: **To search for a QID by category, select the high-level category from the High-Level Category list box. To search for a QID by log source type, select a log source type from the Log Source Type list box. To search for a QID by name, type a name in the QID/Name field.**
 - b. Click **Search**.
 - c. Select the **QID** you want to associate this event with.
7. Click **OK**.

RELATED DOCUMENTATION

[Tuning False Positive Events From Creating Offenses | 112](#)

[Exporting Events | 116](#)

Tuning False Positive Events From Creating Offenses

You can tune false positive events from the **event list** or **event details** page.

You can use the False Positive Tuning function to prevent false positive events from creating offenses.

You must have appropriate permissions for creating customized rules to tune false positives.

For more information about roles, see the *Juniper Secure Analytics Administration Guide*.

1. Click the **Log Activity** tab.
2. Optional. If you are viewing events in streaming mode, click the **Pause** icon to pause streaming.
3. Select the event that you want to tune.
4. Click **False Positive**.
5. In the Event/Flow Property pane on the **False Positive** window, select one of the following options:
 - Event/Flow(s) with a specific QID of <Event>
 - Any Event/Flow(s) with a low-level category of <Event>
 - Any Event/Flow(s) with a high-level category of <Event>
6. In the Traffic Direction pane, select one of the following options:
 - <Source IP Address> to <Destination IP Address>
 - <Source IP Address> to Any Destination
 - Any Source to <Destination IP Address>
 - Any Source to any Destination
7. Click **Tune**.

RELATED DOCUMENTATION

[Exporting Events | 116](#)

[Modifying Event Mapping | 111](#)

PCAP Data

IN THIS SECTION

- [Displaying the PCAP Data Column | 114](#)
- [Viewing PCAP Information | 115](#)
- [Downloading the PCAP File to Your Desktop System | 116](#)

If your JSA Console is configured to integrate with the Juniper Junos OS Platform DSM, then Packet Capture (PCAP) can be received, processed, data can be stored from a Juniper SRX-Series Services Gateway log source.

For more information about the Juniper Junos OS Platform DSM, see the *Juniper Secure Analytics Configuring DSMs Guide*.

Displaying the PCAP Data Column

The **PCAP Data** column is not displayed on the **Log Activity** tab by default. When you create search criteria, you must select the **PCAP Data** column in the Column Definition pane.

Before you can display PCAP data on the **Log Activity** tab, the Juniper SRX-Series Services Gateway log source must be configured with the PCAP Syslog Combination protocol. For more information about configuring log source protocols, see the *Juniper Secure Analytics Configuring DSMs Guide*.

When you perform a search that includes the **PCAP Data** column, an icon is displayed in the **PCAP Data** column of the search results if PCAP data is available for an event. Using the **PCAP** icon, you can view the PCAP data or download the **PCAP** file to your desktop system.

1. Click the **Log Activity** tab.
2. From the **Search** list box, select **New Search**.
3. Optional. To search for events that have PCAP data, configure the following search criteria:
 - a. From the first list box, select **PCAP data**.
 - b. From the second list box, select **Equals**.
 - c. From the third list box, select **True**.
 - d. Click **Add Filter**.
4. Configure your column definitions to include the **PCAP Data** column:
 - a. From the **Available Columns** list in the Column Definition pane, click **PCAP Data**.
 - b. Click the **Add Column** icon on the bottom set of icons to move the **PCAP Data** column to the **Columns** list.
 - c. Optional. Click the **Add Column** icon in the top set of icons to move the **PCAP Data** column to the **Group By** list.
5. Click **Filter**.

6. Optional. If you are viewing events in streaming mode, click the **Pause** icon to pause streaming.
7. Double-click the event that you want to investigate.

For more information about viewing and downloading PCAP data, see the following sections:

- ["Viewing PCAP information" on page 115](#)
- ["Downloading the PCAP file to your desktop system" on page 116](#)

Viewing PCAP Information

From the **PCAP Data** toolbar menu, you can view a readable version of the data in the PCAP file or download the PCAP file to your desktop system.

Before you can view PCAP information, you must perform or select a search that displays the **PCAP Data** column.

Before PCAP data can be displayed, the PCAP file must be retrieved for display on the user interface. If the download process takes an extended period, the **Downloading PCAP Packet information** window is displayed. In most cases, the download process is quick and this window is not displayed.

After the file is retrieved, a pop-up window provides a readable version of the PCAP file. You can read the information that is displayed on the window, or download the information to your desktop system.

1. For the event you want to investigate, choose one of the following options:
 - Select the event and click the **PCAP** icon.
 - Right-click the **PCAP** icon for the event and select **More Options >View PCAP Information**.
 - Double-click the event that you want to investigate, and then select **PCAP Data >View PCAP Information** from the event details toolbar.
2. If you want to download the information to your desktop system, choose one of the following options:
 - Click **Download PCAP File** to download the original PCAP file to be used in an external application.
 - Click **Download PCAP Text** to download the PCAP information in .TXT format
3. Choose one of the following options:
 - If you want to open the file for immediate viewing, select the **Open with** option and select an application from the list box.

- If you want to save the list, select the **Save File** option.

4. Click **OK**.

Downloading the PCAP File to Your Desktop System

You can download the PCAP file to your desktop system for storage or for use in other applications.

Before you can view a PCAP information, you must perform or select a search that displays the PCAP Data column. See **Displaying the PCAP data column**.

1. For the event you want to investigate, choose one of the following options:

- Select the event and click the **PCAP** icon.
- Right-click the PCAP icon for the event and select **More Options >Download PCAP File** .
- Double-click the event you want to investigate, and then select **PCAP Data >Download PCAP File** from the event details toolbar.

2. Choose one of the following options:

- If you want to open the file for immediate viewing, select the **Open with** option and select an application from the list box.
- If you want to save the list, select the **Save File** option.

3. Click **OK**.

RELATED DOCUMENTATION

[Exporting Events | 116](#)

[Modifying Event Mapping | 111](#)

[Tuning False Positive Events From Creating Offenses | 112](#)

Exporting Events

The length of time that is required to export your data depends on the number of parameters specified.

You can export events in Extensible Markup Language (XML) or Comma-Separated Values (CSV) format.

1. Click the **Log Activity** tab.
2. Optional. If you are viewing events in streaming mode, click the **Pause** icon to pause streaming.
3. From the **Actions** list box, select one of the following options:
 - **Export to XML >Visible Columns** - Select this option to export only the columns that are visible on the Log Activity tab. This is the recommended option.
 - **Export to XML >Full Export (All Columns)** - Select this option to export all event parameters. A full export can take an extended period of time to complete.
 - **Export to CSV >Visible Columns** - Select this option to export only the columns that are visible on the **Log Activity** tab. This is the recommended option.
 - **Export to CSV >Full Export (All Columns)** - Select this option to export all event parameters. A full export can take an extended period of time to complete.
4. If you want to resume your activities while the export is in progress, click **Notify When Done**.

When the export is complete, you receive notification that the export is complete. If you did not select the **Notify When Done** icon, the status window is displayed.

RELATED DOCUMENTATION

[Modifying Event Mapping | 111](#)

[Tuning False Positive Events From Creating Offenses | 112](#)

7

CHAPTER

Network Activity Monitoring

[Network Activity Monitoring | 119](#)

[Flow Pipeline | 120](#)

[Flow Sources | 121](#)

[Flow Aggregation | 132](#)

[Flow Direction | 134](#)

[Application Identification | 139](#)

[Superflows | 143](#)

[Deduplication | 146](#)

[View Flow Data on the Network Activity Tab | 147](#)

[VLAN Fields | 155](#)

[Configuring a Flow Processor | 158](#)

[Tuning False Positive Events from Creating Offenses | 162](#)

Network Activity Monitoring

JSA collects information about the way that devices in your network communicate with each other. The record of the communication as it occurs across the network is called a flow.

In JSA, flows appear on the **Network activity** tab, and you must have permissions to view it.

By default, the **Network Activity** tab displays flows in streaming mode. To analyze flows that were already received, you can pause streaming mode, and conduct searches and apply filters to the flows.

If you previously configured a saved search as the default, the results of that search are automatically displayed when you access the **Network Activity** tab.

Difference between events and flows

An event is a record of an activity that occurred on your network at a single point in time. It provides information about an activity that happened and which assets were implicated. For example, if a user tries to authenticate against your firewall and is unsuccessful, the firewall sends a system log to JSA. The authentication attempt is recorded as an event. You view information about events on the **Log Activity** tab.

Unlike an event, a flow is network activity that occurs over time. The flow record shows the actual messages that were sent and received by devices as they communicated with each other over the network. For example, a flow might show when a user sends an email, goes to a web page, downloads a file, or uses social media. The web request might download files such as images and ads over a few seconds, or last a few hours if the user is watching a movie.

Compared to events, flow traffic provides a complete view of what is happening on your network. They show you what was happening before, during, and after the event occurred.

RELATED DOCUMENTATION

[Flow Pipeline](#) | 120

[Flow Sources](#) | 121

[Flow Aggregation](#) | 132

Flow Pipeline

Flows that come into JSA go through an in-depth process to extract additional information about the network communication, looking for indicators that a security incident might have occurred.

Figure 2: Flow Pipeline in JSA



Flow Processor process

The Flow Processor process collects data from a variety of flow sources. It aggregates the data by parsing and normalizing the data, accruing information over a one minute period. It then analyzes the flow to extract additional information such as determining the application and flow direction, and creating superflows before handing it off to the ecs-ec process.

ecs-ec process

The ecs-ec process further parses the flow record and performs additional unification and processing, such as deduplication, asymmetric recombination, licensing, domain tagging, custom flow properties, and flow forwarding. The flow is then passed to the Custom Rule Engine (CRE) to determine if the flow triggers a rule, which might indicate that a security incident has occurred.

RELATED DOCUMENTATION

[Flow Sources | 121](#)

[Flow Aggregation | 132](#)

[Flow Direction | 134](#)

Flow Sources

IN THIS SECTION

- [NetFlow | 123](#)
- [IPFIX | 125](#)
- [sFlow | 129](#)
- [J-Flow | 130](#)
- [Packeteer | 131](#)
- [Network Interface | 131](#)

JSA can receive flows from many different types of flow sources. The flow sources are classified as either internal or external.

Internal Flow Sources

Internal flow sources collect raw packets from either a network tap device or a span or mirror port that is connected to a network interface card. These sources provide raw packet data as it appears on the network, and sends it to a monitoring port on a JSA Flow Processor, which converts the packet data into flow records.

Internal flow sources can be multithreaded. JSA does not keep the entire packet payload. Instead, it captures only some of the packets from the beginning of the communication. This snapshot is referred to as the *payload* or *content capture*.

Figure 3: Packet Capture from an Internal Flow Source

The screenshot displays a network packet capture interface. The top section shows a list of packets with columns for No., Source, Destination, SrcMAC, and Info. Packet 14 is highlighted, showing a GET request for / HTTP/1.1. Below the list, a detailed view of the selected packet is shown, including Ethernet II, Internet Protocol Version 4, and Hypertext Transfer Protocol details. The HTTP request is expanded to show the User-Agent, Host, and Accept headers, and the full request URI is highlighted as http://192.168.115.186:81/. The bottom section shows the raw packet data in hexadecimal and ASCII format.

```

No.    Source          Destination      SrcMAC          Info
11 192.168... 192.16... 00:00:5e:00:53... 52452-8181 [SYN] Seq=0 Win=14600 Len=0 MSS=1460 SACK_PERM=1 TSval=298763421
12 192.168... 192.16... 00:00:5e:00:53... 8181-52452 [SYN, ACK] Seq=0 Ack=1 Win=14400 Len=0 MSS=1460 SACK_PERM=1 TSva
13 192.168... 192.16... 00:00:5e:00:53... 52452-8181 [ACK] Seq=1 Ack=1 Win=14720 Len=0 TSval=2987634219 TSecr=1748481
14 192.168... 192.16... 00:00:5e:00:53... GET / HTTP/1.1
15 192.168... 192.16... 00:00:5e:00:53... 8181-52452 [ACK] Seq=1 Ack=182 Win=15616 Len=0 TSval=1748481214 TSecr=29876
16 192.168... 192.16... 00:00:5e:00:53... HTTP/1.1 200 OK (text/html)
17 192.168... 192.16... 00:00:5e:00:53... 52452-8181 [ACK] Seq=182 Ack=250 Win=15744 Len=0 TSval=2987634230 TSecr=174

> Frame 14: 247 bytes on wire (1976 bits), 247 bytes captured (1976 bits)
> Ethernet II, Src: IcannIam_00:53:00 (00:00:5e:00:53:00), Dst: IcannIam_00:53:01 (00:00:5e:00:53:01)
> Internet Protocol Version 4, Src: 192.168.115.143, Dst: 192.168.115.186
> Transmission Control Protocol, Src Port: 52452, Dst Port: 8181, Seq: 1, Ack: 1, Len: 181
> Hypertext Transfer Protocol
  > GET / HTTP/1.1\r\n
    User-Agent: curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.18 Basic ECC zlib/1.2.3 libidn/1.18 libssh:
    Host: 192.168.115.186:81\r\n
    Accept: */*\r\n
    \r\n
    [Full request URI: http://192.168.115.186:81/]
    [HTTP request 1/1]
    [Response in frame: 16]

0000 00 00 5e 00 53 01 00 00 5e 00 53 00 08 00 45 00 ..^..S... ^..S...E.
0010 00 e9 0b a0 40 00 40 06 c5 d4 c0 a8 73 8f c0 a8 ...@.@. ....S...
0020 73 ba cc e4 1f f5 6a 9f a5 b6 f9 bd 64 6d 80 18 s.....j. ....dm..
0030 00 73 69 76 00 00 01 01 08 0a b2 13 ae 2b 68 37 ..5iv.... ..+h7
0040 b4 be 47 45 54 20 2f 20 48 54 54 50 2f 31 2e 31 ..GET / HTTP/1.1
0050 0d 0a 55 73 65 72 2d 41 67 65 6e 74 3a 20 63 75 ..User-A gent: cu
0060 72 6c 2f 37 2e 31 39 2e 37 20 28 78 38 36 5f 36 rL/7.19. 7 (x86_6
0070 34 2d 72 65 64 68 61 74 2d 6c 69 6e 75 78 2d 67 4-redhat -linux-g
0080 6e 75 29 20 6c 69 62 63 75 72 6c 2f 37 2e 31 39 nu) libc url/7.19
0090 2e 37 20 4e 53 53 2f 33 2e 31 38 20 42 61 73 69 .7 NSS/3 .18 Basi
00a0 63 20 45 43 43 20 7a 6c 69 62 2f 31 2e 32 2e 33 c ECC zl ib/1.2.3
00b0 20 6c 69 62 69 64 6e 2f 31 2e 31 38 20 6c 69 62 libidn/ 1.18 lib
00c0 73 73 68 32 2f 31 2e 34 2e 32 0d 0a 48 6f 73 74 ssh2/1.4 ..Host
00d0 3a 20 31 39 32 2e 31 36 38 2e 31 31 35 2e 31 38 : 192.16 8.115.18
00e0 36 3a 38 31 38 31 0d 0a 41 63 63 65 70 74 3a 20 6:8181.. Accept:
00f0 2a 2f 2a 0d 0a 0d 0a */*....

```

External Flow Sources

JSA also supports external flow sources, such as routers that send common network monitoring protocols, such as NetFlow, IPFIX, sFlow J-Flow, and Packeteer data.

These external flow sources can provide a different level of visibility than internal flow sources. For example, NetFlow records can provide both the router interface that the packets crossed, as well as the ASN record number of the originating network. When using IPFIX, additional fields that are not parsed into normalized fields can be placed into the payload as name value pairs, which can then be used as custom properties.

External sources do not require as much CPU utilization to process so you can send them directly to a Flow Processor. In this configuration, you can have a dedicated Flow Processor, both receiving and creating flow data.

Figure 4: Packet Capture from an External Flow Source

No.	Source	Destination	Src/MAC	Info
1	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (77 bytes) Obs-Domain-ID= 1 [Data-Template:257] [Data:257]
2	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (77 bytes) Obs-Domain-ID= 1 [Data-Template:258] [Data:258]
3	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (93 bytes) Obs-Domain-ID= 1 [Data-Template:259] [Data:259]
4	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (77 bytes) Obs-Domain-ID= 1 [Data-Template:260] [Data:260]
5	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (77 bytes) Obs-Domain-ID= 1 [Data-Template:261] [Data:261]
6	127.0.0.	127.0...	00:00:00:00:00...	IPFIX flow (109 bytes) Obs-Domain-ID= 1 [Data-Template:262] [Data:262]

```

Ethernet II, Src: 00:00:00_00:00:00 (00:00:00:00:00:00), Dst: 00:00:00_00:00:00 (00:00:00:00:00:00)
Internet Protocol Version 4, Src: 127.0.0.1, Dst: 127.0.0.1
User Datagram Protocol, Src Port: 56520, Dst Port: 2055
Cisco NetFlow/IPFIX
  Version: 10
  Length: 77
  Timestamp: Jul 26, 2018 08:16:45.000000000 AEST
  FlowSequence: 1
  Observation Domain Id: 1
  Set 1 [id=2] (Data Template): 257
    FlowSet Id: Data Template (V10 [IPFIX]) (2)
    FlowSet Length: 36
    Template (Id = 257, Count = 7)
      Template Id: 257
      Field Count: 7
      Field (1/7): PROTOCOL
      Field (2/7): IP_SRC_ADDR
      Field (3/7): L4_SRC_PORT
      Field (4/7): IP_DST_ADDR
      Field (5/7): L4_DST_PORT
      Field (6/7): BYTES
      Field (7/7): PKTS
  Set 2 [id=257] (1 flows)
    FlowSet Id: (Data) (257)
    FlowSet Length: 25
    [Template Frame: 1]
    Flow 1
      Protocol: TCP (6)
      SrcAddr: 10.10.10.10
      SrcPort: 1111 (1111)
      DstAddr: 192.168.0.1
      DstPort: 80 (80)
      Octets: 111111
      Packets: 111
  
```

NetFlow

NetFlow is a proprietary accounting technology that is developed by Cisco Systems. NetFlow monitors traffic flows through a switch or router, and interprets the client, server, protocol, and port that is used. It also counts the number of bytes and packets, and sends that data to a NetFlow collector.

The process of sending data from NetFlow is often referred to as a NetFlow Data Export (NDE).

JSA accepts NetFlow Data Exports (NDE) so that it functions as a NetFlow collector. JSA supports NetFlow versions 1, 5, 7, and 9.

While NetFlow expands the amount of the network that is monitored, it uses a connection-less protocol (UDP) to deliver NDEs. After an NDE is sent from a switch or router, the NetFlow record is purged. UDP doesn't guarantee the delivery of data. As a result, inaccurate presentations of both traffic volumes and bidirectional flows, and reduced alerting capabilities, might result with a NetFlow flow source.

For more information about NetFlow, see the [Cisco web site](#).

NetFlow Flow Source Configuration

When you configure an external flow source for NetFlow, you must do the following tasks:

- Make sure that the appropriate firewall rules are configured.

If you change your **External Flow Source Monitoring Port** parameter in the Flow Processor configuration, you must also update your firewall access configuration.
- Make sure that the appropriate ports are configured for your Flow Processor.

For more information, see "[Configuring a Flow Processor](#)" on page 158.

NetFlow Flow Source Template

Juniper suggests that, at minimum, the following fields are included in the NetFlow flow source template:

- FIRST_SWITCHED
- LAST_SWITCHED
- PROTOCOL
- IPV4_SRC_ADDR
- IPV4_DST_ADDR
- L4_SRC_PORT
- L4_DST_PORT
- IN_BYTES or OUT_BYTES
- IN_PKTS or OUT_PKTS
- TCP_FLAGS (TCP flows only)

Supported Fields

The following lists show some of the types of fields that are supported for NetFlow flow sources.

VLAN fields

The following VLAN fields are supported for NetFlow:

- vlanId (IANA Element ID 58)
- postVlanId (IANA Element ID 59)
- dot1qVlanId (IANA Element ID 243)
- dot1qPriority (IANA Element ID 244)
- dot1qCustomerVlanId (IANA Element ID 245)
- dot1qCustomerPriority (IANA Element ID 246)
- postDot1qVlanId (IANA Element ID 254)
- postDotqCustomerVlanId (IANA Element ID 255)
- dot1qDEI (IANA Element ID 388)
- dot1qCustomerDEI (IANA Element ID 389)

MAC Address Fields

The following MAC address fields are supported for NetFlow:

- sourceMacAddress (IANA Element ID 56)
- postDestinationMacAddress (IANA Element ID 57)
- DestinationMacAddress (IANA Element ID 80)
- postSourceMacAddress (IANA Element ID 81)

IPFIX

Internet Protocol Flow Information Export (IPFIX) is an accounting technology that monitors traffic flows through a switch or router. It interprets the traffic to determine the client, server, protocol, and port that is used. It also counts the number of bytes and packets, and sends that data to an IPFIX collector.

The process of sending IPFIX data is often referred to as a NetFlow Data Export (NDE), but IPFIX provides more flow information and deeper insight than NetFlow v9.

JSA accepts NDEs so that it functions as an IPFIX collector. IPFIX uses User Datagram Protocol (UDP) to deliver NDEs. After an NDE is sent from the IPFIX forwarding device, the IPFIX record might be purged.

IPFIX Flow Source Configuration

When you configure an external flow source for IPFIX, you must do the following tasks:

- Add a NetFlow flow source.

NOTE: Your JSA might include a default NetFlow flow source. If it does, JSA can use the default NetFlow flow source to process the IPFIX flows.

To confirm that your system includes a default NetFlow flow source, on the **Admin** tab, select **Flow Sources**. If **default_Netflow** is listed in the flow source list, IPFIX is already configured.

- Ensure that the appropriate firewall rules are configured.

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

- Ensure that the appropriate ports are configured for your Flow Processor.

For more information, see ["Configuring a Flow Processor" on page 158](#).

IPFIX Flow Source Template

Ensure that the IPFIX template from the IPFIX source includes the following IANA-listed Information Elements:

- protocolIdentifier (4)
- sourceIPv4Address (8)
- destinationIPv4Address (12)
- sourceTransportPort (7)
- destinationTransportPort (11)
- octetDeltaCount (1) or postOctetDeltaCount (23)
- packetDeltaCount (2) or postPacketDeltaCount (24)

- tcpControlBits (6) (TCP flows only).
- flowStartSeconds (150) or flowStartMilliseconds (152) or flowStartDeltaMicroseconds (158)
- flowEndSeconds (151) or flowEndMilliseconds (153) or flowEndDeltaMicroseconds (159)

Supported Fields

The following lists show some of the types of fields that are supported for IPFIX flow sources.

To add support for additional IPFIX fields that are not shown by JSA, you can use the `/api/ariel/taggedfields` API to create a new tagged field.

VLAN fields

The following VLAN fields are supported for IPFIX:

- vlanId (IANA Element ID 58)
- postVlanId (IANA Element ID 59)
- dot1qVlanId (IANA Element ID 243)
- dot1qPriority (IANA Element ID 244)
- dot1qCustomerVlanId (IANA Element ID 245)
- dot1qCustomerPriority (IANA Element ID 246)
- postDot1qVlanId (IANA Element ID 254)
- postDot1qCustomerVlanId (IANA Element ID 255)
- dot1qDEI (IANA Element ID 388)
- dot1qCustomerDEI (IANA Element ID 389)

MAC Address Fields

The following MAC address fields are supported for IPFIX:

- sourceMacAddress (IANA Element ID 56)
- postDestinationMacAddress (IANA Element ID 57)
- DestinationMacAddress (IANA Element ID 80)
- postSourceMacAddress (IANA Element ID 81)

Network Address Translation (NAT) fields

The following fields are supported for Network Address Translation (NAT) and Network Address Port Translation (NAPT):

- postNATSourceIPv4Address (IANA Element ID 225)
- postNATDestinationIPv4Address (IANA Element ID 226)
- postNAPTSourceTransportPort (IANA Element ID 227)
- postNAPTDestinationTransportPort (IANA Element ID 228)

MPLS Fields

The following MPLS fields are supported for IPFIX:

- mplsTopLabelType (IANA Element 46)
- mplsTopLabelIPv4Address (IANA Element 47)
- mplsTopLabelStackSection (IANA Element 70)
- mplsLabelStackSection2 (IANA Element 71)
- mplsLabelStackSection3 (IANA Element 72)
- mplsLabelStackSection4 (IANA Element 73)
- mplsLabelStackSection5 (IANA Element 74)
- mplsLabelStackSection6 (IANA Element 75)
- mplsLabelStackSection7 (IANA Element 76)
- mplsLabelStackSection8 (IANA Element 77)
- mplsLabelStackSection9 (IANA Element 78)
- mplsLabelStackSection10 (IANA Element 79)
- mplsVpnRouteDistinguisher (IANA Element 90)
- mplsTopLabelPrefixLength (IANA Element 91)
- mplsTopLabelIPv6Address (IANA Element 140)
- mplsPayloadLength (IANA Element 194)
- mplsTopLabelTTL (IANA Element 200)

- mplsLabelStackLength (IANA Element 201)
- mplsLabelStackDepth (IANA Element 202)
- mplsTopLabelExp (IANA Element 203)
- postMplsTopLabelExp (IANA Element 237)
- pseudoWireType (IANA Element 250)
- pseudoWireControlWord (IANA Element 251)
- mplsLabelStackSection (IANA Element 316)
- mplsPayloadPacketSection (IANA Element 317)
- sectionOffset (IANA Element 409)
- sectionExportedOctets (IANA Element 410)

sFlow

sFlow is a multi-vendor and user standard for sampling technology that provides continuous monitoring of application-level traffic flows on all interfaces simultaneously.

An sFlow combines interface counters and flow samples into sFlow datagrams that are sent across the network to an sFlow collector. sFlow traffic is based on sampled data and, therefore, might not represent all network traffic.

JSA supports flow sources for sFlow versions 2, 4, and 5.

sFlow uses a connection-less protocol (UDP). When data is sent from a switch or router, the sFlow record is purged. UDP doesn't guarantee delivery of the data. As such, inaccurate presentations of both traffic volumes and bidirectional flows, and reduced alerting capabilities, might result when using an sFlow flow source.

For more information, see the [sFlow website](#).

sFlow Flow Source Configuration

When you configure an external flow source for sFlow, you must do the following tasks:

- Ensure that the appropriate firewall rules are configured.
- Ensure that the appropriate ports are configured.

For more information, see ["Configuring a Flow Processor" on page 158](#).

J-Flow

J-Flow is a proprietary accounting technology used by Juniper Networks that allows you to collect IP traffic flow statistics.

J-Flow enables you to export data to a UDP port on a J-Flow collector. You can also enable J-Flow on a router or network interface to collect network statistics for specific locations on your network.

J-Flow uses a connection-less protocol (UDP). When data is sent from a switch or router, the J-Flow record is purged. UDP doesn't guarantee delivery of the data. As such, inaccurate presentations of both traffic volumes and bidirectional flows, and reduced alerting capabilities, might result when using a J-Flow flow source. J-Flow traffic is based on sampled data and, therefore, might not represent all network traffic.

For more information about J-Flow, see the [Juniper Networks website](#).

J-Flow Flow Source Configuration

When you configure an external flow source for J-Flow, you must do the following tasks:

- Ensure that the appropriate firewall rules are configured.
- Ensure that the appropriate ports are configured.

For more information, see ["Configuring a Flow Processor" on page 158](#).

Supported VLAN Fields

The following VLAN fields are supported for J-Flow:

- vlanId
- postVlanId
- dot1qVlanId
- dot1qPriority
- dot1qCustomerVlanId
- dot1qCustomerPriority
- dot1qDEI

- dot1qCustomerDEI
- postDot1qVlanId
- postDotqCustomerVlanId

Packeteer

Packeteer devices collect, aggregate, and store network performance data.

After you configure an external flow source for Packeteer, you can send flow information from a Packeteer device to JSA.

Packeteer uses a connection-less protocol (UDP). When data is sent from a switch or router, the Packeteer record is purged. Because UDP doesn't guarantee delivery of the data, inaccurate presentations of both traffic volumes and bidirectional flows, and reduced alerting capabilities might result when using a Packeteer flow source.

Packeteer Flow Source Configuration

To configure Packeteer as an external flow source, you must do the following tasks:

- Ensure that the appropriate firewall rules are configured.
- Ensure that you configure Packeteer devices to export flow detail records, and configure the Flow Processor as the destination for the data export.
- Ensure that the appropriate ports are configured.
- Ensure that the class IDs from the Packeteer devices can automatically be detected by the Flow Processor.

For more information, see ["Configuring a Flow Processor" on page 158](#).

Network Interface

JSA can monitor packets that come in on any network interface card that is installed on your system. The network interface card must be installed for JSA to show it in the list of configurable packet-based flow sources.

When you configure the **Network Interface** flow source, configure only one log source for each Ethernet interface. To filter the network traffic that comes in on the flow source, specify a Berkeley Packet Filter (BPF) in the **Filter String** field.

After the flow source is configured and the changes are deployed, you can view the network traffic on the **Network Activity** tab.

You can customize the way that JSA processes the network traffic. For example, you can configure the way that applications are detected, superflow thresholds, flow direction configuration, and network hierarchy. You can also write rules, perform queries, and filter the network traffic.

For more information, see "[Configuring a Flow Processor](#)" on page 158.

RELATED DOCUMENTATION

[Flow Aggregation](#) | 132

[Flow Direction](#) | 134

[Application Identification](#) | 139

Flow Aggregation

IN THIS SECTION

- [Flow Capacity Limits](#) | 133

JSA combines information together to give you more information about a single flow without sending more flow records. This process is known as aggregation.

The flow shows a communication session between two hosts by normalizing the packet attributes into a flow record that includes the following information:

- Source IP address
- Source port
- Destination IP address
- Destination port

- Protocol
- Flow ID (flow source dependent)
- VLAN fields (flow source dependent)
- VXLAN fields (flow source dependent)

As the hosts continue to communicate, information such as the byte and packet counters and the payload capture is aggregated into a single flow record. For communications that span more than 1 minute, JSA reports on the current metrics for the flow at the end of each 1-minute interval. The entire communication session is represented by multiple flow records that have the same **First Packet Time**, but with incremental **Last Packet Time** values.

The flow information is updated if the attributes are the same. When one or more attributes change, the flow is assumed to be unique, and a new flow record is created.

Flow Capacity Limits

Flow capacity limits ensure that the Flow Processor process in JSA is not overloaded.

When the Flow Processor process receives more traffic than it can deal with, an overflow record is created for each protocol that is observed in the excess traffic. These records are easily identified because they have a source IP address of 127.0.0.4 and a destination IP address of 127.0.0.5.

For example, JSA determines that the flow capacity limit of your Flow Processor is 100,000 flows. During a peak period, the appliance captures 120,000 flows in a one minute interval. The excess 20,000 flows are not parsed, but instead an overflow record is created for each protocol that is seen in the 20,000 flows. The overflow record includes byte and packet counters, but information such as source or destination IP addresses, ports, and payload capture is not collected and stored.

Flow Capacity Limits

Flow capacity is determined based on a number of different factors:

- **Deployment flow limit** This flow limit is based on the sum of all flow licenses across your deployment.
- **Hardware flow limit** The hardware flow limit is the recommended number of flows calculated based on the available CPUs and memory.
- **User flow limits** You can set the maximum number of flows that you want JSA to process at one time.

If a user flow limit is set, it takes precedence over both the deployment flow limit and the hardware limit.

If no user limit is set, the minimum of either the hardware limit or the deployment limit is used.

NOTE: To ensure that you achieve the maximum benefit from your flows per minute (FPM) license, flow capacity limits are enforced after aggregation. Updates to existing flows within the 1-minute reporting interval do not contribute to your FPM license limit.

RELATED DOCUMENTATION

[Flow Direction | 134](#)

[Application Identification | 139](#)

[Superflows | 143](#)

Flow Direction

IN THIS SECTION

- [Flow Direction Algorithms | 136](#)
- [Displaying the Flow Direction Algorithm Field in Search Results | 137](#)
- [Common Destination Ports | 138](#)

The Flow Processor process analyzes each flow to determine the direction of the network communication.

In some cases, the flow traffic is bidirectional where the client communicates with the server and the server responds to the client. In this scenario, both the client and the server operate as though they are the source and the other is the destination. To address this, JSA sets the flow direction to ensure that the source and destination devices are reported consistently throughout the entire communication session. The flow data is normalized and all flows follow the same convention, where **Destination** always refers to the server, and **Source** always refers to the client.

To determine the flow direction, JSA analyzes the flow to determine whether the source and destination ports match the list of common destination ports that are defined in the JSA configuration. The flow direction is reversed when the following criteria is matched:

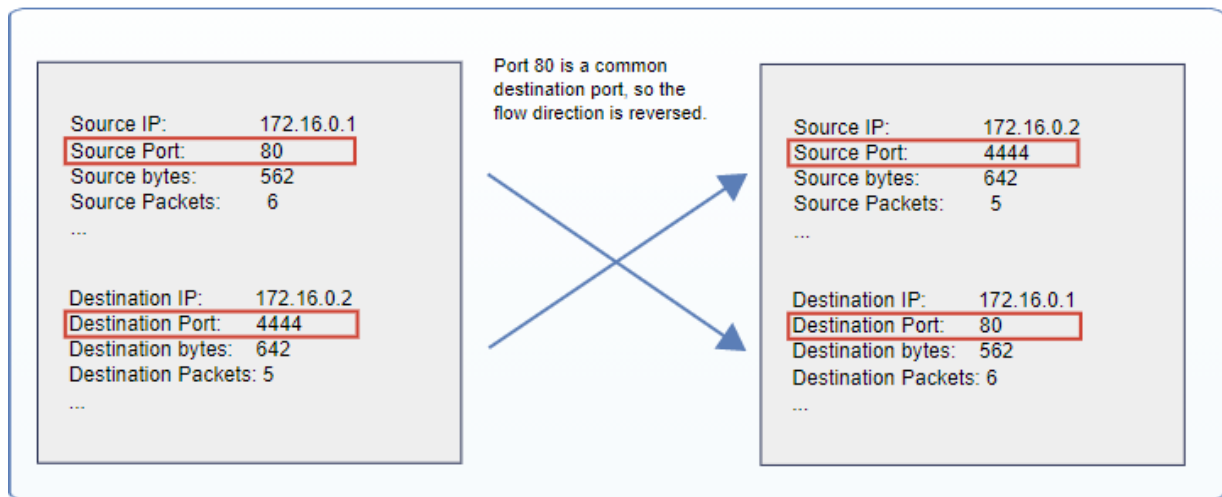
- If the destination port **does not** match the list of common destination ports, reverse the flow direction if **either** of the following conditions are true:
 - The source port is a common destination port.
 - The source port is less than 1024 and the destination port is greater than 1024.
- If the destination port **does** match the list of common destination ports, reverse the flow direction if **both** of the following conditions are true:
 - The source port is a common destination port.
 - The source port is less than 1024 and the destination port is greater than 1024.

If the flow does not match any of the flow direction criteria, JSA uses the flow arrival time to determine the flow direction.


TIP: If you do not want JSA to determine the flow direction, set the **Use Common Destination Port** field to **No** when you configure the Flow Processor. For more information, see "[Configuring a Flow Processor](#)" on page 158.

Example: Flow Direction Reversed by JSA

In this flow, the source port is 80, which is a common destination port. The destination port is higher than 1024 which, according to RFC1700, excludes it as a common destination port. In this case, JSA flipped the flow direction.



On the **Flow Information** window, you can see the flow direction algorithm that was used to set the direction.

Flow Information							
Protocol	tcp_ip		Application	Web.Web.Misc			
Magnitude	 (6)		Relevance	10	Severity	1	
First Packet Time	29 Oct 2018, 08:59:02		Last Packet Time	29 Oct 2018, 08:59:02		Storage Time	29 Oct 2018, 09:00:02
Event Name	Web.HTTPWeb						
Low Level Category	Web						
Application Determination Algorithm	QRadar port based mapping (4)						
Flow Direction Algorithm	Single common destination port (1)						
Domain	Default Domain						
Source and Destination Information							
Source IP	172.16.0.2		Destination IP	172.16.0.1			
Source Asset Name	172.16.0.2		Destination Asset Name	172.16.0.1			
Source Port	4444		Destination Port	80			

Flow Direction Algorithms

Flow direction algorithms are used to detect which side of the communication is more likely to be the destination device, and reverses the flow direction as required. The algorithms provide information on how the traffic originally appeared on the network, and which features of the traffic caused it to be reversed.

The following table displays the values that are used in the flow direction algorithm.

Numeric value	Algorithm	Description
1	Single common destination port (reversed)	Either the source port or the destination port was found in the list of common destination ports, and JSA reversed the flow direction.

(Continued)

Numeric value	Algorithm	Description
2	Both common destination ports but one was RFC 1700 preferred (reversed)	<p>Both the source port and the destination port are defined as common destination ports. According to RFC1700, the source port is a preferred destination port, so JSA reversed the flow direction.</p> <p>The RFC1700 preferred ports are in the range of 0 to 1023, which are controlled and assigned by the Internet Assigned Number Authority (IANA).</p>
3	Arrival time	<p>The flow does not match the criteria for any other flow direction algorithm. JSA used the flow arrival time to determine the flow direction.</p> <p>The Flow Processor process assumes that the request was received before the response, and the flow direction remains as it was received.</p>
4	Flow exporter	The flow direction is set by an external flow exporter, such as a Packeteer device.
5	Single common destination port (unaltered)	Either the source port or the destination port was found in the list of common destination ports. JSA did not alter the flow direction.
6	Both common destination ports but one was RFC 1700 preferred (unaltered)	Both the source port and the destination port are defined as common destination ports. According to RFC1700 , the destination port is a preferred destination port, so JSA did not alter the flow direction.

Displaying the Flow Direction Algorithm Field in Search Results

Use the search feature to add the flow direction algorithm to the **Flow Details** window. You can use the flow direction algorithm to identify the criteria that caused JSA to reverse the flow direction.

1. To display the flow direction algorithm field on the **Flow Details** window, follow these steps:

- a. Click the **Network Activity** tab.
- b. From the **Search** list, select **New Search**.
- c. In the **Column Definition** section, scroll down the list of available columns and add **Flow Direction Algorithm** to the list of columns to display.
- d. Click **Filter**.

The **Flow Direction Algorithm** column appears on the **Network Activity** tab, displaying a value that represents the algorithm that was used.

2. To display the flow direction algorithm in an advanced search, use the LOOKUP function to show the text description for the enumerated flow direction algorithm field. For example, the AQL query might look like this:

```
SELECT sourceip, destinationip, LOOKUP('flow direction algorithm', "flow direction algorithm") FROM flows
```

Common Destination Ports

Common destination ports are used by JSA to determine the direction of a network communication, and to ensure that the source and destination devices are reported consistently throughout the entire communication session.

Use the Common Destination Ports API to view the common destination ports that are configured in your system, and add new common destination ports.

NOTE: Earlier releases of JSA used the **appid_map.conf** configuration file to determine the list of common destination ports.

You can use the Common Destination Port API to do the following actions:

- Get a list of the common destination ports that are defined by default.

config/flow/applications/common_destination_ports/default_configurations

The default configurations are read-only. They are provided as a system backup in case the configuration is mistakenly deleted or changed.

- Get information about the common destination ports that are active in the deployed system. This endpoint is also used to create, update, or delete active common destination ports.

config/flow/applications/common_destination_ports/active_configurations

When you use the API to change the common destination ports, the Flow Processor process automatically loads the configuration. You do not need to deploy the configuration to propagate the changes to the system.

For more information about using the JSA RESTful API, see *RESTful API Overview*.

RELATED DOCUMENTATION

[Application Identification | 139](#)

[Superflows | 143](#)

[Deduplication | 146](#)

Application Identification

IN THIS SECTION

- [Displaying the Application Determination Algorithm Field in Search Results | 142](#)

The Flow Processor process uses algorithms to determine the flow application. Each algorithm relies on different types of information to determine the application.

QRadar Network Insights relies on its own set of inspectors and application detection methods. QRadar Network Insights, the Flow Processor algorithms are used only when QRadar Network Insights cannot identify a specific protocol.

The application determination algorithms are shown in the following table.


Table 23: Application Determination Algorithms

Numeric value	Algorithm name	Description
2	Application signatures	<p>A payload-based algorithm that looks at the way that the payload is structured.</p> <p>This algorithm uses information from the signatures.xml file.</p>
3	State-based decoding	<p>A payload-based algorithm that uses complex internal logic.</p>
4	JSA port-based mapping	<p>A port-based algorithm that uses a pre-defined list of application mappings.</p> <p>This algorithm uses information from the /opt/qradar/conf/appid_map.conf file.</p>
5	User port-based mapping	<p>A port-based algorithm that uses a customizable list of application mappings.</p> <p>Use this algorithm to add new port-based mappings or reclassify existing mappings that come with JSA.</p> <p>This algorithm uses information from the /opt/qradar/conf/user_application_mapping.conf file.</p>
6	ICMP protocol mapping	<p>A protocol-based algorithm that looks at the protocol type and code.</p>
7	Flow exporter	<p>An algorithm that relies on the Flow Exporter to determine the application.</p>
8	QNI Application Signatures	<p>This algorithm is used by QRadar Network Insights.</p>
9	X-Force Web Application Classification	<p>This algorithm is used by QRadar Network Insights.</p>

Table 23: Application Determination Algorithms (Continued)

Numeric value	Algorithm name	Description
10	QNI port heuristics	This algorithm is used by QRadar Network Insights. It indicates that the application is identified by using port based heuristics, and represents a low degree of confidence.
11	QNI initial data	This algorithm is used by QRadar Network Insights. It indicates that the application is identified by using the initial data in the flow session, and represents a medium degree of confidence.
12	QNI parsers	This algorithm is used by QRadar Network Insights. It indicates that the application is identified by parsing the available data, and represents the highest degree of confidence.

You can see which type of application detection algorithm that is used in the **Application Determination Algorithm** field on the **Flow Information** window.

Flow Information							
Protocol	tcp_ip	Application	Web.Web.Misc				
Magnitude	 (6)	Relevance	10	Severity	1	Credibility	10
First Packet Time	29 Oct 2018, 08:59:02	Last Packet Time	29 Oct 2018, 08:59:02	Storage Time	29 Oct 2018, 09:00:02		
Event Name	Web.HTTPWeb						
Low Level Category	Web						
Application Determination Algorithm	QRadar port based mapping (4)						
Flow Direction Algorithm	Single common destination port (1)						
Domain	Default Domain						
Source and Destination Information							
Source IP	172.16.0.2	Destination IP	172.16.0.1				
Source Asset Name	172.16.0.2	Destination Asset Name	172.16.0.1				
Source Port	4444	Destination Port	80				

Custom Applications

If your organization has nonstandard or customized applications, you can add them to the `/opt/qradar/conf/user_application_mapping.conf` or `signatures.xml` files.

You can use the **Application Determination Algorithm** field to check that the correct algorithm was used to identify your customized applications. For example, you might define a custom application based on the port usage. Flows from that application are identified by algorithm 5, which is **User Port Based Mapping**. By verifying the algorithm that is used to identify the application, you can assign a level of confidence to the application mapping.

For more information, see the *Juniper Secure Analytics Application Configuration Guide*.

Displaying the Application Determination Algorithm Field in Search Results

Use the search feature to add the application determination algorithm to the **Flow Details** window. You can use the application identification algorithm to identify the criteria that JSA used to identify which application the flow originated from.

1. To display the application determination algorithm field on the **Flow Details** window, follow these steps:
 - a. Click the **Network Activity** tab.
 - b. From the **Search** list, select **New Search**.
 - c. In the **Column Definition** section, scroll down the list of available columns and add **Application Determination Algorithm** to the list of columns to display.
 - d. Click **Filter**.

The **Application Determination Algorithm** column appears on the **Network Activity** tab, displaying a value that represents the algorithm that was used.

2. To display the application determination algorithm in an advanced search, use the LOOKUP function to show the text description for the enumerated algorithm field.

For example, the AQL query might look like this.

```
SELECT LOOKUP('application determination algorithm', "application determination algorithm"),  
* FROM flows
```

RELATED DOCUMENTATION

[Superflows](#) | 143

[Deduplication](#) | 146

[Flow Direction](#) | 134

Superflows

IN THIS SECTION

- [Superflow Type A: Network Scan](#) | 144
- [Superflow Type B: Distributed Denial of Service \(DDoS\)](#) | 144
- [Superflow Type C: Port Scan](#) | 145

JSA analyzes individual flows to look for indicators that common attack vectors are being used on your network. When the number of flows that match the criteria reaches a specified number, JSA groups the individual flows into a *superflow*. Each superflow counts as only one flow against the flows per minute (FPM) license, regardless of how many flow records are bundled within it.

You can configure the flow threshold to control the number of unique flows that must match the criteria before a superflow is created. Given a threshold of 100, the first 99 flows are sent as normal flow records. The 100th flow and subsequent matching flows are included in the superflow record. JSA continues to report on the superflow every minute, until one full interval passes with no matching traffic. Because a superflow is kept alive even if only one flow record matches, some superflows might appear to be small.








TIP: If you do not want JSA to create superflows, in the Flow Processor configuration settings, change the **Create Super Flows** setting to **No**.

Superflow Type A: Network Scan

A network scan attempts to discover all of the active hosts on your network and map the hosts to an IP address.

JSA looks for flows where one host sends data to many hosts, and flags this type of activity as a Type A superflow. This unidirectional flow is an aggregate of all of the flows that have the same source IP with different destination IPs, but the following parameters are the same:

- Protocol
- Source bytes-to-packets ratio
- Source IP address
- Destination port (TCP and UDP flows only)
- TCP flags (TCP flows only)
- ICMP type and code (ICMP flows only)

Flow Type ▼	First Packet Time	Source IP	Source Port	Destination IP	Destination Port	Protocol
	8 Apr 2016, 06:32:26	10.91.229.20	N/A	Multiple (9)	N/A	icmp_ip
	8 Apr 2016, 06:32:25	10.91.229.20	N/A	Multiple (5)	N/A	icmp_ip
	8 Apr 2016, 06:32:26	10.104.8.249	Multiple (6)	Multiple (6)	47808	udp_ip
	8 Apr 2016, 06:32:49	10.91.229.20	N/A	10.104.8.58	N/A	icmp_ip
	8 Apr 2016, 06:32:36	10.112.38.51	Multiple (13)	Multiple (13)	53213	udp_ip
	8 Apr 2016, 06:32:59	10.106.168.37	Multiple (3)	Multiple (3)	8612	udp_ip
	8 Apr 2016, 06:32:57	10.106.168.37	Multiple (2)	Multiple (2)	8612	udp_ip

Superflow Type B: Distributed Denial of Service (DDoS)

A DDoS attack occurs when multiple systems flood the bandwidth or resources of a targeted system.

JSA looks for flows where many hosts send data to one destination host, and flags this activity as a Type B superflow. This unidirectional flow is an aggregate of all of the flows that have different source IPs with a single destination IP, but the following parameters are the same:

- Protocol
- Source bytes-to-packets ratio
- Destination IP address
- Destination port (TCP and UDP flows only)

- TCP flags (TCP flows only)
- ICMP type and code (ICMP flows only)

Flow Type ▼	First Packet Time	Source IP	Source Port	Destination IP	Destination Port	Protocol
B	8 Apr 2016, 06:32:34	Multiple (64)	Multiple (64)	10.92.229.21	137	udp_ip
B	8 Apr 2016, 06:32:32	Multiple (68)	Multiple (68)	10.91.229.21	137	udp_ip
B	8 Apr 2016, 06:32:54	Multiple (3)	Multiple (3)	10.92.229.20	137	udp_ip
B	8 Apr 2016, 06:32:54	Multiple (4)	Multiple (4)	10.91.229.20	137	udp_ip
B	8 Apr 2016, 06:32:35	Multiple (53)	Multiple (53)	10.92.229.21	137	udp_ip
B	8 Apr 2016, 06:32:32	Multiple (52)	Multiple (52)	10.91.229.21	137	udp_ip
B	8 Apr 2016, 06:32:36	Multiple (27)	Multiple (27)	10.91.0.60	80	tcp_ip
B	8 Apr 2016, 06:33:20	Multiple (3)	Multiple (3)	10.92.229.21	137	udp_ip

Superflow Type C: Port Scan

A port scan attempts to identify the ports that are in use by a specific host on your network.

JSA looks for flows that have one source IP and one destination IP, but many ports. This unidirectional flow is an aggregate of all non-ICMP flows that have different source or destination ports, but the following parameters are the same:

- Protocol
- Source IP address
- Destination IP address
- Source bytes-to-packets ratio
- TCP flags (TCP flows only)

Flow Type ▼	First Packet Time	Source IP	Source Port	Destination IP	Destination Port	Protocol
C	8 Apr 2016, 06:32...	10.91.229.20	Multiple (2)	10.103.24.38	Multiple (2)	tcp_ip
C	8 Apr 2016, 06:32...	10.89.2.34	Multiple (2)	10.89.160.4	Multiple (2)	tcp_ip
C	8 Apr 2016, 06:32...	10.89.2.34	Multiple (3)	10.89.160.1	Multiple (3)	tcp_ip
C	8 Apr 2016, 06:32...	10.89.162.1	Multiple (12)	10.89.2.11	Multiple (12)	udp_ip
C	8 Apr 2016, 06:33...	10.89.2.24	Multiple (5)	10.202.32.116	Multiple (5)	tcp_ip
C	8 Apr 2016, 06:33...	10.89.2.34	Multiple (5)	10.89.160.2	Multiple (5)	tcp_ip

RELATED DOCUMENTATION

[Deduplication](#) | 146

Deduplication

IN THIS SECTION

- [Configuring Flow Deduplication | 146](#)

If your JSA deployment includes multiple Flow Processors that provide data to a Flow Processor, you can configure flow deduplication to remove duplicate flows. The flow deduplication process ensures that flow records with the same unique identifier are reported only once.

Before discarding the duplicate flow, JSA merges the following fields with the original flow record so that you can tell where the duplicate record was observed on the network.

- Flow source ID
- Flow source name
- Source ASN
- Destination ASN
- Input IF index
- Output IF index

Configuring Flow Deduplication

You can configure flow deduplication to take place at either the Flow Processor.

When your Flow Processor receives data from two different flow sources that are monitoring the same part of the network, configure deduplication at the Flow Processor appliance. For example, your Flow Processor might receive data from a 3rd party flow exporter as well as its own network TAP.

When you have two different Flow Processors that monitor overlapping parts of the network, the individual flow processors are unaware of the deduplication. In the scenario, you configure deduplication at the Flow Processor appliance.

You can set the flow deduplication method in the configuration settings for the appliance. For more information, see ["Configuring a Flow Processor" on page 158](#).

RELATED DOCUMENTATION

[View Flow Data on the Network Activity Tab | 147](#)

[Application Identification | 139](#)

[Superflows | 143](#)

View Flow Data on the Network Activity Tab

IN THIS SECTION

- [Viewing Normalized Flow Data | 148](#)
- [Viewing Streaming Flows | 148](#)
- [Viewing Grouped Flows | 149](#)
- [Viewing AWS Flow Log Data | 149](#)
- [Viewing flows that contain MPLS fields | 152](#)
- [Exporting Flows | 154](#)

The default view on the **Network Activity** tab is a stream of real-time events. As the hosts continue to communicate, the flow is updated. The entire communication session is represented by multiple flow records that have the same **First Packet Time**, but with **Last Packet Time** values that increment through time.

The **View** list contains options to also view events for specified time periods. After you choose a specified time period from the **View** list, you can modify the displayed time period by changing the date and time values in the Start Time and End Time fields.

Viewing Normalized Flow Data

Normalization involves preparing flow data to display readable information. You view normalized flow data on the **Network Activity** tab.

1. Click the **Network Activity** tab.
2. In the center of the window, from the **Display** list box, choose the display view that you want to see.

The **Normalized (With IPv6 Columns)** display shows source and destination IPv6 addresses for IPv6 flows.

3. From the **View** list box, select the time frame that you want to display.

- a. To stream flows as they are received, select **Real Time**.

You can click the **Pause** icon to pause streaming.

- b. To reduce the number of flows that you see, select a time frame.

If you select a time frame to display, a time series chart is displayed. Click **Hide Charts** to remove it. If you use the Mozilla Firefox web browser and you have an ad blocker extension installed, the charts do not appear.

For more information about using time series charts, see "[Time Series Chart Overview](#)" on page 206.

4. To view the more information about a particular flow, double-click it to open the **Flow Information** window.

Viewing Streaming Flows

Streaming mode enables you to view flow data entering your system in real time. This mode provides you with a real-time view of your current flow activity by displaying the last 50 flows.

If you apply any filters on the Network Activity tab or in your search criteria before you enable streaming mode, the filters are maintained in streaming mode. However, streaming mode does not support searches that include grouped flows. If you enable streaming mode on grouped flows or grouped search criteria, the Network Activity tab displays the normalized flows.

1. Click the **Network Activity** tab.
2. From the View list box, select **Real Time (streaming)**.

3. Optional. Pause or play the streaming flows. When streaming is paused, the last 1,000 flows are displayed.

NOTE: When you are streaming flows, the status bar displays the average number of results that are received per second. This display is the number of results that the Console successfully received from the Flow processors. If this number is greater than 40 results per second, only 40 results are displayed. The remainder is accumulated in the result buffer. To view more status information, hover over the status bar.

When flows are not streaming, the status bar displays the number of search results that are currently displayed and the amount of time that is required to process the search results.

Viewing Grouped Flows

View flows that are grouped by various options.

The **Display** list box is not displayed in streaming mode because streaming mode does not support grouped flows. If you entered streaming mode by using non-grouped search criteria, this option is displayed.

After you select an option from the **Display** list box, the column layout of the data depends on the chosen group option. Each row in the flows table represents a flow group.

1. Click the **Network Activity** tab.
2. From the **View** list box, select the time frame that you want to display.
3. From the **Display** list box, choose which parameter you want to group flows on.
4. To view the **List of Flows** page for a group, double-click the flow group that you want to investigate.

The **List of Flows** page does not retain chart configurations that you might define on the **Network Activity** tab.

5. To view the details of a flow, double-click the flow that you want to investigate.

Viewing AWS Flow Log Data

Flow logs that are received through Amazon Web Service (AWS) integrations include extra properties in the flow information.

In addition to the standard normalized flow properties, the following properties are shown for AWS flow logs:

- Interface name (available for all IPFIX flows that send this field)
- Region (available for all IPFIX flows that send this field)
- Firewall Name (available for all IPFIX flows that send this field)
- Firewall Event (enumerated, available for all IPFIX flows that send this field)
- AWS Action (enumerated)
- AWS Log Status (enumerated)
- AWS Account ID
- VPC ID
- Subnet ID
- Instance ID

The following table shows the string description for each of the enumerated fields:

Table 24: AWS Enumerated Strings

Enumerated field	String description
Firewall Event	<p>The numerical values for the Firewall Event field map to the following descriptions:</p> <ul style="list-style-type: none"> • 0 = Ignore • 1 = Flow Created • 2 = Flow Deleted • 3 = Flow Denied • 4 = Flow Alert • 5 = Flow Update

Table 24: AWS Enumerated Strings (Continued)

Enumerated field	String description
AWS Action	<p>The numerical values for the AWS Action field map to the following descriptions:</p> <ul style="list-style-type: none"> • 0 = N/A • 1 = Accept • 2 = Reject
AWS Log Status	<p>The numerical values for the AWS Log Status field map to the following descriptions:</p> <ul style="list-style-type: none"> • 0 = N/A • 1 = OK • 2 = No Data • 3 = Skip Data

To include the description for the enumerated property in your query results, you must include the LOOKUP function in your AQL search string.

1. Click the **Network Activity** tab.
2. In the **Advanced Search** box, build the AQL query that includes the LOOKUP for the field that you want to include in your search.

The following examples show the LOOKUP statements for each of the enumerated fields in the AWS flow log:

```
LOOKUP('firewall event', "firewall event")
```

```
LOOKUP('aws action', "aws action")
```

```
LOOKUP('aws log status', "aws log status")
```

For example, the following query uses a LOOKUP in the WHERE clause and groups the accepted flows by application:

```
SELECT APPLICATIONNAME(applicationid), count(*) as NumFlows FROM flows WHERE LOOKUP('aws action', "aws action") == 'Accept' GROUP BY applicationid ORDER BY NumFlows DESC
```

In this example, the query uses a LOOKUP in the SELECT clause to show the number of accepted flows vs rejected flows in the AWS environment:


```
SELECT LOOKUP('aws action', "aws action"), count(*) as NumFlows FROM flows WHERE "aws action" > 0 GROUP BY
"aws action" ORDER BY NumFlows DESC LAST 7 DAYS
```

Viewing flows that contain MPLS fields

Internet Protocol Flow Information Export (IPFIX) is a common protocol that allows exporting of flow information from network devices. Multiprotocol Label Switching (MPLS) is a routing technique that runs on any protocol.

With MPLS support for IPFIX flow records in Flow Processor, you can filter and search for IPFIX flows in JSA that contain MPLS fields and write rules based on the values of these MPLS fields.

For example, an IPFIX flow is exported from a switch on a network that uses MPLS. The IPFIX flow that is exported from the router contains information about the MPLS stack, which is now saved as part of the flow in JSA. The MPLS stack can contain up to 10 layers where each layer shows information about the flow routing. These MPLS fields are included in rules, searches, and filters, and can be viewed in the **Flow Details** window.

Filter on MPLS Fields

Use the **Add Filter** option on the **Network Activity** tab to filter on MPLS fields.

Search for MPLS Fields

Use the **Advanced Search** option on the **Network Activity** tab to search for MPLS fields.

View Information about MPLS Fields

You can view information about MPLS fields by double-clicking a flow in the **Flow Details** window on the **Network Activity** tab.

IPFIX MPLS Information Elements

The following table describes the IPFIX MPLS information elements that are supported. All of these elements have Private Enterprise Number (PEN): 0.

Field	Element ID
mplsTopLabelType	46

(Continued)

Field	Element ID
mplsTopLabelIPv4Address	47
mplsTopLabelStackSection	70
mplsLabelStackSection2	71
mplsLabelStackSection3	72
mplsLabelStackSection4	73
mplsLabelStackSection5	74
mplsLabelStackSection6	75
mplsLabelStackSection7	76
mplsLabelStackSection8	77
mplsLabelStackSection9	78
mplsLabelStackSection10	79
mplsVpnRouteDistinguisher	90
mplsTopLabelPrefixLength	91
mplsTopLabelIPv6Address	140
mplsPayloadLength	194

(Continued)

Field	Element ID
mplsTopLabelTTL	200
mplsLabelStackLength	201
mplsLabelStackDepth	202
mplsTopLabelExp	203
postMplsTopLabelExp	237
pseudoWireType	250
pseudoWireControlWord	251
mplsLabelStackSection	316
mplsPayloadPacketSection	317
sectionOffset	409
sectionExportedOctets	410

Exporting Flows

You can export flows in Extensible Markup Language (XML) or Comma Separated Values (CSV) format. The length of time that is required to export your data depends on the number of parameters specified.

1. Click the **Network Activity** tab.
2. Optional. If you are viewing flows in streaming mode, click the **Pause** icon to pause streaming.
3. From the **Actions** list box, select one of the following options:

- **Export to XML >Visible Columns** - Select this option to export only the columns that are visible on the Log Activity tab. This is the recommended option.
- **Export to XML >Full Export (All Columns)** - Select this option to export all flow parameters. A full export can take an extended period of time to complete.
- **Export to CSV >Visible Columns** - Select this option to export only the columns that are visible on the Log Activity tab. This is the recommended option.
- **Export to CSV >Full Export (All Columns)** - Select this option to export all flow parameters. A full export can take an extended period of time to complete.

4. If you want to resume your activities, click **Notify When Done**.

When the export is complete, you receive notification that the export is complete. If you did not select the **Notify When Done** icon, the **Status** window is displayed.

RELATED DOCUMENTATION

[VLAN Fields | 155](#)

[Superflows | 143](#)

[Deduplication | 146](#)

VLAN Fields

IN THIS SECTION

- [Assign Domains and Tenants to Flows with VLAN Information | 157](#)

JSA retains Virtual Local Area Network (VLAN) information that is exported in external flow records.

VLAN information can be found in flows that are received from IPFIX, NetFlow V9, sFlow V5, or J-Flow V9. It can also be viewed in internal flows, such as those that are received by Network Interface Cards, or a dedicated QRadar Network Insights appliance.

You can use the VLAN information in searches, filters, and custom rules.

Supported VLAN fields

The following VLAN fields are supported for IPFIX, Netflow version 9, and J-Flow flow records:

- vlanId
- postVlanId
- dot1qVlanId
- dot1qPriority
- dot1qCustomerVlanId
- dot1qCustomerPriority
- postDot1qVlanId
- postDotqCustomerVlanId

The following VLAN fields are supported for raw packets and sFlow version 5. QRadar Network Insights also supports these fields.

- dot1qVlanId
- dot1qPriority
- dot1qCustomerVlanId
- dot1qCustomerPriority
- dot1qDEI
- dot1qCustomerDEI

VLAN separation

All flows that contain VLAN information are tagged with the following specific fields to ensure that the network traffic from different groups of users is kept separate:

- Enterprise VLAN ID
- Customer VLAN ID

The values that appear in either field depend on the VLAN configuration:

- If the VLANs are stacked, both the **Enterprise VLAN ID** and **Customer VLAN ID** fields are populated with the relevant VLAN values.
- If the VLANs are nonstacked, the **Enterprise VLAN ID** property is set to 0, and the **Customer VLAN ID** shows the relevant VLAN value.

- If the inner VLAN is set to 0, the **Enterprise VLAN ID** shows the VLAN value and the **Customer VLAN ID** shows as 0.

Assign Domains and Tenants to Flows with VLAN Information

With domain management support for VLAN flows, you can define your domains in JSA based on the VLAN information in your network.

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. You can also filter and query the domains for the VLAN-based domain.

You can assign tenants to domain definitions to achieve multi-tenancy. The VLAN-based domain definitions enable multi-tenancy across different VLANs, if required.

For example, two domain definitions are created and mapped to two network tenants:

- For tenant ABC, traffic is sent on Enterprise VLAN ID = 0, and Customer VLAN ID = 10.
- For tenant DEF, traffic is sent on Enterprise VLAN ID = 0, and Customer VLAN ID = 20.

The first domain definition is created for tenant ABC, which contains a flow VLAN definition of Enterprise VLAN ID = 0 and customer VLAN ID = 10.

A second domain definition is created for tenant DEF, which contains a flow VLAN definition of Enterprise VLAN ID = 0 and customer VLAN ID = 20.

Incoming flows with Enterprise VLAN ID and Customer VLAN ID fields set to 0 and 10 are viewed only by *tenant ABC*. Similarly, incoming flows with Enterprise VLAN ID and Customer VLAN ID fields of 0 and 20 are only viewed by *tenant DEF*. This reflects the traffic ownership for each tenant in the network.

RELATED DOCUMENTATION

[Superflows | 143](#)

[Deduplication | 146](#)

[View Flow Data on the Network Activity Tab | 147](#)

Configuring a Flow Processor

By changing the Flow Processor configuration settings, you can manage the way that JSA collects and processes flows that are received from the device.

The following table describes the Flow Processor configuration parameters:

Table 25: Flow Processor Configuration Parameters

Parameter	Description
Maximum Content Capture	Specify the maximum amount of data (bytes per packet) that you want the Flow Processor to capture and retain in the flow payload.
Maximum Data Capture/Packet	Specify the maximum amount of data (bytes per packet) that you want the Flow Processor to analyze.
Flow buffer size	Specify the maximum number of flows that can be buffered in memory.
Maximum Number of Flows	Specify the maximum number of flows that you want to send from the Flow Processor to an Event Collector.
Alias Autodetection	Set to Yes to allow JSA to auto-detect flow sources. With auto-detection turned on, JSA can automatically create flow source aliases for external flow sources, such as routers.
Remove duplicate flows	Set this to Yes if you want the Flow Processor to remove duplicate flows. If you have asymmetric traffic in your network, set this parameter to No .
Verify NetFlow Sequence Numbers	Set this to Yes if you want the Flow Processor to check the incoming NetFlow sequence numbers to ensure that all packets are present and in order. JSA displays a notification if a packet is missing or received in incorrect sequence.

Table 25: Flow Processor Configuration Parameters *(Continued)*

Parameter	Description
External Flow De-duplication method	<p>Choose the method that you want to use to remove duplicate external flows.</p> <ul style="list-style-type: none"> • Select Source to compare the originating flow sources. <p>This method compares the IP address of the device that exported the current external flow record to the IP address of the device that exported the first external record of the flow. If the IP addresses do not match, the current external flow record is discarded.</p> <ul style="list-style-type: none"> • Select Record to compare the individual external flow records. <p>This method logs a list of every external flow record that is detected by a device, and compares each subsequent record to that list. If the current record is found in the list, the record is discarded.</p> <p>If you choose this method, you must also set the External flow record comparison mask parameter.</p>
Flow Carry-over Window	<p>Specify the number of seconds that the Flow Processor process holds one-sided flows. The default setting is 6 seconds.</p> <p>This setting allows time for JSA to receive the flow response. Flows that fall within the carry-over window are not sent until the next reporting interval.</p>

Table 25: Flow Processor Configuration Parameters *(Continued)*

Parameter	Description
External flow record comparison mask	<p>Specify the method to use to compare external flow records.</p> <p>This parameter is valid only if you chose Record as the method to use for external flow De-duplication.</p> <p>You can choose which flow record fields are to be used when comparing external flow records:</p> <ul style="list-style-type: none"> • D (Direction) • B (ByteCount) • P (PacketCount) <p>You can combine the flow record fields to include the following combinations:</p> <ul style="list-style-type: none"> • The DBP option uses direction, byte count, and packet count. • The XBP option uses byte count and packet count. • The DXP option uses direction and packet count. • The DBX option uses direction and byte count. • The DXX option uses direction. • The XBX option uses byte count. • The XXP option uses packet count.
Create Super Flows	Set this to Yes if you want JSA to group flows that have similar properties into one flow record
Type A Superflows (Network Scan)	Specify the threshold to be reached before JSA creates a Type A (one to many) superflow.
Type B Superflows (DDos)	Specify the threshold to be reached before JSA creates a Type B (many to one) superflow.

Table 25: Flow Processor Configuration Parameters (Continued)

Parameter	Description
Type C Superflows (Port Scan)	Specify the threshold to be reached before JSA creates a Type C (one to one) superflow.
Recombine Asymmetric flows	Set this to Yes if you want JSA to recombine asymmetric flows.
Ignore Asymmetric Superflows	Set this to Yes if you want JSA to create superflows when asymmetric flows are enabled.
Use Common Destination Port	Set this to Yes if you want JSA to determine whether to reverse the flow direction.

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**, and select the Flow Processor that you want to configure.
4. On the **Deployment Actions** menu, click **Edit Host**.
5. Click the gear icon next to **Component Management**.
6. Edit the configuration options and click **Save**.
7. Repeat the configuration steps for each Flow Processor in your deployment.
8. Close the **System and License Management** window.
9. Deploy your changes.

This will restart the Flow Processor process on every managed host that you modified.

RELATED DOCUMENTATION

[Tuning False Positive Events from Creating Offenses | 162](#)

[View Flow Data on the Network Activity Tab | 147](#)

[VLAN Fields | 155](#)

Tuning False Positive Events from Creating Offenses

You can prevent false positive flows from creating offenses. You can tune false positive flows from the flow list or flow details page.

You must have appropriate permissions for creating customized rules to tune false positives.

1. Click the **Network Activity** tab.
2. Optional. If you are viewing flows in streaming mode, click the **Pause** icon to pause streaming.
3. Select the flow that you want to tune.
4. Click **False Positive**.
5. In the Event/Flow Property pane on the **False Positive** page, select one of the following options:
 - Event/Flow(s) with a specific QID of <Event>
 - Any Event/Flow(s) with a low-level category of <Event>
 - Any Event/Flow(s) with a high-level category of <Event>
6. In the Traffic Direction pane, select one of the following options:
 - <Source IP Address> to <Destination IP Address>
 - <Source IP Address> to any Destination
 - Any Source to <Destination IP Address>
 - Any Source to any Destination
7. Click **Tune**.

RELATED DOCUMENTATION

[Exporting Flows](#)

[Network Activity Monitoring](#) | 119

8

CHAPTER

Asset Management

[Asset Management](#) | 164

[Sources Of Asset Data](#) | 166

[Incoming Asset Data Workflow](#) | 168

[Updates to Asset Data](#) | 170

[Identification Of Asset Growth Deviations](#) | 175

[Asset Blacklists and Whitelists](#) | 180

[Asset Profiles](#) | 183

Asset Management

IN THIS SECTION

- [Asset Data | 164](#)
- [Asset Limits | 164](#)
- [Asset Profiles | 166](#)
- [Collecting Asset Data | 166](#)

Collecting and viewing asset data helps you to identify threats and vulnerabilities. An accurate asset database makes it easier to connect offenses that are triggered in your system to physical or virtual assets in your network.

NOTE: JSA Log Manager only tracks 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 7.

Asset Data

An *asset* is any network endpoint that sends or receives data across your network infrastructure. For example, notebooks, servers, virtual machines, and handheld devices are all assets. Every asset in the asset database is assigned a unique identifier so that it can be distinguished from other asset records.

Detecting devices is also useful in building a data set of historical information about the asset. Tracking asset information as it changes helps you monitor asset usage across your network.

Asset Limits

The asset database has a limited capacity. When the asset limit for your hardware is reached, you cannot create any new assets until sufficient space is available in the database. The following table describes the asset limits for each hardware type:

Table 26: Asset Limits for Hardware

Hardware Type	Asset Limit for Console only	Asset Limit for Console with Managed Host
JSA3800	60,000	60,000
JSA7500	300,000	700,000
JSA5800	500,000	1,000,000
JSA7800	500,000	1,000,000

The following table describes the asset limits for Virtual Machines (VM):

Table 27: Asset Limits for Virtual Machines

Virtual Machine (VM) Type	Asset Limit for Console only	Asset Limit for Console with Managed Host
VM with minimum requirements	60,000	60,000
VM with 16-core CPU and 48GB memory	200,000	600,000
VM with 32-core CPU and 128GB memory	500,000	1,000,000
VM with 49-core CPU and 129GB memory	500,000	1,000,000
VM or Software only install	500,000	1,000,000

Asset Profiles

An *asset profile* is a collection of all information that JSA collected over time about a specific asset. The profile includes information about the services that are running on the asset and any identity information that is known.

JSA automatically creates asset profiles from identity events and bidirectional flow data or, if they are configured, vulnerability assessment scans. The data is correlated through a process that is called *asset reconciliation* and the profile is updated as new information comes into JSA. The asset name is derived from the information in the asset update in the following order of precedence:

- Given name
- NETBios host name
- DNS host name
- IP address

Collecting Asset Data

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.

Sources Of Asset Data

IN THIS SECTION

- [Domain-aware Asset Data | 167](#)

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 | 168](#)

[Updates to Asset Data | 170](#)

[Identification Of Asset Growth Deviations | 175](#)

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 5: 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 | 170](#)

[Identification Of Asset Growth Deviations | 175](#)

[Asset Blacklists and Whitelists | 180](#)

Updates to Asset Data

IN THIS SECTION

- [Identity Information | 171](#)
- [Asset Reconciliation Exclusion Rules | 171](#)
- [Example: Asset Exclusion Rules That Are Tuned to Exclude IP Addresses from the Blacklist | 173](#)
- [Asset Merging | 174](#)

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 28: 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

Table 28: Rule Tests and Responses (Continued)

Scenario	Rule response
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
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.

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.

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 | 175](#)

[Asset Blacklists and Whitelists | 180](#)

[Asset Profiles | 183](#)

Identification Of Asset Growth Deviations

IN THIS SECTION

- [DHCP Server Example Of Unnatural Asset Growth in an Asset Profile | 176](#)
- [Threshold Settings | 176](#)
- [System Notifications That Indicate Asset Growth Deviations | 176](#)
- [Example: How Configuration Errors for Log Source Extensions Can Cause Asset Growth Deviations | 177](#)
- [Troubleshooting Asset Profiles That Exceed the Normal Size Threshold | 177](#)
- [New Asset Data is Added to the Asset Blocklists | 178](#)

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][9.21.118.83/- -] [-/- -]
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 allowlist prevents it from inadvertently reappearing on the blocklist.

RELATED DOCUMENTATION

[Asset Blacklists and Whitelists | 180](#)

[Asset Profiles | 183](#)

[Updates to Asset Data | 170](#)

Asset Blacklists and Whitelists

IN THIS SECTION

- [Asset Blacklists | 180](#)
- [Asset Allowlists | 181](#)

JSA uses a group of asset reconciliation rules to determine if asset data is trustworthy. When asset data is questionable, JSA uses asset blacklists and whitelists to determine whether to update the asset profiles with the asset data.

An asset blacklist is a collection of data that JSA considers untrustworthy. Data in the asset blacklist is likely to contribute to asset growth deviations and JSA prevents the data from being added to the asset database.

An asset blacklist 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 blacklist match, it checks the whitelist to see whether the value exists. If the asset update matches data that is on the whitelist, the change is reconciled and the asset is updated. Whitelisted asset data is applied globally for all domains.

Your JSA administrator can modify the asset blacklist and whitelist data to prevent future asset growth deviations.

Asset Blacklists

An *asset blacklist* is a collection of data that JSA considers untrustworthy based on the asset reconciliation exclusion rules. Data in the asset blacklist 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 blacklists. Blacklisted 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 blacklist, 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 29: Reference Collection Names for Asset Blacklist 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.

Your JSA administrator can modify the blacklist 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 blacklists.

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 blacklist 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.

Your JSA administrator can modify the allowlists 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 blacklists 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 blacklist. 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 blacklist 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 30: 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.

RELATED DOCUMENTATION

[Asset Profiles | 183](#)

[Updates to Asset Data | 170](#)

[Identification Of Asset Growth Deviations | 175](#)

Asset Profiles

IN THIS SECTION

- [Vulnerabilities | 183](#)
- [Assets Tab Overview | 184](#)
- [Viewing an Asset Profile | 185](#)
- [Adding or Editing an Asset Profile | 187](#)
- [Searching Asset Profiles | 192](#)
- [Saving Asset Search Criteria | 194](#)
- [Asset Search Groups | 195](#)
- [Asset Profile Management Tasks | 198](#)
- [Research Asset Vulnerabilities | 200](#)

Asset profiles provide information about each known asset in your network, including what services are running on each asset.

Asset profile information is used for correlation purposes to help reduce false positives. For example, if a source attempts to exploit a specific service running on an asset, then JSA determines if the asset is vulnerable to this attack by correlating the attack to the asset profile.

Asset profiles are automatically discovered if you have flow data or vulnerability assessment (VA) scans configured. For flow data to populate asset profiles, bidirectional flows are required. Asset profiles can also be automatically created from identity events. For more information about VA, see the *Juniper Secure Analytics Managing Vulnerability Assessment Guide*.

For more information about flow sources, see the *Juniper Secure Analytics Administration Guide*.

Vulnerabilities

You can use JSA Vulnerability Manager and third-party scanners to identify vulnerabilities.

Third-party scanners identify and report discovered vulnerabilities using external references, such as the Open Source Vulnerability Database (OSVDB), National Vulnerability Database (NVDB), and Critical Watch. Examples of third-party scanners include QualysGuard and nCircle ip360. The OSVDB assigns a

unique reference identifier (OSVDB ID) to each vulnerability. External references assign a unique reference identifier to each vulnerability. Examples of external data reference IDs include Common Vulnerability and Exposures (CVE) ID or Bugtraq ID. For more information on scanners and vulnerability assessment, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

JSA Vulnerability Manager is a component that you can purchase separately and enable using a license key. JSA Vulnerability Manager is a network scanning platform that provides awareness of the vulnerabilities that exist within the applications, systems, or devices on your network. After scans identify vulnerabilities, you can search and review vulnerability data, remediate vulnerabilities, and rerun scans to evaluate the new level of risk.

When JSA Vulnerability Manager is enabled, you can perform vulnerability assessment tasks on the **Vulnerabilities** tab. From the **Assets** tab, you can run scans on selected assets.

For more information, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

Assets Tab Overview

The **Assets** tab provides you with a workspace from which you can manage your network assets and investigate an asset's vulnerabilities, ports, applications, history, and other associations.

Using the **Assets** tab, you can:

- View all the discovered assets.
- Manually add asset profiles.
- Search for specific assets.
- View information about discovered assets.
- Edit asset profiles for manually added or discovered assets.
- Tune false positive vulnerabilities.
- Import assets.
- Print or export asset profiles.
- Discover assets.
- Configure and manage third-party vulnerability scanning.
- Start JSA Vulnerability Manager scans.

For information about the Server Discovery option in the navigation pane, see the *Juniper Secure Analytics Administration Guide*.

For more information about the VA Scan option in the navigation pane, see the *Juniper Secure Analytics Risk Manager User Guide*.

Viewing an Asset Profile

From the asset list on the **Assets** tab, you can select and view an asset profile. An asset profile provides information about each profile.

Asset profile information is automatically discovered through Server Discovery or manually configured. You can edit automatically generated asset profile information.

The **Asset Profile** page provides the information about the asset that is organized into several panes. To view a pane, you can click the arrow (>) on the pane to view more detail or select the pane from the **Display** list box on the toolbar.

The **Asset Profile** page toolbar provides the following functions:

Table 31: Asset Profile Page Toolbar Functions

Options	Description
Return to Asset List	Click this option to return to the asset list.
Display	From the list box, you can select the pane that you want to view on the Asset Profile pane. The Asset Summary and Network Interface Summary panes are always displayed.
Edit Asset	Click this option to edit the Asset Profile. See "Adding or Editing an Asset Profile" on page 187.
View by Network	If this asset is associated with an offense, this option will allow you to view the list of networks that are associated with this asset. When you click View By Network , the List of Networks window is displayed.
View Source Summary	If this asset is the source of an offense, this option will allow you to view source summary information. When you click View Source Summary , the List of Offenses window is displayed.

Table 31: Asset Profile Page Toolbar Functions (Continued)

Options	Description
View Destination Summary	<p>If this asset is the destination of an offense, this option will allow you to view destination summary information.</p> <p>When you click View Destination Summary, the List of Destinations window is displayed.</p>
History	<p>Click History to view event history information for this asset. When you click the History icon, the Event Search window is displayed, pre-populated with event search criteria:</p> <p>You can customize the search parameters, if required. Click Search to view the event history information.</p>
Applications	<p>Click Applications to view application information for this asset. When you click the Applications icon, the Flow Search window is displayed, pre-populated with event search criteria.</p> <p>You can customize the search parameters, if required. Click Search to view the application information.</p>
Search Connections	<p>Click Search Connections to search for connections. The Connection Search window is displayed.</p> <p>This option is only displayed when JSA Risk Manager is been purchased and licensed. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i>.</p>
View Topology	<p>Click View Topology to further investigate the asset. The Current Topology window is displayed.</p> <p>This option is only displayed when JSA Risk Manager is been purchased and licensed. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i>.</p>
Actions	<p>From the Actions list, select Vulnerability History.</p> <p>This option is only displayed when JSA Risk Manager is been purchased and licensed. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i>.</p>

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**
3. Double-click the asset that you want to view.
4. Use the options on the toolbar to display the various panes of asset profile information. See ["Adding or Editing an Asset Profile" on page 187](#).
5. To research the associated vulnerabilities, click each vulnerability in the Vulnerabilities pane. See Table 10-10
6. If required, edit the asset profile. See ["Adding or Editing an Asset Profile" on page 187](#).
7. Click **Return to Assets List** to select and view another asset, if required.

Adding or Editing an Asset Profile

Asset profiles are automatically discovered and added; however, you might be required to manually add a profile.

When assets are discovered using the Server Discovery option, some asset profile details are automatically populated. You can manually add information to the asset profile and you can edit certain parameters.

You can only edit the parameters that were manually entered. Parameters that were system generated are displayed in italics and are not editable. You can delete system generated parameters, if required.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Choose one of the following options:
 - To add an asset, click **Add Asset** and type the IP address or CIDR range of the asset in the **New IP Address** field.
 - To edit an asset, double-click the asset that you want to view and click **Edit Asset** .
4. Configure the parameters in the MAC & IP Address pane. Configure one or more of the following options:
 - Click the **New MAC Address** icon and type a MAC Address in the dialog box.
 - Click the **New IP Address** icon and type an IP address in the dialog box.

- If **Unknown NIC** is listed, you can select this item, click the **Edit** icon, and type a new MAC address in the dialog box.
- Select a MAC or IP address from the list, click the **Edit** icon, and type a new MAC address in the dialog box.
- Select a MAC or IP address from the list and click the **Remove** icon.

5. Configure the parameters in the Names & Description pane. Configure one or more of the following options:

Parameter	Description
DNS	Choose one of the following options: <ul style="list-style-type: none"> • Type a DNS name and click Add. • Select a DNS name from the list and click Edit. • Select a DNS name from the list and click Remove.
NetBIOS	Choose one of the following options: <ul style="list-style-type: none"> • Type a NetBIOS name and click Add. • Select a NetBIOS name from the list and click Edit. • Select a NetBIOS name from the list and click Remove.
Given Name	Type a name for this asset profile.
Location	Type a location for this asset profile.
Description	Type a description for the asset profile.
Wireless AP	Type the wireless Access Point (AP) for this asset profile.
Wireless SSID	Type the wireless Service Set Identifier (SSID) for this asset profile.
Switch ID	Type the switch ID for this asset profile.

(Continued)

Parameter	Description
Switch Port ID	Type the switch port ID for this asset profile.

6. Configure the parameters in the Operating System pane:
 - a. From the **Vendor** list box, select an operating system vendor.
 - b. From the **Product** list box, select the operating system for the asset profile.
 - c. From the **Version** list box, select the version for the selected operating system.
 - d. Click the **Add** icon.
 - e. From the **Override** list box, select one of the following options:
 - **Until Next Scan** Select this option to specify that the scanner provides operating system information and the information can be temporarily edited. If you edit the operating system parameters, the scanner restores the information at its next scan.
 - **Forever** Select this option to specify that you want to manually enter operating system information and disable the scanner from updating the information.
 - f. Select an operating system from the list.
 - g. Select an operating system and click the **Toggle Override** icon.
7. Configure the parameters in the CVSS & Weight pane. Configure one or more of the following options:

Parameter	Description
Collateral Damage Potential	<p>Configure this parameter to indicate the potential for loss of life or physical assets through damage or theft of this asset. You can also use this parameter to indicate potential for economic loss of productivity or revenue. Increased collateral damage potential increases the calculated value in the CVSS Score parameter.</p> <p>From the Collateral Damage Potential list box, select one of the following options:</p> <ul style="list-style-type: none">• None• Low• Low-medium• Medium-high• High• Not defined <p>When you configure the Collateral Damage Potential parameter, the Weight parameter is automatically updated.</p>
Confidentiality Requirement	<p>Configure this parameter to indicate the impact on confidentiality of a successfully exploited vulnerability on this asset. Increased confidentiality impact increases the calculated value in the CVSS Score parameter.</p> <p>From the Confidentiality Requirement list box, select one of the following options:</p> <ul style="list-style-type: none">• Low• Medium• High• Not defined

(Continued)

Parameter	Description
Availability Requirement	<p>Configure this parameter to indicate the impact to the asset's availability when a vulnerability is successfully exploited. Attacks that consume network bandwidth, processor cycles, or disk space impact the availability of an asset. Increased availability impact increases the calculated value in the CVSS Score parameter.</p> <p>From the Availability Requirement list box, select one of the following options:</p> <ul style="list-style-type: none"> • Low • Medium • High • Not defined
Integrity Requirement	<p>Configure this parameter to indicate the impact to the asset's integrity when a vulnerability is successfully exploited. Integrity refers to the trustworthiness and guaranteed veracity of information. Increased integrity impact increases the calculated value in the CVSS Score parameter.</p> <p>From the Integrity Requirement list box, select one of the following options:</p> <ul style="list-style-type: none"> • Low • Medium • High • Not defined
Weight	<p>From the Weight list box, select a weight for this asset profile. The range is 0 - 10.</p> <p>When you configure the Weight parameter, the Collateral Damage Potential parameter is automatically updated.</p>

8. Configure the parameters in the Owner pane. Choose one or more of the following options:

Parameter	Description
Business Owner	Type the name of the business owner of the asset. An example of a business owner is a department manager. The maximum length is 255 characters.
Business Owner Contact	Type the contact information for the business owner. The maximum length is 255 characters.
Technical Owner	Type the technical owner of the asset. An example of a business owner is the IT manager or director. The maximum length is 255 characters.
Technical Owner Contact	Type the contact information for the technical owner. The maximum length is 255 characters.
Technical User	<p>From the list box, select the username that you want to associate with this asset profile.</p> <p>You can also use this parameter to enable automatic vulnerability remediation for Juniper Secure Analytics Vulnerability Manager. For more information about automatic remediation, see the <i>Juniper Secure Analytics Vulnerability Manager User Guide</i>.</p>

9. Click **Save**.

Searching Asset Profiles

You can configure search parameters to display only the asset profiles you want to investigate from the **Asset** page on the **Assets** tab.

When you access the **Assets** tab, the **Asset** page is displayed populated with all discovered assets in your network. To refine this list, you can configure search parameters to display only the asset profiles you want to investigate.

From the **Asset Search** page, you can manage Asset Search Groups. For more information about Asset Search Groups, see ["Asset Search Groups" on page 195](#).

The search feature will allow you to search host profiles, assets, and identity information. Identity information provides more detail about log sources on your network, including DNS information, user logins, and MAC addresses.

Using the asset search feature, you can search for assets by external data references to determine whether known vulnerabilities exist in your deployment.

For example:

You receive a notification that CVE ID: CVE-2010-000 is being actively used in the field. To verify whether any hosts in your deployment are vulnerable to this exploit, you can select **Vulnerability External Reference** from the list of search parameters, select **CVE**, and then type the following:

2010-000

To view a list of all hosts that are vulnerable to that specific CVE ID.

NOTE: For more information about OSVDB, see <http://osvdb.org/> . For more information about NVDB, see <http://nvd.nist.gov/> .

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. On the toolbar, click **Search >New Search**.
4. Choose one of the following options:
 - To load a previously saved search, go to Step 5.
 - To create a new search, go to Step 6.
5. Select a previously saved search:
 - a. Choose one of the following options:
 - Optional. From the **Group** list box, select the asset search group that you want to display in the **Available Saved Searches** list.
 - From the **Available Saved Searches** list, select the saved search that you want to load.
 - In the **Type Saved Search or Select from List** field, type the name of the search you want to load.
 - b. Click **Load** .
6. In the Search Parameters pane, define your search criteria:
 - a. From the first list box, select the asset parameter that you want to search for. For example, **Hostname**, **Vulnerability Risk Classification**, or **Technical Owner**.
 - b. From the second list box, select the modifier that you want to use for the search.

- c. In the entry field, type specific information that is related to your search parameter.
- d. Click **Add Filter**.
- e. Repeat these steps for each filter that you want to add to the search criteria.

7. Click **Search**.

You can save your asset search criteria. See "[Saving Asset Search Criteria](#)" on page 194.

Saving Asset Search Criteria

On the **Asset** tab, you can save configured search criteria so that you can reuse the criteria. Saved search criteria does not expire.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Perform a search.
4. Click **Save Criteria**.
5. Enter values for the parameters:

Parameter	Description
Enter the name of this search	Type the unique name that you want to assign to this search criteria.
Manage Groups	Click Manage Groups to manage search groups. This option is only displayed if you have administrative permissions.
Assign Search to Group(s)	Select the check box for the group you want to assign this saved search. If you do not select a group, this saved search is assigned to the Other group by default.
Include in my Quick Searches	Select this check box to include this search in your Quick Search list box, which is on the Assets tab toolbar.

(Continued)

Parameter	Description
Set as Default	Select this check box to set this search as your default search when you access the Assets tab.
Share with Everyone	Select this check box to share these search requirements with all users.

Asset Search Groups

Using the **Asset Search Groups** window, you can create and manage asset search groups.

These groups allow you to easily locate saved search criteria on the **Assets** tab.

Viewing Search Groups

Use the **Asset Search Groups** window to view a list group and subgroups.

From the **Asset Search Groups** window, you can view details about each group, including a description and the date the group was last modified.

All saved searches that are not assigned to a group are in the **Other** group.

The **Asset Search Groups** window displays the following parameters for each group:

Table 32: Asset Search Groups Window Toolbar Functions

Function	Description
New Group	To create a new search group, you can click New Group . "Creating a New Search Group" on page 196.
Edit	To edit an existing search group, you can click Edit . "Editing a Search Group" on page 196.
Copy	To copy a saved search to another search group, you can click Copy . "Copying a Saved Search to Another Group" on page 197.

Table 32: Asset Search Groups Window Toolbar Functions (Continued)

Function	Description
Remove	To remove a search group or a saved search from a search group, select the item that you want to remove, and then click Remove . "Removing a Group or a Saved Search from a Group" on page 197.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Select **Search >New Search**.
4. Click on **Manage Groups**.
5. View the search groups.

Creating a New Search Group

On the **Asset Search Groups** window, you can create a new search group.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Select **Search >New Search**.
4. Click **Manage Groups**.
5. Select the folder for the group under which you want to create the new group.
6. Click **New Group**.
7. In the **Name** field, type a unique name for the new group.
8. Optional. In the **Description** field, type a description.
9. Click **OK**.

Editing a Search Group

You can edit the **Name** and **Description** fields of a search group.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.

3. Select **Search >New Search**.
4. Click **Manage Groups**.
5. Select the group that you want to edit.
6. Click **Edit**.
7. Type a new name in the **Name** field.
8. Type a new description in the **Description** field.
9. Click **OK**.

Copying a Saved Search to Another Group

You can copy a saved search to another group. You can also copy the saved search to more than one group.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Select **Search >New Search**.
4. Click **Manage Groups**.
5. Select the saved search that you want to copy.
6. Click **Copy**.
7. On the **Item Groups** window, select the check box for the group you want to copy the saved search to.
8. Click **Assign Groups**.

Removing a Group or a Saved Search from a Group

You can use the **Remove** icon to remove a search from a group or remove a search group.

When you remove a saved search from a group, the saved search is not deleted from your system. The saved search is removed from the group and automatically moved to the **Other** group.

You cannot remove the following groups from your system:

- Asset Search Groups
- Other

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Select **Search >New Search** .
4. Click **Manage Groups**.
5. Select the saved search that you want to remove from the group:
 - Select the saved search that you want to remove from the group.
 - Select the group that you want to remove.

Asset Profile Management Tasks

You can delete, import, and export asset profiles using the Assets tab.

Using the **Assets** tab, you can delete, import, and export asset profiles.

Deleting Assets

You can delete specific assets or all listed asset profiles.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. Select the asset that you want to delete, and then select **Delete Asset** from the **Actions** list box.
4. Click **OK**.

Importing Asset Profiles

You can import asset profile information.

The imported file must be a CSV file in the following format:

ip,name,weight,description

Where:

- **IP** Specifies any valid IP address in the dotted decimal format. For example: 192.168.5.34.
- **Name** Specifies the name of this asset up to 255 characters in length. Commas are not valid in this field and invalidate the import process. For example: WebServer01 is correct.

- **Weight** Specifies a number from 0 to 10, which indicates the importance of this asset on your network. A value of 0 denotes low importance and 10 is very high.
- **Description** Specifies a textual description for this asset up to 255 characters in length. This value is optional.

For example, the following entries might be included in a CSV file:

- **192.168.5.34,WebServer01,5,Main Production Web Server**
- **192.168.5.35,MailServ01,0,**

The import process merges the imported asset profiles with the asset profile information you have currently stored in the system.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. From the **Actions** list box, select **Import Assets**.
4. Click **Browse** to locate and select the CSV file that you want to import.
5. Click **Import Assets** to begin the import process.

Exporting Assets

You can export listed asset profiles to an Extended Markup Language (XML) or Comma-Separated Value (CSV) file.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles**.
3. From the **Actions** list box, select one of the following options:
 - Export to XML
 - Export to CSV
4. View the status window for the status of the export process.
5. If you want to use other tabs and pages while the export is in progress, click the **Notify When Done** link.

When the export is complete, the **File Download** window is displayed.

6. On the **File Download** window, choose one of the following options:
 - **Open** Select this option to open the export results in your choice of browser.

- **Save** Select this option to save the results to your desktop.

7. Click **OK**.

Research Asset Vulnerabilities

You can double-click the vulnerability to display more vulnerability details.

The **Research Vulnerability Details** window provides the following details:

Parameter	Description
Vulnerability ID	Specifies the ID of the vulnerability. The Vuln ID is a unique identifier that is generated by Vulnerability Information System (VIS).
Published Date	Specifies the date on which the vulnerability details were published on the OSVDB.
Name	Specifies the name of the vulnerability.
Assets	Specifies the number of assets in your network that have this vulnerability. Click the link to view the list of assets.
Assets, including exceptions	Specifies the number of assets in your network that have vulnerability exceptions. Click the link to view the list of assets.
CVE	Specifies the CVE identifier for the vulnerability. CVE identifiers are provided by the NVDB. Click the link to obtain more information. When you click the link, the NVDB website is displayed in a new browser window.
xforce	Specifies the X-Force identifier for the vulnerability. Click the link to obtain more information. When you click the link, the Internet Security Systems website is displayed in a new browser window.

(Continued)

Parameter	Description
OSVDB	<p>Specifies the OSVDB identifier for the vulnerability.</p> <p>Click the link to obtain more information. When you click the link, the OSVDB website is displayed in a new browser window.</p>
Plugin Details	<p>Specifies the JSA Vulnerability Manager ID.</p> <p>Click the link to view Oval Definitions, Windows Knowledge Base entries, or UNIX advisories for the vulnerability.</p> <p>This feature provides information on how JSA Vulnerability Manager checks for vulnerability details during a patch scan. You can use it to identify why a vulnerability was raised on an asset or why it was not.</p>
CVSS Score Base	<p>Displays the aggregate Common Vulnerability Scoring System (CVSS) score of the vulnerabilities on this asset. A CVSS score is an assessment metric for the severity of a vulnerability. You can use CVSS scores to measure how much concern a vulnerability warrants in comparison to other vulnerabilities.</p> <p>The CVSS score is calculated using the following user-defined parameters:</p> <ul style="list-style-type: none"> • Collateral Damage Potential • Confidentiality Requirement • Availability Requirement • Integrity Requirement <p>For more information about how to configure these parameters, see "Adding or Editing an Asset Profile" on page 187.</p> <p>For more information about CVSS, see http://www.first.org/cvss/.</p>
Impact	<p>Displays the type of harm or damage that can be expected if this vulnerability is exploited.</p>

(Continued)

Parameter	Description
CVSS Base Metrics	<p>Displays the metrics that are used to calculate the CVSS base score, including:</p> <ul style="list-style-type: none"> • Access Vector • Access complexity • Authentication • Confidentiality impact • Integrity impact • Availability impact
Description	<p>Specifies a description of the detected vulnerability. This value is only available when your system integrates VA tools.</p>
Concern	<p>Specifies the effects that the vulnerability can have on your network.</p>
Solution	<p>Follow the instructions that are provided to resolve the vulnerability.</p>
Virtual Patching	<p>Displays virtual patch information that is associated with this vulnerability, if available. A virtual patch is a short-term mitigation solution for a recently discovered vulnerability. This information is derived from Intrusion Protection System (IPS) events. If you want to install the virtual patch, see your IPS vendor information.</p>
Reference	<p>Displays a list of external references, including:</p> <ul style="list-style-type: none"> • Reference Type Specifies the type of reference that is listed, such as an advisory URL or mail post list. • URL Specifies the URL that you can click to view the reference. <p>Click the link to obtain more information. When you click the link, the external resource is displayed in a new browser window.</p>

(Continued)

Parameter	Description
Products	<p>Displays a list of products that are associated with this vulnerability.</p> <ul style="list-style-type: none">• Vendor Specifies the vendor of the product.• Product Specifies the product name.• Version Specifies the version number of the product.

1. Click the **Assets** tab.
2. On the navigation menu, click **Asset Profiles** .
3. Select an asset profile.
4. In the Vulnerabilities pane, click the **ID** or **Vulnerability** parameter value for the vulnerability you want to investigate.

RELATED DOCUMENTATION

[Updates to Asset Data | 170](#)

[Identification Of Asset Growth Deviations | 175](#)

[Asset Blacklists and Whitelists | 180](#)

9

CHAPTER

Chart Management

[Chart Management | 205](#)

[Time Series Chart Overview | 206](#)

[Chart Legends | 208](#)

[Configuring Charts | 208](#)

Chart Management

You can use various chart configuration options to view your data.

If you select a time frame or a grouping option to view your data, then the charts display above the event or flow list.

If you select a time frame or a grouping option to view your data, then the charts display above the event list.

Charts do not display while in streaming mode.

You can configure a chart to select what data you want to plot. You can configure charts independently of each other to display your search results from different perspectives.

Chart types include:

- Bar Chart - Displays data in a bar chart. This option is only available for grouped events.
- Pie Chart - Displays data in a pie chart. This option is only available for grouped events.
- Table - Displays data in a table. This option is only available for grouped events.
- Time Series - Displays an interactive line chart that represents the records that are matched by a specified time interval. For information about configuring time series search criteria, see "[Time Series Chart Overview](#)" on page 206.

After you configure a chart, your chart configurations are retained when you:

- Change your view by using the **Display** list box.
- Apply a filter.
- Save your search criteria.

Your chart configurations are not retained when you:

- Start a new search.
- Access a quick search.
- View grouped results in a branch window.
- Save your search results.

NOTE: If you use the Mozilla Firefox web browser and an ad blocker browser extension is installed, charts do not display. To display charts, you must remove the ad blocker browser extension. For more information, see your browser documentation.

RELATED DOCUMENTATION

[Time Series Chart Overview | 206](#)

[Chart Legends | 208](#)

[Configuring Charts | 208](#)

Time Series Chart Overview

Time series charts are graphical representations of your activity over time.

Peaks and valleys that are displayed in the charts depict high and low volume activity. Time series charts are useful for short-term and long term trending of data.

Using time series charts, you can access, navigate, and investigate log or network activity from various views and perspectives.

NOTE: You must have the appropriate role permissions to manage and view time series charts.

To display time series charts, you must create and save a search that includes time series and grouping options. You can save up to 100 time series searches.

Default time series saved searches are accessible from the list of available searches on the event or flow search page.

You can easily identify saved time series searches on the **Quick Searches** menu, because the search name is appended with the time range specified in the search criteria.

If your search parameters match a previously saved search for column definition and grouping options, a time series chart might automatically display for your search results. If a time series chart does not automatically display for your unsaved search criteria, no previously saved search criteria exists to match your search parameters. If this occurs, you must enable time series data capture and save your search criteria.

You can magnify and scan a timeline on a time series chart to investigate activity. The following table provides functions that you can use to view time series charts.

Table 33: Time Series Charts Functions

Function	Description
View data in greater detail	<p>Using the zoom feature, you can investigate smaller time segments of event traffic.</p> <ul style="list-style-type: none"> • Move your mouse pointer over the chart, and then use your mouse wheel to magnify the chart (roll the mouse wheel up). • Highlight the area of the chart you want to magnify. When you release your mouse button, the chart displays a smaller time segment. Now you can click and drag the chart to scan the chart. <p>When you magnify a time series chart, the chart refreshes to display a smaller time segment.</p>
View a larger time span of data	<p>Using the zoom feature, you can investigate larger time segments or return to the maximum time range. You can expand a time range using one of the following options:</p> <ul style="list-style-type: none"> • Click Zoom Reset at the upper left corner of the chart. • Move your mouse pointer over the chart, and then use your mouse wheel to expand the view (roll the mouse wheel down).
Scan the chart	<p>When you have magnified a time series chart, you can click and drag the chart to the left or right to scan the timeline.</p>

RELATED DOCUMENTATION

[Chart Legends | 208](#)

[Configuring Charts | 208](#)

[Chart Management | 205](#)

Chart Legends

Each chart provides a legend, which is a visual reference to help you associate the chart objects to the parameters they represent.

Using the legend feature, you can perform the following actions:

- Move your mouse pointer over a legend item or the legend color block to view more information about the parameters it represents.
- Right-click the legend item to further investigate the item.
- Click a pie or bar chart legend item to hide the item in the chart. Click the legend item again to show the hidden item. You can also click the corresponding graph item to hide and show the item.
- Click **Legend**, or the arrow beside it, if you want to remove the legend from your chart display.

RELATED DOCUMENTATION

[Configuring Charts | 208](#)

[Chart Management | 205](#)


[Time Series Chart Overview | 206](#)

Configuring Charts

You use configuration options to change the chart type, the object type you want to chart, and the number of objects that are represented on the chart. For time series charts, you can also select a time range and enable time series data capture.

Data can be accumulated so that when you perform a time series search, a cache of data is available to display data for the previous time period. After you enable time series data capture for a selected parameter, an asterisk (*) is displayed next to the parameter in the Value to Graph list box.

NOTE: Charts are not displayed when you view events or flows in Real Time (streaming) mode. To display charts, you must access the **Log Activity** or **Network Activity** tab and perform a grouped search that specifies a time range.

1. Click the **Log Activity** or **Network Activity** tab.
2. To create a grouped search, follow these steps:
 - a. On the toolbar, click **Search >New Search**.
 - b. From the **Available Saved Searches**, select a search and click **Load**.
 - c. Go to the Column Definition pane and if the **Group By** list box is empty, from the **Available Columns** list, select a column.
 - d. Click **Search**.
3. To use a grouped search, on the toolbar, click **Quick Searches** and select a grouped search.
4. In the Charts pane, click the **Configure** icon
(

).
).
).
).
5. Configure the following parameters:

Parameter	Description
Value to Graph	The object type that you want to graph on the Y axis of the chart. Options include all normalized and custom event or flow parameters that are included in your search parameters.
Display Top	The number of objects that you want to view in the chart. The default is 10. If you include more than 10 items in your chart, your data might be illegible.
Chart Type	If your bar, pie, or table chart is based on saved search criteria with a time range of more than 1 hour, you must click Update Details to update the chart and populate the event details.
Capture Time Series Data	Enables time series data capture. When you select this check box, the chart begins accumulating data for time series charts. By default, this option is disabled. This option is available only on Time Series charts.

(Continued)

Parameter	Description
Time Range	<p>The time range that you want to view.</p> <p>This option is only available on Time Series charts.</p>

6. If you selected the **Time Series** chart option and enabled the **Capture Time Series Data** option, in the Charts pane, click **Save** .
7. To view the list of events or flows if your time range is greater than 1 hour, click **Update Details**.

RELATED DOCUMENTATION

[Chart Management | 205](#)

[Time Series Chart Overview | 206](#)

[Chart Legends | 208](#)

10

CHAPTER

Event and Flow Searches

[Event and Flow Searches](#) | 212

[Offense Searches](#) | 235

[Finding IOCs Quickly with Lazy Search](#) | 250

[Deleting Search Criteria](#) | 251

[Using a Subsearch to Refine Search Results](#) | 252

[Managing Search Results](#) | 253

[Managing Search Groups](#) | 254

[Search Example: Daily Employee Reports](#) | 259

Event and Flow Searches

IN THIS SECTION

- [Creating a Customized Search | 212](#)
- [Querying with Dynamic Search | 219](#)
- [Saving Search Criteria | 220](#)
- [Scheduled Search | 222](#)
- [Quick Filter Search Options | 223](#)
- [Advanced Search Options | 226](#)

You can perform searches on the **Log Activity**, **Network Activity**, and **Offenses** tabs.

Use search and index options in JSA that improve search performance and return quicker results. To find specific criteria, advanced searches use AQL search strings.

You can specify filter criteria to search for events, flows, and offenses. After you perform a search, you can save the search criteria and the search results.

If your JSA administrator configured resource restrictions to set time or data limitations on event and flow searches, the resource restriction icon appears next to the search criteria.

Creating a Customized Search

You can search for data that match your criteria by using more specific search options. For example, you can specify columns for your search, which you can group and reorder to more efficiently browse your search results.

The duration of your search varies depending on the size of your database.

You can add new search options to filter through search results to find a specific event or flow that you are looking for.

The following table describes the search options that you can use to search event and flow data:

Table 34: Search Options

Options	Description
Group	<p>Select an event search group or flow search group to view in the Available Saved Searches list.</p> <p>Select an event search group to view in the Available Saved Searches list.</p>
Type Saved Search or Select from List	Type the name of a saved search or a keyword to filter the Available Saved Searches list.
Available Saved Searches	This list displays all available searches, unless you use Group or Type Saved Search or Select from List options to apply a filter to the list. You can select a saved search on this list to display or edit.
Search	The Search icon is available in multiple panes on the search page. You can click Search when you are finished configuring the search and want to view the results.
Include in my Quick Searches	Select this check box to include this search in your Quick Search menu.
Include in my Dashboard	<p>Select this check box to include the data from your saved search on the Dashboard tab. For more information about the Dashboard tab, see "Dashboard Management" on page 20.</p> <p>NOTE: This parameter is only displayed if the search is grouped.</p>
Set as Default	Select this check box to set this search as your default search.
Share with Everyone	Select this check box to share this search with all other users.
Real Time (streaming)	<p>Displays results in streaming mode.</p> <p>NOTE: When Real Time (streaming) is enabled, you are unable to group your search results. If you select any grouping option in the Column Definition pane, an error message opens.</p>
Last Interval (auto refresh)	The Log Activity and Network Activity tabs are refreshed at one-minute intervals to display the most recent information.

Table 34: Search Options (*Continued*)

Options	Description
Recent	<p>After you select this option, you must select a time range option from the list.</p> <p>NOTE: The results from the last minute might not be available. Select the <i><Specific Interval></i> option if you want to see all results.</p>
Specific Interval	<p>After you select this option, you must select the date and time range from the Start Time and End Time calendars.</p>
Data Accumulation	<p>Displayed when you load a saved search.</p> <p>If no data is accumulating for this saved search, the following information message is displayed: Data is not being accumulated for this search.</p> <p>If data is accumulating for this saved search, the following options are displayed:</p> <p>When you click or hover your mouse over the column link, a list of the columns that are accumulating data opens.</p> <p>Use the Enable Unique Counts/Disable Unique Counts link to display unique event and flow counts instead of average counts over time. After you click the Enable Unique Counts link, a dialog box opens and indicates which saved searches and reports share the accumulated data.</p>
Current Filters	<p>Displays the filters that are applied to this search.</p>
Save results when the search is complete	<p>Saves the search results.</p>
Display	<p>Specifies a predefined column that is set to display in the search results.</p>
Name	<p>The name of your custom column layout.</p>
Save Column Layout	<p>Saves a custom column layout that you modified.</p>
Delete Column Layout	<p>Deletes a saved custom column layout.</p>

Table 34: Search Options (Continued)

Options	Description
Type Column or Select from List	<p>Filter the columns that are listed in the Available Columns list.</p> <p>For example, type Device to display a list of columns that include Device in the column name.</p>
Available Columns	<p>Columns that are currently in use for this saved search are highlighted and displayed in the Columns list.</p>
Add and remove column arrows (top set)	<p>Use the top set of arrows to customize the Group By list.</p> <ul style="list-style-type: none"> • To add a column, select one or more columns from the Available Columns list and click the right arrow. • To remove a column, select one or more columns from the Group By list and click the left arrow.
Add and remove column arrows (bottom set)	<p>Use the bottom set of arrows to customize the Columns list.</p> <ul style="list-style-type: none"> • To add a column, select one or more columns from the Available Columns list and click the right arrow. • To remove a column, select one or more columns from the Columns list and click the left arrow.
Group By	<p>Specifies the columns on which the saved search groups the results.</p> <ul style="list-style-type: none"> • To move a column up the priority list, select a column and click the up arrow. You can also drag the column up the list. • To move a column down the priority list, select a column and click the down arrow. You can also drag the column down the list. <p>The priority list specifies in which order the results are grouped. The search results are grouped by the first column in the Group By list and then grouped by the next column on the list.</p>

Table 34: Search Options (Continued)

Options	Description
Columns	<p>Specifies columns that are chosen for the search. You can select more columns from the Available Columns list. You can further customize the Columns list by using the following options:</p> <ul style="list-style-type: none"> • To move a column up the priority list, select a column and click the up arrow. You can also drag the column up the list. • To move a column down the priority list, select a column and click the down arrow. You can also drag the column down the list. <p>If the column type is numeric or time-based and an entry is in the Group By list, then the column includes a list. Use the list to choose how you want to group the column.</p> <p>If the column type is group, the column includes a list to choose how many levels that you want to include for the group.</p>
Move columns between the Group By list and the Columns list	<p>Move columns between the Group By list and the Columns list by selecting a column in one list and dragging it to the other.</p>
Order By	<p>From the first list, select the column by which you want to sort the search results. Then, from the second list, select the order that you want to display for the search results.</p>
Results Limit	<p>Specifies the number of rows that a search returns on the Edit Search window. The Results Limit field also appears on the Results window.</p> <ul style="list-style-type: none"> • For a saved search, the limit is stored in the saved search and re-applied when search is loaded. • When you are sorting a column in the search result that has a row limit, sorting is done within the limited rows, which are shown in the data grid. • For a grouped by search where time series chart is turned on, the row limit applies only to the data grid. The Top N list in the time series chart controls how many time series are drawn in the chart.

1. Choose a search option:

- To search events, click the **Log Activity** tab.

- To search flows, click the **Network Activity** tab.
2. From the **Search** list, select **New Search**.
 3. Select a previously saved search.
 4. To create a search, in the Time Range pane, select the options for the time range that you want to capture for this search.

NOTE: The time range that you select might impact performance, when the time range is large.

5. Enable unique counts in the **Data Accumulation** pane.

NOTE: Enabling unique counts on accumulated data, which is shared with many other saved searches and reports might decrease system performance.

6. In the Search Parameters pane, define your search criteria.
 - a. From the first list, select a parameter that you want to search for.
 - b. From the second list, select the modifier that you want to use for the search.

NOTE: To search for an event or flow whose custom property does not have a value, use the **is N/A** operator. To search for an event or flow whose custom property has a value, use the **is not N/A** operator.

- c. From the entry field, type specific information that is related to your search parameter.
 - d. Click **Add Filter**.
 - e. Repeat these steps for each filter that you are adding to the search criteria.
7. To automatically save the search results when the search is complete, select the **Save results when search is complete** check box, and then type a name for the saved search.
 8. In the Column Definition pane, define the columns and column layout that you want to use to view the results:
 - a. From the **Display** list, select the preconfigured column that is set to associate with this search.
 - b. Click the arrow next to **Advanced View Definition** to display advanced search parameters.

- c. Customize the columns to display in the search results.
- d. In the **Results Limit** field, type the number of rows that you want the search to return.

TIP: If you configure a log source that belongs to multiple log source groups but has only one event that matches your search criteria, the search generates results for each log source group (including the parent group) that the event belongs to. This is expected behavior.

9. Click **Filter**.

Creating a Custom Column Layout

Create a custom column layout by adding or removing columns in an existing layout.

1. On the **Log Activity** or the **Network Activity** tab, click **Search >Edit Search**.
2. In the **Column Definition** pane, select an existing column layout in the **Display** list.

When you modify the layout, the name in the **Display** list is automatically changed to *Custom*.

3. Modify your search grouping.
 - a. To add a column to your search group, select a column from the **Available Columns** list and click the right arrow to move the column to the **Group By** list.
 - b. To move a column from the **Columns** list to your search group, select a column from the **Columns** list and drag it to the **Group By** list.
 - c. To remove a column from your search group, select the column from the **Group By** list and click the left arrow.
 - d. To change the order of your column groupings, use the up and down arrows or drag the columns into place.
4. Modify your column layout.
 - a. To add a column to your custom layout, select a column from the **Available Columns** list and click the right arrow to move the column to the **Columns** list.
 - b. To move a column from the **Group By** list to your custom layout, select a column from the **Group By** list and drag it to the **Columns** list.
 - c. To remove a column from your custom layout, select the column from the **Columns** list and click the left arrow.
 - d. To change the order of your columns, use the up and down arrows or drag the columns into place.

5. In the **Name** field, enter the name of your custom column layout.
6. Click **Save Column Layout**.

Deleting a Custom Column Layout

You can delete an existing user-created column layout.

1. On the **Log Activity** or the **Network Activity** tab, click **Search >Edit Search**.
2. In the **Column Definition** pane, select an existing user-created column layout in the **Display** list.
3. Click **Delete Column Layout**.

Querying with Dynamic Search

Use the dynamic search API to search for data that involves aggregated functions, such as **COUNT**, **SUM**, **MAX**, and **AVG**. For example, you can count the number of asset IDs per asset hostname by using the **COUNT_PER** function.

You can build your query on the following data sources:

- Assets
- Offenses
- Vulninstances

You can add a field without a function as a simple field, or you can add a field with a function as a complex field to build columns. You can also add conditions to filter your data.

1. Click the **Admin** tab.
2. In the **Dynamic Search** section, click **Dynamic Search**.
3. Select a **Data Source**.
4. Complete the **Available Columns** and **Available Filters** sections.
5. To add a name, description, range of the search, retention period, or search type to your query, enable one or more **Extra Search Properties**.
6. To copy your JSON script, click **Generate JSON**.

Your results appear in the **JSON generated by your query** section. Click **Copy to Clipboard** to copy your JSON script.

7. To reset your selections, click **Reset**.

8. Click **Run Query**.

The results of your query are listed in plain text or link format. For example, if you chose to query the **ASSET_ID** field, you can click the results to view the **Asset Summary** window for each asset ID.

Saving Search Criteria

You can save configured search criteria so that you can reuse the criteria and use the saved search criteria in other components, such as reports. Saved search criteria does not expire.

If you specify a time range for your search, then your search name is appended with the specified time range. For example, a saved search named Exploits by Source with a time range of Last 5 minutes becomes Exploits by Source - Last 5 minutes.

If you change a column set in a previously saved search, and then save the search criteria using the same name, previous accumulations for time series charts are lost.

1. Choose one of the following options:

- Click the **Log Activity** tab.
- Click the **Network Activity** tab.

2. Click the **Log Activity** tab.

3. Perform a search.

4. Click **Save Criteria**.

5. Enter values for the parameters:

Option	Description
Parameter	Description
Search Name	Type the unique name that you want to assign to this search criteria.

(Continued)

Option	Description
Assign Search to Group(s)	Select the check box for the group you want to assign this saved search. If you do not select a group, this saved search is assigned to the Other group by default. For more information, see "Managing Search Groups" on page 254 .
Manage Groups	Click Manage Groups to manage search groups. For more information, see "Managing Search Groups" on page 254 .
Timespan options:	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Real Time (streaming) - Select this option to filter your search results while in streaming mode. • Last Interval (auto refresh) Select this option to filter your search results while in auto-refresh mode. The Log Activity and Network Activity tabs refreshes at one-minute intervals to display the most recent information. • Last Interval (auto refresh) Select this option to filter your search results while in auto-refresh mode. The Log Activity and Network Activity tabs refreshes at one-minute intervals to display the most recent information. • Recent Select this option and, from this list box, select the time range that you want to filter for. • Specific Interval- Select this option and, from the calendar, select the date and time range you want to filter for.
Include in my Quick Searches	Select this check box to include this search in your Quick Search list box on the toolbar.

(Continued)

Option	Description
Include in my Dashboard	<p>Select this check box to include the data from your saved search on the Dashboard tab. For more information about the Dashboard tab, see "Dashboard Management" on page 20.</p> <p>NOTE: This parameter is only displayed if the search is grouped.</p>
Set as Default	<p>Select this check box to set this search as your default search.</p>
Share with Everyone	<p>Select this check box to share these search requirements with all users.</p>

6. Click **OK**.

Scheduled Search

Use the Scheduled search option to schedule a search and view the results.

You can schedule a search that runs at a specific time of day or night.

If you schedule a search to run in the night, you can investigate in the morning. Unlike reports, you have the option of grouping the search results and investigating further. You can search on number of failed logins in your network group. If the result is typically 10 and the result of the search is 100, you can group the search results for easier investigating. To see which user has the most failed logins, you can group by user name. You can continue to investigate further.

You can schedule a search on events or flows from the **Reports** tab. You must select a previously saved set of search criteria for scheduling.

1. Create a report

Specify the following information in the **Report Wizard** window:

- The chart type is Events/Logs or Flows.
- The report is based on a saved search.

- Generate an offense.

You can choose the **create an individual offense** option or the **add result to an existing offense** option.

You can also generate a manual search.

2. View search results

You can view the results of your scheduled search from the **Offenses** tab.

- Scheduled search offenses are identified by the **Offense Type** column.

If you create an individual offense, an offense is generated each time that the report is run. If you add the saved search result to an existing offense, an offense is created the first time that the report runs. Subsequent report runs append to this offense. If no results are returned, the system does not append or create an offense.

- To view the most recent search result in the **Offense Summary** window, double-click a scheduled search offense in the offense list. To view the list of all scheduled search runs, click **Search Results** in the **Last 5 Search Results** pane.

You can assign a Scheduled search offense to a user.

Quick Filter Search Options

Search event and flow payloads by typing a text search string that uses simple words or phrases.

Quick filter is one of the fastest methods that you use to search for event or flow payloads for specific data. For example, you can use quick filter to find these types of information:

- Every firewall device that is assigned to a specific address range in the past week
- A series of PDF files that were sent by a Gmail account in the past five days
- All records in a two-month period that exactly match a hyphenated user name
- A list of website addresses that end in .ca

You can filter your searches from these locations:

- **Log Activity toolbar and Network Activity toolbars**--Select **Quick Filter** from the list box on the **Search** toolbar to type a text search string. Click the **Quick Filter** icon to apply your **Quick Filter** to the list of events or flows.

- **Add Filter Dialog box**--Click the **Add Filter** icon on the **Log Activity** or **Network Activity** tab. Select **Quick Filter** as your filter parameter and type a text search string.
- **Flow search pages** --Add a quick filter to your list of filters.

NOTE: 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.

When you view **flows** in real-time (streaming) or last interval mode, you can type only simple words or phrases in the **Quick Filter** field. When you view **events** or **flows** in a time-range, follow these syntax guidelines:

Table 35: Quick Filter Syntax Guidelines

Description	Example
Include any plain text that you expect to find in the payload.	Firewall
Search for exact phrases by including multiple terms in double quotation marks.	"Firewall deny"
Include single and multiple character wildcards. The search term cannot start with a wildcard.	F?rwall or F??ew*
Group terms with logical expressions, such as AND, OR, and NOT. To be recognized as logical expressions and not as search terms, the syntax and operators must be uppercase.	(%PIX* AND ("Accessed URL" OR "Deny udp src") AND 10.100.100.*)
When you create search criteria that includes the NOT logical expression, you must include at least one other logical expression type, otherwise, no results are returned.	(%PIX* AND ("Accessed URL" OR "Deny udp src") NOT 10.100.100.*)
Precede the following characters by a backslash to indicate that the character is part of your search term: + - && ! () {} [] ^ " ~ * ? : \.	"%PIX\ -5\ -304001"

Limitations

Quick filter searches operate on raw event or flow log data and don't distinguish between the fields. For example, quick filter searches return matches for both source IP address and destination IP address, unless you include terms that can narrow the results.

Search terms are matched in sequence from the first character in the payload word or phrase. The search term **user** matches **user_1** and **user_2**, but does not match the following phrases: **ruser**, **myuser**, or **anyuser**.

Quick filter searches use the English locale. *Locale* is a setting that identifies language or geography and determines formatting conventions such as collation, case conversion, character classification, the language of messages, date and time representation, and numeric representation.

The locale is set by your operating system. You can configure JSA to override the operating system locale setting. For example, you can set the locale to **English** and the JSA console can be set to **Italiano (Italian)**.

If you use Unicode characters in your quick filter search query, unexpected search results might be returned.

If you choose a locale that is not English, you can use the Advanced search option in JSA for searching event and payload data.

How Does Quick Filter Search and Payload Tokens Work?

Text that is in the payload is split into words, phrases, symbols, or other elements. These tokens are delimited by space and punctuation. The tokens don't always match user-specified search terms, which cause some search terms not to be found when they don't match the generated token. The delimiter characters are discarded but exceptions exist such as the following exceptions:

- Periods that are not followed by white space are included as part of the token.

For example, **1.2.3.4:56** is tokenized as host token **1.2.3.4** and port token **56**.

- Words are split at hyphens, unless the word contains a number, in which case, the token is not split and the numbers and hyphens are retained as one token.
- Internet domain names and email addresses are preserved as a single token.

1.2.3.4/home/www is tokenized as one token and the URL is not separated.

1.2.3.7/calling1/www2/scp4/path5/fff is tokenized as host **1.2.3.7** and the remainder is one token / **calling1/www2/scp4/path5/fff**

File names and URL names that contain more than one underscore are split before a period (.).

Example of multiple underscores in a file name:

If you use `hurricane_katrina_ladm118.jpg` as a search term, it is split into the following tokens:

- `hurricane`
- `katrina_ladm118.jpg`

Search the payload for the full search term by placing double quotation marks around the search term: `"hurricane_katrina_ladm118.jpg"`

Example of multiple underscores in a relative file path:

The `thumb.ladm1180830/thumb.ladm11808301806.hurricane_katrina_ladm118.jpg` is split into the following tokens:

- `thumb.ladm1180830/thumb.ladm11808301806.hurricane`
- `katrina_ladm118.jpg`

To search for `hurricane_katrina_ladm118.jpg`, which consists of one partial and one full token, place an asterisk in front of the query term, `*hurricane_katrina_ladm118.jpg`

Advanced Search Options

Use the **Advanced Search** field to enter an Ariel Query Language (AQL) that specifies the fields that you want and how you want to group them to run a query.

NOTE: When you type an AQL query, use single quotation marks for a string comparison, and use double quotation marks for a property value comparison.

The **Advanced Search** field has auto completion and syntax highlighting.

Use auto completion and syntax highlighting to help create queries. For information about supported web browsers, see "[Supported Web Browsers](#)" on page 9

NOTE: If you use a quick filter on the **Log Activity** tab, you must refresh your browser window before you run an advanced search.

Accessing Advanced Search

Access the **Advanced Search** option from the **Search** toolbar that is on the **Network Activity** and **Log Activity** tabs to type an AQL query.

Access the **Advanced Search** option from the **Search** toolbar that is on the **Log Activity** tab to type an AQL query.

Select **Advanced Search** from the list box on the **Search** toolbar.

Expand the **Advanced Search** field by following these steps:

1. Drag the expand icon that is at the right of the field.
2. Press Shift + Enter to go to the next line.
3. Press Enter.

You can right-click any value in the search result and filter on that value.

Double-click any row in the search result to see more detail.

All searches, including AQL searches, are included in the audit log.

AQL Search String Examples

The following table provides examples of AQL search strings.

Table 36: Examples Of AQL Search Strings

Description	Example
Select default columns from events.	SELECT * FROM events
Select default columns from flows.	SELECT * FROM flows
Select default columns from events.	SELECT * FROM events
Select specific columns.	SELECT sourceip, destinationip FROM events
Select specific columns and order the results.	SELECT sourceip, destinationip FROM events ORDER BY destinationip
Run an aggregated search query.	SELECT sourceip, SUM(magnitude) AS magsum FROM events GROUP BY sourceip

Table 36: Examples Of AQL Search Strings (Continued)

Description	Example
Run a function call in a SELECT clause.	SELECT CATEGORYNAME(category) AS namedCategory FROM events
Filter the search results by using a WHERE clause.	SELECT CATEGORYNAME(category) AS namedCategory, magnitude FROM events WHERE magnitude > 1
Search for events that triggered a specific rule, which is based on the rule name or partial text in the rule name.	SELECT LOGSOURCENAME(logsourceid), * from events where RULENAME(creeventlist) ILIKE '%suspicious%'
Reference field names that contain special characters, such as arithmetic characters or spaces, by enclosing the field name in double quotation marks.	SELECT sourceip, destinationip, "+field/name+" FROM events WHERE "+field/name+" LIKE '%test%'

The following table provides examples of AQL search strings for X-Force.

Table 37: Examples Of AQL Search Strings for X-Force

Description	Example
Check an IP address against an X-Force category with a confidence value.	select * from events where XFORCE_IP_CONFIDENCE('Spam',sourceip)>3
Search for X-Force URL categories associated with a URL.	select url, XFORCE_URL_CATEGORY(url) as myCategories from events where XFORCE_URL_CATEGORY(url) IS NOT NULL
Retrieve X-Force IP categories that are associated with an IP.	select sourceip, XFORCE_IP_CATEGORY(sourceip) as IPcategories from events where XFORCE_IP_CATEGORY(sourceip) IS NOT NULL

For more information about functions, search fields and operators, see the *Juniper Secure Analytics Ariel Query Language guide*.

AQL Search String Examples

Use the Ariel Query Language (AQL) to retrieve specific fields from the events, flows, and simarc tables in the Ariel database.

NOTE: When you build an AQL query, if you copy text that contains single quotation marks from any document and paste the text into JSA, your query will not parse. As a workaround, you can paste the text into JSA and retype the single quotation marks.

Reporting Account Usage

Different user communities can have different threat and usage indicators.

Use reference data to report on several user properties, for example, department, location, or manager. You can use external reference data.

The following query returns metadata information about the user from their login events.

```
SELECT
REFERENCETABLE('user_data','FullName',username) as 'Full Name',
REFERENCETABLE('user_data','Location',username) as 'Location',
REFERENCETABLE('user_data','Manager',username) as 'Manager',
DISTINCTCOUNT(username) as 'Userid Count',
DISTINCTCOUNT(sourceip) as 'Source IP Count',
COUNT(*) as 'Event Count'
FROM events
WHERE qidname(qid) ILIKE '%logon%'
GROUP BY 'Full Name', 'Location', 'Manager'
LAST 1 days
```

Insight Across Multiple Account Identifiers

In this example, individual users have multiple accounts across the network. The organization requires a single view of a users activity.

Use reference data to map local user IDs to a global ID.

The following query returns the user accounts that are used by a global ID on events that are flagged as suspicious.

```
SELECT
  REFERENCEMAP('GlobalID Mapping',username) as 'Global ID',
  REFERENCETABLE('user_data','FullName', 'Global ID') as 'Full Name',
  DISTINCTCOUNT(username),
  COUNT(*) as 'Event count'
FROM events
WHERE RULENAME(creEventlist) ILIKE '%suspicious%'
GROUP BY 'Global ID'
LAST 1 days
```

The following query shows the activities that are completed by a global ID.

```
SELECT
  QIDNAME(qid) as 'Event name',
  starttime as 'Time',
  sourceip as 'Source IP', destinationip as 'Destination IP',
  username as 'Event Username',
  REFERENCEMAP('GlobalID_Mapping', username)as 'Global User'
FROM events
WHERE 'Global User' = 'John Doe'
LAST 1 days
```

Identify Suspicious Long-term Beaconing

Many threats use command and control to communicate periodically over days, weeks, and months.

Advanced searches can identify connection patterns over time. For example, you can query consistent, short, low volume, number of connections per day/week/month between IP addresses, or an IP address and geographical location.

The following query detects potential instances of hourly beaconing.

```
SELECT sourceip, destinationip,
  DISTINCTCOUNT(DATEFORMAT(starttime,'HH')) as 'different hours',
  COUNT(*) as 'total flows'
FROM flows
WHERE flowdirection = 'L2R'
GROUP BY sourceip, destinationip
```

```
HAVING "different hours" > 20
AND "total flows" < 25
LAST 24 hours
```

TIP: You can modify this query to work on proxy logs and other event types.

The following query detects potential instances of daily beaconing.

```
SELECT sourceip, destinationip,
DISTINCTCOUNT(DATEFORMAT(starttime,'dd'))as 'different days',
COUNT(*) as 'total flows'
FROM flows
WHERE flowdirection='L2R'
GROUP BY sourceip, destinationip
HAVING "different days" > 4
AND "total flows" < 14
LAST 7 days
```

The following query detects daily beaconing between a source IP and a destination IP. The beaconing times are not at the same time each day. The time lapse between beacons is short.

```
SELECT
sourceip,
LONG(DATEFORMAT(starttime,'hh')) as hourofday,
(AVG( hourofday*hourofday) - (AVG(hourofday)^2))as variance,
COUNT(*) as 'total flows'
FROM flows
GROUP BY sourceip, destinationip
HAVING variance < 01 and "total flows" < 10
LAST 7 days
```

The following query detects daily beaconing to a domain by using proxy log events. The beaconing times are not at the same time each day. The time lapse between beacons is short.

```
SELECT sourceip,
LONG(DATEFORMAT(starttime,'hh')) as hourofday,
(AVG(hourofday*hourofday) - (AVG(hourofday)^2)) as variance,
COUNT(*) as 'total events'
FROM events
```



```
WHERE LOGSOURCEGROUPNAME(devicegroup1ist) ILIKE '%proxy%'
GROUP BY url_domain
HAVING variance < 0.1 and "total events" < 10
LAST 7 days
```

The `url_domain` property is a custom property from proxy logs.

External Threat Intelligence

Usage and security data that is correlated with external threat intelligence data can provide important threat indicators.

Advanced searches can cross-reference external threat intelligence indicators with other security events and usage data.

This query shows how you can profile external threat data over many days, weeks, or months to identify and prioritize the risk level of assets and accounts.

```
Select
REFERENCETABLE('ip_threat_data','Category',destinationip) as 'Category',
REFERENCETABLE('ip_threat_data','Rating', destinationip) as 'Threat Rating',
DISTINCTCOUNT(sourceip) as 'Source IP Count',
DISTINCTCOUNT(destinationip) as 'Destination IP Count'
FROM events
GROUP BY 'Category', 'Threat Rating'
LAST 1 days
```

Asset Intelligence and Configuration

Threat and usage indicators vary by asset type, operating system, vulnerability posture, server type, classification, and other parameters.

In this query, advanced searches and the asset model provide operational insight into a location.

The `Assetproperty` function retrieves property values from assets, which enables you to include asset data in the results.

```
SELECT
ASSETPROPERTY('Location',sourceip) as location,
COUNT(*) as 'event count'
FROM events
```

```
GROUP BY location
LAST 1 days
```

The following query shows how you can use advanced searches and user identity tracking in the asset model.

The **AssetUser** function retrieves the user name from the asset database.

```
SELECT
APPLICATIONNAME(applicationid) as App,
ASSETUSER(sourceip, now()) as srcAssetUser,
COUNT(*) as 'Total Flows'
FROM flows
WHERE srcAssetUser IS NOT NULL
GROUP BY App, srcAssetUser
ORDER BY "Total Flows" DESC
LAST 3 HOURS
```

Network LOOKUP Function

You can use the **Network LOOKUP** function to retrieve the network name that is associated with an IP address.

```
SELECT NETWORKNAME(sourceip) as srcnet,
NETWORKNAME(destinationip) as dstnet
FROM events
```

Rule LOOKUP Function

You can use the **Rule LOOKUP** function to retrieve the name of a rule by its ID.

```
SELECT RULENAME(123) FROM events
```

The following query returns events that triggered a specific rule name.

```
SELECT * FROM events
WHERE RULENAME(creEventList) ILIKE '%my rule name%'
```

Full TEXT SEARCH

You can use the TEXT SEARCH operator to do full text searches by using the **Advanced search** option.

In this example, there are a number of events that contain the word "firewall" in the payload. You can search for these events by using the **Quick filter** option and the **Advanced search** option on the **Log Activity** tab.

- To use the **Quick filter** option, type the following text in the **Quick filter** box: 'firewall'
- To use the **Advanced search** option, type the following query in the **Advanced search** box:

```
SELECT QIDNAME(qid) AS EventName, * from events where TEXT SEARCH 'firewall'
```

Custom Property

You can access custom properties for events and flows when you use the **Advanced search** option.

The following query uses the custom property "MyWebsiteUrl" to sort events by a particular web URL:

```
SELECT "MyWebsiteUrl", * FROM events ORDER BY "MyWebsiteUrl"
```

Converting a Saved Search to an AQL String

Convert a saved search to an AQL string and modify it to create your own searches to quickly find the data you want. Now you can create searches faster than by typing the search criteria. You can also save the search for future use.

1. Click the **Log Activity** or **Network Activity** tab.
2. From the **Search** list, select **New Search** or **Edit Search**.
3. Select a previously saved search.
4. Click **Show AQL**.
5. From the AQL window, click **Copy to Clipboard**.
6. In the **Search Mode** section, click **Advanced Search**.
7. Paste the AQL string text into the **Advanced Search** text box.
8. Modify the string to include the data you want to find.
9. Click **Search** to display the results.

RELATED DOCUMENTATION

[Offense Searches | 235](#)

[Finding IOCs Quickly with Lazy Search | 250](#)

[Deleting Search Criteria | 251](#)

Offense Searches

IN THIS SECTION

- [Searching Offenses on the My Offenses and All Offenses Pages | 235](#)
- [Searching Offenses on the By Source IP Page | 242](#)
- [Searching Offenses on the By Destination IP Page | 244](#)
- [Searching Offenses on the By Networks Page | 246](#)
- [Saving Search Criteria on the Offenses Tab | 247](#)
- [Searching for Offenses That Are Indexed on a Custom Property | 249](#)

You can search offenses by using specific criteria to display offenses that match the search criteria in a results list.

You can create a new search or load a previously saved set of search criteria.

Searching Offenses on the My Offenses and All Offenses Pages

On the **My Offenses** and **All Offenses** pages of the **Offense** tab, you can search for offenses that match your criteria.

The following table describes the search options that you can use to search offense data on the **My Offenses** and **All Offenses** pages.

For information about categories, see the *Juniper Secure Analytics Administration Guide*.

Table 38: My Offenses and All Offenses Page Search Options

Options	Description
Group	This list box allows you to select an offense Search Group to view in the Available Saved Searches list.
Type Saved Search or Select from List	This field allows you to type the name of a saved search or a keyword to filter the Available Saved Searches list.
Available Saved Searches	This list displays all available searches, unless you apply a filter to the list using the Group or Type Saved Search or Select from List options. You can select a saved search on this list to display or edit.
All Offenses	This option allows you to search all offenses regardless of time range.
Recent	This option allows you to select a pre-defined time range you want to filter for. After you select this option, you must select a time range option from the list box.
Specific Interval	<p>This option allows you to configure a custom time range for your search. After you select this option, you must select one of the following options.</p> <ul style="list-style-type: none"> • Start Date between Select this check box to search offenses that started during a certain time period. After you select this check box, use the list boxes to select the dates you want to search. • Last Event/Flow between Last Event between Select this check box to search offenses for which the last detected event occurred within a certain time period. After you select this check box, use the list boxes to select the dates you want to search.
Search	The Search icon is available in multiple panes on the search page. You can click Search when you are finished configuring the search and want to view the results.
Offense Id	In this field, you can type the Offense ID you want to search for.
Description	In this field, you can type the description that you want to search for.
Assigned to user	From this list box, you can select the user name that you want to search for.

Table 38: My Offenses and All Offenses Page Search Options (Continued)

Options	Description
Direction	<p>From this list box, you can select the offense direction that you want to search for. Options include:</p> <ul style="list-style-type: none"> • Local to Local • Local to Remote • Remote to Local • Remote to Remote • Local to Remote or Local • Remote to Remote or Local
Source IP	In this field, you can type the source IP address or CIDR range you want to search for.
Destination IP	In this field, you can type the destination IP address or CIDR range you want to search for.
Magnitude	From this list box, you can specify a magnitude and then select to display only offenses with a magnitude that is equal to, less than, or greater than the configured value. The range is 0 - 10.
Severity	From this list box, you can specify a severity and then select to display only offenses with a severity that is equal to, less than, or greater than the configured value. The range is 0 - 10.
Credibility	From this list box, you can specify a credibility and then select to display only offenses with a credibility that is equal to, less than, or greater than the configured value. The range is 0 - 10.
Relevance	From this list box, you can specify a relevance and then select to display only offenses with a relevance that is equal to, less than, or greater than the configured value. The range is 0 - 10.

Table 38: My Offenses and All Offenses Page Search Options (Continued)

Options	Description
Contains Username	In this field, you can type a regular expression (regex) statement to search for offenses containing a specific user name. When you define custom regex patterns, adhere to regex rules as defined by the Java programming language. For more information, you can refer to regex tutorials available on the web.
Source Network	From this list box, you can select the source network that you want to search for.
Destination Network	From this list box, you can select the destination network that you want to search for.
High Level Category	From this list box, you can select the high-level category that you want to search for.
Low Level Category	From this list box, you can select the low-level category that you want to search for.
Exclude	<p>The options in this pane allow you to exclude offenses from the search results. The options include:</p> <ul style="list-style-type: none"> • Active Offenses • Hidden Offenses • Closed Offenses • Inactive offenses • Protected Offense
Close by User	<p>This parameter is only displayed when the Closed Offenses check box is cleared in the Exclude pane.</p> <p>From this list box, you can select the user name that you want to search closed offenses for or select Any to display all closed offenses.</p>

Table 38: My Offenses and All Offenses Page Search Options (Continued)

Options	Description
Reason For Closing	<p>This parameter is only displayed when the Closed Offenses check box is cleared in the Exclude pane.</p> <p>From this list box, you can select a reason that you want to search closed offenses for or select Any to display all closed offenses.</p>
Events	From this list box, you can specify an event count and then select to display only offenses with an event count that is equal to, less than, or greater than the configured value.
Flows	From this list box, you can specify a flow count and then select to display only offenses with a flow count that is equal to, less than, or greater than the configured value.
Total Events/Flows Total Events	From this list box, you can specify a total event and flow count and then select to display only offenses with a total event and flow count that is equal to, less than, or greater than the configured value.
Destinations	From this list box, you can specify a destination IP address count and then select to display only offenses with a destination IP address count that is equal to, less than, or greater than the configured value.
Log Source Group	From this list box, you can select a log source group that contains the log source you want to search for. The Log Source list box displays all log sources that are assigned to the selected log source group.
Log Source	From this list box, you can select the log source that you want to search for.
Rule Group	From this list box, you can select a rule group that contains the contributing rule that you want to search for. The Rule list box displays all rules that are assigned to the selected rule group.
Rule	From this list box, you can select the contributing rule that you want to search for.
Offense Type	From this list box, you can select an offense type that you want to search for. For more information about the options in the Offense Type list box, see Table 2.

The following table describes the options available in the **Offense Type** list box:

Table 39: Offense Type Options

Offense types	Description
Any	This option searches all offense sources.
Source IP	To search for offenses with a specific source IP address, you can select this option, and then type the source IP address that you want to search for.
Destination IP	To search for offenses with a specific destination IP address, you can select this option, and then type the destination IP address that you want to search for.
Event Name	<p>To search for offenses with a specific event name, you can click the Browse icon to open the Event Browser and select the event name (QID) you want to search for.</p> <p>You can search for a particular QID using one of the following options:</p> <ul style="list-style-type: none"> • To search for a QID by category, select the Browse by Category check box and select the high- or low-level category from the list boxes. • To search for a QID by log source type, select the Browse by Log Source Type check box and select a log source type from the Log Source Type list box. • To search for a QID by log source type, select the Browse by Log Source Type check box and select a log source type from the Log Source Type list box. • To search for a QID by name, select the QID Search check box and type a name in the QID/Name field.
Username	To search for offenses with a specific user name, you can select this option, and then type the user name that you want to search for.
Source MAC Address	To search for offenses with a specific source MAC address, you can select this option, and then type the source MAC address that you want to search for.
Destination MAC Address	To search for offenses with a specific destination MAC address, you can select this option, and then type the destination MAC address that you want to search for.

Table 39: Offense Type Options (Continued)

Offense types	Description
Log Source	<p>From the Log Source Group list box, you can select the log source group that contains the log source you want to search for. The Log Source list box displays all log sources that are assigned to the selected log source group.</p> <p>From the Log Source list box, select the log source that you want to search for.</p>
Host Name	To search for offenses with a specific host name, you can select this option, and then type the host name that you want to search for.
Source Port	To search for offenses with a specific source port, you can select this option, and then type the source port that you want to search for.
Destination Port	To search for offenses with a specific destination port, you can select this option, and then type the destination port that you want to search for.
Source IPv6	To search for offenses with a specific source IPv6 address, you can select this option, and then type the source IPv6 address that you want to search for.
Destination IPv6	To search for offenses with a specific destination IPv6 address, you can select this option, and then type the destination IPv6 address that you want to search for.
Source ASN	To search for offenses with a specific Source ASN, you can select the source ASN from the Source ASN list box.
Destination ASN	To search for offenses with a specific destination ASN, you can select the destination ASN from the Destination ASN list box.
Rule	To search for offenses that are associated with a specific rule, you can select the rule group that contains the rule you want to search from the Rule Group list box. The Rule Group list box displays all rules that are assigned to the selected rule group. From the Rule list box, you select the rule that you want to search for.
App ID	To search for offenses with an application ID, you can select the application ID from the App ID list box.

1. Click the **Offenses** tab.

2. From the **Search** list box, select **New Search**.
3. Choose one of the following options:
 - To load a previously saved search, go to Step "4" on page 242.
 - To create a new search, go to Step "7" on page 242.
4. Select a previously saved search using one of the following options:
 - From the **Available Saved Searches** list, select the saved search that you want to load.
 - In the **Type Saved Search** or **Select from List** field, type the name of the search you want to load.
5. Click **Load**.
6. Optional. Select the **Set as Default** check box in the Edit Search pane to set this search as your default search. If you set this search as your default search, the search automatically performs and displays results each time you access the **Offenses** tab.
7. On the Time Range pane, select an option for the time range you want to capture for this search. See Table 1.
8. On the Search Parameters pane, define your specific search criteria. See Table 1.
9. On the Offense Source pane, specify the offense type and offense source you want to search:
 - a. From the list box, select the offense type that you want to search for.
 - b. Type your search parameters. See Table 2.
10. In the Column Definition pane, define the order in which you want to sort the results:
 - a. From the first list box, select the column by which you want to sort the search results.
 - b. From the second list box, select the order that you want to display for the search results. Options include Descending and Ascending.
11. Click **Search**.

["Saving Search Criteria on the Offenses Tab" on page 247](#)

Searching Offenses on the By Source IP Page

This topic provides the procedure for how to search offenses on the **By Source IP** page of the **Offense** tab.

The following table describes the search options that you can use to search offense data on the **By Source IP** page:

Table 40: By Source IP Page Search Options

Options	Description
All Offenses	You can select this option to search all source IP addresses regardless of time range.
Recent	You can select this option and, from this list box, select the time range that you want to search for.
Specific Interval	To specify an interval to search for, you can select the Specific Interval option and then select one of the following options: <ul style="list-style-type: none"> • Start Date between Select this check box to search source IP addresses associated with offenses that started during a certain time period. After you select this check box, use the list boxes to select the dates you want to search for. • Last Event/Flow between Last Event between Select this check box to search source IP addresses associated with offenses for which the last detected event occurred within a certain time period. After you select this check box, use the list boxes to select the dates you want to search for.
Search	The Search icon is available in multiple panes on the search page. You can click Search when you are finished configuring the search and want to view the results.
Source IP	In this field, you can type the source IP address or CIDR range you want to search for.
Magnitude	From this list box, you can specify a magnitude and then select display only offenses with a magnitude that is equal to, less than, or greater than the configured value. The range is 0 - 10.
VA Risk	From this list box, you can specify a VA risk and then select display only offenses with a VA risk that is equal to, less than, or greater than the configured value. The range is 0 - 10.
Events/Flows Events	From this list box, you can specify an event or flow count and then select display only offenses with a magnitude that is equal to, less than, or greater than the configured value.

Table 40: By Source IP Page Search Options (Continued)

Options	Description
Exclude	<p>You can select the check boxes for the offenses you want to exclude from the search results. The options include:</p> <ul style="list-style-type: none"> • Active Offenses • Hidden Offenses • Closed Offenses • Inactive offenses • Protected Offense

1. Click the **Offenses** tab.
2. Click **By Source IP**.
3. From the **Search** list box, select **New Search**.
4. On the Time Range pane, select an option for the time range you want to capture for this search. See Table 1.
5. On the Search Parameters pane, define your specific search criteria. See Table 1.
6. On the Column Definition pane, define the order in which you want to sort the results:
 - a. From the first list box, select the column by which you want to sort the search results.
 - b. From the second list box, select the order that you want to display for the search results. Options include **Descending** and **Ascending**.
7. Click **Search**.

["Saving Search Criteria on the Offenses Tab" on page 247](#)

Searching Offenses on the By Destination IP Page

On the **By Destination IP** page of the **Offense** tab, you can search offenses that are grouped by the destination IP address.

The following table describes the search options that you can use to search offenses on the **By Destination IP** page:

Table 41: By Destination IP Page Search Options

Options	Description
All Offenses	You can select this option to search all destination IP addresses regardless of time range.
Recent	You can select this option and, From this list box, select the time range that you want to search for.
Specific Interval	To specify a particular interval to search for, you can select the Specific Interval option, and then select one of the following options: <ul style="list-style-type: none"> To specify a particular interval to search for, you can select the Specific Interval option, and then select one of the following options: Last Event/Flow between Last Event between Select this check box to search destination IP addresses associated with offenses for which the last detected event occurred within a certain time period. After you select this check box, use the list boxes to select the dates you want to search
Search	The Search icon is available in multiple panes on the search page. You can click Search when you are finished configuring the search and want to view the results.
Destination IP	You can type the destination IP address or CIDR range you want to search for.
Magnitude	From this list box, you can specify a magnitude, and then select display only offenses with a magnitude that is equal to, less than, or greater than the configured value.
VA Risk	From this list box, you can specify a VA risk, and then select display only offenses with a VA risk that is equal to, less than, or greater than the configured value. The range is 0 - 10.
Events/Flows Events	From this list box, you can specify an event or flow count magnitude, and then select display only offenses with an event or flow count that is equal to, less than, or greater than the configured value.

1. Click the **Offenses** tab.
2. On the navigation menu, click **By Destination IP**.

3. From the **Search** list box, select **New Search**.
4. On the Time Range pane, select an option for the time range you want to capture for this search. See Table 1.
5. On the Search Parameters pane, define your specific search criteria. See Table 1.
6. On the Column Definition pane, define the order in which you want to sort the results:
 - a. From the first list box, select the column by which you want to sort the search results.
 - b. From the second list box, select the order in which you want to display the search results. Options include **Descending** and **Ascending**.
7. Click **Search**.

Saving search criteria on the Offense tab On the **Offenses** tab, you can save configured search criteria so that you can reuse the criteria for future searches. Saved search criteria does not expire.

Searching Offenses on the By Networks Page

On the **By Network** page of the **Offense** tab, you can search offenses that are grouped by the associated networks.

The following table describes the search options that you can use to search offense data on the **By Networks** page:

Table 42: Search Options for Search Offense Data on the By Networks Page

Option	Description
Network	From this list box, you can select the network that you want to search for.
Magnitude	From this list box, you can specify a magnitude, and then select display only offenses with a magnitude that is equal to, less than, or greater than the configured value.
VA Risk	From this list box, you can specify a VA risk, and then select display only offenses with a VA risk that is equal to, less than, or greater than the configured value.

Table 42: Search Options for Search Offense Data on the By Networks Page (Continued)

Option	Description
Event/Flows Event	From this list box, you can specify an event or flow count, and then select display only offenses with an event or flow count that is equal to, less than, or greater than the configured value.

1. Click the **Offenses** tab.
2. Click **By Networks**.
3. From the **Search** list box, select **New Search**.
4. On the Search Parameters pane, define your specific search criteria. See Table 1.
5. On the Column Definition pane, define the order in which you want to sort the results:
 - a. From the first list box, select the column by which you want to sort the search results.
 - b. From the second list box, select the order in which you want to display the search results. Options include **Descending** and **Ascending**.
6. Click **Search**.

["Saving Search Criteria on the Offenses Tab" on page 247](#)

Saving Search Criteria on the Offenses Tab

On the **Offenses** tab, you can save configured search criteria so that you can reuse the criteria for future searches. Saved search criteria does not expire.

1. Procedure
2. Perform a search. See Offense searches.
3. Click **Save Criteria**.
4. Enter values for the following parameters:

Option	Description
Parameter	Description
Search Name	Type a name you want to assign to this search criteria.
Manage Groups	Click Manage Groups to manage search groups. See "Managing Search Groups" on page 254 .
Timespan options:	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • All Offenses Select this option to search all offenses regardless of time range. • Recent Select the option and, from this list box, select the time range that you want to search for. • Specific Interval - To specify a particular interval to search for, select the Specific Interval option, and then select one of the following options: <ul style="list-style-type: none"> Start Date between - Select this check box to search offenses that started during a certain time period. After you select this check box, use the list boxes to select the dates you want to search for. Last Event/Flow between - Select this check box to search offenses for which the last detected event occurred within a certain time period. After you select this check box, use the list boxes to select the dates you want to search. Last Event between - Select this check box to search offenses for which the last detected event occurred within a certain time period. After you select this check box, use the list boxes to select the dates you want to search.
Set as Default	Select this check box to set this search as your default search.

5. Click **OK**.

Searching for Offenses That Are Indexed on a Custom Property

Define search criteria to filter the offense list and make it easier to see which offenses you need to investigate. You can use the offense type in your search criteria to find all offenses that are based on a custom property. You can filter the query results to show offenses that have a specific custom property capture result.

The custom property must be used as a rule index. For more information, see ["Offense Indexing" on page 60](#).

1. Click the **Offenses** tab.
2. From the **Search** list, select **New Search**.
3. On the **Offense Source** pane, select the custom property in the **Offense Type** list.

The **Offense Type** list shows only normalized fields and custom properties that are used as rule indexes. You cannot use **Offense Source** to search DateTime properties.

4. To search for offenses that have a specific value in the custom property capture result, type the value that you want to search for in the filter box.
5. Configure other search parameters to satisfy your search requirements.
6. Click **Search**.

All offenses that meet the search criteria are shown in the offense list. When you view the offense summary, the custom property that you searched on is shown in the **Offense Type** field. The custom property capture result is shown in the **Custom Property Value** field in the **Offense Source Summary** pane.

RELATED DOCUMENTATION

[Finding IOCs Quickly with Lazy Search | 250](#)

[Deleting Search Criteria | 251](#)

[Using a Subsearch to Refine Search Results | 252](#)

Finding IOCs Quickly with Lazy Search

Lazy search returns the first 1000 events that are related to the search criterion. For example, if you need to search for a particular MD5 as part of a malware outbreak investigation, you do not need to review every related event. Do a *lazy search* to quickly return a limited result set.

To take advantage of the *lazy search*, you must have the Admin security profile, or a non-administrator security profile that is configured in the following way:

- Permission precedence set to **No Restrictions**.
- Access to all networks and log sources.

Lazy search cannot be used by users with non-administrator security profiles on networks where domains are configured.

You use the *JSA/lazy search* to search for an indicator of compromise (IOC), such as unusual outbound network traffic or anomalies in privileged user account activity.

1. To do a lazy search for quick filters, do these steps:
 - a. On the **Log Activity** tab, in the **Quick Filter** field, enter a value.
 - b. From the **View** list, select a time range.
2. To do a lazy search for basic searches, do these steps:
 - a. On the Log Activity tab, click **Search >New Search**.
 - b. Select a **Recent** time range or set a **Specific Interval**.
 - c. Ensure that **Order by** field value is set to **Start Time** and the **Results Limit** field value is **1000** or less. Aggregated columns must not be included in the search.
 - d. Enter a value for the **Quick Filter** parameter and click **Add Filter**.
3. To disable lazy search completely, do these steps:
 - a. Click the **System Settings** on the **Admin** tab.
 - b. In the **System Settings** window, remove any values from the **Default Search Limit** field.

RELATED DOCUMENTATION

[Deleting Search Criteria | 251](#)

[Using a Subsearch to Refine Search Results | 252](#)

[Managing Search Results | 253](#)

Deleting Search Criteria

You can delete search criteria.

When you delete a saved search, then objects that are associated with the saved search might not function. Reports are JSA objects that use saved search criteria. After you delete a saved search, edit the associated objects to ensure that they continue to function.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. Click the **Log Activity** tab.
3. From the **Search** list box, select **New Search** or **Edit Search**.
4. In the Saved Searches pane, select a saved search from the **Available Saved Searches** list box.
5. Click **Delete**.
 - If the saved search criteria is not associated with other JSA objects, a confirmation window is displayed.
 - If the saved search criteria is associated with other objects, the **Delete Saved Search** window is displayed. The window lists objects that are associated with the saved search that you want to delete. Note the associated objects.
6. Click **OK**.
7. Choose one of the following options:
 - Click **OK** to proceed.
 - Click **Cancel** to close the **Delete Saved Search** window.

If the saved search criteria was associated with other JSA objects, access the associated objects that you noted and edit the objects to remove or replace the association with the deleted saved search.

RELATED DOCUMENTATION

[Using a Subsearch to Refine Search Results | 252](#)

[Managing Search Results | 253](#)

[Managing Search Groups | 254](#)

Using a Subsearch to Refine Search Results

When you define a search that you want to use as a base for subsearching, make sure that Real Time (streaming) option is disabled and the search is not grouped.

You can use a subsearch to search within a set of completed search results. The subsearch is used to refine search results, without searching the database again.

This feature is not available for grouped searches, searches in progress, or in streaming mode.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. Click the **Log Activity** tab.
3. Perform a search.
4. When your search is complete, add another filter:
 - a. Click **Add Filter**.
 - b. From the first list box, select a parameter that you want to search for.
 - c. From the second list box, select the modifier that you want to use for the search. The list of modifiers that are available depends on the attribute that is selected in the first list.
 - d. In the entry field, type specific information that is related to your search.
 - e. Click **Add Filter**.

The Original Filter pane specifies the original filters that are applied to the base search. The Current Filter pane specifies the filters that are applied to the subsearch. You can clear subsearch filters without restarting the base search. Click the **Clear Filter** link next to the filter you want to clear. If you clear a filter from the Original Filter pane, the base search is relaunched.

If you delete the base search criteria for saved subsearch criteria, you still have access to saved subsearch criteria. If you add a filter, the subsearch searches the entire database since the search function no longer bases the search on a previously searched data set.

["Saving Search Criteria" on page 220](#)

RELATED DOCUMENTATION

[Managing Search Results | 253](#)

[Managing Search Groups | 254](#)

Managing Search Results

IN THIS SECTION

- [Canceling a Search | 253](#)
- [Deleting a Search | 254](#)

You can initiate multiple searches, and then navigate to other tabs to perform other tasks while your searches complete in the background.

You can configure a search to send you an email notification when the search is complete.

At any time while a search is in progress, you can return to the **Log Activity** or **Network Activity** tabs to view partial or complete search results.

At any time while a search is in progress, you can return to the **Log Activity** tab to view partial or complete search results.

Canceling a Search

While a search is queued or in progress, you can cancel the search on the **Manage Search Results** page.

If the search is in progress when you cancel it, the results that were accumulated until the cancellation are maintained.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. From the **Search** menu, select **Manage Search Results**.
3. Select the queued or in progress search result you want to cancel.
4. Click **Cancel**.

5. Click **Yes**.

Deleting a Search

If a search result is no longer required, you can delete the search result from the **Manage Search Results** page.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. Click the **Log Activity** tab.
3. From the **Search** menu, select **Manage Search Results**.
4. Select the search result that you want to delete.
5. Click **Delete**.
6. Click **Yes**.

RELATED DOCUMENTATION

[Managing Search Groups | 254](#)

[Search Example: Daily Employee Reports | 259](#)

[Using a Subsearch to Refine Search Results | 252](#)

Managing Search Groups

IN THIS SECTION

- [Viewing Search Groups | 255](#)
- [Creating a New Search Group | 256](#)
- [Editing a Search Group | 257](#)

- Copying a Saved Search to Another Group | 257
- Removing a Group or a Saved Search from a Group | 258

Using the **Search Groups** window, you can create and manage event, flow, and offense search groups.

These groups allow you to easily locate saved search criteria on the **Log Activity**, **Network Activity**, and **Offenses** tabs, and in the Report wizard.

Viewing Search Groups

A default set of groups and subgroups are available.

You can view search groups on the **Event Search Group**, **Flow Search Group**, or **Offense Search Group** windows.

All saved searches that are not assigned to a group are in the **Other** group.

The **Event Search Group**, **Flow Search Group**, and **Offense Search Group** windows display the following parameters for each group.

Table 43: Search Group Window Parameters

Parameter	Description
Name	Specifies the name of the search group.
User	Specifies the name of the user that created the search group.
Description	Specifies the description of the search group.
Date Modified	Specifies the date the search group was modified.

The **Event Search Group**, **Flow Search Group**, and **Offense Search Group** window toolbars provide the following functions.

Table 44: Search Group Window Toolbar Functions

Function	Description
New Group	To create a new search group, you can click New Group . See "Creating a New Search Group" on page 256 .
Edit	To edit an existing search group, you can click Edit . See "Editing a Search Group" on page 257 .
Copy	To copy a saved search to another search group, you can click Copy . See "Copying a Saved Search to Another Group" on page 257 .
Remove	To remove a search group or a saved search from a search group, select the item that you want to remove, and then click Remove . See "Removing a Group or a Saved Search from a Group" on page 258 .

1. Choose one of the following options:

- Click the **Log Activity** tab.
- Click the **Network Activity** tab.

2. Select **Search > Edit Search**.

3. Click **Manage Groups**.

4. View the search groups.

Creating a New Search Group

You can create a new search group.

1. Choose one of the following options:

- Click the **Log Activity** tab.
- Click the **Network Activity** tab.

2. Select **Search > Edit Search**.

3. Click **Manage Groups**.

4. Select the folder for the group under which you want to create the new group.
5. Click **New Group**.
6. In the **Name** field, type a unique name for the new group.
7. Optional. In the **Description** field, type a description.
8. Click **OK**.

Editing a Search Group

You can edit the **Name** and **Description** fields of a search group.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. Select **Search > Edit Search**.
3. Click **Manage Groups**.
4. Select the group that you want edit.
5. Click **Edit**.
6. Edit the parameters:
 - Type a new name in the **Name** field.
 - Type a new description in the **Description** field.
7. Click **OK**.

Copying a Saved Search to Another Group

You can copy a saved search to one or more groups.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.

2. Select **Search > Edit Search**.
3. Click **Manage Groups**.
4. Select the saved search that you want to copy.
5. Click **Copy**.
6. On the **Item Groups** window, select the check box for the group you want to copy the saved search to.
7. Click **Assign Groups**.

Removing a Group or a Saved Search from a Group

You can use the **Remove** icon to remove a search from a group or remove a search group.

When you remove a saved search from a group, the saved search is not deleted from your system. The saved search is removed from the group and automatically moved to the **Other** group.

You cannot remove the following groups from your system:

- Event Search Groups
- Flow Search Groups
- Offense Search Groups
- Other

You cannot remove Event Search Groups from your system.

1. Choose one of the following options:
 - Click the **Log Activity** tab.
 - Click the **Network Activity** tab.
2. Select **Search > Edit Search**.
3. Click **Manage Groups**.
4. Choose one of the following options:
 - Select the saved search that you want to remove from the group.
 - Select the group that you want to remove.

5. Click **Remove**.
6. Click **OK**.

RELATED DOCUMENTATION

[Search Example: Daily Employee Reports | 259](#)

[Using a Subsearch to Refine Search Results | 252](#)

[Managing Search Results | 253](#)

Search Example: Daily Employee Reports

The following example describes how to use a complex advanced search query to see specific employee information.

For identity management purposes, you decide to generate a daily report of the user activity in JSA. The report must include information about the employee, such as their user names, their serial number, their manager, and their activities.

An employee might have multiple user names in JSA. You use the RESTful API to build a reference map that returns all associated user names to the employee's name, `Global_User`. For the serial number and the manager's name, you create another reference data set and add it to the reference map.

Employee activities can range from login failures to JSA tasks, such as deleting objects. These events are recorded by JSA. By specifying the frequency of the events in the map, you can gauge when suspicious activity occurs. You group the data by the employee's name and the event name, and then sort the data by the highest event frequency within a 24-hour time frame.

To see this daily report, you log in to JSA console. In the Advanced Search text box on the **Log Activity** tab, you type the following search query:

```
select REFERENCEMAP('GlobalID_Mapping', username) as Global_User,
QIDNAME(qid) as 'Event Name', count(*) as 'Event Count', FIRST(username) as
UserId, REFERENCETABLE('employee_data', 'SerialNum', Global_user) as 'Serial
Number', REFERENCETABLE('employee_data', 'Manager', Global_User) as Manager
from events where (Global_User IS NOT NULL) GROUP BY Global_user, 'Event
Name' ORDER BY 'Event Count' DESC last 1 DAYS
```

RELATED DOCUMENTATION

[Using a Subsearch to Refine Search Results | 252](#)

[Managing Search Results | 253](#)

[Managing Search Groups | 254](#)

11

CHAPTER

Custom Event and Flow Properties

[Custom Event and Flow Properties | 262](#)

[Creating a Custom Property | 263](#)

[Modifying or Deleting a Custom Property | 265](#)

[Defining Custom Properties by Using Custom Property Expressions | 265](#)

[Use case: Create a report that uses Event Data that is not Normalized | 272](#)

Custom Event and Flow Properties

IN THIS SECTION

- [Extraction-based Properties | 262](#)
- [Calculation-based Properties | 263](#)
- [AQL-based Properties | 263](#)

JSA normalizes standard information that is parsed by the DSM, such as user names, IP addresses, and ports.

Some event sources send unique information that is not normalized. You can use custom properties to extract that data from the event or flow payload, and then use the non-normalized data in custom rules, searches, and reports.

The type of custom property that you create depends on the method that you want to use to define the non-normalized data in the payload.

Extraction-based Properties

Create an extraction-based property when you want to use a regex or JSON expression to parse the property values from the event or flow payloads.

For example, you have a report that shows all the users who changed other user's permissions on an Oracle server. The report uses normalized data to show the list of users who made the permission changes and the number of changes they made. The user account that was changed is not normalized and cannot be shown in the report. You can create a regex-based custom property to extract this information from the logs, and then use the property in searches and reports.

When the event or flow is parsed, the expression pattern is tested against each payload until the pattern matches. The first pattern to match the event or flow payload determines the data to be extracted.

When you define custom regex patterns, follow the regex rules as defined by the Java programming language. To learn more about regex rules, you can view regex tutorials on the web.

Calculation-based Properties

Create a calculation-based property when you want to do calculations on existing numeric event and flow properties. For example, you can create a calculation-based property that divides one numeric property by another numeric property to display a percentage value.

AQL-based Properties

Create an AQL-based property when you want to combine multiple extraction and calculation-based properties into a single property. For example, you can use AQL-based custom properties to combine extraction-based URLs, virus names, or secondary user names into a single property.

```
CONCAT( 'Src=', sourceip, ' | ', 'User=', username, ' | ', 'Domain=', DOMAINNAME(domainid) )
```

NOTE: The AQL expression can include AQL functions.

It does not support expressions that use SELECT, FROM, or database names.

You cannot use aggregate functions, such as SUM or GROUP, or other AQL-based custom properties.

RELATED DOCUMENTATION

[Creating a Custom Property | 263](#)

[Modifying or Deleting a Custom Property | 265](#)

[Defining Custom Properties by Using Custom Property Expressions | 265](#)

[Use case: Create a report that uses Event Data that is not Normalized | 272](#)

Creating a Custom Property

JSA includes a number of existing custom event properties that are not enabled or parsed by default. Ask your administrator to review the custom event property that you want to create to ensure that it does not exist.

To create custom event properties, you must have the **User Defined Event Properties** permission. To create custom flow properties, you must have the **User Defined Flow Properties** permission.

Users with administrative capabilities can create custom event and flow properties by selecting **Custom Event Properties** or **Custom Flow Properties** on the **Admin** tab.

Create a custom property to extract data that JSA does not typically show from the event or flow payloads. Custom properties must be enabled, and extraction-based custom properties must be parsed, before you can use them in rules, searches, reports, or for offense indexing.

Although multiple default custom properties might have the same name and the same log source, they can have different regex expressions, event names, or categories. For example, there are multiple custom properties for Microsoft Windows Security Event Log called **AccountName**, but each one is defined by a unique regex expression.

1. Click the **Log Activity** tab or the **Network Activity** tab.
2. If you are viewing the events or flows in streaming mode, click the **Pause** icon to pause streaming.
3. Double-click the event or flow that contains the data that you want to extract, and then click **Extract Property**.
4. In the **Property Type Selection** pane, select the type of custom property that you want to create.
5. Configure the custom property parameters.
Click the help icon to see information about the custom property parameters.
6. If you are creating an extraction-based custom property that is to be used in rules, search indexes, or forwarding profiles, ensure that the **Parse in advance for rules, reports, and searches** check box is selected.
7. Click **Test** to test the expression against the payload.
8. Click **Save**.

RELATED DOCUMENTATION

[Custom Event and Flow Properties | 262](#)

[Modifying or Deleting a Custom Property | 265](#)

[Defining Custom Properties by Using Custom Property Expressions | 265](#)

[Use case: Create a report that uses Event Data that is not Normalized | 272](#)

Modifying or Deleting a Custom Property

Edit a property when you want to change the property parameters, such as the regex expression or the log source type.

You can search for a specific property by using the Search properties field. The search is not case-sensitive.

Make a copy of a custom property when you want to change it, and then save it using a different name.

To delete a property, you must first remove all dependencies to it. Deleting a custom property does not delete the indexed property fields from the Ariel database.

1. Choose one of the following options:
 - To edit or delete a custom event property, click the **Log Activity** tab.
 - To edit or delete a custom flow property, click the **Network Activity** tab.
2. From the **Search** list box, select **Edit Search**.
3. Click **Manage Custom Properties**.
4. Select the property from the list, and click **Edit**, **Copy**, or **Delete**.
5. Make the required changes to the property, and then click **Save**.

RELATED DOCUMENTATION

[Custom Event and Flow Properties | 262](#)

[Creating a Custom Property | 263](#)

[Defining Custom Properties by Using Custom Property Expressions | 265](#)

[Use case: Create a report that uses Event Data that is not Normalized | 272](#)

Defining Custom Properties by Using Custom Property Expressions

Define a custom property for an event payload by using a custom property expression. Because JSON parsing begins when a valid JSON object is detected, the entire event does not need to be in JSON format. Similarly, LEEF and CEF parsing begins only when a valid LEEF/CEF message is detected within the event. Regex parsing runs through the entire payload.

JSA supports the following custom property expression types:

- Regex
- JSON
- LEEF
- CEF
- Name Value Pair
- Generic List
- XML

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.

1. Log in to JSA and click the **Admin** tab.
2. From the Data Sources section, click **Custom Event Properties**, and then click **Add**.
3. In the Property Type Selection section, select **Extraction Based**.
4. In the Test Field, enter the event payload that you want to use to test your custom property.
5. In the Property Definition section, complete the following steps:
 - a. If you're adding an expression to an existing property, select **Existing Property** and select a property from the list.
 - b. If you're defining a new property, select **New Property** and enter the name of the property.
 - c. To use the property for rules, reports and searches, select the **Parse in advance for rules, reports, and searches** check box.

You must select this check box to use the property for rules and indexes. Selecting the check box increases the efficiency of reports and searches, but you don't need to select it to use the property for reports and searches. When you select the check box, properties are parsed when the event is initially received and before it is stored. As a result, the loads are put on the event collection service.

- d. Select a **Field Type** for the property.

If you choose IP as the type for your custom property, JSA supports only IPv4.
 - e. Optional: Enter a description for the property.
6. In the Property Expression Definition section, complete the following steps:
 - a. Keep the **Enabled** check box selected; otherwise, clear the check box to disable the property.

- b. From the **Log Source Type** list, select a log source type for the property.
- c. If the expression is only evaluated against events for a specific log source, select the log source from the **Log Source** list. If you want it to be evaluated against all log sources, don't select.
- d. If the expression is only evaluated against events with a specific event name or QID, click the **Event Name** and browse for a QID to associate the expression with.
- e. If the expression is evaluated against any event with a specific low-level category, select Category, and select the **High Level Category** and then **Low Level Category** for the event.

NOTE: If the expression is evaluated for all events of the selected log source type and log source, ensure that you set the **Low Level Category** and **High Level Category** to Any.

- f. From the **Extraction using** field, select the extraction method to use for the property.

Table 45: Property Extraction Methods

Extraction Method	Valid Expression Form	Example
Regex	Enter the regex and the capture group number.	

Table 45: Property Extraction Methods *(Continued)*

Extraction Method	Valid Expression Form	Example
JSON Keypath	<p>A valid JSON expression is in the form:</p> <p><i>/"<name of top-level field>"</i></p> <p>For an event in a nested JSON format, a valid JSON expression is in the form:</p> <p><i>/"<name of top-level field>"/"<name of sub-level field>"/.../"<name of sub-level field_n>"</i></p> <p>To extract the 'user' field, type <i>/"user"</i> in the JsonKeypath field.</p> <p>To extract just the 'last_name' value from the 'user' subobject, type this expression:</p> <p><i>/"user"/"last_name"</i></p>	<p>The following example is a simple case of an event for a flat JSON record:</p> <pre>{ "action": "login", "user": "Firstname Lastname" }</pre> <p>The following example is a complex case of an event for a JSON record with nested objects:</p> <pre>{ "action": "login", "user": { "first_name": "Firstname", "last_name": "Lastname" } }</pre>

Table 45: Property Extraction Methods *(Continued)*

Extraction Method	Valid Expression Form	Example
LEEF Key	<p>Valid LEEF expressions are in the form of either a single key reference, or a special LEEF header field reference.</p> <p>To extract the 'usrName' property, type usrName in the LEEF Key field.</p> <p>The possible keys that can be extracted in these examples are:</p> <ol style="list-style-type: none"> i. - devTimeFormat ii. - devTime iii. - usrName iv. - name v. - authType vi. - src <p>To extract a header key property, type the key in the following format in the LEEF Key field:</p> <p>\$eventid\$</p> <p>The LEEF header values can be extracted by using the following expressions:</p> <ol style="list-style-type: none"> i. - \$leefversion\$ ii. - \$vendor\$ iii. - \$product\$ iv. - \$version\$ v. - \$eventid\$ 	<p>The following example is a simple case of an event that is formatted in LEEF V1.0:</p> <pre>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+0 200 usrName=flastname name= Firstname Lastname authType =interactive Password src=192.168.0.1</pre> <p>The following example is a simple case of an event that is formatted in LEEF V2.0 with the caret (^) separator character, and contains the same keys as the LEEF V1.0 example:</p> <pre>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+0 200 ^usrName=flastname^name= Firstname Lastname ^authType =interactive Password^src=192.168.0.1</pre>

Table 45: Property Extraction Methods *(Continued)*

Extraction Method	Valid Expression Form	Example
<p>CEF Key</p>	<p>Valid CEF expressions are in the form of either a single key reference, or a special CEF header field reference.</p> <p>To extract the 'cs1' property, type cs1 in the CEF Key field.</p> <p>The possible keys that can be extracted in the example are:</p> <ol style="list-style-type: none"> i. - start ii. - duser iii. - cs1 iv. - cs1Label v. - cs2 vi. - cs2Label vii. - src <p>To extract a header key property, type the key in the following format in the CEF Key field:</p> <p>\$id\$</p> <p>The CEF header values can be extracted by using the following expressions:</p> <ol style="list-style-type: none"> i. - \$cefversion\$ ii. - \$vendor\$ iii. - \$product\$ iv. - \$version\$ v. - \$id\$ 	<p>The following example shows an event that is in CEF format:</p> <pre> CEF:0 ABC Company SystemDefender 1.13 console_login Console Login 1 start=Oct 18 2017 11:26:03 duser=flastname cs1=Firstname Lastname cs1Label=Person Name cs2=interactivePassword cs2Label=authType src=192.168.0.1 </pre>

Table 45: Property Extraction Methods *(Continued)*

Extraction Method	Valid Expression Form	Example
	<p>vi. - \$name\$</p> <p>vii. - \$severity\$</p>	
Name Value Pair Key	Valid Name Value Pair expressions are in the form of a single key reference.	<p>The following example shows an event that is in Name Value Pair format:</p> <pre>Company=ABC Company;Product=SystemDefender; Version=1.13;EventID=console_login; Username=jsmith;Name=John Smith;authType=interactivePassword;</pre>
Generic List Keypath	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.	<p>The following example shows an event that is in Generic List format:</p> <pre>ABC Company;1.13;console_login;jsmith; John Smith;interactivePassword;</pre>
XML Key	<p>Valid XML expressions are in the form of a single key reference.</p> <p>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.</p>	<p>The following example shows an event that is in XML format:</p> <pre><EPOEvent><MachineInfo> <MachineName>NEPTUNE< / MachineName><MachineName >VALUE23</MachineName> <AgentGUID>9B-B5-A6-A8-37-B3< / AgentGUID><IPAddress someattrib ="someattribvalue">192.0.2.0 </ IPAddress><OSName>Windows 7</ OSName> <UserName>I am a test user <UserName></MachineInfo> </ EPOEvent></pre>

- g. If you chose the Numeric **Field Type** in the Property Definition section, select a number format in the **Extracted Number Format** field in the Format section to define any digit group separators for the locale of the custom property.
 - h. If you chose the Date/Time **Field Type** in the Property Definition section, enter a format in the **Extracted Date/Time Format** and **Locale** fields in the Format section to define the date and time for the locale of the custom property.
 - i. Click **Test** to test the property expression definition.
7. Click Save.

RELATED DOCUMENTATION

[Custom Event and Flow Properties | 262](#)

[Creating a Custom Property | 263](#)

[Modifying or Deleting a Custom Property | 265](#)

[Use case: Create a report that uses Event Data that is not Normalized | 272](#)

Use case: Create a report that uses Event Data that is not Normalized

You can use a custom property to extract data that is not normalized from a payload, and use that data to build a report. For example, you can build a report that is based on the interface information that is in Cisco ASA firewall deny messages.

In this example, we'll use the following sample Cisco ASA firewall events to demonstrate how to extract the interface value from the event payload, and then build a report that uses that data.

```
<162>Sep 02 2014 11:49:41: %ASA-2-106001: Inbound TCP connection denied
from 10.10.10.128/58826 to 10.11.11.11/9100 flags SYN on interface External
<162>Sep 02 2014 11:49:40: %ASA-2-106001: Inbound TCP connection denied
from 10.10.10.128/58826 to 10.11.11.11/9100 flags SYN on interface Loopback
<162>Sep 02 2014 11:49:17: %ASA-2-106001: Inbound TCP connection
denied from 10.10.10.128/58821 to 10.11.11.11/9100 flags SYN on interface Internal
```

1. ["Creating a Custom Property" on page 263.](#)

In the sample events above, you can see that the event payload includes the word interface followed by the value that you want to extract. To capture the interface information from the events above, create an extraction-based custom property and configure it to use the regex expression `interface\s(.*)\b`.

To ensure that the new custom property is available to use in a search, select the **Parse in advance for rules, reports, and searches** check box, and enable the custom property.

2. ["Event and Flow Searches" on page 212](#), and in the **Group By** field, select the new custom event property.

To ensure that the search results include only Cisco ASA events, add the log source as a quick filter option in the search parameters. Save the search criteria so that you can use it in a report. Assign the saved search to a group to make it easier to find later.

3. ["Creating Custom Reports" on page 325](#), and configure the graph content to use the new saved search.

If the report was not configured to run after saving, you can run the report immediately by selecting **ActionsRun Report**.

RELATED DOCUMENTATION

[Custom Event and Flow Properties | 262](#)

[Creating a Custom Property | 263](#)

[Modifying or Deleting a Custom Property | 265](#)

[Defining Custom Properties by Using Custom Property Expressions | 265](#)

12

CHAPTER

Rules

[Custom Rules in JSA | 275](#)

[Custom Rules | 277](#)

[Creating a Custom Rule | 279](#)

[Configuring an Event or Flow As False Positive | 284](#)

[Anomaly Detection Rules | 284](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

[Editing Building Blocks | 295](#)

[Rule Performance Visualization | 296](#)

Custom Rules in JSA

IN THIS SECTION

- [What Are Rules? | 275](#)
- [What Are Building Blocks? | 275](#)
- [How do Rules Work? | 276](#)
- [How is an Offense Created from a Rule? | 276](#)

Rules, sometimes called correlation rules are applied to events, flows, or offenses to search for or detect anomalies. If all the conditions of a test are met, the rule generates response.

What Are Rules?

Custom rules test events, flow, and offenses to detect unusual activity in your network. You create new rules by using AND and OR combinations of existing rule tests.

JSA Event Collectors gather events from local and remote sources, normalizes these events, and classifies them into low-level and high-level categories. For flows, JSA Flow Processor read packets from the wire or receive flows from other devices and then converts the network data to flow records. Each Event Processor processes events or flow data from the JSA Event Collectors. Event Processors examine and correlate the information to indicate behavioral changes or policy violations. The custom rules engine (CRE) processes events and compares them against defined rules to search for anomalies. When a rule condition is met, the Event Processor generates an action that is defined in the rule response. The CRE keeps track of the systems that are involved in incidents, contributes events to offenses, and generates notifications.

What Are Building Blocks?

A building block is a collection of tests that don't result in a response or an action.

A building block groups commonly used tests to build complex logic, so that it can be reused in rules. A building block often tests for IP addresses, privileged user names, or collections of event names. For

example, a building block can include the IP addresses of all DNS servers. Rules can then use this building block.

How do Rules Work?

JSA Event Collectors gather events from local and remote sources, normalize these events, and classify them into low-level and high-level categories. For flows, JSA Flow Processor read packets from the wire or receive flows from other devices and then converts the network data to flow records. Each Event Processor processes events or flow data from the JSA Event Collectors. Flow Processors examine and correlate the information to indicate behavioral changes or policy violations. The custom rules engine (CRE) processes events and compares them against defined rules to search for anomalies. When a rule condition is met, the Event Processor generates an action that is defined in the rule response. The CRE keeps track of the systems that are involved in incidents, contributes events to offenses, and generates notifications.

How is an Offense Created from a Rule?

JSA creates an offense when events, flows, or both meet the test criteria that is specified in the rules.

JSA analyzes the following information:

- Incoming events and flows
- Asset information
- Known vulnerabilities

The rule that created the offense determines the offense type.

The magistrate prioritizes the offenses and assigns the magnitude value based on several factors, including number of events, severity, relevance, and credibility.

NOTE: Building blocks are tested before rules are tested.

For example, you have a building block that is defined to trigger an offense on high magnitude events. The log activity can show that there were high magnitude events, but no offense was triggered. This can happen because when the building block was tested, the events was not at high magnitude. The magnitude of the event did not increase until the rules were tested.

One solution is to set a rule to check for the different in Severity, Credibility, and Relevance rather than to use a building block.

Custom Rules

IN THIS SECTION

- [What Are Custom Rules? | 277](#)
- [Rule Types | 277](#)
- [Managing Rules | 278](#)
- [Domain-specific Rules | 278](#)
- [Rule Conditions | 278](#)

JSA includes rules that detect a wide range of activities, including excessive firewall denials, multiple failed login attempts, and potential botnet activity. You can also create your own rules to detect unusual activity.

What Are Custom Rules?

Customize default rules to detect unusual activity in your network.

Rule Types

Each of the event, flow, common, and offense rule types test against incoming data from different sources in real time. There are multiple types of rule tests. Some check for simple properties from the data set. Other rule tests are more complicated. They track multiple, event, flow, and offense sequences over a period of time and use "counter" that is on one or more parameters before a rule response is triggered.

- **Event rules**--Test against incoming log source data that is processed in real time by the JSA Event Processor. You create an event rule to detect a single event or event sequences. For example, to

monitor your network for unsuccessful login attempts, access multiple hosts, or a reconnaissance event followed by an exploit, you create an event rule. It is common for event rules to create offenses as a response.

- **Flow rules**--Test against incoming flow data that is processed by the JSA Flow Processor. You can create a flow rule to detect a single flow or flow sequences. It is common for flow rules to create offenses as a response.
- **Common rules**--Test against event and flow data. For example, you can create a common rule to detect events and flows that have a specific source IP address. It is common for common rules to create offenses as a response.
- **Offense rules**--Test the parameters of an offense to trigger more responses. For example, a response generates when an offense occurs during a specific date and time. An offense rule processes offenses only when changes are made to the offense. For example, when new events are added, or the system scheduled the offense for reassessment. It is common for offense rules to email a notification as a response.

Managing Rules

You can create, edit, assign rules to groups, and delete groups of rules. By categorizing your rules or building blocks into groups, you can efficiently view and track your rules. For example, you can view all rules that are related to compliance.

Domain-specific Rules

If a rule has a domain test, you can restrict that 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 a response.

To create a rule that tests conditions across the entire system, set the domain condition to **Any Domain**.

Rule Conditions

Most rule tests evaluate a single condition, like the existence of an element in a reference data collection or testing a value against a property of an event. For complex comparisons, you can test event rules by building an Ariel Query Language (AQL) query with WHERE clause conditions. You can use all of the WHERE clause functions to write complex criteria that can eliminate the need to run numerous

individual tests. For example, use an AQL WHERE clause to check whether inbound SSL or web traffic is being tracked on a reference set.

You can run tests on the property of an event, flow, or offense, such as source IP address, severity of event, or rate analysis.

With functions, you can use building blocks and other rules to create a multi-event, multi-flow, or multi-offense function. You can connect rules by using functions that support Boolean operators, such as OR and AND. For example, if you want to connect event rules, you can use **when an event matches any/all of the following rules** function.

RELATED DOCUMENTATION

[Creating a Custom Rule | 279](#)

[Configuring an Event or Flow As False Positive | 284](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

Creating a Custom Rule

JSA includes rules that detect a wide range of activities, including excessive firewall denials, multiple failed login attempts, and potential botnet activity. You can also create your own rules to detect unusual activity.

Before you begin to create a new rule, you must have the **Offenses >Maintain Custom Rules** permission.

When you define rule tests, test against the smallest data possible. Testing in this way helps rule test performance and ensures that you don't create expensive rules. To optimize performance, start with broad categories that narrow the data that is evaluated by the rule test. For example, start with a rule test for a specific log source type, network location, flow source, or context (R2L, L2R, L2L). Any mid-level tests might include IP addresses, port traffic, or any other associated test. The rule must test payload and regex expression last.

Similar rules are grouped by category. For example, Audit, Exploit, DDoS, Recon, and more. When you delete an item from a group, the rule or building block is only deleted from the group; it remains available on the **Rules** page. When you delete a group, the rules or building blocks of that group remain available on the **Rules** page.

1. From the **Offenses**, **Log Activity**, or **Network Activity** tabs, click **Rules**.
2. From the **Display** list, select **Rules** to create a new rule.
3. From the **Display** list, select **Building Blocks** to create a new rule by using building blocks.

4. From the **Actions** list, select a rule type.

Each rule type tests against incoming data from different sources in real time. For example, event rules test incoming log source data and offense rules test the parameters of an offense to trigger more responses.

5. On the **Rule Test Stack Editor** page, in the **Rule** pane, type a unique name that you want to assign to this rule in the **Apply** text box.

6. From the list box, select **Local** or **Global**.

- If you select **Local**, all rules are processed on the Event Processor on which they were received and offenses are created only for the events that are processed locally.
- If you select **Global**, all matching events are sent to the JSA console for processing and therefore, the JSA console uses more bandwidth and processing resources.

Learn more about Local and Global rules:

Global rule tests

Use global rules to detect things like **multiple user login failures** where the events from that user might appear on multiple Event processors. For example, if you configured a **Local** rule for five login failures in 10 minutes from the same username, all 5 of those login failures must appear on the same Event Processor. Therefore, if three login failures were on one Event Processor and 2 were on another, no offense is generated. However, if you set this rule to **Global**, it generates an offense.

7. From the **Test Group** list, select one or more tests that you want to add to this rule. The CRE evaluates rule tests line-by-line in order. The first test is evaluated and when true, the next line is evaluated until the final test is reached.

Learn more about using rules for events that are not detected:

The following rule tests can be triggered individually, but subsequent rule tests in the same rule test stack are not acted upon.

- **when the event(s) have not been detected by one or more of these log source types for this many seconds**
- **when the event(s) have not been detected by one or more of these log sources for this many seconds**
- **when the event(s) have not been detected by one or more of these log source groups for this many seconds**

These rule tests are not activated by an incoming event, but instead are activated when a specific event is not seen for a specific time interval that you configured. JSA uses a watcher task that periodically queries the last time that an event was seen (last seen time), and stores this time for the event, for each log source. The rule is triggered when the difference between this last seen time and the current time exceeds the number of seconds that is configured in the rule.

8. To export the configured rule as a building block to use with other rules, click **Export as Building Block**.
9. On the **Rule Responses** page, configure the responses that you want this rule to generate.
Learn more about rule response page parameters:

Table 46: Event , Flow and Common Rule, and Offense Rule Response Page Parameters

Parameter	Description
Severity	Select this checkbox to assign a severity level to the event, where 0 is the lowest and 10 is the highest. The severity is displayed in the Annotation pane of the event details.
Credibility	Select this checkbox to assign credibility to the log source. For example, is the log source noisy or expensive? The range is 0 (lowest) to 10 (highest) and the default is 10. Credibility is displayed in the Annotation pane of the event details.
Relevance	Select this checkbox to assign relevance to the weight of the asset. For example, how much do you care about the asset? The range is 0 (lowest) to 10 (highest) and the default is 10. Relevance is displayed in the Annotation pane of the event details.
Bypass further rule correlation event	Select this checkbox to match an event or flow to bypass all other rules in the rule engine and prevent it from creating an offense. The event is written to storage for searching and reporting.
Dispatch New Event	Select this checkbox to dispatch a new event in addition to the original event or flow, which is processed like all other events in the system. Dispatches a new event with the original event, and is processed like all other events in the system. The Dispatch New Event parameters are displayed when you select this checkbox . By default, the checkbox is clear.
Email	Select this checkbox to change the Email Locale setting from the System Settings on the Admin tab.
Send to Local Syslog	Select this checkbox to log the event or flow locally. By default, this checkbox is clear. NOTE: Only normalized events can be logged locally on an appliance. If you want to send raw event data, you must use the Send to Forwarding Destinations option to send the data to a remote syslog host.

Table 46: Event , Flow and Common Rule, and Offense Rule Response Page Parameters (Continued)

Parameter	Description
Send to Forwarding Destinations	<p>Select this checkbox to log the event or flow on a forwarding destination.</p> <p>A forwarding destination is a vendor system, such as SIEM, ticketing, or alerting systems. When you select this checkbox, a list of forwarding destinations is displayed.</p> <p>To add, edit, or delete a forwarding destination, click the Manage Destinations link.</p>
Notify	<p>Select this checkbox to display events that are generated as a result of this rule in the System Notifications item on the Dashboard tab.</p> <p>If you enable notifications, configure the Response Limiter parameter.</p>
Add to Reference Set	<p>Select this checkbox to add events that are generated as a result of this rule to a reference set. You must be an administrator to add data to a reference set.</p> <p>To add data to a reference set, follow these steps:</p> <ol style="list-style-type: none"> a. From the first list, select the property of the event or flow that you want to add. b. From the second list, select the reference set to which you want to add the specified data.
Add to Reference Data	<p>To use this rule response, you must create the reference data collection.</p>
Remove from Reference Set	<p>Select this checkbox to remove data from a reference set.</p> <p>To remove data from a reference set:</p> <ol style="list-style-type: none"> a. From the first list box, select the property of the event or flow that you want to remove. Options include all normalized or custom data. b. From the second list box, select the reference set from which you want to remove the specified data. <p>The Remove from Reference Set rule response provides the following function:</p> <p>Refresh: Click Refresh to refresh the first list box to ensure that the list is current.</p>

Table 46: Event , Flow and Common Rule, and Offense Rule Response Page Parameters (Continued)

Parameter	Description
Remove from Reference Data	To use this rule response, you must have a reference data collection.
Execute Custom Action	Select this checkbox to write scripts that do specific actions in response to network events. For example, you might write a script to create a firewall rule that blocks a particular source IP address from your network in response to repeated login failures. You add and configure custom actions by using the Define Actions icon on the Admin tab.
Response Limiter	Select this checkbox to configure the frequency in which you want this rule to respond.

An SNMP notification might resemble the following example:

```
"Wed Sep 28 12:20:57 GMT 2005, Custom Rule Engine Notification - Rule 'SNMPTRAPTst' Fired. 172.16.20.98:0 ->
172.16.60.75:0 1, Event Name: ICMP Destination Unreachable Communication with Destination Host is
Administratively Prohibited, QID: 1000156, Category: 1014, Notes: Offense description"
```

A syslog output might resemble the following example:

```
Sep 28 12:39:01 localhost.localdomain ECS: Rule 'Name of Rule' Fired: 172.16.60.219:12642 ->
172.16.210.126:6666 6, Event Name: SCAN SYN FIN, QID: 1000398, Category: 1011, Notes: Event description
```

To verify that the event triggers the rule test based on your building block, you can create an email response.

RELATED DOCUMENTATION

[Configuring an Event or Flow As False Positive | 284](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

[Editing Building Blocks | 295](#)

Configuring an Event or Flow As False Positive

You might have legitimate network traffic that triggers false positive flows and events that makes it difficult to identify true security incidents. You can prevent events and flows from correlating into offenses by configuring them as false positives.

1. From the, **Log Activity**, or **Network Activity** tabs, click the pause on the upper right to stop real-time streaming of events or flows.
2. Select the event that you want to tune.
3. Click **False Positive**.
4. Select an event or flow property option.
5. Select a traffic direction option.
6. Click **Tune**.

The event or flow that matches the specified criteria will no longer correlates into offenses. To edit false positive tuning, use the **User-BB_FalsePositive: User Defined Positive Tunings building** block in the **Rules** section on the **Offenses** tab.

RELATED DOCUMENTATION

[Creating a Custom Rule | 279](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

[Editing Building Blocks | 295](#)

Anomaly Detection Rules

IN THIS SECTION

- [Anomaly Rules | 285](#)
- [Threshold Rules | 285](#)
- [Behavioral Rules | 286](#)
- [Creating an Anomaly Detection Rule | 288](#)

Anomaly detection rules test the results of saved flow or events searches to detect when unusual traffic patterns occur in your network.

Anomaly detection rules require a saved search that is grouped around a common parameter, and a time series graph that is enabled. Typically the search needs to accumulate data before the anomaly rule returns any result that identifies patterns for anomalies, thresholds, or behavior changes.

Anomaly Rules

Test event and flow traffic for changes in short-term events when you are comparing against a longer timeframe. For example, new services or applications that appear in a network, a web server crashes, firewalls that all start to deny traffic.

Example: You want to be notified when one of your firewall devices is reporting more often than it usually does because your network might be under attack. You want to be notified when you receive twice as many events in 1 hour. You follow these steps:

- Create and save a search that groups by log source, and displays only the count column.
- Apply the saved search to an anomaly rule, and add the rule test, and **when the average value (per interval) of count over the last 1 hour is at least 100% different from the average value (per interval) of the same property over the last 24 hours.**

Threshold Rules

Test events or flows for activity that is greater than or less than a specified range. Use these rules to detect bandwidth usage changes in applications, failed services, the number of users connected to a VPN, and detecting large outbound transfers.

Example: A user who was involved in a previous incident has large outbound transfer

When a user is involved in a previous offense, automatically set the Rule response to add to the Reference set. If you have a watch list of users, add them to the Reference set. Tune acceptable limits within the Threshold rule.

A reference set, WatchUsers, and Key:username are required for your search.

Complete the following search, and then apply it to a Threshold rule.

```
select assetuser(sourceip, now()) as 'srcAssetUser', Applicationname(applicationid)as 'AppName',
long(sum(sourcebytes +destinationbytes)) as 'flowsum' from flows where flowdirection = 'L2R' and
```

```
REFERENCESETCONTAINS('Watchusers', username)group by 'srcAssetUser', applicationid order by 'flowsum' desc last 24
hours
```

Behavioral Rules

Test events or flows for volume changes that occur in regular patterns to detect outliers. For example, a mail server that has an open relay and suddenly communicates with many hosts, or an IPS (intrusion protection system) that starts to generate numerous alert activities.

A behavioral rule learns the rate or volume of a property over a pre-defined season. The season defines the baseline comparison timeline for what you're evaluating. When you set a season of 1 week, the behavior of the property over that 1 week is learned and then you use rule tests to alert you to any significant changes.

After a behavioral rule is set, the seasons adjusts automatically. When the data in the season is learned, it is continually evaluated so that business growth is profiled within the season; you do not have to change your rules. The longer a behavioral rule runs, the more accurate it becomes. You can then adjust the rule responses to capture more subtle changes.

The following table describes the Behavioral rule test parameter options.

Table 47: Behavioral Rule Test Definitions

Rule test parameter	Description
Season	The most important value. The season defines the baseline behavior of the property that you are testing, and which the other rule tests use. To define a season, consider the type of traffic that you are monitoring. For example, for network traffic or processes that include human interaction, 1 week is a good season time frame. For tracking automated services where patterns are consistent, you might want to create a season as short as 1 day to define that pattern of behavior.
Current traffic level	Weight of the original data with seasonal changes and random error accounted for. This rule test asks the question, "Is the data the same as yesterday at the same time?"

Table 47: Behavioral Rule Test Definitions (Continued)

Rule test parameter	Description
Current traffic trend	Weight of changes in the data for each time interval. This rule test asks the question, "How much does the data change when it compares this minute to the minute before?"
Current traffic behavior	Weight of the seasonal effect for each period. This rule test asks the question, "Did the data increase the same amount from week 2 to week 3, as it did from week 1 to week 2?"
Predicted value	Use predicted values to scale baselines to make alerting more or less sensitive.

The forecast for value from the (n+1) th interval is calculated by using the following formula:

$$F_{n+1} = B_n + T_n + T_{n+1-s}$$

Where F is the predicted value, B is the base value for interval n, T is the trend value for interval n, and T is the trend value for season intervals ago and s is the number of intervals within the season.

The base value is calculated by using the following formula:

$$B_{n+1} = (0.2 + 0.3 * (\langle \text{Current traffic level} \rangle / 100.0)) * (\text{value}_{n+1} - T_{n+1-s}) + (1 - (0.2 + 0.3 * (\langle \text{Current traffic level} \rangle / 100.0))) * T_n$$

The trend value is calculated by using the following formula:

$$T_{n+1} = (0.2 + 0.3 * (\langle \text{Current traffic trend} \rangle / 100.0)) * (B_{n+1} - B_n) + (1 - (0.2 + 0.3 * (\langle \text{Current traffic trend} \rangle / 100.0))) * T_n$$

Smoothed deviation D is calculated by using the following formula:

$$D_{n+1} = (0.2 + 0.3 * (\langle \text{Current traffic level} \rangle / 100.0)) * |\text{value}_{n+1} - F_{n+1}| + (1 - (0.2 + 0.3 * (\langle \text{Current traffic level} \rangle / 100.0))) * D_{n+1-s}$$

The behavioral rule produces an alert for the interval if the following expression is false:

$$F - (1 + (\text{sensitivity} / 100.0) * 3) * D \leq \text{value} \leq F + (1 + (\text{sensitivity} / 100.0) * 3) * D$$

During the first season, the behavioral rule learns for future calculations and doesn't produce any alerts.

Creating an Anomaly Detection Rule

Anomaly detection rules test the result of saved flow or event searches to search for unusual traffic patterns that occur in your network. Behavioral rules test event and flow traffic according to "seasonal" traffic levels and trends. Threshold rules test event and flow traffic for activity less than, equal to, or greater than a configured threshold or within a specified range.

To create anomaly detection rules on the **Log Activity** tab, you must have the **Log Activity Maintain Custom Rules** role permission.

To create anomaly detection rules on the **Network Activity** tab, you must have the **Network Activity Maintain Custom Rules** role permission.

To manage default and previously created anomaly detection rules, use the **Rules** page on the **Offenses** tab.

When you create an anomaly detection rule, the rule is populated with a default test stack, based on your saved search criteria. You can edit the default tests or add tests to the test stack. At least one **Accumulated Property** test must be included in the test stack.

By default, the **Test the [Selected Accumulated Property] value of each [group] separately** option is selected on the **Rule Test Stack Editor** page.

An anomaly detection rule tests the selected accumulated property for each event or flow group separately. For example, if the selected accumulated value is **UniqueCount(sourceIP)**, the rule tests each unique source IP address for each event or flow group.

The **Test the [Selected Accumulated Property] value of each [group] separately** option is dynamic. The **[Selected Accumulated Property]** value depends on the option that you select for the **this accumulated property test** field of the default test stack. The **[group]** value depends on the grouping options that are specified in the saved search criteria. If multiple grouping options are included, the text might be truncated. Move your mouse pointer over the text to view all groups.

1. Click the **Log Activity** or **Network Activity** tab.
2. Perform an aggregated search.

You can add a property to the **group by** in a new historical search or select a property from the **Display** list on the current search page.

3. On the search result page, click **Configure**, and then configure the following options:
 - a. Select a property from the **Value to Graph** list.
 - b. Select **time series** as the chart type from the **Value to Graph** list
 - c. Enable the **Capture Time Series Data** check box.

- d. Click **Save**, and then enter a name for your search.
- e. Click **OK**.
- f. Select last 5 minutes from the **Time Range** list, while you wait for the time series graph to load.

You must have time series data for the property that you selected in the **Value to Graph** list to run a rule test on that accumulated property.

4. From the **Rules** menu, select the rule type that you want to create.
 - Add Anomaly Rule
 - Add Threshold Rule
 - Add Behavioral Rule
5. On the **Rule Test Stack Editor** page, in the **enter rule name here** field, type a unique name that you want to assign to this rule.

6. To apply your rule by using the default test, select the first rule in the anomaly **Test Group** list.

You might need to set the accumulated property parameter to the property that you selected from the **Value to Graph** list that you saved in the search criteria. If you want to see the result sooner, set the percentage to a lower value, such as 10%. Change **last 24 hours** to a lesser time period, such as 1 hour. Because an anomaly detection tests on aggregated fields in real time to alert you of anomalous network activity, you might want to increase or decrease events or flows in your network traffic.

7. Add a test to a rule.
 - a. To filter the options in the **Test Group** list, type the text that you want to filter for in the **Type to filter** field.
 - b. From the **Test Group** list, select the type of test that you want to add to this rule.
 - c. To identify a test as an excluded test, click **and** at the beginning of the test in the Rule pane. The **and** is displayed as **and not**.
 - d. Click the underlined configurable parameters to customize the variables of the test.
 - e. From the dialog box, select values for the variable, and then click **Submit**.
8. To test the total selected accumulated properties for each event or flow group, disable **Test the [Selected Accumulated Property] value of each [group] separately**.
9. In the groups pane, enable the groups you want to assign this rule to.
10. In the **Notes** field, type any notes that you want to include for this rule, and then Click **Next**.

11. On the **Rule Responses** page, configure the responses that you want this rule to generate.

Learn more about rule response page parameters for anomaly detection rules:

The following table provides the Rule Response page parameters if the rule type is Anomaly.

Table 48: Anomaly Detection Rule Response Page Parameters

Parameter	Description
Dispatch New Event	Specifies that this rule dispatches a new event with the original event or flow, which is processed like all other events in the system. By default, this check box is selected and cannot be cleared.
Offense Naming	<p>If you want the Event Name information to contribute to the name of the offense, select the This information should contribute to the name of the associated offense(s) option.</p> <p>If you want the configured Event Name to contribute to the offense, select the This information should set or replace the name of the associated offense(s).</p> <p>NOTE: After you replace the name of the offense, the name won't change until the offense is closed. For example, if an offense is associated with more than one rule, and the last event doesn't trigger the rule that is configured to override the name of the offense, the offense's name won't be updated by the last event. Instead, the offense name remains the name that is set by the override rule.</p>
Severity	The severity level that you want to assign to the event. The range is 0 (lowest) to 10 (highest) and the default is 5. The Severity is displayed in the Annotations pane of the event details.

Table 48: Anomaly Detection Rule Response Page Parameters (Continued)

Parameter	Description
Credibility	The credibility that you want to assign to the log source. For example, is the log source noisy or expensive? Using the list boxes, select the credibility of the event. The range is 0 (lowest) to 10 (highest) and the default is 5. Credibility is displayed in the Annotations pane of the event details.
Relevance	The relevance that you want to assign to the weight of the asset. For example, how much do you care about the asset? Using the list boxes, select the relevance of the event. The range is 0 (lowest) to 10 (highest) and the default is 5. Relevance is displayed in the Annotations pane of the event details.
Ensure that the dispatched event is part of an offense	As a result of this rule, the event is forwarded to the magistrate. If an offense exists, this event is added. If no offense was created on the Offenses tab, a new offense is created.
Notify	Events that generate as a result of this rule are displayed in the System Notifications item in the Dashboard tab. If you enable notifications, configure the Response Limiter parameter
Send to Local SysLog	Select this check box if you want to log the event or flow locally. By default, the check box is clear. NOTE: Only normalized events can be logged locally on a JSA appliance. If you want to send raw event data, you must use the Send to Forwarding Destinations option to send the data to a remote syslog host.

Table 48: Anomaly Detection Rule Response Page Parameters (Continued)

Parameter	Description
Add to Reference Set	<p>Adds events that are generated as a result of this rule to a reference set. You must be an administrator to add data to a reference set.</p> <p>To add data to a reference set, follow these steps:</p> <ol style="list-style-type: none"> a. From the first list, select the property of the event or flow that you want to add. b. From the second list, select the reference set to which you want to add the specified data.
Add to Reference Data	<p>To use this rule response, you must create the reference data collection.</p>
Remove from Reference Set	<p>If you want this rule to remove data from a reference set, select this check box.</p> <p>To remove data from a reference set, follow these steps:</p> <ol style="list-style-type: none"> a. From the first list, select the property of the event or flow that you want to remove. b. From the second list, select the reference set from which you want to remove the specified data.
Remove from Reference Data	<p>To use this rule response, you must have a reference data collection</p>

Table 48: Anomaly Detection Rule Response Page Parameters (Continued)

Parameter	Description
Execute Custom Action	<p>You can write scripts that do specific actions in response to network events. For example, you might write a script to create a firewall rule that blocks a particular source IP address from your network in response to repeated login failures.</p> <p>Select this check box and select a custom action from the Custom action to execute list.</p> <p>You add and configure custom actions by using the Define Actions icon on the Admin tab.</p>
Publish on the IF-MAP Server	If the IF-MAP parameters are configured and deployed in the system settings, select this option to publish the offense information about the IF-MAP server.
Response Limiter	Select this check box and use the list boxes to configure the frequency with which you want this rule to respond
Enable Rule	Select this check box to enable this rule. By default, the check box is selected.

An SNMP notification might resemble:

```
"Wed Sep 28 12:20:57 GMT 2005, Custom Rule Engine Notification Rule 'SNMPTRAPTst' Fired. 172.16.20.98:0 ->
172.16.60.75:0 1, Event Name: ICMP Destination Unreachable Communication with Destination Host is
Administratively Prohibited, QID: 1000156, Category: 1014, Notes: Offense description"
```

```
Sep 28 12:39:01 localhost.localdomain ECS: Rule 'Name of Rule' Fired: 172.16.60.219:12642 ->
172.16.210.126:6666 6, Event Name: SCAN SYN FIN, QID: 1000398, Category: 1011, Notes: Event description
```

12. Click **Next**.
13. Click **Finish**.

RELATED DOCUMENTATION

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

[Editing Building Blocks | 295](#)

[Configuring an Event or Flow As False Positive | 284](#)

Configuring a Rule Response to Add Data to a Reference Data Collection

Set up rules that use reference data to alert you to suspicious activity. For example, include a list of privileged users into reference data and then set up a rule that is triggered to alert you when privileged user anomalies occur.

Before you send data to a reference set, your JSA administrator must create the reference set.

JSA supports the following data collection types:

- **Reference set**--A set of elements, such as a list of IP addresses or user names, that are derived from events and flows that are occurring on your network.
- **Reference map**--Data is stored in records that map a key to a value. For example, to correlate user activity on your network, you create a reference map that uses the **Username** parameter as a key and the user's **Global ID** as a value.
- **Reference map of sets**--Data is stored in records that map a key to multiple values. For example, to test for authorized access to a patent, use a custom event property for **Patent ID** as the key and the **Username** parameter as the value. Use a map of sets to populate a list of authorized users.
- **Reference map of maps**--Data is stored in records that map one key to another key, which is then mapped to single value. For example, to test for network bandwidth violations, you create a map of maps. Use 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**--In a reference table, data is stored in a table that maps one key to another key, which is then mapped to single value. The second key has an assigned type. This mapping is similar to a database table where each column in the table is associated with a type. For example, you create a reference table that stores the **Username** parameter as the first key, and has multiple secondary keys that have a user-defined assigned type such as **IP Type** with the **Source IP** or **Source Port** parameter as a value. You can configure a rule response to add one or more keys that are defined in the table. You can also add custom values to the rule response. The custom value must be valid for the secondary key's type.

1. Create the reference data collection by using the **Reference Set Management** widget on the **Admin** tab.

You can also create a reference data collection by using the **ReferenceDataUtil.sh** script.

2. Create a rule by using the **Rules** wizard.
3. Create a rule response that sends data to a reference data collection. You can add the data as either shared data or domain-specific data.

Learn more about Add to Reference Data parameters:

Add to a Reference Map - Sends data to a collection of single key/multiple value pairs. You must select the key and value for the data record, and then select the reference map that you want to add the data record to.

Add to a Reference Map Of Sets - Sends data to a collection of key/single value pairs. You must select the key and the value for the data record, and then select the reference map of sets you want to add the data record to.

Add to a Reference Map Of Maps - Send data to a collection of multiple key/single value pairs. You must select a key for the first map, a key for the second map, and then the value for the data record. You must also select the reference map of maps you want to add the data record to.

Add to a Reference Table - Sends data to a collection of multiple key/single value pairs, where a type was assigned to the secondary keys. Select the reference table that you want to add data to, and then select a primary key. Select your inner keys (secondary keys) and their values for the data records.

RELATED DOCUMENTATION

[Editing Building Blocks | 295](#)

[Configuring an Event or Flow As False Positive | 284](#)

[Creating a Custom Rule | 279](#)

Editing Building Blocks

You can edit any of the default building blocks to use it in multiple rules or to build complex rules or logic. You can save a group of tests as building blocks for use with rules.

For example, you can edit the **BB:HostDefinition: Mail Servers** building block to identify all mail servers in your deployment. Then, you can configure any rule to exclude your mail servers from the rule tests.

1. Click the **Offenses** or **Network Activity** tab.

2. Click **Rules**.
3. From the **Display** list, select **Building Blocks**.
4. Double-click the building block that you want to edit.
5. Update the building block, as necessary.
6. Click **Next**.
7. Continue through the wizard.
8. Click **Finish**.

RELATED DOCUMENTATION

[Configuring an Event or Flow As False Positive | 284](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

[Creating a Custom Rule | 279](#)

Rule Performance Visualization

IN THIS SECTION

- [View the Metrics for a Rule | 297](#)
- [Colors and Bars in the Performance Column on the Rules Page | 297](#)

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 easily determine the efficiency of rules in the JSA pipeline, directly from the Rules page.

NOTE: Note: You must be an Administrator to turn on rule performance visualization. After rule performance visualization is turned on, users can view performance metrics for rules. For more information about turning on rule performance visualization, see the *Juniper Secure Analytics Administration Guide*.

When rule performance visualization is turned on, the **Performance** column is added to the Rules page. The **Performance** column is blank until a performance issue occurs in the custom rule engine.

When events or flows are routed to storage, JSA begins collecting metrics on enabled rules for efficiency measures. Metrics are collected on all event, common, and flow rules. When you save rule updates, the metrics are cleared for the rules that you updated to avoid any confusion around performance and updated rules. This option is configurable by an Administrator.

You can sort rules by their performance metrics and identify the more expensive rules. When you review the rules, you can adjust the tests to optimize each rule, and reduce the load on the system

With rule performance visualization, you see how expensive the rules are. JSA operations teams can monitor any expensive rules and ensure that they do not cause future performance issues.

By having rules run efficiently, the workload on the system can decrease. Over time, this efficiency can help JSA avoid any performance degradations around rules, which cause rules to bypass rule correlation. As a result, potential suspect activity might not trigger a notification, potentially missing future security-related issues.

For more information about tuning rules, see the *Juniper Secure Analytics Tuning Guide*.

View the Metrics for a Rule

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.

From the **Network Activity** tab or the **Log Activity** tab, click **Rules** to display the Rules page. Double-click a rule to open the **Rule Wizard**.

Colors and Bars in the Performance Column on the Rules Page

The number of bars that display is a visual aid for color blindness.

One red bar - The rule is under-performing and needs to be tuned. The EPS/FPS throughput for this rule is below the lower limit. Open the rule and tune the tests.

Two orange bars - The rule might need some tuning.

Three green bars - The rule has a high throughput above the upper limit of the EPS/FPS threshold.

NOTE: The colors and number of bars can't be changed. The definition of a rule that is underperforming is configurable by an Administrator.

RELATED DOCUMENTATION

[Creating a Custom Rule | 279](#)

[Configuring an Event or Flow As False Positive | 284](#)

[Configuring a Rule Response to Add Data to a Reference Data Collection | 294](#)

13

CHAPTER

Historical Correlation

[Historical Correlation | 300](#)

[Historical Correlation Overview | 301](#)

[Creating a Historical Correlation Profile | 302](#)

[Viewing Information About Historical Correlation Runs | 304](#)

Historical Correlation

Use historical correlation to run past events and flows through the custom rules engine (CRE) to identify threats or security incidents that already occurred.

NOTE: You cannot use historical correlation in Log Manager. For more information about the differences between JSA and Log Manager, see "[Capabilities in Your JSA Product](#)" on page 7.

By default, an JSA deployment analyzes information that is collected from log sources and flow sources in near real-time. With historical correlation, you can correlate by either the start time or the device time. *Start time* is the time that the event was received by JSA. *Device time* is the time that the event occurred on the device.

Historical correlation can be useful in the following situations:

- **Analyzing bulk data**--If you bulk load data into your JSA deployment, you can use historical correlation to correlate the data against data that was collected in real-time. For example, to avoid performance degradation during normal business hours, you load events from multiple log sources every night at midnight. You can use historical correlation to correlate the data by device time to see the sequence of network events as they occurred in the last 24 hours.
- **Testing new rules**--You can run historical correlation to test new rules. For example, one of your servers was recently attacked by new malware for which you do not have rules in place. You can create a rule to test for that malware. Then, you can use historical correlation to check the rule against historical data to see whether the rule would trigger a response if it were in place at the time of the attack. Similarly, you can use historical correlation to determine when the attack first occurred or the frequency of the attack. You can continue to tune the rule and then move it into a production environment.
- **Re-creating offenses that were lost or purged**--If your system lost offenses because of an outage or other reason, you can re-create the offenses by running historical correlation on the events and flows that came in during that time.
- **Identifying previously hidden threats**--As information becomes known about the latest security threats, you can use historical correlation to identify network events that already occurred but did not trigger an event. You can quickly test for threats that have already compromised your organization's system or data.

Historical Correlation Overview

IN THIS SECTION

- [Data Selection | 301](#)
- [Rule Selection and Handling | 301](#)
- [Offense Creation | 302](#)

You configure a historical correlation profile to specify the historical data that you want to analyze and the rule set that you want to test against. When a rule is triggered, an offense is created. You can assign the offense for investigation and remediation.

Data Selection

The profile uses a saved search to collect the historical event and flow data to use in the run. Ensure that your security profile grants permission to view the events and flows that you want to include in the historical correlation run.

Rule Selection and Handling

The JSA console processes data against only the rules that are specified in the historical correlation profile.

Common rules test data in both events and flows. You must have permission to view both events and flows before you can add common rules to the profile. When a profile is edited by a user who doesn't have permission to view both events and flows, the common rules are automatically removed from the profile.

You can include disabled rules in a historical correlation profile. When the profile runs, the disabled rule is evaluated against the incoming events and flows. If the rule is triggered, and the rule action is to generate an offense, the offense is created even when the rule is disabled. To avoid generating unnecessary distractions, rule responses, such as report generation and mail notifications, are ignored during historical correlation.

Because historical correlation processing occurs in a single location, the rules that are included in the profile are treated as global rules. The processing does not change the rule from local to global, but handles the rule as if it were global during the historical correlation run. Some rules, such as stateful rules, might not trigger the same response as they would in a normal correlation that is run on a local event processor. For example, a local stateful rule that tracks five failed logins in 5 minutes from the same user name behaves differently under normal and historical correlation runs. Under normal correlation, this local rule maintains a counter for the number of failed logins that are received by each local event processor. In historical correlation, this rule maintains a single counter for the entire JSA system. In this situation, offenses might be created differently compared to a normal correlation run.

Offense Creation

Historical correlation runs create offenses only when a rule is triggered and the rule action specifies that an offense must be created. A historical correlation run does not contribute to a real-time offense, nor does it contribute to an offense that was created from an earlier historical correlation run, even when the same profile is used.

The maximum number of offenses that can be created by a historical correlation run is 100. The historical correlation run stops when the limit is reached.

You can view historical offenses on the **Threat and Security Monitoring** dashboard and on the **Offenses** tab at the same time that you review real-time offenses.

RELATED DOCUMENTATION

[Creating a Historical Correlation Profile | 302](#)

[Viewing Information About Historical Correlation Runs | 304](#)

Creating a Historical Correlation Profile

You create a historical correlation profile to rerun past events and flows through the custom rules engine (CRE). The profile includes information about the data set and the rules to use during the run.

NOTE: You can create historical profiles only in JSA. You cannot create historical profiles in Log Manager.

Common rules test data in both events and flows. You must have permission to view both events and flows before you can add common rules to the profile. When a profile is edited by a user who doesn't have permission to view both events and flows, the common rules are automatically removed from the profile.

You can configure a profile to correlate by either start time or device time. *Start time* is the time when the events arrive at the event collector. *Device time* is the time that the event occurred on the device. Events can be correlated by start time or device time. Flows can be correlated by start time only.

You can include disabled rules in the profile. Rules that are disabled are indicated in the rules list with **(Disabled)** after the rule name.

A historical correlation run does not contribute to a real-time offense, nor does it contribute to an offense that was created from an earlier historical correlation run, even when the same profile is used.

NOTE: If you create too many historical correlation profiles that have many rules that are assigned to them, your offenses can be slow to load. If your offenses are slow to load, you can either delete unneeded profiles or edit them to have fewer rules.

1. Open the Historical Correlation dialog box.
 - On the **Log Activity** tab, click **Actions >Historical Correlation**.
 - On the **Network Activity** tab, click **Actions >Historical Correlation**.
 - On the **Offenses** tab, click **Rules >Actions >Historical Correlation**.
2. Click **Add** and select **Event Profile** or **Flow Profile**.
3. Type a name for the profile and select a saved search.
You can use only non-aggregated saved searches.
4. On the **Rules** tab, select the rules to be run against the historical data, and choose the correlation time.
If you select the **Use all enabled rules** check box, you cannot include disabled rules in the profile. If you want to include both enabled and disabled rules in the profile, you must select them individually from the rules list and click **Add Selected**.
5. On the **Schedule** tab, enter the time range for the saved search and set the profile schedule settings.
6. On the **Summary** tab, review the configuration and choose whether to run the profile immediately.
7. Click **Save**.
The profile is put into a queue to be processed. Queued profiles that are based on a schedule take priority over manual runs.

RELATED DOCUMENTATION

[Viewing Information About Historical Correlation Runs | 304](#)

[Historical Correlation Overview | 301](#)

Viewing Information About Historical Correlation Runs

View the history of a historical correlation profile to see information about past runs for the profile.

You can see the list of offenses that were created during the run and the catalog of events or flows that match the triggered rules in the profile. You can view the history for historical correlation runs that are queued, running, complete, complete with errors, and canceled.

For each rule in the profile that contributes to an offense, a catalog is created for each unique value of the property that the offense is indexed on. For each rule that does not contribute to an offense, a single catalog is created.

The following table shows how a historical correlation profile handles catalog creation under different scenarios. In each scenario, the catalog contains all the events or flows that either fully or partially match the triggered rule.

Table 49: Historical Correlation Catalog Examples

Scenario	Result
A rule generates offenses that are indexed on source IP address. The events that triggered the rule have three different source IP addresses.	The historical correlation profile creates three catalogs.
A rule generates offenses that are indexed on username. The events that triggered the rule have five different usernames.	The historical correlation profile creates five catalogs.
A rule is triggered, but the rule action does not create an offense.	The historical correlation profile creates a single catalog that includes all events that triggered the rule.

You cannot build reports on historical correlation data directly from JSA. If you want to use third-party programs to build reports, you can export the data from JSA.

1. Open the **Historical Correlation** dialog box.
 - On the **Log Activity** or **Network Activity** tab, click **Actions > Historical Correlation**.
 - On the **Offenses** tab, click **Rules >Actions >Historical Correlation**.
2. Select a profile and click **View History**.
 - a. If the historical correlation run status is **Completed** and the **Offense Count** is 0, the profile rules did not trigger any offenses.
 - b. If the historical correlation run created offenses, in the **Offense Count** column, click the link to see a list of the offenses that were created.

If only one offense was created, the offense summary is shown.
3. In the **Catalogs** column, click the links to see the list of events that either fully or partially match the profile rules.

The **StartTime** column in the event list represents the time that JSA received the event.
4. Click **Close**.

RELATED DOCUMENTATION

[Historical Correlation Overview | 301](#)

[Creating a Historical Correlation Profile | 302](#)

14

CHAPTER

Juniper Networks X-Force Integration

X-Force Integration in JSA | 307

X-Force Data on the Dashboard | 307

IP Address and URL Categories | 308

Confidence Factor and IP Address Reputation | 311

Searching Data from Juniper X-Force Exchange with Advanced Search Criteria |
313

X-Force Integration in JSA

Juniper Networks X-Force security experts use a series of international data centers to collect tens of thousands of malware samples, to analyze web pages and URLs, and to run analysis to categorize potentially malicious IP addresses and URLs. You can use this data to identify and remediate 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

X-Force Data on the Dashboard

The **Internet Threat Information Center** widget on the **Threat and Security Monitoring** dashboard uses X-Force data to provide up-to-date advisories on security issues, daily threat assessments, security news, and threat repositories.

The dashboard widget uses an embedded RSS feed to display X-Force data in the dashboard widget. The JSA console must have access to the Internet to receive data from the X-Force update server (www.iss.net).

The dashboard uses four AlertCon threat level images to provide a visual indicator of the current threat level.

Table 50: AlertCon Threat Levels

Level	Type	Description
1	Normal threats	Ordinary activity that compromises unprotected networks, minutes to hours after JSA connects to the Internet.
2	Increased vigilance	Vulnerabilities or online threats to computer networks that requires vulnerability assessment and corrective action.

Table 50: AlertCon Threat Levels (Continued)

Level	Type	Description
3	Focused attacks	Specific weakness and vulnerabilities that are the target of Internet attacks and require immediate defensive action.
4	Catastrophic threats	Critical security situations within a network that dictate an immediate and focused defensive action. This condition might be imminent or ongoing.

For more information about the current threat level, click the **Learn More** link to open the **Current Threat Activity** page on the Juniper X-Force Exchange website.

To view a summary of the current advisories, click the arrow icon next to the advisory. To investigate the full advisory, click the advisory link.

RELATED DOCUMENTATION

[IP Address and URL Categories | 308](#)

[Confidence Factor and IP Address Reputation | 311](#)

IP Address and URL Categories

IN THIS SECTION

- [Finding IP Address and URL Information in X-Force Exchange | 309](#)
- [Creating a URL Categorization Rule to Monitor Access to Certain Types Of Websites | 310](#)

X-Force Threat Intelligence categorizes IP address and URL information.

The IP addresses are grouped into the following categories:

- Malware hosts

- Spam sources
- Dynamic IP addresses
- Anonymous proxies
- Botnet Command and Control
- Scanning IP addresses

The X-Force Threat Intelligence feed also categorizes URL addresses.

Finding IP Address and URL Information in X-Force Exchange

Use right-click menu options in JSA to find information about IP addresses and URLs that is found on X-Force Exchange. You can use the information from your JSA searches, offenses, and rules to research further or to add information about IP addresses or URLs to an X-Force Exchange collection.

You can contribute either public or private information to track data in collections when you research security issues.

A *collection* is a repository where you store the information that is found during an investigation. You can use a collection to save X-Force Exchange reports, comments, or any other content. An X-Force Exchange report contains both a version of the report from the time when it was saved, and a link to the current version of the report. The collection contains a section that has a wiki-style notepad where you can add comments that are relevant to the collection.

1. To look up an IP address in X-Force Exchange from JSA, follow these steps:
 - a. Select the **Log Activity** or the **Network Activity** tab.
 - b. Right-click the IP address that you want to view in X-Force Exchange and select **More Options >Plugin Options >X-Force Exchange Lookup** to open the X-Force Exchange interface.
2. To look up a URL in X-Force Exchange from JSA, follow these steps:
 - a. Select either the **Offenses** tab, or the event details windows available on the **Offenses**.
 - b. Right-click the URL you want to look up in X-Force Exchange and select **>Plugin Options > X-Force Exchange Lookup** to open the X-Force Exchange interface.

Creating a URL Categorization Rule to Monitor Access to Certain Types Of Websites

You can create a rule that sends an email notification if users of the internal network access URL addresses that are categorized as gambling websites.

To use X-Force data in rules, your administrator must configure JSA to load data from the X-Force servers.

To create a new rule, you must have the **Offenses >Maintain Custom Rules** permission.

1. Click the **Offenses** tab.
2. On the navigation menu, click **Rules**.
3. From the **Actions** list, select **New Event Rule**.
4. Read the introductory text on the Rule wizard and click **Next**.
5. Click **Events** and click **Next**.
6. From the **Test Group** list box, select **X-Force Tests**.
7. Click the plus (+) sign beside the **when URL (custom) is categorized by X-Force as one of the following categories** test.
8. In the **enter rule name here** field in the Rule pane, type a unique name that you want to assign to this rule.
9. From the list box, select **Local** or **Global**.
10. Click the underlined configurable parameters to customize the variables of the test.
 - a. Click **URL (custom)**.
 - b. Select the URL property that contains the URL that was extracted from the payload and click **Submit**.
 - c. Click **one of the following categories**.
 - d. Select **Gambling / Lottery** from the X-Force URL categories, click **Add +** and click **Submit**.
11. To export the configured rule as a building block to use with other rules:
 - a. Click **Export as Building Block**.
 - b. Type a unique name for this building block.
 - c. Click **Save**.

12. On the Groups pane, select the check boxes of the groups to which you want to assign this rule.
13. In the **Notes** field, type a note that you want to include for this rule, and click **Next**.
14. On the **Rule Responses** page, click **Email** and type the email addresses that receive the notification.
15. Click **Next**.
16. If the rule is accurate, click **Finish**.

RELATED DOCUMENTATION

[Confidence Factor and IP Address Reputation | 311](#)

[Searching Data from Juniper X-Force Exchange with Advanced Search Criteria | 313](#)

Confidence Factor and IP Address Reputation

IN THIS SECTION

- [Determining a Threshold | 311](#)
- [Tuning False Positives with the Confidence Factor Setting | 312](#)

IP address reputation data is evaluated on the time that it is seen and the volume of messages or data. X-Force categorizes IP address reputation data and assigns a confidence factor value 0 - 100, where 0 represents no confidence and 100 represents certainty. For example, X-Force might categorize a source IP address as a scanning IP with a confidence factor of 75, which is a moderately high level of confidence.

Determining a Threshold

As an example, spam messages with an IP address reputation entry of 0 indicates that the source IP traffic is not spam, whereas an entry of 100 indicates definite spam traffic. Thus, values less than 50 indicate less probability that the message is spam, and values greater than 50 indicate more probability

that the message is spam. A value of 50 or higher is the threshold where you might consider action on a triggered rule.

These probabilities are based on ongoing web-based data that Juniper X-Force Threat Intelligence continuously collects and analyzes from around the world in X-Force data centers. As data is collected, the system evaluates how much spam is received from a particular IP address, or how frequently the flagged IP address is in the IP address reputation category. The more times, the higher the system scores the confidence factor.

Tuning False Positives with the Confidence Factor Setting

Use the confidence factor to limit the number of offenses that are created by triggered rules. Depending on the level of protection that you want, you adjust the confidence values to a level that best matches your network environment.

When you tune rules, consider a scale where 50 is the tipping point. On assets of lower importance, you might weigh an X-Force rule to trigger at a higher confidence factor for specific categories, like spam. For example, tuning a rule to a confidence factor of 75 means the rule triggers only when X-Force sees an IP address at or above a confidence factor of 75. This tuning reduces the number of offenses that are generated on lower priority systems and non-critical assets. However, an important system or critical business asset with a confidence factor of 50 triggers an offense at a lower level and brings attention to an issue more quickly.

For your DMZ, choose a higher confidence value such as 95% or higher. You do not need to investigate many offenses in this area. With a high confidence level, the IP addresses are more likely to match the category that is listed. If it is 95% certain that a host is serving malware, then you need to know about it.

For more secure areas of the network, like a server pool, lower the confidence value. More potential threats are identified and you spend less effort investigating because the threat pertains to a specific network segment.

For optimum false positive tuning, manage your rule triggers by segment. Look at your network infrastructure and decide which assets need a high level of protection, and which assets do not. You can apply different confidence values for the different network segments. Use building blocks for grouping commonly used tests so that they can be used in rules.

1. Click the **Log Activity** tab.
2. On the toolbar, click **Rules >Rules**.
3. Double-click a rule to start the Rule wizard.
4. In the filter box, type the following text:

when this host property is categorized by X-Force as this category with confidence value equal to this amount

5. Click the **Add test to rule (+)** icon.
6. In the Rule section, click the this amount link.
7. Enter a confidence value.
8. Click **Submit**.
9. Click **Finish** to exit the Rules wizard.

RELATED DOCUMENTATION

[Searching Data from Juniper X-Force Exchange with Advanced Search Criteria | 313](#)

[IP Address and URL Categories | 308](#)

Searching Data from Juniper X-Force Exchange with Advanced Search Criteria

For complex queries, you can search and filter data from X-Force Exchange by using Advanced Search expressions.

Advanced searches return data from the **Log Activity** or the **Network Activity** tab in JSA.

URL searches cannot be returned from the **Network Activity** tab because the URL information is provided by the event data.

1. Click the **Log Activity** tab.
2. On the **Search** toolbar, select the **Advanced Search**.
3. Type an AQL query expression.

NOTE: The following table describes some common search expressions.

Table 51: X-Force Advanced Search Expressions

Description	Example
Searches for source IP addresses that have a confidence factor above 50.	<pre>select * from events where XFORCE_IP_CONFIDENCE('Spam',sourceip)>50</pre>
Searches associated with a URL.	<pre>select url, XFORCE_URL_CATEGORY(url) as myCategories from events where XFORCE_URL_CATEGORY(url) IS NOT NULL</pre>
Searches associated with a source IP address.	<pre>select sourceip, XFORCE_IP_CATEGORY(sourceip) as IPcategories from events where XFORCE_IP_CATEGORY(sourceip) IS NOT NULL</pre>

4. Click **Search**.

RELATED DOCUMENTATION

[IP Address and URL Categories | 308](#)

[Confidence Factor and IP Address Reputation | 311](#)

15

CHAPTER

Report Management

Report Management | 316

Report Layout | 317

Chart Types | 317

Report Tab Toolbar | 320

Graph Types | 322

Creating Custom Reports | 325

Editing a Report | 330

Viewing Generated Reports | 330

Deleting Generated Content | 331

Manually Generating a Report | 332

Duplicating a Report | 333

Sharing a Report | 333

Branding Reports | 334

Report Groups | 335

Report Management

IN THIS SECTION

- [Timezone Considerations | 316](#)
- [Report Tab Permissions | 316](#)
- [Report Tab Parameters | 317](#)

You can use the **Reports** tab to create, edit, distribute, and manage reports.

Detailed, flexible reporting options satisfy your various regulatory standards, such as PCI compliance.

You can create your own custom reports or use a default reports. You can customize and rebrand default reports and distribute these to other users.

The **Reports** tab might require an extended period of time to refresh if your system includes many reports.

NOTE: If you are running Microsoft Exchange Server 5.5, unavailable font characters might be displayed in the subject line of emailed reports. To resolve this, download and install Service Pack 4 of Microsoft Exchange Server 5.5. For more information, contact Microsoft support.

Timezone Considerations

To ensure that the Reports feature uses the correct date and time for reporting data, your session must be synchronized with your timezone.

During the installation and setup of JSA products, the time zone is configured. Check with your administrator to ensure your JSA session is synchronized with your timezone.

Report Tab Permissions

Administrative users can view all reports that are created by other users.

Non-administrative users can view reports that they created only or reports that are shared by other users.

Report Tab Parameters

The **Reports** tab displays a list of default and custom reports.

From the **Reports** tab, you can view statistical information about the reports template, perform actions on the report templates, view the generated reports, delete generated content.

If a report does not specify an interval schedule, you must ["Manually Generating a Report" on page 332](#).

You can point your mouse over any report to preview a report summary in a tooltip. The summary specifies the report configuration and the type of content the report generates.

Report Layout

A report can consist of several data elements and can represent network and security data in various styles, such as tables, line charts, pie charts, and bar charts.

When you select the layout of a report, consider the type of report you want to create. For example, do not choose a small chart container for graph content that displays many objects. Each graph includes a legend and a list of networks from which the content is derived; choose a large enough container to hold the data. To preview how each chart displays a data, see ["Graph Types" on page 322](#).

RELATED DOCUMENTATION

[Chart Types | 317](#)

[Report Tab Toolbar | 320](#)

[Graph Types | 322](#)

Chart Types

When you create a report, you must choose a chart type for each chart you include in your report.

The chart type determines how the data and network objects appear in your report.

You can use any of the following types of charts:

Table 52: Chart Types

Chart Type	Description
None	Use this option if you need white space in your report. If you select the None option for any container, no further configuration is required for that container.
Asset Vulnerabilities	Use this chart to view vulnerability data for each defined asset in your deployment. You can generate Asset Vulnerability charts when vulnerabilities have been detected by a VA scan. This chart is available after you install JSA Vulnerability Manager.
Connections	This chart option is only displayed if you purchased and licensed JSA Risk Manager. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i> .
Device Rules	This chart option is only displayed if you purchased and licensed JSA Risk Manager. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i> .
Device Unused Objects	This chart option is only displayed if you purchased and licensed JSA Risk Manager. For more information, see the <i>Juniper Secure Analytics Risk Manager User Guide</i> .
Events/Logs	Use this chart to view event information. You can base a chart on data from saved searches on the Log Activity tab. You can configure the chart to plot data over a configurable period of time to detect event trends. For more information about saved searches, see Searches

Table 52: Chart Types (Continued)

Chart Type	Description
Log Sources	Use this chart to export or report on log sources. Select the log sources and log source groups that you want to appear in the report. Sort log sources by report columns. Include log sources that are not reported for a defined time period. Include log sources that were created in a specified time period.
Flows	Use this chart to view flow information. You can base a chart on data from saved searches on the Network Activity tab. You can configure the chart to plot flow data over a configurable period of time to detect flow trends. For more information about saved searches, see Searches
Top Destination IPs	Use this chart to display the top destination IPs in the network locations you select.
Top Offenses	Use this chart to display the top offenses that occur at present time for the network locations you select.
Offenses Over Time	Use this chart to display all offenses that have a start time within a defined time span for the network locations you select.
Top Source IPs	Use this chart to display and sort the top offense sources (IP addresses) that attack your network or business assets.
Vulnerabilities	The Vulnerabilities option is only displayed when the JSA Vulnerability Manager was purchased and licensed. For more information, see the <i>Juniper Secure Analytics Vulnerability Manager User Guide</i> .

Table 53: Chart Types

Chart Type	Description
None	Use this option if you need white space in your report. If you select the None option for any container, no further configuration is required for that container.
Asset Vulnerabilities	Use this chart to view vulnerability data for each defined asset in your deployment. You can generate Asset Vulnerability charts when vulnerabilities have been detected by a VA scan. This chart is available after you install JSA Vulnerability Manager.
Vulnerabilities	The Vulnerabilities option is only displayed when the JSA Vulnerability Manager was purchased and licensed. For more information, see the <i>Juniper Secure Analytics Vulnerability Manager User Guide</i> .

RELATED DOCUMENTATION

[Report Tab Toolbar | 320](#)

[Graph Types | 322](#)

[Creating Custom Reports | 325](#)

Report Tab Toolbar

You can use the toolbar to perform a number of actions on reports.

The following table identifies and describes the Reports toolbar options.

Table 54: Report Toolbar Options

Option	Description
Group	
Manage Groups	Click Manage Groups to manage " Report Groups " on page 335. Using the Manage Groups feature, you can organize your reports into functional groups. You can share report groups with other users.
Actions	<p>Click Actions to perform the following actions:</p> <ul style="list-style-type: none"> • Create— Select this option to create a new report. • Edit— Select this option to edit the selected report. You can also double-click a report to edit the content. • Duplicate— Select this option to "Duplicating a Report" on page 333 the selected report. • Assign Groups— Select this option to assign the selected report to a "Assign a Report to a Group" on page 338. • Share— Select this option to share the selected report with other users. You must have administrative privileges to "Sharing a Report" on page 333. • Toggle Scheduling— Select this option to toggle the selected report to the Active or Inactive state. • Run Report— Select this option to "Manually Generating a Report" on page 332. To generate multiple reports, hold the Control key and click on the reports you want to generate. • Run Report on Raw Data— Select this option to generate the selected report using raw data. This option is useful when you want to generate a report before the required accumulated data is available. For example, if you want to run a weekly report before a full week has elapsed since you created the report, you can generate the report using this option. • Delete Report— Select this option to delete the selected report. To delete multiple reports, hold the Control key and click on the reports you want to delete. • Delete Generated Content— Select this option to delete all generated content for the selected rows. To delete multiple generated reports, hold the Control key and click on the generate reports you want to delete.

Table 54: Report Toolbar Options (*Continued*)

Option	Description
Hide Interactive Reports	Select this check box to hide inactive report templates. The Reports tab automatically refreshes and displays only active reports. Clear the check box to show the hidden inactive reports.
Search Reports	<p>Type your search criteria in the Search Reports field and click the Search Reports icon. A search is run on the following parameters to determine which match your specified criteria:</p> <ul style="list-style-type: none"> • Report Title • Report Description • Report Group • Report Groups • Report Author User Name

RELATED DOCUMENTATION

[Graph Types | 322](#)

[Creating Custom Reports | 325](#)

[Editing a Report | 330](#)

Graph Types

Each chart type supports various graph types that you can use to display data.

The network configuration files determine the colors that the charts use to depict network traffic. Each IP address is depicted by using a unique color. The following table provides examples of how network and security data is used in charts. The table describes the chart types that are available for each type of graph.

Table 55: Graph Types

Graph type	Available chart types
Line	<ul style="list-style-type: none"> • Events/Logs • Flows • Connections • Vulnerabilities
Stacked Line	<ul style="list-style-type: none"> • Events/Logs • Flows • Connections • Vulnerabilities
Bar	<ul style="list-style-type: none"> • Events/Logs • Flows • Asset Vulnerabilities Connections • Connections • Vulnerabilities
Horizontal Bar	<ul style="list-style-type: none"> • Top Source IPs • Top Offenses • Offenses Over Time • Top Destination IPs
Stacked Bar	<ul style="list-style-type: none"> • Events/Logs • Flows • Connections

Table 55: Graph Types (Continued)

Graph type	Available chart types
Pie	<ul style="list-style-type: none"> • Events/Logs • Flows • Asset Vulnerabilities • Connections • Vulnerabilities
Table	<ul style="list-style-type: none"> • Events/Logs • Flows • Top Source IPs • Top Offenses • Offenses Over Time • Top Destination IPs • Connections • Vulnerabilities <p>To display content in a table, you must design the report with a full page width container.</p>
Aggregate Table	<p>Available with the Asset Vulnerabilities chart.</p> <p>To display content in a table, you must design the report with a full page width container.</p>

The following graph types are available for JSA Log Manager reports:

- Line
- Stacked Line
- Bar
- Stacked Bar

- Pie
- Table

NOTE: When you create bar and stacked bar graph reports, the legend is presented in a fixed format and the bars or bar sections are represented by color coded labels in most cases. If you select time as the value for the x axis, you can create time intervals on the x axis.

RELATED DOCUMENTATION

[Creating Custom Reports | 325](#)

[Editing a Report | 330](#)

[Viewing Generated Reports | 330](#)

Creating Custom Reports

You must have appropriate network permissions to share a generated report with other users.

For more information about permissions, see the *Juniper Secure Analytics Administration Guide*.

Use the Report wizard to create and customize a new report.

The Report wizard provides a step-by-step guide on how to design, schedule, and generate reports.

The wizard uses the following key elements to help you create a report:

- **Layout**— Position and size of each container
- **Container**— Placeholder for the featured content
- **Content**— Definition of the chart that is placed in the container

After you create a report that generates weekly or monthly, the scheduled time must elapse before the generated report returns results. For a scheduled report, you must wait the scheduled time period for the results to build. For example, a weekly search requires seven days to build the data. This search will return results after 7 days.

When you specify the output format for the report, consider that the file size of generated reports can be one to 2 megabytes, depending on the selected output format. PDF format is smaller in size and does not use a large quantity of disk storage space.

1. Click the **Reports** tab.
2. From the **Actions** list box, select **Create**.
3. On the **Welcome to the Report wizard!** window, click **Next**.
4. Select one of the following options:

Option	Description
Manually	By default, the report generates 1 time. You can generate the report as often as you want.
Hourly	<p>Schedules the report to generate at the end of each hour. The data from the previous hour is used.</p> <p>From the list boxes, select a time frame to begin and end the reporting cycle. A report is generated for each hour within this time frame. Time is available in half-hour increments. The default is 1:00 a.m for both the From and To fields.</p>
Daily	<p>Schedules the report to generate at the end of each day. The data from the previous day is used.</p> <p>From the list boxes, select the time and the days of the week that you want the report to run.</p>
Weekly	<p>Schedules the report to generate weekly using the data from the previous calendar week, from Monday to Sunday.</p> <p>Select the day that you want to generate the report. The default is Monday. From the list box, select a time to begin the reporting cycle. Time is available in half-hour increments. The default is 1:00 a.m.</p>
Monthly	<p>Schedules the report to generate monthly using the data from the previous calendar month.</p> <p>From the list box, select the date that you want to generate the report. The default is the first day of the month. Select a time to begin the reporting cycle. Time is available in half-hour increments. The default is 1:00 a.m.</p>

5. In the **Allow this report to generate manually** pane, **Yes** or **No**.
6. Configure the layout of your report:
 - a. From the **Orientation** list box, select **Portrait** or **Landscape** for the page orientation.
 - b. Select one of the six layout options that are displayed on the Report wizard.
 - c. Click **Next** .
7. Specify values for the following parameters:

Parameter	Values
Report Title	The title can be up to 60 characters in length. Do not use special characters.
Logo	From the list box, select a logo.
Pagination Options	From the list box, select a location for page numbers to display on the report. You can choose not to have page numbers display.
Report Classification	Type a classification for this report. You can type up to 75 characters in length. You can use leading spaces, special characters, and double byte characters. The report classification displays in the header and footer of the report. You might want to classify your report as confidential , highly confidential , sensitive , or internal .

8. Configure each container in the report:
 - a. From the **Chart Type** list box, select a chart type.
 - b. On the **Container Details** window, configure the chart parameters.

NOTE: You can also create asset saved searches. From the **Search to use** list box, select your saved search.

- c. Click **Save Container Details**.
- d. If you selected more than one container, repeat steps a to c.
- e. Click **Next** .

9. Preview the **Layout Preview** page, and then click **Next**.
10. Select the check boxes for the report formats you want to generate, and then click **Next**.

NOTE: Extensible Markup Language is only available for tables.

11. Select the distribution channels for your report, and then click **Next**. Options include the following distribution channels:

Option	Description
Report Console	Select this check box to send the generated report to the Reports tab. Report Console is the default distribution channel.
Select the users that should be able to view the generated report.	This option displays after you select the Report Console check box. From the list of users, select the users that you want to grant permission to view the generated reports.
Select all users	This option is only displayed after you select the Report Console check box. Select this check box if you want to grant permission to all users to view the generated reports. You must have appropriate network permissions to share the generated report with other users.
Email	Select this check box if you want to distribute the generated report by email.
Enter the report distribution email address(es)	This option is only displayed after you select the Email check box. Type the email address for each generated report recipient; separate a list of email addresses with commas. The maximum characters for this parameter are 255.

(Continued)

Option	Description
Include Report as attachment (non-HTML only)	This option is only displayed after you select the Email check box. Select this check box to send the generated report as an attachment.
Include link to Report Console	This option is only displayed after you select the Email check box. Select this check box to include a link to the Report Console in the email.

12. On the **Finishing Up** page, enter values for the following parameters.

Option	Description
Report Description	Type a description for this report. The description is displayed on the Report Summary page and in the generated report distribution email.
Please select any groups you would like this report to be a member of	Select the groups to which you want to assign this report. For more information about groups, see "Report Groups" on page 335 .
Would you like to run the report now?	Select this check box if you want to generate the report when the wizard is complete. By default, the check box is selected.

13. Click **Next** to view the report summary.
14. On the **Report Summary** page, select the tabs available on the summary report to preview your report configuration.

The report immediately generates. If you cleared the **Would you like to run the report now** check box on the final page of the wizard, the report is saved and generates at the scheduled time. The report title is the default title for the generated report. If you reconfigure a report to enter a new report title, the report is saved as a new report with the new name; however, the original report remains the same.

RELATED DOCUMENTATION

Editing a Report | 330

[Viewing Generated Reports | 330](#)

[Deleting Generated Content | 331](#)

Editing a Report

Using the Report wizard, you can edit any default or custom report to change.

You can use or customize a significant number of default reports. The default **Reports** tab displays the list of reports. Each report captures and displays the existing data.

NOTE: When you customize a scheduled report to generate manually, select the time span **End Date** before you select the **Start Date**.

1. Click the **Reports** tab.
2. Double-click the report that you want to customize.
3. On the Report wizard, change the parameters to customize the report to generate the content you require.

If you reconfigure a report to enter a new report title, the report is saved as a new report with the new name; however, the original report remains the same.

RELATED DOCUMENTATION

[Viewing Generated Reports | 330](#)

[Deleting Generated Content | 331](#)

[Manually Generating a Report | 332](#)

Viewing Generated Reports

On the **Reports** tab, an icon is displayed in the **Formats** column if a report has generated content. You can click the icon to view the report.

When a report has generated content, the **Generated Reports** column displays a list box. The list box displays all generated content, which is organized by the time-stamp of the report. The most recent

reports are displayed at the top of the list. If a report has no generated content, the **None** value is displayed in the **Generated Reports** column.

Icons representing the report format of the generated report are displayed in the **Formats** column.

Reports can be generated in PDF, HTML, RTF, XML, and XLS formats.

NOTE: The XML and XLS formats are available only for reports that use a single chart table format (portrait or landscape).

You can view only the reports to which you have been given access from the administrator. Administrative users can access all reports.

If you use the Mozilla Firefox web browser and you select the RTF report format, the Mozilla Firefox web browser starts a new browser window. This new window launch is the result of the Mozilla Firefox web browser configuration and does not affect JSA. You can close the window and continue with your JSA session.

1. Click the **Reports** tab.
2. From the list box in the **Generated Reports** column, select the time-stamp of report you want to view.
3. Click the icon for the format you want to view.

RELATED DOCUMENTATION

[Deleting Generated Content | 331](#)

[Manually Generating a Report | 332](#)

[Duplicating a Report | 333](#)

Deleting Generated Content

When you delete generated content, all reports that have generated from the report template are deleted, but the report template is retained.

1. Click the **Reports** tab.
2. Select the reports for which you want to delete the generated content.
3. From the **Actions** list box, click **Delete Generated Content**.

RELATED DOCUMENTATION

[Manually Generating a Report | 332](#)

[Duplicating a Report | 333](#)

[Sharing a Report | 333](#)

Manually Generating a Report

A report can be configured to generate automatically, however, you can manually generate a report at any time.

While a report generates, the Next Run Time column displays one of the three following messages:

- **Generating**— The report is generating.
- **Queued (position in the queue)**— The report is queued for generation. The message indicates the position that the report is in the queue. For example, 1 of 3.
- **(x hour(s) x min(s) y sec(s))**— The report is scheduled to run. The message is a count-down timer that specifies when the report will run next.

You can select the **Refresh** icon to refresh the view, including the information in the **Next Run Time** column.

1. Click the **Reports** tab.
2. Select the report that you want to generate.
3. Click **Run Report**.

After the report generates, you can "[Viewing Generated Reports](#)" on page 330 from the Generated Reports column.

RELATED DOCUMENTATION

[Duplicating a Report | 333](#)

[Sharing a Report | 333](#)

[Branding Reports | 334](#)

Duplicating a Report

To create a report that closely resembles an existing report, you can duplicate the report that you want to model, and then customize it.

1. Click the **Reports** tab.
2. Select the report that you want to duplicate.
3. From the **Actions** list box, click **Duplicate**.
4. Type a new name, without spaces, for the report.

You can "[Editing a Report](#)" on page 330 the duplicated report.

RELATED DOCUMENTATION

[Sharing a Report](#) | 333

[Branding Reports](#) | 334

[Report Groups](#) | 335

Sharing a Report

You can share reports with other users. When you share a report, you provide a copy of the selected report to another user to edit or schedule.

Any updates that the user makes to a shared report does not affect the original version of the report.

You must have administrative privileges to share reports. Also, for a new user to view and access reports, an administrative user must share all the necessary reports with the new user.

You can only share the report with users that have the appropriate access.

1. Click the **Reports** tab.
2. Select the reports that you want to share.
3. From the **Actions** list box, click **Share**.
4. From the list of users, select the users with whom you want to share this report.

RELATED DOCUMENTATION

[Branding Reports | 334](#)

[Report Groups | 335](#)

[Duplicating a Report | 333](#)

Branding Reports

Ensure that the graphic you want to use is 144 x 50 pixels with a white background.

To make sure that your browser displays the new logo, clear your browser cache.

To brand reports, you can import logos and specific images. To brand reports with custom logos, you must upload and configure the logos before you begin using the Report wizard.

Report branding is beneficial for your enterprise if you support more than one logo. When you upload an image, the image is automatically saved as a Portable Network Graphic (PNG).

When you upload a new image and set the image as your default, the new default image is not applied to reports that have been previously generated. Updating the logo on previously generated reports requires you to manually generate new content from the report.

If you upload an image that is larger in length than the report header can support, the image automatically resizes to fit the header; this is approximately 50 pixels in height.

1. Click the **Reports** tab.
2. On the navigation menu, click **Branding**.
3. Click **Browse** to browse the files that are located on your system.
4. Select the file that contains the logo you want to upload. Click **Open**.
5. Click **Upload Image**.
6. Select the logo that you want to use as the default and click **Set Default Image**.

RELATED DOCUMENTATION

[Report Groups | 335](#)

[Duplicating a Report | 333](#)

[Sharing a Report | 333](#)

Report Groups

IN THIS SECTION

- [Creating a Report Group | 336](#)
- [Editing a Group | 336](#)
- [Sharing Report Groups | 337](#)
- [Assign a Report to a Group | 338](#)
- [Copying a Report to Another Group | 338](#)
- [Removing a Report | 339](#)

You can sort reports into functional groups. If you categorize reports into groups, you can efficiently organize and find reports.

For example, you can view all reports that are related to Payment Card Industry Data Security Standard (PCIDSS) compliance.

By default, the **Reports** tab displays the list of all reports, however, you can categorize reports into groups such as:

- Compliance
- Executive
- Log Sources
- Network Management
- Security
- VoIP
- Other

When you create a new report, you can assign the report to an existing group or create a new group. You must have administrative access to create, edit, or delete groups.

For more information about user roles, see the *Juniper Secure Analytics Administration Guide*.

Creating a Report Group

You can create new groups.

1. Click the **Reports** tab.
2. Click **Manage Groups**.
3. Using the navigation tree, select the group under which you want to create a new group.
4. Click **New Group**.
5. Enter values for the following parameters:
 - **Name** Type the name for the new group. The name can be up to 255 characters in length.
 - **Description** Optional. Type a description for this group. The description can be up to 255 characters in length.
6. Click **OK**.
7. To change the location of the new group, click the new group and drag the folder to the new location on the navigation tree.
8. Close the **Report Groups** window.

Editing a Group

You can edit a report group to change the name or description.

1. Click the **Reports** tab.
2. Click **Manage Groups**.
3. From the navigation tree, select the group that you want to edit.
4. Click **Edit**.
5. Update values for the parameters, as necessary:
 - **Name** Type the name for the new group. The name can be up to 255 characters in length.
 - **Description** Optional. Type a description for this group. The description can be up to 255 characters in length. This field is optional.
6. Click **OK**.

7. Close the **Report Groups** window.

Sharing Report Groups

You can share report groups with other users.

You must have administrative permissions to share a report group with other users.

For more information about permissions, see the *Juniper Secure Analytics Administration Guide*.

You cannot use the Content Management Tool (CMT) to share report groups.

For more information about the CMT, see the *Juniper Secure Analytics Administration Guide*.

On the **Report Groups** window, shared users can see the report group in the report list.

To view a generated report, the user must have permission to view the report.

1. Click the **Reports** tab.
2. On the **Reports** window, click **Manage Groups**.
3. On the **Report Groups** window, select the report group that you want to share and click **Share**.
4. On the **Sharing Options** window, select one of the following options.

Option	Description
Default (inherit from parent)	<p>The report group is not shared.</p> <p>Any copied report group or generated report remains in the users report list.</p> <p>Each report in the group is assigned any parent report sharing option that was configured.</p>
Share with Everyone	The report group is shared with all users.

(Continued)

Option	Description
Share with users matching the following criteria...	<p>The report group is shared with specific users.</p> <ul style="list-style-type: none"> • User Roles--Select from the list of user roles and press the add icon (+). • Security Profiles--Select from the list of security profiles and press the add icon (+).

5. Click **Save**.

On the **Report Groups** window, shared users see the report group in the report list. Generated reports display content based on security profile setting.

Assign a Report to a Group

You can use the **Assign Groups** option to assign a report to another group.

1. Click the **Reports** tab.
2. Select the report that you want to assign to a group.
3. From the **Actions** list box, select **Assign Groups**.
4. From the **Item Groups** list, select the check box of the group you want to assign to this report.
5. Click **Assign Groups**.

Copying a Report to Another Group

Use the **Copy** icon to copy a report to one or more report groups.

1. Click the **Reports** tab.
2. Click **Manage Groups**.
3. From the navigation tree, select the report that you want to copy.

4. Click **Copy**.
5. Select the group or groups to which you want to copy the report.
6. Click **Assign Groups**.
7. Close the **Report Groups** window.

Removing a Report

Use the **Remove** icon to remove a report from a group.

When you remove a report from a group, the report still exists on the **Reports** tab. The report is not removed from your system.

1. Click the **Reports** tab.
2. Click **Manage Groups**.
3. From the navigation tree, navigate to the folder that contains the report you want to remove.
4. From the list of groups, select the report that you want to remove.
5. Click **Remove**.
6. Click **OK**.
7. Close the **Report Groups** window.

RELATED DOCUMENTATION

[Duplicating a Report | 333](#)

[Sharing a Report | 333](#)

[Branding Reports | 334](#)