JUNIPER NETWORKS | Engineering Simplicity

Juniper Secure Analytics Architecture and Deployment Guide

Published
2022-05-09

RELEASE
7.5.0

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

Juniper Secure Analytics Architecture and Deployment Guide
7.5.0

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide

Use this guide to understand JSA architecture, assess JSA component functionality in your network, and plan and perform JSA deployment.

1
CHAPTER

JSA Architecture Overview

JSA Architecture Overview

IN THIS SECTION

When you plan or create your JSA deployment, it's helpful to have a good awareness of JSA architecture to assess how JSA components might function in your network, and then to plan and create your JSA deployment.

JSA collects, processes, aggregates, and stores network data in real time. JSA uses that data to manage network security by providing real-time information and monitoring, alerts and offenses, and responses to network threats.

JSA is a modular architecture that provides real-time visibility of your IT infrastructure, which you can use for threat detection and prioritization. You can scale JSA to meet your log and flow collection, and analysis needs. You can add integrated modules to your JSA platform, such as JSA Risk Manager, and JSA Vulnerability Manager.

The operation of the JSA security intelligence platform consists of three layers, and applies to any JSA deployment structure, regardless of its size and complexity. The following diagram shows the layers that make up the JSA architecture.

Figure 1: JSA Architecture

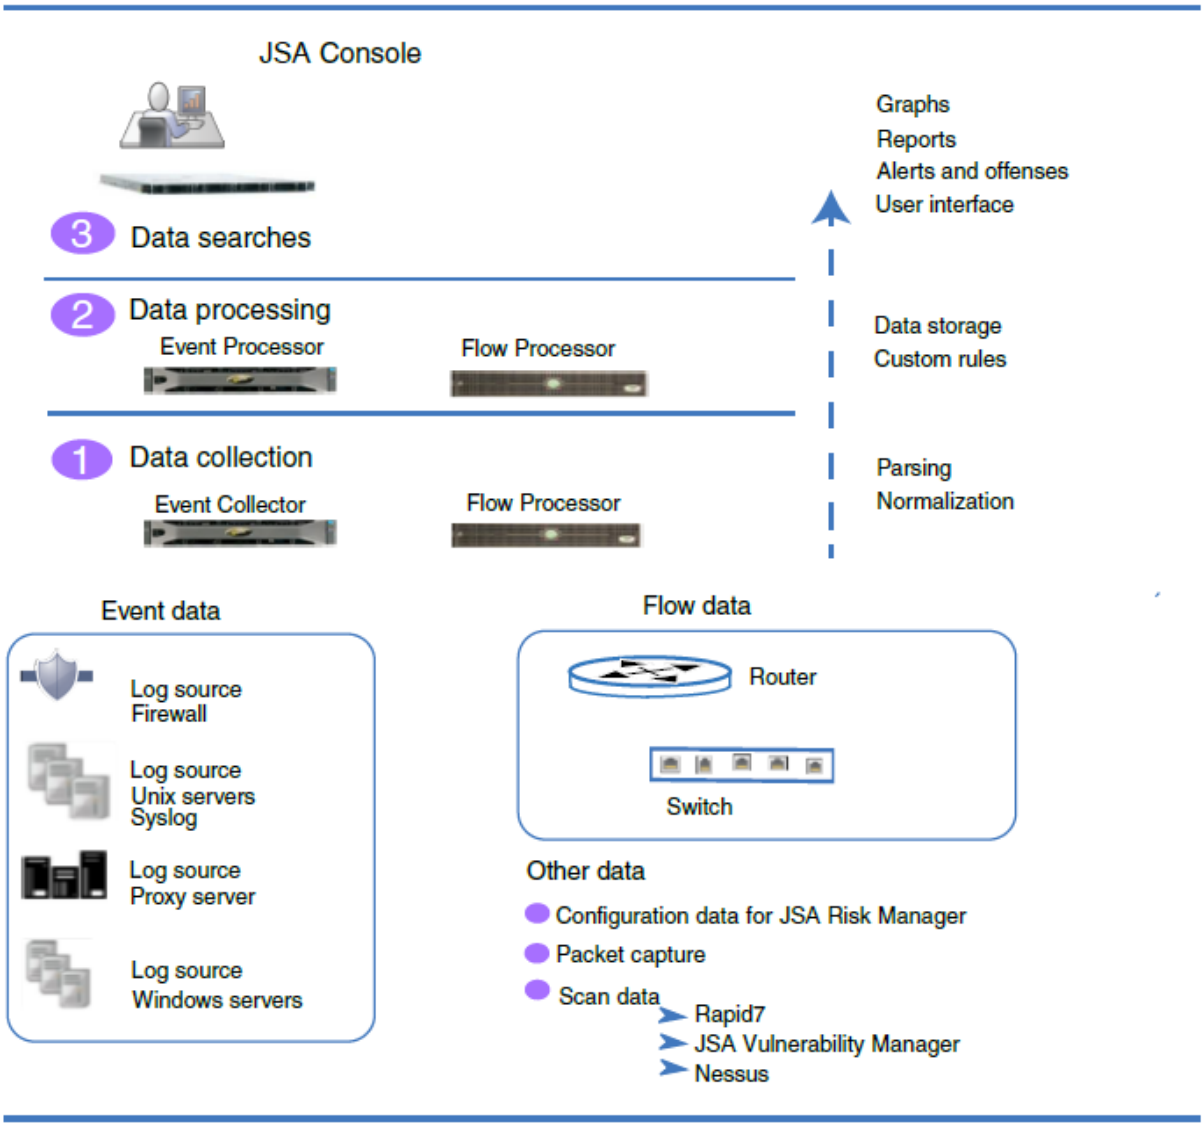

The JSA architecture functions the same way regardless of the size or number of components in a deployment. The following three layers that are represented in the diagram represent the core functionality of any JSA system.

Data Collection

Data collection is the first layer, where data such as events or flows is collected from your network. The All-in-One appliance can be used to collect the data directly from your network or you can use collectors such as JSA Event Collectors or JSA Flow Processor to collect event or flow data. The data is parsed and normalized before it passed to the processing layer. When the raw data is parsed, it is normalized to present it in a structured and usable format.

The core functionality of JSA is focused on event data collection, and flow collection.

Event data represents events that occur at a point in time in the user's environment such as user logins, email, VPN connections, firewall denys, proxy connections, and any other events that you might want to log in your device logs.

Flow data is network activity information or session information between two hosts on a network, which JSA translates in to flow records. JSA translates or normalizes raw data in to IP addresses, ports, byte and packet counts, and other information into flow records, which effectively represents a session between two hosts.

Data Processing

After data collection, the second layer or data processing layer is where event data and flow data are run through the Custom Rules Engine (CRE), which generates offenses and alerts, and then the data is written to storage.

Event data, and flow data can be processed by an All-in-One appliance without the need for adding Event Processors or Flow Processors. If the processing capacity of the All-in-One appliance is exceeded, then you might need to add Event Processors, Flow Processors or any other processing appliance to handle the additional requirements. You might also need more storage capacity, which can be handled by adding Data Nodes.

Other features such as JSA Risk Manager, JSA Vulnerability Manager collect different types of data and provide more functions.

JSA Risk Manager collects network infrastructure configuration, and provides a map of your network topology. You can use the data to manage risk by simulating various network scenarios through altering configurations and implementing rules in your network.

Use JSA Vulnerability Manager to scan your network and process the vulnerability data or manage the vulnerability data that is collected from other scanners such as Nessus, and Rapid7. The vulnerability data that is collected is used to identify various security risks in your network.

Data Searches

In the third or top layer, data that is collected and processed by JSA is available to users for searches, analysis, reporting, and alerts or offense investigation. Users can search, and manage the security admin tasks for their network from the user interface on the JSA console.

In an All-in-One system, all data is collected, processed, and stored on the All-in-One appliance.

In distributed environments, the JSA console does not perform event and flow processing, or storage. Instead, the JSA console is used primarily as the user interface where users can use it for searches, reports, alerts, and investigations.

JSA Components

Use JSA components to scale a JSA deployment, and to manage data collection and processing in distributed networks.

NOTE: Software versions for all JSA appliances in a deployment must be same version and build. Deployments that use different versions of software are not supported because environments that use mixed versions can cause rules not to fire, offenses not to be created or updated, and errors in search results.

JSA deployments can include the following components:

- **JSA console**--The JSA console provides the JSA user interface, and real-time event and flow views, reports, offenses, asset information, and administrative functions.

  In distributed JSA deployments, use the JSA console to manage hosts that include other components.

- **JSA Event Collector**--The Event Collector collects events from local and remote log sources, and normalizes raw log source events to format them for use by JSA. The Event Collector bundles or coalesces identical events to conserve system usage and sends the data to the Event Processor.

  - Use the JSA Event Collector in remote locations with slow WAN links. The Event Collector appliances do not store events locally. Instead, the appliances collect and parse events before they send events to an Event Processor appliance for storage.

  - The Event Collector can use bandwidth limiters and schedules to send events to the Event Processor to overcome WAN limitations such as intermittent connectivity.

  - The Event Collector is assigned to an EPS license that matches the Event Processor that it is connected to.

- **JSA Event Processor**--The Event Processor processes events that are collected from one or more Event Collector components. The Event Processor processes events by using the Custom Rules Engine (CRE). If events are matched to the CRE custom rules that are predefined on the Console, the Event Processor executes the action that is defined for the rule response.

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

  The processing rate for events is determined by your events per second (EPS) license. If you exceed the EPS rate, events are buffered and remain in the Event Collector source queues until the rate drops. However, if you continue to exceed the EPS license rate, and the queue fills up, your system drops events, and JSA issues a warning about exceeding your licensed EPS rate.

  When you add an Event Processor to an All-in-One appliance, the event processing function is moved from the All-in-One to the Event Processor.

- **JSA Flow Processor**--The Flow Processor processes flows from one or more JSA flow processor appliances. The Flow Processor appliance can also collect external network flows such as NetFlow, J-Flow, and sFlow directly from routers in your network. You can use the Flow Processor appliance to scale your JSA deployment to manage higher flows per minute (FPM) rates. Flow Processors include an on-board Flow Processor, and internal storage for flow data. When you add a Flow Processor to an All-in-One appliance, the processing function is moved from the All-in-One appliance to the Flow Processor.

- **JSA Data Node**--Data Nodes enable new and existing JSA deployments to add storage and processing capacity on demand as required. Data Nodes help to increase the search speed in your deployment by providing more hardware resources to run search queries on.

- **QRadar App Host**--An App Host is a managed host that is dedicated to running apps. App Hosts provide extra storage, memory, and CPU resources for your apps without impacting the processing capacity of your JSA Console. Apps such as User Behavior Analytics with Machine Learning Analytics require more resources than are currently available on the Console.

For more information about managing JSA components, see the *Juniper Secure Analytics Administration Guide*.

For more information about JSA appliance specifications, see the *Juniper Secure Analytics Hardware Guide*.

RELATED DOCUMENTATION

JSA Events and Flows | 9

JSA Maximum EPS Certification Methodology

JSA appliances are certified to support a certain maximum events per second (EPS) rate. Maximum EPS depends on the type of data that is processed, system configuration, and system load.

Deployments that significantly deviate from the test parameters that are described in this document might not be able to support the certified rates. The maximum certified EPS rate is absolute. If the load on your system is lighter than the JSA maximum EPS certification load, the EPS maximum rate for your deployment won't increase.

The following information describes the test parameters used to determine the maximum EPS rates of JSA hosts to help you set expectations and plan future JSA deployments with an appropriate EPS goal in mind.

- Event Traffic
  - Unique log sources - 50,000
  - Unique log source types - 17
  - Unique source IP addresses 250,000
  - Unique destination IP addresses - 250,000
  - Unique username - 300,000
  - Coalescing ratio - 15%
  - Average raw event size - 382 B
- Traffic composition specifics: Percentage of the total contribution of data for each device type out of the total dataset. For example, the Microsoft Windows Security events represent 25% of the total dataset used in testing.
  - Microsoft Windows Security - 25%
  - Linux OS - 25%
  - Cisco IOS - 15%
  - Cisco ASA - 10%
  - Linux DHCP - 5%
  - Aruba Mobility controller - 5%
  - Blue Coat SG Appliance - 3%

    - McAfee Web Gateway - 3%
    - Apache HTTP Server - 1%
    - CheckPoint - 1%
    - Cisco IronPort - 1%
    - F5 Networks FirePass - 1%
    - FireEyeMPS - 1%
    - IBM Security Network ProtectionXGS - 1%
    - Palo Alto PA Series - 1%
    - Symantec Endpoint Protection - 1%
    - Websense V Series - 1%
- System configuration
    - Network Hierarchy - 1000 objects
    - Custom properties - 350
    - Custom Rules and Building Blocks - 451
    - Indexes - 20
- Artifacts created as a result of data processing
    - Offenses - 3000
    - Assets - 365,000
    - Reference Data - 11 data structures, 100,000 elements in total
- User load
    - Up to 16 concurrent searches

JSA Events and Flows

IN THIS SECTION

The core functions of JSA are managing network security by monitoring flows and events.

A significant difference between event and flow data is that an event, which typically is a log of a specific action such as a user login, or a VPN connection, occurs at a specific time and the event is logged at that time. A flow is a record of network activity that can last for seconds, minutes, hours, or days, depending on the activity within the session. For example, a web request might download multiple files such as images, ads, video, and last for 5 to 10 seconds, or a user who watches a Netflix movie might be in a network session that lasts up to a few hours. The flow is a record of network activity between two hosts.

Events

JSA accepts event logs from log sources that are on your network. A log source is a data source such as a firewall or intrusion protection system (IPS) that creates an event log.

JSA accepts events from log sources by using protocols such as syslog, syslog-tcp, and SNMP. JSA can also set up outbound connections to retrieve events by using protocols such as SCP, SFTP, FTP, JDBC, Check Point OPSEC, and SMB/CIFS.

Event Pipeline

Before you can view and use the event data on the JSA console, events are collected from log sources and then processed by the Event Processor. JSA All-in-One appliance functions as the Event Collector and Event Processor, in addition to fulfilling the role of the JSA console.

JSA can collect events by using a dedicated Event Collector appliance, or by using an All-in-One appliance where the event collection service and event processing service runs on the All-in-One appliance.

The following diagram shows the layers of the event pipeline.

Figure 2: Event Pipeline

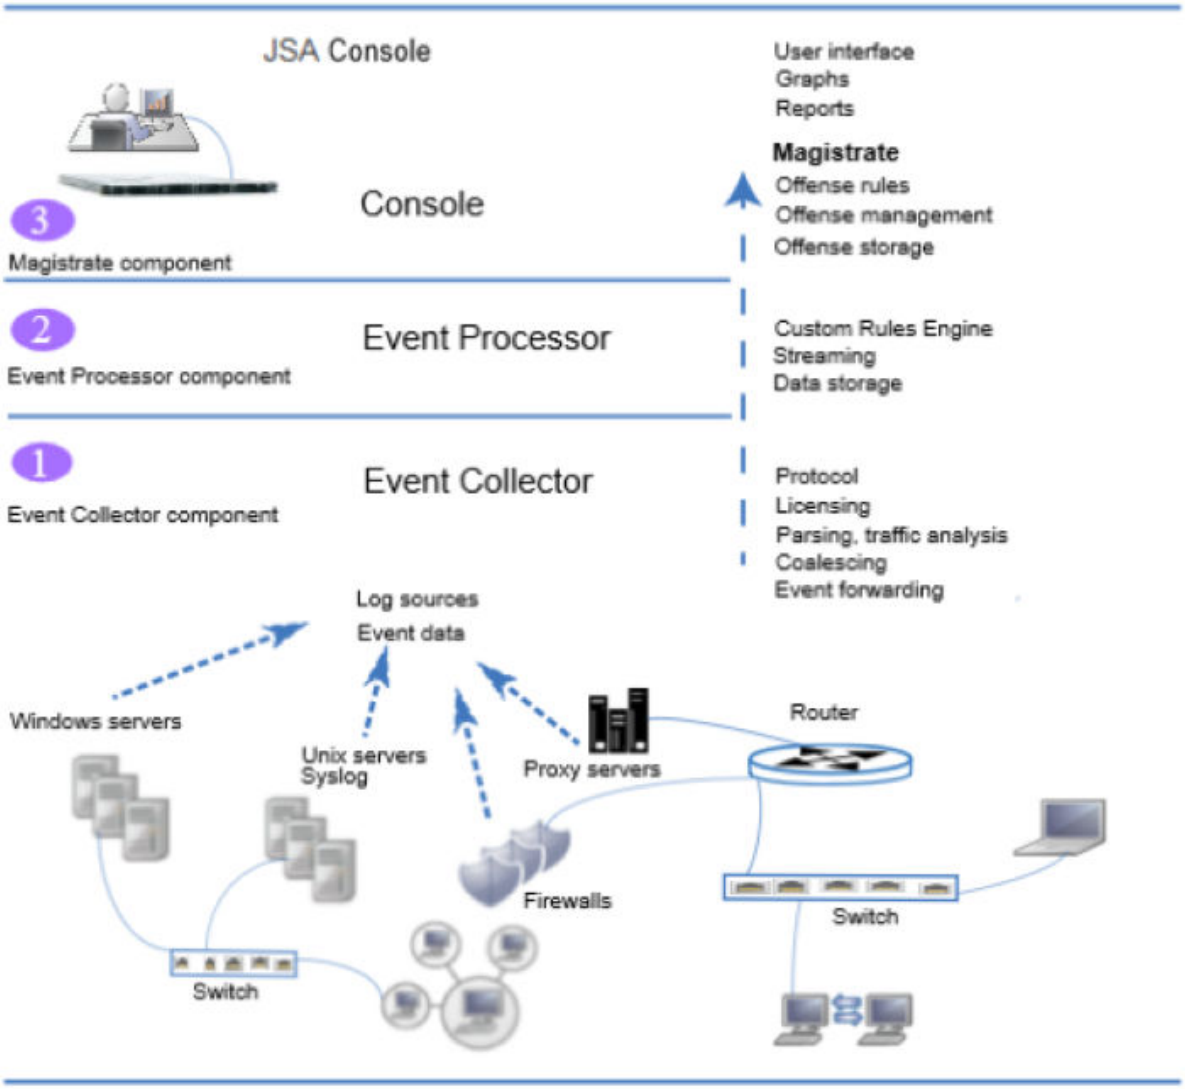

- **Event collection**--The Event Collector component completes the following functions:
  - **Protocol**

    Collects data from log source protocols such as Syslog, JDBC, OPSEC, Log File, and SNMP.

- **License throttling**

  Monitors the number of incoming events to the system to manage input queues and EPS licensing.

- **Parsing**

  Takes the raw events from the source device and parses the fields into a JSA usable format.

- **Log source traffic analysis and auto discover**

  Applies the parsed and normalized event data to the possible DSMs that support automatic discovery.

- **Coalescing**

  Events are parsed and then coalesced based on common attributes across events.

- **Event forwarding**

  Applies routing rules for the system to forward data to offsite targets, external Syslog systems, JSON systems, and other SIEMs.

When the Event Collector receives the events from log sources such as firewalls, the events are placed into input queues for processing.

The queue sizes vary based on the protocol or method that is used, and from these queues, the events are parsed and normalized. The normalization process involves turning raw data into a format that has fields such as IP address that JSA can use.

JSA recognizes known log sources by the source IP address or host name that is contained in the header.

JSA parses and coalesces events from known log sources into records. Events from new or unknown log sources that were not detected in the past are redirected to the traffic analysis (auto detection) engine.

When new log sources are discovered, a configuration request message to add the log source is sent to the JSA console. If auto detection is disabled, or you exceed your log source licensed limit, the new log sources are not added.

- **Event processing**--The Event Processor component completes the following functions:

  - Custom Rules Engine (CRE)

    The Custom Rules Engine (CRE) is responsible for processing events that are received by JSA and comparing them against defined rules, keeping track of systems involved in incidents over time, generating notifications to users. When events match a rule, a notification is sent from the Event Processor to the Magistrate on the JSA console that a specific event triggered a rule. The

Magistrate component on the JSA console creates and manages offenses. When rules are triggered, responses or actions such as notifications, syslog, SNMP, email messages, new events, and offenses are generated.

- Streaming

  Sends real-time event data to the JSA console when a user is viewing events from the **Log Activity** tab with Real time (streaming). Streamed events are not provided from the database.

- Event storage (Ariel)

  A time-series database for events where data is stored on a minute by minute basis. Data is stored where the event is processed.

The Event Collector sends normalized event data to the Event Processor where the events are processed by Custom Rules Engine (CRE). If events are matched to the CRE custom rules that are predefined on the JSA console, the Event Processor executes the action that is defined for the rule response.

- **Magistrate on the JSA console**--The Magistrate component completes the following functions:

  - Offense rules

    Monitors and acts on offenses, such as generating email notifications.

  - Offense management

    Updates active offenses, changes statuses of offenses, and provides user access to offense information from the **Offenses** tab.

  - Offense storage

    Writes offense data to a Postgres database.

  The Magistrate Processing Core (MPC) is responsible for correlating offenses with event notifications from multiple Event Processor components. Only the JSA console or All-in-One appliance has a Magistrate component.

Flows

JSA flows represent network activity by normalizing IP addresses, ports, byte and packet counts, and other data, into flow records, which effectively are records of network sessions between two hosts. The component in JSA that collects and creates flow information is known as Flow Processor.

JSA Flow collection is not full packet capture. For network sessions that span multiple time intervals (minutes), the flow pipeline reports a record at the end of each minute with the current data for metrics

such as bytes, and packets. You might see multiple records (per minute) in JSA with the same "First Packet Time" but the "Last Packet Time" values increment through time.

A flow starts when the flow processor detects the first packet that has a unique source IP address, destination IP address, source port, destination port, and other specific protocol options, including 802.1q VLAN fields.

Each new packet is evaluated. Counts of bytes and packets are added to the statistical counters in the flow record. At the end of an interval, a status record of the flow is sent to a Flow Processor and statistical counters for the flow are reset. A flow ends when no activity for the flow is detected within the configured time.

Flow Processor can process flows from the following internal or external sources:

- External sources are flow sources such as netflow, sflow, jflow. External sources can be sent to a dedicated flow processor or to a Flow Processor appliance. External sources do not require as much CPU processing because every packet is not processed to build flows. In this configuration, you might have a dedicated flow processor and a Flow Processor that both receive and create flow data. In smaller environments (less than 50 Mbps), an All-in-One appliance might handle all the data processing.
- The flow processor collects internal flows by connecting to a SPAN port, or a network TAP. The JSA Flow Processor can forward full packets from it's capture card to a packet capture appliance but it does not capture full packets itself.

The following diagram shows the options for collecting flows in a network.

Figure 3: JSA Flows

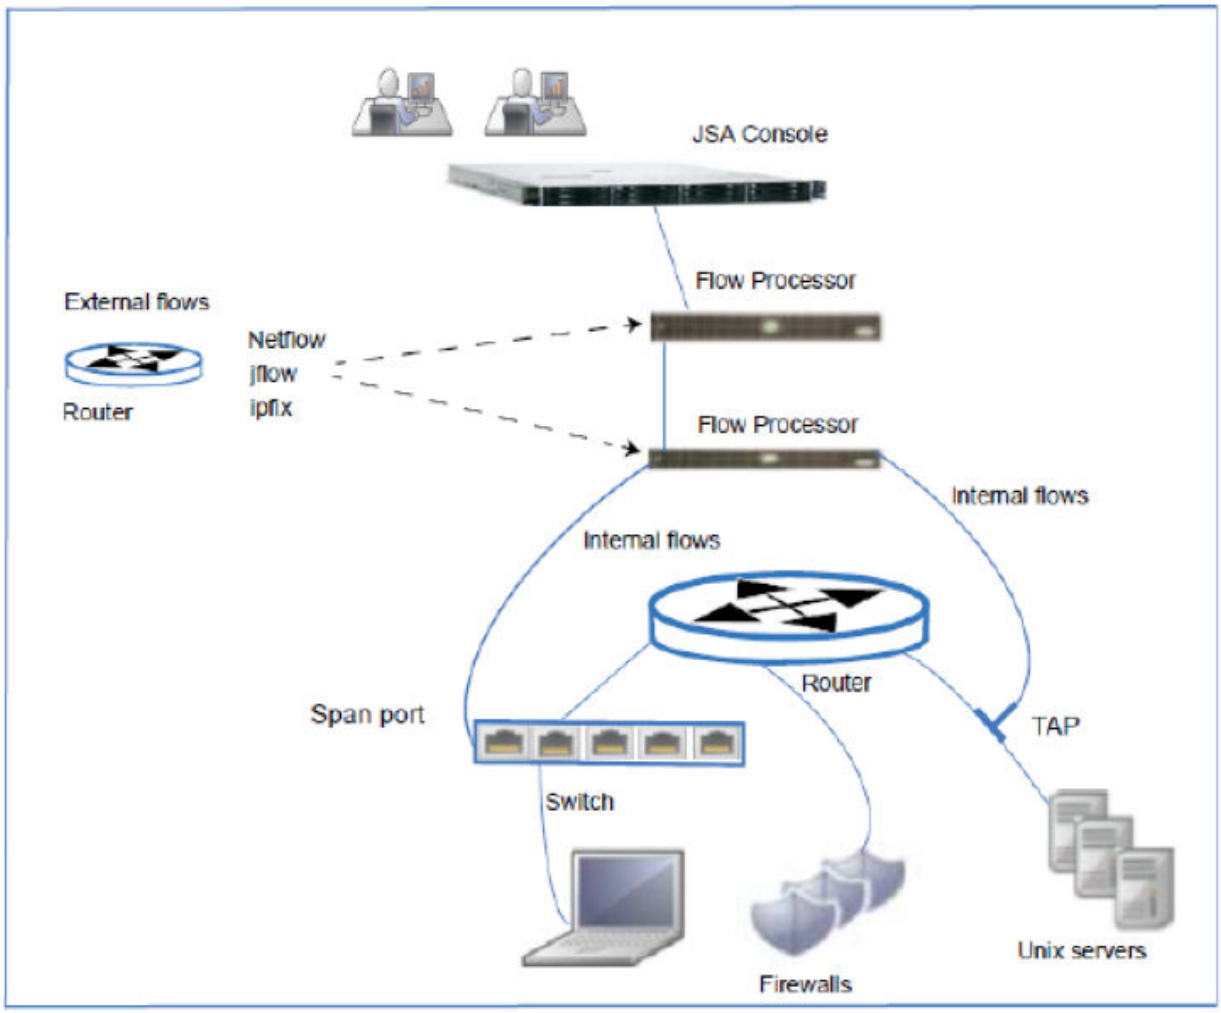

Flow Pipeline

The flow processor generates flow data from raw packets that are collected from monitor ports such as SPANs, TAPs and monitor sessions, or from external flow sources such as netflow, sflow, jflow. This data is then converted to JSA flow format and sent down the pipeline for processing.

The Flow Processor runs the following functions:

- **Flow deduplication**

Flow deduplication is a process that removes duplicate flows when multiple flow processors provide data to Flow Processors appliances.

- **Asymmetric recombination**

  Responsible for combining two sides of each flow when data is provided asymmetrically. This process can recognize flows from each side and combine them in to one record. However, sometimes only one side of the flow exists.

- **License throttling**

  Monitors the number of incoming flows to the system to manage input queues and licensing.

- **Forwarding**

  Applies routing rules for the system, such as sending flow data to offsite targets, external Syslog systems, JSON systems, and other SIEMs.

Flow data passes through the Custom Rules Engine (CRE), and it is correlated against the rules that are configured, and an offense can be generated based on this correlation. You view offenses on the **Offenses** tab.

RELATED DOCUMENTATION

JSA Components | 5

2
CHAPTER

JSA Deployment Overview

JSA Deployment Overview

JSA architecture supports deployments of varying sizes and topologies, from a single host deployment, where all the software components run on a single system, to multiple hosts, where appliances such as event collectors, and flow processors, Data Nodes, an App Host, Event Processors, and Flow Processors, have specific roles.

The primary focus of the first deployment example is to describe a single All-in-One appliance deployment for a medium-size company. Later examples describe the deployment options as the company expands. The examples describe when to add JSA components, such as Flow Processors, event collectors, and Data Nodes, and when you might need to co-locate specific components.

The requirements for your JSA deployment depend on the capacity of your chosen deployment to both process and store all the data that you want to analyze in your network.

Before you plan your deployment, consider the following questions:

- How does your company use the Internet? Do you upload as much as you download? Increased usage can increase your exposure to potential security issues.
- How many events per second (EPS) and flows per minute (FPM) do you need to monitor?

  EPS and FPM license capacity requirements increase as a deployment grows.
- How much information do you need to store, and for how long?

The following diagram shows the JSA components that you can use to collect, process, and store event and flow data in your JSA deployment. An All-in-One appliance includes the data collection, processing, storage, monitoring, searching, reporting, and offense management capabilities.

The Event Collector collects event data from log sources in your network, and then sends the event data to the Event Processor. The flow processor collects flow data from network devices such as a switch SPAN port, and then sends the data to the Flow Processor. Both processors process the data from the collectors and provide data to the JSA console. The processor appliances can store data but they can

also use the Data Nodes to store data. The JSA console appliance is used for monitoring, data searches, reporting, offense management, and administration of your JSA deployment.

Figure 4: JSA Event and Flow Components

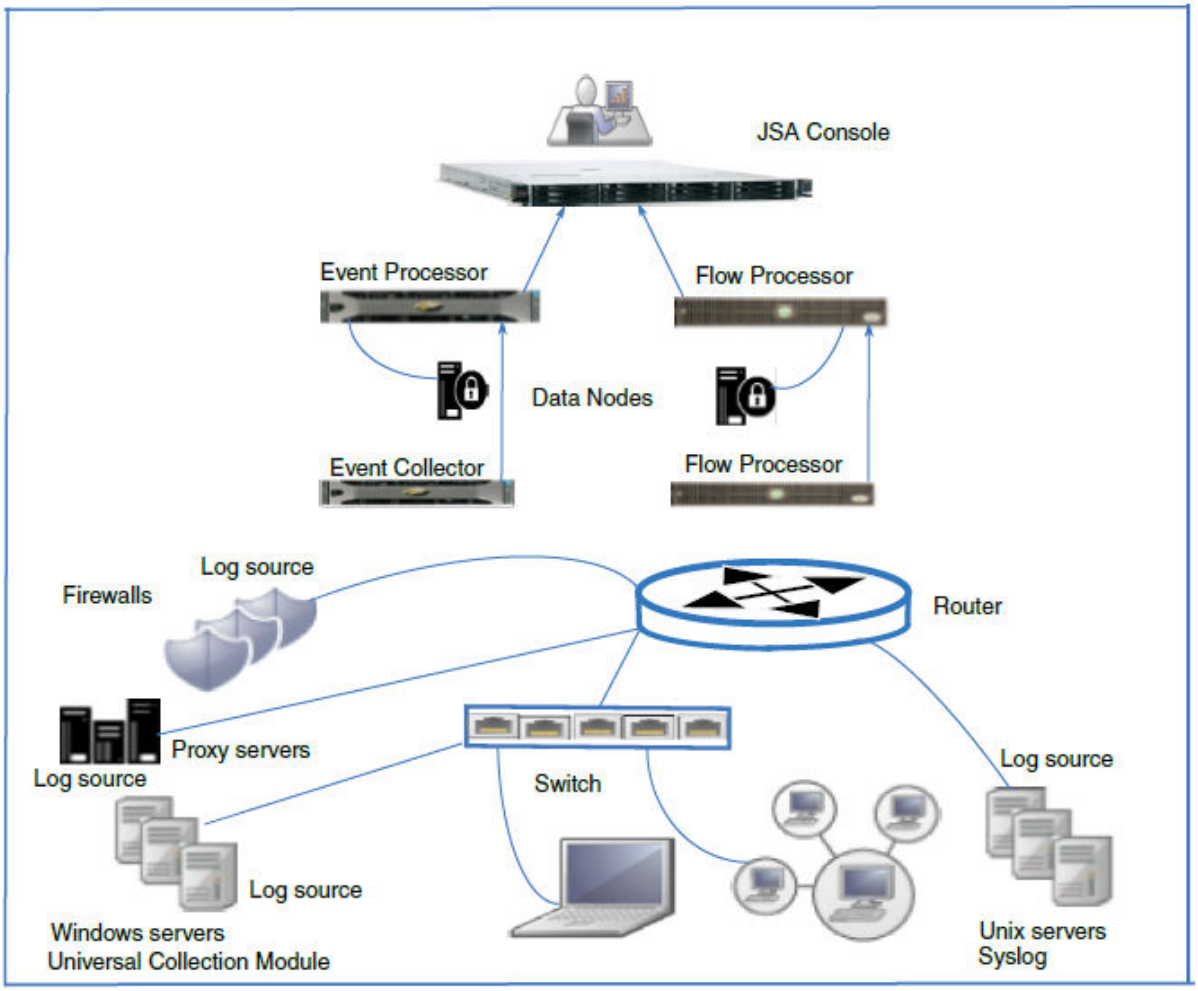

All-in-One Deployment

In a single host JSA deployment, you have an All-in-One JSA appliance that is a single server which collects data, such as syslog event data logs, and Windows events, and also flow data, from your network.

An All-in-One appliance is suitable for a medium-sized company that has low exposure to the Internet, or for testing and evaluation purposes. Single server deployments are suitable for companies that monitor network activity and events such as authentication services and firewall activity.

An All-in-One appliance provides you with the capabilities that you need, up to a specific capacity that is determined by your license and the hardware specifications of the system.

Manufacturing company deploys a single JSA server

You are a medium-sized manufacturing company with less than 1000 employees. You deploy a JSA All-in-One appliance to collect, process, and monitor event and flow data. With that deployment, you can collect up to 5,000 events per second (EPS), and 200,000 flows per minute (FPM).

The following diagram shows an All-in-One appliance, which collects data from event and flow sources, processes the data, and provides a web application where you can search, monitor, and respond to security threats.

Figure 5: All-in-One Deployment

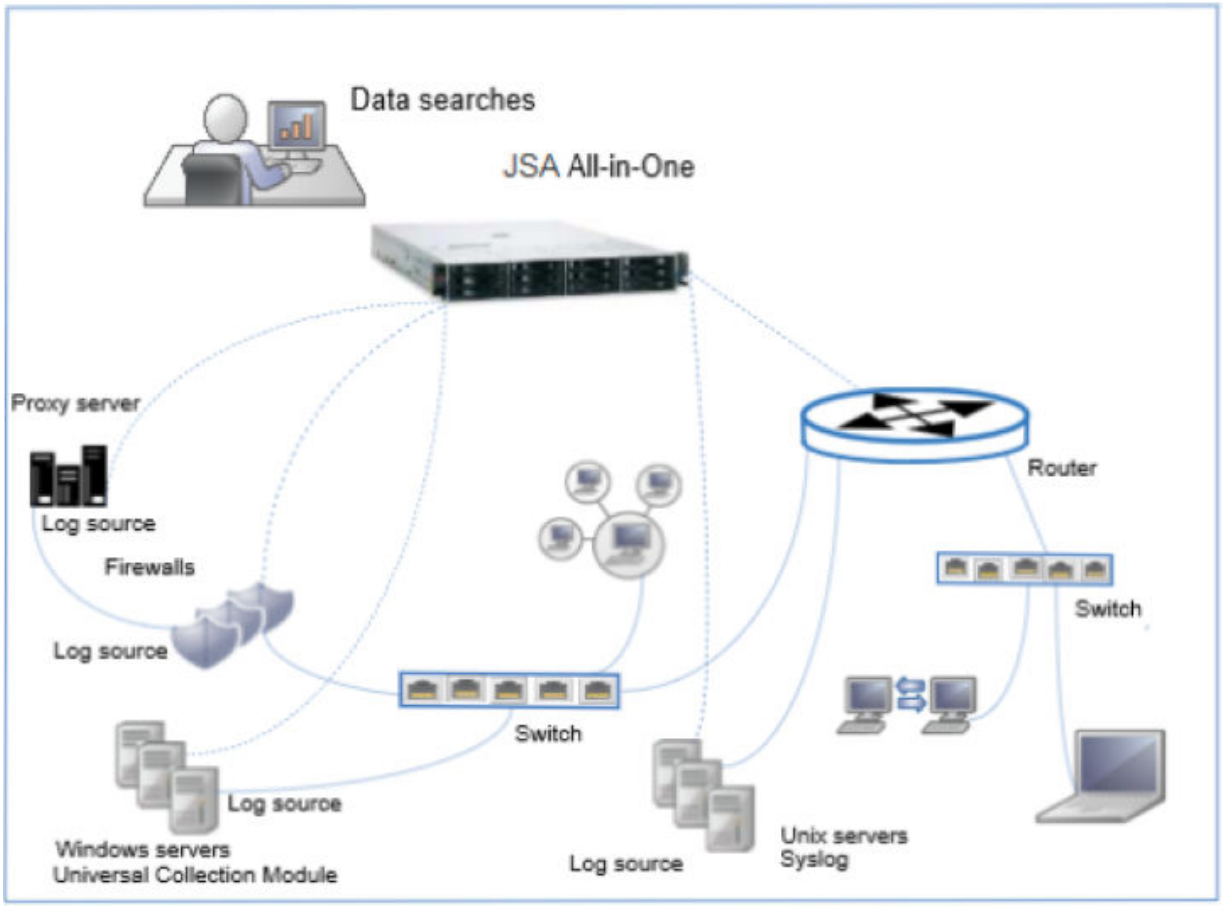

An All-in-One appliance performs the following tasks:

- Collects event and network flow data, and then normalizes the data in to a data format that JSA can use.
- Analyzes and stores the data, and identifies security threats to the company.
- Provides access to the JSA web application.

As your data sources grow, or your processing or storage needs increase, you can add appliances to expand your deployment. For more information, see "Expanding Deployments to Add More Capacity" on page 20.

RELATED DOCUMENTATION

Expanding Deployments to Add More Capacity

IN THIS SECTION

Your business might create or expand a deployment beyond an JSA All-in-One appliance because of the lack of processing or data storage capacity, or when you have specific data collection requirements.

The topology and composition of your JSA deployment are influenced by the capability and capacity of that deployment to collect, process, and store all the data that you want to analyze in your network.

If your processing or storage needs to expand beyond the capacity of your All-in-One appliance, you can reconfigure your JSA environment to a distributed deployment. For more information, see "Adding an Appliance to an All-in-One Console" on page 26.

To get rough estimates of the events per second (EPS) or flows per minute (FPM) that you need to process in your deployment, use the size of your logs that are collected from firewalls, proxy servers, and Windows boxes.

Reasons to Add Event or Flow Processors to an All-in-One Deployment

You might need to add flow or event collectors to your deployment under these conditions:

- Your data collection requirements exceed the collection capability of the All-in-One appliance.
- You must collect events and flows in a different location than where your All-in-One appliance is installed.
- You are monitoring larger, or higher-rate packet-based flow sources that are faster than the 50 Mbps connection on the All-in-One.

An All-in-One appliance can collect up to 15,000 events per second (EPS) and 300,000 flows per minute (FPM). If your collection requirements are greater, you might want to add event collectors and flow processors to your deployment.

An All-in-One appliance processes the events and flows that are collected. By adding event collectors and flow processors, you can use the processing that the All-in-One appliance usually does for searches and other security tasks.

Packet-based flow sources require a flow processor that is connected either to a Flow Processor, or to an All-in-One appliance in deployments where there is no Flow Processor appliance. You can collect external flow sources, such as NetFlow, or IPFIX, directly on a Flow Processor or All-in-One appliance.

Adding Remote Collectors to a Deployment

Add JSA event collectors or JSA flow processors to expand a deployment when you need to collect more events locally and collect events and flows from a remote location.

For example, you are a manufacturing company that has a JSA All-in-One deployment and you add e-commerce and a remote sales office. You now must monitor for security threats and are also now subject to PCI audits.

You hire more employees and the Internet usage changes from mostly downloading to two-way traffic between your employees and the Internet. Here are details about your company.

- The current events per second (EPS) license is 1000 EPS.

- You want to collect events and flows at the sales office and events from the e-commerce platform.
- Event collection from the e-commerce platform requires up to 2000 events-per-second (EPS).
- Event collection from the remote sales office requires up to 2000 events-per-second (EPS).
- The flows per minute (FPM) license is sufficient to collect flows at the remote office.

You take the following actions:

1. You add the e-commerce platform at your head office, and then you open a remote sales office.
2. You install an Event Collector and a flow processor at the remote sales office that sends data over the Internet to the All-in-One appliance at your head office.
3. You upgrade your EPS license from 1000 EPS to 5000 EPS to meet the requirements for the extra events that are collected at the remote office.

The following diagram shows an example deployment of when an Event Collector and a flow processor are added at a remote office.

Figure 6: Collectors in Remote Office

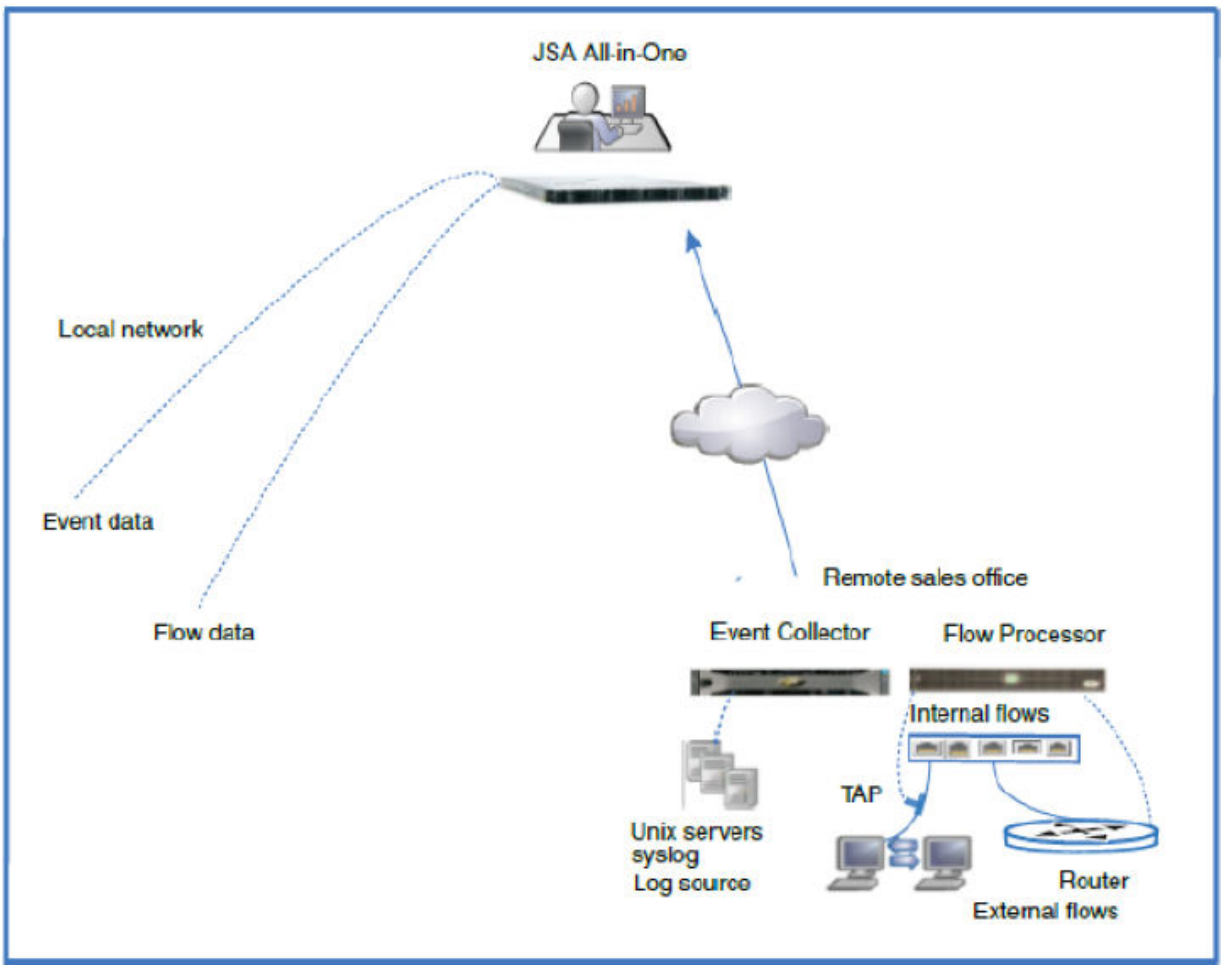

In this deployment, the following processes occur:

- At your remote office, the Event Collector collects data from log sources and the flow processor collects data from routers and switches. The collectors coalesce and normalize the data.
- The collectors compress and send data to the All-in-One appliance over the wide area network.
- The All-in-One appliance processes, and stores the data.
- Your company monitors network activity by using the JSA web application for searches, analysis, reporting, and for managing alerts and offenses.
- The All-in-one collects and processes events from the local network.

Adding Processing Capacity to an All-in-One Deployment

Add Event Processors and Flow Processors to your JSA deployment to increase processing capacity and increase storage. Adding processors frees up resources on your JSA Console by moving the processing and storage load to dedicated servers.

When you add Event Processors or Flow Processors to an All-in-One appliance the All-in-One acts as a JSA Console. The processing power on the All-in-One appliance is dedicated to managing and searching the data that is sent by the processors, and data is now stored on the Event Processors and other storage devices, rather than on the Console.

You typically add Event Processors and Flow Processors to your JSA deployment for the following reasons:

- As your deployment grows, the workload exceeds the processing capacity of the All-in-One appliance.
- Your security operations center employs more analysts who do more concurrent searches.
- The types of monitored data, and the retention period for that data increases, which increases processing and storage requirements.
- As your security analyst team grows, you require better search performance.

Running multiple concurrent JSA searches and adding more types of log sources that you monitor, affects the processing performance of your All-in-One appliance. As you increase the number of searches and the amount of monitored data, add Event Processors and Flow Processors to improve the performance of your JSA deployment.

When you scale your JSA deployment beyond the 15,000 EPS and 300,000 FPM on the most powerful All-in-One appliance, you must add processor appliances to process that data.

Example: Adding a JSA Event Processor to your deployment

You can add a JSA Event Processor 1624, which collects and processes up to 40,000 EPS. You increase your capacity by another 40,000 EPS every time you add a JSA Event Processor 1624 to your deployment. Add a JSA Flow Processor 1724, which collects and processes up to 1,200,000 FPM.

The JSA Event Processor 1624 is a collector and a processor. If you have a distributed network, it's a good practice to add Event Collectors to distribute the load and to free system resources on the Event Processor.

In the following diagram, processing capacity is added when an Event Processor and a Flow Processor are added to a JSA appliance (All-in-One), and the following changes take place:

- Event and flow processing is moved off the All-in-One appliance to the event and flow processors.

- Event processing capacity increases to 40,000 EPS, which includes the 15,000 EPS that was on the All-in-One.
- Flow processing capacity increases to 1,200,000 FPM, which includes the 300,000 FPM that was on the All-in-One.
- Data that is sent by the event and flow processor is processed and stored on the event and flow processors.

Figure 7: Adding Processing Capacity

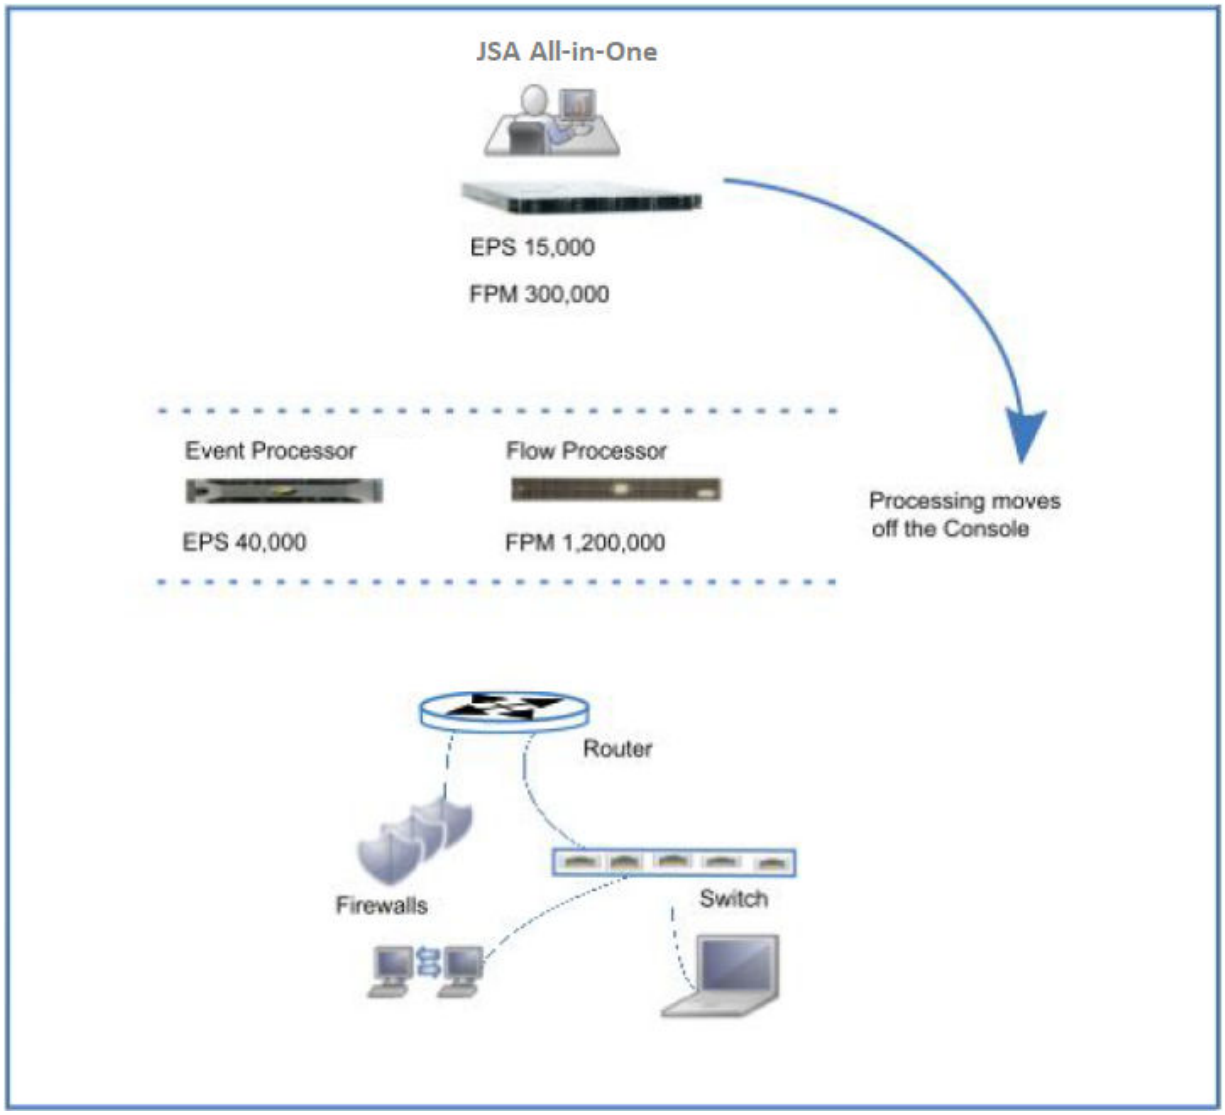

Search performance is faster when you install Event Processors and Flow Processors on the same network as your JSA Console.

Adding processors and collectors expands the processing capacity of your JSA deployment. You can also increase the storage capacity of your deployment. Your company's data retention needs can increase due to more traffic or to changes to retention policies. Adding Data Nodes to your deployment expands your data storage capacity, and improves search performance.

When to add Collectors to Processors

Add Event Collectors and Flow Processors to Event Processors for the same reasons that you add collectors to an All-in-One appliance:

- Your data collection requirements exceed the collection capability of your processor.
- You must collect events and flows at a different location than where your processor is installed.
- You are monitoring packet-based flow sources.

NOTE: Event Collectors can buffer events, but Flow Processors can't buffer flows.

Because search performance is improved when processors are installed on the same network as the console, adding collectors in remote locations, and then sending that data to the processor, speeds up your JSA searches.

Adding an Appliance to an All-in-One Console

Any All-in-One console can become part of a distributed deployment by adding another appliance to extend the resources and performance of the All-in-One console. Adding more appliances can increase storage, process more data, and search faster. The Console appliance manages the other JSA appliances in the network. When you add appliances to an All-in-One console, it becomes a distributed deployment console.

Adding hosts allows users to expand on the capabilities, storage, and resources from an All-in-One appliance to create a distributed deployment.

1. Log in to the All-in One Console as an administrator.
2. From the Admin tab, click the **System and License Management** icon.
3. In the Display list, select **Systems**.
4. On the Deployment Actions menu, select Add Host.
5. Enter the **Host IP**, **Host Password** (root user password). and configure the properties to Encrypt Host Connections or define Network Address Translation parameters.

6. Check either **Encrypt Host Connections** or **Network Address Translation** and configure your choice.
7. Click **Add**.
8. From the Admin tab, click **Deploy Changes**.
9. Click **Continue** to restart services.

The host is added to the deployment and the appliance is added to the user interface. After the host is added, you might need to allocate licenses to the appliance.

RELATED DOCUMENTATION

Geographically Distributed Deployments | 27

JSA Vulnerability Manager Deployments | 29

Geographically Distributed Deployments

In geographically distributed deployments your JSA deployment might be impacted by intermittent or poor connectivity to remote data centers. You might also be impacted by local regulations, such as complying with specific state or country regulations to keep data in the place of origin. Both of these situations require that you keep collectors on site. If you must keep data in the place of origin, then you must keep the processor on site.

For example, your company is growing and this growth not only increases activity in the network, but it also requires that you expand your JSA deployment to other countries. Data retention laws vary from country to country, so the JSA deployment must be planned with these regulations in mind.

You note these following conditions:

- Your company must collect event data from one of the office locations that has intermittent connectivity.

- Your company must comply with data retention regulations in the countries where data is collected. For example, if Germany requires that data remains in-country, then that data must not be stored outside that country.

Figure 8: Geographically Distributed Deployment

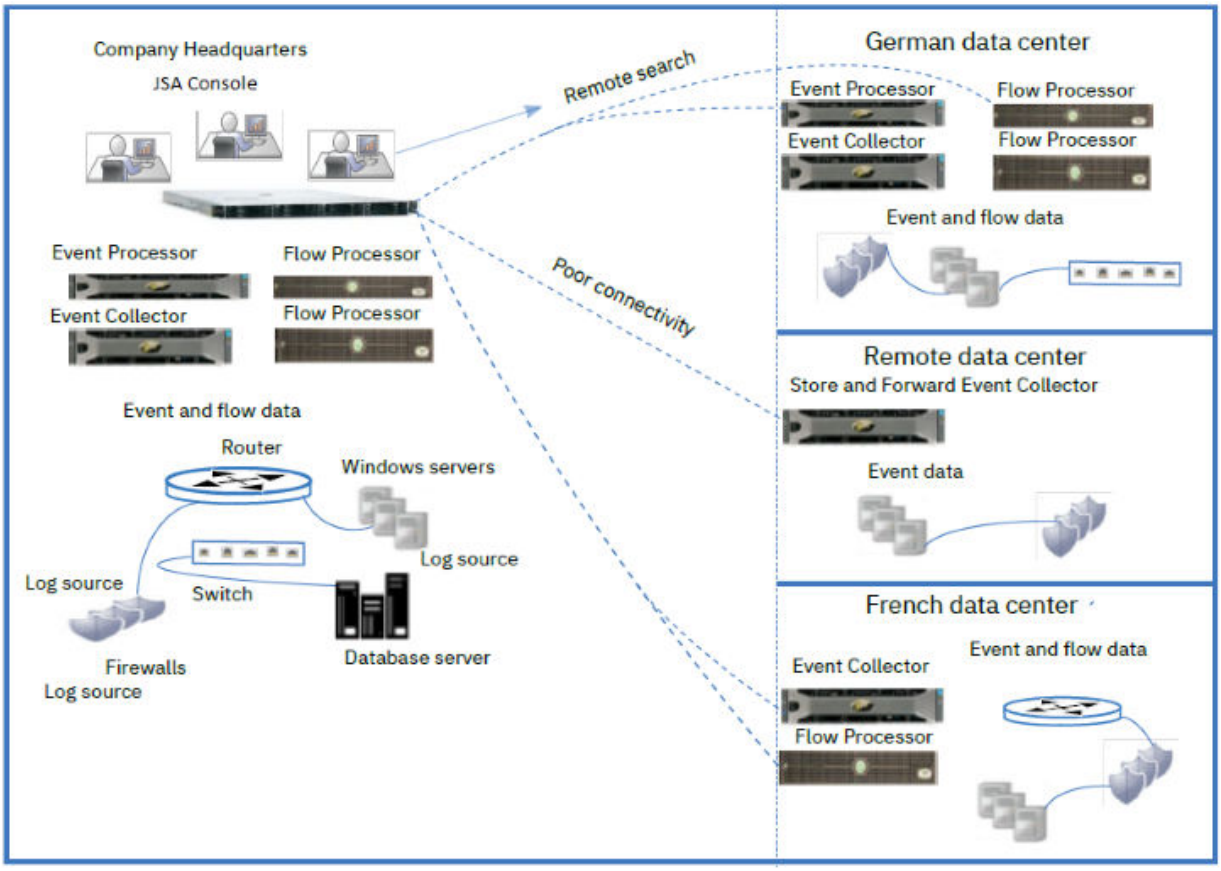

In the geographically distributed deployment, the following processes occur:

- Your company installs collectors and processors in the German data center to comply with local data laws.
- In the French data center, your company installs collectors, so that data is sent by the collectors to the head office, which is processed and stored at the head office. Search speeds are increased by having the processor appliances on the same high-speed network segment as the JSA Console.
- Your company adds a store-and-forward Event Collector that has scheduled and rate-limited forwarding connections in the remote data center. The scheduled and rate-limited connections compensate for intermittent network connectivity and ensures that the requirement for more bandwidth is avoided during regular business hours.

If you're constantly searching for data on a remote processor, it's better to have that processor on the same high-speed network segment as the JSA Console. If the bandwidth between the JSA Console and remote processor is not good, you might experience latency with your searches, especially when you're doing multiple concurrent searches.

RELATED DOCUMENTATION

JSA Vulnerability Manager Deployments | 29

Expanding Deployments to Add More Capacity | 20

JSA Vulnerability Manager Deployments

IN THIS SECTION

Locate and manage the vulnerabilities in your network by deploying JSA Vulnerability Manager.

JSA Vulnerability Manager discovers vulnerabilities on your network devices, applications, and software adds context to the vulnerabilities, prioritizes asset risk in your network, and supports the remediation of discovered vulnerabilities.

You can integrate JSA Risk Manager for added protection, which provides network topology, active attack paths and high-risk assets risk-score adjustment on assets based on policy compliance. JSA Vulnerability Manager and JSA Risk Manager are combined into one offering and both are enabled through a single base license.

Depending on the product that you install, and whether you upgrade JSA or install a new system, the **Vulnerabilities** tab might not be displayed. Access JSA Vulnerability Manager by using the **Vulnerabilities** tab. If you install JSA, the **Vulnerabilities** tab is enabled by default with a temporary license key. If you install Log Manager, the **Vulnerabilities** tab is not enabled. You can use the **Try it Out** option to try out JSA Vulnerability Manager for 30 days. You can purchase the license for JSA Vulnerability Manager separately and enable it by using a license key. For more information about upgrading, see the *Upgrading Juniper Secure Analytics to 7.5.01 Guide*.

JSA Vulnerability Manager Components

The following information describes the JSA Vulnerability Manager Processor.

- The scan processor is responsible for the scheduling and managing scans, and delegating work to the scanners that might be distributed throughout your network.
- You can have only one scan processor in a JSA deployment.
- When you install and license JSA Vulnerability Manager on an All-in-One system, a vulnerability processor is automatically deployed on your JSA console and includes a scanning component.
- The vulnerability processor provides a scanning component by default. If required, you can move the vulnerability processor to a different managed host in your deployment.
- If you add a Vulnerability Processor managed host appliance, and JSA Vulnerability Manager is used for the first time, then the scan processor is assigned to the Vulnerability Processor managed host appliance.
- The scanning processor is governed by the processing license, which determines the maximum number of assets that can be processed by JSA Vulnerability Manager.
- The scan processor can run on the JSA console or a managed host.

The following information describes the JSA Vulnerability Manager scanner.

- You can deploy a scanner on a virtual machine or as software only.
- You can deploy a JSA Vulnerability Manager scanner dedicated scanner appliance, which is a Vulnerability Scanner appliance.
- You can deploy a scanner on a JSA console or on the following managed hosts: Flow Processor, Flow Processor Event Collector, Event Processor, or Data Node.
- The number of assets that you can scan with a scanner is determined by the scanner capacity and is not impacted by licensing.

Components and Scan Process

Scan jobs are completed by a processor and a scanner component. The following diagram shows the scan components and the processes that run.

Figure 9: Scan Components and Process

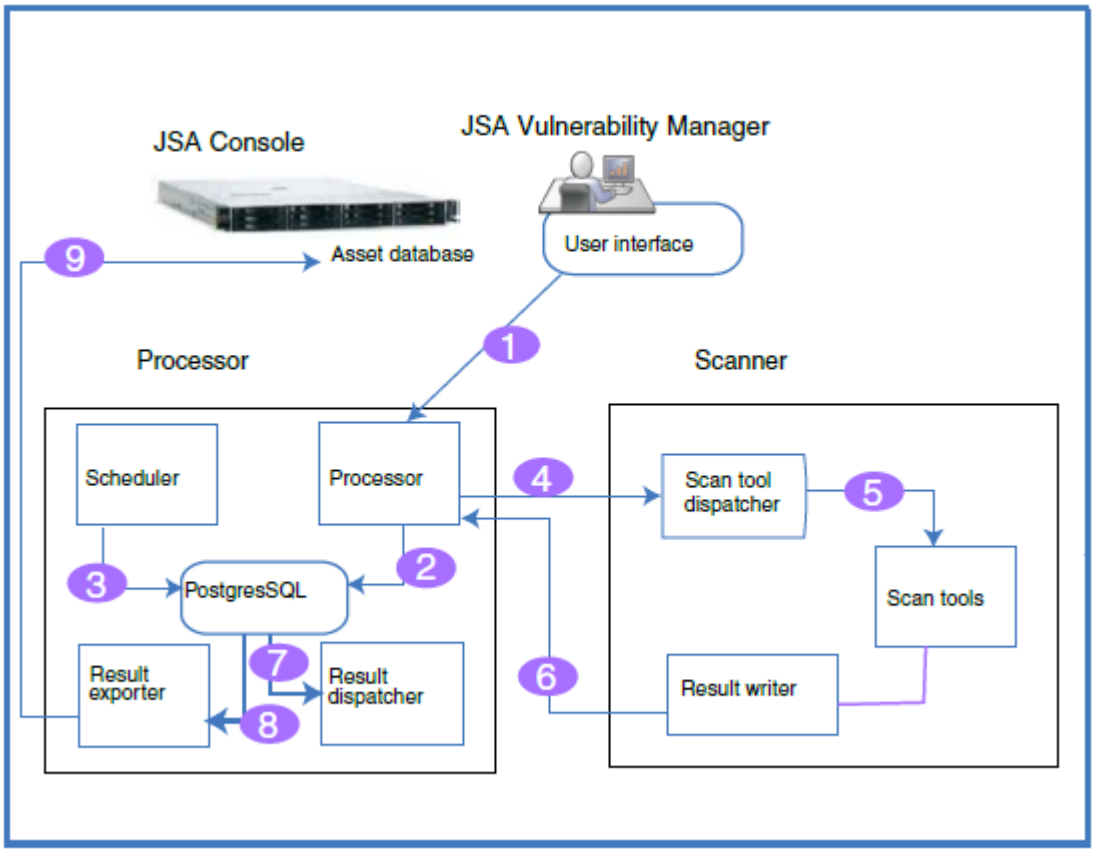

The following list describes the steps in the scan process:

1. You create a scan job by specifying parameters such as IP addresses of assets, type of scan, and required credentials for authenticated scans.
2. The scan job is accepted by the processor, logged, and added to the database along with scheduling information to determine when the job runs.
3. The scheduler component manages the scheduling of scans. When the scheduler initiates a scan, it determines the list of tools that are required and queues them for invocation, and then the tools are assigned to the relevant scanner.

4. Scanners poll the scan processor continuously for scan tools that it must run by sending a unique scanner ID. When the scheduler has queued tools that are relevant to the specific scanner the tools are sent to the scanner for invocation.

   JSA Vulnerability Manager uses an attack tree methodology to manage scans and to determine which tools are launched. The phases are: asset discovery, port/service discovery, service scan, and patch scan.

5. The dispatcher runs and manages each scan tool in the list. For each tool that is run, the dispatcher sends a message to the processor that indicates when a scan tool starts and finishes.

6. The output from the scan tool is read by the result writer, which then passes these results back to the processor.

7. The result dispatcher processes the raw results from the scan tools and records them into the Postgres database.

8. The result exporter finds completed scans in the processor database and exports the results to the JSA console.

9. The exported results are added to the JSA database where users can view and manage the scan results.

All-in-one Deployment

You can run JSA Vulnerability Manager from an All-in-one system, where the scanning and processing functions are on the Console. The following information describes what you can do with a basic setup:

- Scan up to 255 assets.
- Unlimited discovery scans.
- Use hosted scanner for DMZ scanning.
- Manage scan data from third-party scanners that are integrated with JSA.
- Deploy a scanner on any managed host.
- Deploy unlimited stand-alone software or virtual scanners.

Expanding a Deployment

As your deployment grows, you might need to move the processing function off the JSA console to free up resources, and you might want to deploy scanners closer to your assets.

The following list describes reasons to add scanners to your deployment:

- To scan assets in a different geographic region than the JSA Vulnerability Manager processor.
- If you want to scan many assets concurrently within a short time frame.
- You might want to add a scanner to avoid scanning through a firewall that is a log source. You might also consider adding the scanner directly to the network by adding an interface on the scanner host that by-passes the firewall.

The following diagram shows a scanning deployment with external scanning and scanners deployed on managed hosts.

Figure 10: Scanning Deployment

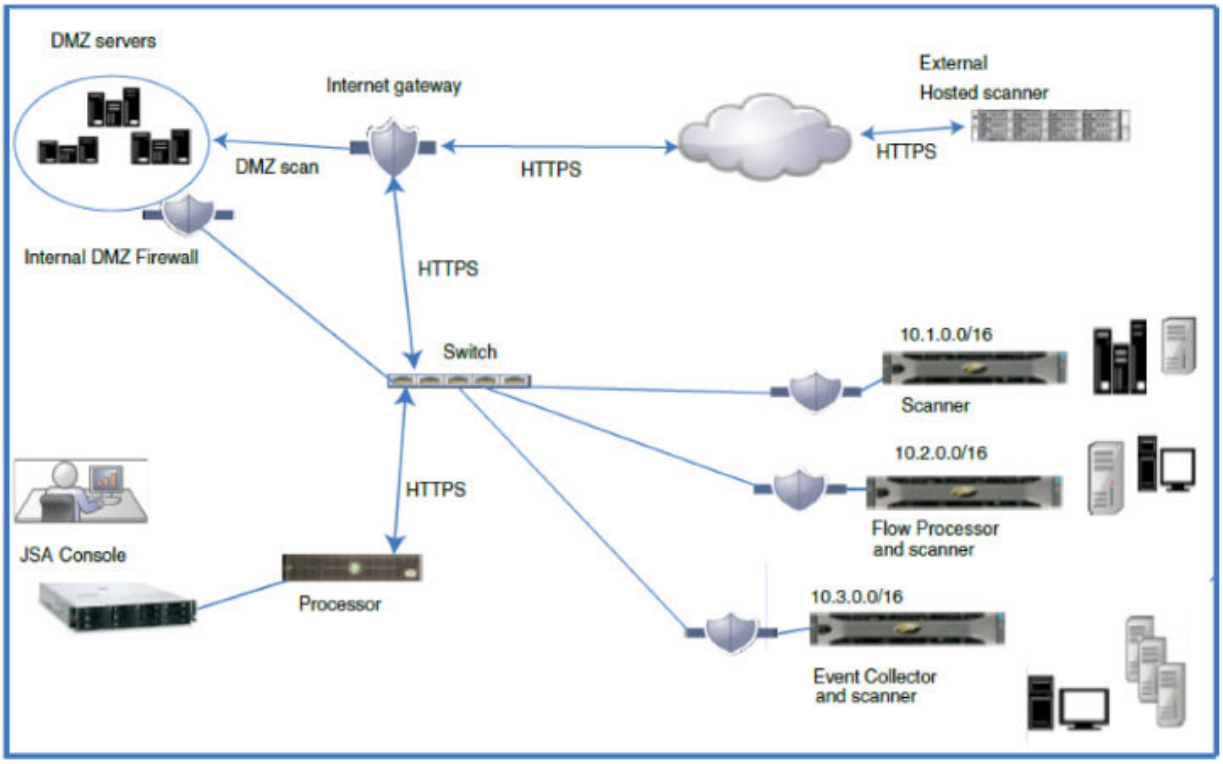

DMZ Hosted Scanner

A hosted scanner scans your DMZ from the Internet by using your public IP address. If you want to scan the assets in the DMZ for vulnerabilities, you do not need to deploy a scanner in your DMZ. You must configure JSA Vulnerability Manager with a hosted IBM scanner that is located outside your network. For more information, see the *Juniper Secure Analytics Vulnerability Manager User Guide*.

JSA Vulnerability Manager Integrations

JSA Vulnerability Manager integrates with HCL BigFix to help you filter and prioritize the vulnerabilities that can be fixed. BigFix provides shared visibility and control between IT operations and security. BigFix applies Fixlets to high priority vulnerabilities that are identified and sent by JSA Vulnerability Manager to BigFix. Fixlets are packages that you deploy to your assets or endpoints to remediate specific vulnerabilities.

JSA Vulnerability Manager integrates with IBM Security SiteProtector to help direct intrusion prevention system (IPS) policy. When you configure IBM Security SiteProtector, the vulnerabilities that are detected by scans are automatically forwarded to IBM Security SiteProtector. IBM Security SiteProtector receives vulnerability data from JSA Vulnerability Manager scans that are run only after the integration is configured. Connecting to IBM Security SiteProtector.

Third-party Scanners

JSA Vulnerability Manager delivers an effective vulnerability management platform, regardless of the source of the scan data. JSA Vulnerability Manager integrates seamlessly with third-party scanners such as Nessus, nCircle, and Rapid 7.

You require JSA Vulnerability Manager scanning to get the following options:

- Event driven and on-demand scanning
- Asset database and watchlist based scanning
- Scanning from existing JSA appliances and managed hosts
- Detection of newly published vulnerabilities that are not present in any scan results

You require JSA Risk Manager to get the following options:

- Asset, vulnerability, and traffic-based vulnerability management

- Adjusted vulnerability scores and context aware risk scoring.

JSA Risk Manager and JSA Vulnerability Manager

Enhance your network security by integrating JSA Risk Manager with JSA Vulnerability Manager. Data sources, such as scan data, enable JSA Risk Manager to identify security, policy, and compliance risks in your network and calculate the probability of risk exploitation.

JSA Vulnerability Manager and JSA Risk Manager are combined into one offering and both are enabled through a single base license.

Add a JSA Risk Manager appliance to get the following capabilities:

- Compliance assessment
- Risk policies that are based on vulnerability data and risk scores that help you quickly identify high-risk vulnerabilities.
- Visibility into potential exploit paths from potential threats and untrusted networks through the network topology view.
- Risk policy-based filtering.
- Topology visualization
- False positives reduction in vulnerability assessments.
- Visibility into what vulnerabilities are blocked by firewalls and Intrusion Prevention Systems (IPS).

JSA Risk Manager Appliance

Install JSA Risk Manager separately on a JSA Risk Manager appliance.

You must install JSA Console before you set up and configure the JSA Risk Manager appliance. It is a good practice to install JSA and JSA Risk Manager on the same network switch.

You require only one JSA Risk Manager appliance per deployment.

The following diagram shows a deployment that has a scanner and JSA Risk Manager.

Figure 11: Scanning Deployment with Risk Manager

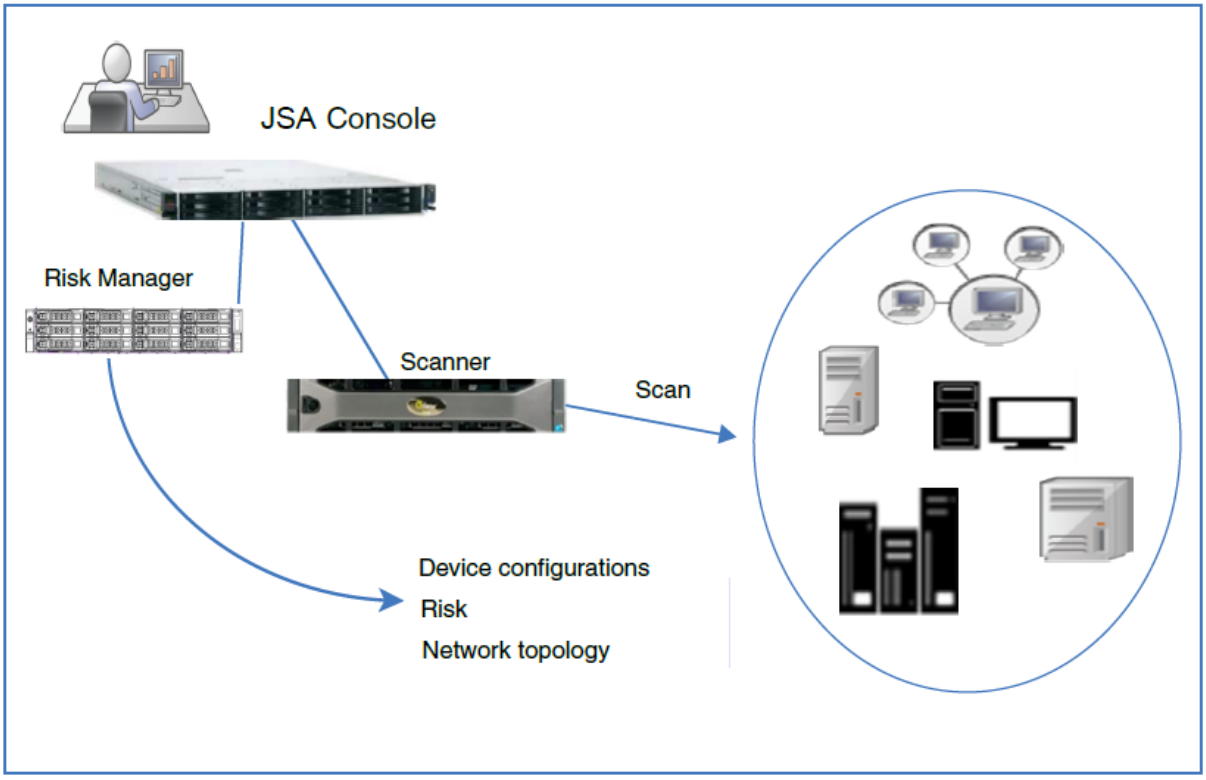

Use Risk Manager to complete the following tasks:

- Centralized risk management.
- View and filter your network topology
- Import and compare device configurations
- View connections between network devices.
- Search firewall rules.
- View existing rules and the event count for triggered rules.
- Search devices and paths
- Query network connections
- Simulate the possible outcomes of updating device configurations.

- Monitor and audit your network to ensure compliance.
- Simulate threats or attacks against a virtual model.
- Search for vulnerabilities.

RELATED DOCUMENTATION

Expanding Deployments to Add More Capacity | 20

Geographically Distributed Deployments | 27

CHAPTER 3

Data Nodes and Data Storage

Data Nodes and Data Storage

IN THIS SECTION

JSA processor appliances and All-in-One appliances can store data but many companies require the stand-alone storage and processing capabilities of the Data Node to handle specific storage requirements and to help with implementing data retention policies. Many companies are impacted by regulations and laws that mandate keeping data records for specific periods.

Data Node Information

The following list describes information about Data Nodes:

- Data Nodes add storage and processing capacity.
- Data Nodes are plug-n-play and can be added to a deployment at any time.
- Data Nodes integrate seamlessly with existing deployments.
- Use Data Nodes to reduce the processing load on processor appliances by removing the data storage processing load from the processor.
- Users can scale storage and processing power independently of data collection.
- As of JSA 2014.7, a new data format with native data compression is used. Data is compressed in memory and is written out to disk in a proprietary binary compressed format. The new data format enables a better search performance and a more efficient use of system resources than the previous data format. The previous data format did not have a native built-in compression in older versions of JSA.

The following diagram shows an example of some uses for Data Nodes in a deployment.

Figure 12: Using Data Node Appliances to Manage Your Data Storage

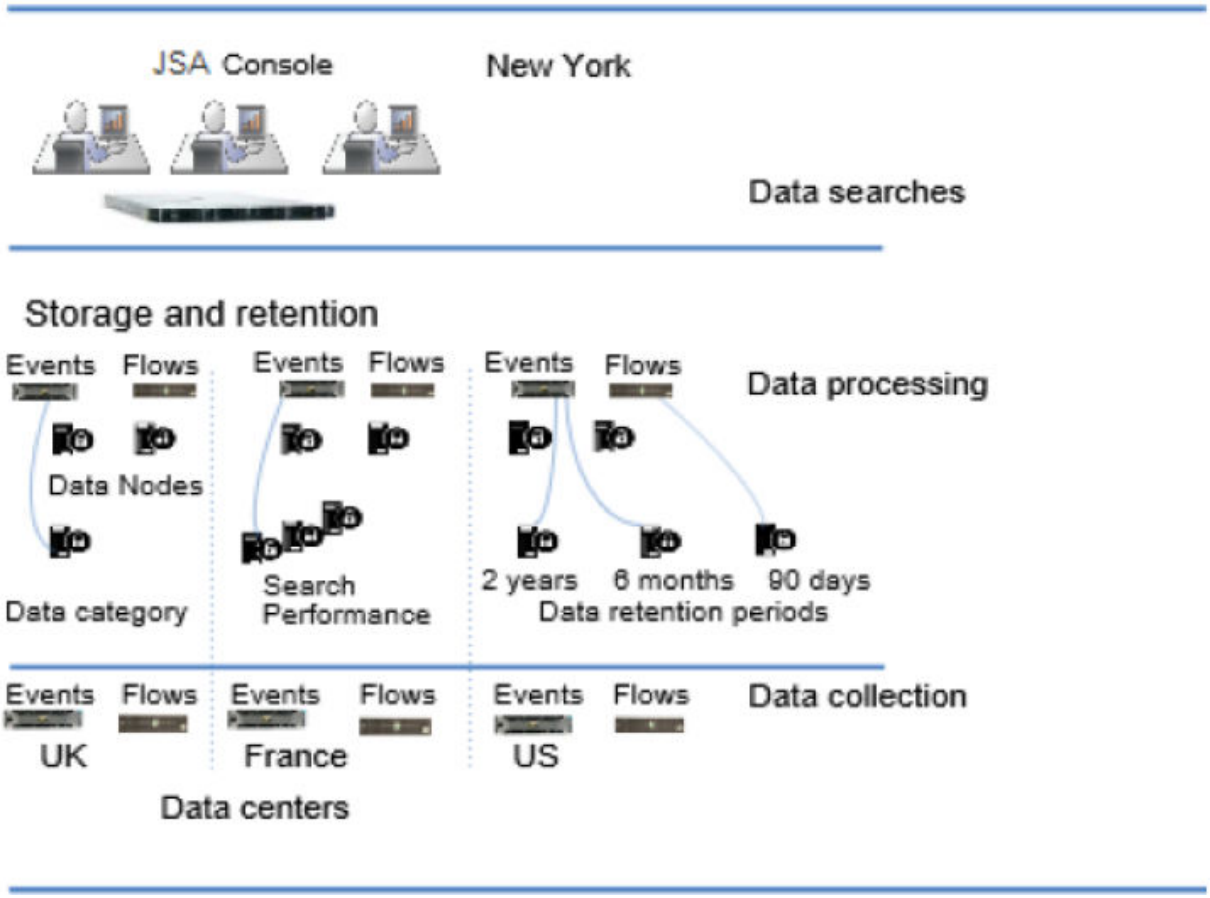

The following list describes the different elements that you need to consider when you deploy Data Nodes.

- **Data clustering**-- Data Nodes add storage capacity to a deployment, and also improve performance by distributing data that is collected across multiple storage volumes. When the data is searched, multiple hosts, or a cluster does the search. The cluster can improve search performance, but doesn't require you to add multiple event processors. Data Nodes multiply the storage for each processor.

  NOTE: You can connect a Data Node to only one processor at a time, but a processor can support multiple Data Nodes.

- **Deployment considerations**-- Keep the following information in mind as you set up Data Nodes in a deployment.

  - Data Nodes are available with JSA 2014.5 and later.
  - Data Nodes perform similar search and analytic functions as event and flow processors in a JSA deployment.

    The operational speed on a cluster is affected by the slowest member of a cluster. Data Node system performance improves if Data Nodes are sized similarly to the Event Processors and Flow Processors in a deployment. To facilitate similar sizing between Data Nodes and event and flow processors, Data Nodes are available on JSA core appliances.
  - Data Nodes are available in three formats: software (on your own hardware), physical, and appliances. You can mix the formats in a single cluster.

- **Bandwidth and latency**-- Ensure that you have a 1 Gbps link and less than 10 ms latency between hosts in the cluster. Searches that yield many results require more bandwidth.
- **Appliance compatibility**-- Data Nodes are compatible with all existing JSA appliances that have an Event Processor or Flow Processor component, including All-In-One appliances.

  Data Nodes support high availability (HA).
- **Installation of Data Nodes**-- Data Nodes use standard TCP/IP networking, and do not require proprietary or specialized interconnect hardware.

  Install each Data Node that you want to add to your deployment the same as you would install any other JSA appliance. Associate Data Nodes with either an event or flow processor. For more information, see *Configuring a managed host* in the *Juniper Secure Analytics Administration Guide*.

  You can attach multiple Data Nodes to a single Event Processor or Flow Processor in a many-to-one configuration.

  When you deploy high availability (HA) pairs with Data Node appliances, install, deploy, and rebalance data with the HA appliances before you synchronize the HA pair. The combined effect of the data rebalancing and the replication process that is utilized for HA results in significant performance degradation. If HA is set up on appliances to which Data Nodes are being introduced, then disconnect HA on the appliances and then reconnect it when the rebalance of the cluster is complete.
- **Decommissioning Data Nodes**-- Use the **System and License Management** window to remove Data Nodes from your deployment, as with any other JSA appliance. Decommissioning does not erase data on the host, nor does it move the data to your other appliances. If you need to retain access to the data that was on the Data Nodes, you must identify a location to move that data to.
- **Data Rebalancing**-- Adding a Data Node to a cluster distributes data to each Data Node. If it is possible, data rebalancing tries to maintain the same percentage of available space on each Data Node. New Data Nodes added to a cluster initiate more rebalancing from cluster event and flow processors to achieve efficient disk usage on the newly added Data Node appliances.

Starting with JSA 2014.5, data rebalancing is automatic and concurrent with other cluster activity, such as queries and data collection. No downtime is experienced during data rebalancing.

Data Nodes offer no performance improvement in the cluster until data rebalancing is complete. Rebalancing can cause minor performance degradation during search operations, but data collection and processing continue unaffected.

NOTE: Encrypted data transmission between Data Nodes and Event Processors is not supported.

- **Management and Operations**-- Data Nodes are self-managed and require no regular user intervention to maintain normal operation. JSA manages activities, such as data backups, high availability, and retention policies, for all hosts, including Data Node appliances.
- **Data Node failure**-- If a Data Node fails, the remaining members of the cluster continue to process data.

  When the failed Data Node returns to service, data rebalancing can occur to maintain proper data distribution in the cluster, and then normal processing resumes. During the downtime, data on the failed Data Node is unavailable, and I/O errors that occur appear in search results from the log and network activity viewers in the JSA user interface.

  For catastrophic failures that require appliance replacement or the reinstallation of JSA, decommission Data Nodes from the deployment and replace them using standard installation steps. Copy any data that is not lost in the failure to the new Data Node before you deploy. The rebalancing algorithm accounts for data that exists on a Data Node, and shuffles only data that was collected during the failure.

  For Data Nodes deployed with an HA pair, a hardware failure causes a failover, and operations continue to function normally.

SAN Overview

To increase the amount of storage space on your appliance, you can move a portion of your data to an offboard storage device. You can move your **/store**, **/store/ariel**, or **/store/backup** file systems.

Multiple methods are available for adding external storage, including iSCSI, and NFS (Network File System). You must use iSCSI to store data that is accessible and searchable in the UI, such as the **/store/ariel** directory, and reserve the use of NFS for data backups only.

Moving the **/store** file system to an external device might affect JSA performance.

After migration, all data I/O to the **/store** file system is no longer done on the local disk. Before you move your JSA data to an external storage device, you must consider the following information:

- Searches that are marked as saved are also in the **/transient** directory. If you experience a local disk failure, these searches are not saved.
- A transient partition that exists before you move your data is likely to remain in existence after the move, and it can be mounted on an iSCSI storage mount.

For more information about offboard storage, see the *Juniper Secure Analytics Configuring Offboard Storage Guide*.

4

CHAPTER

App Host

App Host

An App Host is a managed host that is dedicated to running apps. App Hosts provide extra storage, memory, and CPU resources for your apps without impacting the processing capacity of your JSA Console. Apps such as User Behavior Analytics with Machine Learning Analytics require more resources than are currently available on the Console.

App Host information

The following list describes information about App Hosts:

- The App Host was added in JSA 7.3.2 to replace the deprecated App Node.
- You cannot upgrade to JSA 7.3.2 or later with an App Node in your deployment.
- You can only have one App Host per deployment.
- You can run all of your apps on an App Host, or on the Console. It's not possible to run some apps on the App Host and some others on the Console.
- Port 5000 must be open on your Console.

App Host specifications

The following table shows the minimum requirements and suggested specifications for an App Host.

NOTE: *The suggested specifications for medium and large sized deployments haven't been tested. If you are using some of the larger apps, such as the Pulse Dashboard or User Behavior Analytics with Machine Learning, the minimum requirements are probably insufficient. Consider upgrading your deployment environment.

Table 1: App Host Specifications

	CPU cores	RAM	Disk Space	Description
Small	4	12 GB	256 GB	Minimum requirements for an App Host. You can run most apps with the minimum requirements, but not larger apps such as JSA DNS Analyzer and User Behavior Analytics with Machine Learning.
Medium	12 or more	64 GB or more	500 GB or more	*You can run all apps that exist today, but this specification does not give you room for future apps.
Large	24 or more	128 GB or more	1 TB or more	*You can run all apps that exist today and you would have room for future apps.

5
CHAPTER

HA Deployment Overview

HA Deployment Overview

IN THIS SECTION

Implement high availability (HA) in your JSA deployment to keep JSA functions running, if there is a hardware or software failure in your deployment.

By using high availability, you can continue to collect, store, and process event and flow data, if any failures occur.

To enable HA, JSA connects a primary HA host with a secondary HA host to create an HA cluster.

The following diagram shows a basic HA setup.

Figure 13: Basic HA Setup

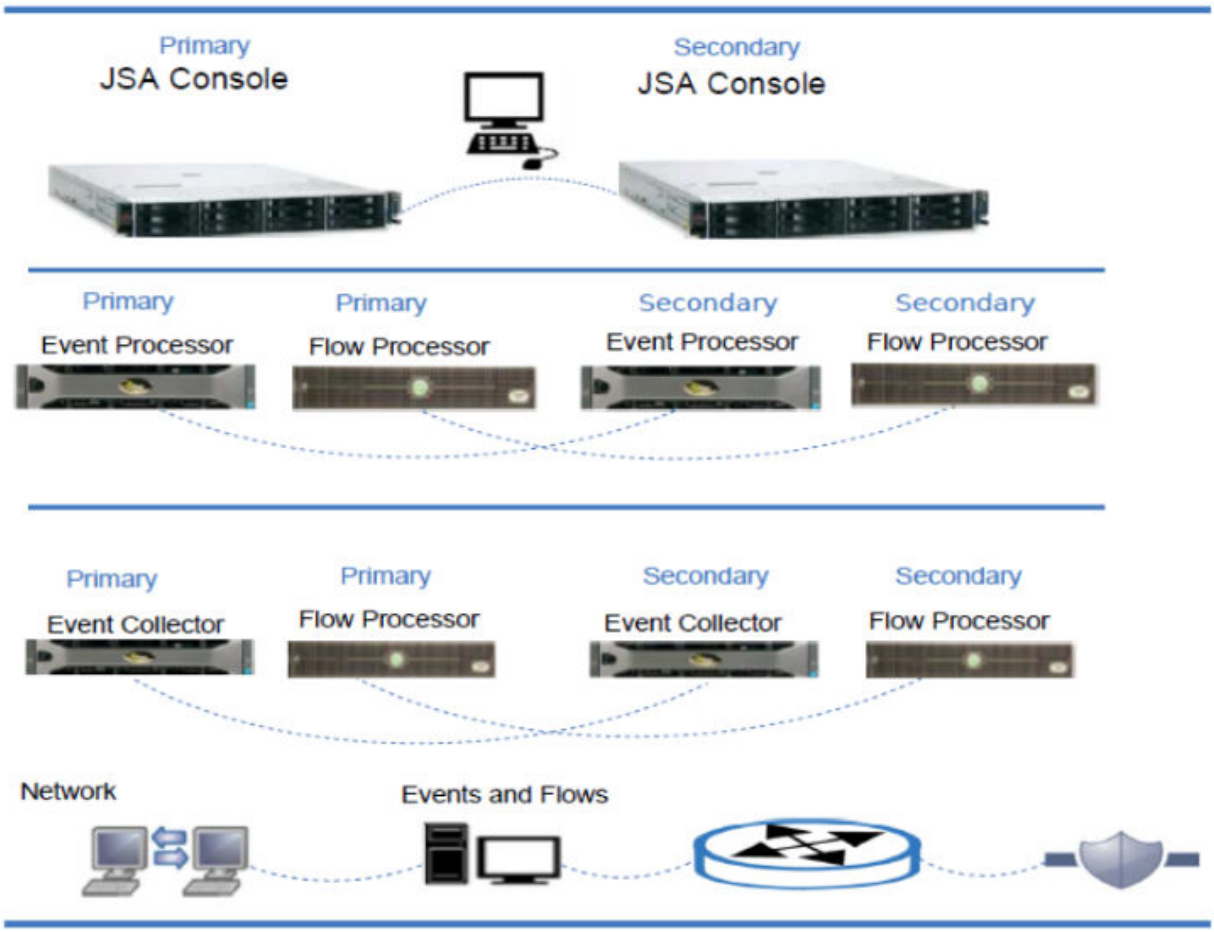

HA Overview

In an HA deployment, you install and configure a second appliance that takes over the role of the device, if the primary appliance fails in one of the following scenarios:

- A power supply failure
- A network failure that is detected by network connectivity tests
- An operating system malfunction that delays or stops the heartbeat ping tests
- A complete RAID failure on the primary HA host

- A manual failover
- A management interface failure on the primary HA host

For best performance in large deployments it is strongly recommended to use a 10 Gbps interface for your HA Crossover. Using a 10 Gbps interface reduces the time needed for system synchronization and ensures optimal performance of the pair. If you do not have a 10 Gbps interface available consider bonding multiple 1 Gbps interfaces for crossover.

For more information about HA, see the *Juniper Secure Analytics High Availability Guide*.

6
CHAPTER

Backup Strategies

Backup Strategies

IN THIS SECTION

Back up your business critical information to safeguard against loss of that data. Different types of data require different backup strategies.

JSA Data Backups

Data classification is an important consideration for backup strategies for the following reasons:

- Data such as personal identity information (PII) needs to be stored securely, and might need to be kept separate from bulk data backups, and retained for longer periods for compliance reasons.
- Keep JSA system configuration data separate from your security data such as events and flows. It is safer to keep the system configuration separate and easier to restore this data if it stored separately.
- Store data such as PCI data in a separate location so that you can easily access this data when auditors want to see it.
- Think about types of data and retention periods when you develop your backup strategies.
- You can back up some types of data more frequently than others and you can use offsite storage for some data to insure against data loss.

Retention Settings

The default setting for JSA backup retention is 7 days. You can also do an on-demand backup after you make major configuration changes. You can give this on-demand backup a descriptive name to easily find your changes if you need to return to this configuration.

Scheduled backups overwrite older scheduled backups. On-demand backups are kept indefinitely. After the JSA backup volume reaches 75% of its capacity, scheduled backups no longer run.

Backup Location

The backup location is also a significant consideration when you deploy JSA. If your backups remain on a host, and that host fails, then all backup data is lost.

You can either create your backups on an external system, or copy backups to an external system.

Store copies of important data locally and remotely for added data security.