

IBM QRadar
7.3.3

Architecture and Deployment Guide



Note

Before you use this information and the product that it supports, read the information in [“Notices” on page 45](#).

Product information

This document applies to IBM® QRadar® Security Intelligence Platform V7.3.3 and subsequent releases unless superseded by an updated version of this document.

© **Copyright International Business Machines Corporation 2016, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- Introduction to QRadar deployments.....V**

- Chapter 1. QRadar architecture overview..... 1**
 - QRadar components..... 4
 - QRadar events and flows..... 5

- Chapter 2. QRadar deployment overview.....11**
 - All-in-One deployment..... 12
 - Expanding deployments to add more capacity.....13
 - Adding remote collectors to a deployment..... 14
 - Adding processing capacity to an All-in-One deployment.....15
 - Geographically distributed deployments..... 17
 - QRadar Vulnerability Manager deployments..... 19
 - QRadar Risk Manager and QRadar Vulnerability Manager 23
 - Forensics and full packet collection..... 25
 - Forwarding packets to QRadar Packet Capture.....31

- Chapter 3. Data Nodes and data storage..... 35**

- Chapter 4. App Host.....39**

- Chapter 5. HA deployment overview..... 41**

- Chapter 6. Backup strategies..... 43**
 - QRadar data backups.....43
 - Retention settings..... 43
 - Backup location..... 43

- Notices.....45**
 - Trademarks..... 46
 - Terms and conditions for product documentation..... 46
 - IBM Online Privacy Statement..... 47
 - General Data Protection Regulation..... 47
 - Privacy policy considerations 48

Introduction to QRadar deployments

The IBM QRadar Deployment Guide helps you plan your QRadar installation.

Intended audience

This information is intended for use by security administrators who are responsible for investigating and managing network security. To use this guide you must have a knowledge of your corporate network infrastructure and networking technologies.

Technical documentation

For information about how to access more technical documentation, technical notes, and release notes, see [Accessing IBM Security Documentation Technical Note](http://www.ibm.com/support/docview.wss?rs=0&uid=swg21612861) (<http://www.ibm.com/support/docview.wss?rs=0&uid=swg21612861>).

Contacting customer support

For information about contacting customer support, see the [Support and Download Technical Note](http://www.ibm.com/support/docview.wss?rs=0&uid=swg21612861) (<http://www.ibm.com/support/docview.wss?rs=0&uid=swg21612861>).

Statement of good security practices

IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Please Note:

Use of this Program may implicate various laws or regulations, including those related to privacy, data protection, employment, and electronic communications and storage. IBM QRadar may be used only for lawful purposes and in a lawful manner. Customer agrees to use this Program pursuant to, and assumes all responsibility for complying with, applicable laws, regulations and policies. Licensee represents that it will obtain or has obtained any consents, permissions, or licenses required to enable its lawful use of IBM QRadar.

Chapter 1. QRadar architecture overview

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

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

IBM QRadar SIEM (Security Information and Event Management) 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 QRadar to meet your log and flow collection, and analysis needs. You can add integrated modules to your QRadar platform, such as QRadar Risk Manager, QRadar Vulnerability Manager, and QRadar Incident Forensics.

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

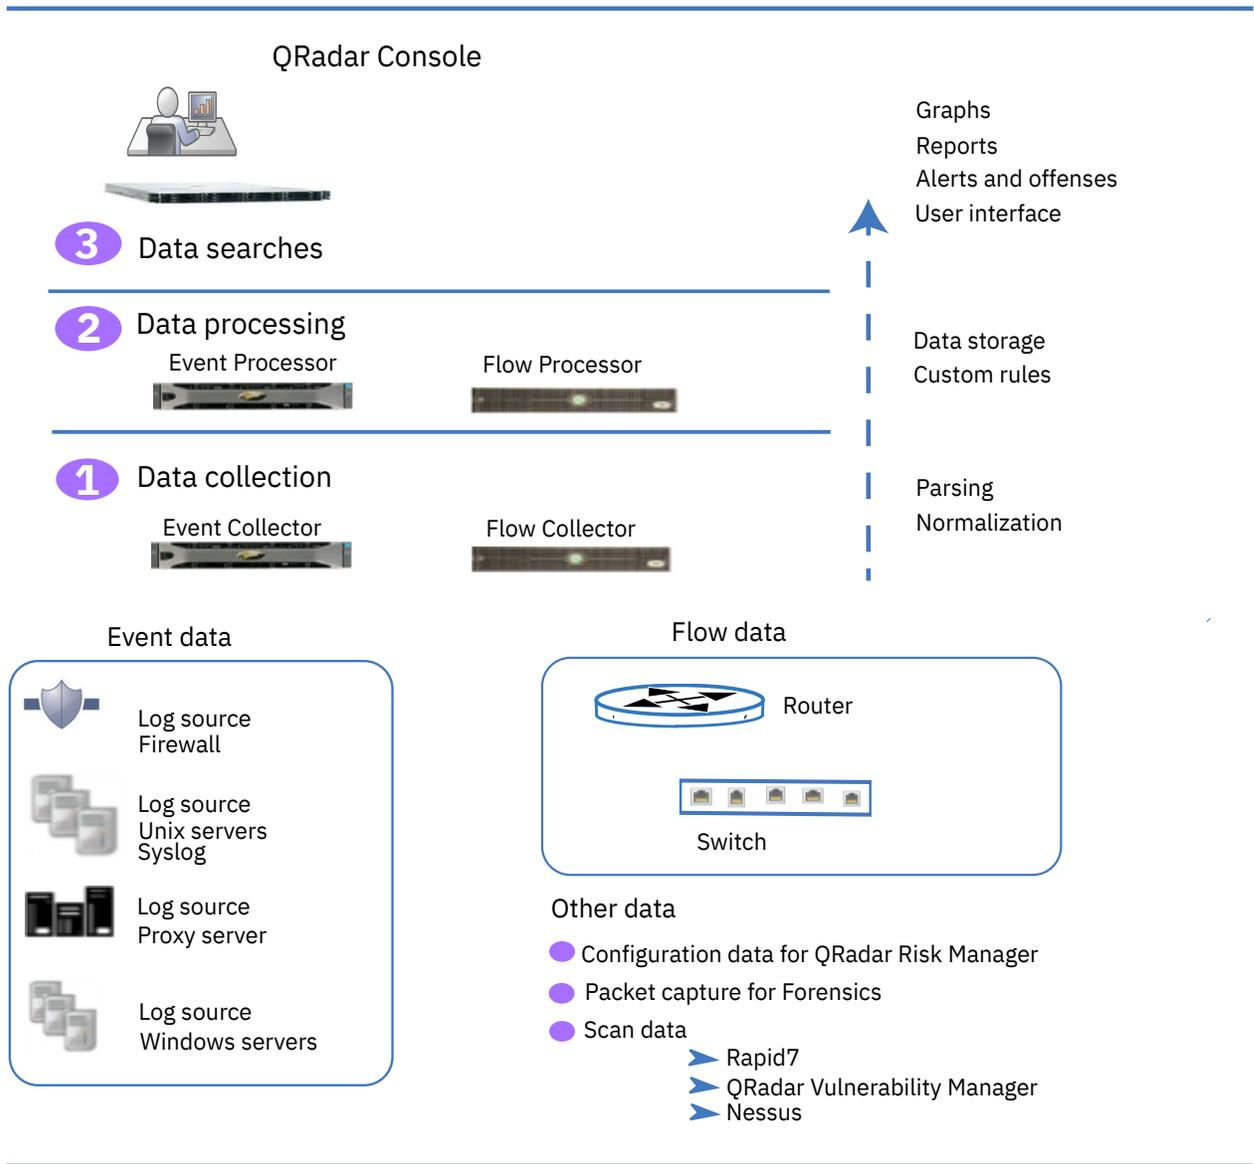


Figure 1. QRadar architecture

The QRadar 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 QRadar 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 QRadar Event Collectors or QRadar QFlow Collectors to collect event or flow data. The data is parsed and normalized before it is 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 QRadar SIEM 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 QRadar translates in to flow records. QRadar 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. In addition to collecting flow information with a Flow Collector, full packet capture is available with the QRadar Incident Forensics component.

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 QRadar Risk Manager (QRM), QRadar Vulnerability Manager (QVM), or QRadar Incident Forensics collect different types of data and provide more functions.

QRadar 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 QRadar 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.

Use QRadar Incident Forensics to perform in-depth forensic investigations, and replay full network sessions.

Data searches

In the third or top layer, data that is collected and processed by QRadar 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 QRadar Console.

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

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

QRadar components

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

Important: Software versions for all IBM QRadar appliances in a deployment must be same version and fix pack level. 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.

QRadar deployments can include the following components:

QRadar Console

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

In distributed QRadar deployments, use the QRadar Console to manage hosts that include other components.

QRadar 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 QRadar. The Event Collector bundles or coalesces identical events to conserve system usage and sends the data to the Event Processor.

- Use the QRadar Event Collector 1501 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.

QRadar 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 QRadar 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.

QRadar QFlow Collector

The Flow Collector collects flows by connecting to a SPAN port, or a network TAP. The IBM QRadar QFlow Collector also supports the collection of external flow-based data sources, such as NetFlow from routers.

QRadar QFlow Collectors are not designed to be full packet capture systems. For full packet capture, review the QRadar Incident Forensics option. The QRadar QFlow Collector 1310 appliance specifically, can forward packets to a QRadar Packet Capture appliance, which allows for flow collection and packet collection from a single packet source.

You can install a QRadar QFlow Collector on your own hardware or use one of the QRadar QFlow Collector appliances.

Restriction: The QRadar Log Manager does not support flow collection or Flow Collectors, which is supported only in QRadar SIEM deployments.

QRadar Flow Processor

The Flow Processor processes flows from one or more QRadar QFlow Collector 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 QRadar 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.

QRadar Data Node

Data Nodes enable new and existing QRadar 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 QRadar 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 QRadar components, see the *IBM QRadar Administration Guide*.

For more information about QRadar appliance specifications, see the *IBM QRadar Hardware Guide*.

QRadar events and flows

The core functions of IBM QRadar SIEM 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

QRadar 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.

QRadar accepts events from log sources by using protocols such as syslog, syslog-tcp, and SNMP. QRadar 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 QRadar Console, events are collected from log sources and then processed by the Event Processor. A QRadar All-in-One appliance functions as the Event Collector and Event Processor, in addition to fulfilling the role of the QRadar Console.

QRadar 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.

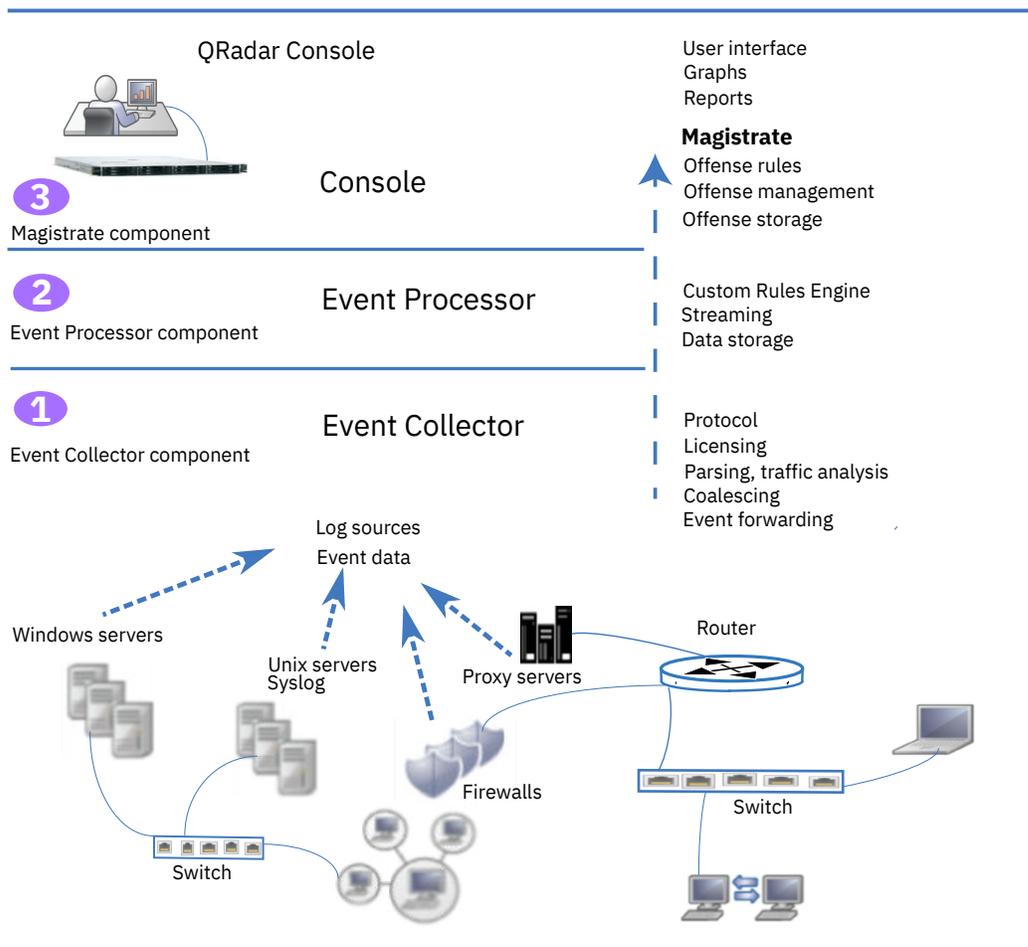


Figure 2. Event pipeline

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 QRadar 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 QRadar can use.

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

QRadar 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 QRadar 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 QRadar 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 QRadar Console that a specific event triggered a rule. The Magistrate component on the QRadar 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 QRadar 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 QRadar Console, the Event Processor executes the action that is defined for the rule response.

Magistrate on the QRadar 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 QRadar Console or All-in-One appliance has a Magistrate component.

Flows

QRadar 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 QRadar that collects and creates flow information is known as QFlow.

QRadar 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 QRadar with the same "First Packet Time" but the "Last Packet Time" values increment through time.

A flow starts when the Flow Collector 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.

QFlow 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 Collector or to a Flow Processor such as the QRadar Flow Processor 1705 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 Collector 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 Collector collects internal flows by connecting to a SPAN port, or a network TAP. The QRadar QFlow Collector 1310 can forward full packets from its 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.

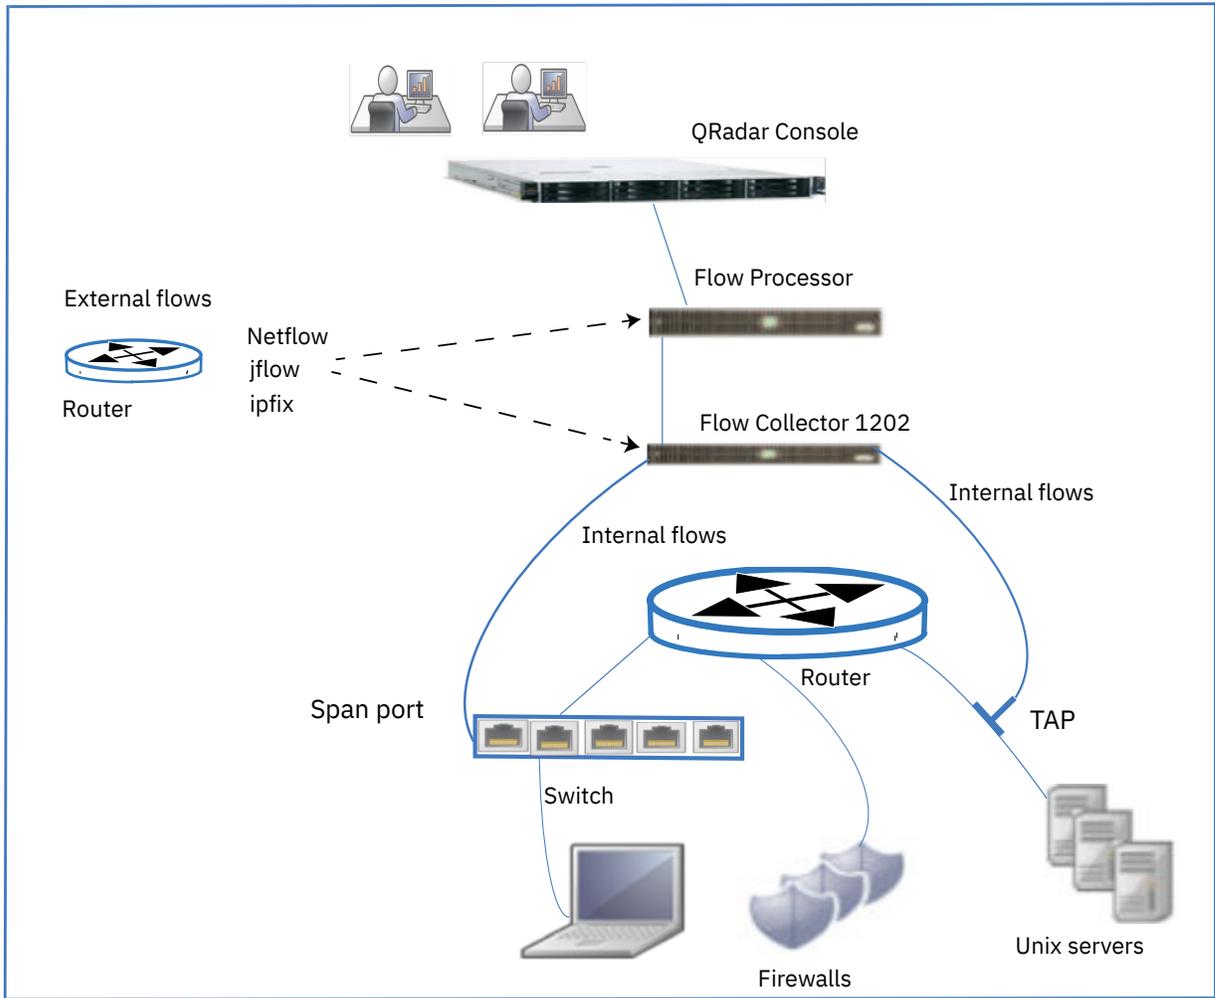


Figure 3. QRadar flows

Flow pipeline

The Flow Collector 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 QRadar 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 Collectors 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.

Chapter 2. QRadar deployment overview

IBM QRadar 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 Collectors, 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 QRadar components, such as Flow Processors, Event Collectors, and Data Nodes, and when you might need to co-locate specific components.

The requirements for your QRadar 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 QRadar components that you can use to collect, process, and store event and flow data in your QRadar 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 Collector 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 QRadar Console. The processor appliances can store data but they can also use the Data Nodes to store data. The QRadar Console appliance is used for monitoring, data searches, reporting, offense management, and administration of your QRadar deployment.

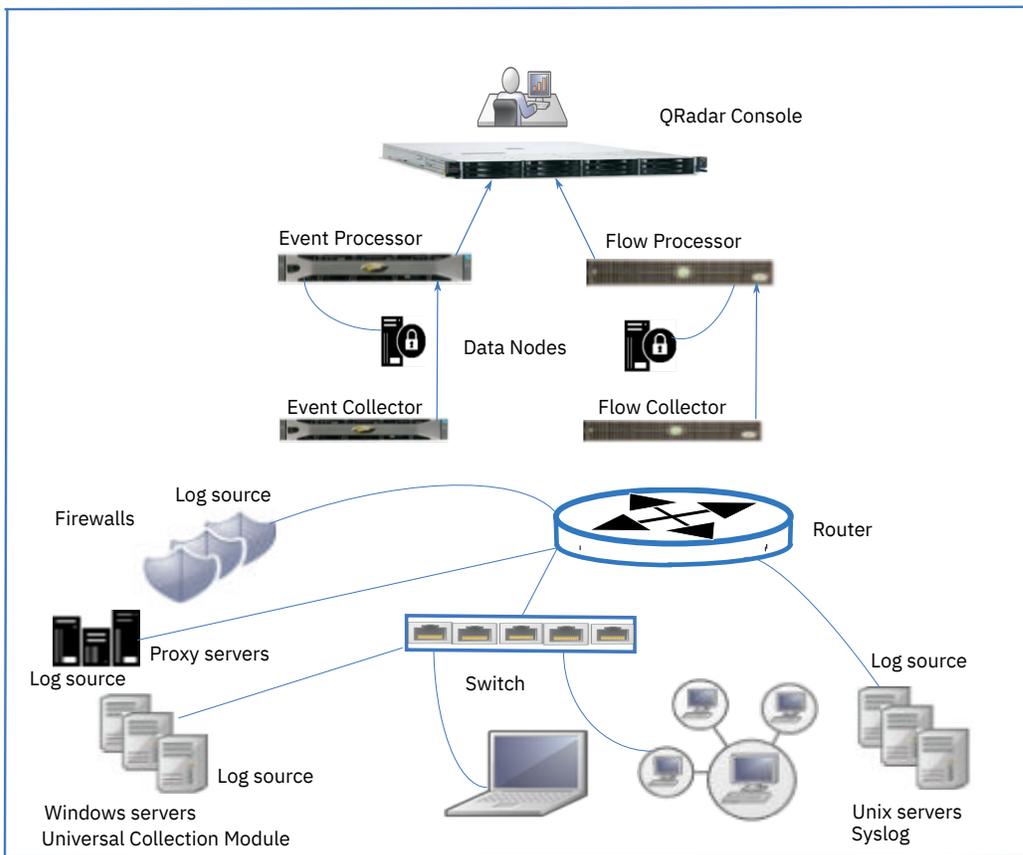


Figure 4. QRadar event and flow components

All-in-One deployment

In a single host QRadar deployment, you have an All-in-One QRadar 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. For example, a QRadar 3105 (All-in-One) typically processes up to 5000 EPS (events per second), and 200,000 FPM (flows per minute), whereas a QRadar 3128 (All-in-One) typically processes up to 15,000 EPS and 300,000 FPM.

Manufacturing company deploys a single QRadar server

You are a medium-sized manufacturing company with less than 1000 employees. You deploy a QRadar 3105 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.

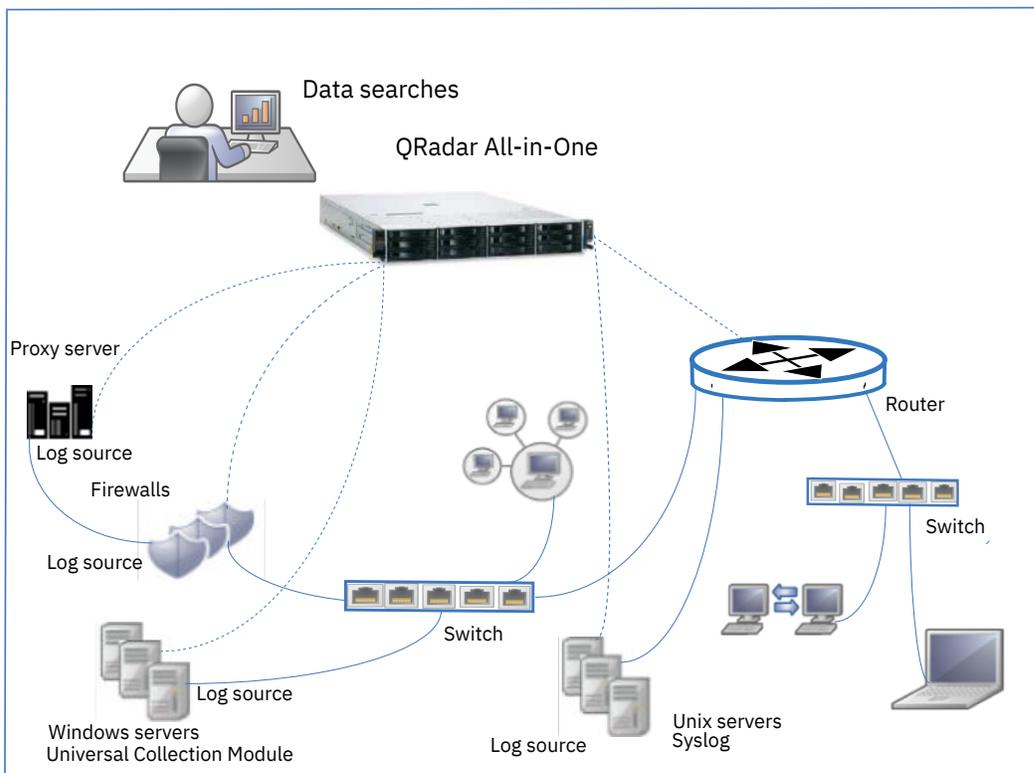


Figure 5. All-in-One deployment

The QRadar 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 QRadar can use.
- Analyzes and stores the data, and identifies security threats to the company.
- Provides access to the QRadar web application.

As your data sources grow, or your processing or storage needs increase, you can add appliances to expand your deployment.

Expanding deployments to add more capacity

Your business might create or expand a deployment beyond an IBM QRadar 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 QRadar 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.

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 collectors 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.

A 3128 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 Collectors to your deployment. For example, you can add a QRadar QFlow Collector 1202, which collects up to 3 Gbps.

An All-in-One appliance processes the events and flows that are collected. By adding Event Collectors and Flow Collectors, 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 Collector 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 QRadar Event Collectors or QRadar Flow Collectors 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 QRadar 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 Collector 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 Collector are added at a remote office.

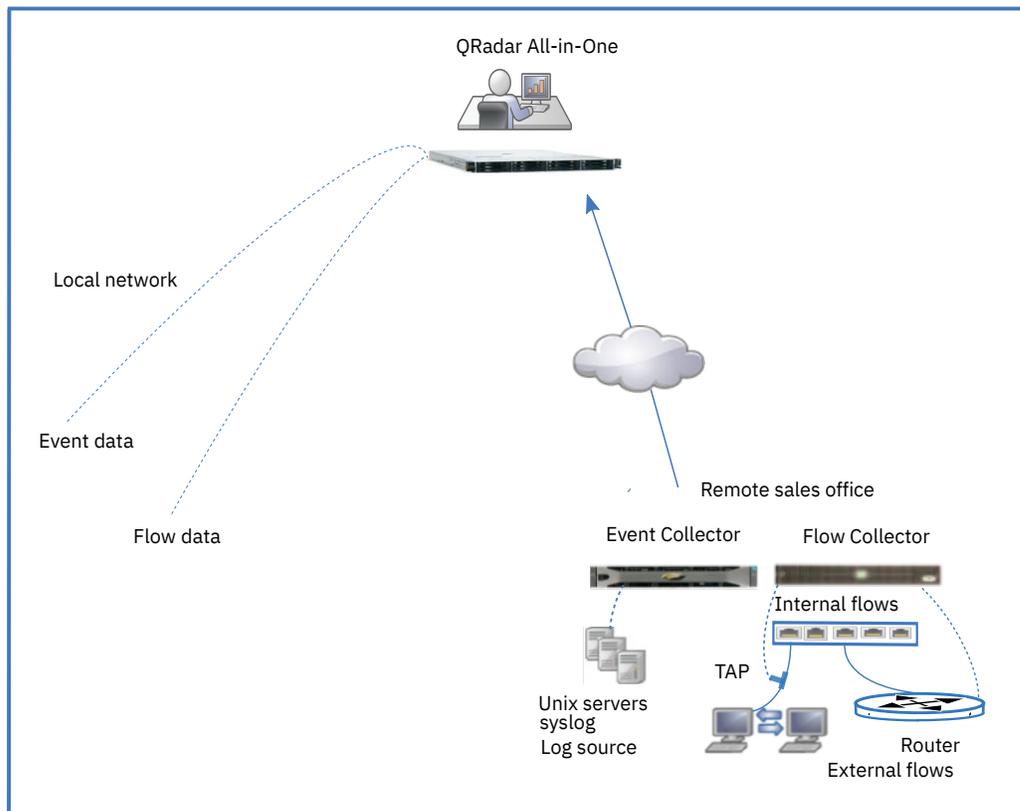


Figure 6. Collectors in remote office

In this deployment, the following processes occur:

- At your remote office, the Event Collector collects data from log sources and the Flow Collector 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 QRadar 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 QRadar deployment to increase processing capacity and increase storage. Adding processors frees up resources on your QRadar 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 QRadar 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 QRadar 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 QRadar 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 QRadar deployment.

When you scale your QRadar 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 QRadar Event Processor to your deployment

You can add a QRadar Event Processor 1628, which collects and processes up to 40,000 EPS. You increase your capacity by another 40,000 EPS every time you add a QRadar Event Processor 1628 to your deployment. Add a QRadar Flow Processor 1728, which collects and processes up to 1,200,000 FPM.

The QRadar Event Processor 1628 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 an QRadar 3128 (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 collectors is processed and stored on the event and flow processors.

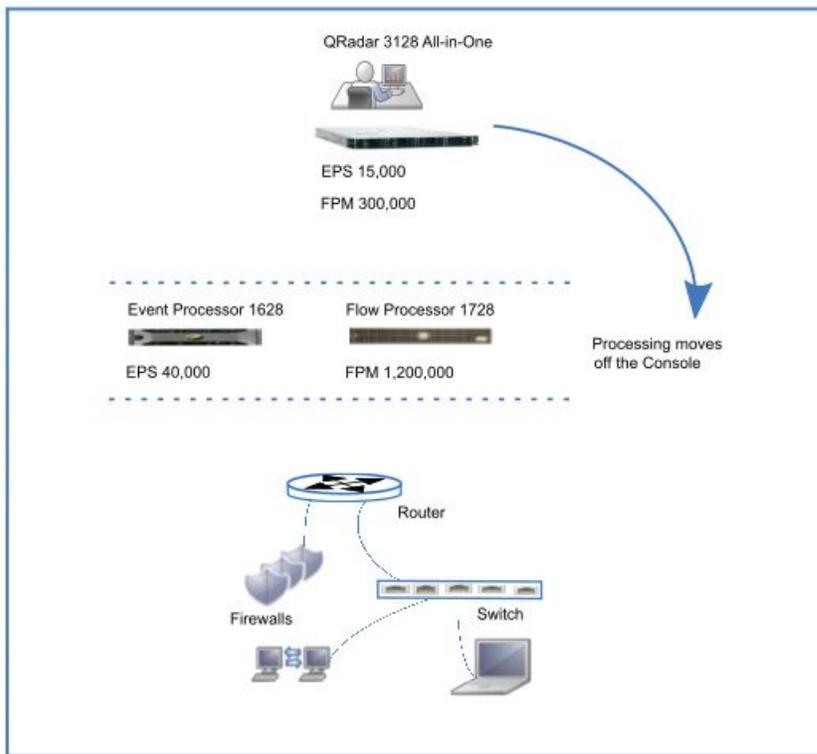


Figure 7. Adding processing capacity

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

Adding processors and collectors expands the processing capacity of your QRadar 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 Collectors to Event Processors and Flow 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 Collectors 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 QRadar searches.

Geographically distributed deployments

In geographically distributed deployments your IBM QRadar 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 QRadar deployment to other countries. Data retention laws vary from country to country, so the QRadar 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.

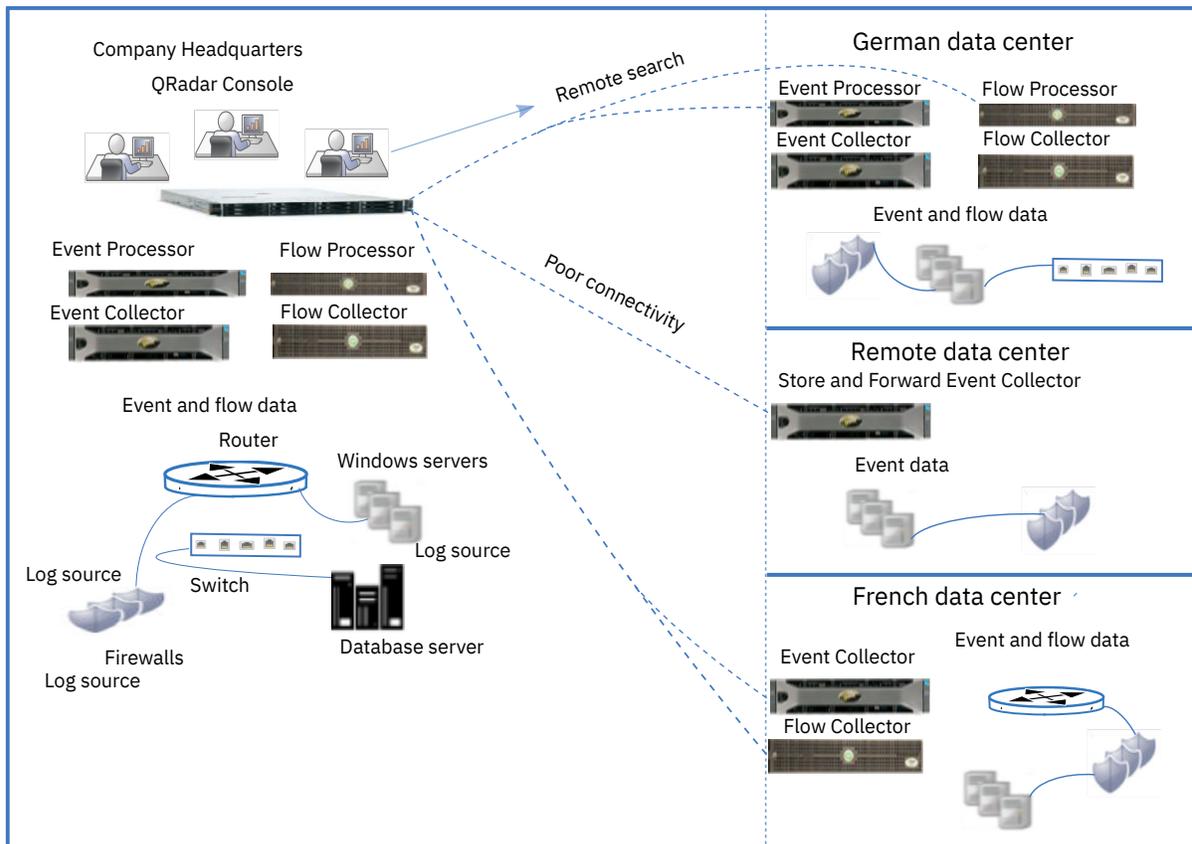


Figure 8. Geographically Distributed Deployment

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 QRadar 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 QRadar Console. If the bandwidth between the QRadar Console and remote processor is not good, you might experience latency with your searches, especially when you're doing multiple concurrent searches.

QRadar Vulnerability Manager deployments

Locate and manage the vulnerabilities in your network by deploying IBM QRadar Vulnerability Manager. Enhance your network security by integrating add-on features such as IBM BigFix® and IBM Security SiteProtector.

IBM QRadar 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 QRadar 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. QRadar Vulnerability Manager and QRadar 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 IBM QRadar or install a new system, the **Vulnerabilities** tab might not be displayed. Access IBM QRadar Vulnerability Manager by using the **Vulnerabilities** tab. If you install IBM QRadar SIEM, the **Vulnerabilities** tab is enabled by default with a temporary license key. If you install QRadar Log Manager, the **Vulnerabilities** tab is not enabled. You can use the **Try it Out** option to try out QRadar Vulnerability Manager for 30 days. You can purchase the license for QRadar Vulnerability Manager separately and enable it by using a license key. For more information about upgrading, see the *IBM QRadar Upgrade Guide*.

QRadar Vulnerability Manager components

The following information describes the QRadar 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 QRadar deployment.
- When you install and license QRadar Vulnerability Manager on an All-in-One system, a vulnerability processor is automatically deployed on your QRadar 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 600 managed host appliance, and QRadar Vulnerability Manager is used for the first time, then the scan processor is assigned to the 600 managed host appliance.
- The scanning processor is governed by the processing license, which determines the maximum number of assets that can be processed by QRadar Vulnerability Manager.
- The scan processor can run on the QRadar Console or a managed host.

The following information describes the QRadar Vulnerability Manager scanner.

- You can deploy a scanner on a virtual machine or as software only.
- You can deploy a QRadar Vulnerability Manager scanner dedicated scanner appliance, which is a 610 appliance.
- You can deploy a scanner on a QRadar Console or on the following managed hosts: Flow Collector, 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.

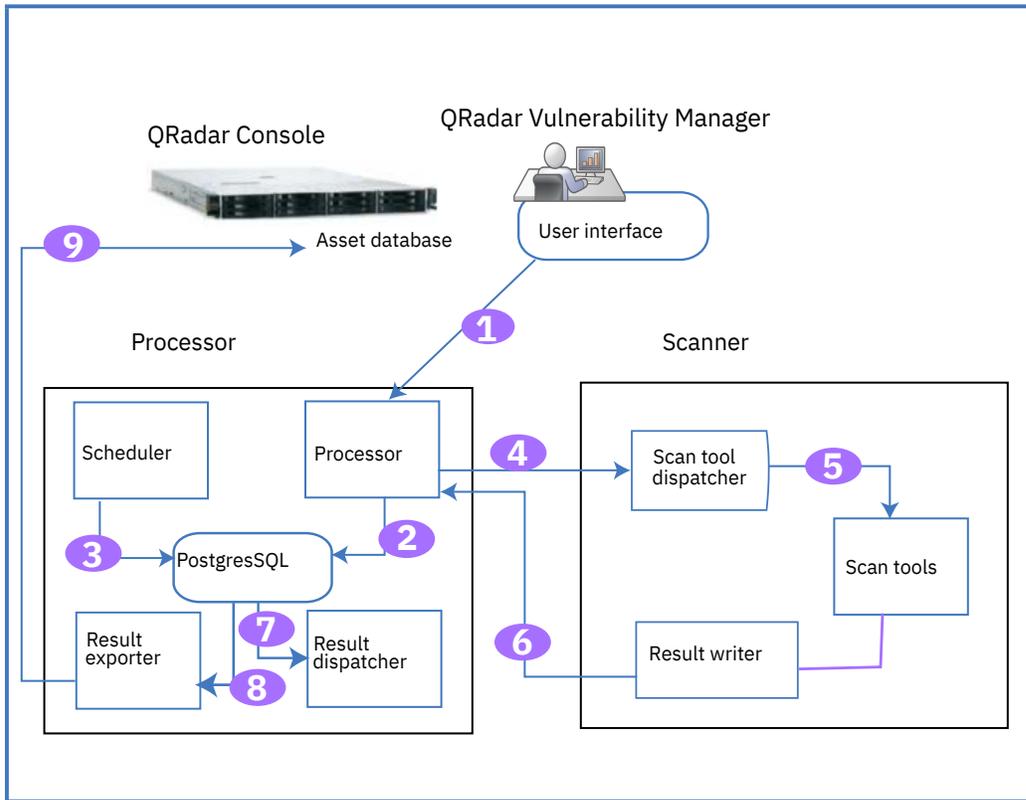


Figure 9. Scan components and process

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.

QRadar 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 PostgreSQL database.
8. The result exporter finds completed scans in the processor database and exports the results to the QRadar Console.
9. The exported results are added to the QRadar database where users can view and manage the scan results.

All-in-one deployment

You can run QRadar 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 QRadar.
- 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 QRadar 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 QRadar 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 bypasses the firewall.

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

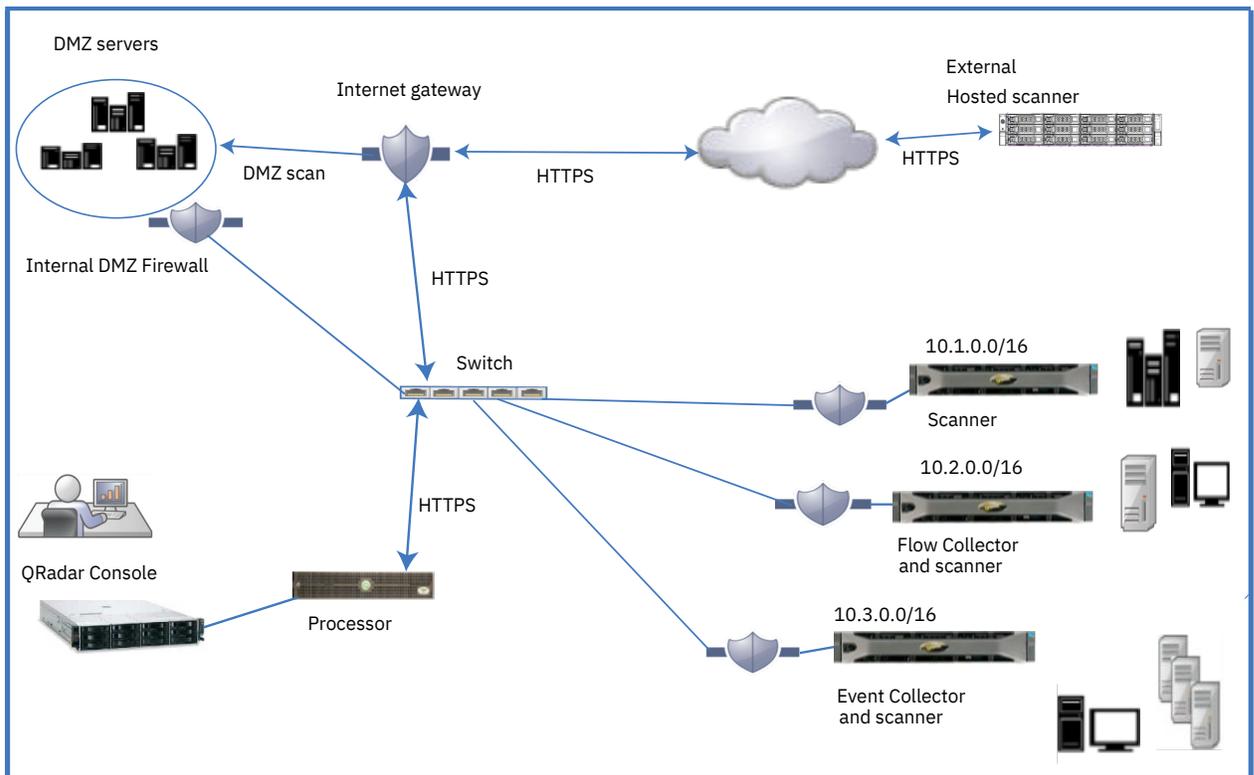


Figure 10. Scanning deployment

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 QRadar Vulnerability Manager with a hosted IBM scanner that is located outside your network. For more information, see the *IBM QRadar Vulnerability Manager User Guide*.

QRadar Vulnerability Manager integrations

IBM QRadar Vulnerability Manager integrates with IBM 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 QRadar Vulnerability Manager to BigFix. Fixlets are packages that you deploy to your assets or endpoints to remediate specific vulnerabilities.

QRadar 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 QRadar Vulnerability Manager scans that are run only after the integration is configured. Connecting to IBM Security SiteProtector.

Third-party scanners

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

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

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

You require QRadar Risk Manager to get the following options:

- Asset, vulnerability, and traffic-based vulnerability management
- Adjusted vulnerability scores and context aware risk scoring.

QRadar Risk Manager and QRadar Vulnerability Manager

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

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

Add a QRadar Risk Manager 700 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).

QRadar Risk Manager appliance

Install QRadar Risk Manager separately on a QRadar Risk Manager 700 appliance.

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

You require only one QRadar Risk Manager appliance per deployment.

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

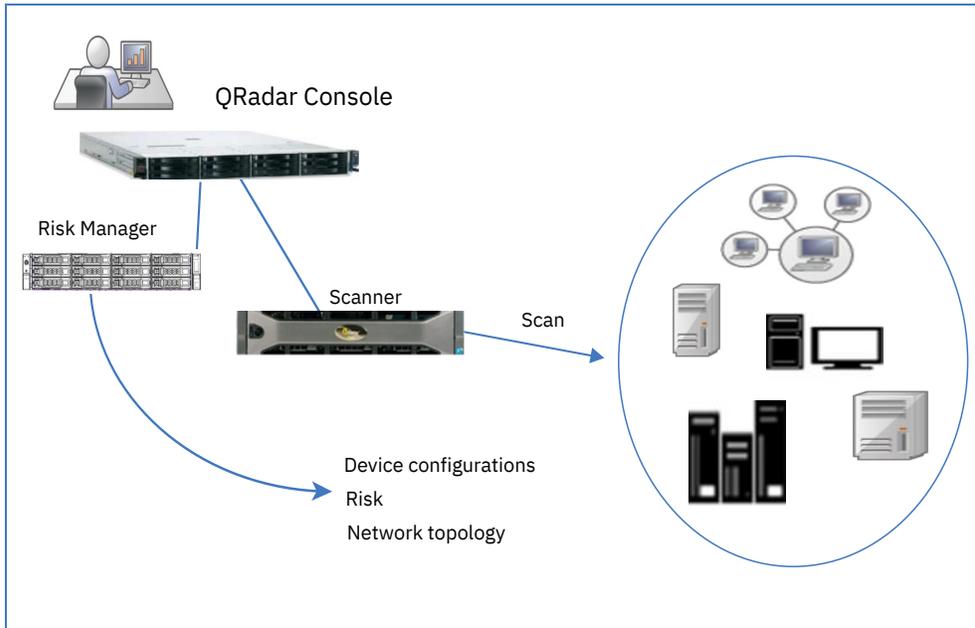


Figure 11. Scanning deployment with Risk Manager

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.

Forensics and full packet collection

Use IBM QRadar Incident Forensics in your deployment to retrace the step-by-step actions of a potential attacker, and conduct an in-depth forensics investigation of suspected malicious network security incidents.

QRadar Incident Forensics reconstructs raw network data that is related to a security incident back into its original form.

QRadar Incident Forensics integrates with the IBM QRadar Security Intelligence Platform and is compatible with many third-party packet capture offerings.

QRadar Incident Forensics offers an optional QRadar Packet Capture appliance to store and manage data that is used by QRadar Incident Forensics if no other network packet capture (PCAP) device is deployed. Any number of these appliances can be installed as a tap on a network or sub-network to collect the raw packet data.

QRadar Packet Capture components

The following components can be included in a QRadar deployment:

QRadar Console

Provides the QRadar product user interface. In distributed deployments, use the QRadar Console to manage multiple QRadar Incident Forensics Processor hosts.

QRadar Incident Forensics Processor

Provides the QRadar Incident Forensics product interface. The interface delivers tools to retrace the step-by-step actions of cyber criminals, reconstruct raw network data that is related to a security incident, search across available unstructured data, and visually reconstruct sessions and events.

You must add QRadar Incident Forensics Processor as a managed host before you can use the security intelligence forensics capability.

QRadar Incident Forensics Standalone

Provides the QRadar Incident Forensics product user interface. Installing QRadar Incident Forensics Standalone provides the tools that you need to do forensics investigations. Only forensics investigative and the related administrative functions are available.

QRadar Packet Capture

You can install an optional QRadar Packet Capture appliance. If no other network packet capture (PCAP) device is deployed, you can use this appliance to store data that is used by QRadar Incident Forensics. You can install any number of these appliances as a network tap or subnetwork to collect the raw packet data.

If no packet capture device is attached, you can manually upload the packet capture files in the user interface or by using FTP.

Depending on your network and packet capture requirements, you can connect up to five packet capture devices to a QRadar Incident Forensics appliance.

QRadar Packet Capture Data Node appliances

For extra storage capacity, you can connect up to two QRadar Packet Capture Data Node appliances to each QRadar Packet Capture master system.

All-in-One deployment

In standalone or all-in-one deployments, you install the IBM QRadar Incident Forensics Standalone software. These single appliance deployments are similar to installing the QRadar Console and QRadar Incident Forensics managed host on one appliance, but without log management, network activity monitoring, or other security intelligence features. For a stand-alone network forensics solution, install the QRadar Incident Forensics Standalone in small to midsize deployments.

The following diagram shows a basic QRadar Incident Forensics All-in-One deployment.

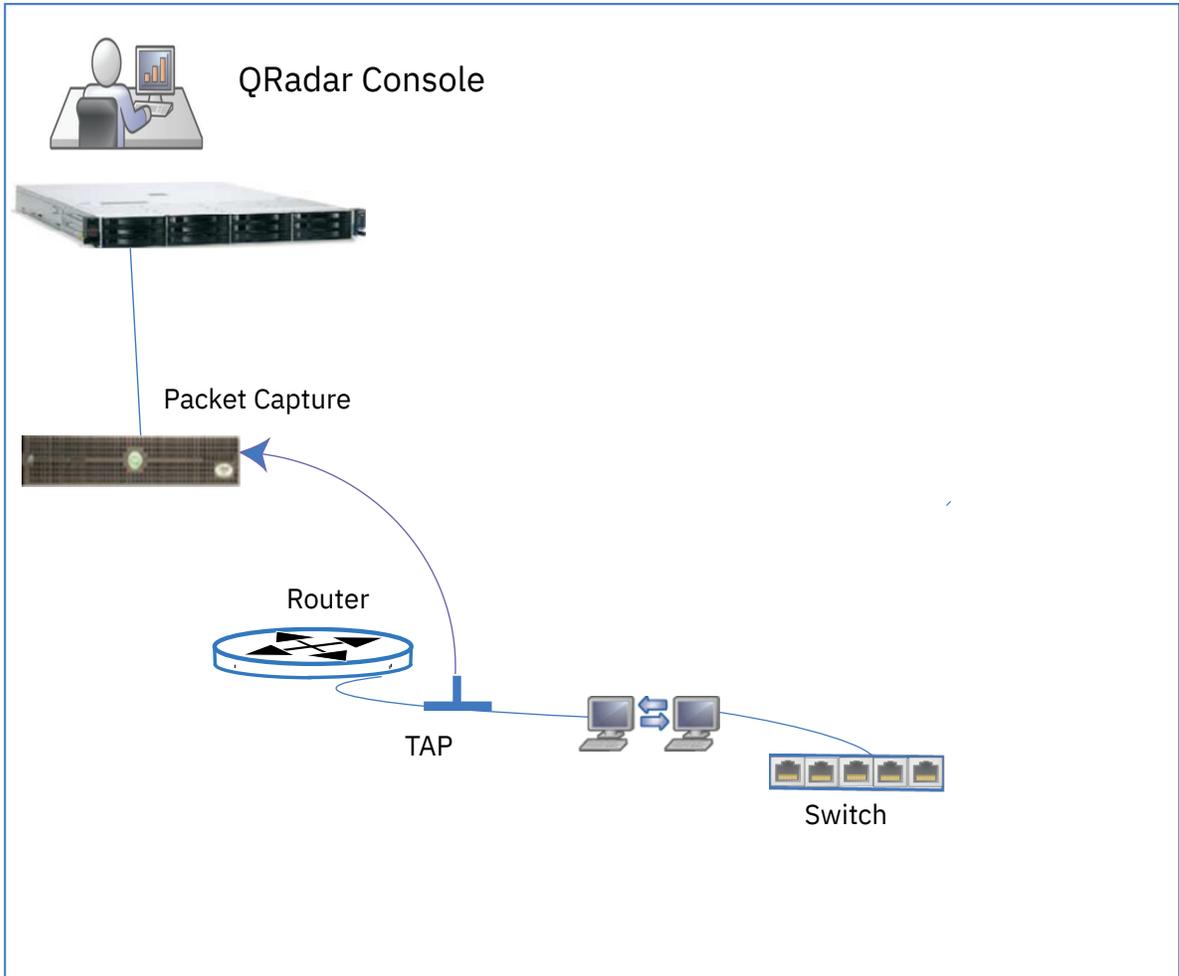


Figure 12. All-in-one deployment

Distributed deployment

In a distributed deployment, you can have the following three appliances:

- QRadar Console
- QRadar Packet Capture managed host (QRadar Packet Capture processor)
- QRadar Packet Capture (optional)

Software versions for all IBM QRadar appliances in a deployment must be the same version and fix level. Deployments that use different versions of software are not supported.

The following diagram shows a QRadar Incident Forensics distributed deployment.

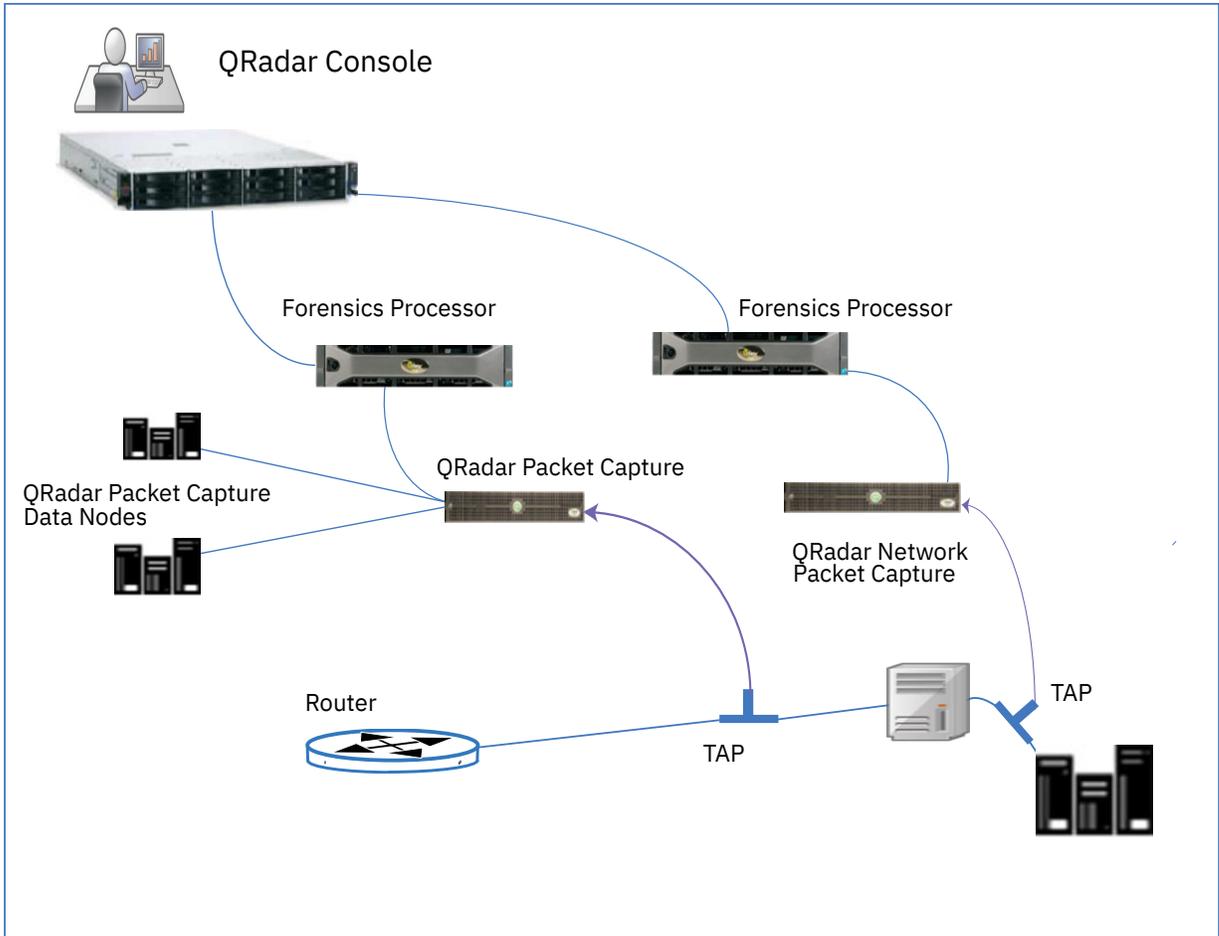


Figure 13. Distributed deployment

The following diagram shows packet forwarding from a IBM QRadar QFlow Collector 1310 with a 10G Napatech network card to a QRadar Packet Capture appliance.

The QRadar QFlow Collector uses a dedicated Napatech monitoring card to copy incoming packets from one port on the card to a second port that connects to a IBM QRadar Packet Capture appliance.

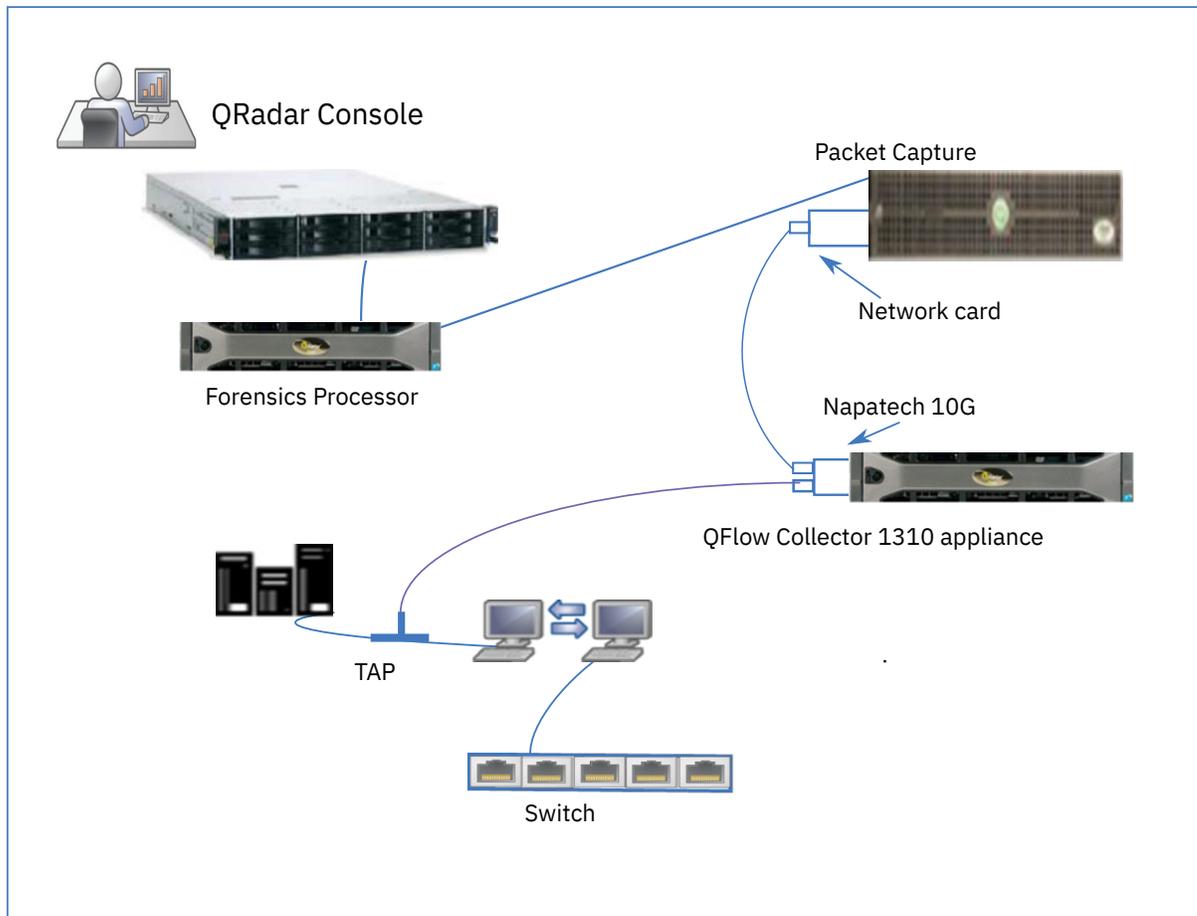


Figure 14. Packet forwarding

Forwarding packets to QRadar Packet Capture

You can monitor network traffic by sending raw data packets to a IBM QRadar QFlow Collector 1310 appliance. The QRadar QFlow Collector uses a dedicated Napatech monitoring card to copy incoming packets from one port on the card to a second port that connects to a IBM QRadar Packet Capture appliance.

If you already have a QRadar QFlow Collector 1310 with a 10G Napatech network card, you can mirror the traffic to QRadar Packet Capture.

As shown in the following diagram, if you already have a QRadar QFlow Collector 1310 with a 10G Napatech network card, you can mirror the traffic to QRadar Packet Capture.

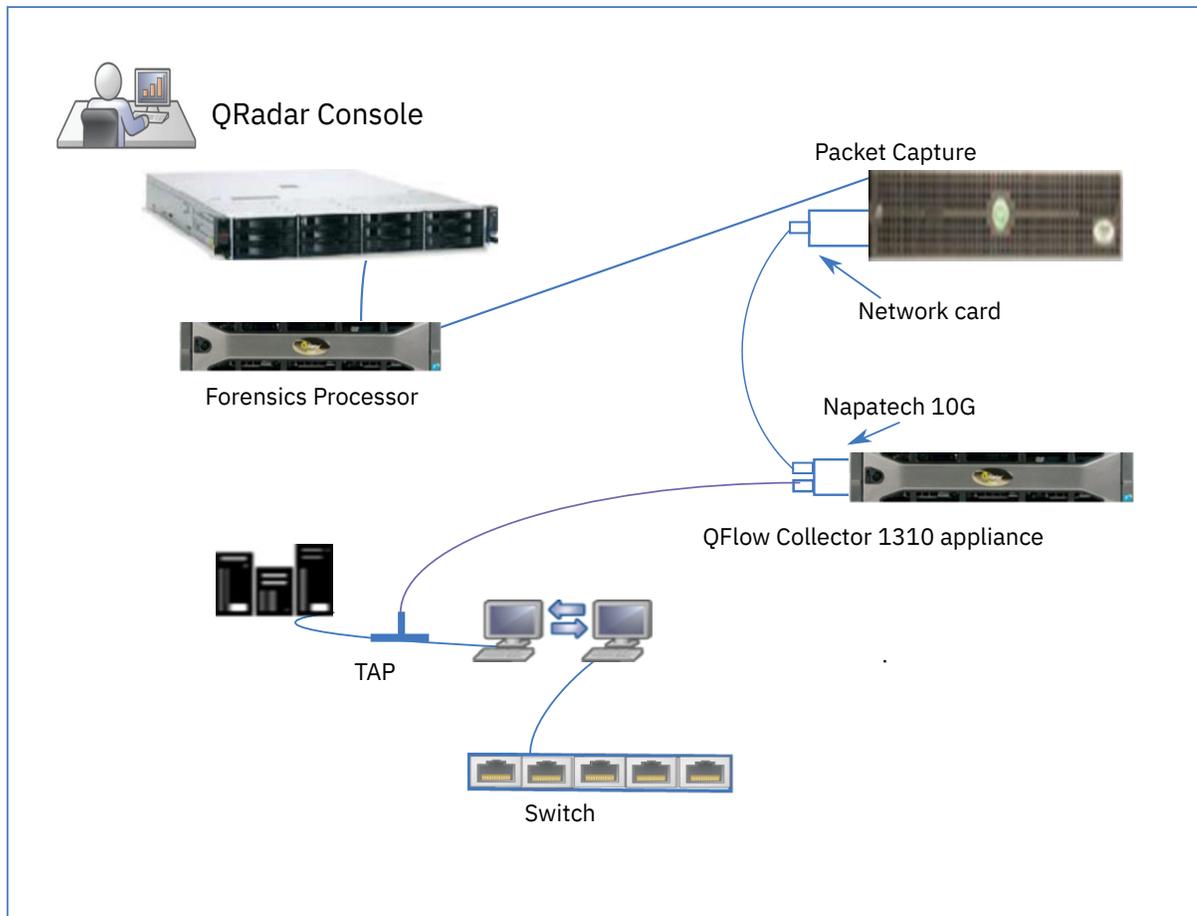


Figure 15. Packet data forwarding from a QRadar QFlow Collector to QRadar Packet Capture by using the Napatech card

Before you begin

Ensure that the following hardware is set up in your environment:

- You attached the cable to port 1 of the Napatech card on the QRadar QFlow Collector 1310 appliance.
- You attached the cable that is connected to port 2 of the Napatech card, which is the forwarding port, to the QRadar Packet Capture appliance.
- Verify layer 2 connectivity by checking for link lights on both appliances.

Procedure

1. Using SSH from your IBM QRadar Console, log in to QRadar QFlow Collector as the root user. On the QRadar QFlow Collector appliance, edit the following file.

```
/opt/qradar/init/apply_tunings
```

- a) Locate the following line, which is around line 137.

```
apply_multithread_qflow_changes()
{
    APPLIANCEID=`$NVABIN/myver -a`
    if [ "$APPLIANCEID" == "1310" ]; then
        MODELNUM=$(/opt/napatech/bin/AdapterInfo 2>&1 | grep "Active FPGA Image" | cut -
d'-' -f2)
        if [ "$MODELNUM" == "9220" ]; then..
```

- b) In the AppendToConf lines that follow the code in the preceding step, add these lines:

```
AppendToConf SV_NAPATECH_FORWARD YES
AppendToConf SV_NAPATECH_FORWARD_INTERFACE_SRCDEST "0:1"
```

These statements enable packet forwarding, and forward packets from port 0 to port 1.

- c) Ensure that *multithreading* is enabled, by verifying that the following line is in the `/opt/qradar/conf/nva.conf` file.

```
MULTI_THREAD_ON=YES
```

2. Run the `apply_tunings` script to update the configuration files on the QRadar QFlow Collector, by typing the following command:

```
./apply_tunings restart
```

3. Restart IBM QRadar services by typing the following command:

```
systemctl restart hostcontext
```

4. Optional: Verify that your Napatech card is receiving and transmitting data.

- a) To verify that the Napatech card is receiving data, type the following command:

```
/opt/napatech/bin/Statistics -dec -interactive
```

The "RX" packet and byte statistics increment if the card is receiving data.

- b) To verify that the Napatech card is transmitting data, type the following command:

```
/opt/napatech/bin/Statistics -dec -interactive
```

The "TX" statistics increment if the card is transmitting data.

5. Optional: Verify that your QRadar Packet Capture is receiving packets from your QRadar QFlow Collector appliance.

- a) Using SSH from your QRadar Console, log in to your QRadar Packet Capture appliance as root on port 4477.

- b) Verify that the QRadar Packet Capture appliance is receiving packets by typing the following command:

```
watch -d cat /var/www/html/statisdata/int0.txt
```

The `int0.txt` file updates as data flows into your QRadar Packet Capture appliance.
For more information about packet capture, see the *IBM QRadar Packet Capture Quick Reference Guide*.

Chapter 3. Data Nodes and data storage

IBM QRadar 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 QRadar V.7.2.7, 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 QRadar.

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

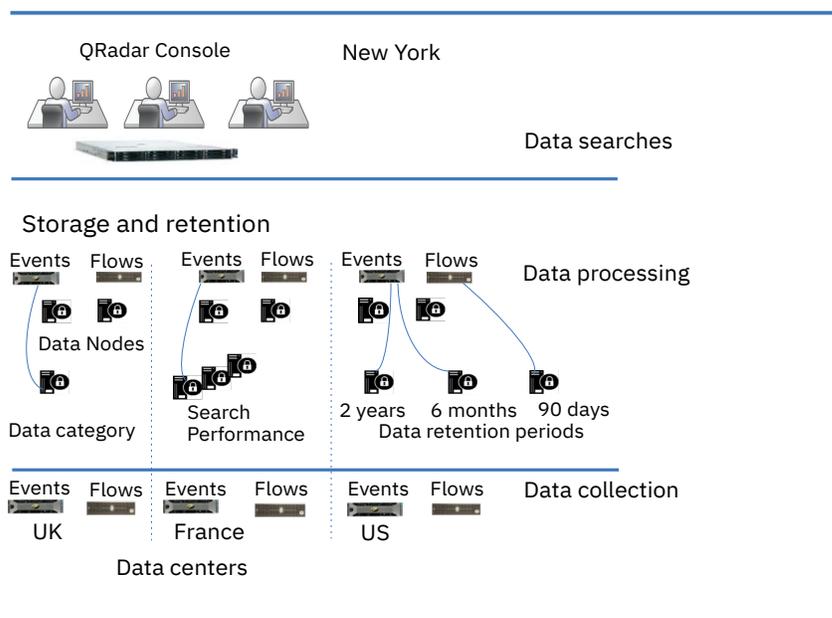


Figure 16. Using Data Node appliances to manage your data storage

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 QRadar V7.2.2 and later.
- Data Nodes perform similar search and analytic functions as event and flow processors in a QRadar 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 XX05, XX28, and XX29 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 QRadar appliances that have an Event Processor or Flow Processor component, including All-In-One appliances. Data Nodes are not compatible with QRadar Incident Forensics PCAP 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 QRadar appliance. Associate Data Nodes with either an event or flow processor. For more information, see *Configuring a managed host* in the *IBM QRadar 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 QRadar 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 QRadar V7.2.3, 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. The following firewall ports must be opened for Data Node communication with the Event Processor:

- Port 32006 between Data Nodes and the Event Processor appliance.
- Port 32011 between Data Nodes and the Console's Event Processor.

Management and Operations

Data Nodes are self-managed and require no regular user intervention to maintain normal operation. QRadar 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 QRadar user interface.

For catastrophic failures that require appliance replacement or the reinstallation of QRadar, 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, Fiber Channel, and NFS (Network File System). You must use iSCSI or Fiber Channel 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 QRadar performance.

After migration, all data I/O to the `/store` file system is no longer done on the local disk. Before you move your QRadar 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, or Fiber Channel storage mount.

For more information about offboard storage, see the *IBM QRadar Security Intelligence Offboard Storage Guide*.

Chapter 4. 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 QRadar 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 QRadar V7.3.2 to replace the deprecated App Node.
- You cannot upgrade to V7.3.2 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.

	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 QRadar 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.

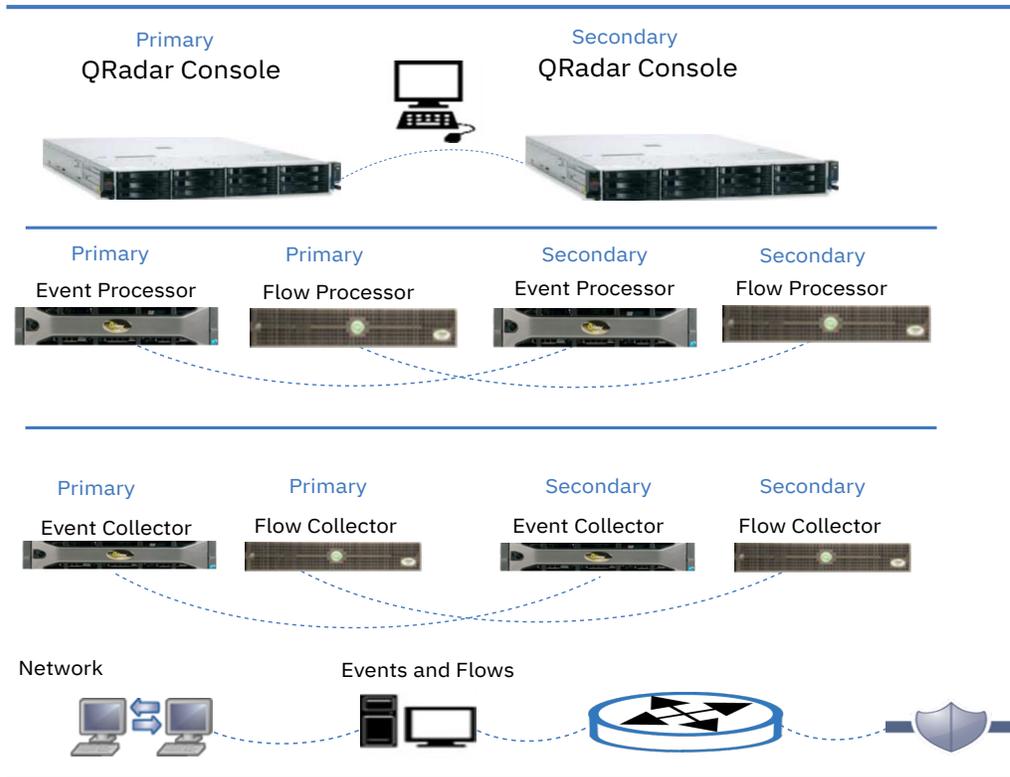
Chapter 5. HA deployment overview

Implement high availability (HA) in your IBM QRadar deployment to keep QRadar 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, QRadar connects a primary HA host with a secondary HA host to create an HA cluster.

The following diagram shows a 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 *IBM Security QRadar SIEM High Availability Guide*.

Chapter 6. Backup strategies

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

QRadar 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 QRadar 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 QRadar 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 QRadar 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 QRadar. 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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions..

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux[®] is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.



Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s session id for purposes of session management and authentication. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

General Data Protection Regulation

Clients are responsible for ensuring their own compliance with various laws and regulations, including the European Union General Data Protection Regulation. Clients are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that may affect the clients’ business and any actions the clients may need to take to comply with such laws and regulations. The products, services, and other capabilities described herein are not suitable for all client situations and may have restricted availability. IBM does not provide legal, accounting or auditing advice or represent or warrant that its services or products will ensure that clients are in compliance with any law or regulation.

Learn more about the IBM GDPR readiness journey and our GDPR capabilities and Offerings here: <https://ibm.com/gdpr>

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s session id for purposes of session management and authentication. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

