

IBM QRadar
7.5.0

High Availability Guide



Note

Before you use this information and the product that it supports, read the information in [“Notices” on page 33](#).

Contents

Introduction.....	V
Chapter 1. HA overview.....	1
Data consistency for HA.....	1
Real-time data synchronization.....	2
Post-failover data synchronization.....	2
High-availability clusters.....	2
Failovers.....	3
Primary HA host failure.....	4
Secondary HA host failure.....	4
HA failover event sequence.....	4
Network connectivity tests.....	4
Heartbeat ping tests.....	5
Primary disk failure.....	5
Manual failovers.....	5
Chapter 2. HA deployment planning.....	7
Firmware update.....	7
Appliance requirements.....	7
System requirements for virtual appliances.....	8
IP addressing and subnets.....	13
Link bandwidth and latency.....	13
Data backup requirements.....	13
Offboard storage requirements for HA.....	14
Chapter 3. HA management.....	15
Status of HA hosts	15
Viewing HA cluster IP addresses.....	17
Creating an HA cluster.....	17
Disconnecting an HA cluster.....	19
Updating the /etc/fstab file.....	20
Editing an HA cluster.....	20
Setting an HA host offline.....	20
Setting an HA host online.....	21
Switching a primary HA host to active.....	21
Chapter 4. Recovery options for HA appliances.....	23
Recovering a failed primary HA host.....	23
Recovering a failed secondary HA host.....	24
Restoring a primary HA host to a previous version or factory default.....	24
Restoring a secondary HA host to a previous version or factory default.....	25
Chapter 5. Troubleshooting QRadar HA deployments.....	27
Restoring a failed secondary HA host.....	27
Restoring a failed primary HA host.....	28
Verifying the status of primary and secondary hosts.....	28
Setting the status of the primary HA host to online.....	29
Chapter 6. Recovery solution for QRadar deployments.....	31

Notices.....	33
Trademarks.....	34
Terms and conditions for product documentation.....	34
IBM Online Privacy Statement.....	35
General Data Protection Regulation.....	35

Introduction to QRadar high-availability deployments

Administrators can protect IBM® QRadar® data by implementing a high-availability (HA) solution.

Intended audience

QRadar SIEM administrators who are responsible for installing and deploying the product must know their corporate network infrastructure, the Linux® operating system, and networking technologies.

Technical documentation

To find IBM QRadar product documentation on the web, including all translated documentation, access the [IBM Knowledge Center](http://www.ibm.com/support/knowledgecenter/SS42VS/welcome) (<http://www.ibm.com/support/knowledgecenter/SS42VS/welcome>).

For information about how to access more technical documentation in the QRadar products library, see [Accessing IBM Security Documentation Technical Note](http://www.ibm.com/support/docview.wss?rs=0&uid=swg21614644) (www.ibm.com/support/docview.wss?rs=0&uid=swg21614644).

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. HA overview

If your hardware or network fails, IBM QRadar can continue to collect, store, and process event and flow data by using high-availability (HA) appliances.

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

If a primary HA host fails, then the secondary HA host maintains access to the same data as the primary by using data synchronization or shared external storage.

The secondary HA host inherits the license from the primary HA host. There is no need to apply a separate license to the secondary host.

For more information about using shared external storage with HA, for example iSCSI, Fibre Channel, or NFS, see the *IBM Security QRadar Offboard Storage Guide*.

Unless otherwise noted, all references to QRadar refer to QRadar SIEM and IBM QRadar Log Manager.

You can use HA on hardware or virtual appliances, and with either appliance or software installations, if you meet the HA requirements. HA is not supported in cloud environments. Security Technical Implementation Guide (STIG) is not supported in QRadar high-availability (HA) deployments.

Related concepts

High-availability clusters

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

Data consistency for HA

When an HA failover occurs, IBM QRadar ensures the consistency of your data.

Data consistency for HA

When an HA failover occurs, IBM QRadar ensures the consistency of your data.

The type of storage that you use determines how HA data consistency is maintained. If you configure HA with external storage, data consistency is maintained by using a component such as an iSCSI or Fibre Channel external storage device. See [“Offboard storage requirements for HA” on page 14](#).

If you do not use external storage devices, then QRadar HA maintains data consistency between a primary and secondary HA host by using Distributed Replicated Block Device.

Data synchronization occurs in the following situations in an HA environment:

- When you initially configure an HA cluster.
- When a primary HA host is restored after a failover.
- During normal HA operation, data is synchronized in real time between the primary and secondary host.

Related concepts

HA overview

If your hardware or network fails, IBM QRadar can continue to collect, store, and process event and flow data by using high-availability (HA) appliances.

Link bandwidth and latency

To configure high-availability (HA), you must consider the bandwidth and latency between the primary and secondary HA hosts.

Status of HA hosts

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

Related information

[How do I set up and synchronize HA hosts? \(Security Learning Academy course\)](#)

Real-time data synchronization

When you configure an HA cluster, the `/store` file system on the primary HA host is automatically synchronized with the `/store` partition on the secondary HA host.

If the primary HA host fails over, the `/store` file system on the secondary HA host is automatically mounted to its local disk, where it continues to read from and write to the data received by the primary HA host before the failover.

After synchronization is complete, the secondary HA host assumes a status of standby.

Depending on the size of the primary `/store` partition and performance, disk synchronization can take an extended time period. Ensure that the connection between the primary and secondary HA host has a minimum bandwidth of 1 Gbps.

Related concepts

Status of HA hosts

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

Related information

[How do I set up and synchronize HA hosts? \(Security Learning Academy course\)](#)

Post-failover data synchronization

Data that is collected by a primary high-availability (HA) host, up to the point of failover, is maintained virtually, in real time, by the secondary HA host.

When the primary HA host is restored after a failure, only the data that is collected by the secondary HA host in the intervening period is synchronized with the primary HA host. Therefore, post-failover disk synchronization is faster than initial disk synchronization, unless the disk on the primary HA host was replaced or reformatted when the host was manually repaired.

When restored from a failover, the status of the primary HA host becomes offline. You must set the primary HA host to an online state, and set the secondary host to an offline state, before it can become the active host. Disk replication with the secondary HA host is enabled while the primary HA host remains offline.

Related tasks

Setting an HA host online

You can set the primary or secondary HA host to Online.

High-availability clusters

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

Primary HA host

The primary HA host is any console or managed host in your IBM QRadar SIEM deployment that requires protection from data loss if there is a failure.

When you create an HA cluster, the IP address of the primary HA host is automatically reassigned to a cluster virtual IP address. Therefore, you must assign an unused IP address to the primary HA host.

When you create an HA cluster, the original host name of your primary HA host becomes the virtual hostname for the HA cluster, and "-primary" is appended to the host name of your primary HA host. You can't change the host name of a host in an HA cluster.

The primary HA host can act as a standby system for the secondary HA host. For example, if the primary HA host is repaired after a failover, the status changes to standby.

Secondary HA host

The secondary HA host is the standby system for the primary HA host.

If the primary HA host fails, the secondary HA host automatically takes over all the responsibilities of the primary HA host.

When you create an HA cluster, the host name of your secondary HA host is changed to `<cluster_host_name>-secondary`. You can't change the host name of a host in an HA cluster.

Virtual IP address

When you create an HA cluster, the cluster virtual IP address takes the IP address of the primary HA host.

Configuring the cluster

Use the HA wizard to configure the primary host, secondary host, and cluster virtual IP address.

The following items are validated when you configure by using the HA wizard:

- The secondary HA host has a valid HA activation key.
- The secondary HA host is not part of another HA cluster
- The software versions on the primary and secondary HA hosts are the same
- If the primary HA host is configured with an external storage device, the secondary HA host is configured to access the same external storage device.
- The primary and secondary HA hosts support the same Device Support Module (DSM), scanner, and protocol RPMs.

Related concepts

[HA overview](#)

If your hardware or network fails, IBM QRadar can continue to collect, store, and process event and flow data by using high-availability (HA) appliances.

[Primary HA host failure](#)

If the secondary high-availability (HA) host detects a primary host failure, it automatically takes over the responsibilities of the primary HA host and becomes the active system.

[Status of HA hosts](#)

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

[IP addressing and subnets](#)

To configure high-availability (HA), you must consider the subnet that is used by the secondary HA host and the virtual IP address.

Related tasks

[Creating an HA cluster](#)

Failovers

When a primary or secondary high-availability (HA) host fails, IBM QRadar maintains data consistency.

The following scenarios cause failover:

- 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 Redundant Array of Independent Disks (RAID) failure on the primary HA host.
- A manual failover.
- NFS volumes that become read-only or not writable.

The following scenarios do not cause an automatic HA failover:

- If a QRadar process develops an error, stops functioning, or exits with an error.
- If a disk on your primary HA host reaches 95% capacity, QRadar data collection stops, but the primary HA host continues to function.

Primary HA host failure

If the secondary high-availability (HA) host detects a primary host failure, it automatically takes over the responsibilities of the primary HA host and becomes the active system.

When a primary HA host is recovered from a failover, it does not automatically take over the active status in the HA cluster. Instead, the secondary HA host remains the active system and the primary host acts as the standby system.

Important: You must switch the primary back to the active status after successfully recovering from a primary failure. See [“Switching a primary HA host to active”](#) on page 21.

Related concepts

[High-availability clusters](#)

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

Secondary HA host failure

If the primary high-availability (HA) host detects a secondary host failure, it automatically assumes the responsibilities of the secondary HA host and becomes the active system.

HA failover event sequence

IBM QRadar initiates a sequence of events when a primary high-availability (HA) host fails.

During failover, the secondary HA host assumes the responsibilities of the primary HA host. The following actions in sequence are completed in sequence:

1. If configured, external shared storage devices are detected and the file systems are mounted. For more information, see the IBM Security *Offboard Storage Guide*.
2. A management interface network alias is created, for example, the network alias for eth0 is eth0:0.
3. The cluster virtual IP address is assigned to the network alias.
4. All QRadar services are started.
5. The secondary HA host connects to the console and downloads configuration files.

Network connectivity tests

To test network connectivity, the IBM QRadar Console automatically pings all existing managed hosts in your QRadar deployment.

If the primary HA QRadar console loses network connectivity to a managed host, but the connection to the secondary HA console remains intact, the HA secondary QRadar console completes another network connectivity test with the managed hosts. If the test succeeds, the primary HA console completes a controlled failover to the secondary HA console. If the test fails, HA failover is not completed because the secondary HA console might also be experiencing network connectivity problems.

Related tasks

[Creating an HA cluster](#)

Heartbeat ping tests

You can test the operation of the primary high-availability (HA) host by configuring the time interval of heartbeat ping tests.

If the secondary HA host does not receive a response from the primary HA host within a preconfigured time period, automatic failover to the secondary HA host is completed.

Related tasks

[Creating an HA cluster](#)

Primary disk failure

If RAID completely fails and all disks are unavailable, the primary HA host completes a shutdown and fails over to the secondary HA host.

After a failover, the primary HA host assumes a status of **Failed**.

Related concepts

[Status of HA hosts](#)

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

Manual failovers

You can manually force a failover from a primary high-availability (HA) host to a secondary HA host.

Manually forcing a failover is useful for planned hardware maintenance on a console or managed host. Ensure the following before you conduct a manual failover:

- The primary and secondary HA hosts are synchronized.
- The secondary HA host has a status of standby.

To perform a manual failover on the primary HA host, set the primary system to offline to make the secondary HA host active. After the secondary host becomes active, you can perform maintenance on the primary host.

To perform a manual failover on the secondary HA host, set the secondary system to offline. After the primary host becomes active, you can perform maintenance on the secondary host.

Do not manually force a failover on a primary HA host when you install patches or install software upgrades. For more information, see the *IBM QRadar Upgrade Guide*.

Related tasks

[Setting an HA host offline](#)

You can set the primary or secondary high-availability (HA) host to **Offline** from the **Active** or **Standby** state.

Chapter 2. HA deployment planning

Plan your high-availability deployment.

Before you implement high-availability (HA), review all the requirements to understand and prepare your IBM QRadar deployment.

Firmware update

Update the firmware on IBM QRadar appliances to take advantage of additional features and updates for the internal hardware components.

For more information about updating firmware, see [Firmware update for QRadar](http://www-01.ibm.com/support/docview.wss?uid=swg27047121) (<http://www-01.ibm.com/support/docview.wss?uid=swg27047121>).

Appliance requirements

Before you add a secondary appliance to your IBM QRadar host, you must review the hardware configuration differences between your primary and secondary appliances.

Appliances that you order as primary and secondary HA pairs are matched to ensure compatibility. However, replacing an appliance or adding HA to an older host with a different hardware configuration can lead to data replication issues. Data replication issues can occur when you replace end of life hardware or create primary and secondary HA pairs that have appliances from different manufacturers.

Partition requirements for /store

The combined size of the /store and /transient partitions on the secondary host must be equal to or larger than the /store partition on the primary host.

For example, do not pair a primary host that uses a 4 TB /store partition to a secondary host that has a 2 TB /store partition and a 1 TB /transient partition.

Storage requirements

Follow these storage requirements when you replace an appliance:

- Ensure that the replacement appliance includes storage capacity that is equal to or greater than the original hardware you replace, and be at least 130 gigabytes (GB).
- Secondary replacement appliances can have larger storage capacity than the primary appliance. If so, partitions on the secondary are resized to match the storage capacity on the primary appliance when you configure the HA pair.
- Primary replacement appliances can have larger storage capacity than the secondary appliance. If so, partitions on the primary are resized to match the storage capacity on the secondary appliance when you configure the HA pair.
- If you replace both primary and secondary appliances, then the system resizes the storage partition that is based on the appliance with the smallest capacity.

Managed interfaces

- The primary host cannot contain more physical interfaces than the secondary.

During a failover, the network configuration of the primary is replicated to the secondary host. If the primary is configured with more interfaces, any additional interfaces cannot be replicated to the secondary during a failover.

- The secondary host must use the same management interface as the primary HA host.

If the primary HA host uses ens192, for example, as the management interface, the secondary HA host must also use ens192.

- The management interface supports one cluster virtual IP address.
- TCP port 7789 must be open and allow communication between the primary and secondary for Distributed Replicated Block Device traffic.

Distributed Replicated Block Device traffic is responsible for disk replication and is bidirectional between the primary and secondary host.

- The QRadar software version must be identical between the primary and secondary host before you pair a primary to a secondary appliance for the first time.

If the QRadar version between your primary and secondary differ, you must patch either the primary or secondary appliance to ensure both appliances use the same software version.

After the primary and secondary appliances are paired together, disk replication ensures that any additional software updates are also applied to the secondary.

- Ensure that the secondary host has a valid HA activation key.

System requirements for virtual appliances

To ensure that IBM QRadar works correctly, you must use virtual appliances that meet the minimum requirements.

For more information about supported hypervisors and virtual hardware versions, see [Creating your virtual machine](#).

QRadar virtual appliances require x86 hardware.

QRadar appliances are certified to support certain maximum events per second (EPS) rates. Maximum EPS depends on the type of data that is processed, system configuration, and system load. For more information, see [QRadar maximum EPS certification methodology](#).

Note: The minimum requirements support QRadar functions with minimum data sets and performance. The minimum requirements support a QRadar system that uses only the default apps. For optimal performance, use the suggested requirements.

For information about running QRadar on 3rd party clouds, see [QRadar cloud marketplace images](#).

QRadar Incident Forensics is installed from a separate ISO than other QRadar appliances. For more information about installing QRadar Incident Forensics as a virtual appliance, see [Virtual appliance installations for QRadar Incident Forensics](#).

QRadar Incident Forensics is installed from a separate ISO than other QRadar appliances. For more information about installing QRadar Incident Forensics as a virtual appliance, see "Virtual appliance installations for QRadar Incident Forensics" in *IBM QRadar Incident Forensics Installation Guide*.

Important: You can change the memory or the CPU of your virtual appliance by shutting down the virtual appliance and making the changes. When you restart the virtual appliance, the system detects the changes and adjusts the performance-related configuration.

Memory requirements

The following table describes the memory requirements for virtual appliances.

Appliance	Minimum memory requirement	Suggested memory requirement
QRadar Flow Virtual 1299	6 GB	6 GB
QRadar Data Node Virtual 1400 appliance	24 GB	64 GB

Table 1. Minimum and suggested memory requirements for QRadar virtual appliances (continued)

Appliance	Minimum memory requirement	Suggested memory requirement
QRadar Event Collector Virtual 1599	12 GB (up to 20,000 EPS) 64 GB (40,000 EPS) 128 GB (80,000 EPS)	16 GB (up to 20,000 EPS) 64 GB (40,000 EPS) 128 GB (80,000 EPS)
QRadar SIEM Event Processor Virtual 1699 up to 20,000 EPS	16 GB	64 GB
QRadar SIEM Event Processor Virtual 1699 20,000 EPS or higher	128 GB	128 GB
QRadar SIEM Flow Processor Virtual 1799 up to 1,200,000 FPM	16 GB	64 GB
QRadar SIEM Flow Processor Virtual 1799 1,200,000 FPM or higher	128 GB	128 GB
QRadar SIEM Event and Flow Processor Virtual 1899 5,000 EPS or less 200,000 FPM or less	16 GB	64 GB
QRadar SIEM Event and Flow Processor Virtual 1899 30,000 EPS or less 1,000,000 FPM or less	128 GB	128 GB
QRadar SIEM All-in-One Virtual 3199 5,000 EPS or less 200,000 FPM or less	32 GB	64 GB
QRadar SIEM All-in-One Virtual 3199 30,000 EPS or less 1,000,000 FPM or less	128 GB	128 GB
QRadar Log Manager Virtual 8099	24 GB	48 GB
QRadar Risk Manager	24 GB	48 GB

Appliance	Minimum memory requirement	Suggested memory requirement
QRadar Vulnerability Manager Processor Important: The IBM QRadar Vulnerability Manager scanner is end of life (EOL) in 7.5.0 Update Package 6, and is no longer supported in any version of IBM QRadar. For more information, see QRadar Vulnerability Manager: End of service product notification (https://www.ibm.com/support/pages/node/6853425) .	32 GB	32 GB
QRadar Vulnerability Manager Scanner Important: The IBM QRadar Vulnerability Manager scanner is end of life (EOL) in 7.5.0 Update Package 6, and is no longer supported in any version of IBM QRadar. For more information, see QRadar Vulnerability Manager: End of service product notification (https://www.ibm.com/support/pages/node/6853425) .	16 GB	16 GB
QRadar App Host	12 GB	64 GB or more for a medium-sized App Host 128 GB or more for a large sized App Host

Processor requirements

The following table describes the CPU requirements for virtual appliances.

QRadar appliance	Threshold	Minimum number of CPU cores	Suggested number of CPU cores
QRadar Flow Virtual 1299	10,000 FPM or less	4	4
QRadar Event Collector Virtual 1599	5,000 EPS or less	8	16
	20,000 EPS or less	19	19
	40,000 EPS or less	40	40
	80,000 EPS or less	80	80

Table 2. CPU requirements for QRadar virtual appliances (continued)

QRadar appliance	Threshold	Minimum number of CPU cores	Suggested number of CPU cores
QRadar SIEM Event Processor Virtual 1699	5,000 EPS or less	8	24
	20,000 EPS or less	16	32
	40,000 EPS or less	40	48
	80,000 EPS or less	56	80
QRadar SIEM Flow Processor Virtual 1799	150,000 FPM or less	4	24
	300,000 FPM or less	8	24
	1,200,000 FPM or less	16	32
	2,400,000 FPM or less	40	48
	3,600,000 FPM or less	56	80
QRadar SIEM Event and Flow Processor Virtual 1899	200,000 FPM or less 5,000 EPS or less	16	32
	300,000 FPM or less 15,000 EPS or less	40	48
	1,200,000 FPM or less 30,000 EPS or less	56	80
QRadar SIEM All-in-One Virtual 3199	25,000 FPM or less 500 EPS or less	4	24
	50,000 FPM or less 1,000 EPS or less	8	24
	100,000 FPM or less 1,000 EPS or less	12	24
	200,000 FPM or less 5,000 EPS or less	16	32
	300,000 FPM or less 15,000 EPS or less	40	48
	1,200,000 FPM or less 30,000 EPS or less	56	80
QRadar Log Manager Virtual 8099	2,500 EPS or less	4	16
	5,000 EPS or less	8	16

Table 2. CPU requirements for QRadar virtual appliances (continued)

QRadar appliance	Threshold	Minimum number of CPU cores	Suggested number of CPU cores
QRadar Vulnerability Manager Processor Important: The IBM QRadar Vulnerability Manager scanner is end of life (EOL) in 7.5.0 Update Package 6, and is no longer supported in any version of IBM QRadar. For more information, see QRadar Vulnerability Manager: End of service product notification (https://www.ibm.com/support/pages/node/6853425) .		4	4
QRadar Vulnerability Manager Scanner Important: The IBM QRadar Vulnerability Manager scanner is end of life (EOL) in 7.5.0 Update Package 6, and is no longer supported in any version of IBM QRadar. For more information, see QRadar Vulnerability Manager: End of service product notification (https://www.ibm.com/support/pages/node/6853425) .		4	4
QRadar Risk Manager		8	8
QRadar Data Node Virtual 1400 appliance		4	16
QRadar App Host		4	12 or more for a medium-sized App Host 24 or more for a large-sized App Host

Storage requirements

Your virtual appliance must have at least 256 GB of storage available.

The following table shows the storage requirements for installing QRadar by using the virtual or software only option.

Table 3. Minimum storage requirements for appliances when you use the virtual or software installation option.

System classification	Appliance information	IOPS	Data transfer rate (MB/s)
Minimum performance	Supports XX05 licensing	800	500
Medium performance	Supports XX29 licensing	1200	1000
High Performance	Supports XX48 licensing	10,000	2000
Small All-in-One (Console) or 1600	Less than 500 EPS	300	300
Event/Flow Collectors	Events and flows	300	300

IP addressing and subnets

To configure high-availability (HA), you must consider the subnet that is used by the secondary HA host and the virtual IP address.

Administrators must ensure that the following conditions are met:

- The secondary host is in the same subnet as the primary host.
- When the IP address of the primary host is reassigned as a cluster virtual IP, the new IP address that you assign must be in the same subnet.
- The secondary HA host that you want to add to the HA cluster is not a component in another HA cluster.

Related concepts

High-availability clusters

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

Link bandwidth and latency

To configure high-availability (HA), you must consider the bandwidth and latency between the primary and secondary HA hosts.

If your HA cluster is using disk synchronization, the following conditions must be met:

- The connection between the primary and secondary HA host has a minimum bandwidth of 1 gigabits per second (Gbps).
- The latency between the primary and secondary HA host is less than 2 milliseconds (ms).

Note: If your HA solution uses a wide area network (WAN) to geographically distribute the hosts in your cluster, latency increases with distance. If latency rises above 2 ms, then system performance is affected.

Related concepts

Data consistency for HA

When an HA failover occurs, IBM QRadar ensures the consistency of your data.

Data backup requirements

There are items to consider for data backup before you configure hosts for High-availability (HA).

If a backup archive originates on an HA cluster, click **Deploy Full Configuration** to restore the HA cluster configuration after the restore is complete. If disk replication is enabled, the secondary HA host immediately synchronizes data after the system is restored.

If the secondary HA host is removed from the deployment after a backup is completed, the secondary HA host displays a **Failed** status on the **System and License Management** window.

For more information about restoring backup archives in an HA environment, see the *IBM QRadar Administration Guide*

Offboard storage requirements for HA

You can implement high-availability (HA) when the IBM QRadar `/store` partition is mounted to an external storage solution, such as an iSCSI or Fibre Channel device.

If you implement an external storage solution, the data that is received by the primary HA host is automatically moved to the external device. It remains accessible for searching and reporting.

If a failover occurs, the `/store` partition on the secondary HA host is automatically mounted to the external device. On the external device, it continues to read and write to the data received by the primary HA host before the failover.

For more information about configuring shared external storage with HA, see the IBM QRadar *Offboard Storage Guide*

Administrators must review the following HA requirements before you implement an offboard storage device:

- The primary HA host must be configured to communicate with the external device. The data in the `/store` partition of the local disk must be moved to the external storage device.
- The secondary HA host must be configured to communicate with the external device. In doing so, when a primary HA host fails over, the secondary HA host can detect the external storage device.
- You must create an HA cluster only after the secondary HA host is configured to access the same external storage device.
- If you must reconfigure your external storage device or HA cluster settings, you must remove the HA cluster between the primary and secondary HA host. For more information, see *Disconnecting an HA cluster*.
- Ensure that there is at least a 1 Gbps connection between each HA host and your external device.

Chapter 3. HA management

If you need to tune, troubleshoot, or update your high-availability (HA) settings, use the **System and License Management** window in the IBM QRadar SIEM **Admin** settings.

Administrators can use the **System and License management** window to complete the following HA tasks:

- Monitor the state of an HA cluster.
- Force the manual failover of a primary HA host to complete maintenance on the primary host.
- Disconnect an HA cluster to alter the partitions of the primary and secondary HA hosts.
- Configure the ping test time period after which automatic failover to a secondary HA host occurs.
- Modify the HA cluster settings that are used to control network connectivity testing.

Status of HA hosts

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

The following table describes the status of each host that is displayed in the **System and License Management** window:

Status	Description
Active	Specifies that the host is the active system and that all services are running normally. The primary or secondary HA host can display the active status. Note: If the secondary HA host displays the active status, the primary HA host failed.
Standby	Specifies that the host is acting as the standby system. In the standby state, no services are running but data is synchronized if disk replication is enabled. If the primary or secondary HA host fails, the standby system automatically becomes the active system.
Failed	Specifies that the primary or secondary host failed. If the primary HA host displays Failed, the secondary HA host assumes the responsibilities of the primary HA host and displays the Active status. If the secondary HA host displays Failed, the primary HA host remains active, but is not protected by HA. A system in a failed state must be manually repaired or replaced, and then restored. If the network fails, you might need access to the physical appliance.
Synchronizing	Specifies that data is synchronizing between hosts. Note: This status is displayed only when disk replication is enabled.
Online	Specifies that the host is online.

Table 4. HA status descriptions (continued)

Status	Description
Offline	<p>Specifies that an administrator manually set the HA host offline. Offline mode indicates a state that is typically used to complete appliance maintenance.</p> <p>When an appliance indicates a status of offline:</p> <ul style="list-style-type: none"> Data replication is functioning between the active and offline HA hosts. Services that process events, flows, offenses, and heartbeat ping tests are stopped for the offline HA host. Failover cannot occur until the administrator sets the HA host online.
Restoring	<p>Specifies that the host is restoring. For more information, see “Verifying the status of primary and secondary hosts” on page 28.</p>
Needs License	<p>Specifies that a license key is required for the HA cluster. In this state, no processes are running.</p> <p>For more information about applying a license key, see your <i>Administration Guide</i>.</p>
Setting Offline	<p>Specifies that an administrator is changing the status of an HA host to offline.</p>
Setting Online	<p>Specifies that an administrator is changing the status of an HA host to online</p>
Needs Upgrade	<p>Specifies that the secondary HA host requires a software upgrade.</p> <p>When the Needs Upgrade status is displayed, the primary remains active, but is not protected against failover. Disk replication of events and flows continues between the primary and the secondary HA hosts.</p>
Upgrading	<p>Specifies that the secondary HA host is being upgraded by the primary HA host.</p> <p>If the secondary HA host displays the Upgrading status, the primary HA host remains active, but is not protected by HA. Heartbeat monitoring and disk replication, if enabled, continue to function.</p> <p>After DSMs or protocols are installed and deployed on a Console, the Console replicates the DSM and protocol updates to its managed hosts. When primary and secondary HA hosts are synchronized, the DSM and protocols updates are installed on the secondary HA host.</p> <p>Only a secondary HA host can display an Upgrading status.</p>

Related concepts

[Real-time data synchronization](#)

When you configure an HA cluster, the /store file system on the primary HA host is automatically synchronized with the /store partition on the secondary HA host.

[High-availability clusters](#)

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

[Primary disk failure](#)

If RAID completely fails and all disks are unavailable, the primary HA host completes a shutdown and fails over to the secondary HA host.

Data consistency for HA

When an HA failover occurs, IBM QRadar ensures the consistency of your data.

Related tasks

Verifying the status of primary and secondary hosts

Viewing HA cluster IP addresses

You can display the IP addresses of all the components in your High-availability (HA) cluster.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. Identify the QRadar primary console.
5. Hover your mouse over the **host name** field.

Creating an HA cluster

Pairing a primary host, secondary high-availability (HA) host, and a virtual IP address creates an HA cluster.

Before you begin

- If external storage is configured for a primary HA host, you must also configure the secondary HA host to use the same external storage options. For more information, see the *QRadar Offboard Storage Guide*.
- Ensure that no undeployed changes exist before you create an HA cluster.

About this task

If disk synchronization is enabled, it might take 24 hours or more for the data in the `/store` partition on the primary HA host `/store` partition to initially synchronize with the secondary HA host.

If the primary HA host fails and the secondary HA host becomes active, the Cluster Virtual IP address is assigned to the secondary HA host.

In an HA deployment, the interfaces on both the primary and secondary HA hosts can become saturated. If performance is impacted, you can use a second pair of interfaces on the primary and secondary HA hosts to manage HA and data replication. Use a crossover cable to connect the interfaces.

Important: You can enable a crossover connection during and after the creation of an HA cluster and this does not cause any event collection downtime.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. Click **System and License Management**.
3. Select the host for which you want to configure HA.
4. From the **Actions** menu, select **Add HA Host** and click **OK**.
5. Read the introductory text. Click **Next**.
6. Type values for the parameters:

Option	Description
Primary Host IP address	<p>A new primary HA host IP address. The new IP address replaces the previous IP address. The current IP address of the primary HA host becomes the Cluster Virtual IP address.</p> <p>The new primary HA host IP address must be on the same subnet as the virtual host IP address.</p> <p>Important: For a NAT deployment, use the private IP address of the primary HA host.</p>
Secondary HA host IP address	<p>The IP address of the secondary HA host. The secondary HA host must be on the same subnet as the primary HA host.</p> <p>Important: For a NAT deployment, use the private IP address of the secondary HA host.</p>
Enter the root password of the host	The root password for the secondary HA host. The password must not include special characters.
Confirm the root password of the host	The root password for the secondary HA host again for confirmation.

7. To configure advanced parameters, click the arrow beside **Show Advanced Options** and type values for the parameters.

Option	Description
Heartbeat Interval (seconds)	<p>The time, in seconds, that you want to elapse between heartbeat pings. The default is 10 seconds.</p> <p>For more information about heartbeat pings, see “Heartbeat ping tests” on page 5.</p>
Heartbeat Timeout (seconds)	The time, in seconds, that you want to elapse before the primary HA host is considered unavailable if no heartbeat is detected. The default is 30 seconds.
Network Connectivity Test List peer IP addresses (comma delimited)	<p>The IP addresses of the hosts that you want the secondary HA host to ping. The default is to ping all other managed hosts in the QRadar deployment.</p> <p>For more information about network connectivity testing, see “Network connectivity tests” on page 4.</p>
Disk Synchronization Rate (MB/s)	<p>The disk synchronization rate. The default is 100 MB/s.</p> <p>Increase this value to 1100 MB/s when you are using 10 G crossover cables.</p> <p>Note: Do not exceed your system's capacity. The limit for Distributed Replicated Block Devices is 4096 MB/s.</p>
Disable Disk Replication	This option is displayed only when you are configuring an HA cluster by using a managed host.
Configure Crossover Cable	<p>Crossover cables allow QRadar to isolate the replication traffic from all other QRadar traffic, such as events, flows, and queries.</p> <p>You can use crossover cables for connections between 10 Gbps ports, but not the management interface.</p>

Option	Description
Crossover Interface	Select the interfaces that you want to connect to the primary HA host. Important: All interfaces with an established link, or an undetermined link, appear in the list. Select interfaces with established links only.
Crossover Advanced Options	Select Show Crossover Advanced Options to enter, edit, or view the property values.

8. Click **Next**, and then click **Finish**.

Important: When an HA cluster is configured, you can display the IP addresses that are used in the HA cluster. Hover your mouse over the **Host Name** field on the **System and License Management** window.

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

10. Click **Advanced > Deploy Full Configuration** to enable network connectivity tests.

Related concepts

High-availability clusters

A high-availability (HA) cluster consists of a primary HA host, a secondary HA host, and cluster virtual IP address. For information in installing an HA appliance, see [Installing a QRadar appliance](#).

Network connectivity tests

To test network connectivity, the IBM QRadar Console automatically pings all existing managed hosts in your QRadar deployment.

Heartbeat ping tests

You can test the operation of the primary high-availability (HA) host by configuring the time interval of heartbeat ping tests.

Disconnecting an HA cluster

By disconnecting an HA cluster, the data on your primary HA console or managed host is not protected against network or hardware failure.

Before you begin

If you migrated the `/store` file system to a Fibre Channel device, you must modify the `/etc/fstab` file before you disconnect the HA cluster. For more information, see [“Updating the /etc/fstab file”](#) on page 20.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. Select the HA host that you want to remove.
5. From the toolbar, select **High Availability > Remove HA Host**.
6. Click **OK**.

Note: When you remove an HA host from a cluster, the host restarts.

Updating the /etc/fstab file

Before you disconnect a Fibre Channel HA cluster, you must modify the /store and /storetmp mount information in the /etc/fstab file.

About this task

You must update the /etc/fstab file on the primary HA host and the secondary HA host.

Procedure

1. Use SSH to log in to your QRadar host as the root user:
2. Modify the etc/fstab file.
 - a) Locate the existing mount information for the /store and /storetmp file systems.
 - b) Remove the **noauto** option for the /store and /storetmp file systems.
3. Save and close the file.

What to do next

[Disconnecting an HA cluster.](#)

Editing an HA cluster

You can edit the advanced options for your HA cluster.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. Select the row for the HA cluster that you want to edit.
5. From the toolbar, select **High Availability > Edit HA Host**.
6. Edit the parameters in the table in the advanced options section.
7. Click **Next**.
8. Review the information.
9. Click **Finish**.

Setting an HA host offline

You can set the primary or secondary high-availability (HA) host to **Offline** from the **Active** or **Standby** state.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. Select the HA host that you want to set to offline.
5. From the toolbar, select **High Availability > Set System Offline**.

Related concepts

[Manual failovers](#)

You can manually force a failover from a primary high-availability (HA) host to a secondary HA host.

Setting an HA host online

You can set the primary or secondary HA host to Online.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. Select the offline HA host that you want to set to Online.
5. From the toolbar, select **High Availability > Set System Online**.

What to do next

On the **System and License Management** window, verify the status of the HA host. Choose from one of the following options:

- If the primary HA host displays a status of **Active**, HA host is restored.
- If you experience a problem, restore the primary or secondary HA host. For more information, see *Restoring a failed secondary HA host* or *Restoring a failed primary HA host*.

Related concepts

[Post-failover data synchronization](#)

Data that is collected by a primary high-availability (HA) host, up to the point of failover, is maintained virtually, in real time, by the secondary HA host.

Switching a primary HA host to active

You can set the primary high-availability (HA) host to be the active system.

Before you begin

The primary HA host must be the standby system and the secondary HA host must be the active system.

About this task

If your primary host is recovered from a failure, it is automatically assigned as the standby system in your HA cluster. You must manually switch the primary HA host to be the active system and the secondary HA host to be the standby system.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click the **System and License Management** icon.
4. In the **System and License Management** window, select the **secondary HA host**.
5. From the toolbar, select **High Availability > Set System Offline**.

Your primary HA host is automatically switched to be the Active HA host.

Note: Your IBM QRadar SIEM user interface might be inaccessible during this time. To return the secondary HA host to standby and return the primary HA host to active run either of the following commands:

- From the active host run:

```
/opt/qradar/ha/bin/ha giveback
```

The host that was active is now the standby host and the standby host is now the active host.

- From the standby host run:

```
/opt/qradar/ha/bin/ha takeover
```

The host that was on standby is now the active host and the host that was the active host is now the standby host.

6. In the **System and License Management** window, select the **secondary HA host**.
7. From the toolbar, select **High Availability > Set System Online**.
Your secondary HA host is now the standby system.

What to do next

When you can access the **System and License Management** window, check the **status** column. Ensure that the primary HA host is the active system and the secondary HA host is the standby system.

Chapter 4. Recovery options for HA appliances

You can reinstall or recover IBM QRadar high-availability (HA) appliances.

If your HA cluster uses shared storage, manually configure your external storage device. For more information, see the IBM QRadar *Offboard Storage Guide*.

Recovering a failed primary HA host

You can recover a failed primary high-availability (HA) IBM QRadar host if the build version of the primary HA host is the same as the QRadar build version installed on the secondary HA host.

Before you begin

Ensure that the following requirements are met:

- The required hardware is installed.
- You need the cluster virtual IP address and the primary HA host IP address. You can identify the IP address in the **System and License Management** window. For more information, see [“Viewing HA cluster IP addresses” on page 17](#).
- A keyboard and monitor are connected by using the VGA connection.

Procedure

1. Type `root` at the login prompt to start the installation wizard.
2. Accept the End User License Agreement.
3. Select the appliance type:
 - Appliance Install
 - Software Install

Important: You must choose the same appliance type as the failed primary. Do not choose an HA standby appliance.
4. In the **Type of Setup** window, select **HA Recovery Setup**.
5. Follow the instructions in the wizard.
6. Configure the QRadar network settings.
 - a) In the **Cluster Virtual IP Address Setup** window, enter the cluster virtual IP address.
 - b) In the **Network Information Setup** window, enter the original hostname and the IP address of the primary HA host.

Note: When an HA cluster is created, "-primary" is appended to the original hostname of the primary HA host. Do not include "-primary" when you enter the original hostname in the **Network Information Setup** window.
7. Configure the QRadar root password.
8. Review your software version. If your secondary HA host patch version is newer than the software on this appliance, download and install the SFS (software fix) from [Fix Central](http://www.ibm.com/support/fixcentral/) (www.ibm.com/support/fixcentral/) to upgrade this appliance to match the software version.
9. Log in to the QRadar user interface.
10. Select **Main menu > Admin > System and License Management > Systems**.
11. Highlight the primary HA host that you are restoring and select **High Availability > Restore System**.

Recovering a failed secondary HA host

You can recover a failed secondary high-availability (HA) IBM QRadar host if the build version of the secondary HA host must be the same as the QRadar build version installed on the primary HA host.

Before you begin

Ensure that the following requirements are met:

- The required hardware is installed.
- You need the secondary HA host IP address. You can identify the IP address in the **System and License Management** window.
- A keyboard and monitor are connected by using the VGA connection.

Procedure

1. Type `root` at the login prompt to launch the installation wizard.
2. Accept the End User License Agreement.
3. Select the appliance type: **High Availability Appliance**.
4. Follow the instructions in the wizard.
5. Configure the QRadar root password.
6. Review your software version. If your primary HA host patch version is newer than the software on this appliance, download and install the SFS (software fix/patch) from [Fix Central](http://www.ibm.com/support/fixcentral/) (www.ibm.com/support/fixcentral/) to upgrade this appliance to match the software version.
7. Log in to the QRadar user interface.
8. Select **Main menu > Admin > System and License Management > Systems**.
9. Highlight the secondary HA host that you are restoring and select **High Availability > Restore System**.

Restoring a primary HA host to a previous version or factory default

Restore an IBM QRadar primary high-availability (HA) host to a previous version or factory default. You can restore a failed QRadar primary HA host that does not include a recovery partition or a USB port to a previous version. You can also restore the system to factory defaults. When you restore the failed primary HA host, all data is removed and the factory default configuration is restored on the host.

Procedure

1. Use SSH to log in to the console as the root user.
2. Copy the `recovery.py` script from the console to the failed primary HA host.

```
scp recovery.py root@<TargetIP_address>:/root
```
3. Obtain the QRadar ISO from the following location: <https://www.ibm.com./support>
4. Copy the ISO file to the target QRadar host.

```
scp <iso_file_name> root@<TargetIP_address>:/root
```
5. Use SSH to log in to the primary HA host.
6. Type the following commands:

```
chmod 755 recovery.py
```

```
./recovery.py -r --default --reboot <iso_file_name>
```

7. Press Enter when prompted to restart the system.
8. When prompted, type `flatten` and press Enter.

Results

The installer repartitions and reformats the hard disk, installs the operating system, and then installs QRadar. Wait for the `flatten` process to complete. This process can take up to several minutes. After the process is complete, the normal installation process continues.

What to do next

[“Recovering a failed primary HA host” on page 23](#)

Restoring a secondary HA host to a previous version or factory default

Restore an IBM QRadar secondary high-availability (HA) host to a previous version or factory default. You can restore a failed QRadar secondary HA host that does not include a recovery partition or a USB port to a previous version. You can also restore the system to factory defaults. When you restore the failed secondary HA host, all data is removed and the factory default configuration is restored on the host.

Procedure

1. Use SSH to log in to the console as the root user.
2. Copy the `recovery.py` script from the console to the failed secondary HA host.

```
scp recovery.py root@<TargetIP_address>:/root
```
3. Obtain the QRadar ISO from the following location: <https://www.ibm.com./support>
4. Copy the ISO file to the target QRadar host.

```
scp <iso_file_name> root@<TargetIP_address>:/root
```
5. Use SSH to log in to the secondary HA host.
6. Type the following commands:

```
chmod 755 recovery.py
```

```
./recovery.py -r --default --reboot <iso_file_name>
```

7. Press Enter when prompted to restart the system.
8. When prompted, type `flatten` and press Enter.

Results

The installer repartitions and reformats the hard disk, installs the operating system, and then installs QRadar. Wait for the `flatten` process to complete. This process can take up to several minutes. After the process is complete, the normal installation process continues.

What to do next

[“Recovering a failed secondary HA host” on page 24](#)

Chapter 5. Troubleshooting QRadar HA deployments

Use the status of the HA hosts in the **System and License Management** window to help you troubleshoot.

Status combinations and possible resolutions

The following table describes the possible status settings for primary and secondary HA hosts. Each status combination requires a different troubleshooting approach.

Primary HA host status	Secondary HA host status	Possible action
Active	Failed or Unknown	Ensure that the secondary host is on, and that you can log on to it as a root user by using SSH. If you can connect, see “Restoring a failed secondary HA host” on page 27.
Failed or Unknown	Active	Ensure that the primary host is on, and that you can log on to it as a root user by using SSH. If you can connect, see “Restoring a failed primary HA host” on page 28.
Unknown	Unknown	If you cannot connect to the primary or secondary HA host by using SSH, ensure that your network and hardware configuration is operational.
Offline	Active	To set the primary host online, see Set the primary HA host online .

Identifying active hosts

You can identify the most recent active host in your HA cluster by using SSH.

1. To display the HA cluster configuration, type the following command:

```
/opt/qradar/ha/bin/ha cstate
```

2. Review the following line: in the output:

```
Local: R:PRIMARY S:ACTIVE/ONLINE CS:NONE P:1:0 HBT:UP RTT:2 1:0 SI:4105589  
Remote: R:SECONDARY S:STANDBY/ONLINE CS:NONE P:1:0 HBC:UP RTT:2 I:11753 SI:1382557
```

- If the output displays `Local: R:PRIMARY S:ACTIVE/ONLINE`, the primary HA host is the active system.
- If the output displays the following text, `Remote: R:SECONDARY S:ACTIVE/ONLINE`, the secondary HA host is the active system.

Restoring a failed secondary HA host

You can restore a failed secondary HA host.

Important: Restore only a failed secondary host, or a secondary host with unknown status. If you reinstall the HA secondary host, the state changes to standby.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click **System and License Management**.
4. Select the secondary HA host that you want to restore.
5. From the **High Availability** menu, click **Restore System**.
6. If the secondary HA host displays a status of **Failed** or **Unknown** in the **System and License Management** window, use SSH to log in to the secondary HA host as the root user to ensure that the host is operational.
7. Restart the secondary HA host by typing `reboot`.
8. After the system is restarted, if the secondary HA host displays a status of **Failed** or **Unknown**, from the **High Availability** menu, click **Restore System**.

Related tasks

[Verifying the status of primary and secondary hosts](#)

Restoring a failed primary HA host

You can restore a failed primary HA host.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click **System and License Management**.
4. Select the primary HA host that you want to restore.
5. From the **High Availability** menu, click **Restore System**.
6. Verify the status of the primary HA host.
7. If the primary HA host displays a status of **Offline**, in the **System and License Management** window, click **High Availability** > **Set System Online**.
8. If the primary HA host displays a status of **Failed** or **Unknown** in the **System and License Management** window, use SSH to log in to the primary HA host as the root user to ensure that the host is operational.
9. Restart the primary HA host by typing the following command: `reboot`

Related tasks

[Setting the status of the primary HA host to online](#)

Verifying the status of primary and secondary hosts

You must verify that the primary and secondary HA hosts are operational.

Procedure

1. Identify whether the primary HA host was configured as a console or managed host.
2. If the primary HA host is configured as a console, use SSH to log in to the Cluster Virtual IP address as the root user:
 - If you can connect to the Cluster Virtual IP address, restore access to the QRadar. For more information, see the *IBM Security QRadar SIEM Troubleshooting Guide*.
 - If you cannot connect to the Cluster Virtual IP address, use SSH to log in to the secondary HA host as the root user to ensure that it is operational.

3. If your secondary host is configured as a managed host, use SSH to log in to the secondary HA host as the root user.
 - If you cannot connect to the primary or secondary HA host by using SSH, ensure that your network and hardware configuration is operational.
 - If you can connect to the primary and secondary HA host, identify the most recently active HA host in your HA cluster.

Related concepts

Status of HA hosts

You can review the status of the primary and secondary host in your high-availability (HA) cluster.

Related tasks

Verifying the status of primary and secondary hosts

Setting the status of the primary HA host to online

If the primary HA host displays a status of offline, you can reset the status to online.

Procedure

1. On the navigation menu (☰), click **Admin**.
2. On the navigation menu, click **System Configuration**.
3. Click **System and License Management**.
4. Select the primary HA host that you want to restore.
5. In the **System and License Management** window, if the primary HA host displays a status of **Offline**, you must restore the primary HA host.

Related tasks

Restoring a failed primary HA host

Chapter 6. Recovery solution for QRadar deployments

Maintaining data redundancy is crucial to resiliency and recovery from data loss. There are a wide variety of solutions that are currently deployed in the field to prevent and recover from data loss, and vary greatly in terms of complexity, cost, and effectiveness. IBM QRadar provides the IBM QRadar Data Synchronization app as a solution to maintain your configuration and data during a failure of your main site.

IBM QRadar Data Synchronization app

The QRadar Data Synchronization app mirrors your data to another identical system. It is possible to maintain configurations and data when you have two identical QRadar systems in separate geographic environments that are a mirror of each other. Data is collected at both sites and ensures operations can continue to function as normally as possible in scenarios when your main site fails.

QRadar Data Synchronization forwards live data, for example, flows and events from the main site's QRadar system to a parallel destination site. You can set up data synchronization with deployments that are in different geographical locations.

To use the QRadar Data Synchronization app, the main site and destination site deployments must be running QRadar 7.4.0 FixPack 3 or later. The destination site must be a fully duplicated deployment (1:1 host ratio) for hosts that contain or collect Ariel (event and flow) data. This includes Event Processors, Flow Processors, All in one Event Processors and Flow Processors, Event Collectors, Flow Collectors, consoles, and data nodes. However, QRadar Risk Manager, QRadar Vulnerability Manager, QRadar Incident Forensics, QRadar Network Insights, and QRadar App Host do not require 1:1 mapping.

A high-availability (HA) cluster is considered one host and the Data Synchronization app supports a HA cluster that is paired with a non-HA host.

Note: App data backup is currently not available using the Data Synchronization app. For more information about app data backup and recovery, see [Backing up and restoring app data](#).

To learn more about the QRadar Data Synchronization app, see [Data Synchronization app](https://www.ibm.com/support/pages/node/6574833) [<https://www.ibm.com/support/pages/node/6574833>].

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.

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”.

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>

