

IBM QRadar
7.4.3

Offboard Storage Guide



Note

Before you use this information and the product that it supports, read the information in [“Notices” on page 35](#).

Contents

Introduction to offboard storage devices for QRadar products.....	v
Chapter 1. Overview.....	1
Appliance storage requirements for virtual and software installations.....	2
File system options.....	2
Performance impact.....	3
Storage expansion.....	3
External storage options.....	3
External storage limitations.....	4
Offboard storage in HA environments.....	5
Chapter 2. iSCSI external storage device.....	7
Configuring the iSCSI volumes.....	7
Moving the /store file system to an iSCSI storage solution.....	9
Moving the /store/ariel file system to an iSCSI storage solution.....	11
Mounting the iSCSI volume automatically.....	12
Configuring iSCSI in an HA deployment.....	12
Verifying iSCSI connections.....	14
Troubleshooting iSCSI issues.....	14
Secondary network interfaces.....	16
Configuring control of secondary interfaces in HA deployments.....	16
Chapter 3. Fibre Channel storage.....	19
Verifying your Emulex adapter installation.....	20
Verifying the Fibre Channel connections.....	20
Moving the /store file system to a Fibre Channel solution.....	22
Moving the /store/ariel file system to a Fibre Channel solution.....	23
Moving the /store file system to a multipath Fibre Channel solution.....	25
Moving the /store file system to a multipath Fibre Channel solution in an HA deployment.....	26
Configuring the mount point for the secondary HA host.....	26
Removing HA from a Fibre Channel solution.....	27
Chapter 4. NFS offboard storage device.....	29
Moving backups to an NFS.....	29
Configuring a mount point for a secondary HA host.....	31
Configuring NFS backup on an existing HA cluster.....	32
Notices.....	35
Trademarks.....	36
Terms and conditions for product documentation.....	36
IBM Online Privacy Statement.....	37
General Data Protection Regulation.....	37

Introduction to offboard storage devices for QRadar products

This guide provides information about how to move the /store, /store/backup, or /store/ariel file systems to an external storage device for IBM® QRadar® products.

Intended audience

System administrators responsible for configuring offboard storage devices must have administrative access to QRadar systems and to network devices and firewalls. The system administrator must know the corporate network 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?uid=swg21616144) (<http://www.ibm.com/support/docview.wss?uid=swg21616144>).

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. Offboard storage overview

To increase the amount of storage space available to your appliance, you can move a portion 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, Fibre Channel, and NFS (Network File System). You must use iSCSI or Fibre Channel to store data that is accessible and searchable, such as the `/store/ariel` directory.



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

You can use offboard storage solutions on any managed host or console, including high-availability (HA) systems. When you use iSCSI or Fibre Channel with HA, the external storage device is mounted by the active HA node, ensuring data consistency for an HA failure. When you use external storage with HA, you must configure these devices on the primary and secondary HA hosts.

Before you implement an offboard storage solution, consider your local storage options, existing hardware infrastructure, and your data retention and fault tolerance requirements.

Offboard storage data encryption

Event data in QRadar® is not encrypted when stored. However, `/store` or `/store/ariel` partitions can be placed on an external device, which uses transparent (to QRadar) cryptography. For more information on external storage options, see [“External storage options”](#) on page 3.

Important: To set up encryption on the storage, see the documentation for your storage solution.

Local storage

Data that is stored locally on a QRadar appliance can be accessed with lower latency than on external storage. When possible, use local storage and Data Node appliances as an alternative to an external storage device.

Multiple appliances

Use multiple appliances if larger storage capacity is required for your QRadar deployment.

When multiple appliances are not feasible, or when an existing deployment can increase capacity by using available external storage, then external storage might be appropriate for your deployment.

Hardware and infrastructure

Your existing infrastructure and experience with storage area networks are important factors in deciding whether to use an offboard storage solution.

Certain offboard devices require less configuration and might be able to use existing network infrastructures. For example, iSCSI uses existing Ethernet networking, while Fibre Channel uses specialized hardware.

Data retention and fault tolerance

Your QRadar data retention policy is important in considering an offboard storage solution. If your data retention settings exceed the capacity of existing storage or if you are planning to expand the retention of existing deployed appliances, you might require an offboard storage solution.

An offboard storage solution can be used to improve your fault tolerance and disaster recovery capabilities.

Appliance storage requirements for virtual and software installations

To install QRadar using virtual or software options the device must meet minimum storage requirements.

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

Note: The minimum required storage size will vary, based on factors such as event size, events per second (EPS), and retention requirements.

Table 1. 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 XX28 and XX29 licensing	1200	1000
High Performance	Supports XX48 licensing	10,000	2000
Small All-in-One or 1600	Less than 500 EPS	300	300
Event/Flow Collectors	Events and flows	300	300

File system options for offboard storage

Use an offboard storage solution to move the `/store` file system or specific subdirectories, such as the `/store/ariel` directory.

You can move the `/store` file system to increase the fault tolerance levels in your IBM QRadar deployment. Each option impacts QRadar performance.

The `/store/ariel` directory is the most common file system that is moved to an offboard storage solution. By moving the `/store/ariel` file system, you can move collected log and network activity data to external storage. The local disk remains used for the PostgreSQL database and for temporary search results.

Administrators can move the following types of QRadar data to offboard storage devices:

- PostgreSQL metadata and configuration information
- Log activity, payloads (raw data), normalized data, and indexes
- Network activity, payloads, normalized data, and indexes
- Time series graphs (global views and aggregates)

Note: Do not move `/transient` or `/storetmp` to an offboard storage device. Moving these directories causes your system to stop functioning properly.

Performance impact of offboard storage solutions

Moving the `/store` file system to an external device might affect QRadar performance.

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

- Maintain your log and network activity searches on your local disk by mounting the `/store/transient` file system to the unused `/store` file partition.
- Searches that are marked as saved are also in the `/store/transient` directory. If you experience a local disk failure, these searches are not saved.

Storage expansion

Storage expansion

By creating multiple external volumes and mounting `/store/ariel/events` and `/store/ariel/flows`, you can expand your storage capabilities past the single file system that is configured by default with IBM QRadar. A single file system supports up to 500 TB.

Store partition

Any subdirectory in the `/store` file system can be used as a mount point for your external storage device. However, only the `/store` and `/store/ariel` file systems are supported for offboard with a high-availability deployment.

If you want to move dedicated event or flow data, you might configure more specific mount points. For example, you can configure `/store/ariel/events/records` and `/store/ariel/events/payloads` as mount points.

More storage expansion options

You can add more data storage to QRadar host or optimize your current storage by using one or more of [these options](https://www.ibm.com/support/pages/qradar-reaching-data-storage-limits) (<https://www.ibm.com/support/pages/qradar-reaching-data-storage-limits>):

- Install a Data Node. Data Nodes enable new and existing QRadar deployments to add storage and processing capacity on demand as required. For more information, see the *IBM QRadar Architecture and Deployment Guide*.
- [Configure your Network File System \(NFS\) storage](#). You can configure NFS for a stand-alone QRadar Console, new QRadar HA deployments, or existing QRadar HA deployments.
- Configure your retention policies to define how long QRadar is required to keep event and flow data, and what to do when that data reaches a certain age. For more information, see the *IBM QRadar Administration Guide*.
- Enable [event coalescing](https://www.ibm.com/support/pages/qradar-how-does-coalescing-work-qradar) (<https://www.ibm.com/support/pages/qradar-how-does-coalescing-work-qradar>) to improve performance, and reduce storage impacts, when a large burst of events is received that match a specific criteria.

External storage options

You can use iSCSI, Fibre Channel, or Network File System (NFS) to provide an offboard storage solution.

Onboard disks provide a low latency and high throughput storage solution, which is tested and validated with various workloads. When multiple appliances are deployed, performance and capacity scale at the same rate.

Fibre Channel

Fibre Channel provides the fastest offboard performance by using storage area network (SAN) speeds of 200 MBps to 3200 MBps, depending on your network configuration.

Fibre Channel performance might be impacted by factors within the SAN implementation, such as the following factors:

- Disk or spindle counts per volume
- Number of concurrent sessions
- Cache capacity in the SAN controllers

iSCSI

iSCSI uses a dedicated storage channel over standard Ethernet infrastructure, rather than a dedicated SAN network. For this reason, iSCSI can be the easiest to implement, most cost effective, and most readily available.

If you implement an iSCSI solution, then network capacity is shared between external storage access and management interface I/O. In this situation, you can configure a secondary network interface on a separate storage network.

QRadar supports 1 Gbit and 10 Gbit connectivity out of the box on many appliances.

NFS



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

Use NFS for tasks during off-peak times, tasks that involve batch file writes, and tasks that involve a limited volume of file I/O. For example, use NFS for daily configuration and data backups.

NFS storage operates over existing management Ethernet networks and is limited by networking performance. It is possible to use a dedicated network interface to improve networking performance when compared to sharing a management interface. The NFS protocol might affect performance for file access, locking, and network permissions.

If NFS is used only for backups, the same NFS share can be used for each host. The backup files contain the system host name, which enables the identification of each backup file. If you are storing a long period of data on your NFS shares, consider a separate share or export for each appliance in your deployment.

External storage limitations

Multiple systems cannot access the same block device in an IBM QRadar deployment.

If you configure iSCSI in an HA environment, do not mount the iSCSI or Fibre Channel volumes on the secondary host while the primary host is accessing the volumes.

An external storage device must be able to provide consistent read and write capacity of 100 MBps to 200 MBps. When consistent read and write capacity is not available, the following issues might occur:

- Data write performance is impacted.
- Search performance is impacted.

If performance continues to degrade, then the processing pipeline can become blocked and QRadar might display warning messages and drop events and flows.

Offboard storage in HA environments

If you choose to move the `/store` file system in a high-availability (HA) environment, the `/store` file system is not replicated by using Disk Replication Block Device.

If you move the `/store/ariel` file system to an offboard storage device and maintain the `/store` file system on local disk, the `/store` file system is synchronized with the secondary HA host. By default, when your environment is configured for HA, Disk Replication Block Device is enabled.

If you are using HA and you move a file system other than `/store` or `/store/ariel` to a multipath offboard storage solution, you must stop **multipathd** before upgrading QRadar.

Chapter 2. iSCSI external storage device

You can configure an iSCSI storage device in a standard or high-availability (HA) IBM QRadar deployment.

When you configure an iSCSI external storage device, you must migrate the QRadar data that is maintained on your `/store` or `/store/ariel` file system and then mount the `/store` or `/store/ariel` file system to a partition on the iSCSI device volume.

Depending on your device configuration, you might be required to create a partition on the volume of your iSCSI disk.

If you configure iSCSI in an HA deployment and your primary HA host fails, your iSCSI device can be used to maintain data consistency with your secondary HA host.

In HA environments, review the `/var/log/messages` file for errors in your iSCSI storage configuration.

iSCSI configuration in standard QRadar deployments

To move data from a QRadar Console or managed host to an iSCSI storage device:

- Follow the instructions in [Configure iSCSI volumes](#).
- Migrate the file system to an iSCSI storage device.
 - [Move the /store/ariel file system to an iSCSI storage solution](#).
 - [Move the /store file system to an iSCSI storage solution](#).
- Follow the instructions in [Mount the iSCSI volume automatically](#).
- Follow the instructions in [Verify iSCSI connections](#).

iSCSI configuration in HA deployments

To move data in an HA deployment to an iSCSI storage device:

- Follow the instructions in [Configure iSCSI volumes](#) on your primary appliance.
- Migrate the file system from your primary appliance to an iSCSI storage device.
 - [Move the /store/ariel file system to an iSCSI storage solution](#).
 - [Move the /store file system to an iSCSI storage solution](#).
- Follow the instructions in [Mount the iSCSI volume automatically](#) on your primary appliance.
- Follow the instructions in [“Configuring iSCSI in an HA deployment”](#) on page 12 on your secondary appliance.
- Follow the instructions in [Verify iSCSI connections](#).

Configuring the iSCSI volumes

You can configure iSCSI for a stand-alone QRadar Console or a QRadar Console that is the primary high-availability (HA) host in an HA deployment.

About this task

Optionally, you can create a partition on the volume of the external iSCSI storage device.

IBM QRadar V7.2.1 and later uses the XFS file system. You can create the partition on your iSCSI device with either an ext4 or XFS file system.

Disk partitions are created by using a GUID Partition Table (GPT). You can use a new device partition as the mount point for the file system, such as `/store` or `/store/ariel` that you migrate.

Important: If you created an iSCSI or Fibre Channel device partition on your external device and QRadar data is stored, then you cannot create a partition or reformat the partition on the volume.

Procedure

1. Using SSH, log in to the QRadar Console as the root user.
2. Edit the `/etc/iscsi/initiatorname.iscsi` file to include the iSCSI qualified name for your host.

```
InitiatorName=<iqn.yyyy-mm>.<reversed_domain_name>:<hostname>
```

Example: `InitiatorName=iqn.2014-11.com.qradar:pl13`

3. Open a session to the iSCSI server by typing the following command:

```
systemctl restart iscsi
```

4. To detect volumes on the iSCSI server, type the following command:

```
iscsiadm -m discovery --type sendtargets --<portal_IP_address>:[<port>]
```

The *IP address* option is the IP address of the iSCSI server. The *port* is optional. Record the initiator name.

5. To log in to the iSCSI server, type the following command:

```
iscsiadm -m node --targetname <initiator_name_from_step_4> --portal <IP_address>:[<port>]> --login
```

6. To find the iSCSI device volume name, type the following command:

```
dmesg | grep "Attached SCSI disk"
```

7. To create a partition, use the GNU parted command:

```
parted /dev/<volume>
```

8. Configure the partition label to use GPT by typing the following command:

```
mklabel gpt
```

9. If the following message is displayed, type Yes.

```
Warning: The existing disk label on /dev/<volume> will be destroyed and all data on this disk will be lost. Do you want to continue?
```

10. Create a partition on the iSCSI disk volume.

- a) To create the partition, type the following command:

```
mkpart primary 0% 100%
```

- b) Set the default units to TB by typing the following command:

```
unit TB
```

- c) Verify that the partition is created by typing the following command:

```
print
```

- d) Exit from GNU parted by typing the following command:

```
quit
```

- e) Update the kernel with the new partition data by typing the following command:

```
partprobe /dev/<volume>
```

After you update the kernel, you might be prompted to restart the appliance. If you are prompted to do so, restart the appliance.

f) To verify that the partition is created, type the following command:

```
cat /proc/partitions
```

11. Reformat the partition and make a file system.

- To create an XFS file system, type the following command:

```
mkfs.xfs -f /dev/<partition>
```

- For an ext4 files system, type the following command:

```
mkfs.ext4 /dev/<partition>
```

What to do next

See [“Moving the /store/ariel file system to an iSCSI storage solution” on page 11](#) or [“Moving the /store file system to an iSCSI storage solution” on page 9](#).

Related tasks

[Troubleshooting iSCSI issues](#)

Moving the /store file system to an iSCSI storage solution

You can migrate the IBM QRadar data that is maintained in the /store file system and mount the /store file system to an iSCSI device partition.

Migrating the /store files system to your offboard storage device can take an extended time.

Before you begin

[Configure iSCSI volumes.](#)

Procedure

1. Stop the QRadar services by typing the following commands in the order specified:

Note: Run the command `systemctl stop solr`, only if you have QRadar Incident Forensics in your deployment.

```
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop hostservices
systemctl stop systemStabMon
systemctl stop crond
systemctl stop soir
```

Note: Run the command `systemctl stop tomcat` on the Console.

2. Unmount the file systems by typing the following commands:

```
umount /store
```

3. Create the /store_old directory by typing the following command:

```
mkdir /store_old
```

4. Derive the iSCSI device partition universal unique identifier (UUID) by typing the following command:

```
blkid /dev/<partition>
```

5. Edit the /etc/fstab file to update the existing /store file system mount point to /store_old.

6. Make a new mount point for the `/store` file system by adding the following text to the `/etc/fstab` file:

- If the file system is XFS, add the following text:

```
UUID=<uuid> /store xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

- If the file system is ext4, add the following text:

```
UUID=<uuid> /store ext4 noatime,noauto,nobarrier 0 0
```

Save and close the file.

7. Mount the `/store` file system to the iSCSI device partition by typing the following command:

```
mount /store
```

8. Mount the `/store_old` file system to the local disk by typing the following command:

```
mount /store_old
```

9. Move the data from the local disk to the iSCSI storage device by typing the following command:

```
cp -af /store_old/* /store
```

10. Unmount `/store_old` by typing the following command:

```
umount /store_old
```

11. Remove the `/store_old` directory by typing the following command:

```
rmdir /store_old
```

12. Edit the `/etc/fstab` file to remove the `/store_old` entry.

13. Start the QRadar services by typing the following commands in the order specified:

Note: Run the command `systemctl start solr`, only if you have QRadar Incident Forensics in your deployment.

```
systemctl start crond
systemctl start systemStabMon
systemctl start hostservices
systemctl start tomcat
systemctl start ecs-ec-ingress
systemctl start hostcontext
systemctl start solr
```

14. Remove the local copy of `/store` from the logical volume manager (LVM) by typing the following commands:

```
lvchange -an /dev/storerhel/store 2>/dev/null
lvrename /dev/storerhel/store /dev/storerhel/storeold 2>/dev/null
```

What to do next

See [“Mounting the iSCSI volume automatically”](#) on page 12.

Related tasks

[Moving the `/store/ariel` file system to an iSCSI storage solution](#)

You can migrate the IBM QRadar data that is maintained in the `/store/ariel` file system and mount the `/store/ariel` file system to an iSCSI device partition.

Moving the `/store/ariel` file system to an iSCSI storage solution

You can migrate the IBM QRadar data that is maintained in the `/store/ariel` file system and mount the `/store/ariel` file system to an iSCSI device partition.

Before you begin

[Configure iSCSI volumes.](#)

Procedure

1. Stop the QRadar services by typing the following commands in the order specified:

```
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop hostservices
systemctl stop systemStabMon
systemctl stop crond
```

2. Move the existing mount point by typing the following commands:

```
cd /store
mv ariel ariel_old
```

3. Verify the Universally Unique Identifier (UUID) of the iSCSI device partition by typing the following command:

```
blkid /dev/partition
```

4. Add the mount point for the `/store/ariel` file system by adding the following text to the `/etc/fstab` file:

- If the file system is XFS, copy the following text into a text editor, remove the line break, and paste as a single line:

```
UUID=uuid /store/ariel xfs inode64,logbsize=256k,noatime,
noauto,nobarrier 0 0
```

- If the file system is ext4, add the following text

```
UUID=uuid /store/ariel ext4 noatime,noauto,nobarrier 0 0
```

5. Create the `ariel` directory for the mount point by typing the following command:

```
mkdir /store/ariel
```

6. Mount `/store/ariel` to the iSCSI device partition by typing the following command:

```
mount /store/ariel
```

7. Verify that `/store/ariel` is correctly mounted by typing the following command:

```
df -h
```

8. Move the data from the local disk to the iSCSI storage device by typing the following command:

```
cp -af /store/ariel_old/* /store/ariel
```

9. Remove the `/store/ariel_old` directory by typing the following command:

```
rmdir /store/ariel_old
```

10. Start the QRadar services by typing the following commands in the order specified:

```
systemctl start crond
systemctl start systemStabMon
systemctl start hostservices
```

```
systemctl start tomcat
systemctl start ecs-ec-ingress
systemctl start hostcontext
```

What to do next

See [“Mounting the iSCSI volume automatically”](#) on page 12.

Related tasks

[Moving the /store file system to an iSCSI storage solution](#)

Mounting the iSCSI volume automatically

You must configure IBM QRadar to automatically mount the iSCSI volume.

Before you begin

Ensure that you moved the /store/ariel or /store file systems to an iSCSI storage solution.

Procedure

1. Add the iSCSI script to the startup information by typing the following commands:

```
systemctl enable iscsi
```

2. Enable the iscsi-mount service by typing the following command:

```
systemctl enable iscsi-mount
```

3. Verify that the iSCSI device is correctly mounted by typing the following command:

```
df -h
```

What to do next

If you are configuring iSCSI in a standard QRadar deployment, see [“Verifying iSCSI connections”](#) on page 14.

If you are configuring a high-availability (HA) environment, you must set up your secondary HA host by using the same iSCSI connections that you used for your primary HA host. For more information, see [“Configuring iSCSI in an HA deployment”](#) on page 12.

Related tasks

[Configuring iSCSI in an HA deployment](#)

To use an iSCSI device in an HA environment, you must configure the primary and secondary HA hosts to use the same iSCSI external storage device. Only iSCSI single path is supported for the HA deployment.

Configuring iSCSI in an HA deployment

To use an iSCSI device in an HA environment, you must configure the primary and secondary HA hosts to use the same iSCSI external storage device. Only iSCSI single path is supported for the HA deployment.

About this task

Ensure that you use a different **initiatorname** on the primary and secondary HA hosts. Your iSCSI device must be configured to enable each **initiatorname** to access the same volume on the iSCSI device.

Important: Configure iSCSI for your secondary HA host before you create your HA cluster.

Procedure

1. Use SSH to log in to the secondary HA host as the root user.
2. To configure your HA secondary host to identify the iSCSI device volume, add the iSCSI qualified name for your host to the `/etc/iscsi/initiatorname.iscsi` file.

```
Initiatorname=iqn.<yyyymm>.{reversed domain name}:<hostname>
```

Important: The **initiatorname** for your secondary HA host must be different than the **initiatorname** for your primary HA host.

Example: `InitiatorName=iqn.2008-11.com.qradar:p114`

3. Restart the iSCSI service to open a session to the server by typing the following command:

```
systemctl restart iscsi
```

4. To detect the volume on the iSCSI server, type the following command:

```
iscsiadm -m discovery --type sendtargets --portal <IP_address>:[<port>]
```

Note: The *port* is optional.

5. Verify the login to your iSCSI server by typing the following command:

```
iscsiadm -m node --targetname <initiator_name_from_the_previous_step> --portal <IP_address>:[<port>] --login
```

6. To find the iSCSI device volume name, type the following command:

```
dmesg | grep "Attached SCSI disk"
```

7. Configure the mount point for the secondary HA host.

- a) If you are moving the `/store` file system, unmount the file systems by typing the following commands:

```
umount /store
```

- b) Identify the UUID of the iSCSI device partition by typing the following command:

```
blkid /dev/<partition>
```

- c) To move the `/store` file system, edit the file settings in the `/etc/fstab` file to be the same as the mount points that might be listed in the `/etc/fstab` file on the HA primary host:

- `/store`
- For the `/store` partition, use the same UUID value that is used for the `/store` partition on the primary.

- d) If you are moving the `/store/ariel` file system, edit the settings in the `/etc/fstab` file to be the same as the mount point that is listed in the `/etc/fstab` file on the HA primary host for `/store/ariel`.

8. Configure the secondary HA host to automatically mount the iSCSI volume.

- a) Add the iSCSI script to the startup information by typing the following commands:

```
systemctl enable iscsi
```

- b) Enable the `iscsi-mount` service by typing the following command:

```
systemctl enable iscsi-mount
```

- c) If you are moving the `/store` file system, rename the local copy of `/store` by typing the following commands:

```
lvchange -an /dev/storerhel/store 2>/dev/null
lvrename /dev/storerhel/store /dev/storerhel/storeold 2>/dev/null
```

What to do next

Create an HA cluster. For more information, see *IBM QRadar High Availability Guide*.

See [“Verifying iSCSI connections”](#) on page 14.

Verifying iSCSI connections

Verify that the connections between a primary HA host or secondary HA host and an iSCSI device are operational

Procedure

1. Using SSH, log in to the primary or secondary HA host as the root user.
2. To test the connection to your iSCSI storage device, type the following command:

```
ping iSCSI_Storage_IP_Address
```

3. Verify the iSCSI service is running and that the iSCSI port is available by typing the following command:

```
telnet iSCSI_Storage_IP_Address 3260
```

Note: The default port is 3260.

4. Verify that the connection to the iSCSI device is operational by typing the following command:

```
iscsiadm -m node
```

To verify that the iSCSI device is correctly configured, you must ensure that the output that is displayed for the primary HA host matches the output that is displayed for the secondary HA host.

If the connection to your iSCSI volume is not operational, the following message is displayed:

```
iscsiadm: No records found
```

5. If the connection to your iSCSI volume is not operational, then review the following troubleshooting options:
 - Verify that the external iSCSI storage device is operational.
 - Access and review the `/var/log/messages` file for specific errors with your iSCSI storage configuration.
 - Ensure that the iSCSI **initiatornames** values are correctly configured by using the `/etc/iscsi/initiatorname.iscsi` file.
 - If you cannot locate errors in the error log, and your iSCSI connections remain disabled, then contact your Network Administrator to confirm that your iSCSI server is functional or to identify network configuration changes.
 - If your network configuration has changed, you must reconfigure your iSCSI connections.

What to do next

Establish an HA cluster. You must connect your primary HA host with your secondary HA host by using the QRadar user interface. For more information about creating an HA cluster, see the *IBM QRadar High Availability Guide*.

Troubleshooting iSCSI issues

To prevent iSCSI disk and communication issues, you must connect QRadar, the iSCSI server, and your network switches to an uninterruptible power supply (UPS). Power failure in a network switch might result in your iSCSI volume reporting disk errors or remaining in a read-only state.

About this task

In a high-availability (HA) environment, if your primary host fails, you must restore your iSCSI configuration to the primary host. In this situation, the `/store` or `/store/ariel` data is already migrated to the iSCSI shared external storage device. Therefore, to restore the primary host iSCSI configuration, ensure that you configure a secondary HA host. For more information, see [“Configuring iSCSI in an HA deployment”](#) on page 12

Procedure

1. Determine whether a disk error exists.
 - a) Using SSH, log in to QRadar Console as the root user.
 - b) Create an empty file `filename.txt` on your iSCSI volume by typing one of the following command:
 - `touch /store/ariel/filename.txt`
 - `touch /store/filename.txt`

If your iSCSI volume is mounted correctly and you have write permissions to the volume, the touch command creates an empty file that is named `filename.txt` on your iSCSI volume.

If you see an error message, unmount and remount the iSCSI volume.

2. Stop the QRadar services.
 - If you migrated the `/store` file system, type the following commands in the specified order:

```
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop hostservices
systemctl stop systemStabMon
systemctl stop crond
```

- If you migrated the `/store/ariel` file system, type the following command:

```
systemctl stop hostcontext
```

3. Unmount the iSCSI volume.

- If you migrated the `/store` file system, type the following commands:

```
umount /store
```

- If you migrated the `/store/ariel` file system, type the following command:

```
umount /store/ariel
```

4. Mount the iSCSI volume.

- If you migrated the `/store` file system, type the following commands:

```
mount /store
```

- If you migrated the `/store/ariel` file system, type the following command:

```
mount /store/ariel
```

5. Test the mount points.

- If you migrated the `/store` file system, type the following command:

```
touch /store/filename.txt
```

- If you migrated the `/store/ariel` file system, type the following command:

```
mount /store/ariel/filename.txt
```

If you continue to receive a read-only error message after you remount the disk, then reconfigure your iSCSI volume.

Alternatively, you can unmount the file system again and run a manual file system check with the following command: `fsck /dev/partition`.

6. Start the QRadar services.

Related tasks

Configuring the iSCSI volumes

You can configure iSCSI for a stand-alone QRadar Console or a QRadar Console that is the primary high-availability (HA) host in an HA deployment.

Secondary network interfaces

You can configure a secondary network interface with a private IP address to connect to an iSCSI storage area network (SAN).

You use secondary network interface to improve performance. If you configure a secondary network interface, you require address information from your SAN network manager. For more information about configuring a network interface, see "Network interface management" in the *IBM QRadar Administration Guide*.

HA systems in iSCSI deployments

For dedicated access to the iSCSI storage network, use the following order to set up high availability (HA), iSCSI, and a network interface:

- __ 1. Configure the primary and secondary appliances.
- __ 2. Set up external iSCSI storage on both hosts.
- __ 3. Configure HA on the primary and secondary hosts.
- __ 4. Configure control of the secondary interfaces for your HA appliances.

The HA process for IBM QRadar controls the all network interfaces. When an HA appliance is in active mode, the HA process enables the interfaces. When HA is in standby mode, the HA process disables the interfaces. If the dedicated network interface for storage is disabled and the HA system goes into failover, the standby host tries to go into active mode. If the HA system is in standby mode, you cannot access the iSCSI storage system. Access issues are caused during the transition of the HA node from standby to active. The HA process brings the secondary interface online, but when the iSCSI system is mounted, the networking is not available and the failover process fails. The standby HA host cannot change to active mode.

To resolve the issue, you must remove control of the iSCSI network interface from the HA system to ensure that network interface is always active. Remove any dependencies that the network interface has on the status of the HA node. The HA primary and secondary hosts must have unique IP addresses on these secondary network interfaces.

Related tasks

Configuring control of secondary interfaces in HA deployments

If you use iSCSI and a dedicated network interface in a high-availability (HA) deployment, you must ensure that the secondary interface is not managed by the HA process. Configure the management of the secondary interface to ensure that if a failover to the secondary HA host occurs, the interface always remains active.

Configuring control of secondary interfaces in HA deployments

If you use iSCSI and a dedicated network interface in a high-availability (HA) deployment, you must ensure that the secondary interface is not managed by the HA process. Configure the management of

the secondary interface to ensure that if a failover to the secondary HA host occurs, the interface always remains active.

Before you begin

Ensure that the following conditions are met:

- Separate IP addresses for the dedicated iSCSI network interface on each of the HA servers. The IP addresses must be on different networks.

Separate IP addresses prevent IP address conflicts when the network interfaces are active on both HA hosts at the same time. The iSCSI software and drivers can access the external storage at startup and during the HA failover. Also, the external volume can be successfully mounted when the HA node switches from standby to active.

For more information about configuring network interfaces, see "Configuring network interfaces" in the *IBM QRadar Administration Guide*.

- The primary and secondary appliances are configured.

For more information, see the *IBM QRadar High Availability Guide*.

- iSCSI storage is configured.
- NetworkManager is disabled by typing the following command.

```
systemctl status NetworkManager
```

Procedure

1. On the primary host, use SSH to log in to the QRadar Console as the root user.
2. Disable the QRadar HA service control of network interface.
 - a) Go to the `/opt/qradar/ha/interfaces/` directory
The directory contains a list of files that have a name that starts with `ifcfg-`. One file exists for each interface that is controlled by QRadar HA processes.
 - b) Delete the file that is used to access your iSCSI storage network.
Deleting the file removes control of the interface from the HA processes.
3. Re-enable operating system-level control of the network interfaces.
 - a) Go to the `/etc/sysconfig/network-scripts` directory.
 - b) Open the `ifcfg-` file for the interface that connects to your iSCSI network.
 - c) To ensure that the network interface is always active, change the value for the `ONBOOT` parameter to `ONBOOT=yes`.
4. To restart the iSCSI services, type the following command:

```
systemctl restart iscsi
```

5. Repeat these steps for the HA secondary appliance.
6. To test access to your iSCSI storage from your secondary appliance, use the ping command:

```
ping <iscsi_server_ip_address>
```

Related concepts

Secondary network interfaces

You can configure a secondary network interface with a private IP address to connect to an iSCSI storage area network (SAN).

Chapter 3. Fibre Channel storage

You can configure Fibre Channel (FC) in a standard QRadar deployment or in a high-availability (HA) environment. You can also configure FC multipath to provide redundancy if your FC switch fails.

When you configure an FC device, you can move the IBM QRadar data in your `/store` or `/store/ariel` file system. Then, mount the `/store` or `/store/ariel` file system to a partition on the FC device.

Frequently searched data must be moved to a faster disk. For example, move recent data or data that is used for security incident investigations. However, deploying high performance offboard disk storage might be costly. Where possible, use lower performance and less expensive offboard storage for activities such as moving older data, archiving, or for reporting purposes.

If you are using FC only for archive purposes, then use the same mount point for every appliance and configure these mount points to correspond with each unique FC volume.

In deployments that use multiple appliances, ensure that each appliance is configured to use a separate FC volume. Failure to use separate volumes can result in two devices that mount the same block device, which can corrupt the block device file system.

Depending on your device configuration, you might be required to create a partition on the volume of your FC disk.

Configuring Fibre Channel (FC) is different for a primary high-availability (HA) host than the secondary HA host. To configure FC, you must ensure that the primary HA host and secondary HA host are not connected in an HA cluster.

Fibre Channel configuration in standard QRadar deployments

To move data from a QRadar Console or managed host to an FC storage device:

- Follow the instructions in [“Verifying your Emulex adapter installation” on page 20](#).
- Follow the instructions in [“Verifying the Fibre Channel connections” on page 20](#).
- Migrate the file system to an FC device:
 - [“Moving the /store file system to a Fibre Channel solution” on page 22](#).
 - [“Moving the /store/ariel file system to a Fibre Channel solution” on page 23](#).

Multipath Fibre Channel configuration in standard QRadar deployments

To move data from a QRadar Console or managed host to multiple FC storage devices:

- Follow the instructions in [“Verifying your Emulex adapter installation” on page 20](#).
- Follow the instructions in [“Verifying the Fibre Channel connections” on page 20](#).
- Follow the instructions in [“Moving the /store file system to a multipath Fibre Channel solution” on page 25](#).

Fibre Channel configuration in HA deployments

To move data in an HA deployment to an FC storage device:

- Follow the instructions in [“Verifying your Emulex adapter installation” on page 20](#).
- Follow the instructions in [“Verifying the Fibre Channel connections” on page 20](#).
- Migrate the file system to an FC device:
 - [“Moving the /store file system to a Fibre Channel solution” on page 22](#).
 - [“Moving the /store/ariel file system to a Fibre Channel solution” on page 23](#)

- Follow the instructions in [“Configuring the mount point for the secondary HA host” on page 26](#) on the secondary appliance.

If you configure FC in an HA deployment and your primary HA host fails, your FC device can be used to maintain data consistency with your secondary HA host. For more information about data consistency and shared storage in an HA environment, see the *IBM QRadar High Availability Guide*.

Multipath Fibre Channel configuration in HA deployments

To move data in an HA deployment to multiple FC storage devices:

- Follow the instructions in [“Verifying your Emulex adapter installation” on page 20](#).
- Follow the instructions in [“Verifying the Fibre Channel connections” on page 20](#).
- Follow the instructions in [“Moving the /store file system to a multipath Fibre Channel solution in an HA deployment” on page 26](#).

Verifying your Emulex adapter installation

You must verify that an Emulex LPe12002 or LPe16002B Host Bus adapter is attached and installed with the correct firmware and driver versions.

Before you begin

To use the Fibre Channel protocol, you must install an Emulex LPe12002 or LPe16002B Host Bus adapter on your IBM QRadar appliance. In a high-availability (HA) deployment, you must install an Emulex LPe12002 or LPe16002B card on the primary and secondary HA host.

The Emulex LPe Host Bus adapter must use the following driver and firmware versions, or later:

- Driver version: 8.3.5.68.5p
- Firmware version: 1.10A5(U3D1.10A5), sli-3

Procedure

1. Using SSH, log in to your QRadar host as the root user:
2. Verify that an Emulex LPe12002 or LPe16002B card is attached by typing the following command:

```
hbacmd listhbas
```

If no result is displayed, contact your system administrator.

3. Verify that the Emulex card is using the correct firmware and driver versions by typing the following command:

```
hbacmd HBAAttrib <device_id>
```

device_id is the Port WWN that is displayed in the preceding step.

What to do next

[“Verifying the Fibre Channel connections” on page 20](#)

Verifying the Fibre Channel connections

You must identify the disk volume on the external Fibre Channel device. If required, you must also create a partition on the volume.

Before you begin

[Verify your Emulex adapter installation.](#)

Procedure

1. Attach both Fibre Channel cables to the Emulex LPe12002 or LPe16002B Host Bus adapter on your QRadar Console appliance.
2. Using SSH, log in to your QRadar Console as the root user.
3. Identify the Fibre Channel volume by typing the following command:

```
ls -l /dev/disk/by-path/*-fc-*
```

If multiple Fibre Channel devices are attached and you cannot identify the correct Fibre Channel volume, contact your system administrator.

4. If there is no partition on the Fibre Channel volume, then create a partition on the Fibre Channel device volume.

Note: QRadar does not support the creation of a partition on a Logical Volume Manager (LVM) drive. The file system fails to mount during a system boot, when a partition is created on an LVM drive.

- a) To create a partition, use the GNU parted command:

```
parted /dev/<volume>
```

- b) Configure the partition label to use GPT by typing the following command:

```
mklabel gpt
```

- c) If the following message is displayed, type Yes.

```
Warning: The existing disk label on /dev/<volume> will be
destroyed and all data on this disk will be lost. Do you want to
continue?
```

- d) To create the partition, type the following command:

```
mkpart primary 0% 100%
```

- e) Set the default units to TB by typing the following command:

```
unit TB
```

- f) Verify that the partition is created by typing the following command:

```
print
```

- g) Exit from GNU parted by typing the following command:

```
quit
```

- h) Update the kernel with the new partition data by typing the following command:

```
partprobe /dev/<volume>
```

You might be prompted to restart the appliance.

- i) To verify that the partition is created, type the following command:

```
cat /proc/partitions
```

5. Reformat the partition and make a file system.

- To create an XFS file system, type the following command:

```
mkfs.xfs -f /dev/<partition>
```

- To create an ext4 file system, type the following command:

```
mkfs.ext4 -f /dev/<partition>
```

- To create an XFS file system for multipath Fibre Channel, type the following command:

```
mkfs.xfs -f -L multiPath /dev/<partition>
```

- To create an ext4 file system for multipath Fibre Channel, type the following command:

```
mkfs.ext4 -f -L multiPath /dev/<partition>
```

What to do next

If you are moving the `/store` file system to a Fibre Channel solution, go to [“Moving the `/store` file system to a Fibre Channel solution”](#) on page 22.

If you are moving the `/store/ariel` file system to a Fibre Channel solution, go to [“Moving the `/store/ariel` file system to a Fibre Channel solution”](#) on page 23.

If you are moving the `/store` file system to a multipath Fibre Channel solution, go to [“Moving the `/store` file system to a multipath Fibre Channel solution”](#) on page 25.

If you are moving the `/store` file system to a multipath Fibre Channel solution in an HA deployment, go to [“Moving the `/store` file system to a multipath Fibre Channel solution in an HA deployment”](#) on page 26.

Moving the `/store` file system to a Fibre Channel solution

You can move the IBM QRadar data that is maintained in the `/store` file system and mount the `/store` file system to a Fibre Channel (FC) device partition.

Before you begin

[“Verifying the Fibre Channel connections”](#) on page 20

Procedure

1. After the QRadar installation, connect QRadar with fibre channel and restart.
2. Stop the QRadar services by typing the following commands in the order specified:

```
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop hostservices
systemctl stop systemStabMon
systemctl stop crond
```

3. Unmount the file systems by typing the following commands:

```
umount /store
```

The `/transient` file system is mounted only when the `/store` file system is XFS.

4. Create the `/store_old` directory by typing the following command:

```
mkdir /store_old
```

5. Derive the device partition universal unique identifier (UUID) by typing the following command:

```
blkid /dev/partition
```

6. Edit the `/etc/fstab` file to update the existing `/store` file system mount point to `/store_old`.
7. Add a mount point for the `/store` file system by adding the following text to the `/etc/fstab` file:

- If the file system is XFS and you are not using HA, add the following text:

```
UUID=uuid /store xfs inode64,logbsize=256k,noatime,nobarrier 0 0
```

- If the file system is XFS and you are using HA, add the following text:

```
UUID=uuid /store xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

- If the file system is ext4 and you are not using HA, add the following text:

```
UUID=uuid /store ext4 noatime,nobarrier 0 0
```

- If the file system is ext4 and you are using HA, add the following text:

```
UUID=uuid /store ext4 noatime,noauto,nobarrier 0 0
```

Save and close the file.

8. Mount the `/store` file system to the FC device partition by typing the following command:

```
mount /store
```

9. Mount the `/store_old` file system to the local disk by typing the following command:

```
mount /store_old
```

10. Copy the data to the Fibre Channel partition by typing the following command:

```
cp -af /store_old/* /store
```

11. Unmount `/store_old` by typing the following command:

```
umount /store_old
```

12. Remove the `/store_old` directory by typing the following command:

```
rmdir /store_old
```

13. Edit the `/etc/fstab` file to remove the `/store_old` mount point entry.

14. Start the QRadar services by typing the following commands in the order specified:

```
systemctl start crond
systemctl start systemStabMon
systemctl start hostservices
systemctl start tomcat
systemctl start ecs-ec-ingress
systemctl start hostcontext
```

15. Remove the local copy of `/store` from the logical volume manager (LVM) by typing the following command:

```
lvchange -an /dev/storerhel/store 2>/dev/null
lvrename /dev/storerhel/store /dev/storerhel/storeold 2>/dev/null
```

16. Verify the Fibre Channel mount point by typing the following command:

```
df -h
```

Moving the `/store/ariel` file system to a Fibre Channel solution

You can move the IBM QRadar data that is maintained in the `/store/ariel` file system and mount the `/store/ariel` file system to a Fibre Channel (FC) device partition.

Before you begin

See [“Verifying the Fibre Channel connections”](#) on page 20.

Procedure

1. Connect QRadar to the Fibre Channel and restart.

2. Stop the QRadar services by typing the following commands in the order specified:

```
systemctl stop systemStabMon
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop hostservices
systemctl stop crond
```

3. Create a temporary directory by typing the following command:

```
mkdir /tmp/fcdata
```

4. Mount the Fibre Channel storage partition to the temporary directory by typing the following command, where *<partition>* is the name of the device partition:

```
mount /dev/<partition> /tmp/fcdata
```

5. Copy the data to the Fibre Channel device by typing the following command:

```
cp -af /store/ariel/* /tmp/fcdata
```

6. Unmount the Fibre Channel partition by typing the following command:

```
umount /tmp/fcdata
```

7. Verify the Universally Unique Identifier (UUID) of the Fibre Channel device partition by typing the following command, where *<partition>* is the name of the device partition:

```
blkid /dev/<partition>
```

8. Edit the `fstab` file by typing the following command:

```
vi /etc/fstab
```

9. Add the mount point for the `/store/ariel` file system by adding the following text, where *<uuid>* is the UUID of the Fibre Channel device partition, to the `/etc/fstab` file.

If the file system is XFS:

```
UUID=<uuid> /store/ariel xfs inode64,logbsize=256k,noatime,nobarrier 0 0
```

If the file system is ext4:

```
UUID=<uuid> /store/ariel ext4 defaults,noatime,nobarrier 0 0
```

10. Save and close the file.

11. Mount the `/store/ariel` file system to the FC device partition by typing the following command:

```
mount /store/ariel
```

12. Start the QRadar services by typing the following commands in the order specified:

```
systemctl start crond
systemctl start hostservices
systemctl start tomcat
systemctl start ecs-ec-ingress
systemctl start hostcontext
systemctl start systemStabMon
```

13. Verify the Fibre Channel mount point by typing the following command:

```
df -h
```

Moving the /store file system to a multipath Fibre Channel solution

In IBM QRadar, you can implement multipath Fibre Channel. If you experience a storage area network or SAN switch issue, multipath provides extra redundancy to prevent disruption to flow and event data.

Before you begin

Ensure that you completed the following tasks:

- [Verify your Emulex adapter installation.](#)
- [Verify the Fibre Channel connections.](#)

Procedure

1. Log in to your QRadar Console as the root user, using SSH.
2. Identify a storage area network (SAN) partition by typing the following command:

```
blkid -o list
```

3. Stop the QRadar services by typing the following commands in the order specified:

```
systemctl stop systemStabMon
systemctl stop hostcontext
systemctl stop ecs-ec-ingress
systemctl stop tomcat
systemctl stop imq
systemctl stop hostservices
systemctl stop crond
```

4. Unmount the file systems by typing the following commands:

```
umount /store
```

5. Create a /store_old directory by typing the following command:

```
mkdir /store_old
```

6. Determine the Universally Unique Identifier (UUID) of the device partition by typing the following command:

```
blkid /dev/<partition>
```

7. Edit the /etc/fstab file.

- a) Replace the /store mount point with /store_old.
- b) Add the new /store mount point.

If the file system is XFS and you are not using HA, add the following text:

```
UUID=<uuid> /store xfs inode64,logbsize=256k,noatime,nobarrier 0 0
```

If the file system is XFS and you are using HA, add the following text:

```
UUID=<uuid> /store xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

If the file system is ext4 and you are not using HA, add the following text:

```
UUID=<uuid> /store ext4 noatime,nobarrier 0 0
```

If the file system is ext4 and you are using HA, add the following text:

```
UUID=<uuid> /store ext4 noatime,noauto,nobarrier 0 0
```

8. Mount the file systems and copy the data to your device by typing the following commands:

```
mount /store
mount /store_old
cp -af /store_old/* /store
umount /store_old
```

9. Start the QRadar services by typing the following commands in the order specified:

```
systemctl start crond
systemctl start hostservices
systemctl start imq
systemctl start tomcat
systemctl start ecs-ec-ingress
systemctl start hostcontext
systemctl start systemStabMon
```

10. Enable Fibre Channel multipath by typing the following command:

```
mpathconf --enable
```

11. Rename the local copy of /store by typing the following command:

```
lvchange -an /dev/storerhel/store 2>/dev/null
lvrename /dev/storerhel/store /dev/storerhel/storeold 2>/dev/null
```

12. Edit the /etc/fstab file to remove the /store_old mount point entry.

Moving the /store file system to a multipath Fibre Channel solution in an HA deployment

To use multipath Fibre Channel storage in a high-availability (HA) environment, you must configure the primary HA host and the secondary HA host to use the same storage partition.

About this task

Important:

- You must configure multipath on the primary and secondary HA appliances before you initiate HA syncing.
- Before you add HA to a Fibre Channel configuration, confirm that /store/backup is local. Create links to /store/backup only after adding HA.

Procedure

1. Verify that the correct Fibre Channel hardware is installed on your secondary HA appliance. For more information, see [“Verifying your Emulex adapter installation”](#) on page 20
2. Verify the HA Fibre Channel connections. For more information, see [“Verifying the Fibre Channel connections”](#) on page 20
3. Configure Fibre Channel on your primary HA appliance. For more information, see [“Moving the /store file system to a multipath Fibre Channel solution”](#) on page 25
4. Configure the file system mount point for the secondary HA host. For more information, see [“Configuring the mount point for the secondary HA host”](#) on page 26.

Configuring the mount point for the secondary HA host

You must configure the mount point on the secondary high-availability (HA) host for the file system that is moved to a Fibre Channel storage device.

Procedure

1. Using SSH, log in to the secondary HA host as the root user.
2. Derive the UUID by typing the following command:


```
blkid /dev/<partition>
```

3. Update the kernel with the Fibre Channel partition data by typing the following command:

```
partprobe
```

Troubleshoot: If you see a warning error message that the kernel cannot read the partition table, type the following command: `ls -l /dev/disk/by-uuid/<UUID>`. If no output is displayed, then restart the secondary HA host by typing `reboot`.

4. If you are moving the `/store` directory, unmount the file systems by typing the following command:

```
umount /store
```

The `/transient` file system is mounted only when the `/store` file system is XFS.

5. If you redirected the `/store` file system to an offboard device, choose one of the following options to edit the `/etc/fstab` file.

- If the `/store` file system is an XFS file system, update the following lines. For each line, copy the text into a text editor, remove any line breaks, and paste as a single line.

```
UUID=<uuid> /store xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

- If the `/store` file system is ext4, update the following line:

```
UUID=<uuid> /store ext4 defaults,noatime,noauto,nobarrier 0 0
```

6. If you are moving the `/store/ariel` file system to an offboard device, choose one of the following options to edit the `/etc/fstab` file.

- If the `/store/ariel` file system is an XFS file system, update the following lines. For each line, copy the text into a text editor, remove any line breaks, and paste as a single line.

```
UUID=<uuid> /store/ariel xfs inode64,logbsize=256k,noatime,  
noauto,nobarrier 0 0
```

- If the `/store/ariel` file system is ext4, update the following line:

```
UUID=<uuid> /store/ariel ext4 defaults,noatime,noauto,nobarrier 0 0
```

7. If you are moving the `/store` file system, rename the local copy of `/store` by typing the following command:

```
lvchange -an /dev/storerhel/store 2>/dev/null  
lvrename /dev/storerhel/store /dev/storerhel/storeold 2>/dev/null
```

What to do next

Create an HA cluster. For more information, see *IBM QRadar High Availability Guide*.

Removing HA from a Fibre Channel solution

Remove the 'noauto' value from the `/store` mount point entry in `/etc/fstab` before you remove HA.

Procedure

1. Edit the `/etc/fstab` file to remove the 'noauto' value from the `/store` mount point entry.

- If the file system is XFS, the mount point entry should have the following text:

```
UUID=uuid /store xfs inode64,logbsize=256k,noatime,nobarrier 0 0
```

- If the file system is ext4, the mount point entry should have the following text:

```
UUID=uuid /store ext4 noatime,nobarrier 0 0
```

Save and close the file.

2. Sign in to the QRadar user interface.
3. Click **Admin**.
4. Click the **System and License Management** icon.
5. Select the HA host that you want to remove.
6. From the toolbar, select **High Availability > Remove HA Host**.
7. Click **OK**.

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

Chapter 4. NFS offboard storage device

You can back up the IBM QRadar data to an external Network File System (NFS).



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

NFS storage with a stand-alone QRadar Console

To move backup files to NFS from a stand-alone QRadar Console, follow the instructions in [“Moving backups to an NFS”](#) on page 29.

NFS storage with a new HA deployment

To move backup files to NFS for a new HA deployment:

- Follow the instructions in [“Moving backups to an NFS”](#) on page 29 for your primary HA appliance.
- Follow the instructions in [“Configuring a mount point for a secondary HA host”](#) on page 31 for your secondary HA appliance.
- Create an HA cluster. For more information, see *IBM QRadar High Availability Guide*.

NFS storage with an existing HA deployment

To move backup files from an existing HA deployment, follow the instructions in [“Configuring NFS backup on an existing HA cluster”](#) on page 32.

Moving backups to an NFS

You can configure Network File System (NFS) for a stand-alone QRadar appliance, or a QRadar appliance that you are making the primary host in an HA deployment.

Before you begin

You must ensure that the QRadar appliance can reach the NFS server.

Important:

Maintain a local copy (`backup.nfs`) of your backup on your system so if the NFS mount fails the backups are still available. Monitor the directory that holds the local backups carefully to ensure the directory you use to hold your backups doesn't cause any disk storage issues.

About this task



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

Even though the risk is low with a Linux OS, ransomware can encrypt all mounted remote drives. If you use an NFS mount, you can reduce your risk by mounting only the NFS drive while you copy data from a local drive to the NFS-mounted drive. Then, remove the NFS-mounted drive for daily operations.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

Procedure

1. Run nightly backups to the local drive, `/store/backup`.
2. Use SSH to log in to the QRadar host as the root user.
3. Start NFS services by typing the following commands:

```
systemctl enable rpcbind
systemctl start rpcbind
```

4. Add the following line to the `/etc/fstab` file.

```
nfsserver:/nfs/export/path /store/backup nfs rw,soft,intr,noac 0 0
```

You might need to adjust the settings for the NFS mount point to accommodate your configuration.

5. Move your backup files from the existing directory to a temporary location by typing the following commands:

```
cd /store/
mv backup backup.local
```

6. Create a new backup directory by typing the following command:

```
mkdir /store/backup
```

7. Set the permissions for the NFS volume by typing the following command:

```
chown nobody:nobody /store/backup
```

8. Mount the NFS volume by typing the following command:

```
mount /store/backup
```

The root user must have read and write access to the mounted NFS volume because the `hostcontext` process runs as root user.

Use the local copy of your backup that you created. See, [Important note](#).

9. Verify that `/store/backup` is mounted by typing the following command:

```
df -h
```

10. Copy the backup files from the temporary location to the NFS volume by typing the following command:

```
cp -f /store/backup.local/* /store/backup
```

11. Verify the files by typing the following commands:

```
sha256sum /store/backup.local/* > backuplocal.sha256.txt
sha256sum /store/backup.nfs/* > backupnfs.sha256.txt
diff backuplocal.sha256.txt backupnfs.sha256.txt
```

If you see differences between the files, stop and determine the reason. One reason might be that your copy filled the destination partition, or another reason might be that there was a network outage during your copy procedure.

12. After you verify the copy procedure was successful, remove the `backup.local` directory by typing the following commands:

```
cd /store
rm -r backup.local
```

When you remove the `backup.local` directory, you prevent the local partition from filling.

What to do next

If you are setting up NFS for a new HA deployment, follow the instructions in [“Configuring a mount point for a secondary HA host”](#) on page 31 for your secondary HA appliance.

Configuring a mount point for a secondary HA host

On your existing secondary high-availability (HA) host, you must configure an NFS mount point for the alternative IBM QRadar backup file location.

Before you begin

Ensure that the HA secondary host can connect with the NFS server.

About this task



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

Procedure

1. Using SSH, log in to the QRadar secondary HA host as the root user:
2. Create a backup file location that matches the backup file location on your primary HA host. The default location for QRadar backups is `/store/backup`.

For more information, see [“Configuring NFS backup on an existing HA cluster”](#) on page 32.

Restriction: Do not create your new backup location under the `/store` file system. Use a different directory, such as `/backup` or `/nfs`.

3. Start the NFS services by typing the following commands:

```
systemctl enable rpcbind
systemctl start rpcbind
```

4. Add the following line to the `/etc/fstab` file:

```
<hostname>:<shared_directory> <backup_location> nfs
rw,soft,intr,clientaddr=<HA_IP_address> 0 0
```

5. Mount the new QRadar backup file location by typing the following command:

```
mount <backup_location>
```

Configuring NFS backup on an existing HA cluster

You can configure Network File System (NFS) for an existing high-availability cluster.

About this task



Warning:

If you use NFS or a Windows share for offboard storage, your system can lock and cause an outage. This practice is not supported by IBM QRadar.

If you choose to use NFS anyway, NFS can be used only for daily backup data, such as the `/store/backup` directory. You cannot use NFS for storing active data, which includes the PostgreSQL and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

Procedure

1. Use SSH to log in to the primary HA host as the root user.
2. Start NFS services by typing the following commands:

```
systemctl enable rpcbind
systemctl start rpcbind
```

3. Add the following line to the `/opt/qradar/ha/fstab.back` file.

```
nfsserver:/nfs/export/path /store/backup nfs rw,soft,intr,noac 0 0
```

You might need to adjust the settings for the NFS mount point to accommodate your configuration.

4. Add the same line to the `/etc/fstab` file, preceded by `#HA`.

```
#HA nfsserver:/nfs/export/path /store/backup nfs rw,soft,intr,noac 0 0
```

You might need to adjust the settings for the NFS mount point to accommodate your configuration.

5. Repeat steps 1 - 4 on the secondary HA host.
6. Move your backup files from the existing directory on the primary HA host to a temporary location by typing the following commands:


```
cd /store/
mv backup backup.local
```

7. Create a new backup directory on the primary HA host by typing the following command:

```
mkdir /store/backup
```

8. Set the permissions for the NFS volume on the primary HA host by typing the following command:

```
chown nobody:nobody /store/backup
```

9. On the navigation menu () , click **Admin**.
10. Click **Advanced > Deploy Full Configuration**.
All services restart.
11. Verify that the `/store/backup` mount point is listed in the output of the following command on the primary and secondary HA hosts:

```
grep MOUNTS /opt/qradar/ha/ha.conf
```

12. Verify that `/store/backup` is mounted on the primary HA host by typing the following command:

```
df -h
```

13. On the primary HA host, move the backup files from the temporary location to the NFS volume by typing the following command:

```
mv -f /store/backup.local/* /store/backup
```

14. Remove the `backup.local` directory by typing the following commands:

```
cd /store  
rm -rf backup.local
```


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.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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>

