

IBM Security QRadar
Version 7.2.1

Offboard Storage Guide



Note: Before using this information and the product that it supports, read the information in [Notices and Trademarks](#) on [page 47](#).

CONTENTS

INTRODUCTION TO IBM SECURITY QRADAR OFFBOARD STORAGE

Intended audience	3
Technical documentation	3
Contacting customer support	3
Statement of good security practices	3

1 OFFBOARD STORAGE OVERVIEW

External storage considerations	5
Local storage	5
Multiple appliances	5
Existing hardware, infrastructure, and expertise	6
Data retention and fault tolerance	6
File system migration	6
Performance impact	6
Storage expansion	7
External storage options	7
Fibre Channel	7
iSCSI	7
NFS	8
External storage limitations	8
Offboarding with High Availability	9

2 OFFBOARD DISK PARTITIONS

Creating a disk partition	11
-------------------------------------	----

3 ISCSI IMPLEMENTATION

iSCSI HA considerations	15
Secondary network interfaces	16
Configure iSCSI in a standard QRadar deployment	16
Connecting QRadar to an iSCSI network	16
Assigning and configuring the iSCSI volumes	17
Migrating /store/ariel to an iSCSI storage solution	18
Migrating /store to an iSCSI storage solution	19
Auto-mounting the iSCSI volume	21
Configure iSCSI in an HA deployment	22
Connecting an HA secondary host to an iSCSI device	22

Assigning and configuring iSCSI volumes for the HA secondary host	23
Configuring the mount point for the secondary HA host	23
Configuring the secondary HA host to auto-mount the iSCSI volume.	24
Verifying iSCSI connections	25
Troubleshoot iSCSI	26
Reconfigure iSCSI when restoring a failed primary HA console	26
Detecting Disk Errors	26
Unmounting and remounting the iSCSI volume	26

4 FIBRE CHANNEL IMPLEMENTATION

Fibre Channel considerations	29
Verifying your Emulex adapter installation	30
Configure Fibre Channel in a standard QRadar deployment	31
Verifying the Fibre Channel connections	31
Migrating /store to Fibre Channel	32
Migrating /store/ariel to Fibre Channel	34
Verifying the Fibre Channel mount point	35
Configure Fibre Channel in an HA deployment	36
Verifying the HA Fibre Channel connections	36
Configuring the mount point for the secondary HA host	36

5 USING NFS FOR QRADAR BACKUPS

Migrate backups from a stand-alone QRadar console	39
Enabling NFS connections	39
Configuring the NFS mount	40
Migrating existing backups to NFS	41
Migrating backups from a new QRadar HA deployment	42
Migrating backups from an existing QRadar HA deployment	42
Configuring a new QRadar backup location	43
Changing the QRadar backup file location	44
Configuring a mount point for a secondary HA host	44

A NOTICES AND TRADEMARKS

Notices	47
Trademarks	49
Privacy policy considerations	49

INDEX

INTRODUCTION TO IBM SECURITY QRADAR OFFBOARD STORAGE

This information is intended for use with *IBM® Security QRadar®* and provides information on how to migrate the /store or /store/ariel file systems using external storage devices.

Unless otherwise noted, all references to QRadar refer to IBM Security QRadar SIEM, IBM Security QRadar Log Manager, and IBM Security QRadar Network Anomaly Detection.

Note: Any references to the *IBM Security QRadar High Availability Guide* are not applicable to IBM Security QRadar Network Anomaly Detection.

Intended audience System administrators responsible for configuring offboard storage devices must have administrative access to *IBM Security QRadar* and your network devices and firewalls. The system administrator must have knowledge of your corporate network and networking technologies.

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

Contacting customer support For information on 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 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.

1

OFFBOARD STORAGE OVERVIEW

You can offboard your /store, /store/ariel, or /store/backup file systems to an iSCSI, Fibre Channel, or Network File System (NFS) external storage solution.

You can implement an offboard storage solution by using a standard IBM Security QRadar primary console or in a High Availability (HA) environment. When you use iSCSI or Fibre Channel with HA, then the external device ensures data consistency if your primary HA host fails.

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

Note: The IBM Security QRadar High Availability Guide is not applicable to IBM Security QRadar Network Anomaly Detection

To offboard a QRadar file system to an external storage device, you must configure your iSCSI, Fibre Channel, or NFS external storage devices. If you implement external storage with HA, you must configure these devices on the primary HA host and the secondary HA host.

External storage considerations

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

Local storage

The disk on your QRadar appliance is faster than external storage and supports up to 16 TB of data. When possible, use local storage 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 an existing deployment can increase capacity by using available external storage, then external storage might be appropriate for your deployment.

Existing hardware, infrastructure, and expertise

Your existing infrastructure and experience with storage area networks (SANs) 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 more specialized hardware.

For more information, see [External storage options](#).

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 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. For more information, see [File system migration](#).

File system migration

An offboard storage solution can be used to migrate the /store file system or specific subdirectories. For example, /store/ariel.

You can migrate the /store file system when you have to increase the fault tolerance levels in your QRadar deployment. Each option has a different impact on QRadar performance. For more information see, [Performance impact](#).

Migrating the /store file system to an external device can provide an alternative resolution to implementing HA. For more information about HA, see the *IBM Security QRadar High Availability Guide*.

The /store/ariel directory is the most commonly offboarded file system. By offboarding the /store/ariel file system, you can migrate collected log and network activity data to external storage and ensure that the local disk remains used for the postgresql database and temporary search results.

Administrators can offboard the following types of QRadar data:

- Postgres 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).

Performance impact

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

After migration, all data I/O to the /store file system is no longer performed on the local disk. Before you migrate 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/ariel/persistent_data file system to the unused /store file partition.
- Searches that are marked as saved are also in the /store/ariel/persistent_data directory. If you experience a local disk failure, these searches are not saved.

Storage expansion Any subdirectory in the /store file system can be used as a mount point for your external storage device.

By creating multiple volumes and mounting /store/ariel/events and /store/ariel/flows, you can expand your storage capabilities past the 16 TB file system limit supported by QRadar.

If you have to migrate dedicated event or flow data, you might configure more specific mount points. For example, /store/ariel/events/records and /store/ariel/events/payloads. Specific mount points can provide up to 32 TB of storage for either Log or Network Activity data.

External storage options

You can use iSCSI, Fibre Channel, or NFS to provide an offboard storage solution. These options must be reviewed against the benefits of local storage.

Onboard disks provide a faster solution than offboard storage devices. Local disk storage on QRadar appliances support read speeds of between 200 Mbps to 400MBps and write speeds of almost 200 Mbps. When multiple appliances are deployed, performance and capacity scale at the same rate.

Fibre Channel Fibre Channel provides the fastest offboard performance by using 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:

- Disk or spindle counts per volume.
- The number of concurrent sessions.
- The cache capacity in the SAN controllers.

For more information, see [Fibre Channel Implementation](#).

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. For more information, see [Secondary network interfaces](#).

Note: Using a dedicated interface, you are limited to 1 Gbps and might experience only 200 MBps to 400 MBps. Your iSCSI storage device might be capable of providing only 25 MBps to 50 MBps I/O performance.

For more information, see [iSCSI Implementation](#).

NFS An NFS solution must not be used to store active QRadar data. You can offboard the /store/backup file system to an external NFS.

If the /store file system is mounted to an NFS solution, postgres data can be corrupted. If the /store/ariel file system is mounted to NFS, QRadar experiences performance issues.

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

NFS storage operates over existing management ethernet networks and is limited to performance levels of 20 MBps to 50 MBps. The NFS protocol incurs extra overhead for file access, locking, and network permissions. Performance impact can be remediated by using a dedicated network interface.

Note: If NFS is used only for backups, the same NFS share can be used for backups. This is because the backup files on each host also contain the systems host name, enabling the identification of each backup file. If you are storing a longer period of data on your NFS shares, then consider a separate share or export for each appliance in your deployment.

For more information, see [Using NFS for QRadar Backups](#).

External storage limitations

If you deploy an external storage solution, you must be aware of the limitations.

QRadar does not support multiple systems that access the same block device. If you are configuring iSCSI in an HA environment, do not mount the iSCSI or Fibre Channel volumes on the secondary host while the primary host is operational and accessing the volumes. For more information, see [Configure iSCSI in an HA deployment](#).

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

- Data write performance can be impacted.
- The performance of user interface searches can be impacted.

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

Offboarding with High Availability

If you choose to offboard /store in an HA environment, the /store file system is not replicated by using Disk Replication Block Device (DRBD).

If you offboard the /store/ariel file system and maintain the /store file system on local disk, then the /store file system is synchronized with the secondary HA host by using DRBD. By default, when your environment is configured for HA, DRBD is enabled.

For more information, see the *IBM Security QRadar High Availability Guide*. The IBM Security QRadar High Availability Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

2

OFFBOARD DISK PARTITIONS

If you configure an iSCSI or Fibre Channel storage solution, then you can create a partition on the volume of the external disk.

Release 7.2.1 of QRadar uses the xfs file system. You can create the partition on your iSCSI device with either an ext4 or xfs file system.

iSCSI and Fibre Channel disk partitions are created by using GUID Partition Table (GPT). A new device partition can be used as the mount point for the file system that you offboard. For example, the /store or /store/ariel file system.

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.

Note: These procedures assume an advanced knowledge of the Linux operating system.

Creating a disk partition

You can create a partition on the volume of the external iSCSI or Fibre Channel storage device.

Before you begin

Before you can create a partition on an iSCSI or Fibre Channel device, you must verify that the device is connected and that you can access the correct disk volume on the device.

If you are configuring Fibre Channel, see [Verifying your Emulex adapter installation](#) and

Procedure

Step 1 Using SSH, log in to the primary QRadar Console as the root user.

Username: `root`

Password: `<password>`

Step 2 Choose one of the following options:

- If you are creating an iSCSI device partition, go to [Step 3](#).

- If you are creating a Fibre Channel device partition, go to [Step 4](#).

Step 3 Identify the iSCSI volume:

- a Clear the kernel ring buffer by typing the following command:

```
dmesg -c
```

- b Reload the iSCSI service by typing the following command:

```
service iscsi restart
```

- c Locate the iSCSI device volume name by typing the following command:

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

Step 4 Identify the Fibre Channel volume by typing the following command:

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

Note: If multiple Fibre Channel devices are attached and you cannot identify the correct Fibre Channel volume, then contact your System Administrator.

Step 5 Start GNU parted by typing the following command:

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

Where: **<volume>** is the iSCSI device volume that you identified in [Step 3](#) or [Step 4](#). For example: `sdb`.

Step 6 Configure the partition label to use GPT by typing the following command:

```
mklabel gpt
```

Step 7 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?
```

Step 8 Create a partition on the iSCSI disk volume by typing the following command:

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

Step 9 Set the default units to TB by typing the following command:

```
unit TB
```

Step 10 Verify that the partition is created by typing the following command:

```
print
```

Step 11 Exit from GNU parted by typing the following command:

```
quit
```

Step 12 Update the kernel with the new partition data by typing the following command:

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

Where: **<volume>** is the iSCSI device volume that you identified in [Step 3](#) or [Step 4](#). For example: `sdb`.

If the following message is displayed, then restart your appliance. Type **reboot**.

Warning: WARNING: the kernel failed to re-read the partition table on /dev/sda (Device or resource busy). As a result, it may not reflect all of your changes until after reboot.

Step 13 Verify that the partition is created.

```
cat /proc/partitions
```

Step 14 Reformat the partition and make a file system. Choose one of the following options:

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

Where `<partition>` is the name of the new partition. For example: `sdb1`.

3

ISCSI IMPLEMENTATION

iSCSI can be configured in a standard IBM Security QRadar deployment or in a High Availability (HA) environment.

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 Fibre Channel disk. For more information, see [Creating a disk partition](#).

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. For more information on data consistency and shared storage in an HA environment, see the *IBM Security QRadar High Availability Guide*.

Administrators can perform the following tasks:

- 1 Configure iSCSI with a standard QRadar console, see [Configure iSCSI in a standard QRadar deployment](#).
- 2 Configure iSCSI in a QRadar HA deployment, see [Configure iSCSI in an HA deployment](#).

Note: iSCSI configuration requires an advanced knowledge of the Linux operating system.

iSCSI HA considerations

Before you implement iSCSI, you must be aware of the configuration options and HA setup requirements.

The procedures for configuring iSCSI are different for a primary HA host and secondary HA host. To configure iSCSI you must ensure that the primary HA host and secondary HA host are not connected in an HA cluster. For more information on HA clustering, see the *IBM Security QRadar High Availability Guide*.

During iSCSI configuration in an HA environment, you can access and review the /var/log/messages file for specific errors with your iSCSI storage configuration.

Initiatornames

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

The initiatorname is configured in the `/etc/iscsi/initiatorname.iscsi` file and is used by QRadar to identify the volume on the iSCSI device where the `/store` or `/store/ariel` file system is mounted.

Secondary network interfaces

You can configure a secondary network interface with a private IP address to connect to an iSCSI Storage Area Network (SAN).

If you choose this option, then configure your SAN using this method to improve performance. If you configure a secondary network interface, you will require address information from your SAN network manager. For more information on configuring a network interface, see your Administration Guide.

Note: Your network configuration might differ, however, this guide assumes that your management interface is `eth0` and your iSCSI interface is `eth1`.

Configure iSCSI in a standard QRadar deployment

You can configure iSCSI in a standard deployment by using a QRadar console.

Administrators must perform the following tasks in sequence:

- 1 Prepare QRadar to connect to the iSCSI network, see [Connecting QRadar to an iSCSI network](#).
- 2 Assign and configure the iSCSI volumes, see [Assigning and configuring the iSCSI volumes](#).
- 3 Choose from one of the following options:
 - [Migrating /store/ariel to an iSCSI storage solution](#).
 - [Migrating /store to an iSCSI storage solution](#)
- 4 Auto-mount the iSCSI volumes, see [Auto-mounting the iSCSI volume](#).
- 5 Verify the iSCSI connections to the primary HA host, see [Verifying iSCSI connections](#).

Connecting QRadar to an iSCSI network

You must prepare the QRadar host to connect to your iSCSI network.

About this task

This procedure applies to the configuration of iSCSI for a stand-alone QRadar console and a QRadar console that is being used as primary HA host in an HA deployment.

Procedure

Step 1 Using SSH, log in to the QRadar Console as the root user.

Username: `root`

Password: `<password>`

Step 2 Configure your system to identify the iscsi device volume:

a Open the `initiatorname.iscsi` file for editing by typing the following command:

```
vim /etc/iscsi/initiatorname.iscsi
```

b Edit the file with the iSCSI qualified name for your host. Type the following:

```
InitiatorName=iqn.<yyyy-mm>.{reversed domain name}:<hostname>
```

For example: `InitiatorName=iqn.2008-11.com.q11labs:p113`

c Save and close the file.

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

```
service iscsi restart
```

What to do next

Perform the steps in the procedure, [Assigning and configuring the iSCSI volumes](#).

Assigning and configuring the iSCSI volumes

You must assign and configure the volumes on the iSCSI device.

Before you begin

Perform the steps in the procedure, [Connecting QRadar to an iSCSI network](#).

Procedure

Step 1 Detect the volume on the iSCSI server by typing the following command:

```
iscsiadm -m discovery --type sendtargets --portal <IP address>:<port>
```

Where:

`<IP address>` is the IP address of the iSCSI server.

`<port>` is the port number of the iSCSI server. This is an optional parameter.

Step 2 Verify the login to the iSCSI server by typing the following command:

```
iscsiadm -m node -l
```

Step 3 Optional. Create a partition on the iSCSI device volume.

For more information, see [Creating a disk partition](#).

What to do next

You can migrate the `/store` or `/store/ariel` file system to an iSCSI device. Choose one of the following options:

- [Migrating /store/ariel to an iSCSI storage solution](#)

- [Migrating /store to an iSCSI storage solution](#)

Note: To retain optimal system performance, migrate /store/ariel.

Migrating /store/ariel to an iSCSI storage solution

You can migrate the 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

Perform the steps in the procedure, [Assigning and configuring the iSCSI volumes](#).

Procedure

Step 1 Stop the hostcontext service by typing the following command:

```
service hostcontext stop
service tomcat stop
service hostservices stop
service systemStabMon stop
service crond stop
```

Step 2 Move the existing mount point aside by typing the following commands:

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

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

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

Where <partition> is the name of the iSCSI device partition. For example: `sdb1`

Step 4 Configure the /store/ariel file system by using the fstab file:

a Open the fstab file for editing by typing the following command:

```
vim /etc/fstab
```

b Add the mount line for the new /store/ariel mount point by typing the following line:

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

Where:

<uuid> is the value derived in [Step 3](#).

c Save and close the file.

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

```
mkdir ariel
```

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

```
mount /store/ariel
```

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

```
df -h
```

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

```
mv /store/ariel_old/* /store/ariel
```

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

```
rmdir /store/ariel_old
```

Step 10 Start the hostcontext service by typing the following command:

```
service crond start
service systemStabMon start
service hostservices start
service tomcat start
service hostcontext start
```

What to do next

Perform the steps in the procedure, [Auto-mounting the iSCSI volume](#).

Migrating /store to an iSCSI storage solution

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

Before you begin

Perform the steps in the procedure, [Assigning and configuring the iSCSI volumes](#).

About this task

Migrating /store to your offboard storage device can take an extended period of time. For assistance with reducing the time taken to migrate your data, contact Customer Support or engage Professional Services.

Procedure

Step 1 Stop QRadar services by typing the following commands in sequence:

```
service hostcontext stop
service tomcat stop
service hostservices stop
service systemStabMon stop
service crond stop
```

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

```
umount /store/tmp
umount /store/ariel/persistent_data
umount /store
```

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

```
mkdir /store_old
```

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

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

Where <partition> is the name of the iSCSI device partition. For example: `sdb1`

- Step 5** Modify the file system mount points by using the `fstab` file:

- a Open the `fstab` file for editing by typing the following command:

```
vi /etc/fstab
```

- b Update the existing `/store` file system mount point to `/store_old`.

- c Add a new mount point for the `/store` file system by typing one of the following lines:

- If the `/store` file system is `ext4` then type:

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

- If the `/store` file system is `xfs` then type:

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

Where <uuid> is the UUID of the `iscsi` device partition.

- d Modify the `/store/tmp` mount line to use the following file system options:

```
noatime,noauto,nobarrier 0 0
```

- e If `/store/ariel/persistent_data` is listed in the `fstab` file, then type the following file system options:

```
xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

- f Save and close the file.

- Step 6** Mount `/store` to the iSCSI device partition by typing the following command:

```
mount /store
```

- Step 7** Mount `/store_old` to the local disk by typing the following command:

```
mount /store_old
```

- Step 8** Move the data to the iSCSI device by typing the following command:

```
mv -f /store_old/* /store
```

- Step 9** Re-mount the file systems by typing the following commands::

```
mount /store/tmp
mount /store/ariel/persistent_data
```

- Step 10** Unmount `/store_old` by typing the following command:

```
umount /store_old
```

- Step 11** Remove the `/store_old` mount point from the `/etc/fstab` file:

- a Open the `/etc/fstab` file for editing by typing the following command:

```
vi /etc/fstab
```

- b Remove the line for the /store_old mount point.
- c Save and close the file.

Step 12 Start QRadar services by typing the following commands in sequence:

```
service crond start
service systemStabMon start
service hostservices start
service tomcat start
service hostcontext start
```

What to do next

Perform the steps in the procedure, [Auto-mounting the iSCSI volume](#).

Auto-mounting the iSCSI volume

You must configure QRadar to auto-mount the iSCSI volume.

Before you begin

Perform the steps in one of the following procedures:

- [Migrating /store/ariel to an iSCSI storage solution](#)
- [Migrating /store to an iSCSI storage solution](#)

Procedure

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

```
chkconfig --add iscsi
chkconfig --level 345 iscsi on
```

Step 2 Create a symbolic link to the iscsi-mount script by typing the following command:

```
ln -s /opt/qradar/init/iscsi-mount /etc/init.d
```

Step 3 Add the iscsi-mount script to the startup by typing the following commands:

```
chkconfig --add iscsi-mount
chkconfig --level 345 iscsi-mount on
```

Step 4 Verify that the iSCSI device is correctly mounted by restarting your system.

- a Restart the system by typing the following command:


```
reboot
```
- b Ensure that the iSCSI mount point is retained by typing the following command:


```
df -h
```

What to do next

Optional. If you are configuring an HA environment, you must setup your secondary HA host using the same iSCSI connections that you used for you

primary HA host. For more information, see [Configure iSCSI in an HA deployment](#).

Configure iSCSI in an HA deployment

To use an iSCSI device in an HA environment, you must configure the primary HA host and secondary HA host to use the same iSCSI external storage device.

Administrators must perform the following tasks in sequence:

- 1 Configure the primary HA host with the iSCSI device, see [Configure iSCSI in a standard QRadar deployment](#)
- 2 Configure the secondary HA host with the iSCSI device, see [Connecting an HA secondary host to an iSCSI device](#).

Connecting an HA secondary host to an iSCSI device

You must prepare the secondary HA Host to connect to your iSCSI network.

Before you begin

You must ensure that you have configured your primary HA host. For more information, see [Connecting QRadar to an iSCSI network](#).

Procedure

- Step 1** Using SSH, log in to the secondary HA host as the root user.

Username: `root`

Password: `<password>`

- Step 2** Configure your HA secondary host to identify the iSCSI device volume:

- a Open the `initiatorname.iscsi` file for editing by typing the following command:

```
vi /etc/iscsi/initiatorname.iscsi
```

- b Edit the file with the iSCSI qualified name for your host by typing the following command:

```
Initiatorname=iqn.<yyyy-mm>.{reversed domain name}:<hostname>
```

For example: `InitiatorName=iqn.2008-11.com.q11labs:p113`

- c Save and close the file.

- Step 3** Restart the iSCSI service to open a session to the server by typing the following command:

```
service iscsi restart
```

What to do next

Perform the steps in the procedure, [Assigning and configuring iSCSI volumes for the HA secondary host](#).

Assigning and configuring iSCSI volumes for the HA secondary host

You must assign and configure the iSCSI volume for the secondary HA host.

Before you begin

Perform the steps in the procedure, [Connecting an HA secondary host to an iSCSI device](#)

Procedure

Step 1 Detect the volume on the iSCSI server by typing the following command:

```
iscsiadm -m discovery --type sendtargets --portal <IP address>:<port>
```

Where:

<IP address> is the IP address of the iSCSI external storage device.

<port> is the port number of the iSCSI device. This is an optional parameter.

Step 2 Verify the login to your iSCSI server by typing the following command:

```
iscsiadm -m node -l
```

Step 3 Identify the iSCSI device volume:

a Clear the kernel ring buffer by typing the following command:

```
dmesg -c
```

b Reload the iSCSI service by typing the following command:

```
service iscsi restart
```

c Locate the device volume by typing the following command:

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

What to do next

Perform the steps in the procedure, [Configuring the mount point for the secondary HA host](#).

Configuring the mount point for the secondary HA host

You must configure the mount point on the secondary HA host for the file system that you offboarded. For example, /store or /store/ariel.

Important: When you configure iSCSI on a secondary HA host in a HA deployment, do not mount the iSCSI volume if it is in use by the primary HA host.

Before you begin

Perform the steps in the procedure, [Assigning and configuring iSCSI volumes for the HA secondary host](#).

Procedure

Step 1 Unmount the file systems by typing the following commands:

```
umount /store/tmp
```

```
umount /store/ariel/persistent_data
```

```
umount /store
```

Step 2 Identify the UUID of the iSCSI device partition by typing the following command:

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

Where: <partition> is the iSCSI device partition. For example: `sdb1`

Step 3 Configure the secondary HA host to identify the partition on the iSCSI volume:

d Open the fstab file for editing by typing the following command:

```
vim /etc/fstab
```

e Add the iSCSI device partition UUID to the /store file system entry.

f Edit the settings for following file systems to replicate the mount points listed in the /etc/fstab file on the HA primary:

- /store
- /store/tmp
- /store/ariel/persistent_data

g Save and close the file.

Important: After you save and close the file, do not remount the file systems.

What to do next

Perform the steps in the procedure, [Configuring the secondary HA host to auto-mount the iSCSI volume](#).

Configuring the secondary HA host to auto-mount the iSCSI volume

You must configure the secondary HA host to auto-mount the iSCSI volume.

Attention: Do not reboot the secondary HA host when iSCSI auto-mount configuration is complete. This will attempt to mount the external storage device and conflict with the existing mounts on the primary HA host.

Procedure

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

```
chkconfig --add iscsi
chkconfig --level 345 iscsi on
```

Step 2 Create a symbolic link to the iscsi-mount script by typing the following command:

```
ln -s /opt/qradar/init/iscsi-mount /etc/init.d
```

Step 3 Add the iscsi-mount script to the startup by typing the following commands:

```
chkconfig --add iscsi-mount
chkconfig --level 345 iscsi-mount on
```

What to do next

Perform the steps in the procedure, [Verifying iSCSI connections](#).

Verifying iSCSI connections

You can verify that the connections between a primary HA host or secondary HA host and an iSCSI device are operational.

Procedure

Step 1 Using SSH, log in to the primary or secondary HA host as the root user.

Username: `root`

Password: `<password>`

Step 2 Test the connection to your iSCSI storage device by typing the following command:

```
ping <iSCSI_Storage_IP_Address>
```

Step 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: Port 3260 is the default port for the iSCSI storage solution.

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

```
iscsiadm -m node
```

Note: To verify that the iSCSI device is correctly configured, you must ensure that the output displayed by typing the `iscsiadm -m node` command, is the same for the primary HA host and secondary HA host.

If the following output is displayed, go to [Step 5](#).

```
iscsiadm: No records found
```

Step 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 initiator names are correctly configured by using the `/etc/iscsi/initiatornames.iscsi` file. For more information, see [iSCSI HA considerations](#).
- 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.

Note: 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 Security QRadar High Availability Guide*.

Note: The IBM Security QRadar High Availability Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

Troubleshoot iSCSI	<p>To prevent iSCSI disk and communication issues, you must connect QRadar, the iSCSI server, and your network switches to a Uninterruptable 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.</p> <p>Administrators should review the following troubleshooting information:</p> <ul style="list-style-type: none"> • Reconfigure iSCSI when restoring a failed primary HA console • Detecting Disk Errors • Unmounting and remounting the iSCSI volume
Reconfigure iSCSI when restoring a failed primary HA console	<p>In an 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, follow the instructions for configuring a secondary HA host. For more information see, Connecting an HA secondary host to an iSCSI device.</p>
Detecting Disk Errors	<p>After a power failure in a network switch you must detect disk errors.</p> <p>Procedure</p> <p>Step 1 Using SSH, log in to QRadar Console as the root user.</p> <p>Username: <code>root</code></p> <p>Password: <code><password></code></p> <p>Step 2 Type the following command:</p> <pre>touch /store/ariel/filename.txt</pre> <p>or</p> <pre>touch /store/filename.txt</pre> <p>What to do next</p> <p>If your iSCSI volume is mounted correctly and you have write permissions to the volume, the touch command creates an empty file named filename.txt on your iSCSI volume.</p> <p>If you receive a read-only error message, then see Unmounting and remounting the iSCSI volume.</p>
Unmounting and remounting the iSCSI volume	<p>If you detect a disk error, such as the file system in a read-only state, you can attempt to correct the disk error by unmounting and remounting the iSCSI volume.</p>

Procedure

Step 1 Using SSH, log in to QRadar Console as the root user.

Username: `root`

Password: `<password>`

Step 2 Stop the QRadar services by choosing one of the following options:

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

```
service hostcontext stop
service tomcat stop
service hostservices stop
service systemStabMon stop
service crond stop
```

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

```
service hostcontext stop
```

Step 3 Unmount the iSCSI volume by choosing one of the following options:

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

```
umount /store/tmp
umount /store
```

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

```
umount /store/ariel
```

Step 4 Mount the iSCSI volume by choosing one of the following options:

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

```
mount /store
mount /store/tmp
```

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

```
mount /store/ariel
```

Step 5 Test the mount points by choosing one of the following options:

- 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:

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

If you continue to receive a read-only error messages after remounting the disk, then reconfigure your iSCSI volume, see [Connecting QRadar to an iSCSI network](#).

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

Where: `<partition>` is the name of the iSCSI device partition. For example:
`sdb1`

If you do not know the drive name, remount the volume, then check the mounted volumes using the following command:

```
mount
```

Step 6 Start the QRadar services by choosing one of the following options:

- If you migrated the /store file system, type the following commands in the specified order:

```
service crond start
service systemStabMon start
service hostservices start
service tomcat start
service hostcontext start
```

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

```
service hostcontext start
```

4

FIBRE CHANNEL IMPLEMENTATION

Fibre Channel can be configured in a standard IBM Security QRadar deployment or in a High Availability (HA) environment.

When you configure a Fibre Channel device, you can migrate the QRadar data in your /store or /store/ariel file system and then mount the /store or /store/ariel file system to a partition on the Fibre Channel device.

Depending on your device configuration, you might be required to create a partition on the volume of your Fibre Channel disk. For more information, see [Creating a disk partition](#).

If you configure Fibre Channel in an HA deployment and your primary HA host fails, your Fibre Channel 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 Security QRadar High Availability Guide*.

Administrators can perform the following tasks:

- 1 Configure Fibre Channel with a standard QRadar console, see [Configure Fibre Channel in a standard QRadar deployment](#).
- 2 Configure Fibre Channel in a QRadar HA deployment, see [Configure Fibre Channel in an HA deployment](#).

Note: Fibre Channel configuration requires an advanced knowledge of the Linux operating system.

Fibre Channel considerations

Before you implement Fibre Channel, you must be aware of the configuration options and HA setup requirements.

The procedures for configuring Fibre Channel are different for a primary HA host and secondary HA host.

To configure Fibre Channel, you must ensure that the primary HA host and secondary HA host are not connected in an HA cluster. For more information about HA clustering, see the *IBM Security QRadar High Availability Guide*.

Frequently searched data must be offboarded to a faster disk. For example, more recent data or data that is used for security incident investigation. However, you must be aware that deploying high performance offboard disk storage might have a significant cost implication. Where possible, use lower performance, less expensive offboard storage for activities such as, migrating older data, archiving, or for reporting purposes.

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

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

Verifying your Emulex adapter installation

You must verify that an Emulex LPe12002 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 Host Bus adapter on your QRadar appliance. In an HA deployment, you must install an Emulex LPe12002 card on the primary and secondary HA host. For information, contact your System Administrator.

About this task

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

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

Procedure

Step 1 Using SSH, log in to your QRadar host as the root user:

Username: `root`

Password: `<password>`

Step 2 Verify that an Emulex LPe12002 card is attached by typing the following command:

```
hbacmd listhbas
```

If no result is displayed, then contact your System Administrator.

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

```
hbacmd HBAAttrib <device id>
```

Where: `<device id>` is the Port WWN that you derived in [Step 2](#).

What to do next

If you are configuring an HA deployment, then you must check that the secondary HA host is using an Emulex LPe12002 Host Bus Adapter with the same driver and firmware versions.

Configure Fibre Channel in a standard QRadar deployment

You can configure Fibre Channel in a standard deployment by using a QRadar console.

Administrators must perform the following tasks in sequence:

- 1 Verify that the correct Fibre Channel hardware is installed, see [Verifying your Emulex adapter installation](#).
- 2 Verify the connections to the Fibre Channel network, see [Verifying the Fibre Channel connections](#).
- 3 Choose one of the following options:
 - [Migrating /store to Fibre Channel](#)
 - [Migrating /store/ariel to Fibre Channel](#)
- 4 Verify that the file system is correctly mounted to Fibre Channel device partition. See [Verifying the Fibre Channel mount point](#).

Verifying the Fibre Channel connections

You must identify the disk volume on the external Fibre Channel device and if required create a partition on the volume.

Before you begin

Perform the steps in the procedure, [Verifying your Emulex adapter installation](#).

Procedure

Step 1 Using SSH, log in to your QRadar Console as the root user:

Username: `root`

Password: `<password>`

Step 2 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.

Step 3 **Optional.** If there is no partition on the Fibre Channel volume, then create a partition on the Fibre Channel device volume. For more information, see [Creating a disk partition](#).

Important: This step is only applicable to a stand-alone QRadar Console and a QRadar console that is being used as primary HA host in an HA deployment.

What to do next

Migrate your data to the Fibre Channel device. Choose one of the following options:

- [Migrating /store to Fibre Channel.](#)
- [Migrating /store/ariel to Fibre Channel.](#)

Migrating /store to Fibre Channel

You can migrate the QRadar data that is maintained in the /store file system and mount the /store file system to a Fibre Channel device.

Before you begin

Perform the steps in the procedure, [Verifying the Fibre Channel connections.](#)

Procedure

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

```
service systemStabMon stop
service hostcontext stop
service tomcat stop
service hostservices stop
service crond stop
```

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

```
umount /store/tmp
umount /store/ariel/persistent_data
umount /store
```

Step 3 Create a /store_old directory by typing the following command:

```
mkdir /store_old
```

Step 4 Determine the Universally Unique Identifier (UUID) of the device partition by typing the following command:

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

Where: <partition> is the name of the device partition. For example: `sdb1`

Note: If there is no partition on your Fibre Channel device volume, you must create a partition. For more information, see [Creating a disk partition.](#)

Step 5 Modify the fstab file:

a Edit the fstab file by typing the following command:

```
vi /etc/fstab
```

b Locate the existing mount line for the /store file system.

c Modify the existing /store file system entry to /store_old.

d Add a new mount point for the /store file system by typing one of the following lines:

- If the /store file system is ext4 then type:

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

- If the /store file system is xfs then type:

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

Where <uuid> is the UUID of the fibre channel device partition.

- e Modify the /store/tmp mount line to use the following file system options:

```
noatime,noauto,nobarrier 0 0
```

- f If /store/ariel/persistent_data is listed in the fstab file, then type the following file system options:

```
xfs inode64,logbsize=256k,noatime,noauto,nobarrier 0 0
```

- g Save and close the file.

- Step 6** Mount the new Fibre Channel /store file system by typing the following command:

```
mount /store
```

- Step 7** Mount the /store_old file system by typing the following command:

```
mount /store_old
```

- Step 8** Copy the data to the Fibre Channel partition by typing the following command:

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

- Step 9** Mount the /store/tmp file system by typing the following command:

```
mount /store/tmp
mount /store/ariel/persistent_data
```

- Step 10** Unmount the /store_old file system by typing the following command:

```
umount /store_old
```

- Step 11** Remove the /store_old mount point from the /etc/fstab file:

- a Open the /etc/fstab file for editing by typing the following command:

```
vi /etc/fstab
```

- b Remove the line for the /store_old mount point.

- c Save and close the file.

- Step 12** Start the QRadar services by typing the following commands in order:

```
service crond start
service hostservices start
service tomcat start
service hostcontext start
service systemStabMon start
```

What to do next

Perform the steps in the procedure, [Verifying the Fibre Channel mount point](#).

Migrating /store/ariel to Fibre Channel

You can migrate the QRadar data that is stored in the /store/ariel file system and then mount /store/ariel to a Fibre Channel device.

Before you begin

Perform the steps in the procedure, [Verifying the Fibre Channel connections](#).

Procedure

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

```
service systemStabMon stop
service hostcontext stop
service tomcat stop
service hostservices stop
service crond stop
```

Step 2 Create a temporary directory by typing the following command:

```
mkdir /tmp/fcdata
```

Step 3 Mount the Fibre Channel storage partition to the temporary directory by typing the following command:

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

Where: <partition> is the name of the device partition. For example: `sdb1`

Step 4 Copy the data to the Fibre Channel device by typing the following command:

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

Step 5 Unmount the Fibre Channel partition by typing the following command:

```
umount /tmp/fcdata
```

Step 6 Determine the UUID of the Fibre Channel device partition by typing the following command:

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

Where <partition> is the name of the Fibre Channel device partition. For example: `sdb1`.

Step 7 Modify the fstab file:

a Edit the fstab file by typing the following command:

```
vi /etc/fstab
```

b Add a new mount point for /store/ariel by typing the following line:

```
UUID=<uuid> /store/ariel ext4 defaults,noatime,nobarrier 1 2
```

Where: <uuid> is the value that you derived in step [Step 6](#).

c Save and close the file.

Step 8 Mount the /store/ariel file system to the Fibre Channel device partition by typing the following command:

```
mount /store/ariel
```

Step 9 Start the QRadar services by typing the following commands in sequence:

```
service crond start  
service hostservices start  
service tomcat start  
service hostcontext start  
service systemStabMon start
```

What to do next

Perform the steps in the procedure, [Verifying the Fibre Channel mount point](#).

Verifying the Fibre Channel mount point

You can verify that the file system that you migrated is correctly mounted to Fibre Channel device partition.

About this task

Note: Do not perform this task when you configure Fibre Channel on a secondary HA host.

Procedure

Step 1 Type the following command:

```
df -h
```

Step 2 Verify that the /store or /store/ariel file system is correctly mounted to the Fibre Channel device partition.

Configure Fibre Channel in an HA deployment

To use Fibre Channel storage in an HA environment, you must configure the primary HA host and the secondary HA host to use the same storage partition.

Administrators must perform the following tasks in sequence:

- 1 Verify that the correct Fibre Channel hardware is installed on your secondary HA host, see [Verifying your Emulex adapter installation](#).
- 2 Configure Fibre Channel on your primary HA host. For more information, see [Configure Fibre Channel in a standard QRadar deployment](#).
- 3 Verify the HA Fibre Channel connections. For more information, see [Verifying the HA Fibre Channel connections](#).
- 4 Configure the file system mount point for the secondary HA host, see [Configuring the mount point for the secondary HA host](#).

Verifying the HA Fibre Channel connections

You must identify the disk volume on the external Fibre Channel device and if required create a partition on the volume.

Before you begin

Perform the steps in the procedure, [Verifying your Emulex adapter installation](#).

Procedure

- Step 1** Using SSH, log in to your QRadar secondary HA host as the root user:

Username: `root`

Password: `<password>`

- Step 2** Identify the Fibre Channel volume by typing the following command:

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

Note: If multiple Fibre Channel devices are attached and you cannot identify the correct Fibre Channel volume, contact your System Administrator.

What to do next

Perform the steps in the procedure, [Configuring the mount point for the secondary HA host](#).

Configuring the mount point for the secondary HA host

You must configure the mount point on the secondary HA host for the file system that is offboarded. For example: `/store` or `/store/ariel`.

Before you begin

Perform the steps in the procedure, [Verifying the HA Fibre Channel connections](#).

Procedure

- Step 1** Derive the UUID for the Fibre Channel device partition by using the primary HA host.
- a Using SSH, log in to your QRadar primary HA host as the root user:

Username: `root`

Password: `<password>`

- b** Derive the UUID by typing the following command:

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

Where: `<partition>` is the name of the device partition. For example: `sdb1`

Step 2 Ensure that the secondary HA host can access the device partition:

- a** Update the kernel with the Fibre Channel partition data by typing the following command:

```
partprobe
```

If the following error message is displayed, go to **b**.

```
Warning: WARNING: the kernel failed to re-read the partition
table on /dev/sda (Device or resource busy). As a result, it
may not reflect all of your changes until after reboot.
```

- b** Identify the Fibre Channel device partition by typing the following command:

```
ls -l /dev/disk/by-uuid/<partition>
```

Where: `<partition>` is the value that you derived in step 1.

If no output is displayed then reboot the secondary HA host, type `reboot`.

Step 3 Unmount the `/store` and `/store/tmp` file systems, by typing the following commands:

```
umount /store/tmp
umount /store
```

Step 4 Edit the `/etc/fstab` file by typing the following command:

```
vi /etc/fstab
```

Step 5 Choose one of the following options:

- a** If you offboarded `/store`, then add the following lines to the `/etc/fstab` file:

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

Where: `<uuid>` is the value that you derived in [Step 1](#).

- b** If you offboarded `/store/ariel` add the following lines to the `/etc/fstab` file:

```
UUID=<uuid> /store/ariel ext4
defaults,noatime,noauto,nobarrier 1 2
```

Where: `<uuid>` is the value that you derived in [Step 1](#).

Step 6 Save and close the file.

What to do next

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

5

USING NFS FOR QRADAR BACKUPS

You can offboard the QRadar backup file system to an external Network File System (NFS).

You cannot use NFS for storing active data, which includes the postgres and ariel databases. If you do use NFS, it might cause database corruption or performance issues.

Important: Depending on your QRadar High Availability (HA) deployment, you might be required to change the location of your QRadar backup files and configure your NFS share with this new location.

You can migrate QRadar backups to NFS from a stand-alone QRadar console, configure a new QRadar HA deployment and migrate backups to NFS or migrate backups from an existing QRadar HA deployment.

Migrate backups from a stand-alone QRadar console

You can configure NFS for a stand-alone QRadar console.

You must perform the following tasks in sequence:

- 1 Enable the connections to NFS. See [Enabling NFS connections](#).
- 2 Configure the NFS mount for /store/backup. See [Configuring the NFS mount](#).
- 3 Migrate backup data to the NFS volume. See [Migrating existing backups to NFS](#).

Enabling NFS connections

You must enable the connections to your NFS server.

About this task

You must enable the connections to your NFS server for any of the following situations:

- To migrate /store/backup to NFS from a stand-alone QRadar console.
- To enable NFS connections on your primary and secondary HA hosts. You must enable these connections for new and existing HA deployments.

Procedure

Step 1 Using SSH, log in to QRadar as the root user:

```
Username: root
```

```
Password: <password>
```

Step 2 Edit the /etc/hosts file by typing the following command:

```
vi /etc/hosts
```

Step 3 Add your NFS server to the /etc/hosts file by typing the following line:

```
<IP address> <hostname>
```

Where:

<IP address> is the IP address of your NFS server.

<hostname> is the name of your NFS server.

Step 4 Save and close the file.

Step 5 Edit the iptables firewall by typing the following command:

```
vi /opt/qradar/conf/iptables.pre
```

Step 6 Add the following line:

```
-A INPUT -i <interface> -s <IP address> -j ACCEPT
```

Where:

<interface> is ETH0 or ETH1 if you have a dedicated NFS network.

<IP address> is the IP address of your NFS server.

Step 7 Save and close the file.

Step 8 Restart iptables by typing the following command:

```
/opt/qradar/bin/iptables_update.pl
```

Step 9 Add NFS to be part of the startup by typing the following commands:

```
cd /etc/rc3.d
chkconfig --level 3 nfs on
chkconfig --level 3 nfslock on
```

Step 10 Start NFS services by typing the following commands:

```
service nfslock start
service nfs start
```

Configuring the NFS mount

You must configure the NFS mount point for the /store/backup file system.

Before you begin

Perform the steps in the procedure, [Enabling NFS connections](#).

About this task

You must configure your NFS mounts for any of the following situations:

- If you are migrating /store/backup to NFS from a stand-alone QRadar console.

- If you are configuring an HA deployment for the first time, then you must configure an NFS mount point for /store/backup on your primary and secondary HA hosts.

Procedure

Step 1 Using SSH, log in to the QRadar as the root user:

Username: `root`

Password: `<password>`

Step 2 Open the fstab file for editing by typing the following command:

```
vi /etc/fstab
```

Step 3 Add the following line to the /etc/fstab file:

```
<hostname>:<shared_directory> /store/backup nfs soft,intr,rw 0 0
```

Where:

`<hostname>` is the name of your NFS server.

`<shared_directory>` is the path to your shared directory on the NFS server.

Note: You might need to adjust the settings for the NFS mount point to accommodate your configuration. For example: `<hostname>: <shared_directory> /store/backup nfs soft,intr,rw,noac 0 0`. For more information about common NFS mount options, type `man nfs` to view the UNIX man page for NFS.

Migrating existing backups to NFS

You can migrate existing backup files to an NFS volume.

Before you begin

Perform the steps in the procedure, [Configuring the NFS mount](#).

About this task

Do not perform this procedure on a secondary HA host. In an HA deployment, QRadar backup data is migrated from the primary HA host.

Procedure

Step 1 Using SSH, log in to the QRadar Console or primary HA host as the root user:

Username: `root`

Password: `<password>`

Step 2 Move your backup files from the existing directory to a temporary location by typing the following commands:

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

Step 3 Create a new backup directory by typing the following command:

```
mkdir /store/backup
```

Step 4 Set the permissions for the NFS volume by typing the following command:

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

Step 5 Mount the NFS volume by typing the following command:

```
mount /store/backup
```

Step 6 Verify that /store/backup is mounted by typing the following command:

```
df- h
```

Step 7 Move the backup files from the temporary location to the NFS volume by typing the following command:

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

Step 8 Remove the backup.local directory by typing the following commands:

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

Migrating backups from a new QRadar HA deployment

You can configure a new HA deployment and migrate QRadar backups to NFS.

Before you begin

The steps in this procedure explain how to configure NFS with a new HA deployment. If you already configured an HA deployment, then you must change the QRadar backup file location and configure your NFS share. For more information, see [Migrating backups from an existing QRadar HA deployment](#).

About this task

To use NFS storage in an HA environment, you must configure the primary HA host and the secondary HA host to use the same NFS share.

Procedure

- Step 1** Enable the connections to NFS on your primary and secondary HA hosts, see [Enabling NFS connections](#).
- Step 2** Configure the NFS mount point on your primary and secondary HA hosts, see [Configuring the NFS mount](#).
- Step 3** Migrate QRadar backups from the primary HA host to the NFS share, see [Migrating existing backups to NFS](#).

Migrating backups from an existing QRadar HA deployment

You can migrate QRadar backups to NFS from an existing HA deployment.

Before you begin

The steps in this procedure explain how to configure NFS with an existing HA deployment. If you have not already configured your HA deployment, and want to use NFS with HA, see [Migrating backups from a new QRadar HA deployment](#).

About this task

To use NFS storage in an HA environment, you must configure the primary HA host and the secondary HA host with the same NFS configurations.

Procedure

- Step 1** Enable the connections to NFS on your primary and secondary HA hosts, see [Enabling NFS connections](#).
- Step 2** Configure a new file location for your QRadar backups, see [Configuring a new QRadar backup location](#).
- Step 3** Change the file location of your QRadar backups, see [Changing the QRadar backup file location](#).
- Step 4** Configure the NFS mount point for new backup file location on your secondary HA host, see [Configuring the NFS mount](#).

Configuring a new QRadar backup location

You can create a new file location to store your QRadar backups.

About this task

If you have an existing HA cluster, then you must change the QRadar backup location on your primary HA host.

Procedure

- Step 1** Using SSH, log in to the QRadar Console as the root user:
Username: `root`
Password: `<password>`
- Step 2** Create a new file location to store your QRadar backups. For example:
`/nfs/backup`
Attention: Do not create your new backup location under the `/store` file system.
- Step 3** Open the `fstab` file for editing by typing the following command:
`vi /etc/fstab`
- Step 4** Add the following line to the `/etc/fstab` file:
`<hostname>:<shared_directory> <backup location> nfs soft,intr,rw 0 0`
Where:
`<hostname>` is the name of your NFS server.
`<shared_directory>` is the path to your shared directory on the NFS server.
`<backup location>` is the backup location that you created in [Step 2](#).
- Step 5** Save and close the file.
- Step 6** Mount the new backup file location to the NFS share by typing the following command:
`mount <backup location>`

Step 7 Copy the existing backup data to the NFS share by typing the following command:

```
mv /store/backup/* <backup location>
```

What to do next

Perform the steps in the procedure, [Changing the QRadar backup file location](#).

Changing the QRadar backup file location

You can change the location for QRadar backup files.

About this task

If you have an existing HA cluster and you want to configure an NFS server on your primary and secondary HA hosts, then you must change the location of your QRadar backups.

Procedure

- Step 1** Log in to the QRadar Console.
- Step 2** Click the **Admin** tab.
- Step 3** On the navigation menu, click **System Configuration**.
- Step 4** Click the **Backup and Recovery** icon.
- Step 5** On the toolbar, click **Configure**.
- Step 6** In the **Backup Repository Path** field, type the location where you want to store your QRadar backup files. For example, `/nfs/backup`.
- Step 7** Click **Save**.
- Step 8** Close the Backup Archives window.
- Step 9** On the **Admin** tab menu, click **Deploy Changes**.

For more information about QRadar backup and recovery settings, see the *IBM Security QRadar SIEM Administration Guide* or the *IBM Security QRadar Log Manager Administration Guide*.

Configuring a mount point for a secondary HA host

You must configure an NFS mount point, on your existing secondary HA host, for the alternative QRadar backup file location.

Before you begin

Perform the steps in the procedures.

- Step 1** Using SSH, log in to the QRadar secondary HA host as the root user:
Username: `root`
Password: `<password>`
- Step 2** Create a backup file location that matches the backup file location on your primary HA host. For more information, see [Configuring a new QRadar backup location](#).
For example: `/nfs/backup`

Attention: Do not create your new backup location under the `/store` file system.

Step 3 Open the fstab file for editing by typing the following command:

```
vi /etc/fstab
```

Step 4 Add the following line to the /etc/fstab file:

```
<hostname>:<shared_directory> <backup location> nfs soft,intr,rw  
0 0
```

Where:

<hostname> is the name of your NFS server.

<shared_directory> is the path to your shared directory on the NFS server.

<backup location> is the backup file location that you created in [Step 2](#)

Note: You might need to adjust the settings for the NFS mount point to accommodate your configuration. For example: <hostname>: <shared_directory> <backup location> nfs soft,intr,rw,noac 0 0.

Step 5 Save and close the file.

Step 6 Mount the new QRadar backup file location by typing the following command:

```
mount <backup location>
```

Where: <backup location> is the new QRadar backup file location. For more information, see [Configuring a new QRadar backup location](#).

A

NOTICES AND TRADEMARKS

What's in this appendix:

- [Notices](#)
- [Trademarks](#)

This section describes some important notices, trademarks, and compliance information.

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/region 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*

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law:

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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply 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 Corporation
170 Tracer Lane,
Waltham MA 02451, USA*

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All 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 the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

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.

The following terms are trademarks or registered trademarks of other companies:

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Privacy policy considerations

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

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

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

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

INDEX

Symbols

/store
 migrating to iSCSI 19
 migrating using fibre channel 32
 when to consider migrating 6
/store/ariel
 migrating to iSCSI 18

A

auto-mount
 iSCSI volumes 21

B

before you begin
 fibre channel 30

C

configure iSCSI for HA
 before you begin 15
conventions 3

E

external storage
 limitations 8
 when to consider 5
external storage options
 fibre channel 7
 iSCSI 7
 NFS 8

F

fibre channel
 before you begin 30
 connecting qradar 31
 in a standard deployment 31
 migrating a subdirectory of /store 34
 more information 7
 using in an HA environment 36
 verifying mount points 35

H

high availability
 before you begin 15
 using iSCSI 22

I

intended audience 3
iSCSI
 assigning volumes 17
 configuring volumes 17
 more information 7
 troubleshooting 26
 usage in a standard QRadar deployment 16
iSCSI connections
 verifying 25
iSCSI HA
 assigning iSCSI volumes 23
 auto-mounting iSCSI volumes 24
 configuring secondary host mount points 23
 connecting a secondary host to iSCSI 22
iSCSI hA
 configuring iSCSI volumes 23
iSCSI network
 connecting QRadar 16
iSCSI network
 connecting a secondary HA host 22
iSCSI volumes
 assigning and configuring 17
 auto-mounting 21
 auto-mounting using iSCSI HA 24
iSCSI with HA
 using iSCSI with HA 22

L

limitations of external storage 8

M

migrating /store
 using fibre channel 32
migrating the /store file system
 when to consider 6

N

NFS
 more information 8

Q

QRadar standard deployment
 using iSCSI 16

S

- secondary HA host
 - assigning and configuring volumes 23
 - auto-mounting iSCSI volumes 24
 - configuring mount points 23
 - connecting to the iSCSI network 22

T

- troubleshooting
 - detecting disk errors 26
 - iSCSI 26
 - mounting iSCSI volumes 26
 - reconfiguring a failed primary host 26
 - unmounting iSCSI volumes 26
- troubleshooting iSCSI
 - verifying connections to iSCSI 25

W

- when to migrate the /store file system 6