

IBM Security QRadar  
Version 7.1.0 (MR1)

## *Offboard Storage Guide*



**Note:** Before using this information and the product that it supports, read the information in [Notices and Trademarks](#) on [page 43](#).

# CONTENTS

---

## ABOUT THIS GUIDE

Intended Audience . . . . .	3
Conventions . . . . .	3
Technical Documentation . . . . .	4
Contacting Customer Support . . . . .	4

---

## 1 OVERVIEW

When to Consider External Storage . . . . .	5
Types of Stored QRadar Data . . . . .	6
Migrating the /store File System . . . . .	6
Migrating the /store/ariel File System . . . . .	7
External Storage Options . . . . .	7
Fibre Channel . . . . .	7
iSCSI . . . . .	8
NFS (Network File System) . . . . .	8
Limitations of Using External Storage . . . . .	8
External Storage Considerations in an HA Environment . . . . .	9

---

## 2 CONFIGURING ISCSI

Before you Begin . . . . .	11
Configuring iSCSI in a Standard QRadar Deployment . . . . .	12
Connect QRadar to the iSCSI Network . . . . .	12
Assign and Configure the iSCSI Volumes . . . . .	13
Migrating Data to the iSCSI Storage Solution . . . . .	14
Configuring the System to Auto-mount the iSCSI Volume . . . . .	18
Configuring iSCSI in an HA Environment . . . . .	18
Connect the HA Secondary Host to the iSCSI Device . . . . .	19
Assign and Configure iSCSI Volumes for the HA Secondary Host . . . . .	19
Configure the Mount Point for the HA Secondary Host . . . . .	20
Configure the HA Secondary Host to Auto-Mount the iSCSI Volume . . . . .	21
Verifying iSCSI Connections . . . . .	22
Connect the Primary and Secondary Host in the QRadar User Interface . . . . .	23
Troubleshooting . . . . .	23
Configuring iSCSI When Restoring a Failed Primary HA Console . . . . .	23
Detecting Disk Errors . . . . .	24
Unmounting and Remounting the iSCSI Volume . . . . .	24

---

<b>3</b>	<b>CONFIGURING FIBRE CHANNEL</b>	
	Best Practices . . . . .	27
	Fibre Channel Performance . . . . .	27
	Fibre Channel Archiving . . . . .	28
	Using Fibre Channel Volumes . . . . .	28
	Before You Begin . . . . .	28
	Fibre Channel Configuration Types . . . . .	30
	Configuring Fibre Channel in a Standard Deployment . . . . .	30
	Configuring Fibre Channel HA . . . . .	30
	Configuring Fibre Channel . . . . .	31
	Preparing QRadar to Connect to Fibre Channel Network . . . . .	31
	Migrating /store to the Fibre Channel Solution . . . . .	33
	Migrating a subdirectory of /store to the Fibre Channel Storage Solution . . . . .	35
	Verifying the Fibre Channel Mount . . . . .	36
<b>4</b>	<b>USING NFS FOR QRADAR BACKUPS</b>	
	NFS Considerations . . . . .	39
	Implementing NFS for Backups . . . . .	39
<b>A</b>	<b>NOTICES AND TRADEMARKS</b>	
	Notices . . . . .	43
	Trademarks . . . . .	45

---

**INDEX**

# ABOUT THIS GUIDE

The *IBM Security QRadar Offboard Storage Guide* 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 HA Guide are not applicable to IBM Security QRadar Network Anomaly Detection.

---

---

**Intended Audience** This guide is intended for all QRadar users responsible for investigating and managing network security. This guide assumes that you have QRadar access and a knowledge of your corporate network and networking technologies.

---

**Conventions** The following conventions are used throughout this guide:

- ▶ Indicates that the procedure contains a single instruction.

**NOTE** Indicates that the information provided is supplemental to the associated feature or instruction.

---



**CAUTION**

---

*Indicates that the information is critical. A caution alerts you to potential loss of data or potential damage to an application, system, device, or network.*

---



**WARNING**

---

*Indicates that the information is critical. A warning alerts you to potential dangers, threats, or potential personal injury. Read any and all warnings carefully before proceeding.*

---

---

**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=swg21614644).  
(<http://www.ibm.com/support/docview.wss?rs=0&uid=swg21614644>)

---

**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>)

# 1

## OVERVIEW

This document provides information on how to initially configure iSCSI, Fibre Channel, and NFS external storage options using QRadar 7.1.

If you are upgrading your IBM Security QRadar deployment, and are required to reconfigure the connections to an existing external storage device, see the *Reconfiguring Offboard Storage During a QRadar Upgrade, Technical Note*.

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.

If you need to configure an external storage solution using QRadar 7.0, you should use the following documentation:

- *Configuring iSCSI* technical note.
- *Configuring Fibre Channel* technical note.
- *Implementing NFS* technical note.

This section includes the following topics:

- [When to Consider External Storage](#)
- [Types of Stored QRadar Data](#)
- [Limitations of Using External Storage](#)
- [External Storage Considerations in an HA Environment](#)

---

### When to Consider External Storage

The local disk on your QRadar appliance is significantly faster than external storage and currently supports up to 16TB of data. For this reason, we recommend local storage as a preferred option. In QRadar deployments where larger storage capacity is required, multiple appliances are recommended.

When multiple appliances are not feasible, or an existing deployment can increase capacity by utilizing available external storage, you need to consider the following before migrating your data to external storage:

- What hardware configuration are you using in your QRadar deployment?
- Do you have expertise in managing external storage devices?

- What existing infrastructure investments do you have within your organization?
- What is your policy on data retention periods?
- Do your configured retention settings exceed the capacity of existing storage?
- Do you require access, via the QRadar user interface, to data that has been migrated to offboard storage?
- Do you need to expand the retention of existing deployed appliances?
- Do you require increased fault tolerance and disaster recovery capabilities?

---

## Types of Stored QRadar Data

QRadar data is located in the /store file system and its subdirectories. An offboard storage solution can be used to migrate the entire /store file system or specific subdirectories. Each option has a different impact on QRadar performance. For more information, see [External Storage Options](#). You can offboard the following QRadar data:

- Postgres meta data 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).

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 16TB file system limit currently supported by QRadar.

If you need 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. This provides up to 32TB of storage for either Log or Network Activity data.

For additional information on expanding your storage capabilities, contact your QRadar technical resource or Customer Support.

This section includes the following topics:

- [Migrating the /store File System](#)
- [Migrating the /store/ariel File System](#)

## Migrating the /store File System

It is common to migrate the /store file system when you need to increase the fault tolerance levels in your QRadar deployment. Migrating this file system to your external device also provides an alternative resolution to implementing an HA (High Availability) environment. For more information on HA, see the *IBM Security QRadar HA Guide*.



**NOTE** The IBM Security QRadar HA Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

---

Migrating the /store file system to an offboard device can negatively affect QRadar performance. After migration, all data I/O to the this file system is no longer performed on the local disk. Before migrating your data, consider the following:

- Maintain your Log and Network Activity searches on your local disk, by mounting the directory, /store/ariel/persistent\_data/, on the unused /store file partition.

**NOTE** Searches marked as saved are also located in /store/ariel/persistent\_data/ directory. If you experience a local disk failure, these searches are not saved. For further assistance, contact Customer Support.

---

### **Migrating the /store/ariel File System**

The /store/ariel directory is the most commonly offboarded file system. By migrating this file system, you can:

- Migrate collected log and network activity data to external storage.
- Ensure the local disk remains used for the postgresql database and temporary search results.

---

## **External Storage Options**

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

This section includes the following topics:

- [Fibre Channel](#)
- [iSCSI](#)
- [NFS \(Network File System\)](#)

### **Fibre Channel**

Fibre Channel provides the fastest offboard performance, with Storage Area Network (SAN) speeds of between 200MBps and 3200MBps, depending on your network architecture. However, while the SAN architecture provides high data rates, performance can 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 on configuring Fibre Channel with QRadar, see [Configuring Fibre Channel](#)

**iSCSI** iSCSI utilizes 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.

However, because iSCSI utilizes existing network and host management interfaces, network capacity is shared between external storage access and management interface I/O. In this situation, it is common to configure a secondary network interface on a separate storage network.

---

**NOTE**

Using a dedicated interface, you are limited to 1Gbps of network capacity and more commonly will experience only 200-400Mbps. Your iSCSI storage device may only be capable of providing 25-50MBps I/O performance.

---

For more information on configuring iSCSI with QRadar, see [Configuring iSCSI](#).

**NFS (Network File System)**

While QRadar supports NFS for external storage, we recommend that you do not use NFS for storing active data. If `/store` is mounted to an NFS solution, postgres data can be corrupted. If `/store/ariel` is mounted to NFS, QRadar will experience performance issues.

NFS is more commonly used for daily configuration and data backups, since these tasks are performed during off-peak times, involve batch file writes, and a limited volume of file I/O.

NFS storage is limited to performance levels of approximately 20MB-50MBps, since it runs over existing management ethernet network. In addition, the NFS protocol incurs additional overhead for file access, locking, and network permissions. This can be remediated by using a dedicated network interface.

---

**NOTE**

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

---

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

---

**Limitations of Using External Storage**

When considering an external storage solution you should consider the following limitations:

- QRadar does not support multiple systems accessing the same block device. If you are configuring iSCSI in an HA environment, you should not attempt to 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 [Configuring iSCSI in an HA Environment](#)

- The performance of local storage in QRadar appliances is significantly faster than external storage. Therefore, local storage is always the preferred, recommended option.

An external storage device should be capable of providing consistent read and write capacity of 100-200MBps. When this is not available, the following can occur:

- Data write performance can be effected.
- The performance of user interface searches can be impacted.
- If capacity drops further, the processing pipeline can become blocked and QRadar might display messages that Events and Flows are being dropped.

---

### External Storage Considerations in an HA Environment

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

If you choose to offboard /store/ariel, and maintain /store on your local QRadar disk, the /store file system will be replicated to the secondary HA device during a primary failure, using DRBD. By default, when your environment is configured for HA, DRBD is enabled.

### NOTE

---

If you migrate QRadar data to an external storage device in a HA environment, the directories you choose to migrate will impact your HA configuration. For more information, see the *IBM Security QRadar HA Guide*. The IBM Security QRadar HA Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

---



# 2

## CONFIGURING ISCSI

iSCSI can be configured in a standard IBM Security QRadar deployment or in a High Availability (HA) environment. If you are using HA, iSCSI can be used to maintain replicated data on a secondary host. Using this approach, you must configure your secondary host with the same external iSCSI device as the primary host. For more information, see [Configuring iSCSI in an HA Environment](#).

Your network configuration may differ, however, this section assumes that your management interface is eth0 and your iSCSI interface is eth1.

### NOTE

---

The procedures described below assume an advanced knowledge of the Linux operating system. For assistance, please contact Customer Support.

---

This section includes the following topics:

- [Before you Begin](#)
- [Configuring iSCSI in a Standard QRadar Deployment](#)
- [Configuring iSCSI in an HA Environment](#)
- [Troubleshooting](#)

---

### Before you Begin

Before you configure iSCSI external storage in an HA environment you should review the following information:

- Ensure that you use a different initiatorname on both the primary host and the HA secondary host. Your iSCSI device should be configured so that each initiatorname can access the same volume on the iSCSI device.

The initiatorname is stored in the `/etc/iscsi/initiatorname.iscsi` file and is used to identify the iSCSI device volume where the `/store` or `/store/ariel` file system should be mounted.

- You cannot configure iSCSI if you have already connected the QRadar primary and secondary host in an HA pairing. To configure iSCSI using HA, ensure the primary and secondary hosts are not paired. For more information on HA pairing, see the *IBM Security QRadar HA Guide*.

**NOTE**


---

The IBM Security QRadar HA Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

---

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

---

## Configuring iSCSI in a Standard QRadar Deployment

Before you can migrate your QRadar data using iSCSI, you must configure QRadar to connect to the iSCSI device and assign and configure the iSCSI volumes.

This section includes the following topics:

- [Connect QRadar to the iSCSI Network](#)
- [Assign and Configure the iSCSI Volumes](#)
- [Migrating Data to the iSCSI Storage Solution](#)
- [Configuring the System to Auto-mount the iSCSI Volume](#)

### Connect QRadar to the iSCSI Network

To prepare QRadar to connect to your iSCSI network:

- Step 1** Optional. From the **Admin** tab, configure a secondary network interface with a private IP address to connect to the iSCSI Storage Area Network (SAN). This is optional, but we recommend that you configure your SAN using this method to improve performance.

**NOTE**


---

You will require network interface address information from your SAN network manager. For more information on configuring a network interface, see your Administration Guide.

---

- Step 2** Using SSH, log in to the QRadar Console as the root user.

Username: `root`

Password: `<password>`

- Step 3** Configure your system 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. 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 4** Open a session to the iSCSI server by typing the following command:

```
service iscsi restart
```

You are now ready to assign and configure the iSCSI volumes. See [Assign and Configure the iSCSI Volumes](#).

### Assign and Configure the iSCSI Volumes

To assign and configure your iSCSI volumes:

**Step 1** Detect volumes 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.

The output should resemble the following:

```
172.16.151.142:3260,1 iqn.2008-10.lab.q1labs:iscsiVol1
```

**Step 2** Verify that the login to the iSCSI server is functional by typing the following command:

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

The output should resemble the following:

```
Logging in to [iface: default, target:
iqn.2008-10.lab.q1labs:iscsiVol, portal: 172.16.151.142,3260]
Login to [iface: default, target:
iqn.2008-10.lab.q1labs:iscsiVol, portal: 172.16.151.142,3260]:
successful
```

**Step 3** Determine the iSCSI device name:

**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"
```

The output should resemble the following:

```
sd 4:0:0:0: [sdb] Attached SCSI disk
```

Where [sdb] is the volume on the iSCSI device.

**Step 4** Reformat the iSCSI device partition, if it has not previously been used:

**CAUTION**


---

*If the partition on the volume has been used before and you need to retain the data in the volume, then you cannot create partitions or reformat the partitions in the volume.*

---

- a Optional. Create a partition.

For information about creating a partition, see your Linux documentation.

- b Reformat the partition by typing the following command:

```
mkfs.ext4 /dev/<device name>
```

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

**NOTE**


---

You can create one or more partitions on the iSCSI volume and mount them separately. If the new volume is larger than 2TB, create a GUID Partition Table (GPT). Using GPT, the new volume is limited to 16TB. If you are using MSDOS partitioning, you are limited to a single 2TB partition.

---

You are now ready to migrate your data to the iSCSI external storage solution. See [Migrating Data to the iSCSI Storage Solution](#).

### Migrating Data to the iSCSI Storage Solution

You can choose which directory level you want to migrate to the external iSCSI device: `/store` or `/store/ariel`. To retain optimal system performance, we recommend that you migrate `/store/ariel`.

This section includes the following topics:

- [Migrating /store/ariel to the iSCSI Storage Solution](#)
- [Migrating /store to the iSCSI Storage Solution](#)

#### Migrating /store/ariel to the iSCSI Storage Solution

To migrate the `/store/ariel` file system to the iSCSI storage solution:

- Step 1** Stop the `hostcontext` service by typing the following command:

```
service hostcontext 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/<device name>
```

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



The output should resemble the following:

```
/dev/sdb1: UUID="89ec181b-dcd1-4698-b1ae-9f1b1b044f62"
```

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

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

```
vi /etc/fstab
```

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

```
UUID=<uuid> /store/ariel <file system>
noatime,noauto,noBarrier 0 0
```

Where:

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

<file system> is the version you used to format the file system.

For example: `ext4`.

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
```

The output should resemble the following:

```
Filesystem              Size  Used Avail Use% Mounted on
/dev/sdb1                20G  172M   19G   1% /store/ariel
```

**Step 8** Move the data from the local volume 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** Restart the Hostcontext service by typing the following command:

```
service hostcontext restart
```

#### NOTE

---

For most situations, you only need to mount a single /store/ariel on your iSCSI storage solution. However, if you need a different configuration for your iSCSI mount points, contact Customer Support.

---

You are now ready to configure the system to automatically mount the iSCSI volume. See [Configuring the System to Auto-mount the iSCSI Volume](#).

### Migrating /store to the iSCSI Storage Solution

To migrate the /store file system to the iSCSI storage solution:

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

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

**Step 2** Unmount /store/tmp by typing the following command:

```
umount /store/tmp
```

**Step 3** Unmount the existing /store directory by typing the following command:

```
umount /store
```

**Step 4** Create the /store\_old directory by typing the following command:

```
mkdir /store_old
```

**Step 5** Verify the UUID of the iSCSI device partition by typing the following command:

```
blkid /dev/<device name>
```

Where **<device name>** is the name of the iSCSI device including the partition number. For example: **sdb1**

The output should resemble the following:

```
/dev/sdb1: UUID="89ec181b-dcd1-4698-b1ae-9f1b1b044f62"
```

**Step 6** Modify the /store and /store/tmp mount point options using the fstab file:

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

```
vi /etc/fstab
```

b Locate the existing /store mount point:

```
UUID=<uuid> /store <file system> defaults,noatime,nobarrier 1
2
```

c Modify the line by typing the following:

```
UUID=<uuid> /store_old <file system>
defaults,noatime,nobarrier 1 2
```

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

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

Where:

**<uuid>** is the value derived in [Step 5](#).

Where: **<file system>** is the version you used to format the file system.

For example: **ext4**.

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

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

f Save and close the file.

**Step 7** Mount /store to the iSCSI device partition by typing the following command:

```
mount /store
```

**Step 8** Mount /store\_old to the local disk by typing the following command:

```
mount /store_old
```

**Step 9** Move the data from the local /store\_old file system to the iSCSI device by typing the following command:

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

#### NOTE

---

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.

---

**Step 10** Re-mount /store/tmp by typing the following command:

```
mount /store/tmp
```

**Step 11** Unmount /store\_old by typing the following command:

```
umount /store_old
```

**Step 12** 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 13** Restart services by typing the following commands in the specified order:

```
service crond restart
service systemStabMon restart
service hostservices restart
service tomcat restart
service hostcontext restart
```

#### NOTE

---

For most situations, you only need to mount a single /store on your iSCSI storage solution. However, if you need a different configuration for your iSCSI mount points, contact Customer Support.

---

You are now ready to configure the system to automatically mount the iSCSI volume. See [Configuring the System to Auto-mount the iSCSI Volume](#).

### Configuring the System to Auto-mount the iSCSI Volume

To configure the system to auto-mount the iSCSI volume:

**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 to your file 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
```

- If you migrated /store the output should resemble the following:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sdb1	20G	1.3G	18G	7%	/store

- If you migrated /store/ariel the output should resemble the following:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sdb1	20G	1.3G	18G	7%	/store/ariel

The /store or /store/ariel file system has been successfully migrated to the external iSCSI storage device.

---

### Configuring iSCSI in an HA Environment

In an HA environment, the primary host and secondary host both use iSCSI shared external storage. Therefore, if you are configuring HA, you must first configure iSCSI with your primary QRadar host and migrate either /store or /store/ariel data to the iSCSI external device.

#### NOTE

These steps must be performed before connecting the HA secondary host to the primary host using the QRadar user interface.

---

This section includes the following topics:

- [Connect the HA Secondary Host to the iSCSI Device](#)
- [Assign and Configure iSCSI Volumes for the HA Secondary Host](#)

- [Configure the Mount Point for the HA Secondary Host](#)
- [Configure the HA Secondary Host to Auto-Mount the iSCSI Volume](#)
- [Connect the Primary and Secondary Host in the QRadar User Interface](#)
- [Verifying iSCSI Connections](#)

### Connect the HA Secondary Host to the iSCSI Device

Prepare the HA Secondary Host to connect to your iSCSI network:

- Step 1** Optional. Configure a secondary network interface with a private IP address to connect to the iSCSI SAN. This is optional, but we recommend that you configure your SAN using this method to improve performance.

#### NOTE

---

You will require network interface address information from your SAN network manager. For more information on configuring a network interface, see your Administration Guide.

---

- Step 2** Using SSH, log in to the HA Secondary Host as the root user.

Username: `root`

Password: `<password>`

- Step 3** 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. 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 4** Restart the iSCSI service to open a session to the server by typing the following command:

```
service iscsi restart
```

You are now ready to assign and configure the iSCSI volumes. See [Assign and Configure iSCSI Volumes for the HA Secondary Host](#).

### Assign and Configure iSCSI Volumes for the HA Secondary Host

To assign and configure iSCSI volumes for the HA Secondary Host:

- Step 1** Detect volumes 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.

The output should resemble the following:

```
172.16.151.142:3260,1 iqn.2008-10.lab.qllabs:iscsiVol1
```

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

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

The output should resemble the following:

```
Logging in to [iface: default, target:
iqn.2008-10.lab.qllabs:iscsiVol, portal: 172.16.151.142,3260]
Login to [iface: default, target:
iqn.2008-10.lab.qllabs:iscsiVol, portal: 172.16.151.142,3260]:
successful
```

**Step 3** Determine the iSCSI device name:

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

```
dmesg -c
```

b To reload the iSCSI service, type the following command:

```
service iscsi restart
```

c To locate the device name, type the following command:

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

The output should resemble:

```
sd 4:0:0:0: [sdb] Attached SCSI disk
```

Where [sdb] is the volume on the device.

You are now ready to configure the /store mount point for the HA secondary host, see [Configure the Mount Point for the HA Secondary Host](#).

### Configure the Mount Point for the HA Secondary Host

To configure the mount point for the HA secondary host:

#### NOTE

---

When configuring iSCSI on a HA Secondary Host in a HA deployment, do not mount the iSCSI volume if it is in use by the QRadar Primary Host.

---

**Step 1** Verify the UUID of the iSCSI device partition by typing the following command:

```
blkid /dev/<device name>
```

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

The output should resemble the following:

```
/dev/sdb1: UUID="89ec181b-dcd1-4698-b1ae-9f1b1b044f62"
```

**Step 2** Configure your host to identify the correct partition on the iSCSI volume:

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

```
vi /etc/fstab
```

b Edit the mount point for the data you have migrated by typing the following line:

```
UUID=<uuid> <directory> <file system>
noatime,noauto,nobarrier 0 0
```

Where:

<uuid> is the value derived in [Step 1](#).

<directory> is /store or /store/ariel.

<file system> is the version you used to format the file system.

For example: `ext4`.

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

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

d Save and close the file.

You are now ready to Configure the HA Secondary Host to auto-mount the iSCSI volume, see [Configure the HA Secondary Host to Auto-Mount the iSCSI Volume](#).

### Configure the HA Secondary Host to Auto-Mount the iSCSI Volume

To configure the HA Secondary Host to auto-mount the iSCSI volume:



#### CAUTION

---

*Do not reboot the HA secondary 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 QRadar primary host.*

---

**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
```

You are now ready to verify the connections to your iSCSI device from both the primary host and secondary HA host. See [Verifying iSCSI Connections](#).

### Verifying iSCSI Connections

To verify that the connection to the iSCSI device is functional:

**Step 1** Using SSH, log in to the QRadar Primary 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 the connection to the iSCSI device by typing the following command:

```
iscsiadm -m node
```

The output should resemble the following:

```
172.16.90.45:3260,1
iqn.2003-10.com.lefthandnetworks:lefthandmgmtgroup1:27:iscsi
```

#### NOTE

---

To verify that iSCSI is correctly configured, you must ensure that the output displayed by typing the `iscsiadm -m node` command, is the same for both the Primary 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, 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 occurring in your iSCSI storage configuration.
- Ensure that the iSCSI initiator names are correctly configured using the `/etc/iscsi/initiatornames.iscsi` file. For more information, see [Before you Begin](#).



- If you do not locate errors in the error log, and your iSCSI connections remain disabled, you should contact your Network Administrator to confirm iSCSI server availability or network configuration changes.

**NOTE** If your network configuration has changed, you must reconfigure your iSCSI connections.

**Step 6** Using SSH, log in to the QRadar Secondary HA Host as the root user.

Username: `root`

Password: `<password>`

**Step 7** Repeat **Step 2** through **Step 5**.

You are now ready to connect your primary QRadar host to your HA secondary host, see [Connect the Primary and Secondary Host in the QRadar User Interface](#).

**Connect the Primary and Secondary Host in the QRadar User Interface**

To establish an HA cluster, you must connect the QRadar primary host with the HA secondary host using the QRadar user interface. For more information about creating an HA cluster, see the *IBM Security QRadar HA Guide*.

**NOTE** The IBM Security QRadar HA Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

**Troubleshooting**

To prevent iSCSI disk and communication issues, we recommend that you connect the QRadar, iSCSI server, and network switches to a Uninterruptable Power Supply (UPS). Power failure in a network switch may result in your iSCSI volume reporting disk errors or remaining in a read-only state.

This section includes the following topics:

- [Configuring iSCSI When Restoring a Failed Primary HA Console](#)
- [Detecting Disk Errors](#)
- [Unmounting and Remounting the iSCSI Volume](#)

**Configuring iSCSI When Restoring a Failed Primary HA Console**

In an HA environment, if your primary host fails, you must restore your iSCSI configuration to the primary host. In this event, your `/store` or `/store/ariel` data has already been migrated to the iSCSI shared external storage device. Therefore, to restore the primary host iSCSI configuration, follow the instructions for configuring an HA secondary host. For more information see, [Connect the HA Secondary Host to the iSCSI Device](#).

**Detecting Disk Errors** After Power failure in a network switch, we recommend that you perform the following test to detect disk errors:

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

Username: `root`

Password: `<password>`

**Step 2** Type the following command:

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

or

```
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 named filename.txt on your iSCSI volume.

If you receive a read-only error message, see [Unmounting and Remounting the iSCSI Volume](#).

### Unmounting and Remounting the iSCSI Volume

If you detected 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:

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

Username: `root`

Password: `<password>`

**Step 2** To stop the services, choose one of the following options:

- If you migrated the /store file system to the iSCSI storage solution, 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 /store/ariel to the iSCSI storage solution, type the following command:

```
service hostcontext stop
```

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

- If you migrated /store to the iSCSI storage solution, type the following commands:

```
umount /store/tmp
```

```
umount /store
```

- If you migrated /store/ariel to the iSCSI storage solution, type the following command:

```
umount /store/ariel
```

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

- If you migrated /store to the iSCSI storage solution, type the following commands:

```
mount /store
```

```
mount /store/tmp
```

- If you migrated /store/ariel to the iSCSI storage solution, type the following command:

```
mount /store/ariel
```

**Step 5** To test the mount points, choose one of the following options:

- If you migrated /store to the iSCSI storage solution, type the following command:

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

- If you migrated /store/ariel to the iSCSI storage solution, type the following command:

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

If you continue to receive a read-only error messages after remounting the disk, reconfigure your iSCSI volume, see [Configuring iSCSI in a Standard QRadar Deployment](#).

Alternatively, you can unmount the file system again and run a manual file system check with the following command: `fsck /dev/<device name>`. Where `<device name>` is the name of the iSCSI volume including the partition number. 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** To start the services, choose one of the following options:

- If you migrated /store to the iSCSI storage solution, 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 /store/ariel to the iSCSI storage solution, type the following command:

```
service hostcontext start
```



# 3

## CONFIGURING FIBRE CHANNEL

Fibre Channel can be configured in a standard IBM Security QRadar deployment or in a High Availability (HA) environment. QRadar on board disk systems can support read write performance at rates of 200MB to 300MB per second. Using Fibre Channel, similar performance can be achieved provided your disk storage system and volume is correctly designed and configured.

If you are unable to achieve the same levels of performance using Fibre Channel this will effect the speed with which data can be stored and searched. In this event, you should consider onboard storage or an alternative external storage solution. For more information, see [Limitations of Using External Storage](#).

The section includes the following topics:

- [Best Practices](#)
- [Before You Begin](#)
- [Fibre Channel Configuration Types](#)
- [Configuring Fibre Channel](#)

---

### Best Practices

If you are using Fibre Channel in a multi-appliance deployment, there are a number of best practices you should consider before configuring Fibre Channel.

This section includes the following topics:

- [Fibre Channel Performance](#)
- [Fibre Channel Archiving](#)
- [Using Fibre Channel Volumes](#)

### Fibre Channel Performance

We recommend that data which is searched more frequently, is offboarded to a faster disk. For example, more recent data or data that is used for security incident investigation. However, you should be aware that deploying high performance offboard disk storage may have a significant cost implication.

Where possible, you should use lower performing, less expensive offboard storage for activities such as, migrated older data, archiving, or for reporting purposes.

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

**Using Fibre Channel Volumes** For QRadar deployments which use multiple appliances, you should ensure that each appliance is configured to use a separate Fibre Channel volume. Failure to do this can result in two devices attempting to mount the same block device, which can corrupt the block device file system.

---

## Before You Begin

To configure Fibre Channel, we recommend that you install an Emulex LPe12002 Host Bus Adapter card, running firmware version 1.10A5 (U3D1.10A5) sli-3.

The following QRadar processors and appliances are compatible with the Emulex LPe12002 Host Bus Adapter card:

- QRadar 2100
- QRadar 3100
- QRadar 1601
- QRadar 1701
- QRadar 1801

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

Username: `root`

Password: `<password>`

**Step 2** To verify that an Emulex LPe12002 Host Bus Adapter card is attached, type the following command:

**hbacmd listhbas**

The output might resemble the following:

```
Manageable HBA List
Port WWN      : 10:00:00:00:c9:d0:92:38
Node WWN      : 20:00:00:00:c9:d0:92:38
Fabric Name   : 00:00:00:00:00:00:00:00
Flags         : 8000f100
Host Name     : angel-primary.qllabs.inc
Mfg           : Emulex Corporation
Serial No.    : BT02461636
Port Number   : 0
Mode          : Initiator
PCI Function  : 0
Port Type     : FC
Model         : LPe12002-M8
Port WWN      : 10:00:00:00:c9:d0:92:39
Node WWN      : 20:00:00:00:c9:d0:92:39
Fabric Name   : 00:00:00:00:00:00:00:00
Flags         : 8000f100
```

```

Host Name      : angel-primary.qllabs.inc
Mfg           : Emulex Corporation
Serial No.    : BT02461636
Port Number   : 1
Mode          : Initiator
PCI Function  : 1
Port Type     : FC
Model        : LPe12002-M8

```

If the command displays no result, there is no Fibre Channel card installed.

**Step 3** To verify the firmware version, type the following command:

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

The output might resemble the following:

```

Host Name      : angel-primary.qllabs.inc
Manufacturer   : Emulex Corporation
Serial Number  : FC10849279
Model         : LPe12002-M8
Model Desc    : Emulex LPe12002-M8 8Gb 2-port PCIe Fibre
Channel Adapter
Node WWN      : 20 00 00 00 c9 b7 67 5e
Node Symname  : Emulex LPe12002-M8 FV1.10A5 DV8.2.0.63.3p
HW Version    : 31004549
Opt ROM Version: 5.03a2
FW Version   : 1.10A5 (U3D1.10A5), sli-3
Vendor Spec ID : 10DF
Number of Ports: 1
Driver Name    : lpfc
Device ID     : F100
HBA Type      : LPe12002-M8
Operational FW : SLI-3 Overlay
SLI2 FW       : 1.10a5
SLI3 FW       : 1.10a5
IEEE Address   : 00 00 c9 b7 67 5e
Boot Code     : Enabled
Boot Version  : 5.03a2
Driver Version : 8.2.0.63.3p; HBAAPI(I) v2.3.b, 07-12-10
Kernel Version : 1.10a0
HBA Temperature: Normal
Function Type  : FC
Sub Device ID : F100
Sub Vendor ID  : 10DF

```

## Fibre Channel Configuration Types

You can configure Fibre Channel for use in a standard deployment or in a High Availability (HA) environment.

This section includes the following topics:

- [Configuring Fibre Channel in a Standard Deployment](#)
- [Configuring Fibre Channel HA](#)

### Configuring Fibre Channel in a Standard Deployment

To configure Fibre Channel in a standard deployment:

- Step 1** Prepare QRadar to connect to the Fibre Channel network. See [Preparing QRadar to Connect to Fibre Channel Network](#).
- Step 2** Migrate the storage directory to the Fibre Channel storage solution. By default, QRadar stores data in the /store directory, however, storing data in subdirectories of /store is supported. Choose one of the following:
  - [Migrating /store to the Fibre Channel Solution](#)
  - [Migrating a subdirectory of /store to the Fibre Channel Storage Solution](#)

Verify that Fibre Channel storage mounts properly. See [Verifying the Fibre Channel Mount](#).

### Configuring Fibre Channel HA

In an HA deployment, the secondary host maintains the same data as the primary host by one of two methods: data replication or shared external storage.

If you use the shared external storage method, your secondary host must be configured with the same external Fibre Channel device as the primary host.

We recommend that the Emulex LPe12002 Host Bus Adapter cards on the primary and secondary hosts are installed with the same driver version. To verify the Fibre Channel driver version, SSH into each host and enter the following command:

```
/sbin/modinfo lpfc | grep description
```

The output might resemble the following:

```
description: Emulex LightPulse Fibre Channel SCSI driver
8.2.0.63.3p
```



To configure Fibre Channel for use with HA, you must:

**Step 1** Configure Fibre Channel on the primary host:



**CAUTION**

---

*This step must be performed before adding the secondary host.*

---

- a Prepare the primary host to connect to the Fibre Channel network. See [Preparing QRadar to Connect to Fibre Channel Network](#).
- b Migrate the `/store` directory on the primary host to the Fibre Channel storage solution. See [Migrating /store to the Fibre Channel Solution](#).
- c Verify the Fibre Channel mount on the primary host. See [Verifying the Fibre Channel Mount](#).

**Step 2** Install the secondary host.

For more information, see the *Installation Guide* for your QRadar product.

**Step 3** Configure Fibre Channel on the secondary host:

- a Prepare the secondary host to connect to the Fibre Channel network. See [Preparing QRadar to Connect to Fibre Channel Network](#).
- b Migrate the `/store` directory on the secondary host to the Fibre Channel storage solution. Only perform [Step 2](#) through [Step 10](#) of the procedure described in [Migrating /store to the Fibre Channel Solution](#).

Access QRadar and configure the HA cluster. For more information about configuring HA, see the *IBM Security QRadar HA Guide*.

**NOTE**

---

The IBM Security QRadar HA Guide is not applicable to IBM Security QRadar Network Anomaly Detection.

---

## Configuring Fibre Channel

This section includes the following topics:

- [Preparing QRadar to Connect to Fibre Channel Network](#)
- [Migrating /store to the Fibre Channel Solution](#)
- [Verifying the Fibre Channel Mount](#)

## Preparing QRadar to Connect to Fibre Channel Network

To prepare QRadar to connect to a Fibre Channel network:

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

Username: `root`

Password: `<password>`

**Step 2** To verify the attached devices, type the following command:

```
dmesg | less
```

**Step 3** When the file is open, type the following command to search for the `lpfc` string:

```
:/lpfc
```

The output might resemble the following:

```
lpfc 0000:06:00.0: 0:1303 Link Up Event x1 received Data: x1 x2
x10 x2 x0 x0 0
Vendor: MAXTOR      Model: ATLAS15K2_146SCA  Rev: JNZ6
Type:   Direct-Access          ANSI SCSI revision: 03
SCSI device sdb: 286749480 512-byte hdwr sectors (146816 MB)
sdb: Write Protect is off
sdb: Mode Sense: bf 00 10 08
SCSI device sdb: drive cache: write through w/ FUA
SCSI device sdb: 286749480 512-byte hdwr sectors (146816 MB)
sdb: Write Protect is off
sdb: Mode Sense: bf 00 10 08
SCSI device sdb: drive cache: write through w/ FUA
sdb: sdb1
sd 3:0:0:0: Attached scsi disk sdb
Vendor: MAXTOR      Model: ATLAS15K2_146SCA  Rev: JNZ6
Type:   Direct-Access          ANSI SCSI revision: 03
```

This example verifies the Fibre Channel link and SCSI drive named `sdb` is connected to the network.

**Step 4** Reformat the Fibre Channel partition, if it has not previously been used.



#### CAUTION

---

*If the volume has been used before and you need to retain the data in the volume, then you cannot create partitions or reformat the partitions.*

---

- a Optional. Create a partition. For information about creating a partition, see your Linux documentation.
- b Reformat the partition by typing the following command:

```
mkfs.ext4 /dev/<device name>
```

Where `<device name>` is the name of the Fibre Channel device including the partition number.

For example: `sdd1`

**NOTE** You can create one or more partitions on the volume and mount them separately. If the new volume is larger than 2TB, you must create a GUID Partition Table (GPT). Using GPT, the new volume is limited to 16 TB. If you are using MSDOS partitioning, you are limited to a single 2TB partition.

You are now ready to migrate the storage directory to the Fibre Channel storage solution. Choose one of the following procedures:

- To migrate /store to the Fibre Channel storage solution, see [Migrating /store to the Fibre Channel Solution](#).

To migrate a subdirectory of /store to the Fibre Channel storage solution, see [Migrating a subdirectory of /store to the Fibre Channel Storage Solution](#).

### Migrating /store to the Fibre Channel Solution

To migrate the /store file system to the Fibre Channel storage solution:

**Step 1** To stop the services, type the following commands in order:

**NOTE** If you are configuring Fibre Channel on a secondary host in an HA deployment, skip this step and go to [Step 2](#).

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

**Step 2** Unmount /store/tmp by typing the following command:

```
umount /store/tmp
```

**Step 3** Unmount your existing /store directory by typing the following command:

```
umount /store
```

**Step 4** Create the /store\_old directory by typing the following command:

```
mkdir /store_old
```

**Step 5** Verify the Universally Unique Identifier (UUID) of the device partition by typing the following command:

```
blkid /dev/<device name>
```

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

The output should resemble the following:

```
/dev/sdb1: UUID="89ec181b-dcd1-4698-b1ae-9f1b1b044f62"
```

**Step 6** Open the fstab file for editing by typing the following command:

```
vi /etc/fstab
```

**Step 7** Locate the existing /store mount line, which resembles the following:

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

**Step 8** Modify the line by typing the following:

```
UID=<uuid> /store_old <file system> defaults 1 2
```

Where: <file system> is the version you used to format the file system.

For example: **ext4**.

**Step 9** If you are configuring Fibre Channel in an HA environment, choose one of the following options:

- If you are migrating /store on a primary host in an HA cluster, add the following line to the /etc/fstab file:

```
UUID<uuid> /store <file system> noatime 1 2
```

Where: <file system> is the version you used to format the file system.

For example: **ext4**.

- If you are migrating /store on a secondary host in an HA cluster, add the following line to the /etc/fstab file:

```
UUID=<uuid> /store <file system> noatime,noauto 1 2
```

**Step 10** Save and close the file.

**Step 11** Mount the new Fibre Channel /store file system by typing the following command:

```
mount /store
```

**Step 12** Mount the old /store file system by typing the following command:

```
mount /store_old
```

**Step 13** Copy the data from the existing /store file system to the Fibre Channel directory by typing the following command:

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

**Step 14** Re-mount /store/tmp by typing the following command:

```
mount /store/tmp
```

**Step 15** Unmount /store\_old by typing the following command:

```
umount /store_old
```

**Step 16** Restart the services by typing these commands in the following order:

```
service hostservices restart
```

```
service tomcat restart
```

```
service hostcontext restart
```

```
service systemStabMon restart
```

```
service crond restart
```

**NOTE** For most situations, you only need to mount a single /store on your Fibre Channel storage solution. If, however, you need a different configuration for your Fibre Channel mount points, contact Customer Support.

---

You are now ready to verify the Fibre Channel mount point. Go to [Verifying the Fibre Channel Mount](#).

### Migrating a subdirectory of /store to the Fibre Channel Storage Solution

To migrate a subdirectory of /store to the Fibre Channel storage solution:

**Step 1** To stop the services, type the following commands in the order specified:

**NOTE** If you are configuring Fibre Channel on a secondary host in an HA deployment, skip this step and go to [Step 2](#).

---

```
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 volume to the temporary directory by typing the following command:

```
mount /dev/<device name> /tmp/fcdata
```

**Step 4** Copy existing data to the Fibre Channel storage volume using the temporary mount point by typing the following command:

```
cp -af <subdirectory_path>/ * /tmp/fcdata
```

Where <subdirectory\_path> is the directory path to the subdirectory you want to migrate.

For example:

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

In this example, /store/ariel is the subdirectory

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

```
umount /tmp/fcdata
```

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

```
blkid /dev/<device name>
```

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

The output should resemble the following:

```
/dev/sdb1: UUID="89ec181b-dcd1-4698-b1ae-9f1b1b044f62"
```

**Step 7** Open the `fstab` file for editing by typing the following command:

```
vi /etc/fstab
```

**Step 8** Add the mount line for the new subdirectory mount point by typing the following line:

```
UUID=<uuid> /store/ariel <file system> noatime,nobarrier 1 2
```

Where: `<file system>` is the version you used to format the file system.

For example: `ext4`.

**Step 9** Save and close the file.

**Step 10** To mount the new Fibre Channel `/store/ariel` subdirectory, type the following command:

```
mount /store/ariel
```

**Step 11** To restart the services, type these commands in the following order:

```
service hostservices restart
service tomcat restart
service hostcontext restart
service systemStabMon restart
service crond restart
```

You are now ready to verify the Fibre Channel mount. Go to [Verifying the Fibre Channel Mount](#).

### Verifying the Fibre Channel Mount

To verify that the Fibre Channel device mounts correctly:

#### NOTE

This procedure is not required when configuring Fibre Channel on a secondary HA host.

**Step 1** Type the following command:

```
df -h
```

**Step 2** Review the screen output and look for the newly added volume.

The migrated directory should be linked to the configured Fibre Channel storage device, such as `/dev/sdc1`, in the following example:

```
# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sda2 12G 5.4G 6.5G 46% /
/dev/sda1 99M 50M 44M 54% /boot
```

```
/dev/sda3 11G 406M 9.7G 4% /var/log  
/dev/sdc1 910G 558M 663G 1% /store  
/dev/sda5 10G 33M 10G 1% /store/tmp
```





# 4

## USING NFS FOR QRADAR BACKUPS

Using NFS, you can store IBM Security QRadar backup data externally using existing network infrastructures and protocols.

This section includes the following topics:

- [NFS Considerations](#)
- [Implementing NFS for Backups](#)

---

### NFS Considerations

While QRadar supports NFS for external storage, we recommend that you do not use NFS for storing active data, including:

- **Postgres Database** - The postgres database is stored in the `/store/postgres/` directory. Database corruption can occur if you write data to a `/store` file system that is mounted on the NFS. You should mount the `/store/postgres` partition on a local disk, not on NFS.
- **Ariel Database** - The ariel database is stored on the `/store/ariel/` directory. Performance issues can occur if ariel data is stored on the NFS. Every minute, a series of distinct files are created by QRadar and this can compromise system performance.

You should only use NFS for QRadar backups, which are stored in the `/store/backup/` directory. To do this, mount your NFS storage on the `/store/backup/` partition. For more information about backing up your data, see your Administration Guide.

---

### Implementing NFS for Backups

To implement NFS for QRadar backups:

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

Username: `root`

Password: `<password>`

**Step 2** Open the `/etc/hosts` file for editing 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> nfsserver
```

Where:

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

**Step 4** Save and close the file

**Step 5** Edit the iptables firewall to enable the connection to your NFS server:

a Open the iptables.pre file for editing by typing the following:

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

b Add the following line:

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

Where:

<interface> is the QRadar interface on your NFS network. This is ETH0, unless you have configured a dedicated NFS network using ETH1.

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

**Step 6** Restart iptables by typing the following command:

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

The NFS services are disabled by default.

**Step 7** Add the 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 8** Manually start NFS services by typing the following commands:

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

**Step 9** Configure the mount point for the /store/backup file system:

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

```
vi /etc/fstab
```

b Add the following line:

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

Where:

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

#### NOTE

---

You may need to adjust the settings for the NFS mount point to accommodate your configuration. For example: `/nfsshare/qradar/backup /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.

---

**Step 10** Migrate existing backup files to the NFS volume:

- a Move your backup files from the existing `/store/backup` directory to a temporary location by typing the following commands:

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

- b Create a new backup directory by typing the following command:

```
mkdir /store/backup
```

- c Set the permissions for the NFS volume by typing the following command:

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

- d Mount the NFS volume by typing the following command:

```
mount /store/backup
```

- e Verify that `/store/backup` is mounted by typing the following command:

```
df- h
```

- f Move the backup files from the temporary location to the NFS volume by typing the following command:

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

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

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

Your NFS backup is now mounted and operational.



# 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 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 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](http://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](http://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.





# INDEX

---

## Symbols

- /store
  - migrating to iSCSI 16
  - migrating using fibre channel 33
  - when to consider migrating 6
- /store/ariel
  - migrating to iSCSI 14
  - when to consider migrating 7

---

## A

- archiving
  - using fibre channel 28
- auto-mount
  - iSCSI volumes 18

---

## B

- before you begin
  - fibre channel 28
- best practices
  - fibre channel 27

---

## C

- configure iSCSI for HA
  - before you begin 11
  - connecting the secondary and primary device 23
- conventions 3

---

## E

- external storage
  - limitations 8
  - types of stored data 6
  - using with HA 9
  - when to consider 5
- external storage options
  - fibre channel 7
  - iSCSI 8
  - NFS 8

---

## F

- fibre channel
  - archiving 28
  - before you begin 28
  - best practices 27
  - configuration types 30
  - configuring 31
  - connecting qradar 31
  - in a standard deployment 30

- migrating a subdirectory of /store 35
- more information 7
- using in an HA environment 30
- using volumes 28
- verifying mount points 36

---

## H

- high availability
  - before you begin 11
  - configuring external storage 9
  - using iSCSI 18

---

## I

- intended audience 3
- iSCSI
  - assigning volumes 13
  - configuring volumes 13
  - migrating data 14
  - more information 8
  - troubleshooting 23
  - usage in a standard QRadar deployment 12
- iSCSI connections
  - verifying 22
- iSCSI HA
  - assigning iSCSI volumes 19
  - auto-mounting iSCSI volumes 21
  - configuring secondary host mount points 20
  - connecting a secondary host to iSCSI 19
  - connecting the primary and secondary host 23
- iSCSI hA
  - configuring iSCSI volumes 19
- iSCSI network
  - connecting QRadar 12
- iSCSI network
  - connecting a secondary HA host 19
- iSCSI volumes
  - assigning and configuring 13
  - auto-mounting 18
  - auto-mounting using iSCSI HA 21
- iSCSI with HA
  - using iSCSI with HA 18

---

## L

- limitations of external storage 8

---

## M

- migrating /store
  - using fibre channel 33
- migrating the /store file system

when to consider 6  
migrating the /store/ariel file system  
when to consider 7

---

## **N**

### NFS

implementing for backups 39  
more information 8  
using with /store/backup 39

---

## **Q**

### QRadar data

migrating to iSCSI 14

### QRadar standard deployment

using iSCSI 12

---

## **S**

### secondary HA host

assigning and configuring volumes 19  
auto-mounting iSCSI volumes 21  
configuring mount points 20  
connecting to the iSCSI network 19

---

## **T**

### troubleshooting

detecting disk errors 24

iSCSI 23

mounting iSCSI volumes 24

reconfiguring a failed primary host 23

unmounting iSCSI volumes 24

### troubleshooting iSCSI

verifying connections to iSCSI 22

---

## **W**

when to migrate the /store file system 6

when to migrate the /store/ariel file system 7