TECHNICAL NOTE

CONFIGURING iSCSI

APRIL 2010

You can use your iSCSI storage network in your STRM or STRM Log Manager deployment. This document provides information on configuring and using iSCSI devices with your deployment. Unless otherwise noted, all references to STRM refer to both STRM and STRM Log Manager.

Caution: *The procedures in this document assume an advanced knowledge of a Linux operating system. For assistance, please contact Juniper Networks Customer Support.*

Note: *Your network configuration may differ, however, for the purposes of the document, we assume that your management interface is ETH0 and your iSCSI interface is ETH1.*

Note: *STRM supports the LeftHand storage solution. Before you prepare STRM, the LeftHand storage solution should be connected and configured. For more information about LeftHand, refer to your LeftHand documentation.*

iSCSI Configuration Types

You can configure iSCSI for use in a standard deployment or in an HA environment.

This section identifies the procedures for the following configuration types:

- Configuring iSCSI in a Standard Deployment
- Configuring iSCSI for use with HA
- Configuring iSCSI When Restoring a Failed Primary HA Console

Configuring iSCSI in a Standard Deployment

To configure iSCSI in a Standard Deployment:

Step 1 Prepare STRM to connect to the iSCSI network. See Preparing STRM to Connect to iSCSI Network.

Step 2 Assign and configure the iSCSI volumes. See Assigning and Configuring iSCSI Volumes.

Step 3 Migrate the /store directory to the iSCSI storage solution. See Migrating /store to the iSCSI Storage Solution.

Step 4 Configure the system to auto-mount the iSCSI volume. See Configuring the System to Auto-mount the iSCSI Volume.

Step 5 Verify the iSCSI mount. See Verifying the iSCSI Mount.

Configuring iSCSI for use with HA

In a High Availability (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 by configured with the same external iSCSI device as the primary host.

To configure iSCSI for use with HA, you must:

Step 1 Configure iSCSI on the primary host:

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

- **a** Prepare the primary host to connect to the iSCSI network. See Preparing STRM to Connect to iSCSI Network.
- **b** Assign and configure the iSCSI volumes on the primary host. See Assigning and Configuring iSCSI Volumes.
- **c** Migrate the /store directory on the primary host to the iSCSI storage solution. See Migrating /store to the iSCSI Storage Solution.
- **d** Configure the primary host to auto-mount the iSCSI volume. See Configuring the System to Auto-mount the iSCSI Volume.
- **e** Verify the iSCSI mount on the primary host. See Verifying the iSCSI Mount.

Step 2 Install the secondary host. See the *STRM Installation Guide* or the *STRM Log Manager Installation* guide.

Step 3 Configure iSCSI on the secondary host:

- **a** Prepare the primary host to connect to the iSCSI network. See Preparing STRM to Connect to iSCSI Network.
- **b** Assign and configure the iSCSI volumes on the primary host. See Assigning and Configuring iSCSI Volumes. Skip Step 5.
- **c** Migrate the /store directory on the primary host to the iSCSI storage solution. See Migrating /store to the iSCSI Storage Solution. Only perform Step 2 through Step 9.
- **d** Configure the secondary host to auto-mount the iSCSI volume. See Configuring the System to Auto-mount the iSCSI Volume.

Step 4 Access STRM and configure the HA cluster. For more information about configuring HA, see the *STRM Administration Guide* or the *STRM Log Manager Administration Guide*.

Configuring iSCSI When Restoring a Failed Primary HA Console

To configure iSCSI when restoring a failed primary HA Console:

Step 1 Prepare the primary host to connect to the iSCSI network. See Preparing STRM to Connect to iSCSI Network.

Step 2 Assign and configure the iSCSI volumes on the primary host. See Assigning and Configuring iSCSI Volumes. Skip Step 5.

Step 3 Migrate the /store directory on the primary host to the iSCSI storage solution. See Migrating /store to the iSCSI Storage Solution. Only perform Step 2 through Step 9.

Step 4 Configure the system to auto-mount the iSCSI volume. See Configuring the System to Auto-mount the iSCSI Volume.

Configuring iSCSI

This section includes:

- Preparing STRM to Connect to iSCSI Network
- Assigning and Configuring iSCSI Volumes
- Migrating /store to the iSCSI Storage Solution
- Configuring the System to Auto-mount the iSCSI Volume
- Verifying the iSCSI Mount

Preparing STRM to Connect to iSCSI Network

To prepare STRM to connect to an iSCSI network:

Step 1 From the Admin interface, configure a secondary network interface with a private IP address to connect to the iSCSI SAN.

Note: *Network interface address information, from your SAN network manager, is required. For more information on configuring a network interface, see the STRM Administration Guide.*

Step 2 Log into STRM as root.

Step 3 Open the following file:

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

Step 4 Add the iSCSI qualified name for your host to the initiatorname.iscsi file in the following format:

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

For example:

```
InitiatorName=iqn.2008-11.com.q1labs:pl13
```

Step 5 Save and close the file.

Step 6 Enter the following command to start/restart the iSCSI service to open a session to the server:

```
service iscsi restart
```

You are now ready to assign and configure the iSCSI volumes. See Assigning and Configuring iSCSI Volumes.

Assigning and Configuring iSCSI Volumes

To assign and configure iSCSI volumes:

Step 1 Detect volumes on the iSCSI server:

```
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.

A list of iSCSI volumes available to your host appears. The output should resemble the following:

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

Step 2 Make sure the login for your iSCSI server is functional:

```
iscsiadm -l -m node -T <target iqn name>
```

The output from this command should resemble:

```
sLogging 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 Make sure your STRM system detects the new disks:

```
dmesg | tail -20 | grep 'Attached'
```

The output from this command should resemble:

```
sd 0:0:1:0: Attached scsi disk sdb
```

Step 4 Determine the iSCSI device name:

a Clear the kernel ring buffer:

```
dmesg -c
```

b Reload the iSCSI service:

```
/etc/init.d/iscsi restart
```

c Locate the device name:

```
dmesg | grep 'Attached scsi disk'
```

The output from this command should resemble:

```
sd 4:0:0:0: Attached scsi disk sdd
```

Note: *If you are configuring iSCSI for a secondary host in an HA deployment, go to* Migrating /store to the iSCSI Storage Solution.

Step 5 To reformat the volume, if it has not previously been used, choose one of the following:

Caution: *If the volume has been used before, do not reformat the volume if you want to retain the data already on the volume.*

a Create a partition.

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

b Reformat the volume using a volume label:

```
mkfs.ext3 -l <volume label> /dev/<device name>
```

Where:

- **<device name>** is the name of the iSCSI device including the partition number.

For example: sdd1

c Reformat the volume without a volume label:

```
mkfs.ext3 <device name>
```

Note: *You can create one or more partitions on the volume and mount them separately. If the new volume is larger than 2 TB, you must create a GPT partition. If you are using Red Hat Enterprise Linux, the new volume must be smaller than 2 TB.*

You are now ready to migrate the /store directory to the iSCSI storage solution. See Migrating /store to the iSCSI Storage Solution.

Migrating /store to the iSCSI Storage Solution

To migrate /store to the iSCSI storage solution:

Step 1 Stop the services in the following order:

Note: *Skip this step if you are configuring iSCSI on a secondary host in an HA deployment. Go to* Step 2.

```
service systemStabMon stop
service hostcontext stop
service tomcat stop
service imq stop
service postgresql stop
```

Step 2 Unmount /store/tmp:

```
umount /store/tmp
```

Step 3 Unmount your existing /store directory:

```
umount /store
```

Step 4 Create the /store_old directory:

```
mkdir /store_old
```

Step 5 Open the following file:

```
/etc/fstab
```

Step 6 Locate the existing /store mount line:

```
LABEL=/store /store ext3 noatime 1 2
```

Step 7 Modify the line to match:

```
LABEL=/store /store_old ext3 noatime 1 2
```

Step 8 Choose one of the following options:

a If you reformatted the volume using a volume label, create the following line:

```
LABEL=<volume label> /store ext3 noatime
```

b If you reformatted the volume without a volume a label, create the following line:

```
<device name> /store ext3 noatime
```

Step 9 Save and close the file.

Note: *If you are configuring iSCSI on a secondary host in an HA deployment, go to* Configuring the System to Auto-mount the iSCSI Volume*.*

Step 10 Mount the new iSCSI /store:

```
mount /store
```

Step 11 Mount the old /store:

```
mount /store_old
```

Step 12 Copy the data from the existing /store to the /san directory.

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

Step 13 Re-mount /store/tmp:

```
mount /store/tmp
```

Step 14 Unmount /store_old:

```
umount /store_old
```

Step 15 Restart the services in the following order:

```
service postgresql restart
service imq restart
service tomcat restart
service hostcontext restart
service systemStabMon restart
```

Note: *For most situations, you only need to mount a single /store on you iSCSI storage solution. If, however, 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 Open the following directory:

```
cd /etc/rc3.d
```

Step 2 Add the iSCSI script to be part of the startup.

For example:

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

Step 3 Create a link to /opt/qradar/init/iscsi-mount script from the /etc/init.d directory:

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

Step 4 Add iscsi-mount script to be part of the startup:

For example:

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

You are now ready to verify the iSCSI mount. See Verifying the iSCSI Mount.

Verifying the iSCSI Mount

To verify that the iSCSI mounts properly:

Note: *This procedure is not required when configuring iSCSI on a secondary HA host.*

Step 1 Enter the following command:

```
df -h
```

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

For 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
#
```

Published: 2010-04-01