IBM System Storage
N series

Linux Host Utilities 6.1 Quick Start Guide

This guide is for experienced Linux users. It provides the basic information required to get the Linux Host Utilities installed and set up on a Linux host.

The steps listed are for a typical installation and cover multiple Linux environments, including Red Hat Enterprise Linux or SUSE Linux and Veritas Storage Foundation with Veritas Dynamic Multipathing (VxDMP). While many steps are common to all environments, some steps apply only to a specific environment.

If you are not an experienced Linux user, see the *Linux Host Utilities Installation and Setup Guide*. That document provides detailed steps and examples. It also includes troubleshooting information as well as instructions for optional tasks, such as setting up a SAN boot LUN.

For information about installing and setting up third-party hardware and applications, such as host bus adapters or Veritas Storage Foundation, see the documentation that accompanies those products.

Note: Occasionally there are known problems that can affect your system setup. Read the *Linux Host Utilities Release Notes* before you install the Host Utilities. The *Release Notes* are updated whenever an issue is found and might contain the information that was discovered after this guide was produced.

Task 1: Make sure the prerequisites for installing and setting up the Host Utilities have been met

To install the Host Utilities and set up your system, you must perform tasks on both the host and the storage system. In some cases, the tasks you perform vary depending on your environment; for example, the protocol you are using and whether you are using DM-Multipath for multipathing or VxDMP.

Note: The IBM N series interoperability matrix website (accessed and navigated as described in Websites) contains the most current information about supported environments for the Host Utilities.

1. Verify that your host system is correct, including:
   - Host operating system version and patches
   - HBAs and drivers, model and version, or software initiator

     You must make sure the HBAs and drivers are installed and configured with the appropriate parameter values. For more information about setting up the HBAs and drivers, see the *Linux Host Utilities Installation and Setup Guide.*
   - Volume management and multipathing, if you are using multipathing
   - Veritas Storage Foundation, if you are using Veritas Storage Foundation

     Additional Veritas Storage Foundation setup: Veritas Storage Foundation requires the Array Support Library (ASL) and Array Policy Module (APM).
     - Veritas Storage Foundation 5.1 includes these modules as part of its installation so they are automatically installed when you install Veritas Storage Foundation.
     - Veritas Storage Foundation 5.0 requires that you download the correct versions from the Symantec website and manually install them on the host.

       For instructions on installing these modules, see the *Linux Host Utilities Installation and Setup Guide*.

     You must also set the following parameters to the specified values:

- **dmp_restore_interval=60**
- **dmp_restore_policy=disabled**
- **dmp_lun_retry_timeout=300**
- **(Veritas Storage Foundation 5.1 SP1 and later) dmp_path_age=120**
- **(iSCSI only) node.session.timeo.replacement_timeout = 120**

Note: (Red Hat Linux 4) ConnFailTimeoutof = 5

(Red Hat Linux 6 and later, SUSE Linux 11 and later) In addition, if the host is running either Red Hat Linux 6 or later or SUSE Linux Enterprise Server 11 or later, you must configure it to work with Veritas Storage Foundation by performing the following steps:

a. **(SUSE Linux 11 and later only)** Install SUSE Linux Enterprise Server 11 with kernel version 2.6.27.45-0.1.1 or later from Novell.

b. **(Red Hat Linux 6 and later, SUSE Linux 11 and later)** Create the file /etc/udev/rules.d/40-rport.rules with the following content line: `KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add",RUN+="/bin/sh -c 'echo 20 > /sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000 >/sys/class/fc_remote_ports/%k/dev_loss_tmo'"`

c. **(Red Hat Linux 6 and later)** Make sure the value of the IOFENCE timeout parameter is set to `30000`

You can use the command **gabconfig -f 30000** to set this value. Because this value is not persistent across reboots, you must check it each time you boot the host.

d. Reboot the host.

2. Verify that your storage system is set up correctly, which includes having:
   - The correct version of Data ONTAP installed.

     Note: With Red Hat Enterprise Linux 6.0 or later, it is recommended that you run Data ONTAP 8.0.1 or higher on the storage systems connected to the host. That way you can use the Block Limits VPD page (0xb0) information, which results in enhanced I/O performance on the N series LUN.
   - The appropriate license for the protocol on which your environment runs.
   - Appropriate HBAs or software initiators set up to work with the host as needed by your protocol.
   - ALUA enabled on all platforms that support it.

     Note: Starting with Data ONTAP 8.0, ALUA is the default igroup. For iSCSI environments, ALUA is not supported.

     (Veritas environments) To use ALUA in a Veritas environment, you must have Veritas Storage Foundation 5.1 installed. In a Veritas environment, the ASL detects whether ALUA is enabled. It sets the array type to A/A-NETAPP for non-ALUA environments and to ALUA for ALUA environments.
   - Working volumes and qtrees (if desired) set up.
3. **(FC environments)** If you are using a fabric connection, verify that the switch is set up correctly, which includes having the switch:
   - Cabled according to the instructions in the *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP.
   - Zoned appropriately, using the supported zoning technique in single initiator zoning from a host's initiator's standpoint.
   - Powered on in the correct order: switch, disk shelves, storage systems, and then the host.

   Note: For information about supported topologies, see the *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP.

4. Confirm that the host and the storage system can communicate.

Task 2: Install and set up the Host Utilities

To install and set up the Host Utilities, you must download the .rpm file containing the Host Utilities software. After you install this software, you may need to perform additional configuration steps for your environment.

You must be logged on as "root" to install or uninstall the Host Utilities.

1. Remove any currently installed version of the Host Utilities.

   For the Linux Host Utilities 5.3 or later, enter the **rpm -e** *package_name* command.

   For any version of the Host Utilities prior to 5.3, go to the directory where that version is installed (the default is `/opt/ontap/santools`) and enter the **./uninstall** command.
2. Get a copy of the .rpm file containing the Host Utilities software for your environment.

   The Linux Host Utilities provide two versions of the software package:
   - A 32-bit version: ontap_linux_host_utilities-6-1.i386.rpm
   - A 64-bit version: ontap_linux_host_utilities-6-1.x86_64.rpm
3. Go to the directory where you downloaded the Host Utilities file.
4. Install the Host Utilities software by entering the following command:

   If you are upgrading the Host Utilities from Linux Host Utilities 5.3 or later, you can use the **rpm -Uvh** *package_name* command.

(iSCSI environments) Task 3: Configure the iSCSI protocol

If you are using iSCSI, you must make sure your system is correctly set up for this protocol. Some of the tasks you perform differ based on whether you have a software iSCSI initiator or a hardware iSCSI HBA.

1. Record the host's iSCSI node name so that you can supply it when you create an igroup.

   The names are stored in the following files. You can open the file for your version of Linux with a text editor:

If you are using...	Open the following file...
Red Hat Enterprise Linux 5 or 6 series SUSE Linux Enterprise Server 10 or 11 series	`/etc/iscsi/initiatorname.iscsi`
Red Hat Enterprise Linux 4 series	`/etc/initiatorname.iscsi`

   Tip: Before you record the node name, you may want to change it to something that is easier to use. You cannot change the standard format of the node name, but you can change the random numbers at the end of the name to something like the host name. So, if the node name is **iqn.2005-03.com.RedHat:012345** or **iqn.1996-04.de.suse:012345**, you can change the numbers that appear at the end (in these examples, 012345). For more details, see the *Linux Host Utilities Installation and Setup Guide*.
2. To use multipathing, you must set the timeout values in the iSCSI configuration file.

   Note: You only need to change the timeout values when you are using the iSCSI protocol. If you are using the FC protocol, the default timeout values are sufficient.

If you are using...	Do the following...
Red Hat Enterprise Linux 5 or 6 series SUSE Linux Enterprise Server 10 or 11 series	1. Using a text editor, open the `/etc/iscsi/initiatorname.iscsi` file. 2. Modify `node.session.timeo.replacement_timeout`. For DM-Multipath environments, set the value to `5`. For VxDMP environments, set the value to `120`.
Red Hat Enterprise Linux 4 series	1. Using a text editor, open the `/etc/initiatorname.iscsi` file. 2. Remove the comment indicator from `ConnFailTimeout` and set it to `5`.

3. **(Optional)** You can set up the CHAP protocol to provide enhanced security.
   For details on doing this, see the *Linux Host Utilities Installation and Setup Guide*.
4. Start the iSCSI service.
   At the Linux host command prompt:

If you are using...	Enter...
Red Hat Enterprise Linux 6 series **Note:** You must execute this command the first time you start the iSCSI service on a host running Red Hat Enterprise Linux 6 series. If you execute **`/etc/init.d/iscsi start`** without previously executing **`service iscsid force-start`**, you will get an error message.	**`service iscsid force-start`**
Red Hat Enterprise Linux 5 or 4 series SUSE Linux Enterprise Server 10 or 11 series	`etc/init.d/iscsi start`

5. Configure the iSCSI initiator to discover the target so that the host can access LUNs on the target.
   The method you use to discover the target depends on your version of the operating system:
   - Red Hat Enterprise Linux 5 or 6 series: Use the **`iscsiadm`** utility.
   - Red Hat Enterprise Linux 4 series: Enter the IP address of an Ethernet interface on the storage system as the value for `DiscoveryAddress` in the `/etc/iscsi.conf` file.
   - SUSE Linux Enterprise Server 10 or 11 series: Use either the **`iscsiadm`** utility or YaST2. If you use YaST2, make sure the port number is 3260.

   For details on performing target discovery, see the *Linux Host Utilities Installation and Setup Guide*.
6. Use the **`chkconfig`** command to configure the iSCSI service to start automatically at system boot:

If you are using...	Enter...
Red Hat Enterprise Linux	**`chkconfig iscsi on`**
SUSE Linux Enterprise Server 10 or 11 series	**`chkconfig open-iscsi on`**

7. Specify whether the system automatically logs in to an iSCSI node at startup or whether you must manually log it in to the node.
   Setting the login mode affects only the nodes that are discovered after the value is set. If you set your login mode to manual, you must log in to the nodes manually the next time the system starts up. If you set your login mode to automatic, the system automatically logs in to the nodes when it starts.

   Note: (Red Hat Enterprise Linux 4 series) When you are running Red Hat Enterprise Linux 4 series, all sessions are logged in automatically when you start the iSCSI service.

   a. Start the iSCSI service.

b. Set the login mode:
   - To set a login mode for all targets and their ports, edit the `/etc/iscsi/iscsid.conf` file.
   - To set the login mode for a specific portal on a target or for all the portals on a target, use the **iscsiadm** command.

   If you are using the **iscsiadm** command, you must specify the command line. Based on the command line you enter, the **iscsiadm** command sets the mode as either manual or automatic for a specific port on a target, all the ports on the target, or all the targets:

To set the login mode for...	Do the following...
A specific port on a target	Enter the appropriate command for your system: **iscsiadm --mode node -T** *targetname* **-p** *ip:port* **-o update -n node.startup -v manual|automatic** or **iscsiadm --mode node -T** *targetname* **-p** *ip:port* **-o update -n node.conn[0].startup -v manual|automatic**
All the ports on a target	Enter the appropriate command for your system: **iscsiadm --mode node -T** *targetname* **-o update -n node.startup -v manual|automatic** or **iscsiadm --mode node -T** *targetname* **-o update -n node.conn[0].startup -v manual|automatic**
All the targets	1. Modify the **/etc/iscsi/iscsid.conf** file to add the following line. You must specify either `manual` or `automatic`: `node.startup = manual|automatic` 2. Rediscover the iSCSI target. 3. Restart the iSCSI service.

Task 4: Set up DM-Multipath for native Linux environments

If you are using a native Linux environment, you must configure DM-Multipath by specifying values in the multipath.conf file.

Note: If you are using Veritas Dynamic Multipathing, you must set that up instead of the native DM-Multipath. For instructions on doing that, see the Veritas documentation and the Veritas notes included in the section *Task 1: Make sure the prerequisites for installing and setting up the Host Utilities have been met.*

The following are some points to remember about DM-Multipath:

- DM-Multipath creates a single device in `/dev/mapper/` that contains all the paths to a single LUN.

  This path is on top of the SCSI devices that Linux creates for each path to a LUN. For example, suppose a single LUN has two paths. Linux creates devices such as `/dev/sdd` and `/dev/sdf` for these paths. DM-Multipath then creates a single multipath device on top of these two devices. So, in this example, the DM-Multipath device `/dev/mapper/360a98000433468525634447 17a513571` appears in `/dev/mapper/` and contains all the paths to that LUN.
- When you are using DM-Multipath, you should create a file system for each LUN and then mount the LUN using the device in `/dev/mapper/`.

To set up DM-Multipath, edit the `/etc/multipath.conf` file to provide the recommended values. If you do not have this file, you can copy the sample configuration file that comes with your operating system. When you edit this file:

1. Use the **`blacklist`** section to exclude all the devices that do not correspond to LUNs configured on the storage systems that are mapped to your host.

   These are the devices that do **not** show up when you enter the command: `sanlun lun show`

2. Make sure you use the correct settings based on whether you have ALUA enabled:

If you are running...	With ALUA...	Without ALUA...
Red Hat Enterprise Linux 6 series SUSE Linux Enterprise Server 11 series SUSE Linux Enterprise Server 10 SP2 or later	Set **`prio`** to: `"alua"`	Set **`prio`** to: `"ontap"`
Other versions of Red Hat Enterprise Linux	**Red Hat Enterprise Linux 5.1 and later:** Set **`prio_callout`** to: `"/sbin/mpath_prio_alua /dev/%n"`	**Red Hat Enterprise Linux 5 series and 4 series:** Set **`prio_callout`** to: `"/sbin/mpath_prio_alua /dev/%n"`
SUSE Linux Enterprise Server 10 SP1 and earlier	Not Applicable	Set **`prio_callout`** to: `"sbin/mpath_prio_ontap /dev/%n"`
All supported Linux operating systems that support ALUA	Set **`hardware_handler`** to: `"1 alua"`	Set **`hardware_handler`** to: `"0"`

   For more details and examples of the values recommended for the multipath.conf file for different versions of Red Hat Enterprise Linux and SUSE Linux Enterprise Server, see the *Linux Host Utilities Installation and Setup Guide* and the *Recommended Host Settings for Linux Host Utilities*.

3. Start DM-Multipath.

   You can either start DM-Multipath manually or configure it to start automatically while booting.

   (Manual start) If you start it manually and configure the LUNS to work with it.

   a. Enter the **`start`** command line:

If you are running...	Enter the following command...
Red Hat Enterprise Linux	`# /etc/init.d/multipathd start`

SUSE Linux Enterprise Server	`# /etc/init.d/boot.multipath start`
	
	or
	
	`# /etc/init.d/multipathd start`

   b. To configure the DM-Multipath devices, enter the following command: `# multipath`

   (Automatic start) If you want to configure DM-Multipath to start automatically when the system boots, you must add the multipath service to the boot sequence:

If you are running...	Enter the following commands...
Red Hat Enterprise Linux	`chkconfig --add multipathdchkconfig multipathd on` **Note:** You should reboot the host if you are configuring a SAN boot LUN on the host.
SUSE Linux Enterprise Server	`chkconfig --add boot.multipathchkconfig --add multipathdchkconfig boot.multipath onchkconfig multipathd on`

Task 5: Set up access between the host and the LUNs on the storage system

You must make sure your host discovers and can work with the LUNs on the storage system.

The method you use when working with LUNs often varies depending on your system environment and factors such as whether you are using multipathing, whether you have an HBA, a hardware iSCSI initiator, or a software iSCSI initiator, whether you are using Veritas Storage Foundation, and which version of Linux you are using.

Keep the following notes in mind as you work with LUNs:

- You must be logged in as "root" to execute the Host Utilities commands, such as **sanlun**.
- If you have more than one path from a host to a LUN, you should use multipathing.

  Without multipathing software, the host cannot distinguish multiple LUNs from multiple paths to the same LUN.
- If you are not using multipathing, you should:
  - limit each LUN to a single path.
  - provide persistent identification for the LUNs (this is a best practice).

1. Create and map igroups and LUNs.

   You must create at least one igroup and at least LUN and then map the LUN to the igroup. The **lun setup** command steps you through this process. For details on creating an igroup and LUNs, see the *SAN Administration Guide (called Block Access Management Guide for iSCSI and FC in Data ONTAP 8.1 and earlier)* for your version of Data ONTAP.
2. **(FC)** If your environment is running the FC protocol and supports ALUA, make sure ALUA is enabled. With Data ONTAP 8.0 and later, ALUA is automatically enabled when you create an igroup in an environment using FC.

   To determine whether ALUA is enabled, enter the command: `igroup show -v igroup_name`

   If ALUA is not enabled, enable it by entering the command: `igroup set igroup_name alua yes`

   Note: You must use ALUA if you are running Data ONTAP in Cluster-Mode.
3. Discover the new LUNs by entering the appropriate command for your environment.

   Note: After LUNs have been discovered, they are automatically added to the DM-Multipath configuration.

Operating system	Method
FC or hardware iSCSI initiator environments using DM-Multipath or VxDMP	
All Linux operating systems	• **(Veritas and Linux operating systems)** Enter: `/usr/bin/rescan-scsi-bus.sh` • **(Veritas only)** If you are using Veritas Storage Foundation, initiate a rescan of the operating system device tree from the Veritas Volume Manager (VxVM) by entering the following command: `vxdisk scandisks`
Software iSCSI initiator environments using DM-Multipath or VxDMP	

Operating system	Method
Red Hat 5 or 6 series **SUSE Linux Enterprise Server 10 or 11 series**	To discover a new LUN on a system running DM-Multipath, enter one of the following commands: • To obtain the list of all the current sessions, enter: `iscsiadm -m session` • To rescan all the sessions, enter: `iscsiadm -m session --rescan` • To rescan using the SCSI **rescan** script, enter: `/usr/bin/rescan-scsi-bus.sh` • **(Veritas only)** To rescan Veritas devices, enter: `vxdisk scandisks`
Red Hat 4	1. Enter the following command on the Linux host to reload the iSCSI service:`/etc/init.d/iscsi reload`

4. Verify that the host has discovered the LUNs.

   Note: (Veritas only) If you want to view LUNs on the VxVM disks, use the **vxdisk list** command.

If you are using...	Enter the following command...
All DM-Multipath and VxDMP environments	`sanlun lun show all`
(iSCSI software initiator) Red Hat Enterprise Linux 5 or 6 series	`iscsiadm --mode session --sid=N -P 3`
(iSCSI software initiator) Red Hat Enterprise Linux 4 series	`iscsi-ls -l`

(iSCSI software initiator) SUSE Linux Enterprise Server 10 or 11 series	• **(SUSE Linux Enterprise Server 10 SP2 or later):**
	`iscsiadm --mode session --sid=N -P 3`
	• **(SUSE Linux Enterprise Server 10 SP1):**
	`iscsiadm --mode session --sid=N -i`

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:

  www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

  www.ibm.com/storage/support/nseries/

  This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

  www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

NA 210-05669_A0, Printed in USA

GA32-2206-01