

## Configuring Tape Devices for EMC NetWorker

### Technical Note

P/N 300-008-352

REV A09

January 2014

---

This technical note describes best practices for configuring tape devices for use by NetWorker servers and storage nodes.

The following topics are covered:

◆ Common Device Interface usage .....	2
◆ Enabling CDI .....	3
◆ Persistent binding and naming .....	3
◆ Automatic library and device configuration .....	5
◆ Tape drives in virtualized environments .....	5
◆ Shoe-shining .....	7
◆ Device ordering .....	8
◆ Multipathing .....	9
◆ SAN best practices .....	9
◆ Device block sizes .....	10
◆ LTO tape drive firmware issues .....	13
◆ Hardware-based encryption .....	13
◆ Device configuration best practices .....	14
◆ Preventing SCSI resets .....	27

## Common Device Interface usage

The Common Device Interface (CDI) media analysis tool was introduced in 2001 with NetWorker release 7.0 to provide a generic passthrough solution across all operating systems for the purposes of tape control and status collection, without depending on the implicit operating system driver mechanisms and without getting in the way of "Read/Write" operations of the devices.

CDI is a newer industry-wide method of talking to drives compared to old-style magnetic tape input/output.

In the past, there were a lot of reports of CDI causing problems and thus it was frequently disabled. In all cases, it was due to faulty tape drive device drivers that did not perform well when SCSI passthrough commands were passed to it.

### Advantages of CDI

The following are advantages of using CDI:

- CDI enables the use of the TapeAlert feature. In case of errors, the tape drive can send a full message back to the NetWorker software instead of simply providing a generic error code. TapeAlert provides intelligent tape drive management, an extension to the SCSI error system.
- CDI enables on-demand autocleaning of tape devices.
- MT-IO by design causes tape drives to drop out of streaming mode to write file markers, thus dropping the overall performance of the drives.
- CDI is required for many new features added to the NetWorker software in recent versions, such as scanning of device serial numbers for early detection of device ordering issues, which is a new feature in NetWorker release 7.4 Service Pack 3 (SP3).
- CDI is required for WORM and future encryption support.
- CDI is required for support of Persistent SCSI Reserve/Release.
- CDI will be required for any future enhancement of device management.

- The NetWorker software provides additional command line tools for diagnostics and control of tape devices that are available only when CDI is enabled. The man pages or command line reference guide provide more information on commands starting with “cdi\_\*”.

When CDI is enabled, use NetWorker release 7.4 SP4 or later when remote storage nodes are present.

## Enabling CDI

In cases where appropriate device drivers are in use, enable CDI (set it to “SCSI commands”) as it enables NetWorker to use a more advanced interface when communicating with tape devices.

The use of CDI is standard in both device tests and Reliability and Availability tests performed by EMC Quality Assurance. It has also been successfully tested with Virtual Tape Library systems like EMC EDL 3.x and EMC DL3D 1.01 or later.

File-type devices (both FTD and AFTD), NDMP-attached devices, and optical devices do *not* use CDI.

The use of CDI is *not* recommended for logical devices presented by AlphaStor.

For all devices of the same type that are present and configured on a specific host, the CDI setting should be the same. It is *not* supported to have CDI enabled for some devices of a specific type and disabled for others of the same type. It is supported to have CDI enabled only for certain specific device types and not others; for example, enabled for new LTO-4 devices, but disabled for older DLT devices.

It is supported to have CDI enabled only on certain hosts while disabled on the others, even in cases where the underlying device is the same (for example when DDS is used). However, this is not the case when NetWorker support for SCSI reserve/release is used; in this case, the CDI setting must be the same for all instances of the shared device.

## Persistent binding and naming

Enabling persistent binding and naming for tape libraries and tape devices is the recommended best practice to prevent device reordering on reboots or plug and play events. If a device reordering occurs, the NetWorker software is unable to use any affected drives until the configuration is manually corrected.

### Persistent binding

Persistent binding is used to ensure that the operating system of a server always sees SAN-presented devices with the same SCSI target ID across reboots by statically mapping a target's WWN address to a desired SCSI address. On some operating systems, this is done by default, while on others it has to be set manually.

In most cases, it is required that persistent binding is also set on the Fibre Channel HBA level through the configuration utility that comes with the fibre channel Host Bus Adapter (HBA). The HBA device driver documentation provides details.

Persistent binding is required for consistent library operations as NetWorker communicates with the library controller over a SCSI address that is chosen during initial library configuration.

If the SCSI address changes, the library becomes unavailable. In this case, disable the library and change the "control port" address to reflect the new SCSI address of the library controller.

### Persistent naming

Persistent naming is used to ensure that the operating system or device driver of a server always creates and uses the same symbolic path for a device (referred to as device file).

Device filenames created as a result of persistent naming depend on the operating system and device drivers used to enable and configure tape devices. The NetWorker software has extensive support for generic device drivers provided with operating systems.

Once persistently named device files are created and present on the host, enable the "use persistent names" option when scanning for tape devices from the NetWorker Management Console (NMC), and perform a re-scan of devices followed by a reconfiguration of the tape library. It is recommended to delete existing devices from the NetWorker configuration before attempting to reconfigure the library.

Details on how to configure persistent naming from the operating system or device driver side is documented in the device configuration best practices section, as follows. For more information on how to configure persistent binding and naming, contact the operating system and device driver vendor.

## Automatic library and device configuration

The NetWorker software provides the ability to scan, detect, and configure library and tape devices directly from the NetWorker Management Console (NMC) GUI. This is the recommended way to configure libraries, and manual configuration through the **jbconfig** or **jbedit** command line tools should be used only in special cases.

Scanning of devices, either from the NMC GUI or through the **inquire** command line tool should *not* be performed while the devices are in use. This applies to both of the following:

- Any device on the host that is being scanned and is currently in use (actively performing backup or restore operations).
- Any device that is shared with a different host and is in use on that host.

The reason for this is that on some operating system platforms, initiating a scan can cause SCSI resets for devices that are currently in use, leading to aborted or unusable backups.

In cases where vendor-provided device drivers are used, NetWorker autodetect operations might not function correctly. This is due to the fact that some vendors do not adhere to operating system implementation guidelines and often implement their own code for common functions, such as creating and assigning a device filename for a device. The result is that autodiscovery and configuration of devices from the NMC GUI does not work. Devices can still be configured through manual procedures (**jbconfig** or **jbedit**). For such reasons, it is recommended to use native operating system device drivers when possible, instead of vendor-provided device drivers. The use of vendor-provided drivers does not limit NetWorker support for core device functionality but does impair additional logic around automatic handling of such devices.

Autodiscovery and configuration of devices is also not possible when the system has multiple different device drivers using separate logic for assignment of device filenames. In such cases, either ensure that there is only one device driver in use for all the tape devices or perform library configuration manually through **jbconfig** or **jbedit**.

## Tape drives in virtualized environments

Although most virtualization systems enable the configuration of tape drives inside a guest container, the use of such configured tape drives has limited support.

For any virtualization system that does not provide dedicated HBA capabilities, the use of tape drives in guest containers is *not* supported for production environments. Such configurations can be used for testing or evaluation but never to store production backups. The reason is that virtualized or passthrough drivers have high interruption rates, are not suitable for sequential input/output, and can cause a high number of SCSI resets, even in the best cases.

At the moment, this includes all editions of:

- Microsoft Hyper-V
- Linux Xen

The use of tape drives is qualified and considered safe for virtualization systems that enable full isolation of input/output. The physical input/output interface *cannot* be shared between guests.

At the moment, systems that have required capabilities are:

- Solaris LDom
- Solaris Zones
- IBM LPar
- HP VPar/NPar
- VMware

In all cases, the supported solution requires the use of a dedicated HBA for tape drive connectivity inside a guest system. The use of shared HBAs in any of the preceding scenarios is *not* supported.

The technical limitation of tape drive connectivity inside virtualized systems does not exist for pure backup-to-disk solutions (using an advanced file type device or DD Boost device).

The *EMC Information Protection Software Compatibility Guide* provides the current list of supported platforms.

### **Tape drives in VMware guest system**

Support for the use of tape drives in a VMware Virtual Machine (VM) requires specific compatible hardware and VMware versions and must be configured through VMDirectPath.

VMDirectPath I/O enables a guest operating system on a VM to directly access the following:

- ♦ Physical PCI and PCIe devices connected to a host in a hardware passthrough mode.
- ♦ An I/O device, bypassing the virtualization layer.

Any other use of tape drives in a VM is *not* supported.

This applies to both physical and virtual tape drives.

VMDirectPath I/O is a vSphere function that requires the following:

- ♦ Compatible I/O MMU on the host. Compatible I/O MMU can be found on servers equipped with Intel VT-d technology (for example, Intel Xeon Nehalem 55xx or later CPUs) or AMD IOMMU technology (for example, AMD Opteron 41xx or later CPUs).
- ♦ VMware vSphere 4.0 Update 2, at a minimum.
- ♦ Use of a compatible HBA. The VMware hardware compatibility list provides a list of supported HBAs.
- ♦ A maximum of two HBAs per VM.

The use of VMDirectPath is incompatible with the following VMWare features:

- ♦ VMotion
- ♦ Fault tolerance
- ♦ Snapshots and VM suspend
- ♦ Hot adding and removing of virtual devices

**Note:** SAN and Hotadd transport modes are *not* supported. Only the network block device (nbd) transport mode is supported.

## Shoe-shining

An effect known as shoe-shining might occur during read/write operations if the data transfer rate falls below the minimum threshold at which the tape drive heads are designed to transfer data to or from a continuously running tape. When the transfer rate becomes too low and streaming is no longer possible, the drive must decelerate and stop the tape, rewind it a short distance, restart it, position back to the point at which streaming stopped, and then resume the operation. The resulting back-and-forth tape motion resembles that of shining shoes with a cloth.

When shoe-shining occurs, it significantly affects the attainable data rate, drive life, and tape life.

The minimum speed at which shoe-shining occurs depends on the type of tape drive.

For example, to prevent shoe-shining on LTO-5 tape drives that can write at a maximum speed of 160 MB/s, the minimum speed should be 40 MB/s (25% of the maximum speed).

## Device ordering

Device ordering issues arise when an operating system assigns a different device filename for a previously configured device. As a result, tape drives configured in NetWorker no longer match the actual state of the operating system, and NetWorker is no longer able to properly use the tape device.

Proper configuration of persistent binding and persistent naming resolves issues regarding device ordering, as the operating system always assigns the same device file regardless of any external events.

The NetWorker software provides a separate command line tool that can check the availability of all configured devices in the NetWorker database, but can be instructed to check only jukeboxes, only stand-alone drives, only local devices or devices belonging to a specific host. The Command Reference Guide or the man page for the **jbverify** command provides details.

Starting with NetWorker 7.4 SP3, functionality has been added to check the serial number of the device before issuing any mount requests. If the serial number of the device does not match the serial number marked during configuration, any tape operation on that device is aborted and an error message is posted.

This functionality is automatically enabled for all installations of NetWorker 7.4 SP3 and later with CDI enabled. However, for upgrade scenarios where a previous version of NetWorker was already present on the server, a reconfiguration of the tape library is needed. This is due to the fact that an initial scan of drive serial numbers is performed during library configuration.

Additionally, a reconfiguration of the library to scan for device serial numbers must be performed after changing the device's CDI attribute from Not Used to SCSI Commands.

Until the library reconfiguration is performed, tape devices work normally, but without new functionality.



**Note:** Upgrading the firmware of a library or a device drive on either a physical tape library or virtual tape library can result in a change to the serial number of a tape drive. It is recommended to perform library rescan/reconfiguration following firmware updates.

## Multipathing

Multipathing is a technique that uses more than one physical path between the server and target device.

Multipathing can be configured for tape devices, but is not directly utilized:

- Most operating systems will maintain separate device files for the same device visible over different paths. In such cases, for each path failure, the device must be reconfigured in NetWorker.
- Automatic path failover is not possible for sequential input/output used by tape drives. Thus, any backup or restore in progress is aborted regardless of the multipathing solution in place.

## SAN best practices

The following are recommendations for best practices in SAN environments:

- Use separate zoning for tape devices. Zones should *not* mix disk and tape devices.
- Use single initiator zoning. Each SAN zone should contain only one HBA, but it can contain a number of targets (tape devices).
- Avoid using 1 Gbit interfaces. Fiber channel protocol used for 1 Gbit interfaces (FCP-1) is *not* tape-safe.
- Avoid using multi-vendor switches in a single SAN cloud. Combining switching equipment from different vendors causes equipment to run in a compatibility mode that does not include FCP-2 tape-safe extensions.
- Use dedicated HBAs for tape devices due to the inherent differences in data streams (short blocks versus sequential).
- Limit the zoning of tape devices to only systems that actively use them. Limiting the size of each zone reduces the chances of external interferences and enables quicker diagnostics in case of any issues.

- When planning new environments or expanding existing ones, take into account all the parts of the data path to avoid bottlenecks. For example, zoning multiple high-speed devices to a single HBA that cannot sustain such data transfer is *not* recommended. Also, having multiple high-speed HBAs in a single system and connected on the same bus results in a bottleneck inside the system bus itself.
- Avoid oversubscription of interswitch links, which can lead to input/output delays and potential timeouts. When applicable, enable congestion control on a fiber channel.
- When performing SAN zoning changes, avoid changes for zones in which the devices are active (performing backup or restore). Changing zones while data transfer is active can interrupt the data transfer, thus causing the backup or restore operation to fail.

## Device block sizes

Block size indicates the size of data on a single transfer to the physical tape. This is dependent on the tape device, SCSI controller, and SCSI controller driver. Default block size varies on different platforms, and as a result, data actually written on tape might be in a format that cannot be easily interpreted by another host. The actual block size is assigned to a tape during the labeling process and marked in the header of the tape.

The NetWorker software has the ability to control block sizes, but only if it is allowed to do so by the operating system. In cases where the operating system assigns a specific value for a block size, NetWorker settings no longer have any impact.

It is strongly recommended that the operating system configuration is set to use a variable block size, thus allowing application control over it. Configuring a variable block size is achieved by setting it to a value of 0. For specific information about configuring a variable block size, consult the operating system documentation.

To check if the operating system is currently configured to allow for an application-controlled block size, use the **cdi\_block\_limits** command line tool:

```
# cdi_block_limits -f /dev/rmt/0cbn
maximum block size allowed is 16776128
minimum block size allowed is 1024
```

For heterogeneous environments where access to tapes is done from different platforms, it is recommended to set the block size in NetWorker to a uniform value across all configured devices. This can be done from

the NMC GUI under advanced settings for each device, or globally for a media type family by setting the `NSR_DEV_BLOCK_SIZE_<media type>` environment variable before the NetWorker server is started:

```
NSR_DEV_BLOCK_SIZE_<media type>=value_in_KB
```

**Example:**

To force a 256 KB block size for an LTO-4 device type on a Solaris platform, modify the `/etc/init.d/networker` script to include the following lines near the beginning of the script:

```
NSR_DEV_BLOCK_SIZE_LTO_ULTRIUM_4=256
export NSR_DEV_BLOCK_SIZE_LTO_ULTRIUM_4
```

To check the volume block size for the currently mounted volume, read the value directly from the NetWorker server.

**Note:** Forcing larger block sizes has a positive performance impact on high-speed tape drives, but it can also waste space on the tape.

Testing has shown the optimal value for LTO 1-3 to be 128 KB, and for LTO-4 to be 256 KB.

After a volume is loaded in a tape drive, the current volume block size and desired block size are noted in NetWorker for reference:

```
nsradmin> show name; volume block size; device block size
nsradmin> print type: nsr device
name: /dev/rmt/0cbn
volume block size: 64 KB;
device block size: handler default;
```

This example shows that the volume was labeled with a 64 KB block size, and the handler default value indicates that NetWorker does *not* attempt to force the block size during label operations but leaves it as the default value.

## Examples of NetWorker behavior

**Example 1:**

- Host A uses a block size of 32 KB and writes the data to tape.
- Host B uses a block size of 64 KB and is later used to read the data from tape.

In this case, NetWorker issues a warning and internally disables some functionality, such as fast positioning, but the restore succeeds.

**Example 2:**

- Host A uses a block size of 64 KB and writes the data to tape.
- Host B uses a block size of 32 KB and is later used to read the data from tape.

In this case, the restore fails because NetWorker is unable to read the original data blocks. An error message such as "Failed to read 65536 byte block with 32768 byte transfer" is displayed.

**Example 3:**

- Host A uses a block size of 64 KB and starts writing the data to tape.
- Host B uses a block size of 32 KB and later appends data writing to the same tape.

This creates a tape with a mixed block size, which creates a scenario that cannot be fully detected unless the **scanner** command line tool is used to analyze the data on the tape. In this case, the restore succeeds from host A (as explained under **Example 1**), but fails from host B (as explained in **Example 2**).

**Default block sizes**

Default block sizes used by NetWorker (handler default):

3480 = 32 KB	9490 = 64 KB
3570 = 256 KB	9840 = 256 KB
3590 = 384 KB	9940 = 256 KB
3592 = 384 KB	CentricStor = 256 KB
4890 = 32 KB	LTO Accelis = 128 KB
4mm = 32 KB	LTO Ultrium = 64 KB
8mm = 32 KB	LTO Ultrium-2 = 64 KB
8mm AIT = 64 KB	LTO Ultrium-3 = 128 KB
8mm AIT-2 = 64 KB	LTO Ultrium-4 = 128 KB
8mm AIT-3 = 192 KB	SAIT-1 = 64 KB
8mm AIT-4 = 128 KB	SAIT-2 = 64 KB
8mm AIT-5 = 128 KB	SD3 = 256 KB
8mm Mammoth-2 = 192 KB	SLR = 64 KB
8mm Mammoth-3 = 256 KB	dtf = 384 KB
himt = 32 KB	file = 32 KB
magnetic = 32 KB	qic = 32 KB
optical = 32 KB	short = 32 KB
vhs = 32 KB	tz85-88 = 32 KB

tzs20 = 64 KB	tz89-90 = 256 KB
travan10 = 32 KB	DLT 20GB = 32 KB
T10000 = 256 KB	DLT = 96 KB
TS1120 = 384 KB	DLT vs160 = 128 KB
TS1130 = 384 KB	DLT-s4 = 128 KB
VXA = 64 KB	DLT-v4 = 128 KB
adv_file = 128 KB	DLT1 = 32 KB
dst (NT) = 992 KB	DLT7000 = 128 KB
dst = 1168 KB	DLT8000 = 96 KB
tk50 = 32 KB	SDLT = 128 KB
tk70 = 32 KB	tkz90 = 384 KB

**Note:** Unless otherwise specified, different generations of the same technology use the same block size.

## LTO tape drive firmware issues

Ensure that HP LTO tape drives have the required firmware:

- ◆ HP LTO3 drives require M6BS or later firmware.  
**Note:** HP LTO3 tape drive firmware M66S and earlier contains a critical defect.
- ◆ HP LTO4 drives require H5AS or later firmware.  
**Note:** HP LTO4 tape drive firmware H58S and earlier contains a critical defect.

Failure to apply the correct hardware results in a tape drive sometimes ignoring the “move forward” commands issued by the backup application, which causes overwritten data and immediate data loss.

**Note:** The same tape drives are shipped in Oracle StorageTek configurations and require an identical firmware upgrade from Oracle.

## Hardware-based encryption

The use of LTO-4 hardware-based encryption is supported by NetWorker when controlled by management utilities that are provided with the LTO-4 hardware or by third-party key management software.

EMC does *not* test or certify these key management utilities. However, the NetWorker application can read from and write to LTO-4 devices

that use hardware-based encryption.

The use of this encryption is transparent to NetWorker. Neither the encryption nor the key management process is managed by the NetWorker application. This includes the ability to turn encryption on and off within NetWorker and the management of encryption keys.

## Device configuration best practices

Following are some of the recommendations regarding LTO (Gen.1 through Gen.4) device configuration for different operating systems.

Since the different setups can also work efficiently, this should be considered as best practice and not as a strict requirement for NetWorker operations. Recommendations refer to both physical drives and virtual drives as presented by virtual tape libraries.

This information is primarily based on built-in operating system drivers or IBM vendor-provided drivers (original or OEM). In the case of HP, Quantum, or other vendor-based configurations, only operating system-related sections apply.

The NetWorker software requires a functional device file created on the operating system. Setup of the driver itself is *not* supported by EMC, and the following information is given as a best practice.

## AIX

- IBM AIX can use either built-in operating system tape drivers or external IBM Atape drivers that are installed as a separate package. Advanced capabilities of modern tape drives are exposed only when using Atape drivers as the built-in drivers with a reduced feature set.
- The minimum recommended version of Atape driver is 9.7.5.0. The recommended driver is from the v11 family, and the latest test version is 11.1.3.0. A new driver can be installed and reloaded online *without* a system reboot.

Avoid using drivers from the Atape v10 family due to multiple issues. Use either the suggested v9 or v11 drivers.

- If the AIX dynamic tracking feature is set to enabled, then the Atape device driver does *not* report any external FC/SCSI resets to the operating system, thus causing a tape rewind during backup without notification.
- AIX automatically keeps device files reserved, so no additional steps are needed unless device files under /dev are manually removed.
- If device files are deleted and re-created, ensure that device ordering within NetWorker is correct.
- A new feature in newer versions of the Atape driver known as dynamic tracking can cause stability issues due to corruption of SCSI ioctl messages, resulting in SCSI resets. At this time, the Atape Dynamic Tracking option should be set to DISABLE.
- The AIX native SMC driver interferes with direct SCSI communication to a changer device. If an AIX system is used for access to the library, it should be disabled.

Use the following command to check the SMC status:

```
# lsdev -Cc tape | grep smc
```

The status should be Defined and not Available. If the status is Available, disable it for each instance of smc by running the following command on each reboot:

```
# rmdev -l smc0
```

If drives are shared (used in the DDS environment), the operating system-level SCSI reserve/release must be disabled. For generic devices, the field name is res\_support, and for newer IBM LTO drives, the field name is retain\_reserve.

To test, run the following command:

```
# lsattr -El rmt0.1 | grep res_support
```

This should be set to No (ignore True at the end of the line). To disable it, either use the SMIT tool under the device config section or use the following command:

```
# /usr/sbin/chdev -l rmt0 -a <field>=no
```

## Solaris

- Solaris provides tape device support through the generic SCSI tape device driver (st). For the minimum revision of the st device driver, refer to the specific Solaris version.
- If an advanced device configuration is required, for example, to force hardware compression or other device-specific options, add the device vendor-provided entries to the /etc/st.conf configuration file and reboot the host system.
- IBM provides a device driver for Solaris (IBMTape, which creates stcbn devices), but EMC does *not* recommend its use at this time as it lacks a full implementation of the required API. CDI should be disabled if IBM drivers are used on Solaris due to known issues with IBM device drivers.
- For Emulex HBAs, use one of the following:
  - ◆ Emulex original driver (lpfc stack)
  - ◆ Older Sun-branded driver, based on the Emulex driver
  - ◆ Newer Sun Leadville class FC driver stack (emlxs)

**Note:** The Sun Leadville class stack (emlxs) is recommended because it integrates best with the Solaris operating system, including proper propagation and rediscovery on unplanned fabric events.
- If the Emulex original driver is used (lpfc), refer to <http://www.emulex.com/support/solaris/lpfc.jsp> for additional required patches and information.



- If the fiber channel stack used is different than Sun Leadville, automatic device detection in NetWorker might not function since fiber channel-connected devices are not properly exposed to the Solaris operating system.

In such cases, refer to the `lus_add_fp_devs` command, and if necessary, manually modify `/usr/kernel/drv/lus.conf`.

- Automatic detection of devices is *not* needed for manual configuration through the `jbconfig` or `jbedit` command.
- Solaris automatically keeps device files reserved, so no additional steps are needed unless device files under `/dev/rmt` are manually removed.
- If device files are deleted and re-created or a system reconfiguration was forced, ensure that device ordering in NetWorker is correct.
- For an example of configuring device file persistency after device files have been re-created, see [http://docs.sun.com/source/819-0139/ch\\_8\\_persist\\_binding.html](http://docs.sun.com/source/819-0139/ch_8_persist_binding.html).
- If large slowdowns are noticed with transfer rates dropping to less than 500 KB/sec, add the following values to the beginning of the NetWorker start script (typically `/etc/init.d/networker`) to disable internal threading by the Solaris kernel when asynchronous I/O is attempted:

```
DISABLE_SOL_ASYNC_IO=yes
export DISABLE_SOL_ASYNC_IO
```

## Solaris 8

- Patch ID 108725 references the Solaris generic st (SCSI tape) kernel module. The minimum required revision is -21.
- If the appropriate patch version is installed, there should *not* be any additional entries in `/etc/st.conf` on the system.
- A patch install can be done online with the only requirement being a reload of the st kernel module. A reboot is *not* required.
- To check the current patch level, use the following command:

```
# showrev -p | grep 'Patch: 108725'
```

## Solaris 9

- ♦ Patch ID 113277 references the Solaris generic st (SCSI tape) kernel module. The minimum required version is -35.

**Note:** This patch also includes the sd (SCSI disk) kernel module.

Therefore, a reboot is required upon installation.

- To check the current patch level, use the following command:  

```
# showrev -p | grep 'Patch: 113277''
```
- If the optional Solaris package “Storage Foundation suite” is installed, it adds SCSI generic (sgen) support to Solaris. In that case, use of the lus driver (Legato User SCSI) provided by EMC is no longer supported, and library control is done through the sgen driver. As a result, NetWorker low-level SCSI jukebox commands (**sjj**) *cannot* be interactively used while the library is configured and enabled in NetWorker.

### Solaris 10

- The minimum revision of the SCSI tape kernel module included in Solaris 10 includes native LTO support. No additional patches are required.
- Library connectivity on Solaris 10 no longer uses the lus driver (Legato User SCSI). Instead, the Solaris SCSI generic (sgen) interface is used.
- As a result, NetWorker low-level SCSI jukebox commands (**sjj**) *cannot* be interactively used while the library is configured and enabled in NetWorker.

### Windows

The following sections describe device configuration issues for Microsoft Windows platforms. Recommendations on settings and device drivers differ depending on the version of the Microsoft Windows operating system, as described in the following sections.

## Windows 2000

- The use of vendor-provided drivers is recommended as Microsoft built-in support for tape drives is not fully present in Windows 2000.
- For IBM LTO drives, the latest available version of IBM Atape drivers for Windows 2000 is 6.1.4.8.

As IBM has stopped development for the Windows 2000 platform, those drivers are lacking critical patches found in newer versions. Their use is safe on stand-alone systems, but avoid using DDS when possible as it can cause SCSI reservation issues.

- The use of CDI is *not* recommended on Windows 2000 as the latest available tape device drivers are not up-to-date.
- When installing Atape drivers:
  - ♦ If upgrading an Atape driver, uninstall the existing version prior to upgrading.
  - ♦ The driver selected during installation should be a non-exclusive version. The NetWorker software does *not* support the exclusive mode.
- The use of CDI is *not* recommended on Windows 2000 as the latest available tape device drivers from both Microsoft and drive vendors are not up-to-date.
- For Windows 2000 platforms, there is no operating system support for persistent naming, TUR, or any other advanced features.
- The Windows Removable Storage Manager service must be disabled at the operating system layer.

## Windows 2003

- Microsoft provides a default built-in device driver for LTO tape drives. Microsoft built-in drivers are generally sufficient, and the use of vendor-provided drivers is *not* needed unless explicitly required.
- Windows Server 2003 does *not* include native support for LTO-3 and LTO-4 tape drives. Either use vendor-provided drivers or install a Microsoft update.

For IBM LTO tape drives, use Microsoft update 941140:

<http://support.microsoft.com/kb/941140>

For Quantum LTO tape drives, use Microsoft update 971649:

<http://support.microsoft.com/kb/971649>

- Windows Server 2003 does *not* include native support for LTO-5 tape

drives. Either use vendor-provided drivers or install a Microsoft update.

For IBM LTO-5 tape drives, use Microsoft update 971650:

<http://support.microsoft.com/kb/971650>

For Quantum LTO-5 tape drives, use Microsoft update 971649:

<http://support.microsoft.com/kb/971649>

- Microsoft Windows 2003 Service Pack 2 addresses many problems with tape device access, so it is highly recommended for NetWorker server and storage node operations.
- For IBM LTO drives, the use of IBM Atape drivers is supported. For versions of IBM Atape device drivers earlier than 6.1.8.4, the use of CDI is *not* recommended. Additionally, if the data path failover (DPF) feature is enabled on the device driver side when using older drivers, it can potentially lead to data loss.
  - ◆ If upgrading an IBM Atape driver, uninstall the existing version prior to upgrading.
  - ◆ The IBM Atape driver selected during installation should be a non-exclusive version. The NetWorker software does *not* support the exclusive mode.
  - ◆ IBM has changed installer options in different versions of device drivers. The README file distributed with the IBM drivers provides details on how to install drivers in non-exclusive mode.
- For HP LTO drives, the use of HP drivers is supported. HP LTO drivers are part of the HP StorageWorks Library and Tape Tools software download.
- It is recommended to enable persistent naming at the operating system layer to stop the operating system from changing device file order on reboots. After completing the changes, a reboot is required. All tape drives should appear with a long name, and the device configuration should use those names.
- ◆ For Microsoft and other vendor-provided device drivers, including Quantum, HP, and Dell but excluding IBM Atape, set the registry key:

```
HKLM\System\CurrentControlSet\Control\Tape\Persistence:  
DWORD=1
```

The Microsoft Knowledge Base article 873337 provides details.

This registry key must *not* coexist with the registry key set for persistent naming of IBM Atape drivers, described as follows.

For systems that use Windows 2003 Service Pack 1, a Microsoft hotfix documented in Knowledge Base article 933682 is also needed. Systems earlier than Windows 2003 Service Pack 1 do *not* have the persistent naming functionality.

The NetWorker software is automatically able to configure newly created persistent names, starting with NetWorker 7.4 SP2. For earlier versions, manual configuration through the **jbconfig** or **jbedit** command is required.

- IBM Atape drivers (ibmtp2k3) do *not* rely on the Microsoft implementation of persistent naming, but require their own registry setting:

```
HKLM\System\CurrentControlSet\Services\ibmtp2k3\Persistent
Naming:DWORD=1
```

The IBM Tape Device Drivers Installation and User's Guide provides details, including the required changes to the system registry, at:

[ftp://ftp.software.ibm.com/storage/devdvr/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdvr/Doc/IBM_Tape_Driver_IUG.pdf)

The registry key must *not* coexist with the registry key set for persistent naming of Microsoft drivers, described earlier.

NetWorker 7.4 SP4 or later is required for automatic configuration of tape drives with newly created persistent names. For earlier versions of NetWorker, perform manual configuration through the **jbconfig** or **jbedit** command.

- Disable Tape Unit Ready (TUR) at the operating system layer, as described in Microsoft Knowledge Base article 842411. This is required for NetWorker operations.
- It is no longer required to disable the Windows Removable Storage Manager (RSM) service at the operating system layer. However, RSM must be prevented from accessing tape drives that are under NetWorker control. To achieve this, start the RSM manager and remove all tape drives from the RSM configuration so it does not attempt to check device status at regular intervals, which causes unwanted SCSI resets.

## Windows 2008

- Microsoft provides a default built-in device driver for LTO tape drives. Microsoft built-in drivers are generally sufficient. The use of vendor-provided drivers is not needed unless explicitly required.

- Windows Server 2008 does *not* include native support for LTO-5 tape drives. Either use vendor-provided drivers or install the Microsoft update.

For IBM LTO-5 tape drives, use Microsoft update 971650:

<http://support.microsoft.com/kb/971650>

For Quantum LTO-5 tape drives, use Microsoft update 971649:

<http://support.microsoft.com/kb/971649>

- For IBM LTO drives, the use of IBM Atape drivers is supported. The minimum required version of IBM Atape drivers for Windows 2008 is 6.1.9.5. When downloading drivers from the IBM website, use drivers marked with `_w08_`.
  - ◆ If upgrading an IBM Atape driver, uninstall the existing version prior to upgrading.
  - ◆ The IBM Atape driver selected during installation should be a non-exclusive version. The NetWorker software does *not* support the exclusive mode.
  - ◆ IBM has changed installer options in different versions of device drivers. The README file distributed with the IBM drivers provides details on how to install drivers in non-exclusive mode.
- For HP LTO drives, the use of HP drivers is supported. HP LTO drivers are part of the HP StorageWorks Library and Tape Tools software download.
- Configuration of persistent naming is recommended at the operating system layer to stop the operating system from changing device file order on reboots.

After completing the changes, a reboot is required. All tape drives should appear with a long name, and the device configuration should use those names.

- ◆ For Microsoft and other vendor-provided device drivers, including Quantum, HP, and Dell but excluding IBM Atape, set the registry key:

```
HKLM\System\CurrentControlSet\Control\Tape\Persistence:  
DWORD=1
```

The Microsoft Knowledge Base article 873337 provides details.

This registry key must *not* coexist with the registry key set for persistent naming of IBM Atape drivers, described as follows.

For systems that use Windows 2003 Service Pack 1, a Microsoft hotfix documented in Knowledge Base article 933682 is also needed. Systems earlier than Windows 2003 Service Pack 1 do *not* have the persistent naming functionality.

The NetWorker software is automatically able to configure newly created persistent names, starting with NetWorker 7.4 SP2. For earlier versions, manual configuration through the **jbconfig** or **jbedit** command is required.

- For IBM Atape drivers, in cases where IBM-provided drivers (ibmtp2k8) are used, automatic detection for persistent naming is currently *not* present and will be added in future versions of NetWorker. As a workaround, create registry keys required by the Atape driver for both Windows 2003 and Windows 2008:

```
HKLM\System\CurrentControlSet\Services\ibmtp2k3\Persistent
Naming:DWORD=1
```

```
HKLM\System\CurrentControlSet\Services\ibmtp2k8\Persistent
Naming:DWORD=1
```

The IBM Tape Device Drivers Installation and User's Guide provides details, including the required changes to the system registry, at:

[ftp://ftp.software.ibm.com/storage/devdvr/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdvr/Doc/IBM_Tape_Driver_IUG.pdf)

The registry key must *not* coexist with the registry key set for persistent naming of Microsoft drivers, as described earlier.

- Disable Tape Unit Ready (TUR) at the operating system layer, as described in Microsoft Knowledge Base article 842411. This is required for NetWorker operations. Although the Microsoft Knowledge Base article refers to Windows 2003, it currently also applies to Windows 2008.
- On Windows 2008, the Windows Removable Storage Manager (RSM) service is no longer installed by default. However, when RSM is present, it must be prevented from accessing tape drives that are under NetWorker control. To achieve this, start the RSM manager and remove all tape drives from the RSM configuration so it does not attempt to check device status at regular intervals, which causes unwanted SCSI resets.

## Linux

- The Linux `st` kernel module has native LTO support by using its

generic functionalities. IBM also provides an Atape driver for Linux, but EMC does *not* recommend its use at this time.

- The minimum recommended version of RedHat Enterprise Linux for use with CDI is RHEL4 Update 4, with kernel version 2.6.9-39. Earlier versions of RHEL do *not* provide reliable tape operations due to masking of SCSI resets and thus can lead to data loss.
- The minimum recommended version of Novell SUSE Linux Enterprise Server for use with CDI is SLES9 SP3, with kernel version 2.6.5-7.267. Earlier versions of SLES do *not* provide reliable tape operations due to masking of SCSI resets and thus can lead to data loss.
- The use of DDS is *not* recommended for Linux distributions earlier than the recommended levels.
- If advanced device configuration is required, for example, to force hardware compression or other device specific options, add device vendor-provided entries to the `/etc/stinit.def` configuration file and re-initialize the operating system tape device support by running the **stinit** command. This should *not* be done while devices are in use, either on this system or any other system where devices are shared.
- Ensure that the `sg` (SCSI Generic) kernel module is loaded before the `st` (SCSI Tape) kernel module. This is done automatically in recommended versions of Linux distributions.
- Linux provides a built-in HBA driver for both Qlogic and Emulex but with limited capabilities only, without monitoring or full plug-and-play functionalities. It is recommended to use vendor-provided drivers for HBAs.
- Current versions of drivers for multi-purpose adapters based on LSI Fusion-MPT technology do *not* provide full SCSI passthrough capabilities and CDI should be disabled.
- Logging of the following kernel messages (dmesg output) is due to the `lpfc` HBA driver not handling the tagged command queuing correctly:

```
kernel: st: Error: 8000002, cmd: a 0 2 0 0 0
kernel: st: Current: sense key: Illegal Request
```

In these cases, upgrade the HBA driver or disable the tagged command queuing. Contact the Linux vendor support for detailed instructions.



- Configuration of persistent naming is recommended at the operating system layer to stop the operating system from changing device file order on reboots.

Persistent naming is achieved by using udev rules that are controlled by `/etc/udev/udev.conf` and included folders or files.

- ◆ Newer versions of the udev subsystem already contain rules for tape devices, so no additional operating system configuration is needed.
- ◆ Check if the device file is already created under `/dev/tape/*`. In that case, simply enable “use persistent names” from NMC when scanning for devices or configuring libraries.
- ◆ The udev subsystem can be upgraded regardless of the version of the underlying Linux distribution, with the minimum kernel version 2.6.
- ◆ If the appropriate udev rule for tape devices is present on the system but device filenames are not being created, check if the system is configured to allow for scanning of all SCSI devices.
- ◆ The following is an example rule that creates entries such as `/dev/tape/by-id/x-xxxx-nst`:  

```
KERNEL=="nst*", ENV{ID_SERIAL}=="?* ", SYMLINK+="tape/by-id/${env{ID_BUS}}-${env{ID_SERIAL}}-nst
```

This is an example *only*. The documentation for the specific Linux distribution provides more details.

- ◆ To test udev rules, use the **udevtest** command with the device name from SysFS as a parameter:  

```
# find /sys -name nst0
# udevtest /sys/class/scsi_tape/nst0
```
- ◆ Any rule that is linked to a static property is a valid rule. To discover all linkable properties of a device, use the following commands (availability depends on the installed version of udev):  

```
# udevinfo -e
# udevinfo -a -p /sys/class/scsi_tape/nst0
```
- ◆ Avoid using rules that use external programs such as **scsi\_id**, **lsscsi**, or **scsiinfo** as they might cause SCSI resets during configuration.
- ◆ NetWorker 7.4 SP2 or later is required for automatic configuration of tape drives with newly created persistent names. For earlier versions of NetWorker, perform manual configuration through the **jbconfig** or **jbedit** command.

## HP-UX

- HP-UX performs persistent binding and naming on all devices by default. No special configuration is needed.
- HP-UX also provides a default built-in driver, STape. Alternatively, you can use the atdd driver provided by IBM.
- If the HP Enterprise Monitoring System (EMS) is enabled, ensure that the dm\_stape module is disabled as it causes SCSI resets when used with NetWorker. One way of doing so is to modify the file `/var/stm/config/tools/monitor/dm_stape.cfg` with the value `POLL_INTERVAL=0` to disable polling and restart EMS.

### HP STape specific information:

- The HP-UX STape driver disables the hardware compression capability of LTO drives.
- The HP-UX hotfix PHKL\_40389 tape cumulative patch is required if operating system-provided drivers are used.
- If the STape driver is used in cases where drives are shared (used in a DDS environment), the operating system-level SCSI reserve/release must be disabled.

Set the kernel tunable parameter:

```
st_san_safe=1
```

Also, run these two commands:

```
# scsimgr set_attr -d estape -a norewind_close_disabled=1  
# scsimgr save_attr -d estape -a norewind_close_disabled=1
```

**IBM atdd specific information:**

- For HP-UX 11 v3, the minimum required version of IBM atdd drivers to support CDI is v6.0.0.76.
- For HP-UX 11 v1 and v2, the IBM atdd drivers do *not* have the complete functionality required for CDI.
- When installing IBM atdd drivers on HP-UX, read the enclosed readme file for the proper procedure to prevent unwanted system crashes.

**Tru64**

- Support for all tape devices is provided through generic operating system built-in drivers with no additional configuration needed for advanced options, such as persistent naming.
- If advanced device configuration is required, for example, to force hardware compression or other device specific options, add device vendor-provided entries to the `/etc/ddr.dbase` configuration file and re-initialize the operating system device database by running the `ddr_config -c` command.

**Preventing SCSI resets**

In most cases, a SCSI reset is not indicative of faulty hardware. SCSI resets occur as part of the SCSI standard in cases where the target being accessed is currently in use (thus returning a busy status). To minimize the risk of unwanted SCSI resets, follow the recommendations in the SAN best practices section of this document.

A SCSI reset manifests in NetWorker differently, depending on whether the operating system itself detected the reset. If the operating system detected the reset, then the NetWorker software also detects it and proactively marks the volume currently in use as full to prevent any possible data loss.

If information on the SCSI reset is filtered on the device driver level, the operating system does *not* have knowledge of it. Thus, the application level, including NetWorker, is unable to act accordingly. In this case, the tape drive can perform a rewind without NetWorker knowing about it. This may cause data loss as older data blocks can be overwritten by the current backup stream.

## Primary causes of SCSI resets

The following are the primary causes of SCSI resets:

- Sharing violations between different hosts.
- Sharing violations between different applications on the same host.
- Fiber channel reconfiguration events in cases of added or removed targets inside the SAN zone (RSCN events).
- Fiber channel link renegotiations in cases of bad links. This is not limited to hard errors (loss of sync, CRC, and so on), but also includes a number of soft errors. For example, a high number of encoding errors causes a high retry count that can break the tape input/output.
- Fiber channel routing discovery in routed SAN environments (FSPF event).

Unwanted SCSI reset commands must be prevented at their originating point, rather than attempting to block it at the transport layer (for example, at the SAN switch), because the reset is a necessary part of the disk arbitration process, especially in cluster environments.

The use of any operating system or third-party vendor-provided monitoring system that attempts to monitor tape devices is *not* supported due to potential SCSI resets causing tapes marked prematurely full or even leading to data loss.

Do *not* use low-level operating system commands that attempt to communicate with a device directly while the device is in use by NetWorker. This includes commands such as **mt**.

## External tools

Operating system or third-party vendor-provided hardware monitoring tools often include the capability to poll all hardware devices in regular intervals as a part of an availability check.

If polling is performed on an active tape drive, a device driver can send an unwanted SCSI reset.

In all cases, polling of tape devices should be disabled without any exceptions.

The following is an example list of some of the known monitoring tools that include hardware polling capabilities:

- ◆ Sun SRS: Sun Resource System Monitor
- ◆ Sun SRC: Sun Resource Center

- ◆ Sun SMC: Sun Management Console
- ◆ Sun ESM: Enterprise SAN Manager
- ◆ Sun Explorer
- ◆ HP EMS: Enterprise Monitoring System
- ◆ HP Top Tools
- ◆ HP OpenView
- ◆ HP Compaq Insight Manager
- ◆ Fujitsu SAN InSite
- ◆ Computer Associates BrightStor SAN Manager
- ◆ Qlogic SANsurfer
- ◆ Dot Hill SANpath
- ◆ InControl
- ◆ EMC Grab for UNIX/Linux
- ◆ EMC Report for Windows

It is safe to use such tools only when there is no tape activity on the system.

### **Reserve/release**

Reserve/release is an attempt to prevent problems with SCSI resets by using existing SCSI commands to prevent more than one host or application from accessing any particular drive while it is in use by the NetWorker system.

Operating system-controlled SCSI reserve/release must be disabled under all circumstances. Most operating systems do not use it by default, but when they do, it interferes with NetWorker tape handling and can cause SCSI resets.

### **NetWorker configuration options**

NetWorker release 7.3 and later includes application-controlled handling of reserve/release, which can be used in environments that experience problems.

### Configuration options in NetWorker

The Reserve/Release attribute in the NetWorker Device resource has the following options:

- None – Default value.
- Simple – Uses old style SCSI reserve and release commands.
- Persistent – Uses the newer Persistent Reserve Out command with a reservation key specified in the Persistent Reserve Key attribute.
- Persistent+APTPL – Same as Persistent *except* the Active Persist Through Power Loss bit is set.

To configure application-controlled reserve/release from the NetWorker software, you must enable CDI on NetWorker. Operating system-controlled reserve/release must be disabled for NetWorker to be able to control it.

Persistent reserve/release requires a compatible tape drive and device driver that understands such commands.

The man pages for the NetWorker commands **cdi\_pr**, **cdi\_reserve**, and **cdi\_release** provide details on manual control of reserve/release. The *EMC NetWorker Administration Guide* provides more information on configuring reserve/release features in NetWorker.

Copyright © 2014 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.