

Configuring Tape Devices for EMC NetWorker

Technical Note

P/N 300-008-352

REV A05

January 2010

This technical note describes best practices to configure tape devices for 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 & device configuration.....	4
♦ Tape drives in virtualized environments.....	5
♦ Device ordering.....	6
♦ Multipathing.....	7
♦ SAN best practices	7
♦ Device block sizes	8
♦ Hardware-based encryption.....	11
♦ Device configuration best practices.....	11
♦ Preventing SCSI resets.....	23

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 drive 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 pass-through commands were passed to it.

Advantages of CDI

The following are advantages to using CDI:

- ♦ CDI enables the use of 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 allows on-demand auto-cleaning 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 drive.
- ♦ 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 (new feature in NetWorker release 7.4 Service Pack 3).
- ♦ CDI is required for WORM and future encryption support.
- ♦ CDI is required for support of Persistent SCSI Reserve/Release.
- ♦ CDI is required for any future enhancement of device management.
- ♦ NetWorker provides additional command line tools for diagnostics and control of tape devices which are available only when CDI is enabled. For reference on the commands, see man pages or command line reference guide for commands starting with "cdi_*".

When CDI is enabled, use NetWorker release 7.4 Service Pack 4 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 allows 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 higher.

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.

All devices of the same type that are present and configured on the 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 type 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 avoid device reordering on reboots or plug and play events. If a device reordering occurs, the NetWorker software is not able to use any affected drives until the configuration is manually corrected.

Persistent binding

Persistent binding is used to ensure 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 using the configuration utility that comes with the fibre channel Host Bus Adapter (HBA). For details refer to the HBA device driver documentation.

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 iscome 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 file names created as a result of persistent naming depend on the operating system and device drivers used to enable and configure tape devices. NetWorker 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, below. For further information on how to configure persistent binding and naming, please contact the operating system and device driver vendor.

Automatic library & device configuration

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

Scanning of devices, either from NMC GUI or using `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 which is being scanned is currently in use (actively performing backup or restore operations)
- ♦ Any device is shared with a different host and in use on that host.

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

In cases where vendor-provided device drivers are used, NetWorker auto-detect operations may not function correctly. This is due to 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 file name for a device. The result of this is that auto-discovery and configuration of devices from the NetWorker Management Console GUI does not work. Devices can still be configured using manual procedures (jbconfig or jbedit). For such reasons, when possible it is recommended to use native operating system device drivers instead of vendor provided device drivers. Using vendor-provided drivers does not limit NetWorker support for core device functionality, but does impair additional logic around automatic handling of such devices.

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

Tape drives in virtualized environments

Although most virtualization systems allow configuration of tape drives inside a guest container, the use of such configured tape drivers 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 pass-through drivers have high interruption rate and as such are not suitable for sequential I/O and can cause high number of SCSI resets even in best cases.

At the moment, this includes all editions of:

- VMWare

- Microsoft Hyper-V
- Sun Solaris Zones
- Linux Xen.

The use of tape drives is qualified and considered safe for virtualization systems which allow for full isolation of I/O. Note that the physical I/O interface cannot be shared between guests.

At the moment, systems which have required capabilities are:

- Solaris LDom
- IBM LPar
- HP VPar/NPar

This technical limitation does not exist for pure backup-to-disk solutions (using an Advanced File Type Device), and support is based on qualification efforts. For the current list of supported platforms, see the EMC Information Protection Software Compatibility Guide.

Device ordering

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

NetWorker provides a separate command line tool that can check 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. For details, see the Command Reference Guide or the man page for the `jbverify` command.

Starting with NetWorker 7.4 Service Pack 3, 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 error message posted.

This functionality is automatically enabled for all installations of NetWorker 7.4 SP3 and higher with CDI enabled. However, for upgrade scenarios where previous version of NetWorker was already present on the server, a reconfiguration of the tape library is needed. This is due to

fact that initial scan of drive serial numbers is performed during library configuration.

Additionally, a reconfiguration of the library must be performed 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 that upgrading firmware of a library or a device drive on either physical tape library or virtual tape library can result in a change to 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 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 I/O 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 number of targets (tape devices).
- ♦ Avoid using 1Gbit interfaces. Fiber channel protocol used for 1Gbit 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 compatibility mode which does not include FCP-2 tape safe extensions.

- ♦ Use dedicated HBAs for tape devices due to inherent differences in data streams (short blocks .vs. sequential).
- ♦ Limit zoning of tape devices to only systems which actively use them. Limiting the size of each zone reduces chances of external interferences and allows for quicker diagnostics in case of any issues.
- ♦ When planning new environments or expanding existing ones, take into account all parts of the data path to avoid bottlenecks. For example, zoning multiple high-speed devices to a single HBA which 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 over-subscription of inter-switch links, which can lead to I/O delays and potential timeouts. When applicable, enable congestion control on fibre 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 may 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.

NetWorker has the ability to control block sizes, but only if it is allowed to do so by 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 variable block size, thus allowing application control over it. Configuring variable block size is achieved by setting it to a value of 0. For specific information about configuring variable block size, consult the operating system documentation.

To check if the operating system is currently configured to allow for 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 NetWorker server is started:

```
NSR_DEV_BLOCK_SIZE_<media type>=value_in_kb
```

Example: To force 256KB block size for LTO-4 device type on Solaris platform, modify the `/etc/init.d/networker` script to include 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 NetWorker server.

Note that forcing larger block sizes has positive performance impact on high-speed tape drives, but it can also waste space on the tape.

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

Once a volume is loaded in a tape drive, the current volume block size as well as desired block size is 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 64 KB block size, and the handler default value indicates that NetWorker does not attempt to force the block size during label operations but leave it as the default value.

Examples of NetWorker behavior

Example 1:

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

In this case, the restore fails because NetWorker is not able 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 block size of 64KB and starts writing the data to tape
- ♦ Host B uses block size of 32KB and later appends data writing to the same tape

This creates a tape with a mixed block size which creates a scenario which 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 fail 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
hmt = 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 same block size.

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

Below 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 is to be considered as best-practice and not as a strict requirement for NetWorker operations. Recommendations refer to both physical drives as well as 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.

NetWorker software requires a functional device file created on the

operating system. Setup of the driver itself is not supported by EMC and the information below is given as best-practice.

AIX

- ♦ IBM AIX can use either built-in OS tape drivers or external IBM Atape drivers which are installed as a separate package. Note that advanced capabilities of modern tape drives are exposed only when using Atape drivers as the built-in drivers have a reduced feature set.
- ♦ Minimum recommended version of Atape driver is 9.7.5.0, recommended is driver from v11 family, latest test version is 11.1.3.0. New driver can be installed and reloaded online without system reboot.

Avoid using drivers from Atape v10 family due to multiple issues (use either suggested v9 or v11 drivers)

- ♦ 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 recreated, make sure that device ordering within NetWorker is correct.
- ♦ New feature present in newer versions of Atape driver known as dynamic tracking is known to cause stability issues due to corruption of SCSI ioctl messages resulting in SCSI resets. At this time, Atape Dynamic Tracking option should be set to DISABLE.
- ♦ 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 SMC status:

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

Define the status as not Available. If the status is Available, disable it for each instance of smc, and reboot the computer:

```
# rmdev -l smc0
```

If drives are shared (used in DDS environment), OS-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 the status, run this 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 run this command:

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

Solaris

- ♦ Solaris provides tape device support using generic SCSI tape device driver (st). For minimum revision of st device driver, refer to specific Solaris version.
- ♦ If advanced device configuration is required (for example, forcing hardware compression or other device-specific options), add device vendor provided entries to /etc/st.conf configuration file and reboot the host system.
- ♦ IBM provides device driver for Solaris (IBMTape which creates stcbn devices), but its use is not recommended by EMC at this time as they lack 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 follow:
 - Emulex original driver (lpfc stack)
 - Older Sun-branded driver (based on Emulex driver)
 - Newer Sun Leadville class FC driver stack (emlxs).
- ♦ 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 Emulex original driver is used (lpfc), refer to <http://www.emulex.com/support/solaris/lpfc.jsp> for additional required patches and information.
- ♦ If the Fibre Channel stack used is different than Sun Leadville, automatic device detection in NetWorker may not function since Fibre Channel connected devices are not properly exposed to the Solaris operating system.

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

- ♦ Note that Automatic detection of devices is not needed for manual configuration via `jbconfig` or `jbedit`.
- ♦ 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 recreated or system reconfiguration was forced, make sure that device ordering in NetWorker is correct.

- ♦ For an example of configuring device file persistency after they have been recreated, see http://docs.sun.com/source/819-0139/ch_8_persist_binding.html.
- ♦ If large slowdowns are noticed (transfer rates drop below 500KB/sec), add following values to the beginning of NetWorker start script (typically /etc/init.d/networker) to disable internal threading done by Solaris kernel when Asynchronous I/O is attempted:

```
DISABLE_SOL_ASYNC_IO=yes
export DISABLE_SOL_ASYNC_IO
```

Solaris 8

- ♦ Patch ID 108725 references Solaris generic st (SCSI tape) kernel module. Minimum required revision is -21.
- ♦ Note that if the appropriate patch version is installed, there should not be any additional entries in /etc/st.conf present on the system.
- ♦ Patch install can be done online with only requirement of reload of st kernel module (no reboot required).
- ♦ To check the current patch level, use

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

Solaris 9

Patch ID 113277 references Solaris generic st (SCSI tape) kernel module. Minimum required version is -35. Note that as this patch also includes sd (SCSI disk) kernel module, therefore a reboot is required upon installation.

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

### Solaris 10

- ♦ The minimum revision of 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 (sji) cannot be interactively used while the library is configured and enabled in NetWorker.

## Windows

The following sections describe device configuration issues for Windows platforms. As recommendations on settings and device drivers differ depending on the version of the Windows operating system, refer to appropriate section below.

### Windows 2000

- ♦ Using 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 perfectly 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 latest available tape device drivers are not up-to-date.
- ♦ When installing Atape drivers:
  - If upgrading Atape driver, uninstall the existing version prior to upgrading.
  - The driver selected during installation should be a non-exclusive version. Exclusive mode is not supported by the NetWorker software.
- ♦ The use of CDI is not recommended on Windows 2000 as latest available tape device drivers from both Microsoft and drive vendors are not up-to-date.
- ♦ For Windows 2000 platform 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. When using of Microsoft's built-in drivers, Microsoft hotfix 941140 is required for full LTO-3 and LTO-4 support.

- ♦ Microsoft Windows 2003 Service Pack 2 addresses many problems regarding 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 below 6.1.9.4, the use of CDI is not recommended. Additionally, if 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. Exclusive mode is not supported by the NetWorker software.

IBM has changed installer options in different versions of device drivers. Refer to the README file distributed with the IBM drivers for details 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 part of "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. Furthermore, all tape drives should appear with 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
```

Details are covered in Microsoft Knowledge Base article 873337.

This registry key must not co-exist with the registry key set for persistent naming of IBM Atape drivers noted below.

- For systems using Windows 2003 Service Pack 1, a Microsoft hotfix documented under KB article 933682 is



also needed. Systems older than Windows 2003 Service Pack 1 do not have persistent naming functionality.

- NetWorker is automatically able to configure newly created persistent names starting with NetWorker 7.4 Service Pack 2. For older versions, manual configuration using jconfig or jedit is required.
- For IBM Atape drivers (ibmtp2k3):
  - IBM drivers do not rely on the Microsoft implementation of persistent naming, but require their own registry setting:

```
HKLM\System\CurrentControlSet\Services\ibmtp2k3\PersistentNaming:DWORD=1
```

For details refer to “IBM Tape Device Drivers Installation and User’s Guide” found at [ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM_Tape_Driver_IUG.pdf), which explains requires changes to system registry.

This registry key must not co-exist with the registry key set for persistent naming of Microsoft drivers noted above.

- NetWorker 7.4 SP4 or higher is required for automatic configuration of tape drives with newly created persistent names. For older versions of NetWorker, perform manual configuration using jconfig or jedit.
- ♦ Disable Tape Unit Ready (TUR) at the operating system layer as per instructions in Microsoft Knowledge Base article 842411 – This is required for NetWorker operations.
- ♦ It is no longer required to disable Windows Removable Storage Manager (RSM) service at the operating system layer, however, RMS must be prevented from accessing tape drives which are under NetWorker control. To achieve that, start RSM manager and remove all tape drives from 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 and the use of vendor-provided drivers is not needed unless explicitly required.  
If Microsoft-provided drivers (tape and/or ltotape) are used, no special hotfixes are needed.
- ♦ For IBM LTO drives, the use of IBM Atape drivers is supported.

Minimum required version of IBM Atape drivers for Windows 2008 is 6.1.9.5. When downloading drivers from IBM site, use drivers marked with \_w08\_.

- If upgrading IBM Atape driver, uninstall the existing version prior to upgrading.
  - The IBM Atape driver selected during installation should be a non-exclusive version. Exclusive mode is not supported by the NetWorker software.
  - IBM has changed installer options in different versions of device drivers. Refer to the README file distributed with the IBM drivers for details 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 part of “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. Furthermore, all tape drives should appear with 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
```

Details are covered in Microsoft Knowledge Base article 873337.

This registry key must not co-exist with the registry key set for persistent naming of IBM Atape drivers noted below.

- For systems using Windows 2003 Service Pack 1, a Microsoft hotfix documented under KB article 933682 is also needed. Systems older than Windows 2003 Service Pack 1 do not have persistent naming functionality.
- NetWorker is automatically able to configure newly created persistent names starting with NetWorker 7.4 Service Pack 2. For older versions, manual configuration using jconfig or jbedit is required.

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

```
HKLM\System\CurrentControlSet\Services\ibmtp2k3\PersistentNaming:DWORD=1
```

```
HKLM\System\CurrentControlSet\Services\ibmtp2k8\PersistentNaming:DWORD=1
```

For details refer to "IBM Tape Device Drivers Installation and User's Guide" found at [ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM\\_Tape\\_Driver\\_IUG.pdf](ftp://ftp.software.ibm.com/storage/devdrv/Doc/IBM_Tape_Driver_IUG.pdf), which explains requires changes to system registry.

This registry key must not co-exist with the registry key set for persistent naming of Microsoft drivers noted above.

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

## Linux

- ♦ Linux st kernel module does have native LTO support by utilizing its generic functionalities. IBM also provides Atape driver for Linux, but its use is not recommended by EMC at this time
- ♦ Minimum recommended version of RedHat Enterprise Linux for use of CDI is RHEL4 Update4 (using kernel version 2.6.9-39). Older versions of RHEL do not provide reliable tape operations due to masking of SCSI resets and thus can lead to data loss.
- ♦ Minimum recommended version of Novell SUSE Linux Enterprise Server for use with CDI is SLES9 SP3 (using kernel version 2.6.5-7.267). Older 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 below recommended levels.
- ♦ If advanced device configuration is required (for example, forcing hardware compression or other device specific options), add device vendor provided entries to /etc/stinit.def configuration file and re-initialize operating system tape device support by executing the stinit command. Note that this should not be done while devices are in use either on this system or any other system where devices are shared.

- ♦ Make sure that the SCSI Generic (sg) kernel module is loaded before the st (SCSI Tape) kernel module. This is done automatically in recommended versions of Linux distributions.
- ♦ Linux provides 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 pass-through capabilities and CDI should be disabled.
- ♦ In cases where the following kernel messages (dmesg output) are logged:

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

This is due to lpfc HBA driver not handling tagged command queuing correctly. Upgrade HBA driver or disable 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 which are controlled by /etc/udev/udev.conf and included folders/files.

- Newer versions of udev sub-system 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.
- Note that udev subsystem can be upgraded regardless of the version of the underlying Linux distribution with minimum kernel version 2.6.
- If the appropriate udev rule for tape devices is present on the system, but devices file names are not being created, check if the system is configured to allow for scanning of all SCSI devices.
- Example rule which creates entries like /dev/tape/by-id/x-xxxx-nst:

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

However, this is an example only - refer to documentation

for the specific Linux distribution for more details.

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

```
find /sys -name nst0
udevtest /sys/class/scsi_tape/nst0
```

- Any rule which 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 which utilize external programs like `scsi_id`, `lsscsi` or `scsiinfo` as they may cause SCSI resets during configuration.
- NetWorker 7.4 SP2 or higher is required for automatic configuration of tape drives with newly created persistent names. For older versions of NetWorker, perform manual configuration using `jbconfig` or `jbedit`.

## HP-UX

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

### HP STape specific information:

- ♦ HP-UX STape driver disables the hardware compression capability of LTO drives.
- ♦ HP-UX hotfix PHKL\_36312 is required for CDI operations.
- ♦ If the STape driver is being used in cases where drives are shared (used in DDS environment), operating system-level SCSI reserve/release must be disabled:

Set kernel tunable parameter:

```
st_san_safe=1
```

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

## Tru64

- Support for all tape devices is provided by generic operating system built-in drivers with no additional configuration needed for advanced options like persistent naming.
- If advanced device configuration is required (for example, forcing hardware compression or other device specific options), add device vendor provided entries to the /etc/ddb.dbase configuration file and re-initialize operating system device database by executing the `ddb_config -c` command.

## Preventing SCSI resets

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

SCSI reset manifests in NetWorker differently depending if the operating system itself detected the reset. If the reset was detected by the operating system, then the NetWorker software also detects it and pro-actively mark the volume currently in use as full to avoid any possible data loss.

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

### Primary causes of SCSI resets

The following are 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/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, etc.), but also includes a number of soft errors as well. For example, a high number of encoding errors causes a high retry count which can break the tape I/O.
- ♦ 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 provided monitoring system which 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 which attempt to communicate with device directly while device is in use by NetWorker (this includes commands like `mt`).

### External tools

Operating system or 3rd party vendor provided hardware monitoring tools often include capability to poll all hardware devices in regular intervals as a part of availability check.

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

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

Below 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



- 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

## Reserve / Release

Reserve/release is an attempt at preventing 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's tape handling and can cause SCSI resets.

## NetWorker configuration options

NetWorker releases 7.3 and later include application controlled handling of reserve/release which can be used in environment which experience problems.

## Configuration options in NetWorker

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

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

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

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

For manual control of reserve/release, refer to man pages for the

NetWorker commands `cdi_pr`, `cdi_reserve`, `cdi_release`. For further information on configuring reserve/release features in NetWorker, refer to NetWorker Administration Guide.

Copyright © 2010 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.