Tape Libraries in SAN

Configured with

OmniBack II & Advanced Tape Services

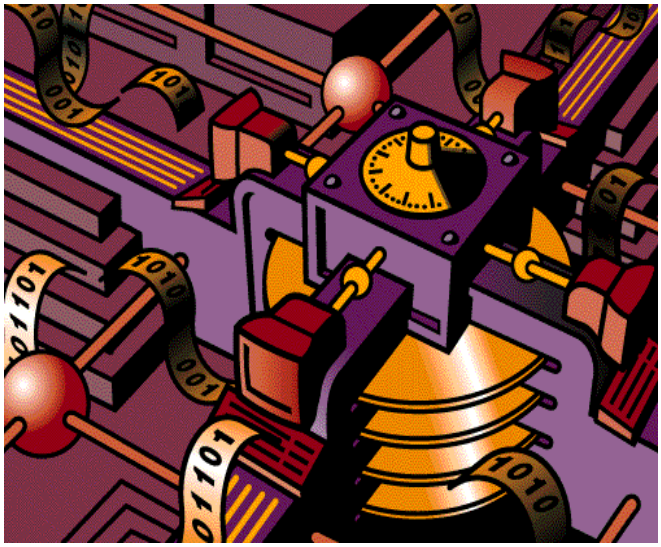

Table Of Contents

1 Introduction

Today's storage market is changing very fast. Old paradigms that have been around for a long time are no longer true. All limitations in the storage world that are tied to SCSI limitations are no longer relevant because of the adoption of the fibre channel as transport medium for the SCSI protocol. By using fibre it will be possible in the near future to build up a complex network between the hosts and the storage devices that makes the storage devices globally accessible like hosts today in a LAN. This network is called Storage Area Network.
Backup utilities need to be aware of these changes because they need to handle devices (library and drives) that are controlled by more than one host.
HP OpenView OmniBack II as of today is Storage Area Network ready. Paradigms of sharing devices such as library or drives have been possible with OmniBack II for some time and are supported by the architecture. That means OmniBack II is ready to support the Storage Area Network today.

1.1 Objective

This whitepaper describes how to configure an OmniBack II library accessed by multiple hosts in a Storage Area Network. The library control path as well as the possible drive configurations are discussed. The library configuration focuses on the Service Guard 11.05 extension Advanced Tape Service and does not give a general view on the library configuration.

1.2 Non Objective

It is not the objective of the whitepaper to describe how to install and configure an OmniBack II online backup in a MC/ServiceGuard cluster.
The paper assumes that the reader is familiar with HP-OpenView OmniBack II terminology and functionality as well as MC/ServiceGuard. For more detail on these products refer to the documentation listed in the appendix.

1.3 System Requirements

1.3.1 Interface Cards

The following table provides a list of fibre channel host based adapters that are referenced in the *Device Sharing White Paper* from ESBU.

Fibre Channel interface cards available for HP 9000	
D-Series	A3591A
K-Series	A3404A
T-Series	A3636A
V-Series	A3740A

(or replacement interface cards supported for FC)

1.3.2 FCS-MUX

A 3511A Fibre Channel to SCSI Multiplexer Firmware Revision
3840 or newer

1.3.3 Operating System

HP-UX	11.x
Patches	`PHKL_18531` (or replacement) This patch includes the release/reserve in a `sopen` or `sclose` call. `PHSS_17581` (or replacement) MC/ServiceGuard 11.05 patch

1.3.4 Tape Devices

For supported tape devices check the list of ATS supported tape devices that are kept in the file `/etc/cmcluster/shared tape/ats_tapelibs` on the cluster.

1.3.5 Software Components

Product	Version
MC/ServiceGuard	11.05
HP OpenView OmniBack II	A.03.00

This functionality does not require any OmniBack II patch released for A.03.00. Check the OmniBack II web pages whether the customer environment requires any released OmniBack II patch.

2 Advanced Tape Services

2.1 About Advanced Tape Service

Advanced Tape Service (ATS) is an integrated part of the HP MC/ServiceGuard release 11.05.

HP MC/ServiceGuard is a long term proven high availability software for HP-UX that enables an application to run in a cluster. The cluster supports up to 16 nodes and an application can be started on each host configured in the package. All resources of the application such as

- Disks
- Volume Groups
- Filesystems
- Processes
- IP address

are automatically assigned and started with the application. Service Guard manages the failover of the application based on resource failures (host panic, IP connection failure…) or service failures. This depends on the configuration.

Even though the application is high available backup still needs to be performed in order to be able to recover from

- Logical data loss
- Logical data corruption
- A complete disaster

In the past MC/ServiceGuard did not manage tape resources as SCSI only allowed a single host connection. With the introduction of the fibre channel to SCSI multiplexer (equivalent to the fibre channel bridge) a tape device can today be accessed by two hosts simultaneously similar to disks in the past. Therefore MC/ServiceGuard added the tape devices in the managed resources list. ATS is the part that manages the tape resources.

2.2 ATS components

2.2.1 Configuration File

The configuration file keeps all information about the setup of the system. The information is divided into the following sections. For reference an example of a configuration file is appended.

Global Information
This information is general information about the ATS server. It determines which cluster (CLUSTER_NAME) the ATS service belongs to. In the cluster definition the owner of the ATS service are defined.

Tape Device Information
Tape device information is listed for each tape device found. This includes stand-alone tape drives, tape library robotic devices, and the individual tape drives in a tape library device. The term autochanger refers to the robotic arm device of the tape library. Selected information can be edited using any HP-UX editor.

No Hardware Match Information
Tape device information is included in the file for each tape device previously known but currently not found. If a device is unavailable to the current active cluster (disconnected from the host e.g.) , it will display as `No Hardware Match` for the current active cluster. All the standard tape device information is included, a note is added, and it is placed at the end of the file. This entry can be completely removed or retained for reference as you choose.

For more information on the configuration file see the example at the end and refer to the ATS documentation.

2.2.2 Commands

The following is a list of ATS functions and their associated ATS commands.

To create or update the ATS configuration file, use the **`stquerycl`** command.

- This command scans all the nodes in the cluster for shared tape devices and shared library robotic devices.
- It then generates an ATS configuration file, assigning cluster-wide identifiers to each shared device.
- If `stquerycl` is executed on a cluster with an existing ATS configuration, the existing ATS configuration is merged, device by device, with any newly detected hardware.

NOTE:	An existing MC/ServiceGuard configuration is required for this command to function.

To verify the ATS configuration information, use the **`stcheckconf`** command.

- This command reads the ATS configuration file and verifies it against the MC/ServiceGuard configuration file.

To verify and apply your ATS configuration, use the **`stapplyconf`** command.

- This command reads the ATS configuration file and verifies it against the MC/ServiceGuard configuration file.

- If the verification is successful, the MC/ServiceGuard configuration is updated with the ATS configuration information.
- This updated configuration is distributed to all nodes in the cluster.

To view the ATS configuration portion of the MC/ServiceGuard configuration file, use the **`stgetconf`** command.

- This command retrieves the ATS configuration information from the MC/ServiceGuard configuration file and displays them in standard out.

To view the current status of the shared devices configured on the cluster, use the **`stviewcl`** command.

- This command uses the ATS configuration configured on the cluster and displays the status of the configured shared tape devices and shared library robotic devices.

To reclaim a shared tape device for use by the cluster, use the **`streclaim`** command.

- This command reclaims the specified device for use by the cluster if a process left the device reserved.

To remove ATS configuration from your MC/ServiceGuard configuration, use the **`stdeleteconf`** command.

- This command deletes the ATS portion from the MC/ServiceGuard configuration file.
- This command does not remove ATS device special files or associated alias files.

3 Library Sharing

This chapter describes how a library in OmniBack II must be configured in order to serve the various needs in a high available, shared environment.

Compared to a single tape drive that is shared by multiple hosts and just configured multiple times a library is more complex to share. The information of the librarie's repository is kept in OmniBack II.

If two SCSI II libraries are configured for two different hosts, OmniBack II would maintain two independent repositories even though they are physically the same. That would result in inconsistent information. For example one library loads a tape from a particular slot into the drive. If the other library were to access the same slot at that time it would mark the slot as empty. In order to avoid this the user must not configure two libraries.
In order to configure one library that is based on multiple hosts, a virtual host needs to control the library. The virtual host provides a single hostname for a cluster of hosts. A virtual host for multiple hosts can be created using a clustering software like MC/ServiceGuard.

Environment

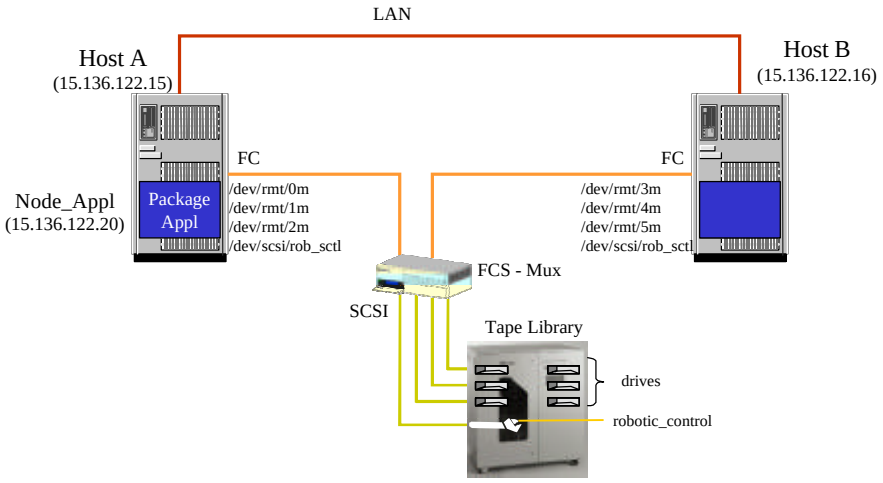

3.1 Installation

The installation procedure assumes that the customer has a pre-configured Service Guard environment that consists of host A as primary host and host B as secondary host. The high available application is configured as package 'Appl' on host A and has the floating hostname Node_Appl.

3.1.1 OmniBack II

1 Install the OmniBack II Media Agent

The OmniBack II media agent must be installed to all hosts the package can be started on (in our case host_A and host_B). The OmniBack II installation window can be used for the installation.

2 Import the virtual hostname

As the media agent needs to be started wherever the package is running OmniBack II needs to be able to talk to the virtual or floating hostname (node_Appl). Therefore the virtual hostname (node_Appl) must be imported to the OmniBack II cell using the installation window. Make sure that the checkbox **Virtual Host** ('Additional hostname' in Motif GUI) is selected in the Import Host window.

NOTE:	Do not install the software on the floating hostname. This will only install the software on the host the package is currently running on. On the other hand the software has been installed already on the static hostnames and the software must not be installed twice on the same host.

3.1.2 ATS

ATS is an integrated part of MC/ServiceGuard. The standard installation of MC/ServiceGuard on host A and host B include all software parts that are needed.

3.2 Library

OmniBack II supports different library concepts like ACSLS (STK) controlled library or direct SCSI controlled. This integration supports only the **SCSI-II Library** model of OmniBack II. The **SCSI-II Library** in OmniBack II supports all major library models in the market such as HP 4546 or STK 9710. Check the device matrix in the product info page on http://www.openview.hp.com/products/omnibackmixed/ for a complete list.

3.2.1 Configuration

The correct configuration of both components MC/ServiceGuard and OmniBack II is important for the usage of the tape library.

3.2.1.1 Service Guard

Service Guard is used to maintain the tape resources with ATS as well as providing a virtual hostname for the robotic control.

ATS
Before the ATS component can be configured the SG configuration must be completed and a cluster must exist. Refer to the MC/ServiceGuard documentation on how to do this. For reference examples of the configuration files for ATS and MC/ServiceGuard are attached to this document.

1 **Create the ATS configuration file**
Run `stquerycl` to gather the configuration of all attached tape devices and robotic control. This will create the configuration file. In the configuration file the new device file names

`/dev/rmt/st#m`	for tape device files
`/dev/rac/sac#`	for robotic control devices

are listed including a usage policy for all devices. New more meaningful names can be selected for the device files and added to the configuration file using the ALIAS_NAME (`/dev/rmt/lib1_drive1`). ATS will create this alias as device file on the cluster hosts.

2 **Check the ATS configuration**
The configuration can be checked using `stcheckconf`.

3 **Apply the ATS configuration**
If no error occurs during the check the configuration can be applied to the cluster using `stapplyconf`.

NOTE: MC/ServiceGuard does not allow ATS changes while the cluster is running. Therefore the cluster must be halted before the ATS configuration can be applied.

4 **Create a Robotic Control Package**
In the same cluster the application is running, a package must be configured providing the virtual hostname for the robotic control. This package does not need to maintain any resources besides an IP address. The package should only be allowed to switch to the hosts that have access to the library. This whitepaper uses **pkg_robotic** as package name.
The nodes, the package can switch to is configured in the **Package Configuration File**. (NODE_NAME). The IP address of the package is configured in the **Package Control Script** (IP, SUBNET).

OmniBack II connects to the virtual hostname and if no library control process (`uma`) is running OmniBack II starts automatically a new process. The package will only switch if the primary node goes down.

5 Change the /etc/hosts file

MC/ServiceGuard just provides the floating IP address. In order to bind the hostname with the IP address the standard Unix tools like /etc/hosts file needs to be used. We assume that the hostname is **node_robotic**.

NOTE: Step 4 and 5 are possible without ATS. Standard MC/ServiceGuard functionality is used to achieve this.

Configuration

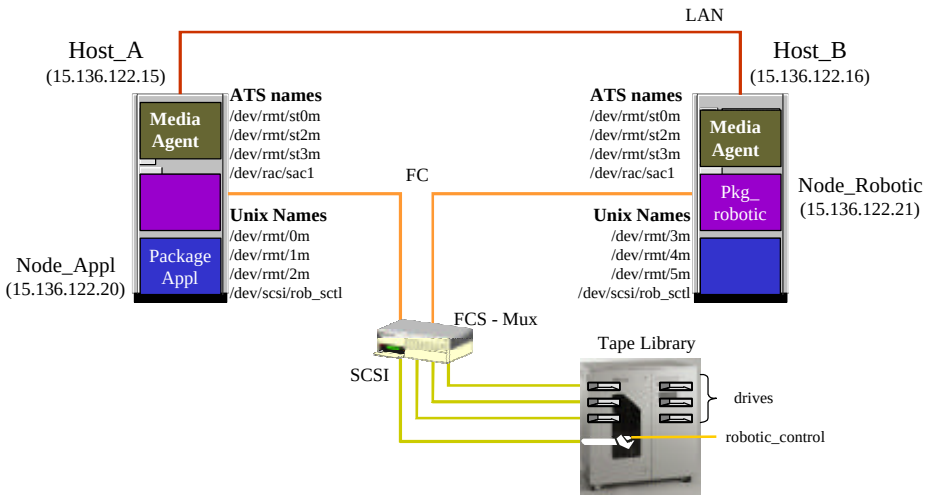

An example of a configuration file and control script is attached at the end.

3.2.1.2 OmniBack II

1 Import the virtual robotic host

Before the virtual robotic hostname can be used it has to be part of the OmniBack II cell. Therefore the hostname **node_robotic** must be imported in the OmniBack II cell using the Installation window. Make sure the checkbox **Virtual Host** is selected.

NOTE: The Service Guard package pkg_robotic must be running during this action.

2 Create a SCSI II Library

In the OmniBack II Media Management window a normal SCSI II library needs to be configured with the following options

Hostname

Use the floating hostname **node_robotic**. OmniBack II will start the `uma` on the host where the package **pkg_robotic** runs.

Control Path

Use the sctl device file created by ATS (`/dev/rac/sac1`). OmniBack II must be able to access the robotic from any host in the cluster. As only one control path can be

configured the control path needs to be identical on all hosts. The ATS extension to MC/ServiceGuard provides this automatically.

Busy Drive Handling
In case the hosts that writes the data to the tape (media agent) fails during the backup no request is sent to the `uma` process to unload the tape and the tape stays in the drive. The next backup using this drive must handle this situation. OmniBack II can react automatically in three different ways. That is configurable through the library option **Busy Drive Handling**. The following options are available

- Abort — (default) The backup will be aborted
- Eject — The uma will eject the media from the drive and put it any empty slot
- Eject to mail slot — The uma will take the media from the drive and put it in the library mailslot (CAP)

If the backup should automatically continue **Eject** should be selected. As the tape is moved to an unknown slot the library should be scanned before the next backup.

Release and Reserve
The majority of the robotic actions are atomic. In case another process wants to access the robotic he can do so without causing any problems. Nevertheless a few robotic commands consists of a sequence of scsi pass through commands. In order to assure the success of these commands OmniBack II can reserve the control path of a library. In OmniBack II A.03.10 an option in the library configuration window will be added for this. OmniBack II A.03.00 just provides a variable in the `/opt/omni/.omnirc` file on the media agent hosts (host A and host B) to be set.

```
OB2SCSILOCK=1 all pass through commands are reserved
OB2SCSILOCK=0 no reserve of pass through commands (default)
```

If OmniBack II is the only application that controls the library there is no need to reserve the SCSI path through commands as OmniBack II internally is avoiding any sharing violation. In case the library control could by any chance accessed outside of OmniBack II the reserve flag (`OB2SCSILOC=1`) should be set.

3.3 Drives

Now a single library is configured. In order to use a library the drives for this library must be configured. The drives are visible from multiple hostnames

- Host_A
- Host_B
- Node_robotic
- Node_Appl

And have multiple device file names

- Unix device files
- ATS device files

This fact allows multiple drive configurations for one single physical drive. In order to avoid sharing violations of the drives the OmniBack II Lockname can be used.

Lockname
In the advanced options for a logical drive a field Lockname can be found. In this field an additional name for the logical device can be entered that will be used for locking this

device in OmniBack II. All logical devices that share the same Lockname can only be used exclusively. OmniBack II automatically takes care of this and will queue any other backup session that is started while the Lockname is already in use.

Release and Reserve
The kernel patch **PHKL_18531** enables that all `sopen` and `sclose` calls, issued to the tape, automatically reserve and release the tape. In this way a normal backup process started from OmniBack II reserves the tape drives and avoids sharing violation from outside OmniBack II (f.e. tar, cpio ..).

When the tape is ejected from the drive, because a new tape is needed for the backup or the backup is over, OmniBack II would close the tape. This will release the scsi path and potentially another application could steal this path. If OmniBack II would then try to continue to use the drive the backup would fail.
This is not an issue as the time of changing a tape is very short and data security is guaranteed as no data can be overwritten.

Drive Index
All drives are configured for the one library that has been created on the floating host name. This also means that one drive index can occur multiple times for a library and the same physical drives need to have the same drive index. The drives do not need to be necessarily connected via the same FC-Mux as the robotic control and to the same hosts. In this case this would not change the configuration it would just mean that more drives would be configured for the library.

3.3.1 Floating drives

Drives that should be accessible from both hosts depending on which host the package is running must be configured based on the virtual host. For Example:

Hostname	node_Appl
Device Control Path	`/dev/rmt/st3m`
Lockname	Lib1_Drive_1

It is important that the device file of the tape is available on all static hosts (Host_A and Host_B) and that the name is identical. By using the ATS device file names this is automatically guaranteed.

3.3.2 Local drives

The drives can still be used in the standard way using the static hostname and the local device file. For Example:

Hostname	Host_A
Device Control Path	`/dev/rmt/0m`
Lockname	Lib1_Drive_1

The examples show the device identified by `/dev/rmt/0m` and `/dev/rmt/st3m`. Both device files are pointing the same physical devices and therefore the Lockname (Lib1_Drive_1) is identical.

4 Usage

This paragraph will explain how the different drives configured can be used for backup and what exactly will happen.

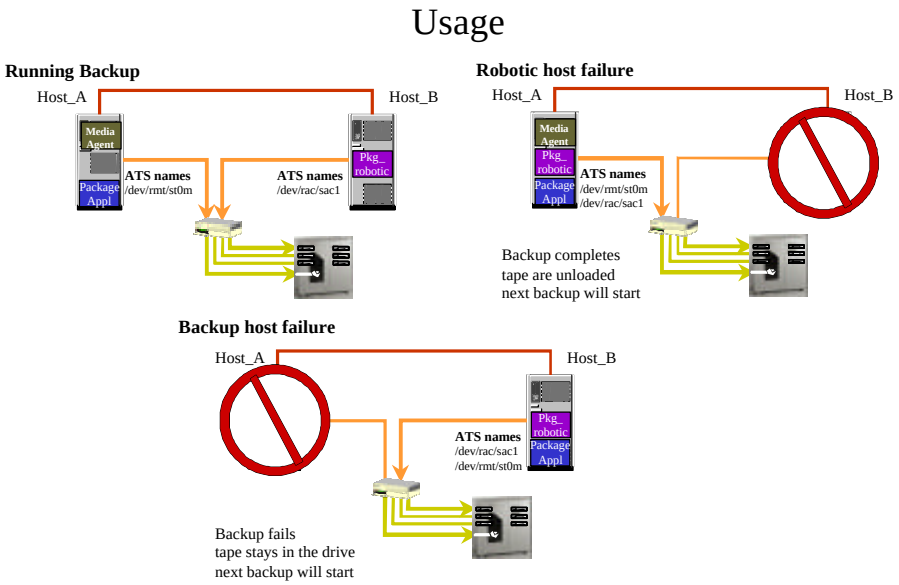

4.1 High available backup

The backup of an application in a cluster should start always no matter whether the application runs on Host_A or Host_B. In this case the customer creates a library as described with floating drives. As the virtual host is used OmniBack II determines during the startup of the backup where the package is running and starts the media agent on this host. The usage of the ATS device file names guarantees that no matter where the package is running the access path to the drive is identical.

In case the package fails during the backup the backup will be aborted. There is no automatic restart of the backup and the tape would stay in the drive. The next backup would start though. The tape drive has been released by ATS and the **Busy Drive Handling** of the library would handle the tape that is still in the drive.

For information on automatic failover of backup read the OmniBack II documentation about the high available OmniBack II cell server.

4.2 Local backup

In order to backup the local data of HP-UX on the static hosts the local drives could be used. This would not conflict with the high available backup as the Lockname and the SCSI release reserve would prevent the double usage of the drives.

In case the robotic control package is different than the host of the backup OmniBack II would start the process to load the tape into the drive on the remote host. At the end of the backup or to change the tape during the backup OmniBack II would contact this

process again. If in the meantime the package failed over to another host. OmniBack II would simply restart this process and the backup would continue and succeed.

4.3 Load Balanced

If the backup option **Load Balanced** ('Load balancing' in the Motif GUI) is used together with sharable drives the backup specification must include the drives that are local for the object. For example:

Application on Node_appl	list Floating drive
Application on Host A	list Local drive

In this way OmniBack II would automatically assign the right drives.

4.4 Multiple Libraries

Configure multiple libraries for one physical library is possible too. In this case the slots needs to be clearly divided between the libraries and the usage of such a concept would be completely different.

Library 1	slot 1-100
Library 2	slot 101-200

5 Licensing

We recommend in a "tape sharing" environment to use the single drive licenses because they will "float" between the different client systems while the Multi-drive server is "tied" to a single system. It also makes sense in that a single drive will have multiple logical drives pointing to it, but obviously only one of them will be running at a time, so the user only needs buy licenses for each physical drive (i.e. 4 physical drive, 4 licenses).

The implementation is such that all license strings are tied to the IP address of the cell manager (CM) system and the "multi-drive server" (B6952/B6954) licenses are then assigned by the OmniBack II admin to the systems he chooses. He can then move them to and from various machines with needing to obtain new licenses.

The single drive licenses (B6951/B6953) work like the "old" OBII licenses in that they will automatically "float" around the cell for drives connected to any systems. I quoted "float" here because all licensing checks are performed on the CM system but for the customer it is in end effect "floating".

If the cell manager is running in a service guard environment configured as a package, the license should be tied to the IP address associated with package (the floating IP address).

Appendix A ATS Configuration File

```
# ********************************************************************
# ************ TAPE SHARING CLUSTER CONFIGURATION FILE ***************
# ***** For complete details about cluster parameters and how to *****
# ***** set them, consult the stquerycl(1m) and atsconf(4) man pages *
# ********************************************************************
CLUSTER_NAME my_cluster
# Generated Tue Nov 10 13:09:01 1998
# The following parameters can be changed to optimize performance of
# automatic reclamation and stviewcl. These are the number of assistant
# processes which will be created on each node in the cluster.

MAX_NUM_RECLAIM_PROCS 2 (range is 2-100)
MAX_NUM_DISPLAY_PROCS 2 (range is 2-100)

# The following parameter specifies the minimum time (in seconds) between
# status queries for stviewcl.

MIN_TIME_FOR_STATUS_UPDATE 1 (range is 1-3600)

# The following device information was returned by stquerycl(1m).
# The cluster-wide identifier (numeric embedded in DEVICE_NAME)
# may be modified but must remain unique across the cluster.

DEVICE_NAME /dev/rmt/st6m

# The following flag can be modified. If "no" is changed to "yes" the device
# will not be reclaimed and its status will not be reported.

HIDE_DEVICE no

# The following information can be added to this file for the above device.
# Remove the "#" and add alias name. Multiple aliases are allowed.
# ALIAS_NAME /dev/rmt/meaningful_name
# ALIAS_NAME /dev/rmt/meaningful_name
# The following information can be added to this file for the above device.
# Remove the "#" and add any description desired. Multiple lines are
allowed.
# DESCRIPTION This is a commented line of device description.
# DESCRIPTION This is another commented line of device description.

# DO NOT ALTER any of the remaining data for this device entry!
DEVICE_TYPE tape
FCMUX_ID 1 0 0 0 0 0 0 0 117 23 56 15 137 28 240 27
BUS 0
TARGET 5
LUN 0
DEVICE_ID 1 128 2 2 0 0 0 0 137 86 36 85 243 163 32 100
VENDOR QUANTUM
PRODUCT_ID DLT7000
NODE_NAME hansel
MAJOR_NUMBER 205
MINOR_NUMBER 0x175000
NODE_NAME heckle
MAJOR_NUMBER 205
MINOR_NUMBER 0x145000

# **** NOTE **** NOTE **** NOTE **** NOTE **** NOTE **** NOTE ****
# The following device entry does not match the hardware scan.
# This entry will be ignored (you may delete this ENTIRE entry).
# The following device information was returned by stquerycl(1m).
# The cluster-wide identifier (numeric embedded in DEVICE_NAME)
# may be modified but must remain unique across the cluster.

DEVICE_NAME /dev/rmt/st16m

# The following flag can be modified. If "no" is changed to "yes" the device
# will not be reclaimed and its status will not be reported.

HIDE_DEVICE no

# The following information can be added to this file for the above device.
# Remove the "#" and add alias name. Multiple aliases are allowed.
```

```
# ALIAS_NAME /dev/rmt/meaningful_name
# ALIAS_NAME /dev/rmt/meaningful_name
# The following information can be added to this file for the above device.
# Remove the "#" and add any description desired. Multiple lines are allowed.
# DESCRIPTION This is a commented line of device description.
# DESCRIPTION This is another commented line of device description.

# DO NOT ALTER any of the remaining data for this device entry!
DEVICE_TYPE tape
FCMUX_ID 1 0 0 0 0 0 0 0 117 23 56 15 17 145 6 45
BUS 0
TARGET 2
LUN 0
DEVICE_ID 1 128 2 2 0 0 0 0 208 217 190 20 5 12 229 249
VENDOR Quantum
PRODUCT_ID DLT4000
NODE_NAME gretel
MAJOR_NUMBER 205
MINOR_NUMBER 0x002000
NODE_NAME hansel
MAJOR_NUMBER 205
MINOR_NUMBER 0x002000
# **** NOTE End of entry with no hardware match NOTE ****
```

Appendix B Package Configuration File

```
PACKAGE_NAME pkg_robotic
# Enter the names of the nodes configured for this package. Repeat
# this line as necessary for additional adoptive nodes.
# Order IS relevant. Put the second Adoptive Node AFTER the first
# one.
# Example : NODE_NAME original_node
# NODE_NAME adoptive_node

NODE_NAME Host_B
NODE_NAME Host_B
RUN_SCRIPT /etc/cmcluster/pkg_robotic/pkg_robotic.cntl
RUN_SCRIPT_TIMEOUT NO_TIMEOUT
HALT_SCRIPT /etc/cmcluster/pkg_robotic/pkg_robotic.cntl
HALT_SCRIPT_TIMEOUT NO_TIMEOU

# optionally specify services. If you specify services you have to
# use the same service names again later in the package control
# script. They MUST match.
# We will configure services later !!! Leave blank or commented
# optionally specify subnet to be monitored. This means if the subnet
# is not available anymore on the node on which the package runs then
# the package will be moved to an adoptive node.

SUBNET 192.6.100.0

# The default for PKG_SWITCHING_ENABLED is YES. In the event of a
# failure, this permits ServiceGuard to transfer the package to an
# adoptive node. Adjust as necessary.

PKG_SWITCHING_ENABLED YES

# The default for NET_SWITCHING_ENABLED is YES. In the event of a
# failure, this permits ServiceGuard to switch LANs locally (transfer
# to a standby LAN card). Adjust as necessary.

NET_SWITCHING_ENABLED YES

# The default for NODE_FAIL_FAST_ENABLED is NO. If set to YES,
# in the event of a failure, ServiceGuard will halt the node on
# which the package is running. Adjust as necessary.

NODE_FAIL_FAST_ENABLED NO
```

Appendix C Package Control Scripts

```
/etc/cmcluster/pkg_robotic/pkg_robotic.cntl

# Set PATH to reference the appropriate directories.

PATH=/sbin:/usr/bin:/usr/sbin:/etc:/bin:/opt/dtcmgr/sbin

# Volume groups, Logical volumes, and File systems.
#
# Make sure that you turn auto-activation in /etc/lvmrc off and
# that you only activate VG's automatically which are independent
# to ServiceGuard packages.
#
# Make sure you did "vgchange -c y /dev/vgX" once to enable
# Shared VG Locking for this Volume Group or that the VG is
# declared in the Cluster Configuration File.
#
#VGCHANGE="vgchange -a e -q n" # In case of disk mirroring
#VGCHANGE="vgchange -a y

VGCHANGE="vgchange -a e" # If no disk mirroring

# Make sure that shared VG's and LV's have the same names on
# all required nodes and that the same mount points exist on all
# nodes.
# Make also sure that you DO NOT SPECIFY an entry for the filesystems
# in /etc/fstab. fsck and mount for Shared VG's is done by the package
VG[0]=
LV[0]= ; FS[0]=
# To find out which subnets are available for SG execute the command
# cmquerycl -v -l net -n node1 -n node2 [ ... -n node_n ]
# IP[0]=... will be a relocatable IP address assigned to the package
# The IP address must be unique - you cannot have duplicate IP's !!!

IP[0]=15.136.122.21
SUBNET[0]=255.255.255.0

# The service names MUST match the ones in the package description
# file. If the command uses parameters don't forget the quotes ""
# SERVICE_NAME[0]=
# SERVICE_CMD[0]=
# SERVICE_RESTART[0]=
```

Appendix D **Documentation**

MC/ServiceGuard

B3936-90032	Using Advanced Tape Services
B3935-90015	MC/ServiceGuard Version A.11.05 Release Notes
B3936-90026	Managing MC/ServiceGuard, Sixth Edition
Intranet	http://haweb.cup.hp.com/HA/
Internet	http://www.datacentersolutions.hp.com/2_3_index.html

HP OpenView OmniBack II

B1953-90000	Concept Guide
B6960-90002	Administrator's Guide using the Motif User Interface
B6960-90002	Administrator's Guide using the Windows User Interface
Intranet	http://ovweb.bbn.hp.com/stor/hp/htdocs/index.htm
Internet	http://www.openview.hp.com/solutions/storage/

ESBU

Intranet

http://essd.boi.hp.com/Products/sans/release%201/ibrsweb/sharingwp.html