



An Oracle White Paper
January 2013

Protecting SPARC SuperCluster—Tape Backup with Symantec NetBackup

Introduction	1
Hardware and Software	2
NetBackup Master Servers	2
NetBackup Media Servers.....	2
Oracle's StorageTek Library.....	3
OS	3
Backup Software	3
Scalability.....	3
Hostname and IP Information.....	3
Half Rack Oracle SPARC SuperCluster	4
NetBackup Backup Servers	4
Fibre Channel Connectivity	5
Fibre Channel Zoning.....	5
InfiniBand Connectivity.....	5
InfiniBand HCA Configuration on Oracle Solaris	6
InfiniBand Switch Configuration.....	8
Adding Backup Server InfiniBand Ports to InfiniBand Network	8
Date Synchronization	10
NetBackup Configuration	10
Modify /etc/hosts Files.....	10
Library Drive OS Verification	11
Install NetBackup	11
Install NetBackup Device Driver and Configure Devices	12
Modify bp.conf Files	13
Link Oracle RMAN to NetBackup	14
Verify Connectivity	14
Configuring Storage Devices.....	15
Configure Storage Unit Groups	15
Create NetBackup Policies.....	16
Verify NetBackup Clients.....	22
NetBackup Host Property Configuration.....	22
NetBackup Alternate Restore Configuration	23
Oracle RMAN Script.....	24
Modifications for Data Files	24
Modifications for Archive Logs	25
Modifications for Control File.....	26
Execute Backup Test	26
Restoring the Database	26

Restore Example	26
NetBackup Tuning	32
Identifying NetBackup Bottlenecks	32
Hardware Compression	33
Recommended Tape Drive Tuning.....	33
NetBackup Performance	33
NetBackup Troubleshooting	34
Slow Transfer Rates.....	34
NetBackup Spawning Excessive Jobs Versus Number of Available Drives	34
Conclusion	34
Appendix.....	35

Introduction

The Oracle SPARC SuperCluster is a highly redundant hardware configuration for storing data, but redundancy is no substitution for data backups. Data backups are needed in case data corruption is replicated to the redundant systems, or a disaster strikes. It is also important to store data at an offsite location to protect against disasters. When considering a backup solution for Oracle SPARC SuperCluster, meeting a suitable backup window is imperative and can be challenging with the large capacities of the Oracle SPARC SuperCluster. Oracle's StorageTek tape libraries and Oracle's StorageTek tape drives in combination with InfiniBand provide a high-performance, cost-effective solution. This paper discusses detailed setup for an Oracle tape backup solution with Symantec NetBackup using an InfiniBand network. The information contained in this document is intended for the system administrator responsible for protecting the Oracle SPARC SuperCluster. The information provided is also intended to fill knowledge gaps in setup procedures for the OS and tape backup software. In addition to detailed configuration instructions, sections on tuning, troubleshooting, and additional resources are provided.

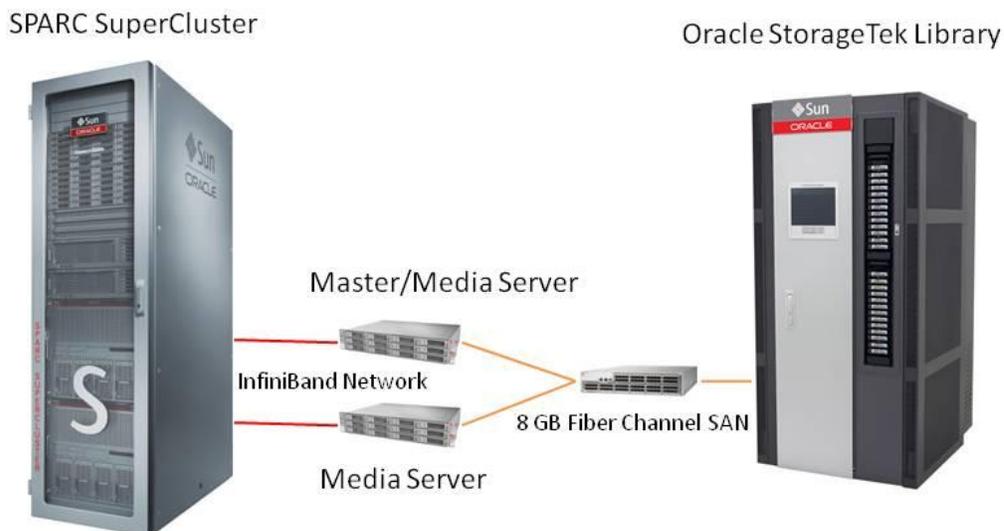


Figure 1: Physical components and connections for Oracle SPARC SuperCluster tape solution.

Hardware and Software

Below is a summary of the hardware and software utilized for testing this configuration.

TABLE1. TEST ENVIRONMENT

HARDWARE/SOFTWARE	VERSION	PURPOSE
Oracle Solaris	11	OS on Backup Servers
Oracle Solaris kernel	Branch: 0.175.0.0.0.2.1	Kernel on Backup Servers
NetBackup	7.0, 7.1, or 7.5	Backup Software
Oracle SunFire x4270 M2	M2 (12 Drive)	NetBackup Servers
Oracle SPARC SuperCluster	Half Rack Configuration	Cluster Environment
Oracle's StorageTek SL3000	Latest Firmware	Tape Library for Backup
Oracle's StorageTek T10000C	Latest Firmware	Tape Drives for Backup
Brocade 5300 8 Gb Fibre Channel Switch	Latest Firmware	SAN Switch
Emulex 8 Gb HBA	SG-XPCIE2FC-EM8-Z	Tape Connectivity on Media Servers
InfiniBand HCA Cards	X4237	InfiniBand Connectivity to Media Servers

The hardware used in this configuration is for example only. The setup procedures are applicable to any environment and any combination of tape libraries, tape drives, and media servers. This document is specifically written for configuring backups over InfiniBand, but can also be leveraged for an Ethernet backup setup as the majority of the NetBackup and Oracle Recovery Manager (Oracle RMAN) configuration will still apply.

NetBackup Master Servers

One NetBackup Master Server can be used for the different sizes of the Oracle SPARC SuperCluster—half rack, full rack, or multiple racks. In a configuration consisting of multiple racks, a dedicated Master Server is recommended. The configuration used in this testing is a half rack Oracle SPARC SuperCluster so a shared Master/Media Backup Server was sufficient.

NetBackup Media Servers

The number and size of each NetBackup Media Backup Server can vary depending on where bus saturation occurs, or other limits (CPU, memory) are reached. The InfiniBand and Fibre Channel interfaces do not consume heavy CPU bandwidth, so most likely the limiting factor is bus saturation.

Oracle's StorageTek Library

StorageTek tape libraries and StorageTek tape drives.

OS

The NetBackup Master/Media and NetBackup Media Servers have Oracle Solaris 11 installed.

Backup Software

NetBackup Enterprise Server software (7.0 or above – 7.5 used for testing) and associated licensing.

Scalability

This solution can scale up or down depending on the size of the environment. Simply apply the principles used in this document for a specific configuration. The ratio of tape drives to NetBackup Media Servers will vary depending on the size of the NetBackup Media Server being used, the speed and capabilities of the bus on the NetBackup Media Server, the amount of data to be backed up, and the particular RPOs and RTOs. Please see the NetBackup Backup Planning and Performance Tuning Guide in the Appendix for detailed information on sizing a NetBackup environment.

Example 1: One NetBackup Media Server and seven tape drives. The configuration changes in this document only need to be made for the single NetBackup Media Server and the Oracle RMAN script should be set up to round robin across the available compute nodes. Oracle RMAN parallelism needs to be set accordingly (match parallelism to the number of drives in the environment). All tape drives will be zoned and set up on the single NetBackup Media Server and a storage unit group is not necessary. All the policy setup and InfiniBand setup on the single NetBackup Media Server will be identical to a multiple NetBackup Media Server environment.

Example 2: Two or more NetBackup Media Servers and additional tape drives. Allocate the maximum number of drives to NetBackup Media Server 1 necessary to saturate the I/O bus. Repeat for each additional NetBackup Media Server until all drives are allocated (each NetBackup Media Server does not have to contain the same number of tape drives). Configure one or more storage unit groups in NetBackup to accommodate the requirements. In most instances only one storage unit group is needed in order to spread the backup load. Replicate the setup of the hosts and bp.conf for each of the additional NetBackup Media Servers as well as modifying the Oracle RMAN script for the number of drives in the environment.

Hostname and IP Information

This section details the Ethernet information and InfiniBand information for the test environment.

Half Rack Oracle SPARC SuperCluster

Compute Node 1 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aiesscdb01 IP 10.80.75.106

InfiniBand:

aiessc01-stor IP 192.168.30.6

Compute Node 2 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aiesscdb02 IP 10.80.75.107

InfiniBand:

aiessc02-stor IP 192.168.30.7

Storage Server 1 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aiessccel01 IP 10.80.75.108

InfiniBand:

aiessccel01-priv IP 192.168.10.14

Storage Server 2 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aiessccel02 IP 10.80.75.109

InfiniBand:

aiessccel02-priv IP 192.168.10.15

Storage Server 3 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aiessccel03 IP 10.80.75.110

InfiniBand:

aiessccel03-priv IP 192.168.10.16

Sun ZFS Storage Appliance server 1 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aie-zba-h1 IP 10.80.74.170

InfiniBand:

aie-zba-h1-stor IP 192.168.30.100

aie-zba-h1-stor IP 192.168.30.102

aie-zba-h1-stor IP 192.168.30.104

aie-zba-h1-stor IP 192.168.30.106

Sun ZFS Storage Appliance server 2 (model per Oracle SPARC SuperCluster specifications):***Ethernet:***

aie-zba-h2 IP 10.80.74.172

InfiniBand:

aie-zba-h2-stor IP 192.168.30.101

aie-zba-h2-stor IP 192.168.30.103

aie-zba-h2-stor IP 192.168.30.105

aie-zba-h2-stor IP 192.168.30.107

NetBackup Backup Servers

Master/Media Server 1:***Ethernet:***

aie-mediaserver IP 10.80.75.5

InfiniBand:

aie-mediaserver-ib IP 192.168.30.200

Media Server 2:***Ethernet:***

aie-mediaserver2 IP 10.80.75.7

InfiniBand:

aie-mediaserver2-ib IP 192.168.30.201

Note: This example uses a half rack Oracle SPARC SuperCluster. Depending on the use of a half rack, a full rack or multiple racks, the number of compute nodes, storage servers, and potentially NetBackup Media Servers, will be different.

Fibre Channel Connectivity

Eight Gb Fibre Channel switch is recommended.

The number of Fibre Channel cards will vary depending on the number of tape drives available, server HBA slots available, and performance requirements.

The FC HBA cards on the NetBackup Media Servers were inserted into slots one and slots five on each host.

Fibre Channel Zoning

For performance reasons, a maximum of two tape drives were zoned to each FC port in the configuration. Depending on the situation, and available resources, there is an option to cable more drives to one port, or less, but keep in mind that doing so may affect performance.

Since one NetBackup media server is also serving the role as the Master Server, the library robotic control path is added into one HBA port which has one drive zoned.

InfiniBand Connectivity

Connection of the NetBackup Backup Servers to the Oracle Sparc SuperCluster InfiniBand network is dependent on the configured environment. The servers can be:

- Connected directly to Oracle Sparc SuperCluster InfiniBand leaf switches—This option is used if the backup servers are the only devices (other than Oracle Sparc SuperCluster expansion) that will be connected to the infrastructure.
- Connected to external InfiniBand leaf switches—This option is used if more appliances or devices will be connected. In this case, two additional leaf switches are introduced into the InfiniBand network, but are not located within an Oracle Sparc SuperCluster or an Oracle Sparc SuperCluster expansion rack. A typical scenario for this configuration would be when both a Oracle's Sun ZFS Backup Appliance and backup servers (connected to tape drives) are needed.

Figure 2 shows the InfiniBand switch configuration with the addition of two external leaf switches. The InfiniBand ports of the two NetBackup Backup Servers used in this testing are connected to the external leaf switches, as there is a Oracle's Sun ZFS Backup Appliance appliance currently in the environment as well.

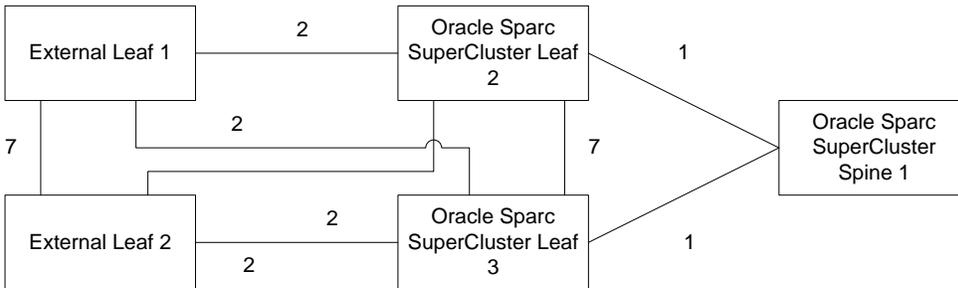


Figure 2: External InfiniBand leaf switches connected to the Oracle Sparc SuperCluster InfiniBand leaf switches.

The port connections on the external leaf switches are not as tightly controlled as are the ports on the Oracle Sparc SuperCluster InfiniBand switches. Thus, specific port selection on the external leaf switches is not critical.

Connection of the two backup servers to the external leaf switches should follow these guidelines:

- Connect Port 1 of the InfiniBand HCA port on NetBackup Master/Media Server to a port on external leaf 1
- Connect Port 2 of the InfiniBand HCA port on NetBackup Master/Media Server to a port on external leaf 2
- Connect Port 1 of the InfiniBand HCA port on NetBackup Media Server to a port on external leaf 1
- Connect Port 2 of the InfiniBand HCA port on NetBackup Media Server to a port on external leaf 2

These connections provide for the maximum throughput and availability.

InfiniBand HCA Configuration on Oracle Solaris

Oracle Solaris 11 InfiniBand packages will be installed from an Oracle Solaris 11 software repository and may or may not match the stack installed on the Oracle SPARC SuperCluster database servers.

Each NetBackup Backup Server contains one dual-port InfiniBand card. Two InfiniBand cables are used for each NetBackup Backup Server.

Note: Second cable for redundancy only.

The InfiniBand HCA cards on the NetBackup Backup Servers were inserted into PCIe 2.0 slot three on each host. Slot zero on each server was intentionally left blank as it shares the same passive riser as slot three. The cards were set up in this manner for performance reasons.

Note: Depending on the server model, the speed of each PCI slot may vary. Choose the fastest PCI slots available for the InfiniBand cards when selecting a PCI slot on the NetBackup Media Server.

1. Install the `server_install` package from the Oracle Solaris 11 repository:

```
pkg install server_install
```

2. Verify that the InfiniBand HCA cards are present on the system:

```
root@aie-mediaserver:~# dladm show-phys |grep Infiniband
LINK      MEDIA      STATE SPEED DUPLEX  DEVICE
net4      Infiniband up      32000 unknown ibp0
net5      Infiniband up      32000 unknown ibp1
```

The above output shows net4 and net5 are active InfiniBand ports

3. Create an InfiniBand partition for each port to be used for the Internet Protocol Multi-Path (IPMP) group (InfiniBand is being set up for redundancy and IPMP is the equivalent of bonding in Oracle Solaris):

```
root@aie-mediaserver:~# dladm create-part -f -l net4 -P 8503 bondib0_0
root@aie-mediaserver:~# dladm create-part -f -l net5 -P 8503 bondib0_1
```

4. Create an interface for each of the InfiniBand partitions:

```
root@aie-mediaserver:~# ipadm create-ip bondib0_0
root@aie-mediaserver:~# ipadm create-ip bondib0_1
```

5. Create an IPMP group out of the two interfaces:

```
root@aie-mediaserver:~# ipadm create-ipmp -i bondib0_0,bondib0_1 bondib0
```

6. Use the ipadm utility to assign an IP address, assign a netmask, and bring the interface online:

```
root@aie-mediaserver:~# ipadm create-addr -T static -a 192.168.30.200/24 bondib0/v4
```

7. Run ifconfig -a to verify the interface is online:

```
root@aie-mediaserver:~# ifconfig -a
lo0: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
net0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 10.80.75.5 netmask fffffe00 broadcast 10.80.75.255
    ether 0:21:28:a5:2c:de
bondib0_0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 65520 index 8
    inet 0.0.0.0 netmask ff000000
    groupname bondib0
    ipib 80:0:0:4a:fe:80:0:0:0:0:0:0:0:21:28:0:1:cf:57:53
bondib0_1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 65520 index 9
    inet 0.0.0.0 netmask ff000000
    groupname bondib0
    ipib 80:0:0:4c:fe:80:0:0:0:0:0:0:0:21:28:0:1:cf:57:54
bondib0: flags=8001000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4,IPMP> mtu 65520
index 10
    inet 192.168.30.200 netmask ffffff00 broadcast 192.168.30.255
    groupname bondib0
```

From the above output the bondib0 group is online with the configured IP address of 192.168.30.200

Note: IPv6 entries removed from ifconfig output for brevity.

- Repeat steps 1-7 for any additional backup servers

InfiniBand Switch Configuration

The InfiniBand switch is partitioned by port GUID.

The port GUID of each InfiniBand port on the backup servers can be identified with `dladm show-ib`:

```
root@aie-mediaserver:~# dladm show-ib
LINK  HCAGUID          PORTGUID          PORT  STATE  PKEYS
net4   21280001CF5752   21280001CF5753   1     up     FFFF,8503
net5   21280001CF5752   21280001CF5754   2     up     FFFF,8503
```

```
root@aie-mediaserver2:~# dladm show-ib
LINK  HCAGUID          PORTGUID          PORT  STATE  PKEYS
net6   21280001A0CDF8   21280001A0CDFA   2     up     FFFF,8503
net5   21280001A0CDF8   21280001A0CDF9   1     up     FFFF,8503
```

Adding Backup Server InfiniBand Ports to InfiniBand Network

The GUIDs of the NetBackup Backup Servers' ports need to be added to the existing Oracle SPARC SuperCluster InfiniBand configuration. These ports must be added to the configuration with the partition key of 8503 in order for the components to communicate with each other.

- Log on to the Oracle SPARC SuperCluster InfiniBand spine switch as root and verify that it is running subnet manager by executing the `enablesm` command

```
login as: root
root@aiesscsw-ib1's password:
Last login: Tue Sep 25 08:19:01 2012 from dhcp-brm-bl5-204-3e-east-10-135-75-254.usdhcp.oraclecorp.com
[root@aiesscsw-ib1 ~]# enablesm
opensm (pid 15906) is already running...
Starting partitiond daemon
/usr/local/util/partitiond is already running
(Users may also perform a 'restart' if wanted)
[root@aiesscsw-ib1 ~]#
```

By default, the spine switch is given a hostname of `<sscid>sw-ib1`, where `<sscid>` is the prefix name given to the entire Oracle SPARC SuperCluster system. In these examples, the `<sscid>` is `aiessc`.

The `enablesm` command will report that the subnet manager “partitiond” is already running, and will attempt to start it if it is not.

2. Enter the command **getmaster** to verify that this is the master switch of the configuration. If the master switch is not running on the spine switch, log out and log in to the designated master switch for the remainder of this procedure

```
[root@aiesscsw-ib1 ~]# getmaster
Local SM enabled and running
20120913 10:16:51 Master SubnetManager on sm lid 13 sm guid 0x2128e8ac27a0a0 : SUN DCS 36P QDR
aiesscsw-ib1.us.oracle.com
```

3. Add the NetBackup Backup Servers to the InfiniBand configuration:

- a. Enter the command **smpartition start** to start a reconfiguration session

```
[root@aiesscsw-ib1 ~]# smpartition start
```

- b. Enter the command **smpartition add** to add the four new GUIDs to the configuration

```
[root@aiesscsw-ib1 ~]# smpartition add -n sto -port 21280001cf5753 21280001cf5754 21280001a0cdf9
21280001a0cdfa
```

- c. Enter the command **smpartition list modified** to verify the new GUIDs have been added correctly

```
[root@aiesscsw-ib1 ~]# smpartition list modified
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 11
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
ic1s10 = 0x0501,ipoib,defmember=full:
0x0021280001ef30f7,
0x0021280001ef33bf,
0x0021280001ef30b7,
0x0021280001ef314b;
ic2s10 = 0x0502,ipoib,defmember=full:
0x0021280001ef30f8,
0x0021280001ef33c0,
0x0021280001ef30b8,
0x0021280001ef314c;
sto = 0x0503,ipoib,defmember=full:
0x0021280001cf5753,
0x0021280001cf5754,
0x0021280001a0cdf9,
0x0021280001a0cdfa,
0x0021280001ef43f8,
0x0021280001ef43b7,
0x0021280001cf90c0,
0x0021280001ef43bb,
0x0021280001ef43bc,
0x0021280001cf90bf,
0x0021280001ef43b8,
0x0021280001ef43f7,
0x0021280001ef3048,
```

```

0x0021280001ef30af,
0x0021280001ef30f8,
0x0021280001ef30f7,
0x0021280001ef33c0,
0x0021280001ef33bf,
0x0021280001ef30cc,
0x0021280001ef342b,
0x0021280001ef30b8,
0x0021280001ef30b7,
0x0021280001ef314c,
0x0021280001ef314b,
0x0021280001efec65,
0x0021280001efec66,
0x0021280001efecb1,
0x0021280001efecb2;

```

d. Enter the command **smpartition commit** to apply the new configuration and propagate configuration changes to all InfiniBand switches in the configuration

```
[root@aiesscsw-ib1 ~]# smpartition commit
```

Date Synchronization

Ensure the Network Time Protocol (NTP) is in use or have all the dates/times between all servers involved in this configuration in sync.

NetBackup Configuration

NetBackup installation consists of modifying the `/etc/hosts` files, library drive OS verification, software installation, `bp.conf` modification, Oracle RMAN library linking, server communication verification, storage device configuration, storage unit group configuration, policy creation, client communication verification, host property modification, and alternate restore configuration. For additional details on NetBackup installation, please refer to NetBackup Documentation link in the Appendix.

Modify `/etc/hosts` Files

The hosts file on both of the compute nodes should already be populated with Ethernet and InfiniBand entries for itself, the other node(s) and Oracle's Exadata Storage Servers.

Copy one of the hosts files from a compute node and append it to the hosts file on each NetBackup Backup Server. Add the InfiniBand addresses and names of the backup servers.

Example: `/etc/hosts` on `aie-mediaserver`

```

::1 aie-mediaserver localhost
127.0.0.1 localhost localhost
10.80.75.5 aie-mediaserver
#NBU Media Server
10.80.75.7 aie-mediaserver2
#InfiniBand
192.168.30.200 aie-mediaserver-ib
192.168.30.201 aie-mediaserver2-ib

```

Edit each of the hosts file on the compute nodes by inserting the Ethernet and InfiniBand addresses with the names of the backup servers.

```
10.80.75.5      aie-mediaserver
192.168.30.200 aie-mediaserver-ib
10.80.75.7     aie-mediaserver2
192.168.30.201 aie-mediaserver2-ib
```

Library Drive OS Verification

Use the `cfgadm` utility to verify the robot and tape drives can be seen on each server in the configuration. The following example shows the robot and six tape drives on NetBackup Master Server and six tape drives on the NetBackup Media Server:

```
root@aie-mediaserver:/# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c8	fc-fabric	connected	configured	unknown
c8::500104f000acc718	med-changer	connected	configured	unknown
c8::500104f000acc752	tape	connected	configured	unknown
c9	fc-fabric	connected	configured	unknown
c9::500104f000acc74c	tape	connected	configured	unknown
c9::500104f000acc74f	tape	connected	configured	unknown
c10	fc-fabric	connected	configured	unknown
c10::500104f000acc75b	tape	connected	configured	unknown
c11	fc-fabric	connected	configured	unknown
c11::500104f000acc755	tape	connected	configured	unknown
c11::500104f000acc758	tape	connected	configured	unknown

```
root@aie-mediaserver2:/# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c1	fc-fabric	connected	configured	unknown
c1::500104f000acc76a	tape	connected	configured	unknown
c2	fc-fabric	connected	configured	unknown
c2::500104f000acc767	tape	connected	configured	unknown
c2::500104f000acc76d	tape	connected	configured	unknown
c4	fc-fabric	connected	configured	unknown
c4::500104f000acc764	tape	connected	configured	unknown
c5	fc-fabric	connected	configured	unknown
c5::500104f000acc75e	tape	connected	configured	unknown
c5::500104f000acc761	tape	connected	configured	unknown

Note: Output excerpted to only show relevant devices.

Install NetBackup

1. Install NetBackup Enterprise Server software on Oracle's Sun Server designated as the NetBackup Master Server (host **aie-mediaserver** in this example). Specify hostname **aie-mediaserver** as the name of the Master

Server. When prompted as to whether or not one has NetBackup Media Servers, specify yes, and use **aiemediaserver2** as the NetBackup Media Server's name.

Example:

Would you like to use "aiemediaserver" as the configured NetBackup server name of this machine? [y,n] (y)

Is aiemediaserver the master server? [y,n] (y)

Media servers can be added during this installation or to a NetBackup environment after installation completes. Refer to the NetBackup Administrator's Guide, Volume I for more information.

Do you want to add any media servers now? [y,n] (n) y

Enter the fully qualified name of a media server (q to quit): aiemediaserver2

Enter the fully qualified name of a media server (q to quit): q

2. Install NetBackup Enterprise Server software on the other Media Server (**aiemediaserver2**). When prompted for the NetBackup Media Servers' names, use itself as the server name. Specify **aiemediaserver** as the name of the NetBackup Master Server.

Example:

Would you like to use "aiemediaserver2" as the configured NetBackup server name of this machine? [y,n] (y)

Is aiemediaserver2 the master server? [y,n] (y) n

What is the fully qualified name of the master server? aiemediaserver

3. Install NetBackup client software on each of the compute nodes. Use the InfiniBand names for each of the NetBackup clients (**aiessc01-stor** and **aiessc02-stor**).

Install NetBackup Device Driver and Configure Devices

For Oracle Solaris Fibre Channel devices it is often necessary to rebuild the NetBackup sg driver in order for NetBackup to detect the devices.

1. root@aiemediaserver:~# rm /kernel/drv/sg.conf
2. root@aiemediaserver:~# cd /usr/opensv/volmgr/bin/driver
3. root@aiemediaserver:/usr/opensv/volmgr/bin/driver# ../sg.build all
The file ./st.conf should be appended to /kernel/drv/st.conf.
A reboot may be necessary to create any new device files.
Created file ./sg.conf.
Created file ./sg.links.
4. root@aiemediaserver:/usr/opensv/volmgr/bin/driver# ./sg.install

```

Copied files to /kernel/drv/amd64.
Doing add_drv of the sg driver
Removing old /dev/sg entries.
Editing /etc/devlink.tab...
Copying original /etc/devlink.tab to /etc/devlink.tab.11-30-12-09:57:33.
Added entry in /etc/devlink.tab file.
Made links in /dev/sg

```

5. root@aie-mediaserver:/usr/opensv/volmgr/bin/driver# ../sgscan

```

/dev/sg/c0tw500104f000acc718l0: Changer: "STK SL3000"
/dev/sg/c0tw500104f000acc74cl0: Tape (/dev/rmt/1): "STK T10000C"
/dev/sg/c0tw500104f000acc74fl0: Tape (/dev/rmt/2): "STK T10000C"
/dev/sg/c0tw500104f000acc752l0: Tape (/dev/rmt/0): "STK T10000C"
/dev/sg/c0tw500104f000acc755l0: Tape (/dev/rmt/5): "STK T10000C"
/dev/sg/c0tw500104f000acc758l0: Tape (/dev/rmt/4): "STK T10000C"
/dev/sg/c0tw500104f000acc75bl0: Tape (/dev/rmt/3): "STK T10000C"

```

6. Repeat steps 1-5 for any additional backup servers.

Modify bp.conf Files

The bp.conf files on the NetBackup Master Server and NetBackup Media Servers do not require modification and should look as follows:

```

root@aie-mediaserver:/usr/opensv/netbackup# pg bp.conf
SERVER = aie-mediaserver
SERVER = aie-mediaserver2
CLIENT_NAME = aie-mediaserver
CONNECT_OPTIONS = localhost 1 0 2
USE_VXSS = PROHIBITED
VXSS_SERVICE_TYPE = INTEGRITYANDCONFIDENTIALITY
EMMSERVER = aie-mediaserver
HOST_CACHE_TTL = 3600
VXDBMS_NB_DATA = /usr/opensv/db/data
LIST_FS_IMAGE_HEADERS = NO
MEDIA_UNMOUNT_DELAY = 1800
MAX_INDEXING_JOBS = 12
TELEMETRY_UPLOAD = YES

```

```

root@aie-mediaserver2:/usr/opensv/netbackup# pg bp.conf
SERVER = aie-mediaserver
SERVER = aie-mediaserver2
CLIENT_NAME = aie-mediaserver2
MEDIA_SERVER = aie-mediaserver2
CONNECT_OPTIONS = localhost 1 0 2
USE_VXSS = PROHIBITED
VXSS_SERVICE_TYPE = INTEGRITYANDCONFIDENTIALITY
EMMSERVER = aie-mediaserver
HOST_CACHE_TTL = 3600
MEDIA_UNMOUNT_DELAY = 1800

```

```
TELEMETRY_UPLOAD = NO
```

The bp.conf file on each of the clients (aiesscdb01 and aiesscdb02) machines should include a SERVER entry for **aie-mediaserver** and MEDIA_SERVER entries for each of the NetBackup Media Servers using the InfiniBand names **aie-mediaserver-ib** and **aie-mediaserver2-ib**. The CLIENT_NAME must be listed using the InfiniBand name – **aiessc01-stor** or **aiessc02-stor**.

```
root@aiesscdb01:/usr/opensv/netbackup# pg bp.conf
SERVER = aie-mediaserver
MEDIA_SERVER = aie-mediaserver-ib
MEDIA_SERVER = aie-mediaserver2-ib
CLIENT_NAME = aiessc01-stor
CONNECT_OPTIONS = localhost 1 0 2
```

Link Oracle RMAN to NetBackup

Link the Oracle RMAN binary to NetBackup (repeat on each database server):

1. Log in as root to the database server
2. Switch user to oracle: **su – oracle**
3. Change to lib directory in Oracle Home: **cd \$ORACLE_HOME/lib**
4. Execute NetBackup link script: **/usr/opensv/netbackup/bin/oracle_link**
5. Verify link: **ls -l \$ORACLE_HOME/lib/libobk.so**

Verify Connectivity

1. From the NetBackup Master Server (aie-mediaserver) launch the NetBackup GUI:
`/usr/opensv/netbackup/bin/jnbSA&`
2. Navigate to the Media Server list: NetBackup Management->Host Properties->Media Servers

Users should see aie-mediaserver and aie-mediaserver2 in the list. Select each one and a green check box should appear indicating connection to each server:

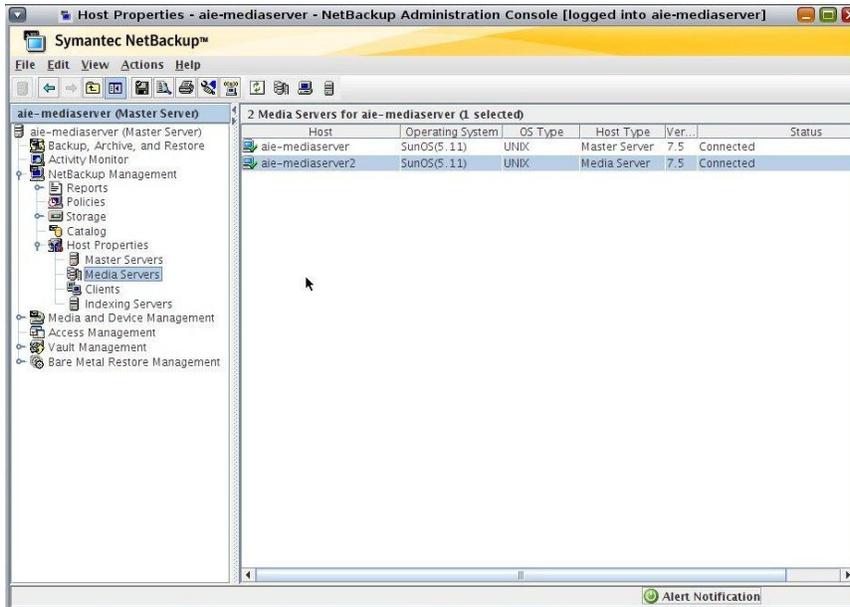


Figure 3: NetBackup Host Properties Media Server list.

Configuring Storage Devices

1. Use the NetBackup Configuration Wizard and follow the prompts to configure the storage devices. For this configuration the robot should be discovered on the NetBackup Master Server, and the tape drives should be split evenly between NetBackup Media Servers **aie-mediaserver** and **aie-mediaserver2**.
2. Navigate to Media and Device Management-> Devices->Drives and then Media Management->Devices->Robots to verify the configuration.
3. Inventory the robot and setup tape pools per organizational rules.

Configure Storage Unit Groups

In order to load balance between NetBackup Media Servers users must configure a NetBackup Storage Unit Group. Navigate to NetBackup Management->Storage->Storage Unit Groups->, right click and select “New Storage Unit Group.” From the “New Storage Unit Group” screen, enter a name and select the Storage Units to be added to the group. Under the “Storage Unit Selection” section, choose the radio button for “Media Server Load Balancing.”

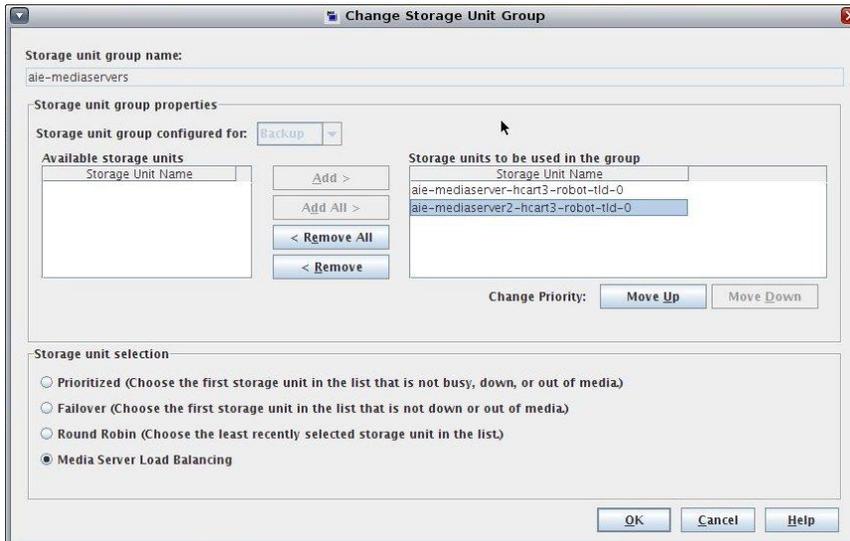


Figure 4: NetBackup New Storage Unit Group screen.

Create NetBackup Policies

For this configuration, two NetBackup policies are setup for Oracle RMAN operations. The first policy which will be called Automatic-Oracle is set up to control scheduled execution of the Oracle RMAN script. A defined schedule is created in the policy (daily full for example), the client which will execute the script is selected, and the path to the Oracle RMAN script on the selected database node (users can set up multiple nodes for script execution in case one is down, but that is not covered in this document) is supplied as the backup selection. The only purpose of the first policy is to execute the Oracle RMAN script on the defined schedule. The second policy, which will be called OracleAppPolicy, contains a Default Application Schedule and the names of each node in the Oracle Real Application Clusters (Oracle RAC). This policy is required to allow access to the NetBackup server from the Oracle RAC nodes as well as defining which storage resources are available to the Oracle RAC nodes. Since the Oracle RMAN script is executed on one of the Oracle RAC nodes, it is considered a user-defined backup and NetBackup needs a method to authorize the nodes to access the NetBackup server resources for backup jobs. The authorization is accomplished via this policy. In the OracleAppPolicy policy users specify a storage unit, define the backup window and specify the clients (the database nodes). No backup selection is specified as that is controlled by the Oracle RMAN script. It is possible to have both schedules in the same policy, but they are set up separately here for a better illustration. Please consult the NetBackup for Oracle Administrator's Guide linked in the Appendix for further details.

1. **Policy 1:** Name = Automatic-Oracle – This policy is the schedule policy that will execute the Oracle RMAN script to initiate the backup jobs
2. Set policy type to **Oracle**
3. Set Policy Storage to Storage Unit Group **aie-mediaservers**
4. Set Policy Volume Pool to the pool in use

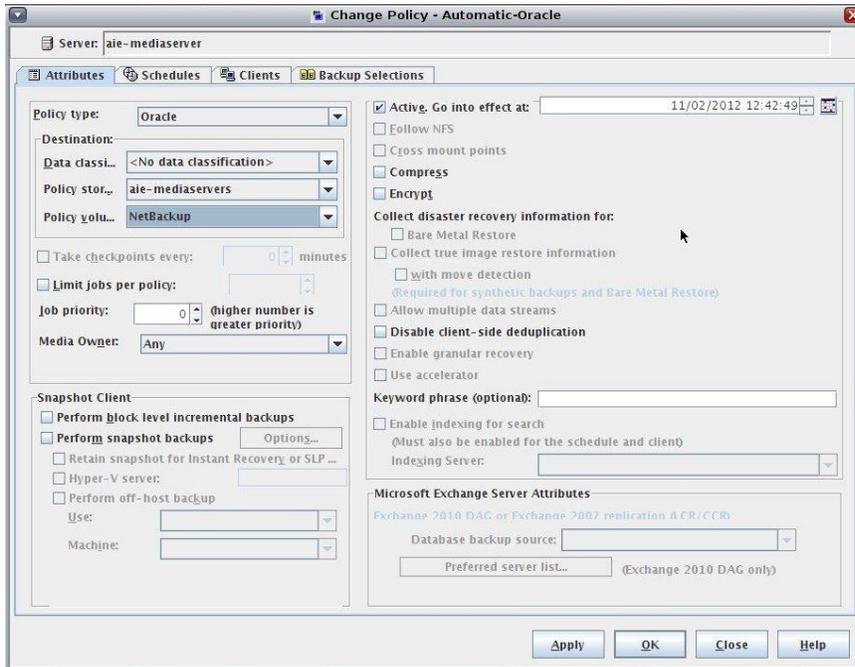


Figure 5: NetBackup Attributes screen for Add a New Policy.

Create schedules for backups. In this example we are creating a full backup with the **Automatic Full Backup** Type:

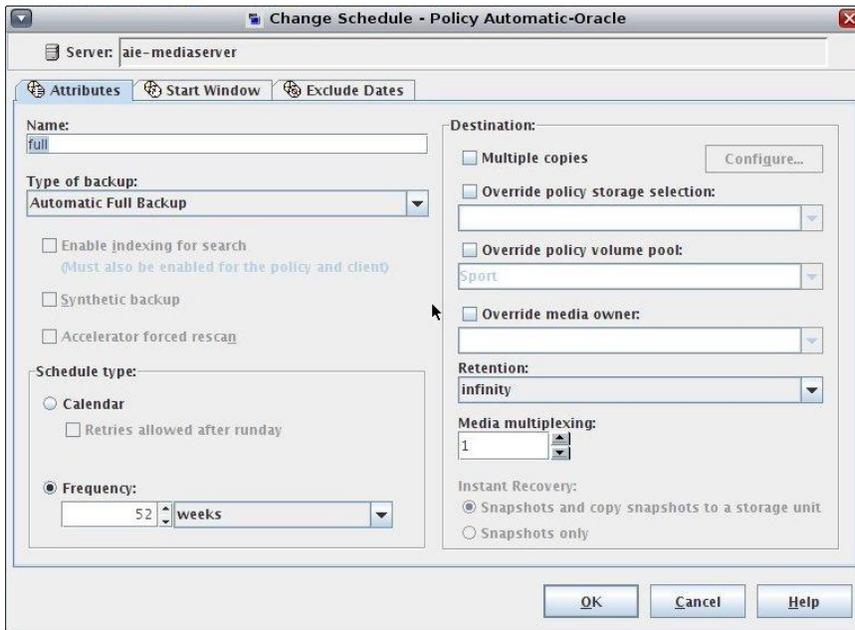


Figure 6: NetBackup Add Schedule attributes screen for Add a New Policy.

On the Start Window tab, define the **Start Window** to the time frame that is desired for backups to run (not pictured).

Define the clients. This is the host which will execute the Oracle RMAN script. In this example, **aiessc01-stor** is used to execute the Oracle RMAN script. If redundancy is desired, users can set up this entry using the SCAN name, and specify the SCAN address here. Note that other setup is also required in hosts files, bp.conf files, and specific Oracle Database files to allow the SCAN address to work. Setting up SCAN is beyond the scope of this document; please refer to Oracle documentation.

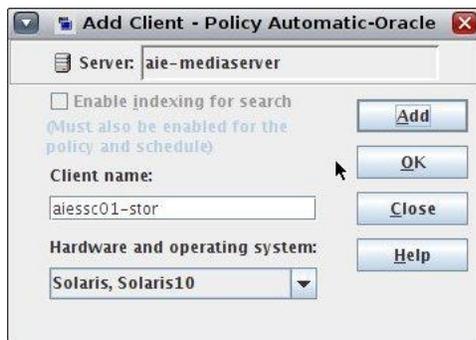


Figure 7: NetBackup Add Client screen for Add a New Policy.

Note: **aiessc01-stor** is the only client being added to the Clients tab in this example.

Define the backup selection to include the path to the Oracle RMAN script on **aiessc01-stor**:

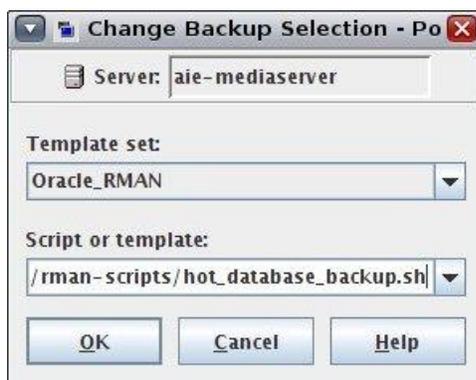


Figure 8: NetBackup Add Backup Selection screen for Add a New Policy.

1. **Policy 2:** Name = OracleAppPolicy – This is the policy that allows Oracle RMAN to execute jobs on the NetBackup server for the specified clients
2. Set policy type to **Oracle**
3. Set Policy Storage to Storage Unit Group **aie-mediaserver2**
4. Set Policy Volume Pool to the pool in use

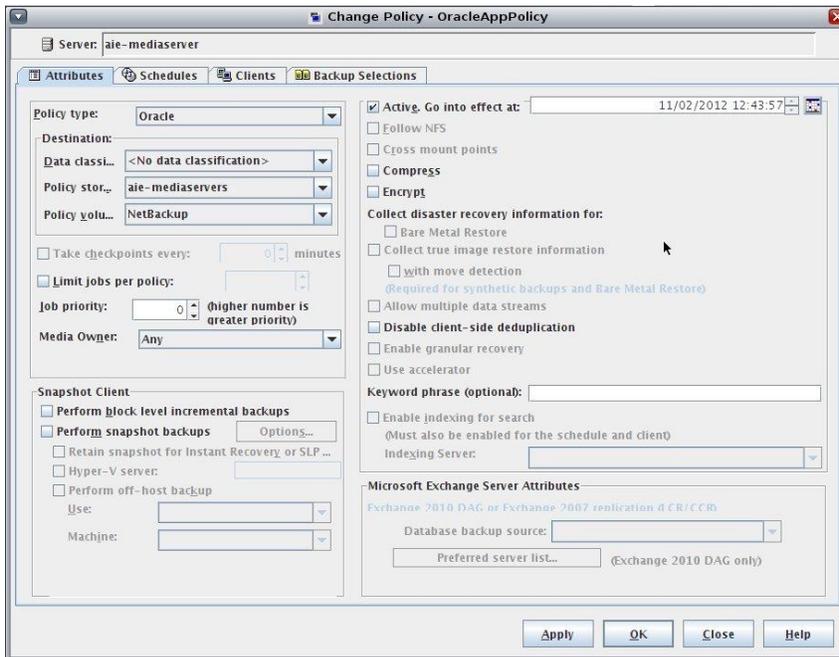


Figure 9: NetBackup Attributes screen for Add a New Policy.

Create an application backup schedule for the backups. Type **Application Backup** for Oracle RMAN must be used for this to work successfully.

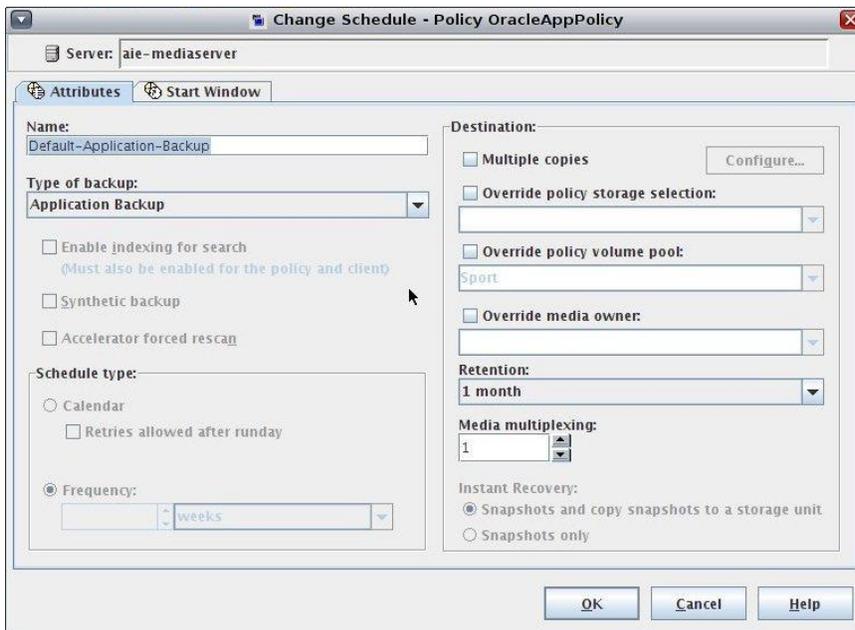


Figure 10: NetBackup Add Schedule attributes screen for Add a New Policy.

On the Start Window tab, define the start window for seven days (always open). The primary schedule is controlled by the first policy created (Automatic-Oracle), but this one should be always open to ensure it covers any time frame set in the other policy.

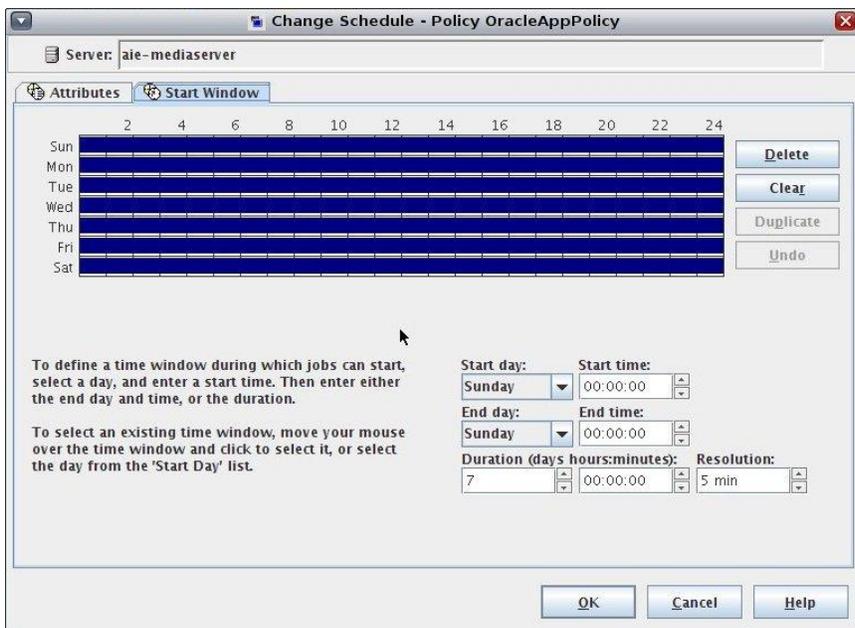


Figure 11: NetBackup Add Client screen for Add a New Policy.

Define the clients. Each compute node needs to be defined as a client in this policy:

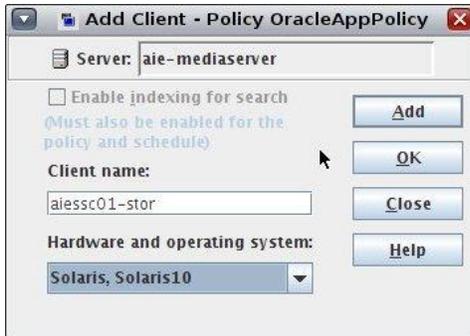


Figure 12: NetBackup Add Client screen for Add a New Policy.

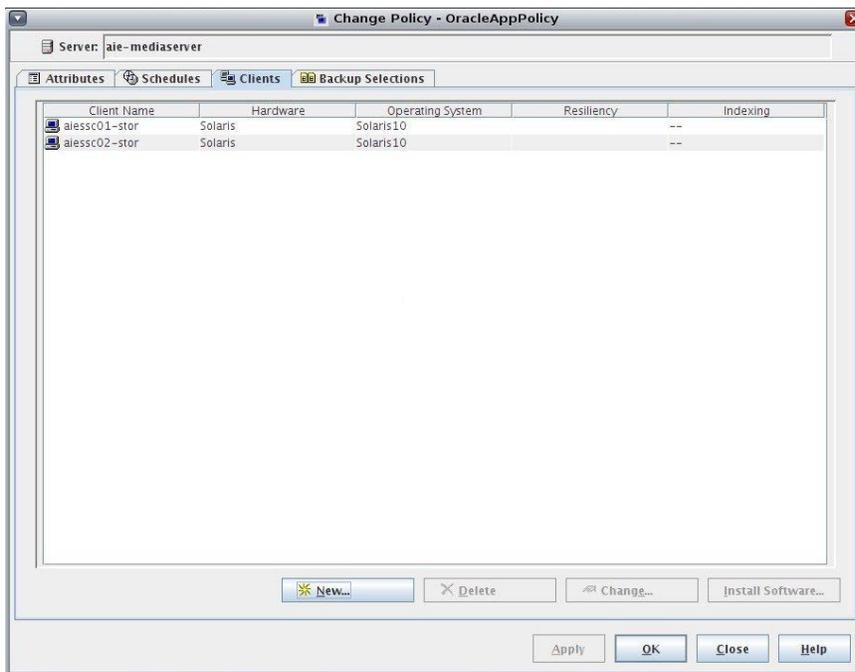


Figure 13: NetBackup Clients screen for Add a New Policy.

The Backup Selections section is left blank, as Oracle RMAN controls what is being backed up based on the settings in the Oracle RMAN script.

Verify NetBackup Clients

Navigate to the client list: NetBackup Management->Host Properties->Clients

aiessc01-stor and aiessc02-stor should be visible. Select each one and a green check box should appear indicating the connection to each client.

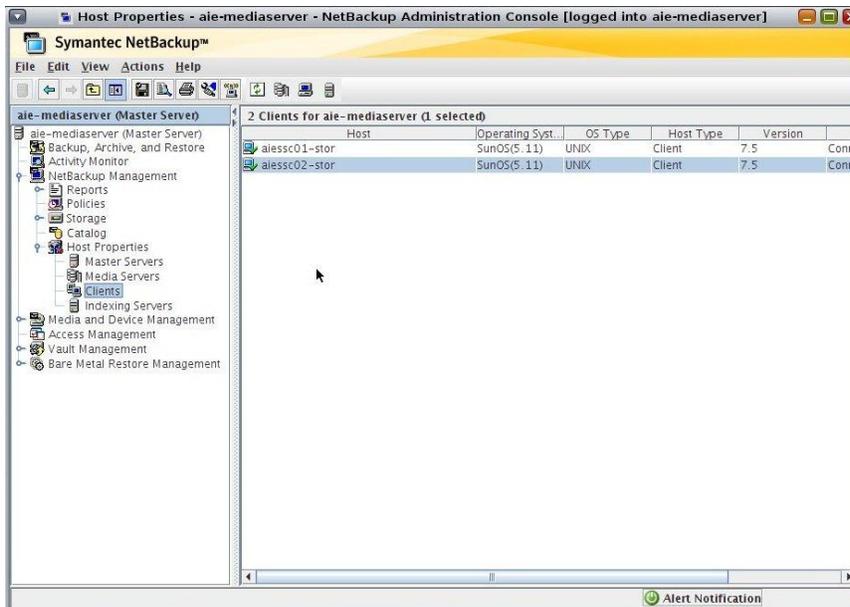


Figure 14: NetBackup Host Properties Clients screen.

NetBackup Host Property Configuration

Navigate to NetBackup Management->Host-Properties->Master Servers and change the “Maximum jobs per client” to be greater than or equal to the maximum number of active Oracle RMAN channels to be used—12 in this configuration.

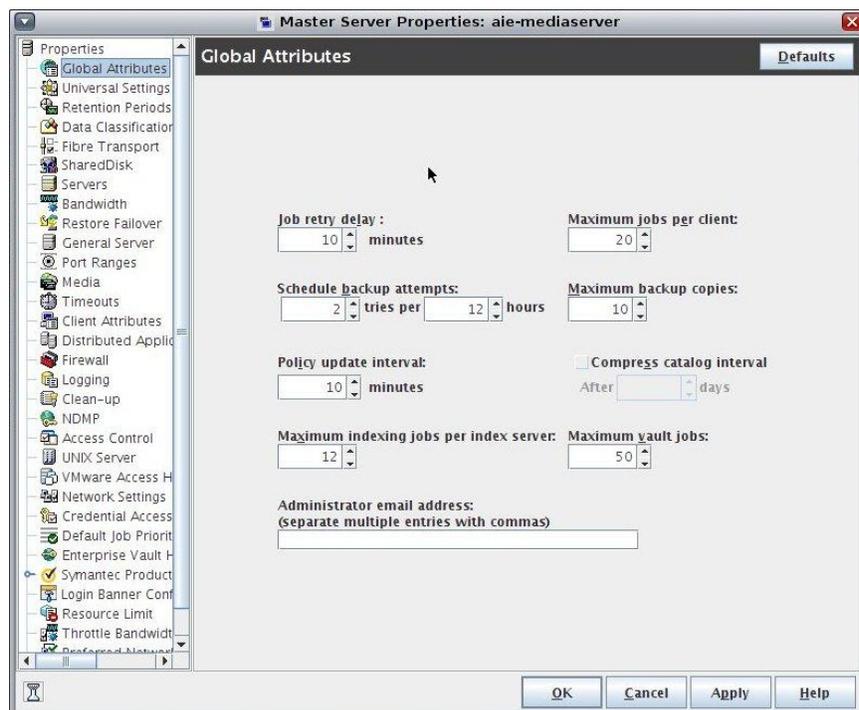


Figure 15: NetBackup Master Server Properties screen.

NetBackup Alternate Restore Configuration

When parallel backups are run across multiple compute nodes, permissions on restores become an issue because as far as NetBackup is concerned different nodes own different pieces of the backup. In the Oracle RMAN script below, four channels are allocated but to two separate nodes (aiessc01-stor and aiessc02-stor). If users try to restore the database backup from a single node, they will encounter errors accessing any pieces owned by the other node. Additionally, if users try to allocate channels to all nodes during restores without conducting the required setup in NetBackup, this will also fail. To allow this configuration to work properly, users must enable alternate restores in NetBackup for the database nodes.

On the NetBackup Master Server, execute the following steps:

1. `cd /usr/opensv/netbackup/db`
2. `mkdir altnames`
3. `cd altnames`
4. `echo aiessc01-stor >> aiessc01-stor`
5. `echo aiessc02-stor >> aiessc01-stor`
6. `cp aiessc01-stor aiessc02-stor`

The end result is two files—aiessc01-stor and aiessc02-stor—each with the same contents:

```
aiessc01-stor
aiessc02-stor
```

There are several different scenarios for alternate restores in NetBackup, and all require the altnames directory to be set up. In the restore example illustrated later in this document, channels are allocated to the same hosts that were used during backup, but this scenario assumes all nodes are present and available for use during the restore.

Users can also link the node names in the NetBackup catalog so every node can restore the other node's pieces, but they must be careful that regular filesystem backups do not get mixed in with the database backups. Please refer to the NetBackup Alternate Restores section in the Appendix for complete information about alternate restores in an Oracle RAC configuration with NetBackup:

Oracle RMAN Script

Copy the hot database backup template script from the NetBackup Oracle RMAN examples on the client to a designated location. Do this operation on one client, modify the script to match the environment, and then copy the script to the other compute node if setting up parallel backups (not covered in this document).

Example: host aiesscdb01 will be used to create the script.

1. `mkdir /rman-scripts`
2. `cp /usr/openv/netbackup/ext/db_ext/oracle/samples/rman/hot_database_backup.sh /rman-scripts/hot_database_backup.sh`
3. Ensure the Oracle RMAN script is executable for owner and group:

```
chmod 770 hot_backup_database.sh
```

Modifications for Data Files

Change the ORACLE_HOME entry to match the specific ORACLE_HOME path:

Example: **ORACLE_HOME=/u01/app/oracle/product/11.2.0.3/test1**

Change the ORACLE_SID to match the SID:

Example: **ORACLE_SID=test11**

Change environment variables to match the Oracle installation:

```
ORACLE_BASE=/u01/app/oracle  
LD_LIBRARY_PATH=$ORACLE_HOME/lib  
export ORACLE_BASE LD_LIBRARY_PATH
```

Change the ORACLE_USER to the user executing RMAN:

Example: **ORACLE_USER=oracle**

Change the TARGET_CONNECT_STR to match the authentication for the environment:

Example: **TARGET_CONNECT_STR=sys/welcome1**

Beginning in the section `CMD_STR=` modify the allocation of channels for the backup pieces. The number of channels allocated needs to match the number of tape drives on the NetBackup Media Servers (four in this example). Also note that a round robin allocation from each compute node is being set up (e.g., first channel to aiessc01-stor connecting to instance test11, second channel to aiessc02-stor connecting to instance test12). For any blank spaces between allocate channel entries there must be a pound comment sign; otherwise, problems will

occur. Also note that the policy name specified for NB_ORA_POLICY must match the policy name created in NetBackup.

```
$RMAN target $TARGET_CONNECT_STR nocatalog msglog $RMAN_LOG_FILE append << EOF
RUN {
  ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
  CONNECT='sys/welcome1@test11';
  ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
  CONNECT='sys/welcome1@test12';
  ALLOCATE CHANNEL ch02 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
  CONNECT='sys/welcome1@test11';
  ALLOCATE CHANNEL ch03 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
  CONNECT='sys/welcome1@test12';
```

Remove or comment out the FILESPERSET (default of 5) and let Oracle RMAN manage the number of files per set for optimization.

Example:

```
# FILESPERSET 5
```

Release each of the tape channels following the backup of the data files:

```
RELEASE CHANNEL ch00;
RELEASE CHANNEL ch01;
RELEASE CHANNEL ch02;
RELEASE CHANNEL ch03;
```

Modifications for Archive Logs

Modify the allocation of channels for the archive logs. Note that it is not necessary to allocate the full number of channels for the archive logs. Experiment to determine the number of drives needed to back up the archive logs in the backup window.

Example: Allocating two drives for the archive logs.

```
ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
CONNECT='sys/welcome1@test11';
ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)'
CONNECT='sys/welcome1@test12';
```

Release the channels for the same number of drives that were allocated for the archive log backups:

```
RELEASE CHANNEL ch00;
RELEASE CHANNEL ch01;
```

Modifications for Control File

For the control file section, no changes are needed. If users are using Oracle RMAN with Recovery Catalog, modify the Oracle RMAN script to work with that feature.

Execute Backup Test

NetBackup setup is now complete. Cycle the daemons on the NetBackup Master and NetBackup Media Server to ensure all configurations changes are refreshed (kill all services, restart master, and then restart media server). Execute a backup by manually executing the automatic backup schedule (**Automatic-Oracle**) and monitor the jobs.

Restoring the Database

In the event of a logical database corruption, or a complete disaster, restoration of the database from tape can be executed. While it is possible to configure and execute a restore from the NetBackup GUI (Refer to NetBackup documentation), in most instances the database administrator (DBA) will be conducting the restoration and will be executing the restore job from the machine where the database resides. The DBA will use Oracle RMAN to request the needed backup pieces from the NetBackup Media Management Layer. The Oracle RMAN backup pieces are cataloged in the controlfile, or Oracle RMAN Catalog if using Oracle Recovery Manager with catalog, and NetBackup also knows about the pieces in its own catalog. If a complete disaster has occurred, the Oracle SPARC SuperCluster has been reinstalled, and the Oracle RMAN catalog is not being used, users will also require the DBID of the Oracle Database to perform the restore.

Restore Example

There are many different restore scenarios available with Oracle RMAN, depending on the type of problem encountered in the Oracle Database (consult Oracle RMAN documentation for complete information). A tablespace will be restored in this example to illustrate how to retrieve tape data from Oracle RMAN. There are many options available for a tablespace restore in Oracle Database (Oracle Flashback, Fast Recovery Area on disk, etc.). This example is only intended to demonstrate how to retrieve Oracle RMAN tape backups. In a real world situation, if a disaster occurred, the spfile, controlfile, all database objects and the archivelogs would be restored.

Example of steps to restore a tablespace in the Oracle Database:

```
oracle@aiesscdb01:/rman-scripts$ rman target / cmdfile restorebackupset
```

Recovery Manager: Release 11.2.0.3.0 - Production on Tue Nov 13 16:46:52 2012

Copyright (c) 1982, 2011, Oracle and/or its affiliates. All rights reserved.

connected to target database: TEST1 (DBID=1166807773)

```
RMAN> RUN {  
2> SQL 'alter tablespace CARD offline immediate';  
3> SET AUTOLOCATE ON;
```

```

4> ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
5> ALLOCATE CHANNEL ch01 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
6> ALLOCATE CHANNEL ch02 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
7> ALLOCATE CHANNEL ch03 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
8> ALLOCATE CHANNEL ch04 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
9> ALLOCATE CHANNEL ch05 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
10> ALLOCATE CHANNEL ch06 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
11> ALLOCATE CHANNEL ch07 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
12> ALLOCATE CHANNEL ch08 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
13> ALLOCATE CHANNEL ch09 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
14> ALLOCATE CHANNEL ch10 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc01-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
15> ALLOCATE CHANNEL ch11 TYPE 'SBT_TAPE' PARMS='ENV=(NB_ORA_CLIENT=aiessc02-
stor,NB_ORA_SERV=aie-mediaserver-ib,NB_ORA_POLICY=OracleAppPolicy)' CONNECT=*;
16> RESTORE TABLESPACE CARD;
17> RECOVER TABLESPACE CARD;
18> sql 'alter tablespace card online';
19> RELEASE CHANNEL ch00;
20> RELEASE CHANNEL ch01;
21> RELEASE CHANNEL ch02;
22> RELEASE CHANNEL ch03;
23> RELEASE CHANNEL ch04;
24> RELEASE CHANNEL ch05;
25> RELEASE CHANNEL ch06;
26> RELEASE CHANNEL ch07;
27> RELEASE CHANNEL ch08;
28> RELEASE CHANNEL ch09;
29> RELEASE CHANNEL ch10;
30> RELEASE CHANNEL ch11;
31> }
32>
33>
34>

```

using target database control file instead of recovery catalog
sql statement: alter tablespace CARD offline immediate

executing command: SET autolocate

allocated channel: ch00
channel ch00: SID=898 instance=test11 device type=SBT_TAPE
channel ch00: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch01
channel ch01: SID=859 instance=test12 device type=SBT_TAPE
channel ch01: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch02
channel ch02: SID=911 instance=test11 device type=SBT_TAPE
channel ch02: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch03
channel ch03: SID=885 instance=test12 device type=SBT_TAPE
channel ch03: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch04
channel ch04: SID=950 instance=test11 device type=SBT_TAPE
channel ch04: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch05
channel ch05: SID=898 instance=test12 device type=SBT_TAPE
channel ch05: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch06
channel ch06: SID=963 instance=test11 device type=SBT_TAPE
channel ch06: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch07
channel ch07: SID=911 instance=test12 device type=SBT_TAPE
channel ch07: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch08
channel ch08: SID=976 instance=test11 device type=SBT_TAPE
channel ch08: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch09
channel ch09: SID=924 instance=test12 device type=SBT_TAPE
channel ch09: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch10
channel ch10: SID=989 instance=test11 device type=SBT_TAPE
channel ch10: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

allocated channel: ch11
channel ch11: SID=937 instance=test12 device type=SBT_TAPE
channel ch11: Veritas NetBackup for Oracle - Release 7.5 (2012020723)

Starting restore at NOV 13, 2012 16:47:17

channel ch00: starting datafile backup set restore
channel ch00: specifying datafile(s) to restore from backup set
channel ch00: restoring datafile 00006 to +DATA_AIESSC/test1/load1/card.dbf
channel ch00: restoring datafile 00021 to +DATA_AIESSC/test1/load1/card15.dbf
channel ch00: restoring datafile 00036 to +DATA_AIESSC/test1/load1/card30.dbf

channel ch00: reading from backup piece mnnq76ef_1_2
channel ch01: starting datafile backup set restore
channel ch01: specifying datafile(s) to restore from backup set
channel ch01: restoring datafile 00010 to +DATA_AIESSC/test1/load1/card4.dbf
channel ch01: restoring datafile 00025 to +DATA_AIESSC/test1/load1/card19.dbf
channel ch01: restoring datafile 00040 to +DATA_AIESSC/test1/load1/card34.dbf
channel ch01: reading from backup piece mrnq76eh_1_2
channel ch02: starting datafile backup set restore
channel ch02: specifying datafile(s) to restore from backup set
channel ch02: restoring datafile 00008 to +DATA_AIESSC/test1/load1/card2.dbf
channel ch02: restoring datafile 00023 to +DATA_AIESSC/test1/load1/card17.dbf
channel ch02: restoring datafile 00038 to +DATA_AIESSC/test1/load1/card32.dbf
channel ch02: reading from backup piece mpnq76eg_1_2
channel ch03: starting datafile backup set restore
channel ch03: specifying datafile(s) to restore from backup set
channel ch03: restoring datafile 00011 to +DATA_AIESSC/test1/load1/card5.dbf
channel ch03: restoring datafile 00026 to +DATA_AIESSC/test1/load1/card20.dbf
channel ch03: restoring datafile 00041 to +DATA_AIESSC/test1/load1/card35.dbf
channel ch03: reading from backup piece msnq76eh_1_2
channel ch04: starting datafile backup set restore
channel ch04: specifying datafile(s) to restore from backup set
channel ch04: restoring datafile 00009 to +DATA_AIESSC/test1/load1/card3.dbf
channel ch04: restoring datafile 00024 to +DATA_AIESSC/test1/load1/card18.dbf
channel ch04: restoring datafile 00039 to +DATA_AIESSC/test1/load1/card33.dbf
channel ch04: reading from backup piece mqnq76eg_1_2
channel ch05: starting datafile backup set restore
channel ch05: specifying datafile(s) to restore from backup set
channel ch05: restoring datafile 00017 to +DATA_AIESSC/test1/load1/card11.dbf
channel ch05: restoring datafile 00032 to +DATA_AIESSC/test1/load1/card26.dbf
channel ch05: restoring datafile 00047 to +DATA_AIESSC/test1/load1/card41.dbf
channel ch05: restoring datafile 00056 to +DATA_AIESSC/test1/load1/card50.dbf
channel ch05: reading from backup piece mnnq76ef_1_2
channel ch06: starting datafile backup set restore
channel ch06: specifying datafile(s) to restore from backup set
channel ch06: restoring datafile 00007 to +DATA_AIESSC/test1/load1/card1.dbf
channel ch06: restoring datafile 00022 to +DATA_AIESSC/test1/load1/card16.dbf
channel ch06: restoring datafile 00037 to +DATA_AIESSC/test1/load1/card31.dbf
channel ch06: reading from backup piece monq76eg_1_2
channel ch07: starting datafile backup set restore
channel ch07: specifying datafile(s) to restore from backup set
channel ch07: restoring datafile 00013 to +DATA_AIESSC/test1/load1/card7.dbf
channel ch07: restoring datafile 00028 to +DATA_AIESSC/test1/load1/card22.dbf
channel ch07: restoring datafile 00043 to +DATA_AIESSC/test1/load1/card37.dbf
channel ch07: restoring datafile 00052 to +DATA_AIESSC/test1/load1/card46.dbf
channel ch07: reading from backup piece minq76ee_1_2
channel ch08: starting datafile backup set restore
channel ch08: specifying datafile(s) to restore from backup set
channel ch08: restoring datafile 00016 to +DATA_AIESSC/test1/load1/card10.dbf
channel ch08: restoring datafile 00031 to +DATA_AIESSC/test1/load1/card25.dbf
channel ch08: restoring datafile 00046 to +DATA_AIESSC/test1/load1/card40.dbf
channel ch08: restoring datafile 00055 to +DATA_AIESSC/test1/load1/card49.dbf

channel ch08: reading from backup piece mlnq76ef_1_2
channel ch09: starting datafile backup set restore
channel ch09: specifying datafile(s) to restore from backup set
channel ch09: restoring datafile 00012 to +DATA_AIESSC/test1/load1/card6.dbf
channel ch09: restoring datafile 00027 to +DATA_AIESSC/test1/load1/card21.dbf
channel ch09: restoring datafile 00042 to +DATA_AIESSC/test1/load1/card36.dbf
channel ch09: restoring datafile 00051 to +DATA_AIESSC/test1/load1/card45.dbf
channel ch09: reading from backup piece mhnq76ee_1_2
channel ch10: starting datafile backup set restore
channel ch10: specifying datafile(s) to restore from backup set
channel ch10: restoring datafile 00014 to +DATA_AIESSC/test1/load1/card8.dbf
channel ch10: restoring datafile 00029 to +DATA_AIESSC/test1/load1/card23.dbf
channel ch10: restoring datafile 00044 to +DATA_AIESSC/test1/load1/card38.dbf
channel ch10: restoring datafile 00053 to +DATA_AIESSC/test1/load1/card47.dbf
channel ch10: reading from backup piece mjnq76ee_1_2
channel ch11: starting datafile backup set restore
channel ch11: specifying datafile(s) to restore from backup set
channel ch11: restoring datafile 00015 to +DATA_AIESSC/test1/load1/card9.dbf
channel ch11: restoring datafile 00030 to +DATA_AIESSC/test1/load1/card24.dbf
channel ch11: restoring datafile 00045 to +DATA_AIESSC/test1/load1/card39.dbf
channel ch11: restoring datafile 00054 to +DATA_AIESSC/test1/load1/card48.dbf
channel ch11: reading from backup piece mknq76ef_1_2
channel ch00: restored backup piece 1
channel ch00: restore complete, elapsed time: 00:07:10
channel ch00: starting datafile backup set restore
channel ch00: specifying datafile(s) to restore from backup set
channel ch00: restoring datafile 00019 to +DATA_AIESSC/test1/load1/card13.dbf
channel ch00: restoring datafile 00034 to +DATA_AIESSC/test1/load1/card28.dbf
channel ch00: restoring datafile 00049 to +DATA_AIESSC/test1/load1/card43.dbf
channel ch00: reading from backup piece munq76pp_1_2
channel ch02: piece handle=mpnq76eg_1_2 tag=TAG20121113T140812
channel ch02: restored backup piece 1
channel ch02: restore complete, elapsed time: 00:08:16
channel ch02: starting datafile backup set restore
channel ch02: specifying datafile(s) to restore from backup set
channel ch02: restoring datafile 00020 to +DATA_AIESSC/test1/load1/card14.dbf
channel ch02: restoring datafile 00035 to +DATA_AIESSC/test1/load1/card29.dbf
channel ch02: restoring datafile 00050 to +DATA_AIESSC/test1/load1/card44.dbf
channel ch02: reading from backup piece mvnq76q9_1_2
channel ch05: piece handle=mmnq76ef_1_2 tag=TAG20121113T140812
channel ch05: restored backup piece 1
channel ch05: restore complete, elapsed time: 00:12:04
channel ch05: starting datafile backup set restore
channel ch05: specifying datafile(s) to restore from backup set
channel ch05: restoring datafile 00018 to +DATA_AIESSC/test1/load1/card12.dbf
channel ch05: restoring datafile 00033 to +DATA_AIESSC/test1/load1/card27.dbf
channel ch05: restoring datafile 00048 to +DATA_AIESSC/test1/load1/card42.dbf
channel ch05: reading from backup piece mtnq76pp_1_2
channel ch08: piece handle=mlnq76ef_1_2 tag=TAG20121113T140812
channel ch08: restored backup piece 1
channel ch08: restore complete, elapsed time: 00:14:11

channel ch03: piece handle=msnq76eh_1_2 tag=TAG20121113T140812
channel ch03: restored backup piece 1
channel ch03: restore complete, elapsed time: 00:14:32
channel ch11: piece handle=mknq76ef_1_2 tag=TAG20121113T140812
channel ch11: restored backup piece 1
channel ch11: restore complete, elapsed time: 00:14:52
channel ch01: piece handle=mrnq76eh_1_2 tag=TAG20121113T140812
channel ch01: restored backup piece 1
channel ch01: restore complete, elapsed time: 00:15:02
channel ch04: piece handle=mqnq76eg_1_2 tag=TAG20121113T140812
channel ch04: restored backup piece 1
channel ch04: restore complete, elapsed time: 00:15:02
channel ch09: piece handle=mhnq76ee_1_2 tag=TAG20121113T140812
channel ch09: restored backup piece 1
channel ch09: restore complete, elapsed time: 00:15:12
channel ch06: piece handle=monq76eg_1_2 tag=TAG20121113T140812
channel ch06: restored backup piece 1
channel ch06: restore complete, elapsed time: 00:15:23
channel ch02: piece handle=mvnq76q9_1_2 tag=TAG20121113T140812
channel ch02: restored backup piece 1
channel ch02: restore complete, elapsed time: 00:08:36
channel ch07: piece handle=minq76ee_1_2 tag=TAG20121113T140812
channel ch07: restored backup piece 1
channel ch07: restore complete, elapsed time: 00:17:03
channel ch10: piece handle=mjnq76ee_1_2 tag=TAG20121113T140812
channel ch10: restored backup piece 1
channel ch10: restore complete, elapsed time: 00:17:23
channel ch00: piece handle=munq76pp_1_2 tag=TAG20121113T140812
channel ch00: restored backup piece 1
channel ch00: restore complete, elapsed time: 00:12:24
channel ch05: piece handle=mtnq76pp_1_2 tag=TAG20121113T140812
channel ch05: restored backup piece 1
channel ch05: restore complete, elapsed time: 00:07:50
Finished restore at NOV 13, 2012 17:07:18

Starting recover at NOV 13, 2012 17:07:18

starting media recovery
media recovery complete, elapsed time: 00:00:05

Finished recover at NOV 13, 2012 17:07:27

sql statement: alter tablespace card online

released channel: ch00
released channel: ch01
released channel: ch02
released channel: ch03
released channel: ch04
released channel: ch05
released channel: ch06

released channel: ch07
 released channel: ch08
 released channel: ch09
 released channel: ch10
 released channel: ch11

Recovery Manager complete.

NetBackup Tuning

The NetBackup default settings for transferring data to tape are set generically and must be tuned to match the environment, in order to achieve maximum transfer rates.

Identifying NetBackup Bottlenecks

Please refer to the Symantec NetBackup Planning and Performance Tuning Guide link in the Appendix for additional details.

- Enable the NetBackup BPTM log on each of the NetBackup Media Servers (mkdir /usr/opensv/netbackup/logs/bptm on each media server and cycle daemons).
- Setup a test policy to backup data from one of the Oracle Database servers to identify the base transfer rate. This can be either filesystem data on the Oracle Database server or an Oracle RMAN backup.
- Once the job completes run the following commands to retrieve the pertinent information from the BPTM log:
 - more /usr/opensv/netbackup/logs/bptm/log.DATE | grep -i kbytes/sec (Note: Using BPTM provides transfer rate minus time required to mount and position).
 - Example:
 write_backup: successfully wrote backup id aiesc01-stor_1293487677, copy 1, fragment 1, 221292576 Kbytes at 50262.248 Kbytes/sec
- Analyze the BPTM log to identify if/when NetBackup is users waiting for “full buffers”
 - more /usr/opensv/netbackup/logs/bptm/log.DATE | grep -i "waited for full"
 - Example:
 fill_buffer: [16711] socket is closed, waited for empty buffer 251258 times, delayed 504798 times, read 241125152 Kbytes
 write data: waited for full buffer 251067 times, delayed 253235 times
- If there is excessive waits for full buffers, the NUMBER_DATA_BUFFERS and SIZE_DATA_BUFFERS files should be modified to reduce the waits. These files are text files stored in the /usr/opensv/netbackup/db/config directory on each NetBackup Media Server. Note that they do not exist by default and should be created. Each file contains a single numeric value for the setting. The NUMBER_DATA_BUFFERS is an integer value and the SIZE_DATA_BUFFERS is also an integer value, but must be specified in increments of 1024 (ex. 32K = 32768).
- Modify the NUMBER_DATA_BUFFERS, SIZE_DATA_BUFFERS and execute additional backup jobs to measure performance. Check BPTM for full buffers and repeat procedure until no further speed improvements are achieved.
- If the BPTM log shows waiting for empty buffers, then data is arriving faster than it can be written to tape. In this scenario tune the NUMBER_DATA_BUFFERS to see if the waits can be reduced. When users are

finished changing tuning parameters, remember to disable the NetBackup logging, as there is overhead associated with logging.

Note: Results after tuning modification where applied.

```
more /usr/opensv/netbackup/logs/bptm/log.DATE | grep -i "waited for full"
```

```
fill_buffer: [7808] socket is closed, waited for empty buffer 118248 times, delayed 120304 times, read 215323680 Kbytes
```

```
write_data: waited for full buffer 12 times, delayed 1673 times
```

```
more /usr/opensv/netbackup/logs/bptm/log.DATE | grep -i kbytes/sec
```

```
write_backup: successfully wrote backup id aiessc01-stor_1293493455, copy 1, fragment 1, 241125152 Kbytes at 345626.425 Kbytes/sec
```

Hardware Compression

By default most tape drives have compression enabled, and it is recommended to compress at the tape drive level. Users can experiment with compression using Oracle RMAN, but if they opt to enable compression at a different layer of the backup architecture, they must disable tape drive compression to avoid doubling compression, which can result in larger backups and slower transfer rates.

Recommended Tape Drive Tuning

The tuning parameters used in this environment are leveraged from the whitepaper, Protecting Exadata – Tape Backup with Symantec NetBackup.

The two variables used to tune NetBackup using the StorageTek T10000C tape drive are:

```
NUMBER_DATA_BUFFERS    256
SIZE_DATA_BUFFERS      2097152
```

NetBackup Performance

Data from the Oracle SPARC SuperCluster can be backed up directly to tape or from Oracle's Sun ZFS Storage Appliance. The Sun ZFS Storage Appliance can be used to back up the data from the Oracle SPARC SuperCluster and the backup sets can then be backed up to tape as well. Restore of the backup sets residing on tape can be restored directly to the Oracle SPARC SuperCluster.

The table below shows the average transfer rates per drive and the aggregate throughput.

TABLE2. PERFORMANCE STATISTICS

DATA TYPE	TOTAL BACKUP AGGREGATE THROUGHPUT WITH HARDWARE COMPRESSION	AVERAGE BACKUP TRANSFER RATE PER TAPE DRIVE WITH HARDWARE COMPRESSION	TOTAL RESTORE AGGREGATE THROUGHPUT WITH HARDWARE COMPRESSION	AVERAGE RESTORE TRANSFER RATE PER TAPE DRIVE WITH HARDWARE COMPRESSION
Database	15.0 TB/hr	347 MB/s	5.8 TB/hr	135 MB/s

Sun ZFS Storage Appliance Backup Set	10.5 TB/hr	245 MB/s	6.0 TB/hr	138 MB/s
--------------------------------------	------------	----------	-----------	----------

NetBackup Troubleshooting

The following is a list of some common issues which may be encountered:

Slow Transfer Rates

- Verify the MTU settings (ex. `ifconfig -ib0 | grep MTU`) on each NetBackup Media Server. The MTU setting should be set to 65520.
 - Run the “`ifconfig -a`” command to verify MTU size.
- If the transfer rate is good when one drive is running, and adding drives slows the transfer rate, then the MTU setting may be set incorrectly. Verify the MTU is set to 65520 on each of the compute nodes.
 - Run “`ifconfig -a`” to verify MTU size.
- Verify the data is transferring through the InfiniBand interface. There should be activity on bond 0 or bond 1 (whichever is active). If there is no traffic, then the InfiniBand interface is not configured correctly. Verify the InfiniBand port light is green. If it is not green, then run `ifconfig up/down` on the interface. If it is still not green, then reinstall all Oracle Solaris InfiniBand packages.
 - `snoop` command
- Make sure application compression is disabled (for NetBackup it is set in the policy).
 - Run “`prstat`” which will display CPU usage. If CPU usage is at, or close to 100 percent for “oracle” then application or Oracle RMAN compression may be enabled.
- If the MTU settings, the InfiniBand connections, and the compression settings are verified, but performance is still not what is expected, refer to the tuning section of this document for further settings to check.

NetBackup Spawning Excessive Jobs Versus Number of Available Drives

- Check the Oracle RMAN script to see if the `FILESERSET` parameter is being defined. If it is defined, remove the setting. This command tells Oracle RMAN how many files to include in each backup piece. When this parameter is not defined, Oracle RMAN will determine the optimal number of pieces to include in each backup file, and performance will be improved. It is recommended to leave the setting undefined.

Conclusion

In conclusion, Oracle’s StorageTek tape products offer cost-effective data protection for the Oracle SPARC SuperCluster and, when coupled with Symantec NetBackup, provide a complete solution for data protection.

Appendix

- [General SPARC SuperCluster T4-4 Information](#)
- [Protecting Exadata-Tape Backup with Symantec NetBackup](#)
- [Symantec NetBackup 7.5 Administrator's Guide for Unix and Linux, Vol. 1](#)
- [Symantec NetBackup 7.5 Device Configuration Guide](#)
- [Symantec NetBackup 7.5 Installation Guide for Unix and Linux](#)
- [Symantec NetBackup 7.5 for Oracle Administrator's Guide](#)
- [Symantec NetBackup 7.0-7.1 Backup Planning and Performance Tuning Guide](#)



Protecting SPARC SuperCluster—Tape Backup
with Symantec NetBackup
January 2013
Author: Hung Lieu – Integrated Storage
Relations

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. NetBackup is a trademark or registered trademark of Symantec Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0113

Hardware and Software, Engineered to Work Together