



---

ZFS STORAGE  
APPLIANCE

An Oracle Technical White Paper  
June 2015

# NDMP Implementation Guide for the Oracle ZFS Storage Appliance

## Table of Contents

Introduction .....	4
Documentation Conventions.....	5
Network Data Management Protocol (NDMP) Overview.....	6
Oracle ZFS Storage Appliance NDMP Support .....	6
Supported NDMP Topologies .....	6
Supported NDMP Features .....	7
Selective File Restore.....	7
Direct Access Recovery (DAR) .....	7
ZFS-NDMP Backup .....	7
Multi-Stream Concurrent Backups .....	8
Incremental and Cumulative Level Backups .....	8
Token Based Backups.....	8
Backup of Replication Targets .....	8
NDMP File Service Extensions .....	9
Data Encryption .....	9
NDMP Formats.....	10
Tar and Dump NDMP Format Use.....	10
Specifying a Backup Path .....	10
Backing Up a Snapshot.....	11
Restoring Backup Data.....	12
ZFS-NDMP Format Use .....	13
Passing NDMP Parameters .....	13
ZFS-NDMP System Snapshot Concepts .....	16
Specifying a Backup Path .....	17
Backing Up a Snapshot.....	18
Backing Up a Replication Target.....	19
Restoring Backup Data.....	20

Configuring NDMP on the Oracle ZFS Storage Appliance .....	24
Configuring the Fibre Channel Interface for NDMP .....	24
Configuring the NDMP Service .....	24
Attaching Fibre Channel Tape Drives and Tape Libraries .....	27
Monitoring Oracle ZFS Storage Appliance NDMP .....	28
Status Dashboard .....	28
Status NDMP .....	30
Analytics .....	31
Supported Backup Applications .....	33
Best Practices .....	35
Oracle ZFS Storage Appliance Cluster Configuration .....	35
Deciding When to Use ZFS-NDMP .....	36
Backing Up ZFS-NDMP Projects or Datasets .....	36
Replication Target Backup Considerations .....	37
Conclusion .....	39
Appendix A: Oracle Secure Backup ZFS-NDMP Configuration.....	40
Adding an NDMP Host.....	40
Setting NDMP Parameters .....	40
Configuring Tape Storage Devices .....	41
Configuring Disk Storage Devices .....	44
Creating a Backup Dataset.....	44
Performing a Backup .....	45
Performing a Restore .....	45
Managing Instances .....	45
Appendix B: Symantec NetBackup ZFS-NDMP Configuration.....	47
Specifying an NDMP Host .....	47
Configuring Storage Devices .....	48
Inventorying the Tape Library .....	52
Creating a Backup Policy.....	53
Setting Maximum Jobs Per Client.....	56
Restoring an NDMP Backup.....	56
Appendix C: EMC NetWorker ZFS-NDMP Configuration .....	58
Configuring a Storage Node .....	58
Labeling Tapes.....	61

Adding an NDMP Client Definition .....	62
Changing NDMP Parameters in the Client Definition .....	65
Restoring from an NDMP Backup.....	66
Appendix D: References.....	71

## Introduction

With the large number of backup configurations and applications available today, trying to decide how best to design a data protection architecture for important business data can be confusing. One data protection technology widely available in backup applications and Network-Attached Storage (NAS) appliances is the Network Data Management Protocol (NDMP). NDMP provides a simplified mechanism to transfer data directly from a source to a target without having to be processed through an intermediary backup application media server, thus potentially increasing efficiency and reducing network utilization. The Oracle ZFS Storage Appliance implements the NDMP standard and has been qualified with many of the leading backup applications. Additionally, the Oracle ZFS Storage Appliance has an enhanced NDMP mode (ZFS-NDMP) which takes advantage of the underlying system architecture. ZFS-NDMP is capable of protecting both shares and block storage, and can potentially greatly increase the backup performance of millions-of-files shares.

This paper explores the NDMP features supported by the Oracle ZFS Storage Appliance as of the OS 8.4.0 software release. It includes information regarding the configuration of the NDMP service, attaching tape devices, and monitoring NDMP processes. The ZFS-NDMP mode and its requirements are described in detail, including Appendices that walk through the configuration of those backup applications which support its implementation.

## Documentation Conventions

This document makes many references to the Oracle ZFS Storage Appliance Browser User Interface (BUI). The navigation directive through the BUI will be in the format:

BUI -> Top Menu Selection -> Sub Menu Selection -> Sub Item 1 -> Sub Item 2 -> ...

“BUI” represents the starting point after logging in to the BUI.

“Top Menu Selection” represents the menu options in the darker grey shaded area near the very top of the BUI (white lettering, which turns yellow when selected).

“Sub Menu Selection” represents the menu options just below the Top Menu Selection corresponding with the Top Menu Selection picked.

“Sub Item” represents the dynamic content that displays below the Sub Menu Selection corresponding with the Top Menu Selection and Sub Menu Selection picked. The Sub Item can either be a known value such as “Snapshots,” or can be a user configuration specific value such as “<share name>”. It can also include a directive such as “double-click.”

For example, to navigate to the screen provided in Figure A, the directive may be:

BUI -> Shares -> Projects -> <double-click desired project> -> Filesystems

which would be interpreted as:

1. Log in to the Oracle ZFS Storage Appliance BUI.
2. Click on “Shares” in the top menu.
3. Click on “Projects” in the sub menu.
4. Double-click on the project that will be viewed.
5. Click on “Filesystems”.

The “->” directive shorthand is used to enhance document readability.



The screenshot shows the Oracle ZFS Storage Appliance BUI interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares' (highlighted), 'Status', and 'Analytics'. Below this, there are sub-menus for 'SHARES', 'PROJECTS', 'ENCRYPTION', and 'SCHEMA'. The main content area is titled 'Projects' and shows 'testproject' selected. Underneath, there are tabs for 'Shares', 'General', 'Protocols', 'Access', 'Snapshots', and 'Replication'. The 'Shares' tab is active, showing a table of filesystems for 'testproject'.

NAME	SIZE	MOUNTPOINT	ENCRYPTED
testshare1	44.9K	/export/testshare1	
testshare2	44.9K	/export/testshare2	
testshare3	44.9K	/export/testshare3	

Figure A. Share listing example

## Network Data Management Protocol (NDMP) Overview

NDMP is an open standard protocol for network-based backup of Network-Attached Storage (NAS). It enables backup application software and NAS systems to interoperate without creating specific integrations for each software vendor and storage array, allowing for the storage array to be “backup-ready” without requiring backup client software to be installed on the NAS system.

The Oracle ZFS Storage Appliance supports NDMP versions 3 and 4. At the time of this writing, most backup applications have switched to NDMP version 4. The testing in support of this document was done using NDMP version 4.

## Oracle ZFS Storage Appliance NDMP Support

NDMP provides an efficient, non-intrusive mechanism for backing up and restoring shares and LUNs on the Oracle ZFS Storage Appliance. It is possible to back up shares without the use of NDMP by having the shares mounted directly to backup application clients or media servers, but this non-NDMP method leads to greater network and server resource usage.

NDMP offers the following advantages:

- Tape drives can be attached, using a Fibre Channel SAN, directly to the Oracle ZFS Storage Appliance for faster performance.
- The tape library robotics control can either be directly attached (Fibre Channel) to the Oracle ZFS Storage Appliance, or separately connected to the backup application’s media server.
- All file system attributes, permissions, share settings, and metadata are backed up.
- The NDMP service on the Oracle ZFS Storage Appliance takes an automatic snapshot to ensure consistency of data. The backup action operates on the snapshot, thereby allowing normal read/write activity to continue on the share(s).
- The backup application media server does not need filesystem-level access to the NAS shares. An NDMP-specific user-defined userid and password grants access to the backup and restore of the share(s), but does not grant access to the data on the shares.

### Supported NDMP Topologies

The Oracle ZFS Storage Appliance supports the following NDMP topologies:

- Local – The tape drive is connected to the Oracle ZFS Storage Appliance through Fibre Channel. Backup data moves directly from the dataset to the tape drive without going over the network to a backup server. This method usually provides the best performance. The tape drive can either be dedicated to the system or shared across a Storage Area Network (SAN).
- Remote – The tape drive or storage device (disk storage, Virtual Tape Library) is connected to a backup application media server. Data from the Oracle ZFS Storage

Appliance travels over the network to the backup media server, which then transfers the data to tape.

- **3-Way** – A tape drive connected to one Oracle ZFS Storage Appliance is used to store data from another Oracle ZFS Storage Appliance. The data travels over the network to the system controlling the tape resource. This is a specific variation of remote backup. However, this topology is only supported from one Oracle ZFS Storage Appliance to another. Oracle does not certify or support this configuration with any other vendor's NAS appliance.
- **Passthrough** – Data is written directly from a backup application media server over a network to the Oracle ZFS Storage Appliance attached tape drive. This is not a common topology and is only supported by a limited number of backup applications.

### Supported NDMP Features

The following sections outline the NDMP features that are supported by the Oracle ZFS Storage Appliance. The backup application must implement the particular feature in order for it to be available. Refer to the "Supported Backup Applications" section for details.

#### Selective File Restore

Selective File Restore enables the backup application to know the file level contents of the backup, and to restore independent files of the backup without having to restore the entire backup image. During the backup, the file history is provided to the backup application. During a selective file (or files) restore, the backup application reads from the beginning of the backup image until it locates the targeted file(s), and restores it to the desired location.

#### Direct Access Recovery (DAR)

DAR is a feature of NDMP that can significantly reduce the time it takes to perform individual file restoration. As an enhancement to Selective File Restore, DAR also provides the backup application with positioning information of the files on tape. Rather than starting at the beginning of the backup image and searching until the item is found, the backup application can directly seek to the file's location on the tape and restore it. This is particularly helpful with multi-tape volume backup image restores.

DAR is typically the default setting for backup applications, but can be disabled if required. It does, however, require that file history information is available in order to be useful. Refer to the backup application's administration guide for more information.

#### ZFS-NDMP Backup

The Oracle ZFS Storage Appliance has implemented a proprietary NDMP format which enables the backup application to store block-level data to tape. This format includes backing up and restoring both filesystem shares and LUNs. It has the potential of greatly improving backup performance as compared to traditional NDMP formats. It is discussed in depth in the "NDMP Formats" section of this document.



### Multi-Stream Concurrent Backups

The Oracle ZFS Storage Appliance allows for multiple backup requests to stream data to the same tape drive simultaneously. Not all backup applications support this feature.

### Incremental and Cumulative Level Backups

The Oracle ZFS Storage Appliance supports full backups (level 0) as well as incremental level (levels 1 – 9) backups. A full backup is taken the first time that the target is backed up. Thereafter, differential incremental levels can be specified to back up only the data that has changed since the last full or comparable level backup.

Traditionally, a full backup is conducted on a weekly basis (for instance, Sunday), and then level backups on subsequent days (for example, Monday level 1, Tuesday level 2, ..., Saturday level 6). This allows for shorter backup windows during the week, since only the changes are backed up. The drawback, however, is that a full restore of the data would require that all the level backups be restored sequentially. For example, using the provided level schedule, if a full restore is required for Thursday, backup levels 0, 1, 2, 3, 4, and 5 would need to be sequentially restored. Depending on the size of the protected volume, the restore process can be fairly long.

Another method is to run cumulative incremental level backups, where the same level is specified after a full backup. A cumulative incremental level backup captures the data that has changed since the last full backup. So, performing another level 1 backup on Tuesday would capture all the changes included in Monday's level 1 backup. A level 1 backup on Wednesday would capture all of the changes included in Monday's and Tuesday's backups. This method increases the amount of data stored, but decreases the time required to restore a full volume, since only the full backup and then the specific day's level 1 backup would need to be restored.

### Token Based Backups

In level backups, the administrator only has ten backup levels (0 – 9) to utilize. Once a level 9 backup has been requested, all subsequent level 9 backups will become cumulative incremental backups. Thus, a full backup (or a level lower than 9) is required before any additional differential incremental backups can be resumed.

Token-based backups use a date-time stamped "token" to extend beyond the level backup limitation. After a full backup is performed, the backup application can then request more than one billion differential incremental backups (up to  $2^{30} - 1$ ) over time. Whereas this is not practical from the sense of level backups, it does enable some backup applications' infinite incremental offerings. These applications are able to synthesize restorable full backups from the incremental NDMP backups. Once the backup application has a full backup of the volume, then it only needs incremental data updates in the future. Refer to the individual backup application's administration guide to configure its specific implementation.

### Backup of Replication Targets

The Oracle ZFS Storage Appliance includes the licensable feature to replicate a volume from one Oracle ZFS Storage Appliance to another, either continuously or according to a

schedule. All NDMP backup formats are supported from the replication source. Only ZFS-NDMP can be used to perform backups of the data at the replication target. For more information specific to this backup method, search for “NDMP Replica Backup” in the online help accessible through the Oracle ZFS Storage Appliance BUI's top right corner.

#### **NDMP File Service Extensions**

NDMP File Service Extensions is an enhancement to the NDMP version 4 specification. These extensions allow an enabled backup application to implement NDMP disk targets on the Oracle ZFS Storage Appliance. The disk targets use the same commands and protocols as the NDMP tape devices, and therefore allows for the easy transfer of backup images between disk and tape, providing for a disk-to-disk-to-tape backup solution with NDMP. The benefits to this implementation include:

- Virtually instantaneous backup response since the disk is always available and does not need to wait for tape mount, dismount, or positioning operations.
- Potentially faster backup speeds limited only by the network bandwidth.
- Potentially better tape drive utilization since the data from disk to tape is transferred internally on the Oracle ZFS Storage Appliance. Data is consolidated and staged in larger sizes to keep the tape drive running at full speed.

Refer to the individual backup application's administration guide to determine if this feature is available.

#### **Data Encryption**

There are several paths available to encrypt data among the Oracle ZFS Storage Appliance, the backup application, and the tape drives. Many backup applications have encryption mechanisms built into their software that encrypt the data based on keys that the backup application manages. Tape drives can either be encrypted through the backup software or through a separate key management system such as Oracle Key Manager (OKM). The Oracle ZFS Storage Appliance can manage its own encryption keys on the appliance or also integrate with OKM.

When using the appliance-level encrypted volumes on the Oracle ZFS Storage Appliance, the data is unencrypted before it is backed up via NDMP. Encryption capable tape drives can be used to store the backup on encrypted media. However, the data restored back to the Oracle ZFS Storage Appliance will not be encrypted unless it is restored to an encrypted project or share.

## NDMP Formats

The Oracle ZFS Storage Appliance supports tar/dump NDMP and ZFS-NDMP formats. The features supported by the NDMP formats are displayed in Table 1.

TABLE 1. NDMP FORMAT FEATURE LIST		
FEATURE	TAR/DUMP NDMP	ZFS-NDMP
NDMP Version	Supports Version 3 and Version 4	Supports Version 4 Only
Volume Backup	Supports Filesystem/Share Backup Only	Supports Filesystem/Share and LUN Backup
Full Volume Restore	Supports Filesystem/Share Restore Only	Supports Filesystem/Share and LUN Restore
Individual File Restore	Supported	Not Supported. Entire volume must be restored.
Incremental Level Backup	Supported	Supported
Project Level Backup	Not Supported	Supported
Token Based Backup	Supported	Supported
Replication Source Backup	Supported	Supported
Replication Target Backup	Not Supported	Supported
NDMP File Service Extensions	Supported	Supported

### Tar and Dump NDMP Format Use

The NDMP service on the Oracle ZFS Storage Appliance accepts both dump and tar backup requests from the backup application. These are the most common NDMP backup type requests and are supported by all of the backup applications currently qualified with the Oracle ZFS Storage Appliance.

#### Specifying a Backup Path

Only shares and filesystems can be backed up using the tar/dump NDMP format, using its Oracle ZFS Storage Appliance mounted name. This is a path which will start with `/export`. This mountpoint can be viewed by using the Oracle ZFS Storage Appliance BUI to navigate to BUI -> Shares -> Filesystems. Each share name will be listed with its respective mountpoint on the right, as shown in Figure 1.

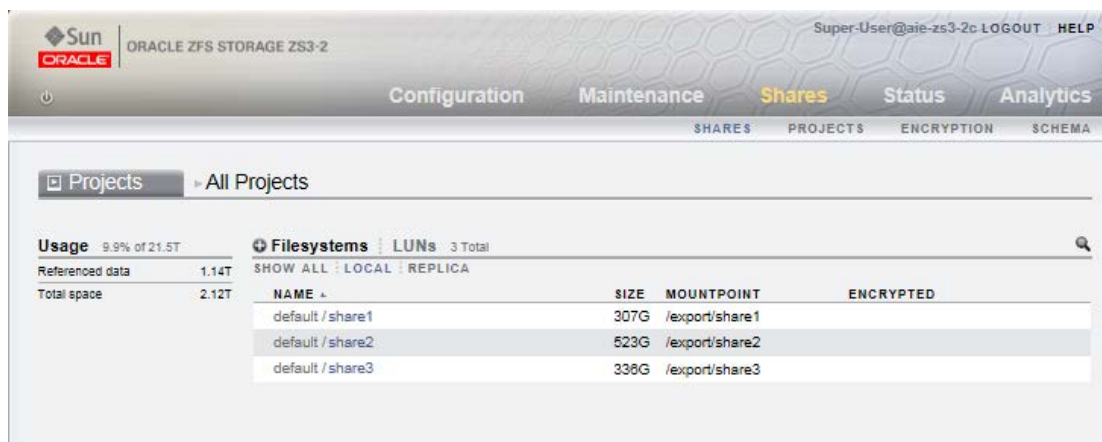


Figure 1. Example Oracle ZFS Storage Appliance share listing

This mountpoint is used as the path for the backup definition in the backup application. Subdirectories of these mountpoints (for example, `/export/share2/subdir3`) can also be specified.

#### Backing Up a Snapshot

Before a backup begins, the NDMP service on the Oracle ZFS Storage Appliance takes a snapshot of the share being backed up to ensure a consistent, point-in-time image. This allows for the continued use of the share while the backup is in progress. After the backup is completed, this system-generated snapshot is deleted. These actions are performed regardless of conducting a full-, level-, or token-based backup.

The Oracle ZFS Storage Appliance allows the administrator to create manually generated or system periodic snapshots. In order to back up from one of these snapshots, the snapshot path must be appended to the mountpoint. To determine the snapshots available for a share, navigate to BUI -> Shares -> Shares -> Filesystems -> <double-click share> -> Snapshots. The screen capture in Figure 2 shows an example of the snapshots available for mountpoint `/export/share3`.

The screenshot shows the Oracle ZFS Storage Appliance web interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' section is active, showing 'default > share3'. The interface displays usage statistics (1.7% of 19.7T), properties (zfs/snapshot visibility: Hidden, Scheduled Snapshot Label), and static properties (Creation date: 2015-3-11, Compression: 1.00x, Case sensitivity: Mixed, Reject non UTF-8: yes, Normalization: None, Encryption: Off). A table lists snapshots and schedules:

NAME	CREATION	UNIQUE	TOTAL	CLONES
Quarterly	2015-5-5 10:19:55	0	336G	-
Monthly	2015-5-5 10:19:46	0	336G	-

Figure 2. Example Oracle ZFS Storage Appliance share snapshot listing

The “Monthly” snapshot, for example, would have a backup specification of:

```
/export/share1/.zfs/snapshot/Monthly
```

The `/.zfs/snapshot/` directive notifies the Oracle ZFS Storage Appliance that the dataset source is a snapshot. The snapshot name is case sensitive.

Since the manual snapshot is already a read-only point-in-time entity, an additional snapshot is not created for the backup.

A backup of a manual snapshot will result in a full (level 0) backup.

### Restoring Backup Data

Using `tar/dump NDMP`, the Oracle ZFS Storage Appliance can restore data back onto the original share or onto a separately designated share. A full backup can be restored to an existing share or to one that has not yet been created. With an existing share, the restore will overlay the current contents. With a non-existent share, the restore will create a new share on the Oracle ZFS Storage Appliance which will have a unique name that starts with “`ndmp-`”, and will be mounted with the user-specified mountpoint. Figure 3 shows an example of this type of restore with the mountpoint of `/export/share-new`.

The screenshot shows the Oracle ZFS Storage Appliance web interface. At the top, there is a navigation bar with tabs for Configuration, Maintenance, Shares, Status, and Analytics. Below this, there are sub-tabs for SHARES, PROJECTS, ENCRYPTION, and SCHEMA. The main content area is titled 'Projects' and shows a selected project named 'testproject'. Underneath, there are sub-tabs for Shares, General, Protocols, Access, Snapshots, and Replication. The 'Shares' sub-tab is active, showing a table of filesystems and LUNs. The table has columns for NAME, SIZE, MOUNTPOINT, and ENCRYPTED. The data rows are:

NAME	SIZE	MOUNTPOINT	ENCRYPTED
ndmp-6275dd21-c626-62c8-f2fe-aec2a0ae9e02	832M	/export/share-new	
share1	832M	/export/test/share1	
share2	825M	/export/test/share2	
share3	832M	/export/test/share3	

Figure 3. Example Tar/Dump NDMP restore to a non-existent share

The restored share can be given a new name by clicking and holding (about 1 second) on the current name. This should highlight the current name and allow it to be modified.

This non-existent share restore will put the created share in the same project from which the original share came. If the project no longer exists, it will be recreated during the restore.

All incremental restores must be applied to an existing share. Otherwise, the restore will fail.

A restore of a manual snapshot must be performed on an existing share. Otherwise, the restore will fail.

### ZFS-NDMP Format Use

The ZFS-NDMP format differs from tar/dump NDMP by transferring block-level information instead of file-level information. This enables ZFS-NDMP to back up Fibre Channel or iSCSI LUNs as well as shares and filesystems. Depending on the structure of the filesystem, ZFS-NDMP can dramatically improve backup performance on shares which include millions of small files. A limitation of ZFS-NDMP, however, is that the entire backup must be restored in order to recover a single file. Since it is block level, file histories are not provided back to the backup application. Reserve space (the size of largest volume backed up) must be maintained on the Oracle ZFS Storage Appliance in order to accommodate a restore from which the desired file(s) can then be copied.

### Passing NDMP Parameters

In order for a backup application to use ZFS-NDMP, it needs to provide a mechanism to allow the backup administrator to send NDMP parameters to the Oracle ZFS Storage Appliance NDMP service during backups and restores. When using tar/dump NDMP, these parameters typically do not need to be modified by the backup administrator. For ZFS-NDMP, however, at a minimum, the TYPE parameter must be set to "zfs". Each backup application provides a slightly different interface for the backup administrator to use to change these parameters. Refer to the backup application's Administration Guide for application-specific information.

The backup applications which provide this interface and have been qualified with Oracle ZFS Storage Appliance ZFS-NDMP are:

- Oracle Secure Backup v10.3.0.2 and above
- Symantec NetBackup v7.0 and above
- EMC NetWorker v7.6 and above

Configuration examples for each of these backup applications with ZFS-NDMP are provided in the Appendices of this document.

The following listed parameters are considered by the Oracle ZFS Storage Appliance NDMP service. Many of these parameters are transparent to the typical backup administrator. However, understanding how these parameters work can lead to more advanced backup and restore configurations.

- `TYPE=(zfs | dump | tar)`

Used for backup and restore operations.

The `TYPE` parameter alerts the NDMP service to which NDMP format is to be used during the backup or restore operation. This is typically set by the backup application as `tar` or `dump`. For ZFS-NDMP, however, `TYPE` needs to be set to `zfs`. There is no default value for this parameter.

- `ZFS_FORCE=(n | y)`

Used only for restore operations.

The `ZFS_FORCE` parameter determines whether the Oracle ZFS Storage Appliance should perform a rollback to the last restored snapshot before restoring the specified incremental backup. If data has changed since the last restored snapshot and no rollback is performed (that is, `ZFS_FORCE=n`), the attempt to restore the next level of incremental backup will fail. If `ZFS_FORCE` is set to `y` in this situation, a rollback will be performed and all data changes since the last restored snapshot will be lost. However, this rollback will allow the incremental restore to succeed. If `ZFS_FORCE` is unset, the default value is `n`.

Note that the backup application setting of `ZFS_FORCE` is only acted upon when `TYPE=zfs` and when the Oracle ZFS Storage Appliance NDMP service setting “ZFS rollback before restore” is configured to “On DMA request.” The `ZFS_FORCE` parameter is ignored when “ZFS rollback before restore” is configured to “Always” or “Never.” Refer to the “Configuring the NDMP Service” section for more details.

This parameter is only acted on when `TYPE=zfs`.

- `UPDATE=(y | n)`

Used only for backup operations.

The `UPDATE` parameter controls how previous ZFS-NDMP level snapshots are handled after a successful ZFS-NDMP backup is completed. In general terms, setting the parameter to `y` will replace the previous snapshot with the one generated for use with the current ZFS-NDMP backup. Setting the parameter to `no` will delete the generated

snapshot after the current ZFS-NDMP backup completes, leaving the previous same level snapshot in place. For detailed handling of snapshots using this parameter, refer to “ZFS-NDMP System Snapshot Concepts” later in this section.

With tar/dump NDMP, this parameter is used to determine whether the Oracle ZFS Storage Appliance NDMP `dumpdates` file should be updated at the conclusion of the backup.

If unset, the default value is `y`.

- `DMP_NAME=(level | token | <application specified>)`

Used for backup and restore operations.

The `DMP_NAME` parameter allows the backup application to provide a unique backup set name so that several series of NDMP level backups can be maintained on the same dataset. The `DMP_NAME` value can only contain alphanumeric characters, the underscore character (`_`), and the hyphen character (`-`). Additionally, it is limited to 31 characters. More details on this parameter are provided in “ZFS-NDMP System Snapshot Concepts” later in this section.

If unset, the default value is `level` for incremental operations and `token` for token-based operations. Both `level` and `token` are reserved names, and should not be user specified in the `DMP_NAME` parameter. This parameter is available to all NDMP types.

- `LEVEL=(0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9)`

Used for backup and restore operations.

The `LEVEL` parameter indicates whether the Oracle ZFS Storage Appliance should perform a full backup, an incremental backup, or a cumulative incremental backup. `L0` indicates a full backup, and `L1 – L9` indicate incremental backups, where only changes since the last level’s backup are backed up. Two incremental backups at the same level (such as `L1`), will generate a cumulative incremental backup where the latest incremental backup will contain all the changes since the last lower level backup (such as `L0`).

If unset, the default value is `0`. This parameter is used by all NDMP formats, and is typically specified by the backup application based on the type of backup or restore requested. It is included here primarily for reference.

- `ZFS_MODE=(recursive | dataset)`

Used for backup and restore operations.

The `ZFS_MODE` parameter selects whether to backup just the basic dataset (`dataset`), or the dataset and its descendants (`recursive`). Either setting can be used to back up a share, LUN, or a snapshot. A `recursive` setting is required when backing up a project in order to back up all of the shares and LUNs within the project. A `dataset` setting for a project backup will succeed, but will have no contents.

If unset, the default value is `recursive`. This parameter is only acted on when `TYPE=zfs`.



- HIST=(y | n)

Used only for backup operations.

The HIST parameter indicates to the NDMP service whether the backup application is expecting file-level history data on the information requested for backup. This parameter is ignored when TYPE=zfs.

- DIRECT=(y | n)

Used for backup and restore operations.

The DIRECT parameter indicates to the NDMP service whether to provide or use DAR backup position data. This parameter is ignored when TYPE=zfs.

### ZFS-NDMP System Snapshot Concepts

Every time a tar/dump NDMP backup is processed, the Oracle ZFS Storage Appliance creates a snapshot of the share and processes the backup from that snapshot. Thus, it allows the share to remain in production while the backup continues. Once the backup is completed, the snapshot is destroyed. ZFS-NDMP system snapshots are more complex.

Every time a ZFS-NDMP backup is processed (excluding backing up a manual snapshot), the Oracle ZFS Storage Appliance creates a snapshot in the form of:

```
.ndmp.$LEVEL.$DMP_NAME.$ZFS_MODE.<timestamp>
```

where the variables correspond with the associated NDMP parameters. For instance, the example:

```
.ndmp.1.daily.r.1430495073
```

would correspond with the NDMP variables being set to:

```
SET TYPE=zfs
SET LEVEL=1
SET DMP_NAME=daily
SET ZFS_MODE=recursive
```

ZFS\_MODE uses the abbreviation of recursive (“r”) and dataset (“d”).

The differential incremental snapshot in this example would have been based on an earlier level 0 (full) backup snapshot, such as:

```
.ndmp.0.daily.r.1430473328
```

Likewise, the next day’s differential incremental backup would generate a snapshot similar to:

```
.ndmp.2.daily.r.1430543221
```

If either the DMP\_NAME or ZFS\_MODE are modified, the system would generate a different series of snapshots that do not affect this series. This allows for multiple series of backup snapshots to be maintained on the volume. For example, the different series could be set up as quarterly, monthly, and weekly.

Due to their block structure, ZFS-NDMP incremental backups must be restored in the order they were taken. To restore the level 2 backup contents to a new location, first the level 0 backup must be restored, then the level 1, and finally the level 2.

By default, the system will overwrite the next “like” incremental backup snapshot when the new one is created. Thus, when the next full backup is requested, a new `.ndmp.0.daily.r.<time stamp>` will replace the previous one, unless `UPDATE=n`. If `UPDATE=n`, the new snapshot will be destroyed after the backup has completed, and the original backup snapshot will remain on the system.

### Specifying a Backup Path

In order to use ZFS-NDMP, the backup path must be specified as the Oracle ZFS Storage Appliance ZFS path to the project, filesystem or share, or LUN. This is different than the path specified for tar/dump NDMP backup. A method to determine the ZFS path of the share is to use the BUI to navigate to BUI -> Shares -> Shares -> Filesystems -> <double-click share> -> General. The ZFS path will be displayed under the share name. Figure 4 shows the screen capture of “share2”, where the ZFS path is `BigPool/local/testproject/share2`.



Figure 4. Determining the ZFS path of a share using the Oracle ZFS Storage Appliance BUI

With this example, the backup path in the backup application would be specified as:

```
/BigPool/local/testproject/share2
```

Even though the leading “/” is not displayed in the BUI, it is required for the backup path.

Similarly, the ZFS path of a LUN can be determined by navigating to BUI -> Shares -> Shares -> LUNs -> <double-click LUN> -> General. Figure 5 shows the screen capture of “testlun1”, where the ZFS path is `BigPool/local/testproject/testlun1`.

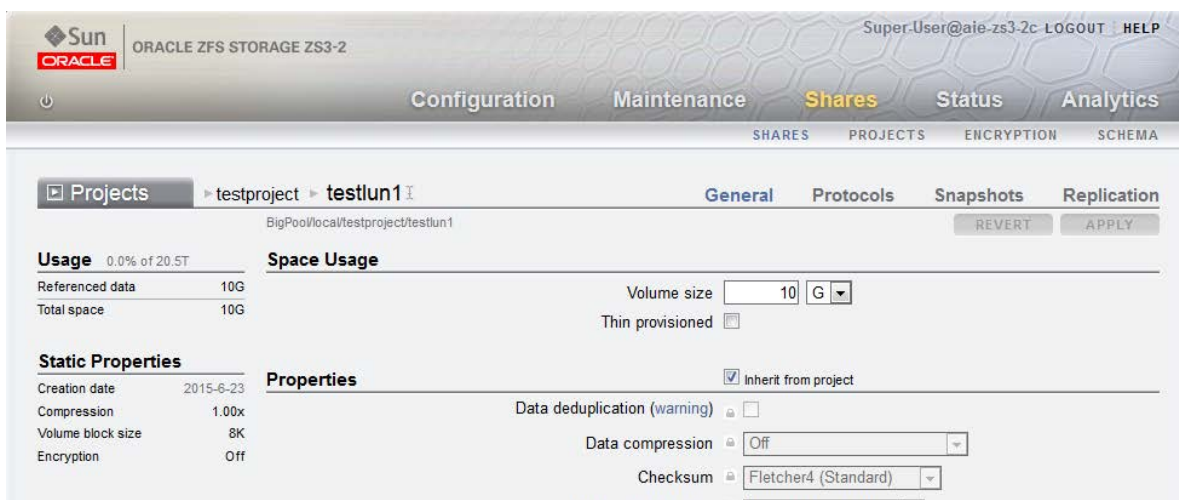


Figure 5. Determining the ZFS path of a LUN using the Oracle ZFS Storage Appliance BUI

With this example, the backup path in the backup application would be specified as:

```
/BigPool/local/testproject/testlun1
```

Again, the leading “/” is required.

With ZFS-NDMP, it is possible to back up the entire project with a single backup specification. Since the previous two examples are both in the same project (testproject), the backup specification could simply be:

```
/BigPool/local/testproject
```

which would back up all of the shares and LUNs within the project.

### Backing Up a Snapshot

Manual snapshots can also be backed up with ZFS-NDMP, but must also be specified using their ZFS path. A method to determine a snapshot’s path is to navigate to BUI -> Shares -> Shares -> Filesystems or LUNs -> <double-click share or LUN> -> Snapshots -> Snapshots. Figure 6 shows the screen capture of “share3”, where the ZFS path is BigPool/local/testproject/share3, and a manual snapshot named “EndOfMonth” was created.

Figure 6. Determining the ZFS path of a manual snapshot using the Oracle ZFS Storage Appliance BUI

The backup path specification in the backup application would be:

```
/BigPool/local/testproject/share3@EndOfMonth
```

The “@” directive notifies the Oracle ZFS Storage Appliance that the dataset source is a snapshot.

Since the manual snapshot is already a read-only point-in-time entity, an additional snapshot is not created for the backup.

The first backup of any manual snapshots of a volume will result in a full (level 0) backup. Subsequent manual snapshots can be used to perform incremental backups provided that the new manual snapshot has a unique name and that the previous level’s manual snapshot still exists on the volume.

### Backing Up a Replication Target

This feature is new in OS8.4, and is delivered as a deferred update. Please verify that the update is applied by navigating to BUI -> Maintenance -> System. If the update has not been applied, a Deferred Updates section will be displayed with “Support for ndmp-zfs replica backup” as one of the options, as shown in Figure 7.

Figure 7. Unapplied Deferred Updates listing

The feature must be applied in order for replica backups to work properly.

Backing up a replication source is the same as backing up any typical volume. Backing up a replication target, however, requires different handling. The first difference is the backup specification, which will now be in the form of `/Pool/<replication id>/Project/share`. The full volume reference can be located on the replication target host BUI by navigating to BUI -> Shares -> Shares -> Filesystems -> Replica -> <double-click share>, as shown in Figure 8, just above the “Space Usage” heading.



Figure 8. Determining the replication target backup specification

In the figure example, the backup specification for the volume would be:

```
/Pool1/nas-rr-7dcb38ee-108b-ca2c-98c1-9ae01179ff94/NDMP/nshare1
```

The backup specification for the project (assuming project-level replication is configured) would be:

```
/Pool1/nas-rr-7dcb38ee-108b-ca2c-98c1-9ae01179ff94/NDMP
```

In either case, the replication target hostname needs to be provided as the backup client.

Manual snapshots created on the original source volume or project can also be specified using the `@snapshot` nomenclature.

For information regarding setting up replication and replication concepts, please refer to the Oracle ZFS Storage Appliance online help.

### Restoring Backup Data

A ZFS-NDMP full (level 0) backup can only be restored to a path that does not already exist. So, if the backup administrator plans to restore back to the original volume, the original volume must first be deleted. It is more likely that the backup administrator will restore to an alternative location to recover the targeted files and copy them back to the original volume.

In the case of restoring a single share or LUN, that path must be specified within a project, such as:

```
/BigPool/local/testproject/recovery-share
```

where the newly recovered share would be named `recovery-share` and would be located in the project `testprojec`”.

In the case of restoring an entire project, the project itself cannot exist, but the shares are not specified. So, the recovery path for a project could simply be provided as:

```
/BigPool/local/recovery-project
```

where the newly recovered project would be created as “`recovery-project`”, and all of the shares within the backup will be restored within this project as their original names (that is, `/BigPool/local/recovery-project/share3`).

Share name conflicts will occur when restoring an entire project to the same Oracle ZFS Storage Appliance where the original project still exists. The newly recovered shares will attempt to mount using the same mountpoints as the original shares, but will fail. This can be observed by navigating to BUI -> Shares -> Projects -> <double-click restored project> -> Filesystems. On the far right of the filesystem listing, yellow circles appear, indicating an issue with the share (as displayed in Figure 9).



The screenshot shows the Oracle ZFS Storage Appliance BUI interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' section is active, and the 'Projects' tab is selected. The project name is 'recovery-pro...'. The 'Filesystems' section shows a table with the following data:

NAME	SIZE	MOUNTPOINT	ENCRYPTED
share1	832M	/export/test/share1	
share2	825M	/export/test/share2	
share3	832M	/export/test/share3	

Each row in the table has a yellow circle icon in the rightmost column, indicating a mount error for each share.

Figure 9. Oracle ZFS Storage Appliance BUI filesystem listing indicating mount errors

The original shares will remain mounted. This condition can be corrected by navigating to BUI -> Shares -> Projects -> <double-click restored project> -> General. Under the “Inherited Properties” section, the mountpoint for the entire project can be modified, as shown in Figure 10.

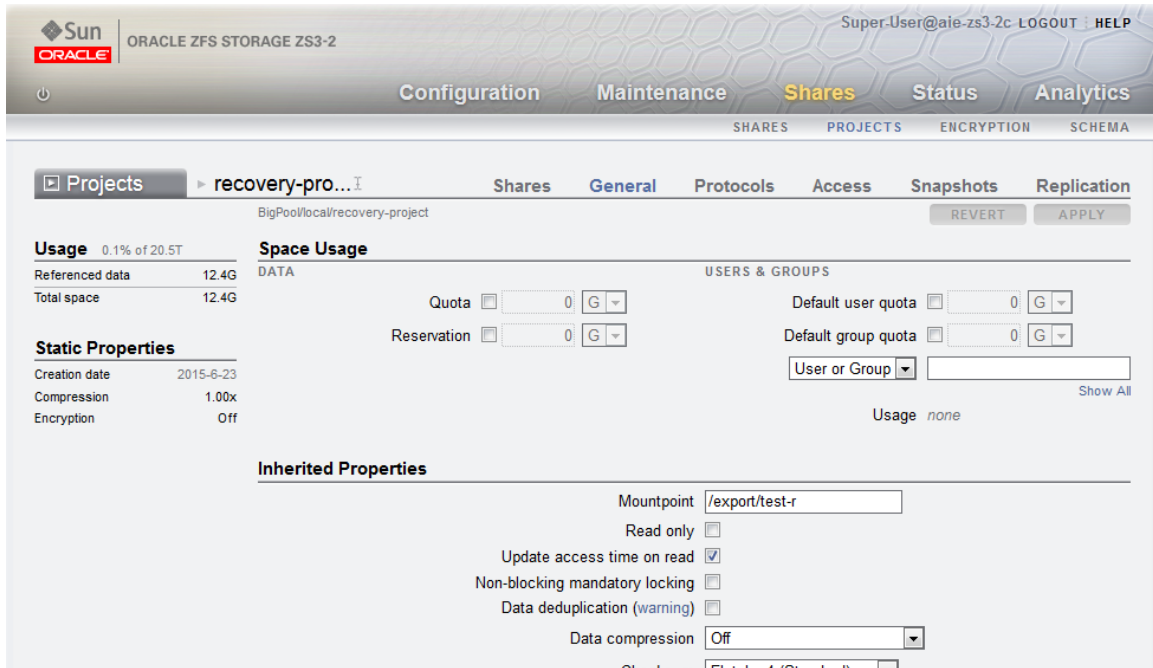


Figure 10. Changing project level share mountpoints on the Oracle ZFS Storage Appliance BUI

In this example, the mountpoint was modified from `/export/test` to `/export/test-r`. Once the field is applied, the system will remount the shares with the new mountpoint, as shown in Figure 11.



Figure 11. Results of applying a new mountpoint to the restored project

When restoring a LUN, the restored LUN will attempt to use the same GUID that was part of the original backup. If that GUID already exists (for instance, the original LUN is still in production), a new GUID will be assigned. It may take several minutes before it appears.

When restoring from a snapshot, the backup application may choose to provide the restore path with or without the snapshot name. If a snapshot name is provided in the restore path (for example, `/Pool/local/Project/myshare@snapshot1`), the Oracle ZFS Storage Appliance will simply strip off the snapshot name and attempt to restore to the designated path. This is particularly helpful with a backup application's "Restore to Original Location" option. The Oracle ZFS Storage Appliance NDMP service will alert the backup application that the `@snapshot` will be ignored. It is up to the backup application to log this message.



## Configuring NDMP on the Oracle ZFS Storage Appliance

Configuration of NDMP on the Oracle ZFS Storage Appliance is performed through the BUI.

### Configuring the Fibre Channel Interface for NDMP

The individual Fibre Channel ports on an Oracle ZFS Storage Appliance can be configured either as targets or initiators. Each port that will be used for NDMP devices must be set to initiator mode. The ports are set as targets by default, so a change will be required.

The Fibre Channel ports can be configured by navigating to BUI -> Configuration -> SAN -> Fibre Channel -> Ports, as shown in Figure 12.



Figure 12. Configuring Fibre Channel ports using the Oracle ZFS Storage Appliance BUI

**NOTE:** Changing the Fibre Channel port from “Target” to “Initiator” will force a system reboot.

The Fibre Channel port WWN is provided under the port designation. Use this WWN to configure WWN-based (alias) zoning on connected Fibre Channel switches.

### Configuring the NDMP Service

The NDMP service on the Oracle ZFS Storage Appliance is configured through the BUI, by navigating to BUI -> Configuration -> Services -> <double-click NDMP>, and is displayed in Figure 13.

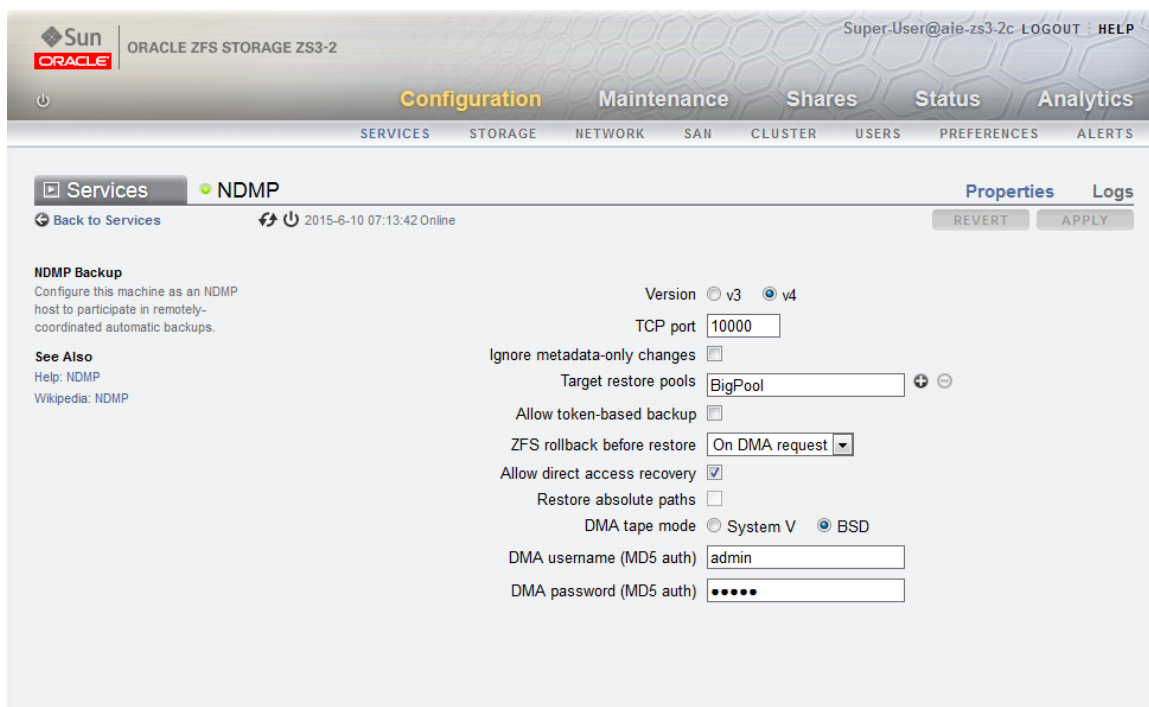


Figure 13. Oracle ZFS Storage Appliance NDMP service configuration screen with typical settings

Verify that the NDMP service is enabled. A green dot should be displayed just to the left of the “NDMP” title, and it should display “Online” to the right of the power icon just below it. If the NDMP service is not online, click the power icon to start it.

The Oracle ZFS Storage Appliance uses the term “DMA” when referring to the backup application. The term “DMA” stands for “Data Management Application,, which is a generic term directly from the NDMP specification.

The following selectable options are available:

- **Version:** This selects the NDMP version to use with the backup application. All of the supported backup applications have been tested using version 4 (v4).
- **TCP port:** The NDMP standard defaults to port 10000 for NDMP control operations. NDMP version 4 allows for the ability to specify the desired port value. This port value must match between the backup application and the Oracle ZFS Storage Appliance NDMP service. Otherwise, communication will fail.
- **Ignore metadata-only changes:** When checked, this selection directs the system to back up only the files in which content has changed, ignoring files for which only metadata, such as permissions or ownership, has changed. This option only applies to incremental tar/dump backup types, and is disabled by default.
- **Target restore pools:** When doing a full restore using tar/dump NDMP types, the system will re-create datasets if there is not already a share mounted at the target. Because the NDMP protocol specifies only the mount point, the system will, by default, choose a pool in

which to recreate any projects and shares. For a system with multiple pools, specify one or more pools to restore to. Multiple pools need only be specified in a cluster with active pools on each head; it is the user's responsibility to make sure that the list is consistent with any storage configuration changes (such as adding a new pool, or deleting a pool). If the pools are offline, or if none of the pools specified exist, then the system will select a default pool at random.

- **Allow token-based backup:** For ZFS-NDMP only. When checked, this setting notifies the backup application that the Oracle ZFS Storage Appliance is capable of performing token-based incremental backups as well as level-based incremental backups for ZFS-NDMP. Left unchecked (the default), the Oracle ZFS Storage Appliance notifies the backup application that only performs level-based incremental backups for ZFS-NDMP. This setting does not affect the behavior of tar/dump NDMP token-based backup, which is always enabled.
- **ZFS rollback before restore:** For ZFS-NDMP only. When restoring an incremental backup, this selection determines whether the system rolls back the target project and share to the snapshot used as the base for the incremental restore. If the project and shares are rolled back, then any changes made since that snapshot will be lost. When this setting is set to "On DMA request", then the NDMP service chooses the appropriate action based on the `ZFS_FORCE` parameter. When set to "Always", the NDMP service will automatically initiate the ZFS rollback regardless of the `ZFS_FORCE` parameter value. A "Never" setting will not perform the ZFS rollback under any circumstance. A restore failure will occur if the NDMP service needs to perform a ZFS rollback but cannot based on these settings. It is recommended that this property is set to "Never" or "On DMA request" so that no unexpected rollback is performed. It can be set to "Always" when needed to restore a particular ZFS-NDMP incremental backup when required, then set back after the restore has been completed.
- **Allow direct access recovery:** For tar/dump NDMP only. When checked, this parameter allows direct access recovery (DAR) data to be used to locate the targeted files on the tape media. Enabling this option can reduce the time it takes to recover a small number of files from many tapes. The backup must have already been DAR enabled in order for a DAR restore to be performed. This feature is enabled by default.
- **Restore absolute paths:** This setting is only valid for NDMP v3 (tar/dump NDMP). When checked, it specifies that the complete absolute path to that file is also restored (instead of just the file itself). This option is disabled by default.
- **DMA tape mode:** This feature is specific to the tape drives directly attached to the Oracle ZFS Storage Appliance and used by the NDMP service. The two possibilities (System V or BSD) primarily affect the way the tape is positioned after certain operations. The preference is backup application specific. Even though the default is set to System V, the majority of the backup applications use BSD. Unless otherwise directed, BSD should be selected. Changing the tape mode does affect the device file naming. With BSD, the device file will be in the format of `/dev/rmt/#bn` whereas the device file for the System V

mode will be in the format of `/dev/rmt/#n`. A backup application device rescan will be necessary after changing the mode.

- **DMA username (MD5-auth):** Along with the TCP port, the backup application also needs to know the username with which to authenticate communications. This is a unique username which does not correspond with any other local users or domain users. It is strictly used for NDMP access. The username specified here must also be configured in the backup application.
- **DMA password (MD5-auth):** This password corresponds with the username (above). It, too, must be specified within the backup application.

When a parameter is changed on the NDMP services screen, the change will need to be enacted by clicking the “Apply” button. Once applied, the NDMP service will be restarted with the new parameters.

NOTE: The NDMP service parameters are shared on a clustered Oracle ZFS Storage Appliance system. Changes made to the NDMP service parameters on one clustered head will also be reflected on the other clustered head.

### Attaching Fibre Channel Tape Drives and Tape Libraries

The Oracle ZFS Storage Appliance supports tape libraries and tape drives that are connected point-to-point with the Fibre Channel ports, or connected through a Fibre Channel switch. The Fibre Channel switch must be zoned properly before the system can detect the attached library and/or tape drives. Navigate to BUI -> Status -> NDMP to display the discovered device, as shown in Figure 14.

TYPE	PATH	VENDOR	MODEL	WWN	SERIAL
Robot	/dev/scsi/changer/c0t500104F000D14CFE0	STK	SL150	50:01:04:F0:00:D1:4C:FE	464970G+1333SY13
Tape drive	/dev/rmt/0bn	HP	Ultrium 6-SCSI	50:01:04:F0:00:D1:4D:01	HU1325W7NM
Tape drive	/dev/rmt/1bn	HP	Ultrium 6-SCSI	50:01:04:F0:00:D1:4C:FE	HU1325W8DA
Tape drive	/dev/rmt/2bn	HP	Ultrium 6-SCSI	50:01:04:F0:00:D1:4D:04	HU1325W8CB
Tape drive	/dev/rmt/3bn	HP	Ultrium 6-SCSI	50:01:04:F0:00:D1:4D:07	HU1325W8CD

Recent activity 0 Total

No active NDMP sessions.

Figure 14. Oracle ZFS Storage Appliance NDMP device display

All of the connected devices should now be configured.

The NDMP Device display will only show ten devices at a time. It may be necessary to navigate to the next screen(s) of devices.

If any expected devices are not displayed, try the following:

- Press the F5 key. The browser may simply need to be refreshed.
- Verify that all of the devices are powered on and cabled properly.
- If using a Fibre Channel switch, verify that all of the switch zoning is correct.
- Return to the NDMP Service screen (Figure 13) and restart the NDMP service. Once restarted, return to the NDMP Device screen (Figure 14) and press the F5 key.
- Reboot the system by clicking the power icon under the Oracle banner and choosing “Reboot”. Once the system has been rebooted, return to the NDMP Device screen.

## Monitoring Oracle ZFS Storage Appliance NDMP

The Oracle ZFS Storage Appliance provides several BUI screens from which NDMP activity can be monitored.

### Status Dashboard

The Oracle ZFS Storage Appliance system dashboard can be configured to include NDMP disk activity, as shown in Figure 15.

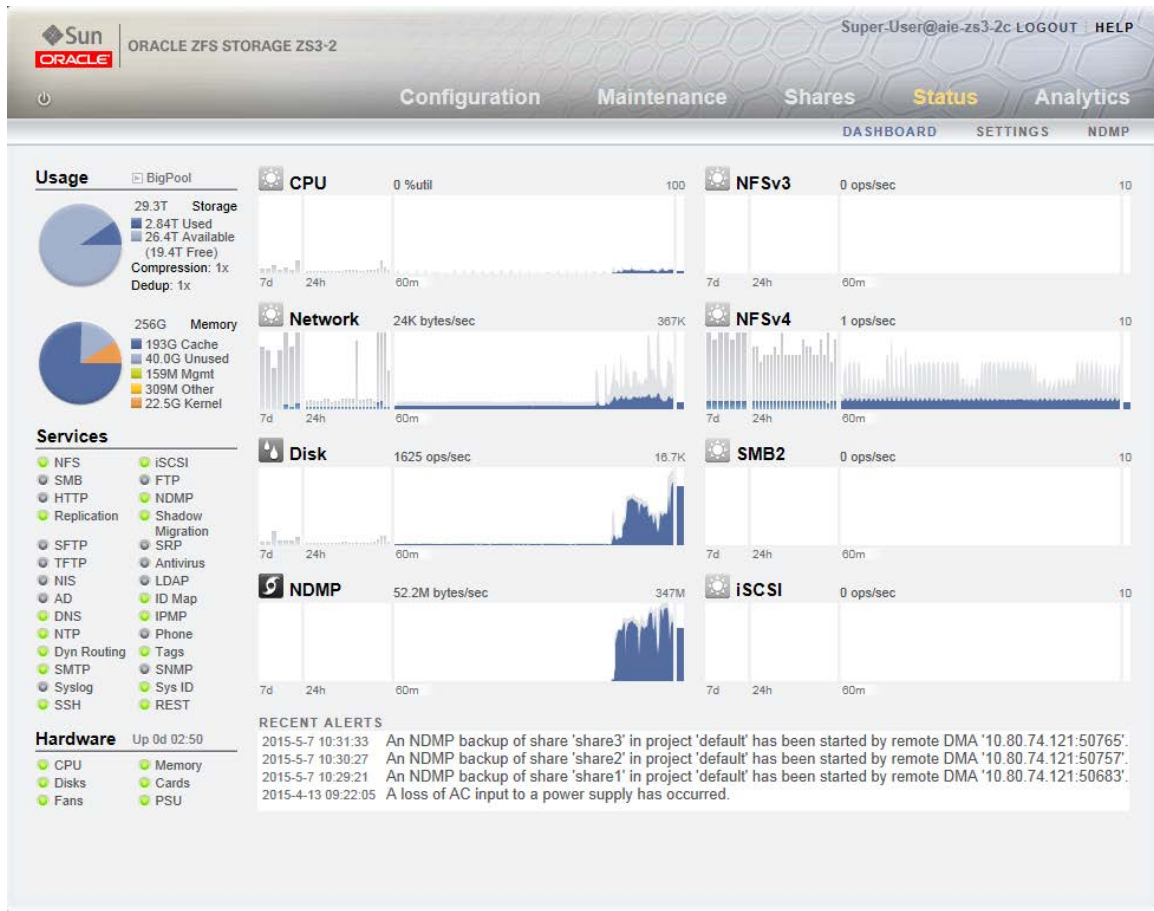


Figure 15. Oracle ZFS Storage Appliance dashboard configured with NDMP disk activity

By default, NDMP is not one of the eight categories displayed. To include NDMP, navigate to BUI -> Status -> Settings -> Layout -> <click drop-down menu> -> <select NDMP> -> <click Apply>. The configuration screen is displayed in Figure 16.

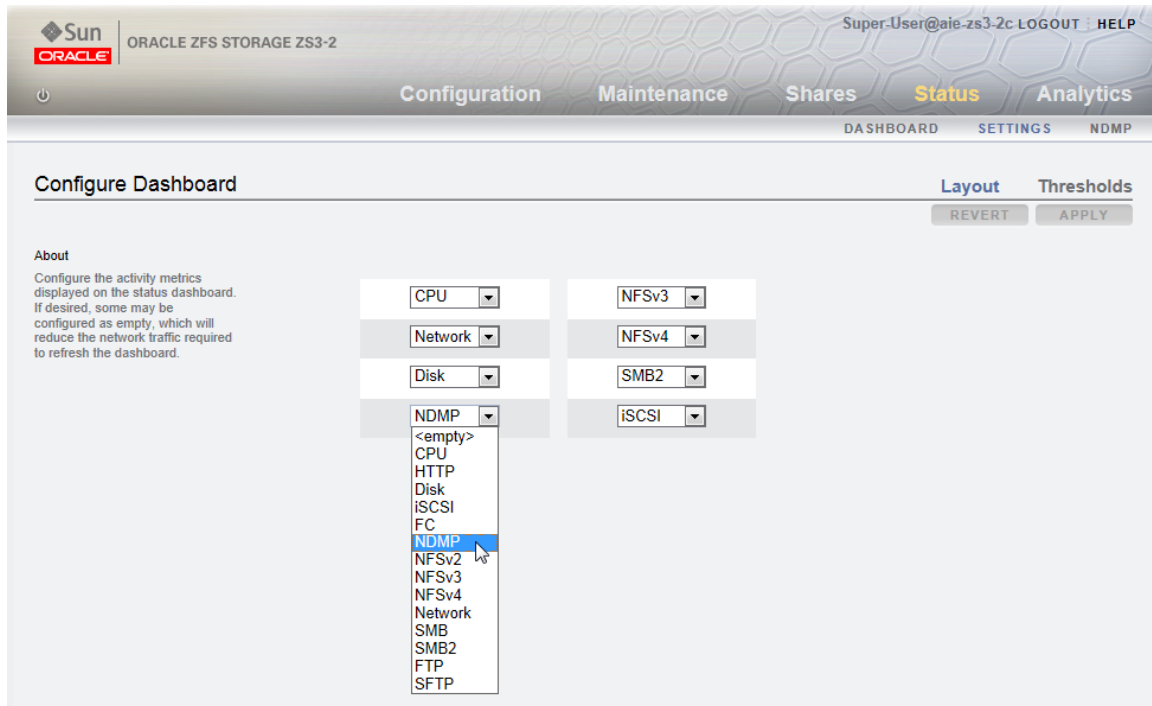


Figure 16. Including NDMP disk activity on the system dashboard

NDMP disk activity will now be included on the system dashboard.

Additionally, NDMP backups and restores will be posted to the “Recent Alerts” section at the bottom of the system dashboard, as demonstrated in Figure 15.

### Status NDMP

The NDMP devices screen also displays information regarding the active NDMP sessions in the “Recent activity” section, as displayed in Figure 17.

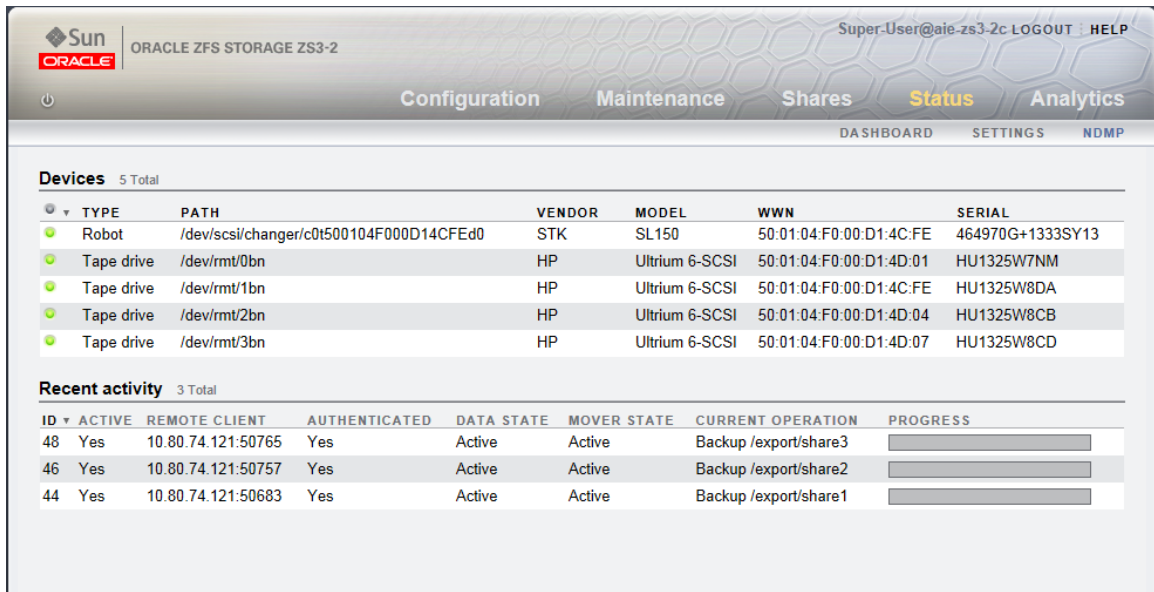


Figure 17. NDMP devices display including active NDMP session detail

### Analytics

The NDMP disk activity displayed on the system dashboard provides a high-level view on system metrics. It is possible to see additional metrics in the Analytics portion of the BUI. An example of an analytic of “Protocol: NDMP bytes per second broken down by device” is displayed in Figure 18.

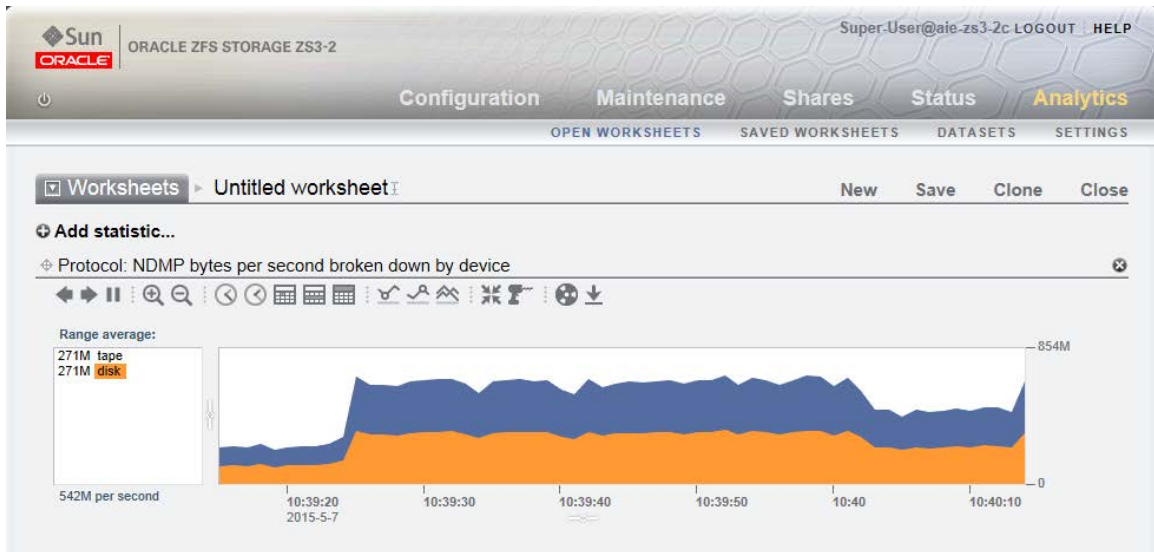


Figure 18. Example NDMP analytic displaying NDMP bytes per second broken down by device

In this example, NDMP local backups are in progress which utilize the disk (data source) and the tape (data target) without using any network resources. Network resources would be reported in three-way and remote NDMP configurations.



The NDMP analytics available can be selected by clicking the plus sign icon next to “Add statistic...” and are located under the “Protocol” section heading. When the default set of analytics are enabled on the system, the available analytics is limited to those displayed in Figure 19.

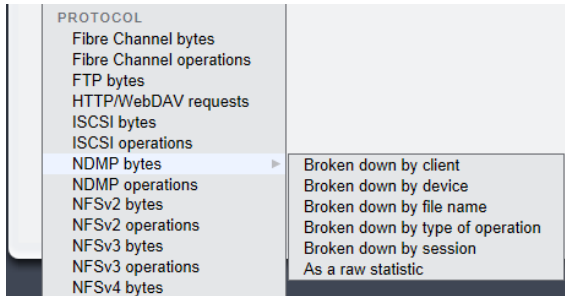


Figure 19. Basic NDMP analytics options

When advanced analytics are enabled on the system, additional options are available, as displayed in Figure 20.

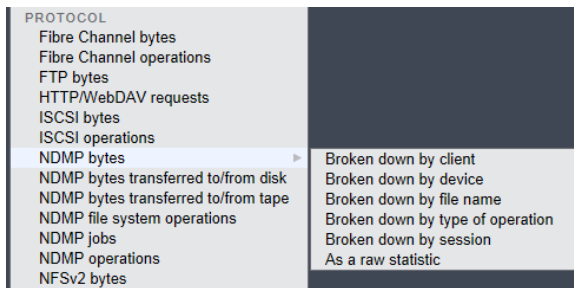


Figure 20. Advanced NDMP analytics options

Advanced analytics can be enabled by navigating to BUI -> Configuration -> Preferences, checking the “Make available advanced analytics statistics” box, and clicking “Apply”. The browser may need to be refreshed (F5 key) to reflect the changes on the Analytics screen.

For more information regarding the use of analytics, please refer to the “Analytics” section of the Oracle ZFS Storage Appliance Help (top right corner of the BUI).

## Supported Backup Applications

Oracle ZFS Storage Appliance engineering has worked with backup application vendors to validate that the products integrate seamlessly with each other. Not all backup applications interoperate with every feature of the Oracle ZFS Storage Appliance. Table 2 provides a list of the backup applications supported with the Oracle ZFS Storage Appliance, along with the features that the backup applications support.

TABLE 2. ORACLE ZFS STORAGE APPLIANCE FEATURES SUPPORTED BY BACKUP APPLICATIONS

BACKUP APPLICATION	FEATURES SUPPORTED
Oracle Secure Backup (OSB)	<ul style="list-style-type: none"> <li>• Local, Remote, Three-way, and Passthrough NDMP Topologies</li> <li>• Tar/Dump NDMP and ZFS-NDMP Formats</li> <li>• Selective File Restore (Tar/Dump NDMP Format Only)</li> <li>• Direct Access Recovery (DAR) (Tar/Dump NDMP Format Only)</li> <li>• Incremental/Level Backup and Restore</li> <li>• Token Based Backup and Restore</li> <li>• Backup of Replication Targets (ZFS-NDMP Only)</li> <li>• NDMP Disk Targets (Using NDMP File Service Extensions)</li> </ul>
Symantec NetBackup	<ul style="list-style-type: none"> <li>• Local, Remote, and Three-way NDMP Topologies</li> <li>• Tar/Dump NDMP and ZFS-NDMP Formats</li> <li>• Selective File Restore (Tar/Dump NDMP Format Only)</li> <li>• Direct Access Recovery (DAR) (Tar/Dump NDMP Format Only)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> <li>• Backup of Replication Targets (ZFS-NDMP Only)</li> </ul>
CommVault Simpana	<ul style="list-style-type: none"> <li>• Local, Remote, and 3 way NDMP Topologies</li> <li>• Tar/Dump NDMP Formats</li> <li>• Selective File Restore</li> <li>• Direct Access Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> </ul>
EMC NetWorker	<ul style="list-style-type: none"> <li>• Local, Remote, and Three-way NDMP Topologies</li> <li>• Tar/Dump NDMP and ZFS-NDMP Formats</li> <li>• Selective File Restore (Tar/Dump NDMP Format Only)</li> <li>• Direct Access Recovery (DAR) (Tar/Dump NDMP Format Only)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backup and Restore</li> <li>• Backup of Replication Targets (ZFS-NDMP Only)</li> <li>• Token Based Backup and Restore</li> </ul>

BACKUP APPLICATION	FEATURES SUPPORTED
Symantec Backup Exec	<ul style="list-style-type: none"> <li>• Local, Remote, and Three-way NDMP Topologies</li> <li>• Tar/Dump NDMP Formats</li> <li>• Selective File Restore</li> <li>• Direct Access Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> </ul>
Dell NetVault	<ul style="list-style-type: none"> <li>• Local, Remote, and Three-way NDMP Topologies</li> <li>• Tar/Dump NDMP Formats</li> <li>• Selective File Restore</li> <li>• Direct Access Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> </ul>
IBM Tivoli Storage Manager (TSM)	<ul style="list-style-type: none"> <li>• Local, Remote, and Three-way NDMP Topologies</li> <li>• Tar/Dump NDMP Formats</li> <li>• Selective File Restore</li> <li>• Direct Access Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> <li>• Token Based Backup and Restore</li> </ul>
HP Data Protector	<ul style="list-style-type: none"> <li>• Local NDMP Topology</li> <li>• Tar/Dump NDMP For</li> <li>• Selective File Restore</li> <li>• Direct Access Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Incremental/Level Backups</li> </ul>
EMC Avamar NDMP Accelerator	<ul style="list-style-type: none"> <li>• Remote NDMP Topology</li> <li>• Tar/Dump NDMP Formats</li> <li>• Selective File Restore</li> <li>• Direct Acces Recovery (DAR)</li> <li>• Multi-Stream Concurrent Backups</li> <li>• Token Based Backups</li> </ul>

## Best Practices

The following sections provide some of the best practices for configuring and using NDMP with the Oracle ZFS Storage Appliance.

### Oracle ZFS Storage Appliance Cluster Configuration

The network and Fibre Channel target mode port resources on the Oracle ZFS Storage Appliance transition from one cluster head to the other during a failover operation. The Fibre Channel initiator mode port resources, however, do not. Each cluster head manages its own attached devices separately from the other. Thus, all of the tape resources of the first cluster head are inaccessible when failed over to the second cluster head. Sharing the tape drives between the two clustered heads can help accessibility, but complicates device administration since the tape drives may become offline to the backup application during failed over conditions.

The recommended configuration for using NDMP with a clustered Oracle ZFS Storage Appliance includes:

- Configure the tape library robotics attached directly to the backup application media server. This enables continued robotics control regardless of which clustered Oracle ZFS Storage Appliance head is failed over, preventing the robotics from going offline. This can be achieved by direct library robotics connection, by connecting through a switch, or by separate control software, such as Oracle StorageTek ACSLS (refer to the backup application administration guide for support).
- Use a Fibre Channel switch to share the tape drives between the two clustered heads. Systematically attach one tape drive at a time to both systems so that they configure with the same device identifier. For example, connect the first tape drive to the zoned switch and verify that each cluster head detects the tape drive as `/dev/rmt/0bn`. Then, connect the second tape drive to the zoned switch and verify that each cluster head detects the tape drive as `/dev/rmt/1bn`. (If the tape drives are not configuring as the same device identifier on both heads, contact support to clear the tape drive configuration from the system, and start again.)
- Configure half of the tape drives in the backup application with the first cluster head as their connection point. Configure the other half of the tape drives in the backup application with the second cluster head as their connection point. This procedure varies based on the backup application. Refer to the backup application's administration guide for procedures.

An NDMP backup or restore may fail during a cluster head failover. It is dependent on which head fails over and which head is performing the operation. If the NDMP operation is running on the second clustered head and that head fails over its resources, then the NDMP operation will be interrupted and fail. The operation will need to be restarted. If the first clustered head fails over its resources and the second clustered head is performing the NDMP operation, though, the NDMP operations will continue uninterrupted on the second clustered head since its resources are unaffected.

## Deciding When to Use ZFS-NDMP

Assuming that the backup application supports ZFS-NDMP, the next step is to determine whether it is beneficial to use its features. Regarding filesystems, it is very dependent on the quantity and type of data that is stored, as well as the recovery objective. Generally, ZFS-NDMP will outperform tar/dump NDMP, since it does not incur the same protocol overhead that is inherent to tar/dump NDMP. The performance benefits of ZFS-NDMP greatly increase as the backed-up file sizes get smaller. It is recommended that each candidate environment runs its own tests to determine the performance improvement.

Another consideration is that tar/dump NDMP will be able to restore a single file while ZFS-NDMP will need the entire filesystem to be restored. The ZFS-NDMP restore duration of the entire filesystem should be slightly less than the backup duration since reads tend to be faster than writes to a tape drive. However, since histories are not provided, it cannot be searched through the backup application catalog. Both ZFS-NDMP and tar/dump NDMP can be used on the same Oracle ZFS Storage Appliance. The backup administrator may choose to run tar/dump NDMP on some filesystems while running ZFS-NDMP on others.

Fibre Channel and iSCSI LUNs can be backed up with ZFS-NDMP, but not at all with tar/dump NDMP. They can, however, be backed up through their mountpoints on the particular clients with the backup application outside of NDMP. One key point with using ZFS-NDMP to back up LUNs is to make sure that the LUN is in a consistent state when the system snapshot is created. How this is achieved will be dependent on the backup application (pre-scripts and post-scripts) and the application that is controlling the content of the LUN. In many cases, the controlling application may need to unmount the LUN, initiate the backup, and then remount it. The procedure will be application dependent.

## Backing Up ZFS-NDMP Projects or Datasets

ZFS-NDMP backs up block data and does not return file history data. Thus, the entire backup must be restored to retrieve its contents. The requirement, then, is that a reserve of disk space is needed in the event that a restore is necessary. That reserve disk space on the Oracle ZFS Storage Appliance must be at least as large as the largest backup.

Even though the entire project can be backed up by a single backup definition (that is, `/Pool/local/Project`), the backup administrator must decide whether it is appropriate for the particular backup environment since:

- The entire project must be restored.
- Reserve disk space for the project restore must be managed.
- The duration of the restore is dependent on the size of the entire project.

For most environments, it is recommended to manage the backup data definitions at the dataset level (that is, `/Pool/local/Project/share1`, `/Pool/local/Project/share2`, and so on) instead of the project level. The dataset definition level in the backup application provides that:

- Only the specific dataset needs to be restored.

- Reserve disk space requirements are reduced to only the largest dataset size.
- Restore time is shortened since only the one dataset needs to be recovered.

### Replication Target Backup Considerations

It can be very convenient to back up a volume at a replication target rather than at the source Oracle ZFS Storage Appliance. Advantages include offloading the backup processing from the primary system as well as possibly moving backup processing to a remote location altogether. The backups can be taken either when replication events are scheduled or set to continuous.

When a backup is requested on a replica target, the Oracle ZFS Storage Appliance NDMP service uses the last replication snapshot as the basis for the backup. You can view active replication snapshots by navigating to BUI -> Shares -> Shares or Project -> Replica -> <double-click selection> -> Snapshots. An example display of a project replication target is in Figure 21.

The screenshot shows the Oracle ZFS Storage Appliance BUI interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' tab is active, and the 'PROJECTS' sub-tab is selected. The main content area shows the 'Snapshots' page for a project replication target. It includes a 'Usage' section showing 10.1% of 3.20T used, a 'Properties' section with a dropdown for '.zfs/snapshot visibility' set to 'Hidden', and a 'Remote Replication' section. Below these is a 'Snapshots' table with 4 total entries, showing columns for NAME, CREATION, UNIQUE, and TOTAL. The table lists four snapshots, with the most recent one having a suffix of -104.

NAME	CREATION	UNIQUE	TOTAL
.rr-fd4554f8-c95b-6dbf-e346-e305619c4bf3-4f9	2015-6-12 13:16:17	0	329G
.rr-fd4554f8-c95b-6dbf-e346-e305619c4bf3-4f8	2015-6-12 13:16:05	10.2K	329G
ProjectSnapshot1	2015-6-12 10:04:01	10.2K	329G
.rr-fd4554f8-c95b-6dbf-e346-e305619c4bf3-104	2015-6-12 09:45:57	10.2K	329G

Figure 21. Example replication target snapshot display

In this example, the snapshot on the bottom with suffix -104 was used to perform a level 0 ZFS-NDMP backup. This snapshot will be maintained until another level 0 backup is performed, at which point the latest snapshot (in this case with suffix -4f9), would be replaced as its basis, and the original snapshot would be released. The maintained snapshot is considered to be in persistent hold. If the subsequent backup were an incremental backup, then the original level 0 backup based snapshot and the latest snapshot used for the incremental backup would both be put into persistent hold. Incremental backups cannot continue after a replication reversal or a sever action. A new level 0 backup must be created to start a new incremental series after such an event.

Persistently held snapshots can be manually deleted through the BUI screen shown in this previous figure. The manual deletion triggers a precursory confirmation dialog window, with a warning about the effect to future NDMP backups if the deletion is committed.

If the backup administrator will only be performing level 0 backups, then it is recommended that the `UPDATE=n` NDMP parameter is provided with the backup request. In this case, the Oracle ZFS Storage Appliance NDMP service will not put a persistent hold on the replication snapshot used for the backup. It simplifies management and potentially reduces held disk space.

## Conclusion

The Oracle ZFS Storage Appliance implementation of the NDMP protocol provides a powerful tool for backing up and restoring system volumes. In addition to backing up filesystems, it also enables the ability to back up and restore filesystems and LUNs at the block level with the ZFS-NDMP enhancement. Lastly, the recent implementation of NDMP File Service Extensions paves the way for faster and more flexible configuration topologies in the future. The Oracle ZFS Storage Appliance NDMP implementation has been qualified with industry-leading backup applications to offer a complete solution for data protection needs.



## Appendix A: Oracle Secure Backup ZFS-NDMP Configuration

OSB sets NDMP parameters at the host definition level. It is unique in that it enables the setting of backup NDMP parameters as well as restore NDMP parameters.

The following sections outline the basic configuration steps needed to configure OSB NDMP operations in a local NDMP topology. It is assumed that the Oracle ZFS Storage Appliance has already been configured for NDMP operations. All of the commands in this section use the OSB CLI `obtool`, which presents itself as `ob>`. For more information about `obtool` and its commands and parameters, refer to the Oracle Secure Backup Administration Guide.

All of the following commands are run using OSB 12.1. Other OSB versions may vary slightly, but the general procedures are the same.

### Adding an NDMP Host

The `mkhost` command is used to add and configure the Oracle ZFS Storage Appliance to OSB as follows:

```
ob> mkhost -a ndmp -o -r mediaserver,client -i 192.168.10.10 -u admin -B zfs -q
aie-zfssa
Password: *****
Password (again): *****
```

where `-a` indicates access type, `-o` indicates “bring online”, `-r` specifies roles, `-i` specifies network ip, `-u` specifies userid, `-B` specifies NDMP type (`zfs` for ZFS-NDMP or `dump` for tar/dump NDMP), `-q` indicates to prompt for password, and `aie-zfssa` is the name of the host to be added.

The host configuration details can be reviewed using the `lshost` command.

```
ob> lshost -l aie-zfssa
aie-zfssa:
  Access mode:          NDMP
  IP names:             192.168.10.10
  NDMP port:           (default)
  NDMP user name:      admin
  NDMP password:       (set)
  NDMP backup type:    zfs
  NDMP protocol version: (default)
  NDMP auth type:      (default)
  Algorithm:           aes192
  Encryption policy:   allowed (Encryption supported only with encrypting
drives)
  Rekey frequency:     1 month (system default)
  Key type:            transparent
  In service:          yes
  Roles:               mediaserver,client
  Trusted host:        yes
  UUID:                fe9be794-9f7b-1032-98af-8c917f7b128a
ob>
```

### Setting NDMP Parameters

OSB sets its NDMP variables at the client level, but is able to set different variables for backup and restore. Under most situations, additional NDMP parameters do no need to be

set for OSB. OSB already sets the `TYPE` parameter and the `LEVEL` parameter with each backup and restore. In the cases where the parameters do need to be explicitly set, the `chhost` command is used. The `-w` argument sets write parameters and the `-Y` argument sets the read parameters. So, the `DMP_NAME` parameter can be set and displayed as follows:

```
ob> chhost -W DMP_NAME=monthly -Y DMP_NAME=monthly aie-zfssa
ob> lshost -l aie-zfssa
aie-zfssa:
  Access mode:          NDMP
  IP names:            aie-zfssa
  NDMP port:           (default)
  NDMP user name:      admin
  NDMP password:       (set)
  NDMP backup type:    zfs
  NDMP protocol version: (default)
  NDMP auth type:      (default)
  Algorithm:           aes192
  Encryption policy:   allowed (Encryption supported only with encrypting
drives)
  Rekey frequency:     1 month (system default)
  Key type:            transparent
  In service:          yes
  Roles:               mediaserver,client
  Trusted host:        yes
  UUID:               fe9be794-9f7b-1032-98af-8c917f7b128a
  Backup environment:  DMP_NAME=monthly
  Restore environment: DMP_NAME=monthly
ob>
```

The parameters can be removed by using the `-x` argument for the write environment and `-z` for the restore environment as follows:

```
ob> chhost -x DMP_NAME -z DMP_NAME aie-zfssa
ob> lshost -l aie-zfssa
aie-zfssa:
  Access mode:          NDMP
  IP names:            aie-zfssa
  NDMP port:           (default)
  NDMP user name:      admin
  NDMP password:       (set)
  NDMP backup type:    zfs
  NDMP protocol version: (default)
  NDMP auth type:      (default)
  Algorithm:           aes192
  Encryption policy:   allowed (Encryption supported only with encrypting
drives)
  Rekey frequency:     1 month (system default)
  Key type:            transparent
  In service:          yes
  Roles:               mediaserver,client
  Trusted host:        yes
  UUID:               fe9be794-9f7b-1032-98af-8c917f7b128a
ob>
```

## Configuring Tape Storage Devices

The `discoverdev` command is used to recognize the available devices and configure them as follows:

```
ob> discoverdev -h aie-zfssa -v -c
```

```

26d4ce32-dd94-1032-82b2-91f6ee891ecd:
  Host:          aie-zfssa
  Device type:   Library
  Model:         SL500
  Serial number: 557040100352
  World Wide Name: WWN[50:01:04:F0:00:8C:C9:47]
  Device name:   aie-zfssa_lib_1
  Existing device: No
  Attachment new:
    Host:          aie-zfssa
    Raw device:    /dev/scsi/changer/c0t500104F0008CC947d0
26def5a6-dd94-1032-82b2-91f6ee891ecd:
  Host:          aie-zfssa
  Device type:   Tape
  Model:         Ultrium 4-SCSI
  Serial number: HU194278P7
  World Wide Name: WWN[50:01:04:F0:00:8C:C9:6F]
  Device name:   aie-zfssa_tape_1
  Existing device: No
  Attachment new:
    Host:          aie-zfssa
    Raw device:    /dev/rmt/2bn
26df023a-dd94-1032-82b2-91f6ee891ecd:
  Host:          aie-zfssa
  Device type:   Tape
  Model:         Ultrium 4-SCSI
  Serial number: HU194278P3
  World Wide Name: WWN[50:01:04:F0:00:8C:C9:72]
  Device name:   aie-zfssa_tape_2
  Existing device: No
  Attachment new:
    Host:          aie-zfssa
    Raw device:    /dev/rmt/1bn
26df0eba-dd94-1032-82b2-91f6ee891ecd:
  Host:          aie-zfssa
  Device type:   Tape
  Model:         Ultrium 4-SCSI
  Serial number: HU194278P2
  World Wide Name: WWN[50:01:04:F0:00:8C:C9:75]
  Device name:   aie-zfssa_tape_3
  Existing device: No
  Attachment new:
    Host:          aie-zfssa
    Raw device:    /dev/rmt/0bn

```

Checking each library to associate discovered drive(s) with DTE...

```
*** th0__warning: number of storage elements default (42) differs from current (84)
```

```
*** th0__warning: number of import export elements default (5) differs from current (10)
```

```
*** th0__warning: number of data transfer elements default (2) differs from current (3)
```

```
Assigning DTE 1 in library aie-zfssa_lib_1 for drive aie-zfssa_tape_1 with serial number: HU194278P7
```

```
Assigning DTE 2 in library aie-zfssa_lib_1 for drive aie-zfssa_tape_2 with serial number: HU194278P3
```

```
Assigning DTE 3 in library aie-zfssa_lib_1 for drive aie-zfssa_tape_3 with serial number: HU194278P2
```

```
ob>
```

where `-h` specifies the appliance host name, `-v` indicates verbose output, and `-c` indicates that the discovered devices should be configured to OSB. The “`th0__warning`” messages can be ignored.

Review the device configuration with the `lsdev` command as follows:

```
ob> lsdev
library          aie-zfssa_lib_1  not in service
  drive 1        aie-zfssa_tape_1 not in service
  drive 2        aie-zfssa_tape_2 not in service
  drive 3        aie-zfssa_tape_3 not in service
ob>
```

You can rename the devices with the `rendev` command as follows:

```
ob> rendev --nq aie-zfssa_lib_1 s1500
ob> rendev --nq aie-zfssa_tape_1 lto4-1
ob> rendev --nq aie-zfssa_tape_2 lto4-2
ob> rendev --nq aie-zfssa_tape_3 lto4-3
ob> lsdev
library          s1500          not in service
  drive 1        lto4-1          not in service
  drive 2        lto4-2          not in service
  drive 3        lto4-3          not in service
ob>
```

where `-nq` indicates not to prompt, then the original device name, and finally the new device name.

Using the `chdev` command will bring the devices online, as follows:

```
ob> chdev -o s1500 lto4-1 lto4-2 lto4-3
ob> lsdev
library          s1500          in service
  drive 1        lto4-1          in service
  drive 2        lto4-2          in service
  drive 3        lto4-3          in service
ob>
```

where `-o` indicates to bring the device online/in service, and then the new device names follow.

The library volumes are identified with the `inventory` command. Then use the `lsvol` command to view them, as follows:

```
ob> inventory -L s1500
ob> lsvol -L s1500
Inventory of library s1500:
  in  4:          barcode RAS678
  in  5:          barcode RAS677
  in  6:          barcode RAS680
  in  7:          barcode RAS674
ob>
```

where `-L` specifies the library name in both commands.

Make the library volumes available for use using the `labelvol` command as follows:

```
ob> labelvol -D lto4-1 -f 4-5
ob> lsvol -L s1500
Inventory of library s1500:
```

```

    in    4:          barcode RAS678, 818384896 kb remaining, mediainfo hw
encryptable
    in    5:          barcode RAS677, 818384896 kb remaining, mediainfo hw
encryptable
    in    6:          barcode RAS680
    in    7:          barcode RAS674
ob>

```

where `-D` indicates the name of the tape drive to use, `-f` specifies to force overwrite any previous contents, and 4-5 indicates the library element slot numbers to start and end with (4 through 5).

### Configuring Disk Storage Devices

OSB provides a feature to set up an Oracle ZFS Storage Appliance share as a disk pool, which can be used alongside or in place of tape. After the NDMP host is added to OSB and the share is created on the Oracle ZFS Storage Appliance, the `mkdev` command is used to create the disk pool as follows:

```

ob> mkdev -t disk -z -a aie-zfssa:/export/osbdiskpool zfssa-disk
ob>

```

where `-t` indicates the type as disk, `-z` indicates to initialize the disk, `-a` provides the network path to share, and the name of the device will be `zfssa-disk`.

The newly configured disk device will display as follows:

```

ob> lsdev
library          sl500           in service
  drive 1        lto4-1         in service
  drive 2        lto4-2         in service
  drive 3        lto4-3         in service
disk pool        zfssa-disk     in service

ob> lsdev -l zfssa-disk
zfssa-disk:
  Device type:      disk pool
  In service:       yes
  Debug mode:       no
  Capacity:         (not set)
  Consumption:      0
  Free space goal:  (system default)
  Concurrent jobs:  25
  Blocking factor:  (default)
  Max blocking factor: (default)
  UUID:             bb5341b8-dfd0-1032-99ef-9a3c0fb95891
  Attachment 1:
    Host:            aie-zfssa
    Directory:       /export/osbdiskpool
ob>

```

The disk pool device can now be used as any tape device.

### Creating a Backup Dataset

Create an OSB Backup Dataset with the `mkds` command, which will open a vi editor and a predefined template. Only two lines are necessary within the flat file: the `include host` directive and the `include path` directive. Two Dataset examples (ZFS-NDMP and tar/dump NDMP) are:

```
ob> mkds aie-zfssa-zfs

include host aie-zfssa
include path /Pool/local/Project/myshare
```

```
ob> mkds aie-zfssa-dump

include host aie-zfssa
include path /export/myshare
```

## Performing a Backup

Use the `backup` command to initiate an on-demand backup as follows:

```
ob> backup -D aie-zfssa-zfs -l 0 -r lto4-2 --go
Info: backup request 1 (dataset aie-zfssa-zfs) submitted; job id is admin/248.
ob>
```

where `-D` specifies the Dataset to use, `-l` specifies the backup level, `-r` (optional) specifies which device to use (which can also be a disk pool), and `-go` indicates that the backup command is complete.

Use the `lsjob` and the `catxcr` commands to monitor the backup progress.

## Performing a Restore

Use the `restore` command to initiate a restore operation as follows (Note: Use the `set host` directive first to set the host's reference catalog):

```
ob> set host aie-zfssa
ob> restore /Pool/local/Project/myshare -a /Pool/local/Project/myshare-new -d
lto4-3 --go
Info: 1 catalog restore request item submitted; job id is admin/253.
ob>
```

where the first path specifies the share that was backed up, `-a` specifies the alternate restore share, `-d` (optional) specifies which device to use, and `-go` indicates that the `restore` command is complete.

Use the `lsjob` and the `catxcr` commands to monitor the restore progress.

## Managing Instances

OSB enables the backup administrator to manage the backup images on disk and tape. Images which have been backed up to a disk pool can have additional copies on tape, or even migrate from disk to tape to more allocated disk space for the next backup cycle. Use the `lsbkcu` command to view available backup images, copy or migrate images using the `cpinstance` command, and view individual copy locations using the `lsinstance` command. The command interaction is demonstrated in the following command sequence:

```
ob> lsbkup -y disk
      Backup Image Name      Client      Type      Created      Size
aie-host1-20150518-135702    aie-host1    FS    2015/05/18.07:57    1.2
GB
aie-host1-20150518-152647    aie-host1    FS    2015/05/18.09:26    1.2
GB
ob> cpinstance -r lto4-2 -m aie-host1-20150518-135702
```

Info: copy instance for aie-host1-20150518-135702.1 submitted; job id is admin/264.

ob> **cpinstance -r lto4-1 aie-host1-20150518-152647**

Info: copy instance for aie-host1-20150518-152647.1 submitted; job id is admin/263.

ob> **lsbkup -y disk**

Backup Image Name	Client	Type	Created	Size
aie-host1-20150518-152647	aie-host1	FS	2015/05/18.09:26	1.2 GB

ob> **lsbkup -y tape**

Backup Image Name	Client	Type	Created	Size
aie-host1-20150518-135702	aie-host1	FS	2015/05/18.07:57	1.2 GB
aie-host1-20150518-152647	aie-host1	FS	2015/05/18.09:26	1.2 GB

ob> **lsinstance**

Instance Name	Created	Container(s)
aie-host1-20150518-152647.1	2015/05/18.09:26	zfssa-disk
aie-host1-20150518-152647.2	2015/05/18.09:37	VOL000006
aie-host1-20150518-135702.2	2015/05/18.09:37	VOL000007

ob>

The `lsbkup` command with `(-y)` disk filter lists the backup images that are on disk. The first `cpinstance` command migrates `(-m)` the first image to `(-r)` tape drive `lto4-2`. The second `cpinstance` command simply copies (no `-m` option) the second image to `(-r)` tape drive `lto4-1`. After the migrate and copy are completed, the `lsbkup` command with the `(-y)` disk filter now shows that only the second image is still on disk, whereas the `lsbkup` command with the `(-y)` tape filter shows that the two images are resident on tape. The `lsinstance` command displays all the individual copies and which containers they are located in. The `.1` and `.2` at the end of the image name refer to the copy number of the image.

## Appendix B: Symantec NetBackup ZFS-NDMP Configuration

Symantec NetBackup allows the backup administrator to add NDMP parameters to the backup specifications. Those same NDMP parameters will be used by the restore.

The following sections outline the basic configuration steps needed to configure Symantec NetBackup for NDMP operations in a local NDMP topology. It is assumed that the Oracle ZFS Storage Appliance has already been configured for NDMP operations. All actions will be performed using the NetBackup Administrative Console of Symantec NetBackup 7.6. Other versions will be similar, but may have some differences.

Further information can be found in the Symantec NetBackup Administrator's Guide and the Symantec NetBackup for NDMP Administrator's Guide.

### Specifying an NDMP Host

The following steps are needed to configure the Oracle ZFS Storage Appliance as an NDMP host to Symantec NetBackup:

1. Open the NetBackup Administration Console and expand Media and Device Management, and Credentials.
2. Right-click NDMP Hosts and click New.
3. Enter the Fully Qualified Domain Name (FQDN) or IP Address for the Oracle ZFS Storage Appliance, and click OK.




Figure 22. Entering the NDMP host name in a dialog window of the Symantec NetBackup Administration console

4. Select the radio button next to "Use the following credentials" and provide the NDMP user name and password that was configured in the Oracle ZFS Storage Appliance NDMP service. Click OK.



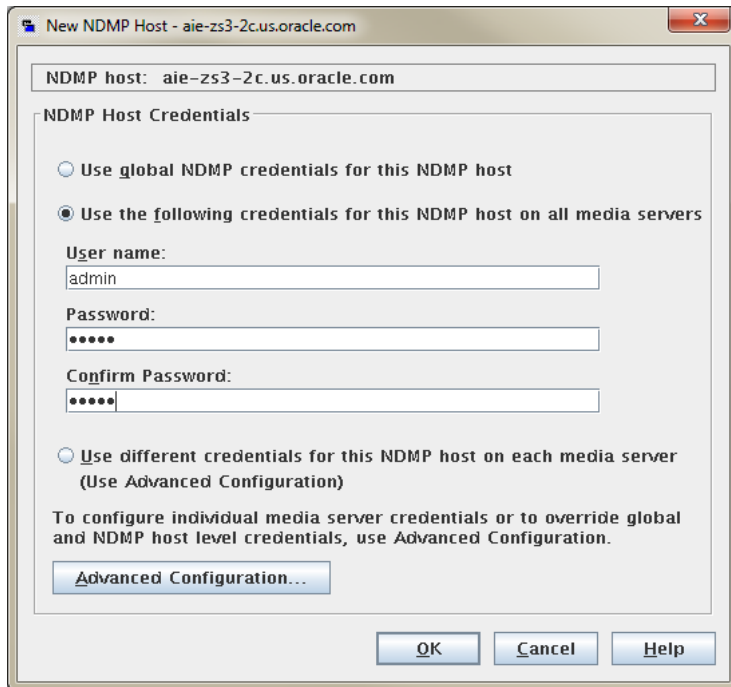


Figure 23. Setting up NDMP host credentials

## Configuring Storage Devices

Perform the following steps to configure the Oracle ZFS Storage Appliance Fibre Channel attached tape devices to Symantec NetBackup:

1. Open the NetBackup Administration Console and select Media and Device Management.
2. In the right windowpane, select Configure Storage Devices.
3. Click Next in the Welcome window.
4. Select the Master Server and click the Change... button.



Figure 24. Selecting device hosts in the NetBackup Administration console's Device Configuration Wizard

5. Check the NDMP server box and click OK.




Figure 25. Designating the NDMP server as device type

6. Click **Next** on the Device Hosts window.
7. Check the box next to the Oracle ZFS Storage Appliance name and click **Next** on the NDMP Hosts window.

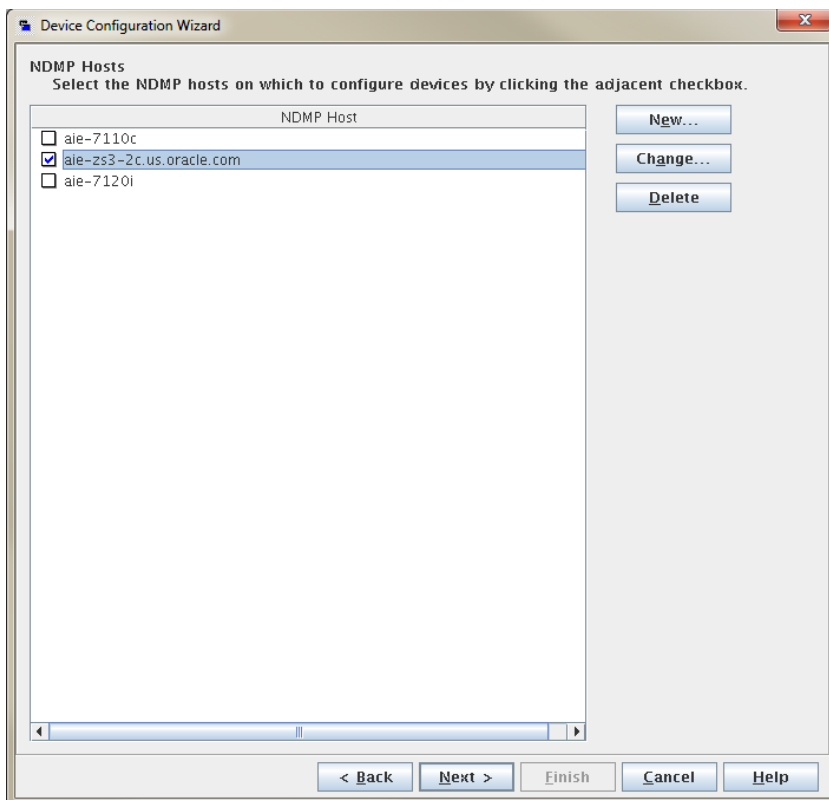


Figure 26. Designating the NDMP host on which to configure devices

8. Click Next on the Scanning Hosts window after it has discovered the tape devices.
9. Click Next on the SAN Clients window.
10. Review the results in the Backup Devices window, and click Next.
11. Select the devices that Symantec NetBackup should use, and click Next.

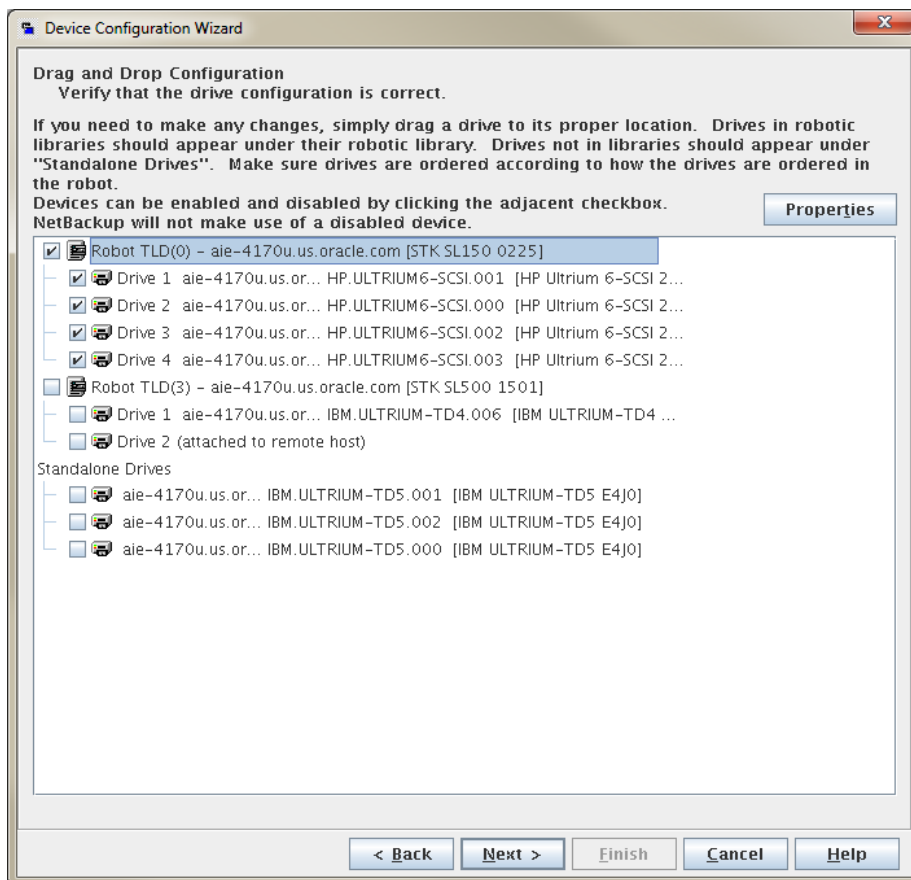


Figure 27. Selecting devices NetBackup will use

12. Click Next on the Updating Device Configuration window.
13. Click on the configured robot on the Configure Storage Units window, and click Properties towards the bottom left of the window.

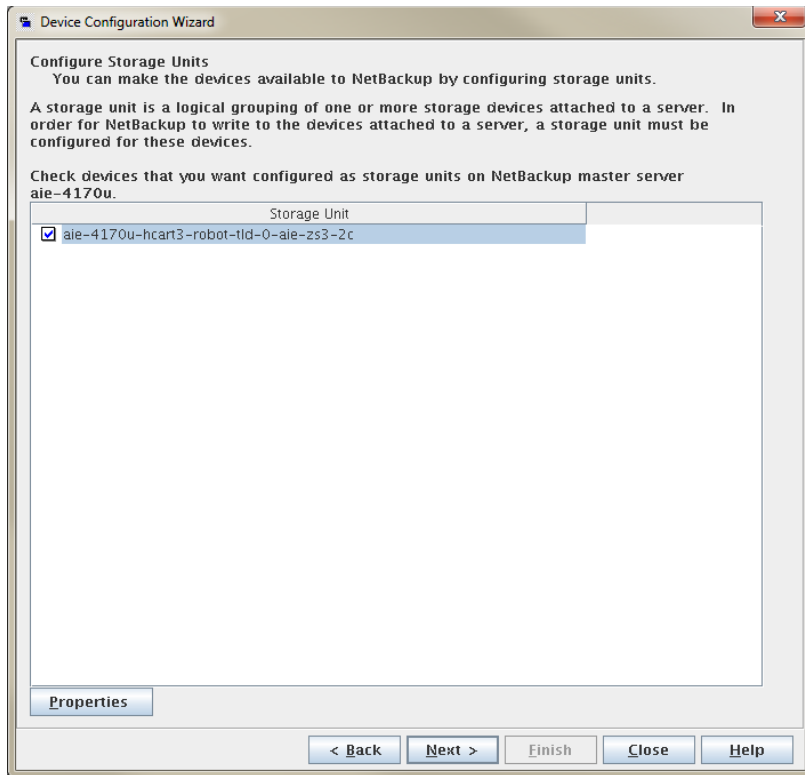


Figure 28. Designating storage units for configuration

14. Enter the preferred Storage unit name for the library and tape drives, and click OK.

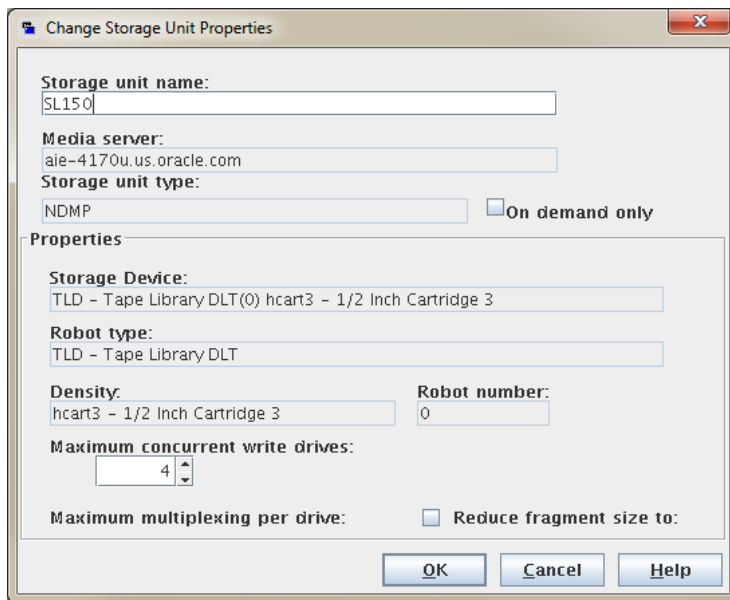


Figure 29. Entering properties for the selected storage unit

15. Click Next on the Configure Storage Units window.
16. Click Finish on the Finished window.

## Inventorying the Tape Library

Use the following steps to inventory the tape library attached to the Oracle ZFS Storage Appliance:

1. Open the NetBackup Administration Console, expand Media and Device Management, expand Devices, and click on Robots.
2. In the right windowpane, right click on the library and select Inventory Robot...

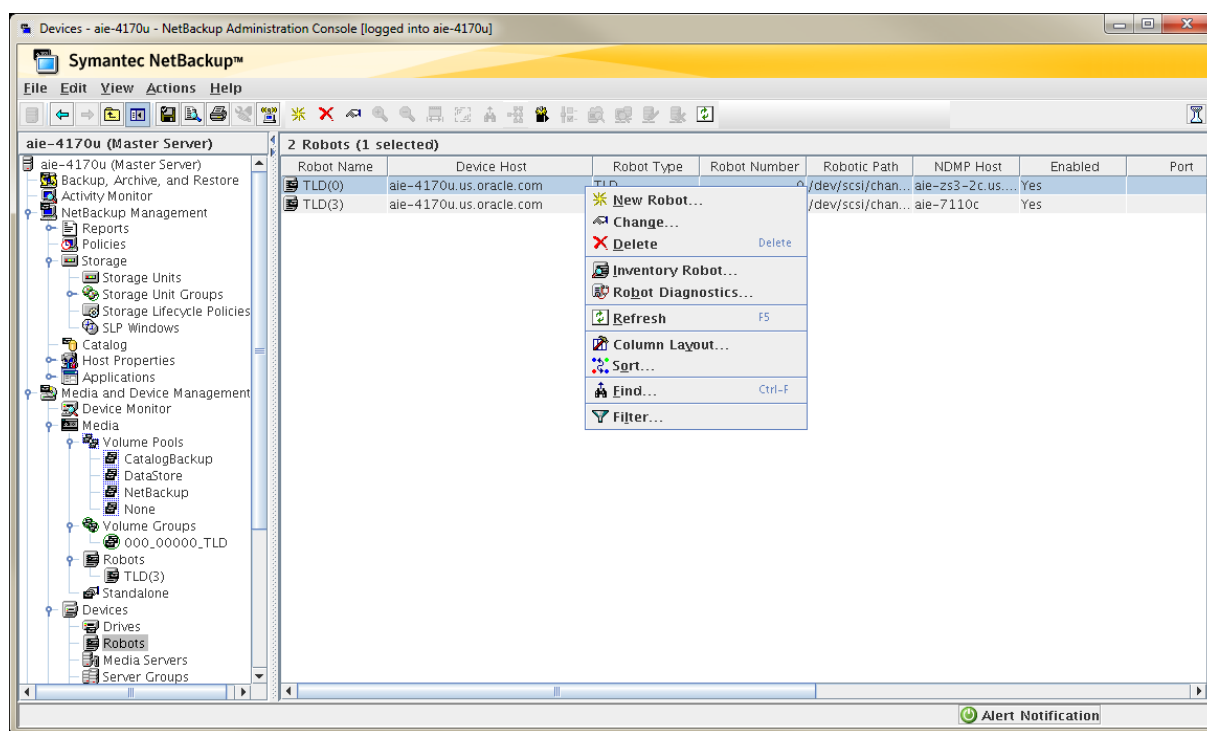


Figure 30. Taking an inventory of the tape library in the NetBackup Administration Console

3. Click the radio button next to Update volume configuration, and then click Start.

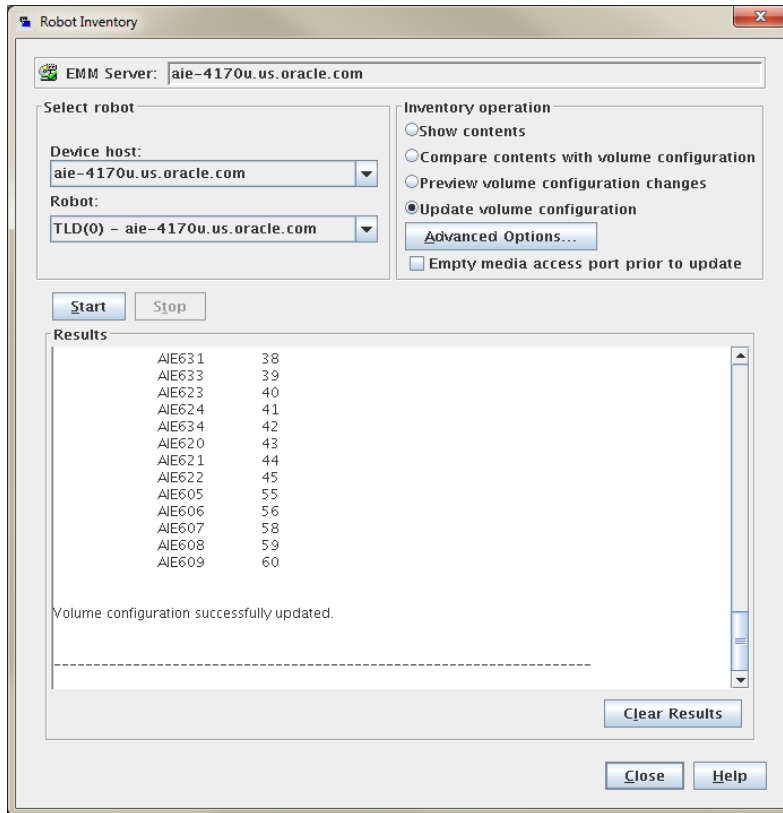


Figure 31. Updating and viewing the volumes in NetBackup

4. Click the **Close** button.

### Creating a Backup Policy

Perform the following steps to create a backup policy specific to NDMP in Symantec NetBackup:

1. Open the NetBackup Administration Console, expand NetBackup Management, right-click Policy, and click New Policy.
2. Enter a name for the new policy and leave the Policy Configuration Wizard unchecked. Click OK.

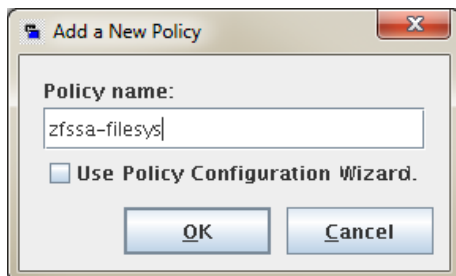


Figure 32. Using the NetBackup Policy Configuration Wizard to enter a new policy

3. For Policy Type, select NDMP.
4. For Policy Storage, select the storage unit name of the library and tape drives created earlier.
5. Do not select any options under “Snapshot Client and Replication Director”. Check the box next to ‘Allow multiple data streams’ if this policy should use more than one tape drive (when multiple volumes are specified).
6. Click Apply.

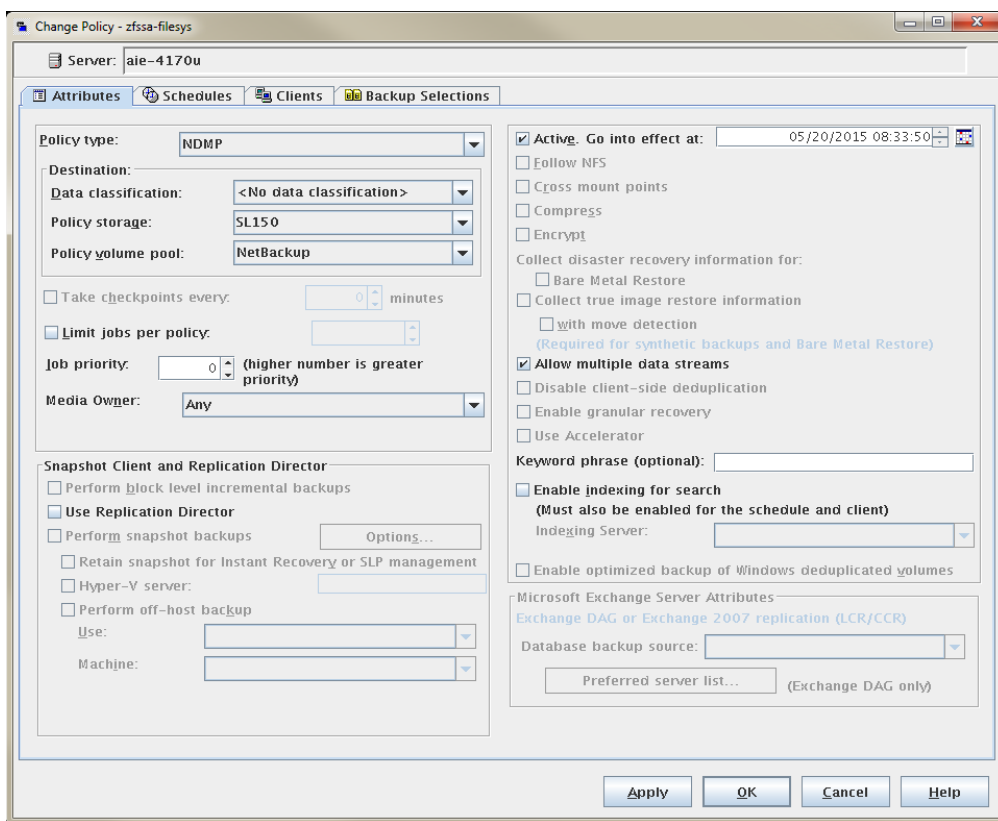


Figure 33. Setting attributes for the new policy

7. Click the Schedules tab, and create the desired backup schedule(s).
8. Click the Clients tab, click New, and enter the client name of the Oracle ZFS Storage Appliance and select 'NDMP, NDMP for the 'Hardware and operating system'. Click OK.

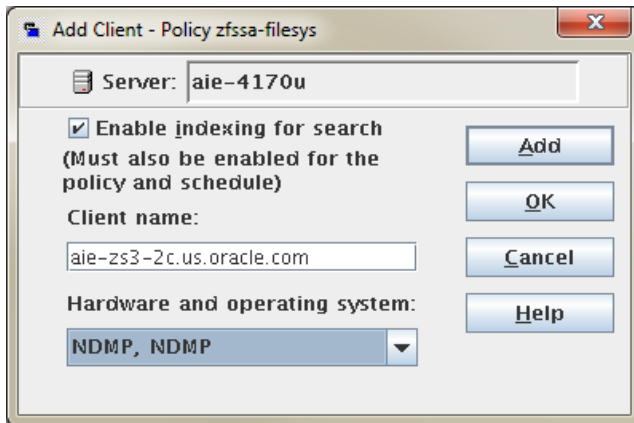


Figure 34. Assigning the client to the policy

9. Click the Backup Selections tab, and click New.

10. FOR TAR/DUMP NDMP SPECIFICATION:

Add the share pathnames in the `/export/sharename` format. Click OK.

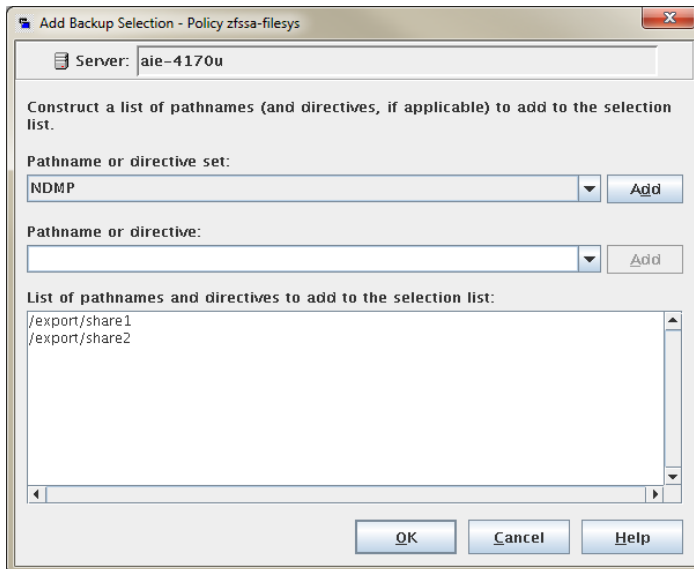


Figure 35. Adding the share pathname for a Tar/Dump NDMP specification

11. FOR ZFS-NDMP SPECIFICATION:

The first entry must be `SET TYPE=zfs` and the subsequent share definitions need to be in the `/Pool/local/Project/share` format. Click OK when done.



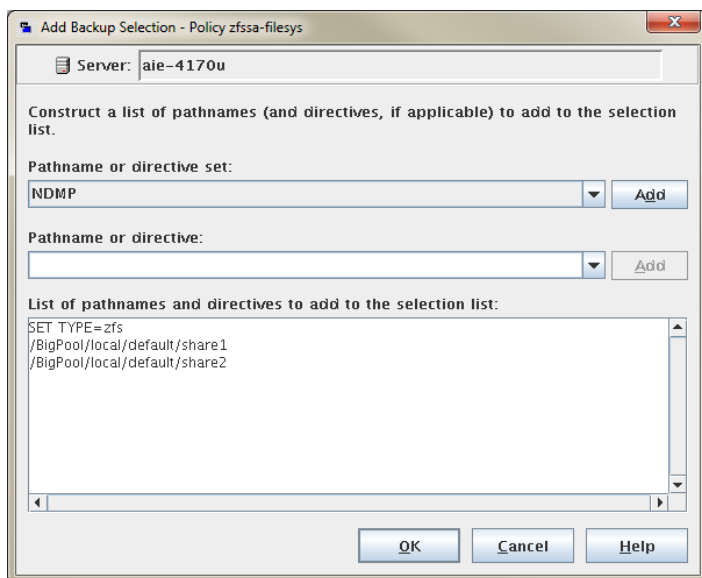


Figure 36. Defining type and path for the share for a ZFS-NDMP specification

Additional NDMP parameters can also be entered, but need to be defined before the share listing.

12. Click OK to complete the policy definition.

### Setting Maximum Jobs Per Client

If configured backup policies contain multiple shares and “Allow multiple data streams” is selected, then the global setting for “Maximum jobs per client” needs to be set on the master server.

1. Open the NetBackup Administration Console, expand NetBackup Management, expand Host Properties, and click Master Servers. Double-click the appropriate Master Server name.
2. Under Properties, click Global Attributes.
3. In the “Maximum jobs per client:” box, increase the number to the desired amount. To simultaneously use all tape drives, change the value to at least the number of tape drives connected to the Oracle ZFS Storage Appliance. Click OK to finish.

### Restoring an NDMP Backup

Restoring a backup policy specific to NDMP in Symantec NetBackup requires the following steps:

1. Open the NetBackup Administration Console and click on Backup, Archive, and Restore.

- On the far right side of the right windowpane, click the pointing finger icon to the right of “Policy type: Standard”.

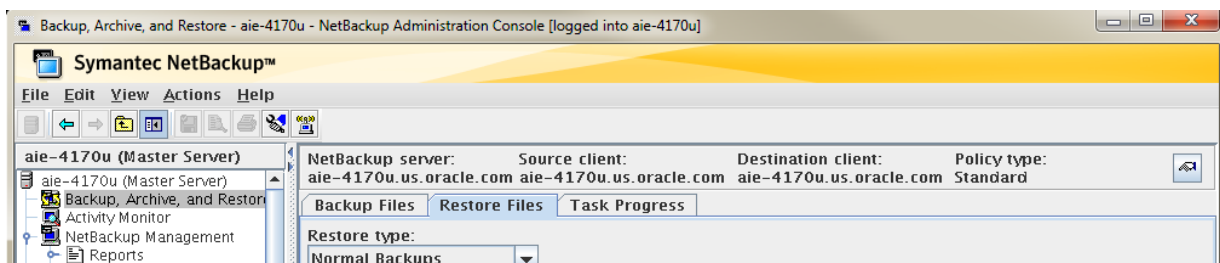


Figure 37. Editing the policy information in the NetBackup Administration Console

- Enter the hostname of the Oracle ZFS Storage Appliance in both the ‘Source client’ and ‘Destination client’ fields. Select NDMP as the “Policy type for restores:”, and click OK.

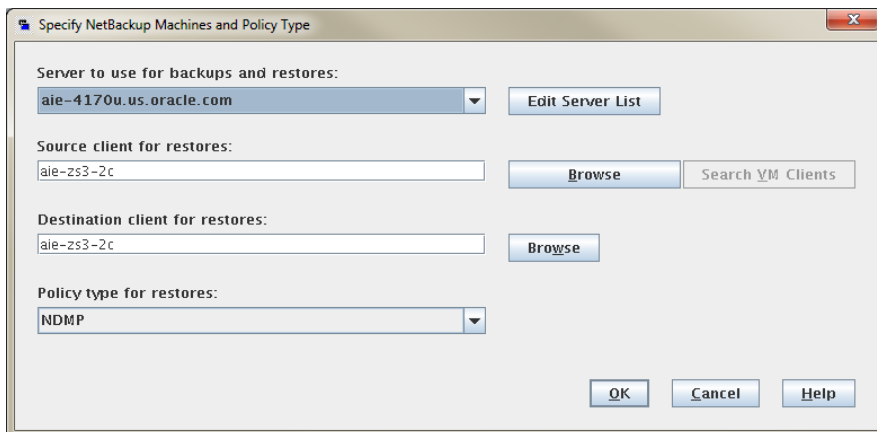


Figure 38. Defining the server and policy type for restores

**NOTE:** Even though the Oracle ZFS Storage Appliance should be resolvable by simple hostname or FQDN, the Symantec NetBackup restore operation seems to prefer simple hostname. Using the FQDN may result in Symantec NetBackup not finding the catalog entries.

- Specify the target directory.

**FOR TAR/DUMP NDMP:**

Specify the target directory in the “Browse directory:” field in the /export/sharename format, and the contents of the directory will be displayed in the right windowpane. Select the items to restore, and click Restore.

**FOR ZFS-NDMP:**

Specify the target directory in the “Browse directory:” field in the /Pool/local/Project/sharename format. Since history data is not provided with

ZFS-NDMP, no contents will be displayed in the right windowpane. Select the share and click Restore.

5. Provide the destination selections on the General tab, and click Start Restore.
6. Enter a name for the new policy and leave the Policy Configuration Wizard unchecked. Click OK.

## Appendix C: EMC NetWorker ZFS-NDMP Configuration

EMC NetWorker allows the backup administrator to add NDMP parameters to the backup specifications (clients). Those same NDMP parameters will be used for the restore.

The following sections outline the basic configuration steps needed to configure EMC NetWorker for NDMP operations in a local NDMP topology. It is assumed that the Oracle ZFS Storage Appliance has already been configured for NDMP operations. The actions will be performed using the EMC NetWorker Administration console included with EMC NetWorker 8.2. Other versions will be similar, but may have some differences.

Further information can be found in the EMC NetWorker Administrator's Guide.

### Configuring a Storage Node

The Oracle ZFS Storage Appliance attached devices must first be configured to EMC NetWorker as follows:

1. Navigate to the EMC NetWorker Administration console for the NetWorker host. Click on the Devices tab in the top row.




Figure 39. Selecting Devices in the EMC NetWorker Administration console

2. Right-click on the NetWorker host in the left pane and select Scan for Devices...
3. Click "Create a new Storage Node", provide the hostname of the Oracle ZFS Storage Appliance, select the ndmp radio button, and provide the NDMP user name and

password that was configured in the Oracle ZFS Storage Appliance NDMP service. Click Start Scan.

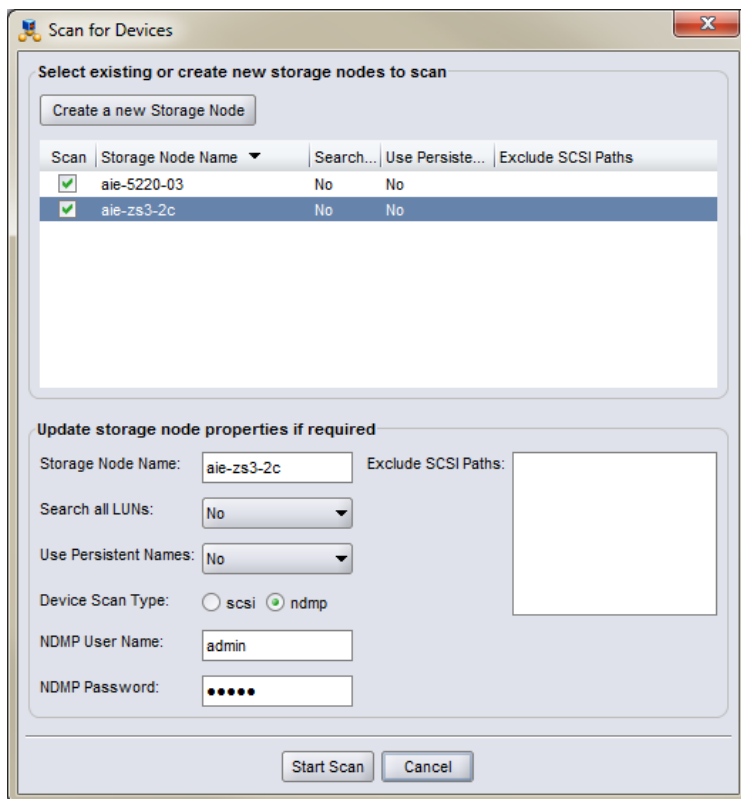


Figure 40. Designating the storage node in the Scan for Devices window

4. The Oracle ZFS Storage Appliance will show up in the left windowpane under "Storage Nodes". (A click on the Storage Nodes folder may be needed to refresh the tree.) Expand the tree to display all of the discovered devices. The library and tape drives will initially be offline and unconfigured.
5. Right click on the library name and select Configure Library... Select Check All to configure all of the tape drives to the library, and click Start Configuration.

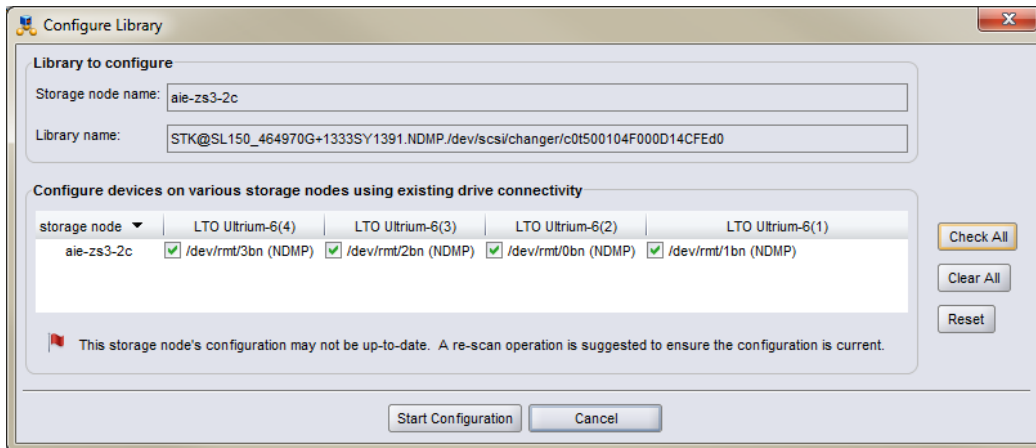


Figure 41. Configuring the library for the selected storage node

The tape library and tape drives should now be online and available for use.

6. Click on the Oracle ZFS Storage Appliance hostname under the Storage Nodes tree, right click the tape library name in the right windowpane, and select Properties.... From the Library Properties window, under the corresponding tabs, the following settings can be changed to match the environment:

- Change the library name (General tab).  
NOTE: The library will need to be disabled before the name can be changed.
- Disable EMC NetWorker auto cleaning (General tab).
- Set tape drive parallelism (Configuration tab).

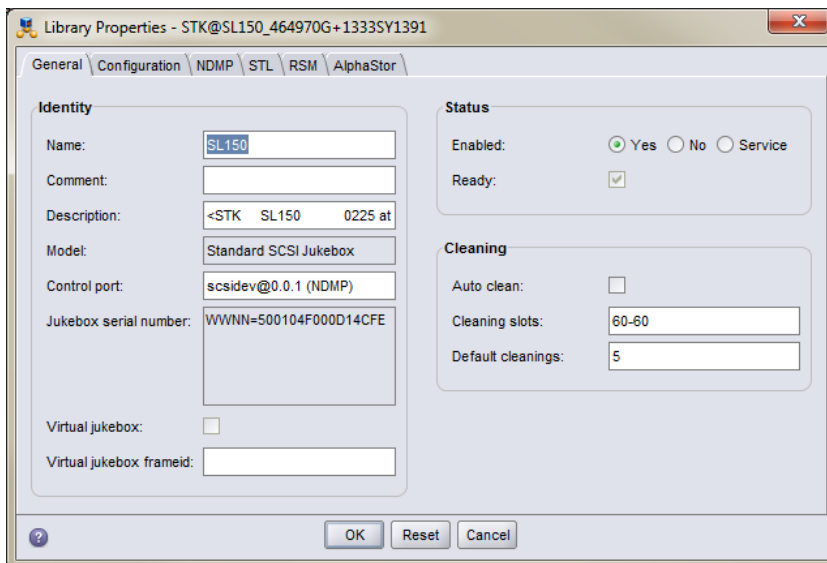


Figure 42. Changing properties for the library, selecting tabs for various properties' entries

7. Click OK to commit the changes.

## Labeling Tapes

After the library and tape drives are added, EMC NetWorker automatically performs a tape inventory of the library. The tapes, however, need to be labeled by performing the following:

1. Navigate to the EMC NetWorker Administration console for the NetWorker host and click on the Devices tab in the top row.
2. Expand the Libraries folder in the navigation tree and select the tape library.

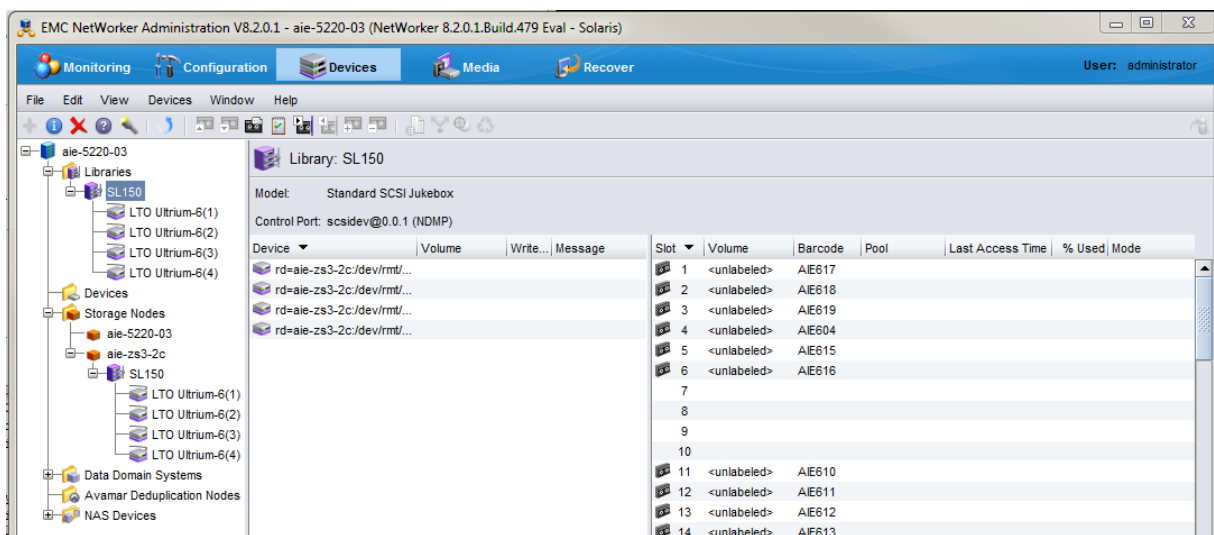


Figure 43. Selecting the tape library in the EMC NetWorker Administration console

3. In the right windowpane where the tape volumes are listed, select the media to be labeled, right-click on the media and select Label...
4. Select the desired target media pool and click OK.



Figure 44. Designating library media for labeling

### Adding an NDMP Client Definition

EMC NetWorker needs a client backup definition for the Oracle ZFS Storage Appliance to specify its backup contents. The NDMP client is created by performing the following steps:

1. Navigate to the EMC NetWorker Administration console for the NetWorker host and click on the Configuration tab in the top row. Right click on Clients in the navigation tree in the left windowpane, and select New Client Wizard.
2. Enter the hostname of the Oracle ZFS Storage Appliance and select the radio button next to "NDMP Client". Click Next.

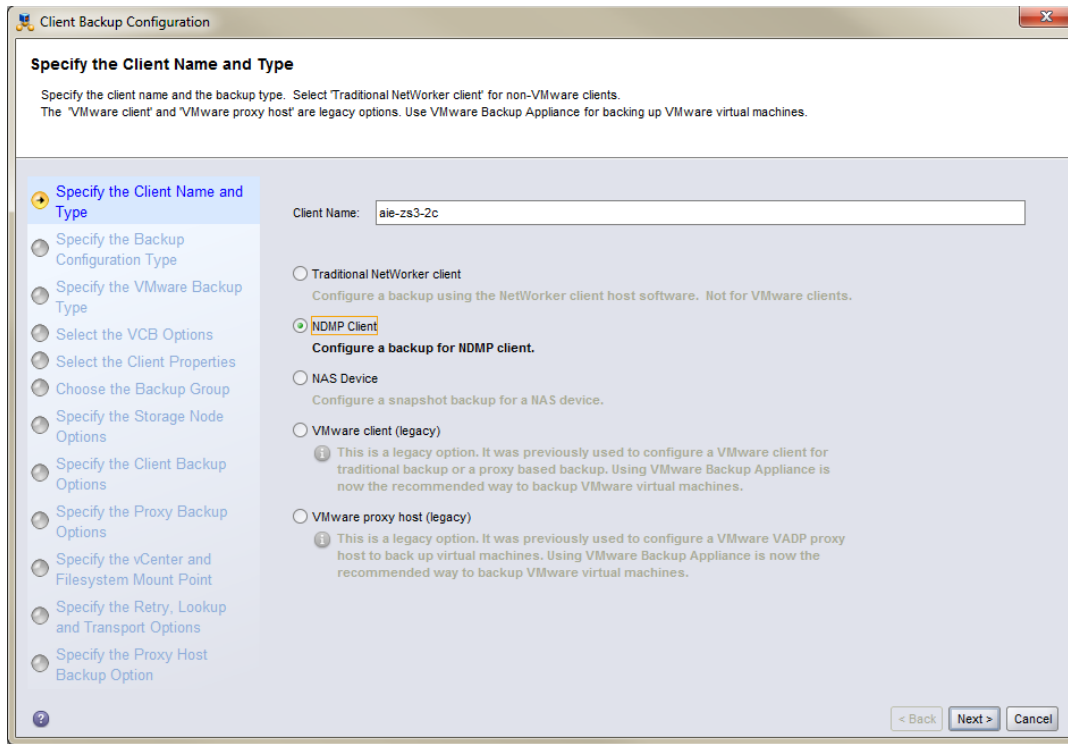


Figure 45. Designating the NDMP client

3. Specify the NDMP user name and password that was configured on the Oracle ZFS Storage Appliance NDMP service, and click Next.
4. Select the NDMP Backup Type.

FOR TAR/DUMP NDMP:

Select “dump” or “tar” as the NDMP Backup Type, enter the Oracle ZFS Storage Appliance hostname, and leave all App Info: checkboxes selected. Click Next.

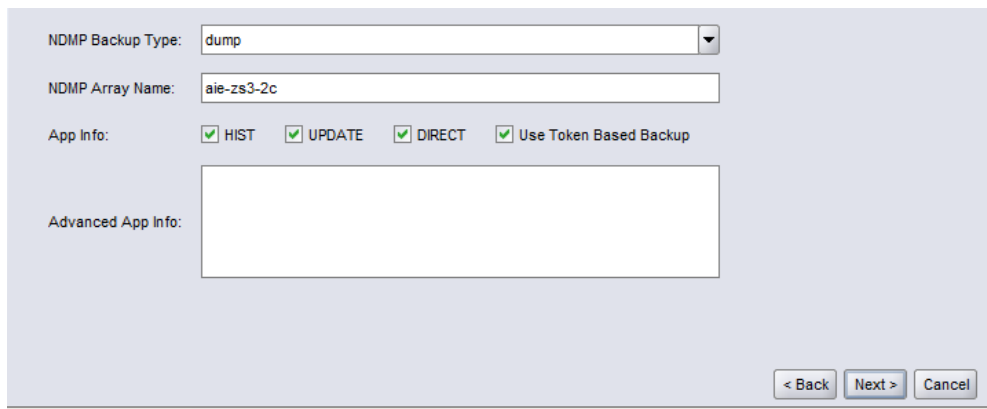


Figure 46. Defining the backup type for Tar/Dump NDMP



**FOR ZFS-NDMP:**

Select “zfs” as the NDMP Backup Type, enter the Oracle ZFS Storage Appliance hostname, and only leave the Update checkbox selected. Click Next.

The screenshot shows a configuration window with the following fields and options:

- NDMP Backup Type:** A dropdown menu with 'zfs' selected.
- NDMP Array Name:** A text input field containing 'ale-zs3-2c'.
- App Info:** A group of four checkboxes:
  - HIST
  - UPDATE
  - DIRECT
  - Use Token Based Backup
- Advanced App Info:** An empty text area.
- Navigation:** Three buttons at the bottom right: '< Back', 'Next >', and 'Cancel'.

Figure 47. Defining the backup type for ZFS-NDMP

5. Select the desired Target Pool, and click Next.
6. Specify the shares on the Oracle ZFS Storage Appliance.

**FOR TAR/DUMP NDMP:**

Specify the shares on the Oracle ZFS Storage Appliance in the `/export/sharename` format, and click Next.

**FOR ZFS-NDMP:**

Specify the shares and/or LUNs on the Oracle ZFS Storage Appliance in the `/Pool/local/Project/volumename` format, and click Next.

7. Make any necessary policy changes (if any), and click Next.
8. Make any necessary group selection changes (if any), and click Next.
9. Only check the box next to the Oracle ZFS Storage Appliance hostname to select its library and tape drive resources as the backup target, and click Next.

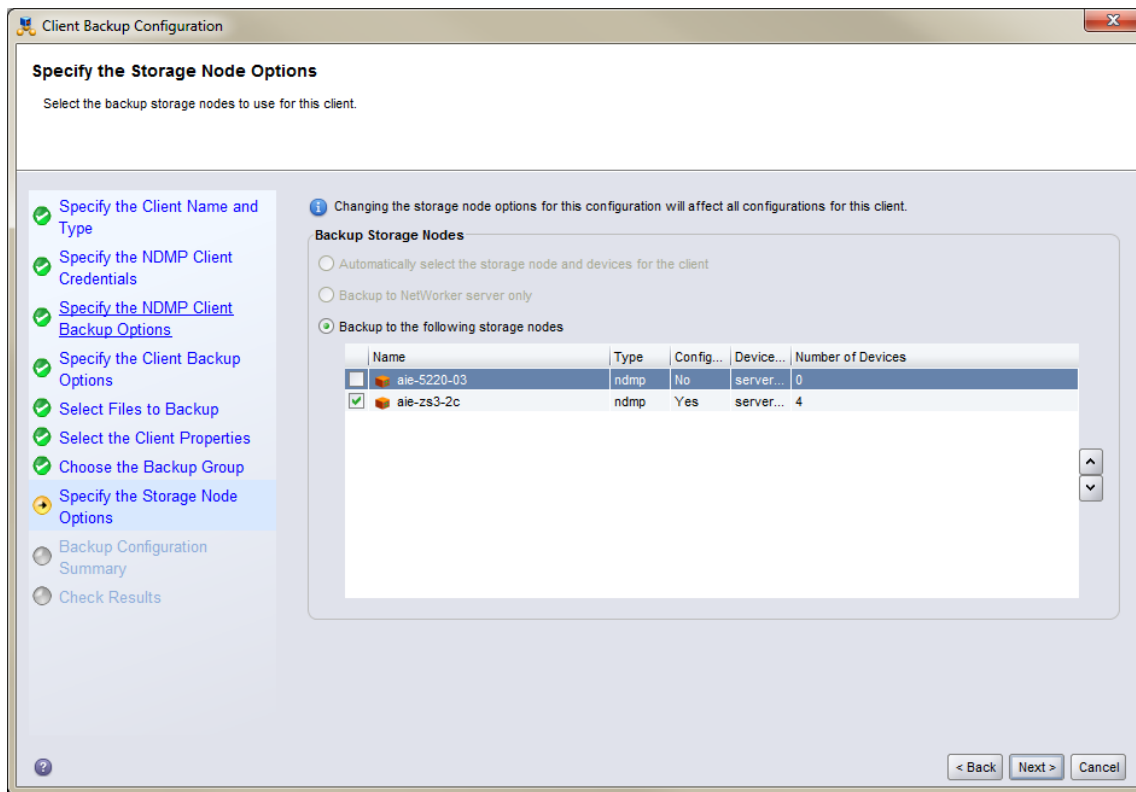


Figure 48. Selecting the backup storage nodes for the client

10. Review the Backup Configuration Summary and click Create.
11. The new client should be successfully added and listed. Click Finish to close the wizard.

### Changing NDMP Parameters in the Client Definition

The NDMP parameters are mostly set during the client definition. However, in the case where additional NDMP parameters need to be set:

1. Navigate to the EMC NetWorker Administration console for the NetWorker host and click on the Configuration tab in the top row. Click on Clients in the navigation tree to display the defined clients in the right windowpane.
2. Right-click on the client that needs to be updated and select Modify Client Properties.
3. Click on the Apps & Modules tab. Some NDMP parameters will already be set in the Application Information section of the property page. This is where additional NDMP parameters can be specified. Click OK or Cancel to exit.

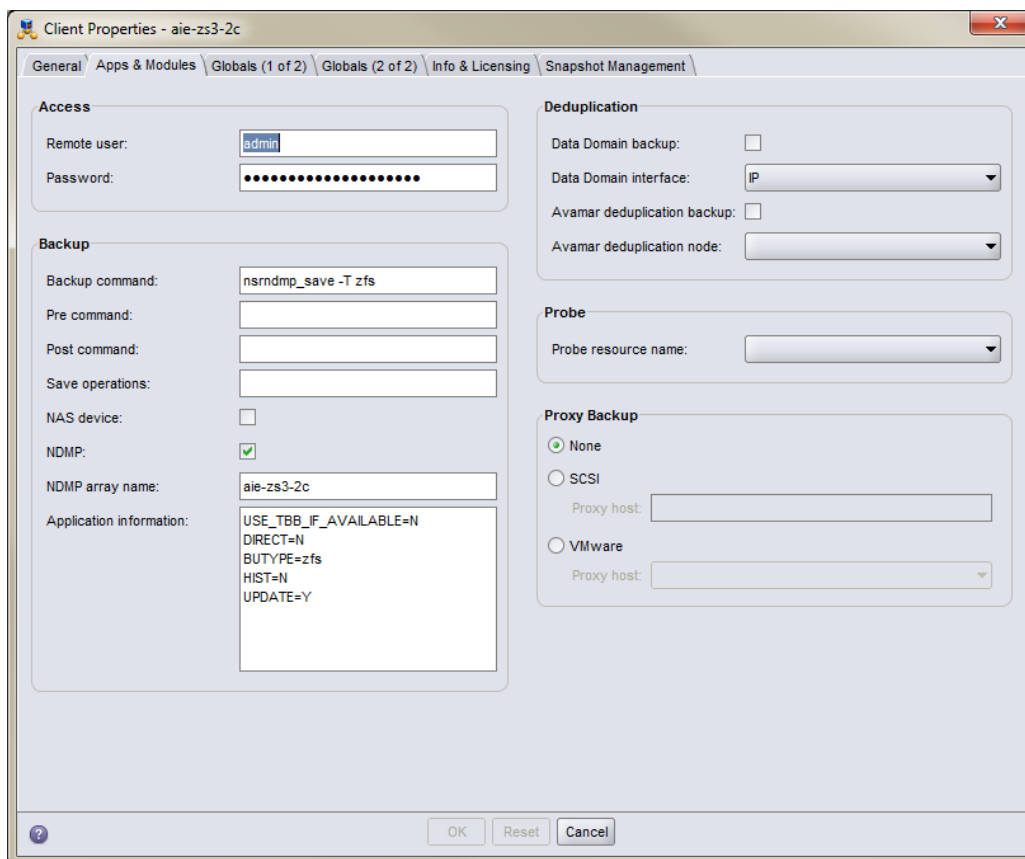


Figure 49. Setting additional NDMP parameters for the defined client

## Restoring from an NDMP Backup

The following steps can be used to restore from NDMP Backups:

1. Navigate to the EMC NetWorker Administration console for the NetWorker host and click on the Recover tab in the top row. Click on the green plus sign (+) (just below the File menu) to create a recovery definition.
2. Provide the hostname of the Oracle ZFS Storage Appliance in the Source Host section and the available recovery types should list in the lower window. Click Next.

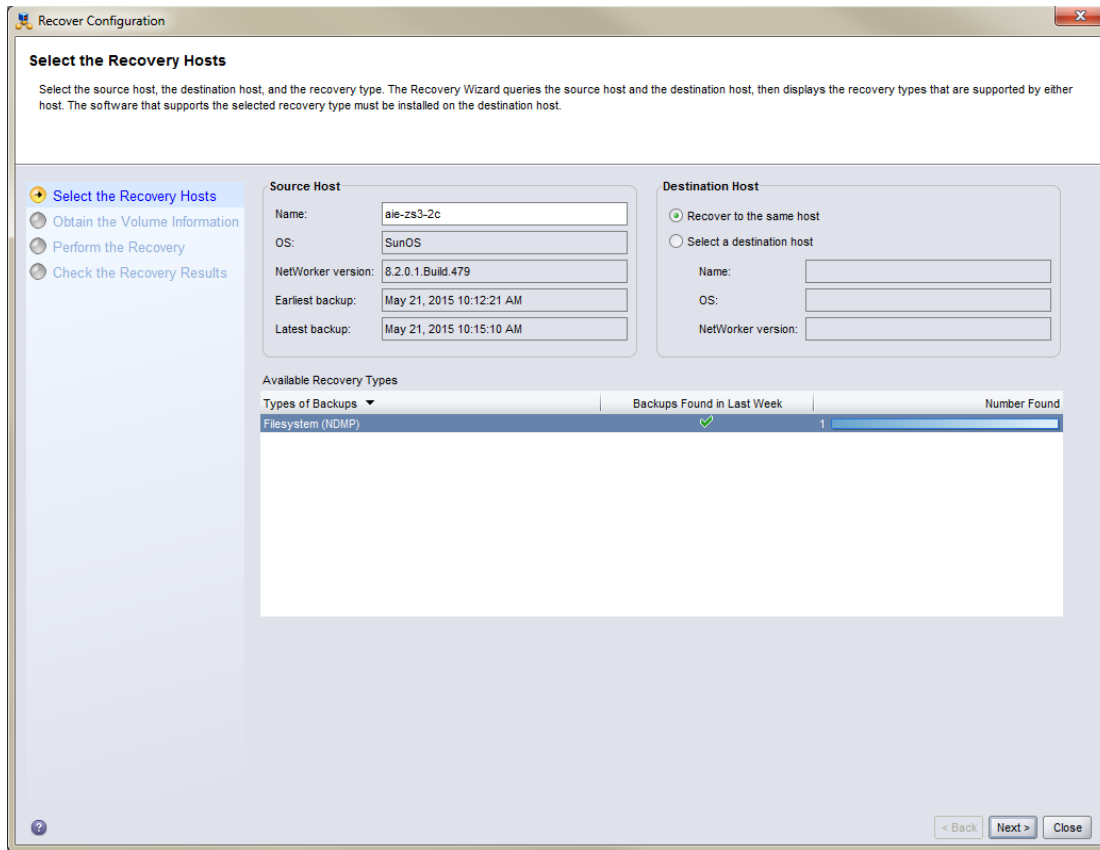


Figure 50. Setting up a restoration in the NetWorker Administration console

3. Select the target for restoration.

FOR TAR/DUMP NDMP:

The Browse tab should already be selected and the Oracle ZFS Storage Appliance hostname should appear in the directory tree. Expand the tree and select the files that will be restored. Click Next.

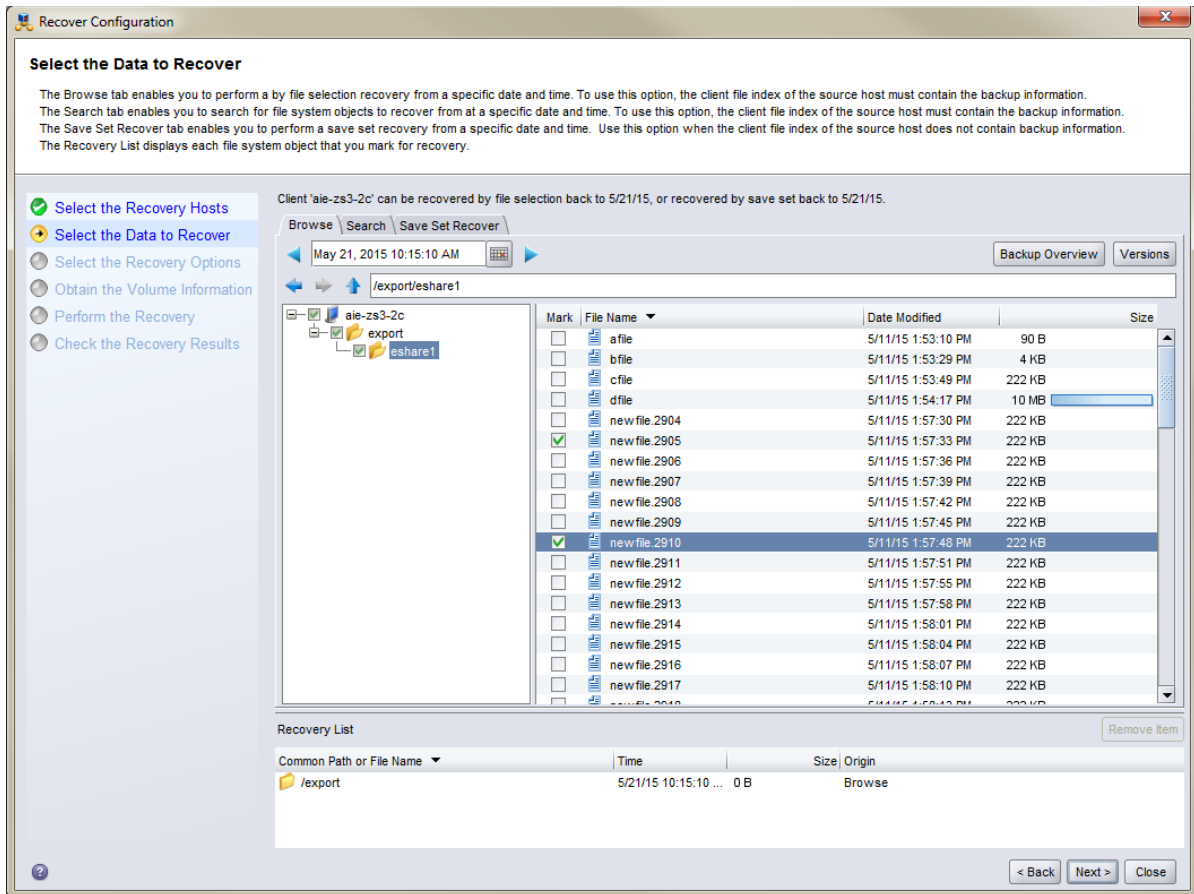


Figure 51. Selecting the restoration/recovery target for Tar/Dump NDMP

FOR ZFS-NDMP:

Click the Save Set Recover tab. Choose the search criteria in the 'Found in' section, and click Query. The available Save Set Names will appear. Select the desired save set name and click **Next**.

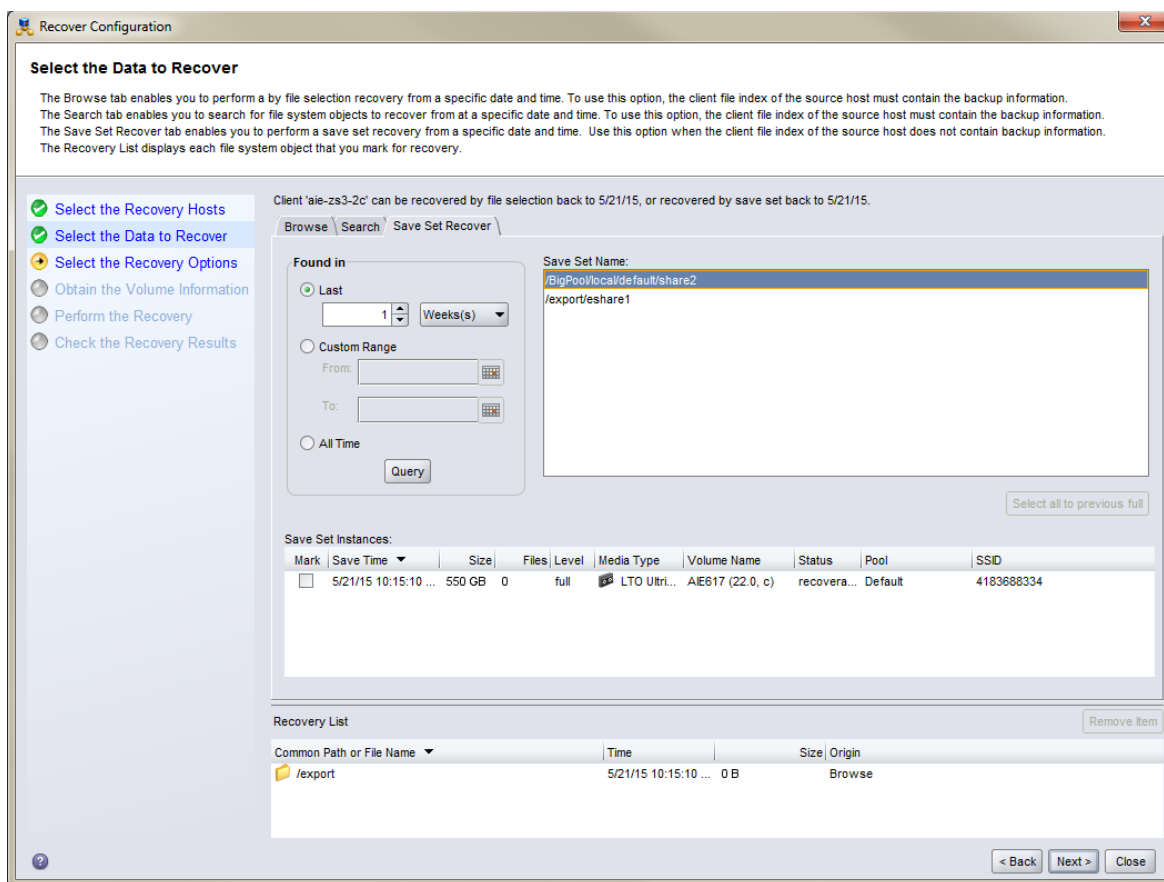


Figure 52. Selecting the restoration/recovery target for ZFS-NDMP

- On the Select the Recovery Options window, select the desired restore path option, and click Next.

FOR TAR/DUMP NDMP: Use the `/export/sharename` format for the restore path.

FOR ZFS-NDMP: Use the `/Pool/local/Project/volumename` format for the restore path.

- On the Obtain Volume Information window, keep the defaults and click Next.
- On the Perform the Recovery window, enter a job name in the 'Recover name:' field, and click Run Recovery.

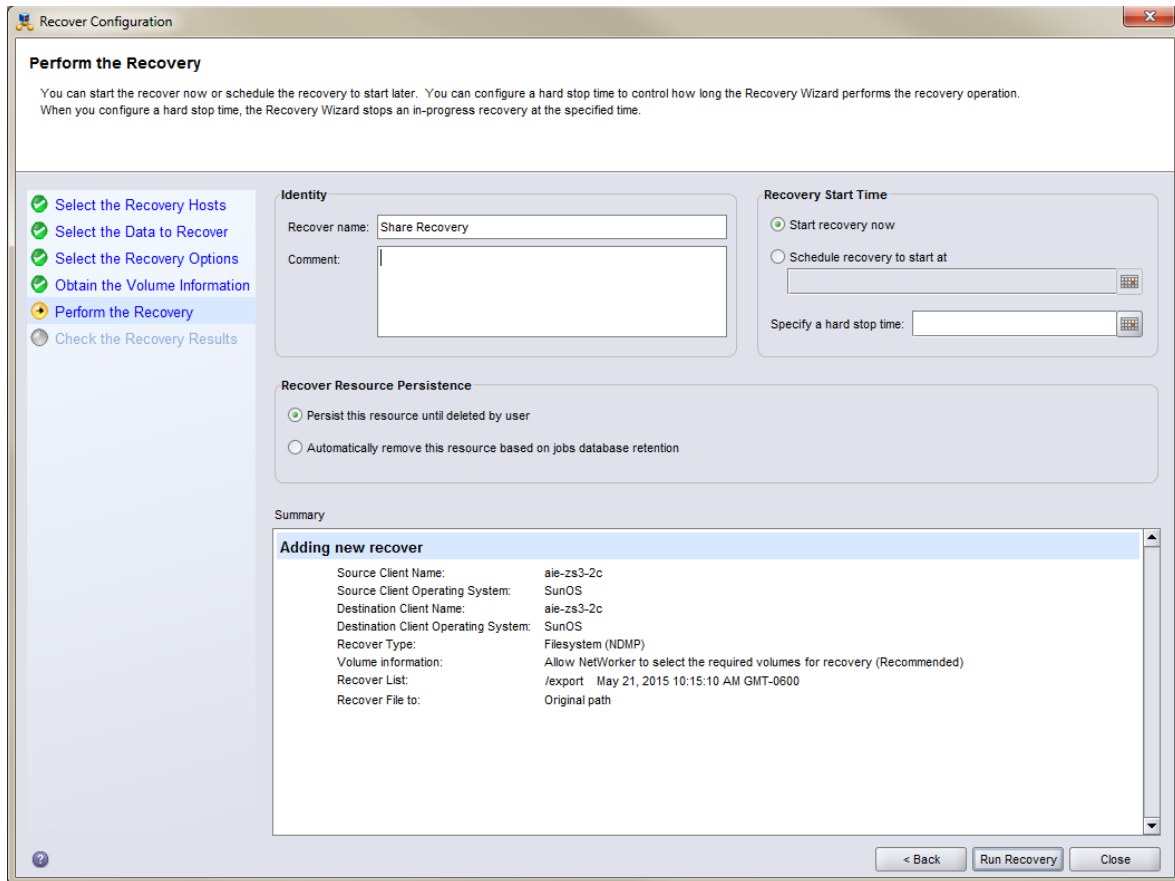


Figure 53. Naming the recovery job and initiating the run

## Appendix D: References

See the following resources for additional information relating to the products covered in this document.

- Oracle ZFS Storage Appliance White Papers and Subject-Specific Resources  
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>
- Oracle ZFS Storage Appliance Product Information  
<https://www.oracle.com/storage/nas/index.html>
- Oracle ZFS Storage Appliance Documentation Library, including Installation, Analytics, Customer Service, and Administration guides:  
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>

The *Oracle ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help context.

The Help function in Oracle ZFS Storage Appliance can be accessed through the browser user interface.





NDMP Implementation Guide for the Oracle ZFS  
Storage Appliance  
June 2015, Version 2.0  
Author: Eric Polednik  
Version 1.0 Author, Ryan Ameson

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 © 2015, 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.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

**Hardware and Software, Engineered to Work Together**

