Technical Report

Open Systems SnapVault (OSSV) Best Practices Guide

TR-3466 | Revised for OSSV 3.0

ABSTRACT

This document is a guide to help aid in the understanding and deployment of Open Systems SnapVault[®] (OSSV). Open Systems SnapVault is a disk-to-disk backup and recovery solution for protecting data residing on third-party storage and platforms.

TABLE OF CONTENTS

1 INTRODUCTION 3

2 OVERVIEW 3

2.1 THEORY OF OPERATIONS 4

3 OSSV FEATURES 5

4 SETUP AND ADMINISTRATION 9

4.1 COMMAND LINE INTERFACE 12

4.2 PROTECTION MANAGER 17

5 PROTECTING MICROSOFT SQL SERVER 21

6 MICROSOFT CLUSTER SERVER 25

7 BEST PRACTICES 26

8 NETAPP PARTNERS 29

LIST OF FIGURES

Figure 1) OSSV transfers data directly to NetApp storage. 3

Figure 2) OSSV Configurator. 11

Figure 3) Add OSSV hosts using the “Add OSSV Host Wizard.” 17

Figure 4) Create a resource pool 18

Figure 5) Create a copy of the “Remote backup only” policy. 18

Figure 6) Create a copy of the schedule template. 19

Figure 7) Create a dataset using the “Add Dataset Wizard.” 19

Figure 8) Manually select the source data to be managed by the dataset. 20

Figure 9) Launch the “Protection Policy Wizard.” 20

Figure 10) Protection Manager lists instances and databases as folders. 24

Figure 11) Change the “NDMP Host ID” and “NDMP Host Name” using the OSSV Configurator. 26

1 INTRODUCTION

Open Systems SnapVault is a disk-to-disk backup and recovery solution to protect data residing on non NetApp storage systems and platforms. This agent-based solution transfers data directly from an OSSV host to a NetApp secondary storage system in the form of block-level incremental backups. These backups are captured as Snapshot™ copies on the NetApp secondary system. The advantage is fast, reliable, space-optimized backups centralized on NetApp technology.

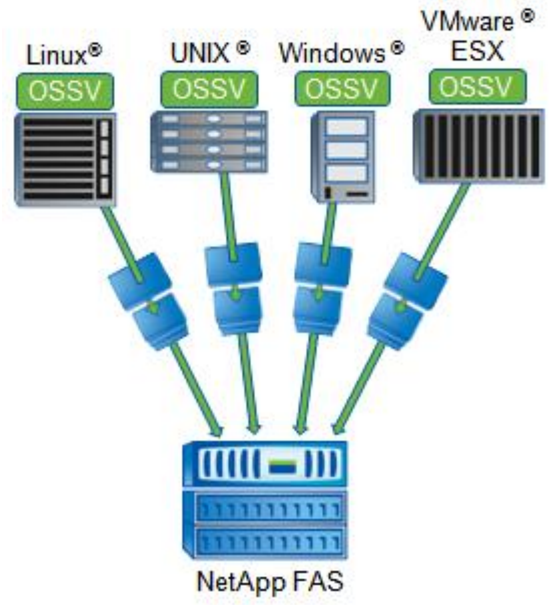

Figure 1) OSSV transfers data directly to NetApp storage.

2 OVERVIEW

The basic components that make up an OSSV architecture are as follows:

- OSSV host
- OSSV agent software
- TCP/IP network
- NetApp storage system

OSSV supports Windows, UNIX, Linux, and VMware ESX clients. For a complete list of supported versions refer to the Open Systems SnapVault® 3.0 Release Notes.

Data on these clients are backed with OSSV. Unlike traditional tape-based backup solutions, OSSV only backs up the blocks that have changed since the last backup. The only time a full backup is performed is during the very first backup.

OSSV backups are captured on the NetApp secondary system in Snapshot copies for retention. Different backup schedules can maintain different levels of retention. For example, a user may want to keep 6 hourly backups, 30 daily backups, 2 years of weekly backups, and 5 years of monthly backups. During a recovery, data is restored directly from the NetApp Snapshot copy.

In addition, the backup contents contained in the Snapshot copies represent a full data backup and are completely accessible and readable by the user since OSSV stores backups in a native format.

2.1 THEORY OF OPERATIONS

Open Systems SnapVault relationships are what map source data to the NetApp secondary system. OSSV source data can reside on locally attached disk or SAN attached disk and can consist of a disk, file system, or subdirectory.

Note: Open Systems SnapVault will not back up CIFS or NFS data.

On the NetApp secondary system, OSSV backups for a given relationship will map to a qtree within a volume. Many OSSV relationships can share the same destination volume, but each relationship will have its own qtree. Likewise, each OSSV host may be a member of multiple relationships.

RELATIONSHIP CREATION AND BASELINE TRANSFER

Creating an OSSV relationship also invokes a baseline transfer (the initial full or level-0 backup). A relationship is created on the NetApp secondary system by issuing **`snapvault`** commands or by using NetApp Protection Manager software. When the relationship is created, the NetApp secondary system contacts the OSSV client and requests the baseline transfer of the source data. It's important to know that backup and restore operations are done as a "pull." Backups are pulled by the NetApp secondary, while restores are pulled by the OSSV host. In order for OSSV to perform block-level incremental transfers after the baseline is complete, checksums are calculated for every 4kB block of source data. This checksum information is stored in an internal database on the OSSV host. The source data is then transferred to a qtree on the NetApp secondary. When the baseline transfer is complete, a Snapshot copy is taken.

INCREMENTAL BACKUPS

OSSV backups performed after the initial baseline transfer has completed are always block-level incremental by default. By using the 4kB block checksums previously calculated, OSSV is able to back up only the blocks within the files that have changed. Because of this built-in efficiency, OSSV is well suited for environments with slower network links. After the incremental backups are complete, a Snapshot copy is taken and kept as a recovery point based on the retention requirement.

There are two phases that take place during a backup. The first phase performs a file system scan on the OSSV host and a directory structure build on the NetApp secondary. Information about file deletions is also sent to the NetApp secondary. In the second phase checksum calculations are performed on modified files, and the backup data is transferred to the NetApp secondary.

INTERNAL DATABASE

Each OSSV host maintains a state database for each of its relationships. This database contains the following files:

- History file
- Block-level incremental checksum file
- Checkpoint file

By default, each relationship's database is backed up as a part of the relationship's normal OSSV backup.

3 OSSV FEATURES

This section covers the features built into OSSV. For more information on these features refer to the Open Systems SnapVault® 3.0 Installation and Administration Guide.

Default file and command locations on OSSV hosts are listed below.

<install_dir> on Windows is `C:\Program Files\netapp\snapvault`

<install_dir> on Linux and UNIX is `/usr/snapvault`

BLOCK-LEVEL INCREMENTALS

OSSV block-level incremental (BLI) backup functionality is designed to minimize the amount of data transferred during a backup by identifying only those blocks within each file that have changed.

BLI transfers allow more frequent backups and improved recovery point objectives by reducing the amount of time required to complete each backup. In addition, storage resources required to retain the backup data are minimal when compared to backup technologies that use file-level incrementals with frequent fulls.

In order for OSSV to identify changed blocks, OSSV first checks file modification times. Checksum values are then calculated and preserved for each 4kB block of data in those new or modified files. These checksum values are stored in the OSSV database. The size of the database will typically be about 2% of the size of the source data.

There are configuration options available in the OSSV Configurator utility on the OSSV host that control specific BLI behavior. These BLI settings are as follows:

- **High -** Always compute checksums (on baseline and incremental)
- **Low -** Compute checksums only on changed files, and do not compute checksums during baseline
- **Off -** Disable BLI functionality, and never compute checksums

The default setting is “high” and is preferred in most cases. A setting of “low” reduces the time and resources it takes to perform the baseline, but incremental updates take longer and storage requirements increase. “Off” completely disables BLI functionality and might make sense when files are very small or when file changes encompass the entire file.

FILTER DRIVER

OSSV 3.0 includes a built-in filter driver. This driver tracks changed block on the file system in real time and reduces backup times by decreasing changed block detection time and reducing the amount of checksum calculations that must be performed. The filter driver is ideal for large file structures such as databases. By default, the filter driver is enabled for application data (such as SQL® Server) and disabled for normal file data. This behavior can be modified in the `snapvautlt.cfg` file on the OSSV host.

The following parameters are available in the `snapvautlt.cfg` for enabling and disabling the filter driver:

`OSSV:UseChangelogsForFileSystems`	True / False (default is False)
`OSSV:UseChangelogsForApps`	True / False (default is True)

NAME-BASED BLI

In some cases, applications modify files by:

1. Creating a temporary copy of the original file
2. Making the necessary changes to that temporary file
3. Deleting the original file
4. Saving the temporary file under the same name as the original file

OSSV can detect this condition and treat the new instance of the renamed temporary file as the updated original file without having to transfer the entire file.

In other cases, applications make changes to files by:

1. Inserting data into or removing data from the middle of the file
2. Rewriting all subsequent data blocks in the file to new positions in the file

Microsoft® Word, Excel, and PowerPoint are some of the applications that are known to exhibit this behavior when saving changes to files. For files that are modified in this manner, OSSV backs up all blocks in the file that have different positions or different checksum values.

OPEN FILE BACKUPS

OSSV integrates with Microsoft Volume Shadow Copy Service (VSS) in order to protect Windows open files. During backup, OSSV on a Windows host triggers a VSS snapshot for the volume it is protecting. It uses this snapshot as the source for the backup. When the backup is complete, OSSV removes the VSS snapshot on the host. In the event that the backup is interrupted, OSSV retains the VSS snapshot and uses it during a checkpoint restart. By default, the VSS snapshot remains available for restarts for 10 minutes. If after 10 minutes a retry has not started, OSSV removes the snapshot. This timeout value can be changed using the OSSV Configurator utility on the OSSV host.

By default, OSSV has a 180-second timeout value for which it will await VSS snapshot creation. This timeout can be changed using the OSSV Configurator utility on the OSSV host.

To disable VSS on a particular volume, populate the “List of Drives/Mount points not to Snapshot” field in the OSSV Configurator utility.

Note: VSS snapshots triggered by OSSV during backup are not available for use by other processes.

SYSTEM STATE BACKUP AND RESTORE

Windows system state data, including the registry and the event logs, can be backed up and restored with OSSV.

Using the keyword “SystemState” in the primary path initiates a system state backup of the OSSV host:

```
snapvault start -S ossv1:SystemState sv_secondary:/vol/sec_vol/sec_qtree
```

Boot files and system files are backed up even when they are on different volumes. Subsequent backups use block-level incremental backups.

Restores also use the keyword “SystemState” in place of the file system path:

```
<install_dir>\bin\snapvault restore -S sv_secondary:/vol/sec_vol/sec_qtree SystemState
```

In certain Active Directory® environments, there are other options for restoring system state data. The keyword may change to **`SystemStatePrimary`** when restoring system data from a backup and marking it primary.

SystemState backups can also be used as part of a disaster recovery plan. To accomplish this, back up the entire system drive, including any other relevant partitions or drives and the SystemState. Be aware that when recovering from a disaster using a complete system backup, OSSV does not support “bare metal restore.” Prior to restore, the base operating system must be installed on identical hardware with identical service packs, names, drive letter mappings, file system types, etc.

CHECKPOINT RESTART

OSSV takes checkpoints every five minutes, by default, during the data transfer phase of the backup. If, during the backup, an interruption occurred, OSSV will restart from the most recent checkpoint. This is true for the baseline transfer as well as all incremental backups.

To modify the checkpoint interval, edit the following parameter in the **`<install_dir>\config\snapvault.cfg`** file. The value is specified in seconds, and the default value is 300. The minimum value is 60.

```
[QSM:Checkpoint Interval]
```

OSSV will retry an interrupted backup based on the “tries” limit set for the relationship. By default, the “tries” limit is 2. Effectively this means 1 retry, since the initial attempt counts as a “try.” To increase the number of retries to 2, set the “tries” limit to 3 using the –t flag in the **`snapvault start`** or **`snapvault modify`** command on the NetApp secondary. For example:

```
snapvault start –t 3 -S ossv1:c:\ fas1:/vol/backups/ossv1
```

```
snapvault modify –t 3 -S ossv1:c:\ fas1:/vol/backups/ossv1
```

Setting the “tries” limit to 0 will effectively disable the backups.

DYNAMIC THROTTLING

OSSV hosts can throttle backups based on a specific schedule. This schedule is defined in the **`wan.cfg`** file on the OSSV host. By default, OSSV will check for changes to this file every 15 minutes and at the beginning of each backup. Throttle settings defined by the **`wan.cfg`** file are systemwide and shared by all transfers from that particular host.

An example entry in the **`wan.cfg`** might look like the following:

```
Value=mon-thu@9-18#100!,fri@9-18#150!,18-21#200
```

This schedule is translated as:

Monday–Thursday	9 a.m. to 6 p.m. 100KB/sec
Friday	9 a.m. to 6 p.m. 150KB/sec
Every day	6 p.m. to 9 p.m. 200KB/sec

For periods of time not defined by the schedule, no throttle is enforced.

Bandwidth units in the **`wan.cfg`** file are in KB/sec. If throttling is also configured for the relationship on the NetApp secondary (using Protection Manager or the **`–k`** flag in the **`snapvault`** command), the lesser of the values will be used.

COMPRESSION

OSSV has the capability to compress data as it is sent across the network during backup and restore operations. Compression can be enabled on the NetApp secondary for individual relationships (Data ONTAP® 7.3 or higher), or it can be enabled globally (Data ONTAP 7.3.1 or higher).

To enable compression for a specific relationship, use the **`–o compression=on`** option within the **`snapvault start`** or **`snapvault modify`** command on the NetApp secondary. For example:

```
snapvault start –o compression=on -S ossv1:c:\ fas1:/vol/backups/ossv1
```

To enable compression globally, use the following command on the NetApp secondary:

```
options snapvault.ossv.compression on
```

For restores, compression can be enabled by adding the **`–c`** flag to the **`snapvault restore`** command on the OSSV host. For example:

```
<install_dir>\bin\snapvault restore –c –S fas1:/vol/backups/ossv1 c:\temp\restore
```

Compression behavior can be tuned in the **`snapvault.cfg`** file on the OSSV host. The following parameters are available:

`QSM:CompressionLevel`	LOW, MEDIUM, or HIGH (default is MEDIUM)
`QSM:CompressionLowPriority`	True / False (default is False)
`QSM:EnableCompression`	True / False (default is True)

The **`QSM:CompressionLevel`** parameter determines how much compression will take place. However, it can also impact the amount of time it takes for the compression operations to run.

The **`QSM:CompressionLowPriority`** parameter determines the CPU priority allowed for compression operations. The default is “False.” Setting this to “True” results in longer compression times.

The **`QSM:EnableCompression`** parameter can be used to allow or disallow compression on a particular OSSV host.

DEDUPLICATION

Storage efficiency is a core value of OSSV. Block-level incremental technology eliminates much of the redundant data that would otherwise be backed up. NetApp deduplication, however, is also integrated with OSSV to provide even more space savings and efficiency on the NetApp destination system. Since NetApp deduplication is performed at the volume level, it is a good idea to direct OSSV relationships with common data to the same volume.

It is best to enable deduplication on the NetApp secondary volume prior to performing the baseline transfer. Deduplication, after being enabled on the secondary volume, runs automatically as backups complete. Because of this integration, deduplication does not run via a deduplication schedule.

Note: When using Protection Manager, a Provisioning Manager license is required in order to automatically provision secondary volumes with deduplication enabled. A provisioning policy that enables “on-demand” deduplication for backups can be created and applied to a dataset.

DATA EXCLUSIONS

OSSV can exclude certain data from backup. The following files on the OSSV host can be populated to exclude specific data:

```
<install_dir>\etc\file-exclude.txt
<install_dir>\etc\path-exclude.txt
<install_dir>\etc\file-system-exclude.txt
```

The file-exclude.txt and path-exclude.txt files accept wildcards. The file-system-exclude.txt file does not.

Note: The file-system-exclude.txt file is not available on Windows or AIX hosts.

IPv6 SUPPORT

OSSV supports IPv4 and IPv6 for backup and restore operations as well as internal communications within the OSSV host. IPv6 requires Data ONTAP 7.3.3 or higher.

OSSV supports IPv6 on the following platforms:

- Microsoft Windows 2003 and 2008
- Red Hat® Enterprise Linux
- Novell® SUSE® Linux Enterprise Server
- Oracle™ Solaris™
- IBM AIX
- HP-UX

IPv6 IP addresses should be enclosed with brackets in `snapvault` commands. For example:

```
snapvault start -S [0ffA::88fe:3456:7654:AA34]:c:\ fas1:/vol/backups/ossv1
```

SUPPORT FOR VOLUME MOUNT POINTS

OSSV can protect data using mounted folders on Windows 2003 and Windows 2008. Data can also be restored to a mounted folder or to a folder within a mounted folder. When restoring both the volume and the mounted folder within the volume, first restore the volume. After restoring the volume, restore the volume mount point.

LREP TOOL

OSSV includes a utility called LREP that can be used to perform baseline transfers and restores out of band. This is useful when network bandwidth is such that baseline transfers and restores would take a significant amount of time to complete. With the LREP utility, a baseline transfer can be written to a portable media device such as a USB disk drive, shipped to the remote location, and used to seed the data on the NetApp secondary. To secure the data as it travels on the portable media, the LREP utility can compress and encrypt the data.

The LREP utility is bundled with OSSV and is also available separately from the Utility ToolChest on the NOW site. For detailed information on using the LREP tool, refer to the LREP documentation available from the Utility ToolChest.

4 SETUP AND ADMINISTRATION

For details and requirements for installing OSSV software, refer to the Open Systems SnapVault® 3.0 Installation and Administration Guide. However, a useful tool for viewing the details of an installation on an OSSV host is the **svinstallcheck** command. This command displays the OSSV version, path information, space availability, NDMP listening port, drives that are suitable for backup, and more.

LICENSING

To use OSSV, a SnapVault secondary license is required on the NetApp secondary system. To determine if the license exists, run the **license** command and look for the **sv_ontap_sec** key.

In addition to the SnapVault secondary license, OSSV requires a license for each platform being protected. For example, to enable backups for Windows hosts an **sv_windows_pri** license must be enabled on the NetApp secondary system. Client licenses are free and can be obtained from the NOW site.

OSSV client licenses are as follows:

Windows – **sv_windows_pri**

Linux – **sv_linux_pri**

UNIX – **sv_linux_pri**

VMware – **sv_vi_pri**

To add a license key on the NetApp secondary, use the **license add** command.

Note: All licensing is done on the NetApp secondary system. There are no licenses on the OSSV host.

In addition to the licenses above, NetApp recommends that a NearStore® Option license be enabled on the NetApp secondary to allow the maximum number of concurrent transfers. To determine if the NearStore Option has been licensed, run the **license** command and look for the **nearstore_option** key.

OSSV can be set up and managed using the command line interface (CLI) or Protection Manager.

ENABLING OSSV

OSSV must be properly enabled before use to allow communications between the OSSV host and the NetApp secondary.

On the NetApp secondary, make sure SnapVault is enabled by running the **options snapvault.enable** command. If it is not enabled, turn it on with the following command.

```
options snapvault.enable on
```

OSSV communicates to the NetApp secondary on port 10566. This port cannot be changed and should be opened on any firewalls in the network path. In addition, OSSV hosts will listen for NDMP-based management applications such as Protection Manager on port 10000 by default. This port is set during the OSSV software installation and can be changed by modifying the “NDMP Listen Port” field in the OSSV Configurator utility on the host. When using Protection Manager, any changes to the NDMP listening port will need to be modified in Protection Manager as well. The port number can be changed within Protection Manager by modifying the properties for the OSSV host.

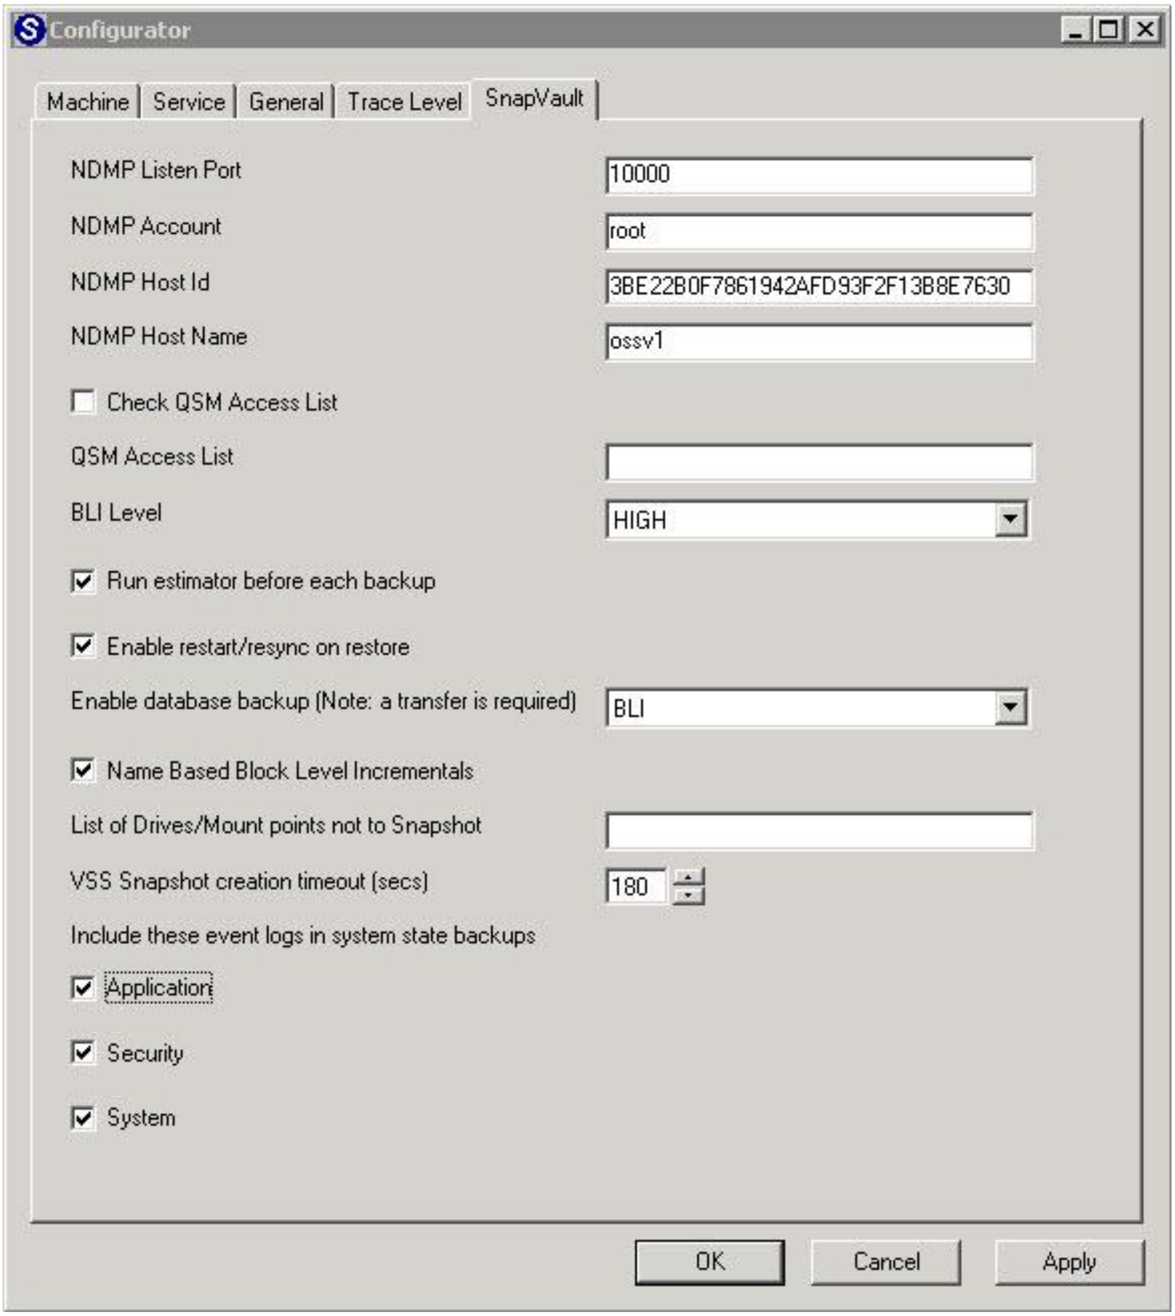

NDMP Listen Port: 10000

NDMP Account: root

NDMP Host Id: 3BE22B0F7861942AFD93F2F13B8E7630

NDMP Host Name: ossv1

□ Check QSM Access List

QSM Access List:

BLI Level: HIGH

☑ Run estimator before each backup

☑ Enable restart/resync on restore

Enable database backup (Note: a transfer is required): BLI

☑ Name Based Block Level Incrementals

List of Drives/Mount points not to Snapshot:

VSS Snapshot creation timeout (secs): 180

Include these event logs in system state backups

☑ Application

☑ Security

☑ System

Figure 2) OSSV Configurator.

In order to restrict the NetApp secondary systems that are allowed to initiate backups from an OSSV host, the QSM Access List field can be populated with the hostname of the secondary using the OSSV Configurator utility. When using this option, make sure the “Check QSM Access List” box is enabled.

The `snapvault.access` option on the NetApp secondary can be used to restrict the OSSV hosts that are allowed to initiate restores from that secondary system. This option can be populated with a list of hostnames for each OSSV host. For example:

```
options snapvault.access host=ossv1,ossv2,ossv3
```

Note: This option is not required for restores done via CIFS or NFS.

On the OSSV host, the OSSV service can be stopped and started using the CLI or the OSSV Configurator. These actions are often required after making OSSV configuration changes on the host. To control the OSSV service using the CLI, the following commands are available.

```
<install_dir>\bin\snapvault stop

<install_dir>\bin\snapvault start

<install_dir>\bin\snapvault restart
```

Using the OSSV Configurator, the service controls are located on the Service tab.

4.1 COMMAND LINE INTERFACE

When using the CLI, there are several commands that are used to perform various operations.

PERFORMING BASELINE TRANSFERS

Before performing a baseline transfer, a target volume should be available on the NetApp secondary. For new volumes that will be used as an OSSV destination, NetApp recommends disabling the normal snapshot schedule as well as snapshot reservations. For example, to create a 1TB volume called “backups” on the aggregate, “aggr1,” run the following command on the NetApp secondary:

```
vol create backups aggr1 1t

snap sched backups 0 0 0

snap reserve backups 0
```

The **`snap sched`** command disables the normal snapshot schedule, and the **`snap reserve`** command disables the snap reserve.

To create a relationship and initiate a baseline transfer, run the **`snapvault start`** command from the NetApp secondary. For example, to protect the C:\ drive on the OSSV host, “ossv1,” run the following command:

```
snapvault start –S ossv1:c:\ /vol/backups/ossv1
```

The **`snapvault start`** command in this example establishes the relationship between the C:\ drive on the OSSV host and the qtree, “ossv1,” in the volume “backups.”

The qtree, “ossv1,” is created by the **`snapvault start`** command and cannot be created manually.

CREATING SCHEDULES

OSSV schedules are based on the volume. Therefore all OSSV relationships that share the same destination volume operate on the same schedule. When scheduling OSSV using the CLI, the **`snapvault snap sched`** command is run on the NetApp secondary system. For example:

```
snapvault snap sched –x backups ossv_daily 30@mon-fri@23
```

In this example, a backup schedule is created for the volume called “backups.” According to this schedule OSSV will run at 11 p.m. Monday through Friday. After all backups initiated by the schedule complete, a Snapshot copy called “ossv_daily.0” is created. The **`snap list`** command can be used to view the Snapshot copies for the volume. The most recent Snapshot copy will have a “.0” suffix.

If a backup takes significantly longer to complete than the other jobs in the schedule, a **`snapvault status`** will show the other jobs sitting in a “quiescing” state until the backup finishes. This is because all of the backups are captured by a single Snapshot copy for the volume. For this reason, it is best to organize relationships and volumes in ways to avoid conflict.

In addition, NetApp secondary systems can handle different numbers of concurrent transfers depending on the model and Data ONTAP version. In large environments, it is necessary to keep this in mind and set up schedules accordingly. To determine the maximum number of concurrent transfers for a particular platform, refer to the "Data Protection Online Backup and Recovery Guide" on the NOW site.

MANUAL BACKUPS

The OSSV schedule will take care of performing incremental backups. However, manual incremental backups can be run from the NetApp secondary system if needed. In order to start a backup manually, use the **`snapvault update`** and **`snapvault snap create`** commands. For example:

```
snapvault update /vol/backups/ossv1
```

This will initiate an incremental transfer for the relationship associated with "/vol/backups/ossv1." The **`snapvault status`** command can be run on the NetApp secondary to identify destination volume and qtree if needed. The **`snapvault status`** command is also used to determine the state of the backup. When the backup completes, the relationship returns to "Idle" status.

Note: An OSSV schedule will initiate backups for all relationships in a volume, while a manual backup will only initiate a backup for a specific relationship in a volume.

After the transfer is complete, the **`snapvault snap create`** command is used to secure the backup data in a Snapshot copy. In order for the **`snapvault snap create`** command to work, a schedule entry for the relationship must exist. If there is no need to have backups run from a schedule, a minimal schedule entry can be created. For example, the following command will create a schedule entry that does not invoke a transfer from the OSSV host, nor will it create any snapshots. However, it will manage the retention of snapshots called "ossv_daily" that are manually created.

```
snapvault snap sched backups ossv_daily 30@-
```

To manually create a snapshot after the incremental transfer has completed, run the **`snapvault snap create`** command. For example:

```
snapvault snap create backups ossv_daily
```

In this example, a Snapshot copy called "ossv_daily.0" is created for the volume "backups."

MONITORING STATUS

OSSV status can be monitored from the NetApp secondary and from the OSSV host. The **`snapvault status`** command is available on both. When run without any options, the **`snapvault status`** command lists each relationship's status and lag time. Lag time is the amount of time that has passed since the beginning of the last successful backup.

The **`snapvault status`** –l command is used to gather more information about the relationships. It's important to know that this output it slightly different on the NetApp secondary and the OSSV host. Therefore it is a good idea to run the command in both places for complete information.

The **`<install_dir>\bin\snapvault status –l`** command run from the OSSV host displays the following fields:

- **`Source`**
- **`Destination`**
- **`Status`**
- **`State`**
- **`Lag`**

- **`Mirror Timestamp`**
- **`Base Snapshot`**
- **`Current Transfer Type`**
- **`Contents`**
- **`Last Transfer Type`**
- **`Last Transfer From`**
- **`Last Transfer Size`** (includes compression ratio if enabled)
- **`Last Transfer Duration`**
- **`Total files to transfer`**
- **`Total files transferred`**
- **`Current File Size`**
- **`Current File Progress`**
- **`Current File Name`**
- **`Transfer Error ID`**
- **`Transfer Error Message`**

The **`snapvault status –l`** output from the NetApp secondary system displays the following fields:

- **`Source`**
- **`Destination`**
- **`Status`**
- **`Progress`**
- **`Compression Ratio`** (if enabled)
- **`State`**
- **`Lag`**
- **`Mirror Timestamp`**
- **`Base Snapshot`**
- **`Current Transfer Type`**
- **`Current Transfer Error`**
- **`Contents`**
- **`Last Transfer Type`**
- **`Last Transfer Size`**
- **`Last Transfer Duration`**
- **`Last Transfer From`**

MONITORING SPACE ON THE OSSV HOST

In order for OSSV to perform backups, there must be sufficient space available on the OSSV host to store the OSSV database. The Free Space Estimator utility can be used to confirm that sufficient space exists. By default, the Free Space Estimator is run prior to each backup and can be disabled with the OSSV Configurator utility. It can also be run manually with the **`<install_dir>\bin\svestimator`** command on the OSSV host. By default, the Free Space Estimator will not cause a backup to fail. However, this behavior can be changed if needed by modifying the **`<install_dir>\config\snapvault.cfg`** file.

An example of running the Free Space Estimator manually follows.

```
C:\Program Files\netapp\snapvault\bin>svestimator c:\

Scanning system volumes...
Volume 'C:\' type Normal NTFS Free Space 82%
Volume 'D:\' type Normal OFFLINE Free Space 0%
Volume 'E:\' type Normal OFFLINE Free Space 0%
Volume 'M:\' type Normal OFFLINE Free Space 0%
Volume 'N:\' type Normal OFFLINE Free Space 0%
Volume 'Q:\' type Normal OFFLINE Free Space 0%



Examining 'c:\'...

Estimated space requirements so far:
Database: 53.93 MB
Temp: 94.98 MB

Analyzing space requirements...
Estimator has found sufficient space for backup
```

PERFORMING RESTORES

There are different ways in which restore operations can be performed. One way is to use the `snapvault restore` command on the OSSV host. The other way is to copy data directly from a Snapshot copy on the NetApp secondary system using a CIFS or an NFS connection.

It is generally best to restore data to an alternate (or temporary) directory on the OSSV host, check the restored data, and then move the data into its proper place. The following is an example of restoring all data from the most recent backup to an alternate (or temporary) directory on the OSSV host using the **`snapvault restore`** command. After the restore completes, the relationship created during the restore can be released.

```
<install_dir>\bin\snapvault restore –S fas1:/vol/backups/ossv1 C:\temp\restore

<install_dir>\bin\snapvault release C:\temp\restore fas1:/vol/backups/ossv1
```

The **`snapvault restore`** command can also be used to restore a single file. For example:

```
<install_dir>\bin\snapvault restore –S fas1:/vol/backups/ossv1/Report.doc
C:\temp\restore\Report.doc

<install_dir>\bin\snapvault release C:\temp\restore\Report.doc fas1:/vol/backups/ossv1
```

The **`snapvault restore`** command can be used to restore to the original location. In order to restore data to the original location, the backup relationship must first be released. Additionally, in order to maintain the ability to continue doing block-level incremental backups after the restore, the “restart/resync” option must be enabled on the OSSV host. This option can be enabled using the OSSV Configurator utility on the OSSV host.

To restore data to the original location, release the backup relationship and then perform the restore. For example:

```
<install_dir>\bin\snapvault release C:\data fas1:/vol/backups/ossv1

<install_dir>\bin\snapvault restore –S fas1:/vol/backups/ossv1 C:\data
```

After the restore is complete, the original backup relationship can be reestablished on the NetApp secondary. This gives the relationship the ability to continue doing incremental backups.

```
snapvault start -r -S ossv1:\c:\data /vol/backups/ossv1
```

Note: To enable compression during restore, refer to the compression section in this document.

The OSSV database for a particular relationship can also be restored if needed. Restore the `.OSSV_DATABASE_BACKUP` file from the root of the destination qtree to a file in any temporary location on the OSSV host. For example:

```
<install_dir>\bin\snapvault restore -S fas1:/vol/backups/ossv1/.OSSV_DATABASE_BACKUP
C:\temp\database
```

The OSSV service running on the host will automatically recognize this as a database restore and complete the operations to put the recovered database in place.

The other way in which restores can be accomplished is by copying data directly from a Snapshot copy on the NetApp secondary. From a CIFS or an NFS client, mount the `~snapshot` (or `.snapshot` on UNIX and Linux) directory for the volume. After mounting the `~snapshot` directory, the list of Snapshot copies will appear as directories. Browse the appropriate Snapshot copy and copy the data as required.

Note: OSSV does not have a bare metal recovery option. To restore an entire system, the operating system as well as OSSV would first need to be installed. Then the C:\ drive could be restored, followed by the System State (on Windows hosts).

MODIFYING RELATIONSHIP OPTIONS

The `snapvault modify` command can be used to modify relationships. For example, to change the “tries” limit run the following command on the NetApp secondary.

```
snapvault modify -t 3 -S ossv1:c:\ fas1:/vol/backups/ossv1
```

To enable compression for a particular relationship, run the following command:

```
snapvault modify -o compression=on -S ossv1:c:\ fas1:/vol/backups/ossv1
```

DELETING RELATIONSHIPS

OSSV relationships can be deleted from the NetApp secondary system using the `snapvault stop` command. Issuing this command permanently destroys the relationship and the qtree to which the data was backed up. It does not, however, remove the Snapshot copies that contain the historical backup data unless the volume is manually destroyed as well.

For example, to delete a relationship run the following command on the NetApp secondary:

```
snapvault stop /vol/backups/ossv1
```

This removes this relationship and destroys the qtree, “ossv1.” The following warning is displayed after running the command, requiring confirmation.

```
Stopping /vol/backups/ossv1 is permanent.
The secondary qtree will be deleted.
Further incremental updates will be impossible.
Data already stored in snapshots will not be deleted.
This may take a long time to complete.
Are you sure you want to do this? Y
```

4.2 PROTECTION MANAGER

OSSV 3.0 supports Operations Manager 3.8 and higher and its equivalent version of Protection Manager. Using Protection Manager to create and maintain OSSV relationships is preferred due to its policy-based management style.

OSSV hosts listen for communications from Protection Manager on port 10000 by default. This is set during the installation of the OSSV software and can be changed by modifying the “NDMP Listen Port” field in the OSSV Configurator utility on the host. Any firewalls that are in the network path will need to have this port open.

The basic steps required to create an OSSV relationship in Protection Manager are as follows:

- Add an OSSV host using the “Add OSSV Host Wizard.”
- Create a resource pool using the “Create Resource Pool Wizard.”
- Copy the “Remote backups only” protection policy and modify schedules.
- Create a dataset using the “Add Dataset Wizard.”

Details about each of these steps are discussed below.

ADD OSSV HOST

The “Add OSSV Host Wizard” is used to add OSSV hosts into Protection Manager. Complete the wizard by supplying the IP address or hostname of the OSSV host and the NDMP credentials established during the OSSV software installation.

Note: Hostname resolution is required even when using an IP address to add the host.

While adding the OSSV host, the wizard will ask about a NetApp Host Agent. The NetApp Host Agent is no longer required or bundled with OSSV. Unless the NetApp Host Agent has been specifically installed on the OSSV host, select “Continue without a NetApp Host Agent.”

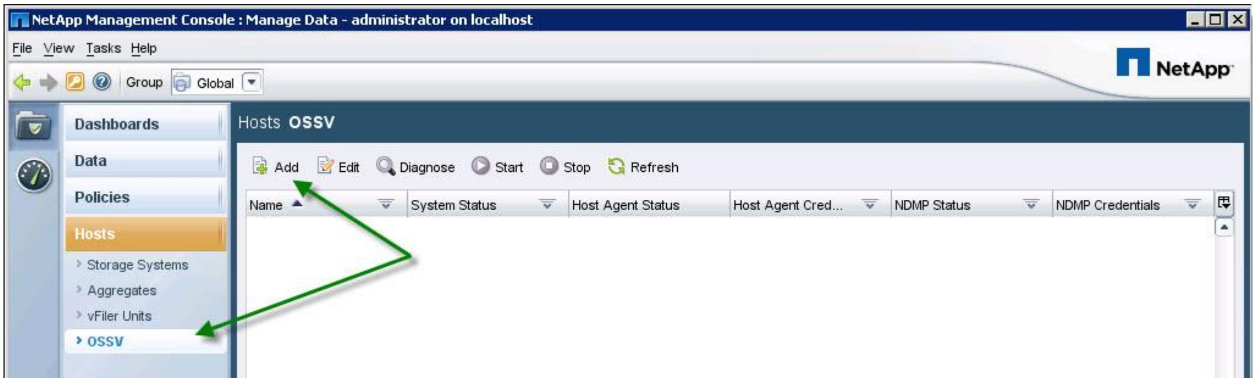

Figure 3) Add OSSV hosts using the “Add OSSV Host Wizard.”

CREATE A RESOURCE POOL

Using resource pools allows Protection Manager to provision the secondary volumes needed for backup. To create a resource pool, use the “Add Resource Pool Wizard” and select one or more aggregates.

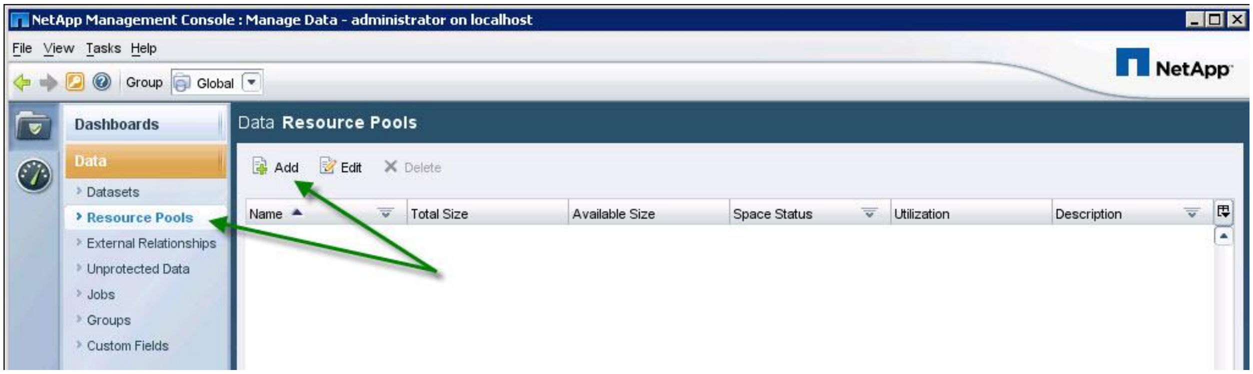

Name	Total Size	Available Size	Space Status	Utilization	Description

Figure 4) Create a resource pool

CREATE A PROTECTION POLICY

Protection policies control how data is protected as well as retained. Protection policies also reference schedule policies to control how often data protection operations occur. When working with protection policies it is good practice to create a copy from the templates and edit the new version. To configure a basic protection policy for OSSV, create a copy of the “Remote backups only” policy.

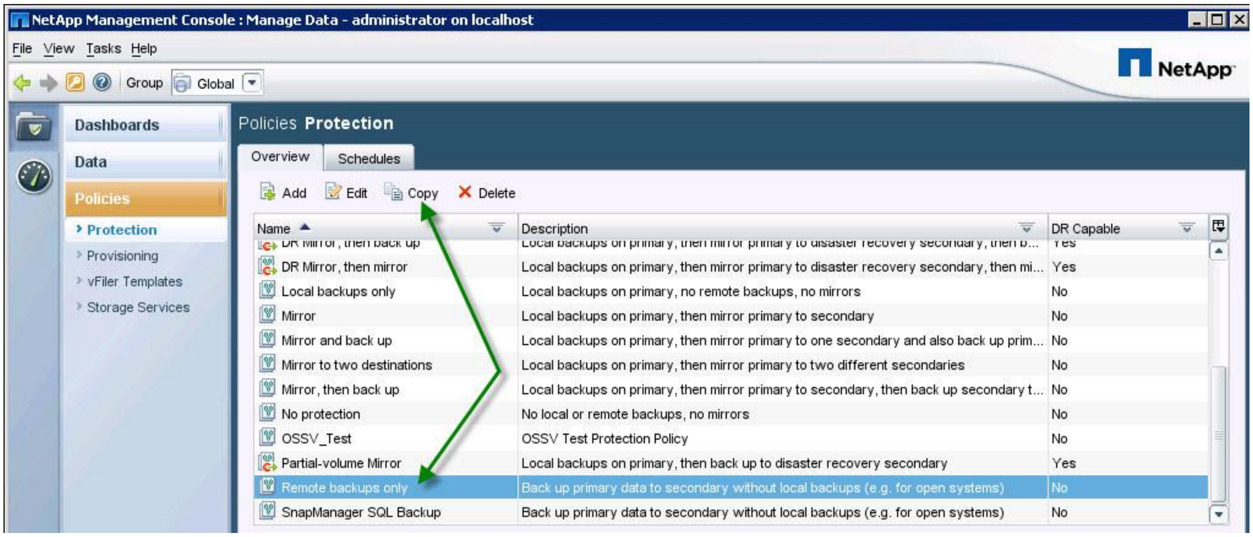

Name	Description	DR Capable
DR Mirror, then back up	Local backups on primary, then mirror primary to disaster recovery secondary, then b...	Yes
DR Mirror, then mirror	Local backups on primary, then mirror primary to disaster recovery secondary, then mi...	Yes
Local backups only	Local backups on primary, no remote backups, no mirrors	No
Mirror	Local backups on primary, then mirror primary to secondary	No
Mirror and back up	Local backups on primary, then mirror primary to one secondary and also back up prim...	No
Mirror to two destinations	Local backups on primary, then mirror primary to two different secondaries	No
Mirror, then back up	Local backups on primary, then mirror primary to secondary, then back up secondary t...	No
No protection	No local or remote backups, no mirrors	No
OSSV_Test	OSSV Test Protection Policy	No
Partial-volume Mirror	Local backups on primary, then back up to disaster recovery secondary	Yes
Remote backups only	Back up primary data to secondary without local backups (e.g. for open systems)	No
SnapManager SQL Backup	Back up primary data to secondary without local backups (e.g. for open systems)	No

Figure 5) Create a copy of the “Remote backup only” policy.

After creating the copy, edit the new policy and change the name as needed. “Nodes and Connections” contains settings for retention and scheduling. For OSSV there will be no schedule or retention set for “Primary data.”

The backup schedule policy can be selected under “Primary data to Backup.” In addition, throttle policies can also be applied if needed. OSSV retention settings are defined under “Backup.”

Note: Throttle policies used in Protection Manager are not the same as dynamic client-side throttling. In the event that both throttle features are used, the lower of the settings will be used.

If the schedule policy templates need to be customized, they can be copied and modified as well.

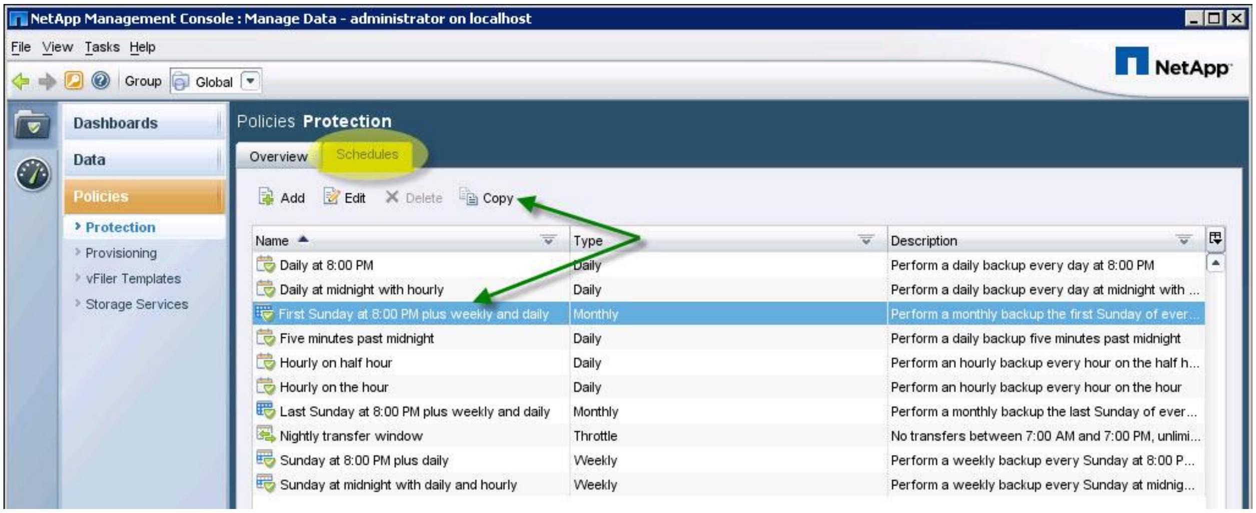

Name	Type	Description
Daily at 8:00 PM	Daily	Perform a daily backup every day at 8:00 PM
Daily at midnight with hourly	Daily	Perform a daily backup every day at midnight with ...
First Sunday at 8:00 PM plus weekly and daily	Monthly	Perform a monthly backup the first Sunday of ever...
Five minutes past midnight	Daily	Perform a daily backup five minutes past midnight
Hourly on half hour	Daily	Perform an hourly backup every hour on the half h...
Hourly on the hour	Daily	Perform an hourly backup every hour on the hour
Last Sunday at 8:00 PM plus weekly and daily	Monthly	Perform a monthly backup the last Sunday of ever...
Nightly transfer window	Throttle	No transfers between 7:00 AM and 7:00 PM, unlimi...
Sunday at 8:00 PM plus daily	Weekly	Perform a weekly backup every Sunday at 8:00 P...
Sunday at midnight with daily and hourly	Weekly	Perform a weekly backup every Sunday at midnig...

Figure 6) Create a copy of the schedule template.

CREATE A DATASET

The dataset pulls everything together into a manageable object. The dataset includes the source data to be protected, the resource pool that will store the backups, and the protection policy. To create a dataset use the “Add Dataset Wizard.”

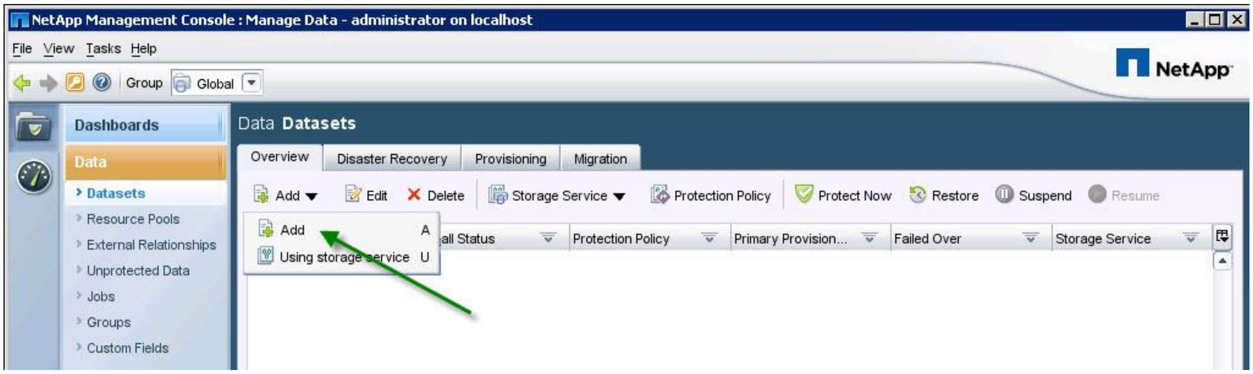

Figure 7) Create a dataset using the “Add Dataset Wizard.”

Give the dataset a meaningful name and click Next. On the next screen choose “Select resources manually” and click Next. Choose the source data intended to be managed by this dataset and click Next.

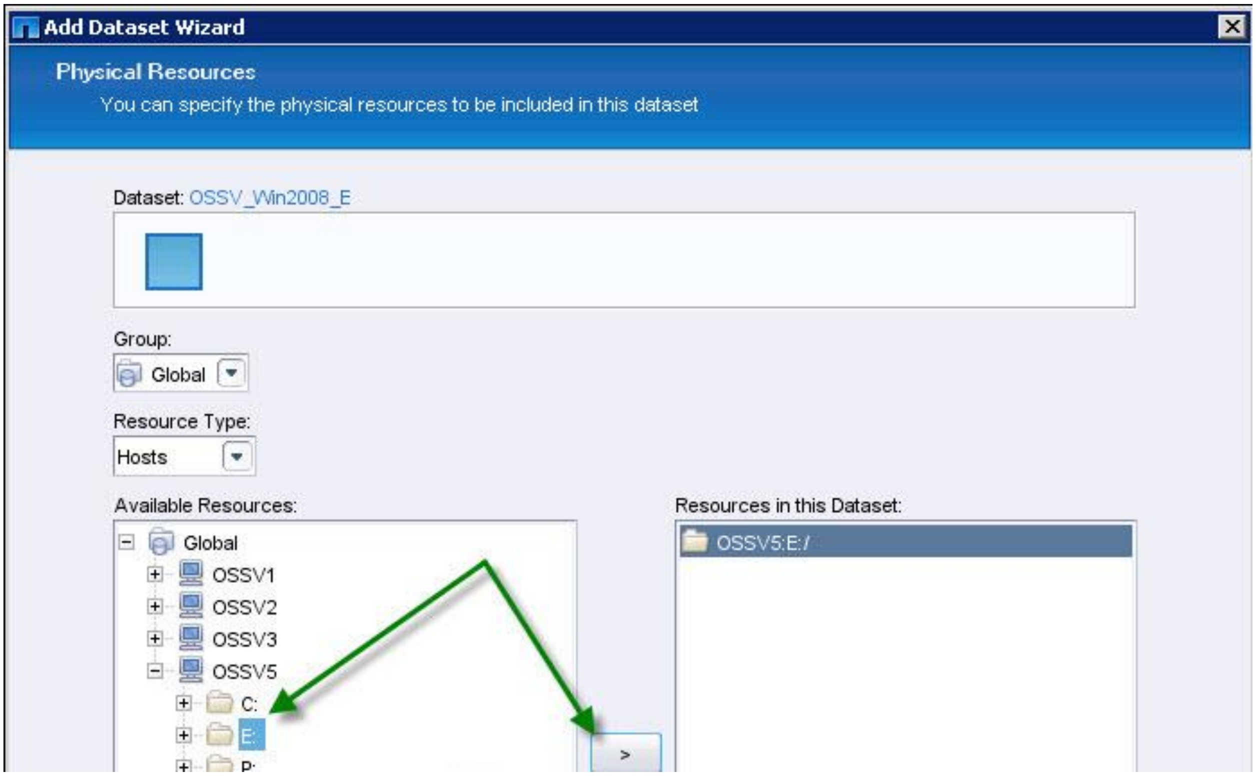
Add Dataset Wizard

Physical Resources

You can specify the physical resources to be included in this dataset

Dataset: OSSV_Win2008_E

Group: Global

Resource Type: Hosts

Available Resources:

Resources in this Dataset: OSSV5:E:/

Figure 8) Manually select the source data to be managed by the dataset.

After selecting the source data, continue through and finish the wizard. At this point the dataset will be created, but it will not contain a protection policy. Highlight the dataset and click the Protection Policy button to launch the “Protection Policy Wizard.”

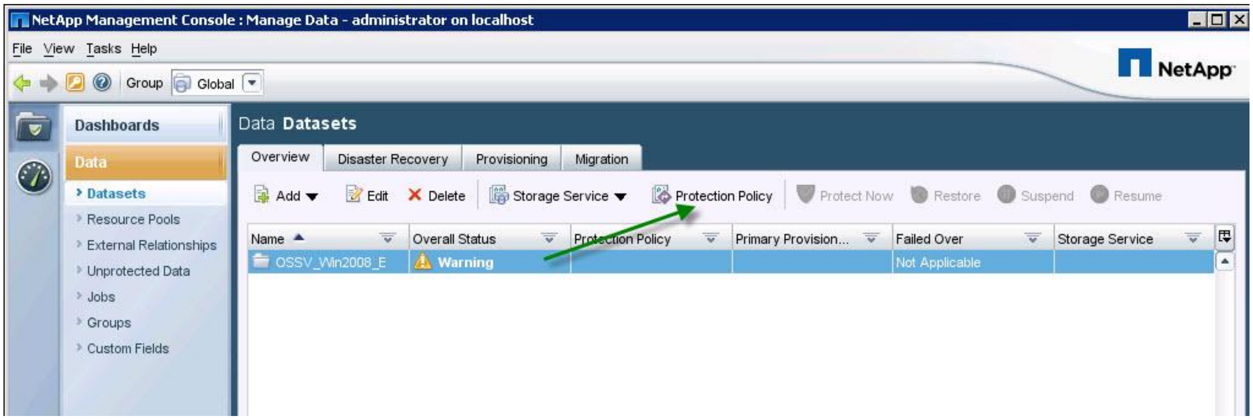
NetApp Management Console : Manage Data - administrator on localhost

Group: Global

Data Datasets

Name	Overall Status	Protection Policy	Primary Provision...	Failed Over	Storage Service
OSSV_Win2008_E	Warning			Not Applicable	

Figure 9) Launch the “Protection Policy Wizard.”

Using the wizard, select the protection policy previously copied and modified and click Next. By selecting “Provision and attach resources using a policy” on the next screen, the resource pool previously created can be used by this protection policy. Click Next.

On the next screen, choose the resource policy previously created. Continue through and finish the wizard. Protection Manager will create the secondary volume and initiate the baseline transfer.

5 PROTECTING MICROSOFT SQL SERVER

OSSV 3.0 introduces the ability to protect Microsoft SQL Server databases that are hosted on non NetApp primary storage. The following versions are supported:

- SQL Server 2005
- SQL Server 2008
- Windows 2003 (32-bit and 64-bit including R2)
- Windows 2008 (32-bit and 64-bit including R2)

As with file backups, OSSV uses block-level incremental backups when protecting SQL Server. However, the OSSV filter driver is enabled by default for SQL Server data in order to decrease backup times. When protecting SQL Server, OSSV leverages VSS to quiesce the SQL Server database for consistency.

The following types of SQL Server backups can be performed with OSSV:

- Database backup
- Transaction log backup
- Local transaction log backup

OSSV protection for SQL Server can be set up and managed using the CLI or Protection Manager. There are no additional license requirements specific to protecting SQL Server.

DATABASE BACKUPS

Database backups protect the entire database and can be scheduled to run as often as every hour. When using the CLI to configure database backups, the **`svapp`** command is useful for listing the instances and database paths available for backup.

Example output from the **`svapp -list mssql -verbose`** command:

```
Database\Tlog       Backup Path                 Protected     FS Paths

-------------       -----------                 ---------     --------

db1                 app:mssql:MSSQLSERVER:db1         No      E:\SQL\data\db1_Data.MDF

                                                              D:\SQL\data\db1_Log.LDF

                                                              E:\SQL\data\db1_Data1.NDF

db1:Tlog            app:mssql:MSSQLSERVER:db1:Tlog    No
      E:\OSSVDB\apps\mssql\backedupTlogs\MSSQLSERVER\db1\db1.trn
```

Relationships for database backups are configured using the following format in the **`snapvault`** command:

```
app:mssql:<instance_name>:<db_name>
```

For example:

```
snapvault start -S ossv1:app:mssql: MSSQLSERVER:db1 fas1:/vol/SQLbackup/MSSQLSERVER_db1
```

TRANSACTION LOG BACKUPS

Transaction log backups can be scheduled to run as often as every hour. When using the CLI to configure transaction log backups, the **svapp** command is useful for listing the instances and transaction log paths available for backup. Relationships for database backups are configured using the following format in the **snapvault** command:

```
app:mssql:<instance_name>:<db_name>:Tlog
```

For example:

```
snapvault start -S ossv1:app:mssql: MSSQLSERVER:db1:Tlog
fas1:/vol/SQLTlogbackup/MSSQLSERVER_db1_Tlog
```

LOCAL TRANSACTION LOG BACKUPS

In order to perform backups more frequently than every hour, OSSV can initiate transaction log backups that are stored locally on the SQL Server host. These local transaction log backups can run as often as every five minutes. To configure local transaction log backups, populate the following files:

- **<install_dir>\config\ossv-mssql-local-Tlog-DBs.cfg**
- **<install_dir>\config\ossv_mssql.cfg**

The **ossv-mssql-local-tlog-DBs.cfg** file contains a list of databases and/or instances for which local transaction log backups will be performed. An example entry is:

```
MSSQLSERVER:db1
```

The backup interval can be set in the **ossv_mssql.cfg** file and must be between 5 and 55 minutes. A value of 0 disables the feature. For example, to perform local transaction log backups every 10 minutes, set the following parameter:

```
[MSSQL:TLog Backup Interval]
```

```
Value=10
```

In clustered configurations, OSSV stores the local transaction log backup file on the same volume as the OSSV database for the relationship. The **svapp** command will show the directory. For example:

```
E:\ossvdb\apps\mssql\backedupTlogs\MSSQLSERVER\db1
```

In nonclustered environments, OSSV stores the local transaction log backup file in the default database location. For example:

```
C:\Program Files\NetApp\snapvault\db\apps\mssql\backedupTlogs\MSSQLSERVER\db1
```

The location for storing these local transaction logs can be changed by adding or modifying the following parameter in the **ossv_mssql.cfg** file:

```
[MSSQL:TLog Backup Directory]
```

Note: The location of the local transaction logs cannot be changed when operating in a clustered environment.

TRUNCATING TRANSACTION LOGS

By default, OSSV will truncate the transaction logs after a database or transaction log backup. To change this behavior, modify the following parameter in the **ossv_mssql.cfg** file:

```
[MSSQL:TLog Truncate]
```

RESTORE BEHAVIOR

There are several ways to modify the behavior of OSSV and how it accomplishes restores of SQL Server databases. By default, after successfully restoring a database, OSSV leaves the database in a “restoring” state. This allows a subsequent transaction log restore. After the transaction log restore completes, OSSV brings the database online. This behavior can be modified by editing the **`ossv_mssql.cfg`** file.

In addition, this file can be modified so that OSSV also restores the local transaction log.

The following parameter is set to FALSE by default, allowing a subsequent transaction log backup. When set to TRUE, OSSV will bring the database online after a database restore and will not wait for a transaction log restore.

```
[MSSQL:Recover After DB Restore]
```

In order for OSSV to bring the database online after a transaction log restore (including a local transaction log restore), the following parameter should be set to TRUE. The default setting is TRUE.

```
[MSSQL:Recover After TLOG Restore]
```

OSSV can restore the local transaction log after the database and transaction log restores have been completed. To modify this behavior, change the following parameter to TRUE. The default setting is FALSE.

```
[MSSQL:Restore Local TLog]
```

If the database administrator needs to roll through the logs to a particular point in time, the following settings are ideal:

```
[MSSQL:Recover After DB Restore]
Value=FALSE

[MSSQL:Recover After TLOG Restore]
Value=FALSE

[MSSQL:Restore Local TLog]
Value=FALSE
```

With this setup, the administrator can restore the database and the transaction logs. OSSV will not bring the database online. The administrator can then roll through local transaction log files. The backup file will have a name in the form of **`<db>-local.trn`**. This **`svapp`** command will show the directory for the transaction log backups.

DATABASE RESTORES

When using the CLI to perform a database restore, the **`snapvault status`** command can be used to identify the path for the relationship. For example:

```
ossv1:app:mssql:MSSQLSERVER:db1b        fas1:/vol/SQLbackup/MSSQLSERVER_db1

ossv1:app:mssql:MSSQLSERVER:db1:Tlog    fas1:/vol/SQLTlogbackup/MSSQLSERVER_db1_Tlog
```

To begin a restore of the most recent database backup, run the **`snapvault restore`** command from the OSSV host. For example:

```
<install_dir>\bin\snapvault restore –S fas1:/vol/SQLbackup/MSSQLSERVER_db1
app:mssql:MSSQLSERVER:db1
```

If the **`MSSQL:Recover After DB Restore`** parameter is false, a transaction log restore from the most recent backup can also be restored. For example:

```
<install_dir>\bin\snapvault restore –S fas1:/vol/SQLbackup/MSSQLSERVER_db1_Tlog
app:mssql:MSSQLSERVER:db1:Tlog
```

Note: The **`–s flag`** can be used in the **`snapvault restore`** command to restore from a specific Snapshot copy.

The local transaction log will be restored automatically depending on the **`MSSQL:Restore Local TLog`** parameter.

PROTECTION MANAGER

NetApp recommends that backup and recovery for SQL Server be managed with Protection Manager. When using Protection Manager, instances and databases are displayed as subdirectories under the “app:mssql” folder on the OSSV host.

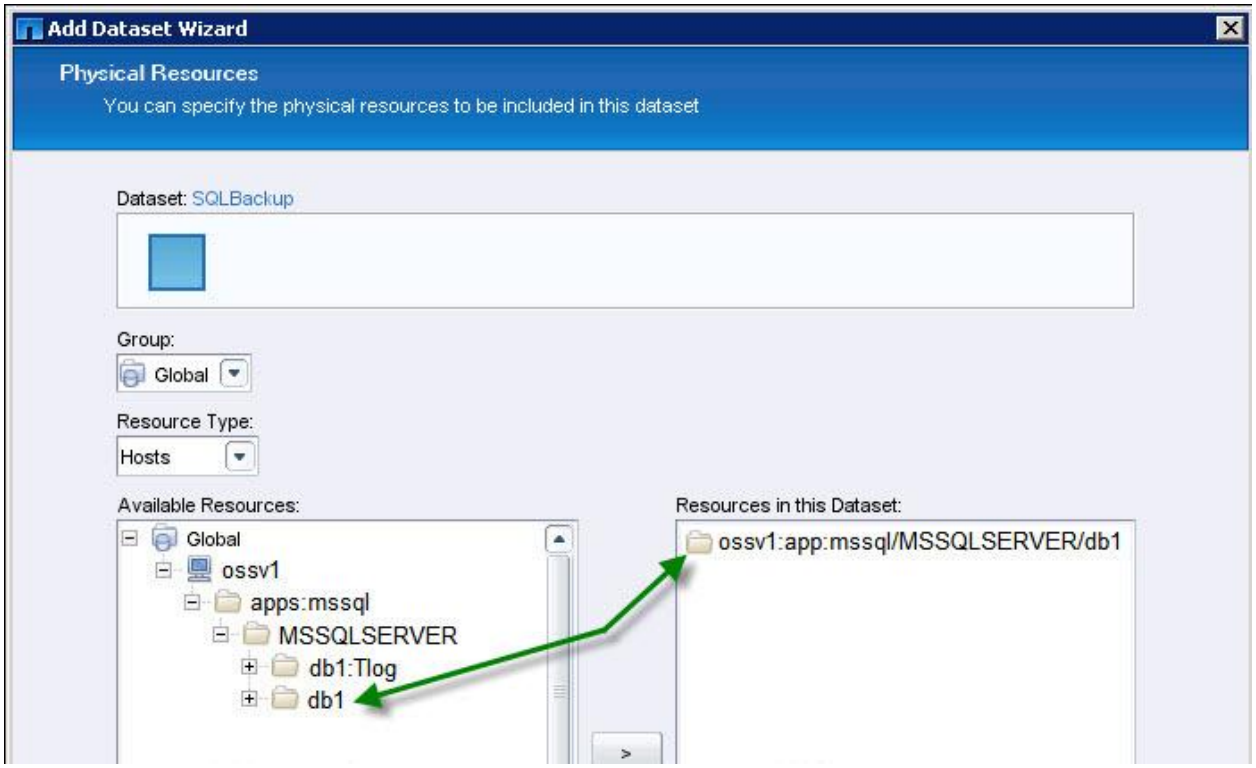
Dataset: SQLBackup

Group: Global

Resource Type: Hosts

Add Dataset Wizard
Physical Resources
You can specify the physical resources to be included in this dataset
Available Resources:
Global
ossv1
apps:mssql
MSSQLSERVER
db1:Tlog
db1
>
Resources in this Dataset:
ossv1:app:mssql/MSSQLSERVER/db1

Figure 10) Protection Manager lists instances and databases as folders.

Since datasets contain their own schedules, NetApp recommends creating separate datasets for database backups and transaction log backups with schedules that do not overlap. This will prevent database backups and transaction log backups from starting at the same time.

Note: For SQL Server hosts, do not select the entire client, the “apps:mssql” object, or the instance as the resource for the dataset. Select only the individual databases and transaction logs as needed.

6 MICROSOFT CLUSTER SERVER

OSSV 3.0 can be installed on 2-node Microsoft Cluster Server (MSCS) clusters. The **`svcluster utility`** on the OSSV host is used to enable cluster support within the OSSV software. To enable cluster support, run the following command on both OSSV hosts in the cluster.

```
<install_dir>\cluster\mscs\svcluster enable
```

The **`svcluster`** utility does two things:

1. It creates a resource type in MSCS called "OSSVResourceType."
2. It changes the behavior of OSSV such that the database location for new relationships resides in the "ossvdb" directory at the root of the volume for the relationship.

The OSSVResourceType needs to be added as a resource in each group protected by OSSV and be made dependent on the disk resources. This resource restarts the OSSV service on the node that initiates an "offline" or "move." This is done to stabilize the OSSV database during failover. As a result, any backups running on that host will fail and restart from the checkpoint. However, backups for the "moved" group will not be able to restart from the checkpoint.

Both OSSV hosts in the cluster should be set up to have the same general configuration settings. This includes settings such as the "QSM Access List," compression, throttling, and any nondefault time-out values that have been set. In addition, disk resources should use the same drive letter for each node.

When configuring OSSV relationships that are part of an MSCS resource group, use the failover IP address or hostname assigned to that resource group. OSSV will be able to perform backups from either node depending on the location of the resource group. Relationships for local data that is not part of the cluster should be configured using the physical hostname of the OSSV host.

Note: If an existing OSSV host with existing relationships is configured for cluster support, all backups (except System State) will need to be reestablished along with a new baseline transfer.

Note: OSSV can run on a cluster node without having cluster support enabled as long as none of the data being protected is controlled by the cluster.

PROTECTION MANAGER

Protection Manager can manage OSSV hosts that have been configured for MSCS support. However, here are a few steps that need to be done in order to add all of the IP addresses assigned to the resource groups into Protection Manager.

1. Disable Host Agent Discovery on the Operations Manager server:

   ```
   dfm option set discoverAgents=no
   ```

2. Disable the NDMP Host Discovery option on the Operations Manager server:

   ```
   dfbm option set discoverNdmp=no
   ```

3. Add both OSSV hosts (physical hostnames) into Protection Manager as normal.
4. For each resource group hostname (failover IP) to be added into Protection Manager:
   a. Open the OSSV Configurator on the respective OSSV host.
   b. Change the "NDMP Host Name" field to match the hostname assigned to the resource group.
   c. Change the last four or five characters of the "NDMP Host Id" field such that it is unique.
   d. Add the hostname into Protection Manager.
5. Enable the Host Agent Discovery on the Operations Manager server:

   ```
   dfm option set discoverAgents=yes
   ```

6. Enable the NDMP Host Discovery option on the Operations Manager server:

```
dfbm option set discoverNdmp=yes
```

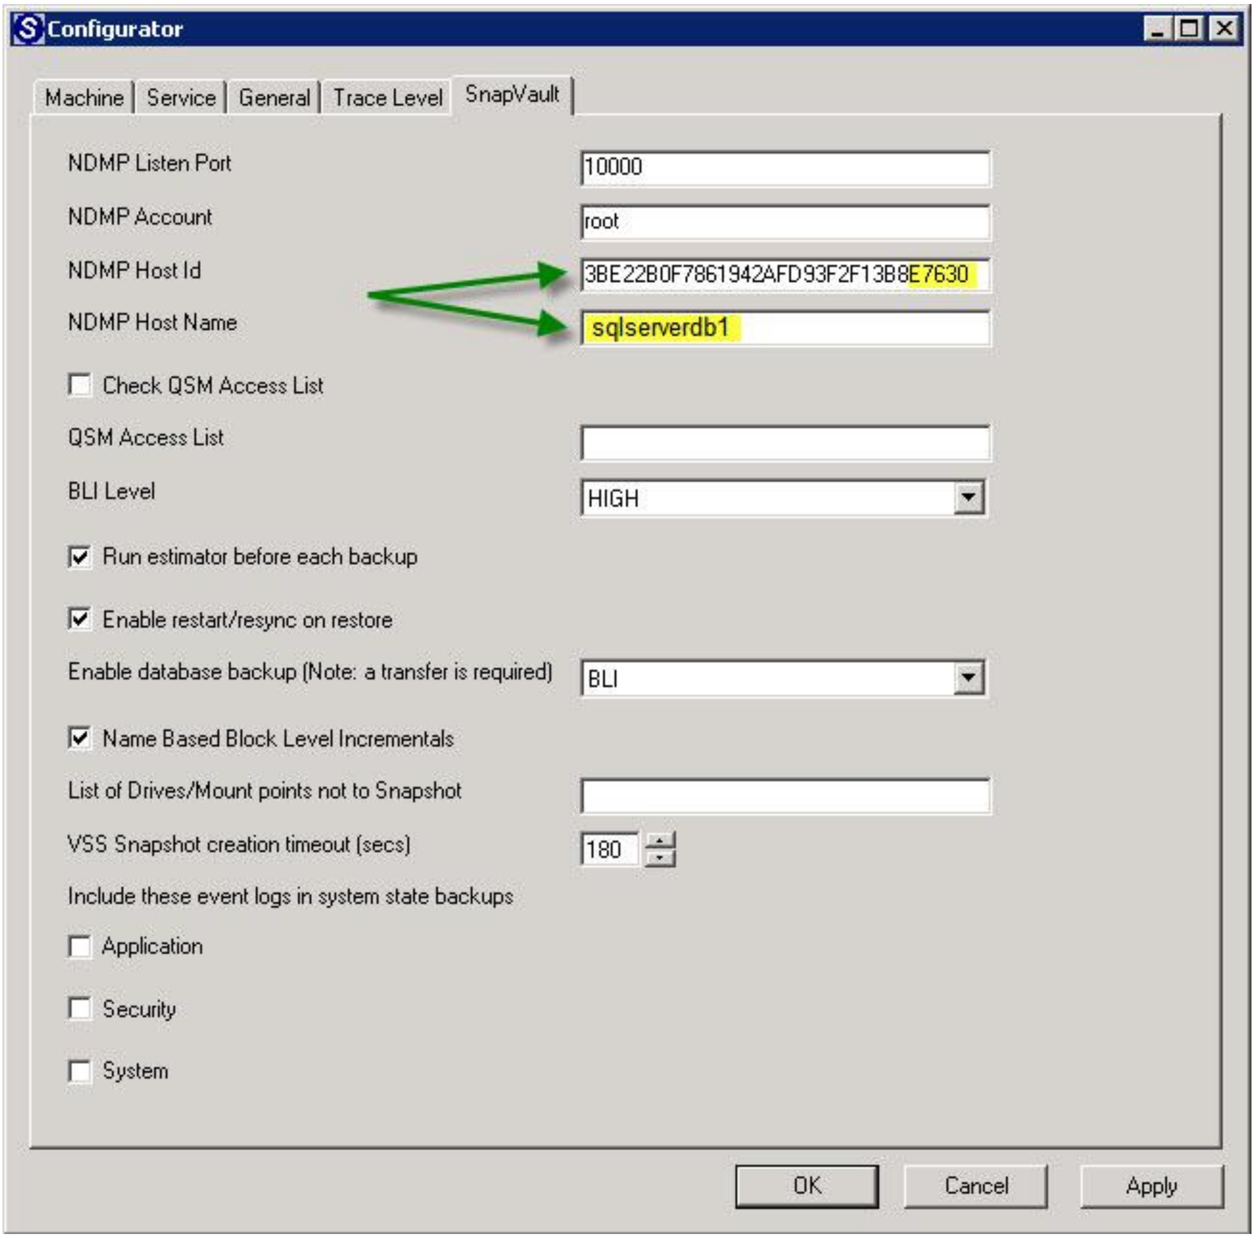

NDMP Listen Port: 10000
NDMP Account: root
NDMP Host Id: 3BE22B0F7861942AFD93F2F13B8E7630
NDMP Host Name: sqlserverdb1
☐ Check QSM Access List
QSM Access List:
BLI Level: HIGH
☑ Run estimator before each backup
☑ Enable restart/resync on restore
Enable database backup (Note: a transfer is required): BLI
☑ Name Based Block Level Incrementals
List of Drives/Mount points not to Snapshot:
VSS Snapshot creation timeout (secs): 180
Include these event logs in system state backups
☐ Application
☐ Security
☐ System

Figure 11) Change the “NDMP Host ID” and “NDMP Host Name” using the OSSV Configurator.

7 BEST PRACTICES

There are a number of best practices mentioned throughout this document. This section lists them together.

KNOW THE DATA

It is best to understand the data that is protected by OSSV:

- Total amount of source data
- File sizes
- Change rates

The OSSV Profiler is a tool that can be used on Windows hosts to simulate OSSV backups prior to installation. This tool can be used to gather statistics about the data and the impact on the OSSV host. The OSSV Profiler can be found in the Utility ToolChest on the NOW site.

Understanding the data will help when architecting the OSSV solution. It is best to group "like" data together. This means that data similar in composition, size, and change rate should be backed up together to common destination volumes. If one backup takes significantly longer to complete than the other backups on the volume, it will cause all of the other backups to wait before they can complete.

SIZE SECONDARY STORAGE APPROPRIATELY

It is important to size secondary volumes appropriately. When using Protection Manager this is not a concern, because Protection Manager provisions secondary volumes automatically. However, when creating volumes manually it is important to plan ahead. Consider the following when manually provisioning secondary volumes.

- The source data (size, change rates, etc.)
- The number of relationships that will share the secondary volume
- Backup schedules
- Retention

It is common for many relationships to share the same destination volume. However, the volume should be sized such that all of the baselines and incremental backups have ample room given the retention requirements.

INTERNATIONAL LANGUAGE SETTINGS

In situations in which OSSV data contains non-ASCII characters, NetApp recommends enabling UTF-8 on the secondary volume. In some cases the NetApp secondary system may be in one locale while the OSSV hosts are in different locales. As long as UTF-8 is enabled it will be able to back up the non-ASCII files.

New volumes inherit the language setting of the root volume. If UTF-8 is enabled on the root volume it will automatically be enabled on new volumes.

When using Protection Manager, different behaviors exist for provisioning secondary volumes, depending on the approach. When not using a Provisioning Policy, the secondary volume will inherit the language setting from the root volume. If the root volume has UTF-8 enabled, then the new volume will also have UTF-8 enabled. The create_ucode and convert_ucode settings for the new volume are not enabled.

When a Provisioning Policy is used to provision the secondary volume, the new volume will inherit the language setting from the root volume. In addition, UTF-8 will be enabled as well as create_ucode and convert_ucode settings.

FAN IN WHEN USING DEDUPLICATION

When using NetApp deduplication on the NetApp secondary volume, more deduplication potential will be gained by sending data from multiple sources to that same volume. For example, when protecting data on the C: drive of several Windows 2003 servers, group those relationships into the same volume for better space savings potential.

USE PROTECTION MANAGER TO MAXIMIZE RETENTION

Unlike the scheduling functionality within Data ONTAP (using the **`snapvault snap sched`** command), Protection Manager includes the ability to create monthly schedules. Using monthly schedules can help lengthen the amount of retention available for backup data. OSSV backup data can be kept for years if needed. As an example, 7 years of backup data could be retained as follows:

- Daily backups for 90 days
- Weekly backups for 2 years
- Monthly backups for 7 years

Using a schedule similar to the following, this retention model would consume around 242 Snapshot copies. A maximum of 250 Snapshot copies is allowed per volume.

- Daily backups, Monday—Saturday at 11 p.m.
- Weekly backups, Sundays at 11 p.m.
- Monthly backups, first Sunday of the month at 11 p.m.

UNDERSTAND STREAM LIMITS

NetApp secondary systems can handle different numbers of concurrent transfers depending on the model and Data ONTAP version. In large environments, it is necessary to keep this in mind and set up schedules accordingly. If backups are scheduled such that they exceed the maximum number available to the system, some of those backups will queue. NetApp recommends a NearStore license on the NetApp secondary system so that the maximum number of streams is available. To determine the maximum number of concurrent transfers for a particular platform, refer to the "Data Protection Online Backup and Recovery Guide" on the NOW site.

Note: In Data ONTAP 7.3 the maximum number of transfers that can be running AND queued is 712. In Data ONTAP 7.3.1 and higher this increases to 1,024.

REPLICATING OSSV BACKUP DATA

OSSV backups cannot be replicated from the NetApp secondary to a tertiary system using SnapVault. They can, however, be replicated using SnapMirror®. Mirroring the backups using SnapMirror creates a duplicate copy of the backup data. This is especially useful when off-site backup copies are required.

OSSV AND MULTISTORE

When deploying OSSV in a MultiStore® environment, OSSV relationships can only be created using the default vFiler® unit (vfiler0). OSSV destination volumes can be owned by nondefault vFiler units, but the relationship must be configured using vfiler0. Configuring OSSV relationships using nondefault vFiler units is not supported.

8 NETAPP PARTNERS

NetApp partner solutions are also available for OSSV. Syncsort, CommVault, and Bakbone are NetApp partners that are able to manage OSSV using their management interface. In addition, these partners bring their own value-add and extend the functionality of OSSV.

Backup Express from Syncsort includes complete support for NetApp SnapVault, including OSSV management for Windows, Linux, and UNIX. In addition to Backup Express, Syncsort also provides its own OSSV agent, which supports additional application backup support with OSSV. Backup Express can be used to manage both the NetApp and Syncsort OSSV agent.

The CommVault Simpana suite, based on CommVault's Common Technology Engine, provides data protection by managing data throughout its lifecycle via integrated backup/recovery, migration, archiving, replication, and storage management. By adding CommVault QiNetix QuickRecovery, you can enable backup and recovery of Exchange, SQL, and Oracle® with the NetApp OSSV agent.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

NetApp®
www.netapp.com