NETAPP TECHNICAL REPORT

Open Systems SnapVault Best Practices Guide

Jeremy Merrill, Darrin Chapman

TR-3466

ABSTRACT

This document is a deployment guide for architecting and deploying Open Systems SnapVault® (OSSV) in a customer environment. It describes backing up and restoring data that resides on systems other than NetApp (open systems) to a NetApp® storage system by using NetApp SnapVault technology. As always, please refer to the release notes on NOW™ (NetApp on the Web) for updates and current requirements, issues, and limitations. This document is intended for field personnel who need assistance in architecting and deploying an OSSV solution.

TABLE OF CONTENTS

1 INTRODUCTION

The Open Systems SnapVault (OSSV) primary agent has extended the reach of NetApp SnapVault technology to the open systems server. OSSV facilitates block-level incremental transfers from the open systems platform directly to a secondary storage system. Various data sets can now be maintained remotely on a common NetApp storage system with or without the NearStore® license.

This document describes the OSSV theory of operation, major features and enhancements, management options, typical deployments, and best practices for deploying OSSV in an enterprise environment.

1.1 REQUIREMENTS AND ASSUMPTIONS

For the methods and procedures described in this document to be useful to the reader, the following assumptions are made:

- The reader has at least basic knowledge of backup and recovery in a tape and/or disk environment.
- The reader has at least basic UNIX® and Windows® administration skills, has access to the administrative login for the server, and has administrative access to the server console.
- The reader has at least basic NetApp administration skills and has administrative access to the destination storage system via the command-line interface.
- The secondary system has the licenses necessary to perform the activities outlined in this document. Specifically, the storage system must have the SnapVault secondary and Open Systems SnapVault primary licenses installed.
- The NetApp secondary system has the required block-level storage or network protocol interconnects to perform the activities outlined in this document.

In the examples in this report, all administrative commands are performed at the server, NetApp storage system console for clarity. Other management tools (Protection Manager, CommVault, Syncsort, NetVault) can be used, but they are not demonstrated in this document.

1.2 INTENDED AUDIENCE

The information in this document is intended for field personnel who are responsible for architecting and deploying Open Systems SnapVault solutions. A brief overview of OSSV basics is presented to establish baseline knowledge before describing the specific features, best practices, and actual installation and configuration.

2 OVERVIEW

Open Systems SnapVault is a heterogeneous disk-to-disk data protection solution that is ideal for use with NetApp storage systems. An OSSV primary system corresponds to a backup client in the traditional backup architecture. The SnapVault secondary is always a data storage system running Data ONTAP®. OSSV software protects data residing on a primary, which can be a storage system from a server running an operating system from leading server vendors such as Solaris™, HP-UX, AIX, Windows, VMware® ESX, and Linux®.

Three main components are installed in the OSSV environment:

- The primary system
- The OSSV agent residing on the primary system
- The secondary system

A predetermined directory or file system is chosen to be backed up to near-line storage. The data set is mapped to a secondary system qtree on a NetApp secondary storage system. Once the data to be protected is identified and a destination volume or qtree is chosen as the secondary, the agent can be installed on the primary system and the basic parameters configured. During installation, the agent installs various subdirectories on the open systems server. Components installed on the primary system include:

- OSSV primary database
- Set of OSSV executables
- OSSV log file
- OSSV exclude list files

2.1 THEORY OF OPERATION

There are two phases to the OSSV backup mechanism:

- Phase 1: File system scan on primary and directory structure built on secondary
- Phase 2: Actual data set transfer

After a successful baseline transfer, OSSV operates by examining files for changes via two methods: modification time and block checksums. The modification time is a coarse estimation of the true amount of changed data, due to the fact that the modification time is updated when at least one block of the file is written. By using 4kB block checksums, OSSV is able to back up only the portions of the file that have changed. This is referred to as block-level incrementals, or BLI. OSSV can back up whole files or changed blocks, depending on user requirements. In all cases, only changed data blocks are sent to the secondary system.

In a block-level incremental deployment, OSSV can significantly reduce the amount of network traffic over traditional backup strategies by sending only incremental changes in increments of 4kB data blocks. Once the initial baseline transfer is complete, OSSV sends only changed blocks, effectively resulting in an “incremental forever” strategy. Remote office backups, especially those over slower wide area networks, are now easier to achieve with the introduction of OSSV agents at the remote office. These remote offices can be backed up to a central location such as a data center.

However, environments with faster network connections and/or high change rates in many large files may benefit more with the block-level incremental option turned *off,* resulting in entire file transfers.

In addition to block-level incremental backups, OSSV introduces other critical features into the backup environment. These include open file backups, health checks, checkpoint restarts, exclude lists, and system state backup and restore.

OSSV deployment consists of several steps, including schedule determination, change rate determination, retention policy, performing a baseline or level-0 backup, volume creation and sizing, and relationship creation. The actual relationship is created during the baseline transfer.

RELATIONSHIP CREATION AND BASELINE TRANSFER

In response to command-line or NDMP-based management interface input, the SnapVault secondary storage system requests initial baseline (entire file system requiring backup) image transfers of directories specified for backup from an open systems platform. These transfers establish SnapVault relationships between the open systems platform *directories* and the SnapVault secondary *qtrees*.

The open systems platform, when prompted by the secondary storage system, transfers initial base images of specified directories to a qtree location on the secondary storage system. When the baseline transfer is complete, the secondary system creates a Snapshot™ copy (baseline) of the volume containing the destination qtree. If multiple transfers are occurring, faster transfers are in a “quiescing” state until *all* transfers are complete.

A new Snapshot copy is created each time a baseline is performed, and up to 250 Snapshot copies can be maintained according to a schedule configured by the backup administrator. Each Snapshot copy consumes an amount of disk space equal to the differences between it and the previous Snapshot copy.

SCHEDULING AND RETENTION POLICY CONSIDERATIONS

In a typical legacy backup environment, incremental backups were usually performed once a day, with full backups once a week. The fastest restore could take hours and required the intervention of a backup operator or a system administrator. In the OSSV or SnapVault configuration, incremental backups are performed as often as once an hour (with daily and weekly options); each incremental backup is usable as if it were a full backup; and most restores can be performed in minutes or less by end users, without the need for backup operator intervention or use of a backup server.

Two scheduling options are currently available: command-line scheduling in Data ONTAP and an NDMP-based scheduling mechanism using Protection Manager or another supported NDMP management tool.

To determine the proper scheduling and retention policies for a particular environment, it is important to understand the backup and restore requirements of the organization. Several factors must be considered, including restore granularity, media costs, data change rates, types of data to be protected, and risk.

An excellent resource for determining schedules and retention policies is the *SnapVault Best Practices Guide*.

3 OSSV FEATURE REVIEW

This section describes several but not all of the features included with Open Systems SnapVault.

3.1 BLOCK-LEVEL INCREMENTALS

Open Systems SnapVault BLI backup is designed to minimize the backup of data that has not changed since a previous backup operation. OSSV uses checksums to identify portions of a file that have changed between a previous and the current backup.

A BLI backup recognizes that a file has changed based on a time stamp and checksum algorithm. Exactly which blocks have changed is determined, and only those blocks are sent to the secondary storage system.

Typically, incremental backups are more frequent, reduce the amount of time required to back up data, and minimize the resources required to perform backups when compared to baseline or full backups. BLI significantly reduces the amount of data that needs to be transferred to backup storage, as well as the amount of data that must be stored on backup storage disk.

Changed blocks are recognized based on checksum values that are calculated and preserved for each block by the OSSV agent. Checksums are calculated on 4kB blocks of file data stored on an internal database. These checksum database files are stored in the OSSV internal DB directory. Each relationship has its own checksum file directories. *Approximately 2%* of the baseline ends up being the checksum file database size.

First, time stamps of files are compared to the time of the last successful backup operation. After being identified, a checksum is performed on that file. By default, every block of every file has a checksum operation performed against it during baseline operations. This is referred to as “high” BLI and results in typically longer transfer times and more CPU and disk consumption on the primary system. You can configure block incremental processing to trade off efficiencies among four variables: primary system CPU utilization, disk consumption, network bandwidth utilization, and OSSV transfer time. Enabling block-level incremental updates normally causes a checksum value to be calculated for every block of every file during the initial OSSV baseline transfer. As a result, baseline transfer execution time, CPU utilization, disk consumption, and network bandwidth utilization are increased compared to incremental transfers that are not block related. In this case, it is possible that significant resources can be consumed by calculating checksum values for static files that never change. The checksum levels can be configured as high, low, or off, by using the `svconfigurator` utility.

- **High.** Always computes checksums, on baseline transfers and incremental updates..
  Primary impact: High CPU utilization on the baseline transfer
  Network impact: Lower amount of data transferred on updates (incremental)
  Secondary impact: Same as network impact
  When to implement: If all files are subject to small changes

- **Low.** Compute checksums on changed files and only on updates; no checksum performed during baseline.
  Primary impact: More CPU utilization during updates, faster baseline transfer
  Network impact: Large amount of data transferred during baseline, lower amounts of data transferred during updates
  Secondary impact: Same as network impact
  When to implement: If a small subset of files is likely to change

- **Off.** No checksums are calculated at any time. Similar to older versions of OSSV. Full files are transferred once they are identified as being changed files.
  Primary impact: Fast baseline transfer, less impact on CPU during file system scan

Network impact: Large amount of data transferred during baseline, potentially large amounts of data (large files) during updates as well
Secondary impact: Same as network impact
When to implement: If a small subset of files is likely to change or if files are changing completely

3.2 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.

3.3 OPEN FILE BACKUP

In a Microsoft Windows environment, there are currently two options for making sure that open files are backed up successfully: Open File Manager (OFM) and Microsoft Virtual Shadow Copy Service (VSS).

The OFM utility allows Windows 2000 files that are open and in use to be backed up with only a very short disruption to users or their current applications. OFM is automatically installed in your system at the same time the proper agent is installed. It is not enabled unless the Windows 2000 Server is rebooted and the OFM component has been licensed on the secondary system.

Note: This is a separate license on the secondary system, and it must be installed in order for OFM to operate properly.

Windows 2003 provides a native snapshot mechanism as part of the VSS. VSS snapshot functionality (called shadow copy) is integrated with the OSSV agent as a standard feature. No VSS installation is required; although there are configurable parameters, there are no required configuration steps for OSSV and VSS integration. All Windows 2003 Open Systems SnapVault agent backups use VSS unless it is specifically disabled from the secondary by using the `back_up_open_files=off` option. This option can be used with OFM as well.

Both open file backup components can be tuned by using the `svconfigurator` tool, which is built into the OSSV primary system.

OFM CONFIGURATION OPTIONS

OFM waits for a set time when write activity is quiet and the system is in a safe state to initiate the backup of open files. OFM listens continuously for a period of write inactivity until it is ready to initiate a synchronize for backup or until the preset synchronization timeout period has expired. This value is configurable in `svconfigurator` under the SnapVault tab; the parameter is referred to as “OFM Write Inactivity Period (seconds)” and it defaults to 3 seconds. The minimum value is 1 second; the maximum value is 60 seconds. OFM waits this period of time for write inactivity; if it cannot find a write inactivity period and has tried for 60 seconds, it fails. This type of failure can result in a blank backup.

This default 60-second timer is known as “Maximum time to wait for OFM synchronization (seconds),” and it can be found under the SnapVault tab in `svconfigurator`. Again, the range is 1 to 60 seconds.

Specific drives can be excluded from open file backup by using OFM. There is an option under the SnapVault tab in `svconfigurator` that allows this exclusion.

Note: OFM supports one active Snapshot copy per disk volume letter. For example, if you try to use OFM to back up two directories on the same disk volume, one of the two directory backups fails. If an active Snapshot copy is encountered during an attempted backup, OFM fails.

Note: Approximately 15% free space is required on the drive being copied by using OFM. If you are running OSSV 2.2 or later, you can use the Free Space Estimator to determine whether there is sufficient free space to run OFM.

VSS CONFIGURATION OPTIONS

Certain conditions must be met before the OSSV agent can acquire a VSS snapshot. You can set the amount of time (snapshot timeout) that the agent waits until it retries a VSS snapshot if the conditions are not right at the time. Setting this parameter avoids unacceptably long waiting periods. The default is 180 seconds (the maximum value). The minimum value is 1 second.

Like OFM, the VSS values can be modified by using `svconfigurator` under the SnapVault tab.

Note: When using OFM/VSS with OSSV, the snapshots are placed in the following “pseudo” file systems:

```
OFM: \\?\Z:
VSS: \\?\GLOBALROOT\Device\Harddisk\VolumeShadowCopy42
```

However, the snapshots in these locations are discarded after the transfer, and OSSV can’t be used to back up the snapshot locations. OSSV handles all interaction with OFM or VSS.

3.4 CHECKPOINT RESTART

Checkpoint restarts allow the user to restart a failed baseline transfer and a failed incremental transfer. Checkpoints are taken every 5 minutes by default. The restarts may occur at these 5-minute intervals, although they may span multiple 5-minute intervals and go back further into the transfer than expected. The unexpected restart of longer than 5 minutes is due to the fact that the restarts occur at file boundaries. If a large file spans multiple 5-minute intervals, the restart goes back further.

Checkpoint restarts occur only during phase 2 (data send phase) of the OSSV transfer.

Note: This option is not a replacement for resyncing relationships in a disaster recovery scenario.

Note: Starting with OSSV 2.5, a checkpoint is taken at specified time intervals (default of 5 minutes), regardless of file boundaries. This functionality is referred to as block-level checkpoints.

3.5 BACKUP EXCLUSION LISTS

This feature allows the user to exclude files and paths from backup. There are two files that are installed once the OSSV agent is in place on the primary system. The `file-exclude.txt` file contains file exclusion configuration information with wildcard capabilities. The `path-exclude.txt` file contains path exclusion information where full directories and their contents can be excluded. Both files are located under the `$INSTALL_DIR\etc` directory by default.

A file or directory is excluded if the filename or any path element matches a file exclusion entry in the list in one of these two files.

On Windows systems, exclusion list files are Unicode text files. On UNIX systems, exclusion list files are multibyte text files. Each entry is on its own line. Wildcard characters are supported. Use an asterisk (*) to specify any number of characters within a single path element. Use a question mark (?) to specify one character within a single path element. Use an exclamation mark (!) to remove the special meaning from * or ?. Use a pound sign (#) at the beginning of a line for a comment.

Note: For details, see the *OSSV Installation and Administration Guide* (available on NOW), or open the exclusion files to see notes and comments.

3.6 SYSTEM STATE BACKUP AND RESTORE

Various components, including the REGISTRY and system files, can now be backed up and restored by using the OSSV agent. The components included in the system state backup vary depending on the operating system, installed applications, and configuration. These components may include:

- REGISTRY
- System files and settings, including the boot files
- System files that are under Windows file protection
- Certificate services database
- IIS metadirectory

- Various performance counters
- Active Directory® and SYSVOL data (domain controllers)
- Event Log, Application Log, Security Log (OSSV 2.5 and later)

Essentially, replacing the primary file system path with the keyword `SystemState` initiates a system state backup of the primary:

```
snapvault start -S ossv_prim: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.

Note: On Windows 2000 systems, before starting system state data backups, make sure that the system has recent Windows service packs installed. You can have problems with restored systems if the backed-up system was running an older version of the operating system.

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

```
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.

For more information on system state restores, see the OSSV release notes, available on NOW.

You can also use system state backups as part of a disaster recovery plan. To create a backup for use in a disaster recovery plan, you essentially back up the entire system drive and any other relevant partitions or drives and back up the system state. Be aware that when recovering from a disaster by using a complete system backup, NetApp does not support “bare metal restore.” The user would need to make sure that the base operating system is installed on an identical hardware configuration with identical service packs, names, drive letter mappings, file system types, and so on. For more details on backing up and restoring a complete system for disaster recovery purposes, refer to the OSSV release notes, available on NOW.

3.7 OSSV DATABASE BACKUP

With older versions of the OSSV client database, the database backup is a manual process. There are two options for backing up the database in older versions. The user can simply include the `$INSTALL_DR/db` path in the backup stream. Or the `svdb` command can be used, which creates a file under the `$INSTALL_DR/db` directory.

Beginning with OSSV 2.2, the OSSV database backup is performed automatically with every OSSV transfer. There are three configuration levels for the client database. “BLI” backs up the history file and its corresponding BLI checksum file; “DB only” backs up only the history file; and “None” disables the automatic database backup. The options can be set from the `svconfigurator` utility, under the SnapVault tab. The database is created under the qtree root as `.OSSV_Database` on the secondary. To restore this database, issue a `snapvault restore` command on the primary. When performing the restore, specify any destination, such as`\temp\database`, and OSSV automatically detects the database, decodes it, and places it in the appropriate directory.

3.8 COMMON SNAPSHOT MANAGEMENT

Prior to OSSV 2.5, whenever the transfer of files failed, the Snapshot copies were deleted and a new Snapshot copy was created during the transfer restart. However, Open Systems SnapVault 2.5 has the ability to retain old Snapshot copies and to use these copies during transfer restarts. This process is called *common Snapshot management.* There are two possible configurations for this parameter, `MaxCPRestartWaitTime` and `FailCPRestartOnNewSnapshot`. For information about setting these parameters, see the *OSSV Installation and Administration Guide*.

3.9 LINK COMPRESSION

With OSSV 2.5 and later, it is possible to transfer the data in a compressed state. This functionality helps reduce bandwidth consumption by compressing the data before sending it across the network. When the data is received on the destination system, it is decompressed and written to disk, so the data is not stored in a compressed state. Link compression is either enabled on a per relationship basis, or it can be enabled globally with Data ONTAP 7.3.1 using the option `options snapvault.ossv.compression on`. In addition, it is possible to turn compression on and off by using the `snapvault modify` command with the `–o compression=on/off` option for each relationship.

Note: To use link compression, it is necessary to have Data ONTAP 7.3 installed on the destination system.

3.10 FREE SPACE ESTIMATOR

Free Space Estimator is a command on the primary (`svestimator`) in OSSV 2.2 that you can use to determine whether there is sufficient disk space available on the primary to perform an OSSV transfer. Free Space Estimator can be run on a primary with or without OSSV already installed. It can help determine whether a specified drive has enough space available to install OSSV on a new client.

There are two modes for Free Space Estimator, built-in mode and standalone mode. In built-in mode, Free Space Estimator runs in the background, at the start of every OSSV transfer, and reports whether there is sufficient space to back up based on the current OSSV configuration. The results are recorded in the `$INSTALL_DIR/etc` directory. If sufficient space isn't found, the operation is not aborted by default. However, this can be changed by modifying the `snapvault.cfg` file. In addition, it is possible to disable Free Space Estimator from running automatically by disabling it in the `svconfigurator` utility under the SnapVault tab. In standalone mode, Free Space Estimator is installed as a standalone application on a system that may or may not have OSSV installed. This utility can be useful if you are planning to deploy OSSV and need to know how much free space is available and which directories can handle the OSSV installation.

Note: Free Space Estimator is set to run before any OSSV transfer. If the file system to be backed up is large, Free Space Estimator may take an extended period of time to estimate space requirements, causing the secondary to time out before the backup can start. In this case, NetApp recommends that Free Space Estimator be disabled before the OSSV transfer.

3.11 UNATTENDED INSTALL

Unattended Install enables you to install or upgrade OSSV software on a primary storage system with minimal user intervention. Although this new feature does assist in creating an

installation image, it is not 100% hands off. This feature is useful for deploying over a large number of primary systems. Unattended Install allows you to set the installation variables noninteractively, and in most cases a reboot is not required after the installation. To perform an unattended installation, an installation script and other supporting files are required. To gather these files, use the utility `svconfigpackager`, available with OSSV. This utility saves the current configuration settings to file when run on a primary. In addition, it creates an installation script that in conjunction with the configuration settings file and other files can be used to perform unattended installations or upgrades.

When creating an Unattended Install package, the installation script and other files created by the `svconfigpackager` utility on an operating system cannot be used for running an unattended installation on a different operating system. For example, an Unattended Install created on a Windows 2000 Server cannot be used on a Windows Server 2003. In addition, there are separate installation scripts for Solaris and HP-UX. For more information, see "Creating an Unattended Installation Image" in section 7.

Note: Unattended Install is supported only on systems running OSSV 2.x; OSSV 1.x agents cannot be upgraded by using this technique.

3.12 RESYNC AFTER RESTORE/BREAK

This feature is available starting with OSSV 2.2, Resync after restore/break allows the user to resynchronize a relationship without requiring a new baseline transfer. Before OSSV 2.2, if a relationship got out of sync, a new baseline would be required. A system is considered synchronized as long as a common snapshot exists between the primary and secondary. If this is lost, the incremental backups begin to fail. There are three ways an OSSV relationship can get out of sync:

- An older version of the OSSV database is restored to the primary
- Data is restored by using the `snapvault restore` command
- The state of the destination qtree in an OSSV relationship is changed to read-writable (even if the contents in the qtree were not modified)

Note: If the contents of the qtree were modified, all data written to the qtree will be lost on the resync.

To resync a relationship, use the `snapvault start -r` command from the secondary.

Note: Resync after restore/break resynchronizes the relationship between the primary and secondary. This does not propagate changes back to the primary; it is not similar to the resync in SnapMirror®.

Note: Resync after restore/break is not available in versions earlier than Data ONTAP 7.1.

3.13 LREP

Offices that are small enough to have their name, print, and file service requirements met by a single server typically have limited WAN bandwidth. OSSV is a strong solution for remote office backup, but the initial transfer can be crippling. LREP can be used to seed the baseline on the secondary. Included with the OSSV installation are the binaries `lrep_reader` and `lrep_writer`.

Use the LREP utility to write the initial transfer (the baseline) to a portable drive, such as a Zip drive. The portable hardware could also be a FAS250 shared between offices. Ship the portable media to the data center and locally write to the secondary system. No network bandwidth is used, only a manual process of moving the media from remote site to data center. Once the data is on the secondary system, modify the OSSV relationship to reflect the real primary → secondary relationship.

The latest version of LREP has built-in support for both compression and encryption. These two components allow compression and encryption of the data while it's in transport to the data center.

For more information on LREP, see the *LREP User Guide*, available on the toolchest page for LREP on NOW.

4 MANAGEMENT OPTIONS

Open Systems SnapVault can be managed from a number of NDMP-based applications. Once the primary agent has been installed, the schedules and retention policies have been determined, the network is in place including any firewall settings, and the directory to qtree mappings has been laid out, a management system can be introduced acting as a third party in a primary-to-secondary relationship. The management system communicates over an available TCP/IP network interconnecting the primary, the secondary, and the management system. In all cases, NDMP is used as a transport for SnapVault messages and commands. All scheduling, baselines, relationship creations, retention policies, backup control, and monitoring can be centrally configured on the management system—all using NDMP over the IP network.

The current supported management applications are:

- NetApp Protection Manager
- BakBone NetVault
- Syncsort Backup Express
- CommVault QiNetix or Simpana

In addition to these GUI-based applications, OSSV can be managed via the Data ONTAP CLI. This document, uses the Data ONTAP CLI.

4.1 PROTECTION MANAGER

Protection Manager is an intuitive and innovative backup and replication management software package for NetApp disk-based data protection environments. Protection Manager delivers greater data protection and higher productivity by providing policy-based management, including automated data protection setup. This management application lets your administrators apply consistent data protection policies across the enterprise, automate complex data protection processes, and pool backup and replication resources for higher utilization. For more information on using Protection Manager and the minimum requirements, see the *Protection Manager Administration Guide* on NOW.

4.2 COMMAND LINE INTERFACE

To manage OSSV relationships via the Data ONTAP CLI, you use the `snapvault` commands. The commands are the same as for SnapVault, except that you identify the OSSV primary as

the source. All backup schedules and relationships are configured on the secondary. In addition, for restores, you use the CLI that is installed on the primary.

4.3 SYNCSORT

Syncsort Backup Express has been certified for NetApp Data ONTAP and is currently in collaborative development on Data ONTAP 7.0. Fully integrated OSSV management is available with Backup Express. Backup Express includes complete support for NetApp SnapVault, including OSSV management for Windows, Linux, and UNIX. In addition to Backup Express, Syncsort provides its own OSSV agent. Backup Express can be used to manage both the NetApp and Syncsort OSSV agents. The Syncsort OSSV agent for Windows also includes integrated application support for Exchange and SQL Server®. In addition, Syncsort provides a bare metal recovery option.

Note: Although Syncsort does provide its own OSSV agent, it does not perform exactly the same as the NetApp OSSV agent. Syncsort creates images that it transfers, which can be mounted as snapshot-based LUNs to select individual files or database objects for restore. For more information, visit the Syncsort Web page.

4.4 COMMVAULT

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

5 BEST PRACTICES AND RECOMMENDATIONS

This section discusses recommendations and common best practices. Many of these items have been uncovered by field testing, others by lab testing. These items are not necessary for OSSV to function completely, but they definitely have an impact in terms of sizing, scalability, manageability, availability, configuration, and overall architecture. Be sure to consider these items in all OSSV deployments. Most are well-tested issues and are highly recommended for production installations of Open Systems SnapVault.

5.1 TAKE STOCK OF YOUR DATA

Gather the following information for every machine to be backed up.

Identify all directories that will be backed up. Obtain estimates of how much data is contained in these directories. How frequently does this data need to be backed up? What is the rate of change of data? How large is the backup window? Can you kick off multiple backups simultaneously? How many volumes are available on the secondary? What is the size of the baseline data set? Does the data set have non-ASCII filenames, such as filenames that contain Japanese characters?

Caution: A very large number of small files (greater than 1 million) can have an adverse impact on performance and can also result in a significantly large amount of overhead data being transferred during backups. SnapVault introduces about 16kB of overhead for every file being backed up.

5.2 SECONDARY CONSIDERATIONS

Use of Data ONTAP 7.x or later is recommended. At a minimum you will need an sv_ontap_sec license and an OSSV primary license. For Windows 2000, NetApp recommends an OFM license. Create secondary volumes and pay close attention to the size of volumes. They must have enough space to hold backups and up to 250 Snapshot copies. Check the `maxfiles` value of the volume. If the backup contains or may contain non-ASCII filenames, the storage system's volume language should be changed to contain UTF-8. Plan your backups—for example, what primary directory will go into what secondary volume? Assign easy-to-read qtree names, consistent qtree names, and names that correspond to source file systems and hosts.

NetApp highly recommends having all relationships that take roughly the *same transfer time* go into the *same volume*. For example, you might save weekly backups of all long-running transfers to one volume and nightly backups of short-running transfers to another volume. Mixing long and short transfers results in the longer transfers holding up the shorter ones from being visible. For faster transfers, output of "quiescing" appears if fast and slow transfers are mixed on the same volume. Those faster transfers remain in a quiescing state until the slowest transfer is complete; this can be viewed as holding up or slowing down the faster transfers.

Remember that schedules for SnapVault are on a per volume basis. One volume with many relationships can potentially cause a flood of backup traffic every time the scheduled backup time is reached. OSSV attempts to back up all relationships in that volume based on that one schedule. Creating more than one volume results in multiple schedules; each volume can have its own schedule; and the schedules can be staggered.

When defining backup schedules, keep in mind SnapVault transfer limits. If all relationships are in one volume, which maintains one schedule, all transfers occur concurrently. The R200 platform supports a maximum of 128 concurrent SnapVault transfers; with small numbers of volumes and large numbers of relationships, you run the risk of exceeding this limit. If the limit is exceeded, the remaining transfers are queued up and started as others complete. Other smaller platforms support only between 4 and 16 concurrent transfers; those platforms are limited, especially when dealing with large numbers of SnapVault relationships.

NEARSTORE PERSONALITY

To enable customers to use FAS storage systems as secondary storage, a new software license option called NearStore Personality (nearstore_option) has been introduced. This license option can be installed only on the FAS3020 and 3050 systems. This option is supported on Data ONTAP 7.1 and later. This license option provides increased concurrent streams when FAS3020 or 3050 storage systems are used as destinations for SnapMirror and SnapVault transfers and to enable SnapVault for NetBackup™. This license option should not be installed on these storage systems if they are intended to handle primary application workloads.

CONCURRENT REPLICATION LIMITS

The default Data ONTAP behavior without the nearstore_option license is to maintain a fixed upper limit for concurrent SnapMirror and SnapVault transfers based on the type of disks the storage system has attached.

Once the nearstore_option license is installed, the storage system switches to NearStore Personality. When the storage systems take on the NearStore Personality, Data ONTAP limits

the maximum concurrent transfers based on the type of the replication operation. For more information on the maximum concurrent streams for your platform, see the *Data Protection Online Backup and Recovery Guide* on NOW.

5.3 SPACE REQUIREMENTS

The OSSV agent requires space on the primary system for various components. All default OSSV subdirectories, including the `trace` and `db` subdirectories, are installed here. Open Systems SnapVault needs disk space for its built-in database. The database disk space requirements depend on the number of files and the average file size and number of directories.

- If block-level incremental backup is set to Off:
  Size of database = number of files and directories in the backup data set X 96 bytes
- If block-level incremental backup is set to: Low or High:
  Checksums require 16 bytes per 4K bytes of data
- Size of database = (number of files and directories in the backup data set x 96 bytes) + [total size of all files in the backup data set / (4K bytes x 16 bytes)]

Temporary space requirements are approximately twice the size of the database just described.

If you are running OSSV 2.2 or later, a simple way to determine whether you have enough space is to use Free Space Estimator, `svestimator`. It's a good idea to run this on a schedule (weekly, monthly, or quarterly) to avoid incidents.

Note: For the latest space requirements for the OSSV primary database, see the OSSV 2.x release notes, available on NOW.

In addition to built-in database space requirements, if you are using Open File Manager to manage open file backups by using Snapshot technology for a particular drive, OFM needs extra space for the drive for which it is creating a point-in-time copy. NetApp recommends a minimum of 15% additional free space on the file systems that are being backed up. If free space is not available, disable OFM for those drives.

SNAPVAULT OVERHEAD

When using OSSV, there is always some sort of overhead to be transferred for files in the OSSV relationship that have been modified. The OSSV primary sends one 4kB header for every file or directory that exists in the relationship. In addition, for files or directories that are larger than 2MB, an additional 4kB header is transferred for every 2MB.

6 OTHER CONSIDERATIONS BEFORE DEPLOYMENT

6.1 MULTIPLE CONCURRENT OSSV TRANSFERS FROM THE SAME OSSV PRIMARY

You should plan your backup schedules so that 16 or fewer transfers occur at the same time from the same primary system. It is best to create multiple schedules on different volumes on the secondary system; or to simply eliminate (by consolidation) the number of file systems being backed up from the primary system.

6.2 LOW-BANDWIDTH NETWORK LINKS

Customers who deploy Open Systems SnapVault in environments where a low-bandwidth network link separates the primary from the secondary have an aversion to repeating baseline transfers, for obvious reasons. These customers should perform frequent backups of the primary OSSV database to minimize the possibility of OSSV database corruption requiring a baseline transfer to be repeated. Resync after restore/break is not available prior to OSSV 2.2, so be sure to use `svdb` and back up the local `$INSTALL_DIR/db`. If a baseline transfer takes an extremely long time, consider backing up the OSSV database after every incremental. Make sure that older database backups are deleted from the primary due to free space considerations. This isn't as much of an issue if you are running OSSV 2.2, because of the resync after restore/break functionality.

6.3 SOURCE DATA CONSIDERATIONS

A large number of small files may degrade performance and also result in a large amount of overhead data sent over the network. The `snapvault status` output reveals higher-than-expected data transfers when such a backup is complete.

If a large number of files are not likely to be modified, consider changing the BLI level to Low. This keeps the OSSV database size in check.

6.4 DATABASE BACKUPS

An application database must be unmounted (shut down) before OSSV backups are initiated. OSSV performs a file-level backup in most database environments, because file modification times are constantly changing. OSSV is not integrated with any database backup APIs or prescripts or postscripts. The database files need to be dismounted (brought to a logically consistent state and closed) before using OSSV to back them up. If users want to use the "hot backup mode" method, they must script it and test it themselves to make sure that the procedure works reliably in their environment. It is crucial that an exact procedure is followed in the script.

In addition to scripts that reside on the hosts, when running DFM 3.2 or later, the pre- and postscripting capabilities can be used to manipulate the database before and after the OSSV transfer.

7 INSTALLATION AND CONFIGURATION

7.1 PRIMARY SYSTEM PLATFORMS

Follow all recommendations in the previous sections for calculating free space on the primary system. Don't forget about OSSV primary database requirements and Open File Manager space considerations (in a Windows environment). For specific memory, disk space, and other requirements, refer to the *OSSV Installation and Administration Guide,* available on NOW.

7.2 LICENSING

All licensing is configured on the secondary system. For every open systems platform being backed up, a matching OSSV primary license must be purchased and installed on the secondary system. The following license pack sizes are available:

- Windows: `sv_windows_pri` (takes care of each supported Windows OS)

- VMware ESX: `sv_vi_pri` (takes care of each supported ESX Server)
- UNIX: `sv_unix_pri` (takes care of each supported UNIX OS)
- Linux: `sv_linux_pri` (takes care of each supported Linux OS)

If you are using Open File Manager in a Windows 2000 environment, a separate license must be installed, in addition to the base Windows license.

- OFM: `sv_windows_ofm_pri`

The secondary systems must also be licensed as secondaries.

- SnapVault Secondary: `sv_ontap_sec`

In addition, if you are using the NearStore Personality on your secondary systems:

- NearStore Option: `nearstore_option`

For more information, refer to the *OSSV Installation and Administration Guide,* available on NOW.

7.3 SECONDARY SYSTEM REQUIREMENTS

The secondary system should be running a minimum of Data ONTAP 6.5. For more information about the latest version of Data ONTAP, refer to the OSSV release notes, available on NOW.

Basic configurations:

```
options snapvault.enable on
```

`options snapvault.access all` (or list of IP addresses or hostnames allowed to back up to this system)

7.4 FIREWALL

If there is a port filter or firewall between the OSSV primary and secondary, make sure that TCP port 10566 (the SVListener port) is open. If you are using DFM or another NDMP-based management tool to manage the OSSV relationship, make sure that TCP port 10000 (the default NDMP port) is open as well.

Port 10566 must be open in both directions during restore operations.

7.5 RUNNING THE FREE SPACE ESTIMATOR

In this example, we run the Free Space Estimator on a Windows 2003 system (the process is the same for other platforms). There are three options for the `svestimator` command:

- `-o` if OSSV is already installed; if so, it uses the current OSSV setting configured in the `svconfigurator` utility .
- `-i` includes the OSSV installation package size in its calculation; this is the option to use when determining whether OSSV can be installed on a new server.
- ‘`-d` outputs debug trace to a directory that is created in the current directory.

In order to run, Free Space Estimator requires two files to properly estimate free disk space: the path and file exclusion lists (located in `$INSTALL_DIR/etc`) and a configuration file named `estimator.cfg`.

The `estimator.cfg` file contains user-defined options that are taken into consideration when estimating free space (located in `$INSTALL_DIR/config`). On a standalone, you must create this file and the path and file exclusion files in the directory in which the standalone space estimator is run.

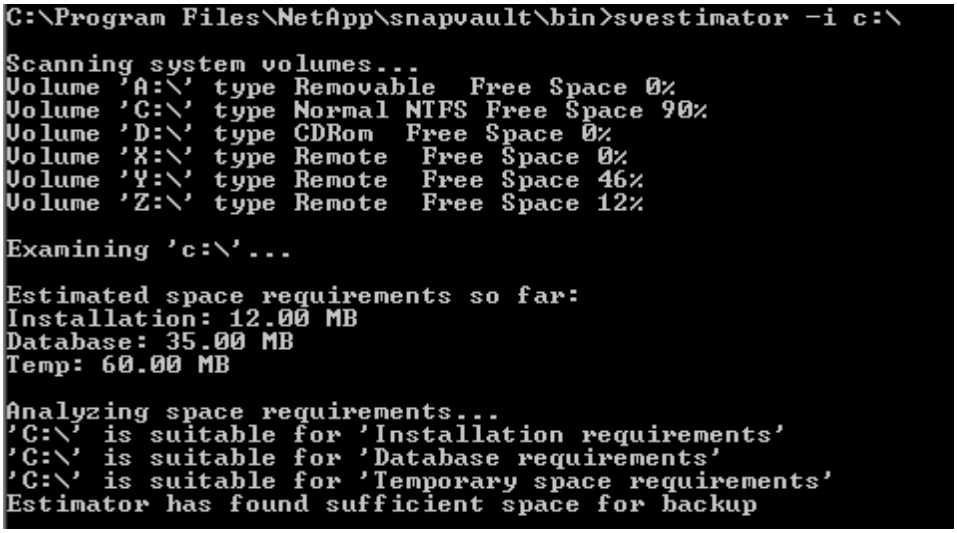

Figure 1) svestimator output (standalone).

Figure 1 is an example of running the Free Space Estimator on a Windows 2000 system that doesn't have OSSV installed. For this we use the `-i` option, which shows that we have 90% free space on the `C:\` drive. After it scans all the system volumes, it determines how much space is required and whether there is enough on the specified path. The final section of the output tells us that we do have enough space for the installation files, the database requirement, and the temporary space requirements.

7.6 SVCONFIGURATOR

Figure 2) svconfigurator.

After installing the OSSV agent on your primary system, you complete the basic configuration by modifying any outstanding parameters using the `svconfigurator` tool (accessible from the Start menu or from the command line, or `$INSTALL_DIR\bin\svconfigurator.exe` (`$INSTALL_DIR/bin/svconfigurator` in UNIX).

When the `svconfigurator` GUI appears, browse the tabs to view information about your particular installation.

Figure 3) svconfigurator, Machine tab.

The Machine tab displays information about the version of OSSV and the primary system machine information (OS, hardware, and so on).

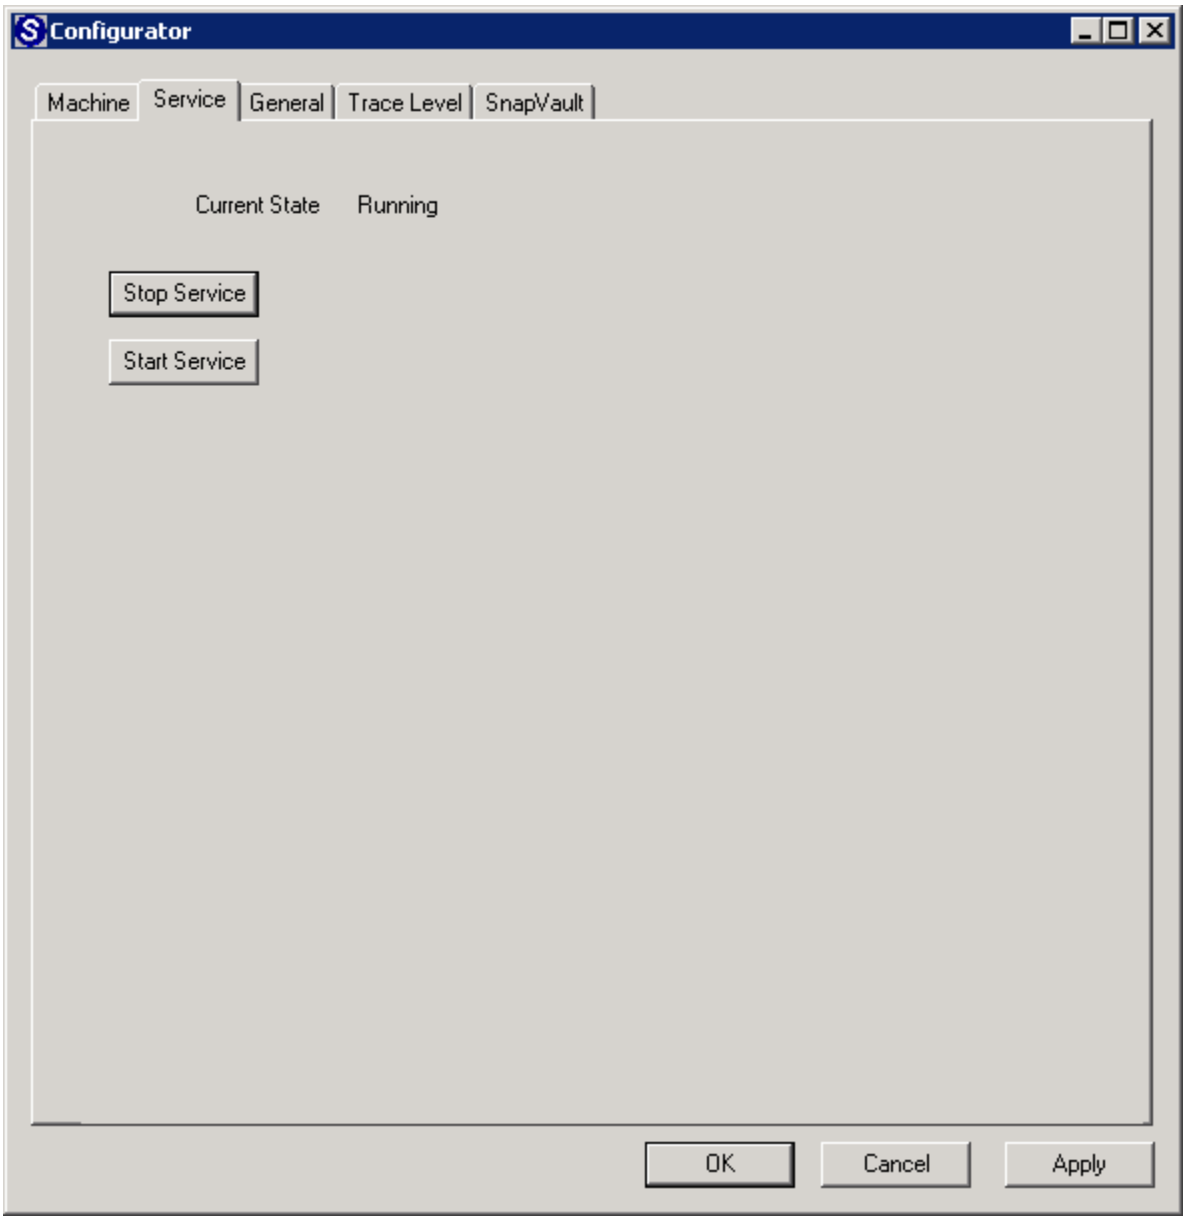

Figure 4) svconfigurator, Service tab.

You use the Service tab to stop and start the OSSV services.

Note: Use this tab when stopping and starting the OSSV service, to make sure that all pertinent services for OSSV are stopped and then restarted. NetApp recommends that you do *not* you stop and start OSSV from Windows Services.

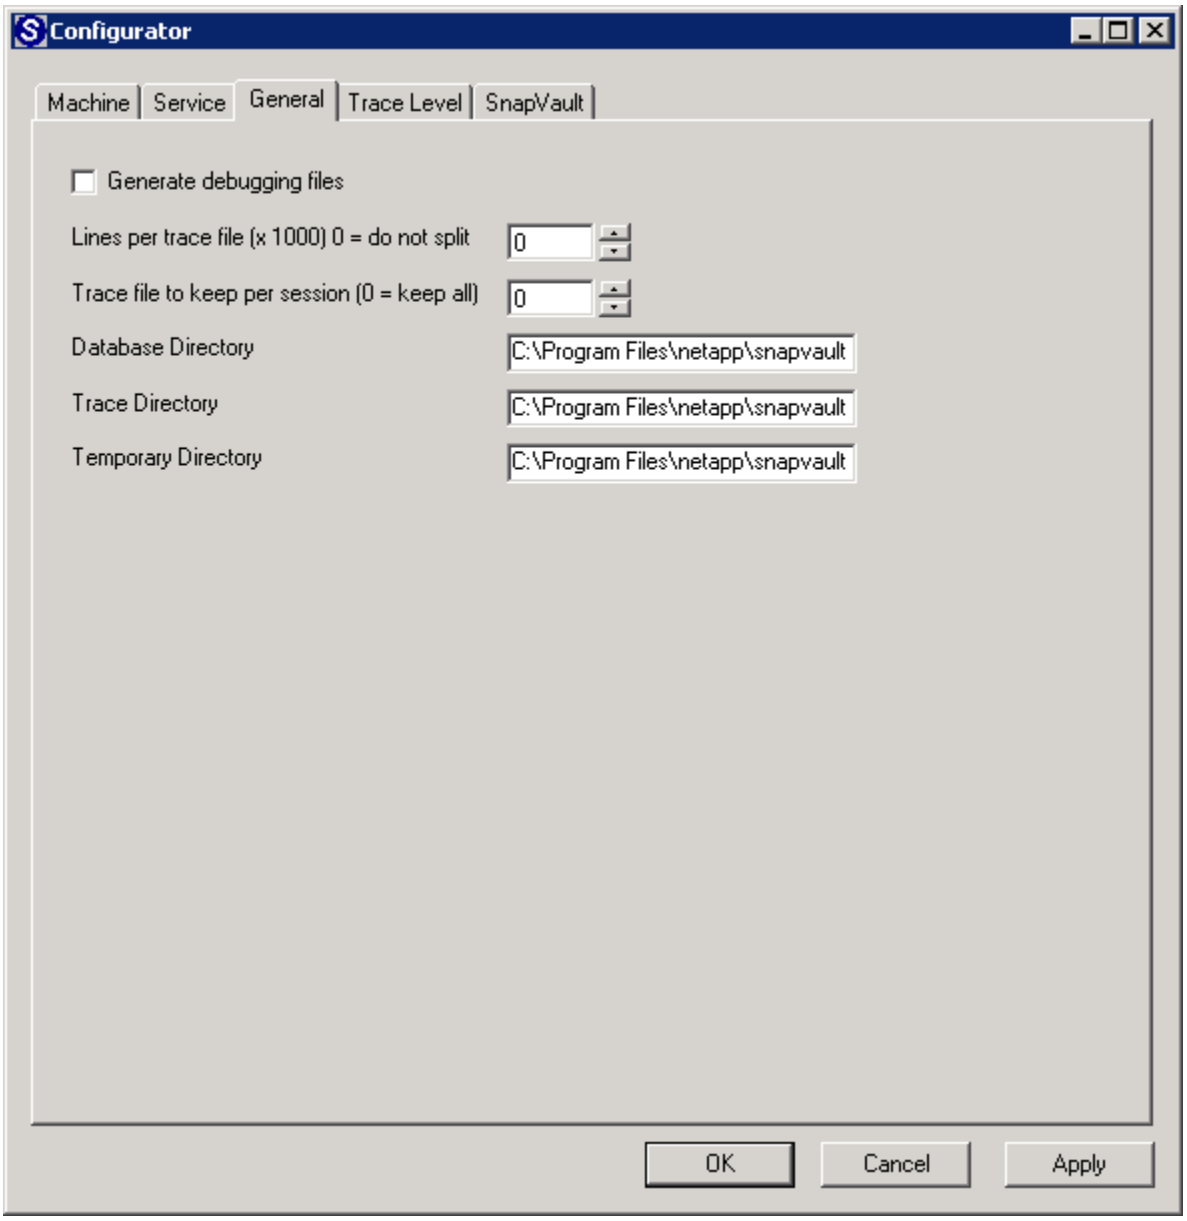

Figure 5) svconfigurator, General tab.

You use the General tab to generate debug files. The General tab also allows default directory locations to be modified, which might be necessary if these default file systems are approaching maximum capacity. Be very clear about the amount of free space in the `$INSTALL_DIR\db` directory.

Note: If you generate debug logs, especially in Verbose mode, you must set the level back to Normal once all relevant data is gathered. Otherwise, the `$INSTALL_DIR\trace` directory will rapidly reach its limit.

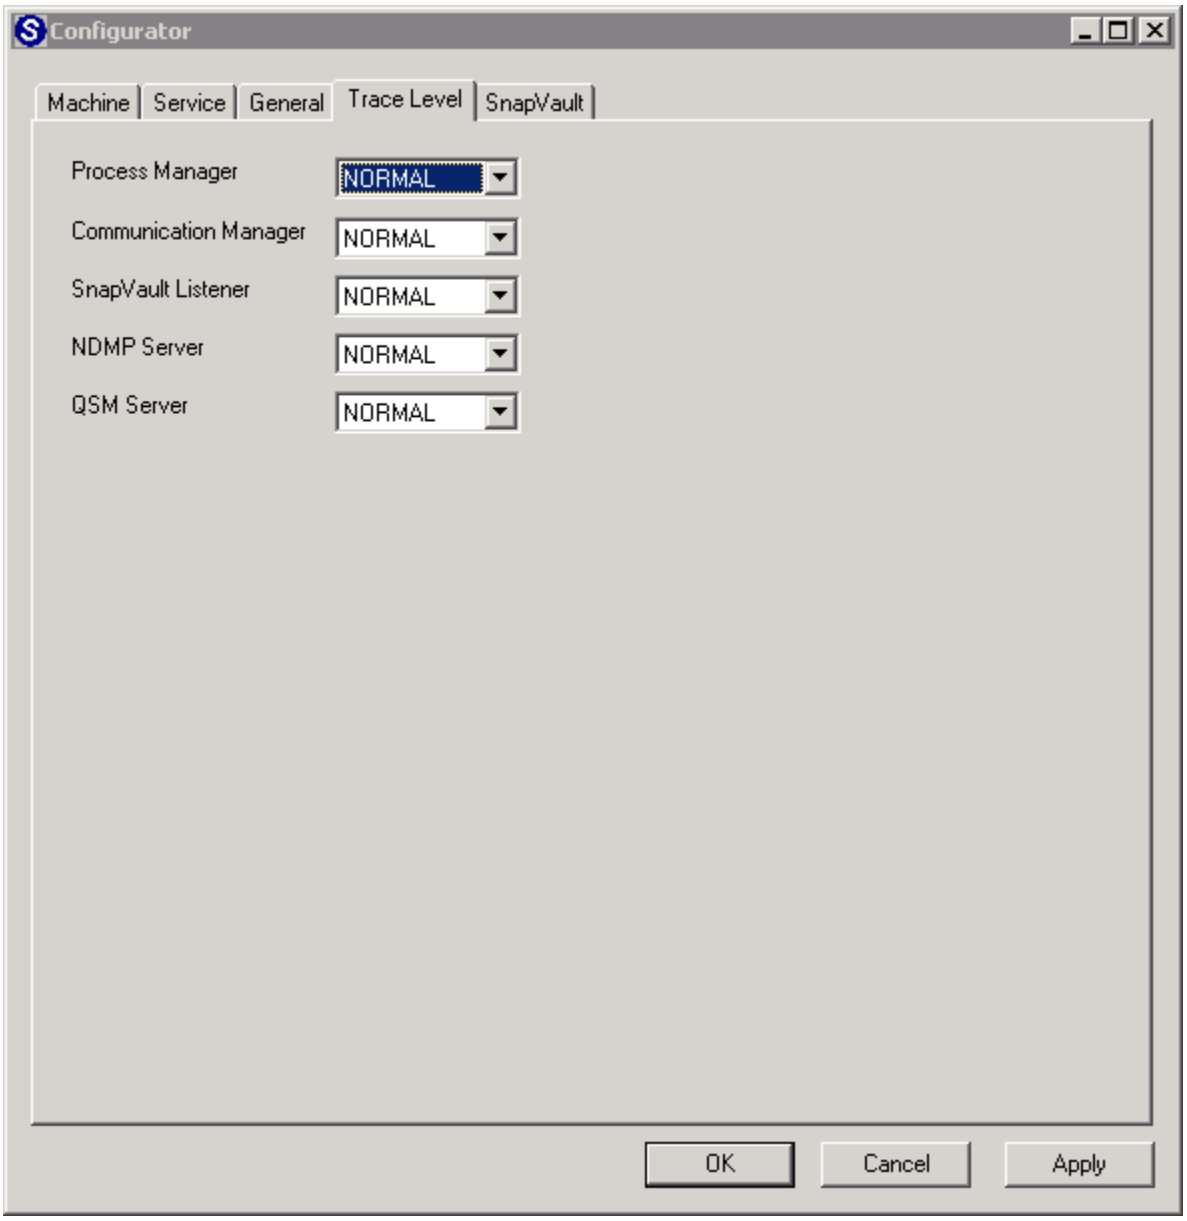

Figure 6) svconfigurator, Trace Level tab.

Use the Trace Level tab to modify default logging output for the OSSV processes. These settings are modified when Generate Debugging Files is selected on the General tab.

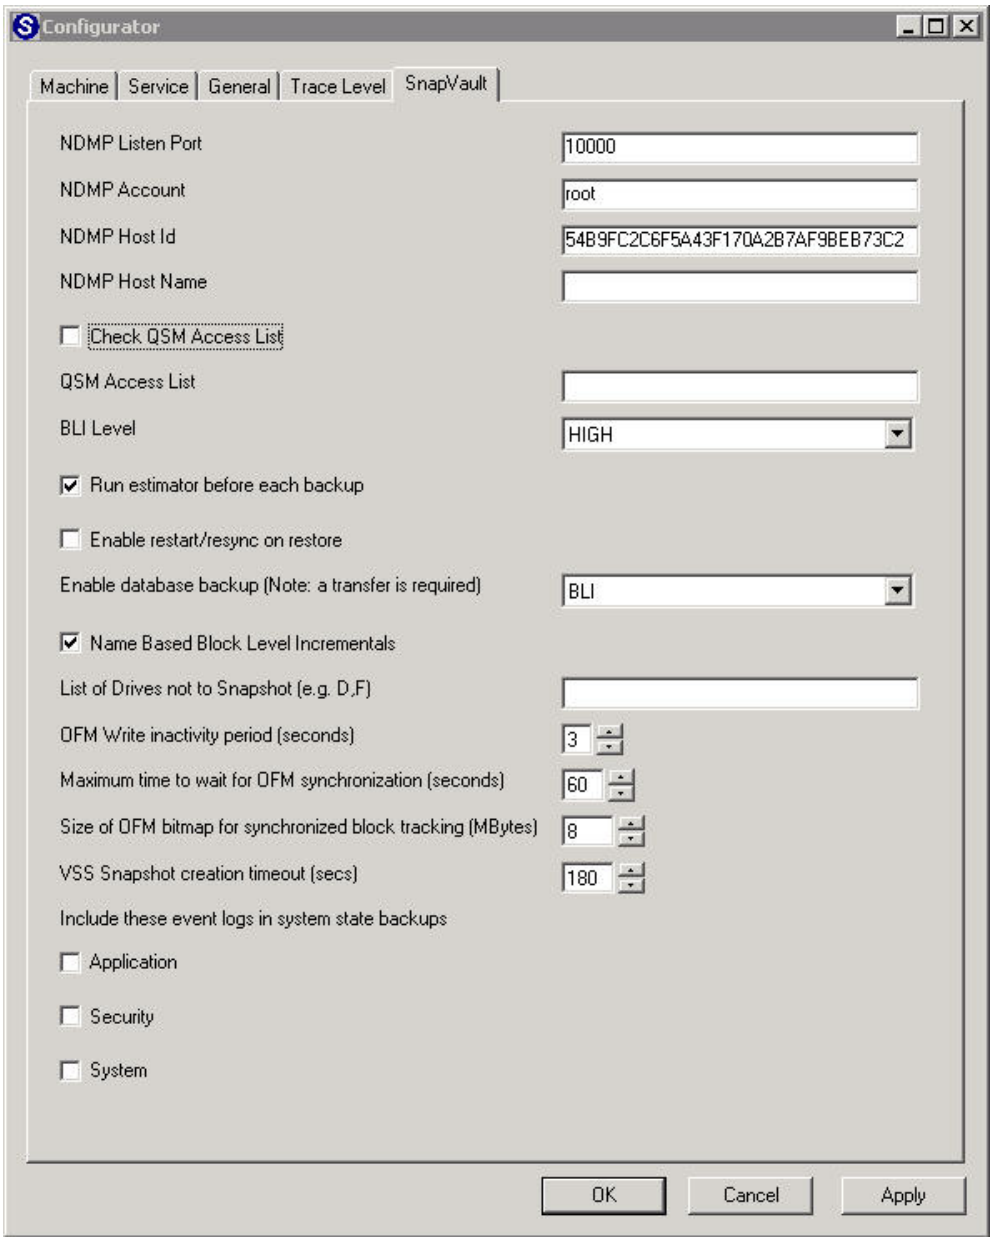

Figure 7) svconfigurator, SnapVault tab.

Use the SnapVault tab to modify multiple OSSV configurations. From this location, you can modify BLI level, OFM default parameters, NDMP parameters, VSS parameters, and security settings (`Check QSM Access List`). When you select `Check QSM Access List`, you must specify the systems (name or IP) to which you are allowing this particular primary system to back up.

Note: Any time a parameter is modified, click Apply to apply settings. It is also often necessary to stop and start the OSSV service after making a change (use the Service tab). For more information about these settings, see the OSSV release notes, located on the NOW site.

7.7 BINARIES

```
C:\Program Files\NetApp\snapvault>dir
 Volume in drive C has no label.
 Volume Serial Number is D8AF-7547

 Directory of C:\Program Files\NetApp\snapvault

03/10/2006  11:31 AM    <DIR>          .
03/10/2006  11:31 AM    <DIR>          ..
03/10/2006  11:17 AM    <DIR>          bin
03/10/2006  12:02 PM    <DIR>          config
03/10/2006  11:17 AM    <DIR>          db
03/10/2006  11:17 AM    <DIR>          etc
03/10/2006  11:17 AM    <DIR>          installfiles
03/10/2006  11:17 AM    <DIR>          lib
03/10/2006  11:17 AM    <DIR>          packages
03/10/2006  11:17 AM    <DIR>          pit
03/10/2006  11:17 AM               812 RELEASEDEF
03/10/2006  11:31 AM                 0 RELEASEDEF.lck
03/10/2006  11:17 AM    <DIR>          replaced
03/10/2006  11:39 AM    <DIR>          tmp
03/10/2006  11:17 AM    <DIR>          trace
03/10/2006  11:17 AM    <DIR>          util
               2 File(s)            812 bytes
              14 Dir(s)  73,699,995,648 bytes free
```

Figure 8) OSSV binaries.

A set of binaries is installed with the primary OSSV agent. These binaries are available to perform various tasks to maintain a successful OSSV environment. Tasks include modifications of core settings, NDMP password changes, checking the health of the installation, OSSV database backup, stopping and starting services, updating the parameters typically modified in `svconfigurator` at the command line, and so on. Most of these binaries can also be executed by using `svconfigurator`. Some environments (typically UNIX shops) do not allow `x` sessions or GUI management, and command line is necessary.

```
C:\Program Files\NetApp\snapvault\bin>svinstallcheck
SnapVault home directory: 'C:/Program Files/netapp/snapvault'
SnapVault database directory: 'C:\Program Files\netapp\snapvault\db'
SnapVault temporary directory: 'C:\Program Files\netapp\snapvault/tmp'
SnapVault Database and Temporary directories have 7% space left (2727Mb)
SnapVault service is running
SnapVault listener is running
NDMP Server, on port 10000, details:
    Vendor      Netapp
    Product     SnapVault
    Version     2_6_2008MAR10
    Host        vmcoe-mgmt
    Host Id     54B9FC2C6F5A43F170A2B7AF9BEB73C2
    OS Type     Windows 2003
    OS Version 5.2
SnapVault QSM Server is responding correctly
Validating filesystems:
   Drive 'A:\' is removable, unsuitable for backup
   Drive 'C:\' is suitable for backup
   Drive 'D:\' is a CDROM, unsuitable for backup
   Drive 'E:\' is suitable for backup
   Drive 'Y:\' is removable, unsuitable for backup
   Drive 'Z:\' is removable, unsuitable for backup

Check Succeeded
```

Figure 9) svinstallcheck.

An important executable is `svinstallcheck`, otherwise known as the `HealthCheck` utility. This tool performs a quick sanity check of system details, file systems suitable for backup, NDMP authentication, and database and temporary space available. This file runs automatically at the end of a new installation or upgrade. NetApp recommends that this utility be run manually or via script on a regular basis to track the SnapVault database and temporary directories remaining space.

```
C:\Program Files\NetApp\snapvault\bin>snapvault
The following commands are available; for more information
type "snapvault help <command>"
abort            destinations     help             release
restore          status           service          diag
```

Figure 10) snapvault command.

To monitor the status of an OSSV backup or restore, use the `snapvault` command from the primary system. This command also allows you to release relationships that have been stopped on the secondary systems, freeing up the primary directory for backup to a different location. This command also provides t the important restore function. Follow the syntax rules displayed by using `snapvault <command> help` on the primary.

```
C:\Program Files\NetApp\snapvault\bin>svpassword
Password:
Password changed
```

Figure 11) svpassword.

To modify the NDMP password, use the `svpassword` executable.
Note: There is no confirmation of password, so type carefully.

7.8 ETC AND TRACE DIRECTORIES

The `$INSTALL_DIR\etc` and `$INSTALL_DIR\trace` directories are also important locations for various files and logs.

```
C:\Program Files\NetApp\snapvault\etc>dir
 Volume in drive C has no label.
 Volume Serial Number is B094-1125

 Directory of C:\Program Files\NetApp\snapvault\etc

06/30/2008  03:07 PM    <DIR>          .
06/30/2008  03:07 PM    <DIR>          ..
06/30/2008  03:07 PM             1,264 file-exclude.txt
03/10/2008  04:21 PM            11,845 license.txt
06/30/2008  03:07 PM                 4 nextpid.dat
06/30/2008  03:07 PM               237 OSSVInstallLog.txt
06/30/2008  03:07 PM             1,772 path-exclude.txt
06/30/2008  03:33 PM               138 snapvault
06/30/2008  03:07 PM                14 SnapVault process manager service.lck
               7 File(s)         15,274 bytes
               2 Dir(s)   2,856,054,784 bytes free
```

Figure 12) etc directory contents.

To find the log files associated with a primary, look in the `etc` directory for the `snapvault` log file. This file contains all backup information, restore information, and error information for the primary system. VSS and OFM information is also noted, as well as time stamps, file systems being backed up, and so on. The log file is automatically archived in this location also and is

created during the first transfer of the following day. Log files for previous days are named snapvault.*yyyymmdd*.

In addition to the log file, you will find the exclude list files (`path-exclude.txt` and `file-exclude.txt`) in this directory. As mentioned earlier, these files allow the exclusion of specific files, ranges of files, and full directories.

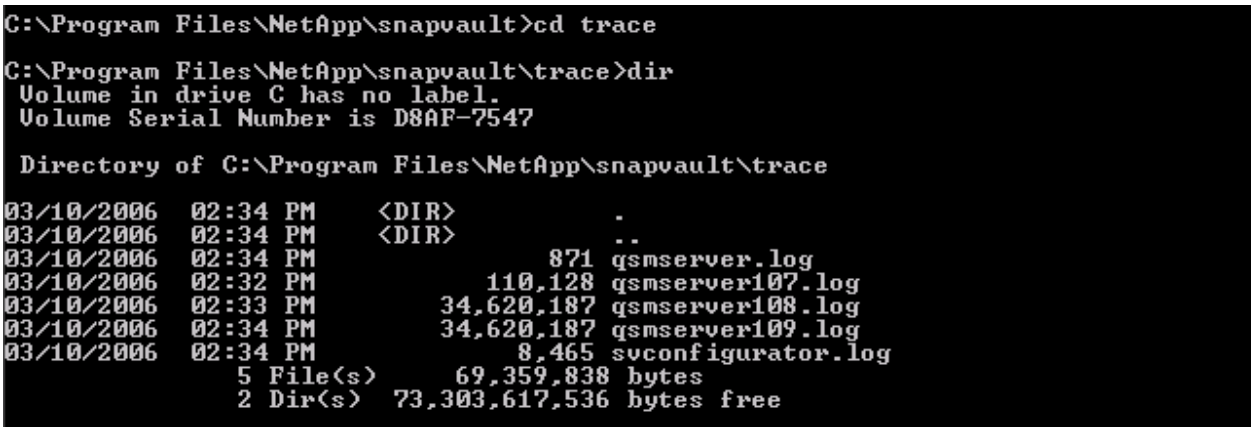

Figure 13) Trace directory.

The trace directory includes log files that were generated during debugging and troubleshooting. Do not enable debugging without the assistance of support. If debugging is enabled, be sure to set the debug level back to its default of Normal when you have finished troubleshooting. Multiple files can exist in this location—multiple files for one process or multiple files for many processes. These files can grow rapidly, so space consumption should be monitored.

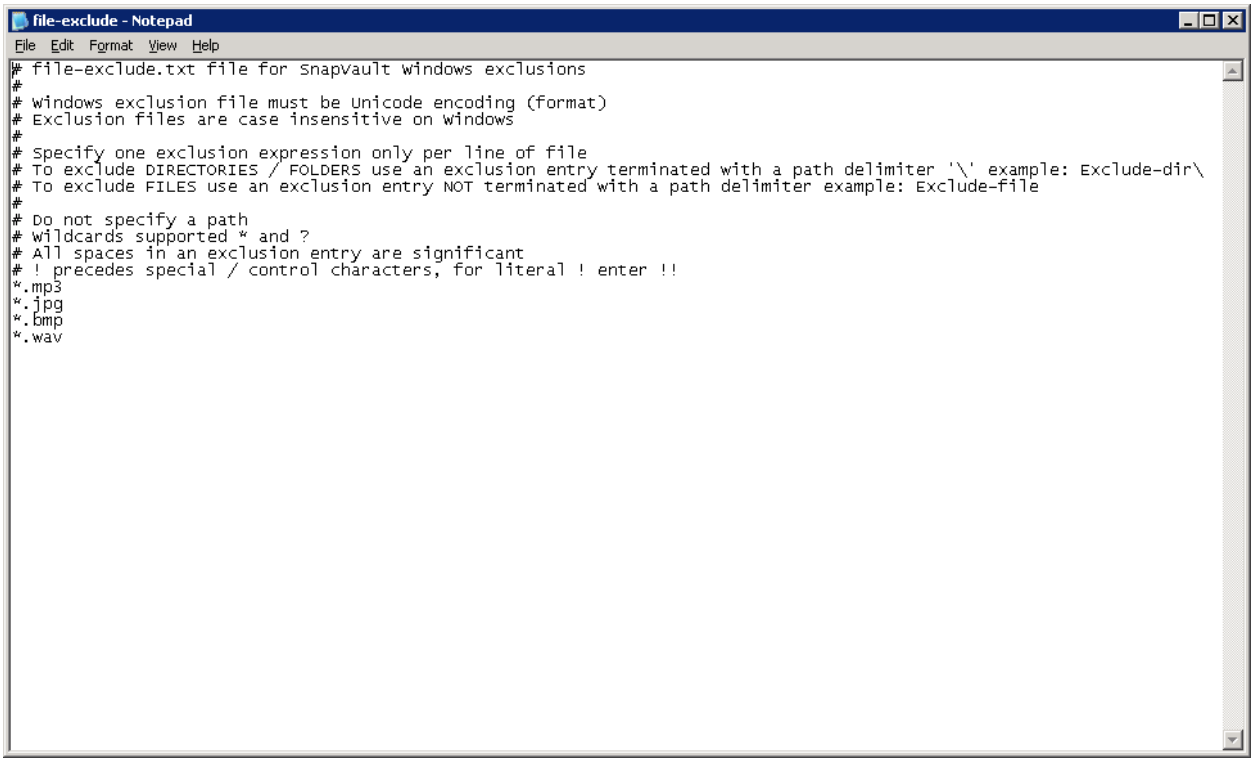

Figure 14) file-exclude.txt.

The exclude files contain various comments on syntax, wildcards, and allowed characters. In addition to the `file-exclude.txt` file displayed in Figure 14, a similar file, `path-exclude.txt`, allows you to exclude full paths and directories. Once an update is made to one or both of these files, the exclusion occurs on the next update. Pay close attention to these parameters and refer to the OSSV release notes, available on NOW, for specific information.

7.9 CREATING AN UNATTENDED INSTALL IMAGE

To help deploy OSSV over a large number of clients, OSSV 2.2 introduced Unattended Install. This section describes how to create an unattended installation on a Windows 2003 system.

Note: Once you create an unattended installation image on a Windows 2003 system, you can deploy it only to other Windows 2003 systems. If the environment also contains Windows 2000 systems, you must create a separate Windows 2000 unattended installation image.

To create an unattended installation image, you must go to a system that has OSSV installed and stop OSSV services by using the Service tab of the `svconfigurator` utility.

When the services have been stopped, use the `svconfigurator` utility to configure all the parameters you want your new installation to have, then close the `svconfigurator` utility.

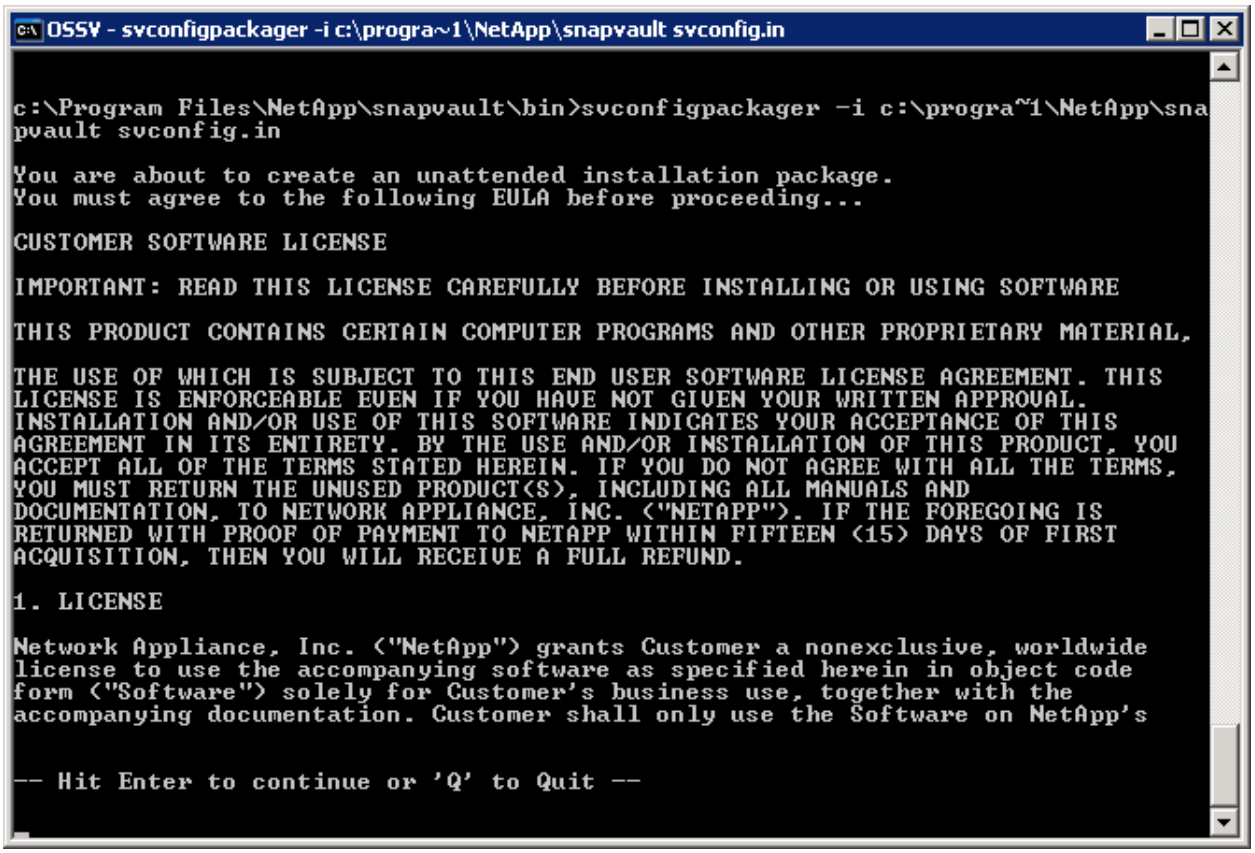

Figure 15) svconfigpackager.

Open a CLI window and change to the `$INSTALL_DIR/bin` directory. To create an installation script and save the configuration settings to a file, use:

```
svconfigpackager -i path_name filename
```

where `path_name` is the location of the destination for the rollout.

Note: If this is an upgrade, be sure to add the `-h` option, which honors the existing configuration parameters.

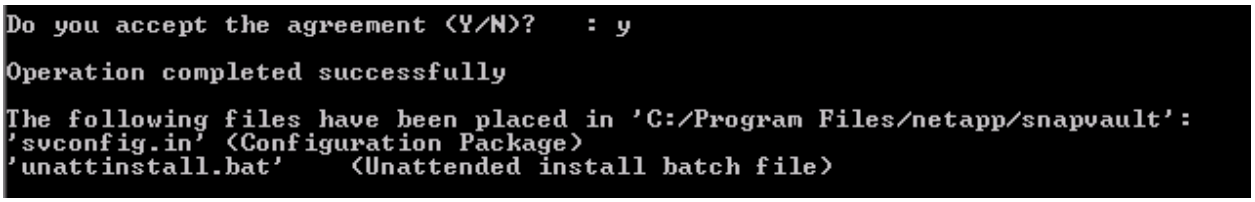

Figure 16) svconfigpackager, completed.

When you accept the terms of the license agreement, you will see that two files have been created, `svconfig.in` and `unattinstall.bat`. These files will be placed on each server, along with the OSSV installation files.

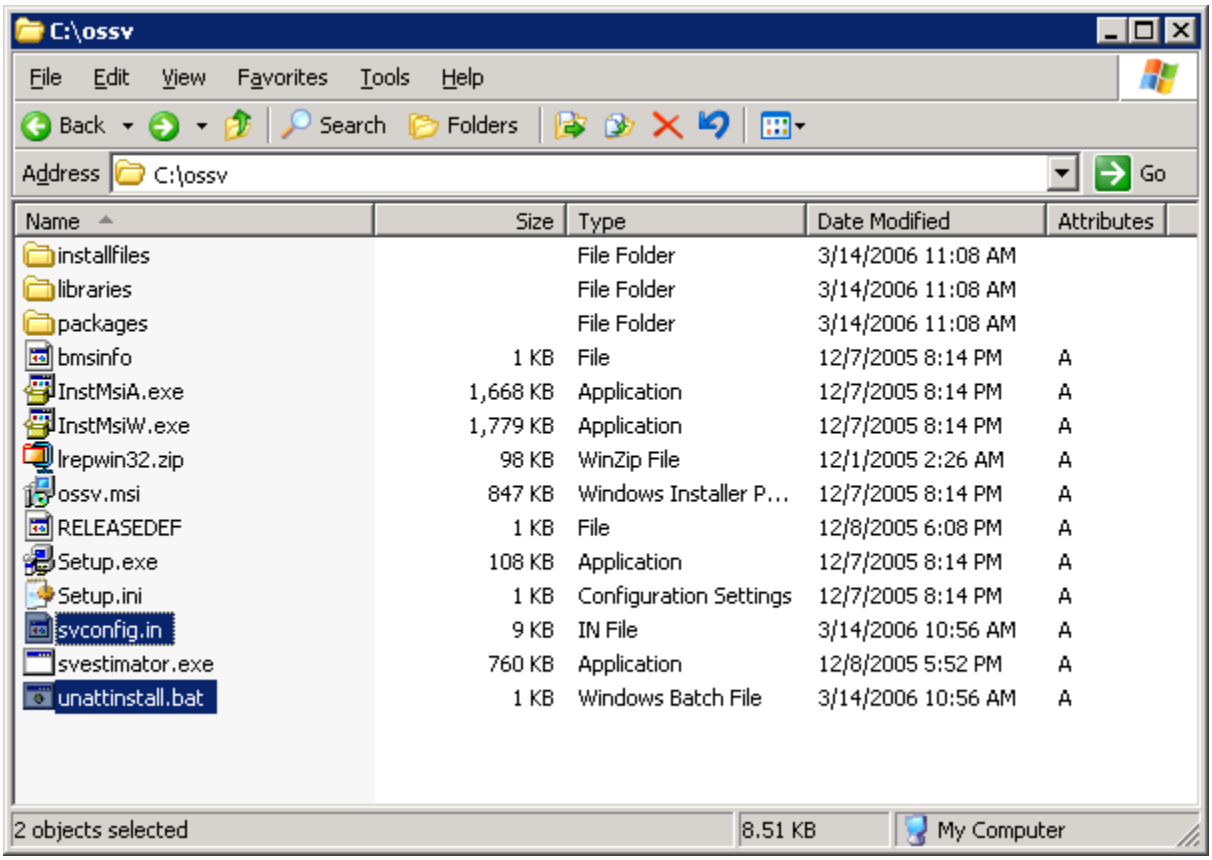

Figure 17) svconfigpackager files.

Here we have unzipped the OSSV package and placed the `svconfig.in` and `unattinstall.bat` files in `c:\ossv` on the server on which we want to perform the unattended installation.

```
C:\ossv>unattinstall.bat

C:\ossv>msiexec /i ossv.msi /qn targetdir="c:\progra~1\NetApp\snapvault" db_dir=
"C:\Program Files\netapp\snapvault\db" trace_dir="C:\Program Files\netapp\snapva
ult/trace" tmp_dir="C:\Program Files\netapp\snapvault/tmp" reboot=ReallySuppress
 UNATTENDED_INSTALL=1 HONOR_EXISTING_CONFIG=0 CONFIG_FILE=svconfig.in

C:\ossv>
```

Figure 18) unattinstall.bat.

When all the files have been copied, open a command window and issue the `unattinstall.bat` command. This reads the configuration file (`svconfig.in`) and installs OSSV on the server without any user interaction. At the end of the installation, the script automatically runs `svinstallcheck` to verify that the installation was successful. If the installation fails, log files are generated and are logged in `%SystemRoot%\Documents and Settings\`*`Current User`*`\Local Settings\Temp`.

7.10 CREATING A BASELINE RELATIONSHIP

When all the previous sections have been successfully completed and reviewed, review the OSSV release notes again to make sure that no settings or recommendations were missed. Now you can create a relationship and initiate the baseline transfer, which allows you o begin scheduled block-level or file-level updates or incrementals moving forward.

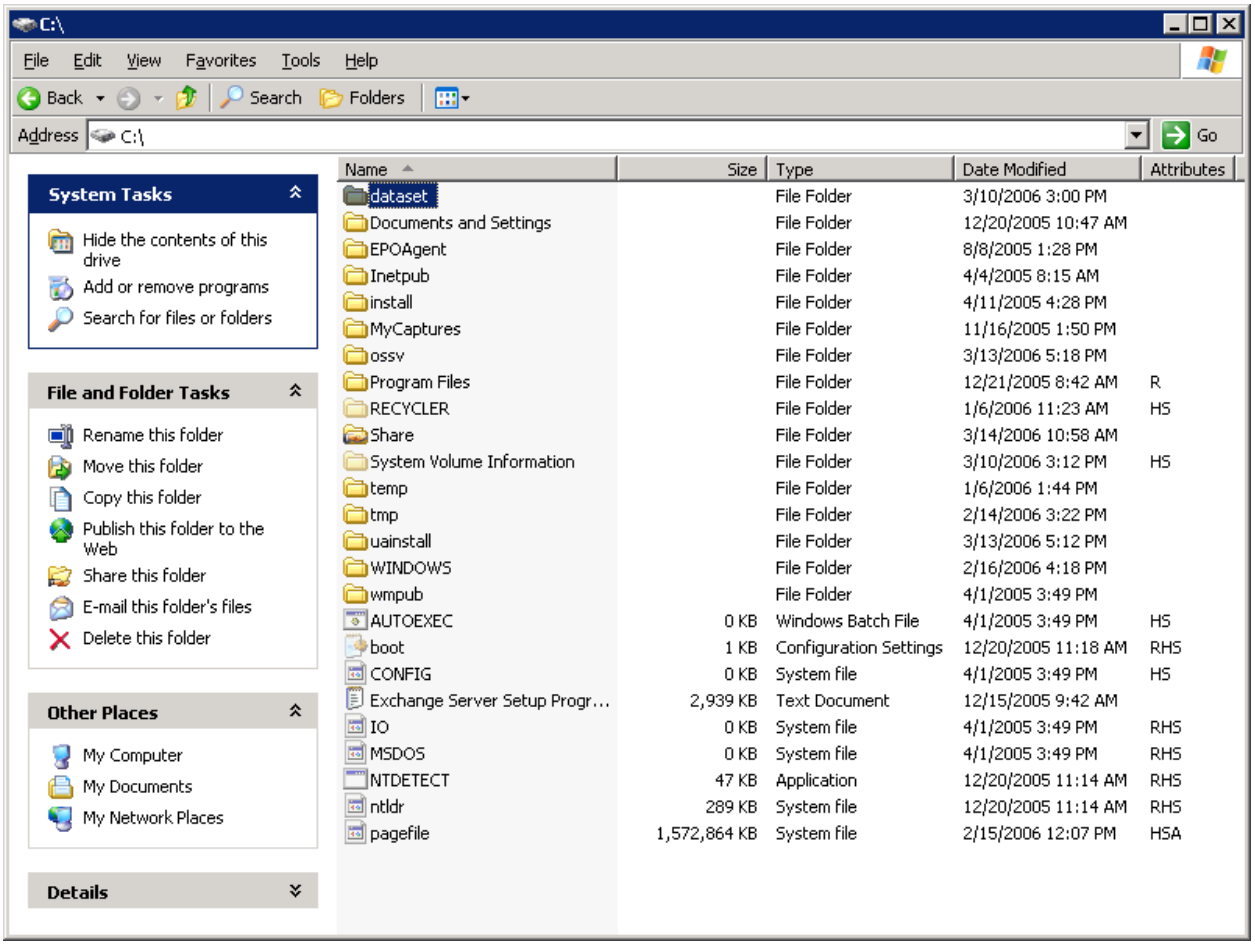

Figure 19) Data set to be backed up.

Select the data set to back up and then create a meaningful name for the secondary qtree. In this example, we are backing up `C:\dataset`, which happens to be a 1.5GB file system.

```
r100-rtp01> snapvault start -S 10.61.132.69:c:\dataset /vol/ossv_flex/dataset
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
r100-rtp01>
```

Figure 20) Running snapvault start from secondary.

Log in to the secondary system and issue the baseline creation command, `snapvault start` (see the man pages for detailed syntax and options). As shown in Figure 20, relationship is backing up `C:\dataset` from the host 10.61.132.69 to the `/vol/ossv_flex/C_dataset` qtree located on r100-rtp01.

```
r100-rtp01> snapvault status -l /vol/ossv_flex/dataset
Snapvault secondary is ON.

Source:                  10.61.132.69:c:\dataset
Destination:             r100-rtp01:/vol/ossv_flex/dataset
Status:                  Transferring
Progress:                488456 KB
State:                   Uninitialized
Lag:                     -
Mirror Timestamp:        -
Base Snapshot:           -
Current Transfer Type:   Initialize
Current Transfer Error:  -
Contents:                Transitioning
Last Transfer Type:      -
Last Transfer Size:      -
Last Transfer Duration:  -
Last Transfer From:      -
r100-rtp01>
```

Figure 21) snapvault status –l.

NetApp recommends monitoring the transfer with the `snapvault status -l` command (long listing). From this output, you can view bytes transferred, time, type of transfer (baseline, update), age of backup (lag), and so on. Key items to take away from this output are `Progress`, `Current Transfer Type`, `Status`, and `Lag time`.

`Progress` is the amount of data (kB) that has been transferred at the time the `snapvault status -l` command was issued. This output is generally higher than the actual amount of data residing on the primary system, due to overhead associated with SnapVault (approximately 16kB for every file updated).

`Current Transfer Type` displays whether this is a baseline, (`Initialize`) transfer, an incremental (`Replica`) transfer, or a `Restore` transfer. As always, refer to the Data ONTAP man pages for details.

`Lag` is the amount of time that has expired since the last successful backup. Basically, a session with `Lag = 23` hours means that the session displayed is 23 hours old, or 23 hours behind the current time. This is normal if nightly backups are being performed; however, if hourly backups are being performed, this reveals that the backups are not occurring every hour and haven't occurred for the last 23 hours.

```
r100-rtp01> snap list ossv_flex
Volume ossv_flex
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  Mar 10 15:04  r100-rtp01(0033604316)_ossv_flex-base.0 (busy,snapvault)
r100-rtp01>
```

Figure 22) snap list.

When the transfer is complete, the secondary creates a baseline Snapshot copy, which is a softlocked Snapshot copy. All subsequent updates are named according to the `snapvault`

`snap sched` input. Snapshot copies accumulate until the retention number specified in the `snapvault snap sched` is reached.

7.11 SCHEDULING OSSV BACKUPS VIA THE SECONDARY SYSTEM

A good rule of thumb is to specify all primaries that have the same class of data for protection to be directed to the same volume on the secondary server. So, for example, data on primaries that need to be backed up every night might go to `/vol/backup_nightly`. Data on primaries that need to be backed up twice a day might go to `/vol/backup_twice_daily`.

The `snapvault snap sched` command is used for scheduling with Data ONTAP (you can choose to schedule by using a supported NDMP tool as well). `snapvault snap sched` is available on the secondary system. This command sets, changes, or lists Snapshot schedules. If no schedule argument is given, the command lists the currently configured Snapshot schedules.

There are two options for scheduling: `-x` is the transfer schedule; `-c` is the create schedule (the default). The `-x` option tells OSSV to *transfer* new data from all primary qtrees residing in that volume before creating the Snapshot copy. If `-x` is not specified, only local Snapshot copies on the secondary are created, without any communication with the primary. To move changed data from the primary on a schedule (incrementals), the `-x` option must be added to the `snapvault snap sched` command input.

```
R100> snapvault snap sched -x backup sv_daily 5@Mon-Fri@23
```

where the schedule itself is:

```
cnt[@day_list][@hour_list]
OR cnt[@hour_list][@day_list]
```

This command schedules incremental backups every weekday at 2300 hours on the volume named `backup`. When Snapshot copies are created and begin to accumulate (toward the specified `cnt` or retention policy) on the secondary for a particular volume, they are numbered from oldest to newest, from 0 to `cnt`-1. When creating a new Snapshot copy, the SnapVault process on the secondary system deletes the oldest Snapshot copies, increments by 1 the number on the remaining Snapshot copies, and finally creates a new 0 Snapshot copy.

For more information about this command, see the SnapVault section of the *Online Backup and Recovery Guide* and the Data ONTAP man pages.

7.12 RECOVERING OSSV DATA BY USING THE COMMAND LINE

If it is possible to NFS mount or CIFS map the secondary volume on the OSSV client, recoveries can be performed by simply copying SnapVault data from the secondary volume. Permissions on SnapVault data are the same as permissions on user data, so the same authentication rules apply.

If it is not possible to mount or map the secondary volume, you can use `snapvault restore`, which is part of the OSSV distribution.

All command-line actions to restore data to a primary are initiated from the primary. Log on to the primary to begin restoring data:

```
C:\>cd Program Files\netapp\snapvault\bin

C:\Program Files\netapp\snapvault\bin>snapvault restore -s sv_daily -S
r100-rtp01:/vol/ossv/my_documents C:\Temp\restored_my_documents
```

The text following `-s` is the Snapshot copy name. The text following `-S` is the name of the secondary server, followed by the location of the backed-up data on the secondary volume. The final argument is the location on the primary where the file is to be restored. Note that the file can be restored to a different location on the primary under a different name. The command line is case sensitive. When files are located multiple subdirectories deep on a file system, be careful with this command, because it is sensitive. Single file restores can be issued by using `snapvault restore`.

Note: The root primary file system name that was being backed up is essentially replaced with the qtree name when performing a restore by using this command.

8 TROUBLESHOOTING

8.1 COMMON TROUBLESHOOTING TASKS

When troubleshooting OSSV issues, it's important to check the status of a few commands, such as the `snapvault status -l` command, which provides a detailed list of the relationship and also identifies any relevant error codes. When the error code is noted, the `snapvault diag` command can be used to get further information about the error. In addition to running the `snapvault` commands, the logs on both the NetApp storage system and the OSSV client should be checked for additional information. All OSSV data is stored in the SnapMirror log on the NetApp storage system and the SnapVault log on the OSSV client.

8.2 OSSVINFO

When opening a case with technical support, a common set of files and output is required. In addition, the executable `OSSVINFO.exe` or `OSSVINFO.sh`, located on the internal engineering pages, is available. This executable automatically executes the following commands and obtains the appropriate log files. `OSSVINFO.exe` is supported on Windows 2000 and Windows 2003, and `OSSVINFO.sh` is supported on Solaris, Linux, IRIX, HP-UX, and AIX systems that have the OSSV agent installed.

Windows OSSVINFO Syntax

```
OSSVINFO.exe [ -s secondary ] [ -l username:password ] outfile.txt
```

UNIX OSSVINFO Syntax`OSSVINFO.sh [-s filer [-l user[:passwd]]]`

The output is saved to a file in the form `ossvinfo-%Y%m%d%k%M%S.log`.

This executable can be run from a Windows command shell or a UNIX command prompt. When reporting problems seen in Data ONTAP, always use the `-s` option to collect information specific to Data ONTAP.

8.3 SECONDARY SYSTEM LOGS

Obtain all the relevant log files from both of the OSSV secondary systems:

- `/etc/log/snapmirror`
- `/etc/messages`

8.4 PRIMARY SYSTEM LOGS AND DATA

Obtain all the relevant log files and system information from the OSSV primary systems.

- `$INSTALL_DIR\etc\snapvault` or `$INSTALL_DIR/etc/snapvault`
- Properties of the file system or drive being backed up
- Properties of the file system or drive containing the OSSV agent installation:
  - On Windows:
    My Computer > right-click the drive icon > Properties > General
  - On UNIX:
    `df -k <directory>`
- Output from `/etc/{v}fstab`
- Error messages seen on the console, if any

8.5 GENERATING DEBUG INFORMATION

When troubleshooting OSSV cases, you may be asked to generate debug information pertaining to specific OSSV processes running on primary systems. Use the following steps.

SET UP DEBUG INFORMATION

Before performing the following procedure, make sure that no updates or transfers are occurring or will occur while debug is enabled.

1. Open `svconfigurator` (Start > Run, or use the command line).
2. On the General tab, select the Generate Debugging Files checkbox.
3. On the Trace Level tab, select Verbose from the SnapVault Listener drop-down list.
4. Click Apply.
5. On the Service tab, select Stop Service.
6. Wait for Current State to display Stopped. .
7. Go to the `$INSTALL_DIR\trace` directory (if present). If there are any files there, delete all of them; these are old debug files.
   Note: Deleting old files frees space; these debug files can grow large quite quickly.
8. Click Start Service.
9. Wait for Current State to display Running.
10. Click OK to close `svconfigurator`.

COLLECT DEBUG FILES

1. Open a command prompt (Windows) or a shell (UNIX).
2. Change the directory to the `$INSTALL_DIR\bin` directory.
3. Run the `snapvault status` command to generate the necessary debug files.
4. Open `svconfigurator`.
5. On the Service tab, click Stop Service.
6. Wait for Current State to display Stopped.

DELETE DEBUG FILES AND DISABLE DEBUG

Delete debug files as soon as enough information has been collected. Turn debug off by reversing the procedure described in “Collect Debug Files.” This is important, because the default location for the trace directory is under the default `$INSTALL_DIR`. A file system can quickly fill up if debug is left on.

9 REFERENCES

Read the following documents to gain a better understanding of SnapVault and OSSV.

- Technical reports
  - TR-3234: Leveraging NetApp SnapVault for Heterogeneous Environments
  - TR-3487: SnapVault Best Practices Guide
  - TR-3252: Enhancing Heterogeneous Backup Environments with SnapVault
- Manuals
  - Data ONTAP 7.x Online Backup and Recovery Guides
  - Protection Manager Administration Guide
  - OSSV 2.x release notes, available on NOW
  - OSSV Installation and Administration Guide

10 APPENDIX: MODIFYING DATA OF AN OSSV DESTINATION

Because the Open Systems SnapVault destination is read-only, there are two methods that you can use to modify the data on the secondary system. The first method is to use the SnapVault/SnapMirror bundle, making the OSSV destination a SnapMirror destination. This method allows the data to reside in the existing qtree, and no extra volumes need to be created. The second method is to use the FlexClone® technology.

9.1 USING THE SNAPMIRROR/SNAPVAULT BUNDLE

SnapVault does not currently have the ability to create a writable destination on the secondary system. However, you can use SnapMirror to convert the Open Systems SnapVault destination to a SnapMirror destination, making it a typical SnapMirror destination that can be quiesced and broken. To use this bundle, you must be running a version of Data ONTAP that supports the resync after restore/break feature introduced in OSSV 2.2. In addition, the secondary system must also have a SnapMirror license enabled.

CONVERTING AND MAKING THE SECONDARY READ/WRITE

To convert an OSSV or SnapVault secondary backup destination to a usable/writable destination (typically for DR situations), perform these steps on the secondary storage system.

1. Turn SnapMirror and SnapVault off.
2. Switch to privileged mode (`priv set diag`).
3. Convert SnapVault qtree to SnapMirror qtree (`snapmirror convert <sec_qtree_path>`).
4. Turn SnapMirror on.
5. Quiesce the qtree.
6. Break the mirror, making it writable.
7. Turn SnapVault on.

REESTABLISHING THE RELATIONSHIP

With OSSV, there is currently no mechanism to propagate any changes made while the qtree is in a read/write state. If it is determined that the same destination qtree needs to be used, then all changes to the qtree are lost. The only other option is to leave the qtree as is and perform a new baseline of the source volume.

Resync the secondary qtree:

```
snapvault start -r -S <pri_system>:<pri_path> <sec_qtree_path>
```

Note: Resync with OSSV (`snapvault start -r` command) is not yet supported. For the latest versions of Data ONTAP that are supported, see KB12138.

9.2 USING FLEXCLONE VOLUMES

FlexClone volumes are a point-in-time, writable copy of the parent volume (OSSV destination). Changes made to the parent volume after the FlexClone volume is created are not reflected in the FlexClone volume. To create a FlexClone volume, the OSSV destination must be a flexible volume and the FlexClone license must be installed.

CREATING A FLEXCLONE VOLUME

To create a FlexClone:

Secondary: `vol clone create clone_vol -b parent_volume [parent_snap]`

If no parent Snapshot is named, Data ONTAP creates a new base Snapshot.

Note: The base Snapshot cannot be deleted as long as a clone of the parent volume exists.

After the FlexClone has been created, you can use CIFS/NFS to mount the FlexClone volume to the host with the data in a readable/writable format.

Note: Any changes made to the FlexClone volume cannot be propagated back to the OSSV primary. The only way to make the changes in the FlexClone volume is to manually copy the files that were created or modified.