UGS Teamcenter Engineering

Deployment and Operational Considerations for NetApp® Filers

Dennis O'Brian, Bob Brandenstein, Bill Halpin, Rajesh Godbole & H.T. Sun | UGS PLM Solutions Inc. and Network Appliance | March 2005 | TR-3171

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Table of Contents

1. Introduction

UGS's product life cycle management (PLM) solutions represent a unified approach to extended enterprise collaboration that enables all participants in your product life cycle to work in concert as you bring products to market and support your customer base. Teamcenter Engineering (TcEng) is an industry-driven, customer-proven Web-centric collaborative product life cycle management system for the engineering enterprise. Teamcenter Engineering (formerly iMAN) provides distributed engineering/manufacturing teams with the vaulting, global sharing, and workgroup management capabilities they need to capture, manage, and leverage geometry and engineering data created by multiple CAD, CAM, and CAE systems. The backup and recovery of this critical data in the event of hardware failure or data corruption is a significant design and operational element of a TcEng environment. Network Appliance's storage solutions offer capabilities called Snapshot™ and SnapRestore® that provide a high-performance and reliable mechanism to create an almost instantaneous backup of the database and data volumes. These capabilities provide an attractive solution set for TcEng environments and should be a key design criterion for future deployments.

Network Appliance™ storage appliances called Filers integrate seamlessly with Teamcenter Engineering into a PLM collaborative engineering environment. Filers are reliable and provide excellent performance, scalability, and data availability. They also deliver native multiprotocol access to design data in a mixed-mode environment of UNIX® and Windows® clients. This provides high-performance access to a single copy of the data, which is shared across all types of clients.

This deployment guide specifically describes how Teamcenter Engineering infrastructure can be configured with NetApp filer storage appliances in a heterogeneous client environment for use of Snapshot and SnapRestore. This deployment guide is a supplement to *iMAN Collaborative Solutions Deployment Guide*, available with the product CD.

2. Teamcenter Engineering PLM Deployments with NetApp Filers

The challenge is to maximize the time-to-market benefits of concurrent engineering while maintaining control of your data and managing access by the people who need it—when they need it. UGS PLM Solutions' Teamcenter Engineering solves this problem by storing the master data as "datasets" in a secure "vault." There, its integrity can be assured and all changes to it controlled, monitored, and recorded.

Many users create or access data directly in the TcEng vault from an application that has been integrated with TcEng (e.g., UG, CATIA...). When an application has been integrated this way, changes to the master data (revisions) are recorded and tracked in the vault each time the user saves his or her work. This way, the vault contains the latest working version, ensuring security of even work-in-process. Reference copies of the master data, on the other hand, can be retrieved by authorized users in various departments for design, analysis, review/approval, or revision. Data sets retrieved for revision can optionally be checked out, which locks them for revision by other users until checked back in. Revised data sets can then be imported back into the vault by authorized users.

A revised data set does not replace the original. Instead, it is stored as a new version, keeping the original version intact. Any number of data set versions can be managed by TcEng, allowing recovery of the information at any stage in its development.

Geometry data, documents, and any other intellectual property are managed in TcEng as revisable "items." Items can have any number of revisions, and each revision can have any number of data sets. Each data set contains one or more related files; for example, there may be multiple representations of the part (3D, jpeg, gif) in the data set. Information about the data sets, their metadata, is managed in an Oracle® database. The metadata includes the locations of data sets in the vault, which consists of one or more file system volumes.

The data management capabilities of the underlying storage infrastructure are critical to the success of a Teamcenter Engineering environment. The following bottlenecks in accessing, sharing, and archiving product design and production information files need to be overcome:

- Customers typically operate TcEng in mixed environments with both Windows and UNIX users accessing data sets controlled by TcEng
- TcEng manages large data sets that can cause significant backup windows that require taking both the data and the network out of operation for hours. This can lead to a substantial loss of productivity for the end users
- And, in the event that data sets need to be recovered from tape due to data loss or corruption, data recovery can take hours or days, again stopping productivity cold

These major challenges can be readily met with the deployment of NetApp technologies in Teamcenter Engineering's PLM environment:

- Multiprotocol eliminates the cost of acquiring separate file servers that are CIFS aware as enterprises start to harness the power of cheaper Windows PCs in a homogeneous UNIX environment where NFS is exclusively adopted
- File systems that directly support multiple protocols also tend to perform better than third party software emulations of those protocols, such as SAMBA
- Snapshot technology drastically reduces the time of major database and data set backup from hours to a matter of seconds
- SnapRestore minimizes the impact of system downtime by ensuring instant and smooth data recovery due to file corruption or hardware failure

NetApp's Native Multiprotocol File Service

Typically, enterprises have environments that include both Windows and UNIX platforms. It is important for these organizations that both the Windows and UNIX platforms share the same data in order to eliminate unnecessary proliferation of data through duplication, and to reduce the resulting data management complexity.

This is a key challenge in that secure access to the same design data must be granted to both sets of clients, which use fundamentally different security mechanisms. Additionally, integrity of the data in the vault must be maintained. NetApp filers allow multiprotocol access to design data, which allows heterogeneous access to the same data using NFS and CIFS protocols.

To provide these features NetApp filers rely on the Data ONTAP™ microkernel and WAFL® file system. These were designed specifically to provide storage services to support NFS, CIFS, DAFS, HTTP, and other file service protocols. Therefore there are no functionality gaps across protocols as with the emulated approaches. Further, kernel-based security and file-locking enforcement are inherently stronger than user-space application software methods.

Backing Up with NetApp Filers:

In practice, the Oracle database is fairly small (5 to 50GB), depending on the data managed,

number of users, and other implementation details, while the data volumes can be expansive (thousands of gigabytes). The data volumes and their corresponding database must be backed up on a regular basis to ensure the TcEng installation can recover from errors, power failures, or other significant disasters. Since the location of data set files in the volumes is managed in the Oracle database, it is critical that the Oracle backup and the volume backup agree—that they are synchronized.

To ensure database/volumes synchronization, backups require the process of quiescing the user environment to ensure the database and data volumes are complete and consistent. Then they can be safely backed up using backup media, such as DLT, for offline or off-site storage. The backup process can be time-consuming depending on the size of the volumes and technology selected. It is not uncommon for this process to take several hours or even overnight to complete, during which time the TcEng PLM system is unavailable. This limits implementations that require 24x7 operation.

The use of NetApp filers and their Snapshot technology offers a significant improvement in this process without sacrificing reliability, data integrity, or data recovery. Once the system is quiesced, a complete backup image of the database and volumes can be accomplished in less than a minute, thereby allowing the system to be returned online and available to users. Once the Snapshot has been taken, normal backup procedures can be used to move this image to tape for archival purposes while the real data is available to users.

Beginning in TcEng version8.0, APIs allow the NetApp filer backup process to be integrated with the system so that users can continue to work even while the Snapshot is being taken. Integrated this way, the backup process can notify TcEng that a Snapshot is about to begin. TcEng then sets all data access to read-only, so that no volume information (e.g., file locations) is changed. When the Snapshot completes, the backup process can notify TcEng, which then allows data in the volumes to be revised (i.e., new data sets created).

Snapshots

Designed into the Data ONTAP kernel and WAFL file system, NetApp's Snapshot technology allows creation of point-in-time copies of Teamcenter Engineering database and data volumes very quickly. These Snapshots can be utilized as a consistent, read-only source for backups and can be used as low-overhead online copies for easy recovery from potential data loss. Snapshots are only available on NetApp filers, and are extremely useful to both users and system administrators.

A Snapshot is a "frozen" (read-only) view of the file system created by preserving the pointers to all the disk blocks currently in use at the time of the Snapshot. After a Snapshot has been taken, changes to files are reflected in updates to the current set of pointers, no differently than if no Snapshots existed. Snapshots can be scheduled to occur automatically on a recurring basis. There can be up to 255 Snapshots at any one time. For more details about Snapshots please see Appendix A and also Section 2 of *NetApp Technical Report 3002* [TR-3002].

Snapshots incur no performance overhead and are designed to minimize disk space consumption. Snapshots are "invisible" to users under ordinary circumstances—except when the user chooses to see them. Pathnames for Snapshots can be slightly different for NFS versus CIFS access.

SnapRestore

The complement to NetApp Snapshot is SnapRestore. The SnapRestore software allows an enterprise to recover almost instantly from disaster scenarios. In seconds, SnapRestore quickly restores lost files using stored Snapshot copies and can recover anything, from an individual file

to a single home directory to a multiterabyte volume. NetApp's SnapRestore software makes recovering your data fast and easy so operations can be quickly resumed.

Efficiency of Snapshot and SnapRestore

NetApp's Snapshot technology makes extremely efficient use of storage by storing only block-level changes between each successive Snapshot. Since the Snapshot process is automatic and virtually instantaneous, backups are significantly faster and simpler. SnapRestore software uses Snapshot technology to perform near-instantaneous data restoration. In contrast, alternative storage solutions copy all of the data and require much more time and disk storage for the backup-and-restore operations.

Additional Benefits

With SnapRestore, data can be restored from any one of the Snapshots stored on the file system. This allows an application development team, for example, to revert to Snapshots from various stages of their design, or test engineers to quickly and easily return data to a baseline state. Restoring to the base environment takes only seconds, and the restored environment is identical to the point at which the Snapshot copy was created.

3. Solution Architecture

The Teamcenter Engineering architecture discussed in this document focuses on the location of data and system processes when used with a NetApp filer to optimize performance, allow both Windows and UNIX clients access to data, and make use of Snapshot and SnapRestore.

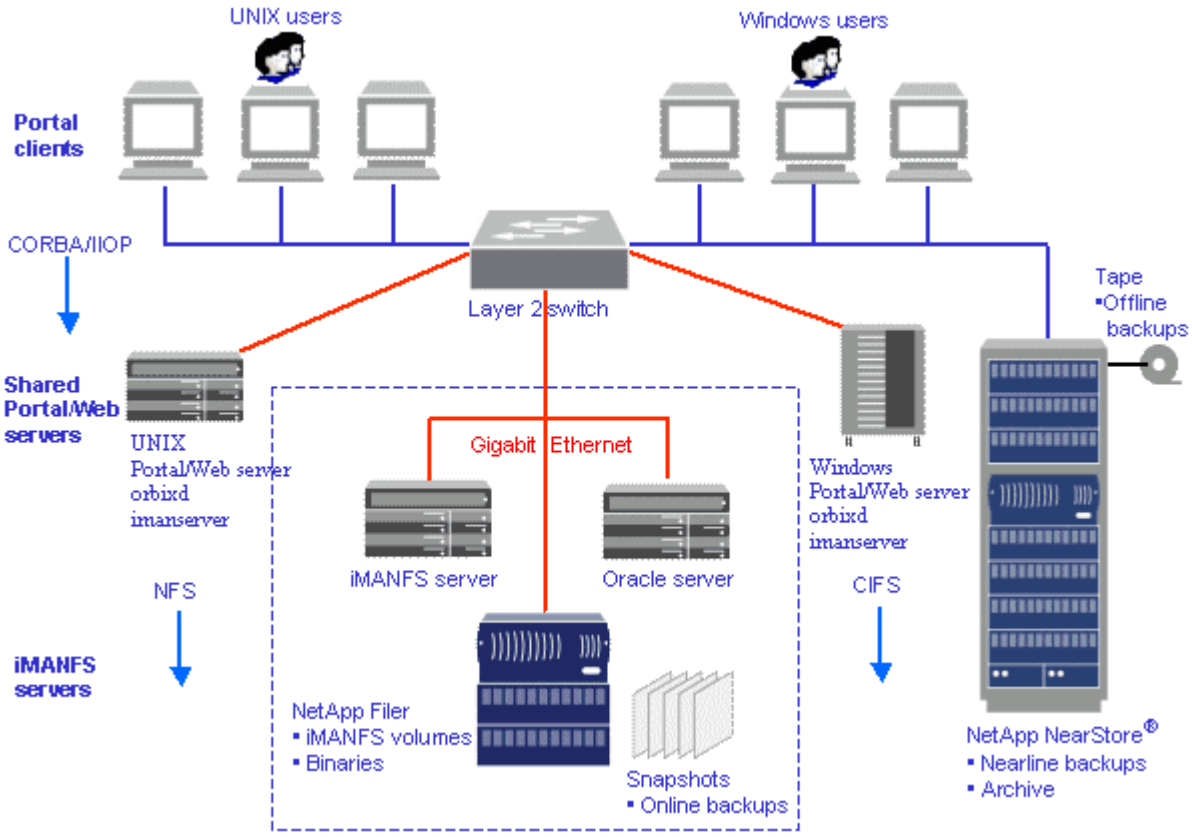

Figure 1. Teamcenter Engineering PLM Example Deployment with Filers

Network Appliance Inc.

Note:

- UNIX or Windows clients can connect to either UNIX or Windows applications servers.
- Portal servers must be able to connect to the filer file systems (i.e., mount for NFS).
- It is important to use UNC when specifying the Windows Path Name in the Volume Definition administration page.
- Traditionally, iMANFS runs on the file (or volume) server machine. When a filer is configured to hold the volume data, another server must be configured to run this lightweight process. It is possible to run the iMANFS server on the Oracle database server, although a separate machine is preferred.

3.1. Alternate Configurations

For portal clients, Teamcenter Engineering offers a configuration option called LSE or Local Server Edition. This configuration is recommended for larger implementations where server resources are limited, but it requires a more robust client configuration. In this configuration, the portal client and the portal server are both installed on each client. This removes the requirement for separate portal server hardware. The client machines must consist of a supported platform and have larger memory and disk requirements.

A typical installation would look as follows:

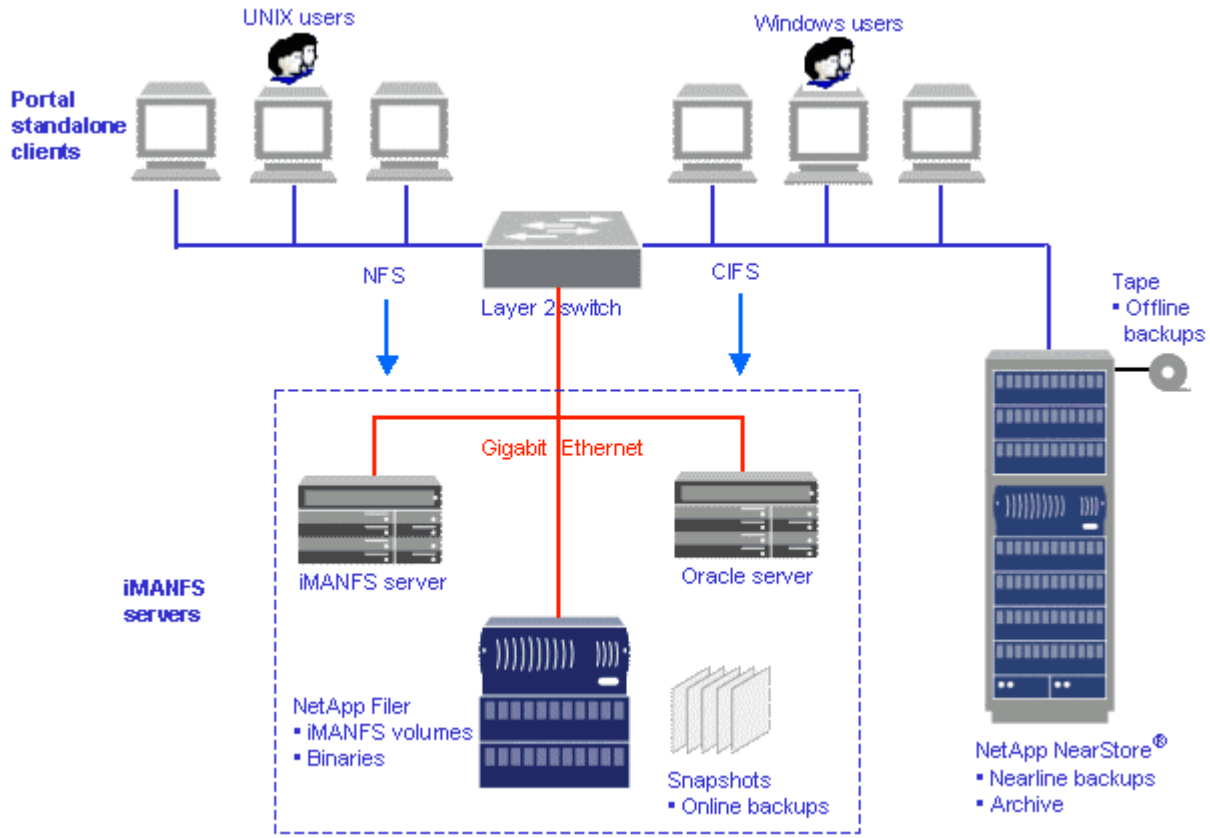

Figure 2. Teamcenter Engineering LSE Configuration Example

3.2. Configuration Considerations

In configurations that use NetApp filers, the location of the iMANFS daemon is a consideration, although not an extremely critical one. The iMANFS daemon is a very lightweight process and can be installed on any available server in the network that is accessible to the portal clients.

Many system administrators with NetApp filers have their environments set up as follows:

- Teamcenter Engineering volume data is sitting on a multiprotocol filer
- The Teamcenter Engineering applications are configured to use IMAN_Security_Level=3 in the `.iman_env` configuration file
- iMANFS daemon is running on a UNIX server (iMANFS running on a Windows server will be tested and documented in the future)
- Teamcenter Engineering application servers (i.e. portal servers) are running on Windows machines

The important consideration for this configuration is to set up the shares, exports and mount points in such a way that when iMANFS makes a comparison check between the Unix and Windows paths as defined in the Teamcenter Engineering volume definition, it can verify that both paths are valid to access files in the Teamcenter Engineering volume. The comparison check is done by finding a matching entry for each volume definition path in the `/etc/mnttab` file. The advantage of this type of configuration is that Windows processes communicate via CIFS and UNIX processes communicate via NFS with the same filer.

In order to achieve the goal described above, the `IMAN_negotiate_nfs_access` preference has to be set to NO and `IMAN_force_nfs_access` has to be set to YES in the `.iman_env` file for Teamcenter Engineering version 7.0.3.10 or later.

3.3. Implementation Details

3.3.1. Creation of File System Volumes on NetApp Filers

File system volumes can be created and shared on filers using the command line interface or the Web-based management tool FilerView®. FilerView can be used by pointing a Web browser to the URL on the filer `http://<filer>/na_admin` where `<filer>` is the network node name of the filer.

For example a new file system volume, vol04 consisting of 14 x 36G drives, can be created in `Volumes -> Manage -> Add New Volume`.

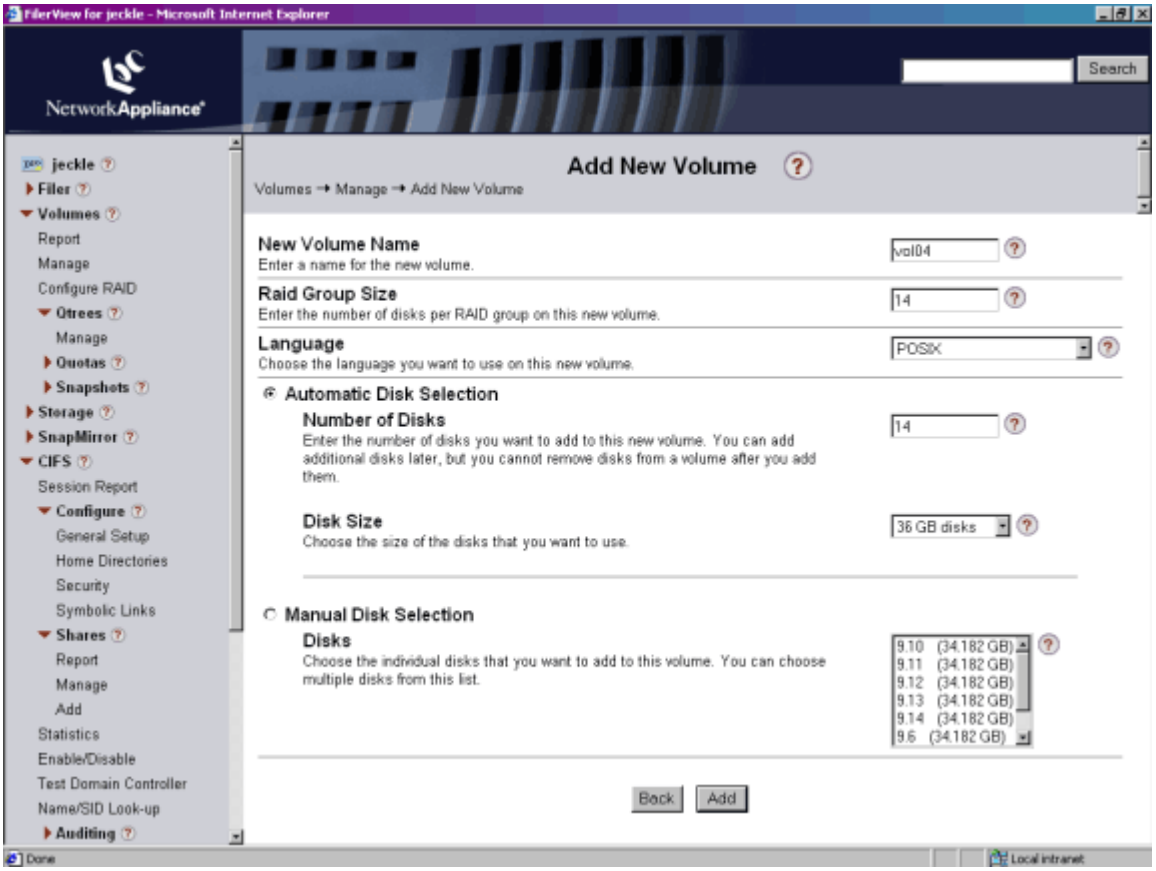

Figure 3. Add New Filer Volume

Logical directories called qtrees can be created on the newly created file system volume. Quotas can be associated with a qtree. A qtree can be grown or shrunk by changing the quota associated with it by editing the `/etc/quotas` file on the filer.

For example, on the newly created volume vol04, a qtree 881c_win_vol10 can be created.

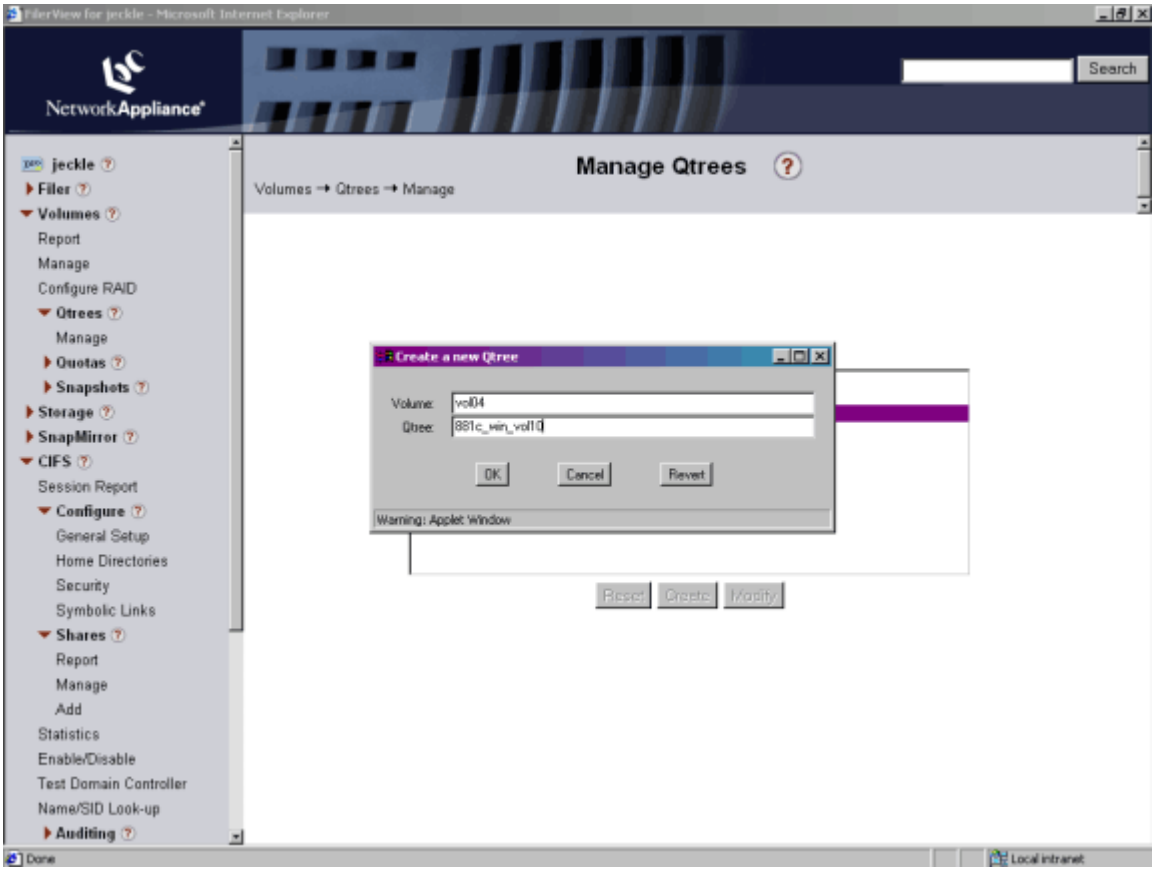

Figure 4. Create a Qtree

Qtrees can be of type unix, ntfs, or mixed. For a multiprotocol environment, qtree type "mixed" should be selected using `Volumes -> Qtree -> Manage` in FilerView.

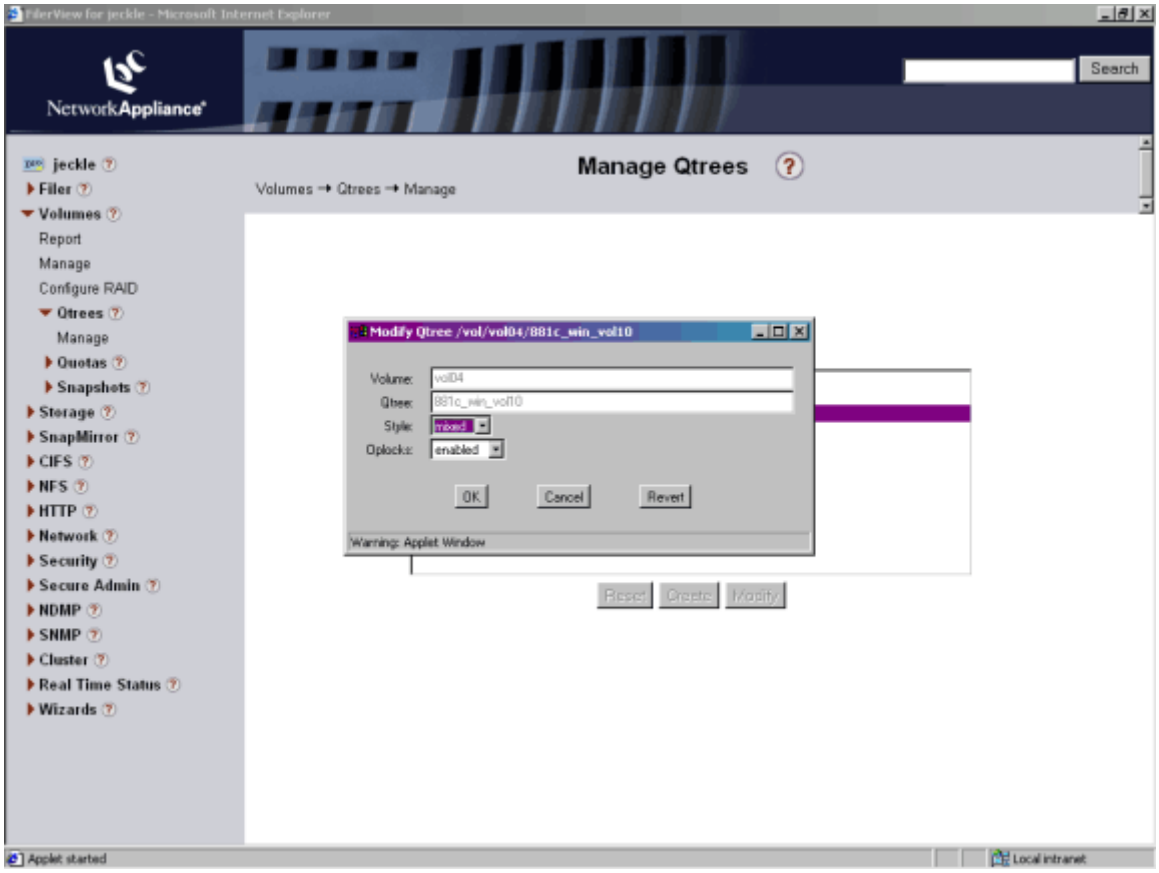

Figure 5. Change Qtree Security to Mixed

The qtree can be shared using `CIFS -> Shares -> Add`

A CIFS share for the qtree 881cin_vol10 can be created as shown.

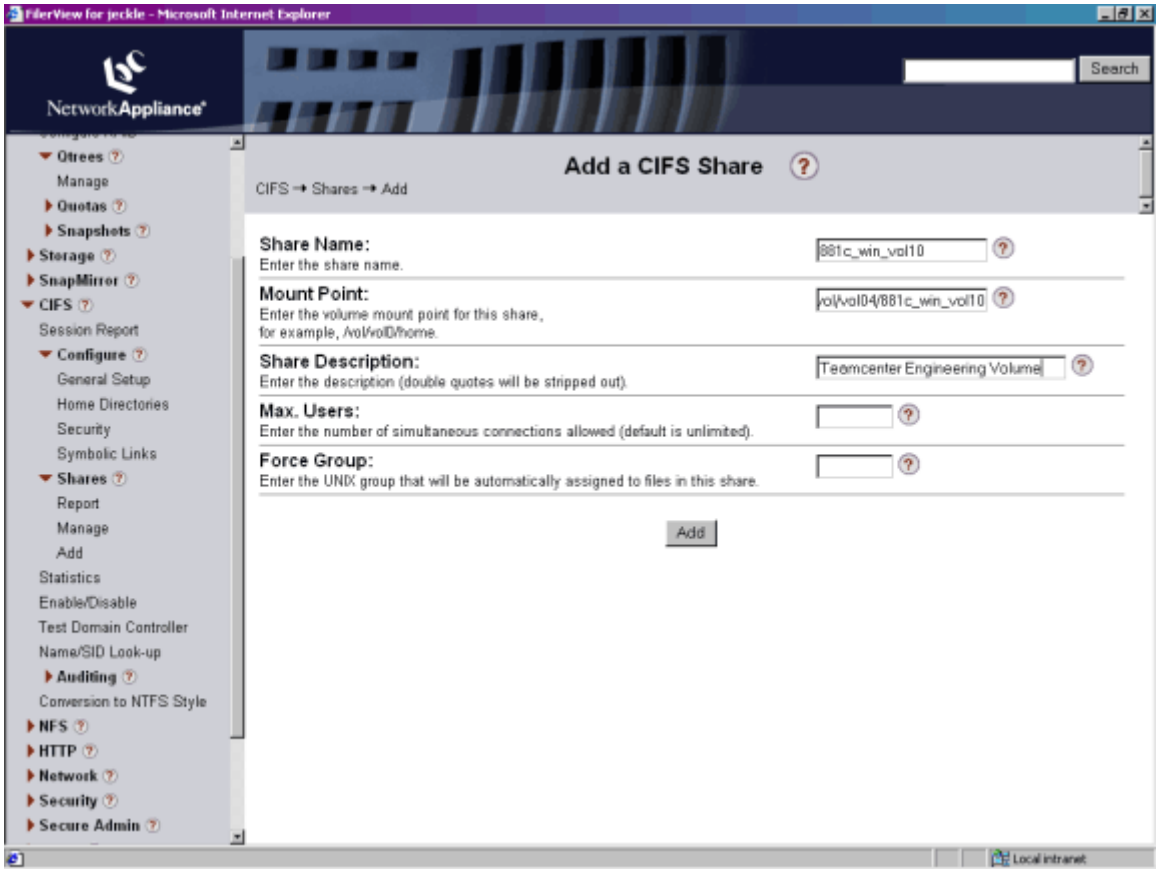

Figure 6. Add CIFS Share

The same qtree will also need to be exported using NFS by entering `NFS -> Manage Exports`.

The next example shows the creation of an NFS export for the qtree 881c_win_vol10 using FilerView.

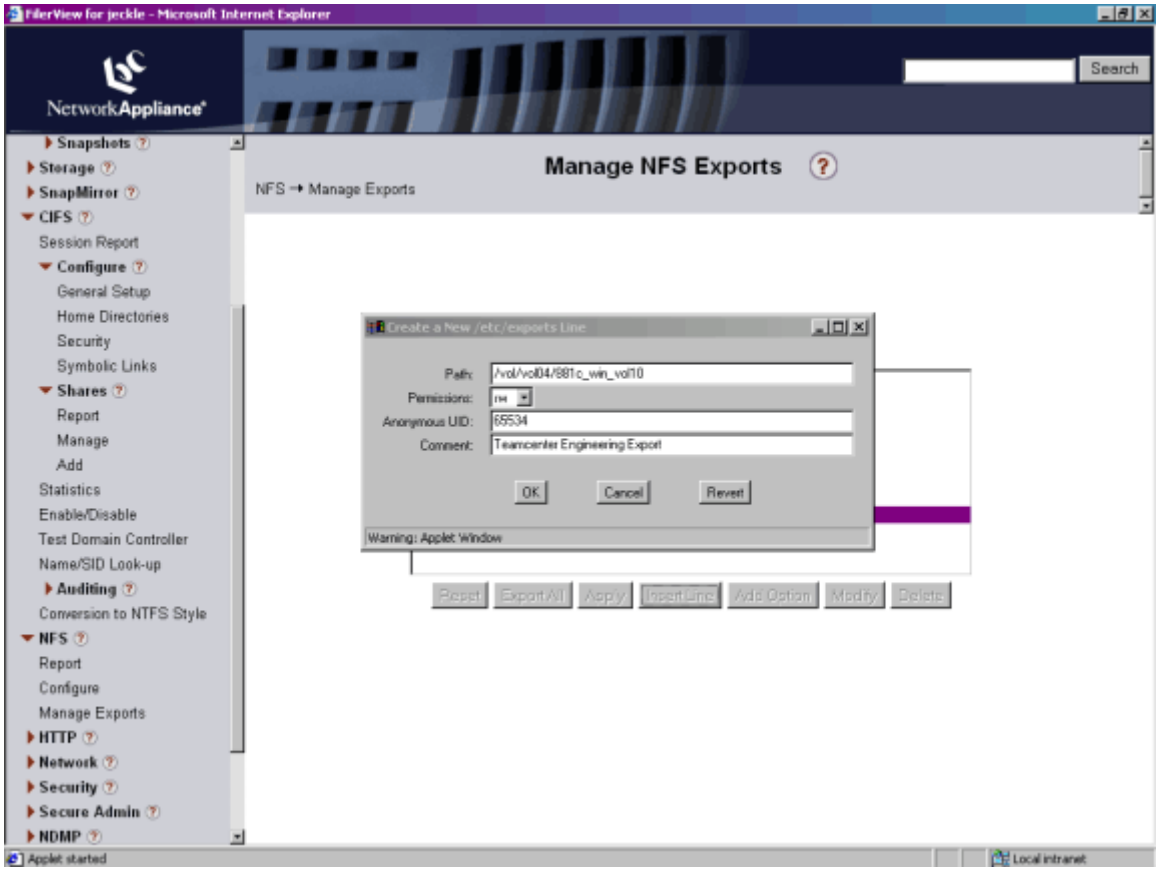

Figure 7. Create NFS Export

3.3.2. Creation of TcEng Volumes

Teamcenter Engineering volumes are defined in the Organization application. The TcEng volume information includes the path to the file system volume (one for UNIX, one for Windows), and is stored in the database. Clients use this information to access data files stored in Teamcenter Engineering using the file system path.

Figure 7 shows how TcEng supports multiprotocol access to NetApp filers. Access to the same volume is specified by supplying an NFS(UNIX) path and a CIFS(Windows UNC) path.

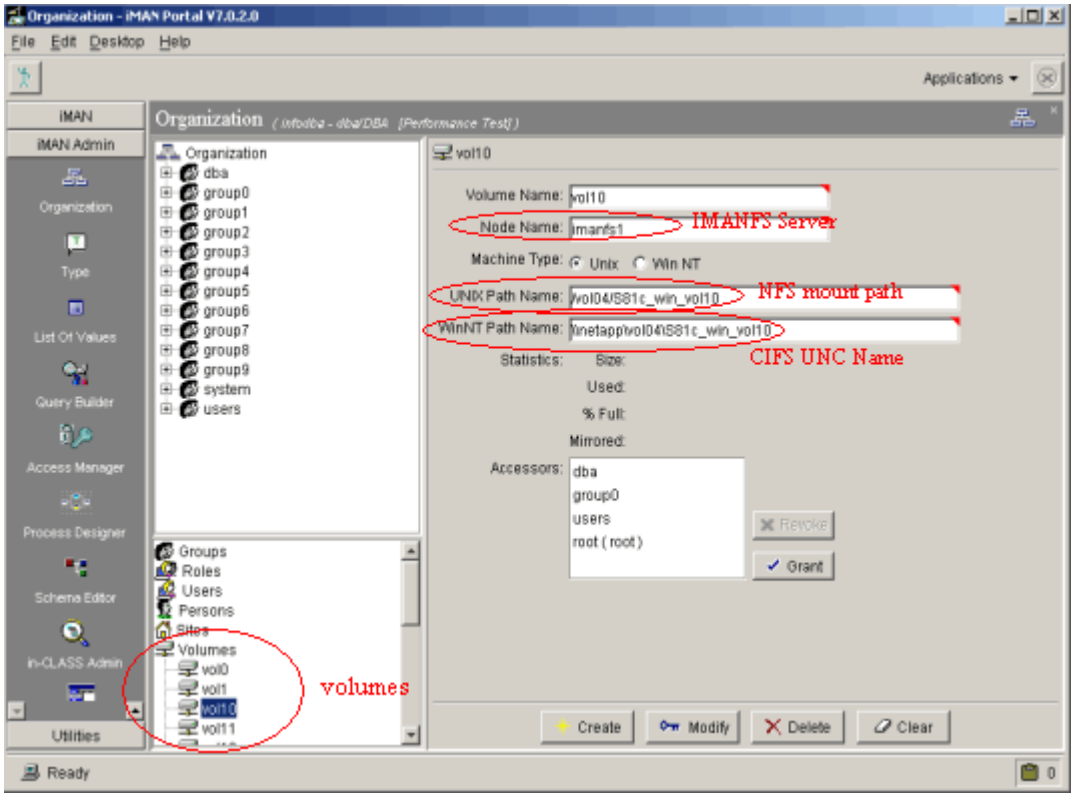

Figure 8. Teamcenter Engineering Create Volume

3.3.3. Installation of Oracle for Use with NetApp Filers

In a Teamcenter Engineering installation, putting the Oracle database on a NetApp filer has not been tested by UGS if NFS is the access protocol. Due to high volume of database update transactions, UGS has validated the configurations where the database is on a disk local to the Oracle server or configured in a high performance filer/SAN configuration.

Best practices on installation of Oracle on a filer should be followed for best results. The whitepaper *Oracle9i for UNIX: Integrating with a Network Appliance Filer for Oracle9i* contains information on new installs and migration of existing Oracle databases onto a filer.

3.3.4. Backup Design Considerations

A significant benefit of NetApp filers is the ability to back up data via Snapshots in a matter of seconds, minimizing downtime. During the backup, data may be locked down by the backup process and thereby be unavailable to the application. Traditional methods could take hours, in which case the TcEng environment is unavailable. This is not acceptable for 24x7 operations. With Snapshot, the data is quiesced, backed up, and available in a very short period of time. Then the logically consistent Snapshot can be backed up to an offline medium while the data continues be available to users. Again, up to 255 Snapshots can be kept for each file system volume at any point in time.

The guidelines for backup of Teamcenter Engineering version older than 8.0 are as follows:

- Notify online users to log out 15 or more minutes prior to initiating the backup.
- Boot out remaining users after the waiting period.

- Shut down iMANFS daemon. Note that doing this carries a risk for internally integrated applications (i.e., UG) that continue to perform writes to TcEng volumes. It is recommended that UG users save their work outside TcEng while iMANFS is not running.
- Suspend writes to Oracle. This is achieved by shutting down the instance (cold backup mode) or by running in archivelog mode (hot backup mode).

This suspends writes to the filer. Next:

- Create a Snapshot for the TcEng and Oracle volumes.

```
rsh <filer> snap create <volume> <snapshot_name>
e.g.:
rsh <filer> snap create imanfs_vol01 backup1
rsh <filer> snap create oracle_vol01 backup1
```

  This step ensures that a point-in-time image of the Teamcenter Engineering data has been created.

- Restart iMANFS daemon and resume normal Oracle operations.

Starting from Teamcenter Engineering version8.0 or greater, backup can be accomplished without having to shut down iMANFS daemon; the guidelines are:

- Fifteen or more minutes prior to initiating the backup, place the TcEng system in Caution Message Mode. This will warn the user to save their work in order to avoid any loss of data during backup.
- After the warning period, place the TcEng system in Read Only Mode. This will place the iMAN application into a read-only state and prevent files from being written to the volume during backup.
- Disable TcEng logins and notify online users to log out.
- Boot out remaining users.
- Suspend writes to Oracle. This is achieved by shutting down the instance (cold backup mode) or by running in archivelog mode (hot backup mode).

This suspends writes to the filer. Next:

- Create a Snapshot for the TcEng and Oracle volumes.

```
rsh <filer> snap create <volume> <snapshot_name>
e.g.:
rsh <filer> snap create imanfs_vol01 backup1
rsh <filer> snap create oracle_vol01 backup1
```

  This step ensures that a point-in-time image of the Teamcenter Engineering data has been created.

- After the Snapshots are complete, place the TcEng system in Normal Mode. This will allow writing files to the volume to resume.

Data can now be backed up from the `/vol/<volume>/.snapshot` directory on the filer independently of Teamcenter Engineering. The backup window was reduced to the time required to create the Snapshot, which is usually a few seconds.

Snapshots allow up to 255 point-in-time online copies of Teamcenter Engineering volumes in addition to the tape backup. SnapRestore allows a volume to be reverted back to any one of the Snapshots. To display all available Snapshots for volume vol1:

```
rsh <filer> snap list vol1
```

To restore to a specific volume Snapshot:

```
rsh <filer> vol snaprestore imanfs_vol01 -s backup1
rsh <filer> vol snaprestore oracle_vol01 -s backup1
```

This will restore volumes `imanfs` and `oracle` to a previously created Snapshot after file system malfunctions or to correct data corruption problems.

3.3.5. Performance Considerations

To ensure best performance from a Teamcenter Engineering infrastructure, the portal clients, portal and iMANFS servers, network, and filer should be configured to maximize throughput. Gigabit Ethernet should be used especially between the filer and Teamcenter Engineering server to ensure high throughput. Network aggregation technologies such as EtherChannel Trunking allow network traffic to be shared between multiple NIC interfaces, such as the interfaces of a Quad FastEthernet NIC, to configure the Teamcenter portal client systems.

1. **UNIX Recommendations**

   - Manually balance NFS traffic to separate network interfaces on the filer.
     Attach multiple interfaces on the filer to the same physical network each with its own interface name. For example, if two Ethernet interfaces (named toaster-0 and toaster-1) on the filer named `toaster` are attached to the same network where four NFS clients reside, point half to one interface and half to the other. Specify in `/etc/vfstab` on client1 and client2 that these clients mount from toaster-0:/home. Specify in `/etc/vfstab` on client3 and client4 that these clients mount from toaster-1:/home. This method can balance the traffic among interfaces if all clients generate about the same amount of traffic. The filer always responds to an NFS request by sending its reply to the interface on which the request was received.

   - Add disks to a disk-bound volume
     If you have a single-volume filer, use the `sysstat -u` command on the filer to determine the fraction of time that the busiest disk is active. If the fraction is greater than 80%, add disks to the volume using the `vol add` command.

   - Maintain adequate free blocks and free inodes
     If the percentage of free blocks or free inodes falls to less than 10% on any volume, the performance of writes and creates can suffer. Check free blocks and inodes using the `df` command and `df -i` command, respectively. If the percentage of used blocks is greater than 90%, increase blocks by adding disks or deleting Snapshots. If the percentage of free inodes is less than 10%, increase inodes by deleting files or using the `maxfiles` command.

   - Determine when to use UDP or TCP transport
     The following are guidelines to determine when you should use the UDP

transport or the TCP transport to improve filer performance:
Use the TCP transport over a WAN network.
Use the UDP transport over a LAN network.
Use the TCP transport if you are using the UDP transport and you experience packet loss, especially during periods of heavy write traffic.
You can specify the transport using the options `nfs.tcp.enable` command, and also explicitly use TCP for transport by specifying the "proto=tcp" option in `/etc/vfstab` on the NFS clients.

The following performance guidelines apply to Sun™ Solaris™ deployments.

- Increase the size of STREAMS synchronized queues on the Sun client. Add the following to its `/etc/system` file:

```
* Increase size of STREAMS synchronized queues to increase
* network performance. Release Notes for qfe card recommend
* setting this to 25 per 64MB of RAM in the system. It also
* prevents receive overrun on the GbE interface.
set sq_max_size=XX
```

Note that a reboot is required after applying these changes.

- Increase size of UDP and/or TCP high-water marks.

```
ndd -set /dev/udp udp_recv_hiwat 65535
ndd -set /dev/udp udp_xmit_hiwat 65535
ndd -set /dev/tcp tcp_recv_hiwat 65535
ndd -set /dev/tcp tcp_xmit_hiwat 65535
```

The commands can be saved permanently by creating a control script in the client's `/etc/rc2.d` directory. Choose a unique file name, like `S99netperf`.

```
case "$1" in
'start')
        echo "Setting local kernel parameters...\c"
        ndd -set /dev/udp udp_recv_hiwat 65535
        ndd -set /dev/udp udp_xmit_hiwat 65535
        ndd -set /dev/tcp tcp_recv_hiwat 65535
        ndd -set /dev/tcp tcp_xmit_hiwat 65535
        echo " "
        ;;
'stop')
        echo "$0: No parameters changed."
        ;;
*)
        echo "Usage: $0 (start|stop)"
        ;;
esac
exit 0
```

- The following driver patch should be installed for the Sun Gigabit/2.0 NIC. Among other things, the patch reduces erratic behavior in some applications like the dd tool and improves performance, especially on the Sbus NIC.

```
106764-XX SunOS 5.6 / Solaris 2.6
106765-XX SunOS 5.7 / Solaris 2.7
108813-XX SunOS 5.8 / Solaris 8
```

Use the command `showrev -p` to display the installed patches.

2. **Windows Recommendations**

   - Enabling Level I and Level II oplocks for CIFS filers
     Oplocks (opportunistic locks) allow CIFS clients to read ahead, write behind, and lock cache data locally. This reduces traffic to the filer and improves performance. See http://support.microsoft.com/default.aspx?scid=KB;EN-US;Q129202& for a detailed explanation of Level I and Level II oplocks and how they may be relevant to your environment. Both Level I and Level II oplocks are on by default on the filer, which is appropriate in most iMAN environments. To ensure that Level I oplocks are on, the command

     ```
     options cifs.oplocks.enable on
     ```

     can be run on the filer. Similarly, to ensure that Level II oplocks are on, the command

     ```
     options cifs.lvl2_opclocks_cap on
     ```

     can be used.

   - CIFS Negotiated Buffer Size
     cifs.neg_buf_size is the option that controls the negotiated I/O buffer size for clients. The recommended setting of this filer option is at least 33028.

   - Increase the TCP window size for CIFS
     Increasing the TCP window size to its maximum setting on both the filer and the CIFS client can improve performance for large transfers. The TCP window size controls the number of TCP messages that can be transmitted between the filer and CIFS client before an acknowledgement is received back from the destination. The filer supports a maximum window size of 64,240.

     Use the `cifs.tcp_window_size 64240` command to maximize the TCP window size on a filer running CIFS. Use the `nfs.tcp.recvwindowsize 64240` command to maximize the TCP window size on a filer running NFS. Change the window size in the Windows registry on a Windows NT client by adding the DWORD value `\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpWindowSize` and set it to 64,240 (0xFAF0 in hexadecimal). A reboot is required before changes take effect.

3.3.6. Operational Testing

To verify the operation and performance of NetApp filers, Snapshot, and SnapRestore, several configurations of iMANFS, Oracle, and imanserver, both on Windows and UNIX, were set up and tested. Performance and scalability of Teamcenter Engineering was not measured because locating Oracle metadata on NFS file system has not been tested by UGS. UGS has validated that all Oracle metadata, tablespaces, and control files be located on local file system or on SAN attached file system.

The time required for creating Snapshots on an operational volume was collected throughout the day. The operational volume consisted of 14x36GB disks. Size of the volume was 300GB and was 78% utilized. The average amount of time for the Snapshot operation for the 300GB volume was 1.24s. Results were excellent.

	Configuration 1	Configuration 2
Oracle binaries	Filer	HPUX Server
Oracle metadata	Filer	Filer
TcEng volume	Filer	Filer
Imanserver binaries	Filer	Local to client*
iMANFS binaries	Local to HP	Local to W2K server
Imanserver process	Local to client (LSE)*	Local to client*
iMANFS process	HPUX Server	W2K server
Oracle process	HPUX Server	HPUX Server
Portal client	Various*	Various*

* Three client configurations were tested:

- Solaris 2.8 workstation
- HPUX 11.1i workstation
- Windows 2000 SVP2 workstation

NetApp volumes and qtrees must be mixed mode.
TcEng must be in Read Only Mode for Snapshot or SnapRestore.

4. Appendix A

NetApp's WAFL file system offers read-only Snapshots that allow users to access earlier versions of their files. Since the Snapshot is offline, backups can be done from Snapshots, even while the current file system is being actively updated by users.

Files in a Snapshot of the file system have the same access privileges as in the active file system. A user with permission to read a file in the active file system will still be able to read it in a Snapshot. Users without such permission will still be unable to read the file. Write access privileges simply do not apply—files within a Snapshot cannot be changed because Snapshots are read-only. Updates to files must be performed on the file's image in the current, active file system.

Backups can be performed via Snapshots, allowing users to continue to write to the live, actual file system without endangering the restorability of the backup in progress. In fact, the NetApp dump command will automatically create a temporary Snapshot for its own use, unless directed otherwise by the system administrator.

Snapshots are whole parallel file systems, from the top of the directory tree on down. There is an entry into this parallel directory tree accessible from every directory—a subdirectory named `.snapshot` for UNIX/NFS (and PC software that can handle long file names) within every directory in the active file system. These .snapshot subdirectories are invisible. Within the Snapshot subdirectories are lower subdirectories named for the interval in which that particular Snapshot was taken (e.g., `hourly.0`).

Also, at the top of the directory tree, the Snapshot may be entered through a visible directory named .snapshot (for UNIX/NFS), `~snapshot` for long-file-name-compatible PC/CIFS clients, and `~SNAPSHT` for PC/CIFS clients that require 8.3 file names.

In order to avoid confusion, except for the top of the directory tree, Snapshot entry point subdirectories are not displayed—are invisible—in ordinary listings of directory contents, nor can they be seen by other methods of inspecting a directory's contents. This is accomplished slightly differently in NFS than in CIFS. Aside from simplifying the user's interface with the file system, concealing Snapshots also prevent problems such as the UNIX recursive remove (rm -rf) from happening because of the read-only nature of the contents of a Snapshot.

NetApp®

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com