Technical Report

Siemens PLM Teamcenter: Deployment and Implementation Guide on Data ONTAP Operating in Cluster-Mode

NetApp and Siemens PLM
August 2012 | TR-4098

Abstract

This technical report describes in detail how to deploy and implement Siemens® PLM® Teamcenter® on NetApp® Data ONTAP® software operating in Cluster-Mode. It provides information on best practices, the performance of Teamcenter across NetApp storage protocols, deployment options, and the benefits and storage efficiencies that Cluster-Mode offers in a Teamcenter environment.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Siemens PLM software provides a unified solution to enterprises that enables all members in the product lifecycle to work in concert to bring products to market. These solutions are used by companies to help them manage product information throughout the process of designing, developing, and building products, ranging from airplanes to semiconductors. Siemens PLM Teamcenter product is an industry-driven Web-centric product lifecycle management system. It allows distributed engineering and manufacturing teams with global sharing and workgroup management capabilities to capture, manage, and leverage engineering data created by multiple CAD, CAM, and CAE systems. Complex products require thousands of engineering designs and drawings, and Siemens PLM Teamcenter solution helps manage and communicate all of this information. Companies deploy PLM solutions such as Siemens PLM Teamcenter with the goal of maintaining accurate product information, enabling better designs, and supporting collaboration across design teams and the supply chain. PLM products integrate with CAD, CAM, and CAE systems as well as Product Data Management (PDM) solutions.

Data ONTAP operating in Cluster-Mode is NetApp's next-generation storage solution that extends the core strengths of the NetApp Unified Storage Architecture, integrated data protection, and storage efficiency with the addition of massive scalability, increased performance, integrated tiered storage, improved operational efficiency, and a single management system. This third-generation clustered storage solution provides a foundation for continuous business operation and improved flexibility.

Siemens PLM software and NetApp solutions integrate seamlessly to provide a PLM collaborative engineering environment. NetApp Cluster-Mode storage with the Siemens PLM solution provides a scalable storage infrastructure and practices that allow engineering and product development teams to minimize or eliminate data loss, accelerate recovery, enable collaboration among distributed design groups, and simplify data management for faster development cycles, improved engineering productivity, and reduced time to market.

Siemens PLM has been a strong NetApp partner for many years. The two companies work together to validate and support solutions, performance testing, and analysis. They collaborate in offering joint customers a development platform that they can trust.

1.1 Scope

This document is intended for use by individuals who are responsible for architecting, designing, managing, and supporting Siemens PLM Teamcenter on Cluster-Mode storage. It gives the reader an understanding of the details for deploying and implementing Teamcenter on Cluster-Mode storage. For information about backup and recovery, see "Siemens PLM Teamcenter: Backup and Recovery on NetApp Data ONTAP Operating in Cluster-Mode."

2 Architecture Overview of Siemens PLM Teamcenter

Siemens PLM Teamcenter provides thin Web-based and rich Java®-based applications that use J2EE and .NET technology. It is composed of clients, Web services, enterprise application, file management system, and database to manage the designs of product developers. Teamcenter is a scalable solution in which multiple Web application servers and multiple business logic server pools can be configured to support numerous users. These application services can be installed on a single system or distributed on separate systems to balance the workload across different systems. Depending on a customer's requirements, Teamcenter-rich clients can be deployed as a four-tier or two-tier configuration.

This section is a high-level description and explanation of the uses of each of these components. For detailed information on Teamcenter, refer to the "System Administration Guide for Teamcenter" in the current release of the Teamcenter documentation section of the Siemens PLM Web site.

2.1 Four-Tier Deployment

In a four-tier configuration, Siemens PLM Teamcenter is comprised of the following components.

- **Client Tier** hosts the client applications and provides the user interface via a Java application or a browser and hosts secure file caches. There are two types of clients:
  - Rich Client, which is a Java application
  - Thin Client, which is browser based
- **Web Tier** is a composition layer holding Web application session state, serving static content, and routing client requests to the business logic server (Enterprise Tier). It has support for J2EE and .NET technology and currently supports IIS, JBoss, WebLogic, and Websphere Web application servers.
- **Enterprise Tier** hosts the business logic server, generates Teamcenter server processes, and serves dynamic content to clients.
- **Resource Tier** manages and stores Teamcenter persistent data, bulk, and metadata in a database or file vaults. It is composed of the following components.
  - **File Management System,** which manages client access to design files on the storage. The File Management System (FMS) allows files to be requested via logical identities, referred to as FMS tickets, instead of physical location. Thus, instead of working directly with file paths, the File Management System uses these FMS tickets or logical identities for the files. There are two processes that run as part of this file management system: FMS Server Cache (FSC) runs on the server hosting the storage and caches the files, and FMS Client Cache (FCC) runs on the rich client host and caches the file locally. Multiple FSCs can be configured in a hierarchical fashion to support local storage for improved performance at remote locations (that is, over a WAN). The File Management System has the ability to do the following.

    a) Retrieve specific parts of a file or the whole file

    b) Compress files during transmission

    c) Optimize the TCP/IP connection for parallel processing of file chunks

    d) Ability to upload large files asynchronously to user saves actions

The File Vault stores the actual designs that developers are working on.

- **Database,** which stores the metadata associated with the designs being managed by the Teamcenter File Management System. Databases that are supported include Oracle®, Microsoft® SQL Server®, and IBM DB2.

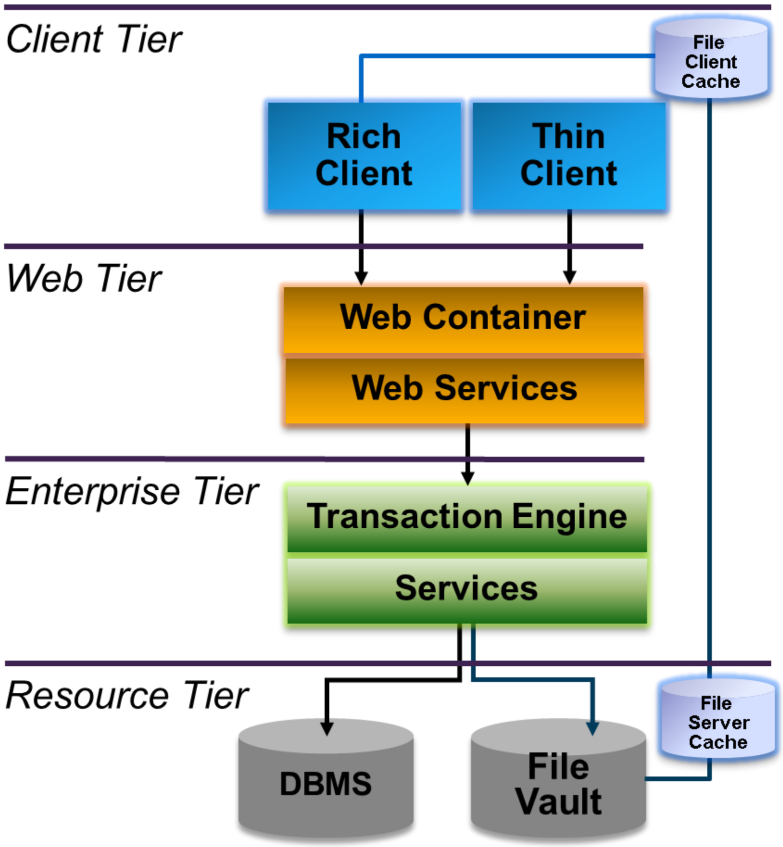

Figure 1) Simplified diagram of a standard four-tier model.

Data Flow For A Four-Tier Teamcenter Configuration

The data flow in a four-tier configuration involves communication between the tiers. When a design file is uploaded to the client for editing or viewing, the data flow is as follows.

1. When a file is requested, the client requests a session from the Web Tier.
2. The Web Tier forwards requests to the server manager.
3. The server manager starts a Teamcenter (TC) server session for each request.
4. The Teamcenter server session that was initiated sends a query to the database and the database returns a unique file ID to the client.
5. The TC server session requests the File Management System to upload the file from the File Vault to the File Server Cache.
6. The file is uploaded from the File Vault to the File Server Cache.
7. Finally the file is uploaded to the File Client Cache in a Rich Client scenario or directly to the Thin Client application via the Web container.

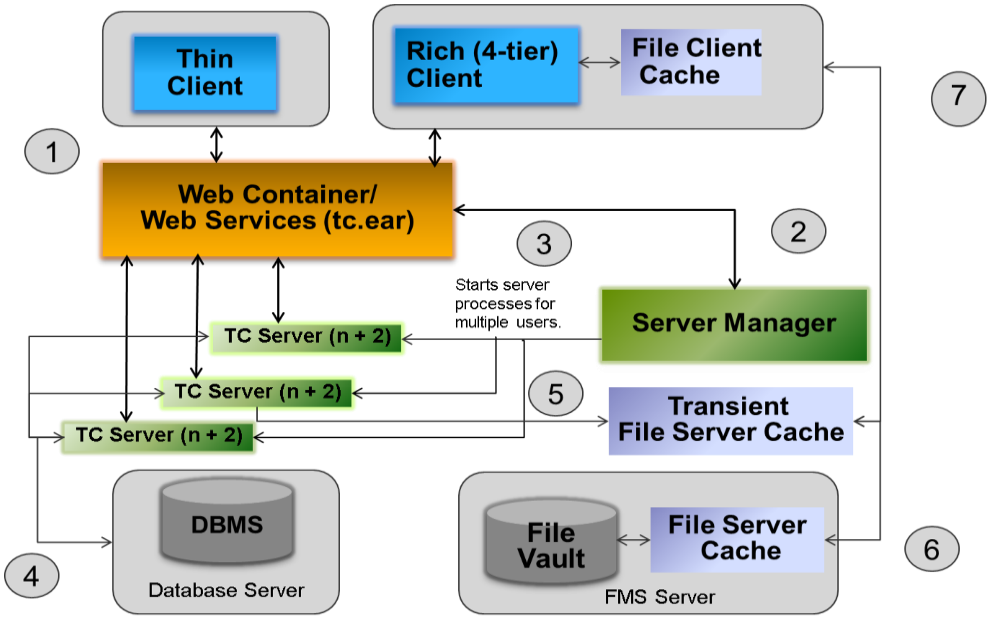

Figure 2- Data Flow In a 4-Tier Configuration

2.2 Two-Tier Deployment

In a two-tier deployment of Teamcenter, there is no Web Tier, and the enterprise and rich client are combined in one tier. The Teamcenter server processes runs on the client host. The client host communicates and requests data directly to the resource tier, which includes the database and file vaults.

Figure 3) Two-tier simplified diagram.

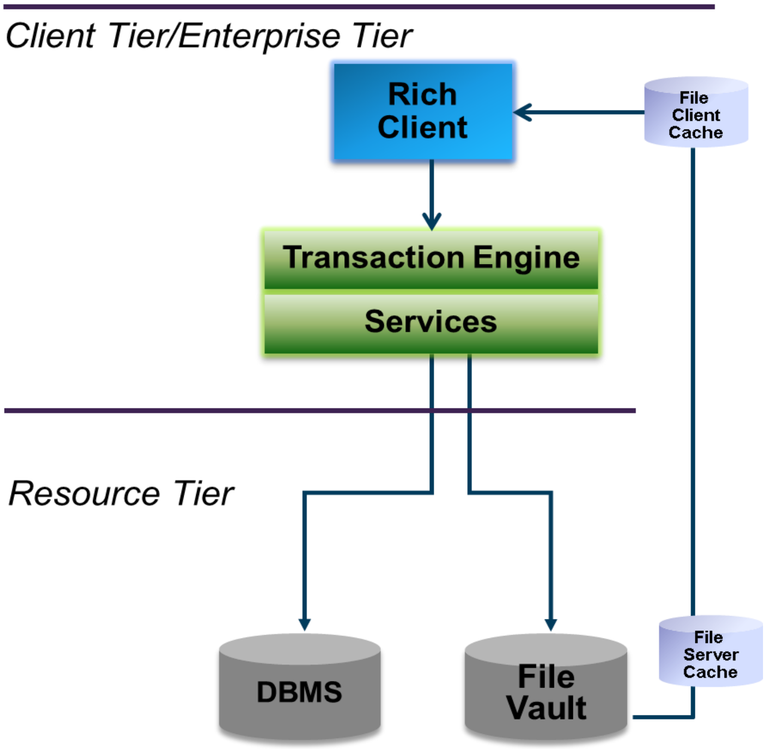

Data Flow In A Two-Tier Configuration

The Teamcenter data flow in a two-tier configuration involves the following.

1. The rich client requests a file from the Teamcenter (TC) server.
2. The TC server requests the object from the database and returns a unique file ID to the TC server.
3. The TC server then communicates with the File Management System to place files from the File Vault into the File Server Cache.
4. The file is then streamed to the File Client Cache to be served to the Rich Client by the File Management System.

Figure 4) Data flow in a two-tier configuration.

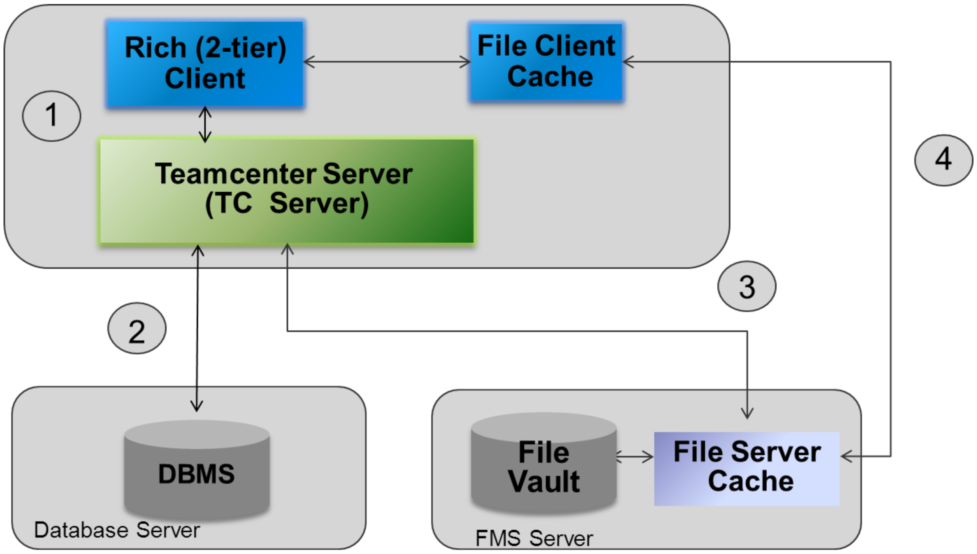

3 Value of Netapp Cluster-Mode

In product development environments, there is a common set of key challenges. Customers deploying Siemens PLM Teamcenter will often operate in a distributed development environment and thus there is a need to have data available 24/7. It is also crucial to control costs, minimize data loss, accelerate recovery, and meet PLM performance requirements. NetApp Data ONTAP 8.1 operating in Cluster-Mode enables product development organizations to couple multiple storage systems in a single integrated cluster under a global namespace. The cluster is managed as a single entity, and data can be easily moved between different storage systems transparent to both Teamcenter and engineers. The key values of NetApp clustered solutions for Siemens PLM Teamcenter environments include:

- Accelerates product development with an always-on infrastructure
- Reduces IT costs by managing the entire Teamcenter storage infrastructure as a single entity
- Decreases risks

3.1 Accelerate Product Development

Project requirements can rapidly change and new projects can increase the data and storage requirements. Being able to quickly provision or move data nondisruptively to accommodate these changes is important. NetApp clustered features such as global namespace, single management system, and nondisruptive volume move can assist in meeting these requirements.

Dynamically rebalance development workloads

With traditional storage systems, if a Teamcenter repository requires additional performance, other projects must be moved off that storage system or that repository must be moved to a system that has higher performance. In either case, the results are an interruption in data availability and new pathnames for the affected projects. With NetApp Cluster-Mode, multiple storage systems are integrated into a single cluster with a global namespace, and projects can be moved between systems transparently to both engineers and the Teamcenter application. With this powerful capability, workloads that affect performance of a hot project can be transparently moved to other systems, or the hot project itself can be moved to another system. In either case, the load is dynamically rebalanced without disrupting either engineers or the Teamcenter application, optimizing performance and helping to accelerate product development.

Dynamically resolve network congestion

Data ONTAP systems operating in Cluster-Mode also offer the ability to dynamically resolve network congestion issues. With traditional systems, if one of the physical ports to a storage system is congested, then the only recourse is to physically move users to an alternate port. This approach unfortunately disrupts engineers who are accessing the data. With Cluster-Mode systems, administrators can simply move the logical interface (LIF) from the congested physical port to another LIF on a less congested port. This action is transparent to engineers and to Teamcenter, again helping to maximize productivity and accelerate the product development process.

Always-on Infrastructure

Cluster-Mode enables an “always-on” storage infrastructure, supporting 24/365 product development. For example, when a storage system for Teamcenter is upgraded, serviced, or retired, the projects on that system can be temporarily moved to other nodes in the system and then later rebalanced after completion of the service or upgrade event, without disrupting the engineer.

Instant Clones for accelerated testing

NetApp FlexClone® technology enables instant, space-efficient clones of production or test data to be created easily, minimizing storage requirements, because physical disk space for these clones is required only when data blocks are changed. With NetApp FlexClone, each engineer can have his or her own virtual copy of production or test data for use in test and development activities.

Instant, consistent backups

Traditional approaches to backups of Teamcenter data can take time; in the case of tape-based backups, the repository might be unavailable for hours. Such lengthy disruptions can significantly affect the product development process. NetApp Snapshot™ technology enables backups to be done in minutes, eliminating long backup windows. By consolidating storage to NetApp, backups are consistent across both the database and the Teamcenter file vault.

3.2 SCALE INFRASTRUCTURE, NOT OPEX

Fast-growing product development organizations often find that as the amount of PLM data increases over time, operating expenses similarly increase. This issue is particularly acute with direct-attached storage (DAS). NetApp Cluster-Mode systems minimize administrative activities and costs associated with growing infrastructures.

A single point of management

A cluster of NetApp systems is managed as an integrated entity, not as a set of independent storage systems. The global management interface enables administrators to move projects between systems and to manage all Teamcenter storage in the cluster from any node. With this approach, the incremental management burden is relatively flat as additional systems are added to the cluster. Transparent data movement simplifies management. With traditional systems, moving data from one storage system to another is management intensive, because users need to be informed of the upcoming interruption and administrators need to move the data and update the Teamcenter configuration with the new path names. There is risk of error throughout this process. With NetApp Cluster-Mode, there's no interruption to manage, no user communications are necessary, and no changes to product development processes are required.

Optimize costs with tiered storage

A NetApp cluster can include nodes of varying levels of performance, as well as different types of storage (SSD, FC, SATA, and SAS) with varying cost, capacity, and performance characteristics. With Cluster-Mode, Teamcenter projects that are less frequently accessed can be transparently moved to lower tiers of storage, enabling costs to be optimized without affecting developer access. Cost savings are further increased with deduplication of CAD files and workspaces across all tiers of storage.

Unified storage with native protocols

Each system in a NetApp cluster supports native file and SAN protocols, including NFS, CIFS, FC, iSCSI, FCoE, and object protocols. This approach eliminates the need to adopt different types of storage systems for different protocols or to deploy protocol emulators on servers. NetApp is the only vendor in the industry to offer unified storage at scale. With unified storage across a broad range of controllers and storage types, a single integrated NetApp cluster can effectively support not only the storage infrastructure for Teamcenter, but also design verification software and other downstream workflows. This approach significantly simplifies management and reduces operating costs.

3.3 Reduce Risks

Your product designs are important intellectual property, and NetApp Cluster-Mode systems offer powerful data protection capabilities and new levels of flexibility to adapt to change.

Easily scale

Whether you have a small development shop that might experience rapid growth or a huge shop with continuing growth, a NetApp cluster gives you the ability to easily, seamlessly grow as your development efforts expand.

Frequent backups

With traditional storage infrastructures for Teamcenter, the disruptions associated with backups are so significant that backups might be done infrequently. Unfortunately, such an approach puts more of your developers' work at risk and makes it very difficult to support aggressive recovery point objectives (RPOs). With NetApp Snapshot technology enabling quick, low-overhead backups, it is now convenient to back up frequently, protecting more of your engineers' work and enabling much more aggressive RPOs.

Simplified disaster protection

With traditional storage infrastructures, the complexity of the environment and the available data mirroring solutions make disaster recovery (DR) solutions difficult or impractical to implement. NetApp SnapMirror® technology makes DR practical through an easy-to-implement, robust mirroring solution. Implementing DR significantly reduces risks for your organization and protects your data in case of a catastrophic event.

4 Performance of Teamcenter on Cluster-Mode

A Teamcenter performance evaluation was conducted across NetApp Cluster-Mode storage protocols: NFS, CIFS, FC, iSCSI, and Split configuration (for example, DB on SAN and File Vaults on NAS) on both Windows® and UNIX® environments. This section describes the details of the Teamcenter performance and scalability tests, results, and analysis. All of the performance data was measured in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. You should verify the applicable data for your specific environment.

Two types of tests were conducted to understand Teamcenter performance across NetApp Cluster-Mode storage protocols. They are:

- Teamcenter FMS component-level throughput benchmark
- Teamcenter system-level scalability benchmark

4.1 Test Environment

The test environment for this performance study consisted of the following hardware and software components. Siemens's PLM Teamcenter v8.3.2 configuration was used for both benchmarks on xServer series machines. All benchmark traffic was isolated on a private, isolated network and all command, control, and monitoring traffic on public networks. The FC benchmarks were done utilizing a 4Gbit SAN and the NFS and iSCSI benchmarks were done via private 1Gbit Ethernet.

Clients: HP DL140 G3 - 1 x Intel® Xeon® 3.00GHz - 2GB RAM - Windows Server® 2008 Client Type: LoadRunner	**FSC/FlexServer**: IBM xServer X3250 M3 - 4 x Intel Core i3 3.07 GHz - 8GB RAM - 16GB swap - SLES 11 SP1/Windows 2008 R2
Web Application Server: IBM xServer X3250 M3 - 4 x Intel Core i3 3.07 GHz - 8GB RAM - 16GB swap - SLES 11 SP1/Windows 2008 R2 - Web Server Version: JBoss 4.2.2.GA	**Database Server**: IBM xServer X3250 M3 - 8 x Intel Xeon X3440 2.53 GHz - 16GB RAM - SLES11 SP1/Windows 2008 R2 - Database Version: Oracle® 11.2.0.1
Business Logic Server: IBM xServer X3250 M3 - 16 x Intel Xeon E5630 2.53 GHz - 64GB RAM - 134GB swap - SLES 11 SP1/Windows 2008 R2	**Storage**: NetApp FAS3240c - 8 x 2.3GHz Intel - 16GB RAM - Data ONTAP v8.1 RC3 - 1TB Flash Cache - 10GBe Nexus 5010 Cluster-Interconnect > 100K Ops/Sec[1] / 1.66 msec ORT[2]) - 16 disk aggregates were used

4.2 Teamcenter System Scalability Benchmark

The Teamcenter system-level scalability benchmark is a thin client system-level benchmark. These tests utilized standard Automated Performance Analysis (APA) scripts that generated 1,000 users executing 50% query/view type of transactions and 50% create/update workflow type of transactions with aggressive login/logout ramp up (>=1 hour steady state). LoadRunner was used to simulate up to 1,000 Teamcenter thin client users. There were 3 user types:

- **Data Analysis User Types,** which did View, View BOM, Create and Delete Folder, Display Designer, Copy, Paste and Cut Item, and Expand PSE
- **Data Review User Types,** which did display BOM report, Open WorkList Scenario, View Image, View Where Ref, View Where Used, Initiate and Complete Review, Save and View Item Properties, Address List
- **Documentation User Types,** which did Create and Delete Form, Create, Edit and Save Dataset, Revise and Delete ItemRev

For this test, the average response time to perform each type of transaction was collected. This benchmark was executed on Teamcenter 8.3 in a 4-tier environment.

Figure 5) Teamcenter system scalability benchmark flow.

1) User selects item(s) with associated file(s)
2) Client requests file information
3) Client receives FMS ticket(s) from database
4) Client requests file(s) from FMS with ticket
5) FMS streams file(s) from NetApp to client

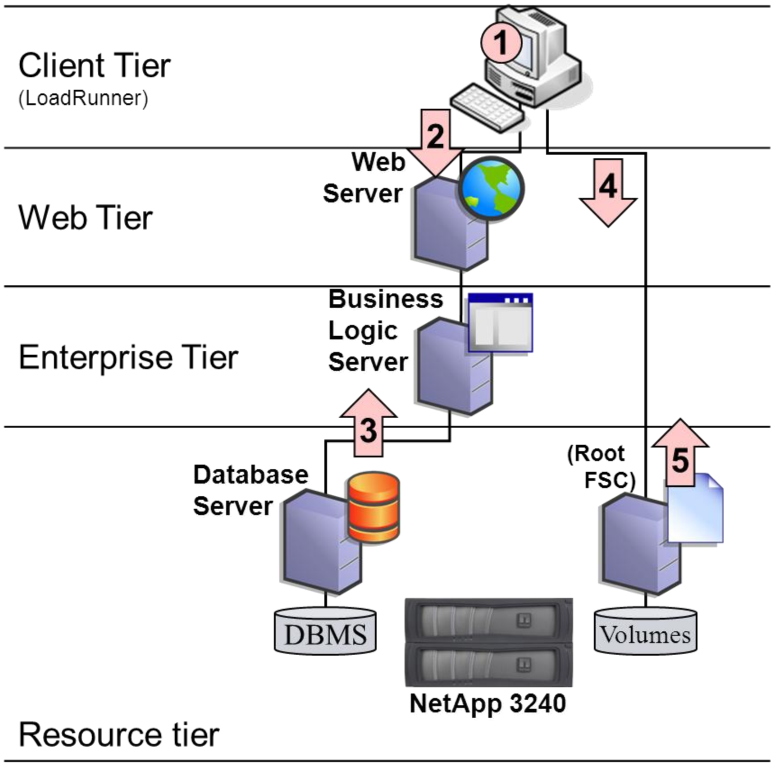

4.3 FMS Component-Level Throughput Benchmark

The Teamcenter File Management System (FMS) component-level throughput benchmark uses an FMSload tool that simulates thousands of Teamcenter FMS requests and generates FMS traffic. In this benchmark, multiple processes are executed and they simulate hundreds or thousands of users doing file uploads, downloads, and deletes accessing the volumes that were assigned to their respective groups.

Figure 6) FMS component-level benchmark flow.

1) Client sends file tag
2) Client receives FMS ticket from database
3) Client requests file from FMS with ticket
4) RMS streams file from NetApp to client

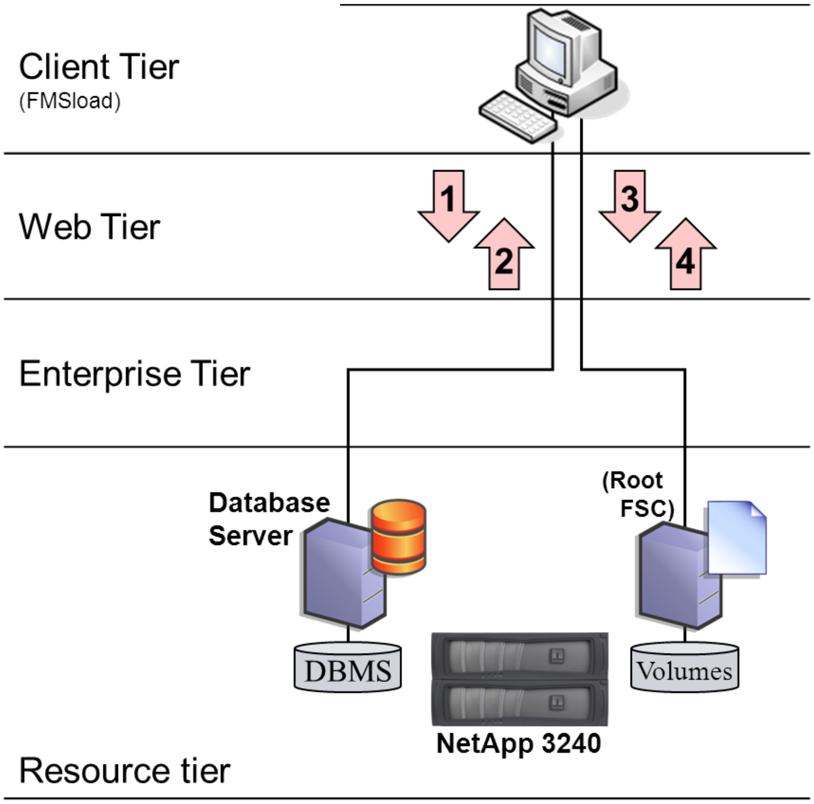

4.4 Test Datasets

Both benchmarks conducted used the same datasets, which involved a 60GB database and 250GB of read and write volumes. The database includes 50,000 users supporting up to 5,000 active users and contains:

- ~100 unique items for each user
  - .prt, .jt, .gif, .tso (true shape volume sweep)….
  - 10 read volumes/10 write volumes; 250GB
- Unique copy of an 87-component benchmark assembly
- 800K items, 1.0M datasets, 2.7M files
- 15.9M pom objects, 6.1M workspace objects
- 2.8M object relations, 32M pom backpointer rows
- 120K BOMs with geometry

- 3,400-line CTS 310, 8,200-line aerospace assembly
- ~100 to 16,000 lines

- 500K occurrences

The File Vaults contained over 2.7 million files.

The figure below describes the file size frequency distribution used for the FMS component-level benchmark. The average file size was 500KB.

Figure 7) File size frequency distribution for FMS component-level benchmark.

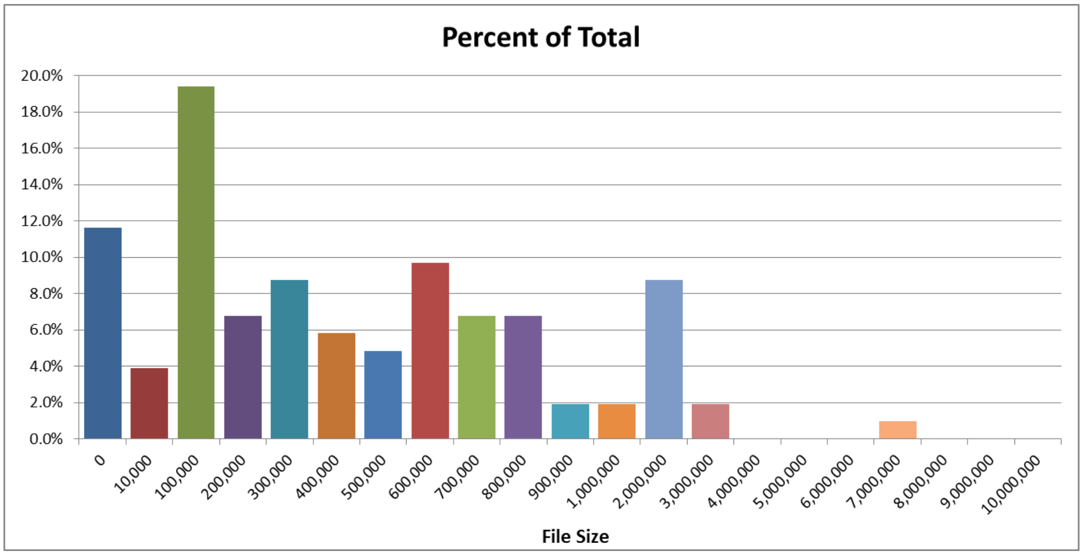

4.5 Results

This section provides the graphical results and analysis from the benchmarks conducted across NetApp Data ONTAP operating in Cluster-Mode storage protocols in both UNIX and Windows environments. For UNIX, in which Teamcenter is running on a UNIX server, the storage protocols tested were NFS v3, iSCSI, FC, and split (database on SAN and pools on NAS) configurations. For the Windows environment in which the Teamcenter server is running on a Windows server, the storage protocols tested were FC, iSCSI, and split configuration (database on SAN and file vaults on CIFS).

UNIX

Figure 8) Teamcenter system scalability (UNIX) benchmark results @ 1,000 users (seconds).

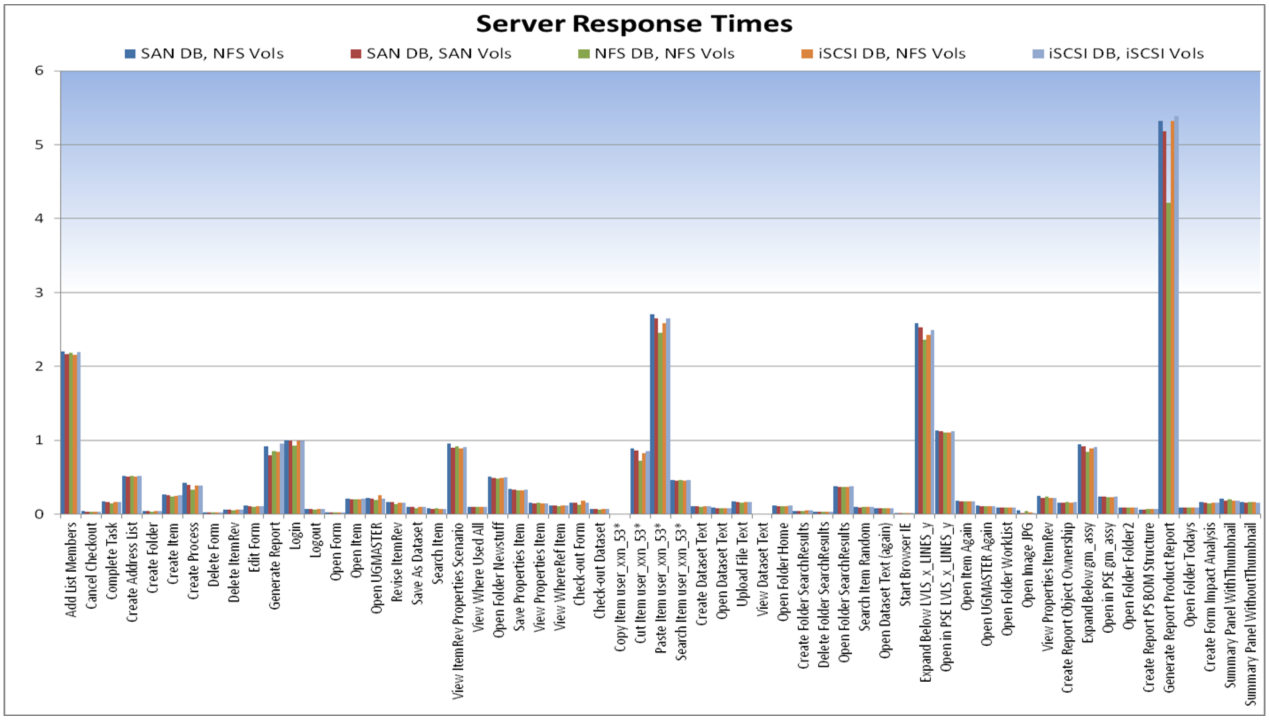

Figure 9) UNIX (Suse) FMS throughput results (seconds).

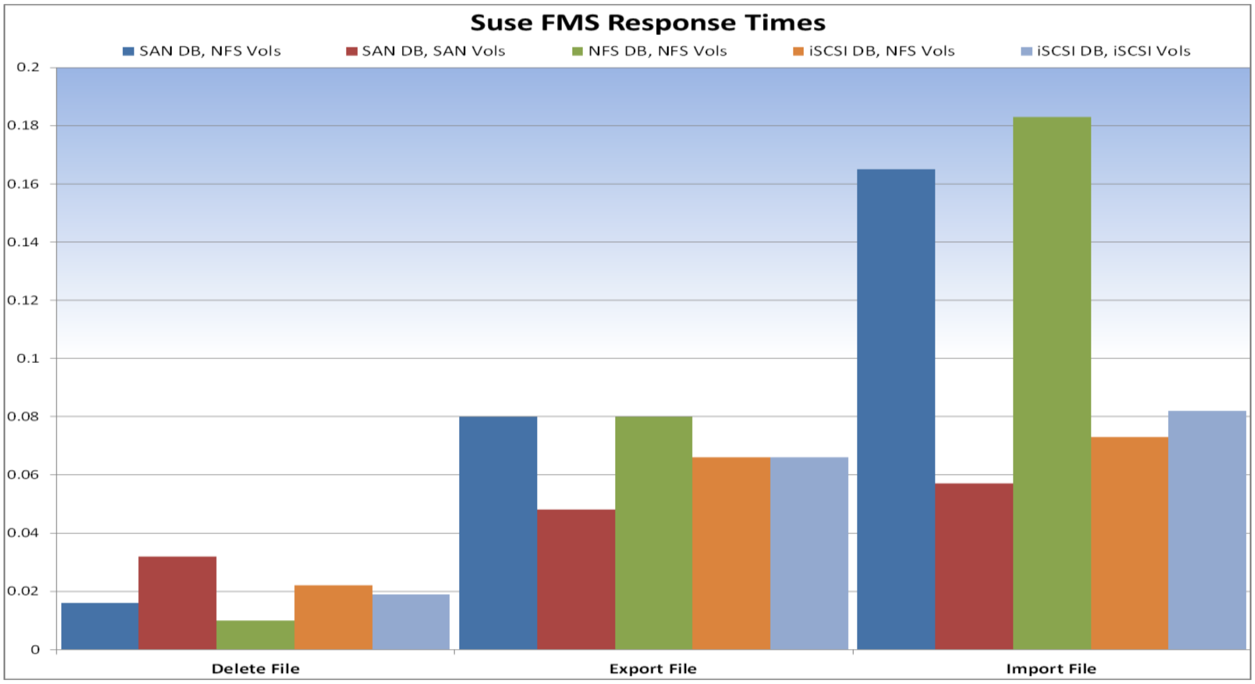

Windows

Figure 10) Windows results @ 500 users.

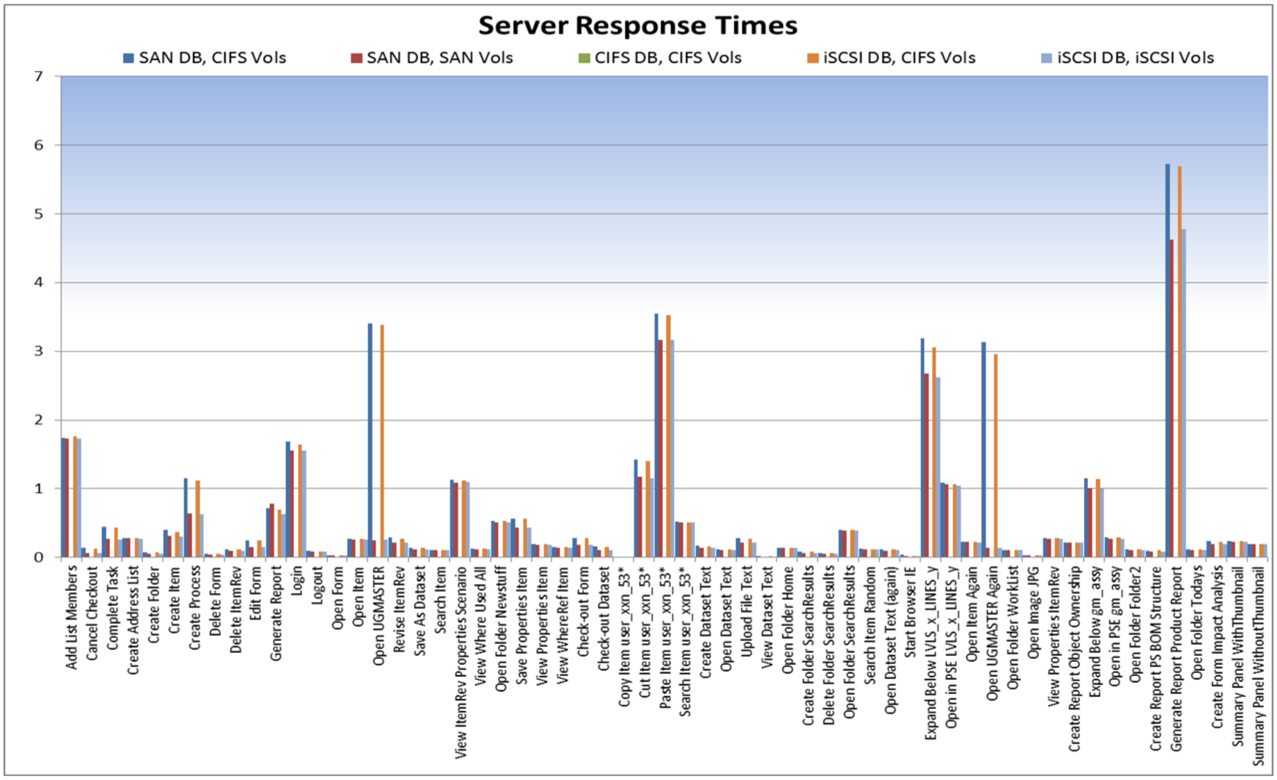

Figure 11) Windows FMS throughput results.

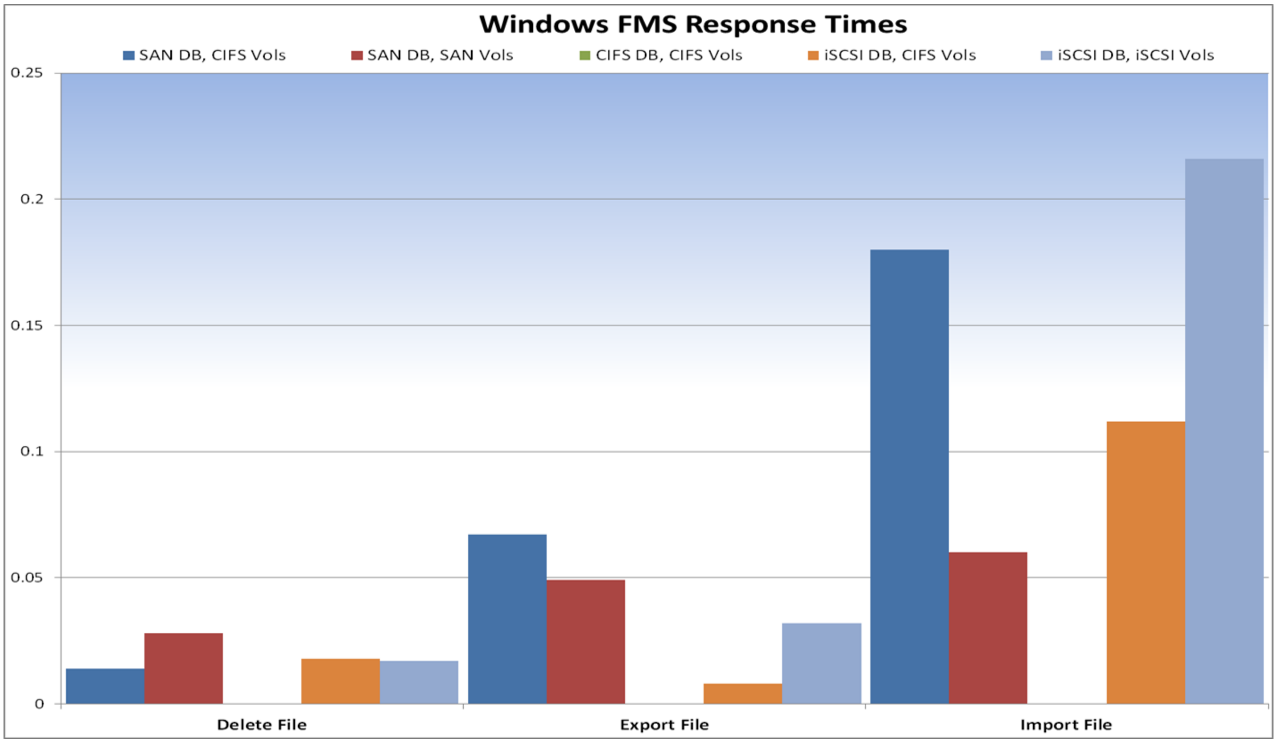

4.6 Analysis

UNIX

The key takeaways from running the benchmarks include:

- For the scalability benchmark, Teamcenter performance across the protocols was close and comparable.
- For the FMS throughput benchmarks, Teamcenter performance is fastest on FC SAN, followed by iSCSI, NFS, and then the split configuration (DB on SAN and file vaults on NAS).

Similar conclusions can be observed when looking at the weighted average of the server response time during the run across protocols, as illustrated in the table below.

Table 1) UNIX server response times.

Protocol DB, File Vaults	Scalability Benchmark Weighted Average (seconds)	FMS Throughput Benchmark Weighted Average (milliseconds)
FC, FC	0.42	67
iSCSI, iSCSI	0.42	103
NFS, NFS	0.41	130
FC, NFS	0.44	107
iSCSI, NFS	0.42	148

Windows

Similarly in the Windows environments:

- For the scalability benchmark, Teamcenter performance across the protocols was fastest when both database and file vaults were on SAN protocols and slowest on the split configuration.
- For the FMS throughput benchmarks, Teamcenter performance is fastest on FC SAN, followed by the split configuration (DB on SAN and file vaults on NAS) and then ISCSI.

Similar conclusions can be observed when looking at the weighted average of the server response time during the run across protocols in the Windows environment, as illustrated in the table below.

Table 2) Windows server response times.

Protocol DB, File Vaults	Scalability Benchmark Weighted Average (seconds)	FMS Throughput Benchmark Weighted Average (milliseconds)
FC, FC	0.48	53
iSCSI, iSCSI	0.42	106
FC, CIFS	0.63	91
iSCSI, CIFS	0.65	27

5 Volume Move Tests

One of the main features of the Data ONTAP system operating in Cluster-Mode is the ability to move volumes nondisruptively. The scalability and FMS throughput benchmarks described in the previous sections were executed while performing a volume move of the database or file vaults to observe the effects or impacts that `vol move` had on performance.

During the scalability benchmark, a volume move of the database was executed about 45 minutes into the benchmark run when about 900+ users had logged in. The following was observed.

- Database volume move from an aggregate in node 1 to an aggregate on node 2 completed within 5 minutes.
- No response time impact was observed.
- No transaction failures were observed.

Similarly, during the FMS throughput benchmark about 5 minutes into the run, the file vaults volume was moved from an aggregate on node 1 to an aggregate on node 2. The following was observed.

- File vaults volume move from an aggregate in node 1 to an aggregate on node 2 completed within 22 minutes.
- Occasional brief spikes in FMS server response time were observed.
- No transaction failures were observed.

These tests were conducted when no other activity other than the benchmarks was occurring on either node. Thus, in general, NetApp recommends scheduling `vol move` operations when the storage is not heavily loaded, because it does take up some storage CPU cycles.

6 Deployment Options

Teamcenter database and file vaults can be deployed on NetApp Data ONTAP operating in Cluster Mode storage for enhanced scalability and flexibility and simplified administration. Using Cluster-Mode, Teamcenter database and file vaults can be deployed on any NetApp storage protocol: FC SAN, iSCSI, NFS, and/or CIFS or in a split configuration in which the database is on a SAN and file vaults on a NAS. The table below highlights the advantages and disadvantages of deploying Teamcenter database and file vaults on each of NetApp's storage protocols.

Table 3) Pros and cons of Teamcenter deployment.

Deployment	Pros	Cons
FC SAN	▪ Performance, generally, is fastest when compared with other protocols.	▪ More expensive because it requires a dedicated fiber-optic storage network. ▪ Extra management of switches. ▪ Backup and recovery are complex and require more steps compared to NAS.
iSCSI	▪ Performance second to FC SAN. ▪ More cost efficient than FC SAN because it uses the existing Ethernet infrastructure.	▪ Backup and recovery are complex and require more steps compared to NAS.

Deployment	Pros	Cons
NFS	▪ Inexpensive. ▪ Simple to manage. ▪ Offers ease of use and granularity in backup and recovery operations.	▪ I/O performance is slower compared to other protocols. However, performance is adequate for most deployments.
Split configuration (database on iSCSI and pools on NAS)	▪ Improves Teamcenter database performance. ▪ Cost efficient if database deployed on iSCSI compared to FC.	▪ Backup and recovery are complex because backups are performed by using two separate protocols while maintaining concurrency between them. However, using NetApp SnapDrive® data management software alleviates this complexity.
CIFS	▪ Provides direct file access to Windows clients. ▪ Inexpensive. ▪ Simple to manage. ▪ Offers ease of use and granularity in backup and recovery operations.	▪ I/O performance is slower compared to other protocols. However, performance is adequate for most deployments.

7 Planning

Selecting the right storage for Teamcenter data to address the needs and requirements of a particular Teamcenter deployment is a complex task. The sizing policies are application specific and vary according to the capacity requirements and workload of every project. However, here are some questions to ask potential joint NetApp and Teamcenter customers to assist in sizing the project and understanding the customer’s needs and requirements.

- What are their main storage criteria: performance, data protection, reliability, manageability, and so on. If performance is the main criterion, ask them to define their expected latency.
- If they have an existing Teamcenter database and file vaults, what are the current sizes and their expected rate of growth?
- How many product developers will access the Teamcenter database and file vaults?
- What is their existing infrastructure: platform (servers and storage devices), protocol, and network infrastructure details. Did they deploy Teamcenter as a two-tier or four-tier configuration? If four-tier, do they have a distributed setup?
- Is storage to be used solely for Teamcenter applications?
- What is their current backup and restore mechanism?

The NetApp Unified Storage Architecture offers customers several options for deploying Teamcenter repositories. Table 6 lists the different NetApp storage protocols to determine which deployment best addresses the challenges of product development, such as performance, heterogeneous access, manageability, cost, and ease of backup and recovery. In terms of performance, FC SAN is the best; however, in terms of manageability, heterogeneous access, and backup and recovery, it may not be the ideal choice. In terms of cost, a NAS solution would be the best choice. However, if the customer wants a balance of all the factors, the best choice is the split configuration in which the database is on iSCSI or FC/SAN for performance and the source pool is on NAS for better access, cost, manageability, and backup and recovery. Thus, there are trade-offs when deciding which type of storage to select.

Table 4) Selection criteria across NetApp storage protocols.

Protocol / Criterion	FC	iSCSI (Software)	iSCSI and NAS (Split Configuration)	NFS/CIFS
Performance	Excellent	Good	Good	Good
Direct file access to hetero clients	No	No	Yes	Yes
Manageability	Fair	Fair	Good	Excellent
Cost	Fair	Excellent	Excellent	Excellent
Granularity and ease of backup and recovery	Fair	Fair	Good	Excellent

8 Storage Efficiency

There are two NetApp technologies that can further improve the storage efficiency of Teamcenter file vaults. The use of NetApp deduplication and FlexClone technology can provide storage savings and thus reduce or defer IT costs.

8.1 Deduplication and Compression

Deduplication is a key feature that can significantly reduce the data footprint of product developer workspaces in Teamcenter environments. NetApp deduplication runs as a background process and the system can perform any other operation during this process. However, NetApp highly recommends scheduling deduplication when the application is quiescent, possibly at night and on weekends.

The main reason for such a huge disk saving in the NetApp storage is because Teamcenter files (files with .prt extensions) normally write the full copy of the same file repeatedly whenever there is a change to that file. The Teamcenter application does not just write the delta changes. Because of this application behavior, many common blocks are written multiple times. Deduplication removes the redundant blocks, saving storage disk space. The percentage of disk savings achieved from this test may actually be higher in a real production environment, because more revisions are likely to be created for the same files.

Deduplication was conducted on the file vaults used for the benchmarks and the observed storage savings was 76%.

Note: Deduplication highly depends on the dataset, and results can vary according to your dataset.

For more information about NetApp deduplication and compression, see TR-3966: NetApp Data Compression and Deduplication Deployment and Implementation Guide: Data ONTAP 8.1 Operating in Cluster-Mode.

8.2 FlexClone

Product development often includes creating writable copies of datasets for testing, debugging, and verification. Without FlexClone thin-cloning technology, creating these copies would consume a lot of space and time. FlexClone technology makes it possible to create writable space-efficient clones in seconds. These clones provide substantial storage savings because they consume space only for data blocks that have been changed.

Here are some examples of how FlexClone can improve storage efficiency and productivity in a Teamcenter environment.

- Test new versions or relevant patches of Teamcenter without affecting the production systems. When the new version has been tested and confirmed, the FlexClone volume can be split and used as the new production volume.
- Run diagnostics on live production environments and preserve the original Teamcenter database and file vaults. The FlexClone software allows debugging work only on the copy.
- Perform performance testing in which Teamcenter database and file vaults can be copied from a baseline and subjected to various performance parameters. When the test run is completed, the FlexClone volume can be destroyed and a new FlexClone volume created from the Snapshot copy.

For more information on NetApp FlexClone, see TR-3347: A Thorough Introduction to FlexClone Volumes.

9 Best Practices

NetApp recommends the following best practices when deploying Teamcenter on NetApp Data ONTAP storage operating in Cluster-Mode.

Disk and Layout Best Practices

- Create and lay out volumes with recovery and manageability in mind; use separate volumes for Teamcenter DB and file vaults. Also, store system tables, user data, and temporary data and archive logs on separate volumes on a single aggregate on the storage system. This will make the recovery process easier since there is data separation.
- Consider the number of spindles for your database. For a database with an I/O-intensive workload, selection of the number of spindles can affect performance. Setting this number too large or too small can affect overall performance and recovery of disks. Using NetApp default values is a recommended best practice to start.
- When planning disk storage space of volumes or LUNs, reserving space for Snapshot copies must be considered.

Teamcenter Application Setting Considerations

Teamcenter applications can be tuned in each tier of its configuration. The Web Tier and Enterprise Tier have parameters that can be tuned, including the following.

- Web Tier
  - Set any JSP or servlet check intervals very long; for instance, once a day or longer.
  - Tune the application Java Virtual Machine (JVM) runtime parameters for high transactional throughput:
    - Set JVM heap sizes to match memory and CPU resources available on the machine.
    - Select the appropriate garbage collection algorithm to match the number of CPUs available.
    - Set JVM generational sizes accordingly based on GC algorithms and available heap size.
- Enterprise Tier—Server Manager

- Server Manager pool-specific parameters can be tuned to make better use of CPU and memory resources by parallelizing Teamcenter processes.

Please refer to the Teamcenter Deployment Guide for more information on Teamcenter tuning considerations.

Database Setting Considerations

In the resource tier a database is required to store Teamcenter metadata. The databases that are currently supported by Teamcenter include Oracle, IBM DB2, and Microsoft® SQL Server®. These are some of the database settings that should be considered when using NetApp storage.

IBM DB2 Database

- DB2_PARALLEL_IO – This parameter parallels the I/O if the number of containers in a tablespace is greater than 1 when reading and writing data in the tablespace containers.
- When using a file system as opposed to raw devices, consider these three parameters:
- EXTENTSIZE – This variable specifies the number of PAGESIZE pages written to a container before going to the next container. Since NetApp's disk segment size is 256KB (32 x 4KB blocks), use extent sizes that are multiples of 256KB. Setting this variable appropriately will allow a more efficient big-block I/O performed by prefetchers.
- PREFETCHSIZE – This defines the prefetch size of a tablespace. Setting the prefetch size reduces query response time since it will read ahead the pages of table data that will be needed for the subsequent query. A best practice is to define a prefetch value that is large enough to parallelize use of available containers. As a rule of thumb, the prefetch size should be a multiple of the extent size used multiplied by the number of data disks used. Thus, for a raid group size of 18, the prefetch size would be 16*256=4096KB.
- NO FILE SYSTEM CACHING - Disabling the File System Caching parameter in the tablespace can improve performance since it does not allow double caching (caching in the kernel file system level and database caching).

Use of the DB2 Configuration Advisor is also helpful in further tuning the configuration parameters.

Oracle11*g*

Oracle parameters that assist I/O performance include the following.

- Direct NFS (DNFS) – If you use the NFS protocol, use of an Oracle DNFS client provides direct access to the database engine without relying on the host's operating system NFS client. This enhances performance, manageability, availability, and scalability.
- DB_FILE_MULTIBLOCK_READ_COUNT – This determines the maximum number of database block reads in one I/O operation during a full table scan. Setting this value to a multiple of the NFS READ/WRITE size (convert the NFS size in bytes to blocks) will limit the number of fragmentation occurrences.
- DB_BLOCK_SIZE – This is the size of Oracle Database blocks. This value should be set to be a multiple of the physical block size at the device level. For the best performance, NetApp recommends that the DB_BLOCK_SIZE be a multiple of the OS block size. Also, if you use the NFS protocol, the NFS rsize and wsize should also be a multiple of the DB_BLOCK_SIZE.
- DISK_ASYNCH_IO – This controls whether I/O to data files, control files, and log files is asynchronous. If asynchronous I/O is supported on the platform, it is advisable to leave the default setting of TRUE in order to parallelize I/O requests with CPU processing during table scans.
- DB_WRITER_PROCESSES – Setting this parameter parallelizes the gathering and writing of buffers. It is usually used in conjunction with DISK_ASYNCH_IO set to TRUE.

These are only some of the best practices related to these databases. For other best practices relating to databases, refer to the NetApp best practices technical reports on Oracle, IBM DB2, and Microsoft SQL.

10 Support Matrix

It is advisable to always check the support matrices for both Siemens PLM and NetApp to check whether the version of the operating system, hardware, and software required to deploy Siemens PLM and NetApp is supported. For Siemens PLM, refer to the Siemens PLM Web site for its support matrix. For NetApp, refer to the NetApp Interoperability Matrix for the current qualifications and system requirements for your specific configuration, including the following:

- Operating system version and patch requirements
- HBA models and drivers
- Supported versions of Data ONTAP
- Supported configurations
- SAN booting requirements

11 Limitations

At the time of this writing, there are limitations when deploying Teamcenter on the current release of NetApp Data ONTAP operating in Cluster-Mode. Future releases of Cluster-Mode will address these limitations, and this document will be updated accordingly. The limitations include the following.

- Deployment of Teamcenter DB or file vaults on iSCSI for an IBM AIX® host is not currently supported.
- For SAN deployment, support for hosts other than Linux® and Windows will not be supported until NetApp Data ONTAP 8.2.
- Cluster-Mode does not currently support the following features. If customers have the following requirements, do not recommend upgrading to Cluster-Mode.
  - MetroCluster™ software
  - SnapVault® software
  - SnapLock® software
  - IPv6
  - Qtree SnapMirror
  - SnapMirror Sync

12 Troubleshooting

Since Teamcenter is a multitier application, log files that are generated can be placed in different locations. However, Teamcenter does provide a mechanism called the Log Manager to centralize and consolidate the log files generated across Teamcenter. For comprehensive information on Log Manager and the specifics on the error logs, refer to the Siemens PLM System Administration Guides. Since Teamcenter also requires a database, please refer to the administration guides for specific databases that you are utilizing in your Teamcenter deployment. This section focuses only on the Teamcenter logs that are useful when troubleshooting Teamcenter on NetApp storage.

12.1 Teamcenter Logs To Check

Teamcenter logs that are helpful in indicating if there are errors with accessing files or data on the NetApp storage include:

- tcserver*pid*.syslog – Tracks actions performed on objects at a session level such as folder creation. This log file is in the directory specified by the TC_TMP_DIR environment variable defined in the tc_profilevars.bat file. Usually it is defined to be /tmp or /var/tmp in a UNIX environment or C:\Temp in Windows.
- security.log – Tracks access to unauthorized data and failed logon attempts. This log file is in the directory specified by the TC_LOG environment variable defined in the tc_profilevars.bat file. Usually it is defined to be /tmp or /var/tmp in a UNIX environment or C:\Temp in Windows.
- *FSC_ID*_startup.log – Contains information on the server runtime operations of the FMS. To be able to access the file vaults, it is important that the file server cache start without errors. This file will provide information on runtime of the FMS. It is located on /tmp in UNIX or %FMS_HOME% on Windows on the server running the file server cache of the FMS.

12.2 Performance Issues

Troubleshooting performance issues requires isolating the bottleneck, which can be on the server, client, network, or storage. For storage performance issues, the Perfstat tool is useful in identifying storage bottlenecks or volumes that are hot.

- Performance and Statistics Monitor (perfstat): Perfstat is a simple Bourne shell script that captures performance and configuration statistics.

13 Conclusion

NetApp continues to offer compelling solutions for Teamcenter customers with its next generation of storage products. NetApp Data ONTAP operating in Cluster-Mode provides a scalable storage infrastructure that helps accelerate product development, improve data protection, and minimize costs for Teamcenter storage infrastructure. NetApp has a unique relationship with Siemens PLM that helps make NetApp the safe choice.

14 References

- Backup and Recovery of Teamcenter on NetApp Data ONTAP Operating in Cluster- Mode, NetApp and Siemens PLM
- Data ONTAP 7G documents on Teamcenter: TR-3658 – Best Practices for Deploying Siemens PLM Software Teamcenter 2005 SR1 and 2007 MP3 (2007.1.3 by Using NetApp Storage Systems, by Bikash R. Choudhury, NetApp, and Bill Halpin, Siemens PLM; and TR3754 – Siemens PLM Software – Teamcenter Backup and Recovery, by Bikash R. Choudhury and Anand Ranganathan, NetApp
- TR-3982 – Data ONTAP 8.1 Operating in Cluster-Mode: An Introduction, by Charlotte Brooks

15 Appendix

15.1 Configuration and Setup

This section describes using the command line to configure and set up Data ONTAP operating in Cluster-Mode storage for use in Teamcenter environments. You can also easily configure Cluster-Mode storage by using the System Manager GUI and the Data ONTAP Element Manager. This is not a comprehensive guide; for more information, see “Teamcenter on NetApp Storage” for further troubleshooting. For comprehensive information, refer to the Siemens PLM System Administration Guide and the NetApp Cluster-Mode Administration Guide. Here are the high-level steps to create volumes for a TeamCenter database and file vaults in Cluster-Mode.

1. Create an aggregate.
2. Create a Vserver.
3. Create volumes.
4. Depending on the type of access (NFS, CIFS, or SAN), create the appropriate protocol service on the Vserver.
5. Create one or more network interfaces for the type of access.
6. Create user groups, UNIX users, and interface groups as needed by the protocol.
7. Access or create the file system.
8. Mount the volume.
9. Install Teamcenter components or database using the volumes mounted.

The following sections provide a quick guide to creating Teamcenter database and file vault volumes for NFS, CIFS, and iSCSI in both Windows and UNIX environments.

15.2 Aggregate Create

Define the storage for the Vserver by creating the aggregate. An aggregate is the physical storage; create it by using:

```
storage aggregate create [-aggregate] <aggregate name> [-diskcount] <integer>
```

The minimum requirements to create an aggregate are the name of the aggregate and the number of disks. However, it is important to note which nodes the disks are physically attached to. Knowing the nodes on which the disks physically reside is necessary in order to understand how to best balance the storage across the nodes.

For example, to create an aggregate named `aggr_test` with a `diskcount` of 5, execute the following command:

```
cl_agnes_cmode::> storage aggregate create -aggregate aggr_test -diskcount 5
```

To view what has been created, execute the following command. For this aggregate, the disks reside on a node named `fas3170c-svl11`.

```
       cl_agnes_cmode::> storage aggr show

Aggregate     Size Available Used% State    #Vols Nodes            RAID Status
--------- -------- --------- ----- ------- ------ ---------------- ----------
aggr_test     0B     0B     - creating   0 fas3170c-svl11    raid_dp, initializing
```

15.3 Vserver Create

1. Create a Vserver to use the aggregate created in the previous section. The following command creates a root volume named `test_root` on aggregate `aggr_test` of `mixed` security style:

   ```
   cl_agnes_cmode::> vserver create -vserver test -rootvolume test_root -aggregate
   aggr_test -ns-switch file -rootvolume-security-style mixed –unix-permissions ---
   rwxrwxrwx
   [Job 6537] Job succeeded: Successful
   ```

2. To view the Vserver, enter:

   ```
   cl_agnes_cmode::> vserver show
                       Admin     Root                  Name    Name
   Vserver     Type    State     Volume     Aggregate  Service Mapping
   ```

```
----------- ------- --------- ---------- ---------- ------- -------
test        cluster running   test_root  aggr_test  file    file
```

3. Modify the root volume test_root permissions so that others can access the volumes created beneath or mounted on top of this Vserver.

```
volume modify -vserver test2 -volume test_vol_vol -unix-permissions ---rwxrwxrwx
```

4. View the instance of Vserver. Creation of a Vserver automatically enables protocols NFS, CIFS, FCP, and iSCSI.

```
cl_agnes_cmode::> vserver show -vserver test -instance
```

Vserver: test

Vserver Type: cluster

Vserver UUID: d368e7c5-25b6-11e1-8a92-123478563412

Root Volume: test_root

Aggregate: aggr_test

Name Service Switch: file

Name Mapping Switch: file

NIS Domain: –

Root Volume Security Style: mixed

LDAP Client: –

Language: C

Snapshot Policy: default

Comment:

Anti-Virus On-Access Policy: default

Quota Policy: default

List of Aggregates Assigned: aggr_test

Limit on Maximum Number of Volumes Allowed: unlimited

Vserver Admin State: running

Allowed Protocols: nfs, cifs, fcp, iscsi

Disallowed Protocols: –

15.4 Volume Create

1. Create a volume in the Vserver:

```
cl_agnes_cmode::> vol create -volume test_vol -vserver test -aggregate aggr_test -
size 20MB -state online -type RW -policy default -unix-permissions ---rwxrwxrwx
(volume create)
[Job 6538] Job succeeded: Successful
```

2. Two volumes are now present in the Vserver. The test_root volume was created during the Vserver create and should not be used for data. This test_root volume is to keep track of Vserver metadata and acts as the root volume. This is the volume on top of which all volumes will be

mounted. The other volume, `test_vol`, was created for data and will be mounted on the root volume.

```
cl_agnes_cmode::> vol show -vserver test
  (volume show)
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
test      test_root    aggr_test    online     RW         20MB    18.90MB    5%
test      test_vol     aggr_test    online     RW         20MB    18.90MB    5%
2 entries were displayed.
```

3. Mount the volume for access:

```
cl_agnes_cmode::> vol mount -vserver test -volume test_vol -junction-path /test_vol
-active true
```

15.5 NFS Access

The example commands in this section are to create an NFS v3 server on Vserver `test` and to set up default policies and export rules. For specific information on how to set up with a DNS server, NFSv4, and local netgroups, see the “Cluster-Mode Administration Guide” on the NetApp Support site (formerly NOW®).

1. Create an NFS service for Vserver `test`. In this scenario, enabled for v3:

```
cl_agnes_cmode::> vserver nfs create -vserver test -access true -v3 enabled
```

2. Create a policy rule for Vserver `test`:

```
cl_agnes_cmode::> vserver export-policy rule create -vserver test -policyname
default -clientmatch 0.0.0.0/0 -rorule any -rwrule any -anon  0 -superuser never
```

3. Create a network interface for clients to access NAS. Items to specify include the Vserver:
   a. LIF name
   b. Home node: `fas3170c-svl11` (it is good to specify the node where the disks reside; for this scenario, the disks in aggregate `aggr_test` reside on node `fas3170c-svl11`)
   c. Home port: `e4a` (this is the home port of the home node specified in previous step b)
   d. IP address: `172.31.8.224`

```
cl_agnes_cmode::> network interface create -vserver test -lif test_lif -role data -
data-protocol nfs,cifs,fcache -home-node fas3170c-svl11 -home-port e4a -address
172.31.8.224 -netmask 255.255.255.0 -status-admin up

Info: Your interface was created successfully; the routing group d172.31.8.0/24 was
created
```

4. Set up the UNIX group and UNIX user for the Vserver `test` for user `root`.

```
cl_agnes_cmode::> services unix-group create -vserver test -name root –id 0

cl_agnes_cmode::> services unix-user create -vserver test2 -user root -id 0 -
   primary-gid 0

cl_agnes_cmode::> services unix-user show -vserver test
```

```
               User            User   Group  Full
Vserver        Name            ID     ID     Name
-------------- --------------- ------ ------ ---------
test           root            0      0      -
2 entries were displayed.
```

5. Mount the volume on the desired host. As root, execute the following:

```
mkdir /test_vol
mount –o vers=3 172.31.8.224:/test_vol /test_vol
```

15.6 CIFS Access

This example creates a CIFS service on Vserver `test`.

1. Create DNS on a Vserver.

```
cl_agnes_cmode::> vserver services dns create -vserver test -domains TEAMCENTER -
state enabled -name-servers 172.31.8.102


cl_agnes_cmode::> dns show
  (vserver services dns show)teamcenter
                                                                       Name
Vserver         State     Domains                            Servers
--------------- --------- ---------------------------------- ----------------
test            enabled   teamcenter.local                    172.31.8.102
```

2. Create a CIFS service on the Vserver. When prompted, enter the user name and password.

```
cl_agnes_cmode::vserver cifs> cifs create -vserver test -cifs-server test_cifs -
domain teamcenter.local
In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "CN=Computers" container within the
"teamcenter.local" domain.

Enter the user name: Administrator

Enter the password:
```

3. Create the network interface (if it has not already been created). Note that CIFS and NFS traffic can use the same network interface and LIF. Therefore, if you have already created an NFS LIF, you can use the same LIF.

   a. Vserver

   b. LIF name

   c. Home node: `fas3170c-svl11` (it is good to specify the node where the disks reside; for this scenario, the disks in aggregate `aggr_test` reside on node `fas3170c-svl11`)

   d. Home port: `e4a` (this is the home port of the home node specified in c)

   e. IP address: `172.31.8.224`

```
   cl_agnes_cmode::> network interface create -vserver test -lif test_lif -role
   data -data-protocol nfs,cifs,fcache -home-node fas3170c-svl11 -home-port e4a -
   address 172.31.8.224 -netmask 255.255.255.0 -status-admin up

   Info: Your interface was created successfully; the routing group d172.31.8.0/24
   was created
```

4. Share the volume:

```
cl_agnes_cmode::> cifs share create -vserver test -share-name test_vol -path
   /test_vol


cl_agnes_cmode::> cifs share show
Vserver        Share         Path        Properties Comment  ACL
---------- ---------- ----------- ---------- -------- -----------
test         test_vol      /test_vol    oplocks    -        Everyone / Full Control
```

5. Create a user group and users for root if they have not already been created:

```
cl_agnes_cmode::> services unix-group create -vserver test -name root –id 0

cl_agnes_cmode::> services unix-user create -vserver test -user root -id 0 -
   primary-gid 0

cl_agnes_cmode::> vserver services unix-group show -vserver test
Vserver        Name                ID
-------------- ------------------- ----------
test           root                0
2 entries were displayed.

cl_agnes_cmode::> vserver services unix-user show -vserver test
               User            User   Group  Full
Vserver        Name            ID     ID     Name
-------------- --------------- ------ ------ --------------------------------
test           root            0      0      -
```

6. Create user mappings. Provide Vserver, direction, position (number order 1–1,024), pattern (Windows login), and replacement (UNIX login):

```
cl_agnes_cmode:: > vserver name-mapping create -vserver test  -direction win-unix -
position 1 -pattern  TEAMCENTER\\Administrator -replacement  root
```

7. Access the CIFS share from a Windows machine. For example, run the following:

```
\\172.31.8.240\test_vol
```

15.7 Interop Access (Both CIFS and NFS Access)

In order to access the volume from both Windows and UNIX machines, the following must be true:

- The security style setting of the volume must be set to `mixed`. If the security style is currently set to UNIX or NTFS, use the following command to modify it:

```
cl_agnes_cmode::> volume modify -vserver test -volume test_vol -security-style
mixed
```

- Mappings for `unix-win` and `unix-win` users should be set appropriately. For example:

  ```
  cl_agnes_cmode:: > vserver name-mapping create -vserver test  -direction win-unix -
  position 1 -pattern  TEAMCENTER\\Administrator -replacement  root
  ```

  ```
  cl_agnes_cmode:: > vserver name-mapping create -vserver test  -direction unix-win -
  position 2 -pattern root -replacement TEAMCENTER\\Administrator
  ```

15.8 iSCSI Access

iSCSI setup differs depending on the host operating system that requires access to the volume via the iSCSI protocol. For details, refer to the Installation and Administration Guide of the host utilities for the specific operating system on the NetApp Support site (formerly NOW). The example setup for iSCSI in this section is for the Linux operating system and Windows 2008 R2 versions.

1. Create an iSCSI service on Vserver test:

   ```
   cl_agnes_cmode::> iscsi create –vserver test
   ```

2. Create a management LIF, which will be used by SnapDrive to access the Vserver and issue ZAPIs:

   ```
   cl_agnes_cmode::> network interface create -vserver test -lif test_mgmt -role data
   -data-protocol none -home-node fas3170c-svl11 -home-port e4b -address 172.31.8.247
   -netmask 255.255.255.0 -status-admin up -firewall-policy mgmt
   ```

3. Unlock the user ID and assign the password `vsadmin` for the Vserver administrator. The password was created when the Vserver was created:

   ```
   cl_agnes_cmode::> security login password -vserver test -username vsadmin

   Please enter a new password:
   Please enter it again:

   cl_agnes_cmode::> security login unlock -vserver test -username vsadmin


   cl_agnes_cmode::volume> security login show -vserver test
    Authentication Acct
   Vserver  UserName  Application Method Role Name  Locked
       -------  ---------  ------------------ ---------  ------
       test    vsadmin   http        password   admin     no
       test    vsadmin   ontapi      password   vsadmi    no
       test    vsadmin   ssh         password   vsadmin   no
   ```

4. Create an iSCSI data LIF for each node in your cluster:

   ```
   cl_agnes_cmode::> network interface create -vserver test -lif test_san09 -role data
   -data-protocol iscsi -home-node  fas3170c-svl09 -home-port e4b -address
   172.31.8.248 -netmask 255.255.255.0
   ```

   ```
   cl_agnes_cmode::> network interface create -vserver test -lif test_san10 -role data
   -data-protocol iscsi -home-node  fas3170c-svl10 -home-port e4b -address
   172.31.8.248 -netmask 255.255.255.0
   ```

   ```
   cl_agnes_cmode::> network interface show -vserver test
                    Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ```

```
----------- ---------- ---------- ------------------ ------------- ------- ----
test
test_mgmt    up/up    172.31.8.247/24    fas3170c-svl11  e4b     true
test_san09   up/up    172.31.8.248/24    fas3170c-svl09  e4b     true
test_san10   up/up    172.31.8.249/24    fas3170c-svl10 e4b     true
```

5. Create volumes for the LUNs:

```
cl_agnes_cmode::volume> volume create -vserver test -volume testsan_vol -aggregate
   aggr_test -size 500MB -state online -space-guarantee none
[Job 7175] Job succeeded: Successful

cl_agnes_cmode::volume> volume show -vserver test
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
test      testsan_vol      aggr_test    online     RW        500MB    474.9MB    5%
```

Configure Linux for iSCSI Access

On the system:

1. Create LUNs for Linux access:

```
cl_agnes_cmode::>  lun create -vserver test -volume testsan_vol -lun linux.lun1  -
size 250MB -ostype linux -space-reserve disabled

Created a LUN of size 250m (262144000)

cl_agnes_cmode::> lun show -vserver test
Vserver   Volume     Qtree      LUN        State   Mapped   Type        Size
--------- ---------- ---------- ---------- ------- -------- -------- --------
test      testsan_vol    ""         linux.lun1 online  unmapped linux       250MB
```

2. Get the name of the initiator on the Linux host:

```
# cat /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.1994-05.com.redhat:a860857dd1c4
```

3. Create an igroup using the initiator name created in step 2:

```
    cl_agnes_cmode::> igroup create -vserver test -igroup ilinux -protocol iscsi -
    ostype linux -initiator iqn.1994-05.com.redhat:a860857dd1c4

    cl_agnes_cmode::> igroup show -vserver test
    Vserver   Igroup       Protocol OS Type  Initiators
    --------- ------------ -------- -------- ------------------------------------
    test      ilinux       iscsi    linux    iqn.1994-05.com.redhat:a860857dd1c4
```

4. Map the LUN to the igroup:

```
    cl_agnes_cmode::> lun map -vserver test -volume testsan_vol -lun linux.lun1 -
    igroup ilinux

    cl_agnes_cmode::> lun show -vserver test
    Vserver   Volume     Qtree      LUN        State   Mapped   Type        Size
    --------- ---------- ---------- ---------- ------- -------- -------- --------
    test      testsan_vol
                         ""         linux.lun1 online  mapped   linux       250MB
```

On the host:

1. Download and install the latest iSCSI host utilities from the NetApp Support site for the Linux host. Refer to the “Host Utilities Manual” from the NetApp Support site (formerly NOW) for complete instructions on how to install the host utilities. The following next steps are quick instructions on how to configure Linux for iSCSI access. This is not meant to be comprehensive; refer to the “Host Utilities Manual” for more details.

2. Verify that the Suse Linux multipath packages are installed on the Linux server:

```
# rpm -q device-mapper
# rpm -q multipath-tools
```

3. Edit `multipath.conf` using an editor and copy and paste the following stanza into `/etc/multipath.conf`:

```
# cp /etc/multipath.conf /etc/multipath.conf.orig

(Copy the following stanza into multipath.conf. This for Suse Linux Enterprise
   Server with SP1 and ALUA)

defaults
{
   user_friendly_names no
   max_fds max
   flush_on_last_del yes
}
blacklist
{
   devnode "^hd[a-z]"
   devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
   devnode "^cciss.*"
}
devices
{
   device
   {
   vendor "NETAPP"
   product "LUN"
   getuid_callout "/lib/udev/scsi_id -g -u -d /dev/%n"
   prio "alua"
   features "1 queue_if_no_path"
   hardware_handler "1 alua"
   path_grouping_policy group_by_prio
   path_selector "round-robin 0"
   failback immediate
   rr_weight uniform
   rr_min_io 128
   path_checker tur
   }
}
```

4. Blacklist local drives by getting the WWID of the local drive and adding it to `/etc/multipath.conf`. For example:

```
# scsi_id -gus /block/sda
SATA_HDS728080PLA380_PFDB32S0R3WHJM
```

Edit `multipath.conf` based on the output of `scsi_id` of the root local drive:

```
blacklist {
        wwid SATA_HDS728080PLA380_PFDB32S0R3WHJM
        devnode "^hd[a-z]"
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^cciss.*"
}
```

5. Enable and start MPxIO for iSCSI:

```
# /etc/init.d/boot.multipath start
# /etc/init.d/multipathd start
```

However, to have the multipath start automatically while booting, enter the following commands:

```
chkconfig --add boot.multipath
chkconfig --add multipathd
chkconfig boot.multipath on
chkconfig multipathd on
```

6. Verify that `multipathd` is running and working:

```
# /etc/init.d/multipathd status
multipathd (pid  2314) is running...

To view a list of multipath devices enter:
   multipath -ll
```

7. Enter the following command to discover the iSCSI target:

```
iscsiadm --mode discovery --op update --type sendtargets –portal <targetip>
```

   targetIP is the IP address of NetApp filer (eg. 172.17.39.172:3260)

8. Enter the following command to see all the active iSCSI sessions.

```
iscsiadm --mode node –l all
```

9. Reboot to enable changes:

```
shutdown --r now
```

10. When the Linux host is back up, obtain the list of all current sessions using:

```
iscsiadm  –-m session
```

   To rescan all the sessions:

```
iscsiadm –m session --rescan
```

11. Use `sanlun` to verify the multipath policy and provider:

```
# sanlun lun show -p
                    ONTAP Path: test:/vol/testsan_vol/linux.lun1
                           LUN: 0
                      LUN Size: 250m
                          Mode: C
```

```
                   Host Device: 3600a09803246696e433f2d2d636e6b39
              Multipath Policy: round-robin 0
            Multipath Provider: Native
--------- ---------- ------- ------------ --------------------------------------
--------
host      vserver
path      path       /dev/   host         vserver
state     type       node    adapter      LIF
--------- ---------- ------- ------------ --------------------------------------
--------
up        primary    sdm     host10       test_san10
up        secondary  sdk     host9        test_san09
```

12. Verify multipaths:

```
  [root@ibmx3455-svl01 etc]# multipath -ll
3600a09803246696e433f2d2d636e6b39 dm-6 NETAPP,LUN C-Mode
[size=250M][features=1 queue_if_no_path][hwhandler=1 alua][rw]
\_ round-robin 0 [prio=50][active]
 \_ 11:0:0:0 sdm 8:192 [active][ready]
\_ round-robin 0 [prio=10][enabled]
 \_ 8:0:0:0  sdj 8:144 [active][ready]
 \_ 9:0:0:0  sdk 8:160 [active][ready]
 \_ 10:0:0:0 sdl 8:176 [active][ready]
```

13. Create a file system on the device:

```
# mkfs -t ext3 /dev/dm-6

mke2fs 1.39 (29-May-2006)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
64000 inodes, 256000 blocks
12800 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=67371008
32 block groups
8192 blocks per group, 8192 fragments per group
2000 inodes per group
Superblock backups stored on blocks:
        8193, 24577, 40961, 57345, 73729, 204801, 221185

Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

This file system will be automatically checked every 22 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

14. Mount the file system and validate the size:

```
#  mkdir /LUN
 # mount /dev/dm-6 /LUN

# df -h /LUN
Filesystem            Size  Used Avail Use% Mounted on
/dev/dm-6             243M  6.1M  224M   3% /LUN
```

Configure Windows for LUN Access

For LUN Windows access, setup can be done either manually or by using SnapDrive for Windows. This example is a manual setup on Windows 2008 R2. Following are instructions for Windows 2008.

1. Install MPIO. If MPIO is not configured in Windows, open the Server Manager. In the left pane of the Server Manager, expand Features and click Add Features. Once the Add Features Wizard launches, select Multipath I/O on the Features page and click Next. Enable MPIO for iSCSI LUNs. On the Confirm Installation page, click Install and, once installation completes successfully, click Close.
2. Enable MPIO. Click on Start menu and then the MPIO icon.

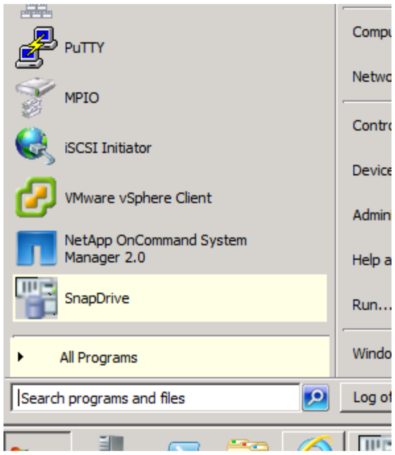

3. Add Support for iSCSI devices by placing a checkbox in the Discover Multi-Paths tab. Click Add, and wait for the Windows 2008 server to reboot.

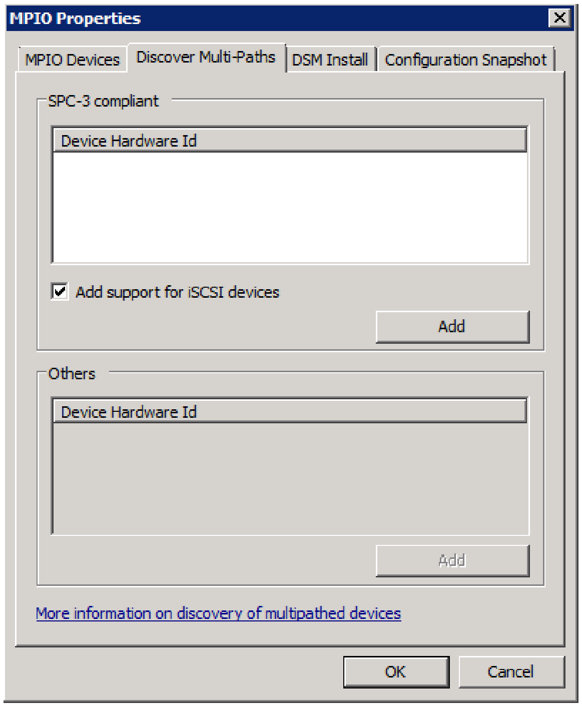

☑ Add support for iSCSI devices

4. After reboot, open the Control Panel and double-click on the iSCSI Initiator icon. If an error message indicates that service is not running, click Yes to start the service.

5. Click on the Discovery tab and Discover Portal.

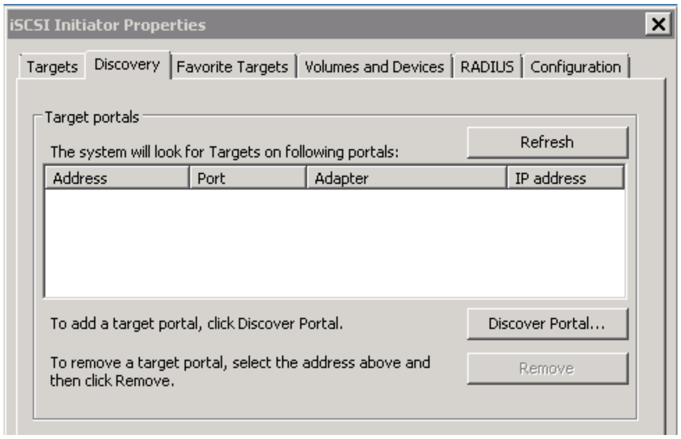

The system will look for Targets on following portals:

To add a target portal, click Discover Portal.

To remove a target portal, select the address above and then click Remove.

6. Enter the IP address of the iSCSI LIF created in section 15.8 and then click OK.

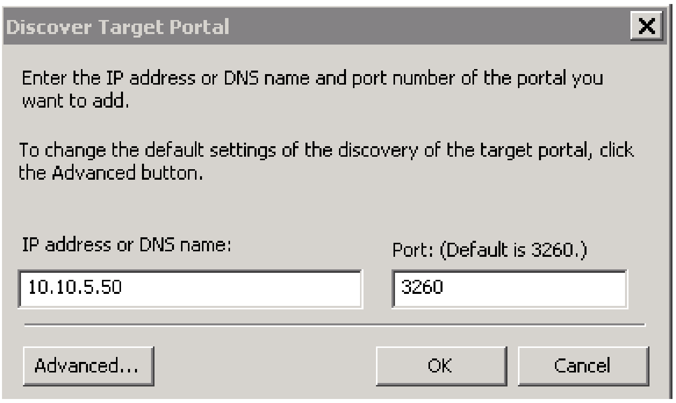

Enter the IP address or DNS name and port number of the portal you want to add.

To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name: 10.10.5.50

Port: (Default is 3260.) 3260

7. In the Targets tab, highlight the new Inactive connection and click Connect.

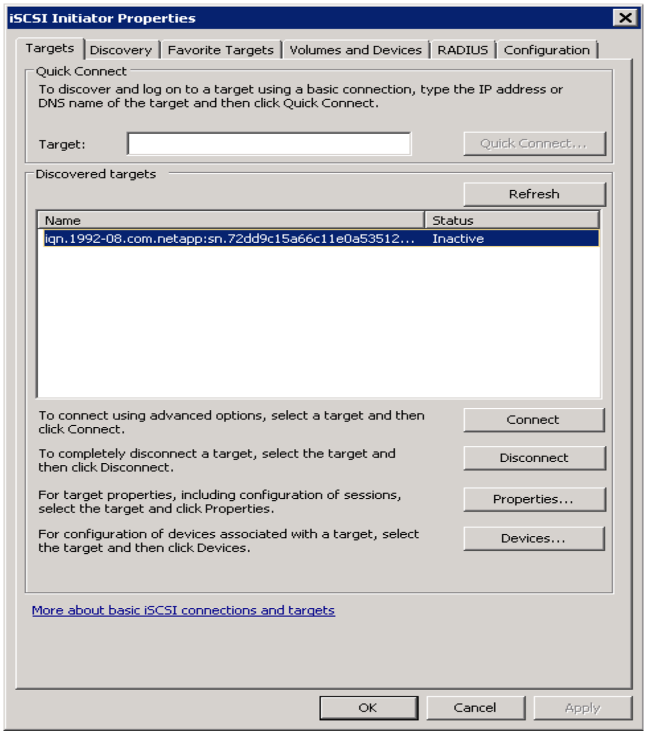

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target:

Name	Status
iqn.1992-08.com.netapp:sn.72dd9c15a66c11e0a53512...	Inactive

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

8. In the Connect To Target dialog, place a check on Enable multi-path, then click Advanced.

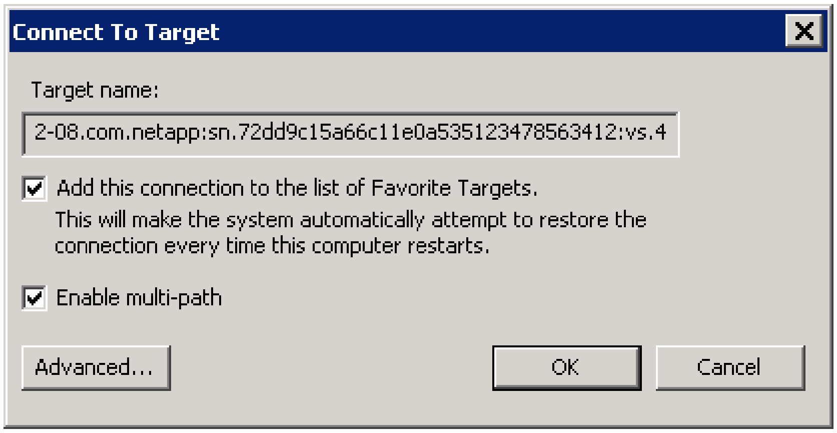

Connect To Target

Target name:

2-08.com.netapp:sn.72dd9c15a66c11e0a535123478563412:vs.4

☑ Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.

☑ Enable multi-path

9. In the Target Portal IP pull-down, select the TCP/IP address of the LIF you entered for the target portal, then click OK. Click OK again in the Connect To Target dialog.

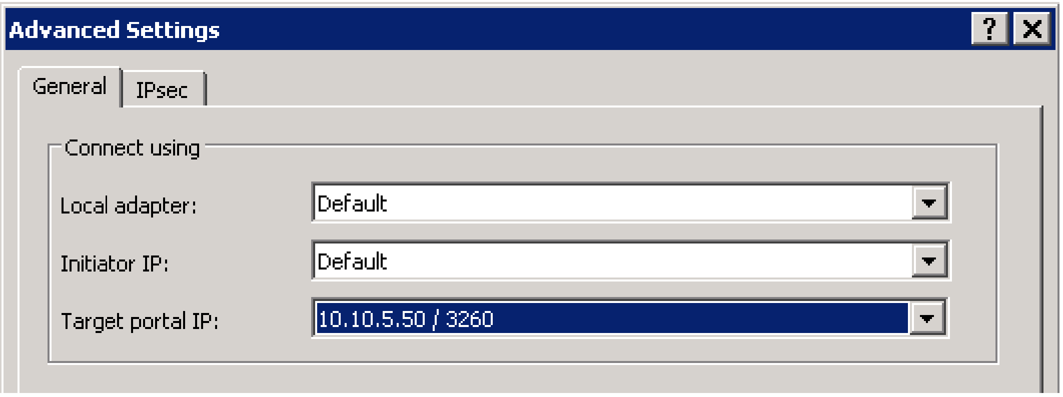

Advanced Settings

Connect using

Local adapter: Default

Initiator IP: Default

Target portal IP: 10.10.5.50 / 3260

10. Back in the Targets tab, highlight the connection again and click Properties.

11. Click Add session.

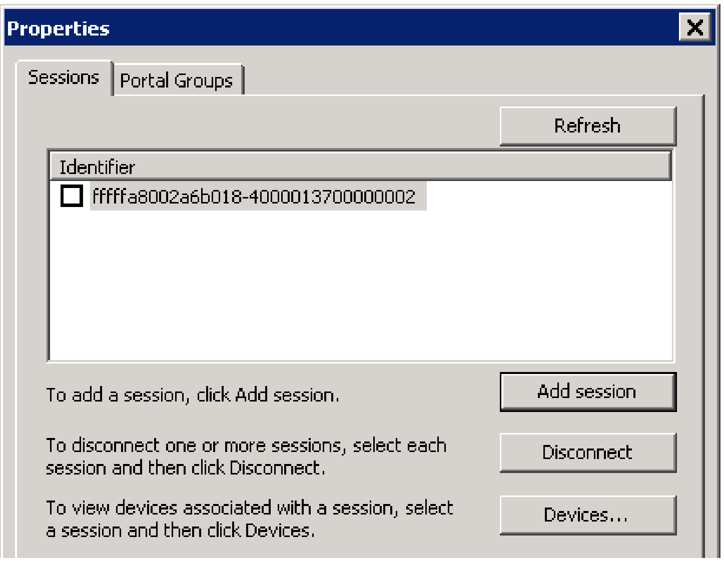

Properties

□ fffffa8002a6b018-4000013700000002

To add a session, click Add session.

To disconnect one or more sessions, select each session and then click Disconnect.

To view devices associated with a session, select a session and then click Devices.

12. In the Connect To Target dialog, check Enable multi-path and click Advanced.
13. This time, select the next IP address of the iSCSI LIFs from the Target Portal IP pull-down. Click OK, then click OK in the Connect To Target dialog.

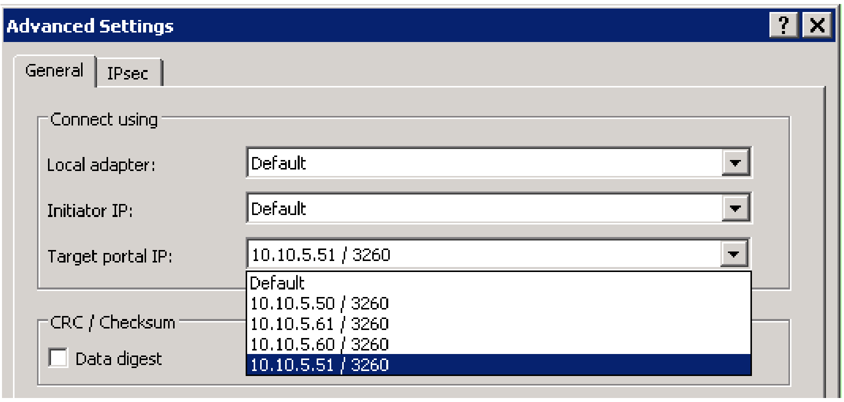

Advanced Settings

Connect using

Local adapter: Default

Initiator IP: Default

Target portal IP: 10.10.5.51 / 3260

14. Repeat the steps for the remaining iSCSI LIFs. When this is done, you should see the Sessions in the Properties tab. Click OK.
15. On the Cluster-Mode command line, create a LUN for Windows:

```
cl_agnes_cmode::> lun create -vserver test -volume testsan_vol -lun windows.lun1 -
ostype windows -size 250MB -space-reserve disabled

Created a LUN of size 251.0m (263208960)
```

16. Validate LUN creation:

```
cl_agnes_cmode::> lun show -vserver test

Vserver   Volume     Qtree      LUN        State   Mapped   Type        Size

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

test      testsan_vol   ""     windows.lun1  online  unmapped windows   251.0MB
```

17. Get the initiator for Windows:

```
cl_agnes_cmode::> iscsi initiator show -vserver test

   Tpgroup       Initiator

Vserver Name     TSIH Name                    ISID              IGroup

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

test    test_san09 85 iqn.1991-05.com.microsoft:rws3.teamcenter.local
```

18. Create an igroup by using the initiator obtained in step 12:

```
cl_agnes_cmode::> igroup create -vserver test -igroup iwin -protocol iscsi -ostype
windows -initiator iqn.1991-05.com.microsoft:rws3.teamcenter.local
```

19. Validate creation:

```
cl_agnes_cmode::> igroup show

Vserver   Igroup       Protocol OS Type  Initiators

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

test      iwin         iscsi    windows  iqn.1991-
   05.com.microsoft:rws3.teamcenter.local
```

20. Map the LUN to the igroup:

```
cl_agnes_cmode::> lun map -vserver test -volume testsan_vol -lun windows.lun1 -
   igroup iwin

cl_agnes_cmode::> lun show -m

Vserver    Volume        Qtree         LUN           Igroup   LUN-ID  Protocol

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

test       testsan_vol   ""            windows.lun1  iwin          0  iscsi
```

21. In Windows, bring up Server Management by right-clicking on Client Desktop and select Manage.

22. In Server Management, navigate to Storage->Disk Management.

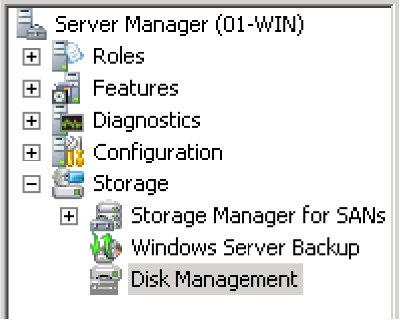

23. Right-click on the box marked Unallocated next to Disk 1 and select New Simple Volume.

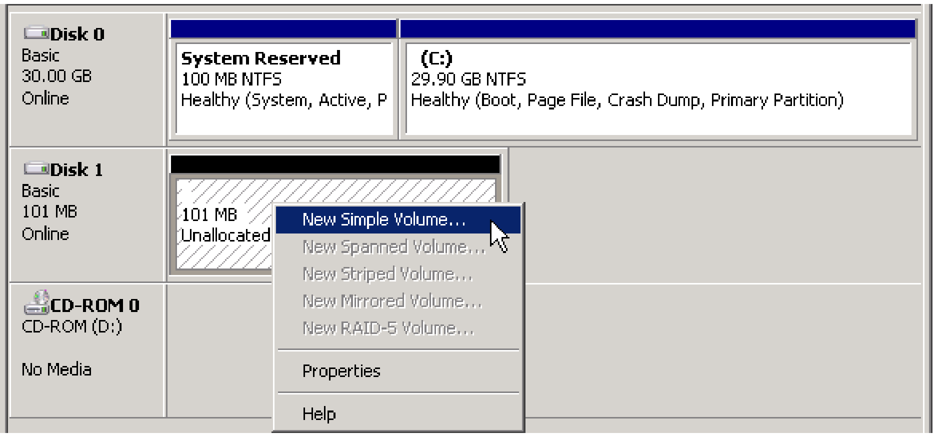

24. Click Next to start the New Simple Volume Wizard. Specify Volume Size, then click Next to accept the defaults.

25. Assign a drive letter and click Next.

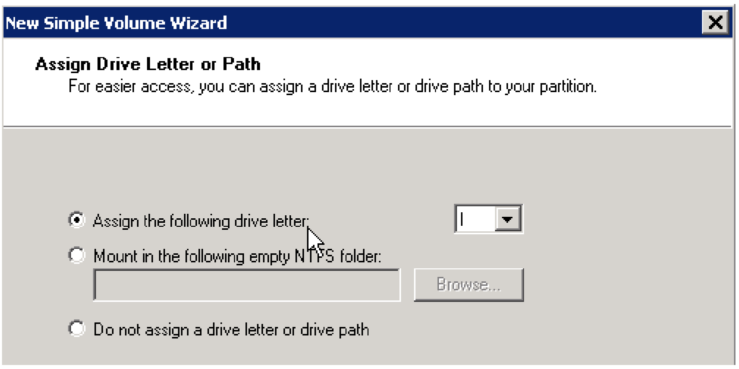

New Simple Volume Wizard

Assign Drive Letter or Path

For easier access, you can assign a drive letter or drive path to your partition.

● Assign the following drive letter: I

○ Mount in the following empty NTFS folder:

○ Do not assign a drive letter or drive path

26. Accept the default formatting options or customize the Volume label. Click Next.

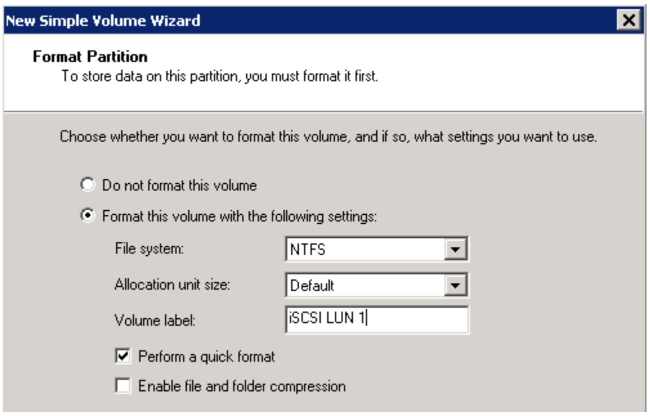

New Simple Volume Wizard

Format Partition

To store data on this partition, you must format it first.

Choose whether you want to format this volume, and if so, what settings you want to use.

○ Do not format this volume

● Format this volume with the following settings:

File system: NTFS

Allocation unit size: Default

Volume label: iSCSI LUN 1

☑ Perform a quick format

□ Enable file and folder compression

27. Click Finish.

15.9 Fibre Channel Access

Configuring the LUN for Fibre Channel (FC) access can differ depending on the operating system. For comprehensive instructions, refer to “FC Host Utilities” on the NetApp Support site (formerly NOW). This example walks through the setup for FC access on Linux. NetApp Data ONTAP operating in Cluster-Mode no longer supports connecting the FC SAN storage directly to the host server. A fiber SAN switch with NPIV support is now required to connect all targets and initiators.

Set Up Storage

1. Set up the storage system to have its FC ports be targets:

```
cl_agnes_cmode::> run -node fas3170c-svl09
Type 'exit' or 'Ctrl-D' to return to the CLI
fas3170c-svl09>
fas3170c-svl09> fcadmin config

                  Local
Adapter Type      State                  Status
---------------------------------------------------
  0c   initiator  CONFIGURED.            online
  0d   initiator  CONFIGURED.            offline

fas3170c-svl09> priv set diag
```

```
Warning: These diagnostic commands are for use by NetApp
         personnel only.
fas3170c-svl09*> fcadmin offline 0c
fas3170c-svl09*> fcadmin config -t target 0c
A reboot is required for the new adapter configuration to take effect.
fas3170c-svl09*>
cl_agnes_cmode::> run -node fas3170c-svl10
Type 'exit' or 'Ctrl-D' to return to the CLI
fas3170c-svl10> fcadmin config

                  Local
Adapter Type      State                  Status
---------------------------------------------------
  0c   initiator  CONFIGURED.            offline
  0d   initiator  CONFIGURED.            offline

fas3170c-svl10> fcadmin config -t target 0c
A reboot is required for the new adapter configuration to take effect.
fas3170c-svl10>
cl_agnes_cmode::> run -node * fcadmin config
2 entries were acted on.

Node: fas3170c-svl09
                  Local
Adapter Type      State                  Status
---------------------------------------------------
  0c   initiator  PENDING (target)       offline
  0d   initiator  CONFIGURED.            offline


Node: fas3170c-svl10


                  Local
Adapter Type      State                  Status
---------------------------------------------------
  0c   initiator  PENDING (target)       offline
  0d   initiator  CONFIGURED.            offline

qntap10g::system> system node reboot -node *
2 entries were acted on.


Connection to 172.17.39.188 closed.
```

Create The WWN and Management LIF

1. On the command line of the storage system, create a management LIF to be used by SnapDrive to access the Vserver and to issue ZAPIs:

```
cl_agnes_cmode::> network interface create -vserver test -lif test_mgmt -role data
-data-protocol none -home-node fas3170c-svl11 -home-port e4b -address 172.31.8.247
-netmask 255.255.255.0 -status-admin up -firewall-policy mgmt
```

2. Unlock the user ID and assign the password `vsadmin` for the Vserver administrator. The password was created when the Vserver was created:

```
cl_agnes_cmode::> security login password -vserver test -username vsadmin

Please enter a new password:
Please enter it again:
```

```
cl_agnes_cmode::> security login unlock -vserver test -username vsadmin


 cl_agnes_cmode::volume> security login show -vserver test
           Authentication Acct
Vserver  UserName  Application Method Role Name  Locked
-------  --------  ------------------ ---------  ------
test    vsadmin   http        password   admin     no
test    vsadmin   ontapi      password   vsadmi    no
test    vsadmin   ssh         password   vsadmin   no
```

3. Create the network interfaces for FC access:

```
cl_agnes_cmode::network> network interface create -vserver test -lif test_fcp09 -
role data -data-protocol fcp -home-node fas3170c-svl09 -home-port 0a

cl_agnes_cmode::network> network interface create -vserver test -lif test_fcp10 -
role data -data-protocol fcp -home-node fas3170c-svl10 -home-port 0a

cl_agnes_cmode::network> net int show -vserver test

 Logical    Status     Network       Currrent    Current Is
Vserver  Interface  Admin/Oper Address/Mask   Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ----
test
test_fcp09 up/down 20:02:00:a0:98:29:02:1e fas3170c-svl09 0a true

          test_fcp10 up/down 20:03:00:a0:98:29:02:1e fas3170c-svl10 0a true
```

4. Create volumes for the LUNs:

```
cl_agnes_cmode::volume> volume create -vserver test -volume testsan_vol -aggregate
aggr_test -size 500MB -state online -space-guarantee none
[Job 7175] Job succeeded: Successful

cl_agnes_cmode::volume> volume show -vserver test
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
test      testsan_vol     aggr_test    online     RW        500MB    474.9MB    5%
```

FC Switch Setup

The setup may vary depending on your FC switch. However, the steps are basically the same.

1. Enable NPIV.
2. Connect all targets (storage nodes) and initiators (server HBAs) to the FC switch.
3. Create an alias for the NetApp cluster nodes that use the WWN created in the previous section.
4. Create a zone for the NetApp cluster node alias and initiators (UNIX or Windows Server HBA: QLogic or Emulex).
5. Save the configuration.

Configure Suse Linux For FC Access

1. Create LUNs for Linux access on the storage system:

```
cl_agnes_cmode::>  lun create -vserver test -volume testsan_vol -lun linux.lun1  -
size 250MB -ostype linux -space-reserve disabled

Created a LUN of size 250m (262144000)

cl_agnes_cmode::> lun show -vserver test
Vserver   Volume     Qtree      LUN        State   Mapped   Type        Size
--------- ---------- ---------- ---------- ------- -------- -------- --------
test      testsan_vol    ""   linux.lun1   online  unmapped linux       250MB
```

2. Get the initiator name for the Linux server:

```
cl_agnes_cmode::>  fcp initiator show
    Logical           Port     Initiator    Initiator
Vserver   Interface         Address  WWNN         WWPN         Igroup
--------- ----------------- -------- ------------ ------------ --------------
qsun277_san
   qsun277_fcp01 40000  20:00:00:00:c9:48:ce:9d 10:00:00:00:c9:48:ce:9d

   qsun277_fcp02 40000 20:00:00:00:c9:48:ce:9d 10:00:00:00:c9:48:ce:9d
```

3. Create an igroup by using the initiator name from step 2:

```
cl_agnes_cmode::> igroup create -vserver test -igroup ilinux2 -protocol fcp -ostype
   linux -initiator
10:00:00:00:c9:48:ce:9d


cl_agnes_cmode::> igroup show -vserver test
Vserver   Igroup       Protocol OS Type  Initiators
--------- ------------ -------- -------- ------------------------------------
test      ilinux2       fcp    linux    10:00:00:00:c9:48:ce:9d
```

4. Map the LUN to the igroup:

```
cl_agnes_cmode::> lun map -vserver test -volume testsan_vol -lun linux.lun1 -igroup
   ilinux2

cl_agnes_cmode::> lun show -vserver test
Vserver   Volume     Qtree      LUN        State   Mapped   Type        Size
--------- ---------- ---------- ---------- ------- -------- -------- --------
test      testsan_vol
                     ""         linux.lun1 online  mapped   linux       250M
```

5. Download and install the latest FC host utilities from the NetApp Support site for the Linux host. Refer to the “Host Utilities Manual” from the NetApp Support site (formerly NOW) for complete instructions on how to install the host utilities and set up FC access. The following next steps are quick instructions on how to configure Linux for FC access. These steps are not comprehensive; refer to the “Host Utilities Manual” for more details.

    a. Verify that the Suse Linux multipath packages are installed on the Linux server:

    ```
    # rpm -q device-mapper
    # rpm -q multipath-tools
    ```

b. Edit multipath.conf using an editor and copy and paste the following stanza into /etc/multipath.conf:

```
# cp /etc/multipath.conf /etc/multipath.conf.orig

(Copy the following stanza into multipath.conf. This for Suse Linux
   Enterprise Server with SP1 and ALUA)

defaults
{
   user_friendly_names no
   max_fds max
   flush_on_last_del yes
}
blacklist
{
   devnode "^hd[a-z]"
   devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
   devnode "^cciss.*"
}
devices
{
   device
   {
   vendor "NETAPP"
   product "LUN"
   getuid_callout "/lib/udev/scsi_id -g -u -d /dev/%n"
   prio "alua"
   features "1 queue_if_no_path"
   hardware_handler "1 alua"
   path_grouping_policy group_by_prio
   path_selector "round-robin 0"
   failback immediate
   rr_weight uniform
   rr_min_io 128
   path_checker tur
   }
}
```

c. Blacklist local drives by getting the WWID of the local drive and adding it to /etc/multipath.conf. For example:

```
# scsi_id -gus /block/sda
SATA_HDS728080PLA380_PFDB32S0R3WHJM
```

Edit `multipath.conf` based on the output of `scsi_id` of the root local drive:

```
blacklist {
        wwid SATA_HDS728080PLA380_PFDB32S0R3WHJM
        devnode "^hd[a-z]"
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^cciss.*"
}
```

d. Enable and start MPxIO for iSCSI:

```
# /etc/init.d/boot.multipath start
# /etc/init.d/multipathd start
```

However, to have the multipath start automatically while booting enter the following commands:

chkconfig --add boot.multipath

chkconfig --add multipathd

chkconfig boot.multipath on

chkconfig multipathd on

e. Verify that multipathd is running and working:

```
# /etc/init.d/multipathd status
multipathd (pid  2314) is running...

To view a list of multipath devices enter:
   multipath -ll
```

f. Reboot to enable changes:

```
shutdown –r now
```

g. When the Linux host is back up, obtain the list of all current sessions using:
`iscsiadm –m session`

To rescan all the sessions:

```
iscsiadm –m session --rescan
```

h. Use `sanlun` to verify the multipath policy and provider:

```
# sanlun lun show -p
                    ONTAP Path: test:/vol/testsan_vol/linux.lun1
                           LUN: 0
                      LUN Size: 250m
                          Mode: C
                   Host Device: 3600a09803246696e433f2d2d636e6b39
              Multipath Policy: round-robin 0
            Multipath Provider: Native
--------- ---------- ------- ------------ --------------------------------
   -------
--------
host      vserver
path      path       /dev/   host         vserver
state     type       node    adapter      LIF
--------- ---------- ------- ------------ --------------------------------
   -------
--------
up        primary    sdm     host10       test_san10
up        secondary  sdk     host9        test_san09
```

i. Verify the multipaths:

```
  [root@ibmx3455-svl01 etc]# multipath -ll
3600a09803246696e433f2d2d636e6b39 dm-6 NETAPP,LUN C-Mode
[size=250M][features=1 queue_if_no_path][hwhandler=1 alua][rw]
\_ round-robin 0 [prio=50][active]
```

```
        \_ 11:0:0:0 sdm 8:192 [active][ready]
       \_ round-robin 0 [prio=10][enabled]
        \_ 8:0:0:0  sdj 8:144 [active][ready]
        \_ 9:0:0:0  sdk 8:160 [active][ready]
        \_ 10:0:0:0 sdl 8:176 [active][ready]
```

6. Reboot to enable changes:

```
shutdown –r now
```

7. Use `sanlun` to verify the multipath policy and provider:

```
sanlun lun show –p
```

8. Verify the multipaths:

```
  [root@ibmx3455-svl01 etc]# multipath -ll
3600a09803246696e433f2d2d636e6b39 dm-6 NETAPP,LUN C-Mode
[size=250M][features=1 queue_if_no_path][hwhandler=1 alua][rw]
\_ round-robin 0 [prio=50][active]
 \_ 11:0:0:0 sdm 8:192 [active][ready]
\_ round-robin 0 [prio=10][enabled]
 \_ 8:0:0:0  sdj 8:144 [active][ready]
 \_ 9:0:0:0  sdk 8:160 [active][ready]
 \_ 10:0:0:0 sdl 8:176 [active][ready]
```

9. Create a file system on the device:

```
# mkfs -t ext3 /dev/dm-6

mke2fs 1.39 (29-May-2006)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
64000 inodes, 256000 blocks
12800 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=67371008
32 block groups
8192 blocks per group, 8192 fragments per group
2000 inodes per group
Superblock backups stored on blocks:
        8193, 24577, 40961, 57345, 73729, 204801, 221185

Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 22 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

10. Mount the file system and validate the size:

```
#  mkdir /LUN
 # mount /dev/dm-6 /LUN

# df -h /LUN
Filesystem            Size  Used Avail Use% Mounted on
/dev/dm-6             243M  6.1M  224M   3% /LUN
```

15.10 Mount DB And File Vaults

After the NAS volume or SAN has been mounted, install the database and Teamcenter to utilize the newly created NAS or SAN volume.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

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

www.netapp.com