Technical Report

Best Practices for Deploying Siemens PLM Software Teamcenter 2005 SR1 and 2007 MP3 (2007.1.3) by Using NetApp Storage Systems

Bikash R. Choudhury, NetApp; Subra Mohan, NetApp; Bill Halpin, Siemens PLM Software
August 2009 | TR-3658

EXECUTIVE SUMMARY

NetApp and Siemens PLM Software conducted a series of performance and scalability benchmarks as a joint effort to illustrate a reference architecture (two-tier) for Teamcenter with NetApp® storage. This report documents performance benchmarks across SAN and NAS with different workloads (low [<100 users] and medium [<1,000] users), along with the performance impact of additional cache memory in the storage controller on read workloads. The intent of this report is to help organizations choose the optimal storage deployment option, based on performance and total cost of ownership. This document discusses the impact of using NetApp storage for Teamcenter files (tcengvault) and Oracle® Database. The scalability test was run on Teamcenter 2005 SR1. Due to a bug in the Teamcenter 2005 SR1 `backup_mode` utility further tests were run for backup and restore, PAM II, and deduplication using the Teamcenter 2007 MP3 release. However, the scalability and performance testing with and without PAM II apply to both the 2005 SR1 and 2007 MP3 releases. Note that the `backup_mode` utility can be used only with 2007 MP3.

TABLE OF CONTENTS

INTRODUCTION

Siemens PLM Software product lifecycle management (PLM) solutions are a unified approach to extended enterprise collaboration that enables all participants in your product lifecycle to work in concert as you bring products to market and support your customer base. Teamcenter is an industry-driven, customer-proven, Web-centric collaborative product lifecycle management system for the engineering enterprise. Teamcenter provides distributed engineering and 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 are significant design and operational elements of a Teamcenter environment. NetApp storage solutions include NetApp Snapshot™ and SnapRestore®, which provide a reliable, high-performance mechanism to create an almost instantaneous backup of the database and data volumes. These capabilities offer an attractive solution set for Teamcenter environments and should be key design criteria for future deployments.

NetApp storage systems integrate seamlessly with Teamcenter into a PLM collaborative engineering environment. NetApp storage systems 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, providing high-performance access to a single copy of the data, which is shared across all types of clients.

NetApp, in a joint effort with Siemens PLM Software, has successfully completed component-level and system-level benchmarks to assess performance and scalability when Teamcenter is integrated with NetApp storage systems.

First, NetApp was able to integrate a NetApp storage system into a production-capacity Teamcenter application environment in a two-tier architecture model in NetApp development labs. During these component-level benchmarks, multiple processes were executed, mimicking hundreds of users performing file upload, download, and deletion in parallel. As configured, 10 processes each of upload, download, and deletion were executed, for 30 processes total. Each process simulates the workload of tens of users, representing an average workload of several hundred users over the course of a day. These user processes work with Teamcenter volumes configured on the NetApp storage assigned to their group and are all nonprivileged processes. To benchmark this, the environment consisted of a database server, a dedicated file management system/Teamcenter file services (FMS/TCFS) server, an API server (load canon), and a NetApp storage system that hosted the volumes.

Separately, Siemens PLM software conducted system-level scalability benchmarks at its Cypress, California, development labs. These benchmarks simulate up to 2,000 users accessing the Teamcenter four-tier infrastructure, which was integrated with the NetApp FAS3020c storage system. Several configurations were tested, including Teamcenter volumes (file storage) on the FAS3020c, as well as locating Oracle data files on the FAS3020c with SAN, NFS, and CIFS access. As configured, these benchmarks simulate more than two dozen user-level transactions with users in three separate roles. The usage profile mimics a traditional work-in-process access model, wherein 50% of the transactions result in updates to the database. This was important to maintain representative load on the Oracle server, because fast database updates to the underlying file system are the key to acceptable Oracle performance. Poor I/O response times from the SAN, NFS, or CIFS storage device result in unacceptable end-user response times.

1.1 PERFORMANCE

Performance issues pertaining to disk bottlenecks and high maintenance costs for direct-attached disks are greatly alleviated through flexible volumes, which spread across a large number of spindles. Large PDM files, which can be a bottleneck on conventional systems, are easily accommodated by using NetApp hot-swappable, scalable storage. The powerful storage system with large capacity holds large simulation and testing files, while allowing the reallocation of space on the fly between each iteration.

Scalability benchmarks with Oracle data files on NetApp storage showed that server response times were very similar to local high-throughput, low-latency direct-attached disk arrays. It was noted that the processing on the database server in these SAN and NFS benchmarks showed a slight increase in CPU utilization (at 1,000 users) with Oracle10*g™* R2 data files hosted on the NetApp storage. However, CPU utilization was so low that it is difficult to conclude that running with Oracle data files on SAN or NFS caused the CPU increase.

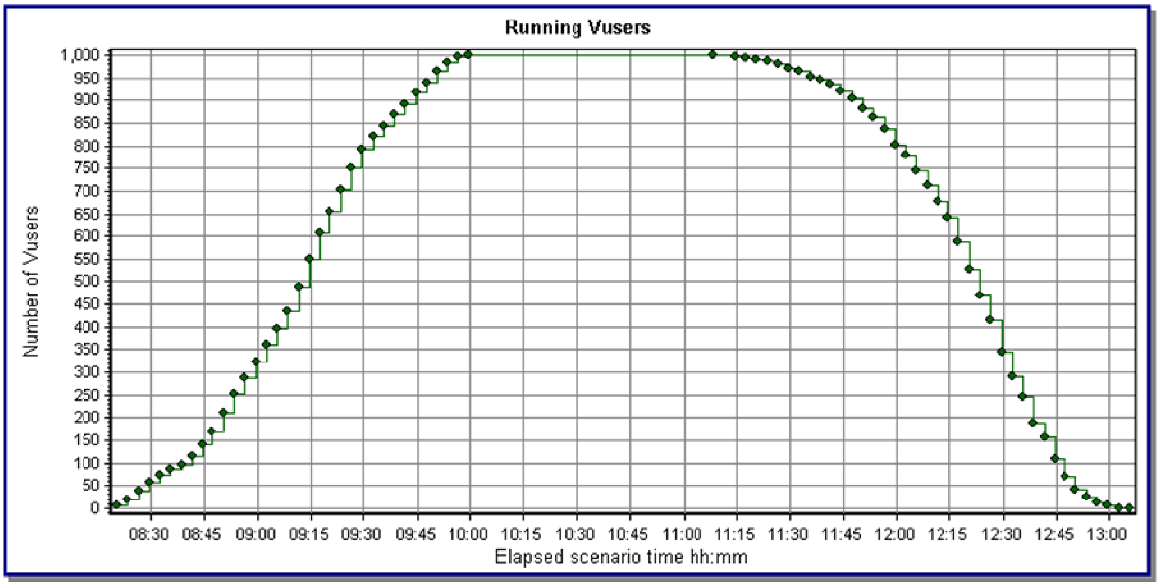

Figure 1) User login profile for 1,000 users.

Figures 2 and 3 illustrate CPU utilization for a server with Oracle data files configured on NetApp storage via NFS and SAN.

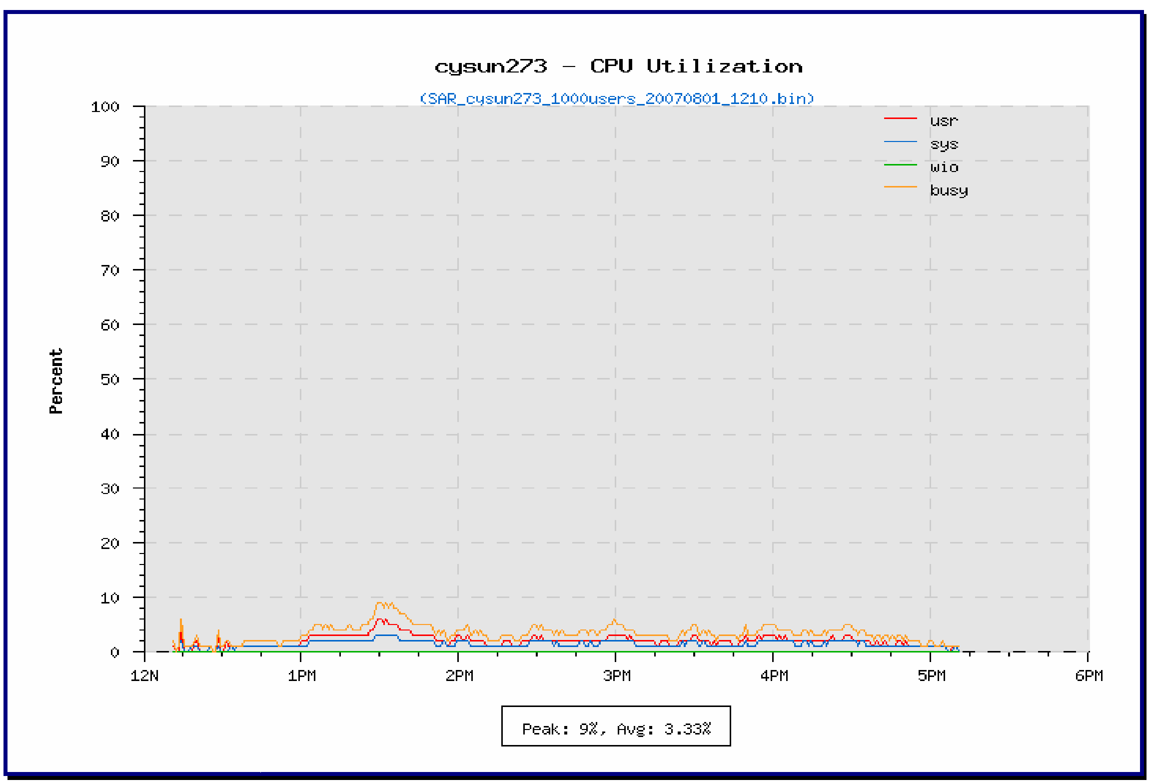

Figure 2) Server utilization with Oracle data files on SAN.

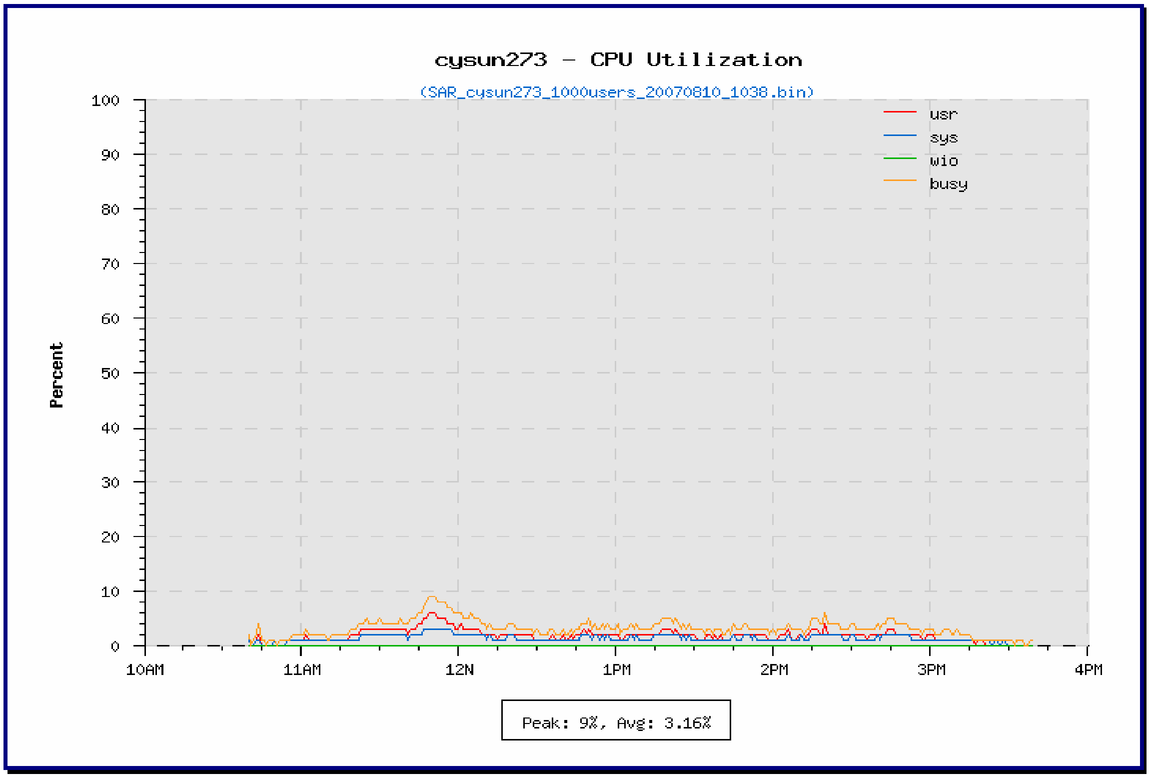

Figure 3) Server CPU utilization with Oracle data files on NFS.

1.2 SCALABILITY

Any application that has been integrated with Teamcenter (for example, NX or CATIA) can create or access data directly in the Teamcenter vault. 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.

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 Teamcenter, allowing recovery of the information at any stage in its development.

Even as the size and the number of the revised data continue to grow along with the number of users, the NetApp storage solution is able to scale to meet user requirements with dynamic online disk expansion that increases storage capacity without rebooting.

1.3 AVAILABILITY

In clustered failover testing, NetApp storage systems were able to successfully fail over in less than 90 seconds with no material impact to the Teamcenter 2005 SR1 or 2007 MP3 application and the Oracle10*g* database. Planned downtime is also minimized with online disk expansion. Plus, integrated NetApp RAID-DP® implementation protects against disk failures to avoid disruption of service to users.

1.4 RECOVERABILITY

Backing up the master database in hot backup mode was made possible by using NetApp Snapshot copies, which copy only incremental data and make backup and recovery fast and seamless. Using this approach, it is possible to recover a 50GB database in a matter of minutes and to keep several Snapshot copies (and therefore multiple recovery points) online.

To implement the new solution, the master environment was staged in a FlexVol® volume inside an aggregate with a large pool of physical disks. A master Snapshot copy of the FlexVol volume supports the creation of copies or clones. Upon request, the Snapshot copy of the master environment can be cloned instantaneously by using FlexClone®. The cloned volumes that are created only consume space incrementally as changes are made to the original image. This saves a significant amount of storage space.

If any files are accidentally deleted in a FlexClone volume, the user uses single-file SnapRestore to retrieve the file without wasting a lot of time. This feature can eliminate 30% of helpdesk requests per user per month and thus helps improve time to market.

In tests at the NetApp development facility that simulated Teamcenter 2005 SR1 failure due to database corruption, the NetApp storage capabilities facilitated recovery in a very short time: It took a few seconds to recover the database using SnapRestore, a couple of minutes to apply the transaction logs to this clean database, and within minutes the Teamcenter 2005 SR1 application was recovered to full throughput capacity.

1.5 MANAGEABILITY

Having the entire resource tier in a central location accessible from anywhere in the world 24x7 proves to be cost effective, scalable, and easier to manage than a distributed environment.

- Using the FMS server and the database over NFS has the following advantages:
  - It reduces the complexity of deploying and managing storage.
  - Using NFS speeds up deployment and simplifies support for a large number of dependent environments.
  - Ethernet-based NAS is much easier to understand and implement,
  - NetApp NAS, which has already been proven in various Oracle Database deployments, met the project's performance requirements of volume file access times under 15 milliseconds. It also provides multihost file system access, file-level data sharing, and the necessary level of security.
- Snapshot technology drastically reduces the time of database and data set backup in the Teamcenter environment, from hours to a matter of seconds. File recovery is fast, and data is accessible from backup if primary storage is out.
- 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 in which NFS is exclusively adopted. Native support of NFS and CIFS within the NetApp kernel allows Windows, UNIX, and Linux® clients to share the same copy of data.

1.6 LOWER TCO

With NetApp storage, less time is spent tuning, managing, and expanding storage. Less administration overhead and higher productivity can result in lower costs and increased profitability.

2 REFERENCE ARCHITECTURE OVERVIEW

2.1 OBJECTIVES

The objective was to define and test the integration of NetApp storage in a Teamcenter production-capacity environment. These tests are designed to provide organizations that have either implemented Teamcenter 2005 SR1 or 2007 MP3 or later or are considering migrating to these releases with options regarding data storage and overall application architecture design. The results of these tests could also be applied to organizations that are currently employing the Teamcenter 2005 SR1 or 2007 MP3 or later application.

The objectives of the tests were achieved by developing and testing the configurations of the storage needs of a Teamcenter deployment of 2005 SR1 or 2007 MP3 in a controlled lab environment. This production-capacity lab environment was tested for concurrent user populations of from a few hundred up to 2,000 users. This range of concurrent users was selected to be representative of a wide range of Teamcenter deployments, with a maximum total named user population of approximately 20,000.

To achieve an objective overview of the NetApp storage solution in the Teamcenter application architecture, a representative NAS configuration over NFS was created for the network storage configuration. By using this as a point of reference for each configuration, it was possible to describe the application architecture on the NetApp storage.

2.2 AUDIENCE FOR THIS REFERENCE ARCHITECTURE

The intended audience for this reference architecture document is technical architects from organizations that fall into one of the following groups:

- Organizations that plan to implement Teamcenter and need to understand their network storage for different aspects of the Teamcenter architecture
- Organizations that are migrating or are planning to migrate to Teamcenter 2005 SR1 or 2007 MP3 or later
- Organizations that have implemented or are currently implementing Teamcenter 2005 SR1 or 2007 MP3 or later and who need to understand how to manage database files or large numbers of FMS files
- Organizations that are currently employing Teamcenter and need to understand their network storage options for different aspects of the Teamcenter architecture.

Organizations in these groups can use the results of this study to understand how to implement NetApp storage technologies in their current Teamcenter application architectures.

3 KEY ELEMENTS OF THE REFERENCE ARCHITECTURE

The reference architecture component-level tests were developed and executed based on the Teamcenter FMS functionality. FMS is the Teamcenter subsystem component responsible for delivering, accepting, and managing Teamcenter files. These areas were derived from the input of a number of groups that include Siemens PLM Software subject matter specialists and NetApp subject matter specialists.

The test operates in two different modes: retrieve (download) and create (upload).

- **Retrieve files (download).** This mode reads an input file that contains a list of Teamcenter data set files to retrieve. For each file in the list, it exports the file using the ITK call `IMF_fmsfile_export`.

  The files exported from Teamcenter engineering were written to a local directory on the system used as a load canon. As configured, the program runs until 60 minutes have passed; or until 10,000 files have been exported; or until the input file has been processed, whichever comes first. With 10 processes running concurrently in retrieve mode, up to 100,000 files are exported in 60 minutes.
- **Create files (upload).** This mode reads an input file that contains a list of files in the file system of the load cannon. Each file is imported into Teamcenter using the ITK call `IMF_fmsfile_import`.

  The input file for this test can contain either just file names or full paths to the files. As configured, the program runs until 60 minutes have passed; or until 2,000 files have been imported; or until the input file has been processed, whichever comes first. After each file has been imported, the program sleeps for two seconds by default. With 10 processes running in create mode, up to 20,000 files are created in 60 minutes.

The system-level scalability benchmarks were developed based on the Teamcenter thin client user interface and were used to help validate the reference architecture for typical end-user operations, as well as to assess the end-user performance impact of locating Oracle data files on the NetApp storage system. The thin client is the browser-based Teamcenter user interface.

This benchmark uses a capture and replay tool to simulate thin client users. HTTP messages from browser to Web application server are captured during recording sessions of end-user transactions. The resulting transaction “scripts” are parameterized with variable data and user names to simulate multiple users. The parameterized scripts are then executed in parallel on several load generation servers to generate a simulated user load. Each simulated user accesses or creates unique data over the course of the five- to six-hour run.

During the execution of the scalability benchmarks, the following were measured from the FMS and Oracle Database server and the NetApp storage:

- The login activities during peak periods for the 1,000-user thin client benchmark
- The CPU and memory utilization on the FMS and the database servers were monitored as the load kept growing
- CPU utilization and NFS operations and latency on the NetApp storage.

This reference architecture document is not intended to act as a sizing guide for hardware to be deployed in a Teamcenter application environment, nor is it intended to be an exhaustive review of the Teamcenter functionality. Additionally, the tests did not develop extensive tuning and performance guidelines for the applications in the architecture, although basic guidelines supplied by Siemens PLM Software were followed.

4 TEAMCENTER 2005 SR1 AND 2007 MP3 OVERVIEW: TWO-TIER ARCHITECTURE MODEL

The reference architecture involving Teamcenter and NetApp storage can be deployed in two-tier and four-tier configurations. The four-tier configurations include client, Web or Web application, enterprise or business logic, and resource tiers. In two-tier configurations, the rich client communicates directly with the enterprise tier; there is no Web tier. The resource tier includes the Oracle or SQL Server® database server, database files, the FMS file server, and the NetApp storage that provides storage volumes at the back end.

The client tier can include Teamcenter engineering rich clients, which is a Java™ application, or thin clients, which are browser based. In two-tier deployments, the Teamcenter business logic server process runs on the client workstation. Browser-based thin clients must be deployed in the four-tier configuration.

This version of Teamcenter introduces the Teamcenter FMS to manage uploading and downloading file data between clients and volumes in both two-tier and four-tier architecture deployments.

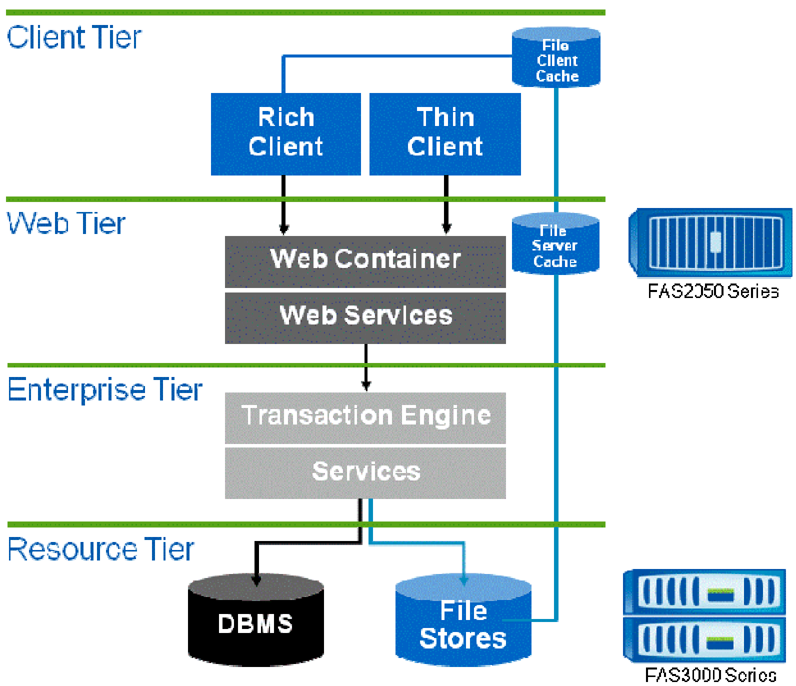

Figure 4) Four-tier Teamcenter architecture.

For improved performance at remote sites over low-speed, high-latency WAN networks, remote file server caches can be deployed on optional NetApp storage systems. Typically, one larger NetApp storage system serves as the central file store, while multiple smaller NetApp storage systems are deployed at each remote site. File client caches are configured on each end user's workstation disk.

This provides a file storage, caching, distribution, and access system with the following benefits:

- **Data distribution.** Administrators can distribute copies of data closer to end users by deploying FMS server caches (FSCs) at remote locations. FSCs can be distributed worldwide, while retaining FMS volume data in central storage.
- **No single point of failure.** FMS provides the capability to administer a fully redundant configuration of configuration servers, volume servers, and cache servers. FMS routing algorithms automatically search for an alternate path in the case of a connection failure.
- **Master configuration server.** FMS provides the capability to administer the FMS deployment configuration with a single master configuration file. FMS automatically distributes the configuration file to all FMS client and server processes.
- **Managed caches.** FMS client and server caches are self-purging. The least recently accessed data is purged first.
- **Secure server caches.** FMS servers do not permit direct access to cached file data. FMS permits file data access when the requestor presents a valid security ticket.
- **Secure volume servers.** FMS servers do not permit direct access to volume file data. FMS permits file data access only when the requestor presents a valid security ticket.
- **Private user caches.** FMS automatically caches data that is downloaded or uploaded by visualization clients in a private user cache, providing fast access to recently accessed files. The user cache automatically purges data to fit within a maximum size.
- **Streamed data delivery.** FMS streams data from volumes down to clients through any number of cache servers. Data becomes available to the user as soon as the first bits stream in, through any number of cache servers as needed. FMS also streams data from the client all the way to the volume on upload.
- **Segment file cache and delivery.** FMS allows applications to transfer only specific parts of a file, improving the overall transfer time and conserving network bandwidth.
- **Secure communication.** FMS supports HTTPS communication between FMS servers and clients and the use of HTTP or HTTPS proxy servers.

TCFS is a legacy file management system, previously called IMANFS, that provides a variety of volume-related services and operates in conjunction with the new FMS.

4.1 TWO-TIER MODEL

In the two-tier model, the Teamcenter engineering rich client (which includes the local server) and the optional applications are integrated with the rich client on the client workstation The database server and file server are typically installed on two separate computers. All of the Oracle data files and logs and the FMS files reside on the NetApp storage volumes mounted via SAN or over NFS.

The Teamcenter engineering FMS and TCFS manage the rich client access to volumes on the NetApp storage:

- The FMS server cache (FSC) process and TCFS run on the server hosting the volume.
- The FMS client cache (FCC) process runs on the rich client host.

Three configurations of Oracle data file locations were evaluated with Teamcenter 2005 SR1 1,000 thin client user benchmarks. This series used the standard automated performance analysis (APA) Teamcenter scalability benchmarks, database, and usage profiles, which are used throughout the Teamcenter development cycle to assess scalability and server utilization. The sole difference between the Oracle configurations was the location of the Oracle tablespace and redo files. Several benchmarks were executed in this experiment and are posted to the APA dashboard.

Oracle data and redo files were located on:

- **Baseline.** High-throughput local disk array
- **SAN.** Fibre Channel (FC) connection to NetApp FAS3020 UFS LUN
- **NFS.** Gigabit Ethernet (NFS) connection to NetApp FAS3020 NFS export

4.2 ADVANTAGES OF NFS

The advantages of using NetApp storage over NFS include.

- **Controlling growth.** In spite of planning the growth of database and FMS files, it becomes quite a challenge to add physical space to local disk arrays without causing downtime. NetApp storage can scale to a very high disk capacity by hot adding disks or shelves on the fly to the existing setup without having to take the storage system down and causing disruptions to the users. The storage can also do thin provisioning of space on demand for users.
- **File I/O utilization.** The disks on the NetApp storage were 144GB Fibre Channel with 2GB back-plane speed. With the 15k RPM drives, the seek time is further reduced. The volumes that held the database and the FMS files were flexible in nature, so that they could be grown and shrunk as needed. These flexible volumes were created on top of large aggregates that had the pool of physical drives. Disks were added as needed to the aggregates as the data size kept growing. The flexible volumes and aggregate architecture provide more read and write concurrency. One single controller was able to handle the different types of workloads (Oracle data files, Oracle logs, FMS files, and so on) efficiently and with less latency.
- **RAID configuration.** The NetApp storage uses RAID-DP for data resiliency. This implementation provides double-parity RAID protection against data loss with negligible performance overhead and no cost penalty compared to single-parity RAID. RAID-DP is a standard feature and prevents data loss in the event of a second drive failure without excessive redundancy costs. This provides a high level of data protection and availability.
- **Backups.** A significant benefit of NetApp storage is the ability to back up data via Snapshot copies in a matter of seconds, minimizing downtime. During the backup, data may be locked down by the backup process and be unavailable to the application. Traditional methods could take hours, during which time the Teamcenter environment is unavailable. This is not acceptable for 24x7 operations. With Snapshot technology, the data is quiesced, backed up, and available again in seconds or minutes. Then the logically consistent Snapshot copy can be backed up to an offline medium while the data continues to be available to users.

4.3 OTHER DEPLOYMENT OPTIONS

The Oracle data and the redo logs can also reside on LUNs while the Teamcenter volumes still reside on NFS volumes; this is another alternative to having both the Teamcenter database and files over NFS. Typically the Teamcenter database constitutes a smaller portion (100–200 GB) of the Teamcenter setup compared to the Teamcenter files (in TBs). With the Oracle components on the LUN, the read and write performance proves to be slightly better than over NFS. Having the Teamcenter database over FCP or iSCSI LUNs on NetApp storage is beneficial for the following reasons.

- Scalability. Data protection with RAID-DP is still possible with LUNs residing on NetApp storage.
- Backup and recovery and archiving can be done for the Teamcenter database residing on LUNs by using NetApp Snapshot technology.
- Microsoft® MPIO setup can help to optimize the I/O traffic to the iSCSI LUNs.
- iSCSI is a natural fit for implementations with Oracle Database upgrading from DAS.
- If it is decided to change the architecture from iSCSI to FCP, the iSCSI LUN can be converted to FCP on the fly without migrating the data.
- Multiwire scaling is straightforward with SAN tools such as VxVM.

FCP implementation may prove to be a challenge for the following reasons:

- High cost of the implementation
- The amount of complexity to manage the LUNs compared to the database residing on NFS volumes

5 REFERENCE ARCHITECTURE CONFIGURATIONS

This section is an overview of the application architecture that was developed for the test. Before the start of each pair of configuration tests, the lab environment was inspected to make sure that the number and size of the database and other files in the NAS environment were appropriately set.

5.1 COMPONENT-LEVEL TEST

For each configuration described in this document, a standard set of tests was executed. These tests consist of the following items.

Table 1) Component-level test description.

Item	Description
Business process	Multiple processes are executed, mimicking thousands of user file access requests, and they are uploading, downloading, and deleting in parallel.
Users	Simulated users execute file operations continuously for a period of 60 minutes. The test runs for 1,000–5,000 users.
Time scale	The results of the tests were gathered from a 60-minute period.

5.1.1 Component-Level Base Configuration

The tests were designed to gather basic metrics regarding the characteristics of a Teamcenter FMS server. Figure 5 shows the basic configuration of the Teamcenter on the NetApp storage.

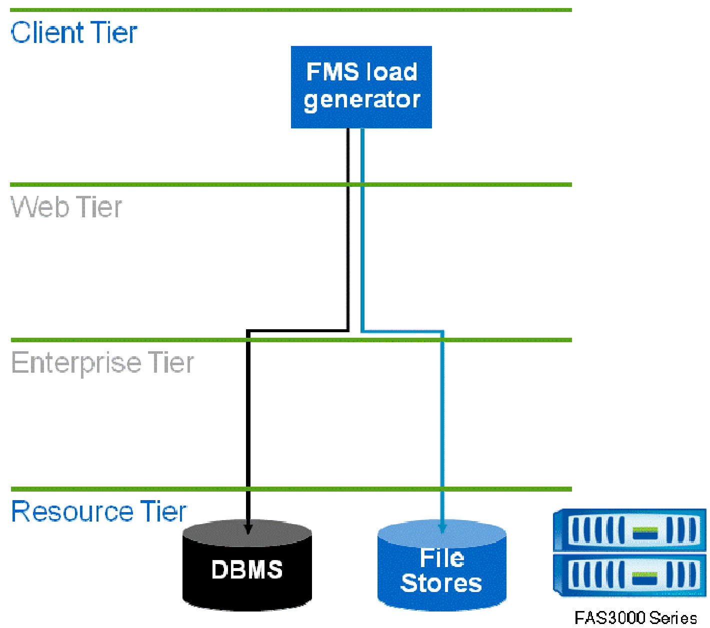

Figure 5) Two-tier architecture with NetApp storage.

5.1.2 Component-Level Objective

The objective of this configuration was to establish FMS response time characteristics for the implementation of Teamcenter file storage on the NetApp storage production environment.

5.1.3 Component-Level Approach

In the network storage configuration, the FMS data files were mounted on a flexible NFS volume for the NFS test and on separate LUNs for the FCP test residing on a 16-disk NFS aggregate on the NetApp storage. The FMS load generator was set to simulate 140,000 read, write, and delete operations in a 60-minute period. These volumes were created on the storage to optimize performance based on NetApp best practices of at least 16 disks per aggregate.

5.1.4 Component-Level Results

The results of these tests show that Teamcenter FMS accessing the NetApp storage architecture over NFS experienced almost the same response time compared to local disks in the overall business process time for hundreds to thousands of users.

5.2 SYSTEM-LEVEL TEST

For each configuration listed in this document, a standard set of tests was executed. These tests consist of the following items.

Table 2) System-level test description.

Item	Description
Business process	Multiple processes are executed, mimicking thousands of thin client users executing more than two dozen typical Teamcenter transactions in parallel.
Users	Simulated users execute Teamcenter operations continuously for a period of over five hours. The test runs for 200–5,000 users.
Time scale	The results of the tests were gathered from a six-hour period.

5.2.1 System-Level Base Configuration

The tests were designed to gather basic metrics regarding the characteristics of a Teamcenter FMS server. Figure 6 shows the basic configuration of the Teamcenter on the NetApp storage.

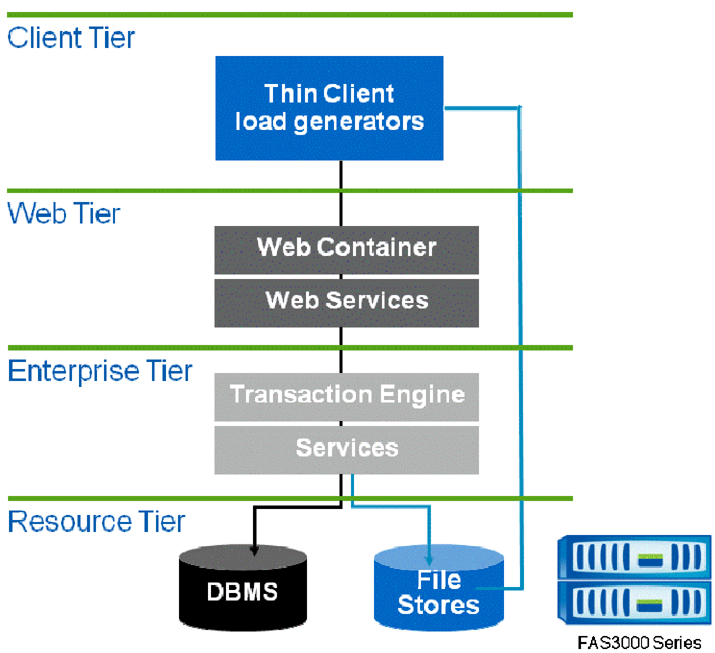

Figure 6) Two-tier architecture with NetApp storage.

5.2.2 System-Level Objective

The objective of this configuration was to validate server response times to establish a blueprint for the implementation of Teamcenter four-tier architecture with Oracle data files and FMS volumes on the NetApp storage production environment. For this test only, the configuration of the NetApp storage was captured, because the installation of the Teamcenter application and the NetApp storage are described in separate documentation provided by each vendor (see the appendix for the location of these documents). Based on this configuration, metrics were captured that describe how the lab environment performed in a NAS environment over NFS.

5.2.3 System-Level Approach

In the network storage configuration, the Oracle data files were mounted on a flexible volume residing on a 16-disk NFS aggregate on the NetApp storage. The Oracle Database data files and the redo logs were stored on flexible volumes in the same 16-disk NFS aggregate. These volumes were created on the storage to optimize performance based on NetApp best practices of at least 16 disks per aggregate.

5.2.4 System-Level Results

The results of these tests show that Teamcenter on NetApp storage architecture over NFS experienced almost the same response time compared to SAN and local disk arrays in the overall business process time for 1,000 users.

Figure 7 illustrates the consistent NFS read/write response time from the NetApp storage over NFS compared to SAN and local disk array (baseline).

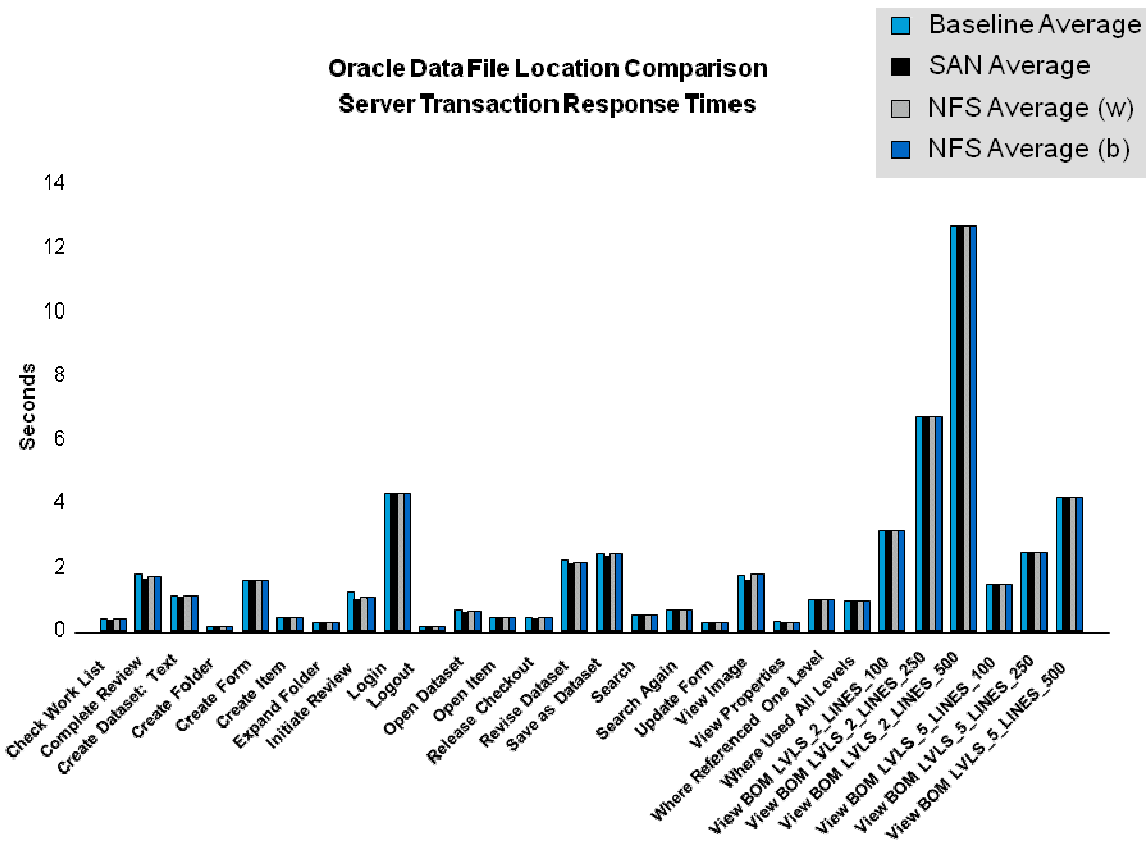

Figure 7) SAN-NFS transaction times comparison.

6 TEAMCENTER 2007 MP3 WITH PAM II OVER NFS

The Teamcenter performance benchmarking effort on NetApp storage was further extended to include quantifying the performance improvements resulting from the use of second-generation Performance Accelerator Modules (PAM IIs); the PAM II is an extended read cache that is available in a PCI-e form-factor within the storage controller. PAM II is an upgrade over the first-generation PAM (DRAM-based) with each module containing either 256GB or 512GB of flash memory. The PAM II is an intelligent read cache that gives you a new way to optimize performance for file services and other workloads that are random read intensive. You can configure up to four 512GB modules in a FAS 6080 as a unified 2TB read cache in the

PCI Express slots of your storage controller running Data ONTAP® 7.3.2. The increased read cache delivers reduced latency (due to more reads being served from memory as opposed to disk) and improved I/O throughput without adding more fast-spinning disk drives. You can tune the PAM II to match your specific workload by using software settings to choose from three modes of operation. With PAM II all three modes—Default/Normal, Metadata, or low-priority— may be enabled. The new-generation card is intelligent enough to determine the appropriate mode with respect to the workload.

- Normal data blocks (Default) mode, which caches both user data that is being read randomly and metadata
- Metadata mode, which caches only metadata
- Low-priority blocks mode, which caches write data blocks and long sequential read blocks in addition to the data that would have been cached by the default mode

6.1 FMSLOAD BENCHMARK TESTING WITH PAM II

The `fmsload` program was a performance benchmarking tool for FMS/TCFS. It was designed to generate FMS/TCFS traffic while minimizing SQL traffic. During the benchmarks, multiple processes are executed, mimicking hundreds or thousands of users uploading, downloading, and deleting files in parallel. As configured, 10 processes execute uploads, downloads, and deletions for a total of 30 processes; each process logs in as a different user. They all work on the volumes that are assigned to their group. These processes are all non-`infodba` processes. To benchmark this process with NetApp storage and PAM II on the back end, the environment consisted of a database server, a dedicated FMS/TCFS server, an API server (load cannon), and a NAS server, which hosts the volumes. We compared performance on the following configurations, while executing the `fmsload` program:

- Config A - FAS3160 configured with dual FC disks without PAM II
- Config B - FAS3160 configured with a single SATA disk shelf with PAM II (256GB)

The fmsload program operates in three different modes: retrieve (download) and create and save (upload). Save mode creates files for a deletion test (`save_files`). Deletion is performed by TCFS.

Retrieve_files (download). This mode reads an input file that contains string representations of tags of file objects. For each tag, it exports the file by using the ITK call `IMF_fmsfile_export`.

The files exported from TC Eng are written to a local directory on the system that is being used as a load cannon. As configured, the program runs until 60 minutes have passed, or until 10,000 files have been exported, or until the input file has been processed, whichever comes first.

Create_files. This mode reads an input file that contains a list of files in the file system of the load cannon. Each file is imported into TC Eng by using `IMF_fmsfile_import`.

The input file for this test can either contain just file names or contain full paths to the files. As configured, the program runs until 60 minutes have passed, or until 2,000 files have been imported or until the input file has been processed, whichever comes first. After each file is imported, the program sleeps for two seconds by default.

Save_files. This mode is the same as the `create_files` mode, except that it saves the file object and writes the tags of these files to a file in the file system of the load cannon. This mode is not for testing; it is only for creating files that will be deleted during a test.

This load cannon program was run on NetApp storage to check the improvement in the response time for reads at the application level with and without PAM II cards.

6.2 RESULTS OF FMSLOAD BENCHMARK TESTS WITH PAM II

- Figure 9 indicates that config B demonstrated a 55% performance improvement over config A for exports (reads) at the application level.
- At the storage sub-system level, we observed that config B delivered twice the I/O throughput of config A.

- Config B demonstrated a 99% reduction in NFS latency. This was made possible because the entire data set was able to fit into the 256GB cache, thus significantly reducing the disk I/Os. Up to 60% of the 256GB cache was utilized during this test.

While testing with PAM II, both normal data blocks mode and low-priority blocks mode were turned on, since the workload contains both random and sequential reads.

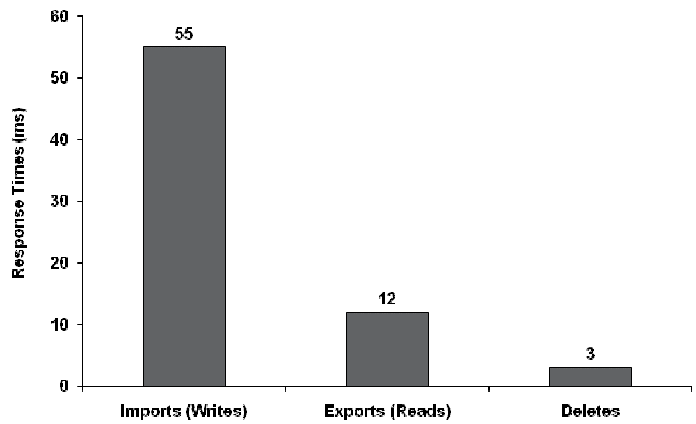

Figure 8) Response times for FMSload Benchmark Tests without PAM II

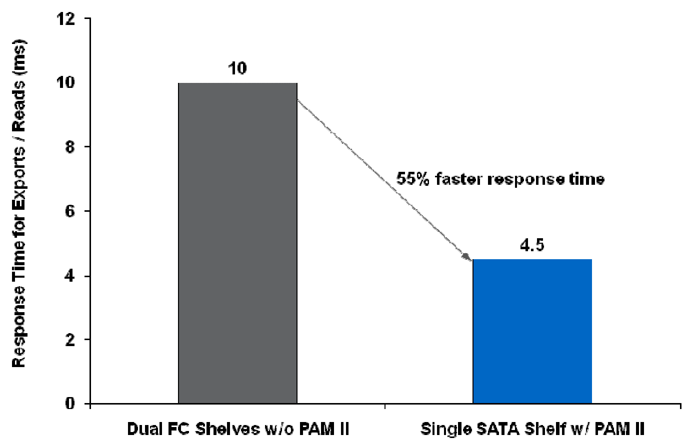

Figure 9) Response time for exports (reads) with PAM II.

6.3 AUTOMATIC PERFORMANCE ANALYSIS (APA) BENCHMARK TESTS WITH PAM II

The APA benchmark test was performed in the Siemens PLM's APA lab on a 70GB Oracle Database over NFSv3. We utilized the same configurations that was used for the `fmsload` test (Config A: FAS3160 with dual FC disk shelves without PAM II; Config B: FAS3160 with single SATA disk shelf with PAM II), at their lab to run the scalability test for 500 users. The results, as illustrated in Figures 10 and 11 showed,

- 96% reduction of disk I/Os
- 97% reduction in NFS latency

This implies that most of the I/O requests were served from the PAM II cache. Up to 60% of the 256GB cache was utilized during this test. The SQL queries shown in Figure 12 that do the metadata operations such as adding a list of 50000 users who are e-mail recipients of workflow information and queries to find out the groups/projects that members are part of could notice up to an 80% improvement in the response time at the application level. The actual percentage of improvement in response time for a SQL query depends on the percentage of overall time taken up by the storage I/O to perform the SQL query. For many of the SQL queries, we don't see improvements in overall response times at the application level that result from PAM II, because the storage I/O takes up a very small percentage of the overall SQL query time.

Figure 13, which shows the read times at the storage subsystem level for the Oracle tablespaces, better illustrates the performance impact of PAM II. For example, we observe that

- The response times for the SYSAUX tablespace read was reduced by 94% with PAM II, because the entire Oracle tablespace fits in the cache.
- The I/O throughput almost doubles because very few I/Os go to the disk.

These results demonstrate that, for Teamcenter PLM environments, you can improve your response times for I/O requests from both rich and thin clients with PAM II while reducing your overall costs by deploying half the number of the less expensive and denser SATA disk shelves. You can reduce data-center footprint and power and cooling costs with SATA disk shelves.

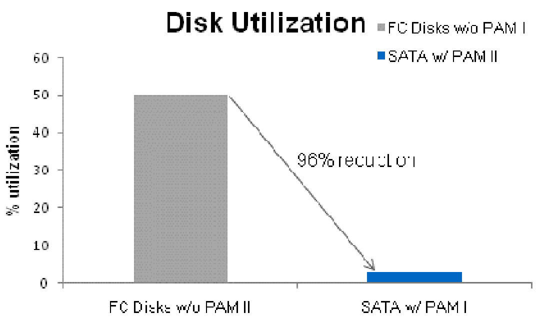

Figure 10) Reduction of disk utilization with PAM II on NetApp storage for Scalability Benchmark.

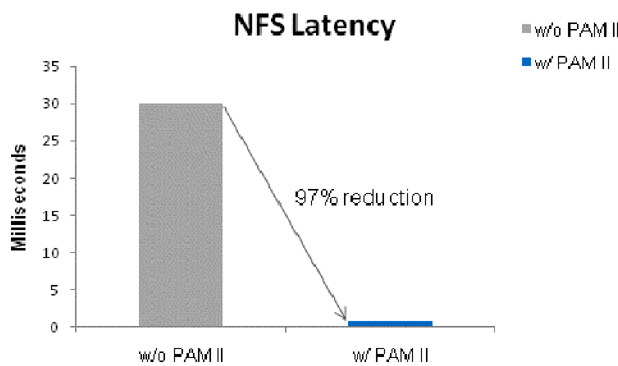

Figure 11) Reduced NFS latency with PAM II on NetApp storage for Scalability Benchmark.

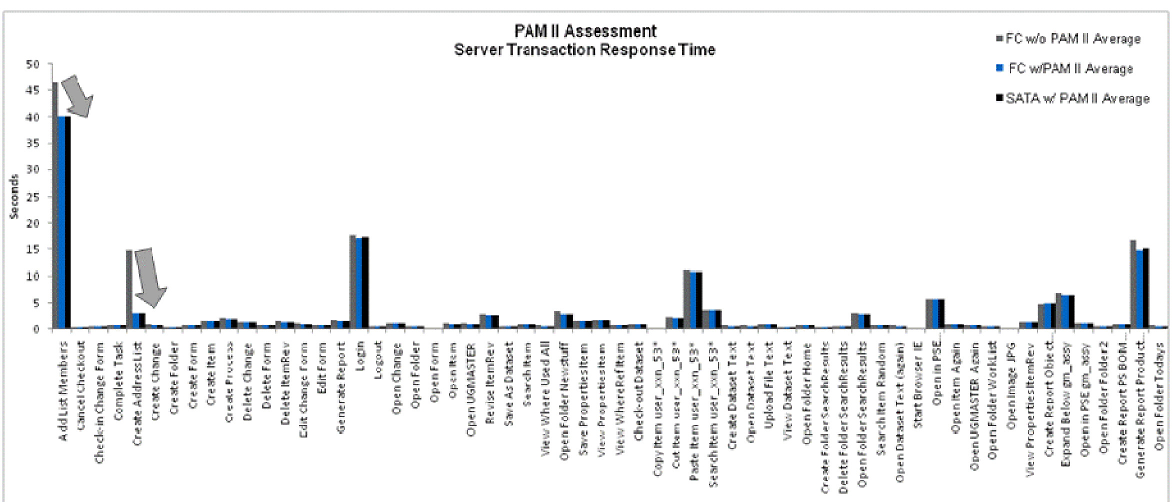

Figure 12) Transaction response time for Teamcenter SQL queries with PAM II.

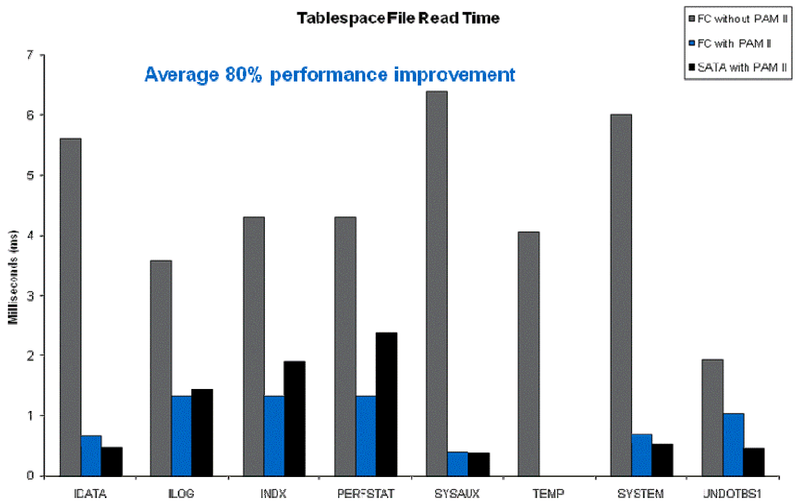

Figure 13) Oracle tablespace read response time during APA (scalability) test.

7 TEAMCENTER 2007 MP3 DISK SPACE SAVING WITH DEDUPLICATION

NetApp deduplication is a general-purpose space reduction feature available on NetApp storage systems. When deduplication is enabled, all data in the specified flexible volume can be scanned at intervals and duplicate blocks removed, resulting in reclaimed disk space. Installing the NearStore® option and A-SIS licenses on a supported system enables the deduplication capability on the primary and secondary storages. Although these licenses are required, there is no charge for either one.

Enabling deduplication on primary and secondary storages in a tiered architecture does not impose any additional performance overhead on the controller. However, NetApp highly recommends not enabling deduplication on primary applications like Teamcenter running on primary storage with high I/O throughput during the day, but instead enabling it when the application is quiescent at night and on weekends.

NetApp deduplication runs as a background process, and the system can perform any other operation during this process. It is also independent of protocols and thus supports all NetApp SAN and NAS protocols: iSCSI, FCP, NFS, and CIFS. For more information, see TR-3505, “NetApp Deduplication for FAS Deployment and Implementation Guide.”

7.1 OBJECTIVES OF DEDUPLICATION TESTING

The objectives of this test are to enable deduplication on the NetApp primary storage during the benchmark as described in section 5 and to evaluate two things:

- The disk space savings on the Teamcenter volume in the NetApp primary storage
- The Teamcenter application response time when the fingerprint database is created for the first time when deduplication is enabled on the storage under heavy load

This test was extended to run the benchmark load on the storage when deduplication is in progress with the PAM card enabled. The response time for the Teamcenter application with deduplication and PAM was compared to the results achieved from tests documented in the earlier sections of this report.

7.2 APPROACH TO DEDUPLICATION TESTING

A similar benchmark workload was used as described in section 6.2 to conduct the deduplication testing with and without PAM.

7.3 RESULTS OF DEDUPLICATION TESTING

The way the benchmark from Siemens PLM works is that it refreshes the benchmark data each time the test is run. During the refreshing cycle, new files are created with different file IDs. The A-SIS algorithm runs the deduplication process each time the test data is refreshed. While the actual benchmarking test is in progress, the deduplication of the data blocks is not complete. The deduplication continues to run in the background even after the Siemens PLM benchmark has completed the run successfully. After the deduplication completes successfully on the `tcengvault` volume (the actual location of the Teamcenter files), we observed ~57% disk savings in that volume. We also observed ~12% disk space savings for the Oracle data files that hold the metadata information for the Teamcenter files. In reality, an ideal Teamcenter environment consists of a small metadata part and a very large number of files (also called a vault). In this benchmark, the database was 40GB, whereas the Teamcenter vault was 206GB before running the deduplication.

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 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..

Even though the use of PAM was tested in this environment, it did not yield any positive results because deduplication still continues in the background, collecting duplicate data blocks even after the benchmark has finished executing. The idea behind testing PAM with deduplication was to find out the benefits of having just one copy of a data block in the buffer rather than having multiple copies. If the benchmark had the capability to run on the same data again, then the PAM would definitely be a great help. However, based on the way the benchmark suite behaves as described earlier, the data has to be refreshed each time by creating new files with different file IDs. The limitation of the test does not rule out the fact that deduplication with PAM not only results in space savings on the disk but also results in only a single copy of a data block in the PAM, allowing more data to be cached. In large Teamcenter deployments with large vault size, the PAM cards can be even more effective by storing a single data block that can be shared by multiple client I/O requests.

The following table shows the percentage of disk space saved on the NetApp storage.

Filesystem	used	saved	%saved
/vol/nas/	2003440	0	0%
/vol/oradata/	38198496	5122960	12%
/vol/orahome/	6990080	448228	6%
/vol/oralogs/	1560320	23052	1%
/vol/tcengvault/	**399831168**	**528940332**	**57%**
/vol/scaletest/	24715984	6430868	21%
/vol/vol0/	1156724	0	0%

While the deduplication process was creating a new hash table during a benchmark run, there was almost no impact on the application read response time.

8 TEAMCENTER 2007 MP3 BACKUP AND RECOVERY USING SNAPSHOT

A virtual read-only copy of the file system is created by taking a Snapshot copy of the root inode. This copy is stored in the active file system as the .snapshot directory in UNIX and snapshot in Windows. Snapshot copies can be created in just a few seconds and require minimal disk capacity. When changes are made to files in the active file system, they are stored in new blocks, and pointers in the file system are adjusted accordingly. A maximum of 255 Snapshot copies can be kept for each file system volume at any point in time. Additional disk capacity may be required, depending on the dynamics of a file system (an average of 10% to 20% of the net capacity for a total of 255 maximum Snapshot copies).

SnapRestore enables an entire file system to be restored from a previous Snapshot copy. SnapRestore copies the inode table for a Snapshot copy back to the inode table for the active file system. SnapRestore requires only a few seconds to restore the data, depending on the file system size.

In this configuration, the entire data structure for the NetApp development lab environment was backed up by using Snapshot technology, and SnapRestore was used to recover a corrupted or missing file.

8.1 OBJECTIVES OF BACKUP TEST

The objective of this configuration is to research how quickly a backup of the Teamcenter environment can be created without affecting the users accessing the database and get it back to a normal read/write state.

Apart from this, a test was also performed to establish the restore times for a complete Teamcenter environment by using a single NetApp storage instance with SnapRestore. There are several reasons that the Teamcenter environment may need to be restored to a previously known good copy:

- Corruption in the underlying Oracle data files, indexes, or rollback segments
- Unplanned power outages to the FMS server or Teamcenter database server
- Misapplication of patches or upgrades to any application in the environment
- Unplanned network outages that cause transactions to be lost between the rich client user and the Teamcenter database server

8.2 APPROACH TO BACKUPS

In this test, the Oracle data files and FMS files were stored on the NetApp storage.

Teamcenter and Oracle Database Backup:

1. The backup script requests Teamcenter engineering to change the state of Teamcenter and freeze all operations on the Teamcenter engineering file system volumes, using the `backup_modes` utility.
2. Teamcenter engineering sends a cautionary message to the users to save all files. Teamcenter engineering is placed in read-only mode once there are no open Teamcenter engineering files in the system.
3. The underlying Oracle Database is put in hot backup mode, allowing Teamcenter to stay available 24x7. This is accomplished by making sure that the database is in archivelog mode and that each tablespace is in hot backup mode.
4. SnapManager® for Oracle (SMO), along with SnapDrive®, makes the backup process for the production database and the Teamcenter vault (that is, the file system repository, TCFS) fast and easy to manage with minimal or no user disruption. The traditional method—switching the TCFS volume to “blobby mode,” writing the user data to a temporary location, and then copying the contents from that location to the TCFS volume when it comes back to normal read/write mode—no longer holds true. SMO uses its own database as a repository for backing up the entire production database (Oracle DB Server) used in the Teamcenter environment, and uses its agent on the TCFS server to backup the TC files from the `tcengvault` volume at the same time. SnapDrive and SMO handle the simultaneous creation of Snapshot copies from the Oracle and the TCFS volumes, keeping both of them consistent. SnapDrive and SMO are independent of any protocol. They can function together on NAS and SAN environments. For details about setting up the backup, see the “Teamcenter Backup/Restore Operational Guide.”
5. After successful installation of SMO, the backup script uses SMO to change the Oracle DB state to a hot backup mode and at the same time uses the Teamcenter `backup_mode` utility to notify Teamcenter users to start saving their open files because the system is getting ready for backup. After a grace time of five minutes (tunable), the script changes the TCFS volume state to readonly and simultaneously takes Snapshot copies on the Oracle DB and TCFS volumes. As soon as the copies are taken on the appropriate volumes, SMO changes the Oracle DB back to a normal state. At the same time the `backup_mode` utility changes the TCFS volume back to a normal state. In our test case, the script was able to create the copies in the respective volumes within two minutes. The actual time taken depends on the size of the Oracle DB and the TCFS volumes.
6. Once the Snapshot copy of the Oracle and TC volumes and/or LUNs is taken and made available to users, the backup software can back up the copies to tape or replication software such as SnapMirror® and SnapVault® and can replicate the Snapshot copies to a secondary storage, either locally for operational backup and recovery or to a different site for disaster recovery or off-site tape archiving. NetApp highly recommends using disk-to-disk backup from the primary to the secondary storage for easy and quick recovery and also for reliability.

8.3 RESULTS OF BACKUP TESTING

After the Snapshot copy creation was completed for the NAS volumes or LUNs on the NetApp storage, the Teamcenter users resumed normal I/O to the files in the volumes. Using the backup scripts, data was backed up in the background to the Snapshot copy that was created on the NetApp storage volume, without affecting the Teamcenter engineering. The entire backup window was reduced to the time required to create the Snapshot copy—a couple of minutes in our test case. Backing up the database by using SnapManager for Oracle was fast, and the Snapshot copies were verified to check the data integrity.

8.4 APPROACH TO RECOVERY

One of the files in the FMS and Oracle data volumes was renamed, and one of the directories was moved to a different location, thus affecting the normal functioning of the environment. SnapManager for Oracle was used to back up the entire volume.

Teamcenter and Oracle Database Restore:

1. Use the Data ONTAP SnapRestore feature if you want to recover Oracle or Teamcenter data that is no longer available or if you are testing a volume or file and want to restore that volume or file to pretest conditions. SnapRestore is used to recover from data corruption. If a primary storage system application corrupts data files in a volume, you can revert the volume or specified files in the volume to a Snapshot copy taken before the data corruption.
2. Single file SnapRestore (SFSR) can be used to revert a single file to a selected Snapshot copy. This is a more practical approach when the file is so large that you cannot copy the previous file version from the Snapshot copy to the active NFS file system.
3. Rapid LUN restore using SnapDrive can restore the LUN, making it available to the host for I/O operations within a few seconds.
4. SnapManager for Oracle also allows quick restore of the database, reducing system and storage overheads and increasing productivity.

Note: All files in a reverted volume have timestamps that are the same as those when the Snapshot copy was created. After a revert operation, incremental backup and restore operations on the file or volume cannot rely on the timestamps to determine what data needs to be backed up or restored.

- Perform a base-level backup of the volume after you restore it.
- When restoring data from tape, use only the backups created after the volume was restored.

8.5 RESULT OF RECOVERY TESTING

SnapRestore was used to revert the volume on the NetApp storage to an earlier state from one of the Snapshot copies that was manually created while initiating the backup process. Single-file SnapRestore was used to restore the single file that was renamed, which the Teamcenter needed for its environment. The single-file SnapRestore process was complete in seconds, and the restore of the 50GB database, using SnapManager for Oracle, was done in a matter of minutes. This greatly reduced the downtime in the environment compared to restoring from tape, which would have taken hours.

9 TEAMCENTER FAILOVER USING NETAPP STORAGE CLUSTERING

A cluster can be created with two NetApp storage heads by connecting the storage heads via a cluster interconnect. This connection is redundant and is used to exchange cluster heartbeats and synchronize the nonvolatile RAM (NVRAM) on both storages. The disk shelves of the cluster partner are connected to the second storage head via a second Fibre Channel loop. If the first storage head fails, the second storage head assumes control of the disk shelves. The IP addresses of both the storage heads are part of the routing table and are maintained in the NVRAM that is mirrored on both storage heads via the cluster interconnect; no data is lost. In case of disruption that leads to a storage head failure, the routing information stored in the partner NVRAM starts to accept network traffic on behalf of the node that went down.

9.1 OBJECTIVE OF FAILOVER AND RECOVERY TIME TEST

The objective of this test configuration was to establish the failover and recovery times for a complete Teamcenter environment using a clustered NetApp storage installation. This failover and recovery can occur for many reasons:

- Unplanned network connectivity outage to a storage head
- Unplanned power outage to a storage head
- Unplanned outage of the storage head

In any of these cases, the storage heads operating in cluster mode fail over to the partner storage head and continue processing transactions.

9.2 APPROACH TO FAILOVER AND RECOVERY TIME TEST

In this test, the Oracle Database files and PRT file were stored on clustered NetApp storage in the NetApp development lab. While the test was being executed, one of the storage heads was physically powered down. Test metrics were taken from the lab environment to establish the time taken to fail over to the new storage and the impact on Teamcenter performance during the failover time.

9.3 RESULT OF FAILOVER AND RECOVERY TIME TEST

In tests conducted with the NetApp storage operating in clustered mode, the complete failover time was 83 seconds. This is within the NetApp guidelines of 90 seconds for a complete failover to take place. The throughput of the Teamcenter application was unaffected by this failover process.

10 MULTISITE TEAMCENTER DEPLOYMENT

To enable users in a globally dispersed enterprise to access and employ product information managed under multiple product definition databases, Teamcenter provides a number of highly versatile data- sharing capabilities using multisite collaboration.

Typically, companies adopt multisite collaboration to accomplish the following tasks.

- Establish a federated network; that is, enable multiple sites that maintain different business rules, user or group controls, and/or workflow-driven processes to exchange and share product definitions while retaining their local autonomy.
  - Lower their development costs
  - Access more specialized skill sets
  - Shorten the development cycle
  - Mitigate financial risk
  - Improve their ability to deliver localized product content
- Overcome network issues by distributing Teamcenter databases and data geographically so that they are closer to the user population or discipline-specific databases, thus optimizing the use of an enterprise's network and delivering higher performance to Teamcenter users.

To facilitate multisite collaboration, Teamcenter leverages the following two key technologies.

SHARED OBJECT REPOSITORY

A shared object repository is a single database instance that is shared across multiple sites and that contains entries for all objects that are intended to be shared across those collaboration sites. This is called Object Directory Services (ODS). It maintains a record of every information object that you intend to share across the Teamcenter network. The shared repository does not physically store the objects themselves, but it does maintain a record of where the information objects reside (that is, their database site), as well as the attributes that users can reference to narrow their database searches. When a site wants to share an object, it "publishes" the object in the shared repository by indicating what network-based sites are authorized to import this particular object. Later, these authorized sites can query the central database, view its published objects, and import appropriate objects into their local database at their discretion.

Sharing Write Access to Shared Data

Generally, data sharing does not involve modifying the shared data. Typically, sites replicate a part for use as an assembly component with no intention of modifying the part itself. There are cases, however, in which a remote site needs to modify data owned by another site. In these situations, multisite collaboration provides two methods for sharing write access to shared data: transferring ownership and remote check-in and checkout.

Transferring Ownership

During the transferring ownership process, the remote site imports the object with transfer of site ownership. For items, this requires transferring site ownership of all revisions and most attachments and files. Ownership access by remote users is controlled by the owning site using site preferences and access management rules.

If an item owned by Site1 is replicated to Site2, and the item's site ownership is transferred to Site3, the site ownership of the replica at Site2 is not updated to show the new owning site. Using the data_sync utility at Site3 does not update the replica at Site2, because the last modification date of the master copy at Site3 has not changed. It is not necessary to sync the owning site property, because the replica at Site2 has not changed. To sync the replica at Site2, run the `data_share` utility at Site3 or perform a remote import at Site2.

Remote Checkin and Checkout

The remote site checks out the object by first replicating the item and then checking out the specific portion of the item requiring modification, such as an attached data set. When the replica is checked out, a remote checkout is performed at the item's owning site, so that no other user in the multisite collaboration network can modify it.

When all modifications have been made to the replica, it is checked in to the owning site. All changes are sent to the owning site and the remote check out status is removed. Any new objects created are owned by the item's owning site.

This method avoids transferring site ownership of an entire item when write access is required only for portions of the item. For performance reasons, Siemens PLM Software recommends using this method whenever possible.

MULTISITE DATA REPLICATION

Multisite data replication provides a network-based environment with import/export capabilities at the Teamcenter application layer that is different from the multisite collaboration described earlier, enabling sites to push or pull product information in accordance with established business rules and workflows. Typically, authorized sites leverage the system's import capabilities to transfer a reference copy of the shared object into their local database. By enabling you to "replicate" shared objects at multiple sites close to your user communities, Teamcenter lets you maximize the Teamcenter network's data access performance. This increases the cost of management because there is no proper tracking of the files in shared environments. Locking the file is not possible. This kind of setup is familiar in smaller Teamcenter environments in which a directory or a group of database objects is exclusively exported to remote users and there is no sharing between the sites. The directory with preset authorization or permission that is exported by the Teamcenter server in the central location can be imported only by the remote Teamcenter server. The remote Teamcenter server cannot import any other object from the central location without proper authorization from the central server.

11 APPENDIX

ENVIRONMENT VARIABLES

Software Component	Parameter Setting	Value	Notes
Siemens PLM Software	Teamcenter	2007 MP3	GA release of Siemens PLM Software PLM application
Teamcenter 2005 SR1/2007 MP3	User login rate	100–1,000 users	Usage profile includes: • Number of users • Types and categories of users • Named versus concurrent versus active user ratios • User login rate • Features used by each user type • How frequently users use features
NetApp StorageView	Oracle Database Volume	Oradata: 40Gb flexible volume, 16-disk aggregate, 2-disk parity (RAID-DP)	Settings for Oracle Database volume: • Style: UNIX • Oplocks: enabled • Quotas enabled: off • Minimal read ahead: off • Update access time: on • Snapshot copies: off • Snapshot directory visible: off • Aggregate RAID group size: 16 (RAID-DP)
NetApp StorageView	Oracle Database logs, Oracle home	Oralog/Orahome: 16Gb each flexible volume, 16-disk aggregate, 2-disk parity (RAID-DP)	Settings for Oracle logs/home volumes: • Style: UNIX • Oplocks: enabled • Quotas enabled: off • Minimal read ahead: off • Update access time: on • Snapshot copies: off • Snapshot directory visible: off • Aggregate RAID group size: 16 (RAID-DP)
NetApp StorageView	Teamcenter: tcengvault volume	TCengvault: 819Gb flexible volume, 16-disk aggregate 2-disk parity (RAID-DP)	Settings for Teamcenter tcengvault volume: • Style: UNIX • Oplocks: enabled • Quotas enabled: off • Minimal read ahead: off • Update access time: off • Snapshot copies: on • Snapshot directory visible: on • Aggregate RAID group size: 16 (RAID-DP)
NetApp StorageView	Teamcenter: scaletest	Scaletest: 80Gb flexible volume, 16-disk aggregate, 2-disk parity (RAID-DP)	Settings for Teamcenter scaletest volume: • Style: UNIX • Oplocks: enabled • Quotas enabled: off • Minimal read ahead: off • Update access time: off • Snapshot copies: on • Snapshot directory visible: on • Aggregate RAID group size: 16 (RAID-DP)

Sun Fire™ V440	Teamcenter Load Generator	Loadfiles	A different partition of 72GB on the local hard drive on the client for the load files. This reduces the NFS traffic over the wire. Create an fmsload_files_ 2_import directory under this partition.
Teamcenter	Weblogic	License	Not required for component-level tests.
Teamcenter load generator	All business cases: Teamcenter user login/password	Randomly selected from low to medium number of users	The login activities during peak periods for the 1,000-user thin client benchmark.
Load generator: component-level benchmark	Business case 1 retrieve mode	Reads an input file that contains string representations of tags of file objects	The program runs until 60 minutes have passed, or until 10,000 files have been exported, or until the input file has been processed, whichever comes first.
Load generator: component-level benchmark	Business case 2 create mode	Reads an input file that contains a list of files in the file system of the load generated	The program runs until 60 minutes have passed, or until 2,000 files have been imported, or until the input file has been processed, whichever comes first.
Load generator: component-level benchmark	Business case 3 delete mode	Reads tags from an input file and deletes the corresponding file object from Teamcenter engineering	The program runs for 60 minutes, or until 2,000 files have been deleted, or until the input file has been processed, whichever comes first.
Backup and recovery on 2007 MP3	Snapshot, SnapRestore, single file SnapRestore	On the NetApp storage volume	Snapshot copies and single file SnapRestore are done in seconds.
Rich client	Version	2007 MP3	No additional patches were applied.
Oracle	Database	Oracle10*g* Release 2	No additional patches were applied.
Solaris™	OS version	10	No additional patches were applied.
Windows	OS version	XP Professional	No additional patches were applied.

ORACLE RDBMS SETTING

```
###############################################################################
# Copyright (c) 1991, 2001, 2002 by Oracle Corporation
###############################################################################

###########################################
# Cache and I/O
###########################################
db_block_size=8192
db_cache_size=25165824
db_file_multiblock_read_count=16

###########################################
# Cursors and Library Cache
###########################################
open_cursors=300

###########################################
# Database Identification
###########################################
db_domain=""
db_name=netapp
```

```
###########################################
# Diagnostics and Statistics
###########################################
background_dump_dest=/u01/app/oracle/oracle/product/10.2.0/db_5/admin/netapp/bd
ump
core_dump_dest=/u01/app/oracle/oracle/product/10.2.0/db_5/admin/netapp/cdump
user_dump_dest=/u01/app/oracle/oracle/product/10.2.0/db_5/admin/netapp/udump

###########################################
# File Configuration
###########################################
control_files=("/u02/oradata/netapp/control01.ctl",
"/u02/oradata/netapp/control02.ctl", "/u02/oradata/netapp/control03.ctl")
db_recovery_file_dest=/u01/app/oracle/oracle/product/10.2.0/db_5/flash_recovery
_area
db_recovery_file_dest_size=2147483648

###########################################
# Instance Identification
###########################################
instance_name=netapp

###########################################
# Job Queues
###########################################
job_queue_processes=10

###########################################
# Miscellaneous
###########################################
compatible=10.1.0.2.0

###########################################
# Network Registration
###########################################
local_listener=LISTENER_NETAPP

###########################################
# Optimizer
###########################################
optimizer_index_caching=95
optimizer_index_cost_adj=10
optimizer_mode=CHOOSE
query_rewrite_enabled=TRUE
query_rewrite_integrity=TRUSTED

###########################################
# Pools
###########################################
java_pool_size=0
large_pool_size=8388608
shared_pool_size=536870912

###########################################
# Processes and Sessions
###########################################
processes=2048
sessions=2258
```

```
##########################################
# Security and Auditing
##########################################
remote_login_passwordfile=EXCLUSIVE

##########################################
# Sort, Hash Joins, Bitmap Indexes
##########################################
pga_aggregate_target=25165824
sort_area_size=65536

##########################################
# System Managed Undo and Rollback Segments
##########################################
undo_management=AUTO
undo_retention=900
undo_tablespace=UNDOTBS1
Database Server
```

STARTING TEAMCENTER SERVICES

There are several processes that must be started in order for the Teamcenter engineering environment to function properly. Run the following scripts to start these services if necessary after a reboot. All processes should run as root except for TCFS, FSC, and the server manager, which should be run by the infodba user.

License Manager	`/scaletest/released/sol/tc2005_SR1/ugflexlm/rc.ugs.flex`
TCFS	`/scaletest/released/sol/tc2005_SR1/bin/rc.ugs.tcfs`
FSC	`/scaletest/released/sol/tc2005_SR1/fms/rc.ugs.FSC_netapp`
Server Manager	`/scaletest/released/sol/tc2005_SR1/pool_manager/startpm_netapp`
WebLogic Server	`/usr/local/bea/user_projects/domains/scaletest/startWebLogic.sh`
Apache[1]	`/usr/local/apache2/bin/apachectl start`
RMIRegistry[1]	`/scaletest/released/sol/tc2005_sr1/portal_otw/server/webapp_root/start_rmi`
Distribution Server[1]	`/scaletest/released/sol/tc2005_sr1/portal_otw/server/webapp_root/start_server`

1. The OTW services are not required at runtime. However, they must be running during installation of a rich client.

Siemens PLM Software License Server (FlexLM) Must Be Running

```
sunv210-svl04 # ps -ef|grep flex
root   748   747   0   Jun 26 ?          2:32 uglmd -T sunv210-svl04 10.8 4 -c
/scaletest/released/sol/tc2005_sr1/ugflexlm/ps
root   747     1   0   Jun 26 ?          2:01
/scaletest/released/sol/tc2005_sr1/ugflexlm/lmgrd -c /scaletest/released/sol/tc
```

Teamcenter Engine Must Be Running

```
sunv210-svl04 (/scaletest/released/sol/tc2005_sr1/ pool_manager)#
./startpm_netapp
Checking OS ...OK
Set descriptor limit
1024
Checking pool_manager directory ...OK
Checking pool_manager processes ...OK
Checking /var/tmp for old tcserver log and syslogs ...
Removing old tcserver log and syslogs  ....
Removing old serverPool log ....
```

```
Starting the pool manager ....
sunv210-svl04 (/scaletest/released/sol/tc2005_sr1/pool_manager) # ps -ef|grep
tc
root 16607 16598   1 13:46:13 pts/3       0:03
/scaletest/released/sol/tc2005_sr1/install/install/jre/bin/java -Xmx512m -Djaco
root   222     7   0   Jun 26 console     0:00 /usr/lib/saf/ttymon -g -d
/dev/console -l console -T sun -m ldterm,ttcompat -h
root   693     1   0   Jun 26 ?           0:00 /usr/lib/snmp/snmpdx -y -c
/etc/snmp/conf
infodba   762     1   0   Jun 26 ?           0:03
/scaletest/released/sol/tc2005_sr1/bin/tcfs -t=tcpip -s=tcfs
root   748   747   0   Jun 26 ?           2:25 uglmd -T sunv210-svl04 10.8 4 -c
/scaletest/released/sol/tc2005_sr1/ugflexlm/ps
infodba   743     1   0   Jun 26 ?           0:03
/scaletest/released/sol/v9125/bin/imanfs -t=tcpip -s=imanfs
root 16610 16607   9 13:46:22 pts/3       0:10
/scaletest/released/sol/tc2005_sr1/bin/tcserver id=tcserver2@pool_netapp_sunv21
 root   747     1   0   Jun 26 ?           1:55
/scaletest/released/sol/tc2005_sr1/ugflexlm/lmgrd -c /scaletest/released/sol/tc
root 16613 16607   9 13:46:25 pts/3       0:10
/scaletest/released/sol/tc2005_sr1/bin/tcserver id=tcserver3@pool_netapp_sunv21
root 16608 16607   8 13:46:19 pts/3       0:10
/scaletest/released/sol/tc2005_sr1/bin/tcserver id=tcserver1@pool_netapp_sunv21
root 16598     1   0 13:46:13 pts/3       0:00 /bin/sh
/scaletest/released/sol/tc2005_sr1/pool_manager/mgrstartnetapp
root 16622 15529   0 13:46:53 pts/3       0:00 grep tc

sunv210-svl04 (/scaletest/released/sol/tc2005_sr1/pool_manager) # ps -ef| grep
java
root 16607 16598   0 13:46:13 pts/3       0:04
/scaletest/released/sol/tc2005_sr1/install/install/jre/bin/java -Xmx512m -Djaco
root   395     1   0   Jun 26 ?          12:23 java -enableassertions -
javaagent:/opt/NTAPsmo/webapp/WEB-INF/lib/aspectj-1.5.2
infodba   738     1   0   Jun 26 ?          12:43 java -
Dfms.config=./fmsmaster_FSC_netapp.xml -Xmx128M -server -Dfsc.config=./FS
```

HARDWARE SPECIFICATIONS

Figure 14 shows the physical layout of the NetApp development lab.

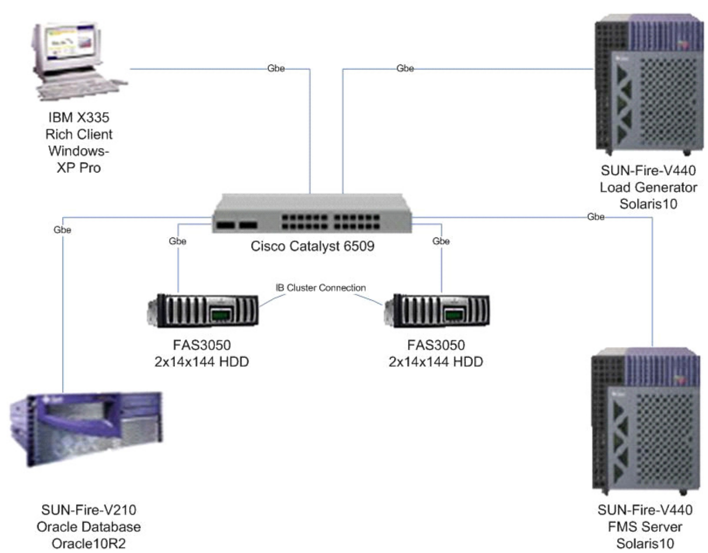

Figure 14) Physical network diagram of lab environment.

All of the physical servers and storage arrays were hosted at the NetApp facilities in Sunnyvale, California. During performance testing, network traffic and general environmental load were monitored so that no external factors influenced the test results.

DATABASE Server/Load Generator

Model	Sun Fire V440
CPU	2 x sparcv9 processor 1593MHz with a sparcv9 floating point processor
Memory	8GB
Hard drive	2 x 72GB
OS	Solaris 5.10 generic
FMS Server Model	Sun Fire V440
CPU	2 x sparcv9 processor 1593MHz with a sparcv9 floating point processor
Memory	8GB
Hard drive	2 x 72GB
OS	Solaris 5.10 generic

Rich Client

Model	IBM X335
CPU	Xeon™ processor 2.80GHz
Memory	1GB
Hard drive	40GB
OS	Windows XP Professional

NetApp Storage for FMS Benchmark

Model	FAS3050c / FAS3070c
OS	Data ONTAP 7.3.1
Disk shelves	DS-14 Fibre Channel (2 total)
Disk drives	144GB: 15k RPM
NIC	Gigabit
FC-AL disk adapters	2

NetApp Storage for PAM II Testing

Model	FAS3170c
OS	Data ONTAP 7.3.2
Disk shelves	DS-14 Fibre Channel (2 total)
Disk drives	144GB: 15k RPM
NIC	Gigabit

FC-AL disk adapters	2
Performance Accelerator Module(PAM)	1

Network Connectivity

Lab Environment HW Component	Connectivity
Sun Fire V440	Gigabit
Sun Fire V440	Gigabit
IBM X335	Gigabit
Storage array	Fibre Channel

References

- UGS Teamcenter Engineering – Deployment and Operational Considerations for NetApp Filers
  http://www.netapp.com/library/tr/3171.pdf
- Teamcenter 2007 MP3 Deployment Guides
  Teamcenter Engineering - http://support.ugs.com/docs/tc_eng/
  Teamcenter - http://support.ugs.com/docs/teamcenter/
- Siemens PLM Software Certification / Support Matrix
  http://support.ugs.com/online_library/certification/index.php?interface=external_interop&p=p&v1=9&s