



NetApp™
Go further, faster

Technical Report

Best Practices for Deploying Rational ClearCase Using NetApp Storage

Anil G., Subra Mohan, NetApp
June 2009 | TR-3774

Updated June 2012

TABLE OF CONTENTS

1	INTRODUCTION TO CLEARCASE AND CLEARCASE MULTISITE.....	4
2	ROLE OF SCM IN SOFTWARE DEVELOPMENT ENVIRONMENTS	4
3	CLEARCASE ARCHITECTURE AND DEPLOYMENTS	4
3.1	VOB STORAGE AND SERVER	4
3.2	VIEW STORAGE AND SERVER.....	5
3.3	REGISTRY SERVER.....	5
3.4	BACKUP REGISTRY SERVER.....	5
3.5	LICENSE SERVER.....	5
4	CUSTOMER CHALLENGES.....	6
5	ADVANTAGES OF USING NETAPP STORAGE FOR RATIONAL CLEARCASE	6
5.1	IMPROVED STORAGE UTILIZATION	6
5.2	ENHANCED PERFORMANCE AND SCALABILITY	7
5.3	MINIMUM ADMINISTRATION OVERHEAD	7
5.4	DATA PROTECTION AND HIGH AVAILABILITY	7
5.5	OPTIMIZED BUILD ENVIRONMENTS	7
5.6	EFFICIENT TEST ENVIRONMENTS.....	7
6	CONFIGURING RATIONAL CLEARCASE WITH NETAPP STORAGE.....	8
6.1	DATA FLOW IN CLEARCASE ENVIRONMENT	8
7	BEST PRACTICES.....	9
7.1	RESILIENCY BEST PRACTICES.....	9
7.2	DISK AND DATA LAYOUT PRACTICES.....	10
7.3	OTHERS.....	10
8	SIZING GUIDELINES	10
9	THIN PROVISIONING	11
10	HIGH AVAILABILITY OF RATIONAL CLEARCASE WITH NETAPP	11
11	PERFORMANCE BENCHMARKS ACROSS PROTOCOLS	12
11.1	SUMMARY OF PERFORMANCE BENCHMARK RESULTS	14
12	BEST-FIT DEPLOYMENT OPTIONS.....	15
12.1	UNIX VOB SERVER OPTIONS.....	15
12.2	WINDOWS VOB SERVER OPTIONS.....	16
12.3	VIEW DEPLOYMENT ON NAS	17
13	DECISION TREE	17
14	PERFORMANCE-RELATED ADD-ON OPTIONS	20
14.1	PERFORMANCE ACCELERATION MODULE (PAM).....	20

15 DEDUPLICATION IN RATIONAL CLEARCASE ENVIRONMENTS USING NETAPP STORAGE	20
16 DATA PROTECTION OPTIONS	21
16.1 ONLINE BACKUP AND RECOVERY	21
16.2 DISASTER RECOVERY	22
16.3 DISK-TO-DISK BACKUP	23
17 CONCLUSION	23
18 ADDENDUM	24

1 INTRODUCTION TO CLEARCASE AND CLEARCASE MULTISITE

IBM Rational ClearCase is an enterprise software configuration management tool that enables development teams to capture, control, and securely manage software changes through the entire cycle of software development. It provides version control, workspace management, build management, and process configurability to simplify the development process.

In today's global scenario, development happens simultaneously across multiple geographical regions. IBM Rational ClearCase MultiSite facilitates such parallel development capabilities by providing developers at different geographical locations with simultaneous access to the same software project assets through replication and synchronization of ClearCase repositories. With the help of a reliable mechanism, it replicates repositories to multiple locations and provides geographically distributed developers with local copies, called replicas, to work with. It also manages synchronization and coordination between different replicas to avoid version conflict.

2 ROLE OF SCM IN SOFTWARE DEVELOPMENT ENVIRONMENTS

Software development is a complex and dynamic process that involves managing multiple different versions of the same instance before a released version can be established. The development process is often spread across geographical boundaries and requires seamless coordination between local and dispersed development teams.

Software Configuration Management (SCM) tracks and controls changes in your software assets and helps you to effectively control the evolution of a software project by providing version control, build management, release management, and process control.

IBM Rational ClearCase is a leading enterprise SCM solution that simplifies the process of change management in large development environments, provides workspace management, and provides parallel development support to improve productivity and accelerate time to market.

3 CLEARCASE ARCHITECTURE AND DEPLOYMENTS

IBM Rational ClearCase is a distributed application with client/server architecture. It is comprised of ClearCase hosts, clients and multiple ClearCase application services such as: VOB service, View service, registry service, and license management architecture.

Each of these services can be installed individually on separate systems or together in a consolidated system. In large Clearcase deployments you might see several VOB service servers just to host VOB services.

3.1 VOB STORAGE AND SERVER

At the heart of a ClearCase environment is a secure data repository referred to as a Versioned Object Base (VOB). It stores versions of file elements, directory elements, derived objects, and metadata associated with these objects. VOBs are hosted on the VOB server and stored on the local disk.

A VOB is comprised of four components: a database and three storage pools.

- The database is an embedded Raima database that records the evolution of version-controlled file system objects and stores the associated metadata, including version labels, hyperlinks, and configuration records.
- The source pool is used for permanent storage of the content of files under Rational ClearCase control. To minimize disk consumption, every text file has a single container file in the source pool.
- The cleartext pool is used for temporary storage of specific versioned files. Though many versions of a file are stored in a VOB, only a few versions are used by the development community at a particular time. These currently "in-use" versions are cached in the cleartext storage pool.

- The derived object pool contains data containers for derived objects and is used only when the Clearmake build utility is utilized in a ClearCase environment. Instead of rebuilding objects, the derived object pool enables the reuse of an object by other work areas.

3.2 VIEW STORAGE AND SERVER

To access VOB data, users need a workspace called View. A View is like a filter that provides users access to specific versions of one or more elements in single VOB or multiple VOBs. Views use a set of rules called configuration specifications to extract the specific version of the artifact that the user sees in a workspace. Using different Views, users can work on different change activities.

Views display specific files and their selected version (according to config spec rules) contained in the VOB. They also maintain copies of files that have been checked out and files that are yet to be checked in the VOB (View private files). Unlike a VOB, which is a permanent repository of software assets, a View stores data temporarily, until the time a user works on it.

Views are hosted locally on the client or on a centralized location on the View server. Though Rational recommends View storage to be local, numerous customers opt for a dedicated host server for Views. This is primarily driven by the need to back up Views from a centralized location since it reduces administrative overhead.

ClearCase provides flexibility to work with both Snapshot and Dynamic Views:

- Snapshot View enables users to work off line by loading copies of file elements from the VOBs into the local system. Users need to access the VOB repository or the Views only when they need to upload their changes.
- Dynamic Views are always connected to the network, enabling transparent, real-time access to the VOB data. When working dynamically, the directory structure that a user uses is controlled by Multiversion File System. Dynamic Views enable fast and immediate updates, build auditing, and derived objects sharing, and provide a global view of code base.

3.3 REGISTRY SERVER

The Rational ClearCase registry is a central repository of information about shared resources such as VOBs and Views. Like centralized “yellow pages,” the Rational ClearCase registry helps users get relevant information in distributed environments.

3.4 BACKUP REGISTRY SERVER

A backup registry host takes periodic snapshots of the primary registry host’s registry files. If the primary registry server fails for some reason, Rational ClearCase provides the ability to “switch over” to a backup registry server.

3.5 LICENSE SERVER

Every user needs a license for using cleartool and multitool commands. These license requests are served by the license server host. There can be multiple license servers in the ClearCase network for serving different sets of licenses.

There are primarily two types of licensing in a ClearCase environment: Atria licensing and FlexLM licensing.

MODES AND CONFIGURATIONS

UNIX MODE: Developers working on UNIX® desktops access VOBs and Views hosted on UNIX servers.

WINDOWS MODE: Developers working on Windows® desktops access VOBs and Views hosted on Windows servers.

INTEROP MODE: Different developers working on UNIX or Windows desktops can access the common source code stored on UNIX servers by configuring either Samba or TAS, because they allow UNIX servers to interact with Windows clients. Rational ClearCase, however, does not support the Interop mode where VOBs and Views are hosted on Windows and accessed by UNIX.

4 CUSTOMER CHALLENGES

Software developers across the globe leverage IBM Rational ClearCase software for version control, workspace management, parallel development support, and build auditing. With developers needing continuous and fast access to ClearCase data, enabling high data availability and consistent storage performance has become a business imperative. The standard VOB backup procedures, however, require locking the VOB database during backup for a longer duration, a paradigm unacceptable in today's 24x7 global environment.

Software development teams often need to access data simultaneously, leading to disk I/O bottlenecks and degraded performance. The complexity of the ClearCase environment, however, makes it difficult to track and analyze the root cause of the problem on time. To add to the challenge, ClearCase uses extremely chatty RPC protocol, which leads to high latency.

Fast growing data environments require effective management of the VOB and View database. Because it is not easy to accurately forecast the storage requirements of VOBs and Views, administrators often end up overprovisioning storage or running out of disk space.

Development environments often require multiplatform support where developers can work in heterogeneous environments. This need to support a mix of Windows, Linux[®], and UNIX environments increases cost and complexity.

To summarize, ClearCase customers face challenges such as:

- Ensuring uninterrupted data availability
- Maintaining optimal data protection
- Optimizing VOB database performance
- Overcoming latency issues caused by RPC protocol
- Easing storage provisioning
- Increasing overall disk utilization
- Providing multiplatform support
- Creating copies of data sets for testing ClearCase upgrades, patches, and diagnostics in a test environment

5 ADVANTAGES OF USING NETAPP STORAGE FOR RATIONAL CLEARCASE

NetApp provides storage solutions that can deliver the right combination of availability, protection, performance, and scalability to address the challenges of Rational ClearCase environments cost-effectively. The NetApp storage platform that is "Ready for IBM Rational software" certified, assures its users that their storage environment is safe, tested, and scalable.

With NetApp[®] solutions you get a range of benefits, including:

5.1 IMPROVED STORAGE UTILIZATION

By consolidating VOB and View data on NetApp, you can leverage FlexVol[®] and deduplication technologies to optimize storage utilization. FlexVol is a storage virtualization technology that pools physical storage and enables you to create and resize virtual volumes per requirement. Free space is aggregated and shared to increase storage utilization, reducing space wastage.

When consolidating your Views on NetApp storage, you are likely to have many duplicate copies of files. NetApp deduplication eliminates redundant data copies and controls data proliferation, can provide substantial storage cost savings. ClearCase users have the option to store source file versions as whole files in conjunction with Data ONTAP[®] deduplication. Whole copies provide can improve versioning times because ClearCase spends less time and needs fewer CPU cycles to assemble the necessary version of each source code file compared to the time taken by delta files. By leveraging deduplication, it is possible to minimize the storage capacity growth that would otherwise result from storing whole files in other storage systems, such as DAS or third-party SANs.

5.2 ENHANCED PERFORMANCE AND SCALABILITY

In a typical Rational ClearCase environment, you need emulation products such as Samba or PC-NFS to allow different platforms to share network files. NetApp unified storage provides better interoperability between UNIX and Windows because it serves NFS and CIFS protocols natively, freeing up CPU resources for ClearCase processes and RPC requests, can improve overall VOB server performance. The NetApp storage system also offers flexibility to support the ClearCase database using FC SAN or iSCSI while supporting source pools on NFS and CIFS protocols, and can provide optimal performance and the flexibility to support a mixed UNIX and Windows desktop environment.

NetApp enables rapid provisioning and seamless scaling. With FlexVol technology, you can increase storage on the fly, with minimal downtime. Non-disruptive capacity expansion and quick performance upgrades help you eliminate the time and complexity associated with storage allocation in complex, shared-access environments.

5.3 MINIMUM ADMINISTRATION OVERHEAD

Storage proliferation and inherent complexities in mixed Windows/UNIX environments add to the administration overhead. Administrators face an uphill task in managing numerous VOB and View servers. NetApp's unified storage architecture simplifies storage administration, enabling you to consolidate your VOB and View storage on a single storage system. Backup of VOBs and Views can be taken from a single source and you can leverage technologies like thin provisioning and deduplication to simplify management, increase storage utilization, and eliminate capital and management costs of multiple, separate storage environments.

5.4 DATA PROTECTION AND HIGH AVAILABILITY

NetApp provides cost-effective data protection solutions to provide uninterrupted ClearCase data availability to users against all odds. NetApp Snapshot™ technology combined with SnapMirror® and SnapRestore® software can provide robust data protection and disaster recovery. Using NetApp Snapshot technology, you can create near-instantaneous backups of VOB and View storage, with minimal impact to performance and data availability. Snapshot copies drastically reduce backup windows, freeing the VOB from being locked for hours. You can leverage SnapMirror and SnapVault® replication software to back up data to a secondary storage, enabling comprehensive data protection.

NetApp brings inherent reliability to development environments with its active-active controller configuration, which could provide up to 99.999% data availability (based on percentage of uptime of NetApp storage systems monitored on a rolling basis between July 2006 and January 2008).

5.5 OPTIMIZED BUILD ENVIRONMENTS

Build environments have specific storage requirements, which can affect the storage performance of the ongoing development process. NetApp solutions help you create an environment in which even the most intense build activity does not slow down storage performance or impede development activity. With NetApp FlexCache®, you can deploy caching storage systems between the primary storage repository and build servers. This relieves the load from the central storage and accelerates read performance automatically without risky duplication.

The NetApp Performance Acceleration Module is another technology that can improve the performance of random read-intensive workloads such as software builds by providing an intelligent, extended read cache in the storage controller.

5.6 EFFICIENT TEST ENVIRONMENTS

Application development often entails creating copies of data sets for testing ClearCase upgrades and patches, running diagnostics, fine-tuning, etc., in a test environment. NetApp FlexClone® technology helps you create space-efficient clones of VOB and View data in seconds, without affecting the production environment. These clones consume space only as changes are made to them, and can provide substantial storage savings and accelerating time to market.

6 CONFIGURING RATIONAL CLEARCASE WITH NETAPP STORAGE

ClearCase customers can put all the data of their VOB, View, and Registry servers—traditionally stored on local disk or on DAS—onto networked-attached storage for enhanced scalability, improved flexibility, and simplified administration.

NetApp was the first NAS platform to offer Rational ClearCase support and also the first storage vendor to achieve the “Ready for Rational” certification. Using NetApp, you can consolidate VOB and View data on NetApp storage and leverage high-performance storage like a dedicated FAS6000 or Data ONTAP GX cluster for software build farm. Though ClearCase is typically deployed on NetApp using NFS, you can also use combinations of NAS and SAN/iSCSI if required.

Using NetApp, you can deploy a ClearCase VOB on various protocols like FC SAN, iSCSI, NFS, and CIFS and also have a split configuration. Though NFS is widely used for ClearCase installations, in rare cases in which database files are changed very frequently, NFS is not ideal because the cache efficiency of a NAS-deployed environment decreases in write-intensive scenarios. For VOB hosts with a lot of memory and large database files, `fdsync ()` and `close ()` are slow for NFS files, extending the VOB database lock times and slowing down overall performance.

To eliminate these drawbacks, the split configuration can be leveraged. Easy to implement, the split configuration facilitates splitting the VOB storage directory (.vbs) to have the database and storage pools residing at different locations. While the database is served through the SAN protocol (iSCSI/FC), the storage pools can be served by either NFS or CIFS.

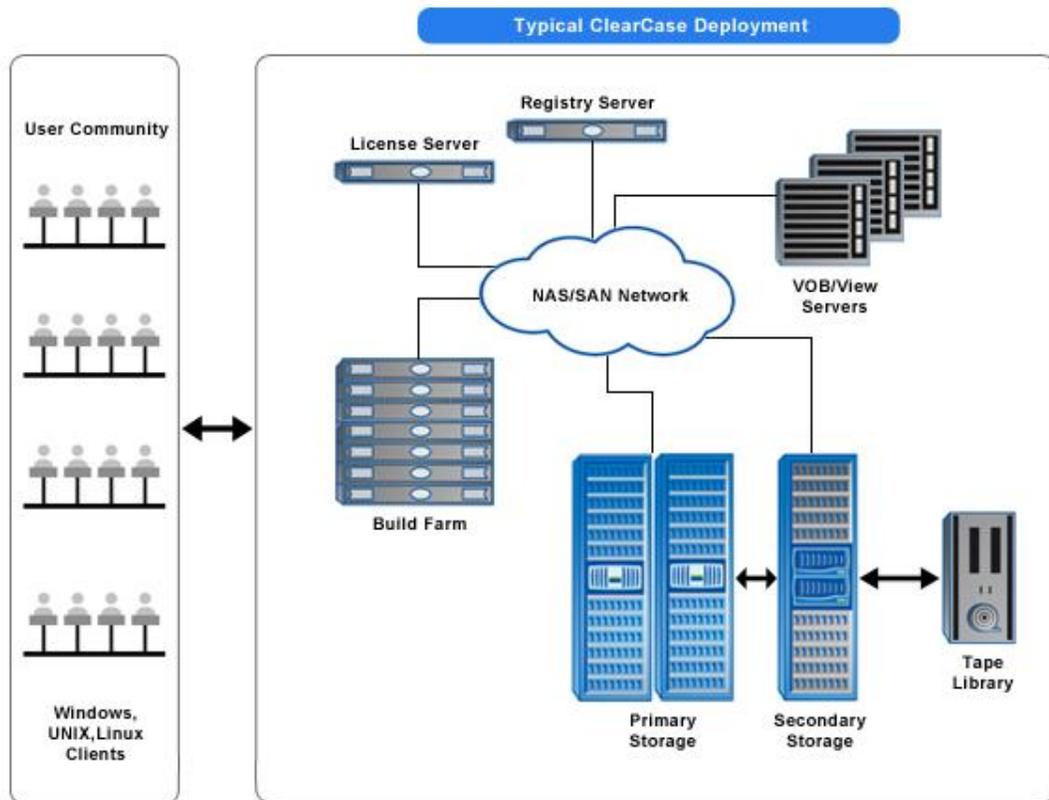


Figure 1) Typical ClearCase deployment.

6.1 DATA FLOW IN CLEARCASE ENVIRONMENT

Rational ClearCase uses an extremely chatty RPC protocol for most of its operations and also uses NFS or CIFS for accessing the files.

When a client program wants to access a service (a VOB or View server or a license request or a registry lookup) on a ClearCase server host, it uses a remote procedure call (RPC) to send a request to the albd_server process on that host (Port number 371). The albd_server initiates the requested service if it is not already started and then issues a response by giving the client the service's port number. (UDP calls with source and destination port information.) Thereafter, the client communicates directly (TCP) with the specific service, without involving the albd_server.

Since the VOB and View storage directories reside on the NetApp storage, the files are served to the clients directly by the storage (native file serving capability) in response to the LOOKUP / ACCESS / GETATTR / WRITE calls from the client.

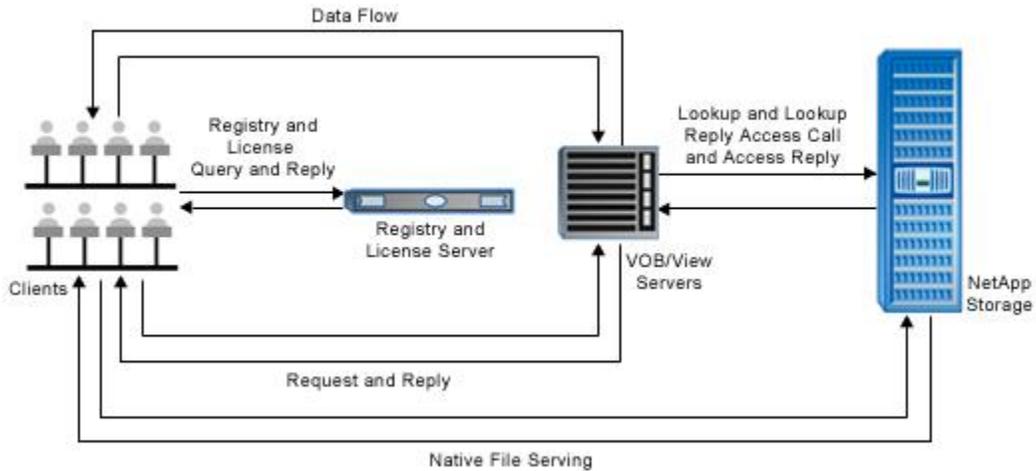


Figure 2) Data flow in ClearCase environment.

7 BEST PRACTICES

This section describes the best practices for configuring Rational ClearCase with NetApp storage.

7.1 RESILIENCY BEST PRACTICES

- Use RAID-DP[®], the NetApp high-performance implementation of RAID 6, for better data protection.
- Use an active-active storage controller configuration (clustered failover) to eliminate single points of failure (SPOFs).
- Use multipath HA with active-active storage configurations to improve overall system availability as well as to promote higher performance consistency.
- Use the default RAID group size when creating aggregates.
- Allow Data ONTAP to select disks automatically when creating aggregates or volumes.
- Use the latest certified Data ONTAP version for IBM Rational ClearCase with appropriate recommended patches, available on the NOW[™] (NetApp on the Web) site.
- Use the latest storage controller, shelf, and disk firmware available on the NOW site.
- Maintain at least one hot spare for each type of disk drive in the storage system:
 - Disk drive differences include FC, SAS, SATA disk drive types; disk size; and rotational speed (RPM)
- Maintain two hot spares for each type of disk drive in the storage system to take advantage of Maintenance Center.
- Reserve space for at least one week of Snapshot backups.
- Do not put user data into the root volume.
- Replicate data with SnapMirror or SnapVault for disaster recovery (DR) protection.
- Replicate to remote locations to increase data protection levels.

7.2 DISK AND DATA LAYOUT PRACTICES

- Use separate volumes for ClearCase VOB and View storage.
- Use FC disks for build-intensive environments.
- Reference VOBs (read only) can be hosted on SATA drives—NearStore®.
- Disable quotas on volumes to be used for VOB storage.
- Interdependent / Admin / Hyperlinked / UCM VOBs should be grouped together.
- Disable oplocks.
- Specify qtree security style. You can use UNIX / NTFS / Mixed, depending upon the requirement.

7.3 OTHERS

- Interop consideration. All Windows user names and their primary groups must have matching user names and primary groups on UNIX (and/or Linux) in a Rational ClearCase Interop environment.
- On UNIX (and Linux), Rational ClearCase relies on group IDs to implement workgroup sharing and to maintain strict access control. To access UNIX VOBs and Views correctly, each Windows user name and primary group name must match its respective user name and primary group name on UNIX.
- Storage used for VOB and View should be on a robust, nonlossy network.
 - Use Gig E between the servers and the appliances.
 - Use TCP.
 - Make sure flow control, speed, and duplex settings match both the switch and the server settings.
- UNIX hosts must mount the appliance using an NFS hard mount.

8 SIZING GUIDELINES

There is no “one size fits all” sizing guideline for a ClearCase environment. The sizing policies are application specific and vary according to the capacity needs and workload characteristics of every project. These policies are determined on the basis of various parameters, including:

- Data footprint and the anticipated growth of VOB and View data over a period of time. To decide a sizing policy, you need to answer certain questions, such as:
 - Do you want VOB, View storage directories, and registry files to be served from the storage? For example, clustered environments require VOB, View, and Registry data to be made available through common storage.
 - In case of a MultiSite environment, are shipping bays being created on storage? MultiSite shipping bays can be configured to hold the shipping packets on a separate volume rather than on the local disk.
 - What is the nature and frequency of build jobs and the resultant data growth?
- Retention of online copies. The sizing policy depends on the number of online copies that need to be retained in a storage device.
- Latency and throughput. Anticipated IOPs and latency time have an impact on the storage size of a particular project. Performance usually gets affected as the size of the VOB increases. However, you can leverage performance-related add-on products like Performance Acceleration Module and FlexCache to enhance your storage performance.
- Protocol or configuration in use. Different protocols like NFS, CIFS, iSCSI, and FCP and configuration modes like split mode have their specific sizing policies and requirements.

9 THIN PROVISIONING

Thin provisioning is a way of presenting more storage space to the hosts connected to a storage system than what has been physically allocated. With thin provisioning, the storage administrator takes an aggregate and creates volumes for different consumers/applications to use, but does not actually allocate any space up front to those volumes. The space gets allocated only when the host application needs it. The unused aggregate space is used for expanding the existing thinly provisioned volumes and for creating new volumes. By allowing “as-needed” provisioning and space reclamation, thin provisioning can result in better storage utilization and reduced capital expenditure on storage infrastructure.

Volume-level thin provisioning can be achieved through any of the following:

- With the volume space guarantee set to none, Data ONTAP allows an administrator to create a large flexible volume (or multiple flexible volumes) that can be greater than the aggregate size. Space only gets allocated to a volume when data is written into the volume.
- An administrator can also, initially, create a smaller-size volume with Data ONTAP and later resize the volume to accommodate growing data.
- The `vol autosize` (autogrow feature) command can also be used with Data ONTAP to automatically increase the volume size once the used space reaches a predefined threshold.

By leveraging various provisioning approaches, you can optimally utilize your resources by allocating less capacity up front and then growing that capacity per need, manually or automatically. This results in seamless scalability and efficient management of the fast-growing ClearCase environment.

You can increase or shrink the ClearCase VOB or View storage volumes on the fly, with minimal downtime. Thus, if VOBs or Views grow much faster than anticipated, you can easily reconfigure the relevant flexible volumes. You can also enable the autogrow option and set thresholds based on the amount of data that is expected to be put in the VOBs.

All these operations take place at the storage level and have no performance impact on the working of ClearCase.

10 HIGH AVAILABILITY OF RATIONAL CLEARCASE WITH NETAPP

IBM Rational ClearCase now provides host-side clustering for high availability (HA). It supports active-standby high-availability configurations in which the Rational ClearCase server is active on only one cluster node at a time. The failover node has ClearCase installed but in a standby mode. In case of failure of the active node, Rational ClearCase is activated on the failover node to avoid downtime and business disruption. To support high-availability configurations, IBM Rational recommends moving VOB, View storage, and registry files on NAS storage.

Along with host-side clustering, a ClearCase environment also needs storage-side cluster failover capability for enhanced data availability and reliability. NetApp storage systems provide high-availability configurations that have the potential to deliver up to 99.999% proven uptime (see source cited in section 5.4). An active-active controller configuration provides fault tolerance by transferring the data service of an unavailable storage device to the other available storage device in the cluster. This can provide continuous data availability and access to VOBs, Views, and Registry data even when one storage system in the cluster undergoes a system failure or an upgrade. An active-active configuration also helps in load balancing of nodes.

You can further improve data availability and performance consistency of your VOBs and Views through the NetApp multipath HA storage configuration option, which provides multiple redundant paths to storage shelves. It avoids unnecessary failovers by providing controller-to-storage-shelf data path redundancy and helps in enhancing storage performance.

NetApp provides storage HA support to ClearCase deployments when using NFS, iSCSI, and FC protocols. However, it does not support CIFS protocol in the HA mode.

11 PERFORMANCE BENCHMARKS ACROSS PROTOCOLS

NetApp evaluated deployment of a ClearCase VOB server across all storage architectures (FC SAN, iSCSI, NAS, and split configuration) on both UNIX and Windows server environments to determine the best-fit deployment options.

This section provides you with the performance overview across NAS, FC SAN, iSCSI, or split configuration options to help you determine the best-fit solution for your business.

CLEARCASE PERFORMANCE BENCHMARKS ON UNIX SERVER

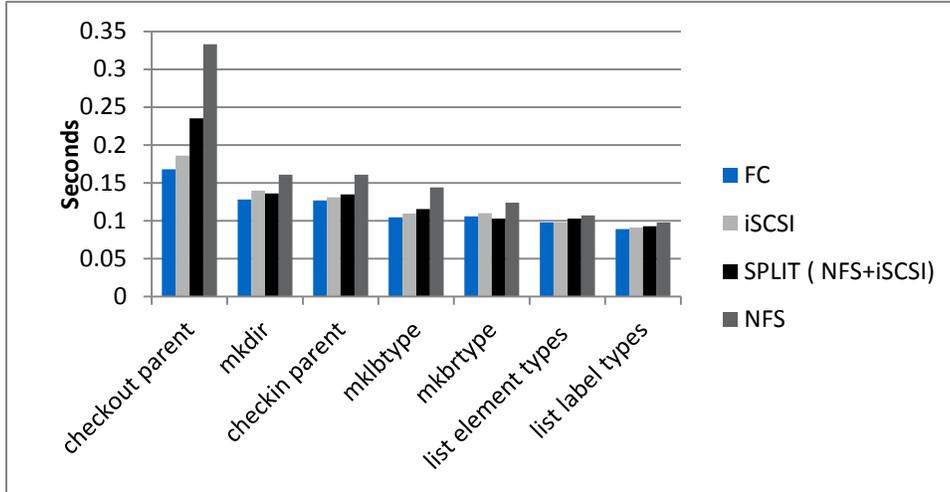


Figure 3) Performance across various protocols for directory and metadata manipulation.

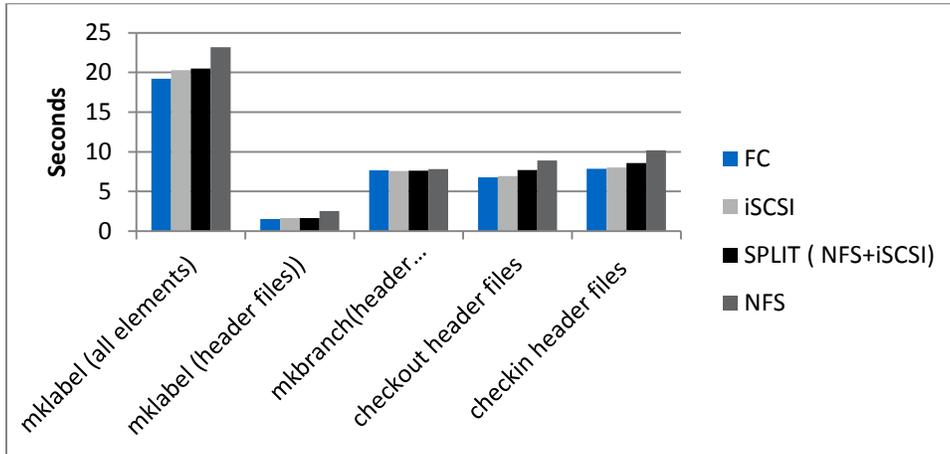


Figure 4) Performance across various protocols for file and metadata manipulation.

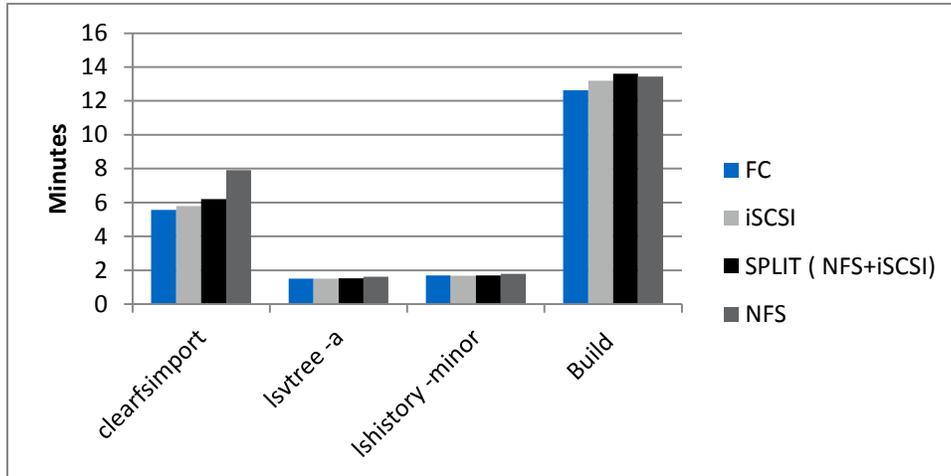


Figure 5) Performance across various protocols for read write test.

CLEARCASE PERFORMANCE BENCHMARKS ON WINDOWS SERVER

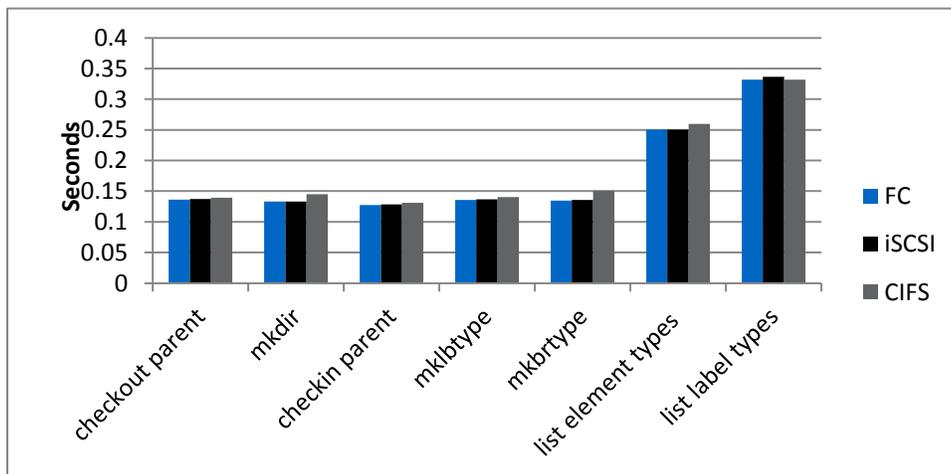


Figure 6) Performance across various protocols for directory and metadata manipulation.

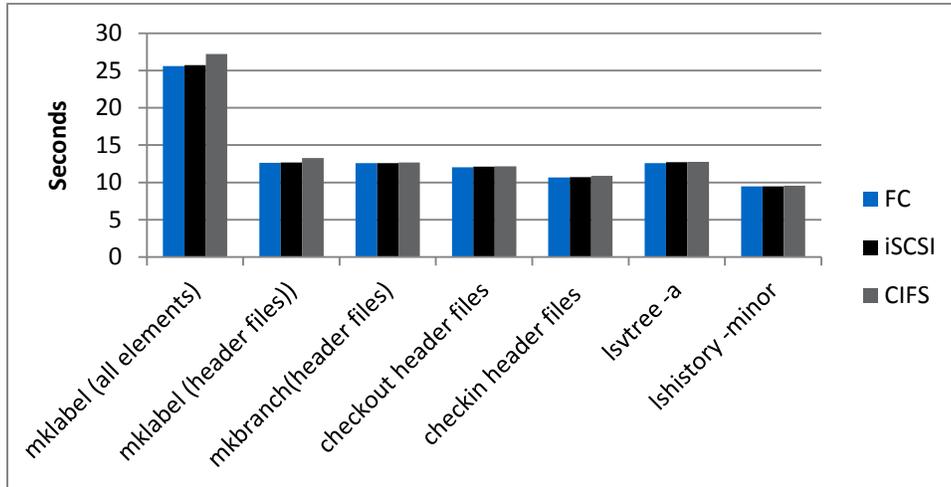


Figure 7) Performance across various protocols for file and metadata manipulation.

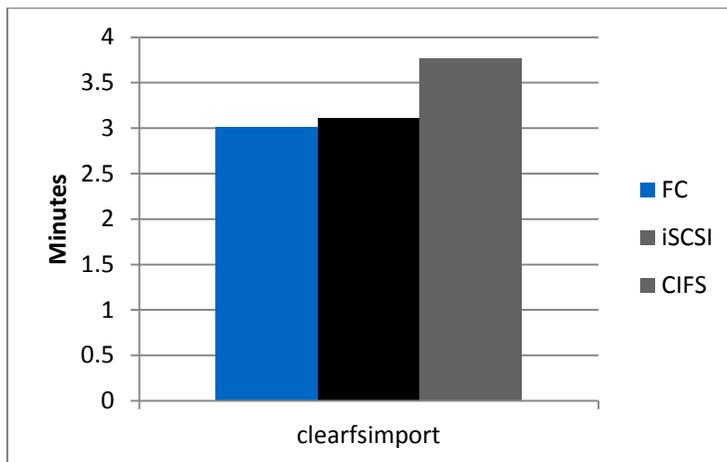


Figure 8) Performance across various protocols for read write test.

11.1 SUMMARY OF PERFORMANCE BENCHMARK RESULTS

A performance test was run that included various ClearCase commands for importing files (writing), listing metadata, metadata creation, applying metadata, file manipulation, and directory manipulation.

Based on the performance tests conducted, we arrived at the following conclusions:

11.1.1 KEY PERFORMANCE TAKEAWAYS

VOB on UNIX Server

- The performance of VOB is fastest on FC SAN, followed by iSCSI, split configuration, and NFS, in that order.
- On average there is only a 7.7% performance difference between split configuration and FC SAN, and a 16% performance difference between NFS and FC SAN.

VOB on Windows Server

- The performance of VOB is fastest on FC SAN, followed by iSCSI and CIFS, in that order.
- On average there is only a 1.4% performance difference between iSCSI and FC SAN, and an 18.6% performance difference between CIFS and FC SAN.

11.1.2 OTHER DECISION CRITERIA FOR DEPLOYMENT ARCHITECTURE

VOB on UNIX Server

- Split configuration provides the same multiprotocol advantage as NAS, enabling both Windows and UNIX/Linux clients to access files via NFS/CIFS without the need for Samba or TAS.
- The Interop environment (Windows + UNIX/Linux clients) is supported.

VOB on Windows Server

- Split configuration is supported.
- The Interop environment is not supported.
- CIFS-based storage cluster failover is not supported.

12 BEST-FIT DEPLOYMENT OPTIONS

12.1 UNIX VOB SERVER OPTIONS

VOB on UNIX server can be deployed on any of the following protocols: FC SAN, iSCSI, and NFS. ClearCase also supports a split configuration in which the VOB database resides on FC SAN / iSCSI and the source pool resides on NAS. Given below are the pros and cons of various connectivity options to help you determine the best-fit deployment option.

12.1.1 VOB DEPLOYMENT ON FC SAN

Pros:

- Performance of VOB is fastest on FC SAN compared to other deployment options.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- Does not provide native file access to heterogeneous clients. You require products like the CIFS emulator on the UNIX server to provide heterogeneous access to VOB files.
- Requires management of hubs and switches.
- Expensive when compared with other protocols because they require a dedicated fiber optic storage network.
- Backup and recovery are complex, requiring more steps compared to file-level protocols like NFS.

12.1.2 VOB DEPLOYMENT ON iSCSI

Pros:

- Read write performance is second only to FC SAN.
- It is highly cost efficient because of its Ethernet infrastructure.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- Does not provide native file access to heterogeneous files and requires emulation products like Samba and TAS.
- Is easier to manage compared to FC SAN, but NAS remains the easiest to set up and manage.
- Backup and recovery are complex, requiring more steps compared to file-level protocols like NFS.

12.1.3 VOB DEPLOYMENT ON NFS

Pros:

- It provides direct file access to heterogeneous clients and eliminates the need for TAS/Samba for Windows client access.
- It is inexpensive and simple to manage because it is based on Ethernet.

- It allows granularity in backup and recovery operations, requiring fewer steps.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- I/O performance is slow compared to other protocols.

Note: You still get reasonably fair performance that is adequate for most deployments.

12.1.4 SPLIT CONFIGURATION

In a split configuration, the VOB database is deployed on SAN, usually iSCSI, and the source pool is on NAS.

Pros:

- It can improve VOB server performance, speeding up disk I/O performance for metadata operations.
- It provides direct file access to heterogeneous clients, enabling both Windows and UNIX/Linux clients to access files via NFS/CIFS without the need for interops like Samba or TAS.
- They can provide excellent cost savings when deployed on iSCSI and NAS, with minimal impact to performance.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- Backup and recovery are a bit more complex since backups are performed using two separate protocols while maintaining concurrency between them.

12.2 WINDOWS VOB SERVER OPTIONS

VOB on a Windows server can be deployed on any of the following protocols: FC SAN, iSCSI, and CIFS. The split configuration is currently supported in this setup. For more information on split configuration for Windows, refer to the IBM Rational white paper titled [“Remote VOB Storage Pools on Windows.”](#) Users need to keep in mind that VOB on a Windows server does not support UNIX/Linux desktop clients. Given below are the pros and cons of various connectivity options to help you determine the best-fit deployment option.

12.2.1 VOB DEPLOYMENT ON FC SAN

Pros:

- Performance of VOB is fastest compared to other deployment options.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- It does not provide direct file access to Windows clients.
- It involves administrative overhead because it requires management of hubs and switches.
- It is expensive and complex because it involves dedicated hubs and switches.
- Backup and recovery are complex, requiring more steps, compared to file-level protocols like CIFS.

12.2.2 VOB DEPLOYMENT ON ISCSI

Pros:

- The read write performance is close to that obtained on FC SAN.
- It is highly cost efficient because of its Ethernet infrastructure.
- Storage can be configured as a high-availability cluster pair in a ClearCase environment.

Cons:

- It does not provide direct file access to Windows clients.
- Backup and recovery are complex, requiring more steps, compared to file-level protocols like CIFS.

- It requires deployment of a software initiator (when not using HBAs with TOE), which adds to the load on the CPU.

12.2.3 VOB DEPLOYMENT ON CIFS

Pros:

- It provides direct file access to Windows clients.
- It is cost effective and easy to manage.
- It provides granular backup and recovery.

Cons:

- I/O performance is slow compared to other protocols.
- Since CIFS is a stateful protocol, transparent storage cluster failover is not supported.

12.3 VIEW DEPLOYMENT ON NAS

Views provide users access to specific software releases for development and testing purposes. Views are usually deployed on NAS irrespective of where the VOB resides.

Pros:

- It is cost effective and easy to manage because it uses regular Ethernet.
- It reduces management overhead by providing backup from a single source. Private files can be recovered easily without requiring the backup administrator's assistance.

13 DECISION TREE

NetApp supports unified storage and provides different modes and protocols for deployment. To help you choose an appropriate mode and protocol for configuring your Rational ClearCase on NetApp storage, NetApp provides a decision tree that takes into account the primary and secondary key points that influence decision making.

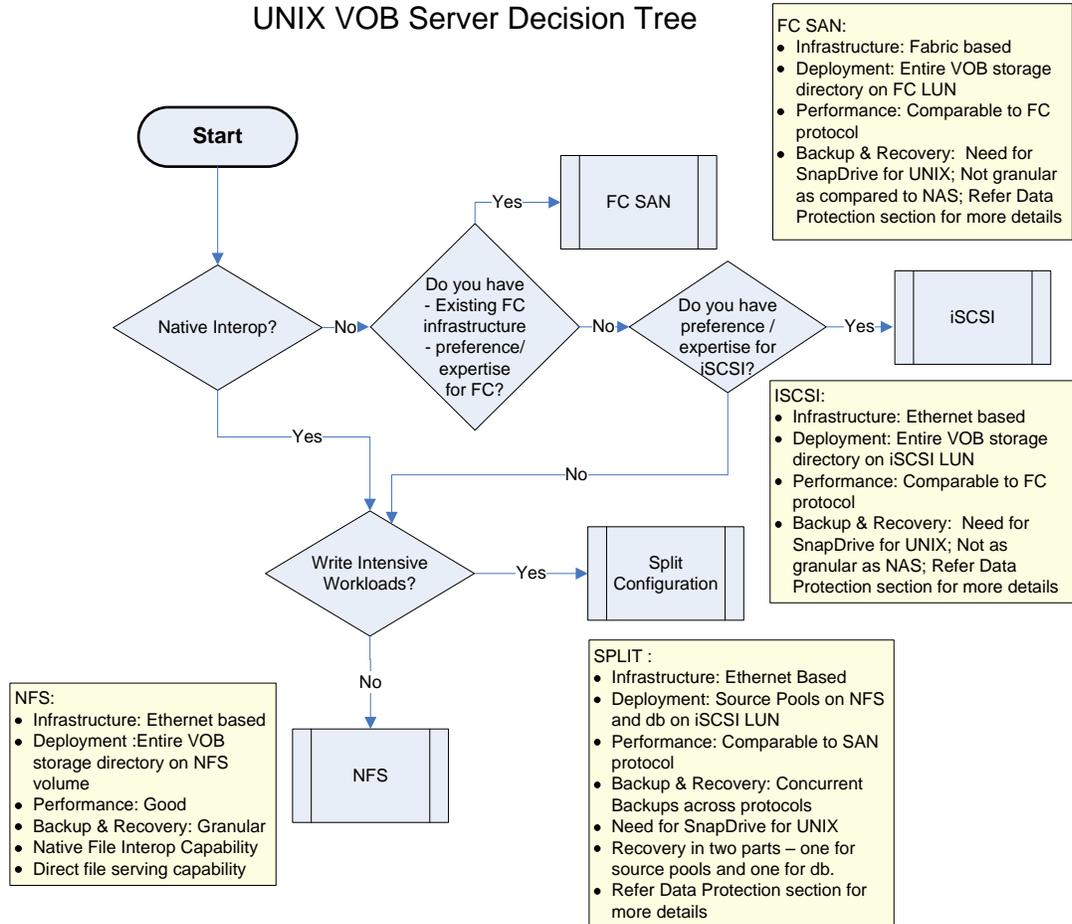
PRIMARY POINTS OF INFLUENCE

- Selected server platform. The server platform is short listed based on the project requirements and recommendations from IBM Rational.
- VOB storage/View storage.
- Support for transparent storage cluster failover.
- Native Interop support. Native Interop involves native serving of NFS and CIFS protocols by NetApp storage for Rational ClearCase dynamic views. This eliminates the need for emulation products (they can either be NFS clients on Windows or SMB server products for UNIX and Linux).

OTHER ASSOCIATED POINTS OF INFLUENCE

- Existing Infrastructure
- IT expertise and preference
- Performance
- Backup and recovery
- Other associated facts

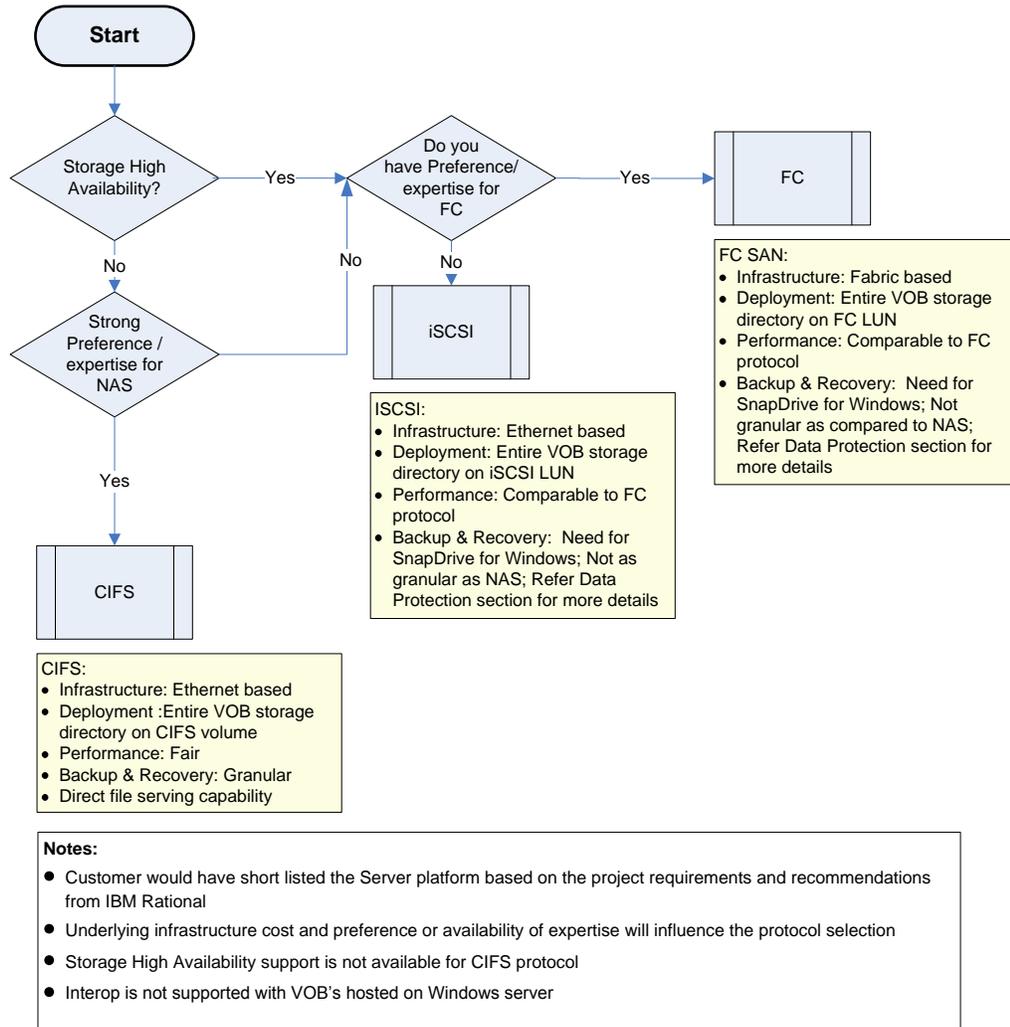
UNIX VOB Server Decision Tree



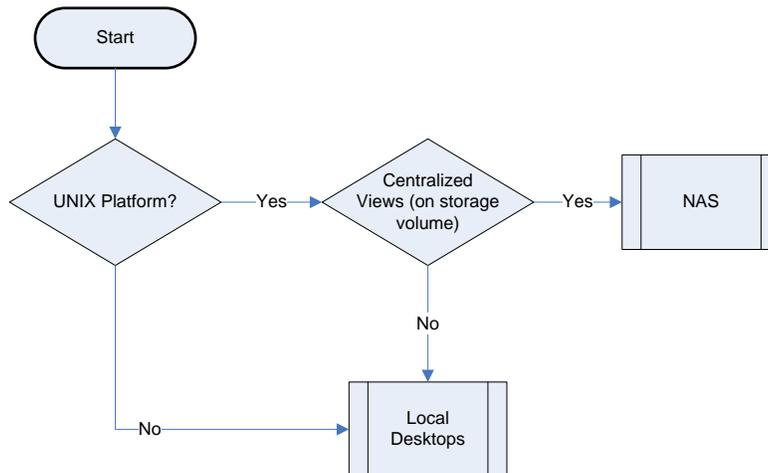
Notes:

- Customer would have short listed the Server platform based on the project requirements and recommendations from IBM Rational
- Underlying infrastructure cost and preference or availability of expertise will influence the protocol selection
- Storage High Availability support is available across all protocols for UNIX based VOB servers
- Native Interop involves native serving of CIFS protocol by NetApp storage for Rational ClearCase dynamic views. This will eliminate the need for emulation products(can either be NFS clients on Windows or SMB server products for UNIX and Linux)
- Write Intensive Workloads: (rare scenarios) For more details refer the section – 5.0 Configuring Rational ClearCase with NetApp storage)

Windows VOB Server Decision Tree



Views: Decision Tree



Notes:

- Centralized Views: View Backup and restoration possible
- Restoration of View Private files
- Degradation in performance seen with Centralized views on Windows platform, due to CIFS redirector issues – recommended to have local views

14 PERFORMANCE-RELATED ADD-ON OPTIONS

14.1 PERFORMANCE ACCELERATION MODULE (PAM)

Any kind of latency or performance degradation can have an adverse effect on the entire software development process. NetApp Performance Acceleration Module optimizes the performance of your storage system and provides intelligent read cache to improve throughput and help reduce latency, without requiring the addition of fast-spinning disk drives. It can improve response time of random read-intensive workloads and help reduce latency by a factor of 10 compared to disk. Since PAM provides an increased read cache, it serves most of the reads from the memory as opposed to the slow access from disks. This can reduce latency and eliminate the need to configure large numbers of disks, thereby decreasing your requirements for rack space, power, and cooling.

The benchmark tests with PAM in NFS environments yielded improved response times from 30% to 55% for metadata operations and by 47% for Clearmake operations, which are read intensive. These results were obtained by running the PAM in default mode where it caches both metadata and user data. The PAM used for the tests had 16GB of DRAM. In subsequent iterations of the test, data was read from the PAM module instead of the disks providing additional efficiency. It must be noted that actual performance improvements achieved with PAM depend on the working data set sizes as well as the amount of host server memory.

15 DEDUPLICATION IN RATIONAL CLEARCASE ENVIRONMENTS USING NETAPP STORAGE

NetApp deduplication is a core Data ONTAP technology that helps you to dramatically improve your storage utilization by eliminating duplicate blocks of data. In a software development environment, engineers working on the same project are likely to have the same set of files in their workspace, with each of them

modifying just a few of the files. You can consolidate all workspaces (ClearCase Views) on NetApp storage and leverage deduplication technology to help reduce the data footprint.

Deduplication is a supported option for Clearcase 7.1.1 and later. Please refer to the following support document.

<http://www-01.ibm.com/support/docview.wss?uid=swg21381946>

16 DATA PROTECTION OPTIONS

Comprehensive data protection is imperative to enable business continuity and high data availability. NetApp provides the following seamless and robust data protection solutions for the ClearCase environment:

- **Snapshot:** NetApp Snapshot technology helps you create frequent, quick, and reliable backups of data with minimal performance impact and consumption of storage space. You can copy important data frequently and can save up to 255 instantaneous, point-in-time versions of each data volume.
- **SnapRestore:** NetApp SnapRestore software can provide rapid restores and immediate online data recovery. Using a simple single command, you can recover entire volumes or individual files in seconds.

16.1 ONLINE BACKUP AND RECOVERY

The ClearCase environment must be backed up to safeguard its business-critical components, VOBs, and Views. NetApp Snapshot technology can provide quick and reliable backups and near-instantaneous online data recovery to help keep your business up and running.

Backup and recovery of VOBs and Views can be made simpler with NetApp. View private files and specific directories can be easily recovered from Snapshot copies, without requiring the storage administrator's intervention. To recover View private files, select the Snapshot copy from which you need to recover the files, drill down to the required View directory, list all the files in the View storage directory, and identify the View private file that you want to restore to the production area.

You can easily restore one or more VOBs from the latest Snapshot copy using simple commands. Interdependent VOBs should be grouped in the same volume or qtree and a group of VOBs that are connected by hyperlinks should be backed up and restored together. This section outlines the steps involved in taking backups and performing recovery of VOBs using NetApp Snapshot technology across different protocols. For the Split Configuration in which the VOB is on the SAN and the source pools are on a NAS, snap create via SnapDrive can be done both on the SAN and NAS on the same command line to maintain consistency. For example, in UNIX the command would be:

```
snapdrive snap create -fs /samba_vob /samba_src -snapname my_snap
```

Table 1) Backup and recovery steps for VOB across protocols.

	NFS and CIFS	Split Configuration	iSCSI and FC
BACKUP	<ul style="list-style-type: none"> • Lock all the VOBs • Snap create • Unlock all the VOBs 	<ul style="list-style-type: none"> • Lock all the VOBs • Snap create (done through SnapDrive to provide consistency at O/S level) <ul style="list-style-type: none"> ○ Create concurrent Snapshot copies for volumes containing source pools and db • Unlock all the VOBs 	<ul style="list-style-type: none"> • Lock all the VOBs • Snap create (done through SnapDrive to provide consistency at O/S level) • Unlock all the VOBs
RESTORE	<ul style="list-style-type: none"> • You must ensure that no server processes are 	<ul style="list-style-type: none"> • You must ensure that no server processes are alive 	<ul style="list-style-type: none"> • You must ensure that no server

	<p>alive for the said VOB(s)</p> <ul style="list-style-type: none"> • Unregister and untag the VOB (delisting from the registry) • Rename the corrupted VOB storage directory • Identify the latest Snapshot copy • Copy the VOB storage directory from the snapshot to the original location (use tools that preserve the file permissions) <ul style="list-style-type: none"> ○ Use snap-restore command if you need to restore the entire volume (very rare) • Register and tag the said VOB • Run dbcheck and checkvob utilities as prescribed by IBM Rational • Unlock the VOB 	<p>for the said VOB(s)</p> <ul style="list-style-type: none"> • Unregister and untag the VOB (delisting from the registry) • Rename the corrupted VOB storage directories containing pools and the db • Identify the Snapshot copy (both volumes) • Copy the VOB storage directory containing pools to the original location from the snapshot • For db snap-connect (using SnapDrive) to the relevant Snapshot copy of the LUN to create a LUN clone • Mount the LUN clone to a mount point • Copy the VOB storage directory containing db to the original location from the cloned LUN <ul style="list-style-type: none"> ○ Use snap-restore command if you need to restore the entire volume/LUN (very rare) • Register and tag the said VOB • Run dbcheck and checkvob utilities as prescribed by IBM Rational • Unlock the VOB 	<p>processes are alive for the said VOB</p> <ul style="list-style-type: none"> • Unregister and untag the VOB (delisting from the registry) • Rename the corrupted VOB storage directory • Identify the latest Snapshot copy • Connect disk (using SnapDrive), pointing to the relevant Snapshot copy of the LUN to create a LUN clone • Mount the LUN clone to a mount point • Copy the VOB storage directory to the original location (use tools that preserve the file permissions) from the cloned LUN <ul style="list-style-type: none"> ○ Use snap-restore command if you need to restore the entire volume/LUN (very rare) • Register and tag the said VOB • Run dbcheck and checkvob utilities as prescribed by IBM Rational • Unlock the VOB • Disconnect the disk
--	--	---	---

16.2 DISASTER RECOVERY

NetApp SnapMirror technology provides you with an affordable and reliable disaster recovery solution. It efficiently replicates data across global networks at high speeds, helping you create mirrored images of ClearCase VOBs on off-site networked storage systems. The mirrored volume is an exact copy of the source volume at the block level, and any changes made in the source volume are updated to the mirrored volume using SnapMirror. Updates are transferred asynchronously and you can schedule updates to the mirror or perform manual updates. In the event of a disaster, the mirrored volume can be turned into a production volume (it would involve some overheads from Rational ClearCase's perspective), helping you minimize your downtime costs. You can also set up a reverse SnapMirror relationship to replicate data back to the production system.

SnapMirror is block-based technology and is unidirectional. The data synchronization happens only between the source and the destination. On the other hand, Rational ClearCase MultiSite is file based; it creates a replica of the VOB (the database and the source files are copied onto the destination) at the remote site and the synchronization mechanism is bidirectional based on the epoch count.

Both SnapMirror and ClearCase MultiSite can be used for backing up ClearCase VOB data depending upon the appropriate use case.

16.3 DISK-TO-DISK BACKUP

A mirrored volume (SnapMirror destination) is an exact replica of the main volume and, if you delete the Snapshot copy from the production site, the mirrored volume Snapshot copy on the DR site also gets deleted. To archive several Snapshot copies dating back several days or months, you need to use SnapVault. SnapVault can be enabled on the volumes containing VOBs and Views. It leverages Snapshot copies to transfer data from the NetApp primary storage system to the NetApp secondary storage. Unlike SnapMirror, SnapVault secondary data is read only—you cannot fail over to a SnapVault secondary. When you need access to SnapVault secondary data, you can easily restore the data using a single command.

17 CONCLUSION

IBM Rational ClearCase is an enterprise-critical Software Configuration Management tool that is leveraged by developers across the globe. Developers require SCM systems to be available all the time because downtime adversely affects the development process, leading to potential revenue losses. A robust and resilient storage system such as NetApp's offers high-performance and high-availability options for geographically distributed product development environments. NetApp provides a unified and scalable storage system to help improve interoperability and eliminate the costs of multiple, separate storage environments. Further, data protection and replication technologies like Snapshot, SnapMirror, and SnapVault can help provide protection of your critical data assets, while shrinking your backup windows. Innovative technologies like PAM optimize SCM and build performance, and technologies like FlexClone can speed up the testing process to accelerate releases. With seamless deployment and a highly scalable and available software development infrastructure, NetApp, in collaboration with IBM Rational ClearCase, can accelerate the global development process while helping reduce overall costs.

18 ADDENDUM

18.1 TEST CONFIGURATION

a) Protocol Performance

UNIX:

Clients: SunFire V125, 2GB RAM; Solaris™ 10 update 3; IBM Rational ClearCase 7.1

VOB Server: SunFire V445, 4GB RAM; Solaris 10 update 3; IBM Rational ClearCase 7.1

Windows

Clients: IBM system x3250, 2GB RAM; Windows 2003 Enterprise Edition; IBM Rational ClearCase 7.1

VOB Server: IBM system x3550, 2GB RAM; Windows 2003 Enterprise Edition; IBM Rational ClearCase 7.1

NetApp Storage: FAS3170 with Data ONTAP 7.3

For NFS protocol, Ethernet connectivity was used with the NetApp storage and the VOB server.

For iSCSI protocol, native s/w initiators were used for both UNIX and Windows.

For FC protocol, HBA cards (Emulex LightPulse LP11002 and QLE 2432) for UNIX and Windows, respectively, were configured on the VOB server to establish direct FC connection to the NetApp storage.

NetApp Storage

b) PAM Performance

Clients: SunFire V125, 2GB RAM; Solaris 10 update 3; IBM Rational ClearCase 7.1

VOB Server: SunFire V445, 2GB RAM; Solaris 10 update 3; IBM Rational ClearCase 7.1

NetApp Storage: FAS3170 with Data ONTAP 7.3

PAM Module: 16GB

User Workload: 700 concurrent ClearCase users

c) Test Files

Samba source code was used as the files for these tests. There were 2,241 files in 178 directories. All of them imported into the VOB and tested across different protocols.

18.2 CLEARCASE SUPPORTED RELEASES

The IBM website has the information on supported (certified) Network Appliance devices for Clearcase

For Clearcase 7.1:

<http://www-01.ibm.com/support/docview.wss?rs=984&uid=swg21302974#Supported%20NAS%20Products>

For Clearcase 7.0.x.x:

http://www-01.ibm.com/support/docview.wss?rs=0&context=SW750&q1=NetApp&q2=support+matrix+netapp&uid=swg21239315&loc=en_US&cs=utf-8&cc=us&lang=en#Supported%20NAS%20Products

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.



© 2012, NetApp Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. Results may vary and depend on customers' systems configurations and specifications. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexCache, FlexClone, FlexVol, NearStore, NOW, RAID-DP, SnapDrive, SnapMirror, SnapRestore, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Windows is a registered trademark of Microsoft Corporation. UNIX is a registered trademark of The Open Group. TR-3774