



Technical Report

Deployment and Implementation Guide: IBM Rational ClearCase on NetApp Data ONTAP Operating in Cluster-Mode

Agnes Jacob, NetApp and Michael Donati, IBM Rational

January 2012 | TR-4032

ABSTRACT

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

TABLE OF CONTENTS

- 1 INTRODUCTION 4**
 - 1.1 SCOPE4
- 2 OVERVIEW OF IBM RATIONAL CLEARCASE ARCHITECTURE..... 4**
 - 2.1 VERSIONED OBJECT BASE SERVER AND STORAGE.....5
 - 2.2 VIEW SERVER AND STORAGE6
 - 2.3 REGISTRY SERVER7
 - 2.4 LICENSE SERVER7
 - 2.5 CHANGE MANAGEMENT SERVER (OPTIONAL)7
 - 2.6 DATA FLOW.....7
- 3 VALUE OF NETAPP CLUSTER MODE 8**
 - 3.1 ACCELERATE DEVELOPMENT8
 - 3.2 SCALE INFRASTRUCTURE, NOT OPEX9
 - 3.3 REDUCE RISKS9
- 4 PERFORMANCE OF CLEARCASE ON CLUSTER-MODE 10**
 - 4.1 TEST ENVIRONMENT.....10
 - 4.2 TESTS11
 - 4.3 RESULTS12
 - 4.4 ANALYSIS.....17
- 5 DEPLOYMENT OPTIONS 17**
 - 5.1 INTEROP VOB DEPLOYMENT18
 - 5.2 WINDOWS VOB DEPLOYMENT19
 - 5.3 VIEW DEPLOYMENT19
- 6 IMPLEMENTATION PLANNING 19**
- 7 STORAGE EFFICIENCY 20**
 - 7.1 DEDUPLICATION AND COMPRESSION.....21
 - 7.2 NETAPP FLEXCLONE21
- 8 BEST PRACTICES 22**
- 9 CONFIGURATION AND SETUP 22**
 - 9.1 AGGREGATE CREATE.....23
 - 9.2 VSERVER CREATE23
 - 9.3 VOLUME CREATE24
 - 9.4 NFS ACCESS25
 - 9.5 CIFS ACCESS26
 - 9.6 INTEROP ACCESS (BOTH CIFS AND NFS ACCESS).....27
 - 9.7 ISCSI ACCESS.....27

9.8	FIBRE CHANNEL ACCESS	38
9.9	MAKE VOB OR VIEW.....	43
10	SUPPORT MATRIX.....	43
11	LIMITATIONS.....	44
12	TROUBLESHOOTING	44
12.1	CLEARCASE LOGS TO CHECK.....	44
12.2	PERMISSIONS ISSUES	44
12.3	PERFORMANCE ISSUES	45
13	REFERENCES	45
13.1	TECHNICAL REPORTS AND SOLUTION BRIEF	45
13.2	NETAPP UNIVERSITY WEB-BASED COURSES	45

LIST OF TABLES

Table 1)	ClearCase tests.	11
Table 2)	Interop: Total times for each test	17
Table 3)	Windows: Total time results.....	17
Table 4)	Pros and cons of VOB deployment on a UNIX server.....	18
Table 5)	Pros and cons of VOB deployment on Windows.....	19
Table 6)	Selection criteria across NetApp storage protocols.....	20

LIST OF FIGURES

Figure 1)	Typical ClearCase deployment.	5
Figure 2)	Data flow.....	7
Figure 3)	Interop: Single-user base ClearCase results.	12
Figure 4)	Interop: Single-user base ClearCase results. (Continued)	13
Figure 5)	Interop: Single-user UCM results.	13
Figure 6)	Interop: Single-user UCM results. (Continued)	14
Figure 7)	Interop: 100-user Web view results (Linux CM server).	14
Figure 8)	Interop: 100-users Web view results (Linux CM server). (Continued)	15
Figure 9)	Interop build results.	15
Figure 10)	Windows:100-user Web view results (Windows CM Server, Windows VOB).....	16
Figure 11)	Windows :100-user Web view results (Windows CM Server, Windows VOB). (Continued)	16
Figure 12)	Deduplication of ClearCase repositories.	21

1 INTRODUCTION

IBM Rational ClearCase is a leading software control management application. It is widely deployed by engineering and development teams in midsize and large enterprises to manage and track different revisions of software components as they are being developed. Key features of Rational ClearCase include version control, workspace and build management, and process configurability.

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

Cluster-Mode storage with the IBM Rational ClearCase solution provides a scalable storage infrastructure and practices that allow engineering and development teams to minimize or eliminate data loss, accelerate recovery, enable collaboration among distributed design groups, and simplify data management to for faster development cycles, improved engineering productivity, and reduced time to market.

IBM Rational has been a strong NetApp partner since 1998. We continue to work together in certifications and support of our solutions, performance testing, and analysis, and we collaborate in offering our joint customers a development platform that they can trust.

1.1 SCOPE

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

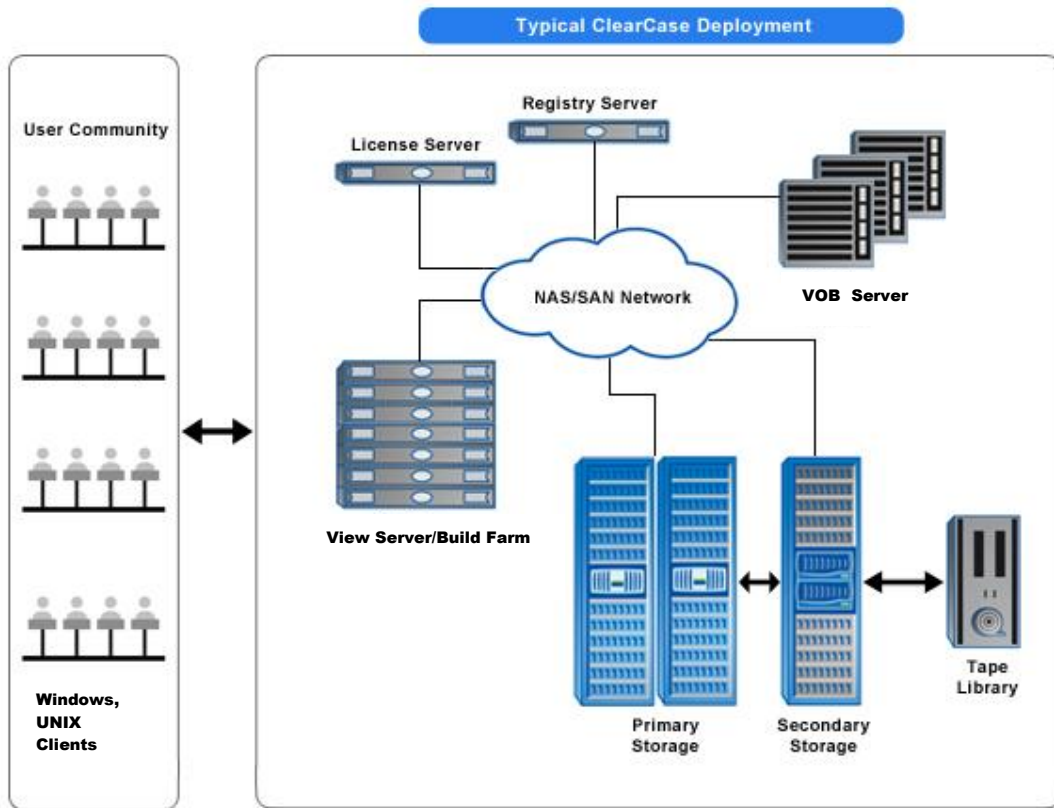
2 OVERVIEW OF IBM RATIONAL CLEARCASE ARCHITECTURE

IBM Rational ClearCase is a distributed and scalable application with a client and server architecture. It is composed of ClearCase hosts, clients, and multiple ClearCase applications services to manage the use of the main data repositories and developer workspaces. These application services, which include a License server, Registry server, VOB server, View server, and Web server, can be installed on a single system or distributed on separate systems to balance workload across different systems.

This section is a high-level description and explanation of the uses of each of these components. For detailed information on ClearCase, refer to "[Change and Configuration Management](#)" on the IBM Rational Web site.

Figure 1 shows a typical ClearCase environment with clients, servers, and storage.

Figure 1) Typical ClearCase deployment.



2.1 VERSIONED OBJECT BASE SERVER AND STORAGE

The main component of Rational ClearCase is the Versioned Object Base (VOB) server, which manages all the software that is being developed. The ClearCase VOB server processes include:

- **albd_server.** Long-lived process that runs on every ClearCase host and spawns other processes below as needed or required. The `albd_server` process assists in connecting the client requests to ClearCase servers via RPC calls (listens for RPC requests on port 371).
- **admin_server.** Short-lived process that is invoked by the `albd_server` process to perform administrative support functions.
- **vobrpc_server.** Handles requests, both metadata and system data, from the `view_server` process.
- **db_server.** Responsible for VOB database transactions (typically metadata related) on the host in response to requests from clients.
- **lockmgr.** Manages the locking activities on a per-VOB basis.

The VOB server is responsible for providing the security of the software development *repository*, referred solely as the VOB, by controlling the data that goes in or out of the repository. The VOB consists of a specialized database and three storage pools:

- The **database** is an embedded Raima database that contains all metadata associated with directories and files under version control. These include version labels, hyperlinks, and configuration records.

This is generally the smallest component of the VOB storage and it usually grows at a small, constant pace over time.

- The **source** pool contains all the files that are under ClearCase control. There is a container file in the source pool for every text file. Source pool storage usually grows depending on the rate at which the project grows.
- The **cleartext** pool contains temporary storage for file versions that are currently in use. This pool acts as a cache for specific versions of files in the VOB that are currently in use by developers. Disk I/O is the heaviest on average for this pool. This area is periodically scrubbed of older versions that are not being accessed by any view work area. Thus the storage for the cleartext pool can fluctuate over time depending on the different revisions being worked on. This area grows at a faster rate than the source pool.
- The **derived** object pool contains the objects created as a result of the Clearmake build utility. This allows developers to reuse (wink in) the executables built by others in their workspaces, instead of rebuilding objects. The storage area for the derived pool can also fluctuate dramatically, depending on the frequencies of builds or different releases, versions being worked on, and the size of the binaries involved. The derived pools are also periodically scrubbed of objects that are no longer referenced by any view.

Users access the VOB data only through ClearCase commands. Users create a view of the VOB data, providing a context by which the data is accessed. There can be one VOB per project, or multiple projects can be consolidated in a single VOB.

2.2 VIEW SERVER AND STORAGE

A *view* is a developer workspace that is used to access the VOB. The View server runs either locally on the client hosts or on a centralized location. It manages and populates the user workspace with specific versions based on the view config spec rules. The view service that runs on a client host is the `view_server` process, a long-lived process that manages the activity of a particular view, reads the config spec rules, and tracks the modifications of the view's private files. The View storage is hosted on the View server host that runs the `view_server` process. There are three components of the view:

- **View-private storage** contains versions of files that are currently checked out in the view, unshared derived objects, and files that are not under version control. Disk I/O can be heavy in this area during builds. The view-private storage can grow with the number and size of files checked out, the number and size of private files, and the number and size of derived objects created during builds.
- **Database** is a Raima database that contains information related to the view state. This area is generally very small or negligible compared to the view-private storage.
- **Configuration specification** contains rules and information related to the versions of the files that developers see or access in their workspace.

For flexibility, Rational ClearCase offers three types of views:

- **Snapshot views** allow users to work off line by loading copies of file elements from the VOB onto the local system. Manual updates are required to see file updates in a VOB. Snapshot views use ClearCase File Server (CCFS), which is a TCP/IP-based mechanism to transfer files between VOB servers and clients.
- **Dynamic views** provide real-time access to the VOB data. Dynamic views use Multi-Versioned File System (MVFS), which gives users quick, transparent access to ClearCase data, fast immediate updates, global view of code base, derived object sharing, and audited builds.

- **Web views** provide remote access to the VOB data via the ClearCase Remote Client.

2.3 REGISTRY SERVER

The Registry server is a centralized “yellow pages” of information about shared resources such as VOBs and views available across the network. It stores the details of the corresponding path and tags of each of the VOBs and views registered on the network.

2.4 LICENSE SERVER

The License server grants users a license to use cleartool and multitool commands. The licenses are bound to license server hosts and thus the host ID is given when requesting a set of licenses. Two types of licensing are available for Rational ClearCase: Atria Licensing and Flexlm Licensing. For more information, see the [IBM Rational Licensing](#) Web site.

2.5 CHANGE MANAGEMENT SERVER (OPTIONAL)

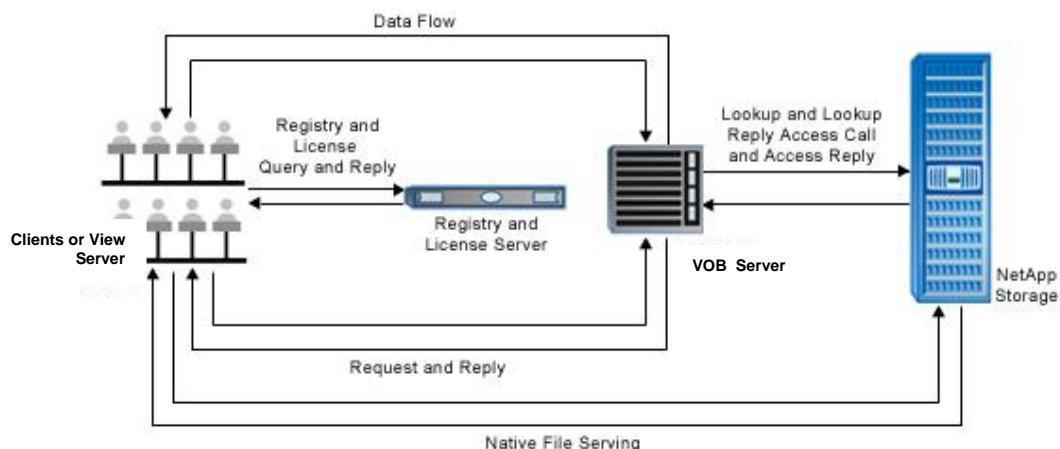
The ClearCase Change Management server (an IBM WebSphere® Application Server that is bundled with ClearCase) provides application support for the ClearCase Web interface, ClearCase Remote Client (CCRC). This allows ClearCase users to access ClearCase data via CCRC, typically over a wide area network.

2.6 DATA FLOW

Rational ClearCase uses the Remote Procedure Call (RPC) protocol for communication between the clients and VOB and View servers, and it also uses NFS and/or CIFS to access the files. The RPC communication in ClearCase should not be confused with the RPC used in the NFS protocol. The ClearCase RPCs do not interact with or invoke NFS RPCs. Figure 2 shows a view of the data flow between client or View servers, VOB servers, and storage. When a ClearCase user queries or accesses the VOB server via an RPC request, the flow is as follows:

1. The request first goes to the registry server to locate the VOB server and license servers to validate the user license.
2. The VOB server checks the VOB database and passes that information back to the View server with the container location (a filehandle pointing to the right location on the NetApp system that holds the VOBs).
3. Depending on the client request, a LOOKUP/ACCESS/GETATTR/WRITE is served natively to the client, if using NAS; or it is served via the VOB server if using the SAN protocol.

Figure 2) Data flow.



3 VALUE OF NETAPP CLUSTER MODE

NetApp Data ONTAP operating in Cluster-Mode enables software development organizations to couple multiple storage systems together into a single integrated cluster that can support part of or the entire development infrastructure. The cluster is managed as a single entity, and data can be easily moved between different storage systems transparently to both ClearCase and software developers. The Cluster-Mode solution offers the following key advantages for IBM Rational ClearCase customers:

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

3.1 ACCELERATE DEVELOPMENT

In development environments, project requirements can change rapidly and new projects can increase data and storage requirements. The ability to quickly provision or move data nondisruptively to accommodate these changes is important. Cluster-Mode features such as global namespace, single management system, and nondisruptive volume move can assist in facilitating these requirements.

DYNAMICALLY REBALANCE DEVELOPMENT WORKLOADS

With traditional storage systems, if a ClearCase VOB requires additional performance, other projects must be moved off of that storage system, or the VOB itself must be moved to a system that has higher performance. In either case, the result is an interruption in data availability and new pathnames for the affected projects.

With Cluster-Mode, multiple storage systems are integrated into a single cluster under a global namespace, and projects can be moved between systems transparently to both end users and the ClearCase application. With this powerful capability, projects that affect the performance of a “hot” VOB can be transparently moved to other systems, or the VOB itself can be moved to another system. In either case, the load is dynamically rebalanced without disrupting either developers or the ClearCase application, optimizing performance for the VOB and helping to accelerate project development.

DYNAMICALLY RESOLVE NETWORK CONGESTION

Data ONTAP operating in Cluster-Mode systems also provide the ability to dynamically resolve network congestion issues. With traditional systems, if one of the physical ports to a storage system is congested, the only recourse is to physically move developers to an alternate port. This approach unfortunately disrupts developer access.

With Cluster-Mode systems, administrators can simply move the logical interface (LIF) from the congested physical port to another LIF on a less congested port. This action is transparent to developers and to ClearCase, again helping to maximize developer productivity and accelerate the development process.

ALWAYS-ON INFRASTRUCTURE

Data ONTAP operating in Cluster-Mode enables an always-on storage infrastructure that supports development work 24x365. For example, when a storage system for ClearCase is upgraded, serviced, or retired, the projects on that system can be temporarily moved to other nodes in the system and then rebalanced after completion of the service or upgrade event, without disruption to users.

CLONES FOR EVERY DEVELOPER

NetApp FlexClone® thin-cloning technology enables instant, space-efficient clones of production or test data. These clones provide substantial storage savings because they consume space only for blocks that have been changed. With FlexClone, each developer can have their own virtual copy of production or test data for use in their test and development activities.

INSTANT BACKUPS

Traditional approaches to backups of ClearCase data can take substantial time; in the case of tape-based backups, VOBs may be unavailable for hours. Such lengthy disruptions can significantly affect the development process. With NetApp Snapshot™ technology, backups can be done in minutes, eliminating long backup windows.

3.2 SCALE INFRASTRUCTURE, NOT OPEX

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

A SINGLE POINT OF MANAGEMENT

A cluster of NetApp systems is managed as an integrated entity, not as a set of independent storage systems. The global management interface enables administrators to move projects between systems and to manage all ClearCase storage in the cluster from any node. With this approach, the incremental management burden remains relatively flat as additional systems are added to the cluster.

TRANSPARENT DATA MOVEMENT SIMPLIFIES MANAGEMENT

With traditional systems, moving data from one system to another is management intensive. Users need to be informed of the upcoming interruption, administrators need to copy the data, and users must then to be informed of the new pathnames. There is also a risk of error throughout this process.

With Cluster-Mode, there's no interruption to manage, no user communications are required, and no changes to developer processes are required.

OPTIMIZE COSTS WITH TIERED STORAGE

A NetApp cluster can include nodes of varying levels of performance, as well as different types of disks (FC, SATA, and SAS) with varying cost, capacity, and performance characteristics. With Cluster-Mode, old ClearCase projects can be transparently moved to lower tiers of storage, optimizing costs without affecting developer access.

UNIFIED STORAGE WITH NATIVE PROTOCOLS

Each system in a NetApp cluster supports native file and SAN protocols, including NFS, CIFS, FC, iSCSI, FCoE, and object protocols. This approach eliminates the server overhead associated with protocol emulators for NFS or CIFS, and it also eliminates the management complexity associated with deployment and maintenance of these protocol emulators. NetApp is the only vendor in the industry to offer unified storage at scale.

ENTIRE DEVELOPMENT INFRASTRUCTURE ON A SINGLE CLUSTER

With unified storage across a broad range of controllers and storage types, a single integrated NetApp cluster can effectively support the storage infrastructure for ClearCase, and it can also build farm software and other downstream workflows. This approach significantly simplifies management and reduces operating costs.

3.3 REDUCE RISKS

Your software is important intellectual property, and Data ONTAP operating in Cluster-Mode systems provide powerful data protection capabilities and new levels of flexibility to adapt to change.

EASILY SCALE

Whether you have a small development shop that may experience rapid growth, or a huge shop with continuing growth, a NetApp cluster makes it easy to grow seamlessly as your development efforts expand.

FREQUENT BACKUPS

With NetApp Snapshot technology enabling quick backups, it is convenient to make frequent backups, protecting more of your developers' work and enabling much more aggressive recovery point objectives.

SIMPLIFIED DISASTER PROTECTION

With traditional storage infrastructures, the complexity of the environment and the available data mirroring solutions make disaster recovery (DR) solutions difficult or impractical to implement.

NetApp SnapMirror® replication technology makes DR practical via an easy-to-implement, robust mirroring solution. Implementing DR can significantly reduce risks for your organization and protect your data in case of a catastrophic event.

4 PERFORMANCE OF CLEARCASE ON CLUSTER-MODE

In ClearCase terms, there are two environment modes:

- **Interop**, which has a UNIX® VOB server and can have UNIX and/or Microsoft® Windows® clients
- **Windows-only**, which has a Windows VOB server and Windows clients

ClearCase VOB performance was conducted across Cluster-Mode storage protocols (NFS, CIFS, FC, and iSCSI) in both Interop and Windows environment modes. This section describes the details of the ClearCase VOB performance tests, results, and analysis. All of the performance data was measured in a controlled environment. Therefore the results obtained in other operating environments may vary significantly. You should verify the applicable data for your specific environment.

4.1 TEST ENVIRONMENT

The test environment for this performance study consisted of the following hardware and software components.

HARDWARE

- ClearCase components
 - Interop environment
 - 1 to 5 UNIX or Window clients
 - 1 VOB server (Solaris 10 update 9)
 - 1 Linux® CM Server
 - Windows environment
 - 1 or 2 Windows clients
 - 1 VOB server (Microsoft Windows Server® 2008 R2)
 - 1 Windows CM server
 - License or Registry server
- NetApp components
 - 2-node FAS 3240
 - Cisco Nexus® 5010 (Cluster Interconnect)
 - 2 DS4243 600GB SAS drives

- Aggregate composed of 9 disks
- Network infrastructure
 - Ethernet: 1GbE for NFS and software iSCSI
 - FC Switch: 4GbE for FC SAN

SOFTWARE

- ClearCase components
 - ClearCase 7.1.2
 - ClearCase test scripts (Perl)
- NetApp components
 - Data ONTAP 8.1 operating in Cluster-Mode (NFS, iSCSI, FC, CIFS, SnapRestore®, SnapMirror, FlexClone, deduplication)

4.2 TESTS

Four types of tests were conducted to understand ClearCase VOB performance across NetApp Cluster-Mode storage protocols in an Interop environment. The tests were single-user base ClearCase; single-user UCM; 100-user Web views (change management server test); and build tests. In the Windows environment, we ran only the 100-user Web views test. Table 1 outlines the types of ClearCase tests conducted and their operations. The dataset used for these tests included 100 VOBs, totaling about 500GB of data. All tests were done using local views, except for builds, in which the views were on NAS storage. During these tests, response times were measured for ClearCase operations.

Table 1) ClearCase tests.

Test	ClearCase Operations	Client
Single-user base ClearCase	<ul style="list-style-type: none"> • Mkdirtype, Mkdirtype, Mkdirlabel, Mkdirlabel-rep • FindMerge, Merge • Checkin, Checkout, Uncheckout • Clearfsimport • Lsco, Lsprivate • Lsbrtype, Lslbtype, Lshistory, Lsvtree • Mkdirbtype-local, Mkdirlabel local 	<ul style="list-style-type: none"> • 1 Linux client for Interop
Single-user UCM	<ul style="list-style-type: none"> • Mkstream • Deliver, Rebase • Mkbl, Diffbl • Mkact, Setact, Lsact • Lstream, Lsbl 	<ul style="list-style-type: none"> • 1 Linux client for Interop
100-user Web views (7-hour test with users running 15 transactions per hour – 60ms/20Mb network)	<ul style="list-style-type: none"> ▪ IsVob, IsHistory, IsAct, IsView ▪ MakeActivity, SetActivity ▪ MakeFolder, CheckOutFolder ▪ Makefile ▪ CheckInFile, CheckOutFile, checkinFolder, UnDoCheckoutFile ▪ CCRCLoad ▪ MkStream ▪ UpdateView 	<ul style="list-style-type: none"> • 1 Linux CM server (uses RPT as test driver); IBM WebSphere app (CCRC) • 1 Windows CM server

Test	ClearCase Operations	Client
Multiplatform build (12-hour C++ build)	<ul style="list-style-type: none"> Build_lib (clearmake) 	<ul style="list-style-type: none"> 5 UNIX (Solaris Sparc, Solaris x86, HP, AIX, Linux) build servers and 1 Windows build server

4.3 RESULTS

This section provides the graphical results from the VOB performance tests conducted across NetApp Data ONTAP operating in Cluster-Mode storage protocols in both Interop and Windows environments. For Interop, in which the VOB server is running on a UNIX server, the storage protocols tested were NFS v3, iSCSI, FC, and hybrid (database on iSCSI and pools on NAS). For the Windows environment in which the VOB server is running on a Windows server, the storage protocols tested were CIFS, FC, and iSCSI.

INTEROP ENVIRONMENT RESULTS

The following figures illustrate the performance of ClearCase operations on VOB storage in an Interop environment across Cluster-Mode storage protocols.

Figure 3) Interop: Single-user base ClearCase results.




Figure 4) Interop: Single-user base ClearCase results. (Continued)




Figure 5) Interop: Single-user UCM results.

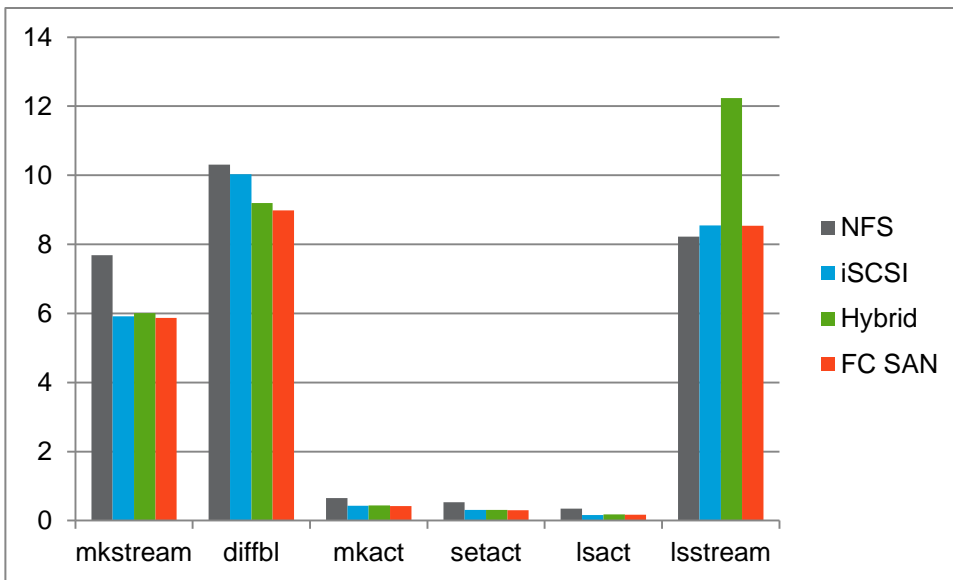


Figure 6) Interop: Single-user UCM results. (Continued)




Figure 7) Interop: 100-user Web view results (Linux CM server).




Figure 8) Interop: 100-users Web view results (Linux CM server). (Continued)

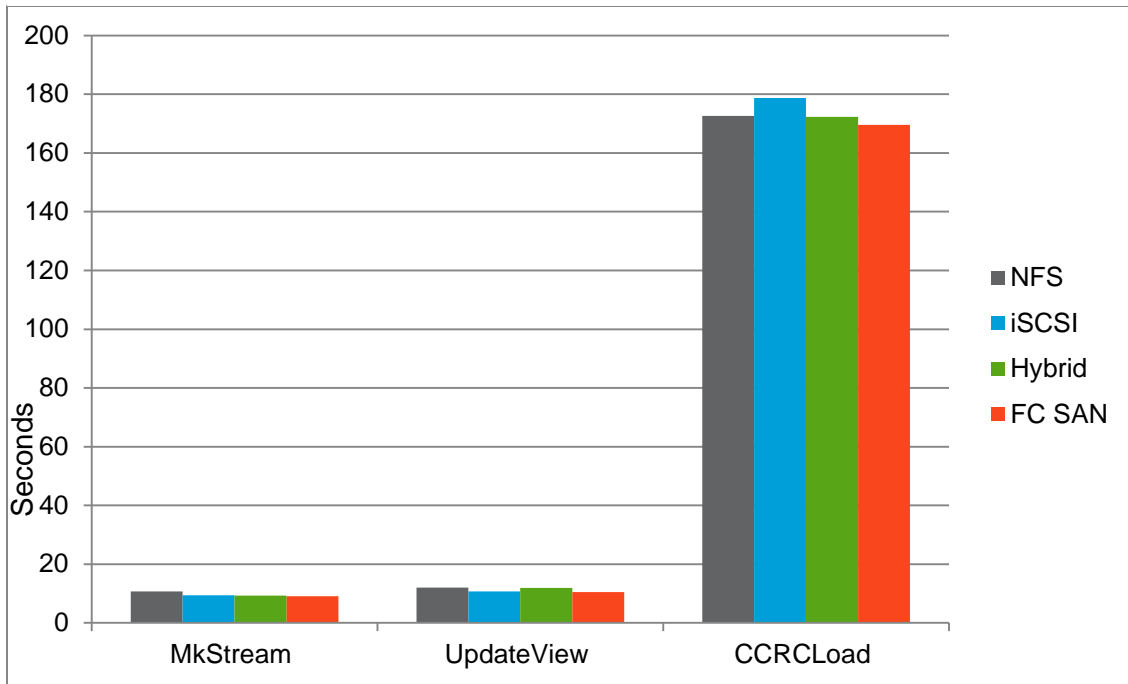
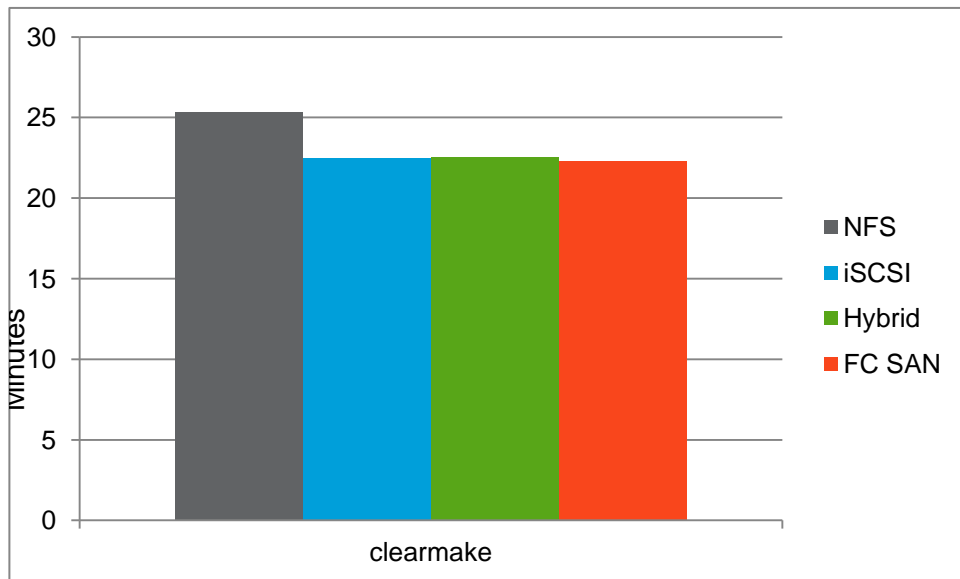


Figure 9) Interop build results.



WINDOWS ENVIRONMENT RESULTS

The following figures illustrate the performance of ClearCase operations on VOB storage in a Windows environment across Cluster-Mode storage protocols that include CIFS, iSCSI, and FC SAN.

Figure 10) Windows: 100-user Web view results (Windows CM Server, Windows VOB).


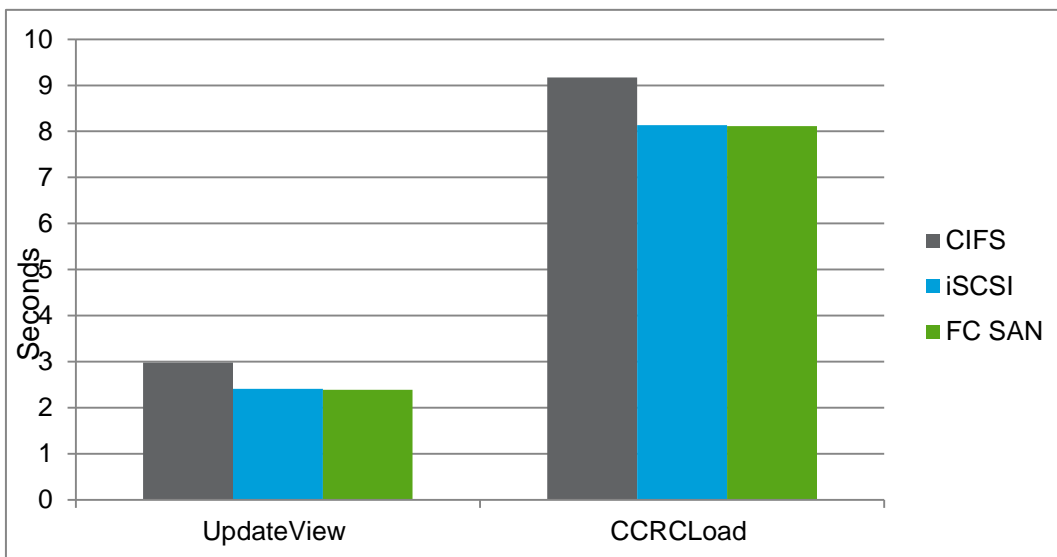


Figure 11) Windows :100-user Web view results (Windows CM Server, Windows VOB). (Continued)



4.4 ANALYSIS

This section summarizes the key conclusions derived from the performance tests in both Interop and Windows ClearCase environments.

INTEROP

- VOB performance is fastest on FC SAN, followed closely by iSCSI and the hybrid configuration and then NFS.
- The total time difference between VOB on NFS and FC SAN was 6% for the single user base ClearCase test; 8% for single user UCM; 8% for 100 Web users; and 12% for builds.
- The average difference between FC SAN and the hybrid configuration was 1% for the Single User Base ClearCase test; 2.8% for Single User UCM; 3.39% for 100 Web User; and 1% for Builds.

Table 2) Interop: Total times for each test.

Note: These numbers were used to calculate percentage differences.

Protocol	Single-User ClearCase Base (seconds)	Single-User UCM (seconds)	100 Web Users (seconds)	Build (minutes)
FC SAN	589.96	1642.83	36.67	22.32
iSCSI	608.54	1681.59	38.91	22.49
NFS	627.74	1793.48	39.86	25.36
Hybrid configuration (database on iSCSI and pools on NAS).	586.04	1690.68	37.96	22.54

WINDOWS

- VOB performance is fastest on FC SAN, followed by iSCSI and then CIFS.
- The total time difference between VOB on FC SAN and CIFS was 11% for 100-user Web tests.

Table 3) Windows: Total time results.

Protocol	100 Web Users (seconds)
FC SAN	29.22
iSCSI	29.32
CIFS	33.05

5 DEPLOYMENT OPTIONS

ClearCase VOB, View, and Registry repositories can be deployed on NetApp Data ONTAP operating in Cluster Mode storage for enhanced scalability and flexibility and simplified administration. Using Cluster-Mode, ClearCase VOBs and views can be deployed on any NetApp storage protocol: FC SAN, iSCSI,

NFS, and/or CIFS. As mentioned previously, there are two environment modes for ClearCase: Interop and Windows environments. This section examines the VOB deployment in each of these environments and discusses the advantages and disadvantages of each deployment option.

5.1 INTEROP VOB DEPLOYMENT

A VOB server that runs on a UNIX server can be deployed on the following NetApp storage protocols: FC SAN, iSCSI, NFS, and a hybrid configuration in which the VOB database resides on a SAN (iSCSI or FC) and pools reside on NAS (NFS). To set up a hybrid configuration, refer to the IBM technote [“Properly Using SAN storage for VOB database access while using a NAS filer for ClearCase Client Access.”](#)

Table 4 highlights the advantages and disadvantages of deploying VOB on each of NetApp’s storage protocols.

Table 4) Pros and cons of VOB deployment on a UNIX server.

Deployment	Pros	Cons
FC SAN	<ul style="list-style-type: none"> Performance is fastest when compared with other protocols. 	<ul style="list-style-type: none"> Requires emulation software such as PC NFS or Samba to provide heterogeneous access to VOB files. More expensive because it requires a dedicated fiber-optic storage network. Extra management of switches. Backup and recovery are complex and require more steps compared to NAS.
iSCSI	<ul style="list-style-type: none"> Performance second to FC SAN. More cost efficient than FC SAN because it uses the existing Ethernet infrastructure. 	<ul style="list-style-type: none"> Requires emulation software such as PC NFS or Samba to provide heterogeneous access to VOB files. Backup and recovery are complex and require more steps compared to NAS.
NFS	<ul style="list-style-type: none"> Provides direct file access to heterogeneous clients and eliminates the need for emulation software. Inexpensive. Simple to manage. Offers ease of use and granularity in backup and recovery operations. 	<ul style="list-style-type: none"> I/O performance is slower compared to other protocols. However, performance is adequate for most deployments.
Hybrid configuration (database on iSCSI and pools on NAS)	<ul style="list-style-type: none"> Improves VOB performance. Provides direct access to heterogeneous clients and eliminates the need for emulation software. If the database is deployed on iSCSI, can be very cost efficient. 	<ul style="list-style-type: none"> Backup and recovery are complex because backups are performed by using two separate protocols while maintaining concurrency between them. However, using NetApp SnapDrive® data management software alleviates this complexity.

5.2 WINDOWS VOB DEPLOYMENT

VOBs can be deployed on Windows by using any of the following NetApp storage protocols: FC, iSCSI, or CIFS. Hybrid configuration is available in Windows environments. For information about hybrid configuration for Windows, refer to the IBM Rational white paper titled [“Remote VOB Storage Pools on Windows.”](#) Table 5 highlights the advantages and disadvantages of deploying ClearCase VOBs on each of NetApp’s storage protocols in a Windows environment.

Table 5) Pros and cons of VOB deployment on Windows.

Deployment	Pros	Cons
FC SAN	<ul style="list-style-type: none"> Performance is fastest when compared with other protocols. 	<ul style="list-style-type: none"> Does not provide direct file access to Windows clients. More expensive because it requires a dedicated fiber optic storage network. Requires extra management of switches. Backup and recovery are complex and require more steps compared to CIFS.
iSCSI	<ul style="list-style-type: none"> Performance second to FC SAN. More cost efficient than FC SAN because it uses the existing Ethernet infrastructure 	<ul style="list-style-type: none"> Does not provide direct file access to Windows clients. Backup and recovery are complex and require more steps compared to CIFS.
CIFS	<ul style="list-style-type: none"> Provides direct file access to Windows clients. Inexpensive. Simple to manage. Offers ease of use and granularity in backup and recovery operations. 	<ul style="list-style-type: none"> I/O performance is slower compared to other protocols. However, performance is adequate for most deployments.

5.3 VIEW DEPLOYMENT

Views are usually deployed on each client’s host local disk. However, some customers may prefer to centralize developers’ views in order to back up source code or designs that developers are working on and have not yet checked into the VOB repository. Views can be centralized on any Cluster-Mode storage protocol. The advantages and disadvantages of deploying views on each protocol are similar to those for VOB deployment, as described in section 5.2. However, for granularity and ease of backup and recovery, the NAS protocol is the best deployment.

6 IMPLEMENTATION PLANNING

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

- What are their main storage criteria: performance, data protection, reliability, manageability, and so on? If performance is the main criterion, ask them to define their expected latency.

- Is the VOB server hosted on Windows or UNIX? If UNIX, will developers access VOB by both Windows and UNIX clients?
- If they have existing VOBs, what is the current size of the VOBs and their expected rate of growth?
- How many developers will access the VOBs?
- What is their existing infrastructure: platform (servers and storage devices), protocol, and network infrastructure details? Do they have multiple VOB servers? How many? Separate License or Registry server?
- Are views local to each client, or are they in a centralized location? Number of View servers, if any? Types of views—dynamic or snapshot?
- What is the nature and frequency of builds and the resulting data growth? Do they use ClearCase clearmake, or do they use other builds like make? If they use other build makes, then derived objects will be not be populated in ClearCase VOB.
- Is storage to be used solely for ClearCase applications?
- What is their current backup and restore mechanism?

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

Table 6)Selection criteria across NetApp storage protocols.

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

7 STORAGE EFFICIENCY

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

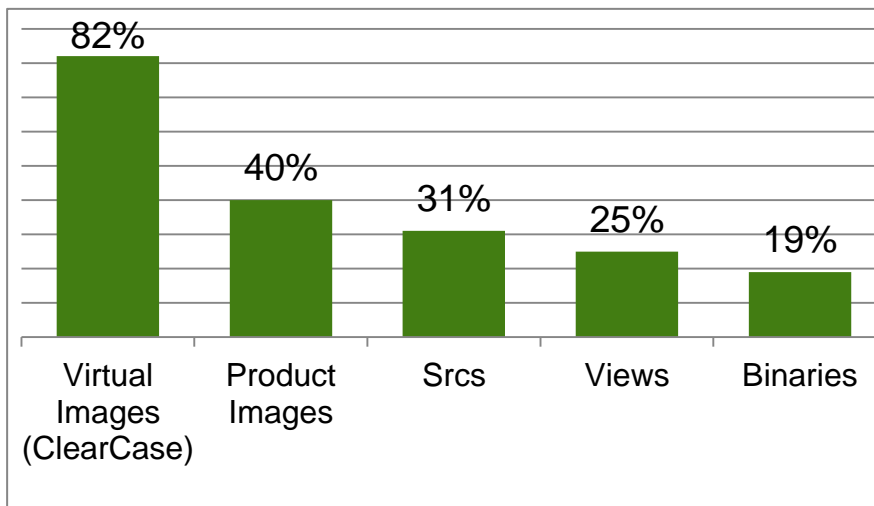
7.1 DEDUPLICATION AND COMPRESSION

Deduplication is a key feature that can significantly reduce the data footprint of developer workspaces in ClearCase environments, because developers tend to share or duplicate a lot of common libraries and binaries in their workspaces.

Deduplication was conducted on VOB and View repositories as well as ClearCase images; Figure 12 describes the savings observed. In our product environment and virtual testing of ClearCase, we found that deduplication improved storage utilization. We observed storage space savings of 82% for virtual images; 40% for product images, which are flat images; 31% for sources (srcs), which were stored in the VOB database; 25% for views; and 19% for binaries stored in the VOB database. This level of data reduction can help IT organizations contain costs by deferring the purchase of additional storage.

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

Figure 12) Deduplication of ClearCase repositories.



Further investigation will be conducted to understand the benefits of using deduplication and compression of whole file copies instead of ClearCase file delta copies and compressions. Using file delta copies can exponentially increase retrieval times because more versions of files are checked into the VOB repository.

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

7.2 NETAPP FLEXCLONE

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

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

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

For more information on NetApp FlexClone, see [TR-3347 A Thorough Introduction to FlexClone Volumes](#).

8 BEST PRACTICES

The following best practices are recommended when deploying ClearCase on NetApp Data ONTAP operating in Cluster-Mode storage:

- Create and lay out volumes with recovery and manageability in mind; use separate volumes for ClearCase VOB and view storage.
- Disable quotas on volumes to be used for VOB storage,
- Interdependent, admin, hyperlinked, and UCM VOBs should be grouped together.
- Disable oplocks.
- Specify the correct security style for your environment. In Interop mode, specify mixed security style when there is a mix of Windows and UNIX clients accessing the VOBs. Specify UNIX security style when there are only UNIX clients. For Windows environments, specify NTFS.
- Rational ClearCase relies on group IDs to implement workgroup sharing and to maintain strict access control. In Interop mode, 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, to avoid permission denied errors.
- Do `vol move` operations when the storage is not heavily loaded, because it does take up some storage CPU cycles.

9 CONFIGURATION AND SETUP

This section describes using the command line to configure and set up Data ONTAP operating in Cluster-Mode storage for use in ClearCase environments. Cluster-Mode storage can also be easily configured by using the System Manager GUI and the Data ONTAP Element Manager. This is not a comprehensive guide; for more information, see the [ClearCase Administrative Guide](#) and the [NetApp Cluster Mode Administration Guide](#). Here are the high-level steps to create a VOB or view storage in Cluster-Mode:

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

8. Mount the volume.
9. Run ClearCase `mkvob` or `mkview` on the mounted volume to create VOB or View repositories for ClearCase use.

The following sections are a quick guide to creating VOB storage for NFS, CIFS, and iSCSI in both Windows and UNIX environments.

9.1 AGGREGATE CREATE

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

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

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

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

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

To view what has been created, execute the following command. For this aggregate, the disks reside on a node named `fas3170c-sv111`.

```
cl_agnes_cmode::> storage aggr show
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID	Status
aggr_test	0B	0B	-	creating	0	fas3170c-sv111	raid_dp,	initializing

9.2 VSERVER CREATE

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

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

2. To view the Vserver, enter:

```
cl_agnes_cmode::> vserver show
```

Vserver	Type	Admin State	Root Volume	Aggregate	Name Service	Name Mapping
test	cluster	running	test_root	aggr_test	file	file

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

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

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

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

```
Vserver: test
Vserver Type: cluster
Vserver UUID: d368e7c5-25b6-11e1-8a92-123478563412
Root Volume: test_root
Aggregate: aggr_test
Name Service Switch: file
Name Mapping Switch: file
NIS Domain: -
Root Volume Security Style: mixed
LDAP Client: -
Language: c
Snapshot Policy: default
Comment:
Anti-Virus On-Access Policy: default
Quota Policy: default
List of Aggregates Assigned: aggr_test
Limit on Maximum Number of Volumes Allowed: unlimited
Vserver Admin State: running
Allowed Protocols: nfs, cifs, fcp, iscsi
Disallowed Protocols: -
```

9.3 VOLUME CREATE

1. Create a volume in the Vserver:

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

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

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


3. Mount the volume for access:

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

9.4 NFS ACCESS

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

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

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

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

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

3. Create a network interface for clients to access NAS. Items to specify include the Vserver:

- a. LIF name
- b. Home node: `fas3170c-sv111` (It is good to specify the node where the disks reside. For this scenario, the disks in aggregate `aggr_test` reside on node `fas3170c-sv111`.)
- c. Home port: `e4a` (This is the home port of the home node specified in c.)
- d. IP Address: `172.31.8..224`

```
cl_agnes_cmode::> network interface create -vserver test -lif test_lif -role data -
data-protocol nfs,cifs,fcache -home-node fas3170c-sv111 -home-port e4a -address
172.31.8.224 -netmask 255.255.255.0 -status-admin up
```

```
Info: Your interface was created successfully; the routing group dl172.31.8.0/24 was
created
```

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

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

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

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

Vserver	User Name	User ID	Group ID	Full Name
test	root	0	0	-

2 entries were displayed.

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

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

9.5 CIFS ACCESS

This example creates a CIFS service on Vserver test.

1. Create DNS on a Vserver.

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

```
cl_agnes_cmode::> dns show
(vserver services dns show)
```

Vserver	State	Domains	Name Servers
test	enabled	clearcase.local	172.31.8.102

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

```
cl_agnes_cmode::vserver cifs> cifs create -vserver test -cifs-server test_cifs -
domain clearcase.local
```

In order to create an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges to add computers to the "CN=Computers" container within the "clearcase.local" domain.

```
Enter the user name: Administrator
```

```
Enter the password:
```

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

- a. Vserver
- b. LIF name
- c. Home node: fas3170c-sv111 (It is good to specify the node where the disks reside. For this scenario, the disks in aggregate aggr_test reside on node fas3170c-sv111.)
- d. Home port: e4a (This is the home port of the home node specified in c.)
- e. IP Address: 172.31.8..224

```
cl_agnes_cmode::> network interface create -vserver test -lif test_lif -role
data -data-protocol nfs,cifs,fcache -home-node fas3170c-sv111 -home-port e4a -
address 172.31.8.224 -netmask 255.255.255.0 -status-admin up
```

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

4. Share the volume:

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

```

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

```

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

```

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

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

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

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

```

6. Create user mappings. Provide Vserver, direction, position (number order 1-1024), pattern (Windows login), and replacement (UNIX login):

```

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

```

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

```

\\172.31.8.240\test_vol

```

9.6 INTEROP ACCESS (BOTH CIFS AND NFS ACCESS)

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

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

```

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

```

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

```

cl_agnes_cmode:: > vserver name-mapping create -vserver test -direction win-unix -
position 1 -pattern CLEARCASE\\Administrator -replacement root
cl_agnes_cmode:: > vserver name-mapping create -vserver test -direction unix-win -
position 2 -pattern root -replacement CLEARCASE\\Administrator

```

9.7 ISCSI ACCESS

ISCSI setup differs depending on the host operating system that requires access to the volume via the iSCSI protocol. For details, refer to the Installation and Administration Guide of the host utilities for the

specific operating system on the NetApp Support site, [NOW](#). The example setup for iSCSI in this section is for the Linux 5.6 operating system and Windows 2003 versions.

1. Create an iSCSI service on Vserver test:

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

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

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

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

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

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

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

```
cl_agnes_cmode::volume> security login show -vserver test
```

```
Authentication Acct
Vserver  UserName  Application Method Role Name  Locked
-----  -
```

Vserver	UserName	Application	Method	Role	Name	Locked
test	vsadmin	http	password	admin	no	
test	vsadmin	ontapi	password	vsadmi	no	
test	vsadmin	ssh	password	vsadmin	no	

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

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

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

```
cl_agnes_cmode::> network interface show -vserver test
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
test						
test_mgmt	up/up	172.31.8.247/24	fas3170c-svl11	e4b	true	
test_san09	up/up	172.31.8.248/24	fas3170c-svl09	e4b	true	
test_san10	up/up	172.31.8.249/24	fas3170c-svl10	e4b	true	

5. Create volumes for the LUNs:

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

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

CONFIGURE LINUX 5.6 FOR ISCSI ACCESS

On Filer:

1. Create LUNs for Linux access:

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

Created a LUN of size 250m (262144000)

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

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

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

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

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

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

4. Map the LUN to the igroup:

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

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

On Host:

1. Download and install the latest iSCSI host utilities from the NetApp Support site for the Linux host. Refer to the Host Utilities Manual from the [NOW](#) site for complete instructions on how to install the host utilities. The following next steps are quick instructions on how to configure Linux 5.6 for iSCSI access. This is not meant to be comprehensive and thus refer to Host Utilities Manual for more details.
2. Verify that the RedHat Linux multipath packages are installed on the Linux server:

```
# rpm -q device-mapper
device-mapper-1.02.55-2.el5
device-mapper-1.02.55-2.el5

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

```
device-mapper-multipath-0.4.7-42.el5
```

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

```
# cp /etc/multipath.conf /etc/multipath.conf.orig
defaults {
    user_friendly_names no
    queue_without_daemon no
    flush_on_last_del yes
    max_fds max
    pg_prio_calc avg
}
devices {
device {
    vendor "NETAPP"
    product "LUN"
    path_grouping_policy group_by_prio
    features "1 queue_if_no_path"
    prio_callout "/sbin/mpath_prio_alua /dev/%n"
    path_checker directio
    path_selector "round-robin 0"
    failback immediate
    hardware_handler "1 alua"
    rr_weight uniform
    rr_min_io 128
    getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
}
}
```

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

```
# scsi_id -gus /block/sda
SATA_HDS728080PLA380_PFD32S0R3WHJM
```

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

```
blacklist {
    wwid SATA_HDS728080PLA380_PFD32S0R3WHJM
    devnode "^hd[a-z]"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
```

5. Enable and start MPxIO for iSCSI:

```
# chkconfig multipathd on
```

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

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

7. Reboot to enable changes:

```
shutdown -r now
```

8. When the Linux host is back up, enable discovery `sendtargets` to the iSCSI LIF IP address:

```
[root@ibmx3455-sv101 etc]# iscsiadm -m discovery -t sendtargets -p 172.31.8.248
172.31.8.248:3260,1056 iqn.1992-08.com.netapp:sn.d368e7c525b611e18a92123478563412:vs.15
172.31.8.249:3260,1057 iqn.1992-08.com.netapp:sn.d368e7c525b611e18a92123478563412:vs.15
```

9. Log into the iSCSI targets that have been discovered for verification:

```
[root@ibmx3455-sv101 etc]# iscsiadm -m node -L automatic
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.d368e7c525b611e18a92123478563412:vs.15, portal: 172.31.8.249,3260]
successful.
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.d368e7c525b611e18a92123478563412:vs.15, portal: 172.31.8.248,3260]
successful.
```

10. Use sanlun to verify the multipath policy and provider:

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

11. Verify multipaths:

```
[root@ibmx3455-sv101 etc]# multipath -ll
3600a09803246696e433f2d2d636e6b39 dm-6 NETAPP,LUN C-Mode
[size=250M][features=1 queue_if_no_path][hwhandler=1 alua][rw]
\_ round-robin 0 [prio=50][active]
  \_ 11:0:0:0 sdm 8:192 [active][ready]
\_ round-robin 0 [prio=10][enabled]
  \_ 8:0:0:0 sdj 8:144 [active][ready]
  \_ 9:0:0:0 sdk 8:160 [active][ready]
  \_ 10:0:0:0 sdl 8:176 [active][ready]
```

12. Create a file system on the device:

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

mke2fs 1.39 (29-May-2006)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
64000 inodes, 256000 blocks
12800 blocks (5.00%) reserved for the super user
First data block=1
```

```

Maximum filesystem blocks=67371008
32 block groups
8192 blocks per group, 8192 fragments per group
2000 inodes per group
Superblock backups stored on blocks:
    8193, 24577, 40961, 57345, 73729, 204801, 221185

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

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

```

13. Mount the file system and validate the size:

```

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

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


```

CONFIGURE WINDOWS FOR LUN ACCESS

For LUN Windows access, setup can be done either manually or by using SnapDrive for Windows. This example is a manual setup on Windows 2003. The same approach can be applied to Windows 2008. However, in Windows 2008, the Microsoft Software Initiator is already installed; it just needs to be enabled.

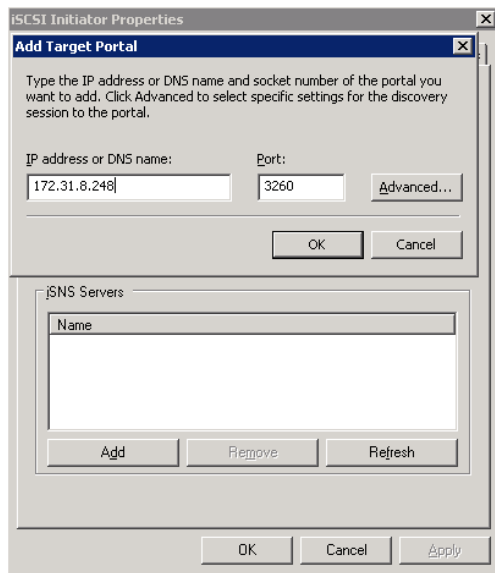
Here are the instructions for Windows 2003.

1. Download, install, and enable the [Microsoft Software Initiator](#) components: Initiator Service, Software Initiator, and MPIO onto the Windows 2003 server.

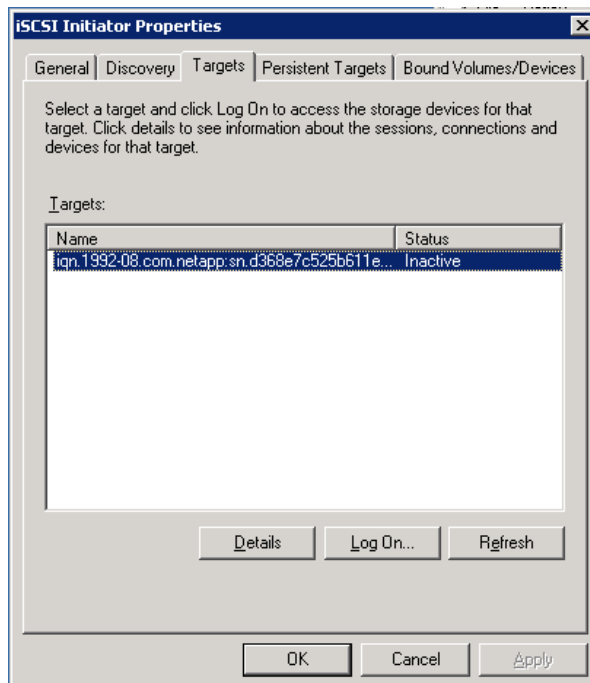


2. Reboot the system.
3. Click the Microsoft iSCSI Initiator icon. If there is no icon, start it from Start > All Programs.

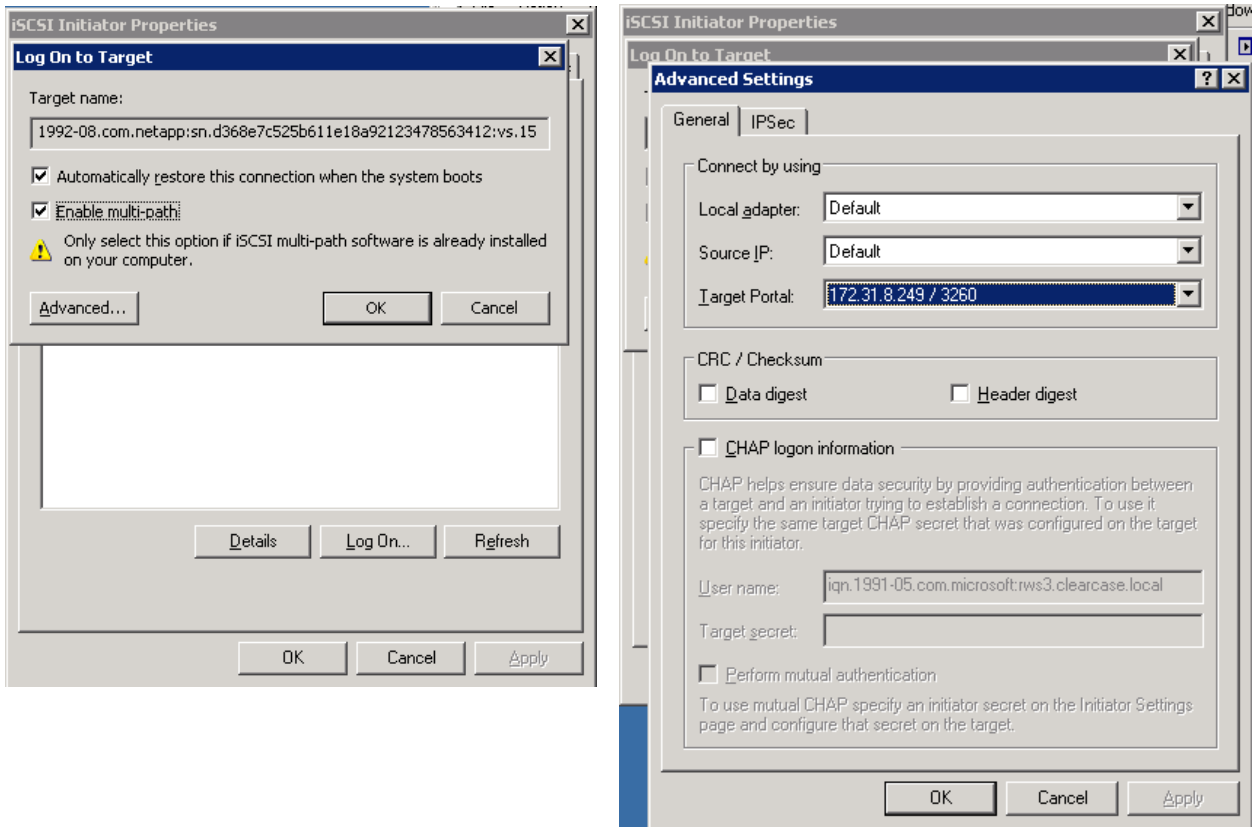
4. Connect to the NetApp system. Select the Discovery tab and then add the IP address of the iSCSI LIF. In this example, it is 172.31.8.248.



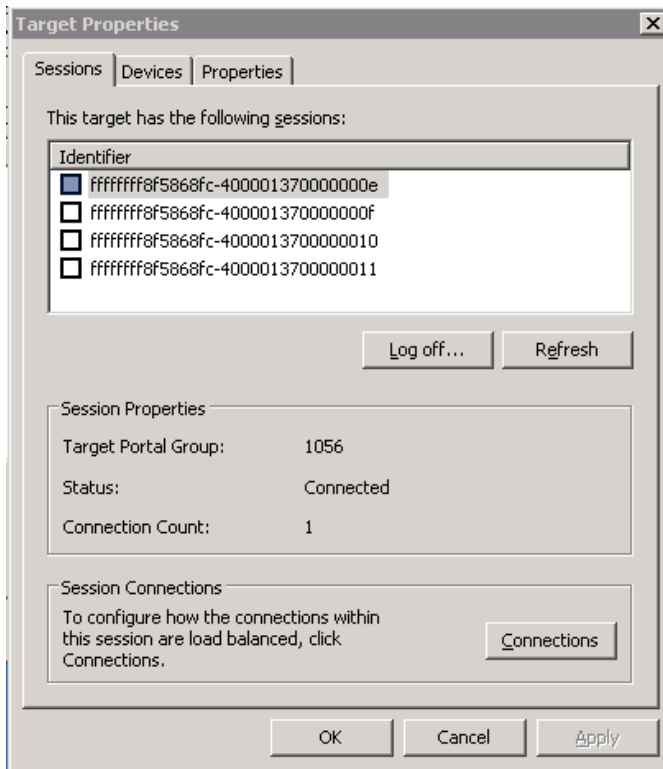
5. Connect to the target or volume. Select the Targets tab, select the target, and click Log On to connect to the target.



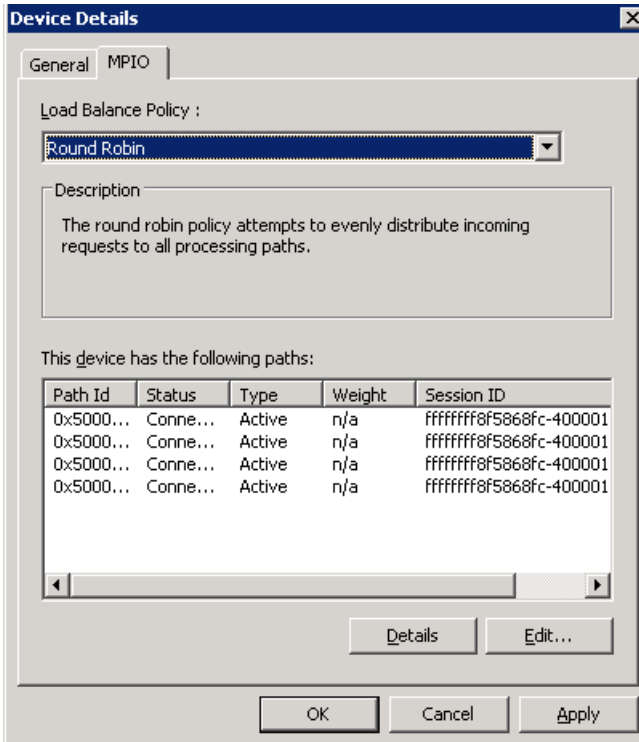
- For each node, click Log On, select the Enable Multi-path and Automatically Restore checkboxes, and click Advanced to select the target name for each node. Then click OK.



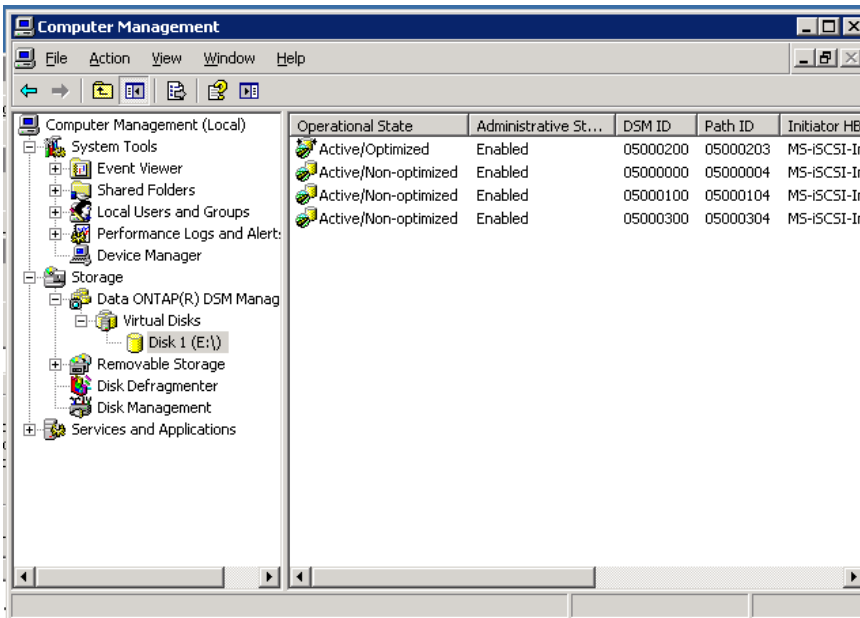
- On the Targets tab, click Details to see an Identifier entry for each node.



- On the Devices tab, click Advanced and then select the MPIO tab to see the following connections, depending on the MPIO policy:



- Under Computer Management > DSM, note that one path is Active/Optimized and the other paths are Active but not optimized.



10. On the Cluster-Mode command line, create a LUN for Windows:

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

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

11. Validate LUN creation:

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

Vserver   Volume      Qtree      LUN          State  Mapped  Type      Size
-----
test      testsan_vol  ""         windows.lun1 online  unmapped windows  251.0MB
```

12. Get the initiator for Windows:

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

Tpgroup      Initiator
Vserver Name  TSIH Name      ISID          IGroup
-----
test         test_san09 85  iqn.1991-05.com.microsoft:rws3.clearcase.local
```

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

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

14. Validate creation:

```
cl_agnes_cmode::> igroup show

Vserver   Igroup      Protocol OS Type  Initiators
-----
test      iwin        iscsi   windows iqn.1991-
05.com.microsoft:rws3.clearcase.local
```

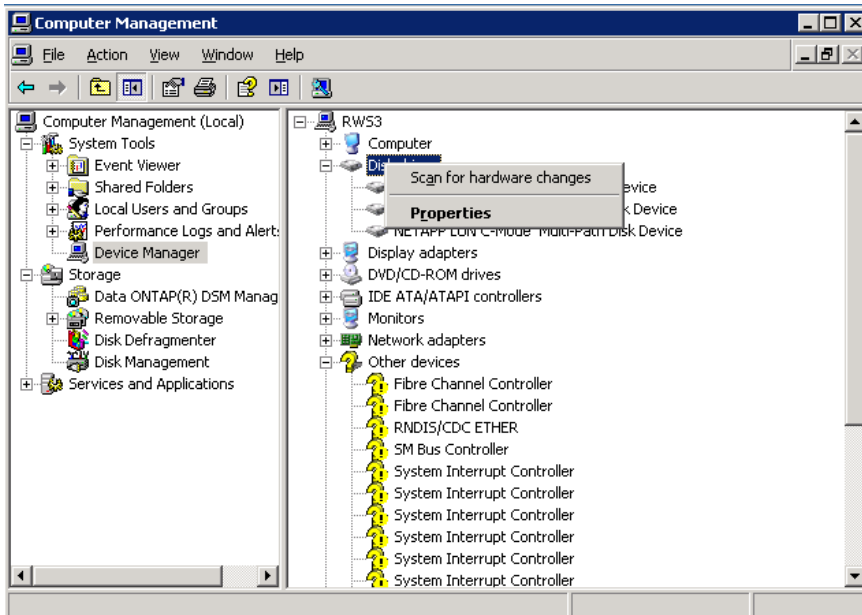
15. Map the LUN to the igroup:

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

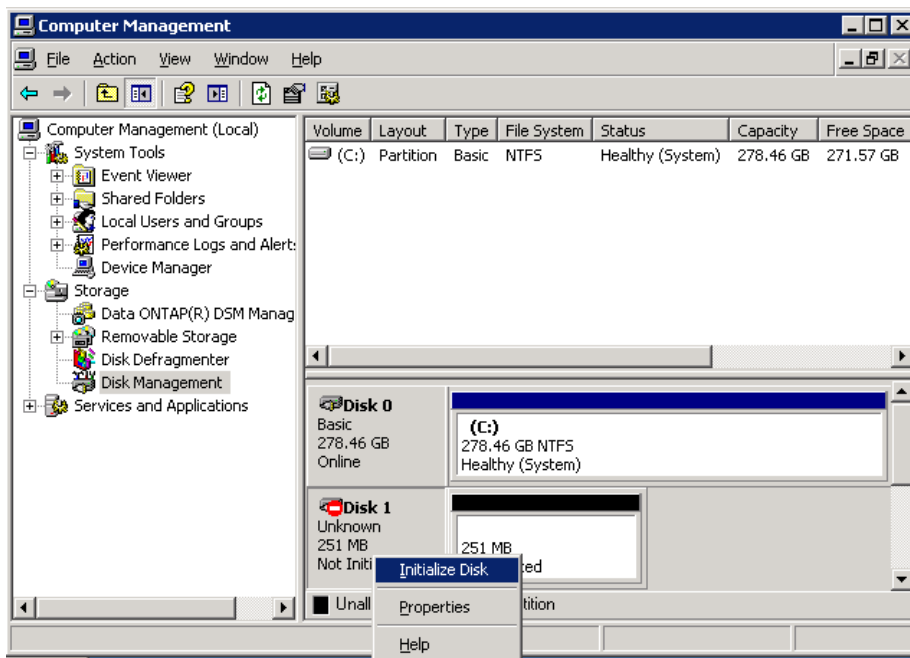
cl_agnes_cmode::> lun show -m

Vserver   Volume      Qtree      LUN          Igroup  LUN-ID  Protocol
-----
test      testsan_vol  ""         windows.lun1 iwin    0       iscsi
```

- In Windows, rescan the hardware to recognize the LUNs created. To start up Computer Management, select Start > My Computer. Right-click and select Manage. Select Disk Manage > Scan for Hardware Changes.



- Select Disk Management. Right-click Disk 1 and select Initialize Disk.



9.8 FIBRE CHANNEL ACCESS

Configuring the LUN for fibre channel (FC) access can differ depending on the operating system. For comprehensive instructions, refer to FC Host Utilities on the NetApp Support site, [NOW](#). This example walks through the setup for FC access on Linux 5.6. NetApp Data ONTAP operating in Cluster Mode no longer supports connecting the FC SAN storage directly to the host server. A fibre SAN switch with NPIV support is now required to connect all targets and initiators.

SET UP STORAGE

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

```
cl_agnes_cmode::> run -node fas3170c-sv109
Type 'exit' or 'Ctrl-D' to return to the CLI
fas3170c-sv109>
fas3170c-sv109> fcadmin config
```

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

```
fas3170c-sv109> priv set diag
Warning: These diagnostic commands are for use by NetApp
        personnel only.
fas3170c-sv109*> fcadmin offline 0c
fas3170c-sv109*> fcadmin config -t target 0c
A reboot is required for the new adapter configuration to take effect.
fas3170c-sv109*>
cl_agnes_cmode::> run -node fas3170c-sv110
Type 'exit' or 'Ctrl-D' to return to the CLI
fas3170c-sv110> fcadmin config
```

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

```
fas3170c-sv110> fcadmin config -t target 0c
A reboot is required for the new adapter configuration to take effect.
fas3170c-sv110>
cl_agnes_cmode::> run -node * fcadmin config
2 entries were acted on.
```

Node: fas3170c-sv109

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

Node: fas3170c-sv110

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

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

Connection to 172.17.39.188 closed.

CREATE THE WWN AND MANAGEMENT LIF

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

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

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

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

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

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

```
cl_agnes_cmode::volume> security login show -vserver test
```

```
Authentication Acct
```

Vserver	UserName	Application	Method	Role	Name	Locked
test	vsadmin	http	password	admin	no	no
test	vsadmin	ontapi	password	vsadmi	no	no
test	vsadmin	ssh	password	vsadmin	no	no

3. Create the network interfaces for FC access:

```
cl_agnes_cmode::network> network interface create -vserver test -lif test_fcp09 -
role data -data-protocol fcp -home-node fas3170c-svl109 -home-port 0a
```

```
cl_agnes_cmode::network> network interface create -vserver test -lif test_fcp10 -
role data -data-protocol fcp -home-node fas3170c-svl110 -home-port 0a
```

```
cl_agnes_cmode::network> net int show -vserver test
```

Logical	Status	Network	Current	Current	Is		
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port	Home	
test							
test_fcp09	up/down	20:02:00:a0:98:29:02:1e	fas3170c-svl109	0a	true		
test_fcp10	up/down	20:03:00:a0:98:29:02:1e	fas3170c-svl110	0a	true		

4. Create volumes for the LUNs:

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

```

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

```

FC SWITCH SETUP

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

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

CONFIGURE LINUX 5.6 FOR FC ACCESS

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

```

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

```

Created a LUN of size 250m (262144000)

```

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

```

2. Get the initiator name for the Linux server:

```

cl_agnes_cmode::> fcp initiator show
Logical      Port      Initiator   Initiator
Vserver     Interface Address  WWNN      WWPN      Igroup
-----
qsun277_san
  qsun277_fcp01 40000  20:00:00:00:c9:48:ce:9d 10:00:00:00:c9:48:ce:9d
  qsun277_fcp02 40000  20:00:00:00:c9:48:ce:9d 10:00:00:00:c9:48:ce:9d

```

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

```

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

```

```

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

```

4. Map the LUN to the igroup:

```

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

```

```

cl_agnes_cmode::> lun show -vserver test

```


Vserver	Volume	Qtree	LUN	State	Mapped	Type	Size
test	testsan_vol	""	linux.lun1	online	mapped	linux	250M

- Download and install the latest FC Host Utilities from the NetApp Support site for the Linux host. Refer to the Host Utilities Manual from the [NOW](#) site for complete instructions on how to install the host utilities and set up FC access. The following next steps are quick instructions on how to configure Linux 5.6 for FC access. These steps are not comprehensive and thus refer to the Host Utilities Manual for more details.
- Verify that the RedHat Linux multipath packages are installed on the Linux server:

```
# rpm -q device-mapper
device-mapper-1.02.55-2.el5
device-mapper-1.02.55-2.el5

# rpm -q device-mapper-multipath
device-mapper-multipath-0.4.7-42.el5
```

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

```
# cp /etc/multipath.conf /etc/multipath.conf.orig
defaults {
    user_friendly_names no
    queue_without_daemon no
    flush_on_last_del yes
    max_fds max
    pg_prio_calc avg
}
devices {
device {
    vendor "NETAPP"
    product "LUN"
    path_grouping_policy group_by_prio
    features "1 queue_if_no_path"
    prio_callout "/sbin/mpath_prio_alua /dev/%n"
    path_checker directio
    path_selector "round-robin 0"
    failback immediate
    hardware_handler "1 alua"
    rr_weight uniform
    rr_min_io 128
    getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
}
}
```

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

```
# scsi_id -gus /block/sda
SATA_HDS728080PLA380_PFD32S0R3WHJM
```

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

```
blacklist {
    wwid SATA_HDS728080PLA380_PFD32S0R3WHJM
    devnode "^hd[a-z]"
}
```

```

        devnode "(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
        devnode "^cciss.*"
    }

```

9. Enable and start MPxIO for FC:

```
# chkconfig multipathd on
```

10. Verify that multipathd is running and working:

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

11. Reboot to enable changes:

```
shutdown -r now
```

12. Use sanlun to verify the multipath policy and provider:

```
sanlun lun show -p
```

13. Verify multipaths:

```

[root@ibmx3455-sv101 etc]# multipath -ll
3600a09803246696e433f2d2d636e6b39 dm-6 NETAPP,LUN C-Mode
[size=250M][features=1 queue_if_no_path][hwhandler=1 alua][rw]
\_ round-robin 0 [prio=50][active]
  \_ 11:0:0:0 sdm 8:192 [active][ready]
\_ round-robin 0 [prio=10][enabled]
  \_ 8:0:0:0 sdj 8:144 [active][ready]
  \_ 9:0:0:0 sdk 8:160 [active][ready]
  \_ 10:0:0:0 sdl 8:176 [active][ready]

```

14. Create a file system on the device:

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

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

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

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

15. Mount the file system and validate the size:

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

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

9.9 MAKE VOB OR VIEW

After the NAS volume or SAN has been mounted, use ClearCase tools to create VOBs and Views by using the newly created NAS or SAN volume. The following example is for a UNIX VOB server using a NAS device.

MAKE VOB

1. Log in to the UNIX VOB server (for example, `rus1`) as the VOB owner and execute:

```
cleartool mkvob -tag /var/tmp/rus1_test -c "test" -host rus1 -hpath
/netapp/vobstg/test.vbs -gpath /netapp/vobstg/test.vbs /netapp/vobstg/test.vbs
```

2. Mount the VOB on all UNIX hosts:

```
cleartool mount /var/tmp/rus1_test
```

MAKE VIEW

1. Log in to a UNIX ClearCase host as a member of the developer group:

```
cleartool mkview -tag test_view -host rus1 -hpath /netapp/viewstg/test_view.vws -
gpath /netapp/viewstg/test_view.vws /netapp/viewstg/test_view.vws
```

10 SUPPORT MATRIX

The ClearCase System Test Suite provided by IBM Rational was used to validate ClearCase on NetApp Data ONTAP 8.1 operating in Cluster-Mode. This test suite uses a combination of C programs and Performix to create and exercise ClearCase VOBs and views. The ClearCase System Tests simulated ClearCase users distributed across multiple test machines. The tests were run in a Windows-only environment and in an Interop mixed environment, in which the VOB server resided on the UNIX server and the clients accessing the VOBs and views were a combination of Windows and UNIX clients.

Logs generated from these tests were submitted to IBM Rational for review and verification of the tests. After IBM Rational confirmed that all of the tests passed, the IBM Rational ClearCase support matrix was updated to indicate support for NetApp Data ONTAP operating in Cluster-Mode. For more information, refer to [IBM's support matrix](#).

Refer to the [NetApp Interoperability Matrix](#) for the current qualifications and system requirements for your specific configuration, including the following:

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

11 LIMITATIONS

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

- Deployment of VOBs or views on iSCSI for an IBM AIX® host is not currently supported.
- For SAN deployment, support for hosts other than Linux will not be supported until NetApp Data ONTAP 8.1.1.
- Cluster-Mode does not currently support the following features. If customers have the following requirements, do not recommend upgrading to Cluster-Mode:
 - MetroCluster™
 - SnapVault®
 - SnapLock®
 - IPv6
 - Qtree SnapMirror
 - SnapMirror Sync

12 TROUBLESHOOTING

This section offers troubleshooting guidelines that may provide insight into some common problems when deploying ClearCase on NetApp storage. For comprehensive information, refer to the [IBM ClearCase Administration Guide](#) and [NetApp Cluster-Mode Administration guides](#).

12.1 CLEARCASE LOGS TO CHECK

ClearCase VOB log messages are located on the VOB server in the `/var/adm/rational/clearcase/log` directory. There are several log files that you can look at to diagnose issues on the ClearCase side. The following logs provide information on the ClearCase processes:

- `db_server_log`: Stores logs for the `db_server` process.
- `albd_log`: Stores logs for the `albd_server` process.
- `vobrpc_servr_log`: Stores logs for the `vobrpc_server` process.

12.2 PERMISSIONS ISSUES

Some common problems that have been observed in this solution relate to multiprotocol access in the Interop environment. If storage is not configured properly, and if there is a mismatch in user names and primary groups, permission-denied errors occur. If your developers require access to the VOBs and views from both Windows and UNIX clients, there are several important Interop considerations.

All Windows and UNIX user names and primary groups must match in both environments in order to access VOBs and views in a mixed environment. You must also specify the mixed security style of the volume in which the VOB or view resides.

The following knowledgebase articles on the NetApp Support site offer some useful tips.

[KB ID: 3012797](#): What are some important considerations when setting up CIFS in Cluster-Mode?

[KB ID: 1012788](#): How to troubleshoot CIFS in Data ONTAP 8x Cluster-Mode

[KB ID: 1012935](#): How to setup Windows to UNIX usermapping over LDAP

[KB ID: 1012801](#): How to diagnose Data ONTAP 8 Cluster-Mode NFS problems

[KB ID: 3011350](#): Troubleshooting NFS issue on the client side: Where to begin

Also, it is important that the ClearCase `albd process` user and group must map to the appropriate user and group in UNIX and Windows.

12.3 PERFORMANCE ISSUES

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

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

13 CONCLUSIONS

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

14 REFERENCES

14.1 TECHNICAL REPORTS AND SOLUTION BRIEF

- Backup and Recovery of IBM Rational ClearCase on NetApp Data ONTAP Operating in Cluster-Mode, by Agnes Jacob and Michael Donati
- Data ONTAP 7G documents on ClearCase: [TR 3774 – Best Practices for Deploying Rational ClearCase Using NetApp](#) and [TR 3677 – Data Protection IBM Rational ClearCase Environment](#), by Anil G and Subra Mohan
- [TR-3982- Data ONTAP 8.1 Operating in Cluster-Mode: An Introduction](#), by Charlotte Brooks
- Solution Brief: [NetApp Cluster Mode with IBM Rational ClearCase](#)

14.2 NETAPP UNIVERSITY WEB-BASED COURSES

- [Value of Data ONTAP Cluster-Mode for IBM Rational ClearCase Environments, CRS-001976-01-A](#)
- [Technical Overview of IBM Rational ClearCase Solution on Data ONTAP Cluster Mode, CRS-001978-01-A](#)

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 and IBM provide 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 and IBM products discussed in this document.



www.netapp.com

© 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. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, MetroCluster, NOW, SnapDrive, SnapLock, SnapMirror, SnapRestore, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Cisco Nexus is a registered trademark of Cisco Systems, Inc. IBM, AIX, ClearCase, Rational, and WebSphere are registered trademarks of IBM Corporation. Linux is a registered trademark of Linus Torvalds. Microsoft, Windows, and Windows Server are registered trademarks of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4032-0212