NETAPP TECHNICAL REPORT

Data Protection: IBM Rational ClearCase Environments

Anil G, Subra Mohan, NetApp

May 2008 | TR-3677

NETAPP DATA PROTECTION AND RECOVERY SOLUTIONS FOR IBM RATIONAL CLEARCASE ENVIRONMENTS

The old adage of "time is money" couldn't be more accurate, particularly when determining the value of storing and recovering data. Data explosion, resulting in ever increasing data footprint, combined with shrinking data backup windows has added complexity to the backup and restore process. In addition to this, given the global nature of software development, the SCM infrastructure needs to be available around the clock and cannot be unavailable for long stretches during tape backups NetApp Snapshot technology has revolutionized the backup and recovery scenarios in Rational ClearCase environments, resulting in drastic reduction in VOB lock times, addressing larger data footprints and frequent full backups. In the following sections we will be exploring NetApp technology and combination of solutions that bring in a radical change in the backup and restore process in IBM Rational ClearCase environments.

TABLE OF CONTENTS

1 OVERVIEW

The old adage of "time is money" couldn't be more accurate, particularly when determining the value of storing and recovering data. Data explosion, resulting in ever increasing data footprint, combined with shrinking data backup windows has added complexity to the backup and restore process.

It's important to understand how technology choices can impact business. The service-level agreements (SLAs) are usually specified by the business application owners and typically include the recovery point objective (RPO) and the recovery time objective (RTO). The RPO is the amount of data that might be lost, and the RTO is the amount of time it takes to perform the recovery. Based on a recent study done by Forrester Consulting on behalf of NetApp, the most critical concern cited by survey respondents was being able to complete all their backups within defined windows. Factors cited by survey respondents that drove them to adopt disk-based backup solutions were reliability, followed by improved recovery times and speed of backups. The restore process has been the Achilles heel of tape backup, resulting in users waiting for their data, which was either corrupted or accidentally deleted, to be restored by the backup administrator. Disk-based backup is the most optimal solution to meet the most demanding service-level agreements.

Traditionally a tape-based solution has been considered to be much cheaper than a disk-based solution from the standpoint of acquisition cost. But the cost disparity is diminishing with the space savings possible in a disk-based backup solution resulting from backing up only block-level incremental changes as opposed to backing up entire files. Disk costs are further reduced by utilizing deduplication technologies, along with the usage of high-capacity, low-cost drives (such as SATA). And, when you take into account the operational and administration costs, on top of the reliability issues associated with a tape-based solution, the total cost of ownership of a disk-based solution is much lower than that of tape. Based on a Mercer Management Consulting study done in 2006, a disk-to-disk backup solution based on NetApp® SnapVault® is 48% less expensive than a similar capacity tape-based backup solution.

Most of the enterprise deployments of software configuration management applications such as IBM Rational ClearCase and IBM Rational ClearQuest are now looking to stipulate significantly higher levels of SLAs defined by RPO and RTO. In addition to this, given the global nature of software development, the SCM infrastructure needs to be available around the clock and cannot be unavailable for long stretches during tape backups, which typically necessitate the VOB to be locked during the backup process. It's impractical, if not impossible, to meet these RPO and RTO levels, as well as rapid restore needs, while shortening the backup window using tape as the primary means for backing up data that might need to be recovered in the near term. A combination of approaches is typically warranted, whereby a disk-based backup solution is used for saving recent work, and tape is used as a longer-term archival solution.

In the following sections we will be exploring NetApp technology and combination of solutions that bring in a radical change in the backup and restore process in IBM Rational ClearCase environments.

2 TYPICAL RATIONAL CLEARCASE DEPLOYMENT

As shown in Figure 1, a typical IBM Rational ClearCase environment consists of a single or multiple VOB servers, view servers, registry server, backup registry server, license server, Rational ClearQuest server, build farm, storage, and client network. VOB and view data, registry files, and ClearQuest database are candidates for data backups.

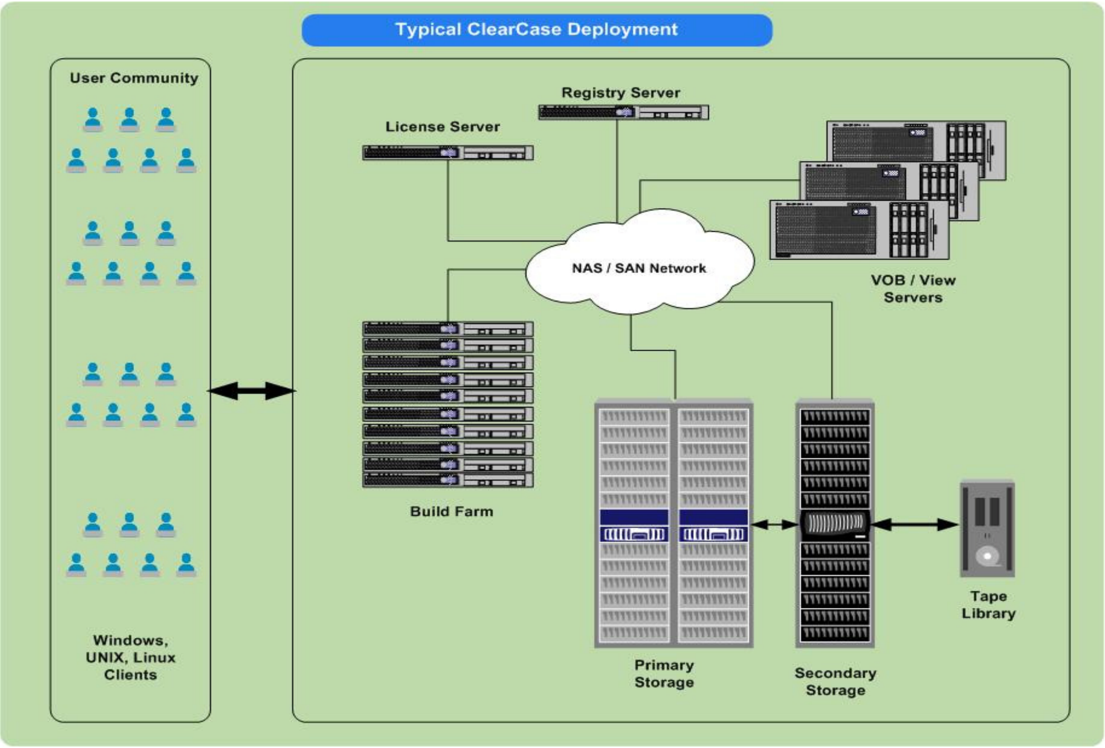

Figure 1) Typical Rational ClearCase deployment

2.1 Backup Requirements

Key factors to be considered while planning your IBM Rational ClearCase backup environment are:

VOB locking

Should back up files even if they are open for writing

Able to back up the largest file in the VOB

Preserve file access times and permissions

Perform full backups

View backup

Backup registry and client list files

The standard VOB backup procedure requires you to lock the VOB, back up the entire VOB storage directory, and then unlock the VOB. Locking a VOB prevents checkins, checkouts, and other operations that affect VOB data and metadata. These include UCM deliver operations and any UCM rebase operations that require merging.

For VOB servers based on UNIX® you just have to lock the VOBs for backup. When backup is done from a Windows® VOB server, often the ClearCase processes have to be stopped. Although VOBs might be locked, some ClearCase processes still hold VOB database files open for write. And very few backup programs can really read files that are open for write for another program. Shutting down ClearCase processes for backup means that during that period the scheduler is not running, so no scheduled activities will be started, like the daily cleanup routines or MultiSite replication.

The backup tool should preserve file access times. On some UNIX platforms, a few utilities, which are often used for backups, reset file access times. This can disrupt derived objects (DOs) and cleartext storage pool scrubbing patterns and might prevent these pools from ever being scrubbed.

File size limitations add to the level of dependency. VOBs that are at schema version 54 can include string files that are larger than 2GB. VOBs that are at feature level 5 can include container files that are larger than 2GB.

IBM Rational recommends going for a full backup strategy for backing up VOBs. Due to the way ClearCase stores the data in the storage structure, the incremental backups are often almost as large as the full backups. And restoring from an incremental backup takes longer than from a full backup.

The content of views, unlike that of VOBs, can usually be reconstructed easily. With the exception of changes to checked-out versions and other view-private files, the contents of any view can be recovered by recreating the view. Regular backups of views can still be important, especially if users are not in the habit of checking in their work regularly. Backing up a view is similar to backing up a VOB.

Rational ClearCase registry data is kept in a group of files in the rgy directory on the registry server host. These are ordinary files that can, and should, be backed up regularly.

2.2 Planning Appropriate Strategy

Going further, we also need to factor in the other specific requirements (as mentioned below) before arriving at a reliable and efficient strategy:

- Accommodate rapid data growth
- Back up interdependent VOBs, both in base ClearCase and UCM-based environments
- Multinodal backup
- Specific file system requirements
- Frequent backups and quick restores
- Planning for multitier backups
- Validity of backups
- Existing SCM methodologies

Often, however, when the amount of the data in ClearCase repositories grows, there is not enough time to create the backup the normal way. When using ClearCase with ClearQuest, the ClearQuest databases should be backed up at the same time as the UCM VOBs. In addition to normal backup from the ClearCase administration file area, it is recommended to have a registry backup that you can have online immediately.

Sometimes it might be needed to dig out a removed element, for example. You probably know that it is bad configuration management to remove anything from an SCM system, but these things do happen.

Data validity is one more sticky point: Is the data valid? The answer to this would include the entire gamut of related points and would start with defining the data set, how long could it last, and how soon it could be restored and defining an approach to arrive at an appropriate solution.

2.3 Different Ways of Backing Up

Choosing between different ways of backup: The standard backup procedure is to lock the VOB and back up the entire VOB (db and pools) and move the contents onto a tape. The VOB must remain locked during the entire duration of backup. This is recommended for sites that can accept the duration of VOB locks.

Semilive backup involves backing up the two major pieces of the VOB separately. The vob_snapshot utility (provided by IBM Rational) can be configured to periodically lock the VOB and copy the VOB database to another location. The VOB storage pools are backed up routinely, without locking the VOB. Backing up the database component separately is fast and can be unlocked soon. Hence, semilive backup is one of the ways to reduce VOB lock times. The backed up VOB storage directory might include the copy of VOB database. However, because it is backed up while the VOB is unlocked, this database is useless and is discarded when the VOB is restored. On the whole, the semilive backup process increases complexity and uncertainties during restore.

Rational ClearCase MultiSite is also used as a VOB backup strategy. A replica copy of the VOB is created and can be used as a backup copy. This can be looked at in two ways: taking a standard backup off the replica or using the replica as a backup copy to avoid locking a VOB. Using MultiSite for backups means that the backup replica needs to remain online so that it can be updated frequently from the original. Almost twice as much disk space is required. (You do not need exactly twice as much space, because derived objects are not replicated, and the cleartext pool for the backup replica is smaller or nonexistent.)

However, the most important thing to note is that a MultiSite replica is not a complete copy of a VOB; derived objects, triggers, and nonobsolete locks are not replicated. In case of a backup and recovery scenario derived objects and pool assignments are the ones that really matter. After a recovery from backup, developers must rebuild derived objects associated with the VOB. Checked-in derived objects are replicated, so they are backed up. Also, pool assignments are specific to a replica, so recreating the replica from a backup replica can undo changes made to them. Any changes or corruptions at the primary are reflected in the replica.

The pros and cons of each of the above mentioned solutions are debatable. Furthermore, if you are also considering view and registry backup, you need to come up with a strategy that would accommodate all the candidates (VOB, view, and registry) for backup and restore in a given ClearCase environment.

NetApp Snapshot™ technology mitigates almost all of the shortcomings involved and brings in a radical change to the data protection scenario.

3 NETAPP SNAPSHOT TECHNOLOGY

The NetApp file system can copy itself at any point in time and make the copied versions of the file system available via special subdirectories that appear in the current (active) file system. Each copied version of the file system is called a Snapshot copy.

A Snapshot copy appears as though it is a read-only copy of the file system and usually requires a small disk space premium.

Snapshot copies are maintained as pointers to disk blocks containing data. When the WAFL® file system creates a Snapshot copy, it makes a copy of the set of pointers from the active file system, but does not actually copy data blocks. While the active file system changes, Snapshot copies continue to

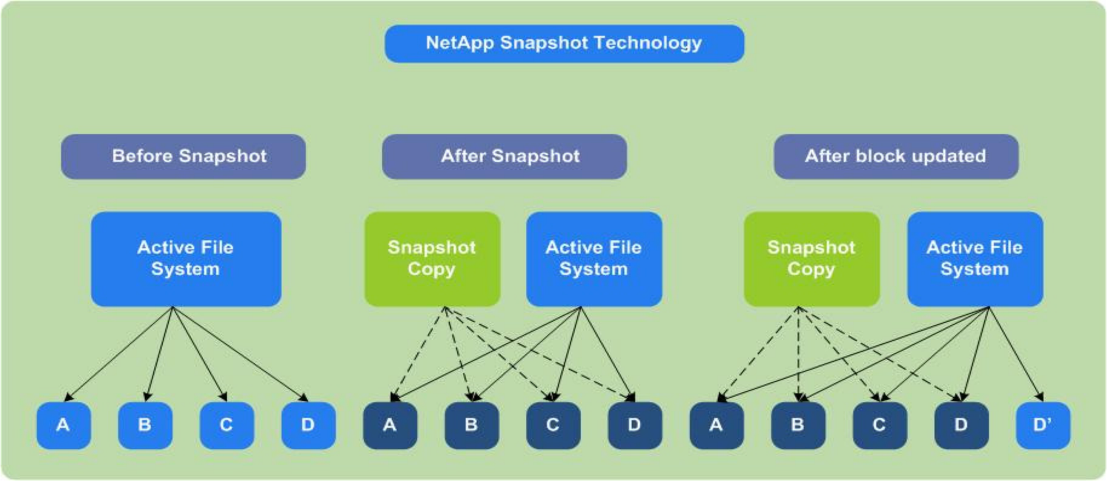

Figure 2) NetApp Snapshot

point to the original content of the changed disk blocks, holding these blocks from the file system's free space, thereby using incremental disk space.

The foundation for protection of data is Snapshot technology. Every NetApp storage system includes this time-proven capability, providing point-in-time static views. When regular Snapshot copies of data are created, information is protected against accidental deletion or corruption.

SnapRestore®, SnapVault, and SnapMirror®, which are described later in this document, are built on top of the underlying Snapshot functionality.

3.1 NetApp Snapshot Copies in Rational ClearCase Environments

NetApp Snapshot technology can be leveraged very effectively in IBM Rational ClearCase environments. As stated earlier, Snapshot copies can be created in seconds and consume no**[[NOTE: Should this be "little" instead of "no"?]]** additional disk space (Snapshot copies consume space only when the file system changes).

In Rational ClearCase environments, the Snapshot feature can be enabled on volumes containing VOBs and views.

A Snapshot copy represents a full backup of the file system at that given point in time. Rational ClearCase requires the VOBs to be locked before triggering the Snapshot copy. The Snapshot mechanism is extremely fast, and it takes only a few seconds to create a Snapshot copy. This in turn drastically reduces the VOB lockout periods, typically from few hours to a few seconds. It generally takes longer to execute the Rational ClearCase lock and unlock commands.

Snapshot copies can be scheduled to occur automatically or be created manually. Automatic schedules can be created on an hourly, nightly, or weekly basis. A volume can maintain 255 Snapshot copies concurrently.

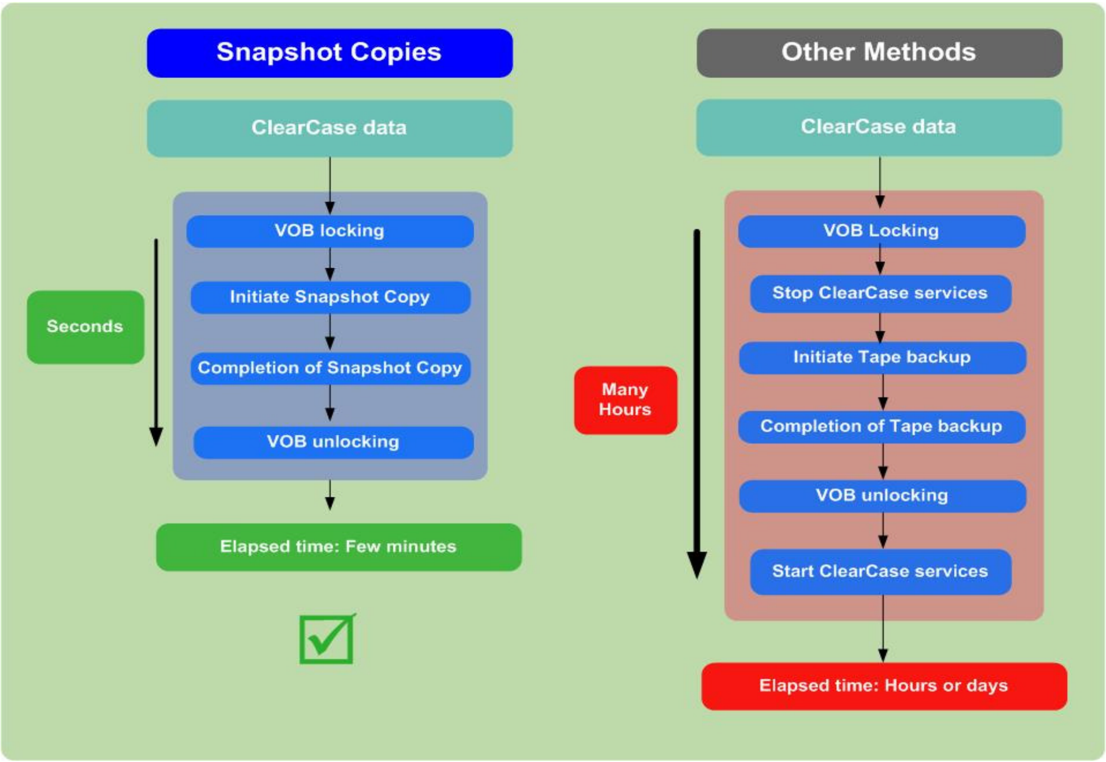

Figure 3) Snapshot copies and other Rational ClearCase backup processes

View private files can also be recovered from the Snapshot copies.

Note: NetApp Snapshot copies are different from the ClearCase vob_snapshot utility.

4 NETAPP SOLUTIONS TO DATA PROTECTION

Smart deployment is all about risk management: planning reduces risk. And properly planning for an appropriate backup approach will result in a reliable and quicker backup and restore process. Moreover, the RPO and the RTO also drive the need to go for either a specific or a combination of backup approaches.

The NetApp approach to protecting data begins with robust and highly reliable hardware and software. Clustered NetApp storage systems deliver up to 99.999% uptime as measured across our installed base of storage systems. NetApp also pioneered the use of double-parity RAID (RAID-DP®, a RAID 6 implementation) on general-purpose production storage. Prior to the introduction of NetApp RAID-DP, protection against multiple disk failures (without incurring a severe performance penalty) required disk mirroring (RAID 1), which provides protection but doubles your disk costs.

The proven capabilities of NetApp hardware are complemented by a unique set of data protection and data availability solutions. You can choose either a specific approach or a combination of the following approaches to Rational ClearCase data protection:

Online backups

Disk-to-disk backup/archival and disaster recovery

Disk-to-tape archival

Each of these approaches defines the purpose and mode of data availability. NetApp solutions make it simple to protect all data types and make sure of their rapid recovery. In the following sections we will discuss each of the backup approaches and associated NetApp solution in detail.

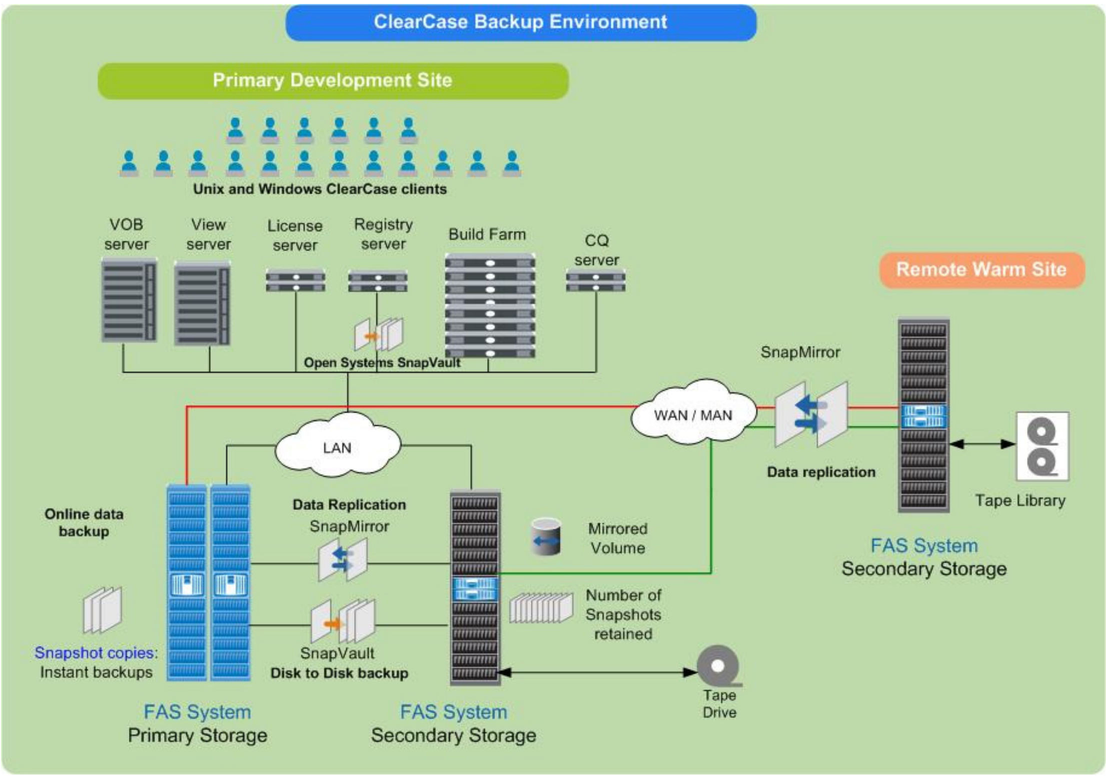

Figure 4) Rational ClearCase backup environment based on NetApp solution

5 ONLINE BACKUPS

Snapshot copies augment and simplify an overall enterprise data protection strategy. Snapshot copies can serve as daily backups and are considered as the first level of backup since they are located on the primary storage and within the same data center.

When regular Snapshot copies of data are created, information is protected against accidental deletion or corruption. Snapshot copies can be configured for an hourly, nightly, or weekly schedule. Users can self-recover their own files in a few seconds from any of the online Snapshot copies without requiring administrator intervention.

The Rational ClearCase configuration management system stores the source data in the form of data containers. Instead of modifying an existing data container when a new version of an element is checked in, Rational ClearCase creates a new container at a different pathname within the source storage pool. (It then deletes the old container.) This results in availability of the base and all deltas associated with a file element and its versions and the related metadata (Raima database) in a given Snapshot copy.

Typically, in case of a VOB database corruption issue, all that one would need is the latest Snapshot copy available. You can restore the said VOB or set of VOBs residing in the volume from the Snapshot copy.

In IBM Rational ClearCase environments, besides the VOB storage, views (dynamic and Snapshot view directories) too can be stored onto NetApp storage. View private files and particular directories can be instantaneously recovered from the online Snapshot copies. In this case, the user or the ClearCase administrator simply uses the copy command, native to the operating system on a host that can access the volumes on the storage device, to move the data from the Snapshot directory to the production area of the volume.

Further, the entire file system can be reverted to the state of any existing Snapshot copy in a matter of seconds using SnapRestore.

The IBM Rational dbcheck utility can be run on the Snapshot copies, thus assuring the consistency and reliability of the said backup.

Note: IBM Rational strongly recommends locking the VOB before initiating NetApp Snapshot copies. The typical sequence would be lock the VOB, then initialize NetApp Snapshot, then unlock the VOB.

6 DISK-TO-DISK BACKUP/ARCHIVAL AND DISASTER RECOVERY

Having an online backup copy makes sure of quick recovery in case of accidential file deletions or corruption. Even though NetApp storage systems provide a high level of availability with technologies such as RAID-DP and redundancies built at several levels, you still need to protect yourself against operator errors or site disasters such as power outages, fires, cyber attacks, or terrorist attacks, as well as regional disasters such as earthquakes, hurricanes, and floods depending on your geographic location. This can be accomplished either by backing up data to tape and shipping it off-site, or to a secondary storage system located within the same site, or at a remote site. Backups are only meaningful if you have the capability to recover the backed up VOBs or views. Making sure that a backup was sucessful is vital to the success of a data protection strategy.

Verification that a tape-based backup was successful is a cumbersome process. Once backups are verified, restores can be sucessfully performed. There can be problems associated with physically locating media, cataloging and archiving become very critical, and maintaining tapes off-site brings in additional complexity to the system.

Data recovery from a tape can be a very complex process. Transporting the tape from off-site, loading it, and repositioning it to the correct locations and thereafter reading the data are a complex and time consuming procedure. In a majority of the tape-only environments, recoveries take longer due to the large number of tape catridges involved and are subject to lower restore rates.

Morever, it is not feasible for performing hourly backups. It is clear that uptime requirements are becoming stricter and the backup windows are shrinking. The VOBs just can't be kept locked for hours together. In some scenarios, such as a global development environment, it might be nearly impossible to lock the VOBs even for the given shortest time period. To further compound these problems, the amount of VOB data can grow rapidly, resulting in larger backup footprints.

Given the above, it has been always a challenge to make sure of frequent and reliable backups of Rational ClearCase environments.

The combination of the NetApp Snapshot feature and disk-based backup methodologies has made a significant improvement to this process. Companies do not regard disk-based backup primarily as a means of reducing costs, but rather as a flexible intermediate layer for backing up data faster and, above all, for accelerating the restore process of Rational ClearCase data.

6.1 SnapMirror

As mentioned in the earlier section, online backups form the first level of defense; there are several approaches to increasing data availability in case of nonavailability of online backups. Mirroring is one of the proven mechanisms to make sure of data availability and minimize downtime. NetApp SnapMirror provides a fast and flexible enterprise solution for mirroring and replicating data over local area, wide area, and Fibre Channel networks.

NetApp SnapMirror can be a key component in implementing enterprise data protection strategies. The data available in the NetApp primary volume is mirrored to the secondary read-only volume. If disaster strikes the primary sites, the online backups might no longer be available. If critical VOB data is mirrored to a different physical location, a serious disaster will not necessarily translate to extended periods of data unavailability. The data can be restored from the mirrored backup.

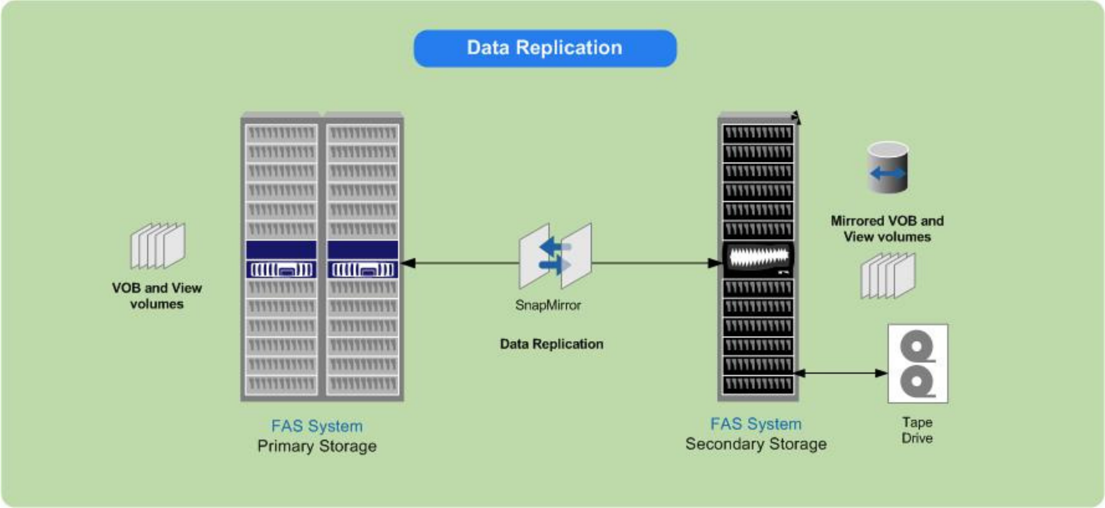

Figure 5) Data replication using NetApp SnapMirror

The mirrored target volume is an exact copy of the source volume at the block level. SnapMirror updates the mirror to reflect incremental changes and can be configured to replicate either using a schedule or via manual update. Updates are transferred asynchronously, thus avoiding data latency issues inherent in other mirroring techniques.

Assuming that the primary volume is unavailable, the mirrored volume can be turned into a production volume. It would involve some overheads from the Rational ClearCase perspective. The advantage of making the mirror to the production volume is that the recovery times are minimized since there is no data transfer involved. Breaking the mirror relationship converts the read-only replica on the destination to a writable. This is to accommodate the further growth as the mirrored volume will now become a production volume.

By setting up a reverse SnapMirror relationship between the new production volume and the former, you will be able to synchronize the data back to the original system.

On the other side, the regular Snapshot copies (online) resident in the volume are also mirrored as part of the SnapMirror transfer. These Snapshot copies are consistent since the VOBs would have been locked before initializing them. Using these Snapshot copies for restoration is recommended. The required Snapshot copy can be copied from the mirrored volume to the original and restored back in the same way as the online Snapshot copies (using SnapRestore). It is recommended that the mirroring update happens after the independent Snapshot schedule.

If the amount of data is less, a mirrored volume can be mounted, and the required data can be copied back to the production volume.

The SnapMirror destination can be either in the same data center or in a remote data center.

6.2 SnapVault and Open Systems SnapVault

Typically, users don't want to use up too much capacity on their primary storage for saving their Snapshot copies, for cost reasons. SnapMirror maintains a replica of the primary storage on a secondary storage system. So, if a Snapshot copy gets deleted on the primary storage system, then it automatically gets deleted on the secondary storage system as well. If you want to archive several Snapshot copies dating back to several days or months, then you will need to use SnapVault or Open Systems SnapVault. It might be rare that you would need to go back to a Snapshot copy created several days to a few months back, since the ClearCase database itself tracks changes to the source code over time. But in some cases it might be needed to dig out a removed element (rmelem), for example. You probably know that it is bad configuration management to remove anything from an SCM system, but these things do happen. Also, regulatory policies might force you to archive those Snapshot copies that go back in time.

Also, administrators are always on the lookout to further build multiple lines of defense protecting the online backups, usually multiple copies of backup triggered at different points in time. Tape-based backup has its own inadequacies (as mentioned earlier). Once again, disk-based backups enhance the data backup and restore capabilities.

SnapVault is a data protection solution for NetApp storage environments, providing online disk-to-disk

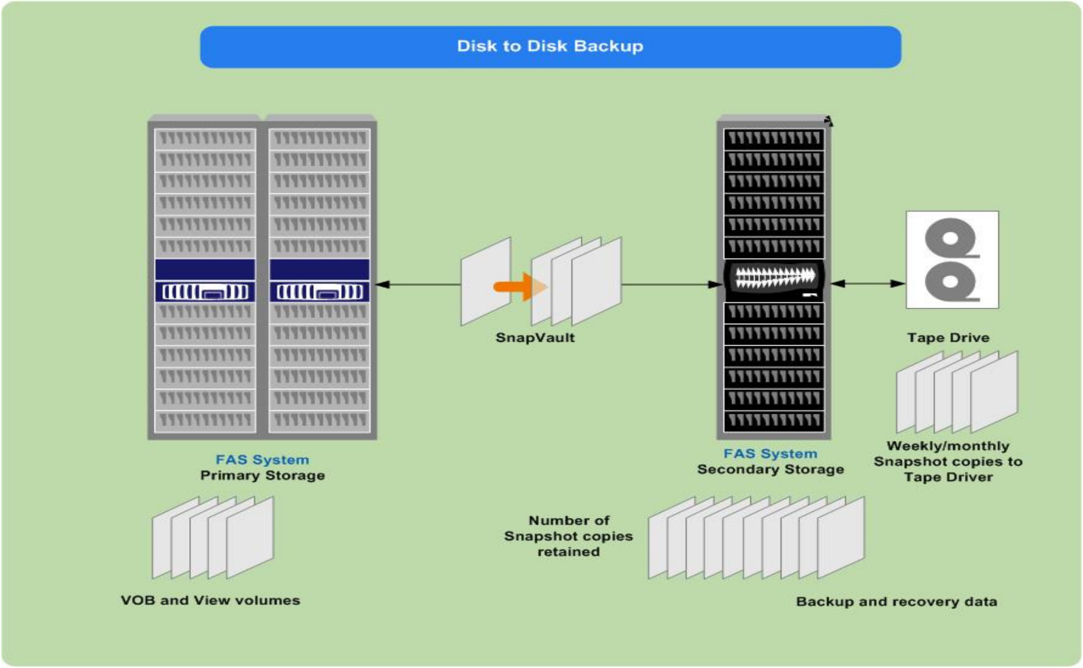

Figure 6) Disk-to-disk backup using SnapVault

backup and recovery for NetApp primary storage systems. SnapVault provides intelligent data movement, reducing network traffic and impact on production systems and performing frequent backups to make sure of superior data protection using field-proven NetApp Snapshot technology to efficiently store days' and potentially weeks' worth of backups online. Snapshot copies can then be backed up to tape using any NDMP-compliant backup software application.

SnapVault can be enabled on the volumes containing VOBs and views. It leverages Snapshot copies to transfer data from a NetApp primary storage system (referred to as SnapVault primaries) to a NetApp secondary storage system (referred to as SnapVault secondary). Unlike SnapMirror, the SnapVault secondary data is totally read-only. You cannot fail over to a SnapVault secondary.

Note: Locking the VOB before initiating the SnapVault base transfer and scheduling SnapVault updates is recommended.

SnapVault has been extended to support heterogeneous systems via Open Systems SnapVault. With Open Systems SnapVault, servers running Windows or various flavors of UNIX and Linux® operating systems, connected to any back-end storage system, can be backed up to a NetApp secondary storage directly. It especially makes sense to use Open Systems SnapVault to back up data in remote offices, where you don't have IT staff to handle backup tasks, to a NetApp storage system located in your primary data center.

The Rational ClearCase registry server can be configured as an Open Systems SnapVault primary. The rgy directory (/var/adm/rational/clearcase/rgy) and the client list file (/var/adm/rational/clearcase/client_list.db), which are essentially part of the Rational ClearCase backup, can be backed up to the SnapVault secondary.

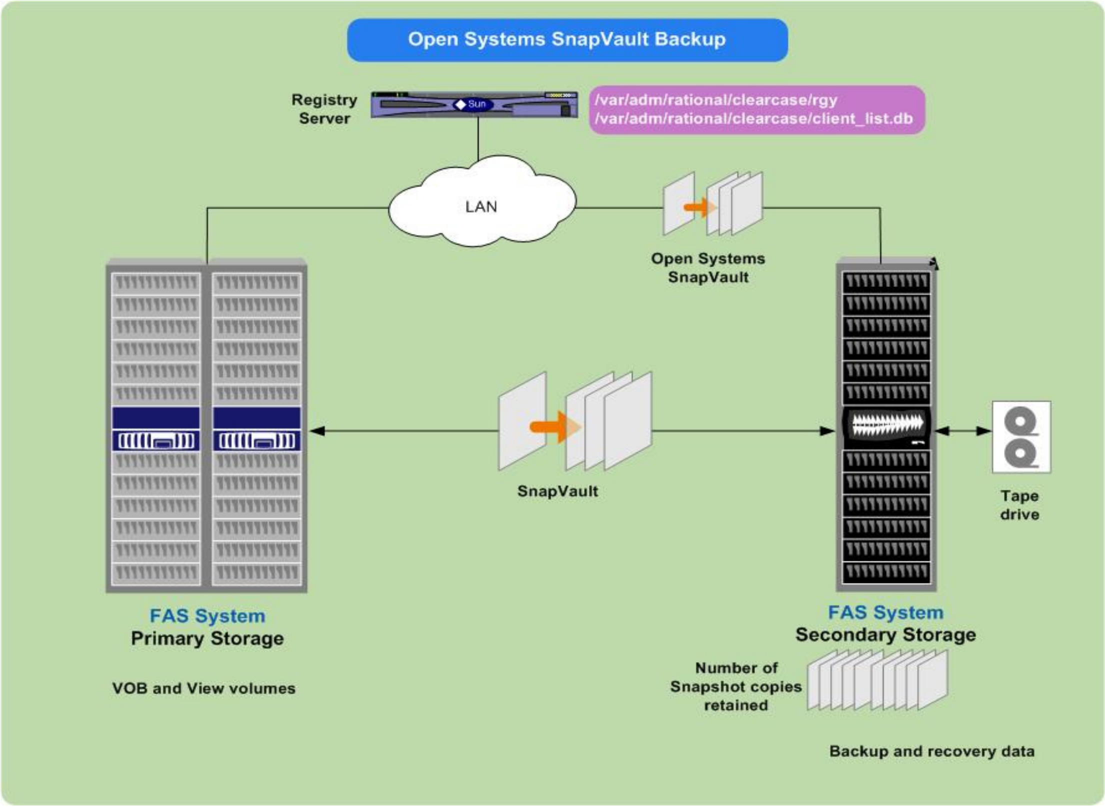

Figure 7) Open Systems SnapVault and Rational ClearCase registry backup

Typically, if files on an open systems platform directory are open when a scheduled SnapVault transfer takes place, they are not backed up until the next scheduled transfer.

The data on the secondary is accessible to users in read-only mode. SnapVault makes a baseline transfer of data (comparable to a full tape backup) and incremental updates thereafter. It is block-based incremental for NetApp storage primaries and file-based incremental for open system primaries.

The SnapVault updates can be configured at specified time intervals through a predetermined schedule or can be initiated manually via a command-line interface.

The SnapVault secondary can be situated either within the local data center or in a remote data center. As part of a multilevel backup strategy, the Snapshot copies can be further backed up to tape using any NDMP-compliant backup software application (discussed in the following sections).

Restoring from SnapVault is simple. The "snapvault restore" command will transfer the specified version of the qtree back to a qtree on the primary system that requests it.

Note: SnapVault backup and restoration are qtree based.

On the other hand, just as the original file was accessed via NFS mount or CIFS share, the SnapVault secondary might be configured with NFS exports and CIFS shares. As long as the destination qtrees are accessible to the users, restoring data from the SnapVault secondary is as simple as copying from a local Snapshot copy.

7 DISK-TO-TAPE BACKUPS

SnapVault provides superior recovery and off-site archival options. One option is to back up data from one or more SnapVault primaries to a remote SnapVault secondary storage system across a wide area network for off-site archival.

The other option is to back up the SnapVault secondary to tape for offline storage in the event of the SnapVault secondary not being available.

This deployment would add a tape backup of the SnapVault secondary storage system and would serve the following purposes:

It enables the storage of an unlimited number of Rational ClearCase backups, while keeping the more recent backups available online in secondary disk storage for a quick recovery, if required. The weekly or monthly Snapshot copies can be further pushed onto the tape.

If tape backup is generated off the SnapVault secondary storage system, the primary and open systems platforms are not subject to performance degradation, system unavailability, and the complexity of direct tape backup of multiple systems. Moreover, the amount of tape media required is substantially reduced. In this instance, tape augments the backup experience.

Another variation to the basic SnapVault deployment is to protect the backed up data on the SnapVault secondary against a site disaster or a secondary storage system failure. The data available on the SnapVault secondary is mirrored to a system configured as a SnapMirror partner, or destination. NetApp SnapMirror provides a fast, flexible solution for mirroring or replicating data over local or wide area networks.

If the SnapVault secondary storage system fails or in case of a site disaster, the SnapMirror destination can be converted to be the secondary storage system and used to continue the SnapVault backup operation with minimum disruption to the environment.

Organizations with multiple development sites, each of which might have a sizable number of developers, might choose to have a SnapVault secondary storage system at each site. But instead of having a tape infrastructure at each site, which would be expensive from both an acquisition and an operational perspective, they could choose to replicate these backups to a central site, where they could periodically back up these consolidated SnapVault copies to tape.

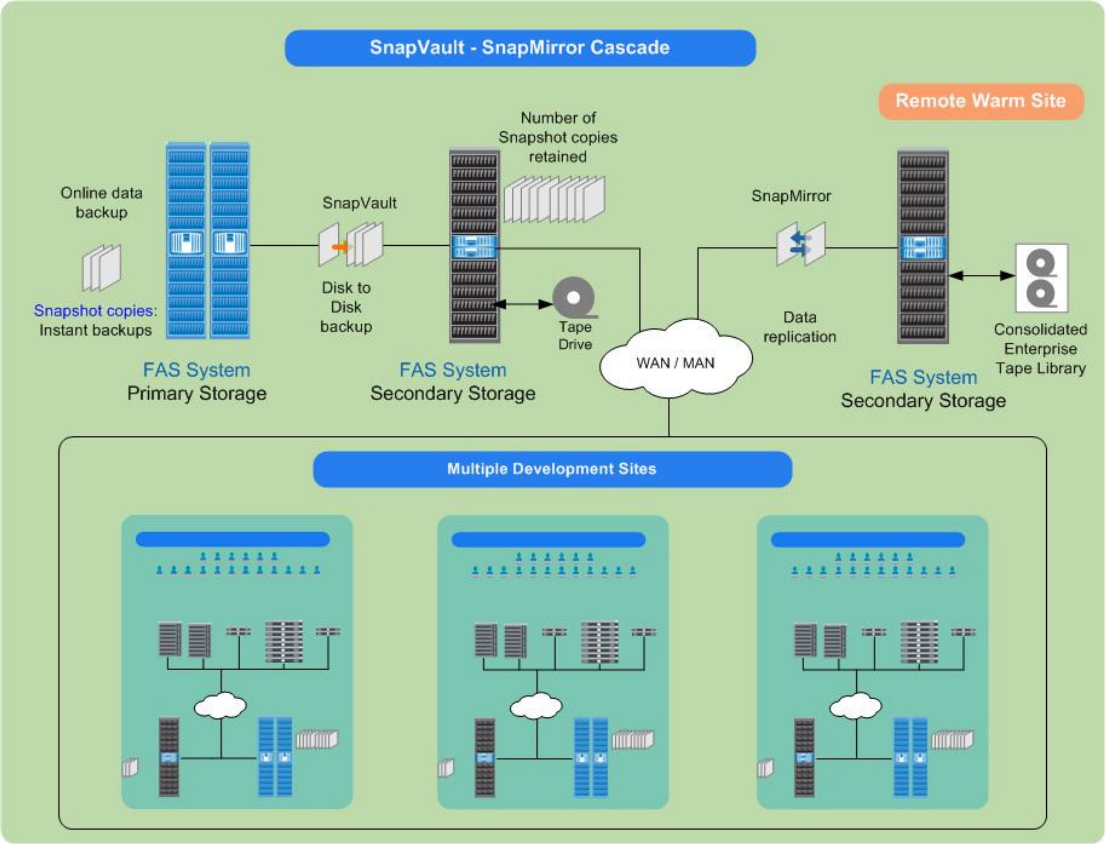

Figure 8) SnapVault and SnapMirror

7.1 Tape Backups

Traditional tape backup still exists as a last line of defense. NetApp native dump and restore commands form the foundation of scalable tape backup strategy. Dump and restore, bundled components of Data ONTAP®, are optimized in its kernel for efficiency and reliability. These commands can be used in simple console-based backups to high-capacity tapes. Or they can be invoked from NDMP-based product for complex enterprise-wide Rational ClearCase backups of multiple volumes on multiple NetApp storage systems.

Dump saves a consistent view of the file system by copying data from a Snapshot copy. Restoring a NetApp dump tape results in a consistent file system or a subset of a file system.

On the other hand, an NDMP (an open standard for centralized control of enterprise-wide data management)-compliant solution separates the flow of backup and control information from the flow of data to and from backup media. These solutions invoke Data ONTAP native dump and restore to back up and restore from NetApp storage.

It can direct a NetApp storage system to back itself up to an attached tape drive, without sending the backup data over the network. NDMP-based solutions are designed to assure data protection and efficient restoration in event of data loss and include many control and management features not available with a NetApp storage system’s native dump and restore commands such as discovery, configuration, scheduling, media management, tape library control, and user interface.

8 SCHEDULING STRATEGIES

Typically, in enterprise-wide Rational ClearCase environments, there might be dependency involved between VOBs or groups of VOBs or Rational ClearQuest databases.

A VOB might be part of a group of VOBs that are connected by hyperlinks.

If UCM and Rational ClearQuest are in use, a PVOB and a Rational ClearQuest database are likely to hold references to each other.

UCM operations such as joining a project, making an activity, and delivering or rebasing a stream are likely to change the data in multiple related databases and, in some cases, the relationships themselves. Even if UCM is not in use, changes in an administrative VOB hierarchy (for example, creating a global type and a local copy) can affect multiple VOBs with a single operation. The longer the interval between the time that the first member group of related databases is backed up and the time that the last member is backed up, the greater the chance that these relationships will be skewed when any of the backups is restored.

A regular backup strategy for a configuration such as the one mentioned earlier would consist of these steps:

1. Lock all the VOBs and the Rational ClearQuest database.
2. Back up all the VOBs and the Rational ClearQuest database.
3. Unlock all the VOBs and the Rational ClearQuest database.

Using NetApp Snapshot technology, all the related VOBs and Rational ClearQuest database volumes can be backed up at a given time, thus making sure of consistency between them. Even if all of the related databases could be backed up in a single event, it is unlikely that they would be restored together. It is far more likely that one or two databases from the set would be restored into an environment where the others were intact, which would again result in skew between the restored and intact databases.

NetApp Snapshot technology complemented with data layout and scheduling strategy will help in maintaining consistency and overcoming issues related to any resulting skew during the restoration process.

Data layout: VOB, view, and ClearQuest database need to be kept in separate volumes and different qtrees within them. The first thing would be to lay the related VOBs in a specific volume or qtree.

Qtree-level granularity is highly recommended in Rational ClearCase environments. All related VOBs (whether they are part of UCM or base ClearCase environment) can be grouped within a single qtree. This will make sure that you back up and restore the related set of VOBs together. You can create up to 4,995 qtrees within a volume.

Scheduling: Snapshot copies across different volumes can be scheduled to be in sync with each other, so that consistency between related data spread across different volumes is maintained.

Scenario:

A software development organization has implemented a UCM-based Rational ClearCase environment consisting of VOB storage and Rational ClearQuest database. All the related VOBs are stored in separate qtrees (ntapucmstore) within a volume. The ClearQuest database is stored in a separate volume (ntapcqstore).

The scheduling is configured in such a manner that Snapshot copies for a volume containing a set of related UCM VOBs (ntapucmvobstore) and a volume containing a Rational ClearQuest database (ntapcqstore) can be scheduled to happen at the same given time.

➢ Snapshot scheduling:

  Default schedule:

  - snap sched vol*x 0 2 6 @8,12,16, 20*

Scheduling Snapshot copies for all relevant volumes at the same time:

- snap sched ntapucmvobstore *0 2 6 @8,12,16,20*
- snap sched ntapcqstore *0 2 6 @8,12,16,20*

➢ SnapVault scheduling:

Scheduling Snapshot copies on the SnapVault primary storage system:

- snapvault snap sched ntapucmvobstore sv_nightly 2@mon-fri@19
- snapvault snap sched ntapcqstore sv_nightly 2@mon-fri@19

Scheduling Snapshot copies on the SnapVault secondary storage system:

- snapvault snap sched –x ntapucmvobstore1 sv_nightly 60@mon-fri@20
- snapvault snap sched –x ntapcqstore sv_nightly 60@mon-fri@20

Note: The recommended granularity of Rational ClearCase data storage is qtree. SnapVault base transfer is configured to happen at the qtree level. Scheduling, however, is done on the volume level. SnapVault setup details are given in the addendum.

Restored data (ntapucmvobstore and ntapcqstore) is consistent and in sync since the backup of the said volumes has happened at a point in time. Besides, the entire qtree, consisting of a set of interdependent VOBs, can be recovered (instead of restoring a single VOB) to mitigate the problem of probable skew between them.

9 FILE SYSTEM REQUIREMENTS

In Rational ClearCase environments, VOB locking and unlocking provide the application-level quiescing. SAN (iSCSI or FC)-based protocols also call for requirements for freezing all the I/O operations to LUNs in the file system. Besides, the file system caches must be first committed before initiating any backup operations (triggering the NetApp Snapshot copies), resulting in truly consistent backups.

NetApp provides the SnapDrive® utility, which addresses the above-mentioned functionality for SAN deployments. Moreover, SnapDrive helps administrators to provision storage and manage it directly from the host (based on UNIX or Windows), enable and define backup policies, and resize storage on fly.

In a nutshell, SnapDrive simplifies storage and data management by using the host operating system and NetApp technologies and hiding the complexity of the steps that must be executed on both the storage and host system, primarily resulting in consistent data Snapshot copies and rapid application recovery from them.

10 CONCLUSION

NetApp Snapshot technology has revolutionized the backup and recovery scenarios in Rational ClearCase environments, resulting in drastic reduction in VOB lock times, addressing larger data footprints and frequent full backups.

NetApp provides a complete suite of cost-effective tools that allows you to back up a Rational ClearCase environment within the shortest possible time frame. Data is highly available, so users are able to work around the clock to meet time-to-market objectives.

Moreover, the restoration procedures are as simple as the backup, resulting in dramatic improvements in RPO and RTO, as well as huge savings in operational costs.

11 ADDENDUM (COMMAND REFERENCE)

11.1 NetApp Snapshot Copies (Online)

Assumptions:

- In the following examples the volume storing Rational ClearCase VOBs and views is named as ntapvobstore2.
- The disk usage numbers shown in the below-mentioned command listing are just for illustration.
- Appropriate licenses are enabled on the NetApp storage.

Check volume status:

```
> vol status
Volume State     Status            Options
 vol0 online     raid4, trad       root
 vol1 online     raid_dp, flex
 ntapvobstore1 online     raid_dp, flex     guarantee=file
 ntapvobstore online     raid_dp, flex
 ntapvobstore2 online     raid_dp, flex
```

Display the Snapshot schedule for a volume:

```
> snap sched ntapvobstore2
Volume ntapvobstore2: 0 2 6@8,12,16,20
```

Display Snapshot schedule for all volumes:

```
> snap sched
Volume vol0: 0 2 6@8,12,16,20
Volume ntapvobstore: 0 2 6@8,12,16,20
Volume ntapvobstore1: 0 2 6@8,12,16,20
Volume ntapvobstore2: 0 2 6@8,12,16,20
Volume vol1: 0 2 6@8,12,16,20
```

Disable and enable client access to Snapshot copies in a volume:

```
> vol options ntapvobstore2 nosnapdir on
> vol options ntapvobstore2 nosnapdir off
```

List the VOBs available in the said volume:

```
>cleartool lsvob |grep /ccasestore/ntapvobstore2

* /vobs/ironcity       /ccasestore/ntapvobstore2/ironcity.vbs public

* /vobs/gb            /ccasestore/ntapvobstore2/gb.vbs public
```

VOB locking:

```
>cleartool lock vob: <vobname>
```

Initializing NetApp Snapshot copies:

```
> snap create ntapvobstore2 ictest
```

VOB unlocking:

```
>cleartool unlock: <vobname>
```

List the Snapshot copies for the said volume:

```
> snap list ntapvobstore2
Volume ntapvobstore2
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 15 09:43 vobtest
  0% ( 0%)    0% ( 0%)  May 15 08:00 hourly.0
  0% ( 0%)    0% ( 0%)  May 15 00:00 nightly.0
  1% ( 0%)    0% ( 0%)  May 14 20:00 hourly.1
  1% ( 0%)    0% ( 0%)  May 14 16:00 hourly.2
  2% ( 0%)    0% ( 0%)  May 14 12:00 hourly.3
  2% ( 0%)    0% ( 0%)  May 14 08:00 hourly.4
  2% ( 0%)    0% ( 0%)  May 14 00:00 nightly.1
  8% ( 6%)    0% ( 0%)  May 13 20:00 hourly.5
  36% (31%)    0% ( 0%)  May 13 06:13 ictest
```

Check the Snapshot disk usage:

```
> df /vol/ntapvobstore2

Filesystem              kbytes       used      avail capacity Mounted on
/vol/ntapvobstore2/  671088640     337532  670751108       0%  /vol/ntapvobstore2/
/vol/ntapvobstore2/.snapshot  167772160     185756  167586404       0%  /vol/ntapvobstore2/.snapshot
```

Display the Snapshot schedule for a volume:

```
> snap sched ntapvobstore2
Volume ntapvobstore2: 0 2 6@8,12,16,20
```

Disable and enable automatic Snapshot copies:

```
> vol options ntapvobstore2 nosnap on
> snap sched ntapvobstore2 0 0 0
> snap sched ntapvobstore2
Volume ntapvobstore2: 0 0 0
       > vol options ntapvobstore2 nosnap off
```

Restoring volume from the Snapshot copy:

```
> snap restore ntapvobstore2 vobtest
```

Restoring selected files online:

Select the files from the path /ccasestore/ntapvobstore2/.snapshot/vobtest/gb.vbs.

11.2 Procedure to Restore View Private File

Rational ClearCase view private files can be restored from the appropriate volume Snapshot copy. Steps and associated command listing for the same are given below.

Mandatory:

- Storage device should support NFS V3 and long filenames.

- Views storage directories reside on NetApp volume /qtree.
- Snapshot copies need to be enabled on the said volume.

Assumptions:

- Volume name/qtree name: ntapvobstore2
- Latest Snapshot copy name: vobtest
- The view storage directory name and private filename to be restored: ntapview.vws and ppe.txt

How do I create dynamic views residing on NetApp storage?

```
>cleartool mkview –tag <dynamic view tag> -host <hostname> -hpath < storage path> -gpath < storage
path> < dynamic view storage path>
```

The hpath, gpath, and the dynamic view path will be same if the view storage resides on the NetApp storage**.**

What are view private files?

1) These are the files that have been newly created that are yet to be checked in the VOB. They are residing in your own workspace (that is, view).

2) Those have been checked out to the view (from the VOB).

Scenario: Usually, the developer creates a view private file (through 1 or 2 mentioned above), and those files lie in the developer's workspace for n number of days. The developer might be working on that file for those many days (EDA design file might take more time to develop) or might have forgotten about it. Or the developer intentionally wants some files in views that should not have been part of the ClearCase network.

Need to recover view private file?

In all the above scenarios, there are chances of accidental deletion of these view private files. In such instances (if we have the view storage residing on NetApp storage) we can leverage the NetApp Snapshot technology to recover them.

Identify the said Snapshot copy from the .snapshot directory:

```
> cd  /ccasestore/ntapvobstore2/.snapshot

> pwd
/ccasestore/ntapvobstore2/.snapshot

> ls -lu
total 152
drwxr-xr-x   8 root     root        4096 May 15 13:30 hourly.0
drwxr-xr-x   8 root     root        4096 May 15 01:30 hourly.1
drwxr-xr-x   8 root     root        4096 May 14 21:30 hourly.2
drwxr-xr-x   8 root     root        4096 May 14 17:30 hourly.3
drwxr-xr-x   8 root     root        4096 May 14 13:30 hourly.4
drwxr-xr-x   8 root     root        4096 May 14 01:30 hourly.5
drwxr-xr-x   7 root     root        4096 May 13 11:43 ictest
drwxr-xr-x   8 root     root        4096 May 15 15:13 vobtest
```

Identify the required view storage directory:

```
>cd vobtest
>pwd
/ccasestore/ntapvobstore2/.snapshot/vobtest
>ls -al
total 64
drwxr-xr-x   8 root     root        4096 May 13 13:51 .
drwxrwxrwx  21 root     root        4096 May 15 15:13 ..
drwxr-xr-x   8 root     root        4096 May 13 14:46 gb.vbs
drwxr-xr-x   8 root     root        4096 May 13 11:10 ironcity.vbs
```

```
drwxr-xr-x   6 root     root        4096 May 13 20:51 ntapview.vws
```

Get into the appropriate view directory:

```
>cd ntapview.vws
>pwd
/ccasestore/ntapvobstore2/.snapshot/vobtest/ntapview.vws
>ls -al

total 80
drwxr-xr-x   6 root     root        4096 May 13 20:51 .
drwxr-xr-x   8 root     root        4096 May 13 13:51 ..
-r--r--r--   1 root     root         310 May 13 14:55 .access_info
-rw-r--r--   1 root     root         221 May 13 11:07 .compiled_spec
-r--r--r--   1 root     root          12 May 13 11:07 .hostname
drwx------   2 root     root        4096 May 15 04:47 .identity
drwxr-xr-x  59 root     root        4096 May 15 04:47 .s
-r--r--r--   1 root     root         104 May 13 11:07 .view
drwxr--r--   3 root     root        4096 May 15 04:47 admin
-rw-r--r--   1 root     root          44 May 13 11:07 config_spec
drwxr-xr-x   2 root     root        4096 May 15 04:47 db
-r--r--r--   1 root     root         373 May 13 11:07 readme.txt
-rw-r--r--   1 root     root           2 May 13 18:31 view_db.state
```

The .s file stores all the source containers:

```
>cd .s
>pwd
/ccasestore/ntapvobstore2/.snapshot/vobtest/ntapview.vws/.s
```

ClearCase view container files are seen:

```
>ls -al
total 472
drwxr-xr-x  59 root     root        4096 May 15 04:47 .
drwxr-xr-x   6 root     root        4096 May 13 20:51 ..
drwxr-xr-x   2 root     root        4096 May 13 11:07 00000
drwxr-xr-x   2 root     root        4096 May 13 11:07 00001
drwxr-xr-x   2 root     root        4096 May 13 11:07 00002
drwxr-xr-x   2 root     root        4096 May 13 11:07 00003
drwxr-xr-x   2 root     root        4096 May 13 11:07 00004
drwxr-xr-x   2 root     root        4096 May 13 11:07 00005
drwxr-xr-x   2 root     root        4096 May 13 11:07 00006
drwxr-xr-x   2 root     root        4096 May 13 11:07 00007
drwxr-xr-x   2 root     root        4096 May 13 11:07 00008
drwxr-xr-x   2 root     root        4096 May 13 11:07 00009
drwxr-xr-x   2 root     root        4096 May 13 11:07 00010
drwxr-xr-x   2 root     root        4096 May 13 11:07 00011
drwxr-xr-x   2 root     root        4096 May 13 11:07 00012
drwxr-xr-x   2 root     root        4096 May 13 11:07 00013
drwxr-xr-x   2 root     root        4096 May 13 11:07 00014
drwxr-xr-x   2 root     root        4096 May 13 11:07 00015
drwxr-xr-x   2 root     root        4096 May 13 11:07 00016
drwxr-xr-x   2 root     root        4096 May 13 11:07 00017
drwxr-xr-x   2 root     root        4096 May 13 11:07 00018
drwxr-xr-x   2 root     root        4096 May 13 11:07 00019
drwxr-xr-x   2 root     root        4096 May 13 11:07 00020
drwxr-xr-x   2 root     root        4096 May 13 11:07 00021
drwxr-xr-x   2 root     root        4096 May 13 11:07 00022
drwxr-xr-x   2 root     root        4096 May 13 11:07 00023
```

```
drwxr-xr-x   2 root     root        4096 May 13 11:07 00024
drwxr-xr-x   2 root     root        4096 May 13 11:07 00025
drwxr-xr-x   2 root     root        4096 May 13 11:07 00026
drwxr-xr-x   2 root     root        4096 May 13 11:07 00027
drwxr-xr-x   2 root     root        4096 May 13 11:07 00028
drwxr-xr-x   2 root     root        4096 May 13 11:07 00029
drwxr-xr-x   2 root     root        4096 May 13 11:07 00030
drwxr-xr-x   2 root     root        4096 May 13 11:07 00031
drwxr-xr-x   2 root     root        4096 May 13 11:07 00032
drwxr-xr-x   2 root     root        4096 May 13 11:07 00033
drwxr-xr-x   2 root     root        4096 May 13 11:07 00034
drwxr-xr-x   2 root     root        4096 May 13 11:07 00035
drwxr-xr-x   2 root     root        4096 May 13 14:54 00036
drwxr-xr-x   2 root     root        4096 May 13 11:07 00037
drwxr-xr-x   2 root     root        4096 May 13 11:07 00038
drwxr-xr-x   2 root     root        4096 May 13 11:07 00039
drwxr-xr-x   2 root     root        4096 May 13 11:07 00040
drwxr-xr-x   2 root     root        4096 May 13 11:07 00041
drwxr-xr-x   2 root     root        4096 May 13 11:07 00042
drwxr-xr-x   2 root     root        4096 May 13 11:07 00043
drwxr-xr-x   2 root     root        4096 May 13 11:07 00044
drwxr-xr-x   2 root     root        4096 May 13 11:07 00045
drwxr-xr-x   2 root     root        4096 May 13 11:07 00046
drwxr-xr-x   2 root     root        4096 May 13 11:07 00047
drwxr-xr-x   2 root     root        4096 May 13 11:07 00048
drwxr-xr-x   2 root     root        4096 May 13 11:07 00049
drwxr-xr-x   2 root     root        4096 May 13 11:07 00050
drwxr-xr-x   2 root     root        4096 May 13 11:07 00051
drwxr-xr-x   2 root     root        4096 May 13 11:07 00052
drwxr-xr-x   2 root     root        4096 May 13 11:07 00053
drwxr-xr-x   2 root     root        4096 May 13 11:07 00054
drwxr-xr-x   2 root     root        4096 May 13 11:07 00055
drwxr-xr-x   2 root     root        4096 May 13 11:07 00056
-rwx------   1 root     root           0 May 13 11:07 view_db.crs_file
```

Identify the view private file that you want to restore back:

```
>du -a * |grep ppe
    16      00004/8000001e482c17b3ppe.txt
```

Rational ClearCase prefixes the relevant container notations with the said file.

Copy the view private file to the view /other location:

```
>cp 00004/8000001e482c17b3ppe.txt /ccasestore/restored
>cd /ccasestore/restored
>ls -al
total 33
drwxr-xr-x   3 root     root        4096 May 15  2008.
dr-xr-xr-x   3 root     root           3 May 15 16:18 ..
-rw-r--r--   1 root     root        6958 May 15  2008 8000001e482c17b3ppe.txt
```

Rename the view private file to original:

```
>mv 8000001e482c17b3ppe.txt   ppe.txt
>ls -al
total 33
drwxr-xr-x   3 root     root        4096 May 15  2008 .
dr-xr-xr-x   3 root     root           3 May 15 16:18 ..
-rw-r--r--   1 root     root        6958 May 15  2008 ppe.txt
```

Verify the contents of the view private file:

```
>cat ppe.txt
```

11.3 Mirroring (SnapMirror)

Assumptions:

- Appropriate licenses are enabled.
- The disk usage numbers shown in the below mentioned command listing are just for illustration.
- Volume name at the SnapMirror primary (on NetApp storage) is ntapvobstore2.
- Volume name at the SnapMirror secondary (on NetApp storage) is smirror.

Display volume status at source:

```
> vol status
        Volume State      Status            Options
          vol0 online     raid4, trad       root
          vol1 online     raid_dp, flex
  ntapvobstore1 online     raid_dp, flex     guarantee=file
  ntapvobstore online     raid_dp, flex
  ntapvobstore2 online     raid_dp, flex
```

Enable SnapMirror and configure destination at source:

```
> options snapmirror.enable on

> options snapmirror.enable
snapmirror.enable            on

> options snapmirror.access host =10.75.68.112

> options snapmirror.access
snapmirror.access            host=10.75.68.112
```

SnapMirror status at source (after SnapMirror initialization):

```
> snapmirror status
Snapmirror is on.
Source                    Destination          State         Lag       Status
FILER13:ntapvobstore2  FILER12:smirror  Source        00:01:31   Idle
```

Snapshot copy list at source (**after SnapMirror initialization):**

```
> snap list ntapvobstore2
Volume ntapvobstore2
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 20 07:35  FILER12(0118042259)_smirror.1 (snapmirror)
  0% ( 0%)    0% ( 0%)  May 20 00:00  nightly.0
  1% ( 0%)    0% ( 0%)  May 19 20:00  hourly.0
  1% ( 0%)    0% ( 0%)  May 19 16:00  hourly.1
  2% ( 0%)    0% ( 0%)  May 19 12:00  hourly.2
  2% ( 0%)    0% ( 0%)  May 19 08:00  hourly.3
  2% ( 0%)    0% ( 0%)  May 19 00:00  nightly.1
  2% ( 0%)    0% ( 0%)  May 18 20:00  hourly.4
  2% ( 0%)    0% ( 0%)  May 18 16:00  hourly.5
  3% ( 0%)    0% ( 0%)  May 15 11:17  vobtest2
  3% ( 0%)    0% ( 0%)  May 15 09:43  vobtest
 36% (34%)    0% ( 0%)  May 13 06:13  ictest
```

Snapshot list at source (after SnapMirror update):

```
> snap list ntapvobstore2
Volume ntapvobstore2
working...

  %/used       %/total  date          name
  ----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 20 07:40  FILER12(0118042259)_smirror.2 (snapmirror)
  0% ( 0%)    0% ( 0%)  May 20 00:00  nightly.0
  1% ( 0%)    0% ( 0%)  May 19 20:00  hourly.0
  1% ( 0%)    0% ( 0%)  May 19 16:00  hourly.1
  2% ( 0%)    0% ( 0%)  May 19 12:00  hourly.2
  2% ( 0%)    0% ( 0%)  May 19 08:00  hourly.3
  2% ( 0%)    0% ( 0%)  May 19 00:00  nightly.1
  2% ( 0%)    0% ( 0%)  May 18 20:00  hourly.4
  3% ( 0%)    0% ( 0%)  May 18 16:00  hourly.5
  4% ( 0%)    0% ( 0%)  May 15 11:17  vobtest2
  4% ( 0%)    0% ( 0%)  May 15 09:43  vobtest
 36% (34%)    0% ( 0%)  May 13 06:13  ictest
```

Enable SnapMirror at destination:

```
> options snapmirror.enable on

> options snapmirror.enable
snapmirror.enable            on
```

Volume creation and status at destination:

```
> vol create smirror rational 850g
Creation of volume 'smirror' with size 850g on containing aggregate 'rational' has completed.

> vol status ( smirror volume created)
         Volume State      Status            Options
             vol0 online     raid_dp, flex     root
           home online     raid_dp, flex
        registry online     raid_dp, flex
         smirror online     raid_dp, flex

> vol restrict smirror
Volume 'smirror' is now restricted.

> vol status ( post restriction)
         Volume State      Status            Options
             vol0 online     raid_dp, flex     root
           home online     raid_dp, flex
        registry online     raid_dp, flex
        smirror restricted raid_dp, flex

> vol status smirror ( shows the containing aggregate)
         Volume State      Status            Options
        smirror restricted raid_dp, flex
                Containing aggregate: 'rational'
```

SnapMirror initialization at the destination:

```
> snapmirror initialize -S FILER13:ntapvobstore2 smirror Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.
```

SnapMirror status and Snapshot copy list at the destination:

```
> snapmirror status
Snapmirror is on.
Source                  Destination         State        Lag       Status
FILER13:ntapvobstore2  FILER12:smirror  Snapmirrored   -00:00:11  Idle

> vol status
        Volume       State                  Status           Options
        smtest       restricted raid_dp, flex
          vol0        online      raid_dp, flex     root
          home        online     raid_dp, flex
      registry        online     raid_dp, flex
       smirror        online     raid_dp, flex     snapmirrored=on,
                                 snapmirrored      fs_size_fixed=on

> snap list smirror
Volume smirror
working...

  %/used       %/total  date         name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 20 07:35  FILER12(0118042259)_smirror.1
  0% ( 0%)    0% ( 0%)  May 20 00:00  nightly.0
  0% ( 0%)    0% ( 0%)  May 19 20:00  hourly.0
  1% ( 0%)    0% ( 0%)  May 19 16:00  hourly.1
  2% ( 0%)    0% ( 0%)  May 19 12:00  hourly.2
  2% ( 0%)    0% ( 0%)  May 19 08:00  hourly.3
  2% ( 0%)    0% ( 0%)  May 19 00:00  nightly.1
  2% ( 0%)    0% ( 0%)  May 18 20:00  hourly.4
  2% ( 0%)    0% ( 0%)  May 18 16:00  hourly.5
  3% ( 0%)    0% ( 0%)  May 15 11:17  vobtest2
  3% ( 0%)    0% ( 0%)  May 15 09:43  vobtest
 36% (34%)    0% ( 0%)  May 13 06:13  ictest
```

SnapMirror updates:

```
> snapmirror update -S FILER13:ntapvobstore2 smirror
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.
```

Snapshot copy list and volume properties postupdate at the destination:

```
> snap list smirror
Volume smirror
working...

  %/used       %/total  date         name
----------  ----------  ------------  --------
  0% ( 0%)    0% ( 0%)  May 20 07:40  FILER12(0118042259)_smirror.2
  0% ( 0%)    0% ( 0%)  May 20 07:35  FILER12(0118042259)_smirror.1
  0% ( 0%)    0% ( 0%)  May 20 00:00  nightly.0
  1% ( 0%)    0% ( 0%)  May 19 20:00  hourly.0
  1% ( 0%)    0% ( 0%)  May 19 16:00  hourly.1
  2% ( 0%)    0% ( 0%)  May 19 12:00  hourly.2
  2% ( 0%)    0% ( 0%)  May 19 08:00  hourly.3
  2% ( 0%)    0% ( 0%)  May 19 00:00  nightly.1
  2% ( 0%)    0% ( 0%)  May 18 20:00  hourly.4
  2% ( 0%)    0% ( 0%)  May 18 16:00  hourly.5
   3% ( 0%)    0% ( 0%)  May 15 11:17  vobtest2
  3% ( 0%)    0% ( 0%)  May 15 09:43  vobtest
 36% (34%)    0% ( 0%)  May 13 06:13  ictest
```

```
> vol options smirror
nosnap=off, nosnapdir=off, minra=off, no_atime_update=off, nvfail=off,
ignore_inconsistent=off, snapmirrored=on, create_ucode=off,
convert_ucode=off, maxdirsize=335462, schedsnapname=ordinal,
fs_size_fixed=on, compression=off, guarantee=volume, svo_enable=off,
svo_checksum=off, svo_allow_rman=off, svo_reject_errors=off,
no_i2p=off, fractional_reserve=100, extent=off, try_first=volume_grow,
snapshot_clone_dependency=off
```

Restoring from the mirrored copy:

```
> snapmirror quiesce smirror
```

```
> snapmirror break smirror
snapmirror break: Destination smirror is now writable.
Volume size is being retained for potential snapmirror resync. If you would like to grow the volume and
do not expect to resync, set vol option fs_size_fixed to off.snapmirror break /vol/smirror.
```

Note: With setting up a reverse SnapMirror relationship between the new production volume and the former, you will able to sync back the data to the original system. If the amount of data is less, a mirrored volume can be mounted, and the required data can be copied back to the production volume.

11.4 Disk to Disk (SnapVault)

Assumptions:

- Appropriate licenses are enabled.
- Volume name/qtree name at the SnapVault primary (on NetApp storage) is ntapvobstore2/project1.
- Volume name/qtree name at the SnapVault secondary (on NetApp storage) is ntapvobstore2/sv_project1.

SnapVault initialization (primary side)

```
> ndmpd status
ndmpd ON.
No ndmpd sessions active.

> options snapvault.enable on

> options snapvault.access host= 10.75.68.112 (added the NetApp secondary)
> options snapvault.access
snapvault.access            host=10.75.68.112
```

SnapVault initialization (secondary side):

```
> ndmpd status
ndmpd ON.
No ndmpd sessions active.

> options snapvault.enable on

> options snapvault.access host= 10.75.68.113 (added the NetApp primary)
> options snapvault.access
snapvault.access            host=10.75.68.113
```

SnapVault baseline transfer (to be initiated from secondary side):

```
> snapvault start -S FILER13:/vol/ntapvobstore2/project1 FILER12:/vol/ntapvobstore2/sv_project1
Transfer started.
```

```
Monitor progress with 'snapvault status' or the snapmirror log.

> snapvault status                                  Snapvault secondary is ON.
Source                              Destination                         State         Lag       Status
FILER13:/vol/ntapvobstore2/project1     FILER12:/vol/ntapvobstore2/sv_project1  Snapvaulted    -
00:00:27  Idle

> qtree status
Volume   Tree     Style Oplocks  Status
-------- -------- ----- -------- ---------
vol0             unix  enabled  normal
home             unix  enabled  normal
registry         unix  enabled  normal
registry ccreg    unix  disabled snapvaulted
registry sv_project1 unix  enabled  snapvaulted
```

Schedule SnapVault updates:

```
>snapvault snap sched ntapvobstore2 sv_nightly 2@mon-fri@19 (on the primary side: 2 Snapshot
copies retained)
>snapvault snap sched –x ntapvobstore2 sv_nightly 90@mon-fri@20 (on the secondary side: 90 nightly
Snapshot copies retained)
```

Restoration from SnapVault:

```
>snap vault restore –s sv_nightly.4 FILER12:/vol/ntapvobstore2/svdqtree
FILER13:/vol/ntapvobstore2/svbqtree (restores the entire qtree)
```

Note: The SnapVault secondary can be mounted and a single VOB can be copied back to the production.

11.5 Open Systems SnapVault

Assumptions:

- Rational ClearCase registry server host name based on UNIX: RAR1(10.75.68.12)
- Directory to be backed up using Open Systems SnapVault (on SnapVault primary): /var/adm/rational/clearcase
- Open Systems SnapVault agent installed on the SnapVault primary: the ClearCase registry server
- Volume name/qtree name at the SnapVault secondary (on NetApp storage): /registry/ccreg
- Appropriate licenses are enabled

On the SnapVault secondary:

```
>options snapvault.enable on
>options ndmpd.enable on
>options snapvault.access host=10.75.68.12
```

Initialize the baseline copy:

```
> snapvault start -S RAR1:/var/adm/rational/clearcase FILER12:/vol/registry/ccreg
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
> snapvault status -l
Snapvault secondary is ON.

Source:              RAR1:/var/adm/rational/clearcase
Destination:         FILER12:/vol/registry/ccreg
Status:              Idle
```

```
Progress:              -
State:                Snapvaulted
Lag:                  00:22:30
Mirror Timestamp:       Mon Apr 14 10:17:25 GMT 2008
Base Snapshot:         FILER12(0118042259)_registry-base.0
Current Transfer Type:  -
Current Transfer Error: -
Contents:              Replica
Last Transfer Type:     Initialize
Last Transfer Size:     1612 KB
Last Transfer Duration: 00:00:04
Last Transfer From:     RAR1:/var/adm/rational/clearcase
```

Scheduling the updates:

```
> snapvault snap sched -x registry sv_hourly 20@mon-fri@7-22
> snapvault snap sched
create registry  0@-
xfer   registry sv_hourly 20@mon-fri@7-22
```

Monitoring the updates:

```
> snap list registry
Volume registry
working...


  %/used       %/total  date          name
----------  ----------  ------------  --------
  5% ( 5%)    0% ( 0%)  Apr 14 10:39 FILER12 (0118042259) _registry-base.0 (busy,snapvault)

> qtree status
Volume   Tree     Style Oplocks  Status
-------- -------- ----- -------- ---------
vol0             unix  enabled  normal
ntapvobstore1         unix  enabled  normal
home             unix  enabled  normal
ntapvobstore2         unix  enabled  normal
ntapviewstore1         unix  enabled  normal
ntapviewstore2         unix  enabled  normal
registry          unix  enabled  normal
registry ccreg    unix  disabled snapvaulted

> snap list registry
Volume registry
working...

  %/used       %/total  date          name
----------  ----------  ------------  --------
  4% ( 4%)    0% ( 0%)  Apr 14 14:05  sv_hourly.0
  9% ( 5%)    0% ( 0%)  Apr 14 14:05  FILER12(0118042259)_registry-base. 0 (busy,snapvault)
 17% (10%)    0% ( 0%)  Apr 14 13:05  sv_hourly.1
 24% (10%)    0% ( 0%)  Apr 14 12:05  sv_hourly.2
 30% (10%)    0% ( 0%)  Apr 14 12:00  hourly.0
 32% ( 4%)    0% ( 0%)  Apr 14 11:05  sv_hourly.3
```

Restoration procedure (at the open system platform console):

```
snapvault restore –s sv_hourly.0 –S FILER12:/vol/registry/ccreg   /var/adm/rational/clearcase
```

11.6 Disk to Tapes

Assumptions:

- NDMP services are initialized.

Write Snapshot copies to tape:

```
>dump 0fb rst0a 63 /vol/ntapvobstore2/.snapshot/snaptape.0
```

List files on the tape:

```
>restore tf rst0a
```

Restoration from tape:

```
>restore rf rst0a (entire volume)
```

```
>restore xf rst0a /ntapvobstore2/src.vbs (single VOB)
```

www.netapp.com