Storage Replication Solutions with the IBM TotalStorage™ Enterprise Storage Server® (including the DS 6000 and DS 8000 families) and VERITAS Software

<u>Jointly Authored by IBM and VERITAS</u>

Phil Mitchell (IBM) [mphil@us.ibm.com]

Girish Venkatachaliah (IBM) [girishv@us.ibm.com]

Ibrahim "Ibby" Rahmani (VERITAS) [irahmani@veritas.com]

Sheri Atwood (VERITAS) [satwood@veritas.com]

Jim Edmonds (VERITAS) [jim.edmonds@veritas.com]

<u>Updated November, 2004</u>

Jeff Huxtable (IBM) [huxtable@us.ibm.com]

Ibrahim "Ibby" Rahmani (VERITAS) [irahmani@veritas.com]

TABLE OF CONTENTS

1.0 Executive Summary

In the current information age, the amount of data is exponentially increasing, and storing and managing this vast amount of data is critically important and challenging for customers. More and more businesses are transforming their operations into information-centric models and demanding uninterrupted availability of data. In addition to tremendous growth in the amount and the value of data, businesses are unwilling to tolerate outages for even routine maintenance and backups.

Unplanned events that inhibit the availability of this data can seriously damage business operations. Any permanent data loss, from disaster or any other source, may have serious negative consequences for the continued viability of a business. Therefore, before disaster strikes, companies must be prepared to minimize or eliminate data loss.

As part of a disaster recovery strategy, companies are increasingly replicating their data to a remote site to protect against numerous disaster scenarios. Customers are further leveraging these distributed copies to minimize the downtime needed to backup the data.

IBM and VERITAS have joined forces to test and identify interoperable advanced copy solutions based on the IBM's TotalStorage™ Enterprise Storage Server® (ESS) and VERITAS' Volume Replicator™ (VVR). This effort builds upon previous collaborative efforts between the two companies to test and identify interoperable storage solutions using both companies' products.

IBM and VERITAS each offer storage solutions for a wide range of environments, from workgroups to enterprise. Both companies are focused on a cross-platform strategy, offering customers' flexibility in selecting their preferred operating platforms and storage networking hardware vendors. The storage solutions are aimed at helping customers experience rapid deployment, interoperability, broad scalability and support for the latest technologies. Cooperative support agreements are in place between the two companies to assist customers who have deployed products from both companies in a mixed environment. These cooperative support agreements broaden each company's ability to provide enterprise-class disaster recovery services to customers, and support faster problem resolution.

IBM and VERITAS are working together to provide greater benefits to end-users by promoting greater interoperability in the solutions that each bring to the marketplace. This paper offers an overview of the replication technologies available today to assist customers in meeting their disaster recovery needs. It addresses the challenges of replicating data, presents best-of-breed solution options based on the ESS and VVR, and the relative merits of each solution option.

The ESS is targeted at open system environments and is designed to provide tremendous reliability, ease of management and exceptional performance. Further, the

ESS provides a rich set of Advanced Copy Services to enable fast backups, synchronous copies, remote mirroring and off-host processing for disaster recovery.

VERITAS is a market leader in storage software. VERITAS Volume Replicator (VVR) is the leading host-based replication software that can reliably, efficiently and consistently mirror data to remote locations over any IP network for maximum business continuity. VVR provides an elegant, flexible, and storage-independent solution to deliver true disaster recovery when data availability is paramount.

2.0 Customer Requirements

As companies become more reliant on information systems and data, loss of data or loss of access to data could significantly disrupt the business. No longer is data part of the business - it is the business.

Gartner stated, "Two out of five enterprises that experience a disaster go out of business within five years. Business continuity plans and disaster recovery services ensure continuing viability." (Gartner, *Disaster Recovery Plans and Systems Are Essential*, by Roberta Witty, Donna Scott, 12 September 2001).

Customers are implementing procedures and technologies to proactively manage recovery and availability. Customers require high availability, minimal application downtime for maintenance, and the ability to perform data backups with minimal or no application outage in order to support the viability of the business. Traditionally, tape backups have been the standard for disaster recovery. Even with the latest high-end tape based solution, a restore procedure from tape could typically take anywhere from hours to days and inevitably result in hours of data loss.

Customers often have different disaster recovery requirements for each business application. Some scenarios may require that the data changes at the primary site be immediately reflected at the disaster recovery site. Certain scenarios may not require having such stringent requirements, and it may be acceptable if the disaster recovery site lags the primary during periods of time where the change rate may exceed the available bandwidth. Certain scenarios may require that the disaster recovery site be consistent at particular specified points in time, so that customers may protect against logical errors or allow for remote backups.

In order to choose the appropriate disaster recovery technologies, customers need to determine their most critical processes and the associated Recovery Point Objective (RPO) and Recovery Time Objective (RTO) requirements. RPO is the point in time to which applications data must be recovered to resume business transaction. RTO is the maximum elapsed time allowed before lack of business function severely impacts an organization. Additionally, customers need to determine the Network Recovery Objective (NRO) as well. NRO indicates the time required to recover or fail over network operations. Comprehensive network failover planning is another dimension to data recovery in a disaster recovery scenario.

3.0 Replication

Data replication provides a solution to the customer requirements described in the previous section. Achieving short recovery times with minimal data loss for critical applications require enterprises to build data replication architectures into their disaster recovery plans. Replication is a technology designed to maintain a duplicate data set on a completely independent storage system at a different geographical location. Depending on business requirements, this could be across campus, across town or across continents. In addition to having data at a disaster-safe location, the ideal replication solution would make sure that the replicated data are current (fully up-to-date), complete (including all applications, logs, configuration data, etc.), and recoverable (consistent data that is free from errors).

Copies of critical data can be produced in real time (synchronous), as fast as possible (asynchronous) or at distinct points in time (periodic), by replication technologies. These technologies are either embedded in the storage hardware itself (i.e. ESS's Peer-to-Peer Remote Copy (PPRC)) or run on hosts attached to storage servers (such as VVR). The data can be transferred over networks connecting the production site and the disaster recovery site. Generally customers opt for a private network between two sites with business critical data.

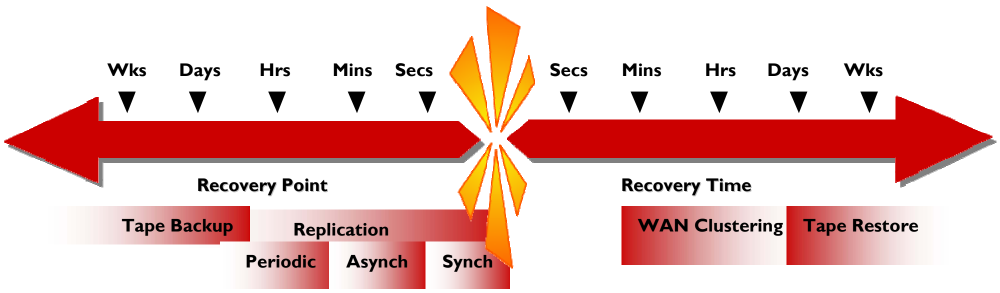

The above diagram provides a good overview of where the technologies fit in terms of the RPO/RTO and from that it will be easier to determine the best technologies to implement for different business requirements. This paper describes the benefits of various replication solutions, the criteria for choosing one solution over another and the pros and cons of each solution discussed. The paper also describes the steps taken to validate such solutions for typical customer scenarios and the best practices for implementing them.

3.1 Synchronous Replication

Synchronous replication is an excellent solution for offsite mirroring and for creating a disaster recovery site. PPRC Metro Mirror and VVR are technologies for synchronous replication. This type of replication requires all changes made within the primary site to be recorded at both the primary and secondary site before the application or storage server returns the acknowledgment that the change has occurred. Synchronous replication is designed so that no data will be lost between the primary and the secondary. This form of replication also means that there typically will be no data mismatch between the primary and the secondary sites. The application has to incur the network round-trip penalty and this could impact application performance.

3.1.1 Synchronous replication using ESS-PPRC Metro Mirror

PPRC is an established data mirroring and replication technology that has been used for many years in mainframe environments and has been available more recently in open systems environments. ESS-PPRC Metro Mirror is a synchronous protocol that allows real-time mirroring of data from one Logical Unit (LUN) to another LUN in a second ESS. The second ESS can be located at another site several miles away. PPRC is application and operating system independent and offloads the processing from the host. Since the copying function occurs at the disk subsystem level, it is completely transparent to the application and the operating system. It requires no resources from the host, and does not tax the network resources, because it uses its own PPRC link. The PPRC protocol keeps the secondary copy up-to-date by confirming that the primary copy will be written only if the primary system receives acknowledgement that the secondary copy has been written. PPRC also allows for customers to replicate only a subset of the data.

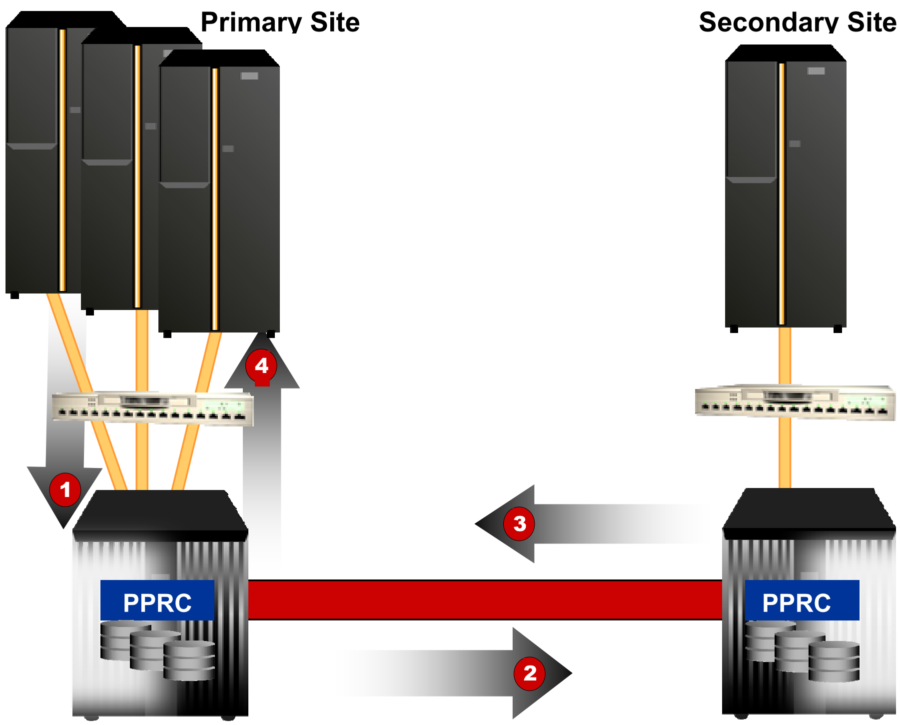

Synchronous replication using ESS PPRC Metro Mirror

Overview of the steps:

Step 1: Application host issues a write request to the primary storage.
Step 2: The primary storage then sends the write to the secondary storage
Step 3: The secondary storage acknowledges the write to the primary
Step 4: The primary storage acknowledges to the application that the write is successful.

3.1.2 Synchronous replication using VERITAS Volume Replicator

This is a solution for customers who want to migrate or replicate from a non-ESS storage array and to the ESS array. VVR is a host-based solution, which is completely integrated with VERITAS Volume Manager and enables replication to occur with minimal additional host resources.

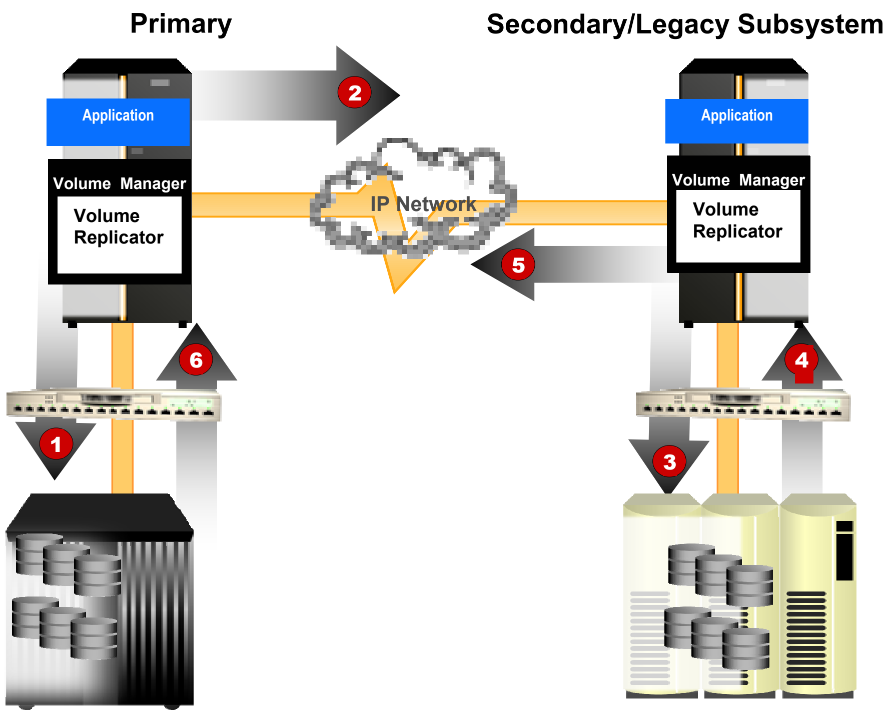

Synchronous replication using VVR, ESS and Legacy Subsystem as Secondary

Brief writeup about the steps:

Step 1: Application issues a write request to VVR
Step 2: The write request is written to VVR log and sent to the secondary host
Step 3: The secondary host acknowledges the write request and writes to storage
Step 4: The secondary storage commits write to disk
Step 5: Write acknowledgement sent to primary site.
Step 6: Write request committed to disk, VVR acknowledges the write to the application and primary host purges write request from VVR log

3.2 Asynchronous Replication

Asynchronous replication is a replication mode for customers who can absorb some minimal data loss but cannot afford to have their application to be impacted by the network round-trip impact. In asynchronous replication mode a write request is made to a primary site and committed immediately. As the network becomes available, the write requests are sent to the secondary site in the same order that they were written to disk at the primary site. At that point, the write request is committed at the secondary site.

3.2.1 Asynchronous Replication using VERITAS Volume Replicator

VVR gives customers the flexibility to replicate over long distances. Using VVR a customer can replicate or migrate data from an existing array to any storage array especially IBM's high-end and mid-range offerings like the ESS. Like IBM PPRC, VVR allows the customer to replicate only important data. In addition as part of an overall disaster recovery solution, the customer can use IBM's high-end array at the primary site and mid-range array at the secondary site. With VVR the existing network infrastructure can be used when replicating at extended distances very similar to PPRC Global Copy, which is explained in detail in section 3.3.1.

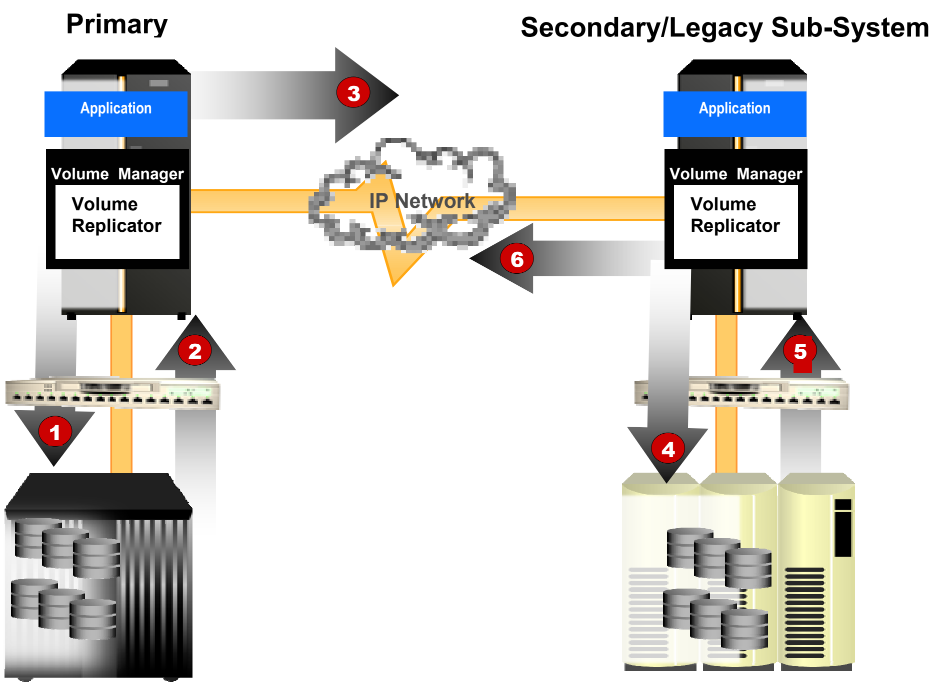

Asynchronous Replication using VVR, ESS and Legacy Subsystems as Secondary

Overview of the steps:

Step 1: Application issues a write request to VVR
Step 2: The write request is written to VVR log, committed to disk, and VVR acknowledges the write to the application
Step 3: Write request is sent to secondary host
Step 4: The secondary host acknowledges the write request and VVR issues a command to issue write request to secondary storage array.

Step 5: Write is committed to secondary storage array
Step 6: Write commitment sent to primary host and primary host purges write request from VVR log

3.2.2 Asynchronous Replication using ESS-PPRC Global Mirror

Asynchronous PPRC (PPRC Global Mirror) is designed to enable a two-site disaster recovery and backup solution for zSeries and open systems environments. In this mode of operation, PPRC copies data from the primary volume to the secondary volume in a *non-synchronous* manner. I/O write completions are returned to the application once they have been committed to the primary ESS. Updates on the secondary volume are performed at a later point in time. A point-in-time FlashCopy (see section 4.4.1) is periodically invoked at the secondary site to provide a continuously updating, nearly up-to-date backup. The data consistency is managed across multiple ESS machines configured in a single session, with data currency at the secondary site lagging behind the primary site an average of 3 to 5 seconds for many environments. PPRC Global Mirror operates over high-speed, Fibre Channel communication links and is designed to maintain a consistent and restartable copy of data at a secondary site that can be located at virtually unlimited distance from the primary site.

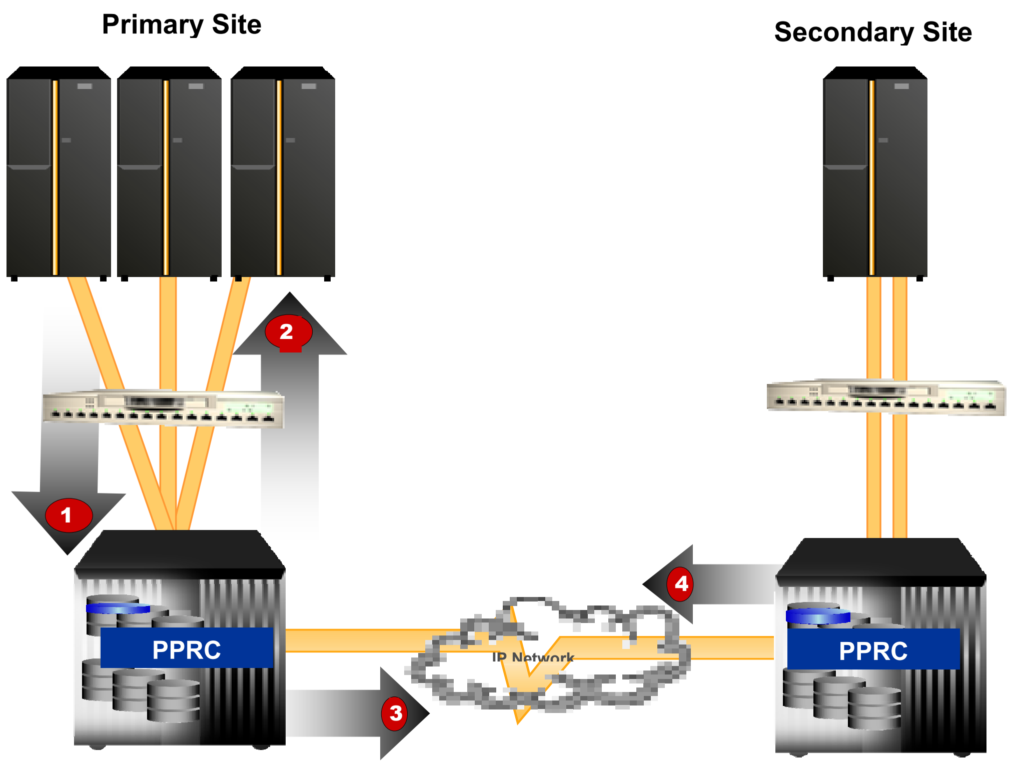

Asynchronous Replication Using PPRC Global Mirror

Overview of the steps:

Step 1: Application host issues a write request to the primary storage.
Step 2: The primary storage acknowledges to the application that the write is successful.
Step 3: The write is then replicated to the secondary storage.
Step 4: The secondary storage acknowledges to the primary storage that the write was successful. FlashCopy is periodically invoked on all secondary storage in the Global Mirror session to provide a consistent point-in-time image of the data.

Comment: I don't know if you want to use the word 'image' here, but I think most people think of 'backup' as another copy of the data taken from the FlashCopy target volumes.

3.3 Periodic Replication

Periodic replication is an excellent choice for customers who want to leverage available IP bandwidth and need to create copies of data at distances in the range of thousands of miles. In this mode of replication, the data is synched up at distinct points in time to maintain a consistent and current copy of data. It is very similar to Asynchronous mode of replication in its operation. This mode of replication typically transfers changed data blocks or similar unit of data like tracks.

3.3.1 Periodic Replication using IBM PPRC Global Copy

IBM's PPRC Global Copy offers an excellent proposition for replicating at unlimited distances for customers requiring periodic point-in-time copies at a remote location. With PPRC Global Copy, PPRC mirrors the primary volumes' updates onto the secondary volumes in a non-synchronous manner while the application is running. In this way, the application's write operations are free of the typical synchronous-like overheads. The non-synchronous characteristic of PPRC Global Copy combined with its throughput-efficient technique for copying only track updates makes it a recommended option for remote copy solutions at very long distances, even continental distances, and with minimal application performance impact.

PPRC Global Copy, unlike other hardware options, does not need any additional storage for staging purposes. The PPRC Global Copy option opens new approaches to mirroring implementations. The consistency of secondary volumes is periodically built during planned time windows.

PPRC Global Copy uses include:

1) Remote Data Migration
2) Transmission of Database log files (Log shipping)
3) Remote backup (Offsite vaulting)

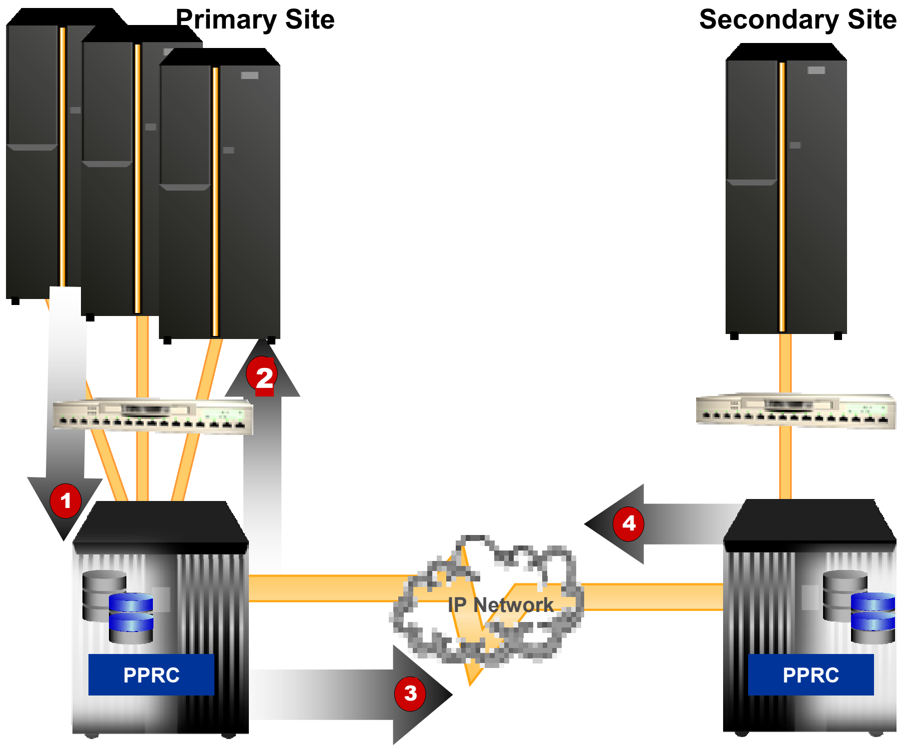

Periodic replication/migration using ESS PPRC Global Copy

Brief writeup about the steps:

Step 1: Application issues write to Storage array
Step 2: The storage array acknowledges the write back to the application.
Step 3: The storage array then sends over only the changed blocks (not entire tracks) to secondary site (this greatly reduces communication costs)
Step 4: At predetermined intervals in time the secondary storage array is sync'ed up with the primary storage array.

3.3.2 ESS-PPRC Metro/Global Copy

PPRC Metro/Global Copy is not a mode of PPRC operation, but it is a combination of Synchronous PPRC (PPRC Metro Mirror) and Periodic Replication (PPRC Global Copy) modes. PPRC Metro/Global Copy provides a long-distance remote copy solution for zSeries and open systems environments by allowing a PPRC secondary volume (involved in a PPRC Metro Mirror relationship) to also simultaneously serve as a PPRC primary volume in a PPRC Global Copy relationship to the remote site. This new capability enables the creation of three-site or two-site Asynchronous Cascading PPRC configurations.

4.0 Replication Based Disaster Recovery Planning – Best Practices

A well thought out disaster recovery plan would require leveraging advanced functions provided by ESS and VERITAS. To begin with all of the data should be built on the hardware RAID-5 or RAID-10 provided by ESS. This helps protection from single disk failures and most double and triple failures. Having ESS-based RAID protection is typically better than mirroring the primary volumes as even point-in-time/snapshot copies are RAID protected against single disk failure. In traditional split-mirror approaches where the third mirror is snapped off to create a business volume, the copy created is not protected against even single disk failures. For VVR customers who need RAID-5, a subsystem-based RAID-5 provided by ESS is the recommended technology to support maximum performance. Having the ESS advanced copy services create both synchronous mirrors and on-disk point-in-time copies at times when the server is busy would provide for exploitation of the ESS advanced hardware functions.

Customers with disk-based backups are typically able to recover faster from disasters. However, disk-based backups should always be in addition to tape backups and not in lieu of.

A complete disaster preparedness plan would include local clustering (for instance, VERITAS Cluster Server or IBM's HACMP) and geographically dispersed clustering (for instance, VERITAS 's Global Cluster Manager or IBM's HAGEO) along with protection provided by storage arrays and replication. An optimal solution would leverage the strengths of both advanced hardware functions and software features for best protection should a disaster occur.

4.1 Comparing Different Modes of Replication

4.1.1 Synchronous Replication (PPRC Metro Mirror and VVR)

Strengths
- Maximum data currency (RPO)
- Data available immediately at secondary (RTO)

Drawbacks
- Potential to impact performance
- Data writes lag by network latency
- Write commit latency across nodes

4.1.2 Asynchronous (VVR Asynchronous mode and PPRC Global Mirror)

Strengths
- The application write performance is not impacted by network latency
- Data is available immediately at secondary (RTO)

Drawbacks
- Secondary may lag primary (RPO)
- Snapshots required for using the secondary data

4.1.3 Periodic Replication (PPRC Global Copy)

Strengths
- Data available immediately at secondary
- Performance is superior as only changed data blocks are transferred.

Drawbacks
- Requires snapshots to ensure that the data is consistent on the secondary, however it is common practice to do this even when using asynchronous replication.

4.2 Comparing Different Methods of Replication

4.2.1 Storage Subsystem Based Replication

Strengths:
- Host Performance is not impacted at all with using Hardware array as the host CPU cycles are not used for replication.
- Can replicate multiple operating systems at the same time (Heterogeneous Operating System support)

Drawbacks:
- Can only replicate to another storage array of the same make and model.
- Synchronous replication possible to only to pairs, more than one target for source is not supported.

4.2.2 Host-Based Replication

Strengths:
- Can replicate between dissimilar storage arrays
- Can replicate over standard IP

Drawbacks:
- Can only replicate to other systems running the same operating system.
- Impact on host performance

4.3 Replication Usage Scenarios

Different customers have different needs and should carefully choose the type of replication for their particular environment. The following are the most common reasons for a replication-based Disaster Recovery solution:
- ❑ Hot Site
- ❑ Automated multi-site data processing
- ❑ Remote or Off Host Processing including Backups

Hot Site: A hot site is defined as "a data center facility with sufficient hardware, communications interfaces and environmentally controlled space capable of providing relatively immediate backup data processing support." A customer could be interested in application resumption or fail over at the secondary site; this is referred to as a Hot Site solution. This method can take hours to days if done manually, which would be unacceptable to customers with low RTO requirements. If application resumption or fail over is a requirement then additional components such as VERITAS Cluster Server (VCS) or IBM HACMP on AIX for application fail over and possibly VERITAS Global Cluster Manager or IBM HAGEO for the management of the clustering solution over the wide area may be needed. When used in conjunction with replication, the recovery time could be reduced to minutes. This is typically the quickest way of getting critical applications up and running after a disaster.

Automated multi-site data processing: In this complex scenario different sites are converted into primary sites based on majority of the usage. For example, during peak hours in the US, the New York site could be the primary one, and during peak hours in Japan, the Tokyo site could takeover as the primary site. It involves sequential application fail over as well as the repeated promotion and demotion of nodes. IBM PPRC/VVR, IBM HACMP on AIX/VCS, & IBM HAGEO/GCM can be used to meet such requirements.

Remote/Off Host Backup processing: Data can be replicated to another site for long-term storage. A point-in-time snapshot of the replicated data is required for backup or vaulting purposes due to the fact that during active replication, file systems at the secondary are not mountable for backup, off-host reporting or any other purpose. A snapshot using VERITAS FlashSnap functionality or ESS FlashCopy is the way to make this data stable and usable. PPRC Global Copy can be cascaded with IBM FlashCopy as an excellent solution for creating remote backups.

4.4 Making Point In Time Copies of Data

Point-in-time copies or (sometimes referred to as snapshots) are useful for creating copies of systems, and databases for backup and off-host processing as in database warehousing and decision support purposes. Point-in-time copies are an important part of an overall Disaster Recovery strategy because they help protect against logical corruption and enable the ability to utilize the data at the secondary location.

4.4.1 Point in Time Copy using ESS FlashCopy

ESS FlashCopy provides an instantaneous copy, or "view", of what the original data/ESS logical volume looked like at a specific point-in-time. This is known as the T0 (time-zero) copy. When the ESS FlashCopy is invoked, the command returns to the operating system as soon as the FlashCopy pair has been established and the necessary control bitmaps have been created. This process takes only a few seconds to complete. Thereafter, a T0 copy of the source volume is available for use. As soon as the pair has been established, you can read and write to both the source and the target volumes. The copy looks exactly like the original source volume and is an instantly available, binary copy. ESS also offers the ability to suppress the background copy task using the

Do not perform background copy (NOCOPY) option. It results in a logical copy of the data with only the blocks that are about to be changed being recorded before the change occurs, commonly referred to as the "copy-on-write" technique. This may be useful if the copy is needed for only for a short time, such as making a backup to tape.

FlashCopy supports the ability to establish a relationship where the FlashCopy target is also a PPRC primary volume. This enables you to create full or incremental point-in-time copy of a LUN, volume, or data set at a local site, and then use PPRC to copy the data to a remote site. It also allows you to use the flexibility of FlashCopy on a volume that is already set up for automatic copy to a remote site using PPRC.

4.4.2 Point in Time Copy using VERITAS FlashSnap

VERITAS FlashSnap technology is similar to the ESS FlashCopy. To create a snapshot of a volume, some disks are allocated in the same disk group. This volume is required to be the same size as the source volume. After the snapshot volume is attached to the source volume, the syncing operation starts. There are commands to enquire the status of the syncing. Once syncing is complete then the volume can be detached and used as a Point-in-Time copy. VERITAS FlashSnap also provides a mechanism to merge the volume back into the volume that it was attached to. FlashSnap is a host-based approach.

4.4.3 Point in Time Copy using VERITAS Space-Optimized snapshot

VERITAS VVR also enables you to take instant space-optimized snapshots. Unlike instant full snapshots, the instant space-optimized snapshots usually require less storage space than the original volume because space-optimized snapshots store only the data that has changed between the original volume and the snapshot. Typically, the data that has changed between the original volume and the snapshot volume over the lifetime of the snapshot is minimal compared to the total data on the volume. Thus, the space-optimization achieved can be significant. The snapshot data is managed by VVR using a cache object that functions as a space-optimized persistent store. You must create the cache object before you take the instant space-optimized snapshots or specify the size of the cache object. Multiple snapshots can be created on the same cache object. The cache object can be created with an autogrow option set to on to allow it to grow automatically if the cache volume size is not enough for the specified writes.

5.0 IBM ESS and VERITAS Solution Validation

ESS and VERITAS products have been extensively tested in the IBM San Jose Lab in typical customer scenarios. VERITAS products have been tested on IBM Enterprise Storage Server. VERITAS Volume Replicator has achieved both IBM TotalStorage Proven and IBM ServerProven (for pSeries) status. IBM is also a participant in the VERITAS-Enabled program. The scenarios testing included Exception Testing, Disaster Scenarios over a Wide Area Network in addition to typical usage validations.

The hardware setup for the solution validation included

1) 2 Sun 450 machines
2) Multiple ESS Logical Sub-Systems
3) In-Range Switch
4) Multi-Pathing using 2 Emulex 8000 HBAs and VERITAS DMP on each host machine
5) Replication over WAN/LAN between the machines

The software setup included

1) VERITAS File System 3.4 (and applicable options for Oracle)
2) VERITAS Volume Manager 3.2 (and applicable options for Oracle)
3) VERITAS Volume Replicator 3.2 (and applicable options for Oracle)
4) Oracle 8.1.6
5) Solaris 2.8

5.1 Setup Details

The system was setup with Multiple Logical Sub-Systems from an ESS connected to two Solaris machines. A VERITAS File System was laid out over volumes (ESS LUNs) under the control of VERITAS Volume Manager. The volumes were intentionally not mirrored as ESS supports reliability of data by applying RAID-5 at the storage level (Note: RAID 10 is also supported by ESS). Volume Replication was setup between the two hosts over a Wide Area Network (WAN). VVR was setup as a 2-node system. VVR-RLink was setup in an asynchronous replication mode. The replicated volumes and the SRL were setup as concatenated volumes relying on the RAID-5 protection offered by the ESS Array, which offers an equivalent level of protection to volume mirroring. VERITAS Volume Manager was configured and names of volumes were consistent between the primary and secondary sites as outlined in the VERITAS best practices guide.

5.2 Replication Scenarios

A number of sample solution scenarios were validated including the following listed below, each has an example application.

Solution Scenario 1: Replicate file systems between two machines both having data residing in the ESS.
Solution Application: A file system hosted on ESS can be replicated across large distances.

Solution Scenario 2: Replicate databases between two machines both mounting disks from ESS.
Solution Application: Databases hosted on ESS LUNs can be replicated across significant distances.

Solution Scenario 3: Migrate Primary role from one database server host to the secondary host with both hosts mounting their replicated volumes from ESS LUNs.

Solution Application: To move the processing from one site to another and enable serial application deployment model.

Solution Scenario 4: Force the secondary host to takeover the primary role with both hosts mounting their replicated volumes from ESS LUNs.
Solution Application: Enable processing from a secondary site when the primary is down.

Solution Scenario 5: Pause the replication on the secondary host, force transactions on the primary, and trace the SRL buildup. Resume the replication and ensure that SRL frees up.
Solution Application: Allow for network downtime while the primary is still servicing requests.

5.3 Point in Time Copy Scenarios using ESS FlashCopy

Solution Scenario 6: Validate the ESS FlashCopy for a File System.
Solution Application: To ensure that the snapshot is consistent and can be used for backup and off-host processing.

Solution Scenario 7: Validate the ESS FlashCopy for an Oracle database.
Solution Application: To ensure that the snapshot is consistent and can be used for backup and off-host processing.

5.4 Point in Time Copy Scenarios using VERITAS FlashSnap

Solution Scenario 8: Validate the FlashSnap of a FileSystem residing on ESS LUNs.
Solution Application: To ensure that the snapshot is consistent and can be used for backup and off-host processing.

Solution Scenario 9: Validate the snapshot of an Oracle database.
Solution Application: To ensure that the snapshot is consistent and can be used for backup and off-host processing.

All of the above scenarios were validated in lab settings and the results were successful**.**

6.0 Summary

Disaster recovery planning is a strategic necessity for business. The various options for disaster recovery presented in this paper can be evaluated to find the best fit for disaster recovery requirements. For example, requirements vary based on the applications involved and the importance of particular data.

One can use an appropriate mix of synchronous replication, asynchronous replication and periodic replication to meet disaster recovery requirements. In addition, there are several emerging methods to replicate data, such as replication using database technologies and replication coupled with virtualization layers.

The ESS and VERITAS configurations described in the paper are a sample of the scenarios that have been tested. There is additional work underway to integrate IBM ESS functions with VERITAS software even more deeply, with the goal of further reduced backup times and additional flexibility.

Disaster recovery planning and implementation can be very complex. In these situations, consider using professional services, such as IBM Global Services.

7.0 References:

1) IBM TotalStorage Enterprise Storage Server: Implementing ESS Copy Services in Open Environments -- IBM Redbook, Document Number SG24-5757-04, July 2004 http://www.redbooks.ibm.com

2) IBM TotalStorage Enterprise Storage Server: Implementing ESS Copy Services with IBM eServer zSeries -- IBM Redbook, Document Number SG24-5680-04, July 2004 http://www.redbooks.ibm.com

3) IBM TotalStorage Enterprise Storage Server: PPRC Extended Distance – IBM Redbook, Document Number SG24-6568-00, June 2002 http://www.redbooks.ibm.com

4) IBM Enterprise Storage Server - IBM Document Number: SG24-6424-01, November 2002 http://www.redbooks.ibm.com

5) IBM TotalStorage Enterprise Storage Server: Implementing the Enterprise Storage Server in your Environment -- IBM Redbook, Document Number SG24-5420-00, March 2002 http://www.redbooks.ibm.com

6) IBM TotalStorage Enterprise Storage Server Performance Monitoring and Tuning Guide - - IBM Redbook, Document Number SG24-6422-00, July 2003 http://www.redbooks.ibm.com

7) DB2 UDB Backup and Recovery with ESS Copy Services – IBM Redbook, Document Number SG24-6557-00, August 2002 http://www.redbooks.ibm.com

8) Disaster Recovery Strategies with Tivoli Storage Management – IBM Redbook, Document Number SG24-6844-01, November 2002 http://www.redbooks.ibm.com

9) VERITAS Volume Replicator White Paper: Successful Replication and Disaster Recovery
http://www.veritas.com/products/listing/ProductDownloadList.jhtml;jsessionid=XM3NNS4HGQDMDQFIYBSSFEY?productId=volumereplicator#datasheets

10) VERITAS Disaster Recovery Seminar Series Part I Presentation
http://products.veritas.com/products/campaigns/dr/presentations.htm

11) VERITAS Disaster Recovery Seminar Series Part II Presentation
http://products.veritas.com/products/campaigns/dr/presentations.htm

12) VERITAS Volume Replicator Sales Guide
http://products.veritas.com/products/high_avail/volume_rep/datasheets.htm

13) Management Guide to Replicating Critical Data
http://products.veritas.com/products/high_avail/volume_rep/datasheets.htm

8.0 Related Websites

1) IBM Website
http://www.ibm.com

2) IBM Storage Website
http://www.storage.ibm.com

3) IBM TotalStorage Proven Website.
http://www.storage.ibm.com/proven/index.html

4) IBM Redbooks Website
http://www.redbooks.ibm.com/

5) IBM Sales Website
http://www.storage.ibm.com/readytobuy.html

6) IBM Enterprise Storage Server Product Page
http://www-1.ibm.com/servers/storage/disk/enterprise/index.html

7) IBM Enterprise Storage Server Product Specifications
http://www-1.ibm.com/servers/storage/disk/enterprise/index.html

8) VERITAS Website
http://www.veritas.com

9) VERITAS Sales website
http://vnet.veritas.com

10) VERITAS Volume Replicator Product Page
http://www.veritas.com/products/category/ProductDetail.jhtml?productId=volumereplicator&_requestid=82895

11) VERITAS Volume Replicator Data Sheet
http://www.veritas.com/products/listing/ProductDownloadList.jhtml;jsessionid=XM3NNS4HGQDMDQFIYBSSFEY?productId=volumereplicator#datasheets

DISCLAIMERS

The information in this publication is provided "AS IS" WITHOUT WARRANTY OF ANY KIND, implied or otherwise, including but not limited to any warranties of merchantability or fitness for a particular purpose. The responsibility for use of this information or implementation of any of these techniques lies with the customer and depends on the customer's ability to evaluate and integrate these techniques into their operating environment.

Accordingly, the recipient assumes all risks arising from the use of this information. In no event shall IBM or VERITAS be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. While IBM or VERITAS may have reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The data contained in this document were obtained under a controlled, isolated environment. The results obtained in your operating environment may vary significantly. Accordingly, this data does not constitute a performance guarantee or warranty.

Any statements regarding IBM's and VERITAS's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The following terms are either registered trademarks or trademarks of the IBM Corporation in the United States or other countries or both:
IBM, Enterprise Storage Server, FlashCopy, Peer-to-Peer Remote Copy, PPRC and TotalStorage.

The following terms are either registered trademarks or trademarks of VERITAS Software, Inc in the United States and/or other countries:
VERITAS Volume Manager, VERITAS File System, VERITAS Volume Replicator, VERITAS Cluster Server, VERITAS Global Cluster Manager, VERITAS NetBackup and VERITAS FlashSnap

Other company and product names mentioned herein may be trademarks of their respective companies.