StorageTek®

TECHNICAL BRIEF

Introduction to Open Systems PPRC

APRIL 2004

1 ABSTRACT

This document provides a brief introduction to StorageTek's PPRC software (Peer-to-Peer Remote Copy, a synchronous remote mirroring application) and in particular its implementation in an Open Systems, non-Multiple Virtual System (MVS) environment.

In September 2002, StorageTek® announced support for Open Systems PPRC, its latest enhancement to the Virtual Power Suite™ software. Open Systems PPRC allows nonmainframe clients to enjoy the data-protection capabilities that have been available to mainframe customers since the early 1990s.

PPRC allows multiple discrete virtual disk subsystems to contain and exchange the same information, providing superior protection against data loss from catastrophic events. Once enabled, PPRC runs virtually unattended, allowing data from a primary subsystem at one location to be copied in real time to a secondary (or recovery) subsystem at another location.

With PPRC, records at the recovery site are continuously updated to match those at the primary site, providing an additional level of data redundancy beyond that available in a single V-Series Shared Virtual Array® (SVA™) disk system. PPRC minimizes the potential for data loss during catastrophic events and reduces the impact of such disasters by allowing faster recovery of critical data.

1.1 THE BUSINESS VALUE

Businesses today are dependent on information availability and access. The majority of enterprises acknowledge that they would suffer losses of business impact if they were to lose access to their data for extended periods of time. To some of the larger financial institutions, the cost of information unavailability can run into millions of dollars.

As these businesses become more and more dependent on information technology (IT) to conduct their operations, data availability has probably become the most critical customer requirement today.

Of all the IT resources within a company, data is the most important. Assets such as processing power, software, hardware and facilities can all be replaced in the event of a disaster. However, data is the most volatile and complex of all resources and is certainly the most critical to the effective recovery of a company.

Companies are now seeking to deploy disaster recovery (DR) strategies to enable quick recovery of their production services with minimal impact to business.

StorageTek's PPRC software is designed to assist with the quest for accurate and rapid recovery of your environments in the event of a disaster.

1.2 CONSIDERATIONS WHEN SELECTING A REMOTE COPY METHODOLOGY

There are a number of key issues when selecting the appropriate remote copy methodologies to incorporate into an effective disaster recovery strategy. However, two probably stand alone. These are the speed of data recovery and the integrity of the data itself.

The speed of recovery will be dictated by the business. It depends on how long the business can manage without its data before there is an unacceptable impact on customers and business.

However, the most critical of all requirements has to be data integrity. Data without integrity is potentially worse than having no data at all; this is due to the unpredictability of the results obtained by processing erroneous data.

Open PPRC

Applications implement logic to permit data integrity and allow for successful transaction blackout and recovery in the event of a failure, often referred to as dependent writes or transactional processing. No updates will be issued by the application unless the previous update has completed successfully. Any remote copy solution must be able to help make sure that the logic implemented by the application is maintained. Therefore the data at the primary and secondary locations must be updated in exactly the same sequence as defined by the application. If the solution is not able to apply the logic implemented by the application, data integrity cannot be protected.

For controller-based remote copy implementations, maintaining consistency and integrity of the data across multiple controllers can present a challenge. This is because the application has no knowledge of the remote copy and the controller has no knowledge of the application-dependent write. When the application has completed its update, it is the responsibility of the controller to help make sure that the update is reflected on the secondary subsystem. Since the controller only has knowledge of itself and the secondary subsystem it is connected to, maintaining secondary update sequence consistency across multiple controllers can present a serious issue.

Disasters seldom occur instantly. They are normally preceded by a number of smaller discrete failures before the effect of the disaster is felt. This effect is known as a rolling disaster. A successful DR solution must be able to handle the rolling disaster effectively.

For controller-based remote copy implementations, particularly semisynchronous and asynchronous operations, rolling disasters present a serious exposure since there is no central point of control and since the host has no knowledge of the remote copy. The host will continue to update the primary volumes even though all the remote links or connections may be lost. Not only are the remote volumes out of sync, but worse still, you do not know that they are out of sync. Should a subsequent disaster occur, bad data will be used for recovery since there is no knowledge of problems from the remote site. Synchronous replication — which helps make sure that both the primary and secondary writes are complete before posting I/O complete to the host — can avoid this problem.

However, one of the limiting features of synchronous replication is the performance overheads associated with the added levels of protection utilized in protecting data integrity. This is due to its design, where each update is not reflected back to the system as completed until the secondary controller has confirmed the update. The effect on an application of this performance penalty increases with distance, block size and update intensity. Care must be taken prior to implementing a synchronous copy solution to understand the effects of this performance overhead and to allow you to be able to maintain your service-level agreements.

1.3 SUMMARY

If the performance overheads of synchronous replication are acceptable for your workload, this should always be selected as the mechanism to copy data to a remote site.

Synchronous replication offers the highest levels of data currency after a disaster. With asynchronous replication, data loss will occur as updates buffered in controllers awaiting transmission will be lost in the event of a disaster.

Nonsynchronous controller-based solutions cannot protect sequential integrity in multicontroller environments, and therefore the integrity of the data is at risk.

Data integrity cannot be guarded in the event of a disaster when controller-based nonsynchronous data replication is implemented.

Only synchronous replication can offer the levels of data integrity and consistency required to implement a truly effective DR solution.

2 UNIQUE VALUE ADDS OF THE STORAGETEK PPRC

The architecture of StorageTek virtual disk products allows for a number of unique business benefits provided by StorageTek's implementation of PPRC when compared to other competitive remote copy solutions.

2.1 ONLY COMPRESSED/COMPACTED DATA REPLICATED

Due to the architecture of StorageTek's virtual disk products and their capability to compress and compact data as it is written into the subsystem, all data sent between the subsystems is already in a compressed format, significantly reducing overheads and potential network bandwidth requirements.

2.2 ONLY VALID DATA SENT

As only data is stored on the virtual disk subsystems (we have the capability not to store white space), we are able to transfer only valid data between mirrored volumes. This significantly reduces overhead and potential network bandwidth requirements.

2.3 INTEGRATION OF SNAPSHOT WITH PPRC

StorageTek has integrated SnapShot software, the most efficient data replication software available, into PPRC. This means that customers can now enjoy the benefits of SnapShot technology over distance.

All of the above means that StorageTek is able to offer the most efficient synchronous data replication solution available in the marketplace today!

2.4 POTENTIAL CUSTOMERS

In the past, PPRC was only available to mainframe customers. (Proxy PPRC was also available but it, too, required mainframe for control.) With the introduction of Open Systems PPRC, however, we now have a solution for customers who do not have or require mainframe connectivity.

This is a tremendous enhancement to our current virtual disk portfolio. Customers who require the additional levels of protection provided by PPRC who could not previously exploit these functions due to the lack of Open Systems support can now do so.

Potential customers of StorageTek's PPRC solution are:

> Existing virtual customers who can now enhance their current DR strategy
> New Open Systems customers who see the value of virtual but require the added levels of protection
> Customers looking to implement DR solutions
> Customers looking to complement existing DR solutions who can exploit the advantages of PPRC SnapShot.

3 EVOLUTION OF STORAGETEK'S PPRC

3.1 STANDARD PPRC

In the early 1990s, IBM developed Peer-to-Peer Remote Copy (PPRC) to address the issues of data protection and continuous availability. Early implementations were expensive and lacked acceptable performance and operational flexibility.

In 1998, StorageTek provided PPRC support for its mainframe virtual disk customers. The command set and operational interfaces were completely compatible with IBM's 3990-6 PPRC product. However, this StorageTek-developed PPRC product exploited the internal data compression feature and other aspects of virtual architecture inherent in the SVA and Ramac Virtual Array (IBM RVA base) subsystems and was superior in every way to IBM's original offering. This PPRC product was realistically limited to operational distances of three kilometers or less, due to ESCON cable limitations and performance considerations relative to the response-time impacts on production transactions. If connected to a long-distance medium via a channel extension or converter box, this product suffered severe performance penalties. This was due primarily to the extremely "chatty" nature of the ESCON protocol and the inefficient use of the PPRC link(s).

3.2 POWERPPRC DIRECT MODE

PowerPPRC software was introduced to overcome many of the limitations inherent in standard PPRC. The PowerPPRC product significantly improved the efficiency of the SVA PPRC ESCON protocol, providing you with:

> Dramatic reduction in PPRC protocol overheads
> Increased utilization of PPRC links
> Reduced performance penalty associated with synchronous replication
> Significant improvements on initial synchronization times.

The main reason for the significant improvements in performance and link utilization were due to the implementation of bridge volumes. These bridge volumes consist of specially configured volume pairs and a dedicated ESCON path. This eliminates any arbitration for the link between the subsystems and allows for less queuing and higher link utilizations. (Further information on bridge volumes can be found later in the document.)

3.3 POWERPPRC WAN MODE

For those who require longer distances than are provided by direct-mode PowerPPRC, wide area network (WAN) mode was introduced. This allowed customers to use the benefits of PPRC over hundreds of kilometers (the previous distance limitation was three kilometers). The ESCON connection is still used to connect the subsystems; the additional distance is provided by the CNT Ultra Net Storage Directors.

With the implementation of PowerPPRC software, StorageTek introduced two new volume types, status and bridge volumes. The bridge volume is utilized in both direct and WAN modes. Those looking to replicate data over longer distances than direct mode provides require the addition of a second volume, known as the status volume.

When transferring data over long distances, it is possible that multiple tracks may be transferred over the network at any one time. To keep the network link busy and I/O turnaround times to a minimum, the status volume was introduced. Each time a track is received at the remote site, a status packet is returned back to the primary on a different link, notifying the primary that the update is complete.

3.4 PPRC SNAPSHOT SOFTWARE

Until the introduction of PPRC SnapShot software with the combined advantages of SnapShot in a PPRC environment, users have been limited by a restriction whereby a SnapShot replication could only be directed to a non-PPRC volume. When a PPRC customer wanted to use SnapShot to do a full-volume copy for backup, the customer had to suspend the PPRC pairs, Vary Online the secondary volume, execute the copy with SnapShot, and then resynchronize the volume pair. Another way was to execute a copy with SnapShot from the primary volume onto a non-PPRC volume and then establish a PPRC pair with a secondary and perform the synchronization. Although these approaches work well, there is the need for a simpler, quicker and more efficient way to accomplish it. With PPRC SnapShot software, PowerPPRC users can now create SnapShot copies for backup on both subsystems without moving any data and without using additional physical space.

With PPRC SnapShot software, initially both the source volume and the target volume need to be paired with their respective secondary volumes on the other SVA. When a SnapShot command is sent to replicate a copy of the production primary onto the backup primary, PPRC SnapShot detects that they are both primaries, which means there are two secondary volumes somewhere. PPRC SnapShot verifies that both secondary volumes are on the same remote SVA. From that, PPRC SnapShot knows that it can allow the execution of the SnapShot command on the Local and Remote SVA as it knows data integrity will be preserved (production primary is equivalent to production secondary and backup primary is equivalent to backup secondary). Now, we know that whatever we do in both primary volumes will be replicated on both.

3.5 SNAP-TO-PRIMARY

Thanks to the performance and efficiency of of the StorageTek's PowerPPRC software, many of the DR requirements once considered asynchronous can now be satisfied with a synchronous solution and PPRC SnapShot software. However, for those business requirements that allow for an asynchronous solution, StorageTek now has a feature called Snap-to-Primary that provides users with the capability of propagating consistent SnapShot copies of data volumes to remote subsystems while using minimal capacity and bandwidth. This feature is part of the PPRC SnapShot product. Additionally, each subsequent SnapShot will propagate only data that have been updated, thereby rendering the most efficient remote replication functionality in the industry today. The PowerPPRC Snap-to-Primary feature provides the most efficient remote disk backup capability using standard ESCON cables for extended distances over telecom facilities (WAN mode). Furthermore, it can be used concurrently with the synchronous mode. Snap-to-Primary combines the benefits of StorageTek's SnapShot and PowerPPRC software to offer an efficient, economical and automatic asynchronous disaster recovery solution for those who have great distances between their primary and remote sites and/or have limited bandwidth connecting the two sites. To initialize a Snap-to-Primary configuration, a PowerPPRC pair is established between a volume at the primary site and a volume at the secondary site. Then, with the same nondisruptive, instantaneous feature of SnapShot, a Snap operation is performed from a source volume (non-PPRC) to the PPRC primary volume at the primary site. Immediately, while production continues, PowerPPRC begins copying the data to the PPRC secondary volume at the remote site. When the PPRC secondary volume at the remote site is fully synchronized with the PPRC primary volume at the primary site, a backup SnapShot can be taken of the PPRC secondary volume to preserve a consistent copy of the data through any future iterative process of PowerPPRC synchronization that could occur. *Procedures for accomplishing these backups vary depending on whether or not you have host connectivity to the secondary subsystem.* Upon the next identical Snap-to-Primary operation, only the changed data is propagated to the secondary volume. As you can see, multiple generations can be kept at the secondary site by merely snapping to different target volumes on the secondary. See **Figure 1**, “Multiple Generation Snaps” below.

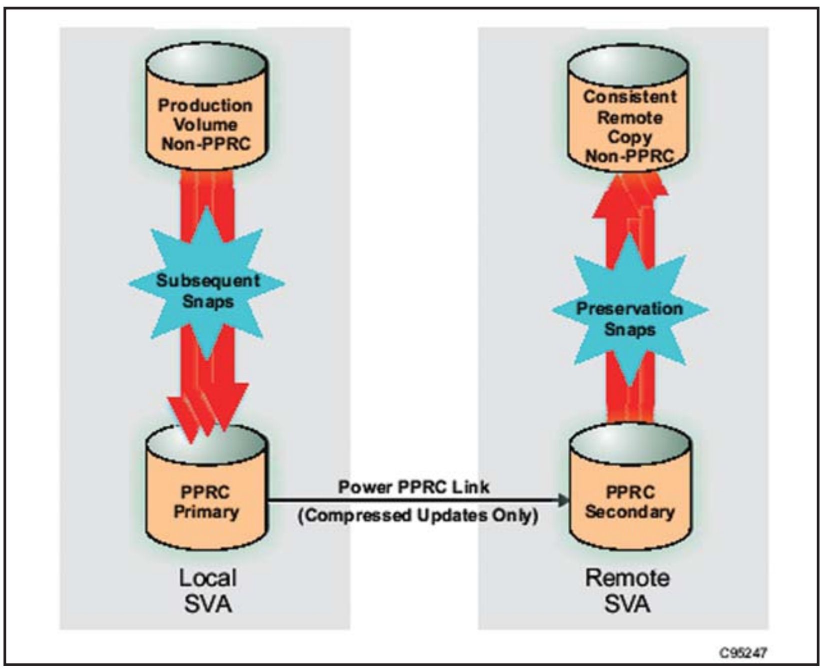

Figure 1. Multiple-generation snaps.

A snap from the remote host (secondary side) would require terminating the pair. **Figure 2** on page 10 shows that pairs can remain active without a termination.

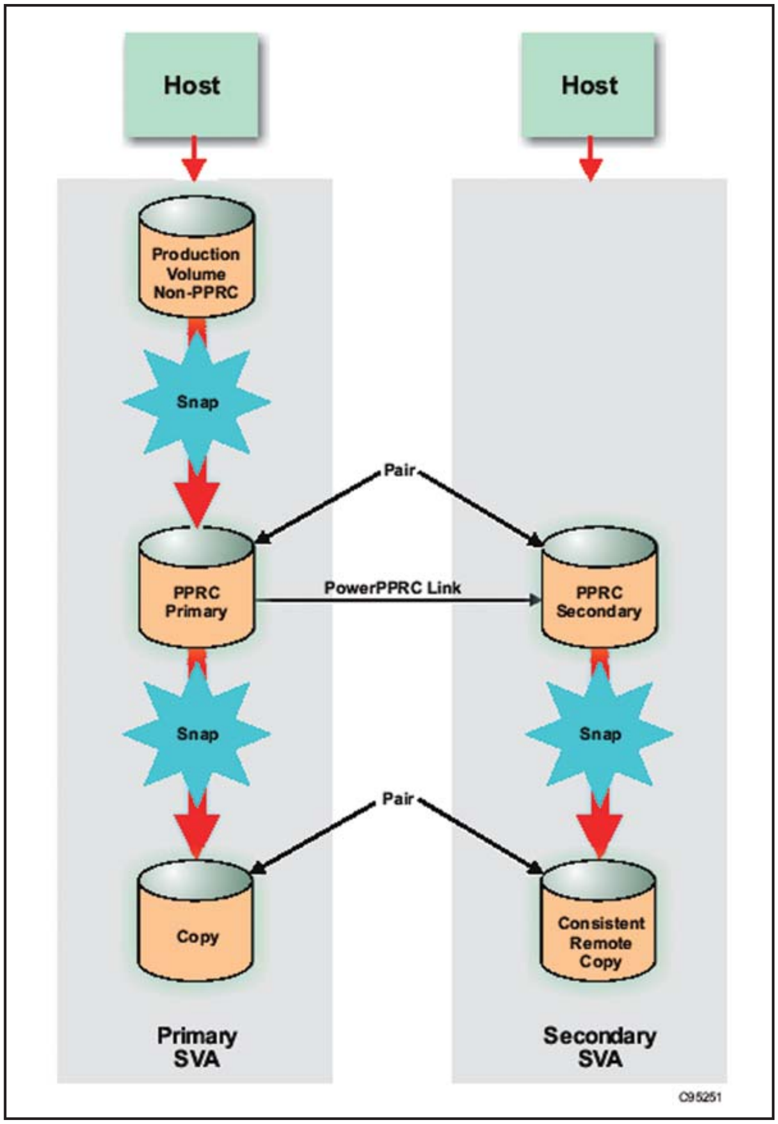

Figure 2. Nontermination snap.

3.6 PROXY PPRC SOFTWARE

Proxy PPRC software utilizes the facilities provided by MVS to provide remote disk mirroring functionality to non-mainframe clients. Open Systems volumes designated for PPRC use are replicated from the primary system to the secondary system as though there were standard MVS volumes, even though they contain nonmainframe data.

No actual PPRC functionality was required by the clients to operate in this environment.

3.7 OPEN SYSTEMS PPRC

Prior to the release of Open Systems PPRC in September 2002, Open Systems clients requiring the data protection facilities provided by PPRC had to use mainframe Proxy PPRC to manage and control the PPRC sessions via MVS.

With this announcement, Open Systems clients are now able to utilize the benefits of PPRC without an OS/390 attachment.

The SVA(™ Administrator (SVAA) software has been enhanced to support additional commands to provide the facilities to configure and manage an Open Systems PPRC environment.

The following platforms support native Open Systems PPRC directly: AIX and W2K.

4 PPRC — BASIC ARCHITECTURE

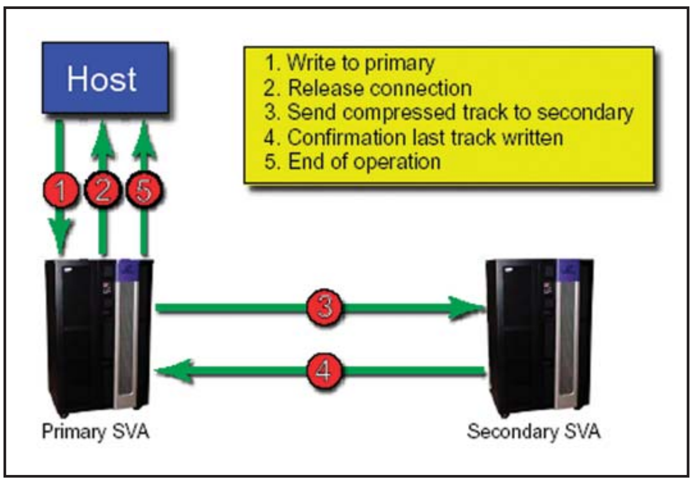

Figure 3. Basic PPRC operation.

1. **Write to primary** — All writes to the primary virtual subsystem are compressed and stored to nonvolatile cache.
2. **Release connection** — Once the I/O has been received by the primary subsystem, the channel (SCSI Bus) is released and made available to the host for other I/O operations.
3. **Send compressed track to secondary** — Once the Update to the primary has completed, it sends full tracks of compressed data to the secondary. The RVA does not decompress the data prior to sending the data to the secondary. These compressed tracks once received by secondary subsystem are then written directly to non-volatile cache on the secondary subsystem.
4. **Confirm last track written** — The secondary system notifies the primary of a “write complete” for the last track.
5. **End of operation** — After the transfer has completed and acknowledgment is received at the primary from the secondary, the primary subsystem notifies the host that the I/O has completed.

5 POWERPPRC — BRIDGE AND STATUS VOLUMES

5.1 BRIDGE VOLUMES

Dramatic improvements in PowerPPRC link utilization are achieved by exploiting a new technique called "throughput queuing." PowerPPRC software accomplishes this by using "bridges" to transfer data to the secondary subsystem. These bridges consist of specially configured PPRC pairs and a dedicated ESCON path that the pair will use. The volumes that make up the "bridge pairs" are not accessible from any host. The bridges are dedicated to handling all the data and status traffic that goes between subsystems. Because the bridge volumes do the transfer work for the other PPRC volume pairs, overhead is greatly reduced. There is no longer any arbitration for the ESCON path among the pairs. A volume that wants to send data to the secondary just puts the track on a queue. Tracks are pulled off the queue by the bridge and sent over one after another in a long chain. This chaining effect also reduces ESCON overheads. The bridge is able to chain tracks from all the different devices in the subsystem.

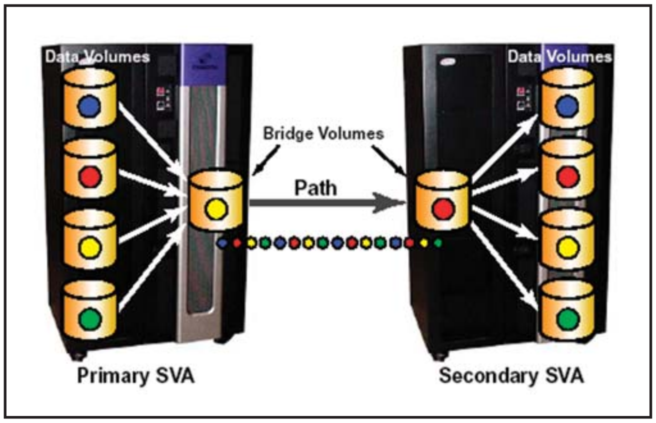

Figure 4. PowerPPRC operation.

Yet another benefit to using bridges involves synchronizing the volumes. When a PPRC pair is established with the "copy all" option, a volume is defined to be "duplex-pending." This means that the volume on the primary is not identical to the volume on the secondary. In order to get the volumes to be exactly the same, all the tracks on the primary must be copied onto the secondary. This process is referred to as "volume synchronization." Once the volumes are synchronized, they are referred to as "duplex." In the standard PPRC implementation, all duplex-pending volumes tried to synchronize at the same time. This consumed subsystem uses resources such as cache space, data paths and processor cycles. When a large number of volumes were synchronizing at the same time, host performance could be impacted.

In order to address this problem, PowerPPRC uses a synchronization queue when bridges are used. When a non-bridge volume is established as duplex-pending, it is placed on the synchronization queue.

The primary bridge volume pulls volume numbers off the synchronization queue (one at a time) and does the synchronization work for the volume. The bridge gets the data from the original volume by doing an internal

SnapShot copy of the volume selected. Once the bridge has sent over all the data, it restarts the original volume and gets the next device off the synchronization queue. The original volume will do a second pass through all of its tracks and send over any tracks that may have been changed while the bridge was synchronizing that volume. Since only a small number of volumes (four or fewer) will be doing synchronizations at the same time, the host response time will not be degraded due to resource consumption by the synchronization processes.

5.2 STATUS VOLUMES

When transferring over long distances such as a wide area network, it is possible to have multiple tracks in the WAN at one time due to the long transmission delays. In order to keep the WAN busy, tracks are sent one after another without receiving status from the secondary. In order to know if a track actually made it to the secondary successfully, a status packet is sent back on a different ESCON link to notify the primary that the track was successful. This means that in a long-distance configuration, there are two types of bridges, data bridges and status bridges. These need to be configured in pairs — one status bridge pair for each data bridge pair. The status bridge pulls status packets off a queue and sends them to the primary system, which notifies the initiating volume of the transfer status of a track.

It is possible to use bridges in a PowerPPRC direct mode configuration that uses direct ESCON connections. However, since the shorter distance is not long enough to have multiple tracks in the connecting cable, a status bridge is not required in this configuration, only a data bridge. Because the data bridge chains the tracks as just mentioned, there is still a large performance benefit to using a data bridge in a non-WAN or direct mode configuration.

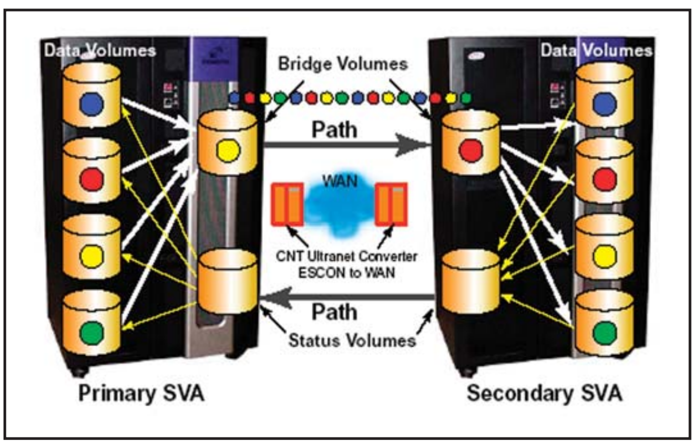

Figure 5. PPRC WAN operation.

6 OPEN SYSTEMS PPRC — HOW TO CONFIGURE

PowerPPRC in the Open Systems environment is achieved through the the SVA™ Administrator (SVAA) software. The end user utilizes SVAA software to configure and maintain the SVA and PPRC environment.

The main steps involved in initiating PPRC communication between two subsystems are as follows:

> Connect the links between the subsystems (ESCON cables/CNT USDs).
> Define a data bridge on each subsystem for each link.

Note: If utilizing WAN mode, define a status bridge on each subsystem for each link.

> Via SVAA, establish the logical paths between primary and secondary subsystem.
> Via SVAA, establish the connection between the data bridge volumes on the primary and secondary subsystems.
> Via SVAA, define the data volumes on the primary and secondary subsystems.
> Via SVAA, establish the connection between the data volumes.
> Via SVAA, the following additional enhancements have been provided to allow for management of the PPRC operation in an Open Systems environment:

- Display information on logical path status
- Delete previously defined logical path definitions
- Display status of paired data volumes
- Display status of data bridge volumes
- Suspend updates between a primary and secondary data volume
- Resume updates between a primary and secondary data volume
- Delete the logical connection between a pair of data volumes.

These commands have a large number of of subparameters, and it is recommended that scripts be written to automate both the setup and shutdown of the environment.

With the SVA, you have the ability to create large LUNs. A large LUN is a combination of a number of functional devices presented back to the host as a single logical entity. Care must be taken when setting up and managing large LUNs in a PPRC environment, as PPRC manages the devices at a function device level and not the LUN level.

For further information on how to configure PPRC in an open environment, please review "Peer-to-Peer Remote Copy Configuration Guide," StorageTek part number MP-4007F.

7 OPEN SYSTEMS PPRC — VCU AND FDID SETUP

The V2X Shared Virtual Array® (SVA™) disk system is made up of 16 virtual control units (VCUs), whereas previous versions of the virtual disk subsystem supported 4 VCUs. These virtual control units act as separate logical controllers within the subsystem.

Prior to any data being mirrored between the subsystems the user must establish the required logical paths between the devices. A path is a logical link between a VCU on a primary subsystem and a VCU on a secondary subsystem.

A path between VCUs is uni-directional. For bi-directional transfers, multiple paths have to be defined over multiple physical connections.

Each VCU supports up to 256 devices, with VCU 0 supporting FDID 000 – 0FF, VCU 1 supporting FDID 100 – FDID 1FF and so on, through to VCU 15.

Virtual Control Unit ID	Supported device range
VCU 0	FDID 000 – FDID 0FF
VCU 1	FDID 100 – FDID 1FF
VCU 2	FDID 200 – FDID 2FF
VCU 3	FDID 300 – FDID 3FF
VCU 4	FDID 400 – FDID 4FF
VCU 5	FDID 500 – FDID 5FF
VCU 6	FDID 600 – FDID 6FF
VCU 7	FDID 700 – FDID 7FF
VCU 8	FDID 800 – FDID 8FF
VCU 9	FDID 900 – FDID 9FF
VCU 10	FDID A00 – FDID AFF
VCU 11	FDID B00 – FDID BFF
VCU 12	FDID C00 – FDID CFF
VCU 13	FDID D00 – FDID DFF
VCU 14	FDID E00 – FDID EFF

It is essential that a degree of planning be performed at this stage and standards be implemented to allow for ease of management and to allow for future growth. A simple recommendation is that when configuring the pairs, configure them with similar addresses, so that if the primary PPRC volume on SVA1 has a functional device address of 0100, the secondary PPRC volume on SVA2 also utilizes device address 0100.

If large LUNs are to be utilized, then all the functional devices, both parent and child, that are associated with that LUN must have a path established for each associated VCU.

The current restriction on the number of available logical paths between primary and secondary subsystems is 16, for the V960 and previous versions of the product, and 64 for the V2X.

8 OPEN SYSTEMS PPRC — SUPPORTED CONFIGURATIONS

The following platforms support Open Systems PPRC:
Windows 2000
AIX

The following have been tested and are running in Toulouse:
HP-UX
Solaris

The following configuration restrictions currently apply to Open Systems PPRC configurations:

> Open Systems PPRC support is only provided for PowerPPRC direct and WAN modes
> ESCON is the only supported protocol for inter-subsystem communications
> ICE3 cards (four-port ESCON) are not supported, only ICE2 cards (two-port ESCON)
> Only CNT USDs are supported for PowerPPRC WAN mode
> Uni- and bi-directional PPRCs are supported between subsystems.
> Only single hop configurations are supported.

9 OPEN SYSTEMS PPRC — HOW TO ORDER

The following feature codes in ICOA can be used to order the features associated with Open Systems PPRC:

ICOA Model Number	ICOA Description	Additional Information
V2X01Ax-0000	V2X SVA:x ARRAY/0 GB CACHE/ 0 CHN	Where x equates to the number of arrays and potential capacity of the subsystem valid values are 1,2,3,4
V2X01Ax-yyyZ	GB CAPACITY @ 4.0 COMPRESSION	Where yyy equates to the capacity of the subsystem, values range from 430GB and 7730GB at a compression ratio of 4:1
V2X01Ax-9957	LINE FILTER CORD 50HZ	50Hz Power Cable
V2X01Ax-ACzz	GB EFFECTIVE CACHE	Where zz is the effective cache size (3:1 compression ratio) , valid values are 8,16 and 24
V2X01Ax-ES04	4 ESCON LINKS	ICE2 Cards to provide up to 4 PPRC links, minimum allowed
V2X01Ax-FMff	FC W/MULTI-MODE GBICS	Where ff is the number of fibre channel interfaces installed, valid values range between 0 and 12
127OPN-0000	SVAA 3.1 FOR Open SystemS	Shared Virtual Array Administrator, required for configuration and management of SVA SVAA MLC is per subsystem per platform SVAA ILC is per subsystem per platform No DSLO offering
1127OPN-CDRM	CD-ROM	Software distribution method
1127OPN-AIXX	AIX	Licensed for AIX
1127OPN-HPXX	HP-UX	Licensed for HP-UX
1127OPN-LINX	SVAA FOR LINUX INTEL	Licensed for Linux
1127OPN-NTXX	NT	Licensed for NT
1127OPN-SOLX	SOLARIS	Licensed for Solaris
1127OPN-WINX	WIN2000	Licensed for Windows 2000
PPRC002-0000	POWER/PPRC LICENSE	Power PPRC LC and MLC per CPU
PPRC002-HD35	3.5 DISK HIGH DENSITY	Software Distribution Method
PPRCSNP-0000	PPRC ShapShot	PPRC ShapShot LC and MLC per subsystem
PPRCSNP-HD35	3.5 DISK HIGH DENSITY	Software Distribution Method

10 STORAGETEK PROFESSIONAL SERVICES — PPRC IMPLEMENTATION

StorageTek offers PPRC implementation service by StorageTek Professional Services. A description of this service can be found on the Global Services Intranet Web site.

11 HINTS, TIPS AND THINGS TO LOOK OUT FOR

There are some additional considerations when implementing Open Systems PPRC:

11.1 DISABLE FILE SYSTEM/VOLUME LEVEL CACHING

As the main purpose for synchronous PPRC is to provide data consistency between data centers, it is recommended that file system/volume caching be disabled. If this is not disabled, then updates that have not been committed to disk and are still buffered in the host may be lost to the secondary subsystem in the event of a failure.

11.2 LARGE LUN SUPPORT

PPRC handles devices at the functional device level. With large LUNs, all of the parent-and-child functional devices that comprise the large LUN have to be managed/monitored separately.

11.3 NO NOTIFICATION OF LOSS OF SYNCHRONIZATION

Once PPRC pairs have been established for volumes between two subsystems, if synchronization is lost between the two subsystems for reasons such as network failure or hardware error, no notification is passed back to the host. This means that a potential failure could occur causing a loss of synchronization between subsystems, and subsequent attempts to recover applications on the secondary volumes could be performed on old data. Users familiar with mainframe PPRC will be aware of the CRIT parameter that can be used to indicate whether writes should continue in the event of loss of synchronization. This is not available with Open Systems PPRC.

11.4 MULTIPLE VOLUME/LUN CONSISTENCY

One of the key issues with remote copy is the ability to maintain consistency across multiple volumes/LUNs spanning multiple subsystems. There are no controls within Open Systems PPRC to establish coherency across multiple volumes aside from halting all applications that are accessing them.

12 REFERENCE INFORMATION

There is limited “referenceability” for Open Systems PPRC.

No SE Tools to model effect on performance

There are currently no StorageTek tools available to model the performance effect of applying PPRC to an existing workload. It is therefore difficult to quantify and demonstrate the effect.

12.1 OTHER SOURCES OF INFORMATION

Additional information on Open Systems PPRC and other related information can be found at the following locations:

> Customer Resource Centre: http://www.support.storagetek.com.

12.2 IBM ONLINE DOCUMENTATION/REDBOOKS

http://www.redbooks.ibm.com/
SG24-5338-xx RAMAC Virtual Array: Implementing PPRC
SG24-2595-xx Planning for IBM Remote Copy

STORAGETEK®

ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK), a $2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151

MS 9153 B 04/04