StorageTek®

TECHNICAL BRIEF

PowerPPRC and StorageTek's V-Series disk systems

APRIL 2004

1 ABSTRACT

Never before have data protection and continuous availability been as important as they are today. We are in an "e-world" where either your business is online, going online or going out of business. That being the case:

> The business records of today's online transactions must be protected, as these records may be the only evidence of a transaction.

> Historically, if you could not process a transaction, the transaction was usually deferred until you were back online. With the fluidity of today's commerce, if you are not online, you may have lost that transaction to a competitor or you may have lost the customer entirely.

> A recent study at Oracle claims that it only takes seven seconds of wait time to lose an online customer.

For mainframe and Open Systems/distributed environments, various options are available to provide mirrored copies. These solutions include host-based software solutions that are disk-agnostic but CPU-intensive:

> XRC from IBM

> Logical volume mirroring contained within VERITAS Volume Manager and most Open Systems operating systems

> Disk-based solutions that are vendor-specific but operating system-agnostic — meaning they can support many operating and logical volume manager environments. This provides one way to manage an entire shop or global enterprise with one disk vendor.

IBM developed Peer-to-Peer Remote Copy (PPRC) in the early 1990s to address the issues of data protection and continuous availability. PPRC is a synchronous remote disk mirroring solution that protects the integrity of data. The early implementations were expensive and lacked both acceptable performance and operational flexibility. They were not universal solutions. Competitive products were built, but they were hampered by a lack of integration into the OS/390 architecture.

Experience with and analysis of those original efforts have led to more robust solutions while the continued falling prices of disk storage and communication lines have made the cost much more affordable.

In June of 2000, StorageTek® introduced PowerPPRC software, an implementation of PPRC that operates totally within the OS/390 architecture and exploits the integrated OS/390 controls. StorageTek's latest implementation is capable of remotely mirroring:

1. Mainframe environments OS/390, z/OS, VM (virtual machine), and 2/VM
2. Open Systems/distributed environments Windows NT, Windows 2000, AIX (advanced interactive processing), Solaris, HP-UX, and Linux
3. Mainframe and Open Systems/distributed environments.

The third environment mirrors Open Systems data through OS/390 facilities. This feature, called Proxy PPRC, provides disaster recovery and other benefits of PowerPPRC to your Open Systems applications. StorageTek recently broadened the scope with Open PPRC, a new functionality that uses SVA™ Administrator (SVAA) software or SVA console to remotely manage the mirroring of Open Systems data in an open-only environment. See the StorageTek technical brief "Introduction to Open Systems PPRC," MS9153B.

Another recent feature, SnapShot software, allows you to make simultaneous SnapShot copies of PPRC volumes (or data sets) on both the primary and remote SVA subsystems. Since it uses SnapShot, it does not take additional disk space, is instantaneous and does not send data over the PPRC link. Now, StorageTek engineers have created a new and exciting feature called Snap-to-Primary. This feature allows users to direct a SnapShot copy (of any SVA V2X volume) to the primary volume of a PPRC pair, thus invoking PowerPPRC to mirror this volume remotely without any manual intervention or performance impact on the source volume. A PPRC volume pair consists of exactly two identical disk volumes each under PPRC control and residing on a separate storage subsystem. Additionally, any subsequent snap operation will cause PowerPPRC to send only updated data to the remote mirror. Please see the section 2.3, Snap-to-Primary, for a more detailed view of this noteworthy feature. It is part of the PPRC SnapShot software.

Additionally, the latest version of the OS/390 SVA management software product, SVAA, offers you the ability to further simplify the configuration and management of PowerPPRC (OS/390 and open PPRC).

PowerPPRC:

> In synchronous mode maintains an exact copy of your primary disk volumes at a remote site.
> Promotes simplified recovery scenarios.
> Allows an expedient restart after a disaster.
> Runs virtually unattended.
> Offers outstanding efficiency and performance.
> Is transparent to applications and systems.
> Maintains the sequence of writes across any number of hosts, control units or volumes.

In addition, the Snap-to-Primary feature of PowerPPRC replicates consistent point-in-time disk copies of your data to a remote site without manual procedures, tape handling or transporting and without affecting production performance.

After implementing the hardware solution, you are able to maintain and modify the configuration without additional expense or vendor assistance. PowerPPRC software is an exceptionally flexible and powerful remote disk mirroring product in the industry.

2 BUSINESS NEEDS AND CONSIDERATIONS

As businesses become more and more dependent on information technology to conduct their operations and stay competitive, the availability of their processing facilities becomes more crucial. Today, most businesses require a high level of availability, which extends to continuous availability 24 hours a day and seven days a week operation.

Governmental regulations including the Patriot Act require some industries to have strategies that ensure against outages and guarantee that business will continue uninterrupted. A lengthy outage could lead to:

> Significant financial losses
> Loss of credibility with customers
> Possibly a total failure of business.

Therefore, the ability to provide continuous availability for business-critical applications is more often than not a necessity for business survival.

Different disaster scenarios require different approaches to data backup and recovery. Table 1 lists the requirements that most customers face.

Requirement	Description/Comment
Short recovery time	The ability to recover from a disaster within an acceptable amount of time
No data loss or a small amount of data loss which allows fro greater distance and still ensures recover-ability to a consistent useable state	An imperative requirement in today's e-world
Nondisruptive DR testing	The operational requirement of testing DR procedures at the remote site without disrupting the production DR configuration
Workload migration	The operational requirement of migrating workloads from one installation to another without disruption
Device migration	A process for migrating data onto new technology without affecting the most important applications
Outage limitation	Outage limitation for normal housekeeping activities
Simple Management	User-friendly configuration and management capability

Table 1.

In most cases, you cannot predict when a disaster will occur. The challenge is to be prepared and have a functional disaster recovery (DR) procedure in place at all times.

Of all the information technology (IT) resources within a company, data is the most important. Other resources, such as processing power, vendor software, storage and building facilities are all ultimately replaceable; but some data is not replaceable! Data is also the most volatile and complex of all resources.

Therefore, the most important step in the disaster recovery process is data recovery. The methods used to replicate and protect critical data, the manner in which it is transported and stored offsite and the techniques used for recovering the data all place restrictions and requirements on the technical specification of the remote site and equipment therein. There are many different processes and products for backing up and recovering data, each having its relative merits.

Synchronous remote disk mirroring products are hardware solutions designed to assist with the age-old quest for accurate and rapid disaster recovery. They also provide valuable workload and disk migration functions. In the past, however, the cost of synchronous remote disk mirroring eliminated it as an alternative for most customers.

Synchronous remote disk mirroring is clearly the most desirable with respect to ease of recovery, recovery time required, data consistency and integrity. With PowerPPRC, StorageTek has the best solution in the industry. It is now easier, more protective, and less expensive than ever to use PowerPPRC.

More companies are using synchronous remote disk mirroring than ever before because of:

> Increased need for data security and availability.
> Increasingly global nature of business enterprises.
> Lower communications costs.
> Improved performance and management technology.

In the mainframe environment, synchronous PowerPPRC helps to ensure data integrity and consistency grouping, giving you the opportunity to maintain compatibility with IBM's Geographically Dispersed Parallel Sysplex (GDPS) application recovery service.

Operational simplicity is an extremely desirable attribute of a sound disaster recovery plan. The simpler a plan is to execute, the higher the probability of successful execution. This issue, along with growing requirements that businesses be available 24 x 7, puts the focus on synchronous remote disk mirroring products to meet those requirements. PowerPPRC software provides provides the performance, scalability, flexibility and manageability to be considered the world-class solution for all customers, large and small. The ability to monitor pair status and perform some configuration tasks can now be accomplished through the SVA Administrator software (a subsystem management software product), mainframe ISPF (Interactive System Productivity Facility) panels, Open Systems SVAA software or the SVAC product.

2.1 PPRC SNAPSHOT

Previously, the combined advantages of SnapShot in a PPRC environment were limited by a restriction whereby a SnapShot replication could only be directed to a non-PPRC volume. A PPRC customer that wanted to use SnapShot to do a full volume copy for backup had to suspend the PPRC pairs, vary the secondary volume online, execute the copy with SnapShot, vary it back offline and then resynchronize the volume pair. Another way was to execute a copy with SnapShot from the primary volume onto a non-PPRC volume, establish that volume as part of a PPRC pair with a secondary and then do the synchronization. Although these approaches worked well, there was the need for a simpler, quicker and more efficient way to accomplish this replication requirement.

With the PPRC SnapShot software, PowerPPRC users can create identical synchronized SnapShot copies for backup on both subsystems without moving any data and without using additional physical space. There is an exciting new feature that comes with this product called Snap-to-Primary. We'll discuss that feature after we cover basic PPRC SnapShot.

Following are some principles of PPRC SnapShot and what you need to know when planning its installation.

PPRC SnapShot is built upon SnapShot and PowerPPRC functionality. It works with SnapShot (SNAP-001 or SNAP-002), PowerPPRC (PPRC-002) and PPRC SnapShot (PPRC-SNP) installed in all participating SVAs connected for remote copy.

With PPRC SnapShot, both the source volume and the target volume need to be paired initially with their respective secondary volumes on the other SVA (see **Figure 1**).

When a SNAP 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 snap command on the local and remote SVA as it knows data integrity will be preserved (production primary is equal to production secondary and backup primary is equal to backup secondary). Now, we know that whatever we do in both primary volumes will be replicated on both secondary volumes. This is shown in FIgure 2.

Figure 1. PPRC pairing.

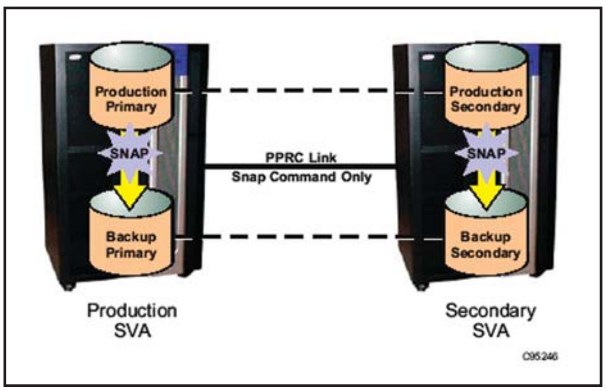

Figure 2. PPRC SnapShot operation.

Note: The dashed horizontal lines in Figure 2 represent PPRC pairs.

2.2 PPRC SNAPSHOT AND PROXYPPRC

PPRC SnapShot works for Multiple Virtual System (MVS) OS/390 as well as Open Systems volumes (when either open PPRC or ProxyPPRC is used to duplex or mirror Open System volumes).

2.3 SNAP-TO-PRIMARY

Thanks to the performance and efficiency of the PowerPPRC, many of the DR requirements once considered to be asynchronous can now be satisfied with a synchronous solution and PPRC SnapShot. 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 software. Additionally, each subsequent SnapShot will propagate only data that have been updated, thereby rendering an extremely efficient remote replication functionality in the industry today. The Power PPRC Snap-to-Primary feature provides the most efficient remote disk backup capability using standard ESCON cables or extended distances over telecom facilities (wide area network [WAN] mode). Furthermore, it can be used concurrently with the synchronous mode. Snap-to-Primary combines the benefits of the SnapShot software and PowerPPRC to offer an efficient, economical and automatic asynchronous disaster recovery solution for those who have great distances between the 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 3.

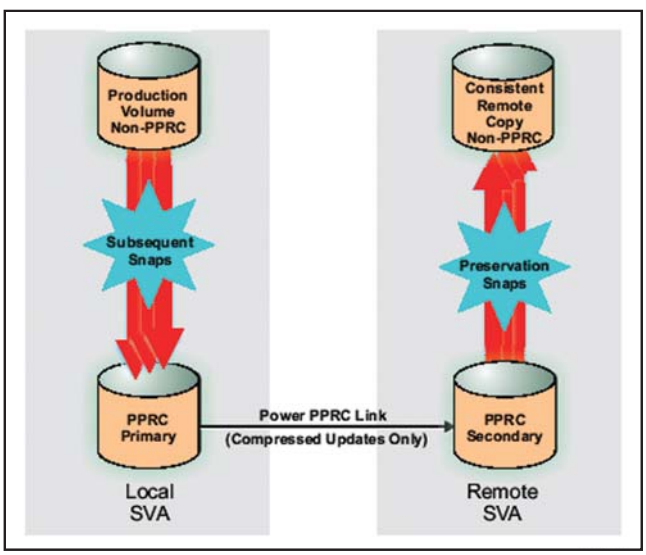

Figure 3.
Multiple-generation snaps.

A snap from the remote host (secondary side) would require terminating the pair. **Figure 4** allows pairs to remain active without a termination.

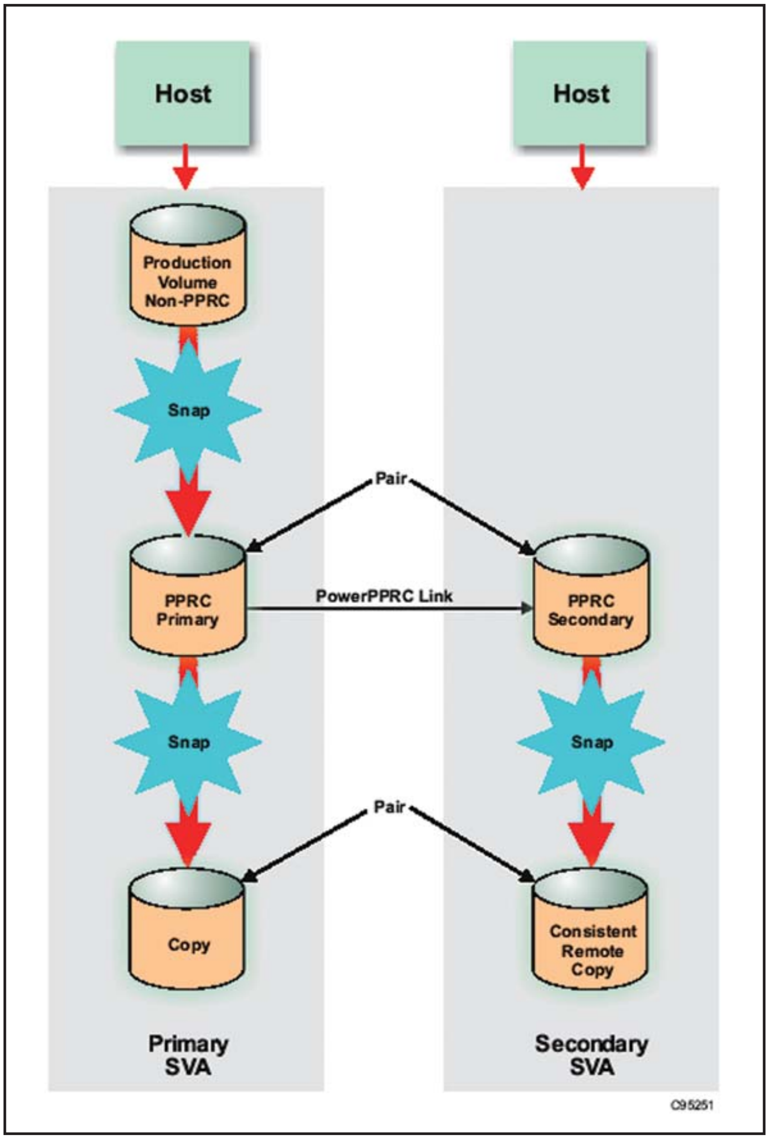

Figure 4. Nontermination snaps.

2.4 SNAP-TO-PRIMARY BENEFITS

The synchronous mirroring of PPRC does not affect response time on the production volumes. Snap-to-Primary offers several benefits:

> Provides all the benefits of PowerPPRC for those who require an asynchronous solution
> Transfers only written data on initial copy

< Transfers only changed data after the first copy is taken
< Transfers data in compressed mode (usually 4:1)
< Uses a new architecture called "throughput queuing" to further improve bandwidth utilization
< Allows you to determine how often sync points need to be taken
< Preserves at least one consistent copy of the data while a new synchronization is being accomplished
< Allows you to determine how many generations of remote consistent copies should be maintained on disk at the secondary site, requiring space only for the differences between the consistent remote copy and the PPRC secondary
< Provides easy, flexible customer management without vendor involvement in configuration changes.

Basically, the difference between PPRC SnapShot and a Snap-to-Primary operation is that with PPRC SnapShot the source volume in the snap process happens to be a PPRC primary volume as well. The results, however, are much different. With a PPRC SnapShot, no data is copied over the link. With Snap-to-Primary, data snapped to the primary is copied/synchronized over to the secondary. There is no specific setup required from the host standpoint. The function is automatically driven from the SVA code.

2.5 SVA HARDWARE AND PPRC MODES AVAILABLE

PPRC Mode	Disk Subsystems	Connectivity	Maximum Distance	Comments
Standard PPRC	9393 and 9500-1	ESCON only	3 Km	Efficient, easily managed
PowerPPRC, direct mode	9500-1.5+, V960 and V2X	ESCON only	3 Km	High-performance, efficient, easily managed
PowerPPRC, WAN mode	9500-1.5+, V960 and V2X	UltraNet, WDM telecom and	2,500+ Km	Long-distance, high-perform ance, efficient, easily managed
Proxy PPRC, PPRC SnapShot; direct or WAN mode	9500-1.5+, V960 and V2X	ESCON or UltraNet, WDM and telecom	Same as PowerPPRC	Same as Power PPRC but performed from OS/390 on behalf of Open Systems host(s)
Open PPRC, direct or WAN mode	V960 and V2X	ESCON or UltraNet, WDM and telecom	Same as PowerPPRC	Same as PowerPPRC but performed from Open Systems host(s) without OS/390 connectivity
Snap-to-Primary	V2X only as primary; 95001-1.5+ and V960 OK as secondary	ESCON or UltraNet, WDM and telecom	Same as PowerPPRC	This is a special feature of PPRC SnapShot that only works on the new V2X.
1. 9500 not available in open				

Table 2.

2.6 ADDITIONAL ADVANTAGES OF POWERPPRC IN EITHER DIRECT OR WAN MODE

Additional advantages could include:

> StorageTek's exploitation of data compression, write queuing and line utilization techniques reduces the number of PowerPPRC links required for synchronization and thereby increases the number available for connectivity purposes.

> StorageTek provides a hardware feature on the SVA that when activated automatically releases all volumes from PowerPPRC control. This feature expedites and simplifies the recovery process at the remote site by replacing numerous PowerPPRC commands with the simple touch of a button.

> It is a simple operation to migrate workloads from one installation to another without a major outage.

> The ability to easily define bridge volumes is provided by SVA Administrator software (version 3.1 and above), user-friendly command line interface ISPF panels, or the SVAC product. These options simplify the initial configuration of PowerPPRC and obviate the use of the operator panel for such tasks.

> PPRC SnapShot allows the testing of DR procedures at the remote site (using a point-in-time copy of production data) without disrupting the production DR configuration or processing.

> Snap-to-Primary combines the benefits of SnapShot and PowerPPRC to offer an efficient, economical and automatic asynchronous disaster recovery solution for those whose distance requirements and/or limited bandwidth eliminate the possibility of synchronous mirroring.

In addition, there may be other financial and operational benefits of the high-performance SVA V-Series virtual architecture.

2.7 THE EVOLUTION

In 1998, StorageTek provided virtual disk customers with a synchronous remote disk copy product called PPRC. 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 to the SVA and RVA subsystems and was superior in every way to IBM's original offering. (The RVA was our SVA marketed and sold by IBM during a multi-year OEM agreement. RVA stands for RAMAC Virtual Array.) This PPRC product was realistically limited to operational distances of 3 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).

2.8 MAJOR ADVANTAGES OF POWERPPRC OVER STANDARD PPRC

The major advantages of PowerPPRC over standard PPRC include:

< A dramatic reduction in PPRC protocol overhead

< A drastic reduction in (or elimination of) PPRC link contention by exploiting throughput queuing (read on for more detail on how it works)

< The achievement of high utilization of telecom facilities

< An increase in maximum distance by two orders of magnitude (100X) or more.

The PowerPPRC product drastically improves the efficiency of the SVA PPRC ESCON protocols for use with long-distance, high-speed communication lines (for example, T3 lines) and optical fibre loops (for example, OC3, SONET) to transfer data between local (primary) and remote (secondary) subsystems. This product incorporates significant advances in the SVA microcode as well as hardware and firmware in CNT's (Computer Network Technology's) UltraNet product. The UltraNet can extend ESCON channels over standard DS (DS1, DS3) or optical (OC1, OC3) telecommunications facilities. This CNT UltraNet connection uses SONET framing and Asynchronous Transfer Mode (ATM) format in optical configurations. The minimal configuration is a pair of single-mode fibres (one to transmit and one to receive) between the UltraNet units. This configuration also supports the connection through Fibre Channel switches. There are also UltraNet boards to connect to multimode fibre if only 62.5-micron fibre is available. The transmission mode of the particular dark fibre will determine the appropriate UltraNet OC3 board.

The CNT UltraNet units are configured with two dual-port ESCON boards (two connectors on each board), for a total of four ESCON links. This provides PowerPPRC WAN with two data links/bridges and two status links/bridges. For planning purposes, you can expect approximately 10 megabytes per second (36 gigabytes per hour) per data link over an OC3 circuit. Only compressed and compacted data is transmitted; therefore you would multiply this number by the appropriate compression factor to estimate the data throughput rate (typically 4:1 or 144 gigabytes per hour).

It is possible to configure PowerPPRC from multiple primary SVAs to a single secondary SVA, and also possible to configure from a single SVA to multiple secondary SVAs. Please note that a PPRC volume pair consists of exactly two volumes — that is to say, you cannot mirror one primary volume to multiple secondary volumes.

When configuring more than a single SVA primary subsystem, three or more ESCON data links can be directed through the same OC3 link. A configuration of six ESCON links to two OC3 links in the same CNT UltraNet frame provides a reasonable balance of ESCON throughput to OC3 bandwidth.

A critical factor in improving performance over long-distance communication lines is minimizing the number of link turnarounds (acknowledgment from the receiving link prior to continuing). Microcode for both PowerPPRC and the UltraNet device has been modified to reduce the number of link turnarounds compared with ESCON. Another very important performance factor for long-distance communications is link utilization (keeping the link busy transferring data). The new PowerPPRC WAN mode microcode allows several devices to transfer data to the remote subsystem while waiting for confirmation from the first device. With Standard PPRC, each primary virtual device sends tracks to its corresponding secondary device. All primary devices on the primary control unit compete for use of the PPRC paths and links to send tracks to the secondary. In the PowerPPRC WAN mode, all tracks from all primary devices are sent from the primary subsystem to the secondary subsystem using a special pair of devices called data bridge volumes. This process is called "throughput queuing." All status commands from the secondary subsystem are sent to the primary subsystem using another different special pair of devices called status bridge volumes, and each subsystem must have at least one data bridge volume and one status bridge volume. This implementation alleviates the competition issue for the PPRC paths that existed in the old standard PPRC product. The next section describes this further.

2.9 HOW BRIDGE VOLUMES AND SYNCHRONIZATION WORK

Dramatic improvements in PowerPPRC link utilization are achieved by exploiting throughput queuing. PowerPPRC 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 overhead. The bridge is able to chain tracks from all the different devices in the subsystem. See **Figure 5**.

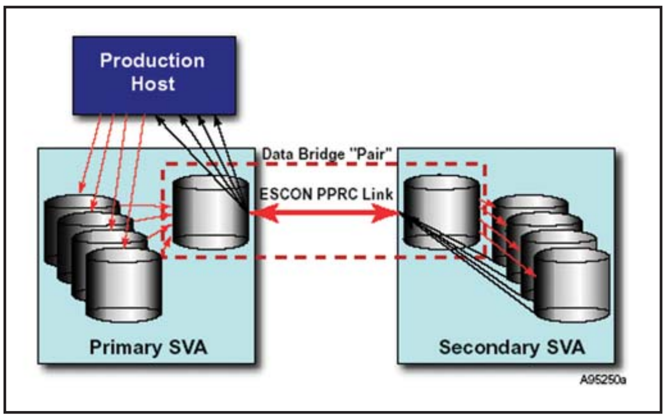

Figure 5. PowerPPRC direct mode.

When transferring over long distances such as a wide area network (WAN), 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, the Primary SVA sends tracks one after another without receiving status from the Secondary SVA. 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 replication 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. See **Figure 6**.

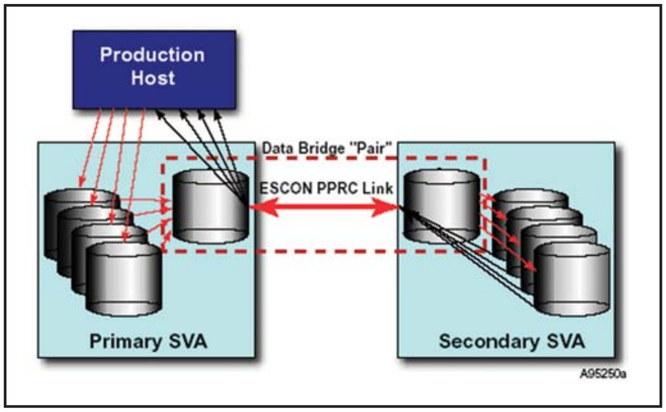

Figure 6. PowerPPRC — wide area network.

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. There is still a significant performance benefit (over standard PPRC) to using a data bridge in a direct mode (non-WAN) configuration since the data bridge chains the tracks to improve link efficiency as mentioned previously. See **Figure 5** on page 13.

There is yet another benefit to using bridges, which 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, a process called "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 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.

To address this problem, PowerPPRC uses a synchronization queue with the bridges. When a volume pair 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 that 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 doing synchronization work for 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.

Some are more interested in synchronization time than host response time. In order to get a faster total synchronization time for all the volumes, they must all be "syncing" at once. This requires that each volume do its own work and that volumes not be serialized on the bridges. The "establishment state" parameter of the "establish pair" (CESTPAIR) command ("alter device" command in open PPRC) give you a choice between faster host response time and faster total synchronization time. If the data bridge is established with the "copy all" option, all of the devices do their synchronizing work simultaneously, resulting in faster total synchronization times. If the bridge is established with the "no copy" parameter, then the bridge utilizes the SnapShot synchronization feature, resulting in a serialized synchronization process with less impact on host response time and subsystem resources.

2.10 PERFORMANCE EXPECTATIONS AND CONSIDERATIONS

In the past, performance and distance limitations severely limited the implementation and exploitation of standard PPRC. StorageTek's PowerPPRC significantly reduces these limitations.

There are two basic metrics regarding the difference in performance of standard PPRC and PowerPPRC: the synchronization time and the overhead encountered.

2.10.1 Synchronization

Because StorageTek's virtual disk architecture stores and transfers data in a compressed format, synchronization time with standard PPRC or PowerPPRC is usually significantly less than for that of the competition. Additionally, because of StorageTek's virtual architecture, only used (and compressed) space is transferred, whereas competing implementations copy at the volume level whether the space is used or not. Granted, when using WAN solutions, network-attachment products, such as UltraNet, compress the data. The StorageTek solution has the additional benefits of:

> Moving only allocated and used space (as opposed to moving complete volumes)
> Moving compressed data
> Supplying any additional compression provided by the network-attachment products.

Volume synchronization can be achieved in one to two minutes, based on customer data. This performance is typically much better than that of any competing implementation.

2.10.2 PPRC overhead and distance considerations

The basic arithmetic of disk subsystem performance analysis centers on the four basic I/O operations, which are read hits, read misses, write hits and write misses. With the high cache hit rates of most mainframe workloads, performance is still quite good even though there are some I/Os that take a long time in every data center. For example, **Table 3** depicts a V2X disk subsystem performing 1,000 I/Os per second with 90 percent reads, 90 percent read hits and 90 percent write hits. Assuming that cache hits take 0.6 milliseconds and cache misses take 6 milliseconds, we can easily compute the response time of this subsystem (**Table 3**).

Type of I/O	Number	Response time(ms) Per I/O	Total time (ms)
Read hits	810	0.6	486
Read misses	90	6.0	540
Write hits	90	0.7	63
Write misses	10	6.0	60
Total	1,000		1,149
Average response time (ms)			1.2

Table 3.

If PowerPPRC were imposed on this workload, what would be the new response time? With standard or PowerPPRC, only writes are sent to the secondary subsystem after the volumes are initially synchronized. If we assume the distance between the two SVAs is 100 kilometers, each write incurs a 3-millisecond overhead due to cable propagation delay and additional overhead incurred. Remember, though, only the write data are changed.

In this case, the average response time increases from 1.2 ms to 1.5 ms (**Table 4**).

Type of I/O	Number	Response Time per I/O (ms)	Total Time (ms)
Read hits	810	0.6	486
Read misses	90	6.0	540
Write hits	90	3.7	333
Write misses	10	91	90
Total	1,000		1,449
Average response time (ms)			1.5
1. Add 3 ms (2 protocol overhead +1 distance latency) to each write (5 ms for the write miss in Table 3) for a total of 9.			

Table 4.

That's an increase of less than one-third of a millisecond to provide the utmost in data protection! Response times can be calculated for your installation using data from Resource Measurement Facility (RMF) and the cache reporter.

The performance of the StorageTek PowerPPRC is outstanding! PowerPPRC necessarily adds some overhead to the system for the writes that go to the remote control unit. If the primary and secondary units are side by side, this overhead is usually between 1 millisecond and 2 milliseconds; 2 milliseconds will be used for the remainder of this discussion. Cable propagation delay adds about 5 microseconds for each kilometer. So, for a distance of 100 kilometers (that's 200 kilometers round trip), there will be a delay of approximately 3 millisecond for writes (2 millisecond basic overhead plus 1 millisecond for distance delay).

For a 500-kilometer configuration (1,000 kkilometer round trip), there will be a delay of only approximately 7 milliseconds for writes (2 milliseconds basic overhead plus 5 milliseconds for distance delay). Referring to the response time analysis presented in Table 4, the overall impact of PowerPPRC on this 500-kilometer operation would be less than a millisecond (0.7 milliseconds). As in all systems, care must be taken to ensure a properly balanced system in orderto avoid the bottlenecks. Every system has bottlenecks. The real issue is whether you encounter and observe these limitations.

Before we leave the topic of performance, we should note the significance of PowerPPRC's performance improvements in comparison with other products in the industry. IBM developed an asynchronous alternative to PPRC called eXtended Remote Copy (XRC) in which the primary and secondary volumes are not always kept exactly alike. XRC was developed in part to overcome the performance penalties imposed by the old version of PPRC. StorageTek's high-performance PowerPPRC solution reduces the need for many asynchronous alternatives while providing the highest data integrity. PowerPPRC also removes the copy burden incurred by XRC from the host and its channels. With PowerPPRC, the disk subsystem, not the host system, handles the data traffic. IBM has recently introduced a product called PPRC-XD, which allows for asynchronous mirroring between subsystems. As with XRC, "in-flight" updates are still subject to loss in certain circumstances.

Figure 7 shows a comparison of PowerPPRC direct mode at differing distances with a uniform time-sharing option (TSO) workload (no skew). As you can see from Figure 7, the SVA V960 provides a broad plateau of consistent performance that supports most environments. The small increase in response time supports most campus and many "across-town" DR requirements. The new SVA V2X with its switched-fibre back end will provide even lower response times and will therefore extend the acceptable boundary of performance depicted by the shaded area in FIgure 7. When you combine this level of performance with PPRC SnapShot and Snap-to- Primary features, it is arguably the best solution available in the industry today.

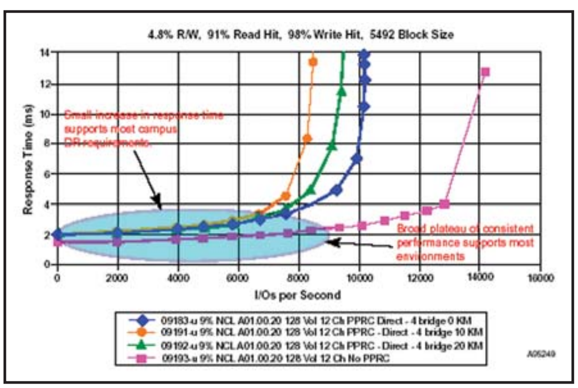

Figure 7. Performance comparisons. The benchmark data used to produce this graph was created using an SVA V960.

3 CASE STUDIES

3.1 CASE STUDY ONE: POWERPPRC DIRECT MODE

Although theoretical and laboratory data are useful, the real proof is found in what this product is doing for you. Based on 12 hours of batch runs, a large customer measured less than one-half of a millisecond average I/O elongation while using PowerPPRC direct mode. Although this test was performed over a short distance, the results are indicative of the low overhead associated with PowerPPRC direct mode. Increased distance would add only 5 microseconds per kilometer. Of course, your implementation will be different, but this example can provide you with a sense of what to expect. See **Table 5**.

PowerPPRC direct mode impact on natch response time (10% write rate)		
Method of operation	**I/Os per second**	**Service time (ms)**
Without PPRC	270	2.45
With PowerPPRC	285	2.94
Power PPRC Overhead		0.49

Table 5.

3.2 CASE STUDY TWO: POWERPPRC WAN MODE

If you separate the SVAs by a distance of 23 kilometers, Figure 6 on page 13 would depict the customer's actual environment in a PowerPPRC WAN mode configuration. The customer also evaluated another vendor's configuration but finally selected the SVA and PowerPPRC configuration. The customer is currently mirroring around nine terabytes using this configuration.

4 CONCLUSION

PowerPPRC now offers both synchronous and asynchronous solutions to your DR requirements. It offers features like PPRC SnapShot and Snap-to-Primary that cannot be duplicated or even emulated by our competition. It is the simplest, most efficient and most cost-effective business continuity alternative on the market today.

When you consider the total cost per mirrored megabyte, significant performance improvements of the latest version of the SVA, configuration flexibility, ease of management and the other tangible cost savings derived from virtual architecture, the clear choice for business continuity is StorageTek with SVA and the PowerPPRC product suite!

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 9152 B 04/04