BROCADE

WHITE PAPER

Enhanced Data Availability and Business Continuity with the IBM TS7700 and Brocade 7840 Extension Solution

Data availability and business continuity offer a vital competitive edge that is crucial to an organization's success. From finance and healthcare to transportation and retail, to state, local, and federal government, all such organizations must adopt proven business continuity and recovery management strategies and storage technologies to successfully address operational risk, availability, and security challenges. One way to do this is to plan and implement a resilient mainframe and extension solution that gives an organization the ability to respond and adapt to a wide variety of external and internal demands, disruptions, disturbances, and threats, while continuing business operations without any significant impact.

A key component in a resilient IT architecture for IBM z Systems customers is the IBM TS7700 with its multicluster grid configuration. The TS7700 Grid configuration is a series of clusters connected by a network to form a high-availability, resilient virtual tape storage architecture. Logical volume attributes and data are replicated via Internet Protocol (IP) across these clusters, which are joined by the grid network. However, to the host, the grid configuration looks like a single storage subsystem. This ensures high availability and ensures that production work continues, even if an individual cluster becomes unavailable.

The Brocade® 7840 Extension Switch and Brocade Fabric Vision™ technology enable TS7700 Grid end users to gain significant increases in performance, security, reliability, and availability between clusters in the grid network. This paper details the features and advantages provided by Brocade Extension technology in a TS7700 Grid network.

Introduction: Resilient IT Architectures

A resilient IT architecture gives an organization the ability to respond and adapt to a wide variety of external and internal demands, disruptions, disturbances, and threats while continuing business operations without any significant impact. Although related to planning for disaster recovery, planning for a resilient IT architecture is much broader in scope. A resilient IT architecture requires organizations to go beyond planning for recovery from an unplanned outage. In fact, a resilient IT architecture enables organizations to avoid outages entirely in other words, to ensure business continuance.

In any business continuity architecture, the most expensive component typically is the network bandwidth for connectivity between sites. The goal in designing a TS7700 Grid architecture is to maximize IT resiliency at the lowest cost. In other words, this architecture maximizes network efficiency by means of connectivity devices that offer the highest performance, highest availability, and lowest operating costs, and that are easy to manage. The Brocade 7840 Extension Switch is such a device.

With an IBM TS7700 Grid network that is connected to the Brocade 7840, an IBM z Systems customer can achieve a highly resilient architecture that is performance optimized and highly available, with low operating costs. This architecture simultaneously simplifies management of the entire TS7700 network environment by using a "single pane of glass" in conjunction with Brocade Network Advisor and Brocade Fabric Vision management software.

This white paper explains the basics of the IBM TS7700 Grid solution and describes the features and technology available from Brocade that help ensure an optimal TS7700 Grid solution.

Challenges in Designing a Resilient IT Architecture

Few events can generate a stronger adverse business impact as an IT outage, even when it lasts for only a few minutes. Even worse is the negative publicity that such events often generate in today's "always connected" news media world. Clients, partners, and the market that a business operates in are driving the need for continuously available, resilient IT architectures. The potential revenue loss from an outage and the damage to a company's reputation can be costly. Add to those problems the possible government sanctions in an increasingly complex regulatory environment, creating even more challenges.

Companies across a diverse spectrum of industries face these growing challenges when planning a resilient IT architecture. Executive focus on resilient IT architecture solutions, such as the IBM TS7700 Grid, is increasing, because it is no longer enough to merely plan for recovery from a disaster. Enterprises that run IBM z Systems need a greater level of availability to deal with a range of contingencies. These contingencies can be as simple as inadvertent power loss or common configuration errors, or as complex as major natural disasters (for instance, hurricanes or earthquakes) or human-induced disasters (for instance, terrorism). A business that plans for a resilient IT architecture must of necessity implement both traditional disaster recovery planning and additional planning for continuous availability.

As data centers and data stores grow, tape operations become more complex. Growth can also lead to increased tape processing times, high management expenses, and skyrocketing hardware costs. For high availability and disaster recovery, the IBM TS7700 solution can be deployed in several networked, multisite grid configurations. Each grid configuration is optimized to help eliminate downtime from planned and unplanned outages, upgrades, and maintenance. The TS7700 Grid enhances the resilience of an organization's IT architecture.

The network infrastructure that supports a TS7700 Grid solution faces challenges and requirements of its own. First, the network components need to individually provide reliability, high availability, and resiliency. The overall solution is only as good as its individual parts. A TS7700 Grid network requires nonstop, predictable performance with components that have "five-9s" availability, in other words, 99.999% uptime. Second, a TS7700 Grid network must be designed with highly efficient components that minimize operating costs. Third, today's rapidly changing and growing amounts of data require a TS7700 Grid network that meets these specifications:

- Uses highly scalable components that support business and data growth and application needs
- Helps accelerate the deployment of new technologies as they become available

IBM TS7700 Grid Basics

The technology of a prior generation produced the IBM Virtual Tape Server (VTS), which had a feature called Peer-to-Peer (PtP) VTS capabilities. PtP VTS was a multisite-capable business continuity and disaster recovery solution. PtP VTS was to tape what PtP Remote Copy (PPRC) was to Direct Access Storage Devices (DASDs). PtP VTS-to-VTS data transmission was originally done by Enterprise Systems Connection (IBM ESCON), then IBM FICON®, and finally Transmission Control Protocol/Internet Protocol (TCP/IP).

With the TS7700, the virtual tape controllers and remote channel extension hardware for the PtP VTS of the prior generation were eliminated. This change provided the potential for significant simplification in the infrastructure that is needed for a business continuity solution and for simplified management. Hosts attach directly to the TS7700s. Instead of FICON or ESCON, the connections between the TS7700 clusters use standard TCP/IP. Similar to the PtP VTS of the previous generation, with the new TS7700 Grid configuration, data can be replicated between the clusters, based on the customer's established policies. Any data can be accessed through any of the TS7700 clusters, regardless of which system the data is on, if the grid contains at least one available copy.

As a business continuity solution for high availability and disaster recovery, multiple TS7700 clusters are interconnected by using standard Ethernet connections. Local and geographically separated connections are supported to provide a great amount of flexibility to address customer needs. This IP network for data replication between TS7700 clusters is more commonly known as a TS7700 Grid. A TS7700 Grid refers to 2 to 6 physically separate TS7700 clusters that are connected to each other with a customer-supplied IP network. The TCP/IP infrastructure that connects a TS7700 Grid

is known as the grid network. The grid configuration is used to form a high-availability disaster recovery solution and to provide metro and remote logical volume replication. The clusters in a TS7700 Grid can be, but do not need to be, geographically dispersed. In a multiple-cluster grid configuration, two TS7700 clusters are often located within 100 kilometers (km) of each other. The remaining clusters can be more than 1,000 km away. This solution provides a highly available and redundant regional solution. It also provides a remote disaster recovery solution outside of the region.

The TS7700 Grid configuration introduces new flexibility for designing business continuity solutions. Peer-to-peer communication capability is integrated into the base architecture and design. No special hardware is required to interconnect the TS7700s. The Virtual Tape Controllers (VTCs) of the previous generations of PtP VTS are eliminated, and the interconnection interface is changed to standard IP networking. If configured for high availability, host connectivity to the virtual device addresses in two or more TS7700s is required to maintain access to data, if one of the TS7700s fails. If the TS7700s are at different sites, channel extension equipment is required to extend the host connections.

With the TS7700 Grid, data is replicated and stored in a remote location to support truly continuous uptime. The IBM TS7700 includes multiple modes of synchronous and asynchronous replication. Replication modes can be assigned to data volumes by using the IBM Data Facility Storage Management Subsystem (DFSMS) policy. This policy provides flexibility in implementing business continuity solutions, so that organizations can simplify their storage environments and optimize storage utilization. This functionality is similar to IBM Metro Mirror and Global Mirror with advanced copy services support for IBM z Systems customers.

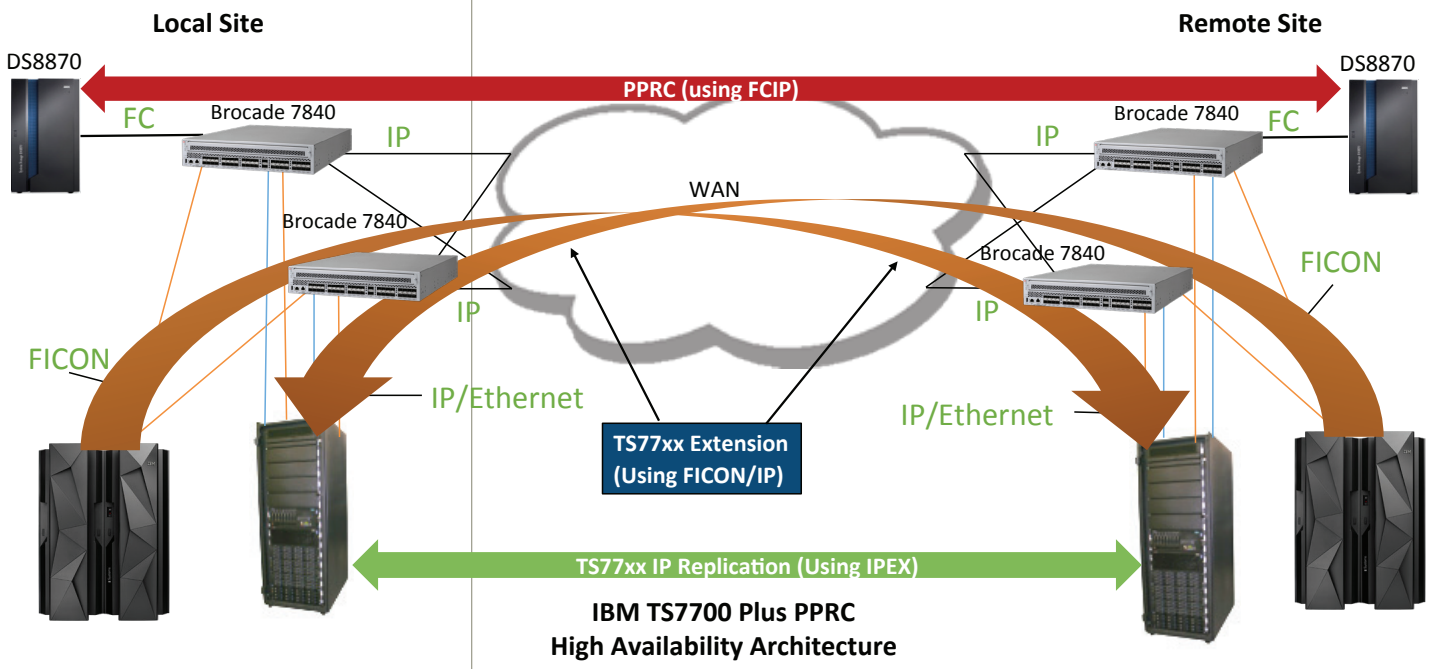

Figure 1: IBM TS7700 Grid plus PPRC high-availability architecture.

The TS7700 Grid is a robust business continuity and IT resilience solution. By using the TS7700 Grid, organizations can move beyond the inadequacies of on-site backup (disk-to-disk or disk-to-tape) that cannot protect against regional (nonlocal) natural or human-induced disasters. By using the TS7700 Grid, data can be created and accessed remotely through the grid network. Many TS7700 Grid configurations rely on this remote access to further increase the importance of the TCP/IP fabric.

With increased storage flexibility, an organization can adapt quickly and dynamically to changing business environments. Switching production to a peer TS7700 can be accomplished in a few seconds with minimal operator skills. With a TS7700 Grid solution, z Systems customers can eliminate planned and unplanned downtime. This approach can potentially save thousands of dollars in lost time and business and can address today's stringent government and institutional data protection regulations.

A TS7700 Grid network requires nonstop predictable performance with components that have "five-9s" availability. A TS7700 Grid network must be designed with highly efficient components that minimize operating costs. These components must also be highly scalable, to support business and data growth and application needs and to help accelerate the deployment of new technologies. The Brocade 7840 with Brocade Fabric Vision technology provides exactly these benefits.

Brocade 7840 Extension Switch Basics

The Brocade 7840 is a purpose-built extension solution. It is an enterprise-class product that is characterized by an essential feature set: excellent performance, increased security, high reliability, proactive monitoring, flow visibility, and diagnostic tools. The Brocade 7840 is an ideal platform for building a high-performance data center extension infrastructure for replication and backup solutions. It leverages any type of inter-data center Wide Area Network (WAN) transport to extend open systems and mainframe storage applications over any distance. Without the use of extension, those distances are often impossible or impractical. In addition, the Brocade 7840 addresses the most demanding disaster recovery requirements. With twenty-four 16 gigabit-per-second (Gbps) Fibre Channel (FC)/FICON ports, sixteen 1/10 Gigabit Ethernet (GbE) ports, and two 40 GbE ports, customers can achieve the bandwidth, port density, and throughput required for maximum application performance over WAN connections.

Brocade technology integrates perfectly into any IP network and provides a highly efficient data transport capable of full bandwidth utilization across great distances. The defining features that bring value to Brocade Extension in IBM TS7700 Grid environments include Extension Hot Code Load (HCL), Extension Trunking, WAN-Optimized TCP (WO-TCP), IP security (IPsec), Adaptive Rate Limiting (ARL), Brocade Fabric Vision technology, and Brocade Network Advisor. Brocade also provides a full spectrum of security features and connectivity validation tools. Overall, Brocade Extension products leverage 20 years of distance connectivity innovation and thought leadership, as demonstrated by the fact that they are the market's preferred extension solution.

The Brocade 7840 is capable of both Fibre Channel over IP (FCIP) extension and IP Extension. Therefore, the Brocade 7840 provides an ideal platform for TS7700 Grid connectivity, for both high-availability and disaster recovery configurations. This section

of the white paper discusses some basics of the Brocade 7840. The remainder of the white paper does not focus on the FCIP Extension capabilities of the 7840; rather, it focuses on IP Extension specifics as they pertain to the IBM TS7700 Grid replication.

Brocade 7840 Extension HCL

Extension HCL was introduced to the storage industry with the Brocade 7840. Firmware upgrades can be done without tunnel disruption. A firmware update can take considerable time, too much time to allow a large extension connection to stay down. Years ago, WAN links had much less bandwidth, and it was not paramount to maintain connectivity during firmware updates. The interim backlog of data was acceptably small. However, by today's standards the amount of backlog data during a firmware upgrade can be significant, on the order of half a terabyte or more when using one 10 Gbps connection. At many enterprises, to comply with Recovery Point Objective (RPO) policy and to maintain a comfort level for Storage Administrators, nonstop operations of TS7700 Grid replication are required. The Brocade 7840 is the only product on the market that maintains extension connectivity during a firmware upgrade. Extension HCL from Brocade is lossless and always keeps data in-order. During the firmware update process no data is lost, and all data sent to Upper Layer Protocol (ULP) is consistent and in order. This means that Extension HCL can be used in mainframe environments without causing Interface Control Checks (IFCCs), which is a testament to the underlying advancements to this technology.

Security of IP Extension Flows: IPsec

Unsecured data leaving the data center potentially could cause data breaches and even unwanted publicity—for an enterprise. Increasingly, end users are facing requirements to encrypt all data that leaves the data center (encryption in flight). These requirements are typically driven by government regulation, internal audit requirements, or a combination of the two. Any data that leaves the safe confines of the data center should be protected by using encryption. Encryption applies not only to the public Internet. Private WAN connections are not secure outside of the data center.

Brocade has developed hardware-based IPsec for secure data in-flight across Brocade Extension Inter-Switch Links (ISLs). Brocade IPsec operates at line-rate and introduces only a couple of microseconds (µs) of added latency, making it useful for synchronous applications. Brocade IPsec uses AES-GCM-256, Diffie-Hellman 2048-bit Modular Exponential (MODP), Internet Key Exchange version 2 (IKEv2), Hashed Message Authentication Mode Secure Hash Algorithm 512 (HMAC-SHA2-512), and Transport Mode, and it is rekeyed every few hours without disruption. A Pre-Shared Key (PSK) is configured per tunnel and trunk on each side.

Best practice with TS7700 Grid networks is to use Brocade IPsec for extension. Brocade IPsec is part of circuit formation. It protects data from virtually every type of attack, including sniffers, data modification, identity spoofing, man-in-the-middle, and Denial of Service (DoS). Brocade IPsec requires no additional licenses or costs and is simple to configure. IPsec plus Extension Trunking provides the ability to granularly load balance encrypted storage flows across all the trunk's member circuits. Up to 20 Gbps is supported for a single trunk, and two such trunks are supported per Brocade 7840 switch. This is a large amount of encrypted load-balanced data bandwidth (40 Gbps) for

a single box. IPsec also prevents the need for costly and complex firewalling. Because firewalls are software-based, they tend to provide poor performance.

The Brocade 7840 with IPsec encourages the use of encrypted data protection and does so with no performance penalty. The Brocade IPsec capability is hardware-implemented and operates at line rate (20 Gbps) per data processor, with 5 µs of added latency. IPsec is included in the Brocade 7840 base unit, with no additional licenses or fees. The Brocade IPsec implementation offers better performance than the TS7700 native-based encryption solution, while providing prudent security at no additional cost.

Acceleration of IP Extension Flows: WO-TCP

TCP is central to the high-speed transport of large data sets that are common in storage extension. Through years of experience, Brocade has developed an aggressive TCP stack called WO-TCP. WO-TCP is a specialized transport that cannot be outperformed by competing WAN optimization products. In other words, WAN optimization provides a negligible benefit when using the Brocade 7840. Overall, Brocade technology is comparable from the perspective of the data transport bottom line. The total bytes transferred within the same period of time, over the same bandwidth, will be virtually the same compared to competing WAN optimization products. All of these benefits are provided, plus the added satisfaction that the cost of purchasing Brocade Extension is considerably less compared to WAN optimization products.

Acceleration of flows across the WAN improves IP storage performance dramatically. Long distance increases latency and is prone to packet loss. Tested applications have demonstrated improvements of up to 50 times, due to the ability to handle latency and packet loss without performance degradation. This performance has nothing to do with compression; any compression achievable is in addition to flow acceleration. Flow acceleration is purely a function of enhanced protocol efficiency across the network.

IP Extension on the Brocade 7840 terminates IP storage TCP flows locally and transports the data across the WAN using WO-TCP. The primary benefit here is the local ACK. By limiting ACKs to the local data center, TCP that originates from an end IP storage device has to be capable of merely high-speed transport within the data center. Most native IP storage TCP stacks are capable only of high speeds over short distances. Beyond the limits of the data center, "droop" becomes a significant factor. Droop refers to the inability of TCP to maintain line rate across distance. Droop worsens progressively as distance increases.

WO-TCP is an aggressive TCP stack designed for Big Data movement, operating on the purpose-built hardware of the Brocade 7840 switch. The Brocade 7840 offers 64 processors and 128 gigabytes (GB) of RAM to support WO-TCP. In comparison, WO-TCP has no droop across two 10 Gbps connections, up to 160 milliseconds (ms) Round-Trip Time (RTT) per data processor. This is equivalent to two fully utilized 10 Gbps WAN connections (OC-192) between Los Angeles and Hong Kong.

Another benefit is the absence of Head of Line Blocking (HoLB) or Slow Drain Device (SDD) problems with WO-TCP streams. WO-TCP on the Brocade 7840 implements a feature called "streams." Streams are used to mitigate HoLB and SDD problems across the IP network and WAN. If all IP storage flows are flow-controlled using a single

TCP Receiver Window (rwnd), then all those flows slow or halt in the event that the TCP Receiver Window is shut or closed. This situation is detrimental to all applications except for the one that flow control is meant for.

It is not practical to create a separate TCP connection for each flow, as this consumes excessive resources. Instead, autonomous streams are created using virtual TCP windows for each stream. The Brocade 7840 can accommodate hundreds of streams per data processor. Two data processors are present for each Brocade 7840 switch. Because a virtual TCP window is used for each stream, if a flow needs to slow down or stop, no other flows are affected; they continue to run at their full rate.

WO-TCP integrates with ARL, and the synergy of these two technologies creates an industry-dominating transport for storage. No similar transport exists on any storage array or virtual tape system-based native IP replication. Clearly, WO-TCP demonstrates the enterprise-class performance superiority of the Brocade 7840 for TS7700 Grid implementations.Table 1. PoE technologies and applications.

Table 1: WO-TCP vs. native TS7700 replication.

Latency (ms) RTT	No Packet Drop		1% Packet Drop	
	TS7700 with IP Extension	TS7700 Native	TS7700 with IP Extension	TS7700 Native
2	590	590	560	16
20	590	590	560	12
100	560	65	540	5
260	540	25	540	3

* Results in megabytes per second

Bandwidth Management and Pooling

Bandwidth management and bandwidth pooling form a feature set that provides aggregate bandwidth from multiple sources (including high availability) and management of that bandwidth. Bandwidth management and pooling use the exclusive Extension Trunking technology on the Brocade 7840 switch. The Brocade 7840 switch also supports jumbo frames. Even if the IP network or WAN does not support jumbo frames, replication devices can still use Local Area Network (LAN)-side jumbo frames, which should offload CPU workload from the device and accelerate replication.

Brocade Extension Trunking

Extension Trunking is a Brocade technology originally developed for mainframes and now broadly used in open systems environments. Extension Trunking has evolved to include IP Extension flows. Extension Trunking bundles multiple circuits together into a single logical trunk. Those circuits can span multiple service providers and different data center LAN switches for redundancy. Bandwidth is managed in such a way that if a data center LAN switch goes offline or encounters any disruption along the path, the bandwidth of the remaining paths adjusts to compensate for the offline path. With the

proper design, bandwidth can be maintained during outages of various devices in the pathway. Extension Trunking shields end devices from IP network disruptions, making network path failures transparent to replication traffic. Multiple circuits (two or more) from the Brocade 7840 are applied to various paths across the IP network. With each added circuit, even more bandwidth is added to the pool. Extension Trunking performs a Deficit Weighted Round Robin (DWRR) schedule when placing batches into the egress. Batches are an efficiency technique used by Brocade to assemble frames into compressed byte streams for transport across the WAN. A feature of Extension Trunking called Lossless Link Loss (LLL) ensures lossless data transmission across the trunk in the event that data is lost in-flight due to an offline circuit, and WO-TCP is no longer operational across that circuit. WO-TCP itself recovers lost or corrupted data across a link, if that circuit is still operational. All data is delivered to the ULP in-order.

The most common example of this is redundant Data Center LAN (DC-LAN) switches. For instance, one circuit goes over DC-LAN Switch A, and the other circuit goes to DC-LAN Switch B. This is a simple and effective architecture for redundancy and increased availability. Of course, as needs dictate, the application of circuits over various paths and service providers (up to eight circuits per IP Extension tunnel) can establish a highly available infrastructure for IP storage.

Extension Trunking performs failover and failback, and no data is lost or delivered out-of-order during such events. Circuits can be designated as backup circuits, which are passive until all the active circuits within the failover group have gone offline. This protects users against a WAN link failure and avoids a restart or resync event. Extension Trunking supports aggregation of multiple WAN connections with different latency or throughput characteristics (up to a 4:1 ratio), allowing WAN circuits to be procured from multiple service providers with different physical routes, to ensure maximum availability. If all WAN circuits are from the same service provider, then chances are high that a single failure event (for example, equipment failure, power loss, or cable cut) can take down all WAN circuits at one time. With Extension Trunking, organizations can protect their replication traffic from these kinds of outage events.

Extension Trunking offers more than the ability to load balance and fail over or fail back data across circuits. Extension Trunking is always a lossless function, providing in-order delivery within an extension trunk (defined by a Virtual Expansion Port, or VE_Port). Even when data in-flight is lost due to a path failure, data is retransmitted over remaining circuits via TCP and placed back in-order before it is sent to ULP. IP storage applications are never subjected to lost data or out-of-order data across the WAN.

Normally, when IP Extension is not used, packet loss in the IP network results in extremely low performance on native IP storage TCP/IP stacks. These stacks have little to no tolerance for packet loss across the IP network. On popular Network-Attached Storage (NAS) platforms during periods of 0.1 percent packet loss and 5 ms RTT latency, the reduction in throughput is 95 percent or more.

The use of the Brocade Extension Trunking feature with IP Extension, for TS7700 Grid implementations offers significant performance and availability enhancements that are not available with other solutions.

Adaptive Rate Limiting (ARL)

Rate Limiting (ARL) is used to maximize WAN utilization while sharing a link with other nonstorage applications. Native TCP/IP stacks on IP storage applications do not have the ability to adaptively alter rate limiting based on conditions in the IP network, but the Brocade 7840 Extension Switch does have this ability. Brocade Adaptive Rate Limiting dynamically adjusts shared bandwidth between minimum and maximum rate limits and drives maximum I/O even with a downed redundant system.

Brocade ARL automatically adjusts the rate limiting on all associated circuits that replicate across the IP network, regardless of the ingress FC device and the WAN path or paths. ARL automatically adjusts rate limiting when other Brocade Extension circuits go online and offline or when the available IP bandwidth that is being experienced changes. ARL works across all Brocade Extension products, using the same WAN infrastructure.

ARL is designed to work on WAN connections that are shared with other IP storage or nonstorage applications. Array auto-adjust rate limiting was not designed for such instances. In fact, the Brocade 7840 can be configured so that during an outage, high-priority applications maintain their bandwidth, while lower-priority devices sacrifice their bandwidth. ARL dynamically adjusts rate limits independent to each circuit, permitting efficient use of WO-TCP across a variety of ever-changing WAN environments. In this example, during the WAN service outage the overall bandwidth is halved, and the Brocade ARL, integrated with WO-TCP, best utilizes the available bandwidth, while maintaining nonstop operations.

ARL is used with Extension Trunking to maintain available bandwidth to storage applications. For example, if DC-LAN Switch A goes down, then as long as DC-LAN Switch B remains online and has ample bandwidth connectivity, it should be able to maintain all of the original bandwidth to the application. In this case, it is necessary for rate limiting to readjust upward and compensate for the lost pathway. Rate limiting is used to prevent oversubscribing the WAN and any associated contention or congestion. Congestion events force TCP to perform flow control, which is extremely inefficient, slow to react, and results in poor performance. ARL adjusts from a normal condition that is not oversubscribed to an outage condition that maintains the same bandwidth. Clearly, this is essential to continuous availability. ARL, in conjunction with the Extension Trunking feature, is an ideally unique performance optimizer for TS7700 Grid implementations.

Prioritizing IP Extension Flows (QoS)

Frequently, the IP network does not have QoS configured, at least for storage applications. Therefore, at a minimum, it is important to deliver data to the IP network that is sequenced according to the Storage Administrator's priorities. Brocade Extension is located at the endpoints of the data transport, the TCP points of origin and termination. These endpoints are the most effective place to QoS-mark data and apply it to various applications. Prioritization of flows across the WAN using QoS can be achieved in various ways. The first and simplest method is to configure priorities on the Brocade 7840 and feed the prioritized flows into the IP network. Three priorities exist for FCIP: high, medium, and low, and three exist for IP Extension: high, medium, and low, for a total of six

priorities. In addition, the percentage of bandwidth during contention that is apportioned to IP Extension and FCIP is configurable. When there is no contention for bandwidth, all available bandwidth can be utilized by a flow. This first method prioritizes flows that are connected to the LAN side of the Brocade 7840. The Brocade 7840 manages bandwidth when sending data to the WAN. The only network device that requires QoS configuration with this method is the Brocade 7840. It is up to the IP network and WAN to perform First In, First Out (FIFO) operations, so that QoS works properly.

The second method is to mark data flows as they exit the Brocade 7840 and enter the IP network. Data flows can be marked with IEEE 802.1P, which is part of 802.1Q Virtual LAN (VLAN) tagging or Differentiated Services Code Point (DSCP), or end-to-end IP-based QoS. The difficulty in using this method is that it requires the IP network to be configured to perform the proper actions based on the marking. If the IP network is not configured to do so, it does not prioritize the data flows. This usually involves a complex and sizable project on the IP network side, to categorize a diverse number of flows and assign priorities to the flows. A problem with this approach is that QoS in the IP network does not remain stable. Applications and priorities change over time.

On the Brocade 7840, features such as QoS, 802.1P, and DSCP marking are fully supported. Typically, Brocade 7840 users prioritize their flows within the FCIP+IP Extension tunnel. It is then possible to create a Service-Level Agreement (SLA) for the tunnel itself, so that QoS can be deployed in the manner best suited for the environment.

When the IP Extension features of the Brocade 7840 are used with the IBM TS7700 Grid, then enhanced reliability, scalability, security, and performance is achieved compared to simple IP switching connectivity on its own.

Brocade Fabric Vision Technology for TS7700 Grid Networks

The Brocade 7840 is one component of an overall system that works together to guard against disruption. The Brocade 7840 has certain features that can facilitate the quick resolution of support issues and determine the root cause of faults or degradation. Extending Brocade Fabric Vision technology over distance for TS7700 Grid implementations delivers increased visibility, pinpoints problems, and accelerates troubleshooting to maximize performance.

RPOs for mission-critical data typically require less than five seconds. Such RPOs are difficult to maintain when data rates are 10 Gbps or higher and when network problems are also present. Often, multiple vendors are involved, such as storage vendors, storage network vendors, IP network vendors, and WAN service providers. These situations can cost organizations considerable sums of money and potentially expose the business to data loss.

To address these problems within the specific context of extension, Brocade introduced Brocade Fabric Vision technology with advanced capabilities for the Brocade 7840. Brocade Fabric Vision includes monitoring, alerting, and reporting tools that are specific

to Brocade Extension. Diagnostic tools are also available that are useful for determining IP network validation and overall health. The objective is to quickly determine a root cause of degraded situations or outages and to expedite a return to normal operations as quickly as possible. Brocade Fabric Vision provides an innovative solution that enables management of the TS7700 Grid network in a much more proactive manner than previously possible. This section of the paper lists and briefly discusses these features and functionalities that are incorporated into Brocade Extension. For a more detailed discussion, please refer to the Brocade white paper, "The Benefits of Brocade Fabric Vision Technology for Disaster Recovery."

Brocade Network Advisor Dashboard

Brocade Network Advisor uses a customizable dashboard, as shown in Figure 2 on the following page. The Brocade Network Advisor dashboard makes visible on a "single pane of glass" the monitors, counters, and status indicators that are most important for the Brocade 7840 and TS7700 Grid environment. Organizations can choose from over one hundred dashboard items, or if the item they need does not exist, they can create the item. The Brocade Network Advisor dashboard management paradigm is a proven significant time saver for performance management and troubleshooting. When what happens in the TS7700 Grid network environment is clear, the goal of continuous uptime becomes a reality.

Figure 2: The Brocade Network Advisor dashboard.

Monitoring and Alerting Policy Suite (MAPS)

Brocade introduced MAPS for Brocade Fabric OS® (Brocade FOS) and Brocade Network Advisor to provide a comprehensive suite of monitors, alerts, actions, and reporting. MAPS assists operations in achieving higher availability, quicker troubleshooting, and infrastructure planning. MAPS provides a prebuilt, policy-based threshold monitoring and alerting tool that proactively monitors the storage extension network health, based on a comprehensive set of metrics at tunnel, circuit, and QoS layers. Administrators can configure multiple fabrics at a time, using predefined or customized rules and policies for specific ports or switch elements.

MAPS monitors utilization, packet loss, RTT, jitter, and state changes for tunnels and trunks,circuits, and PTQ. Each PTQ priority (class-F, low, medium, and high) is monitored independently and includes throughput, duplicate Acknowledgments (ACKs), WO-TCP packet count, packet loss, and slow-starts. MAPS is simple and easy to deploy with preset threshold levels and responses (conservative, moderate, and aggressive) based on Brocade best practices. As needed (though not required), virtually every element is customizable in MAPS. By leveraging prebuilt rules and policy-based templates, MAPS simplifies threshold configuration, monitoring, and alerting. Organizations can configure one, multiple, or all fabrics at once using common rules and policies, or they can customize policies for specific ports, switch elements and items—all through a single dialog. The integrated dashboard displays an overall switch health report, along with details on any out-of-range conditions. Administrators can quickly pinpoint potential issues and easily identify trends and other aberrant behaviors occurring within their fabric.

Flow Vision

The Flow Vision diagnostic tool enables administrators to identify, monitor, and analyze specific application and data flows in order to maximize performance, avoid congestion, and optimize resources. Flow Vision includes Flow Monitor, MAPS for Flow Monitor, and Flow Generator.

Visualization of flows through tunnels via Flow Vision is another advantage of using the Brocade 7840 in a TS7700 Grid environment. Not all flows are created equal, and a tunnel managed by Brocade allows administrators to visualize each application. To ensure that SLAs are being met, Storage Administrators monitor network and flow behavior. This is very difficult to accomplish if managed from each originating device and port.

Troubleshooting network flows often a difficult and daunting endeavor. To make matters worse, Storage Administrators are not familiar with IP networks, and IP Network Administrators are not familiar with storage. These two groups have very different cultures and operating guidelines. It is difficult for Storage Administrators to depend solely on Network Administrators to maintain their replication environment, which makes flow, TCP, circuit, and tunnel monitoring and visualization considerably more important to manage.

When troubleshooting storage flows, imagine that the flows fall into one of two categories: victims or perpetrators. If something goes wrong in the network, every flow becomes a victim. However, sometimes nothing is wrong with the network, and flows fall victim to

perpetrators. Perpetrator flows are flows that utilize excessive resources to the point that other flows fall victim. This frequently happens downstream from the storage handoff to IP networking. Brocade Extension provides features, functionality, and tools to deal with storage SLAs. Flows within the protection of Brocade Extension tunnels meet their SLAs when they face perpetrator flows.

The feature set called Flow Vision enables administrators to identify, monitor, and analyze specific application flows in order to simplify troubleshooting, maximize performance, avoid congestion, and optimize resources. The Brocade 7840 has the capability to monitor specific flows between F_Ports that are communicating end-to-end across the extension network. It is also possible to monitor flows that come in from an E_Port. Flow Vision is a component of Brocade FOS and the Gen5 ASIC. At Logical Unit Number (LUN) level granularity, Input/Output Operations per Second (IOPS) and data rate can be monitored. Flow Vision includes the following features:

- **Flow Monitor:** Provides comprehensive visibility into flows within the fabric, including the ability to automatically learn flows and nondisruptively monitor flow performance. Administrators can monitor all flows from a specific host to multiple targets/LUNs, from multiple hosts to a specific target/LUN, or across a specific ISL. Additionally, they can perform LUN-level monitoring of specific frame types to identify resource contention or congestion that is impacting application performance.
- **Flow Generator:** Provides a built-in traffic generator for pretesting and validating the data center infrastructure—including route verification and integrity of optics, cables, ports, back-end connections, and ISLs—for robustness before deploying applications.

Wtool

Wtool is an IP-specific tool for testing the WAN-side infrastructure on the Brocade 7840. Wtool creates data flows that use the same circuits configured in a tunnel or trunk. Since Wtool uses the same circuit, all the characteristics of that circuit remain viable during testing, including jumbo frames/Path Maximum Transmission Unit (PMTU), VLAN, IPv4/IPv6, and IPsec. If a circuit in a trunk is selected with Wtool, the trunk's other circuits remain online and operational while the selected circuit is decommissioned for testing.

Wtool runs in the background and for a specified amount of time. The amount of time that can be set is nearly limitless. Users can disconnect during the interim without halting or losing the test in progress. Also, multiple test sessions can run simultaneously. On command, Wtool reports the new results, as well as the results of the previous run: timestamp, throughput, RTT, packet loss, and out-of-order packets.

Brocade 7840 for Consolidation

The Brocade 7840 is capable of both traditional FICON FCIP and IP Extension. Therefore, it is an ideal platform to use in both TS7700 Grid disaster recovery and high-availability configurations. In the TS7700 high-availability architecture, it is necessary for the local TS7700 cluster to communicate with the remote cluster or clusters, and possibly for the remote cluster to communicate with the local cluster. In the event that either the local or remote cluster is offline, tape processing can continue for mission-critical applications, such as SAP. Both hosts have connectivity to the remote

cluster. The FICON connection across the WAN uses FCIP Extension over the same tunnel. Bandwidth and prioritization are managed by the Brocade 7840, to ensure reliable operation.

As for the IP connection between local and remote TS7700 clusters, instead of sharing the WAN connection with FCIP, which often is contentious and must be continuously monitored and managed, the TS7700 Grid IP connectivity is managed by a single WAN scheduler and is joined into the extension tunnel on the Brocade 7840. The Brocade 7840 optimally manages both the FCIP and IP Extension flows for optimal performance without contention or oversubscription on the WAN. Flows can further be managed with QoS, compression, IPsec, and ARL without involving long, complex projects on the network side that might cause ongoing operational issues. The Brocade 7840 accelerates data transfers across the WAN by using the Brocade technology called WO-TCP.

Finally, the Brocade 7840 has the ability to manage applications that operationally use a combination of FC+IP or FICON+IP. For example, the Brocade 7840 could be used for FICON based z/OS Global Mirror extension with the Brocade Advanced FICON Accelerator emulation technology, in conjunction with managing the TS7700 Grid IP replication via its IP Extension functionality and features.

Summary

The IBM TS7700 Grid solution is a key component to z Systems disaster recovery and business continuity implementations. The Brocade 7840 offers an innovative, unique IP Extension solution for TS7700 Grid implementations to gain significant increases in performance, security, and availability between data centers. These solutions, combined with the sophisticated tools of Brocade Fabric Vision technology, enable you to distinguish trouble with the network from storage array application problems. These effective tools facilitate more efficient support calls and faster problem resolution. The TS7700 Grid solution with the Brocade 7840 Extension Switch offers the end user a highly available, resilient mechanism for improving RPO and Recovery Time Objectives (RTO).

For more information about Brocade solutions, visit www.brocade.com..

About Brocade

Brocade networking solutions help organizations achieve their critical business initiatives as they transition to a world where applications and information reside anywhere. Today, Brocade is extending its proven data center expertise across the entire network with open, virtual, and efficient solutions built for consolidation, virtualization, and cloud computing. Learn more at www.brocade.com.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE