BROCADE

WHITE PAPER

The Benefits of Brocade Fabric Vision Technology for Disaster Recovery

TABLE OF CONTENTS

Extending Brocade® Fabric Vision™ technology over distance for disaster recovery delivers increased visibility, pinpoints problems, and accelerates troubleshooting to maximize the performance of critical applications. This tech brief provides an overview of the Fabric Vision features and functionality that are included within Brocade Extension solutions.

Introduction

Significant challenges face Storage Administrators in today's data centers. Storage consumes a large quantity of IP network resources. Those IP networks lack adequate visibility of storage flows and are sometimes unreliable, overly complex, and very inflexible. Greater complexity creates more opportunities for issues to occur. Storage and network administration provide different roles and responsibilities in an enterprise. Storage Administrators do not manage IP networks, and Network Administrators do not manage storage. When issues arise, each role can tend to blame the other for the problems.

Recovery Point Objectives (RPOs) for mission-critical data typically require a couple of seconds or less. Such RPOs are difficult to maintain when data rates are high (10 gigabits per second [Gbps] or above) and when network problems are also present. Such degraded states are difficult to troubleshoot. Often multiple vendors are involved, such as storage vendors, storage network vendors, IP network vendors, and Wide-Area Network (WAN) service providers. Such situations cost organizations considerable money and expose the business to data loss. The end result is that user expectations are not met.

To address these problems within the specific context of extension, Brocade has introduced the Brocade 7840 Extension Switch and the Brocade SX6 Extension Blade with advanced Fabric Vision capabilities. Fabric Vision includes a number of monitoring, alerting, and reporting tools specific to Brocade Extension. Additionally, diagnostic tools are available that are useful for determining IP network validation and health. The objective is to quickly determine the root cause of degraded situations or outages, to expedite a return to normal operations as quickly as possible.

The Situation

Brocade customers pose the question, "How can we resolve support issues more quickly and effectively?" Often storage vendor support organizations report that they are contacted after the Remote Data Replication (RDR) application is already down, resulting in emergency measures for disaster recovery. This situation is further aggravated by the inability to quickly pinpoint whether the problem is a network or storage issue. Both customers and storage vendors are greatly interested in the ability of Brocade to proactively monitor and effectively troubleshoot the local Fibre Channel (FC) connections, network device health, and the ability of the IP network to meet its Service Level Agreements (SLAs). Storage arrays are not capable of providing proactive warnings or identifying network problems.

Storage Administrators face conditions such as these:

- **Unreliable IP networks and the unknown within:** Storage Administrators have to rely on the experience, capabilities, availability, accuracy, and willingness of Network Administrators to provide feedback about their networks during performance assessment and troubleshooting. Often, this exchange of information can prolong the length of a degraded state.
- **More can go wrong, more often, than in the past:** Data center environments have grown in both scale and complexity. Virtualization of every type is used at every level, more now than ever. Simply, more opportunities arise for things to go wrong, and each issue affects other issues. Pinpointing what went wrong is not always a simple task.
- **Asymmetric behavior:** Does outbound data take one path and data returns take another path? Are the paths equal in quality? Can the IP network properly load balance?
- **Different administration groups:** At least two administrative groups (Storage and Network) exist within an enterprise, and each group has different roles, training, expertise, loyalties, goals, and responsibilities—possibly even different management personnel. These differences in culture can cause difficulties for Storage Administrators and Network Administrators alike. Storage personnel might not understand the network infrastructure and requirements, and Network personnel might not understand the storage infrastructure and requirements.
- **RPO at high data rates:** RPO has changed over time. The amount of data and number of applications managed has grown significantly. This means that RPO times must shrink. The amount of data processed in even a short amount of time has become significant and represents critical business transactions.
- **Catch-up time from outage backlogs:** Now that data rates are in the 10 Gbps range for many companies, when an outage or degraded state occurs, it is more difficult to reestablish a fully synchronized data consistent state useful for database restart. If the amount of time it takes to regain a consistent state is exceedingly long, the exposure to disaster becomes a liability.

Now, if replication is stressed, is slower than expected, and might even suspend from time to time, what does the timeline to resolution look like? Refer to Figure 1. Generally, a Storage Administrator is not aware of any problems until an error condition occurs. Storage Administration might place a request into network operations to investigate why replication is not performing optimally across the WAN.

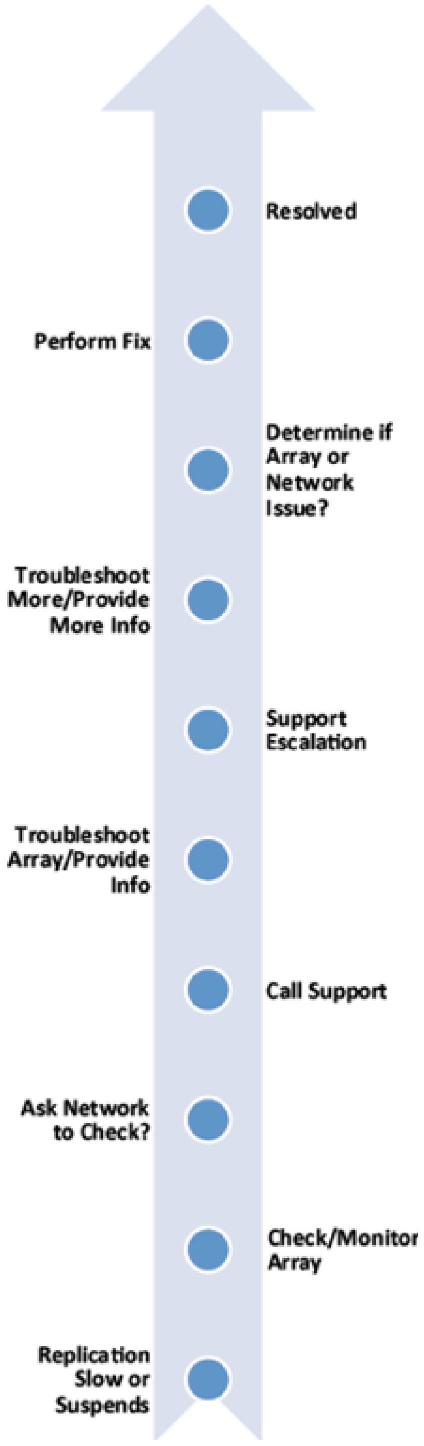
Resolved
Perform Fix
Determine if Array or Network Issue?
Troubleshoot More/Provide More Info
Support Escalation
Troubleshoot Array/Provide Info
Call Support
Ask Network to Check?
Check/Monitor Array
Replication Slow or Suspends

Figure 1: Typical outage timeline.

More often than not, the IP network will immediately be cleared of any issues. After all, if the IP network is not down, other applications are fine and a ping can get through, and if links are not fully utilized, then the assumption is that the network must be functional. To solve the quandary, it is likely that the Storage Administrator will open a support call with the storage array vendor to determine the cause of the issue. Rudimentary troubleshooting to validate basic issues will establish a baseline for the investigation. After that, any call that involves replication across an IP network requires establishing whether the problem exists in the network or in the array. At that point, the process of ruling out issues with the IP network commences.

How does the Brocade 7840 Extension Switch or the Brocade SX6 Extension Blade address this type of situation? Given the many aspects of storage networks, every aspect needs to be continuously monitored and inspected for proper operation. In this case, the Brocade 7840 and the Brocade SX6 monitor the IP network in which its managed tunnel flows, then alerts the Storage Administrator proactively when conditions arise that indicate degraded states. Degraded states within the IP network include: transient congestion events, chronic out of order delivery, instability and frequent network rerouting, oversubscription, data integrity problems, excessive latency, excessive jitter, and more. The capabilities of the Brocade 7840 and the Brocade SX6 enable Storage Administrators to gain visibility into an infrastructure on which they rely, but have had no visibility into before now.

This raises another problem: Storage Administrators cannot spend all of their time monitoring networks. The network itself must have the intelligence to create an alert when an issue becomes suspect. Brocade has incorporated decades of expertise and experience into improving network intelligence. When an alert is created by the Monitoring and Alerts Policy Suite (MAPS), it is based on the experience and expertise of the MAPS technology. This allows Storage Administrators to regain valuable time.

Quicker Time to Resolution

Brocade Fabric Vision technology provides an innovative hardware and software solution that simplifies monitoring, maximizes network availability, and dramatically reduces costs. Featuring monitoring, management, and diagnostic capabilities, Fabric Vision technology enables administrators to preempt problems before they impact operations. This all helps organizations meet Service Level Agreements (SLAs), primarily the SLA of maintaining the desired RPO. Refer to Figure 2 for improved timeline.

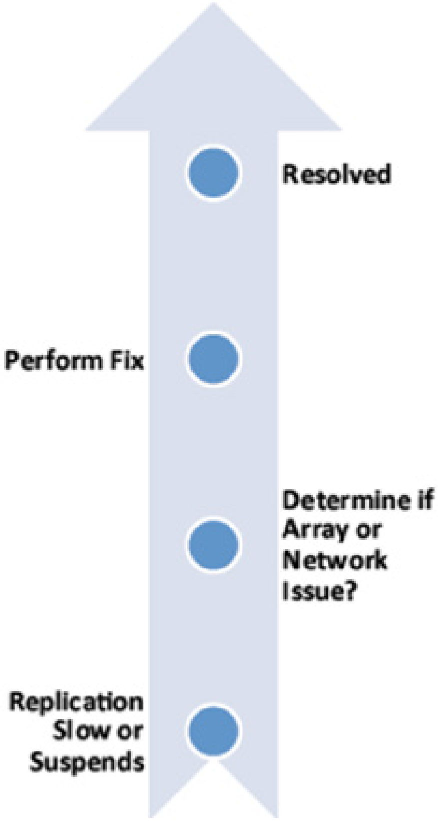
Resolved
Perform Fix
Determine if Array or Network Issue?
Replication Slow or Suspends

Figure 2: Preemptive outage timeline.

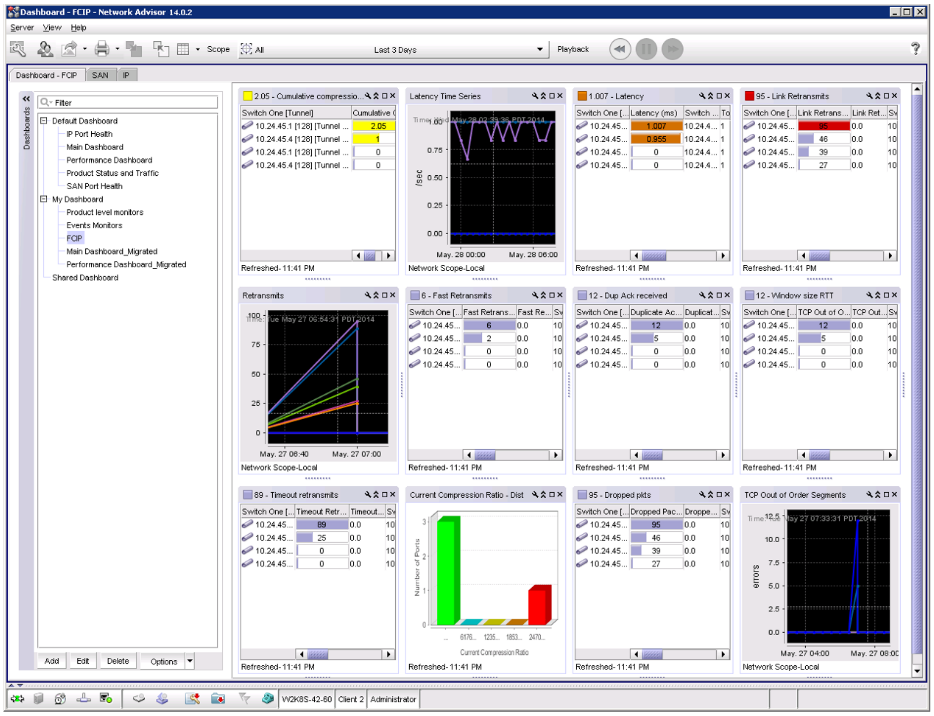

2.05 - Cumulative compressio...

Switch One [Tunnel]	Cumulative C
10.24.45.1 [128] [Tunnel ...	2.05
10.24.45.4 [128] [Tunnel ...	1
10.24.45.1 [128] [Tunnel ...	0
10.24.45.4 [128] [Tunnel ...	0

Refreshed- 11:41 PM

1.007 - Latency

Switch One [...	Latency (ms)	Switch ...	To
10.24.45...	1.007	10.24.4...	1
10.24.45...	0.955	10.24.4...	1
10.24.45...	0	10.24.4...	1
10.24.45...	0	10.24.4...	1

Refreshed- 11:41 PM

95 - Link Retransmits

Switch One [...	Link Retrans...	Link Ret...	Sv
10.24.45...	95	0.0	10
10.24.45...	46	0.0	10
10.24.45...	39	0.0	10
10.24.45...	27	0.0	10

Refreshed- 11:41 PM

6 - Fast Retransmits

Switch One [...	Fast Retrans...	Fast Re...	Sv
10.24.45...	6	0.0	10
10.24.45...	2	0.0	10
10.24.45...	0	0.0	10
10.24.45...	0	0.0	10

Refreshed- 11:41 PM

12 - Dup Ack received

Switch One [...	Duplicate Ac...	Duplicat...	Sv
10.24.45...	12	0.0	10
10.24.45...	5	0.0	10
10.24.45...	0	0.0	10
10.24.45...	0	0.0	10

Refreshed- 11:41 PM

12 - Window size RTT

Switch One [...	TCP Out of O...	TCP Out...	Sv
10.24.45...	12	0.0	10
10.24.45...	5	0.0	10
10.24.45...	0	0.0	10
10.24.45...	0	0.0	10

Refreshed- 11:41 PM

89 - Timeout retransmits

Switch One [...	Timeout Retr...	Timeout...	Sv
10.24.45...	89	0.0	10
10.24.45...	25	0.0	10
10.24.45...	0	0.0	10
10.24.45...	0	0.0	10

Refreshed- 11:41 PM

95 - Dropped pkts

Switch One [...	Dropped Pac...	Droppe...	Sv
10.24.45...	95	0.0	10
10.24.45...	46	0.0	10
10.24.45...	39	0.0	10
10.24.45...	27	0.0	10

Refreshed- 11:41 PM

Figure 3: Brocade Network Advisor dashboard.

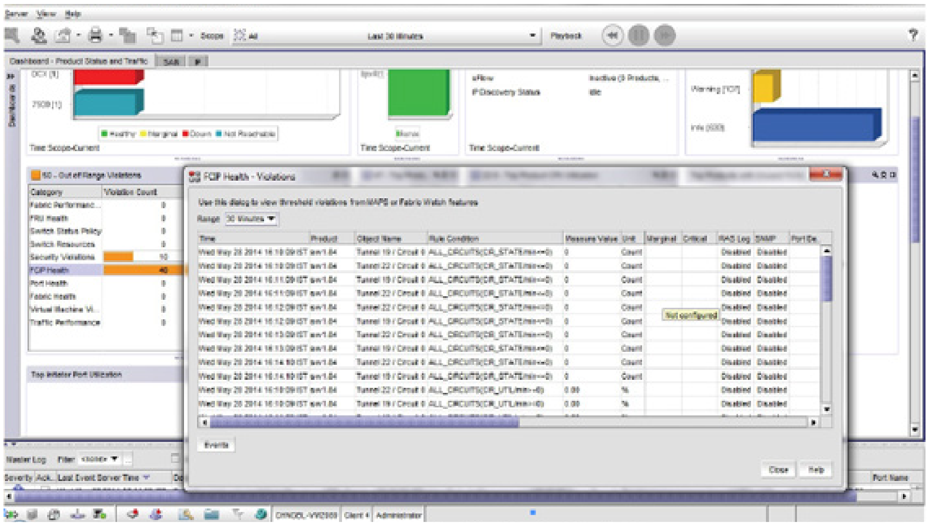

50 - Out of Range Violations

Category	Violation Count
Fabric Performanc...	0
FRU Health	0
Switch Status Policy	0
Switch Resources	0
Security Violations	10
FCIP Health	40
Port Health	0
Fabric Health	0
Virtual Machine Vi...	0
Traffic Performance	0

FCIP Health - Violations

Use this dialog to view threshold violations from MAPS or Fabric Watch features

Range 30 Minutes

Time	Product	Object Name	Rule Condition	Measure Value	Unit	Marginal	Critical	RAS Log	SNMP	Port De
Wed May 28 2014 16:10:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:10:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:11:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:11:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:12:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:12:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:13:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:13:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:14:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:14:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_STATE/min>=0)	0	Count			Disabled	Disabled	
Wed May 28 2014 16:10:09 IST	sw1.84	Tunnel 22 / Circuit 0	ALL_CIRCUITS(CIR_UTIL/min>=0)	0.00	%			Disabled	Disabled	
Wed May 28 2014 16:10:09 IST	sw1.84	Tunnel 19 / Circuit 0	ALL_CIRCUITS(CIR_UTIL/min>=0)	0.00	%			Disabled	Disabled	

Figure 4: Brocade Network Advisor dashboard, FCIP violations pane.

This section lists the features and functionality that Brocade has incorporated into Brocade Extension. These features help preempt support issues altogether—or, when an issue cannot be totally preempted, these features help assure its quick resolution.

Brocade Network Advisor Dashboard

Every Storage Administrator approaches the task of managing and troubleshooting the environment with a particular style. Brocade Network Advisor has a customizable dashboard, as shown in Figure 3. The dashboard makes visible on one pane of glass the monitors, counters, and status indicators that interest you most. You can choose from over one hundred dashboard items, or if the item you need does not exist, you can create it. When you know exactly what is happening in your network, the goal of continuous uptime becomes a reality.

If a dashboard item is indicating an event or status, clicking that item will drill down into more specific information. In the example in Figure 4, a window showing FCIP Health Violations was opened to gather more detailed information. The detailed information shown here includes timestamp, device name, tunnel and circuit designation, the rule that was violated, the offending value that was measured, and the units of that measurement. You can also see if it was registered in the RASlog, Simple Network Management Protocol (SNMP), or e-mail, as well as the level of violation (marginal or critical), and other information.

CLI Dashboard

The CLI dashboard from Brocade is for users that prefer to use the CLI instead of Brocade Network Advisor. All information from each monitoring category can be obtained from the CLI, including any out-of-range conditions and the rules that were triggered. Error statistics provide at-a-glance views of status and the various contributing conditions to that status. Historical switch information for up to the last seven days provides a variety of error counters, giving users instant visibility into problem areas and facilitating the decision process towards proper resolution and planning.

Monitoring and Alerts Policy Suite (MAPS)

MAPS is an easy-to-configure and easy-to-use solution for policy-based threshold monitoring and alerting. MAPS proactively monitors the health and performance of infrastructure to ensure application uptime and availability. 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.

Which extension conditions are monitored and generate alerts? Brocade MAPS for extension monitors:

- Tunnel/Trunk State Change
- Tunnel/Trunk Overall Throughput
- Tunnel/Trunk PTQ Throughput
- Tunnel/Trunk PTQ DupAck (Duplicate Acknowledgements)
- Tunnel/Trunk PTQ Packet Loss
- Tunnel/Trunk PTQ Slow Starts
- Circuit State Change
- Circuit Utilization
- Circuit Packet Loss
- Circuit Round Trip Time (RTT)
- Circuit Jitter

Note: For the PerPriority-TCP-QoS (PTQ) monitors, there is a separate monitor for each priority: class F, High, Medium, and Low.

This set of monitors enables the detection of just about any degraded IP network condition that might occur. If the network goes down for a period longer than the KATOV (Keepalive Time Out Value) for a circuit, that event is detected, and an alert is processed. If the IP network has fits of transient congestion that results in either excessive jitter or packet loss, that event is detected and an alert is processed. If any one of the PTQ priorities suffers low throughput, packet loss, or out-of-order events, that issue is detected and an alert is processed. If the network reroutes, causing the RTT to change for the worse, that event is detected and an alert is processed. MAPS provides comprehensive detection of various network events and the ability to make those events known to Storage Administrators. Storage Administrators can leverage this information to enforce IP network SLAs.

Brocade tunnels are a managed transport between two Brocade Extension endpoints. There are many reasons for using a Brocade tunnel as a data transport across infrastructure, including granular load balancing, lossless failover/failback, higher availability, encryption, bandwidth pooling, protocol optimization, congestion management, Quality of Service (QoS) marking/enforcement, monitoring, reporting, and network diagnostics. A Brocade tunnel uses Extension Trunking and may consist of multiple member circuits, which may traverse one or more service providers, each with a distinct SLA. Brocade provides Storage Administrators with a single point of management and service provider SLA validation. Tunnels are point-to-point, and the endpoints of a tunnel are the same endpoints as the trunk.

There is a hierarchy of extension connectivity: Virtual E_Ports (VE_Ports, or tunnel endpoints) are in the top tier of the network and define the endpoints of tunnels. VE_Ports contain one or more circuits. Each circuit has four WAN Optimized TCP (WO-TCP) sessions, one for each priority (class F, High, Med, Low). MAPS is integrated into each of these tiers. Any indication of a problem in a lower tier poses a problem for its associated upper tiers. For example, if a problem exists with the IP network path that the medium PTQ WO-TCP session is using, the tunnel in which that circuit is a member will, in general, also have a problem. This alerts administrators when thresholds are exceeded.

MAPS for tunnels or trunks (VE_Ports) monitors and performs actions based on throughput and state change.

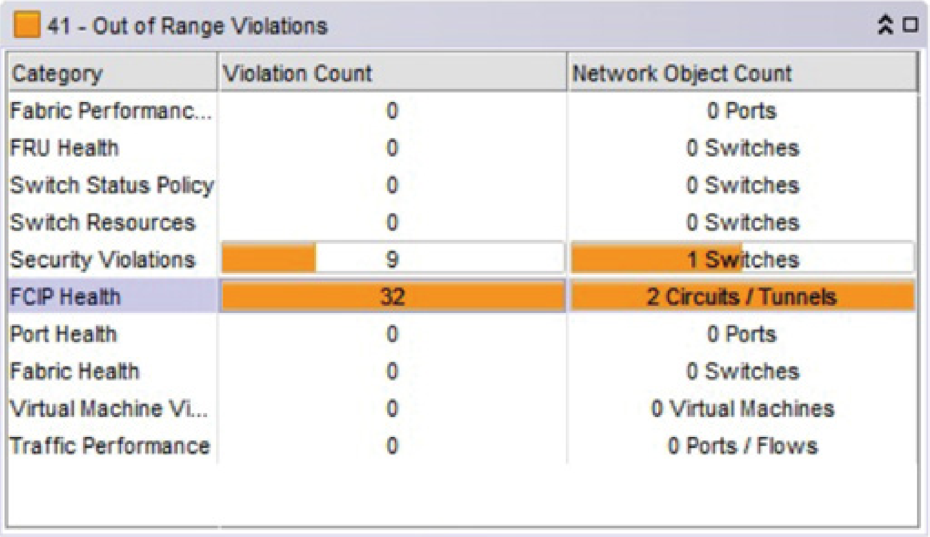

41 - Out of Range Violations

Category	Violation Count	Network Object Count
Fabric Performanc...	0	0 Ports
FRU Health	0	0 Switches
Switch Status Policy	0	0 Switches
Switch Resources	0	0 Switches
Security Violations	9	1 Switches
FCIP Health	32	2 Circuits / Tunnels
Port Health	0	0 Ports
Fabric Health	0	0 Switches
Virtual Machine Vi...	0	0 Virtual Machines
Traffic Performance	0	0 Ports / Flows

Figure 5: MAPS Summary Pane

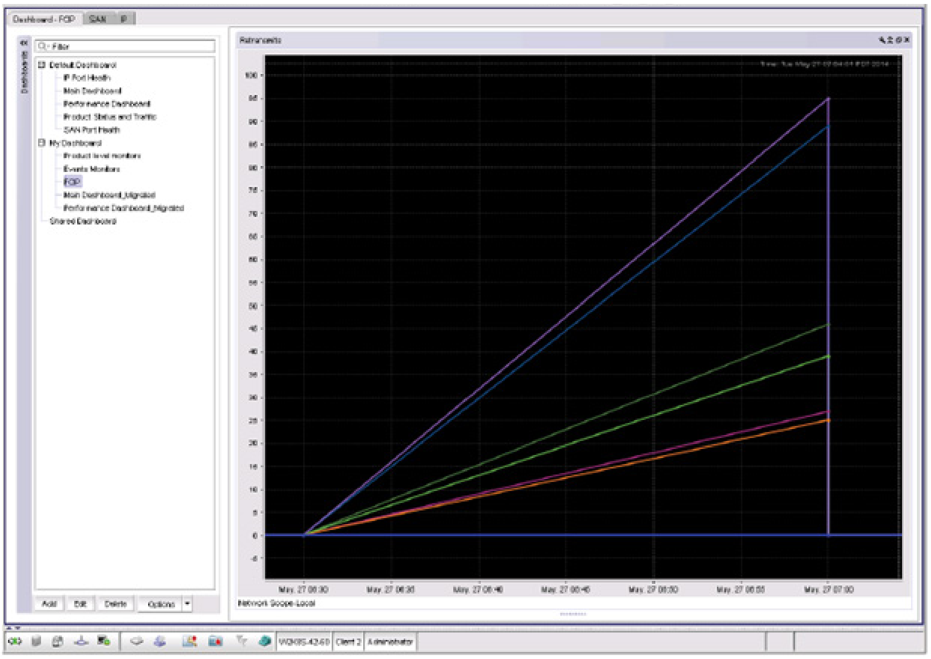

Figure 6: Increasing extension retransmits over half an hour.

Sudden Failures and Gradual Degradation Detection

MAPS can monitor both sudden failures and gradually deteriorating conditions. For example, MAPS can detect and alert users if a Cyclic Redundancy Check (CRC) error counter suddenly increases to five per minute or gradually increases to five per day. This is very useful for monitoring service provider SLAs. Service provider infrastructure is often immense, and it is difficult for service providers to operationally keep tabs on every optic, cable, and device in their network. Fabric Vision provides you with the tools you need to ensure your paths are within the promised SLA. Refer to Figure 6, which illustrates FCIP TCP retransmits.

Policy-Based Monitoring

Policy-Based Monitoring involves predefined monitoring groups with prevalidated monitoring policies. Multiple monitoring categories enable monitoring of the following: tunnels, circuits, VE_Ports, overall switch status, switch ports, small form-factor pluggables (SFPs), port blades, core blades, switch power supplies, fans, temperature sensors, security policy violations, fabric reconfigurations, scaling limits, CPU, memory utilization, traffic performance, and more.

Predefined monitoring policies are tiered to provide the best starting place for your particular environment and administrative style: aggressive, moderate, and conservative. A monitoring group's policy tier can be set individually and, if needed, you can fine-tune your specific environment. Individual items within a group can be set to the appropriate tier for customization. For example, when setting a style for FICON-associated extension connections you might select an aggressive style, and for Remote Data Replication (RDR) extension connections you might select a moderate style.

Custom Monitoring Groups

You can create custom monitoring groups, such as switch ports attached to high-priority applications and switch ports attached to low-priority applications. Monitoring of each group happens according to the group's unique rules, high-priority applications, and low-priority applications.

Flexible Rules

Rule flexibility means you can monitor a given counter for various threshold values, taking different actions when each value is crossed. For example, you can monitor a CRC error counter at a switch port and generate a RASlog entry when the error rate reaches two per minute, send an e-mail notification when the error rate is at five per minute, and fence a port when the error rate exceeds ten per minute.

MAPS: Actions and Alerts

MAPS generates various alerts and actions. Actions include Circuit Fencing, Port Fencing, Port Decommissioning, and setting status. Alerts include RASlog messages, SNMP traps, and e-mail notifications.

Circuit Fencing

Specific to the Brocade 7840 Extension Switch, the Brocade 7800 Extension Switch, the Brocade SX6 Extension Blade, and the Brocade FX8-24 Extension Blade, when certain thresholds are reached, a circuit can be taken offline. This is an important action for Brocade Extension, because a degraded circuit causes all circuits belonging to the same VE_Port to be degraded as well. This might result in a total throughput that is less than it would be without the degraded circuit. It is necessary to isolate a degraded circuit from the remaining clean circuits to maintain overall optimal throughput.

Because data is delivered to the Upper Layer Protocol (ULP) in order, a degraded circuit will cause all member circuits to degrade as well. If a trunk has two circuits, and one circuit is degraded such that it requires retransmits to complete successful transmission, ultimately the data sent on the clean circuit must wait for the retransmissions before delivering to the ULP. This means that both the clean and degraded circuits will go no faster than the degraded circuit. The degraded circuit may effectively be delivering only a small fraction of the bandwidth apportioned to it, depending on the degree of degradation. By fencing the degraded circuit, the clean circuit operates at the full bandwidth apportioned to it.

Port Fencing

The Port Fencing action takes the port offline immediately when user-defined thresholds are exceeded. Supported port types include VE_Ports as well as physical ports (E_Ports and F_Ports). This action is valid only for conditions evaluated by the actual port.

Port Decommissioning

Port Decommissioning acts in addition to Port Fencing. Port Decommissioning allows ports to be gracefully shut down. When certain monitored statistics cross-defined thresholds, ports are decommissioned, similar to Port Fencing but without the abrupt traffic disruption.

SFP Marginal

The SFP marginal action sets the state of the affected SFP transceiver in the MAPS dashboard to "down." This action does not bring the SFP transceiver down, rather, it affects only what is displayed in the dashboard. This action is valid only in the context of Advanced SFP groups.

RASLog Messages

Following an event, MAPS adds an entry to the internal event log for each switch involved. The RASlog stores detailed event information but does not actively send alerts.

SNMP Traps

In environments where you have a high number of messages coming from a variety of switches, you might want to receive them in a single location and view them using a Graphical User Interface (GUI). In this type of scenario, SNMP notifications may be the most efficient notification method. You can avoid logging in to each switch individually, as you need to do for error log notifications.

When specific events occur on a switch, SNMP generates a message (called a "trap") that notifies a management station using SNMP. Log entries can also trigger SNMP traps if the SNMP agent is configured. When the SNMP agent is configured to a specific error message level, error messages at that level trigger SNMP traps.

An SNMP trap forwards the following information to an SNMP management station:

- Name of the element whose counter registered the event
- Class, area, and index number of the threshold that the counter crossed
- Event type
- Value of the counter that exceeded the threshold
- State of the element that triggered the alarm
- Source of the trap

E-Mail Alert

An e-mail alert sends information about the event to one or more specified e-mail addresses. The e-mail alert specifies the device or devices and the threshold and describes the event, much like an error message.

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.

Flow Monitor

Flow Monitor enables you to monitor all the traffic passing through E_Ports, EX_Ports, F_Ports, and xISL_Ports by using hardware-supported flow parameters. Users gain comprehensive visibility into application flows within a fabric, including the ability to learn (discover) flows automatically. Define your own flows to monitor using combinations of ingress and egress ports, source and destination devices, Logical Unit Number (LUN), and frame types. The monitoring provided on the Brocade 7840 and the Brocade SX6 identifies resource contention or congestion that is impacting application performance.

Flow Monitor provides support for learning or manually defining the following types of extension flows:

- Top Talkers
- Flow within a fabric from a host to a target or LUN on a given port
- Flows inside Virtual Switch Logical Fabrics
- Flows passing through E_Ports (Inter-Switch Links, or ISLs)
- Flows passing through F_Ports (end device)
- Frame-based flows
- EX_Ports FC Routing (FCR) flows (routed)
  - Edge-to-Edge
  - Edge-to-Backbone
- Interfabric flows passing through backbone E_Ports
- Interfabric FCR flows passing through xISL_Ports
- NPIV flows from a host, so that you can monitor VM-to-LUN performance

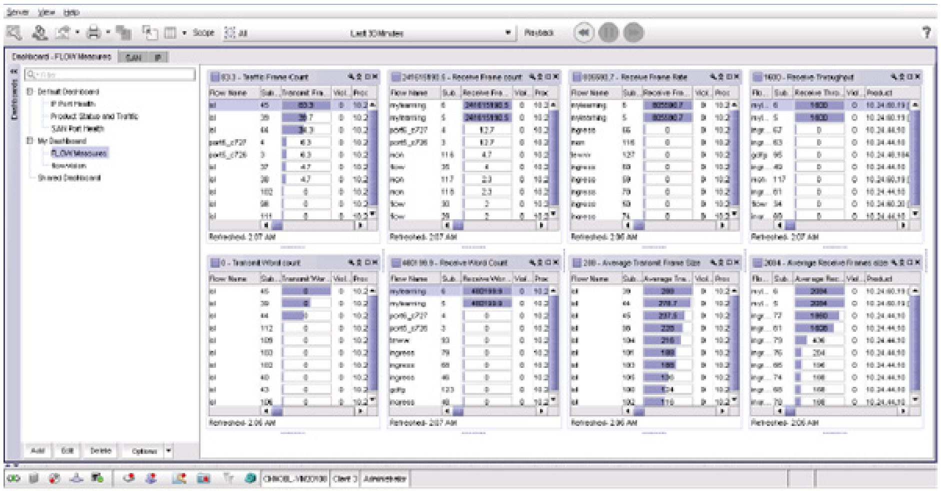

Figure 7: Flow Monitor measurements dashboard.

For specified flows, captured statistics provide insight into application performance. Monitoring various frame types at switch ports gives deeper insights into storage I/O patterns for a LUN, reservation conflicts, and I/O errors. SCSI read/write frame counts and SCSI read/write data statistics are supported on F_Ports when either the source or destination device is directly connected to the switch. Integration with the MAPS enables threshold-based monitoring and alerting.

Statistics include the following (refer to Figure 7):

- Transmitted and received frame counts
- Transmitted and received throughput rates
- SCSI read and write frame counts
- SCSI reads and writes per second (IOPS)
- Monitored frame types:
  - SCSI Aborts
  - SCSI Read
  - SCSI Write
  - SCSI Reserve
  - Rejected Frames
  - Many others

Below is the dashboard for flow measurements from Flow Monitor. Multiple flows are monitored simultaneously, and their various characteristics are displayed.

The flow monitor shown in Figure 8 charts various flows simultaneously. Note that the number of violations (left side) and the respective scales (right side) per the data are displayed.

MAPS for Flow Monitor

MAPS monitors flows that are established within Flow Vision and generates alerts based on user-defined rules. This enables users to monitor and be alerted when established thresholds are exceeded for a particular flow. You can use MAPS for Flow Monitor to identify slow-drain devices so that you can deal with them appropriately.

Flow Generator

Flow Generator is a traffic generator for pretesting and validating infrastructure. Flow Generator is hardware-specific (Condor 3 and later) and can be used across the Brocade Extension infrastructure. Various components and devices can be tested, including internal switch connections, cable plant, new switches and blades, servers, storage, and IP networks.

Flow Generator allows users to:

- Configure a FC/FICON-capable port as a simulated device that can transmit frames at 16 Gbps line rate
- Emulate a FC/FICON fabric without actually having any hosts, targets, or testers
- Pretest the entire SAN

Flow Generator creates a special port type called SIM ports, which are used to process simulated traffic. SIM ports behave like normal E_Ports, EX_Port, or F_Ports but are used only for testing purposes. SIM ports originate and terminate Flow Generator traffic and ensure that this traffic does not leave the destination switch. Flow Generator generates standard or custom frames with user-specified sizes and patterns. Flow Generator supports predefined or learned flows to generate traffic between configured SIM ports. Once you activate a flow, the flow remains active until it is deactivated. The flow lies in wait when the Source ID (SID)/ Destination ID (DID) SIM ports are offline. As soon as SID/DID SIM ports are brought online, the traffic proceeds. As an example use case, create a traffic flow between a SID and DID to validate routing and throughput across an FCIP ISL.

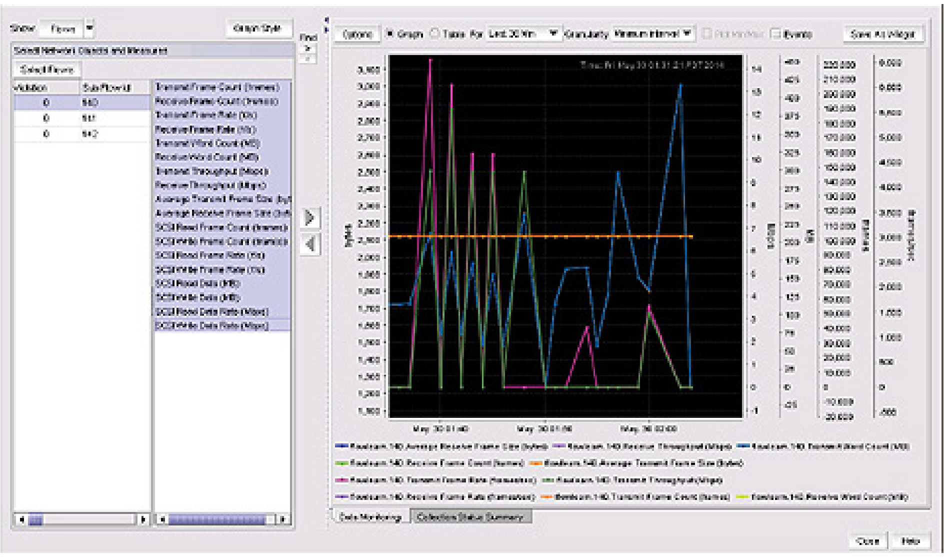
Show
Flows
Graph Style
Find
Select Network Objects and Measures
Select Flows
Violation
Sub Flow Id
0
0
0
Transmit Frame Count (frames)
Receive Frame Count (frames)
Transmit Frame Rate (f/s)
Receive Frame Rate (f/s)
Transmit Word Count (MB)
Receive Word Count (MB)
Transmit Throughput (Mbps)
Receive Throughput (Mbps)
Average Transmit Frame Size (byt
Average Receive Frame Size (byt
SCSI Read Frame Count (frames)
SCSI Write Frame Count (frames)
SCSI Read Frame Rate (f/s)
SCSI Write Frame Rate (f/s)
SCSI Read Data (MB)
SCSI Write Data (MB)
SCSI Read Data Rate (Mbps)
SCSI Write Data Rate (Mbps)
Options
● Graph
○ Table
Granularity
□ Events
Save As Widget
bytes
May 30 01:40
May 30 01:50
May 30 02:00
Data Monitoring
Collection Status Summary
Close
Help

Figure 8: Monitoring flows with Flow Monitor.

WAN Optimized TCP

WAN Optimized TCP (WO-TCP) from Brocade is a high-powered and aggressive TCP stack used for the fast and efficient transport of large data sets across enterprise WAN IP infrastructures. WO-TCP offers many advanced features, but specific to this tech brief, WO-TCP provides an exclusive streams-based mechanism. Multiple data streams (flows) can be transported autonomously without the need for separate TCP stacks for each individual stream. Because each stream is autonomous, there is no Head of Line Blocking (HoLB) in the event that one stream becomes slow and TCP windowing performs flow control. All other flows can continue to run at rate, even when one flow is running well below rate. Additionally, when data needs to be sent without a delivery guarantee—for example, during tests and other special conditions—WO-TCP can designate streams as "nonguaranteed," and the behavior becomes identical to User Datagram Protocol (UDP). This special mode of WO-TCP is essential for testing IP networks while maintaining the same TCP headers used by Brocade Extension. This ensures that IP network testing sees the same extension traffic that it normally would see. Wtool uses WO-TCP.

Wtool

Wtool is an IP-specific tool for testing the WAN-side infrastructure on the Brocade 7840 Extension Switch and the Brocade SX6 Extension Blade. 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.

WO-TCP is an advanced TCP stack with unique abilities. WO-TCP can logically behave just like UDP. Testing an IP network for proper SLA requires a UDP-like transport to provide constant bit rates, no reordering, and no retransmits. If the IP network reorders, you will see it. If the IP network drops packets, you will see it. Traditional TCP hides these issues, making any diagnostic tool useless. Wtool uses the same TCP headers as the tunnel, except that all retransmit mechanisms are disabled, which exposes aberrant network conditions. To test a specific protocol such as FCIP, you need that protocol's TCP headers to traverse the network and pass through security devices that may exist. Additionally, when used with Wtool, WO-TCP prevents traffic windowing due to network error conditions. This means that Wtool always drives traffic at circuit rates despite network errors that would normally cause traditional TCP to close its windows. Brocade provides WO-TCP technology in the world's leading data transport TCP stack.

- Wtool sees all packet drops.
- Wtool sees all packet reordering.
- Wtool drives test traffic at rate.

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 be 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.

Bottleneck Detection

Brocade offers bottleneck detection for FC and FICON fabrics.

ClearLink Diagnostics

ClearLink Diagnostics, sometimes referred to as D_Port, is a Brocade Gen5 Fibre Channel exclusive feature.

FEC and BBC Recovery

Brocade provides Forward Error Correction (FEC) and Buffer-to-Buffer Credit (BBC) recovery, which are important technologies for high-speed Fibre Channel communications. Both are used on the FC/FICON side of the Brocade 7840 and the Brocade SX6.

For more information about these features, refer to: http://www.brocade.com/solutions-technology/technology/san-fabric-technology/fabric-vision.page

Fabric Vision Extension Use Cases

Below are a couple of use case examples. This is not an exhaustive list.

Determine Trouble Location

Isolate IP network issues from storage network or storage device issues. Figure 9 shows a routed architecture. Routed architectures usually involve edge fabrics, which is why they are routed. FCR protects the edge fabrics from IP network and WAN anomalies. Depending on the edge fabric design and deployment, there is a greater possibility that HoLB may occur as a normal part of flow control. This example is a worst-case scenario showing a more complex network, however, this example pertains to simpler architectures as well, such as architectures that do not involve edge fabrics or FCR. The concept is to monitor flows traversing the trunk over the IP network. Users want to monitor interfabric flows through the backbone to make sure that traffic is not

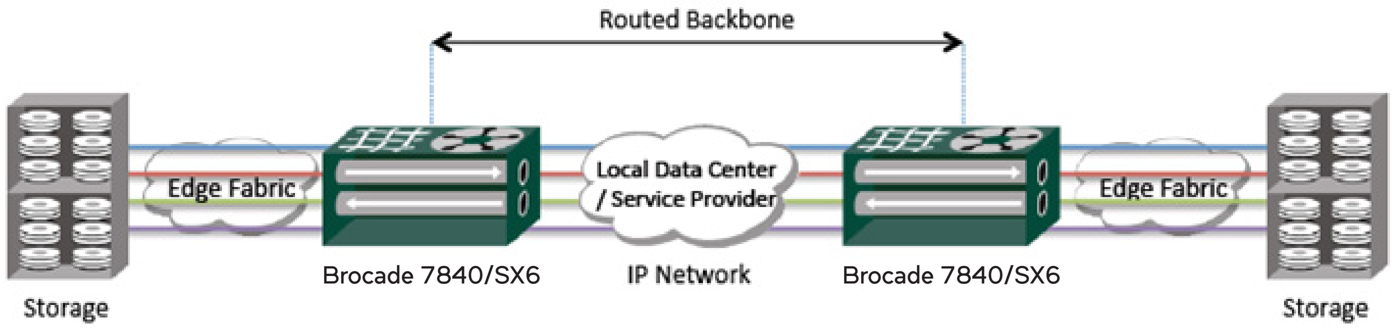

Figure 9: Routed extension architecture.

experiencing less than ideal conditions. MAPS monitors and alerts about WAN conditions using FCIP group policy based on comprehensive multilayer metrics, and VE_Ports, tunnels, circuits, and QoS are monitored for irregularities. Flow Vision monitors individual replication session data and I/O rates, plus RTT and jitter (latency variance).

Advance warning of replication under stress and effective troubleshooting is critical, given the short time and the impact to operations.

Is the IP network providing its SLA? Users receive an alert from MAPS in the event that one of the following occurs:

- Circuit Retransmits: Retransmits grow above a certain rate, indicating congestion or packet loss or that the network is excessively delivering data out-of-order.
- Circuit State Change: A circuit goes down because the KATOV has expired some number of times within a period.
- Circuit RTT: Round Trip Time has increased well beyond the typical operating value.
- Circuit Jitter: Jitter has increased well beyond the typical operating value.

Any of these is an indicator of a problem in the IP network that will negatively affect the performance and impede the goal of safe RDR. The information provided by the Brocade 7840 and the Brocade SX6 can be used to enforce SLAs between Storage Administration and Network Administration and/or service providers.

Determine Maximum Stress

The potential drop of a replication session due to exceeded threshold indicates maximum "stress" during the monitoring interval. If the delta set cycle time is 5 or 10 seconds (Figure 10 shows the cycle time set to 30 seconds), the expectation is to complete the transmission of that delta set across the extension network within that period of time. If that is not happening reliably, there are a few paths to resolution. First, is there enough bandwidth for the amount of data that needs to be sent? Use Flow Vision to monitor the replication flows and determine if they are fully utilizing the bandwidth that is available. If they are, then there is not enough bandwidth. If they are not, the next step is to determine if the IP network is providing its SLA (refer to the previous case). If the bandwidth is not being fully consumed, and the IP network is clean, investigate the storage array for possible issues.

In Figure 10, the delta set cycle time is set to 30 seconds. Cycle times will not be less than 30 seconds. The expectation is that complete data transfer of the delta set will take place within the 30-second interval. As shown, this is not happening regularly. In many instances the amount of time to transfer the data set exceeds 30 seconds and may take as long as 43 seconds. A determination must be made if this is due to a network problem or if there is just too much data to transmit for the bandwidth and compression available.

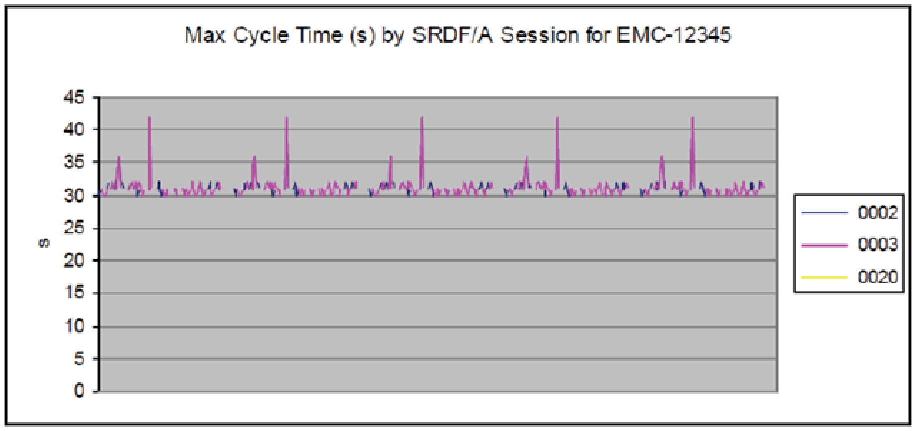

Figure 10: Data delta set cycle time to completion.

Does RDR appear to be stressed, while the IP network administrators claim it should not be? Use these tools to gain visibility into your RDR environment:

- Use Flow Vision to visualize the RDR flows and determine if they are consuming all the available bandwidth.
- Use Flow Vision to monitor individual replication session data and I/O rates.
- Use Flow Vision to monitor RTT and jitter. If this is excessive, it can cause droop across the WAN connection. Protocol optimization such as Brocade FastWrite may be required, if applicable.
- Use MAPS to alert on events that indicate degraded IP network states, which elongate delta set transmission times and hinder recovery from Delta Set Extension or backlog journaling when the network is running behind.

Configuration Validation and IP Network Assessment

Setting up an RDR network is a daunting task. The technology is sophisticated and uses a significant number of moving parts. At fruition, when the user is at the validation point in the setup process, it can be difficult without the proper tools to make sure everything is functioning properly. This applies either when the project is performed in-house or when it is turned over from Brocade (or the OEM or reseller) Professional Services. Once everything is installed and configured, it is time for validation testing. You can do end-to-end testing of the storage array, the storage network, the extension platform, and the IP network.

The Brocade 7840 and the Brocade SX6 provide tools to simplify this process considerably. These tools include ping and traceroute at a rudimentary level. There is a PMTU discovery tool for determining the IP network MTU in cases where it is not already known or has to be verified. For a comprehensive test of the IP network, run Wtool for any desired period of time. For example, run Wtool for the duration of a typical work day. Generate line-rate extension TCP flows and gather IP network statistics for analysis and SLA evaluation. Upon completion, Wtool will provide information about the IP network, including maximum throughput, congestion, loss percentage, out-of-order datagrams, latency (RTT), and other network conditions.

Summary

Storage Administrators are facing challenges with infrastructure they do not manage or control, specifically the IP network, across which many of their applications pass. Because Storage Administrators have little familiarity with the IP network and vice-versa, this can prolong time to resolution when there are operational problems. This is exacerbated when support and multiple vendors are involved. The process requires a fast and efficient way for Storage Administrators to pinpoint where a problem may be occurring.

But how can you more quickly pinpoint problems? Storage arrays have no visibility into the extension network, plus that is not really within the expertise of Storage Administrators. The Brocade 7840 and the Brocade SX6 extension platforms provide unique tools for quickly resolving hard-to-diagnose issues. Fabric Vision offers MAPS, Flow Vision, Flow Generator, and Wtool, which are all enterprise-class tools. MAPS has the ability to detect either sudden changes in the network or errors that present themselves slowly over time. Monitoring is easy to set up and customize with pre-established policies and rules that incorporate thresholds determined by Brocade over more than a decade of experience. When a monitored network characteristic crosses marginal or critical thresholds, various alerts and actions are available. Beyond the automated monitoring, alerts, and actions provided by MAPS, another valuable tool for visualizing flows is available, called Flow Vision.

Flow Vision offers the Flow Monitor tool. Flow Monitor can discover flows automatically or can be configured manually. Upon monitoring a particular flow, MAPS is integrated into Flow Monitor so that thresholds of interest can be set with corresponding alerts and actions. Another tool that Flow Vision provides is called Flow Generator. Flow Generator can generate test flows originating at the FC ASIC level across and traverse Brocade Extension. The characteristics of the flows are either Brocade preconfigured or learned from the application.

WO-TCP is an advanced TCP stack with the ability to treat each individual flow autonomously and eliminate HoLB. WO-TCP also provides a special functionality for generating UDP-like traffic with extension headers. This permits Wtool to perform accurate testing of network conditions while the IP network sees actual extension traffic flows. A couple of cases were discussed in this brief: One was how to determine where trouble may be originating. Another was how to determine stress on array replication applications. Overall, the Brocade 7840 and the Brocade SX6 come with sophisticated tools that 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.

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

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.