Dell EqualLogic Best Practices Series

Creating a DCB Compliant EqualLogic iSCSI SAN with Mixed Traffic

A Dell Technical Whitepaper

Storage Infrastructure and Solutions Engineering

Dell Product Group

August, 2011

Table of Contents

Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions team in August 2011 at the Dell Labs facility in Austin, Texas.

The team that created this whitepaper:

Jerry Daugherty, **Tony Ansley**, and **Camille Daily**

We would like to thank the following Dell team members for providing significant support during development and review:

Chris Almond, **Mike Kosacek**, **Jeremy Hitt**, **Kirt Gillum**, and **Daryl Proctor**

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information. You can submit feedback as follows:

Use the "Start a new thread" link here:

http://www.delltechcenter.com/page/Creating+a+DCB+Compliant+EqualLogic+iSCSI+SAN+with+Mixed+Traffic

1 Introduction

As Dell's EqualLogic PS Series storage matures, additional features continue to add value to the product. EqualLogic controller firmware version 5.1 (and later) supports a new set of networking standards called Data Center Bridging (DCB). DCB unifies the communications infrastructure for the data center, supporting design of a single, shared networking infrastructure that meets the communications needs for all IT operations. This includes those traditionally associated with the following:

Local Area Network: (LAN) End user's access to important application servers such as email, database, and Intra/Internet web servers

Management Network: Required to centrally manage data center resources without impacting other networking requirements.

Storage Area Network: (SAN) Used to connect application servers to a shared storage solution that may use Fibre Channel/Fibre Channel over Ethernet (FC/FCoE) or iSCSI

DCB is a set of new IEEE standards that help ensure performance and delivery. These standards are:

Priority-based Flow Control: (PFC; IEEE 802.1Qbb) Expands the function of the standard class of service structure of Ethernet to provide a mechanism to allow for lossless classes of service since a non-lossless class cannot be paused.

Enhanced Transmission Selection: (ETS; IEEE 802.1Qaz) Provides administrators with the ability to group multiple classes of service together and then define a guaranteed minimum bandwidth allocation from the shared network connection.

Datacenter Bridging Capability Exchange: (DCBx) The glue that binds all of the standards by allowing networking components to understand the settings required to operate within the DCB environment.

Congestion Notification: (CN; IEEE 802.1Qau) Enables DCB switches to identify primary bottlenecks and take preventative action to ensure that these primary points of congestion do not spread to other parts of the network infrastructure.

1.1 Whitepaper purpose

The purpose of this paper is twofold: to provide test-based results showing the benefits of DCB in both dedicated iSCSI environments as well as converged networks (where multiple networks are migrated into one DCB network); and to provide best practice information to help network and storage administrators deploy DCB in new implementations in conjunction with their Dell™ EqualLogic™ PS Series SAN deployment.

2 DCB terminology

2.1 Priority-based Flow Control

Priority-based Flow control (PFC) is an evolution of the concept of Flow Control originally implemented in the MAC Pause feature of Ethernet (IEEE 802.3x). That feature was a simplistic control of traffic made by requesting that the sender stop transmitting for a specified period of time. With no granularity applied to this request; all Ethernet frames were stopped. One unfortunate consequence of this standard was that manufacturers were not required to implement it. If a receiving station requested that traffic be paused, and the sender did not implement pause, then traffic would continue to be transmitted. The noncompliance resulted in the receiver dropping incoming frames. This issue can also arise during what is referred to as "Head-of-Line Blocking", where multiple ingress ports traffic is squeezed into a single egress port.

PFC also asks the sender to stop transmitting, but it in addition leverages the idea of classes of traffic to apply granularity to the process. In a DCB environment, all traffic is tagged with a Class of Service (CoS) using the VLAN "Q-tag". PFC can then request that a specific CoS be paused for a time, while other classes can continue unhindered as shown in Figure 1. In a storage environment, this may mean that "normal" TCP/IP traffic is dropped, while storage traffic tagged with a higher priority can be managed using PFC. As a benefit of PFC, once its use has been negotiated both receiver and sender must adhere to it.

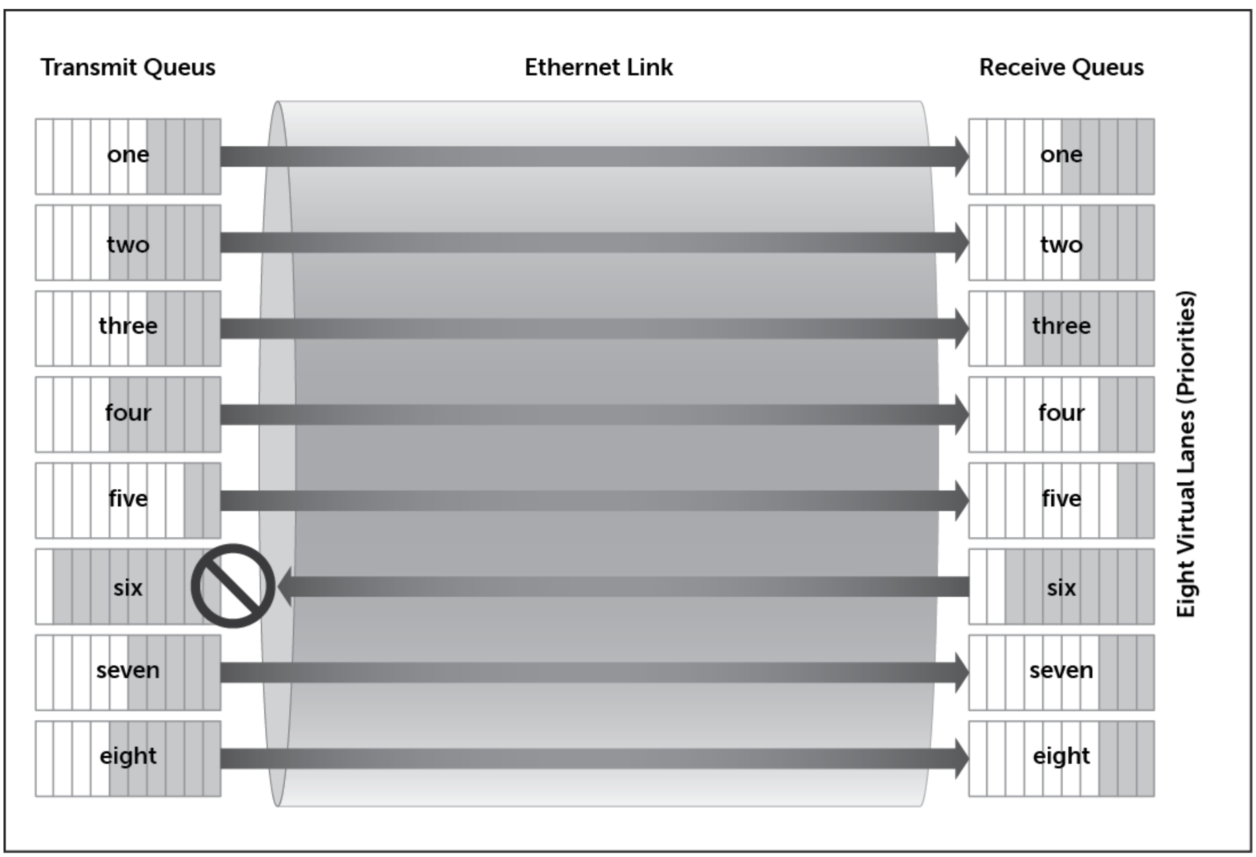

Figure 1 Example PFC Diagram

2.2 Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) is a mechanism for guaranteeing a percentage of bandwidth to a traffic class. A traffic class contains one or more Classes of Service from the VLAN Q-tag. Each traffic class is then assigned a percentage of bandwidth with a granularity of 1%. All traffic class bandwidths must add up to 100%; no oversubscription is allowed.

The bandwidth percentage defined is a minimum guaranteed bandwidth for that traffic class. If a traffic class is not using its entire minimum amount, it can be utilized by other traffic classes that may need it. However, as soon as the original traffic class requires its bandwidth again, the other traffic flow must be throttled to allow the bandwidth to be recovered. This is accomplished through the use of PFC discussed earlier. PFC will issue a pause for the required traffic classes in a manner to allow the bandwidth to be regained with a minimum of dropped frames for the throttled traffic class. An important note to consider when deciding on bandwidth percentages for each traffic class is that the required accuracy of the ETS algorithm is only plus or minus 10%. If the setting for a particular traffic class is not set with this in mind, it may not receive all the bandwidth expected.

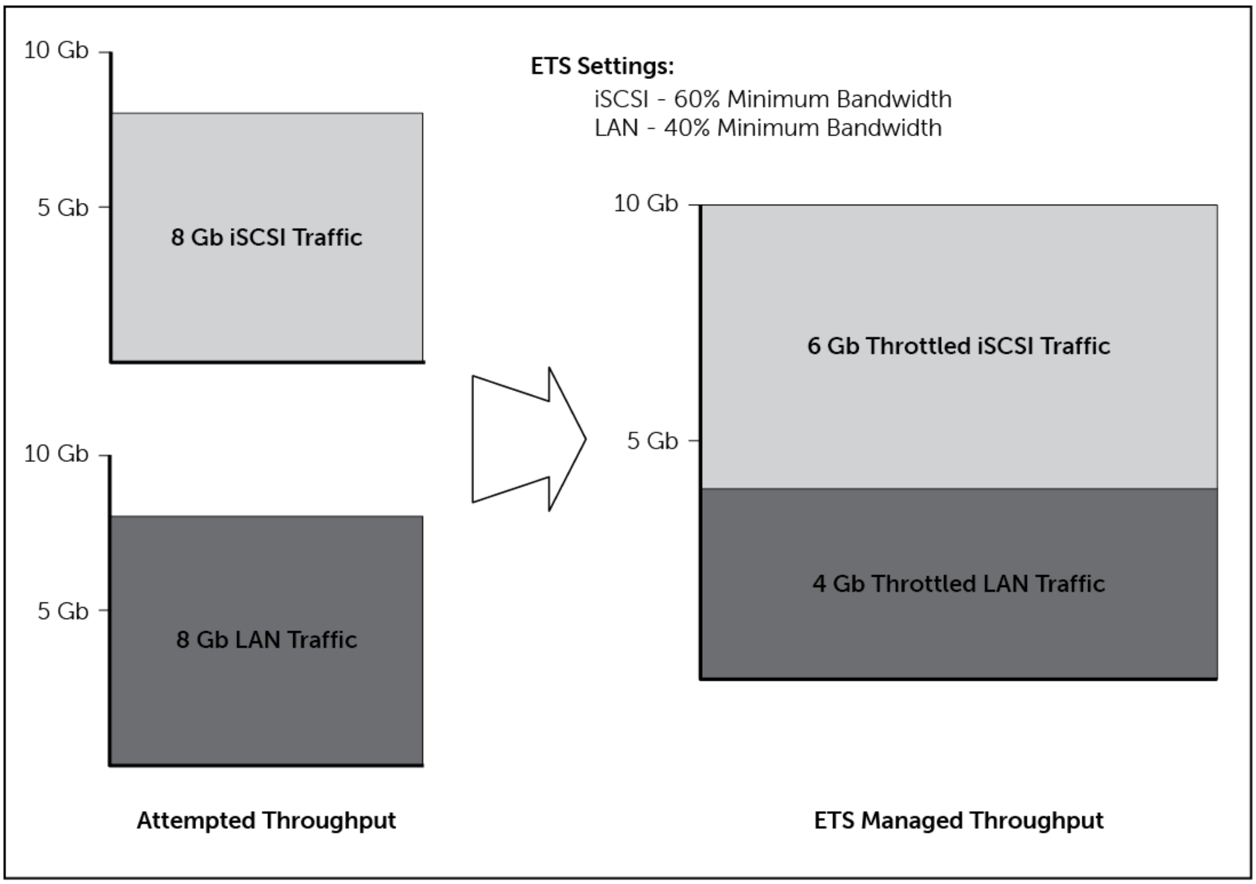

Figure 2 Example ETS diagram

2.3 Congestion Notification

Congestion Notification (CN) is a mechanism for managing congestion throughout a DCB fabric or domain. Ideally, That fabric would consist of all interconnected switches and end-devices that conform to the same settings for PFC, ETS and CN. Frames in a fabric that are conforming to CN will be "tagged" with a Flow Identifier. CN then relays messages between two types of devices called Congestion Points (CPs) and Reaction Points (RP) to control the flows. CPs are generally switches that have the capability to determine that they are experiencing congestion. Once detected, a CP then sends a CN message to the originating RP. When an RP receives the CN message, it begins to throttle the output for the designated flow until the CN messages stop.

This mechanism is a way of moving the congestion from the core of the network to the edge. CN is generally more effective for longer lived traffic flows, as opposed to small bursts of traffic. CN can work hand-in-hand with PFC to control congestion and overruns throughout the DCB fabric.

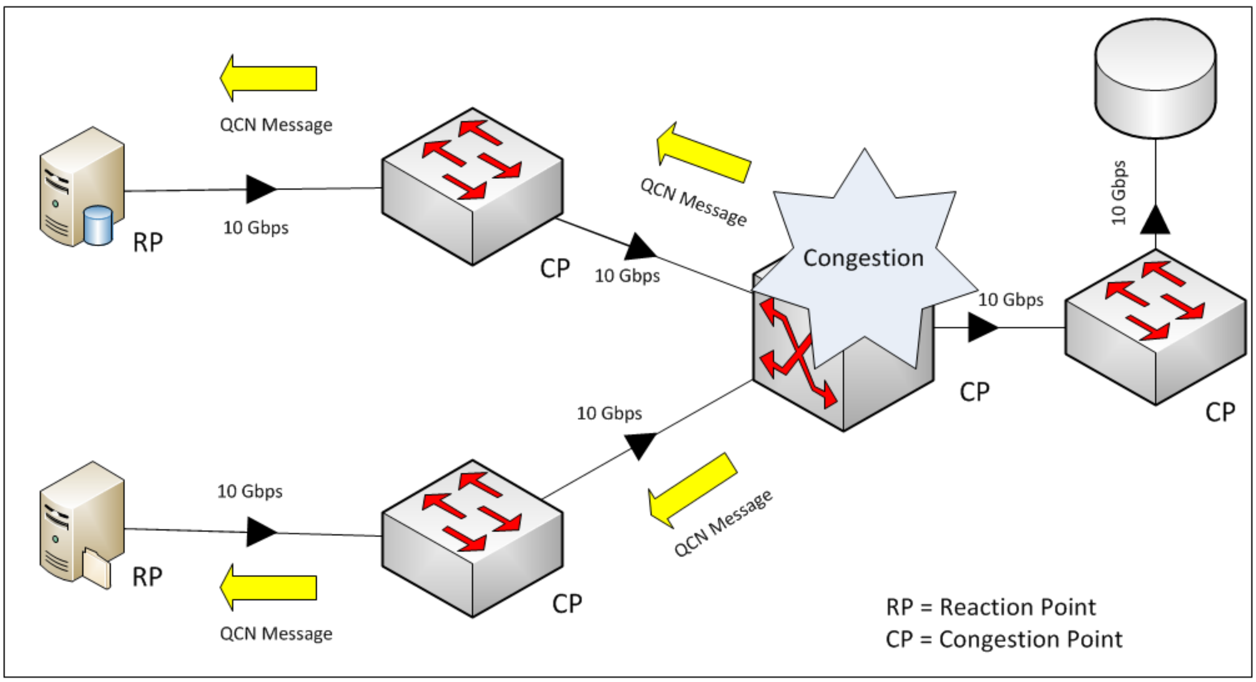

Figure 3 Example topology for congestion notification

2.4 Data Center Bridging Capability Exchange

Datacenter Bridging Capability Exchange (DCBx) is an extension of the IEEE standard 802.1AB for Link Layer Discovery Protocol (LLDP). It uses the existing LLDP framework for network devices to advertise their identity and capabilities. LLDP relies on the use of Type-Length-Values (TLV) to advertise the device capabilities for a multitude of Ethernet functions, as well as its identity. DCBx defines new TLVs specific to the DCB functionalities.

PFC and ETS have specific TLVs defining items such as:

- Whether PFC is to be used
- Priorities that will be managed using PFC
- Which priorities belong to a specific traffic class
- Bandwidth minimums for each defined traffic class

The standard also defines "Application" TLVs. These TLVs allow for the definition of which protocols will be managed, as well as the priorities to be assigned to each. Currently FCoE and iSCSI have Application TLVs; NAS will have a future TLV that will allow it to be DCB aware.

For EqualLogic environments using DCB, support for the iSCSI TLV is required. For FCoE, it is easy for the network to assign a higher priority to its frames, as they have a separate EtherType. For iSCSI however, the end station needs to know how to identify iSCSI frames from the other TCP/IP traffic. The iSCSI TLV identifies what TCP port iSCSI traffic is using, so that the end station can properly assign the class of service to it. Once this has been outlined, PFC and ETS can manage the iSCSI traffic as its own class.

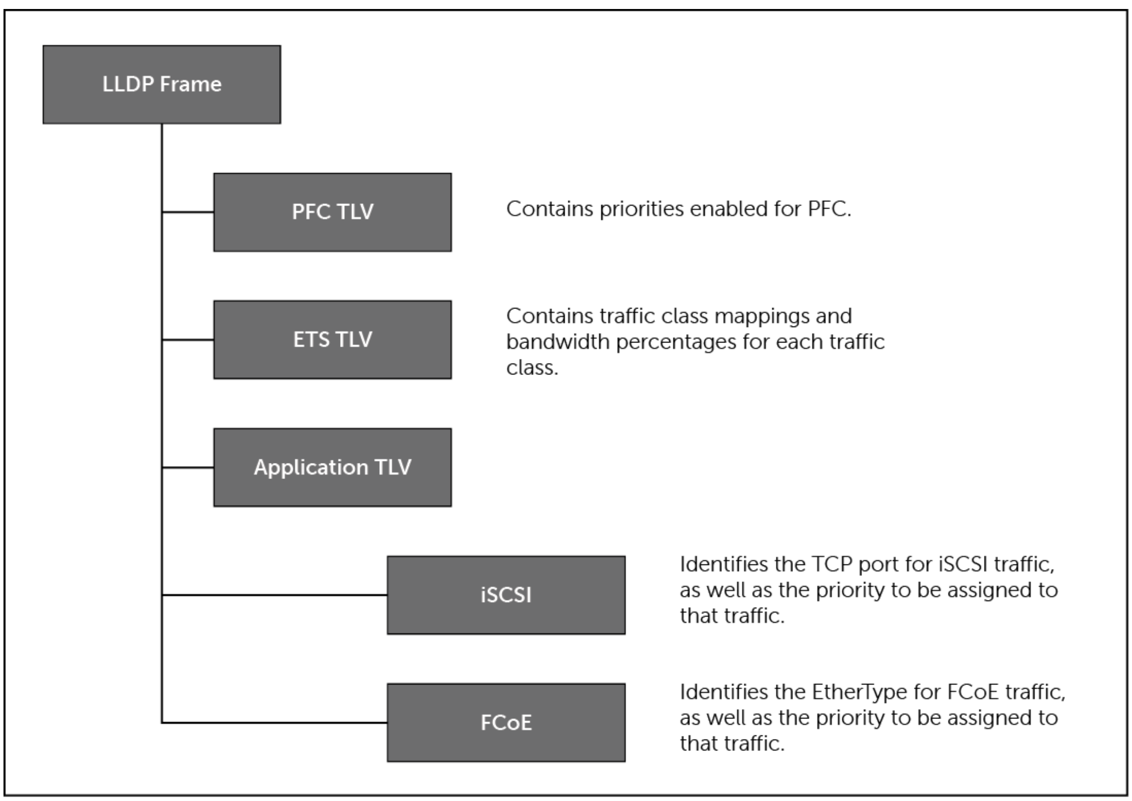

Figure 4 LLPD hierarchy supporting DCBx

3 Benefits of DCB

3.1 Better utilization of bandwidth

In traditional Local Area Network (LAN) environments, fair use of the available bandwidth is not generally an issue. TCP/IP and applications are written to manage the varying latencies and available bandwidths to make issues less impactful to the user. For storage applications however, there is an expectation of a certain amount of bandwidth. The ability to guarantee each application or traffic type a minimum amount of networking resources is an essential part of the ETS features within the DCB environment.

Through the use of ETS, networked devices such as hosts or storage targets can be guaranteed a minimum percentage of bandwidth, while at the same time the ability to access the full bandwidth when it is not in use by other applications. PFC manages the multiple flows of network data to ensure frames are not dropped for lossless priorities.

3.2 Heavy virtualization

As the move to virtualize grows, more and more workloads are being moved onto fewer servers. As the number of virtual machines per host grows, the demands placed on the shared network connections grow. When this is multiplied with the need to run separate and possibly disparate networks for storage and LAN traffic, the result is either poor performance for all applications or a plethora of network ports and cables in each host to support the various network requirements of the virtualized environment.

The potential to merge network flows into fewer 10 Gb connections, while still maintaining the needed performance characteristics, is compelling. Through the use of DCB and converging flows onto fewer connections, administrators can more easily maintain the datacenter and manage the networking resources required in a virtualized environment.

3.3 Power and cooling

The benefits realized from reducing the number of cables and network cards virtualization extends to the power and cooling needs of the datacenter. As the traffic flows converge onto one network instead of several, the resulting number of network switches goes down. Along with the switch count, the power and cooling requirements for the datacenter also fall.

The airflow of each rack is also immediately impacted through the use of fewer cables. One or two 10 Gb fiber connections per server takes up significantly less volume than four or more 1 Gb copper cables per server. With more cables multiplied across the servers in a rack, airflow becomes severely inhibited.

3.4 Cost

As the cost of 10 Gb Converged Network Adapters (CNAs) continues to fall, the economic benefits of converging multiple traffic flows onto 10 Gb will continue to grow. With the same ideas from the areas above, it will be far more cost-effective to purchase fewer ports of 10 Gb DCB Ethernet than to purchase many ports of 1 Gb non-DCB Ethernet.

4 DCB testing

The following workloads and tests were developed to provide empirical evidence of the benefits of DCB in two common scenarios: Dedicated iSCSI networks and Converged traffic networks. In the Dedicated iSCSI tests, traffic was only running from the iSCSI host to the target. No other traffic was being injected. This simulates environments where the customer is implementing DCB, but continuing to separate storage and LAN traffic through the use of dedicated, separate physical networks. For the Converged Traffic tests, iSCSI traffic was still run, and non-iSCSI TCP traffic was also streamed from a TCP source host to the same ports on the iSCSI host. This allowed the impacts of DCB to be directly measured on the initiator links. This scenario provides a high throughput test of both storage and LAN traffic on the same network segment.

4.1 Workload definitions

During each test we gathered throughput (MB/s) and IOs per second (IOPS) data at the host initiator. We also monitored TCP retransmission rates to ensure that it never went above 0.5% during the test runs. We used Medusa Labs Test Tools Suite (MLTT) from JDSU Corporation for I/O generation. MLTT was selected because of its ability to run multiple initiators simultaneously while rolling up the results into a single performance metric, as well as the granularity and detail of the metrics it provides. MLTT was also the only tool that could simultaneously run a TCP workload over the same links from a single, consistent interface.

We tested three IO workload types: sequential read, sequential write, and a random read/write mix. These workloads were selected to imitate three common real-world scenarios:

- **Seq Read**: Designed to imitate large-block, high throughput scenarios such as video streaming or backups
- **Seq Write**: Designed to imitate medium block Seq Write scenarios such as File transfers or Day-to-day OS operations
- **Random R/W**: Designed to imitate small block, high IOPs scenarios such as databases and OLTP transactions

Each test consisted of a single workload type running for 15 minutes. We completed three separate test runs for each workload type, and then computed an average of the results of all three runs. Details on the workload types are shown in Table 1.

Table 1 Workloads

Workload Type	Block Size	Read/Write Ratio
Sequential Read	256 K	100%/0%
Sequential Write	64 K	0%/100%
Random Read/Write	8 K	67%/33%

After comparison runs early in the testing process, it was determined that the sequential write and random read/write workloads did not provide sufficient throughput to cause a noticeable effect. As can be seen from the Table 2, virtually no difference was seen from a DCB versus non-DCB standpoint.

Table 2 DCB and non-DCB Comparison

	Seq Write-No DCB	Seq Write-DCB	Random RW-No DCB	Random RW-DCB
IOPS	29145	29992	78734	78549
Throughput (MB/s)	1821	1874	615	614

Note: For the remainder of this paper, only the results derived from the sequential read workload will be discussed.

4.2 Topologies

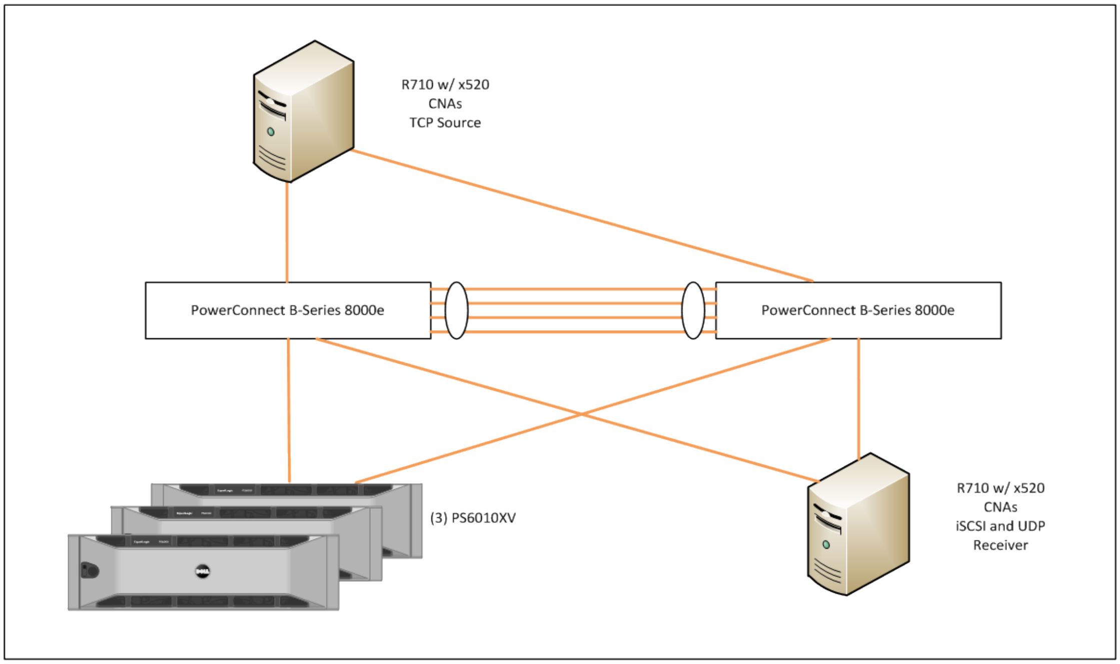

Figure 5 Test topology

In this topology, three EqualLogic arrays were used to ensure that the throughput for the host ports could be maximized. All physical connections were made in a redundant manner, with ports connected across separate switches. These switches were then connected through a link aggregation group of four ports, providing sufficient bandwidth for any traffic traversing the inter-switch link. A separate server was used to provide TCP traffic to the test host, ensuring that the converged traffic would occur on the proper links being measured for the test.

For a complete listing of the components used, and the firmware and driver version, please refer to the appendices.

5 Results and analysis

5.1 DCB versus non-DCB in a dedicated SAN network

For pure iSCSI environments, tests were run to determine if there was any benefit to running DCB protocols. This may be for new datacenter environments where the customer wishes to implement DCB-capable infrastructure while preserving a separation of storage and LAN traffic. It may also present an opportunity for the customer to upgrade existing infrastructure without changing the logical layout of the network.

The most compelling point for the movement of a pure iSCSI environment to DCB is in TCP retransmissions. Under workloads that can generate line rate for 10 Gb connections, TCP retransmissions are always a battle. The graph below illustrates that with the throughput close to line rate on both 10 Gb ports for an initiator TCP retransmissions are measurable. This level of retransmissions would not likely cause noticeable degradation in performance on the network, but could be a harbinger of future problems as the network scales. The drop of retransmissions to zero is indicative of the improvement in response between traditional Ethernet pause and the new PFC in DCB.

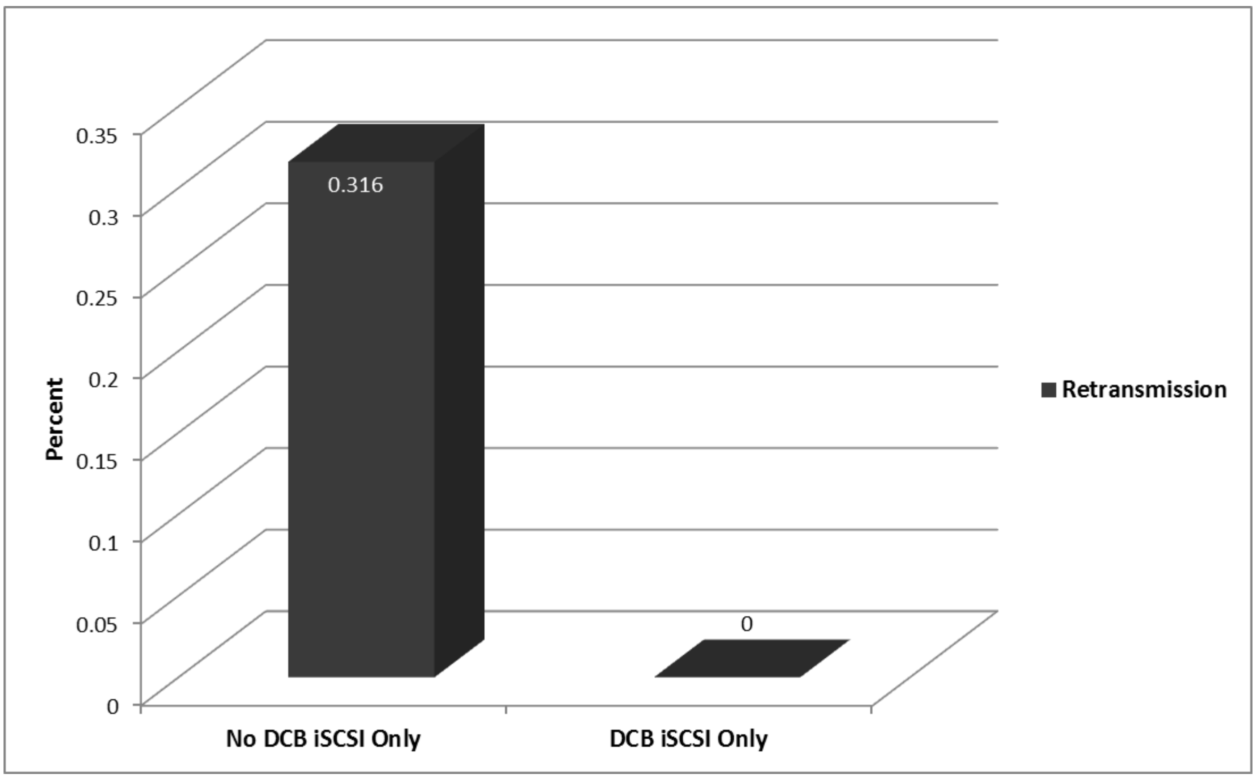

Figure 6 TCP Retransmission Comparison for DCB vs. non-DCB for iSCSI only Traffic

For throughput there is an approximate 5% improvement for DCB versus non-DCB. This increase is likely due to the reduction of TCP retransmission noted in the previous section. In both cases, near line-rate performance was achieved from both ports on the initiator.

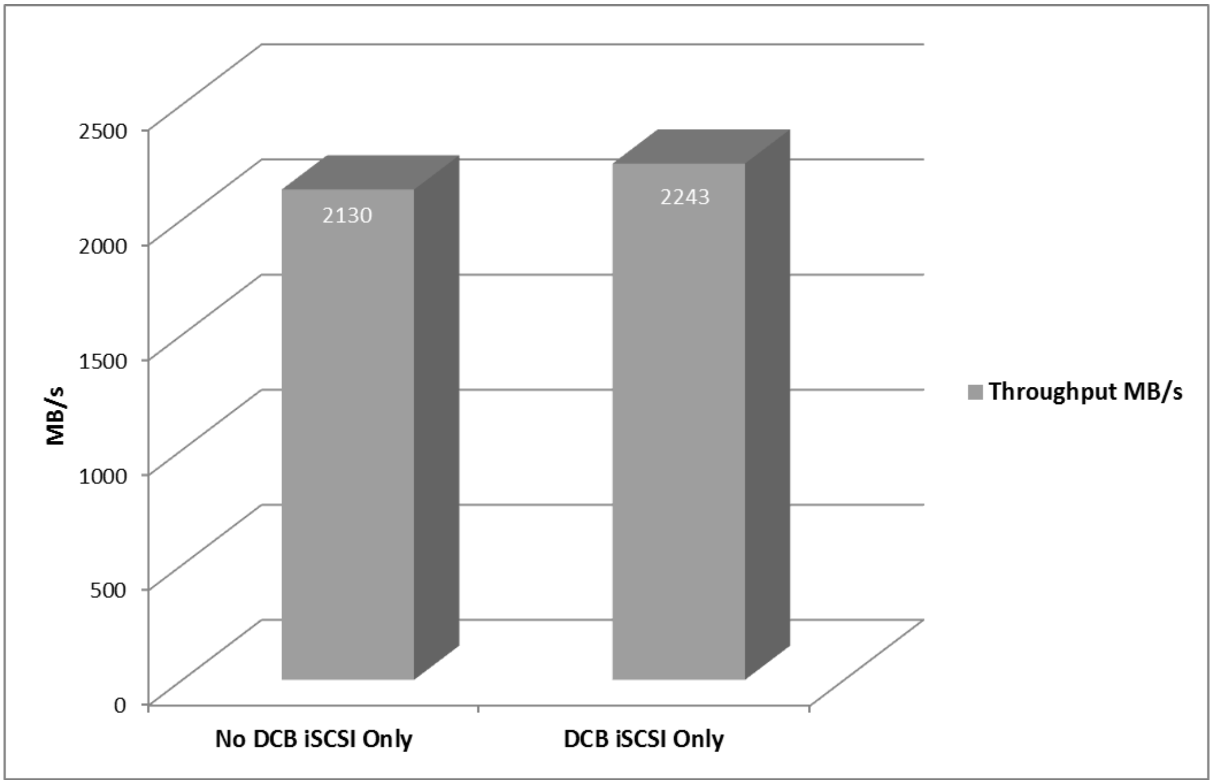

Figure 7 Throughput comparison for DCB vs. non-DCB for iSCSI only traffic

For IOs per second, the result was consistent with a difference of approximately 5%.

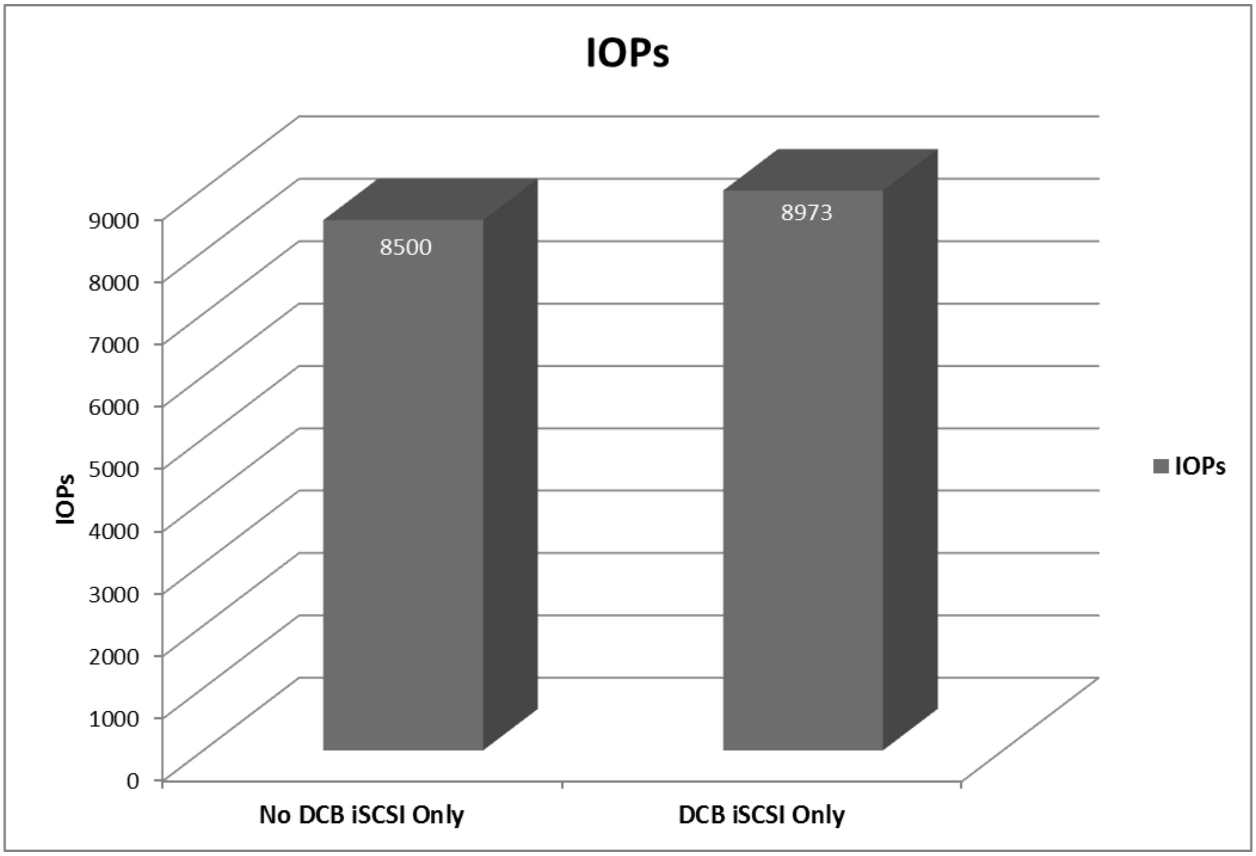

Figure 8 IOs per second comparison for DCB vs. non-DCB for iSCSI only traffic

5.2 DCB versus non-DCB in a shared network

Finally, the throughput for both DCB and non-DCB environments was compared for fully converged traffic. In this series of tests, the same iSCSI load was placed on the network while simultaneously injecting TCP traffic to another server. This load was run at a rate that would consume approximately 80% of line rate if it were the only traffic on the network. The ETS settings on the switch were configured so that iSCSI should receive a minimum of 50% of the available bandwidth at any time. Tests were run for each of the traffic types alone to determine baseline utilization. These numbers were then divided by 2 to gain the expected bandwidth when converged, and compared to an actual test with both traffic types active.

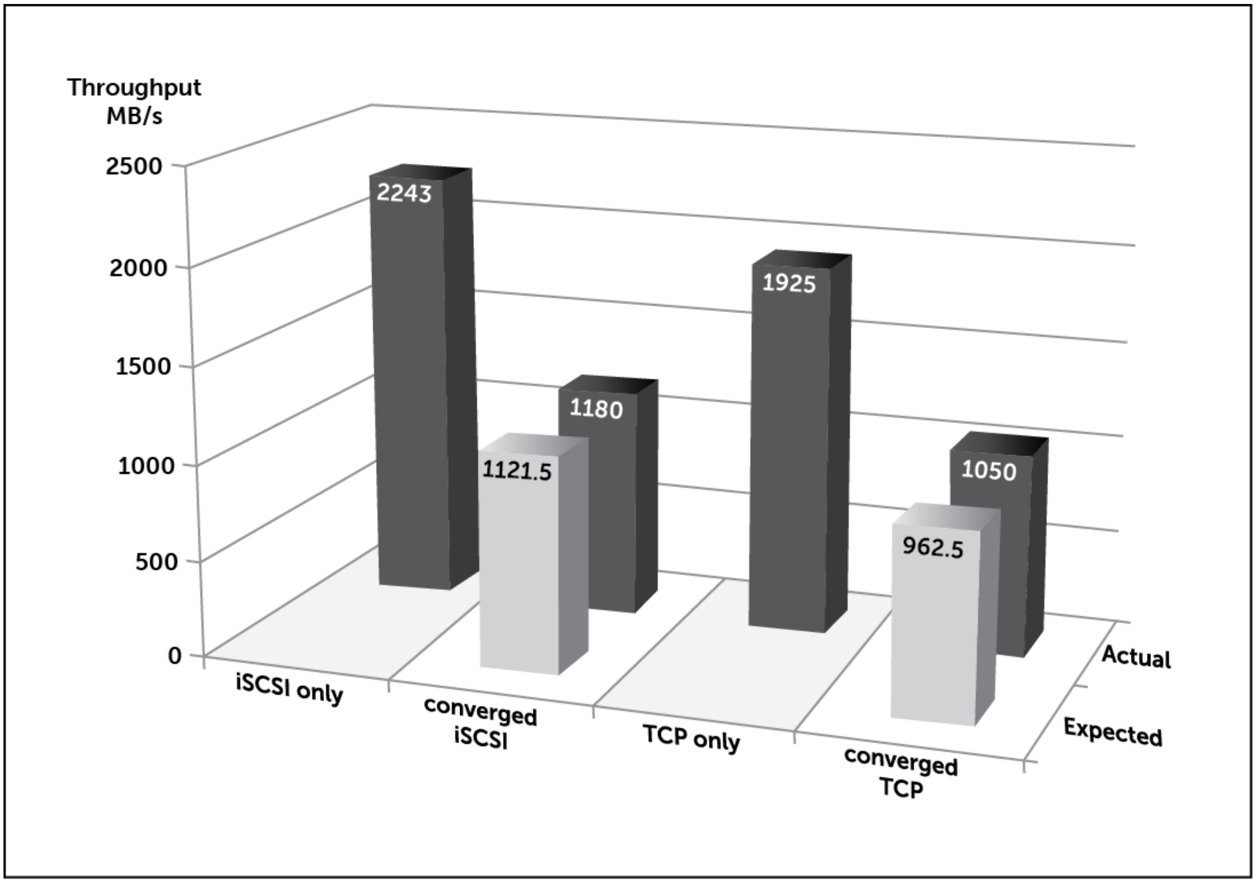

Figure 9 Throughput comparison between pure and converged traffic types

6 Implementation

6.1 Planning is paramount

As with any decision to move to new technology, understanding your needs and planning for them can be the difference between an unsuccessful migration and a successful one. There are many factors to be considered before moving to DCB. For each of the network streams to be converged, a complete understanding of the requirements is required. This includes the amount of bandwidth needed and future traffic needs, and is especially true for storage traffic streams such as iSCSI and FCoE.

One situation to consider is the amount of bandwidth guaranteed to storage through the use of ETS. When initially considering the numbers, the storage administrator should strongly consider future needs of the network. This can be compounded by the fact that the bandwidth sharing protocol in DCB has a margin of error of up to 10%. If you were to decide on a minimum target bandwidth allocation too close to what is needed, actual performance may not meet your expectations as throughput is squeezed to allow for other network traffic.

6.2 New data center environment

A prime example of an appropriate time to move to DCB can be when designing and implementing a new datacenter. With the infrastructure and cable plant optimizations of DCB, there can be significantly fewer cables to purchase, run and manage. While the use of fiber interconnects will be higher from a per-cable standpoint, the lower number can offset the initial costs quickly. Also, as 40 Gb and 100 Gb connections begin to be introduced in the future, this cable plant can remain viable for many years.

Green, environmentally conscious datacenters are on every company's agenda today. As traffic converges on the network, the need for separate fabrics with various flavors of storage, as well as LAN traffic becomes unnecessary. With the implementation of fewer DCB enabled switches, power and cooling costs can be reduced in the new data center as well.

6.3 Multiple storage protocols

As users storage needs continue to explode over the coming years, the management of multiple SANs and their associated fabrics continues to become more and more unwieldy. The advent of FCoE, as well as the tremendous growth of iSCSI in the datacenter provides a perfect chance to converge and collapse those fabrics. DCB is an essential piece of that puzzle. DCB is required for FCoE to even be implemented, and the bandwidth sharing properties of ETS will finally allow the various flavors of storage to coexist on the same network.

When planning to implement disparate storage protocols on the same fabric, ensure that network equipment that can handle the unique demands of this environment is purchased. One feature that can help prepare a network for future needs is the ability to manage multiple priority queues. Very few CNAs available today can perform this feature in hardware, which can have adverse effects in the situation where there are more than two priorities in use on the fabric. A prime example of this is the use of FCoE (Priority 3) and iSCSI (Priority 4) on the same converged fabric. The ability to advertise and understand the iSCSI TLV (described in the "DCB terminology" section on page 3) is going to have a huge impact on the manageability of this type of environment. All of these changes occur automatically because of the iSCSI TLV instead of manually.

6.4 Best practices

Based on the information presented above, there are several Dell recommendations:

As can be seen from the results, converging traffic in a non-DCB environment is not recommended. If the infrastructure is not DCB-ready today, Dell recommends that storage traffic be placed on a separate Ethernet network from other Ethernet traffic such as LAN or Management networks. Barring the ability to separate the traffic physically, the use of VLANs to separate the storage and LAN traffic is encouraged and should only be used for small SAN implementations of no more than 3-4 arrays.

When preparing to deploy a DCB-ready environment, planning of both an infrastructure and DCB environment are paramount. Expected bandwidth requirements of both the storage and LAN traffic that is intended to be converged must be taken into account. When calculating these numbers, be sure to provide for future growth requirements of both types. While ETS settings can be changed on the fly in the future, this can have unexpected results if the change removes too much guaranteed bandwidth from another traffic class.

Even in dedicated iSCSI SAN environments, there are benefits to be gained from implementing a Data Center Bridging environment. The use of PFC provides a better flow control model for storage traffic, and prepares the environment for the introduction of converged traffic in the future with minimal disruption to existing traffic. For these environments, first generation DCB switches that may only support PFC are great options as ETS is not required for a non-shared network infrastructure such as an iSCSI only DCB network.

For deployment and proper operation of a DCB-enabled Ethernet infrastructure using EqualLogic storage, support for the iSCSI TLV is required (check the manufacturer's documentation). EqualLogic firmware requires the advertisement of the iSCSI priority using the iSCSI TLV function in order to function in a DCB environment. Also ensure that the DCB-capable switches chosen support both PFC and ETS when there is a need to converge multiple traffic types. While PFC alone is acceptable in a separated Ethernet environment, the lack of ETS in a converged environment will result in less than ideal results.

Appendix A Component versions

A.1 Components

The following table lists the components used, as well as the firmware and driver versions at the start of the test cycle.

Table 3 Component versions

Component	Purpose / Description	Version number
Server		
Dell PowerEdge R710	Initiator/TCP Sink	BIOS 2.2.11
Windows 2008 R2		Current updates at test cycle start
Storage		
EqualLogic PS6010E	DCB Compliant Storage	5.1.0
Switch		
Brocade B-8000e	DCB Compliant Switch	FOS 6.4.1a (Brocade GA, Dell supported March 2011)
Initiator		
MS iSCSI Initiator		
Intel x520 CNA	DCB Converged Network Adapter	2.8.32.0
Software		
EqualLogic Host Integration Toolkit	MS MPIO Device Specific Module	3.5.1
Test Tools		
Medusa Labs Test Tools	iSCSI storage traffic generation	5.0.0.119262
Management and Monitoring		
SAN Headquarters	Monitoring of EqualLogic array statistics	2.1
EQLMonitor script	Monitoring of TCP retransmit %	
Perfmon	Monitoring of host statistics	

Appendix B Switch configuration

B.1 Switch settings

The following table summarizes the settings used on the PowerConnect B-Series 8000-E DCB capable switch. These settings were required to enable DCB features as well as normal EqualLogic optimizations in the non-DCB tests. There is also a sample configuration from a switch used in the testing.

Table 4 Switch settings

Switch model	PowerConnect B-Series 8000e

Switch inter-connection	• Dynamic Link Aggregation Group (LACP - LAG)
	  o channel-group 1 mode active type standard
	• Rapid Spanning Tree cost (only on 4 switch config)
	  o Interface Port-channel 2
	  o Spanning-tree cost 200000000

Global Switch Settings	• Define PFC/ETS Settings
	  o cee-map iscsi
	  o priority-group-table 0 weight 50
	  o priority-group-table 1 weight 50 pfc
	  o priority-table 0 0 0 0 1 0 0 0
	• Define LLDP Settings
	  o protocol lldp
	  o system-name B8000-sw1
	  o system-description Brocade 8000
	  o advertise optional-tlv port-description
	  o advertise optional-tlv system-name
	  o advertise optional-tlv system-capabilities
	  o advertise optional-tlv system-description
	  o advertise dcbx-iscsi-app-tlv
	• Enable Rapid Spanning-Tree
	  o protocol spanning-tree rstp

Individual Port Settings	• Set Jumbo MTU
	  o mtu 9208
	• Set Layer 2 mode
	  o switchport
	• Set for converged traffic
	  o switchport mode converged
	• Enable Port
	  o no shutdown
	• Define iSCSI CoS
	  o lldp iscsi-priority-bits 0x10
	• Apply defined CEE-MAP named iscsi
	  o cee iscsi
	• Activate MAC Pause Flowcontrol (only if not using DCB)
	  o qos flowcontrol tx on rx on

Switch Firmware	v6.4.1a
Host-Switch Cable Type	Brocade 10G Active 1m (58-1000026-01) 1 meter cable

Switch model	PowerConnect B-Series 8000e
Array-Switch Cable Type	Array: SFP+ SR Optical Transceiver ((DP/N 0N743D); Switch: SFP+ SR Optical Transceiver (10G-SFPP-SR-8); LC-LC Fiber Optic Cable
Switch-Switch Cable Type	Switch: SFP+ SR Optical Transceiver (10G-SFPP-SR-8); LC-LC Fiber Optic Cable

B.2 Sample switch configuration

```
sh run
!
protocol spanning-tree rstp
!
cee-map iscsi
 priority-group-table 0 weight 50
 priority-group-table 1 weight 50 pfc
 priority-table 0 0 0 0 1 0 0 0
!
interface Vlan 1
!
interface TenGigabitEthernet 0/0
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/1
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/2
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/3
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/4
```

```
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/5
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/6
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/7
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/8
 mtu 9208
 channel-group 1 mode active type standard
 no shutdown
 lldp iscsi-priority-bits 0x10
 lacp timeout long
 cee iscsi
!
interface TenGigabitEthernet 0/9
 mtu 9208
 channel-group 1 mode active type standard
 no shutdown
 lldp iscsi-priority-bits 0x10
 lacp timeout long
 cee iscsi
!
interface TenGigabitEthernet 0/10
 mtu 9208
 channel-group 1 mode active type standard
 no shutdown
 lldp iscsi-priority-bits 0x10
 lacp timeout long
 cee iscsi
!
```

```
interface TenGigabitEthernet 0/11
 mtu 9208
 channel-group 1 mode active type standard
 no shutdown
 lldp iscsi-priority-bits 0x10
 lacp timeout long
 cee iscsi
!
interface TenGigabitEthernet 0/12
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/13
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/14
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/15
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/16
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/17
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
```

```
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/18
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/19
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/20
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/21
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/22
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
 cee iscsi
!
interface TenGigabitEthernet 0/23
 mtu 9208
 switchport
 switchport mode converged
 switchport converged allowed vlan all
 no shutdown
 lldp iscsi-priority-bits 0x10
 spanning-tree edgeport
```

```
 cee iscsi
!
interface Port-channel 1
 mtu 9208
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 no shutdown
!
protocol lldp
 system-name B8000-sw1
 system-description Brocade 8000
 advertise optional-tlv port-description
 advertise optional-tlv system-name
 advertise optional-tlv system-capabilities
 advertise optional-tlv system-description
 advertise dcbx-fcoe-app-tlv
 advertise dcbx-iscsi-app-tlv
 advertise dcbx-fcoe-logical-link-tlv
!
line console 0
 login
line vty 0
 exec-timeout 0 0
 login
line vty 1 31
 login
!
end
```

Appendix C Initiator configuration

C.1 Initiator settings

The following table lists the server-specific changes made to the Converged Network Adapter. Note that the QoS Packet Scheduler *must* remain enabled in the Adapter Properties dialog box for DCB to work properly.

Table 5 Initiator settings

Host Converged Network Adapter (CNA)	
Model	Intel x520 – Dual Port
Advanced Network Services Driver	Installed
FCoE using Data Center Bridging Driver	Not Installed
iSCSI using Data Center Bridging Driver	Installed
iSCSI Initiator	Microsoft Windows Server 2008R2 SP1
Jumbo Frames	Enabled at 9000 bytes per frame
DCBx Setting	Use values provided by switch
MPIO Configuration	
Dell EqualLogic Host Integration Toolkit	Version 3.5.1
Dell EqualLogic MPIO Device Specific Module	Default settings: • Maximum Sessions per Slice: 2 • Maximum Sessions per Volume: 6

Related publications

The following Dell publication is a recommended source for additional information.

EqualLogic Configuration Guide:

http://www.delltechcenter.com/page/EqualLogic+Configuration+Guide

EqualLogic, PS Series Group Administration Online Help:

http://psonlinehelp.equallogic.com/V5.1/groupmanager.htm