Transition to the Data Center Bridging Era with EqualLogic PS Series Storage Solutions

A Dell Technical Whitepaper

Storage Infrastructure and Solutions Engineering
Dell Product Group
February, 2012

Table of Contents

Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions team between November 2011 and February 2012 at the Dell Labs facility in Round Rock, Texas.

The team that created this whitepaper:

Lakshmi Devi Subramanian, Chris Almond, and Tony Ansley

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

Mike Kosacek and Jerry Daugherty

1 Introduction

Ethernet is the most widely deployed networking technology today. This is a standard technology and deployed in virtually every data center throughout the world. Despite Ethernet's traditional advantages, special purpose switching fabric technologies needed to be developed to address special requirements, such as storage and data area networking. Using distinct networks for data, management, and storage can be more complex and costly than using a single converged network. Organizations are looking for ways to combine their various storage and data/client networks into a single physical, converged fabric to help reduce the total cost of ownership of the data center infrastructure, connect multiple storage islands, and enhance storage scalability while still providing the ability to meet individual service level agreements for each network type. The introduction of 10Gb Ethernet provides enough throughput to meet this goal. Data Center Bridging (DCB) provides the enhancements to Ethernet that are necessary to complete this convergence process.

DCB provides a set of standards-based extensions to traditional Ethernet, offering a lossless data center transport layer that allows the convergence of different networks onto a single unified network fabric. For network designers ready to begin using DCB in the data center, the challenge now is knowing how to design and implement this technology.

This document is intended for use by network administrators, storage administrators, sales engineers and IT professionals who are interested in an introduction to Data Center Bridging, and who expect to begin planning a transition to converged DCB in their data centers. This document also focuses on operations and design considerations for EqualLogic iSCSI SAN network/storage administrators.

2 Data Center Bridging overview

2.1 Why DCB was created

Ethernet, being a pervasive technology, continues to scale in its bandwidth to 10/40/100Gb and also drop its cost. This provides a great opportunity for other technologies-such as Fibre Channel and Remote Direct Memory Access networking solutions to leverage Ethernet as a converged universal fabric.

For example, as 10GbE becomes more common, 8Gb Fibre channel has been shipping for couple of years, and 16Gb is beginning to ship now. Ethernet continues to develop the next generation focusing on standard bandwidths of 40/100Gb. Fibre Channel over Ethernet (FCoE) was developed to more fully leverage the features of 10G Ethernet and beyond and provide more options to end users for SAN connectivity and networking. However, Fibre channel was not designed to work on unreliable network but was developed with an idea of having dedicated fabrics to avoid packet losses. So, extensions needed to be added to Ethernet and these extensions collectively are called Data Center Bridging. The addition of these standards to the Ethernet protocol suite provides reliability without incurring the penalties traditionally found in Ethernet, while providing faster transmission rates and the ability to control bandwidth allocation.

2.2 Benefits of DCB

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 availability. The ability to guarantee bandwidth for any application or traffic type is an essential part of the DCB environment.

Through the use of Enhanced Transmission Selection (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. Priority-based Flow Control (PFC) manages the multiple flows of network data to ensure frames are not dropped for traffic types that require minimal packet loss (such as storage or voice).

Heavy virtualization

As the move to virtualization 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 10Gb connections while still maintaining the required performance characteristics is compelling. Through the use of DCB you can realize the benefits of a converged network, such as fewer physical connection paths, which can lead to easier maintenance and simplified management.

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 decrease. By using fewer cables, the airflow characteristics and cooling efficiency in data center racks improves. One or two 10Gb fibre connections per server take up significantly less volume than four or more 1Gb copper cables per server. With more cables multiplied across the servers in a rack, cooling airflow paths can be inhibited.

Cost
As the cost of 10Gb Converged Network Adapters (CNAs) continues to fall, the economic benefits of converging multiple traffic flows onto 10Gb Ethernet will increase. As the size and complexity of your data center increases, it will become far more cost-effective to purchase and deploy fewer ports of 10Gb DCB Ethernet than many more ports of 1Gb non-DCB Ethernet.

2.3 DCB terminology

Data Center Bridging (DCB) is a set of IEEE 802.1Q extensions to standard Ethernet, that provide an operational framework for unifying Local Area Networks (LAN), Storage Area Networks (SAN) and Inter-Process Communication (IPC) traffic between switches and endpoints onto a single transport layer.

For an end-to-end DCB network to be properly configured there are various 802.1Q components utilized. The standards include:

- **Priority-based Flow Control:** (PFC; IEEE 802.1Qbb) Expands the function of the standard class of service structure of Ethernet to apply pause functionality on traffic, based on the class of service settings. This ensures that one type of traffic does not affect other types of traffic that have differing classes of service.

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

- **Datacenter Bridging Capability Exchange:** (DCBx; IEEE 802.1Qaz) Allows for communication of DCB capabilities between network components. This protocol enables network components to sense the DCB capabilities of neighboring devices and ensure consistent configuration.

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

2.3.1 Priority-based Flow Control (PFC)

Priority-based Flow control is an evolution of the concept of Flow Control originally implemented in the MAC Pause frame feature of Ethernet (IEEE 802.3x). Pause frames provide a simple way for control of segment traffic, by allowing a NIC to send a request to an adjacent port to stop transmitting for a specific time period. With no granularity applied to this request; all Ethernet frames between two ports are stopped during the pause. But Ethernet NIC manufacturers were not always required to implement it on all devices. Because of this, you can have a situation where a sending device does not understand what to do with the pause frame sent by the receiving device. Frames would continue to be transmitted in this case, which would lead to dropped frames at the receiving device. A similar problem can also arise during what is referred to as "Head-of-Line Blocking", where traffic from multiple ingress ports is squeezed into a single egress port. When this port's buffer fills up, and it is unable to accept any more incoming packets, those packets must be discarded (dropped). This can eventually lead to degradation of the network's overall performance.

In PFC a request is sent to the sender to stop transmitting, but it in addition PFC leverages the idea of classes of traffic to apply more granularity to the process. In a DCB environment, all traffic is tagged with a Class of Service (CoS) using the virtual LAN (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 some (lower priority) TCP/IP traffic is dropped, while iSCSI traffic is tagged with a different CoS, giving it a higher priority. As a benefit of PFC, once its use has been negotiated both receiver and sender must adhere to it.

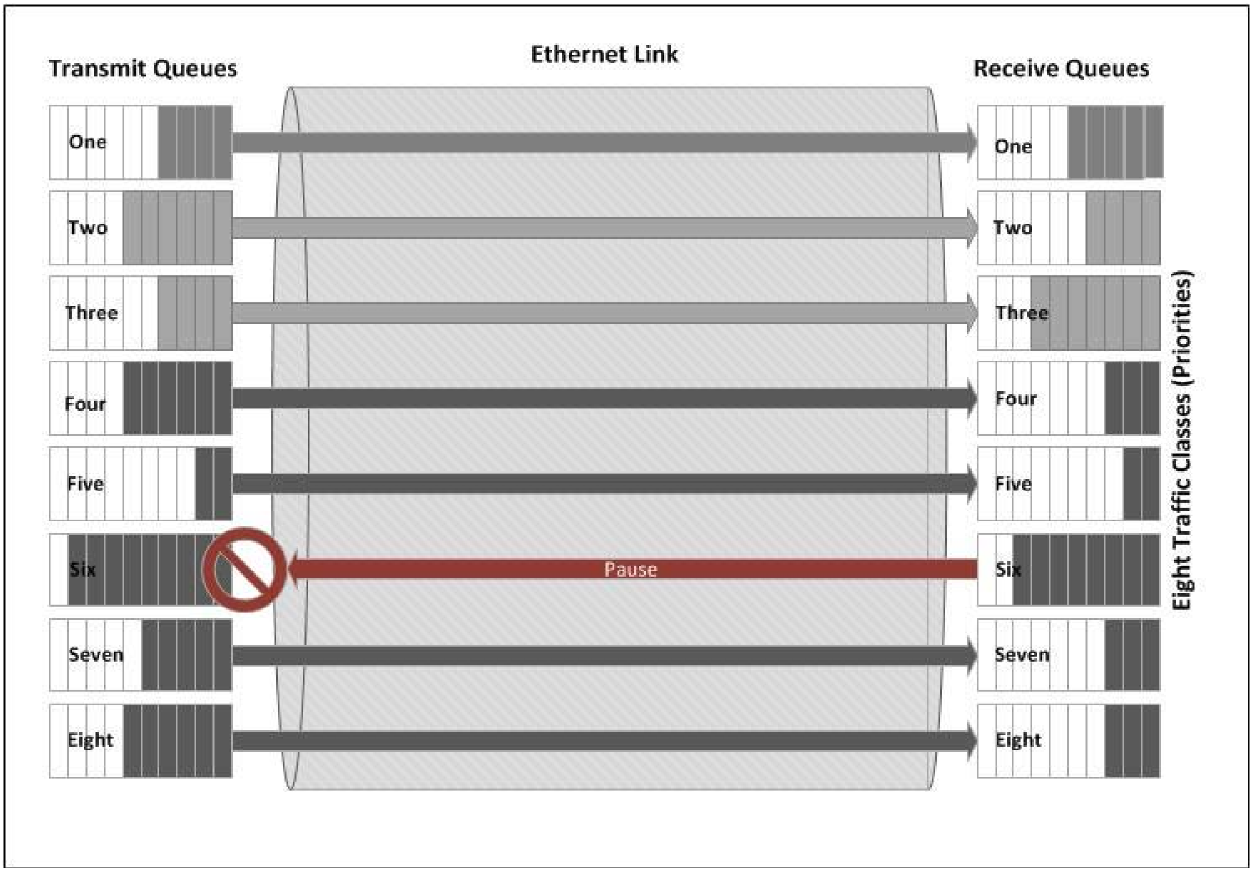

Figure 1 Priority-based Flow Control

2.3.2 Enhanced Transmission Selection (ETS)

Enhanced Transmission Selection is a mechanism for guaranteeing a minimum percentage of bandwidth to a traffic class. A traffic class contains one or more CoS's defined using VLAN Q-tag. Each traffic class is then assigned a percentage of bandwidth (with setting granularity down to 10%). 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 currently using all of its minimum allocation, then other traffic classes may use it. However, as soon as the original traffic class requires its bandwidth again, the other traffic flow may be throttled to allow recovery to the minimum level allocated for that class. This is accomplished through the use of PFC. PFC will issue a pause for the required traffic classes in a manner to allow the bandwidth to be regained while minimizing the probability of dropped frames occurring in the throttled traffic classes. 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.

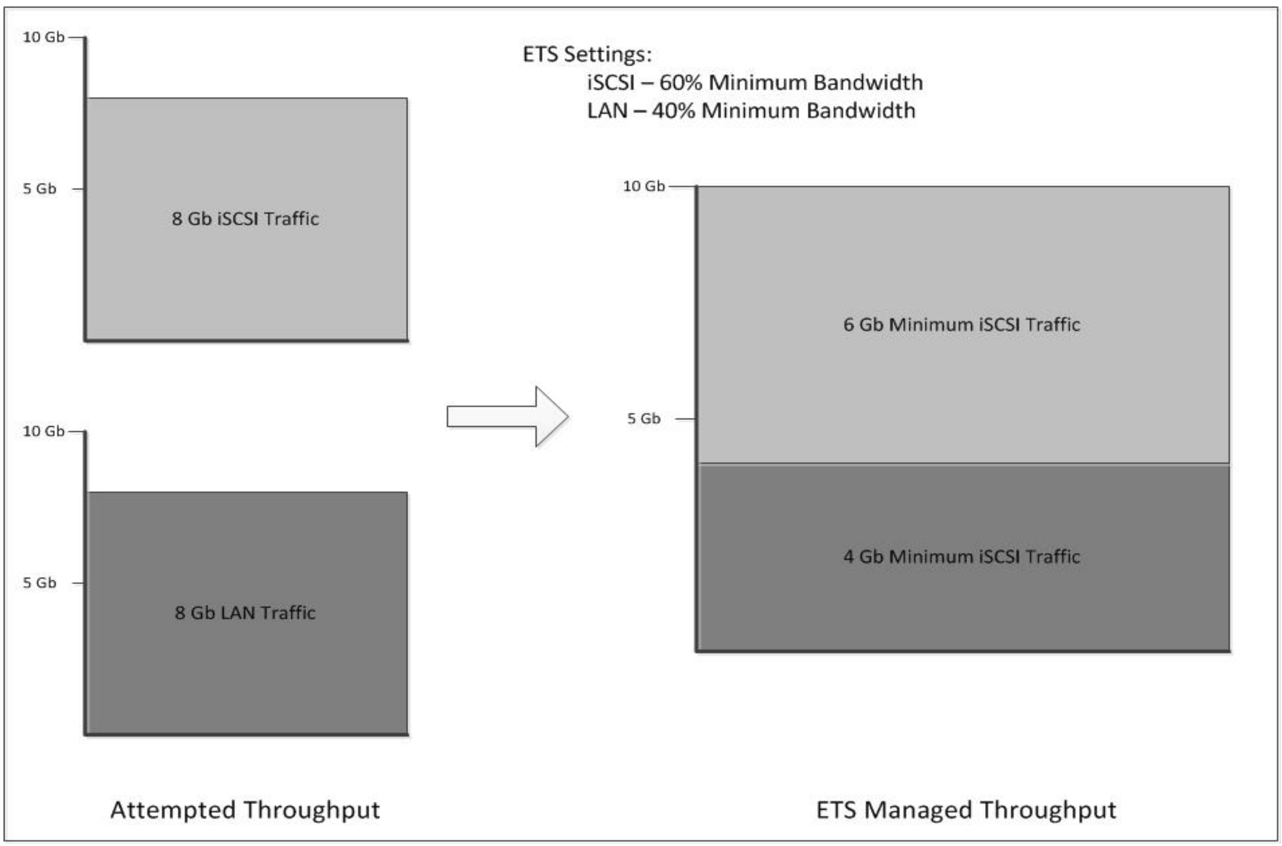

Figure 2 Enhanced Transmission Selection

2.3.3 Congestion Notification (CN)

Congestion Notification is a mechanism for managing congestion throughout a DCB fabric or domain. Ideally, that fabric would consist of interconnected switches and end-devices that all 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. For instance a switch would likely use PFC Pause frames to immediately slow the traffic, but if it determines over time that the traffic is continuing to flow at a high rate, then it can begin using CN to throttle the whole thing down.CN can work hand-in-hand with PFC to control congestion and overruns throughout the DCB fabric.

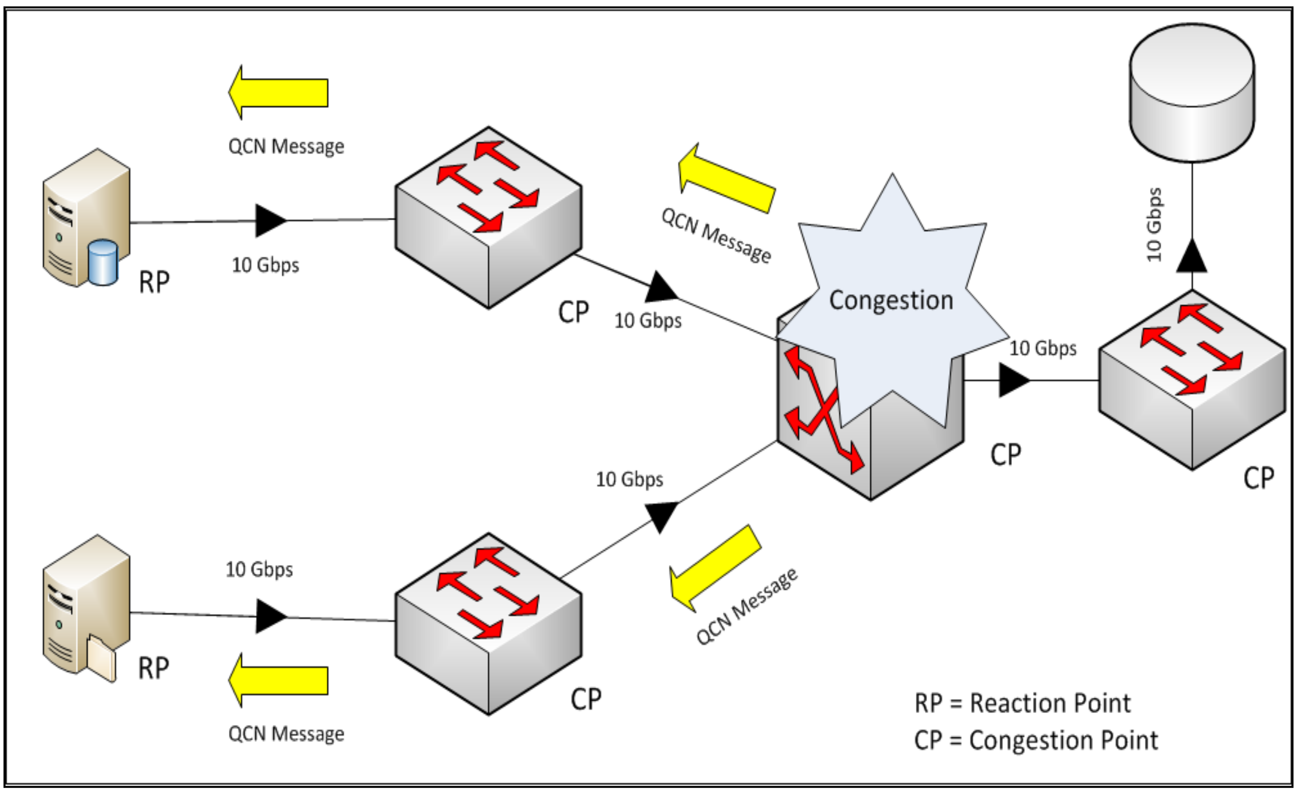

Figure 3 Congestion notification

2.3.4 Data Center Bridging Capability Exchange (DCBx)

Datacenter Bridging Capability Exchange 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 a device's 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 and isolate 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 desired class of service to all iSCSI packets. Once this has been defined, PFC and ETS can manage the iSCSI traffic as its own class.

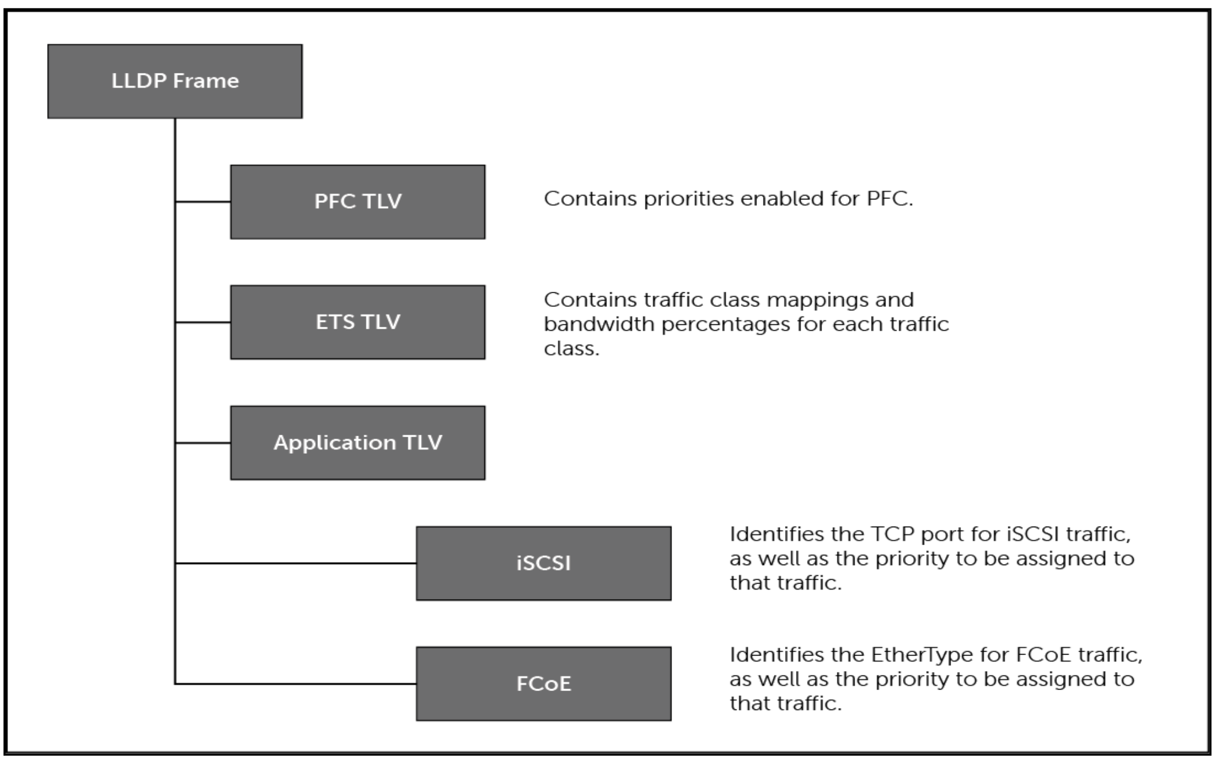

Figure 4 Data Centre Bridging Capability Exchange

3 iSCSI in a converged data center

3.1 Challenge for iSCSI in a shared network

In a traditional, non-DCB Ethernet iSCSI environment, it is a recommended practice to have iSCSI packets flow in a physically isolated SAN (with dedicated switches/ cabling, etc.), so that the iSCSI SAN network traffic is minimally affected by other, lower priority Ethernet traffic. In virtualized environments, it is a recommended practice to employ multiple NICs for traffic isolation (iSCSI, vMotion, production, management, etc.) The main reason for this is to avoid latency issues, unpredictable performance and security concerns. The rationale behind this is that sharing iSCSI network traffic with other traffic causes congestion and competition for limited physical resources in the network. This results in:

- Packet loss, creating non-deterministic performance
- TCP/IP error recovery is expensive from a predictable latency perspective

3.2 iSCSI TLV's role in Enterprise iSCSI

For iSCSI SAN traffic to realize the maximum benefit within a DCB based network infrastructure, one additional requirement needs to be implemented on top of support for PFC and ETS. This additional capability is part of the DCBx standard and is called "TCP Application Type-Length-Value (TLV)" support. This form of TLV allows the DCB infrastructure to apply unique ETS and PFC settings to specific subsets of the TCP/IP traffic stream. This is accomplished by having the switch identify these various sub-streams based on their TCP "socket" or "port" number which is always available in a TCP/IP frame.

TLVs based on TCP sockets are not a required part of the standard, but it is very important to iSCSI storage solutions. iSCSI – as a TCP-based application – has a standard TCP port number (port 3260) and is specifically known as the "iSCSI TLV". Over time, switch, NIC, CNA and router vendors who are dedicated to providing the most efficient and flexible solutions for iSCSI network design, will implement the iSCSI TLV.

3.3 Enterprise iSCSI solution example

Below is an example for enterprise iSCSI Solution that fully supports DCB. The solution includes the elements below.

- PowerEdge servers (R710) in VMware cluster
- EqualLogic array (PS6010XVS)
- PowerConnect switch (B-8000e)
- Emulex Enterprise iSCSI HBAs

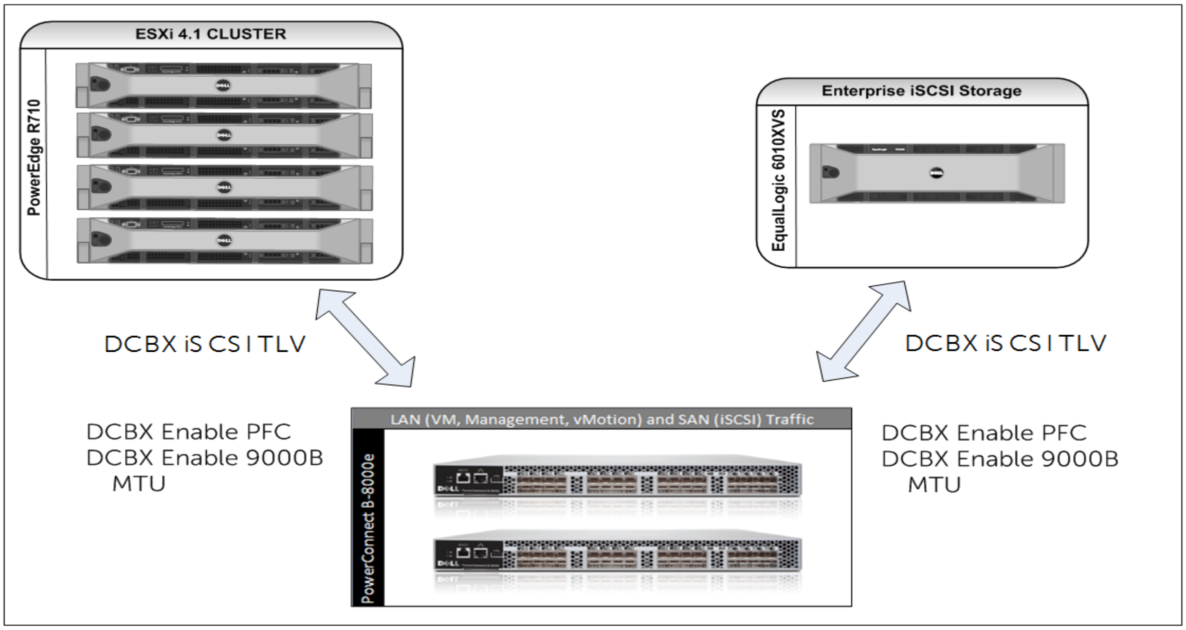

Figure 5 Sample enterprise iSCSI solution with DCB

3.4 EqualLogic DCB support

For EqualLogic environments using DCB over a shared network infrastructure, support for PFC, ETS and iSCSI TLV are 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 non-iSCSI TCP/IP traffic. The iSCSI TLV identifies what TCP port iSCSI traffic is using, so that the end station can properly assign the desired class of service. Once a class of service has been assigned, PFC and ETS can manage the iSCSI traffic throughout DCB shared network infrastructure.

EqualLogic PS Series 10Gb arrays automatically configure ETS and PFC settings provided by the DCB network infrastructure to which they are attached. As befits our ease of use goals around any design feature added to the EqualLogic PS series family, iSCSI TLV provides a “hands-off” method of configuring PS Series arrays to use a DCB network infrastructure, thus ensuring accurate array configuration and minimizing errors that are prevalent in manual network and array configuration methods.

With the release of the EqualLogic Array Software 5.1.1, used in conjunction with the appropriate host connectivity and switching products, EqualLogic PS Series customers now have the ability to deploy their storage into DCB converged data center networks. By using this new set of standards, EqualLogic iSCSI storage can now be deployed into shared networks without concern for quality of service while gaining near “lossless” packet delivery. Dell EqualLogic is one of the first storage vendors to support Enterprise iSCSI over DCB.

4 Transitioning to DCB

4.1 Classes of DCB switches

DCB capable switches can be categorized into different classes. "Core DCB" switches, "Bridge DCB" switches and "non-DCB" switches. Figure 6 illustrates the relationship of these different switch classes when used together to provide a transitional DCB datacenter environment.

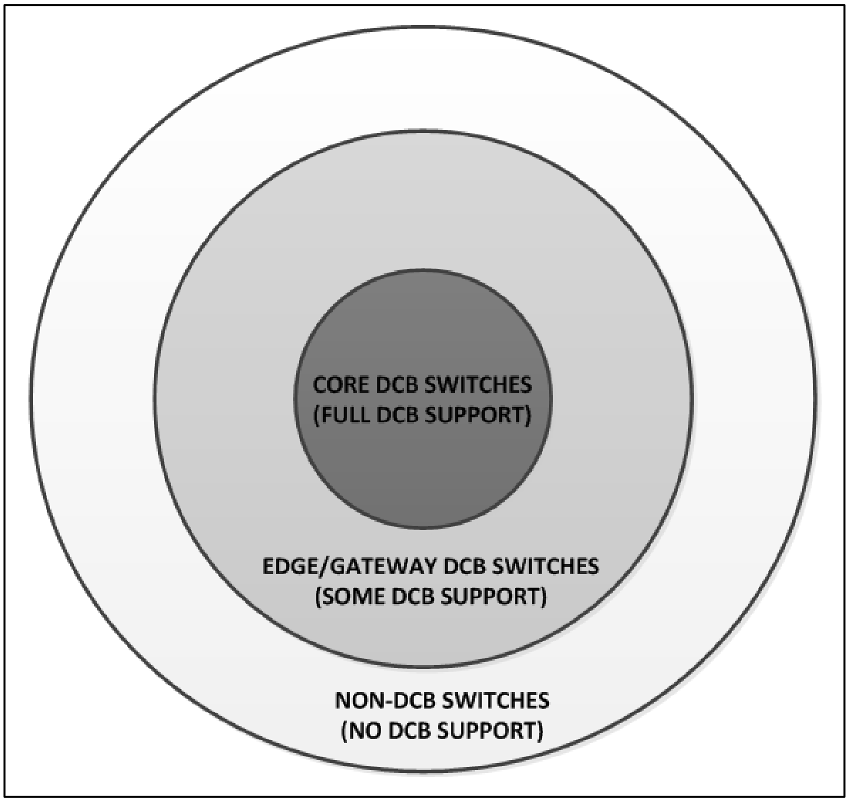

Figure 6 DCB switch classes

4.1.1 Core DCB switch

DCB "core" switches must have full DCB feature support, as specified in section 2.3. In order to achieve full DCB support, a core switch must be capable of supporting the full set of DCB standards. As a core DCB switch, this type of switch is expected to manage multiple, disparate streams of Ethernet traffic. To be able to provide this service, this class of switch must be able to manage flow control (through PFC), bandwidth guarantees for specific types of traffic (through ETS), and optionally congestion awareness (through Congestion Notification). In addition, it must be able to apply these controls across all types of traffic streams including those based on application TLVs like iSCSI.

4.1.2 Bridge DCB switch

"Bridge" DCB switches are switches that have, at a minimum, support for PFC, but may not support all features of DCB. These switches can be introduced as gateway, or bridge switches that work with the core DCB switches to achieve DCB functionality, without disrupting the existing network architecture containing non-DCB switches. In this role, these bridge switches manage only a single type of traffic–much like a non-DCB switch, but have the ability to honor upstream DCB management as it forwards and receives traffic from core DCB switches that are assigned to the specific devices directly attached to that bridge DCB switch (such as non-DCB aware iSCSI devices, or non-DCB downstream switches).In the transitional DCB datacenter, the non-DCB aware devices and switches can be attached to bridge DCB switches to gain access to the shared datacenter network or the non-DCB switches can be completely replaced by either a core or bridge switch (with addition of core switch) to achieve DCB in a non-DCB environment. As such, bridge switches simply “bridge” the non-DCB Ethernet network(s) with DCB Ethernet networks allowing non-DCB devices to leverage the new DCB functionality where it is most important i.e. traversing the shared datacenter core network.

4.1.3 Traditional non-DCB switch

Non-DCB switches are the switches we have all come to know and love and those that implement none of the DCB standards.

4.2 Data center of today

At present, most servers connect to LAN, NAS and iSCSI SAN networks using Ethernet based NICs. Many environments today are still 1Gb Ethernet, in which multiple server adapters and cables are used. A multiplicity of 1GbE connection paths leads to higher power and cooling costs. Figure 7 illustrates a sample datacenter with no DCB support, using separate NICS on each server to connect to the SAN and LAN switches.

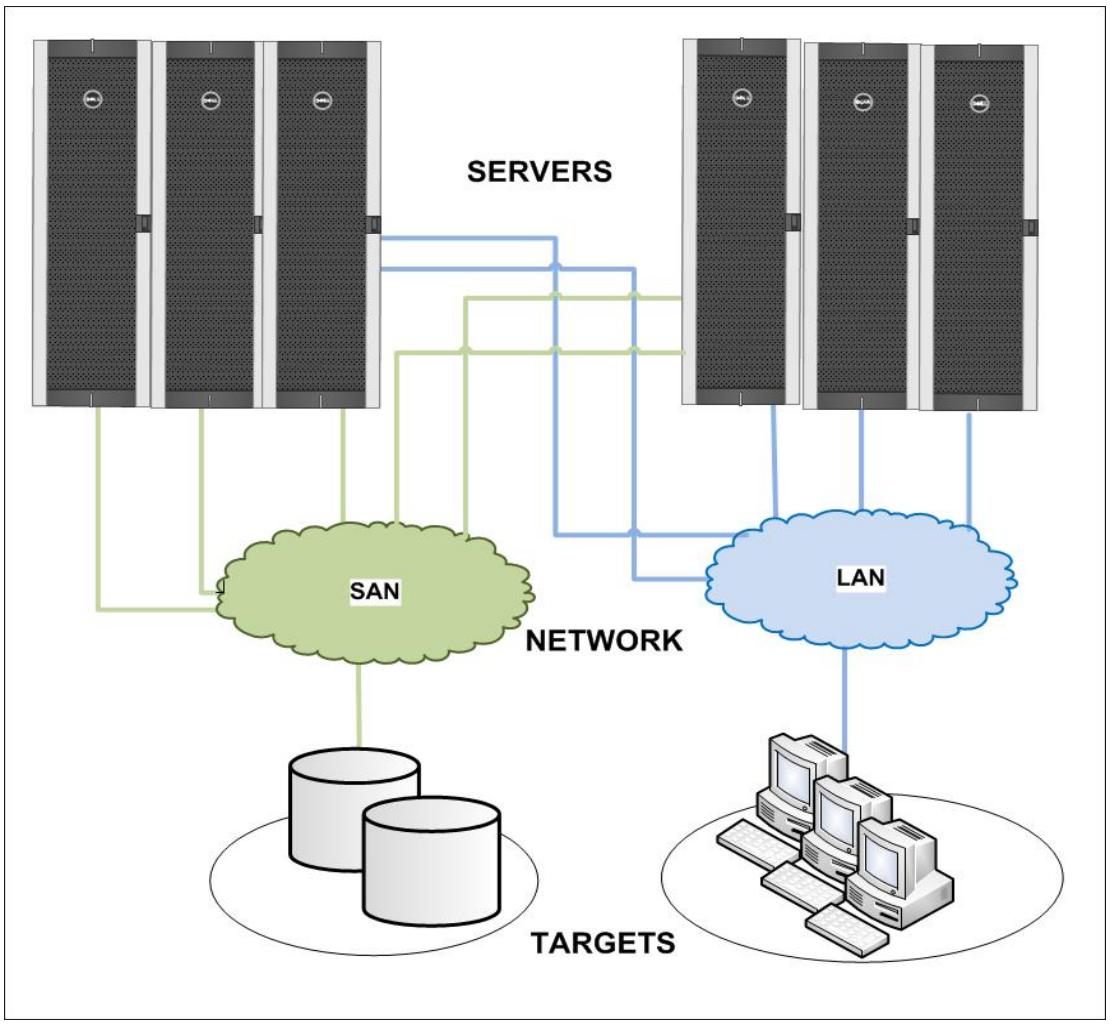

Figure 7 Data center with no DCB support

4.3 Data center in transition

4.3.1 Deployment considerations for DCB network

Consider the following before deploying a DCB network in the data center:

Goals: Identify why DCB needs to be implemented in the existing network infrastructure. For instance, some areas where DCB fits are,

- **Highly virtualized environments**
  As the virtualized environments grow, where a single server may host many distinct OS and application stacks, using higher bandwidth 10GbE becomes increasingly important. Use of 10GbE networking technology provides an opportunity to reduce the number of NICs, cables, and switches, creating greater efficiencies.

- **Reducing costs**
  The use of converged network adapters and DCB enabled switches would alleviate the operational costs involved in maintaining separate management networks requiring the use of multiple cables, NICS and switches.

- **Gaining competitive edge by implementing newer technology**
  An organization that is able to establish new services quickly without interruption to existing IT services with the ability to dynamically throttle bandwidth could be a competitive advantage.

Changes to accommodate DCB: Identify the changes that would be required to be done in the existing setup to support DCB. For example some of the changes are,

- **Changes in network topology**: Consider how the network topology would change with the introduction of DCB. This includes both SAN and LAN networks in the data center. This would include changes in networking parameters in the new DCB network to enable connectivity to existing LAN and SAN infrastructure.

- **Application requirements:** Consider application requirements for servers to be in the DCB network. This includes defining traffic types like Ethernet, FCoE etc. and also the bandwidth requirements for each traffic type.

- **Configuration Changes:** Consider the configuration changes in servers and applications to support a converged network. This may include enabling support for different traffic types and enabling redundant connectivity to SAN and LAN networks.

4.3.2 Deployment example scenarios

As discussed in section 4.3.1, how you transition an existing data center to DCB, depends on your goals. Below are few example scenarios and steps involved in transitioning to DCB.

4.3.2.1 Scenario 1: Simplify server (host) network connections

If the goal is to simplify the host network connections to several networks, following the below steps would lead to a converged DCB network. Figure 8 shows a host's connection path to LAN and SAN with no DCB support.

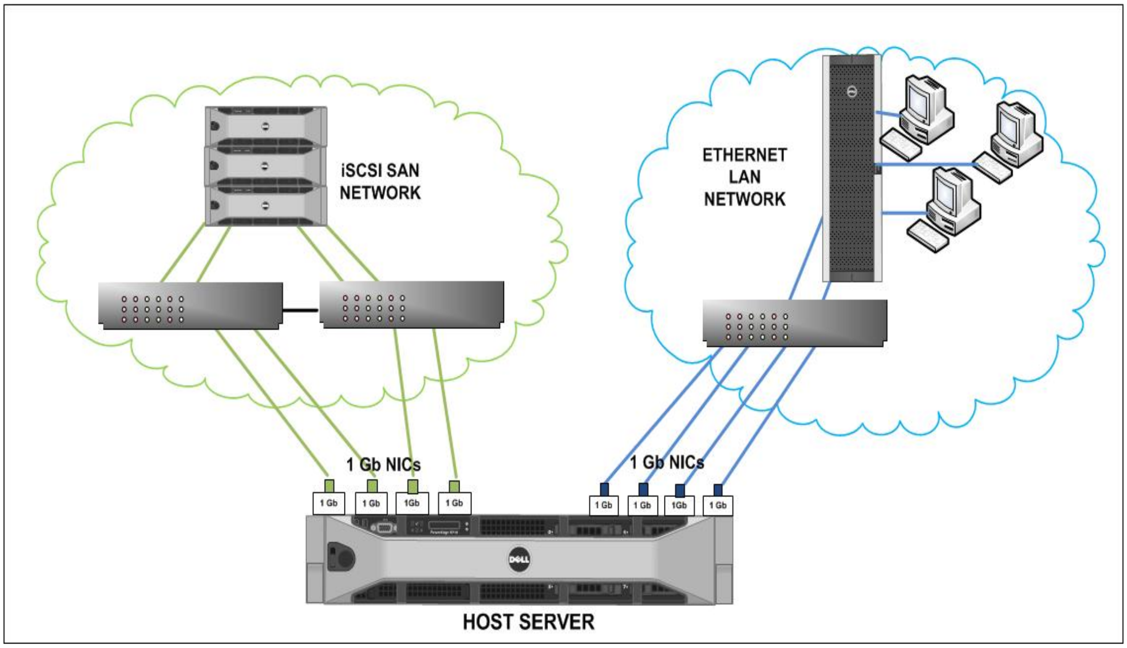

Figure 8 Host with separate NICS to different networks

Step 1: Build a DCB core

The first step is adding DCB core switches in the existing network infrastructure to enable the DCB transition.

Step 2: Add CNA's at the host servers

In this phase adding 10Gb Converged Network Adapters (CNAs) will provide converged connectivity between servers and the core-DCB switch before splitting the traffic to the legacy LAN and SAN infrastructure. (Refer Figure 9)

Step 3: Update the bridge /gateway switch to support DCB

Assuming the arrays in SAN network are 10Gb arrays with DCB support, upgrading the firmware on DCB capable switches or making hardware changes at the switch (the changes at the bridge switch to enable DCB standards depends on the switch vendor and the specific switch's capability) would enable these bridge switches for DCB compliance. Refer figure 9.

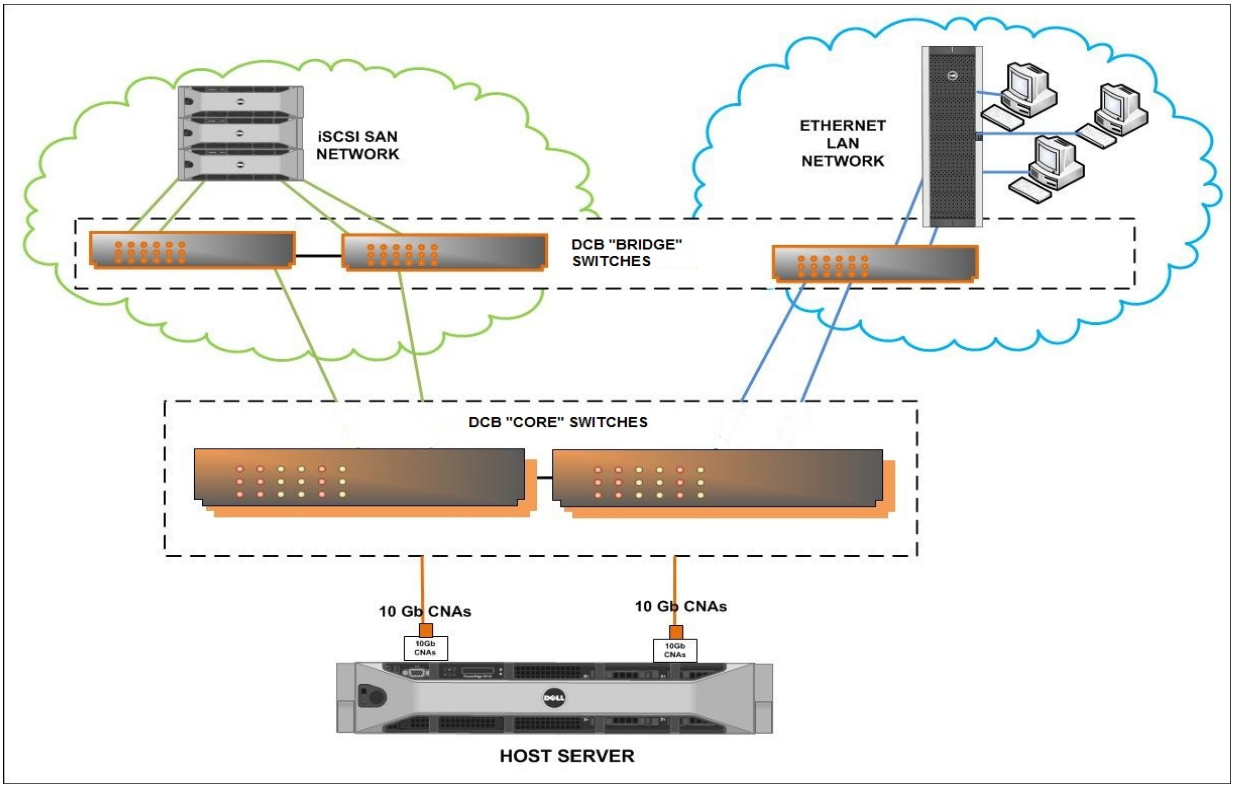

Figure 9 Data center that has transitioned to DCB, minimizing required host connection paths

4.3.2.2 Scenario 2: DCB support in 1Gb SAN storage array network

If the goal of the data center is to transition to DCB while still maintaining the use of 1GbE iSCSI SAN, the steps below would lead to a converged DCB network. While an in-depth discussion on integration of 1GbE and 10GbE iSCSI SANs is out of the scope of this paper, detailed information can be found in *Best Practices for Deploying a Mixed 1 Gb/10Gb Ethernet SAN using Dell EqualLogic Storage Arrays*, which is available on http://en.community.dell.com/techcenter/storage/w/wiki/creating-a-dcb-compliant-equallogic-iscsi-san-with-mixed-traffic.aspx

Step 1: Build a DCB core

The first step is adding DCB core switches to the existing network infrastructure to enable the DCB transition.

Step 2: Add CNA's at the host servers

In this phase adding 10Gb CNAs will provide converged connectivity between servers and the core-DCB switch before splitting the traffic to the legacy LAN and SAN infrastructure. (Refer Figure 10)

Step 3: Add DCB aware bridge /gateway switches to connect to the DCB core

Assuming the arrays in SAN network have no DCB support and both SAN and LAN are connected to non-DCB switches, addition of a layer of DCB aware bridge/gateway switches is required to connect to the core DCB switch. (Refer figure 10).

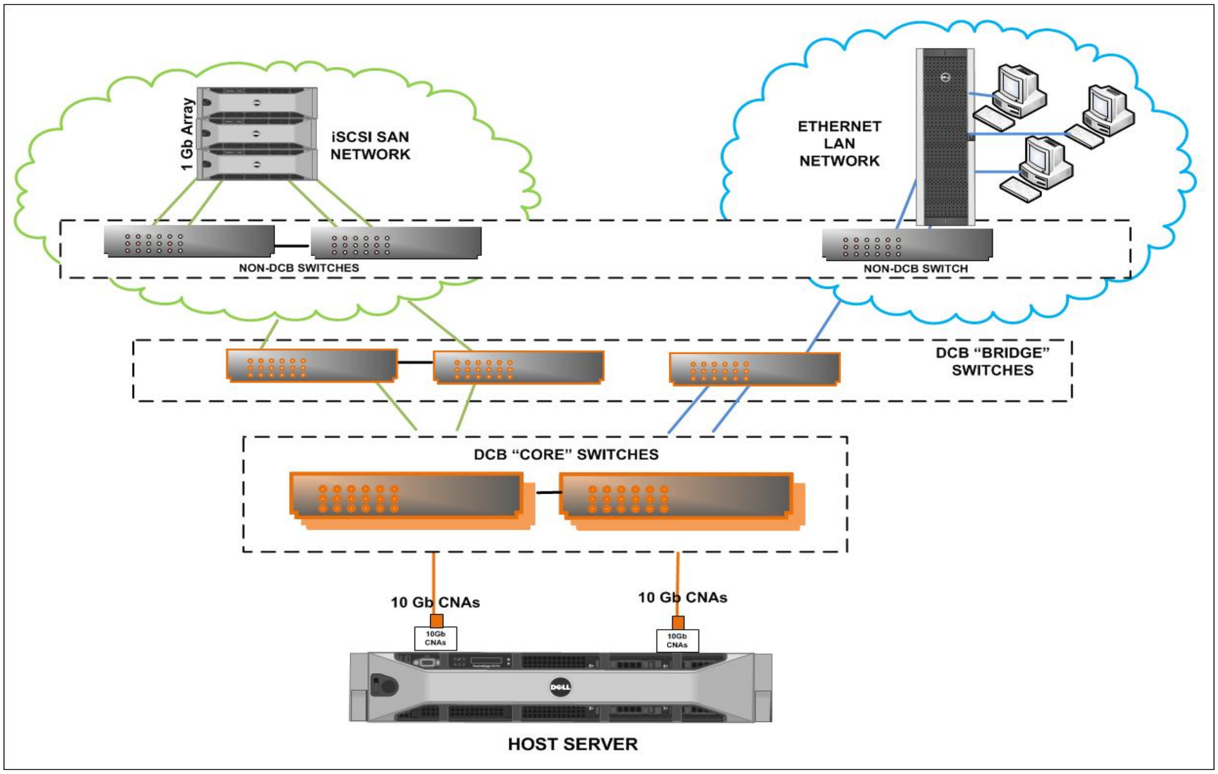

Figure 10 Data center transitioned to DCB

4.4 Future network design based on end-to-end DCB and 10GbE

In this scenario, all different types of networks can share a single network based on 10Gb Ethernet. The enablers here are 10GbE bandwidth combined with the lossless capability provided by full DCB support; CNA's and core DCB switches. The result is an end-to-end converged DCB aware datacenter. An example of future Datacenter is shown in figure 11.

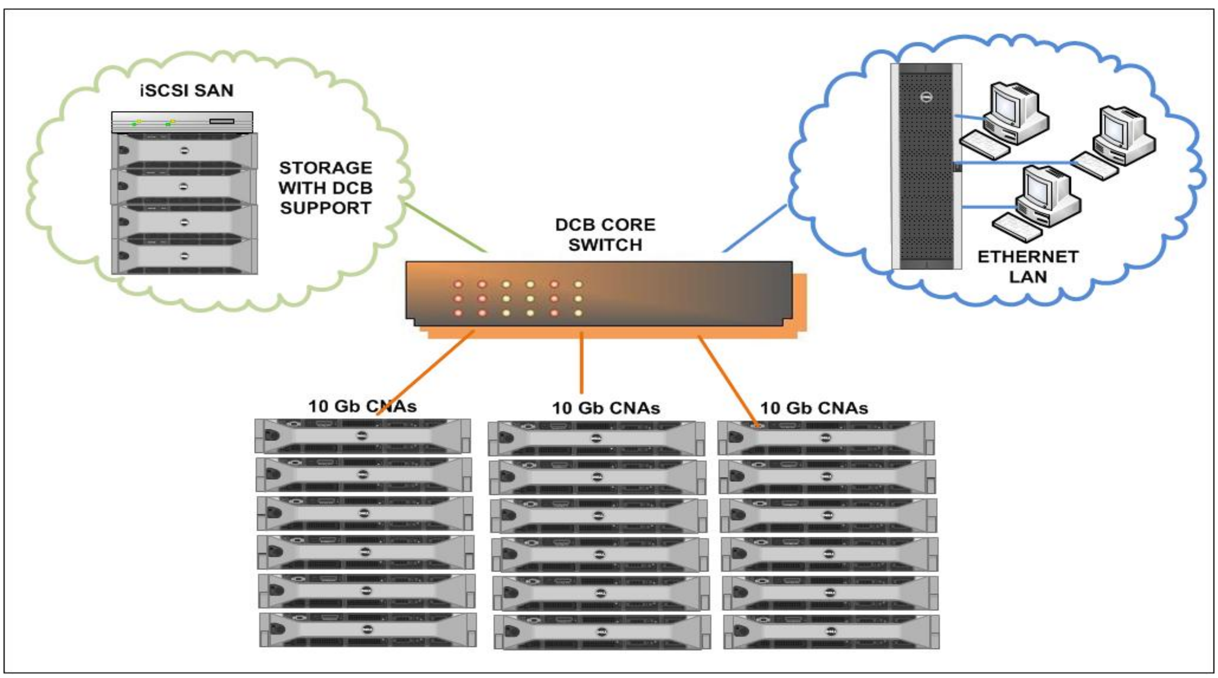

Figure 11 Data center of the future

5 Conclusion

Data Center Bridging is a new network standard that can bring to networks the same consolidation benefits that storage and servers have enjoyed in recent years—higher utilization rates, simpler management, and lower total cost of ownership. It is reliable, offers predictable performance, and can segregate and prioritize traffic types. Administrators can now implement standard Ethernet, Data Center Bridging, or a combination of both.

By taking a phased approach to consolidating data centers around Ethernet, customers can economically build out their Ethernet infrastructure over time while protecting previous infrastructure investments. When preparing to deploy a DCB-ready environment, planning for a network device infrastructure that provides full DCB support is paramount. Expected bandwidth requirements for both storage and LAN traffic to be converged must be taken into account. When calculating these requirements, be sure to provide for future growth 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.

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 standalone network, the lack of ETS in a converged environment will result in less than ideal results.

6 Glossary

iSCSI: Internet Small Computer System Interface. iSCSI is a protocol that uses TCP to transport SCSI commands for a storage network, enabling existing TCP/IP Infrastructure to be used as a SAN

FCoE: Fibre Channel over Ethernet

DCB: Data Center Bridging

DCBX: DCB eXchange protocol

ETS: Enhanced Transmission Selection

PFC: Priority-based Flow Control

LAN: Local Area Network

SAN: Storage Area Network

CNA: Converged Network Adapter

Ethernet: Ethernet is a physical and data link layer technology for LANs

MAC: Media Access Control address is a unique identifier assigned to network interfaces for communications on the physical network segment

CoS: Class of Service is a way of managing traffic in a network by grouping similar types of traffic (for example, e-mail, streaming video, voice, large document file transfer) together and treating each type as a class with its own level of service priority. Unlike Quality of Service (QoS) traffic management, Class of Service technologies do not guarantee a level of service in terms of bandwidth and delivery time; they offer a "best-effort."

QoS: Quality of Service (QoS) is a set of technologies for managing network traffic in a cost effective manner to enhance user experiences for home and enterprise environments. QoS technologies allow you to measure bandwidth, detect changing network conditions (such as congestion or availability of bandwidth), and prioritize or throttle traffic.

TLV: Type-Length-Value or Tagged-Length-Value

7 References

The following resources are referenced in this document.

- "iSCSI Design Considerations and Deployment Guide," VMware 2007, http://www.vmware.com/files/pdf/iSCSI_design_deploy.pdf

- "Creating a DCB Compliant EqualLogic iSCSI SAN with Mixed Traffic," Dell Inc. August, 2011, http://en.community.dell.com/techcenter/storage/w/wiki/creating-a-dcb-compliant-equallogic-iscsi-san-with-mixed-traffic.aspx

- "Deploying the Brocade 8000 Switch to Introduce Server I/O Consolidation into the Data Center," Brocade 2009 http://www.brocade.com/downloads/documents/technical_briefs/Brocade8000_GA-SG-161-01.pdf

- "Network Convergence in the Data Center Makes Sense," Emulex 2010 http://www.emulex.com/company/media-center/analyst-papers/print.html

- "Data Center Bridging Tutorial, "University of New Hampshire February 2009 http://www.iol.unh.edu/services/testing/dcb/training/DCB-Tutorial.pdf

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.