



ALCATEL-LUCENT ENTERPRISE AND QLOGIC PRESENT A VIRTUALIZED AND CONVERGED NETWORK FOR ENTERPRISE DATACENTERS

APPLICATION NOTE

TABLE OF CONTENTS

Introduction / 1

The case for Ethernet / 2

The components of a converged network / 3

Data Center Bridging (DCB) / 3

Storage convergence -FC/ FCoE / 4

Use cases / 5

Configuration guidelines / 7

Mandatory configuration / 8

Sample setup / 8

OmniSwitch Virtual Chassis Configuration steps / 9

QLogic 5800 FC Zoning Configuration / 11

Key benefits of virtualized and converged data centers / 11


INTRODUCTION

This document examines and presents the advancements in Ethernet standards that are enabling the deployment of modern IT solutions using Alcatel-Lucent Enterprise and QLogic technologies.

Enterprises are experiencing a multitude of changes due to industry trends such as datacenter consolidation, mobility, bring your own device (BYOD) and big data analytics. These key business challenges can only be effectively addressed by IT transformation, however, in many enterprises the IT infrastructure and delivery model has not caught up. IT continues to build services using discrete legacy elements, including separate local and storage area network (LAN and SAN), server, and application platforms. As a result, up to 70 percent of IT spending is consumed maintaining legacy infrastructure. To address current requirements and to be flexible for future business demands, IT must transform and provide virtualized, converged network environments, delivering a singular, unified infrastructure for networks, servers, storage and applications.

A virtualized network platform allows IT to view all resources from a higher logical level, enabling improved asset utilization, which in turn enhances operational performance. A converged network platform lets IT consolidate assets and reduce costs while improving flexibility and agility to deliver dynamic on-demand services.

Figure 1. Traditional LAN/ Discrete Architecture




THE CASE FOR ETHERNET

Ethernet is a proven technology and the dominant networking medium of choice, however, it is only a best-effort medium that does not guarantee data delivery. Although the introduction of Transmission Control Protocol/Internet Protocol (TCP/IP) enhanced Ethernet's robustness, the downside is that TCP/IP tends to introduce variable latency. As a result, Ethernet and TCP/IP have not proven suitable for all applications, such as SAN. SAN technologies, including Fibre Channel (FC) and Infiniband offer differentiated, low-latency services that require specialized network transport equipment and IT personnel to manage them. As application mobility across the datacenter has increased, so has the requirement for storage mobility, resulting in ever more expensive and complex architectures.

Ethernet speeds of 40GbE are being deployed today and 100GbE is in early adoption phase. Latency in Ethernet networks is measured in nanoseconds, and the latest IEEE standards deliver end-to-end lossless behavior that removes the traditional limitations associated with Ethernet. To maximize utilization and monetize investments, IT teams plan to use converged Ethernet to carry server, storage and application data on a single infrastructure, improving the efficiency and quality of delivered services. The combined benefits of lossless and low-latency Ethernet make a strong case for deploying a single converged fabric.

The network is the core transport layer, but advanced technologies need to be enabled and interoperable – starting from the server end-point to the storage array – to deliver an end-to-end converged solution. QLogic Converged Network Adapters (CNAs) support concurrent LAN (TCP/IP) and SAN (Fibre Channel over Ethernet [FCoE] and Internet Small Computer System Interface [iSCSI]) traffic over a shared 10GE link. This is achieved by QLogic Network Interface Card partitioning (NPAR) technology, which helps divide a single physical 10GE port into multiple Peripheral Component Interconnect (PCI) physical functions or partitions with flexible bandwidth capacity allocation. Each partition can support concurrent networking and storage protocols, enabling flexible application provisioning, and network functions such as quality of service (QoS) to be offloaded from the central processing unit (CPU). This improves input/output (I/O) performance, and resource consolidation and virtualization, while lowering the solution total cost of ownership (TCO).

Figure 2. Converged Ethernet



THE COMPONENTS OF A CONVERGED NETWORK

A series of standards based and advanced network features have been introduced to carry converged traffic flows. These combine to make Ethernet a lossless infrastructure that understands and adapts to the requirements of the underlying applications.

Data Center Bridging (DCB)

The IEEE 802.1 Datacenter Bridging standard (DCB) enables Ethernet lossless capabilities to reliably transport storage protocol standards. It allows traditional LAN as well as SAN traffic on the same physical network, reducing complexity and cost while providing a unified transport and management framework. DCB encompasses a suite of protocols that enable this convergence of the network fabric.


1. Priority-based Flow Control (PFC: 802.1Qbb) defines the process of configuring and pausing traffic on a per-priority application basis rather than the traditional per-port basis of Ethernet. It is designed to provision specific traffic classes for a lossless channel. As a result, sensitive applications such as SAN traffic are not affected by pauses in transmission of lower priority applications traversing the same port or network segment.
2. Enhanced Transmission Selection (ETS: 802.1Qaz) uses 802.1p to classify traffic into groups with similar service requirements and to configure bandwidth limits for those groups. This facilitates the profiling of sensitive applications that need to be prioritized over the lossless fabric.
3. Data Center Bridging Exchange Protocol (DCBX: 802.1Qaz) uses existing network Link Layer Discovery Protocol (LLDP) capabilities so neighboring devices that support DCB can communicate and negotiate the configuration of their DCB capabilities, including the priority requirements of higher-level applications such as FCoE or iSCSI. This ensures consistent configuration for DCB can be easily provisioned across the entire network.

Prior to the development of IEEE DCBX, various manufacturers defined and used the CEE DCBX v1.0 protocol. Most vendors are beginning to support IEEE DCBX, but the majority of the current install base still supports CEE DCBX. QLogic CNAs enable converged networking by supporting DCB enhancements and have been validated for interoperability with Alcatel-Lucent OmniSwitch™ platform. OmniSwitch platforms support both CEE DCBX and IEEE DCBX versions, auto-detecting a peer's DCBX version. The IEEE DCBX version is the default setting on OmniSwitch lossless Ethernet interfaces. When a CEE-enabled peer is detected, the OmniSwitch will automatically start using CEE DCBX. The DCBX version setting can be manually configured to either IEEE or CEE DCBX on an interface, as required. The OmniSwitch Alcatel-Lucent operating system (AOS) simplifies user network configuration by providing up to 128 user profiles to configure lossless port properties. DCB profiles DCB-1 to DCB-10 are based on the 802.1Q-REV/D1-5 Appendix I standard, but administrators can build custom profiles if specific requirements are not provided in the defaults.

Storage convergence –FC/ FCoE

The Fibre Channel Backbone-5 (T11 FC-BB-5) standard specifies how FCoE should operate over a lossless Ethernet network.

Figure 3. FC Mapping: Fibre Channel levels and sublevels over IEEE 802.3 layers




The OmniSwitch virtualizes the network infrastructure and can support the transport of both IP and lossless storage protocols, such as iSCSI or FCoE. This convergence simplifies the datacenter and ensures the longevity of existing IT investments. Servers can connect to any network element using a CNA that carries storage and data traffic (IP) simultaneously over a single interface. This converged network provides enhanced Ethernet capabilities with QoS support for traditional server-to-storage transport or newly ratified FC-BB-6 standard FCoE VN2VN transport, delivering reliable and flexible solutions that Enterprise datacenter deployments require.

The OmniSwitch offers FCoE Initialization Protocol (FIP) snooping technology, enabling seamless, end-to-end lossless transport and FCoE convergence technologies. Together, these facilitate the extension and interconnection of an FC SAN across an Ethernet infrastructure without having to purchase or manage additional, costly FC switching equipment.

In a traditional FC SAN, the connections between hosts and FC switches are point-to-point. The FC switch controls all communication between hosts and targets. In contrast, Ethernet networks accommodate point-to-multipoint connections and there can be any number of Ethernet switches between a host and an FC switch (FCF) in the SAN. FIP is the process an FCoE-capable node in the Ethernet network uses to login and establish a session with the SAN. FIP-enabled Ethernet switches monitor – or “snoop” – this login and session process. FIP snooping can monitor: the FCoE virtual LAN (VLAN) assigned to the host; the FCF discovery process; and the session fabric login/fabric logout (FLOGI/FLOGO) and keep-alive messages. Snooping of the FIP process allows the Ethernet switch to dynamically apply Access Control Lists (ACL) that effectively create secure point-to-point tunnels over the Ethernet fabric from an E_Node to the FCF in the SAN.

In addition, existing native FC Host Bus Adapters (HBA) can be directly connected to the OmniSwitch, allowing complete convergence and simplifying IT.

Figure 4. Converged Ethernet/ resilient and plug-and-play architecture



USE CASES

DCB and FIP technologies enable the OmniSwitch to be positioned in two different roles for FCoE transport:

1. As an FCoE transit switch - the OmniSwitch supports the FCoE technology used to tunnel FC frames encapsulated within Ethernet Media Access Control (MAC) frames. To provide the necessary FCoE transit switch functionality, the OmniSwitch uses FIP snooping and DCB. A transit switch is essentially a DCB-capable switch that bridges encapsulated FCoE traffic over the Ethernet fabric between FCoE nodes.
2. As an FCoE/FC gateway switch - the OmniSwitch serves as an FCoE forwarder to: connect FCoE nodes to FC switches; connect FC nodes to an FCoE forwarder; and connect native FC fabrics transparently across an FCoE converged network.

To provide the necessary FCoE/FC switching gateway functionality, the OmniSwitch supports the following operational modes:

- N_Port proxy operation aggregates FCoE Node (ENode) logins over a single OmniSwitch FC port that is connected to an FC switch.
- F_Port proxy operation connects FC nodes to an FCoE forwarder or another gateway switch through an FCoE network or via the same gateway switch.
- E_Port proxy operation provides a transparent point-to-point FC link between native E_Ports. This enables Inter-Switch Link (ISL) tunneling between FC fabrics over an FCoE network.

N_Port proxy is essentially a limited form of an FCF function that extends the FC SAN via FCoE. The OmniSwitch port configured as an N_Port proxy (NP_proxy) performs a login to the FC fabric. If this is successful, the N_Port is assigned a Fibre Channel ID (FCID). Once there is a virtual SAN (VSAN) to FCoE VLAN mapping, then the OmniSwitch gateway uses FIP snooping to monitor ENode traffic and provide them with their FCoE VLAN and FCF. As part of the discovery process, the ENode elects an FCF based on the received data. The N_Port function proxies subsequent ENode Fabric login connection and disconnection requests.


When there is more than one N_Port on an OmniSwitch in the same VSAN, auto load-balancing is performed based on the least number of logins serviced. It is applied when an FLOGI request is received from an ENode. Various load-balancing techniques, such as dynamic (default), dynamic-reorder, dynamic-ENode based and static mapping are supported.

The F_Port proxy function enables native FC HBAs to connect with the FC SAN over the FCoE lossless fabric via N_Port proxy. The port on which the HBA connects on the OmniSwitch gateway is referred to as an F_Port. The FCoE network port mapped to the FCoE VLAN is referred to as a VN_Port. The port connected to the FC gateway is referred to as a VF_Port. The VN_Port transports the HBA traffic over the FCoE network to FCoE gateways. As such, the OmniSwitch gateway provides connectivity for the ENodes across the FCoE lossless fabric. When traffic reaches a gateway switch that is directly connected to the FC SAN, then N_Port proxy provides the virtual connection into the SAN. The F_Port proxy function may reside on the same node performing the N_Port proxy function.

The F_Port proxy operation begins by discovering FCF or N_Port proxies reachable via the configured FCoE VLAN. OmniSwitch gateways listen for advertisements from other FCF or N_Port proxies. The F_Port proxy gateway selects an FCF based on its priority and for those FCF whose 'Available for Login' flag is set. If there are multiple FCF with the same priority, then the F_Port proxy gateway will choose the FCF discovered first.

If both the FC switch and the HBA are directly connected to the same OmniSwitch in the same VSAN, then the HBA login is performed directly to the FC switch without involving the OmniSwitch; its FCF function is not required.

Figure 5. Converged Ethernet - FC/FCoE with N-Port/F-Port proxy



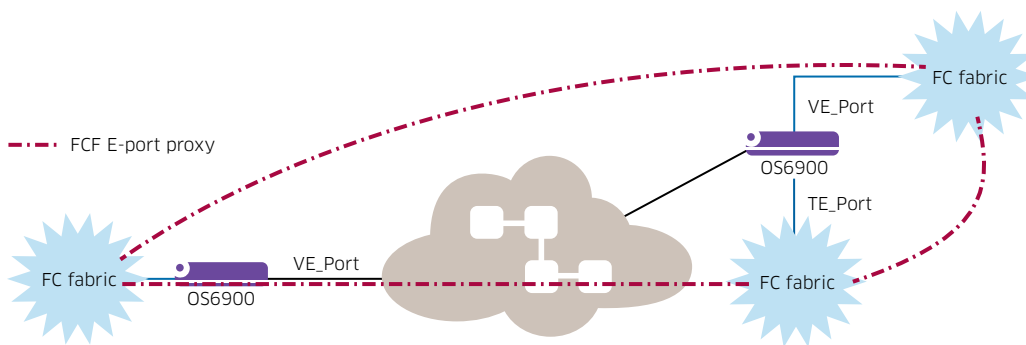
The E_Port proxy function allows two or more FC switches to transparently connect across a lossless FCoE fabric. The OmniSwitch port to which the FC switch E-Port connects is called a tunnel expansion port (TE_Port). The port facing the lossless FCoE network is called a virtual expansion port (VE_Port). This results in two possible tunnel modes:

- TE-TE tunnel, directly between two E_Ports connected to the same OmniSwitch gateway.
- VE-VE tunnels E_Ports across an FCoE lossless fabric.

Each VE_Port is a tunnel end-point mapped to the TE_Port via an FCoE VLAN. From the E_Port perspective, they are directly connected and the TE-TE tunnel is transparent.

The VE_Ports discover FCF that are available for login, achieved either by listening for multicast FCF advertisements or by specific unicast discovery requests. The VE_Ports are then ready to accept FIP Exchange Link Parameters (ELP) messages or relay FC ELP requests received from one of the E_Ports connected to the FC switches.

Figure 6. Converged Ethernet - FC/FCoE resiliency and scale



CONFIGURATION GUIDELINES

OmniSwitches must be configured as follows to support FCoE data:

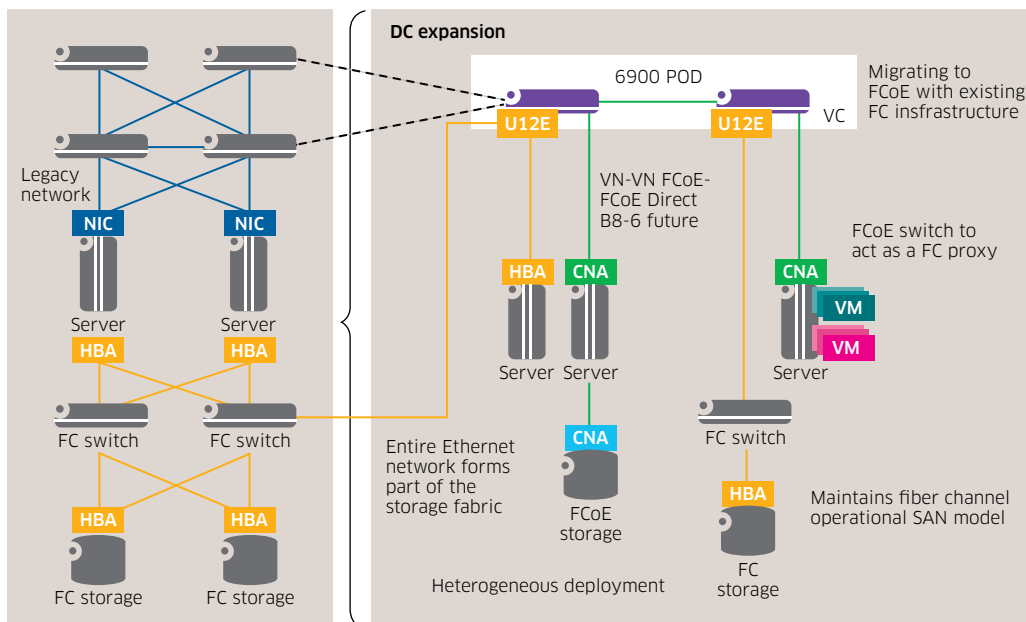
- Apply datacenter licenses to all switches.
- Configure QoS to trust all ports and default classification mode based on ingress 802.1p value.
- Configure a lossless profile on all participating ports. The participating ports are referred to as:
 - Edge ports, which connect to servers or ENodes.
 - NP ports, which connect directly to FC Forwarders (FCFs).
 - FCF, ENode switch interconnection ports.
- DCB configurations can be applied to each port of a switch by applying a DCB profile on the port. Creating a custom profile is recommended, with desired FCoE priority configured as lossless.
- Configuring switches according to desired network topology is recommended, with the settings propagated to the servers and FCFs via DCB protocol.
- Configure the LLDP protocol on participating edge-ports to signal application type-length-value (TLV) to the peer device (ENode). The application TLV specifies the FCoE EtherType and the matching FCoE priority.
- Configure FIP-snooping on the switch.

Mandatory configuration

- Enable FIP-snooping globally on the switch.
- Configure a VLAN for FCoE traffic.
- Configure an FCoE priority. (Default is 3)
- Configure FCoE priority protection (optional, this command tells the switch to either drop or remark non-FCoE traffic ingressing on FCoE ports at the same priority).
- Create a tagged VLAN port association for all participating ports.
- Configure a default VLAN on the path between the FCF and the ENode. This is useful for FIP VLAN discovery request and response.
- Configure port roles for all participating ports. The port-roles are based on the topology and could be one of the following:
 - a) Edge – directly connected to ENode. Maximum security and maximum number of ACLs.
 - b) Enode-only.
 - c) FCF-only.
 - d) Mixed – In meshed/highly redundant topologies where FCFs and ENodes can reside on the same side of a link.
 - e) Trusted – used to save ACLs.

Sample setup

Figure 7. Sample set-up



OmniSwitch 6900 (with OS-XNI-U12E Universal Card)
 QLE8300/8200 CNA
 QLE2562 FC HBA
 QLE5800 FC-Switch
 NetApp FAS2030 FC Storage

OmniSwitch Virtual Chassis Configuration steps

On Switch#1:

```
! Virtual Chassis Manager:
virtual-chassis chassis-id 1 configured-chassis-id 1
virtual-chassis chassis-id 1 vf-link 0 create
virtual-chassis chassis-id 1 vf-link 0 member-port 1/1/1
```

On Switch#2:

```
! Virtual Chassis Manager:
virtual-chassis chassis-id 2 configured-chassis-id 2
virtual-chassis chassis-id 2 vf-link 0 create
virtual-chassis chassis-id 2 vf-link 0 member-port 2/1/1
```

EMP Configuration for RDP:

```
! IP:
ip interface master emp address 10.255.92.20 mask 255.255.255.0
```

VC_of_2- > show ip emp-interfaces

Total 3 interfaces

Name	IP Address	Subnet Mask	Status Forward	Device
EMP-CMMA-CHAS1	10.255.92.6	255.255.255.0	UP	NO EMP
EMP-CMMA-CHAS2	10.255.92.7	255.255.255.0	UP	NO EMP
EMP-VC	10.255.92.20	255.255.255.0	UP	NO EMP

License Information:

Advanced license is required for Virtual Chassis setup and datacenter license is required for Fibre Channel lossless setup.

VC_of_2- > show license-info

VC	device	License	Time (Days)	
			Type	Remaining
1	0	Advanced	PERM	NA
1	0	Data-Center	PERM	NA
2	0	Advanced	PERM	NA
2	0	Data-Center	PERM	NA

VC_of_2-> show virtual-chassis topology

Local Chassis: 1

Oper	Config	Oper				
Chas	Role	Status	Chas ID	Pri	Group	MAC-Address
1	Master	Running	1	100	2	e8:e7:32:36:1e:f5
2	Slave	Running	2	100	2	e8:e7:32:07:96:d5

FCoE application priority settings via LLDP

! LLDP:

lldp nearest-bridge port 1/1/2 tlv application enable

lldp port 1/1/2 tlv application fcoe priority 4

Enabling Lossless Configuration:

qos port 1/1/2 trusted default classification 802.1p

qos qsi port 1/1/2 dcb dcbx version cee [also ieee, default is auto]

qos qsi port 1/1/2 dcb dcbx ets willing no

qos qsi port 1/1/2 dcb dcbx pfc willing no

qos qsi port 1/1/2 qsp dcb 9

FCoE edge port settings:

! FCOE:

fcoe fip-snooping admin-state enable

fcoe port 1/1/2 role edge

fcoe vlan 46 admin-state enable

fcoe vlan 46 name "FCoE VLAN"

vlan 46 members port 1/1/2 tagged

fibre channel & VSAN configuration:

fibre-channel port 2/2/1 mode NP

fcoe vsan-map vsan 46 vlan 46

fibre-channel vsan 46 members port 2/2/1

Steps to create a Fibre channel port

qos qsp dcb 18 import qsp dcb 9 802.3x-pause [create a custom profile]

qos qsi port 2/2/1 dcb dcbx pfc willing no

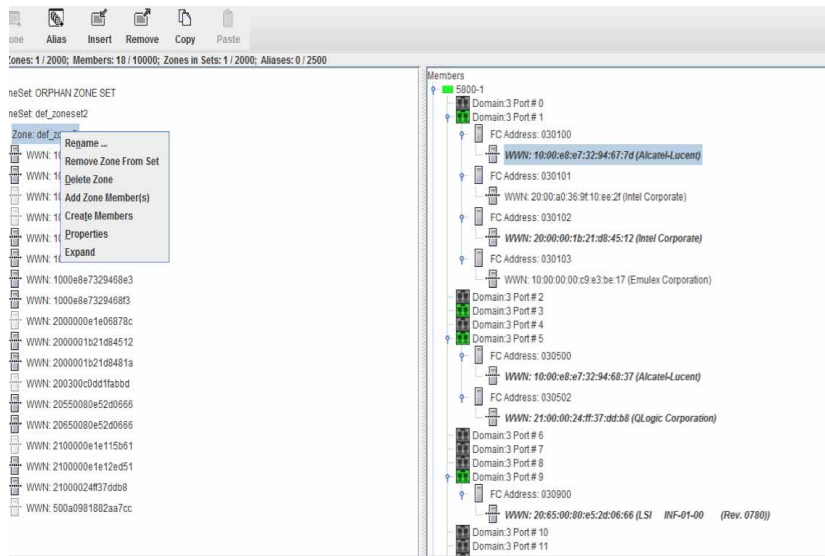
qos qsi port 2/2/1 dcb dcbx pfc tlv disable

qos qsi port 2/2/1 qsp dcb 18

QLogic 5800 FC Zoning Configuration

Zoning a fabric enables the ports and devices of the fabric to be divided into zones for more efficient and secure communication among functionally grouped nodes. Zoning divides the fabric for the purpose of controlling discovery and inbound traffic. A zone is a named group of ports or devices. Members of the same zone can communicate with each other and transmit outside the zone, but cannot receive inbound traffic from outside the zone. Zoning is hardware-enforced; refer to the QLogic 5800 series user manual for additional configuration details and applicable restrictions.

Figure 8. Editing QLogic 5800 FC zoning configuration



KEY BENEFITS OF VIRTUALIZED AND CONVERGED DATA CENTERS

Reduced capital and operational expenditures:

- Enables hardware consolidation by adopting an end-to-end Converged Ethernet network for LAN/SAN transport.
- Maximizes network utilization while providing end-to-end lossless channel for SAN traffic; simplifying IT.
- CNA NPAR technology allows for higher NIC utilization, in turn reducing adapter attachment rate and lowering cabling costs.

Simplified deployment and management:

- Provides concurrent support for Ethernet, FCoE, and iSCSI, easing deployment and minimizing maintenance disruption.

Improved ability to scale virtualized data center:

- The virtualized platform enabled with Virtual Chassis technology and a multipath MESH facilitates deploying scalable, geographically dispersed data centers.
- Enables high I/O performance in virtualized environments.

