FIBRE CHANNEL (FC)/ FC OVER ETHERNET (FCOE) BENEFITING DATA CENTER CONVERGED NETWORK

APPLICATION NOTE

Alcatel·Lucent Enterprise

TABLE OF CONTENTS

INTRODUCTION

FC is an expensive technology that typically requires purpose-built networks with a separate infrastructure and networking staff. While data networks use Ethernet technology, storage networks use FC and iSCSI in addition to Ethernet. Ethernet networks are a viable alternative for storage transport because they are more economical and scalable. In addition, Ethernet speeds are increasing faster, to 40 Gb/s and 100 Gb/s, compared to 8 Gb/s and 16 Gb/s for FC.

Because of their limited performance, smaller market size and high support requirements, FC networks require higher levels of both capital and operational expenditure (CAPEX and OPEX). Nevertheless, organizations selected FC for their storage networks because it can ensure higher levels of performance.

FC has lower bandwidth requirements and limits the number of nodes permitted on each FC network. One of the drivers for the shift to 10 Gb/s Ethernet (10GE) on the server side was the rapid adoption of FCoE in data centers. The switch to 10GE was prompted by the move to virtualization as well as the need to support the bandwidth of the Ethernet network interface controller (NIC) and the FC host bus adapter (HBA) (1, 2, 4 or 8 Gb/s) on a single NIC card.

Servers and storage are not always co-located, nor do they share a common network technology. The trend towards server virtualization means that servers are typically on Ethernet networks, while storage is traditionally on FC networks. These two technologies are different and cannot be connected easily. Ethernet is traditionally a connectionless, best-effort technology, whereas FC is connection-oriented, with a dedicated bandwidth connection.

A number of standards activities map these two technologies. Data Center Bridging (DCB) is a set of standards providing lossless Ethernet capability on Ethernet links to mimic FC network characteristics. It was developed so FC frames could travel over Ethernet seamlessly, with the data frames defined as FCoE. The FCoE protocol specification replaces the FC0 and FC1 layers of the FC stack with Ethernet.

Converged Network Adaptor (CNA) technology combines the functionality of an FC HBA with an Ethernet NIC. If – as seems to be the case – FCoE and the associated Data Center Ethernet (DCE) are here to stay, then it’s only a matter of time before the combined traffic of TCP/IP and FCoE on servers exceed the capacity of 10GE and forces a move to 40GE.

As the transition from FC to FCoE occurs in an effort to reduce costs and simplify network design, enterprises need to select a flexible switching solution that can handle all technologies – FCoE, FC and Ethernet – as needed.

APPLICATION

FCoE is an evolutionary step towards getting Ethernet networks to communicate with the installed FC base. FCoE needs to provide seamless connectivity between the server CNA and the devices attached to the storage area network (SAN), usually storage. At first glance, FCoE simply tunnels a full FC frame onto Ethernet. But this is only a small portion of what is needed. As mentioned in the Introduction, FC networks are connection-oriented and have dedicated bandwidth allocation for data movement, whereas Ethernet is a connectionless, best-effort media. Data center bridging (DCB) consists of Priority-based Flow Control (PFC), Enhanced Transmission Control (ETS), and Congestion Notification (CN). PFC (IEEE 802.1Qbb) provides a link-level flow-control mechanism that can be controlled independently for each frame priority. The goal of this mechanism is to ensure zero loss under congestion in DCB networks. ETS (IEEE 802.1Qbb) provides a common management framework to assign bandwidth to frame priorities. CN (IEEE 8021Qau) provides end-to-end congestion management for protocols that are capable of transmission rate limiting to avoid frame loss. CN is expected to benefit protocols such as TCP that have native congestion management, as it reacts to congestion more quickly.

At least two devices are required to implement FCoE and unified fabric: a lossless Ethernet switch with FC interfaces that can receive FC frames and forward FCoE frames to an Ethernet network, and a multifunction server adapter that supports the local area network (LAN) and the SAN. The switch is very important and must do more than just translate frames.

Table 1. Common terms used in FC and FCoE technology

FIP	FC initiation protocol
FLOGI	Fabric login process
FLOGO	Fabric logout process
HBA	FC host bus adapter
CNA	Converged NIC adapter
FC forwarder (FCF)	An FC switching element able to forward FCoE frames across one or more FCF-MAC addresses and optionally includes one or more lossless Ethernet bridging elements and/or an FC Fabric interface
ENode	An FCoE node/host able to transmit FC frames using one or more ENode MAC address
FC_Port	A port capable of transmitting and receiving FC frames
VN_Port	An instance of FC that operates as an N_Port on the ENode dynamically instantiated on the successful completion of an FIP FLOGI or FIP N_Port ID Virtualization (NPIV) FDISC Exchange
VF_Port	An instance of FC that communicates with one or more VN_Ports on the FCF dynamically instantiated on the successful completion of a FIP FLOGI Exchange
VE_Port	An instance of FC that communicates with another VE_Port on the FCF that is dynamically instantiated on successful completion of a FIP ELP Exchange

WHAT IS FIP SNOOPING AND WHY IS IT NEEDED?

The principle function of FIP snooping is to provide secure, FC-like point-to-point tunnels over the Ethernet fabric from ENode to FCF. In an FC network, the connections from hosts to FC switches are always point–to-point. This allows the switch to control all communication to and from the hosts. FCoE provides much more flexibility, as there can be any number of Ethernet switches in-between the hosts and the FCF. FIPS enables Ethernet switches to emulate an end-to-end secure channel that an FC fabric provisions natively. FIPS monitors FIP virtual LAN (VLAN) discovery, FCF discovery, FLOGI/FLOGO and keep-alive messages to provision and maintain secure tunnels. Access control lists (ACLs) are dynamically created between the host ENode and FCF to provide point-to-point secure communication tunnels.

Use cases

The converged data and storage networking trend requires switches to have NPIV Gateway functionality to enable communication between disparate network types and to simplify the operational efficiency.

An NPIV gateway enables a SAN to expand via FCoE without needing to configure and manage a new FCF. A NPIV gateway emulates an FCF to the Ethernet network while simultaneously emulating an FC node (NPIV-enabled node) to the FC network. FCoE hosts think they are logging into an FCF, and the NPIV gateway logs in each FCoE device into the FC fabric by treating them as locally attached individual virtual node ports (VN_Ports). The NPIV gateway facilitates this by relaying the fabric login and login accept messages between the FCoE devices and FC switches. An NPIV gateway can also add value by allowing multiple links with an FC switch and load balancing logins from FCoE device across those links. An NPIV gateway does not switch frames between FC interfaces or from one FCoE device to another (within the same virtual SAN [VSAN])

An NPIV gateway can be used for standalone systems or with a virtual chassis. It can support a single FC interface connecting to a single SAN or to multiple SANs. It can directly attach to an FC switch or an FC node.

Here are some examples. The primary focus of an NPIV gateway is to switch frames to/from Ethernet to FC media. An NPIV gateway emulates an FCF to the Ethernet network while at the same time it emulates a NPIV enabled node to at attached FC Switch.

Figure 1. NPIV gateway connected to Ethernet and FC networks

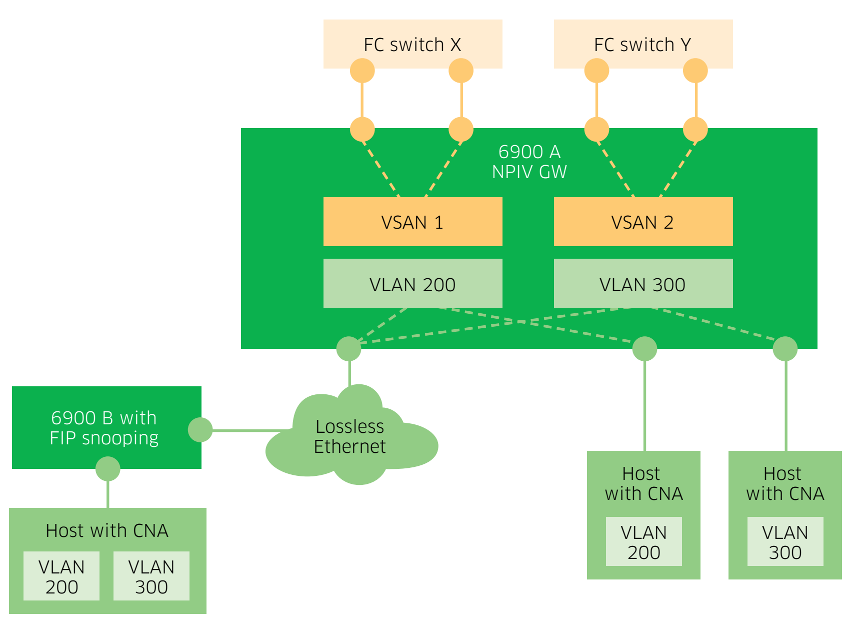

Figure 1 shows a standalone switch operating as a NPIV gateway connecting the Ethernet and FC networks. VSAN 1 is mapped to VLAN 200, and VSAN 2 is mapped to VLAN 300. The switch acts a proxy for all the Ethernet devices to the fiber channel switch. The switch can support up to one SAN per interface

REVERSE NPIV

While the purpose of an NPIV gateway is to extend an FC SAN by using FCoE, the primary function of a Reverse NPIV gateway is to switch FC frames from a host attached to a switch across an FCoE LAN to a standard NPIV gateway or FCF.

A Reverse NPIV gateway emulates an FC switch to a FC node. FC nodes think they are logged into an FC fabric that is directly connected.

The Reverse NPIV gateway alters and forwards the login requests from the FC node across an FCoE LAN towards the target gateway or FCF. Accept messages are relayed back to the originating host. Reverse NPIV gateways do not switch frames between FC interfaces, except for a virtual chassis, as the virtual chassis interconnect links represent an FCoE LAN.

Figure 2. Reverse NPIV

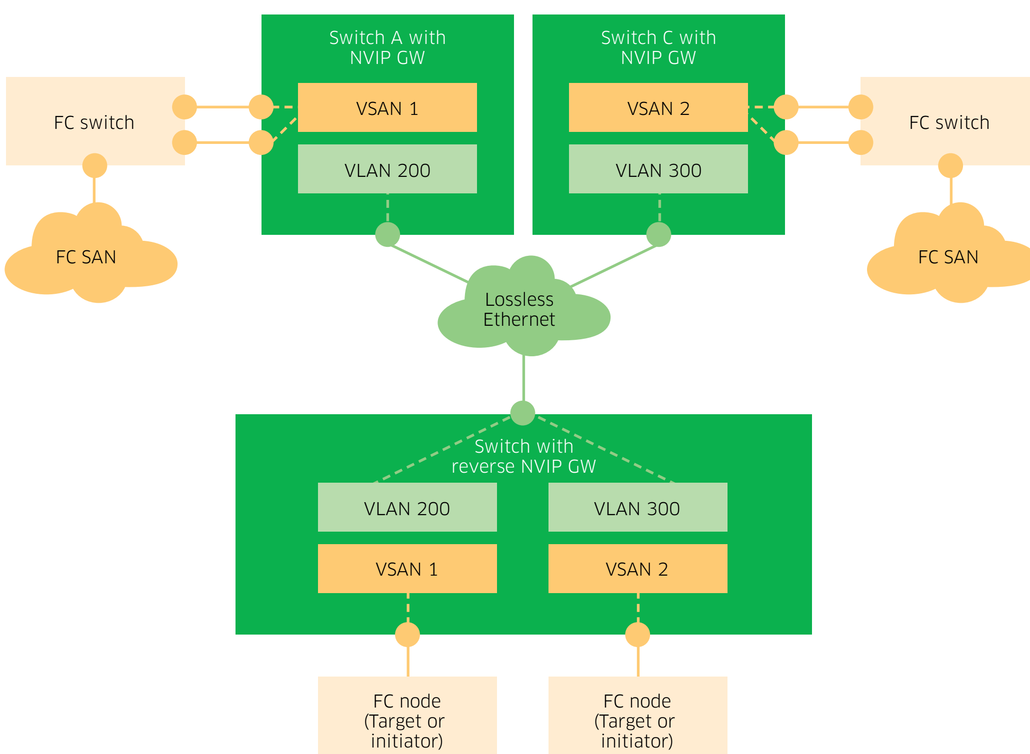

E2E_PORT TUNNELING

The purpose of an E_port tunnel is to extend a FC SAN by using FCoE to transparently connect two FC switches, or an FC switch to an FCF with FCoE interfaces.

Figure 3. E2E_Port tunneling

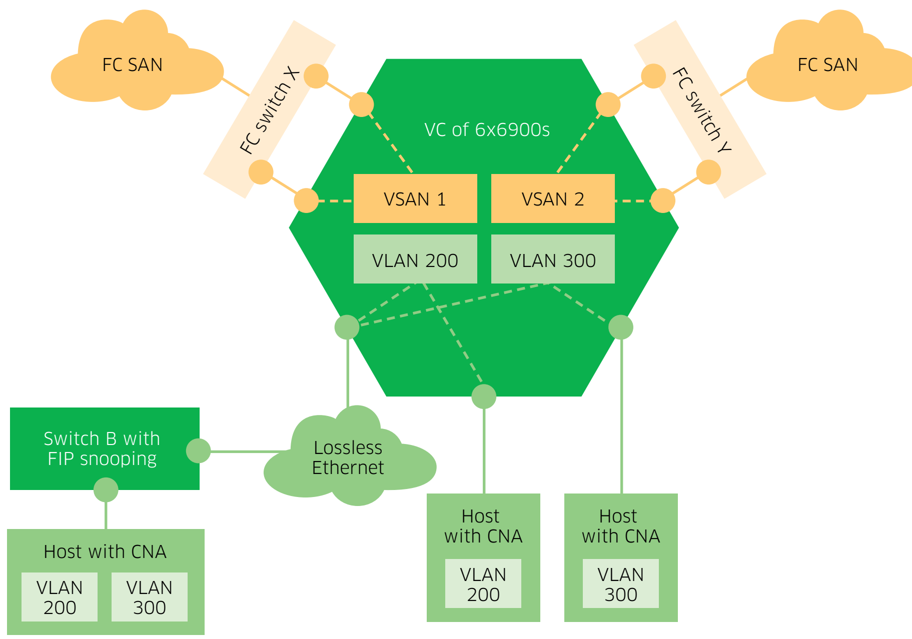

An Alcatel-Lucent OmniSwitch™ 6900 with an FC interface performing the tunnel operation would be termed an initiator; the remote endpoint (either FCF or E_port tunnel) is referred to as a target.

- An E_port tunnel emulates a point-to-point link between two FCFs/switches.
- FC forwarding switches think they are directly connected.
- The E_port tunnel transforms FC protocol messages into FCoE-compliant messages and vice-versa, so two FC forwarders can establish a connection.
- Data frames are passed without modification beyond the L1 header conversion.
- The E_port tunnel initiators must monitor the Ethernet links between the endpoints and maintain them through keep-alives (as described in the FC-BB-5/6 standard)
- An E_port tunnel initiator does not switch frames between FC interfaces, except in the case of a virtual chassis, as the VFL links represent an FCoE LAN.
- One E_port tunnel only connects two endpoints. Only two available endpoints can be connected in an FCoE VLAN.
- Switches in the middle of an FCoE LAN should not add a VSAN interface to a VLAN that is used for an E_port tunnel.

Figure 4. E_Port Tunneling to FCF

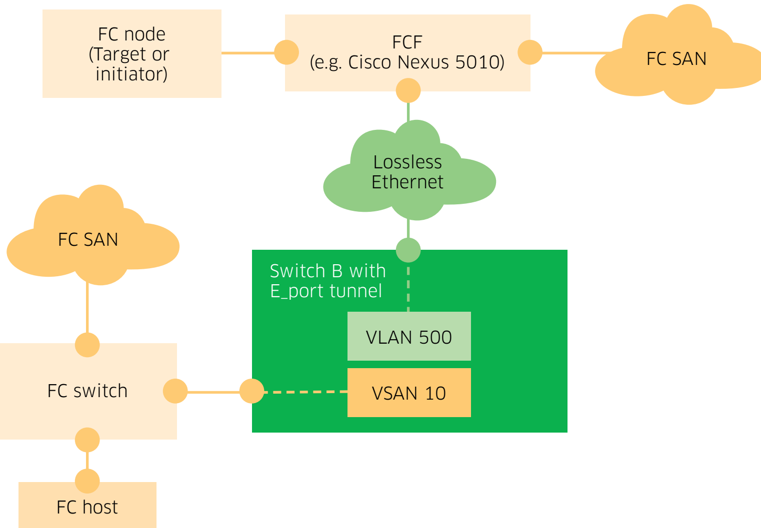

CONVERGED DATA AND STORAGE

There is a clear advantage to converging Ethernet and FC networks to reduce cost and simplify the network while optimizing both equipment and personnel. Ethernet is a more economical solution that scales to higher bandwidths. It allows common hardware such as blade servers and network devices to be used, as well as software in the form of virtualization.

FC is a natural progression towards the deployment of the Intelligent Fabric.

THE ALCATEL-LUCENT APPROACH

Alcatel-Lucent OmniSwitch 6900 (OS6900) Stackable LAN Switches are compact, high-density 10GE and 40GE platforms designed for the most demanding networks – such as those found in SANs or data centers.

The OS6900 can be used to create an Alcatel-Lucent advanced data center switching fabric, which employs “Mesh” network technology that facilitates virtualization. The Mesh is formed by connecting smaller “Pods,” consisting of several directly connected servers, to each other and to core switches. It can scale linearly from several hundred to more than 14,000 server-facing ports, as the Mesh provides resilient and low-latency any-to-any connectivity. Moreover, with 40GE (and, coming soon, 100GE) connections, along with Shortest Path Bridging (SPB), Virtual Chassis, Loss-less Ethernet, and thousands of 10 GigE server ports, the Alcatel-Lucent Mesh architecture ensures the quality delivery of real-time applications for the future, while reducing costs and operational complexity.

Their modular approach of the OS6900 also includes a daughter card with 12 ports that can be configured for Ethernet or FC on a per-port basis. Enterprise network planners can use the OS6900 to support FC, Ethernet or FCoE applications today and manage the transition to all-Ethernet in the future.

Alcatel·Lucent
Enterprise