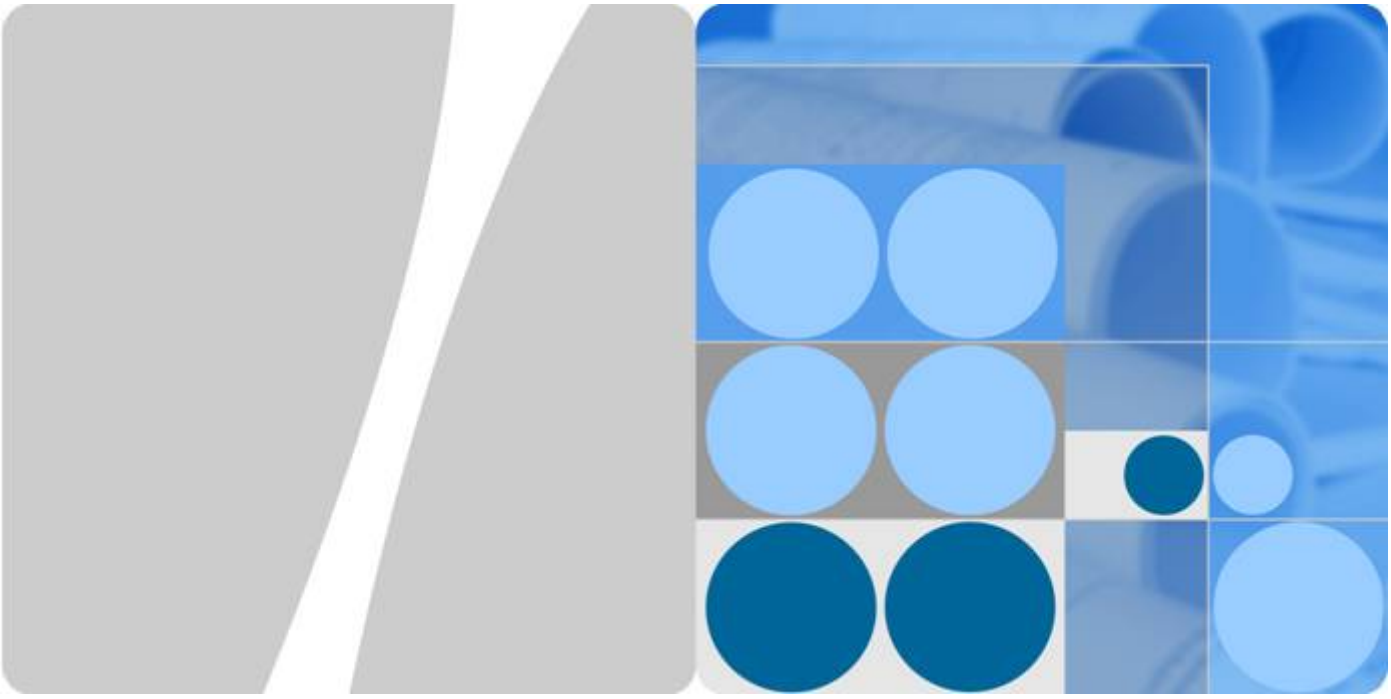

CloudEngine 12800 Series Switches

FCoE and DCB Technology White Paper

Issue 03

Date 2016-01-15

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

Contents

1 Introduction to FCoE

This section describes the definition and functions of FCoE.

Data Center Network Convergence Trend

As shown in **Figure 1-1**, the local area network (LAN) and storage area network (SAN) of a traditional data center are deployed and maintained independently. The LAN transmits services between servers and between servers and clients, and the SAN transmits services between servers and storage devices.

As data centers develop rapidly and increasing servers are deployed, independent deployment of LANs and SANs results in the following problems:

- Complex network: Service deployment is inflexible, network expansion is difficult, and network maintenance and management costs are high.
- Low energy efficiency: Each server is configured with at least 4 to 6 network adapters, including network interface cards (NICs) connected to LANs and host bus adapters (HBAs) connected to SANs. Such settings increase power consumption and cooling costs.

Figure 1-1 Before and after data center network convergence

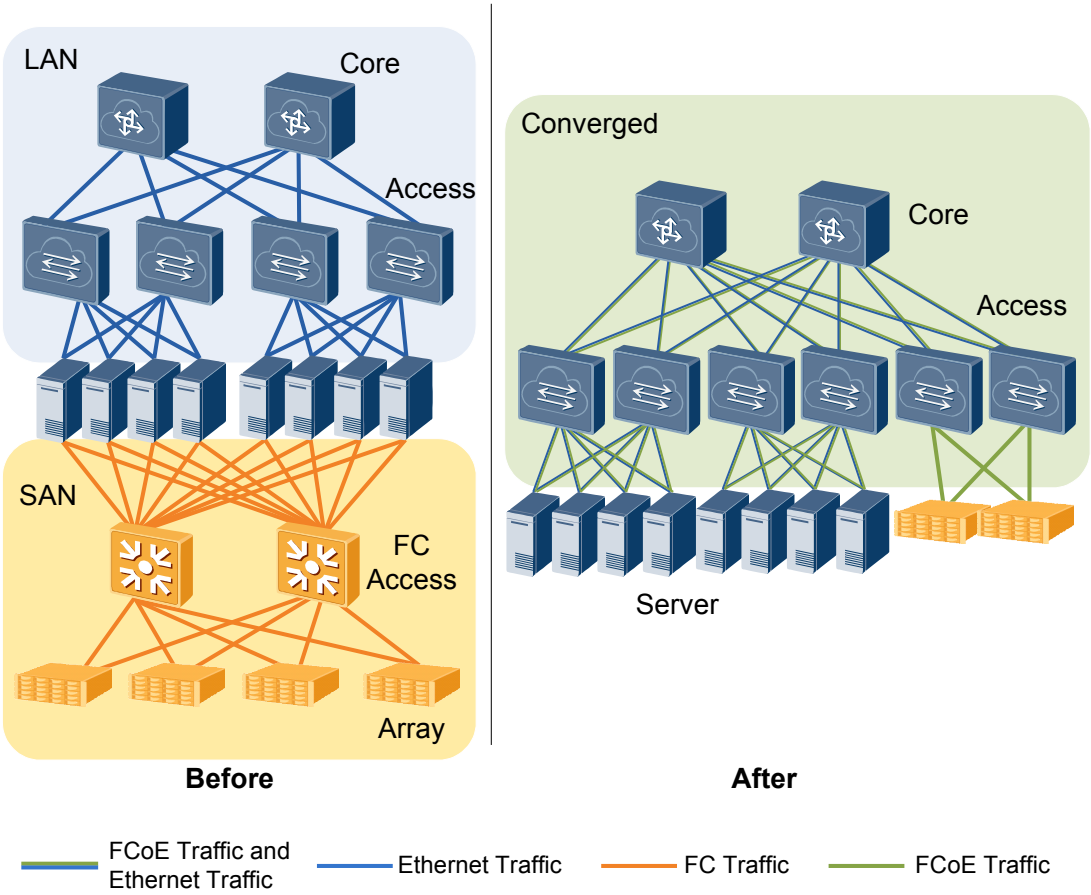

After the LAN and SAN are converged, the SAN and Ethernet LAN share the same integrated network infrastructure, simplifying network infrastructure.

Fiber Channel over Ethernet (FCoE) is used for network convergence.

Definition

FCoE is a network convergence technology defined by the American National Standards Institute (ANSI), and is also an I/O consolidation solution based on the FC protocol.

Purpose

After SAN and LAN networks are converged, the following issues occur:

- FC traffic cannot be forwarded over Ethernet.
- Ethernet cannot ensure lossless forwarding as the FC network.

FCoE and Data Center Bridging (DCB) address these issues:

- FCoE encapsulates FC frames into Ethernet frames so that LANs and SANs share network resources. With FCoE technology, LAN and SAN networks can be converged.
- DCB builds a lossless Ethernet network on a data center network. This technology enables traditional Ethernet to implement congestion control as on the FC SAN and provides QoS guarantee for FCoE services.

Benefits

FCoE brings in the following benefits:

- Reduces Total Cost of Ownership (TCO): FCoE allows LANs and SANs to share network resources. It integrates and fully uses distributed resources, reduces investments on SAN network infrastructure, simplifies network topology, and reduces network management and maintenance costs. Servers use converged network adapters (CNAs), which reduce electricity and cooling costs in data centers.
- Saves investment: FCoE seamlessly integrates existing Ethernet and FC infrastructure on data center networks to maximize the return on investment on FC SAN infrastructure, including various FC SAN tools, and established FC SAN facilities and management architecture.
- Enhances service flexibility: FCoE provides capabilities for all servers to access storage devices and allows virtual machine (VM) migrations. This implementation improves system flexibility and availability.

2 FCoE Principles

About This Chapter

This section describes the implementation of FCoE.

2.1 Basic Concepts of FCoE

As shown in **Figure 2-1**, FCoE involves the following entities: **ENode**, **FCF**, **FIP**, **NPV**, **FSB**, **Fabric**, **FCoE Virtual Link**, **Interface Role**, **FCoE VLAN**, **Zone**, **WWN**, **FC_ID**, **Domain ID**, and **FC-MAP**.

Figure 2-1 FCoE networking

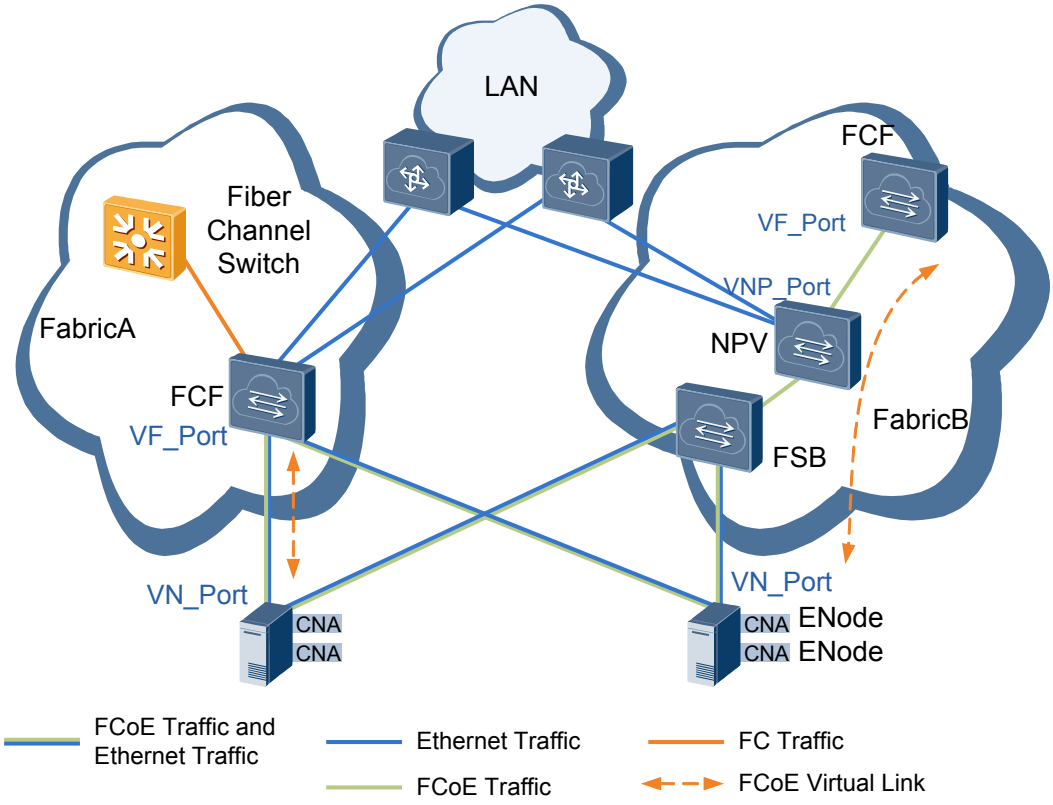

- ENode

  An ENode is a converged network adapter (CNA) that supports Ethernet and FC. In **Figure 2-2**, a traditional server has two network adapters installed: network interface card (NIC) connected to a LAN and a host bus adapter (HBA) connected to a SAN. The CNA provides both NIC and HBA functions. It can forward Ethernet data, process FCoE packets, and encapsulate or decapsulate FCoE packets.

Figure 2-2 Difference between a traditional server and FCoE server

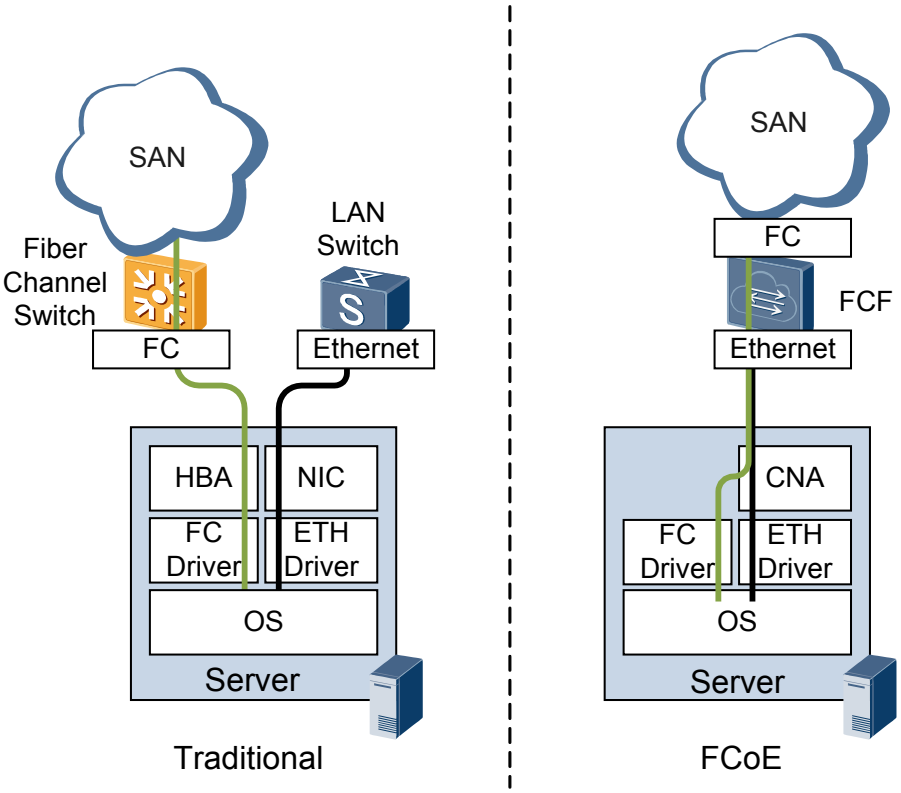

- FCF

  An FCoE forwarder (FCF) is a switch supporting both FCoE and FC. An FCF can forward FCoE packets, and is used to connect the SAN and LAN.

- FIP

  The FCoE Initialization Protocol (FIP) is a Layer 2 protocol that discovers FC terminals on an FCoE network, implements fabric login, and establishes FCoE virtual links. An ENode can log in to the fabric using FIP to communicate with the target FC device. FIP can also maintain FCoE virtual links.

  For details about FIP, see **2.3 FIP Protocol**.

- NPV

  An NPort Virtualization (NPV) switch is located at the edge of a fabric and between an ENode and an FCF. The NPV switch forwards traffic from the ENode to the FCF.

- FSB

  A FIP Snooping Bridge (FSB) is a switch running FIP snooping. The FSB itself does not support FC. FIP snooping enables the FSB to obtain FCoE virtual link information by listening on FIP packets. This function is used to control FCoE virtual link setup and prevent malicious attacks.

- Fabric

  A fabric is the network topology where network nodes are connected through one or more switches.

- FCoE virtual link

  An FCoE virtual link is a point-to-point logical link between FCoE devices, for example, between an ENode and FCF. The connection between an ENode and FCF is not point-to-point when the ENode and FCF are connected through a lossless Ethernet network. The FCoE virtual link is used to solve this problem.

- Interface role

  Figure 2-3 Interface roles

  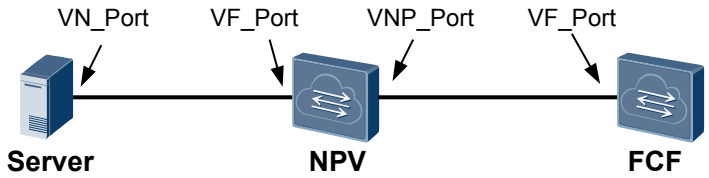

  On the traditional FC network, FC devices are connected through FC interfaces. FC interfaces are classified into node ports (N_Ports) and fabric ports (F_Ports):

  - N_Port: FC device interface that connects to an FC switch. An FC device can be a server or storage device.
  - F_Port: FC switch interface that connects to an FC device and provides fabric access services for the FC device.
  - N_Port Proxy (NP_Port): NPV switch's interface that connects to the FCF.

  FCoE inherits the interface roles of FC. On an FCoE virtual link between an ENode and an FCF, the ENode interface is a VN_Port and the FCF interface is a VF_Port.

- FCoE VLAN

  As defined by FC-BB-5, FCoE packets are forwarded in specified VLANs. In the FC protocol stack, one FC device supports multiple virtual storage area networks (VSANs), which are similar to Ethernet VLANs. FC traffic in different VSANs is identified by FCoE VLANs during FCoE encapsulation. By doing this, FCoE packets carry only FCoE VLAN information and do not need to carry VSAN information, because VSANs are differentiated by FCoE VLANs.

  An FCoE virtual link corresponds to one FCoE VLAN. An FCoE VLAN carries only FCoE traffic and does not carry any Ethernet traffic such as IP traffic.

- Zone

  You can configure zones and add different N_Ports to the zones based on the destination to isolate N_Ports in different zones. This setting achieves access control.

  A zone set is a collection of zones, and a zone is a collection of zone members. Each zone member is an N_Port.

- WWN

  A World Wide Name (WWN) identifies an entity on the fabric and SAN. A World Wide Node Name (WWNN) identifies an ENode and the World Wide Port Name (WWPN) identifies an interface. Each entity on the SAN has been allocated with a WWN before delivery.

- FC_ID

  The FC_ID is called FC address. On a SAN, FC protocols use FC_IDs to access entities. An FC_ID can identify an N_Port on an ENode uniquely.

- Domain_ID

  On a SAN, a domain ID identifies an FC switch uniquely. During packet transmission, FC switches use domain IDs to route and forward packets.

- FC-MAP

  An FCoE frame uses a locally unique MAC address (unique only within the local Ethernet subnet). The locally unique MAC address is dynamically assigned to an ENode

by an FCF or specified on an ENode and acknowledged by an FCF. The locally unique MAC address that is dynamically assigned to an ENode by an FCF is called a fabric-provided MAC address (FPMA). An FPMA has an FC_ID and a 24-bit FCoE MAC Address Prefix (FC-MAP).

2.2 FCoE Encapsulation

FCoE encapsulates FC frames into Ethernet frames so that FC traffic can be transmitted on an Ethernet network. From the FC perspective, FCoE is a different way of transmitting FC traffic. From the Ethernet perspective, FCoE is just another upper layer protocol to carry Ethernet frames.

FCoE Protocol Stack

As shown in **Figure 2-4**, the FC protocol stack is divided into five layers:

- FC-0: defines the media type.
- FC-1: defines the frame encoding mode.
- FC-2: defines the frame format and flow control functions.
- FC-3: defines universal services.
- FC-4: defines mapping from the upper-layer protocol to the FC protocol.

Figure 2-4 Mapping from FC to FCoE

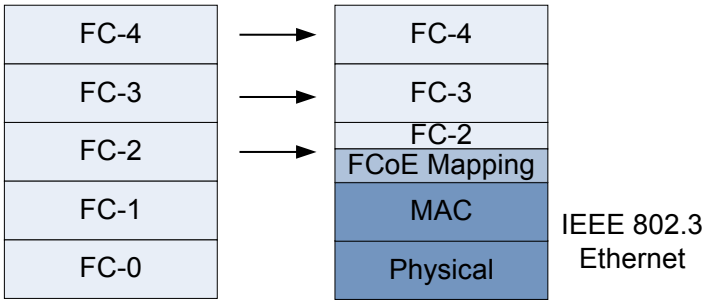

As shown in **Figure 2-4**, FC-0 and FC-1 in the FCoE protocol stack map Physical and MAC layers in IEEE 802.3 Ethernet respectively. The FCoE protocol stack adds an adaptation layer between the upper-layer FC protocol stack and lower-layer Ethernet protocol stack.

FCoE Frame Encapsulation

FCoE encapsulates an FC frame into an Ethernet frame. **Figure 2-5** shows FCoE frame encapsulation.

Figure 2-5 FCoE frame encapsulation

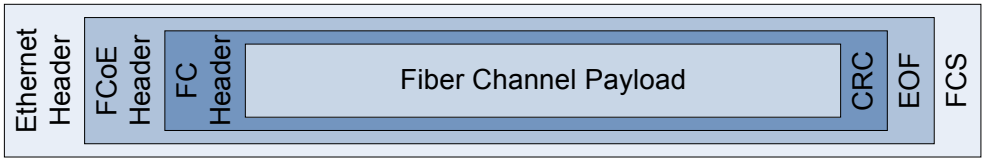

- The Ethernet Header defines the source and destination MAC addresses, Ethernet frame type, and FCoE VLAN. The Ethernet frame type value is FCoE_TYPE (8906h).
- The FCoE Header specifies the FCoE frame version number and control information.
- Similar to an FC frame, the FC Header in a FCoE frame carries the source and destination addresses.

2.3 FIP Protocol

Principles

The FCoE Initialization Protocol (FIP) establishes and maintains FCoE virtual links between FCoE devices, for example, between ENodes and FC forwarders (FCFs).

An FCoE virtual link is established as follows:

1. FIP discovers an FCoE VLAN and the FC virtual interface of the remote device.
2. FIP completes initialization tasks such as fabric login (FLOGI) and fabric discovery (FDISC) for the FCoE virtual link.

After an FCoE virtual link is set up, FIP maintains the FCoE virtual link in the following way:

- Periodically detects whether FC virtual interfaces at both ends of the FCoE virtual link are reachable.
- Tears down the FCoE virtual link through Fabric logout (FLOGO).

FCoE Virtual Link Setup

Figure 2-6 shows the process of setting up an FCoE virtual link between an ENode and an FCF. The ENode and FCF exchange FIP frames to establish the FCoE virtual link. After the FCoE virtual link is set up, FCoE frames are transmitted on the link. FIP frames and FCoE frames have different Ethernet types and encapsulation modes. FCoE frame encapsulation is defined in traditional FC protocol, whereas FIP frame encapsulation is not defined in traditional FC protocol. In FIP implementation, an ENode initiates all protocol packets. An FCF also initiates unsolicited FIP Advertisement packets, as described in **FIP FCF discovery**.

Figure 2-6 FCoE virtual link setup

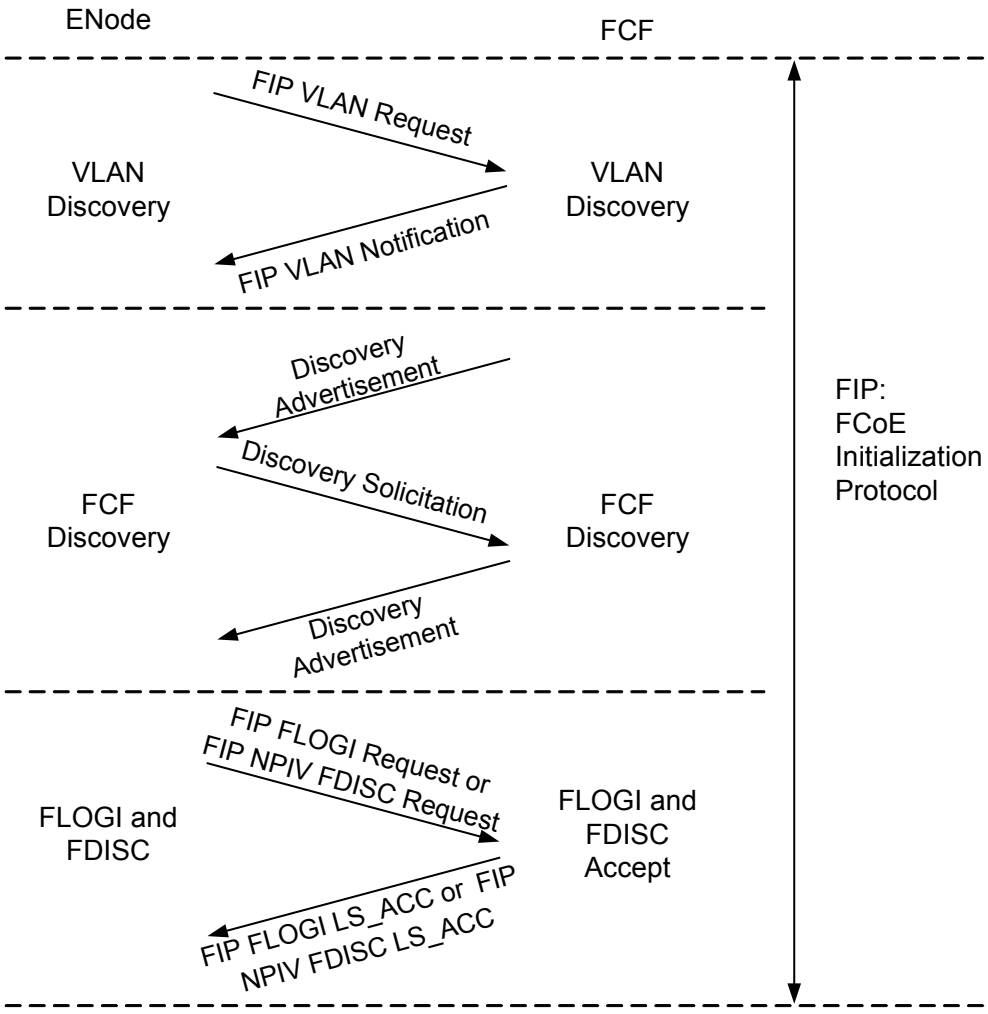

An FCoE virtual link is set up through three phases: FIP VLAN discovery, FIP FCF discovery, and FIP FLOGI and FDISC. The FIP FLOGI and FDISC processes are similar to FLOGI and FDISC processes defined in traditional FC protocol.

1. **FIP VLAN discovery**

   FIP VLAN discovery discovers the FCoE VLANs that will transmit FCoE frames. In this phase, an ENode can discover all the potential FCoE VLANs but does not select an FCF.

   The FIP VLAN discovery process is as follows:

   a. An ENode sends an FIP VLAN discovery request to a multicast MAC address called ALL-FCF-MAC (01-10-18-01-00-02). All FCFs listen on packets destined for this MAC address.

   b. All FCFs that are reachable in a common VLAN of the ENode report one or more FCoE VLANs to the ENode. The FCoE VLANs are available for the ENode's VN_Port login.

   FIP VLAN discovery is an optional phase as defined in FC-BB-5. An FCoE VLAN can be manually configured by an administrator, or dynamically discovered using FIP VLAN discovery.

2. **FIP FCF discovery**

   ENodes use FIP FCF discovery to locate FCFs that allow logins.

   The FIP FCF discovery process is as follows:

   a. Each FCF periodically sends Discovery Advertisement messages in each configured FCoE VLAN. The Advertisement messages are destined for the multicast MAC address ALL-ENode-MAC (01-10-18-01-00-01) on which all ENodes listen. The FIP FCF discovery Advertisement message contains the FCF MAC address and FCoE virtual link parameters such as the FCF priority and timeout interval of FIP packets.

   b. The ENode obtains FCF information from the received Discovery Advertisement messages, selects an FCF with the highest priority, and sends a unicast Discovery Solicitation message to the selected FCF.

   c. After receiving the Discovery Solicitation message, the FCF sends a unicast Discovery Advertisement message, allowing the ENode to log in.

   FCFs send Discovery Advertisement messages periodically, but new ENodes joining a network do not want to wait for Discovery Advertisement messages from all FCFs. Therefore, FC-BB-5 allows ENodes to send Discovery Solicitation messages to the multicast MAC address ALL-FCF-MAC. FCFs that receive the solicitation message send a unicast Discovery Advertisement message to the requesting ENode. Based on the received Discovery Advertisement messages, the ENode selects an FCF with the highest priority to set up a virtual link.

3. **FIP FLOGI and FDISC**

   After discovering all FCFs and selecting one for login, an ENode sends FIP FLOGI or FIP FDISC packets for establishing an FCoE virtual link with the VF_Port on the selected FCF. Then FCoE frames can be exchanged on the established FCoE virtual link. FIP FLOGI and FIP FDISC packets are unicast packets and correspond to FLOGI and FDISC packets in FC respectively. FIP FLOGI and FIP FDISC packets are used for allocating MAC addresses to ENodes so that the ENodes can log in to the fabric.

   FIP FLOGI is similar to FIP FDISC. The difference is as follows: FIP FLOGI refers to FCoE virtual ink setup when an ENode first logs in to the fabric. FIP FDISC refers to FCoE virtual link setup for each VM when multiple VMs exist on an ENode. FIP FLOGI is used an example.

   The FIP FLOGI process is as follows:

   a. An ENode sends an FIP FLOGI Request to the FCF.

   b. The FCF responds to the FIP FLOGI Request of the ENode and allocates a locally unique MAC address: Fabric Provided MAC Address (FPMA) to the ENode. Alternatively, the FCF responds to the FLOGI Request of the ENode, agreeing that the ENode uses its locally unique MAC address: Server Provided MAC Address (SPMA).

FCoE MAC address

An ENode uses different source MAC addresses to encapsulate FCoE and FIP frames. An FIP frame uses a globally unique MAC address (ENode MAC address) assigned to a converged network adapter (CNA) during manufacturing, whereas an FCoE frame uses a locally unique MAC address (unique only within the local Ethernet subnet) dynamically assigned to an ENode by an FCF during FCoE virtual link setup. For details, see **FIP FLOGI and FDISC**.

An ENode uses different source MAC addresses to encapsulate FCoE and FIP frames. An FIP frame uses a globally unique MAC address (ENode MAC address) assigned to a converged

network adapter (CNA) during manufacturing, whereas an FCoE frame uses a locally unique MAC address (unique only within the local Ethernet subnet). The locally unique MAC address is dynamically assigned to an ENode by an FCF or specified on an ENode and acknowledged by an FCF. For details, see **FIP FLOGI and FDISC**.

The locally unique MAC address is called a fabric-provided MAC address (FPMA).

In **Figure 2-7**, an FPMA has an FC_ID and a 24-bit FCoE MAC Address Prefix (FC-MAP). FC-BB-5 defines 256 FC-MAPs to facilitate FCoE deployment. In most cases, the default FC-MAP value 0E-FC-00 can meet deployment requirements. If FC_IDs on an Ethernet VLAN are not unique, FC_IDs may overlap, such as when different fabric or virtual storage area networks (VSANs) map to the same Ethernet VLAN. The use of different FC-MAPs solves this problem.

NOTE

Map one FC fabric to the same Ethernet VLAN. If multiple FC fabrics run on the same Ethernet, map the FC fabrics to different VLANs.

Figure 2-7 FPMA format

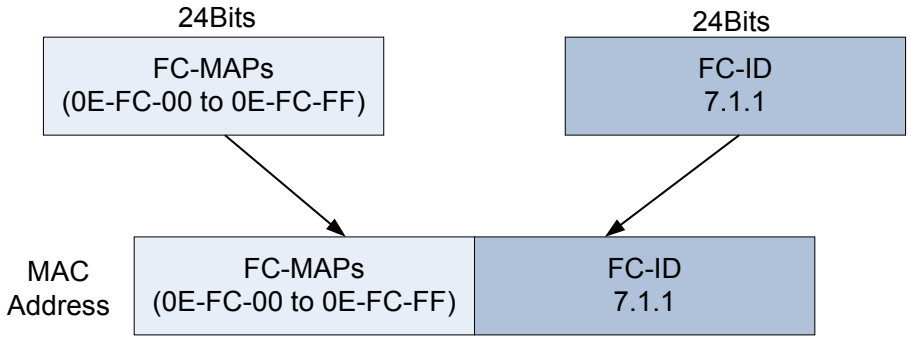

FCoE Virtual Link Maintenance

FCoE virtual link monitoring

On the traditional FC network, FC can immediately detect faults on a physical link. In FCoE, FC cannot immediately detect faults on a physical link because of Ethernet encapsulation. FIP provides a Keepalive mechanism to solve the problem.

FCoE monitors an FCoE virtual link as follows:

- An ENode periodically sends FIP Keepalive packets to an FCF. If the FCF does not receive FIP Keepalive packets within 2.5 times the keepalive interval, the FCF considers the FCoE virtual link faulty and terminates the FCoE virtual link.
- An FCF periodically sends multicast Discovery Advertisement messages with the destination MAC address as ALL-ENode-MAC to all ENodes. If an ENode does not receive multicast Discovery Advertisement messages within 2.5 times the keepalive interval, the ENode considers the FCoE virtual link faulty and terminates the FCoE virtual link.

FLOGO

If an FCF does not receive FIP Keepalive packets from an ENode, the FCF sends an FIP Clear Virtual Link message, requesting FCoE virtual link teardown. If the ENode logs out, the ENode can send a Fabric Logout request to the FCF, requesting the FCF to delete the virtual link.

2.4 FCF

An FCF supports both the FCoE and FC protocol stacks. The FCF encapsulates FC frames into Ethernet frames and uses FCoE virtual links to replace FC physical links. The FCF connects a SAN to a LAN.

Figure 2-8 FCF networking

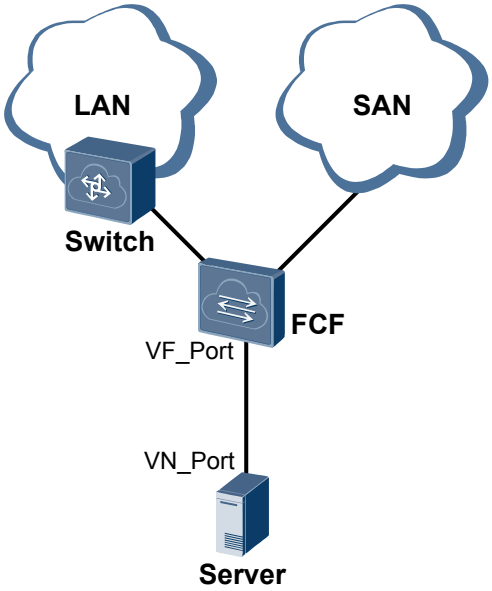

In **Figure 2-8**, an FCoE virtual link connects the VN_Port of the ENode to the VF_Port of the FCF. Each FCF has been allocated with a domain ID, and each SAN supports a maximum of 239 domain IDs. That is, a SAN has a maximum of 239 switches.

2.5 NPV

A SAN allows a maximum of 239 switches. When more edge switches are required, NPV switches can be deployed.

NOTE

The CE12800 does not support NPV.

Figure 2-9 NPV networking

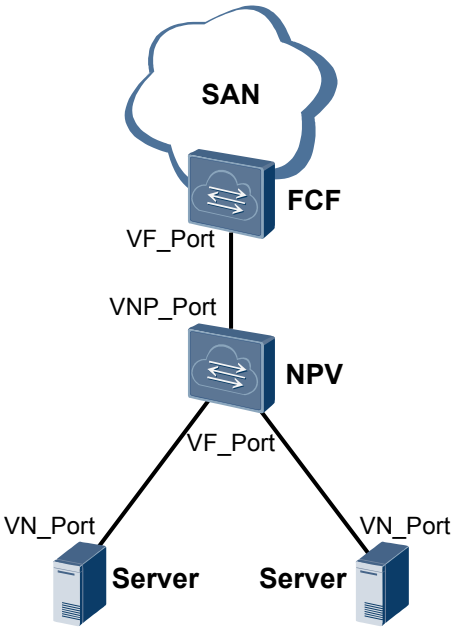

As shown in **Figure 2-9**, an NPV switch is located at the edge of a fabric and between an ENode and an FCF. The NPV switch uses a VF_Port to connect to a VN_Port of an ENode and uses a VNP_Port to connect to a VF_Port of an FCF. The ENode therefore connects to the fabric through the NPV switch, which forwards traffic from all ENodes to the core switch.

For an ENode, an NPV switch is equivalent to an FCF and uses the VF_Port. For an FCF, an NPV switch is equivalent to an ENode and uses the VN_Port.

2.6 FSB

An ENode and an FCF can establish a direct connection or remote connection. FIP snooping solves security problems in remote connection mode.

Direct Connection

As shown in **Figure 2-10**, when an ENode is directly connected to an FCF, the FCoE virtual link and its mapping physical link are point-to-point. Although packets forwarded on the physical link are encapsulated with FCoE, FCoE frame forwarding process is similar to FC frame forwarding because both ends of the physical link support FC.

Figure 2-10 Direct connection

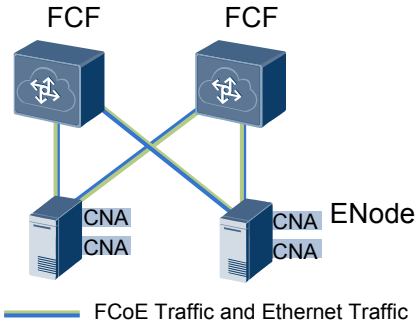

In direct connection mode, FCoE frame processing complies with FC except for data encapsulation at the data link layer. In this mode, FCoE has the same security as FC.

The direct connection mode allows SAN administrators to use original software to manage the SAN when FCoE is used.

Remote Connection

Because the FCF cost is high and a large number of servers are deployed in a data center, establishing direct connections between all servers and FCFs is impractical. As shown in **Figure 2-11**, access switches are deployed between FCFs and ENodes in remote connection mode. Access switches function as FCoE switches and cannot provide some FCF functions, such as FIP snooping bridge (FSB). In remote connection mode, one or more FCoE switches are deployed between ENodes and FCFs.

Figure 2-11 Remote connection

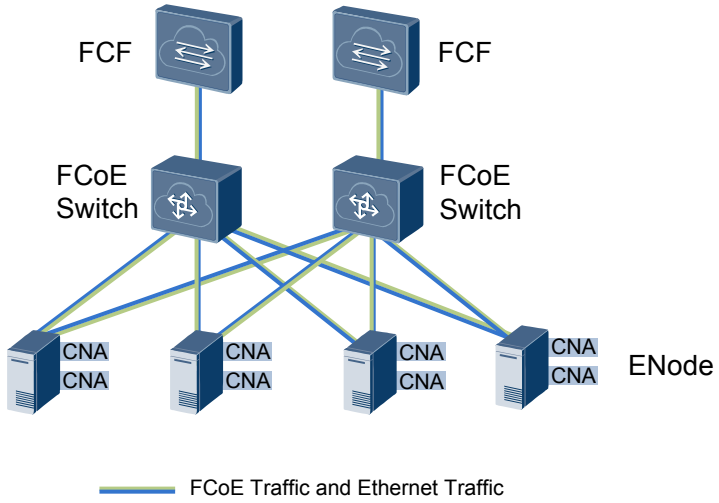

FIP Snooping

On an FC network, an FC switch is considered a trusted device. Other FC devices such as ENodes must get addresses assigned by the FC switch before they can connect to the FC network. The FC devices then log in to the FC switch. FC links are point-to-point, and an FC switch can completely control traffic received and sent by FC devices. Therefore, an FC switch ensures that devices use the assigned addresses to exchange packets and protect FC devices against malicious attacks.

When an FCoE switch is deployed between an ENode and an FCF, FCoE frames are forwarded on the FCoE switch based on the Ethernet protocol because the FCoE switch does not support the FC protocol. In this case, FCoE frames may not be destined for the FCF, and the point-to-point connection between the ENode and FCF is terminated.

To achieve equivalent robustness as an FC network, the FCoE switch must forward FCoE traffic from all ENodes to the FCF. FIP snooping enables the FSB to obtain FCoE virtual link information by listening on FIP packets. This function is used to control FCoE virtual link setup and prevent malicious attacks.

The FCoE switch running FIP snooping is called an FIP snooping bridge (FSB).

2.7 Zone

On an FCoE network, communication between ENodes can be controlled using zones to improve the network security.

Zone

A zone contains multiple members, and an ENode can join different zones. Members in the same zone can communicate with each other, whereas members in different zones cannot.

You can use the following items to define zone members:

- Zone alias: After a zone alias is added to a zone, members associated with the zone alias are also added to the zone.
- FC_ID: On an FC network, ENodes use FC_IDs to communicate with each other.
- FCoE interface: On an FCoE network, ENodes use FCoE interfaces to exchange with each other.
- World Wide Node Name (WWNN): A WWNN is a 64-bit address that identifies an ENode on an FC network.
- World Wide Port Name (WWPN): A WWPN is a 64-bit address that identifies an ENode interface on an FC network.

Figure 2-12 Networking of zones

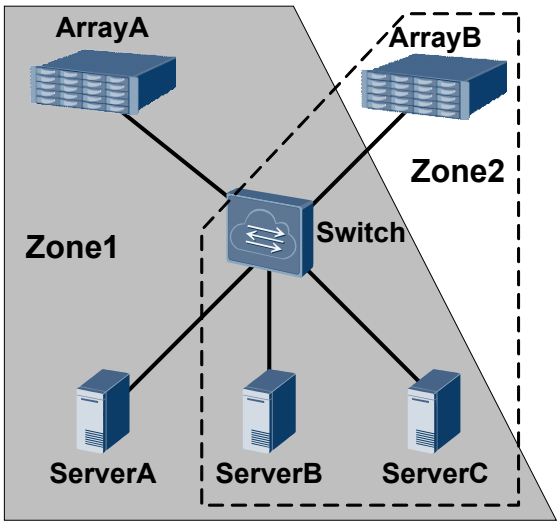

As shown in **Figure 2-12**, ENodes are added to different zones to control access of ENodes. For example, ArrayB can only communicate with ServerB and ServerC, but cannot communicate with ServerA.

Zone Set

A zone set contains multiple zones, and a zone can join different zone sets. Zones in a zone set take effect only after the zone set is activated, and only one zone set can be activated each time in one instance.

Zone Alias

A zone alias is used to simplify the zone configuration. When multiple member nodes need to be added to multiple zones, create a zone alias and associate the member nodes with the zone alias and add the zone alias to the zones. In this case, you do not need to add each node to each zone one by one.

Figure 2-13 Networking of the zone alias

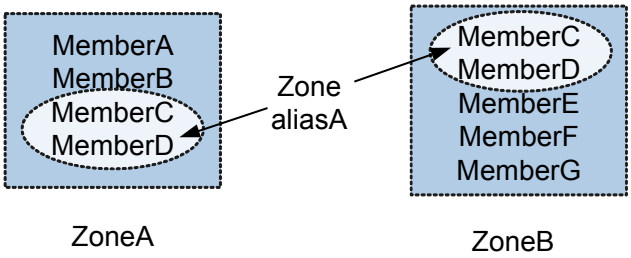

As shown in **Figure 2-13**, if MemberC and MemberD need to join ZoneA and ZoneB, to simplify configurations, add MemberC and MemberD to Zone aliasA and then add Zone aliasA to ZoneA and ZoneB.

3 Applications

About This Chapter

This section describes the applications of FCoE.

3.1 FCF Application

On a converged data center network, FCoE is used to implement network convergence. As shown in **Figure 3-1**, FCFA functions as an access switch and needs to forward LAN, SAN, and Inter-Process Communication (IPC) traffic. To ensure link reliability between ServerA and Array, two links are deployed: ServerA-FCFA-Array and ServerA-FCFB-Array.

Figure 3-1 Typical FCF networking

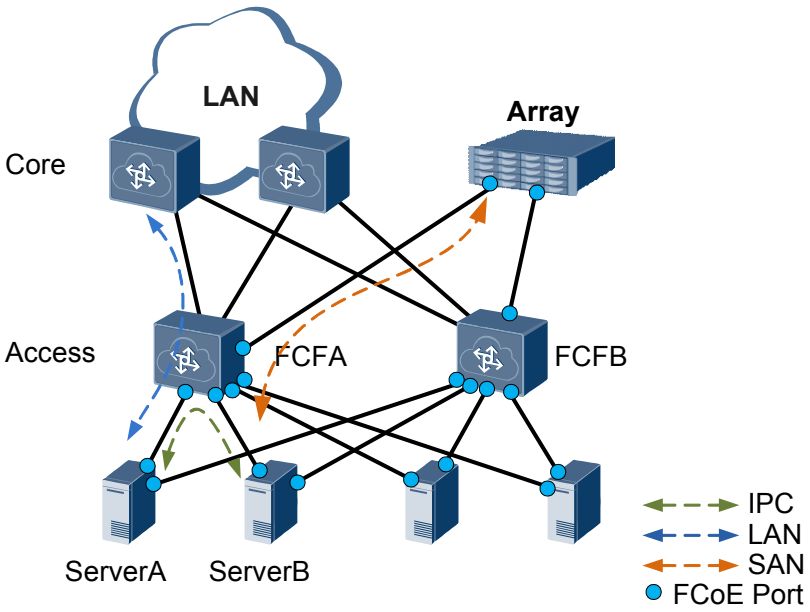

3.2 FCF+NPV Networking

On a converged data center network, FCoE is used to implement network convergence. To reduce network construction costs, access switches are often deployed between servers and FCFs. A SAN allows a maximum of 239 switches. When more edge switches are required, NPV switches can be deployed. As shown in **Figure 3-2**, NPVA functions as an access switch and needs to forward LAN, SAN, and Inter-Process Communication (IPC) traffic. FCFA serves as an aggregation switch and needs to forward SAN traffic. To ensure link reliability between ServerA and Array, two links are deployed: ServerA-NPVA-FCFA-Array and ServerA-NPVB-FCFB-Array. CE series switches only support single-hop FCF.

Figure 3-2 Typical networking of the FCF and NPV switch

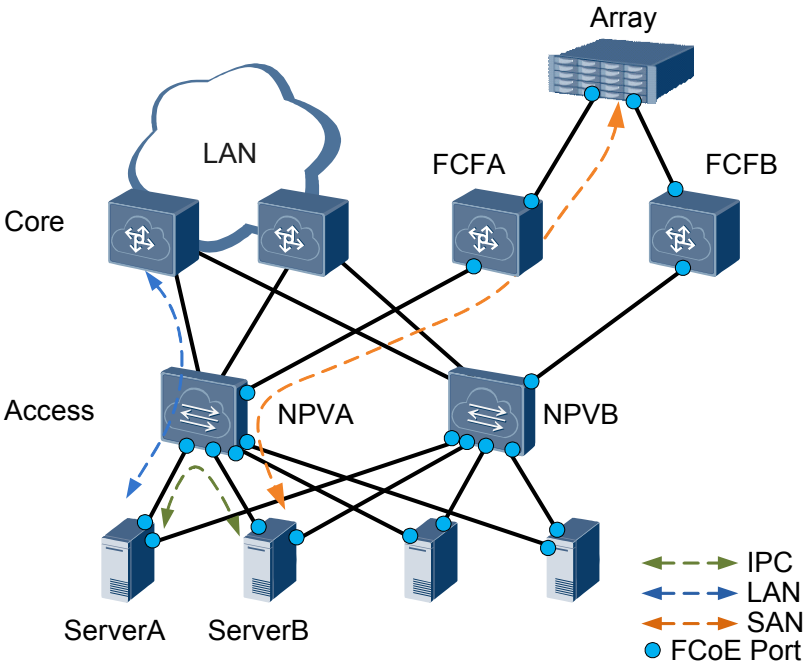

3.3 FCF+FSB Networking

On a converged data center network, FCoE is used to implement network convergence. To reduce network construction costs, access switches are often deployed between servers and FCFs. As shown in **Figure 3-3**, FSBA is used as an access switch and needs to forward LAN, SAN, and Inter-Process Communication (IPC) traffic; FCFA serves as a core switch and needs to forward SAN traffic. To ensure that SAN traffic is correctly forwarded, FIP snooping is configured on FSBA. To ensure link reliability between ServerA and Array, two links are deployed: ServerA-FSBA-FCFA-Array and ServerA-FSBB-FCFB-Array. CE series switches only support single-hop FCF.

When both the FCF and FSB are used, the dual-homing networking is recommended. In single-homing networking, when the FCoE-enabled device connects to some network adapters, the link recovery period is long.

Figure 3-3 Typical FCF+FSB networking

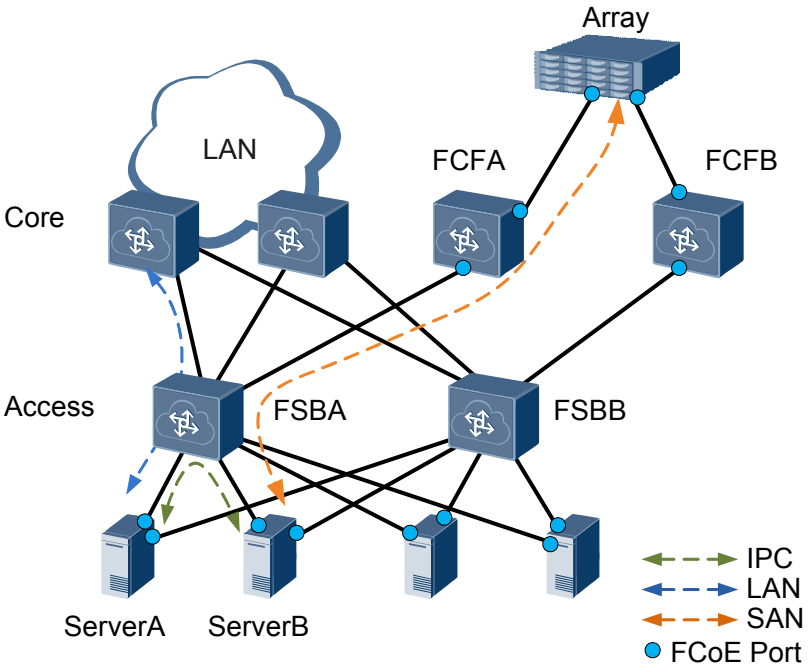

4 Configuration Task Summary

The device supports two FCoE modes: FCF and FSB. You can configure an FCoE mode as needed.

Table 4-1 describes the FCoE configuration tasks.

Table 4-1 FCoE configuration tasks

Scenario	Description	Task
Configure an FCF	An FCF supports both the FCoE and Fiber Channel (FC) protocol stacks. The FCF encapsulates FC frames into Ethernet frames and uses FCoE virtual links to replace FC links. The FCF connects a SAN to a LAN.	**8 Configuring an FCF**
Configure an FSB	An ENode and an FCF can establish a direct or remote connection. The FSB solves security problems in remote connection mode.	**9 Configuring FIP Snooping on the FSB**

5 Configuration Notes

This section describes the product models that support Fibre Channel over Ethernet (FCoE) and notes about configuring FCoE.

Involved Network Element

An FCoE network involves the following network elements (NEs):

- Server: needs to obtain large-capacity storage environment from a storage device.
- FCoE switch: transmits traffic on the FCoE network.
- Fiber Channel (FC) switch: transmits traffic on the storage area network (SAN).
- Storage device: provides large-capacity storage environment for servers.

License Support

- The FCoE forwarder (FCF) of the switch is controlled by the license. By default, the FCF is disabled on the switch. To use the FCF of the switch, purchase the license from the Huawei local office.
- The FCF is controlled by two types of licenses. The license CE-LIC-FCF-ALL supports FCF configuration on all the interfaces. The license CE-LIC-FCF-PORT supports FCF configuration on a total of 48 interfaces on the same card or on different cards. You can load multiple CE-LIC-FCF-PORT licenses to meet capacity expansion requirements.
- FIP snooping bridge (FSB) is a basic feature of a switch and is not under license control.
- If you load the licenses on multiple devices and then set up a stack system, the features activated for the stack are a combination of features activated for these member devices.
- You are advised to first configure the stack system, and then request the license based on the number of interfaces. If you load the license CE-LIC-FCF-PORT and then set up a stack system, the maximum number of FCF-enabled interfaces on a single device in the stack can be configured.
- Even if a license is uploaded to only one of these devices, the license is effective for the stack system. If the stack system breaks up, the licenses loaded to the devices respectively take effect. In addition, the activated licensed features can still be used on the devices even if these features are not activated using the devices' own license.
- If the license of Switch1 supports feature A and the license of Switch2 supports feature B, the two switches support features A and B after they set up a stack. If the stack splits, feature B on Switch1 can still run even after Switch1 restarts but new configurations cannot be added to feature B.

Version Support

Table 5-1 Products and minimum version supporting FCoE

Series	Function	Product	Minimum Version Required
CE12800	FSB	CE12804/CE12808/ CE12812	V100R003C00
		CE12816	V100R003C00
		CE12804S/ CE12808S	V100R005C00
	FCF	CE12804/CE12808/ CE12812	V100R005C00
		CE12816	V100R005C00
		CE12804S/ CE12808S	V100R005C00

Feature Dependencies and Limitations

Table 5-2 FCoE specifications

Item	Specification
Maximum number of instances	● VS: 32 ● System: 512 The total number of FCF and FSB instances cannot exceed the maximum value.
Maximum number of FCoE interfaces	512
Maximum number of VF_Ports	512
Maximum number of VN_Ports	● VS: 4096 ● System: 4096
Maximum number of zones	8192
Maximum number of zone sets	768
Maximum number of zone members	20000
Maximum number of zone aliases	512
Maximum number of members in each zone	64
Maximum number of concurrent online VN_Ports in each zone	64

- In V100R005C00 version, FCoE and TRILL cannot be used together. In V100R005C10 and later versions, by default, FCoE and TRILL cannot be used together. To use both of them, run the **trill adjacency-check disable** command first. The TRILL function has a higher priority than FCoE. If FCoE is configured before TRILL, only TRILL takes effect.
- If an FCF instance has been configured on the device, configuring an FSB instance will cause FCF traffic loss.
- When FCoE is used, the dual-homing networking is recommended. In single-homing networking and both the FCF and FSB are used, when the FCoE-enabled device connects to some network adapters, the link recovery period is long.
- In V100R006C00 and later versions, an SVF system supports FCoE. When a CE6810LI or a CE5800 functions as a leaf switch, an SVF system does not support FCoE.
- When the FCoE-enabled device connects to the QLogic network adapter, the QLogic network adapter that does not receive responses retransmits FLOGI requests every 20s. The FCoE-enabled device enters the VLAN Request state after a period of time. That is, the link recovery period may be as long as 3 minutes.
- When the FCoE-enabled device connects to the Brocade 1020 network adapter, you must configure independent enhanced transmission selection (ETS) priority groups for FCoE and Internet Small Computer Systems Interface (iSCSI) services and enable Priority-based Flow Control (PFC) for queues with priorities. Otherwise, Data Center Bridging eXchange protocol (DCBX) negotiation may fail.
- The Cisco Nexus 5000 that functions as the FCoE NPort Virtualization (NPV) device cannot connect to FCoE forwarders (FCFs) from other vendors.
- The peer-link interfaces of the Multichassis Link Aggregation Group (M-LAG) must be removed from the FCoE VLAN.
- In V100R005C100 and later versions, when the FSB and FCF are configured on the device, M-LAG is unavailable.
- The device does not support outbound traffic statistics collection on an interface.
- The device does not support in-service software upgrade (ISSU) for FCoE.
- FCoE configuration depends on successful DCB negotiation.
- The PFC priorities of member interfaces must be the same in the same instance.

6 Compatibility List

This section describes the compatibility when the Fibre Channel over Ethernet (FCoE) device connects to the third-party network adapter, server, storage device, or switch.

Y indicates compatibility; N indicates incompatibility; N/A indicates that the function is not supported or there is no such a scenario. - indicates that interworking compatibility is not ensured.

NOTE

- Interworking compatibility is not ensured in the scenarios that are not listed in the following table.
- Only the CE6850U-HI supports Fiber Channel (FC) interfaces.

Table 6-1 Network adapter

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810 LI and CE6850 U-HI)	CE6850U-HI	Remarks
Emulex	LPe16002 HBA	Windows 2012 RedHat Enterprise Linux 7 SUSE Linux Enterprise Linux 12 VMWARE 5.5	FC	N/A	N/A	N/A	N/A	Y	The network adapter uses the built-in drive of VMWARE 5.5, and there are many problems about this network adapter. The network adapter goes online slowly, the rate is about half of the full bandwidth, and the disk cannot be found after the storage devices are registered with the FCF again. To solve the problems, upgrade the drive to 10.0.725.203.

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	LPe12002 HBA	VMWARE 5.1.0 Windows 2008 SP2 Windows 2008 R2 SP1 RedHat Enterprise Linux 6 SUSE Linux Enterprise Linux 11 Solaris 11	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	LPe1250 HBA	Windows 2008 Windows 2008 R2 SP1 RedHat Enterprise Linux 6 SUSE Linux Enterprise Linux 11 VMware 5.1	FC	N/A	N/A	N/A	N/A	Y	The network adapter users a later version of VMWARE 5.1u2. It can be registered normally. After the network adapter initiates IO, the optical module does not work and the interface becomes Down.

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	LPe1150 HBA	Windows 2008 SP2 RedHat Enterprise Linux 5 SUSE Linux Enterprise Linux 10 VMware 5.0	FC	N/A	N/A	N/A	N/A	Y	None
	LPe111 HBA	Windows 2008 SP2 RedHat Enterprise Linux 5 SUSE Linux Enterprise Linux 10 VMware 5.0	FC	N/A	N/A	N/A	N/A	Y	The network adapter does not support VMWARE operating system.

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810 LI and CE6850 U-HI)	CE6850U-HI	Remarks
	OCe11100 CNA	VMWARE Win Server 2008	FCoE	Y	Y	Y	Y	N/A	None
Qlogic	QLE2462 HBA	VMWARE 5.1.0 VMWARE 5.0 Windows 2008 SP2 RedHat Enterprise Linux 5 SUSE Linux Enterprise Linux 10 Solaris 10	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	QLE2562 HBA	VMware 5.1 Windows 2008 R2 SP1 RedHat Enterprise Linux 6 SUSE Linux Enterprise Linux 11 Solaris 11	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	QLE2672 HBA	Windows 2012 RedHat Enterprise Linux 7 SUSE Linux Enterprise Linux 12	FC	N/A	N/A	N/A	N/A	Y	The destination FC_ID in PLOGI packets sent by the network adapter is the FC_ID of the NPV switch, and the NPV switch directly discards the PLOGI packets. As a result, the host where the network adapter is installed cannot discover the remote host. Restart the host where the network adapter is installed. Then PLOGI packets do not use the FC_ID of the NPV switch as the destination FC_ID. The Cisco switch used as the NPV

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
									switch does not respond to PLOGI packets with the FC_ID of the NPV switch as the destination FC_ID. In FCF+NPV networking, it is recommended that members on the FCF join a zone in WWN mode.

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810 LI and CE6850 U-HI)	CE6850U-HI	Remarks
	8200 Series CNA	VMWARE Win Server 2008	FCoE	Y	Y	Y	Y	Y	● When the FCoE-enabled device connects to the QLogic network adapter, the QLogic network adapter that does not receive responses retransmits FLOGI requests every 20s. The FCoE-enabled device enters the VLAN Request state after a period of time. That is, the link recovery period may be as long as 3 minutes. ● When both the FCoE forwarder (FCF) and FIP

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
									snooping bridge (FSB) are used, the dual-homing networking is recommended. In single-homing networking, when the FCoE-enabled device connects to some network adapters, the link recovery period is long.

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810 LI and CE6850 U-HI)	CE6850U-HI	Remarks
Intel	X520 Series	VMWARE Win Server 2008	FCoE	Y	Y	Y	Y	-	● Data Center Bridging (DCB) parameters need to be manually set. ● When the Intel X520-2 network adapter connects to an FCoE FCF, there is a low probability that the following problem occurs. After receiving FLOGI ACC packets, the FCF enters the VLAN Request state and does not send PLOGI packets. This problem is

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
									solved after the network adapter is disabled.
Brocade	18602 HBA	VMWARE 5.1.0	FC	N/A	N/A	N/A	N/A	Y	None
	1860 HBA	VMWARE 5.5 Windows 2012 RedHat Enterprise Linux 7 SUSE Linux Enterprise Linux 12	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810 LI and CE6850 U-HI)	CE6850U-HI	Remarks
	425 HBA	VMWARE 5.0 Windows 2008 SP2 RedHat Enterprise Linux 5 SUSE Linux Enterprise Linux 10	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	825 HBA	VMWARE 5.1 Windows 2008 R2 SP1 RedHat Enterprise Linux 6 SUSE Linux Enterprise Linux 11	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Operating System	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	1000 Series	VMWARE Win Server 2008	FCoE	Y	Y	Y	Y	Y	When the FCF connects to the Brocade 1020 network adapter, you must configure independent ETS priority groups for FCoE and iSCSI services and enable PFC for queues with priorities. Otherwise, DCBX negotiation may fail.
ATTO	FC-42ES HBA	Windows 2008 SP2	FC	N/A	N/A	N/A	N/A	Y	None
	FC-82EN HBA	Windows 2008 R2 SP1	FC	N/A	N/A	N/A	N/A	Y	None
HP	AH401A	HP UX 11.31	FC	N/A	N/A	N/A	N/A	Y	None
	AT094－60001	HP UX 11.32	FC	N/A	N/A	N/A	N/A	Y	None

Table 6-2 Blade server

Manufacturer	Model	Network Adapter	SFU Function	CE12800	CE8800	CE7800	CE6800 (excluding CE68100LI and CE6850U-HI)	CE6850U-HI	Remarks
Huawei	E9000	CX311-10GE/CNA	FCoE	Y	Y	Y	Y	N/A	Brocade network adapter provides FCoE functions and supports FCF and NPV modes.
		Brocade FC SFU	FCF	N/A	N/A	N/A	N/A	Y	Brocade network adapter provides FCoE functions and supports FCF and NPV modes. To use the NPV mode, load a license. By default, the FCF mode is used.
	E6000	Brocade FC SFU	NPV	N/A	N/A	N/A	N/A	Y	
	T8000	Qlogic FC SFU	FCF/NPV	N/A	N/A	N/A	N/A	Y	None
IBM	Blade Center E	HBA	FC	N/A	N/A	N/A	N/A	Y	None

Manufacturer	Model	Network Adapter	SFU Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
Dell	PowerEdge M1000e	HBA	FC	N/A	N/A	N/A	N/A	Y	None
	Force 10 mxl	CNA	FSB	Y	Y	Y	Y	-	None
H3C	C8000	Emulex FlexFabric 10GB-2-PORT 554FLB Adapter	FSB	Y	Y	Y	Y	-	None

Table 6-3 Storage device

Manufacturer	Model	Network Adapter	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
Huawei	OceanStor S5500T v3	Brocade CNA	FCoE	Y	Y	Y	Y	N/A	The switch functions as the FCF.
		Chelsio T4 HBA	FC	N/A	N/A	N/A	N/A	Y	None
		8G PMC	FC	N/A	N/A	N/A	N/A	Y	None
		SmartIO FC	FC	N/A	N/A	N/A	N/A	Y	None
		QLogic 16G FC	FC	N/A	N/A	N/A	N/A	Y	The network adapter does not send registration messages (FLOGI) repeatedly. When the FC interface Up event is reported slowly, the storage device may fail to go online.
	OceanStor S5600T	CNA	FCoE	Y	Y	Y	Y	Y	None
EMC	VNX4500	CNA	FCoE	Y	Y	Y	Y	-	None

Manufacturer	Model	Network Adapter	Function	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	VNX5300	HBA	FC	N/A	N/A	N/A	N/A	Y	None
	VNX5500	CNA	FCoE	Y	Y	Y	Y	-	The switch functions as the FCF.
NetApp	FAS2000	HBA	FC	N/A	N/A	N/A	N/A	Y	None
	FAS8040	CNA	FCoE	Y	Y	Y	Y	-	The switch functions as the FCF.
HP	EVA8100	HBA	FC	N/A	N/A	N/A	N/A	Y	None
IBM	DS4300	HBA	FC	N/A	N/A	N/A	N/A	Y	None
SUN	ST6180	HBA	FC	N/A	N/A	N/A	N/A	Y	None

Table 6-4 Switch

Manufacturer	Model	Device Role	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
HUAWEI	SNS2120	FC FCF	N/A	N/A	N/A	N/A	Y	None
	SNS5120	FC FCF	N/A	N/A	N/A	N/A	Y	
	SNS2124	FC FCF	N/A	N/A	N/A	N/A	Y	
	SNS2248	FC FCF	N/A	N/A	N/A	N/A	Y	
Cisco	N5K	FCoE FCF	N/A	Y	Y	Y	Y	The CE12800 does not support FCoE NPort Virtualization (NPV).
		FCoE NPV	N	N	N	N	N	The Cisco Nexus 5000 that functions as the FCoE NPV device cannot connect to FCFs from other vendors.
		FC FCF	N/A	N/A	N/A	N/A	Y	When a server and a storage device connect to the N5K FCF, the server cannot detect the storage device.

Manufacturer	Model	Device Role	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
		FC NPV	N/A	N/A	N/A	N/A	Y	When the N5K functions as the FC NPV device and connects to the CE FCF, if an ENode is online and the **shutdown** and **undo shutdown** commands are configured on the FC interface of the CE, the FC interface in Up state lasts for more than 90s. When no ENode is online, the FC interface in Up state lasts for less than 5s.
	MDS 9K	FCoE FCF	N/A	Y	Y	Y	-	When the CE switch functions as the FCoE NPV device and connects to the MDS9250i, do not use the Eth-Trunk. Use single interfaces and dual planes to ensure reliability.
		FC FCF	N/A	N/A	N/A	N/A	-	None
		FC NPV	N/A	N/A	N/A	N/A	-	None
Brocade	VDX 6730	FCoE FCF	N/A	Y	Y	Y	Y	When a server and a storage device connect to an NPV device (CE switch) and Brocade 8000
	6510	FC FCF	N/A	N/A	N/A	N/A	-	

Manufacturer	Model	Device Role	CE12800	CE8800	CE7800	CE6800 (excluding CE6810LI and CE6850U-HI)	CE6850U-HI	Remarks
	300	FC FCF	N/A	N/A	N/A	N/A	-	FCF, and are registered, a zone is configured on the Brocade 8000 FCF. In this case, the server cannot detect the storage device.
	8000	FC FCF	N/A	N/A	N/A	N/A	Y	

7 Default Configuration

This section provides the default Fibre Channel over Ethernet (FCoE) configuration.

Table 7-1 lists the default FCoE configuration.

Table 7-1 Default FCoE configuration

Parameter	Default Setting
FCF instance	Not configured
FSB instance	Not configured
Zone	Not configured
Domain ID	2
FC-MAP value	0x0efc00
Timeout interval at which FIP packets are exchanged	5 minutes
Interval of FIP Keepalive packets in the FCF instance and NPV instance	8000 milliseconds

8 Configuring an FCF

About This Chapter

On a converged network, an FCoE forwarder (FCF) connects a traditional storage area network (SAN) to a local area network (LAN). An FCF can forward Fibre Channel over Ethernet (FCoE) packets and encapsulate or decapsulate FCoE packets.

Pre-configuration Tasks

Before configuring an FC FCF, complete the following task:

- Configuring a link layer protocol so that interfaces at both ends of a link are Up.

Before configuring an FCoE FCF, complete the following task:

- Configuring a link layer protocol so that interfaces at both ends of a link are Up.
- Configure Data Center Bridging (DCB). After DCB negotiation is successful, configure the FCoE forwarder (FCF).

Configuration Process

Figure 8-1 FCF configuration flowchart

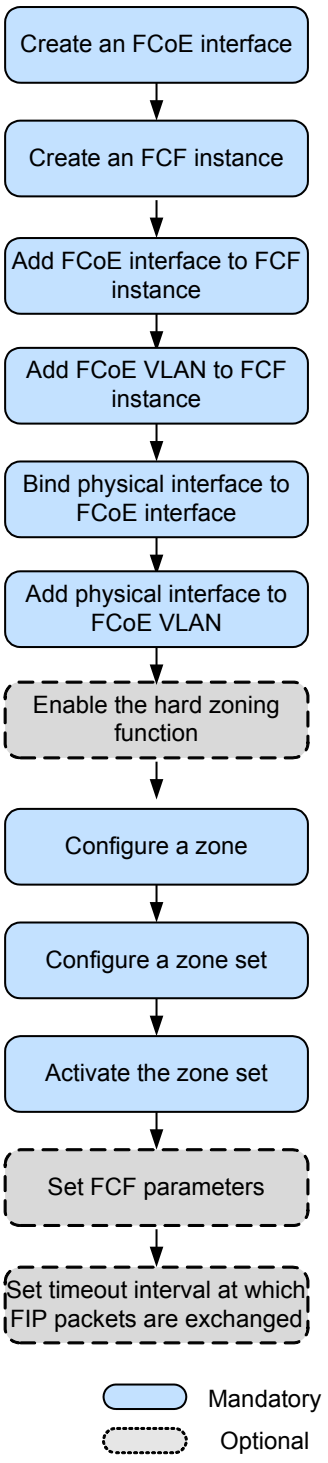

8.1 Creating an FCoE Interface and FCF Instance

8.2 Configuring a Zone

8.3 (Optional) Setting FCF Parameters

8.1 Creating an FCoE Interface and FCF Instance

Context

On a traditional Fiber Channel (FC) network, all FC nodes are interconnected through a fabric. FC nodes belonging to a fabric have the same attributes. To control Fibre Channel over Ethernet (FCoE) traffic forwarding, FCoE uses FC instances because FCoE transmits storage area network (SAN) traffic over the Ethernet. An FC instance defines fabric attributes such as the FCoE VLAN. A VF_Port is a virtual logical interface that is manually created on an FCoE forwarder (FCF), and provides functions of a physical FC interface. The VF_Port can work properly only when the VF_Port is bound to a physical Ethernet interface.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface fcoe-port interface-number
```

An FCoE interface is created and the FCoE interface view is displayed.

Step 3 Run:

```
quit
```

Return to the system view.

Step 4 Run:

```
fcoe fc-instance-name fcf
```

An FCF instance is created and the FCF instance view is displayed.

By default, no FCF instance is configured. Each virtual system (VS) supports a maximum of 32 FCF instances, and the device supports a maximum of 512 instances.

NOTE

An FC instance name is case sensitive. For example, fc1 and FC1 identify different FC instances.

Step 5 Run:

```
member interface fcoe-port fcoe-port1 [ to fcoe-port2 ]
```

The FCoE interface is added to the FCF instance.

By default, no FCoE interface is added to an FCF instance.

Step 6 Run:

```
vlan vlan-id
```

An FCoE VLAN is added to the FCF instance.

By default, no FCoE VLAN is configured.

An FCoE VLAN is only used to forward FCoE packets and FIP packets. An FCoE VLAN belongs to only one FCF instance.

Step 7 Run:

```
quit
```

Return to the system view.

Step 8 Run:

```
interface interface-type interface-number
```

The view of a physical interface to be added to an FCoE interface is displayed.

Step 9 Run:

```
fcoe-port fcoe-port-id
```

The physical interface is added to the FCoE interface.

By default, no physical interface is added to an FCoE interface.

Step 10 Run:

```
port link-type { trunk | hybrid }
```

The link type of the interface is configured.

By default, the link type of an interface is access.

Step 11 Run:

```
port trunk allow-pass vlan { { vlan-id1 [ to vlan-id2 ] }&<1-10> | all }
```

Or,

```
port hybrid tagged vlan { { vlan-id1 [ to vlan-id2 ] }&<1-10> | all }
```

The interface is added to the FCoE VLAN.

Step 12 Run:

```
commit
```

The configuration is committed.

----End

8.2 Configuring a Zone

Context

On an Fibre Channel over Ethernet (FCoE) network, communication between ENodes can be controlled using zones to improve the network security.

A zone contains multiple members, and an ENode can join different zones. Members in the same zone can communicate with each other, whereas members in different zones cannot.

A zone set contains multiple zones, and a zone can join different zone sets. Zones in a zone set take effect only after the zone set is activated, and only one zone set can be activated each time in one instance.

A zone alias is used to simplify the zone configuration. When multiple member nodes need to be added to multiple zones, create a zone alias and associate the member nodes with the zone alias and add the zone alias to the zones. In this case, you do not need to add each node to each zone one by one.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 (Optional) Run:

```
zone hard-zoning enable
```

The hard zoning function is enabled.

By default, the hard zoning function is disabled.

Step 3 (Optional) Run:

```
zone alias alias-name
```

A zone alias is created and the zone alias view is displayed.

By default, no zone alias is configured on the device.

Step 4 (Optional) Run:

```
member { fcid fc-id | interface fcoe-port fcoe-port-id | wwnn wwnn | wwpn wwpn }
```

A member node is associated with the zone alias.

By default, no member node is associated with a zone alias.

Step 5 (Optional) Run:

```
quit
```

Return to the system view.

Step 6 Run:

```
zone zone-name
```

A zone is created and the zone view is displayed.

By default, no zone is configured on the device.

Step 7 Run:

```
member { alias alias-name | fcid fc-id | interface fcoe-port fcoe-port-id | wwnn
wwnn [ auto | both | initiator | target ] | wwpn wwpn [ auto | both | initiator |
target ] }
```

A member node is added to the zone.

By default, no member node is added to a zone.

Step 8 Run:

```
quit
```

Return to the system view.

Step 9 Run:

```
zoneset zoneset-name
```

A zone set is created and the zone set view is displayed.

By default, no zone set is configured on the device.

Step 10 Run:

```
member zone-name
```

The zone is added to the zone set.

By default, no zone is added to a zone set.

Step 11 Run:

```
quit
```

Return to the system view.

Step 12 Run:

```
fcoe fc-instance-name fcf
```

The FCoE forwarder (FCF)F instance view is displayed.

Step 13 (Optional) Run:

```
smart-zoning enable
```

The Smart Zone function is enabled.

By default, the Smart Zone function is disabled.

Step 14 Run:

```
active zoneset zoneset-name
```

The zone set is activated.

By default, no zone set is activated.

Step 15 Run:

```
commit
```

The configuration is committed.

----End

8.3 (Optional) Setting FCF Parameters

Procedure

- Configure a domain ID.
  a. Run:

     ```
     system-view
     ```

     The system view is displayed.
  b. Run:

     ```
     fcoe fc-instance-name fcf
     ```

     An FCoE forwarder (FCF) instance is created and the FCF instance view is displayed.
  c. Run:

     ```
     domain-id domain-id
     ```

     A domain ID is configured for an FCF instance.

     By default, the domain ID is 2.
  d. Run:

     ```
     commit
     ```

The configuration is committed.

- Configure an FC-MAP value.
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  fcoe fc-instance-name fcf
  ```
  An FCF instance is created and the FCF instance view is displayed.
  c. Run:
  ```
  fcmap fc-map
  ```
  An FC-MAP value is configured.

  By default, the FC-MAP value is 0x0efc00.
  d. Run:
  ```
  commit
  ```
  The configuration is committed.
- Enable the FCF to retain the FC_ID after relogin of the ENode.
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  fcoe fc-instance-name fcf
  ```
  An FCF instance is created and the FCF instance view is displayed.
  c. Run:
  ```
  fcid persistent
  ```
  The FCF is enabled to retain the FC_ID after relogin of the ENode.

  By default, the FCF is disabled from retaining the FC_ID after relogin of the ENode.
  d. Run:
  ```
  commit
  ```
  The configuration is committed.

----End

8.4 (Optional) Setting the Interval for Sending FIP Keepalive Packets

Context

FIP Keepalive packets is used to detect the Fibre Channel over Ethernet (FCoE) virtual link status. If the switch used as the FCoE forwarder (FCF) does not receive FIP Keepalive packets after the interval multiplied by 2.5, the switch considers that the virtual link fails and terminates the virtual link.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
fcoe fc-instance-name [ fc-instance-type ]
```

The FCF instance view is displayed.

Step 3 Run:

```
fip fka-adv-period interval_value
```

The interval at which FIP Keepalive packets are sent is set.

By default, the interval at which FIP Keepalive packets are sent is 8000 ms.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

8.5 Checking the Configuration

Prerequisites

All configurations of the FCoE forwarder (FCF) are complete.

Procedure

- Run the **display fcoe port** { **brief** | **instance** *fc-instance-name* | **interface fcoe-port** *fcoe-port-id* } command to check the FCoE interface configuration.
- Run the **display fcoe instance** [**fcf** [*fc-instance-name*]] command to check the FCF instance configuration.

----End

9 Configuring FIP Snooping on the FSB

About This Chapter

Fibre channel Initialization Protocol (FIP) snooping enables the FIP snooping bridge (FSB) to obtain Fibre Channel over Ethernet (FCoE) virtual link information by listening on FIP packets. This function is used to control FCoE virtual link setup and prevent malicious attacks.

Pre-configuration Tasks

Before configuring FIP snooping, complete the following task:

- Configuring a link layer protocol so that interfaces at both ends of a link are Up

Configuration Process

Figure 9-1 FSB configuration flowchart

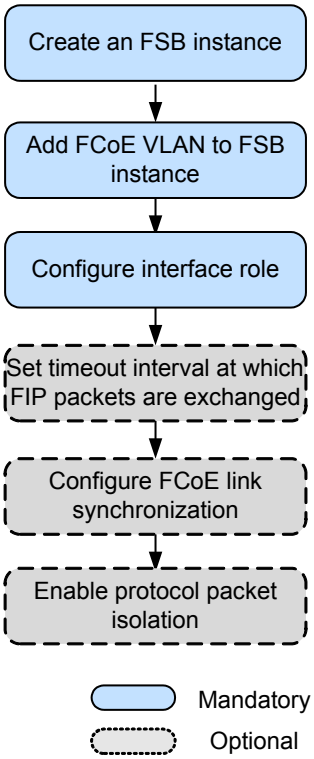

9.1 Configuring an FC Instance

9.2 Configuring a Role for an Interface

9.3 (Optional) Setting the Timeout Interval for Exchanging FIP Packets

9.4 (Optional) Configuring FCoE Link Synchronization

9.5 Checking the Configuration

9.1 Configuring an FC Instance

Context

On a traditional Fiber Channel (FC) network, all FC nodes are interconnected through fabrics. FC nodes belonging to a fabric have the same attributes. Fibre Channel over Ethernet (FCoE) carries storage area network (SAN) traffic over the Ethernet and uses FC instances to control FCoE traffic forwarding; therefore, an FC instance defines fabric attributes such as the FCoE VLAN.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
fcoe fc-instance-name [ fsb ]
```

An FSB instance is created, and the FSB instance view is displayed.

By default, no FSB instance is configured. Each virtual system (VS) supports a maximum of 32 FSB instances, and the device supports a maximum of 512 FSB instances.

NOTE

An FC instance name is case sensitive. For example, fc1 and FC1 identify different FC instances.

Step 3 Run:

```
vlan vlan-id
```

An FCoE VLAN is added to the FSB instance.

An FCoE VLAN is only used to forward FCoE traffic and FIP packets. An FCoE VLAN belongs to only one FSB instance.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

9.2 Configuring a Role for an Interface

Context

On a traditional Fiber Channel (FC) network, initiators and targets of storage area network (SAN) traffic are explicitly specified. Because traditional Ethernet does not define the traffic transmission direction, roles of FIP snooping bridge (FSB) interfaces must be defined during network planning.

Figure 9-2 Roles of FSB interfaces

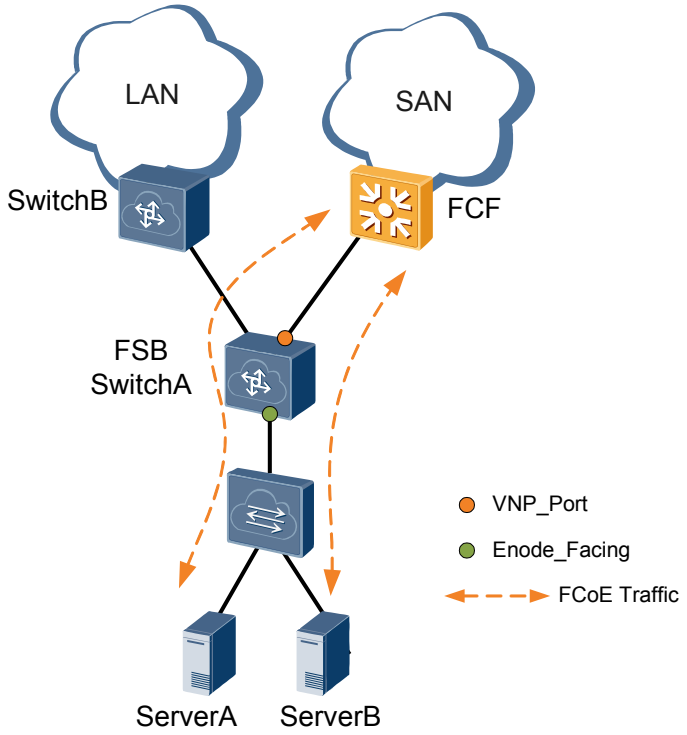

Roles of FSB interfaces are as follows:

- VNP_Port: FCoE switch port connected to an FCF
- ENode_Facing port: FCoE switch port connected to a server

NOTE

Before configuring a role for an interface, add the interface to an FCoE VLAN.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface interface-type interface-number
```

The Ethernet interface view is displayed.

Step 3 Run:

```
port link-type { trunk | hybrid }
```

The link type of the interface is configured.

By default, a port is an access port.

Step 4 Run:

```
port trunk allow-pass vlan { { vlan-id1 [ to vlan-id2 ] }&<1-10> | all }
```

Or

```
port hybrid tagged vlan { { vlan-id1 [ to vlan-id2 ] }&<1-10> | all }
```

The interface is added to an FCoE VLAN.

Step 5 Run:

```
fcoe role vnp
```

The interface is configured as a VNP_Port.

By default, an Ethernet interface is an ENode_Facing port.

Step 6 Run:

```
commit
```

The configuration is committed.

----End

9.3 (Optional) Setting the Timeout Interval for Exchanging FIP Packets

Context

Fibre channel Initialization Protocol (FIP) packets are FIP Keepalive packets used to detect the Fibre Channel over Ethernet (FCoE) virtual link status. When the FIP snooping bridge (FSB) does not receive FIP Keepalive packets within the timeout interval, the FSB considers the FCoE virtual link faulty and terminates the FCoE virtual link.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
fcoe fc-instance-name [ fsb ]
```

The FSB instance view is displayed.

Step 3 Run:

```
fip fka-adv-period interval_value
```

The timeout interval for exchanging FIP packets is set.

By default, the timeout interval for exchanging FIP packets is 5 minutes.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

9.4 (Optional) Configuring FCoE Link Synchronization

Context

In a Fibre channel Initialization Protocol (FIP) snooping scenario shown in **Figure 9-3**, the FCoE forwarder (FCF) cannot immediately detect a link failure between the ENode and switch and can perform link switching only after the Keepalive timer expires. This may result in traffic interruption.

Figure 9-3 Application scenario of FCoE link synchronization

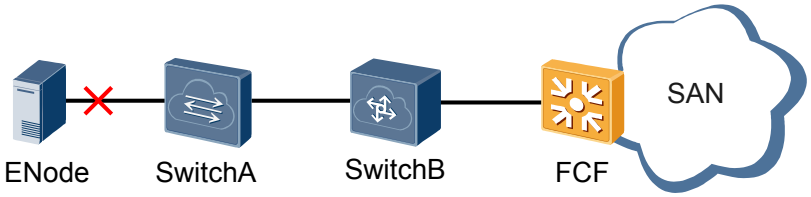

After Fibre Channel over Ethernet (FCoE) link synchronization is enabled, the switch instructs the FCF or ENode to perform link switching immediately after detecting a link failure. This function ensures uninterrupted traffic forwarding.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
fcoe link sync
```

FCoE link synchronization is enabled.

By default, FCoE link synchronization is disabled.

Step 3 Run:

```
commit
```

The configuration is committed.

----End

9.5 Checking the Configuration

Prerequisites

The configurations of Fibre channel Initialization Protocol (FIP) snooping are complete.

Procedure

- Run the **display fcoe vlan** [*vlan-id*] command to check the FCoE VLAN configuration and the FCoE port roles.
- Run the **display current-configuration** | **include fcoe** command to check whether FCoE link synchronization is enabled and whether protocol packet isolation is enabled.

- Run the **display current-configuration configuration fcoe** command to check information about FSB instances.

----End

10 Maintaining FCoE

About This Chapter

This section describes how to check and clear Fibre channel Initialization Protocol (FIP) packet statistics, monitor Fibre Channel over Ethernet (FCoE) running status, and debug FCoE.

10.1 Displaying Statistics on FIP Packets

Context

After the Fibre Channel over Ethernet (FCoE) configuration is complete, run the following command in any view to check statistics on Fibre channel Initialization Protocol (FIP) packets.

Procedure

Run the **display fcoe fip statistics** [*fc-instance-name*] command to check statistics on FIP packets of the specified FIP snooping bridge (FSB) instance or all FSB instances.

10.2 Clearing Statistics on FIP Packets

Context

NOTICE

After the **reset fcoe fip statistics** command is executed, statistics on FIP packets and the latest ENode logout information are cleared and cannot be restored. Confirm your action before you use this command.

Procedure

- Run the **reset fcoe fip statistics** [*fc-instance-name*] command to clear statistics on Fibre channel Initialization Protocol (FIP) packets of the specified FIP snooping bridge (FSB) instance or all FSB instances.
- Run the **reset fcoe last-offline** command to clear the latest ENode logout information of FSB instance.

----End

10.3 Displaying Statistics on FC and FIP Packets

Context

After the Fibre Channel over Ethernet (FCoE) configuration is complete, run the following command in any view to view statistics on Fibre Channel (FC) and Fibre channel Initialization Protocol (FIP) packets.

Procedure

- Run the **display fcoe fc statistics** { **interface fcoe-port** *fcoe-port-id* | **brief** } command to view statistics on FC packets on a specified FCoE interface or all FCoE interfaces of FCF instances.

- Run the **display fcoe virtual-link statistics** { **interface fcoe-port** *fcoe-port-id* | **brief** } command to view statistics on FIP packets on a specified FCoE interface or all FCoE interfaces of FCF instances.

----End

10.4 Clearing Statistics on FC and FIP Packets

Context

NOTICE

After the following commands are executed, statistics on FC and FIP packets are cleared and cannot be restored. Exercise caution when you use the commands.

Procedure

- Run the **reset fcoe fc statistics** [**interface fcoe-port** *fcoe-port-id*] command to clear statistics on FC packets of FCoE forwarder (FCF) instances.
- Run the **reset fcoe virtual-link statistics** [**interface fcoe-port** *fcoe-port-id*] command to clear statistics on FIP packets of FCF instances.

----End

10.5 Monitoring the FCoE Running Status

Context

In routine maintenance, you can run the following commands in any view to view the Fibre Channel over Ethernet (FCoE) running status.

Procedure

- Run the **display fcoe last-offline** command to view the latest ENode logout information of FIP snooping bridge (FSB) instances.
- Run the **display fcoe session** [*fcoe-mac*] [**enode** *enode-mac* | **fcf** *fcf-mac* | **vlan** *vlan-id* | **interface** *interface-type interface-number*] * command to view the FCoE session information of FSB instances.
- Run the **display fcoe port** { **brief** | **instance** *fc-instance-name* | **interface fcoe-port** *fcoe-port-id* } command to view the FCoE interface configuration of FCF instances.
- Run the **display fcoe instance** [**fcf** [*fc-instance-name*]] command to view the FCF instance information.
- Run the **display fcoe name-server** { **brief** | **interface fcoe-port** *fcoe-port-id* } command to view the ENode register information on FCoE interfaces of FCF instances.

- Run the **display fcoe link-status** { **instance** *instance-name* | **interface fcoe-port** *fcoe-port-id* | **brief** } { **online** | **offline** } command to view the ENode registration information on an FCoE link of FCF instances.

----End

11 Configuration Examples

About This Chapter

This section provides Fibre Channel over Ethernet (FCoE) configuration examples.

11.1 Example for Configuring an FCF (FCoE Interfaces)

Networking Requirements

During deployment of a converged data center network, the access network convergence solution is used to protect investments on the existing storage area network (SAN). FCF functions as the access switch and connects to SwitchA and **Array** through its uplink interfaces, and directly connects to ServerA through its downlink interface, as shown in **Figure 11-1**. ServerA connects to FCFA and FCFB to constitute dual-plane networking, ensuring network reliability.

Figure 11-1 FCF networking

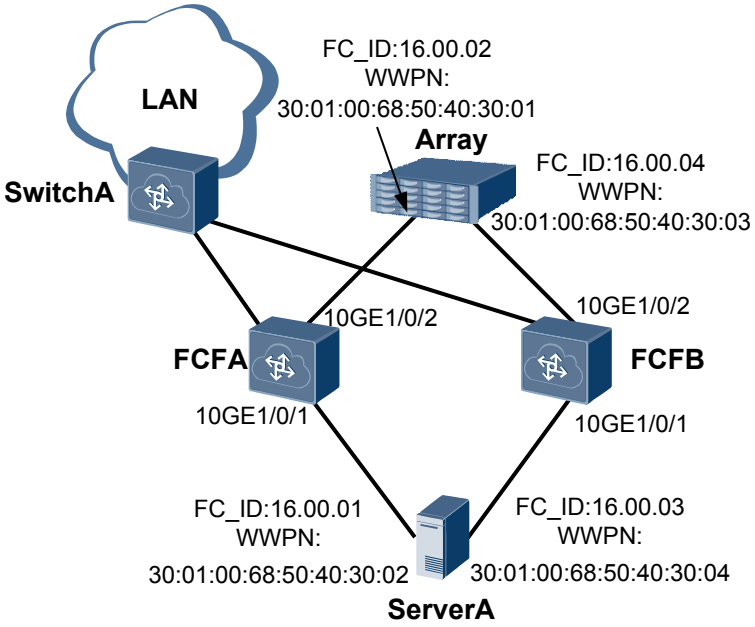

NOTE

- The device supports a maximum of 512 FCoE interfaces and allows a maximum of 4096 VN_Ports to log in.
- In each zone, 32 VN_Ports can be online concurrently by default, and a maximum of 64 VN_Ports can be online.

Configuration Roadmap

The configuration roadmap is as follows:

1. Create a Fibre Channel over Ethernet (FCoE) interface.
2. Create an FC instance, specify an FCoE VLAN, and add an FCoE interface to the Fiber Channel (FC) instance.
3. Add a physical interface to the FCoE interface and FCoE VLAN.
4. Enable the hard zoning function.
5. Configure a zone and add ServerA and **Array** to the zone.

6. Configure and activate a zone set.

NOTE

Configure Data Center Bridging (DCB). After DCB negotiation is successful, configure the FCoE forwarder (FCF).

Data Preparation

The configuration roadmap is as follows:

Device Name	Interface Number	WWPN	FC_ID	Interconnected Device
FCFA	FCoE-Port1	-	-	ServerA
	FCoE-Port2	-	-	Array
FCFB	FCoE-Port1	-	-	ServerA
	FCoE-Port2	-	-	Array
ServerA	-	30:00:00:68:50:40:30:02	16.00.01	FCFA
	-	30:00:00:68:50:40:30:04	16.00.03	FCFB
Array	-	30:00:00:68:50:40:30:01	16.00.02	FCFA
	-	30:00:00:68:50:40:30:03	16.00.04	FCFB

Procedure

Step 1 Create an FCoE interface.

Create FCoE-Port 1 and FCoE-Port 2 on FCFA.

```
<HUAWEI> system-view
[~HUAWEI] sysname FCFA
[*HUAWEI] commit
[~FCFA] interface fcoe-port 1
[*FCFA-FCoE-Port1] quit
[*FCFA] interface fcoe-port 2
[*FCFA-FCoE-Port2] quit
[*FCFA] commit
```

Step 2 Create an FC instance.

Create an FC instance named **fcf1** on FCFA, specify VLAN 2094 as the FCoE VLAN, and add FCoE-Port 1 and FCoE-Port 2 to **fcf1**.

```
[~FCFA] fcoe fcf1 fcf
[*FCFA-fcoe-fcf-fcf1] vlan 2094
[*FCFA-fcoe-fcf-fcf1] member interface fcoe-port 1
[*FCFA-fcoe-fcf-fcf1] member interface fcoe-port 2
[*FCFA-fcoe-fcf-fcf1] quit
[*FCFA] commit
```

Step 3 Configure a physical interface.

Configure 10GE1/0/1 on FCFA to allow VLAN 2094.

```
[~FCFA] interface 10ge 1/0/1
[~FCFA-10GE1/0/1] port link-type trunk
[*FCFA-10GE1/0/1] port trunk allow-pass vlan 2094
```

Add 10GE1/0/1 on FCFA to FCoE-Port 1.

```
[*FCFA-10GE1/0/1] fcoe-port 1
[*FCFA-10GE1/0/1] quit
[*FCFA] commit
```

Configure 10GE1/0/2 on FCFA to allow VLAN 2094.

```
[~FCFA] interface 10ge 1/0/2
[~FCFA-10GE1/0/2] port link-type trunk
[*FCFA-10GE1/0/2] port trunk allow-pass vlan 2094
```

Add 10GE1/0/2 on FCFA to FCoE-Port 2.

```
[*FCFA-10GE1/0/2] fcoe-port 2
[*FCFA-10GE1/0/2] quit
[*FCFA] commit
```

Step 4 Enable the hard zoning function.

Enable the hard zoning function on FCFA.

```
[~FCFA] zone hard-zoning enable
[*FCFA] commit
```

Step 5 Configure a zone.

Create a zone named **Zone1** on FCFA and add ServerA and **Array** to **Zone1**. You can add member nodes to a zone using the WWPN, WWNN, FC_ID, and FC/FCoE interface. The WWPN mode is recommended.

```
[~FCFA] zone zone1
[*FCFA-zone-zone1] member wwpn 30:00:00:68:50:40:30:01
[*FCFA-zone-zone1] member wwpn 30:00:00:68:50:40:30:02
[*FCFA-zone-zone1] quit
[*FCFA] commit
```

Step 6 Configure and activate a zone set.

Create a zone set named **Zoneset1** on FCFA and add **Zone1** to **Zoneset1**.

```
[~FCFA] zoneset zoneset1
[*FCFA-zoneset-zoneset1] member zone1
[*FCFA-zoneset-zoneset1] quit
[*FCFA] commit
```

Activate **Zoneset1** in the FCF instance **fcf1** on FCFA.

```
[~FCFA] fcoe fcf1 fcf
[~FCFA-fcoe-fcf-fcf1] active zoneset zoneset1
[*FCFA-fcoe-fcf-fcf1] quit
[*FCFA] commit
```

Step 7 Verify the configuration.

Check registration information of the server and storage device.

```
[~FCFA] display fcoe name-server brief
The Name-Server Information:
-------------------------------------------------------------------------------
Interface           FC-ID             WWPN
```

```
--------------------------------------------------------------------------------
FCoE1/0/1           16.00.01          30:00:00:68:50:40:30:02
FCoE1/0/2           16.00.02          30:00:00:68:50:40:30:01
--------------------------------------------------------------------------------
Total: 2
```

Step 8 Configure FCFB.

Repeat steps 1 to 7. For details, see the configuration file of FCFB.

----End

Configuration Files

- Configuration file of FCFA

```
#
sysname FCFA
#
zone hard-zoning enable
#
zone
zone1
 member wwpn 30:00:00:68:50:40:30:01
 member wwpn
30:00:00:68:50:40:30:02
#

zoneset
zoneset1
 member
zone1
#
fcoe fcf1
fcf
 vlan
2094
 member interface FCoE-
Port1
 member interface FCoE-Port2
 active zoneset zoneset1
#
interface
10GE1/0/1
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
1
#
interface
10GE1/0/2
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
2
#
interface FCoE-Port1
#
interface FCoE-Port2
#
return
```

- Configuration file of FCFB

```
#
sysname FCFB
```

```
#
zone hard-zoning enable
#
zone
zone1
 member wwpn 30:00:00:68:50:40:30:03
 member wwpn
30:00:00:68:50:40:30:04
#

zoneset
zoneset1
 member
zone1
#
fcoe fcf1
fcf
 vlan
2094
 member interface FCoE-
Port1
 member interface FCoE-Port2
 active zoneset zoneset1
#
interface
10GE1/0/1
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
1
#
interface
10GE1/0/2
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
2
#
interface FCoE-Port1
#
interface FCoE-Port2
#
return
```

11.2 Example for Configuring the FCF and FSB (FCoE interfaces)

Networking Requirements

During deployment of a converged data center network, the access network convergence solution is used to protect investments on the existing SAN. The FCoE forwarder (FCF) functions only as the aggregation device because of its high cost. The FSB used as the access switch directly connects to the switch and FCF through its uplink interfaces, and directly connects to servers through its downlink interfaces, as shown in **Figure 11-2**. ServerA connects to FSBA and FSBB to constitute dual-plane networking, ensuring network reliability.

Figure 11-2 Networking of the FCF and FSB

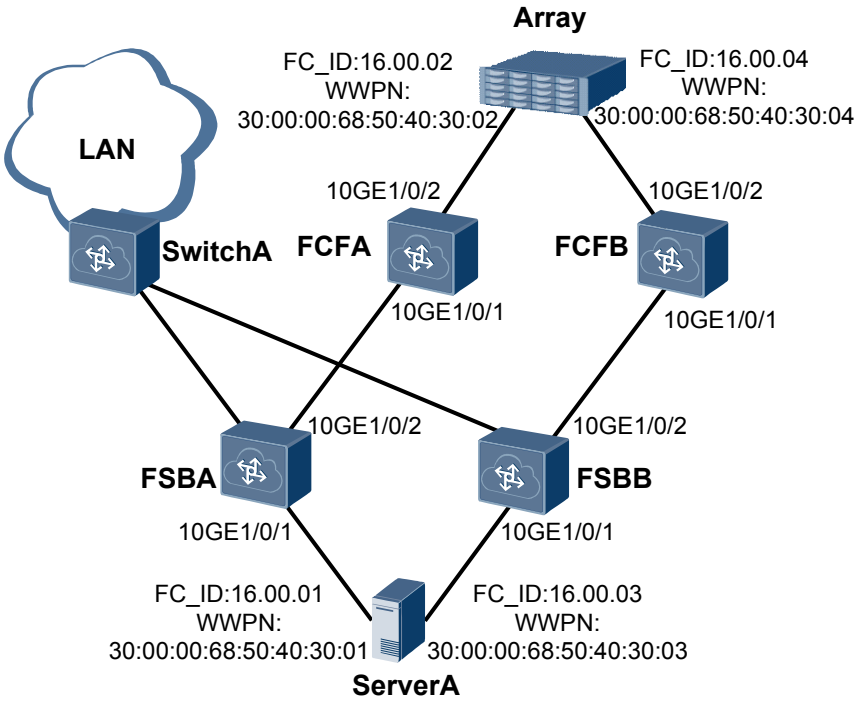

NOTE

- The device supports a maximum of 512 FCoE interfaces and allows a maximum of 4096 VN_Ports to log in.
- In each zone, 32 VN_Ports can be online concurrently by default, and a maximum of 64 VN_Ports can be online.

Configuration Roadmap

The configuration roadmap is as follows:

1. Create a Fibre Channel over Ethernet (FCoE) interface on FCF.
2. On the FCF, create an FCF instance, specify an FCoE VLAN, and add the FCoE interface to the Fiber Channel (FC) instance.
3. On the FCF, add a physical interface to the FCoE interface and FCoE VLAN.
4. Enable the hard zoning function.
5. Configure a zone on the FCF and add ServerA and Array to the zone.
6. Configure and activate a zone set.
7. Configure an FIP snooping bridge (FSB) instance and specify an FCoE VLAN on FSB.
8. On FSB, configure 10GE1/0/2 as a VNP_Port, and 10GE1/0/1 as ENode_Facing ports.

NOTE

Configure Data Center Bridging (DCB). After DCB negotiation is successful, configure the FCF and FSB.

Data Preparation

The configuration roadmap is as follows:

Device Name	Interface Number	WWPN	FC_ID	Interconnected Device
FCFA	FCoE-Port1	-	-	FSBA
	FCoE-Port2	-	-	Array
FCFB	FCoE-Port1	-	-	FSBB
	FCoE-Port2	-	-	Array
FSBA	10GE1/0/1	-	-	ServerA
	10GE1/0/2	-	-	FCFA
FSBB	10GE1/0/1	-	-	ServerB
	10GE1/0/2	-	-	FCFB
ServerA	-	30:00:00:68:50:40:30:01	16.00.01	FCFA
	-	30:00:00:68:50:40:30:03	16.00.03	FCFB
Array	-	30:00:00:68:50:40:30:02	16.00.02	FCFA
	-	30:00:00:68:50:40:30:04	16.00.04	FCFB

Procedure

Step 1 Create an FCoE interface on FCFA.

Create FCoE-Port 1 and FCoE-Port 2 on FCFA.

```
<HUAWEI> system-view
[~HUAWEI] sysname FCFA
[*HUAWEI] commit
[~FCFA] interface fcoe-port 1
[*FCFA-FCoE-Port1] quit
[*FCFA] interface fcoe-port 2
[*FCFA-FCoE-Port2] quit
[*FCFA] commit
```

Step 2 Create an FCF instance on FCFA.

Create an FC instance named **fcf1** on FCFA, specify VLAN 2094 as the FCoE VLAN, and add FCoE-Port 1 and FCoE-Port 2 to **fcf1**.

```
[~FCFA] fcoe fcf1 fcf
[*FCFA-fcoe-fcf-fcf1] vlan 2094
[*FCFA-fcoe-fcf-fcf1] member interface fcoe-port 1
[*FCFA-fcoe-fcf-fcf1] member interface fcoe-port 2
[*FCFA-fcoe-fcf-fcf1] quit
[*FCFA] commit
```

Step 3 Configure a physical interface on FCFA.

Configure 10GE1/0/1 on FCFA to allow VLAN 2094.

```
[~FCFA] interface 10ge 1/0/1
[~FCFA-10GE1/0/1] port link-type trunk
[*FCFA-10GE1/0/1] port trunk allow-pass vlan 2094
```

Add 10GE1/0/1 to FCoE-Port 1.

```
[*FCFA-10GE1/0/1] fcoe-port 1
[*FCFA-10GE1/0/1] quit
[*FCFA] commit
```

Configure 10GE1/0/2 on FCFA to allow VLAN 2094.

```
[~FCFA] interface 10ge 1/0/2
[~FCFA-10GE1/0/2] port link-type trunk
[*FCFA-10GE1/0/2] port trunk allow-pass vlan 2094
```

Add 10GE1/0/2 to FCoE-Port 2.

```
[*FCFA-10GE1/0/2] fcoe-port 2
[*FCFA-10GE1/0/2] quit
[*FCFA] commit
```

Step 4 Enable the hard zoning function.

Enable the hard zoning function on FCFA.

```
[~FCFA] zone hard-zoning enable
[*FCFA] commit
```

Step 5 Configure a zone.

Create a zone named **Zone1** on FCFA and add ServerA and Array to **Zone1**. You can add member nodes to a zone using the WWPN, WWNN, FC_ID, and FCoE interface. The WWPN mode is recommended.

```
[~FCFA] zone zone1
[*FCFA-zone-zone1] member wwpn 30:00:00:68:50:40:30:01
[*FCFA-zone-zone1] member wwpn 30:00:00:68:50:40:30:02
[*FCFA-zone-zone1] quit
[*FCFA] commit
```

Step 6 Configure and activate a zone set.

Create a zone set named **Zoneset1** on FCFA and add **Zone1** to **Zoneset1**.

```
[~FCFA] zoneset zoneset1
[*FCFA-zoneset-zoneset1] member zone1
[*FCFA-zoneset-zoneset1] quit
[*FCFA] commit
```

Activate **Zoneset1** in the FCF instance **fcf1**.

```
[~FCFA] fcoe fcf1 fcf
[~FCFA-fcoe-fcf-fcf1] active zoneset zoneset1
[*FCFA-fcoe-fcf-fcf1] quit
[*FCFA] commit
```

Step 7 Configure an FSB instance on FSBA.

Configure an FC instance **FSB** on FSBA and specify VLAN 2094 as the FCoE VLAN.

```
<HUAWEI> system-view
[~HUAWEI] sysname FSBA
[*HUAWEI] commit
[~FSBA] fcoe FSB
[*FSBA-fcoe-fsb-FSB] vlan 2094
[*FSBA-fcoe-fsb-FSB] commit
[~FSBA-fcoe-fsb-FSB] quit
```

Step 8 Configure interface roles on FSBA.

Configure 10GE1/0/2 on FSBA to allow VLAN 2094. The configuration of 10GE1/0/1 is similar to the configuration of 10GE1/0/2, and are not mentioned here.

```
[~FSBA] interface 10ge 1/0/2
[~FSBA-10GE1/0/2] port link-type trunk
[*FSBA-10GE1/0/2] port trunk allow-pass vlan 2094
```

Configure 10GE1/0/2 on FSBA as the VNP_Port.

```
[*FSBA-10GE1/0/2] fcoe role vnp
[*FSBA-10GE1/0/2] commit
[~FSBA-10GE1/0/2] quit
```

An interface is an ENode_Facing port by default, so the configurations of 10GE1/0/1 is not mentioned.

Step 9 Verify the configuration.

On FCFA, check registration information of the server and storage device.

```
[~FCFA] display fcoe name-server brief
The Name-Server Information:
-------------------------------------------------------------------------------
Interface           FC-ID             WWPN
-------------------------------------------------------------------------------
FCoE1/0/1           16.00.01          30:00:00:68:50:40:30:01
FCoE1/0/2           16.00.02          30:00:00:68:50:40:30:02
-------------------------------------------------------------------------------
Total: 3
```

Check the FCoE VLAN configuration on FSBA.

```
[~FSBA] display fcoe vlan
-------------------------------------------
VLAN     Interface           Interface Role
-------------------------------------------
2094     10GE1/0/1           ENode-Facing
         10GE1/0/2           VNP-Port
-------------------------------------------
```

Check the FCoE session on FSBA.

Run the **display fcoe session** command to check the FCoE session. If FIP snooping is successfully configured, FCoE session information is generated.

```
[~FSBA] display fcoe session
FCoE Session Information:
-------------------------------------------------------------------------
FCoE MAC        ENode MAC       ENode Interface   VLAN
FCID            FCF MAC         FCF Interface     Keep Alive Packet(s)
-------------------------------------------------------------------------
0efc-0010-06ab  0010-9400-0009  10GE1/0/1         2094
0x1006ab        200b-c723-4201  10GE1/0/2         5
-------------------------------------------------------------------------
Total: 1
```

Step 10 Configure FCFB and FSBB.

Repeat steps 1 to 9 on FCFB and FSBB. For details, see the configuration files of FCFB and FSBB.

----End

Configuration Files

- Configuration file of FCFA

```
#
sysname FCFA
```

```
#
zone hard-zoning enable
#
zone
zone1
 member wwpn 30:00:00:68:50:40:30:01
 member wwpn 30:00:00:68:50:40:30:02
#

zoneset
zoneset1
 member
zone1
#
fcoe fcf1
fcf
 vlan
2094
 member interface FCoE-
Port1
 member interface FCoE-Port2
 active zoneset zoneset1
#
interface
10GE1/0/1
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
1
#
interface
10GE1/0/2
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
2
#
interface FCoE-Port1
#
interface FCoE-Port2
#
return
```

- Configuration file of FSBA

```
#
sysname FSBA
#
fcoe FSB
 vlan 2094
#
interface 10GE1/0/1
 port link-type trunk
 port trunk allow-pass vlan 2094
#
interface 10GE1/0/2
 port link-type trunk
 port trunk allow-pass vlan 2094
 fcoe role vnp
#
return
```

- Configuration file of FCFB

```
#
sysname FCFB
#
zone hard-zoning enable
```

```
#
zone
zone1
 member wwpn 30:00:00:68:50:40:30:03
 member wwpn 30:00:00:68:50:40:30:04
#

zoneset
zoneset1
 member
zone1
#
fcoe fcf1
fcf
 vlan
2094
 member interface FCoE-
Port1
 member interface FCoE-Port2
 active zoneset zoneset1
#
interface
10GE1/0/1
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
1
#
interface
10GE1/0/2
 port link-type
trunk
 port trunk allow-pass vlan
2094
 fcoe-port
2
#
interface FCoE-Port1
#
interface FCoE-Port2
#
return
```

- Configuration file of FSBB

```
#
sysname FSBB
#
fcoe FSB
 vlan 2094
#
interface 10GE1/0/1
 port link-type trunk
 port trunk allow-pass vlan 2094
#
interface 10GE1/0/2
 port link-type trunk
 port trunk allow-pass vlan 2094
 fcoe role vnp
#
return
```

12 Introduction to DCB

Data Center Bridging (DCB) is used to build lossless Ethernet, meeting quality of service (QoS) requirements on a converged data center network.

On a converged data center network, local area network (LAN) traffic, storage area network (SAN) traffic, and inter-process communication (IPC) traffic have different QoS requirements:

- SAN traffic is sensitive to packet loss and relies on in-order delivery, which means that packets are delivered in the same order in which they were sent.
- LAN traffic allows packet loss and is delivered on a best-effort (BE) basis.
- IPC traffic is exchanged between servers and requires low latency.

A converged network has high requirements for link sharing and common Ethernet cannot meet the preceding requirements.

Data Center Bridging (DCB) is a set of extensions to Ethernet for use in a data center environment, which is defined by the IEEE 802.1 working group. DCB is used to build lossless Ethernet, meeting QoS requirements on a converged data center network.

Related Documents

Video: **HomeVideo LibraryHuawei CloudEngine Series Switches FCoE&DCB Feature Introduction**

13 DCB Principles

About This Chapter

This section describes the implementation of DCB.

Data Center Bridging (DCB) is a set of extensions to Ethernet for use in a data center environment, which is defined by the IEEE 802.1 working group. DCB is used to build lossless Ethernet, meeting QoS requirements on a converged data center network.

DCB feature include:

13.1 PFC

Background

SAN traffic is sensitive to packet loss on a converged network.

The Ethernet Pause mechanism ensures the lossless transmission service. When a downstream device detects that its receive capability is lower than the transmit capability of its upstream device, it sends Pause frames to the upstream device, requesting the upstream device to stop sending traffic for a period of time. The Ethernet Pause mechanism stops all traffic on a link, whereas FCoE requires link sharing:

- Burst traffic of one type cannot affect forwarding of traffic of other types.
- A large amount of traffic of one type in a queue cannot occupy buffer resources of traffic of other types.

Priority-based Flow Control (PFC) addresses the contradiction between the Ethernet Pause mechanism and link sharing.

Principles

PFC is also called Per Priority Pause or Class Based Flow Control (CBFC). It enhances the Ethernet Pause mechanism. PFC is a kind of flow control mechanism based on priority. As shown in **Figure 13-1**, eight priority queues on the transmit interface of DeviceA correspond to eight buffers on the receive interface of DeviceB. When a receive buffer on DeviceB is congested, DeviceB sends a backpressure signal to DeviceA, requesting DeviceA to stop sending packets in the corresponding priority queue.

Figure 13-1 PFC working mechanism

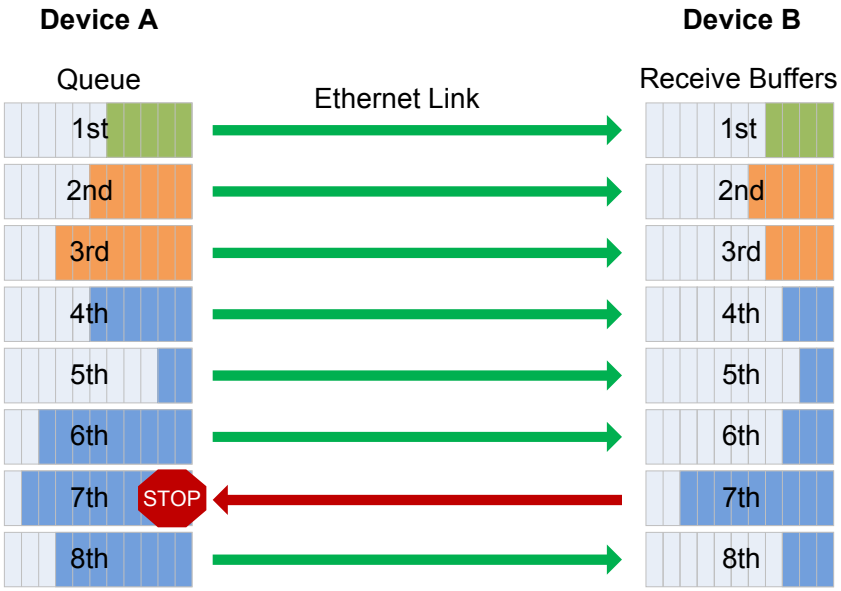

Table 13-1 describes the mappings between packet priorities and interface queues.

Table 13-1 Mappings between packet priorities and interface queues

Packet Type	Priority	Queue
Unicast	0	0
	1	1
	2	2
	3	3
	4	4
	5	5
	6	6
	7	7

A backpressure signal is an Ethernet frame. **Figure 13-2** shows the PFC frame format.

Figure 13-2 PFC frame format

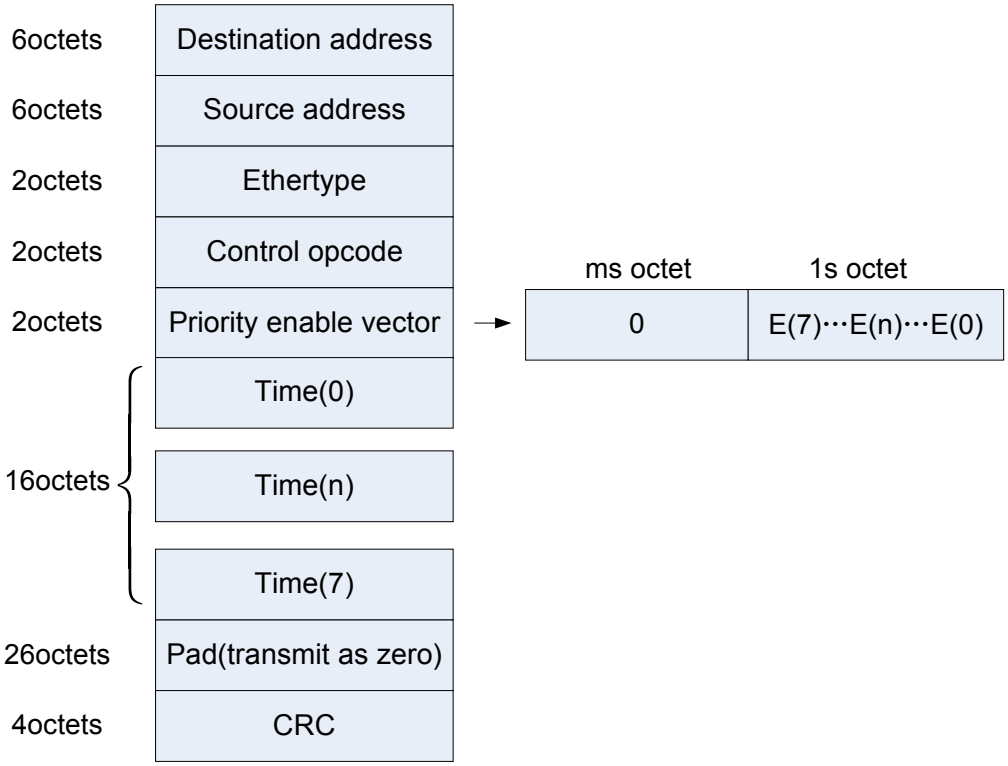

Table 13-2 Fields in a PFC frame

Item	Description
Destination address	Destination MAC address, which has a fixed value of 01-80-c2-00-00-01.

Item	Description
Source address	Source MAC address.
Ethertype	Ethernet frame type. The value is 88-08.
Control opcode	Control code. The value is 01-01.
Priority enable vector	Priority-enable vector. E(*n*) corresponds to queue *n* and determines whether backpressure is enabled for queue *n*. When E(*n*) is 1, backpressure is enabled for queue *n* and the backpressure time is Time(n). When E(*n*) is 0, backpressure is disabled for queue *n*.
Time(0) to Time(7)	Backpressure timer. If Time(n) is 0, backpressure is canceled.
Pad(transmit as zero)	Reserved. The value is 0 during PFC frame transmission.
CRC	Cyclic Redundancy Check (CRC).

When receiving backpressure signals, a device only stops traffic in one or several priority queues, but does not stop traffic on the entire interface. PFC can pause or restart any queue, without interrupting traffic in other queues. This feature enables traffic of various types to share one FCoE link. The system does not apply the backpressure mechanism to the priority queues with PFC disabled and directly discards packets in these queues when congestion occurs.

In an FCoE environment, an administrator can apply PFC to queues of FCoE traffic to ensure lossless transmission of FCoE service.

13.2 ETS

Background

A converged data center network has LAN traffic, SAN traffic, and IPC traffic. The converged network has high QoS requirements. For example, SAN traffic is sensitive to packet loss and relies on in-order deliver. IPC traffic is exchanged between servers and requires low latency. LAN traffic allows packet loss and is delivered on a best-effort (BE) basis. Traditional QoS cannot meet requirements of the converged network, whereas ETS uses a hierarchical flow control mechanism to implement QoS on the lossless Ethernet.

Principles

ETS provides two-level scheduling: scheduling based on priority groups and scheduling based on queues, as shown in **Figure 13-3**. An interface first schedules priority groups, and then schedules priority queues.

Figure 13-3 ETS Process

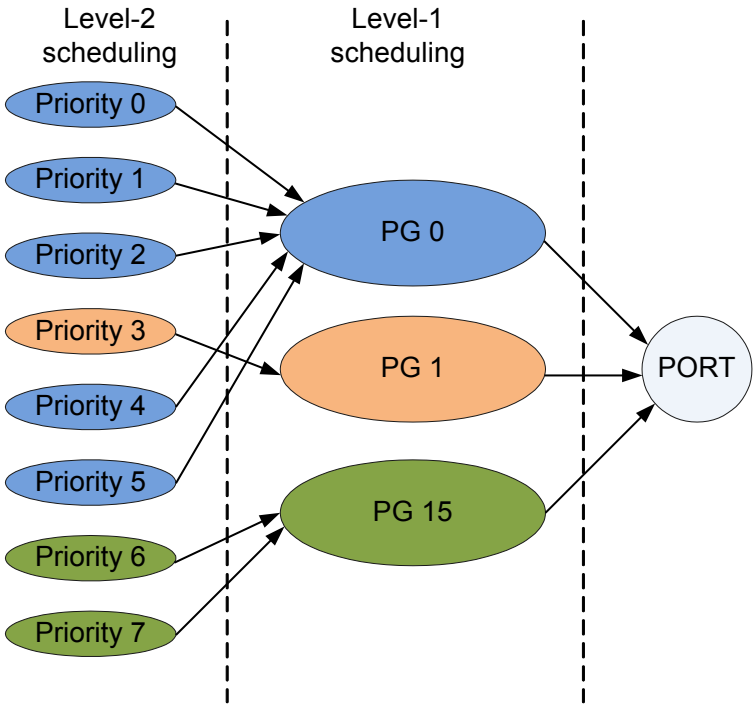

Compared with common QoS, ETS provides scheduling based on priority groups. ETS adds traffic of the same type to a priority group so that traffic of the same type obtains the same CoS.

Scheduling Based on Priority Groups

A priority group is a group of priority queues using the same scheduling mode. You can add queues with different priorities to a priority group. Scheduling based on the priority group is called level-1 scheduling.

ETS defines three priority groups by default: PG0, PG1, and PG15. PG0, PG1, and PG15 process LAN traffic, SAN traffic, and IPC traffic respectively.

Table 13-3 describes attributes of priority groups by default.

Table 13-3 Scheduling based on priority groups

Priority Group ID	Priority Queue	Scheduling Mode	Bandwidth Usage	Traffic Type
PG0	0, 1, 2, 4, 5	WFQ	50%	LAN flow
PG1	3	WFQ	50%	SAN flow
PG15	6, 7	PQ	-	IPC flow

As defined by ETS,PG0 and PG1 use WFQ and PG15 use PQ. PG15 uses priority queuing (PQ) to schedule delay-sensitive IPC traffic. PG0 and PG1 use Weighted Fair Queue (WFQ). Bandwidth can be allocated to priority groups based on actual networking.

NOTE

The scheduling mode of each priority group cannot be changed.

As shown in **Table 13-3**, the queue with priority 3 carries FCoE traffic, so this queue is added to the SAN group (PG1). Queues with priorities 0, 1, 2, 4, and 5 carry LAN traffic, so these queues are added to the LAN group (PG0). The queue with priorities 6, 7 carries IPC traffic, so this queue is added to the IPC group (PG15). The total bandwidth of the interface is 10 Gbit/s. PG15 obtains 2 Gbit/s. PG1 and PG0 each obtain 50% of the total bandwidth, 4 Gbit/s.

Figure 13-4 Congestion management based on priority groups

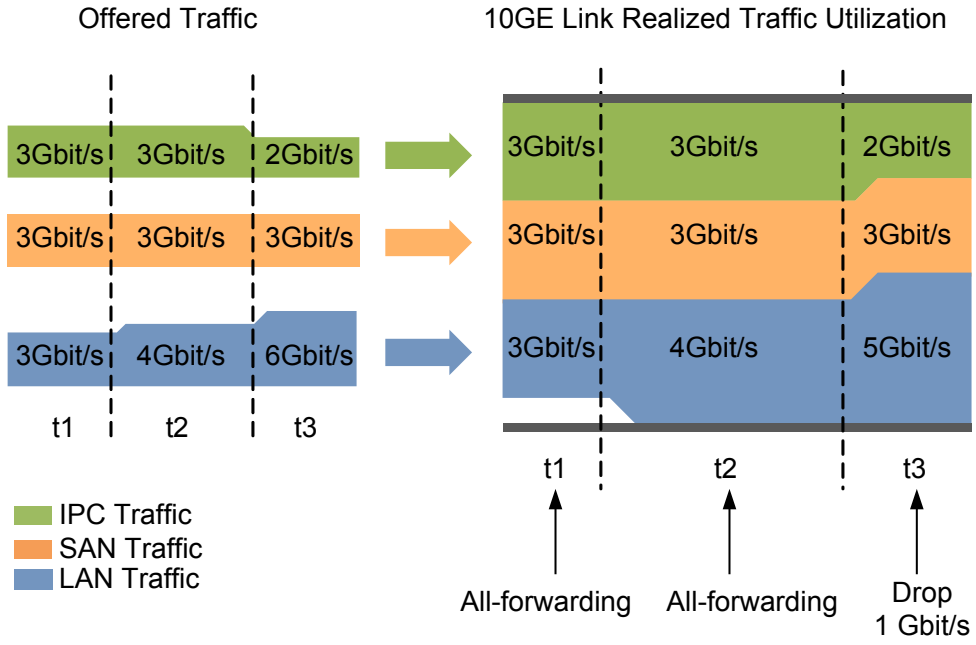

As shown in **Figure 13-4**, all traffic can be forwarded because the total traffic on the interface is within the interface bandwidth at t1 and t2. At t3, the total traffic exceeds the interface bandwidth and LAN traffic exceeds given bandwidth. At this time, LAN traffic is scheduled based on ETS parameters and 1 Gbit/s LAN traffic is discarded.

ETS also provides traffic shaping based on priority groups. This traffic shaping mechanism limits traffic bursts in a priority group to ensure that traffic in this group is sent out at an even rate. For details, see Traffic Shaping in *CloudEngine 12800 Series Switches Configuration Guide - QoS Configuration*.

Priority-based Scheduling

ETS also provides priority-based queue scheduling, level-2 scheduling.

In addition, ETS provides priority-based queue congestion management, queue shaping, and queue congestion avoidance. For details, see *CloudEngine 12800 Series Switches Configuration Guide - QoS Configuration*.

13.3 DCBX

Background

To implement lossless Ethernet on a converged data center network, both ends of an FCoE link must have the same PFC and ETS parameter settings. If PFC and ETS parameters are manually configured, the administrator's workload is heavy and configuration errors may occur. DCBX, as a link discovery protocol, enables DCB devices at both ends of a link to discover and exchange DCB configurations, reducing workloads of administrators.

Principles

DCBX provides the following functions:

- Discovers the DCB configuration of the remote device.
- Detects the DCB configuration errors of the remote device.
- Configures the remote device if permitted.

DCBX enables DCB devices at both ends to exchange the following DCB configurations:

- ETS priority group
- PFC

DCBX encapsulates DCB configurations into Link Layer Discovery Protocol (LLDP) TLVs so that devices at both ends of an FCoE virtual link can exchange DCB configurations. For details about LLDP, see LLDP Configuration in *CloudEngine 12800 Series Switches Configuration Guide - Network Management Configuration*.

In **Figure 13-5**, PFC is used as an example to describe DCBX implementation through LLDP.

Figure 13-5 DCBX implementation through LLDP

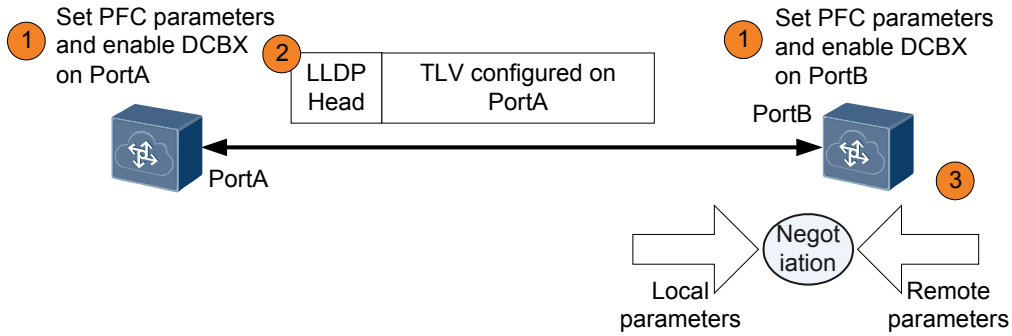

As shown in **Figure 13-5**, LLDP is enabled globally and on PortA and PortB, and PortA is configured to send DCBX TLVs. The implementation is as follows:

1. Set PFC parameters on PortA and PortB, and enable DCBX. The DCBX module instructs PortA and PortB to encapsulate their PFC parameters into LLDPDUs and send the LLDPDUs to each other.
2. The LLDP module of PortA sends LLDPDUs with DCBX TLVs to PortB at intervals.
3. PortB parses the DCBX TLVs in the received LLDPDUs and sends PFC parameters of PortA to the DCBX module. The DCBX module compares PFC parameters of PortA with its PFC parameters. Through negotiation, PFC parameters on the two ends are consistent, and a configuration file is then generated.

DCBX TLV

As shown in **Figure 13-6**, the DCB configuration is encapsulated into specified TLVs. The Type field has a fixed value of 127, and the OUI field varies depending on the protocol type. The OUI field of the IEEE DCBX is 0x0080c2, and the OUI field of the INTEL DCBX is 0x001b21.

Figure 13-6 DCBX TLV format

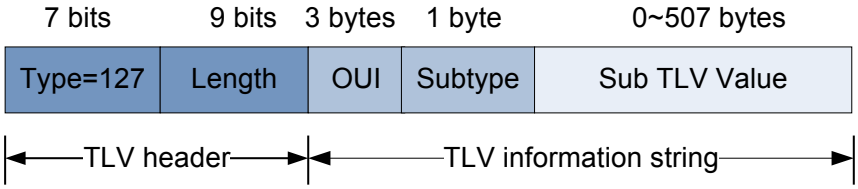

DCBX TLVs include the ETS Configuration TLV, ETS Recommendation TLV, PFC Configuration TLV and App TLV. **Table 13-4** describes the DCBX TLVs.

Table 13-4 IEEE DCBX TLVs

TLV	Subtype	Length	Description
ETS Configuration TLV	09	25	Local ETS configuration: ● Priority group configuration: priority group ID and bandwidth usage of a priority group ● Priority queue configuration: priority queue ID and its priority group ID
ETS Recommendation TLV	0A	25	Recommended ETS configuration, used for ETS configuration negotiation between both ends of an FCoE virtual link: ● Priority group configuration: priority group ID and bandwidth usage of a priority group ● Priority queue configuration: priority queue ID and its priority group ID
PFC Configuration TLV	0B	6	Local PFC configuration: ● Priority queue ID ● Whether PFC is applied to a queue
App TLV	0C	Unfixed value	Carried only when PFC is configured to work in auto mode for interconnection between products and between NICs.

DCBX TLVs include the ETS Configuration TLV, ETS Recommendation TLV, PFC Configuration TLV and App TLV. **Table 13-5** and **Table 13-6** describe the DCBX TLVs.

Table 13-5 INTEL DCBX v1.00 TLVs

TLV	Subtype	Length	Description
DCBX Control Sub-TLV	01	10	Information of DCBX packets.
Priority Group Sub-TLV	02	28	Recommended ETS configuration, used for ETS configuration negotiation between both ends of an FCoE virtual link: ● Bandwidth usage of a priority group ● Priority group ID
Priority Flow Control Sub-TLV	03	5	Local PFC configuration: ● Priority queue ID ● Whether PFC is applied to a queue

Table 13-6 INTEL DCBX v1.01 TLVs

TLV	Subtype	Length	Description
DCBX Control Sub-TLV	01	10	Information of DCBX packets.
Priority Group Sub-TLV	02	17	Recommended ETS configuration, used for ETS configuration negotiation between both ends of an FCoE virtual link: ● Priority group configuration: priority group ID and bandwidth usage of a priority group ● Priority queue configuration: priority queue ID and its priority group ID
Priority Flow Control Sub-TLV	03	6	Local PFC configuration: ● Priority queue ID ● Whether PFC is applied to a queue

14 Applications

This section describes the applicable scenario of DCB.

In data center network convergence scenarios, FCoE implements network convergence. To reduce investments, access and aggregation switches are often deployed between servers and FCFs. As shown in **Figure 14-1**, SwitchA is an access switch and needs to forward LAN, SAN, and IPC traffic; SwitchC is an aggregation switch and needs to forward SAN traffic. To guarantee QoS of LAN, SAN, and IPC traffic, DCB is configured on SwitchA and SwitchC. To ensure link reliability between ServerA and array, multiple links are deployed.

Figure 14-1 FSB networking

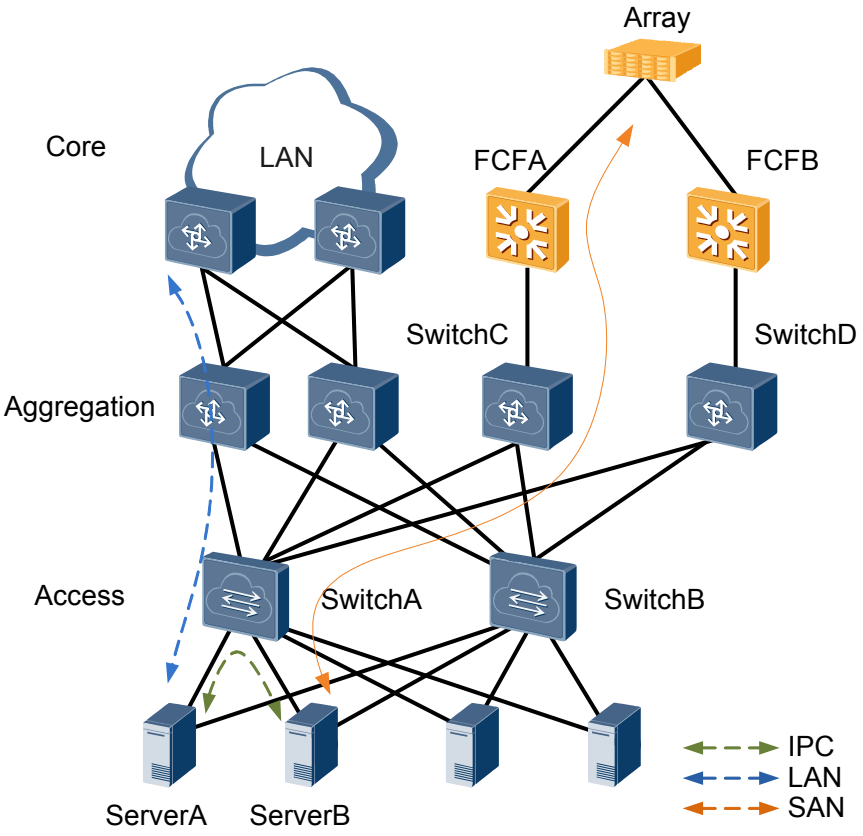

15 Configuration Task Summary

The device supports Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and Data Center Bridging eXchange protocol (DCBX).

Among the following DCB configuration tasks, **Configuring PFC** and **Configuring ETS** are mandatory and can be performed in any sequence. **Configuring DCBX** is optional. When PFC is configured to work in **auto** mode, perform the operation of **Configuring DCBX**.

Table 15-1 DCB configuration task summary

Scenario	Description	Task
Configure PFC	When traffic of a certain type is congested, PFC stops such traffic without interrupting traffic of other types.	**18 Configuring PFC**
Configure ETS	ETS provides flexible and hierarchical scheduling management to prevent service interference and conflicts after data center networks are converged.	**19 Configuring ETS**
Configure DCBX	DCBX allows Data Center Bridging (DCB) devices at both ends of a link to discover and exchange DCB configurations.	**20 Configuring DCBX**

16 Configuration Notes

This section describes the product models that support Data Center Bridging (DCB) and notes about configuring DCB.

Involved Network Element

Other network elements are required to support DCB functions.

License Support

DCB are basic features of a switch and are not under license control.

Version Support

Table 16-1 Products and minimum version supporting DCB

Series	Product	Minimum Version Required
CE12800	CE12804/CE12808/ CE12812	V100R003C00
	CE12816	V100R003C00
	CE12804S/CE12808S	V100R005C00

Feature Dependencies and Limitations

- If you configure both priority mapping and Priority-based Flow Control (PFC), exercise caution when modifying the mappings between the priorities of PFC-enabled queues and priorities in the DiffServ domain. Otherwise, PFC may not work.
- In V100R005C00 and later versions, when there are more than three priority groups, multiple priority groups share system resources. When a priority group has the remaining bandwidth, scheduling of the remaining bandwidth may be inaccurate.
- The Fibre Channel over Ethernet (FCoE) service takes effect only after DCB negotiation is successful.

- In V100R005C00 version, the device does not support in-service software upgrade (ISSU) for DCB.

17 Default Configuration

This section provides the default Data Center Bridging (DCB) configuration.

Table 17-1 lists the default DCB configuration.

Table 17-1 Default DCB Configuration

Parameter	Default Setting
PFC	Disable
ETS	Disable
DCBX	Disable

18 Configuring PFC

When a link is congested by traffic of a certain type, Priority-based Flow Control (PFC) controls transmission of such traffic without interrupting traffic of other types.

Pre-configuration Tasks

Before configuring PFC, complete the following tasks:

- Configuring a link layer protocol so that interfaces at both ends of a link are Up
- Completing **20 Configuring DCBX** before enabling PFC on the interface to set the PFC working mode to **auto**

NOTE

If you configure both priority mapping and PFC, exercise caution when modifying the mappings between the priorities of PFC-enabled queues and priorities in the DiffServ domain. Otherwise, PFC may not work.

The PFC priority of member interfaces in the same FC instance must be the same.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
dcb pfc [ pfc-profile-name ]
```

A PFC profile is created and the PFC profile view is displayed.

The system provides a default PFC profile **default**. This profile can be modified but cannot be deleted. You can create a new PFC profile or modify the default PFC profile according to service requirements.

Step 3 Run:

```
priority { start-priority [ to end-priority ] } &<1-8>
```

PFC for queues with a specified priority is enabled.

By default, PFC is enabled for priority 3.

Step 4 Run:

```
quit
```

Exit from the PFC view.

Step 5 Run:

```
interface interface-type interface-number
```

The interface view is displayed.

Step 6 Run:

```
dcb pfc enable [ pfc-profile-name ] [ mode { auto | manual } ]
```

A PFC profile is applied to the interface and the PFC working mode is specified.

By default, no PFC profile is applied to an interface.

Step 7 Run:

```
commit
```

The configuration is committed.

----End

Checking the Configuration

- Run the **display dcb** [**interface** *interface-type interface-number*] command to check the Data Center Bridging (DCB) configuration and negotiation status.
- Run the **display dcb pfc-profile** [*profile-name*] command to check the PFC profile configuration.

19 Configuring ETS

About This Chapter

Enhanced transmission selection (ETS) provides flexible, hierarchical scheduling management to prevent interference and conflicts between services after Ethernet and storage networks of a data center is converged.

Pre-configuration Tasks

Before configuring ETS, complete the following tasks:

- Configuring a link layer protocol so that interfaces at both ends of a link are Up
- Configuring Priority Mapping to map packet priorities to per-hop behaviors (PHBs)/ colors

19.1 Configuring an ETS Profile

Context

Enhanced transmission selection (ETS) provides two-level flow control:

- Flow control based on the priority group: congestion management and traffic shaping based on the priority group
- Flow control based on the priority: queue congestion management, queue shaping, and queue congestion avoidance

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
dcb ets-profile profile-name
```

An ETS profile is created and the ETS profile view is displayed.

The system provides a default ETS profile **default**. This profile can be modified but cannot be deleted. You can create a new ETS profile or modify the default ETS profile according to service requirements.

Step 3 (Optional) Run:

```
priority-group pg-value
```

A priority group is created.

By default, an ETS profile defines three priority groups: PG0, PG1, and PG15.

Step 4 Run:

```
priority-group pg-value queue { start-queue [  to end-queue ] } &<1-8>
```

The specified interface queues are added to a priority group.

By default, queues 0, 1, 2, 4, and 5 belong to PG0, queue 3 belongs to PG1, and queues 6 and 7 belong to PG15.

Step 5 Configure flow control based on the priority group.

- Run:
  ```
  priority-group pg-value wfq weight weight-value
  ```
  The Weighted Fair Queue (WFQ) weight of a priority group is set.

  By default, the weights of PG0 and PG1 are both 50.

  NOTE

  If there are multiple queues in a priority group, when traffic in queues all exceeds the bandwidth, scheduling is accurate. When some traffic in queues exceeds the bandwidth and the total traffic exceeds the interface bandwidth, scheduling is inaccurate.

- Run:
  ```
  priority-group pg-value shaping cir cir-value [ cbs cbs-value ]
  ```

Traffic shaping is enabled for a priority group and shaping parameters are set.

By default, traffic shaping is disabled for a priority group.

NOTE

During scheduling, traffic within the CIR in PG15 is preferentially scheduled, and then traffic in other priority groups is scheduled.

Step 6 Configure flow control based on the priority.

- Run:
  ```
  queue { start-queue [ to end-queue ] } &<1-8> shaping cir cir-value [ cbs cbs-value ]
  ```
  Traffic shaping is enabled for interface queues and shaping parameters are set.

  By default, traffic shaping is disabled for an interface queue.
- Run:
  ```
  queue { start-queue [ to end- queue ] } &<1-8> wred drop-profile-name
  ```
  A drop profile is bound to interface queues.

  By default, no drop profile is bound to an interface queue.

Step 7 Run:
```
commit
```

The configuration is committed.

----End

19.2 Applying an ETS Profile

Context

You can apply an enhanced transmission selection (ETS) profile to an interface so that the interface can provide differentiated services.

Procedure

Step 1 Run:
```
system-view
```

The system view is displayed.

Step 2 Run:
```
interface interface-type interface-number
```

The interface view is displayed.

Step 3 Run:
```
dcb ets enable ets-profile-name
```

An ETS profile is applied to the interface.

By default, no ETS profile is applied to an interface.

Step 4 Run:
```
commit
```

The configuration is committed.

----End

19.3 Checking the Configuration

Prerequisites

All the enhanced transmission selection (ETS) configurations are complete.

Procedure

- Run the **display dcb ets-profile** [*profile-name*] command to check the ETS profile configuration.
- Run the **display dcb** [**interface** *interface-type interface-number*] command to check the DCB configuration and negotiation status.

----End

20 Configuring DCBX

Data Center Bridging (DCB) allows DCB devices at both ends of a link to discover and exchange DCB configurations.

Context

Data Center Bridging eXchange protocol (DCBX) encapsulates DCB configurations into a Link Layer Discovery Protocol (LLDP) Type, Length, and Value (TLV) so that devices at both ends of an Fibre Channel over Ethernet (FCoE) virtual link can exchange DCB configurations.

When the device connects to a non-Huawei device, the service priorities in the APP TLVs must be the same.

Pre-configuration Tasks

Before configuring Data Center Bridging eXchange protocol (DCBX), complete the following task:

Configuring a link layer protocol so that interfaces at both ends of a link are Up.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
lldp enable
```

The Link Layer Discovery Protocol (LLDP) is enabled globally.

Step 3 Run:

```
interface interface-type interface-number
```

The interface view is displayed.

Step 4 Run:

```
lldp tlv-enable dcbx
```

The interface is configured to advertise DCBX TLVs.

Step 5 (Optional) Run:

```
dcb compliance intel-oui
```

The version number in the outgoing DCBX TLV is set.

Two standards are available for the DCBX protocol: IEEE DCBX and Intel DCBX.

By default, an interface sends DCBX TLVs of the IEEE standard version.

Step 6 Run:

```
quit
```

Exit from the interface view.

Step 7 (Optional) Run:

```
dcb app-profile profile-name
```

An APP profile is created and the APP profile view is displayed.

By default, no APP profile is configured.

Step 8 (Optional) Run:

```
application { fcoe | fip | iscsi | ethtype ethtype | [ tcp | udp ] port port-id }
priority priority-value
```

A service priority in the APP TLV of DCB packets is set.

By default, the priorities of Fibre Channel over Ethernet (FCoE) and Fibre channel Initialization Protocol (FIP) services are both 3, and the ISCSI service priority is 4.

NOTE

PFC must be enabled for queues with a priority; otherwise, DCB negotiation may fail.

Step 9 (Optional) Run:

```
quit
```

Exit from the APP profile view.

Step 10 (Optional) Run:

```
interface interface-type interface-number
```

The view of the interface to which the APP profile needs to be applied is displayed.

Step 11 (Optional) Run:

```
dcb app-profile enable app-profile-name
```

The APP profile is applied to the interface.

By default, no APP profile is applied to an interface.

Step 12 Run:

```
commit
```

The configuration is committed.

----End

Checking the Configuration

- Run the **display lldp tlv-config** [**interface** *interface-type interface-number*] command to check TLV types supported by the system or an interface.

- Run the **display dcb app-profile** [*profile-name*] command to check the APP profile configuration.

21 Maintaining DCB

About This Chapter

This section describes how to check and clear Data Center Bridging (DCB) statistics and monitor DCB running status.

21.1 Monitoring the DCB Running Status

Context

During routine maintenance, you can run the following commands in any view to monitor the Data Center Bridging (DCB) running status.

Procedure

- Run the **display dcb** [**interface** *interface-type interface-number*] command to check the DCB configuration and negotiation status.
- Run the **display dcb ets-profile** [*profile-name*] command to check the enhanced transmission selection (ETS) profile configuration.
- Run the **display dcb pfc-profile** [*profile-name*] command to check the Priority-based Flow Control (PFC) profile configuration.
- Run the **display dcb app-profile** [*profile-name*]command to check the APP profile configuration.
- Run the **display dcb pfc** [**interface** *interface-type interface-number*] command to check statistics about PFC frames.
- Run the **display dcb fail-record** [**interface** *interface-type interface-number*] command to check DCB negotiation failure records on an interface.

----End

21.2 Clearing DCB Statistics

Context

When diagnosing and locating Data Center Bridging (DCB) faults, collect DCB statistics on interfaces in a specified period. Before re-collecting DCB statistics on all interfaces or a specified interface, clear existing DCB statistics.

NOTICE

After the following reset commands are executed, DCB statistics on interfaces are cleared and cannot be restored. Confirm your action before you use these commands.

Procedure

- Run the **reset dcb pfc** [**interface** *interface-type interface-number*] command in the user view to clear Priority-based Flow Control (PFC) statistics on an interface.
- Run the **reset dcb fail-record** [**interface** *interface-type interface-number*] command in the user view to clear DCB negotiation failure records on an interface.

----End

22 Configuration Examples

About This Chapter

This section provides a Data Center Bridging (DCB) configuration example.

22.1 Example for Configuring DCB

Networking Requirements

During deployment of a converged data center network, the access network convergence solution is used to protect investments in the existing storage area network (SAN). The FCoE forwarder (FCF) functions only as the aggregation device because of its high cost. SwitchA functions as the access switch. SwitchA directly connects to Ethernet switch SwitchB and FCF through its uplink interfaces, and directly connects to servers through its downlink interfaces, as shown in **Figure 22-1**. On SwitchA, Fibre channel Initialization Protocol (FIP) snooping is configured to ensure correct SAN traffic forwarding.

SwitchA carries common Ethernet traffic, Fibre Channel over Ethernet (FCoE) traffic, and inter-process communication (IPC) traffic among server clusters. The 802.1p priorities of FCoE traffic and IPC traffic are 3 and 7 respectively. The quality of service (QoS) requirements of the three types of traffic are as follows:

- FCoE traffic requires no packet loss.
- IPC traffic requires low latency.
- Common Ethernet traffic is transmitted in best-effort (BE) mode.

Figure 22-1 Data center networking

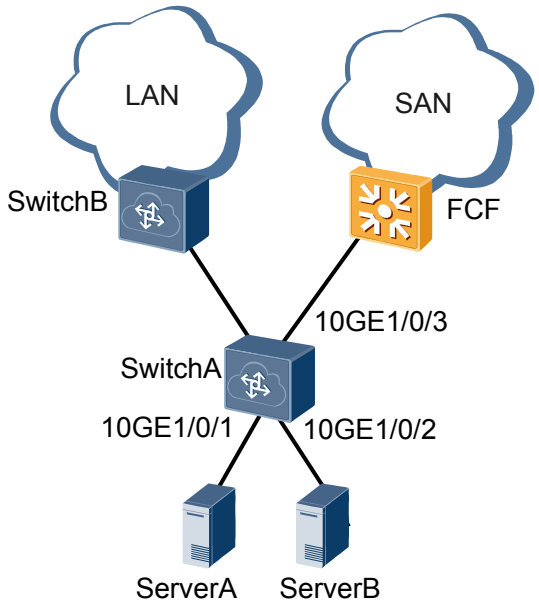

Configuration Roadmap

To meet QoS requirements in the preceding scenario, configure Data Center Bridging (DCB) on SwitchA.

The configuration roadmap is as follows:

1. If the Data Center Bridging eXchange protocol (DCBX) version is **intel-oui**, switch the DCBX version on SwitchA.

2. Configure DCBX on SwitchA to implement DCB capability negotiation at both ends of a link.
3. Configure Priority-based Flow Control (PFC) on SwitchA to ensure lossless transmission of SAN traffic.
4. Configure enhanced transmission selection (ETS) on SwitchA to ensure low latency of inter-process communication (IPC) traffic and bandwidth of SAN traffic.

NOTE

FCoE must be configured on SwitchA. For details on how to configure FCoE, see FCoE Configuration.

Procedure

Step 1 Switch the DCBX version.

The default DCBX version is **IEEE DCBX**. If the DCBX version of the remote device is **intel-oui**, switch the DCBX version on 10GE1/0/1, 10GE1/0/2, and 10GE1/0/3.

```
[~SwitchA] interface 10ge 1/0/1
[~SwitchA-10GE1/0/1] dcb compliance intel-oui
[*SwitchA-10GE1/0/1] quit
[*SwitchA] interface 10ge 1/0/2
[*SwitchA-10GE1/0/2] dcb compliance intel-oui
[*SwitchA-10GE1/0/2] quit
[*SwitchA] interface 10ge 1/0/3
[*SwitchA-10GE1/0/3] dcb compliance intel-oui
[*SwitchA-10GE1/0/3] quit
[*SwitchA] commit
```

Step 2 Configure DCBX.

```
[~SwitchA] lldp enable
[*SwitchA] interface 10ge 1/0/1
[*SwitchA-10GE1/0/1] lldp tlv-enable dcbx
[*SwitchA-10GE1/0/1] quit
[*SwitchA] interface 10ge 1/0/2
[*SwitchA-10GE1/0/2] lldp tlv-enable dcbx
[*SwitchA-10GE1/0/2] quit
[*SwitchA] interface 10ge 1/0/3
[*SwitchA-10GE1/0/3] lldp tlv-enable dcbx
[*SwitchA-10GE1/0/3] quit
[*SwitchA] commit
```

Step 3 Configure PFC.

Enable PFC for queues with priority 3. By default, PFC has been enabled for queues with priority 3. The procedure is not mentioned here.

Enable PFC on an interface.

```
[~SwitchA] interface 10ge 1/0/1
[~SwitchA-10GE1/0/1] dcb pfc enable mode auto
[*SwitchA-10GE1/0/1] quit
[*SwitchA] interface 10ge 1/0/2
[*SwitchA-10GE1/0/2] dcb pfc enable mode auto
[*SwitchA-10GE1/0/2] quit
[*SwitchA] interface 10ge 1/0/3
[*SwitchA-10GE1/0/3] dcb pfc enable mode auto
[*SwitchA-10GE1/0/3] quit
[*SwitchA] commit
```

Step 4 Configure ETS.

Create an ETS profile.

```
[~SwitchA] dcb ets-profile ets1
```

Map queue 3 to PG1, queue 7 to PG15, and other queues to PG0. Queue 3 maps PG1, and queues 6 and 7 map PG15 by default, so you only need to add queue 6 to PG0.

```
[*SwitchA-ets-ets1] priority-group 0 queue 6
```

Configure flow control based on the priority group and set DRR weights of PG0 and PG1 to 60% and 40% respectively.

```
[*SwitchA-ets-ets1] priority-group 0 wfq weight 60
[*SwitchA-ets-ets1] priority-group 1 wfq weight 40
[*SwitchA-ets-ets1] quit
```

Apply the ETS profile to 10GE1/0/1 and 10GE1/0/2.

```
[*SwitchA] interface 10ge 1/0/1
[*SwitchA-10GE1/0/1] dcb ets enable ets1
[*SwitchA-10GE1/0/1] quit
[*SwitchA] interface 10ge 1/0/2
[*SwitchA-10GE1/0/2] dcb ets enable ets1
[*SwitchA-10GE1/0/2] quit
[*SwitchA] commit
```

Step 5 Verify the configuration.

Check the ETS configuration.

```
[~SwitchA] display dcb ets-profile ets1
ETS Profile: ets1
-------------------------------------------------------------------------------
 PG(renum)    Queue   Schedule   Weight      CIR(kbps)     CBS(Bytes)    WRED
-------------------------------------------------------------------------------
  0(0)            -   WFQ        60(60)              -              -    --
                  0   PQ              -              -              -    --
                  1   PQ              -              -              -    --
                  2   PQ              -              -              -    --
                  4   PQ              -              -              -    --
                  5   PQ              -              -              -    --
                  6   PQ              -              -              -    --
  1(1)            -   WFQ        40(40)              -              -    --
                  3   PQ              -              -              -    --
 15(15)           -   PQ              -              -              -    --
                  7   PQ              -              -              -    --
 -------------------------------------------------------------------------------
```

Check the DCB configuration and negotiation status.

Run the **display dcb** command to check the DCB configuration and negotiation status. If the DCB configuration is correct, the following negotiation result is displayed:

```
[~SwitchA] display dcb
M:Manual;    A:Auto
-------------------------------------------------------------------------------
Interface         PFC Name     PFC Status  ETS Name    ETS Status App-Profile
-------------------------------------------------------------------------------
10GE1/0/1         default      ENABLE(A)   ets1        SUCCEED     -
10GE1/0/2         default      ENABLE(A)   ets1        SUCCEED     -
10GE1/0/3         default      ENABLE(A)   -           -           -
-------------------------------------------------------------------------------
```

----End

Configuration Files

```
#
sysname SwitchA
#
dcb pfc default
#
dcb ets-profile ets1
 priority-group 0 queue 0 to 2 4 to 6
 priority-group 15 queue 7
 priority-group 0 wfq weight 60
 priority-group 1 wfq weight 40
#
lldp enable
#
interface 10GE1/0/1
```

```
 port link-type trunk
 port trunk allow-pass vlan 2094
 lldp tlv-enable dcbx
 dcb pfc enable mode auto
 dcb ets enable ets1
#
interface 10GE1/0/2
 port link-type trunk
 port trunk allow-pass vlan 2094
 lldp tlv-enable dcbx
 dcb pfc enable mode auto
 dcb ets enable ets1
#
interface 10GE1/0/3
 port link-type trunk
 port trunk allow-pass vlan 2094
 lldp tlv-enable dcbx
 fcoe role vnp
 dcb pfc enable mode auto
#
return
```