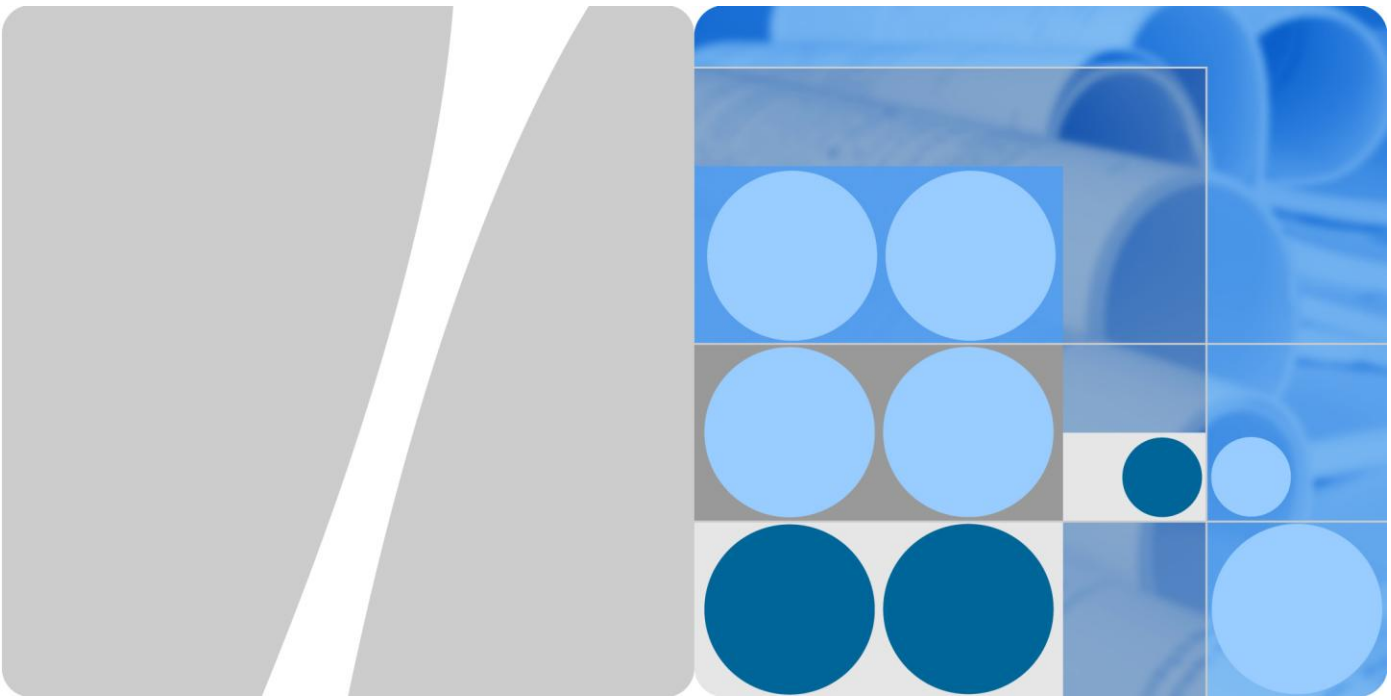

CloudEngine 12800 Series Switches Switches

SVF Technology White Paper

Issue 06

Date 2016-07-20

HUAWEI TECHNOLOGIES CO., LTD.

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 SVF Overview

This section describes the purpose and functions of SVF.

In the traditional access layer networking used in data centers, servers connect to the network through a large number of fixed switches, as shown in Figure 1-1. As the network size expands, increasing network devices are deployed. Network management becomes an important issue in data center infrastructure management.

Figure 1-1 Traditional data center access layer networking

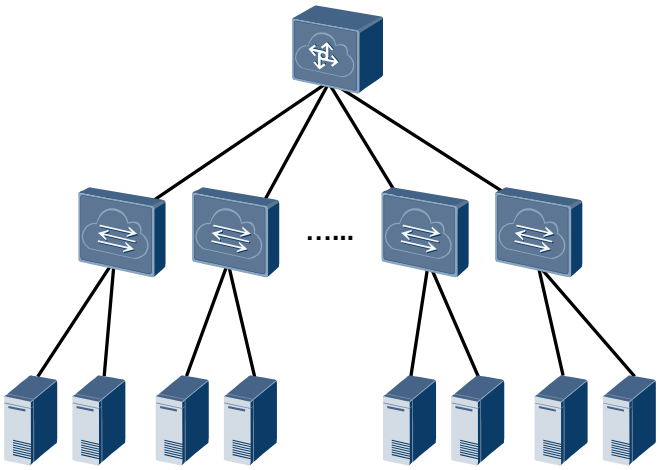

Super virtual fabric (SVF) is a vertical virtualization technology that virtualizes access and aggregation switches into one device to simplify network configuration and management, as shown in Figure 1-2.

Figure 1-2 SVF networking at the access layer

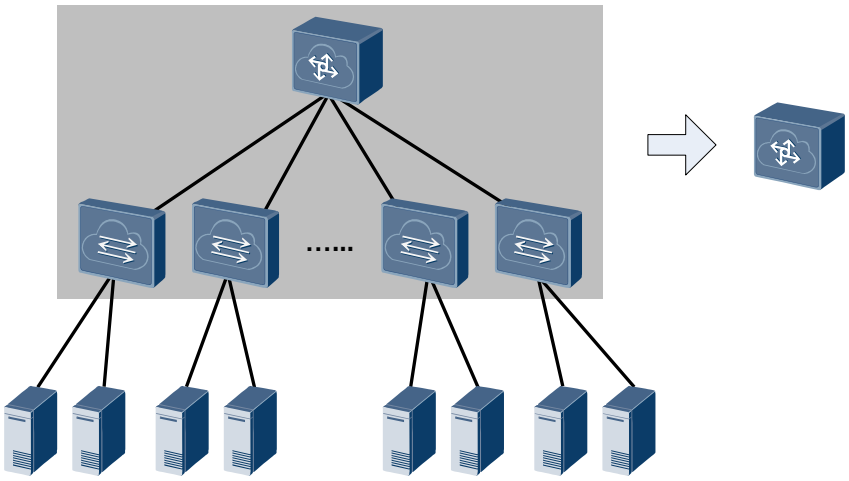

Compared with traditional access layer networking, SVF networking has the following advantages:

- Lower network construction costs: Low-cost switches are used as access switches, so network construction costs are reduced.
- Simpler configuration and management: SVF virtualizes multiple devices into one, reducing the number of nodes to manage. You do not need to deploy complex loop prevention protocols, simplifying network configuration and management.
- Higher scalability and more flexible deployment: When more access ports are required on the network, you only need to add low-cost fixed switches to the network. Moreover, these low-cost switches are deployed near servers, making network deployment more flexible.

2 Principles

2.1 Concepts

Figure 2-1 SVF-related concepts

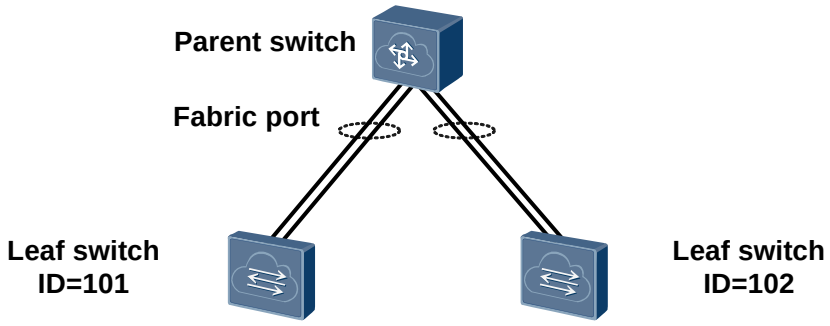

- **Parent/Leaf switch**

  Devices in an SVF system play either of the following roles:

  - Parent switch: controls and manages the entire SVF system.
  - Leaf switch: an extended device that acts as a remote device to connect to the parent switch and servers.

- **Leaf ID**

  A leaf ID is the identifier of a leaf switch. Leaf IDs are used to identify and manage leaf switches in an SVF system. Each leaf switch has a unique leaf ID.

- **Fabric port**

  A fabric port is a logical port that connects a parent switch and a leaf switch. One or more member ports can be added to a fabric port.

2.2 SVF Topology

Figure 2-2 shows the SVF topology where a leaf switch connects to the parent switch through member ports of a fabric port. One leaf switch can connect to only one parent switch.

Figure 2-2 SVF topology

The parent switch must be a standalone switch but not a stack of two devices.

To ensure uplink traffic reliability, you can enable a server to connect to leaf switches of two SVF systems through a multichassis link aggregation group (M-LAG) and each leaf switch to connect to a different parent switch, as shown in Figure 2-3. For details about M-LAG, see M-LAG Configuration in the *CloudEngine 12800 Series Switches Configuration Guide*.

Figure 2-3 M-LAG application in SVF

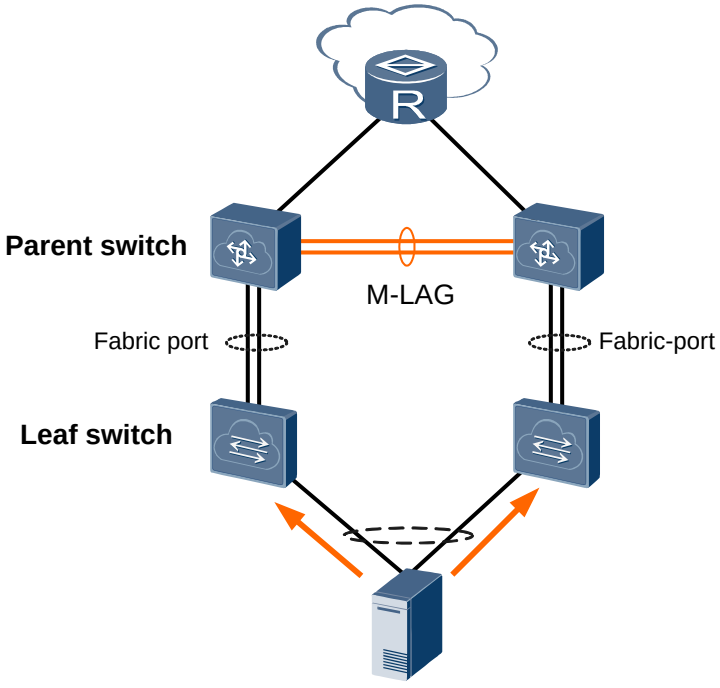

M-LAG ensures high reliability through link backup in addition to loading balancing server traffic. As shown in Figure 2-4, when a link or switch fails, server traffic is switched to another link.

Figure 2-4 Traffic switchover in the case of an M-LAG fault

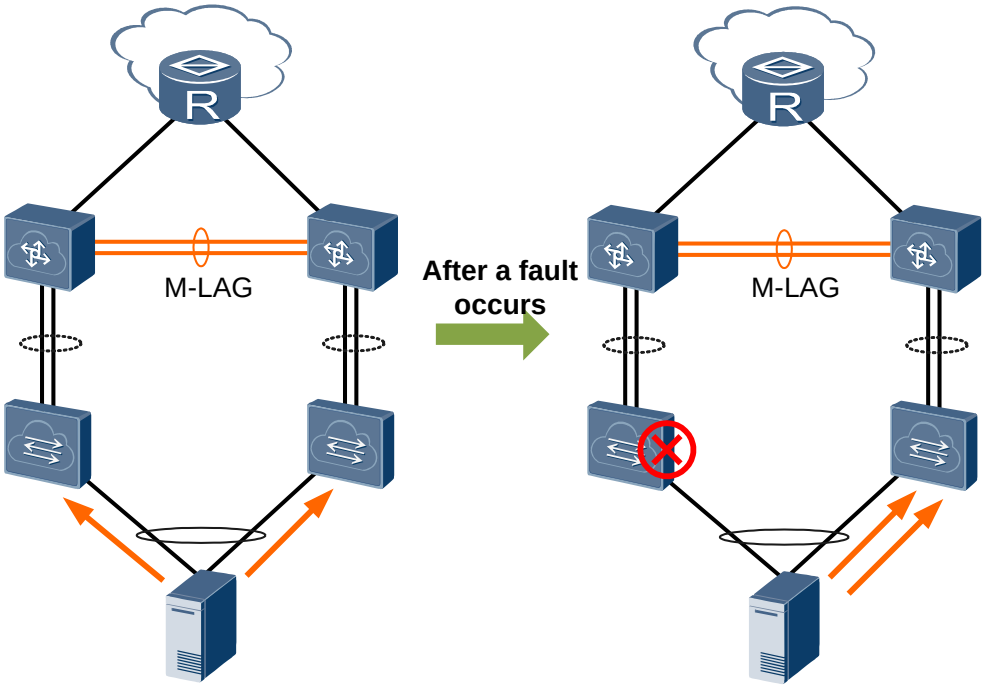

Figure 2-5 shows the network topologies that are not supported in SVF.

Figure 2-5 Unsupported network topologies

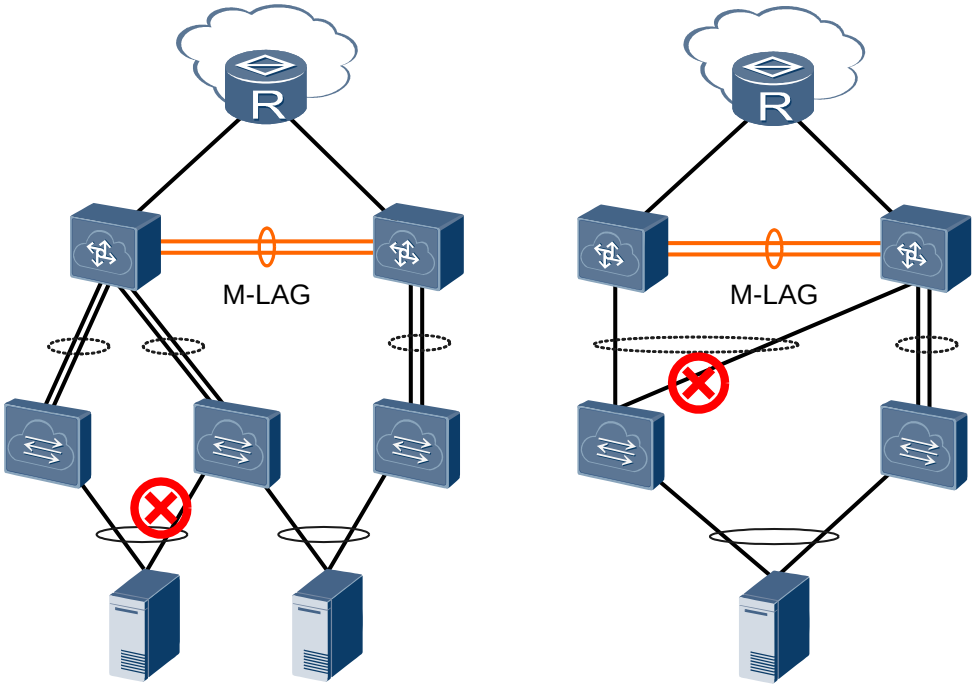

Ports of different leaf switches in the same SVF system cannot be added to the same Eth-Trunk

A leaf switch cannot be dual-homed to two parent switches

2.3 SVF Setup

Switch Working Modes

Switches that can be used as leaf switches in an SVF system support the following working modes:

- Stack mode: runs as an independent switch. A switch starting in this mode has complete system software and saves a configuration file.
- Leaf mode: runs as a leaf switch in an SVF system. A switch starting in this mode can only connect to a parent switch as an extended interface card. It has no configuration or management functions, and does not save any configuration.
- Auto-negotiation: determines whether to work in stack or leaf mode through auto-negotiation.

By default, a switch negotiates with the parent switch to determine its working mode, as shown in Figure 2-6. If a switch is powered on without any configuration (no startup configuration file), it initiates an SVF negotiation initiates an SVF negotiation by sending link probe packets and attempts to join an SVF system. If the switch has a configuration file, it starts in stack mode.

You can set the working mode of a switch. After you set the working mode of a switch to stack or leaf, the switch automatically starts in this mode and does not perform the negotiation.

Figure 2-6 Switch working mode determination

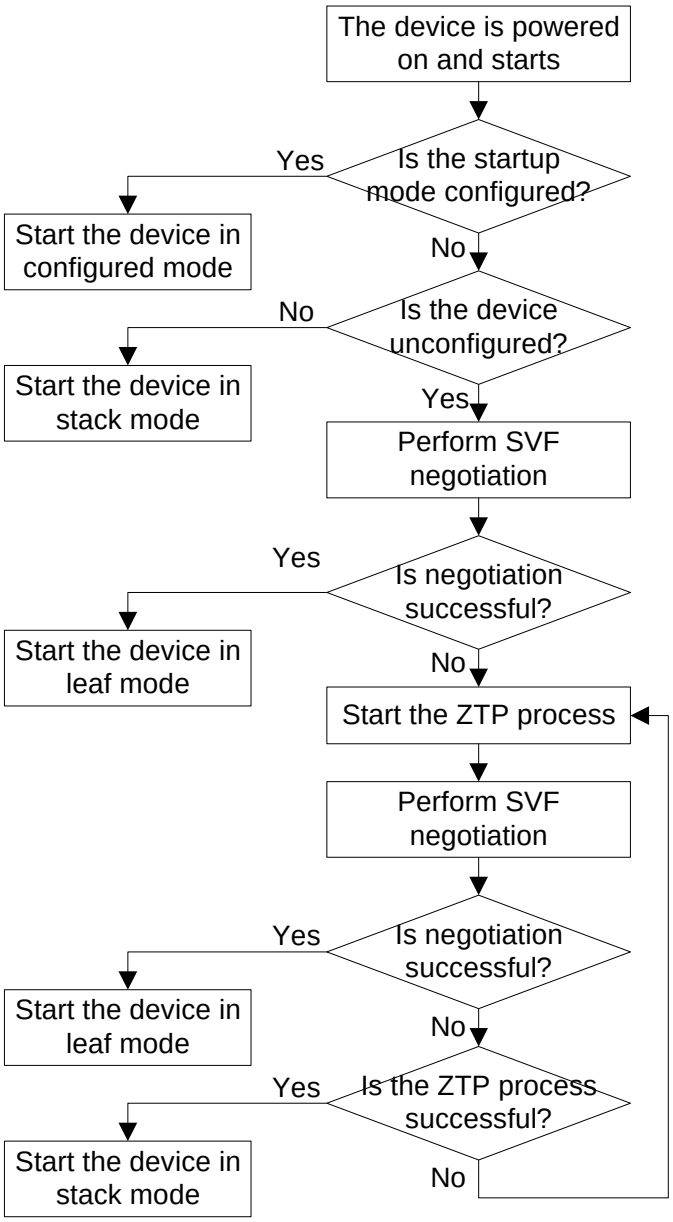

For details about ZTP, see ZTP Configuration in the *CloudEngine 8800&7800&6800&5800 Series Switches Configuration Guide - Basic Configuration*.

SVF Setup Process

After a switch starts in leaf mode, it obtains required information such as the leaf ID and software version from the parent switch. An SVF system is set up after the following steps:

1. Leaf ID allocation: After leaf switches are connected to the parent switch, they send protocol packets to the parent switch. The parent switch then allocates leaf IDs to the leaf switches.
2. Software version synchronization: The leaf switches obtain software version information from the parent switch. If the software version of a leaf switch is different from that of the parent switch, the leaf switch automatically downloads the system software from the parent switch. After the software download is complete, the leaf switch restarts with the new version and registers with the parent switch.

   Leaf switches only obtain a software package from the system software of the parent switch, not the complete system software. The system software of the parent switch contains two software packages with different functions. One is used for running the parent switch itself, and the other is synchronized to the leaf switches.
3. Configuration delivery: The parent switch delivers configuration information to the leaf switches.

2.4 SVF Management and Maintenance

Interface Numbering Rules

In an SVF system, interfaces on a parent switch are numbered in the format *1/slot ID/subcard ID/interface number*, and interfaces on a leaf switch are numbered in the format *leaf ID/1/subcard ID/interface number*, as shown in Table 2-1.

Table 2-1 Interface numbering rules

Role	Interface Numbering Rules	Remarks
Parent switch	1/slot ID/subcard ID/interface number **NOTE** 1 indicates the chassis ID.	After the SVF function is enabled on a parent switch, interfaces on the parent switch are numbered in the format *1/slot ID/subcard ID/interface number* but not *slot ID/subcard ID/interface number*.
Leaf switch	Leaf ID/1/subcard ID/interface number **NOTE** 1 indicates the slot ID.	The leaf ID ranges from 101 to 254.

Logging In to an SVF System

As a virtual device, an SVF system is configured and managed through the parent switch. You can use either of the following methods to log in to an SVF system:

- Through the console interface, management interface, or a Layer 3 interface of the parent switch

- Through the console interface or management interface of any leaf switch When users connect to the management interfaces of the parent and leaf switches of an SVF system, only the user connecting to the parent switch can log in to the SVF system.

You must specify the IP address of the parent switch when logging in to the SVF system remotely.

Configuration and Management

In an SVF system, the leaf switches do not have configuration and management functions, and they are configured and managed on the parent switch. After you configure the leaf switches on the parent switch, the configuration information is saved on the parent switch. The leaf switches do not save any configuration. If a leaf switch is restarted or replaced by a new leaf switch, the parent switch delivers configuration to the leaf switch through the SVF link. Similar to interface cards of chassis switches, leaf switches are plug-and-play.

2.5 SVF Upgrade

During the SVF system upgrade, the system software for next startup is specified on the parent switch. After the parent switch restarts, leaf switches automatically synchronize the system software version with the software version on the parent switch. To minimize service interruption time, leaf switches download the new system software from the parent switch at the same time the parent switch specifies the system software for next startup. After the new system software is downloaded to leaf switches, the leaf switches and parent switch restart simultaneously. Table 2-2 describes the specifications for the system restart time during the SVF system upgrade.

Table 2-2

Software Configuration	System Restart Time
A few configurations • Software configuration: There are a few VLAN configurations, and STP is disabled. • Hardware configuration: The parent switch has two MPUs and six 48*10GE LPUs installed and is fully configured with SFUs. Sixteen CE6810EI leaf switches exist. Only a few VLAN configurations exist, and STP is disabled.	1320 seconds
Typical configuration • Software configuration: A total of 2048 VLANs are configured; Half of ports are added to VLANs and are Up; STP is disabled. • Hardware configuration: The parent switch has two MPUs and six 48*10GE LPUs installed and is fully configured with SFUs. Sixteen CE6810EI leaf switches exist.	2400 seconds
Notes • The preceding system restart time is used only as reference and varies depending on the	

Software Configuration	System Restart Time
software and hardware configurations.	

If M-LAG has been configured between two SVF systems, you can upgrade the two SVF systems one by one to ensure uninterrupted service traffic forwarding. As shown in Figure 2-7, you can upgrade SVF 1 and then upgrade SVF 2 after SVF 1 is upgraded.

Figure 2-7 SVF system upgrade when M-LAG is configured

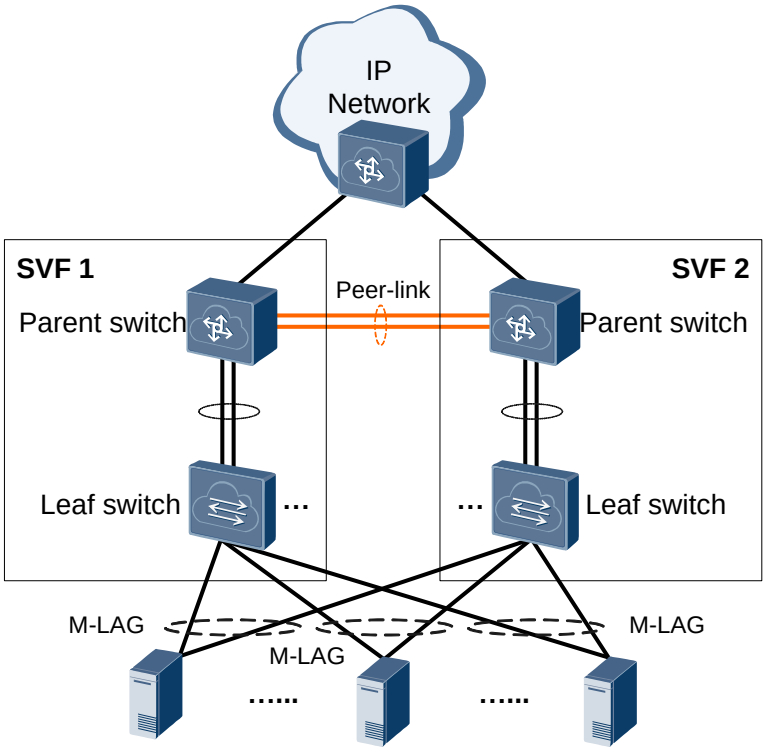

3 Application

This section describes a typical SVF application scenario.

As shown in Figure 3-1, low-cost leaf switches are deployed in server racks to provide servers with network access. This deployment reduces network construction costs. The parent switches are deployed in a centralized manner, which simplifies device management and allows network cables to be routed more neatly.

Figure 3-1 SVF application at the data center access layer

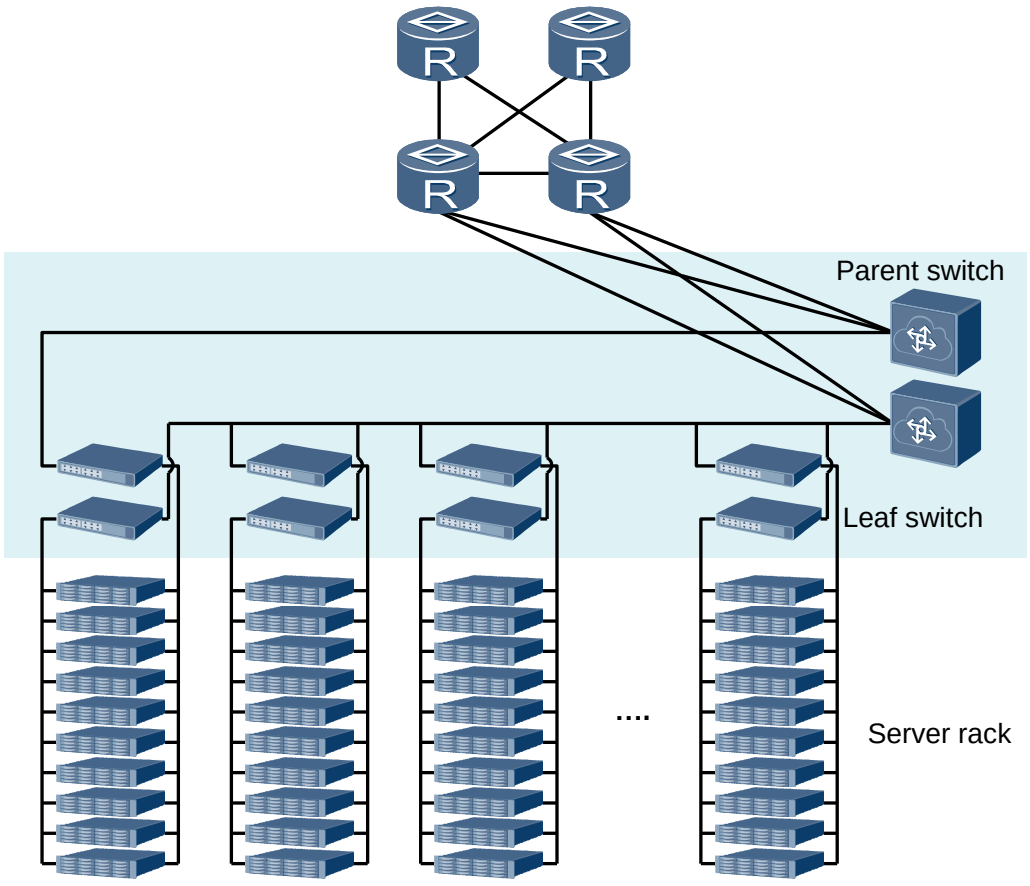

4 Configuration Notes

This section provides the points of attention when configuring SVF.

License Support

SVF is a basic feature of a switch and is not under license control.

Device and Version Requirements

Table 4-1 lists the device models and versions used to set up SVF systems.

Table 4-1 Device models and versions used to set up SVF systems

Version	Parent Switch Model	Leaf Switch Model	Remarks
V100R005C00	• CE12804 • CE12808 • CE12812 • CE12816	• CE6850-48T4Q-EI • CE6810-48S4Q-EI • CE5810-48T4S-EI • CE5810-24T4S-EI	• Leaf switches in an SVF system can be switches of different series and models. For example, CE6810-48S4Q-EI and CE5810-48T4S-EI can belong to the same SVF system. • Versions of leaf switches must be the initial version that supports this model of switches as leaf switches or higher versions.
V100R005C10	• CE12804 • CE12808 • CE12812 • CE12816 • CE12804S • CE12808S	• CE6850-48T4Q-EI • CE6810-48S4Q-EI • CE5810-48T4S-EI • CE5810-24T4S-EI • CE5850-48T4S2Q-EI • CE6810-48S-LI • CE6810-48S4Q	

Version	Parent Switch Model	Leaf Switch Model	Remarks
		-LI • CE6810-24S2Q-LI • CE6810-32T16S4Q-LI	
V100R006C00 V200R001C00	• CE12804 • CE12808 • CE12812 • CE12816 • CE12804S • CE12808S	• CE6850-48T4Q-EI • CE6810-48S4Q-EI • CE5810-48T4S-EI • CE5810-24T4S-EI • CE5850-48T4S2Q-EI • CE5855-48T4S2Q-EI • CE5855-24T4S2Q-EI • CE6810-48S-LI • CE6810-48S4Q-LI • CE6810-24S2Q-LI • CE6810-32T16S4Q-LI • CE6850-48S6Q-HI • CE6850-48T6Q-HI • CE6851-48S6Q-HI	

Device Quantity and Port Requirements

Table 4-2 lists the requirements for member switch quantity and ports during SVF setup.

Table 4-2 SVF member switch quantity and port requirements

Item	Description	Remarks
Number of parent switches	1	The parent switch must be a standalone switch but not a stack of two devices.

Item	Description	Remarks
Number of leaf switches	1-32	The leaf switch must be a standalone switch but not a stack of multiple devices.
Types of member ports of a fabric port	10GE electrical ports, 10GE optical ports, and 40GE optical ports	40GE ports and 10GE ports split from 100GE or 40GE ports of parent switches can be used to connect to leaf switches.
Number of physical member ports in a fabric port	1-8	It is recommended that you add at least two member ports in each fabric port to ensure high reliability.

To learn about the default and configured leaf switch ports to be connected to a parent switch, see Table 4-3.

Table 4-3 List of leaf switch ports to be connected to a parent switch

Version	Leaf Switch Model	Default Port	Configured Port
V100R005C00	• CE5810-48T4S-EI • CE5810-24T4S-EI	Last four 10GE ports	Configured ports are not supported.
	CE6810-48S4Q-EI	Last four 40GE ports or last eight 10GE ports perform auto-negotiation: • If only 10GE ports are connected to a parent switch, 10GE ports are preferred. • If only 40GE ports are connected to a parent switch, 40GE ports are preferred. • If both 10GE and 40GE ports are connected to a parent switch, 40GE ports are preferred.	• Auto-negotiation • 10GE: last eight 10GE ports • 40GE: last four 40GE ports
	CE6850-48T4Q-EI	Last four 40GE ports	• 10GE: last eight 10GE ports • 40GE: last four 40GE ports
V100R005C10	• CE5810-48T4S	4*10GE	Configured ports are not

Version	Leaf Switch Model	Default Port	Configured Port
V100R006C00 V200R001C00	-EI • CE5810-24T4S-EI		supported.
	• CE5850-48T4S2Q-EI • CE5855-48T4S2Q-EI • CE5855-24T4S2Q-EI	2*40GE	• 4*10GE • 2*40GE
	CE6810-48S4Q-EI	8*10GE and 4*40GE ports perform auto-negotiation: • If only 10GE ports are connected to a parent switch, 10GE ports are preferred. • If only 40GE ports are connected to a parent switch, 40GE ports are preferred. • If both 10GE and 40GE ports are connected to a parent switch, 40GE ports are preferred.	• auto-negotiation • 4*10GE • 8*10GE • 4*40GE
	• CE6810-48S4Q-LI • CE6810-32T16S4Q-LI	4*40GE	• 4*10GE • 8*10GE • 4*40GE
	CE6850-48T4Q-EI	4*40GE	• 4*10GE • 8*10GE • 4*40GE
	CE6810-24S2Q-LI	2*40GE	• 4*10GE • 8*10GE • 2*40GE
	CE6810-48S-LI	8*10GE	• 4*10GE • 8*10GE
	• CE6850-48S6Q-HI • CE6850-48T6Q-HI • CE6851-48S6	6*40GE	• 4*10GE • 8*10GE • 4*40GE • 6*40GE

Version	Leaf Switch Model	Default Port	Configured Port
	Q-HI		

Notes

- N*10GE/N*40GE indicates the last N 10GE/40GE interfaces on the device.
- When auto-negotiated or configured N*10GE or N*40GE ports are connected to a parent switch, these ports cannot be used as common service ports again. For example, when any one of 40GE1/0/1 to 40GE1/0/4 on a CE6810-48S4Q-EI is connected to a parent switch, the other ports of 40GE1/0/1 to 40GE1/0/4 cannot be used as common service ports.
- In V100R005C00 and later versions, the auto-negotiation function takes effect on a CE6810-48S4Q-EI switch only one time. Leaf switches can use only the ports negotiated for the first time to connect to a parent switch. To change the leaf switch ports, you must manually configure the type of leaf switch ports.

Cable Requirements

High-speed cables, AOC cables, optical modules and fibers, or network cables can be used to connect stack member switches. Table 4-4 describes the cables applicable to different ports.

Table 4-4 Requirements for stack cables

Port Type	High-Speed Cable	AOC Cable	Optical Module	Network Cable
10GE optical port	SFP+ high-speed cable	SFP+ AOC cable.	10GE SFP/SFP+ optical module The required optical fibers are determined by the optical modules you select.	-
10GE electrical port	-	-	-	10GE electrical ports use Category 6, Category 6A, or Category 7 cables that comply with IEEE 802.3an.
40GE optical port	QSFP+ high-speed cable	QSFP+ AOC cable.	40GE QSFP optical module The required optical fibers are determined by the optical	-

Port Type	High-Speed Cable	AOC Cable	Optical Module	Network Cable
			modules you select.	
Notes • 10GE optical ports can be used for connection only when they have 10GE optical modules installed. If the GE optical or copper modules are installed on 10GE optical ports, the ports cannot be used to set up an SVF system.				

Feature Dependencies and Limitations

Precautions

Pay attention to the following points when configuring SVF:

- Huawei-certified optical or copper modules must be used. If high-speed cables or AOCs are used, you must purchase cables from Huawei.

  Non-Huawei-certified optical or copper modules or cables that are not purchased from Huawei cannot ensure transmission reliability and may affect service stability. Huawei is not liable for any problem caused by the use of non-Huawei-certified optical or copper modules, or cables not purchased from Huawei, and will not fix such problems.
- One leaf switch can connect to only one parent switch.
- You can configure fabric ports in VSs and set the same ID for fabric ports in different VSs. Ports of leaf switches that a fabric port in a VS connects to belong to that VS.
- If the negotiated working of a switch is leaf, the switch starts in leaf mode. To start the switch in other modes, manually configure the working mode of the switch.
- To reduce CPU loads on LPUs and ensure reliability, you are advised to connect a leaf switch to multiple LPUs of the parent switch.

  One LPU can connect to a maximum of 16 leaf switches. If there are more than 16 leaf switches, an alarm is generated. Determine the number of leaf switches connected to each LPU based on the total number of leaf switches in an SVF system. For example, if an SVF system has a total of 10 leaf switches and each leaf switch is dual-homed to two LPUs of the parent switch, connect 5 leaf switches to each LPU.

  If an SVF system has more than 16 leaf switches, you are advised to connect the leaf switches to more than 2 LPUs to ensure redundancy and reliability.
- A function configured on a leaf switch must also be supported by the parent switch's LPU connected to the leaf switch. If the LPU does not support the function, the function does not work on the leaf switch.

  For example, if a leaf switch is connected to a BA series Layer 2 LPU on the parent switch, Layer 3 functions configured on the leaf switch do not work.
- If the LPUs of the parent switch connected to the same leaf switch provide different service performance or entry capacities, packets sent from the leaf switch may fail to be forwarded or processed on the LPU with lower performance or capacity.

  It is recommended that a leaf switch be connected to LPUs of the same specifications on the parent switch.
- If an existing leaf switch is replaced by a new switch of a different model, the new switch's ports with the same port types and port numbers as the existing switch's ports will inherit the configurations of the existing switch's ports. If the inherited

configurations contain the services that are not supported by the new switch, the configurations of these services cannot be deleted from the new switch. To delete these unsupported configurations, enable the new switch to go offline and then delete these configurations in offline mode.

- Multicast resources are shared by multiple services including VLAN, MAC, Eth-Trunk, M-LAG, Layer 2 protocol transparent transmission, SVF, Layer 3 physical interface, and multicast. If multicast resources in the system are insufficient for any of these services you are configuring, the system will display a configuration failure message. To solve this problem, you can delete some unnecessary service configuration, for example, delete unused VLANs.
- If the total traffic volume of high-priority (priority 7) user unknown unicast, multicast, and broadcast packets to be forwarded by a fabric port exceeds 90% of the fabric port bandwidth, these packets will preempt bandwidth of SVF management packets. This situation will result in the loss of management packets, or even worse, the restart of the leaf switch.

  When a CE5810EI functions as a leaf switch, it will also be affected by high-priority unicast packets.
- In V100R005C00, if a fabric port that is incorrectly connected to an external network receives Ethernet packets of the type 0X5100 or 0X88CA, the SVF system may become faulty. If a common service port receives the two types of packets, the SVF system may also become faulty.

  In V100R005C10 and later versions, if a fabric port that is incorrectly connected to an external network receives Ethernet packets of the type 0X5100 or 0X88CA, the SVF system may become faulty. When a common service port receives the two types of packets, it will discard the packets.
- Ports on the CE-FWA/CE-IPSA/CE-L48XS-FDA/CE-L36LQ-FD/CE-L36CQ-FD card cannot be used to establish an SVF system.

Feature Support in SVF

CE12800 series switches support different features than CE6800&5800 series switches. The following table shows the features supported by an SVF system consisting of CE12800 and CE6800&5800. The configuration of features that are not listed in the table is the same as that on CE12800.

Table 4-5 Feature support in SVF

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
CSS	Supported	Not supported	SVF and CSS conflict and cannot be configured together. If a switch does not have CSS enabled but already has CSS configuration, the switch cannot have SVF enabled. To check whether a switch has CSS configuration, run the **display stack configuration** command.
ISSU	Supported	Supported	An SVF system cannot be upgraded using ISSU.
Aging time of dynamic MAC entries	Supported	Supported	Before enabling SVF, run the **mac-address aging-time** *aging-time* command in the system view to set the aging time of MAC entries to

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
			1800s or longer. By default, the aging time of MAC entries is 300s. If VSs in group mode have been configured, you also need to set the aging time of MAC entries for these VSs to 1800s or longer.
Eth-Trunk	Supported	Supported	Eth-Trunk member interfaces must be located on the same leaf switch. Before adding an interface on a leaf switch to an Eth-Trunk, run the **extend enable** command in the Eth-Trunk interface view.
Traffic statistics on the VLANIF interface	Supported	Supported	SVF does not support the statistics on outgoing traffic from the leaf switch.
QinQ	Supported	Supported	SVF and QinQ conflict and cannot be configured together.
VLAN Mapping	Supported	Supported	SVF and VLAN mapping conflict and cannot be configured together.
VLAN	Supported	Supported	• SVF cannot be used with MAC address-based, subnet-based, or protocol-based VLAN assignment simultaneously. • SVF and MUX VLAN conflict and cannot be configured together.
MSTP	Supported	Supported	In an SVF composed of fixed and modular switches, the number of VLANs multiplied by the number of MSTIs cannot exceed 8000.
Sub-interface	Supported	Supported	Leaf switches do not support sub-interfaces, but parent switches support sub-interfaces.
Port isolation	Supported	Supported	In an SVF system consisting of modular and fixed switches, port isolation can be enabled only on interfaces of the same leaf switch.
Port split	Supported	Supported	Ports on leaf switches cannot be split after being assigned to VSs.
Port aggregate	Not supported	Supported by only the CE6810EI, CE6810LI, CE6850-48S6Q-HI, CE6851HI	If leaf switches in an SVF system support interface aggregation, you can configure interface aggregation for the leaf switches on the parent switch.

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
		, and CE5855-24T4S2Q-EI	
Hardware-based forwarding of BPDUs	Supported	Supported	The interfaces that connect the parent switch and leaf switches do not support hardware-based forwarding of BPDUs, while other interfaces support hardware-based forwarding of BPDUs.
IP multicast	Supported	Layer 3 multicast not supported on the CE6810LI	• In an SVF system, the physical interfaces that have been switched to Layer 3 mode using the **undo portswitch** command do not support Layer 3 multicast. • The SVF function cannot be configured together with Bidir-PIM, PIM-DM, or multicast VPN. • When a CE5855EI switch acts as a leaf switch in an SVF system, its interfaces have IPv6 capability.
ARP	Supported	Supported	• If an ARP entry on a Layer 3 interface of a leaf switch does not have a corresponding MAC address entry, the leaf switch may broadcast traffic from the Layer 3 interface in a VLAN. • If the number of ARP entries on a leaf switch exceeds the maximum number supported by a standalone leaf switch, the leaf switch may fail to forward Layer 3 protocol packets or cannot be pinged. It is recommended that the number of ARP entries configured on a leaf switch be no more than the maximum number supported by it. • In V100R005C00 and V100R005C10, the extend mode of ARP resource allocation is not supported by an SVF system.
IP tunnel	A 6to4 auto tunnel is not supported.	The CE6810LI does not support IP tunnels.	• Layer 3 traffic of leaf switches can be transmitted in an IP tunnel, but leaf switches cannot establish IP tunnels with other Layer 3 devices. • Parent switches support IP tunnels.
GRE tunnel	Supported	The CE6810LI does not support GRE	• Layer 3 traffic of leaf switches can be transmitted in a GRE tunnel, but leaf switches cannot establish GRE tunnels with other Layer 3 devices. • Parent switches support GRE tunnels.

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
		tunnels.	
MPLS	Supported	Only the CE6850HI supports MPLS.	• Layer 3 traffic of leaf switches can be transmitted in MPLS LDP tunnels, but leaf switches cannot establish MPLS LDP tunnels with other Layer 3 devices. • Parent switches support MPLS.
MPLS TE	Supported	Not supported	• Layer 3 traffic of leaf switches can be transmitted in MPLS TE tunnels, but leaf switches cannot establish MPLS TE tunnels with other Layer 3 devices. • Parent switches support MPLS TE.
VPLS	Supported	Not supported	Parent switches support VPLS, but leaf switches do not.
VLL	Supported	Not supported	Parent switches support VLL, but leaf switches do not.
BGP VPN	Supported	Not supported	• Layer 3 traffic of leaf switches can be transmitted in BGP VPN tunnels, but leaf switches cannot establish BGP VPN tunnels with other Layer 3 devices. • Parent switches support BGP VPN. • After you run the **traffic-statistics enable** command in the VPN instance view to enable traffic statistics collection, the SVF cannot collect incoming traffic that is directly forwarded to leaf switches and outgoing traffic from leaf switches.
BFD	Supported	Supported	Parent switches support BFD, but leaf switches do not.
URPF	Supported	CE6810LI does not support URPF.	Parent switches support URPF, but leaf switches do not.
EVN	Supported	Not supported	SVF and EVN conflict and cannot be configured together.
VXLAN	supported	Not supported	SVF and VXLAN conflict and cannot be configured together.
TRILL	Supported	CE5810EI and CE6810LI do not support TRILL.	In V100R005C00, Leaf switches do not support the TRILL function, that is, they cannot function as access or network devices. Parent switches support TRILL. In V100R005C10, Leaf switches can only function as access devices. Parent switches support TRILL.

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
NLB	Supported	CE6810LI does not support NLB.	Layer 3 traffic in the leaf switch side cannot access the NLB servers connected to the parent switch.
DCB	Supported	Supported CE5800 does not support DCB.	In V100R005C00 version, an SVF system does not support DCB. In V100R006C00 and later version, an SVF system supports DCB. When a CE6810LI or a CE5800 functions as a leaf switch, an SVF system does not support DCB.
FCoE	Supported	Supported CE5800 does not support FCoE.	In V100R005C00 version, an SVF system does not support FCoE. In V100R006C00 and later version, an SVF system supports FCoE. When a CE5800 functions as a leaf switch, an SVF system does not support FSB. When a CE6810LI or a CE5800 functions as a leaf switch, an SVF system does not support FCF.
Redirection	Supported	Supported	In an SVF system composed of fixed and modular switches, when a traffic policy containing redirection is applied to a VLANIF interface of a leaf switch, local Layer 3 traffic (inbound and outbound interfaces of the traffic are located on the same leaf switch) cannot be redirected to the interface.
NetStream	Supported	CE6810LI only supports exported packets carrying Layer 2 flow statistics. CE6800&5800 series switches do not support inbound NetStream sampling using snoop resources.	• After the leaf switch collects statistics on the traffic to be sent to the parent switch for Layer 3 forwarding, the flow table generated for the traffic does not contain routing information. To collect such routing information, you are advised to configure NetStream on the Fabric-port of the parent switch. However, the routing information on the traffic forwarded by Layer 3 main interface cannot be collected. • If routing information (next-hop IP address and AS) is specified in a NetStream Top Talkers template applied to an interface of a leaf switch, the NetSteram Top Talkers template cannot collect statistics on or sort the traffic to be sent to the parent switch for routing. To collect statistics on and sort such traffic, you are advised to apply the NetStream Top Talkers template to the FNI interface of the parent switch. • When inbound NetStream sampling uses snoop resources, the configuration takes effect only on the parent switch, but cannot

Feature	CE12800	CE6800&5800	Feature Support After SVF Is Configured
			take effect on leaf switches. • A CE5855EI functioning as a leaf switch supports NetStream IPv6 function.
sFlow	Supported	Supported	• The IP address of the sFlow agent must be manually set; otherwise, the **Agent IP** field may be incorrect. • A leaf switch does not support sFlow sampling for routing information of the traffic to be sent to the parent switch for Layer 3 forwarding. To sample routing information for such traffic, you are advised to configure sFlow on the Fabric-port of the parent switch. However, the routing information on the traffic passing Layer 3 main interface cannot be collected. • A CE5855EI functioning as a leaf switch supports sFlow IPv6 function.
Mirroring	Supported	Supported	• An SVF system does not support 1:N mirroring. • Layer 3 remote mirroring between different leaf or parent switches is not supported. • Eth-Trunk does not support mirroring between different leaf or parent switches. • Traffic on interfaces of leaf switches cannot be mirrored to an observing interface group and do not support Layer 3 remote mirroring. • Port mirroring: If the mirrored port is located on a leaf switch, the leaf switch can be configured with only one observing port. If the mirrored port is located on a parent switch, the parent switch can be configured with a maximum of eight observing ports. • VLAN mirroring and flow mirroring: Only one observing port can be configured in the system. • In port mirroring, if the mirrored port is located on a parent switch, it is recommended to configure the observing port on the parent switch rather than a leaf switch.
Open system	Supported	Supported	The SVF function and open system cannot be configured simultaneously.

5 Default Configuration

This section describes the default configuration of SVF.

Table 5-1 Default SVF configuration

Parameter	Default Setting
Whether SVF is enabled	No

6 Configuring a Parent Switch

6.1 Enabling the SVF Function on a Parent Switch

Context

The SVF function can be configured only after the SVF function is enabled on the parent switch.

Pre-configuration Tasks

Before enabling SVF, complete the following tasks:

- Delete the configuration of features that conflict with SVF. The SVF function conflicts with some other functions. If these functions are configured, the SVF function cannot be enabled.

  For the conflicting functions, see "Feature Support in SVF" in 4 Configuration Notes.
- Run the **mac-address aging-time** *aging-time* command in the system view to set the aging time of MAC entries to 1800s or longer. By default, the aging time of MAC entries is 300s.

  If VSs in group mode have been configured, you need to set the aging time of MAC entries for these VSs to 1800s or longer.

Procedure

- After the parent switch starts, run commands to enable the SVF function:
  a. Run the **system-view** command to enter the system view.
  b. Run the **stack** command to enter the stack management view.
  c. Run the **switch svf enable** command to enable the SVF function.

     By default, the SVF function is disabled.

     After the SVF function is enabled, the parent switch automatically restarts. When the parent switch completes startup, interface number format is changed from slot ID/subcard ID/port number to 1/slot ID/subcard ID/port number.
- During the startup of the parent switch, set the working mode menu to enable the SVF function:
  a. Log in to the parent switch through the console port.

b. During the startup process, press **Ctrl+Y** to enter the working mode menu. The message "Press CTRL+Y to modify the switch mode" is displayed.

c. Select "1. Set SVF enable" in the working mode menu to enable the SVF function.

The working mode menu is as follows:

```
Press CTRL+Y to modify the switch mode ..................... 2

       Switch SVF Enable Menu

       1. Set SVF enable    //Enable the SVF function.
       2. Continue to boot   //Continue to start the switch.

You have 30 seconds to select the menu, the system will continue to boot when
time is up.
Enter your choice(1-2):1
```

----End

6.2 Configuring a Fabric Port

Context

A parent switch connects to leaf switches using fabric ports. One or more member ports can be added to a fabric port to improve SVF link bandwidth and reliability.

NOTE

- A fabric port can have a maximum of eight member ports.
- You can configure fabric ports in VSs and set the same ID for fabric ports in different VSs. Ports of leaf switches that a fabric port in a VS connects to belong to that VS.

Procedure

Step 1 Create a fabric port and bind a leaf ID to it.

1. Run the **system-view** command to enter the system view.
2. Run the **interface fabric-port** *port-id* command to create a fabric port.

   By default, no fabric port exists in the system. The value of *port-id* ranges from 1 to 32.

NOTE

A total of 32 fabric ports can be created for VSs.

3. Run the **port bind member** *member-id* command to bind a leaf ID to the fabric port.

   By default, no leaf ID is bound to a fabric port. The value of *member-id* ranges from 101 to 254.

   Leaf IDs are used to identify and manage leaf switches in an SVF system. Each leaf switch has a unique leaf ID. Leaf IDs are bound to fabric ports. When a leaf switch connects to the parent switch, it obtains its leaf ID from the parent switch.

NOTE

- Each fabric port can have only one leaf ID bound to it, and each leaf ID can be bound to only one fabric port.
- If you change or delete the leaf ID bound to a fabric port, the leaf switch connected to the fabric port will restart.

4. (Optional) Run the **description** *description* command to configure the description of the fabric port.

   By default, a fabric port does not have a description.

   To facilitate fabric port management and identification, you can configure descriptions for fabric ports. For example, you can describe the ID of a leaf switch that connects to a fabric port.

5. Run the **commit** command to commit the configuration.

Step 2 Add service ports to the fabric port.

NOTE

The configuration can be performed in the fabric port view or interface view. The two configuration methods obtain the same result.

Service ports are automatically configured as physical member ports after being added to a fabric port. Alternatively, run the **port mode stack interface** *interface-type* { *interface-number1* [**to** *interface-number2*] } &<1-32> command in the stack management view or run the **port mode stack** command in the interface view to configure service ports as physical member ports and then add the physical member ports to a fabric port.

- Configuration in the fabric port view:
  a. Run the **system-view** command to enter the system view.
  b. Run the **interface fabric-port** *port-id* command to enter the fabric port view.
  c. Run the **port member-group interface** *interface-type* { *interface-number1* [**to** *interface-number2*] } &<1-8> command to add member ports to the fabric port.
  d. Run the **commit** command to commit the configuration.
- Configuration in the interface view:
  a. Run the **system-view** command to enter the system view.
  b. Run the **interface** *interface-type interface-number* command to enter the interface view.
  c. Run the **fabric-port** *port-id* command to add the port to the fabric port.

     To add multiple ports to the fabric port, repeat steps b and c.
  d. Run the **commit** command to commit the configuration.

Step 3 (Optional) Disable auto-negotiation on the fabric port.

1. Run the **system-view** command to enter the system view.
2. Run the **interface fabric-port** *port-id* command to enter the fabric port view.
3. Run the **leaf negotiation disable** command to disable auto-negotiation on the fabric port.

   By default, auto-negotiation is enabled on a fabric port.

   To prevent the switch connected to a fabric port from connecting to the SVF system mistakenly, you can disable auto-negotiation on the fabric port. After auto-negotiation is disabled, the unconfigured switch connected to the fabric port cannot join the SVF system through auto-negotiation. However, the switch configured to work in **Leaf** mode can still join the SVF system.
4. Run the **commit** command to commit the configuration.

Step 4 (Optional) Configure a load balancing mode for uplink of leaf switches connected to a fabric port.

1. Run the **system-view** command to enter the system view.
2. Run the **interface fabric-port** *port-id* command to enter the fabric port view.

3. Run the **leaf uplink-port load-balance** { **dst-ip** | **dst-mac** | **src-dst-ip** | **src-dst-mac** | **src-ip** | **src-mac** } command to configure a load balancing mode for uplink of leaf switches connected to the fabric port.

   By default, the load balancing mode for uplink of leaf switches is **src-dst-ip**.

   NOTE

   The configured load balancing mode takes effect only on uplink of leaf switches but not on fabric ports. You can run the **load-balance profile** *profile-name* command to configure the load balancing mode for fabric ports on parent switches.

4. Run the **commit** command to commit the configuration.

----End

6.3 (Optional) Configuring the Device Type or Serial Number of a Leaf Switch Allowed to Join an SVF System

Context

To prevent devices from incorrectly joining an SVF system and ensure SVF security, you can configure the device type or serial number of a leaf switch allowed to join an SVF system. When the device type or serial number of a leaf switch is consistent with the configured one, the leaf switch can join the SVF system. Otherwise, the leaf switch cannot join the SVF system.

By default, all the devices that can function as leaf switches can join an SVF system.

NOTE

When a leaf switch has joined an SVF system, it leaves the SVF system if the changed device type or serial number is inconsistent with the configured one.

Procedure

Step 1 Run the **system-view** command to enter the system view.

Step 2 Run the **stack** command to enter the stack management view.

Step 3 (Optional) Configure the device type or serial number of a leaf switch allowed to Join an SVF system.

- Run the **leaf member** *member-id* **type** *type* command to configure the device type of a leaf switch allowed to join an SVF system.

  By default, the device type of a leaf switch allowed to join an SVF system is not configured.

- Run the **leaf member** *member-id* **serial-number** *serial-number* command to configure the serial number of a leaf switch allowed to join an SVF system.

  By default, the serial number of a leaf switch allowed to join an SVF system is not configured.

Step 4 Run the **commit** command to commit the configuration.

----End

7 (Optional) Configuring a Leaf Switch

7.1 Configuring the Working Mode of Leaf Switches

Context

By default, if a leaf switch starts with no configuration, it performs a negotiation with the parent switch and attempts to join the SVF system in leaf mode, and you do not need to perform any configuration on the leaf switch.

If a switch starts with a configuration file, it enters the stack mode by default and cannot join the SVF system as a leaf switch. In this case, set the working mode of the switch to leaf mode using the BIOS menu or working mode menu during the startup process or using commands after the switch starts. The switch will then join the SVF as a leaf switch after each restart.

Procedure

- An unconfigured switch is plug-and-play and does not require any manual configuration.

  By default, the working mode of a switch is auto-negotiation. That is, an unconfigured switch is plug-and-play and does not require any manual configuration.

  NOTE

  If the negotiated working of a switch is leaf, the switch starts in leaf mode. To start the switch in other modes, manually configure the working mode of the switch.

- If a switch starts with a configuration file, use either of the following methods to configure the leaf mode on the switch:
  - Method 1: Use commands to set the leaf mode after the switch starts.
    i. Run the **system-view** command to enter the system view.
    ii. Run the **stack** command to enter the stack management view.
    iii. Run the **switch mode leaf member** { *member-id* | **all** } command to set the working mode of the switch to leaf mode.

    By default, the working mode of a switch is auto-negotiation. That is, the working mode of a switch is determined through auto-negotiation.

  NOTE

  The configured working mode takes effect after the switch restarts.

  This command can also be used on parent switches and takes effect only on the leaf switches that have joined an SVF system.

  - Method 2: Use the BIOS menu to set the leaf mode during the startup process.

i. Log in to the switch through the console port.

ii. During the startup process, press **Ctrl+B** to enter the BIOS menu. The prompt information is "Press CTRL+B to enter BIOS menu".

iii. In the BIOS menu, choose **6. Modify stack parameters** > **2. Modify stack configuration** > **3. Leaf mode** > **3. Return** > **1. Continue to boot** to set the working mode of the switch to leaf mode, and then choose to continue the startup process.

- Method 3: Use the working mode menu to set the leaf mode during the startup process.

i. Log in to the switch through the console port.

ii. During the startup process, press **Ctrl+Y** to enter the working mode menu. The prompt information is "Press CTRL+Y to modify the switch mode".

iii. In the working mode menu, choose **3. Leaf mode** to set the working mode of the switch to leaf mode, and then the switch automatically restarts.

The working mode menu is as follows:

```
Press CTRL+Y to modify the switch mode ..................... 3

If you modify the switch mode, the switch will reboot.

         Switch Mode Menu

      1. Auto-negotiation mode      //Set the working mode to
auto-negotiation.
      2. Stack mode               //Set the working mode to stack.
      3. Leaf mode                //Set the working mode to leaf.
      4. Continue to boot          //Continue to start.

You have 30 seconds to select the menu, the system will continue to boot
when time is up.
Enter your choice(1-4):
```

----End

7.2 Configuring the Type of Leaf Switch Ports to Be Connected to the Parent Switch

Context

By default, leaf switches connect to a parent switch using some default ports. You can also configure the leaf switch ports to be connected to a parent switch according to networking requirements. Subsequently, leaf switches can use only the configured ports to connect to a parent switch.

To learn about the default and configured leaf switch ports to be connected to a parent switch, see Table 7-1.

Table 7-1 List of leaf switch ports to be connected to a parent switch

Version	Leaf Switch Model	Default Port	Configured Port
V100R005C00	• CE5810-48T4S-EI • CE5810-24T4S-EI	Last four 10GE ports	Configured ports are not supported.
	CE6810-48S4Q-EI	Last four 40GE ports or last eight 10GE ports perform auto-negotiation: • If only 10GE ports are connected to a parent switch, 10GE ports are preferred. • If only 40GE ports are connected to a parent switch, 40GE ports are preferred. • If both 10GE and 40GE ports are connected to a parent switch, 40GE ports are preferred.	• Auto-negotiation • 10GE: last eight 10GE ports • 40GE: last four 40GE ports
	CE6850-48T4Q-EI	Last four 40GE ports	• 10GE: last eight 10GE ports • 40GE: last four 40GE ports
V100R005C10 V100R006C00 V200R001C00	• CE5810-48T4S-EI • CE5810-24T4S-EI	4*10GE	Configured ports are not supported.
	• CE5850-48T4S2Q-EI • CE5855-48T4S2Q-EI • CE5855-24T4S2Q-EI	2*40GE	• 4*10GE • 2*40GE
	CE6810-48S4Q-EI	8*10GE and 4*40GE ports perform auto-negotiation: • If only 10GE ports are connected to a parent switch, 10GE ports are preferred. • If only 40GE ports are connected to a	• auto-negotiation • 4*10GE • 8*10GE • 4*40GE

Version	Leaf Switch Model	Default Port	Configured Port
		parent switch, 40GE ports are preferred. • If both 10GE and 40GE ports are connected to a parent switch, 40GE ports are preferred.	
	• CE6810-48S4Q-LI • CE6810-32T16S4Q-LI	4*40GE	• 4*10GE • 8*10GE • 4*40GE
	CE6850-48T4Q-EI	4*40GE	• 4*10GE • 8*10GE • 4*40GE
	CE6810-24S2Q-LI	2*40GE	• 4*10GE • 8*10GE • 2*40GE
	CE6810-48S-LI	8*10GE	• 4*10GE • 8*10GE
	• CE6850-48S6Q-HI • CE6850-48T6Q-HI • CE6851-48S6Q-HI	6*40GE	• 4*10GE • 8*10GE • 4*40GE • 6*40GE

Notes

- N*10GE/N*40GE indicates the last N 10GE/40GE interfaces on the device.
- When auto-negotiated or configured N*10GE or N*40GE ports are connected to a parent switch, these ports cannot be used as common service ports again. For example, when any one of 40GE1/0/1 to 40GE1/0/4 on a CE6810-48S4Q-EI is connected to a parent switch, the other ports of 40GE1/0/1 to 40GE1/0/4 cannot be used as common service ports.
- In V100R005C00 and later versions, the auto-negotiation function takes effect on a CE6810-48S4Q-EI switch only one time. Leaf switches can use only the ports negotiated for the first time to connect to a parent switch. To change the leaf switch ports, you must manually configure the type of leaf switch ports.

Procedure

- Method 1: Use commands to set the type of leaf switch ports to be connected to the parent switch after the switch starts.
  a. Run the **system-view** command to enter the system view.

b. Run the **stack** command to enter the stack management view.

c. Run the **leaf uplink-port type** { **auto-negotiation** | **4*10ge** | **8*10ge** | **2*40ge** | **4*40ge** | **6*40ge** } **member** { *member-id* | **all** } command to set the type of leaf switch ports to be connected to the parent switch.

By default, leaf switches connect to a parent switch using ports listed in Table 7-1.

NOTE

After you configure the type of leaf switch ports to be connected to a parent switch, restart the switch for the configuration to take effect.

This command can also be used on parent switches and takes effect only on the leaf switches that have joined an SVF system.

- Method 2: Use the BIOS menu to set the type of leaf switch ports to be connected to the parent switch during the startup process.

a. Log in to the switch through the console port.

b. During the startup process, press **Ctrl+B** to enter the BIOS menu. The prompt information is "Press CTRL+B to enter BIOS menu".

c. In the BIOS menu, choose **6. Modify stack parameters** > **2. Modify stack configuration** > **1. Auto negotiation mode/3. Leaf mode** > **1. Auto negotiation mode/2. 4*10GE/3. 8*10GE/4. 4*40GE** > **3. Return** > **1. Continue to boot** to set the type of leaf switch ports to be connected to the parent switch, and then choose to continue the startup process.

This step uses a CE6810-48S4Q-EI as an example and may vary according to device models.

- Method 3: Use the working mode menu to set the working mode during the startup process.

a. Log in to the switch through the console port.

b. During the startup process, press **Ctrl+Y** to enter the working mode menu. The prompt information is "Press CTRL+Y to modify the switch mode".

c. In the working mode menu, choose **4. Continue to boot** > **1. Auto negotiation mode/2. 4*10GE/3. 8*10GE/4. 4*40GE** to set the type of leaf switch ports to be connected to the parent switch.

This step uses a CE6810-48S4Q-EI as an example and may vary according to device models.

The working mode menu is as follows:

```
Press CTRL+Y to modify the switch mode ..................... 3

If you modify the switch mode, the switch will reboot.

         Switch Mode Menu

      1. Auto-negotiation mode      //Set the working mode to auto-negotiation.
      2. Stack mode                //Set the working mode to stack.
      3. Leaf mode                 //Set the working mode to leaf.
      4. Continue to boot           //Continue to start the device.

You have 30 seconds to select the menu, the system will continue to boot when
time is up.
Enter your choice(1-4):4

Modify the uplink port type if needed.
```

```
         Uplink Port Menu

      1. Auto-negotiation mode      //Configure the uplink port to work in
auto-negotiation mode.
      2. 4*10GE                   //Configure the last four 10GE ports as
uplink ports.
      3. 8*10GE                   //Configure the last eight 10GE ports as
uplink ports.
      4. 4*40GE                   //Configure the last four 40GE ports as
uplink ports.
      5. Continue to boot          //Continue to start the device.
Enter your choice(1-5):
```

----End

8 Connecting Parent and Leaf Switches

Context

After completing SVF configurations, connect the leaf switches to the parent switches to set up an SVF system. If you change the working mode of a leaf switch through commands or change the type of the port that connects a leaf switch to a parent switch, restart the leaf switch to make the modification take effect.

NOTE

- One leaf switch can connect to only one SVF system through only one fabric port.
- During the setup of an SVF system, there is no order between the software configuration and cable connection. You can perform the software configuration and then connect cables; alternatively, connect cables and then perform the software configuration.

9 Checking the Configuration

Context

After completing SVF configuration and connecting the parent and leaf switches, you can log in to a parent switch to check whether the SVF system is set up successfully.

Procedure

- Run the **display svf** [**member** *member-id*] command to check information about SVF member switches.
- Run the **display svf configuration** [**all** | **leaf** | **virtual-system** *vs-name* [**leaf**]] command to check the SVF configuration.
- Run the **display svf topology** [**all** | **link** | **neighbor** | **virtual-system** *vs-name* [**link** | **neighbor**]] command to check SVF topology information.
- Run the **display svf troubleshooting** [**all** | **virtual-system** *vs-name*] [**history**] command to check the SVF fault events.
- Run the **display svf link-state last-down-reason** [**all** | **virtual-system** *vs-name*] command to check the reason why SVF link protocol is Down.

----End

Follow-up Procedure

In an SVF system, downlink service ports of leaf switches have the BPDU protection function. If a downlink service port receives BPDU packets, the service port will enter the Error-Down state to prevent loops (The device records the status of an interface as Error-Down when it detects that a fault occurs. The interface in Error-Down state cannot receive or send packets and the interface indicator is off.). In the following example, a service port is in Error-Down state.

```
<HUAWEI> display interface ge 101/1/0/1
GE101/1/0/1 current state : ERROR DOWN(leaf-mstp) (ifindex: 2280)
Line protocol current state : DOWN
......
```

After a service port of a leaf switch enters the Error-Down state, check whether the peer port of the service port runs STP. If so, disable STP on the peer port or enable STP on the service port and leaf switch.

You can recover ports from the Error-Down state using either of the following methods:

- Manually recover ports from the Error-Down state (after the ports become Error-Down).

  Run the **shutdown** and then **undo shutdown** commands in the interface view to recover ports from the Error-Down state.
- Enable ports to automatically recover from the Error-Down state (before the ports become Error-Down).

  To enable service ports to automatically recover from the Error-Down state, run the **error-down auto-recovery cause leaf-mstp interval** *interval-value* command in the system view to enable ports in Error-Down state to become Up automatically and set the delay after which ports become Up automatically.

NOTE

This method takes effect only for the ports that become Error-Down after this command is executed but not for those that have been in Error-Down state before this command is executed.

10 Maintaining an SVF System

10.1 Configuring a Heartbeat Timeout Interval for SVF Links

Context

In an SVF system, the parent and leaf switches exchange heartbeat packets to maintain link connections between them. If no heartbeat packet is transmitted on an SVF link within the timeout interval, the SVF link is considered abnormal. If SVF links are unstable, you can increase the heartbeat timeout interval of SVF links to maintain SVF system stability. If SVF links are stable, you can shorten the heartbeat timeout interval of SVF links to ensure that traffic can be switched in a timely manner in case of link failures.

Procedure

Step 1 Run the **system-view** command to enter the system view.

Step 2 Run the **interface fabric-port** *port-id* command to enter the fabric port view.

Step 3 Run the **port holdtime** *hold-time* command to set a heartbeat timeout interval for SVF links of the fabric port.

By default, the heartbeat timeout interval of SVF links is 40 seconds.

Step 4 Run the **commit** command to commit the configuration.

----End

10.2 Configuring a Description for a Leaf or Parent Switch

Context

You can configure a description for a leaf or parent switch to help you manage and identify the switch.

Procedure

Step 1 Run the **system-view** command to enter the system view.

Step 2 Run the **stack** command to enter the stack management view.

Step 3 Configure a description for a leaf or parent switch.

- Run the **leaf member** *member-id* **description** *description* command to configure a description for a leaf switch.

  By default, a leaf switch does not have a description.
- Run the **spine description** *description* command to configure a description for a parent switch.

  By default, a parent switch does not have a description.

Step 4 Run the **commit** command to commit the configuration.

----End

10.3 Configuring an Alarm Threshold for the Number of Available Fabric Port Links

Context

To ensure fabric port link bandwidth and reliability, you can set an alarm threshold for the number of available fabric port links. When the number of available fabric port links is smaller than the alarm threshold because some links fail, the system generates an alarm and sets downlink service ports of leaf switches to the error-down state. When the number of available fabric port links is larger than or equal to the alarm threshold, the system generates a clear alarm and releases these downlink service ports from the error-down state.

Procedure

Step 1 Run the **system-view** command to enter the system view.

Step 2 Run the **interface fabric-port** *port-id* command to enter the fabric port view.

Step 3 Run the **fabric min-links threshold** *alarm-threshold* command to set an alarm threshold for the number of available fabric port links.

The default alarm threshold is 1.

Step 4 (Optional) Run the **port fabric-link-threshold trigger error-down** command to configure a port to enter the Error-Down state when the number of available fabric port links is smaller than the threshold.

By default, a leaf switch port enters the Error-Down state when the number of available fabric port links is smaller than the threshold.

This command takes effect only on the fabric port where it is executed.

Step 5 Run the **commit** command to commit the configuration.

----End

Follow-up Procedure

By default, when the number of available fabric port links is smaller than the configured alarm threshold, service ports connecting to leaf switches in the fabric port will enter the Error-Down state (The device records the status of an interface as Error-Down when it detects that a fault occurs. The interface in Error-Down state cannot receive or send packets and the interface indicator is off.). In the following example, a service port is in Error-Down state.

```
<HUAWEI> display interface 10ge 101/1/0/1
10GE101/1/0/1 current state : ERROR DOWN(fabric-uplink-threshold) (ifindex: 12)
Line protocol current state : DOWN
......
```

After service ports of a leaf switch enter the Error-Down state, check whether a link fault causes the number of available fabric port links to be smaller than the configured alarm threshold. If so, rectify the link fault. After the link recovers, service ports of the leaf switch will recover from the Error-Down state automatically. If not, recover service ports from the Error-Down state using either of the following methods:

- Add fabric port links or run the **fabric min-links threshold** *alarm-threshold* command to reduce the alarm threshold for the number of available fabric port links. This configuration ensures that the number of available fabric port links is larger than or equal to the configured alarm threshold. Subsequently, service ports of the leaf switch can automatically recover from the Error-Down state.
- Run the **undo port fabric-link-threshold trigger error-down** command to disable the function that configures service ports of a leaf switch to enter the Error-Down state when the number of available fabric port links is smaller than the threshold. After this function is disabled, service ports of leaf switches will automatically recover from the Error-Down state and will not enter the Error-Down state because the number of fabric port links is smaller than the threshold.

10.4 Enabling the Single-Homing Alarm Function for Leaf Switches

Context

On a network requiring high reliability, each leaf switch needs to connect to multiple LPUs of the parent switch. After the single-homing alarm function is enabled, the SVF system generates an alarm if a leaf switch connects to only one LPU of the parent switch.

Procedure

Step 1 Run the **system-view** command to enter the system view.

Step 2 Run the **stack** command to enter the stack management view.

Step 3 Run the **leaf single-homed alarm enable** command to enable the single-homing alarm function for leaf switches.

By default, the single-homing alarm function is disabled.

Step 4 Run the **commit** command to commit the configuration.

----End

11 Configuration Examples

11.1 Example for Deploying M-LAG Between SVF Systems Through V-STP

Networking Requirements

A new data center network has many access devices, which complicate management and configuration of the access layer. Unified management and configuration of access devices are required to reduce the management cost.

SVF systems can be set up on the network to meet the preceding requirements. SwitchA and SwitchB in Figure 11-1 are two CE12800 switches and function as parent switches. Each parent switch connects to multiple leaf switches to form an SVF system. In addition, M-LAG needs to be configured between the two SVF systems to improve traffic transmission reliability.

In this example, leaf switches are CE5810-48T4S-EI switches.

Figure 11-1 SVF networking

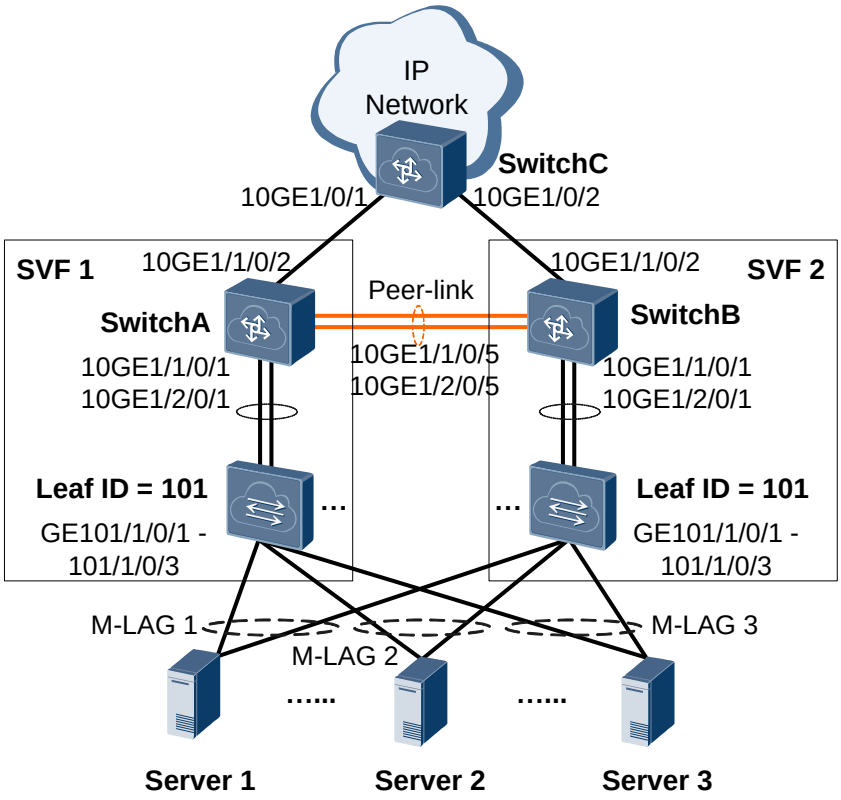

Configuration Roadmap

The configuration roadmap is as follows:

1. Enable the SVF function on parent switches and restart the parent switches.
2. Configure fabric ports on parent switches to connect to leaf switches.
3. Connect leaf switches to the parent switches.
4. Configure M-LAG between the two SVF systems to ensure network reliability.

Procedure

Step 1 Enable the SVF function on parent switches and restart the parent switches.

Enable the SVF function on SwitchA and SwitchB and restart the switches.

```
<HUAWEI> system-view
[~HUAWEI] sysname SwitchA
[*HUAWEI] commit
[~SwitchA] mac-address aging-time 1800
[*SwitchA] commit
[~SwitchA] stack
[~SwitchA-stack] switch svf enable
Warning: Current configuration will be converted to the next startup saved-configuration
file of svf mode.
System will reboot. Continue? [Y/N]: y
```

```
<HUAWEI> system-view
[~HUAWEI] sysname SwitchB
[*HUAWEI] commit
[~SwitchB] mac-address aging-time 1800
[*SwitchB] commit
[~SwitchB] stack
[~SwitchB-stack] switch svf enable
Warning: Current configuration will be converted to the next startup saved-configuration
file of svf mode.
System will reboot. Continue? [Y/N]: y
```

Step 2 After the parent switches restart, log in to them to configure the SVF function.

Configure fabric port 1 on each of SwitchA and SwitchB, add 10GE1/1/0/1 and 10GE1/2/0/1 to fabric port 1, and bind fabric port 1 to leaf ID 101.

```
<SwitchA> system-view
[~SwitchA] interface fabric-port 1
[*SwitchA-Fabric-Port1] port bind member 101
[*SwitchA-Fabric-Port1] port member-group interface 10ge 1/1/0/1 1/2/0/1
[*SwitchA-Fabric-Port1] quit
[*SwitchA] commit

<SwitchB> system-view
[~SwitchB] interface fabric-port 1
[*SwitchB-Fabric-Port1] port bind member 101
[*SwitchB-Fabric-Port1] port member-group interface 10ge 1/1/0/1 1/2/0/1
[*SwitchB-Fabric-Port1] quit
[*SwitchB] commit
```

Step 3 Connect leaf switches to the parent switches and power on the leaf switches. The leaf switches connect to the parent switches through uplink 10GE ports.

NOTE

- This example assumes that the leaf switches are unconfigured switches, work in auto-negotiation mode, and can join the SVF system through auto-negotiation. No manual configuration is required on the leaf switches.
- If the leaf switches have the startup configuration file, you need to clear the configuration in the next startup configuration file of the leaf switches or configure the working mode of the leaf switches to the leaf mode before restarting the switches (in auto-negotiation mode). For details, see 7.1 Configuring the Working Mode of Leaf Switches.

Step 4 Check whether the SVF systems are successfully set up.

Check information about the SVF member switches. The command output on SwitchA is used as an example.

```
[~SwitchA] display svf
Total Spine Number : 1
Total Leaf Number  : 1
-------------------------------------------------------------------------
MemberID Role    MAC            Device Type       Description
-------------------------------------------------------------------------
1       Spine    0019-7459-3300   CE12804
101      Leaf    04f9-388d-e680   CE5810-48T4S-EI
-------------------------------------------------------------------------
```

Step 5 Configure the V-STP, DFS group, peer-link, and M-LAG interface on SwitchA and SwitchB.

An uplink interface of a server connected to a leaf switch needs to be bound to an aggregated link and the link aggregation mode of the server needs to be consistent with that of the leaf switch.

Configure SwitchA.

```
[~SwitchA] stp mode rstp
[*SwitchA] stp v-stp enable
[*SwitchA] interface loopback 0
[*SwitchA-LoopBack0] ip address 10.1.1.1 32
[*SwitchA-LoopBack0] quit
[*SwitchA] dfs-group 1
[*SwitchA-dfs-group-1] source ip 10.1.1.1
[*SwitchA-dfs-group-1] priority 150
[*SwitchA-dfs-group-1] quit
[*SwitchA] interface eth-trunk 0
[*SwitchA-Eth-Trunk0] trunkport 10ge 1/1/0/5
[*SwitchA-Eth-Trunk0] trunkport 10ge 1/2/0/5
[*SwitchA-Eth-Trunk0] mode lacp-static
[*SwitchA-Eth-Trunk0] peer-link 1
[*SwitchA-Eth-Trunk0] quit
[*SwitchA] vlan batch 11
[*SwitchA] interface eth-trunk 1
[*SwitchA-Eth-Trunk1] extend enable
[*SwitchA-Eth-Trunk1] trunkport ge 101/1/0/1
[*SwitchA-Eth-Trunk1] mode lacp-dynamic
[*SwitchA-Eth-Trunk1] port link-type access
[*SwitchA-Eth-Trunk1] port default vlan 11
[*SwitchA-Eth-Trunk1] dfs-group 1 m-lag 1
[*SwitchA-Eth-Trunk1] quit
[*SwitchA] commit
```

Configure SwitchB.

```
[~SwitchB] stp mode rstp
[*SwitchB] stp v-stp enable
[*SwitchB] interface loopback 0
[*SwitchB-LoopBack0] ip address 10.1.1.2 32
[*SwitchB-LoopBack0] quit
[*SwitchB] dfs-group 1
[*SwitchB-dfs-group-1] source ip 10.1.1.2
[*SwitchB-dfs-group-1] priority 120
[*SwitchB-dfs-group-1] quit
[*SwitchB] interface eth-trunk 0
[*SwitchB-Eth-Trunk0] trunkport 10ge 1/1/0/5
[*SwitchB-Eth-Trunk0] trunkport 10ge 1/2/0/5
[*SwitchB-Eth-Trunk0] mode lacp-static
[*SwitchB-Eth-Trunk0] peer-link 1
[*SwitchB-Eth-Trunk0] quit
[*SwitchB] vlan batch 11
[*SwitchB] interface eth-trunk 1
[*SwitchB-Eth-Trunk1] extend enable
[*SwitchB-Eth-Trunk1] trunkport ge 101/1/0/1
[*SwitchB-Eth-Trunk1] mode lacp-dynamic
[*SwitchB-Eth-Trunk1] port link-type access
```

```
[*SwitchB-Eth-Trunk1] port default vlan 11
[*SwitchB-Eth-Trunk1] dfs-group 1 m-lag 1
[*SwitchB-Eth-Trunk1] quit
[*SwitchB] commit
```

Step 6 On SwitchA and SwitchB, configure a VRRP gateway as the dual-active gateway of access devices.

Both devices in the VRRP group are master devices, and the two devices use the same VRID and virtual IP address so that M-LAG devices use the same virtual IP address and virtual MAC address.

Configure SwitchA.

```
[~SwitchA] interface vlanif 11
[*SwitchA-Vlanif11] ip address 10.2.1.1 24
[*SwitchA-Vlanif11] vrrp vrid 1 virtual-ip 10.2.1.111
[*SwitchA-Vlanif11] quit
[*SwitchA] commit
```

Configure SwitchB.

```
[~SwitchB] interface vlanif 11
[*SwitchB-Vlanif11] ip address 10.2.1.2 24
[*SwitchB-Vlanif11] vrrp vrid 1 virtual-ip 10.2.1.111
[*SwitchB-Vlanif11] quit
[*SwitchB] commit
```

Step 7 Configure OSPF on SwitchA, SwitchB, and SwitchC to ensure Layer 3 connectivity.

The Layer 3 main interface is recommended. If a VLANIF interface is used, you need to configure the multi-instance and add the VLAN to different instances or disable STP on the interface and remove the interface from other VLANs except the VLAN corresponding to the VLANIF interface.

Configure SwitchA.

```
[~SwitchA] interface 10ge 1/1/0/2
[~SwitchA-10GE1/1/0/2] undo portswitch
[*SwitchA-10GE1/1/0/2] ip address 10.3.1.1 24
[*SwitchA-10GE1/1/0/2] quit
[*SwitchA] ospf 1
[*SwitchA-ospf-1] area 0
[*SwitchA-ospf-1-area-0.0.0.0] network 10.1.1.1 0.0.0.0
[*SwitchA-ospf-1-area-0.0.0.0] network 10.2.1.0 0.0.0.255
[*SwitchA-ospf-1-area-0.0.0.0] network 10.3.1.0 0.0.0.255
[*SwitchA-ospf-1-area-0.0.0.0] quit
[*SwitchA-ospf-1] quit
[*SwitchA] commit
```

Configure SwitchB.

```
[~SwitchB] interface 10ge 1/1/0/2
[~SwitchB-10GE1/1/0/2] undo portswitch
[*SwitchB-10GE1/1/0/2] ip address 10.4.1.1 24
[*SwitchB-10GE1/1/0/2] quit
[*SwitchB] ospf 1
```

```
[*SwitchB-ospf-1] area 0
[*SwitchB-ospf-1-area-0.0.0.0] network 10.1.1.2 0.0.0.0
[*SwitchB-ospf-1-area-0.0.0.0] network 10.2.1.0 0.0.0.255
[*SwitchB-ospf-1-area-0.0.0.0] network 10.4.1.0 0.0.0.255
[*SwitchB-ospf-1-area-0.0.0.0] quit
[*SwitchB-ospf-1] quit
[*SwitchB] commit
```

Configure SwitchC.

```
<HUAWEI> system-view
[~HUAWEI] sysname SwitchC
[*HUAWEI] commit
[~SwitchC] interface 10ge 1/0/1
[~SwitchC-10GE1/0/1] undo portswitch
[*SwitchC-10GE1/0/1] ip address 10.3.1.2 24
[*SwitchC-10GE1/0/1] quit
[*SwitchC] interface 10ge 1/0/2
[*SwitchC-10GE1/0/2] undo portswitch
[*SwitchC-10GE1/0/2] ip address 10.4.1.2 24
[*SwitchC-10GE1/0/2] quit
[*SwitchC] ospf 1
[*SwitchC-ospf-1] area 0
[*SwitchC-ospf-1-area-0.0.0.0] network 10.3.1.0 0.0.0.255
[*SwitchC-ospf-1-area-0.0.0.0] network 10.4.1.0 0.0.0.255
[*SwitchC-ospf-1-area-0.0.0.0] quit
[*SwitchC-ospf-1] quit
[*SwitchC] commit
```

Step 8 On SwitchA and SwitchB, associate uplink and downlink interfaces with the Monitor Link group.

Configure SwitchA.

```
[~SwitchA] monitor-link group 1
[*SwitchA-mtlk-group1] port 10ge 1/1/0/2 uplink
[*SwitchA-mtlk-group1] port eth-trunk 1 downlink 1
[*SwitchA-mtlk-group1] quit
[*SwitchA] commit
```

Configure SwitchB.

```
[~SwitchB] monitor-link group 1
[*SwitchB-mtlk-group1] port 10ge 1/1/0/2 uplink
[*SwitchB-mtlk-group1] port eth-trunk 1 downlink 1
[*SwitchB-mtlk-group1] quit
[*SwitchB] commit
```

Step 9 Verify the configuration.

Run the **display dfs-group** command to check M-LAG information.

Check DFS group information.

```
[~SwitchA] display dfs-group 1 m-lag
*             : Local node
Heart beat state : OK
Node 1 *
  Dfs-Group ID   : 1
```

```
  Priority      : 150
  Address       : ip address 10.1.1.1
  State         : Master
  Causation     : -
  System ID     : 0019-7459-3301
  SysName       : SwitchA
  Version       : V100R006C00
  Device Type    : CE12800
Node 2
  Dfs-Group ID   : 1
  Priority      : 120
  Address       : ip address 10.1.1.2
  State         : Backup
  Causation     : -
  System ID     : ac94-8400-df01
  SysName       : SwitchB
  Version       : V100R006C00
  Device Type    : CE12800
```

Check M-LAG information on SwitchA.

```
[~SwitchA] display dfs-group 1 node 1 m-lag brief
* - Local node

M-Lag ID    Interface     Port State   Status
      1    Eth-Trunk 1    Up           active(*)-active
```

Check M-LAG information on SwitchB.

```
[~SwitchA] display dfs-group 1 node 2 m-lag brief
* - Local node

M-Lag ID    Interface     Port State   Status
      1    Eth-Trunk 1    Up           active-active(*)
```

In the preceding information, the value of **Heart beat state** is **OK**, indicating that the heartbeat is normal. SwitchA is used as Node 1, its priority is 150, and its status is **Master**. SwitchB is used as Node 2, its priority is 120, and its status is **Backup**. The value of **Causation** is -, and the values of **Port State** of Node 1 and Node 2 are both **Up**, and the M-LAG status of Node 1 and Node 2 is both **active**, indicating that the M-LAG configuration is correct.

----End

Configuration Files

- Configuration file of SwitchA

```
#
sysname SwitchA
#
dfs-group 1
 priority 150
 source ip 10.1.1.1
#
vlan batch 11
#
mac-address aging-time 1800
```

```
#
stp mode rstp
stp v-stp enable
#
stack
 #
 svf enable
#
interface Vlanif11
 ip address 10.2.1.1 255.255.255.0
 vrrp vrid 1 virtual-ip 10.2.1.111
#
interface Eth-Trunk0
 mode lacp-static
 peer-link 1
#
interface Eth-Trunk1
 port default vlan 11
 extend enable
 mode lacp-dynamic
 dfs-group 1 m-lag 1
#
interface Fabric-Port1
 port bind member 101
#
interface GE101/1/0/1
 eth-trunk 1
#
interface 10GE1/1/0/1
 fabric-port 1
#
interface 10GE1/1/0/2
 undo portswitch
 ip address 10.3.1.1 255.255.255.0
#
interface 10GE1/1/0/5
 eth-trunk 0
#
interface 10GE1/2/0/1
 fabric-port 1
#
interface 10GE1/2/0/5
 eth-trunk 0
#
interface LoopBack0
 ip address 10.1.1.1 255.255.255.255
#
monitor-link group 1
 port 10GE1/1/0/2 uplink
 port Eth-Trunk1 downlink 1
#
ospf 1
 area 0.0.0.0
  network 10.1.1.1 0.0.0.0
  network 10.2.1.0 0.0.0.255
```

```
 network 10.3.1.0 0.0.0.255
#
return
```

- Configuration file of SwitchB

```
#
sysname SwitchB
#
dfs-group 1
 priority 120
 source ip 10.1.1.2
#
vlan batch 11
#
mac-address aging-time 1800
#
stp mode rstp
stp v-stp enable
#
stack
 #
 svf enable
#
interface Vlanif11
 ip address 10.2.1.2 255.255.255.0
 vrrp vrid 1 virtual-ip 10.2.1.111
#
interface Eth-Trunk0
 mode lacp-static
 peer-link 1
#
interface Eth-Trunk1
 port default vlan 11
 extend enable
 mode lacp-dynamic
 dfs-group 1 m-lag 1
#
interface Fabric-Port1
 port bind member 101
#
interface GE101/1/0/1
 eth-trunk 1
#
interface 10GE1/1/0/1
 fabric-port 1
#
interface 10GE1/1/0/2
 undo portswitch
 ip address 10.4.1.1 255.255.255.0
#
interface 10GE1/1/0/5
 eth-trunk 0
#
interface 10GE1/2/0/1
 fabric-port 1
```

```
#
interface 10GE1/2/0/5
 eth-trunk 0
#
interface LoopBack0
 ip address 10.1.1.2 255.255.255.255
#
monitor-link group 1
 port 10GE1/1/0/2 uplink
 port Eth-Trunk1 downlink 1
#
ospf 1
 area 0.0.0.0
  network 10.1.1.2 0.0.0.0
  network 10.2.1.0 0.0.0.255
  network 10.4.1.0 0.0.0.255
#
return
```

- Configuration file of SwitchC

```
#
sysname SwitchC
#
interface 10GE1/0/1
 undo portswitch
 ip address 10.3.1.2 255.255.255.0
#
interface 10GE1/0/2
 undo portswitch
 ip address 10.4.1.2 255.255.255.0
#
ospf 1
 area 0.0.0.0
  network 10.3.1.0 0.0.0.255
  network 10.4.1.0 0.0.0.255
#
return
```