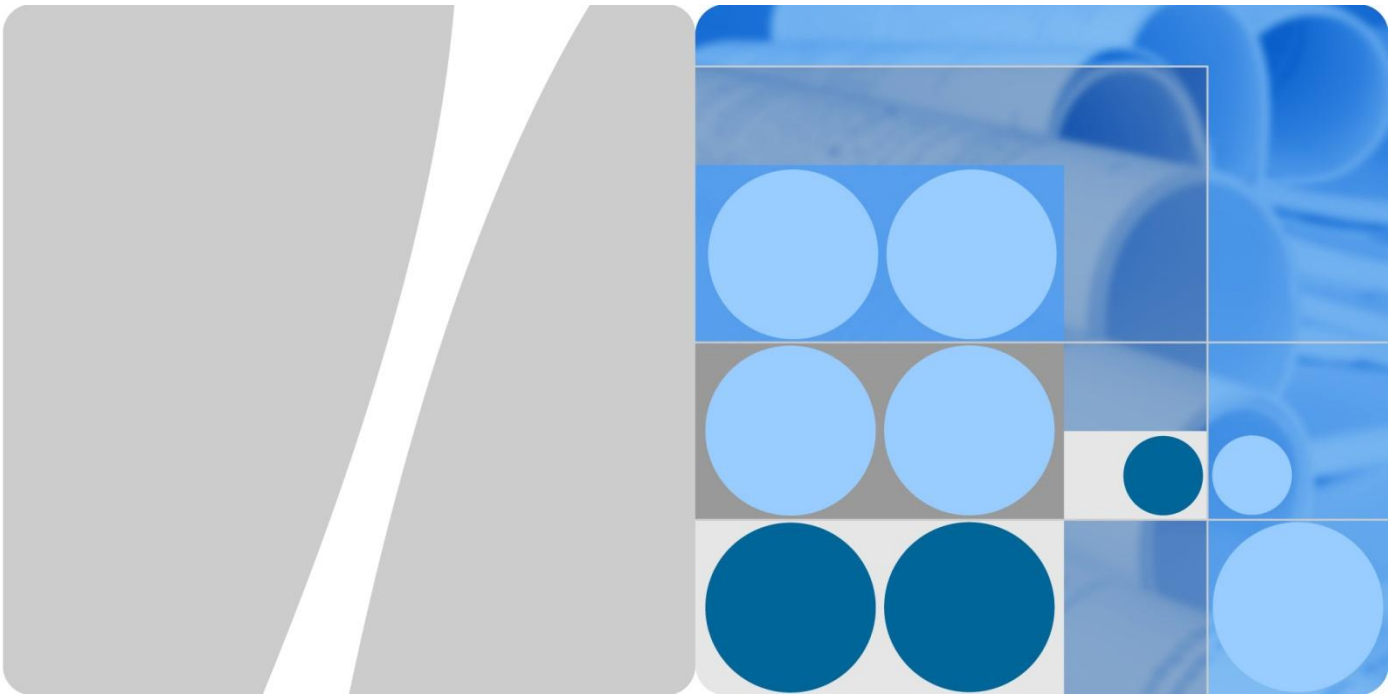

Smart Link & Monitor Link Technology White Paper

Issue 01

Date 2012-10-30

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://enterprise.huawei.com

About This Document

Keywords:

Smart Link, Monitor Link, control VLAN, flush message, rapid switchover, dual uplink, link redundancy

Abstract:

Smart Link is a feature developed to provide effective, reliable link redundancy, load sharing, and fast convergence for dual-uplink networks. Monitor Link is a feature developed to complement the link backup mechanism of Smart Link. By monitoring the uplink, and synchronizing the downlink with the uplink, Monitor Link triggers the switchover between the primary and backup links in a Smart Link group. This document mainly describes the basic concepts, mechanisms, and typical application scenarios of Smart Link and Monitor Link.

Acronyms:

Acronyms	Full spelling
RRPP	Rapid Ring Protection Protocol
STP	Spanning Tree Protocol
SMLK	Smart Link
MTLK	Monitor Link

Contents

Table List

Figure List

1 Smart Link Overview

1.1 Background

Dual-uplink networks are widely used to improve network reliability. As shown in Figure 1-1, Switch A is dual homed to Switch D through Switch B and Switch C.

Figure 1-1 Dual-uplink network

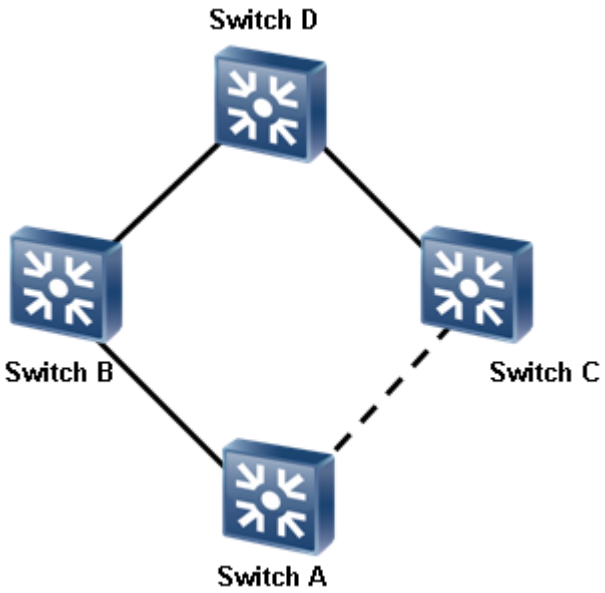

Although dual uplinks can provide link redundancy on the network, the possible loop (Switch A - Switch B - Switch D - Switch C -Switch A) can also cause broadcast storms. Therefore, you need to take measures to prevent loops on the network. In most cases, STP is used to eliminate network loops. However, STP has a long convergence time, and a large amount of traffic is lost during the convergence time. Therefore, STP is not suitable for scenarios demanding short convergence. RRPP is another effective solution used to remove loops. It has shorter convergence time, but it is mostly used on ring networks. Besides, RRPP is hard to configure. To address this problem, HUAWEI developed the Smart Link solution.

1.2 Benefits

Smart Link is dedicated to dual-uplink networks. It delivers the following benefits:

- Keeps one link connected and the other blocked when both links on a dual-uplink network work properly to prevent broadcast storms caused by network loops.
- Switches traffic to the backup link within a few subseconds when the primary link fails to ensure nonstop service forwarding.
- Is easy to configure.

2 Smart Link Implementation

2.1 Basic Concepts in Smart Link

2.1.1 Smart Link Group

A Smart Link group consists of two member ports: master port and slave port. A port can belong to multiple Smart Link groups. Typically, only one port forwards traffic in active state, whereas the other port is blocked in standby state. When a link failure occurs on the port in active state due to port shutdown or Ethernet OAM link errors, the port in standby state becomes active while the former port in active state goes to the blocked state.

Figure 2-1 Smart Link application scenario

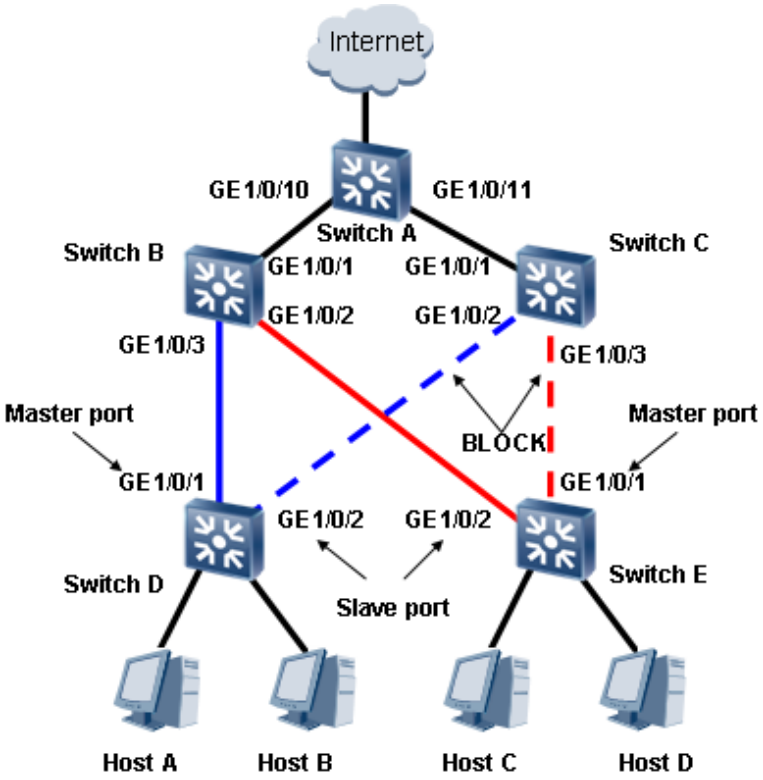

As shown in Figure 2-1, GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 of Switch D form a Smart Link group (marked in blue). GigabitEthernet 1/0/1 is in forwarding state (marked by a continuous line), and GigabitEthernet 1/0/2 is in blocked state (marked by a broken line). GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 of Switch E form another Smart Link group (marked in red). GigabitEthernet 1/0/1 is in blocked state (marked by a broken line), and GigabitEthernet 1/0/2 is in forwarding state (marked by a continuous line).

2.1.2 Master Port

The master port of a Smart Link group is specified using a command. This port can be an Ethernet port (electrical or optical), or an aggregate interface.

In Figure 2-1, the master port in the Smart Link group configured on Switch D is the active port GigabitEthernet 1/0/1. The master port in the Smart Link group on Switch E is the blocked port GigabitEthernet 1/0/1.

2.1.3 Slave Port

The slave port of a Smart Link group is also specified using a command. This port can be an Ethernet port (electrical or optical), or an aggregate interface. The link on which the slave port resides is called the backup link.

In Figure 2-1, the slave port in the Smart Link group on Switch D is the blocked port GigabitEthernet 1/0/2. The slave port in the Smart Link group on Switch E is the active port GigabitEthernet 1/0/2.

2.1.4 Protected VLAN

A protected VLAN carries user data traffic within a Smart Link group. A port can belong to multiple Smart Link groups in different protected VLANs. Each Smart Link group independently calculates the state of its own member ports.

In Figure 2-1, Smart Link Group 1 and Smart Link Group 2 are created on Switch D. Smart Link Group 1 protects VLANs 1 to 10, while Smart Link Group 2 protects VLANs 11 to 20. In this way, traffic of the two VLAN groups is transmitted over different ports.

2.1.5 Control VLAN

Transmit control VLAN

A transmit control VLAN is a VLAN used by a Smart Link group to broadcast flush messages. For more information about flush messages, see section 2.1.6 "Flush Message."

In Figure 2-1, if Switch D and Switch E are enabled to send flush messages and a link switchover occurs, the switches broadcast flush messages in the transmit control VLANs using the new primary links.

Receive control VLAN

A receive control VLAN is used by an upstream device to receive and process flush messages.

In Figure 2-1, if the upstream devices (Switch A, Switch B, and Switch C) of Switch D and Switch E can identify flush messages, and are enabled to receive and process flush messages, when link switchover occurs, the upstream devices process flush messages from the receive control VLANs, and refresh their MAC address forwarding entries and ARP entries.

2.1.6 Flush Message

When a link switchover occurs in a Smart Link group, the original forwarding entries become invalid. Therefore, all devices on the network need to update their MAC address entries and ARP entries. Smart Link sends flush messages to notify devices of updating their MAC address entries and ARP entries.

A flush message uses the IEEE802.3 encapsulation format.

Figure 2-2 Flush message format

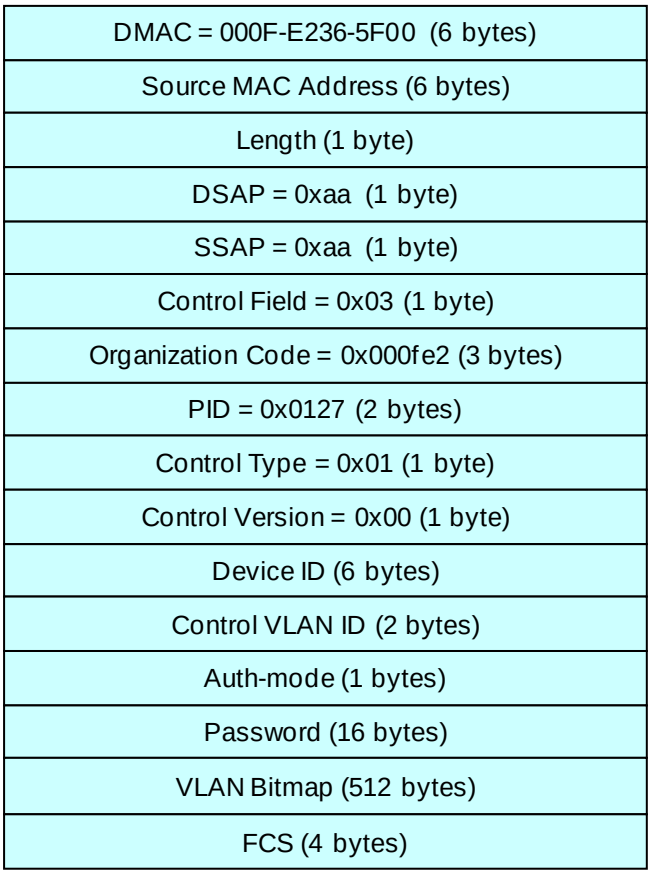

A flush message has the following fields:

- DMAC: indicates an unknown multicast address. 0x010F-E200-0004 is the destination MAC address of a flush message.
- Source MAC Address: indicates the bridge MAC address of the device that sends the flush message.
- Control Type: indicates the control type. Currently, only one control type, 0x01, is available. It is used to delete MAC address entries and ARP entries.
- Control Version: indicates the version number. The current control version is 0x00.
- Device ID: indicates the bridge MAC address of the device that sends a flush message.
- Control VLAN ID: indicates the ID of a transmit control VLAN.
- Auth-mode: indicates the authentication mode. It is used with the password.
- VLAN bitmap: carries a list of VLANs whose MAC address entries need to be updated.

- FCS: is used to check bit errors in frames.

2.2 Smart Link Mechanism

The network shown in Figure 2-3 helps illustrate the descriptions that follow, showing how the Smart Link mechanism causes the link status to change from normal, to faulty, and then to recovery.

Figure 2-3 Smart Link mechanism

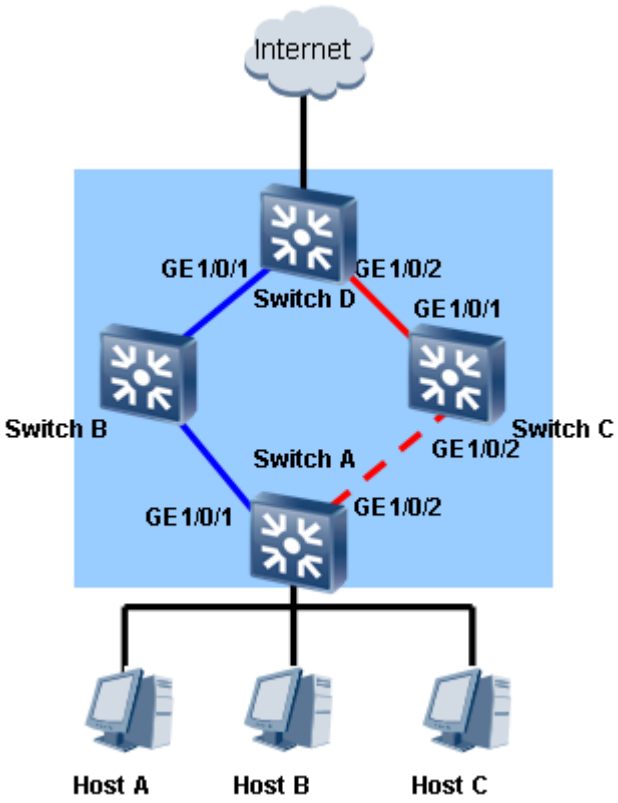

2.2.2 Smart Link Mechanism When Uplinks Work Properly

GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 on Switch A form a Smart Link group, with the former as the master port and the latter as the slave port. When both uplinks are working properly, the master port is in forwarding state,the slave port is in standby state, and the links on which the two ports reside respectively are called the primary link and the backup link. Data is transmitted along the link indicated by the blue line. There is no loop on the network, so there is no broadcast storm.

2.2.3 Smart Link Mechanism When a Link Switchover Occurs

When the primary link on Switch A fails, the master port GigabitEthernet 1/0/1 transits to the standby state, while the slave port GigabitEthernet 1/0/2 transits to the forwarding state. After the link switchover, the MAC address entries and ARP entries kept on the devices on the network may become incorrect, and need to be updated. You can use the following two methods to update MAC address entries and ARP entries:

Flush message sending

This method is used when the upstream devices (such as Switch B, Switch C, and Switch D in Figure 2-3) support Smart Link and are able to identify flush messages.

To implement a rapid link switchover, enable Switch A to send flush messages, and enable all upstream devices' ports on the dual-uplink network to receive and process flush messages. Switches process the flush messages as follows:

1. After a link switchover occurs on Switch A, Switch A sends flush messages along the new primary link of GigabitEthernet 1/0/2. The VLAN Bitmap field is filled in by the protected VLAN IDs of the Smart Link group where GigabitEthernet 1/0/1 is in forwarding state before the link switchover. The Control VLAN ID field is filled in by the transmit control VLAN ID of the Smart Link group.
2. When an upstream device receives a flush message, the upstream device determines whether the transmit VLAN ID carried in the flush message is on the receive control VLAN list configured on the receive port. If not, the upstream device forwards the flush message. If the transmit VLAN ID is on the switch's list, the device retrieves the VLAN Bitmap information from the flush message, and deletes the MAC address entries and ARP entries learned in the VLANs specified in the VLAN Bitmap field.

After these changes, Switch D behaves as follows after receiving a data packet destined for Switch A:

- If Layer 2 forwarding is required, Switch D broadcasts the packet at Layer 2.
- If Layer 3 forwarding is required, Switch D updates ARP entries using the ARP probe mechanism, and then forwards the packet.

NOTE

- All ports on the dual-uplink network must be added to the transmit control VLAN so that flush messages can be correctly transmitted within the transmit control VLAN. Otherwise, flush messages cannot be transmitted or forwarded.
- It is recommended that flush messages be sent in tagged mode. To send flush messages in untagged mode, ensure that the default VLAN of the peer port is the same as the transmit control VLAN. Otherwise, flush messages are not transmitted within the transmit control VLAN.

Uplink Traffic Triggering

This method is used when Huawei switches connect to non-Huawei devices or devices not supporting Smart Link devices that do not support Smart Link.. Updating MAC address entries and ARP entries is triggered by uplink traffic.

- If no uplink traffic from Switch A arrives at Switch D to trigger updating of MAC address entries and ARP entries, Switch D continues to forward traffic out of GigabitEthernet 1/0/1 that are supposed to go to Switch A. The MAC address entries and ARP entries will not be updated until their aging timers expire.
- Because the MAC address entries and ARP entries on Switch A are also incorrect, Switch A cannot send uplink traffic until its MAC address entries and ARP entries age out and new MAC address entries and ARP entries are learned. When uplink traffic from Switch A arrives at Switch D through GigabitEthernet 1/0/2, Switch D updates its own MAC address forwarding entries and ARP entries. After that, when Switch D receives a data packet destined for Switch A, it forwards the packet out if GigabitEthernet 1/0/2. This time the packet can reach Switch A through Switch C.

When the first mechanism is used, the upstream devices update their MAC address entries and ARP entries based on received flush messages, but not wait until these entries age out. In this way, the time required for updating MAC address forwarding entries and ARP entries is

reduced. Normally, a link switchover is completed within a few subseconds Without traffic loss.

2.2.4 Smart Link Mechanism When a Link Recovers

Smart Link supports two working modes in which link recovery mechanisms are different:

- Role preemption: When the primary link recovers, the master port enters the forwarding state and takes over the traffic, whereas the slave port enters the standby state. The slave port transits from standby to forwarding only when the primary link fails.
- Non-role preemption: When the primary link recovers, the slave port remains in forwarding state, whereas the master port remains in standby state to keep the traffic stable.

In Figure 2-3, if role preemption is configured on the Smart Link group on Switch A, when the link of GigabitEthernet 1/0/1 on Switch A recovers, GigabitEthernet 1/0/2 is immediately blocked and transits to the standby state, whereas GigabitEthernet 1/0/1 transits to the forwarding state. If non-role preemption is configured, when the link of GigabitEthernet 1/0/1 on Switch A recovers, GigabitEthernet 1/0/1 remains in standby state, and no link switchover occurs, keeping the traffic stable.

2.3 Load Balancing Through Smart Link

Data traffic of multiple VLANs may exist on a dual-uplink network. By configuring Smart Link, you can implement load sharing by transmitting traffic of different VLANs along different paths. This goal is achieved by configuring the two ports on the dual uplinks as members of two Smart Link groups (each having its own protected VLANs). Each group has its own protected VLANs, and each port has opposite roles in different Smart Link groups..

You are recommended to configure role preemption in the Smart Link groups. Otherwise, traffic cannot be load balanced among two links when a link switchover occurs. When non-role preemption is configured in Smart Link groups, traffic is load balanced on the two uplinks in the beginning. When either uplink fails, all traffic is transmitted along one path and continues that way even after the failed uplink recovers. Therefore, load balancing cannot last on a network using non-role preemption.

In Figure 2-3, two Smart Link groups are created on Switch A and have different protected VLANs. Both of them are configured with the role preemption mode. Smart Link Group 1's master port is GigabitEthernet 1/0/1, slave port is GigabitEthernet 1/0/2, and protected VLANs are VLAN 1 through VLAN 10; Smart Link Group 2's master port is GigabitEthernet 1/0/2, slave port is GigabitEthernet 1/0/1, and protected VLANs are VLAN 11 through VLAN 20. Both Smart Link groups' master ports are in forwarding state. The traffic of VLAN 1 through VLAN 10 is transmitted along the link indicated by the blue line, and the traffic of VLAN 11 through VLAN 20 is transmitted along the link indicated by the red line. In this way, VLAN traffic is load balanced on the two uplinks.

3 Monitor Link Overview

3.1 Background

Figure 3-1 Monitor Link

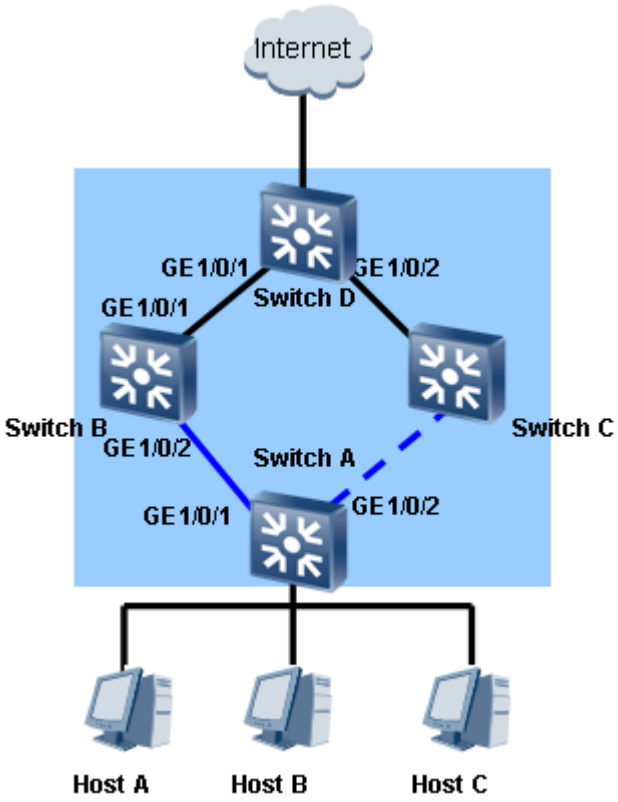

In Figure 3-1, a Smart Link group is configured on Switch A for link redundancy, with GigabitEthernet 1/0/1 as the master port and GigabitEthernet 1/0/2 as the slave port.

When the primary link on which GigabitEthernet 1/0/1 resides fails, traffic switches to the backup link on which GigabitEthernet 1/0/2 resides within a few subseconds. Smart Link delivers reliable link redundancy and rapid convergence.

However, when the link on Switch B's GigabitEthernet 1/0/1 fails, link switchover will not happen in the Smart Link group configured on Switch A because the link on which the master port GigabitEthernet 1/0/1 resides is working properly. However, traffic of Switch A can no longer reach Switch D through GigabitEthernet 1/0/1. Monitor Link technology solves this problem.

3.2 Benefits

Monitor Link complements the Smart Link feature. By monitoring the uplink and synchronizing the downlink with the uplink, Monitor Link triggers a switchover between the primary and backup links in a Smart Link group, completing the link redundancy mechanism of Smart Link.

4 Monitor Link Implementation

4.1 Basic Concepts in Monitor Link

4.1.1 Monitor Link Group

A Monitor Link group is a set of uplink and downlink ports. Downlink ports adapt to the state changes of uplink ports.

Figure 4-1 Monitor Link

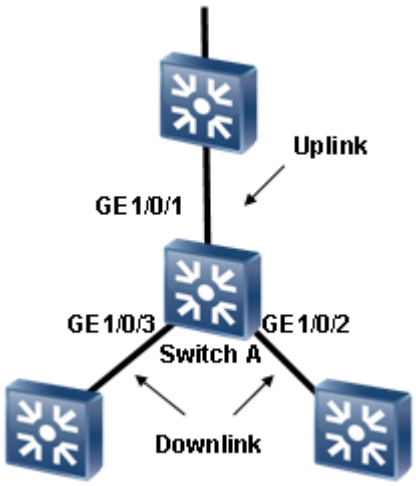

In Figure 4-1, ports GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3 of Switch A form a Monitor Link group.

4.1.2 Uplink Port

An uplink port is a monitored port in a Monitor Link group, and is specified using a command. The port can be an Ethernet port (electrical or optical), or an aggregate interface.

In Figure 4-1, GigabitEthernet 1/0/1 of Switch A is the only uplink port of the Monitor Link group configured on the device.

For a Monitor Link group that has multiple uplink ports, as long as at least one of its uplink ports is in forwarding state, the Monitor Link group is Up. However, if all uplink ports of the

Monitor Link group fail, the Monitor Link group goes Down and all the downlink ports are shut down. If no uplink port is specified in a Monitor Link group, the system considers the Monitor Link group's uplink ports faulty, and shuts down all the downlink ports in the Monitor Link group.

4.1.3 Downlink Port

A downlink port is a monitoring port in a Monitor Link group, and is specified using a command. The port can be an Ethernet port (electrical or optical), or an aggregate interface.

In Figure 4-1, GigabitEthernet 1/0/2 and GigabitEthernet 1/0/3 of Switch A are two downlink ports of the Monitor Link group configured on the device.

NOTE

- When a Monitor Link group's uplink ports recover, only downlink ports that were blocked due to uplink port failure will be brought up. Downlink ports manually shut down will not be brought up automatically. The failure of a downlink port does not affect the uplink ports or other downlink ports.

4.2 Monitor Link Mechanism

In Figure 4-2, a Smart Link group is configured on Switch A. This configuration helps provide reliable access to the Internet for the hosts. GigabitEthernet 1/0/1 is the master port of the Smart Link group, and is in forwarding state. GigabitEthernet 1/0/2 is the slave port.

Figure 4-2 Monitor Link mechanism

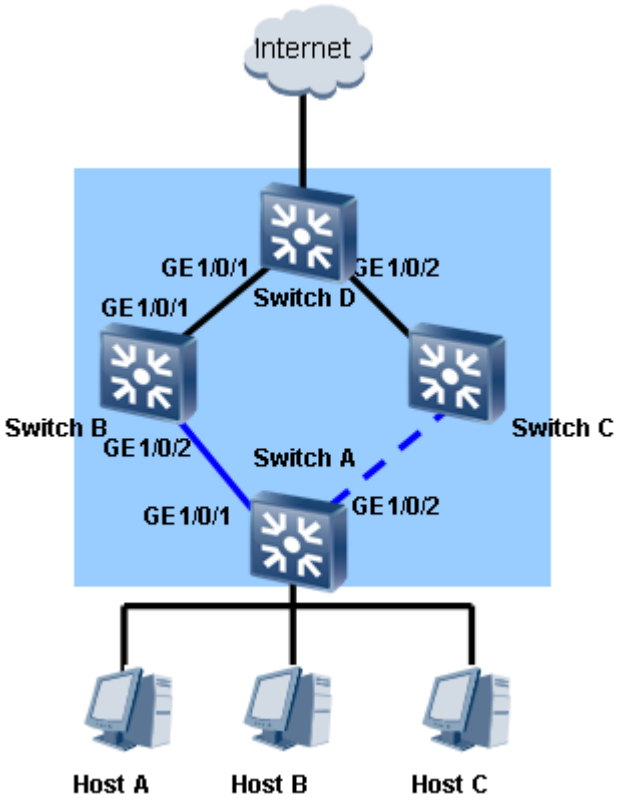

To prevent traffic interruption due to a failure of the link on which GigabitEthernet 1/0/1 of Switch B resides, configure a Monitor Link group on Switch B, and specify GigabitEthernet 1/0/1 as the uplink port, and GigabitEthernet 1/0/2 as the downlink port.

When the link on which GigabitEthernet 1/0/1 of Switch B resides fails, the Monitor Link group shuts down its downlink port GigabitEthernet 1/0/2, triggering a link switchover in the Smart Link group configured on Switch A.

When the link on which GigabitEthernet 1/0/1 of Switch B resides recovers, the downlink port GigabitEthernet 1/0/2 is also brought up. This change triggers another link switchover in the Smart Link group if role preemption is configured in the Smart Link group on Switch A.

Monitor Link and Smart Link implement reliable link redundancy and fast convergence for dual-uplink networks.

5 Application Scenarios

5.1 Smart Link Combined with Monitor Link

Figure 5-1 shows a typical dual-uplink network for Smart Link and Monitor Link.

Figure 5-1 Collaboration between Smart Link and Monitor Link

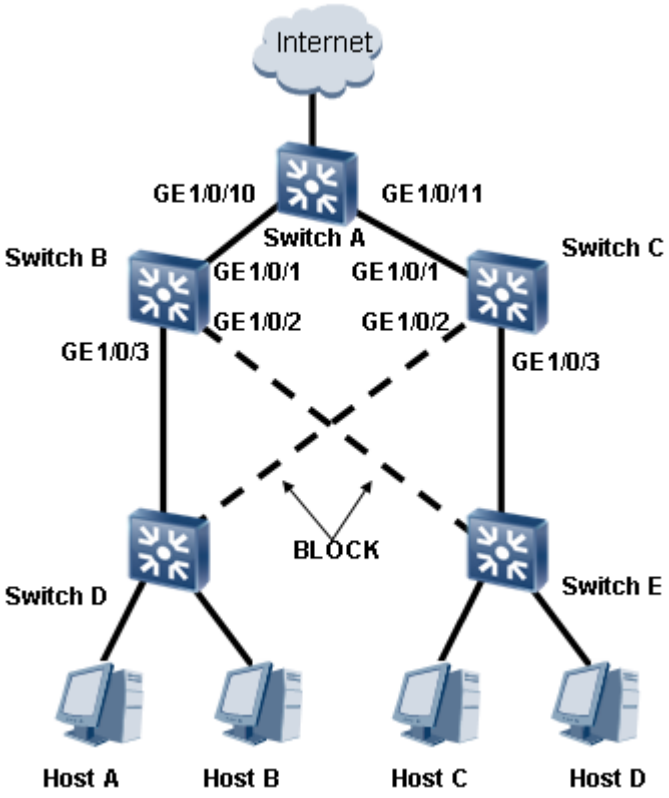

In Figure 5-1, Smart Link is configured on Switch D and Switch E. You can configure multiple Smart Link groups on the devices and assigning different protected VLANs to them so that traffic of protected VLANs that belong to different Smart Link groups is transmitted along different paths, achieving load sharing. When the link between Switch B and Switch D or between Switch C and Switch E fails, the Smart Link groups can immediately detect the link failure and switch services over links.

To enable Switch D or Switch E to detect the failure of the link between Switch A and Switch B or Switch C, configure a Monitor Link group on Switch B or Switch C, with GigabitEthernet 1/0/1 as the uplink port, and GigabitEthernet 1/0/2 and GigabitEthernet 1/0/3 as the downlink ports.

Upon detecting the failure of the link on which the uplink port resides, the Monitor Link group shuts down its downlink ports, triggering a link switchover in the Smart Link group configured on Switch D or Switch E. When the uplink port or link failure recovers, the downlink ports are brought up automatically so that Switch D or Switch E detects the state changes of the link between Switch A and Switch B or Switch C.

5.2 Cascaded Networking with Smart Link and Monitor Link

Figure 5-2 shows the cascaded networking in which Smart Link and Monitor Link are configured to achieve reliable link redundancy.

A Smart Link group member port can be assigned to a Monitor Link group as its uplink member port. You can use Smart Link and Monitor Link technologies together to build cascaded networking.

You need to configure the two member ports of a Smart Link group as the uplink ports of a Monitor Link group, with the peer of the Monitor Link group's downlink port as the master or slave port of another Smart Link group, as shown in Figure 5-2.

Figure 5-2 Cascaded networking

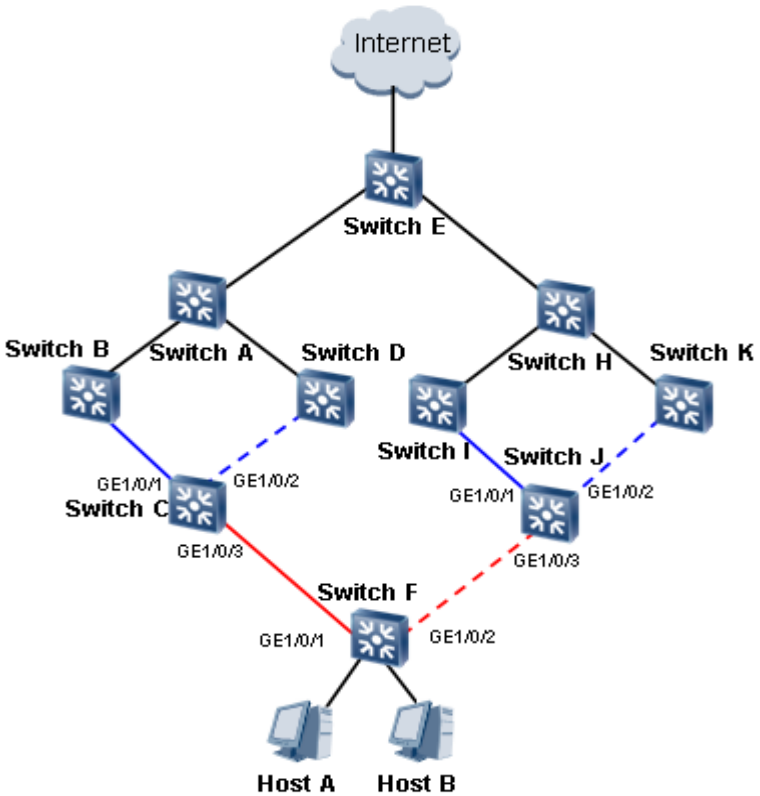

Table 5-1 lists the roles of ports on Switch C, Switch J, and Switch F in Smart Link and Monitor Link shown in Figure 5-2.

Table 5-1 Smart Link and Monitor Link port roles

Switch	Smart Link Group 1		Monitor Link Group 1	
	Master Port	Slave Port	Uplink Port	Downlink Port
Switch C	GE1/0/1	GE1/0/2	GE1/0/1, GE1/0/2	GE1/0/3
Switch J	GE1/0/1	GE1/0/2	GE1/0/1, GE1/0/2	GE1/0/3
Switch F	GE1/0/1	GE1/0/2	No Monitor Link configured	

In Figure 5-2, the red lines indicate the first-level Smart Link backup uplinks, and the blue lines indicate the second-level Smart Link backup uplinks.

5.3 Smart Link Combined with RRPP

Figure 5-3 shows a hybrid network of Smart Link and RRPP.

RRPP is enabled on Switch A, Switch B, Switch C, and Switch D to provide link redundancy.

If STP is used to provide link redundancy, you need to enable STP on all ports connecting Switch C, Switch D, and Switch E. RRPP cannot be used with STP. Because RRPP is enabled on the two ports connecting Switch C and Switch D, you can achieve link redundancy for Switch E by configuring a Smart Link group on Switch E. This configuration is simpler than configuring RRPP subrings on Switch C, Switch D, and Switch E, and applies to the scenario where Switch E does not support RRPP.

Figure 5-3 Smart Link-RRPP hybrid network