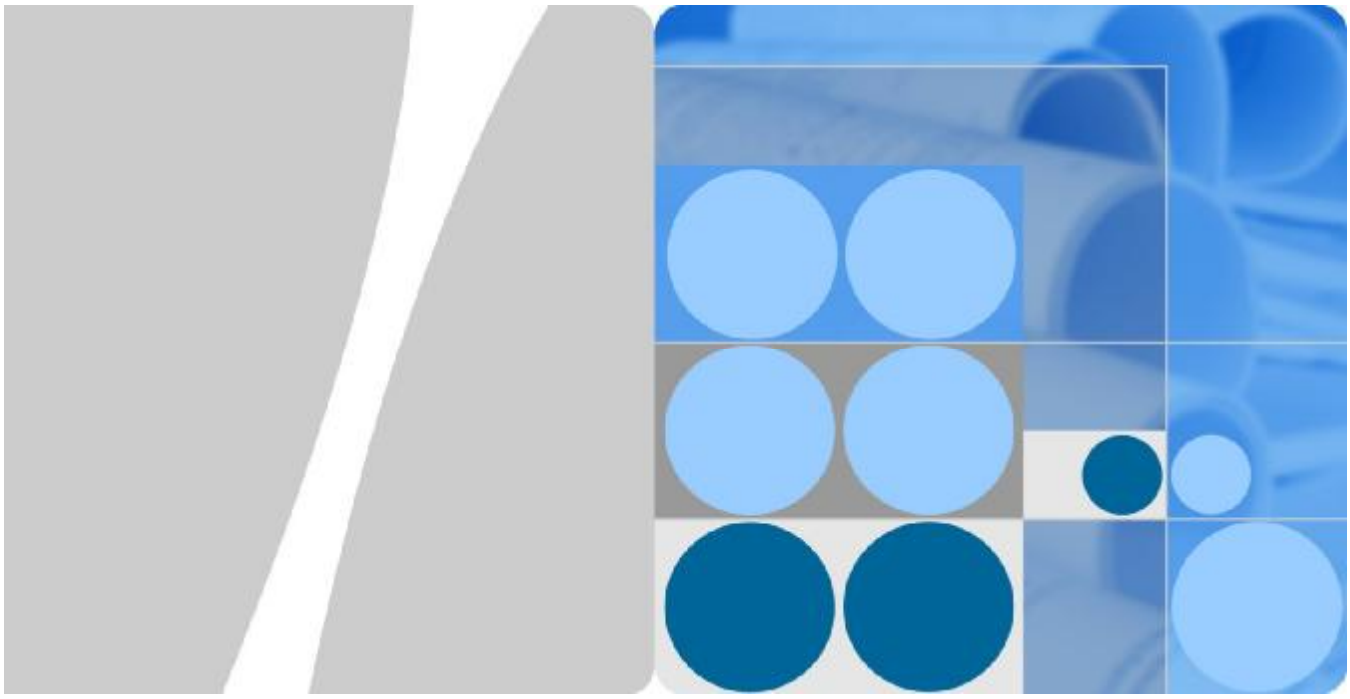

S Series Switches

SSP Technology White Paper

Issue 01

Date 2015-11-16

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

SSP Technology White Paper

Keyword: service splitting platform, SSP, service splitting, service splitting device

Abstract: Service splitting redirects network traffic to the specified destination. It resolves traffic over backbone links to the granularities that can be processed by the packet analysis system. Huawei service splitting platform (SSP) uses ACL to classify packets based on five IP attributes (source/destination IP addresses, source/destination ports, and protocol type), MPLS label, and characteristics field, forwards packets according to the specified action (for example, remove MPLS, GRE, and GTP labels), and redirects packets to the packet analysis system through a group of outbound interfaces.

Acronyms and Abbreviations:

Acronyms and Abbreviations	Full Name
SSP	Service Splitting Platform
GRE	Generic Routing Encapsulation
GTP	GPRS Tunneling protocol
ACL	Access Control List
MQC	Modular QoS Command-Line
ECMP	Equal-Cost Multi-Path Routing

Contents

1 Overview

Thanks to the development of IP and transport technologies, the number of users and network traffic on the Internet is continuously increasing. It is important to control and manage the Internet to prevent attackers from threatening national, enterprise, or home networks.

Huawei Service Splitting Platform (SSP) captures packets transmitted on the network, and forwards traffic information to the analyser for processing. A single analyser cannot process all traffic information because backbone network links provide high bandwidth. Therefore, the SSP splits traffic information into small pieces and forwards them to different analysers.Generally, the SSP is located at the international egress, provincial egress, or MAN egress of a backbone network or the industrial monitoring network for government, army, or school. Figure 1-1 shows a typical network of SSP.

Figure 1-1 SSP network diagram

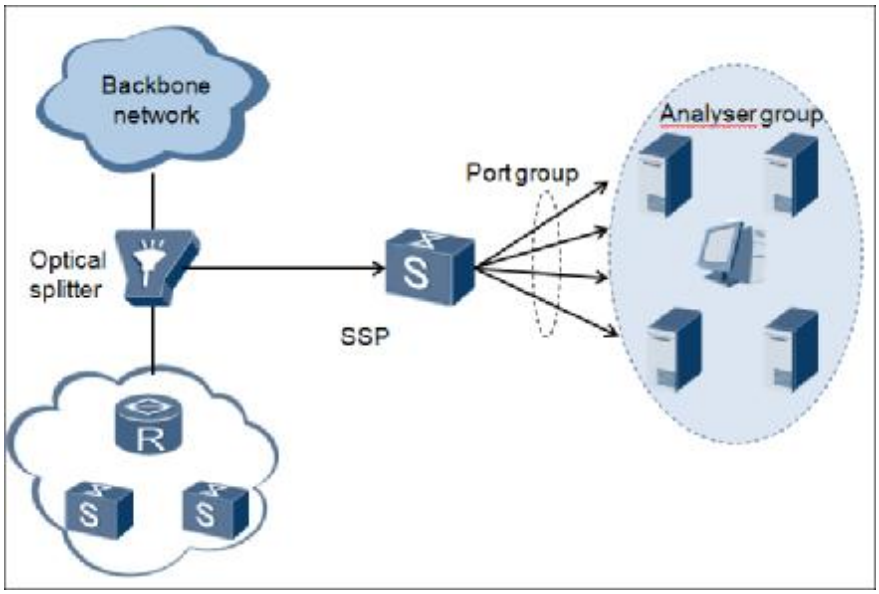

When the ingress interface of the SSP receives packets (including IPv4/IPv6 /VLAN/QINQ/PPPOE/MPLS/GRE/GTP packets) from the optical splitter, the SSP classifies packets based on ACL rules or other conditions, and then:

- Redirects the packets matching the rules to the specified port group
- Discards the packets not matching rules or forwards these packets to public interfaces.

In this way, traffic is imported, distributed, and load balanced. The analysers can thus analyze and monitor the traffic.

2 SSP Working Mechanism

2.1 Concepts

- Unidirectional single-fiber communication

The input and output interfaces of a fiber between the service splitting device and analyser only transmit or receive packets. This function allows the analysers to only receive packets and not to send packets, which ensures data security on the analysers and the packet copies will not affect service traffic.

- Flow pair distribution to the same analyser

In order to fully analyze traffic information, an analyser needs to analyze not only the unidirectional traffic of two communication parties, but also the bidirectional traffic between the two parties. When analyzing the bidirectional traffic, the service splitting device must be able to distribute all the packets between the two communication parties to the same analyser. Therefore, the flow pair sent over the same connection must be sent out through the same outbound interface of the service splitting device.

- Integrity of data packets

The service splitting device is not allowed to modify packets, for example, change the source MAC, destination MAC, or TTL, when forwarding them to the analysers. This ensures the integrity of the data packets and the accuracy of the traffic analysis.

- Packet capturing

In some situations, users only want to obtain a certain field in each packet, for example, the protocol header field, but do not want to obtain the payload. This function allows you to specify the certain fields in packets that the service splitting device will obtain, reducing loads on analysers and saving bandwidth. If a packet is shorter than the specified length, the whole packet is sent to the analyser.

- Removing packet header

Tunnel packets include GRE, GTP, MPLS, IPSec, IPv4-in-IPv4, IPv6 over IPv4, IPv6-in-IPv6, and IPv4 over IPv6 packets. Generally, the analyser only processes the payload data within the tunnel header. Therefore, the service

splitting device needs to remove packet headers to improve packet analysis efficiency.

- Adding the specified information to packets

The analyser may need to verify the inbound interfaces of packets and rules the packets match and know the service splitting device and port group through which the packets are forwarded. The service splitting device replaces the original source MAC addresses with inbound interface information (source card slot ID and source port number), and replaces the destination MAC address with the service splitting device sequence number, outbound interface group number, and rule information (ACL number and rule number). Both the service splitting device and analyser must know the carried information.

2.2 SSP Forwarding

Figure 2-1 SSP forwarding process

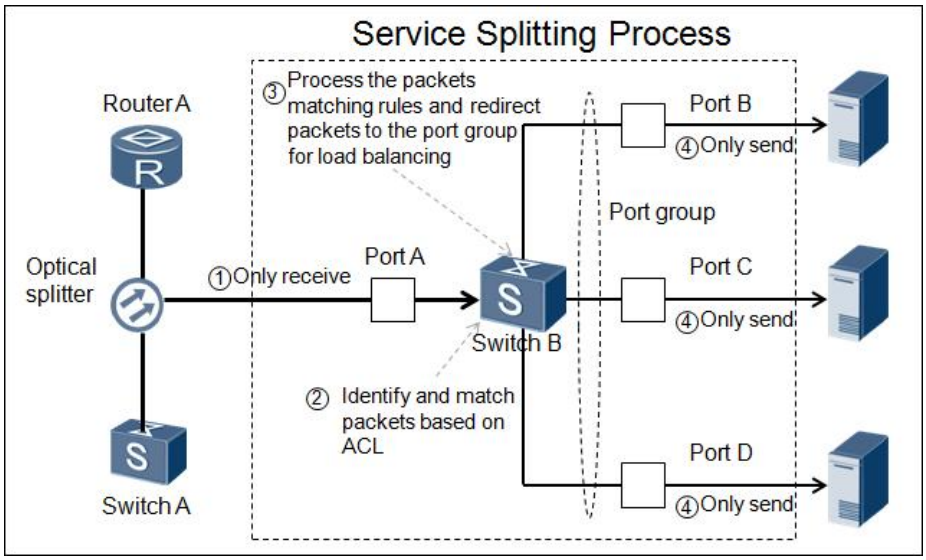

As shown in Figure 2-1, an optical splitter is located between Router A and Switch A to send a copy of traffic to the service splitting device, Switch B.

When the ingress interface of the service splitting device receives packets from the optical splitter, the service splitting device matches packets against the ACL rules or other conditions. If the packets match the rules, the service splitting device takes the specified action, for example, removes packet headers, extracts certain information from the packets, and inserts the specified information to the packets. If the packets do not match the rules, the service splitting device redirects the packets to the specified port group, and the port group load balances the packets to different member ports based on hash algorithm. Then the packets are forwarded to analysers through the member ports. The flow pair distribution to the same analyser function needs

to be configured on the service splitting device so that bidirectional traffic between communicating parties will be sent to the same analyser. In the preceding process, the service splitting device requires the following functions:

- Unidirectional single-fiber communication
- Flow pair distribution to the same analyser
- ACL for classifying and matching packets
- Data forwarding
- Load balancing in port group

The first two functions have been described. The following sections describe the later three functions.

2.2.2 ACL for Classifying and Matching Packets

ACL rules can identify VLAN, MPLS, GRE, GTP, IPv4/IPv6, TCP/UDP, and IP layer tunnel packets, and match character fields in packets. The character fields include five commonly used IP attributes (with or without masks), TCP flag, Layer 2 protocol field, and character code rules with a specified range. These fields are used to match service flows flexibly.

In a character code rule, the user can specify the matching range in a packet (an offset behind the TCP/UDP header) and specify a length of packet character field (for example, an offset behind the payload header). The matching range is not shorter than 128 bytes. The character code length is 2-64 bytes. The user can also specify a certain length of the packet character field within the window size or entire packet. In the entire packet, the character code length ranges from 3 to 13 bytes. The character code rule poses high requirement on ACL matching capability on the service splitting device.

2.2.3 Data Forwarding

When packets match an ACL rule, the service splitting device discards or forwards the packets according to the action setting. The service splitting device uses either of the following methods to process packets:

1. Forward the packets to the port group without modifying the packets.
2. Capture packets, remove the headers, or insert the specified information to the packets, and then forward the packets to the port group.

The service splitting device is not allowed to modify the original packets, for example, change the source MAC, destination MAC, or TTL, when forwarding them to the analysers. This ensures the integrity of the data packets and the accuracy of the traffic analysis.

2.2.4 Load Balancing in Port Group

Generally, the service splitting device forwards traffic according to the traffic characteristics through the port group where Ethernet aggregation or ECMP is configured. The outbound interface is calculated through the hash value calculated based on source/destination MAC addresses, and five IP attributes (protocol type, source IP address, source port, destination IP address, and

destination port). The hash algorithm including five IP attributes load balances traffic based on the flow pair distribution to the same analyser function. For a fragmented packet, the initial and non-initial fragments must be sent to the same port.

When the flow pair distribution to the same analyser function is configured, the five attributes in the flows between two communicating parties are not changed, except that the source/destination IP addresses and source/destination port numbers are swapped. This change does not affect hash calculation result. Therefore, the flows in a session between two parties can be sent to the same analyser.

When distributing the packets encapsulated into Layer 3 tunnel, such as GRE and GTP, the service splitting device performs a hash algorithm based on five internal and external attributes. If the internal IP address changes frequently, the service splitting device performs hash algorithm using the internal IP address. This facilitates load balancing in the port group. For GRE packets, the service splitting device can perform hash algorithm based on the key field.

3 SSP Implementation on Huawei Switches

3.1 Splitting Modes

The modular switches S12700, S9700, and S7700 can function as the service splitting device. They support three SSP modes: Eth-Trunk, ECMP, and multi-trunk. SSP based on Eth-Trunk and multi-trunk uses Layer 2 forwarding, and the SSP based on ECMP uses Layer 3 forwarding. The switches used as service splitting devices must support unidirectional single-fiber communication, flow pair distribution to the same analyser, and data integrity. The SSP function is available only when the license is loaded.

Users can choose a traffic distribution mode based on the actual network structure. The process of SSP based on Eth-Trunk is similar to the process of SSP based on multi-trunk. In the later mode, the port group can be bound to multiple Eth-Trunks, so the maximum number of connected analysers is greatly increased. The ECMP-based SSP distributes IPv4 packets. If the connected analysers are sufficient, the SSP based on Eth-Trunk is recommended.

3.1.1 SSP Based on Eth-Trunk

Figure 3-1 SSP based on Eth-Trunk

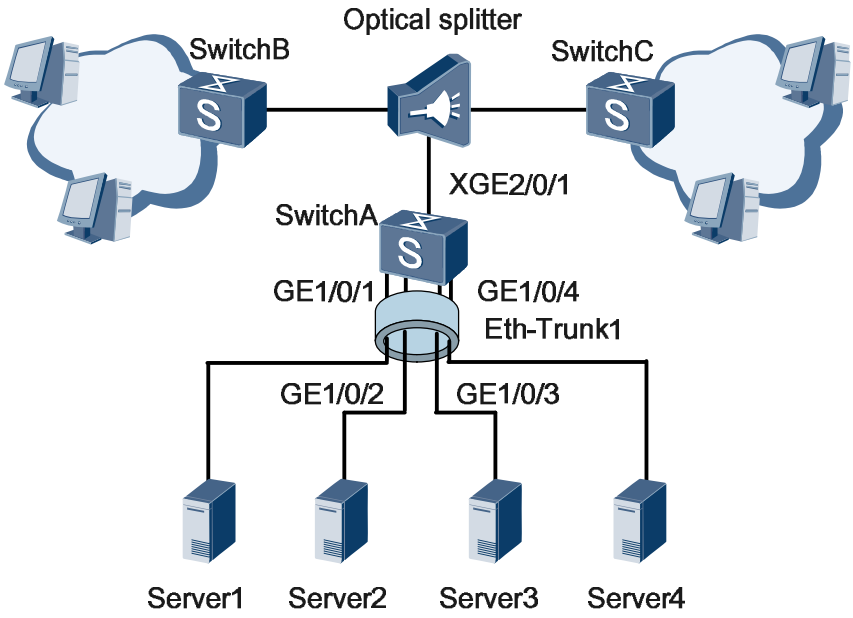

In Figure 3-1, an Eth-Trunk is used as the port group for complex traffic classification (MQC). MQC has three elements: classifier, behavior, and policy. The SSP policy configured based on MQC classifies packets based on protocol types, IP addresses, port numbers, and fragments. The traffic behavior redirects traffic to the specified Eth-Trunk. The traffic policy is applied to the ingress interface to redirect traffic to member interfaces of the Eth-Trunk.

3.1.2 SSP Based on ECMP

Figure 3-2 SSP based on ECMP

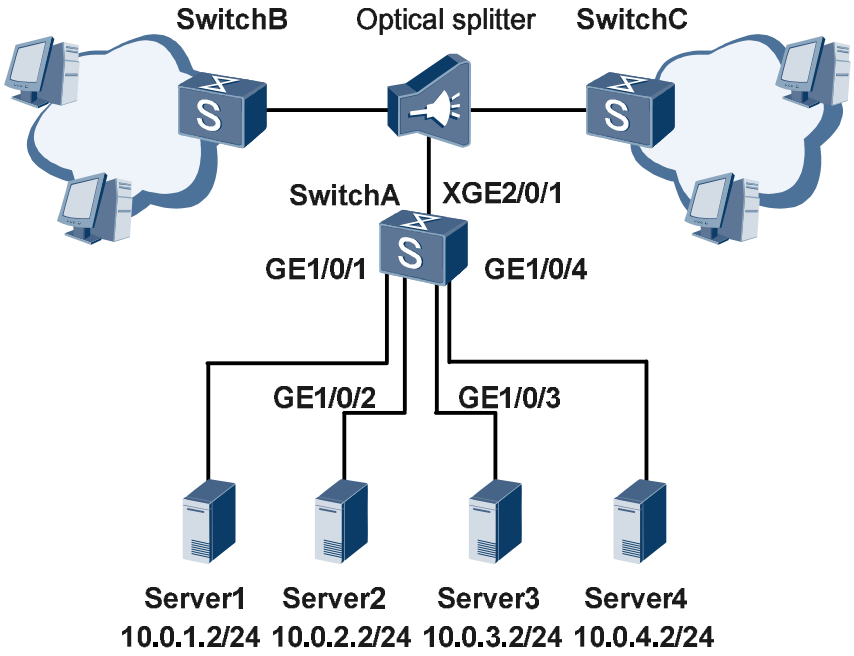

In Figure 3-2, traffic is load balanced based on equal-cost multi-path routing (ECMP). This mode uses Layer 3 forwarding; therefore, the ingress and egress interfaces are Layer 3 interfaces. Multiple static routes destined for 0.0.0.0 are configured, and the next hop addresses are the analysers' IP addresses. Static ARP is configured for next hop addresses. In this way, Layer 3 traffic is load balanced to analysers based on ECMP.

Data integrity must be ensured during Layer 3 forwarding. In a normal Layer 3 forwarding, packet attributes may be modified in any of the following situations:

- Before forwarding a packet, the device compares the destination MAC address of the packet against the device's system MAC address. If they are the same, the device forwards the packet. Otherwise, the device discards the packet. Obviously, the destination MAC addresses of service packets are different from the system MAC address of the service splitting device, so the device changes the destination MAC addresses of the packets, avoiding packet discarding.
- During Layer 3 forwarding, the device replaces the source MAC address of an IP packet with its own system MAC address and replaces the destination MAC address of the packet with the system MAC address of the next-hop device.
- After Layer 3 forwarding is complete, the device should subtract at least one from the TTL value of the packet. The TTL value is carried in IP

protocol packets. Network devices determine whether to discard data packets according to the TTL value.

The service splitting device is not allowed to perform these modifications when forwarding packets to the analysers. This ensures the integrity of the data packets and the accuracy of the traffic analysis. The switch distributes traffic based on ECMP. To avoid modification on packets, the SSP function prevents a device from checking MAC addresses of packets, replacing the source or destination MAC addresses of packets, or changing the TTL values of packets.

3.1.3 SSP Based on Multi-Trunk

Figure 3-3 SSP based on multi-Trunk

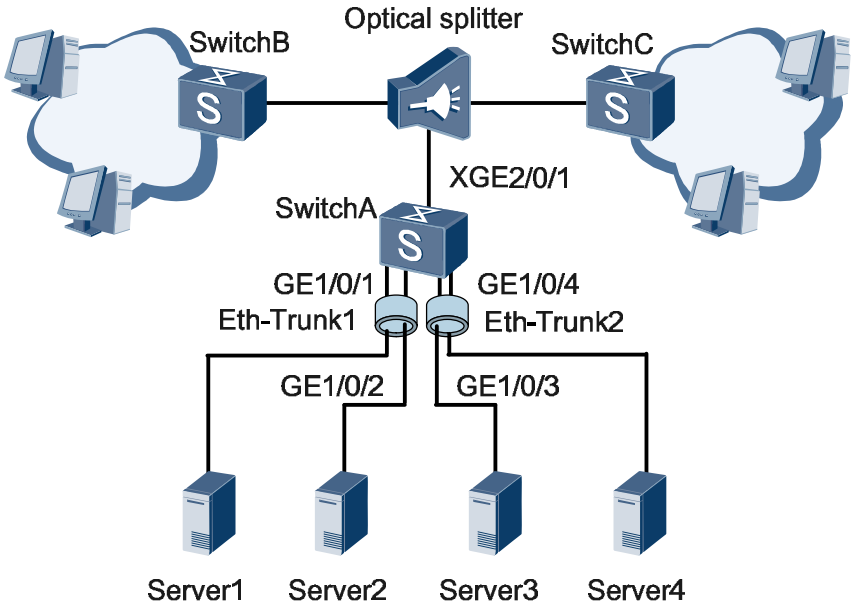

Figure 3-3 shows the SSP process based on multi-trunk. This mode also uses MQC to redirect traffic to the port group. The difference between the two modes is as follows:

- SSP based on Eth-Trunk: load balances traffic to member interfaces in the same Eth-Trunk.
- SSP based on multi-trunk: load balances traffic to member interfaces in different Eth-Trunks. That is, hash calculation is performed using member interfaces of multiple Eth-Trunks.

If one Eth-Trunk contains a maximum of 8 member interfaces, four Eth-Trunks can be specified as the redirect-to ports. Therefore, in SSP based on multi-trunk, a service splitting device can connect to 32 analysers. On some LPUs, a single Eth-Trunk supports a maximum of 32 member interfaces. Therefore, SSP based on Eth-Trunk is recommended.

4 Typical Network

4.1 SSP Based on Eth-Trunk

4.1.1 Network Requirements

As shown in Figure 4-1, Switch A is connected to the upstream optical splitter through XGE2/0/1 to receive traffic. Unidirectional single-fiber communication is configured between Switch A and optical splitter, and between Switch A and analysers. The analysers only receive but do not send data to ensure data security. In addition, Switch A does not send data to the optical splitter, avoiding affect on network services.

It is required that traffic sent from the service splitting device is distributed to four analysers and only the UDP packets with the destination port number being 10000 are matched. The flows in a session are sent to the same analyser.

Figure 4-1 SSP based on Eth-Trunk

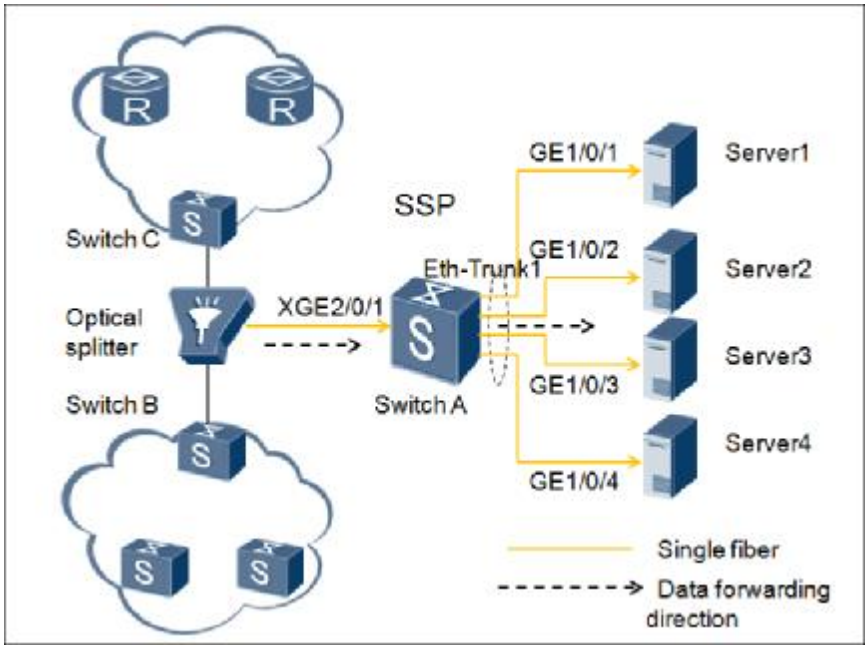

4.1.2 Procedure

Step 1 Configure an Eth-Trunk, add interfaces to the Eth-Trunk, and configure the flow pair to the same analyser function on the Eth-Trunk.

Create Eth-Trunk 1, and add GE 1/0/1, GE 1/0/2, GE 1/0/3, and GE 1/0/4 to Eth-Trunk 1.

```
[SwitchA] interface eth-trunk 1
[SwitchA-Eth-Trunk1] trunkport gigabitethernet 1/0/1 to 1/0/4
```

Configure the flow pair to the same analyser function on Eth-Trunk 1.

```
[SwitchA-Eth-Trunk1] load-balance diffluence
[SwitchA-Eth-Trunk1] quit
```

Step 2 Configure a traffic policy to redirect the UDP packets with the destination port number being 10000 to Eth-Trunk 1.

Create ACL 3000 to match the UDP packet whose destination port number is 10000.

```
[SwitchA] acl 3000
[SwitchA-acl-adv-3000] rule 5 permit udp destination-port eq 10000
[SwitchA-acl-adv-3000] quit
```

Define the traffic classifier **test** and configure ACL 3000.

```
[SwitchA] traffic classifier test
[SwitchA-classifier-test] if-match acl 3000
[SwitchA-classifier-test] quit
```

Define the traffic behavior **test** and set the action to redirection.

```
[SwitchA] traffic behavior test
[SwitchA-behavior-test] redirect interface eth-trunk 1
[SwitchA-behavior-test] quit
```

Define the traffic policy **test** and bind the traffic policy to the traffic classifier and behavior.

```
[SwitchA] traffic policy test
[SwitchA-trafficpolicy-test] classifier test behavior test
[SwitchA-trafficpolicy-test] quit
```

Apply the traffic policy to the inbound interface.

```
[SwitchA] interface xgigabitethernet 2/0/1
[SwitchA-XGigabitEthernet2/0/1] traffic-policy test inbound
[SwitchA-XGigabitEthernet2/0/1] quit
```

Step 3 Configure unidirectional single-fiber communication.

Configure unidirectional single-fiber communication on the interfaces connected to the optical splitter and analysers.

```
[SwitchA]int XGigabitEthernet 2/0/1
[SwitchA -XGigabitEthernet2/0/1]single-fiber  enable
[SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] undo negotiation auto
[SwitchA-GigabitEthernet1/0/1] single-fiber enable
[SwitchA-GigabitEthernet1/0/1] quit
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] undo negotiation auto
[SwitchA-GigabitEthernet1/0/2] single-fiber enable
[SwitchA-GigabitEthernet1/0/2] quit
[SwitchA] interface gigabitethernet 1/0/3
[SwitchA-GigabitEthernet1/0/3] undo negotiation auto
[SwitchA-GigabitEthernet1/0/3] single-fiber enable
[SwitchA-GigabitEthernet1/0/3] quit
[SwitchA] interface gigabitethernet 1/0/4
[SwitchA-GigabitEthernet1/0/4] undo negotiation auto
[SwitchA-GigabitEthernet1/0/4] single-fiber enable
[SwitchA-GigabitEthernet1/0/4] quit
```

Step 4 Configure load balancing based on source address, destination address, source port, destination port, and protocol type.

```
[SwitchA]load-balance-profile huawei
[SwitchA  -load-balance-profile-huawei]ipv4  field  sip  dip  l4-sport
l4-dport protocol
```

Step 5 Run the **display out-interface** command to view the outbound interfaces of the packets carrying the five attributes and check whether the flows of the same session are sent to the same analyser. Then you can find the analyser to which packets are distributed. Run the **display interface brief | include up** command to check whether load balancing takes effect.

4.2 SSP Based on Multi-Trunk

4.2.1 Network Requirements

As shown in Figure 4-2, Switch A is connected to the upstream optical splitter through XGE2/0/1 to receive traffic. Unidirectional single-fiber communication is configured between Switch A and optical splitter, and between Switch A and analysers. The analysers only receive but do not send data to ensure data security. In addition, Switch A does not send data to the optical splitter, avoiding affect on network services.

It is required that traffic sent from the service splitting device is distributed to four analysers and only the UDP packets with the destination port number being 10000 are matched. The flows in a session are sent to the same analyser.

Figure 4-2 SSP based on multi-trunk

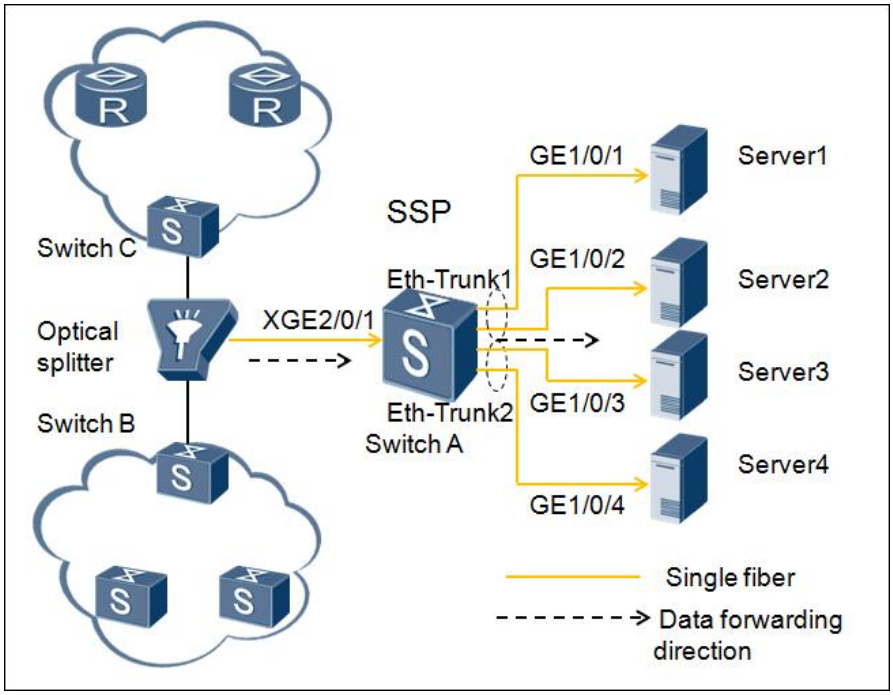

4.2.2 Procedure

Step 1 Configure an Eth-Trunk, add interfaces to the Eth-Trunk, and configure the flow pair to the same analyser function on the Eth-Trunk.

Create Eth-Trunk 1, and add GE 1/0/1 and GE 1/0/2 to Eth-Trunk 1.

```
[SwitchA] interface eth-trunk 1
[SwitchA-Eth-Trunk1] trunkport gigabitethernet 1/0/1 to 1/0/2
```

Create Eth-Trunk 2, and add GE 1/0/3 and GE 1/0/4 to Eth-Trunk 2.

```
[SwitchA] interface eth-trunk 2
[SwitchA-Eth-Trunk1] trunkport gigabitethernet 1/0/3 to 1/0/4
```

Configure load balancing based on service splitting.

```
[SwitchA] ecmp load-balance diffluence
```

Step 2 Configure a traffic policy to redirect the UDP packets with the destination port number being 10000 to Eth-Trunk 1 and Eth-Trunk 2.

Create ACL 3000 to match the UDP packet whose destination port number is 10000.

```
[SwitchA] acl 3000
[SwitchA-acl-adv-3000] rule 5 permit udp destination-port eq 10000
[SwitchA-acl-adv-3000] quit
```

Define the traffic classifier **test** and configure ACL 3000.

```
[SwitchA] traffic classifier test
[SwitchA-classifier-test] if-match acl 3000
[SwitchA-classifier-test] quit
```

Define the traffic behavior **test** and set the action to redirection.

```
[SwitchA] traffic behavior test
[SwitchA-behavior-test] redirect multi-trunk eth-trunk 1 eth-trunk 2
[SwitchA-behavior-test] quit
```

Define the traffic policy **test** and bind the traffic policy to the traffic classifier and behavior.

```
[SwitchA] traffic policy test
[SwitchA-trafficpolicy-test] classifier test behavior test
[SwitchA-trafficpolicy-test] quit
```

Apply the traffic policy to the inbound interface.

```
[SwitchA] interface xgigabitethernet 2/0/1
[SwitchA-XGigabitEthernet2/0/1] traffic-policy test inbound
[SwitchA-XGigabitEthernet2/0/1] quit
```

Step 3 Configure unidirectional single-fiber communication.

Configure unidirectional single-fiber communication on the interfaces connected to the optical splitter and analysers.

```
[SwitchA]int XGigabitEthernet 2/0/1
[SwitchA -XGigabitEthernet2/0/1]single-fiber  enable
[SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] undo negotiation auto
[SwitchA-GigabitEthernet1/0/1] single-fiber enable
[SwitchA-GigabitEthernet1/0/1] quit
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] undo negotiation auto
[SwitchA-GigabitEthernet1/0/2] single-fiber enable
[SwitchA-GigabitEthernet1/0/2] quit
[SwitchA] interface gigabitethernet 1/0/3
[SwitchA-GigabitEthernet1/0/3] undo negotiation auto
[SwitchA-GigabitEthernet1/0/3] single-fiber enable
[SwitchA-GigabitEthernet1/0/3] quit
[SwitchA] interface gigabitethernet 1/0/4
[SwitchA-GigabitEthernet1/0/4] undo negotiation auto
[SwitchA-GigabitEthernet1/0/4] single-fiber enable
[SwitchA-GigabitEthernet1/0/4] quit
```