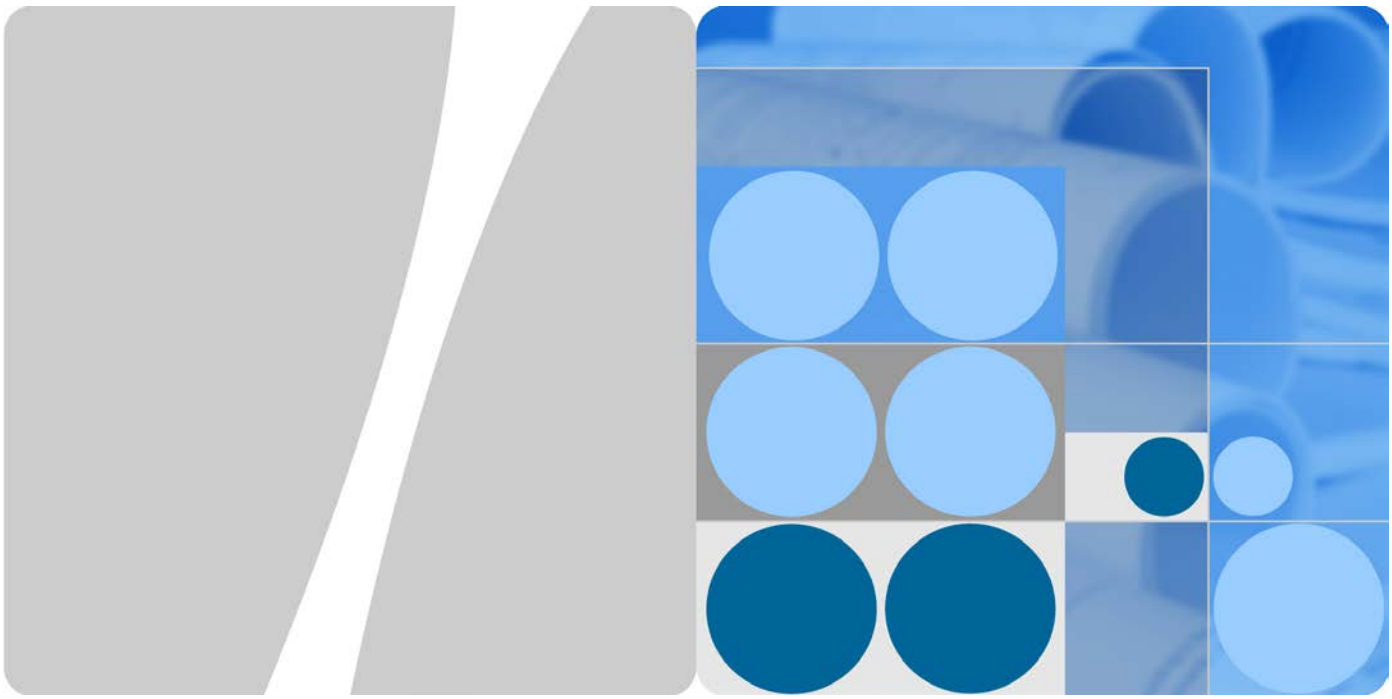

S7700&S9700 V200R012C00 Feature Description - Free Mobility and Service Chain

Issue 01

Date 2018-07-10

Contents

1 Free Mobility

About This Chapter

After the free mobility function is configured, a user can obtain the same network access policy regardless of the user's location and IP address changes.

1.1 Free Mobility Overview

Definition

Free mobility is a solution that allows a user to obtain the same network access policy regardless of the user's location and IP address changes on an agile network.

Background

On an enterprise network, different network access policies can be deployed for users on access devices to meet different network access requirements. On traditional campus networks, users' network access rights are controlled using the NAC technology with VLAN and ACL technologies. Requirements of these technologies are as follows:

1. Employees must connect to the campus network through specified switches, VLANs, or network segments, so that they have the same network access rights.
2. ACLs for controlling users' network access rights need to be preconfigured. In the ACLs, at least the destination IP addresses that are prohibited or allowed to access are configured. Therefore, if the IP address of a user is not fixed and the user's host is both a source and a destination, an ACL is not applicable.
3. The association between ACLs and users only takes effect on the authentication device. Therefore, for a non-authentication device, such as the firewall deployed at the boundary of an enterprise campus network, IP address-based policies must be configured.
4. VLANs and ACLs need to be preconfigured on a large number of authentication switches, bringing huge workload for deployment and maintenance.

Mobile office requires that these limitations be removed and employees access the network from any location, any VLAN, or any IP network segment with controlled network access rights. Therefore, free mobility is introduced. By using the Agile Controller and agile switches, network access rights can automatically migrate when user locations change, improving mobile office experience.

The free mobility solution solves problems faced by traditional campus networks from the following perspective:

1. Decoupling of service policies and IP addresses

   Using the Agile Controller, the administrator can divide users and resources on the entire network into different security groups based on different dimensions. In addition, agile devices in the free mobility solution use an innovative software and hardware design. An agile device can match the source and destination IP addresses of packets with source and destination security groups, and then finds the matching inter-group policy based on the source and destination groups.

   Through the innovative design, all the user- and IP address-based service policies used on traditional networks can be migrated to security group-based policies. When predefining service policies, the administrator does not need to consider users' actual IP addresses, decoupling service policies from IP addresses.

2. Centralized management of user information

   The Agile Controller centrally manages authentication and online information about users and obtains mappings between network-wide users and IP addresses. Non-authentication devices on the network can actively obtain information about source and destination security groups from the Agile Controller based on the source and destination IP addresses of packets.

3. Centralized management of policies

   The Agile Controller is not only the authentication center on campus networks, but also the management center of service policies. The administrator can use the Agile Controller to centrally manage service policies on network-wide policy enforcement devices. After being configured for one time, these service policies can be automatically delivered to enforcement devices on the entire network. These policies include rights policies (for example, group A is forbidden to access group B) and experience guarantee polices (for example, traffic forwarding bandwidth and priority of group A are controlled).

Benefits

1. Simplified network planning: The administrator does not need to consider IP addresses of users when configuring policies.
2. Enhanced control capability: User authentication information can be synchronized between network devices.
3. Improved management efficiency: The administrator does not need to configure devices one by one.

1.2 Understanding Free Mobility

1.2.1 Network Architecture

Figure 1-1 shows the overall architecture of free mobility.

Figure 1-1 Network architecture of free mobility

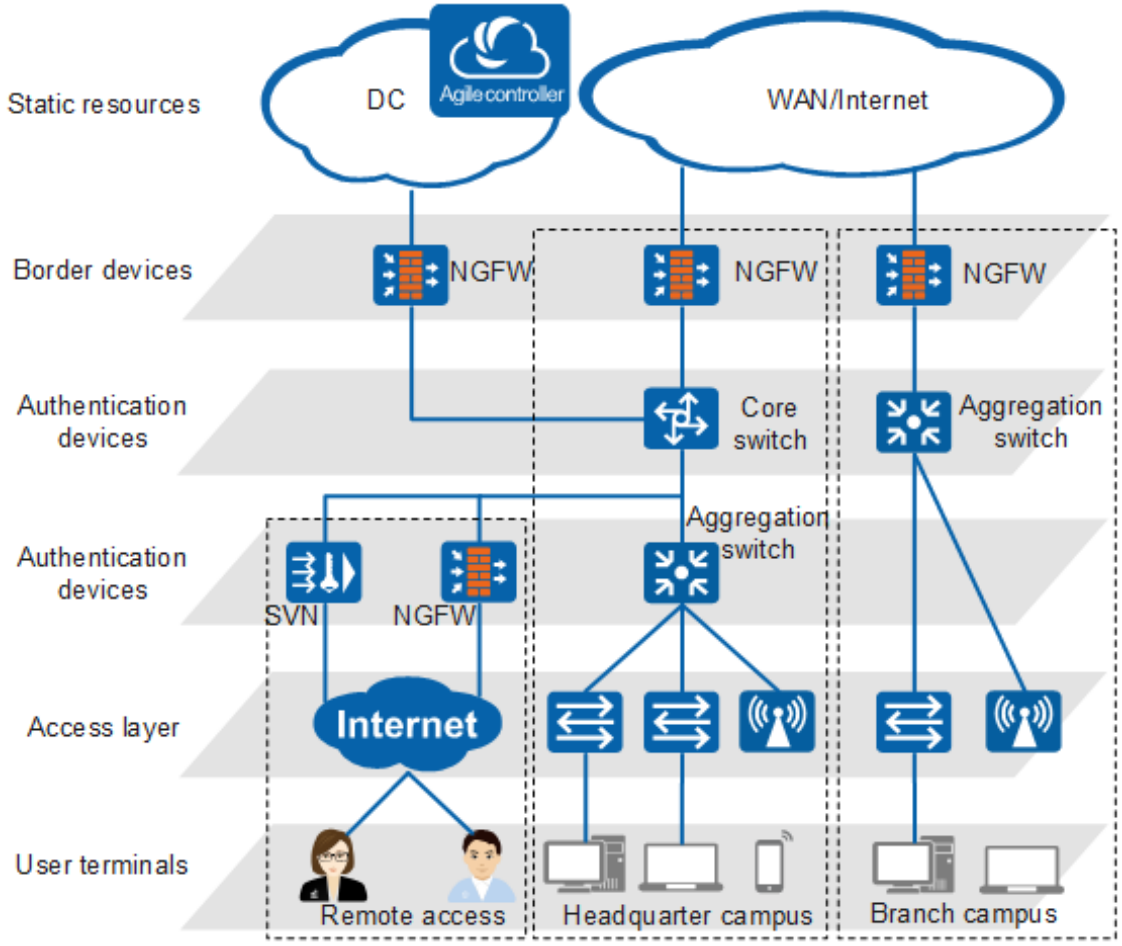

- User terminals initiate authentication.
- The access layer transparently forwards user traffic at Layer 2.
- Authentication devices authenticate users and control their access rights.
- Border devices ensure the forwarding priority of specific users at the egress.
- Static resources are server resources that users can access. These resources can be managed as security groups on the Agile Controller.

The free mobility solution involves the following key roles:

- Controllers: work with network devices to authenticate users and deliver policies, decoupling service policies from IP addresses. Only Huawei Agile Controller-Campus is supported in the free mobility solution.
- Access devices: authenticate terminals, determine whether to allow the terminals to access the network, and control the terminals' network access rights. The access devices include switches and firewalls. Only switches are mentioned in this manual.
- Terminals: provide human-machine interfaces for user authentication and resource access. The terminals include PCs, laptops, smartphones, tablets, and dumb terminals.

1.2.2 Implementation

Figure 1-2 shows the implementation of free mobility.

Figure 1-2 Implementation of free mobility

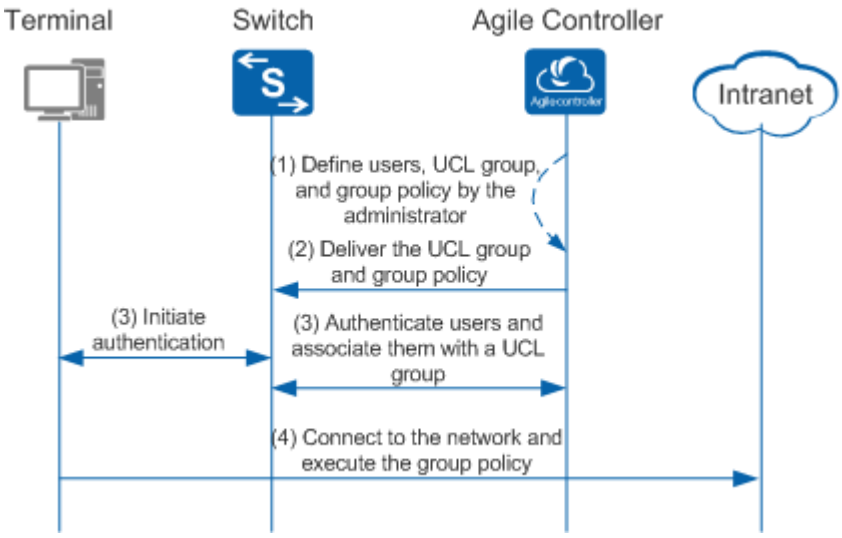

1. An administrator creates a user account and User Control List (UCL) group, adds the user account to the UCL group, and defines the network access policy (that is, group policy) for the user based on the UCL group. All users can access the network only after being authenticated.
2. The Agile Controller delivers the UCL group and network access policy configured by the administrator to all associated switches, so that the switches can identify the UCL group to which the user belongs. The administrator can also deploy some service policies based on security groups on these switches.
3. The user starts authentication. During the authentication, the Agile Controller associates the user with the UCL group based on the user login information. After the authentication succeeds, the Agile Controller delivers the group to which the user belongs as the authorization result to the authentication device. The Agile Controller collects IP addresses of all online users.
4. The user accesses the network. After receiving user packets, the switch attempts to identify security groups to which the source and destination IP addresses of the packets correspond, and enforces UCL group-based policies for packets. If the enforcement device is also the authentication device, the device obtains access users' security group information during authentication.

NOTE
A UCL group is called the security group on the Agile Controller-Campus.

UCL groups identify the user types. The administrator can add the users requiring the same network access policy to the same UCL group, and configure a network access policy for the group. Compared with the solution in which network access policies are deployed for each user, the UCL group-based network control solution greatly reduces the administrator's workload.

1.2.3 Security Group

A security group is an abstracted and logical set of communicating objects on the network. Security group members can be network terminals such as PCs and smartphones. They can be statically added by the administrator or dynamically added during authentication.

The administrator can add the users requiring the same network access policy to the same UCL group, and configure a network access policy for the group. Network objects are added to the same security group based on their similarities in network access, and obtain the same rights based on the policy configured for the security group. For example, an R&D group is a set of individual hosts, a printer group is a set of all printers on the entire network, and a database server group is a set of server IP addresses and ports. Compared with the solution in which network access policies are deployed for each user, the security group-based network control solution greatly reduces the administrator's workload.

Division of Security Groups

The administrator can define security groups to describe and organize the sources or destinations of network traffic, such as user hosts, IP phones, servers, and interfaces of network devices that have IP addresses and can send or receive IP packets. To control mutual access between these devices, the administrator needs to define them on the Agile Controller first.

Security groups are classified into the following types:

- Dynamic user group: users and terminals that can access the network after authentication.
- Static resource group: devices hat have fixed IP addresses. These resources include data center servers, interfaces of network devices, special terminals that use fixed IP addresses to access the network without authentication, and all the other network objects using available IP addresses.

When a security group is bound to multiple authorization rules, it is a dynamic user group. When a security group is bound to multiple IP addresses or IP network segments, it is a static resource group. Differences between dynamic users and static resources are as follows:

- The IP address of a dynamic user is not fixed, and it is dynamically associated with a security group after the user is authenticated. After the user logs out, the association is dynamically canceled. The mappings between user IP addresses and security groups are valid only when users are online. A network device can obtain the mappings only when it functions as the authentication device of the users or actively queries the mappings from the Agile Controller.
- The IP address of a static resource is fixed and configured by the administrator. In the pre-deployment phase, when the Agile Controller synchronizes policies to network devices through the Extensible Messaging and Presence Protocol (XMPP), the binding between security groups and IP addresses is synchronized to all the enforcement devices.

If an IP address is added to different groups in both authentication and static modes, the dynamic group authorized in authentication mode is preferred.

Joining a Security Group

Members can join a security group in the following ways:

- User authentication: The administrator configures user authorization rules and associates them with security groups on the Agile Controller. After a user is authenticated, the Agile Controller automatically associates the user's IP address with the security group to which the user belongs.

  An authorization rule is composed of two parts:

  - Authorization condition: Users are described based on their login conditions.

- Authorization result: Users matching the authorization conditions are associated with the security group specified in the authorization result on their logins. That is, the users are assigned certain identities and rights.

The administrator can configure the following user authorization conditions:

- User information
  - User group: such as HR department, marketing department, and logistics department
  - Role: such as R&D personnel, service personnel, sales personnel, and finance personnel
  - Account: user name used to access the network
- Location information
  - Access device group: physical locations of access devices, such as an office and office building
  - Terminal IP address range: IP address used to connect to the network
  - SSID: SSID used to access a WLAN
- Other information
  - Terminal device group: terminal used to connect to the network, for example, PCs with Windows operating systems or smart terminals with Android operating systems
  - Time range: user online time range
  - Customized condition: RADIUS attributes carried in authentication packets, which are used to determine the current login environment

- Static configuration: The administrator associates IP addresses, network segments, or IP addresses + ports with security groups on the Agile Controller. The administrator can add objects that can access the network without authentication such as dumb terminals and servers to specified security groups.
- Third-party system: Through open interfaces of the Agile Controller, a third-party software system can dynamically add members to or remove members from security groups based on its application scenario and algorithm.

A key technology in the free mobility solution is to synchronize the association between users and security groups to other devices. Common group information synchronization methods are as follows:

1. The Agile Controller synchronizes group information to all relevant devices on the network.
2. A dedicated protocol is used to synchronize group information between devices.
3. User traffic carries group information and is sent to other devices.

1.3 Application Scenarios for Free Mobility

1.3.1 Free Mobility Deployment in Campus Access Scenarios

In campus access scenarios, free mobility controls users' access rights based on accounts, terminal types, and access modes to ensure consistent access rights regardless of users' locations.

In Figure 1-3, the aggregation switch is recommended as the authentication control point in a campus access scenario. After you configure security groups and access control policies on the Agile Controller-Campus, the Agile Controller-Campus delivers the configuration to the aggregation switch.

Figure 1-3 Free mobility deployment in campus access scenarios

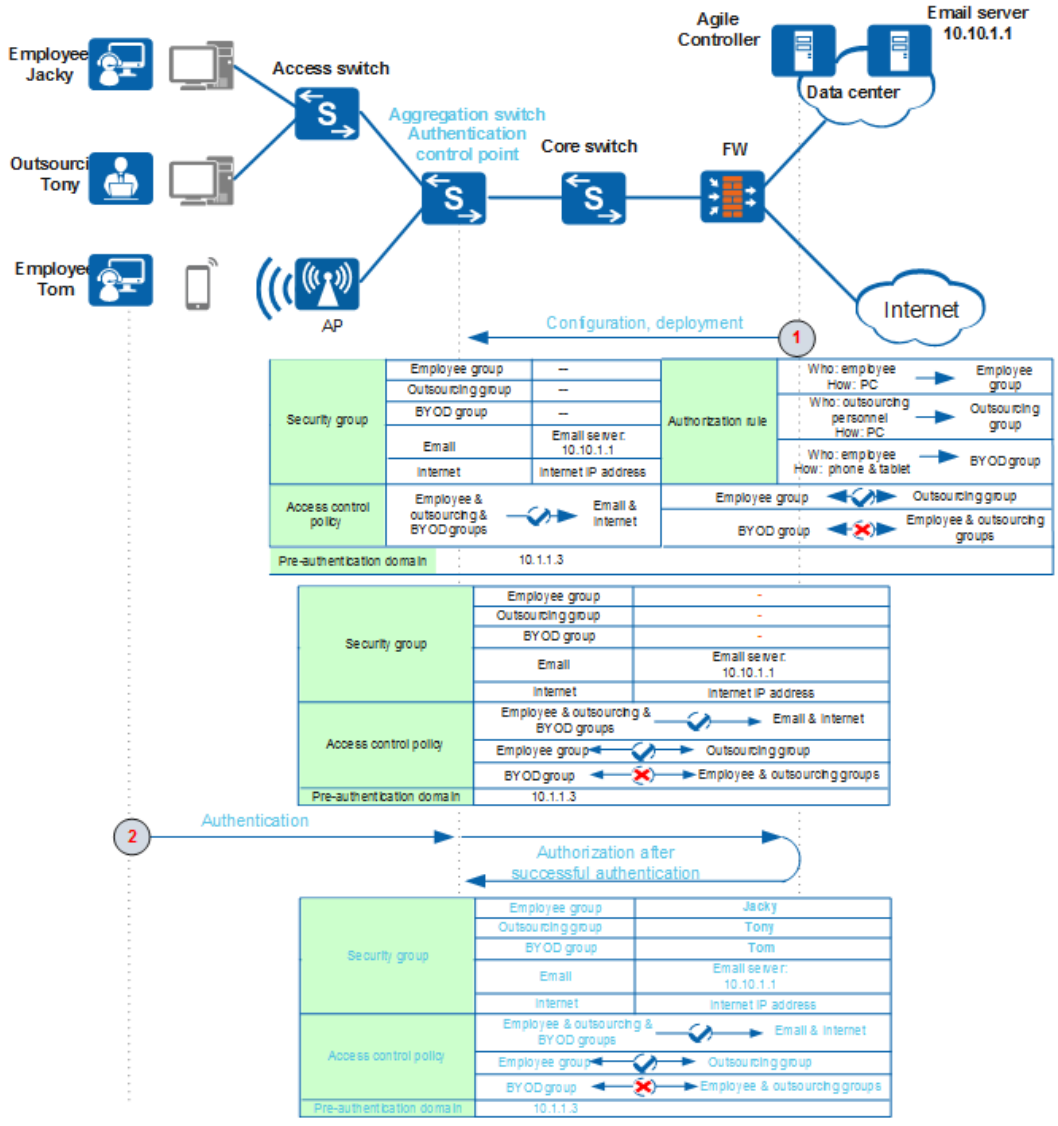

1. You can configure security groups, authorization rules, access control policies, and a pre-authentication domain on the Agile Controller-Campus. The Agile Controller-Campus will deliver the access control policies along with specified security groups and the pre-authentication domain to network-wide devices.
   - Plan the employee group, outsourcing group, and BYOD group for employees logging in through PCs, outsourcing personnel logging in through PCs, and employees logging in through mobile terminals, as well as the email group and Internet group that are bound to the email server IP address and Internet address respectively.

- Define authorization rules to map employees logging in through PCs, outsourcing personnel logging in through PCs, and employees logging in through mobile terminals to the employee group, outsourcing group, and BYOD group.
- Define access control policies to permit or deny access among security groups.
- Specify the pre-authentication domain to permit access to resources in the domain.

2. A user sends an authentication request to the Agile Controller-Campus. After successful authentication, the Agile Controller-Campus adds the user to the specified user group based on an authorization rule and sends the authentication result to the aggregation switch. The aggregation switch permits or denies access from the security group to which the user belongs to target resources based on access control policies.

Access control policies for security groups can control communication between users connected to the same aggregation switch. How can terminals connected to different aggregation switches in Figure 1-4 to communicate with each other?

By default, a switch only saves information of users authenticated on it, and cannot send a user's IP address to the Agile Controller-Campus to query the security group to which the user belongs.

Figure 1-4 Control of communication between user groups on different authentication devices

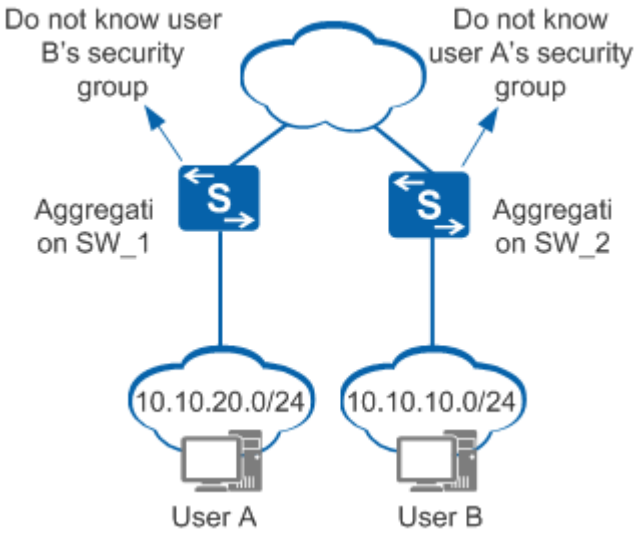

In this case, you are advised to create two security groups and bind them to network segments of **Aggregation SW_1** and **Aggregation SW_2** respectively.

- Security group 1: is bound to 10.10.20.0/24 and contains network resources managed by the aggregation switch SW_1.
- Security group 2: is bound to 10.10.10.0/24 and contains network resources managed by the aggregation switch SW_2.

Then permit access from user A's security group to **Security group 2** and access from user B's security group to **Security group 1** to allow communication between the two users connected to different aggregation switches.

1.3.2 Free Mobility Deployment in VPN Access Scenarios

VPN access applies to large enterprises that have multiple branches. As shown in Figure 1-5, an enterprise has two branches located in Shenzhen and Beijing. Each branch has two VPN networks: the green area and the yellow area. Users in the green area can access the network from fixed and mobile terminals, while users in the yellow area can only access the network from fixed terminals. The aggregation switches (usually PE devices on the MPLS VPN network) function as the authentication control points.

Figure 1-5 Free mobility deployment in VPN access scenarios

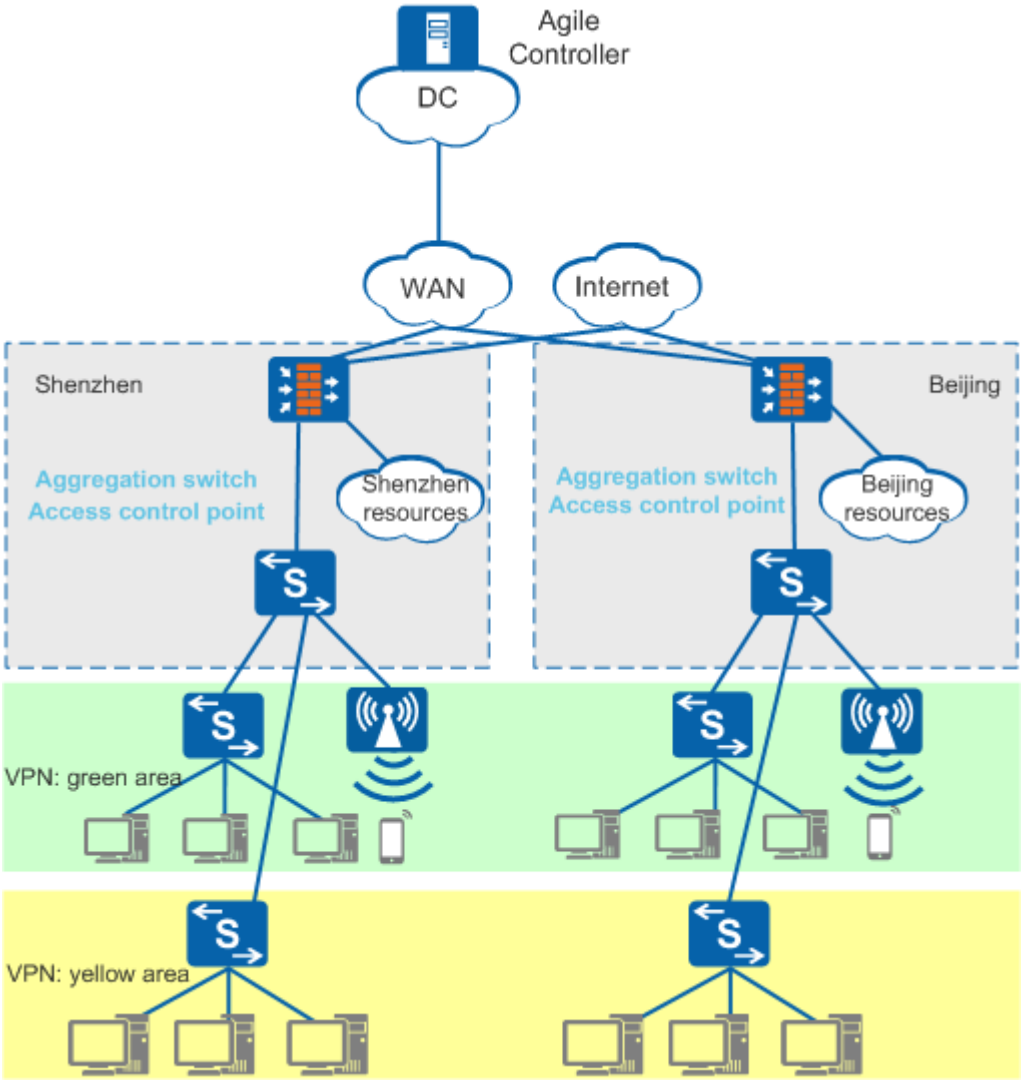

In a VPN access scenario, plan security groups in a unified manner and configure different policies for different VPNs. One user may have different network access rights when the user moves from one VPN to another. Apart from unified policies for the entire network, you can also configure policies on a specific device separately.

The Agile Controller-Campus uses global policies and local policies to deploy universal policies on network-wide devices and special policies on specific devices.

- Global policy: takes effect on network-wide devices or devices in a specified device group. In a VPN access scenario, you can deploy global policies on devices in the green area and yellow area separately.
- Local policy: takes effect on a specific device only. You need to configure local policies when the inherited global policies do not meet service requirements.

2 Service Chain

About This Chapter

This chapter describes how to configure a service chain on a network. This configuration can help redirect different service flows to value-added service devices such as firewalls for processing.

2.1 Principle

Service Chain

On a typical campus network, value-added service devices, such as firewalls, antivirus expert systems, and application security gateways, are often deployed at the edge of an important service department, demilitarized zone (DMZ), campus egress, and data center. Value-added service devices in Figure 2-1 are firewalls. The scheme that deploys an independent value-added service device in each network zone has the following disadvantages:

- Increases investment because too many value-added service devices need to be deployed.
- Wastes resources because value-added service devices are not fully used.
- Complicates device deployment and maintenance because different service processing policies need to be configured on each value-added service device.

Figure 2-1 Typical campus network

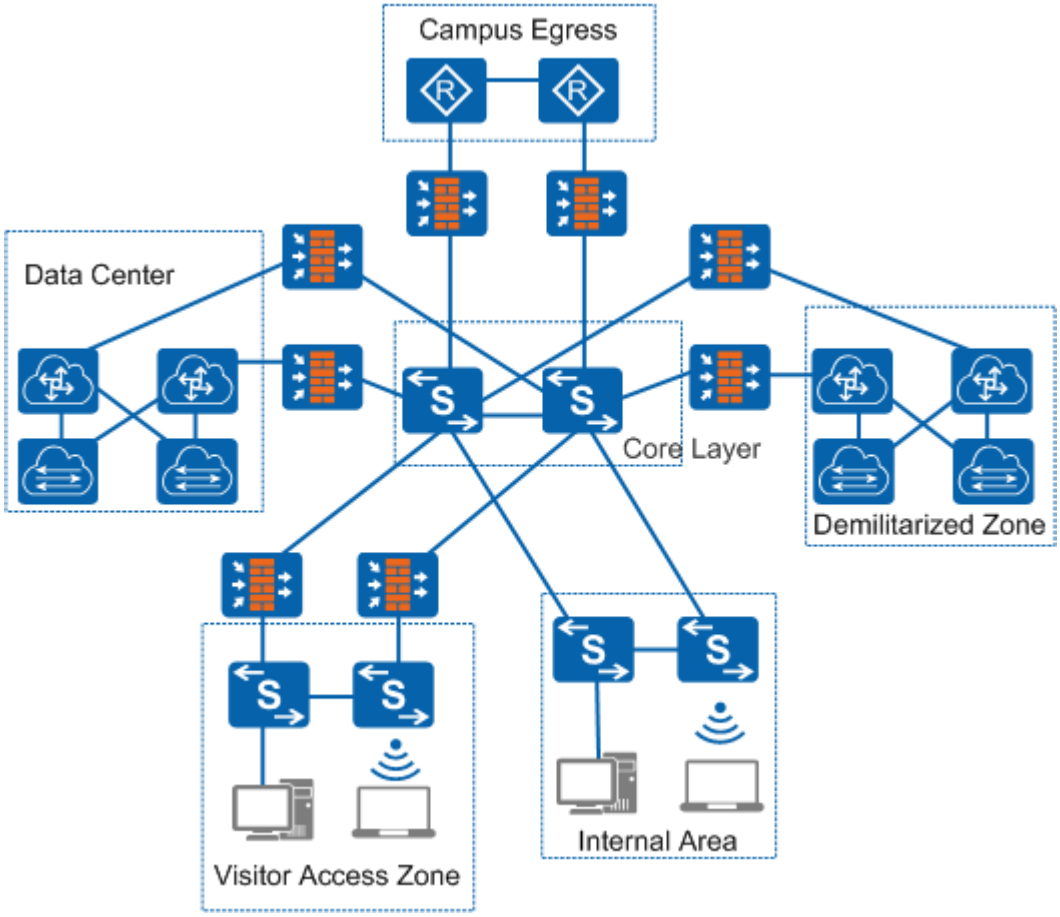

To address the preceding issues, Huawei offers the service chain solution. As shown in Figure 2-2, the service chain solution consists of the policy controller, core switches, and security resource pool. Core switches classify service traffic and then redirect the traffic to different value-added service devices. In the security resource pool, you can deploy one device that has multiple value-added service capabilities or multiple devices that have independent value-added service capabilities. The service chain solution allows value-added service devices to be concentrated in a physical zone. In this solution, you do not need to deploy an independent value-added service device for each network, reducing device costs and improving device utilization. On the campus network, the policy controller controls which service traffic needs to be processed by value-added service devices, improving deployment and maintenance efficiency.

NOTE

Currently, the service chain solution supports three types of value-added service devices: firewall, antivirus expert system, and application security gateway.

Figure 2-2 Campus network in the service chain solution

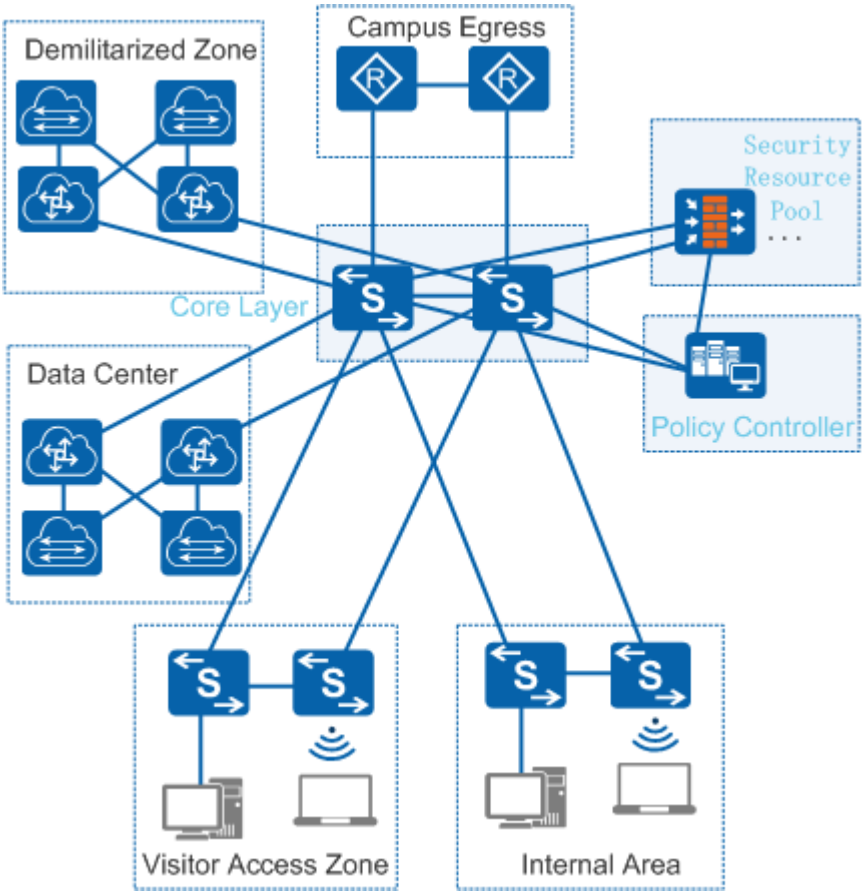

In the service chain solution shown in Figure 2-2, service traffic matching specified rules is forwarded from a core switch to a specified value-added service device through a GRE tunnel, and sent back to the core switch through the GRE tunnel after being processed by the value-added service device. The core switch then forwards the processed service traffic to the destination network. In this manner, different service flows can be imported to different value-added service devices, and a specified service flow can be imported to different value-added service devices in sequence according to policies.

As shown in Figure 2-3, the policy controller delivers instructions to the core switch, firewall, and antivirus expert system so that service traffic between the user and server meets the following requirements:

- Traffic A from the user to the server needs to be processed by the antivirus expert system and then the firewall. Then the processed traffic is forwarded to the server by the core switch.
- Traffic B from the server to the user only needs to be processed by the antivirus expert system and then is forwarded to the user by the core switch.

Figure 2-3 Service chain deployment diagram

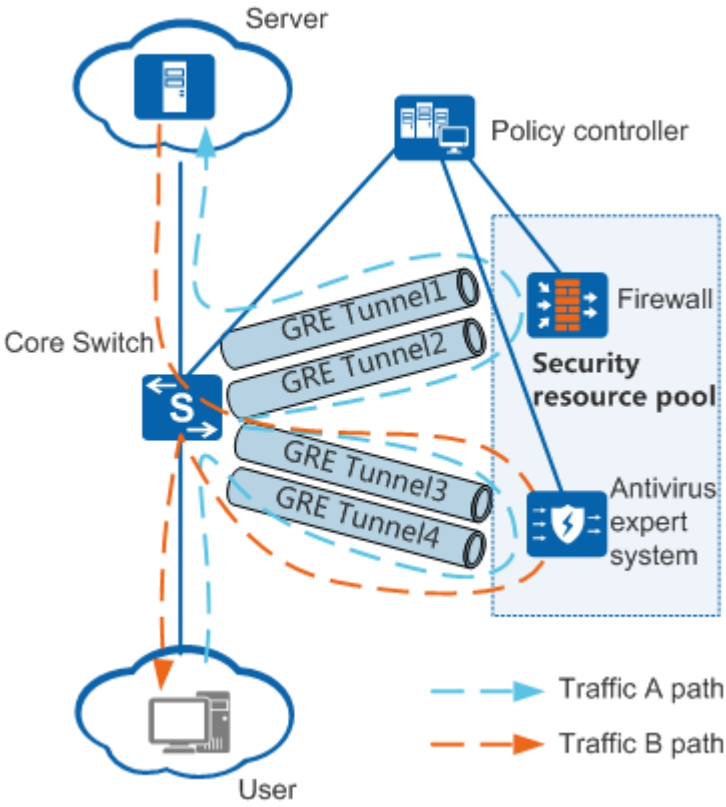

Service Traffic Filtering

In the service chain solution, service flows are filtered using the ACL and UCL rules delivered by the Controller.

- Service flows that match ACL rules can be redirected to a specified GRE tunnel.
- UCL rules allow users in the same user group to have the same filtering policy.

  The Controller delivers the mappings between filtering polices and user groups to switches and value-added service devices. These devices obtain source and destination user group numbers based on source and destination IP addresses of service traffic, and redirect the service traffic to specified GRE tunnels based on the source and destination user group numbers.

Service Chain Implementation

On a campus network, there are three types of service traffic according to traffic directions: service traffic from users to external networks, service traffic from external networks to users, and service traffic between users. Service chain implementation is the same regardless of the traffic direction.

Figure 2-4 describes how a service chain is established for the traffic that is sent from users to external networks and must pass through the application security gateway, antivirus expert system, and firewall.

Figure 2-4 Path for traffic from an user to the external network

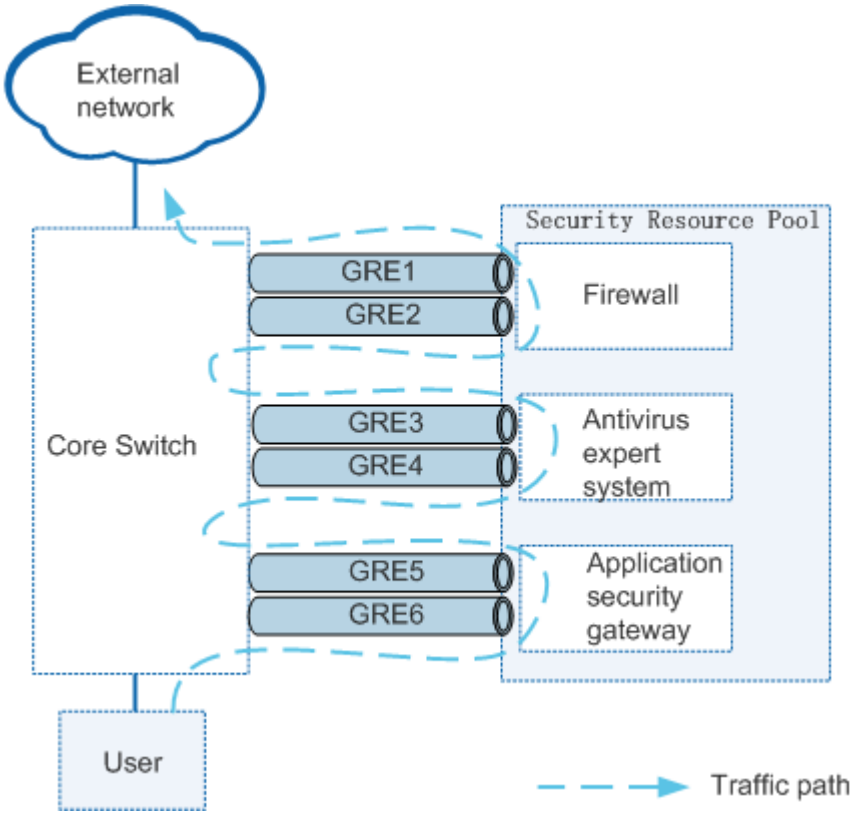

1. After traffic from an user to the external network is forwarded to the core switch, the core switch redirects the traffic to a specified tunnel interface according to matching rules, and then forwards the traffic to the application security gateway through GRE6.
2. After the application security gateway processes the traffic, it redirects the traffic to a specified tunnel interface and returns the traffic to the core switch through GRE5.
3. The core switch redirects traffic received from GRE5 to the tunnel interface where GRE4 resides, and then forwards the traffic to the antivirus expert system. Service traffic is processed by the antivirus expert system and firewall and then returned to the core switch.
4. The core switch receives service traffic through GRE1. Because the traffic does not need to be processed by other value-added service devices, the core switch forwards the traffic to the external network according to the routing table.

Measures Taken by a Service Chain to Respond to Network Faults

When a physical link between value-added service devices or between the core switch and a value-added service device fails, the Controller delivers configurations to cancel the configuration of the GRE tunnel between the two devices and delivers a traffic dropping or forwarding policy.

As shown in Figure 2-5, when the GRE4 tunnel between a core switch and the antivirus expert system fails, the switch sends the failure information to the Controller. The Controller then delivers configurations to the core switch and antivirus expert system to cancel the

configurations of GRE3 and GRE4. Subsequently, traffic received by the core switch from GRE5 is discarded or forwarded along routes based on a specified policy.

Figure 2-5 Direct traffic forwarding in the case of a fault