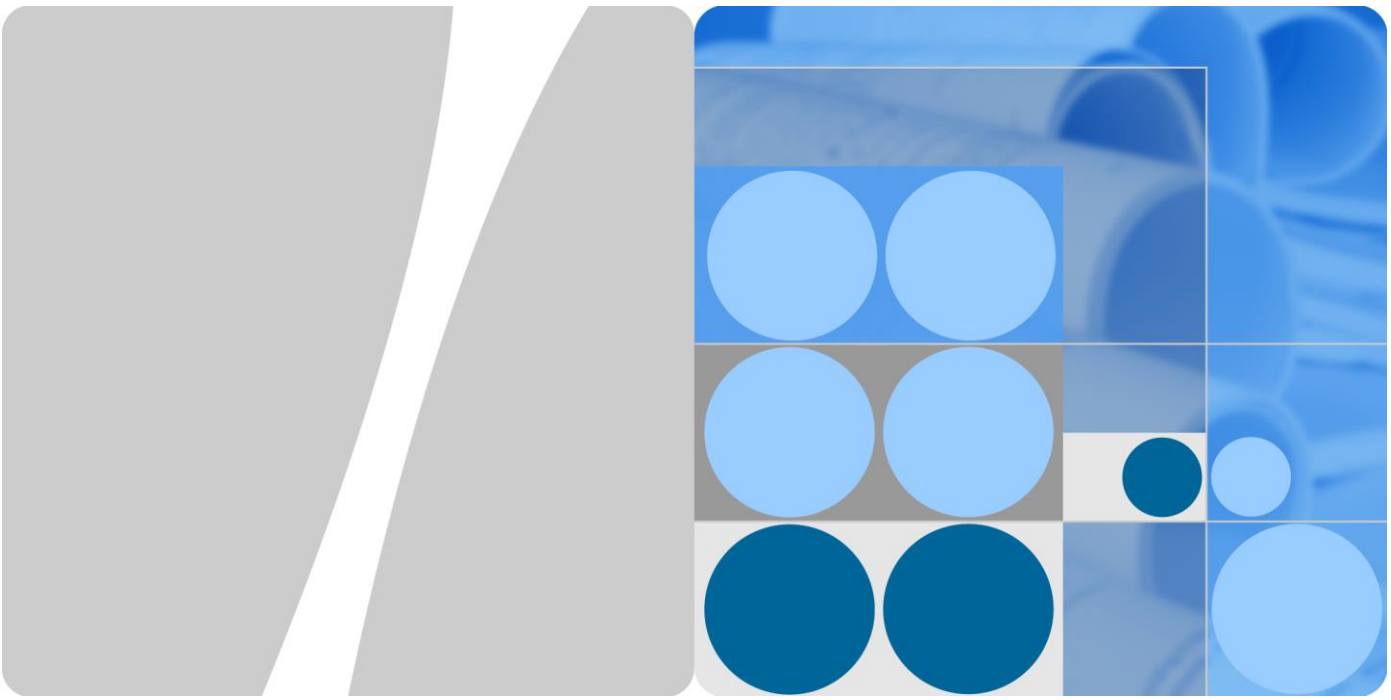

Agile Controller-Campus
V100R002C10

Service Chain Technical White Paper

Issue 01

Date 2016-04-15

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 Technical Background

1.1 New Challenges Brought by Network Changes

Figure 1-1 New challenges brought by network changes

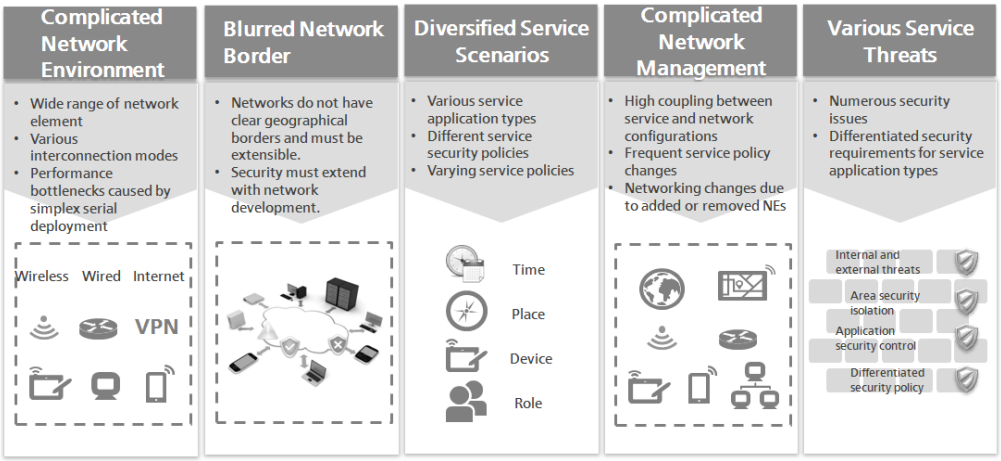

As services and information applications increase rapidly, the network environment becomes complicated, and the network border tends to be blurred. Diversified service application scenarios introduce new security threats and bring more challenges to network management.

- **Complicated network environment:** Modern industrial networks have a wide range of network elements with quite different performance. In actual deployment, simplex serial deployment probably causes performance bottlenecks.
- **Blurred network border:** With the popularity of mobile working, employees can access enterprise networks at any time and any place. This implementation blurs network borders. Deploying security devices at borders is no more a good idea. Security capabilities must extend along with network development.

- **Diversified service scenarios:** Diversified service application types lead to differentiated service processing requirements and service policies and ultimately various service deployment scenarios.
- **Various service threats:** Diversified service application types also cause a variety of security issues and threat types. Each type of service flow requires different processing procedures.
- **Complicated network management:** The previous changes and requirements increase service and network configuration coupling, make policies change more frequently, and require dynamic service adaptation and removal. These are new challenges for network management.

1.2 Shortcomings in Traditional Networking

In traditional networking, service devices are connected in in-line or off-line mode. The in-line and off-line modes have different shortcomings.

1.2.1 In-line Service Device Networking

Figure 1-2 In-line networking

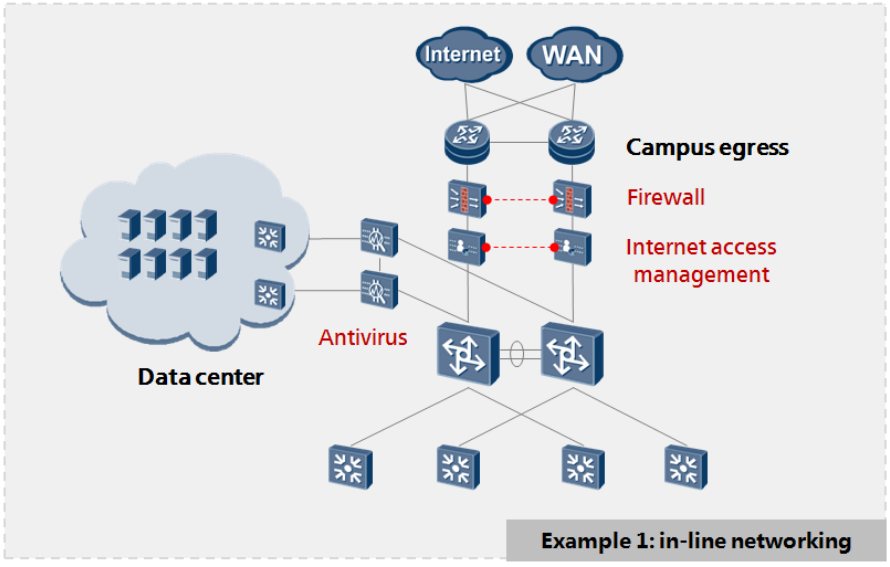

In tradition networking, service devices are deployed in in-line mode. This deployment is simple and intuitive but has the following shortcomings:

- **Performance bottleneck:** The performance of the weakest service device determines the network performance.
- **Undistinguishable traffic:** All traffic flows through service devices. Irrelevant traffic consumes performance resources of the service devices and interferes in fault locating.
- **Required network planning:** IP addresses must be reserved and routes are designed for Layer 3 services in in-line networking.

- **Difficult to add or remove devices:** IP addresses and routes must be re-planned if service devices need to be added or removed. Network communication must be interrupted for cutovers, affecting services.

1.2.2 Policy-based Routing Off-line Service Device Networking

Figure 1-3 Off-line networking

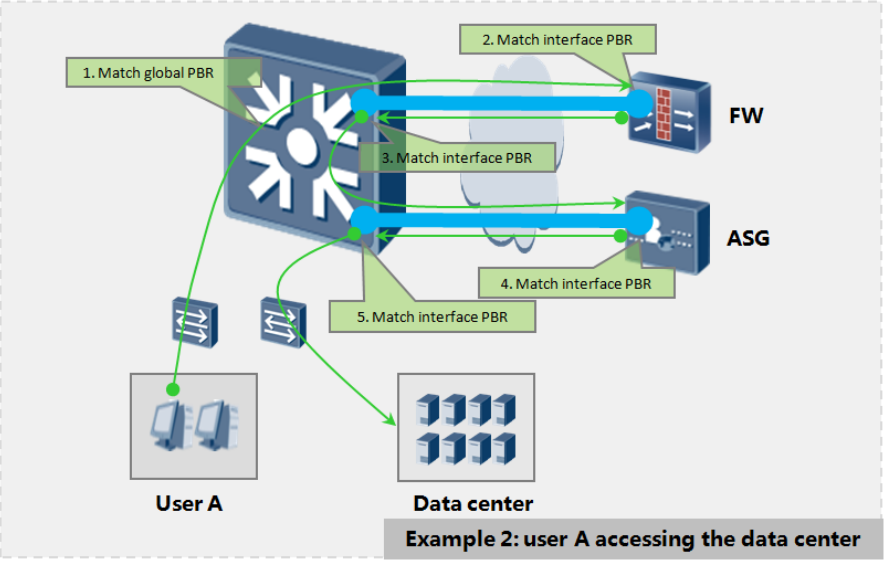

In off-line networking, service devices use policy-based routing to divert traffic. This mode resolves the issues encountered in in-line mode but has the following shortcomings:

- **Complicated configuration:** As shown in Figure 1-3, two service devices are deployed in off-line mode. Policies must be configured globally and on each interface. At least 5 policies are required.
- **Required analysis:** Traffic directions, incoming interfaces, outgoing interfaces, and next-hop addresses must be analyzed.
- **Insufficient reliability:** When two interfaces are used to connect to service devices, these interfaces are not associated. If one interface fails, the other interface is unaware of the fault. For example, if the incoming link to the switch fails, the switch still sends traffic to the service device through the outgoing link. As a result, the traffic is discarded, and services are interrupted.

1.3 Huawei Agile Controller Service Chain Solution

To cope with the previous challenges and changes, Huawei provides the Agile controller Service Chain solution.

1.3.1 Service Chain Principles

1. Service devices and switches are on a Layer 3 network. Each service device connects to a switch through a GRE tunnel for adjacency.

2. The switch identifies service traffic, processes services in a user-defined order, and diverts traffic to appropriate service devices.

The Service Chain solution is based on GRE tunnels and policy-based routing. The agile controller provides visible operations for agile management on enterprise campus networks.

1.3.2 Benefits to Customers

The Service Chain solution brings the following values to customers:

- Easy to understand

  Service logic is abstracted, and the actual networking is transparent to users. Customers only need to focus on service flow definition and service chain logic, which is close to the service plane, easy to understand and use.

- Flexible networking

  Service devices interconnect with the switch through Layer 3 GRE tunnels, allowing flexible service device networking and deployment locations.

- Extensible

  To add or remove service devices, you only need to add or remove GRE logical links without modifying routes or physical topologies.

- Simplified management

  The agile controller automatically interconnects with service devices, adds and deletes service policies, greatly reducing the policy configuration workload.

  The live network modification caused by adding or removing service devices is minimized, reducing basic network management complexity.

2 Concepts and Principles

2.1 Service Chain Concepts

- Service Flow: a flow containing packets with specified service characteristics (for example, packets with a specified field) or packets matching certain service policies (for example, packets accessing a specified IP address or port).
- Chain node: a standalone value-added service device on a service chain, such as a firewall or an antivirus security device.
- Service Chain: a chain of chain nodes providing a composite service. Packets of several service flows travel through one service chain.

2.2 Service Chain Principles

2.2.1 Networking Structure and System Functions

Figure 2-1 Networking structure

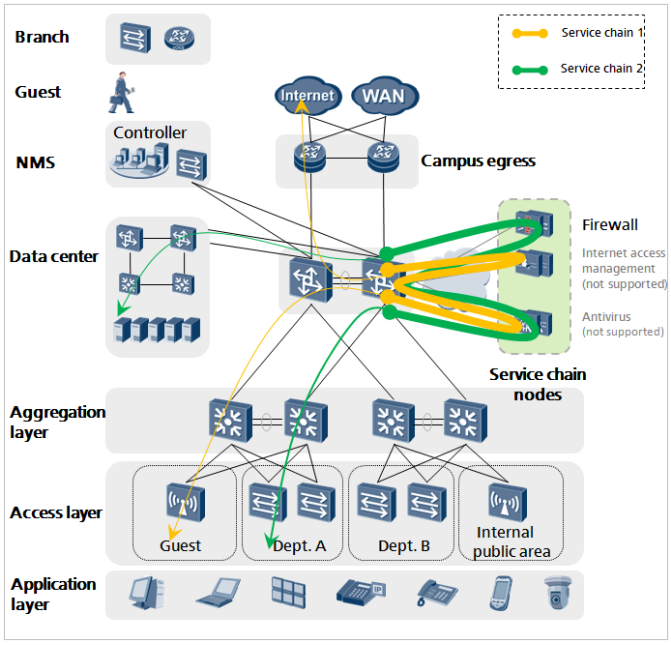

The agile controller, switches, and service devices work at Layer 3.

Layer 3 GRE tunnels are established between service devices and switches. The switches divert traffic to service devices through GRE tunnels based on ACLs or UCLs.

In the previous networking, the Service Chain solution supports the following features:

- Service logic abstraction: Service chain functions are uniformly configured on the agile controller.
- Layer 3 service devices: Layer 3 GRE tunnels are established for interconnection with service devices.
- Layer 3 IPv4 unicast packet Chain: The Chain switch identifies and diverts IPv4 unicast packets forwarded based on Layer 3 routes.
- Service flow definition based on 5-tuple ACLs or dynamic user group UCLs.
- HA configuration on GRE tunnels. If a GRE tunnel fails, packets can be directly forwarded based on routes or discarded.

2.2.2 General Architecture

Figure 2-2 General architecture

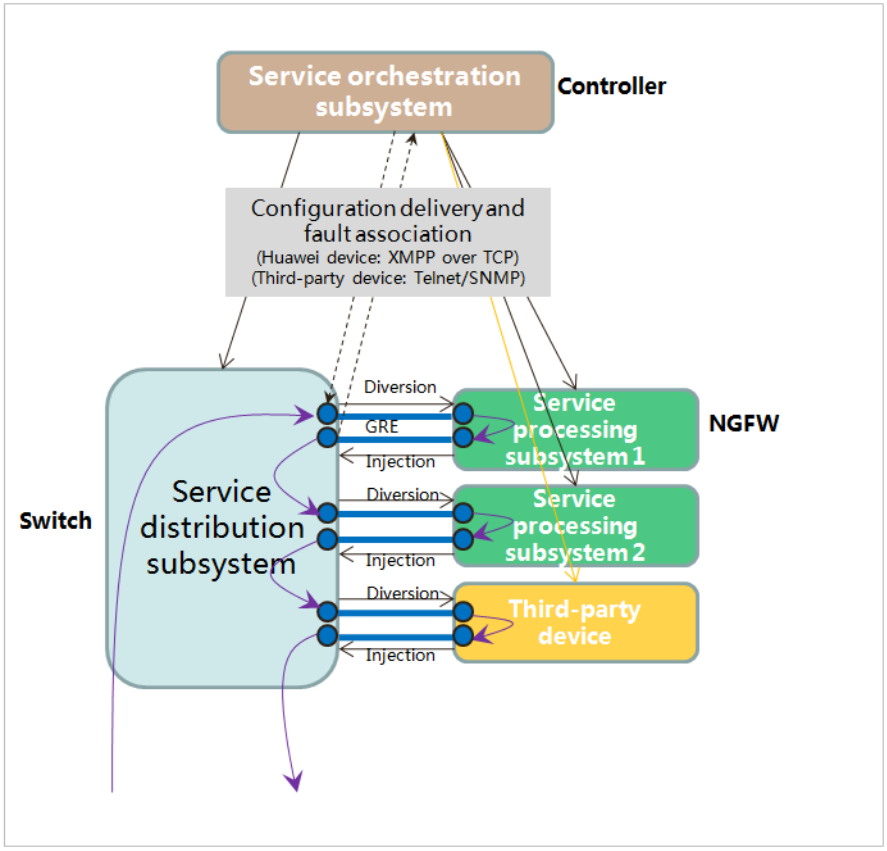

The Service Chain solution consists of three parts:

- Service Chain subsystem: agile controller, configuring service logic for service chains.
- Service distribution subsystem: switch, identifying and redirecting service traffic.
- Service processing subsystem: service devices, processing diverted service traffic.

The Service Chain solution uses the following technologies:

- Machine-to-machine interface configuration:
  - XMPP: used for interconnection between Huawei Chain devices (or service devices) and the agile controller.
  - Telnet/SNMP: used for interconnection between third-party service devices and the agile controller.
- Traffic diversion: GRE tunnel-based policy-based routing

2.2.3 General Process

Figure 2-3 General process

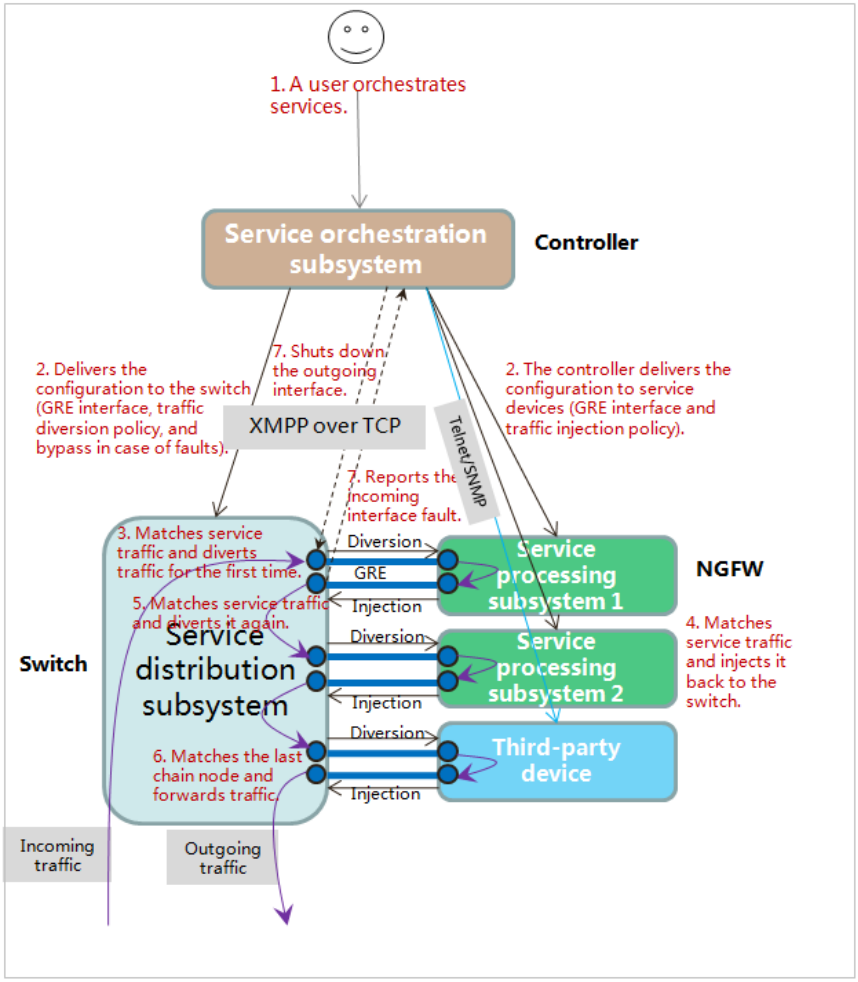

The process is as follows:

1. A user orchestrates services on the agile controller. The configuration includes the service flows, chain nodes, Chain switch, and service chains.
2. The agile controller translates service logic into a machine language and uses XMPP to deliver the configuration to the switch and service devices.
3. When service traffic enters the switch for the first time, the switch queries global policy-based routing. If the traffic matches a service flow rule, the switch diverts the traffic to the first service device based on the rule.
4. The service device processes the traffic and injects it back to the switch based on policy-based routing configured on the incoming interface.
5. After receiving the traffic, the switch finds the interface policy-based routing with a higher priority and diverts the traffic to the second service device based on the rule.
6. After the last service device injects the service traffic back to the switch, the switch finds an appropriate interface policy-based routing, permits the traffic, and forwards it based on a route.

7. The switch uses the GRE Keepalive mechanism to detect the status of GRE tunnels. If one GRE link between the switch and a service device fails, the failure is reported to the agile controller, and the agile controller shuts down the other GRE tunnel on the switch to prevent traffic forwarding abnormalities.

3 Service Chain Functions and Processes

Figure 3-1 shows Service Chain functions.

Figure 3-1 Service Chain functions

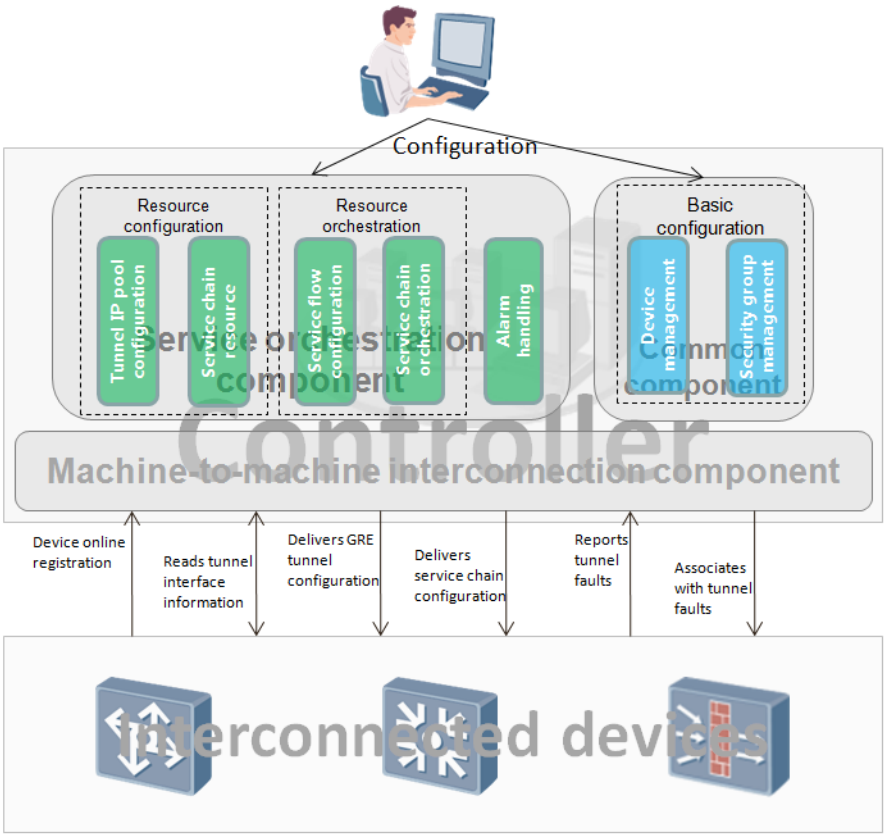

Service Chain functions are as follows:

- Basic configuration: The agile controller basic component provides device management and security group management functions.
  - Device management: The agile controller manages the online and offline of the switch and service devices.

  - Security group management: The agile controller manages user identification policies and security group dynamic mapping.
- Resource configuration: The agile controller manages the resource relationships of service chains.
  - IP pool configuration for tunnels: When automatically creating a tunnel interface on an interconnected device, the agile controller specifies an IP address range for the interface.
  - Service chain resource: You can specify service device resources that can be used by the Chain device and manage GRE tunnel relationships between them.
- Resource Chain: Service flows are diverted to service devices in a certain order.
  - Service flow configuration: You can specify the characteristics of the traffic passing through a service chain.
  - Service chain Chain: You can arrange the Chain device and several service devices in a certain order to form a service chain.
- Alarm handling: The agile controller receives tunnel faults reported by the switch and disables (or enables) related interfaces.
- Machine-to-machine interconnection: The agile controller translates service configuration into a machine language and uses XMPP or Telnet/SNMP to deliver the configuration to devices.

3.1 Basic Configuration

Figure 3-2 Basic configuration process

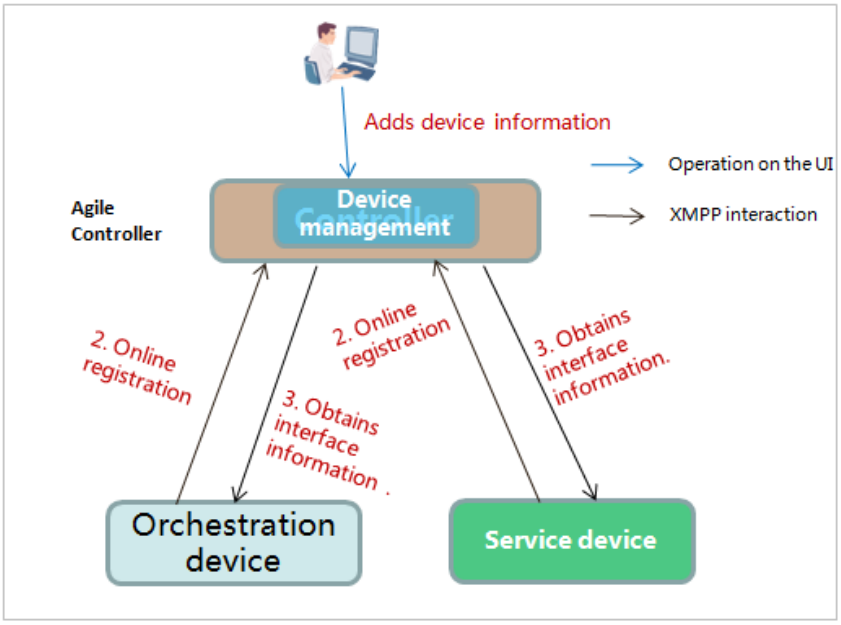

As shown in Figure 3-2, the basic configuration service process is as follows:

1. A user add device information on the agile controller, including the device name and protocol.
2. The Chain and service devices register with the agile controller. The device management module verifies the validity of the devices based on the configured information. If the verification succeeds, the device management module shows the list and information of registered devices.
3. The agile controller obtains loopback interface information from the Chain and service devices for resource configurations.

3.1.1 Resource Configuration

Figure 3-3 Resource configuration process

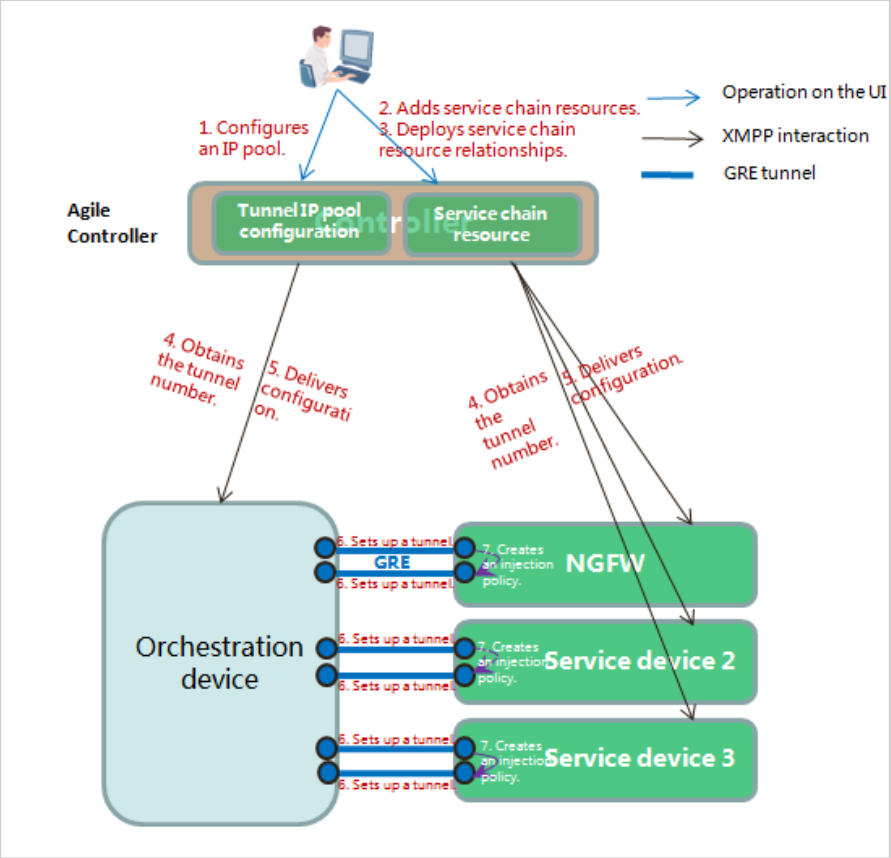

As shown in Figure 3-3, the resource configuration service process is as follows:

1. A user configures the range of IP addresses used by GRE tunnels.
2. The user adds service devices to be used by the Chain device to establish service chain resource relationships.
3. The user deploys the service chain resource relationships to devices.
4. The agile controller obtains idle tunnel interface numbers from devices to generate GRE tunnel configurations.
5. The agile controller delivers GRE tunnel configurations to devices.
6. Based on the received configurations, the Chain and service devices establish GRE tunnels.
7. The user configures policy-based routing between GRE tunnels on each service device to form a loop between the service device and the switch.

3.2 Resource Chain

Figure 3-4 Resource Chain process

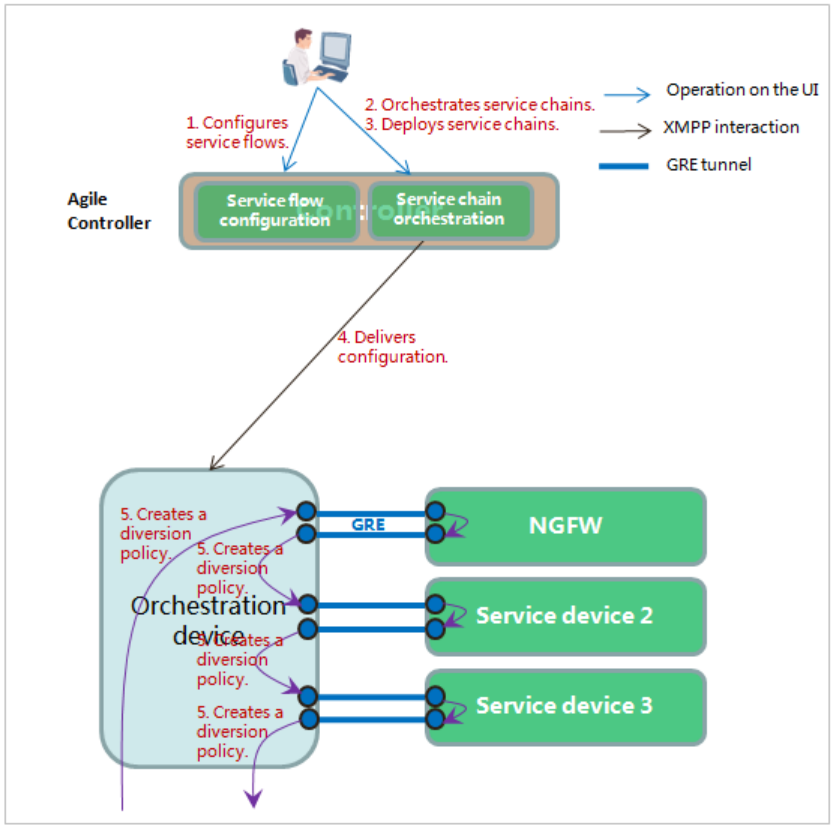

As shown in Figure 3-4, the resource Chain service process is as follows:

1. A user configures service flows based on ACLs or UCLs.
2. The user arranges the Chain device and service devices in sequence to create service chains.
3. The user deploys service chains to the devices.
4. The agile controller delivers service chain configurations to the Chain device.
5. The user configures switch traffic diversion policies globally and on GRE tunnel interfaces to guide service flows to specified service devices in certain orders.

3.3 GRE Tunnel Alarm Handling

Two GRE tunnels are established between the switch and each service device. If the link connecting a service device to the switch fails, traffic flows from the switch to the service device but is discarded by the service device, interrupting services.

The agile controller provides the alarm handling function to resolve the previous issue.

Figure 3-5 GRE tunnel alarm handling process

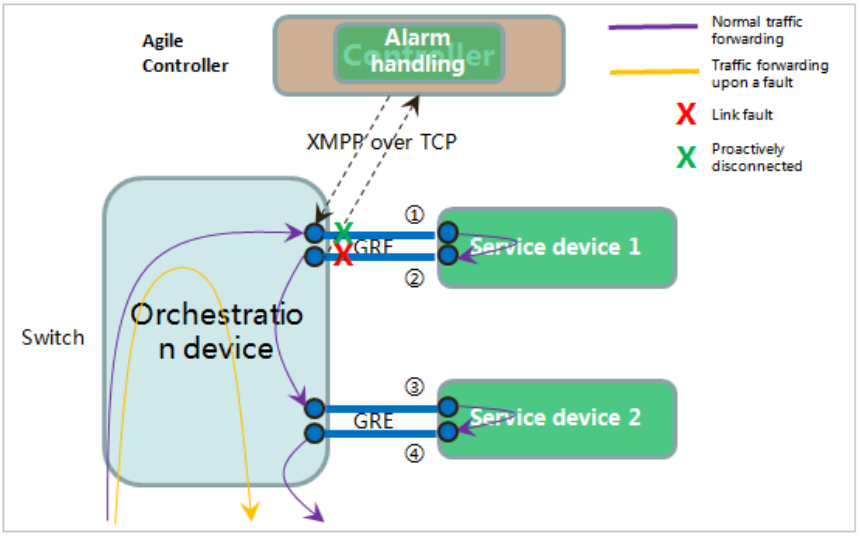

As shown in Figure 3-5, if GRE tunnel 1 fails, the switch processes traffic as follows:

1. The Keepalive mechanism is enabled on GRE tunnels by default. After the switch detects the GRE tunnel fault, it uses XMPP to report the fault to the agile controller.
2. After receiving the alarm, the agile controller shuts down the other GRE tunnel to prevent traffic forwarding abnormalities.
3. If the switch finds that the faulty GRE tunnel recovers, it reports the recovery to the agile controller.
4. After receiving the report, the agile controller restores the other GRE tunnel.

3.4 Machine-to-Machine Interconnection

The Service Chain solution enables the agile controller to provide the adaptable machine-to-machine interconnection function.

Figure 3-6 Machine-to-machine interconnection architecture

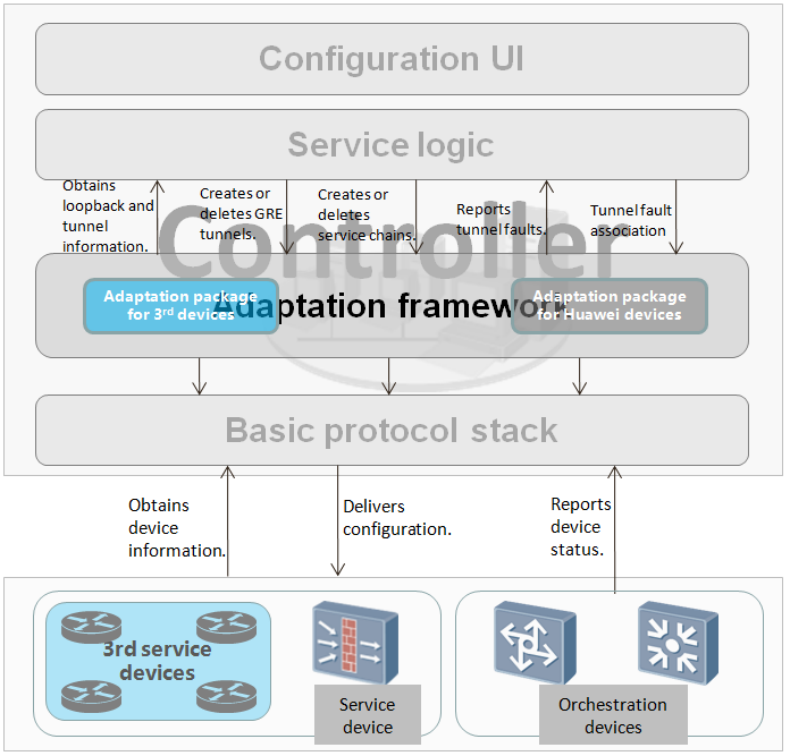

The operations of the agile controller on devices are abstracted as multiple operation interfaces. The adaptation framework is dynamically associated with adaptation packages to support service device extension.

In addition to Huawei NGFW series service devices, the agile controller allows third-party developers to adapt their service devices based on development specifications.

The requirements of service device adaptation are as follows:

- Each third-party service device must have the following capabilities:
  - Allows the agile controller to use SNMP to obtain sysObjectID for device type identification.
  - Supports GRE tunnel configuration using commands.
  - Supports policy-based routing between GRE tunnels for traffic diversion.
- The agile controller provides the following adaptation functions:
  - Extensible third-party service device adaptation mechanism
  - Support for Telnet/SNMP basic protocol stack
  - Abstract service adaptation framework and interfaces to allow the query of idle tunnels and loopback interfaces and the creation/deletion of tunnels.

4 Application Implementation

In the light of agile network characteristics, agile controller Service Chain applies to the following typical scenarios:

- Intranet users accessing a data center
- Intranet users accessing the Internet
- Internet users accessing a data center

4.1 Scenario 1: Intranet Users Accessing a Data Center

When an intranet user accesses a data center, the upstream traffic flows through the access switch, aggregation switch, and core switch and finally reaches the data center. One of the switches can be selected as the Chain device. As service devices are deployed in a security zone to serve the entire network, if the access switch acts as the Chain device, traffic from the aggregation layer bypasses to the core layer, consuming bandwidth resources and increasing delays. Therefore, using the access switch as the Chain device is not preferred. The aggregation or core switch can be used as the Chain device.

4.1.1 Mode 1: Core Switch Acting as the Chain Device

Figure 4-1 Networking for intranet users to access the data center in mode 1

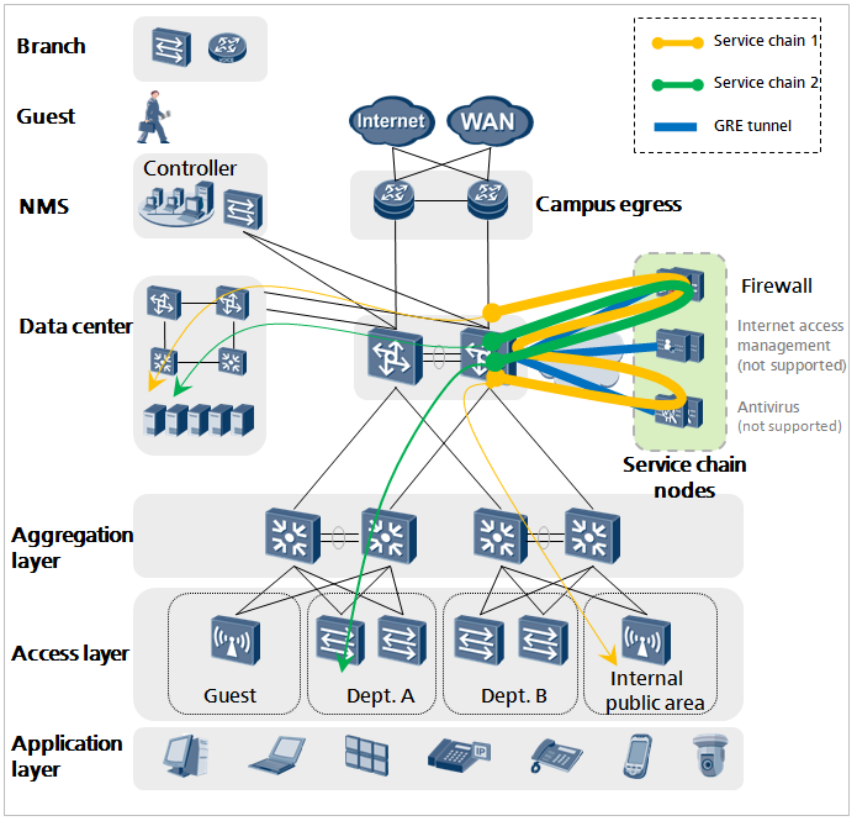

As shown in Figure 4-1, the core switch is selected as the Chain device. Service devices interconnect with the Chain device through Layer 3 GRE tunnels.

For departments on the campus network, the following typical potential risks and requirements exist:

- Access control and antivirus processing are required for the access of public users to the data center. As public users and terminals are not easy to control, USB disks or own devices may bring viruses and Trojan horses. Therefore, access control policies must be deployed to control the access to the data center, and antivirus gateways are also required for data center server security.
- Only access control is required for the access of users in the R&D department to the data center because the enterprise has strictly hardened and implemented security admittance on hosts used by the R&D department.

Chain device selection:

- Generally, the core switch forwards the user traffic destined for the data center. The core switch is recommended as the Chain device for less bypass traffic and shorter forwarding delays.

Service device location:

- Deploying service devices at the core layer is recommended, so that service chains can share these devices.

Service flow definition:

- Switch of any type: Service flows can be defined using ACLs based on protocols, source IP addresses, source ports, destination IP addresses, and destination ports.
- Switch serving as the authentication node and containing only ENP boards: Service flows can be defined using UCLs based on the user information of protocols, source IP addresses, source ports, destination IP addresses, and destination ports.

Deployment description:

- Each service chain is specified for one requirement. In this example, three service chains are defined.
- The service flows from public users to the data center use service chain 1 (antivirus -> firewall).
- The service flows from the data center to public users use service chain 1-reverse (firewall -> antivirus).
- The service flows from the R&D department to the data center use service chain 2 (firewall).
- The service flows from the data center to the R&D department use service chain 2 (firewall).

4.1.2 Mode 2: Aggregation Switch Acting as the Chain Device

Figure 4-2 Networking for intranet users to access the data center in mode 2

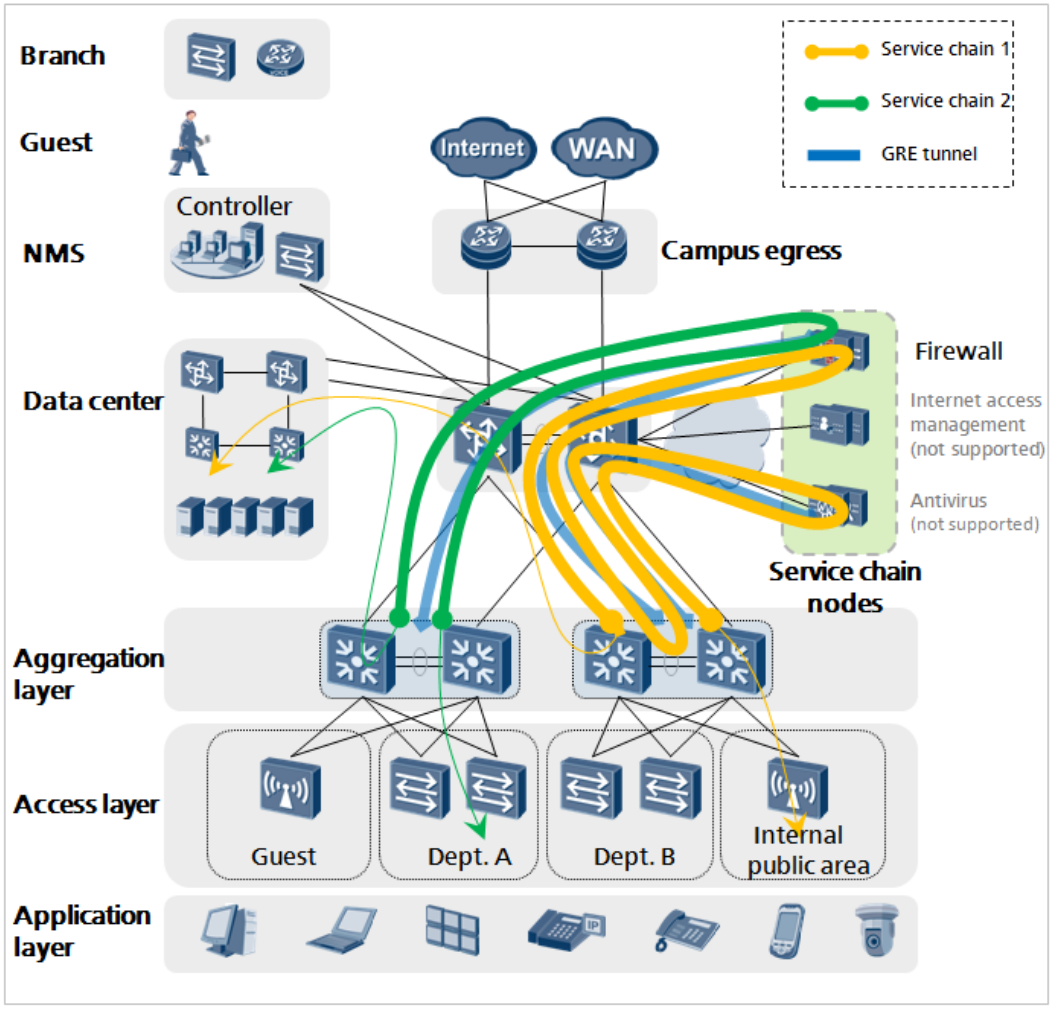

As shown in Figure 4-2, the aggregation switch is selected as the Chain device. Service devices interconnect with the Chain device through Layer 3 GRE tunnels.

For departments on the campus network, the following typical potential risks and requirements exist:

- Access control and antivirus processing are required for the access of public users to the data center. As public users and terminals are not easy to control, USB disks or own devices may bring viruses and Trojan horses. Therefore, access control policies must be deployed to control the access to the data center, and antivirus gateways are also required for data center server security.
- Only access control is required for the access of users in the R&D department to the data center because the enterprise has strictly hardened and implemented security admittance on hosts used by the R&D department.

Chain device selection:

- Users are authenticated on aggregation switches. When service devices are deployed at the core layer, the orchestrated traffic is diverted to the aggregation layer for authentication. If traffic diversion and forwarding delays are tolerable and you want to define services flows based on user information, using aggregation switches as the Chain devices is recommended.

Service device location:

- Deploying service devices at the core layer is recommended, so that service chains can share these devices.

Service flow definition:

- Switch of any type: Service flows can be defined using ACLs based on protocols, source IP addresses, source ports, destination IP addresses, and destination ports.
- Switch serving as the authentication node and containing only ENP boards: Service flows can be defined using UCLs based on the user information of protocols, source IP addresses, source ports, destination IP addresses, and destination ports.

Deployment description:

- Each service chain is specified for one requirement. In this example, three service chains are defined.
- The service flows from public users to the data center use service chain 1 (antivirus -> firewall).
- The service flows from the data center to public users use service chain 1-reverse (firewall -> antivirus).
- The service flows from the R&D department to the data center use service chain 2 (firewall).

4.1.3 Deployment Mode Comparison

The following table lists the comparison between the previous two deployment modes.

Deployment Mode Item	Aggregation Switch Acting as the Chain Device	Core Switch Acting as the Chain Device
ACL-defined service flow	Supported (The switch must be able to distinguish between users based on IP segments.)	Supported (The switch must be able to distinguish between users based on IP segments.)
UCL-defined service flow	Conditionally supported (Condition: The switch has only ENP boards and serves as the authentication node.)	Conditionally supported (Condition: The switch has only ENP boards and serves as the authentication node.)
Service device location	Core layer, shared by aggregation switches	Core layer, shared by aggregation switches
Traffic Chain path	Aggregation layer -> core layer -> service device -> core layer -> aggregation layer	Core layer -> service device -> core layer

Deployment Mode Item	Aggregation Switch Acting as the Chain Device	Core Switch Acting as the Chain Device
Applicable scenario	Traffic bypass is tolerable. Defining service flows based on user information is expected.	Traffic bypass should be avoided. Service flows can be defined based on IP information.

4.2 Scenario 2: Intranet Users Accessing the Internet

Figure 4-3 Networking for intranet users to access the Internet

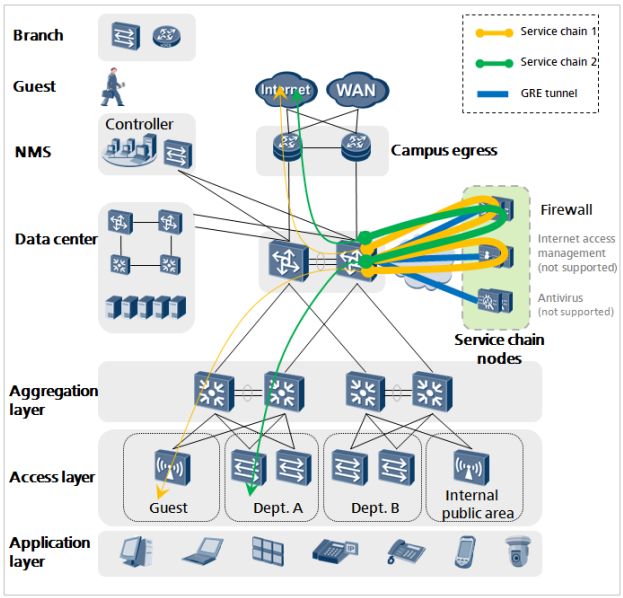

As shown in Figure 4-3, the core switch is selected as the Chain device. Service devices interconnect with the Chain device through Layer 3 GRE tunnels.

For departments on the campus network, the following typical potential risks and requirements exist:

- When guests access the Internet, their access must be controlled and NAT is required. As guests access the Internet from an intranet, their Internet access permissions are under control; their online behaviors (including comments and posts) are monitored and filtered based on keywords; their online application bandwidths are restricted to prevent adverse impact on enterprise services.
- When employees access the Internet, only access control is required. As the enterprise has strictly hardened and implemented security admittance on hosts used by the R&D department, only NAT is required for Internet access traffic in most cases.

Chain device selection:

- Generally, the core switch forwards the user traffic destined for the Internet. The core switch is recommended as the Chain device for less bypass traffic and shorter forwarding delays.

  Users are authenticated on aggregation switches. When service devices are deployed at the core layer, the orchestrated traffic is diverted to the aggregation layer for authentication. If traffic diversion and forwarding delays are tolerable and you want to define services flows based on user information, using aggregation switches as the Chain devices is recommended. The networking mode is the same as that in section 4.1.2.

Service device location:

- Deploying service devices at the core layer is recommended, so that service chains can share these devices.

Service flow definition:

- Switch of any type: Service flows can be defined using ACLs based on protocols, source IP addresses, source ports, destination IP addresses, and destination ports.
- Switch serving as the authentication node and containing only ENP boards: Service flows can be defined using UCLs based on the user information of protocols, source IP addresses, source ports, destination IP addresses, and destination ports.

Deployment description:

- Each service chain is specified for one requirement. In this example, three service chains are defined.
- The service flows from guests to the Internet use service chain 1 (Internet access control -> firewall).
- The service flows from the Internet to guests use service chain 1-reverse (firewall -> Internet access control).
- The service flows from employees to the Internet use service chain 2 (firewall).
- The service flows from the Internet to employees use service chain 2 (firewall).

4.3 Scenario 3: Internet Users Accessing the Data Center

Figure 4-4 Networking for Internet users to access the data center

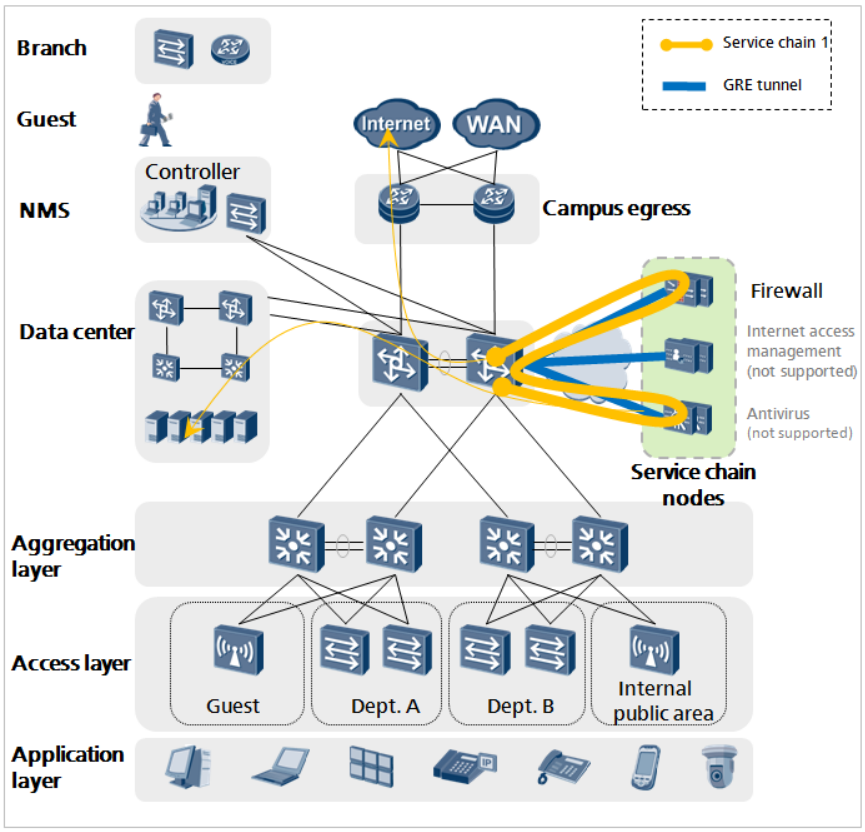

As shown in Figure 4-4, the core switch is selected as the Chain device. Service devices interconnect with the Chain device through Layer 3 GRE tunnels.

For servers in the data center, the following typical potential risks and requirements exist:

- NAT and antivirus processing are required for Internet users to access the data center.

Chain device selection:

- In most cases, the core switch forwards traffic from Internet users to the data center. Using the core switch as the Chain device is recommended.

Service device location:

- Deploying service devices at the core layer is recommended, so that service chains can share these devices.

Service flow definition:

- NAC authentication does not apply to Internet users. Therefore, service flows can be defined only using ACLs based on protocols, source IP addresses, source ports, destination IP addresses, and destination ports.

Deployment description:

- The service flows from Internet users to the data center use service chain 1 (firewall -> antivirus).

5 Reference

None.