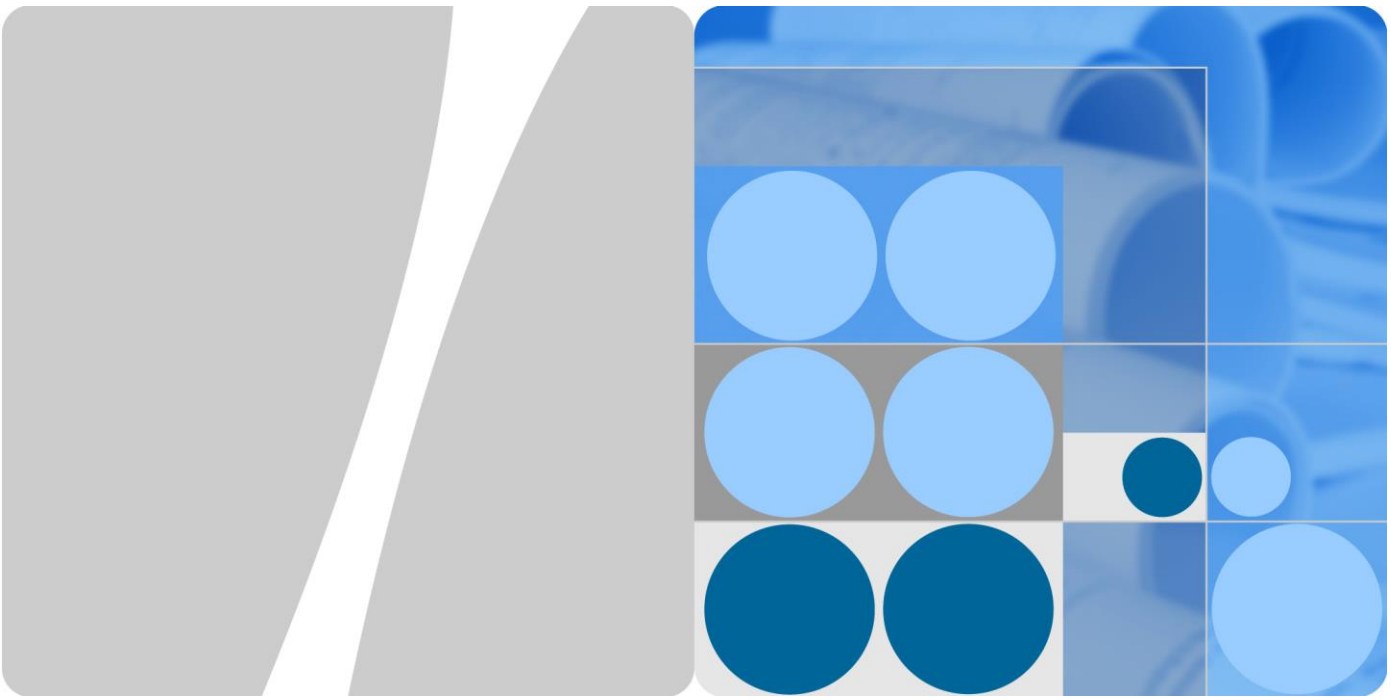

Agile Controller-Campus
V100R002C10

Free Mobility Technology White Paper

Issue 01

Date 2016-04-15

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

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 Overview 1
1.1 Background 1
1.2 Technologies 1
2 Access Policy Control 3
2.1 Access Permission Control 3
2.1.1 Security Group Planning 7
2.1.2 Security Group Access Control Policy Planning 12
2.1.2.1 Special Security Groups 12
2.1.2.2 Domain-based Hierarchical Policy Management 16
2.2 Network Experience Insurance 18
3 Typical Networking Applications 20
3.1 Central Access Permission Control 20
3.2 Distributed Access Permission Control 21
3.3 Network Experience Insurance 22
4 References 23

1 Overview

1.1 Background

1.2 Technologies

1.1 Background

1.2 Technologies

1.1 Background

For the enterprises that are running business in multiple countries, their global employees and partners need to access the enterprise networks for daily work. The increasingly developing wireless networks and remote access technologies, such as VPN, are blurring enterprise network borders. Employees may access enterprise networks from variable locations. Therefore, BYOD has become the trend of network development. With BYOD devices, employees may access enterprise networks at any location. Then how to ensure that employees can conduct daily work without being affected poses a great challenge to enterprise network management and network security. For security reasons, access control policies must be executed regardless of accesses at any place, in any way, at any time, and with any device. On the other hand, when VIPs access the enterprise network, their network access experience must be guaranteed.

To meet the challenges for enterprise network development, Huawei launches the Agile Controller solution to function with other network devices to control the network access from both inside and outside the enterprise network and ensure Free Mobility for employees.

1.2 Technologies

Huawei agile network controls user access based on security groups rather than VLAN+ACL. All policies are based on security groups, not the 5-tuple. Security groups, for sake of security, dynamically categorize users and the resources that the users need to access. The dynamic process is implemented through user authentication and flexible matching of the configured rules on the Agile Controller server. With security groups, policies do not rely on the traditional 5-tuple. Therefore, network IP addresses are decoupled from policies, and network access is no longer implemented by IP address. In this way, the access policy for each user remains the same even when the user's IP address changes.

In actual applications, administrators can define security groups in the Agile Controller center, configure policies for these security groups, and deliver the policies to devices. When users access the network, the Agile Controller can dynamically categorize the users into different security groups based on the configured rules. Then the devices can implement dynamic network access control on users based on the configured policies. Therefore, the services of users can be dynamically adjusted no matter how users access the network.

2 Access Policy Control

2.1 Access Permission Control

2.2 Network Experience Insurance

2.1 Access Permission Control

2.2 Network Experience Insurance

2.1 Access Permission Control

For the sake of information asset security, enterprises categorize users and grant different permissions for them to access information assets. Specifically, the enterprises define the IP address ranges of the data center servers that users are permitted to access.

Some enterprises may also require controlling mutual access between users, because if users of different types can access each other, a user can use another as the springboard to access the servers that the user does not have permission on.

On traditional campus networks, the NAC, VLAN, and ACL technologies are used to control users' network access permissions. However, with the development of campus networks and the emergence of BYOD, these technologies are no longer applicable because they have weaknesses.

Table 2-1 Weaknesses of NAC, VLAN, and ACL

Technology	Mechanism	Weaknesses
Dynamic VLAN+static ACL (access location-sensitive)	Configure static ACLs in advance on the switch and bind the ACLs to VLANs. When users log in, the authentication center assigns the users to different VLANs based on the user identities and access locations. That is, the authentication center may assign users of different types to different	• Multiple authentication policies must be configured based on users' possible access locations. • The pre-configuration requires a lot of time, which is not location-irrelevant from the management perspective. • VLANs and static ACLs

Technology	Mechanism	Weaknesses
	VLANs even if they access the same switch for authentication and may assign users of the same type to different VLANs even if they access different switches for authentication. The network access permissions of the users on different VLANs vary because the users are bound with different ACLs.	must be pre-configured on each switch in advance, which requires heavy workloads. In addition, the ACLs are difficult to maintain. • Each user type corresponds to one or more VLANs. • Each time a user type is added, each access switch needs to add a VLAN, and the corresponding gateway needs to reserve IP addresses. • However, in actual deployment, adding and deleting user types are difficult, and a large number of VLANs and IP address segments are idle during network operation.
Dynamic VLAN+static ACL (access location-insensitive)	Configure static ACLs in advance on the switch and bind the ACLs to VLANs. When users log in, the authentication center assigns the users to different VLANs based on the user identities. That is, the authentication center may assign users of the same type to the same VLAN even if they access different switches for authentication. The network access permissions of the users on different VLANs vary because the users are bound with different ACLs.	• The VLANs must be deployed across switches. • If multiple access switches share the same gateway, they also share the same VLAN. Therefore, a big broadcast domain will be formed. • VLANs and static ACLs must be pre-configured on each switch in advance, which requires heavy workloads. In addition, the ACLs are difficult to maintain. • A large number of IP addresses on the VLAN are not in use. • For example, if an access switch supports a maximum of 48 users, and 10 access switches connect to each gateway. • Then each gateway

Technology	Mechanism	Weaknesses
		needs to reserve a minimum of 480 IP addresses for each type of users. • However, when different types of users access the same gateway, a lot of IP addresses are not in use.
Dynamic VLAN+dynamic ACL	When users log in, the authentication center assigns VLANs and bound ACLs based on user identities. The VLANs isolate different types of users. ACLs are bound to each user. The ACL rules can be pre-configured on the switches or dynamically delivered by the authentication center. The network access permissions of the users vary because the users are bound with different ACLs.	• The VLANs must be deployed across switches. • If multiple access switches (such as switches A and B in the example) share the same gateway, they also share the same VLAN. Therefore, a big broadcast domain will be formed. • Although ACLs can be delivered by the authentication center to reduce the ACL configuration workload, the ACLs in this solution must be bound one by one to users because switches use hardware to match ACLs. Therefore, users cannot share ACLs even if they are the same type. • As a result, the number of rules in each ACL is limited. Otherwise, the ACLs may fail to take effect because of the limit of the switch processing chip on the number of rules.
Static VLAN+dynamic ACL	Statically create VLANs based on interfaces on the switches to prevent broadcast domains and idle WLANs. When users log in, the authentication center assigns the bound ACLs based on user identities.	Failed to isolate mutual access on the same VLAN, because users of the same type may use addresses on the same subnet, and the ACLs for isolating mutual access cannot be pre-configured. For the firewalls that are

Technology	Mechanism	Weaknesses
	ACLs are bound to each user. The ACL rules can be pre-configured on the switches or dynamically delivered by the authentication center. The network access permissions of the users vary because the users are bound with different ACLs.	deployed at network borders instead of authentication points, the traditional IP address-based packet-filtering policies are difficult to configure. Although ACLs can be delivered by the authentication center to reduce the ACL configuration workload, the ACLs in this solution must be bound one by one to users because switches use hardware to match ACLs. Therefore, users cannot share ACLs even if they are in the same group. As a result, the number of rules in each ACL is limited. Otherwise, the ACLs may fail to take effect because of the limit of the switch processing chip on the number of rules.

The analysis shows that the traditional NAC technology has the following key technical problems regardless of whether ACLs are dynamically delivered or statically bound:

- The ACLs must be configured in advance by the administrator and must have the destination IP address ranges that are permitted to access or not. In other words, the ACLs can be decoupled only from sources, not from destinations. Therefore, if the IP address ranges are not fixed, ACLs cannot be used to control mutual access between users.
- The binding between ACLs and users takes effect only on the authentication devices. That is, only the authentication device can identify users. Non-authentication devices, such as the firewalls deployed at the campus network borders, can configure only IP address-based policies. Therefore, if users' IP address ranges are not fixed, firewalls will fail to configure policies. The existing campus networks still need to assign different types of users to different VLANs so that the different types of users can have fixed IP addresses.
- VLANs and ACLs in the traditional NAC technology still need to be configured in advance on the authentication switches, which requires heavy deployment and maintenance workloads. If any configuration needs any change, there would be lots of work for the network administrator.

To resolve these problems, Huawei launches the Free Mobility solution.

In access permission control of the Free Mobility solution, the Agile Controller functions as the control center of the campus network and is responsible for user access control and user permission control on the campus network.

The solution divides the campus network into three logical planes, and the Agile Controller resides in the center of the service management plane. The Agile Controller supports the standard wireless and wired 802.1x/Portal/MAC Bypass authentication, categorizes network access users into security groups based on predefined rules, and dynamically delivers them to authentication devices. If a non-authentication policy execution device fails to identify the security group of a certain IP address, the device proactively queries from the Agile Controller (only the agile firewall supports proactive query). The following figure shows the logical architecture.

Figure 2-1 Logical architecture of the Free Mobility system

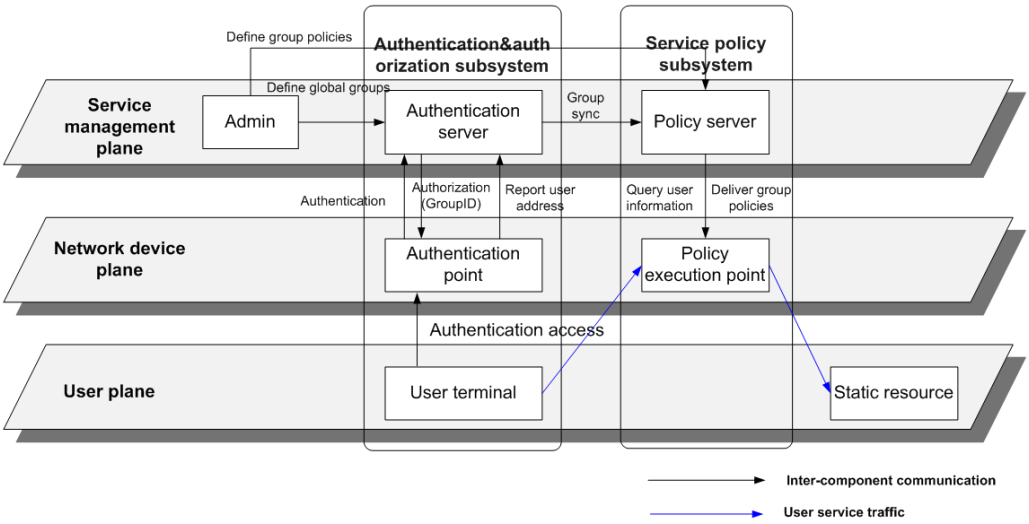

2.1.1 Security Group Planning

Security groups are classified into two categories: static security group (bound with static IP addresses) and user security group (dynamically allocated by the Agile Controller based on authentication parameters). User security groups determine users' permissions and identities, while static security groups classify static server resources.

In addition to the user security groups and static security groups, the Agile Controller supports resource groups. Similar to static security groups, resource groups classify static server resources. The differences between a resource group and a static security group are as follows:

1. Limited by devices' forwarding mechanism, only IP addresses and masks but not protocols and port numbers can be bound to static security groups. Besides, any two static security groups cannot have overlapping IP address segments.
2. IP addresses, masks, protocols, and port numbers can be bound to resource groups, and the groups can be configured with overlapping IP address segments. However, resource groups occupy a large number of ACL rules (Each entry bound to a resource group requires an ACL rule).

Generally, resource groups are used to classify resources of specified protocols and port numbers, and overlapping IP addresses. Resource groups supplement static security groups when static security groups cannot meet service requirements.

Security groups are the foundation in the Free Mobility solution. In practical applications of Free Mobility, user security groups and static security group must be planned before other configurations.

Static Resource Security Group Planning

The security group planning of static resources refers to the planning of application system servers in the data center based on enterprise security requirements.

This type of security groups must be bound to IP address segments, and each security group can bind a maximum of 2048 IP addresses or IP address segments. A security group can contain multiple discontinuous IP address segments.

- Example 1: Company A does not control server access permissions. Users can access all servers of the company as long as they can access the Intranet. In this case, you need to define a security group for all servers of company A and bind the IP address segments of all servers to the security group.
- Example 2: Company B has deployed servers respectively for marketing and R&D. Users can access corresponding servers based on their identities after they access the Intranet. In this case, you need to define two security groups for the servers of company B and bind the IP address segments of the servers for marketing and R&D respectively to the security groups.

When you plan static resource security groups, you need to determine whether a security group needs to be planned to represent the Internet.

Most enterprise networks have two types of special traffic:

- The traffic generated when Intranet users access the Internet
- The traffic generated when Internet hosts access Intranet servers if the enterprise provides network services externally.

Internet hosts access Intranet resources without being authenticated. Therefore, if the administrator does not add the Internet address to a static resource security group, the previous two types of traffic will match the "Intranet user accesses unknown group" and "unknown group accesses server" policies.

Authentication switches identify only the users authenticated on their own. Therefore, the traffic from a user authenticated by one switch to a user authenticated by another switch will match the default policy.

If the actions for the two types of traffic are different, for example, when the enterprise allows cross-authentication point mutual access but prohibits Intranet users from accessing the Internet, the actions of the "Intranet user accesses unknown group" policy conflict.

To resolve this problem, define Internet as a static resource group and configure a policy for Intranet users to access the Internet.

RFC791 defines the range of IPv4 addresses, and RFC1918 reserves some among the IPv4 unicast addresses as private addresses, which cannot be used on the Internet.

One segment for class A private addresses: 10.0.0.0 to 10.255.255.255

16 segments for class B private addresses: 172.16.0.0 to 172.31.255.255

256 segment for class C private addresses: 192.168.0.0 to 192.168.255.255

Therefore, if you need to define the Internet as a security group, exclude these private addresses from the three classes of IPv4 unicast addresses. Then you can have 33 network segments.

Class A	Class B	Class C
64.0.0.0/2	128.0.0.0/3	208.0.0.0/4
32.0.0.0/3	176.0.0.0/4	200.0.0.0/5
16.0.0.0/4	160.0.0.0/5	196.0.0.0/6
0.0.0.0/5	168.0.0.0/6	194.0.0.0/7
12.0.0.0/6	174.0.0.0/7	193.0.0.0/8
8.0.0.0/7	173.0.0.0/8	192.0.0.0/9
11.0.0.0/8	172.128.0.0/9	192.192.0.0/10
	172.64.0.0/10	192.128.0.0/11
	172.32.0.0/11	192.176.0.0/12
	172.0.0.0/12	192.160.0.0/13
		192.172.0.0/14
		192.170.0.0/15
		192.169.0.0/16

Based on the traditional campus network deployment, a typical enterprise usually needs to plan the following static resource security groups:

- Pre-authentication domain server (servers accessible before users are authenticated)

To complete user authentication, hosts need to access some necessary servers, such as the domain server, portal server, DHCP server, DNS server, or patch server, before authentication or during the authentication. These servers must be defined separately for unauthenticated users to use.

When the traffic from an authenticated user to a server in the pre-authentication domain passes through the firewall or SVN device that supports proactive query, the device does not need to look for the security group of the IP address used by this user, because the user is not authenticated yet, and the query result from the Agile Controller must be "unknown group".

The firewall or SVN device that supports proactive query starts refreshing and query 10 minutes after the latest successful query. If the device queries the security group of the IP address during the user authentication phase, the device will regard the user IP address as "unknown group" when the user's traffic passes through in the 10 minutes after the user is authenticated.

Therefore, you need to specify the address ranges (pre-authentication domain) accessible before users are authenticated on the Agile Controller. Then the firewall or SVN can conduct special processing on the traffic to servers in the pre-authentication domain and the traffic returned by the servers from the pre-authentication server. The actions include:

- Permit traffic.
- Do not trigger security group query based on packet source or destination IP addresses.
- Do not generate mapping entries between source or destination IP addresses and security groups.
- Public server

All types of users can access low-security level servers, such as the Intranet websites and mail server.

On the basis of permission granularities and application system types, public servers can be divided into multiple security groups.

- Internet group

If the Intranet and Internet are mutually reachable, you can define an Internet group to describe the Internet users and public servers.

- DMZ server

DMZ servers refer to those that enterprises provide for extranet users to access, such as the company's portal and the address of the SSL VPN gateway deployed on the Intranet. Such servers are usually deployed in the Internet egress area on campus networks or in the DMZ of data centers.

If NAT Server is configured to translate Intranet server address to a public IP address to provide services for extranet users and the public IP address is accessible to all Internet users, add the server to the Internet group, not the DMZ server group. Use application-layer control (such as account and password checks) on the server itself to allow only authorized users to access it.

To allow only specific Internet users to access the public IP address, for example, to provide a dedicated server for partners, you can define an independent security group. In this case, you need to exclude the public IP address of the server from the Internet group.

You can deploy NAT Server to either translate source addresses or not.

If you deploy NAT Server to not translate source addresses and the traffic from Internet users to DMZ servers reaches the Intranet, the packet source addresses are not translated. Usually, NAT Server is deployed to not translate source addresses. In this case, you need to permit mutual access between the Internet group and DMZ server group when you configure access control policies.

If you deploy NAT Server to translate source addresses and the traffic from Internet users to DMZ servers goes through the border gateway firewall, the firewall translates the packet source addresses from public to private addresses. This deployment prevents Intranet from advertising public routes. In this way, you only need to permit only the mutual access between the source NAT address pool (private addresses) and DMZ servers when you configure access control policies.

- Network device group

The IP addresses of network devices on an enterprise network include those of both physical and logical interfaces. Because network devices need to communicate, you need to allow communication between network device groups.

If an enterprise has enabled in-band management (use enterprise service network for network device management without establishing dedicated network device management networks), you are advised to permit only administrator or bastion host IP addresses to log in to network devices through Telnet or SSH for configuration management for security reasons.

To simplify security group configuration in network design, the IP address ranges used by network devices must be separated from those used by user hosts and servers. For example, when you allocate network segment 10.1.0.0/16 to network devices, you are advised to allocate network segments 10.2.0.0/16 and 10.3.0.0/16 respectively to user hosts and servers. In this way, when you define network device groups, you only need to describe one network segment.

- Employee-dedicated application system

Some types of users have dedicated high-security level application systems, such as the code server, test server, and ERP system.

On the basis of permission granularities and application system types, dedicated application systems can be divided into multiple security groups.

If firewall access control or QoS policy control is configured and you have planned static resource security groups, you also need to know the network segments of the entire Intranet, because usually the firewall does not function as an authentication point and needs to query the security group of each IP address from the Agile Controller. The range of IP addresses that the firewall needs to query is usually all network segments of the Intranet.

User Security Group Planning

In user security group planning, you do not need to bind any IP address. For user access, 5W1H (Who, Where, When, Whose, What, and How) is used to plan security groups. In actual applications, the 5W1H is described as follows:

- Who: department, role, and account
- Where: access device, IP address range, and SSID for wireless access
- When: specific network access time
- Whose: whose device, the company's or personal device
- What: access terminal type, operating system, security checks for breaches
- How: wireless or wired access

The following provides three examples of user security group planning:

- Example 1: Company A permits a user to access the network as long as the user has a company account. In this case, you only need to configure one security group. All users that successfully connect to the network are categorized into this group, and the device will execute policies based on the security group. The users that fail the authentication will not be categorized into any group, and the device executes the policy for authentication failure (prohibiting users from access resources by default).
- Example 2: Company B enables security checks for access terminals. Devices with breaches can access only the servers in the isolation domain. In this case, you can configure two security groups respectively for those that can access the Intranet and the isolation domain. For the authentication success group, you can create policies for users to access enterprise Intranet. For the isolation group, you can create policies for users to access isolated network resources, such as the antivirus server and patch server.
- Example 3: Company C has different permissions for marketing employees when they access the Intranet from inside and outside the Intranet. When marketing employees access the Intranet from inside the Intranet, all servers are accessible. However, if they access the Intranet from outside the Intranet, they can access only the mail server and BBS server. In this case, you can configure two security groups for marketing employees respectively for marketing employees accessing the Intranet from inside and outside the Intranet.

Example 1 uses Who to divide security groups, example 2 uses Who and How, and example 3 uses Who and Where. The previous analysis shows that user security group planning is very flexible and enterprises can configure security groups based on actual security control requirements.

2.1.2 Security Group Access Control Policy Planning

2.1.2.1 Special Security Groups

Before you configure access control policies, understand the two predefined special security groups in the system.

The Agile Controller has two default security groups, Any (Groupid: 0) and Unknown (Groupid: 65535).

The Unknown group is used to represent the IP addresses that fail to be identified and does not have any actual member. Network devices use the Unknown group to match service policies for passing traffic. After receiving a packet, if the network device fails to identify the security group of the source or destination IP address, the network device matches service policies to the packet based on the Unknown group.

A network device may fail to identify the security group of an IP address due to multiple reasons. For example:

- Authentication switches can identify only the security groups of the users logging in to themselves. If a local user accesses another user logging in from other authentication points, the authentication switch can identify only the source group of the packet. In this case, the authentication switch will execute the "known group accesses unknown group" policy.
- Non-authentication devices need to proactively query packet source and destination groups from the Agile Controller, because if any network fault between non-authentication devices and the Agile Controller or any fault on the Agile Controller causes query failure, the non-authentication devices will execute the "unknown group accesses unknown group" policy.
- For the traffic from users to the Internet, because public addresses are not used on Intranet and the administrator does not define any static resource security group to represent the public addresses, the network device will execute the "source group accesses unknown group" policy.

The Any group is used to represent any IP address on the network. Network devices use the Any group to match service policies for passing traffic, because in most cases, the permissions for users to access different resource groups vary. You can add users to the Any group to simplify policy configuration.

Assume that an enterprise's servers are divided into 100 security groups. The employees of department A are permitted to access all servers except some. Then the excluded servers are divided into one or more security groups, and the permission to these security groups is configured as "deny" and to the other security groups (Any) as "permit" for the employees in department A. As the Any security groups have the lowest policy priority, the effect of "no permission to some, but permission to all the others" can be achieved. In this way, there is no need to configure the "permit" policy for the employees in department A to access other static resources.

When security group A accesses security group B, the access packets are matched with policies in the following priorities:

- The packets are precisely matched with the policy for accesses from security group A to security group B. If an action is not configured in the policy matrix, the policy matching will continue.

- The packets are matched with the policy for accesses from security group A to security group Any. If an action is not configured in the policy matrix, the policy matching will continue.
- The packets are matched with the policy for accesses from security group Any to security group Any. There is no configuration of this policy in the policy matrix and an action of this policy is built in the network device. Specifically, the switch has the action of "permit", and it forwards traffic even when there is no policy matched. The firewall has the default action of "deny" and it drops traffic by default when there is no policy matched. This default action can be changed to "permit".

Permission Policy Planning

After planning security groups, you can plan access permission policies for security groups.

The Free Mobility solution configures permission policies by logical groups and the policies are decoupled from IP addresses. When planning security group permission policies, you only needs to consider the mutual accesses between two logical groups.

When planning security group permission policies, note the directions of policies. Generally, two terminals transmit packets in both directions.

- For a switch, the A-to-B traffic is not associated with the B-to-A traffic. There are mapping A-to-B policy and B-to-A policy for the traffic in respective directions, which will determine whether the switch forwards the traffic.
- For a firewall, there are sessions. Specifically, the firewall records the IP addresses and port numbers that the communication parties are using, and therefore associates the traffic in both directions and executes policies according to a unified action. Assume that A attempts to access B. When packets reach the firewall, the firewall matches them with the A-to-B policy. If the matched action is "permit", the firewall forwards the packets and establishes a session for communication. When the reply packets from B to A reach the firewall, the firewall will not match them with the B-to-A policy but directly forward them based on the session. Assume that B attempts to access A. When packets reach the firewall, the firewall matches them with the B-to-A policy.

The table below lists the recommended policy matrix configurations.

Requirement	A-to-B Action	A-to-Any Action	B-to-A Action	B-to-Any Action
A is permitted to access B, and B is also permitted to access A.	Permit	-a	Permit	-b
A is not permitted to access B, and B is not permitted to access A, either.	Deny	-	Deny	-
A is permitted to access B, but B is not permitted to	Permit	-	Null - d	Null

Requirement	A-to-B Action	A-to-Any Action	B-to-A Action	B-to-Any Action
access A.				
B is permitted to access A, but A is not permitted to access B.	Null	Null	Permit	-

- a: When the A-to-Any policy is configured as "permit" according to A's overall permissions, the A-to-B policy can be left null.
- b: When the B-to-Any policy is configured as "permit" according to B's overall permissions, the A-to-B policy can be left null.
- c: An agile firewall shall be deployed between A and B to execute policies.
- d: The B-to-A and B-to-Any policies are left null. For the B-to-A traffic that matches the Any-to-Any policy, the action is "permit" on the switch but "deny" on the firewall.

For the requirements of unidirectional accesses only (for example, A-to-B access only), the policy for the reverse direction is left null as follows:

- When A attempts to access B??
  - When the A-to-B packets reach the switch, the action in the matched A-to-B policy is "permit".
  - When the A-to-B packets reach the firewall, the action in the matched A-to-B policy is "permit". In the meantime, the firewall establishes a session.
  - When the B-to-A reply packets reach the firewall, the firewall matches them with the session and forwards them.
  - When the B-to-A reply packets reach the switch, the switch matches them with the Any-to-Any policy as the B-to-A and B-to-Any policies are left null, and therefore forwards them.
- When B attempts to access A??
  - When the B-to-A packets reach the switch, the switch matches them with the Any-to-Any policy as the B-to-A and B-to-Any policies are left null, and therefore forwards them.
  - When the B-to-A packets reach the firewall, the firewall matches them with the Any-to-Any policy as the B-to-A and B-to-Any policies are left null, and therefore blocks them.

The mutual accesses between security groups can be converted as a policy matrix. You can plan the policy matrix as follows:

- Classify users in two types: loose and strict.
  - Loose: no access limits except the explicit ones. Employees in a company are typical loose users.
  - Strict: deny to all accesses except explicit ones. Outsourcing staff and visitors are typical strict users.
- Classify servers in three types: passive, active, and special.
  - Passive: no active access but only passive reply.

  - Active: active access (at special ports).
- Configure user-sourced policies.

  Configure the user-to-Any action ("permit" for a loose user and "deny" for a strict user), and then configure a policy with an action different from the Any policy. Permitting the traffic from a user to an unknown security group is recommended, to avoid that the switch blocks the traffic of mutual user access through different authentication points.
- Configure resource-sourced policies.

  Configure a resource-to-Any action (no action for a passive server but "permit" for an active server), and then configure a policy with "deny" action.
- Configure unknown-group-sourced policies.
  - It is recommended not to configure polices for accesses from an unknown group to Any groups. In other words, the switch, which is not capable of active query, forwards the traffic, and the firewall, which is capable of active query, implements refined controls on the traffic. If the firewall determines an IP address as a known group after a query, the IP address has not been authenticated. Then the firewall uses the default "deny" policy to block the traffic.
  - Do not configure policies (no "permit" or "deny") for accesses between unknown groups, enterprise web servers, Internet groups, and Any groups. A portal authentication user has to access the web page for authentication. If the preceding policies are configured, the traffic for web page access will be forwarded or dropped by the switch, and the portal authentication process will not start.

The following table lists a typical policy matrix.

Table 2-2 A typical policy matrix

	0 - Unknown	**A - Employee**	**B - Outsourcing**	**C - Pre-Authentication**	**D - Common**	**E – Employee Application**	**F - DMZ**	**G - Network Device**	**H – Public Network**	**Any**
0 - Unknown										
A - Employee			×							√
B - Outsourcing	√			√	√				√	×
C - Pre-Authentication									×	√
D -									×	√

	0 - Unknown	A - Employee	B - Outsourcing	C - Pre-Authentication	D - Common	E – Employee Application	F - DMZ	G - Network Device	H – Public Network	Any
Common										
E – Employee Application			×						×	√
FDMZ										
G - Network Device				√				√		×
H – Public Network							√		√	

2.1.2.2 Domain-based Hierarchical Policy Management

Various policy management modes are available to meet enterprise requirements in different scenarios. Possible scenarios are as follows:

1. Network policies are uniformly configured, and all the devices on a campus network are configured with the same policies.
2. Network policies are configured by device group. Devices are classified by the device type and region, and different device groups are configured with different policies.
3. The entire network is divided into multiple VPNs, and policies are configured for each VPN separately.
4. In addition to the preceding scenarios, special policies are configured for each single device separately.

The Agile Controller provides the domain-based hierarchical policy management function to meet requirements in the preceding scenarios.

Domain-based policy management can be implemented from three dimensions:

1. Uniform policies over the entire network: All the devices are configured with the same policies. Modified policies take effect on all the devices.

2. By device group: The administrator adds devices to different groups, and each group is configured with unique policies. Modified policies take effect on only devices in the corresponding group.
3. By VPN: The administrator creates multiple VPN groups, and specifies roles (PE, MCE, and CE) to devices. The PE and MCE devices can be added to multiple VPN groups, while the CE devices can be added to only one VPN group. Devices can obtain policies of the corresponding VPN group.

Besides domain-based policy management, the Agile Controller also supports hierarchical policy management. There are global policies and local policies in hierarchical policy management.

1. Global policies: refer to uniformly configured policies for devices in a domain based on the preceding domain division modes.
2. Local policies: refer to special policies that are configured additionally for a single device, apart from the inherited global policies.

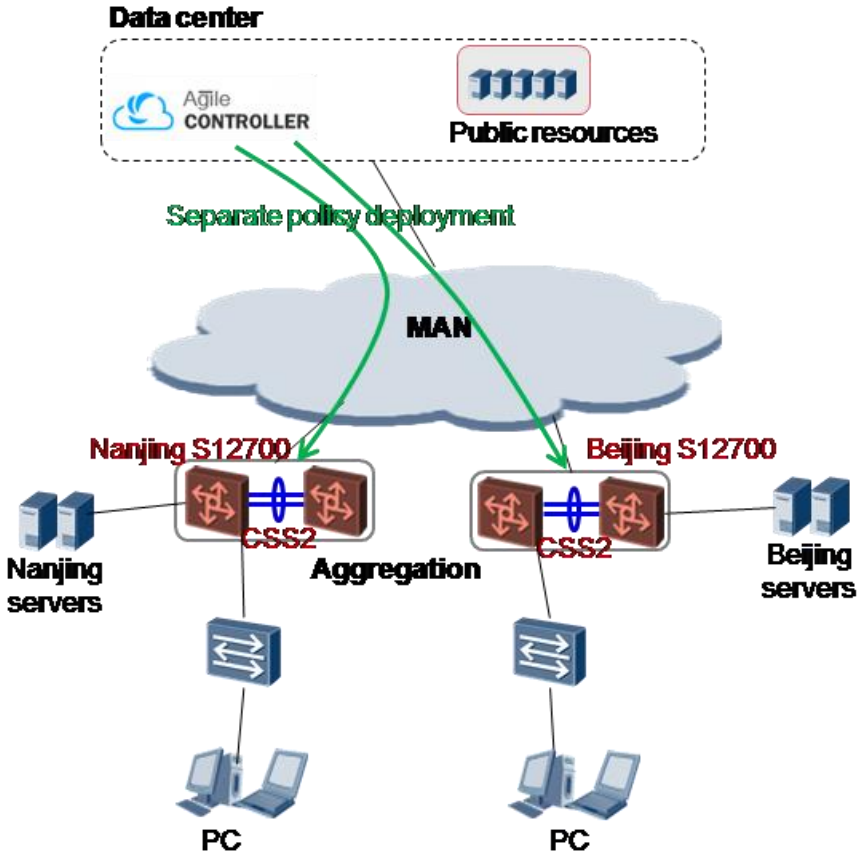

On an enterprise network shown in xxx, there is one data center and two campuses (Nanjing and Beijing). Public resources are deployed in the data center and server resources are deployed independently in the two campuses.

The recommended policy configuration is as follows:

1. Use the uniform policies over the entire network mode, and configure global policies for public resources.

	Public Resources	Any
PC	√	×

2. Configure local policies for the S12700 switches and servers in the Nanjing campus, in addition to the inherited global policies.

	Public Resources	Nanjing Server	Any
PC	√	√	×

3. Configure local policies for the S12700 switches and servers in the Beijing campus, in addition to the inherited global policies.

	Public Resources	Beijing Server	Any
PC	√	√	×

The devices match policies in a descending order from left to right, as shown in the preceding table. To preferentially use a local policy, move the desired policy to the left of global policies.

2.2 Network Experience Insurance

Good experience of network accesses focuses on smooth network service running without disconnections or congestion. The legacy QoS uses techniques like token bucket and queue scheduling to improve key service indicators of delay, jitter, and packet loss ratio.

The Free Mobility solution ensures network experience using the following functions:

- Limits user-specific rates: Based on the unified configurations on the Controller, the Controller delivers the user-specific bandwidth limit to the authentication device when a user attempts to access. This function limits the total Intranet bandwidth for each user and avoids campus network congestion when users are running bandwidth-eating applications.
- Schedules queues by priority: Based on the unified configurations on the Controller, the network device forwards traffic of some security groups by the specified priorities.
- Allows preferred accesses to SSL VPN gateway: Based on unified configurations on the controller, an SSL VPN gateway allows users in some security groups to preempt resources of online common users and access the gateway with preference.

When planning experience insurance policies for security groups, focus on the following issues:

- Which security groups need experience insurance (these security groups will be configured as VIP groups).

  Choose VIP groups based on management demands of the enterprise. VIP groups have higher priorities in forwarding than common users, and VIP group traffic may preempt the bandwidth of common users. When there are many VIP groups, the services of common users may be impacted. To avoid this issue, the Controller allows a maximum of 10 VIP groups. It is also recommended to limit the overall access bandwidth for different types of users (configure the limits in authorization rules).
- Which priorities will apply to forward VIP group traffic.

The Controller provides six priorities: EF, AF4, AF3, AF2, AF1, and BE, which are aligned in a descending order. A network device forwards packets of lower priorities only after forwarding all the packets of higher priorities. It is recommended the set the VIP groups' priority to EF.

- Which devices will be selected as ones to execute the experience insurance policies.

  Select border firewalls that are connected to WAN or Internet, any other border firewalls that have interfaces with bandwidth bottlenecks, or all SSL VPN gateways as the devices to execute the experience insurance policies. If a firewall is not connected to a MAN or carrier network, configure a rate limit for the firewall interfaces that are connected to the WAN or Internet. Specifically, configure the rate limit as the WAN or Internet link bandwidth provided by the carrier.

After you finish the plan, the Agile Controller provides simple configurations for you to deploy QoS policies on devices.

Figure 2-2 QoS policies

3 Typical Networking Applications

3.1 Central Access Permission Control

3.2 Distributed Access Permission Control

3.3 Network Experience Insurance

3.1 Central Access Permission Control

3.2 Distributed Access Permission Control

3.3 Network Experience Insurance

3.1 Central Access Permission Control

Figure 3-1 Networking for central access permission control

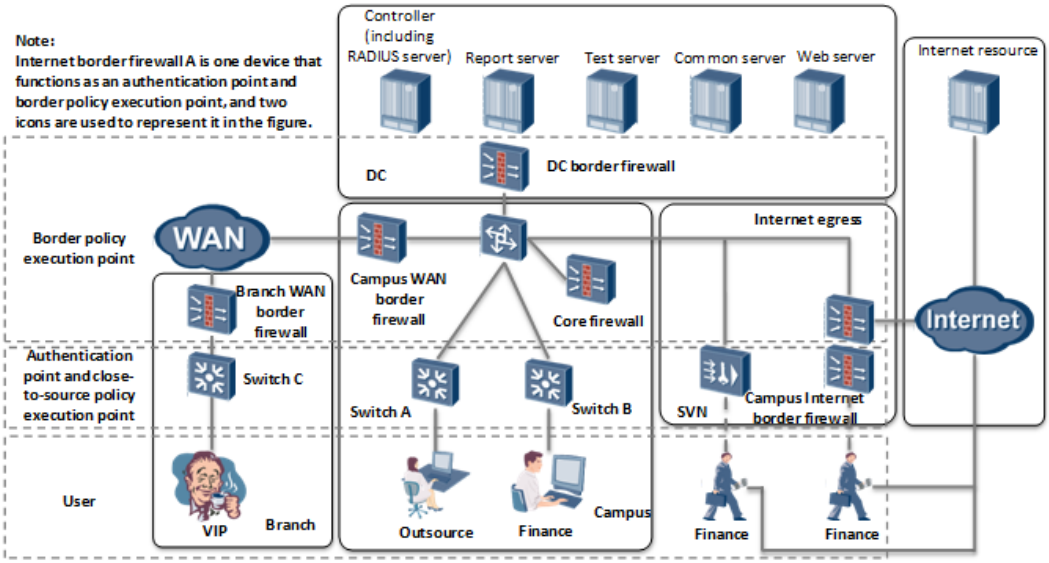

In the Free Mobility solution, the points where access permission policies are executed can be authentication points and network borders. The Agile Controller (including the RADIUS

server) are deployed in the data center. The Agile Controller shall be connected to the policy execution points.

3.2 Distributed Access Permission Control

Figure 3-2 Networking for distributed access permission control

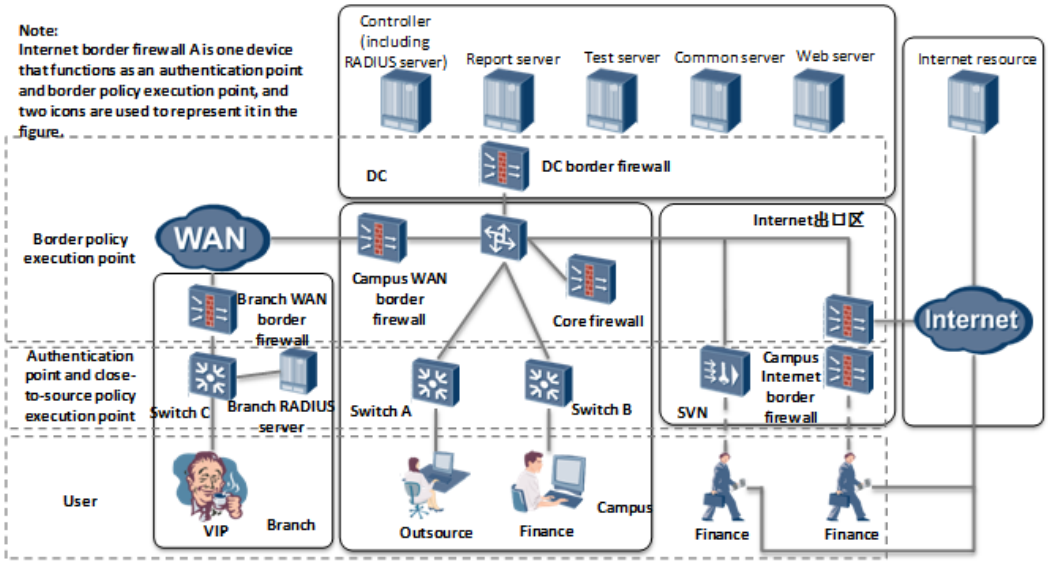

In the Free Mobility solution, the points where access permission policies are executed can be authentication points and network borders. The Agile Controller and campus RADIUS server are deployed in the data center, and other RADIUS servers can be deployed at branches. The Agile Controller shall be connected to the policy execution policies, and to the branch RADIUS servers.

3.3 Network Experience Insurance

Figure 3-3 Networking for network experience insurance

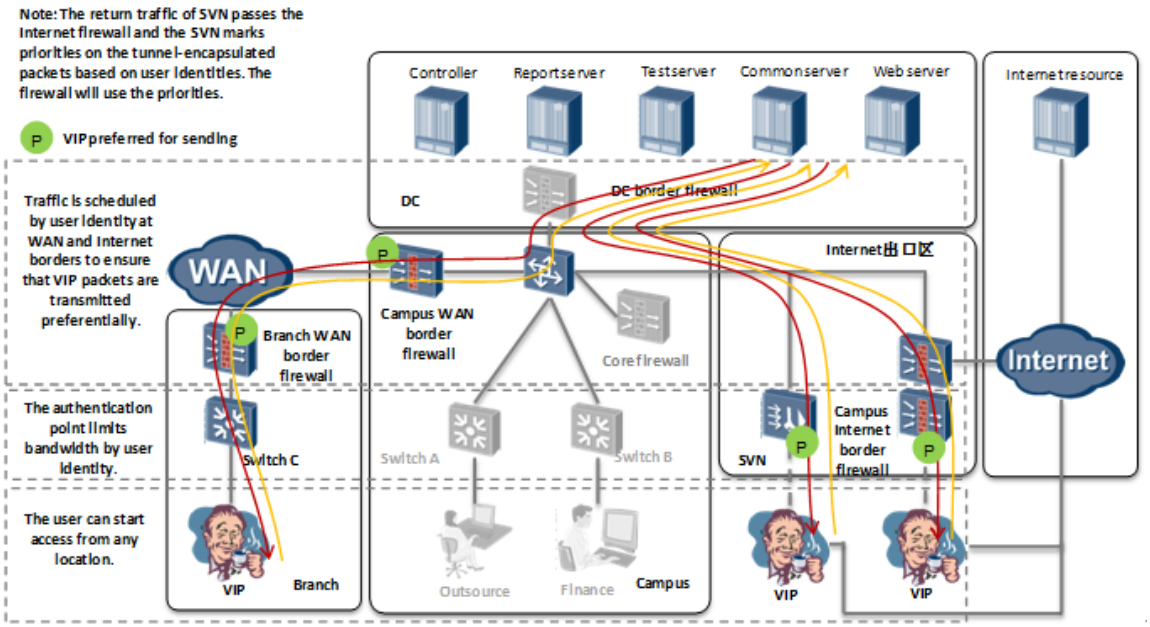

In the Free Mobility solution, the points where network experience insurance policies are executed can be network borders. The Agile Controller can be deployed in either central or distributed mode.

4 References