FortiGate Infrastructure Study Guide

for FortiOS 7.2

FORTINET®
Training Institute

Fortinet Training

https://training.fortinet.com

Fortinet Document Library

https://docs.fortinet.com

Fortinet Knowledge Base

https://kb.fortinet.com

Fortinet Fuse User Community

https://fusecommunity.fortinet.com/home

Fortinet Forums

https://forum.fortinet.com

Fortinet Support

https://support.fortinet.com

FortiGuard Labs

https://www.fortiguard.com

Fortinet Network Security Expert Program (NSE)

https://training.fortinet.com/local/staticpage/view.php?page=certifications

Fortinet | Pearson VUE

https://home.pearsonvue.com/fortinet

Feedback

Email: askcourseware@fortinet.com

6/13/2022

TABLE OF CONTENTS

In this lesson, you will learn about the routing capabilities and features available on FortiGate.

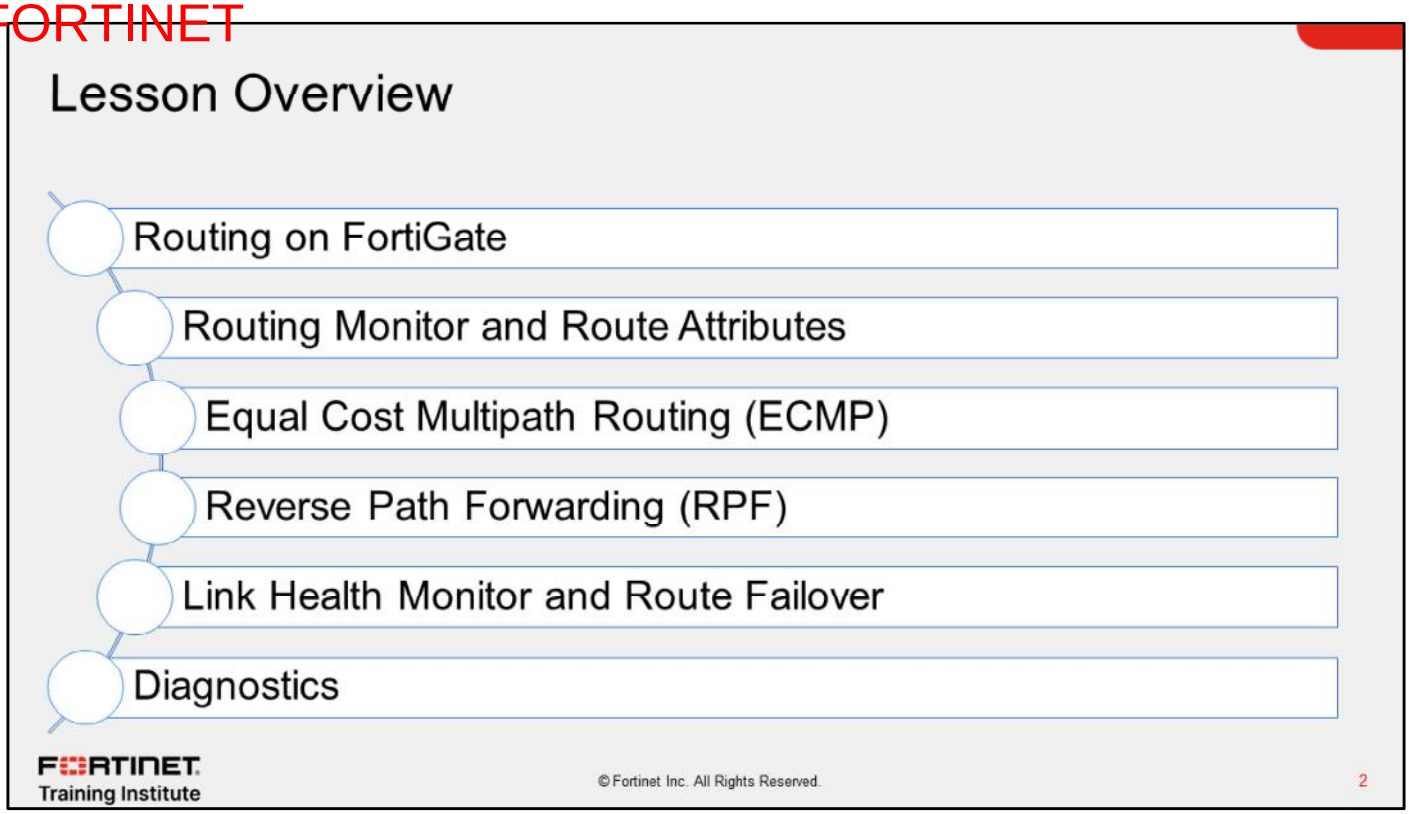

In this lesson, you will learn about the topics shown on this slide.

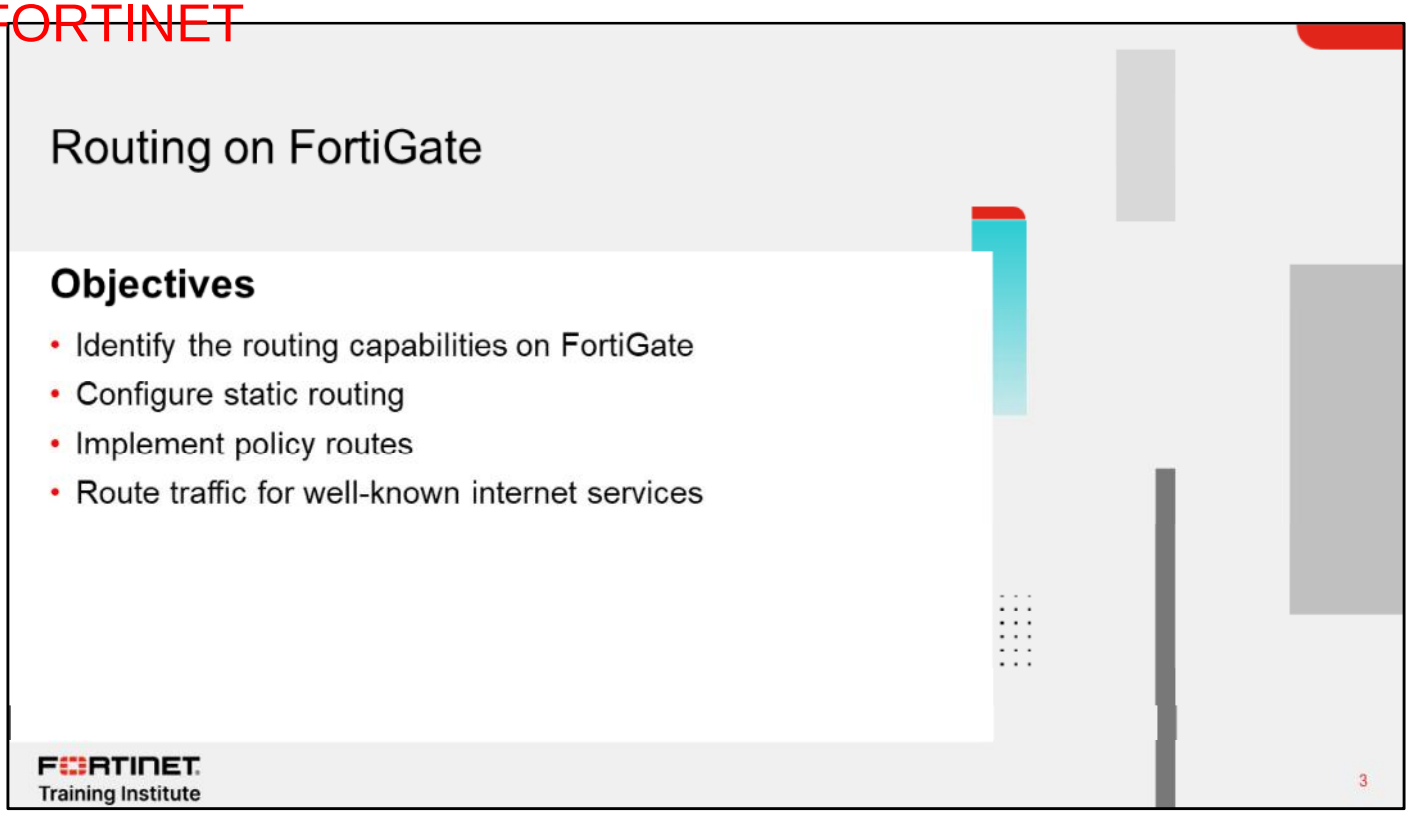

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in routing on FortiGate, you should be able to implement static and policy routing. You will also be able to route traffic for well-known internet services.

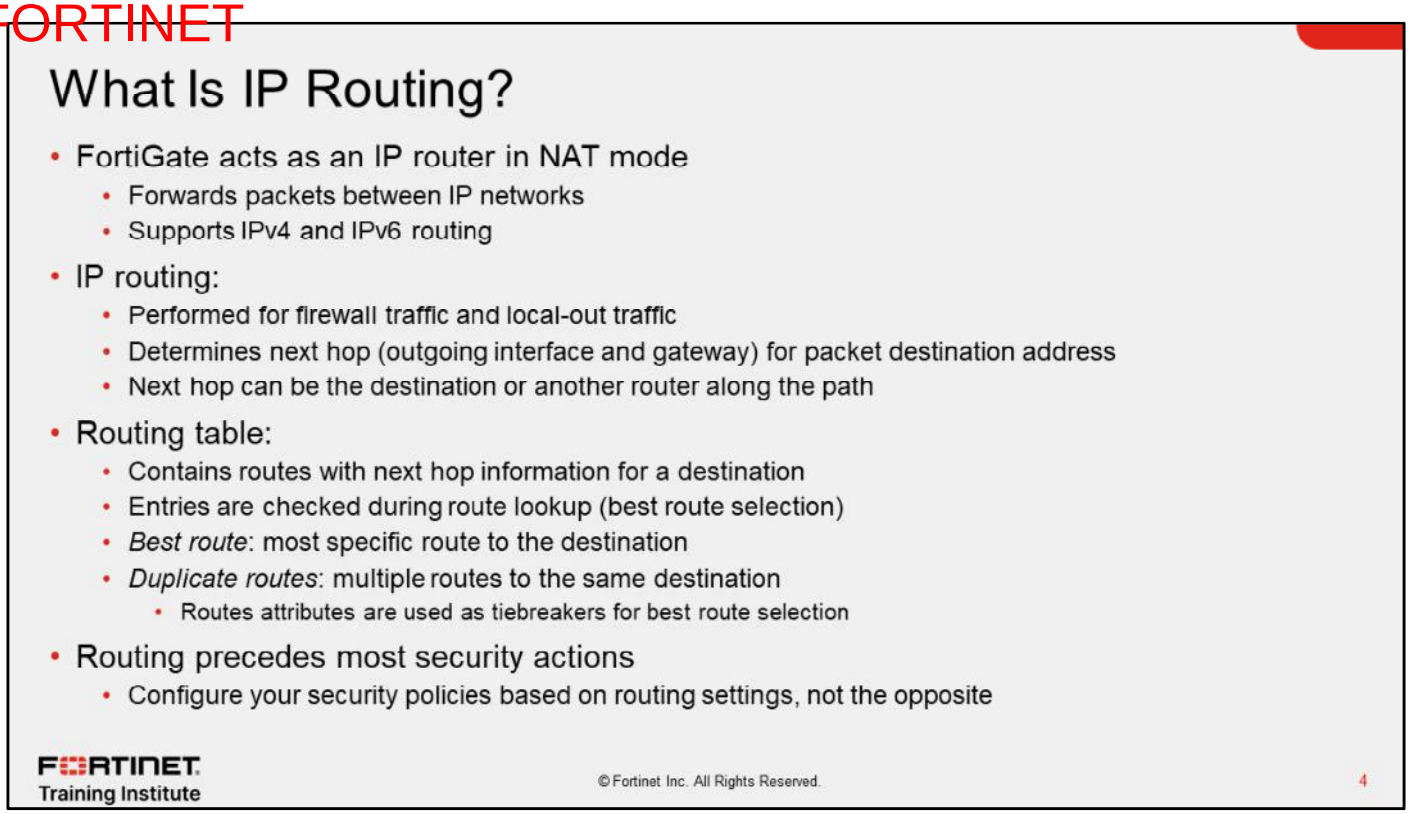

When FortiGate operates in NAT mode—the default operation mode—FortiGate behaves as an IP router. An IP router is a device that forwards packets between IP networks. For that, a router performs IP routing, which is the process of determining the next hop to forward a packet to based on the packet destination IP address. FortiGate supports both IPv4 and IPv6 routing.

FortiGate performs routing for both firewall traffic (also known as user traffic) and local-out traffic. Firewall traffic is the traffic that travels through FortiGate. Local-out traffic is the traffic generated by FortiGate, usually for management purposes. For example, when you ping a device from FortiGate, that's local-out traffic. When FortiGate connects to FortiGuard to download the latest definitions, that's also local-out traffic.

Routers maintain a routing table. A routing table contains a series of entries, also known as routes. Each route in the routing table indicates the *next hop* for a particular destination. The next hop refers to the outgoing interface and gateway to use for forwarding the packet. The next hop can be the destination of the packet or another router along the path to the destination. If the next hop isn't the destination, the next router in the path routes the packet to the next hop. The routing process is repeated on each router along the path until the packet reaches its destination.

To route packets, FortiGate performs a route lookup to identify the best route to the destination. The best route is the most specific route to the destination. If FortiGate finds duplicate routes—that is, multiple routes to the same destination—it uses various route attributes as a tiebreak to determine the best route.

Routing takes place before most security features. For example, routing precedes firewall policy evaluation, content inspection, traffic shaping, and source NAT (SNAT). This means that the security actions that FortiGate performs depend on the outgoing interface determined by the routing process. This also means that your security policy configuration must follow your routing configuration, and not the opposite.

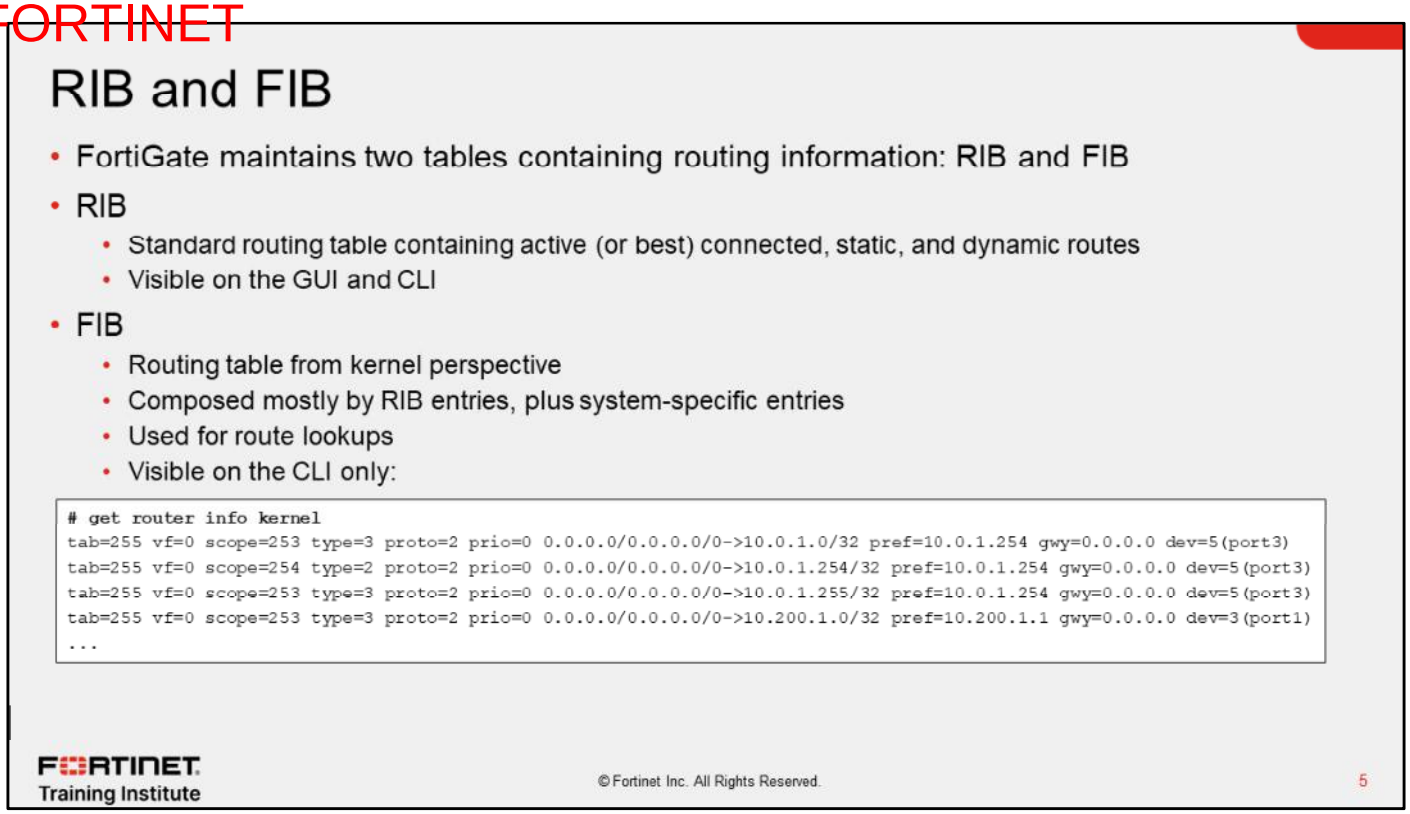

FortiGate maintains its routing information in two tables: RIB and FIB. The routing table, also known as the routing information base (RIB), is a standard routing table containing active (or the best) connected, static, and dynamic routes. The forwarding information base (FIB) can be described as the routing table from the kernel point of view, and is built mostly out of RIB entries plus some system-specific entries required by FortiOS.

When FortiGate performs a route lookup, it checks the FIB and not the RIB. However, because the FIB is composed mostly by RIB entries, then the route lookup mainly involves checking routes from the RIB. For this reason, the route lookup is often referred to as the routing table lookup process. Nonetheless, a more accurate statement is to refer to it as the FIB lookup process.

You can display the RIB entries on the FortiGate GUI and CLI. However, for the FIB, you can display its entries on the FortiGate CLI only. The output on this slide shows the CLI command that displays the FIB. Note that the output has been cut to fit the slide. You will learn how to display the routing table entries in this lesson.

This lesson focuses on the RIB (or routing table) only, and you will learn more about it, including how to monitor its entries, in this lesson.

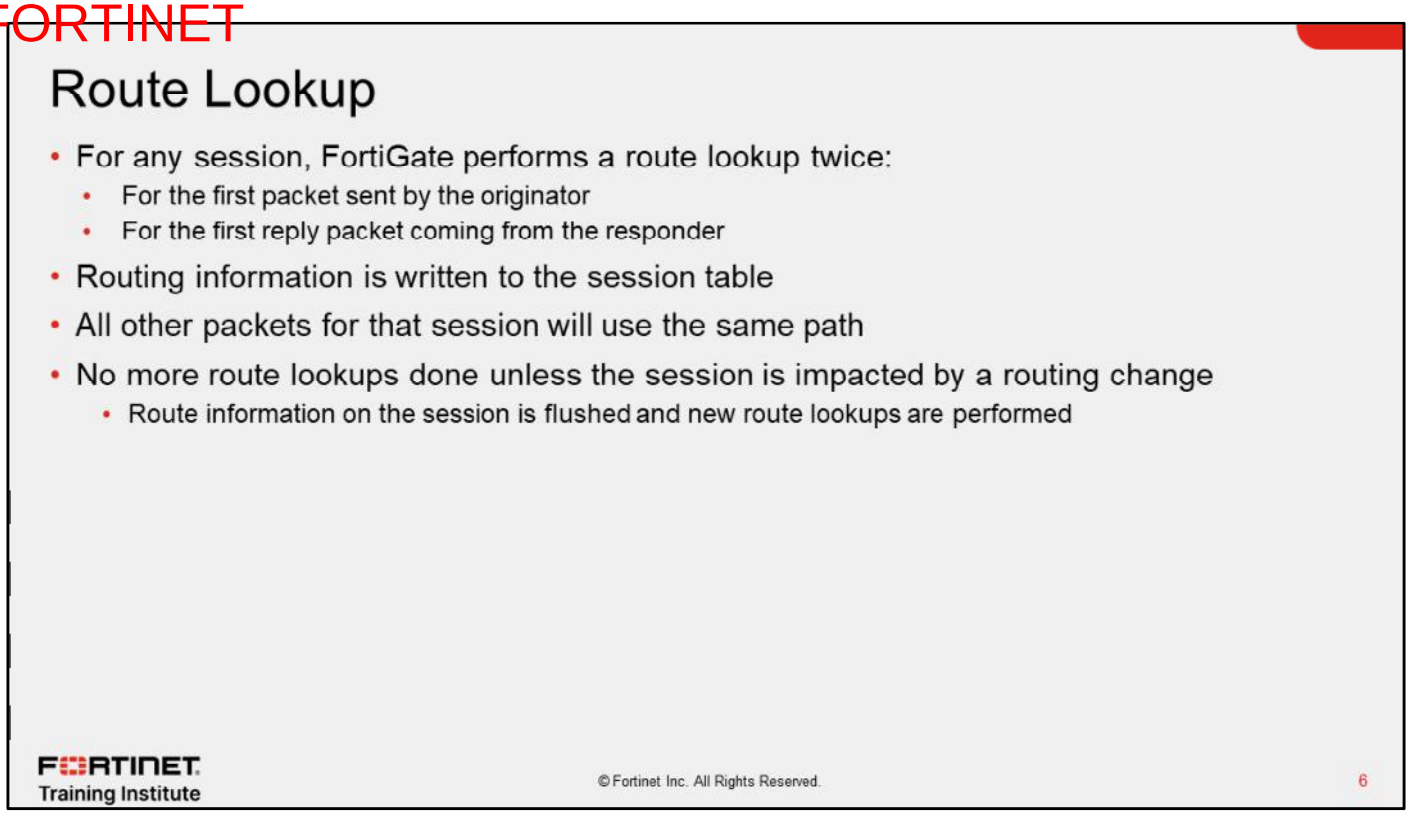

For each session, FortiGate performs two route lookups:

- For the first packet sent by the originator
- For the first reply packet coming from the responder

After completing these two lookups, FortiGate writes the routing information to its session table. Subsequent packets are routed according to the *session table*, not the routing table. So, all packets that belong to the same session follow the same path. However, there is an exception to this rule: if there is a change in the routing table that impacts the session, then FortiGate removes the route information for the session table, and then performs additional route lookups to rebuild this information.

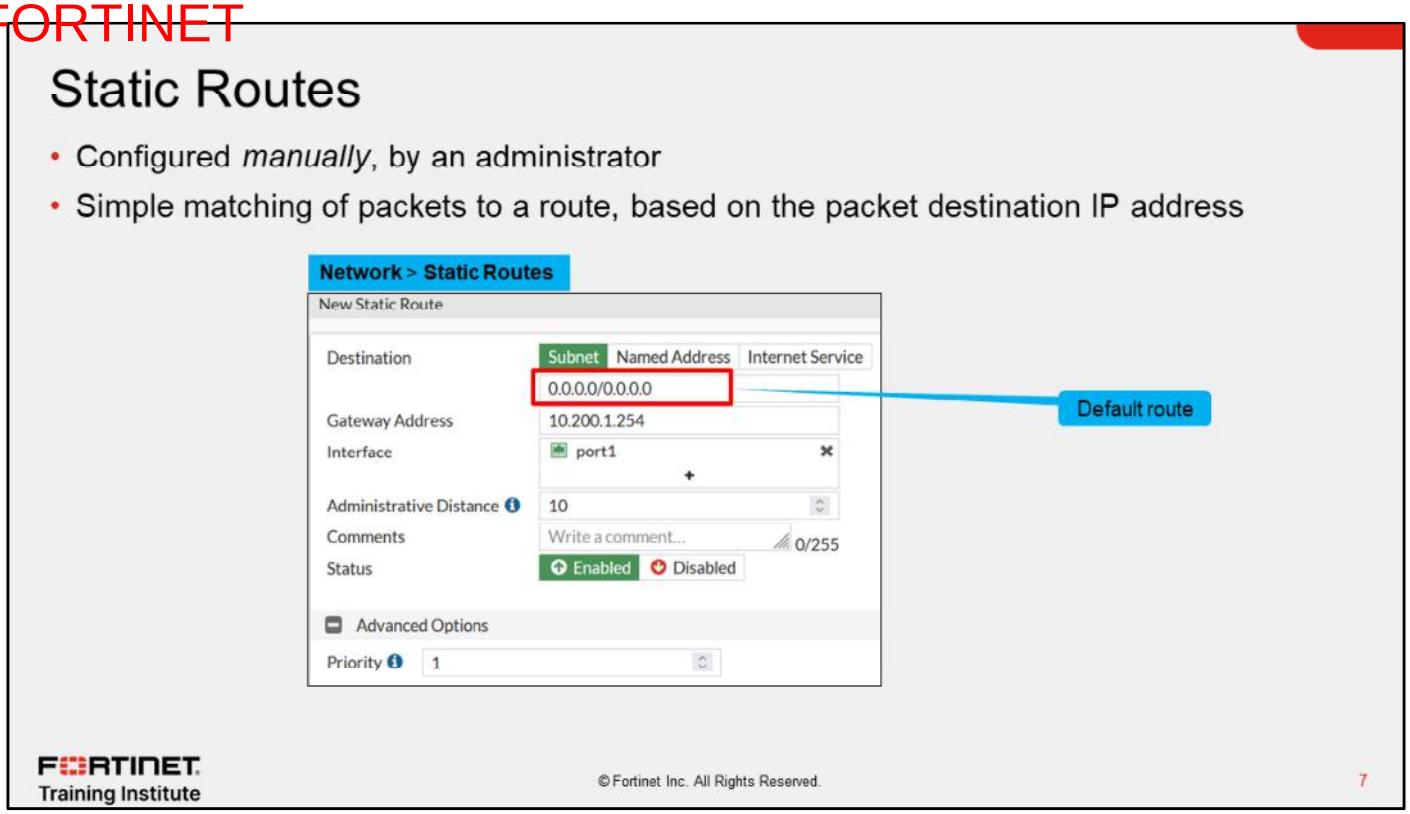

One type of manually configured route is called a static route. When you configure a static route, you are telling FortiGate, “When you see a packet whose destination is within a specific range, send it through a specific network interface, towards a specific router.” You can also configure the distance and priority so that FortiGate can identify the best route to any destination matching multiple routes. You will learn about distance and priority in this lesson.

For example, in simple home networks, DHCP automatically retrieves and configures a route. Your modem then sends all outgoing traffic through your ISP internet router, which can relay packets to their destination. This is typically referred to as a default route, because all traffic not matching any other routes will, by default, be routed using this route. The example shown on this slide is a default route. The destination subnet value of `0.0.0.0/0.0.0.0` matches all addresses within any subnet. Most FortiGate devices deployed at the edge of the network have at least one of these default routes to ensure internet traffic is forwarded to the ISP network.

Static routes are not needed for subnets to which FortiGate has direct Layer 2 connectivity.

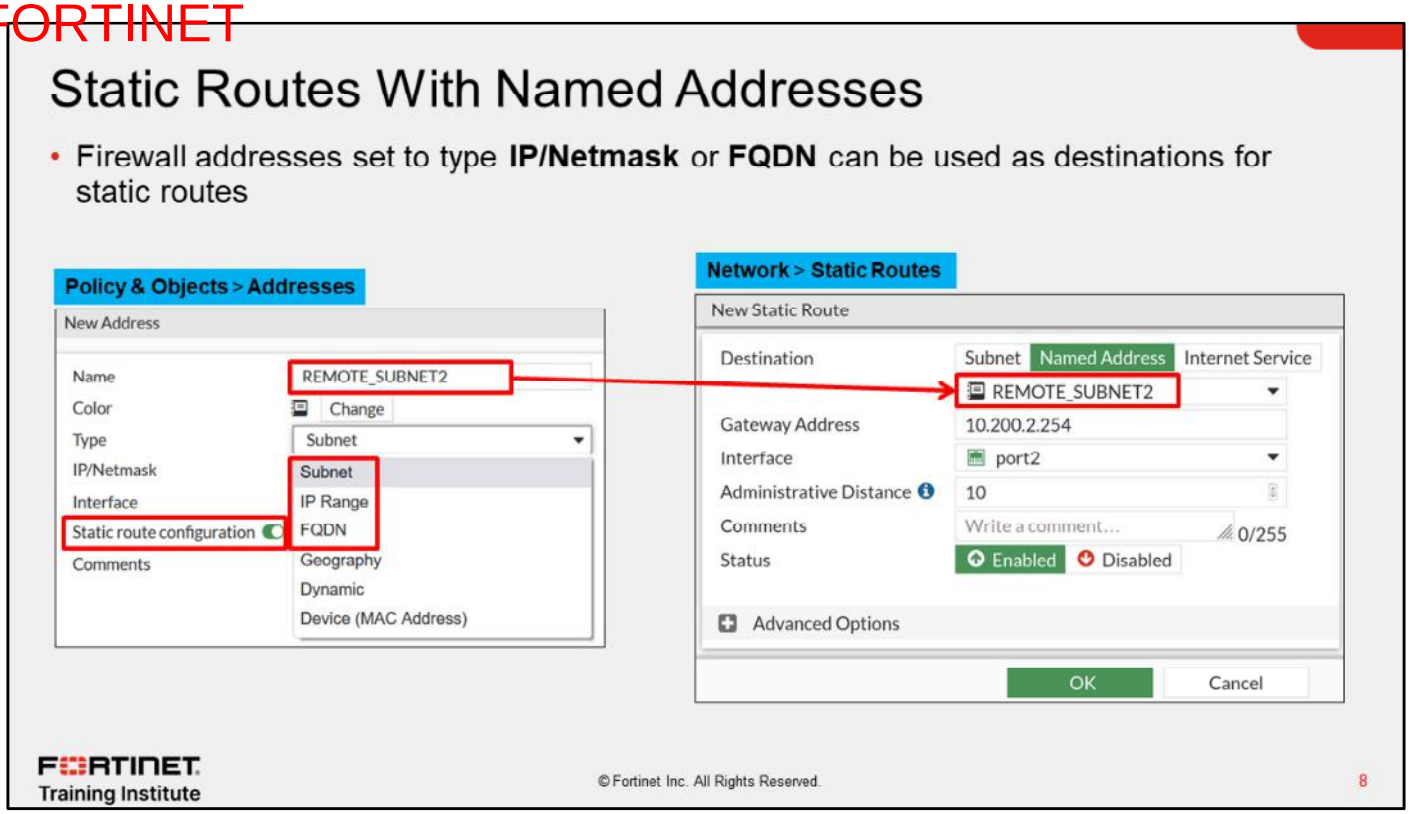

If you create a firewall address object with the type **IP/Netmask** or **FQDN**, you can use that firewall address as the destination of one or more static routes. First, enable **Static route configuration** in the firewall address configuration. After you enable it, the firewall address object becomes available for use in the **Destination** drop-down list for static routes with named addresses.

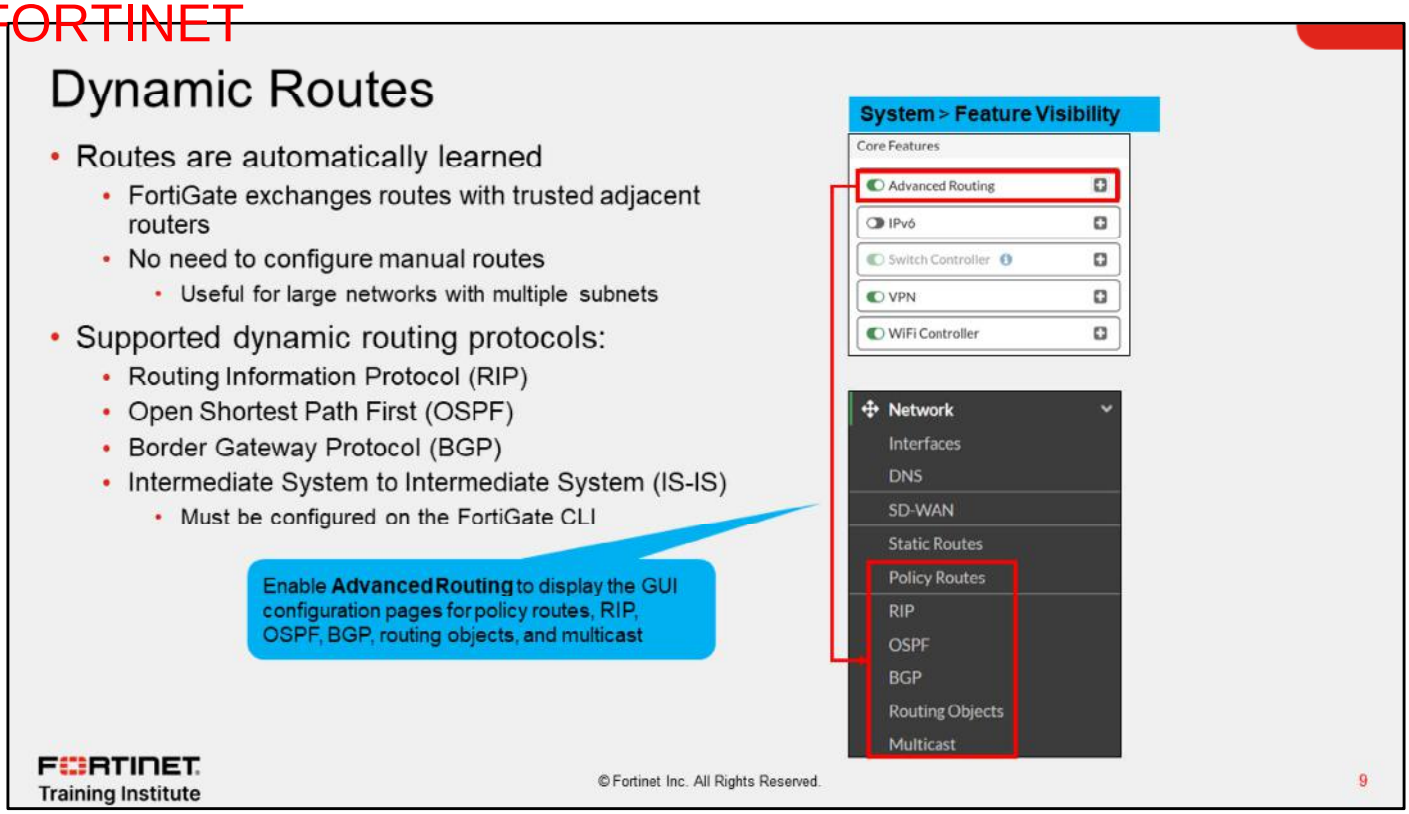

For large networks, manually configuring hundreds of static routes may not be practical. Your FortiGate can help, by learning routes automatically. FortiGate supports several dynamic routing protocols: RIP, OSPF, BGP, and IS-IS.

In dynamic routing, FortiGate communicates with trusted adjacent routers to exchange routing information about their known networks. Then, FortiGate adds the learned routes into its local routing table and considers them during the route lookup process.

You can configure dynamic routing for RIP, OSPF, and BGP protocols using the FortiGate GUI. You just need to make sure that the **Advanced Routing** option in the **Feature Visibility** page is enabled—it's enabled by default. However, for configuring IS-IS, you must use the FortiGate CLI.

Note that when you enable **Advanced Routing** on the **Feature Visibility** page, you also enable the configuration pages for other advanced routing features such as **Policy Routes**, **Routing Objects**, and **Multicast**. You will learn more about policy routes in this lesson.

Larger networks also may need to balance the routing load among multiple valid paths and detect and avoid routers that are down. You will learn more about that in this lesson.

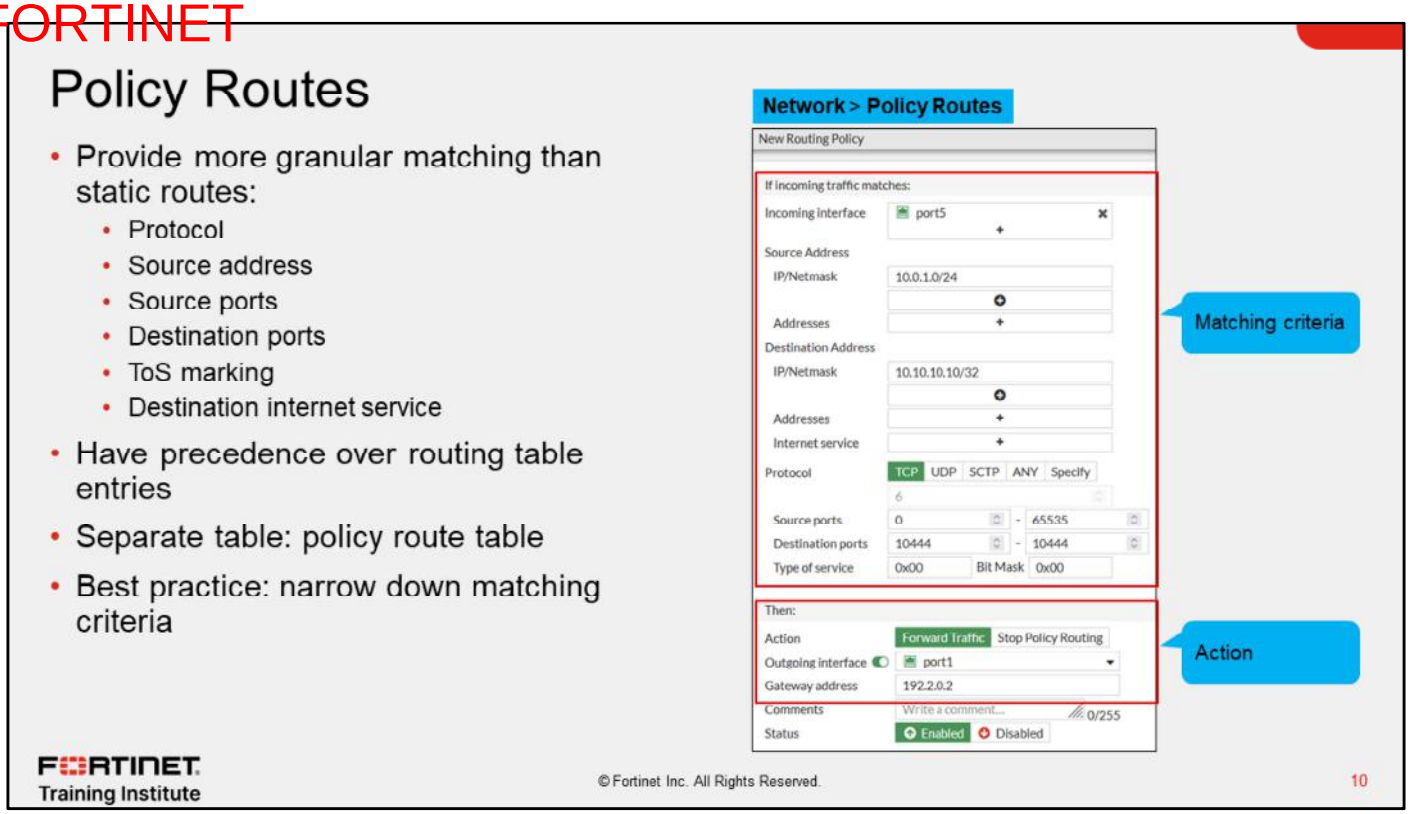

Static routes are simple and are often used in small networks. Policy routes, however, are more flexible because they can match more than just the destination IP address. For example, you can configure as matching criteria the incoming interface, the source and destination subnets, protocol, and port number.

Policy routes are maintained in a separate routing table by FortiGate and have precedence over the entries in the routing table. Because of its precedence, it is a best practice to narrow down the matching criteria of policy routes as much as possible. Otherwise, traffic that is expected to be routed using standard routing, that is, based on the destination address only and the routing table entries, could be handled by policy routes instead.

This slide shows an example of a policy route configured using the FortiGate GUI. The policy route instructs FortiGate to match traffic received at **port5**, sourced from `10.0.1.0/24` and destined to the host `10.10.10.10`. The traffic must also be destined to TCP port 10444 for the policy route to match. FortiGate then forwards the traffic—**Forward Traffic** action—to **port1** through the gateway `192.2.0.2`.

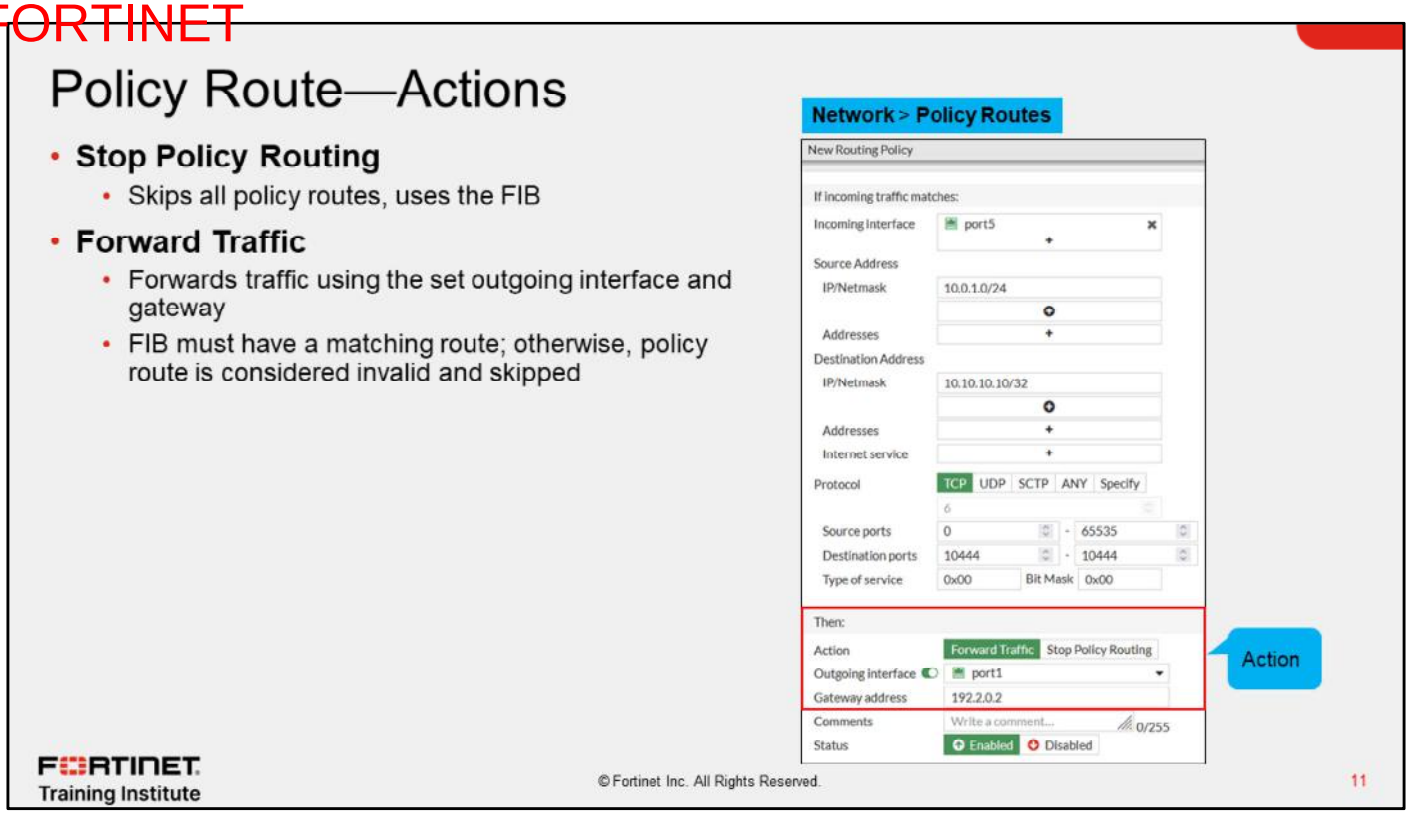

When a packet matches a policy route, FortiGate takes one of two actions. Either it routes the packet to the configured outgoing interface and gateway—**Forward Traffic** action—or it stops checking the policy routes—**Stop Policy Routing** action—so the packet is routed based on the routing table.

Note that when you configure **Forward Traffic** as the action, the **Destination Address**, **Outgoing interface**, and the **Gateway address** settings must match a route in the FIB. Otherwise, the policy route is considered invalid and, as a result, skipped.

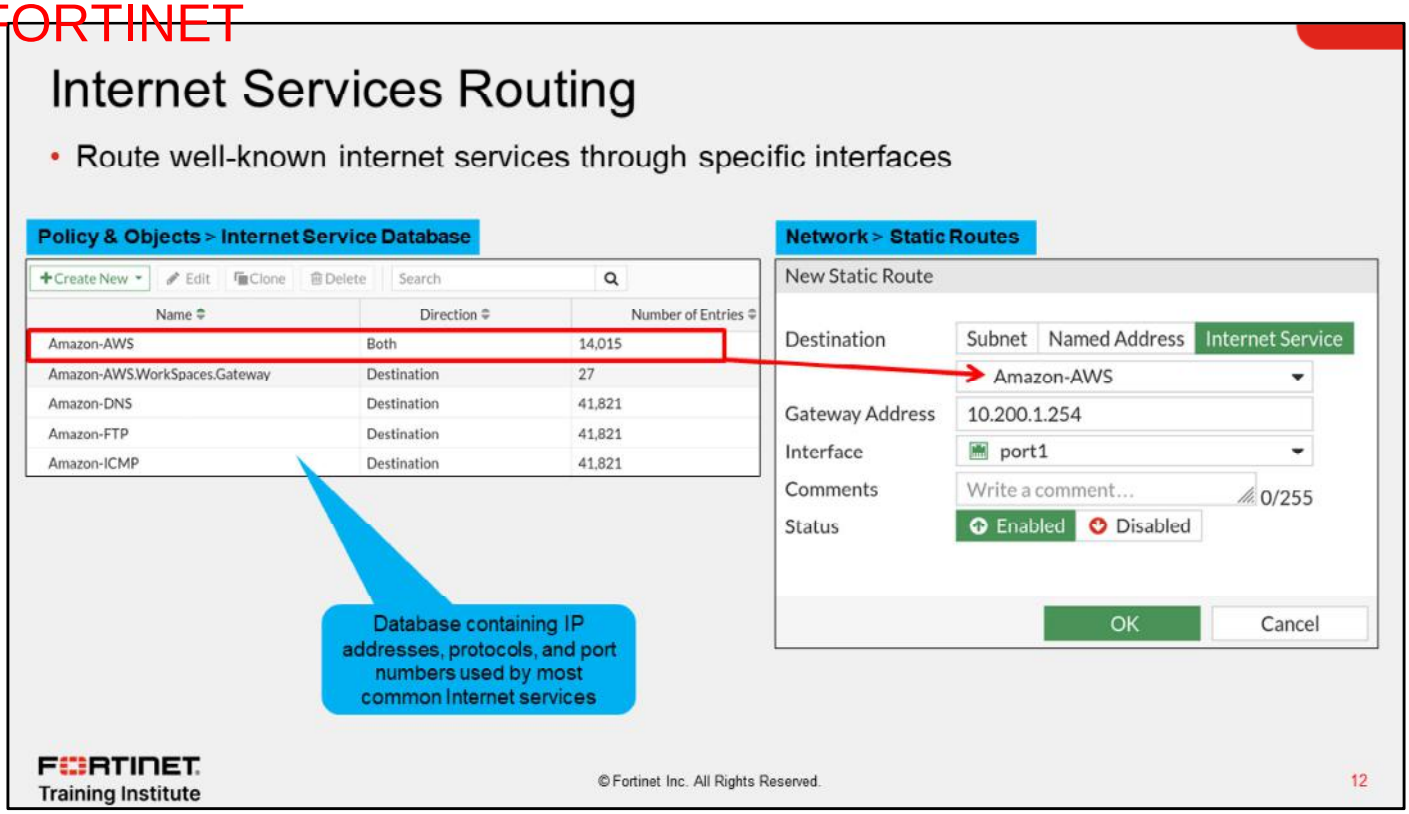

What happens if you need to route traffic to a public internet service (such as Amazon-AWS or Apple Store) through a specific WAN link? Say you have two ISPs and you want to route Netflix traffic through one ISP and all your other internet traffic though the other ISP. To achieve this goal, you need to know the Netflix IP addresses and configure the static route. After that, you must frequently check that none of the IP addresses have changed. The internet service database (ISDB) helps make this type of routing easier and simpler. ISDB entries are applied to static routes to selectively route traffic though specific WAN interfaces.

Even though they are configured as static routes, ISDB routes are actually policy routes and take precedence over any other routes in the routing table. As such, ISDB routes are added to the policy routing table.

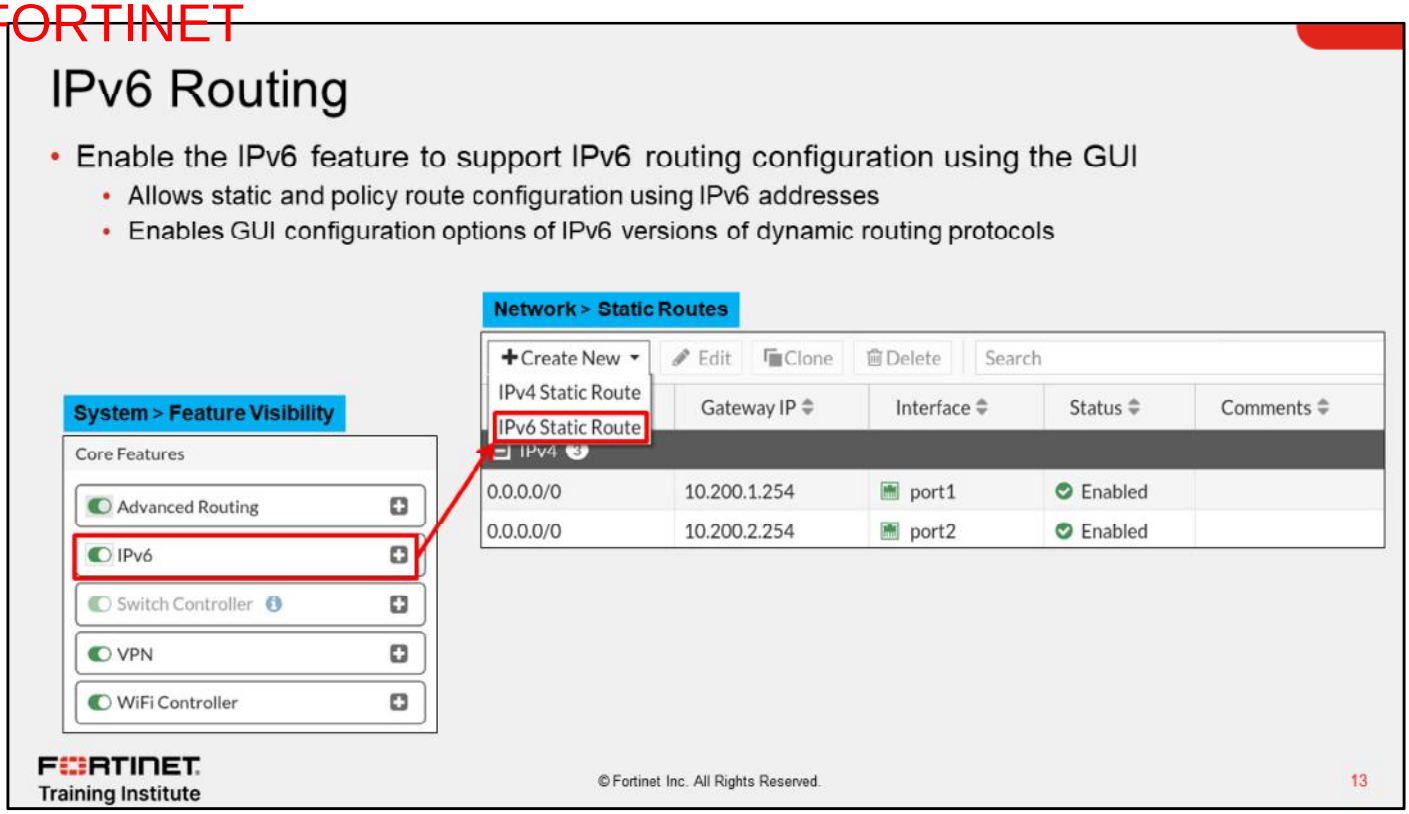

To enable routing configuration for IPv6 addresses using the GUI, you must enable **IPv6** in the **Feature Visibility** menu. Then, you can create static routes and policy routes with IPv6 addresses. Enabling the IPv6 feature also enables GUI configuration options for IPv6 versions of the dynamic routing protocols.

Knowledge Check

1. Which objects can you use to create static routes?
   - ✓ A. ISDB objects
   - B. Service objects

2. When the **Stop policy routing** action is used in a policy route, which behavior is expected?
   - A. FortiGate skips over this policy route and tries to match another in the list.
   - ✓ B. FortiGate routes the traffic based on the regular routing table.

FORTINET Training Institute

14

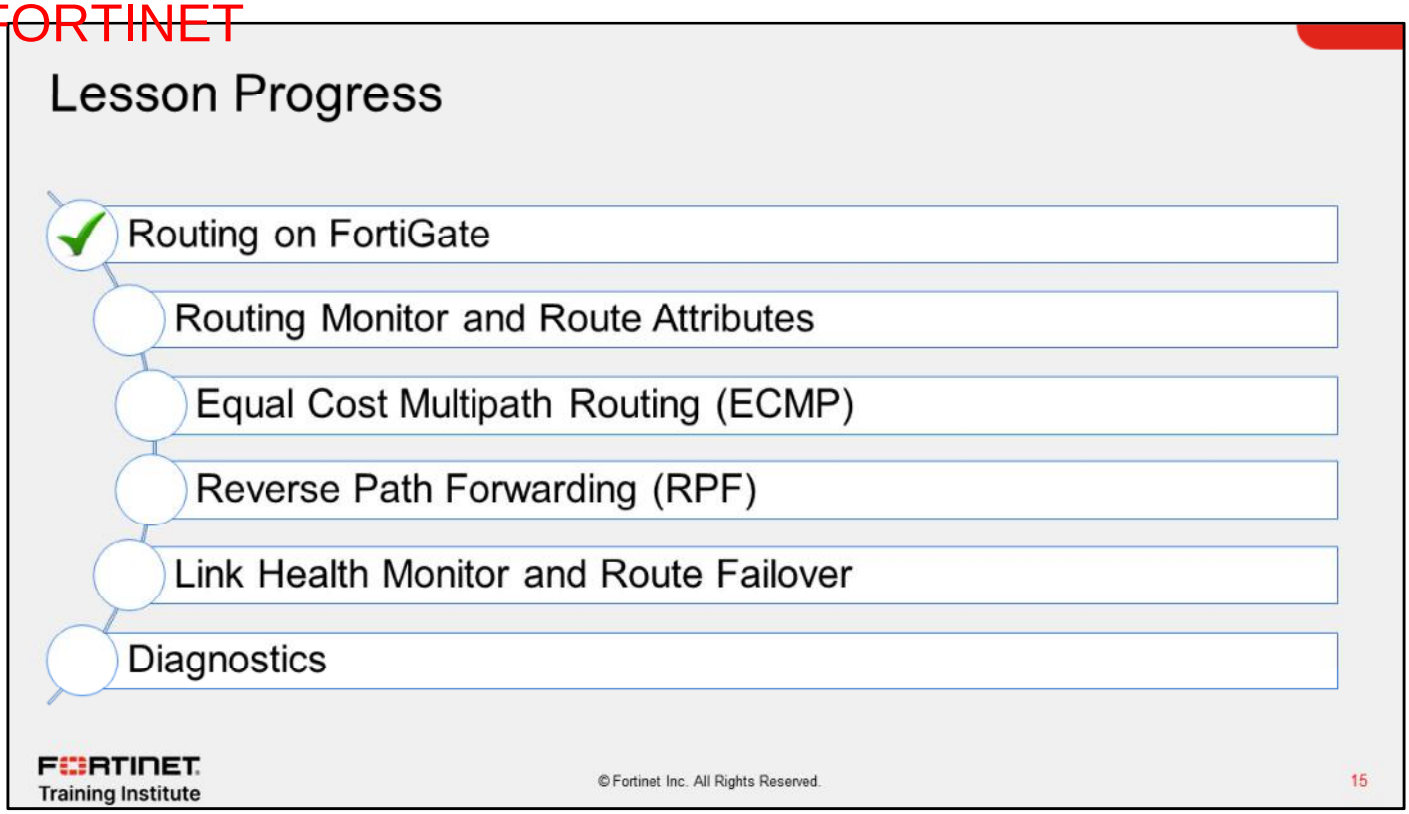

Good job! You now understand routing on FortiGate.

Now, you will learn about routing monitor and route attributes.

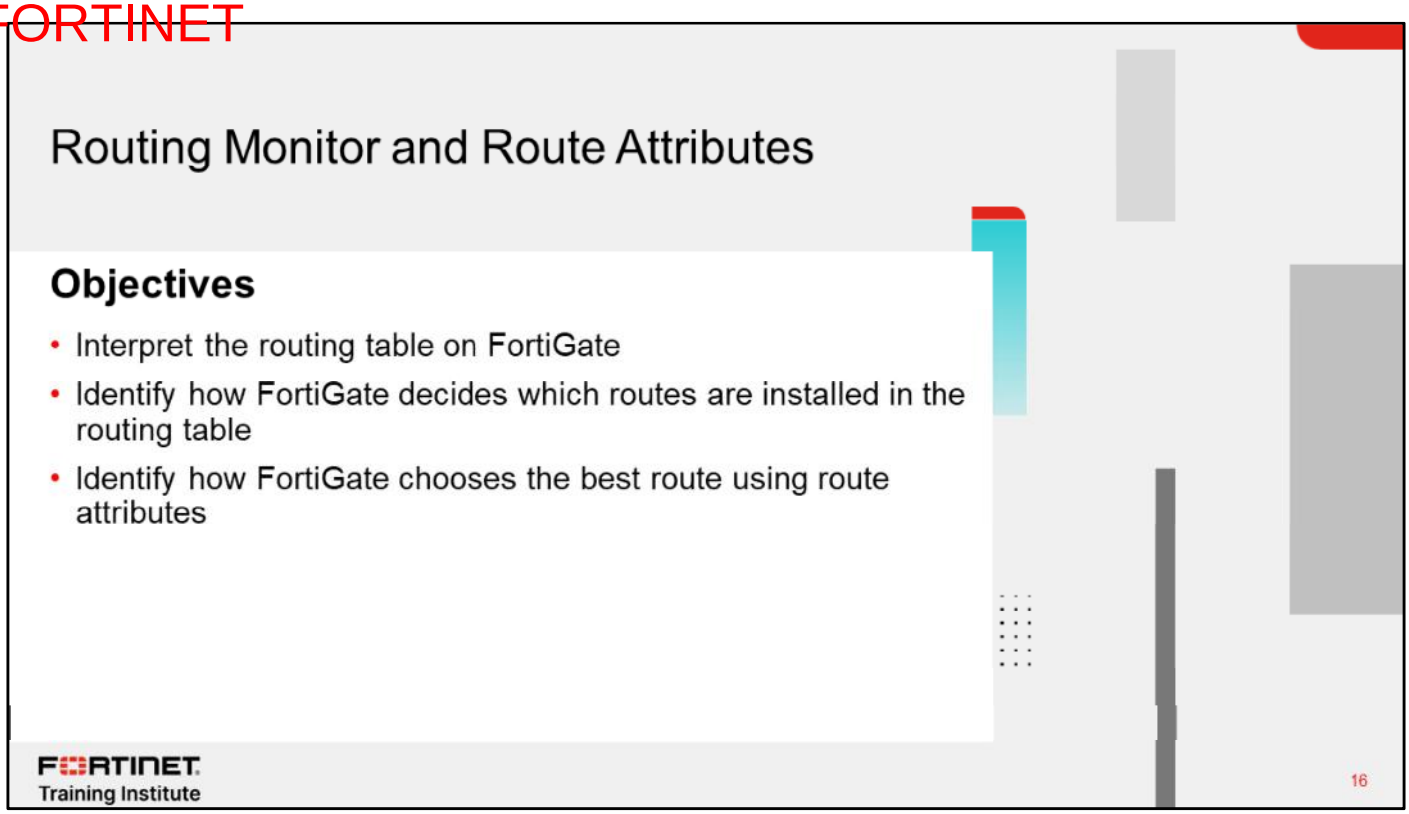

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding the routing monitor and route attributes, you should be able to interpret the routing table, identify which routes are installed in the routing table, and identify how FortiGate chooses the best route using route attributes.

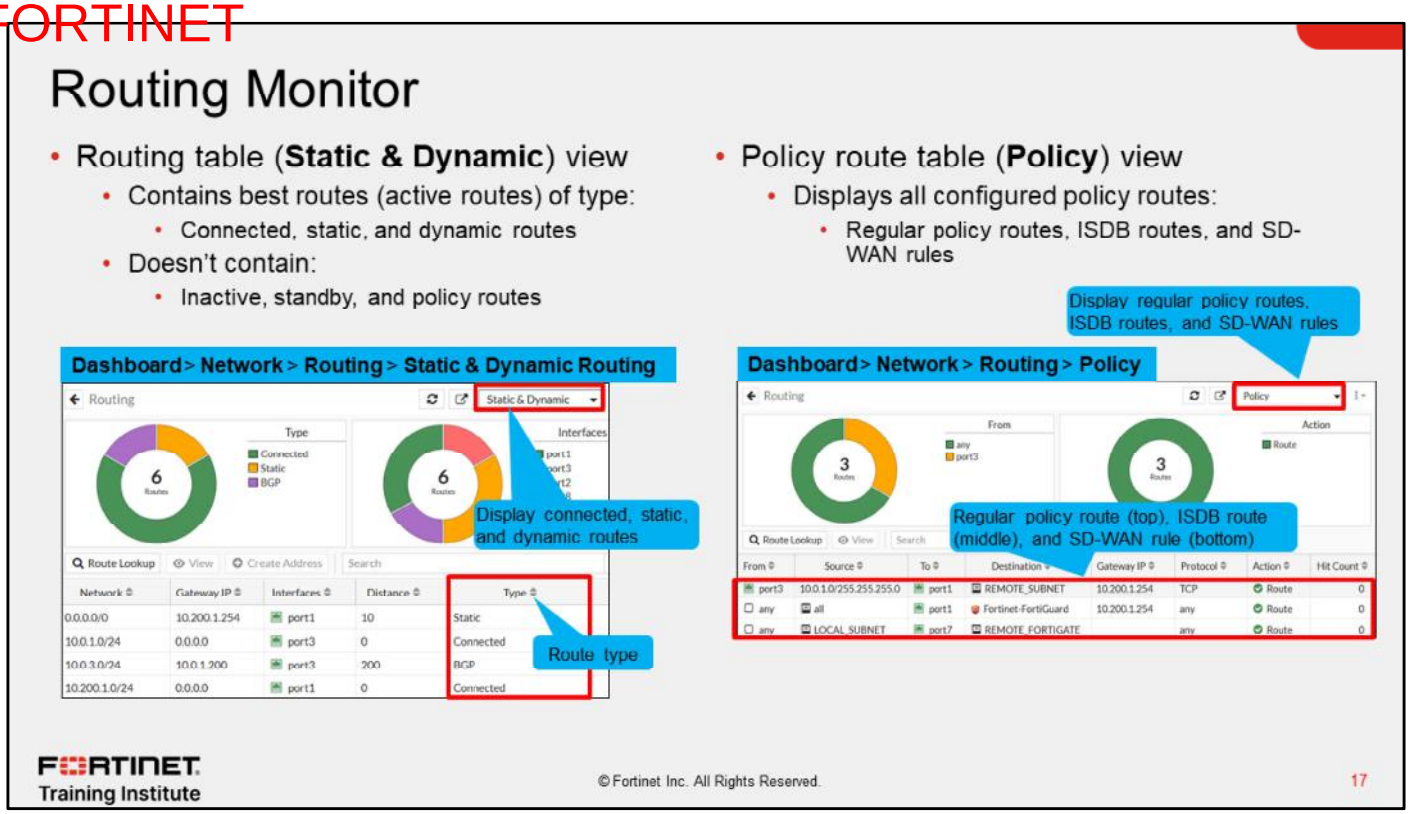

The routing monitor widget on the dashboard page enables you to view the routing table and policy route table entries. The routing table contains *the best routes* (or active routes) of the following type:

- Static: manual routes that are configured by the administrator.
- Connected: automatic routes added by FortiOS after an interface is assigned an IP address. A connected route references the interface IP address subnet.
- Dynamic: routes learned using a dynamic routing protocol such as BGP or OSPF. FortiGate installs these routes automatically in the routing table and indicates the dynamic routing protocol used.

To view the routing table entries, select **Static & Dynamic**, as shown on this slide. However, keep in mind that the routing table doesn't contain the following routes:

- Inactive routes: static and connected routes whose interfaces are administratively down or whose links are down. Static routes are also marked inactive when their gateway is detected as dead by the link health monitor.

- Standby routes: These are active routes that are removed from the routing table because they are duplicate and have higher distances. For instance:
  - A second static default route with a higher distance than another static default route.
  - A dynamic route such as BGP or OSPF, to the same destination as another static route. However, the dynamic route is not displayed in the routing table because the static route has a lower distance.

- Policy routes: These include regular policy routes, ISDB routes, and SD-WAN rules. Policy routes are viewed in a separate table—the policy route table. To view the policy route table entries, select **Policy**.

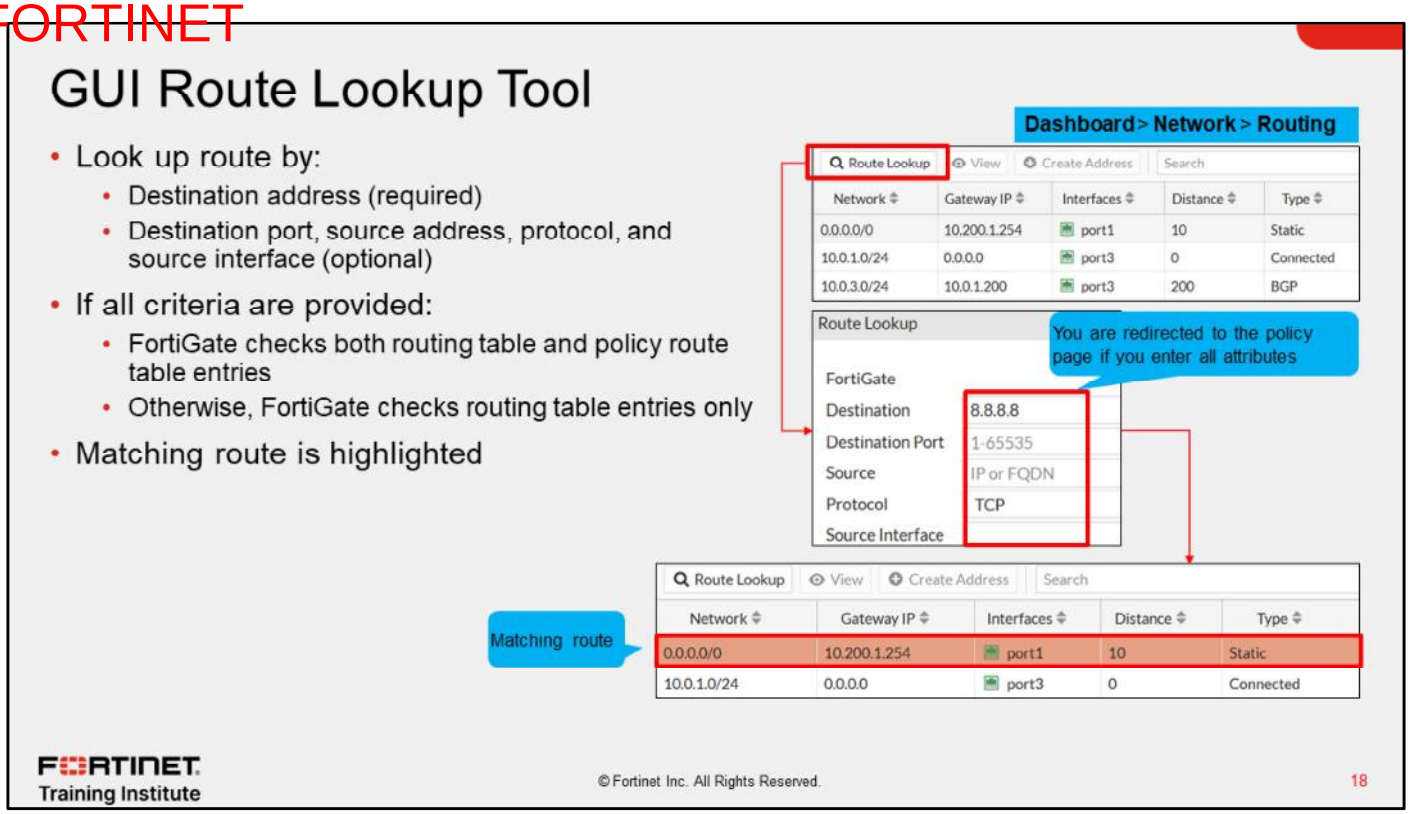

You can perform a route lookup on the routing monitor widget by clicking **Route Lookup**. Then, you must indicate at least the destination address to look up for, and optionally, the destination port, source address, protocol, and source interface.

The way the route lookup works is as follows:

- If you don't provide all lookup criteria, FortiGate considers only the routing table entries. FortiGate then highlights the matching route, if any.
- If you provide all lookup criteria, FortiGate considers both routing table and policy table entries. If the lookup matches a policy route, the GUI redirects you to the policy route page, and then highlights the corresponding matching policy route.

The example on this slide shows a route lookup tool for `8.8.8.8` and TCP as destination address and protocol, respectively. Because the administrator doesn't provide all criteria, FortiGate considers the routing table entries only. Then, the route lookup highlights the static default route as the matching route.

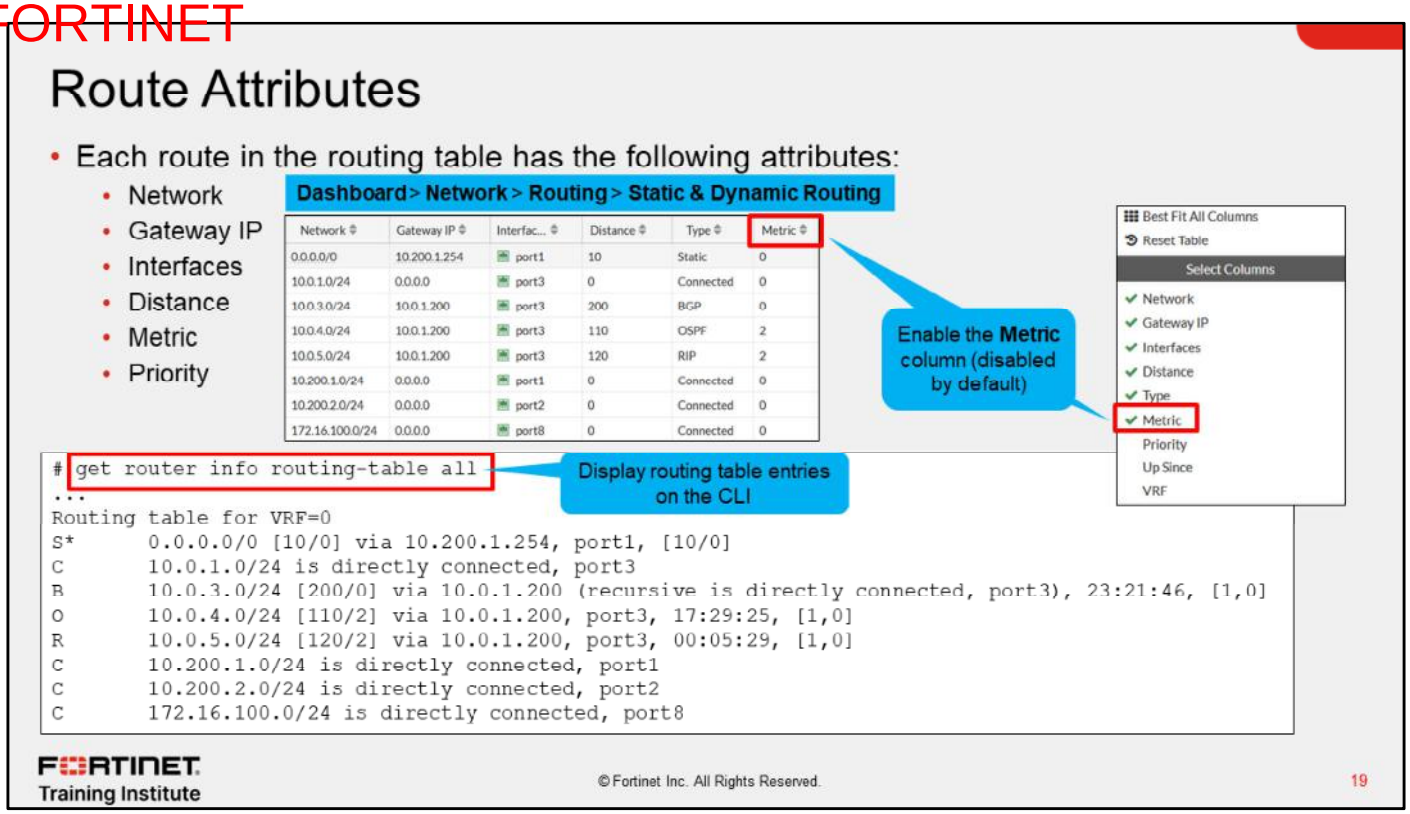

Each of the routes listed in the routing table includes several attributes with associated values.

The **Network** column lists the destination IP address and subnet mask to match. The **Interfaces** column lists the interface to use to deliver the packet.

The **Distance**, **Metric**, and **Priority** attributes are used by FortiGate to make various route selection decisions. You will learn about each of these in this lesson.

This slide also shows the command you can run to display the routing table on the FortiGate CLI. The `get router info routing-table all` command displays the same route entries as the routing monitor widget on the FortiGate GUI.

Distance

- First tiebreaker for duplicate routes (best route selection)
  - The lower the distance, the higher the preference
  - Set by the administrator (except connected routes)
- Best route selection:
  - Route with lowest distance is installed in the RIB
  - Standby routes (higher distance) are not installed in the RIB
    - They are installed in the routing table database
  - Multiple equal-distance duplicate routes but different protocol:
    - FortiGate keeps the route that was learned last (*avoid*)
- Default distance per route type:

Dashboard > Network > Routing > Static & Dynamic

Network	Gateway IP	Interfac...	Distance	Type	Metric
0.0.0.0/0	10.200.1.254	port1	10	Static	0
10.0.1.0/24	0.0.0.0	port3	0	Connected	0
10.0.3.0/24	10.0.1.200	port3	200	BGP	0
10.0.4.0/24	10.0.1.200	port3	110	OSPF	2
10.0.5.0/24	10.0.1.200	port3	120	RIP	2
10.200.1.0/24	0.0.0.0	port1	0	Connected	0
10.200.2.0/24	0.0.0.0	port2	0	Connected	0
172.16.100.0/24	0.0.0.0	port8	0	Connected	0

Connected*	Static (SD-WAN zone)	Static (DHCP)	Static (Manual)	Static (IKE)	EBGP	OSPF	IS-IS*	RIP	IBGP
0	1	5	10	15	20	110	115	120	200

* Hardcoded

FORTINET Training Institute

20

Distance, or administrative distance, is the first tiebreaker that routers use to determine the best route for a particular destination. If there are two or more routes to the same destination (duplicate routes), the lowest-distance route is considered the best route and, as a result, is installed in the routing table. Other lower-distance routes to the same destination are standby routes and, as a result, are not installed in the routing table. Instead, they are installed in the routing table database. You will learn more about the routing table database in this lesson.

You can set the distance for all route types except connected and IS-IS routes. This slide shows the default values per type of route.

In case FortiGate learns two equal-distance routes to the same destination but that are sourced from different protocols, then FortiGate installs in the routing table the route that was learned *last*. For example, if you set the distance of BGP routes to 110, and there is another OSPF route to the same destination using the default administrative distance (110), then FortiGate keeps whichever route was learned last in the routing table. Because this behavior can lead to different results based on the timing of events, then it's not recommended to configure different-protocol routes with the same distance.

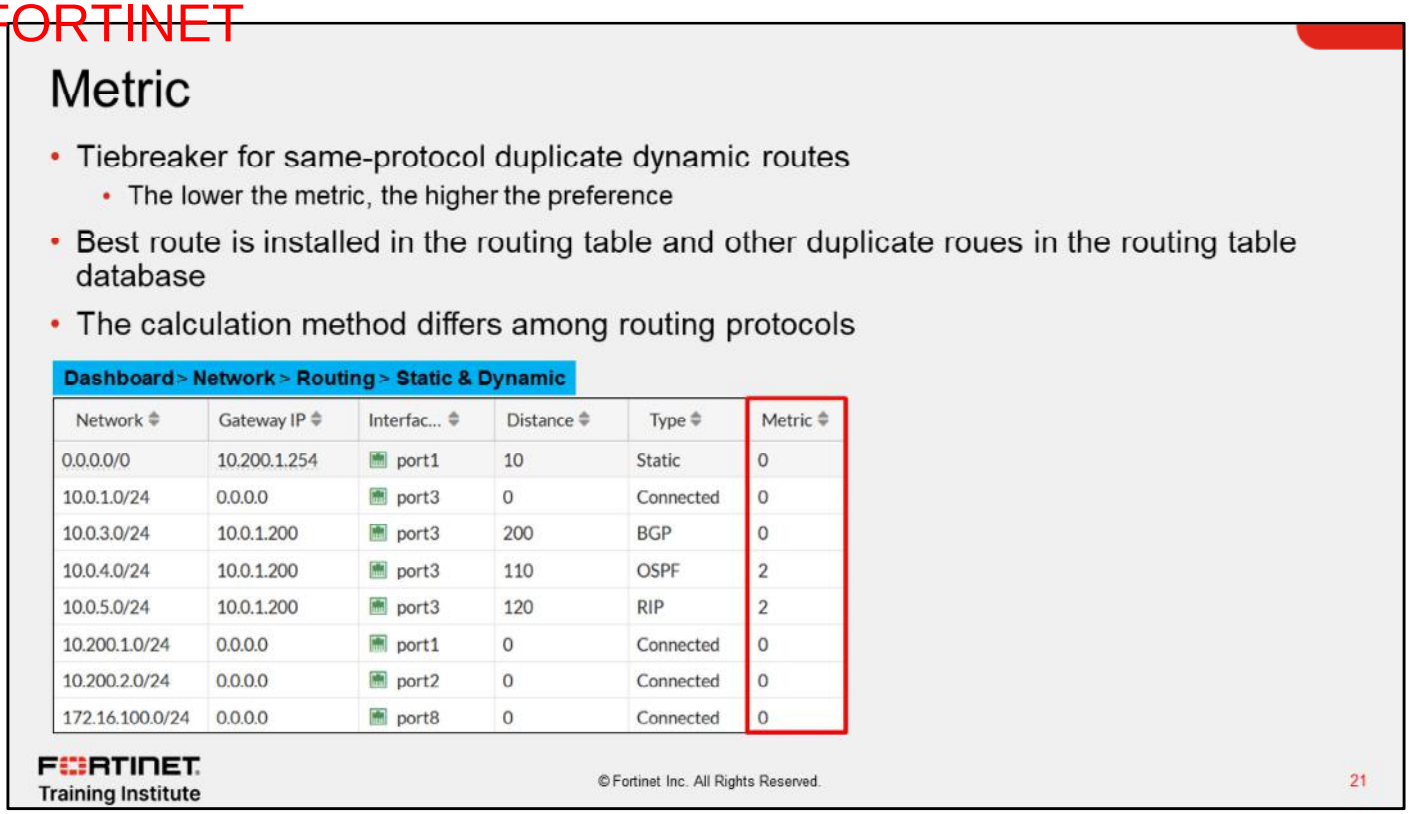

When a dynamic route protocol learns two or more routes to the same destination, it uses the metric as a tiebreaker to identify the best route. The lower the metric, the higher the preference. The dynamic routing protocol then installs the best route in the routing table and the higher-metric routes in the routing table database. Note that the metric is used as tiebreaker for same-protocol dynamic routes, and *not* between different-protocol dynamic routes.

The metric calculation differs among routing protocols, and the details are not covered in this course. For example, RIP uses the hop count, which is the number of routers the packet must pass through to reach the destination. OSPF uses cost, which is determined by the link bandwidth.

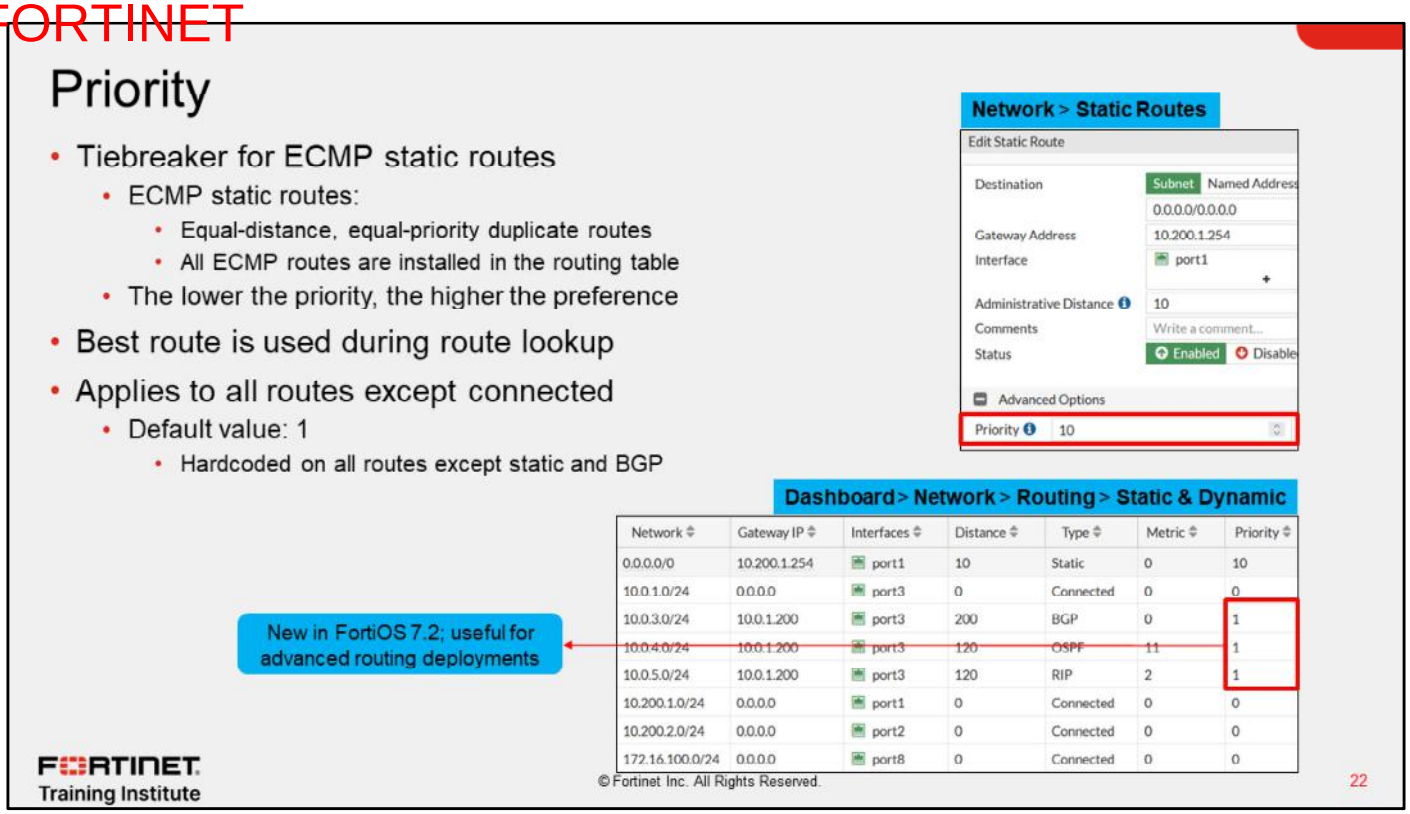

When there are two or more duplicate static routes that have the same distance, FortiGate installs all of them in the routing table. If they also have the same priority, then the routes are known as ECMP static routes, and you will learn more about them in this lesson.

The priority setting enables administrators to break the tie among ECMP *static* routes. The result is that, during the route lookup process, FortiGate selects as the best route the static route with the lowest priority among all the equal-distance duplicate static routes. The lower the priority value, the higher the preference.

Starting FortiOS 7.2, the priority attribute applies to all routes except connected routes and is set to 1 by default. Before FortiOS 7.2, the attribute applied to static routes only and was set to 0 by default. When you upgrade to FortiOS 7.2, FortiOS automatically increases the priority of static routes by 1, and a value of 0 is not longer valid.

For dynamic routes, you can change the priority of BGP routes only. The priority of other dynamic routes is hardcoded to 1. The use of the priority value in dynamic routes is useful for advanced routing deployments involving SD-WAN and multiple virtual routing and forwarding (VRF) IDs. The details on how the priority attribute is beneficial for such cases is outside the scope of this course.

For static routes, you can configure the priority setting under the **Advanced Options** on the FortiGate GUI, as shown on this slide.

To view the priority in the routing monitor widget, you must enable the priority column (disabled by default). You can also view the priority on the routing table on the FortiGate CLI, which you will learn about later in this lesson.

Knowledge Check

1. The priority attribute applies to which type of routes?
   - ✓ A. Static
   - B. Connected

2. Which attribute does FortiGate use to determine the *best* route for same-protocol duplicate dynamic routes?
   - A. Priority
   - ✓ B. Metric

3. Which routes are installed in the routing table?
   - ✓ A. Best active routes
   - B. Policy routes

FORTINET Training Institute

23

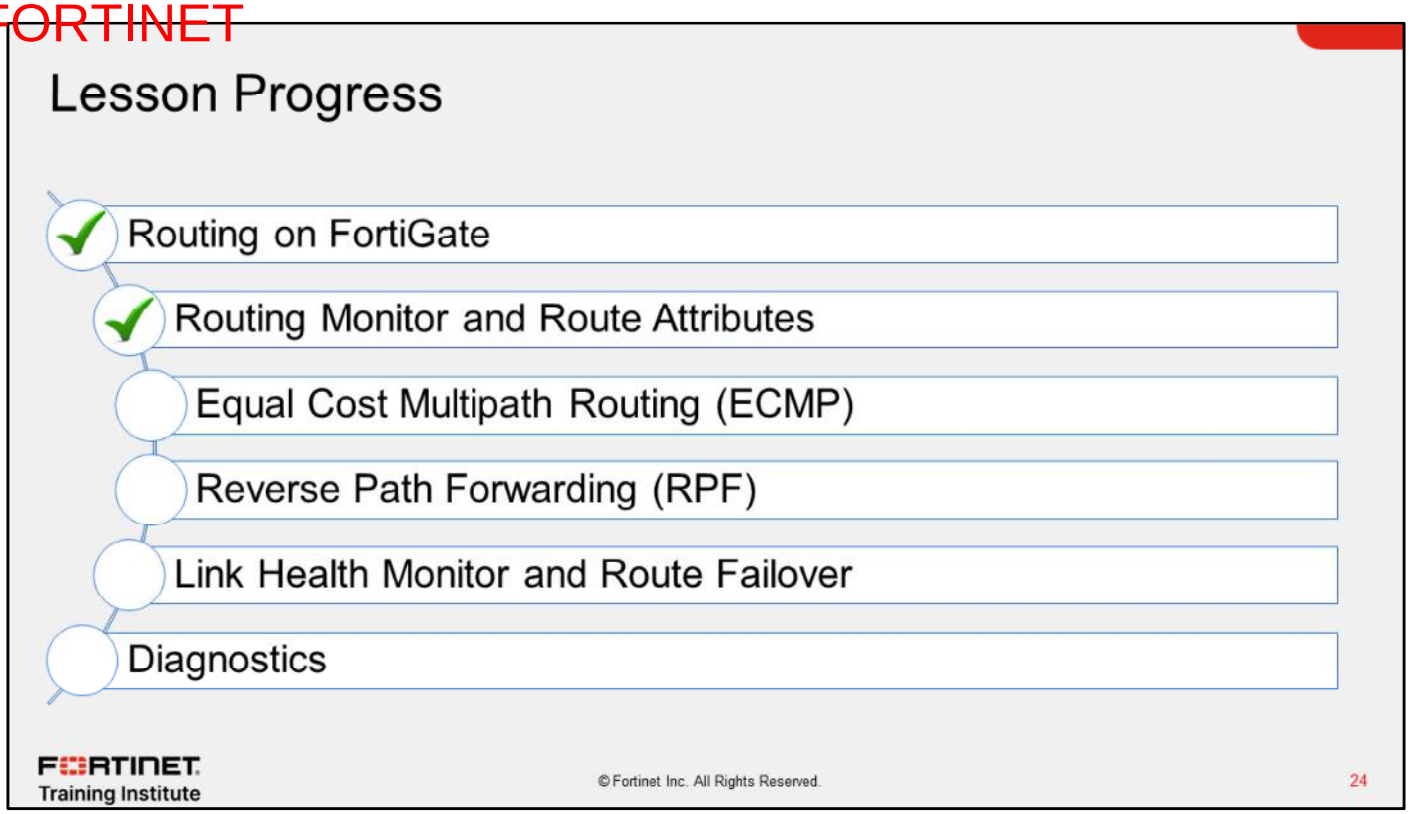

Good job! You now understand the routing monitor and route attributes.

Now, you will learn about ECMP routing.

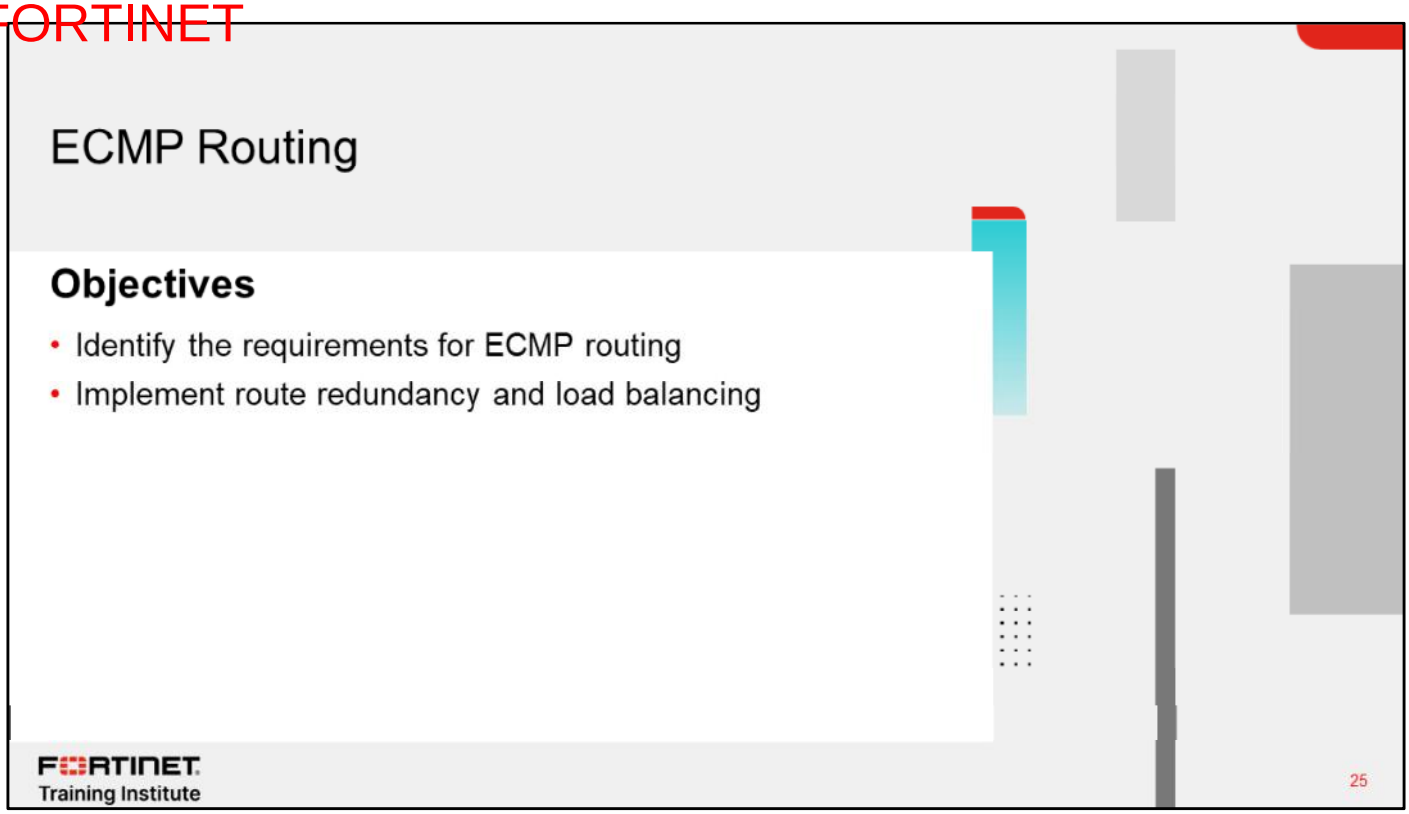

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in ECMP, you should be able to identify the requirements for implementing ECMP and ECMP load balancing.

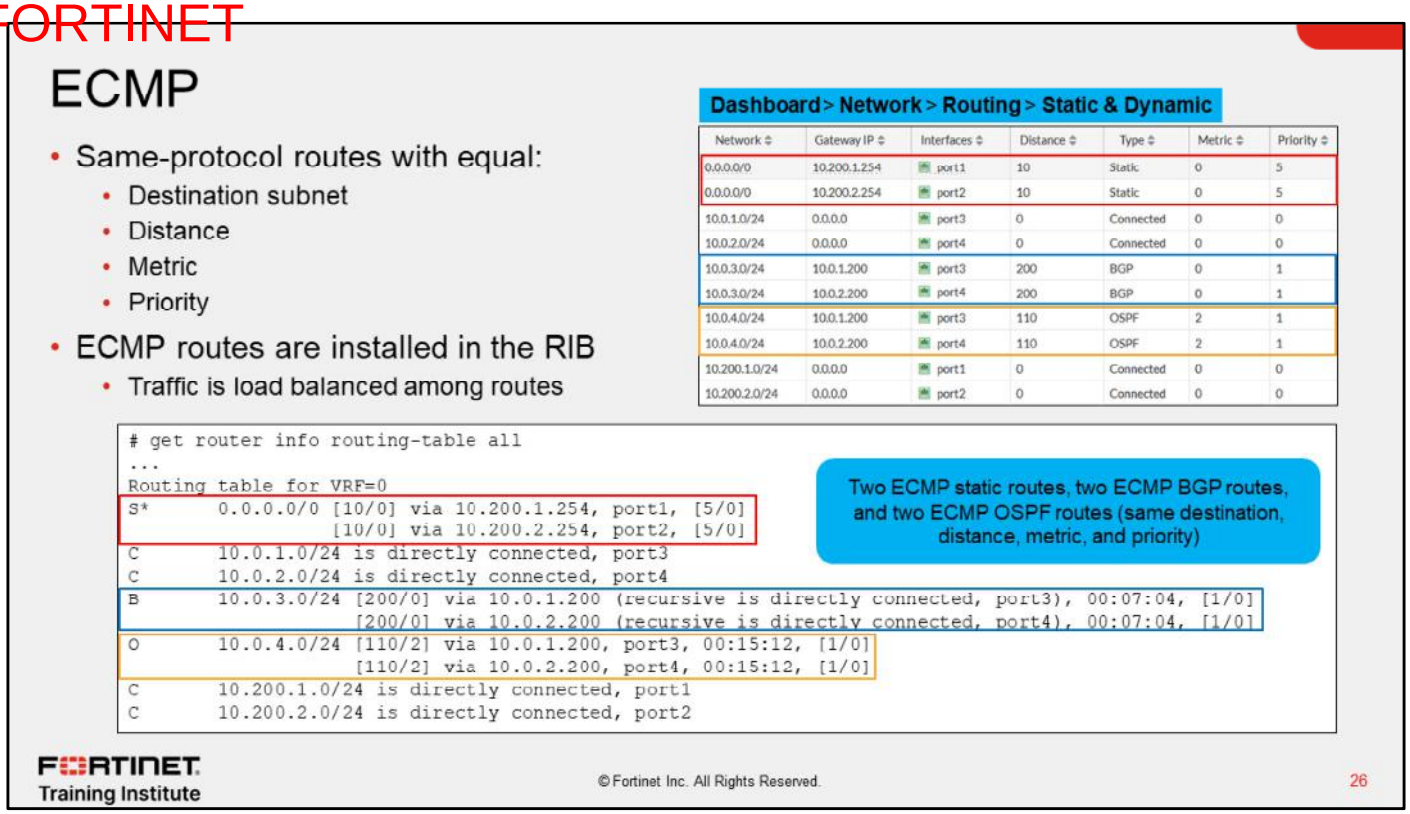

So far, you've learned about the different route attributes that FortiGate looks at to identify the best route to a destination.

But what happens when two or more routes of the same type have the same destination, distance, metric, and priority? These routes are called equal cost multipath (ECMP) routes, and FortiGate installs all of them in the routing table. FortiGate also load balances the traffic among the ECMP routes.

The example on this slide shows two ECMP static routes, two ECMP BGP routes, and two ECMP OSPF routes. For each ECMP group, the destination subnet, distance, metric, and priority are the same.

The result is that FortiGate installs both routes of each ECMP group in the routing table. This lesson, however, focuses on ECMP static routes only.

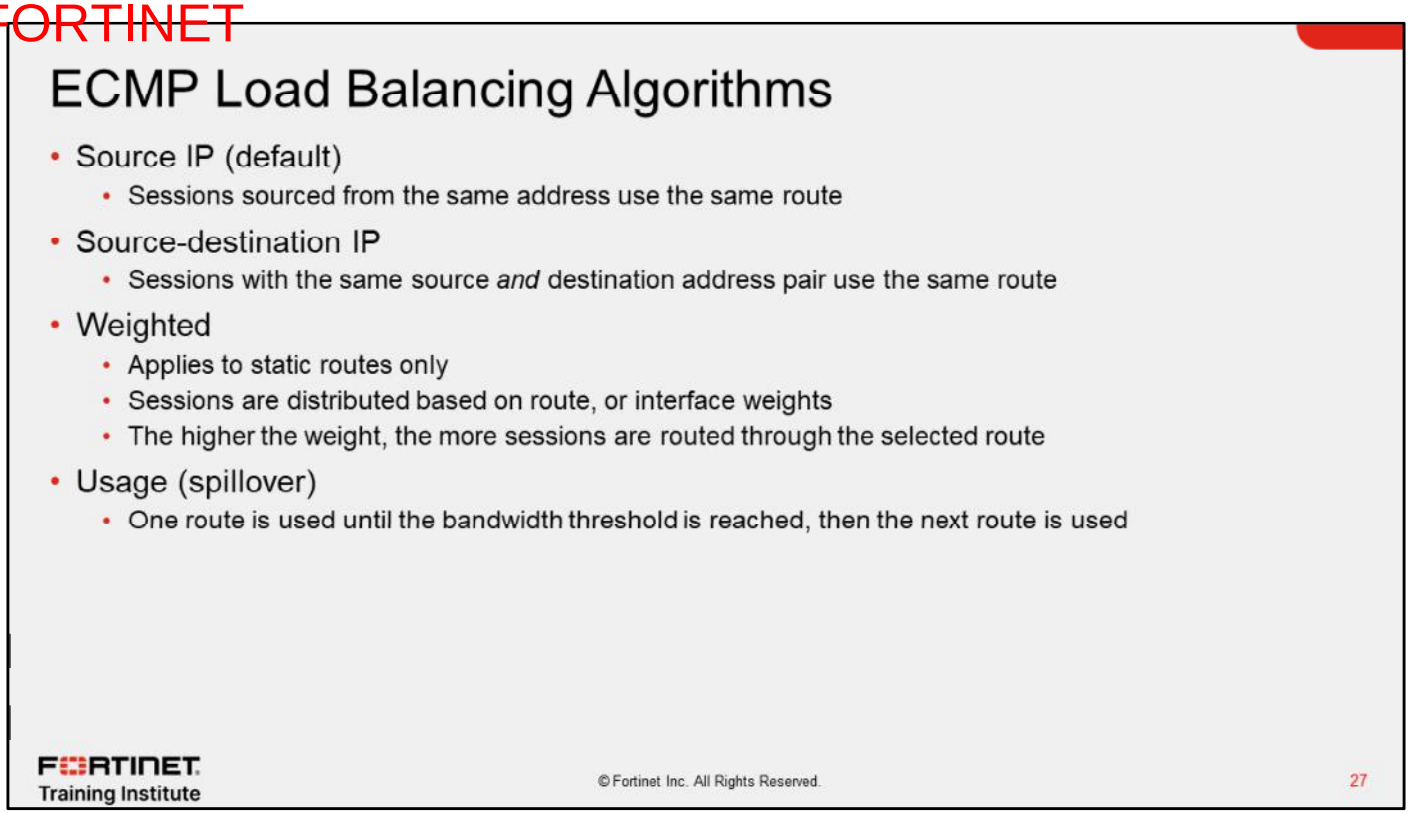

ECMP can load balance sessions using one of the following four algorithms:

- Source IP: This is the default algorithm. FortiGate uses the same ECMP route to route sessions sourced from the same address.
- Source-destination IP: FortiGate uses the same ECMP route to route sessions with the same source-destination IP address pair.
- Weighted: Applies to static routes only. FortiGate load balances sessions based on the route weight or the respective interface weight. The higher the weight, the more sessions FortiGate routes through the selected route.
- Usage (spillover): FortiGate sends sessions to the interface of the first ECMP route until the bandwidth of the interface reaches the configured spillover limit. After the spillover limit is reached, FortiGate uses the interface of the next ECMP route.

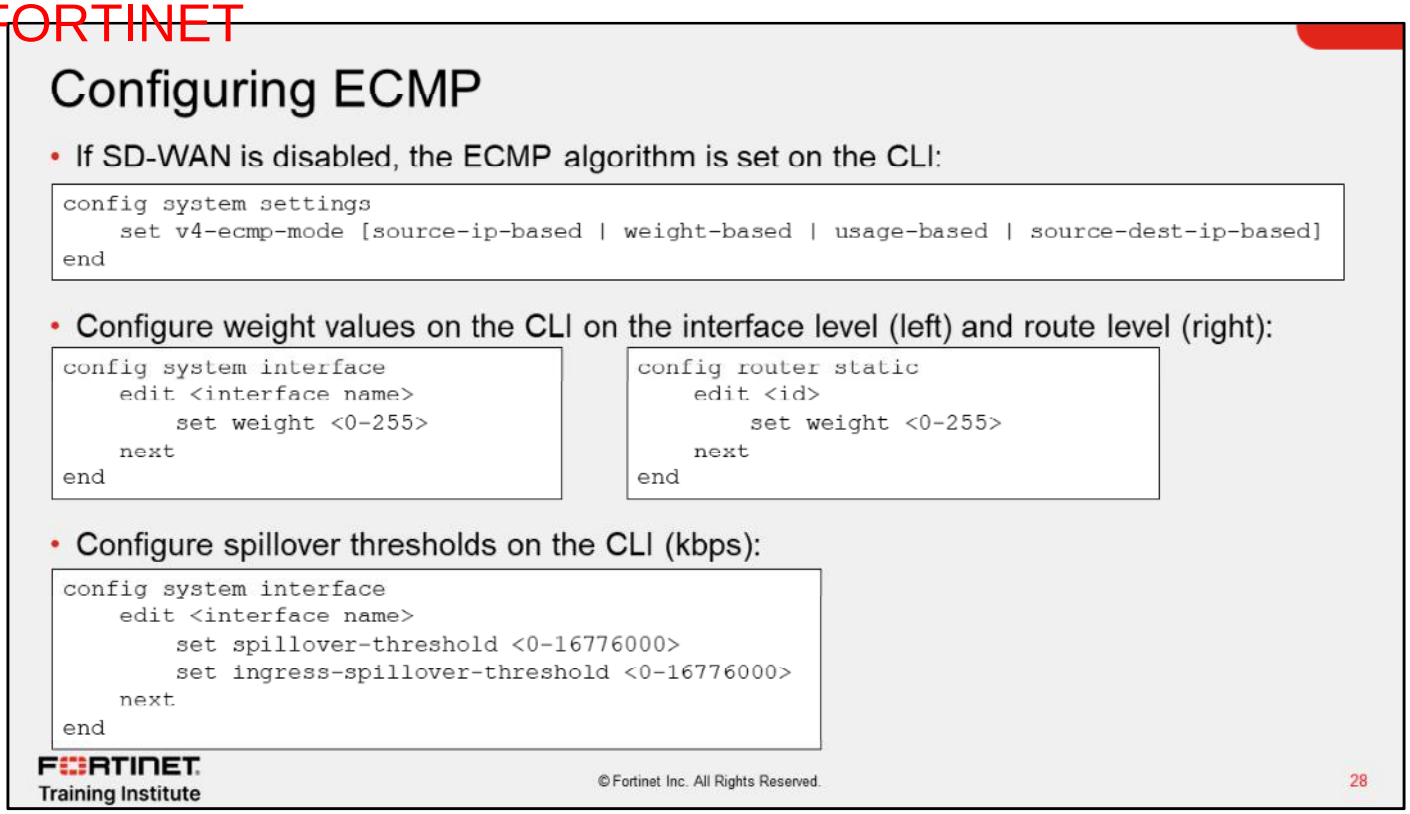

If SD-WAN is disabled, you can change the ECMP load balancing algorithm on the FortiGate CLI using the commands shown on this slide.

When SD-WAN is enabled, FortiOS hides the `v4-ecmp-mode` setting and replaces it with the `load-balance-mode` setting under `config system sdwan`. That is, when you enable SD-WAN, you control the ECMP algorithm with the `load-balance-mode` setting.

For spillover to work, you must also configure the egress and ingress spillover thresholds, as shown on this slide. The thresholds are set to `0` by default, which disables spillover check. For weighted algorithm, you must configure the weights on the interface level or route level, as shown on this slide.

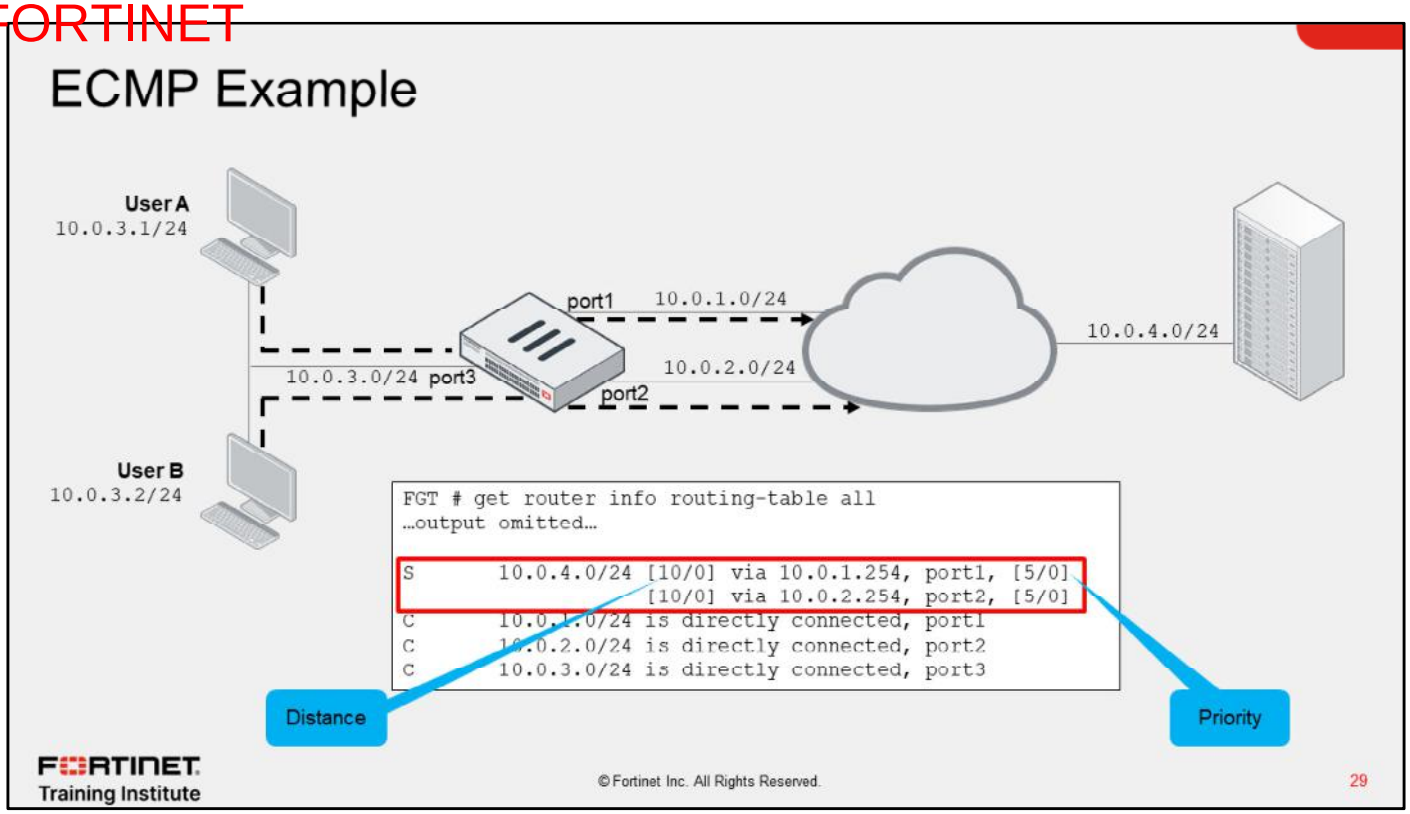

In the scenario shown on this slide, FortiGate has ECMP routes for the `10.0.4.0/24` subnet on port1 and port2. Using the default ECMP algorithm (source IP based), FortiGate may use any of the two routes to route traffic from user A and user B.

In the example shown on this slide, FortiGate selects the route over port1 for user A, and the route over port2 for user B. FortiGate continues to use the same selected routes for the same traffic. In the route over port1 is removed from the routing table, FortiGate automatically starts to forward the traffic sourced from both users and destined to `10.0.4.0/24` through port2.

ECMP enables you to use multiple paths for the same destination, as well as provide built-in failover. Usually, you want to use ECMP for mission-critical services that require high availability. Another reason to use ECMP is for bandwidth aggregation. That is, you can leverage the bandwidth of multiple links by load balancing sessions across them.

While ECMP enables you to leverage multiple WAN links on FortiGate, you may want to use SD-WAN because of the additional benefits.

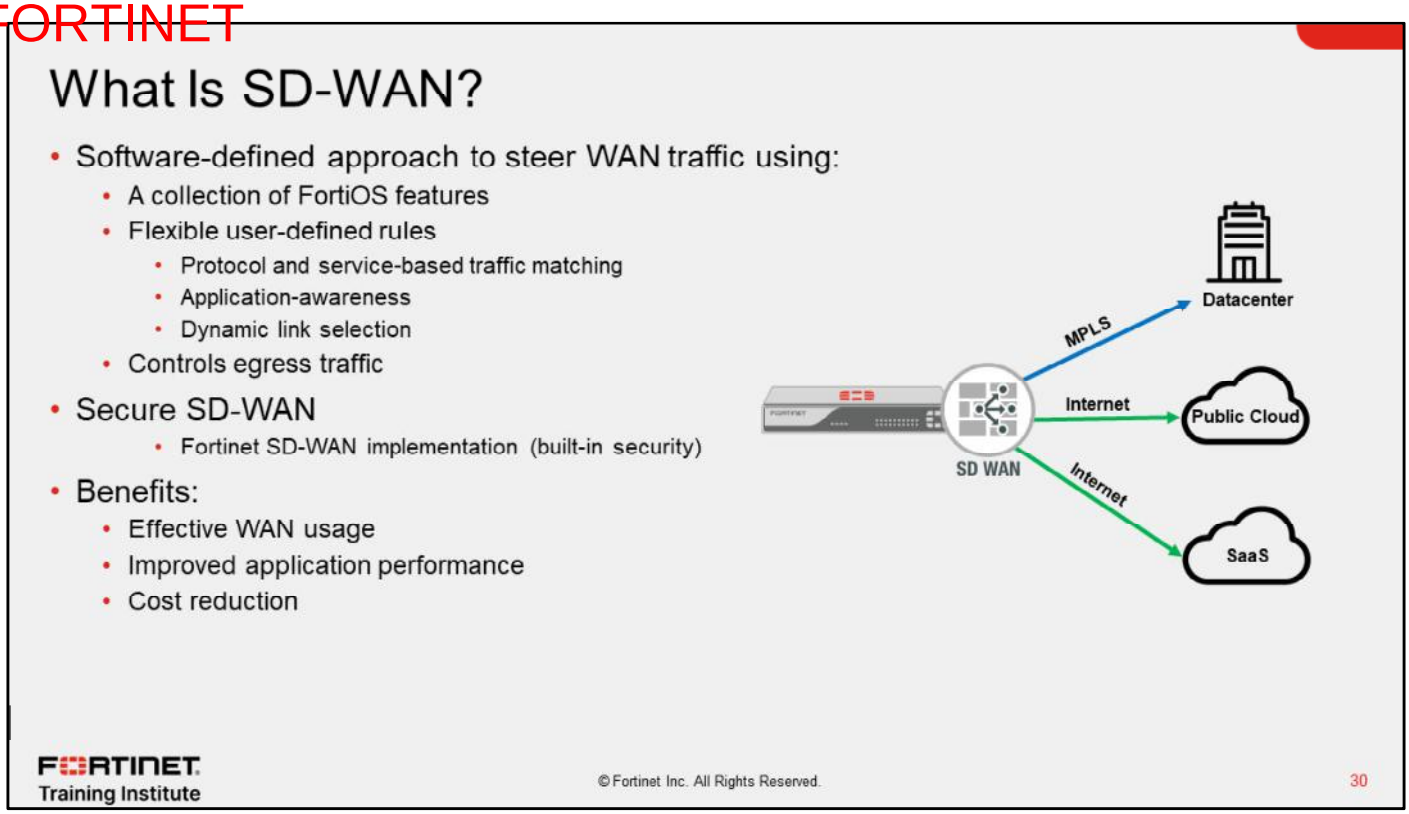

According to Gartner, software-defined WAN (SD-WAN) provides dynamic, policy-based, application path selection across multiple WAN connections and supports service chaining for additional services such as WAN optimization and firewalls. Fortinet implementation of SD-WAN is called secure SD-WAN because it also provides security by leveraging the built-in security features available in FortiOS.

Secure SD-WAN relies on well-known FortiOS features such as IPsec, auto-discovery VPN (ADVPN), link monitoring, advanced routing, internet services database (ISDB), traffic shaping, UTM inspection, and load balancing. The administrator can then combine these features and set rules that define how FortiGate steers traffic across the WAN based on multiple factors such as the protocol, service, or application identified for the traffic, and the quality of the links. Note that SD-WAN controls egress traffic, not ingress traffic. This means that the return traffic may use a different link from the one SD-WAN chose for egress.

One benefit of SD-WAN is effective WAN usage. That is, you can use public (for example, broadband, LTE) and private (for example, MPLS) links to securely steer traffic to different destinations: internet, public cloud, private cloud, and the corporate network. This approach of using different types of links to connect sites to private and public networks is known as hybrid WAN. A hybrid WAN reduces costs mainly because administrators usually steer more traffic over low-cost fast internet links than high-cost slow private links. The result is that private links, such as MPLS links, are often used to steer critical traffic only, or as failover links for high availability.

Another benefit of SD-WAN is an improved application performance because you can steer traffic through the best link that meets the application requirements. During congestion, you can leverage traffic shaping to prioritize sensitive and critical applications over less important ones. Also, the support of ADVPN shortcuts enables SD-WAN to use direct IPsec tunnels between sites to steer traffic, resulting in lower latency for traffic between the sites (spokes), and less load on the central locations (hubs).

Direct Internet Access With SD-WAN

- Traffic steered across multiple internet links
- Typical operation:
  - Critical/sensitive traffic expedited and steered over best performing links
  - Costly links used for critical traffic or failover
  - Static default routing

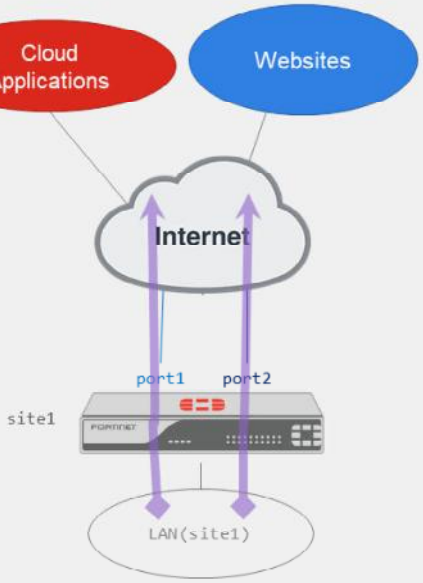

FORTINET Training Institute

31

Direct internet access (DIA), also known as local breakout, is arguably the most common use case for SD-WAN. A site has multiple internet links (also known as underlay links), and the administrator wants FortiGate to steer internet traffic across the links (also known as members). The links are connected to FortiGate using different types of physical interfaces: physical port, VLAN, link aggregation (LAG), USB modem, or through FortiExtender.

Usually, sensitive traffic is expedited and steered over the best performing links, while non-critical traffic is distributed across one or more links using a best effort approach. Costly internet links are commonly used as backup links, or to steer critical traffic only.

For routing, a typical configuration makes use of static default routes. However, in some cases, BGP is used between the ISP and FortiGate, especially if the site must advertise a public IP prefix.

The example on this slide shows a basic DIA deployment. FortiGate has two internet links. One link is connected to port1 and the other to port2. FortiGate uses both links to steer traffic sourced from the LAN and destined to cloud applications and websites on the internet.

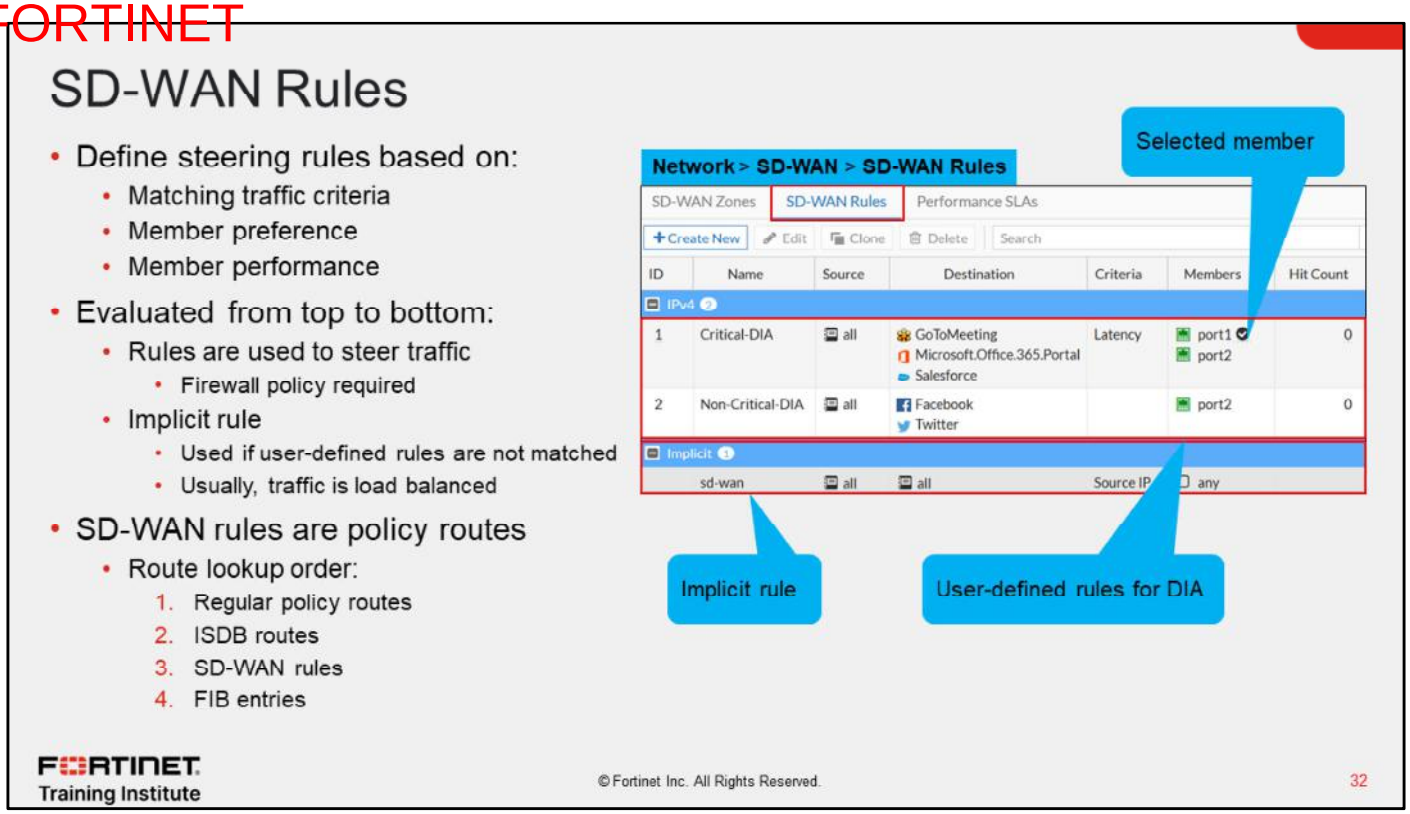

SD-WAN rules represent the intelligence of the SD-WAN solution and the software-defined aspect of it. When you configure an SD-WAN rule, you first define the application or traffic pattern to match. After that, you indicate the preferred members and/or zones to steer the matching traffic to, and in some cases, the performing metrics that the member must meet to be eligible for steering traffic.

SD-WAN rules are evaluated in the same way as firewall policies: from top to bottom, using the first match. However, unlike firewall policies, they are used to steer traffic, and *not* to allow traffic. That is, you must configure corresponding firewall policies to allow the SD-WAN traffic. If none of the user-defined SD-WAN rules are matched, then the implicit rule is used. The implicit rule instructs FortiGate to perform standard routing on traffic. Because SD-WAN deployments usually have multiple routes to the same destination—that is, ECMP routes—then traffic that matches the implicit rule is usually load balanced across multiple SD-WAN members.

SD-WAN rules are essentially policy routes. Like regular policy routes, SD-WAN rules route traffic based on multiple criteria. That is, when you configure an SD-WAN rule, the kernel installs a corresponding policy route that reflects the source, destination, service, and outgoing interfaces configured in the SD-WAN rule. When FortiGate performs a route lookup, it checks the routes in the order of sequence shown on this slide. For example, SD-WAN rules have precedence over FIB entries, but not over regular policy routes.

The example on this slide shows two user-defined rules named **Critical-DIA** and **Non-Critical-DIA**, which are used to steer traffic in our basic DIA setup. The **Critical-DIA** steers **GoToMeeting**, **Microsoft.Office.365.Portal**, and **Salesforce** traffic to the member with the lowest latency, between **port1** and **port2**. The example shows that **port1** is selected because it is the member with the check mark beside it. The **Non-Critical-DIA** rule steers Facebook and Twitter traffic to **port2**. The implicit rule, located at the bottom of the list, is used if none of the two user-defined rules are matched.

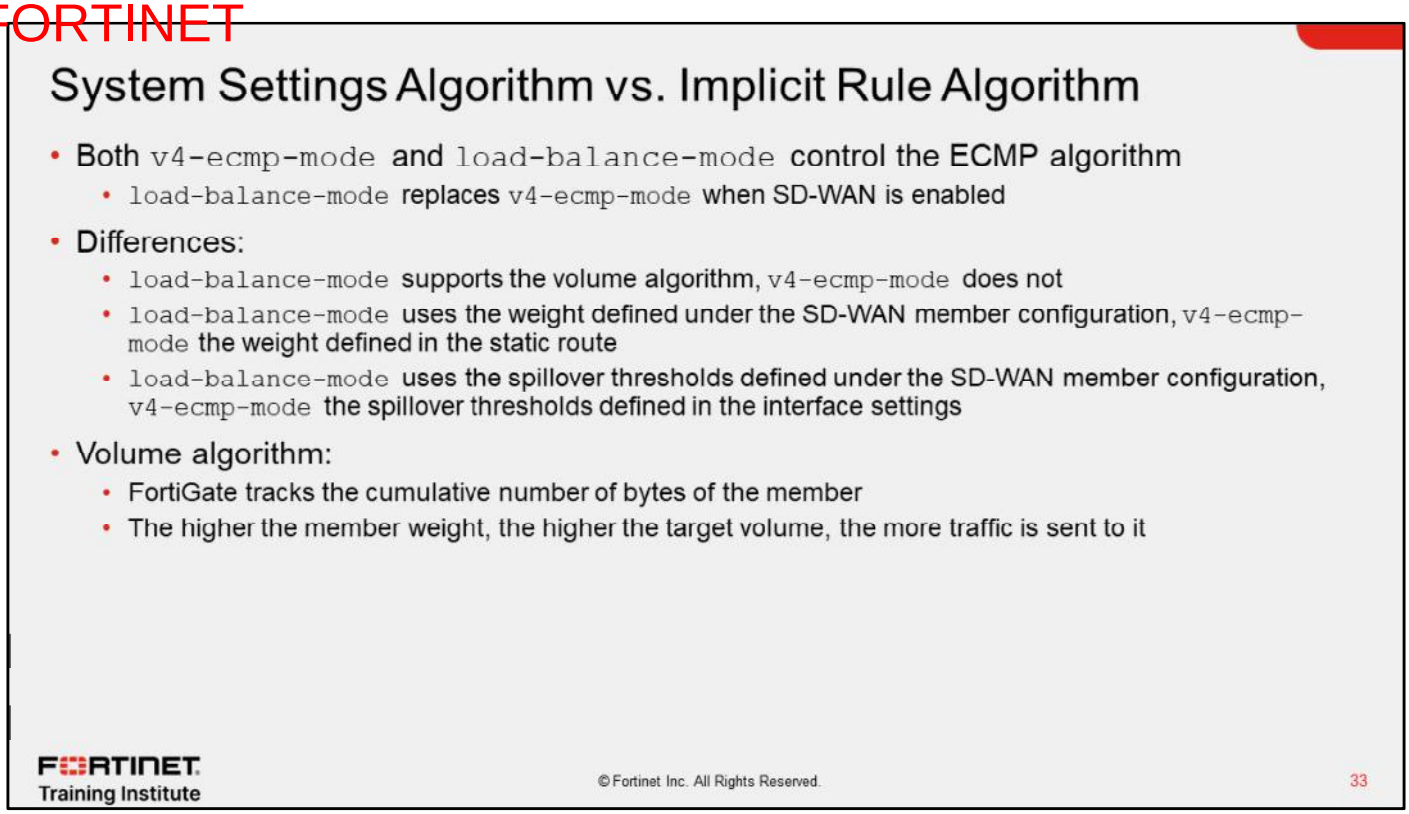

When you enable SD-WAN, FortiOS hides the `v4-ecmp-mode` setting and replaces it with the `load-balance-mode` setting under `config system sdwan`. That is, after you enable SD-WAN, you now control the ECMP algorithm with the `load-balance-mode` setting.

There are some differences between the two settings. The main difference is that `load-balance-mode` supports the volume algorithm, and `v4-ecmp-mode` does not. In addition, the related settings such as weight and spillover thresholds are configured differently. That is, when you enable SD-WAN, the weight and spillover thresholds are defined on the SD-WAN member configuration. When you disable SD-WAN, the weight and spillover thresholds are defined on the static route and interface settings, respectively.

When you set the ECMP algorithm to volume, FortiGate load balances sessions across members based on the measured interface volume and the member weight. That is, the volume algorithm instructs FortiGate to track the cumulative number of bytes of each member and to distribute sessions based on the weight. The higher the weight, the higher the target volume of the interface and, as a result, the more traffic FortiGate sends to it.

Knowledge Check

1. What is the default ECMP algorithm on FortiGate?
   - A. Weighted
   - ✓ B. Source IP

2. How does FortiGate load balance traffic when using the spillover algorithm in ECMP routing?
   - ✓ A. Sessions are distributed based on interface threshold.
   - B. Sessions are distributed based on route weight.

FORTINET Training Institute

34

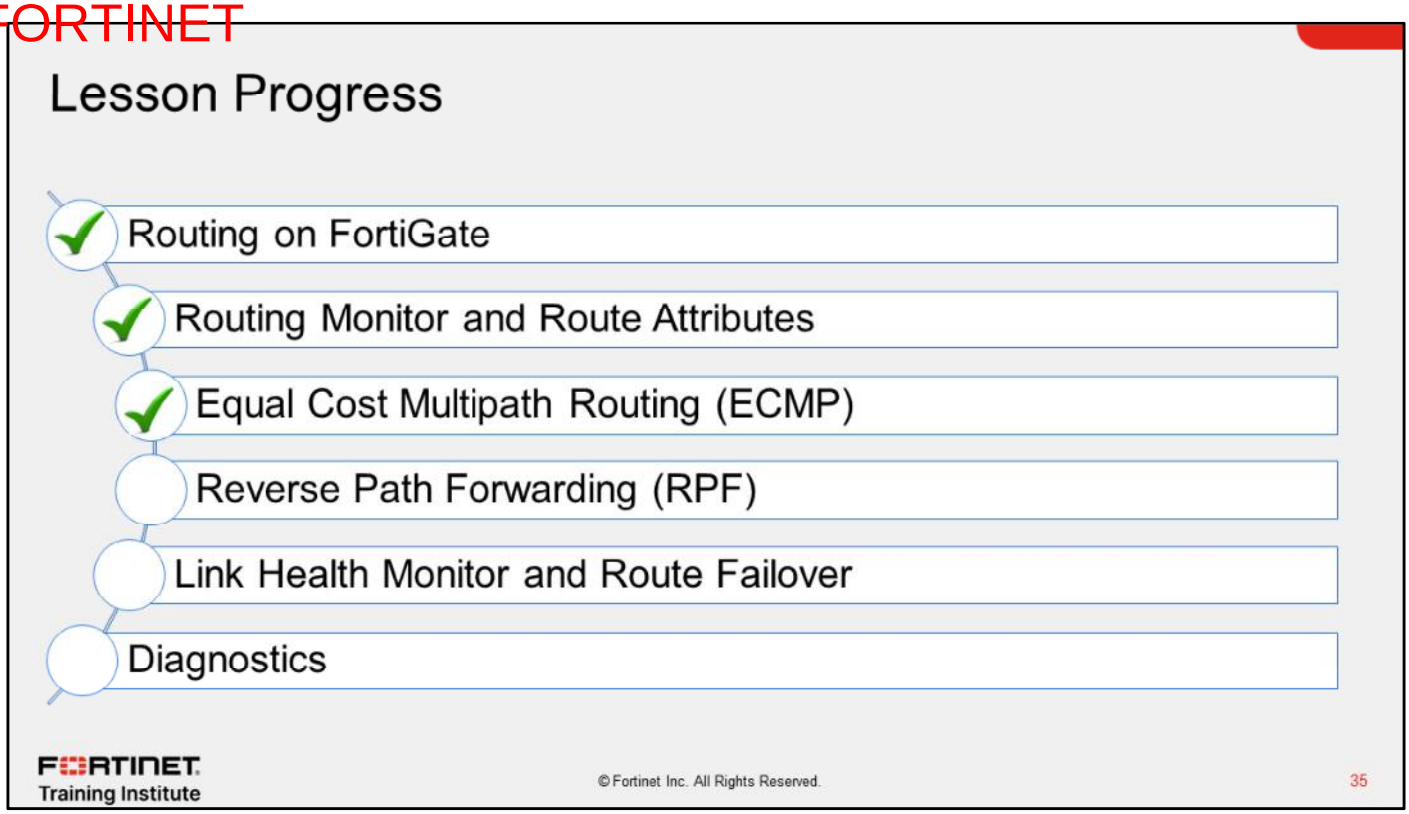

Good job! You now understand ECMP routing.

Now, you will learn about reverse path forwarding.

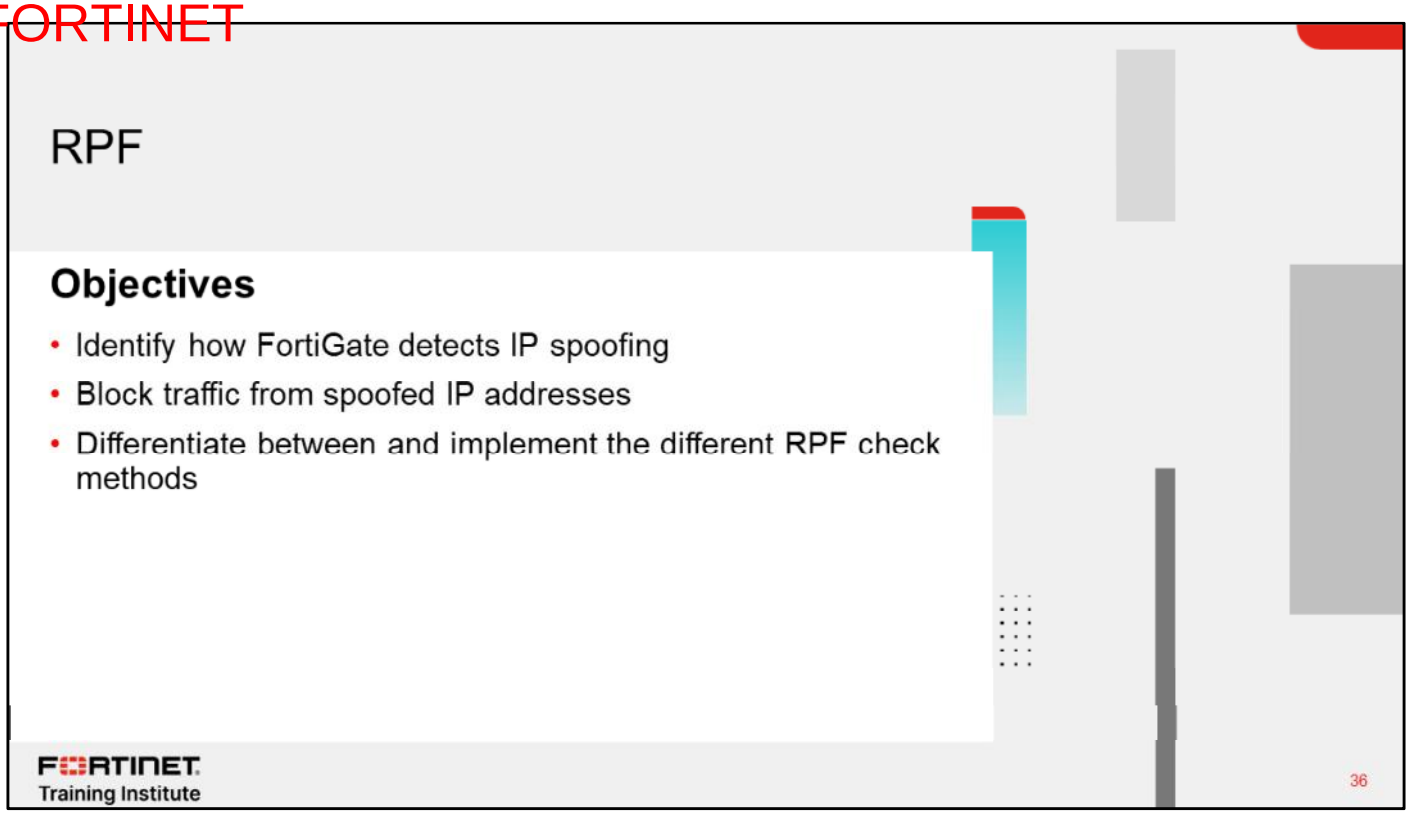

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in RPF, you should be able to identify and block IP spoofing attacks in your network.

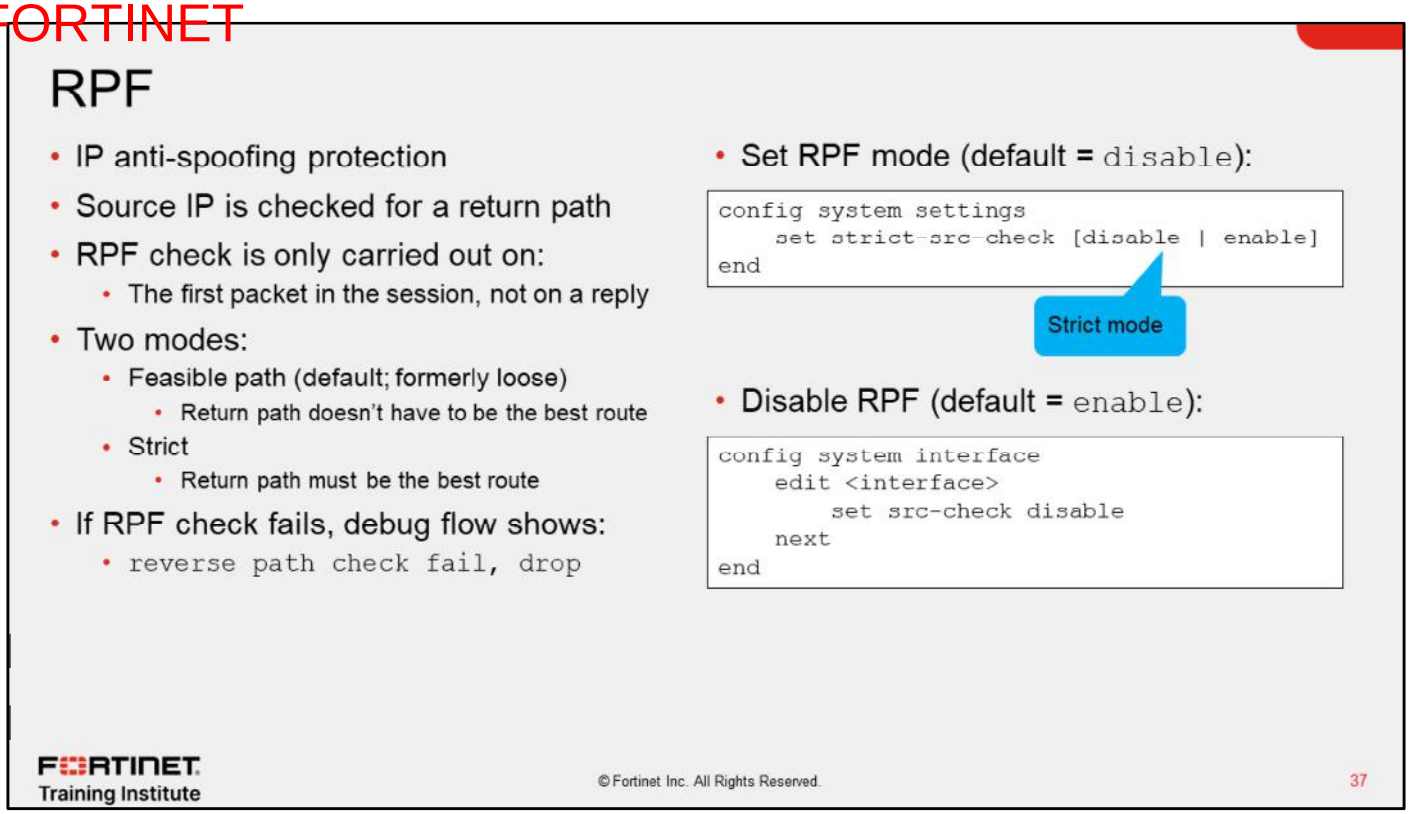

The RPF check is a mechanism that protects FortiGate and your network from IP spoofing attacks by checking for a return path to the source in the routing table.

The premise behind the RPF check is that if FortiGate receives a packet on an interface, and FortiGate doesn't have a route to the packet source address through the incoming interface, then the source address of the packet could have been forged, or the packet was routed incorrectly. In either case, you want to drop that unexpected packet, so it doesn't enter your network.

FortiGate performs an RPF check only on the first packet of a new session. That is, after the first packet passes the RPF check and FortiGate accepts the session, FortiGate doesn't perform any additional RPF checks on that session.

There are two RPF check modes:

- Feasible path: Formerly known as loose, it's the default mode. In this mode, FortiGate verifies that the routing table contains a route that matches the source address of the packet and the incoming interface. The matching route doesn't have to be the best route in the routing table for that source address. It just has to match the source address and the incoming interface of the packet.

- Strict: In this mode, FortiGate also verifies that the matching route is the best route in the routing table. That is, if the routing table contains a matching route for the source address and incoming interface, but there is a better route for the source address through another interface, then, the RPF check fails.

This slide also shows how to change the RPF check mode on the FortiGate CLI, as well as how to disable the RPF check on the interface level.

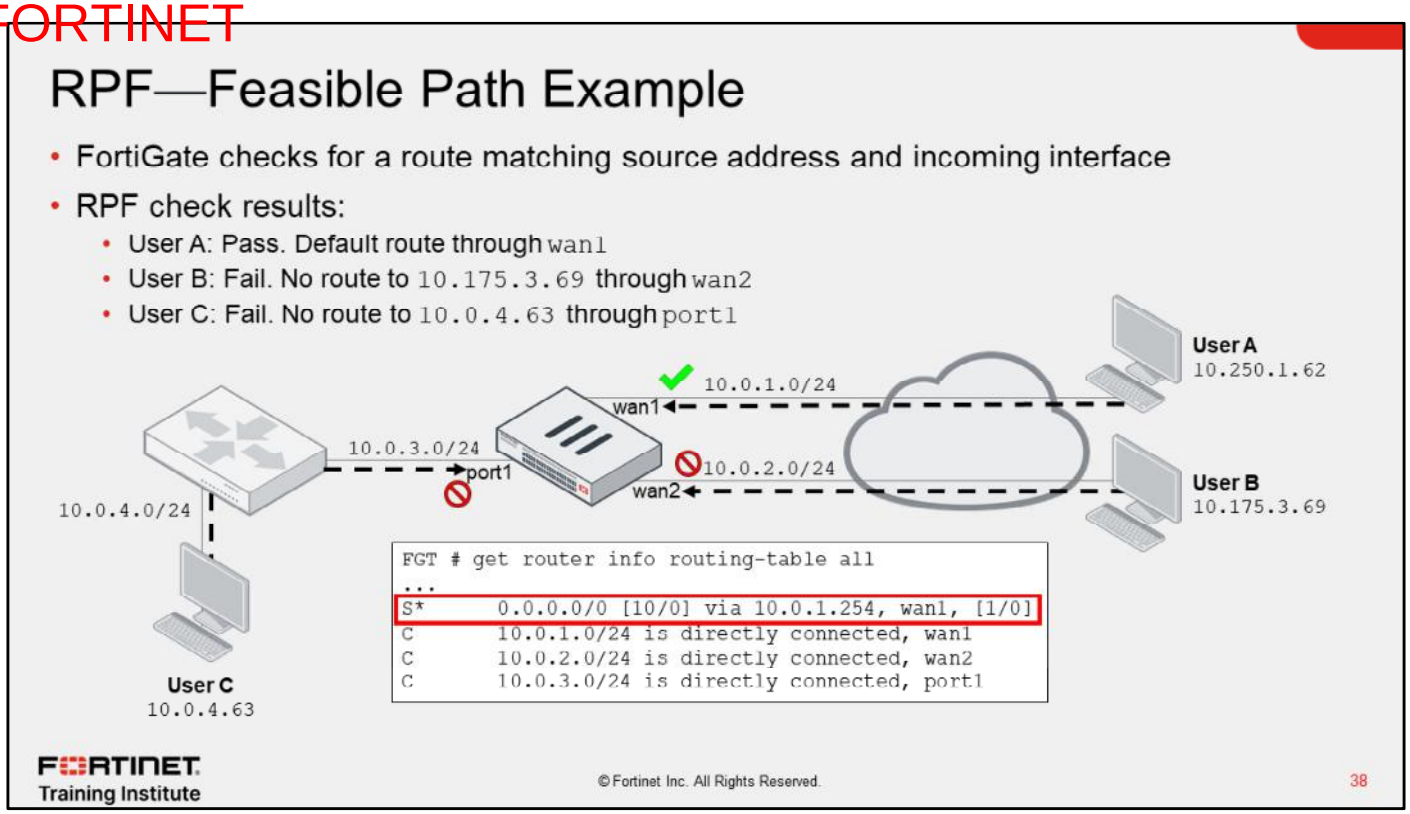

The example on this slide shows a FortiGate device using the feasible path RPF check mode. When FortiGate performs RPF check, it checks in the routing table for a route that matches the source address and the incoming interface of the first original packet.

Based on the topology and routing table shown on this slide, the RPF check results for traffic sourced from each user are:

- User A: Pass. There is a default route through `wan1`. This means that, all packets received at `wan1` pass the RPF check regardless of the source address.
- User B: Fail. FortiGate doesn't have a route to `10.175.3.69` through `wan2` in its routing table.
- User C: Fail. Like the user B case, FortiGate doesn't have a route to `10.0.4.63` through `port1` in its routing table.

RPF—Feasible Path Example (Contd)

- Solution:
  - For user B, add a second static default route, with the same distance, through wan2
    - Use different priority values if you don't want ECMP
  - For user C, add a static route to `10.0.4.0/24` through port1

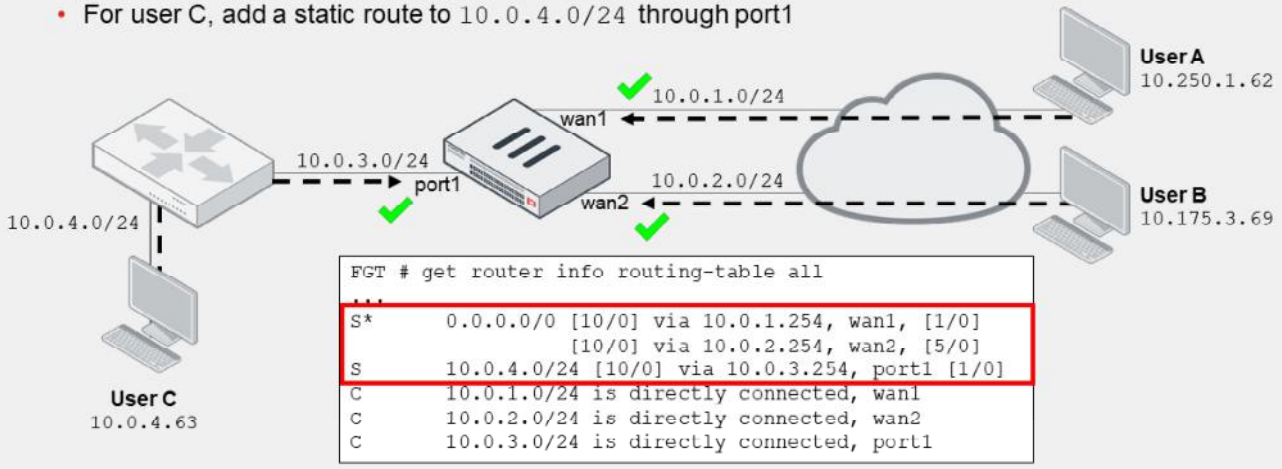


```
FGT # get router info routing-table all
...
S*      0.0.0.0/0 [10/0] via 10.0.1.254, wan1, [1/0]
                 [10/0] via 10.0.2.254, wan2, [5/0]
S       10.0.4.0/24 [10/0] via 10.0.3.254, port1 [1/0]
C       10.0.1.0/24 is directly connected, wan1
C       10.0.2.0/24 is directly connected, wan2
C       10.0.3.0/24 is directly connected, port1
```

FORTINET Training Institute

39

If you consider the packets from user B and user C to be legit packets, you can solve the RPC check fail issue by making sure the routing table contains routes for the return path.

In the example shown on this slide, the administrator adds two new static routes. The static route through `wan2` is a duplicate default route of `wan1`, but has a lower priority. The two default routes are not ECMP routes because of the priority difference, but FortiGate keeps both routes in the routing table. The result is that packets from user B now pass the RPF check.

The static route through `port1` references the `10.0.4.0/24` subnet. The subnet includes user C address (`10.0.4.63`), and as result, packets from user C also pass the RPF check.

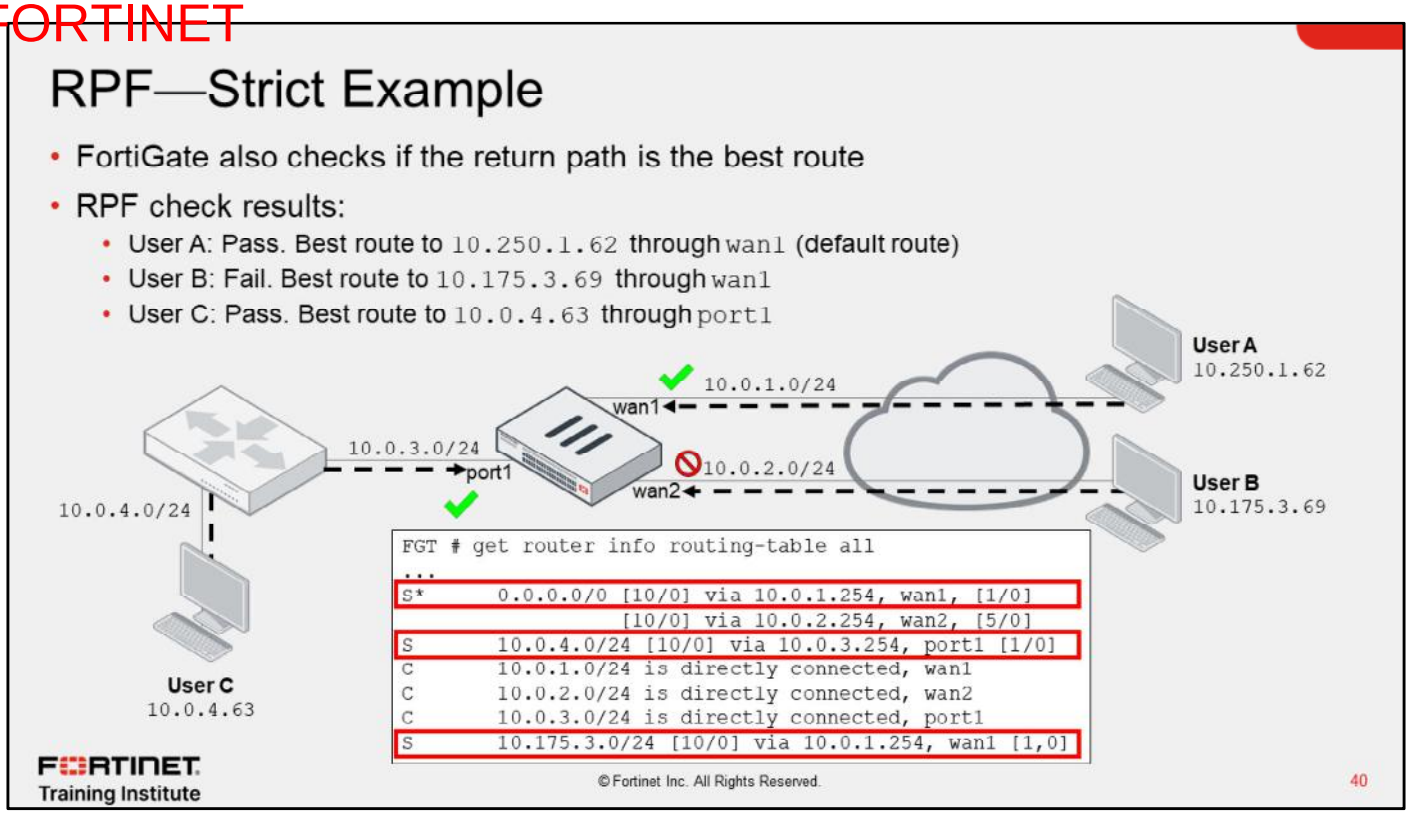

The example on this slide shows a FortiGate device using the strict RPF check mode. In strict mode, FortiGate also checks if the matching route is the best route to the source.

Based on the topology and routing table shown on this slide, the RPF check results for traffic sourced from each user are:

- User A: Pass. There is a default route through `wan1`. The route is also the best (and only) route to `10.250.1.62`.
- User B: Fail. There is a default route through `wan2`. However, there is better (more specific) static route to `10.175.3.69` through `wan1`.
- User C: Pass. FortiGate has a route to `10.0.4.63` through `port1` in its routing table. Although the default routes through wan1 and wan2 are also valid routes for `10.0.4.63`, the best route to user C is the route through port1.

Like the feasible path example, you can solve the RPF fail issue for user B by making the respective changes in the routing table so the best route to user B is through wan2.

Knowledge Check

1. What is the default RPF check method on FortiGate?
   - ✓ A. Feasible path
   - B. Strict

2. Which route lookup scenario satisfies the RPF check for a packet?
   - A. Routing table has a route to the destination IP of the packet through the incoming interface.
   - ✓ B. Routing table has a route for the source IP of the packet through the incoming interface.

FORTINET Training Institute

41

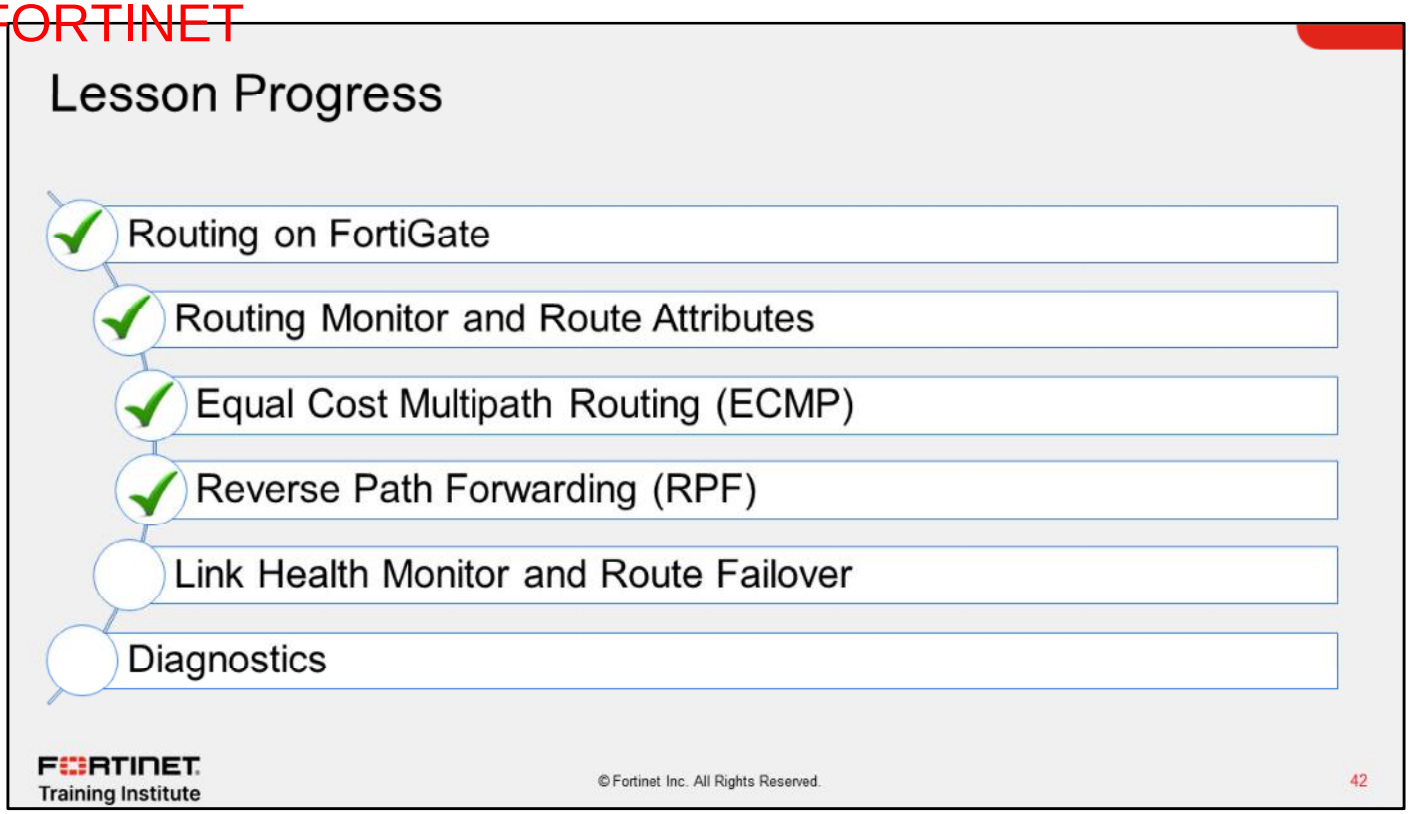

Good job! You now understand RPF.

Now, you will learn about the link health monitor and route failover.

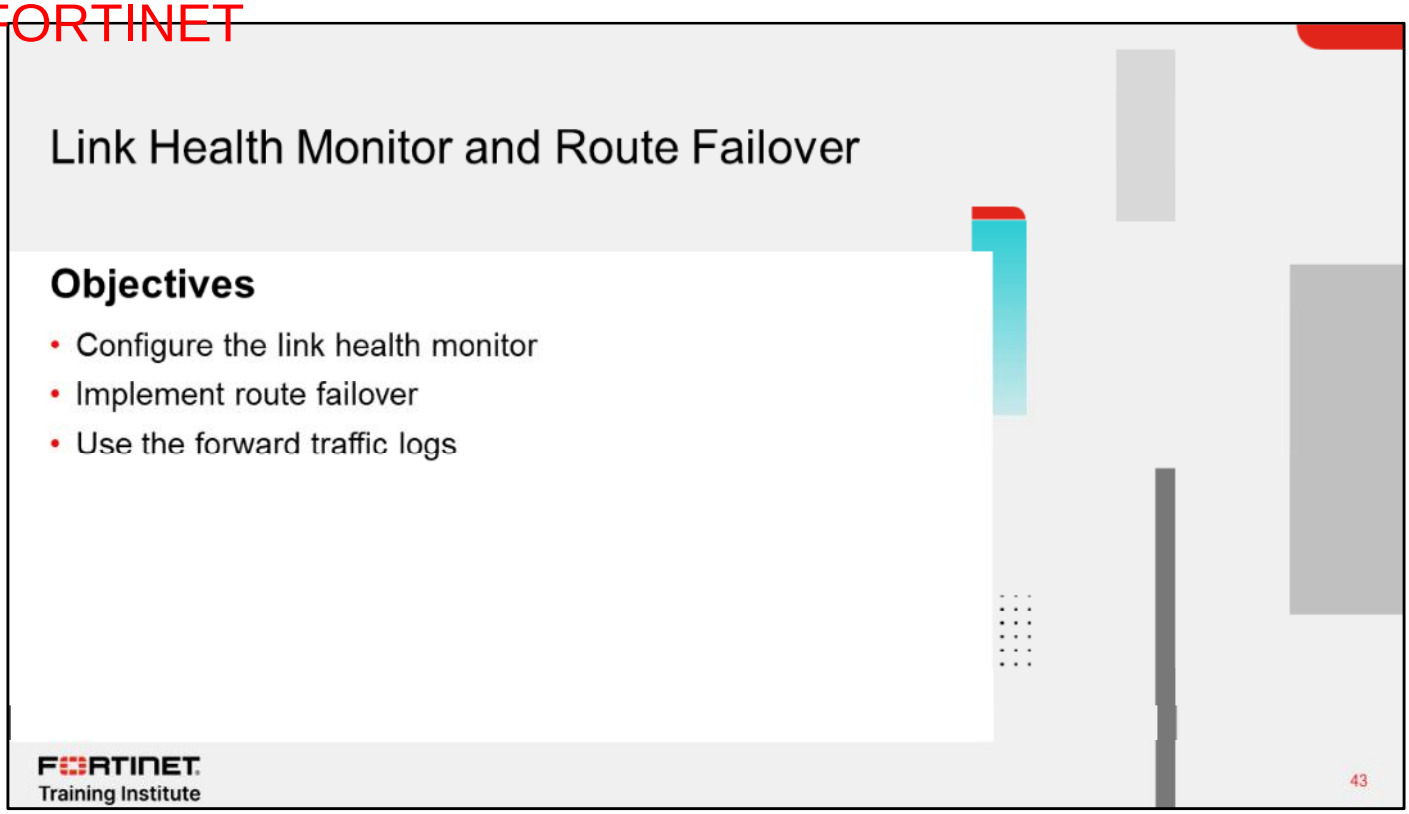

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in configuring link health monitor and implementing route failover, you should be able to monitor the health of your interfaces and then, when a link is detected as dead, configure FortiGate to fail over the traffic to healthy links to minimize service disruption.

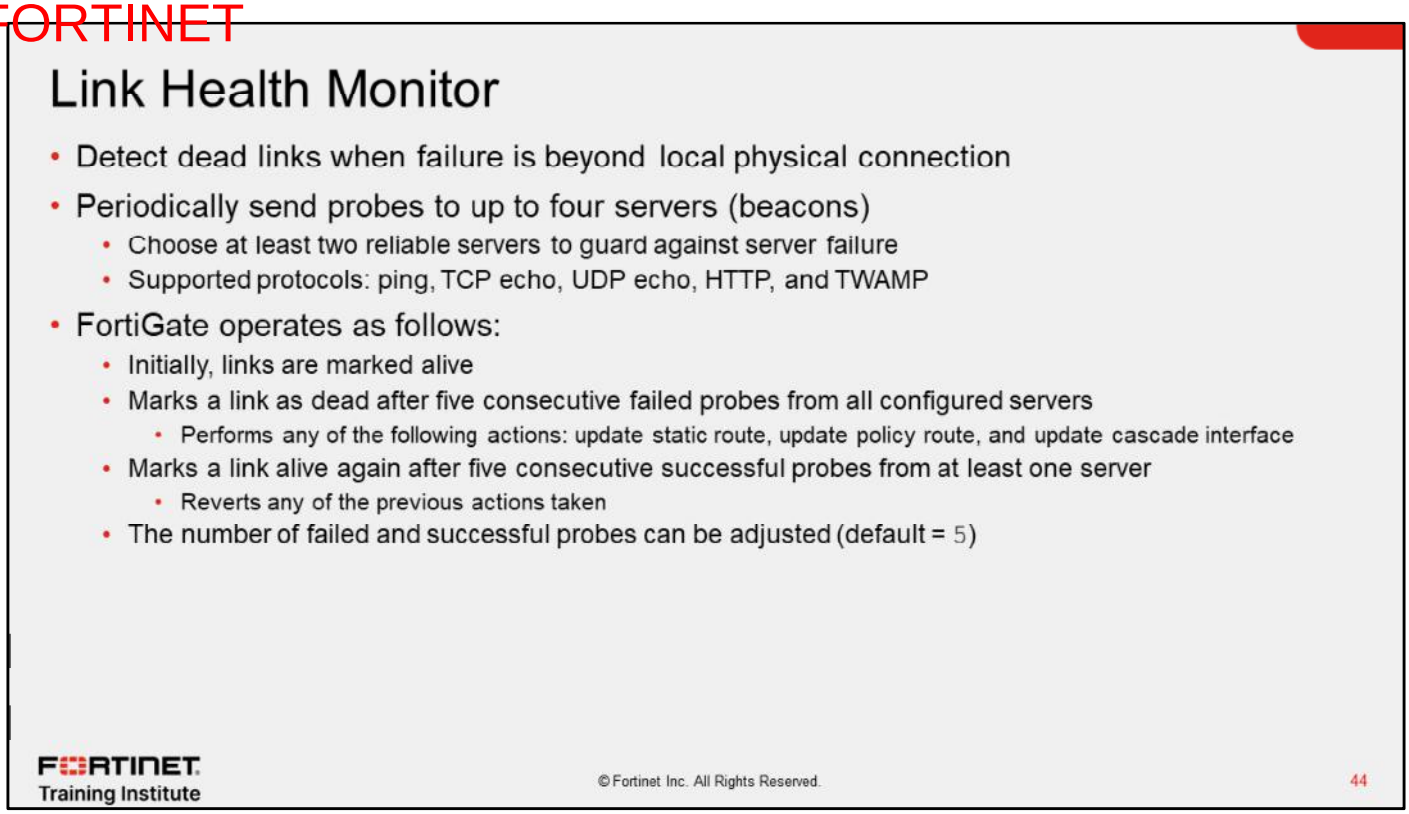

Static routes are kept in the routing table unless the associated interface is administratively down, its link goes down, or there is a duplicate route with a lower distance. Because is possible that the link circuit is dead somewhere along the path to the destination, even though the interface link is up, then it is also possible that FortiGate continues to route traffic though a dead link, which would result in service impact. A common example is the Ethernet connection provided by your ISP modem. The Ethernet connection remains physically up even though the upstream ISP network is down. The devices behind your modem will continue to use the internet connection but they won't receive any replies.

Link health monitor enables FortiGate to detect dead links when the failure is beyond the local physical connection. FortiGate periodically sends probes through the configured gateway and interface to up to four servers that act as beacons. A server can be any host that is normally reachable through that path. It's best practice to configure at least two reliable servers to guard against false positives caused by the server being at fault, and not the link. For probes, you should also use a protocol that the server normally responds to.

Initially, FortiGate considers a link as alive. However, if FortiGate detects five consecutive failed probes from each of the configured servers, FortiGate marks the link as dead. FortiGate considers a failed probe a probe for which it does not receive a reply, or whose reply isn't valid. After FortiGate detects the link as dead, it performs any of the actions shown on this slide. The goal of these actions is to redirect the impacted traffic to other healthy links.

After FortiGate detects the link as dead, it continues to monitor the link. As soon as FortiGate receives five successful replies from at least one of the configured servers, it marks the link as alive again, and then reverts any of the previous actions taken on that link.

The number of failed and successful probes is set to five by default, but can be changed if required.

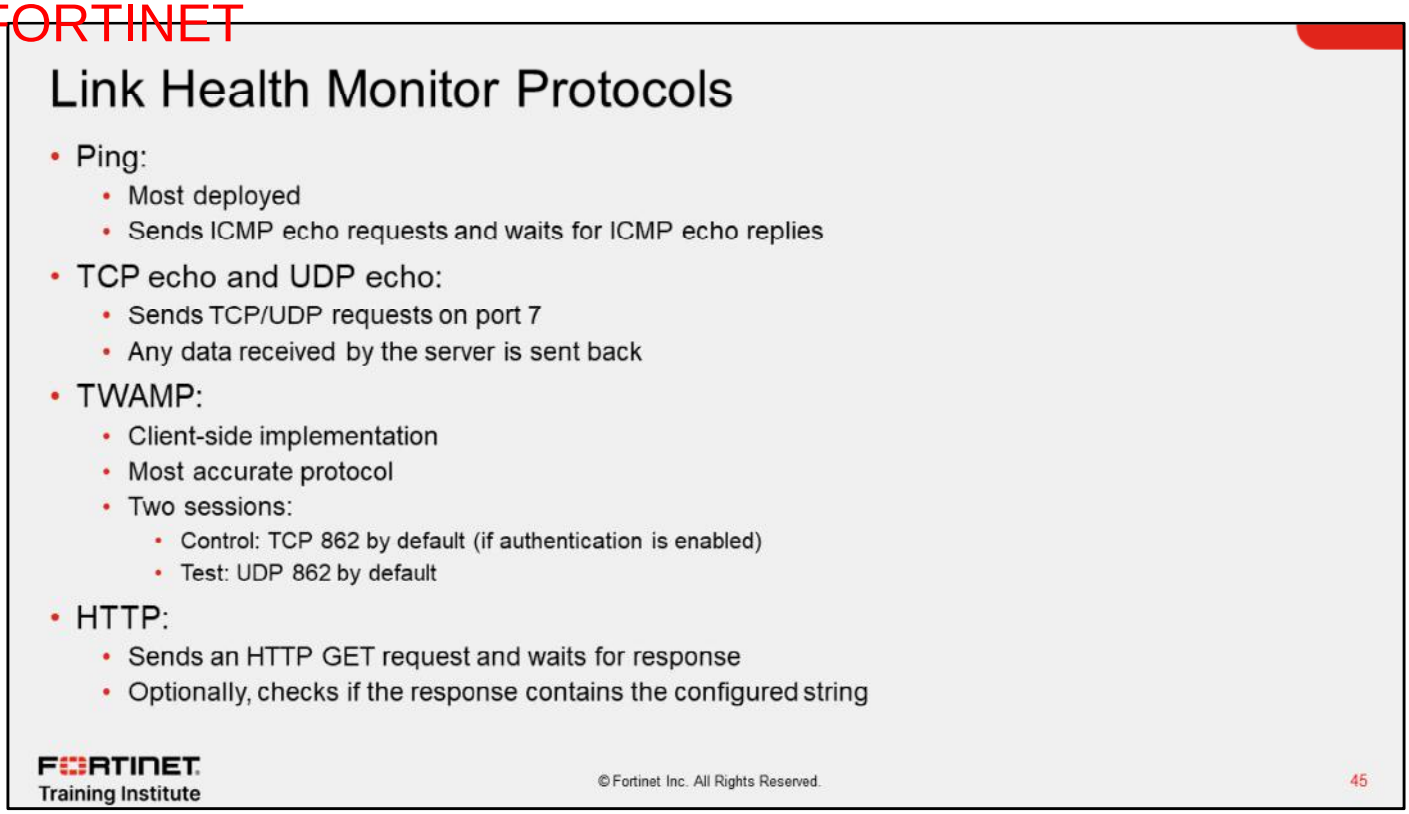

This slide describes the probe protocols supported by link health monitor.

Ping is the most used network monitoring protocol because it is supported by virtually all network devices. When you use ping, FortiGate sends ICMP echo requests to the configured target servers and waits for the respective ICMP echo replies. Because some ISPs and content providers block or limit ICMP traffic on their network, you may want to switch to TCP echo, UDP echo, or TWAMP.

When you use TCP echo and UDP echo, FortiGate sends periodic packets to the configured target servers, which are listening for connections on port 7 for both TCP and UDP. Upon reception of the packets, the server sends back an identical copy of the data it received from FortiGate.

Two-Way Active Measurement Protocol (TWAMP) is the most accurate protocol among the five. Link health monitor uses the client-side implementation of TWAMP. There are two sessions used in TWAMP: control and test. The former is used to authenticate the endpoints, and the latter to exchange packets used to measure the performance. Note that if authentication is disabled—it is disabled by default—FortiGate generates the test session only. FortiGate uses port 862 as default port for both control and test sessions, but you can configure a different port.

When you configure HTTP as the protocol, FortiGate sends periodic HTTP GET requests to the target server, and then waits for a response. Optionally, you can configure FortiGate to check if the response contains a specific string in the HTML content.

Link Health Monitor Actions

Action	Dead	Alive	Effect during dead state
Update static route*	Flag associated static routes as inactive	Flag associate static routes as active	Static routes are removed from routing table
Update policy route**	Disable associated policy routes	Re-enable associated policy routes	Policy routes are skipped
Update cascade interface***	Bring down alert interfaces	Bring back up alert interfaces	Route LAN-originated traffic to a different device

* Associated static routes match the configured gateway and interface in the link health monitor settings
** Associated policy routes match the configured gateway and interface in the link health monitor settings
*** Require the configuration of alert interfaces (usually, your LAN-facing interfaces)

FORTINET Training Institute

46

This slide describes the actions taken by link health monitor when the state of an interface changes from alive to dead, and vice-versa. All three actions are enabled by default.

When you enable update static route and link health monitor detects an interface as dead, FortiGate marks the associated static routes—those matching the configured gateway and interface—as inactive. The result is that the inactive static routes are removed from the routing table. The absence of such routes can then force FortiGate to redirect the traffic to other valid routes, if any. Note that this action applies to static routes only.

The update policy route action works the same as the update static route action, except that instead of marking the associated static routes as inactive after an interface is detected as dead, FortiGate disables the associated policy routes. For that, FortiGate checks the policy route table and disables the policy routes whose outgoing interface and gateway match the configured interface and gateway in the link health monitor settings. Like the update static route action, the goal is for FortiGate to skip the disabled policy route during the route lookup process, so the traffic matches another policy route or FIB route in the system.

The update cascade interface action requires you to configure one or more alert interfaces. FortiGate then brings down the alert interfaces after the monitoring interface is detected dead. The goal is to force the traffic from networks behind the alert interfaces to be routed through a different device after an important interface, such as the internet-facing interface, is dead, which could mean that FortiGate is unable to forward traffic to the WAN. For example, if you are using dynamic routing or Virtual Router Redundancy Protocol (VRRP) on your LAN interface, which is configured as an alert interface, then bringing down the interface can trigger a routing failover to a backup gateway.

If FortiGate detects the interface as alive again, it reverts any action taken so far for the link. That is, FortiGate restores static routes, re-enables policy routes, and brings back up alert interfaces.

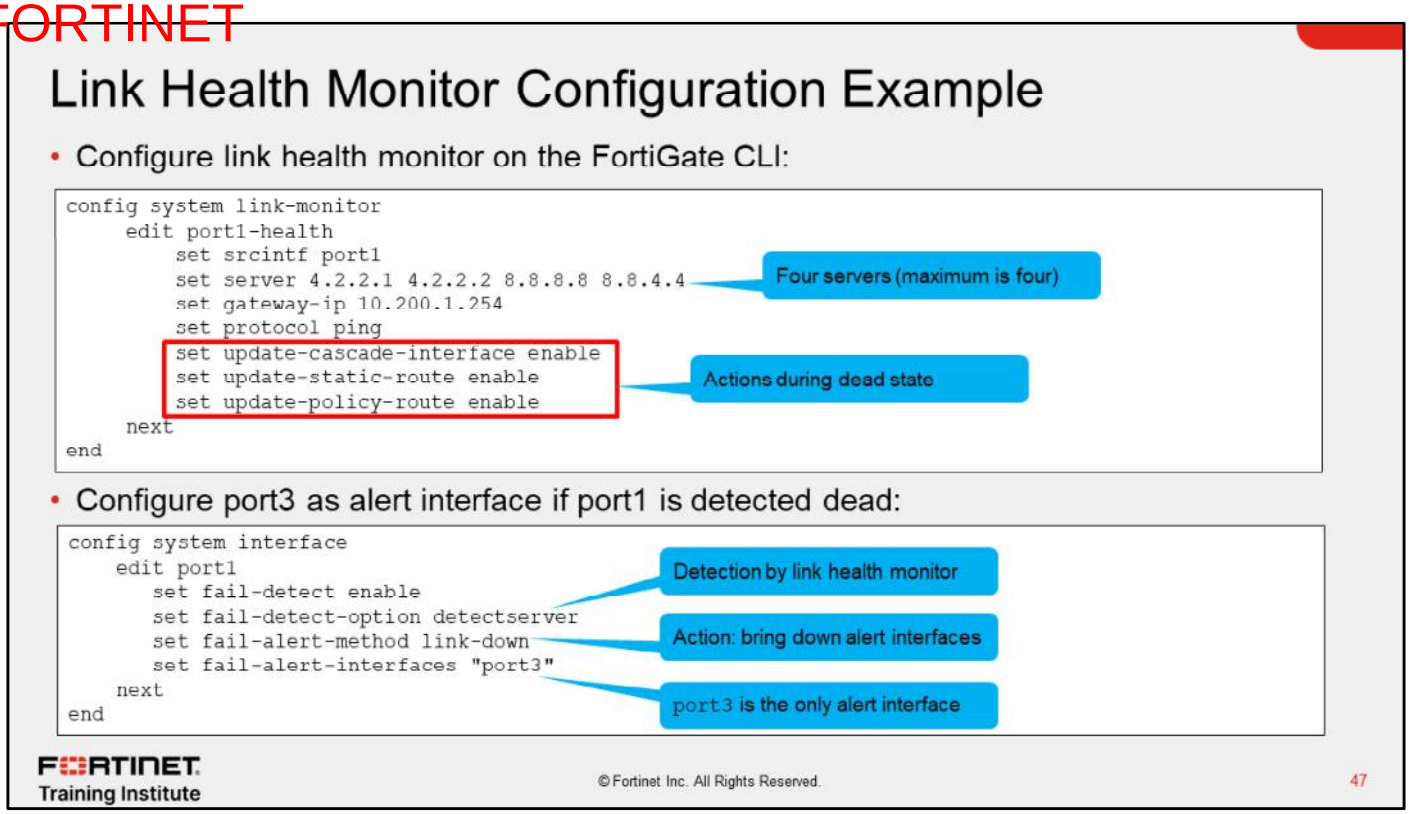

This slide shows a configuration example for link health monitor. FortiGate monitors the health of port1 against Level3 and Google DNS servers (four in total). For sending the probes, FortiGate uses `10.200.1.254` as gateway and ping as protocol.

When the state of port1 changes, FortiGate updates cascade interfaces, static routes, and policy routes. For the update cascade interface action to work, you must configure the alert interfaces. This slide also shows an example of the alert interface configuration required on the monitoring interface (port1). The configuration instructs FortiGate to bring down port3 if port1 is detected dead by the link health monitor feature.

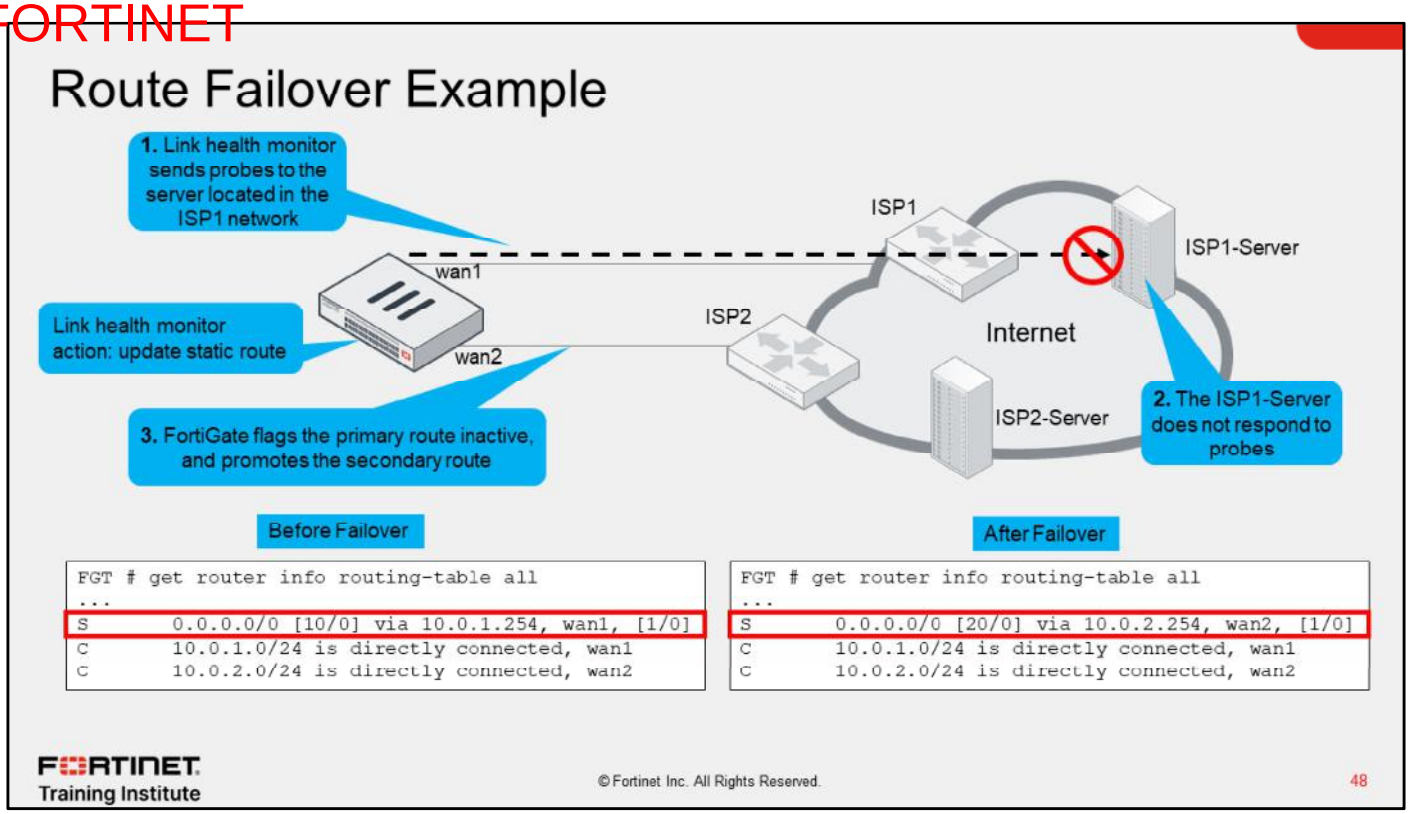

In the example shown on this slide, FortiGate has two internet connections. wan1 is connected to ISP1, and wan2 to ISP2. Within each ISP network, there is a server that FortiGate sends probes to for link health monitoring purposes. For link health monitor, the update static route action is enabled. The administrator configured two static default routes, one through wan1 and the other wan2. The static default routes are assigned a distance of 10 and 20, respectively.

Before failover, the default route over wan1 is installed in the routing table because it has a lower distance, and the default route through wan2 is present in the routing table database as a standby route. The link health monitor sends probes to ISP1-Server located within the ISP1 network through wan1. When FortiGate detects five consecutive failed probes for ISP1-Server, FortiGate flags the default route over wan1 as inactive, which results in the route being removed from the routing table. This also results in the standby default route through wan2 to be installed in the routing table. Then, FortiGate starts using the new default route to route traffic to the internet.

The example shown on this slide makes use of different distance values to control the primary and standby routes. The result is that one default route only is installed in the routing table at any time. In case you always need to have both routes installed in the routing table, you can configure the same distance on both routes, but different priorities. You assign a lower priority number to your primary route, and a higher priority number to your standby route. Having both routes in the routing table is required if you use the interfaces to terminate IPsec VPN tunnels and you want to speed up failover by ensuring the tunnel over the secondary ISP link is already up before failover.

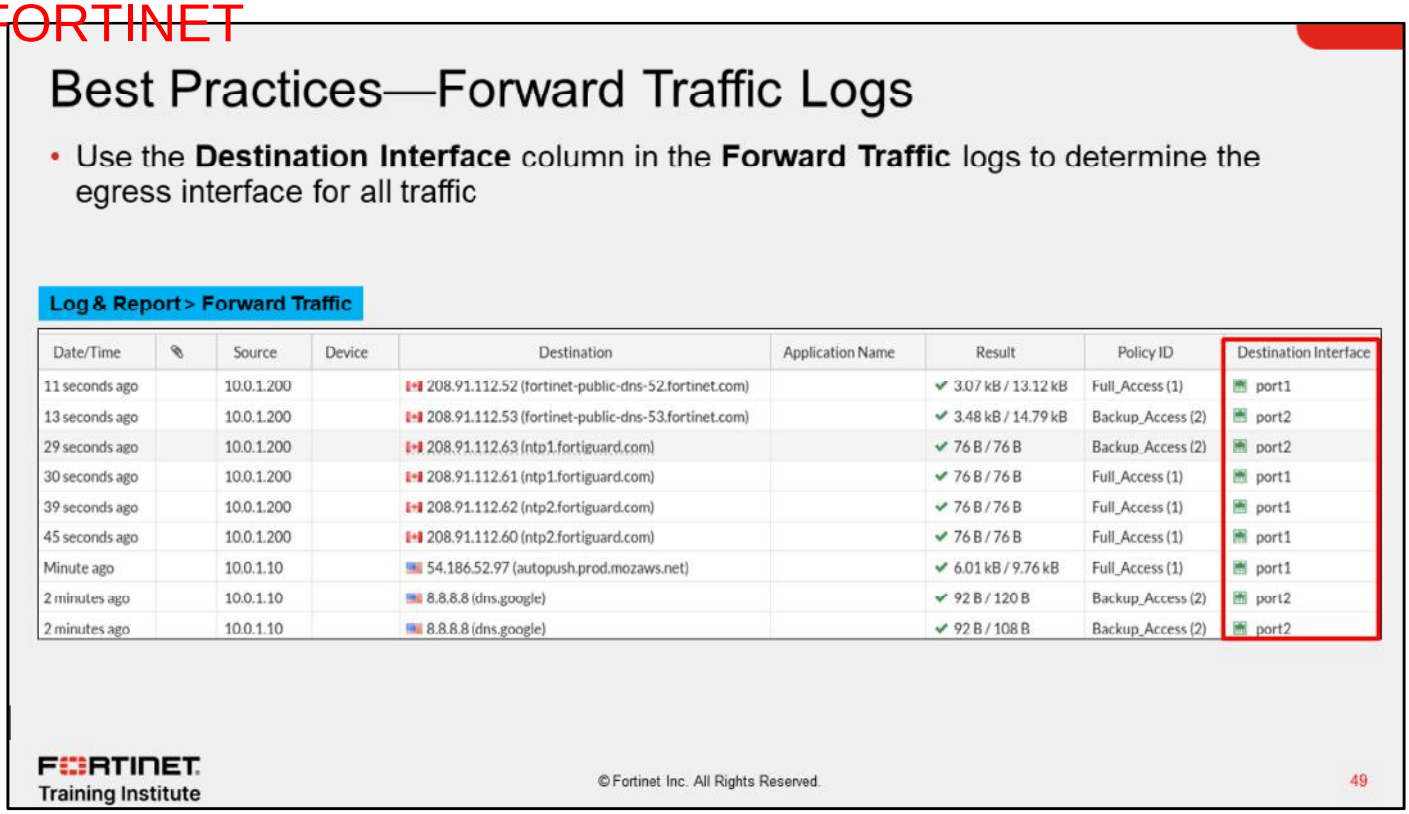

If you enable the **Destination Interface** column in the **Forward Traffic** logs, you can view the egress interface for traffic passing through your FortiGate device. You can use this information to determine which route is applied to which traffic stream, as well as identify any routing configuration issues.

If your firewall policies do not have any security profiles applied, you should enable logging for all sessions in your policies; otherwise, FortiGate does not generate any **Forward Traffic** logs. Use this feature with some caution, since enabling all sessions logging can generate a lot of logs if the firewall policy is handling a high volume of traffic. You should enable it when necessary, and disable it immediately afterwards.

Knowledge Check

1. What is the purpose of the link health monitor setting `update-static-route`?
   - A. It creates a new static route for the backup interface.
   - ✓ B. It removes all static routes associated with an interface detected as dead by the link health monitor.

2. When using link health monitoring, which route attribute can you configure to achieve route failover protection?
   - ✓ A. Distance
   - B. Metric

FORTINET Training Institute

50

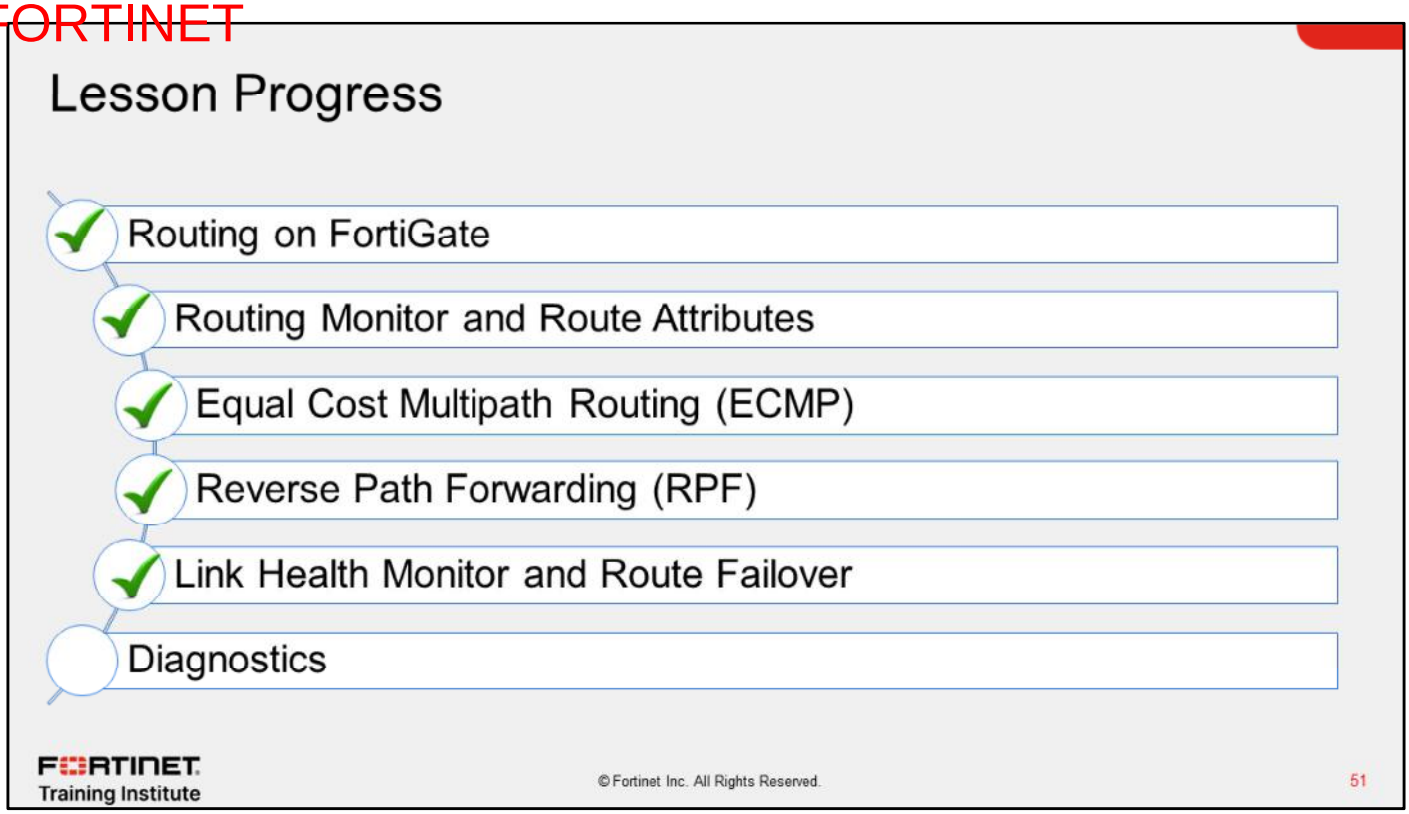

Good job! You now understand the link health monitor and route failover.

Now, you will learn about routing diagnostics.

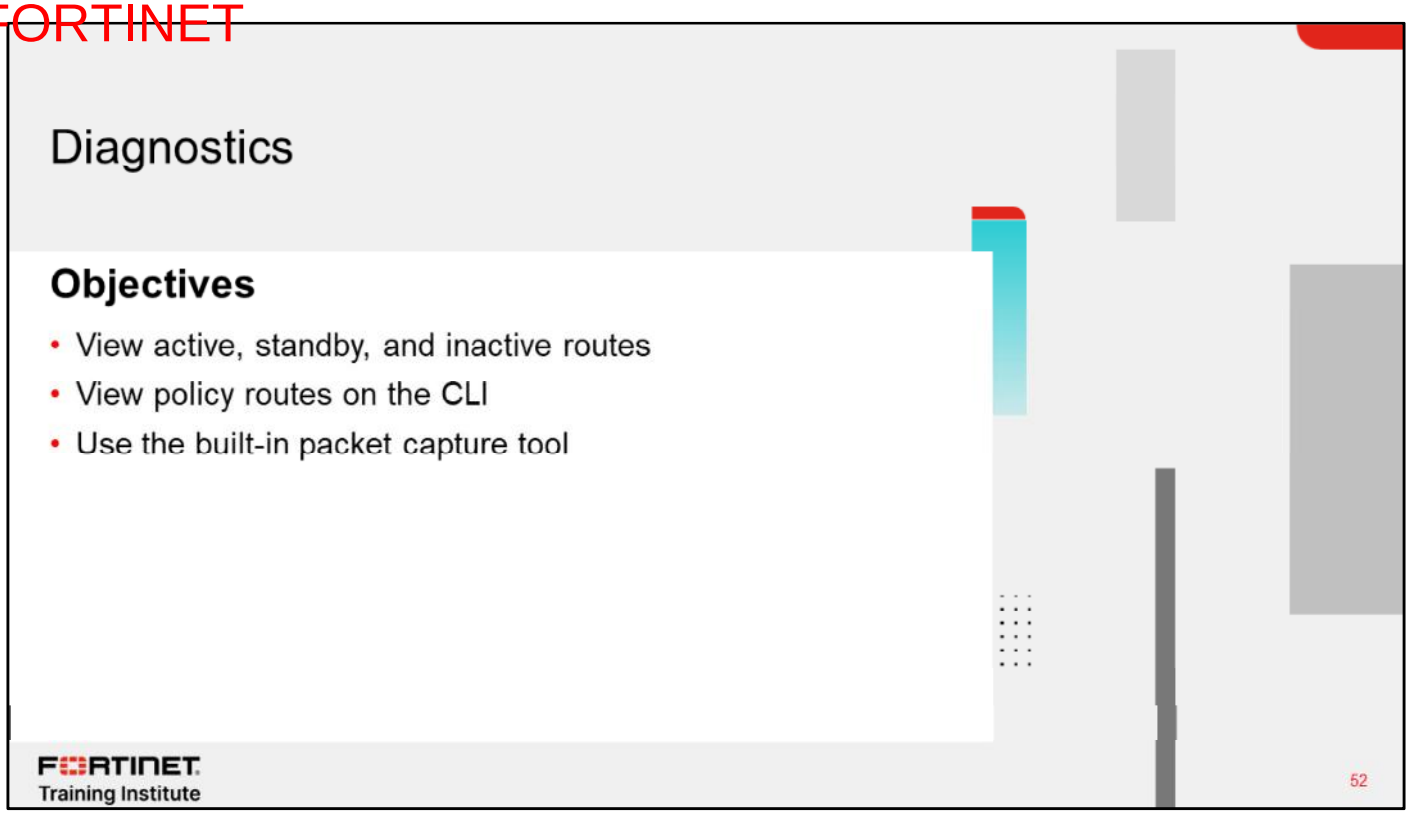

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in routing diagnostics, you should be able to view the entries in the routing table and routing table database, as well as to identify how packets flow across FortiGate.

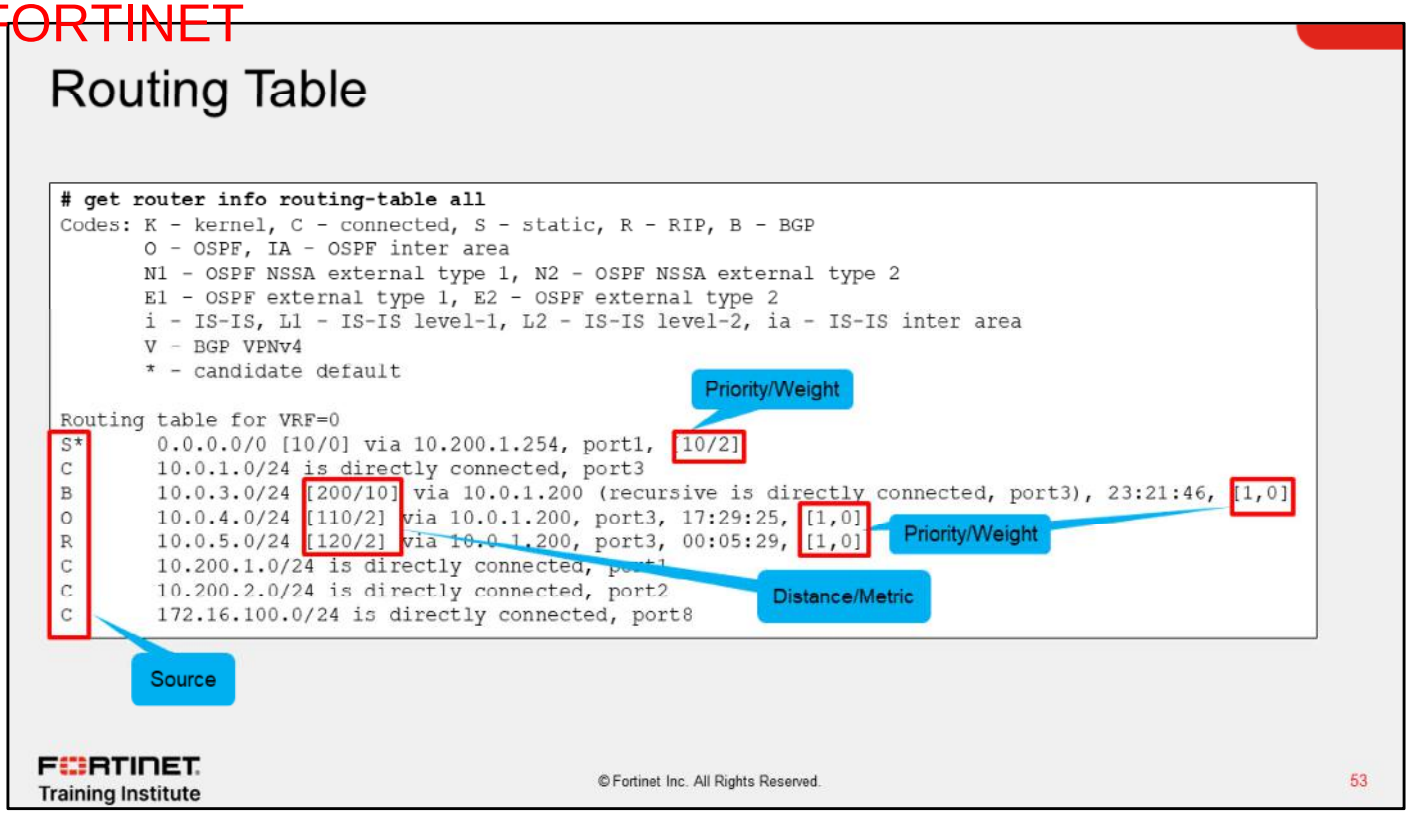

The CLI command shown on this slide displays all entries in the routing table. The routing table displays the routes that make it to the FIB. That is, the best active routes to a destination.

The left-most column indicates the route source. Route attributes are shown inside square brackets. The first number, in the first pair of attributes, is distance, which applies to both dynamic and static routes. The second number is metric, which applies to dynamic routes only.

Static routes and dynamic routes also have priority and weight attributes, which are shown as the last pair of attributes for the respective route. In the case of dynamic routes, the weight is always zero.

This command doesn't show standby or inactive routes, which are present in the routing table database only. For example, when two static routes to the same destination subnet have different distances, the one with the lower distance is installed in the routing table, and the one with the higher distance in the routing table database.

Routing Table Database

```
# get router info routing-table database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       V - BGP VPNv4
       > - selected route, * - FIB route, p - stale info

Routing table for VRF=0
S     *> 0.0.0.0/0 [10/0] via 10.200.1.254, port1, [1/10]
S        0.0.0.0/0 [20/0] via 10.200.2.254, port2, [5/0]
S        8.8.8.8/32 [10/0] via 172.16.100.254, port8 inactive, [1/0]
O        10.0.1.0/24 [110/1] is directly connected, port3, 00:05:47, [1/0]
C     *> 10.0.1.0/24 is directly connected, port3
O        10.0.2.0/24 [110/1] is directly connected, port4, 00:05:47, [1/0]
C     *> 10.0.2.0/24 is directly connected, port4
B     *> 10.0.3.0/24 [200/10] via 10.0.1.200 (recursive is directly connected, port3), 00:05:40, [1/0]
O     *> 10.0.4.0/24 [110/2] via 10.0.1.200, port3, 00:05:27, [1/0]
B        10.0.4.0/24 [200/10] via 10.0.1.200 (recursive is directly connected, port3), 00:05:40, [1/0]
C     *> 10.200.1.0/24 is directly connected, port1
C     *> 10.200.2.0/24 is directly connected, port2
```

Active route

Inactive route

Standby route

FORTINET Training Institute

54

If you want to view active, standby, and inactive routes, use the CLI command shown on this slide to display the routing table database entries.

In the example on this slide, the command shows two standby routes, one static and the other BGP. Both standby routes are standby because there are better routes—lower distance—to the same destination. The better routes show an asterisk next to the route source to indicate they are FIB entries, and therefore, are used for routing traffic.

The output also shows one inactive route. Routes are marked as inactive where the corresponding interface is administratively down, has its link down, or when the interface is detected dead by link health monitor and the update static route action is enabled.

Policy Route Table

```
# diagnose firewall proute list
list route policy info(vf=root):

id=1 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0-0 iif=7 dport=0-65535
path(1) oif=21(T_MPLS_0)
source(1): 10.0.1.0-10.0.1.255
destination(1): 10.0.0.0-10.255.255.255
hit_count=18 last_used=2022-02-23 05:47:21

id=2113929223 static_route=7 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0-0
iif=0 dport=1-65535 path(1) oif=3(port1) gwy=192.2.0.2
source wildcard(1): 0.0.0.0/0.0.0.0
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(1): Fortinet-FortiGuard(1245324,0,0,0)
hit_count=0 last_used=2022-02-23 06:39:07

id=2130903041(0x7f030001) vwl_service=1(Critical-DIA) vwl_mbr_seq=1 2 dscp_tag=0xff 0xff flags=0x0
tos=0x00 tos_mask=0x00 protocol=0 sport=0-65535 iif=0 dport=1-65535 path(2) oif=3(port1) oif=4(port2)
source(1): 10.0.1.0-10.0.1.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(3): GoToMeeting(4294836966,0,0,0, 16354)
Microsoft.Office.365.Portal(4294837474,0,0,0, 41468) Salesforce(4294837976,0,0,0, 16920)
hit_count=0 last_used=2022-02-23 05:46:43
```

This is a regular policy route (ID ≤ 65535)

This is an ISDB route (ID > 65535 and no vwl_service field)

This is an SD-WAN rule (ID > 65535 and the vwl_service field is present)

FORTINET Training Institute

55

FortiOS maintains a policy route table that you can view by running the `diagnose firewall proute list` command.

There are three types of policy routes displayed in the policy route table: regular policy routes, ISDB routes, and SD-WAN rules. Follow these rules to identify each type of policy route in the table:

- Regular policy routes are assigned an ID no higher than 65535. In the output shown on this slide, the first entry is assigned ID 1, which makes it a regular policy route.
- ISDB routes and SD-WAN rules are assigned an ID higher than 65535. However, SD-WAN rule entries include the `vwl_service` field, and ISDB route entries don't. The `vwl_service` field indicates the ID and the name of the rule from the SD-WAN configuration perspective. In the output shown on this slide, the second entry is an ISDB route and the third entry an SD-WAN rule.

Note that although IDs for regular policy routes are in the 1 to 65535 range, the maximum number of regular policy routes that you can configure are much lower and varies among models. For example, you can configure up to 512 regular policy routes in a FortiGate 300D device. For more information about the maximum supported values per model, refer to the `FortiOS Maximum Values Table` on `docs.fortinet.com`. Alternatively, you can run the `print tablesize` command on the FortiGate CLI to get the maximum values for your device.

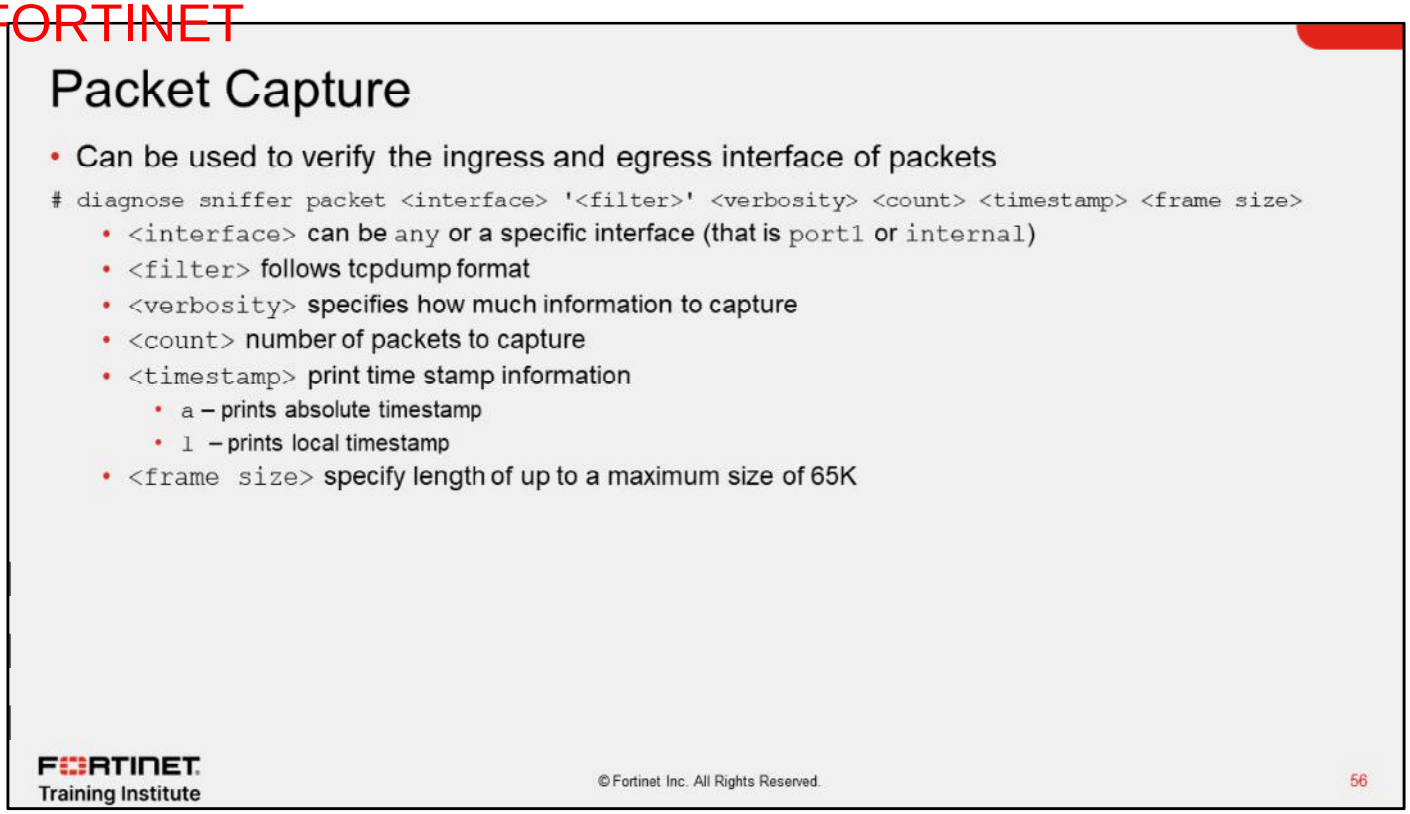

Packet captures, or *sniffers*, are one of the most useful sources of information for debugging routing problems. FortiGate includes a built-in traffic sniffer tool. You can use it to verify the ingress and egress interfaces of packets as they pass through. You can run the built-in sniffer from either the GUI or the CLI. The syntax of the CLI command is shown on this slide.

The `<interface>` option is the name of the physical or logical interface to run the sniffer on. Most of the times, you want to indicate `any` to capture packets on all interfaces. This enables you to see how packets flow across the different interfaces. Another option is to indicate the name of the interface, which is useful when you want to narrow down the packet capture to that interface. Indicating the name of the interface is also required if you want the tool to capture the MAC address information. That is, when you use `any`, the sniffer doesn't capture the real MAC addresses used by the packet.

The filter follows the Berkeley Packet Filter (BPF) syntax used by the well-known tcpdump tool. You should configure specific filters to ensure you're only capturing what you need. You can also specify a `<count>` value to automatically stop the sniffer after capturing a specific number of packets. Otherwise, the sniffer continues capturing packets until you manually stop it using `Ctrl + C`. You can use the `<time stamp>` option to print the time stamp information. Use `a` to print the absolute time stamp, or `l` (lowercase L) to print the local time-zone based time stamp. Time stamp information is particularly useful when correlating sniffer output to debug flow messages. You will learn more about debug flow in another lesson.

By default, the sniffer uses the MTU configured on the interface to limit the packet length during the capture. Using the `<frame size>` argument, you can specify a length larger or smaller than the interface MTU. Note that f you use the `any` interface, the sniffer will default to 1600 bytes.

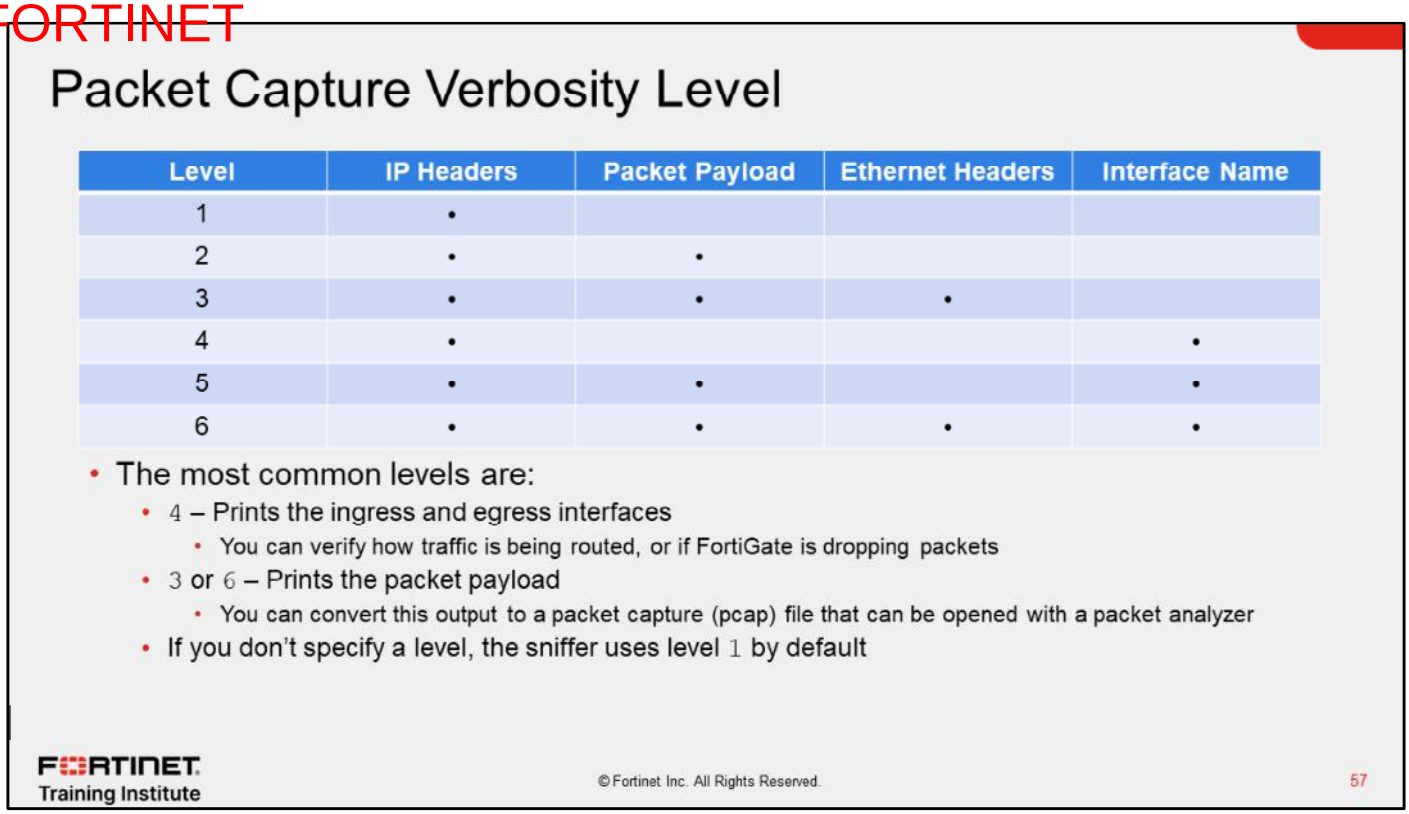

The verbosity level specifies how much information you want to display. There are six different levels and this table shows which ones display the IP headers, packet payload, Ethernet headers, and interface names.

Use verbosity level `4` to take a quick look at how the traffic is flowing through FortiGate (if packets are arriving and how FortiGate is routing them out). You can also use level 4 to check if FortiGate is dropping packets.

Verbosity levels `3` and `6` provide the most output. Both show the IP payloads and Ethernet headers. You can save the output and export it to a packet capture (pcap) file using a Perl script. The pcap file can then be opened with a packet analyzer, such as Wireshark, for further investigation. You can locate the Perl script that converts the sniffer output to pcap on the Fortinet Knowledge Base website (`kb.fortinet.com`).

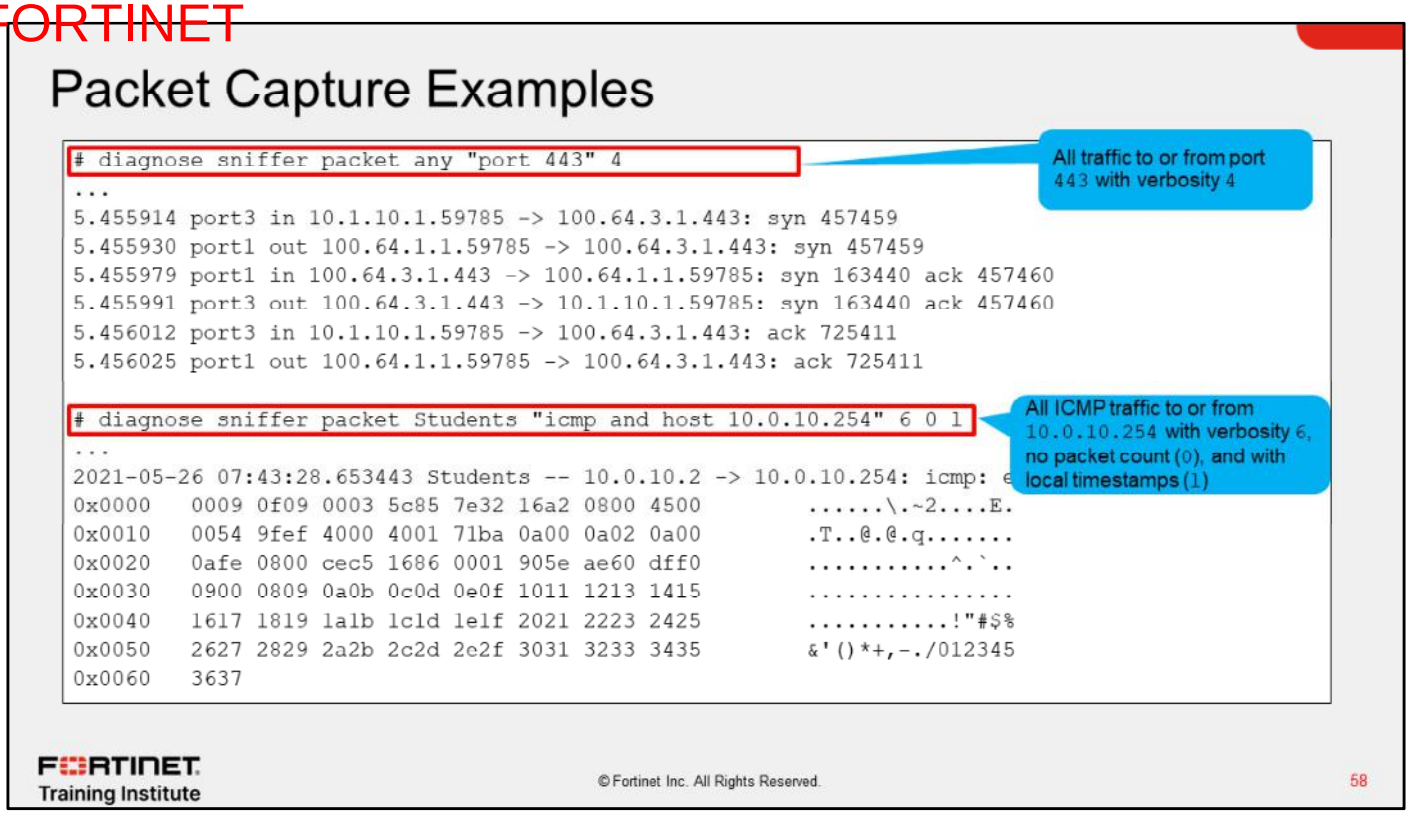

This slide shows two examples of packet capture outputs.

The first example captures all traffic to and from port `443`. It uses verbosity `4`, so the information is easy to read. It displays one packet per line, containing the incoming and outgoing interface, IP addresses, port numbers, and type of packet (`SYN`, `SYN/ACK`, and so on). Note that the interface is set to `any`, which is useful to capture packets that enter or exit multiple interfaces in the device. This enables you to have a better understanding of how packets flow through the firewall. For example, the output shows a three-way handshake established across FortiGate. From the packet capture, you can conclude that the connection is initiated by `10.1.10.1`, which is behind `port3`, and is destined to `100.64.3.1`, which is behind `port1`. You can also conclude that FortiGate performs SNAT for the connection. That is, in the original direction, FortiGate translates the source address to `100.64.1.1` when packets leave `port1`. FortiGate then translates the reply packets back to `10.1.10.1` when they exit `port3`.

The second example captures all ICMP traffic coming from or going to `10.0.10.254`. Unlike the first example, which captures packets on any interface, this example limits the capture to packets that enter or leave the `Student` interface. Although not shown on this slide, the `Students` interface is a VLAN interface. In addition, the verbosity level is set to `6`, which includes the full packet IP payload details. The output is longer and more difficult to read. However, this is one of the two verbosity levels to use (`3` being the other one) if you need to export the output to pcap format. You can then view the pcap file using Wireshark or any other compatible packet analysis tool. Moreover, the additional arguments in the command instruct the sniffer to not set a packet count limit (`0`) and to print the local timestamp for each packet (`1`).

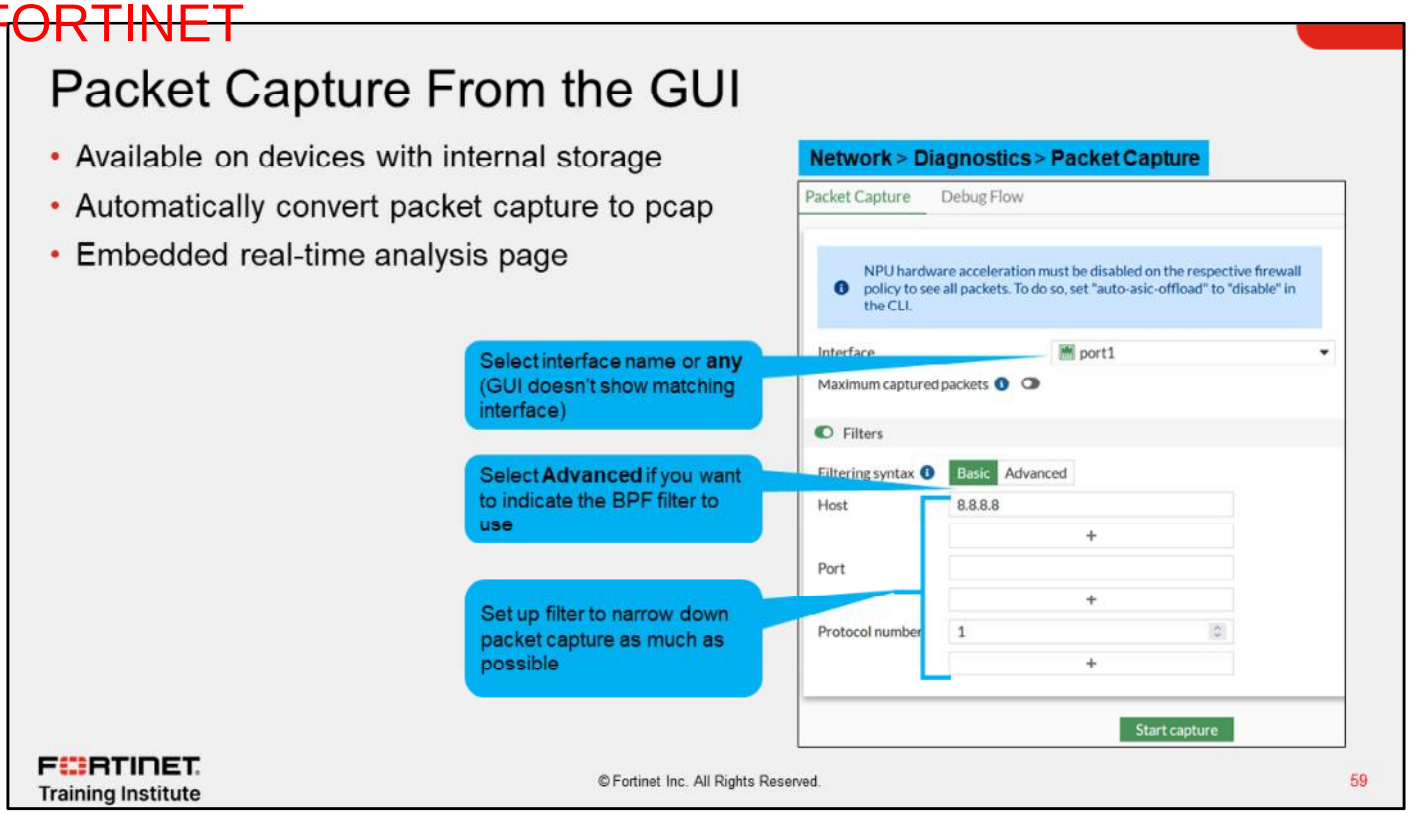

If your FortiGate model has internal storage, you can capture packets on the GUI. Starting FortiOS 7.2, the GUI packet capture tool was improved to also include a real-time analysis tool that enables you to examine the packet capture details directly on the GUI. You also download the respective pcap file in case you prefer to review it using Wireshark or your preferred packet analysis tool.

Before starting the packet capture, you should set up the packet capture filter by using either **Basic** or **Advanced** filter options. When you choose **Basic**, you to indicate basic filter options such as host address, port number, and protocol number. In case you want to use your own BPF filter like you do in the CLI, you can choose **Advanced**.

Regardless of which method you use (CLI or GUI), packet capture filters should be very specific to make sure only the relevant packets are captured, and large amounts of data are not being written to the disk.

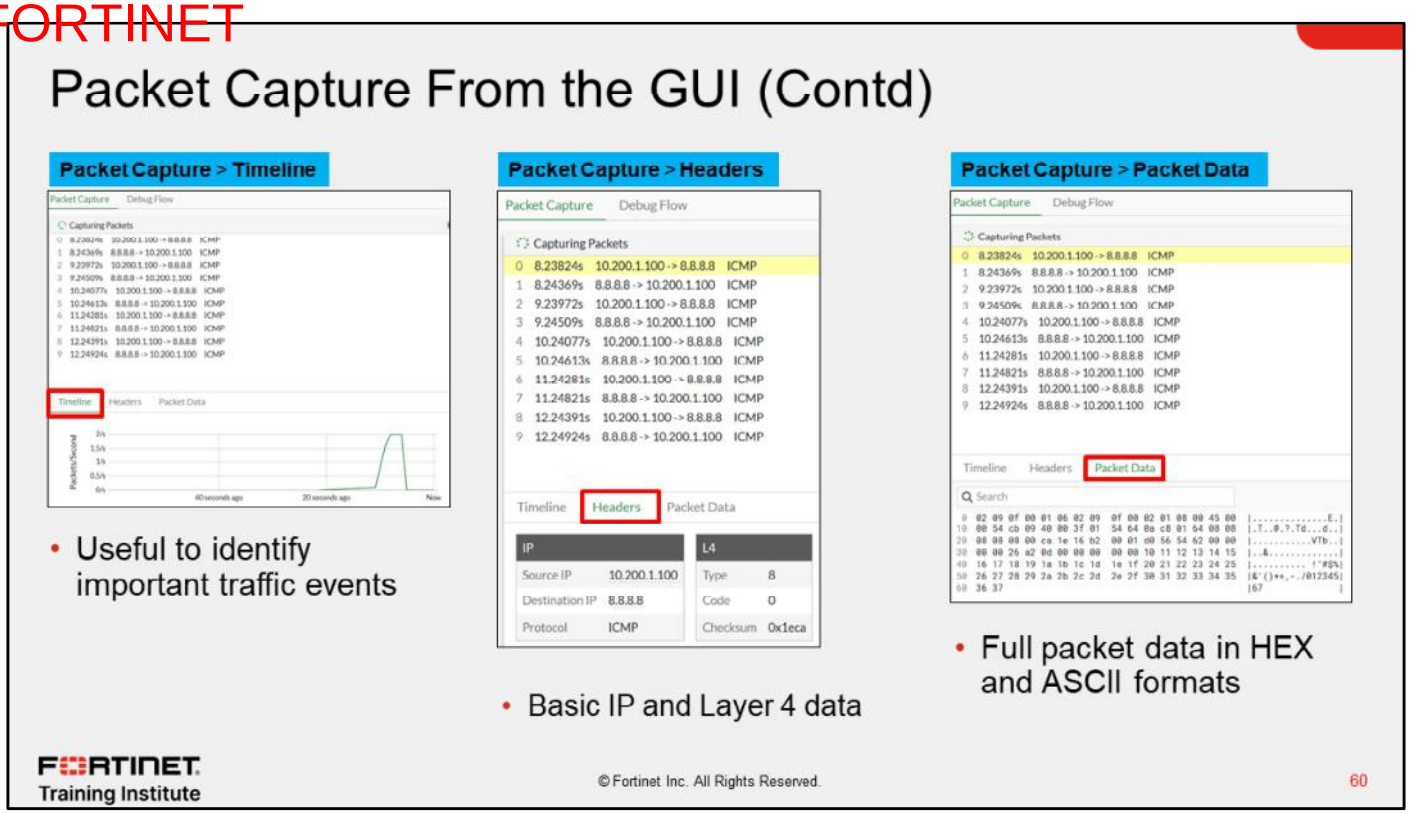

This slide shows an example of the embedded real-time analysis tool included in the GUI packet capture tool starting FortiOS 7.2. After you start the packet capture, the GUI starts displaying the captured packets based on the filter set.

The **Timeline** tab displays a graph with the number of captured packets per second. The graph is useful to quicky identify peaks of traffic related by important events in the network.

The **Headers** tab enables you to examine basic IP (Layer 3) and Layer 4 information on the packet.

The **Packet Data** tab enables you to examine the full packet data using hexadecimal format. Next to the hexadecimal packet data, FortiOS displays the equivalent output in ASCII format.

Knowledge Check

1. What is the distance value for this route?

```
10.200.2.0/24 [110/2] via 10.200.2.254, [25/0]
```

✓A. 110
B. 2

2. Which CLI command can you use to view standby and inactive routes?

A. `get router info routing-table all`
✓B. `get router info routing-table database`

3. Which CLI packet capture verbosity level prints interface names?

A. 3
✓B. 4

FORTINET Training Institute

61

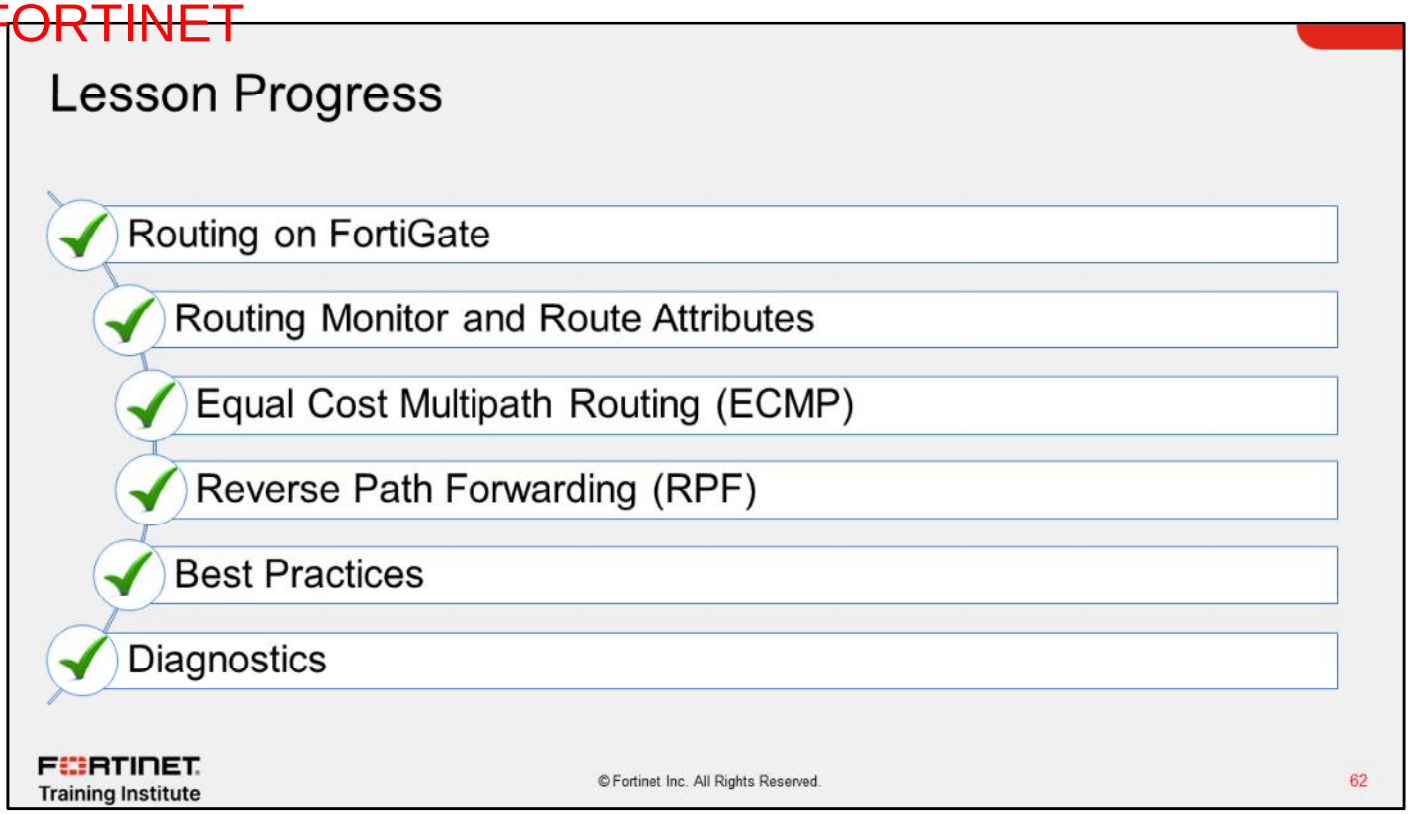

Congratulations! You have completed this lesson.

Now you will review the objectives that you covered in this lesson.

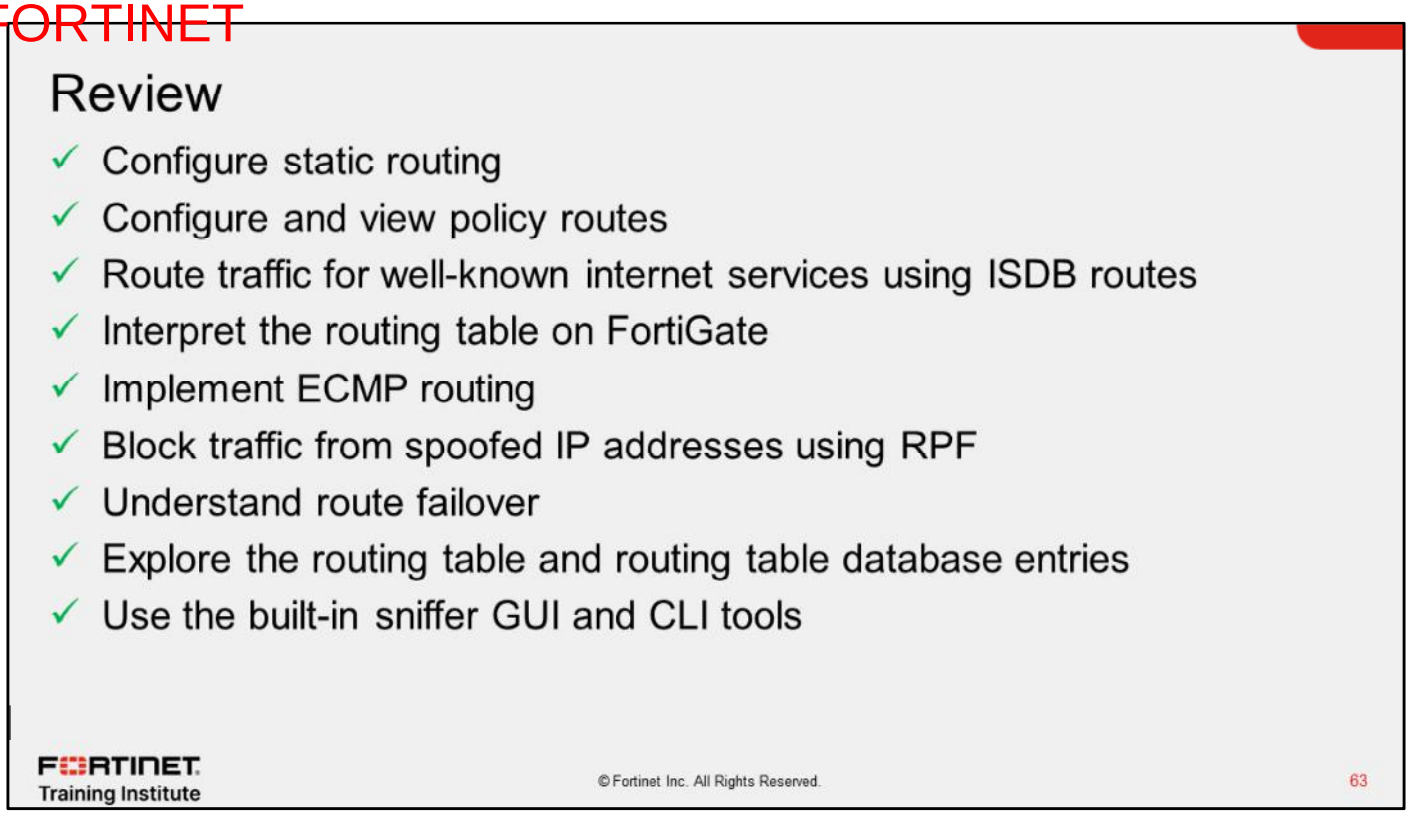

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure, maintain, and troubleshoot the FortiGate routing configuration.

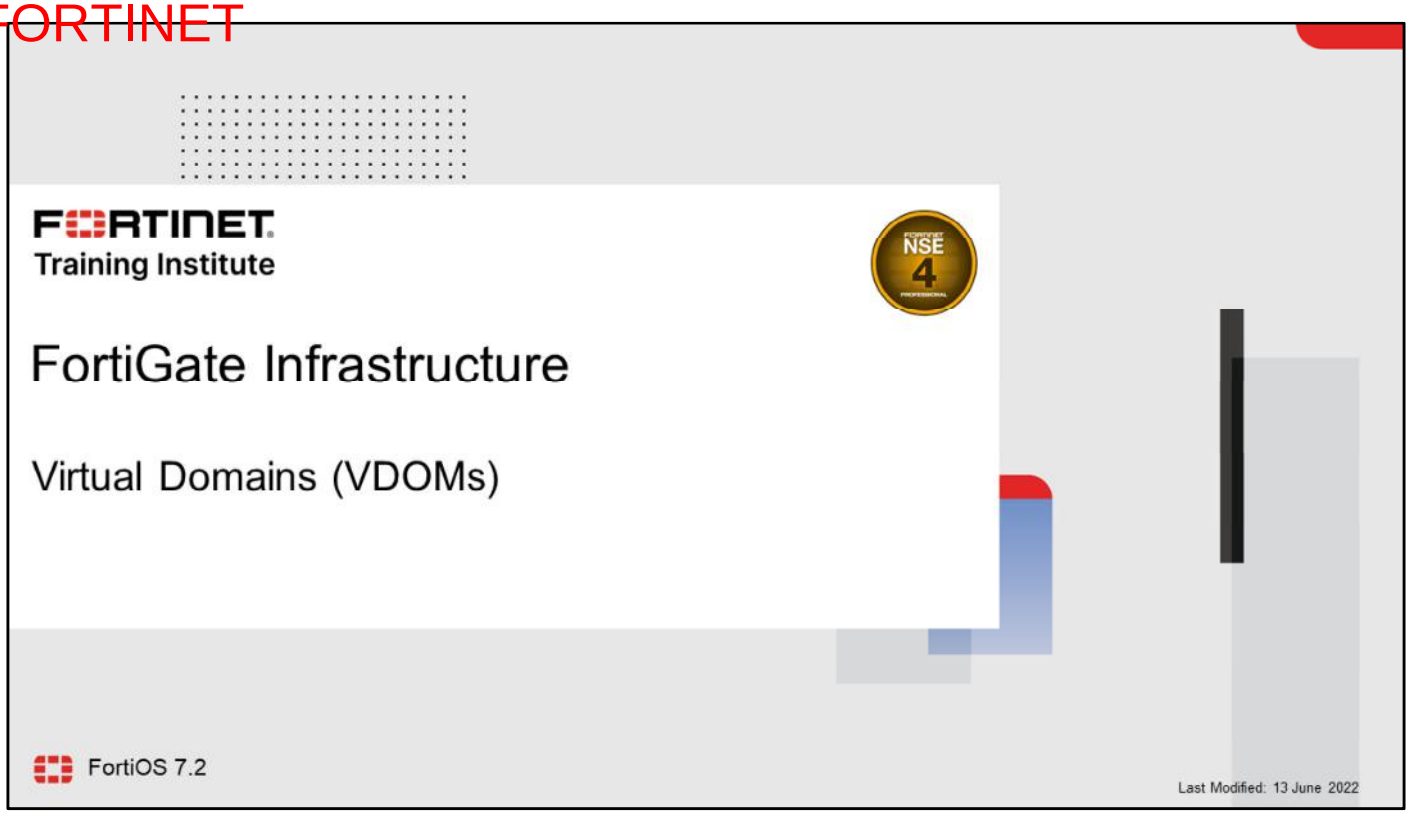

In this lesson, you will learn how to configure VDOMs, and examine examples of common use.

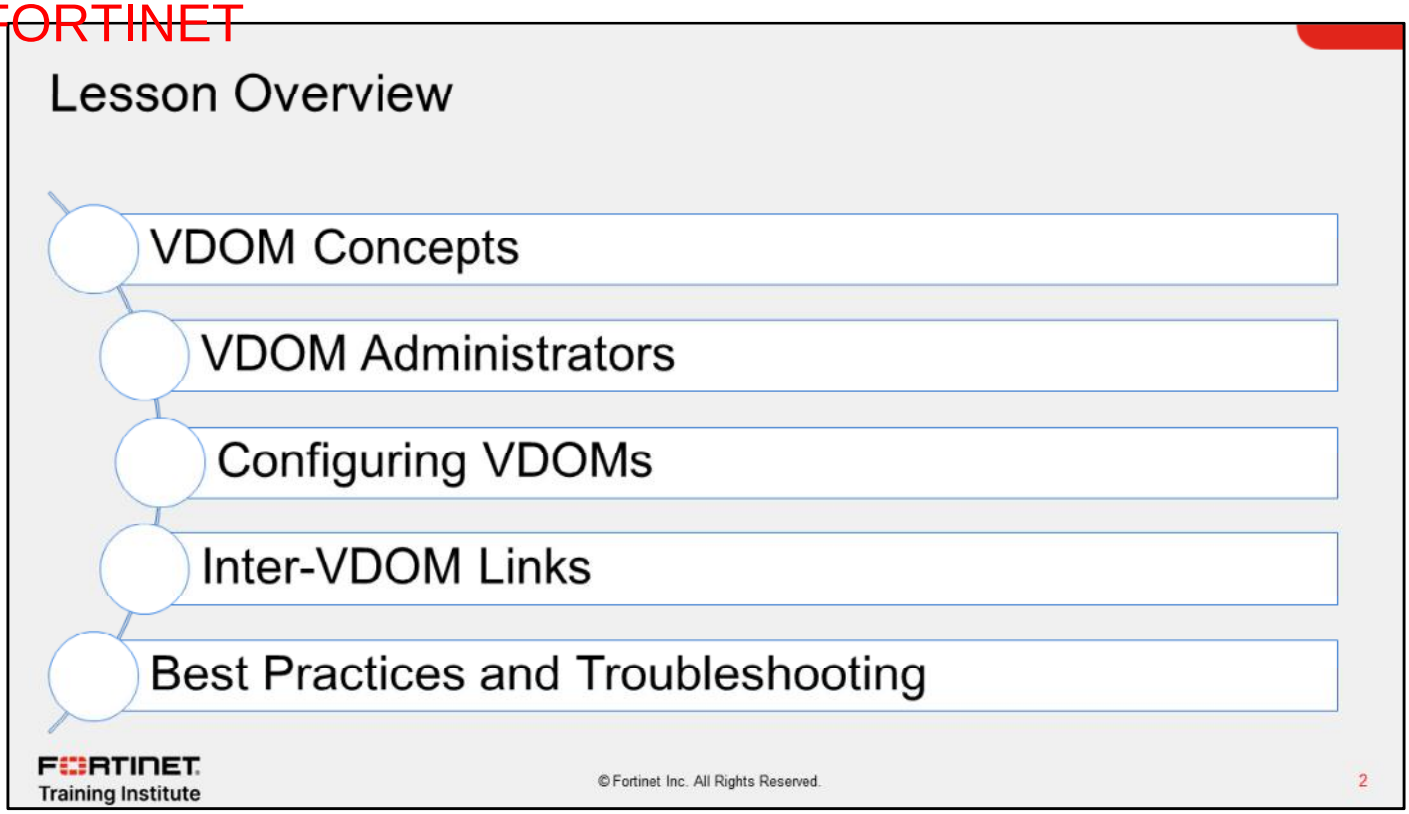

In this lesson, you will learn about the topics shown on this slide.

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in VDOMs, you will be able to understand the key benefits and use cases for VDOMs.

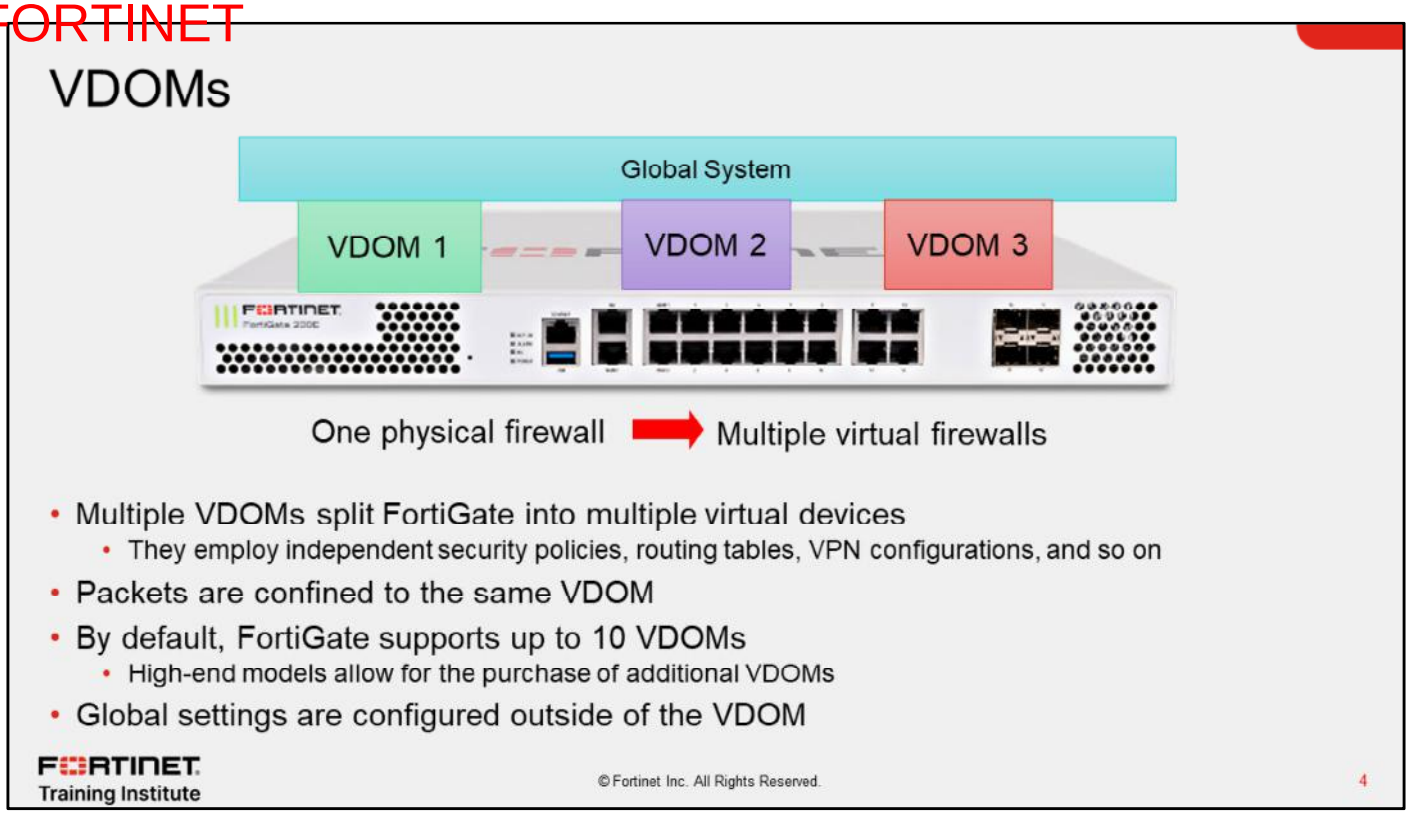

What if a campus wants to keep its departments separate? A datacenter wants to implement various security implementations in a cost-effective manner that maintains all customer traffic separate and secure while also reducing space and making configuration easier? What if you want to segment your network, and subdivide policies and administrators into multiple security domains?

The best solution is to enable FortiGate VDOMs.

A VDOM splits your FortiGate into multiple logical devices and divides one security domain into multiple security domains.

Each VDOM has independent security policies and routing tables. Also, and by default, traffic from one VDOM cannot go to a different VDOM. This means that two interfaces in different VDOMs can share the same IP address, without any overlapping subnet problems.

When you use VDOMs, a single FortiGate device becomes a virtual data center of network security, UTM inspection, and secure communication devices.

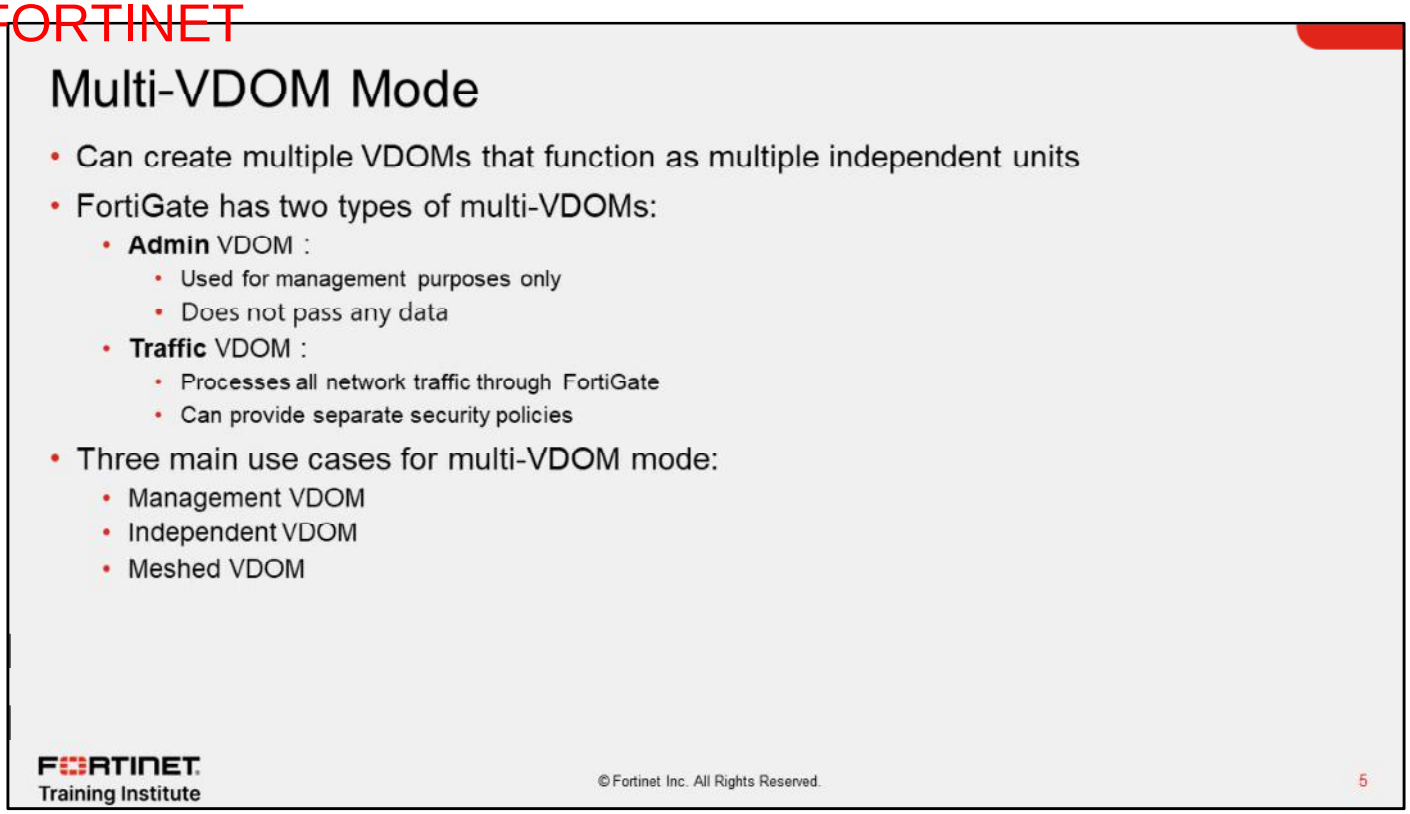

Use multi-VDOM mode when you want to create multiple logical firewalls from a single FortiGate. Each VDOM acts as an independent FortiGate.

Multi-VDOM mode works well for managed service providers leveraging multi-tenant configurations, or large enterprise environments that desire departmental segmentation. You can give each individual tenant or department, visibility and control of their VDOM, while keeping other VDOMs independent and unseen.

Two types of VDOMs can created in multi-VDOM Mode: An admin VDOM and a traffic VDOM.
Admin VDOMs are for FortiGate administration, and traffic VDOMs permit traffic to travel through FortiGate.

Upon upgrade, if a FortiGate is in split-vdom mode, it is converted to multi-vdom mode. The FG-traffic VDOM becomes a traffic type VDOM. The root VDOM becomes an admin VDOM.

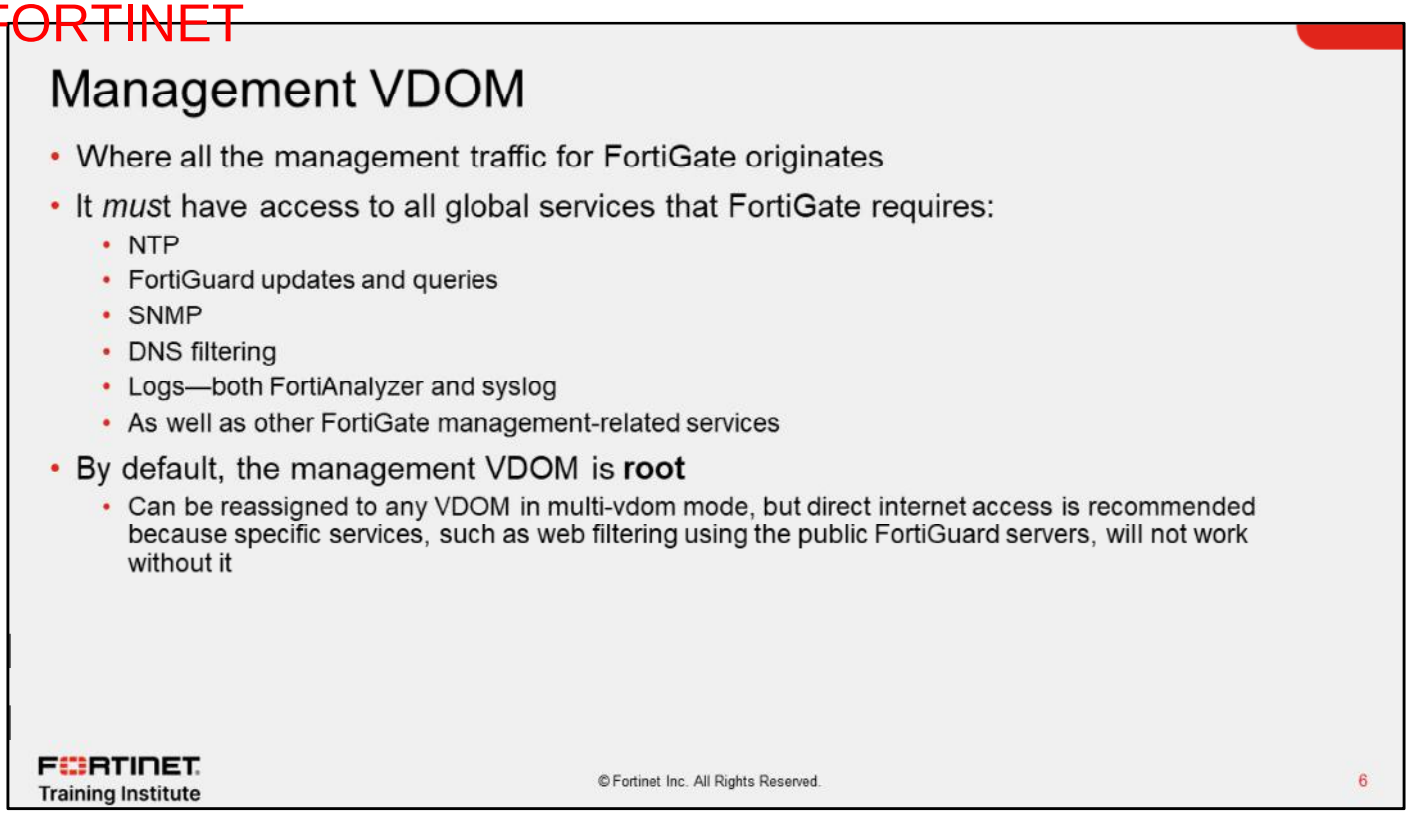

Until now, you've learned about traffic passing *through* FortiGate, from one VDOM to another.

What about traffic originating *from* FortiGate? Some system daemons, such as NTP and FortiGuard updates, generate traffic coming from FortiGate.

Traffic coming from FortiGate to those global services originates from the *management* VDOM. One, and only one, of the VDOMs on a FortiGate device is assigned the role of the management VDOM.

By default, the root VDOM acts as the management VDOM, but you can manually reassign this task to a different VDOM in multi-vdom mode.

It is important to note that the management VDOM designation is solely for traffic originated by FortiGate, such as FortiGuard updates, and has no effect on traffic passing through FortiGate. As such, the management function can be performed by any designated VDOM.

Similar to FortiGate without VDOMs enabled, the administrative VDOM should have outgoing internet access. Otherwise, features such as scheduled FortiGuard updates, fail.

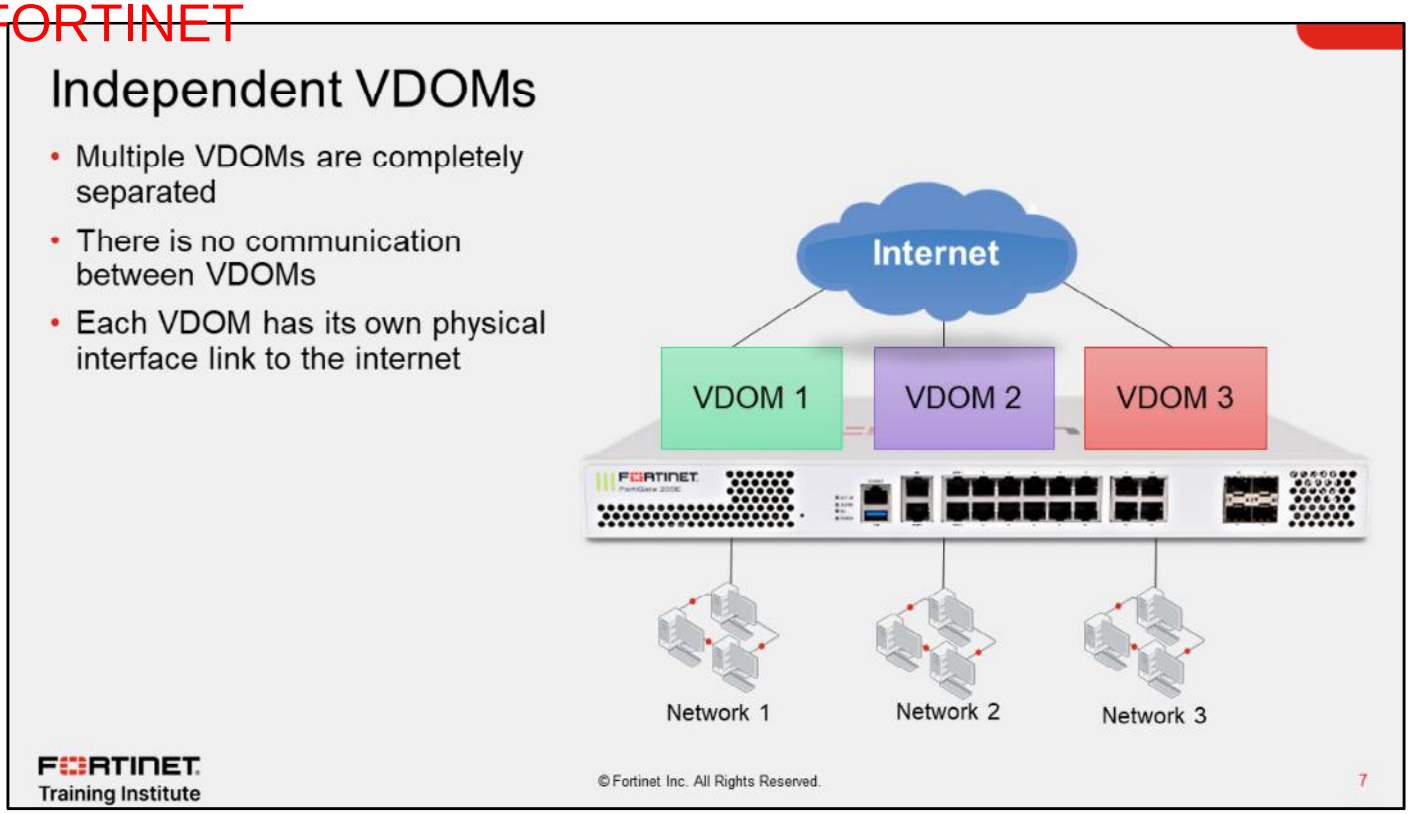

There are a few ways you can arrange your VDOMs. In the topology shown on this slide, each network accesses the internet through its own VDOM.

Notice that there are no inter-VDOM links. So, inter-VDOM traffic is not possible unless it physically leaves FortiGate, toward the internet, and is rerouted back. This topology would be most suitable in a scenario where multiple customers are sharing a single FortiGate, each in their own VDOM, with physically separated ISPs.

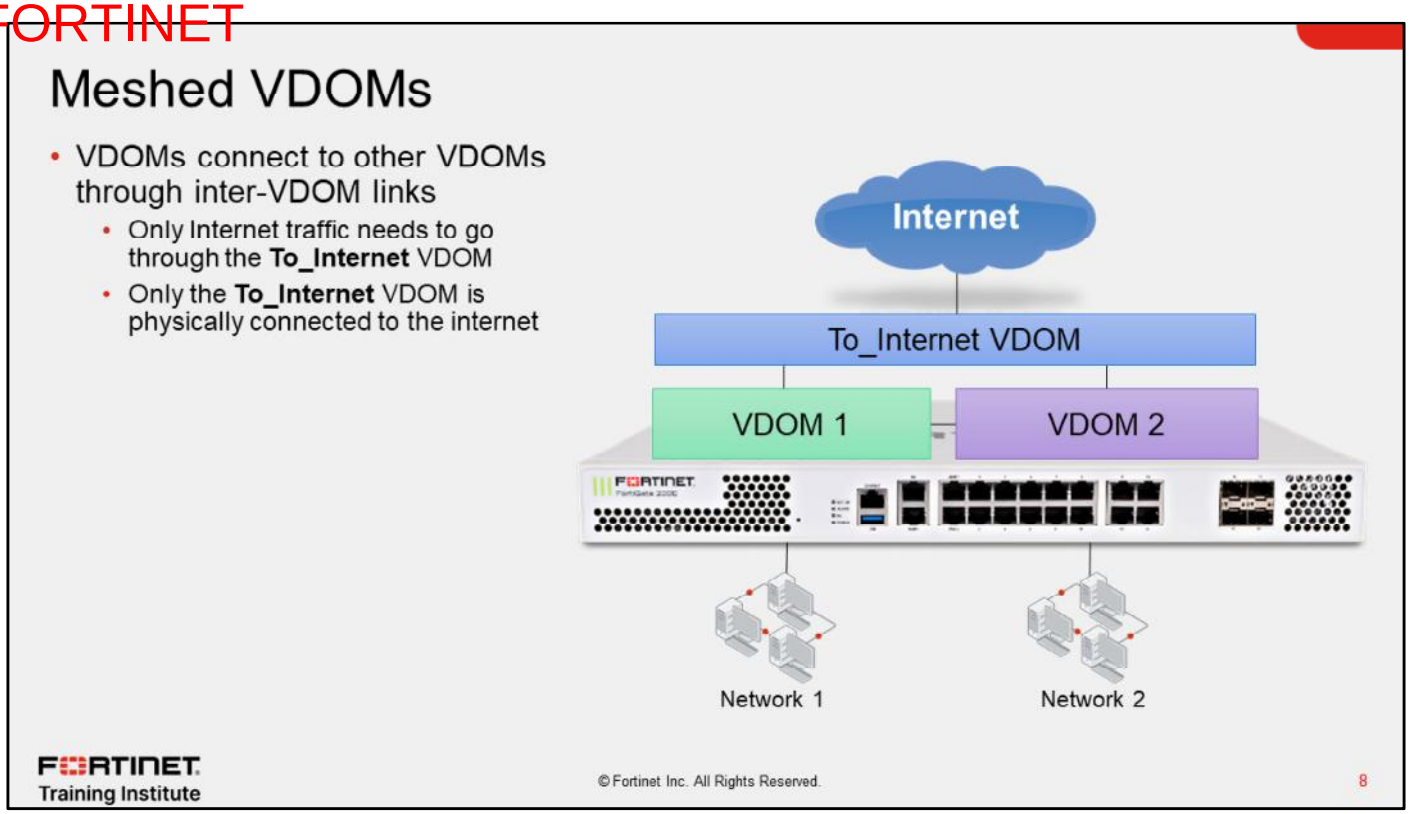

In the example topology shown on this slide, traffic again flows through a single pipe in the **To_Internet** VDOM toward the internet. Traffic between VDOMs doesn't need to leave FortiGate.

However, now inter-VDOM traffic doesn't need to flow through the **To_Internet** VDOM. Inter-VDOM links between VDOMs allow more direct communication.

Similar to the previous example topology, inspection can be done by either the **To_Internet** or originating VDOM, depending on your requirements.

Because of the number of inter-VDOM links, the example shown on this slide is the most complex, requiring the most routes and firewall policies. Troubleshooting meshed VDOMs can also be more time consuming.

However, meshed VDOMs also provide the most flexibility. For large businesses, inter-VDOM communication may be required. Also, inter-VDOM traffic performance may be better because of a shorter processing path, which bypasses intermediate VDOMs.

Routing Through a Single VDOM

- Traffic destined to the internet will *always* be routed through *the designated VDOM* (**To_Internet** in this example)
  - The **To_Internet** VDOM is connected to other VDOMs using inter-VDOM links
  - Only the **To_Internet** VDOM is physically connected to the Internet

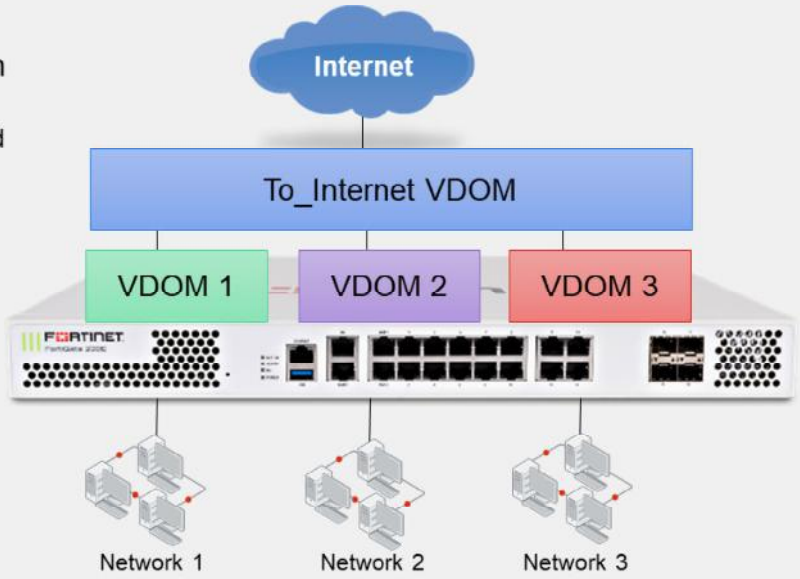

FORTINET Training Institute

9

Like the topology shown on the previous slide, each network in the example topology shown on this slide sends traffic through its VDOM. However, after that, traffic is routed through the **To_Internet** VDOM. So, internet-bound traffic flows through a single pipe in the **To_Internet** VDOM.

This could be suitable in a scenario where multiple customers are sharing a single FortiGate, each in their own VDOM. In this case, the internet-facing VDOM could log and monitor traffic, or provide standard services like antivirus scanning, or both.

The topology shown on this slide has inter-VDOM links. VDOMs are linked only with the **To_Internet** VDOM, but not with each other. If **VDOM1** needs to communicate with **VDOM3**, this traffic would need to be routed through the **To_Internet** VDOM through IP routing decisions and is subject to all firewall policies.

Inspection could be done by either the internet-facing or originating VDOM, depending on your requirements. Alternatively, you could split inspection so that some scans occur in the internet-facing VDOM—ensuring a common security baseline—while other more intensive scans occur in the originating VDOM.

Knowledge Check

1. Which traffic is always generated from the management VDOM?
   - A. Link Health Monitor
   - ✓ B. FortiGuard

2. Which statement about the management VDOM is true?
   - A. It is **root** by default and cannot be changed in multi-vdom mode.
   - ✓ B. It is **root** by default, but can be changed to any VDOM in multi-vdom mode.

FORTINET Training Institute

10

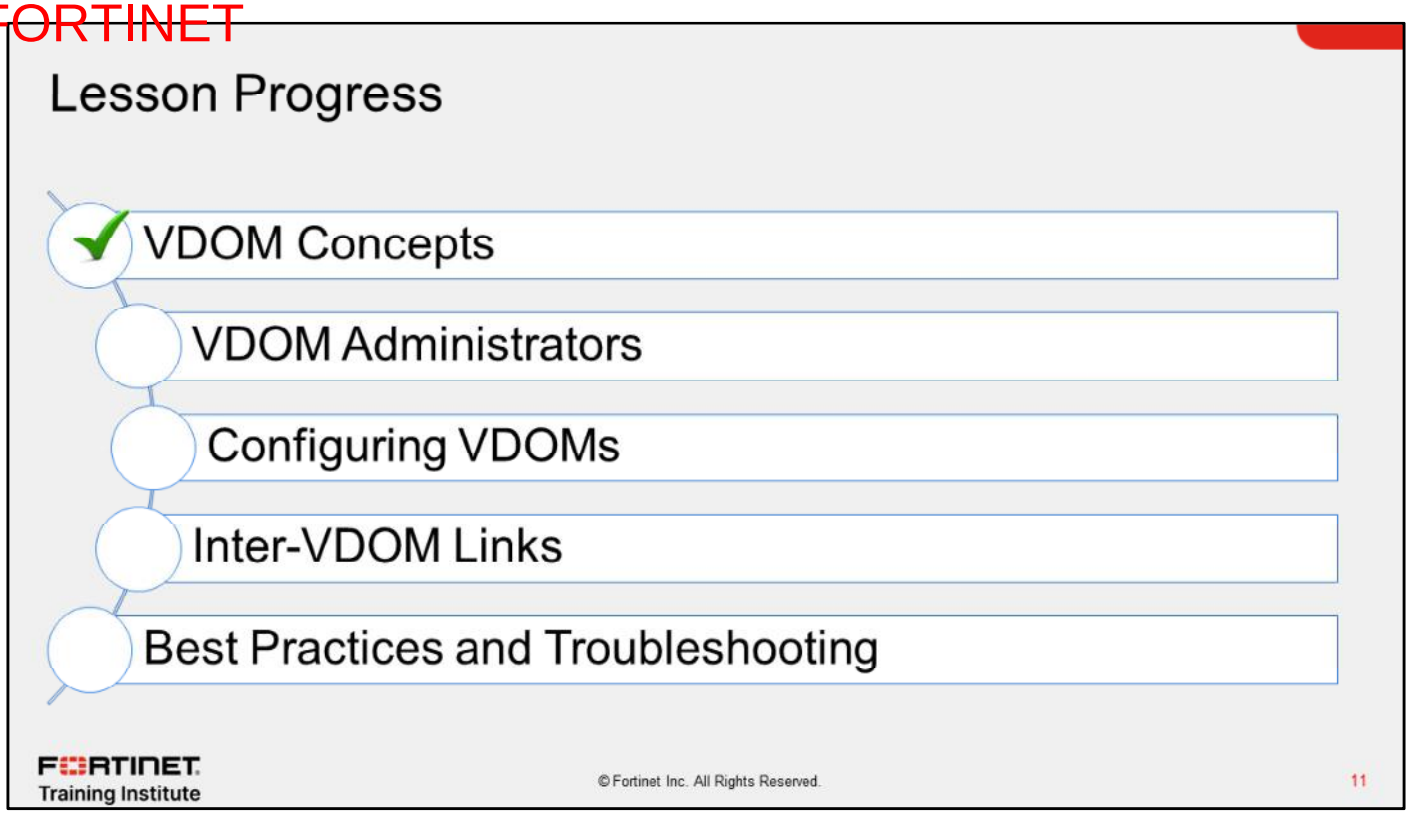

Good job! You now understand some basic concepts about VDOMs.

Now, you'll learn about VDOM administrators.

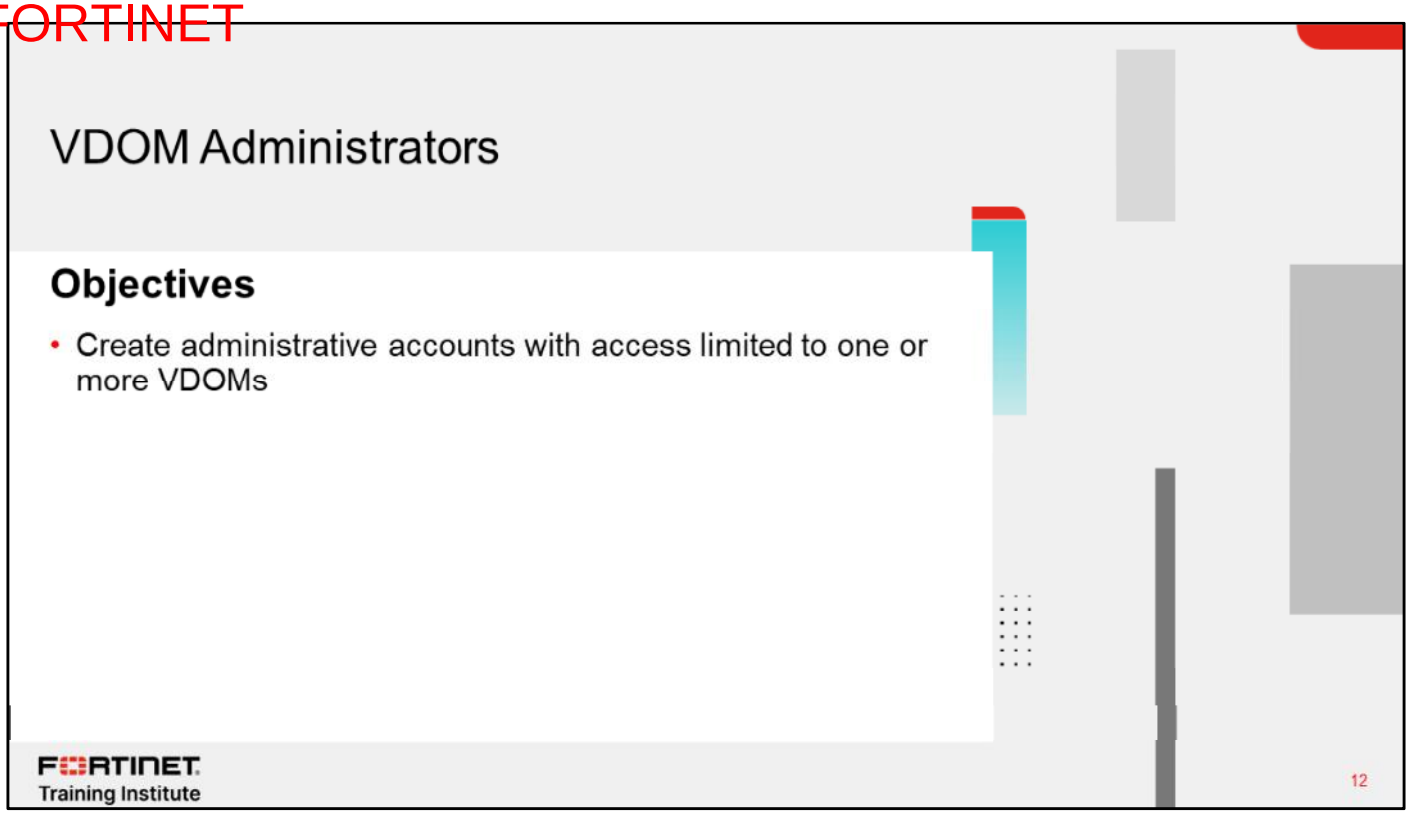

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in creating VDOM administrative accounts, you will be able to understand the differences between the various levels and types of VDOM administrators.

VDOM Administration

- Only the account named **admin** or accounts with the **super_admin** profile can configure and back up all VDOMs

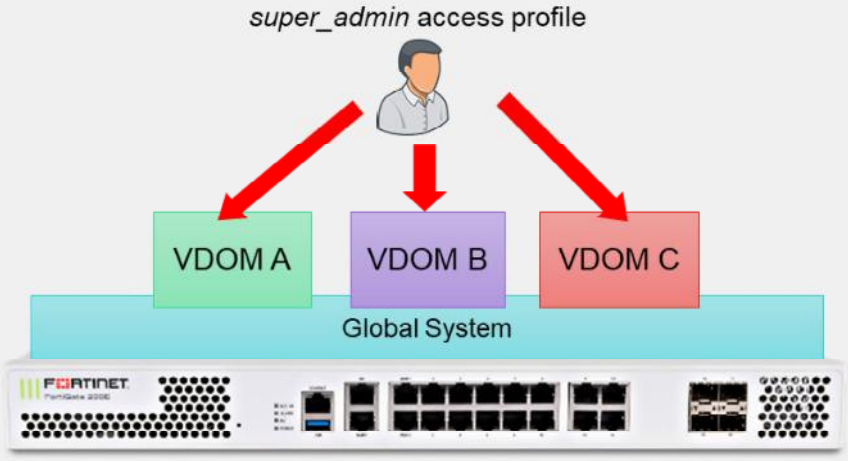

FORTINET Training Institute

13

If you want to grant access to all VDOMs and global settings, select **super_admin** as the access profile when configuring the administrator account. Similar to the account named **admin**, this account can configure all VDOMs.

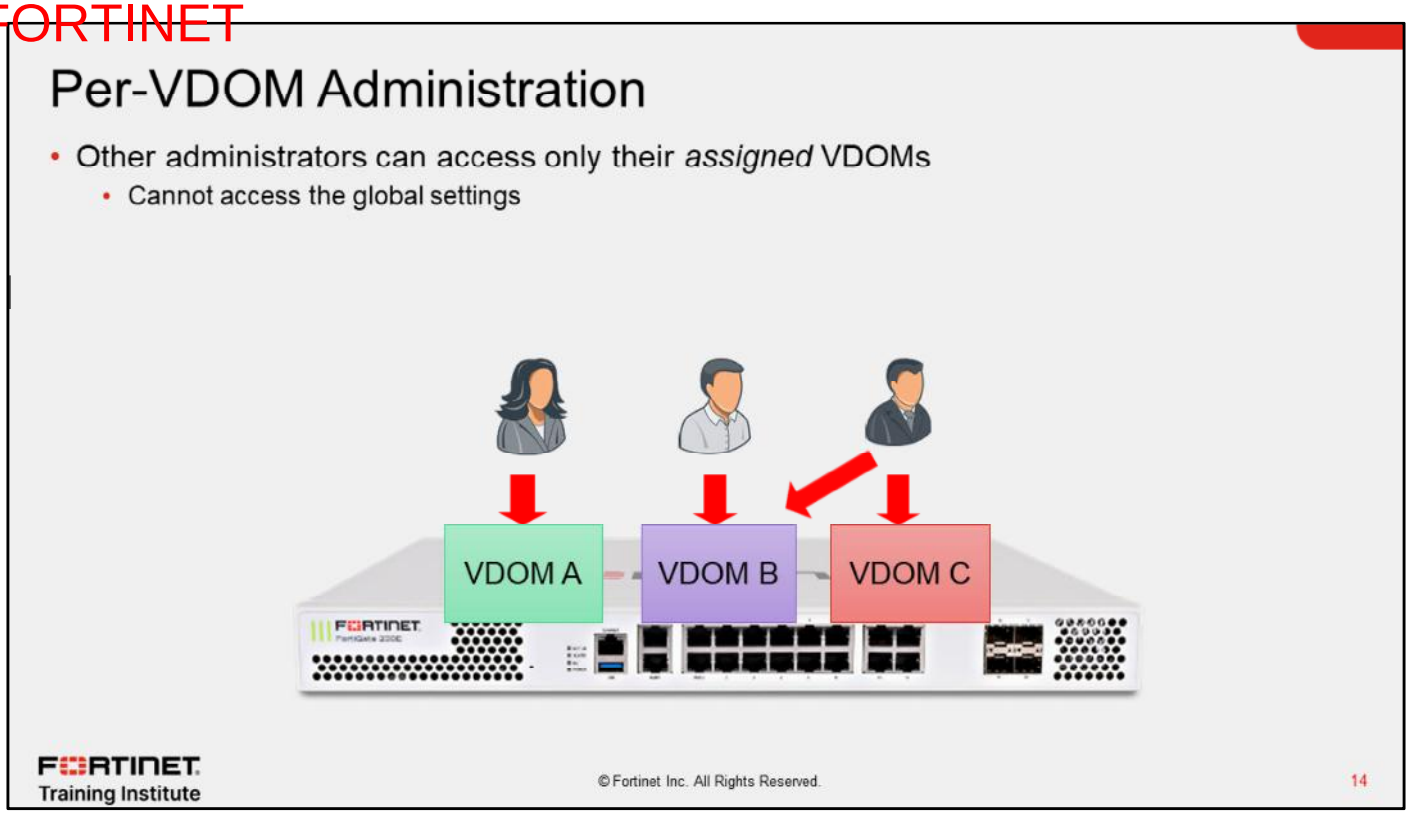

In most cases, you start by creating one administrator account per VDOM. That administrator is chiefly responsible for that domain, including the configuration backups of that VDOM. In larger organizations, you may need to make multiple VDOM administrators. You can assign multiple administrators to each VDOM. You can subdivide permissions using access profiles, in order to follow best practices for segregation of duties.

The converse is also possible. If required, you can assign an administrator to multiple VDOMs.

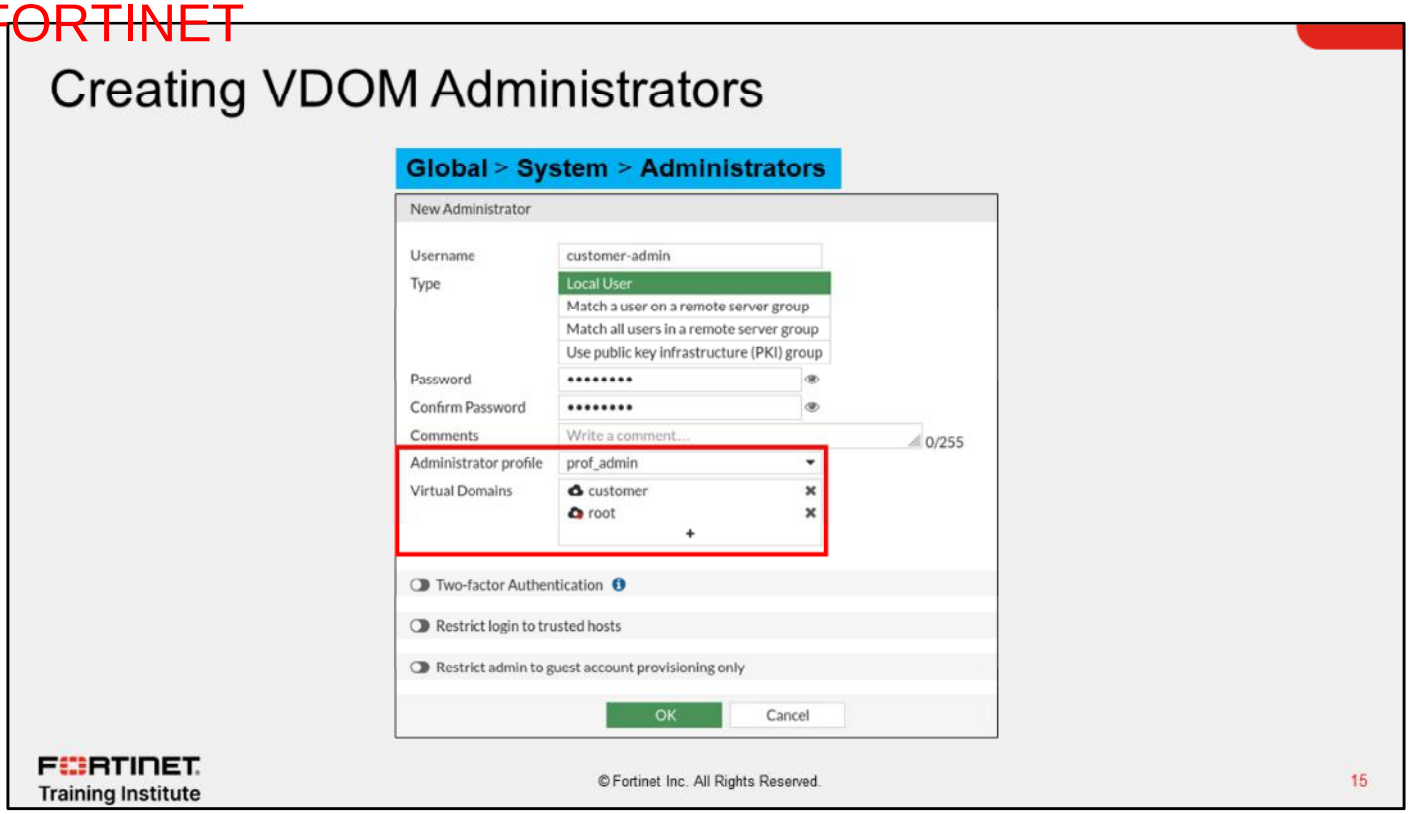

To create new administrator accounts and assign them to a VDOM, click **Global** > **System** > **Administrators**.

Knowledge Check

1. Which type of administrator can make changes to all VDOMs?
   - A. A custom VDOM administrator
   - ✓ B. An administrator with the **super_admin** profile

2. Which statement about VDOM administrators is true?
   - A. There can be only one administrator per VDOM.
   - ✓ B. Each VDOM can have multiple administrators.

FORTINET Training Institute

16

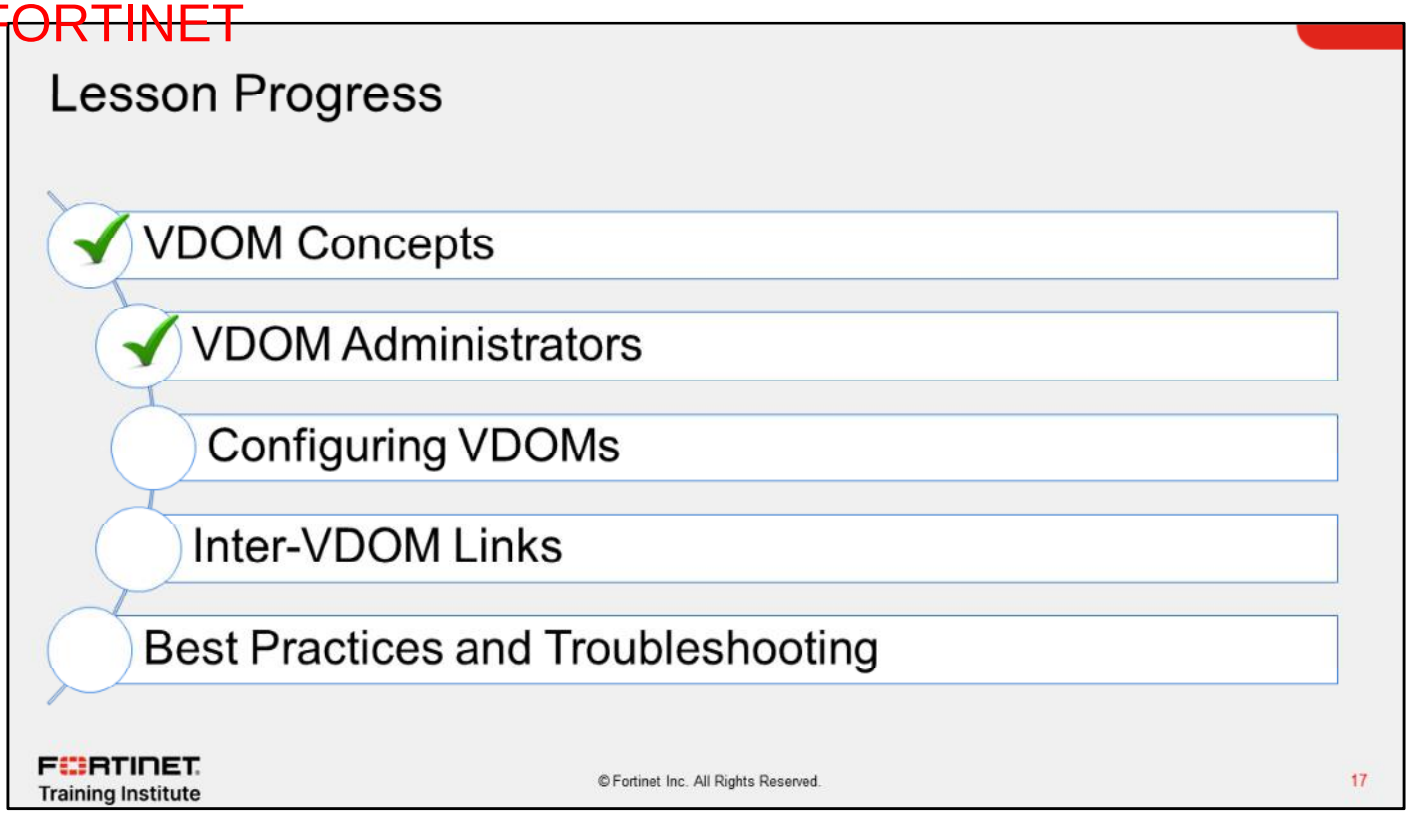

Good job! You now understand VDOM administrators.

Now, you'll learn how to configure VDOMs.

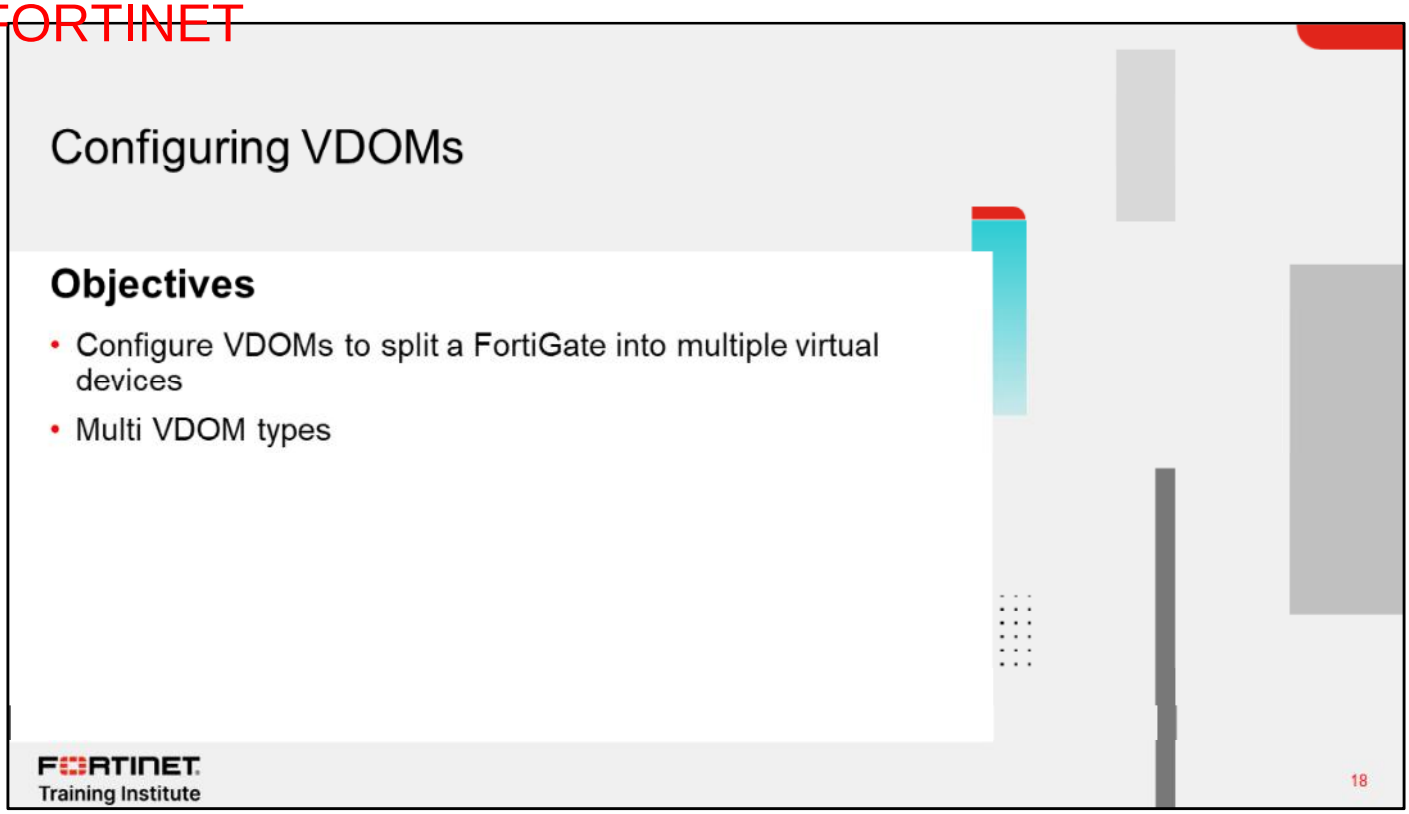

After completing this section, you will be able to achieve the objective shown on this slide.

By demonstrating competence in configuring VDOMs, you will be able to effectively implement VDOMs on your FortiGate.

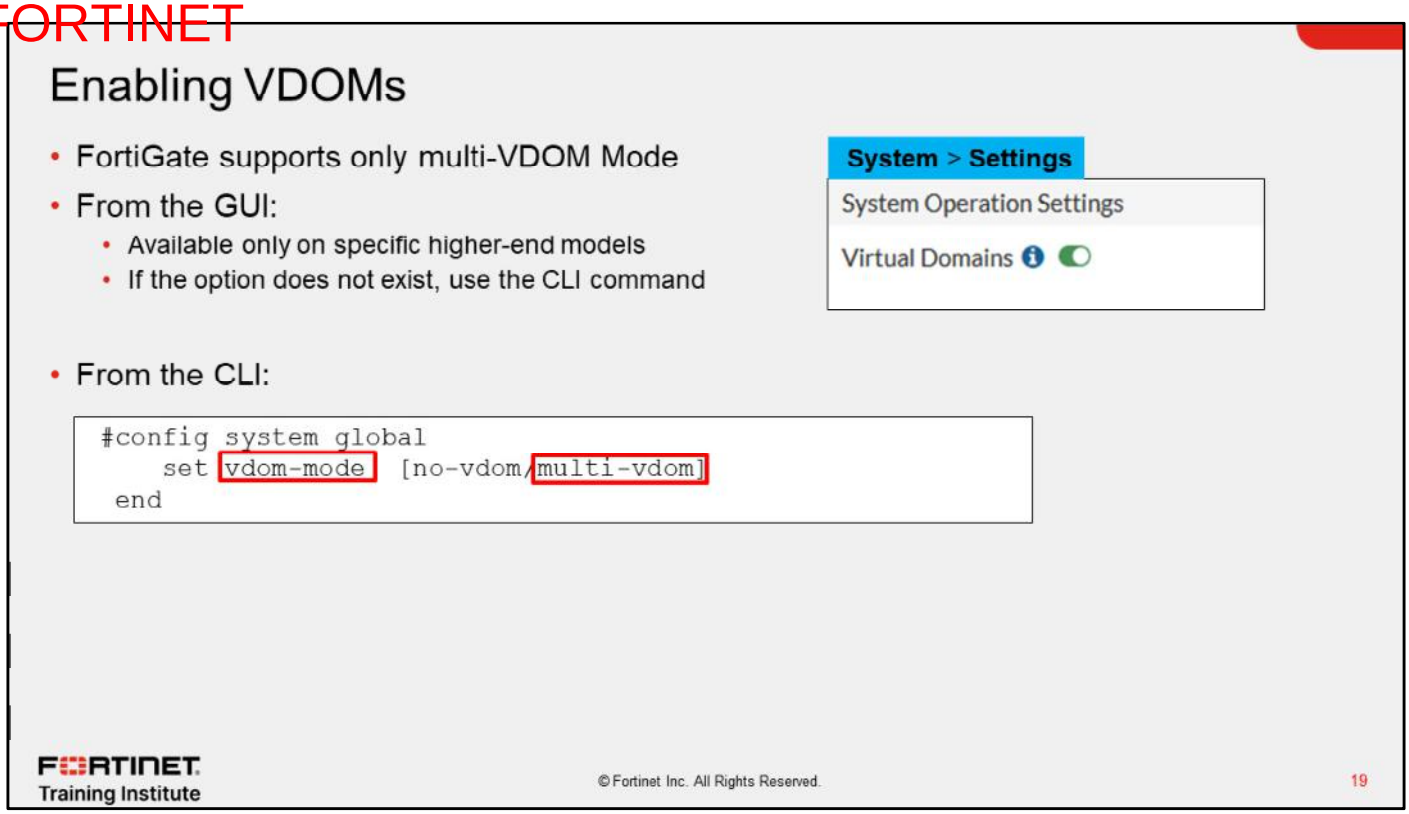

On the GUI, you can enable VDOMs under **System** > **Settings**. The GUI option is available only on higher-end FortiGate Models. Most of the FortiGate models, you can enable VDOMs on the CLI only.

Enabling VDOMs does not cause your FortiGate device to reboot, but it does log out all active administrator sessions. Traffic continues to pass through FortiGate.

Enabling VDOMs restructures both the GUI and CLI, which you will see when you log in again.

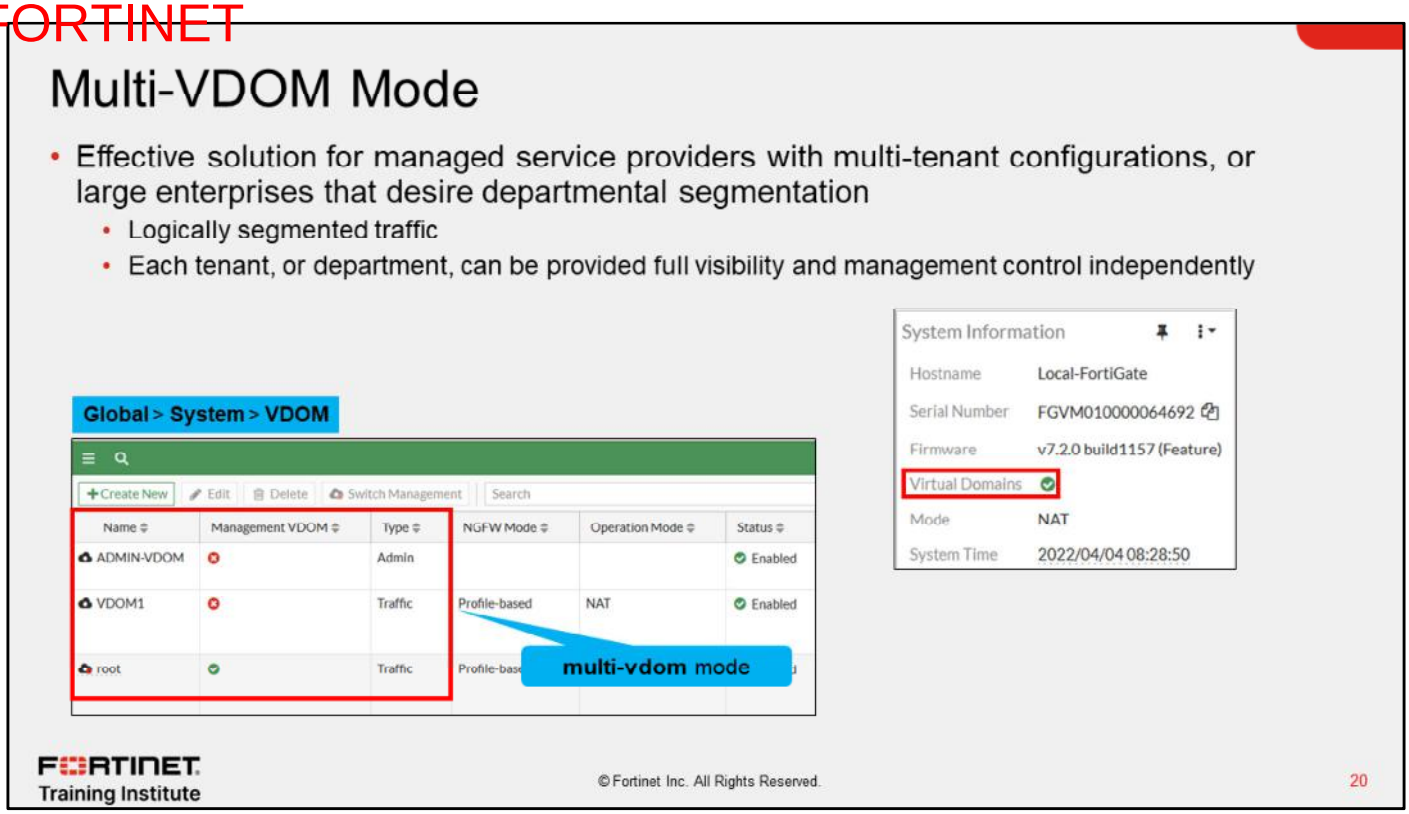

In *multi-vdom* mode, you can create multiple VDOMs that function as multiple independent units.
By default, the root is the management VDOM and can be used to do both management tasks and allow other traffic. You can select any VDOM to act as the management VDOM.

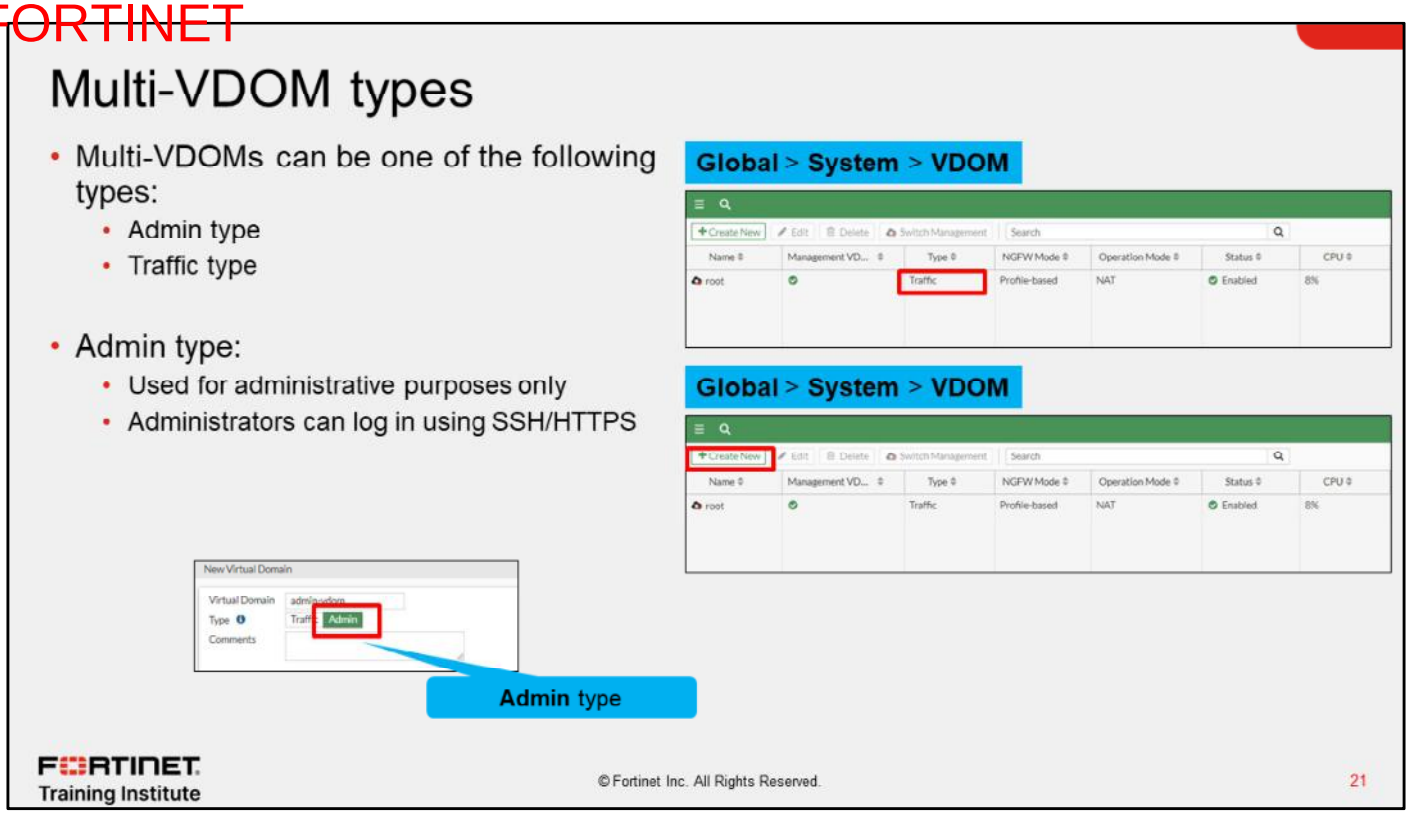

When you enable multi-vdom mode, the root VDOM exists. It is the default management VDOM and is a traffic VDOM. You can create another VDOM (traffic or admin). FortiGate supports only one admin VDOM.

Multi -VDOM types (Contd)

- Traffic type:
  - Can pass traffic like regular VDOMs

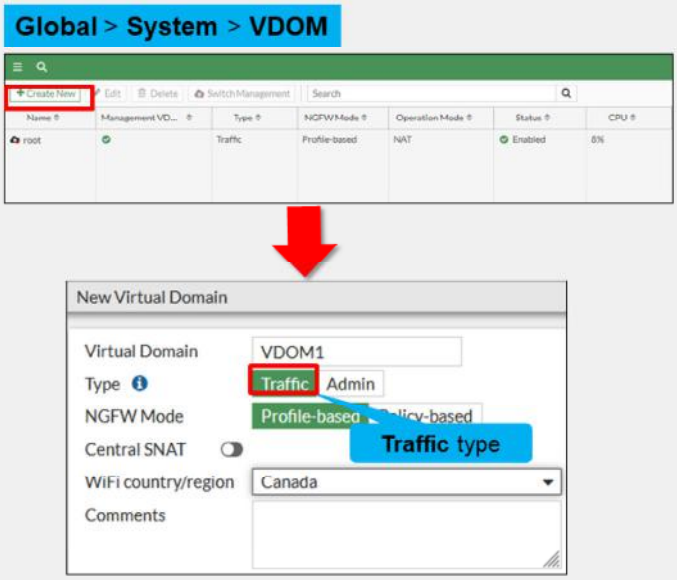

- From CLI:

```
config vdom
edit <vdom>
   config system settings
      set vdom-type [traffic/admin]
   end
```

FORTINET Training Institute

22

When the VDOM type is set to Traffic, the VDOM can pass traffic like a regular VDOM. If an admin VDOM exists, all newly created VDOMs are configured as traffic VDOMs.

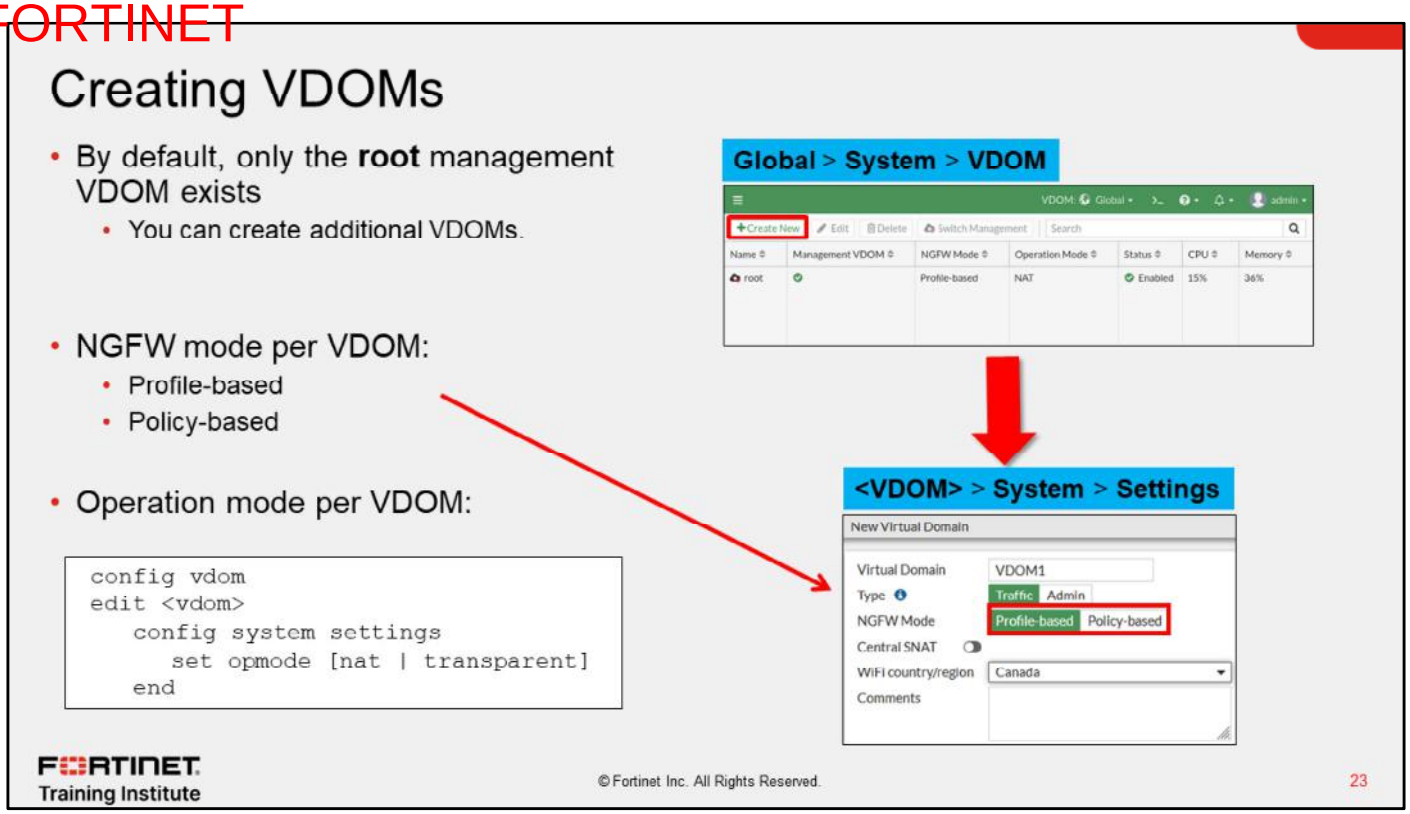

After enabling VDOMs in multi-vdom mode, by default, only one VDOM exists: the root VDOM. It's the default management VDOM.

You need to add a VDOM for each of your security domains. If you're an MSSP, for example, you might add one VDOM for each client company. If you are an enterprise business, you might add one VDOM for each division of your company.

The default inspection-mode is flow, so you can change **NGFW Mode** from **Profile-based** (default) to **Policy-based** directly in **System** > **Settings** for the VDOM.

The **profile-based** NGFW is the traditional mode and you must create antivirus, web filter, and IPS profiles, which are then applied to the policy. **Policy-based** mode is actually a new policy mode. You can add applications and web filtering categories directly to a policy without having to first create and configure application control or web filtering profiles. NGFW mode is a per-VDOM setting. If you set NGFW mode to **Profile-based**, you can configure policies in that VDOM for either flow or proxy inspection. However, if NGFW mode is **Policy-based**, then the inspection mode for all policies in that VDOM is always flow and there is no option available in the policy to change it.

Switching between NGFW modes results in *the loss of all current policies* configured in the VDOM. If you don't want this to happen, or you just want to experiment with a particular NGFW mode, consider creating a new VDOM for testing purposes. You could also back up your configuration before switching modes.

Operation mode is a per-VDOM setting. You can combine transparent mode VDOMs with NAT mode VDOMs on the same physical FortiGate.

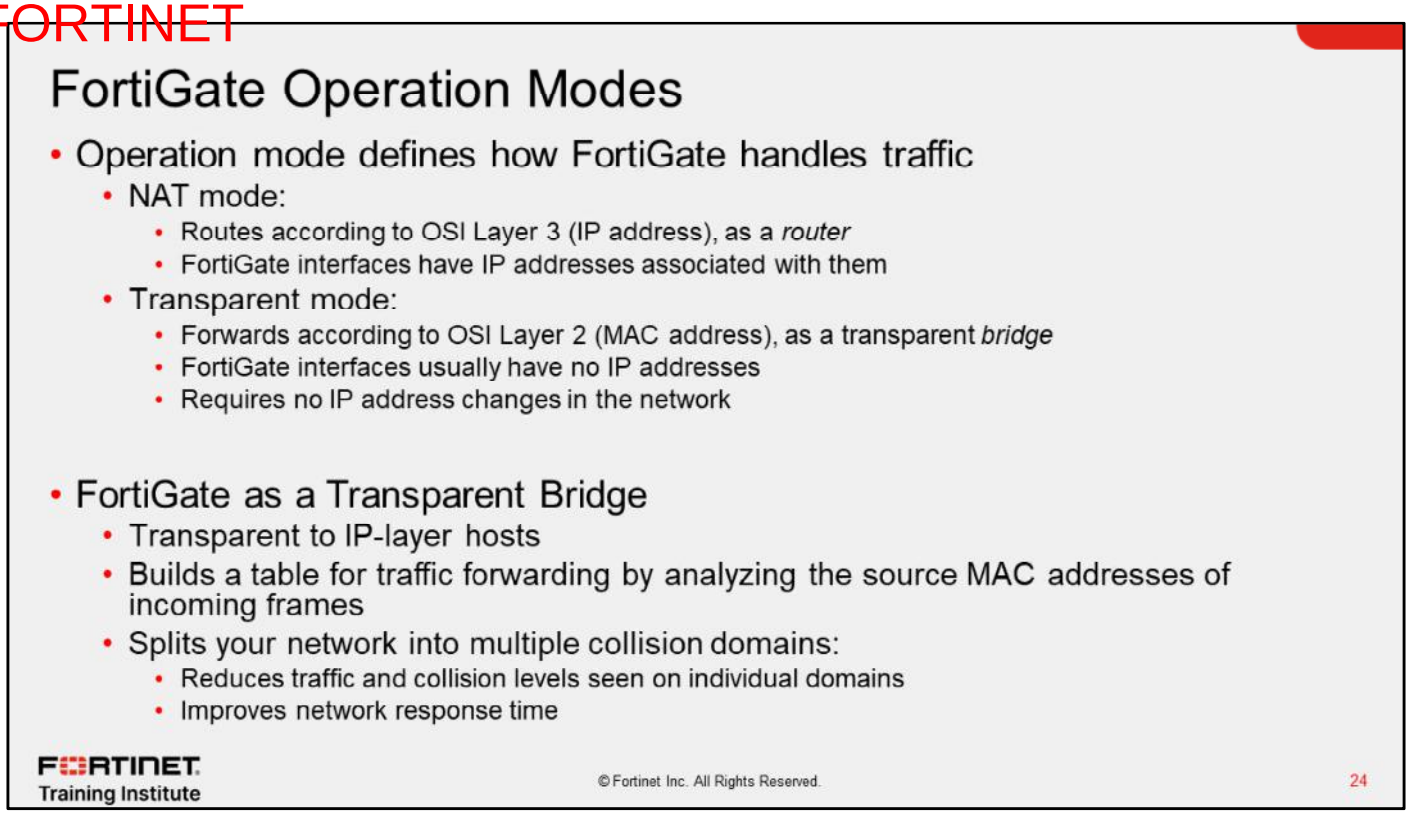

Traditional IPv4 firewalls and NAT mode FortiGate devices handle traffic the same way that routers do. Each interface must be in a different subnet and each subnet forms a different broadcast domain. FortiGate routes IP packets based on the IP header information, overwriting the source MAC address. So, if a client sends a packet to a server connected to a different FortiGate interface, the packet arrives at the server with a FortiGate MAC address, instead of the client MAC address.

In transparent operation mode, FortiGate forwards frames without changing the MAC addresses. When the client receives a packet from a server connected to a different FortiGate interface, the frame contains the real MAC address of the server—FortiGate doesn't rewrite the MAC header. FortiGate acts as a Layer 2 bridge or switch. So, the interfaces do not have IP addresses and, by default, all belong to the same broadcast domain.

This means that you can install a transparent mode FortiGate in a customer network without having to change the customer's IP address plan. Some customers, especially large organizations, don't want to reconfigure thousands of devices to define a new internal network that is separate from their external network.

A transparent mode FortiGate device acts as a transparent bridge. What does that mean? It means that FortiGate has a MAC address table that contains, among other things, the interface that must be used to reach each MAC address. FortiGate populates this table with information taken from the source MAC address of each frame.
FortiGate, as a transparent switch, splits the network into multiple collision domains, reducing the traffic in the network and improving the response time.

Forward Domains

- By default, *all* interfaces on a VDOM belong to the same broadcast domain; even interfaces with different VLAN IDs
  - Broadcast domains that contain multiple interfaces can be very large and add unnecessary broadcast traffic to some LAN segments
- Use this command to subdivide a VDOM into multiple broadcast domains:

```
config system interface
   edit <interface_name>
       set forward-domain <domain_ID>
end
```

- Interfaces with the same domain ID belong to the same broadcast domain

FORTINET Training Institute

25

By default, in transparent operation mode, each VDOM forms a separate forward domain; however, interfaces do not. How does this affect the network?

Until you change the initial VDOM configuration, all interfaces, regardless of their VLAN ID, are part of the same broadcast domain. FortiGate broadcasts from every interface in the VDOM in order to find any unknown destination MAC address. On large networks, this could generate massive broadcast traffic and overwhelming replies—a broadcast storm.

Confirmation Prompt When Creating VDOMs

- VDOM confirmation prompt added
  - So that users do not create new VDOMs accidently in CLI

```
config system global
 set edit-vdom-prompt [enable | disable]
end
```

  - Disabled by default
  - When enabled, if administrator creates a new VDOM, FortiGate displays prompt:

```
# config vdom
    edit student
    The input VDOM name doesn't exist.
    Do you want to create a new VDOM?
    Please press 'y' to continue, or press 'n' to cancel. (y/n)y

    current vf=student:3
```

Prompt to confirm before the new VDOM is created

FORTINET Training Institute

26

A VDOM confirmation prompt has been added so users do not create new VDOMs accidentally on the CLI. This setting is disabled by default. Once enabled, when an administrator creates a new VDOM, FortiGate displays a prompt to confirm before the VDOM is created.

Assigning Interfaces to a VDOM

- You can assign an interface to each VDOM you create
- From CLI:

```
config global
config system interface
  edit <interface_name>
     set vdom <vdom-name>
   end
```

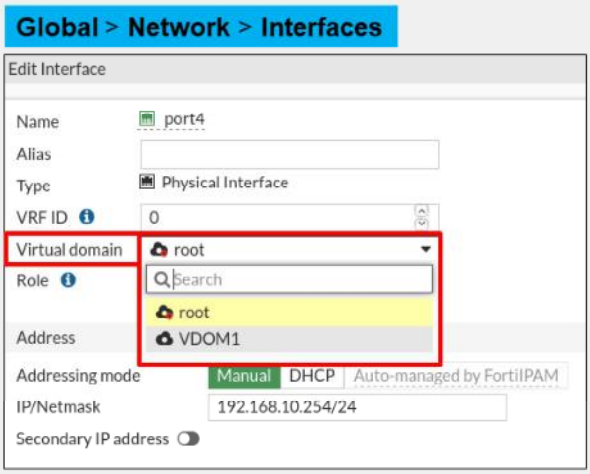

FORTINET Training Institute

27

After adding a VDOM, you can specify which interface belongs to it. Each interface (physical or VLAN) can belong to only one VDOM.

You can move an interface from one VDOM to another, provided it is not associated with any references, such as firewall policies.

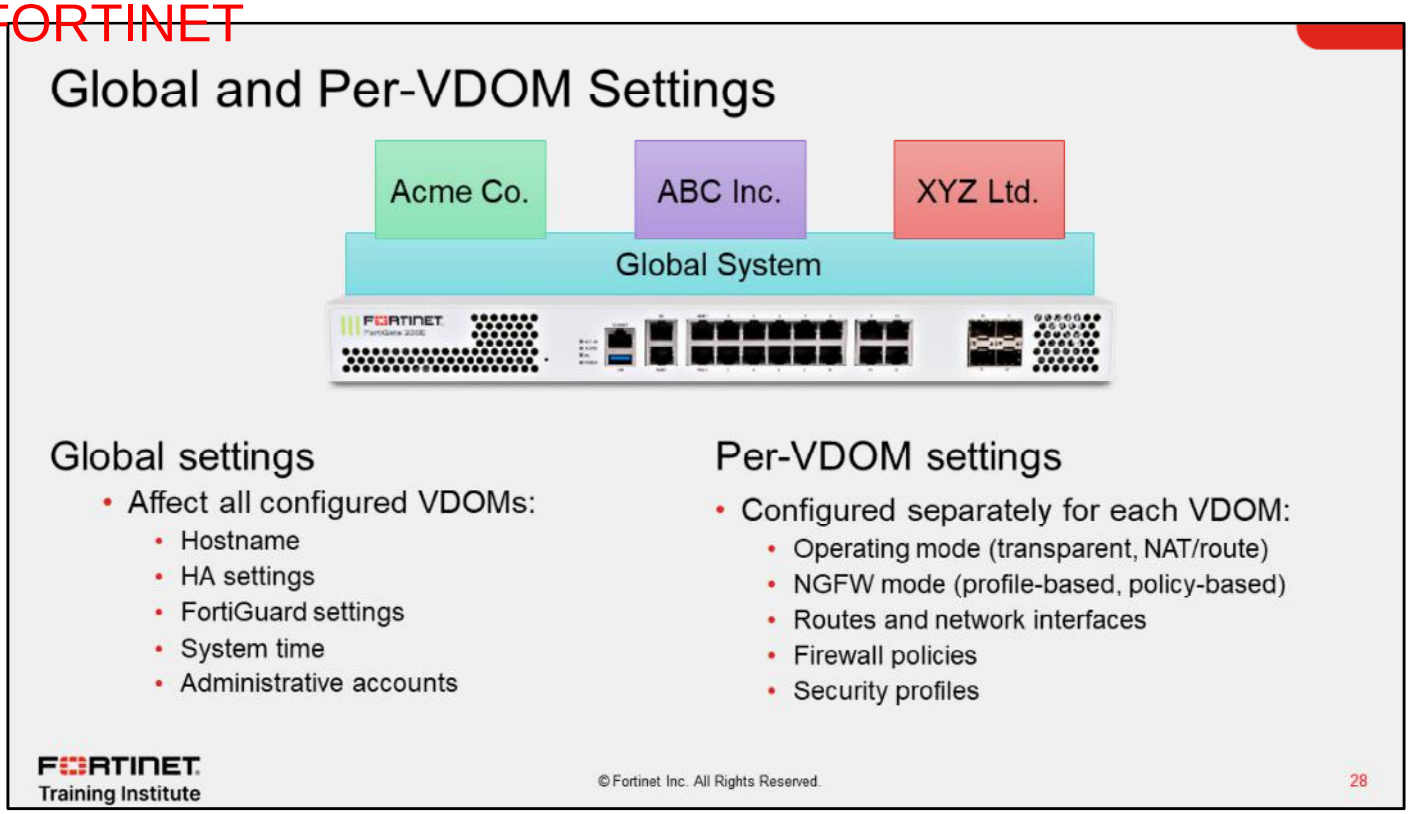

Global resource limits are an example of global settings. The firmware on your FortiGate device and some settings, such as system time, apply to the entire device—they are not specific to each VDOM.

However, you can configure most settings differently for each VDOM. Some examples are firewall policies, firewall objects, static routes, and protection profiles.

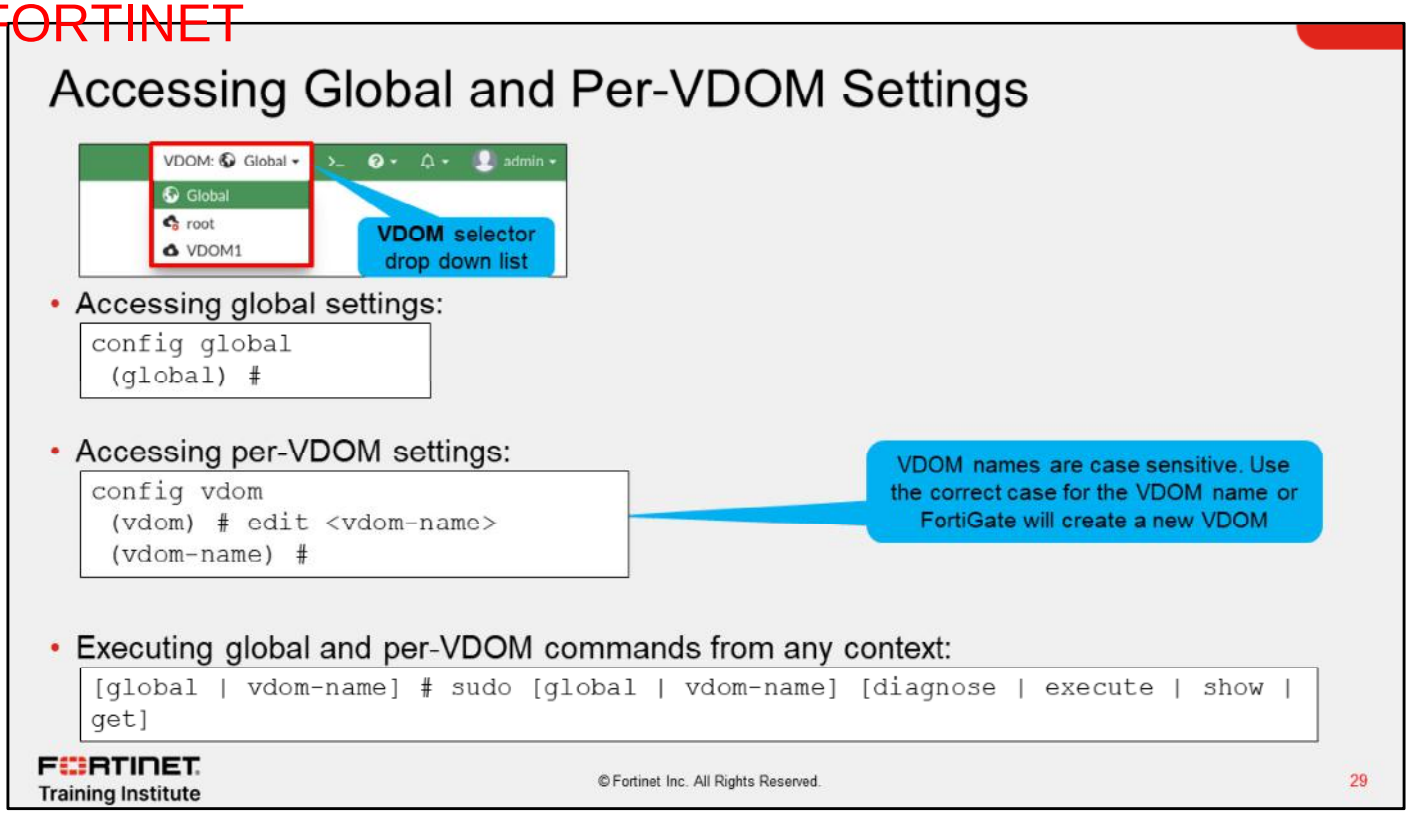

When you log with a regular administrator account, you automatically enter the VDOM associated with that account.

When you log in with the account named admin, you have access to all VDOMs. To access a specific VDOM, select it in the drop-down list at the top of the page.

The VDOM submenu should be familiar; it is essentially the same navigation menu from before you enabled VDOMs. However, the global settings are moved to the Global menu.

To access the global configuration settings on the CLI, you must enter config global to enter into the global context. After that, you can run global commands and change global configuration settings.

To access per-VDOM configuration settings on the CLI, you must enter config vdom, then enter edit followed by the VDOM name. From the VDOM context, you can run VDOM-specific commands and change per-VDOM configuration settings. It is important to note that VDOM names are case sensitive. If you enter the name using the incorrect case, FortiGate creates a new VDOM.

Regardless of which context you are in (global or VDOM), you can use the sudo keyword to run diagnostics commands in a context different from your current one. This allows you to run global and per-VDOM commands, for example, without switching back and forth between the global and per-VDOM contexts.

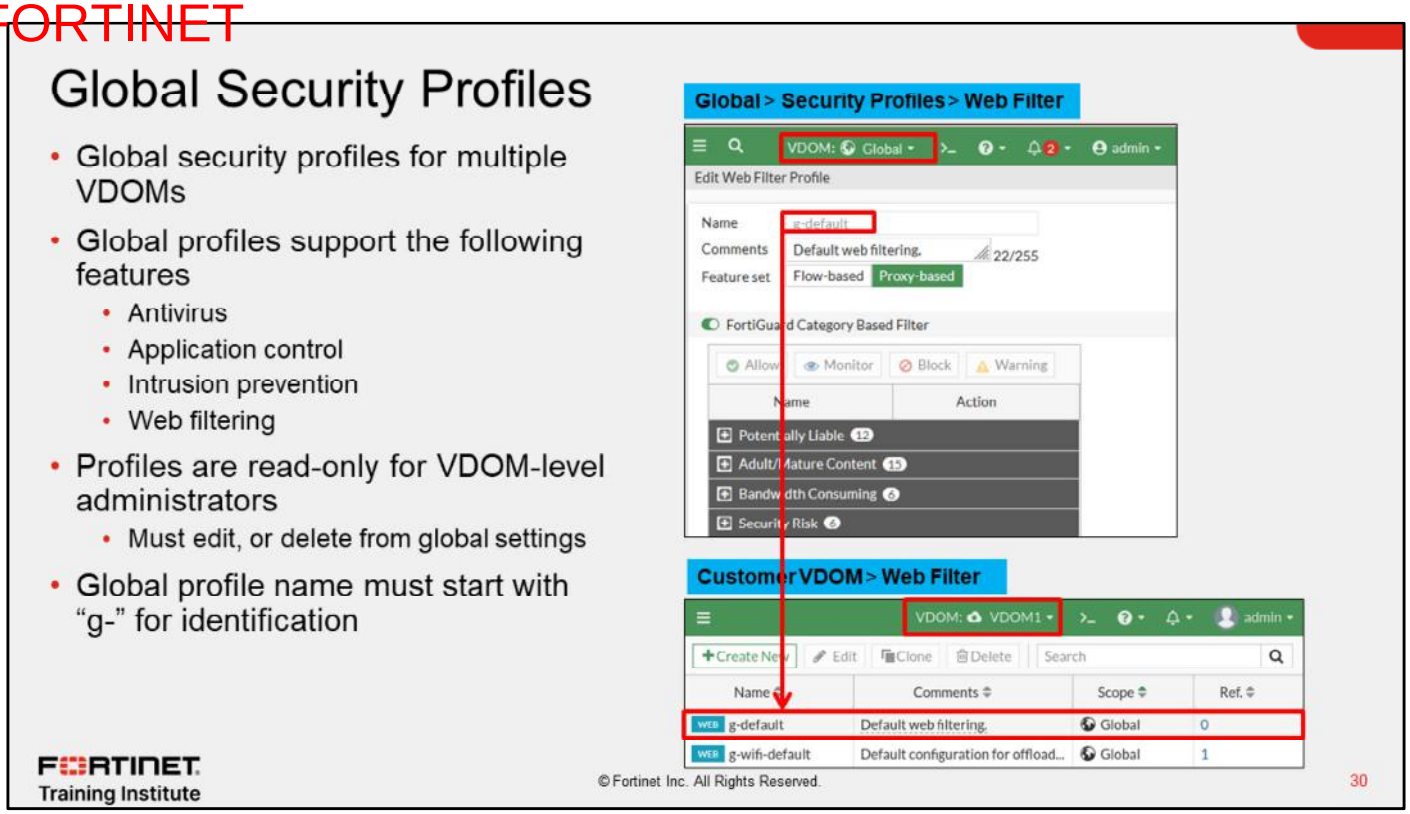

You can configure security profiles globally for use by multiple VDOMs, to avoid creating identical profiles for each VDOM separately. Global profiles are available for the following security features:

- Antivirus
- Application control
- Intrusion prevention
- Web filtering

Some security profile features, such as URL filters, are not available for use in a global profile. The name for any global profile must start with "g-" for identification. Global profiles are available as read-only for VDOM-level administrators and can be edited or deleted only in the global settings. Each security feature has at least one default global profile.

Knowledge Check

1. Which configuration settings are global settings?
   - A. Firewall policies
   - ✓B. FortiGuard settings

2. Which configuration settings are per-VDOM settings?
   - A. Host name
   - ✓B. NGFW mode

FORTINET Training Institute

31

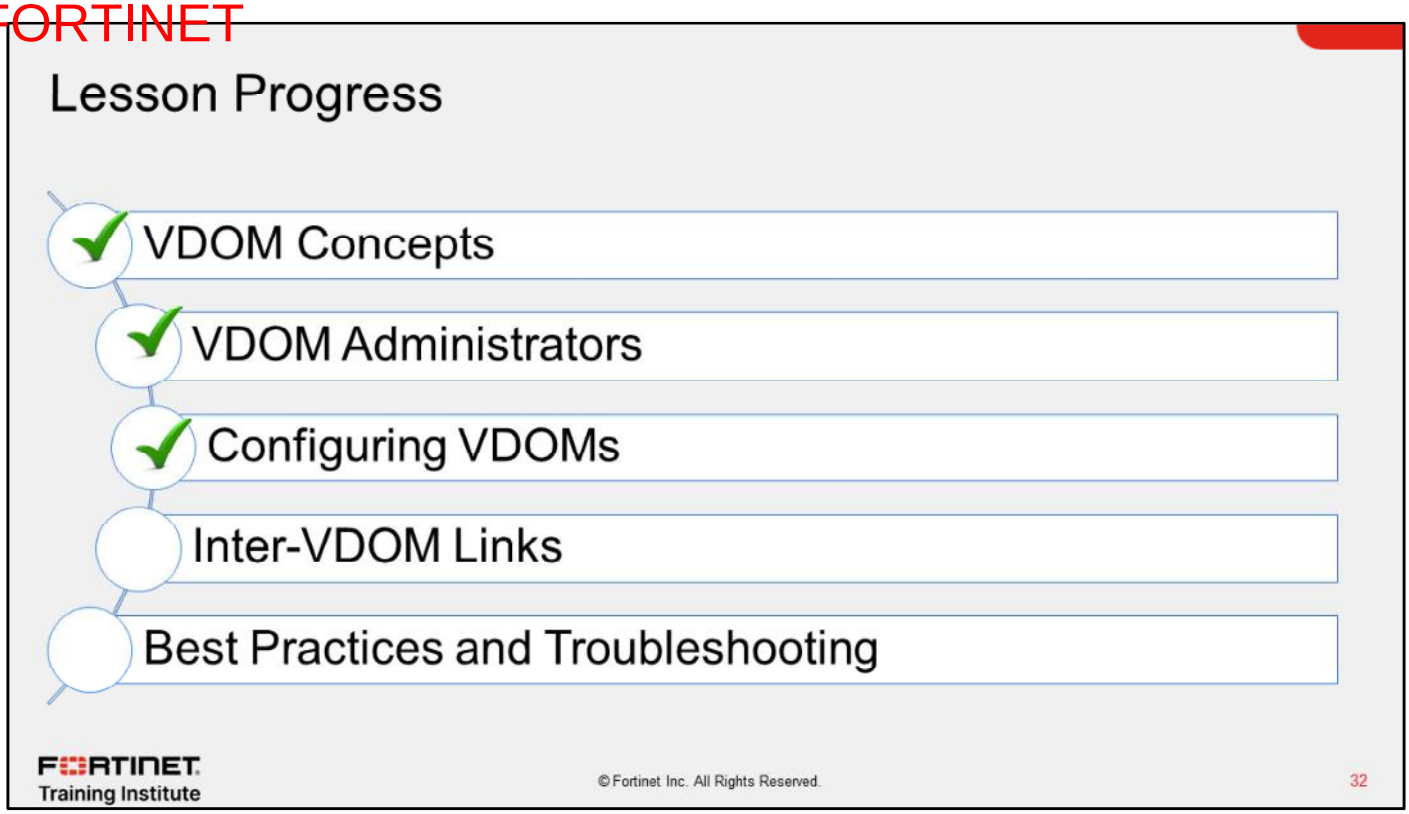

Good job! You now understand how to configure VDOMs.

Now, you'll learn about inter-VDOM links.

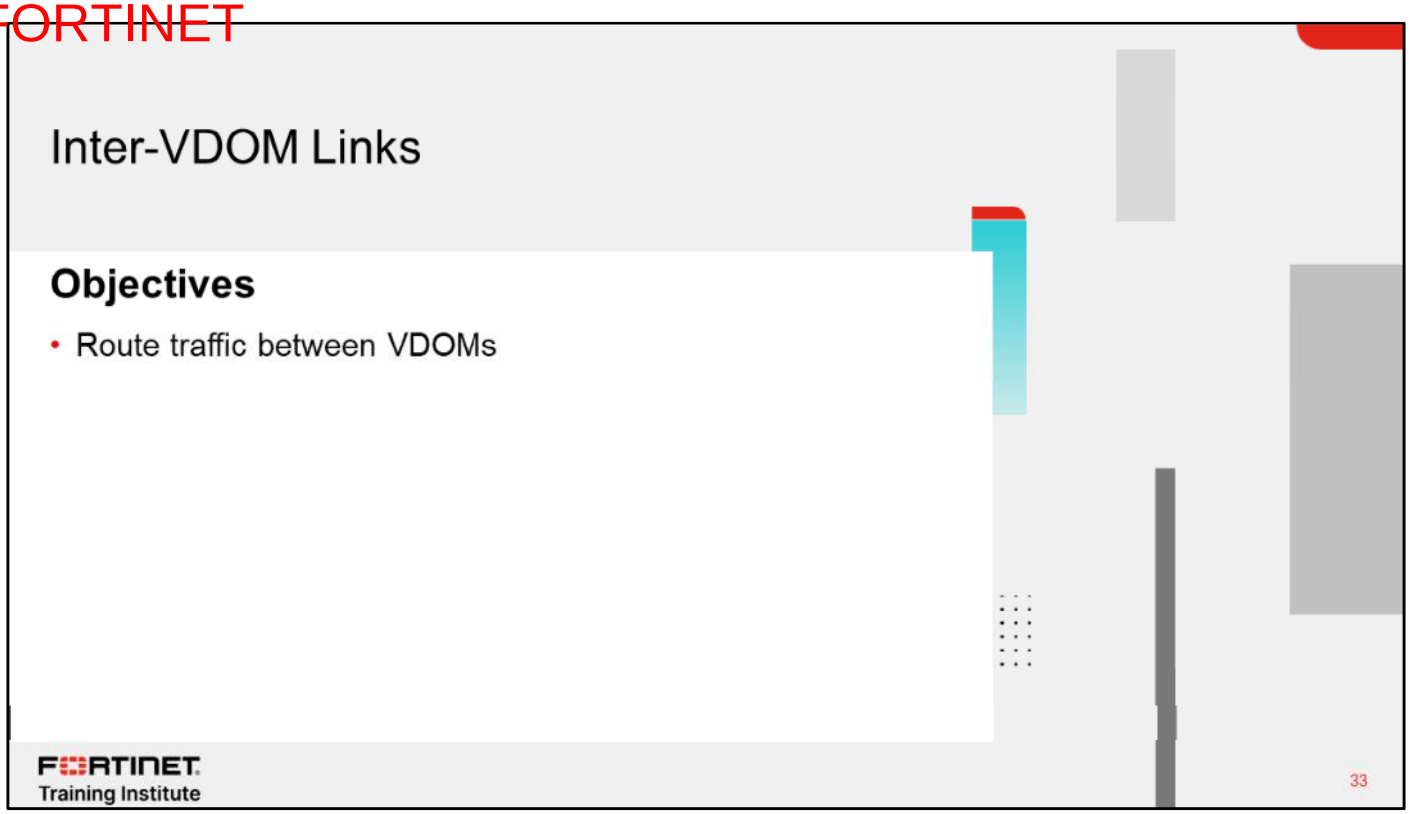

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in inter-VDOM links, you will be able to effectively and efficiently route traffic between VDOMs on FortiGate.

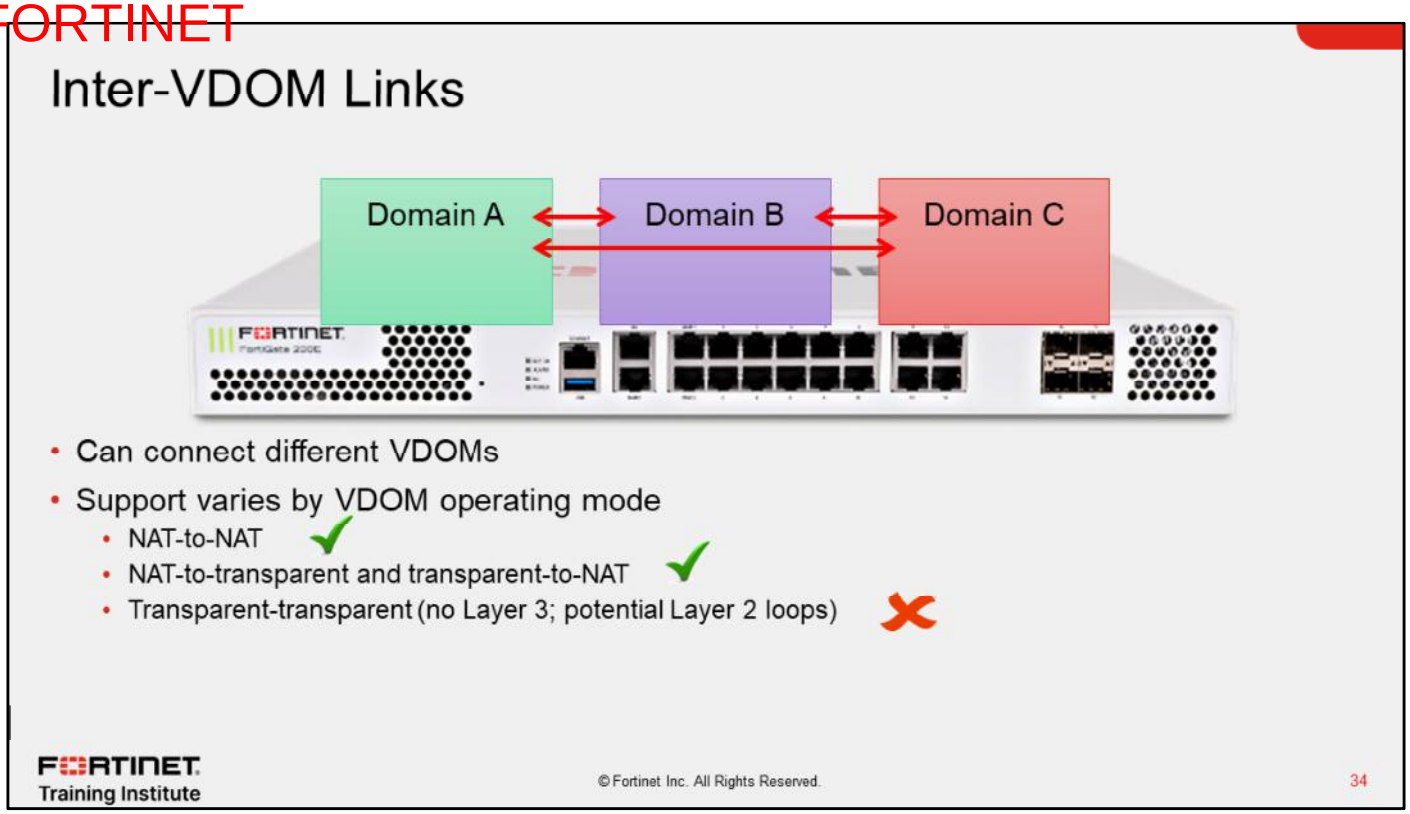

To review, each VDOM behaves like it is on a separate FortiGate device. With separate FortiGate devices, you would normally connect a network cable and configure routing and policies between them. But VDOMs are on the same FortiGate. So, how should you route traffic between them?
The solution is inter-VDOM links. Inter-VDOM links are a type of virtual interface that route traffic between VDOMs. This removes the need to loop a physical cable between two VDOMs.

In the case of a NAT-to-NAT inter-VDOM link, both sides of the link must be on the same IP subnet, because you are creating a point-to-point network connection.

Note that like using inter-VLAN routing, Layer 3 must be involved—you cannot create an inter-VDOM link between Layer 2 transparent mode VDOMs. At least one of the VDOMs must be operating in NAT mode. This, among other benefits, prevents potential Layer 2 loops.

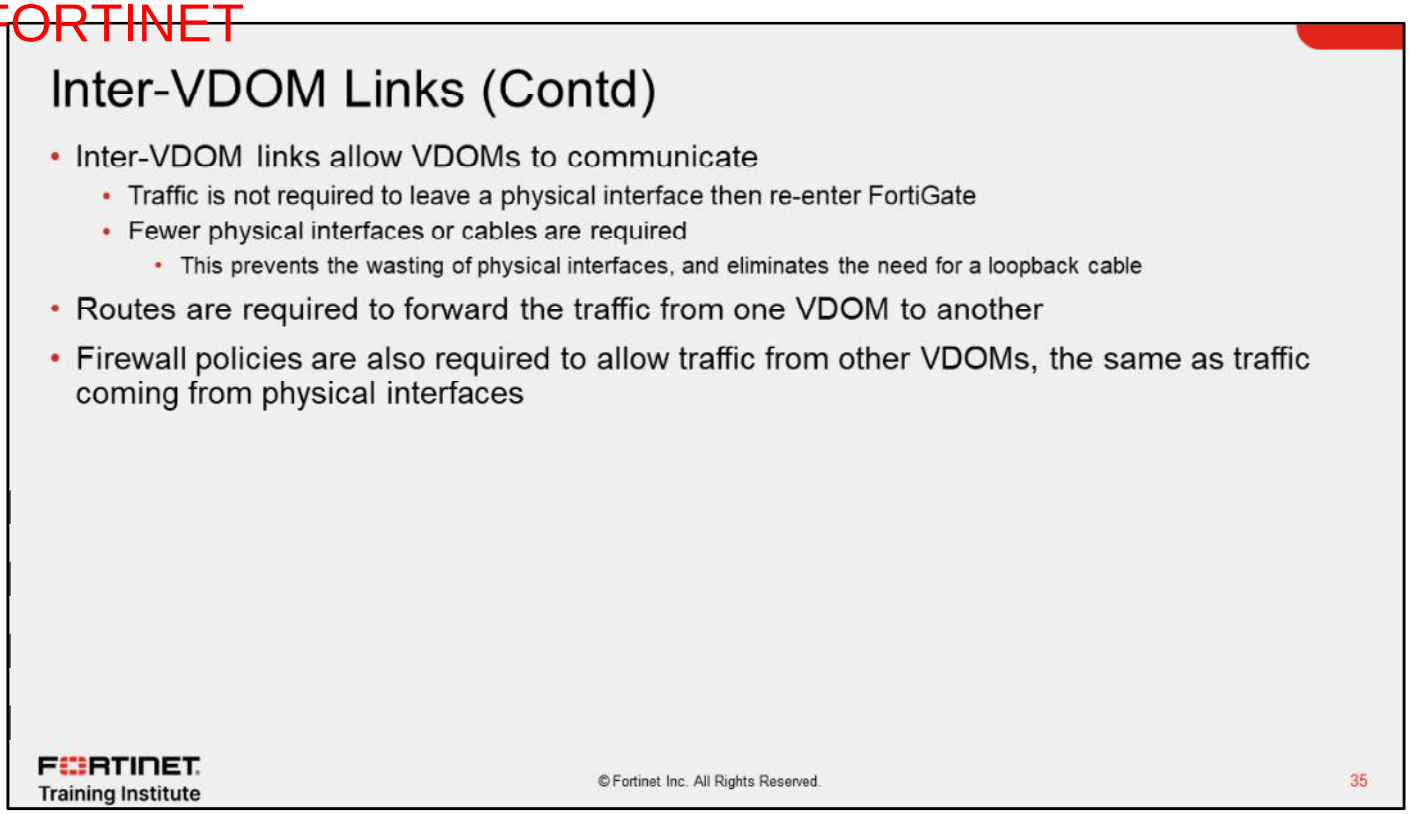

When creating inter-VDOM links, you must create the virtual interfaces. You must also create the appropriate firewall policies in each VDOM, just as you would if the traffic were arriving on a network cable, otherwise, FortiGate will block it.

Additionally, routes are required to correctly route packets between two VDOMs.

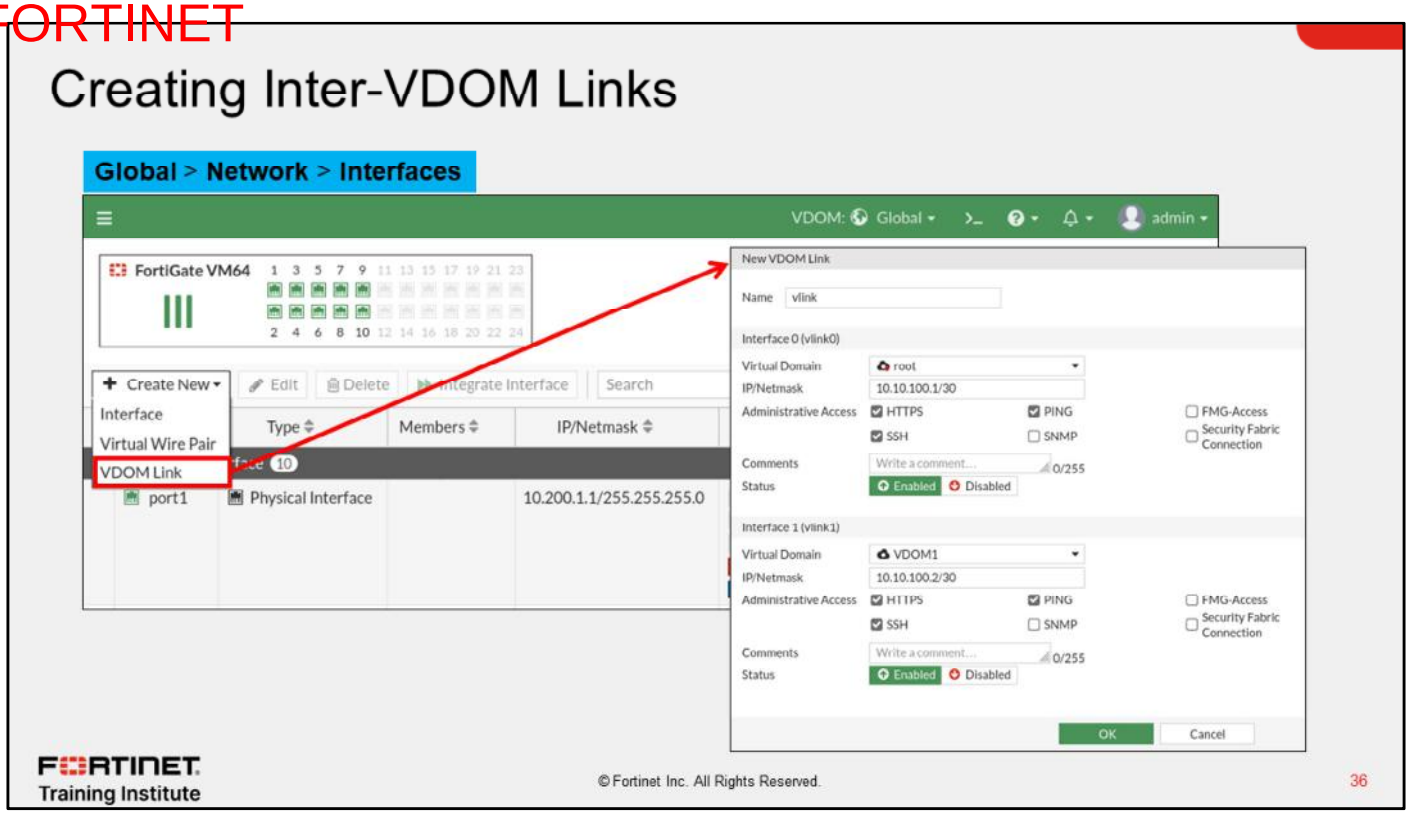

On the GUI, you create a network interface in the **Global** settings. To create the virtual interface, click **Create New**, and then select **VDOM Link**.

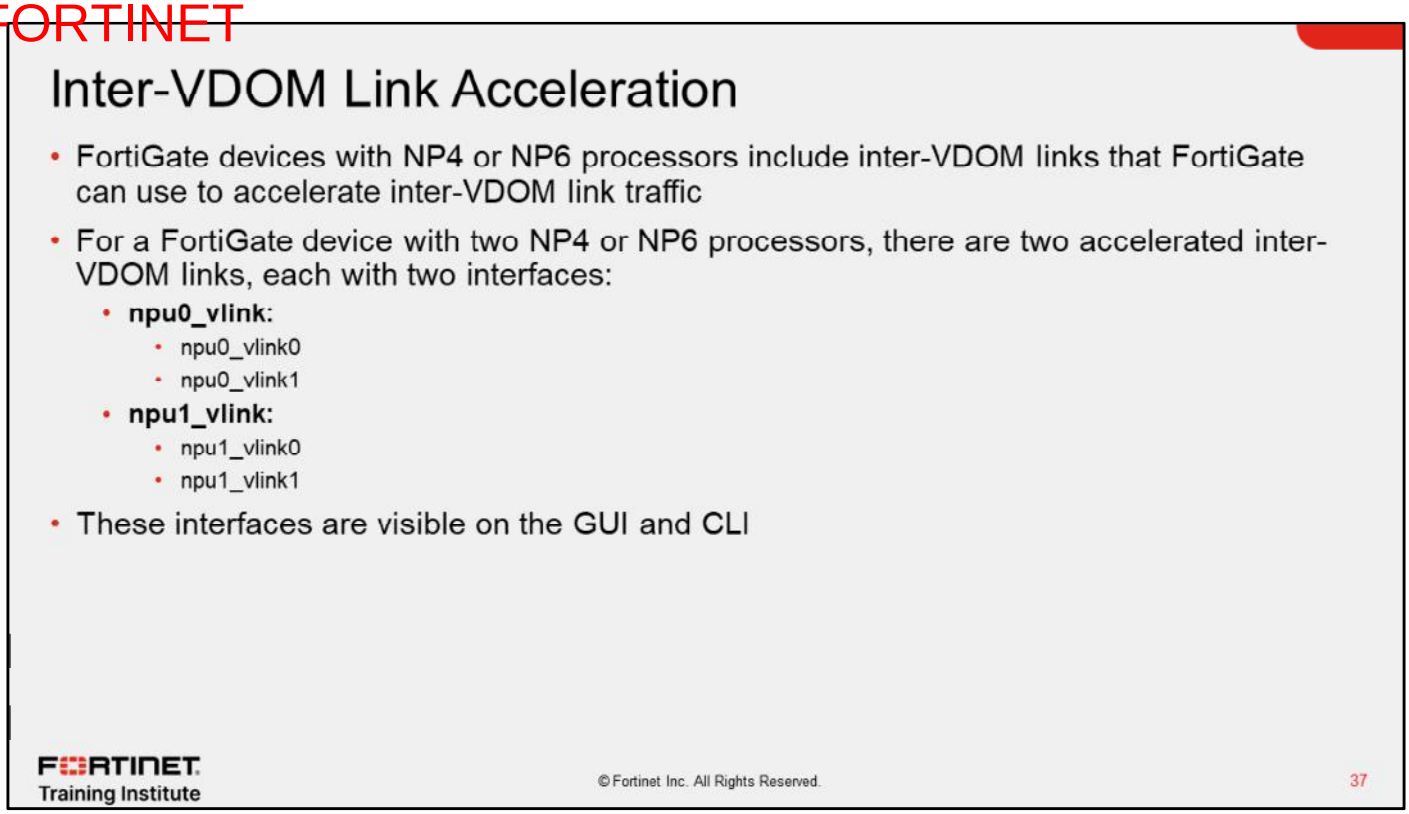

FortiGate devices with NP4 or NP6 processors include inter-VDOM links that FortiGate can use to accelerate inter-VDOM link traffic. For a FortiGate with two NP4 or NP6 processors, there are two accelerated inter-VDOM links, each with two interfaces:

- **npu0_vlink:**
  - npu0_vlink0
  - npu0_vlink1
- **npu1_vlink:**
  - npu1_vlink0
  - npu1_vlink1

These interfaces are visible on the GUI and CLI. By default, the interfaces in each inter-VDOM link are assigned to the root VDOM. To use these interfaces to accelerate inter-VDOM link traffic, assign each interface in the pair to the VDOMs that you want to offload traffic between. For example, if you have added a VDOM named *New-VDOM* to a FortiGate with NP4 processors, you can click **System** > **Network** > **Interfaces** and edit the **npu0-vlink1** interface and set the VDOM to *New-VDOM***.** This results in an accelerated inter-VDOM link between *root* and *New-VDOM.*

Knowledge Check

1. What is a requirement for creating an inter-VDOM link between two VDOMs?
   - A. The NGFW mode of at least one VDOM must be profile based.
   - ✓ B. At least one of the VDOMs must be operating in NAT mode.

2. Which type of VDOM link requires that both sides of the link be assigned an IP address within the same subnet?
   - A. NAT-to-transparent
   - ✓ B. NAT-to-NAT

FORTINET Training Institute

38

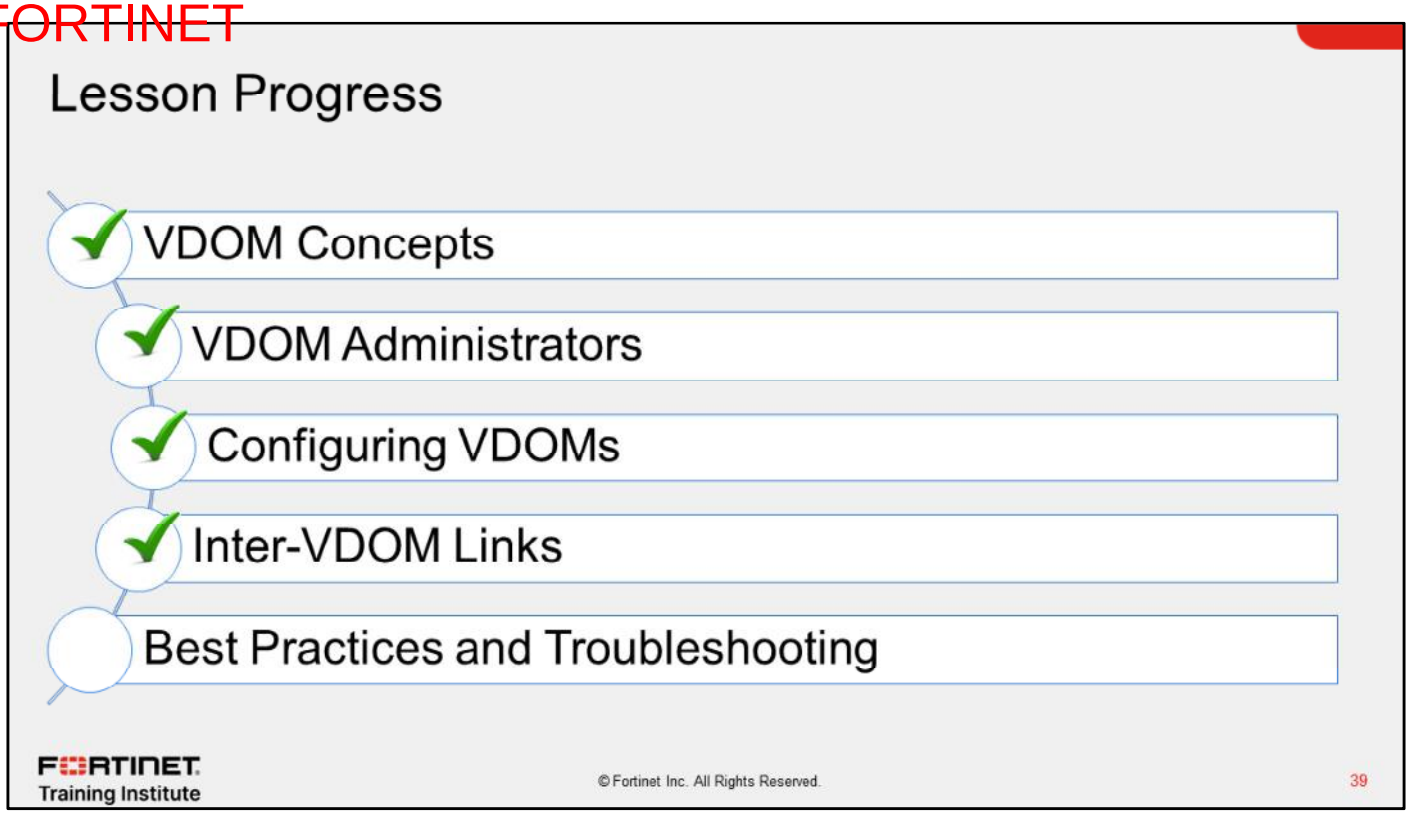

Good job! You now understand inter-VDOM Links.

Now, you'll learn about VDOM best practices and troubleshooting.

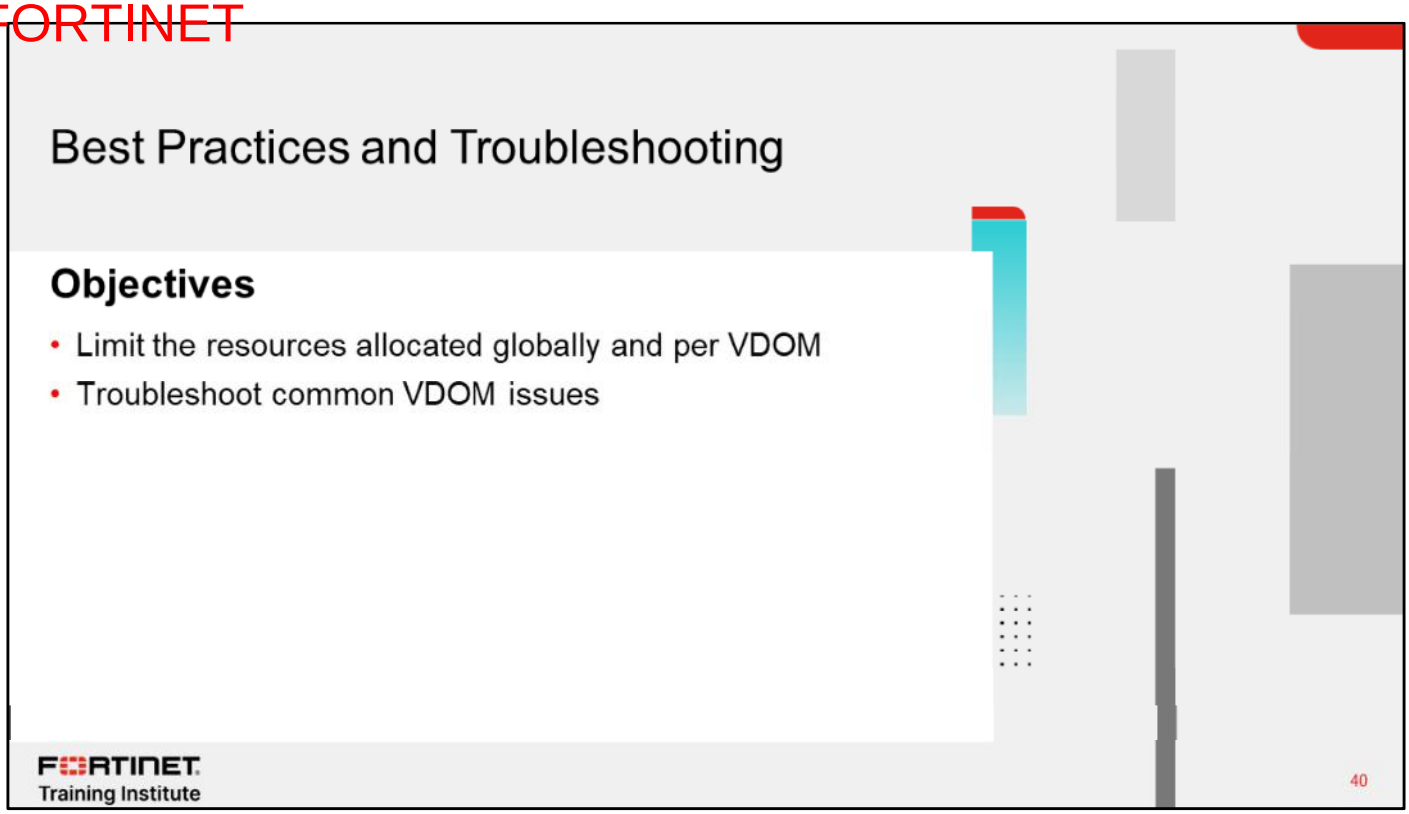

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in VDOM best practices and troubleshooting, you will be able to prevent, identify, and solve common VDOM issues.

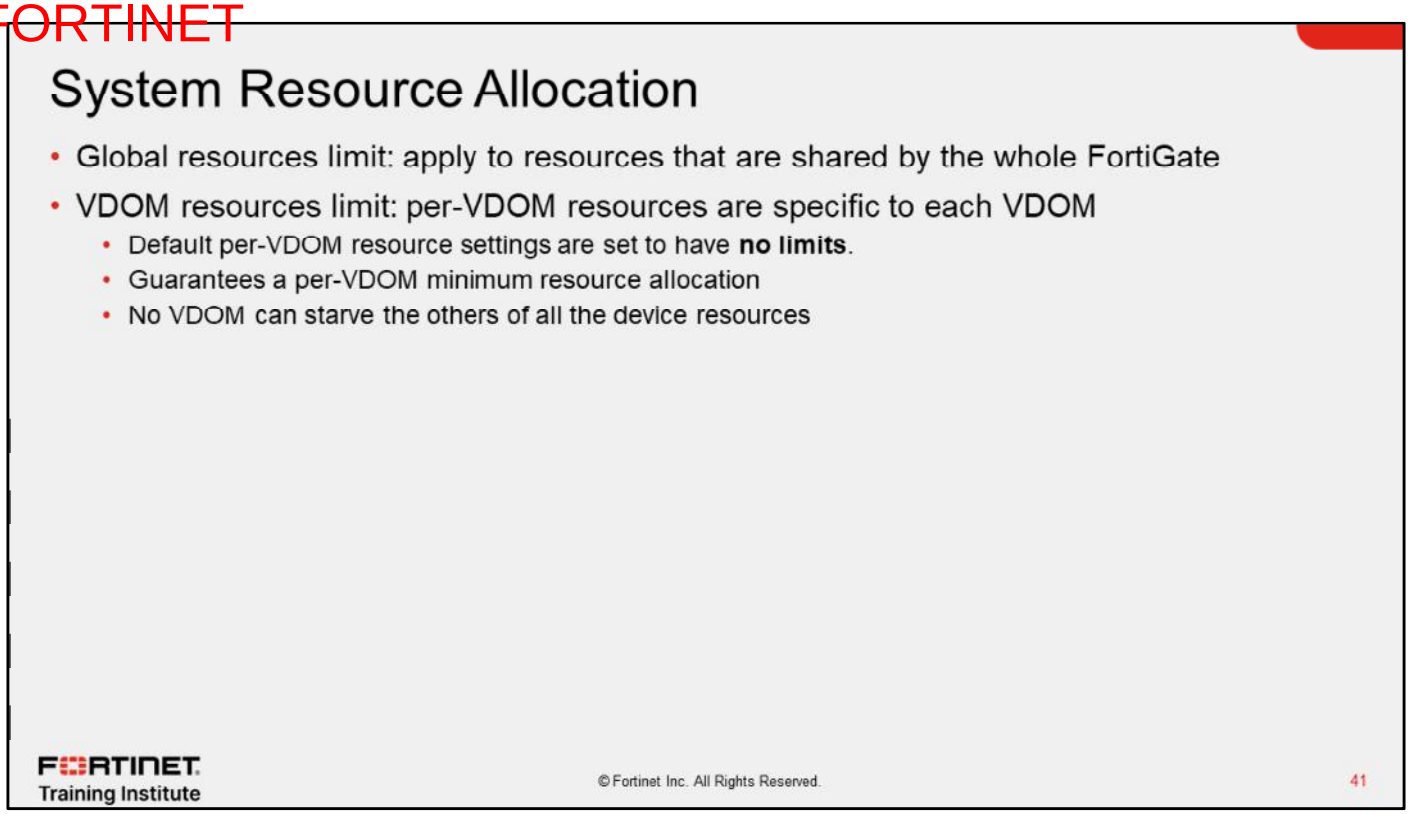

Remember, VDOMs are only a *logical* separation—each VDOM shares physical resources with the others.

By default, all per-VDOM resource settings are set to have no limits. This means that any single VDOM can use all of the FortiGate device's resources. This could deprive other VDOMs of the resources that they require, to the point that could be unable to function.

Unlike FortiGate-VM, VDOMs are not allocated and balanced with weighted vCPU cores, vRAM, and other virtualized hardware.
To fine-tune performance, you can configure resource limits for each feature—IPsec tunnels, address objects, and so on—at the global level and at each VDOM level. This controls the ratio of the system resource usage of each VDOM to the total available resources.

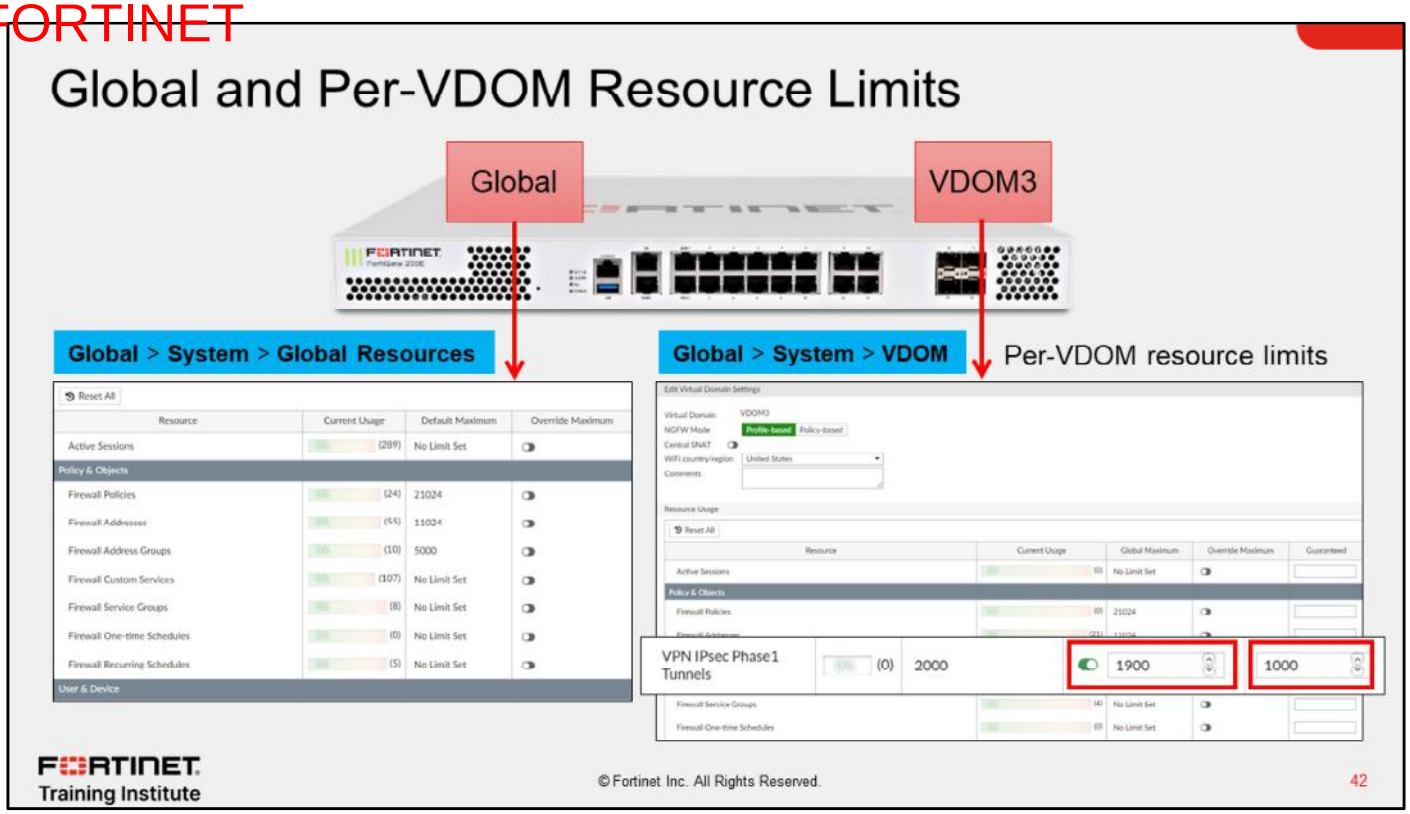

For example, a FortiGate with hardware powerful enough to handle up to 2000 IPsec VPN tunnels and configured with three VDOMs, could be configured as follows to meet specific criteria: VDOM1 and VDOM2 don't use IPsec VPN tunnels often. So, they are allowed to have up to 50 tunnels each. VDOM3, however, uses VPN extensively. Therefore, this FortiGate device is configured to allow VDOM3 to have up to 1900 tunnels, with 1000 guaranteed.

Configure your FortiGate device with global limits for critical features, such as sessions, policies, and so on. Then, configure each VDOM with its own quotas and minimums, within the global limits.

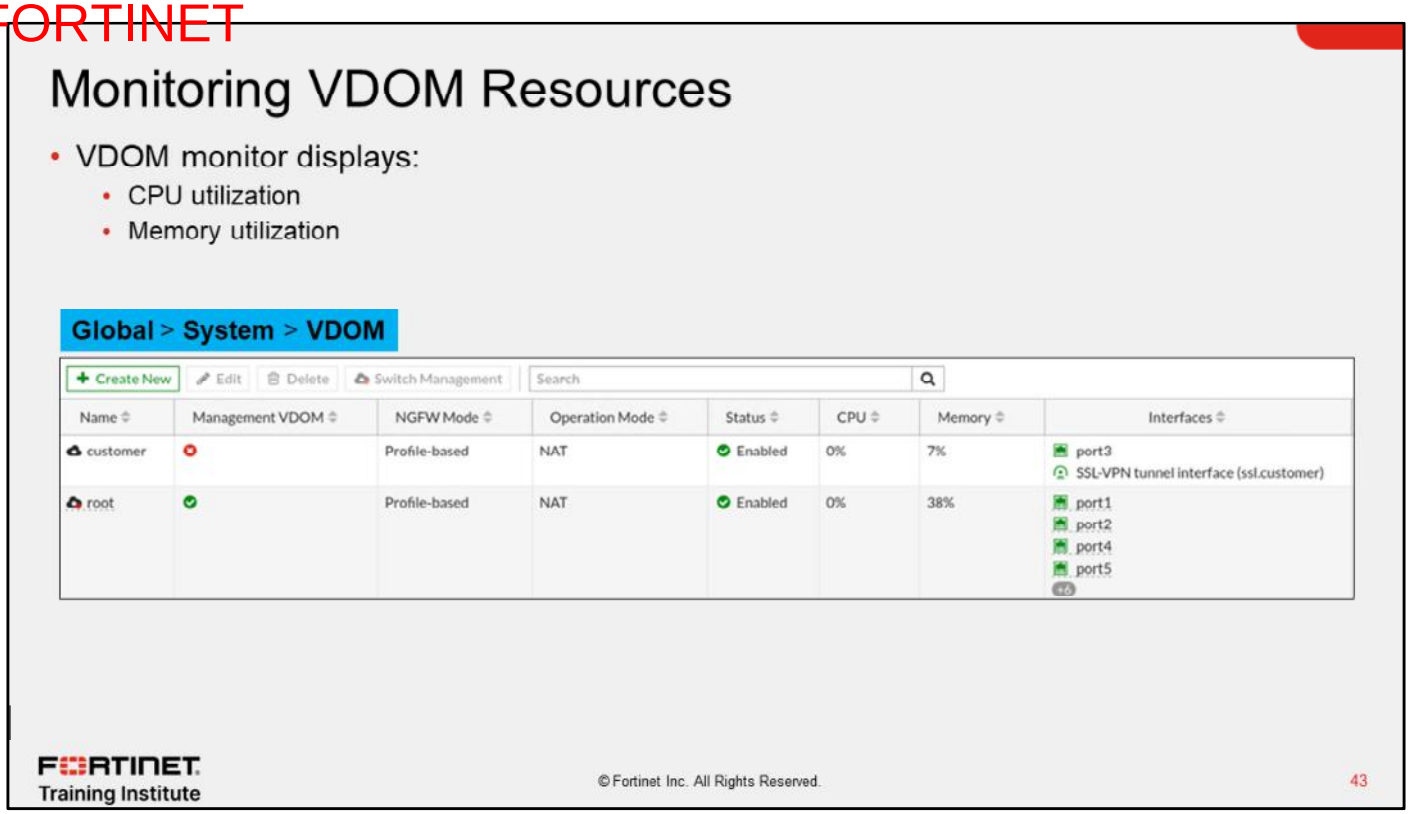

On the GUI, you can click **Global** > **System** > **VDOM** to see the VDOM monitor. It displays the CPU and memory usage for each VDOM.

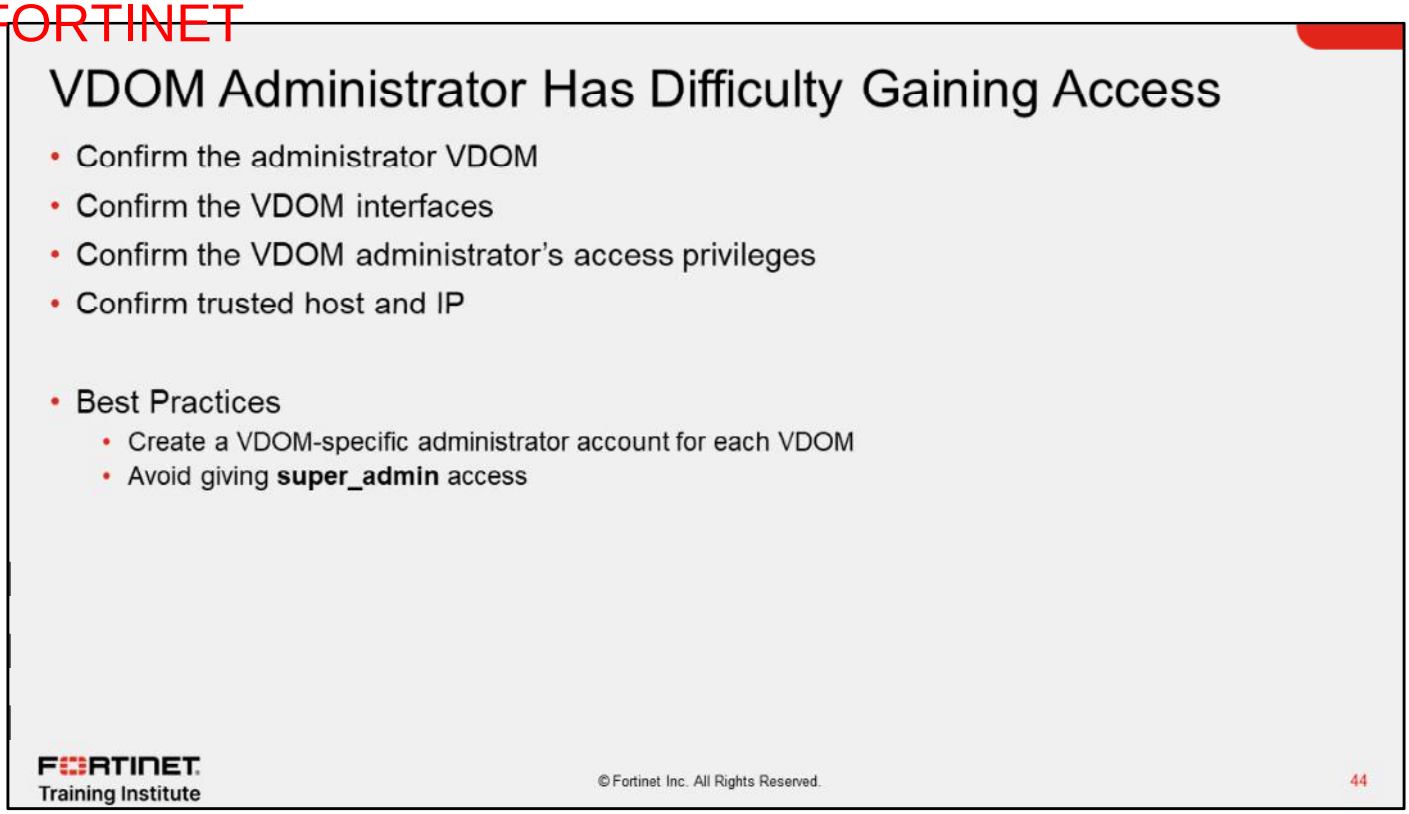

With VDOMs configured, administrators have an extra layer of permissions and may have problems accessing the desired information. If an administrator cannot gain access, check the following:

- Confirm the administrator's VDOM: each administrator account, other than the **super_admin** account, is tied to one or more specific VDOMs. That administrator is not able to access any other VDOM. It may be possible they are trying to access the wrong VDOM (one that they do not have permissions for).

- Confirm the VDOM interfaces: an administrator can access their VDOM only through interfaces that are assigned to that VDOM. If interfaces on that VDOM are disabled or unavailable, there will be no method of accessing that VDOM by its local administrator. The **super_admin** is required to either bring up the interfaces, fix the interfaces, or move another interface to that VDOM to restore access.

- Confirm the VDOM admin access: as with all FortiGate devices, administration access on the VDOM's interfaces must be enabled for the administrators of that VDOM to gain access. For example, if SSH is not enabled, that is not available to administrators. To enable admin access, the **super_admin** clicks **Global** > **Network** > **Interfaces** and enables administrator access for the interface in question.

- Confirm trusted host and IP: if trusted hosts are enabled on the administrator account, ensure the user is connecting from the correct, specified host address, and that no intermediate devices are performing NAT functions on the connection.

Best practice dictates that you should usually avoid unnecessary security holes. Do not provide **super_admin** access, if possible. Instead, restrict each administrator to their relevant domain. That way, they cannot accidentally or maliciously impact other VDOMs, and any damage or mistakes will be limited in scope.

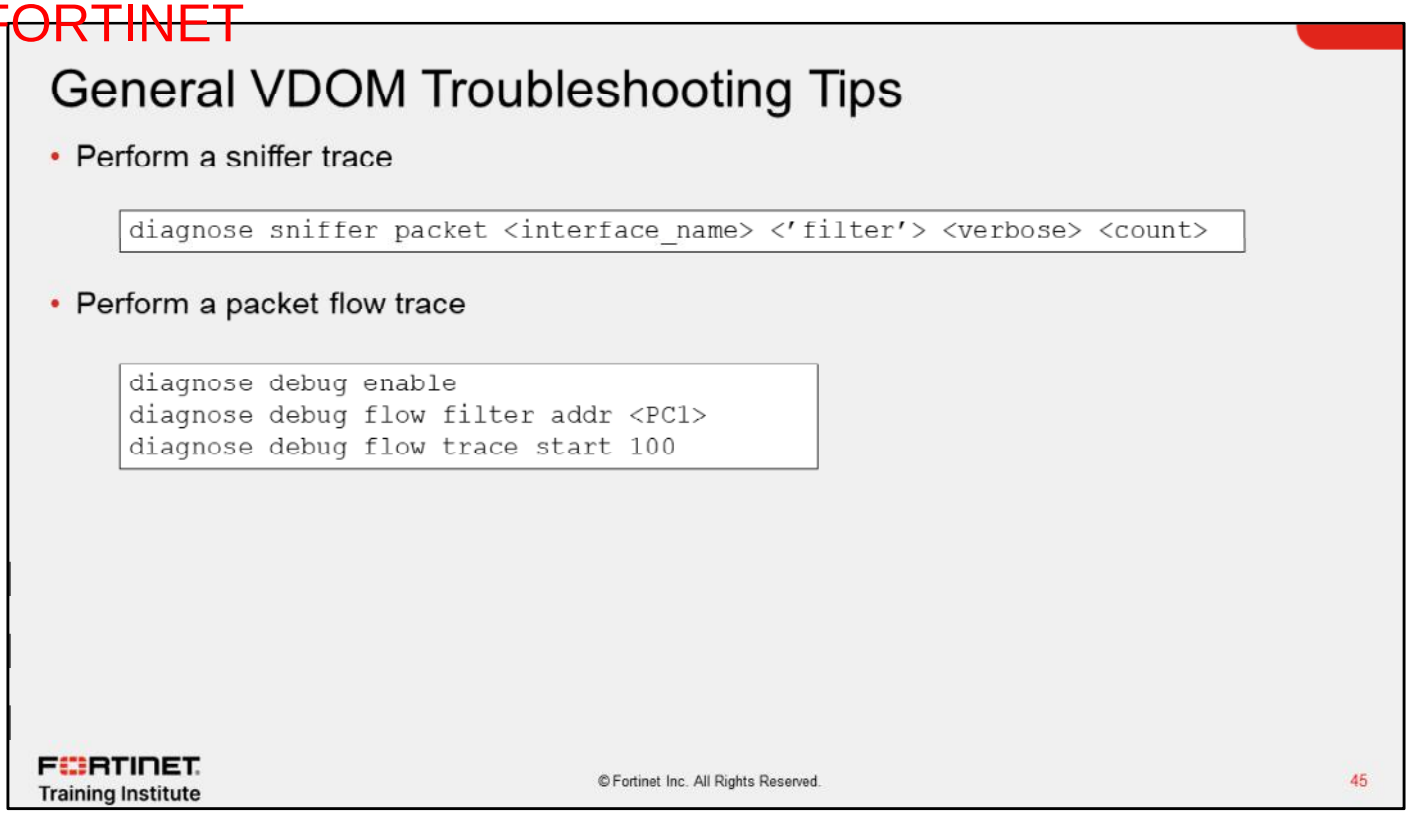

Besides ping and traceroute, there are additional tools for troubleshooting your VDOM configurations. The primary tools for VDOM troubleshooting include packet sniffing and debugging the packet flow.

- Perform a sniffer trace: when troubleshooting networks, it helps to look inside the headers of packets to determine if they are traveling along the expected route. Packet sniffing can also be called a network tap, packet capture, or logic analyzing. The sniffer also indicates what traffic is entering or leaving the egress and ingress interfaces in all VDOMS. This makes it extremely useful for troubleshooting inter-VDOM routing issues.

- Debug the packet flow: traffic should enter and leave the VDOM. If you have identified that network traffic is not entering and leaving the VDOM as expected, debug the packet flow. You can debug only using CLI commands. This tool provides more granular details for help in troubleshooting inter-VDOM traffic because it gives details of routing selection, NAT, and policy selection.

Knowledge Check

1. Of these options, what is a possible reason why an administrator might not be able to gain access to a specific VDOM?
   - ✓ A. The administrator is using an IP address that is not specified as a trusted host.
   - B. The administrator is using the super_admin profile.

2. Which troubleshooting tool is most suitable when trying to verify the firewall policy used by an inter-VDOM link?
   - A. Sniffer trace
   - ✓ B. Packet flow trace

FORTINET Training Institute

46

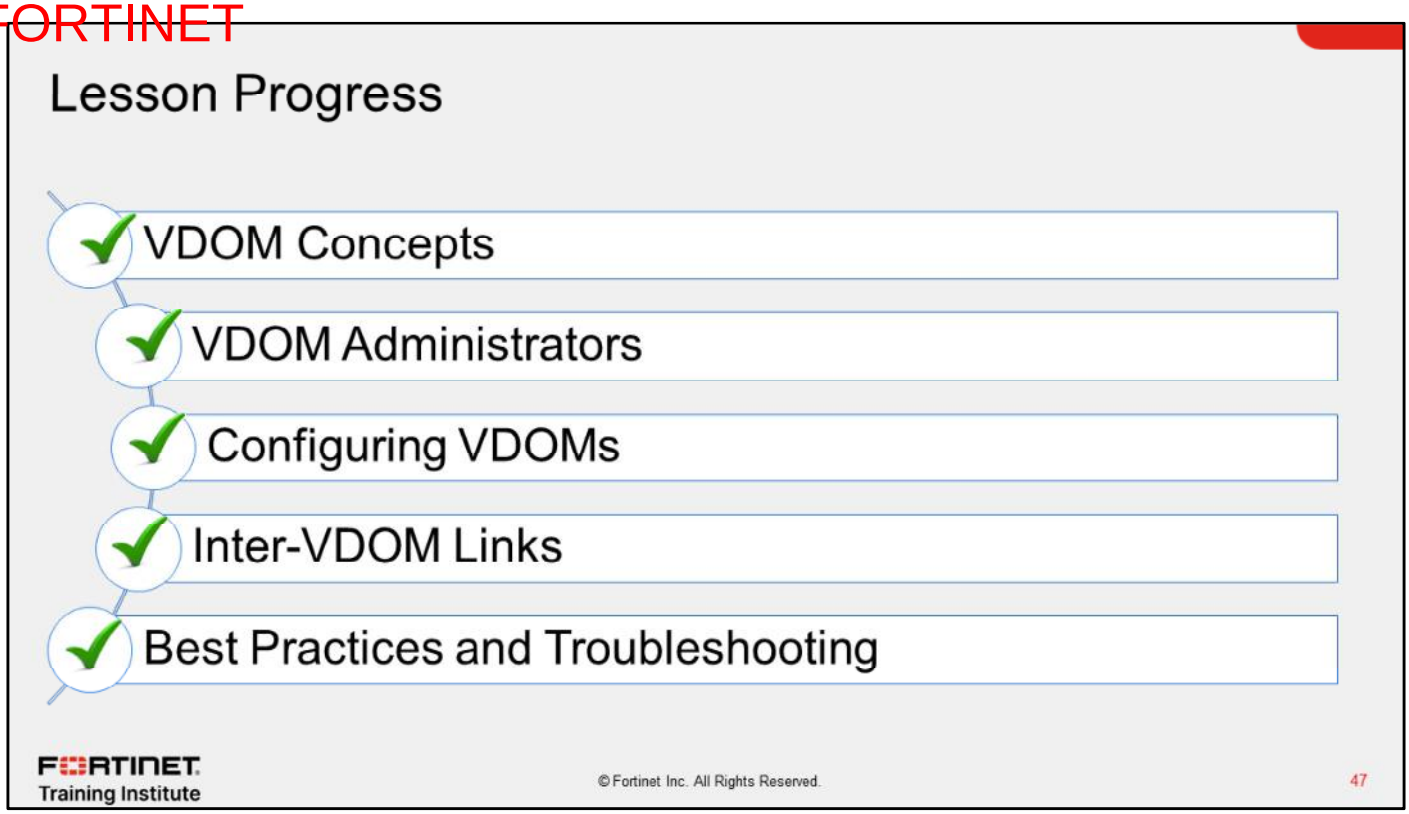

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

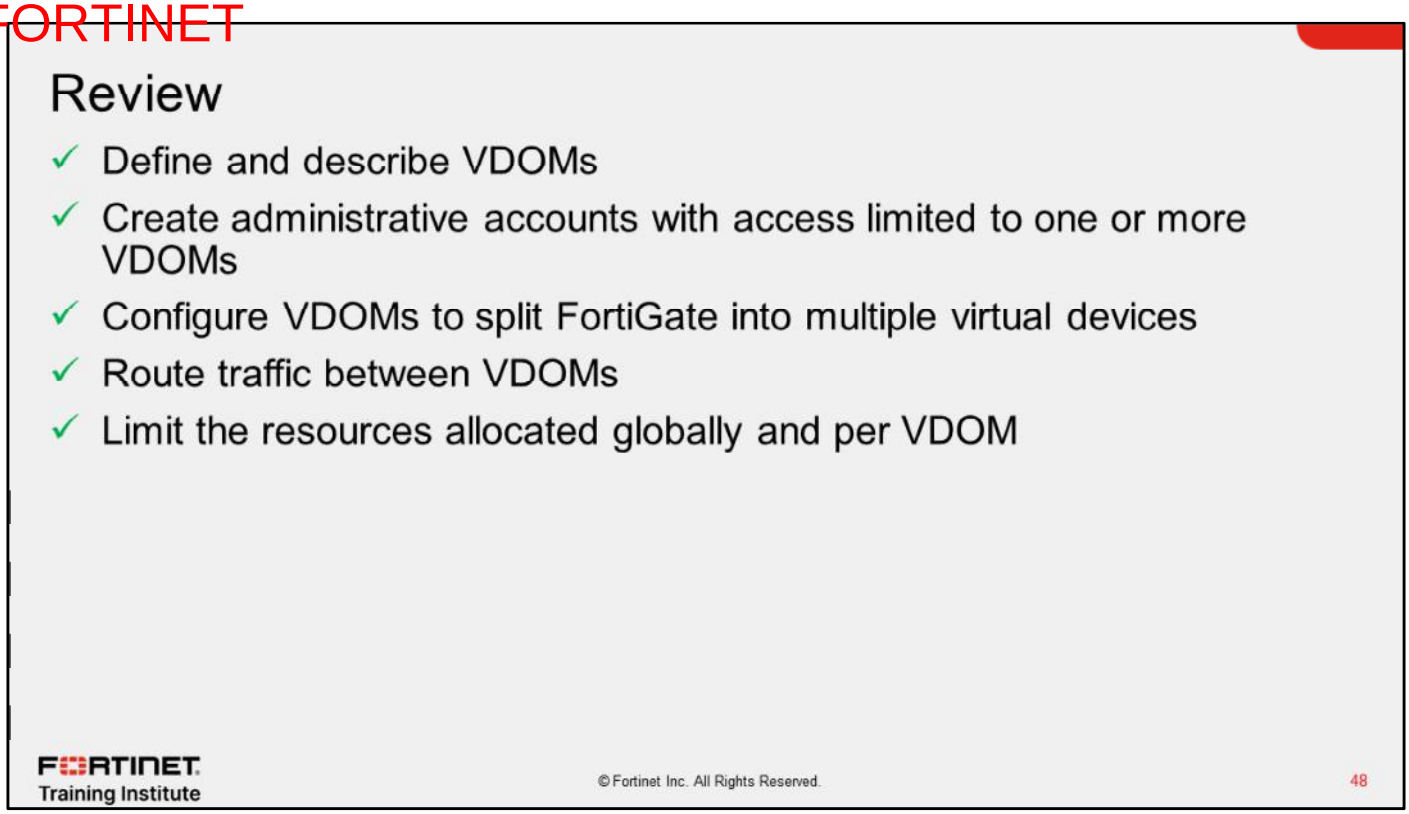

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure VDOMs, and examined examples of common use.

In this lesson, you will learn about Fortinet single sign-on (FSSO). When you use this feature, your users don't need to log on each time they access a different network resource.

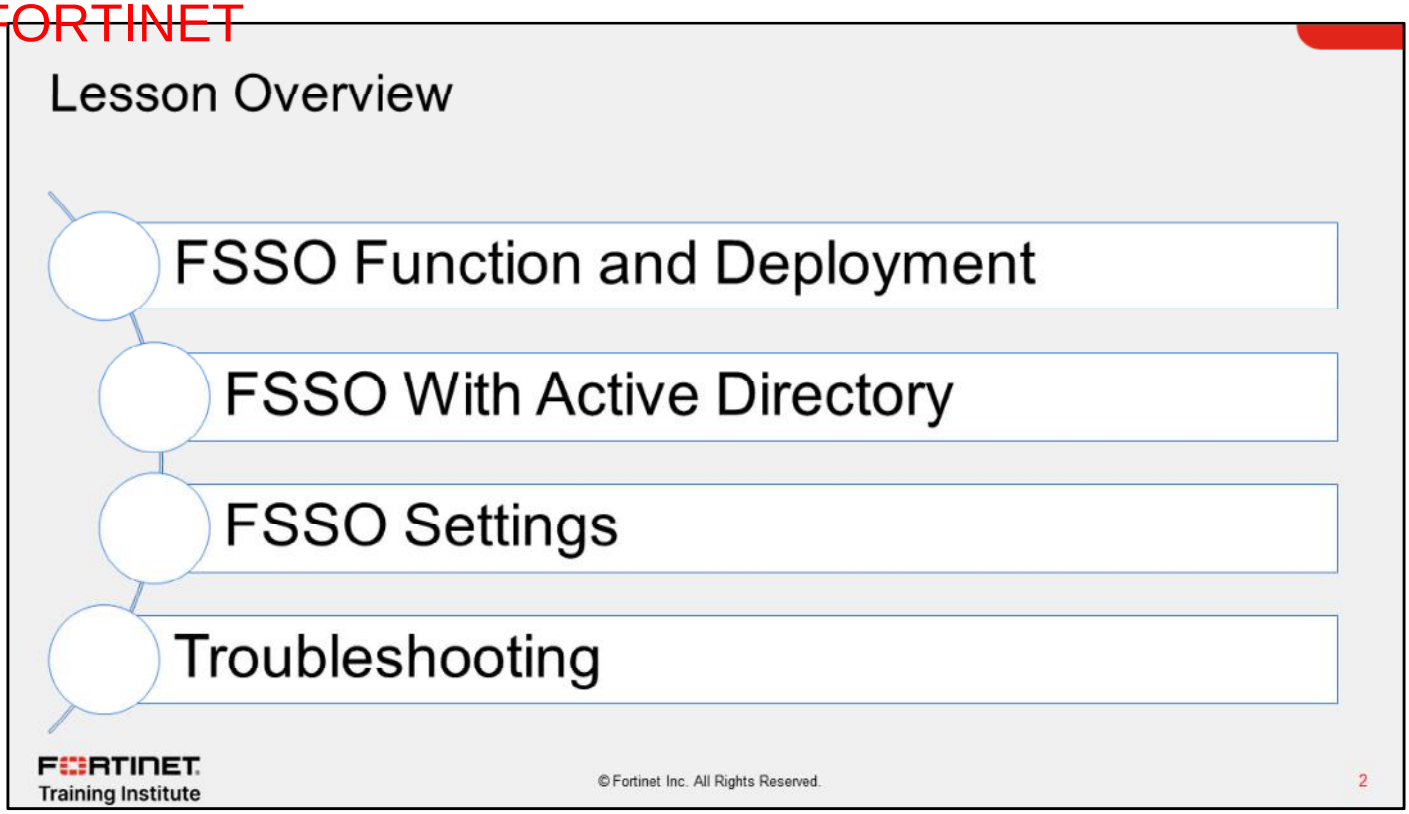

In this lesson, you will learn about the topics shown on this slide.

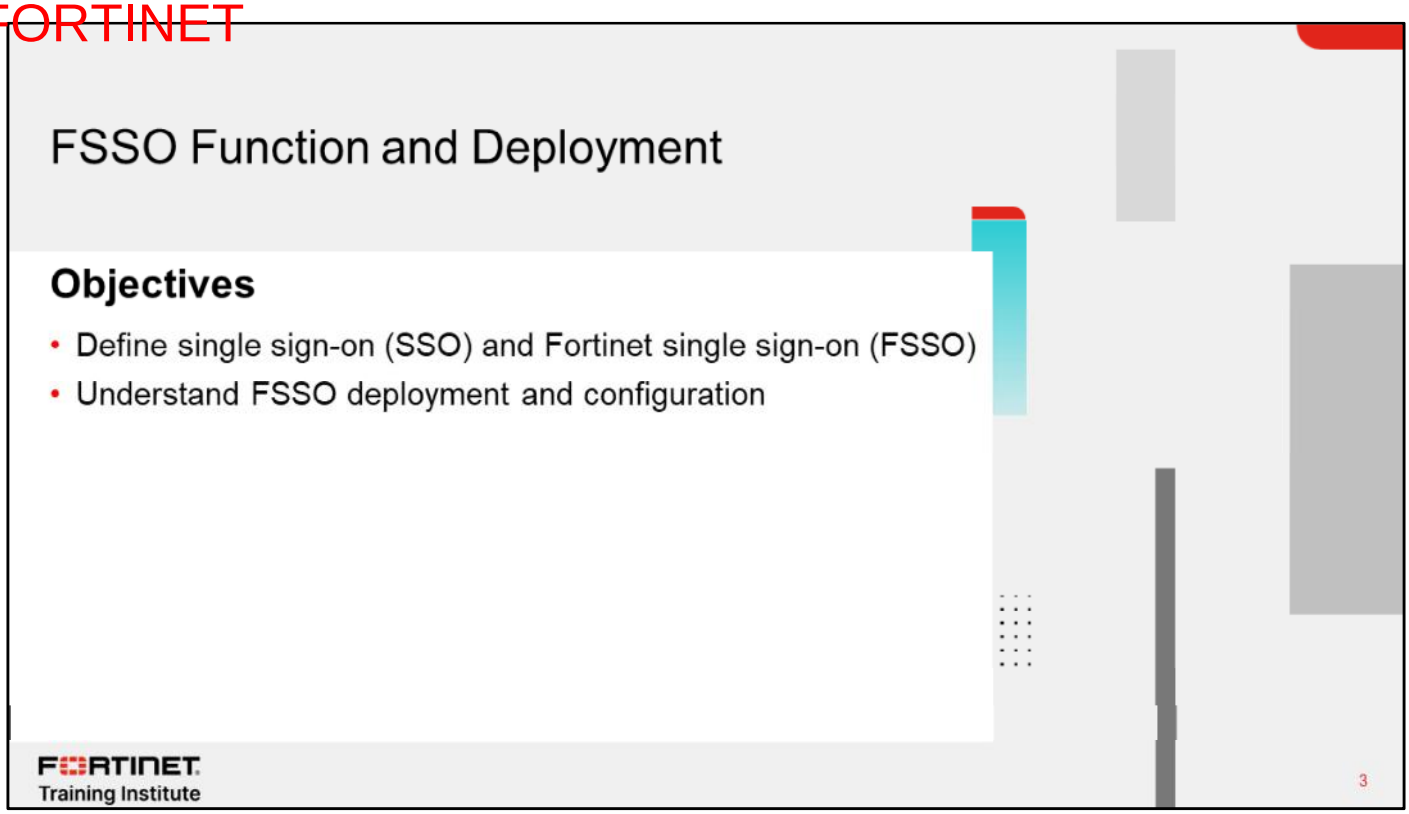

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding SSO concepts, you will be able to more effectively understand FSSO methods.

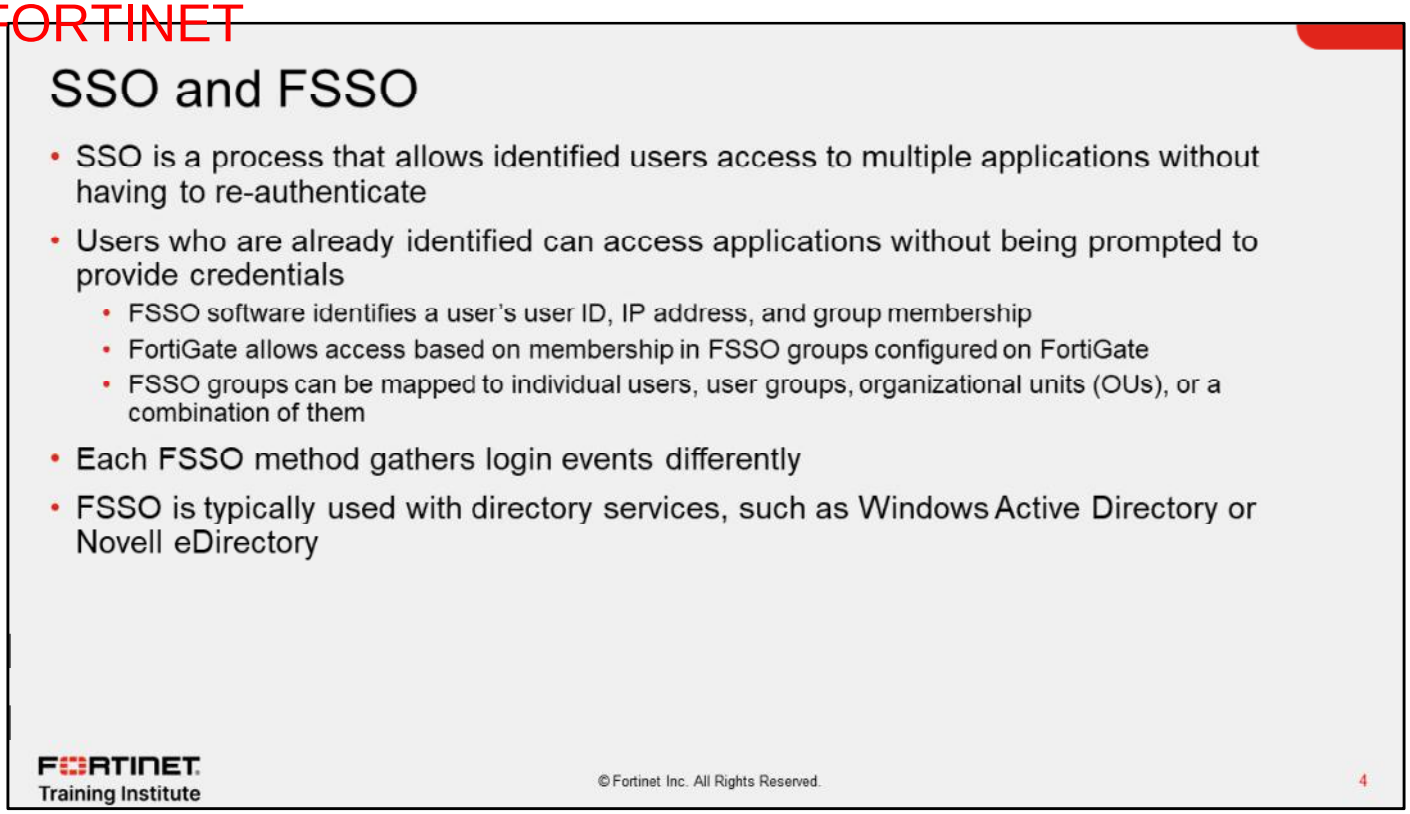

SSO is a process that allows users to be automatically logged in to every application after being identified, regardless of platform, technology, and domain.

FSSO is a software agent that enables FortiGate to identify network users for security policies or for VPN access, in advanced deployments with FortiAuthenticator, without asking for their username and password. When a user logs in to a directory service, the FSSO agent sends FortiGate the username, the IP address, and the list of groups that the user belongs to. FortiGate uses this information to maintain a local database of usernames, IP addresses, and group mappings.

Because the domain controller authenticates users, FortiGate does not perform authentication. When the user tries to access network resources, FortiGate selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

FSSO is typically used with directory service networks such as Windows Active Directory or Novell eDirectory.

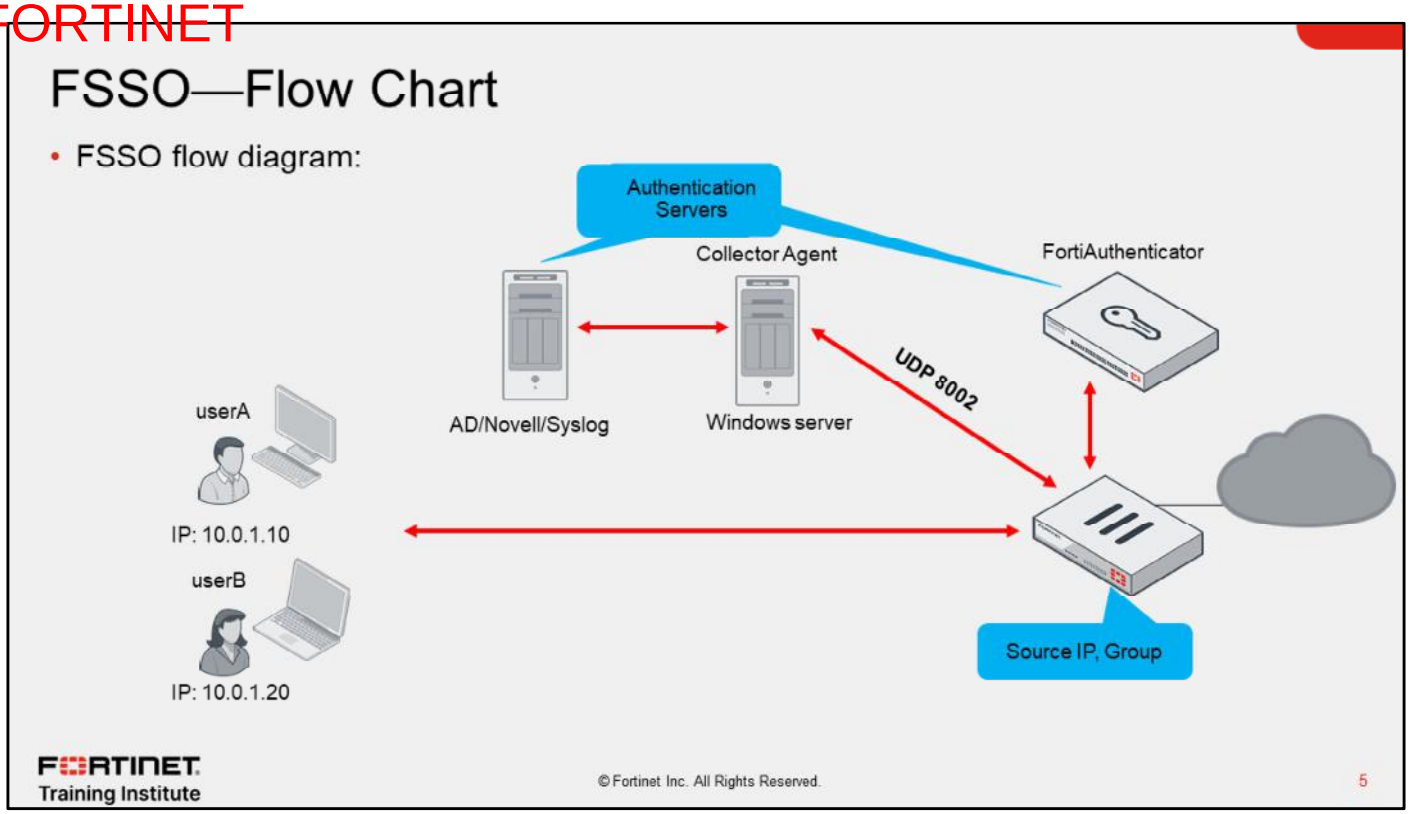

This slide shows the FSSO flow we discussed in the previous slide.

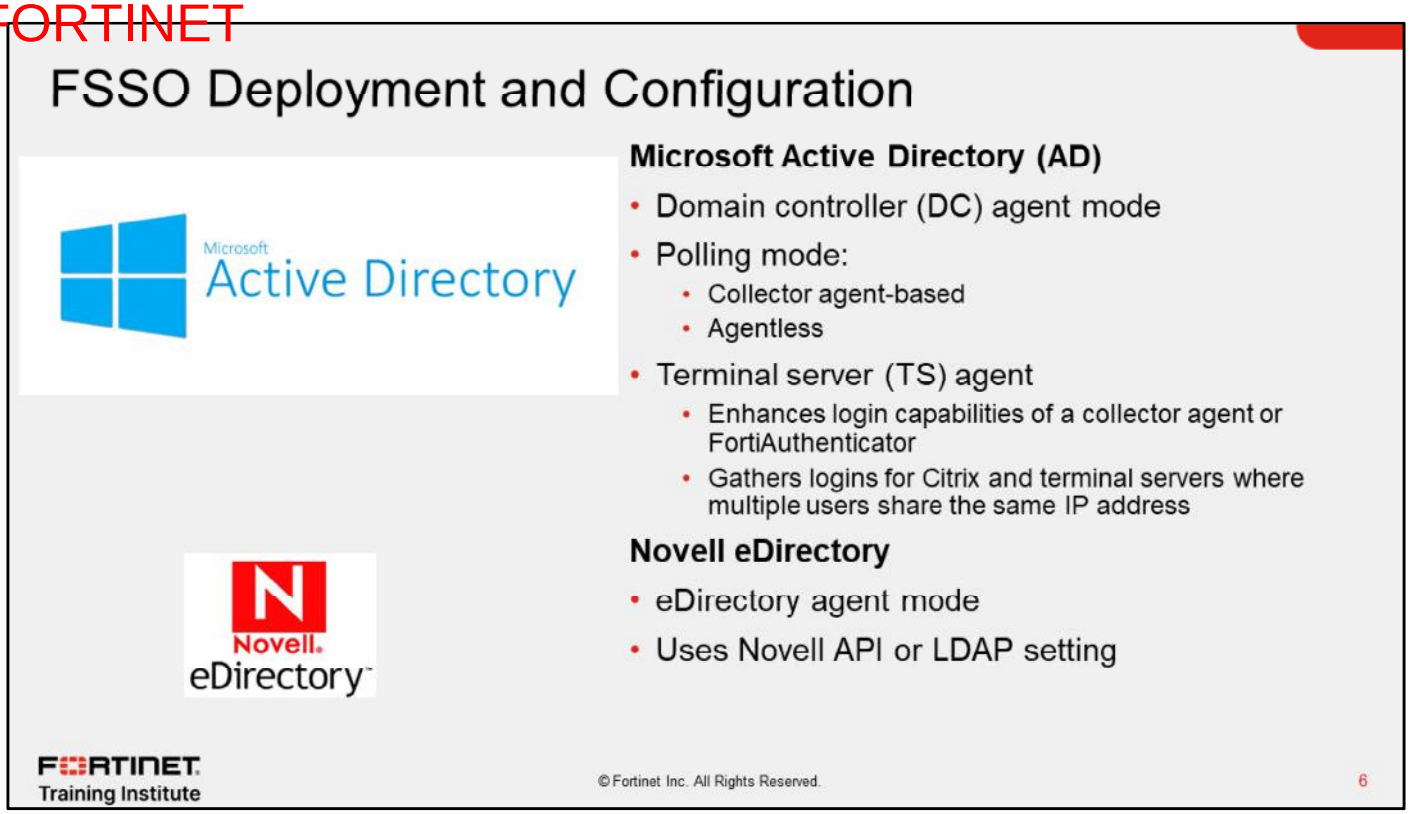

How you deploy and configure FSSO depends on the server that provides your directory services.

FSSO for Windows Active Directory (AD) uses a collector agent. Domain controller (DC) agents may also be required, depending on the collector agent working mode. There are two working modes that monitor user sign-on activities in Windows: DC agent mode and polling mode. FortiGate also offers a polling mode that does not require a collector agent, which is intended for simple networks with a minimal number of users.

There is another kind of DC agent that is used exclusively for Citrix and terminal services environments: terminal server (TS) agents. TS agents require the Windows Active Directory collector agent or FortiAuthenticator to collect and send the login events to FortiGate.

The eDirectory agent is installed on a Novell network to monitor user sign-ons and send the required information to FortiGate. It functions much like the collector agent on a Windows AD domain controller. The agent can obtain information from the Novell eDirectory using either the Novell API or LDAP.

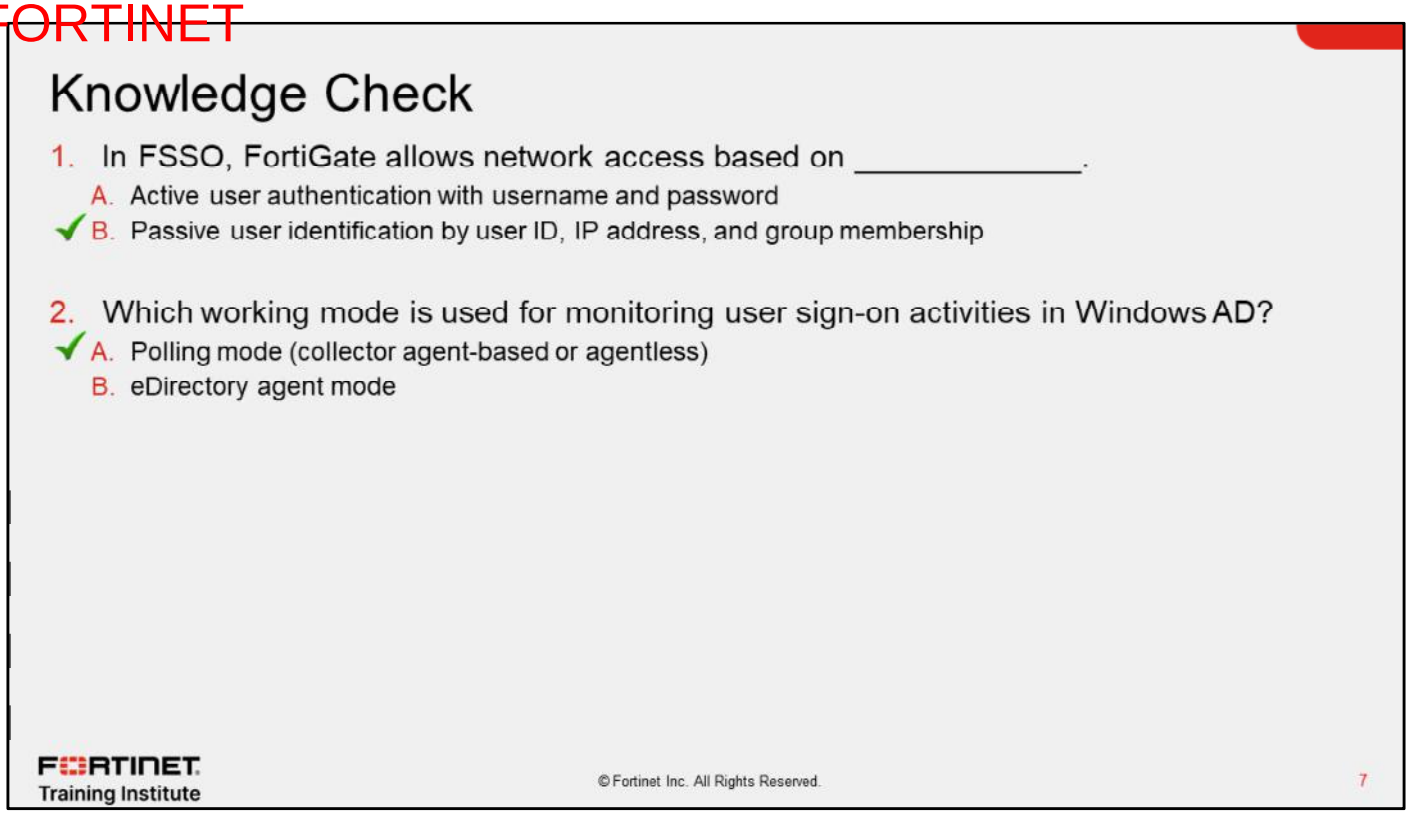

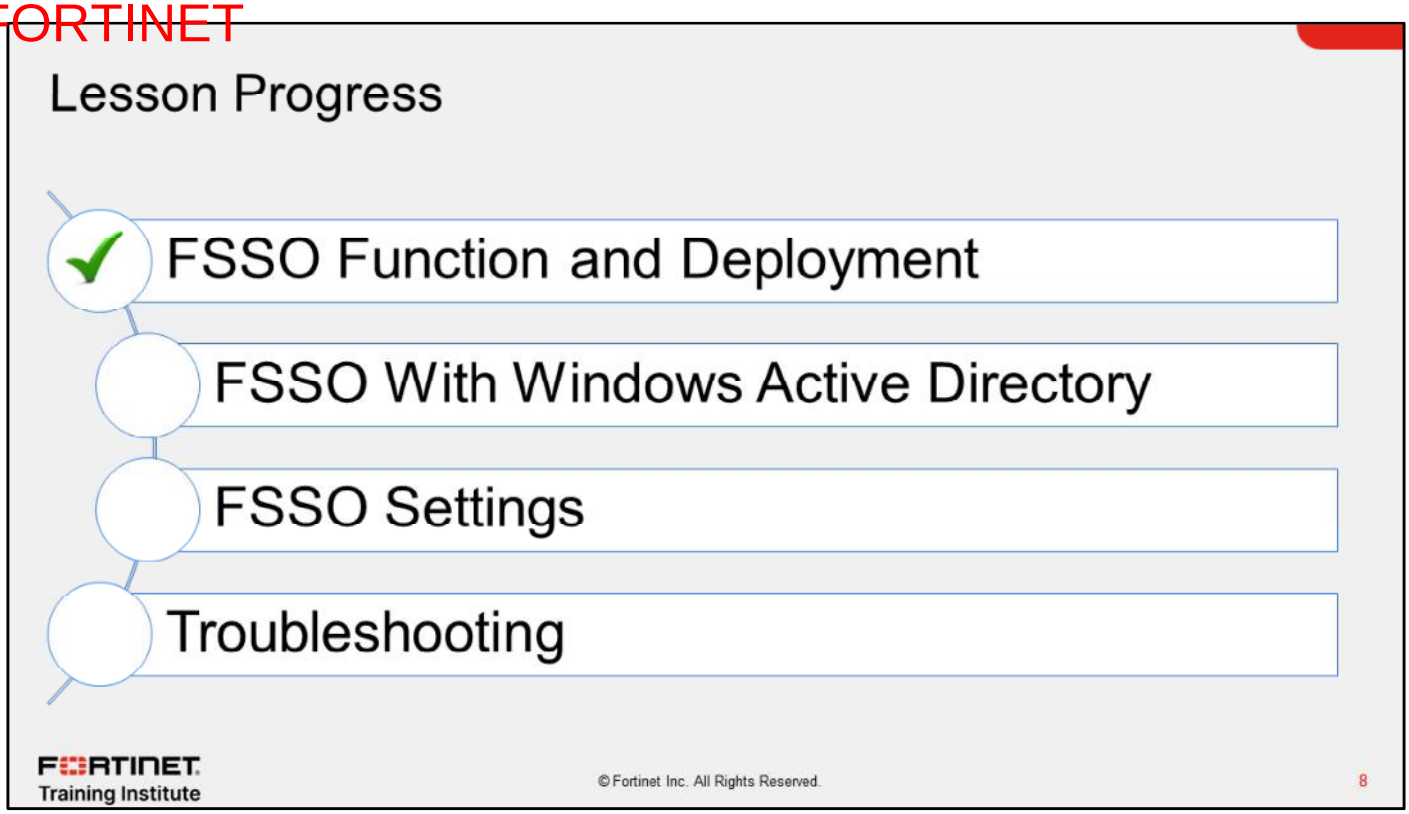

Good job! You now understand basic concepts about the function of FSSO and how it is deployed.

Now, you'll learn about user login events in Windows Active Directory using FSSO.

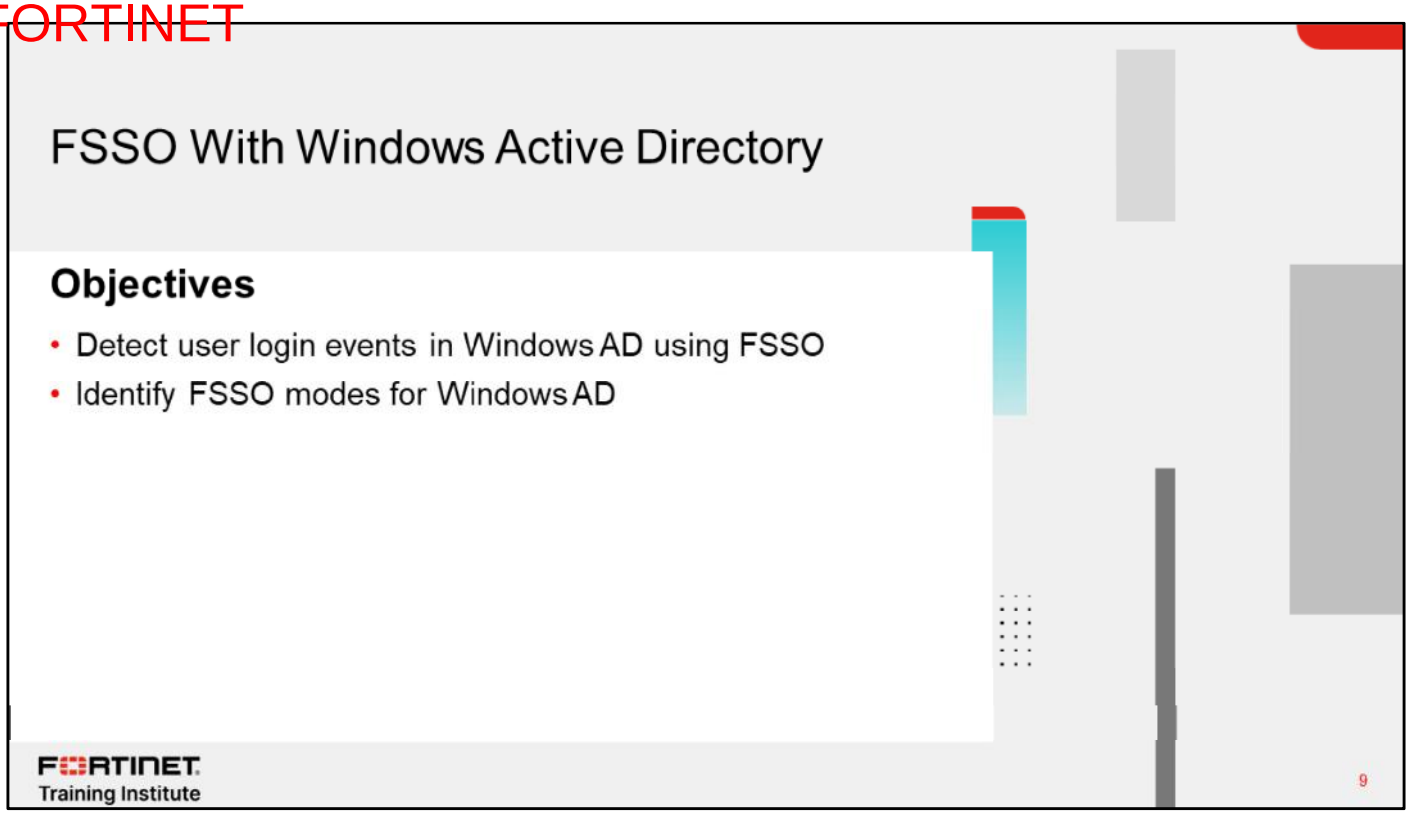

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding the different ways you can configure FSSO for Windows AD, you will be better able to design the architecture of your SSO system.

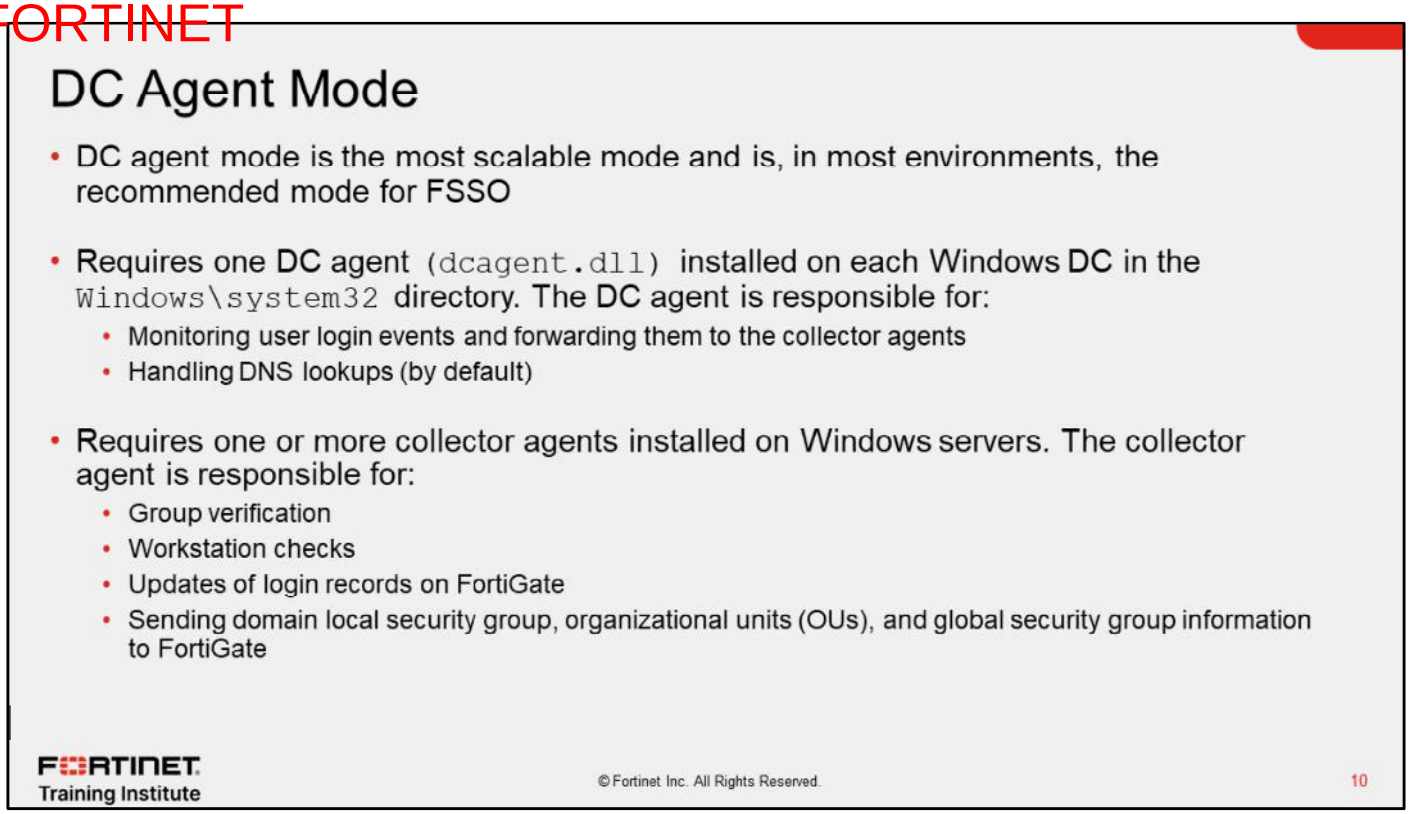

DC agent mode is considered the recommended mode for FSSO.

DC agent mode requires:

- One DC agent installed on each Windows DC
  If you have multiple DCs, this means that you need multiple DC agents. DC agents monitor and forward user login events to the collector agents.

- A collector agent, which is another FSSO component
  The collector agent is installed on a Windows server that is a member of the domain you are trying to monitor. It consolidates events received from the DC agents, then forwards them to FortiGate. The collector agent is responsible for group verification, workstation checks, and FortiGate updates of login records. The FSSO collector agent can send domain local security group, organizational units (OUs), and global security group information to FortiGate devices. It can also be customized for DNS lookups.

  When the user logs on, the DC agent intercepts the login event on the domain controller. It then resolves the DNS of the client, and sends it to the collector agent.

  The collector agent receives it and then performs a DNS resolution in order to check if the IP of the user has changed.

  In some configurations, double DNS resolution is a problem. In this case, you may configure a registry key on the domain controller that hosts the DC agent in order not to resolve the DNS:
  `donot_resolve = (DWORD) 1 at HKLM/Software/Fortinet/FSAE/dcagent`

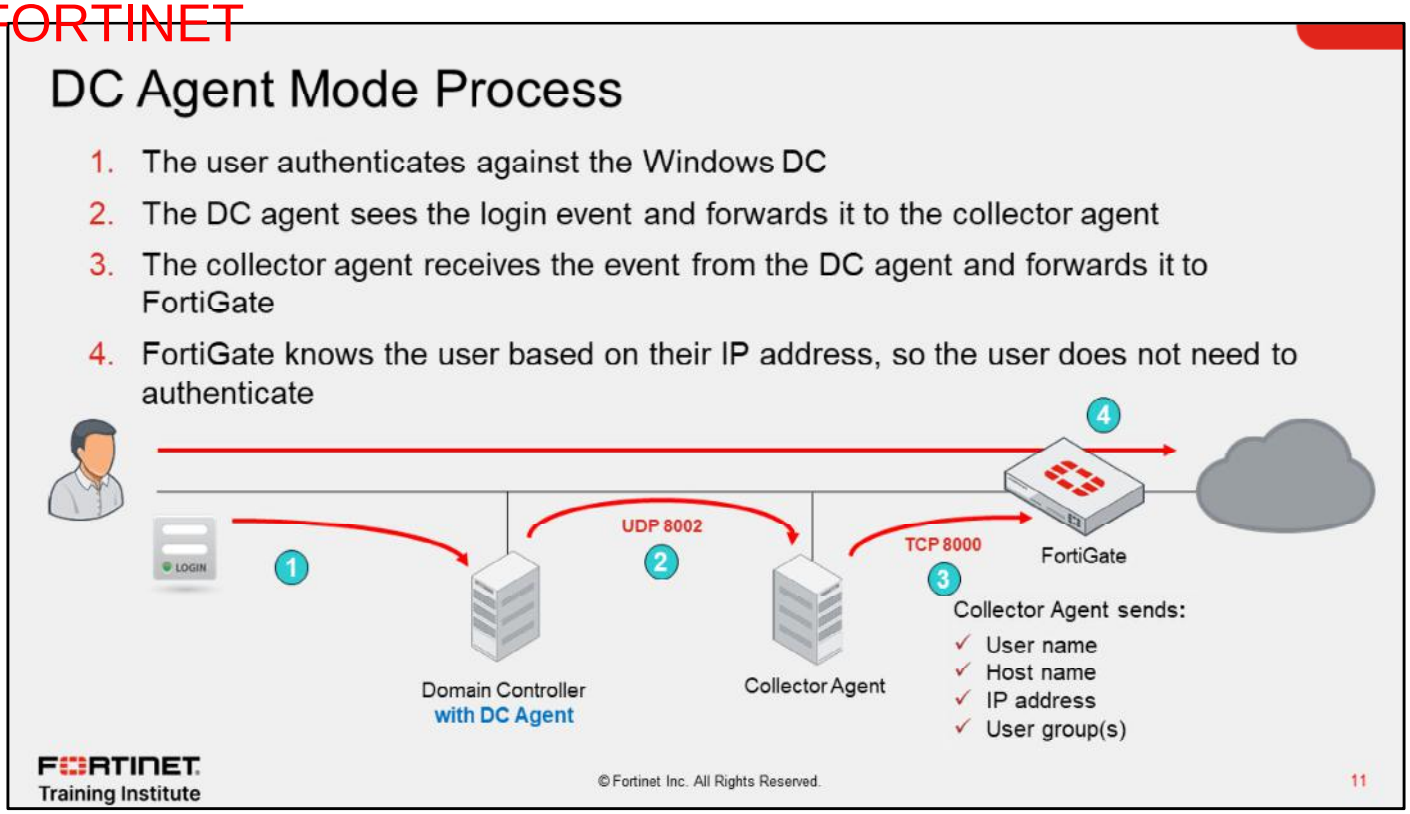

This slide shows the process of information passing between DC agents, the collector agent, and a FortiGate configured for FSSO authentication.

1. When users authenticate with the DC, they provide their credentials.

2. The DC agent sees the login event, and forwards it to the collector agent.

3. The collector agent aggregates all login events and forwards that information to FortiGate. The information sent by the collector agent contains the user name, host name, IP address, and user group(s). The collector agent communicates with FortiGate over TCP port 8000 (default) and it listens on UDP port 8002 (default), for updates from the DC agents. The ports are customizable.

4. FortiGate learns from the collector agent who the user is, their IP address, and some of the AD groups that the user is a member of. When a user tries to access the internet, FortiGate compares the source IP address to its list of active FSSO users. Because the user in this case has already logged in to the domain, and FortiGate already has their information, FortiGate doesn't prompt the user to authenticate again. Rather it allows or denies the traffic based on the matching firewall policy.

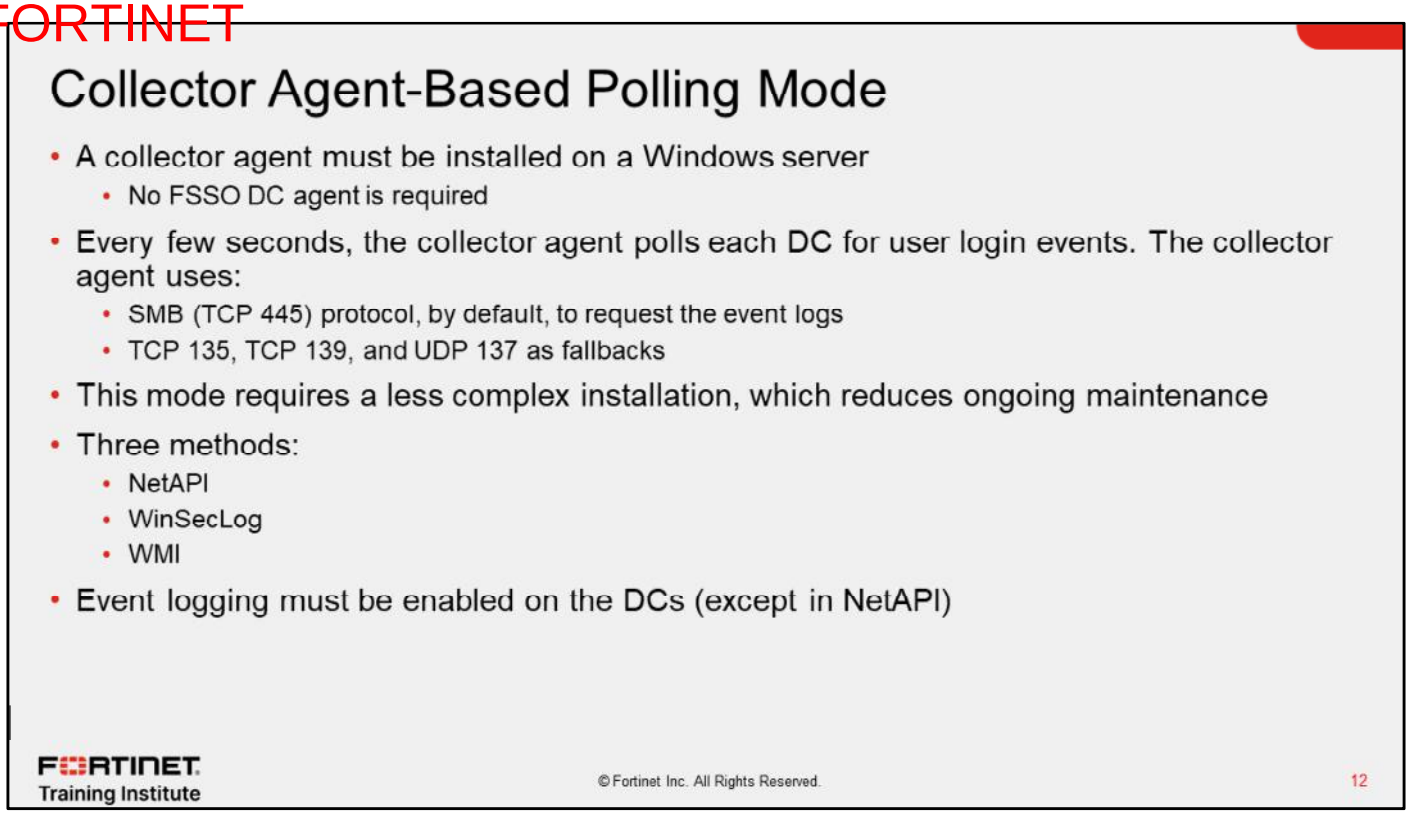

Polling mode can be collector agent-based or agentless.

First, you'll look at the collector agent-based polling mode. Like DC agent mode, collector agent-based mode requires a collector agent to be installed on a Windows server, but it *doesn't* require DC agents to be installed on each DC. In collector agent-based polling mode, the collector agent must be more powerful than the collector agent in DC agent mode, and it also generates unnecessary traffic when there have been no login events.

In Windows Event Log Polling, the most commonly deployed polling mode, the collector agent uses the SMB (TCP port 445) protocol to periodically request event logs from the domain controllers. Other methods may gather information differently, but after the login is received by the collector agent, the collector agent parses the data and builds the user login database, which consists of usernames, workstation names/IP addresses, and user group memberships. This information is then ready to be sent to FortiGate.

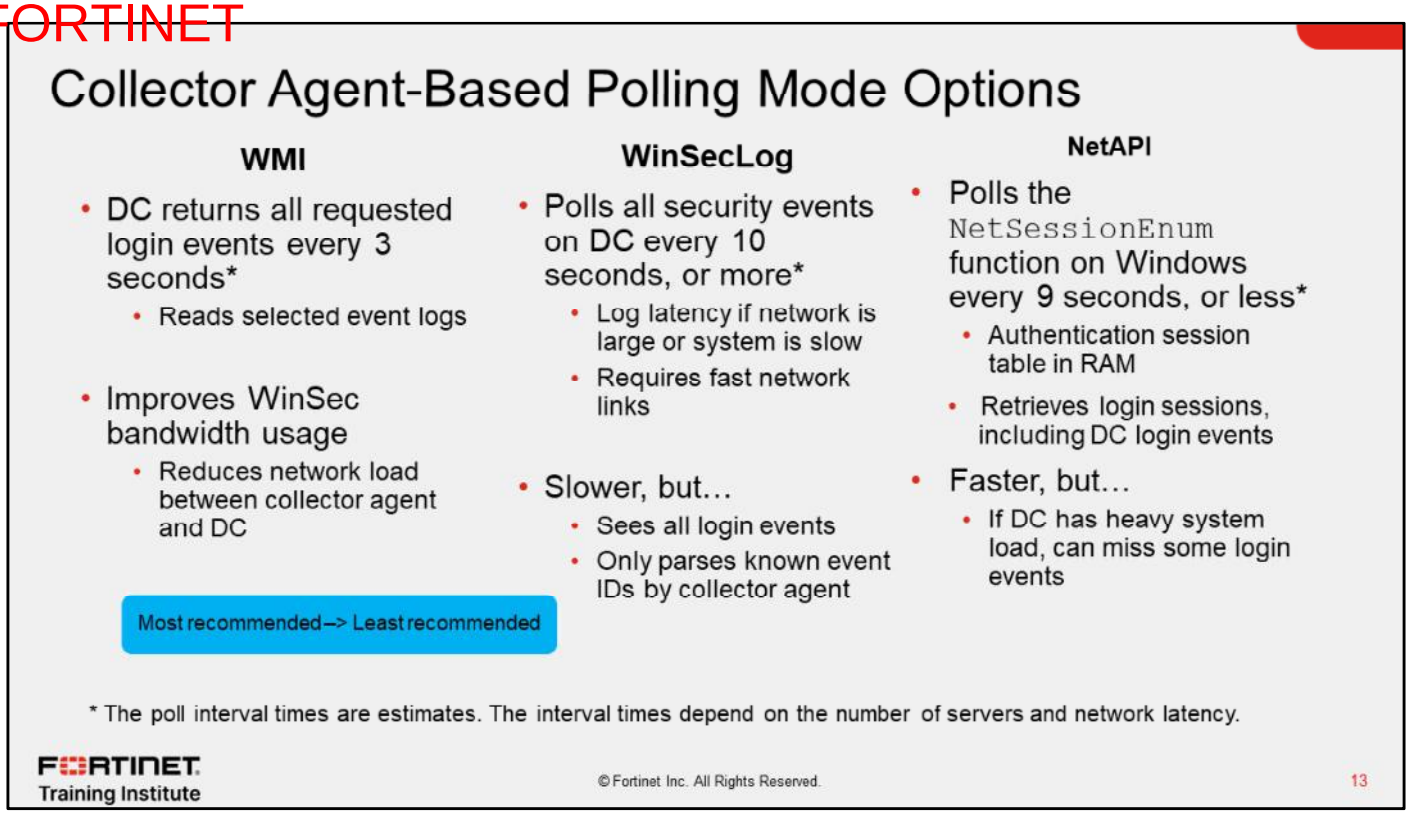

As previously stated, collector agent-based polling mode has three methods (or options) for collecting login information. The order on the slide from left to right shows most recommend to least recommended:

- **WMI**: is a Windows API that gets system information from a Windows server. The DC returns all requested login events. The collector agent is a WMI client and sends WMI queries for user login events to the DC, which, in this case, is a WMI server. The collector agent doesn't need to search security event logs on the DC for user login events; instead, the DC returns all requested login events. This reduces network load between the collector agent and DC.

- **WinSecLog**: polls all the security event logs from the DC. It doesn't miss any login events that have been recorded by the DC because events are not normally deleted from the logs. There can be some delay in FortiGate receiving events if the network is large and, therefore, writing to the logs is slow. It also requires that the audit success of specific event IDs is recorded in the Windows security logs. For a full list of supported event IDs, visit the Fortinet Knowledge Base (http://kb.fortinet.com).

- **NetAPI**: polls temporary sessions created on the DC when a user logs in or logs out and calls the `NetSessionEnum` function on Windows. It's faster than the WinSec and WMI methods; however, it can miss some login events if a DC is under heavy system load. This is because sessions can be quickly created and purged from RAM, before the agent has a chance to poll and notify FortiGate.

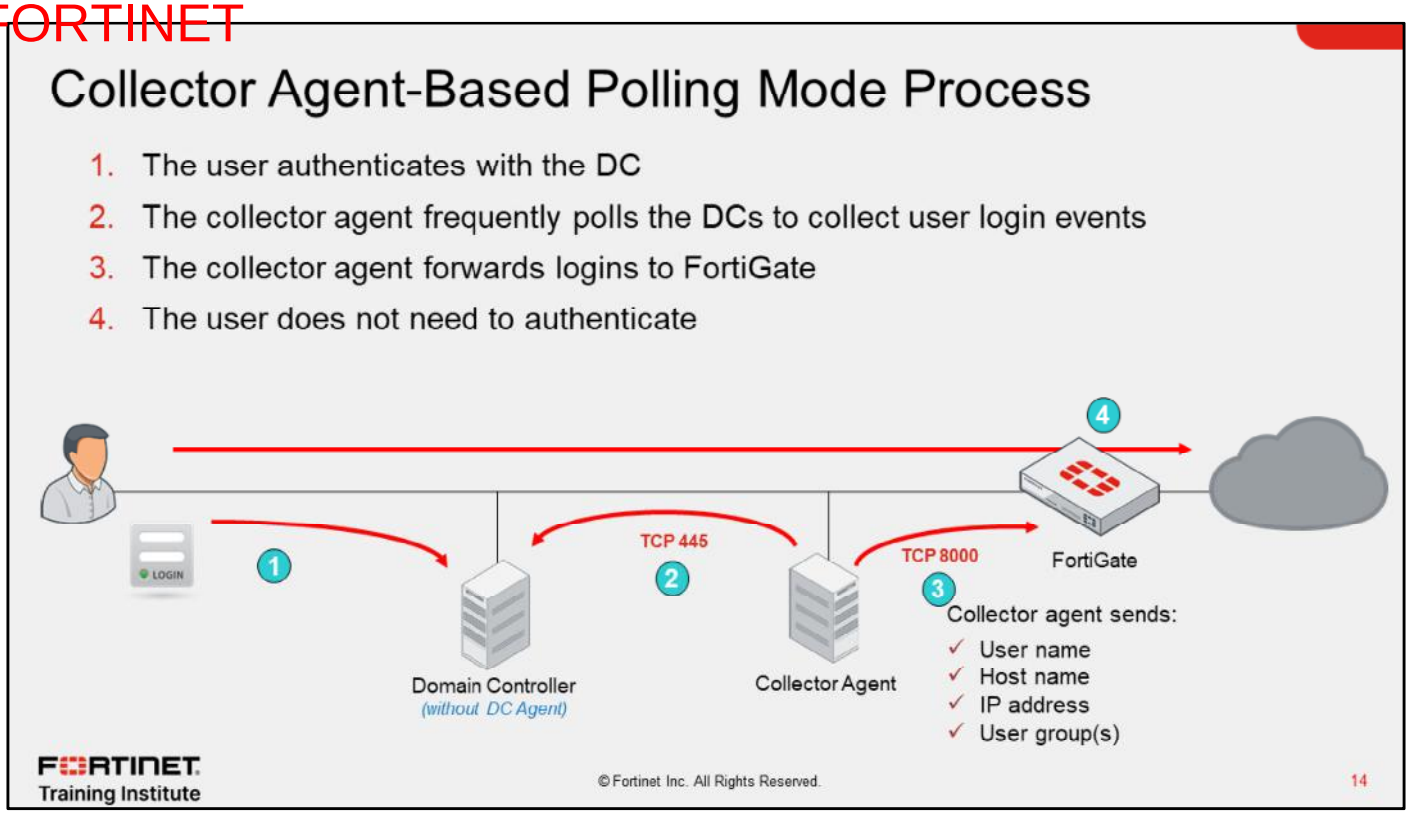

This slide shows an example of FSSO using the collector agent-based polling mode. This example includes a DC, a collector agent, and FortiGate, but the DC doesn't have the dcagent (or, alternatively, dcagent.dll) installed.

1. The user authenticates with the DC, providing their credentials.
2. The collector agent periodically (every few seconds) polls TCP port 445 of each DC directly, to ask if anyone has logged in.
3. The collector agent sends login information to FortiGate over TCP port 8000. This is the same information that is sent in DC agent mode.
4. When user traffic arrives at FortiGate, FortiGate already knows which users are at which IP addresses, and no repeated authentication is required.

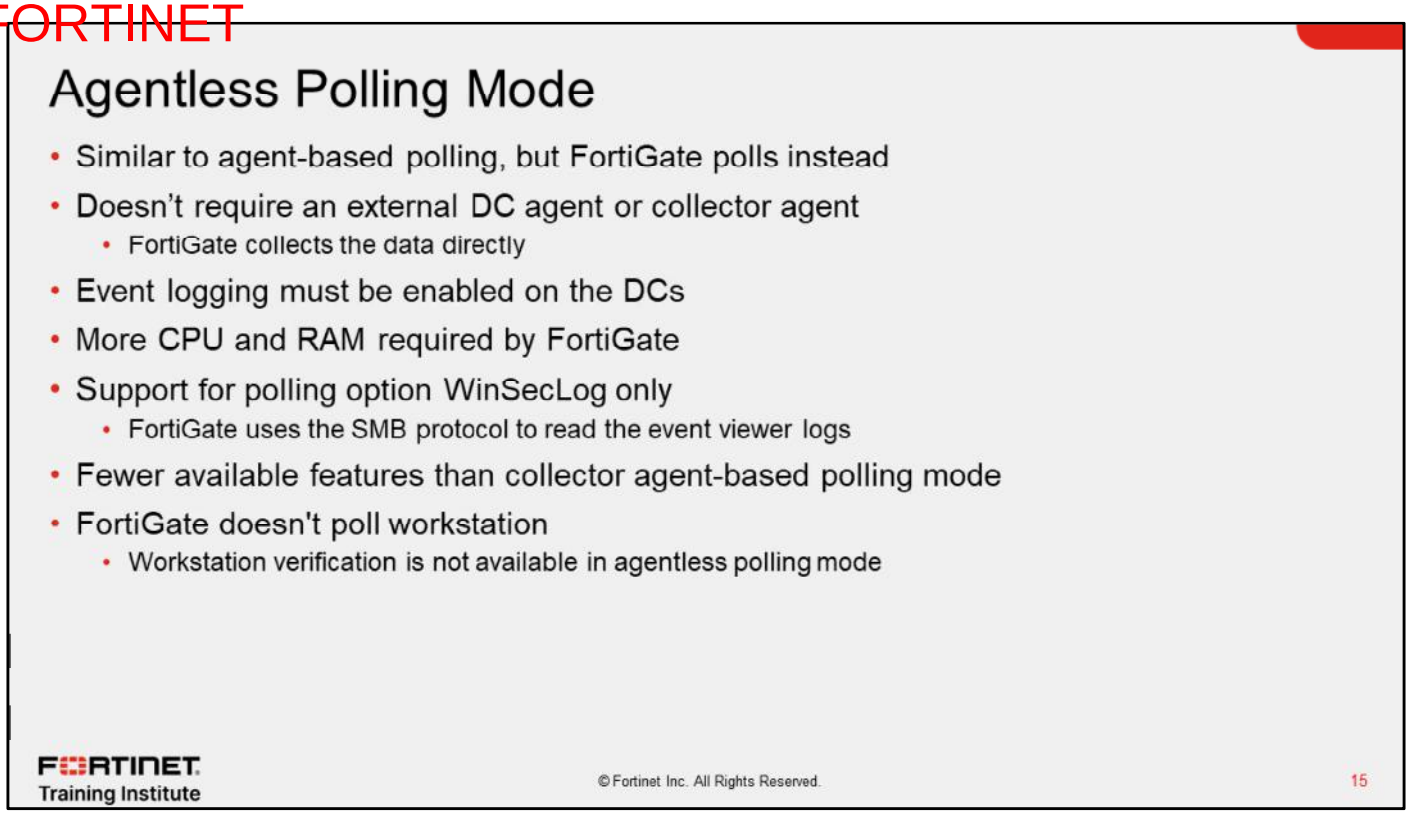

You can deploy FSSO without installing an agent. FortiGate polls the DCs directly, instead of receiving login information indirectly from a collector agent.

Because FortiGate collects all of the data itself, agentless polling mode requires greater system resources, and it doesn't scale as easily.

Agentless polling mode operates in a similar way to WinSecLog, but with only two event IDs: 4768 and 4769. Because there's no collector agent, FortiGate uses the SMB protocol to read the event viewer logs from the DCs.

In agentless polling mode, FortiGate acts as a collector. It is responsible for polling on top of its normal FSSO tasks but does not have all the extra features, such as workstation checks, that are available with the external collector agent.

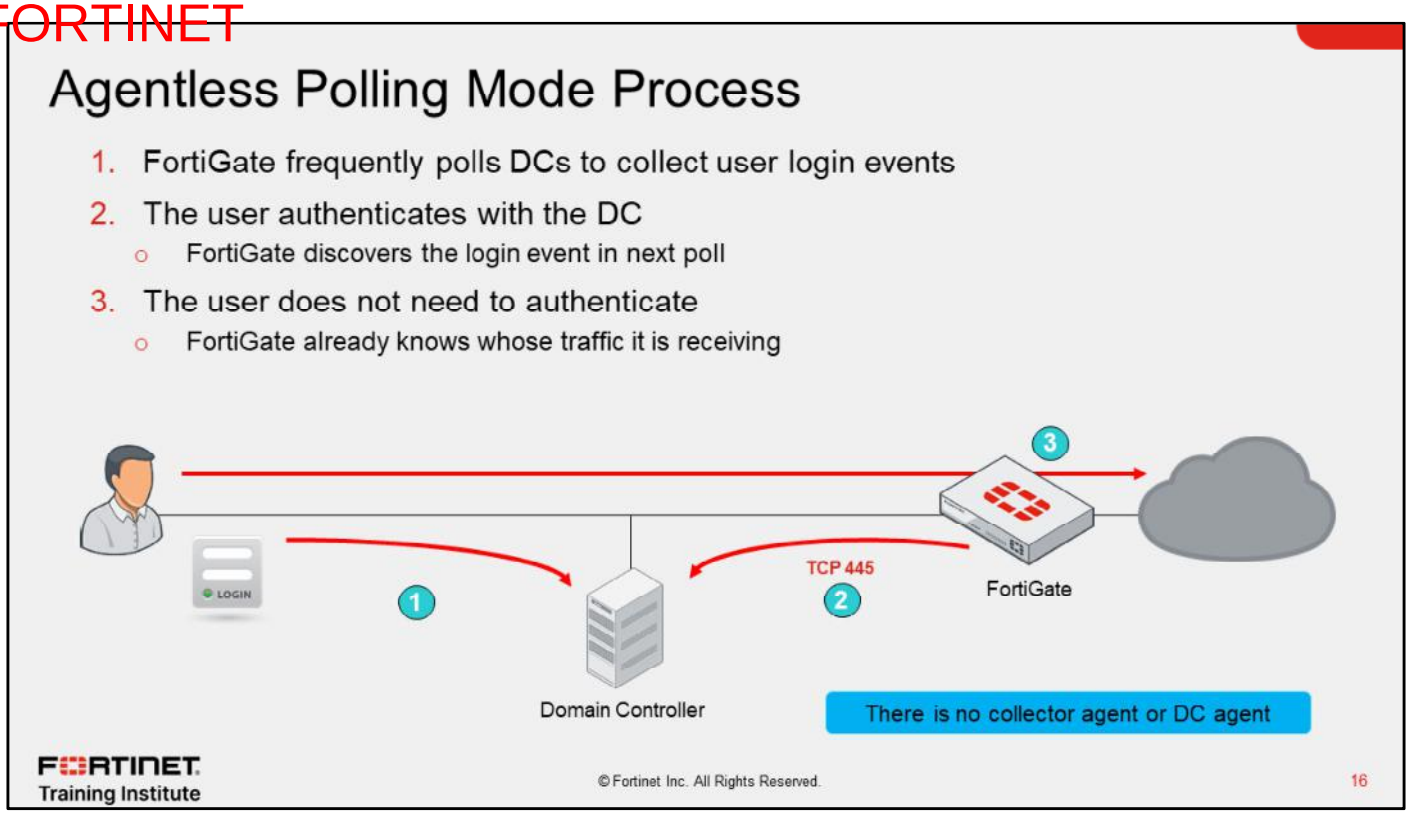

This slide shows how communication is processed without agents. (There is no collector agent or DC agent.)

1. FortiGate polls the DC TCP port 445 to collect user login events.

2. After the user authenticates with the DC, FortiGate registers a login event during its next poll, obtaining the following information: the user name, the host name, and the IP address. FortiGate then queries for the user's user group(s).

3. When the user sends traffic, FortiGate already knows whose traffic it is receiving; therefore, the user does not need to authenticate.

Comparing Modes

	DC agent mode	Polling mode
Installation	Complex—multiple installations (one per DC). Requires reboot.	Easy—one or no installations. No reboot required.
DC agent required	Yes	No
Resources	Shares with DC agents	Has own resources
Scalability	Higher	Lower
Redundancy	Yes	Yes
Level of confidence	Captures all logins	Might miss a login (NetAPI), or have a delay (WinSecLog)

FORTINET Training Institute

17

This table summarizes the main differences between DC agent mode and polling mode.

DC agent mode is more complex. It requires not only a collector agent, but also a DC agent for each monitored domain controller. However, it is also more scalable because the work of capturing logins is done by the DC agents who pass their information directly to the collector.

In polling mode, the collector needs to query every domain controller, every few seconds. So, with each DC that is added, the number of queries grows. If you want to add a second collector agent for redundancy in polling mode, both collector agents need to query every DC individually.

In DC agent mode, the DC agent just has to collect the log once, and send a copy of the necessary information to all the collector agents. In comparison, if you use polling mode, some login events might be missed or delayed, depending on the polling option used.

You do not have to install a collector agent on the DC, you can install it on any Windows machine on the network.

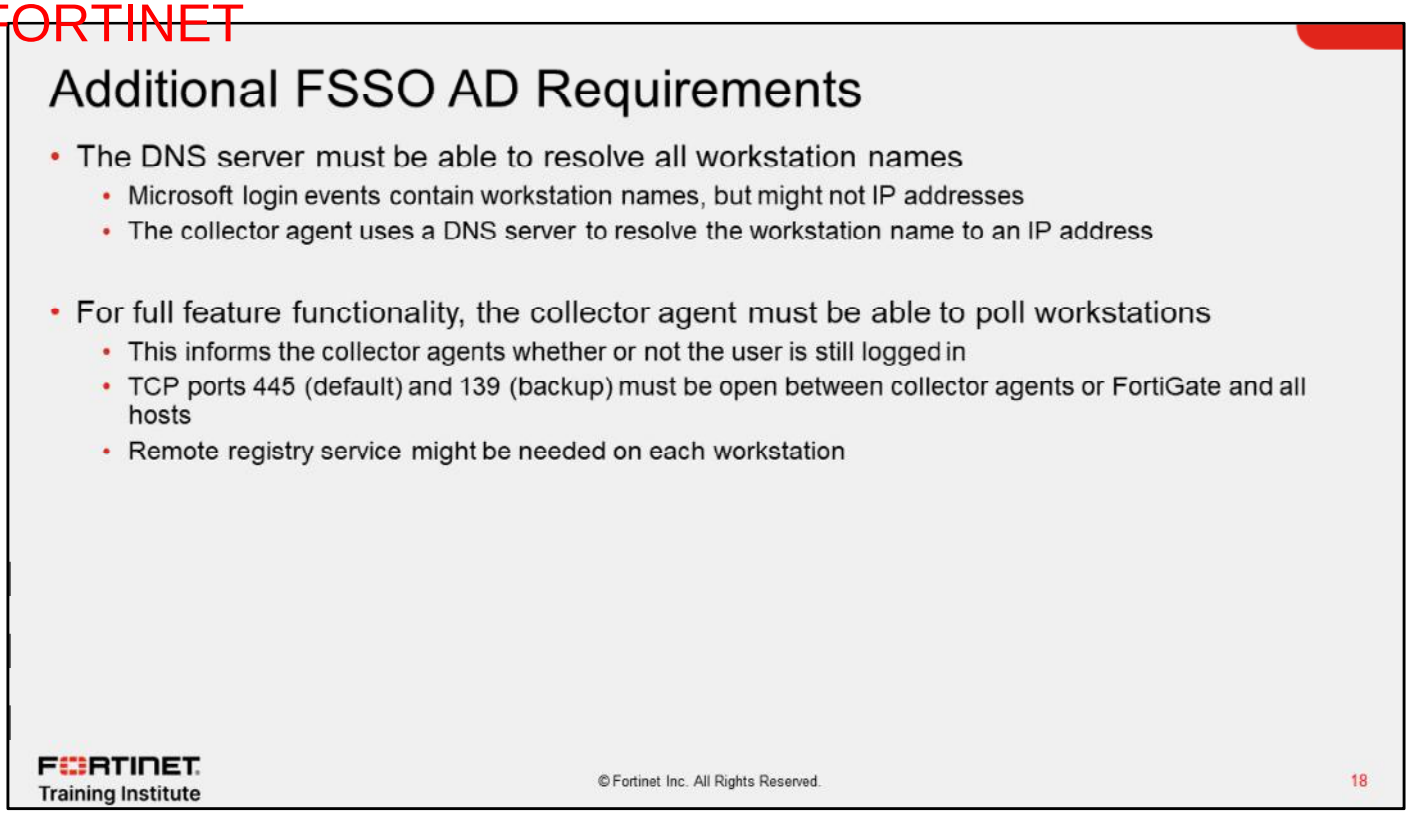

Regardless of the collector method you choose, some FSSO requirements for your AD network are the same:

- Microsoft Windows login events have the workstation name and username, but not the workstation IP address. When the collector agent receives a login event, it queries a DNS server to resolve the IP address of the workstation. So, FSSO requires that you have your own DNS server. If a workstation IP address changes, DNS records must be updated immediately in order for the collector agent to be aware of the change and report them to FortiGate.

- For full feature functionality, collector agents need connectivity with all workstations. Since a monitored event log is not generated on logout, the collector agent (depending on the FSSO mode) must use a different method to verify whether users are still logged in. So, each user workstation is polled to see if users are still there.

- The DC agent, when the user logs in, intercepts the login event on the domain controller. It then resolves the DNS of the client, and sends it to the collector agent.

  The collector agent receives the DNS and then performs a DNS resolution in order to check if the IP of the user has changed.

Knowledge Check

1. Which is the recommended mode for FSSO deployments?
   - ✓ A. DC agent mode
   - B. Polling mode: Agentless

2. Which FSSO mode requires more FortiGate system resources (CPU and RAM)?
   - A. Polling mode: Collector agent-based
   - ✓ B. Polling mode: Agentless

FORTINET Training Institute

19

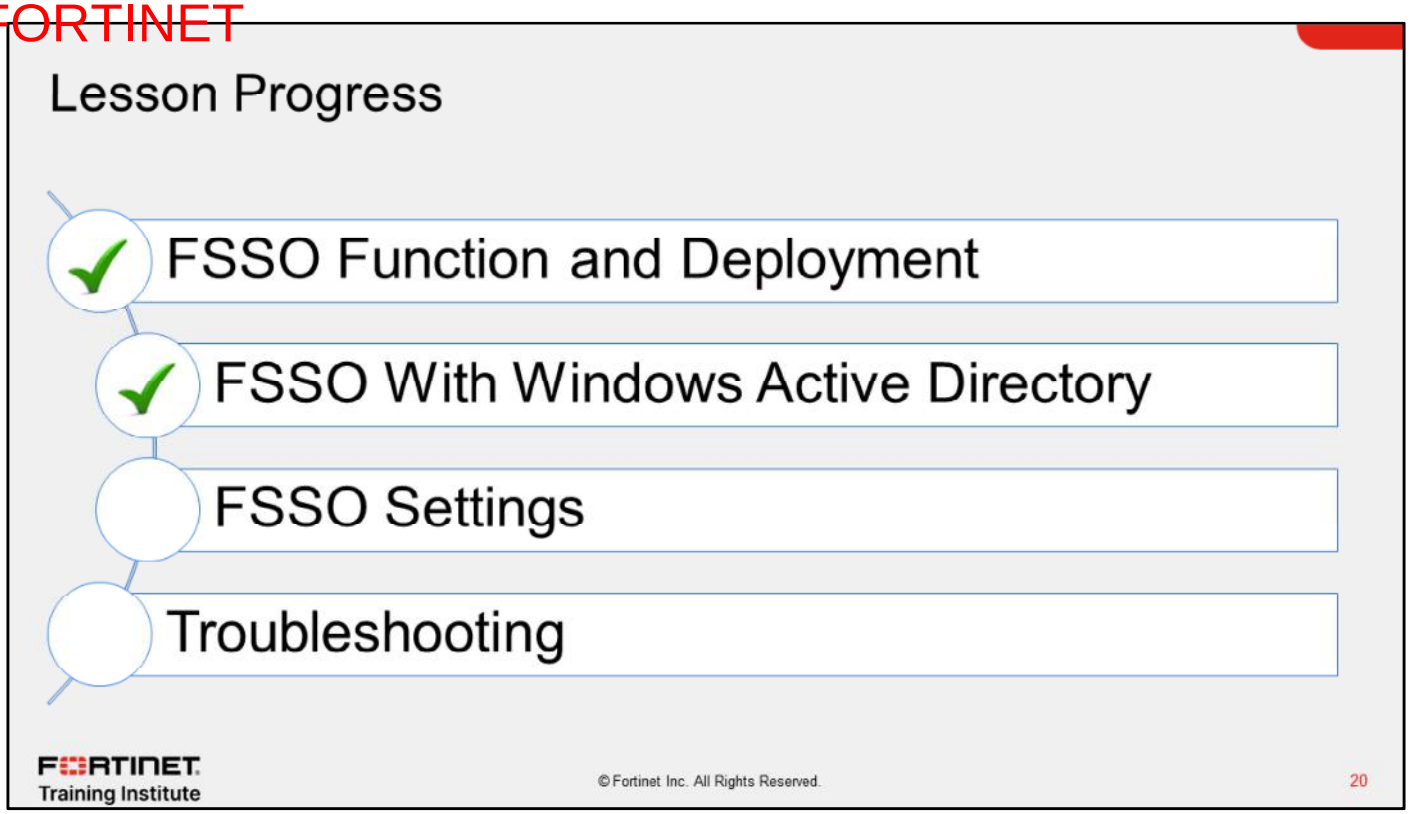

Good job! You now understand how FortiGate detects login events in Windows Active Directory (AD) using FSSO.

Now, you'll lean how to configure FSSO settings.

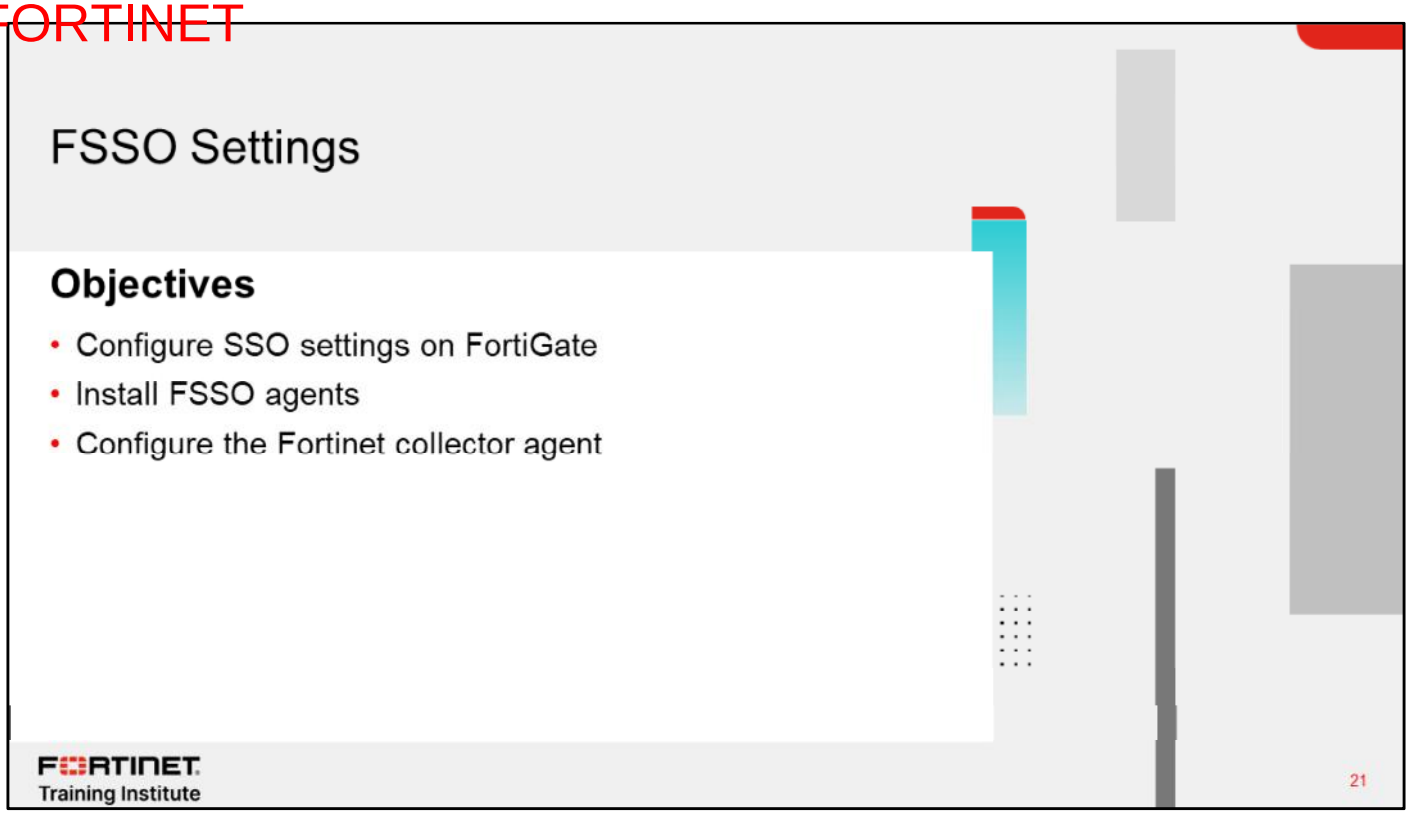

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in configuring the FSSO settings on FortiGate, and installing and configuring the FSSO agents, you will be able to implement FSSO within your network.

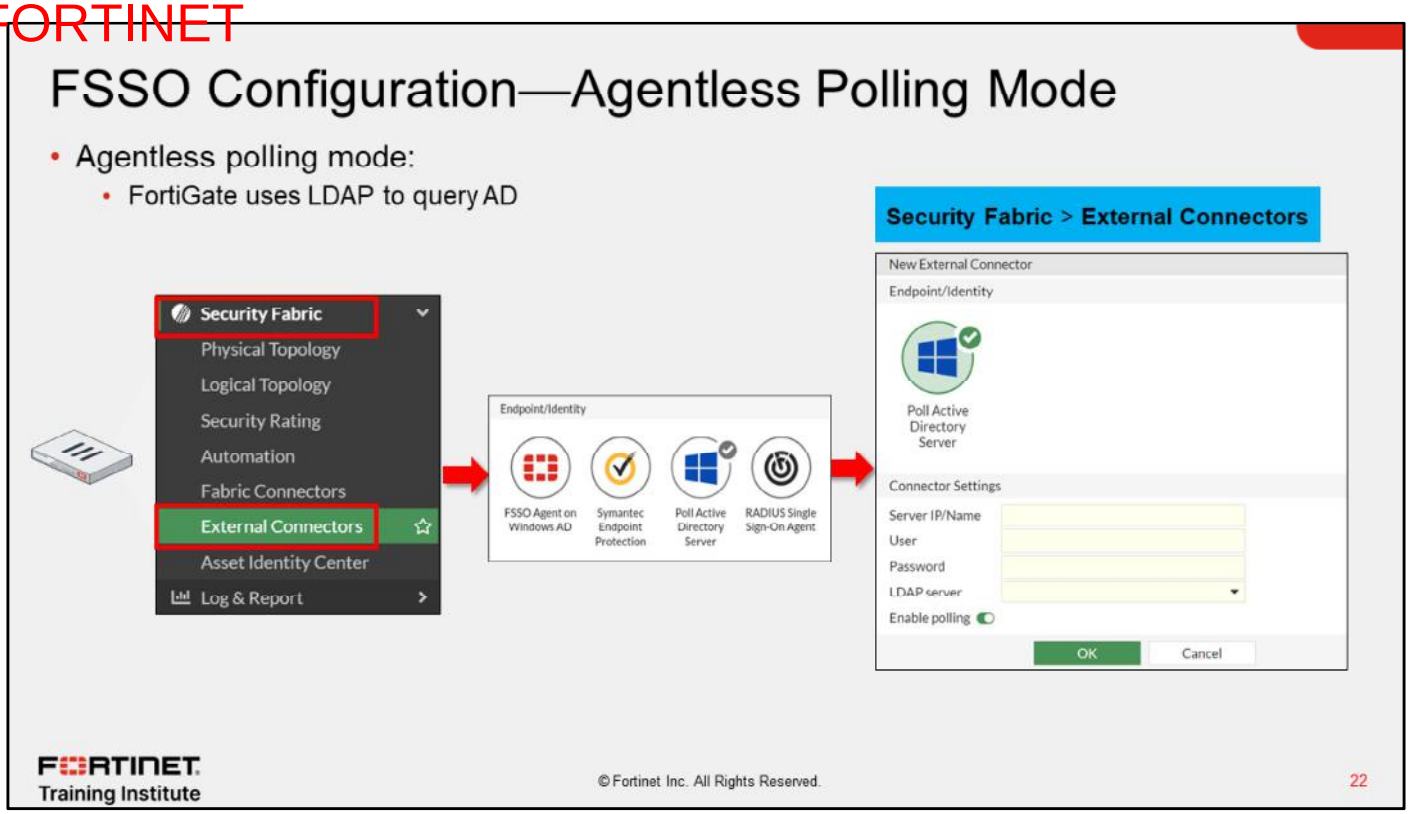

FortiGate FSSO configuration is straightforward.

If FortiGate is acting as a collector for agentless polling mode, you must select **Poll Active Directory Server** and configure the IP addresses and AD administrator credentials for each DC.

FortiGate uses LDAP to query AD to retrieve user group information. For this to happen, you must add the LDAP server to the **Poll Active Directory Server** configuration.

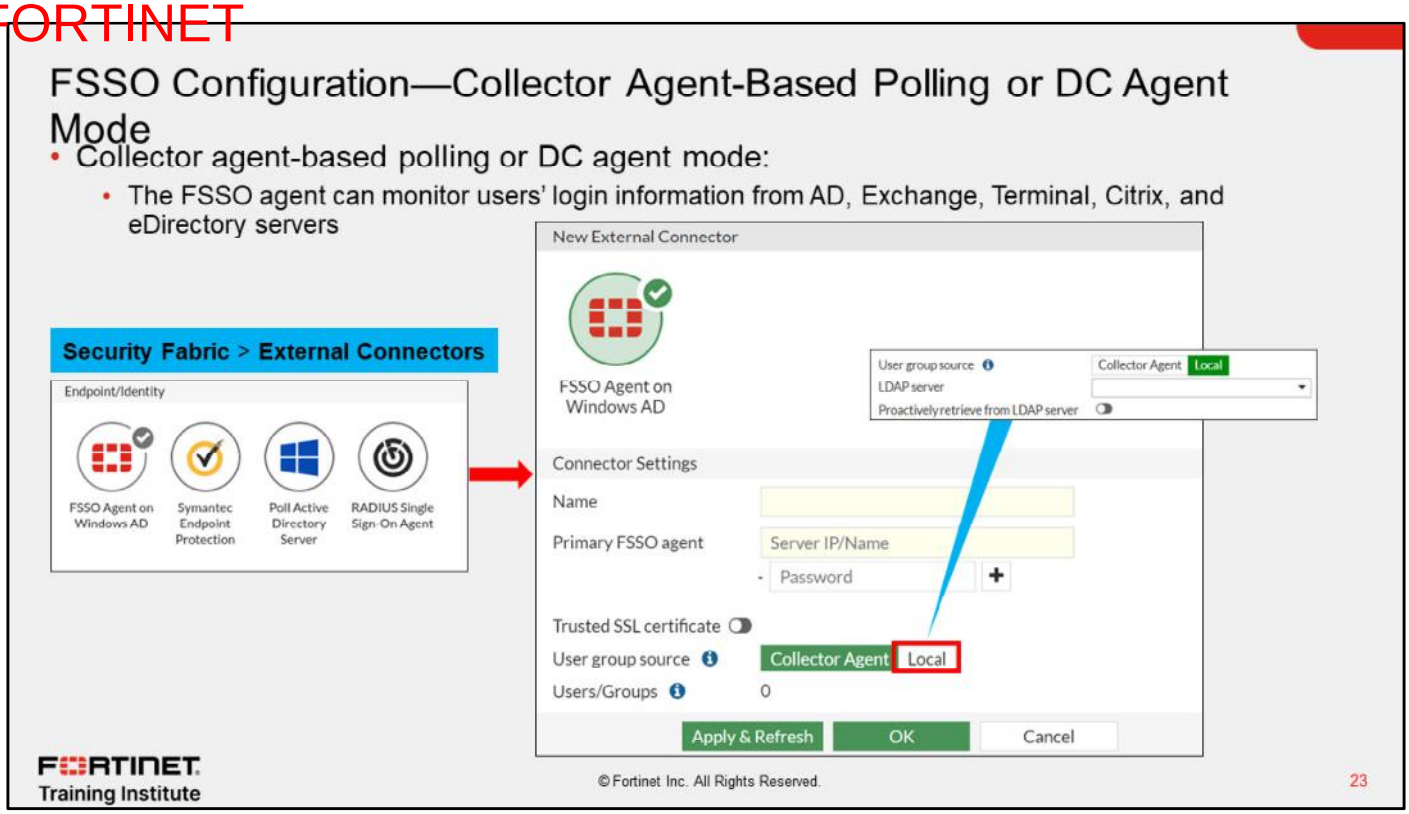

If you have collector agents, using either the DC agent mode or the collector agent-based polling mode, you must select **Fortinet Single-Sign-On Agent** and configure the IP address and password for each collector agent.

The FSSO collector agent can access Windows AD in one of two modes:

- **Collector Agent**: You create group filters are created on the collector agent. You can set FortiGate to **Collector Agent** mode, and the collector agent can still use **Advanced** mode to access nested groups.
- **Local**: You create group filters on FortiGate, using the LDAP server. If you set FortiGate to **Local** mode, you must set the collector agent to **Advanced** mode, otherwise the collector agent does not recognize the group filter sent by FortiGate and does not pass down any user logins.

FSSO Agent Installation

1. Visit the Fortinet support website:
   - https://support.fortinet.com
2. Click **Download** > **Firmware Images**

Available agents:

- DC agent: DCAgent_Setup
- CA for Microsoft servers: FSSO_Setup
- CA for Novell: FSSO_Setup_edirectory
- TS Agent: TSAgent_Setup

3. Select **FortiGate**, then click **Download**.
4. Click **v7.00** > **7.2** > **7.2.0** > **FSSO**

Example image below:

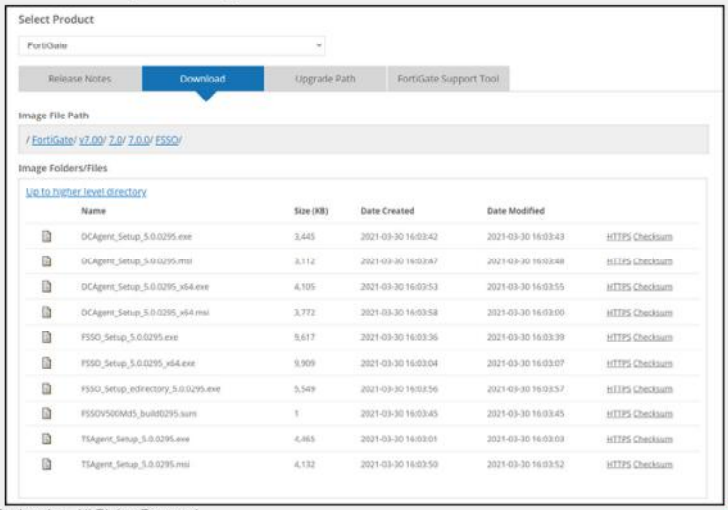

FORTINET Training Institute

24

The FSSO agents are available on the Fortinet Support website. There you will find the following:

- The DC agent
- The collector agent for Microsoft servers: FSSO_Setup
- The collector agent for Novell directories: FSSO_Setup_edirectory
- The terminal server agent (TSagent) installer for Citrix and terminal servers: TSAgent_Setup

Also, for each agent, there are two versions: the executable (`.exe`) and Microsoft Installer (`.msi`).

Notice that you do not need to match the FSSO version with your exact FortiGate firmware version. When installing FSSO, grab the latest collector agent for your major release. You do however, need to match the DC agent version to the collector agent version.

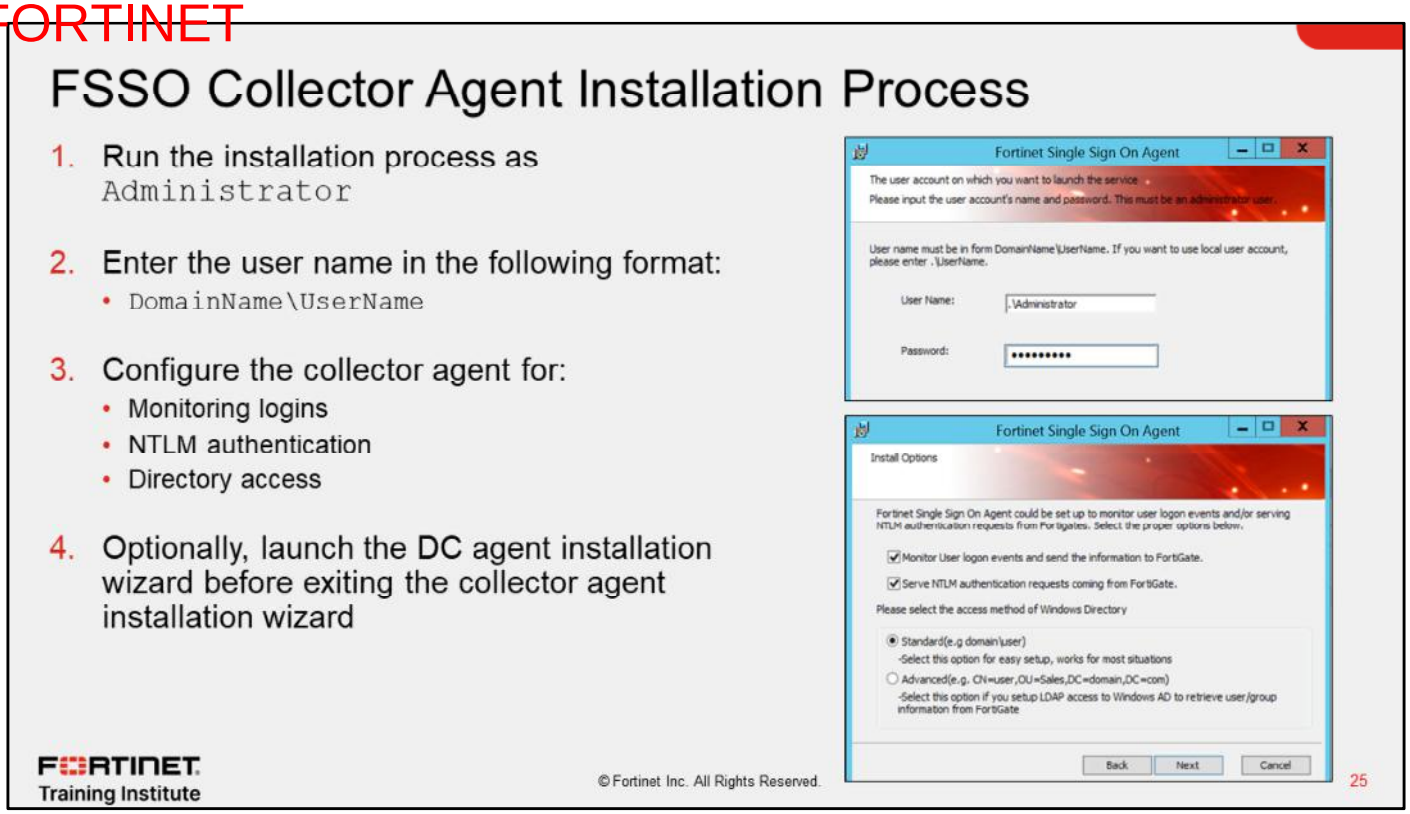

After you've downloaded the collector agent, run the installation process as Administrator and follow these steps in the installation wizard:

1. Read and accept the license agreement.
2. Optionally, change the installation location. The default folder is named `FSAE` (Fortinet Server Authentication Extension).
3. Enter the username. By default, the agent uses the name of the currently running account; however, you can change it using the format: **DomainName\UserName**.
4. Alternatively, configure your collector agent for monitoring, NTLM authentication, and directory access. These options are also customizable after installation. Although the default is **Standard** mode, when doing new FSSO setups it is always a best practice to install in **Advanced** mode. You will look at some of the advantages in this lesson.
5. If you want to use DC agent mode, make sure that **Launch DC Agent Install Wizard** is selected. This automatically starts the DC agent installation.

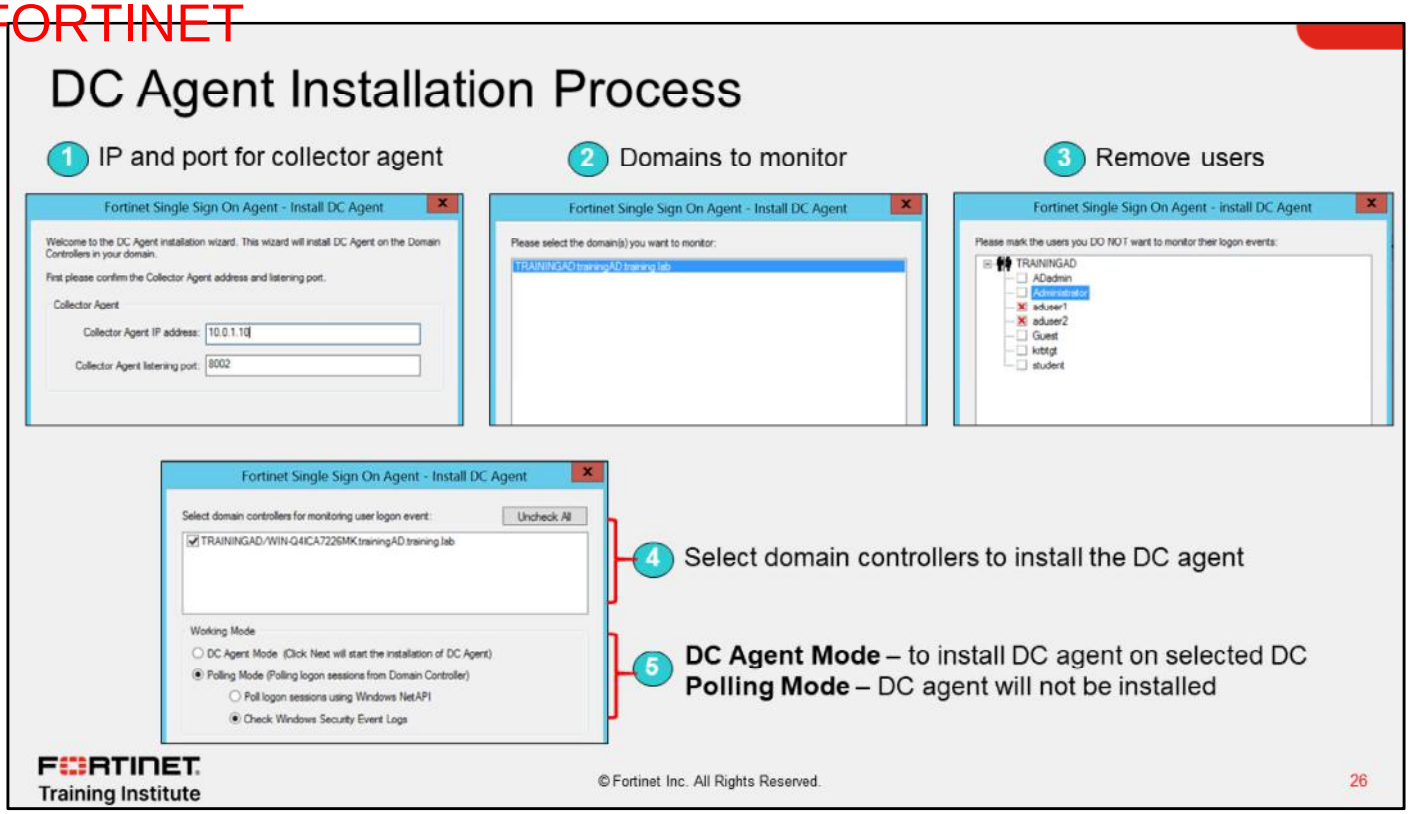

If you have just installed the collector agent and you selected **Launch DC Agent Install Wizard,** the installation process for domain controller agent automatically starts.

1. Enter the IP address for the collector agent. Optionally, you can customize the listening port, if the default value is already used by another service.
2. Select the domains to monitor. If any of your required domains are not listed, cancel the wizard and set up the correct trusted relationship with the domain controller. Then, run the wizard again. Note that this could also be a result of using an account without all the necessary permissions.
3. Optionally, select users that you do not want to monitor; these users' login events are not recorded by the collector and therefore are not passed to FortiGate. While these users are still able to generate login events to the domain, when they are detected by the collector agent, they are discarded so as to not interfere with the logged in user. This is especially useful in environments with a centrally managed antivirus solution, or a scheduled backup service that uses an AD account to start. These accounts can create login events for the collector agent that overwrite existing user logins. This may result in FortiGate applying the incorrect policies and profiles based on the overriding account. You can also customize the option to ignore users after installation is complete.
4. Optionally, clear the checkboxes of domain controllers that you don't want to install the DC agent on. Remember, for DC agent mode FSSO, at least one domain controller must have the DC agent installed. Also remember that installing the DC agent requires a reboot of the DC before it will start gathering login events. You can add or remove the DC agent to DCs at any time after the installation is complete.
5. Select **DC Agent Mode** as the working mode. If you select **Polling Mode**, the DC agent will not be installed.

Finally, the wizard requests a system reboot.

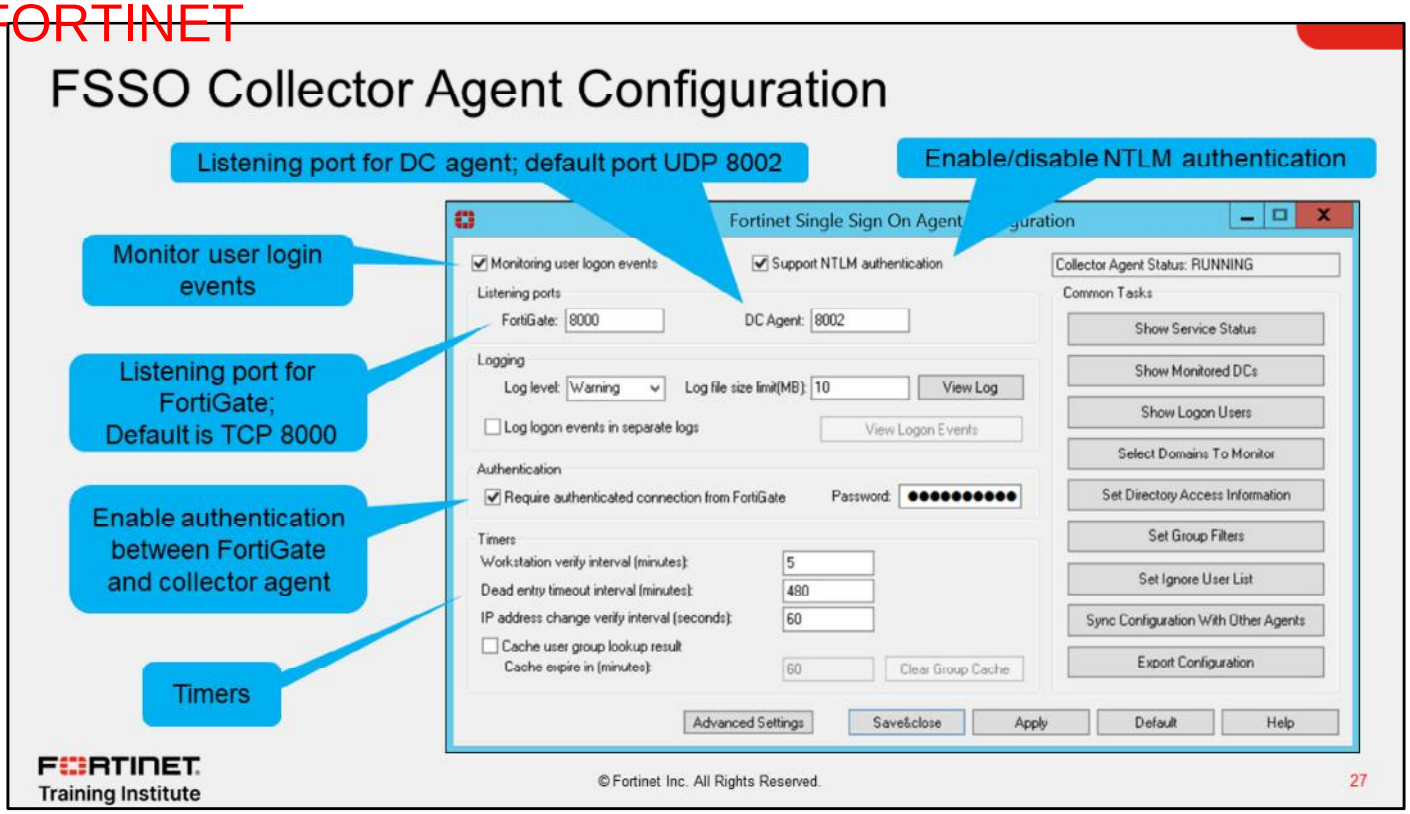

On the FSSO agent configuration GUI, you can configure settings such as:

- The listening port for the communication with the DC agents (UDP)
- The listening port for the communication with FortiGate (TCP)
- NTLM authentication support
- Password authentication between the collector agent and FortiGate
- Timers

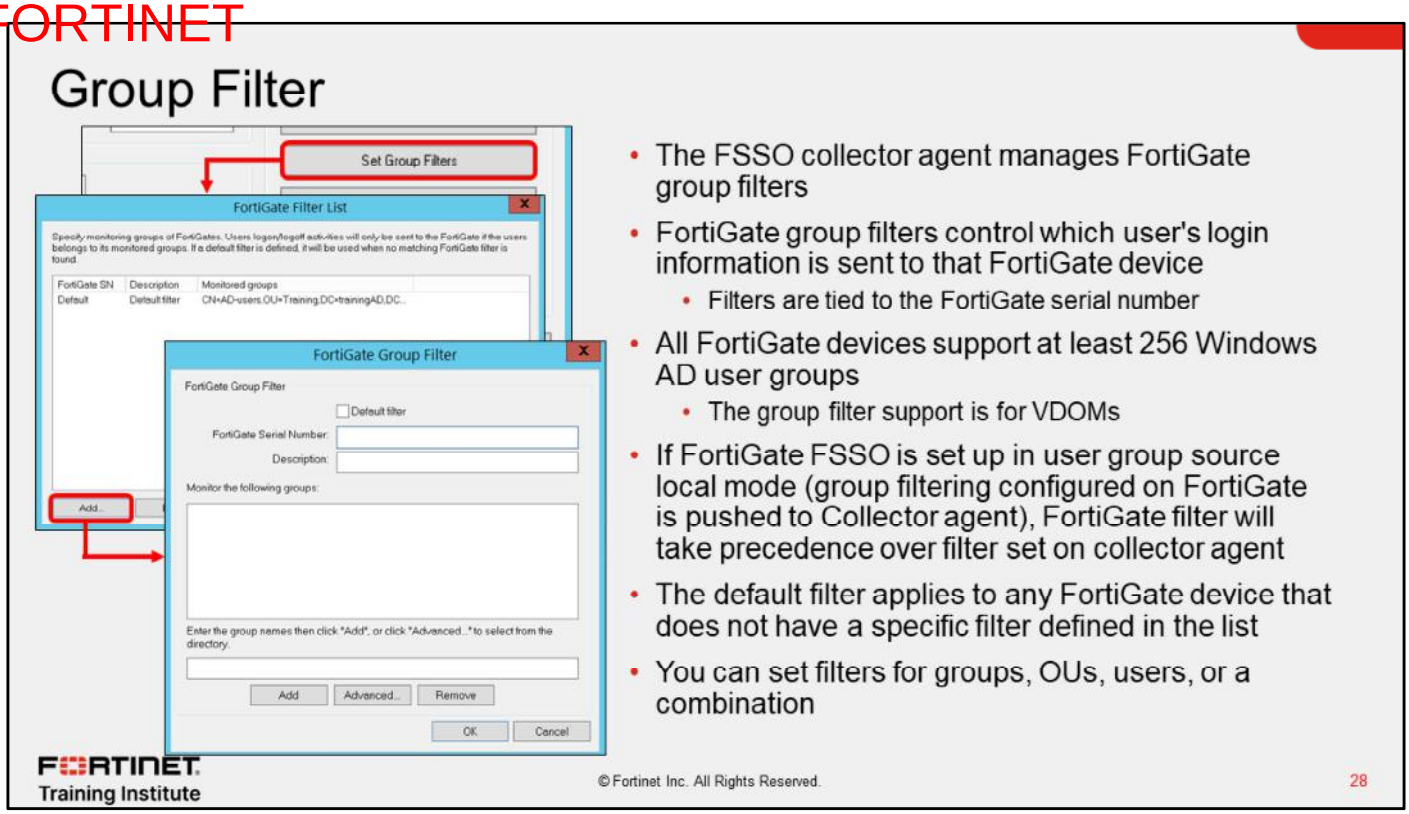

The FSSO collector agent allows you to configure a FortiGate group filter, which actively controls what user login information is sent to each FortiGate device. So, you can define which groups the collector agent passes to individual FortiGate devices.

Monitoring the entire group list in a large AD structure is highly inefficient, and a waste of resources. Most FSSO deployments need group segmentation (at least four or five groups), with the intention of assigning varying levels of security profile configurations to the different groups, using identity-based policies.

Group filters also help to limit the traffic sent to FortiGate. The maximum number of Windows AD user groups allowed on FortiGate depends on the model. Low-end FortiGate models support 256 Windows AD user groups. Mid-range and high-end models can support more groups. This is per VDOM, if VDOMs are enabled on FortiGate.

You can filter on FortiGate instead of the collector agent, but only if the collector agent is operating in advanced mode. In this case, the collector agent uses the list of groups you selected on FortiGate as its group filter for that device.

The filter list is initially empty. At a minimum, you should create a default filter that applies to all FortiGate devices without a defined filter.

Note that if you change the AD access mode from **Standard** to **Advanced** or **Advanced** to **Standard**, you must recreate the filters because they vary depending on the mode.

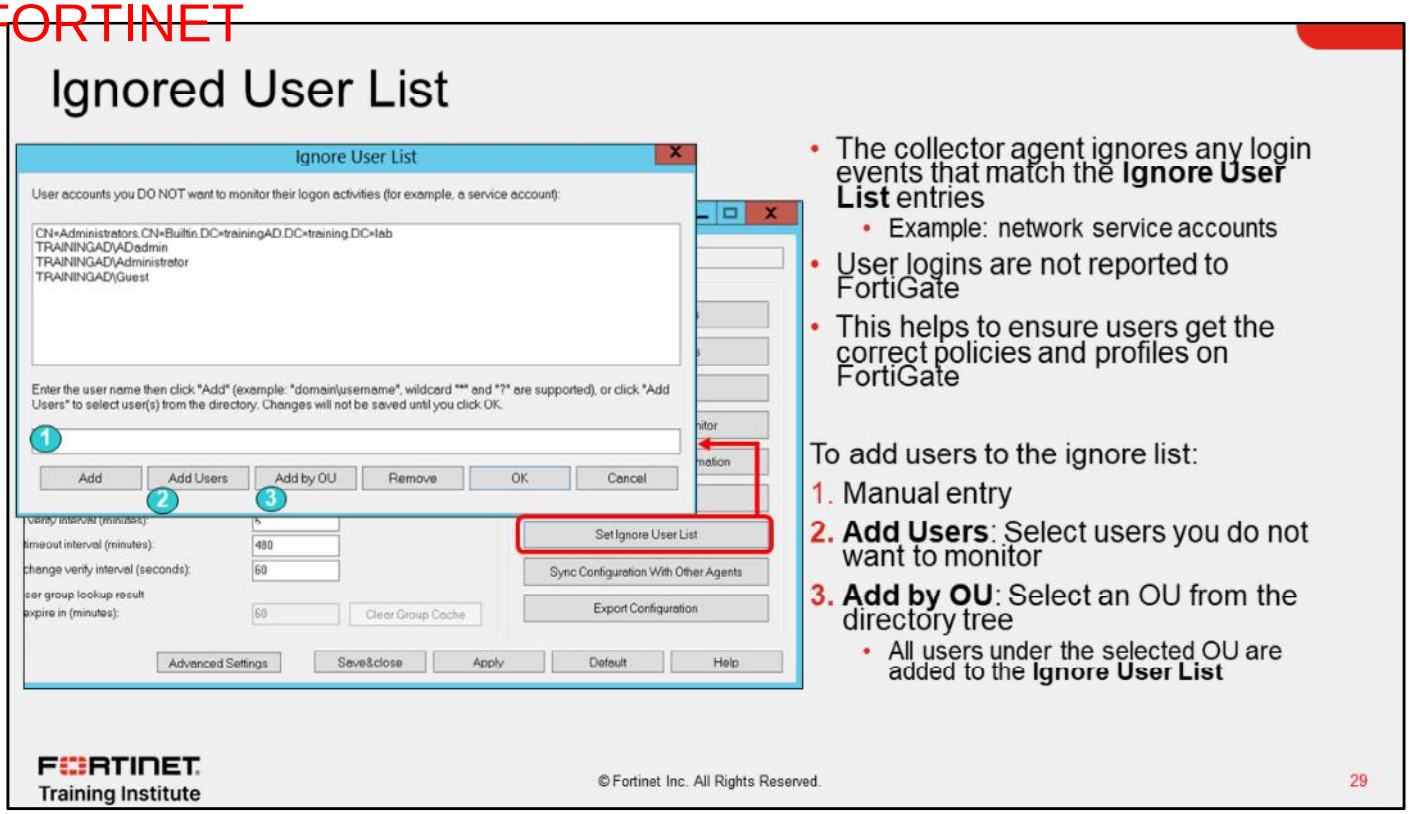

The FSSO collector agent ignores any login events that match the **Ignore User List** entries. Therefore, these login events are not recorded by the collector agent, nor are they reported to FortiGate.

It is a good practice to add all network service accounts to the **Ignore User List**. Service accounts tend to overwrite user login events, and create issues with identity-based policy matching.

You can add users to the **Ignore Users List** in the following ways:
- Manually enter the username.
- Click **Add Users**, and then choose the users you do not want to monitor.
- Click **Add by OU**, and then select an OU from the directory tree.

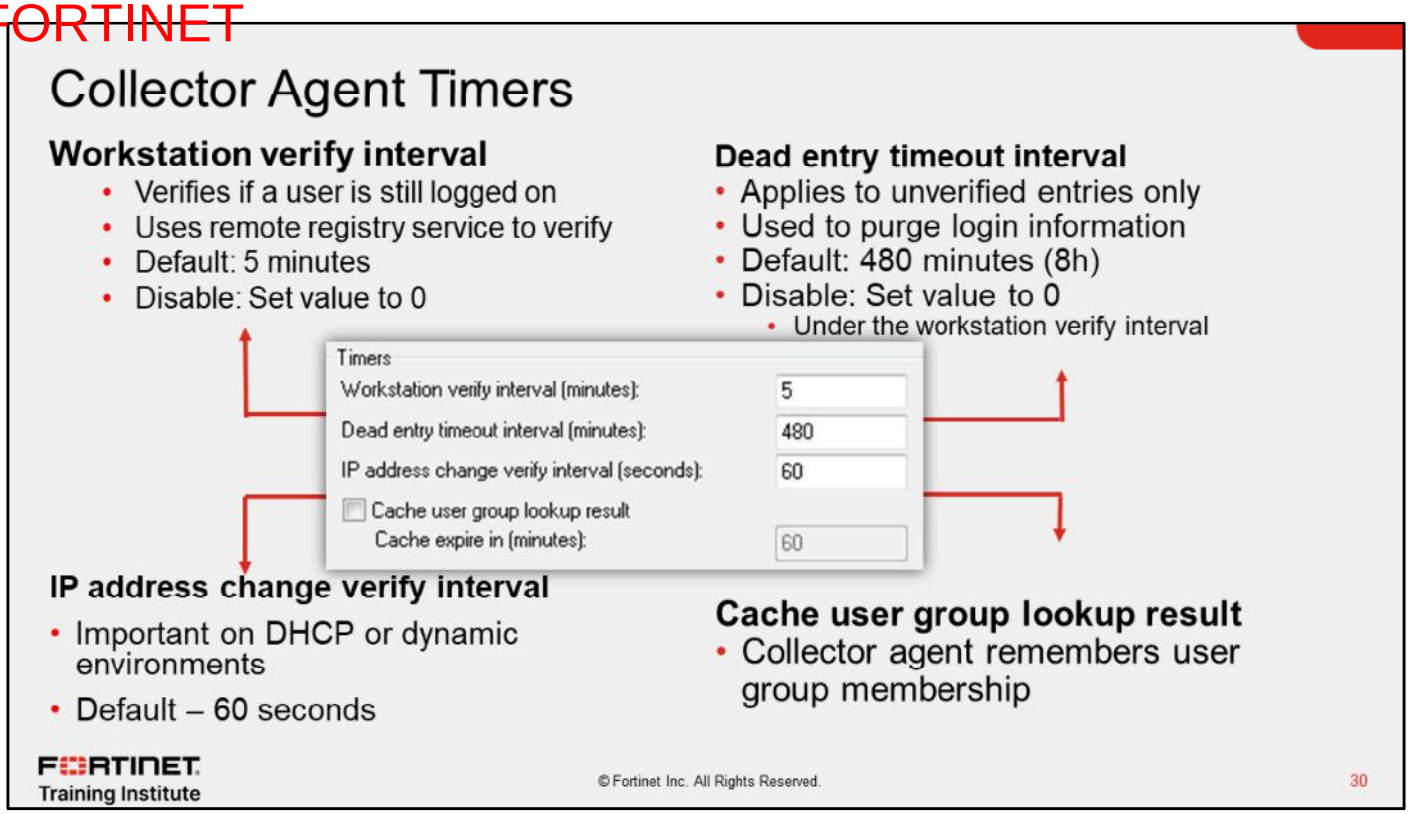

The FSSO collector agent timers play an important role in ensuring the correct operation of FSSO.

Now, you'll take a look at each one and how they work.

- **Workstation verify interval**. This setting controls when the collector agent connects to individual workstations on port 139 (or port 445), and uses the remote registry service to verify if a user is still logged in to the same station. It changes the status of the user under **Show login User**, to **not verified** when it cannot connect to the workstation. If it does connect, it verifies the user and the status remains **OK**. To facilitate this verification process, you should set the remote registry service to auto start on all domain member PCs.

- **Dead entry timeout interval**. This setting applies only to entries with an unverified status. When an entry is not verified, the collector starts this timer. It's used to age out the entry. When the timer expires, the login is removed from the collector. From the perspective of FortiGate, there is no difference between entries that are **OK** and entries that are **not verified**. Both are considered valid.

- **IP address change verify interval**. This setting checks the IP addresses of logged in users and updates FortiGate when a user's IP address changes. This timer is especially important in DHCP or dynamic environments to prevent users from being locked out if they change IP address. The domain DNS server should be accurate; if the DNS server does not update the affected records promptly, the collector agent's IP information is inaccurate.

- **Cache user group lookup result**. This setting caches the user group membership for a defined period of time. It is not updated, even if the user changes group membership in AD.

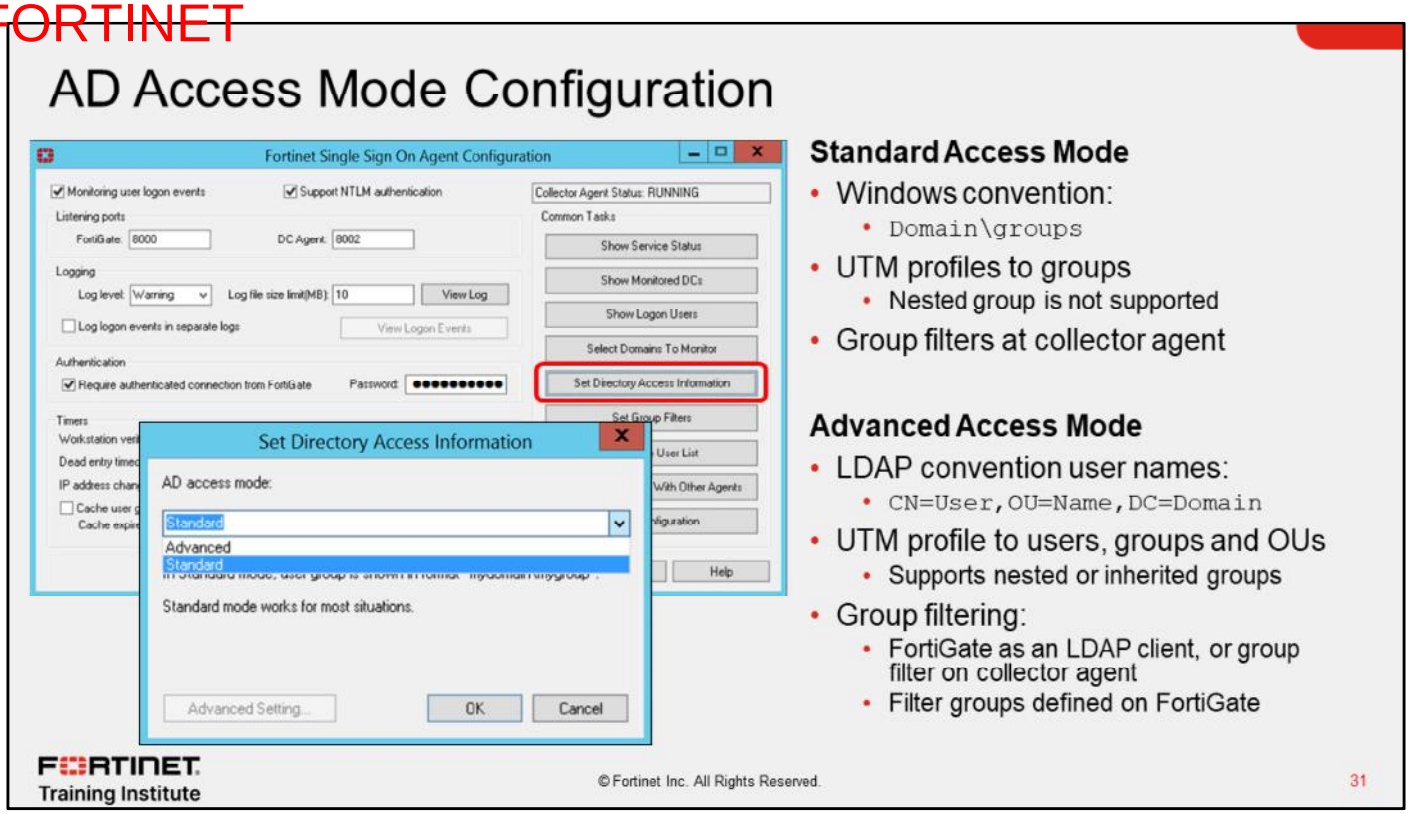

Another important FSSO setting is the AD access mode. You can set the AD access mode by clicking **Set Directory Access Information**. The AD access mode specifies how the collector agent accesses and collects the user and user group information. There are two modes that you can use to access AD user information: **Standard** and **Advanced**.

The main difference between modes is the naming convention used:
- **Standard** mode uses the Windows convention, NetBios: `Domain\groups`, while
- **Advanced** mode uses the LDAP convention: `CN=User, OU=Name, DC=Domain`.

Also, advanced mode supports nested or inherited groups; that is, users can be members of subgroups that belong to monitored *parent* groups. Additionally, in advanced mode, FortiGate can apply security profiles to individual users, user groups, and OUs.

In comparison, in standard mode, you can apply security profiles only to user groups, not individual users.

In advanced mode, you can configure FortiGate as an LDAP client and configure the group filters on FortiGate. You can also configure group filters on the collector agent.

If the LDAP on the collector agent fails, it doesn't matter what the LDAP on the FortiGate says, FSSO won't work. If the FortiGate LDAP fails, but the LDAP on the collector agent is still running, the FortiGate may not be able to collect logs, but the collector agent still collects logs.

Fortinet strongly encourages users to create filters from the collector agent.

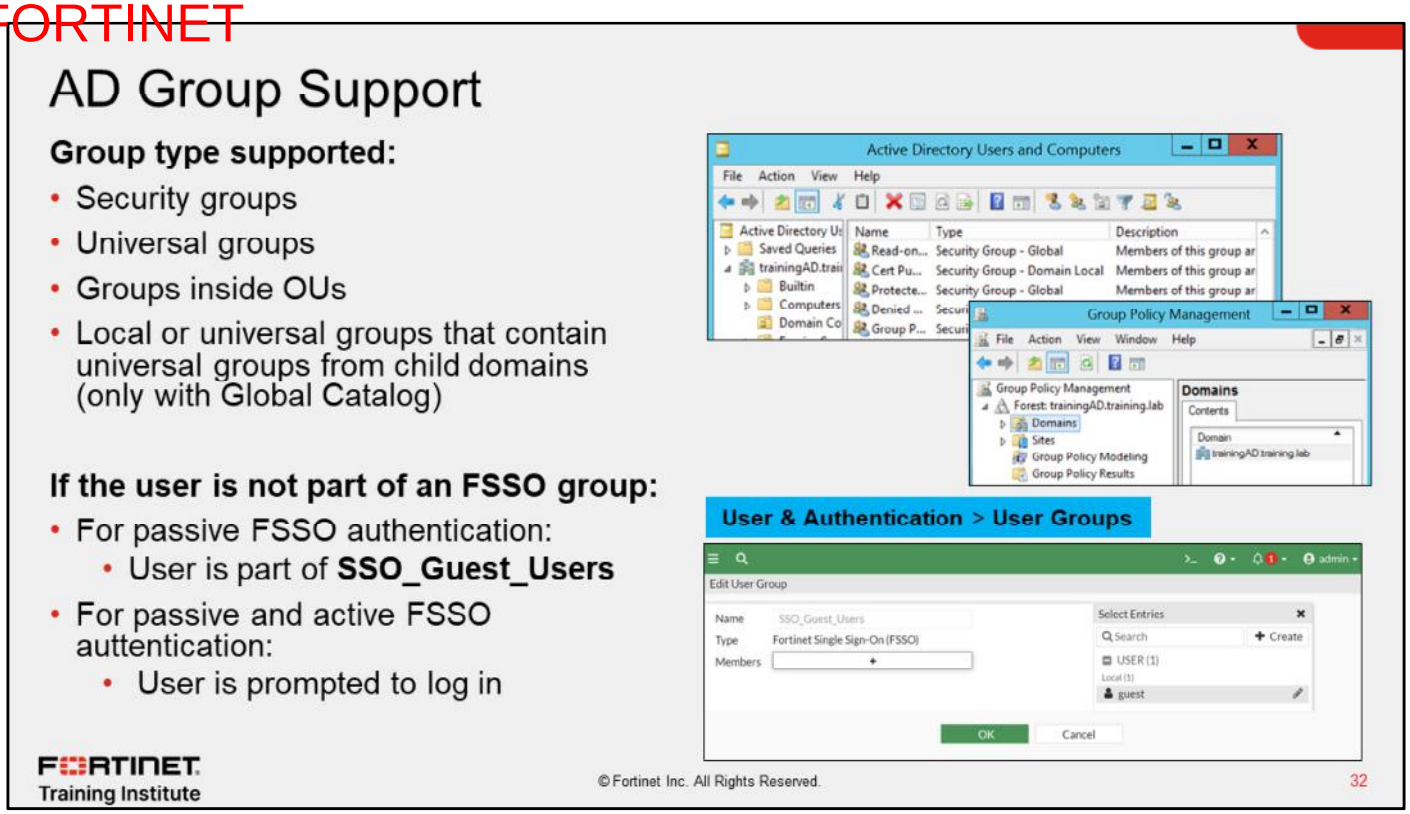

In AD settings, not all group types are supported. AD settings supports filtering groups only from:

- Security groups
- Universal groups
- Groups inside OUs
- Local or universal groups that contain universal groups from child domains (only with Global Catalog)

All FortiGate configurations include a user group called **SSO_Guest_Users**. When only passive authentication is used, all the users that do not belong to any FSSO group are automatically included in this guest group.

This allows an administrator to configure limited network access to guest users that do not belong to the Windows AD domain.

However, if both passive and active authentication are enabled for specific traffic, you cannot use **SSO_Guest_Users**, because traffic from IP addresses not on the FSSO user list must be prompted to enter their credentials.

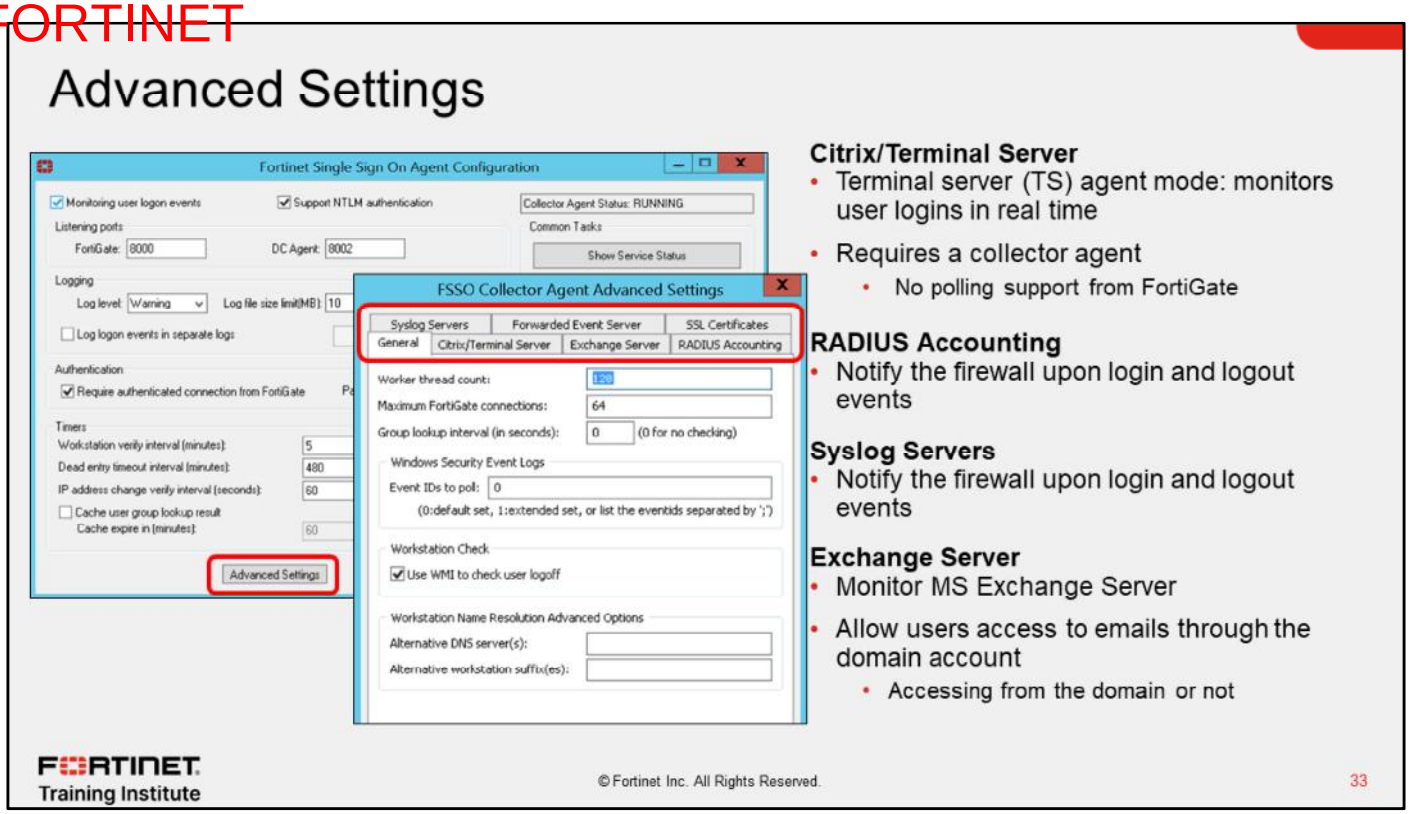

Depending on your network, you might need to configure advanced settings in your FSSO collector agent.

Citrix servers support FSSO. Terminal server (TS) agent mode allows the server to monitor user logins in real time. The TS agent is like a DC agent, it also needs the collector agent to collect and send the login events to FortiGate. It then uses the same ports to report the logins back to the collector agent.

The collector agent on its own can get accurate login events only from Citrix servers if each user gets their own IP address. Otherwise, if multiple users share the same IP address, the TS agent is needed so that it can report to the collector agent the user, IP address, and source port range assigned to that user. The TS agent cannot forward logs directly to FortiGate, the logs first have to be gathered by a collector. This does not work with polling from FortiGate.

A RADIUS server configured as a RADIUS-based accounting system can interact in your network by sending accounting messages to the collector agent. The FSSO collector agent also supports integration with syslog servers for the same purpose.

The FSSO collector agent can also monitor a Microsoft Exchange server, which is useful when users access their email using their domain account.

For **Windows Security Event Logs** polling mode, you can configure **Event IDs to poll** here. For specific event IDs, visit the Fortinet Knowledge Base (http://kb.fortinet.com).

Knowledge Check

1. If you have collector agents using either the DC agent mode or the collector agent-based polling mode, which fabric connector should you select on FortiGate?
   - A. Poll Active Directory Server
   - ✓ B. Fortinet Single Sign-On Agent

2. Which naming conventions does the FSSO collector agent use to access the Windows AD in **Standard** access mode?
   - ✓ A. Windows convention - NetBios: `Domain\groups`
   - B. LDAP convention: `CN=User,OU=Name,DC=Domain`

FORTINET Training Institute

34

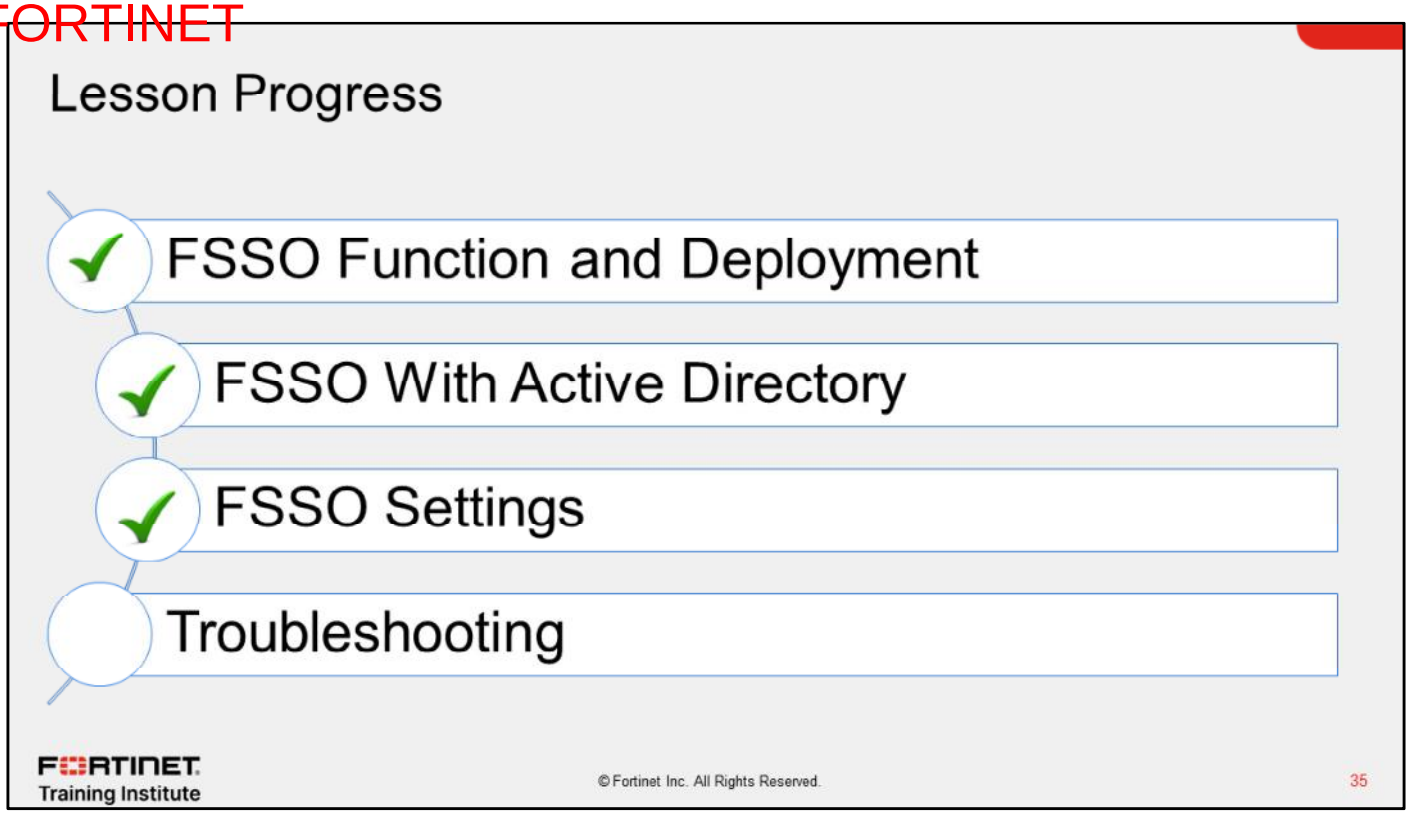

Good job! You now understand how to configure the SSO settings on FortiGate and the FSSO collector agent.

Now, you'll learn about some basic troubleshooting options.

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in FSSO monitoring and troubleshooting, you will be able to prevent, identify, and solve common issues related to FSSO.

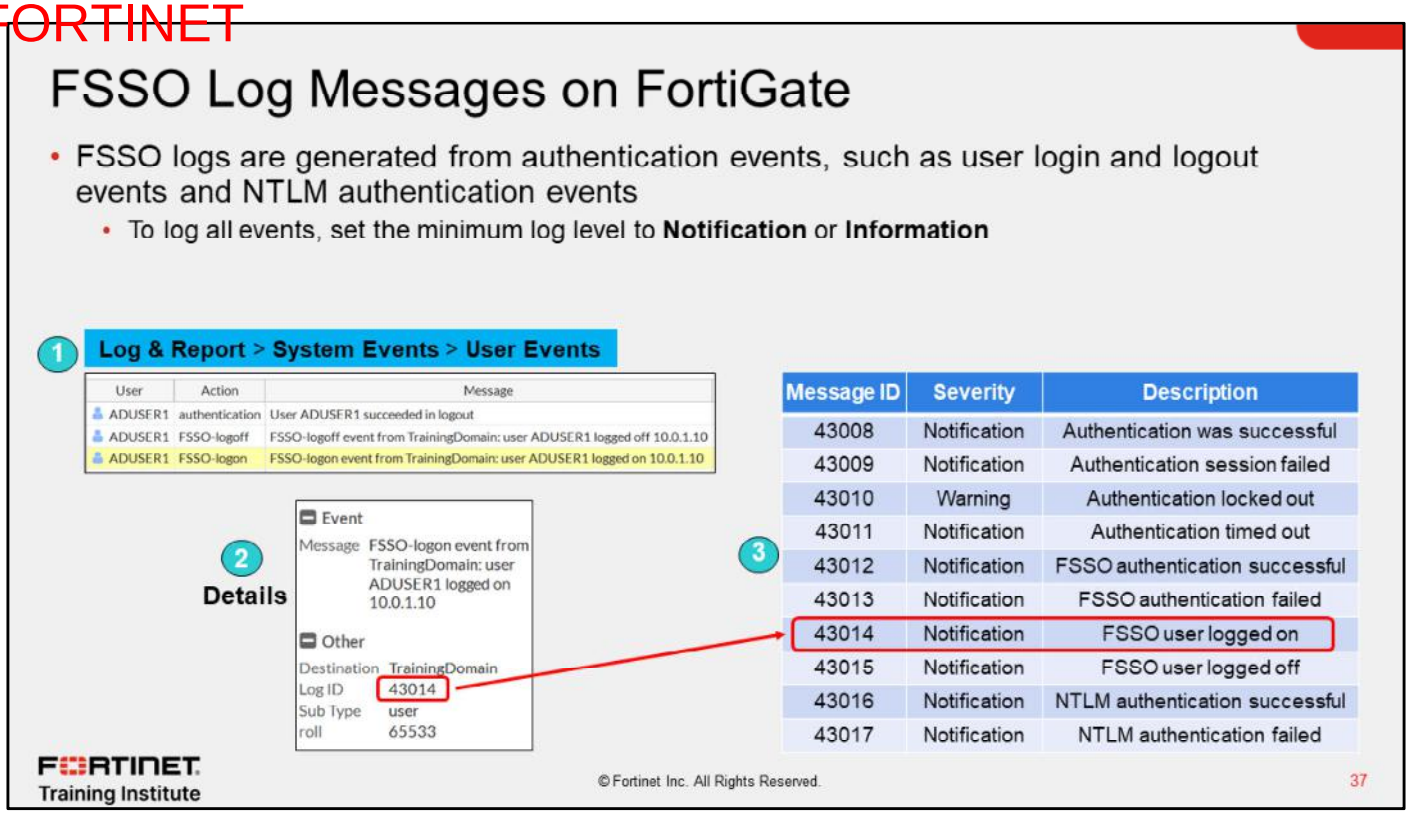

FSSO-related log messages are generated from authentication events. These include user login and logout events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues.

To ensure you log all the events needed, set the minimum log level to **Notification** or **Information**. Firewall logging requires **Notification** as a minimum log level. The closer the log level is to **Debug** level, the more information is logged.

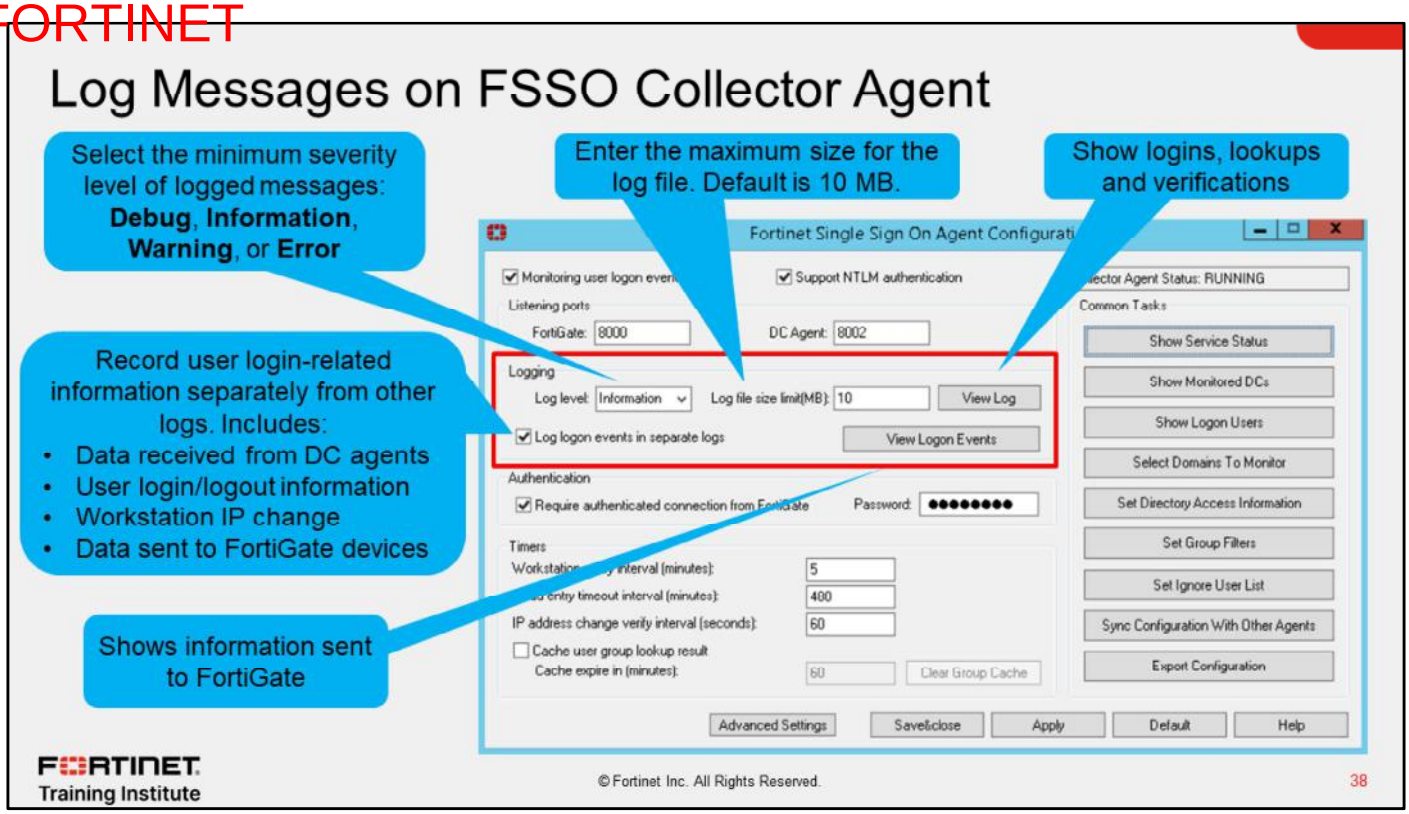

When troubleshooting FSSO agent-based deployments, you might want to look at the log messages generated directly on the FSSO collector agent.

The **Logging** section of the FSSO collector agent allows the following configurations:

- **Log level**: Select the minimum severity level of logged messages. Includes these levels:
  - **Debug**: the most detailed log level. Use it when actively troubleshooting issues.
  - **Information**: includes details about login events and workstation checks. This is the recommended level for most troubleshooting.
  - **Warning**: the default level. It provides information about failures.
  - **Error**: lists only the most severe events.
- **Log file size limit (MB)**: Enter the maximum size for the log file in MB. The default is 10.
- **View Log**: View all FSSO agent logs.
- **Log login events in separate logs**: Record user login-related information separately from other logs. The information in this log includes: data received from DC agents, user login/logout information, workstation IP change information, and data sent to FortiGate devices. When selected, a summary of events sent and removed from FortiGate is listed under **View login Events**, while all other information remains under **View Log**.
- **View login Events**: If **Log login events in separate logs** is enabled, you will can view user login-related information.

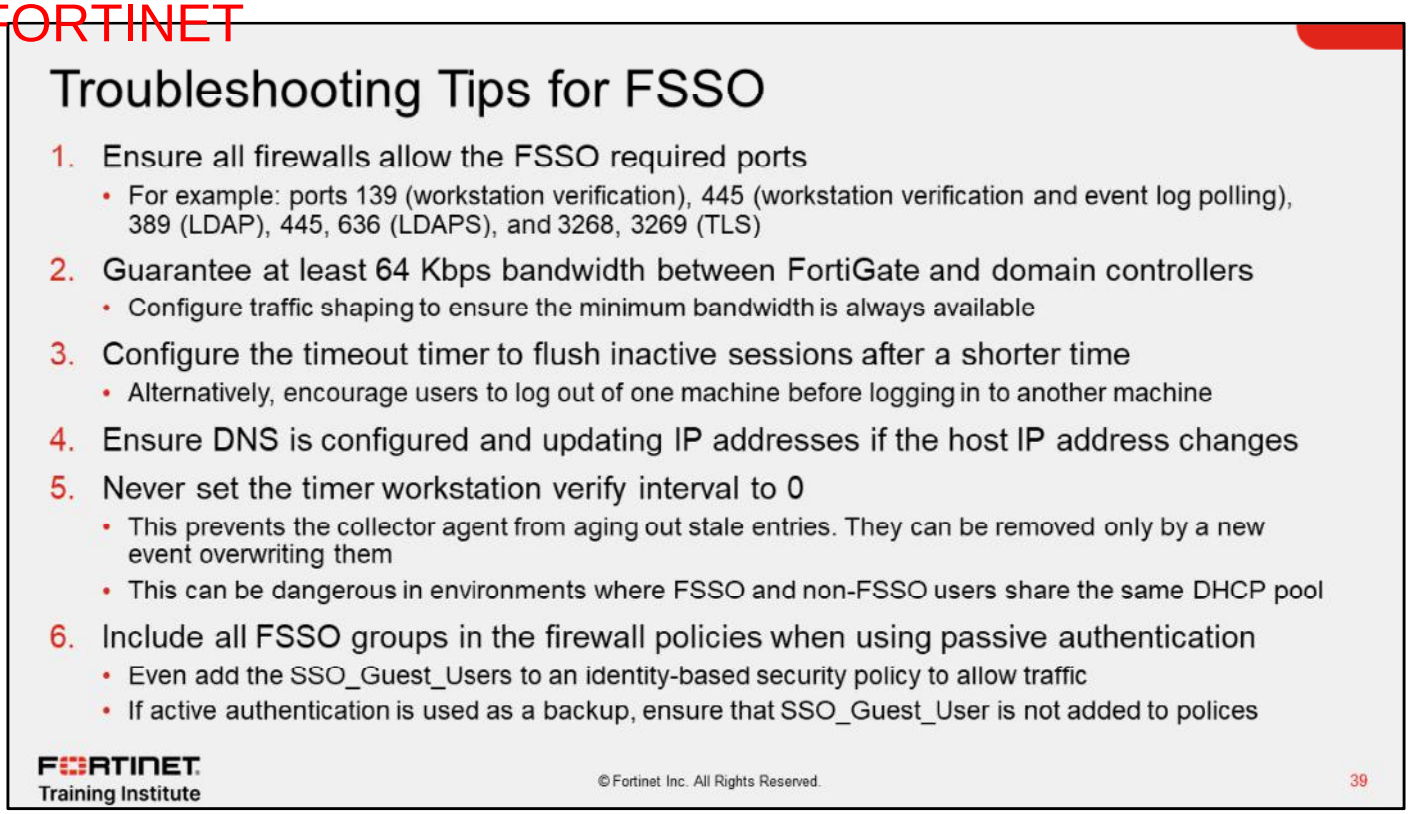

Begin with the following tips, which are useful in many FSSO troubleshooting situations:

- FSSO has a number of required ports that you must allow through all firewalls, or connections will fail. These include ports: 139 (workstation verification), 445 (workstation verification and event log polling), 389 (LDAP), and 445 and 636 (LDAPS).
- Configure traffic shaping between FortiGate and the domain controllers to ensure that the minimum bandwidth is always available. If there is insufficient bandwidth, some FSSO information might not reach FortiGate.
- In an all-Windows environment, flush inactive sessions. Otherwise, you can have a session for a non-authenticated machines go out as an authenticated user. This can occur if the DHCP lease expires for the authenticated user with the collector agent being able to verify that the user has indeed logged out.
- Ensure DNS is configured correctly and updating IP addresses, if workstation IP addresses change.
- Never set the workstation verify interval to 0. This prevents the collector agent from aging out stale entries. They can be removed only by a new event overwriting them. This can be especially dangerous in environments where FSSO and non-FSSO users share the same DHCP pool.
- When using passive authentication only, include the group of guest users in a policy and give them access. Associate their group with a security policy. If you use active authentication as a backup, ensure you do not add SSO_Guest_User to polices. SSO_Guest_User and active authentication are mutually exclusive.

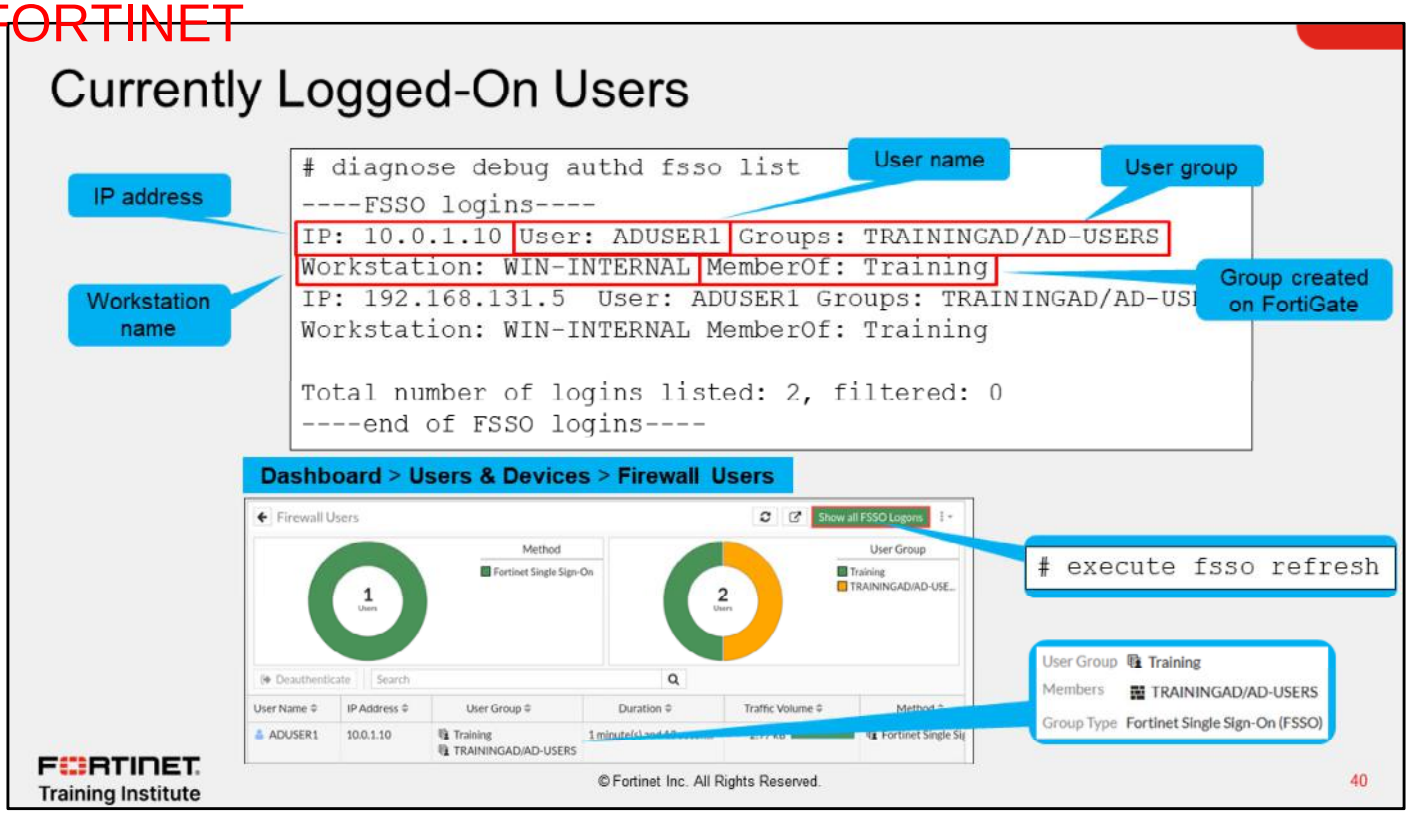

If applying the tips from the previous slide didn't solve your FSSO issues, you may need to apply some `debug` commands.

To display the list of FSSO users that are currently logged in, use the CLI command `diagnose debug authd fsso list`.

For each user, the user name, user group, IP address, and the name of the workstation from which they logged in shows. The `MemberOf` section shows the group that was created on the firewall, to which you mapped the AD group. The same group should show in the **User group** screen on the GUI.

Also, use `execute fsso refresh` to manually refresh user group information from any directory service servers connected to FortiGate, using the collector agent.

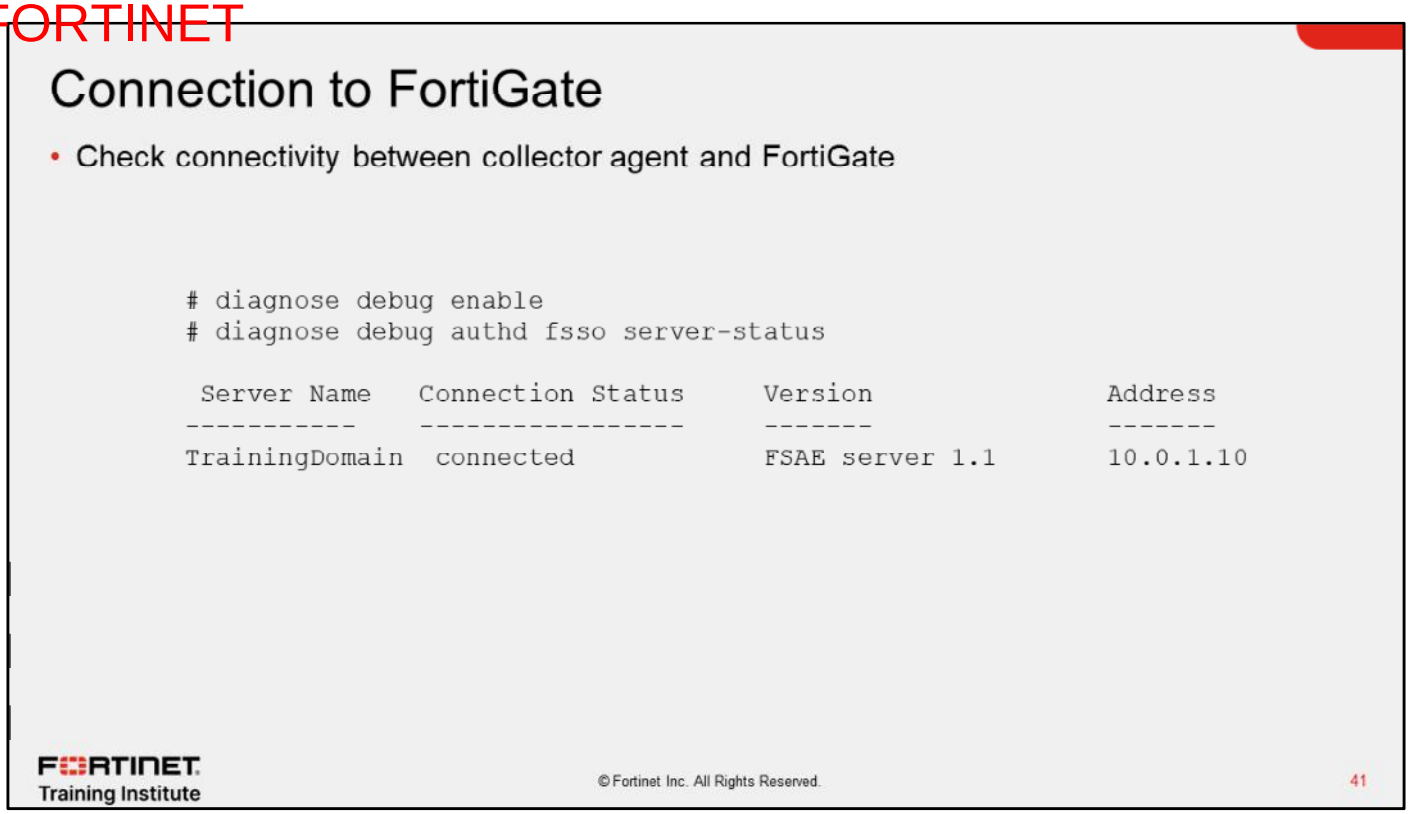

To show the status of communication between FortiGate and each collector agent, you can use the CLI command `diagnose debug authd fsso server-status`.

However, before you use that command, you must first run the command `diagnose debug enable`.

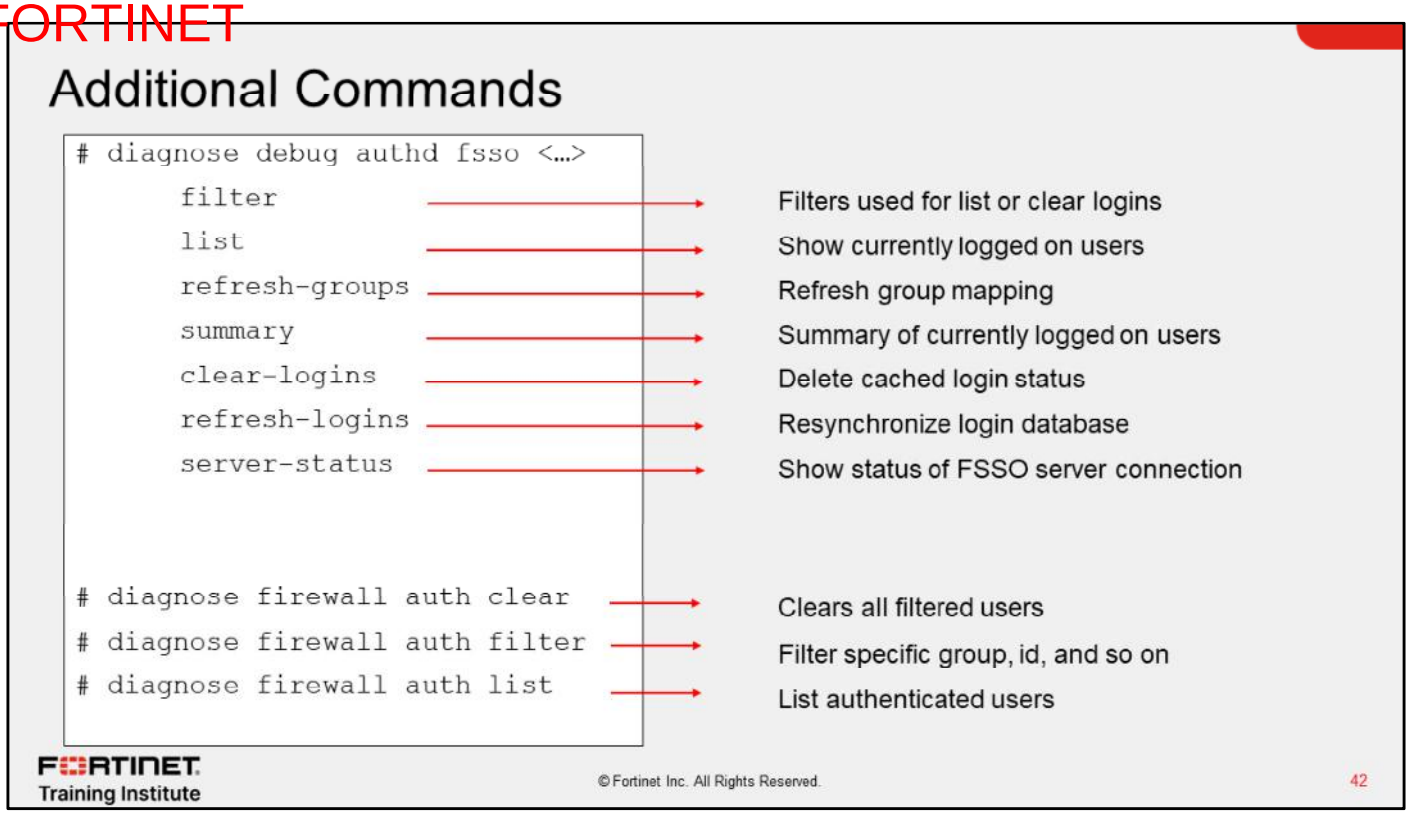

Also, available under `diagnose debug authd fsso` are commands for clearing the FortiGate cache of all currently logged in users, filtering the display of the list of logged in users, and refreshing the login and user group information.

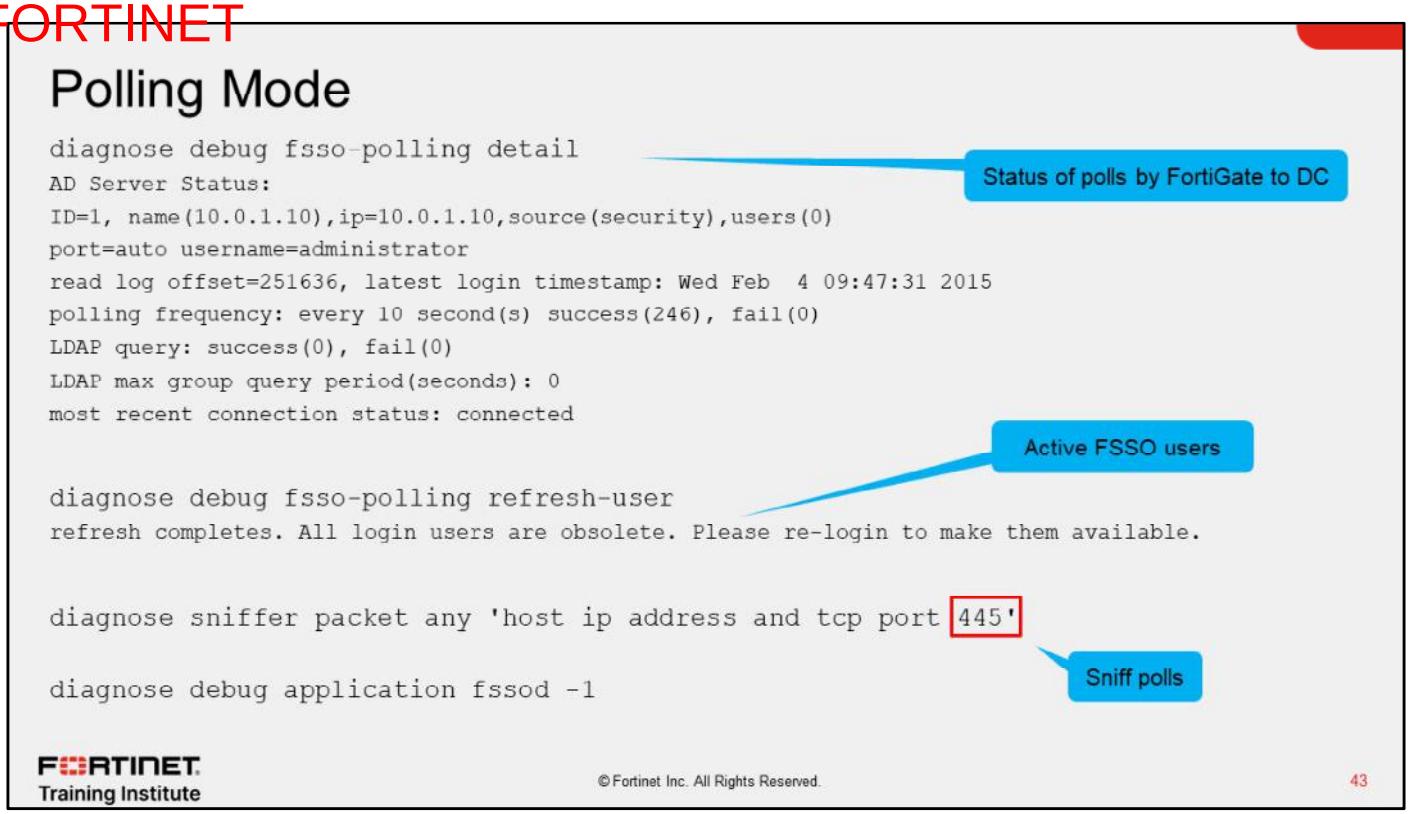

The command `diagnose debug fsso-polling detail` displays status information and some statistics related to the polls done by FortiGate on each DC in agentless polling. If the `read log offset` is incrementing, FortiGate is connecting to and reading the logs on the domain controller. If the `read log offset` is incrementing but you are not getting any login events, check that the group filter is correct and that the domain controller is creating the correct event IDs.

The command `diagnose debug fsso-polling refresh-user` flushes information about all the active FSSO users.

In agentless polling mode, FortiGate frequently polls the event viewer to get the login events. You can sniff this traffic on port 445.

Also, there is a specific FortiGate daemon that handles polling mode. It is the `fssod` daemon. To enable agentless polling mode real-time debug, use the `diagnose debug application fssod -1` command.

Knowledge Check

1. Which logging level shows the login events on the collector agent?
   - ✓ A. Information
   - B. Warning

2. The command `diagnose debug fsso-polling detail` displays information for which mode of FSSO?
   - ✓ A. Agentless polling
   - B. Collector agent-based polling

FORTINET Training Institute

44

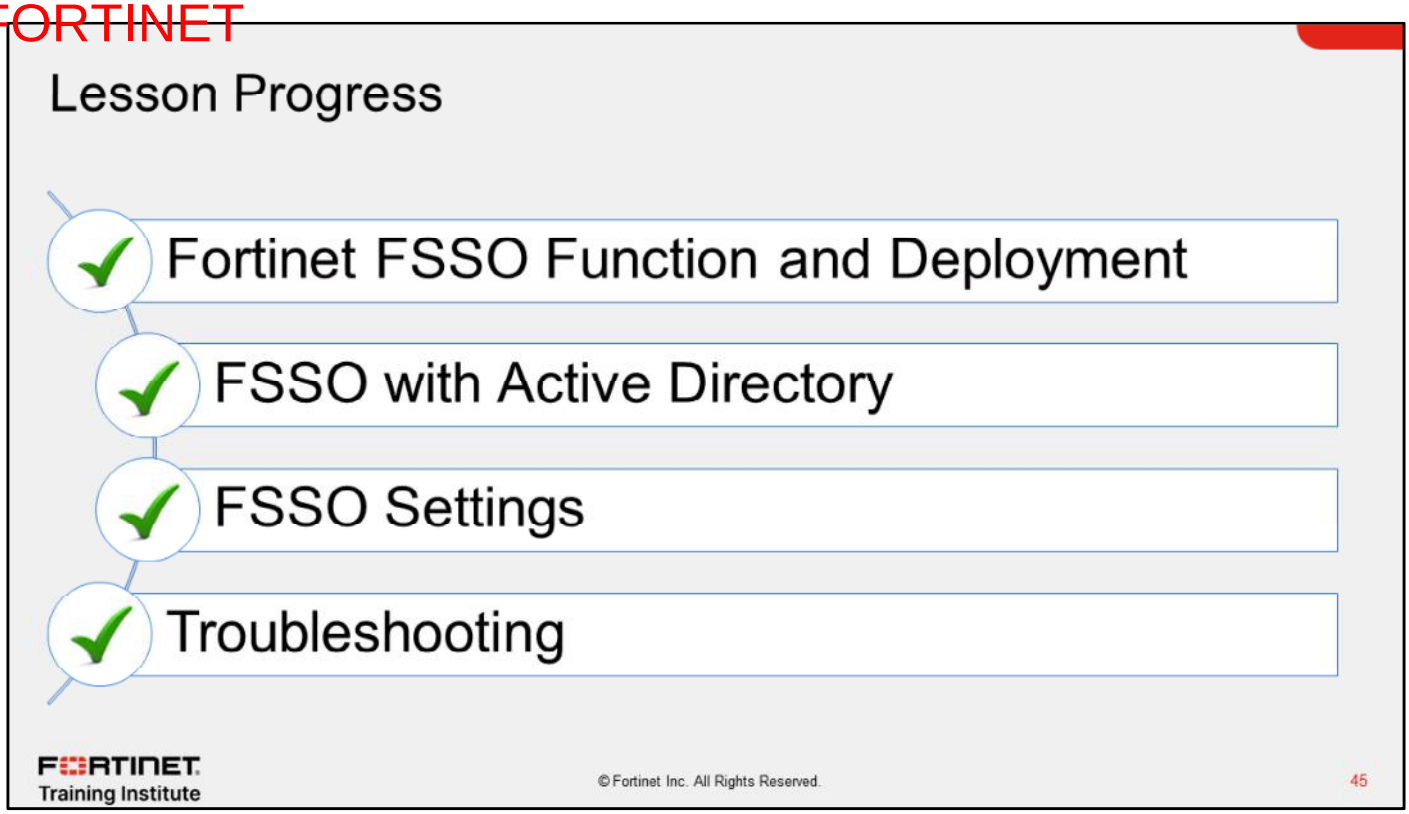

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

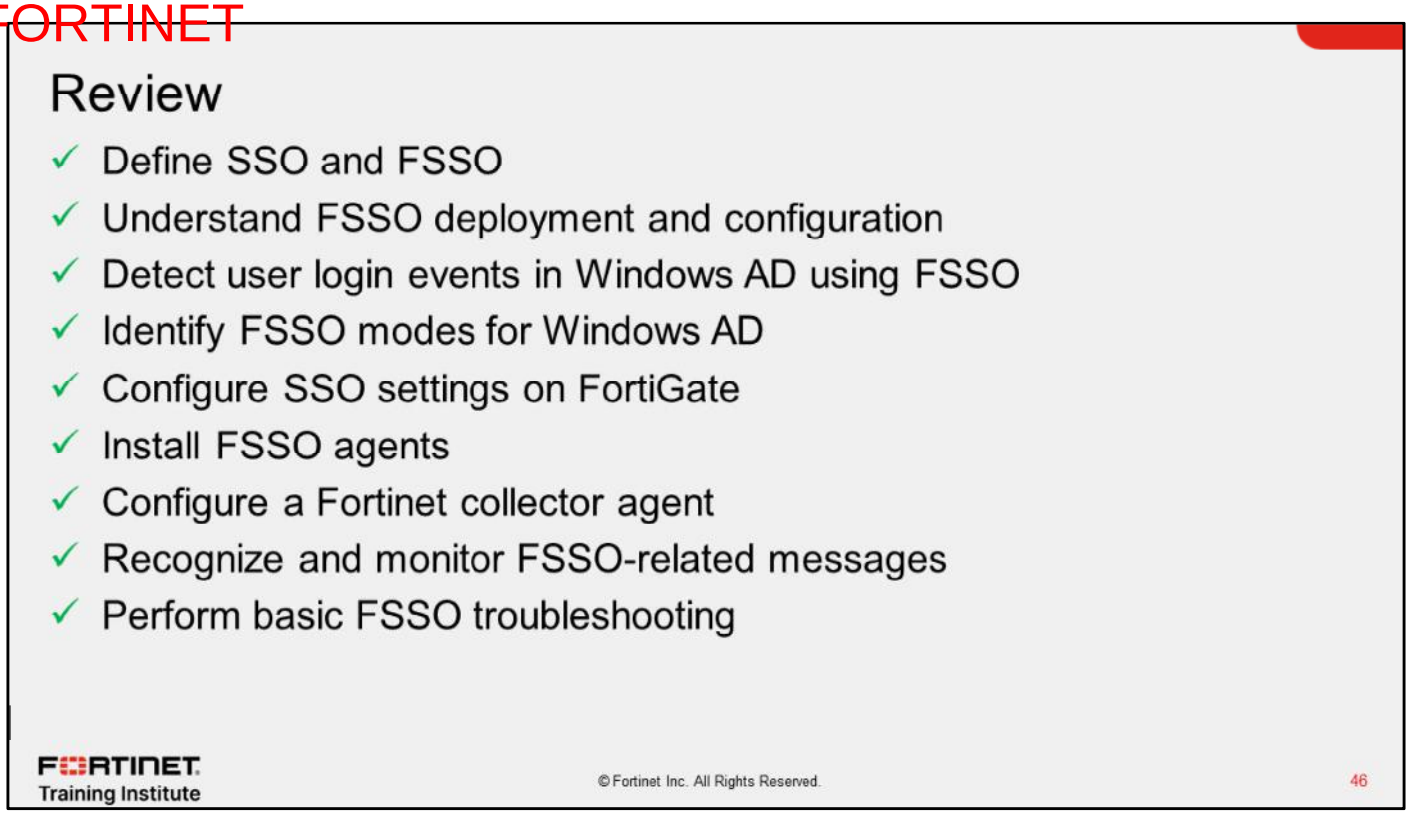

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use FSSO so that your users don't need to log in each time they access a different network resource.

In this lesson, you will learn about zero-trust network access (ZTNA).

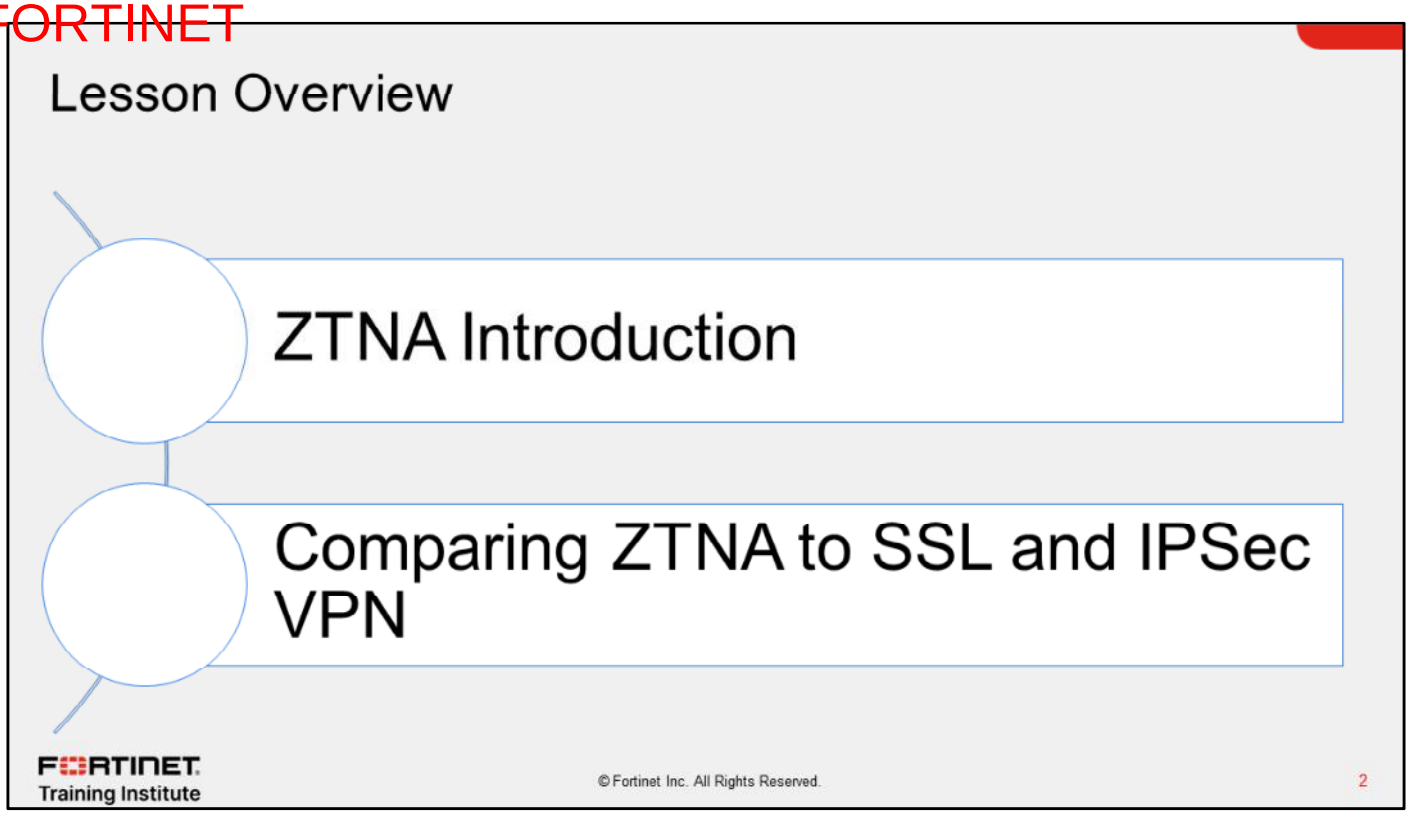

In this lesson, you will learn about the topics shown on this slide.

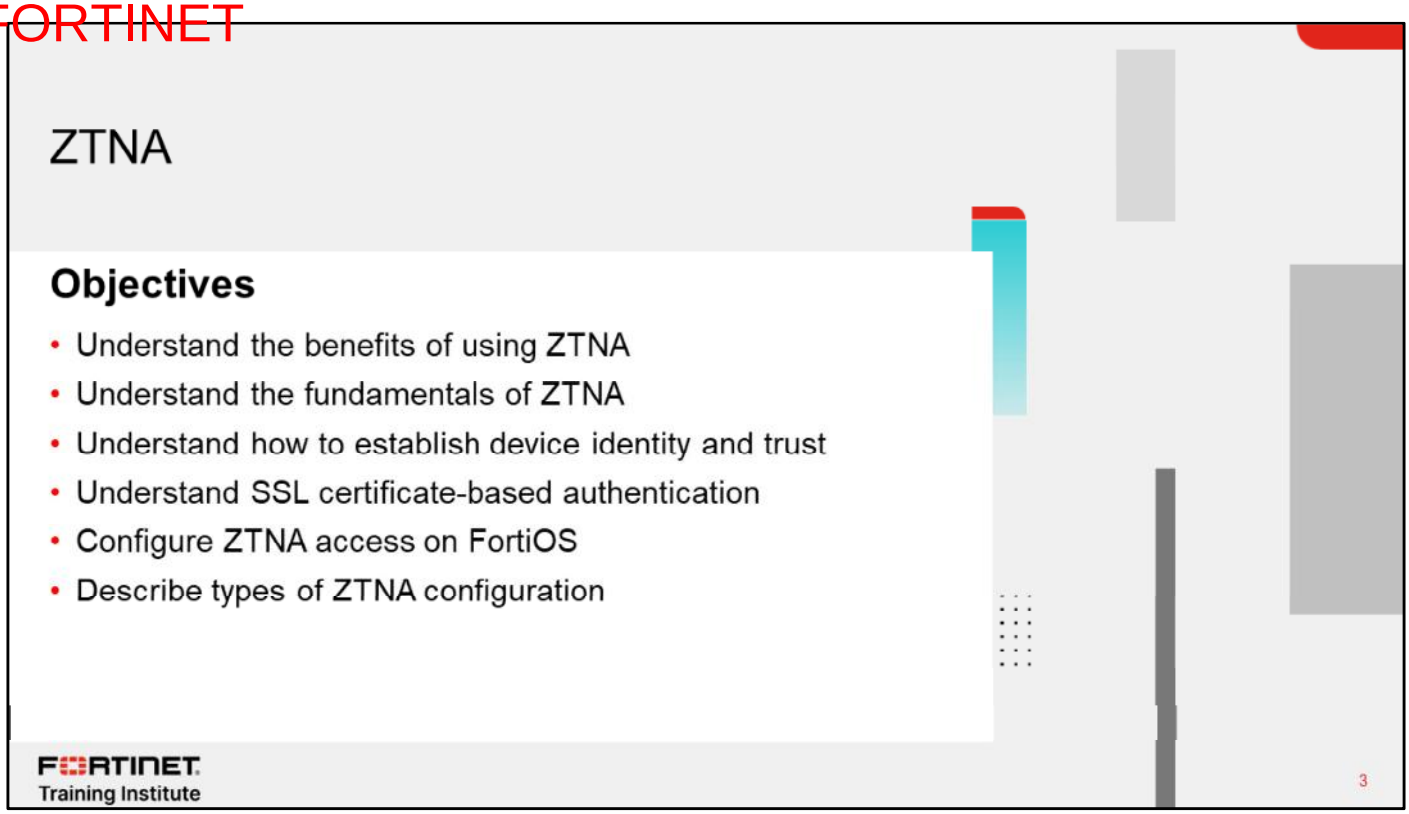

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in ZTNA, you will be able to understand key ZTNA concepts and how to configure ZTNA.

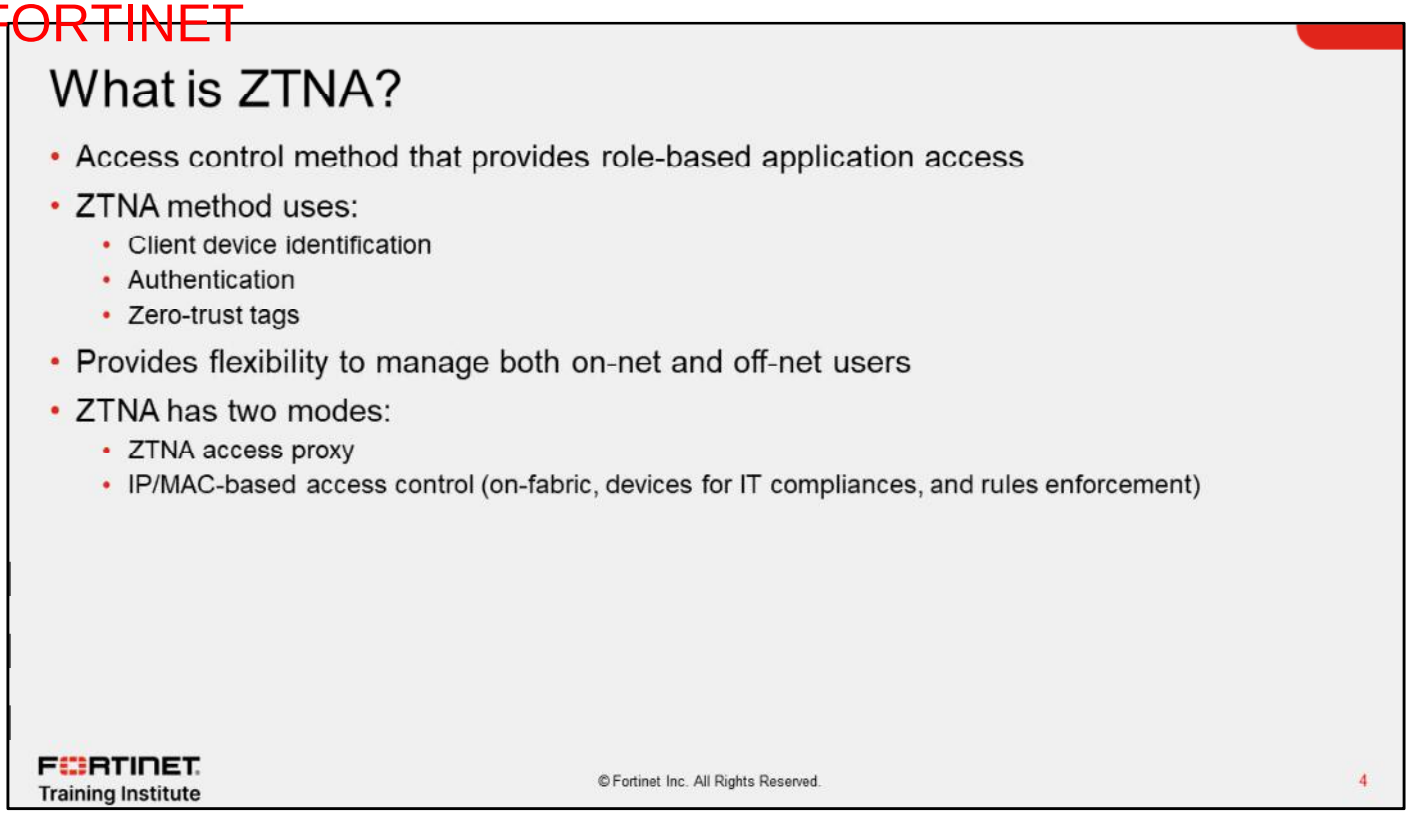

ZTNA is an access control method that uses client device identification, authentication, and zero-trust tags to provide role-based application access. ZTNA gives administrators the flexibility to manage network access for on-fabric local users and off-fabric remote users. ZTNA grants access to applications only after a device verification, authenticating the user's identity, authorizing the user, and then performing context-based posture checks using zero-trust tags.

Traditionally, a user and a device have different sets of rules for on-fabric access and off-fabric VPN access to company resources. With a distributed workforce, and access that spans company networks, data centers, and the cloud, managing the rules can be complex. User experience is also affected when an organization needs multiple VPNs to access various resources.

ZTNA has two modes:

- ZTNA access proxy allows users to securely access resources through an SSL-encrypted access proxy. This simplifies remote access by eliminating the use of VPNs.
- IP/MAC filtering uses ZTNA tags to provide an additional factor for identification, and a security posture check to implement role-based zero-trust access. IP/MAC-based access control enhances security when endpoints are physically located on the corporate network, whereas ZTNA access proxy focuses on access for remote users. IP/MAC-based access control combines IP/MAC with ZTNA tags for identification and security posture check to implement role-based zero-trust access. Firewall policies are configured that use ZTNA tags to control access between on-net devices and an internal webserver. This mode does not require the use of the access proxy, and only uses ZTNA tags for access control.

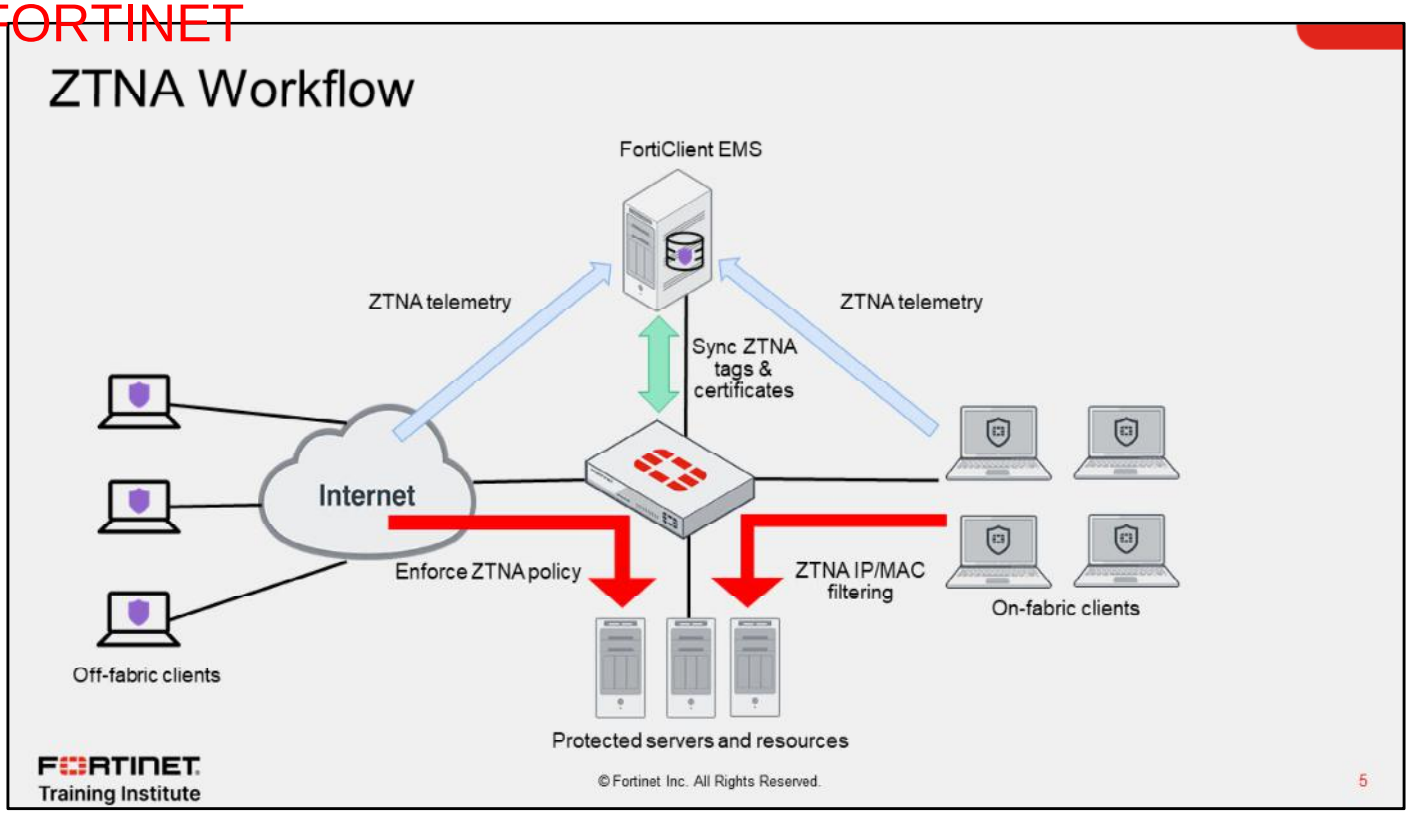

This slide demonstrates ZTNA telemetry, tags, and policy enforcement. You configure ZTNA tag conditions and policies on FortiClient EMS. FortiClient EMS shares the tag information with FortiGate through Security Fabric integration. FortiClient communicates directly with FortiClient EMS to continuously share device status information through ZTNA telemetry. FortiGate can then use ZTNA tags to enforce access control rules to incoming traffic through ZTNA access.

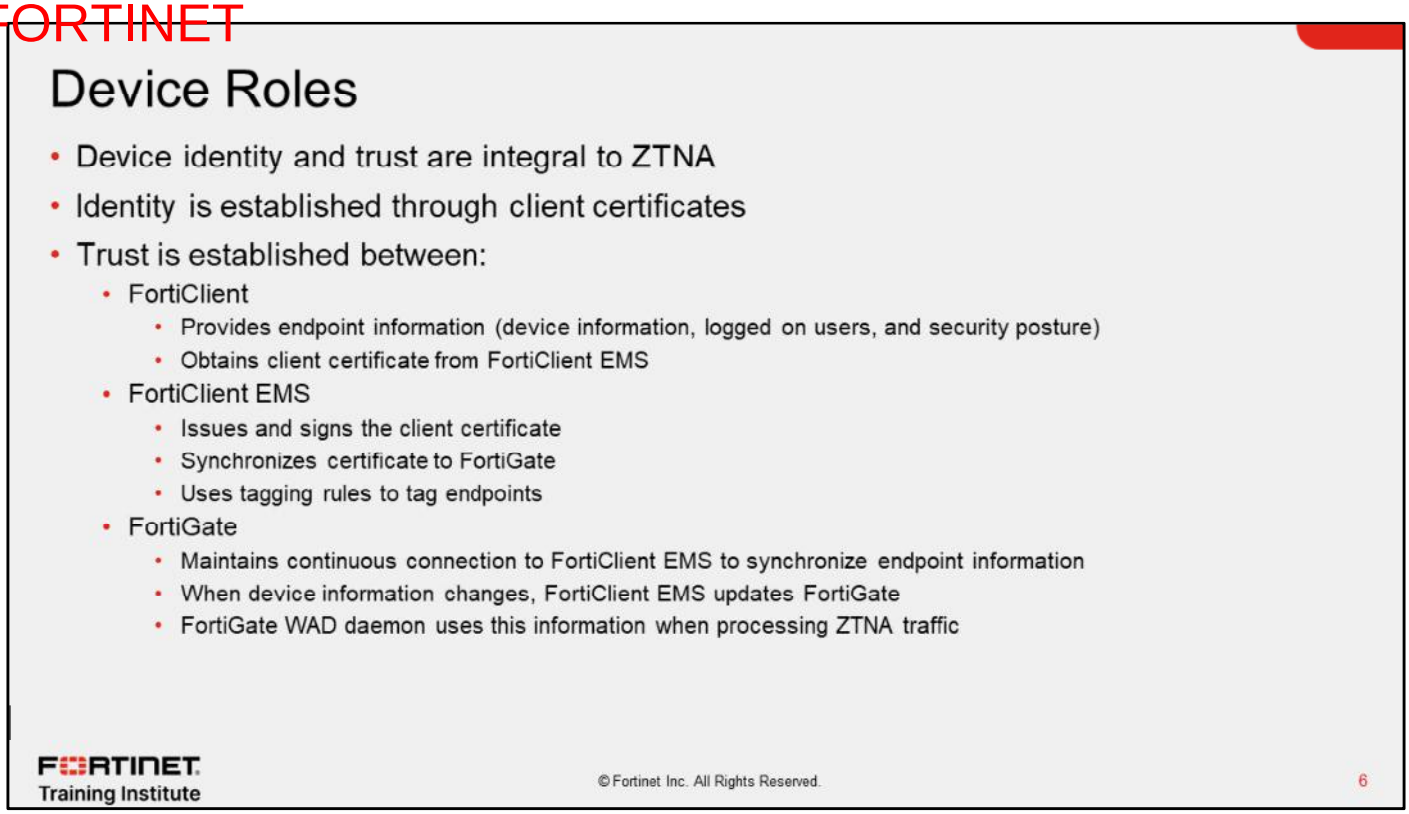

Device identity and trust are integral to ZTNA. Device identity is established through client certificates, and trust is established among FortiClient, FortiClient EMS, and FortiGate devices. In ZTNA, devices perform specific roles.

FortiClient provides the following information to FortiClient EMS when it registers:
- Device information (network details, operating system, model, and so on)
- Logged in user information
- Security posture (on-fabric and off-fabric, antivirus software, vulnerability status, and so on)

FortiClient also requests and obtains a client device certificate from the EMS ZTNA Certificate Authority (CA) on its first attempt to connect to the access proxy. The client uses this certificate to identify itself to FortiGate.

FortiClient EMS issues and signs the client certificate with the FortiClient UID, certificate serial number, and EMS serial number. FortiClient EMS then synchronizes the certificate with FortiGate. FortiClient EMS also shares its EMS ZTNA CA certificate with FortiGate, so that FortiGate can use it to authenticate the clients. FortiClient EMS uses zero-trust tagging rules to tag endpoints based on the information that it has on each endpoint. FortiClient EMS also shares the tags with FortiGate.

FortiGate maintains a continuous connection to FortiClient EMS to synchronize endpoint device information such as FortiClient UID, client certificate SN, FortiClient EMS SN, network details (IP and MAC address), and so on. When device information changes, such as when a client moves from on-fabric to off-fabric, or their security posture changes, FortiClient EMS updates the device information, and then updates the FortiGate.

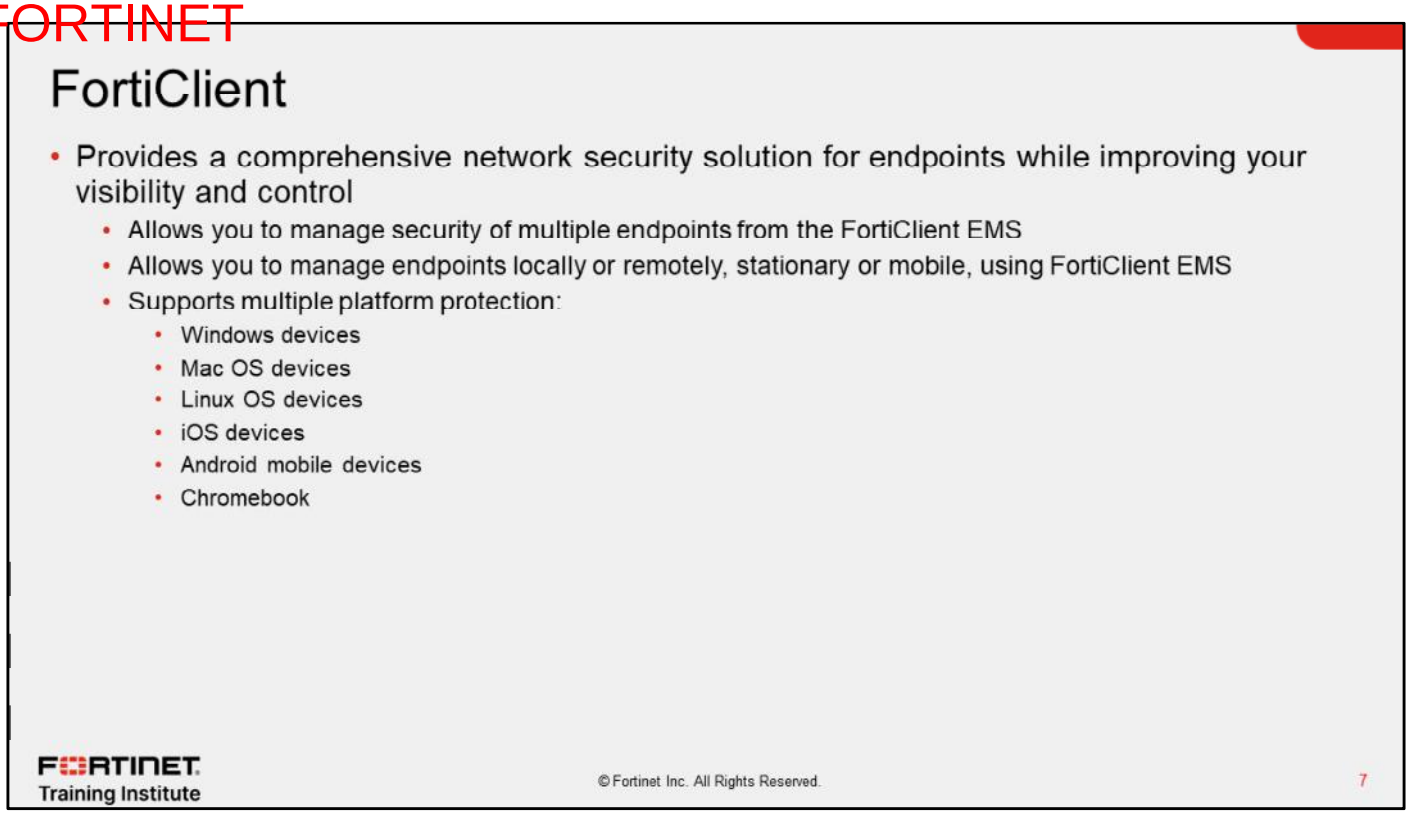

FortiClient provides comprehensive endpoint protection for your Windows-based, Mac-based, and Linux-based desktops, laptops, file servers, and mobile devices such as iOS and Android. It helps you to safeguard your systems with advanced security technologies, all of which you can manage from a single management console.

FortiClient enables every device—local or remote, stationary or mobile—to integrate with your FortiClient EMS and FortiGate. FortiClient supports Windows, Mac OS, Linux, iOS, Android mobile devices and Chromebook, and also integrates your home offices, mobile workers, and visiting partners.

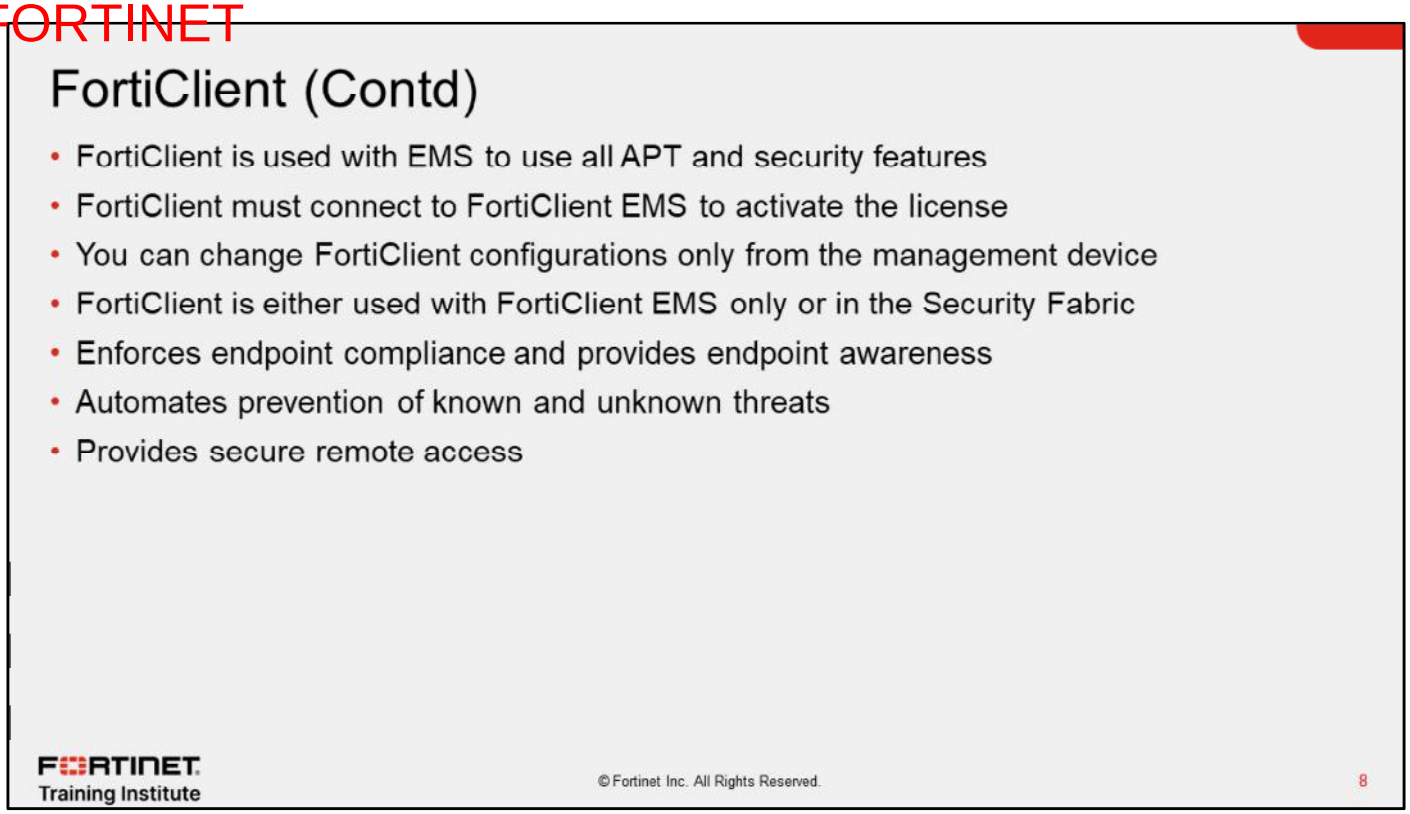

FortiClient must be used with FortiClient EMS. FortiClient must connect to FortiClient EMS to activate its license and become provisioned by the endpoint profile that the administrator configured in FortiClient EMS. You cannot use any FortiClient features until FortiClient is connected to FortiClient EMS and licensed.

When FortiClient is connected only to FortiClient EMS, FortiClient EMS provisions and manages FortiClient. FortiClient EMS also sends zero-trust tagging rules to FortiClient, and uses the results from FortiClient to dynamically group endpoints in EMS. Only FortiClient EMS can control the connection between FortiClient and FortiClient EMS. However, FortiClient cannot participate in the Fortinet Security Fabric.

FortiClient in the security fabric connects to FortiClient EMS to receive a profile of configuration information as part of an endpoint policy. FortiClient EMS is connected to FortiGate to participate in the Security Fabric. FortiClient EMS sends FortiClient endpoint information to FortiGate. FortiGate can also receive dynamic endpoint group lists from FortiClient EMS and use them to build dynamic firewall policies.

FortiClient also provides secure remote access to corporate assets through VPN.

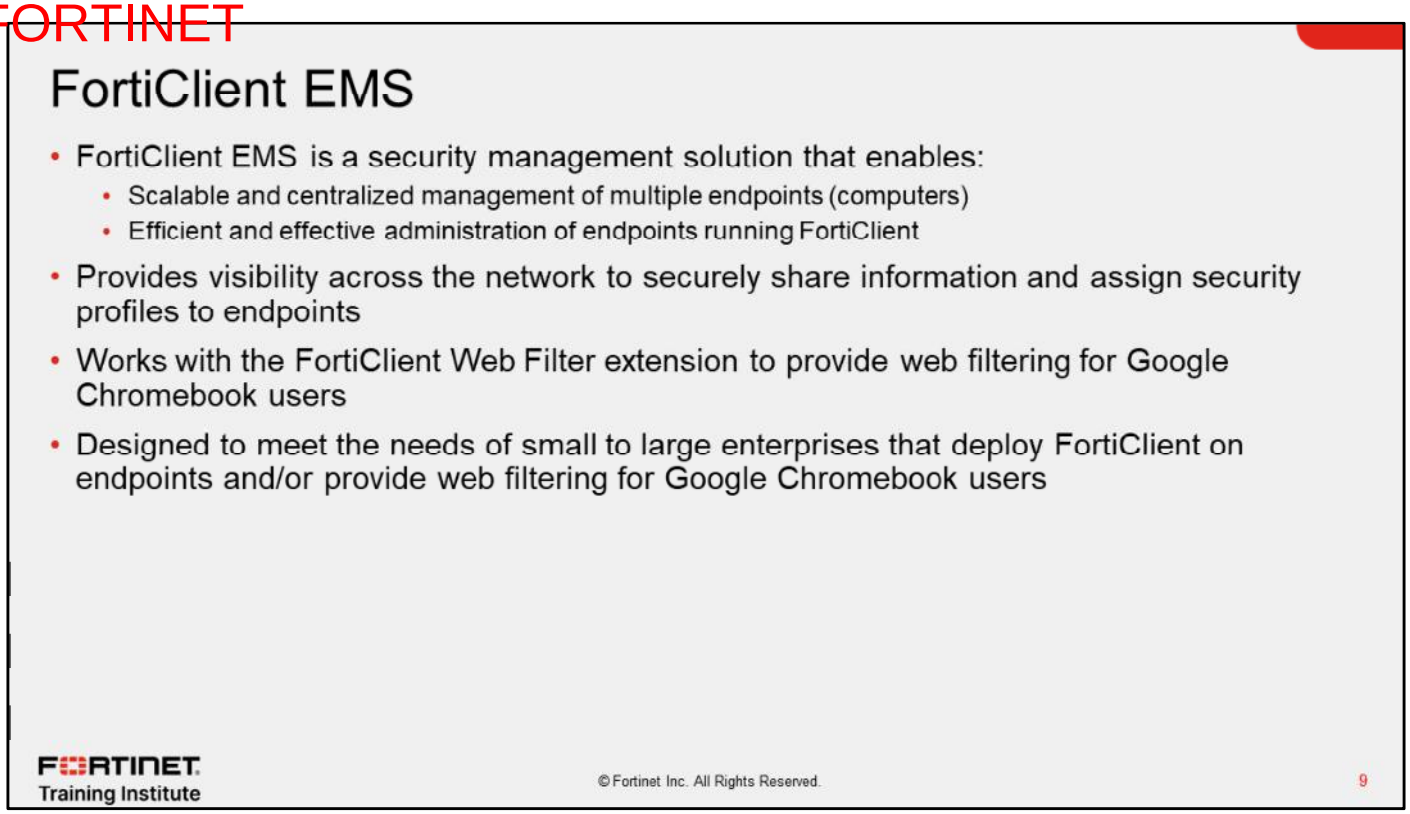

FortiClient EMS is a security management solution that enables scalable and centralized management of multiple endpoints (computers). It also provides efficient and effective administration of endpoints running FortiClient, and visibility across the network to securely share information and assign security profiles to endpoints. It is designed to maximize operational efficiency and includes automated capabilities for device management and troubleshooting.
FortiClient EMS also works with the FortiClient Web Filter extension to provide web filtering for Google Chromebook users.

The benefits of deploying FortiClient EMS include:
- Remotely deploying FortiClient software to Windows computers
- Updating profiles for endpoint users regardless of access location
- Administering FortiClient endpoint connections, such as accepting, disconnecting, and blocking connections
- Managing and monitoring endpoints, such as status, system, and signature information
- Identifying outdated versions of FortiClient software
- Defining web filtering rules in a profile, and remotely deploying the profile to the FortiClient Web Filter extension on Google Chromebook endpoints

You can manage endpoint security for Windows and macOS platforms using a unified organizational security policy. An organizational security policy provides a full, understandable view of the security policies defined in the organization. You can see all policy rules, assignments, and exceptions in a single unified view. FortiClient EMS is part of the Fortinet Endpoint Security Management suite, which ensures comprehensive policy administration and enforcement for an enterprise network.

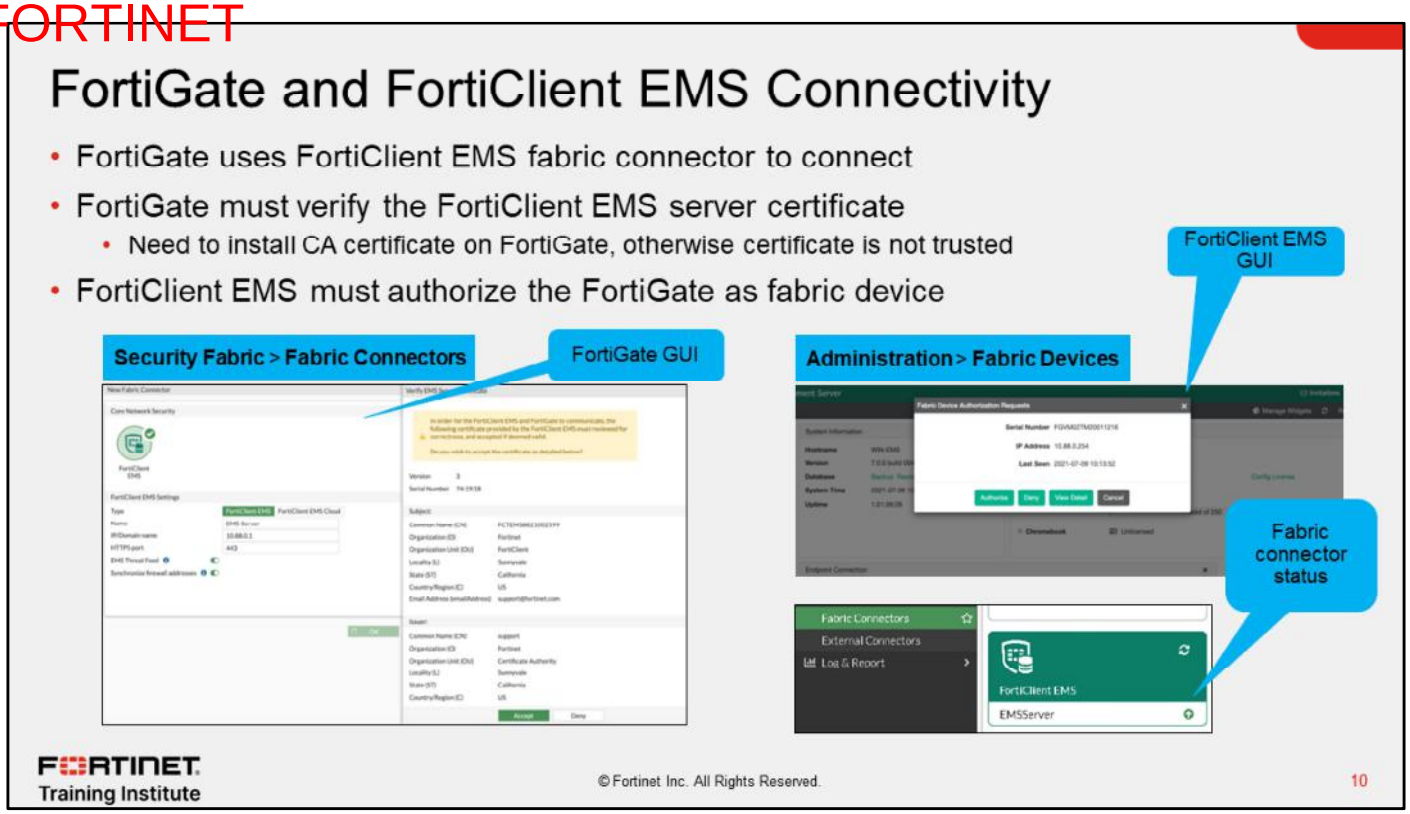

You can configure the on-premises FortiClient EMS connector on FortiGate by clicking **Security Fabric** > **Fabric Connectors**. After applying the FortiClient EMS settings, FortiGate must accept the FortiClient EMS server certificate. However, when you configure a new connection to FortiClient EMS server, the certificate might not be trusted. To resolve, you must manually export and install the root CA certificate on FortiGate. The FortiClient EMS certificate that is used by default for the SDN connection is signed by the CA certificate that is saved on the Windows server when you first install FortiClient EMS. This certificate is stored in the **Trusted Root Certification Authorities** folder on the server. For more information about exporting and installing certificates on FortiGate, refer to the *FortiOS-7.0.1 Administration Guide*.

Next, you must authorize FortiGate on FortiClient EMS. If you log in to FortiClient EMS, a pop-up window opens, requesting you to authorize FortiGate. If you do not log in, you can click **Administration** > **Devices**, select the FortiGate device, and then authorize it. Note that the FortiClient EMS connector status appears down until you authorize FortiGate on FortiClient EMS.

FortiGate automatically synchronizes ZTNA tags after it connects to FortiClient EMS.

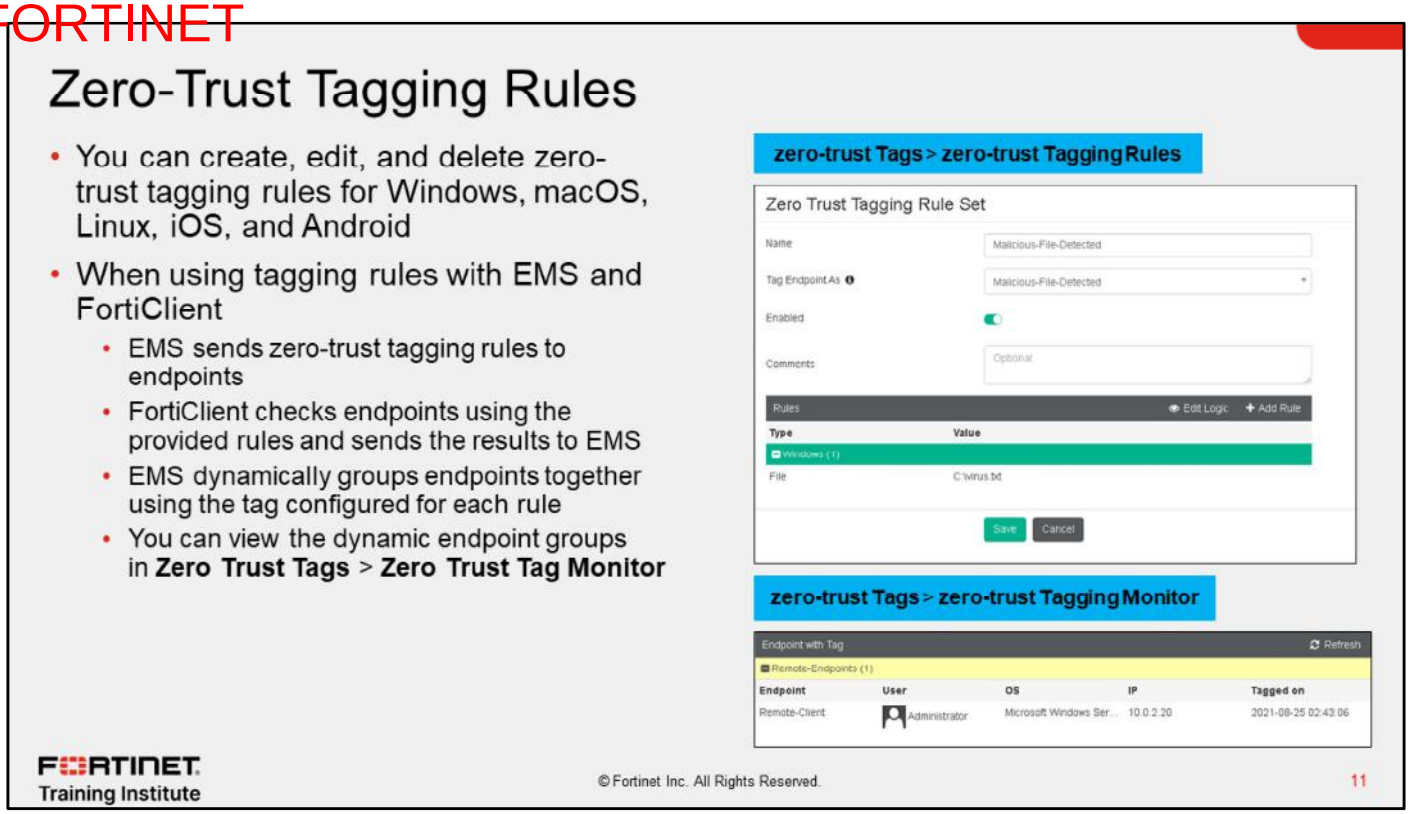

You can create, edit, and delete zero-trust tagging rules for Windows, macOS, Linux, iOS, and Android endpoints. The following happens when using zero-trust tagging rules with FortiClient EMS and FortiClient:

- FortiClient EMS sends zero-trust tagging rules to endpoints through telemetry communication.
- FortiClient checks endpoints using the provided rules and sends the results to FortiClient EMS.
- FortiClient EMS receives the results from FortiClient.
- FortiClient EMS dynamically groups endpoints together using the tag configured for each rule. You can view the dynamic endpoint groups by clicking **Zero Trust Tags** > **Zero Trust Tag Monitor**.

Note that when the endpoint network changes or user login and logout events occur, FortiClient triggers an X-FFCK-TAG message to EMS, even if there are no tag changes. After FortiClient EMS receives the tags, it processes them immediately, and updates the FortiOS tags within five seconds of the REST API response. For other tag changes, FortiClient sends the information to FortiClient EMS regularly.

FortiClient EMS Certificate Management

- FortiClient EMS has a default root CA certificate
- ZTNA CA uses root certificate to sign CSRs from the FortiClient endpoints
- You can revoke and update root CA
  - Force updates to the FortiGate and FortiClient endpoints by generating new certificates
- FortiClient EMS manages individual client certificates

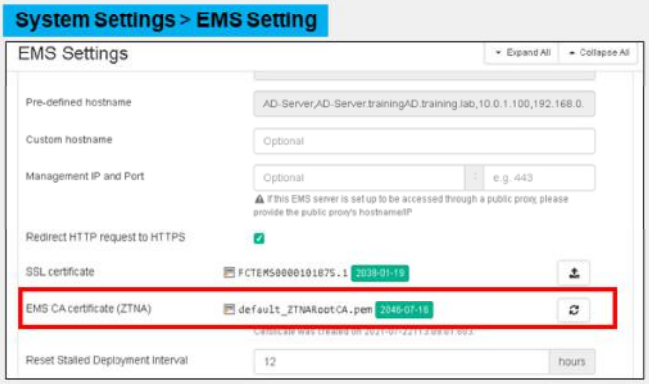

FORTINET Training Institute

12

FortiClient EMS has a **default_ZTNARootCA** certificate generated by default that the ZTNA CA uses to sign CSRs from the FortiClient endpoints. Clicking the refresh button revokes and updates the root CA, forcing updates to the FortiGate and FortiClient endpoints by generating new certificates for each client. FortiClient EMS can also manage individual client certificates. You can also revoke the certificate that is used by the endpoint when certificate private keys show signs of being compromised. Click **Endpoint** > **All Endpoints**, select the client, and then click **Action** > **Revoke Client Certificate**.

Do not confuse the FortiClient EMS CA certificate (ZTNA) with the SSL certificate. The latter is the server certificate that is used by FortiClient EMS for HTTPS access and fabric connectivity to the FortiClient EMS server.

FortiClient EMS Certificate Management (Contd)

- On Windows endpoints, FortiClient automatically installs certificates in the certificate store
  - Certificate information, such as certificate UID and SN, should match the information on FortiClient EMS and FortiGate
  - **Certificates** > **Personal** > **Certificates**
- You can verify by CLI command on FortiGate
  - `diagnose endpoint record list <optional IP address>`

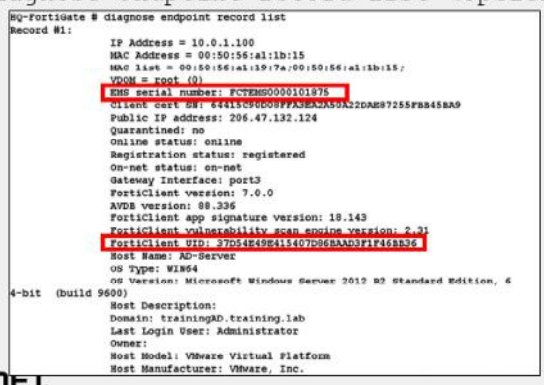

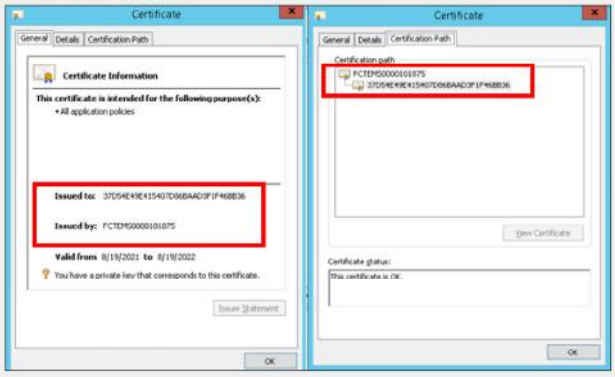

FORTINET Training Institute

13

In Windows, FortiClient automatically installs certificates in the certificate store. The certificate information in the store, such as certificate UID and SN, should match the information on FortiClient EMS and FortiGate. To locate certificates on other operating systems, consult the vendor documentation.

You can use the CLI command `diagnose endpoint record list` a to verify the presence of a matching endpoint record, and information such as the client UID, client certificate SN, and EMS certificate SN on the FortiGate. If any of the information is missing or incomplete, client certificate authentication might fail because FortiClient cannot locate the corresponding endpoint entry.

This slide shows that client certificate information is synchronized with FortiGate.

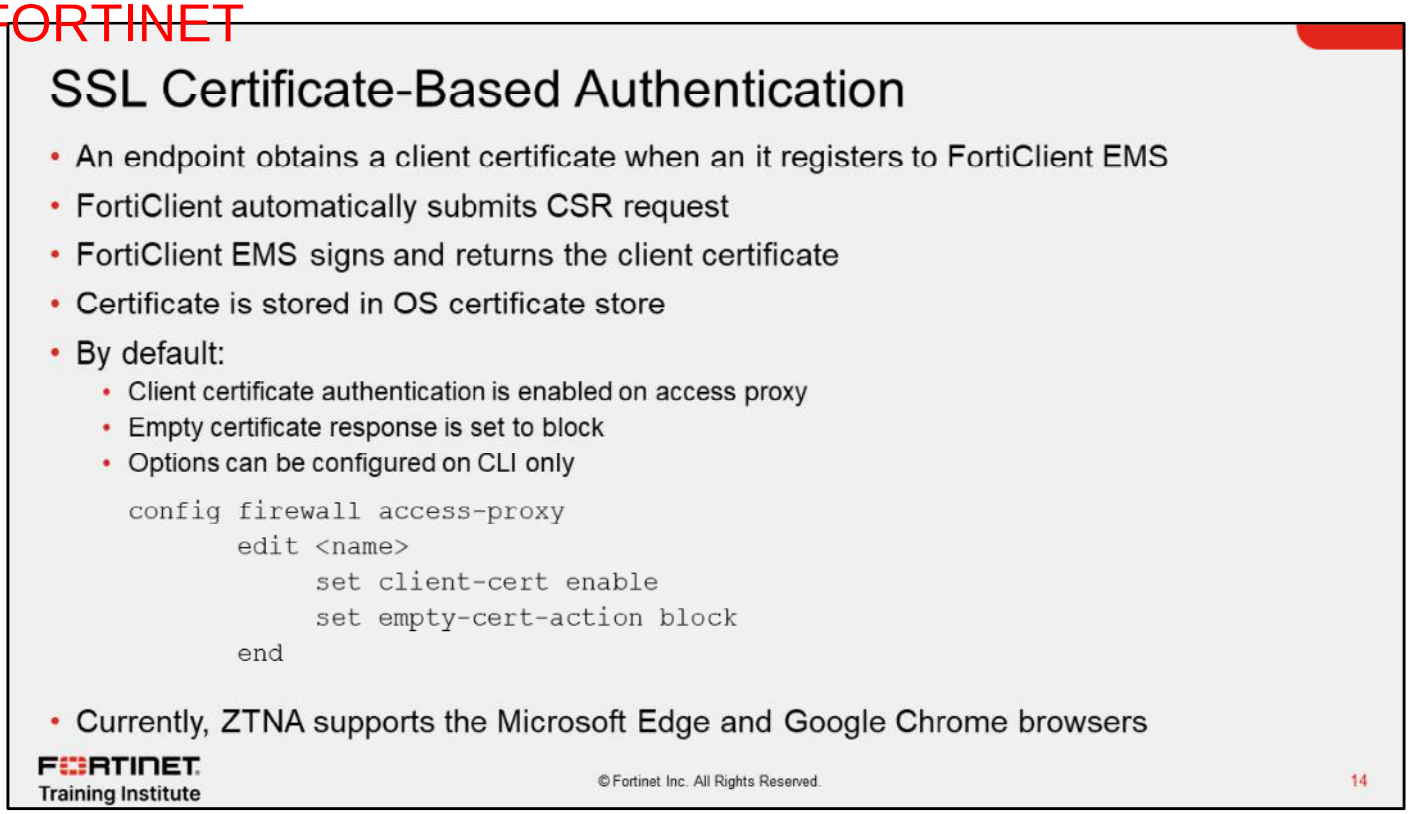

Endpoint obtains a client certificate when it registers to FortiClient EMS. FortiClient automatically submits a CSR request and the FortiClient EMS signs and returns the client certificate. This certificate is stored in the operating system certificate store for subsequent connections. The endpoint information is synchronized between FortiGate and FortiClient EMS. When an endpoint disconnects or is unregistered from FortiClient EMS, its certificate is removed from the certificate store and revokes on FortiClient EMS. The endpoint obtains a certificate again when it reconnects to the FortiClient EMS.

By default, client certificate authentication is enabled on the access proxy, so when FortiGate receives the HTTPS request, the FortiGate WAD process challenges the client to identify itself with its certificate. The FortiGate makes a decision based on specific possibilities.

If the client responds with the correct certificate that the client UID and certificate SN can be extracted from:
- If the client UID and certificate SN match the record on FortiGate, the client is allowed to continue with the ZTNA proxy rule processing.
- If the client UID and certificate SN do not match the record on FortiGate, the client is blocked from further ZTNA proxy rule processing.

If the client cancels and responds with an empty client certificate, the client is allowed to continue with ZTNA proxy rule processing when you can set `empty-cert-action` to `accept`. If `empty-cert-action` is set to `block`, FortiGate blocks the client from further ZTNA proxy rule processing.

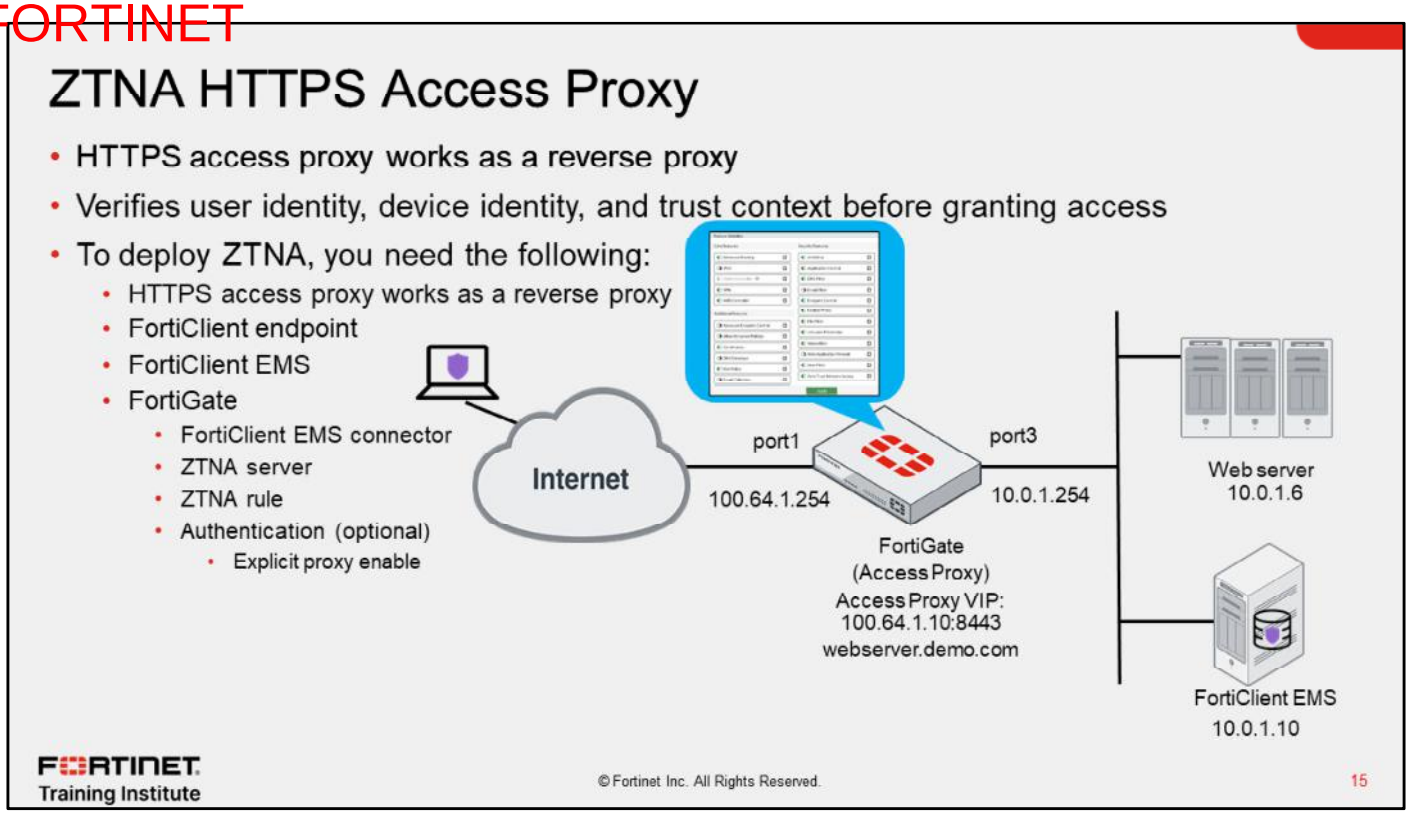

The FortiGate HTTPS access proxy works as a reverse proxy for the HTTP server. When a client connects to a web page hosted by the protected server, the address resolves to the FortiGate access proxy VIP (`100.64.1.10:8443`), as shown on this slide. FortiGate proxies the connection and takes steps to authenticate the device. It prompts the user for the endpoint certificate on the browser, and verifies this against the ZTNA endpoint record that is synchronized from the FortiClient EMS.

To enable ZTNA on the GUI, you must click **System** > **Feature Visibility**, and then enabling **Zero Trust Network Access**.

ZTNA configuration on FortiGate requires the following configuration:

- FortiClient EMS adds a fabric connector in the Security Fabric. FortiGate maintains a continuous connection to the EMS server to synchronize endpoint device information, and also automatically synchronizes ZTNA tags. You can create groups and add tags to use in the ZTNA rules and firewall policies.
- The ZTNA server defines the access proxy VIP and the real servers that clients connect to. You can also enable authentication.
- A ZTNA rule is a proxy policy used to enforce access control. You can define ZTNA tags or tag groups to enforce zero-trust role-based access. You can configure security profiles to protect this traffic.

You can also configure authentication to the access proxy. ZTNA supports basic HTTP and SAML methods.

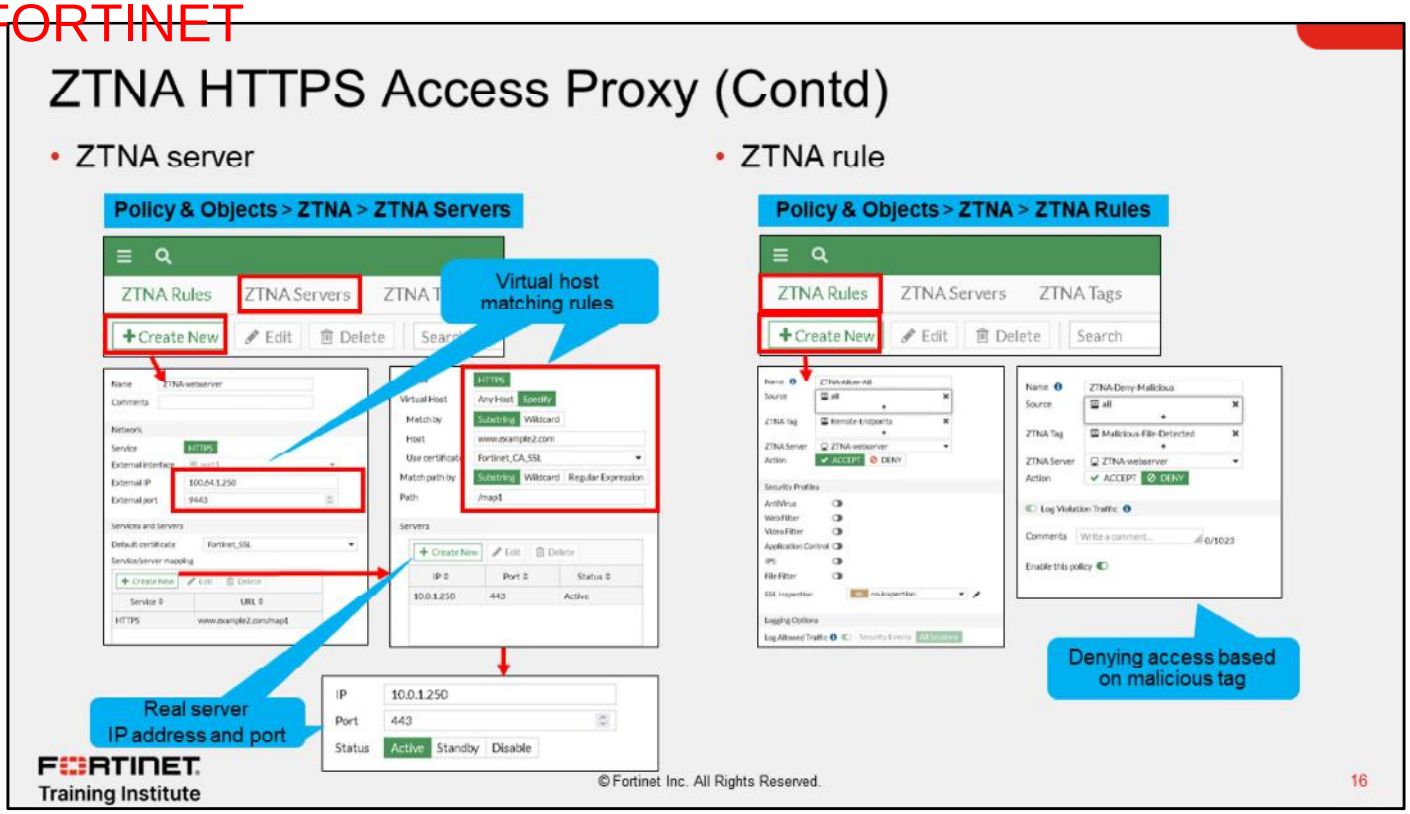

After you configure FortiClient EMS as the fabric connector and you sync ZTNA tags with FortiGate, you must create a ZTNA server or access proxy. The access proxy VIP is the FortiGate ZTNA gateway that clients make HTTPS connections to. The service and server mappings define the virtual host matching rules and the real server mappings of the HTTPS requests.

A ZTNA rule is a proxy policy used to enforce access control. You can define ZTNA tags or tag groups to enforce zero-trust role-based access. To create a rule, type a rule name, and add IP addresses and ZTNA tags or tag groups that are allowed or blocked access. You also select the ZTNA server as the destination. You can also apply security profiles to protect this traffic.

Note that UTM processing of the traffic happens at the ZTNA rule.

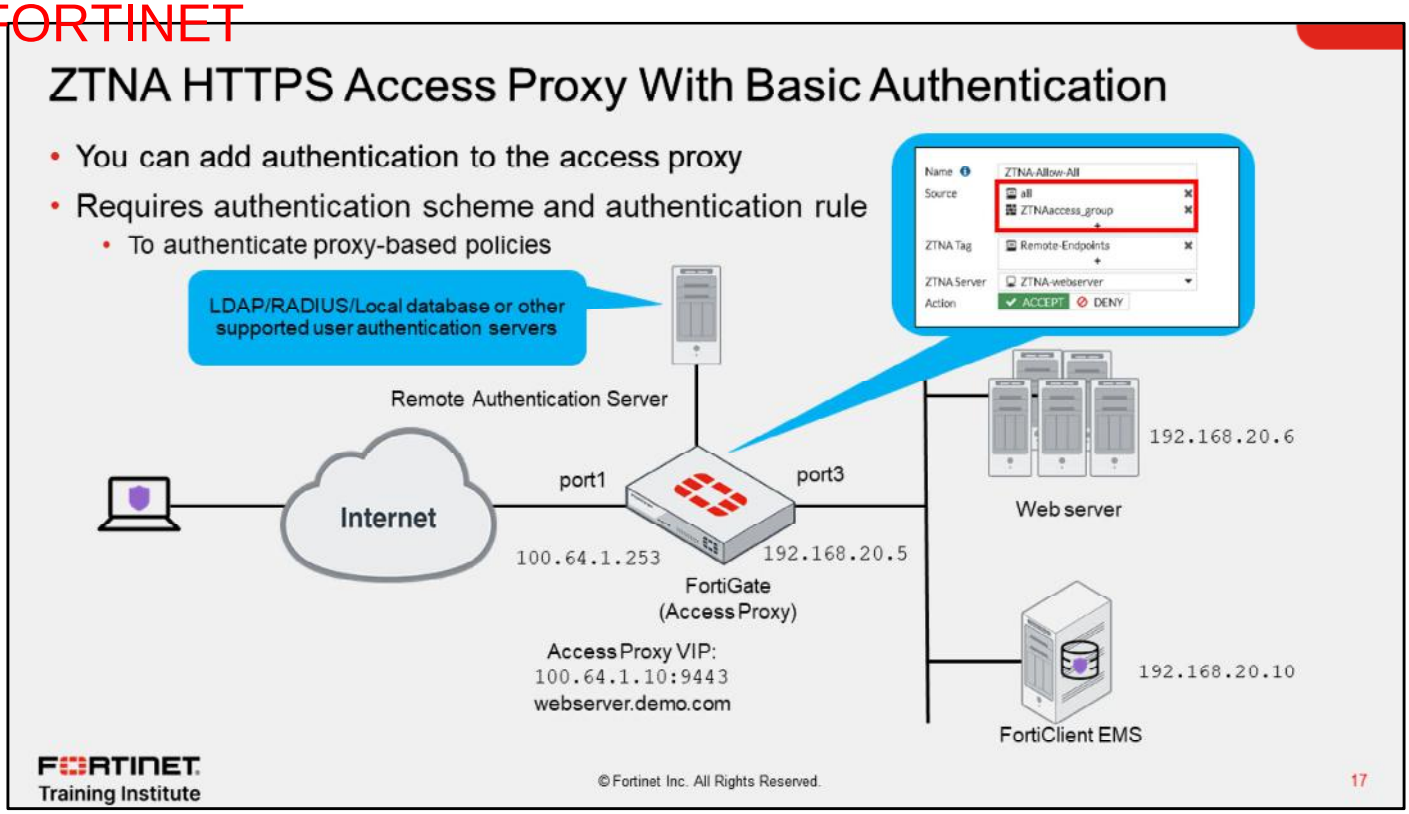

You can add authentication to the access proxy, which requires you to configure an authentication scheme and authentication rule on the FortiGate CLI. You use authentication schemes and authentication rules to authenticate proxy-based policies, similar to configuring authentication for explicit and transparent proxy.

The authentication scheme defines the method of authentication that is applied. ZTNA supports basic HTTP and SAML methods. Each method has additional settings to define the data source. For example, with basic HTTP authentication, a user database can reference an LDAP server, RADIUS server, local database, or other supported authentication servers that the user is authenticated against.

The authentication rule defines the proxy sources and destinations that require authentication, and which authentication scheme to apply. ZTNA supports the active authentication method. The active authentication method references a scheme where users are actively prompted for authentication, as they are with basic authentication. After the authentication rule triggers the method to authenticate the user, a successful authentication returns the groups that the user belongs to.

In the ZTNA rule and proxy policy, you can define a user or user group as the allowed source. Only users that match that user or group are allowed through the proxy policy. This slide shows the ZTNA rule example that user group **ZTNAaccess_group** was added to the authentication configuration after the authentication scheme and authentication rule were added to FortiGate.

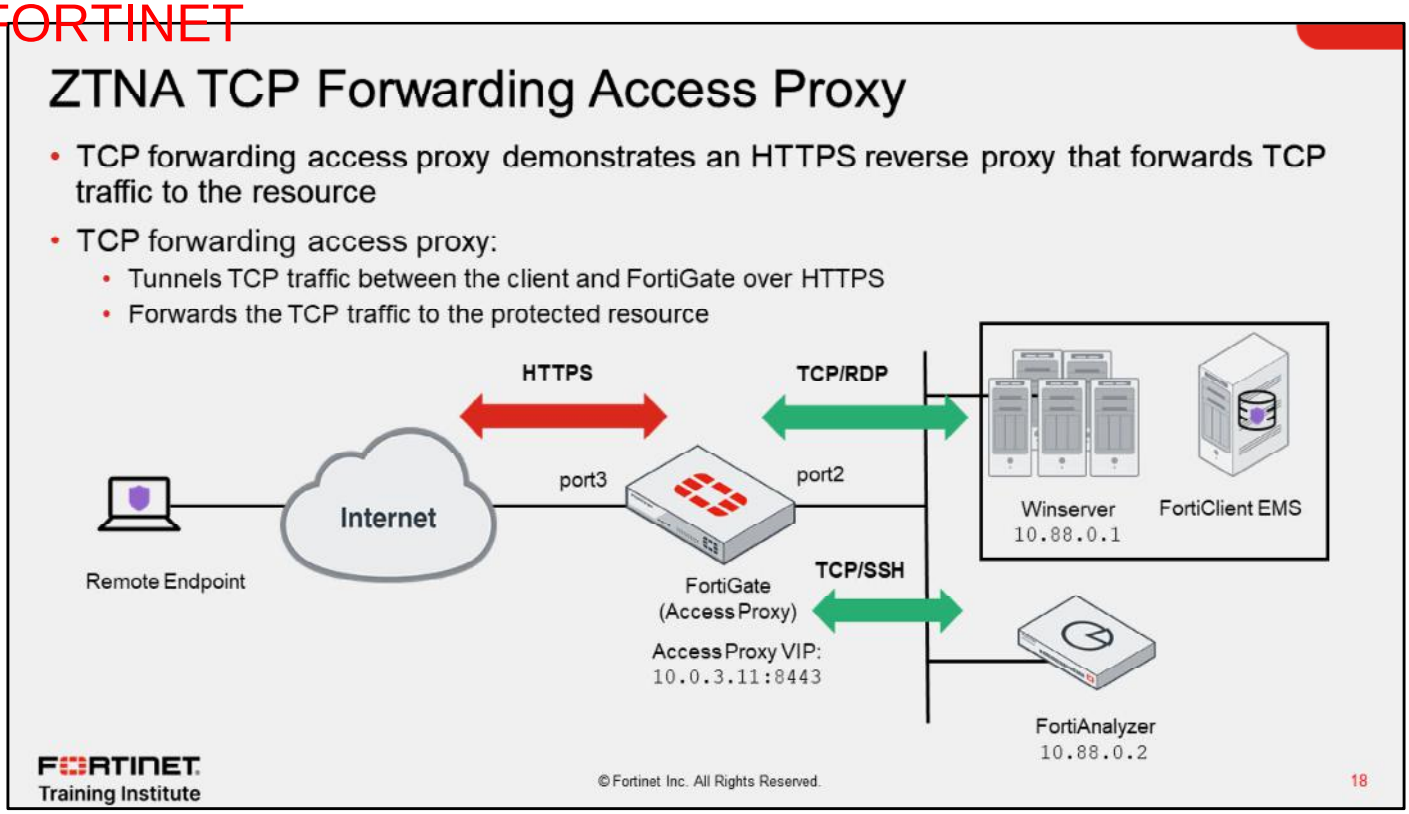

In the example shown on this slide, a TCP forwarding access proxy (TFAP) is configured to demonstrate an HTTPS reverse proxy that forwards TCP traffic to the designated resource. The access proxy tunnels TCP traffic between the client and FortiGate over HTTPS, and forwards the TCP traffic to the protected resource. It verifies user identity, device identity, and trust context, before granting access to the protected source.

RDP access is configured to Winserver, and SSH access to FortiAnalyzer. The topology shown on this slide uses IP address `10.0.3.11` and port-8443 for the external access proxy VIP.

You can also add authentication and a security posture check for TCP Forwarding Access Proxy, which you learned about earlier in this lesson.

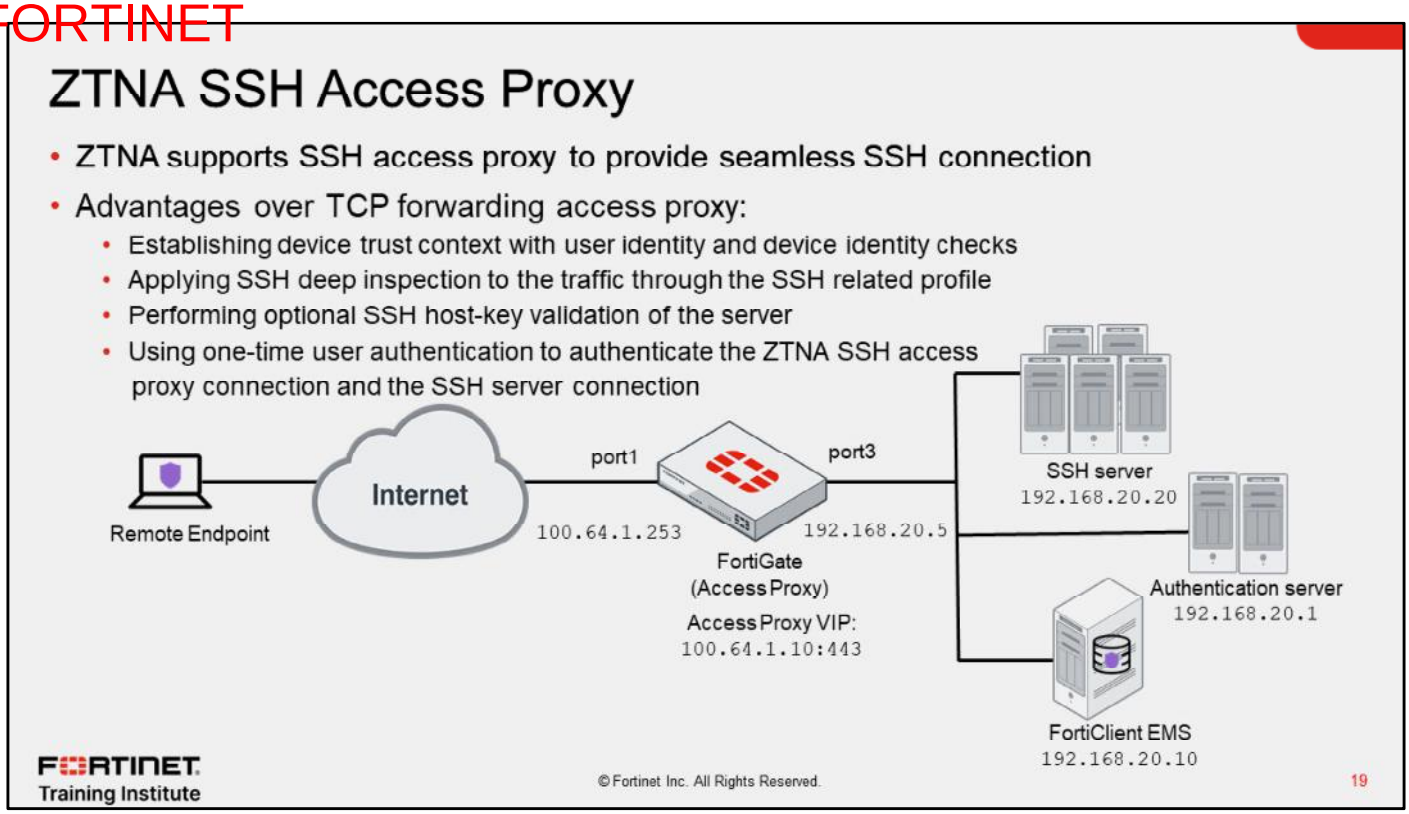

You can configure ZTNA with an SSH access proxy to provide a seamless SSH connection to the server.

Advantages of using an SSH access proxy instead of a TCP forwarding access proxy include:
- Establishing device trust context with user identity and device identity checks
- Applying SSH deep inspection to the traffic through the SSH related profile
- Performing optional SSH host-key validation of the server
- Using one-time user authentication to authenticate the ZTNA SSH access proxy connection and the SSH server connection

To act as a reverse proxy for the SSH server, FortiGate must perform SSH host-key validation to verify the identity of the SSH server. FortiGate does this by storing the public key of the SSH server in its SSH host-key configurations. When endpoint makes a connection to the SSH server, if the public key matches one that is used by the server, then the connection is established. If there is no match, then the connection fails.

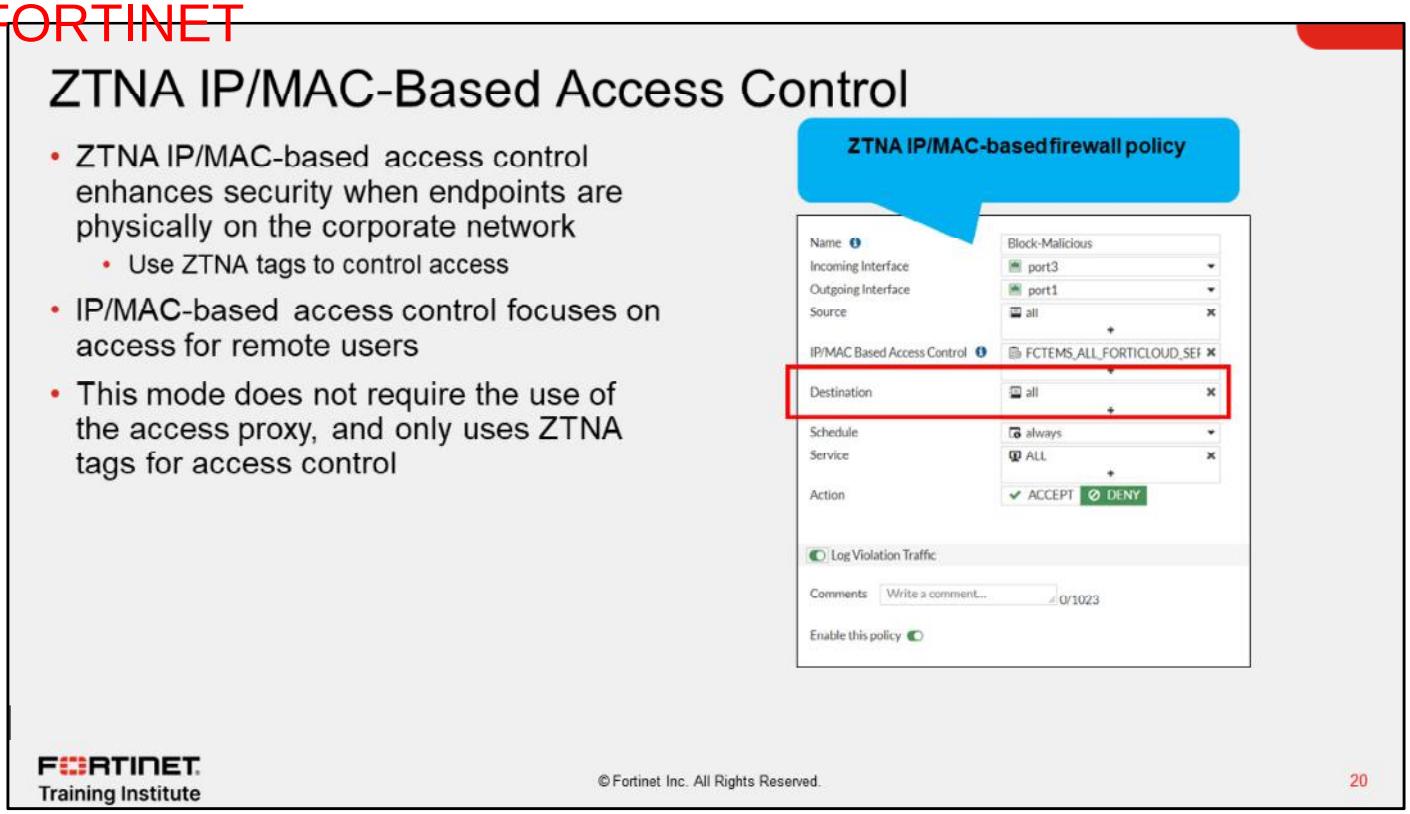

ZTNA IP/MAC-based access control enhances security when endpoints are physically located on the corporate network, whereas ZTNA access proxy focuses on access for remote users. IP/MAC-based access control combines IP/MAC with ZTNA tags for identification and security posture check, to implement role-based zero-trust access. Firewall policies are configured that use ZTNA tags to control access between on-net devices and an internal webserver. This mode does not require the use of the access proxy, and only uses ZTNA tags for access control.

The example firewall policy on this slide uses the existing tag to control access. Traffic is denied to the internet when the FortiClient endpoint is tagged with **FCTEMS_ALL_FORTICLOUD_Malicious**.

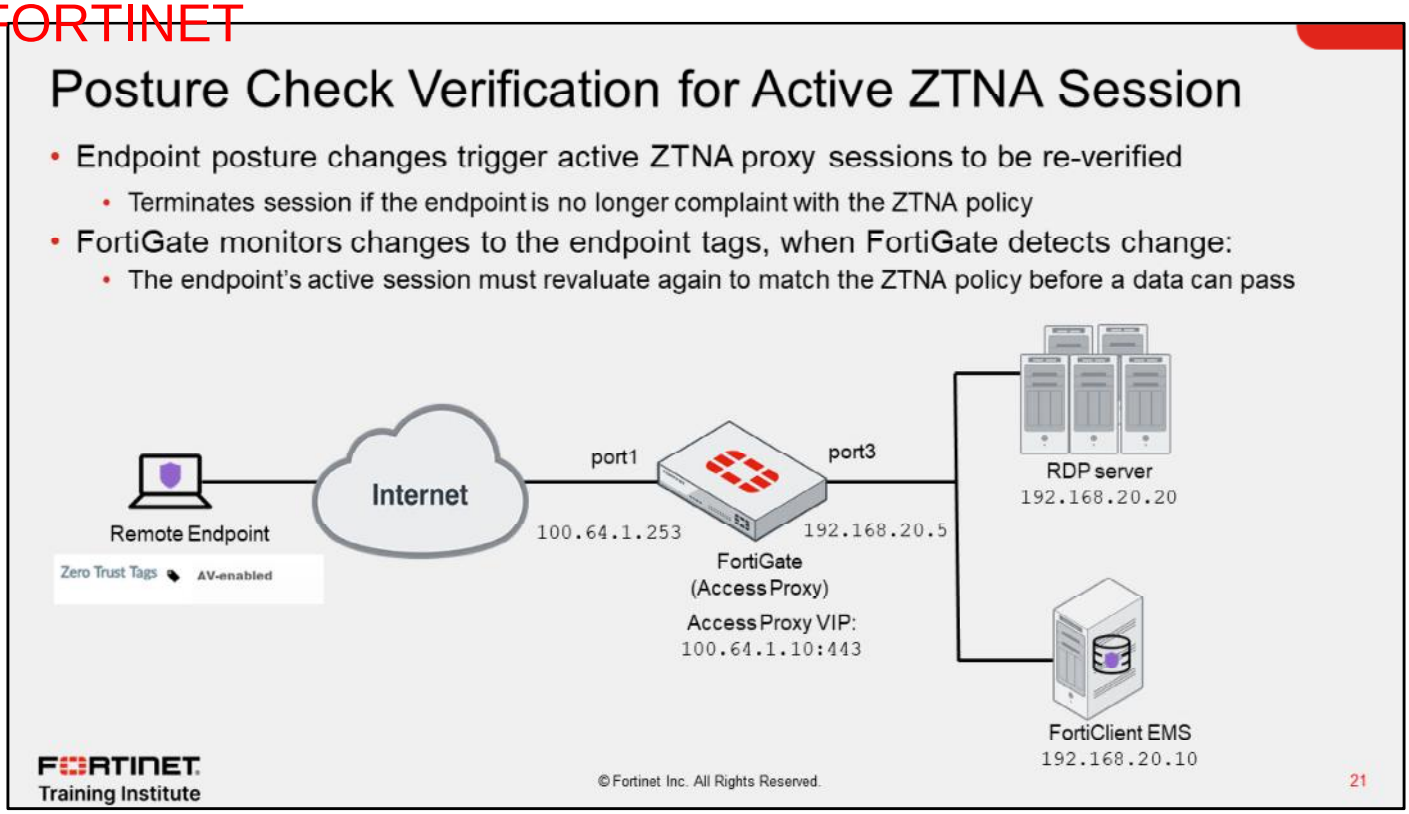

Endpoint posture changes trigger active ZTNA proxy sessions to be re-verified and terminated if the endpoint is no longer compliant with the ZTNA policy. The FortiGate monitors changes to the endpoint tags that are updated by FortiClient EMS. When a change is detected, the endpoint's active ZTNA sessions must match the ZTNA policy again before data can pass.

Note that changes to the ZTNA policy, such as changing the ZTNA tag matching logic, will also trigger re-verification of the client device against the policy.

In the example on this slide, a ZTNA rule is configured to allow access for endpoints that have the *AV-enabled* tag. After an RDP sessions is established, Windows antivirus is disabled on the remote endpoint. The FortiGate re-verifies the session and the active RDP session is removed from the FortiGate session table, causing the RDP session to be disconnected.

Knowledge Check

1. Which component issues and signs the client certificate?
   - ✓ A. FortiClient EMS
   - B. FortiClient

2. Which internet browser supports Fortinet ZTNA?
   - A. Firefox
   - ✓ B. Chrome

FORTINET Training Institute

22

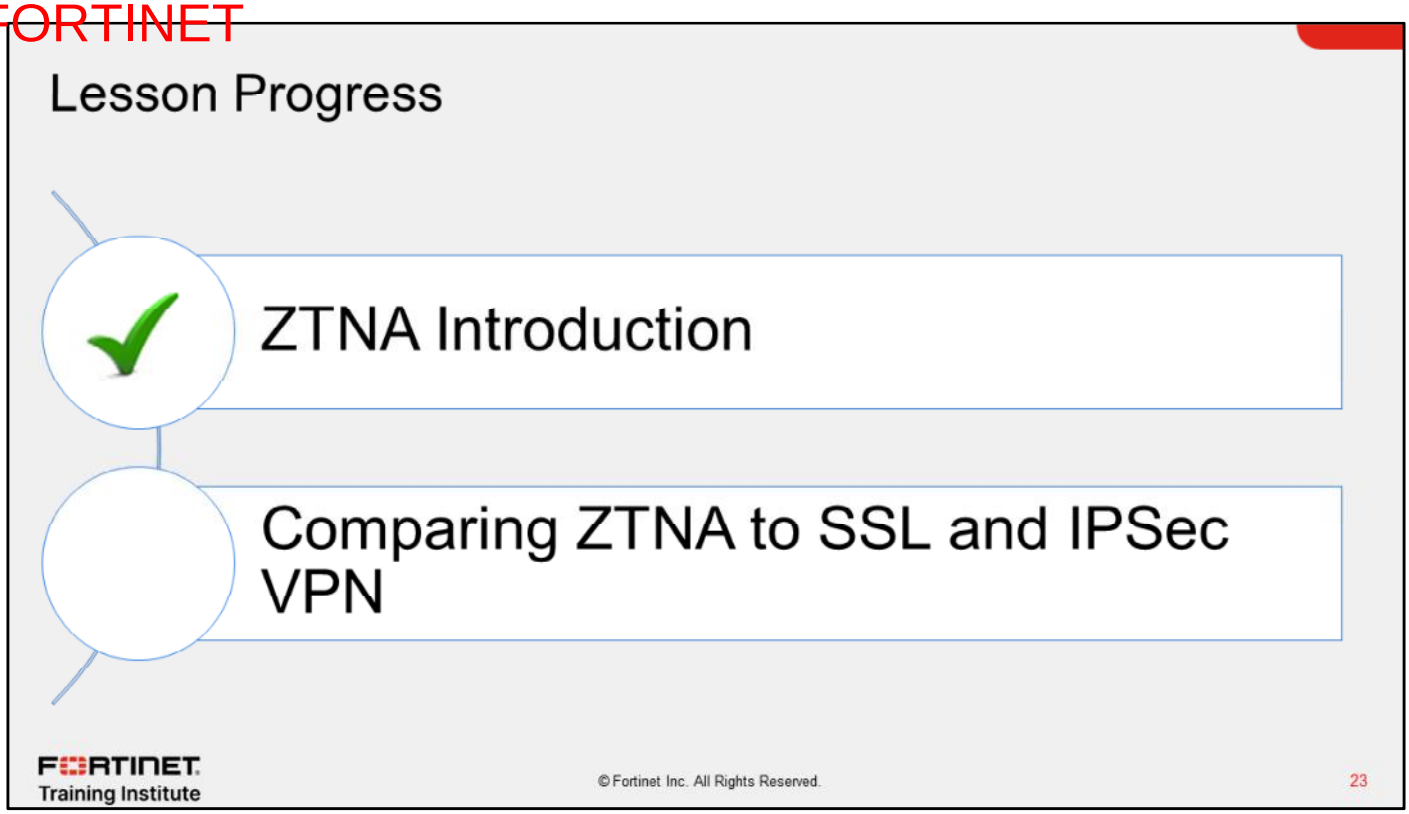

Good job! You now understand key ZTNA concepts and how to configure ZTNA

Now, you will compare ZTNA to SSL and IPsec VPN.

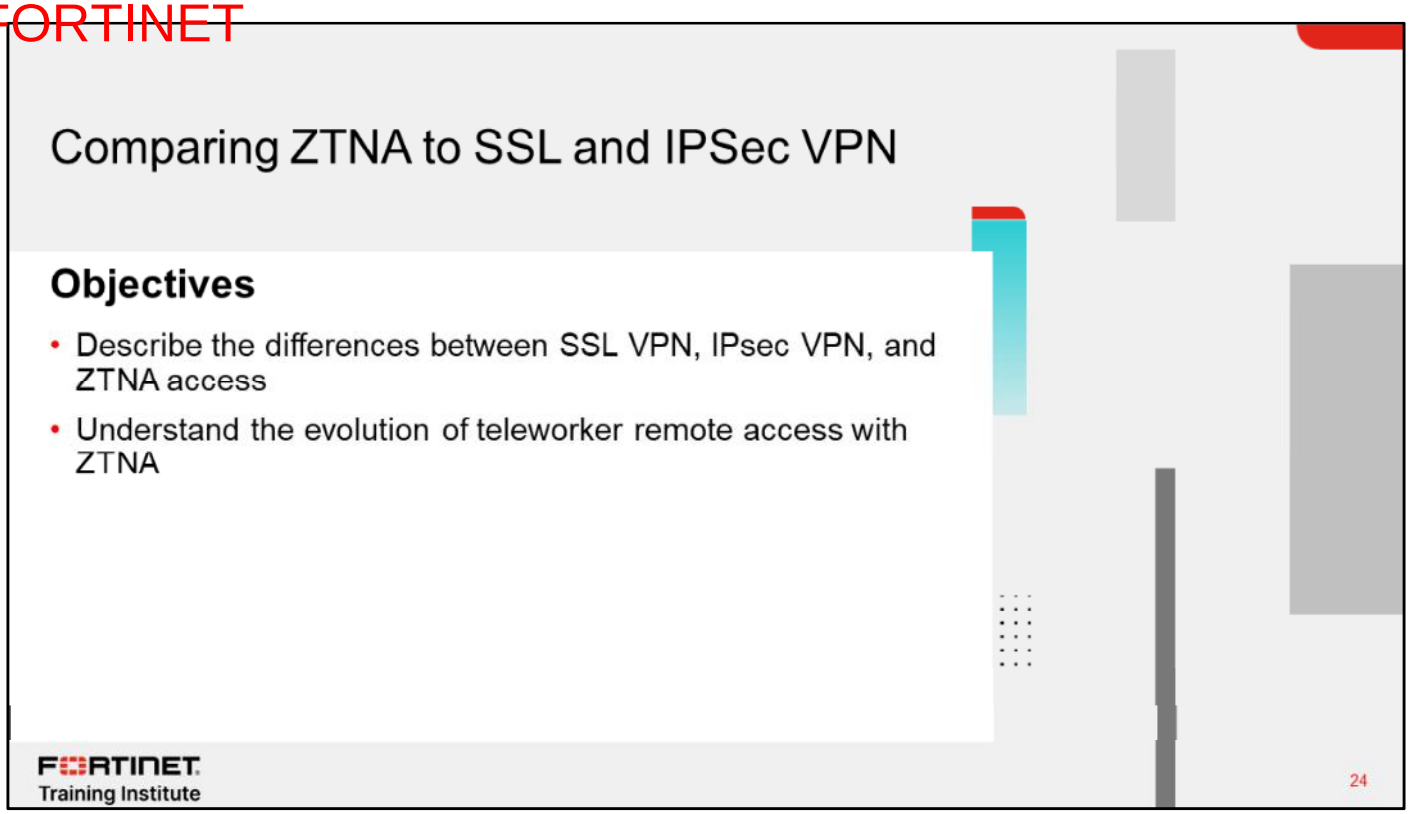

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding the evolution of remote access with ZTNA, you will be able to migrate from VPN to ZTNA HTTPS access proxy.

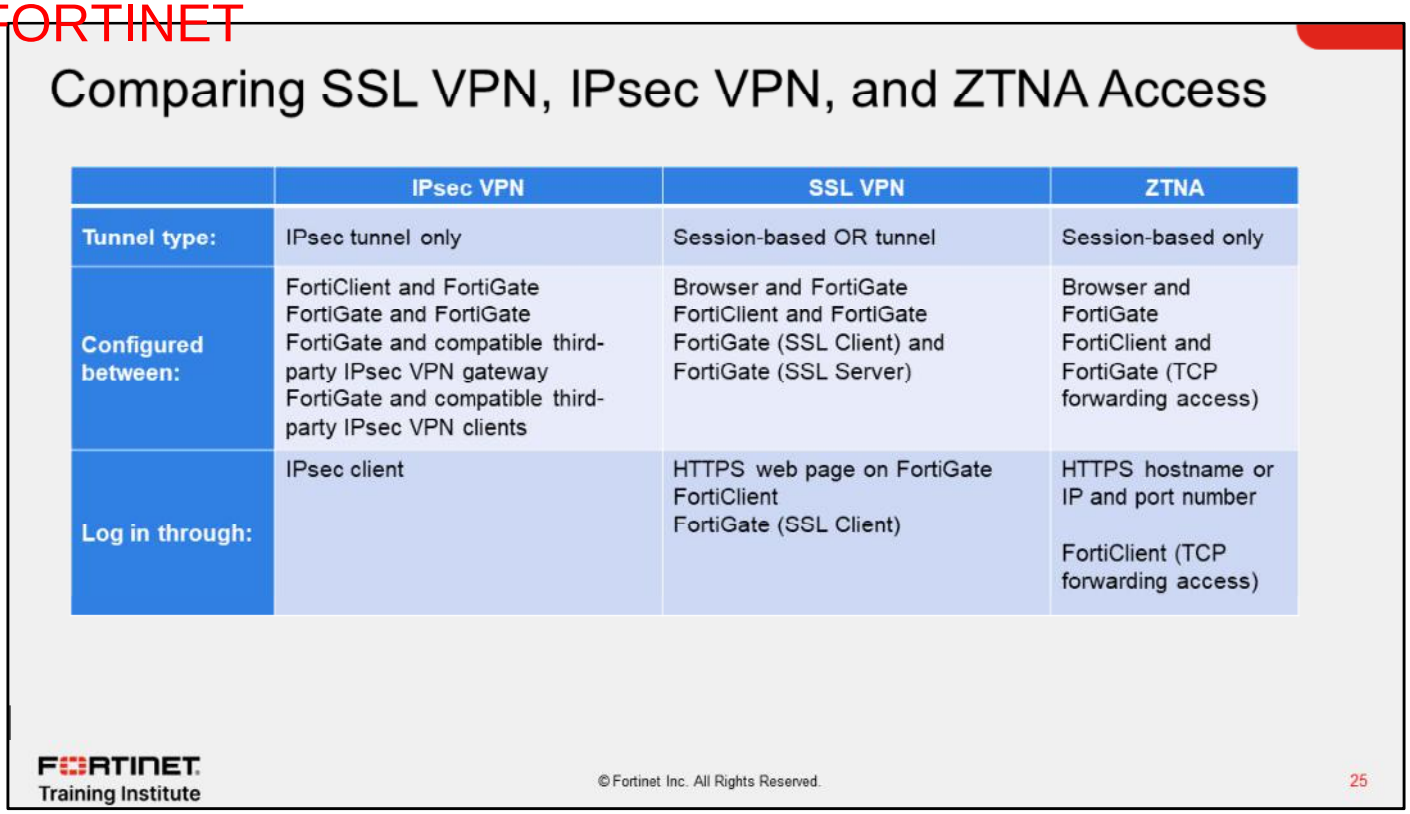

How are SSL VPN and ZTNA access different from IPsec VPNs?

SSL and TLS are commonly used to encapsulate and secure e-commerce and online banking on the internet (HTTP). SSL VPNs and ZTNA use a similar technique, and support non-HTTP protocol encapsulation as well. SSL resides higher up on the network stack than IP and, therefore, it usually requires more bits—more bandwidth—for SSL VPN headers. In comparison, IPsec uses some different methods to provide confidentiality and integrity. The primary protocol used in IPSec is ESP, which encapsulates and encrypts UDP, RDP, HTTP, or other protocols inside the IPsec tunnel.

IPSec is also an industry-standard protocol that can work with multiple vendors and supports peers that are devices and gateways—not just user clients with FortiGate only, like SSL VPN or ZTNA does.

The client software is also different. In an SSL VPN or ZTNA, your web browser might be the only client software you need. You can go to the FortiGate SSL VPN portal (an HTTPS web page) and then log in. Alternatively, you can install FortiClient or configure FortiGate as an SSL VPN client. In comparison, to use IPsec VPN, install special client software or have a local gateway, such as a desktop model FortiGate, to connect to the remote gateway. You might also need to configure firewalls between VPN peers to allow IPsec protocols.

Comparing SSL VPN, IPsec VPN, and ZTNA Access (Contd)

	IPsec VPN	SSL VPN	ZTNA
Category:	Industry standard	Vendor specific	Vendor specific

Ease of use (Configuration):	• Requires installation
	• Flexible setup
	  • Mesh and star topologies
	  • For clients or peer gateways
	  • Performance based: IPsec encryptography is faster in FortiOS
		• Does not require installation
		• Simpler setup
		  • Client-to-FortiGate
		  • FortiGate-to-FortiGate
		  • No user-configured settings
		  • Technical support less requested
			• Does not require installation
			• Simpler setup
			  • Only client-to-FortiGate
			  • No user-configured settings
			  • Technical support less requested

Better for:	Office-to-office traffic Data centers	Provides flexibility tunnel-mode or session-based access	Session-based access only

Attack surface protection	• Traditional perimeter protection:
	  • Defends against external threats only
	  • Doesn't address threat inside the network
		• Traditional perimeter protection:
		  • Defends against external threats only
		  • Doesn't address threat inside the network
			• zero-trust philosophy
			  • No one inside or outside should be trusted
			  • Based on identity authentication

FORTINET Training Institute

26

After you log in, the SSL VPN connects your computer to your private network. No user-configured settings are required, and firewalls are typically configured to allow outgoing HTTP, so technical support calls are less likely. Simplicity makes ZTNA and SSL VPN ideal for non-technical users, or users who connect from public computers, such as those found in public libraries and internet cafés. ZTNA takes this a step further and makes it easier for administrators to perform device compliance checks and configuration. ZTNA also provides an additional authentication mechanism for access control without any interaction required from the end user.

ZTNA follows the zero-trust philosophy to protect the attack surface that states no one inside or outside the network should be trusted unless their identification has been thoroughly checked. zero-trust also assumes that every attempt to access the network or an application is a threat.

Both IPSec and SSL VPN are traditional perimeter-based security approach that only distrusts factors outside the existing network and fail to address threats that already exist within the network.

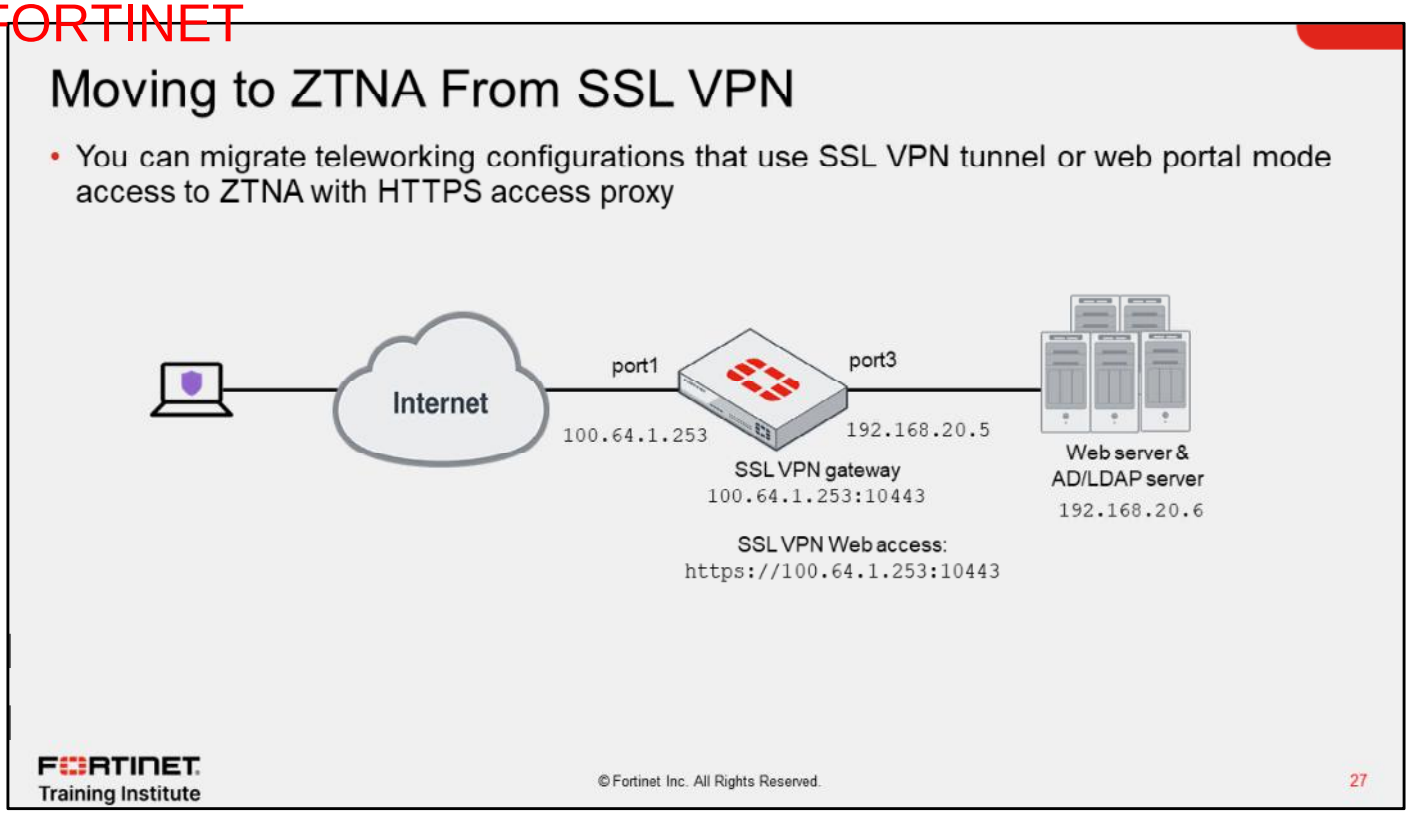

You can use ZTNA to replace VPN-based teleworking solutions. The example on this slide shows that you can migrate teleworking configurations that use SSL VPN tunnel or web portal mode access to ZTNA with HTTPS access proxy, and continue to use the same authentication server and groups to authenticate your remote users.

In addition, by integrating with FortiClient EMS, you can also ensure that FortiGate performs device identification is using client certificates, and checks the security posture before allowing the remote user into the website. This provides granular control over who can access the web resource using role-based access control. It also gives the user transparent access to the website using only their browser. You can even configure ZTNA IP/MAC filtering mode for on-fabric devices to provide similar access control while users are on the network.

Knowledge Check

1. Which remote access solution proxies HTTP and TCP over a secure HTTPS connection?

✓A. ZTNA
B. IPSec

2. What does FortiClient EMS integration ensure?

✓A. Device identification
B. User identification

FORTINET Training Institute

28

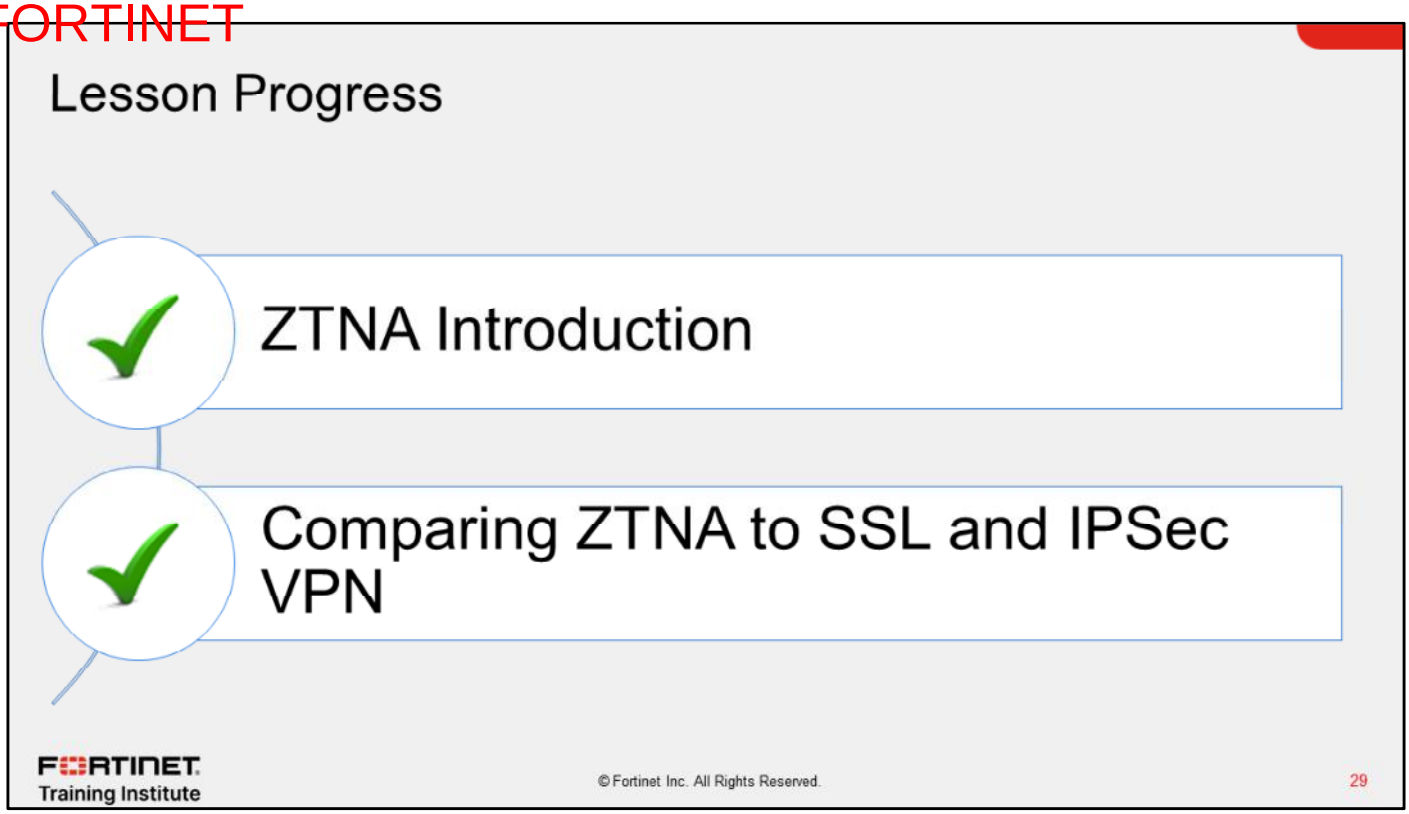

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

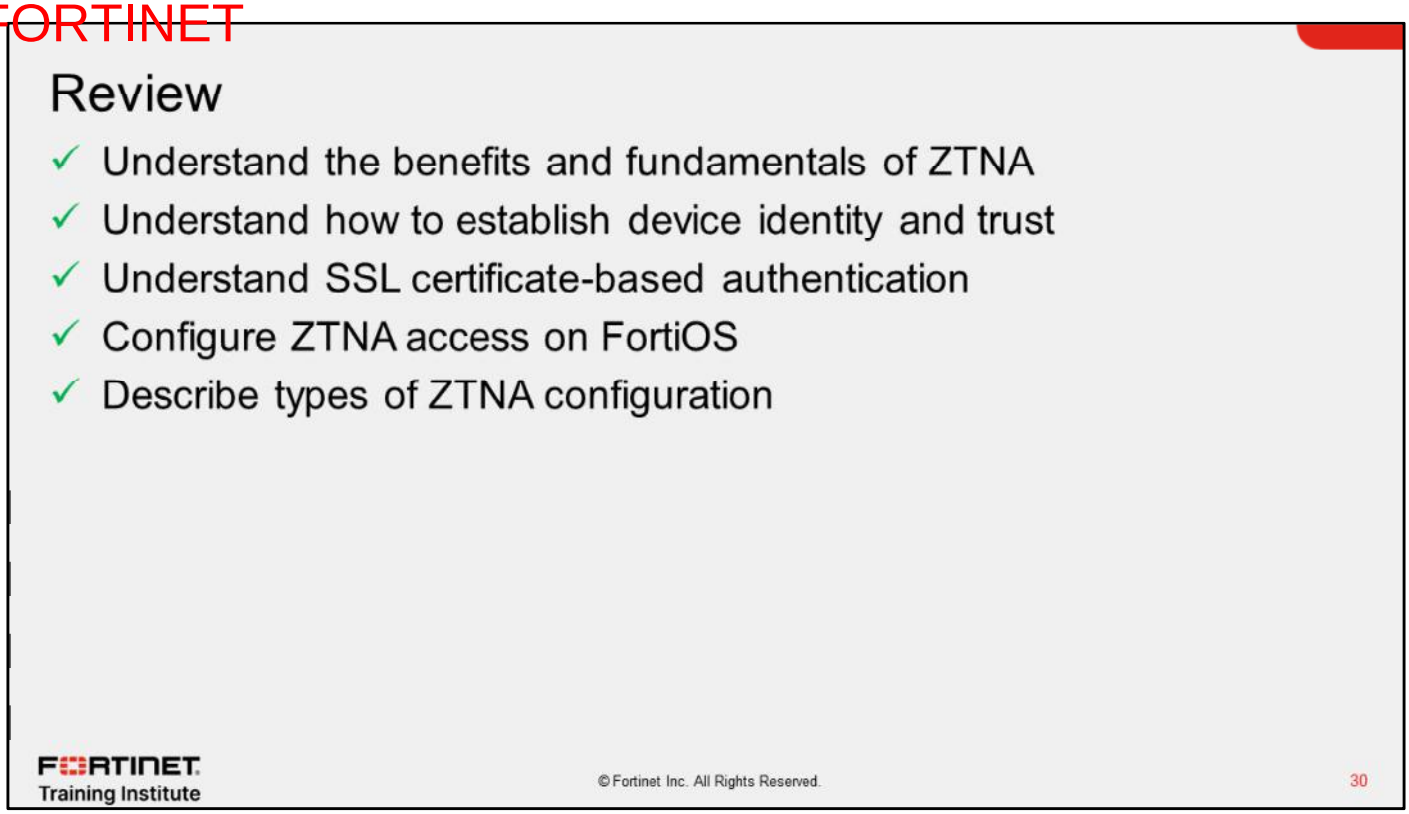

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned about how to configure and use ZTNA.

In this lesson, you will learn how to configure and use SSL VPNs. SSL VPNs are an easy way to give remote users access to your private network.

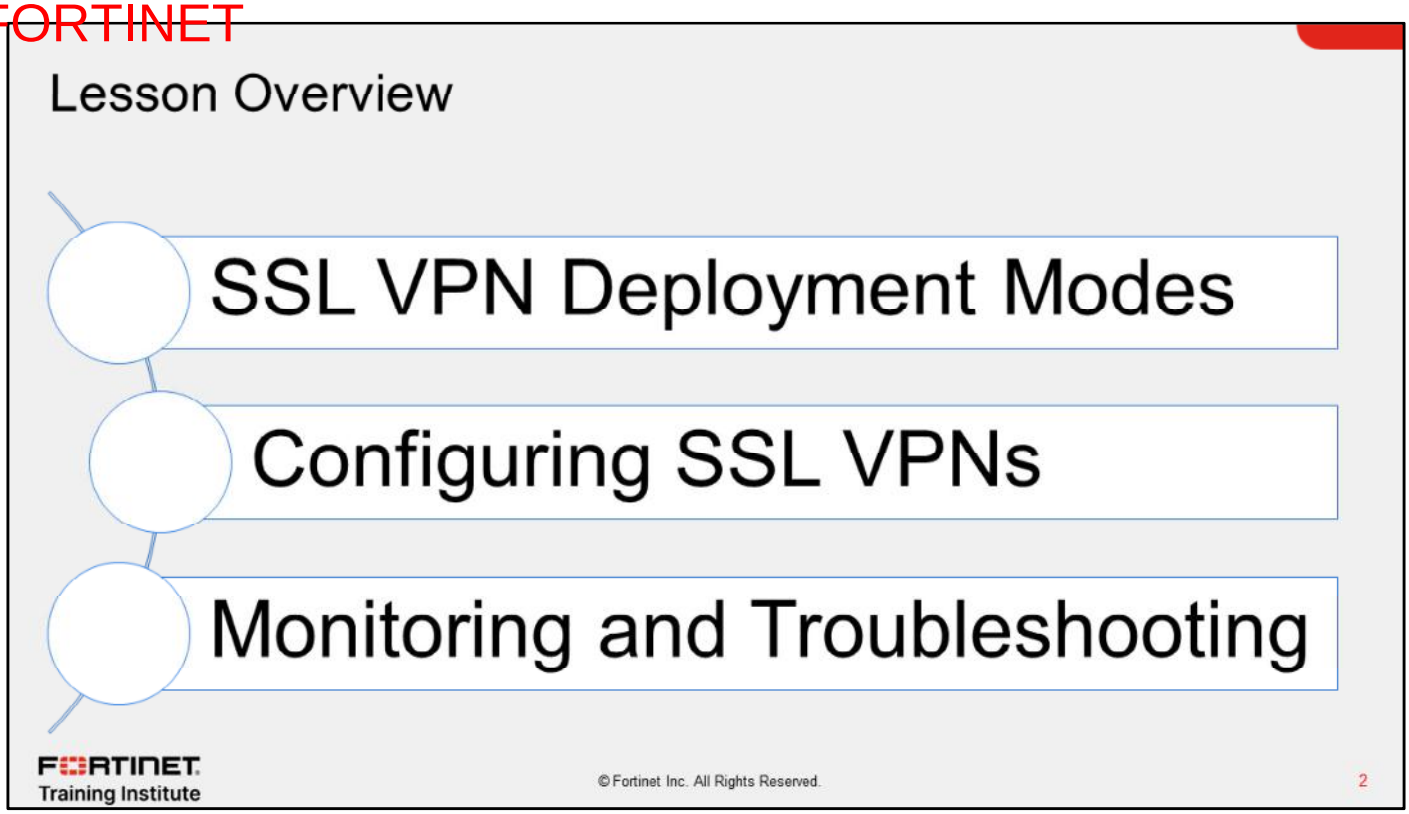

In this lesson, you will learn about the topics shown on this slide.

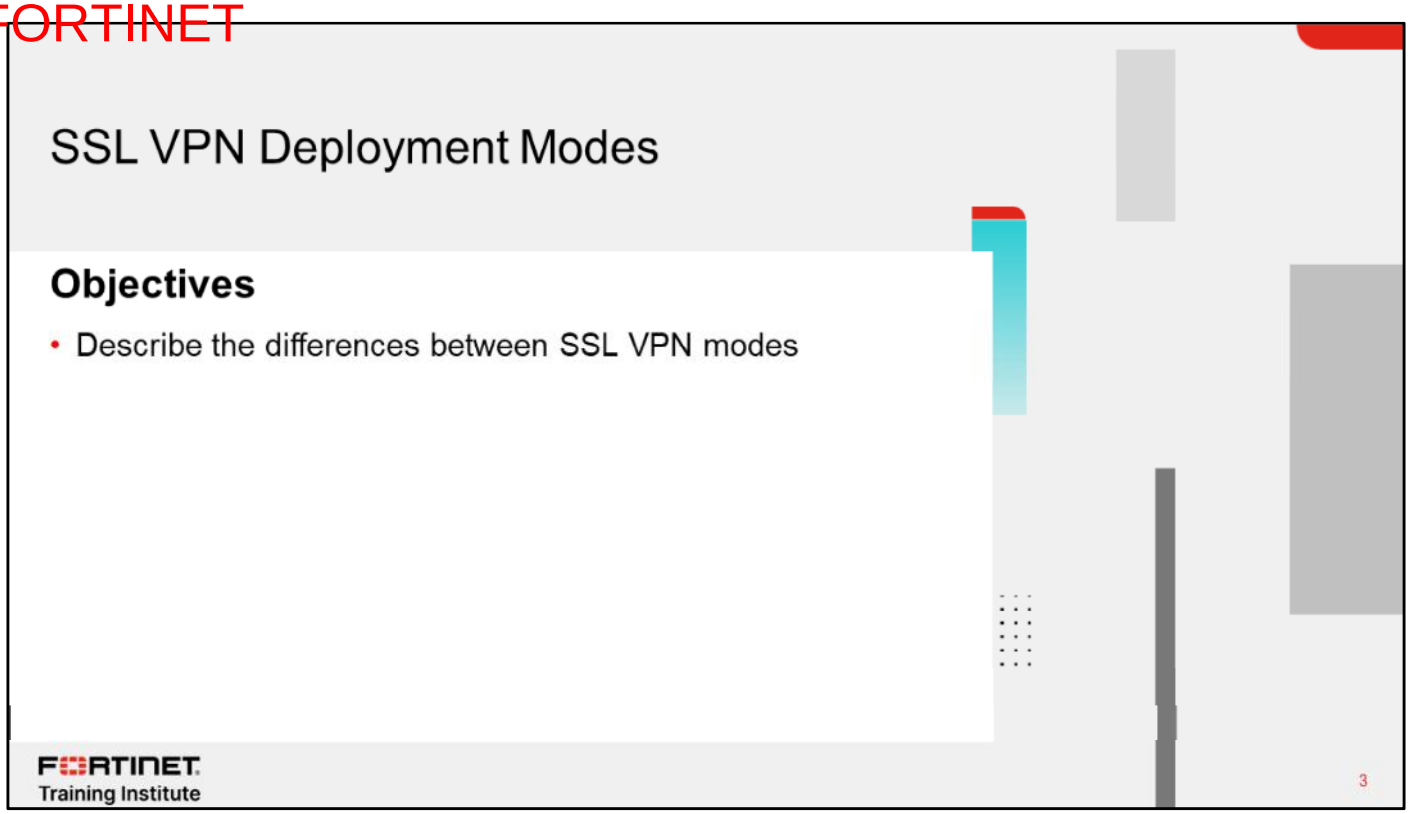

After completing this section, you should be able to achieve the objective shown on this slide. By demonstrating competence in understanding the different ways FortiGate allows SSL VPN connections, you will be able to better design the configuration of your SSL VPN.

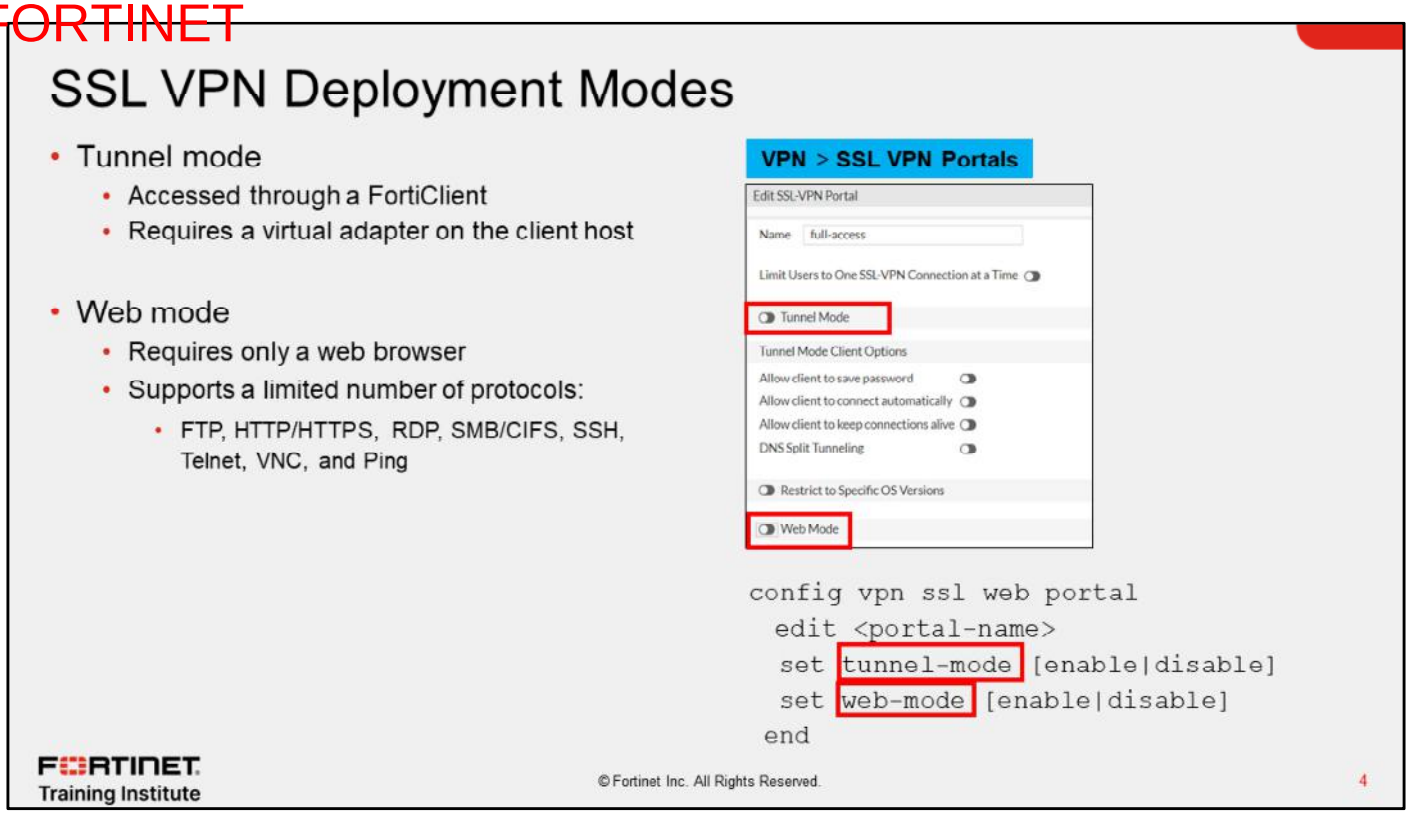

There are two modes you can use to access an SSL VPN. Both can build an SSL VPN connection, but they don't support the same features.

Which should you choose?
It depends on which applications you need to send through the VPN, the technical knowledge of your users, and whether or not you have administrative permissions on their computers.

Tunnel mode supports the most protocols, but requires the installation of a VPN client, or more specifically, a virtual network adapter. To tunnel traffic using the virtual adapter, you must use the FortiClient remote access feature or FortiClient VPN-only client.

Web mode requires only a web browser, but supports a limited number of protocols.

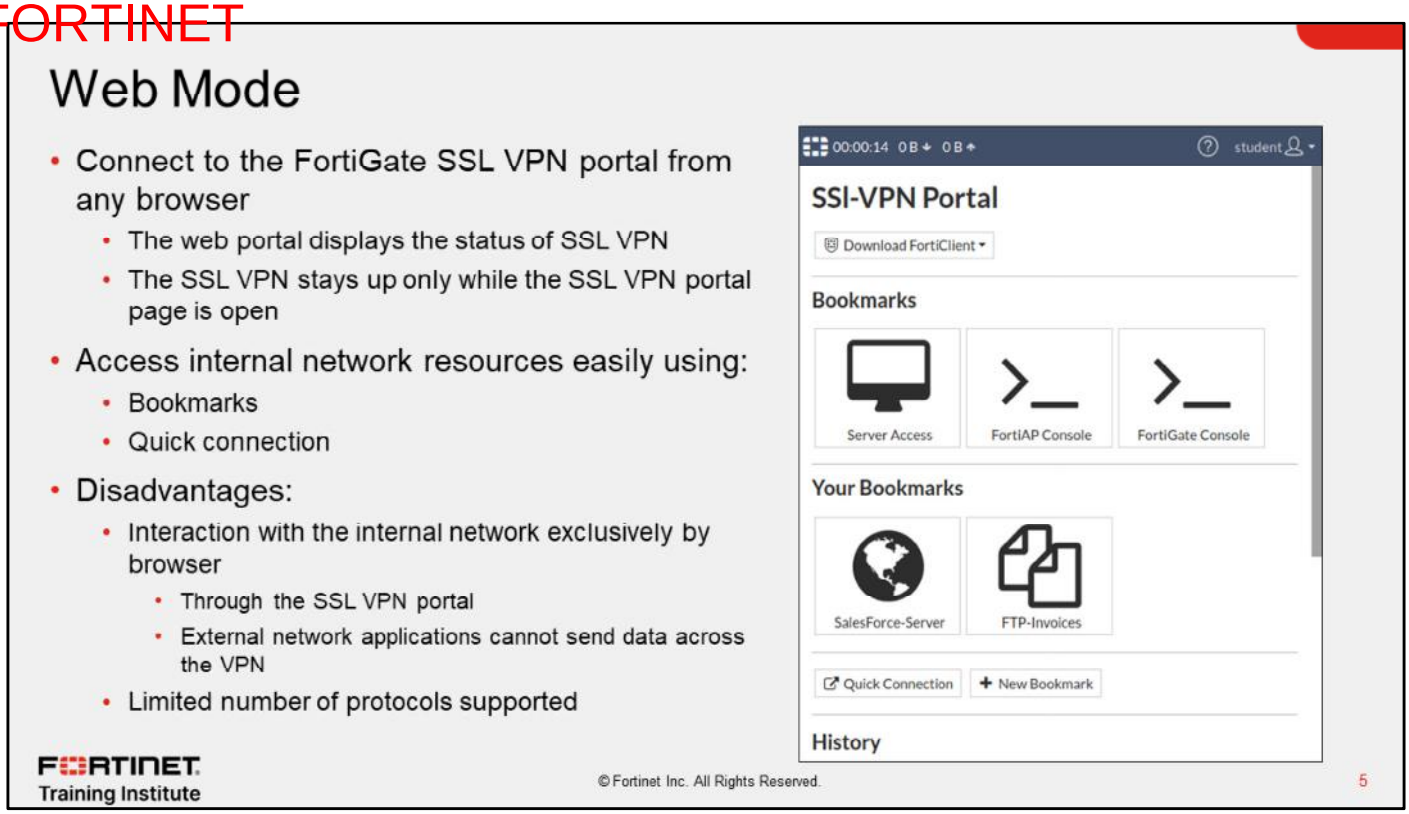

Web mode is the simplest SSL VPN mode.

Like you would with any other HTTPS website, you simply log in to the SSL VPN portal web page on FortiGate. It acts like a server-side reverse proxy, or a simple secure HTTP/HTTPS gateway, that connects you with the applications on the private network.

The **Bookmarks** section on the **SSL VPN Portal** page contains links to all or some of the resources available for the user to access. The **Quick Connection** widget allows users to type the URL or IP address of the server they want to reach. A web SSL VPN user makes use of these two widgets to access the internal network. The main advantage of web mode is that it does not usually require you to install extra software.

Web mode has two main disadvantages:

- All interaction with the internal network must be done using the browser exclusively (through the web portal). External network applications running on the user's PC cannot send data across the VPN.
- This a secure HTTP/HTTPS gateway mechanism that doesn't work for accessing everything, but just few popular protocols, such as HTTP, FTP, and Windows shares.

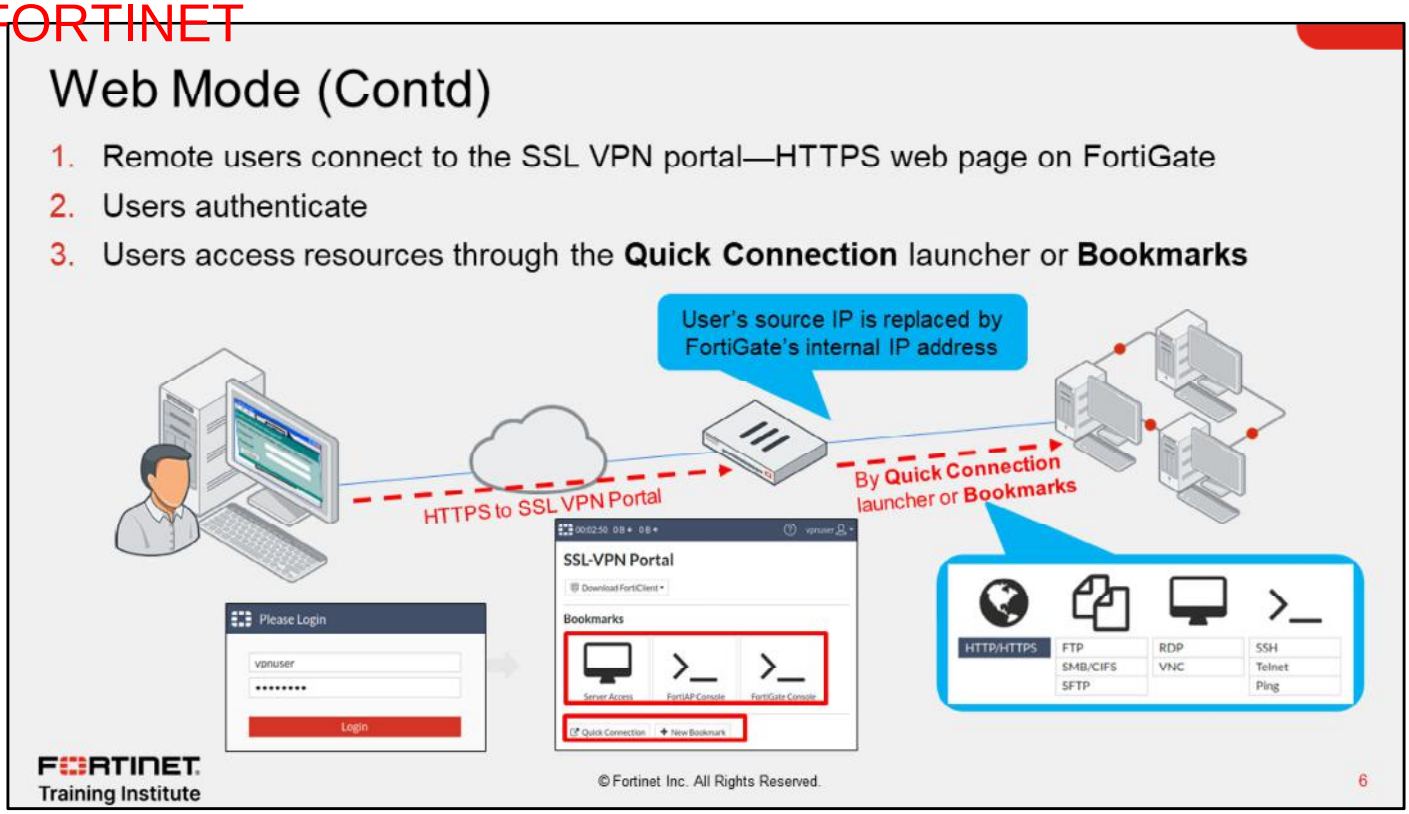

How does web mode work?

1. Remote users establish a secure connection between the SSL security in the web browser and the FortiGate SSL VPN portal, using HTTPS.
2. Once connected, users provide credentials in order to pass an authentication check.
3. Then, FortiGate displays the SSL VPN portal that contains services and network resources for users to access.

Different users can have different portals with different resources and access permissions. Also notice the source IP seen by the remote resources is FortiGate's internal IP address and not the user's IP address.

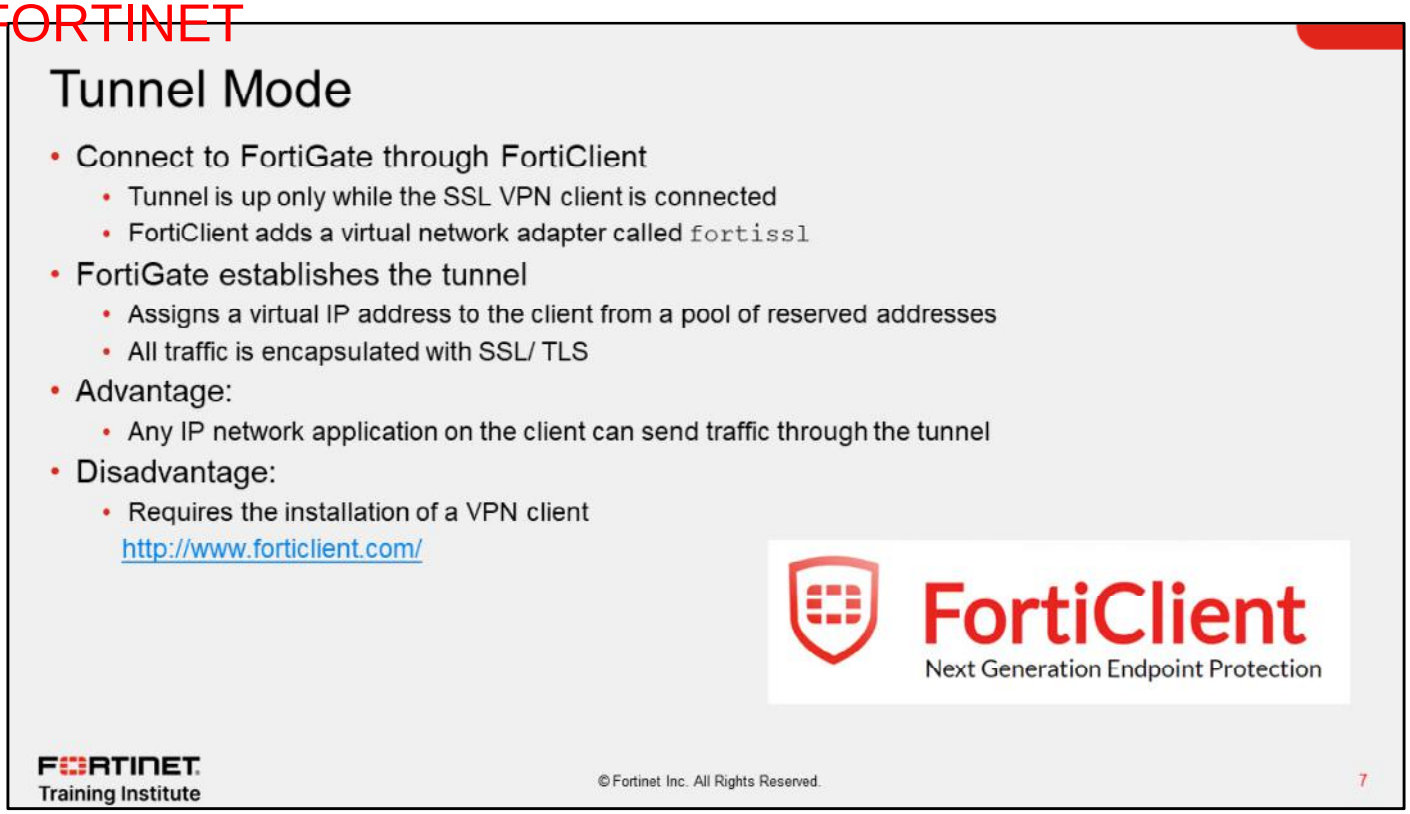

Tunnel mode is the second option FortiGate provides to access resources within an SSL VPN.

Tunnel mode requires FortiClient to connect to FortiGate. FortiClient adds a virtual network adapter identified as `fortissl` to the user's PC. This virtual adapter dynamically receives an IP address from FortiGate each time FortiGate establishes a new VPN connection. Inside the tunnel, all traffic is SSL/TLS encapsulated.

The main advantage of tunnel mode over web mode is that after the VPN is established, any IP network application running on the client can send traffic through the tunnel. The main disadvantage is that tunnel mode requires the installation of a VPN software client, which requires administrative privileges.

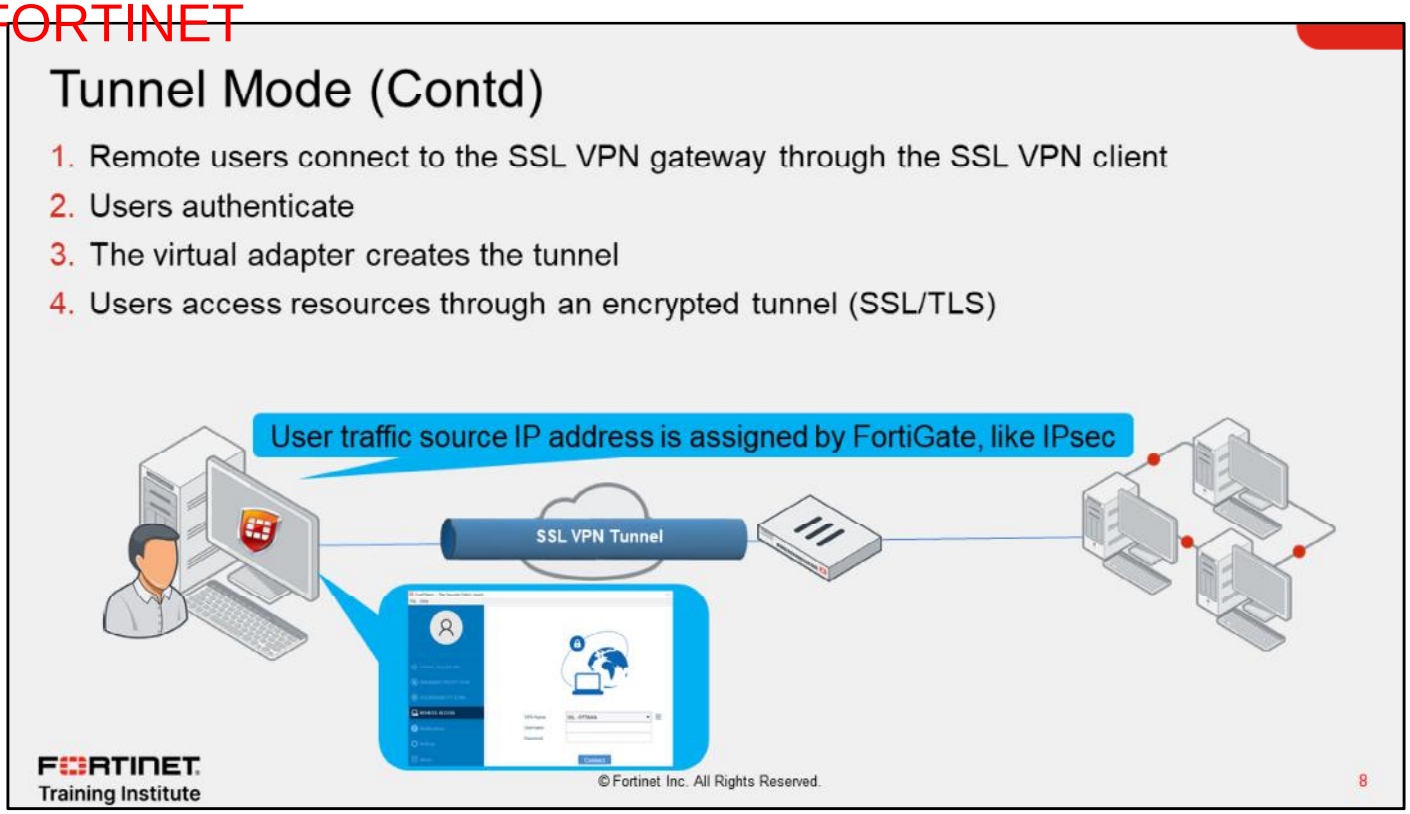

How does tunnel mode work?

1. Users connect to FortiGate through FortiClient.
2. Users provide credentials to successfully authenticate.
3. FortiGate establishes the tunnel and assigns an IP address to the client's virtual network adapter (`fortissl`). This is the client's source IP address for the duration of the connection.
4. Then, users can access services and network resources through the encrypted tunnel.

FortiClient encrypts all traffic from the remote computer and sends it over the SSL VPN tunnel. FortiGate receives the encrypted traffic, de-encapsulates the IP packets, and forwards them to the private network as if the traffic originated from inside the network.

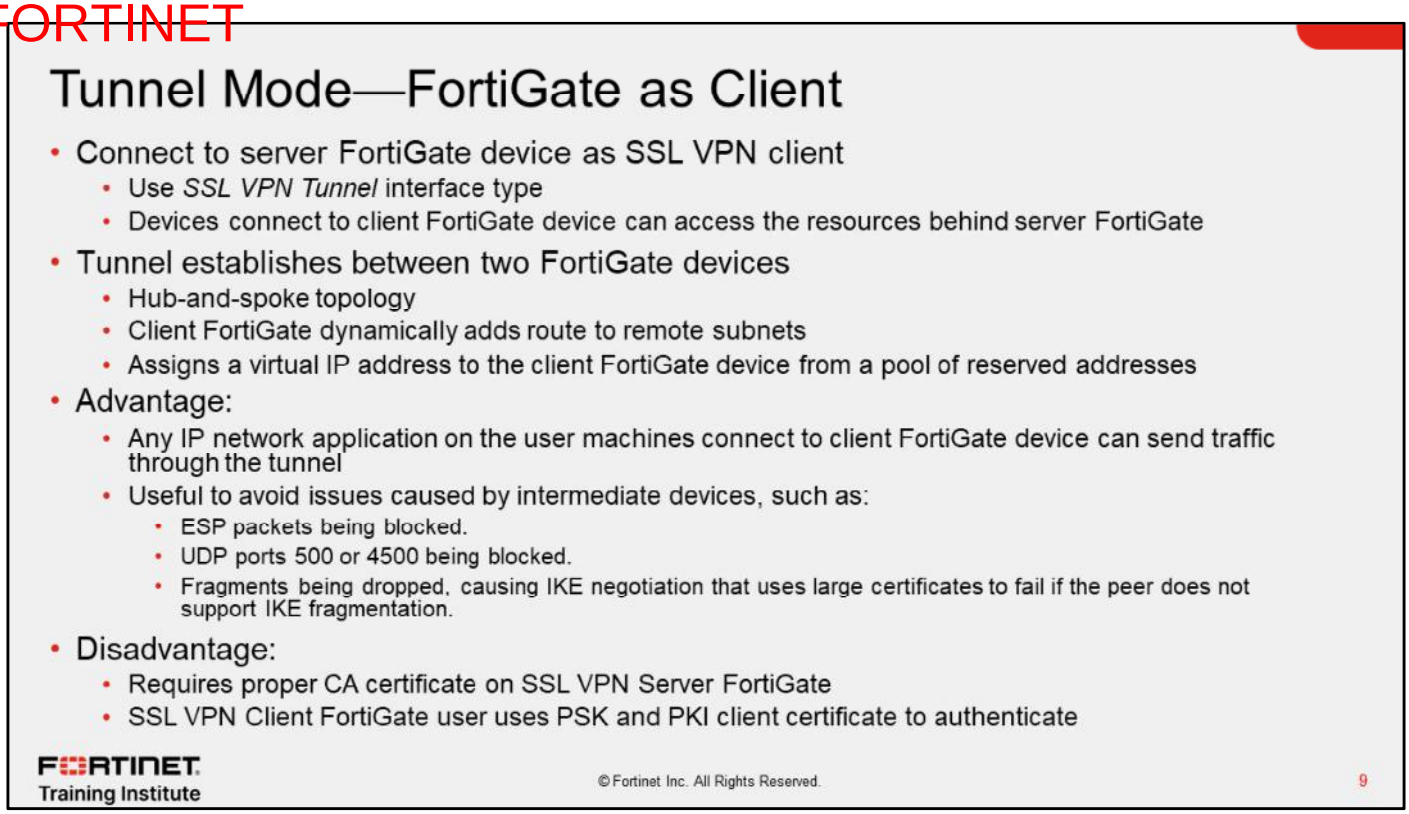

The FortiGate can be configured as an SSL VPN client, using an *SSL-VPN Tunnel* interface type. When an SSL VPN client connection is established, the client dynamically adds a route to the subnets that are returned by the SSL VPN server. Policies can be defined to allow users that are behind the client to be tunneled through SSL VPN to destinations on the SSL VPN server.

This setup provides IP-level connectivity in tunnel mode and allows hub-and-spoke topologies to be configured with FortiGates as both the SSL VPN hub and spokes. This can be useful to avoid issues caused by intermediate devices, such as:

- ESP packets being blocked.
- UDP ports 500 or 4500 being blocked.
- Fragments being dropped, causing IKE negotiation that uses large certificates to fail if the peer does not support IKE fragmentation.

If the client specified destination is *all*, a default route is effectively dynamically created on the SSL VPN client, and the new default route is added to the existing default route in the form of ECMP. You can modify route's distance or priority according to your requirements. To avoid a default route being learned on the SSL VPN client, on the SSL VPN server define a specific destination. Split tunneling is used so that only the destination addresses defined in the server's firewall policies are routed to the server, and all other traffic is connected directly to the internet.

This configuration requires proper CA certificate installation as the SSL VPN client FortiGate/user uses PSK and a PKI client certificate to authenticate. The FortiGate devices must have the proper CA certificate installed to verify the certificate chain to the root CA that signed the certificate.

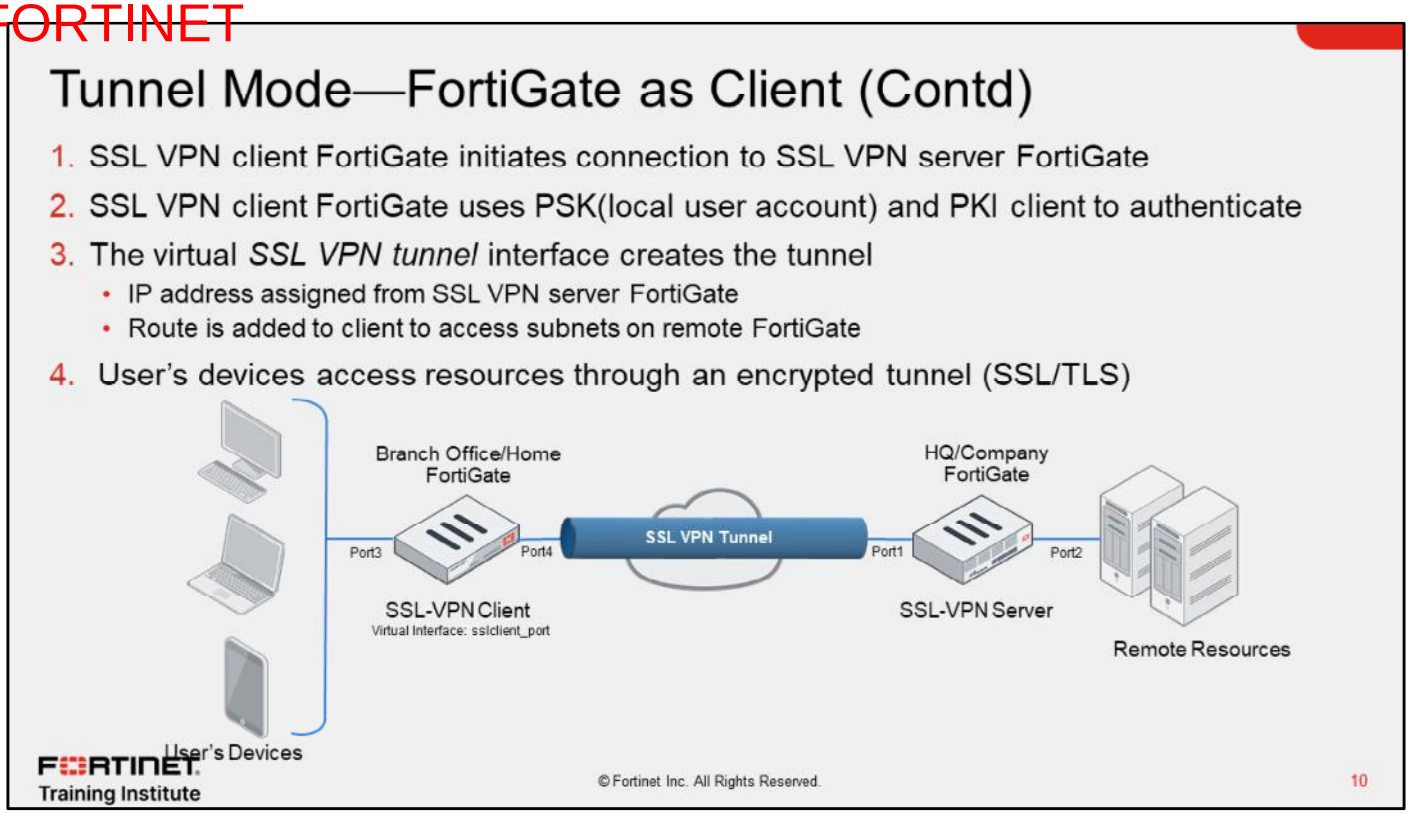

How does tunnel mode work when FortiGate is configured as client?

1. Client FortiGate connects to server FortiGate using SSL/TLS
2. Client FortiGate provides credentials to successfully authenticate. It includes both PSK (local or remote user account) and PKI (certificate) accounts.
3. Server FortiGate establishes the tunnel and assigns an IP address to the client's virtual network adapter. This is the client's source IP address for the duration of the connection.
4. Then, users can access services and network resources through the encrypted tunnel behind client FortiGate.

SSL VPN client FortiGate device encrypts all traffic from the remote computer and sends it over the SSL VPN tunnel. SSL VPN server FortiGate receives the encrypted traffic, de-encapsulates the IP packets, and forwards them to the private network as if the traffic originated from inside the network.

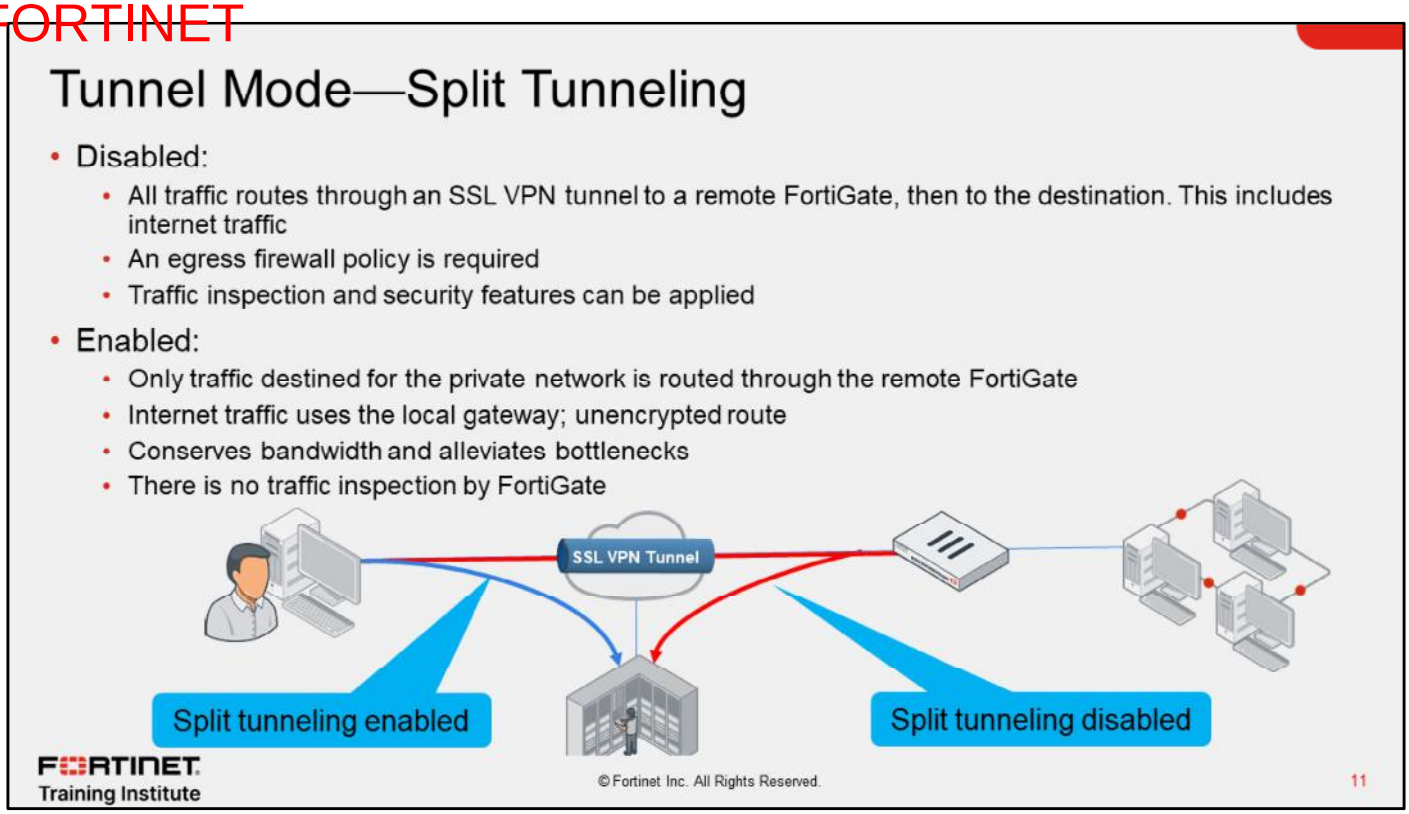

Tunnel mode also supports split tunneling.

When split tunneling is disabled, all IP traffic generated by the client's computer—including internet traffic—is routed across the SSL VPN tunnel to FortiGate. This sets up FortiGate as the default gateway for the host. You can use this method in order to apply security features to the traffic on those remote clients, or to monitor or restrict internet access. This adds more latency and increases bandwidth usage.

In a FortiGate (client) to FortiGate (server) setup, a default route is effectively dynamically created on the SSL VPN client FortiGate, and the new default route is added to the existing default route in the form of ECMP. The following options are available to configure routing:

- To make all traffic default to the SSL VPN server and still have a route to the server's listening interface, on the SSL VPN client, set a lower distance for the default route that is learned from the server.
- To include both default routes in the routing table, with the route learned from the SSL VPN server taking priority, on the SSL VPN client, set a lower distance for the route learned from the server. If the distance is already zero, then increase the priority on the default route.

When split tunneling is enabled, only traffic that is destined for the private network behind the remote FortiGate is routed through the tunnel. All other traffic is sent through the usual unencrypted route. There is no traffic inspection by FortiGate.

Split tunneling helps to conserve bandwidth and alleviates bottlenecks.

Knowledge Check

1. A web-mode SSL VPN user connects to a remote web server. What is the source IP address of the HTTP request the web server receives?
   - A. The remote user IP address
   - ✓ B. The FortiGate device internal IP address

2. Which statement about tunnel-mode SSL VPN is correct?
   - ✓ A. It supports split tunneling.
   - B. It requires bookmarks.

3. A web-mode SSL VPN user uses _____ to access internal network resources.
   - ✓ A. bookmarks
   - B. FortiClient

FORTINET Training Institute

12

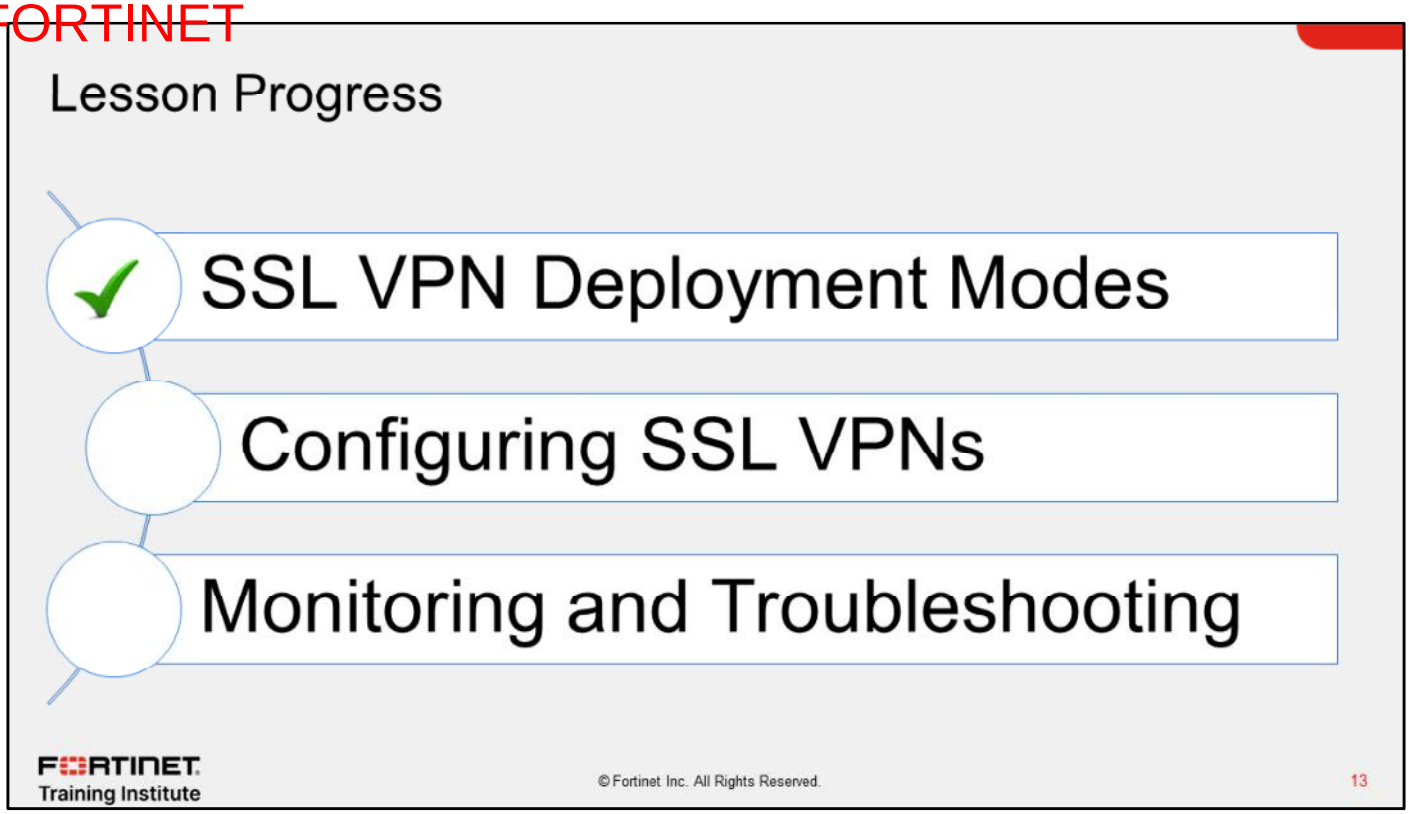

Good job! You now understand the SSL VPN operation modes supported by FortiGate.

Now, you will learn about how to configure SSL VPNs.

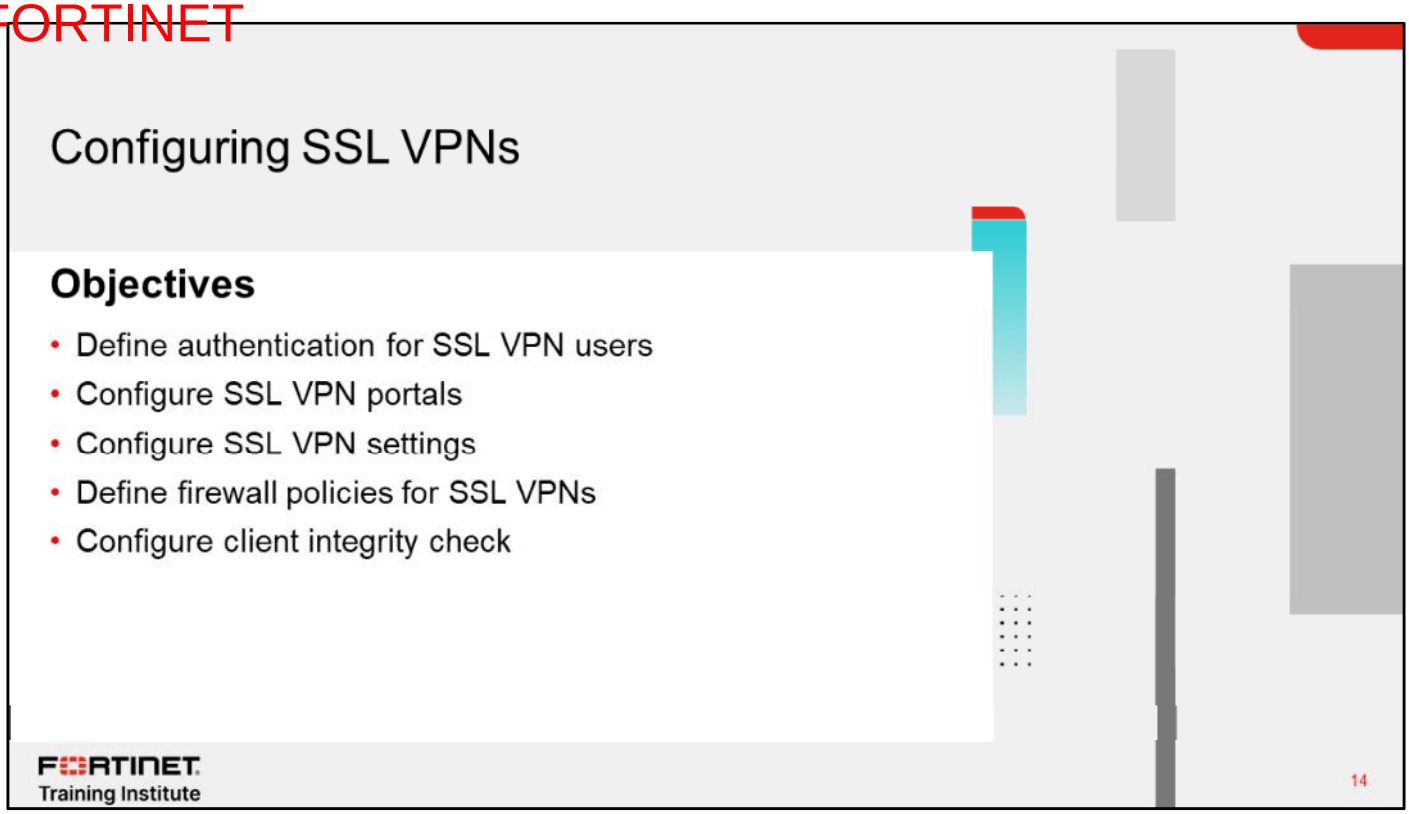

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in configuring the SSL VPN settings on FortiGate, you will be able to better design the architecture of your SSL VPN tunnels.

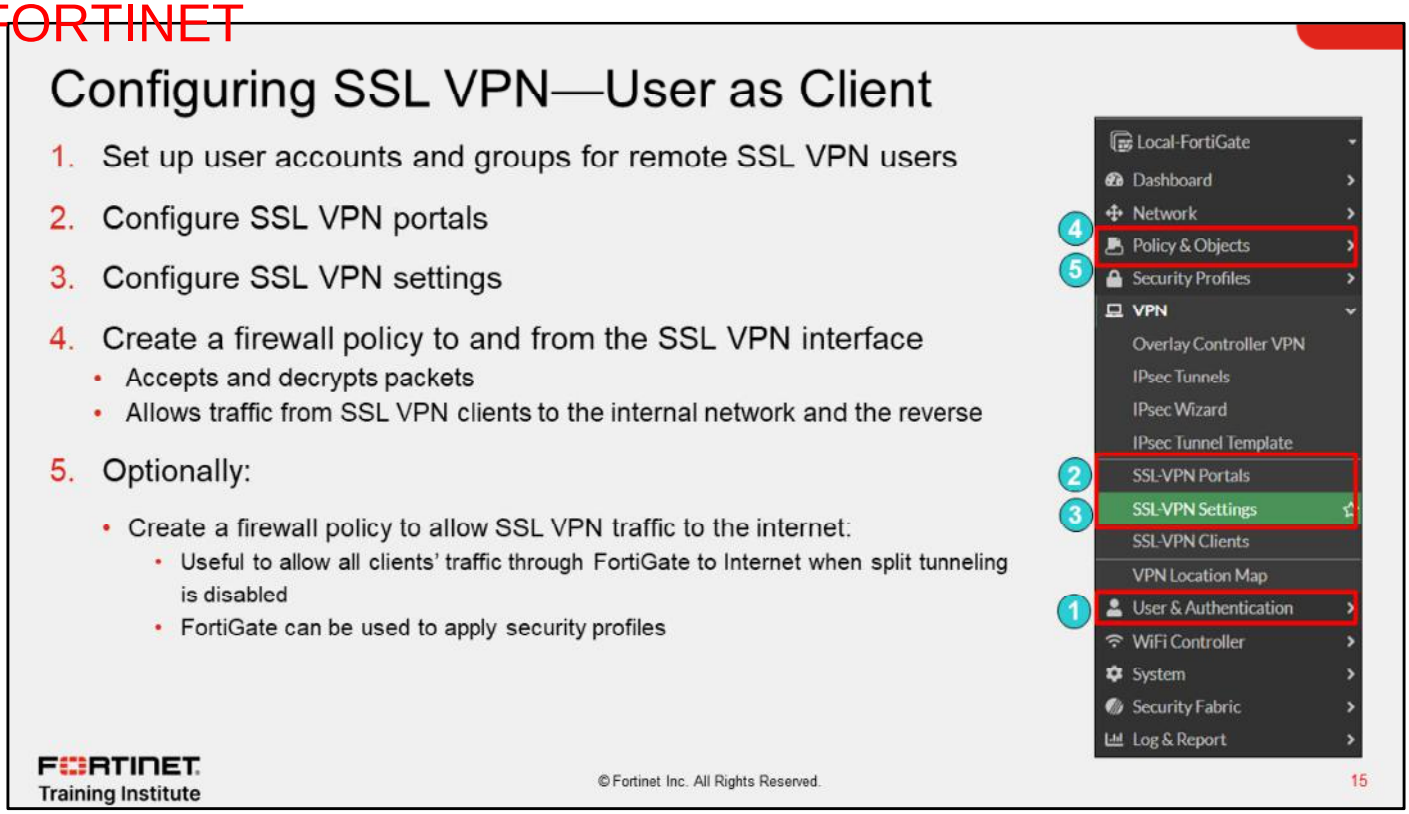

To configure SSL VPN, you must take these steps:

1. Configure user accounts and groups.
2. Configure the SSL VPN portal.
3. Configure SSL VPN settings.
4. Create a firewall policy to accept and decrypt packets. This policy is also used to provide access to internal networks.
5. Optionally, configure a firewall policy to allow traffic from the SSL VPN client to the internet and apply security profiles. User traffic will go to the internet through FortiGate, where you can monitor or restrict client access to the internet.

The first step is to create the accounts and user groups for the SSL VPN clients.

All FortiGate authentication methods, with the exception of remote password authentication using the Fortinet Single Sign-On (FSSO) protocol, can be used for SSL VPN authentication. This includes local password authentication and remote password authentication (using the LDAP, RADIUS, and TACACS+ protocols).

Some steps can be configured in a different order than what is shown on this slide.

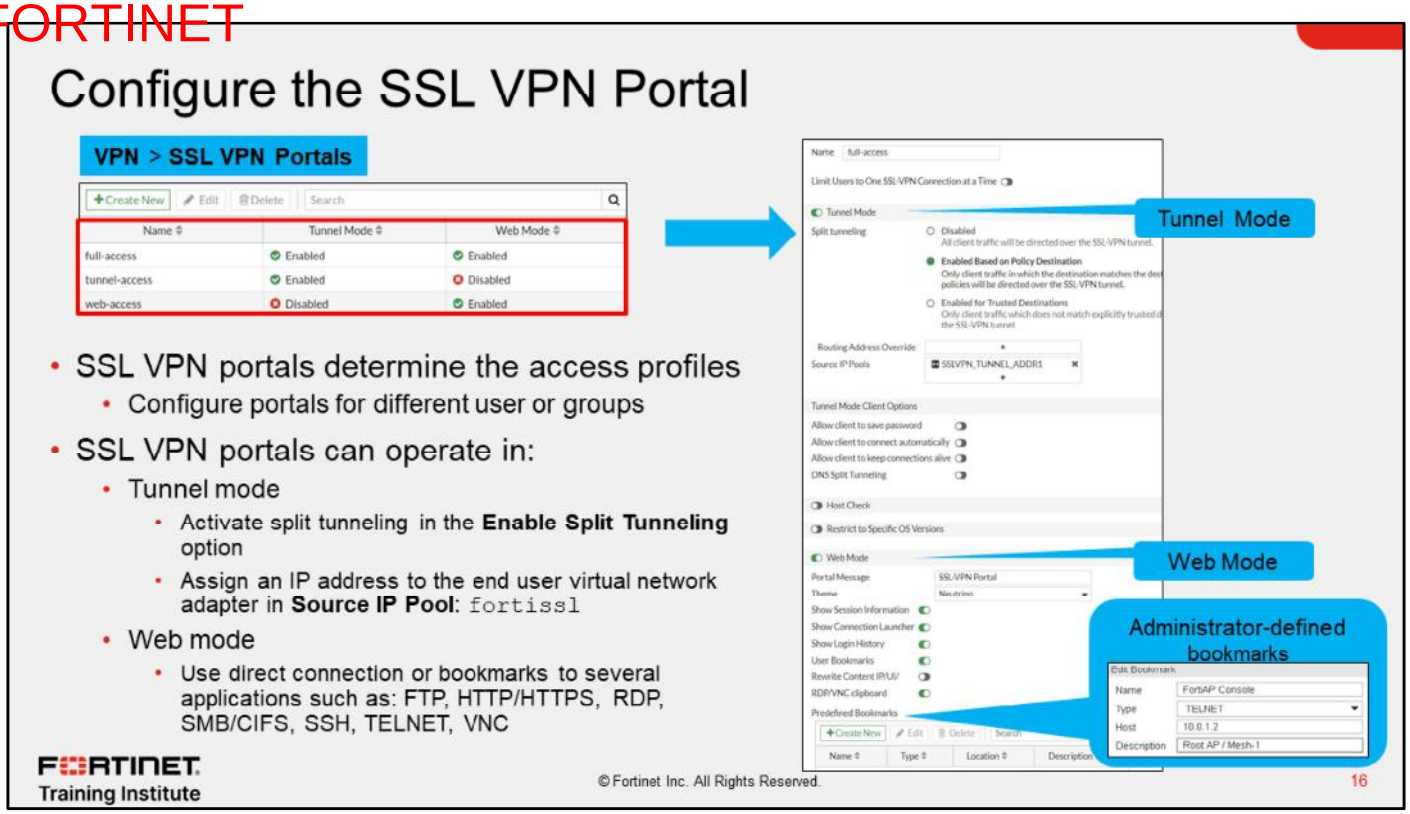

The next step is to configure the SSL VPN portal(s). An SSL VPN portal contains tools and resource links for the users to access.

In tunnel mode, when you enable split tunneling, you need to select either **Enabled Based on Policy Destination** or **Enabled for Trusted Destination** setting, which usually specifies networks behind the FortiGate for the SSL VPN users to access. **Enabled Based on Policy Destination** allows client traffic in which destination is matched with the destination configured on the SSL VPN firewall policy where as **Enabled for Trusted Destination** allows client traffic that does not match the explicitly trusted destination.

Also, for tunnel mode you need to select an IP pool for users to acquire an IP address when connecting. There is a default pool available within the address objects if you do not create your own.

If you enable web mode, you can customize the SSL VPN portal and preconfigure bookmarks to appear for all users who log in to the SSL VPN portal. Also, you can individually configure and link each portal to a specific user or user group, so they have access to only required resources.

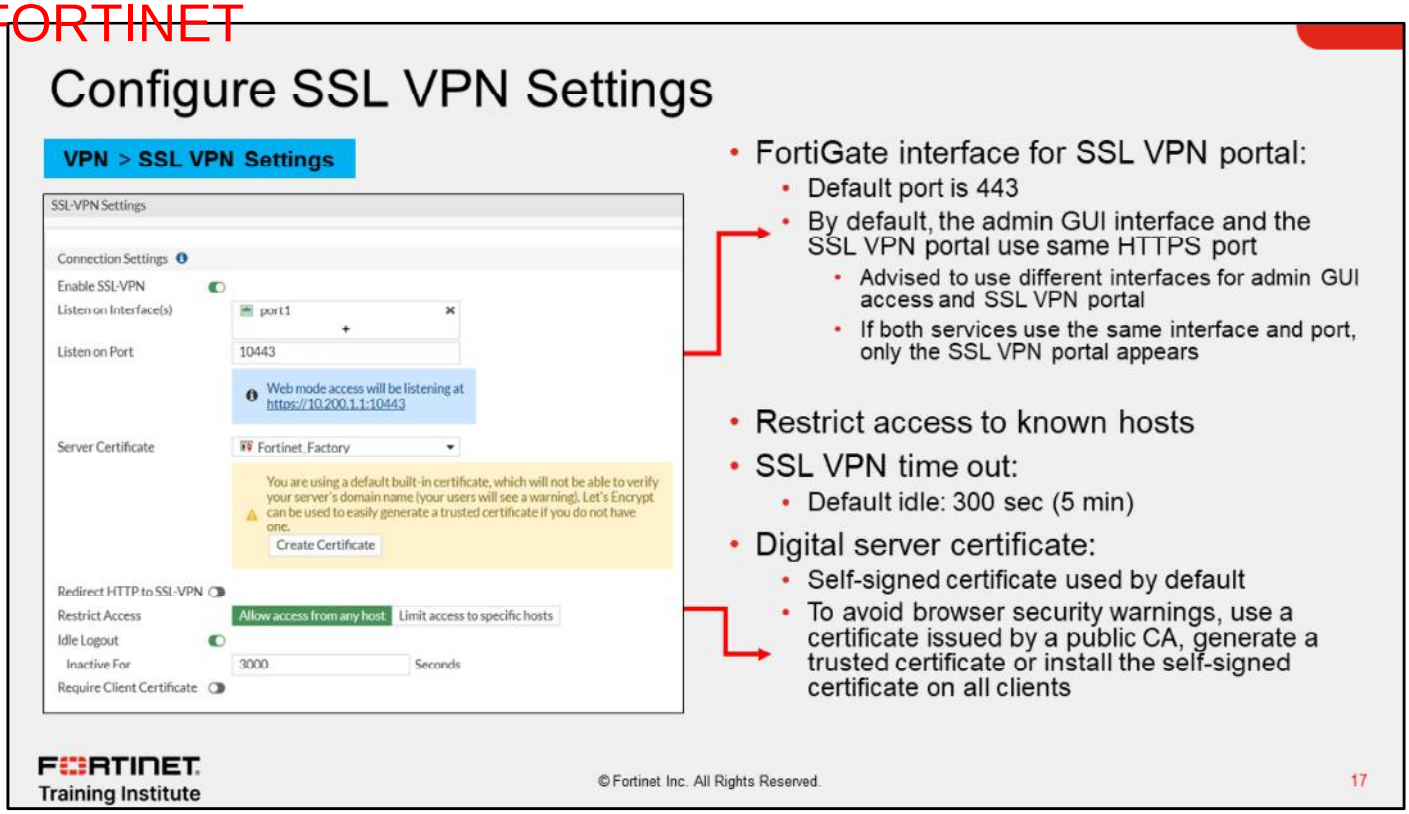

After you configure the SSL VPN portal, the next step is to configure the SSL VPN settings.

Let's start with the **Connection Settings** section. Here, you need to map a FortiGate interface to the SSL VPN portal. The default port for the SSL VPN portal is 443. This means users need to connect to the IP address of the FortiGate interface mapped to the SSL VPN portal, using port443 HTTPS. If you enable **Redirect HTTP to SSL VPN**, users who connect using HTTP (TCP port 80) will be redirected to HTTPS.

Port 443 is the standard default port for administration of the HTTPS protocol. This is convenient because users do not need to specify the port in their browsers. For example, `https://www.example.com/` automatically uses port443 in any browser. This is considered a valid setup on FortiGate because you usually don't access the SSL VPN login through every interface. Likewise, you generally don't enable administrative access on every interface of your FortiGate. So, even though the ports may overlap, the interfaces that each one uses to access may not. However, if the SSL VPN login portal and HTTPS admin access both use the same port, and are both enabled on the same interface, only the SSL VPN login portal will appear. To have access to both portals on the same interface, you need to change the port number for one of the services. If you change the administrator access port, this will affect the port number for that service on all interfaces.

Also, an inactive SSL VPN is disconnected after 300 seconds (5 minutes) of inactivity. You can change this timeout using the **Idle Logout** setting on the GUI.

Finally, like other HTTPS websites, the SSL VPN portal presents a digital certificate when users connect. By default, the portal uses a self-signed certificate, which triggers the browser to show a certificate warning. To avoid the warning, you should use a digital certificate signed by a publicly known certificate authority (CA). You can also generate a certificate for interface. Alternatively, you can load the FortiGate self-signed digital certificate into the browser as a trusted authority.

Configure SSL VPN Settings (Contd)

- Define the IP range for the SSL VPN
  - IPs are assigned to clients' virtual adapters while joined to VPN
  - IP allocation has two methods:
    - First-available (default) or Round robin
    - CLI only

```
conf vpn ssl settings
 set tunnel-addr-assigned-method first-available/round-robin
end
```

- Resolve names by DNS server
  - Use internal DNS if resolving internal domain names
  - Optionally, resolve names by WINS servers
- Specify authentication portal mapping
  - Specify portals for each user or group
  - Define portal for all other users or groups
    - It cannot be deleted

FORTINET Training Institute

18

Define the tunnel-mode client settings and the authentication rules that map users to the appropriate portal.

When users connect, the tunnel is assigned an IP address. You can choose to use the default range or create your own range. The IP range determines how many users can connect simultaneously. There are two IP allocation methods and only available in CLI as shown in the slide:

- First-available (default setting)
- Round robin

Please note when round-robin is used, address pools defined in web portal is ignored, and the `tunnel-ip-pools` or `tunnel-ipv6-pools` under `ssl vpn` setting must be set. Only one set of IP pool address is allowed.

DNS server resolution is effective only when the DNS traffic is sent over the VPN tunnel. Generally, this will be the case only when split tunnel mode is disabled and all traffic is being sent from the user's computer across the tunnel.

Finally, you can allow different groups of users to access different portals. In the example shown on this slide, teachers have access only to the web portal. Accountants can use FortiClient to connect in tunnel mode.

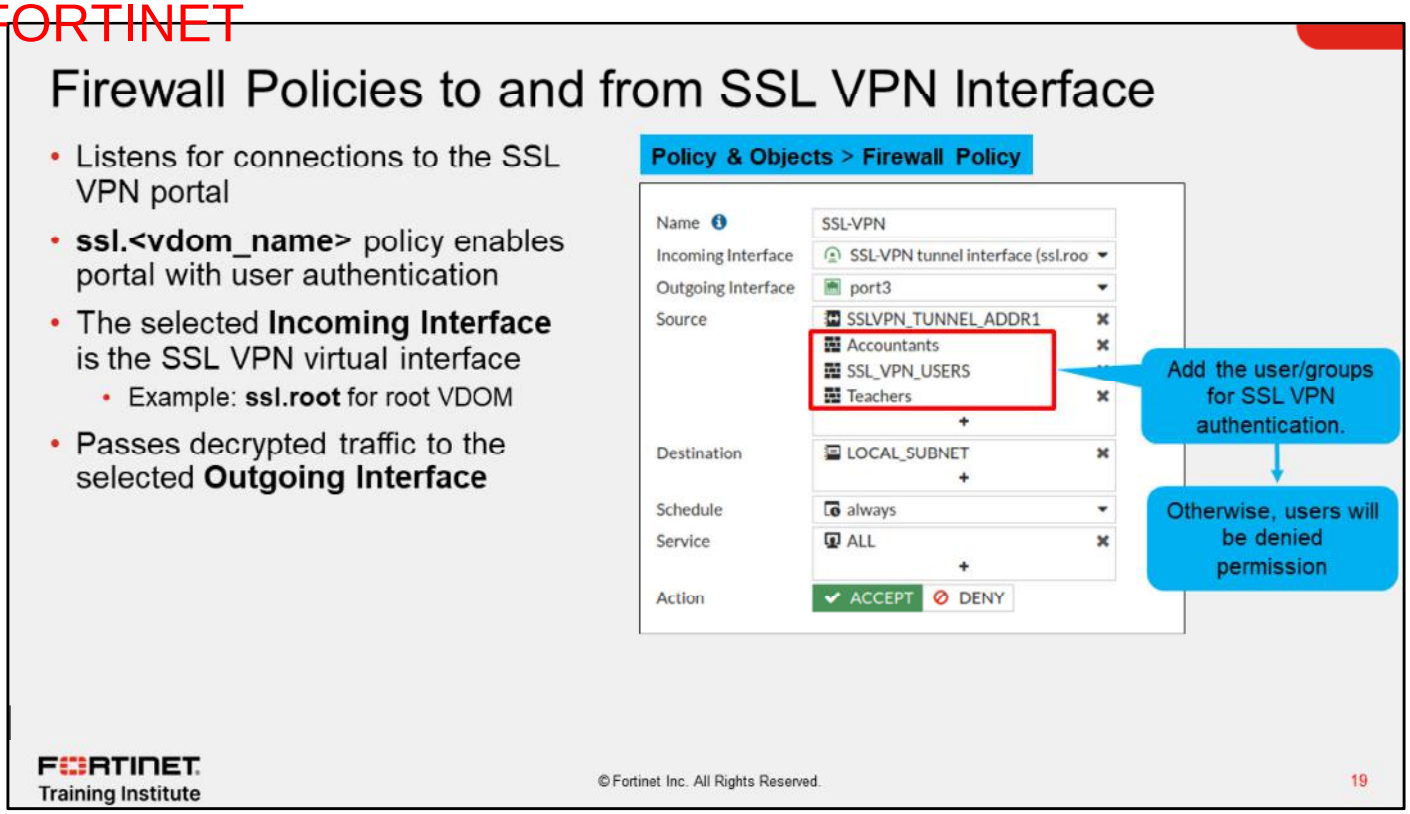

The fourth, and last, mandatory step involves creating firewall policies for logging on.

SSL VPN traffic on FortiGate uses a virtual interface called `ssl.<vdom_name>`. Each virtual domain (VDOM) contains a different virtual interface based on its name. By default, if VDOMs are not enabled, then the device operates with a single VDOM called *root*.

To activate and successfully log in to the SSL VPN, there must be a firewall policy from the SSL VPN interface to the interface to which you want to allow access for the SSL VPN users, including all of the users and groups that can log in as the source. Without a policy like this, no login portal is presented to users.

If there are resources behind other interfaces that users need access to, then you need to create additional policies that allow traffic from `ssl.root` to exit those interfaces.

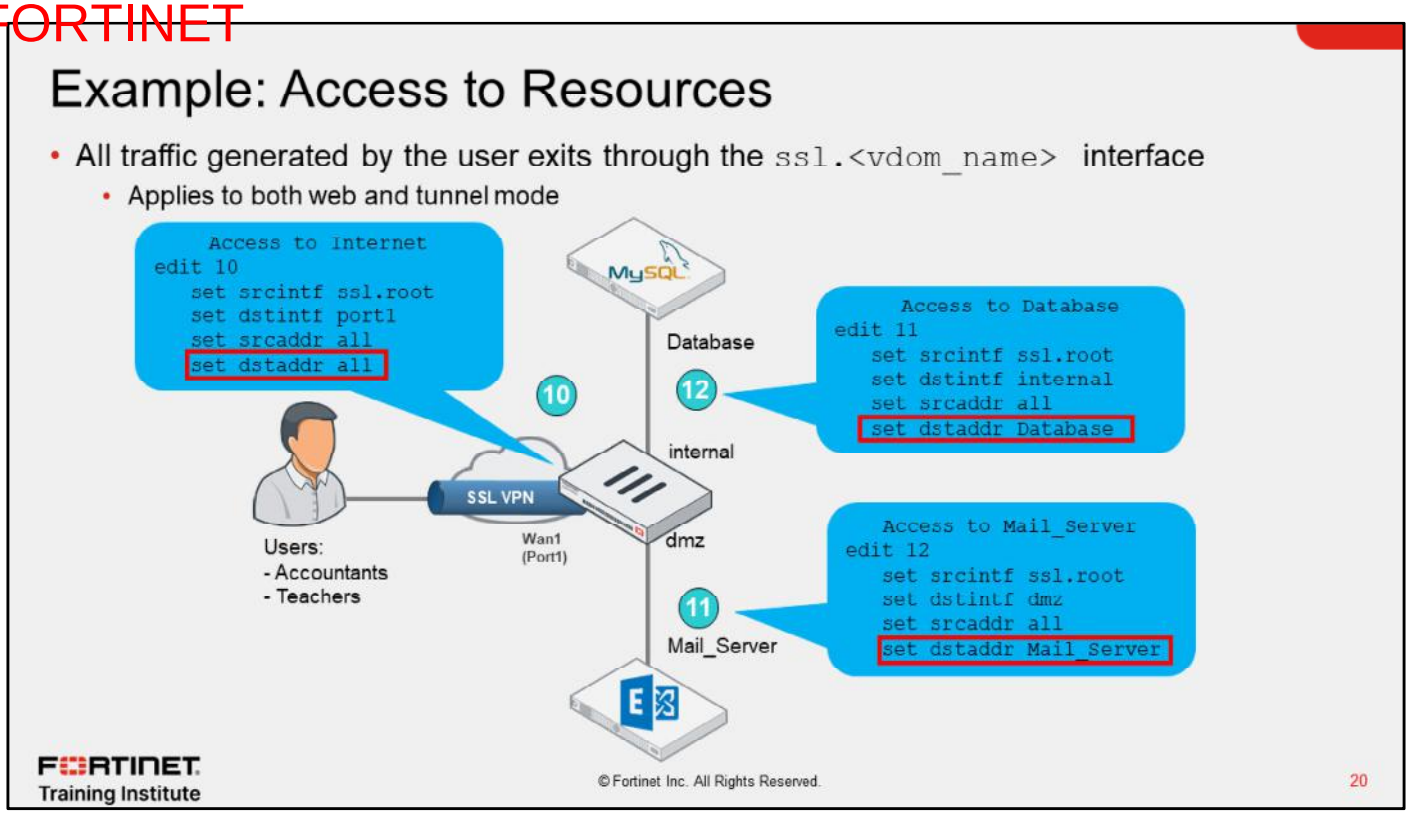

Any traffic from SSL VPN users, whether in web portal or tunnel mode, exits from the `ssl.<vdom_name>` interface.

This slide shows an example of firewall policies that are configured to allow access to resources behind other interfaces that users need access to when connected through SSL VPN.

Optionally, if split tunneling is disabled, you need to create an additional firewall policy from `ssl.root` to the egress interface to allow clients access to the internet.

You can also apply security profiles to this firewall policy to restrict user access to the internet.

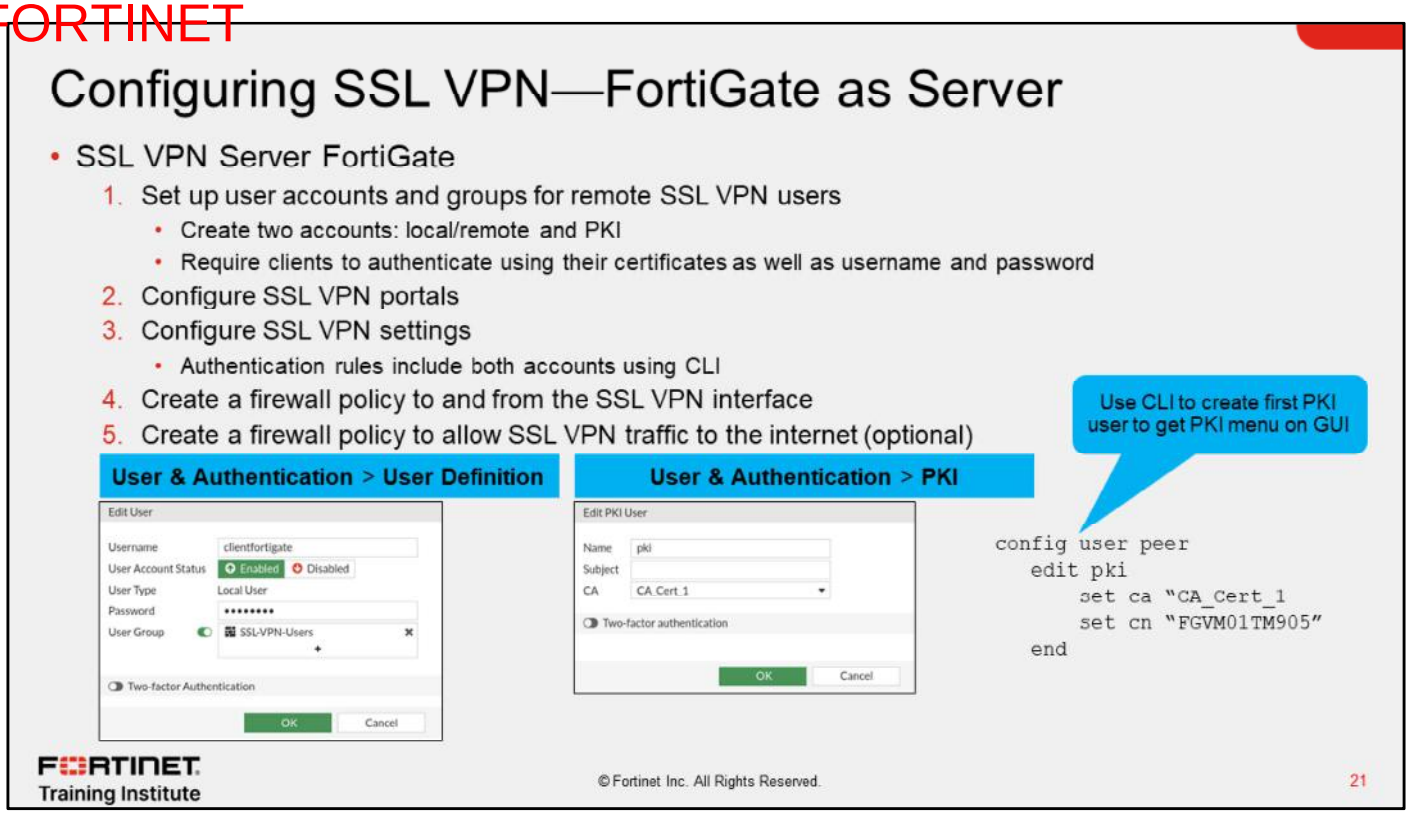

To configure SSL VPN, you must take these steps:
SSL VPN server FortiGate:

1. Set up user accounts and groups for remote SSL VPN users.
   - Create two accounts: local/remote and PKI. The PKI menu is only available in the GUI after a PKI user has been created using the CLI, and a CN can only be configured in the CLI. If no CN is specified, then any certificate that is signed by the CA will be valid and matched.
   - Require clients to authenticate using their certificates as well as username and password.
2. Configure SSL VPN portals.
3. Configure SSL VPN settings.
   - Authentication rules include both accounts using CLI.
4. Create a firewall policy to and from the SSL VPN interface.
5. Create a firewall policy to allow SSL VPN traffic to the internet (optional).

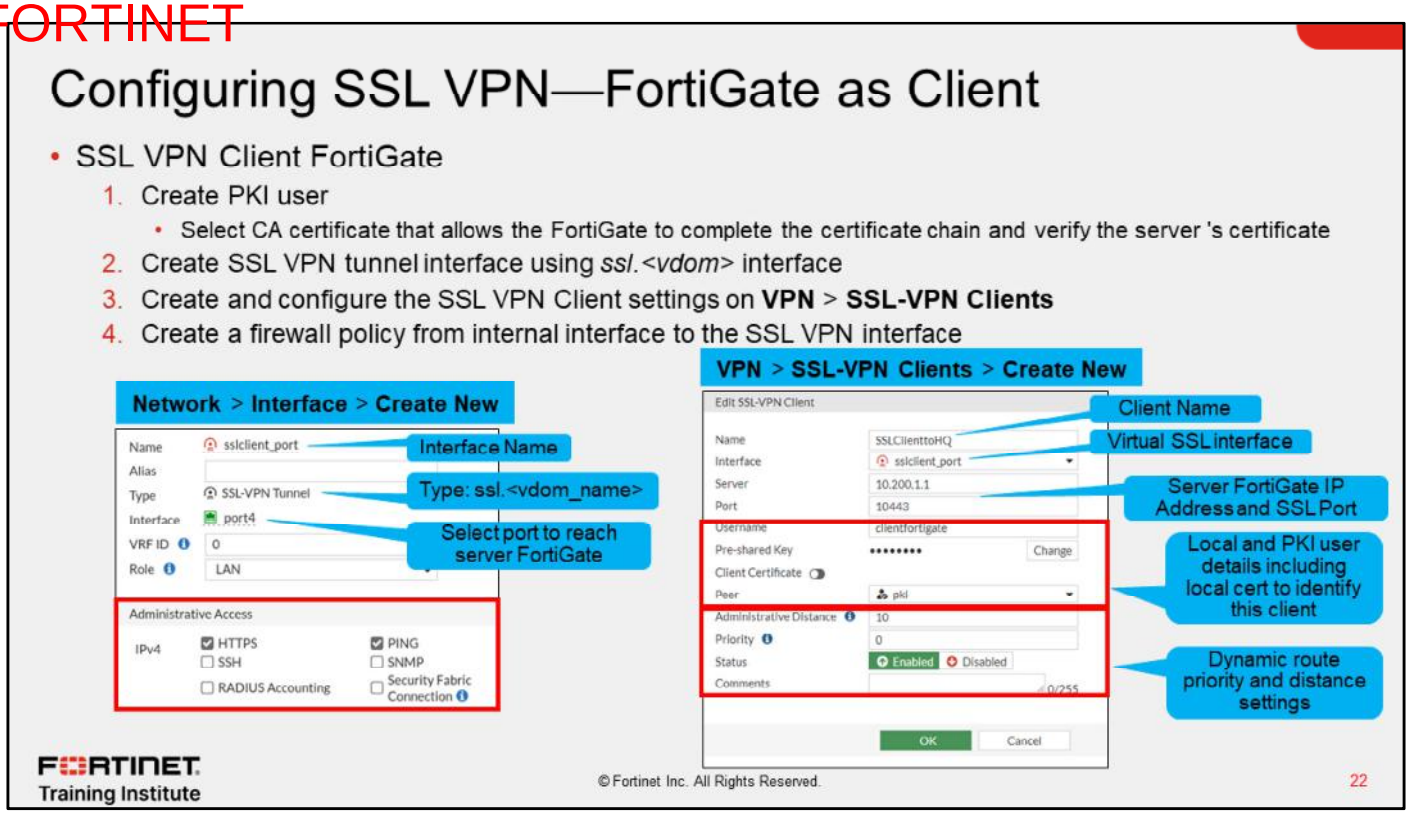

To configure SSL VPN, you must take these steps:
SSL VPN Client FortiGate:

1. Create PKI user:
   - Set the same CN using CLI if PKI user on server FortiGate has CN configured.
   - Select CA certificate that allows the FortiGate to complete the certificate chain and verify the server 's certificate.
2. Create SSL VPN tunnel interface using *ssl.<vdom>* interface.
3. Create and configure the SSL VPN client settings on **VPN** > **SSL-VPN Clients**, it includes:
   - Client name
   - Virtual SSL VPN interface
   - SSL VPN server FortiGate IP address and SSL port number
   - Local username and password and PKI(Peer) user. The **Client Certificate** is the local certificate that is used to identify this client, and is assumed to already be installed on the FortiGate. The SSL VPN server requires it for authentication.
   - When split tunnel is disabled, new default route is added and priority and distance plays an important role.
4. Create a firewall policy to allow traffic from internal interface to the SSL VPN interface.

Client Integrity Checking

- SSL-VPN gateway checks client integrity
  - Requires Microsoft Windows
  - Supported in SSL VPN tunnel mode only
- Detects client security applications recognized by the Windows Security Center
  - Antivirus and firewall software
  - Security attributes recorded on the client's computer
- Checks the status of applications through their globally unique identifier (GUID)
  - Custom host checks
- Determines the state of the applications
  - Active/inactive
  - Current version number
  - Signature updates

FORTINET Training Institute

23

When a user connects to your network through an SSL-VPN, a portal is established between your network and the user's PC. The VPN session is secured natively in two ways: the connection is encrypted and the user must log in with their credentials, such as a username and password. However, you can configure additional checks to increase the security of the connection.

One method of increasing your security is by using client integrity checking. Client integrity ensures that the connecting computer is secure by checking whether specific security software, such as antivirus or firewall software, is installed and running. This feature supports only Microsoft Windows clients, because it accesses the Windows Security Center to perform its checks. Alternatively, you can customize this feature to check the status of other applications using their GUIDs. A GUID is a unique ID in the Windows Configuration Registry that identifies each Windows application. Client integrity can also check the current software and signature versions for the antivirus and firewall applications.

Client integrity checking is applicable to tunnel mode only.

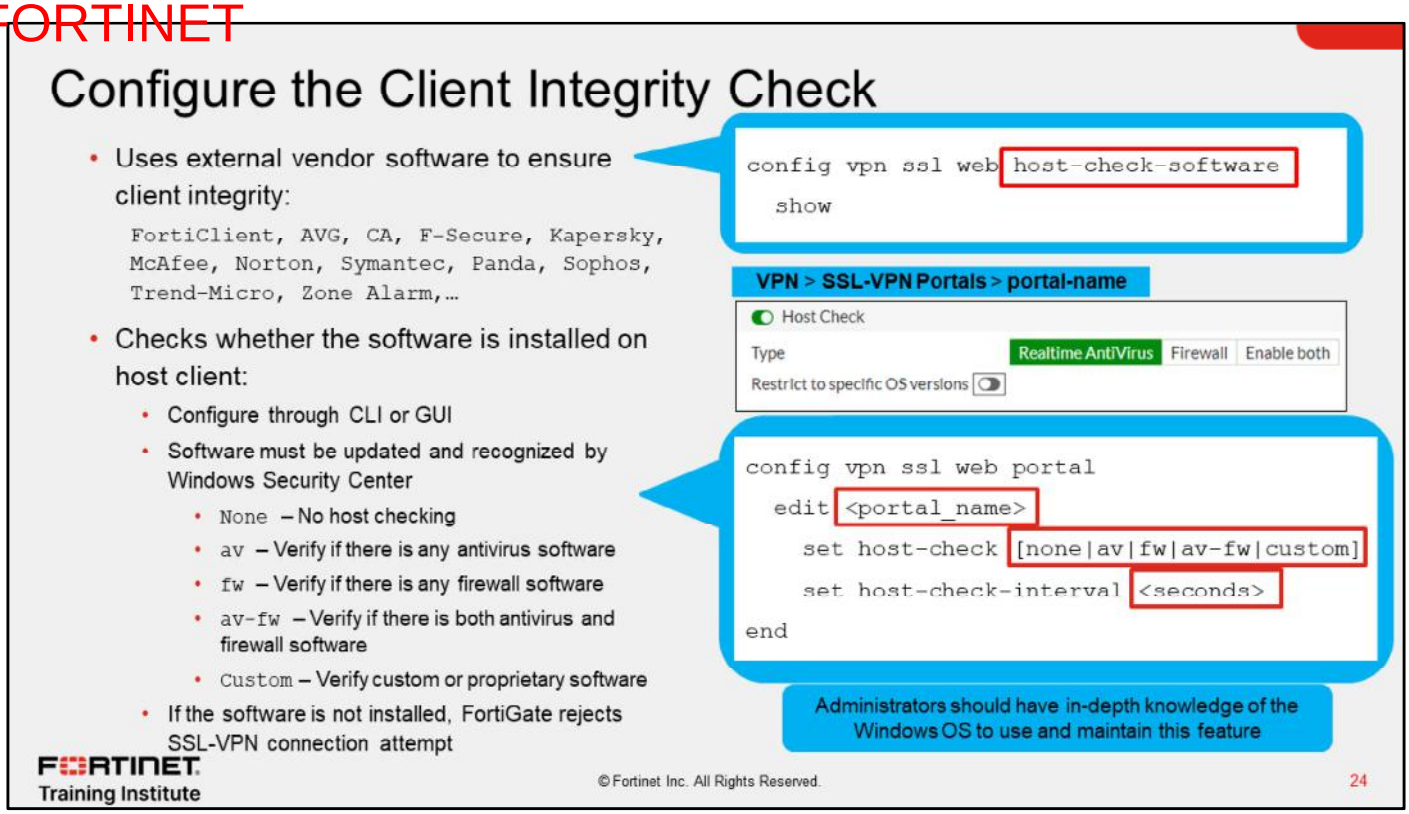

FortiGate performs the client integrity check while the VPN is still establishing, just after user authentication has finished. If the required software is not running on the user's PC, FortiGate rejects the VPN connection attempt, even with valid user credentials. You enable client integrity for each web portal, and you configure it using CLI commands or the FortiGate GUI.

The list of recognized software, along with the associated registry key value, is available on the CLI only. Software is split into three categories: antivirus `(av)`, firewall `(fw)`, and `custom`. Custom is used for customized or proprietary software that an organization may require. Administrators can configure `av`, `fw`, or both settings on the GUI or CLI, but the custom setting is available only on the CLI.

Administrators can also configure OS versions and patch settings to allow or deny VPN connections from specific OS versions.

The disadvantage of enabling client integrity checking is that it can result in a lot of administrative overhead because of the following factors:

- All users must have their security software up to date in order to successfully establish a connection.
- Software updates can result in a change to the registry key values, which can also prevent a user from successfully connecting.

As such, administrators must have in-depth knowledge of the Windows operating system and subsequent registry behavior in order to properly make extended use of and maintain this feature.

Knowledge Check

1. Which step is necessary to configure SSL VPN connections?
   - ✓A. Create a firewall policy from the SSL VPN interface to the resource's interface.
   - B. Enable event logs for SSL VPN traffic: users, VPN, and endpoints.

2. Which action may allow internet access in tunnel mode, if the remote network does not allow internet access to SSL VPN users?
   - ✓A. Enable split tunneling
   - B. Configure the DNS server to use the same DNS server as the client system DNS

FORTINET Training Institute

25

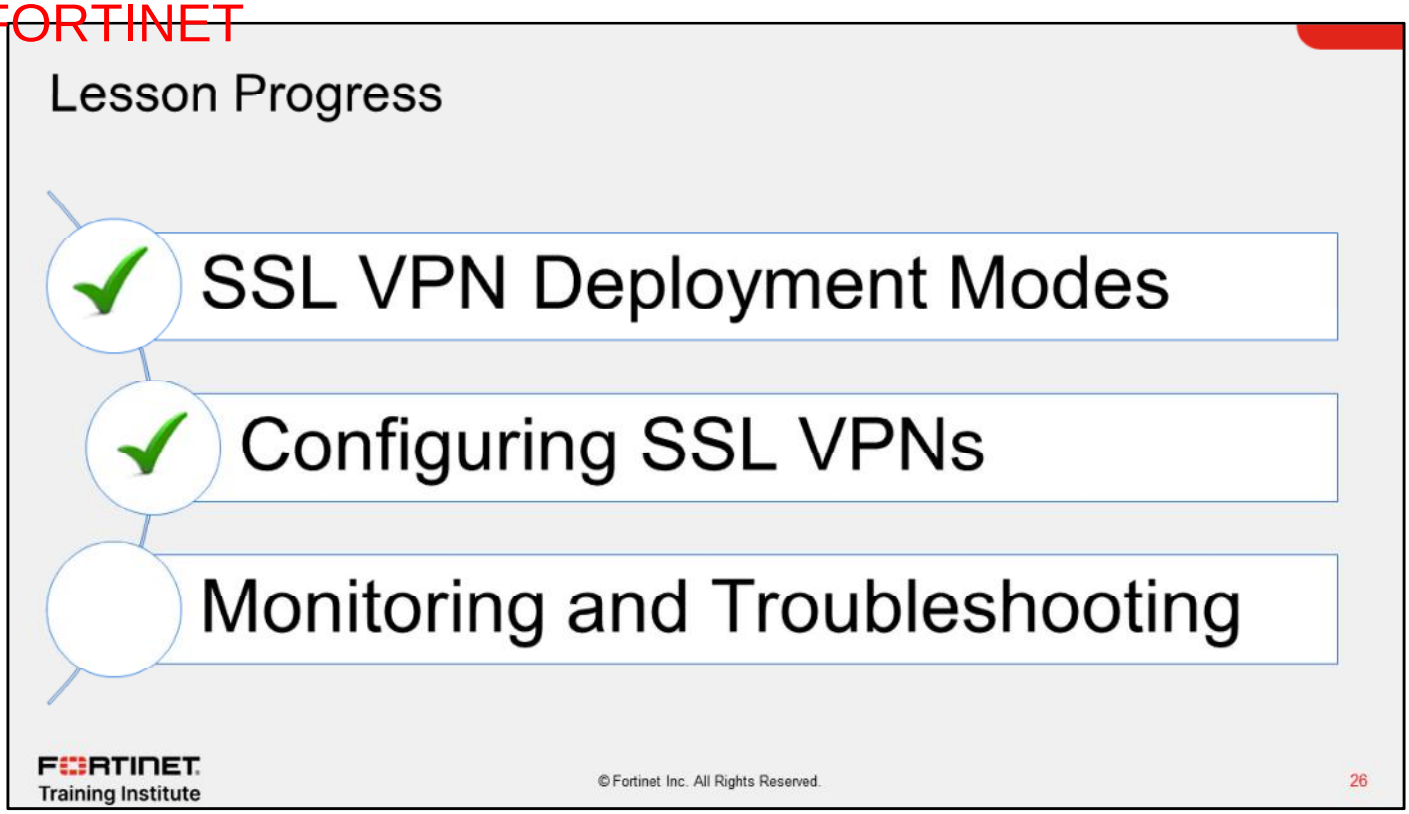

Good job! You now understand how to configure FortiGate for SSL VPN connections.

Now, you'll learn how to monitor SSL VPN sessions, review logs, configure SSL VPN timers, and troubleshoot common issues.

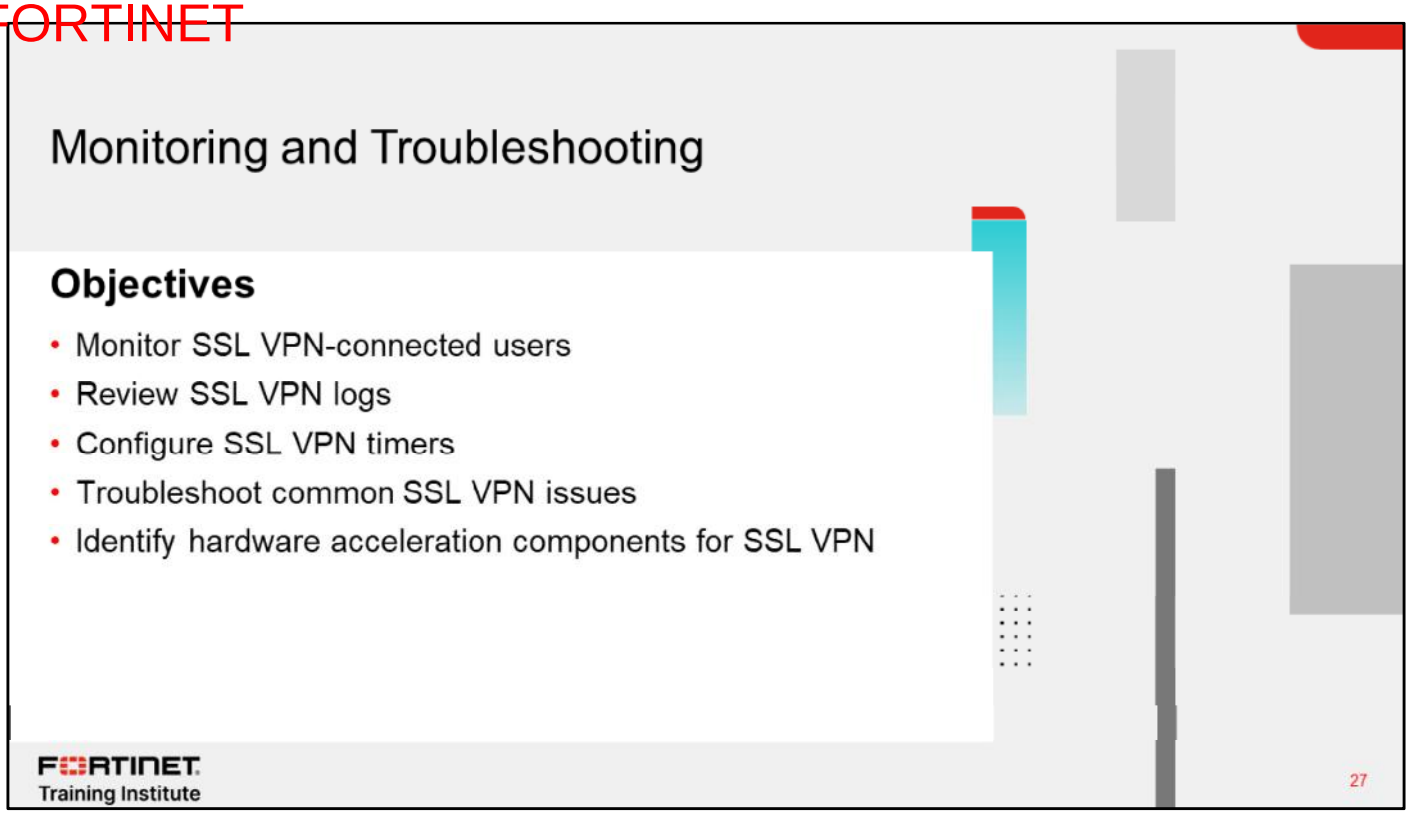

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in SSL VPN monitoring and troubleshooting, you will be able to avoid, identify, and solve common issues and misconfigurations.

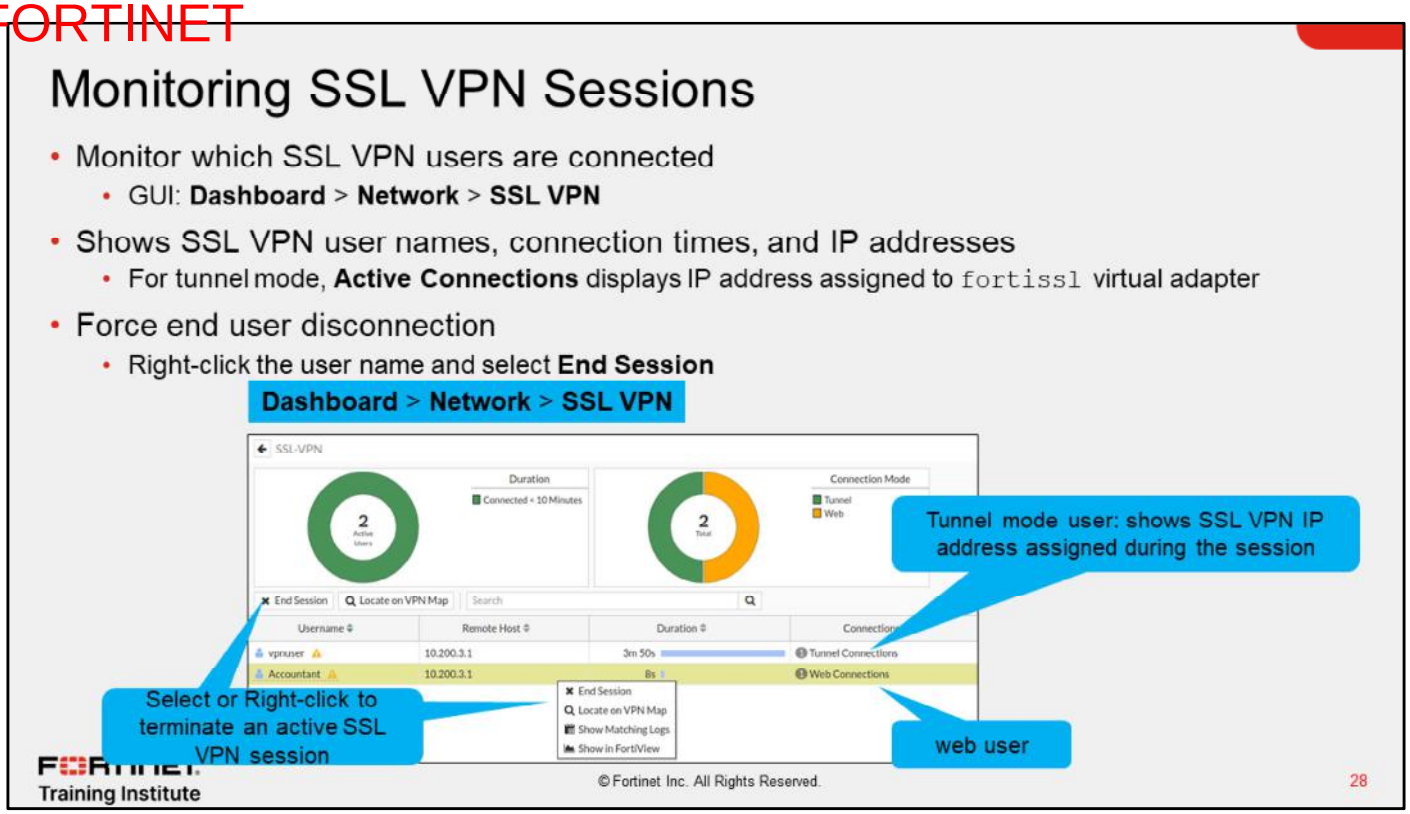

You can monitor which SSL VPN users are connected on the **SSL VPN** widget. This shows the names of all SSL VPN users that are currently connected to FortiGate, their IP addresses (both inside the tunnel and outside), and connection times.

When a user connects using tunnel model, the **Active Connections** column shows the IP address assigned by FortiGate to the `fortissl` virtual adapter on the client's computer. Otherwise, the user is connected only to the web portal page.

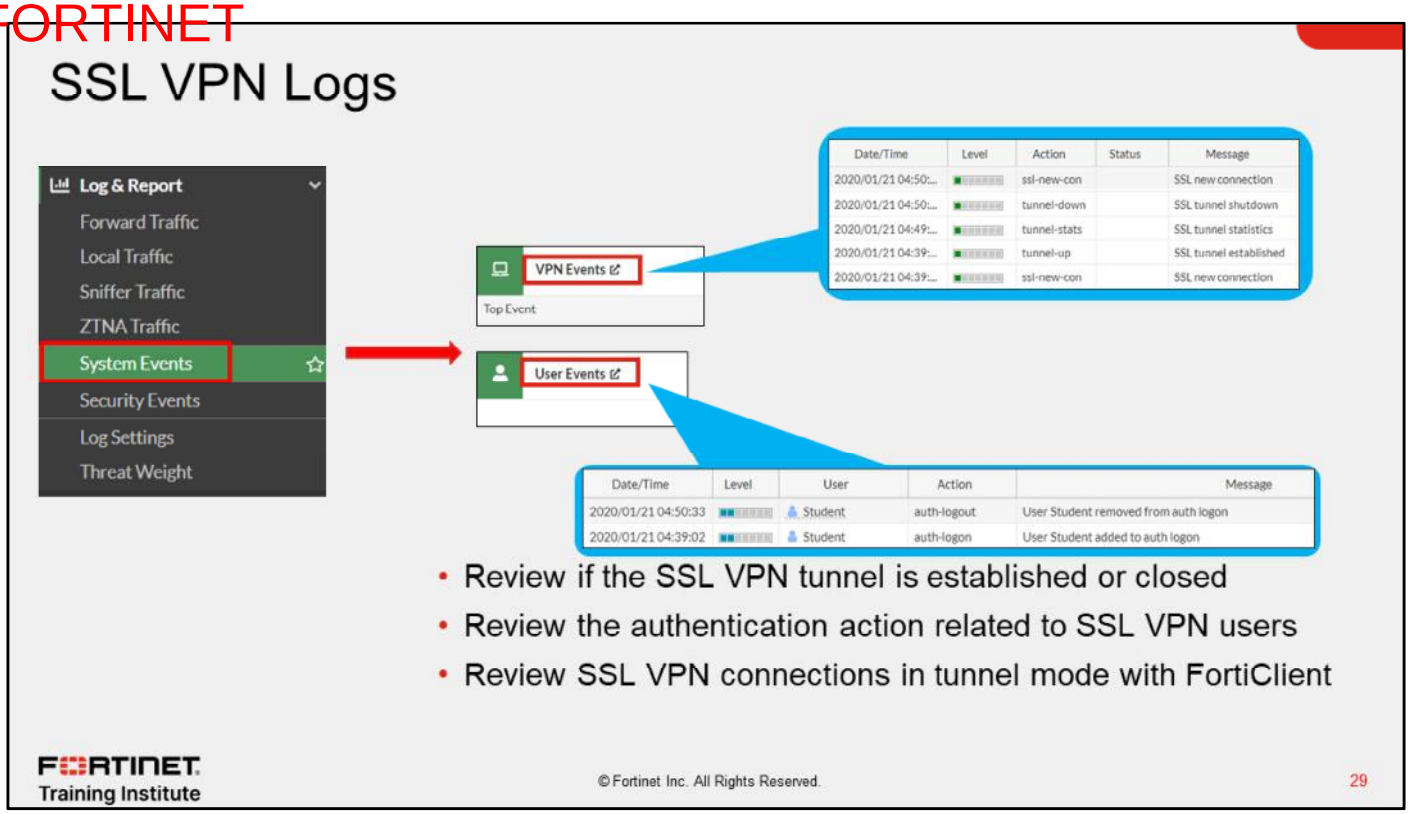

You can also review SSL VPN logs. On **Log & Report** > **System Events**:

- Select **VPN Events** widget to show new connection requests, and if the SSL VPN tunnel is established or closed.
- Select **User Events** widget to see the authentication action related to SSL VPN users.

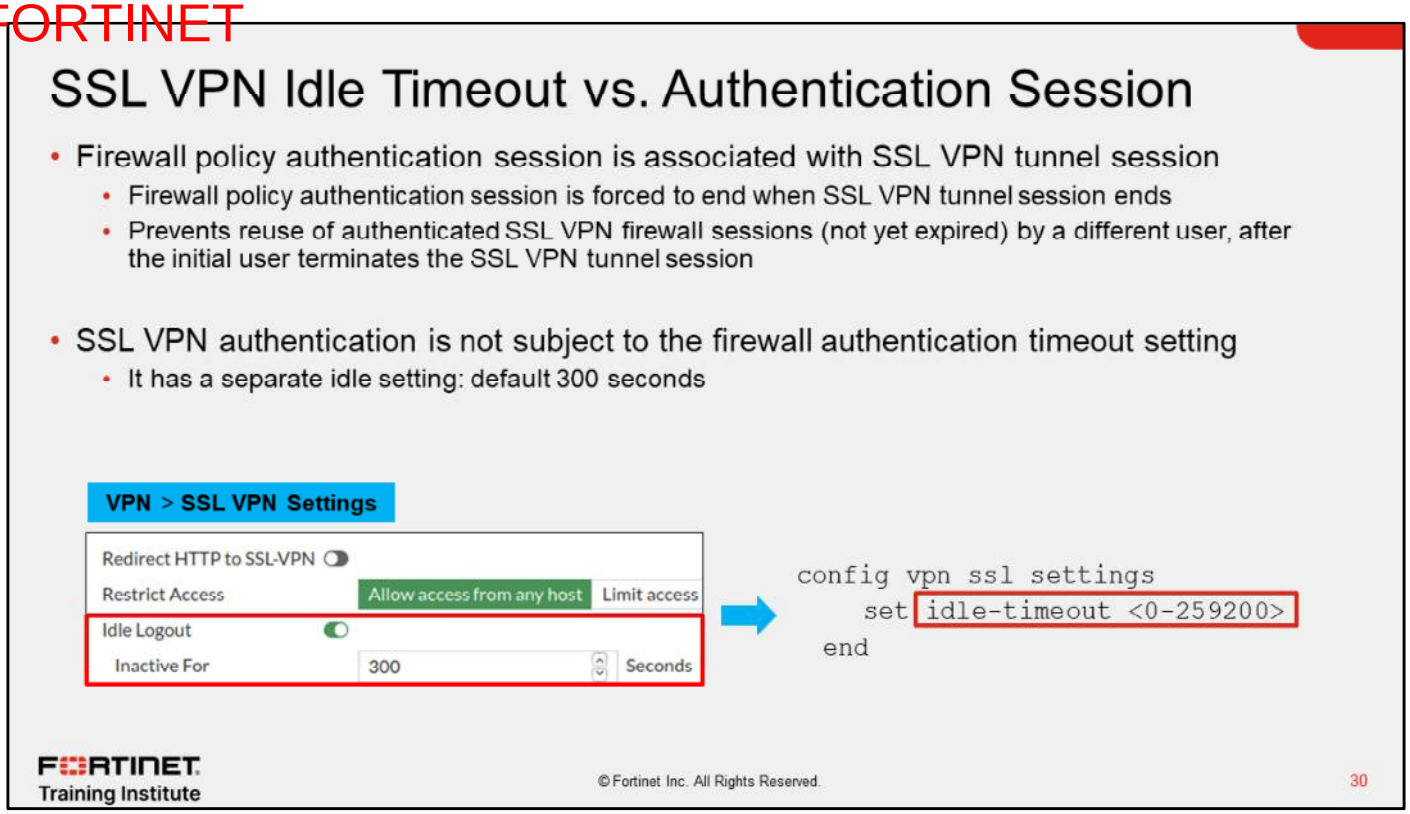

When an SSL VPN is disconnected, either by the user or through the SSL VPN idle setting, all associated sessions in the FortiGate session table are deleted. This prevents the reuse of authenticated SSL VPN sessions (not yet expired) after the initial user terminates the tunnel.

The SSL VPN user idle setting is not associated with the firewall authentication timeout setting. It is a separate idle option specifically for SSL VPN users. A remote user is considered idle when FortiGate does not see any packets or activity from the user within the configured timeout period.

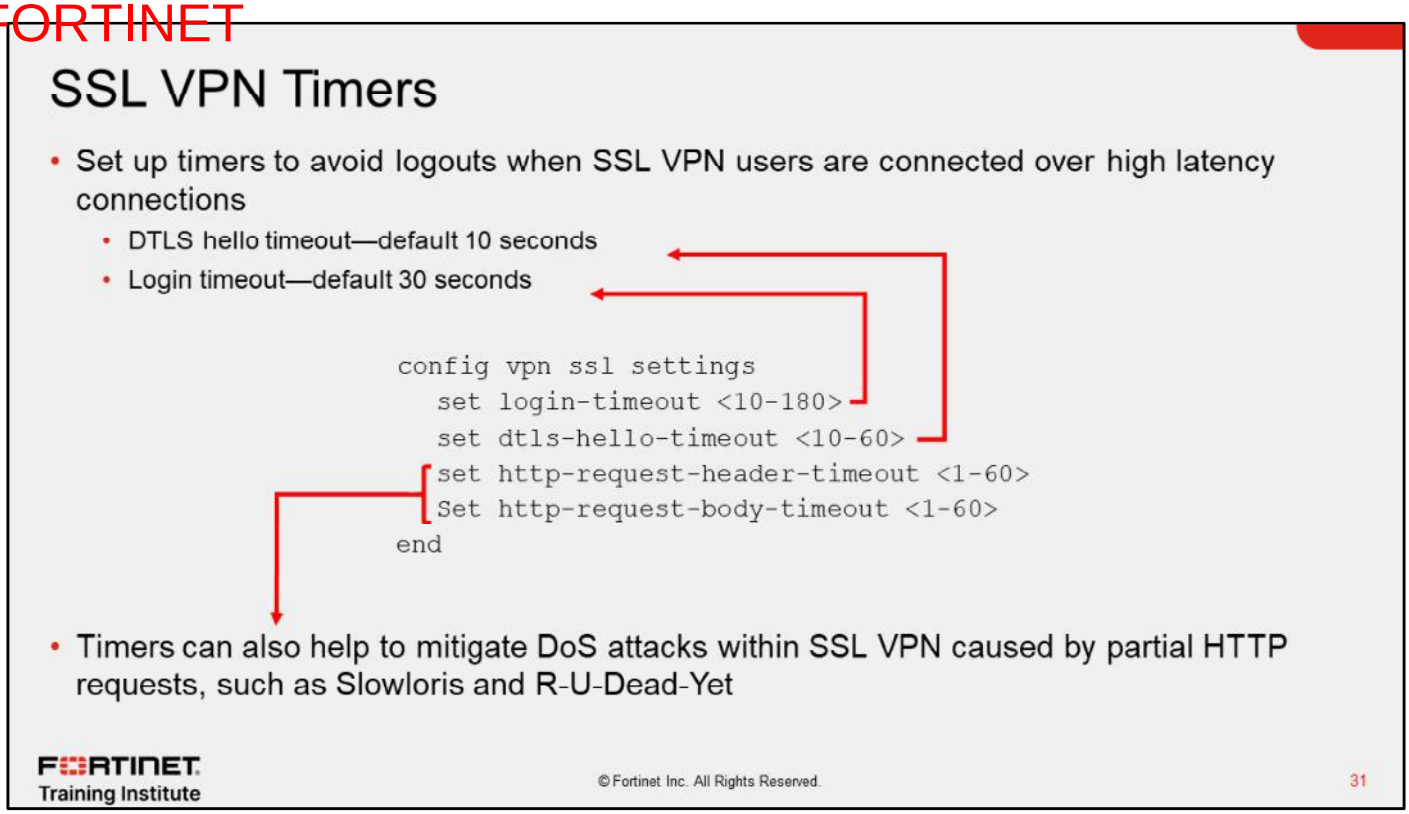

When connected to SSL VPN over high latency connections, FortiGate can time out the client before the client can finish the negotiation process, such as DNS lookup and time to enter a token. Two new CLI commands under `config vpn ssl settings` have been added to address this. The first command allows you to set up the login timeout, replacing the previous hard timeout value. The second command allows you to set up the maximum DTLS hello timeout for SSL VPN connections.

Also, timers can help you to mitigate vulnerabilities such as Slowloris and R-U-Dead-Yet, that allow remote attackers to cause a denial of service through partial HTTP requests.

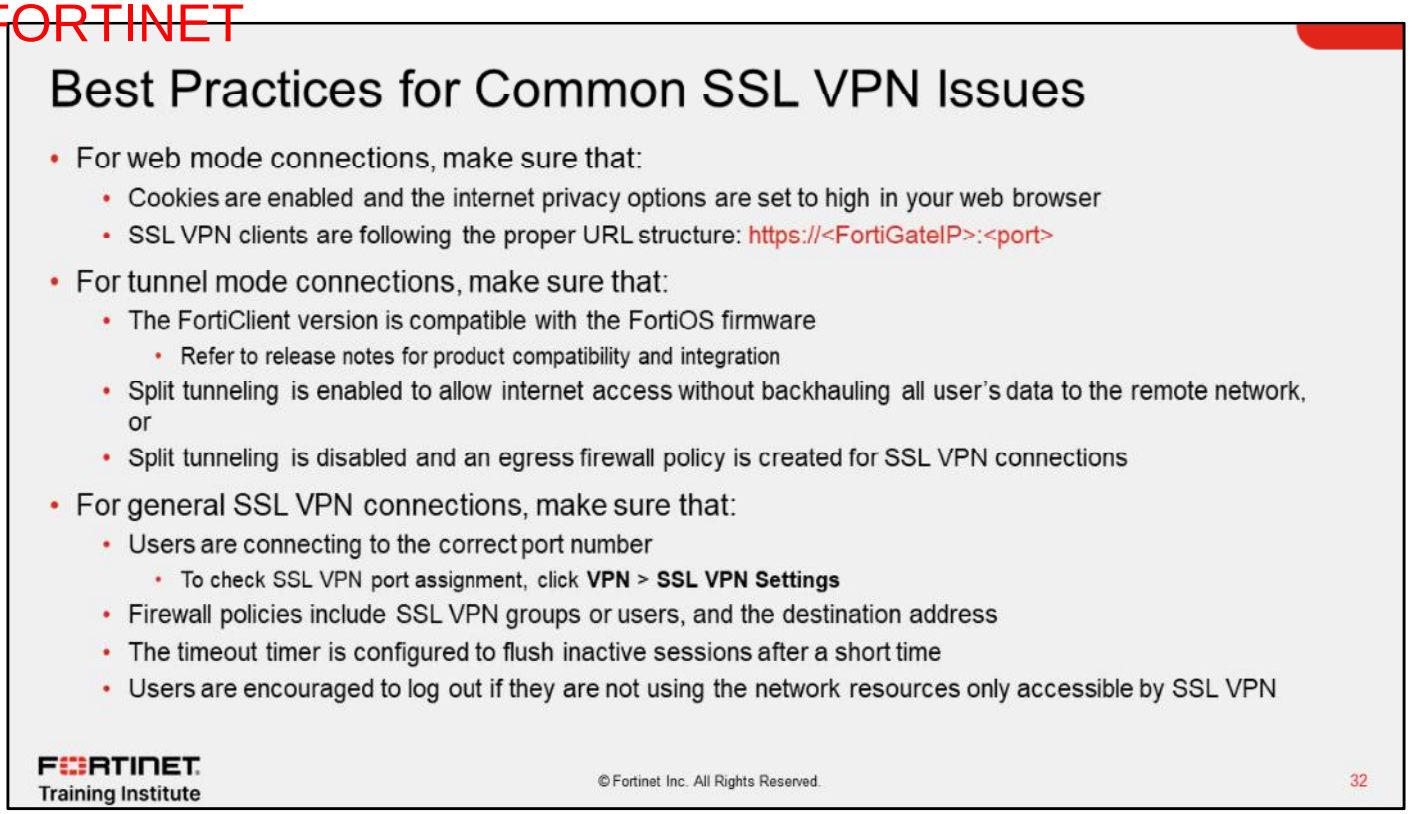

The following are some best practices to keep in mind when using SSL VPNs. These best practices can also be helpful in many SSL VPN troubleshooting situations:

- Enable cookies in your web browser
- Set internet privacy options to high in your web browser
- Use a FortiClient version that is compatible with your FortiOS firmware
- Enable split tunneling or create an egress firewall policy for SSL VPN connections in order to allow access for external resources
- Connect to the correct port number
- Add SSL VPN groups, SSL VPN users, and destination addresses to the firewall policies
- Flush inactive sessions by timeout

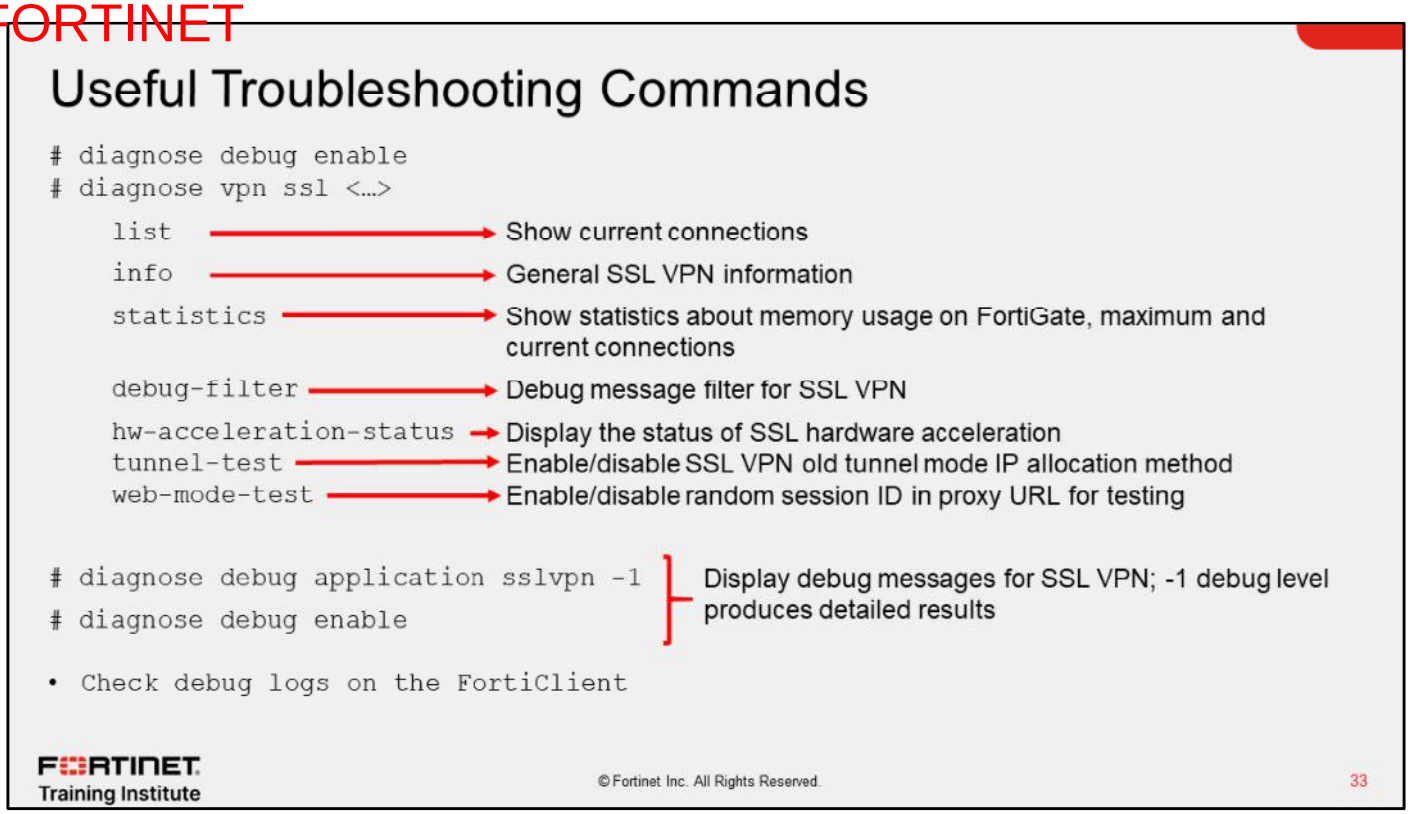

There are several useful troubleshooting commands available under `diagnose vpn ssl`. They include:

- `list`: Lists logged-on users
- `info`: Shows general SSL VPN information
- `statistics`: Shows statistics about memory usage on FortiGate
- `hw-acceleration-status`: Displays the status of SSL hardware acceleration
- `tunnel-test`: Enables or disables SSL VPN old tunnel mode IP allocation method
- `web-mode-test`: Enables or disables random session ID in proxy URL for testing

The command `diagnose debug application sslvpn` shows the entire list of debug messages for SSL VPN connections.

Remember, to use the commands listed above, you must first run the `diagnose debug enable` command. Also, check SSL VPN debug logs on FortiClient.

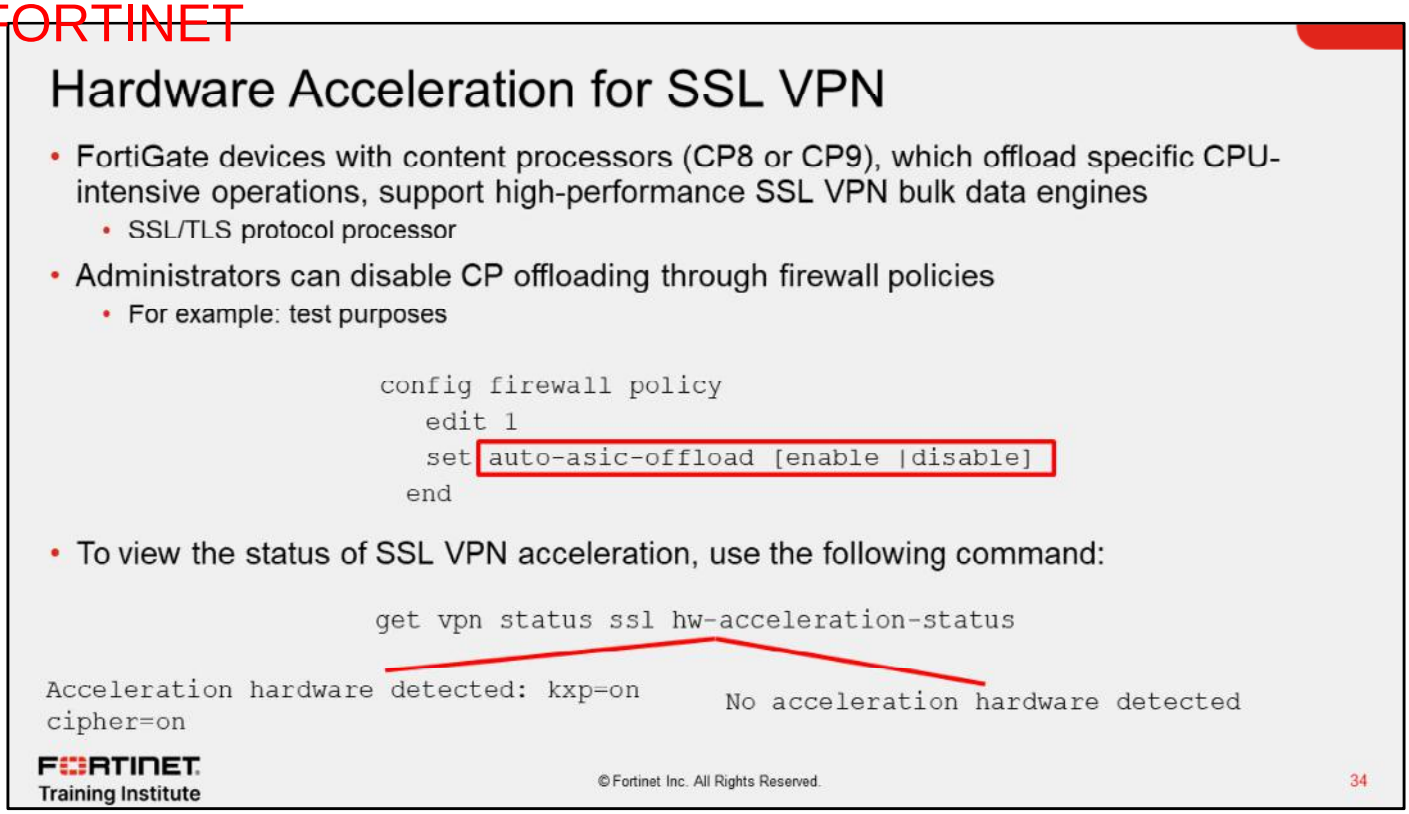

FortiGate devices that have CP8 or CP9 content processors, which accelerate many common resource-intensive, security-related processes, can offload SSL VPN traffic to a high-performance VPN bulk data engine.

This specialized IPsec and SSL/TLS protocol processor processes most of the latest well-known algorithms for encryption.

By default, the offloading process is set up. If, for testing purposes you want to disable it, you can do it using the CLI only at the firewall policy configuration level.

You can also view the status of SSL VPN acceleration using the CLI.

Knowledge Check

1. What does the SSL VPN monitor feature allow you to do?
   - A. Monitor SSL VPN user actions, such as authentication
   - ✓ B. Force SSL VPN user disconnections

2. Which statement about SSL VPN timers is correct?
   - ✓ A. SSL VPN timers can prevent logouts when SSL VPN users experience long network latency.
   - B. The login timeout is a non-customizable hard value.

FORTINET Training Institute

35

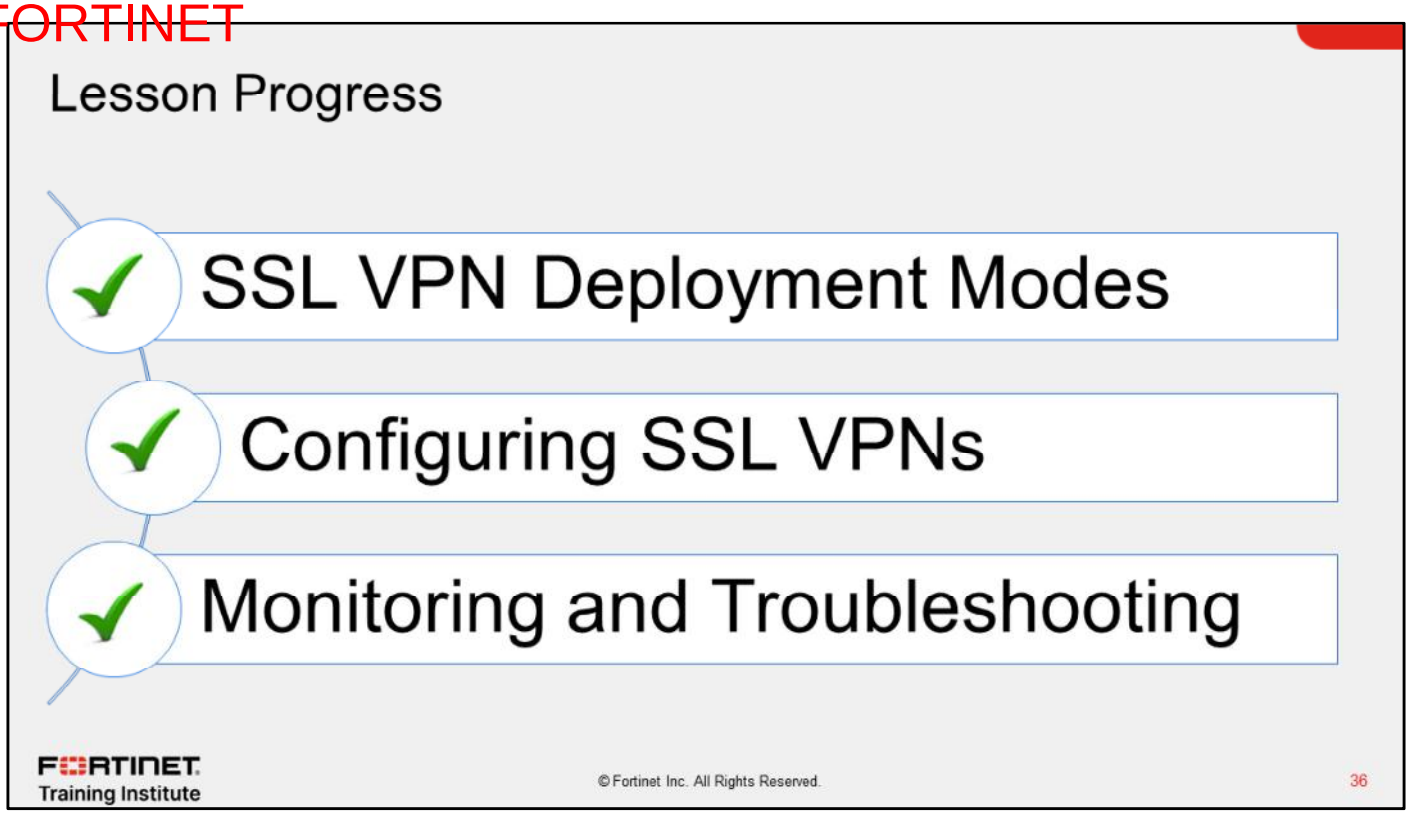

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

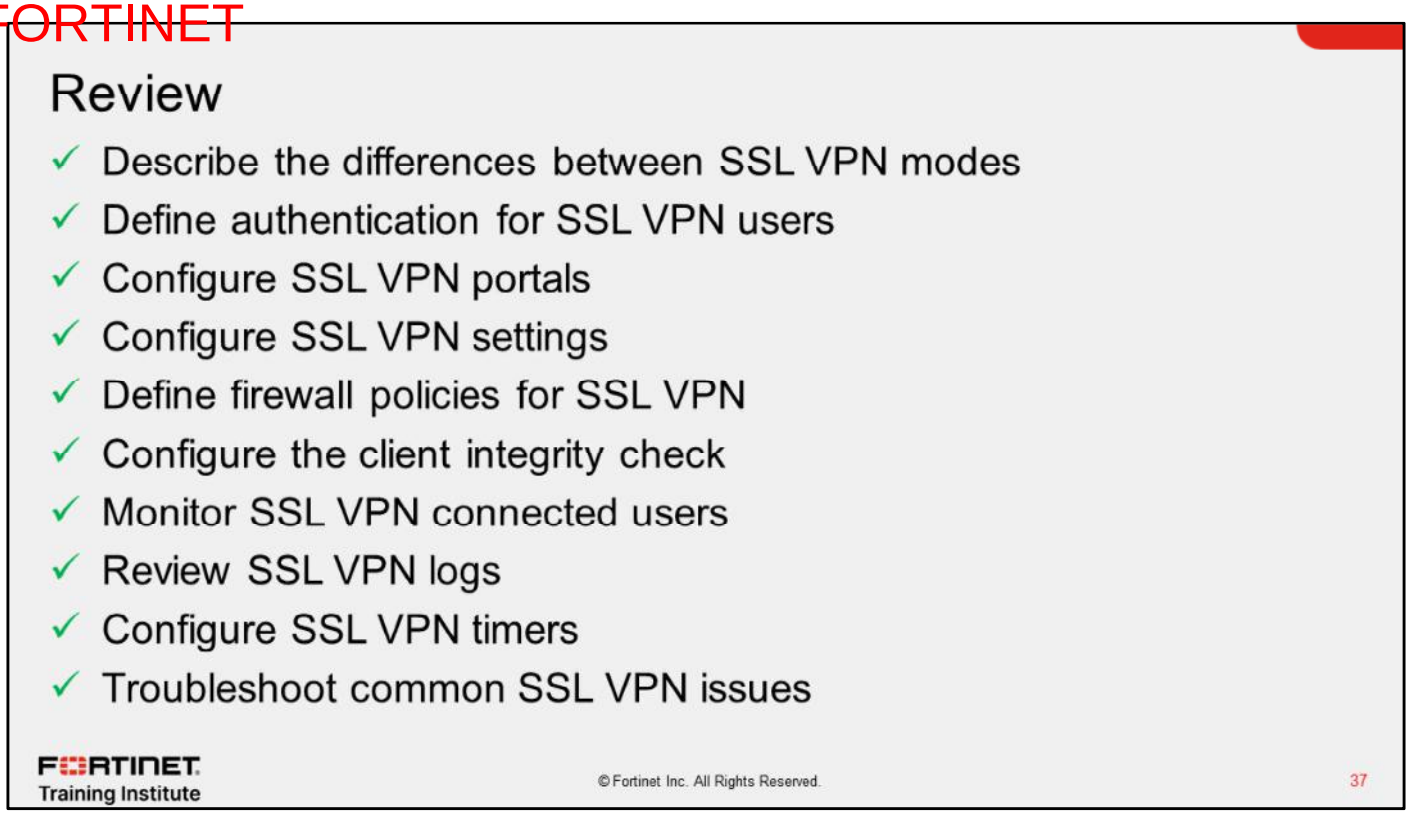

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure and use SSL VPNs to give remote users access to your private network.

In this lesson, you will learn about the architectural components of IPsec VPN and how to configure them.

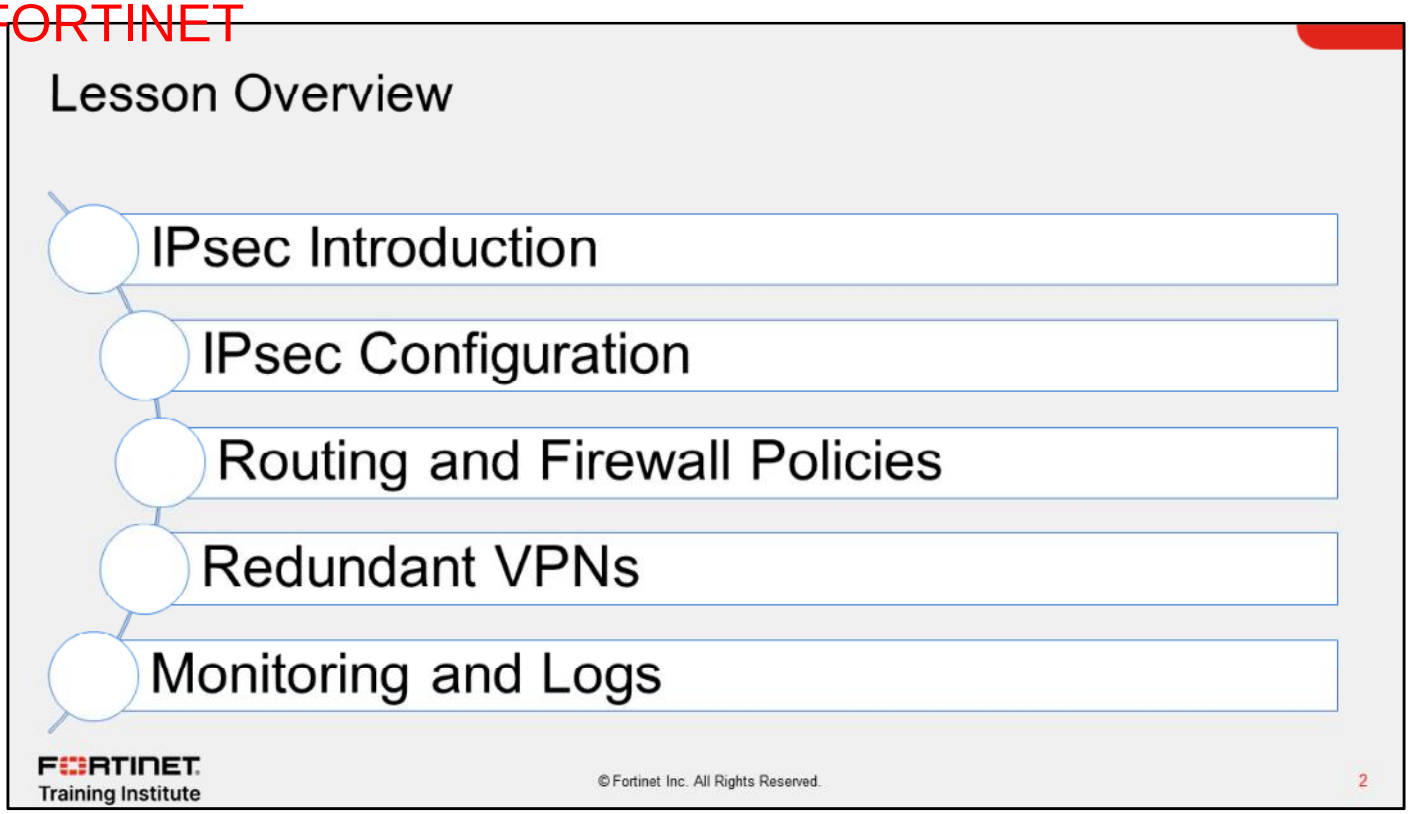

In this lesson, you will learn about the topics shown on this slide.

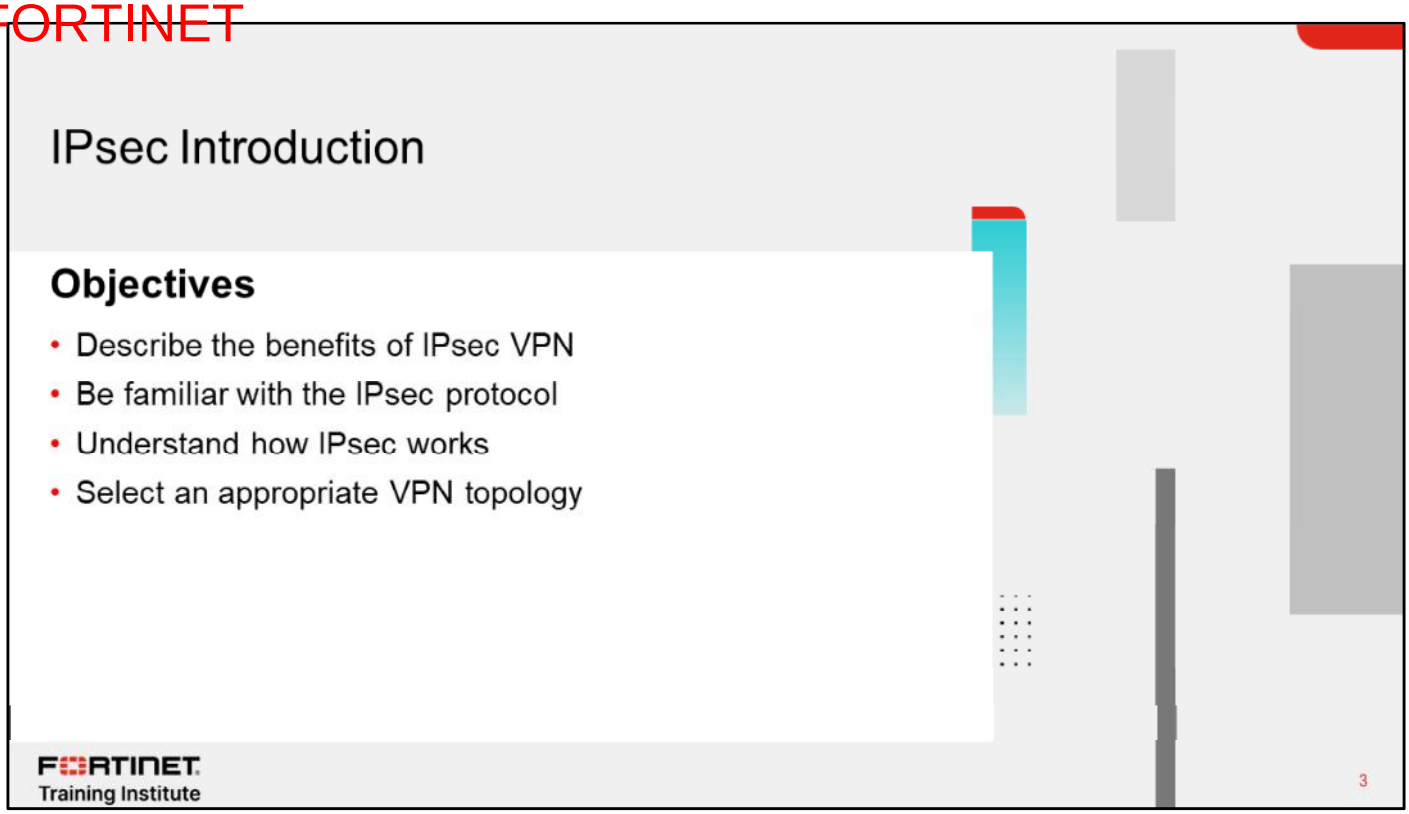

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in IPsec basics, you will be able to understand IPsec concepts and benefits.

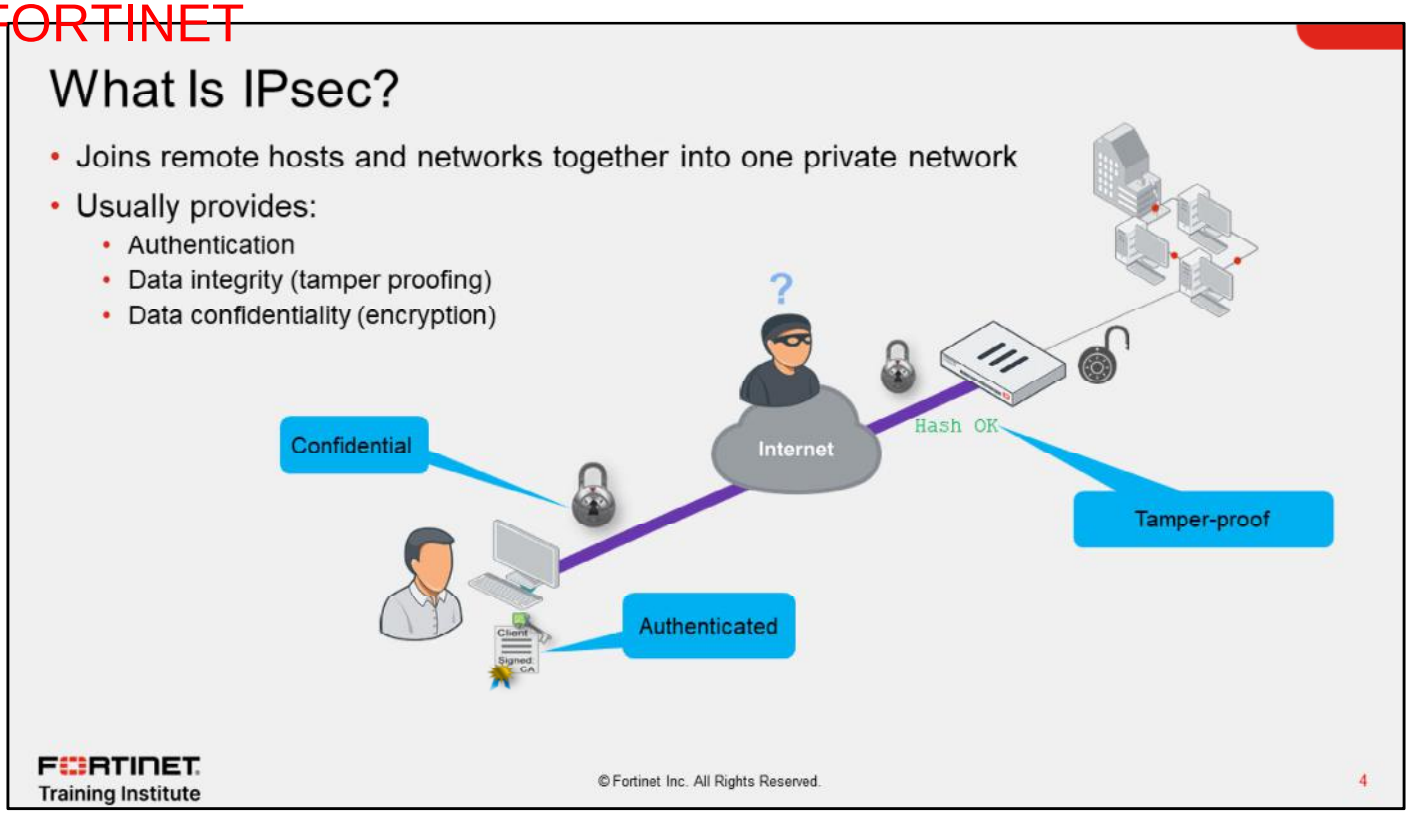

What is IPsec? When should you use it?

IPsec is a vendor-neutral set of standard protocols that is used to join two physically distinct LANs. The LANs are joined as if they were a single logical network, despite being separated by the internet.

In theory, IPsec *does* support null encryption—that is, you can make VPNs that don't encrypt traffic. IPsec also supports null data integrity. But does that provide any advantages over plain traffic? No. No one can trust traffic that may have had an attack injected by an attacker. Rarely do people want data sent by an unknown source. Most people also want private network data, such as credit card transactions and medical records, to remain private.

Regardless of the vendor, IPsec VPNs almost always have settings that allow them to provide three important benefits:

- Authentication: to verify the identity of both ends
- Data integrity (or HMAC): to prove that encapsulated data has not been tampered with as it crosses a potentially hostile network
- Confidentiality (or encryption): to make sure that only the intended recipient can read the message

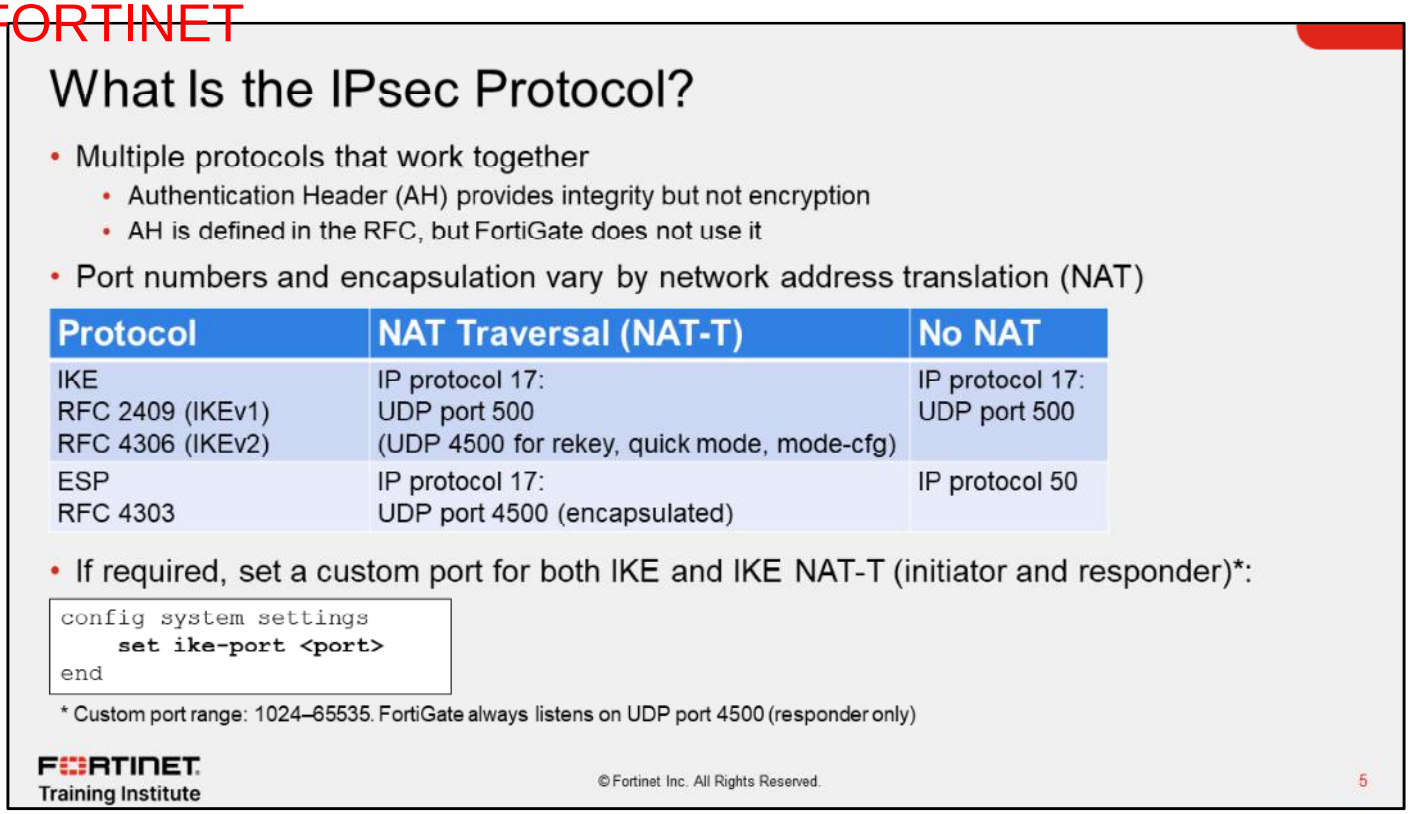

If you're passing your VPN through firewalls, it helps to know which protocols to allow.

IPsec is a suite of separate protocols, which includes:
- Internet Key Exchange (IKE): used to authenticate peers, exchange keys, and negotiate the encryption and checksums that will be used—essentially, it is the *control channel*
- AH: contains the authentication header—the checksums that verify the integrity of the data
- Encapsulating Security Payload (ESP): the encapsulated security payload—the encrypted payload, which is essentially the *data channel*

So, if you must pass IPsec traffic through a firewall, remember that allowing only one protocol or port number is usually not enough.

Note that the IPsec RFC mentions AH, however, AH does not offer encryption, which is an important benefit. Therefore, FortiGate does not use AH. As a result, you don't need to allow the AH IP protocol (51).

To set up a VPN, you must configure matching settings on both ends of the VPN—whether the VPN is between two FortiGate devices, FortiGate and FortiClient, or a third-party device and FortiGate. If the settings don't match, the tunnel setup fails.

The default ports for standard IKE traffic and IKE NAT-T traffic is UDP 500 and UDP 4500, respectively. You can use the CLI command shown on this slide to configure a custom port for both IKE and IKE NAT-T. The custom port is used to initiate and respond to tunnel requests. If NAT is detected, then the custom port can be used for both IKE and UDP-encapsulated ESP traffic. Note that FortiGate always listens for port UDP 4500 regardless of the custom port settings. This enables FortiGate to negotiate NAT-T tunnels on custom and standard ports.

How Does IPsec Work?

- Encapsulation
  - Other protocols wrapped inside IPsec
  - What's inside? Varies by mode:
    - Transport mode—TCP/UDP
    - Tunnel mode—additional IP layer, then TCP/UDP
- Negotiation
  - Authentication
  - Handshake to exchange keys, settings

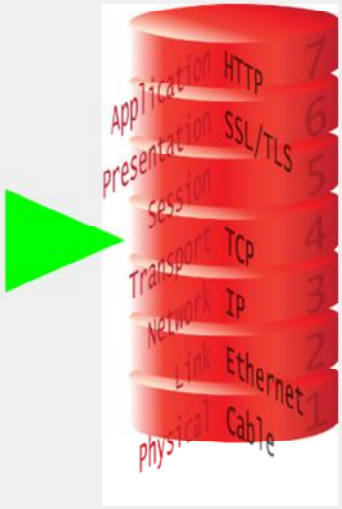

FORTINET Training Institute

6

IPsec provides services at the IP (network) layer. During tunnel establishment, both ends negotiate the encryption and authentication algorithms to use.

After the tunnel has been negotiated and is up, data is encrypted and encapsulated into ESP packets.

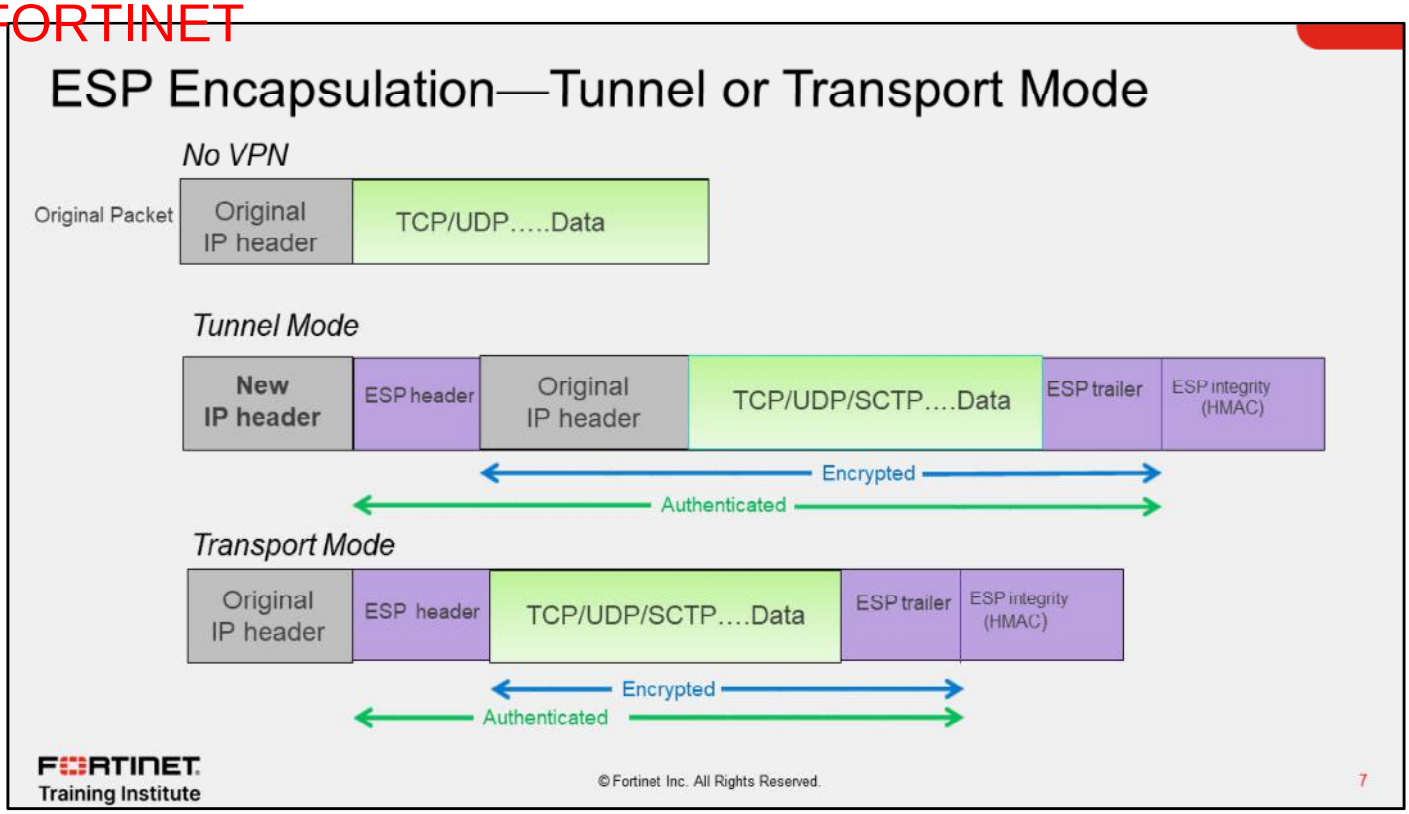

What's encapsulated? It depends on the encapsulation mode being used. IPsec can operate in two modes: transport mode and tunnel mode.

- Transport mode directly encapsulates and protects the fourth layer (transport) and above. The original IP header is not protected and no additional IP header is added.
- Tunnel mode is a true tunnel. The whole IP packet is encapsulated and a new IP header is added at the beginning. After the IPsec packet reaches the remote LAN and is unwrapped, the original packet can continue on its journey.

Note that after you remove the VPN-related headers, a transport mode packet can't be transmitted any further; it has no second IP header inside, so it's not routable. For that reason, this mode is usually used only for end-to-end (or client-to-client) VPNs.

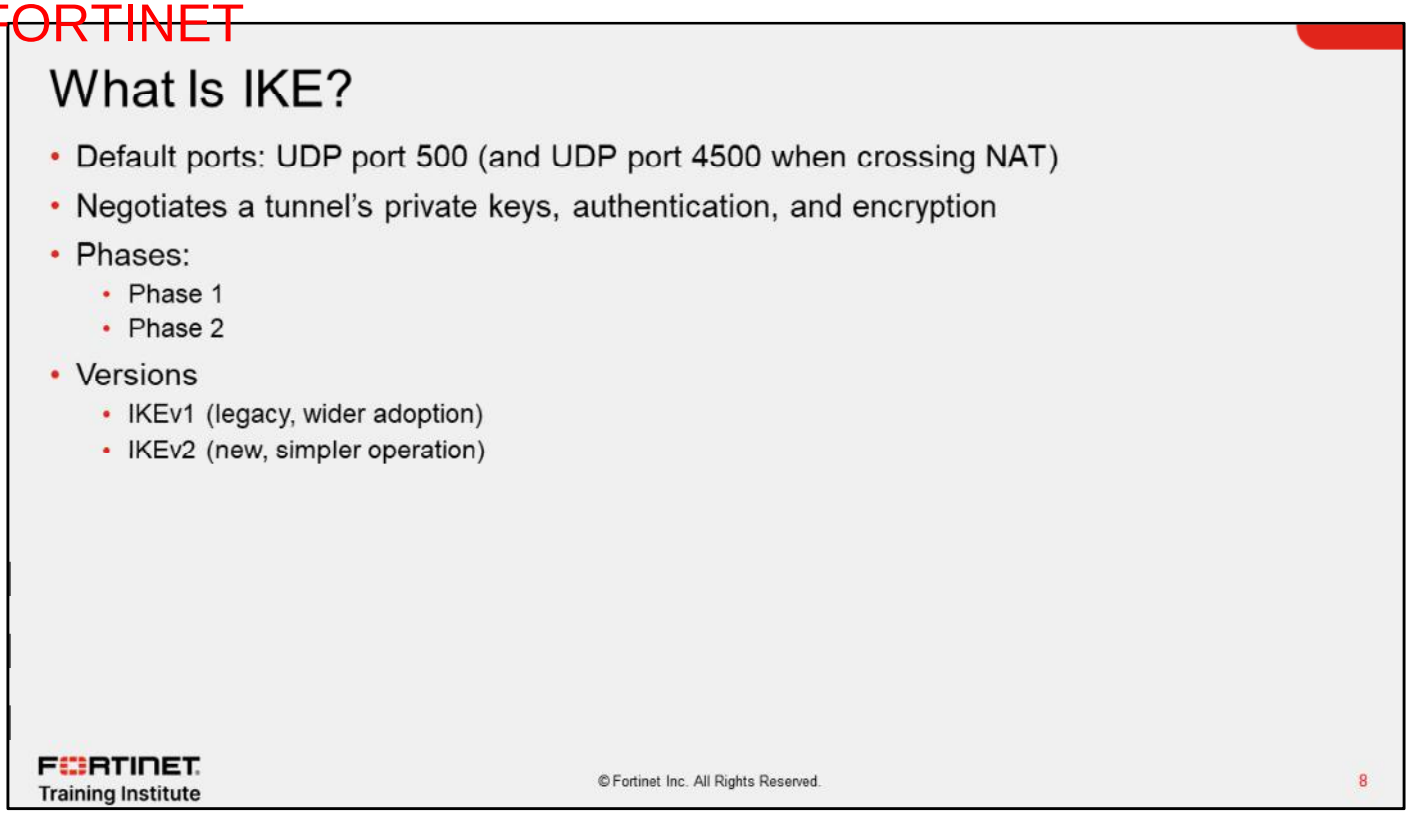

IKE uses UDP port 500. If NAT-T is enabled in a NAT scenario, IKE uses UDP port 4500.

IKE establishes an IPsec VPN tunnel. FortiGate uses IKE to negotiate with the peer and determine the IPsec security association (SA). The IPsec SA defines the authentication, keys, and settings that FortiGate uses to encrypt and decrypt that peer's packets. It is based on the Internet Security Association and Key Management Protocol (ISAKMP).

IKE defines two phases: phase 1 and phase 2.

The are two IKE versions: IKEv1 and IKEv2. Even though IKEv2 is a newer version and features a simpler protocol operation, this lesson focuses on IKEv1 only, because of its much wider adoption.

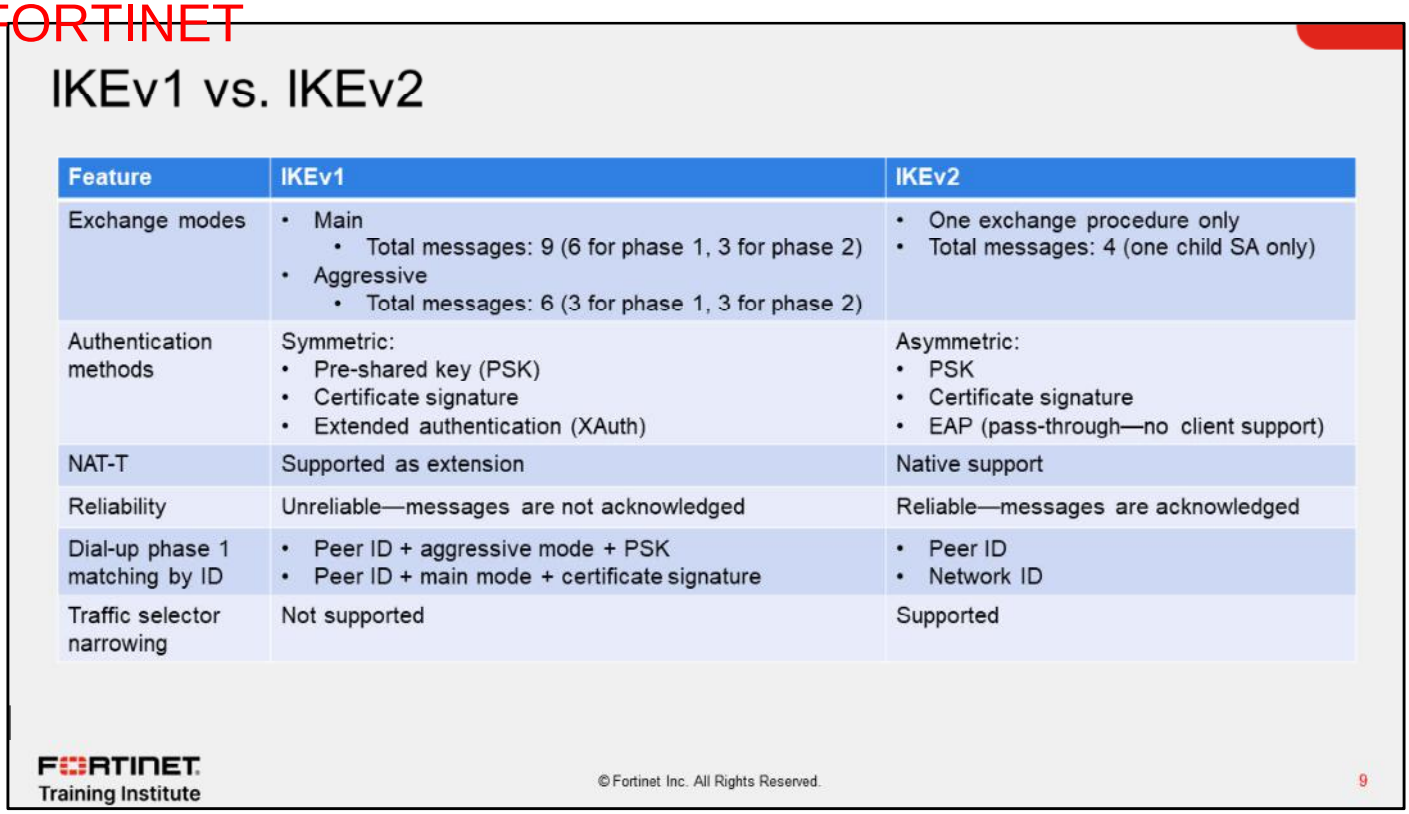

This slide shows a table comparing some of the IKEv1 and IKEv2 features that FortiOS supports. IKEv2 provides a simpler operation, which is the result of using a single exchange mode and requiring less messages to bring up the tunnel.

Authentication-wise, both versions support PSK and certificate signature. Although only IKEv1 supports XAuth, IKEv2 supports EAP, which is equivalent to XAuth. However, the FortiOS IKEv2 EAP implementation is pass-through only. That is, FortiOS doesn't support EAP as a client, which means that you cannot revoke access to peers using IKEv2 unless you use a certificate signature. With IKEv1, you can deny access to VPN peers without having to use a certificate signature by using XAuth. IKEv2 also supports asymmetric authentication, which enables you to configure each peer to use a different authentication method.

Both IKE versions support NAT-T. However, IKEv2 supports NAT-T natively, while IKEv1 supports NAT-T as an extension. Also, IKEv2 is a more reliable protocol than IKEv1 because, like TCP, peers must acknowledge the messages exchanged between them. IKEv1 doesn't support such a mechanism.

When you configure multiple dial-up IPsec VPNs, IKEv2 makes it simpler to match the intended gateway by peer ID. With IKEv2, you can either use the standard peer ID attribute or the Fortinet proprietary network ID attribute to indicate the phase 1 gateway to match on the dial-up server, regardless of the authentication mode in use. However, with IKEv1, you can use the peer ID only, and then combine it with aggressive mode and pre-shared key authentication, or with main mode and certificate signature authentication.

Finally, IKEv2 allows the responder to choose a subset of the traffic the initiator proposes. This is called traffic selector narrowing and enables you to have more flexible phase 2 selector configurations. Traffic selector narrowing enables a peer to automatically narrow down its traffic selector addresses, so it agrees with the traffic selector the remote peer proposes.

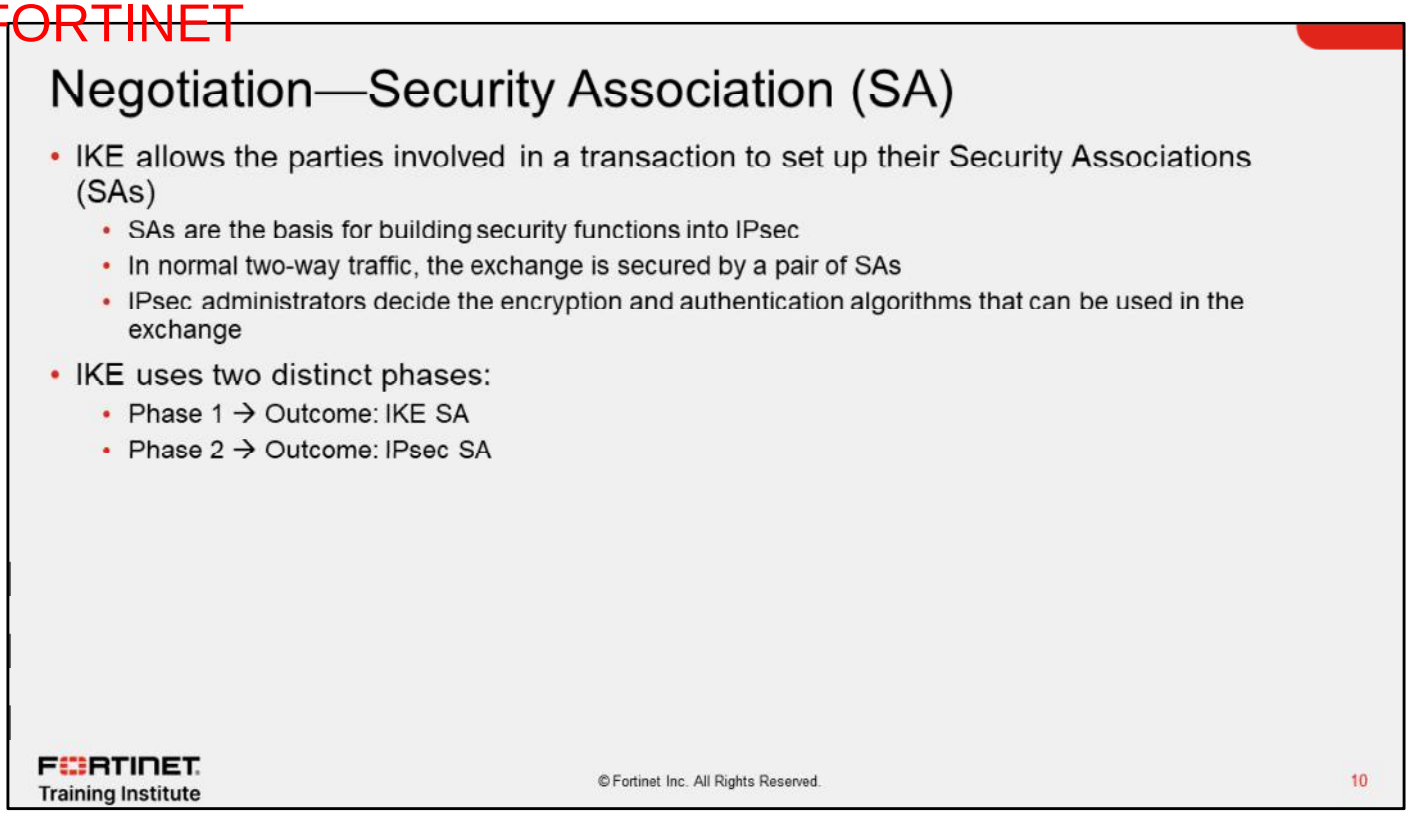

In order to create an IPsec tunnel, both devices must establish their SAs and secret keys, which are facilitated by the IKE protocol.

The IPsec architecture uses SAs as the basis for building security functions into IPsec. An SA is the bundle of algorithms and parameters being used to encrypt and authenticate data travelling through the tunnel. In normal two-way traffic, this exchange is secured by a pair of SAs, one for each traffic direction. Essentially, both sides of the tunnel must agree on the security rules. If both sides cannot agree on the rules for sending data and verifying each other's identity, then the tunnel is not established. SAs expire and need to be renegotiated by the peers after they have reached their lifetime.

IKE uses two distinct phases: phase 1 and phase 2. Each phase negotiates different SA types. The SA negotiated during phase 1 is called IKE SA, and the SA negotiated during phase 2 is called IPsec SA. FortiGate uses IKE SAs for setting up a secure channel to negotiate IPsec SAs. FortiGate uses IPsec SAs for encrypting and decrypting the data sent and received, respectively, through the tunnel.

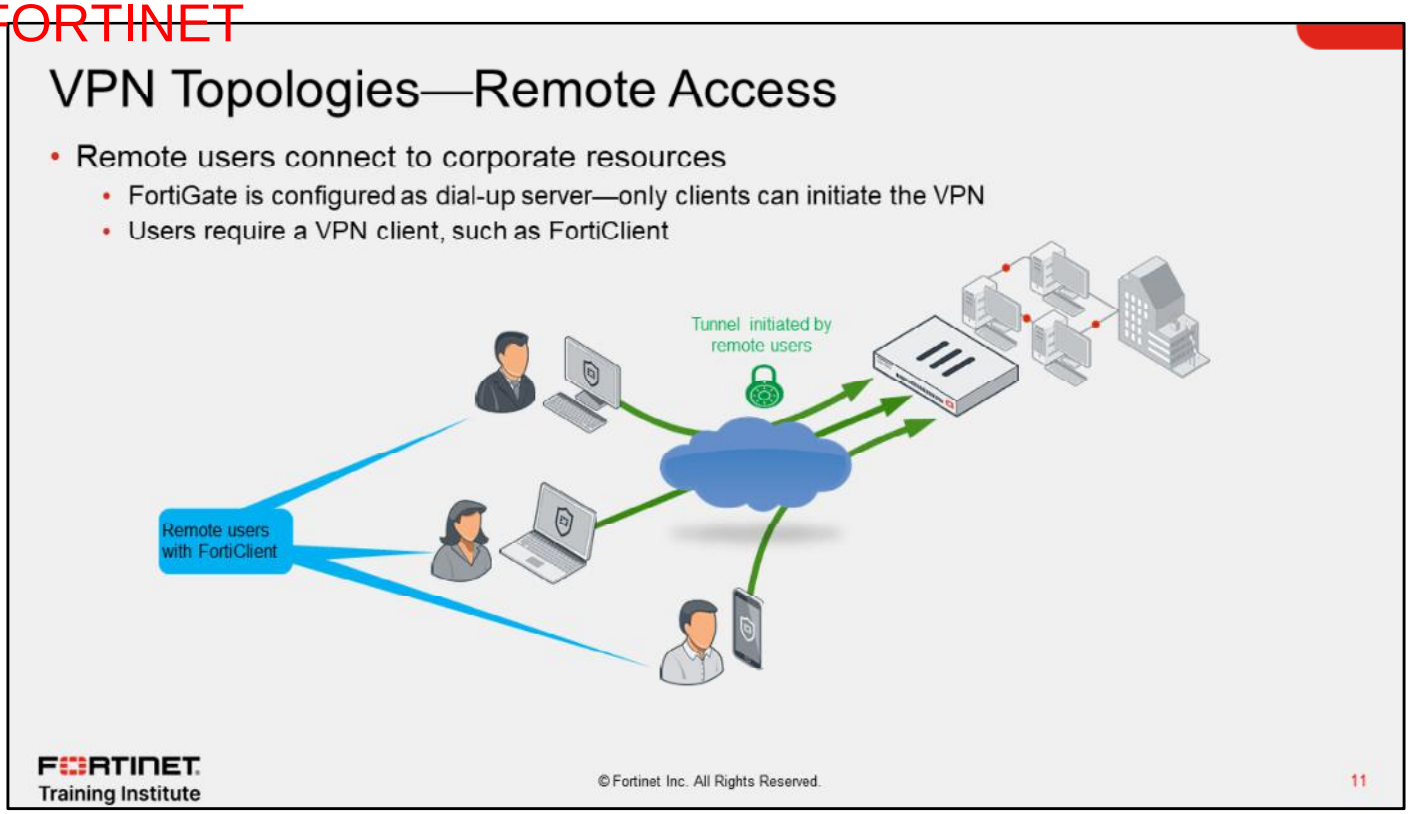

Use remote access VPNs when remote internet users need to securely connect to the office to access corporate resources. The remote user connects to a VPN server located on the corporate premises, such as FortiGate, to establish a secure tunnel. After the user is authenticated, FortiGate provides access to network resources, based on the permissions granted to that user.

In a remote access VPN, FortiGate is usually configured as a dial-up server. You will learn more about dial-up VPNs in this lesson. The IP address of the remote internet user is usually dynamic. Because FortiGate does not know the IP address of the remote user, only the remote user can initiate a VPN connection request.

The remote user side needs a VPN client, such as FortiClient. You must configure FortiClient to match the VPN server settings. FortiClient takes care of establishing the tunnel, as well as routing the traffic destined to the remote site through the tunnel.

In addition, you can use one remote access VPN configuration on your FortiGate device for many remote users. FortiGate establishes a separate tunnel for each of them.

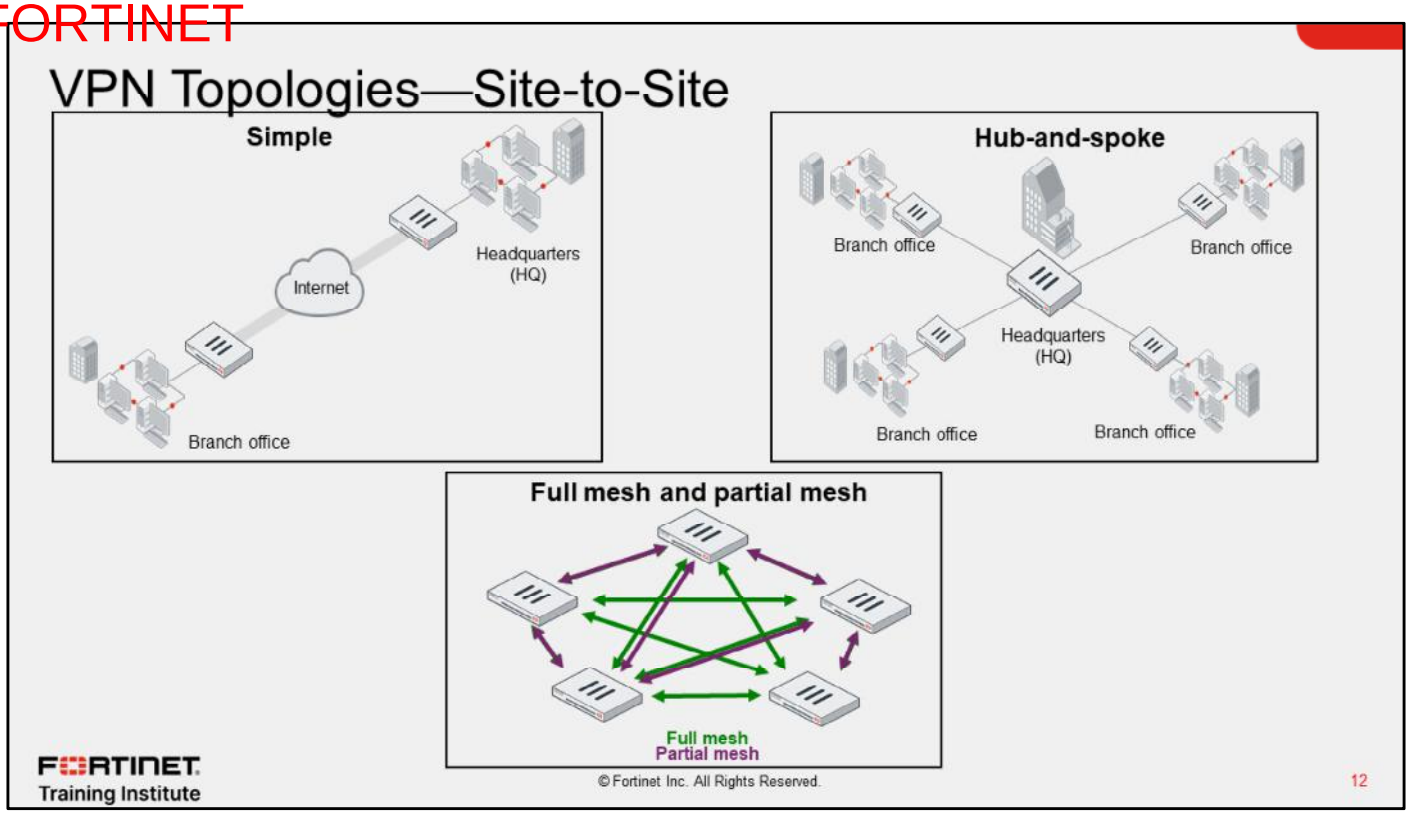

Site-to-site VPN is also known as LAN-to-LAN VPN. A simple site-to-site deployment involves two peers communicating directly to connect two networks located at different offices.

When you need to connect more than two locations, you can use a hub-and-spoke topology. In hub-and-spoke, all clients connect through a central hub. In the example shown on this slide, the clients—spokes—are branch office FortiGate devices. For any branch office to reach another branch office, its traffic must pass through the hub. One advantage of this topology is that the configuration needed is easy to manage. Another advantage is that only the FortiGate at HQ must be very powerful because it handles all tunnels simultaneously, while the branch office FortiGate devices require much fewer resources because they maintain only one tunnel. One disadvantage is that communication between branch offices through HQ is slower than in a direct connection, especially if your HQ is physically distant. Also, if the FortiGate device at HQ fails, VPN failure is company-wide.

In a mesh topology, you can connect FortiGate devices directly and therefore bypass HQ. Two variations of mesh topology exist: full mesh and partial mesh. Full mesh connects every location to every other location. The higher the number of FortiGate devices, the higher the number of tunnels to configure on each FortiGate device. For example, in a topology with five FortiGate devices, you would need to configure four tunnels on each device, for a total of 20 tunnels. This topology causes less latency and requires much less HQ bandwidth than hub-and-spoke, but requires each FortiGate device to be more powerful. Partial mesh attempts to compromise, minimizing required resources but also latency. Partial mesh can be appropriate if communication is not required between every location. However, the configuration of each FortiGate device is more complex than in hub-and-spoke. Routing, especially, may require extensive planning.

Generally, the more locations you have, hub-and-spoke will be cheaper, but slower, than a mesh topology. Mesh places less strain on the central location. It’s more fault-tolerant, but also more expensive.

VPN Topologies—Comparison

Hub-and-Spoke	Partial Mesh	Full Mesh
Easy configuration	Moderate configuration	Complex configuration
Few tunnels	Medium number of tunnels	Many tunnels
High central bandwidth	Medium bandwidth in hub sites	Low bandwidth
Not fault tolerant	Some fault tolerance	Fault tolerant
Low system requirements on average, but high for center	Medium system requirements	High system requirements
Scalable	Somewhat scalable	Difficult to scale
No direct communication between spokes	Direct communication between some sites	Direct communication between all sites

FORTINET Training Institute

13

To review, this slide shows a high-level comparison of VPN topologies. You should choose the topology that is most appropriate to your situation.

Auto-Discovery VPN

- Dynamically negotiates on-demand direct VPNs between spokes
  - Provides the benefits of a full mesh topology over a hub-and-spoke or partial mesh deployment
  - Dynamic routing is recommended to learn routes between hub and spokes and scale up easier
  - Static routing also works, but should be used for small deployments only

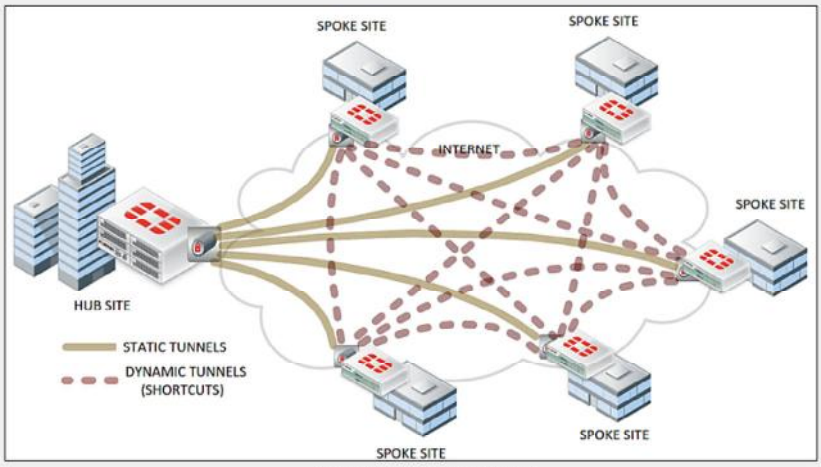

FORTINET Training Institute

14

Each VPN topology has its advantages and disadvantages.

Auto-discovery VPN (ADVPN) is a FortiGate feature that achieves the benefits of a full-mesh topology with the easier configuration and scalability benefits of hub-and-spoke and partial-mesh topologies.

First, you add the VPN configurations for building either a hub-and-spoke or a partial-mesh topology, to the FortiGate devices. Then, you enable ADVPN on the VPNs. ADVPN dynamically negotiates tunnels between spokes (without having them preconfigured) to get the benefits of a full-mesh topology.

You can use dynamic routing and static routing to deploy ADVPN. A dynamic routing protocol, such as BGP, is usually deployed in large networks because it enables you to exchange routing information between spokes and hub easier, and as a result, to scale up. You can also use static routing to deploy ADVPN, but it is recommended to do so in small networks that are not likely to grow considerably.

Whether you use dynamic routing or not, after a shortcut is negotiated, FortiGate automatically adds routes through the shortcut to redirect spoke-to-spoke traffic through it.

Knowledge Check

1. Which IPsec protocol is not supported by FortiGate?
   - A. IKEv2
   - ✓ B. AH

2. Which VPN topology is the most fault tolerant?
   - ✓ A. Full mesh
   - B. Hub-and-spoke

FORTINET Training Institute

15

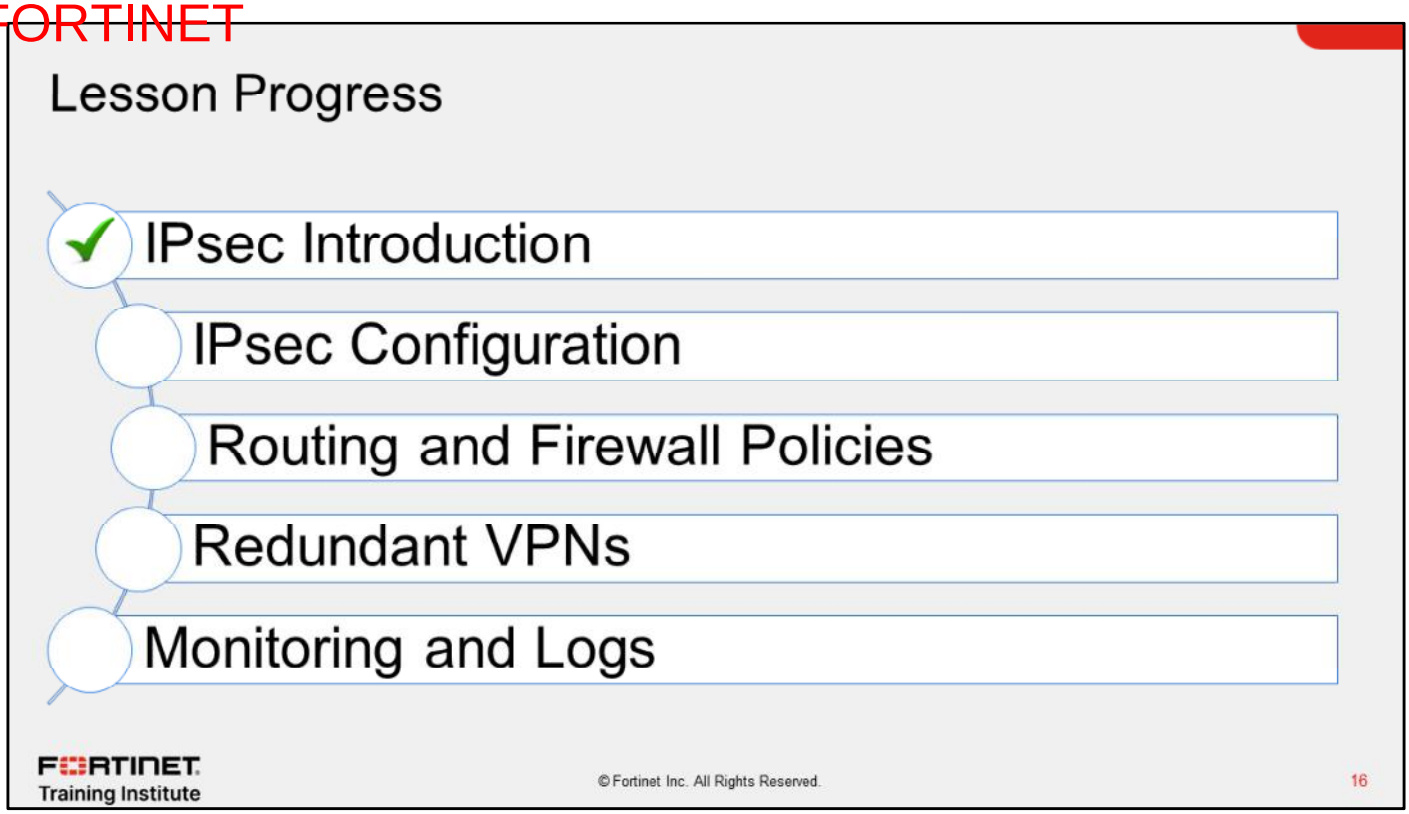

Good job! You have now been introduced to IPsec.

Now, you will learn about IPsec configuration.

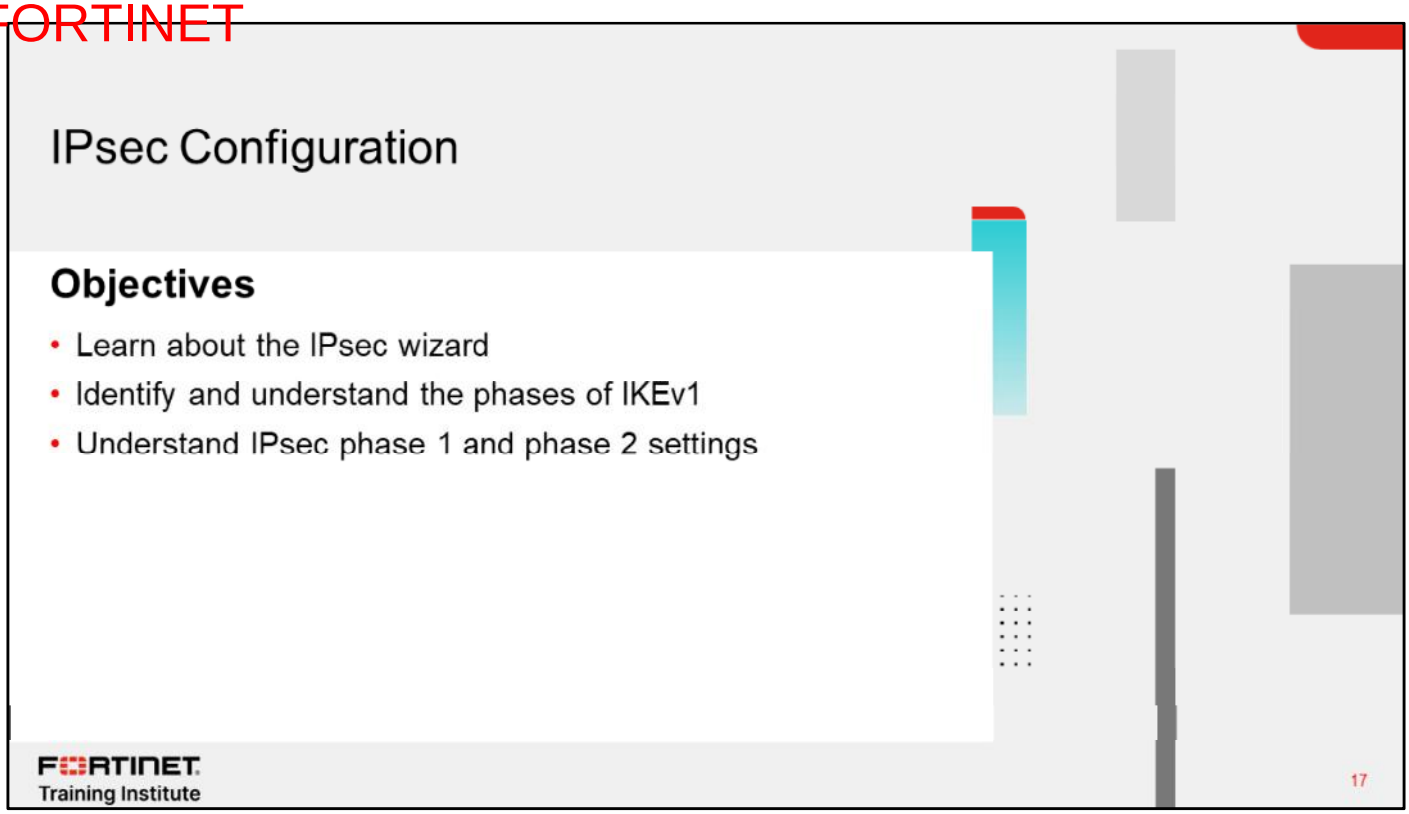

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in IPsec configuration, you will be able to successfully determine the settings required for your IPsec VPN deployment.

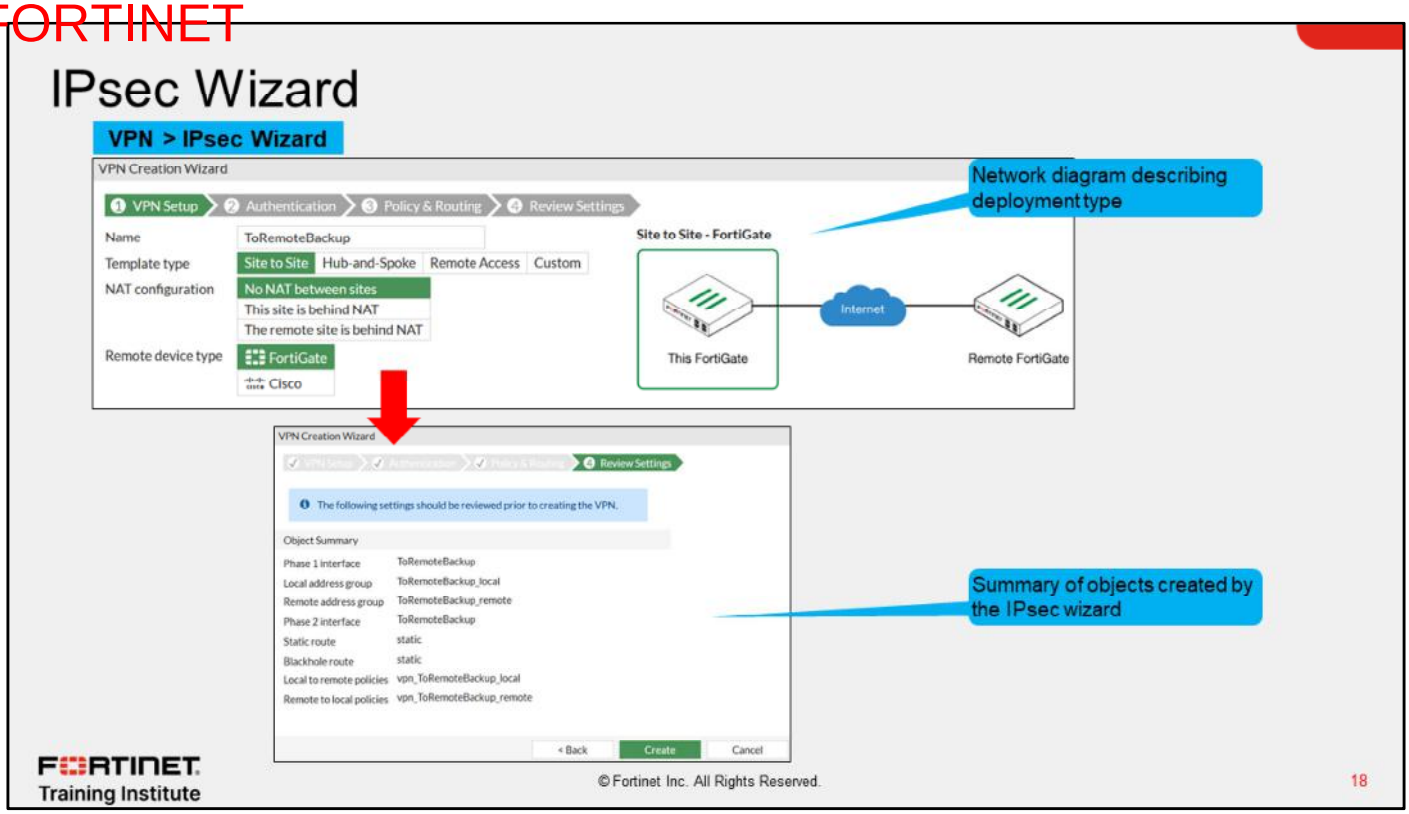

When you create an IPsec tunnel on the GUI, FortiGate redirects you to the **IPsec Wizard**. The wizard simplifies the creation of the new VPN by walking you through a four to five-step process. The first step is to select a template type. If you want to manually configure your VPN, you can select **Custom** as **Template type**, upon which FortiGate takes you directly to the phase 1 and phase 2 settings of the new VPN.

If you want the wizard to configure the VPN for you, then select the template type (**Site to Site**, **Hub-and-Spoke**, or **Remote Access**) that best matches your VPN. After that, the wizard asks you for key information such as the remote gateway information, authentication method, interfaces involved, and subnets. Based on the input you provide, the wizard applies one of the preconfigured IPsec tunnel templates comprising IPsec phase 1 and 2 settings and other related firewall address objects, routing settings, and firewall policies needed for the new tunnel to work.

In addition, the wizard shows a network diagram that changes based on the input provided. The purpose of the diagram is for the administrator to have a visual understanding of the IPsec VPN deployment that the wizard configures based on the input received.

At the end of the wizard, the wizard provides a summary of the configuration changes made in the system, and that the administrator can review if needed.

If you are new to FortiGate, or don't have much experience with IPsec VPNs, using the IPsec wizard is recommended. You can later adjust the configuration applied by the wizard to match your specific needs.

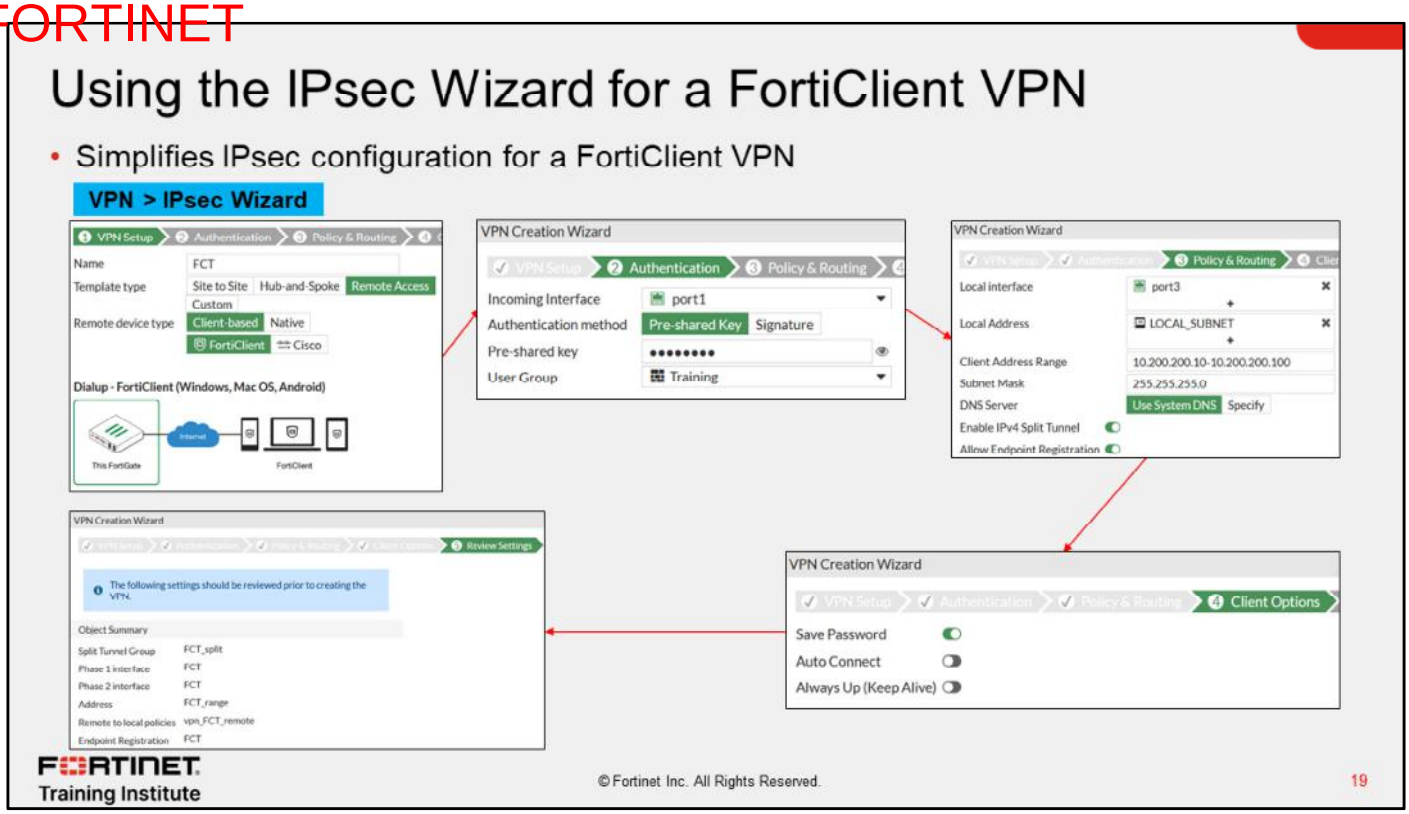

A common use of the IPsec wizard is for configuring a remote access VPN for FortiClient users. The wizard enables IKE mode config, XAuth, and other appropriate settings for FortiClient users. You will learn more about IKE mode config and XAuth in this lesson.

The images on this slide show the four-step process used by the IPsec wizard for assisting the administrator on the FortiClient VPN configuration.

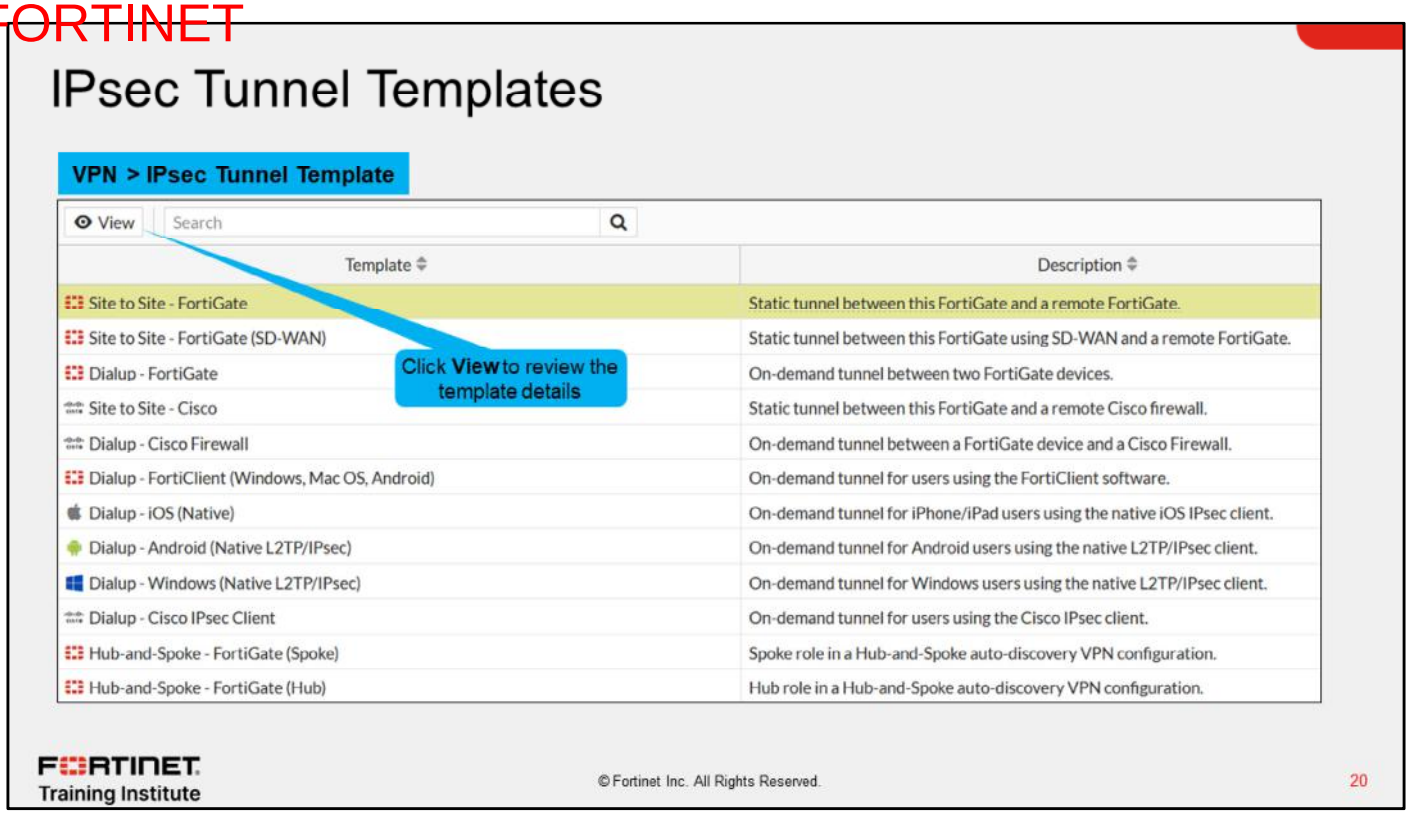

The IPsec wizard uses one of the templates shown on this slide when applying the configuration for the new IPsec tunnel. You can review the settings of a template by selecting the template, and then clicking **View**. You cannot change the template settings.

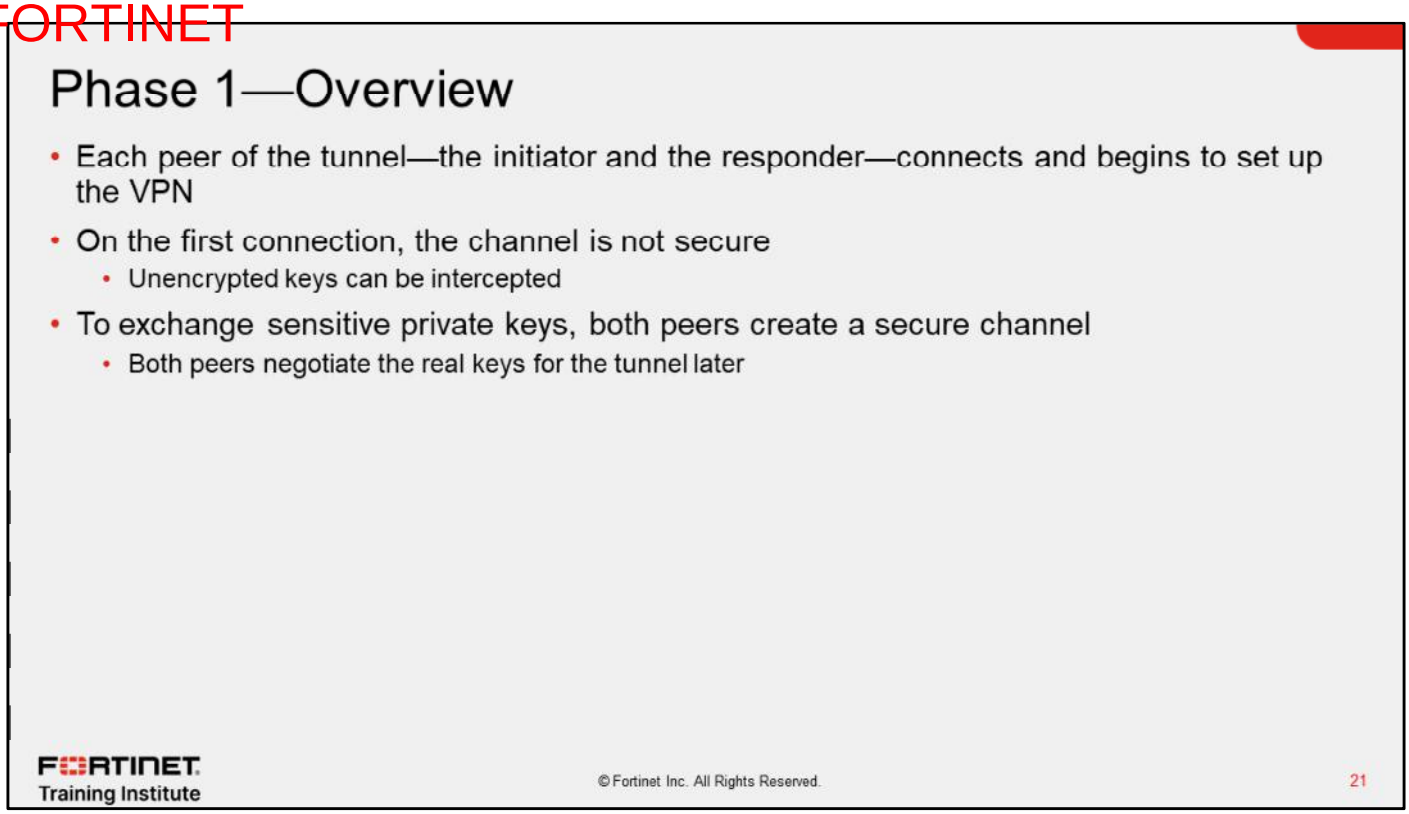

Phase 1 takes place when each peer of the tunnel—the initiator and the responder—connects and begins to set up the VPN. The initiator is the peer that starts the phase 1 negotiation, while the responder is the peer that responds to the initiator request.

When the peers first connect, the channel is not secure. An attacker in the middle could intercept unencrypted keys. Neither peer has a strong guarantee of the other peer's identity, so how can they exchange sensitive private keys? They can't. First, both peers create a secure tunnel. Then, they use this secure tunnel to negotiate the real keys for the tunnel later.

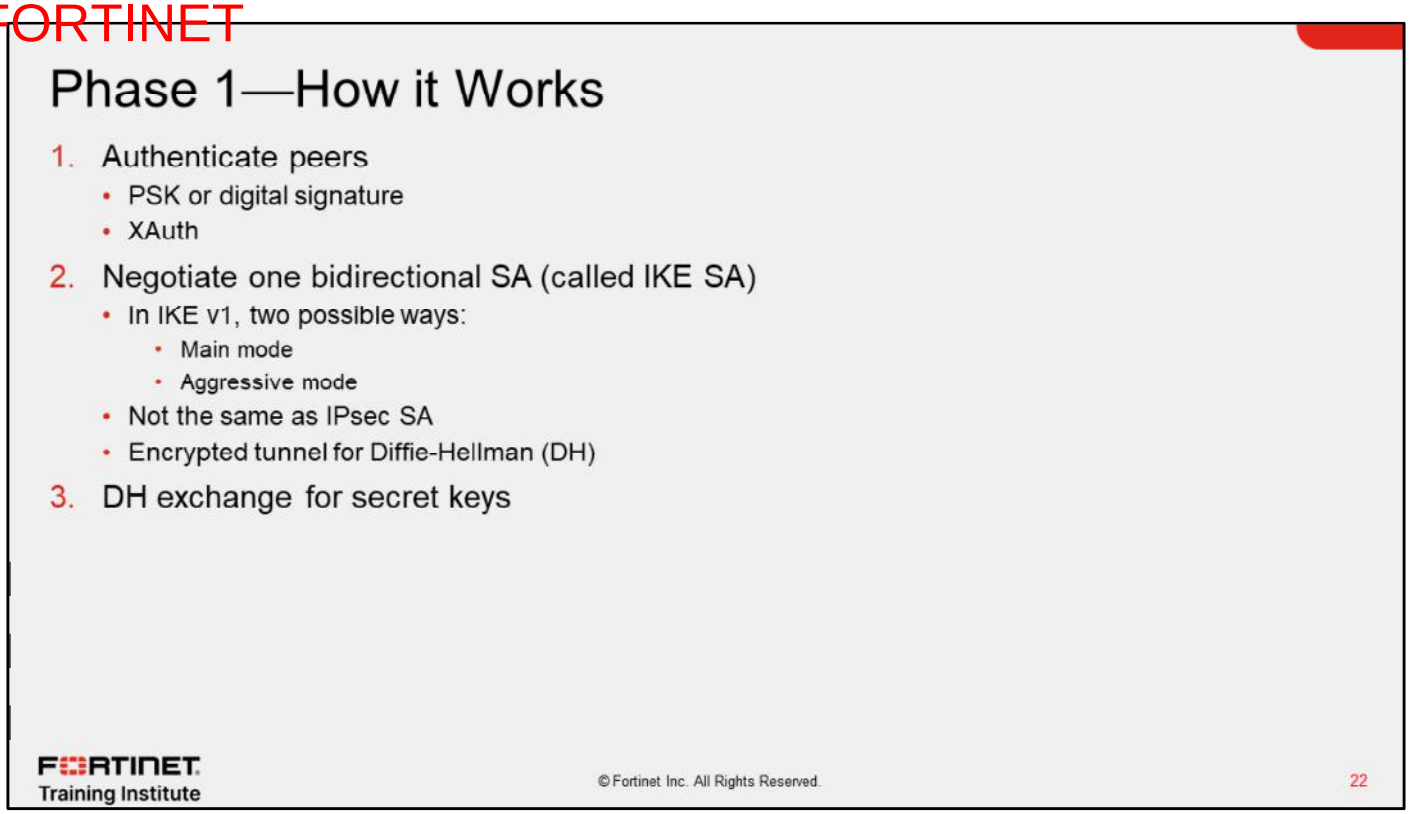

Now you'll examine how phase 1 works.

The purpose of phase 1 is to authenticate peers and set up a secure channel for negotiating the phase 2 SAs (or IPsec SAs) that are later used to encrypt and decrypt traffic between the peers. To establish this secure channel, the peers negotiate a phase 1 SA. This SA is called the IKE SA and is bidirectional.

To authenticate each other, the peers use two methods: pre-shared key or digital signature. You can also enable an additional authentication method, XAuth, to enhance authentication.

In IKEv1, there are two possible modes in which the IKE SA negotiation can take place: main, and aggressive mode. Settings on both ends must agree; otherwise, phase 1 negotiation fails and both IPsec peers are not able to establish a secure channel.

At the end of phase 1, the negotiated IKE SA is used to negotiate the DH keys that are used in phase 2. DH uses the public key (that both ends know) plus a mathematical factor called a nonce, in order to generate a common private key. With DH, even if an attacker can listen to the messages containing the public keys, they cannot determine the secret key.

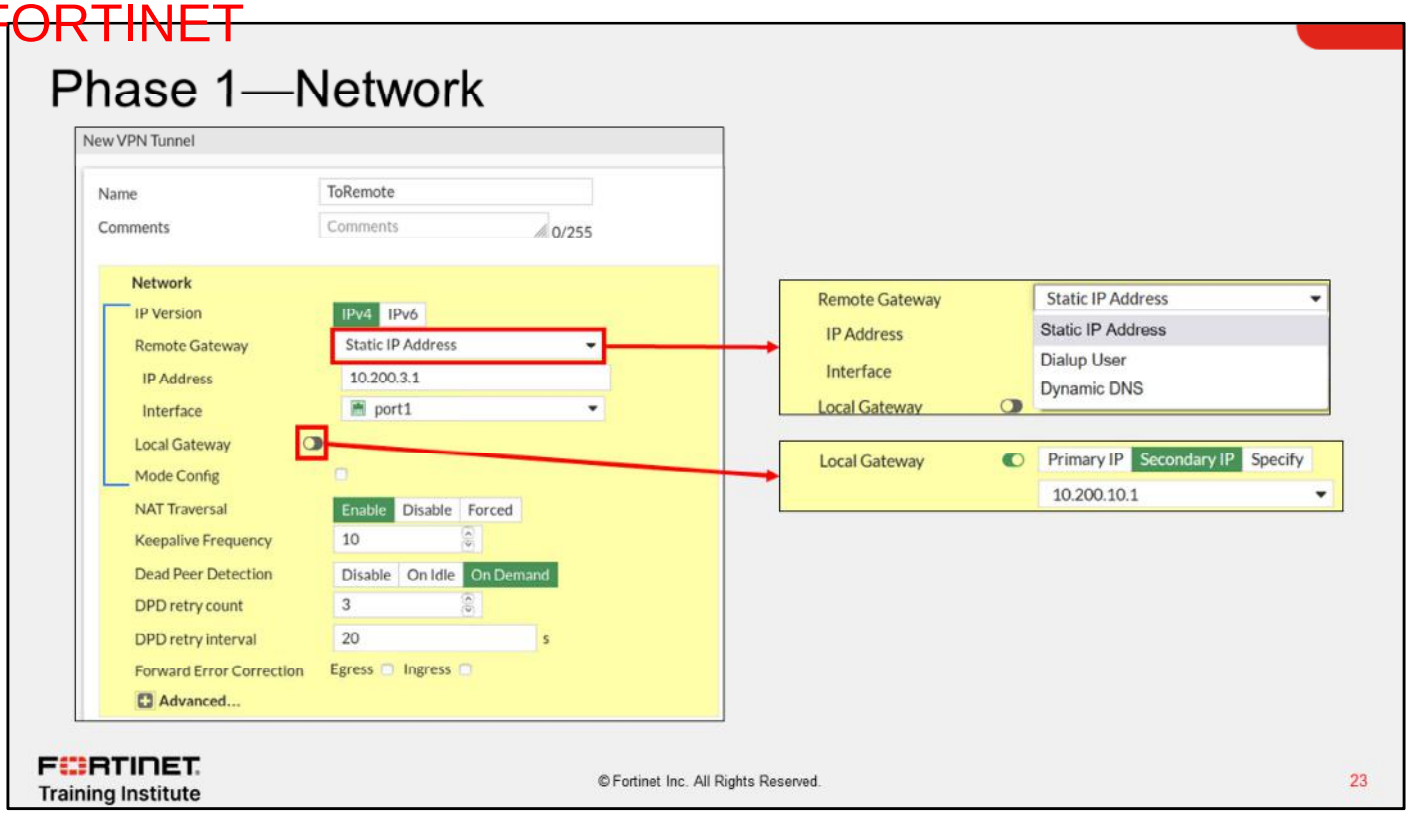

Phase 1 configuration is broken down on the GUI into four sections: **Network**, **Authentication**, **Phase 1 Proposal**, and **XAUTH**. You will learn about the settings available on each section. You will learn about some of these settings in more detail on separate slides.

The section shown on this slide corresponds to the **Network** settings. The section includes the settings related to the connectivity of the IPsec tunnel:

- **IP Version**: select the IP version to use for the IPsec tunnel. Note that this defines only the IP version of the outer layer of the tunnel (after encapsulation). The packets being encapsulated (protected traffic) can be IPv4 or IPv6, and their IP version is defined in the phase 2 selectors.
- **Remote Gateway**: defines the type of the remote gateway. There are three types: **Static IP Address**, **Dialup User**, and **Dynamic DNS**. You will learn more about these types later in this lesson.
- **IP Address**: the IP address of the remote gateway. This field appears only when you select **Static IP Address** as **Remote Gateway**.
- **Interface**: refers to the interface where the IPsec tunnel terminates on the local FortiGate. Usually, this is the interface connected to the internet or the WAN. You need to make sure there is an active route to the remote gateway through this interface, otherwise the tunnel won't come up.
- **Local Gateway**: enable this setting when the interface where the tunnel terminates has multiple addresses assigned, and you want to specify which address to use for the tunnel. When you enable this setting, you see three options: **Primary IP**, **Secondary IP**, and **Specify**. Select **Specify** if you want to use an IP address different from the primary or secondary IP address.
- **Mode Config**: Enables automatic configuration through IKE. FortiGate acts as an *IKE mode config client* when you enable **Mode Config** and you set **Remote Gateway** to either **Static IP address** or **Dynamic DNS**. If you set **Remote Gateway** to **Dialup User**, FortiGate acts as an *IKE mode config server,* and more configuration options appear on the GUI. You will learn more about **Mode Config** in this lesson.

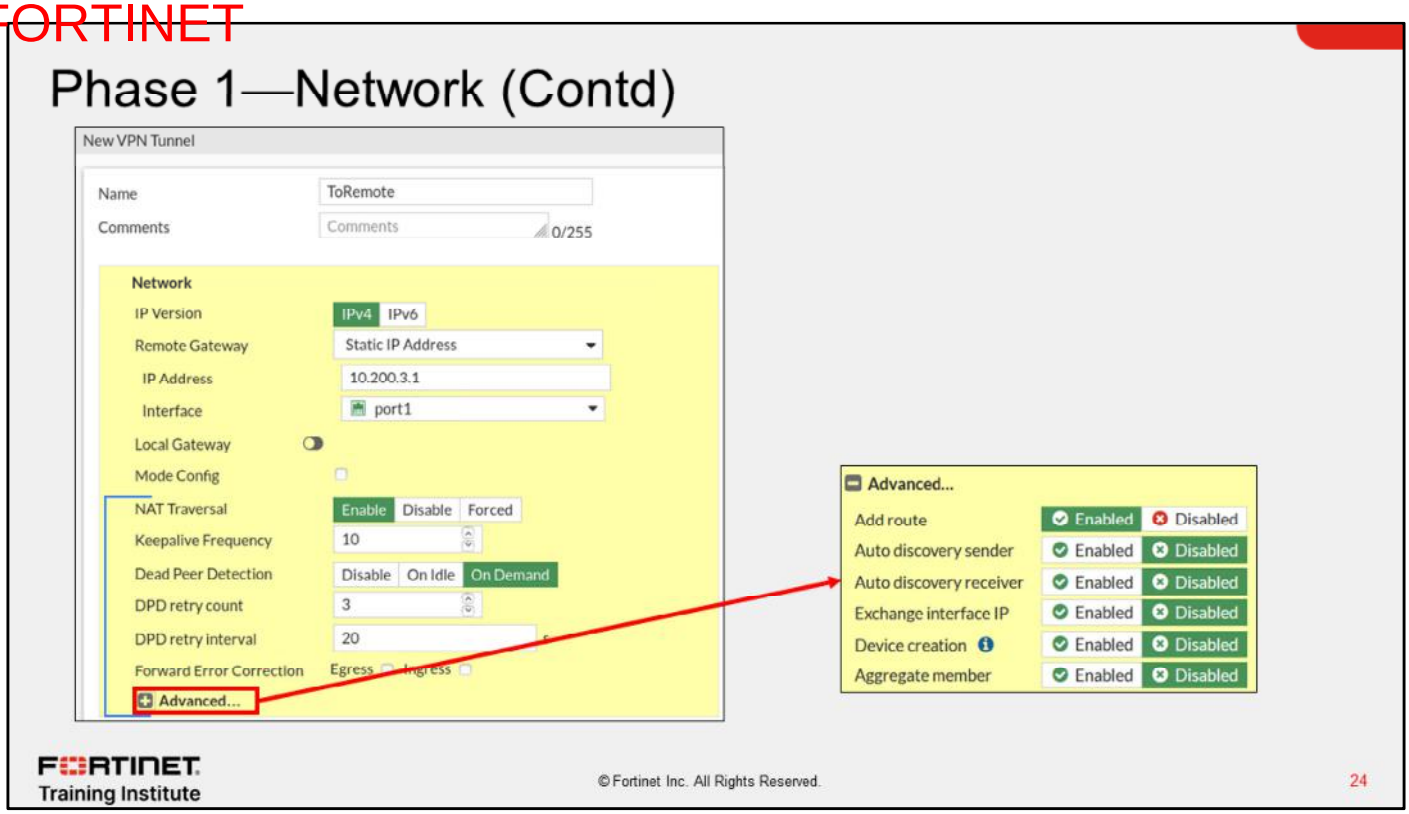

The following are the other options available on the GUI in the **Network** section:

- **NAT Traversal**: The option controls the behavior for NAT traversal. You will learn more about NAT traversal later in this lesson.
- **Keepalive Frequency**: When you enable NAT traversal, FortiGate sends keepalive probes at the configured frequency.
- **Dead Peer Detection**: Use dead peer detection (DPD) to detect dead tunnels. There are three DPD modes. **On Demand** is the default mode. You will learn more about DPD later in this lesson.
- **Forward Error Correction**: Forward error correction (FEC) is a technique that you can use to reduce the number of retransmissions in IPsec tunnels established over noisy links, at the expense of using more bandwidth. You can enable FEC on egress and ingress, and it is only supported when you disable IPsec hardware offloading. You will learn more about IPsec hardware offloading later in this lesson.
- **Advanced**:
  - **Add route**: Disable this setting if you are using a dynamic routing protocol over IPsec and do not want FortiGate to automatically add static routes.
  - **Auto discovery sender**: Enable this setting on a hub if you want the hub to facilitate ADVPN shortcut negotiation for spokes. When enabled, the hub sends a shortcut offer to the spoke to indicate that it can establish a shortcut to the remote spoke.
  - **Auto discovery receiver**: Enable this setting on a spoke if you want the spoke to negotiate an ADVPN shortcut.
  - **Exchange interface IP**: Enable this setting to allow the exchange of IPsec interface IP addresses. This allows a point-to-multipoint connection between the hub and spokes..
  - **Device creation**: Enable this setting to instruct FortiOS to create an interface for every dial-up client. To increase performance, disable this setting in dial-up servers with many dial-up clients.
  - **Aggregate member**: FortiGate allows you to aggregate multiple IPsec tunnels into a single interface. Enable this option if you want the tunnel to become an aggregate member.

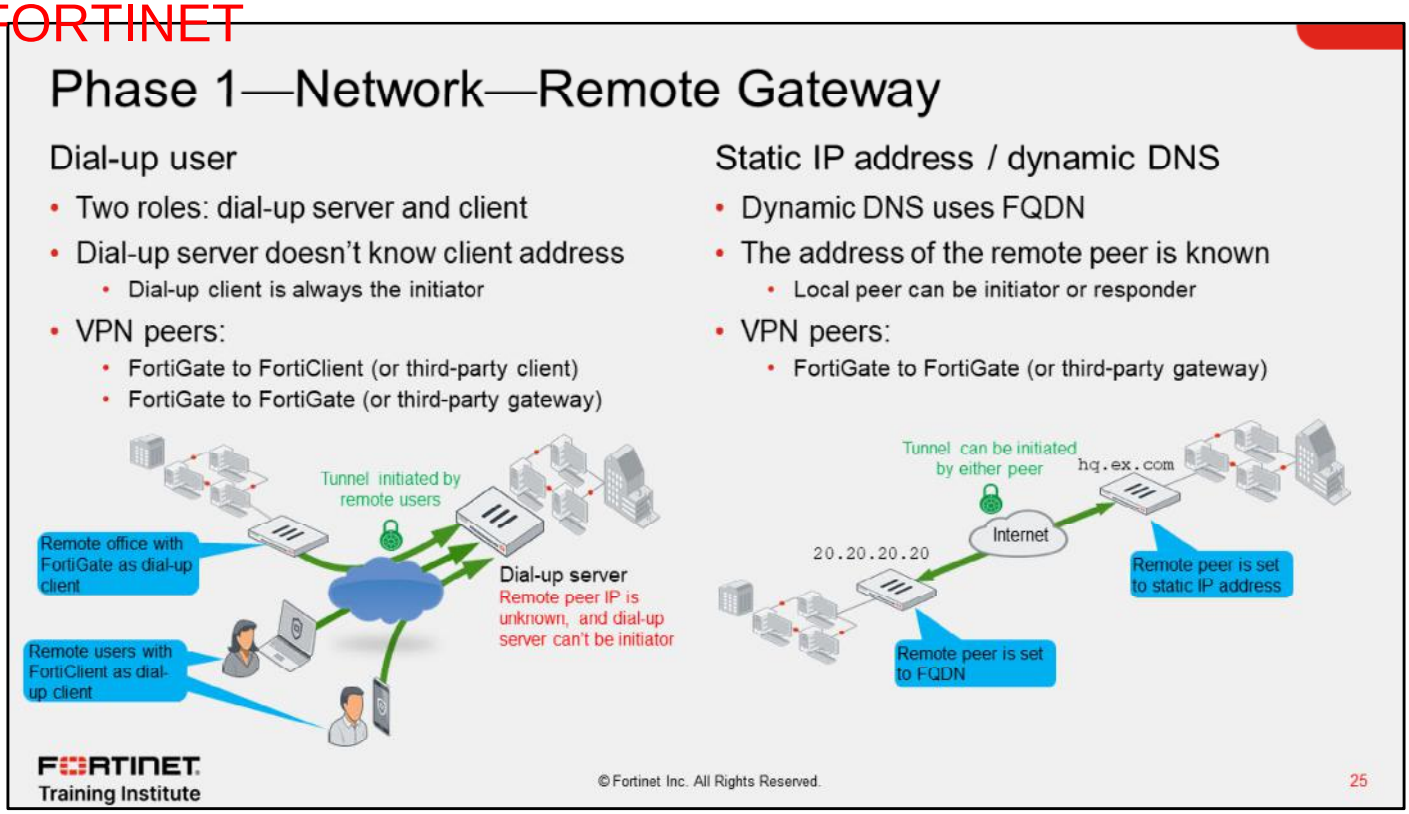

You have three options when configuring the remote gateway type of your VPN: **Dialup User**, **Static IP Address**, and **Dynamic DNS**.

Use **Dialup User** when the remote peer IP address is unknown. The remote peer whose IP address is unknown acts as the dial-up client, and this is often the case for branch offices and mobile VPN clients that use dynamic IP addresses, and no dynamic DNS. The dial-up client must know the IP address or FQDN of the remote gateway, which acts as the dial-up server. Because the dial-up server doesn't know the remote peer address, only the dial-up client can initiate the VPN tunnel.

Usually, dial-up clients are remote and mobile employees with FortiClient on their computer or handheld devices. You can also have a FortiGate device acting as a dial-up client for a remote office. One dial-up server configuration on FortiGate can be used for multiple IPsec tunnels from many remote offices or users.

Use **Static IP Address** or **Dynamic DNS** when you know the remote peer address. If you select **Static IP Address**, then you need to provide an IP address. If you select **Dynamic DNS**, then you need to provide a fully qualified domain name (FQDN), and make sure FortiGate can resolve that FQDN. When both peers know the remote peer address, that is, the remote gateway on both peers is set to **Static IP Address** or **Dynamic DNS**, then any peer can initiate the VPN tunnel.

Note that in a dial-up setup, the dial-up client is just a VPN peer with the remote gateway set to **static IP address** or **dynamic DNS**. When setting your VPN, you can combine different types of remote gateways. For obvious reasons, a tunnel in which both peers has the remote gateway set to **Dialup user** won't work.

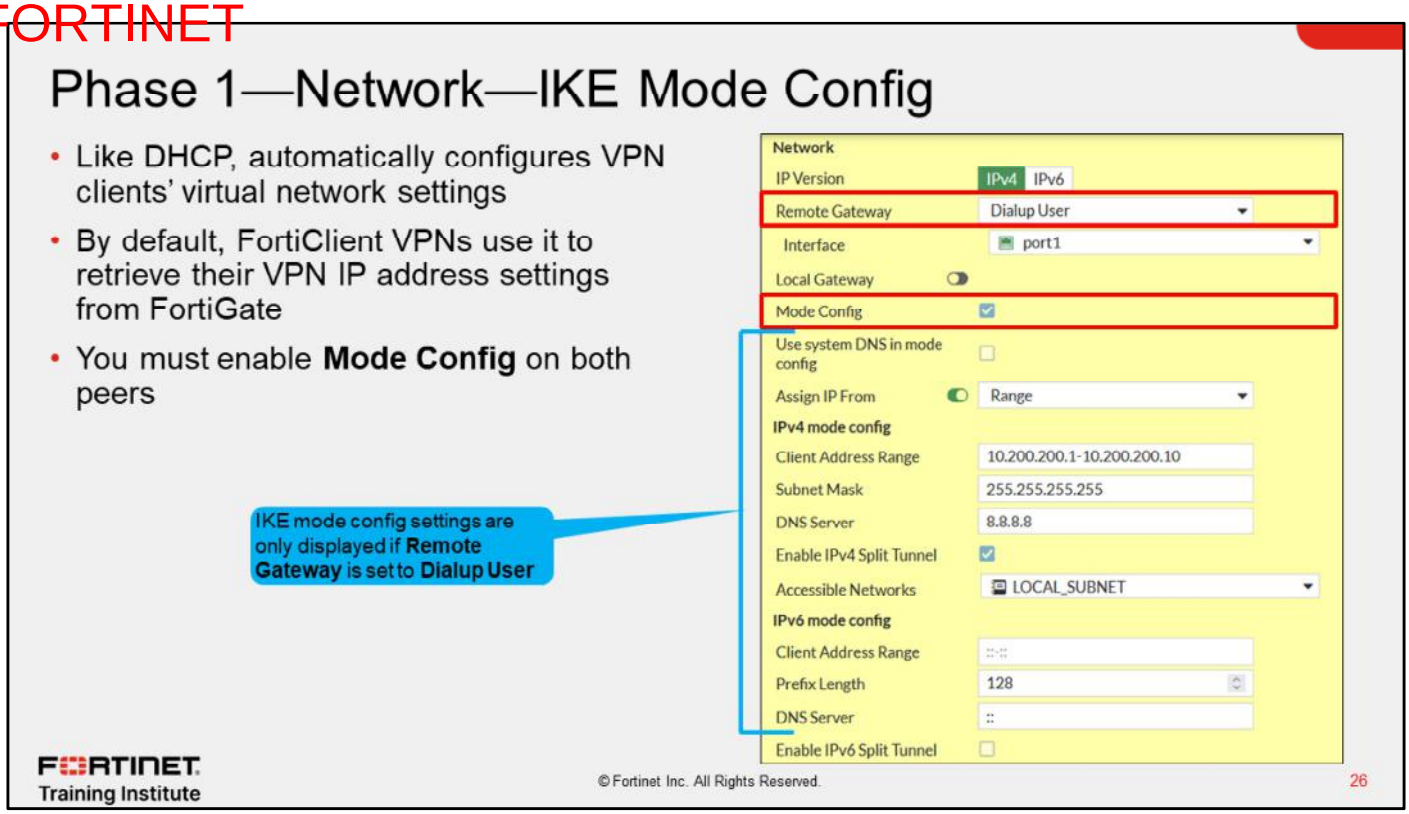

IKE **Mode Config** is similar to DHCP because a server assigns network settings such as IP address, netmask, and DNS servers, to clients. This assignment takes place over IKE messages.

When you enable **Mode Config** on a FortiGate device acting as dial-up server, it pushes network settings to dial-up clients. The dial-up clients are usually FortiClient peers, but they can also be FortiGate peers.

For IKE mode config to work, you must enable the feature on both peers. On FortiClient, **Mode Config** is enabled by default, but on FortiGate, you must manually enable it.

Note that the IKE **Mode Config** settings, are displayed on the GUI only when you set **Remote Gateway** to **Dialup User**. On the FortiGate device acting as dial-up client, you can select **Mode Config** on the GUI, but the additional settings are not displayed.

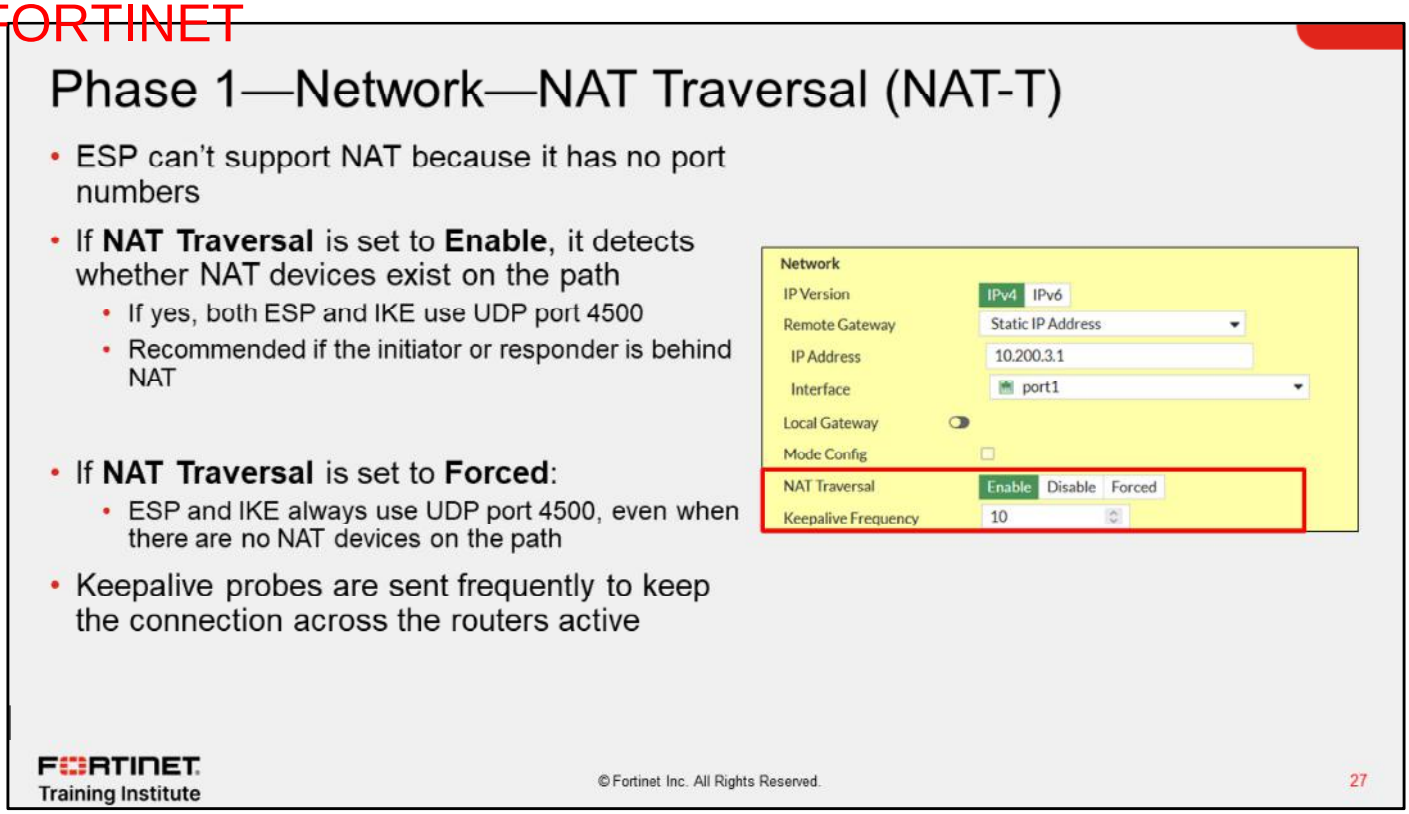

The ESP protocol usually has problems crossing devices that are performing NAT. One of the reasons is that ESP does not use port numbers, like TCP and UDP do, to differentiate one tunnel from another.

To solve this, NAT transversal (NAT-T) was added to the IPsec specifications. When NAT-T is enabled on both ends, peers can detect any NAT device along the path. If NAT is found, then the following occurs on both peers:

- IKE negotiation switches to using UDP port 4500.
- ESP packets are encapsulated in UDP port 4500.

So, if you have two FortiGate devices that are behind, for example, an ISP modem that performs NAT, you will probably need to enable this setting.

When you set the **NAT Traversal** setting to **Forced**, UDP port 4500 is always used, even when there is no NAT device along the path.

When you enable NAT-T, the **Keepalive Frequency** option shows the interval (in seconds) at which FortiGate sends keepalive probes. You need NAT-T when there is one or more routers along the path performing NAT. The purpose of the keepalive probes is to keep the IPsec connection active across those routers along the path.

Phase 1—Network—Dead Peer Detection (DPD)

- Mechanism to detect a dead tunnel
- Useful in redundant VPNs, where multiple paths are available
- Three modes:
  - **On Demand**: DPD probes are sent when there is no inbound traffic
  - **On Idle**: DPD probes are sent when there is no traffic
  - **Disabled**: only reply to DPD probes—don't send probes

FORTINET Training Institute

28

After the peers negotiate the IPsec SAs of a tunnel and, therefore, the tunnel is considered up, the peers usually don't negotiate another IPsec SA until it expires. In most cases, the IPsec SA expires every few hours. This means that if there is a network disruption along the path of the tunnel before the IPsec SA expires, the peers will continue to send traffic through the tunnel even though the communication between the sites is disrupted.

When you enable DPD, DPD probes are sent to detect a failed (or dead) tunnel and bring it down before its IPsec SAs expire. This failure detection mechanism is very useful when you have redundant paths to the same destination, and you want to fail over to a backup connection when the primary connection fails to keep the connectivity between the sites up.

FortiGate supports three DPD modes:

- **On Demand**: FortiGate sends DPD probes if there is only outbound traffic through the tunnel, but no inbound. Because network applications are usually bidirectional, observing only traffic on the outbound direction could be an indication of a network failure.
- **On Idle**: FortiGate sends DPD probes when no traffic is observed in the tunnel. An idle tunnel does not necessarily mean the tunnel is dead. Avoid this mode if you have many tunnels, because the overhead introduced by DPD can be very resource intensive.
- **Disabled**: FortiGate replies only to DPD probes received. FortiGate never sends DPD probes to the remote peer and therefore cannot detect a dead tunnel.

The default DPD mode is **On Demand**. In terms of scalability, **On Demand** is a better option than **On Idle**.

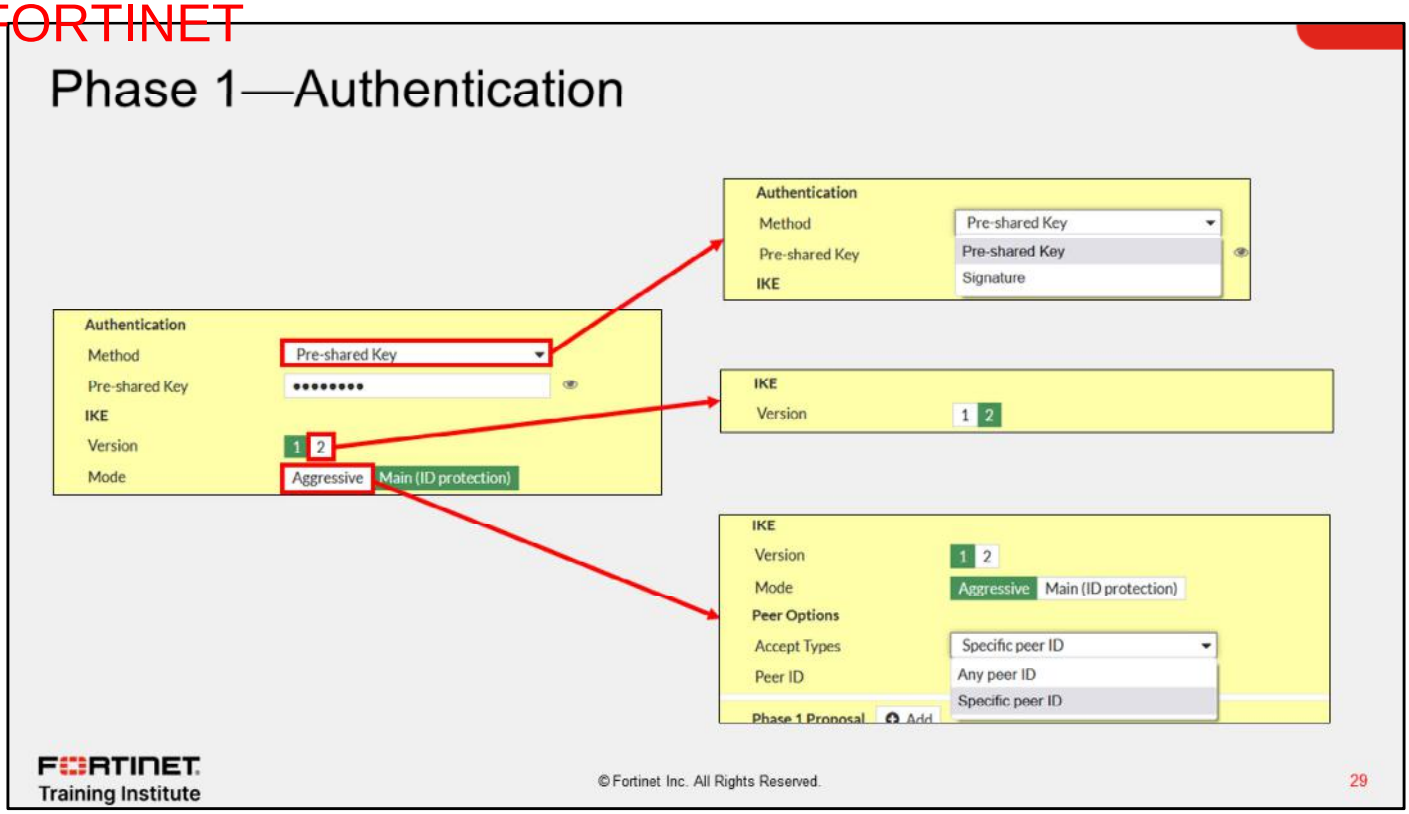

Now, you will learn about the **Authentication** section in phase 1 configuration:

- **Method**: FortiGate supports two authentication methods: **Pre-shared Key** and **Signature**. When you select **Pre-shared Key**, you must configure both peers with the same pre-shared key. When you select **Signature**, phase 1 authentication is based on digital certificate signatures. Under this method, the digital signature on one peer is validated by the presence of the CA certificate installed on the other peer. That is, on the local peer, you need to install both the local peer's certificate and the CA certificate that issued the remote peer certificate.
- **Version**: allows you to select the IKE version to use. When selecting version **2**, aggressive and main modes disappear because they don't apply to IKEv2.
- **Mode**: refers to the IKEv1 mode. Two options are available: **Aggressive** and **Main (ID protection)**. You will learn more about these modes in this lesson.

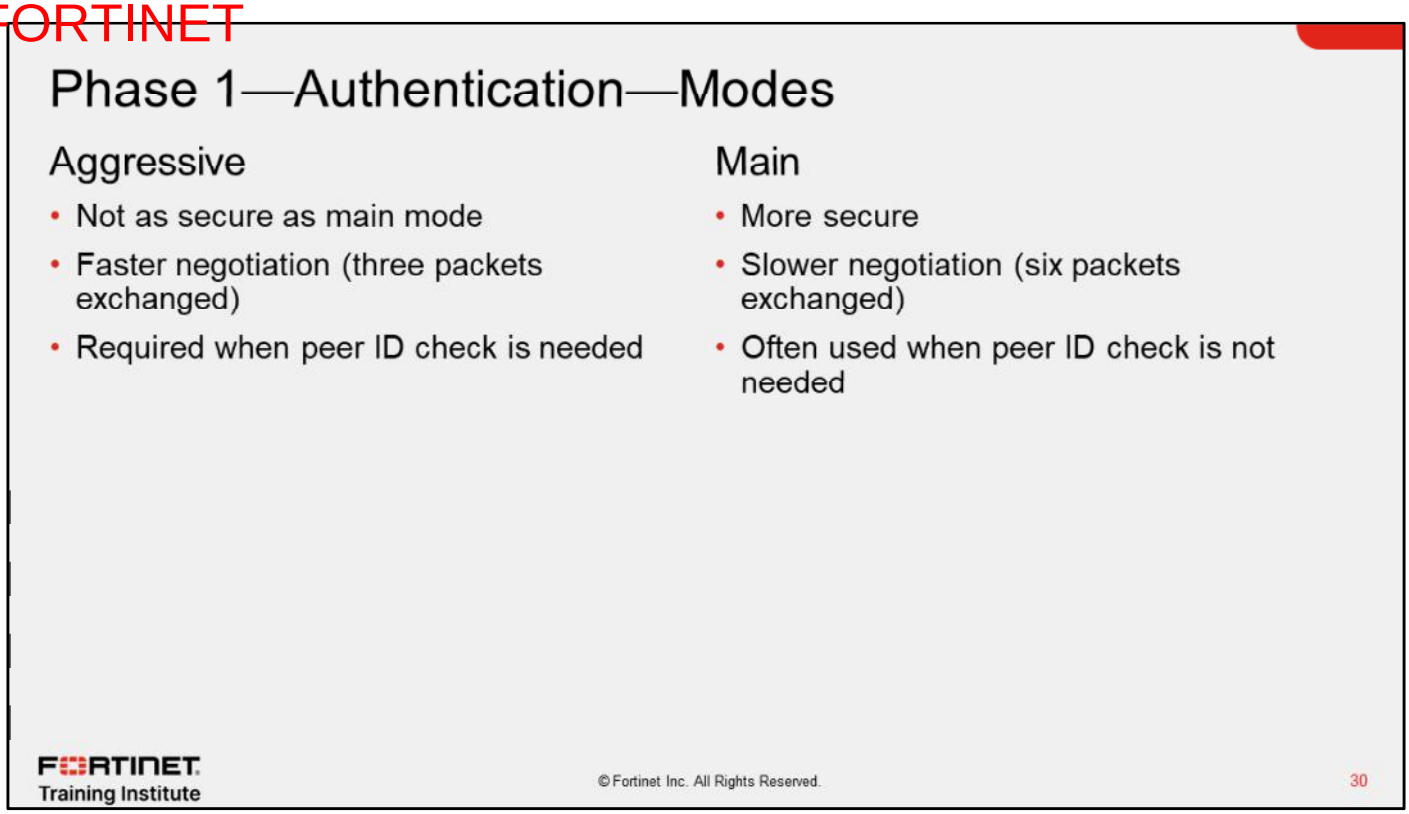

IKE supports two different negotiation modes: main and aggressive. Which one should you use?

To answer that question, we can analyze three categories: security, performance, and deployment.

Security wise, main mode is considered more secure because the pre-shared key hash is exchanged encrypted, whereas in aggressive mode, the hash is exchanged unencrypted. Although the attacker would still have to guess the cleartext pre-shared key for the attack to be successful, the fact that the pre-shared key hash has been encrypted in main mode reduces considerably the chances of a successful attack.

In terms of performance, aggressive mode may be a better option. This is because the negotiation is completed after only three packets are exchanged, whereas in main mode, six packets are exchanged. For this reason, you may want to use aggressive mode when a great number of tunnels terminate on the same FortiGate device, and performance is a concern.

Another use case for aggressive mode, is when there is more than one dial-up tunnel terminating on the same FortiGate IP address, and the remote peer is authenticated using a peer ID because its IP address is dynamic. Because peer ID information is sent in the first packet in an aggressive mode negotiation, then FortiGate can match the remote peer with the correct dial-up tunnel. The latter is not possible in main mode because the peer ID information is sent in the last packet, and after the tunnel has been identified.

When both peers know each other's IP address or FQDN, you may want to use main mode to take advantage of its more secure negotiation. In this case, FortiGate can identify the remote peer by its IP address and, as a result, associate it with the correct IPsec tunnel.

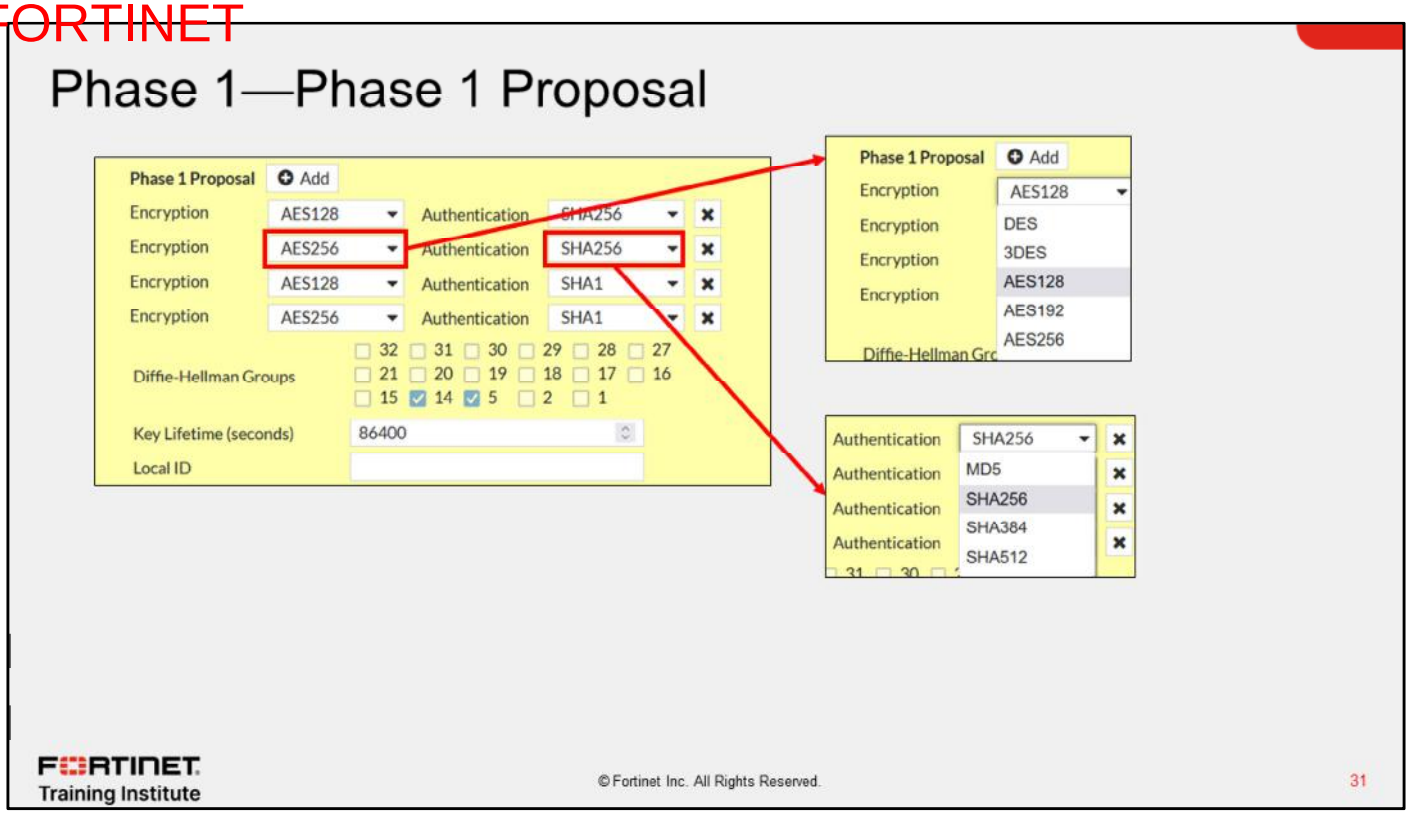

Now, you will learn about the **Phase 1 Proposal** section of phase 1 configuration. This section allows you to enable the different proposals that FortiGate supports when negotiating the IKE SA (or phase 1 SA). You can combine different parameters to suit your security needs. You must at least configure one combination of encryption and authentication algorithms, or several.

- **Encryption**: select the algorithm to use for encrypting and decrypting the data.
- **Authentication**: select the authentication algorithm to use for verifying the integrity and authenticity of the data.
- **Diffie-Hellman Groups**: The Diffie-Hellman (DH) algorithm is used during IKE SA negotiation. The use of DH in phase 1 is mandatory and can't be disabled. You must select at least one DH group. The higher the DH group number, the more secure the phase 1 negotiation is. However, a higher DH group number also results in a longer compute time.
- **Key Lifetime**: defines the lifetime of the IKE SA. At the end of the lifetime, a new IKE SA is negotiated.
- **Local ID**: if the peer accepts a specific peer ID, type that same peer ID in this field.

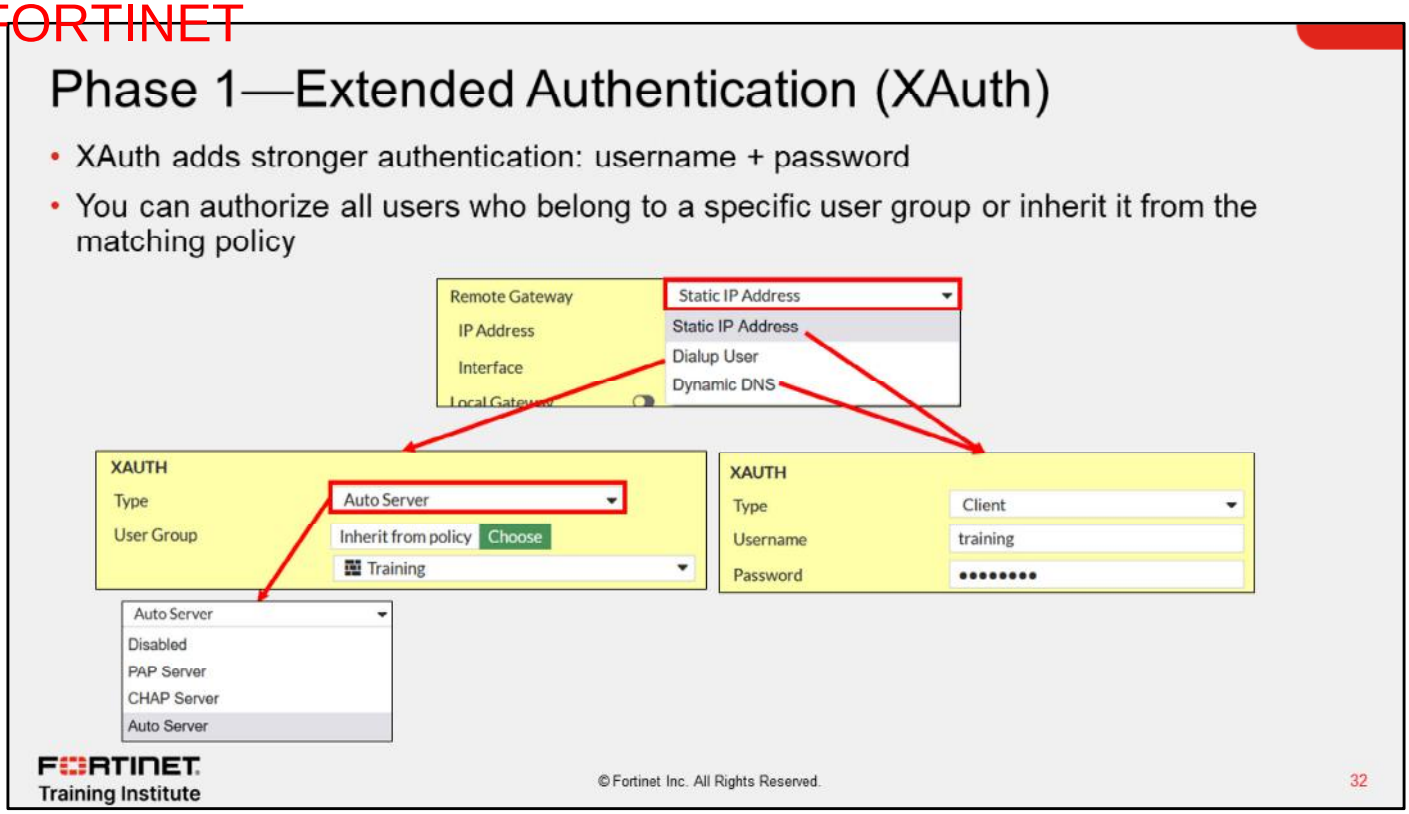

Phase 1 supports two types of authentication: pre-shared keys and digital signatures. The XAuth extension, sometimes called phase 1.5, forces remote users to authenticate additionally with their credentials (username and password). So, additional authentication packets are exchanged if you enable it. What is the benefit? Stronger authentication.

When you set **Remote Gateway** to **Dialup User**, FortiGate acts as the authentication server. The **XAUTH** section shows the authentication server type options: **PAP Server**, **CHAP Server**, and **Auto Server**. In the example shown on this slide, **Auto Server** is selected, which means that FortiGate automatically detects the authentication protocol used by the client.

After you select the authentication server type, you configure how user group matching is performed. There are two options: **Inherit from policy** and **Choose**. The latter is used in the example on this slide, and allows you to select one of the user groups available on FortiGate. Note that, when you select **Choose**, you must configure a separate dial-up VPN for every group of users that require a different network access policy.

The other way to authenticate VPN users with XAuth is by selecting **Inherit from policy**. When you select this option, FortiGate authenticates users based on their matching IPsec policy and, as a result, the configuration for controlling network access is simpler. That is, you control network access by configuring multiple policies for different user groups, instead of configuring multiple tunnels for different user groups. The **Inherit from policy** option follows a similar authentication approach used for SSL VPN remote users. You will learn more about SSL VPN in another lesson.

When **Remote Gateway** is set to **Static IP Address** or **Dynamic DNS**, FortiGate acts as the client, and the **XAUTH** section shows the **Client** option as **Type**. You can then set the credentials that FortiGate uses to authenticate against the remote peer through XAuth.

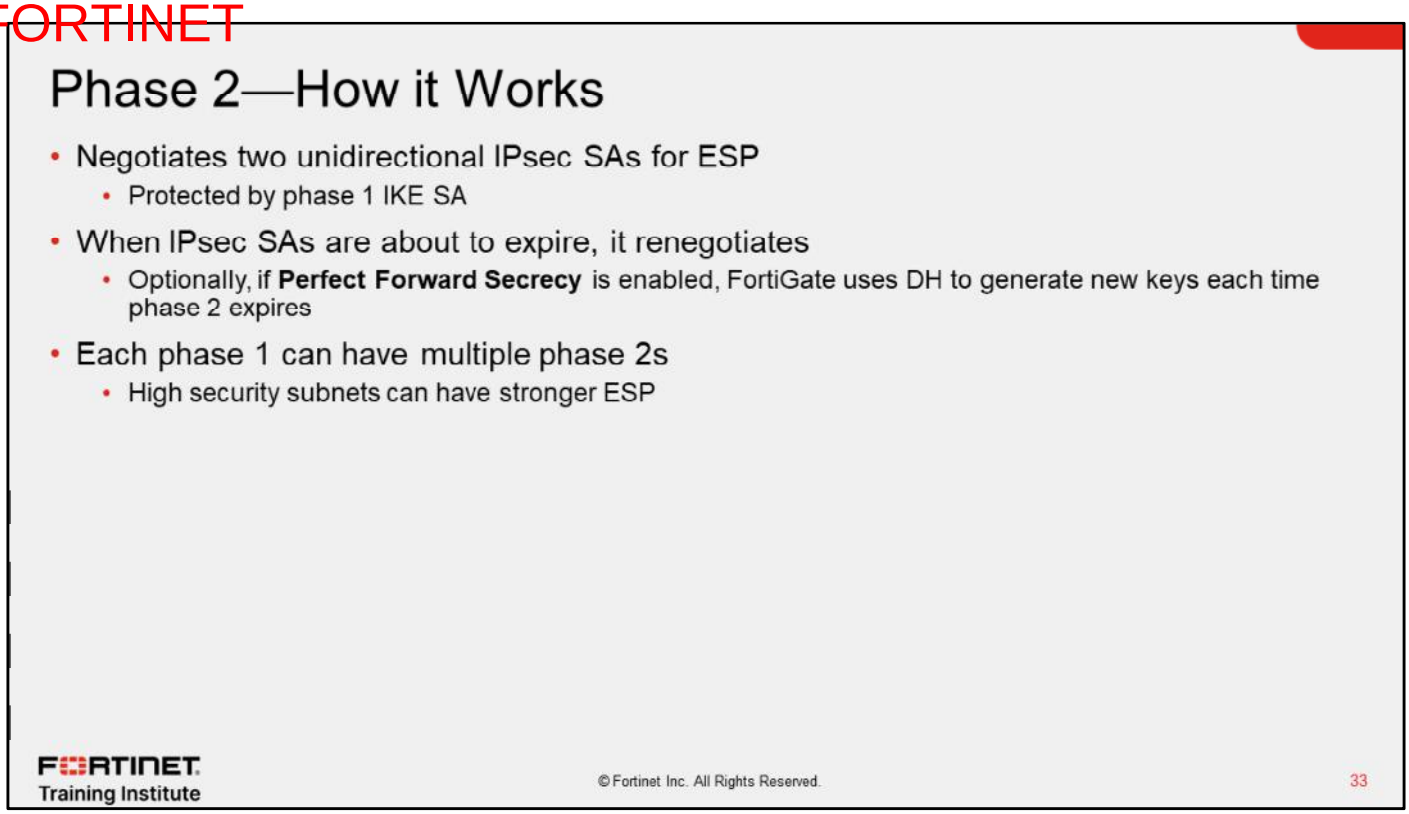

After phase 1 has established a secure channel to exchange data, phase 2 begins.

Phase 2 negotiates security parameters for two IPsec SAs over the secure channel established during phase 1. ESP uses IPsec SAs to encrypt and decrypt the traffic exchanged between sites.

Phase 2 does not end when ESP begins. Phase 2 periodically renegotiates IPsec SAs to maintain security. If you enable **Perfect Forward Secrecy**, each time phase 2 expires, FortiGate uses DH to recalculate new secret keys. In this way, new keys are not derived from older keys, making it much harder for an attacker to crack the tunnel.

Each phase 1 can have multiple phase 2s. When would this happen? For example, you may want to use different encryption keys for each subnet whose traffic is crossing the tunnel. How does FortiGate select which phase 2 to use? By checking which phase 2 selector (or quick mode selector) matches the traffic.

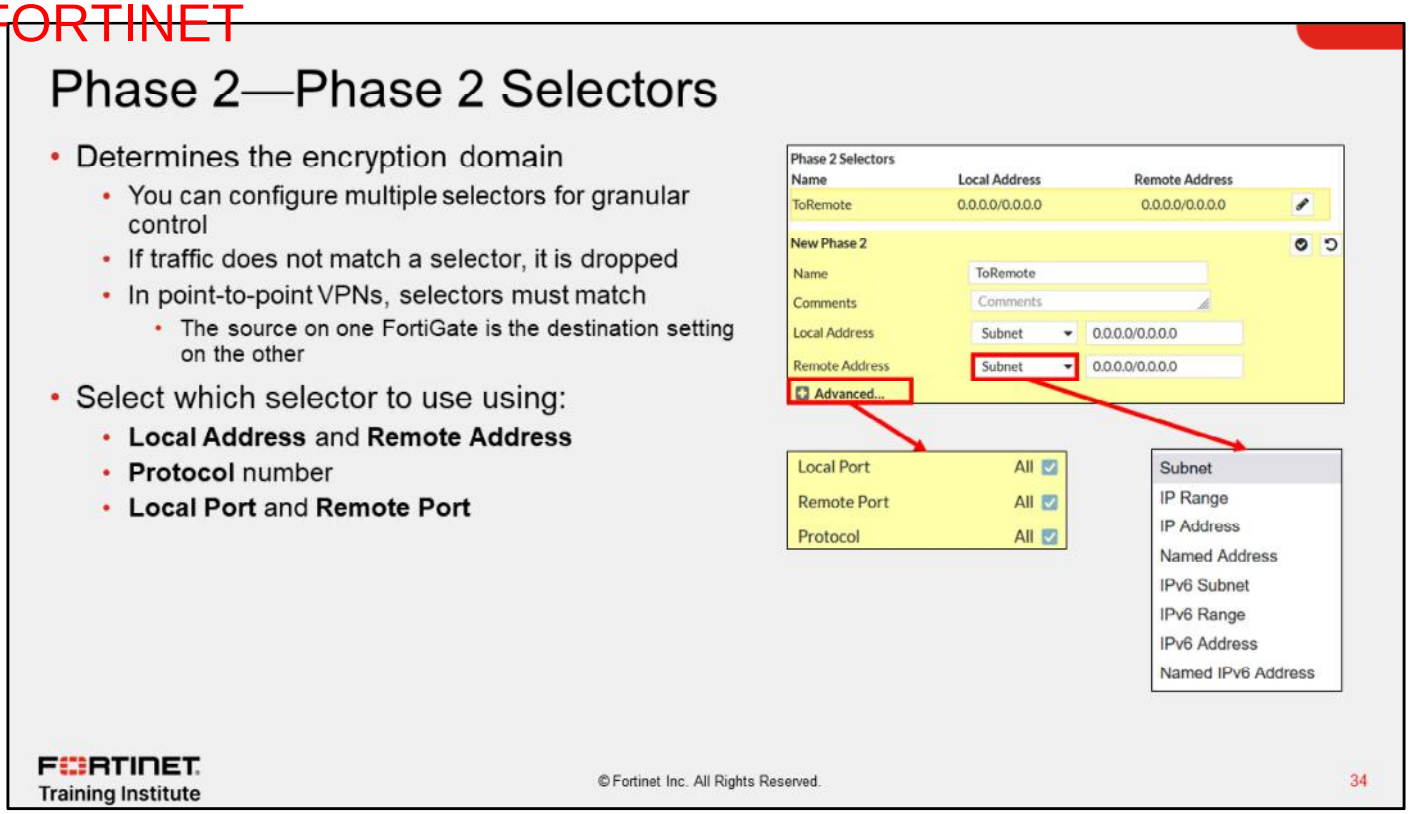

In phase 2, you must define the encryption domain (or interesting traffic) of your IPsec tunnel. The encryption domain refers to the traffic that you want to protect with IPsec, and it is determined by your phase 2 selector configuration.

You can configure multiple selectors to have more granular control over traffic. When you configure a phase 2 selector, you specify the encryption domain by indicating the following network parameters:

- **Local Address** and **Remote Address**: as seen in the example shown on this slide, you can define IPv4 or IPv6 addresses using different address scopes. When selecting **Named Address** or **Named IPv6 Address**, FortiGate allows you to select an IPv4 or IPv6 firewall address object, respectively, configured in the system.
- **Protocol**: is in the **Advanced** section, and is set to **All** by default.
- **Local Port** and **Remote Port**: are also shown in the **Advanced** section, and are set to **All** by default. This applies only to port-based traffic such as TCP or UDP. You will learn more about the **Advanced** section later in this lesson.

Note that after the traffic is accepted by a firewall policy, traffic is dropped before entering the IPsec tunnel if the traffic does not match any of the phase 2 selectors configured. For this reason, usually, it's more intuitive to filter traffic with firewall policies. So, if you don't want to use phase 2 selector filtering, you can just create one phase 2 selector with both the local and remote addresses set to any subnet, like in the example shown on this slide, and then use firewall policies to control which traffic is accepted on the IPsec tunnel.

In addition, the phase 2 selector network parameters on both peers must match if the tunnel is point-to-point, that is, when the remote gateway is *not* set to dial-up user.

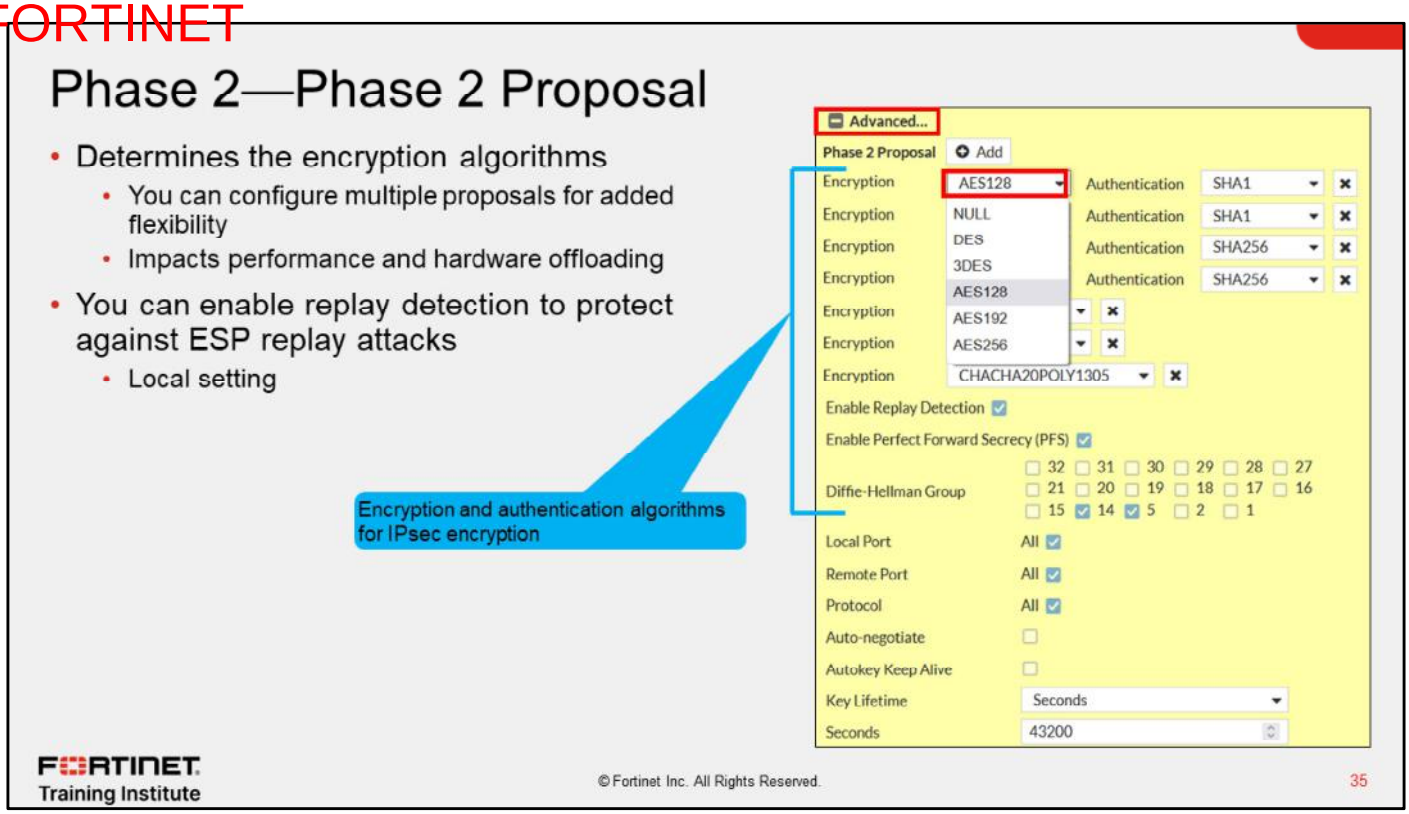

For every phase 2 selector, you need to configure one or more phase 2 proposals. A phase 2 proposal defines the algorithms supported by the peer for encrypting and decrypting the data over the tunnel. You can configure multiple proposals to offer more options to the remote peer when negotiating the IPsec SAs.

Like in phase 1, you need to select a combination of encryption and authentication algorithms. Some algorithms are considered more secure than others, so make sure to select the algorithms that conform with your security policy. However, note that the selection of the algorithms has a direct impact on FortiGate IPsec performance. For example, **3DES** is known to be a much more resource-intensive encryption algorithm than **DES** and **AES**, which means that your IPsec throughput could be negatively impacted if you select **3DES** as the encryption algorithm. Also, note that if you select **NULL** as the encryption algorithm, traffic is not encrypted.

In addition, some encryption algorithms, such as **CHACHA20POLY1305**, are not supported for hardware offload. That is, if you have a FortiGate device that contains network processor (NP) units, you can achieve higher IPsec performance if you select an algorithm that is supported for IPsec offload by your NP unit model, such as AES or DES. For a list of supported encryption algorithms for IPsec hardware offloading, refer to `https://docs.fortinet.com`.

When configuring the phase 2 proposal, you can select **Enable Replay Detection** to detect antireplay attacks on ESP packets. Note that this is a local setting and, therefore, it is not included as part of the proposals presented by the peer during phase 2 negotiation.

Also, if you enable **Perfect Forward Secrecy**, FortiGate uses DH to enhance security during the negotiation of IPsec SAs.

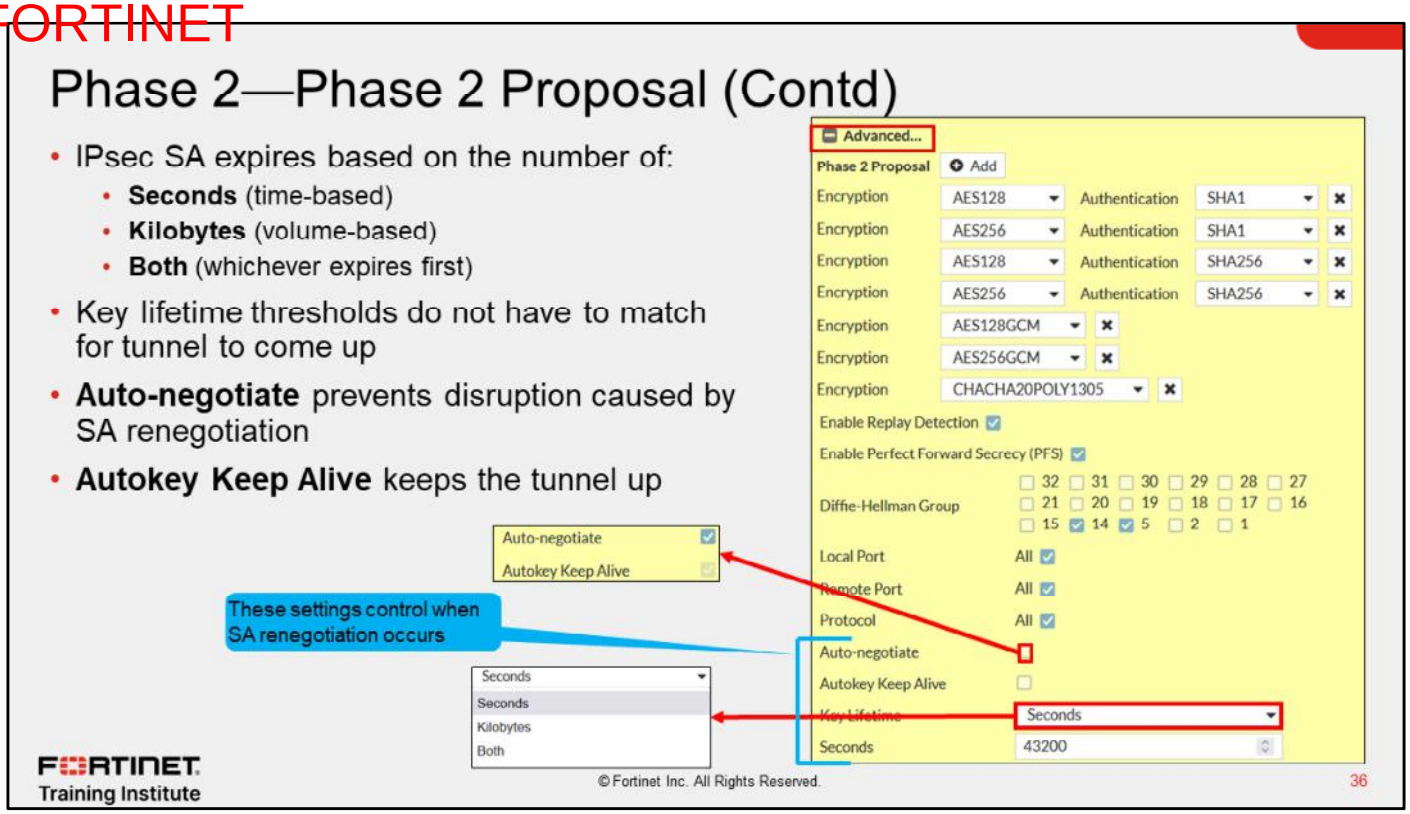

IPsec SAs are periodically renegotiated to improve security, but when does that happen? It depends on the key lifetime settings configured on the phase 2 proposal.

The expiration of an IPsec SA is determined by the lifetime type and threshold configured. By default, **Key Lifetime** is set to **Seconds** (time-based). This means that when the SA duration reaches the number of seconds set as **Seconds**, the SA is considered expired. You can also set the key lifetime to **Kilobytes** (volume-based), upon which the SA expires after the amount of traffic encrypted and decrypted using that SA reaches the threshold set. Alternatively, you can select **Both** as the key lifetime type, upon which FortiGate tracks both the duration of the SA and the amount of traffic. Then, when any of the two thresholds is reached, the SA is considered expired. Note that the key lifetime thresholds do not have to match for the tunnel to come up. When thresholds are different, the peers agree on using the lowest threshold value offered between the two.

When IPsec SAs expire, FortiGate needs to negotiate new SAs to continue sending and receiving traffic over the IPsec tunnel. Technically, FortiGate deletes the expired SAs from the respective phase 2 selectors, and installs new ones. If IPsec SA renegotiation takes too much time, then FortiGate might drop interesting traffic because of the absence of active SAs. To prevent this, you can enable **Auto-negotiate**. When you do this, FortiGate not only negotiates new SAs before the current SAs expire, but it also starts using the new SAs right away. The latter prevents traffic disruption by IPsec SA renegotiation.

Another benefit of enabling **Auto-negotiate** is that the tunnel comes up and stays up automatically, even when there is no interesting traffic. When you enable **Autokey Keep Alive** and keep **Auto-negotiate** disabled, the tunnel does not come up automatically unless there is interesting traffic. However, after the tunnel is up, it stays that way because FortiGate periodically sends keep alive packets over the tunnel. Note that when you enable **Auto-negotiate**, **Autokey Keep Alive** is implicitly enabled.

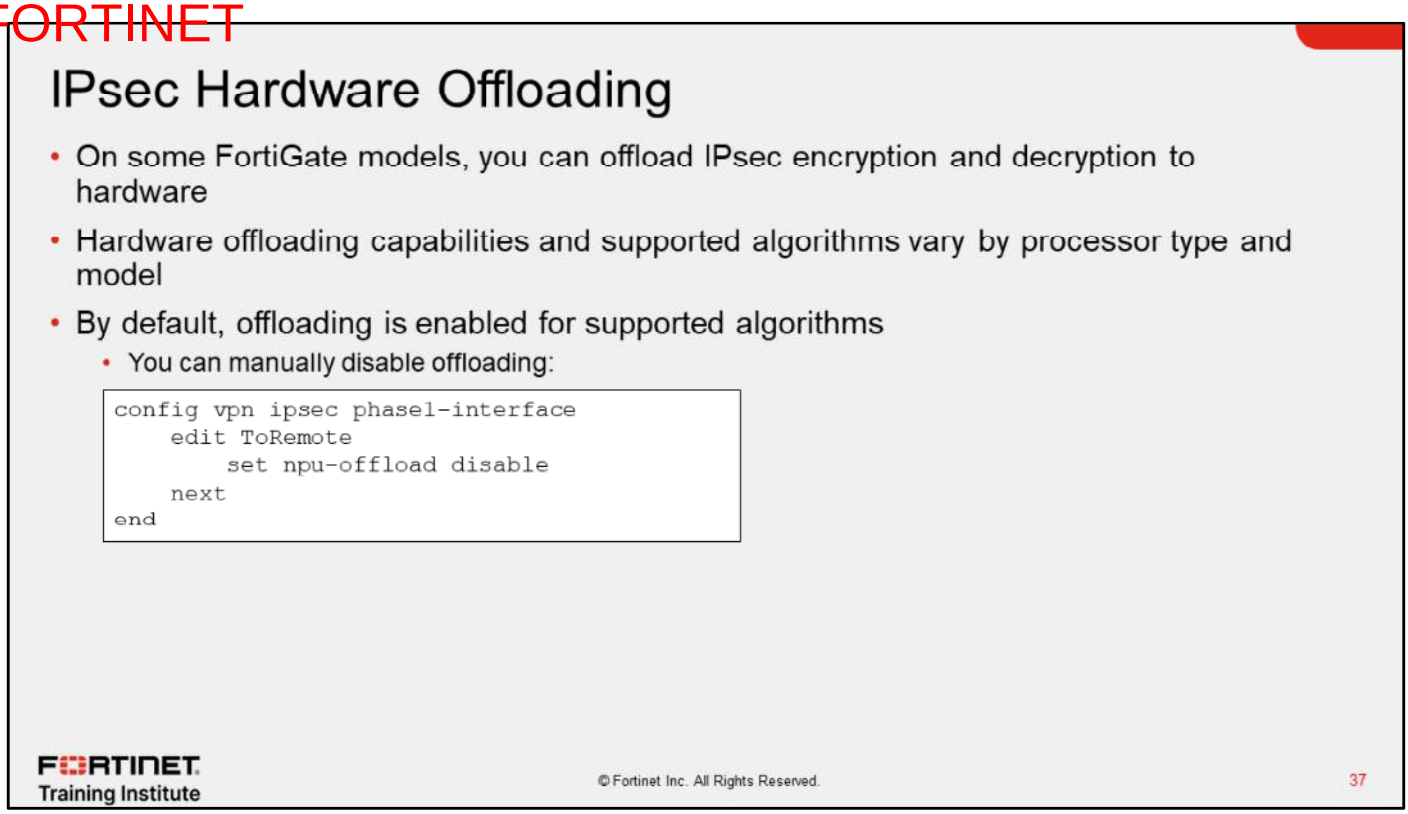

On some FortiGate models, you can offload the encryption and decryption of IPsec traffic to hardware. The algorithms that are supported depend on the NP unit model present on FortiGate. For a list of supported encryption algorithms for IPsec hardware offloading, refer to `https://docs.fortinet.com`.

By default, hardware offloading is enabled for the supported algorithms. This slide shows the commands you can use to disable hardware offloading per tunnel, if necessary.

Knowledge Check

1. Which type of VPN peer can initiate a VPN tunnel?
   - A. Dial-up server
   - ✓ B. Dial-up client

2. On which phase do you configure the algorithms used for traffic encryption?
   - A. Phase 1
   - ✓ B. Phase 2

3. Which IKEv1 negotiation mode is faster?
   - ✓ A. Aggressive
   - B. Main

FORTINET Training Institute

38

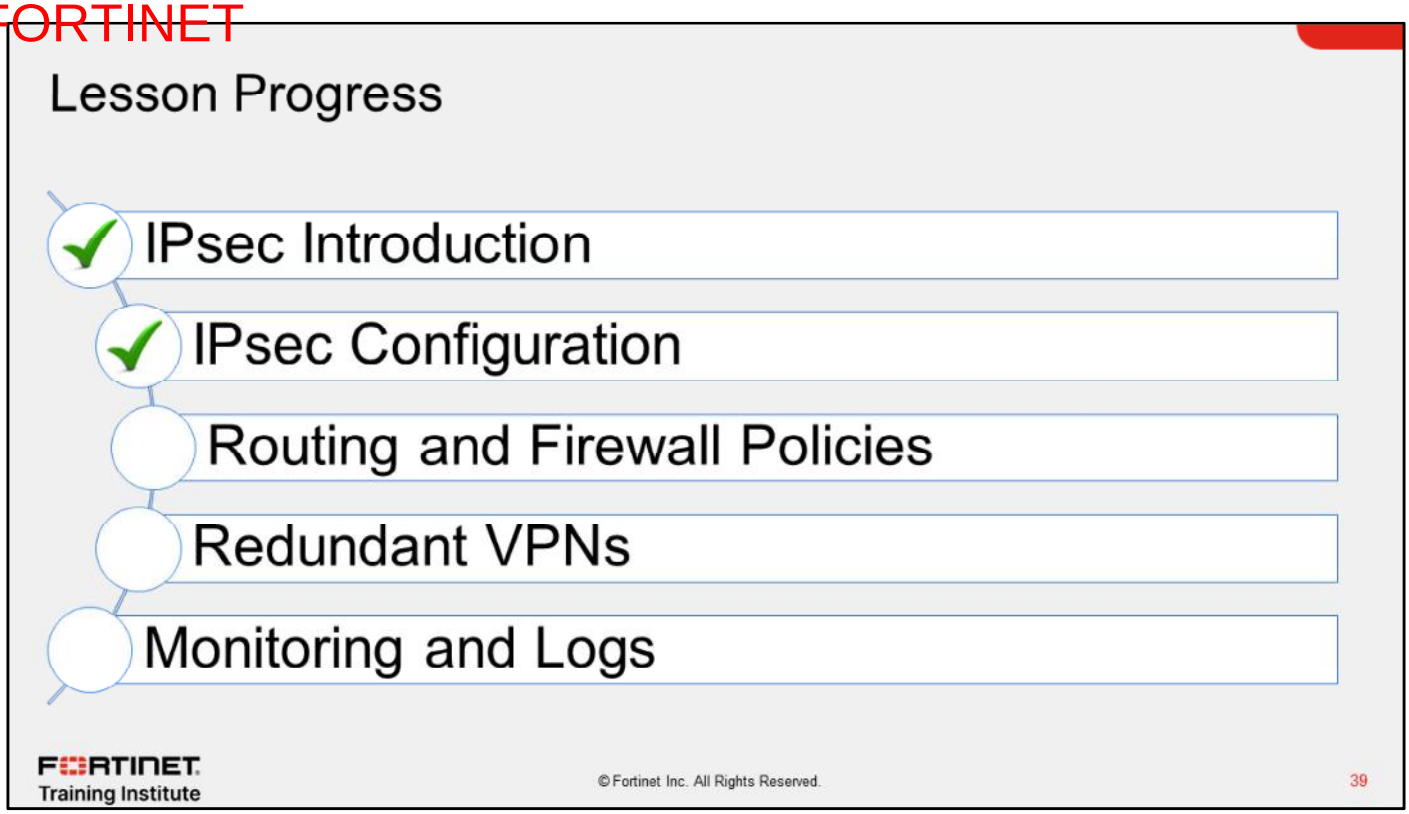

Good job! You now understand IPsec configuration.

Now, you will learn about routing and firewall policies for IPsec traffic.

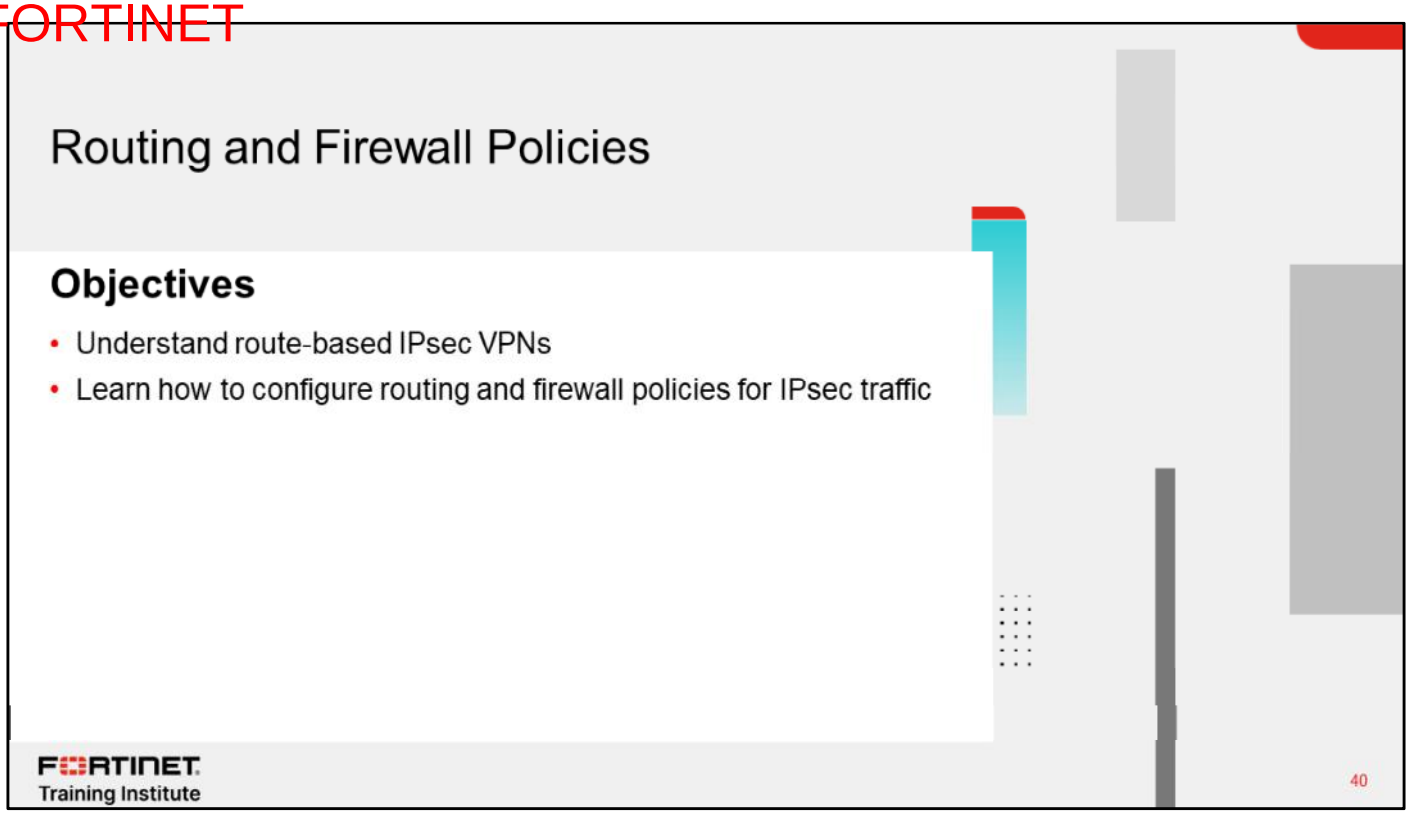

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in routing and firewall policies for IPsec VPNs, you will be able to set up appropriate routing and firewall policies for your IPsec VPN deployment.

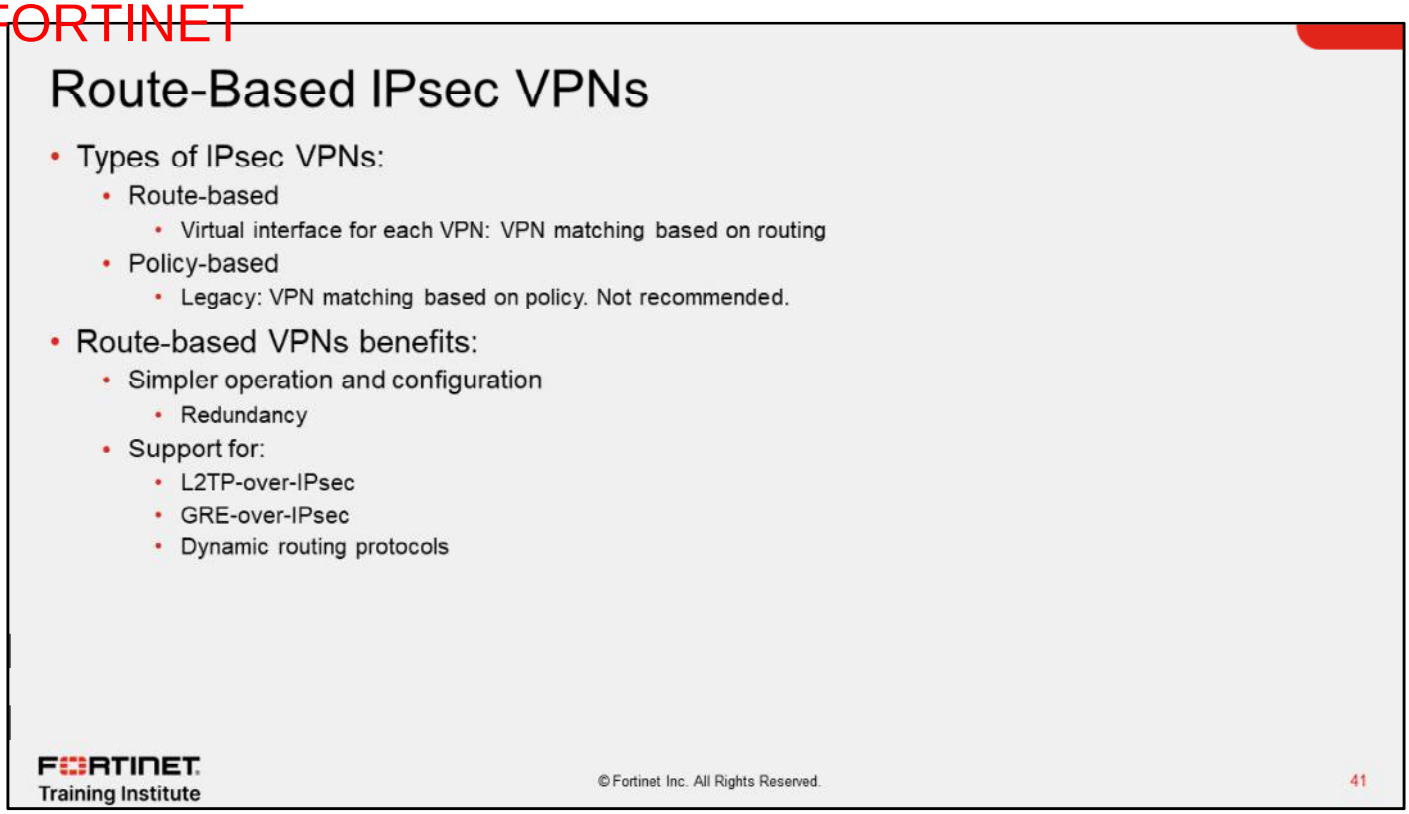

FortiGate supports two types of IPsec VPNs: route-based and policy-based. Policy-based is a legacy IPsec VPN that is supported only for backward compatibility reasons, and its use *is not recommended* for new deployments. Unless otherwise stated, all IPsec VPN references in this lesson are for route-based IPsec VPNs.

In a route-based IPsec VPN, FortiGate automatically adds a virtual interface with the VPN name. This means that not only can you configure routing and firewall polices for IPsec traffic in the same way you do for non-IPsec traffic, but you also can leverage the presence of multiple connections to the same destination to achieve redundancy.

Another benefit of route-based IPsec VPNs is that you can deploy variations of IPsec VPNs such as L2TP-over-IPsec and GRE-over-IPsec. In addition, you can also enable dynamic routing protocols for scalability purposes and best path selection.

Routes for IPsec VPNs

Dial-up user

```
config vpn ipsec phase1-interface
    edit "Dialup"
        set add-route enable | disable
    next
end
```

- `add-route` is enabled (default)
  - No need to configure static routes
  - Static routes are added after phase 2 is up
    - The destination is the local network presented by the dial-up client during phase 2 negotiation
    - The default route distance is 15
  - Static routes are deleted after phase 2 is down
- `add-route` is disabled
  - Useful when dynamic routing protocol is used
  - Dynamic routing protocol takes care of routing updates

Static IP address / dynamic DNS

- Static routes are needed

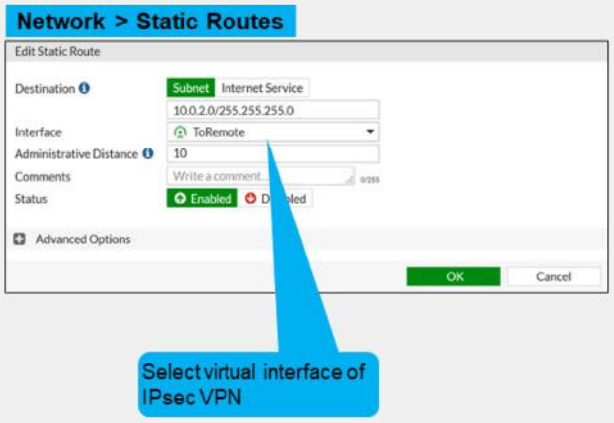

FORTINET Training Institute

42

Although you can use dynamic routing protocols for IPsec VPNs, this lesson covers only the use of static routes.

The routing configuration needed for your IPsec VPN depends on the type of remote gateway configured. When you set the remote gateway to **Dialup User** and enable `add-route`, FortiGate automatically adds a static route for the local network presented by the remote peer during phase 2 negotiation. In addition, the route is added to the routing table only after phase 2 is up. If phase 2 goes down, the static route is removed from the routing table.

When you set the remote gateway to **Dialup User** and disable `add-route`, FortiGate does not add static routes automatically. In this case, a dynamic routing protocol is used between the remote peers to exchange routing information.

When the remote gateway is set to **Static IP Address** or **Dynamic DNS**, you must configure static routes. When you configure a static route, you select the virtual interface of the IPsec tunnel as the outgoing interface.

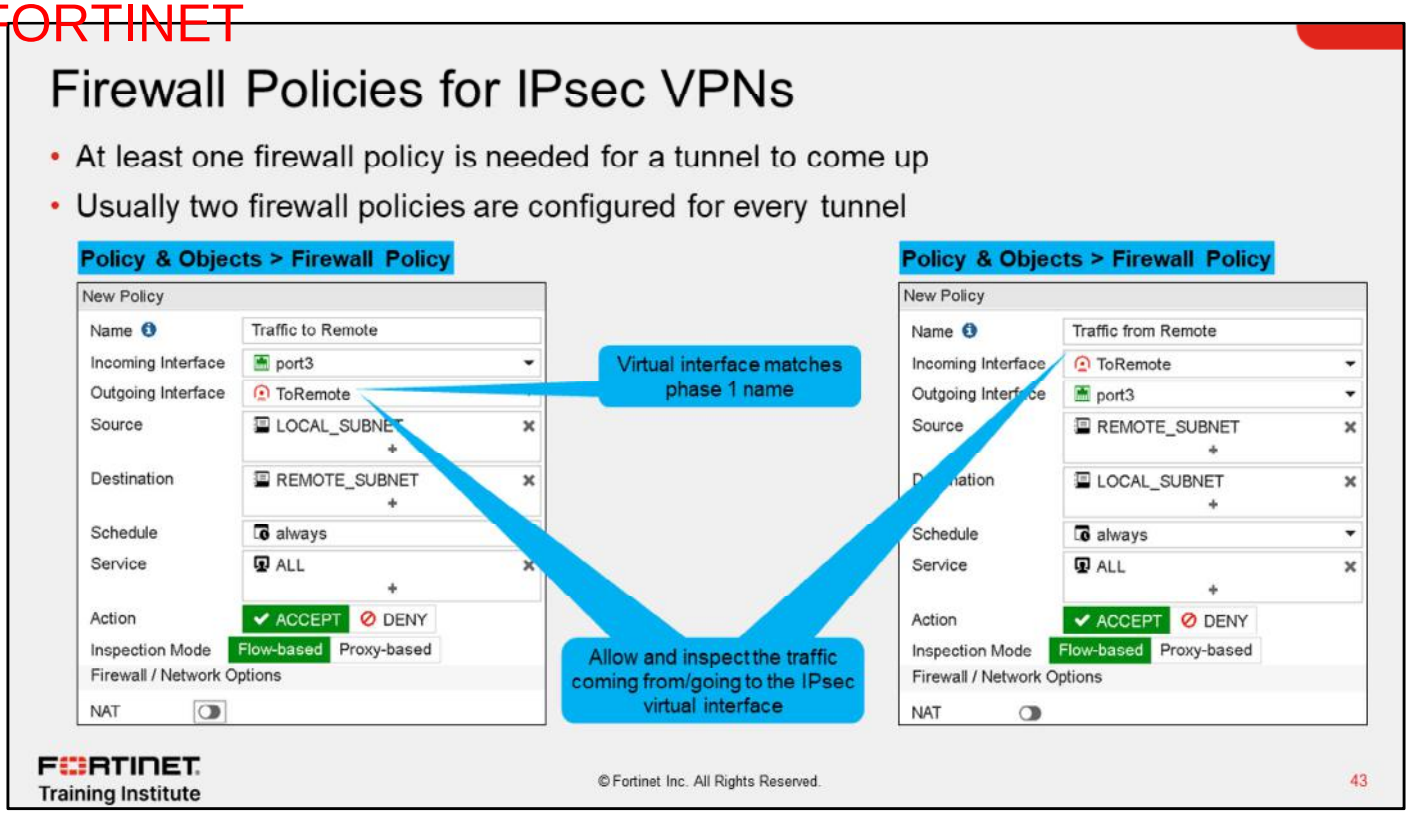

You must configure at least one firewall policy that accepts traffic on your IPsec tunnel. Otherwise, the tunnel will not come up.

When you configure firewall policies for non-IPsec traffic, the policy determines the direction of the traffic that initiates sessions. The same applies to IPsec traffic. For this reason, you usually want to configure at least two firewall policies for your IPsec VPN: one incoming policy and one outgoing policy. The incoming policy allows traffic initiated from the remote site, while the outgoing policy allows traffic to be initiated from the local network.

Note that the policies are configured with the virtual tunnel interface (or phase 1 name) as the incoming or outgoing interface.

Knowledge Check

1. Which IPsec VPN type is legacy and not recommended for new deployments?
   - A. Route-based IPsec VPN
   - ✓ B. Policy-based IPsec VPN

2. What is a configuration requirement for an IPsec tunnel to come up?
   - ✓ A. A firewall policy accepting traffic on the IPsec tunnel
   - B. A route for IPsec traffic

FORTINET Training Institute

44

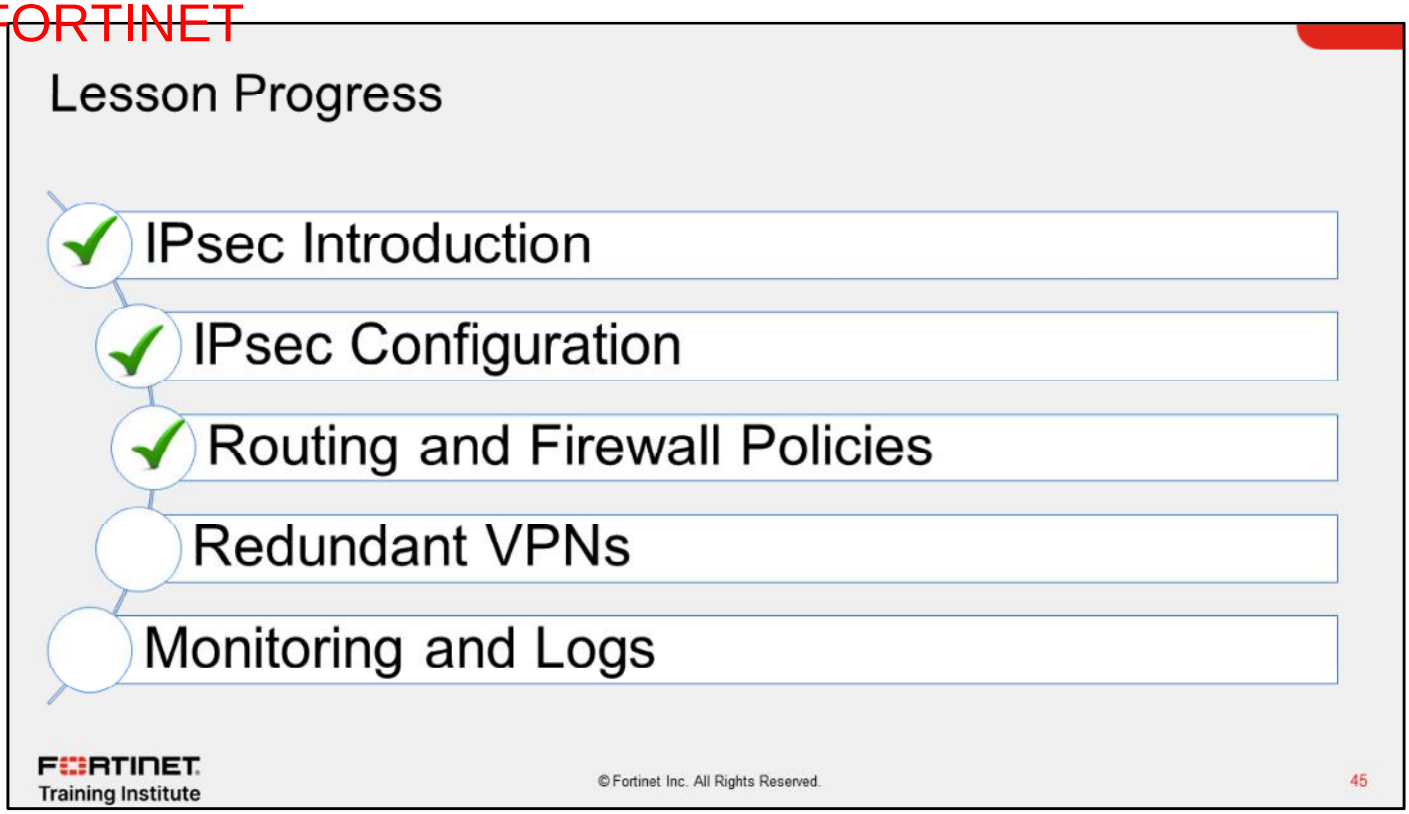

Good job! You now understand routing and firewall policies for IPsec traffic.

Now, you will learn about redundant VPNs.

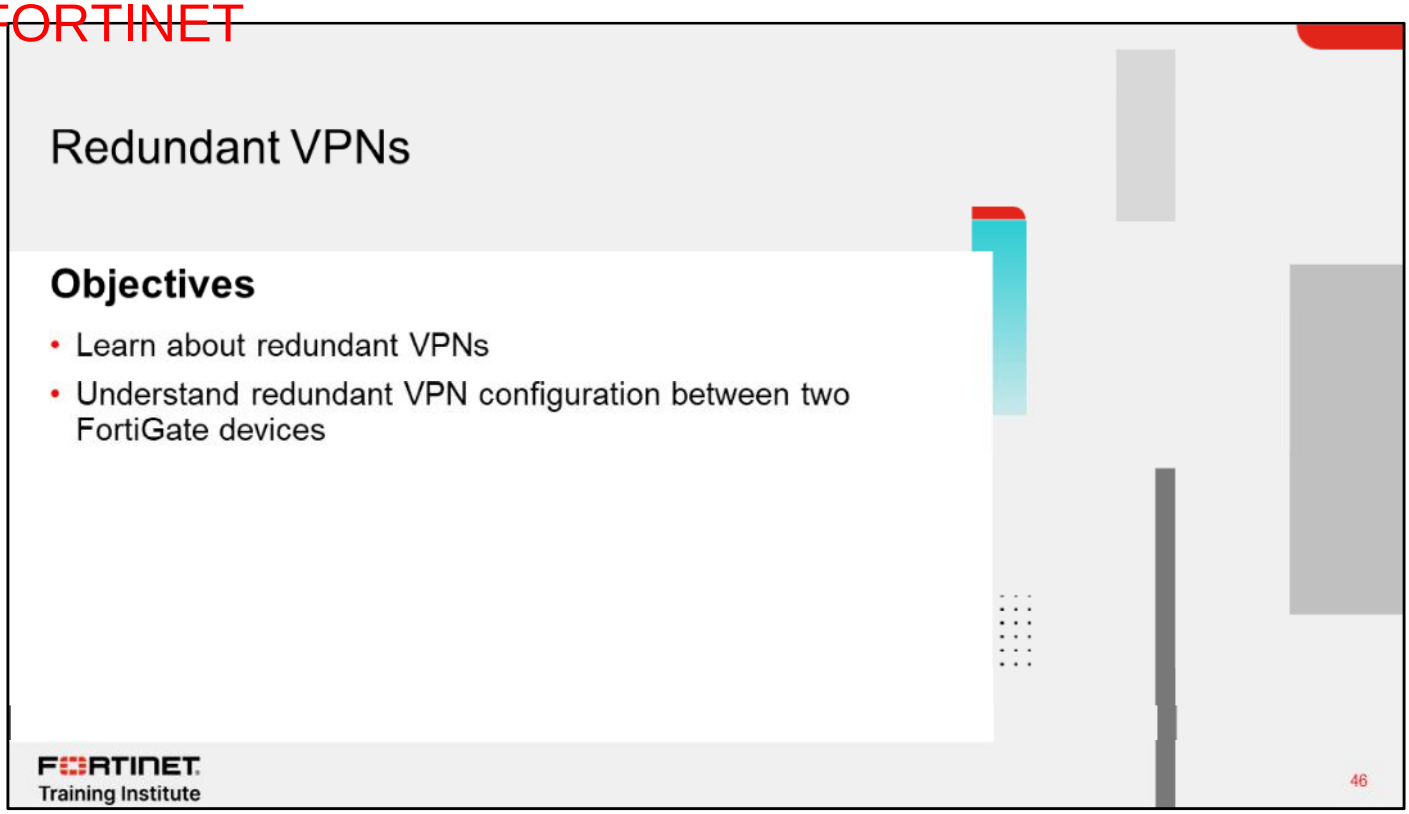

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in redundant VPNs, you will able to add redundancy to your IPsec VPN deployment.

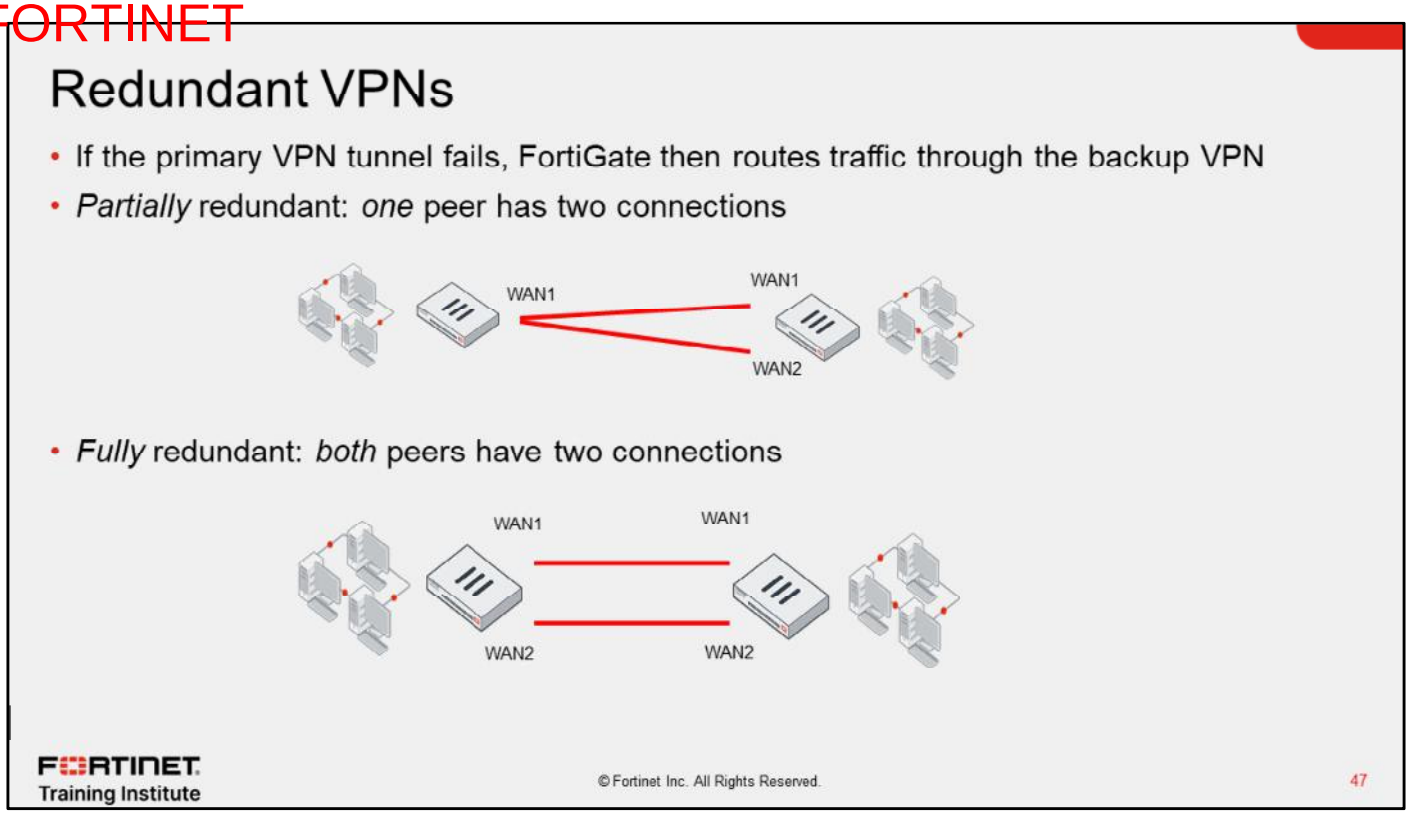

How can you make your IPsec VPN deployment more resilient? Provide a second ISP connection to your site and configure two IPsec VPNs. If the primary IPsec VPN fails, another tunnel can be used instead.

There are two types of redundant VPNs:

- Partially redundant: on one peer (usually the hub, where a backup ISP is available if the main ISP is down), each VPN terminates on *different* physical ports. That way, FortiGate can use an alternative VPN. On the other peer, each VPN terminates on the *same* physical port—so the spoke is not fault tolerant.
- Fully-redundant: both peers terminate their VPNs on different physical ports, so they are both fault tolerant.

Redundant VPN Configuration

- Add one phase 1 configuration for each tunnel. You should enable DPD on both ends.
- Add at least one phase 2 definition for each phase 1
- Add one static route for each path
  - Use distance or priority to select primary routes over backup routes
  - Alternatively, use dynamic routing
- Configure firewall policies for each IPsec interface

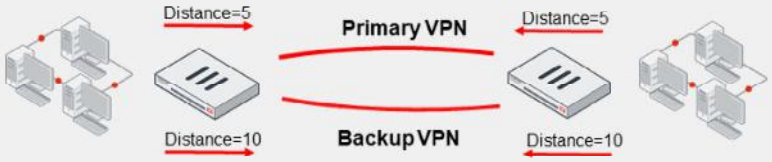

FORTINET Training Institute

48

So, how do you configure a partially or fully redundant VPN?

First, create one phase 1 for each path—one phase 1 for the primary VPN and one for the backup VPN. You should also enable DPD on both ends.

Second, create at least one phase 2 definition for each phase 1.

Third, you must add at least one static route for each VPN. Routes for the primary VPN must have a lower distance (or lower priority) than the backup. This causes FortiGate to use the primary VPN while it's available. If the primary VPN fails, then FortiGate automatically uses the backup route. Alternatively, you could use a dynamic routing protocol, such as OSPF or BGP.

Finally, configure firewall policies to allow traffic through both the primary and backup VPNs.

Knowledge Check

1. Which feature should be enabled in a redundant IPsec VPN deployment?
   - ✓ A. DPD
   - B. XAuth

2. Which setting determines whether a tunnel is used as primary or backup?
   - ✓ A. Routing
   - B. Firewall policies

FORTINET Training Institute

49

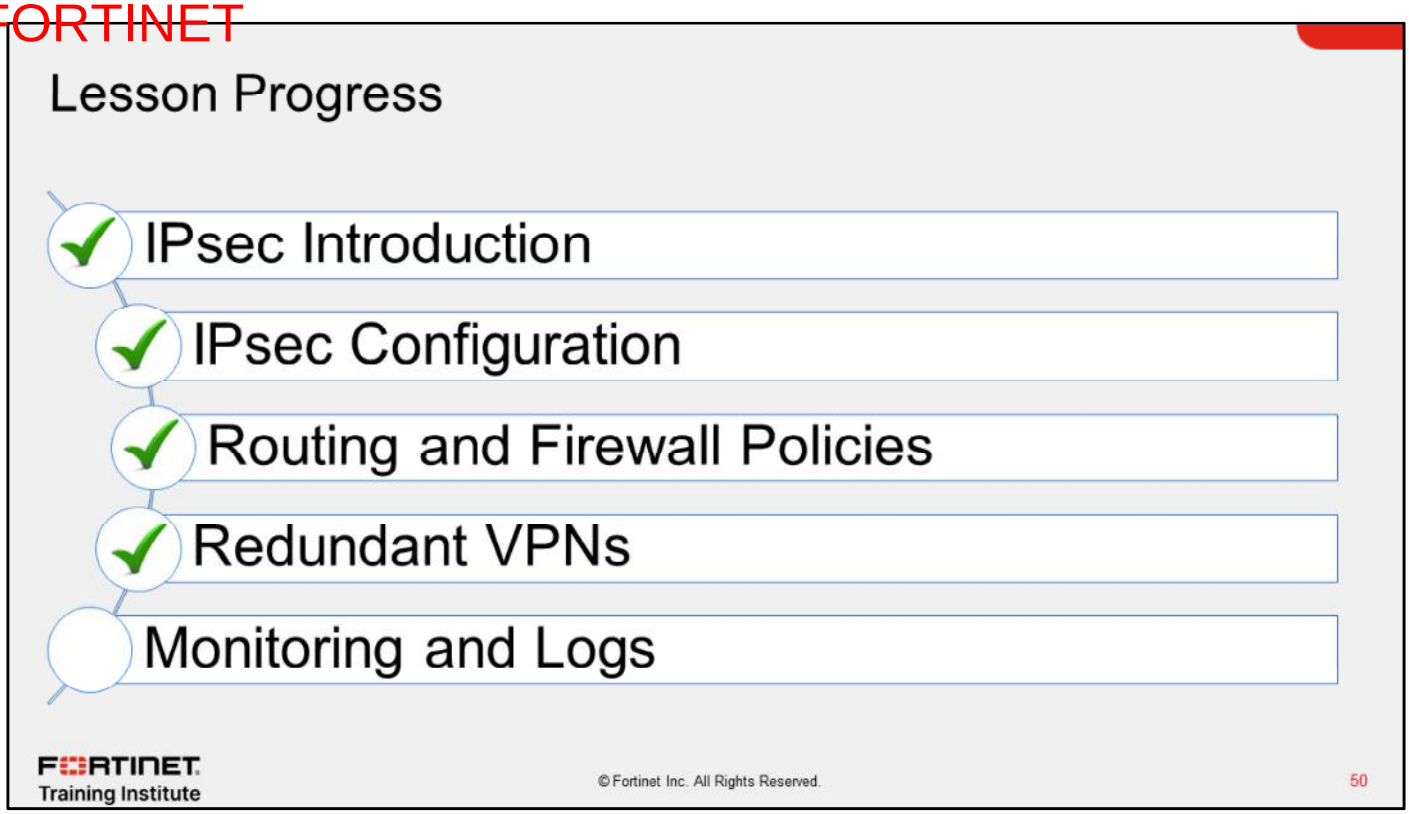

Good job! You now understand redundant VPNs.

Now, you wil l learn about monitoring IPsec VPNs and reviewing their logs.

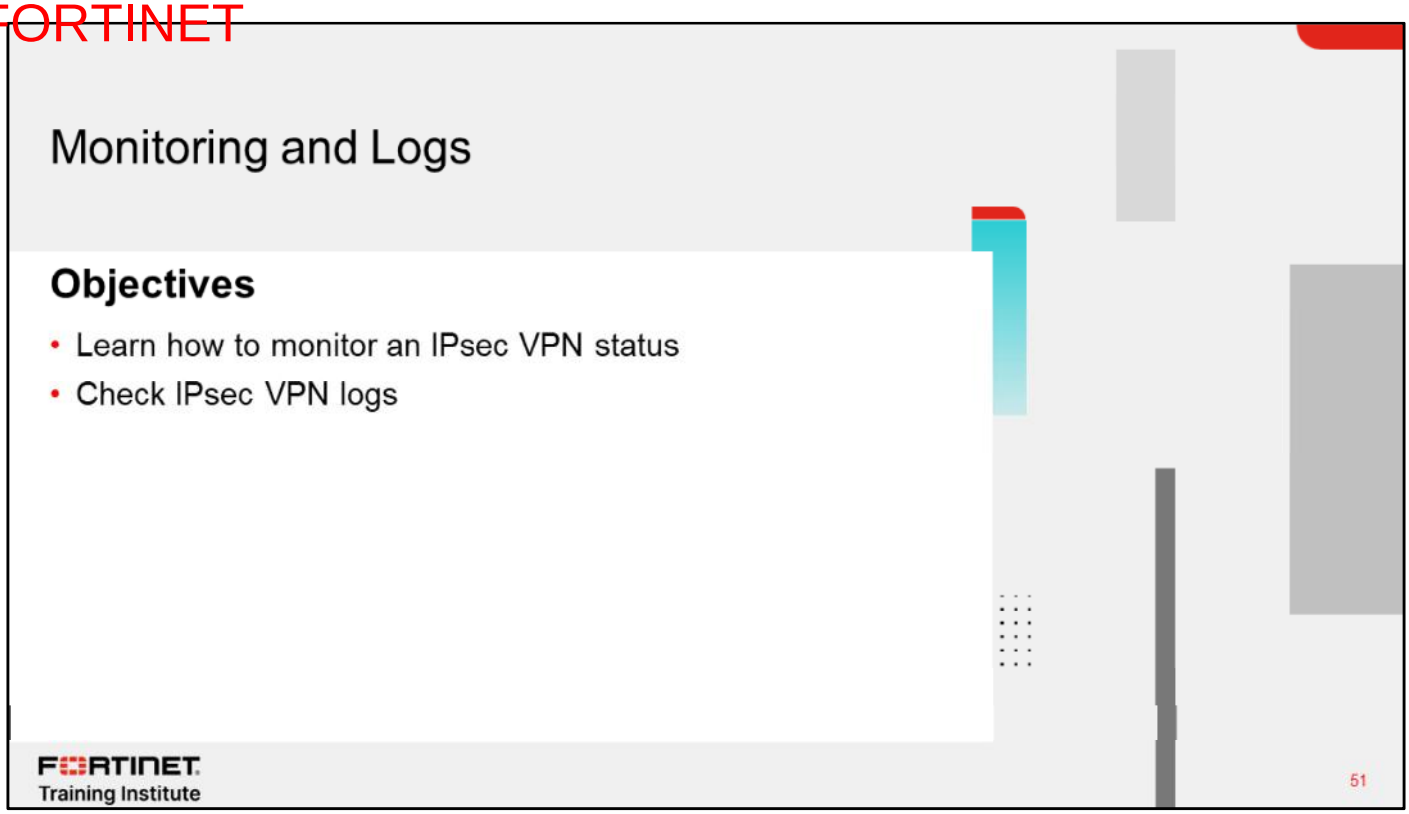

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in monitoring and logs, you will be able to monitor IPsec VPN and review past events.

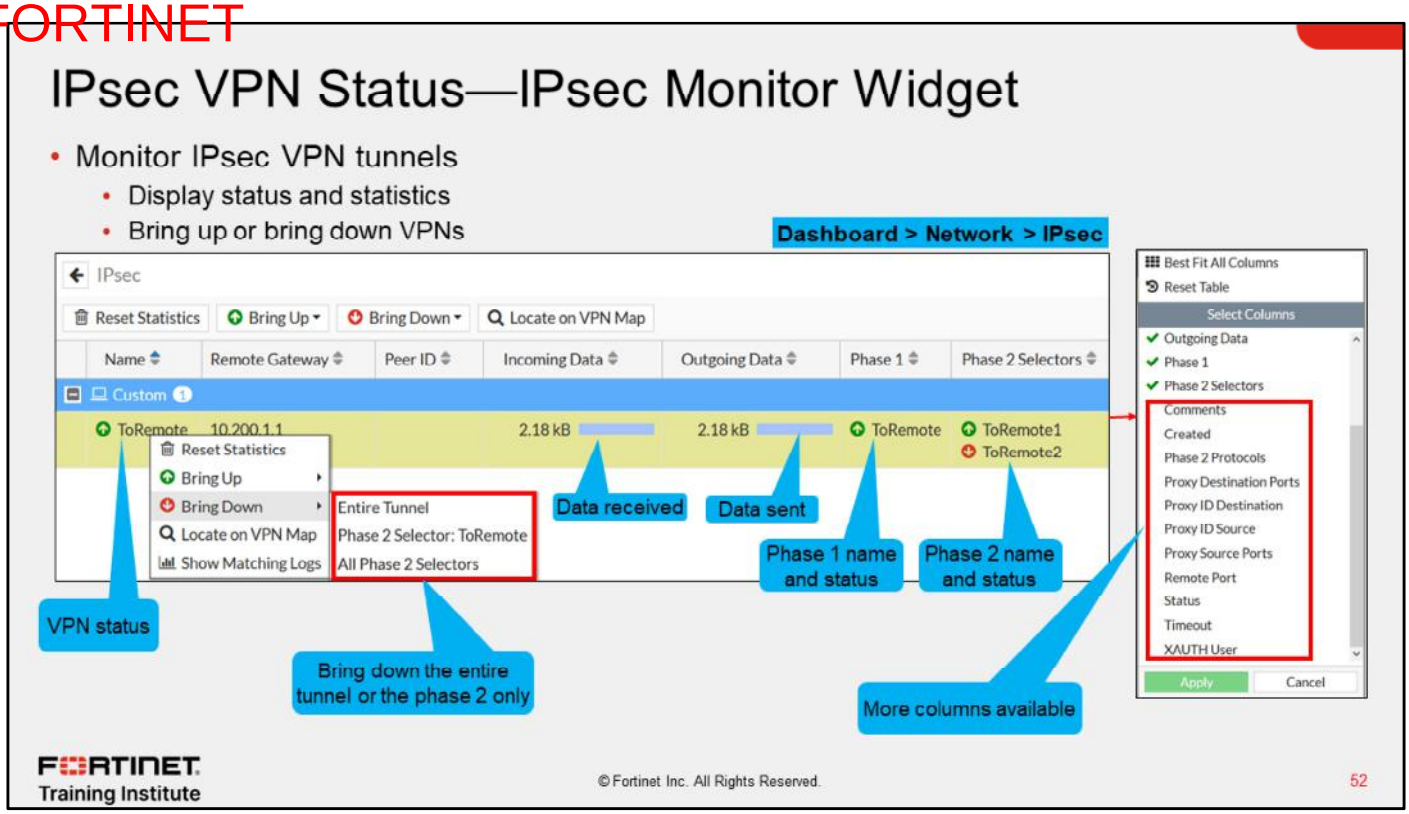

On the GUI dashboard, you can use the IPsec widget to monitor the status of your IPsec VPNs. The widget shows the phase 1 and phase 2 status of an IPsec VPN.

You can also bring up or bring down individual VPNs, and get additional details. When you bring up an IPsec VPN using the IPsec widget, you can choose between bringing up a particular phase 2 selector or all phase 2 selectors in that VPN. Because bringing up a phase 2 selector requires bringing up its phase 1 first, then bringing up a phase 2 selector results in its phase 1 also coming up.

To bring down the VPN, you can choose between bringing down a particular phase 2 selector, all selectors, or the entire tunnel. When you bring down the entire tunnel, you bring down all phase 2 selectors as well as the phase 1.

The **Name** column indicates the VPN status. The VPN is up when at least one of its phase 2 selectors is up. If all phase 2 selectors are down, the VPN status is also down. The **Phase 1** and **Phase 2 Selectors** columns indicate the status of phase 1 and phase 2 selectors, respectively.

The IPsec widget also displays the amount of data sent and received through the tunnel. When you right-click any of the columns, a menu opens with a list of all the columns available. You can enable additional columns to get further details about the IPsec tunnels.

In the example shown on this slide, the **ToRemote** VPN is up because at least one of its phase 2 selectors (**ToRemote1**) is up.

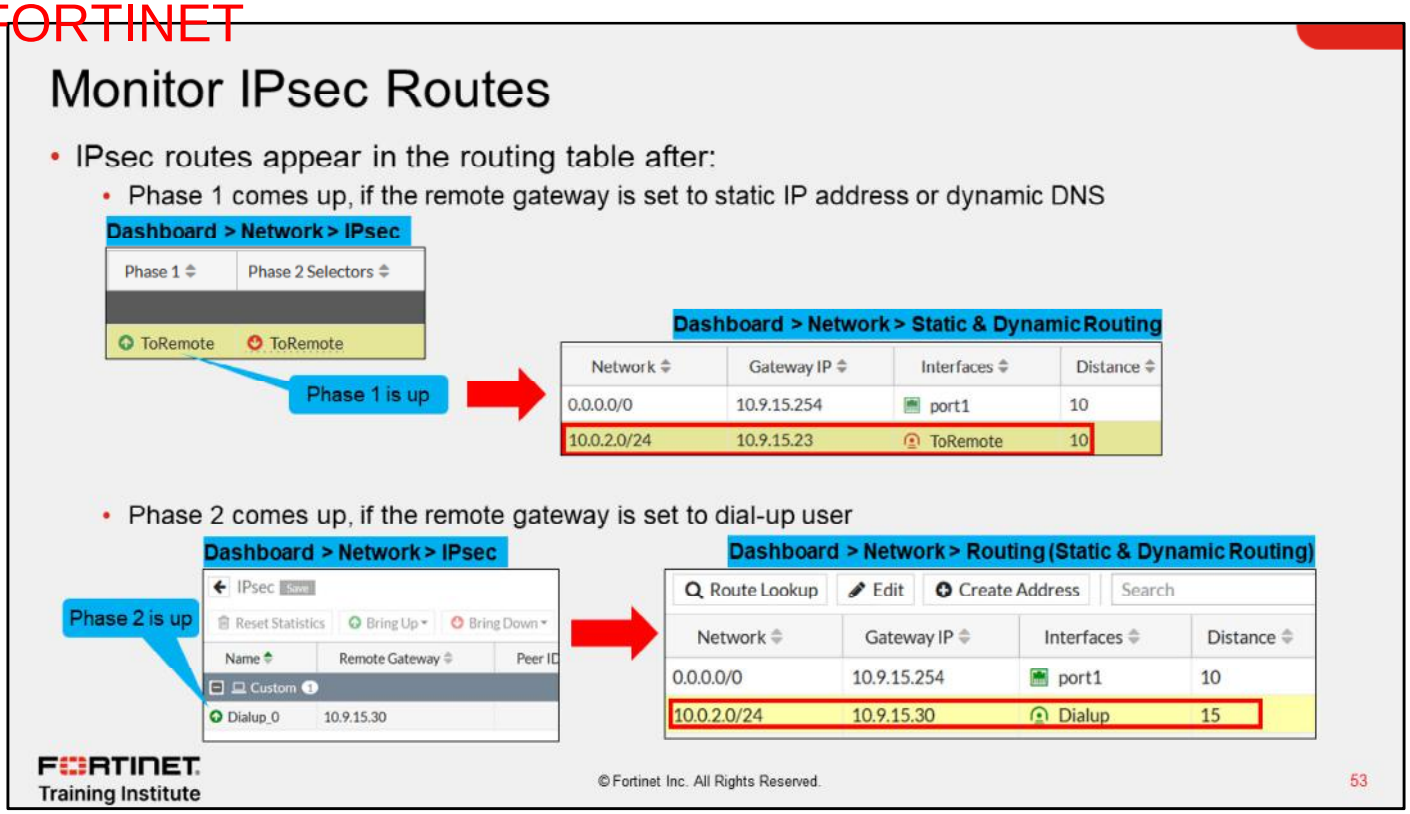

If you set the remote gateway to **Static IP Address** or **Dynamic DNS**, the static routes for these tunnels become active in the routing table after phase 1 comes up. Phase 1 negotiation is started automatically because automatic negotiation is enabled on phase 1 by default. This behavior allows FortiGate to match interesting traffic to the right tunnel. Moreover, if phase 2 is not up, traffic matching the static route triggers a phase 2 negotiation, which eventually results in the tunnel (or phase 2) to come up.

When you set the remote gateway to **Dialup User**, by default, a static route for the destination network is added after phase 2 comes up. The distance set for the static route is 15. If phase 2 goes down, the route is removed from the routing table.

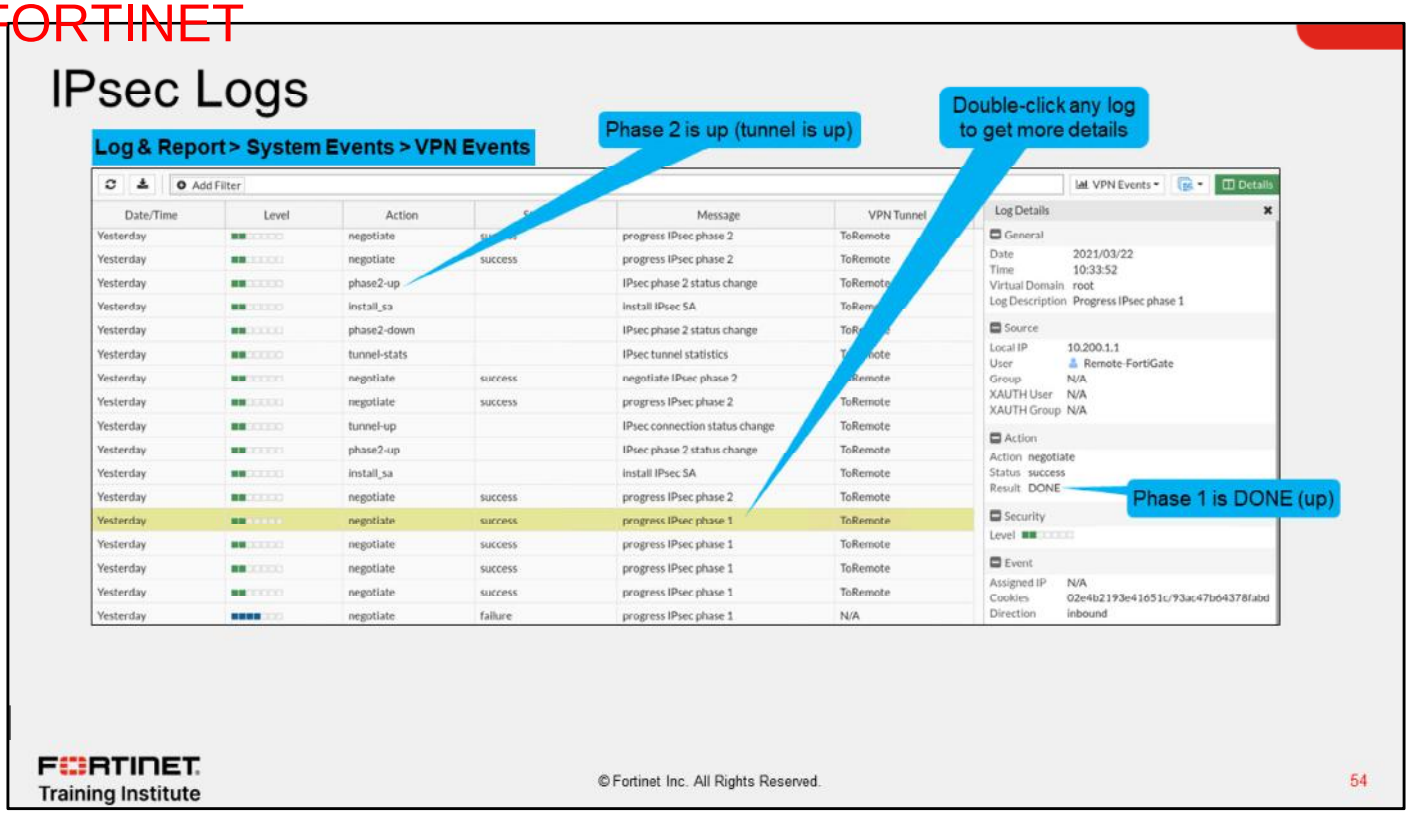

FortiGate logs IPsec VPN events by default. To view IPsec VPN event logs, click **Log & Report** > **System Events** > **VPN Events**.

The logs track the progress of phase 1 and phase 2 negotiations, and report on tunnel up and down events and DPD failures, among other events. For more information about IPsec logs, visit `https://docs.fortinet.com`.

Knowledge Check

1. The IPsec monitor widget enables you to bring down the _____ of an IPsec VPN.
   - A. Phase 1
   - ✓ B. Entire tunnel

2. When the remote gateway is set to dial-up user, a static route to the remote network is added to the routing table after _____.
   - A. Phase 1 comes up
   - ✓ B. Phase 2 comes up

FORTINET Training Institute

55

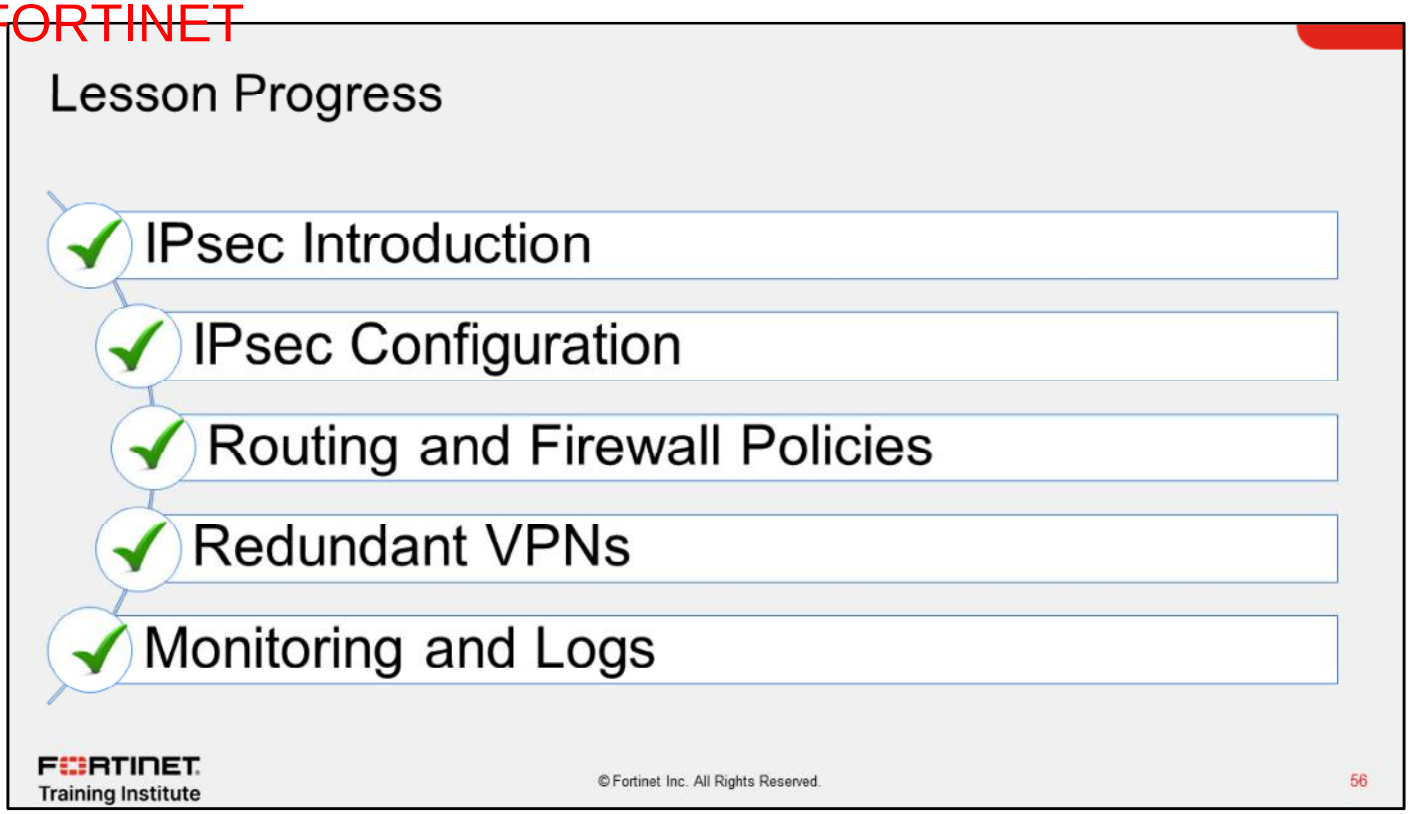

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

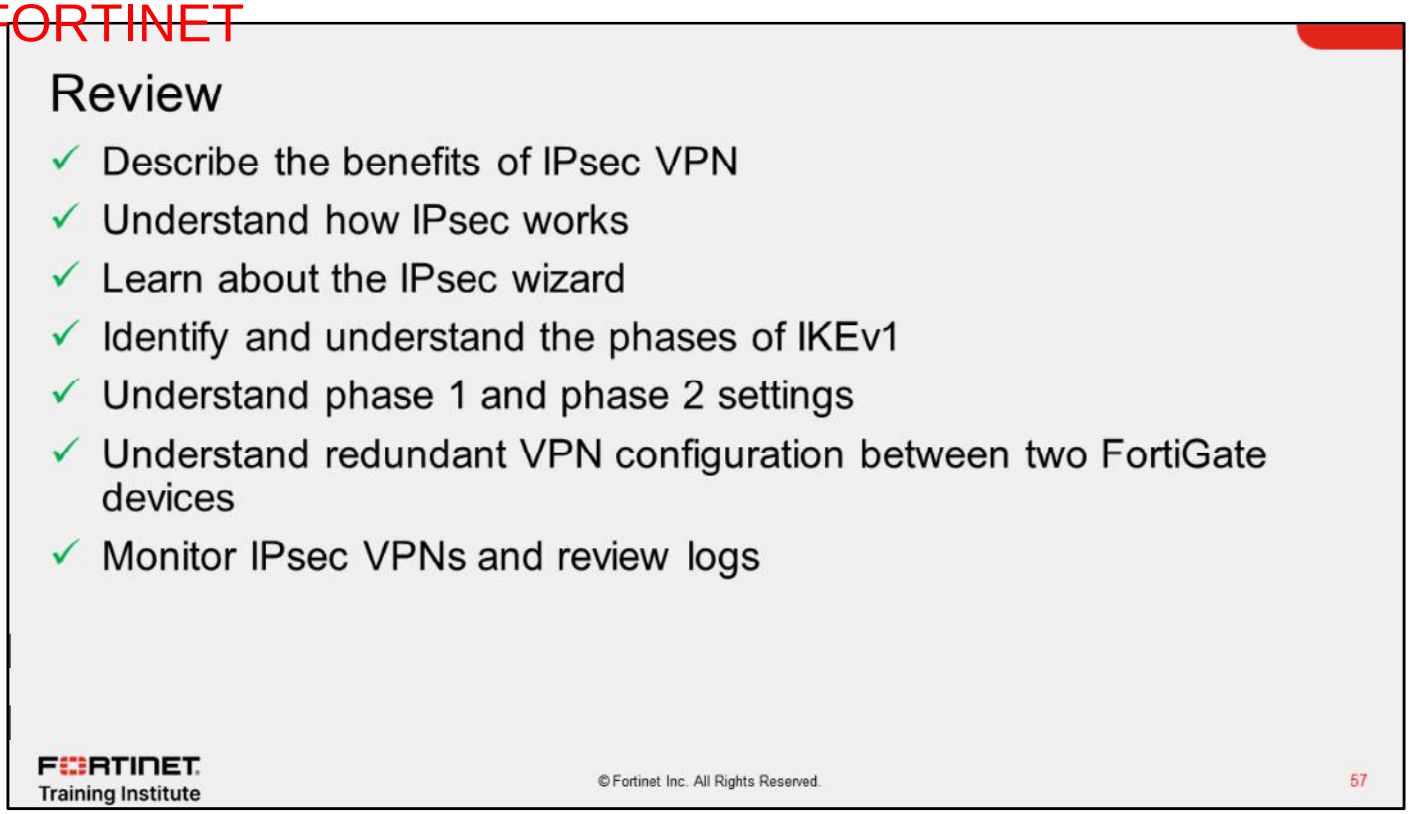

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how the IPsec protocol works, and how to configure and monitor IPsec VPNs on FortiGate.

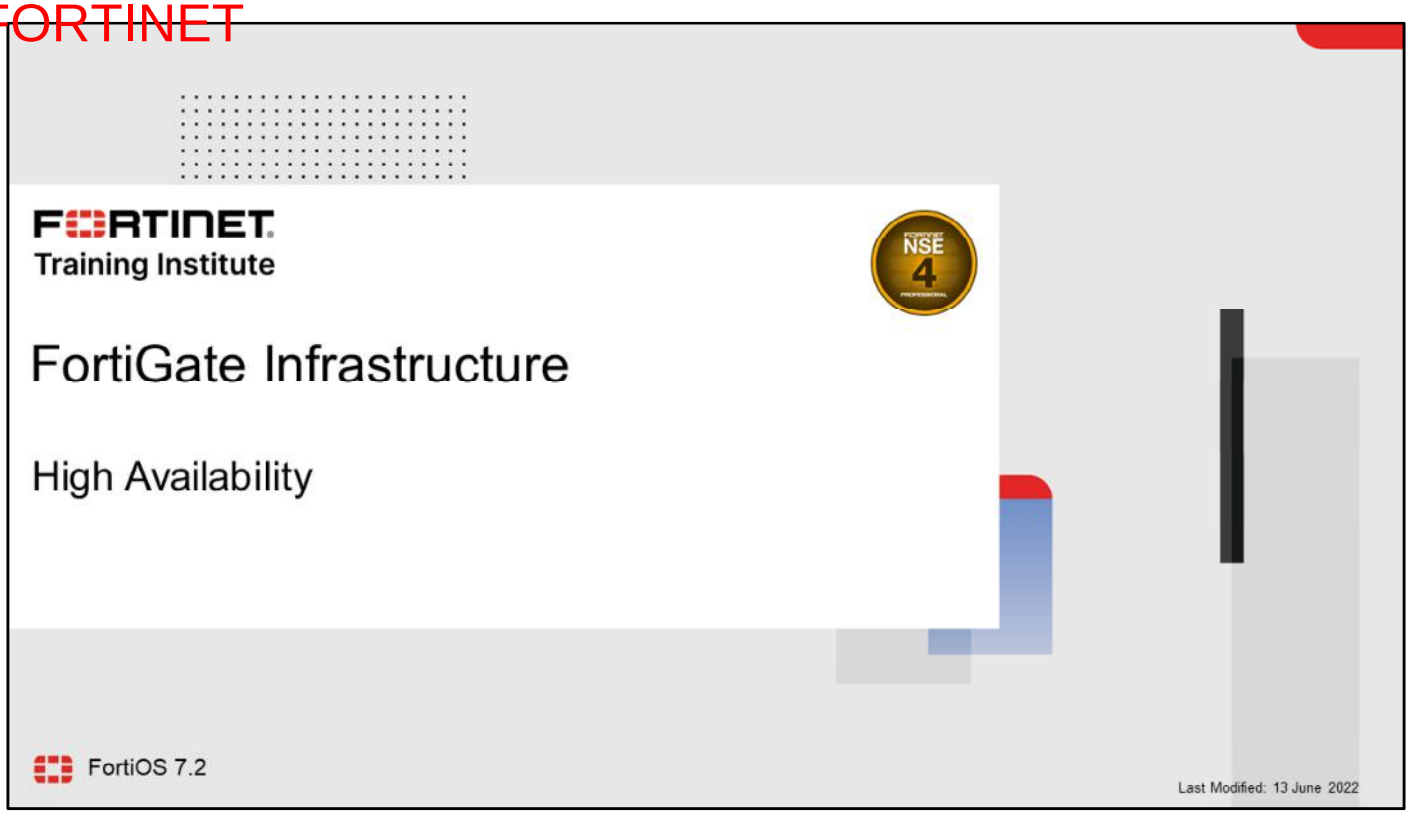

In this lesson, you will learn about the fundamentals of FortiGate high availability (HA) and how to configure it. FortiGate HA provides a solution for enhanced reliability and increased performance.

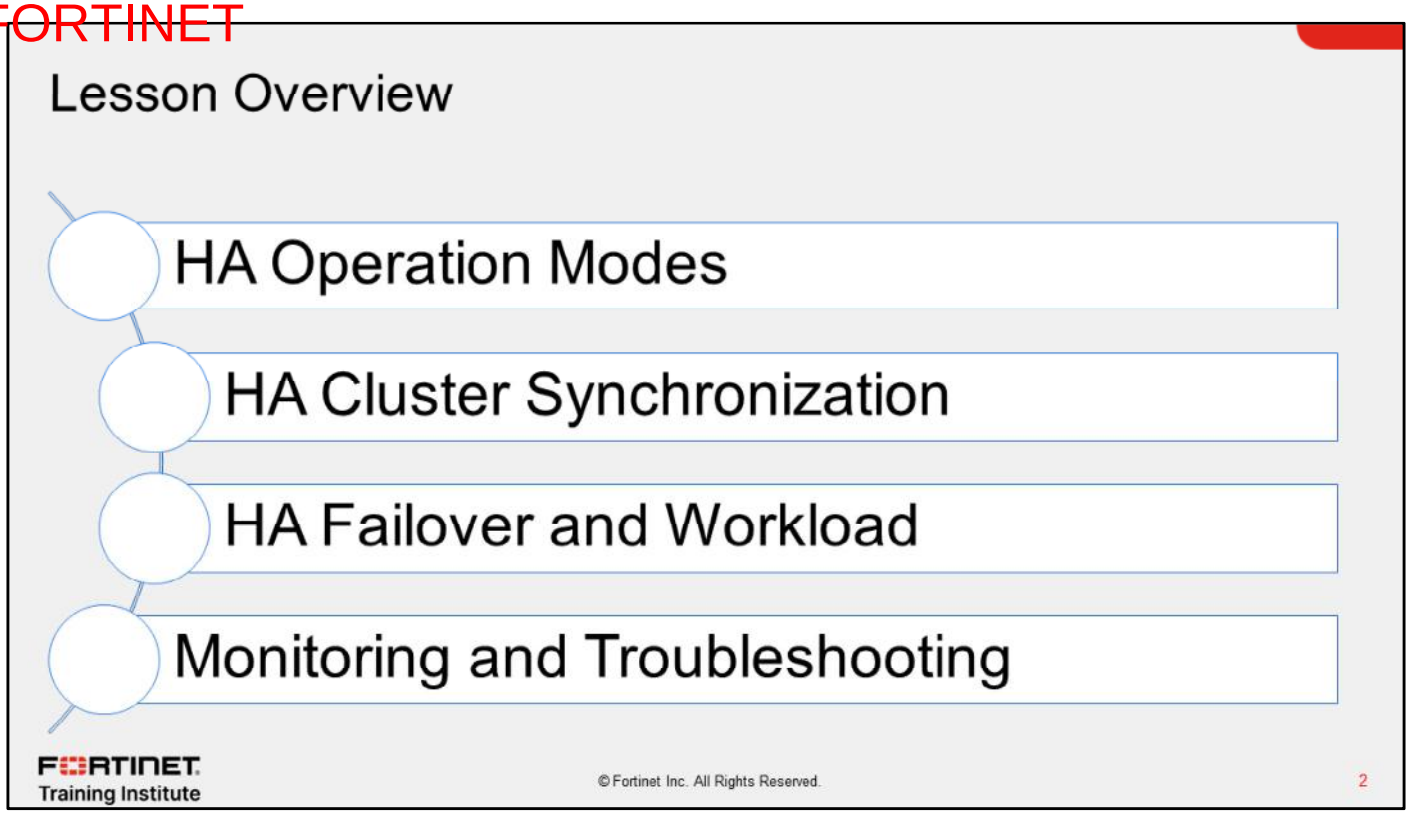

In this lesson, you will learn about the topics shown on this slide.

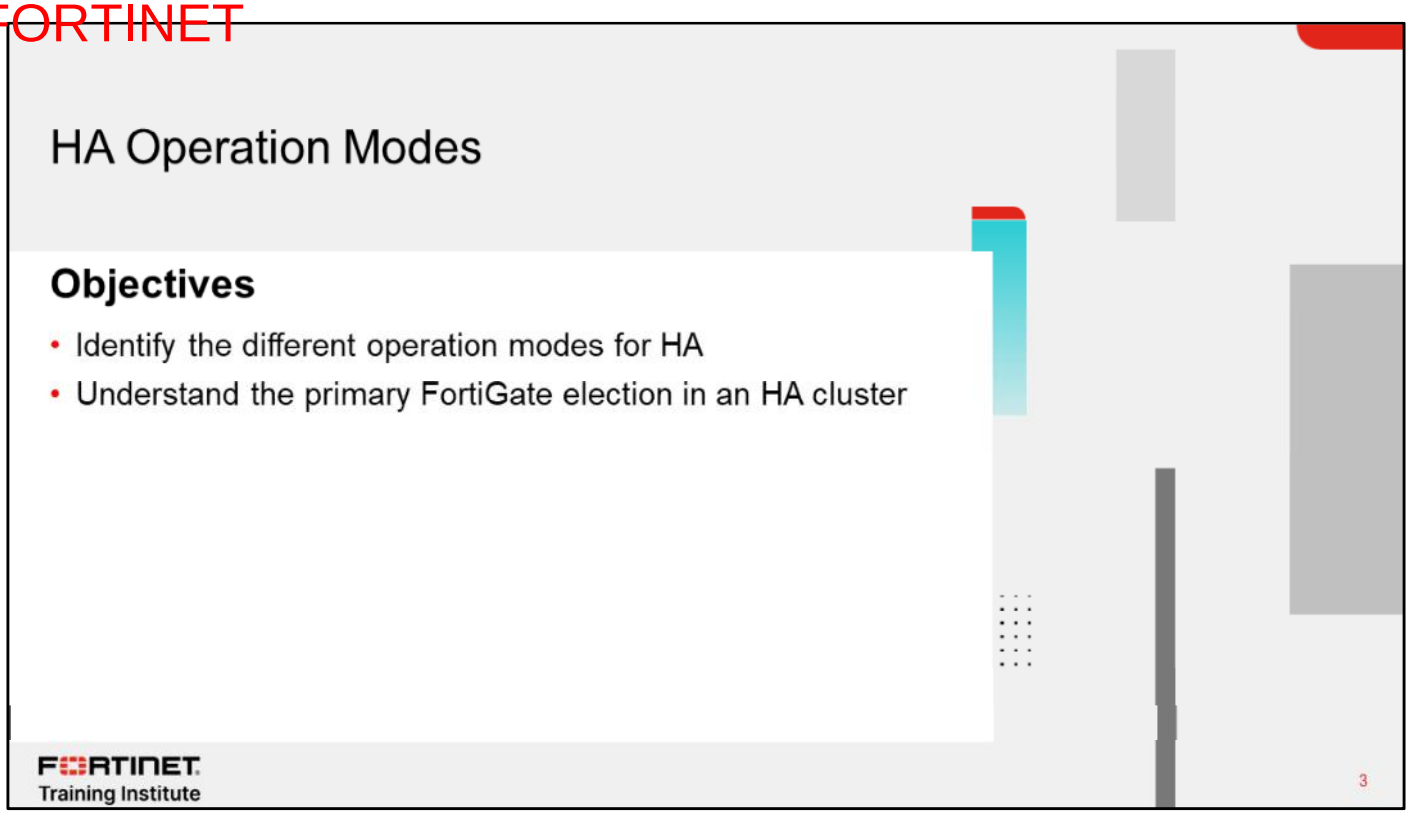

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in HA operation modes and primary FortiGate election, you will be able to choose and implement the right HA operation mode in your network based on your requirements. You will be able to use FortiGate devices effectively in your network.

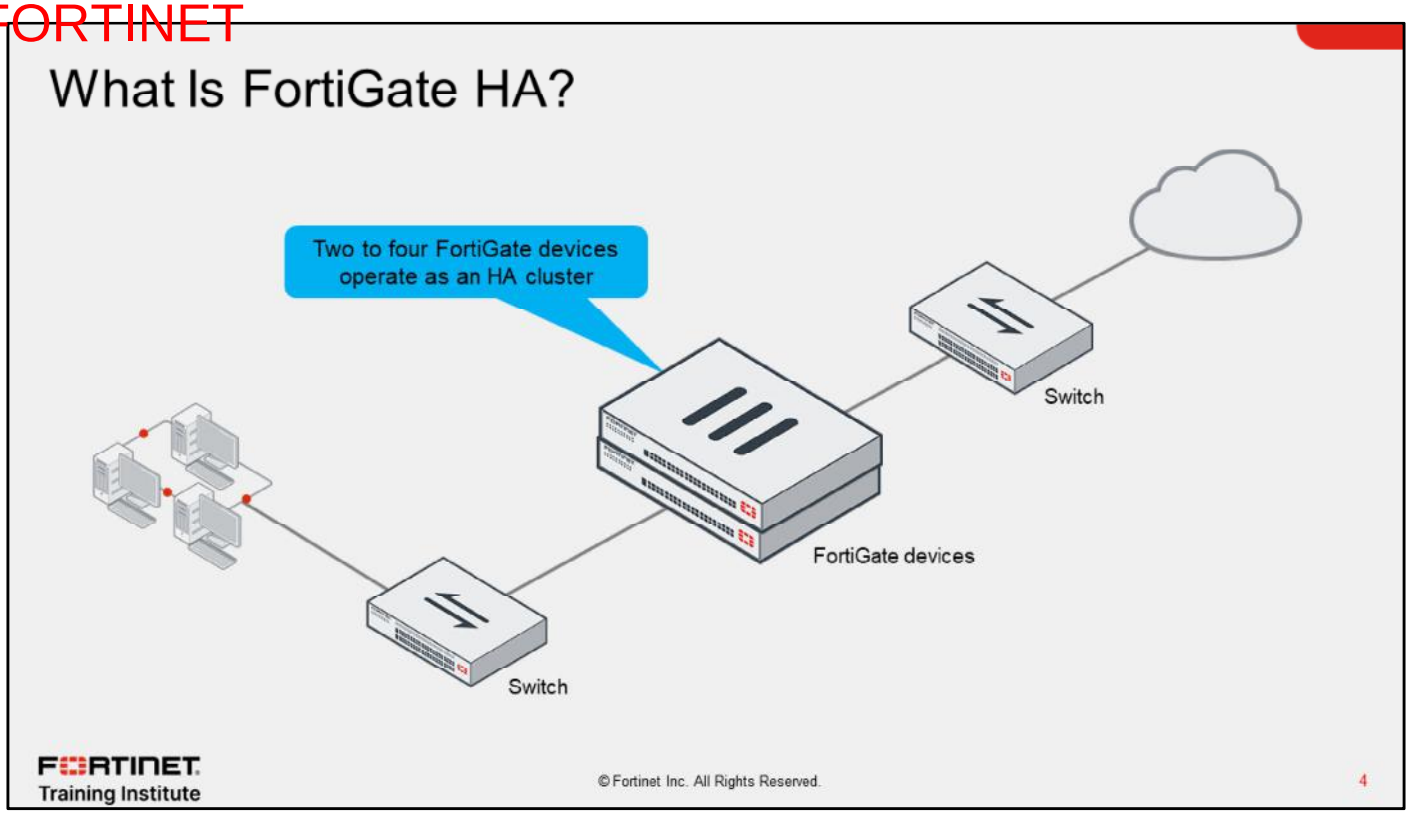

The idea of HA is simple. HA links and synchronizes two to four FortiGate devices to form a cluster for redundancy and performance purposes.

A cluster includes one device that acts as the primary FortiGate (also called the active FortiGate). The primary synchronizes its configuration, session information, FIB entries, FortiGuard definitions, and other operation-related information to the secondary devices, which are also known as standby devices.

The cluster shares one or more heartbeat interfaces among all devices—also known as members—for synchronizing data and monitoring the health of each member.

There are currently two HA operation modes available: active-active (A-A) and active-passive (A-P). Now, you will examine the differences.

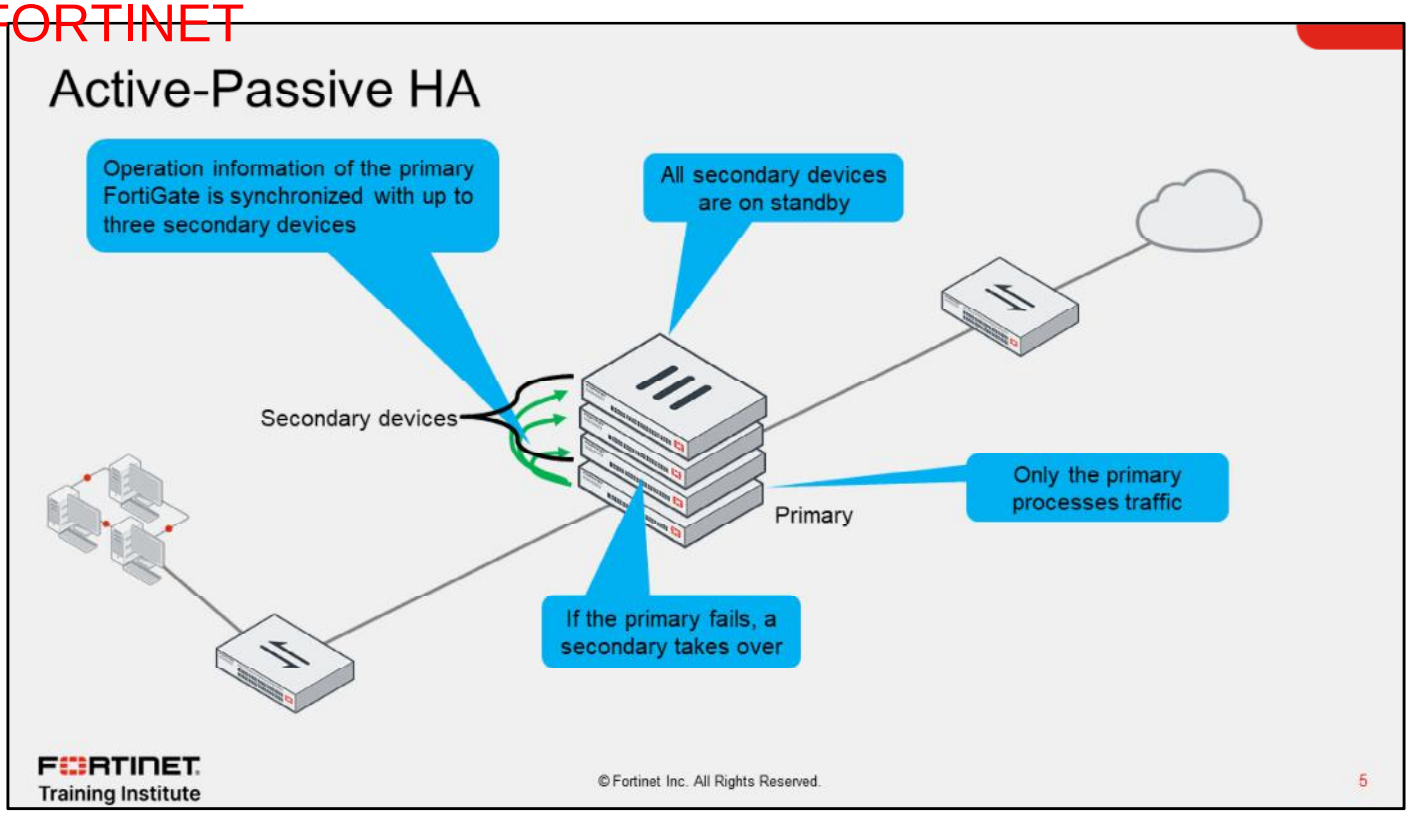

First, take a look at active-passive mode. In either of the two HA operation modes, the operation information (configuration, sessions, FIB entries, and so on) of the primary FortiGate is synchronized with secondary devices.

In active-passive mode, the primary FortiGate is the only FortiGate that actively processes traffic. Secondary FortiGate devices remain in passive mode, monitoring the status of the primary device.

If a problem is detected on the primary FortiGate, one of the secondary devices takes over the primary role. This event is called an *HA failover*.

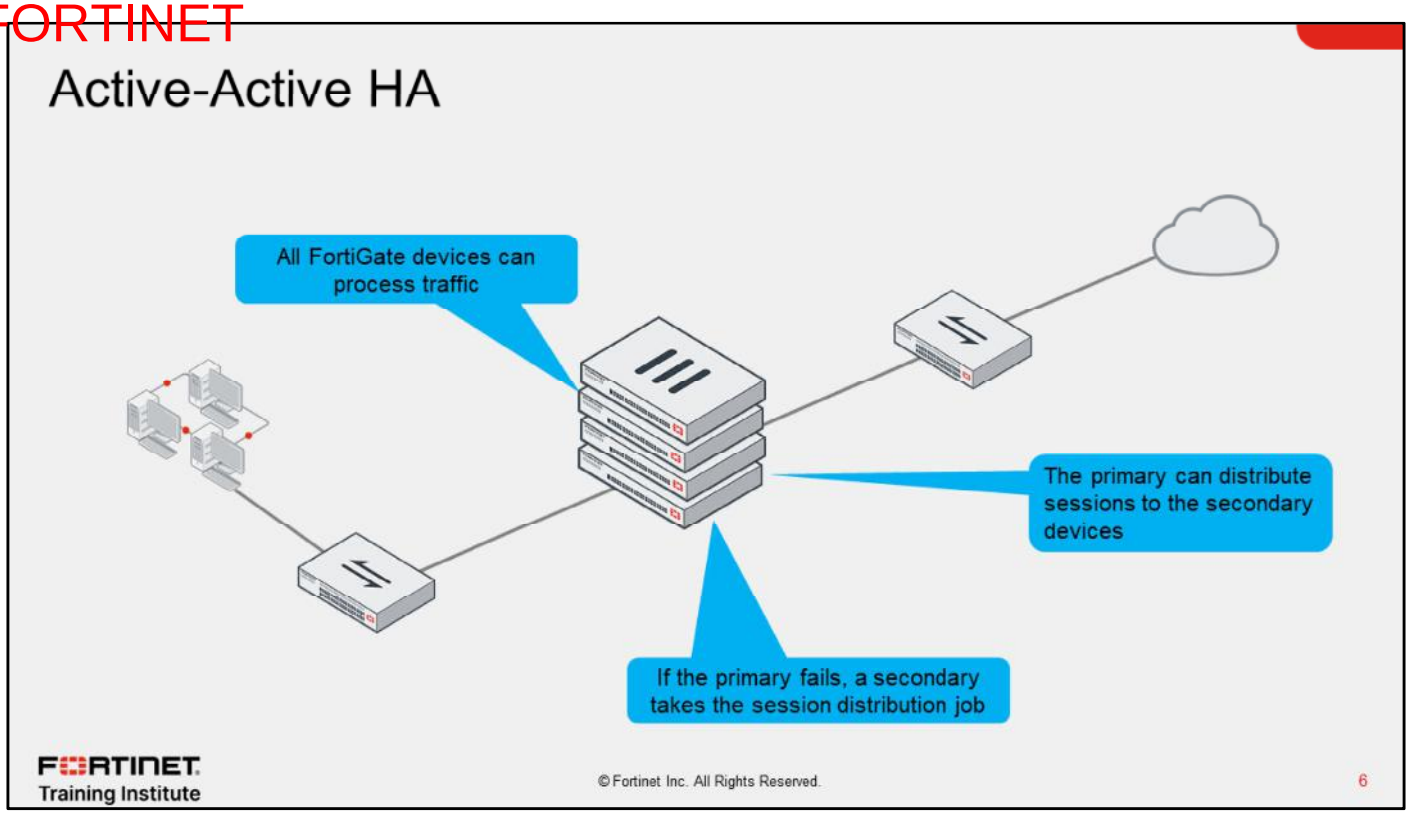

The other HA mode is active-active.

Like active-passive HA, in active-active HA, the operation-related data of the primary FortiGate is synchronized to the secondary FortiGate devices. Also, if a problem is detected on the primary device, one of the secondary devices takes over the role of the primary, to process the traffic.

However, one of the main differences from active-passive mode is that in active-active mode, all cluster members can process traffic. That is, based on the HA settings and traffic type, the primary FortiGate can distribute sessions to the secondary devices.

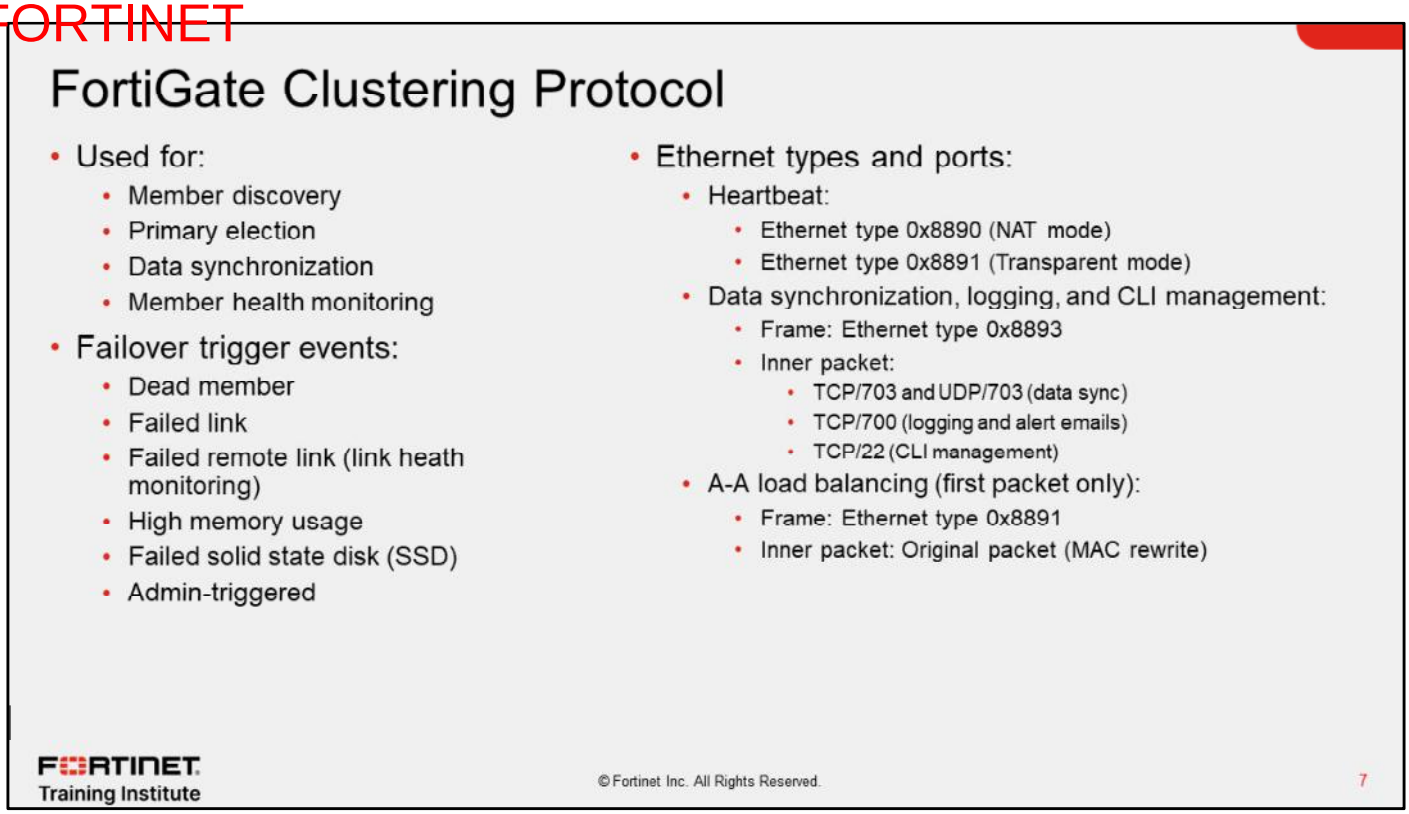

FortiGate HA uses the Fortinet-proprietary FortiGate Clustering Protocol (FGCP) to discover members, elect the primary FortiGate, synchronize data among members, and monitor the health of members.

To discover and monitor members, the members broadcast heartbeat packets over all configured heartbeat interfaces. If the cluster operates in NAT mode, the heartbeat frames are type 8890. In transparent mode, the heartbeat frames are type 8891. If the cluster operates in active-active mode, the first packet of a session distributed to the secondary is encapsulated in Ethernet frames type 8891.

The members also exchange frames type 8893 for data synchronization, local CLI management, and logging purposes. For data synchronization, the inner packet can be TCP port 703 or UDP port 703, depending on the type of data to synchronize. The primary also relays logs and alert emails from secondary devices over TCP port 700. For local HA management using the CLI, the inner packet is SSH.

You can configure the cluster to perform HA failover based on the following events:

- Dead member: The primary FortiGate is unresponsive.
- Failed link: The link of one or more monitored interfaces on the primary FortiGate goes down.
- Failed remote link: FortiGate uses the link health monitor feature to monitor the health of one or more interfaces. The primary fails if the accumulated penalty of all failed interfaces reaches the set threshold.
- High memory usage: The primary fails if its memory utilization reaches the configured threshold.
- Failed SSD: FortiOS detects a failure in an SSD. Only available for devices with SSDs.
- Admin-triggered: The administrator issues a manual failover.

For any of the failover trigger events, the result is that the cluster promotes one of the secondary devices to the new primary FortiGate role.

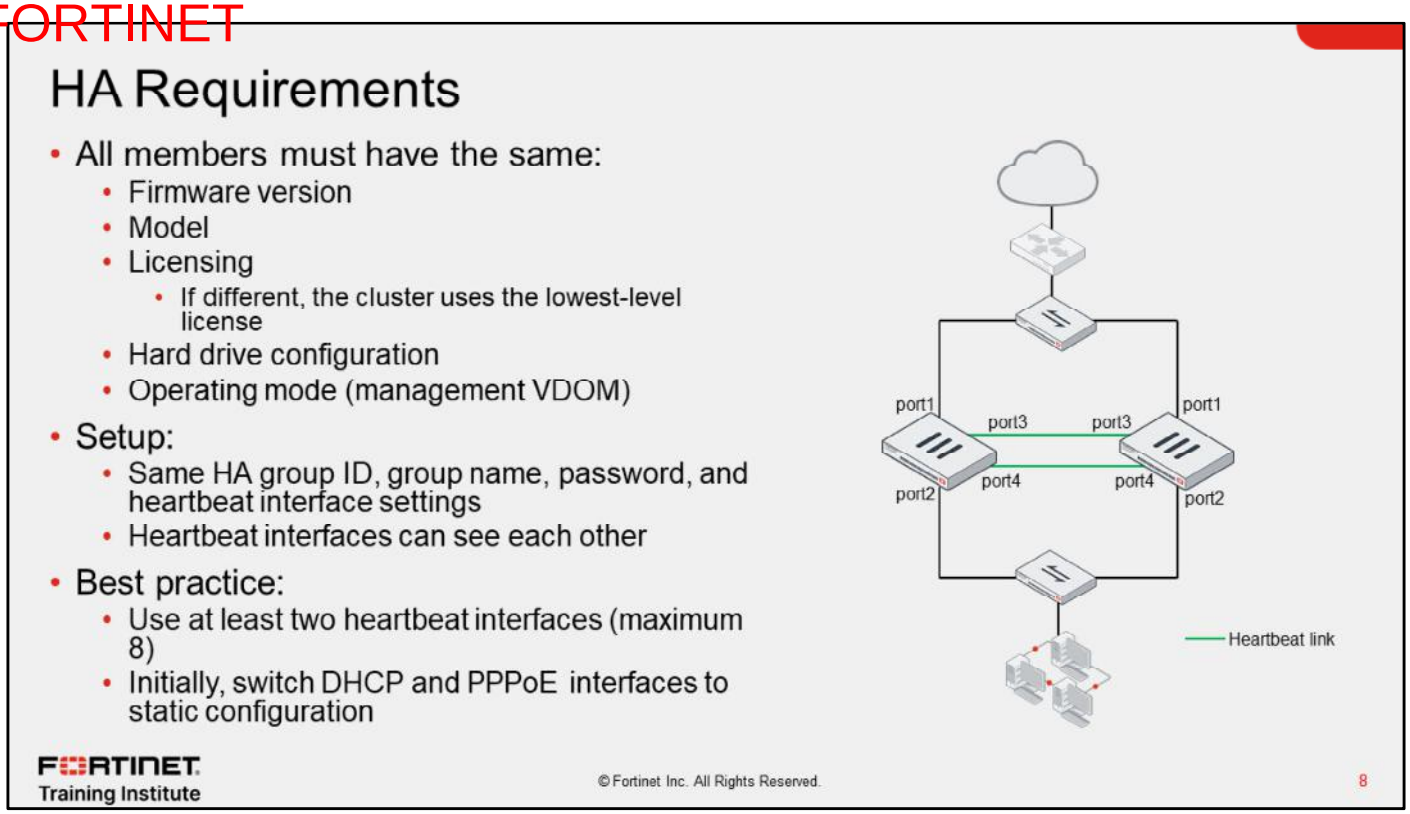

To successfully form an HA cluster, you must ensure that the members have the same:

- Firmware version
- Model: the same hardware model or VM model
- Licensing: includes the FortiGuard license, VDOM license, FortiClient license, and so on
- Hard drive configuration: the same number and size of drives and partitions
- Operating mode: the operating mode—NAT mode or transparent mode—of the management VDOM

If the licensing level among members isn't the same, the cluster resolves to use the lowest licensing level among all members. For example, if you purchase FortiGuard Web Filtering for only one of the members in a cluster, none of the members will support FortiGuard Web Filtering when they form the cluster.

From a configuration and setup point of view, you must also make sure that:

- The HA settings on each member have the same group ID, group name, password, and heartbeat interface settings.
- The heartbeat interfaces on each member can see each other. This usually means placing all heartbeat interfaces in the same broadcast domain, or for two-member clusters, connecting them directly.

It's also a best practice to:

- Configure at least two heartbeat interfaces for redundancy purposes. This way, if one heartbeat link fails, the cluster uses the next one, as indicated by the priority and position in the heartbeat interface list.
- If using DHCP or PPPoE interfaces, use static configuration during the cluster initial setup to prevent incorrect address assignment. After the cluster is formed, you can put back the original interface settings.

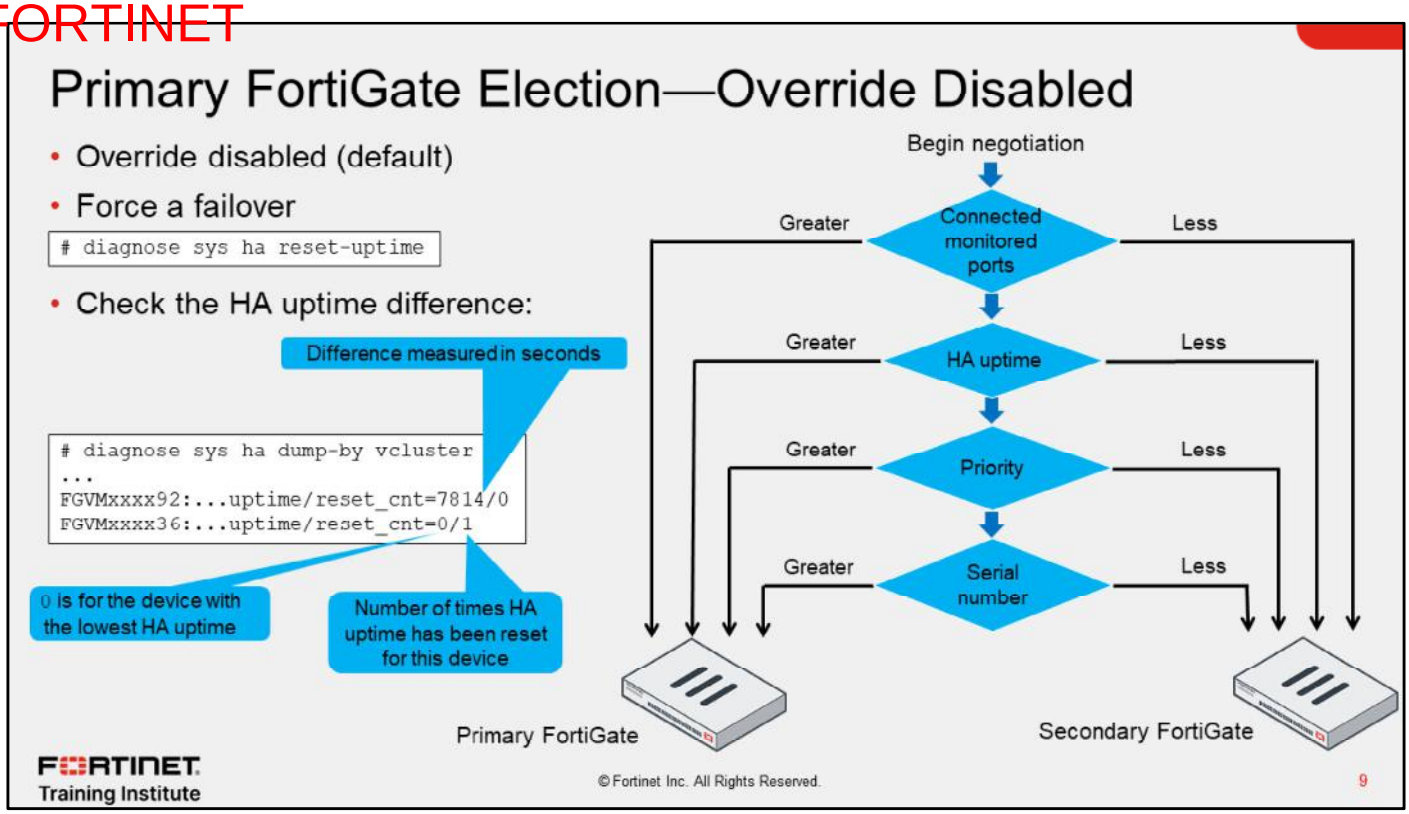

This slide shows the different criteria that a cluster considers during the primary FortiGate election process. The criteria order evaluation depends on the HA override setting. This slide shows the order when the HA override setting is disabled, which is the default behavior. Note that the election process stops at the first matching criteria that successfully selects a primary FortiGate in a cluster.

1. The cluster compares the number of monitored interfaces that have a status of up. The member with the most available monitored interfaces becomes the primary.
2. The cluster compares the HA uptime of each member. The member with the highest HA uptime, by at least five minutes, becomes the primary.
3. The member with the highest priority becomes the primary.
4. The member with the lowest serial number becomes the primary.

When HA override is disabled, the HA uptime has precedence over the priority setting. This means that if you must manually fail over to a secondary device, you can do so by reducing HA uptime of the primary FortiGate. You can do this by running the `diagnose sys ha reset-uptime` command on the primary FortiGate, which resets its HA uptime to 0.

Note that the `diagnose sys ha reset-uptime` command resets the HA uptime and not the system uptime. Also, note that if a monitoring interface fails, or a member reboots, the HA uptime for that member is reset to 0.

This slide also shows how to identify the HA uptime difference between members. The member with `0` in the `uptime` column indicates the device with the lowest uptime. The example shows that the device with the serial number ending in `92` has an HA uptime that is `7814` seconds higher than the other device in the HA cluster. The `reset_cnt` column indicates the number of times the HA uptime has been reset for that device.

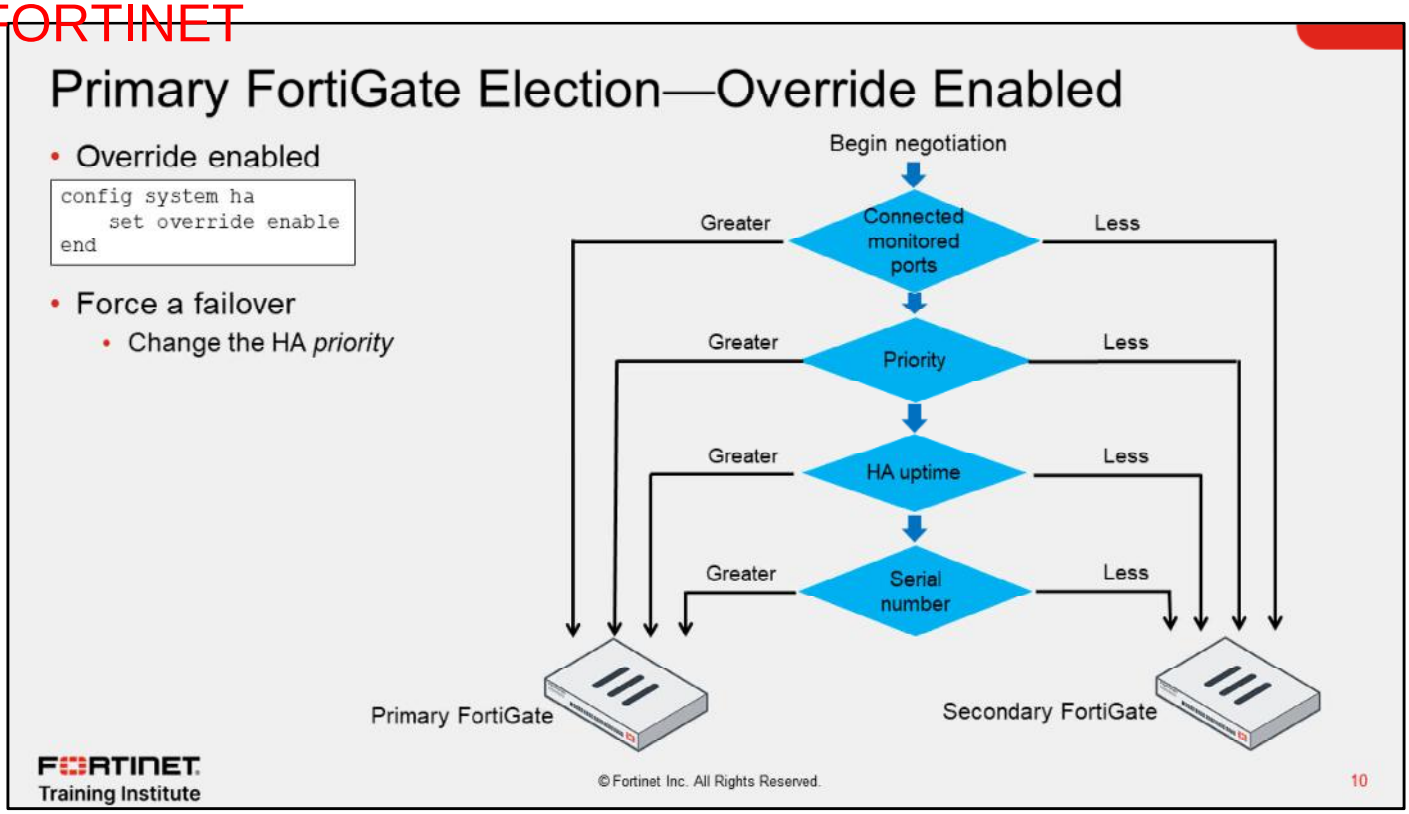

If the HA override setting is enabled, the priority is considered before the HA uptime.

The advantage of this method is that you can specify which device is the preferred primary every time (as long as it is up and running) by configuring it with the highest HA priority value. The disadvantage is that a failover event is triggered not only when the primary fails, but also when the primary is available again. That is, when the primary becomes available again, it takes its primary role back from the secondary FortiGate that temporarily replaced it.

When override is enabled, the easiest way of triggering a failover is to change the HA priorities. For example, you can either increase the priority of one of the secondary devices, or decrease the priority of the primary.

The override setting and device priority values are not synchronized to cluster members. That is, on each member, you must manually enable override and adjust the priority.

Knowledge Check

1. What is a requirement for members to form an HA cluster?
   - A. They must have same host name
   - ✓ B. They must run the same firmware version

2. What is the default order criteria (override disabled) for selecting the primary in an HA cluster?
   - ✓ A. Connected monitored ports > HA uptime > priority > serial number
   - B. Priority > HA uptime > connected monitored ports > serial number

FORTINET Training Institute

11

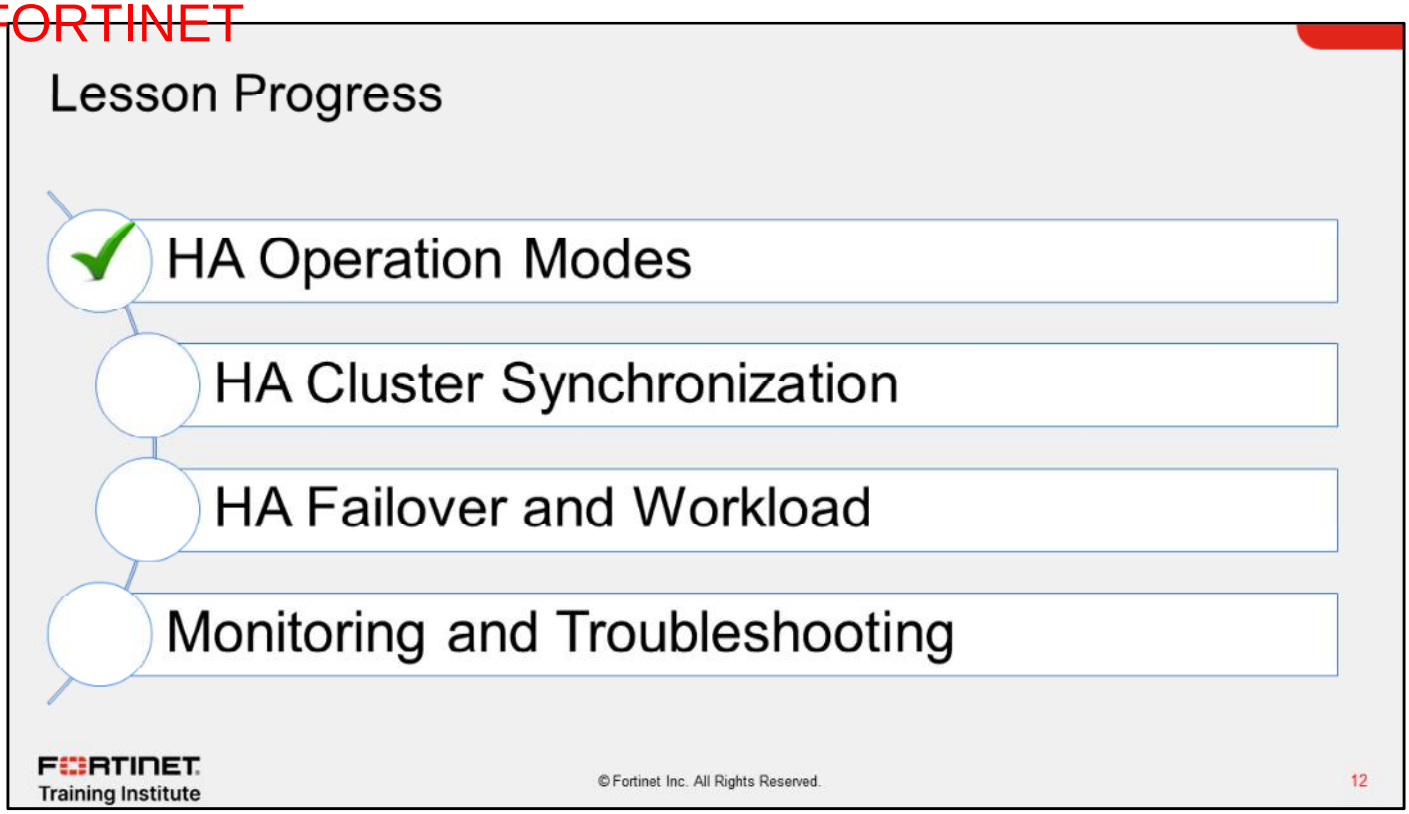

Good job! You now understand HA operation modes and the election of the primary FortiGate in an HA cluster.

Now, you will learn about HA cluster synchronization.

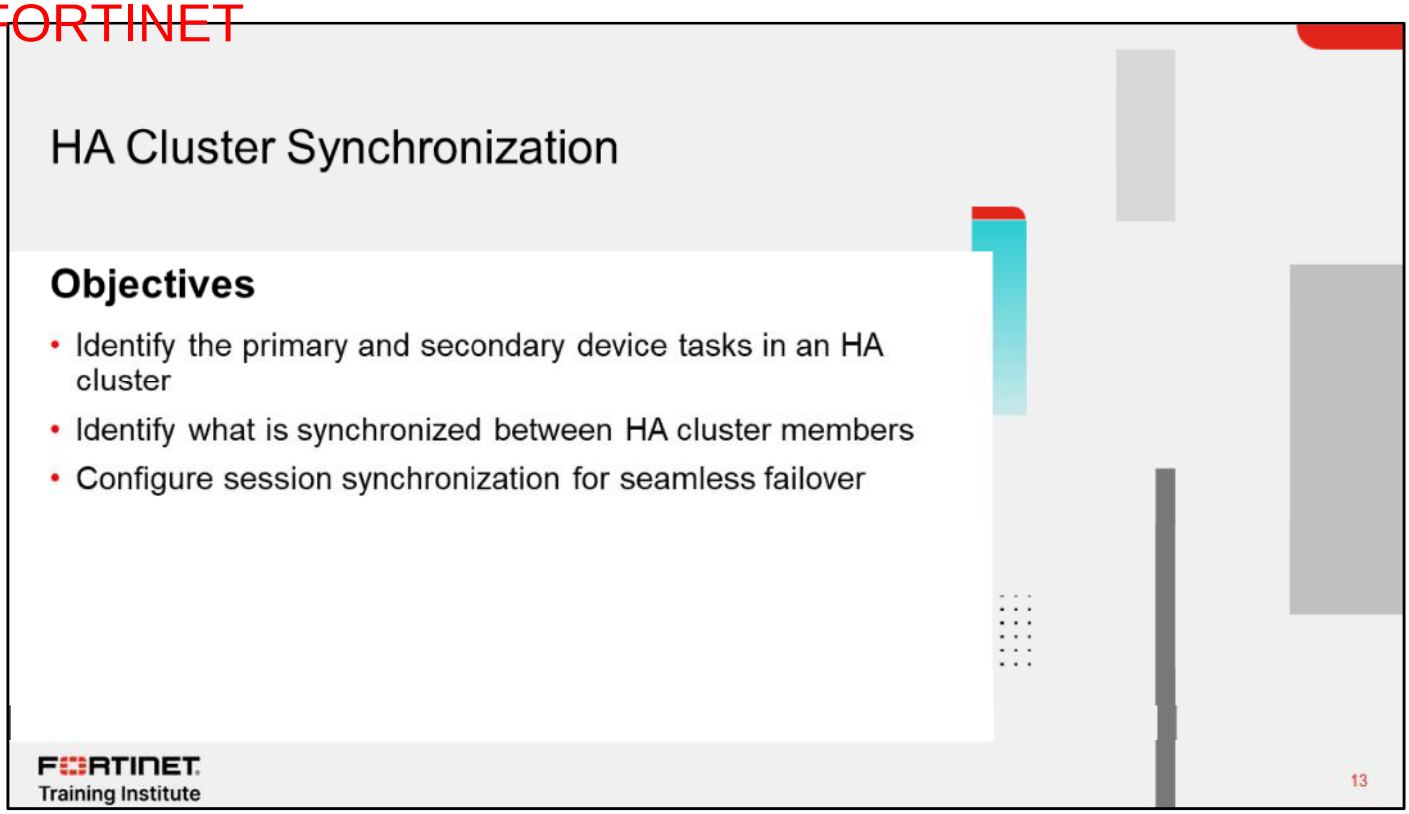

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in cluster synchronization, you will be able to identify the tasks assigned to members based on their roles, as well as what information is synchronized between members. You will also learn how to configure session synchronization to perform session failover to specific types of traffic.

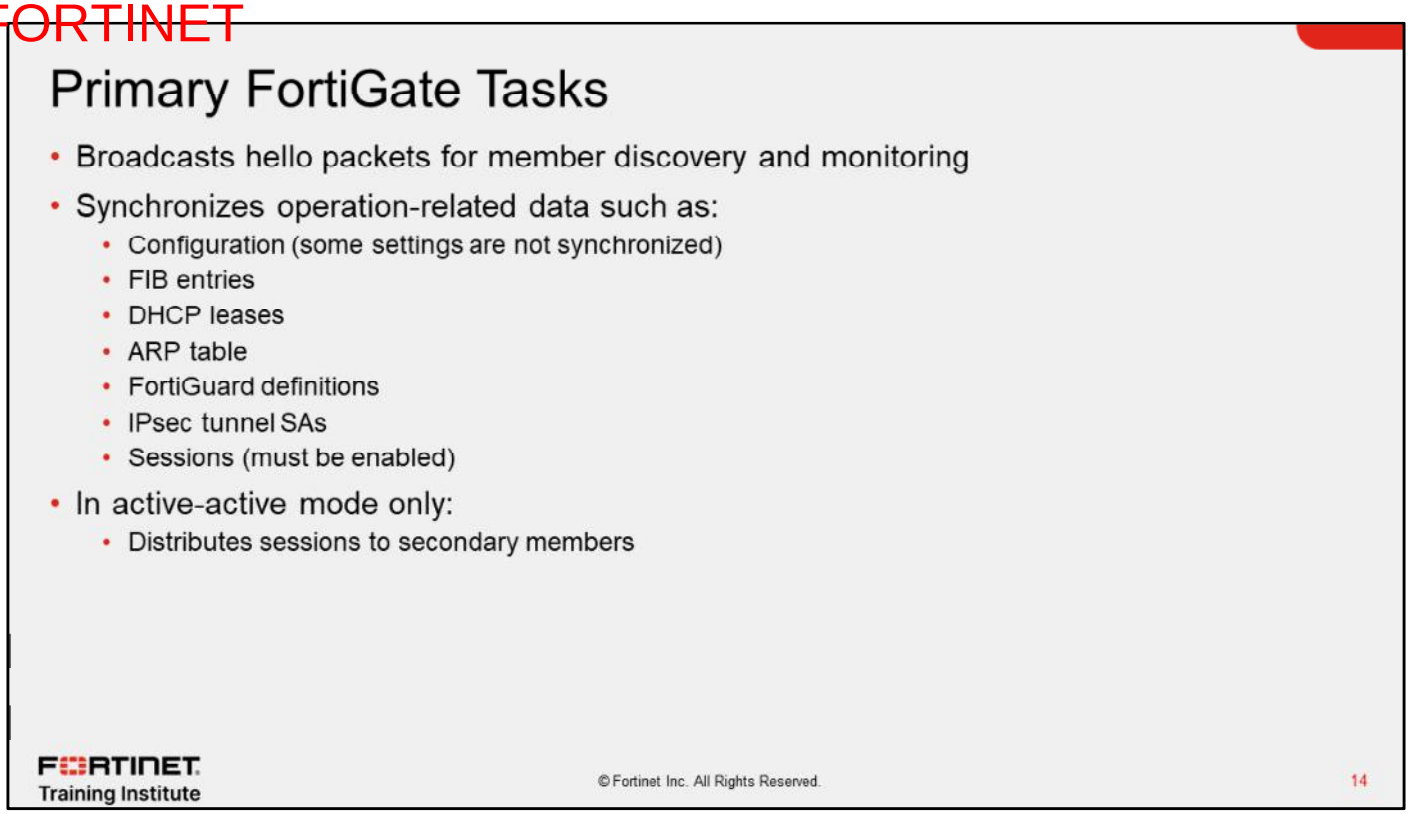

So, what are the tasks of a primary FortiGate?

It monitors the cluster by broadcasting hello packets and listening for hello packets from other members in the cluster. The members use the hello packets to identify if other FortiGate devices are alive and available.

The primary FortiGate also synchronizes its operation-related data to the secondary members. Some of the data synchronized includes its configuration, FIB entries, DHCP leases, ARP table, FortiGuard definitions, and IPsec tunnel security associations (SAs). Note that some parts of the configuration are not synchronized because they are device-specific. For example, the host name, HA priority, and HA override settings are not synchronized.

Optionally, you can configure the primary FortiGate to synchronize qualifying sessions to all the secondary devices. When you enable session synchronization, the new primary can resume communication for sessions after a failover event. The goal is for existing sessions to continue flowing through the new primary FortiGate with minimal or no interruption. You will learn which types of sessions you can enable synchronization for later in the lesson.

In active-active mode only, a primary FortiGate is also responsible for distributing sessions to secondary members.

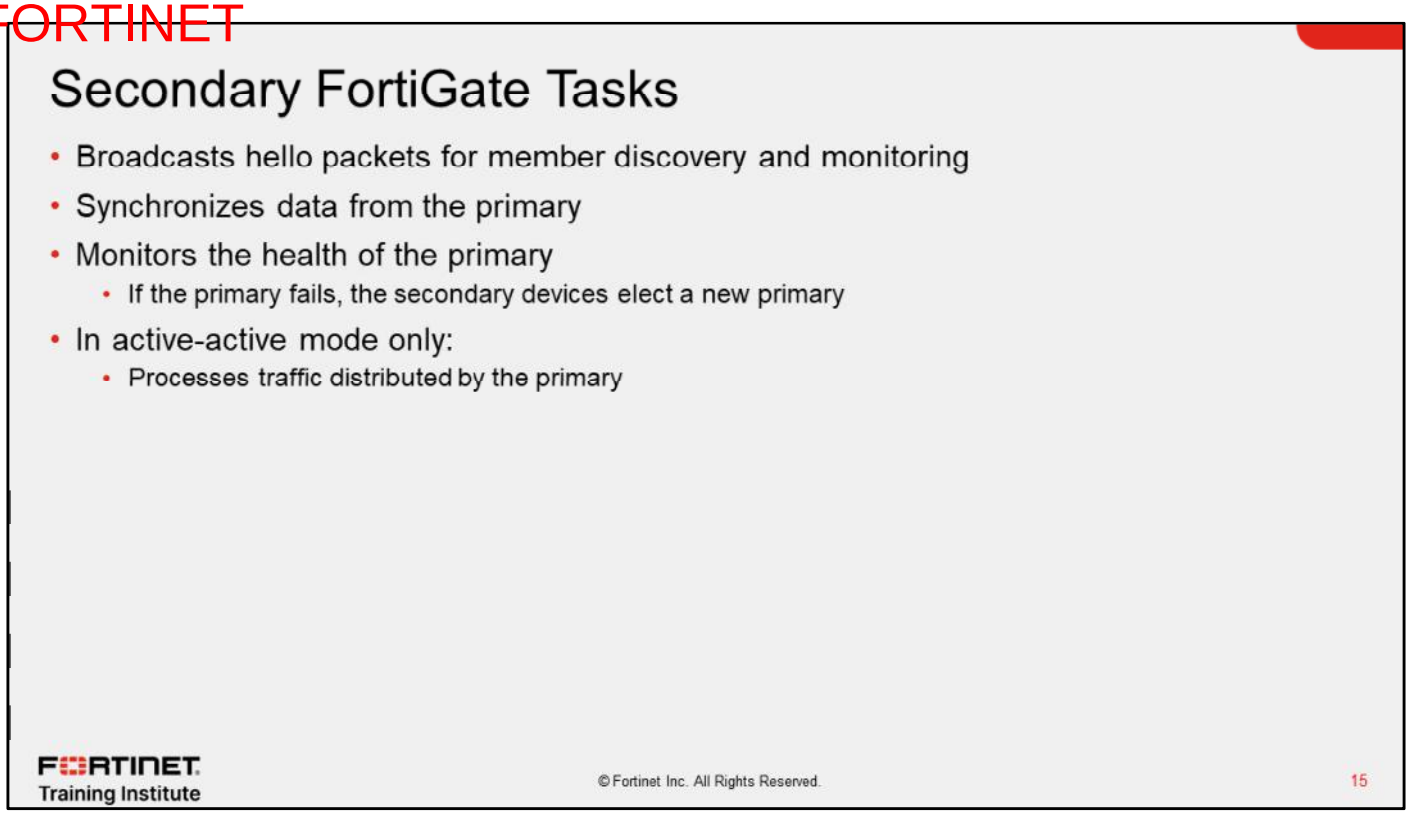

Now, take a look at the tasks of secondary FortiGate devices.

Like the primary, secondary members also broadcast hello packets for discovery and monitoring purposes.

In addition, in active-passive mode, the secondary devices act as a standby device, receiving synchronization data but not actually processing any traffic. If the primary FortiGate fails, the secondary devices elect a new primary.

In active-active mode, the secondary devices don’t wait passively. They process all traffic assigned to them by the primary device.

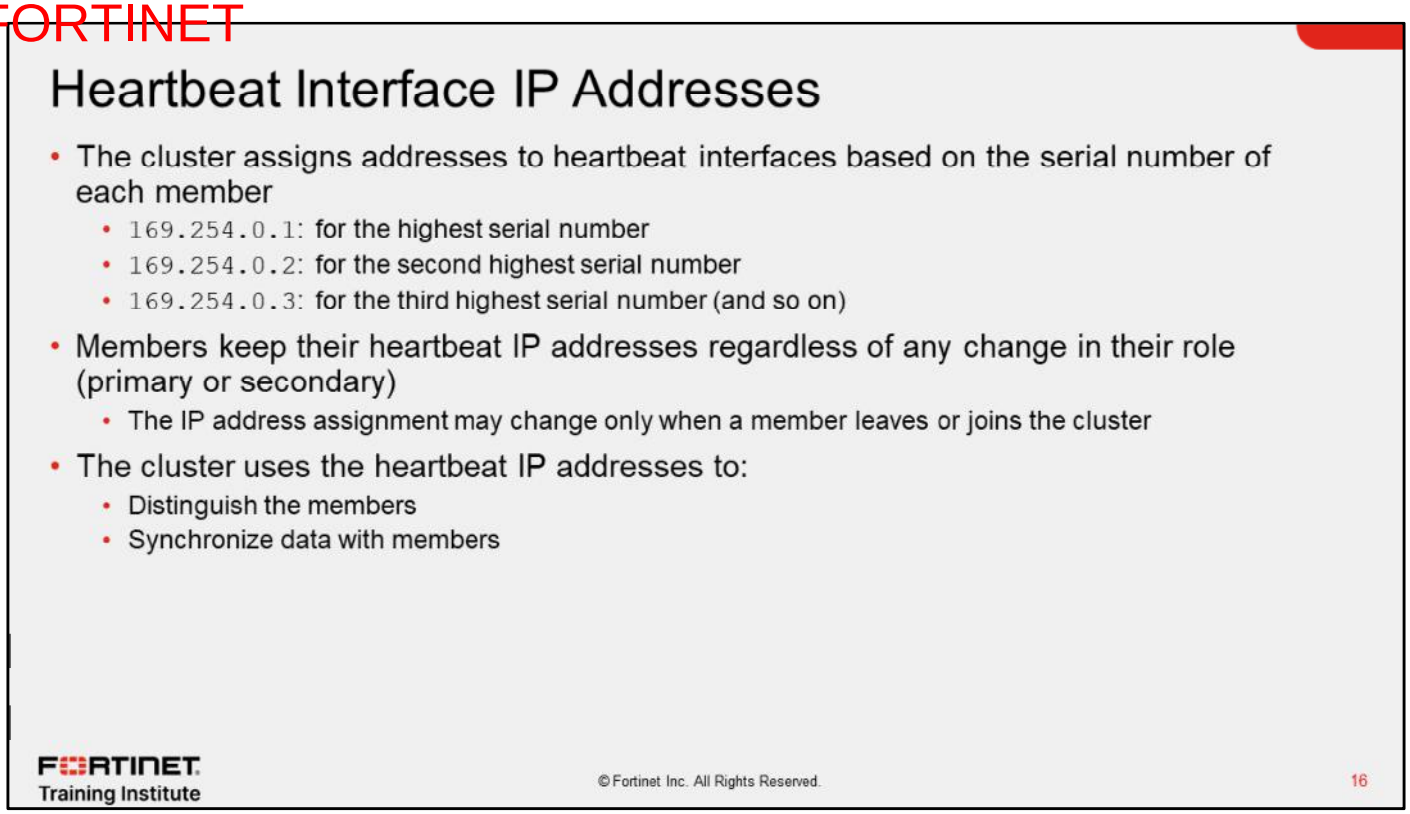

FGCP automatically assigns the heartbeat IP addresses based on the serial number of each device. The IP address 169.254.0.1 is assigned to the device with the highest serial number. The IP address 169.254.0.2 is assigned to the device with the second highest serial number, and so on. The IP address assignment does not change when a failover happens. Regardless of the device role at any time (primary or secondary), its heartbeat IP address remains the same.

A change in the heartbeat IP addresses may happen when a FortiGate device joins or leaves the cluster. In those cases, the cluster renegotiates the heartbeat IP address assignment, this time taking into account the serial number of any new device, or removing the serial number of any device that left the cluster.

The HA cluster uses the heartbeat IP addresses to distinguish the cluster members and synchronize data.

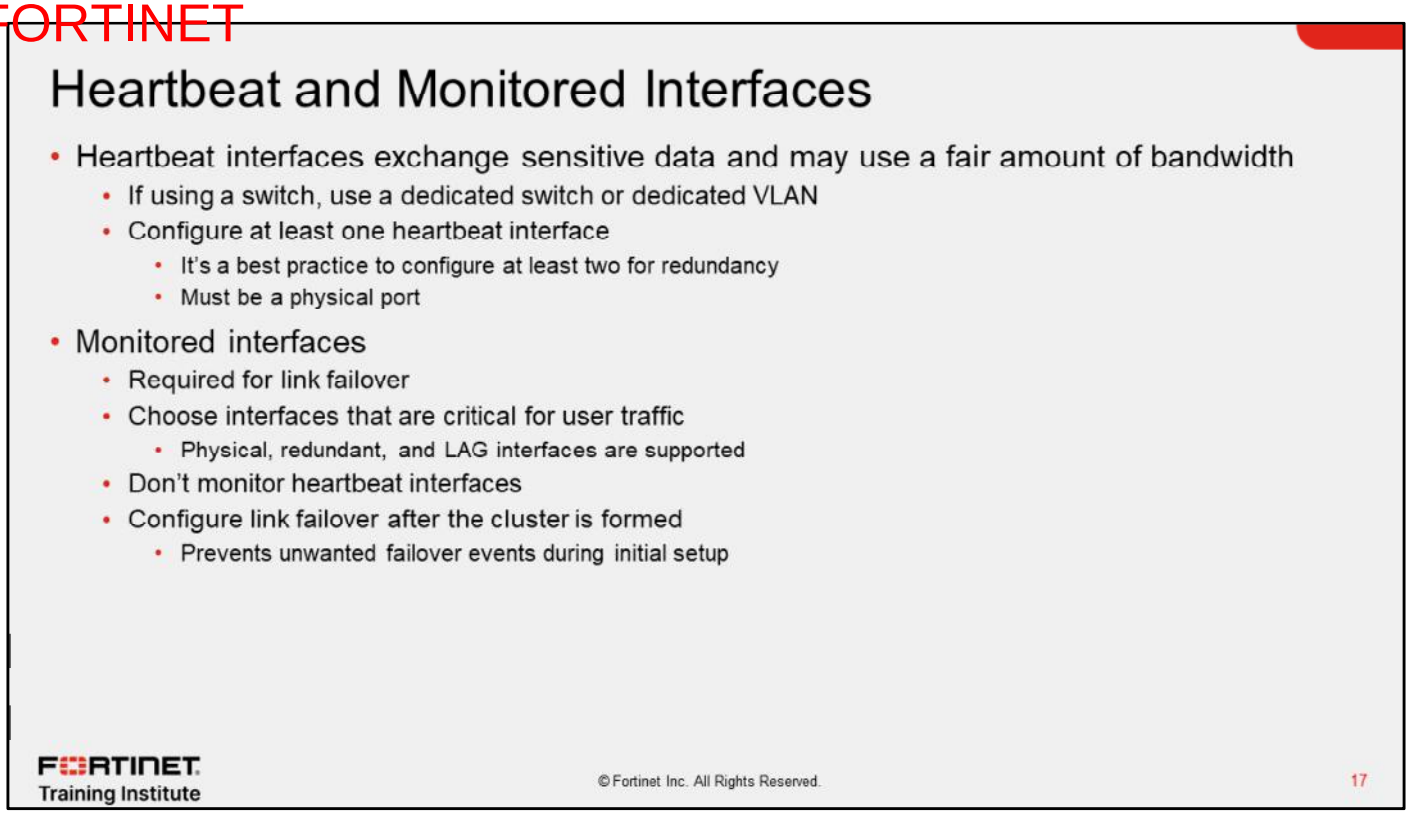

Heartbeat interfaces exchange sensitive information about the cluster operation and may require a fair amount of bandwidth for data synchronization. For this reason, if you use a switch to connect the heartbeat interfaces, it's recommended that you use a dedicated switch or, at least, that you place the heartbeat traffic on a dedicated VLAN.

In addition, you must configure at least one port as a heartbeat interface, but preferably two for redundancy. For heartbeat interfaces, you can use physical interfaces only. That is, you can't use VLAN, IPsec VPN, redundant, or 802.3ad aggregate interfaces. You cannot use FortiGate switch ports either.

For link failover to work, you must configure one or more monitored interfaces. A monitored interface should be an interface whose failure has a critical impact in the network, and therefore, should trigger a device failover. For example, your LAN or WAN interfaces are usually good choices for monitored interfaces. Heartbeat interfaces, however, should not be configured as monitored interfaces because they are not meant to handle user traffic. Note that you can monitor physical ports, redundant interfaces, and link aggregation group (LAG) interfaces.

As a best practice, wait until a cluster is up and running and all interfaces are connected before configuring link failover. This is because a monitored interface can be disconnected during the initial setup and, as a result, trigger a failover before the cluster is fully configured and tested.

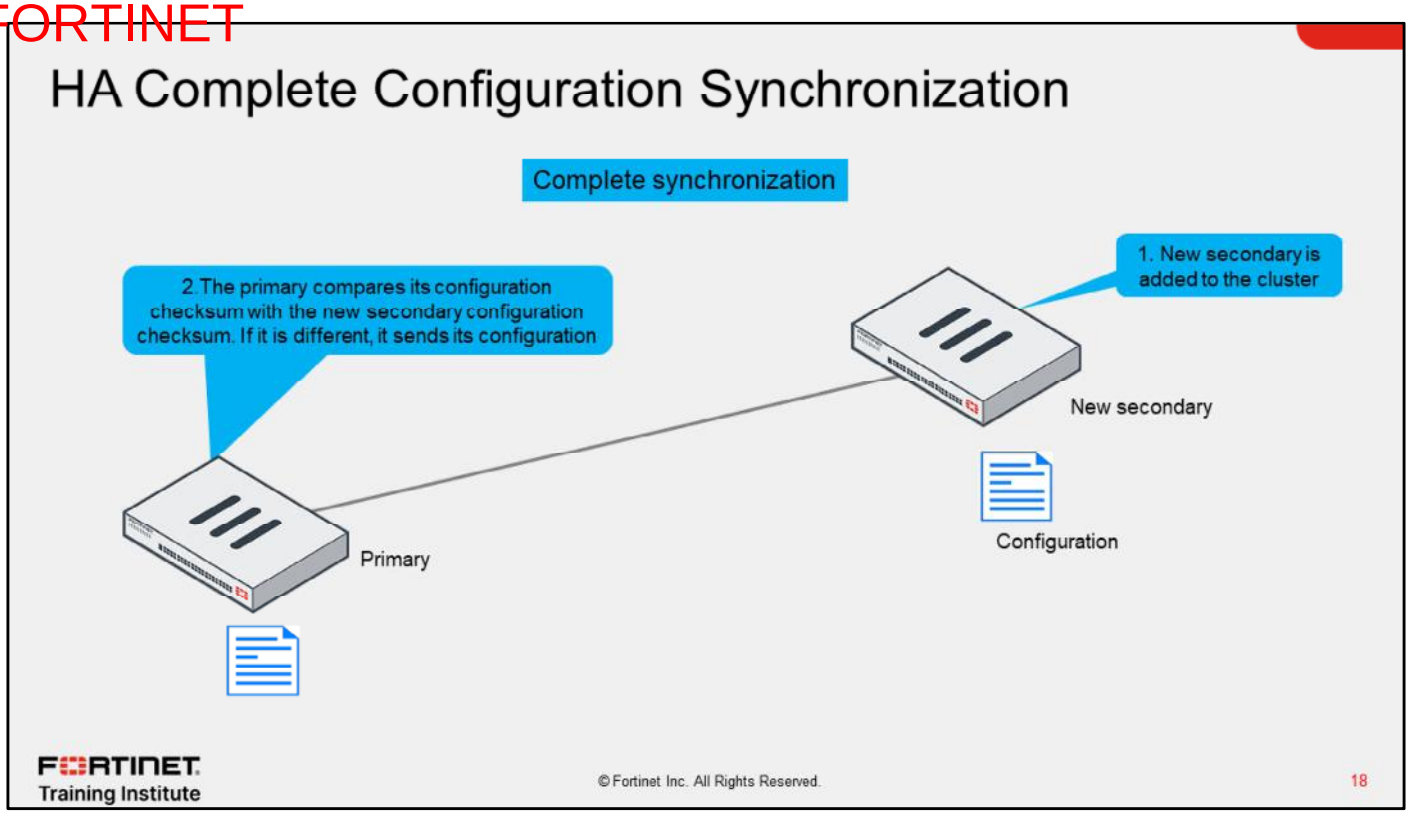

To prepare for a failover, an HA cluster keeps its configurations in sync. You will explore that now.

FortiGate HA uses a combination of both incremental and complete synchronizations.

When you add a new FortiGate to the cluster, the primary FortiGate compares its configuration checksum with the new secondary FortiGate configuration checksum. If the checksums don't match, the primary FortiGate uploads its complete configuration to the secondary FortiGate.

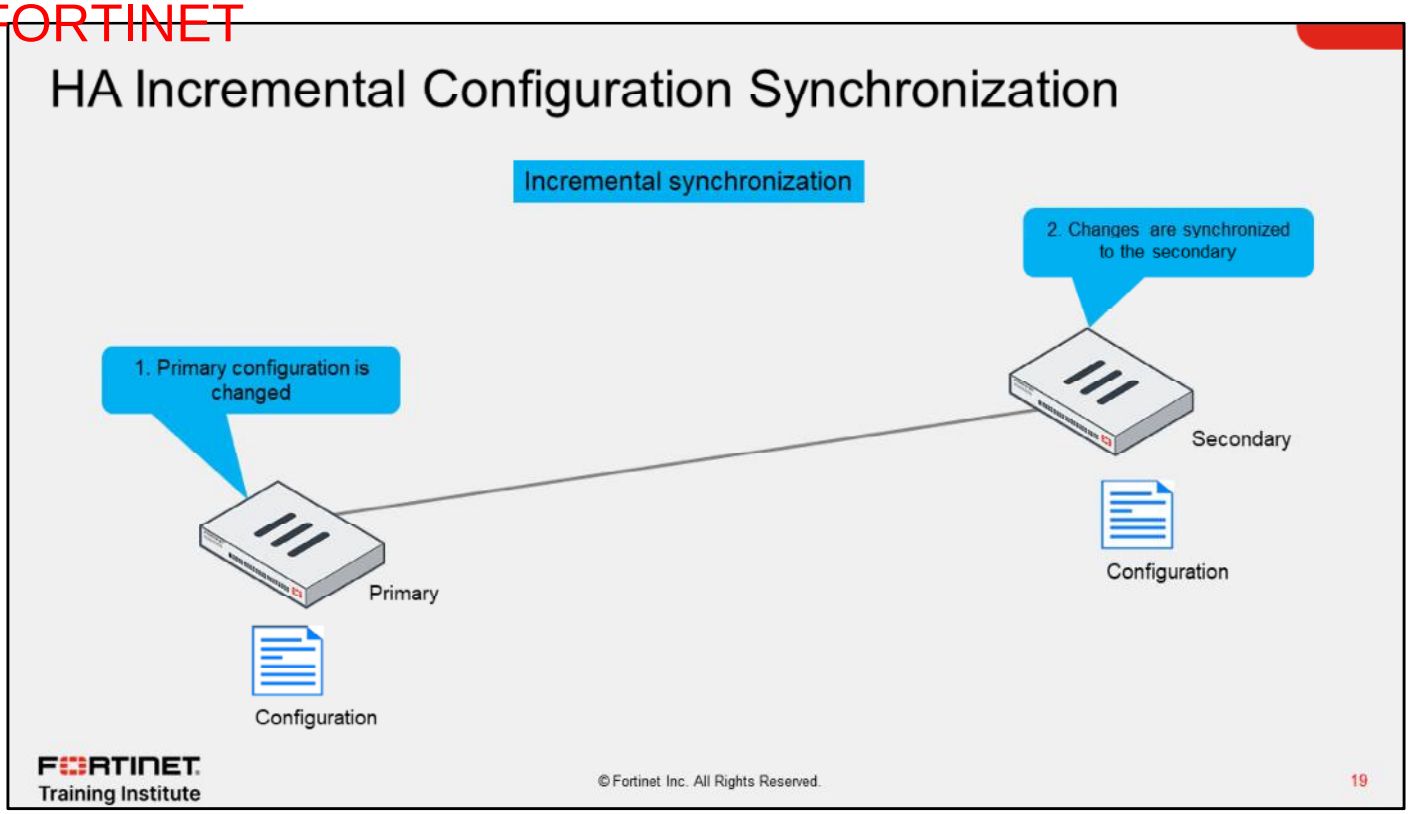

After the initial synchronization is complete, the primary sends any further configuration changes made by an administrator to all the secondary devices. For example, if you create a firewall address object, the primary doesn't resend its complete configuration—it sends only the new object.

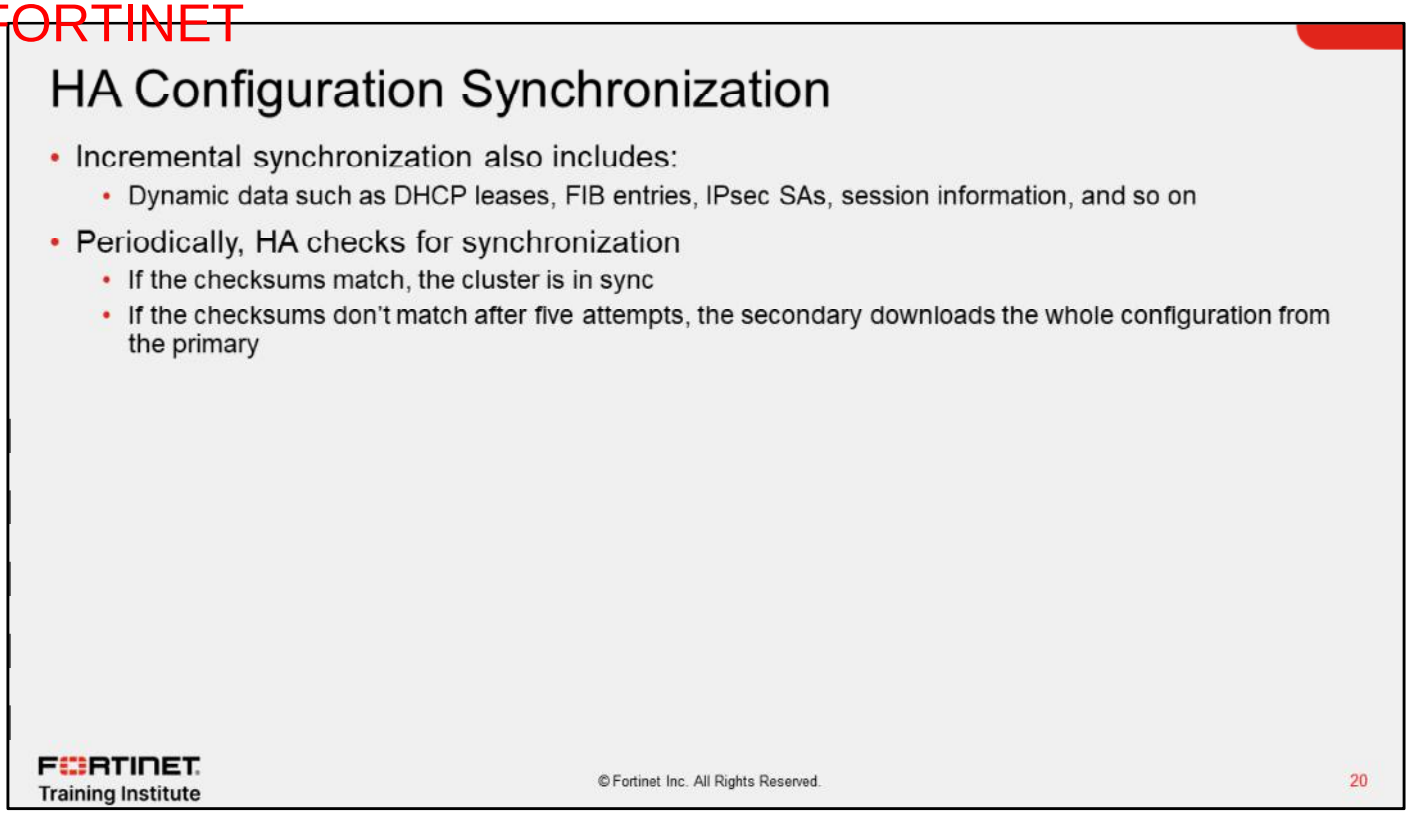

HA propagates more than just configuration details. Some runtime data, such as DHCP leases and FIB entries, are also synchronized.

By default, the cluster checks every 60 seconds to ensure that all devices are synchronized. If a secondary device is out of sync, its checksum is checked every 15 seconds. If the checksum of the out-of-sync secondary device doesn't match for five consecutive checks, a complete resynchronization to that secondary device is done.

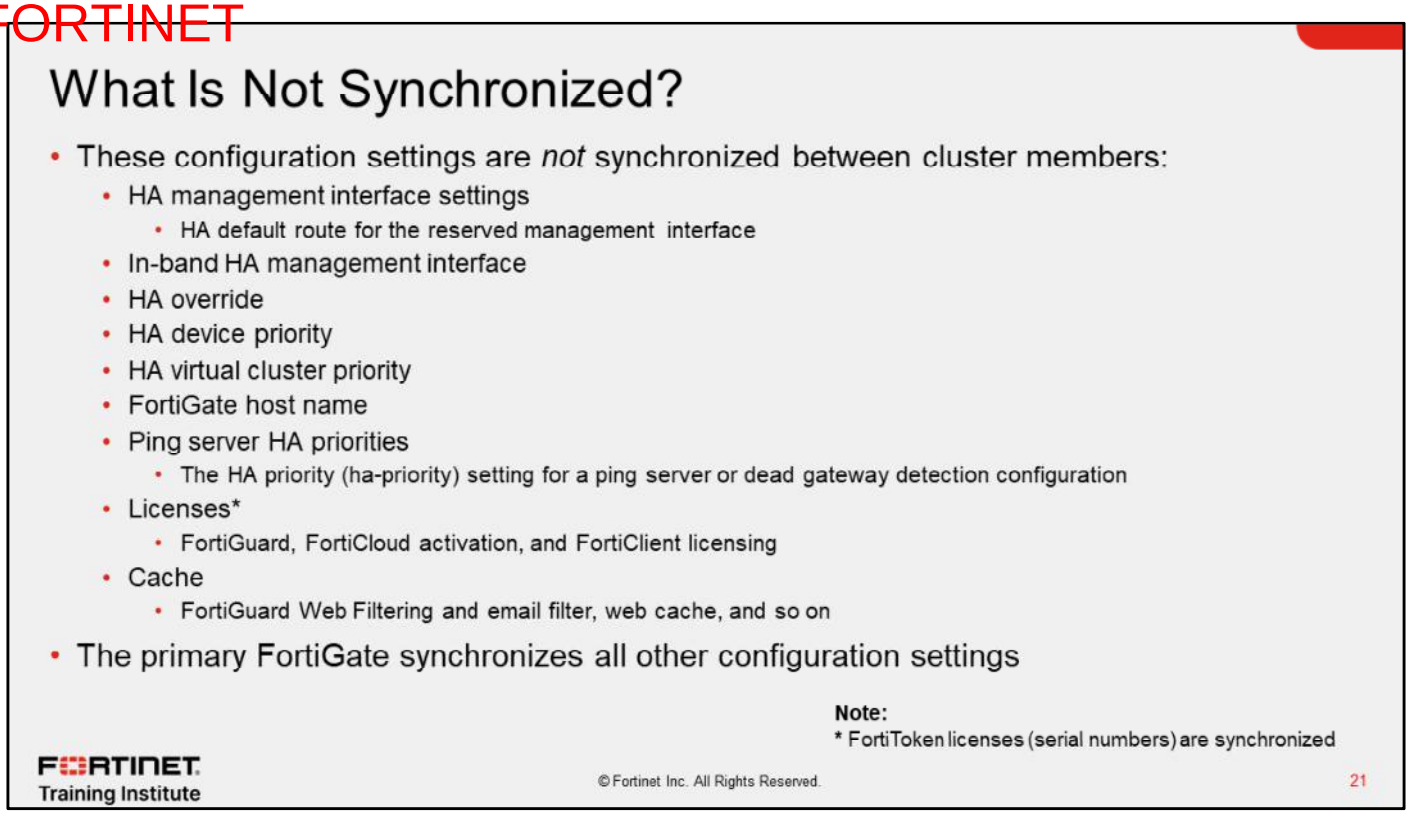

Not all the configuration settings are synchronized. There are a few that are not, such as:

- System interface settings of the HA reserved management interface and the HA default route for the reserved management interface
- In-band HA management interface
- HA override
- HA device priority
- Virtual cluster priority
- FortiGate host name
- HA priority setting for a ping server (or dead gateway detection) configuration
- All licenses except FortiToken licenses (serial numbers)
- Cache

The primary FortiGate synchronizes all other configuration settings, including all other HA settings.

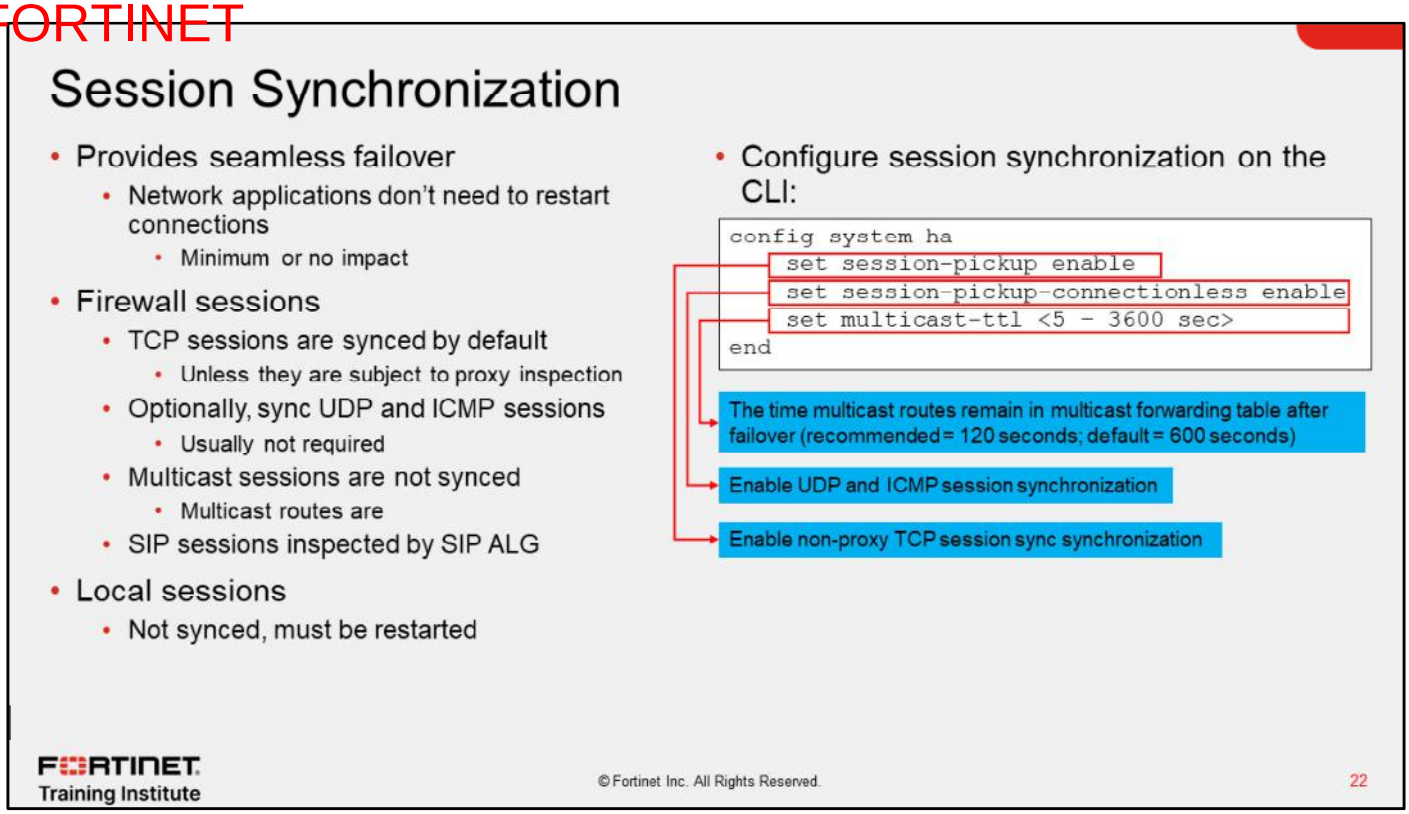

Session synchronization provides seamless session failover. When the primary fails, the new primary can resume traffic for synchronized sessions without network applications having to restart the connections.

By default, the feature synchronizes TCP firewall sessions that are not subject to proxy-based inspection. An exception to this rule is TCP SIP sessions inspected by SIP ALG. Even though SIP ALG performs proxy-based inspection on SIP sessions, FortiGate can still synchronize such SIP sessions. Firewall sessions, also known as pass-through sessions, are user traffic sessions that travel across FortiGate. TCP firewall sessions that are subject to flow-based inspection or no inspection at all, are synchronized to secondary members.

You can also enable the synchronization of UDP and ICMP sessions. Although both protocols are connectionless protocols, FortiGate still allocates sessions for UDP and ICMP connections in its session table. Usually, the synchronization of UDP and ICMP sessions is not required because most UDP and ICMP connections can resume communication if their session information is lost.

For multicast traffic, FortiGate synchronizes multicast routes only. That is, FortiGate doesn't synchronize multicast sessions, which should be fine because multicast sessions are mostly UDP-based and, as mentioned before, UDP sessions can usually resume communication if their session information is lost. To ensure the multicast routing information across members is accurate, you can adjust the multicast time to live (TTL) timer. The timer controls how long the new primary keeps the synced multicast routes in the multicast forwarding table. The smaller the timer value, the more often the routes are refreshed, and so the more accurate the multicast forwarding table is. The recommended timer value is 120 seconds.

Local-in and local-out sessions, which are sessions that are terminated at or initiated by FortiGate, respectively, are not synchronized either. For example, BGP peerings, OSPF adjacencies, as well as SSH and HTTPS management connections must be restarted after a failover.

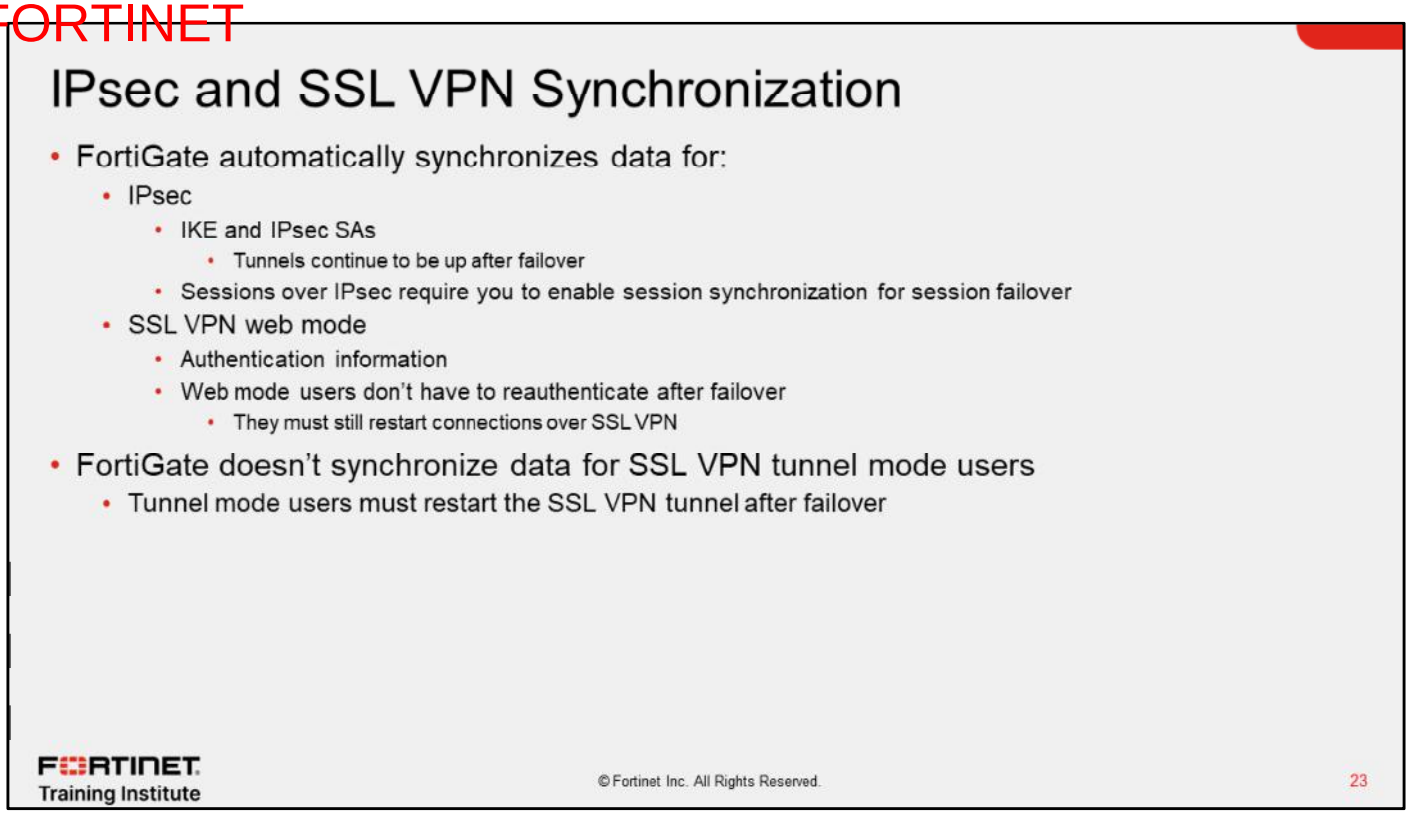

The primary FortiGate automatically synchronizes all IKE and IPsec security associations (SAs) to secondary members. This enables the new primary to resume existing IPsec tunnels after a failover. Note that you must also enable session synchronization if you want the new primary to also resume existing IPsec sessions. Otherwise, after a failover, you must still restart existing TCP connections made over IPsec tunnels, even though the IPsec tunnels continue to be up on the new primary.

For SSL VPN, the primary FortiGate synchronizes the authentication information for SSL VPN web mode users only. That is, they are using SSL VPN web mode, the SSL VPN users don't have to authenticate again after a failover. However, the users must still restart the connections made using SSL VPN web mode to regain access to protected resources. Note that FortiGate doesn't synchronize any information for SSL VPN tunnel mode. That is, after a failover, SSL VPN tunnel mode users must restart their SSL VPN tunnel connection, as well as any connection made through the tunnel.

Knowledge Check

1. Which information is synchronized in an HA cluster?
   - ✓ A. Firewall policies and objects
   - B. FortiGate host name

2. Which one of the following session types can be synchronized in an HA cluster?
   - A. BGP peerings
   - ✓ B. Non-proxy TCP sessions

FORTINET Training Institute

24

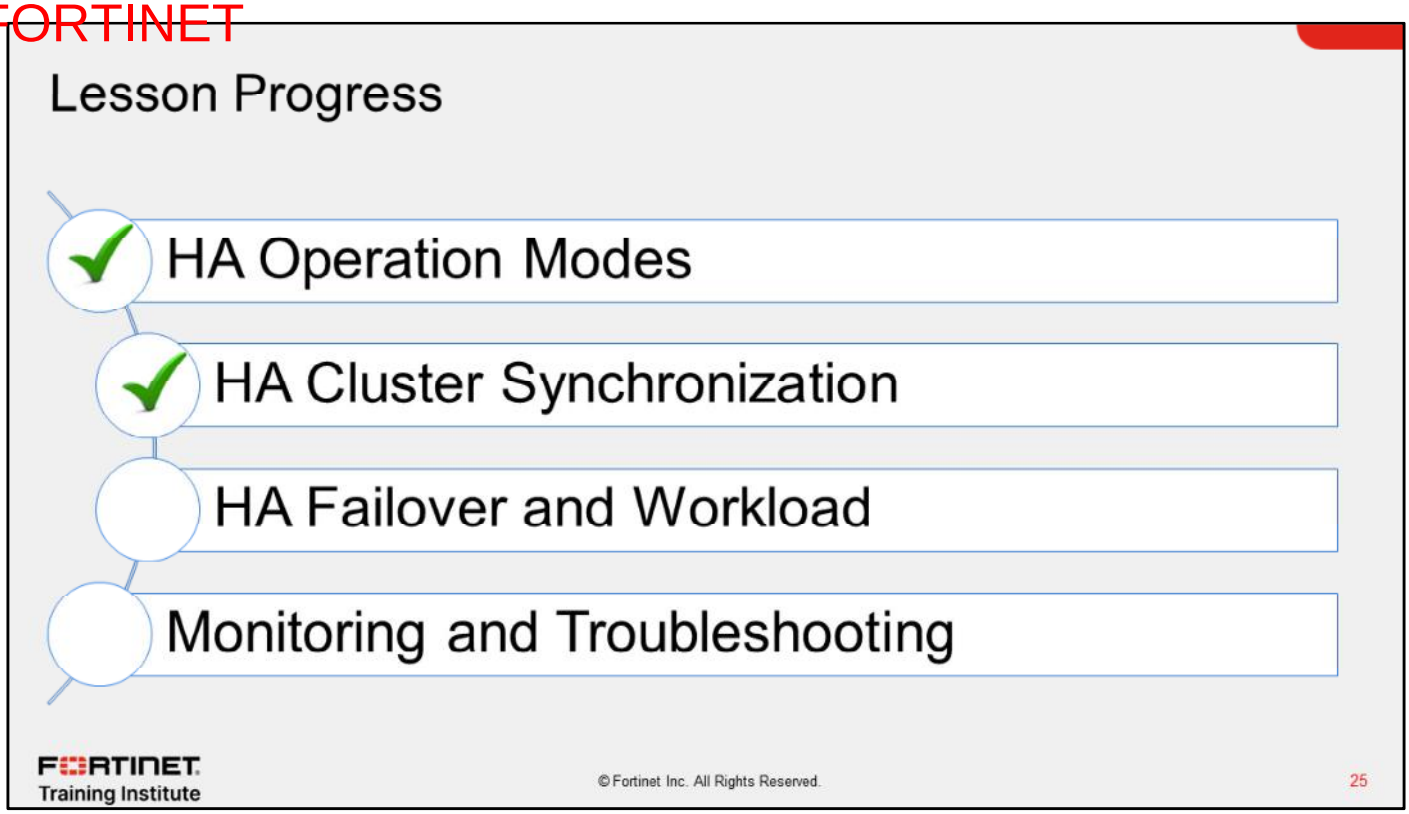

Good job! You now understand HA cluster synchronization.

Now, you will learn about HA cluster failover protection types and workload for primary and secondary FortiGate devices in an HA cluster.

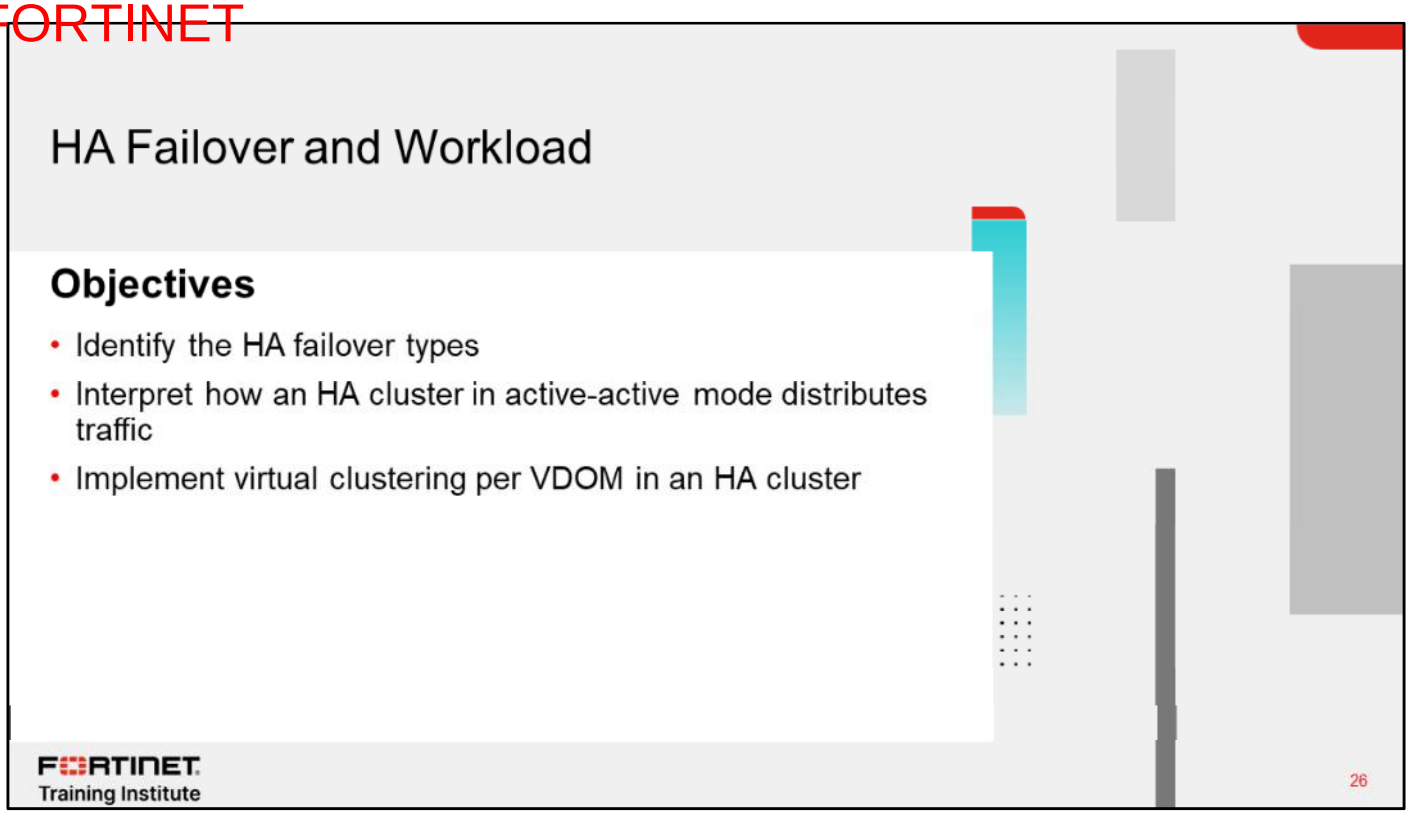

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in failover types and workload, you will be able to identify how enhanced reliability is achieved through HA failover protection. You will also learn about the distribution of traffic in an active-active cluster and distributing traffic using virtual clustering.

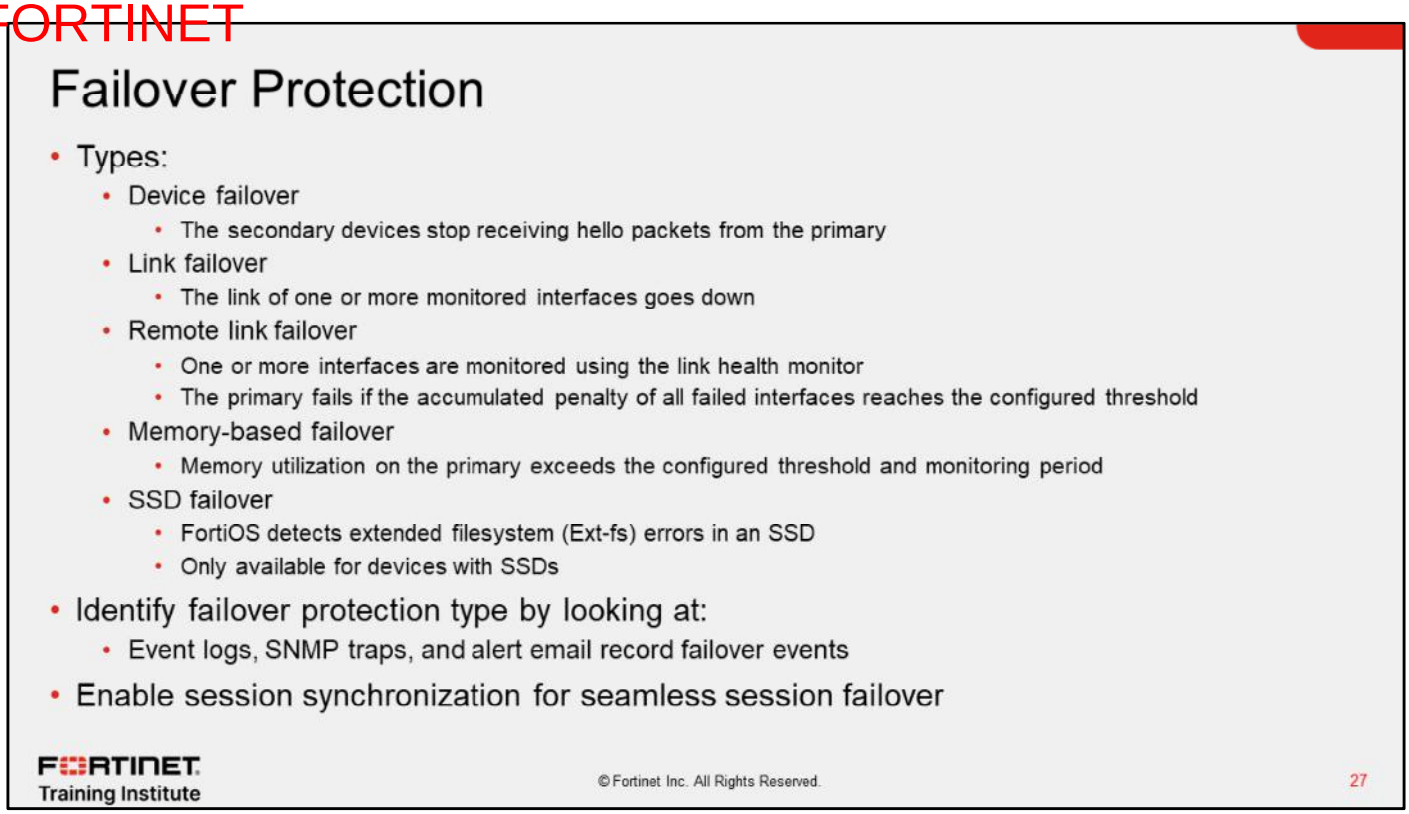

The most common types of failovers are device failovers and link failovers. However, you can also configure remote link failover and memory-based failover. When a failover event is triggered, the secondary devices elect a new primary.

A device failover is triggered when the secondary devices stop receiving the heartbeat hello packets from the primary.

A link failover occurs when the link status of a monitored interface on the primary FortiGate goes down. You can configure an HA cluster to monitor one or more interfaces. If a monitored interface on the primary FortiGate is unplugged, or its link status goes down, a new primary FortiGate is elected.

When you configure remote link failover, FortiGate uses the link health monitor feature to monitor the health of one or more interfaces against one or more servers that act as beacons. The primary FortiGate fails if the accumulated penalty of all failed interfaces reaches the configured threshold.

If you enable memory-based failover, an HA failover is triggered when the memory utilization on the primary FortiGate reaches the configured threshold for the configured monitoring period. You can also enable SSD failover, which triggers a failover if FortiOS detects Ext-fs errors on an SSD on the primary FortiGate.

There are multiple events that might trigger an HA failover, such as a hardware or software failure on the primary FortiGate or an issue on one of the interfaces on the primary. When a failover occurs, an event log is generated. Optionally, you can configure the device to also generate SNMP traps and alert emails.

Make sure that you enable session pickup for sessions you want to protect from a failover event. This way, the new primary can resume traffic for these sessions.

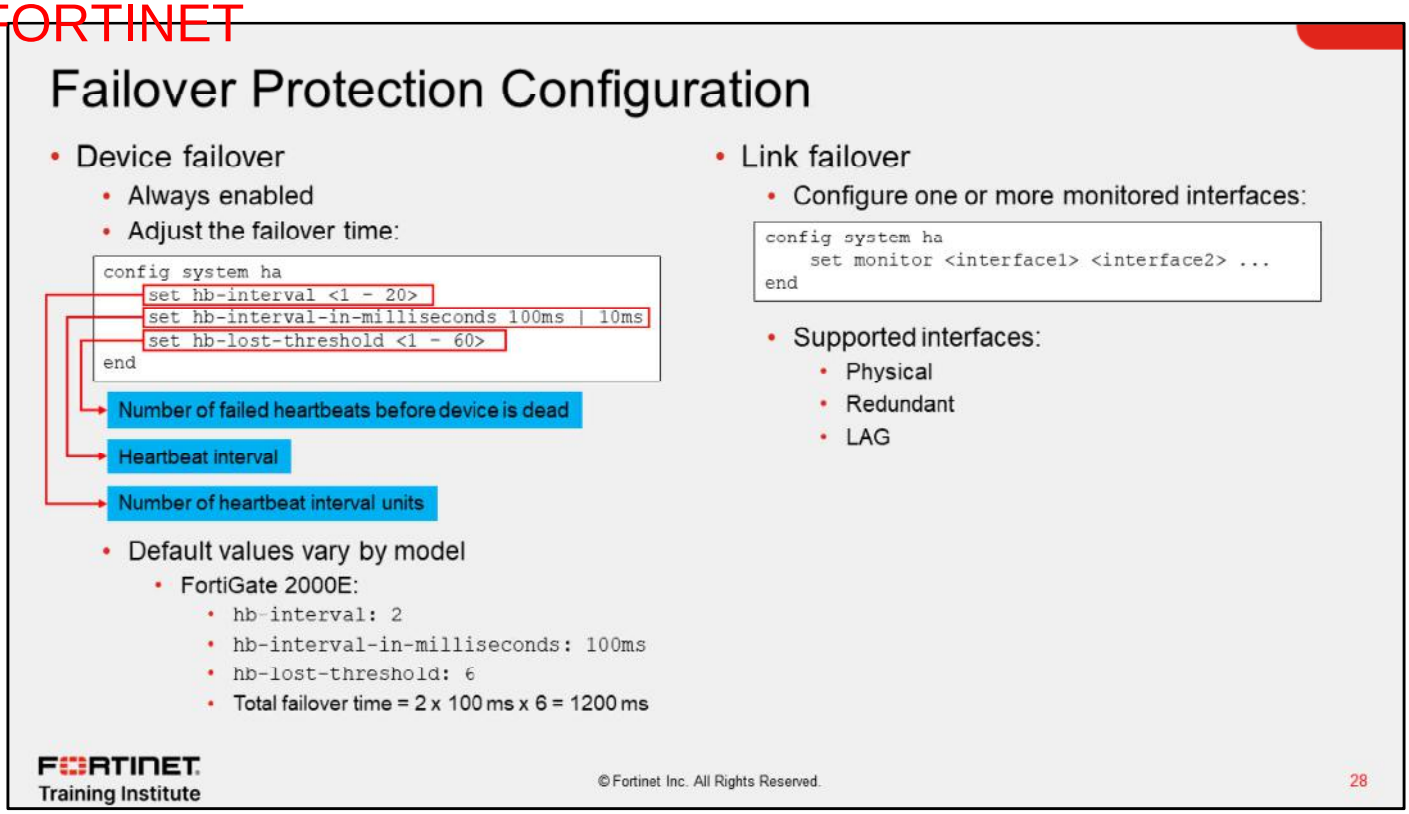

When you configure HA, device failover is always enabled. However, you can adjust the settings that dictate the failover time. To speed up failover, you can reduce the values for all three settings shown on this slide. To reduce false positives, increase their values.

The default values for the three settings vary by model. For example, using the default values on a FortiGate 2000E model results in a device failover time of 1200 milliseconds (1.2 seconds). Note that the 10-millisecond heartbeat interval is supported on NP6 platforms only.

To configure link failover, you must configure one or more monitored interfaces, as shown on this slide. Note that you can configure only physical, redundant, and LAG interfaces as monitored interfaces.

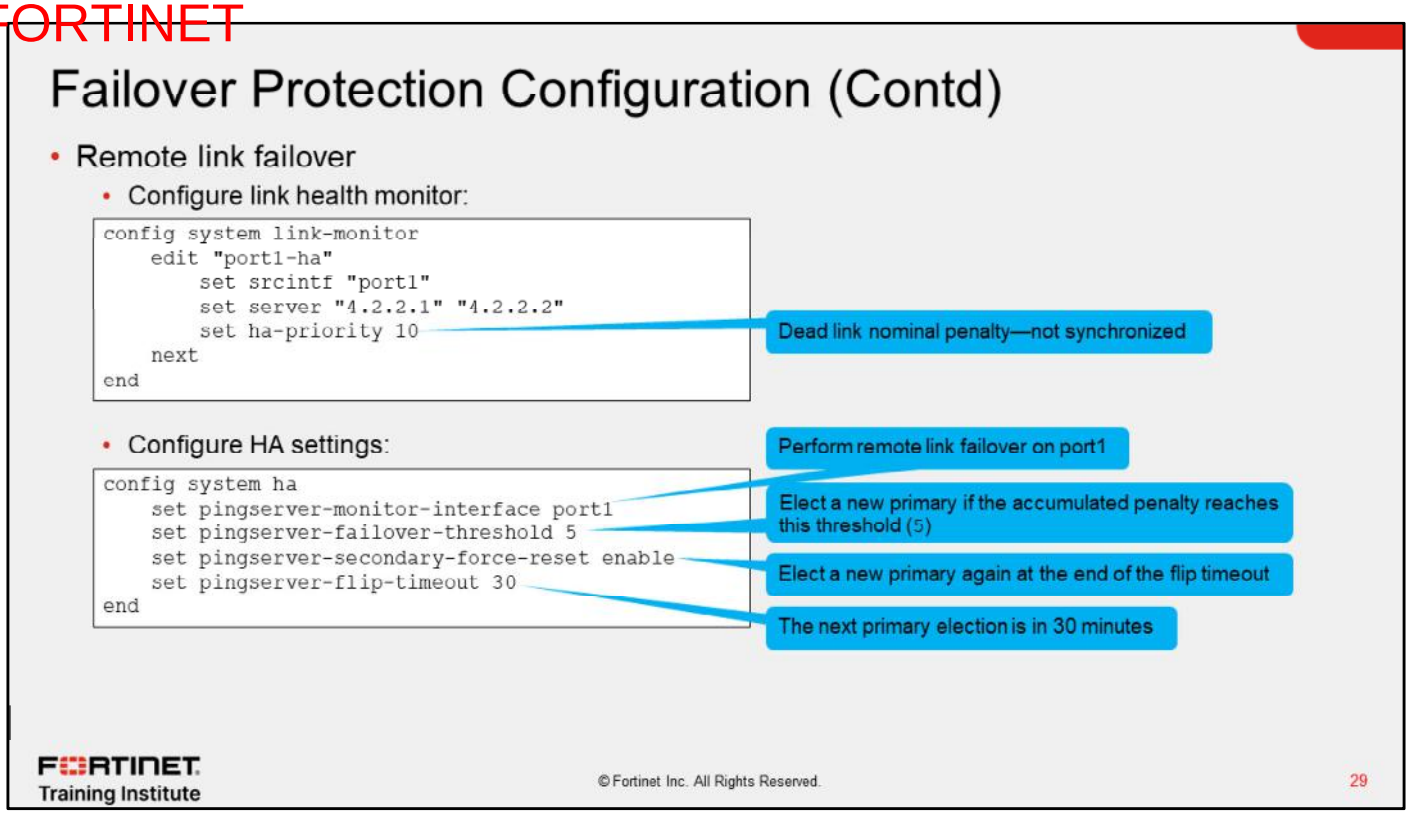

This slide shows a configuration example for remote link failover.

First, you configure link health monitor, as shown in the *Routing* lesson. The `ha-priority` setting in the link health monitor configuration defines the penalty applied to the member after the link is detected as dead. Note that the `ha-priority` setting has local significance only, and therefore, is not synchronized to other members.

The next step is to configure the HA settings related to remote link failover. The configuration on this slide instructs FortiGate to perform remote link failover on port1 as follows:

- When port1 is detected as dead, the nominal penalty (10) is added to the global penalty, which is initially set to 0.
- If the accumulated penalty reaches the penalty threshold (5), then the cluster elects a new primary. A failover occurs when a secondary member has a lower accumulated penalty than the primary. If so, the secondary member with the lowest accumulated penalty becomes the new primary.
- The cluster doesn't elect a new primary again until the flip timeout (30 minutes) has passed.

If during the primary election, the accumulated penalty of all members is the same, then other criteria, such as monitored interfaces, priority, uptime, and so on, are used as tiebreakers to elect the new primary.

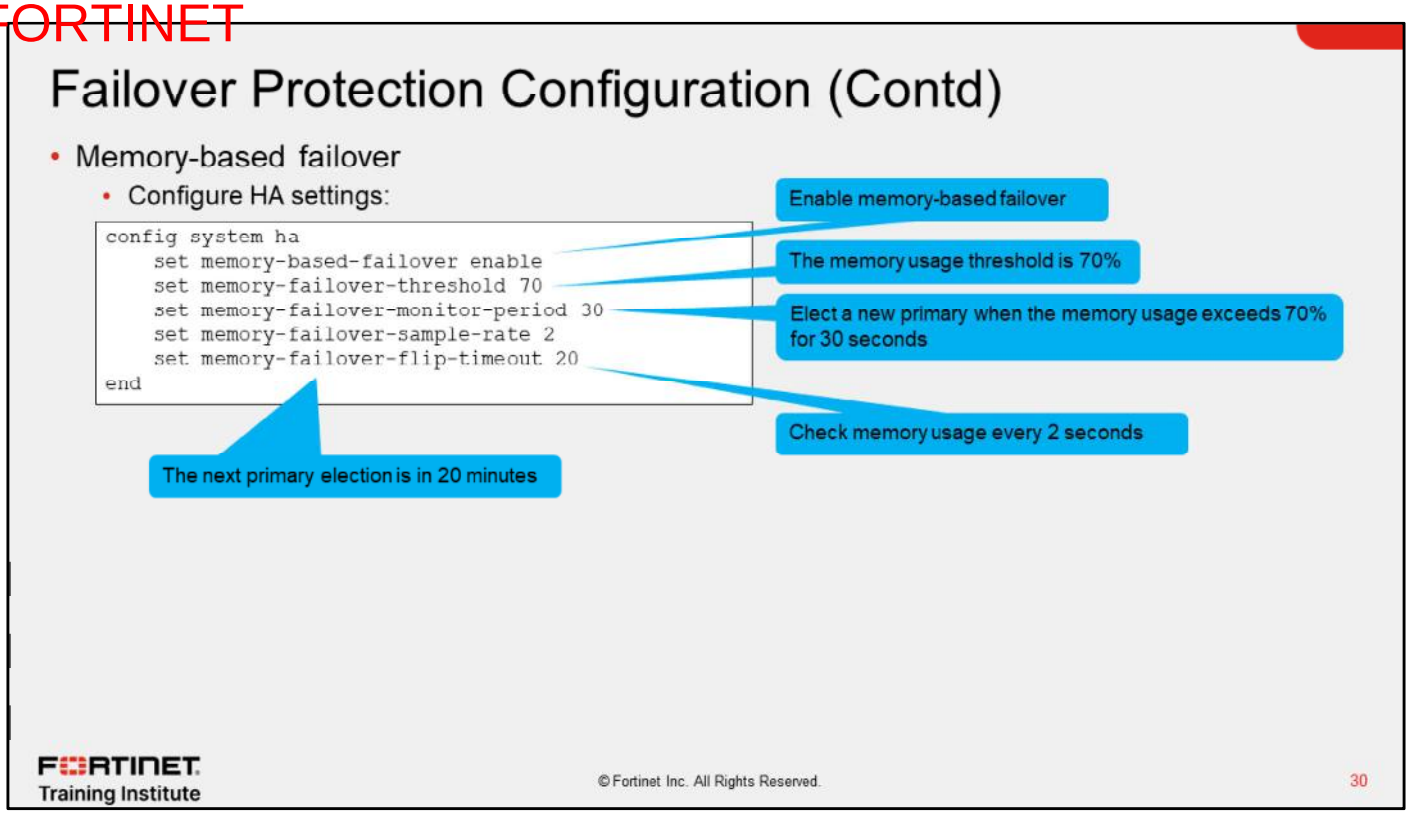

The HA configuration shown on this slide instructs FortiGate to perform memory-based failover as follows:

- When the memory on the primary reaches the threshold (70%) and stays like that for 30 seconds, then the cluster elects a new primary.
- During primary election, a failover occurs when the memory usage on a secondary member is lower than the configured memory threshold (70%). If so, the secondary member becomes the new primary.
- The cluster doesn't elect a new primary again until the flip timeout (20 minutes) has passed.
- Each member in the cluster checks its memory usage every 2 seconds.

If during the primary election, the memory usage of all members are below or above the threshold, then other criteria, such as monitored interfaces, priority, uptime, and so on, are used as tiebreakers to elect the new primary.

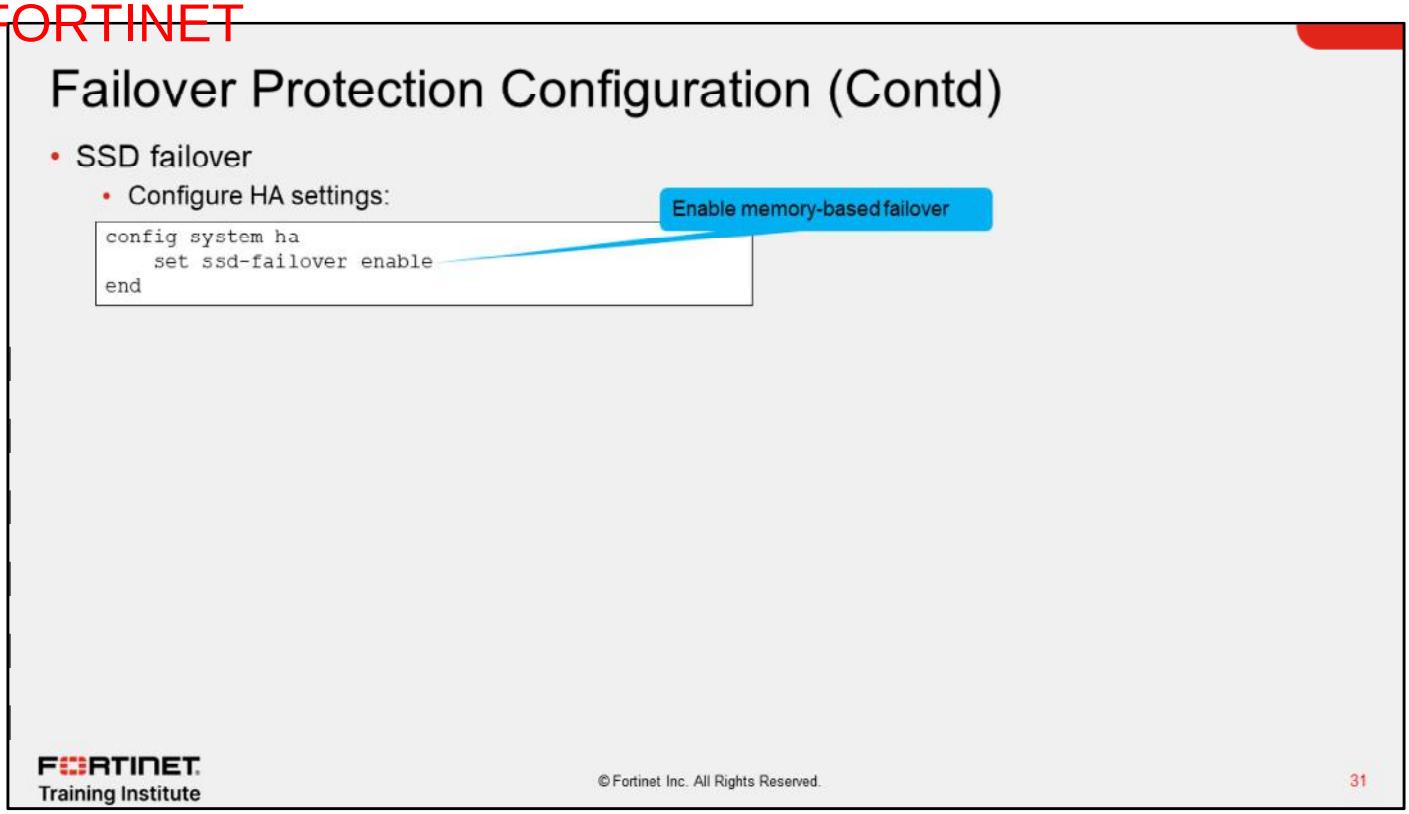

The HA configuration shown on this slide instructs FortiGate to perform a failover when any of the SSD disks on the primary FortiGate report Ext-fs errors. Note that this feature is supported only on FortiGate models with SSD disks.

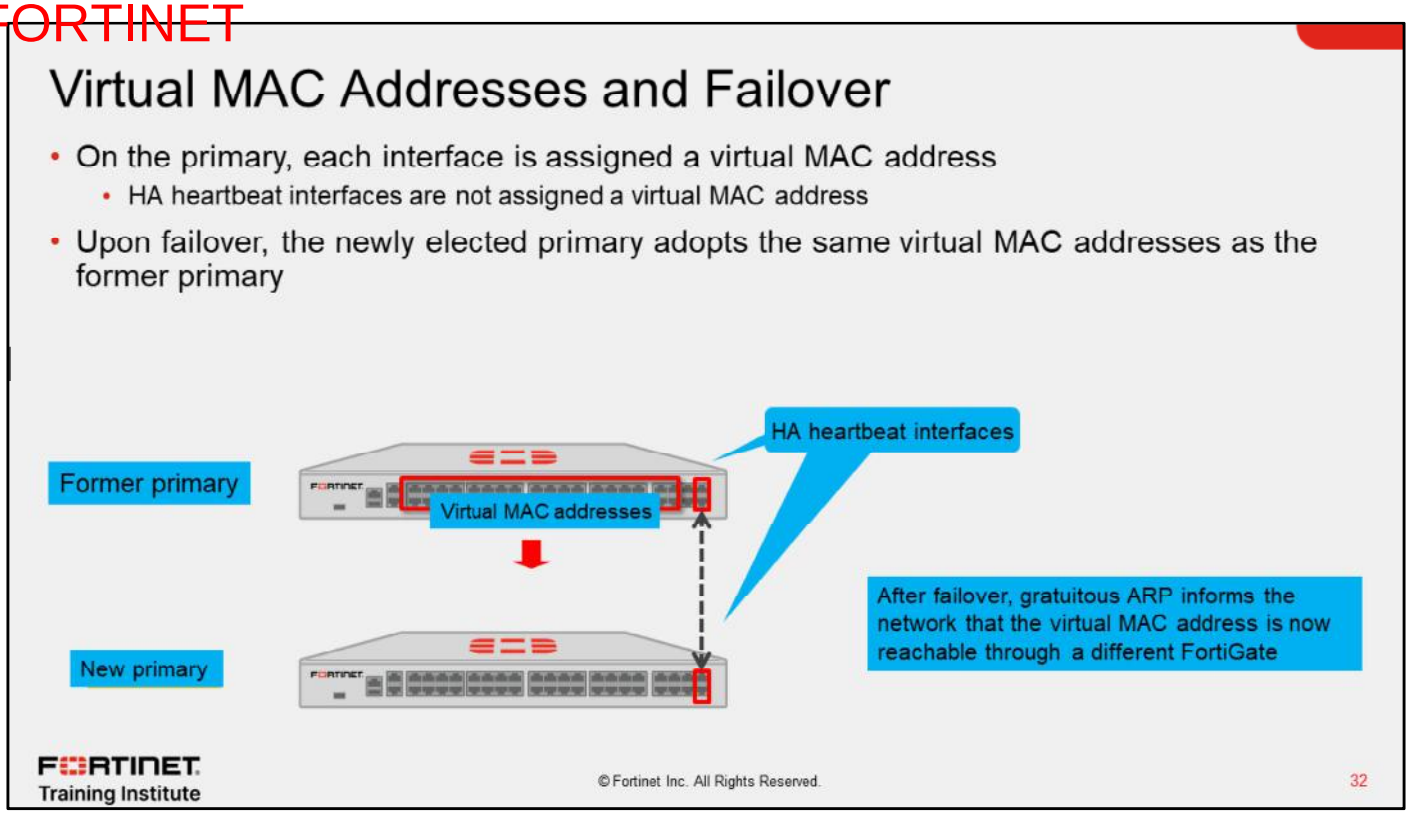

To forward traffic correctly, a FortiGate HA solution uses virtual MAC addresses. When a primary joins an HA cluster, each interface is assigned a virtual MAC address. The HA group ID is used in the creation of virtual MAC addresses assigned to each interface. So, if you have two or more HA clusters in the same broadcast domain, and using the same HA group ID, you might get MAC address conflicts. For those cases, it is strongly recommended that you assign different HA group IDs to each cluster.

Through the heartbeats, the primary informs all secondary devices about the assigned virtual MAC address. Upon failover, a secondary adopts the same virtual MAC addresses for the equivalent interfaces.

The new primary broadcasts gratuitous ARP packets, notifying the network that each virtual MAC address is now reachable through a different switch port.

Note that the MAC address of a reserved HA management interface is not changed to a virtual MAC address. Instead, the reserved management interface keeps its original MAC address.

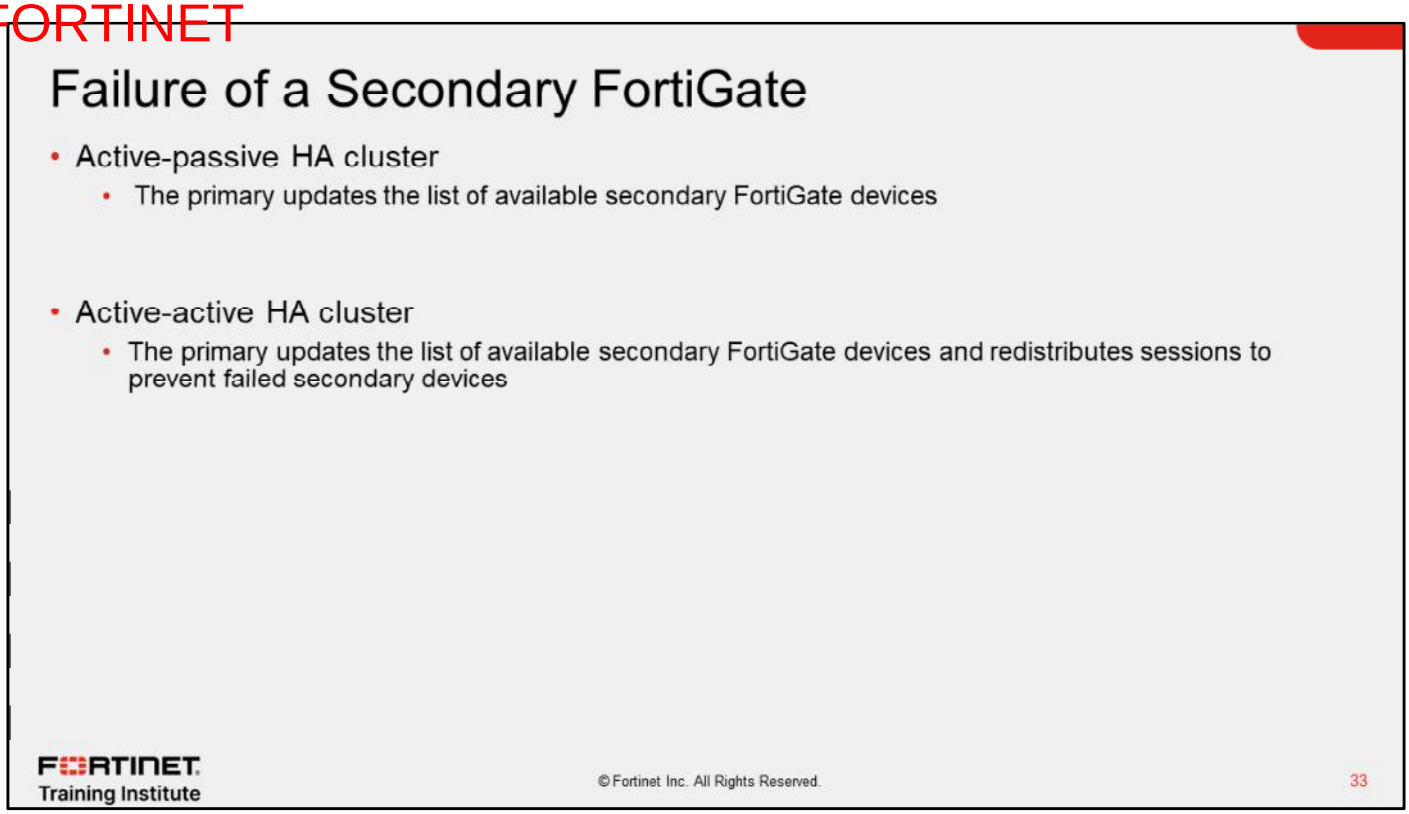

As you learned earlier in this lesson, if a primary fails, a new primary is elected. But what happens if a secondary FortiGate device fails? It depends on the HA mode.

In an active-passive cluster, the primary only updates its list of available secondary FortiGate devices. It also starts monitoring for the failed secondary, waiting for it to come online again.

However, in an active-active cluster, the secondary devices can handle traffic. So, the primary (which tracks and assigns sessions to each secondary) must not only update its list of available secondary FortiGate devices, but also reassign sessions from the failed FortiGate to a different secondary FortiGate.

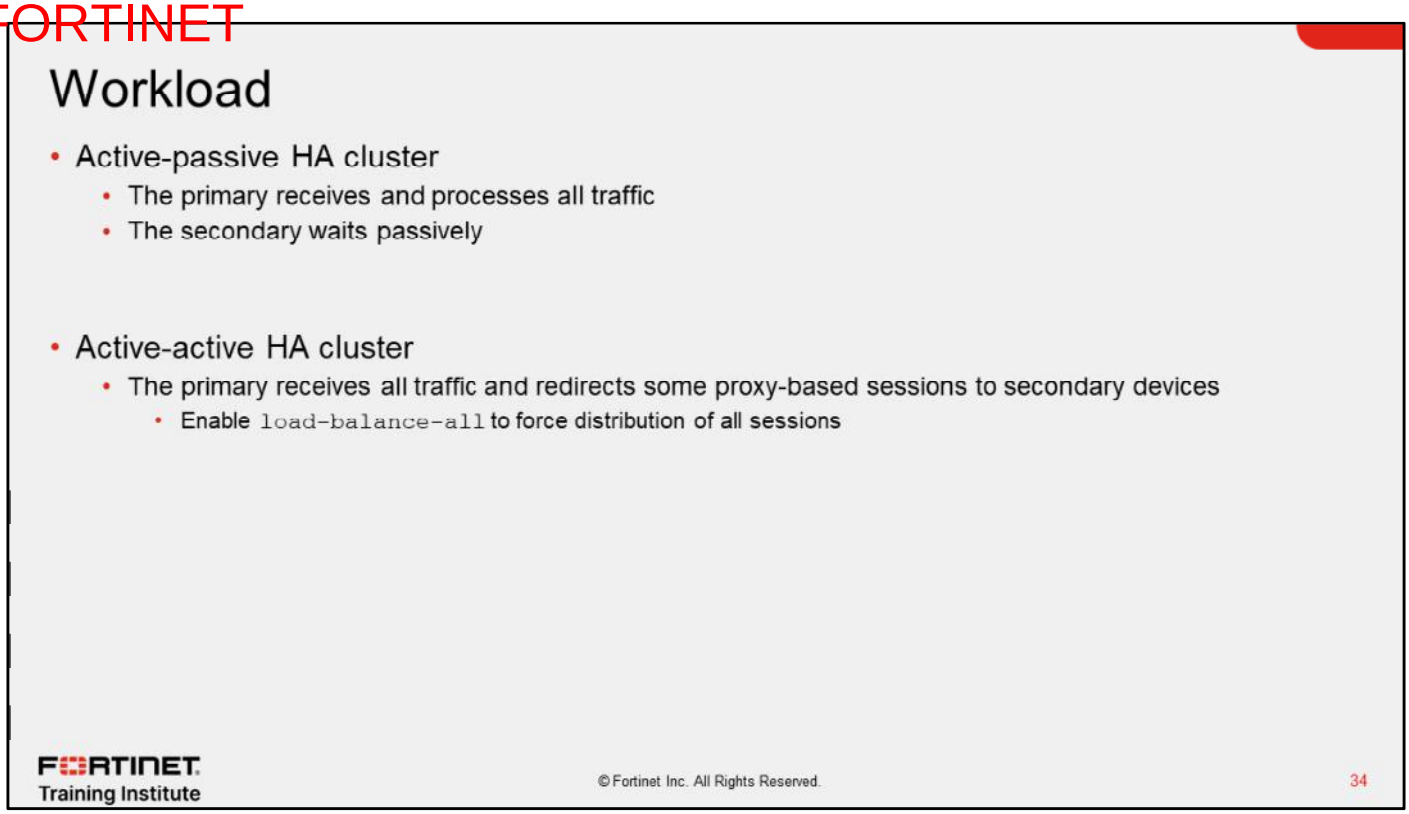

This is how the workload is distributed between roles, depending on the HA mode.

Notice that traffic workload is distributed in active-active mode only. However, keep in mind that by default, only sessions that are subject to proxy inspection are distributed to secondary devices. If you want to force the distribution of sessions that are subject to flow inspection or no inspection at all, then you must enable the `load-balance-all` setting under HA configuration—this setting is disabled by default.

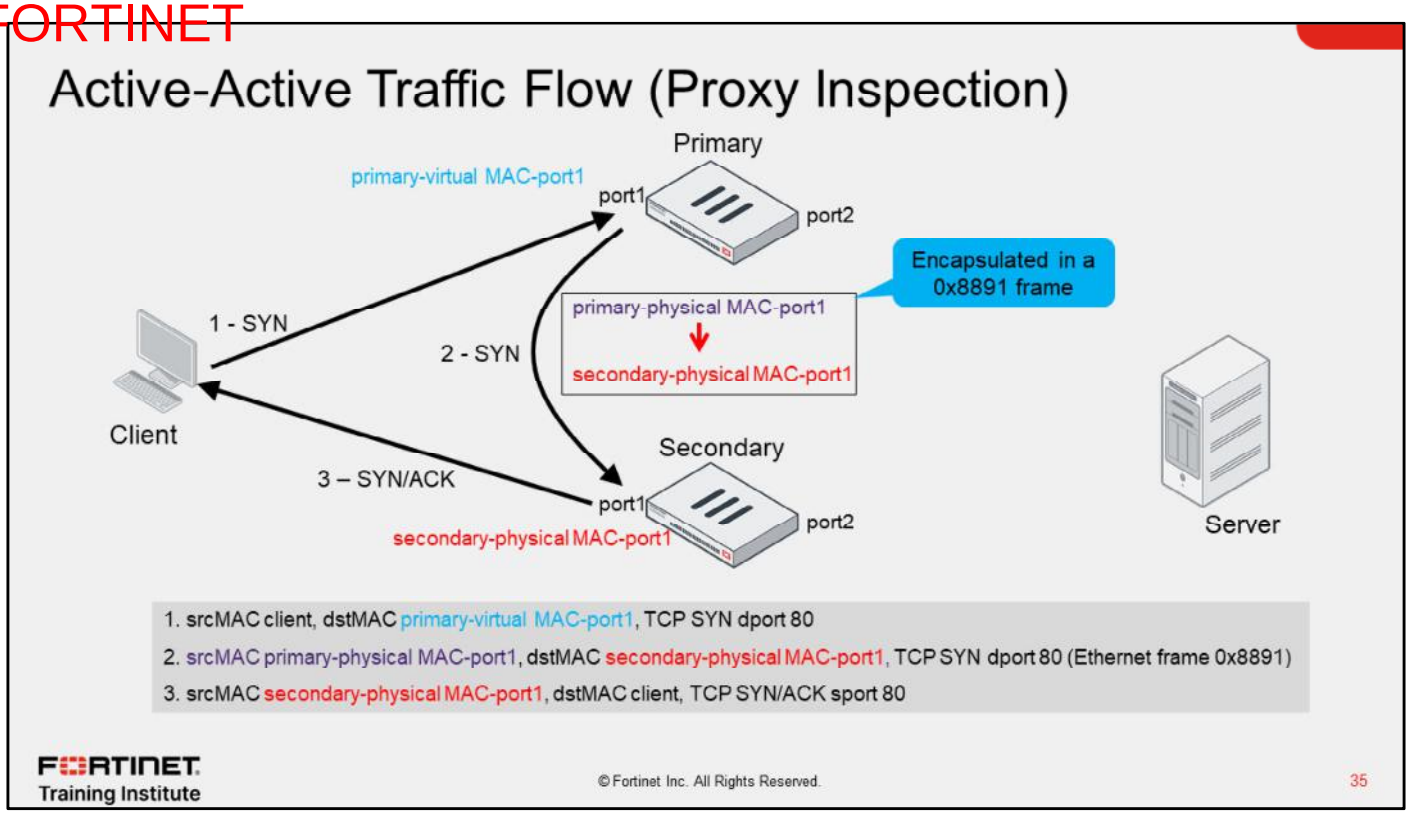

In active-active mode, the following occurs:

- The traffic destined to the cluster is sent to the primary. Because all network ports on the primary—except the heartbeat ports—are assigned a virtual MAC address, the traffic is destined to the virtual MAC address of the receiving port on the primary FortiGate.

- For traffic that is distributed to the secondary, the traffic destined to the endpoints is sent by the secondary. The traffic is sourced from the physical MAC address of the egressing port on the secondary.

This slide shows the flow for distributed traffic that is subject to proxy inspection:

1. The client sends a SYN packet, which is forwarded to port1 on the primary. The packet destination MAC address is the virtual MAC address on port1.

2. The primary forwards the SYN packet to the selected secondary. In this example, the source MAC address of the packet is changed to the physical MAC address of port1 on the primary and the destination MAC address to the physical MAC address of port1 on the secondary. This is also known as MAC address rewrite. In addition, the primary encapsulates the packet in an Ethernet frame type 0x8891. The encapsulation is done only for the first packet of a load balanced session. The encapsulated packet includes the original packet plus session information that the secondary requires to process the traffic.

3. The secondary responds to the client with a SYN/ACK packet that contains the physical MAC address of port1 on the secondary as the source and the MAC address of the client as the destination.

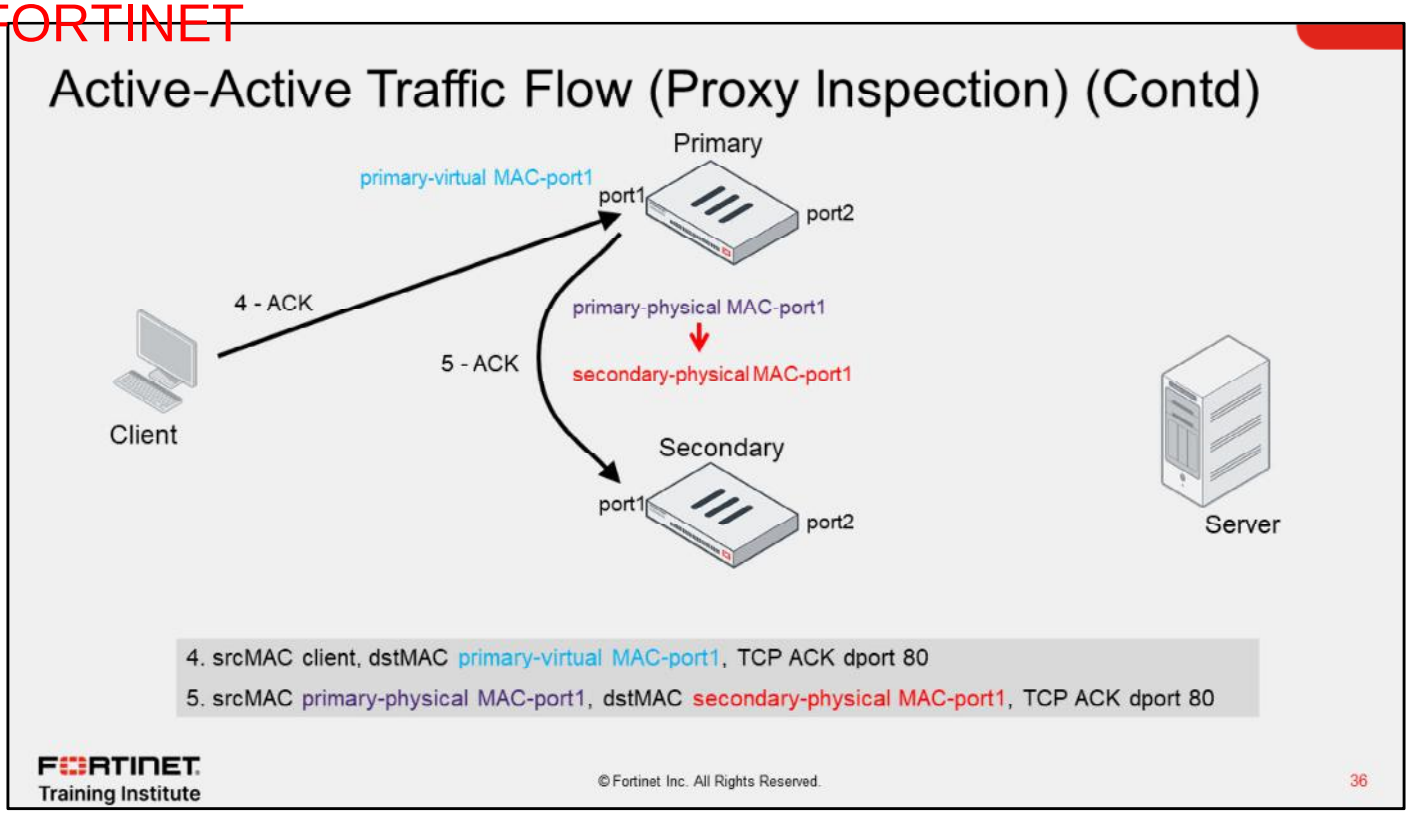

4. The client acknowledges the SYN/ACK by sending an ACK to the cluster. The ACK packet is destined to port1 on the primary.

5. The primary receives the packet and knows that it matches a session that was previously distributed to the secondary. As a result, the primary forwards the ACK packet to the corresponding secondary FortiGate. The packet is sourced from the physical MAC address of port1 on the primary and destined to the physical MAC address of port1 on the secondary. The three-way handshake on the client side is complete.

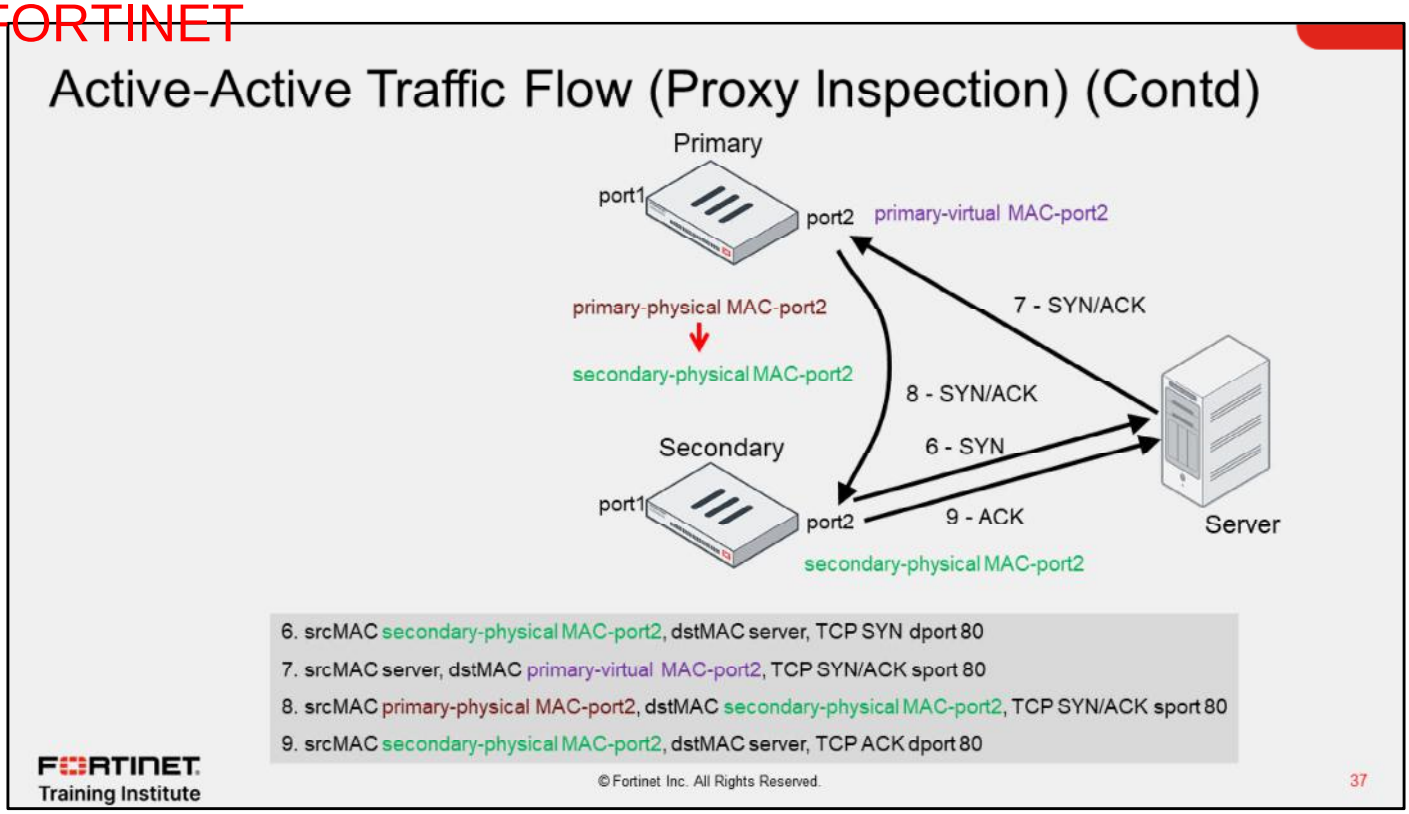

6. The secondary starts the connection with the server by sending a SYN packet using the physical MAC address of port2 as the source. Note that FortiGate contacts the server after it finishes the three-way handshake to the client, not before. The same behavior is seen when FortiGate operates in standalone mode and performs proxy-based inspection.

7. The SYN/ACK packet from the server is sent to port2 on the primary. The destination MAC address is the virtual MAC address of port2.

8. The primary receives the packet and knows that it matches a session that was previously distributed to the secondary. The primary forwards the SYN/ACK packet to the corresponding secondary FortiGate. The packet is sourced from the physical MAC address of port2 on the primary and destined to the physical MAC address of port2 on the secondary.

9. The secondary responds to the server with an ACK packet that contains the physical MAC address of port2 on the secondary as the source and the MAC address of the server as the destination.

The three-way handshake on the server side is also complete. From now on, packets that the client sends follow the same flow. For example, an HTTP GET request packet from the client is first received by the primary, which then forwards it to the secondary for proxy-based inspection. If the packet is allowed, the secondary forwards the packet to the server. Any server response packets to the client HTTP GET request are sent to the primary, which then forwards the packets to the secondary for inspection, and so on.

Note that the goal of active-active mode is to leverage unused CPU and memory resources on secondary devices. The intention is not really to load balance traffic. In fact, because the traffic from endpoints is always sent to the primary, you usually see more traffic on the primary than any secondary devices.

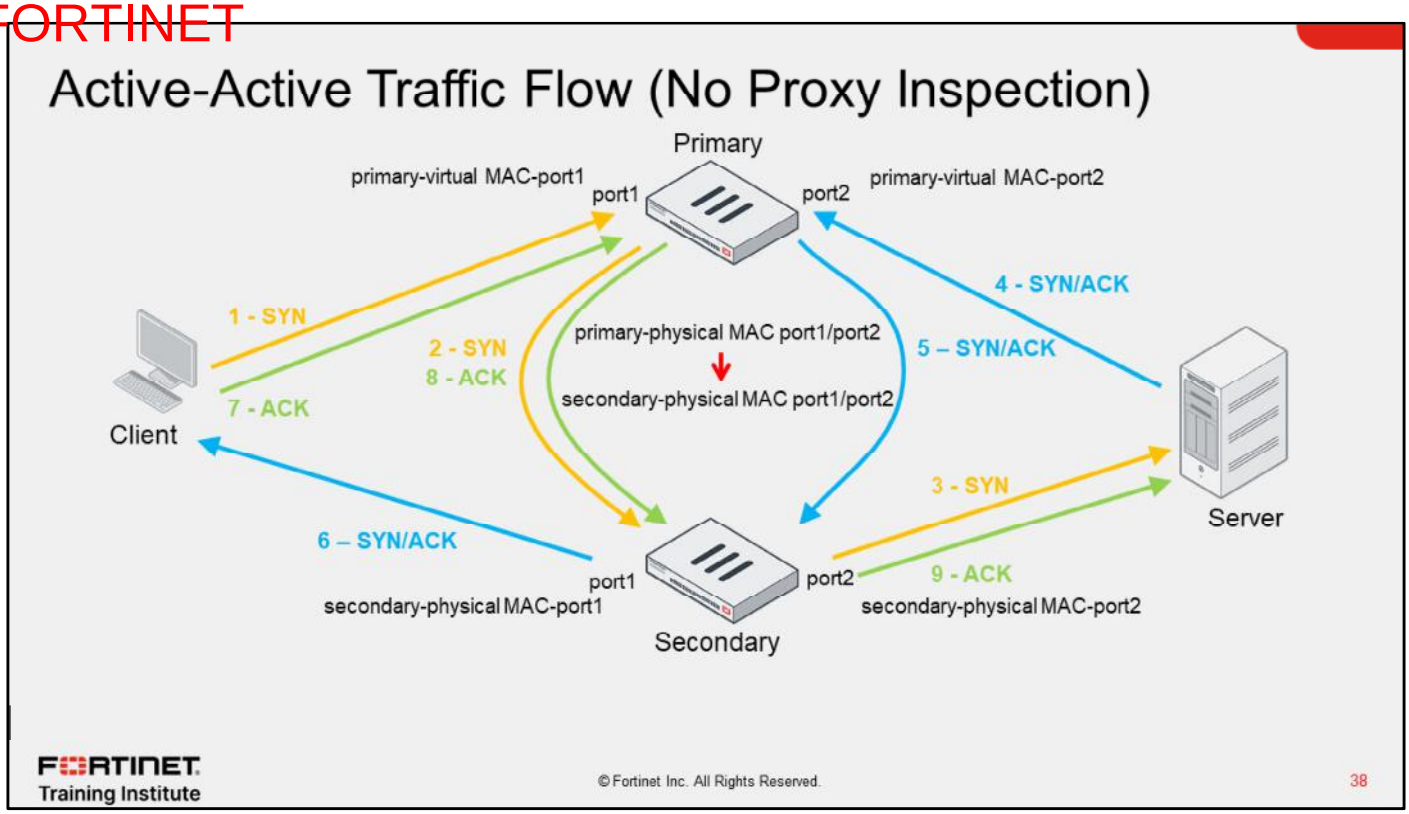

When there is no proxy inspection, that is, when traffic is either subject to flow inspection or no inspection at all, sessions are distributed to the secondary FortiGate only if you enable the `load-balance-all` setting (which is disabled by default) under HA configuration. In addition, as in proxy inspection, you will also see the following behavior:

- Traffic sourced from the client or server and destined to the FortiGate cluster is sent to the primary FortiGate. The source and destination MAC addresses are the endpoint (client or server) and the primary FortiGate virtual MAC address, respectively.
- The primary FortiGate may, in turn, forward the traffic to the secondary if the session is to be load balanced.
- When distributing the traffic to the secondary, FortiGate uses the physical MAC addresses of the primary and secondary devices interfaces as the source and destination MAC addresses, respectively.
- If traffic is load balanced to the secondary FortiGate, any traffic sourced from the cluster and destined to the endpoint is sourced from the secondary FortiGate. This means that the source MAC address is the physical address of the secondary egress interface.

When compared to proxy inspection, the difference is that FortiGate does not reply to packets on behalf of the client or server. For example, instead of replying to the SYN packet that the client sends, FortiGate forwards the packet to the server through the secondary. Similarly, FortiGate forwards packets that the server sends to the client through the secondary.

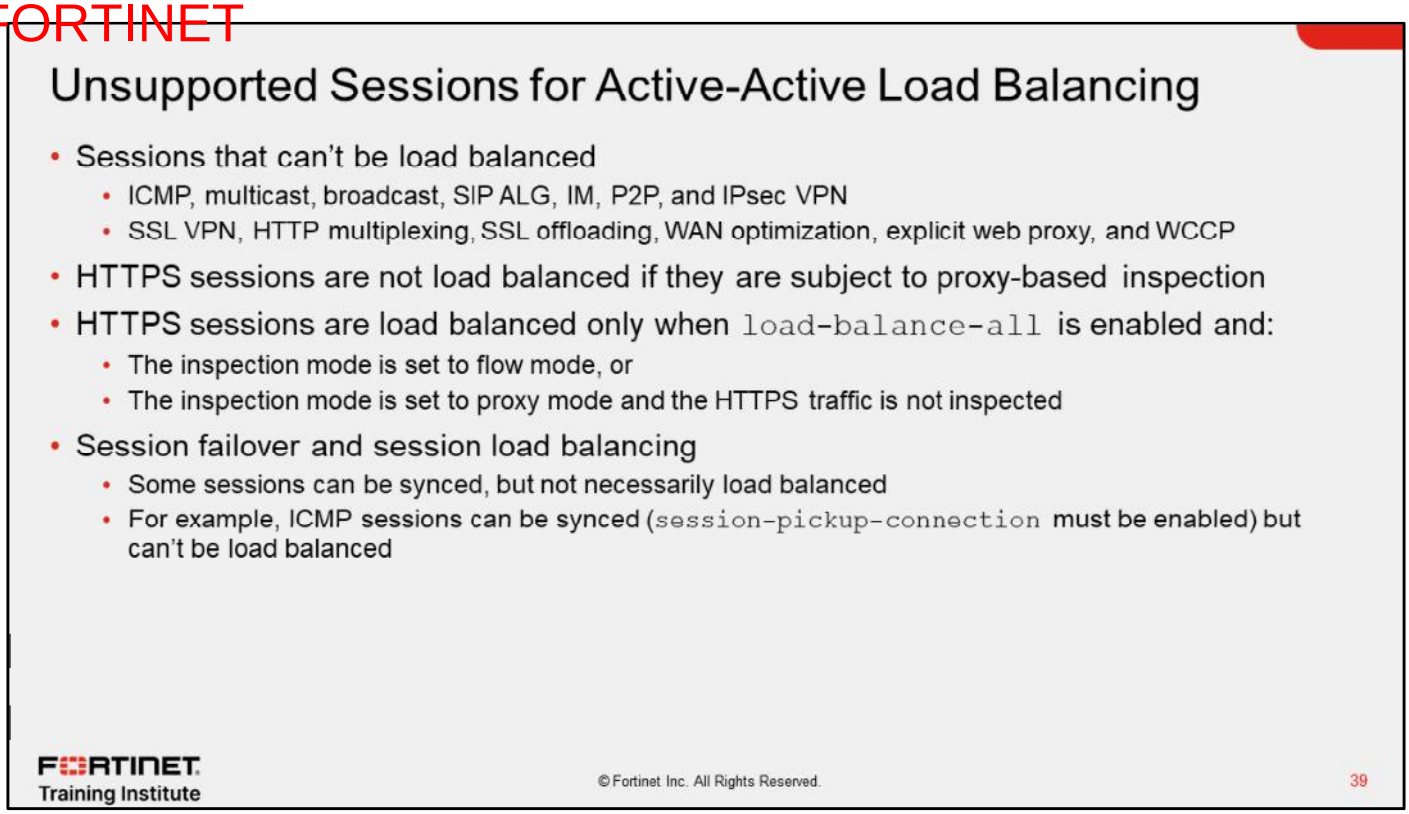

In active-active mode, not all sessions qualify for active-active load balancing. This slide shows a list of sessions that can't be load balanced.

Most of the internet traffic nowadays is HTTPS. For this reason, it is important to understand the limitations for HTTPS traffic load balancing. You must know that HTTPS sessions are not load balanced if they are subject to proxy-based inspection. In fact, the only two scenarios in which HTTPS sessions are load balanced is when the `load-balance-all` setting is enabled and:

- The inspection mode is set to flow mode, or
- The inspection mode is set to proxy mode and the HTTPS traffic is not inspected.

Do not confuse session failover with session load balancing. While some sessions can be synchronized to secondary members for session failover protection, those same sessions aren't necessarily supported for active-active load balancing. For example, ICMP sessions can be synchronized to secondary members if you enable the `session-pickup-connectionless` setting, but they cannot be load balanced.

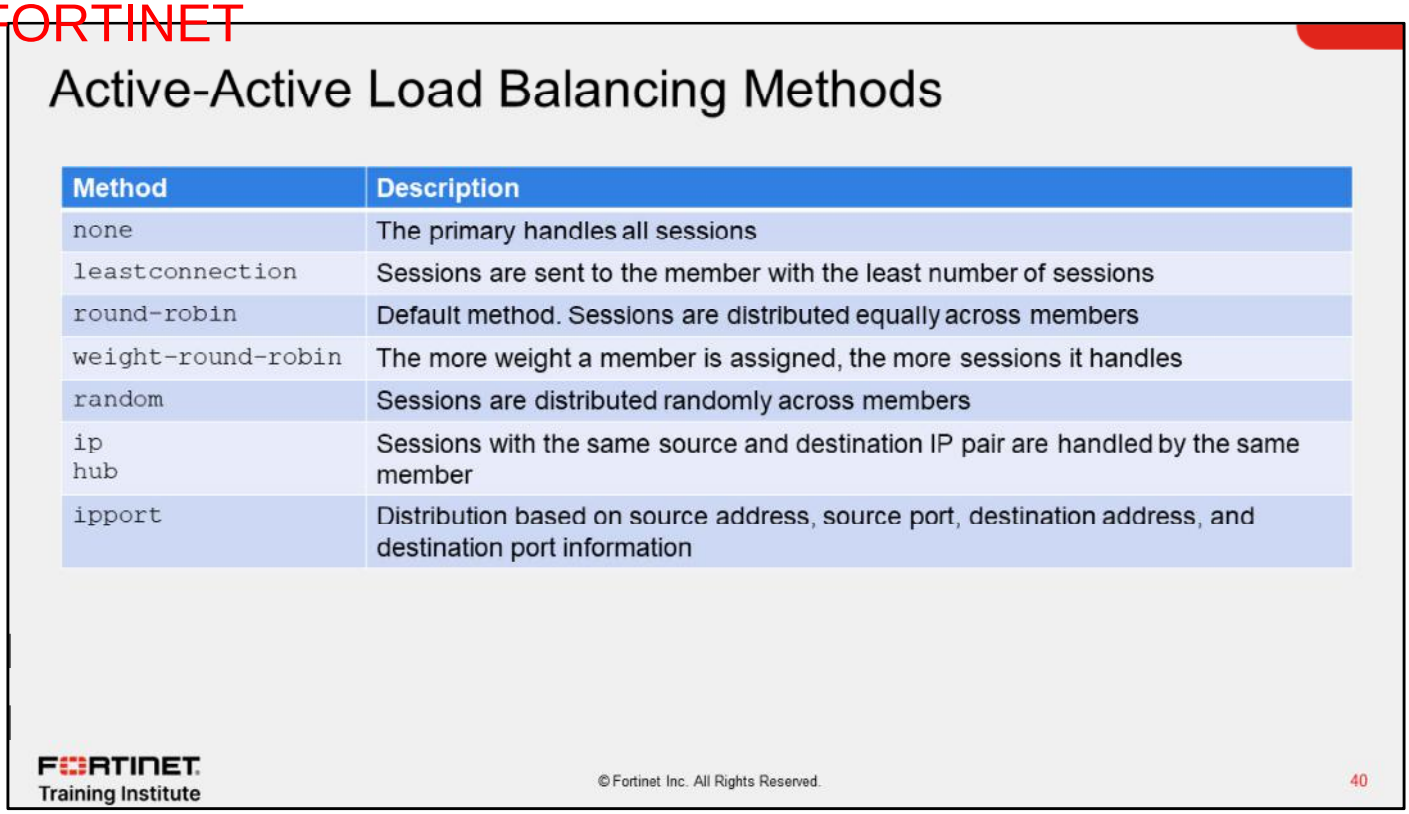

In active-active mode, when the primary device distributes sessions, it uses one of the following load balancing methods:

- `none`: Load balancing is turned off. The primary handles all sessions.
- `leastconnection`: The primary distributes sessions to the member with the least number of sessions.
- `round-robin`: This is the default method. The primary distributes sessions equally across members.
- `weight-round-robin`: The primary distributes sessions across members based on the member weight. The higher the member weight, the more sessions are distributed to that member.
- `random`: The primary distributes sessions randomly across members.
- `ip` and `hub`: The primary distributes sessions with the same source and destination IP pair to the same member. Both methods, `ip` and `hub`, work the same way. Both names in the configuration were kept for legacy compatibility purposes. The `hub` schedule will be removed in a future FortiOS version.
- `ipport`: The primary distributes sessions based on the source address, source port, destination address, and destination port information. The more diverse the traffic is, the more evenly the traffic is distributed across members.

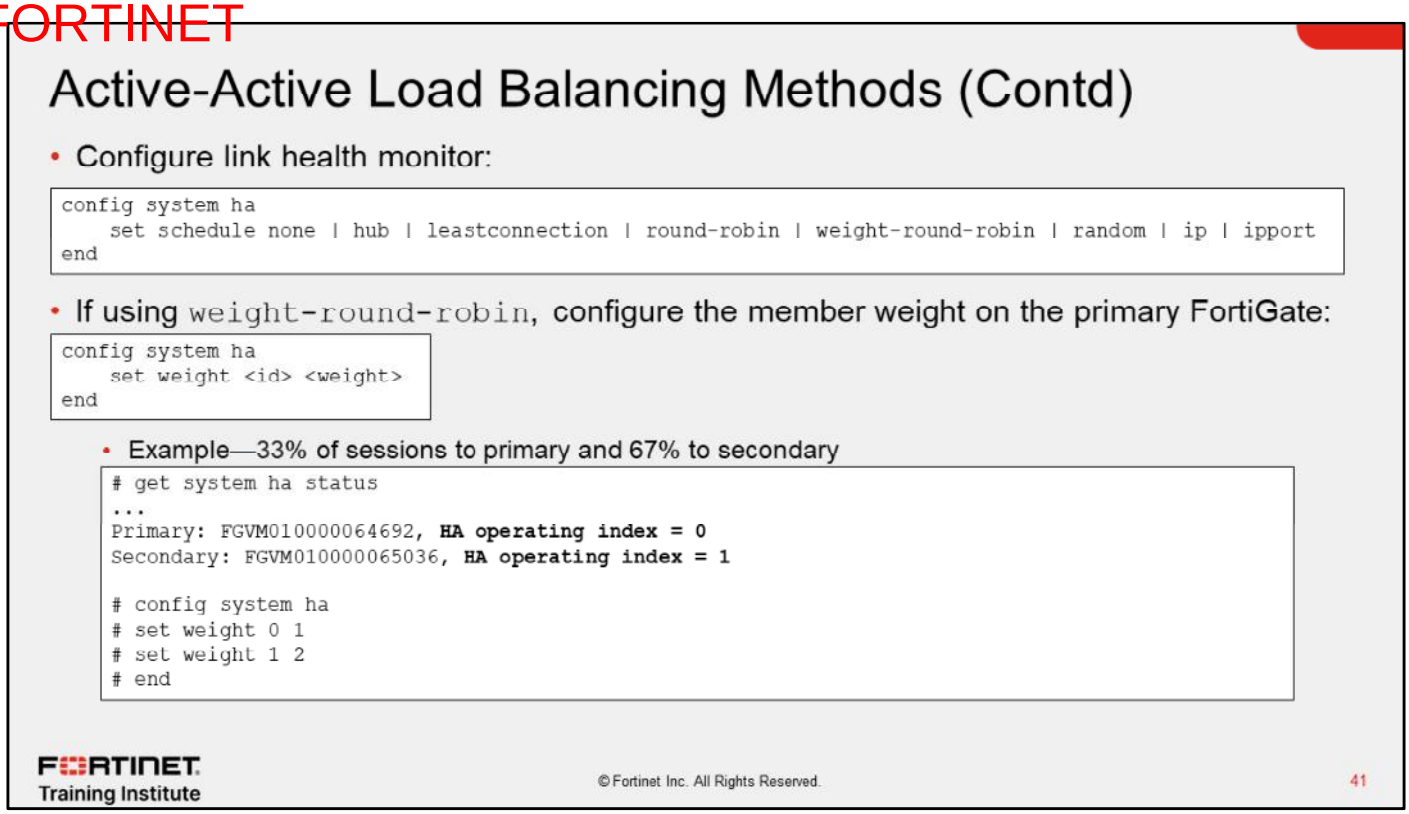

You set the load balancing method by configuring the `schedule` setting, as shown on this slide.

When you select the `weight-round-robin` method, you must also configure the weight for each member, as shown on this slide. You indicate the member ID followed by its weight. The higher the member weight, the more sessions are distributed to that member. You can obtain the member ID from the output of the `get system ha status` command.

This slide also shows a configuration example for a weight-based distribution of 67% of sessions to the secondary FortiGate and 33% of sessions to the primary device. That is, for every three connections that qualify for load balancing, two of them are distributed to the secondary, and one of them to the primary.

Note that you apply the member weight configuration for all members on the primary device. That is, you don't have to apply the weight on each member individually. The cluster will synchronize the configuration to each member for you.

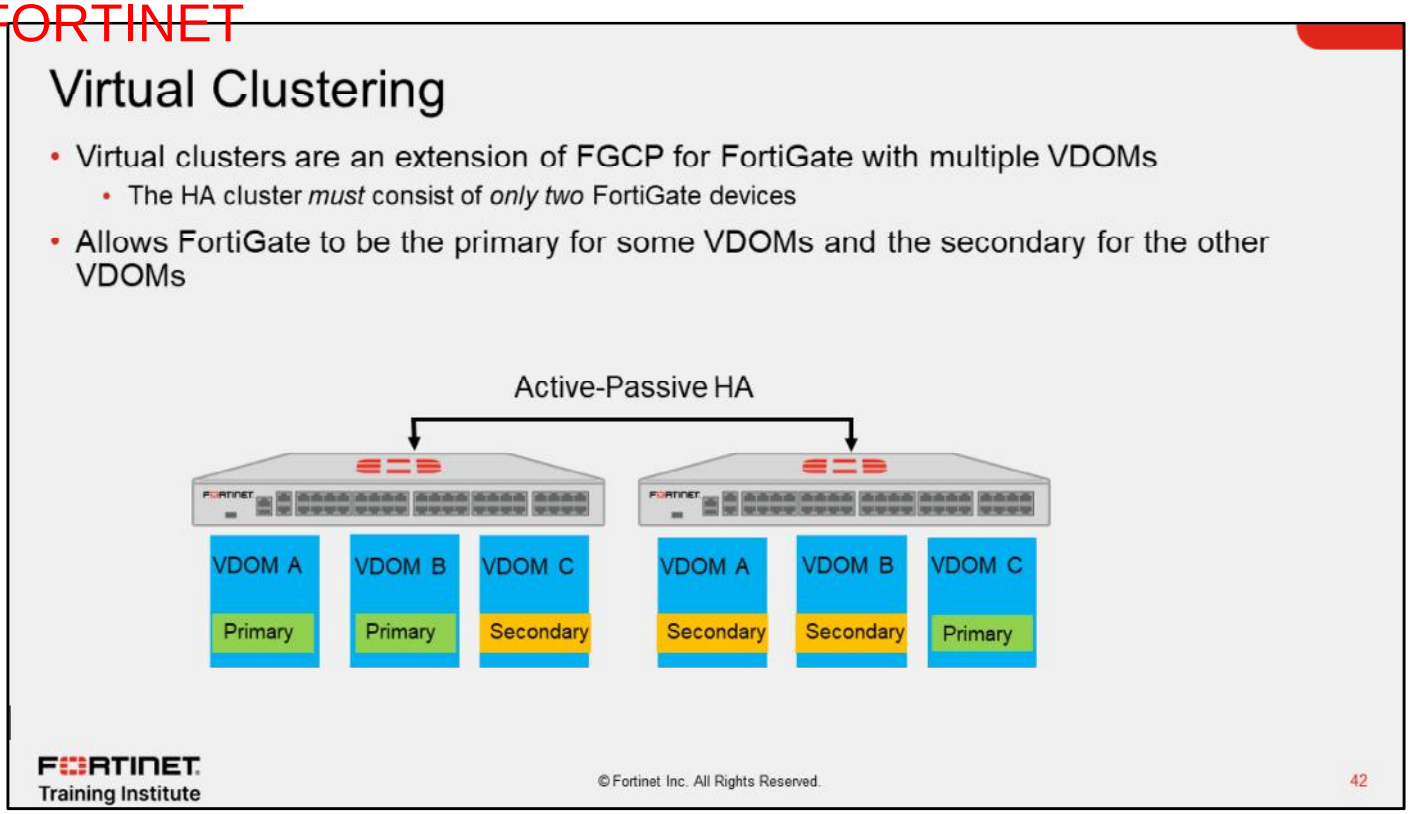

So far, you've learned about HA clustering where each FortiGate device acts as a whole security domain.

But, if you have an HA cluster with multiple VDOMs, you can configure *virtual clusters*.

Virtual clusters allow you to have one device acting as the primary for one VDOM, and as the secondary for a different VDOM. Each VDOM has a primary and a secondary FortiGate. Any device can act as the primary for some VDOMs, and the secondary for the other VDOMs, at the same time. Because traffic from different VDOMs can go to different primary FortiGate devices, you can use virtual clustering to manually distribute your traffic between the two cluster devices, and allow the failover mechanism for each VDOM between two FortiGate devices.

Note that you can configure virtual clustering between *only two* FortiGate devices with multiple VDOMs.

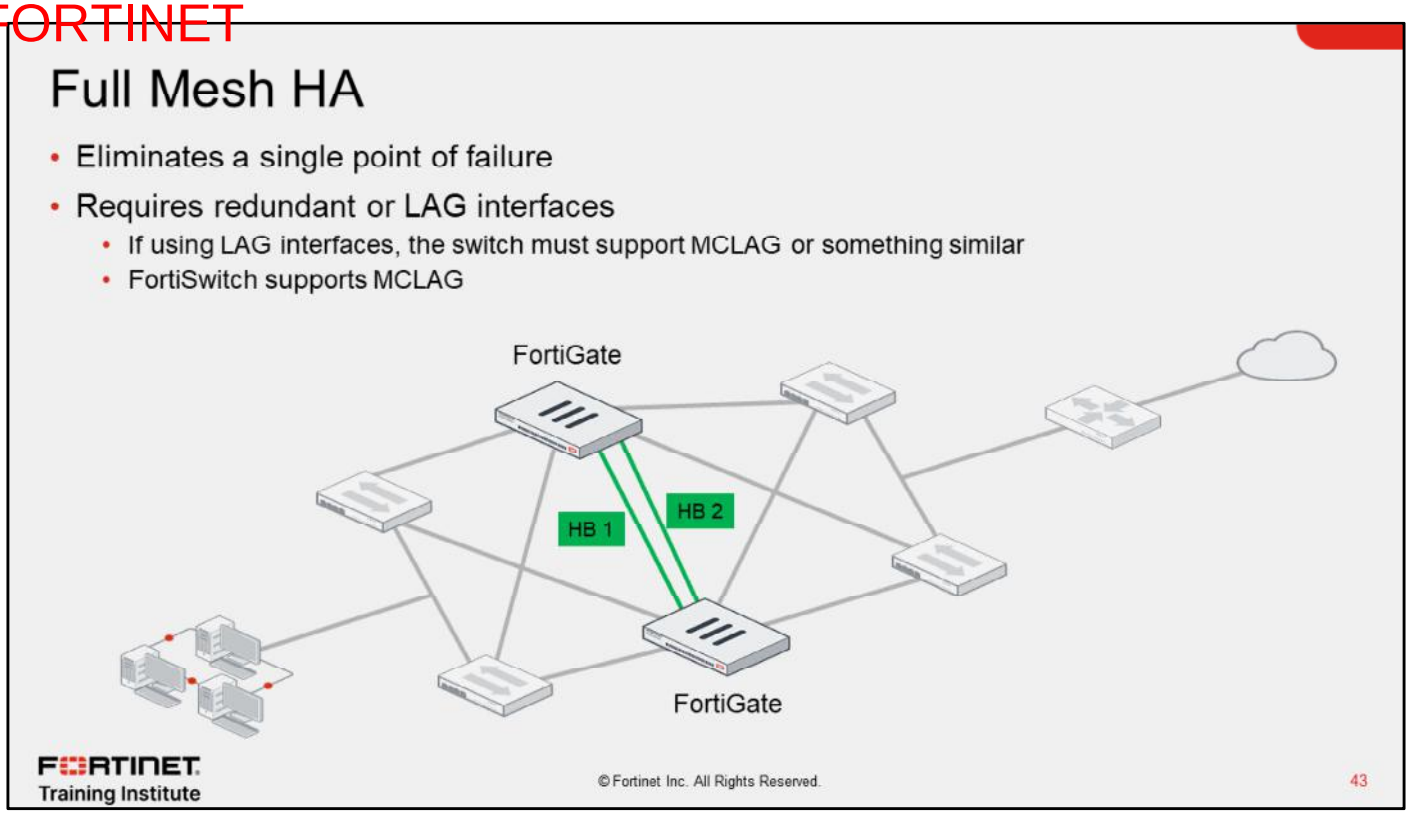

At the beginning of this lesson, you reviewed a simple HA topology. Now, take a look at a more robust topology. It is called *full mesh HA*.

The goal of a full mesh HA topology is to eliminate a single point of failure, not only by having multiple FortiGate devices forming a cluster, but also by having redundant links to the adjacent switches. The goal is to have two switches for both upstream and downstream links, and then connect the redundant links to different switches. For example, the topology on this slide shows two FortiGate devices forming a cluster, and each FortiGate is connected to two redundant switches, using two different interfaces.

To achieve redundancy with adjacent switches, you must deploy redundant or LAG interfaces. If you use redundant interfaces, only one interface remains active. This avoids a Layer 2 loop and a standard switch should suffice. However, if you want to use LAG interfaces, then you must ensure that the switch supports multichassis link aggregation group (MCLAG) or a similar virtual LAG technology that enables you to form a LAG whose interface members connect to different switches. FortiSwitch, which is a Fortinet Ethernet switch, supports MCLAG. You can use FortiSwitch as the adjacent switch to deploy a full mesh HA topology with FortiGate.

Knowledge Check

1. An HA failover occurs when the link status of a monitored interface on the _____ goes down.
   - ✓ A. Primary FortiGate
   - B. Secondary FortiGate

2. In an active-passive HA cluster, you can configure virtual clustering between only ___ FortiGate devices with multiple VDOMs.
   - ✓ A. Two
   - B. Four

FORTINET Training Institute

44

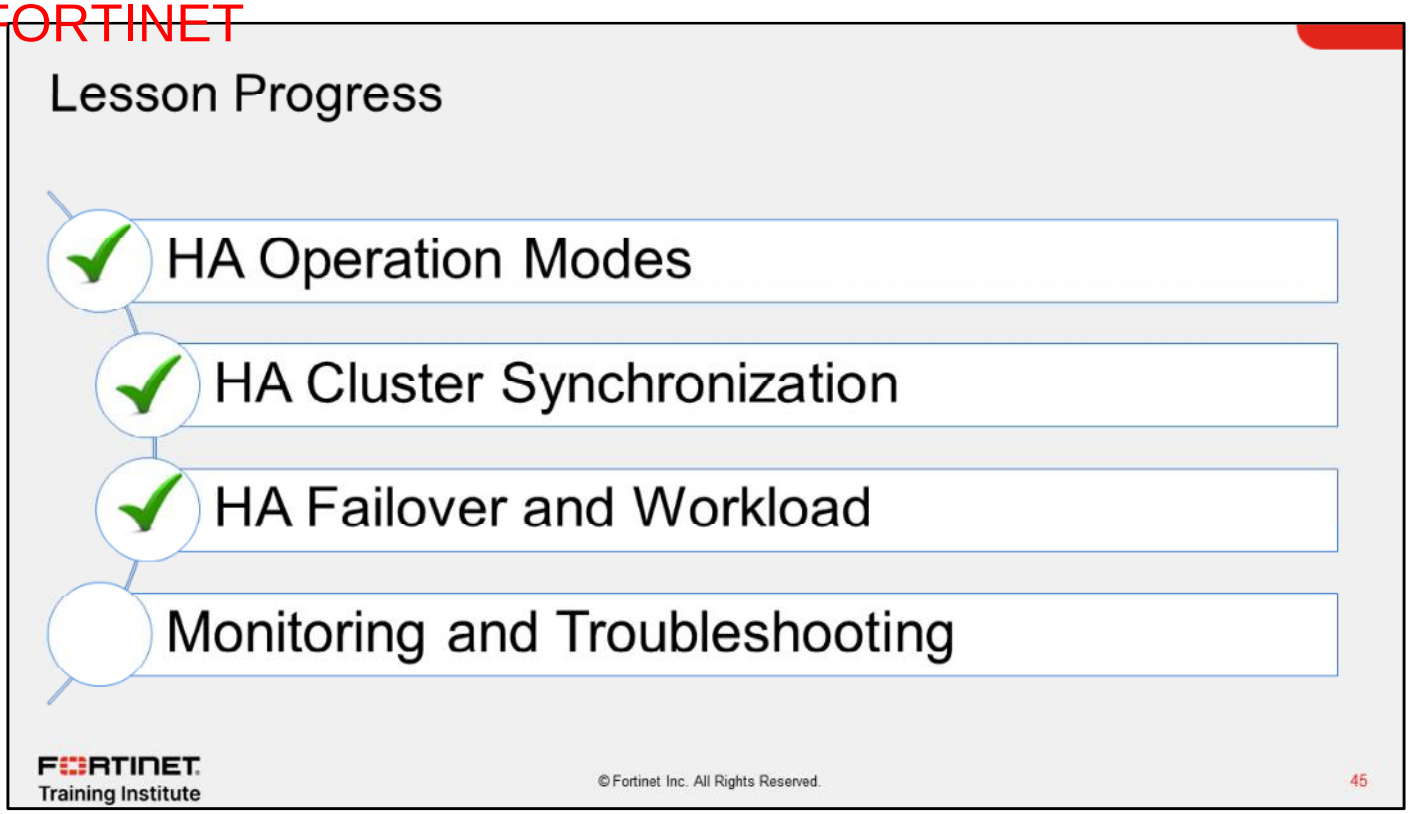

Good job! You now understand HA failover and workload.

Now, you will learn about monitoring and troubleshooting an HA cluster.

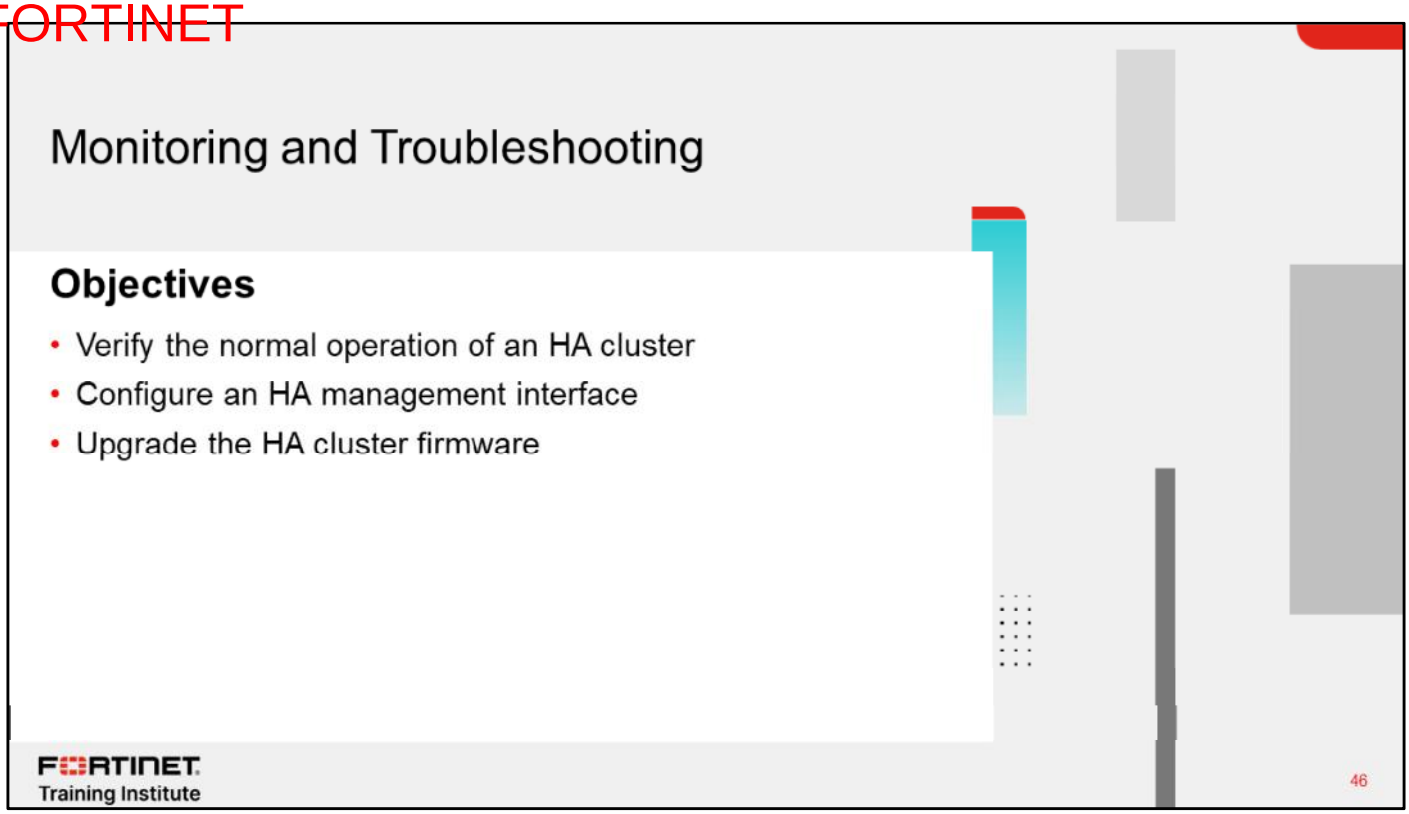

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in monitoring and troubleshooting, you will be able to make sure the cluster is synchronized properly. You will also learn how to configure and access secondary devices in an HA cluster and how to upgrade the firmware on the HA cluster.

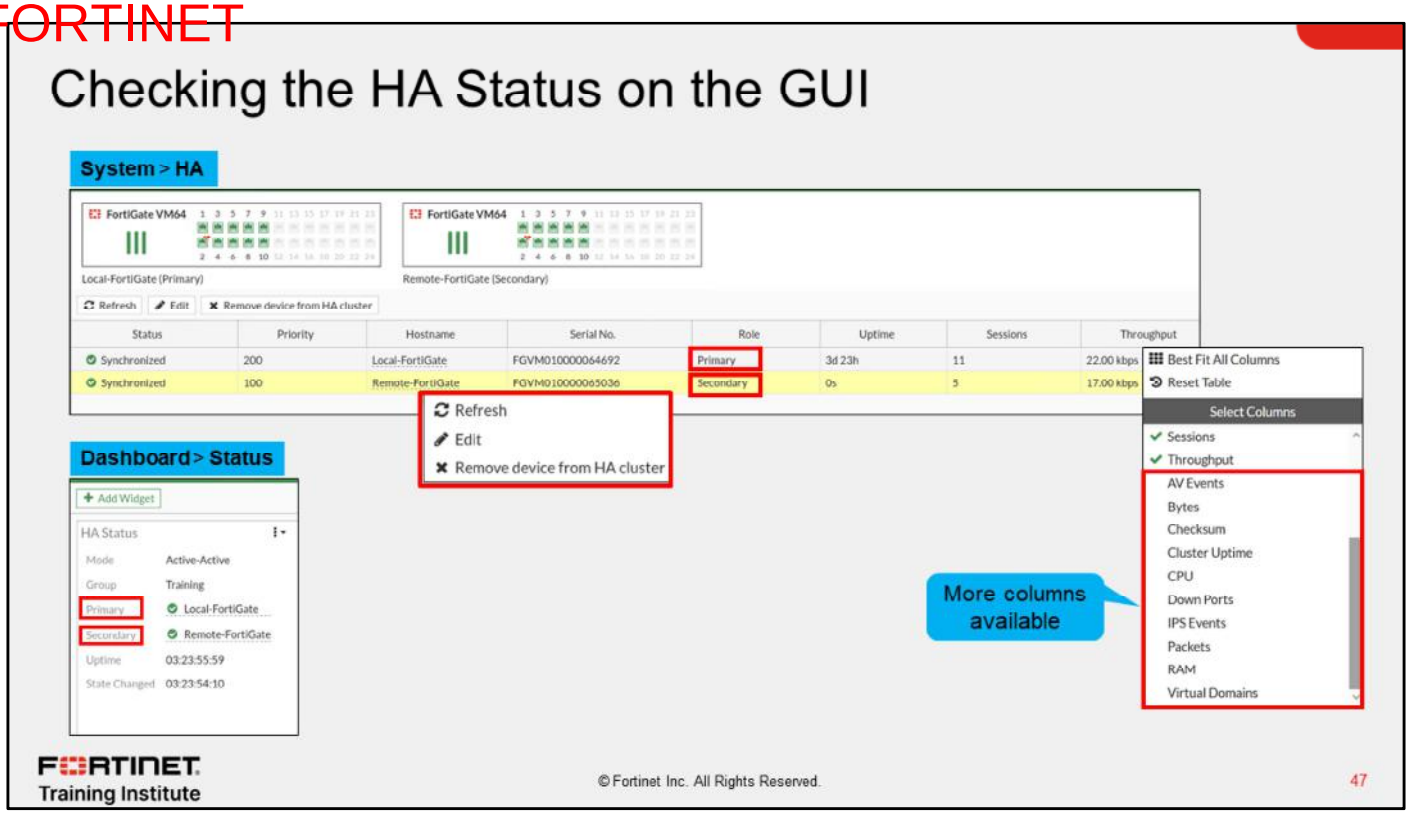

The **HA** page on the FortiGate GUI shows important information about the health of your HA cluster. For each cluster member, the page shows whether the member is synchronized or not, and its status, host name, serial number, role, priority, uptime, and active sessions.

On the **HA** page, you can remove a device from a cluster. When you remove a device from HA, the device operation mode is set to standalone. You can also enable more columns that display other important information about each member such as the checksum, CPU, and memory.

You can also add the **HA Status** widget on the **Dashboard** page. The widget provides a summary of the HA status on the device.

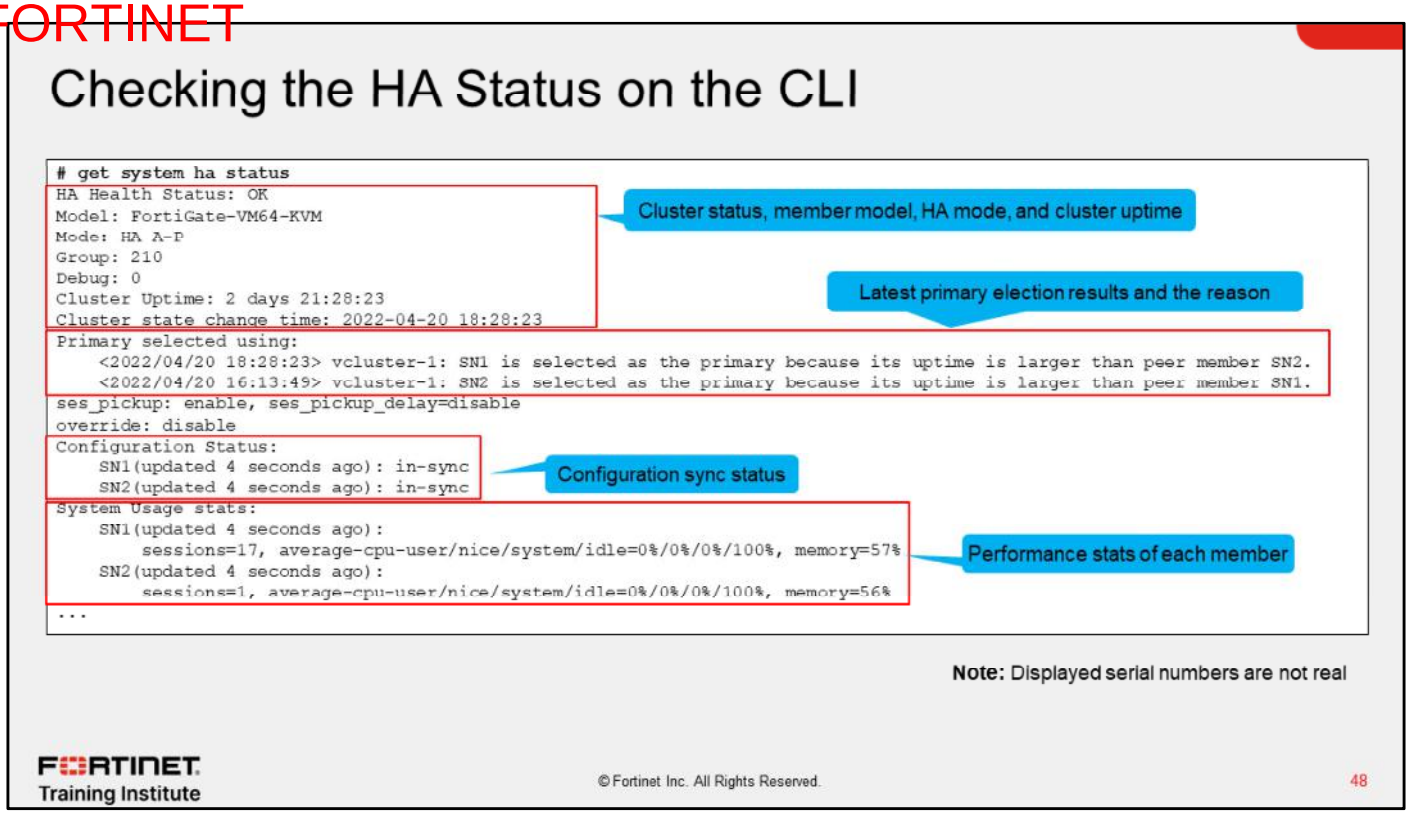

You can get more information about the HA status on the FortiGate CLI by using the `get system ha status` command.

The command displays comprehensive HA status information in a user-friendly output and is usually executed as the first step when troubleshooting HA. This slide shows the first part of an example output that the command provides. Note that the serial numbers of members have been replaced by fake ones (SN1 and SN2), so the output fits on this slide.

At the beginning of the output, you can see the cluster status, the member model, the HA mode in use, and the cluster uptime. The example output shows that the cluster status is good, the member model is FortiGate-VM64-KVM, and the HA mode is active-passive. The cluster has also been up for almost three days.

Next, you can see the latest primary election events, the result, and the reason. The output indicates that a different member was elected as the primary during the last two election events. In both cases, the member was elected because it had a higher HA uptime.

The configuration status information is displayed next. It indicates the configuration sync status for each member. For both members, the configuration is in sync.

Following the configuration status information, you can see the system usage statistics, which report on performance statistics for each member. They indicate the number of sessions that each member handles, as well as the average CPU and memory usage. Note that the `sessions` field accounts for any sessions that the member handles, and not only the sessions that are distributed when the HA mode is active-active.

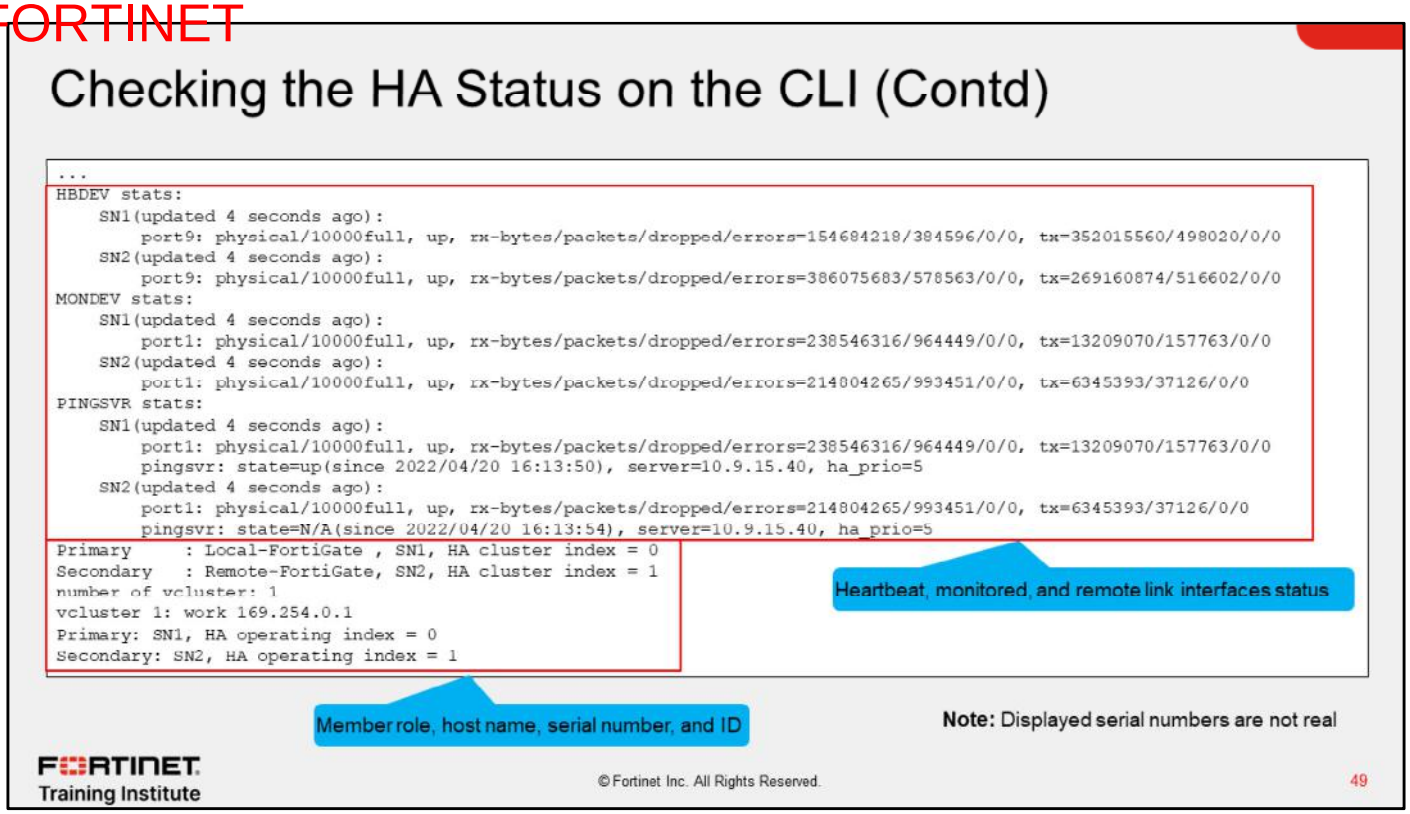

This slide shows the second part of the example output that the `diagnose system ha status` command provides.

The output begins with the status information for the configured heartbeat, monitored, and remote link interfaces. These interfaces enable the cluster to perform device failover, link failover, and remote link failover protection, respectively.

Next, the output shows the role, host name, serial number, and ID information for each member of the cluster. The output indicates that the `Local-FortiGate` and `Remote-FortiGate` devices are primary and secondary members, respectively.

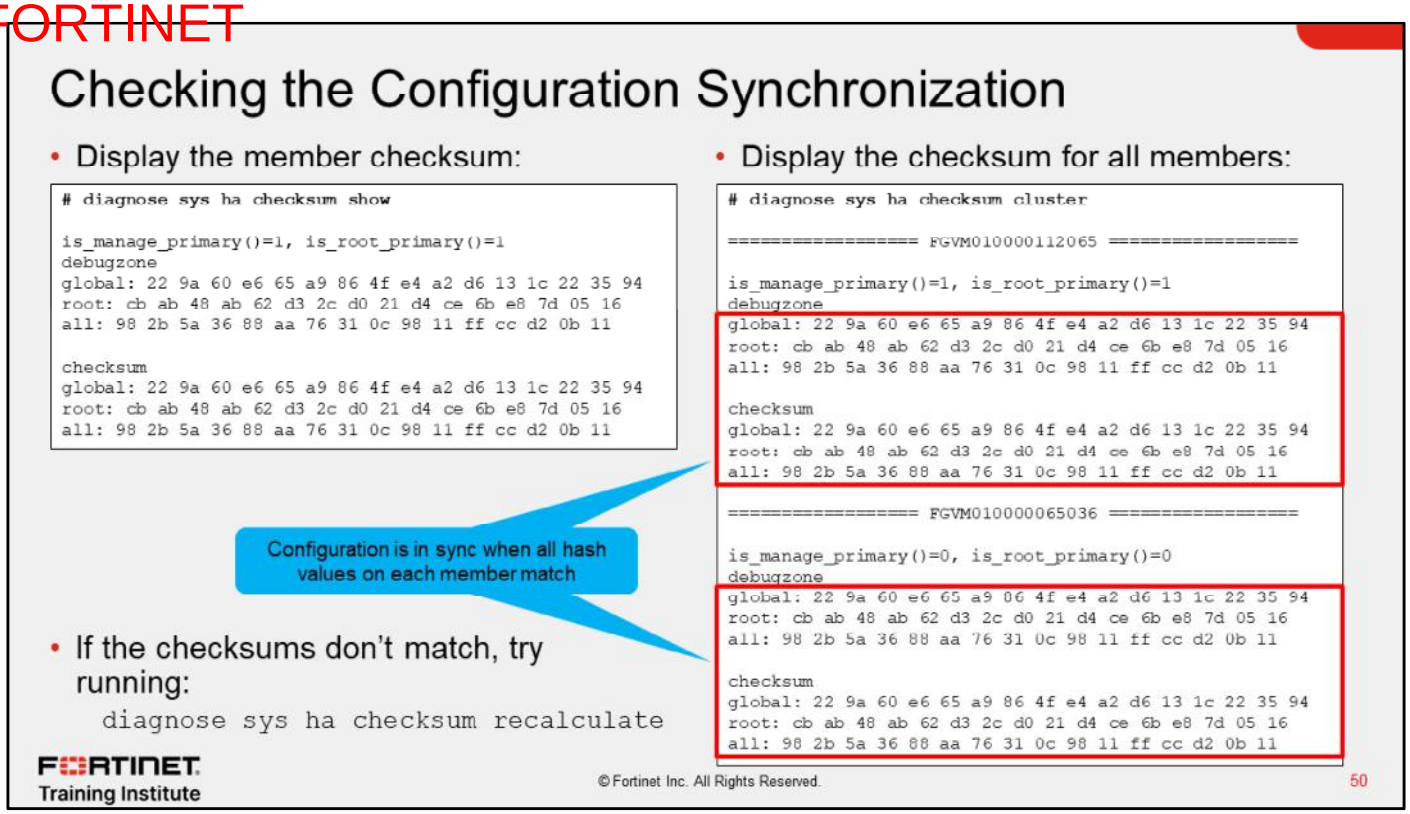

The `diagnose sys ha checksum` command tree enables you to check the cluster configuration sync status. In most cases, you want to use the `diagnose sys ha checksum cluster` command to view the cluster checksum. The output includes the checksum of each member in the cluster.

When you run the `diagnose sys ha checksum cluster` command, the checksum is polled from each member using the heartbeat interface. If HA is not working properly, of if there are heartbeat communication issues, then the command may not show the checksum for members other than the one you run the command on. An alternative is to connect to each member individually and run the `diagnose sys ha checksum show` command instead. This command displays only the checksum of the member you are connected to.

After you obtain the checksums of each member, you can identify the configuration sync status by comparing the checksums. If all members show the exact hash values for each configuration scope, then the configuration of all members is in sync.

To calculate checksums, FortiGate computes a hash value for each of the following configuration scopes:

- `global`: global configuration, such as global settings, FortiGuard settings, and so on
- `root`: settings and objects specific to the root VDOM—if you configure multiple VDOMs, FortiGate computes hash values for each VDOM
- `all`: global configuration plus the configuration of all VDOMs

In some cases, the configuration of members is in sync even though the checksums are different. For these cases, try running the `diagnose sys ha checksum recalculate` command to recalculate the HA checksums.

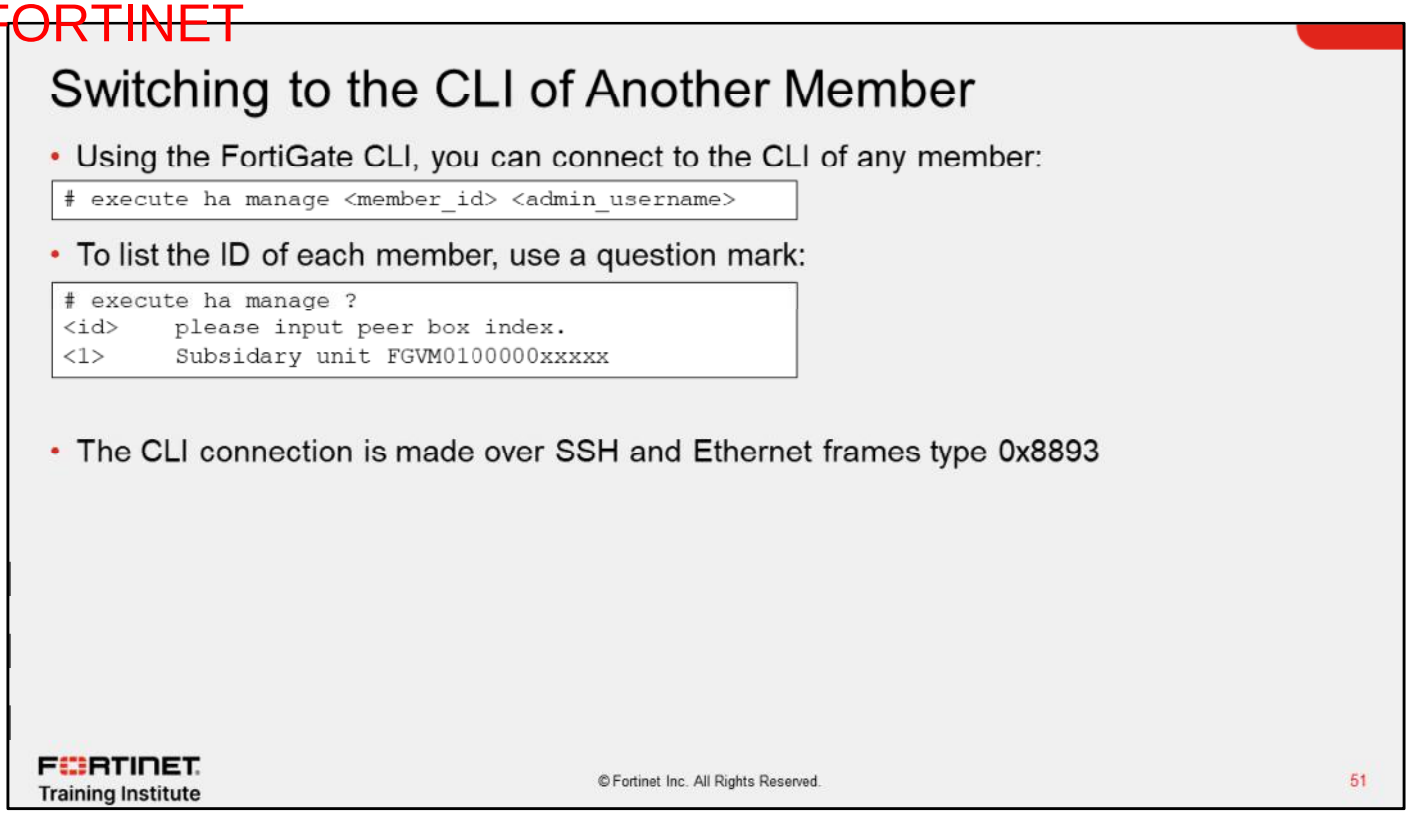

When troubleshooting HA, you may need to connect to the CLI of another member from the CLI of the member you are currently connected to. You do this by using the `execute ha manage` command to connect to the other member.

For example, when you connect to the cluster over SSH using any of the cluster virtual IP addresses, you connect to the primary member. If you then want to connect to another member, you can use the `execute ha manage` command to access its CLI.

This command requires you to indicate the ID of the member you want to connect to and the username you will use to log in. To get the list of member IDs, you can add a question mark to the end of the `execute ha manage` command, as shown on this slide.

Note that when you switch to the CLI of another member, FortiGate establishes an SSH session to that member over the heartbeat interface. The SSH session is then encapsulated in Ethernet frames type 0x8893.

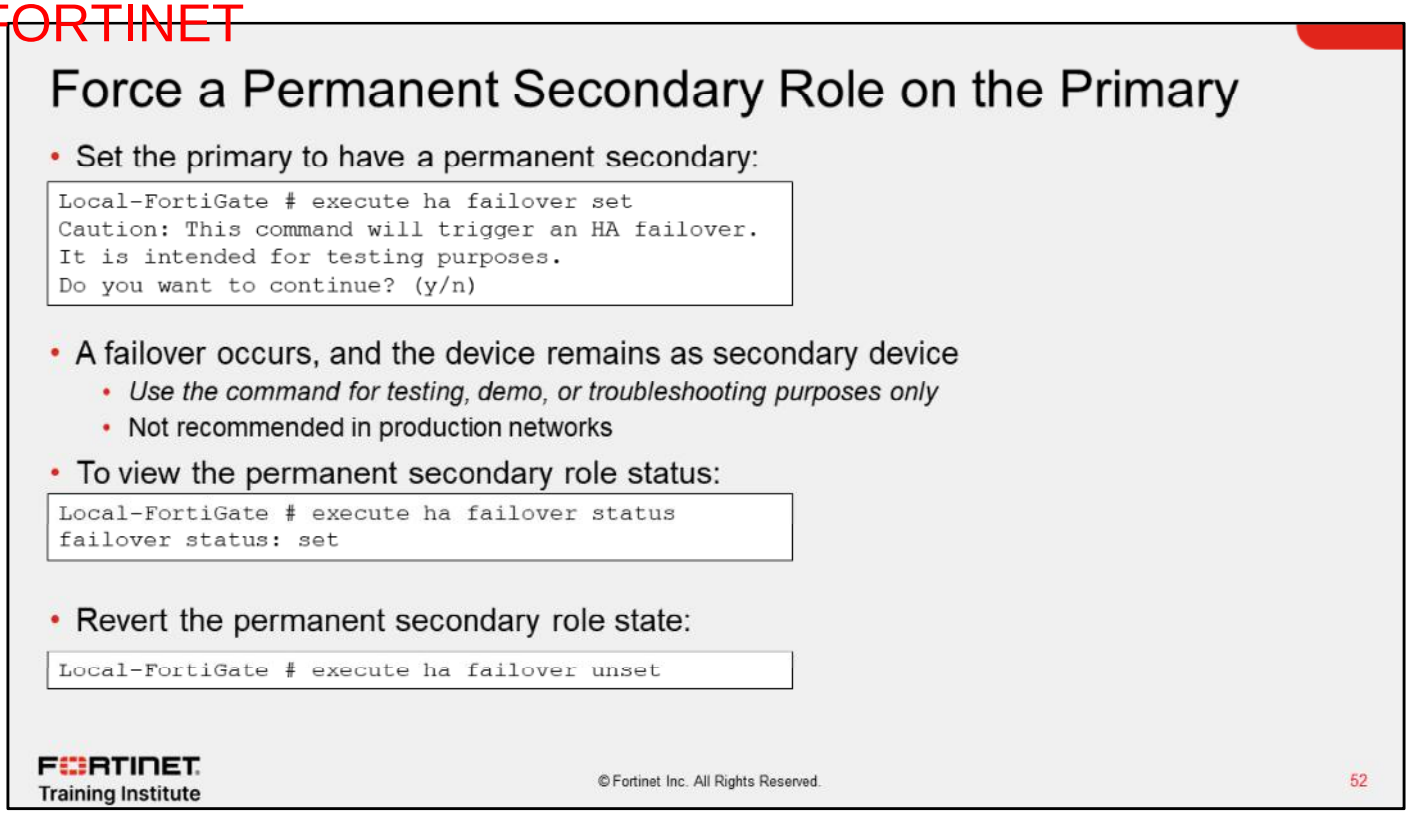

You can set the primary FortiGate to have a permanent secondary role using the `execute ha failover set` command. When you do this, a failover occurs, and the former primary member remains as a secondary member permanently, regardless of the status of other members in the cluster. That is, the impacted member never takes over the cluster even if it's the best candidate for the primary role.

You can revert the permanent secondary role state by running the `execute ha failover unset` command. Note that you should set the primary member to a permanent secondary role for testing, troubleshooting, and demonstration purposes only. Do not use this feature in production networks.

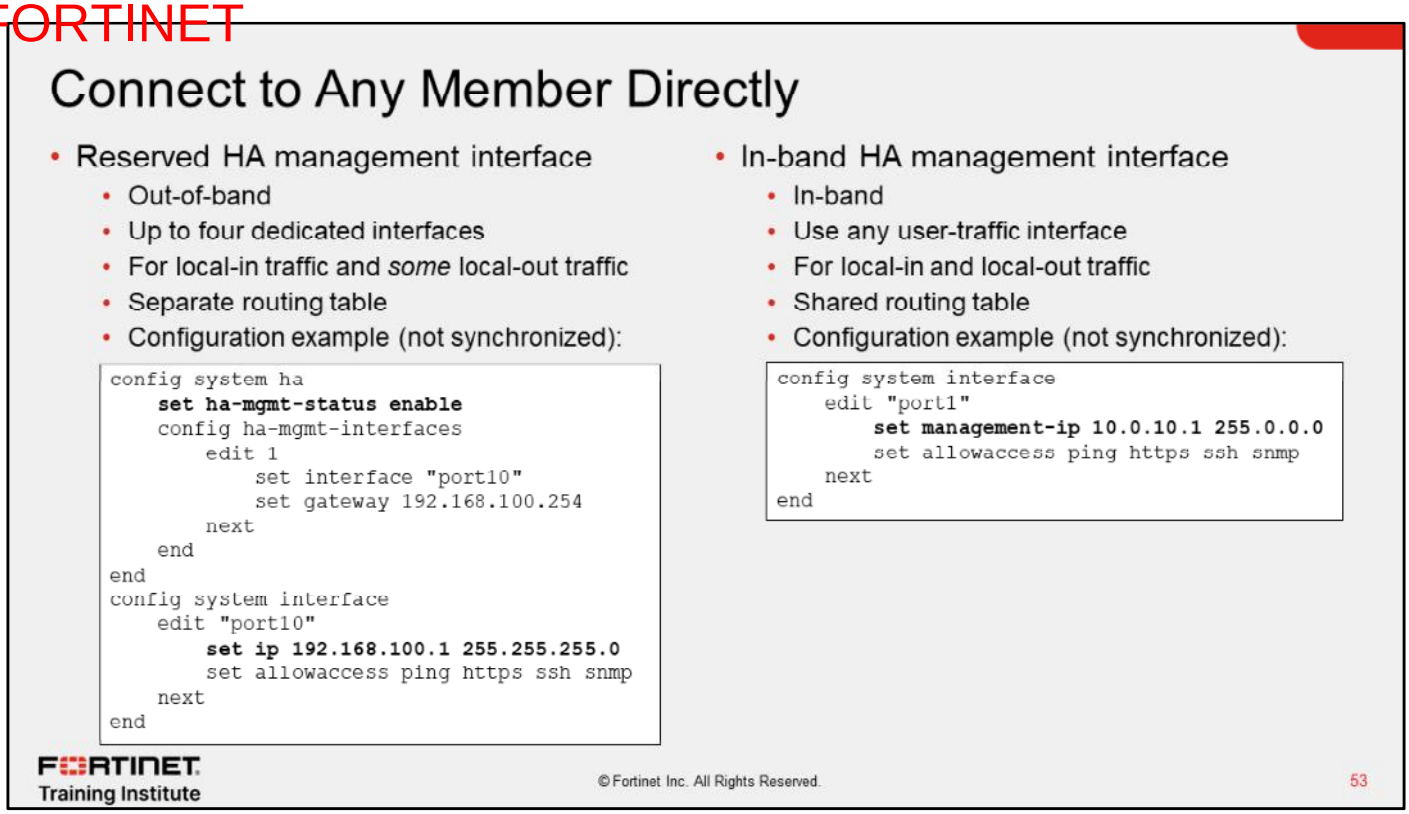

When you connect to a cluster using any of its virtual IP addresses, you always connect to the primary. You can then switch to the CLI of any member in the cluster by using the `execute ha manage` command. But what if you want to access the GUI of a secondary member or maybe poll data from it using SNMP? For this, you need a way to access each member directly regardless of its role in the cluster.

FortiGate provides two ways for the administrator to connect to a member directly no matter what the member role is. The reserved HA management interface is the out-of-band option. You configure up to four dedicated management interfaces, and you assign them a unique address on each member. You can then use the unique address assigned to each member to connect to them directly. You can also instruct FortiGate to use the dedicated management interface for some outbound management services such as SNMP traps, logs, and authentication requests.

Alternatively, you can configure in-band HA management, which enables you to assign a unique management address to a member without having to set aside an interface for that purpose. You assign the management address to any user-traffic that the member uses, and then connect to the member using that unique management address.

If you have unused interfaces, then it's generally more convenient to use a reserved HA management interface because the user and management traffic don't have to compete. Many FortiGate models come with a management interface that you can use for this purpose. Also, the routing information for a reserved HA management interface is placed in a separate routing table, which means that you don't see the interface routes in the FortiGate routing table. This allows for segmentation between data and management traffic.

This slide also shows configuration examples for both management options. For both options, the configuration you apply on a member is not synchronized to other members in the cluster.

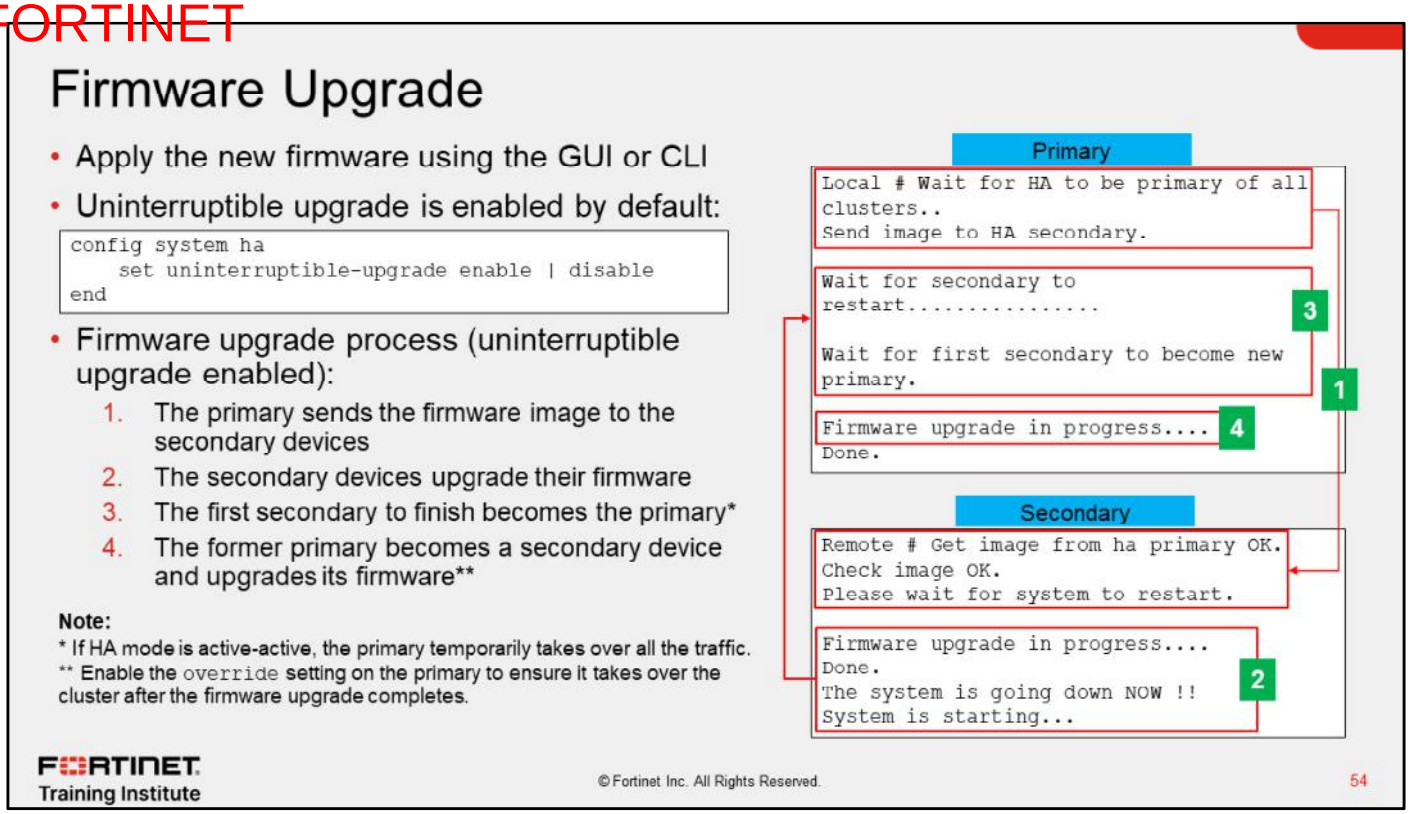

You upgrade an HA cluster in the same way you do for standalone FortiGate devices. That is, you can apply the new firmware using the GUI firmware upgrade tool. In HA, this usually means connecting to the primary FortiGate GUI to apply the new firmware. You can also use the CLI if you prefer.

Also, like on standalone FortiGate devices, the device must reboot to apply the new firmware. However, by default, members in a cluster are upgraded one at a time to minimize service disruption. This feature is called uninterruptible upgrade and is enabled by default. After the administrator applies the new firmware on the primary, uninterruptible upgrade works as follows:

1. The primary sends the firmware to all secondary members using the heartbeat interface.
2. The secondary devices upgrade their firmware first. If the cluster is operating in active-active mode, the primary temporarily takes over all traffic.
3. The first secondary that finishes upgrading its firmware takes over the cluster.
4. The former primary becomes a secondary device and upgrades its firmware next.

Note that depending on the HA settings and uptime, the original primary may remain as a secondary after the upgrade. Later, if required, you can issue a manual failover. Alternatively, you can enable the `override` setting on the primary FortiGate to ensure it takes over the cluster again after it upgrades its firmware, as long as the device is assigned the higher priority.

If you want the cluster to upgrade all members at the same time to speed up the firmware upgrade process, you can disable uninterruptable upgrade, as shown on this slide. Just keep in mind this will result in a service impact during the firmware upgrade.

Knowledge Check

1. Which member is the heartbeat interface IP address `169.254.0.1` assigned to?
   - ✓ A. The member with the highest serial number
   - B. The member with the highest priority

2. Which statement about the firmware upgrade process on an HA cluster is true?
   - ✓ A. You upload the new firmware to the primary FortiGate only.
   - B. The members do not reboot.

FORTINET Training Institute

55

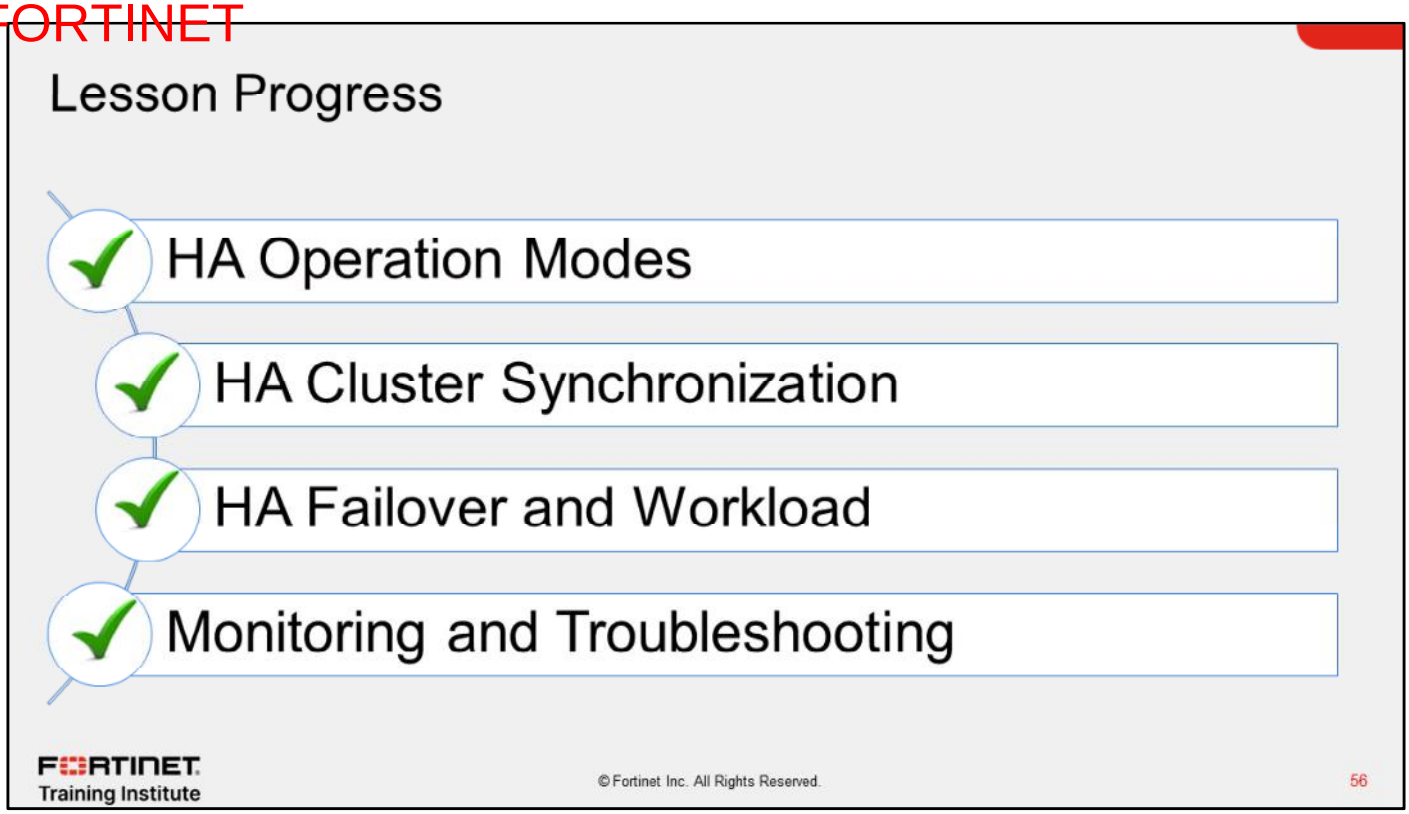

Congratulations! You have completed this lesson.

Now, you will review the objectives that you covered in this lesson.

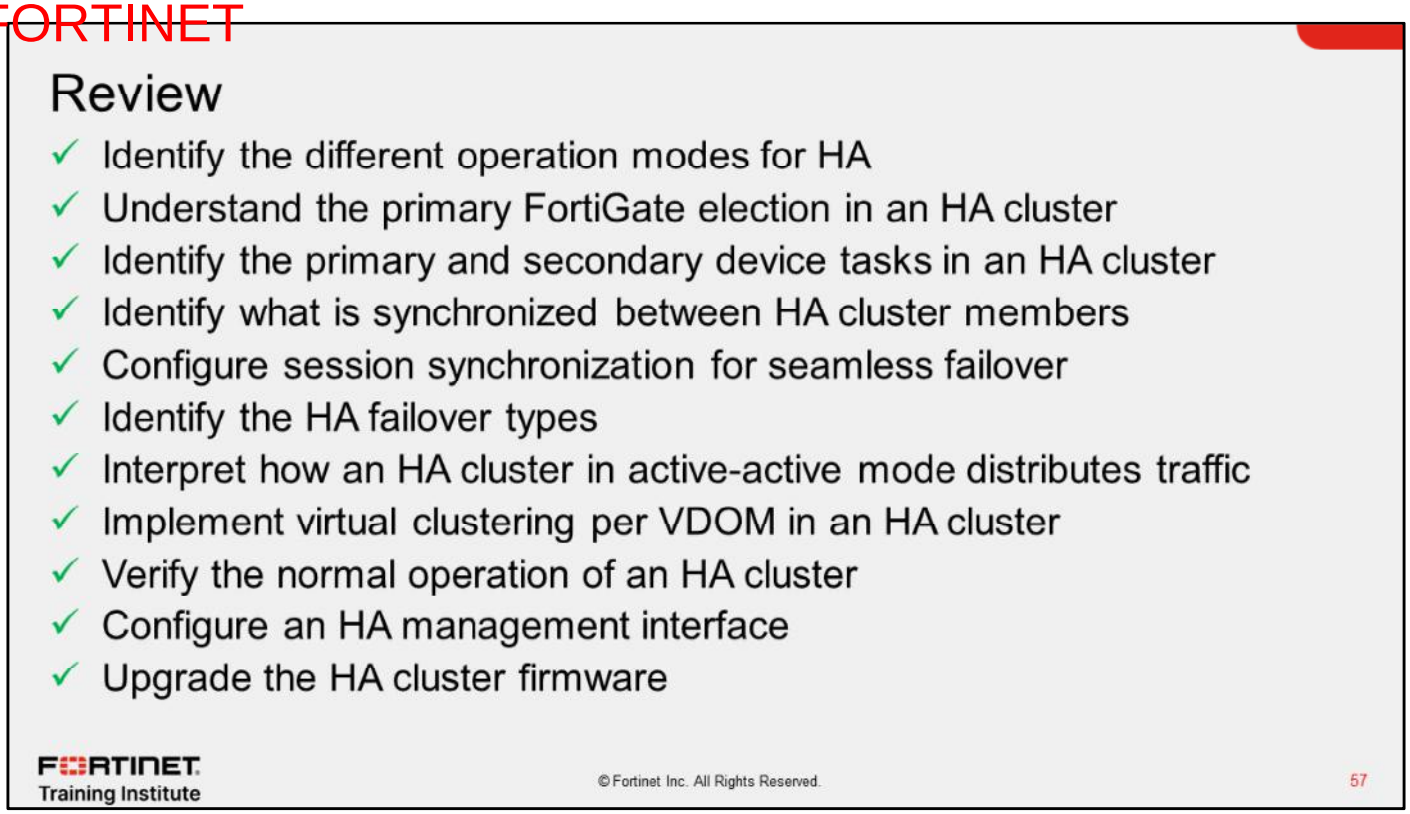

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned about the fundamentals of FortiGate HA and how to configure it.

In this lesson, you will learn about using diagnostic commands and tools.

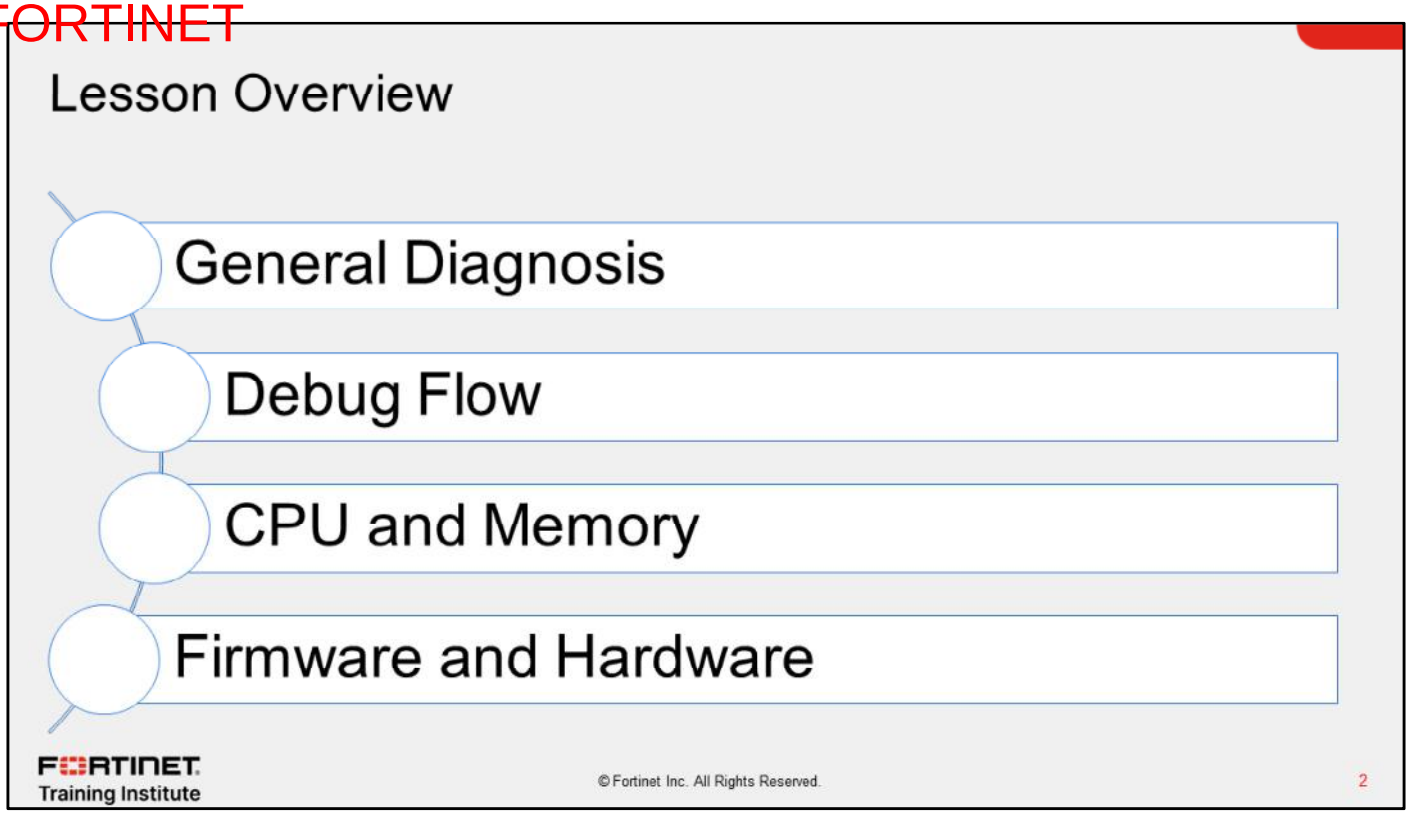

In this lesson, you will learn about the topics shown on this slide.

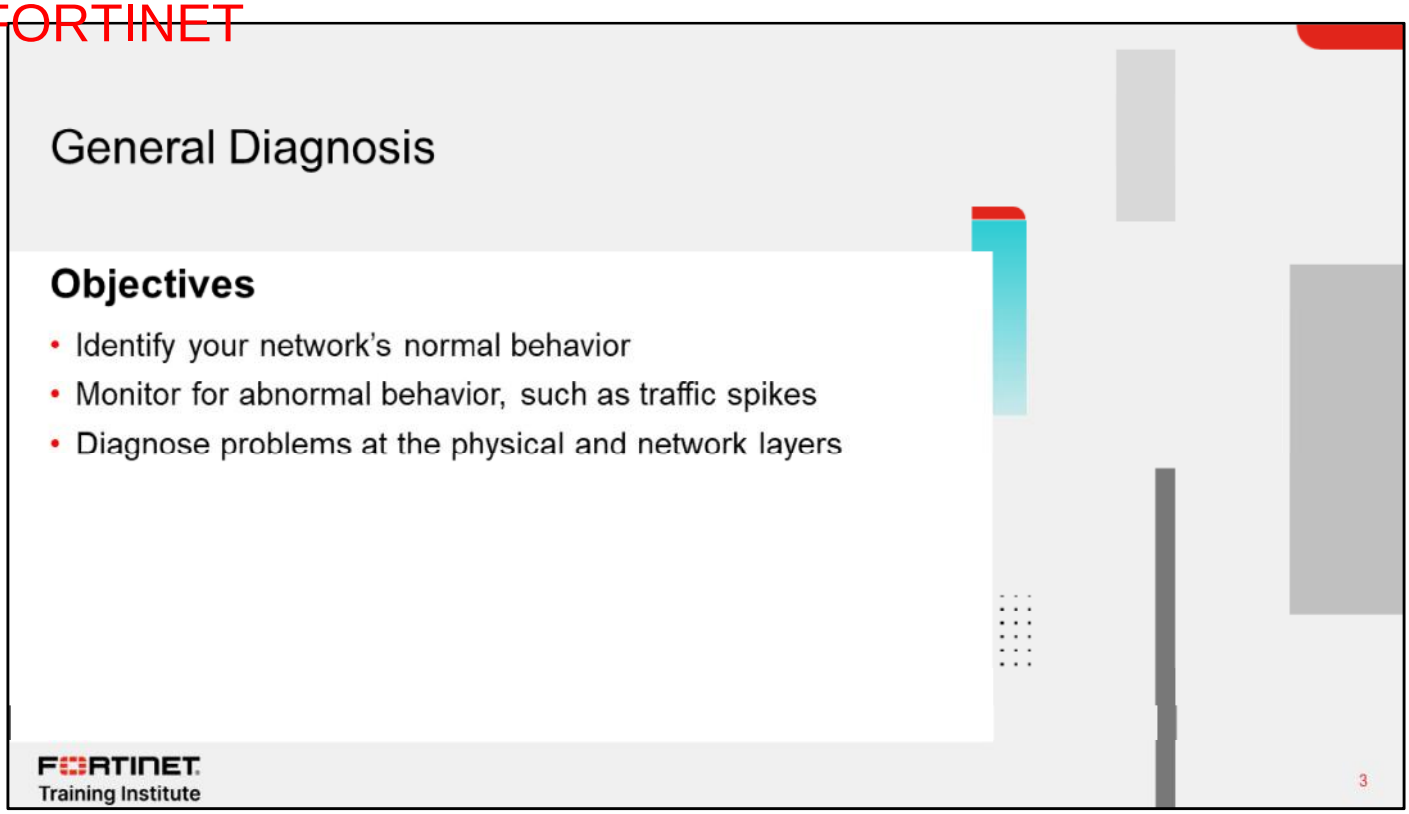

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in general diagnosis, you will be able to discover general information about the status of FortiGate.

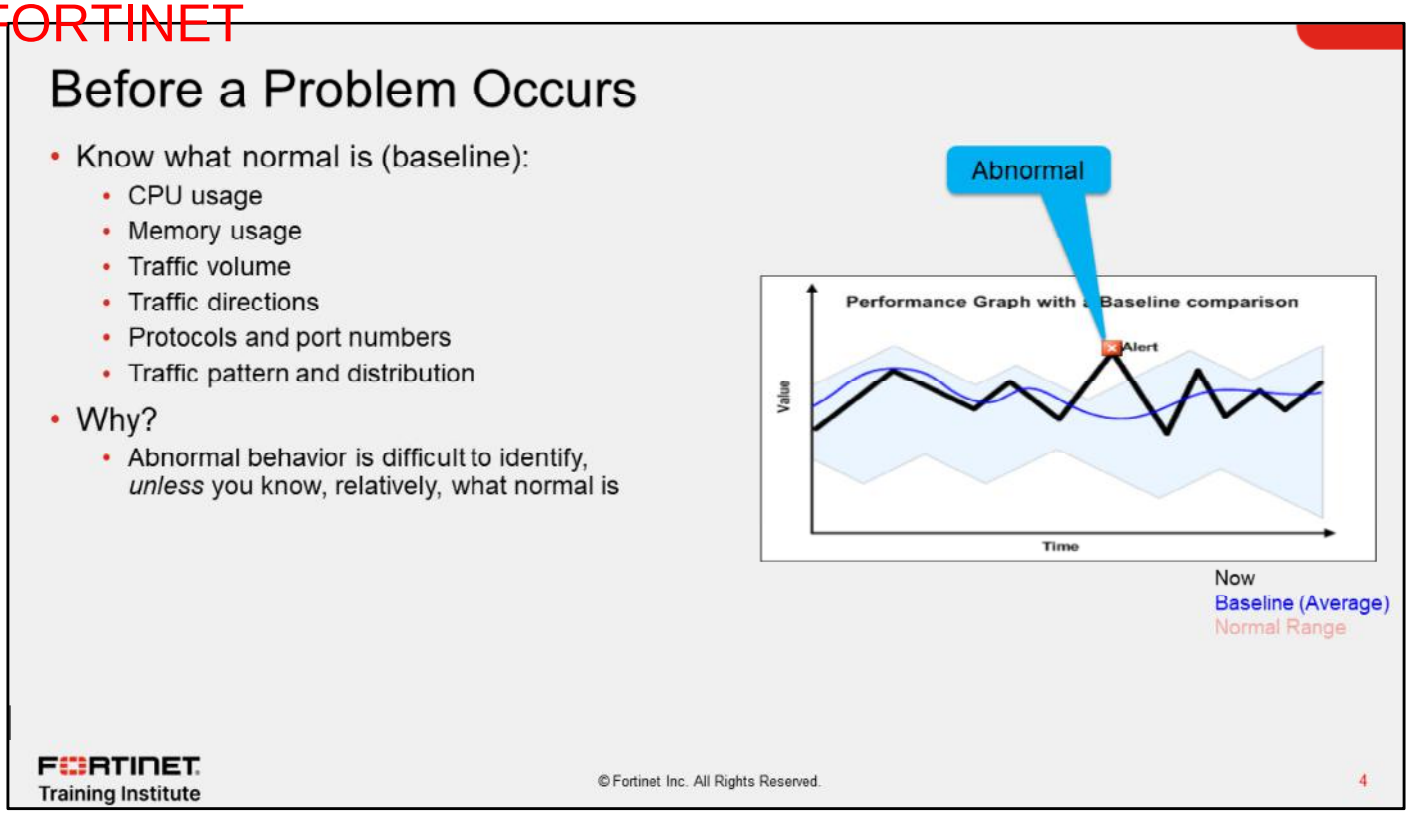

Diagnosis is the process of finding the underlying cause of a problem.

In order to define any problem, first you must know what your network's *normal* behavior is.

In the graph shown on this slide, the range that indicates *normal* is shown in blue. What exactly is this blue line? It indicates the averages—our baseline. What is the thick black line? It's the current behavior. When the current behavior (black line) leaves the normal range, an abnormal event is happening.

Normal is measured and defined in many ways. It can be performance: the expected CPU and memory utilization, bandwidth, and traffic volumes. But, it can also be your network topology: which devices are normally connected at each node. It is also behavior: traffic flow directions, which protocols are blocked or proxied, and the distribution of protocols and applications used during specific times of the day, week, or year.

Network Diagrams

- Why?
  - Explaining or analyzing complex networks is difficult and time-consuming without them
- Physical diagrams:
  - Include cables, ports, and physical network devices
  - Show relationships at Layer 1 and Layer 2
- Logical diagrams:
  - Include subnets, routers, logical devices
  - Show relationships at Layer 3

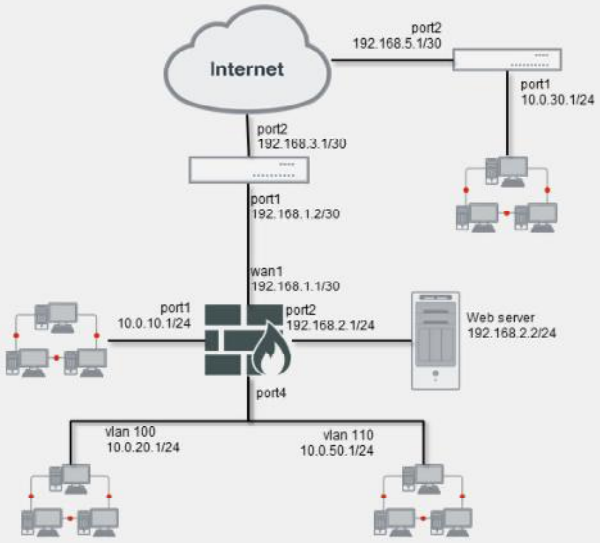

FORTINET Training Institute

5

What is the first way to define what is *normal* for your network?

Flows and other specifications of *normal* behaviour are derived from topology. So, during troubleshooting, a network diagram is essential. If you create a ticket with Fortinet Technical Support, a network diagram should be the first thing you attach.

Network diagrams sometimes combine the two types of diagrams:
- Physical
- Logical

A physical diagram shows how cables, ports, and devices are connected between buildings and cabinets.
A logical diagram shows relationships (usually at OSI Layer 3) between virtual LANs, IP subnets, and routers. It can also show application protocols such as HTTP or DHCP.

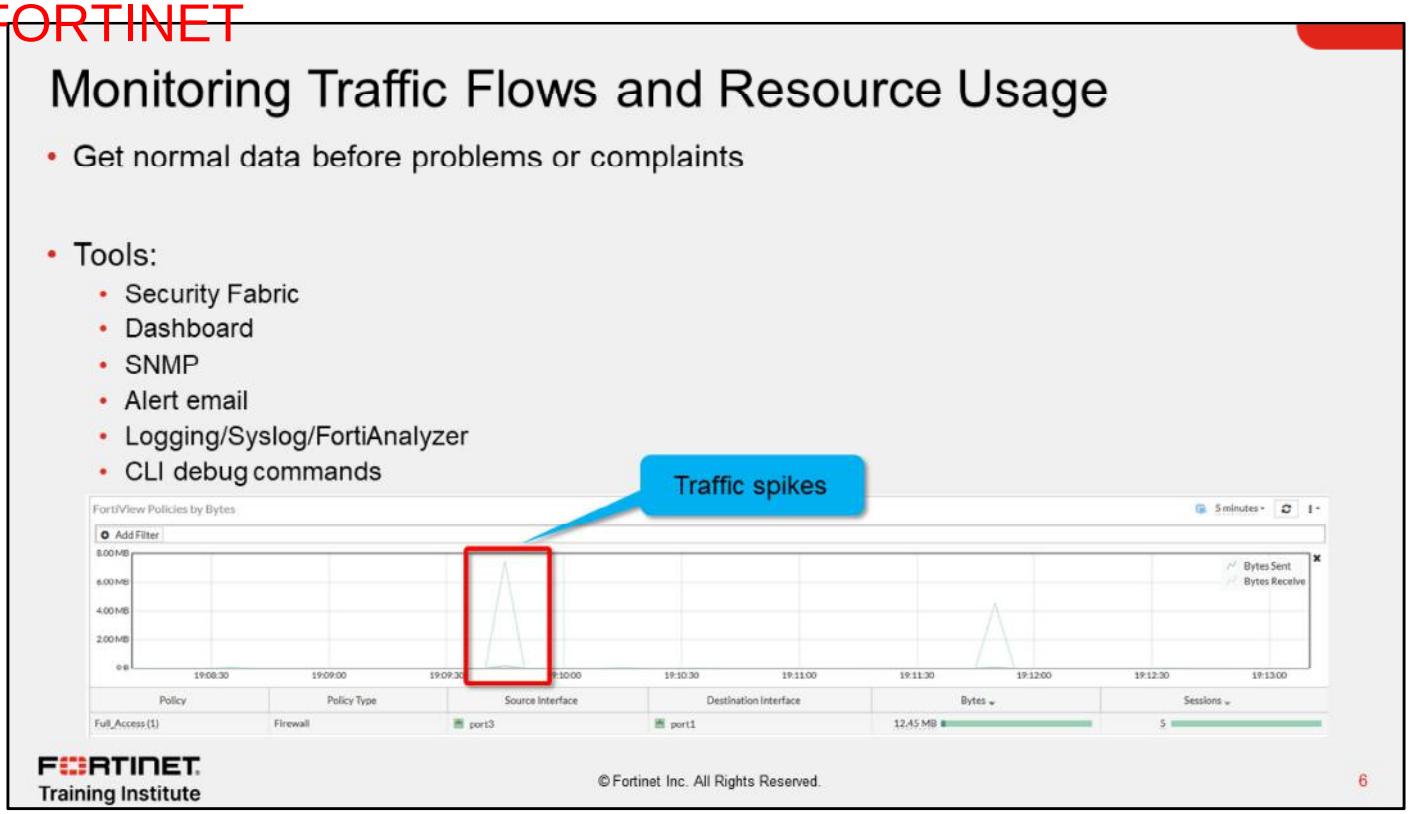

Another way to define normal is to know the average performance range. On an ongoing basis, collect data that shows normal usage.

For example, if traffic processing is suddenly slow, and the FortiGate CPU use is 75%, what does that indicate? If CPU use is usually 60-69%, then 75% is probably still normal. But if normal is 12-15%, there may be a problem.

Get data on both the typical maximum and minimum for the time and date. That is, on a workday or holiday, how many bits per second should ingress or egress each interface in your network diagrams?

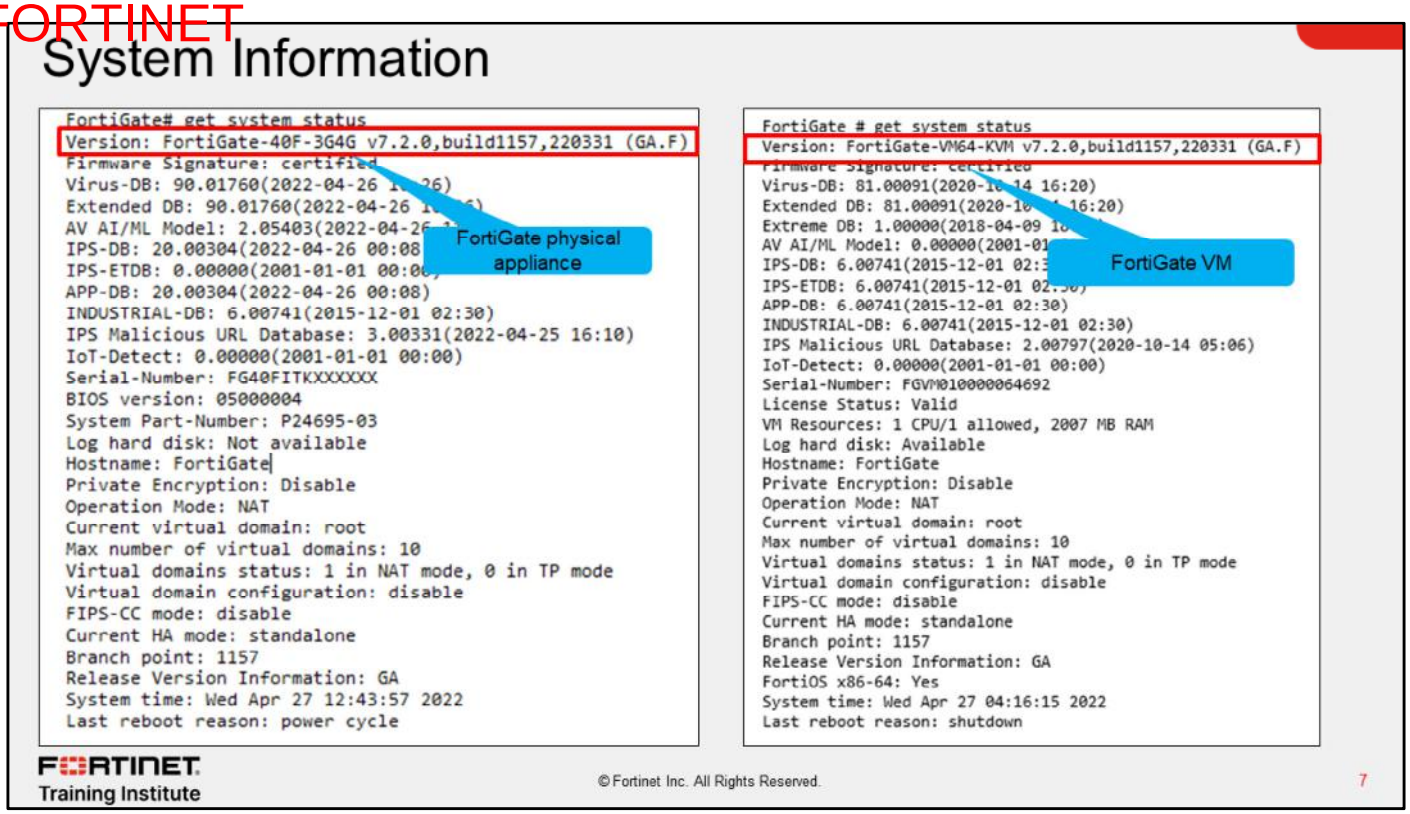

How can we get information about the current status? First, look at CLI commands; you can use them through a local console, even if network issues make GUI access slow or impossible.

A few commands provide system statuses. The `get system status` command provides mostly general-purpose information. The output shows:

- Model
- Serial number
- Firmware version
- Host name
- FortiGuard license status
- System time
- Version of the FortiGuard antivirus, IPS, and IP reputation databases, and others

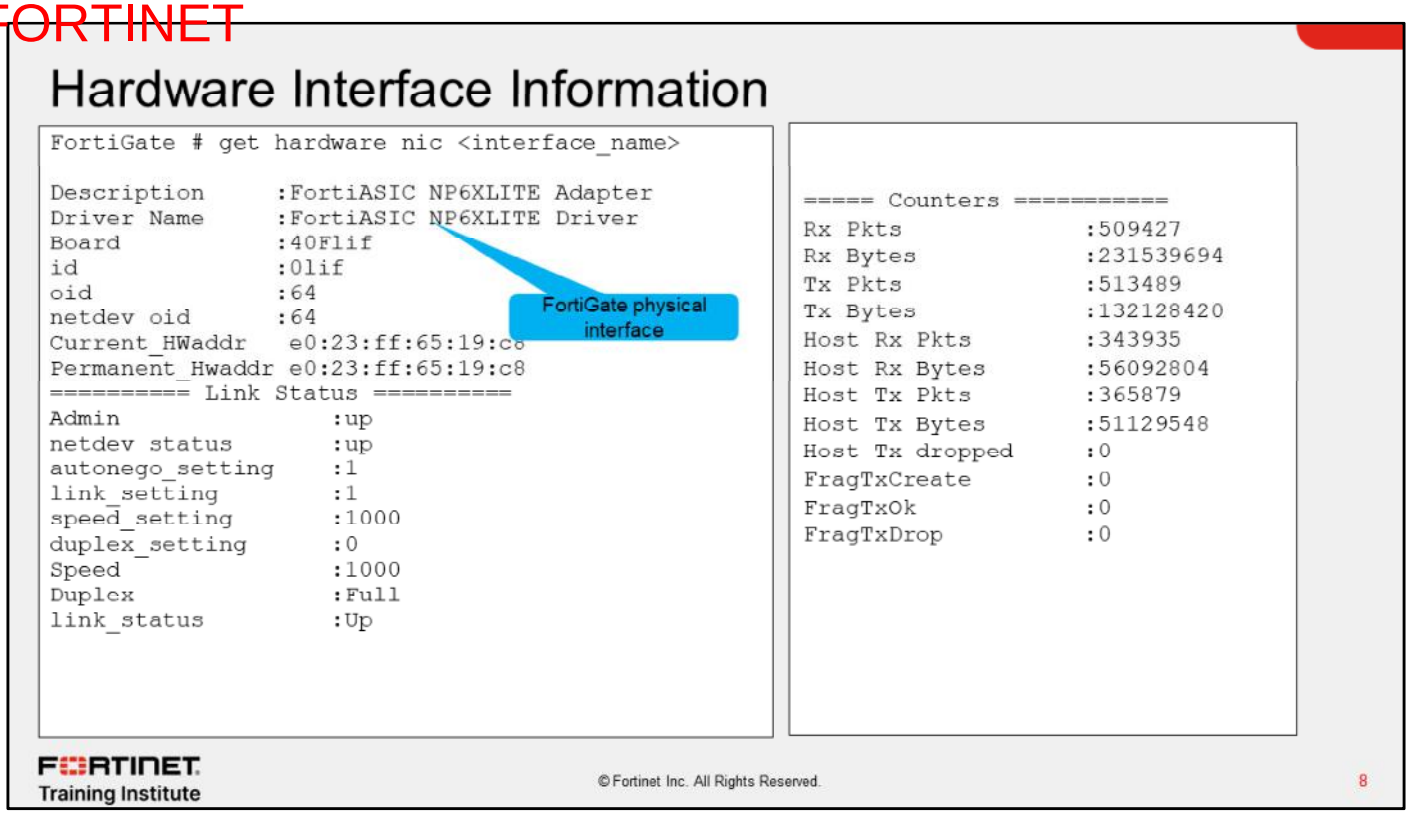

At the physical layer, troubleshooting analyzes which ports are plugged in, media capacity, and negotiated speed and duplex mode.

At the data link layer, diagnostics often analyze how many frames are being dropped because of CRC errors or collisions.

The get hardware nic command is used to display the FortiGate interface hardware and status information. The output might vary depending on the model and NIC driver version.

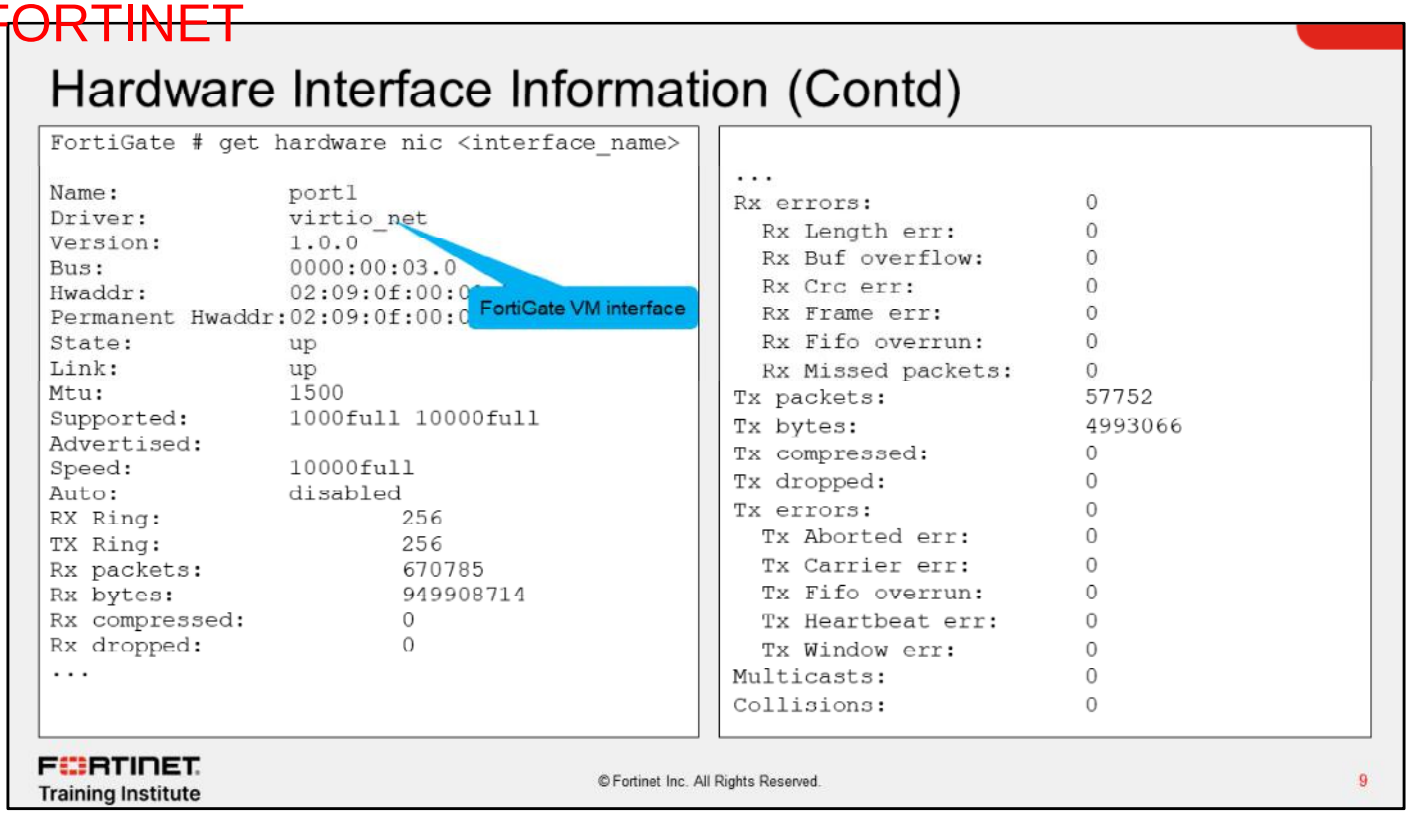

The output on this slide shows the driver name, hardware address, administrative status, and link status, along with send and receive packets and errors.

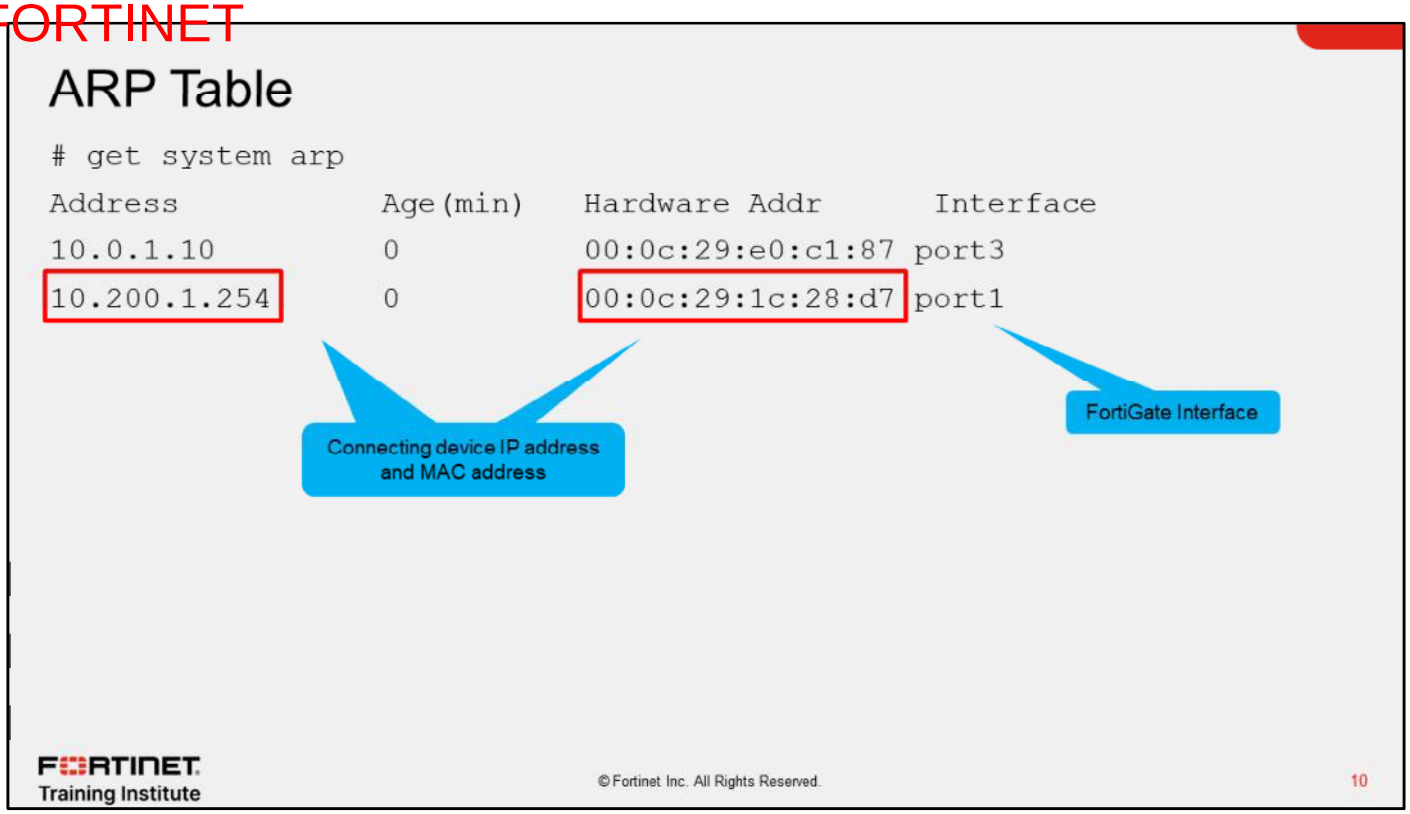

If you suspect that there is an IP address conflict, or that an IP has been assigned to the wrong device, you may need to look at the ARP table. The `get system arp` command is used for that purpose. It shows the FortiGate interface, IP address, and associated MAC address. This command lists the information for all external devices connected to the same LAN segments where FortiGate is connected. The current IP and MAC addresses of FortiGate are not included.

Network Layer Troubleshooting

```
# execute ping-options
adaptive-ping     Adaptive ping <enable|disable>.
data-size         Integer value to specify datagram size in bytes.
df-bit            Set DF bit in IP header <yes | no>.
interface         Auto | <outgoing interface>.
interval          Integer value to specify seconds between two pings.
pattern           Hex format of pattern, e.g. 00ffaabb.
repeat-count      Integer value to specify how many times to repeat PING.
...

# execute ping <ip> IP address or domain name

# execute traceroute <dest> IP address or hostname
```

FORTINET Training Institute

11

Say that FortiGate can contact some hosts through port1, but not others. Is the problem in the physical layer or the link layer? Neither. Connectivity has been proven with at least part of the network. Instead, you should check the network layer. To test this, as usual, start with ping and traceroute.

The same commands exist for IPv6: `execute ping` becomes `execute ping6`, for example.

Remember: location matters. Tests are accurate only if you use the same path as the traffic that you are troubleshooting. To test from FortiGate (to FortiAnalyzer or FortiGuard, for example), use the FortiGate `execute ping` and `execute traceroute` CLI commands. But, to test the path through FortiGate, also use `ping` and `tracert` or `traceroute` from the endpoint—from the Windows, Linux, or Mac OS X computer—not only from the FortiGate CLI.

Because of NAT and routing, you might need to specify a different ping source IP address—the default address is the IP of the outgoing interface. If there is no response, verify that the target is configured to reply to ICMP echo requests.

Knowledge Check

1. Which CLI command can be used to determine the MAC address of a FortiGate default gateway?

✓A. `get system arp`
B. `get hardware nic`

2. Which CLI command can be used to diagnose a physical layer problem?

A. `execute traceroute`
✓B. `get hardware nic`

FORTINET Training Institute

12

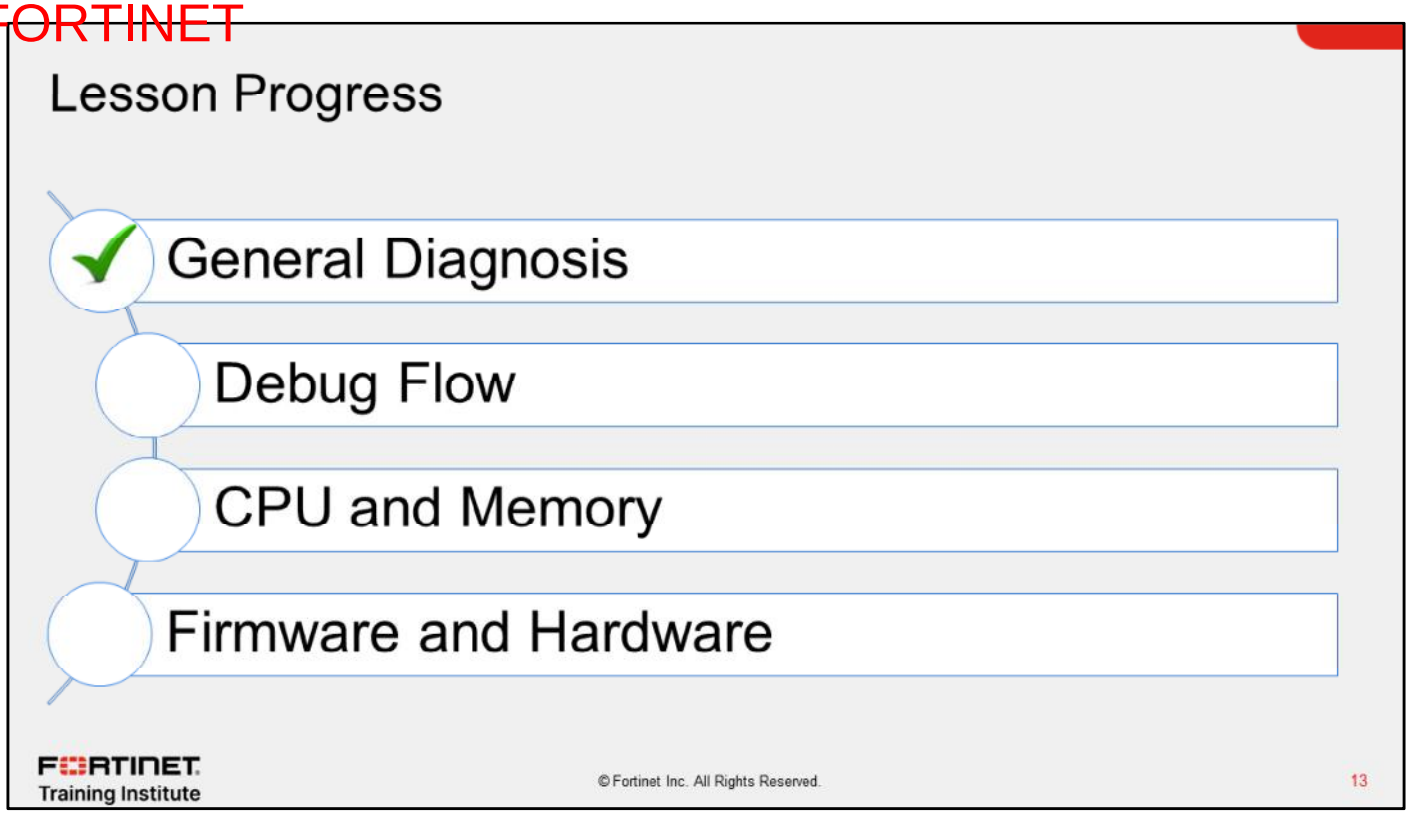

Good job! You now understand general diagnostics.

Now, you will learn about debug flow.

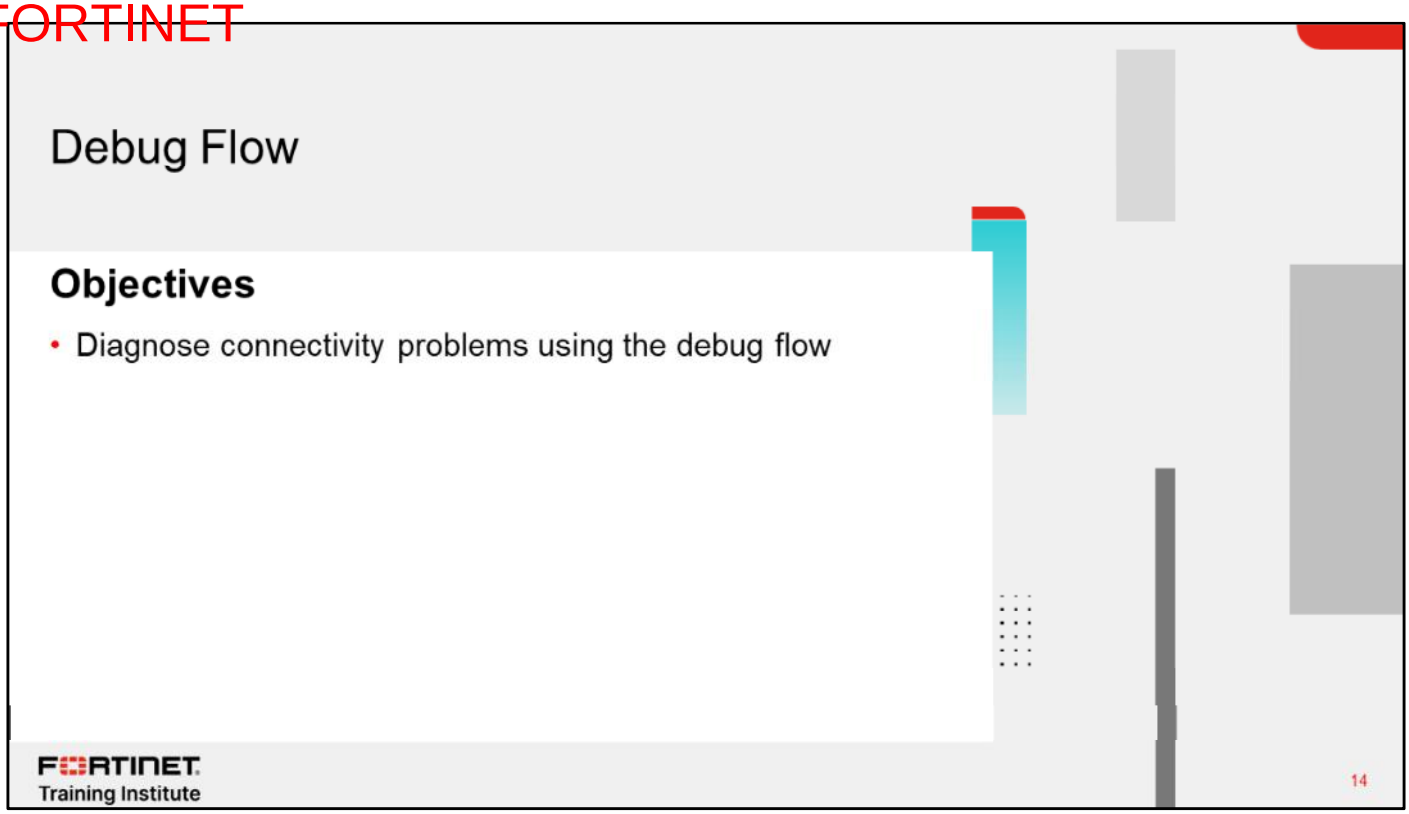

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the debug flow, you will be able to diagnose connectivity problems.

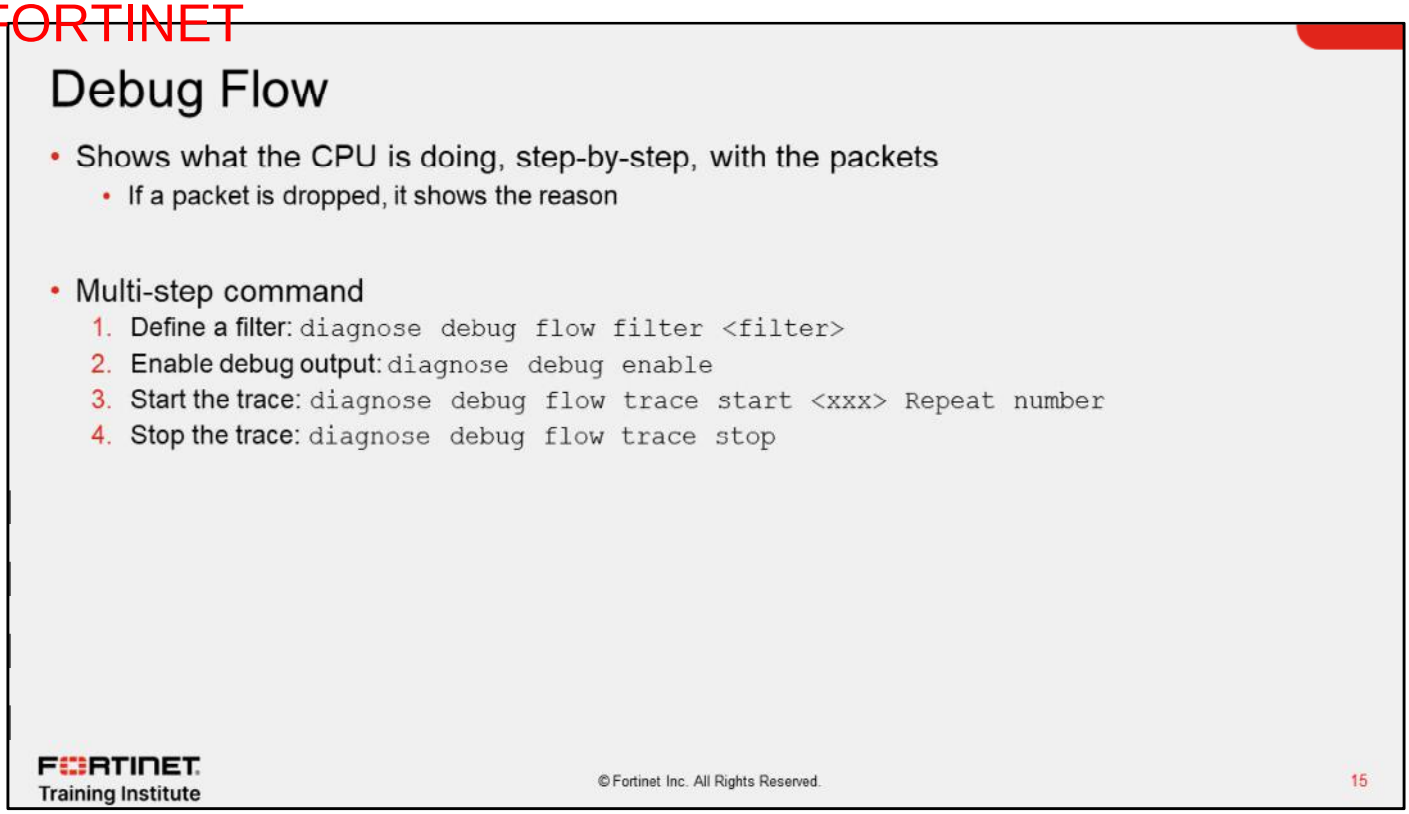

If FortiGate is dropping packets, can a packet capture (sniffer) be used to identify the reason? To find the cause, you should use the debug (packet) flow.

The debug flow shows, step-by-step, how the CPU is handling each packet.

To use the debug flow, follow these steps:
1. Define a filter.
2. Enable debug output.
3. Start the trace.
4. Stop the trace when it`s finished.

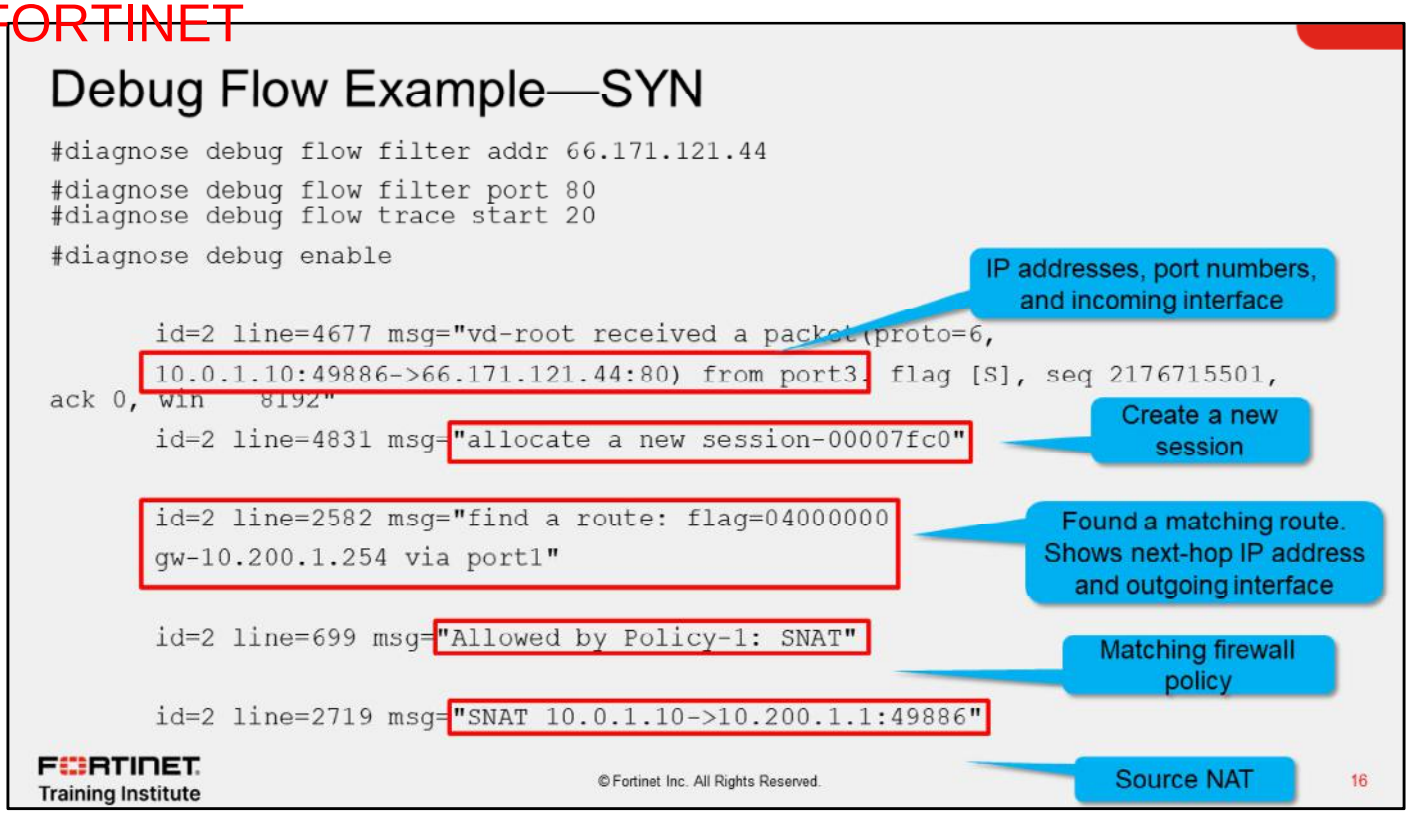

This slide shows an example of a debug flow output of the above `diagnose debug flow` commands,which captures the first packet of a TCP three-way handshake, the SYN packet. It shows:

- The packet arriving at FortiGate, indicating the source and destination IP addresses, port numbers, and incoming interface
- FortiGate creating a session, indicating the session ID
- The route to the destination, indicating the next-hop IP address and outgoing interface
- The ID of the policy that matches and allows this traffic
- How the source NAT is applied

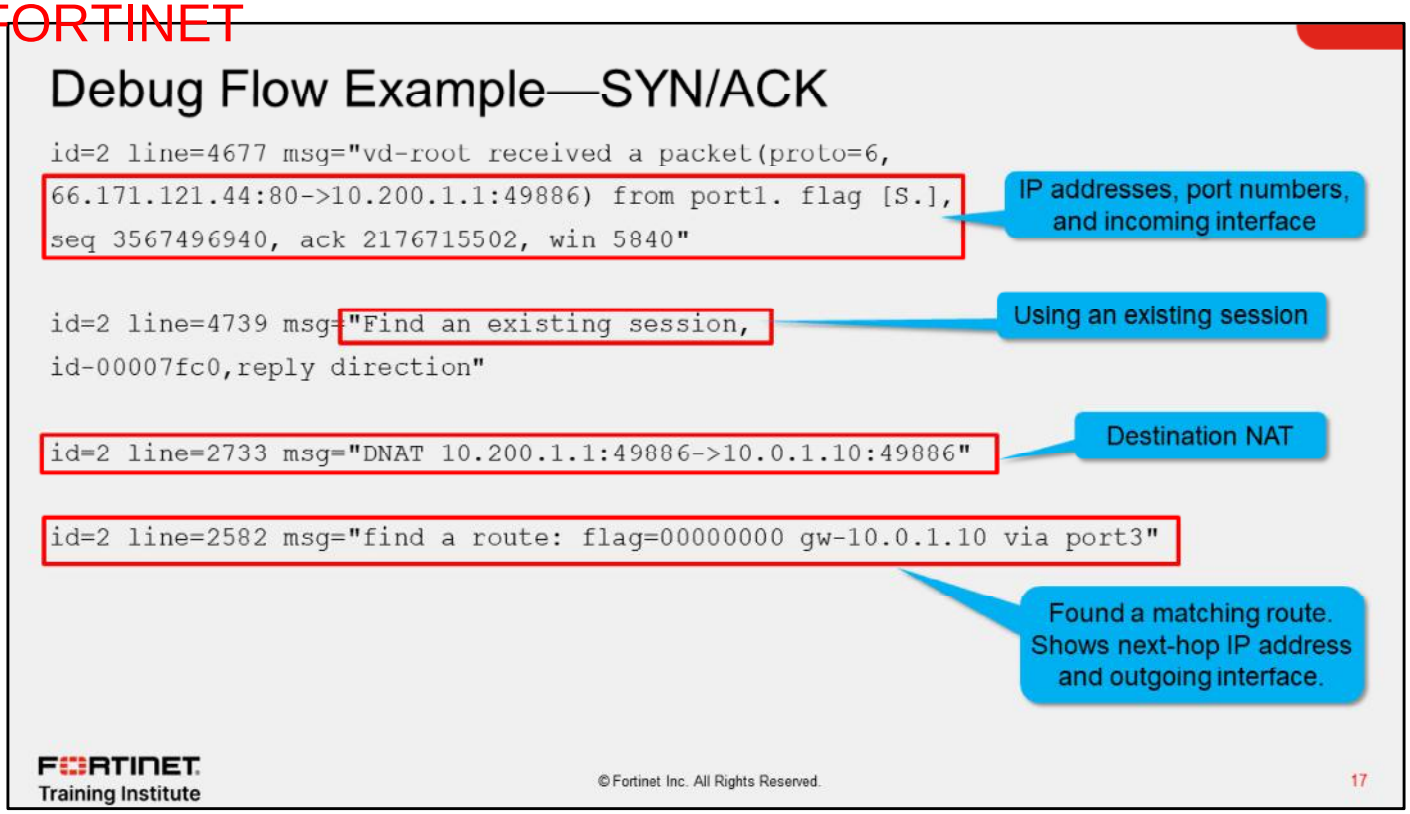

This slide shows the output for the SYN/ACK packet, which is from the same `diagnose debug` command shown on the previous slide. It shows:

- The packet arrival, indicating again the source and destination IP addresses, port numbers, and incoming interface
- The ID of the existing session for this traffic. This number matches the ID of the session created during the SYN packet. The ID is unique for each session, and useful to trace the request/reply packets of the session.
- How the destination NAT is applied
- The route to the destination, indicating again the next-hop IP address and outgoing interface.

If the packet is dropped by FortiGate, this debug shows the reason for that action.

This tool is useful for many other troubleshooting cases, including when you need to understand why a packet is taking a specific route, or why a specific NAT IP address is being applied.

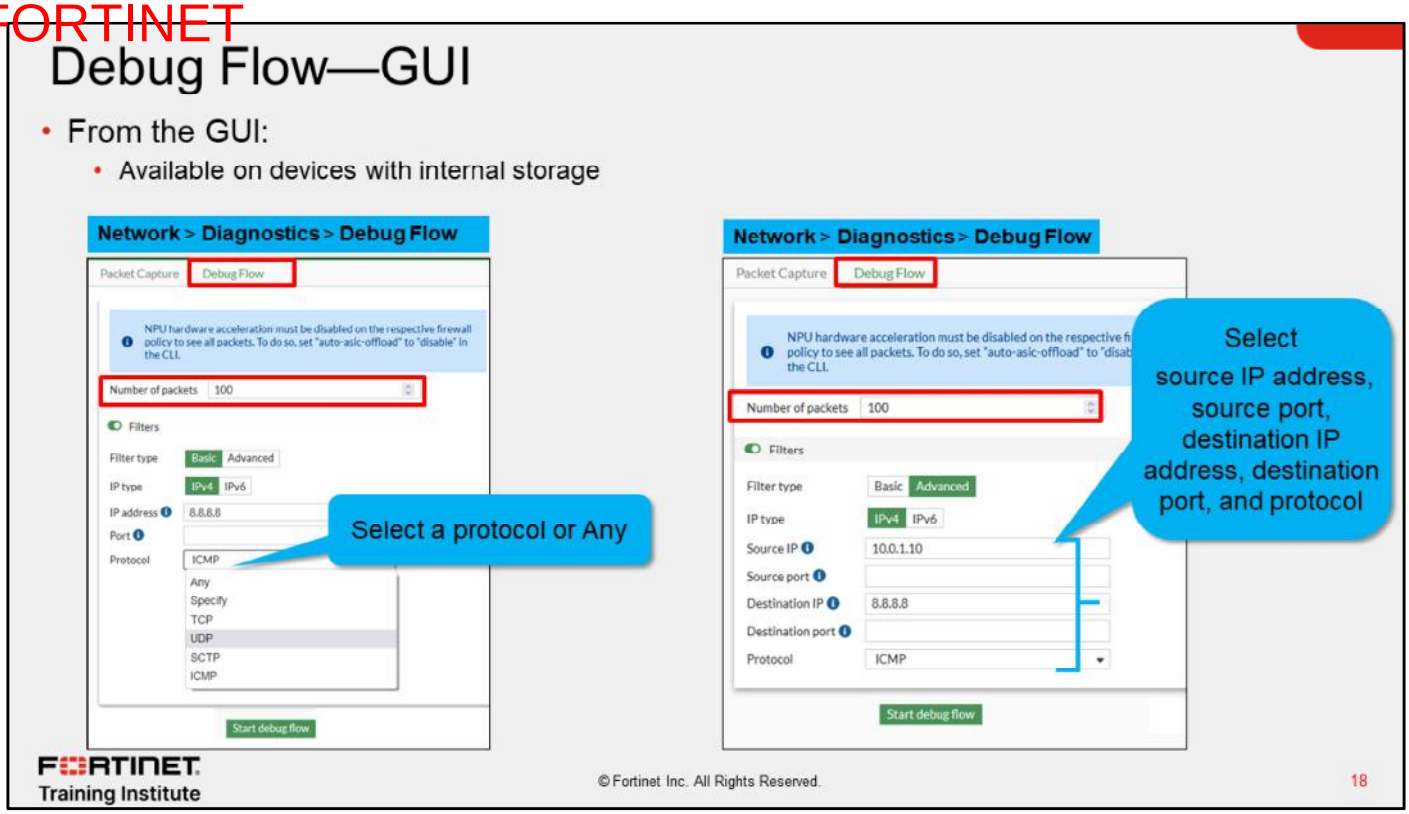

The Debug Flow tool allows you to view debug flow output on the GUI in real time until you stop the debug process.

This tool helps you to examine the packet flow details directly on the GUI.

After you stop the debug flow, you can view the completed output, and filter it by time, message, or function. You can also export the output as a CSV file.

You can set up the Debug Flow tool to use either Basic or Advanced filter options. **Basic** allows you to filter using basic criteria such as host address, port number, and protocol name. **Advanced** allows you to filter by source IP address, source port, destination IP address, destination port, and protocol.

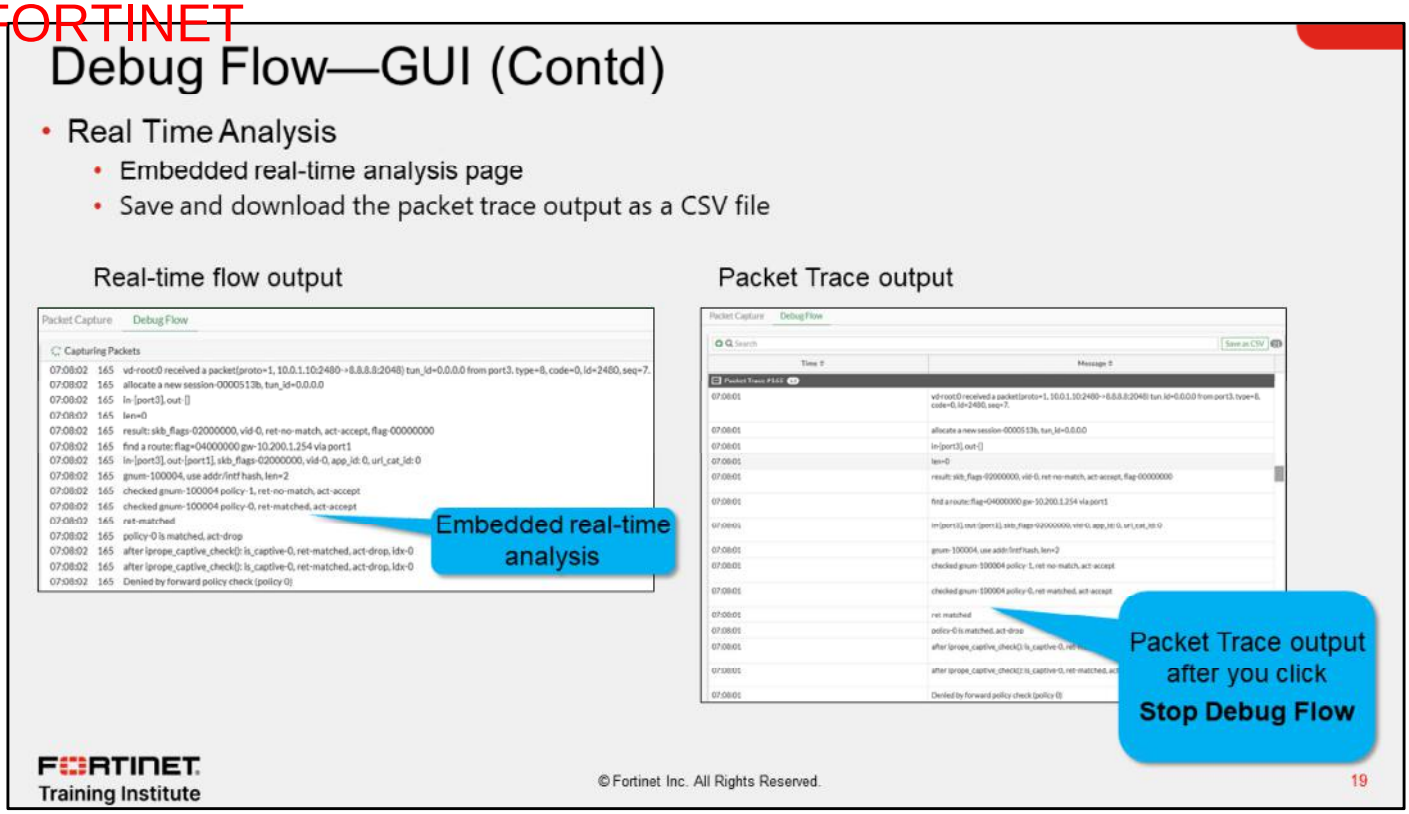

After you start the debug flow, the GUI starts displaying the captured packets based on the filter.

When you stop the debug flow, FortiGate displays a Packet Trace output that you can download and save as a CSV file.

The main difference between these two outputs is that real-time messages are displayed for real-time analysis, but you can save the packet trace outputs and download them for future reference.

Knowledge Check

1. Which information is displayed in the output of a debug flow?
   - ✓A. Incoming interface and matching firewall policy
   - B. Matching security profile and traffic log

2. When is a new TCP session allocated?
   - ✓A. When a SYN packet is allowed
   - B. When a SYN/ACK packet is allowed

FORTINET Training Institute

20

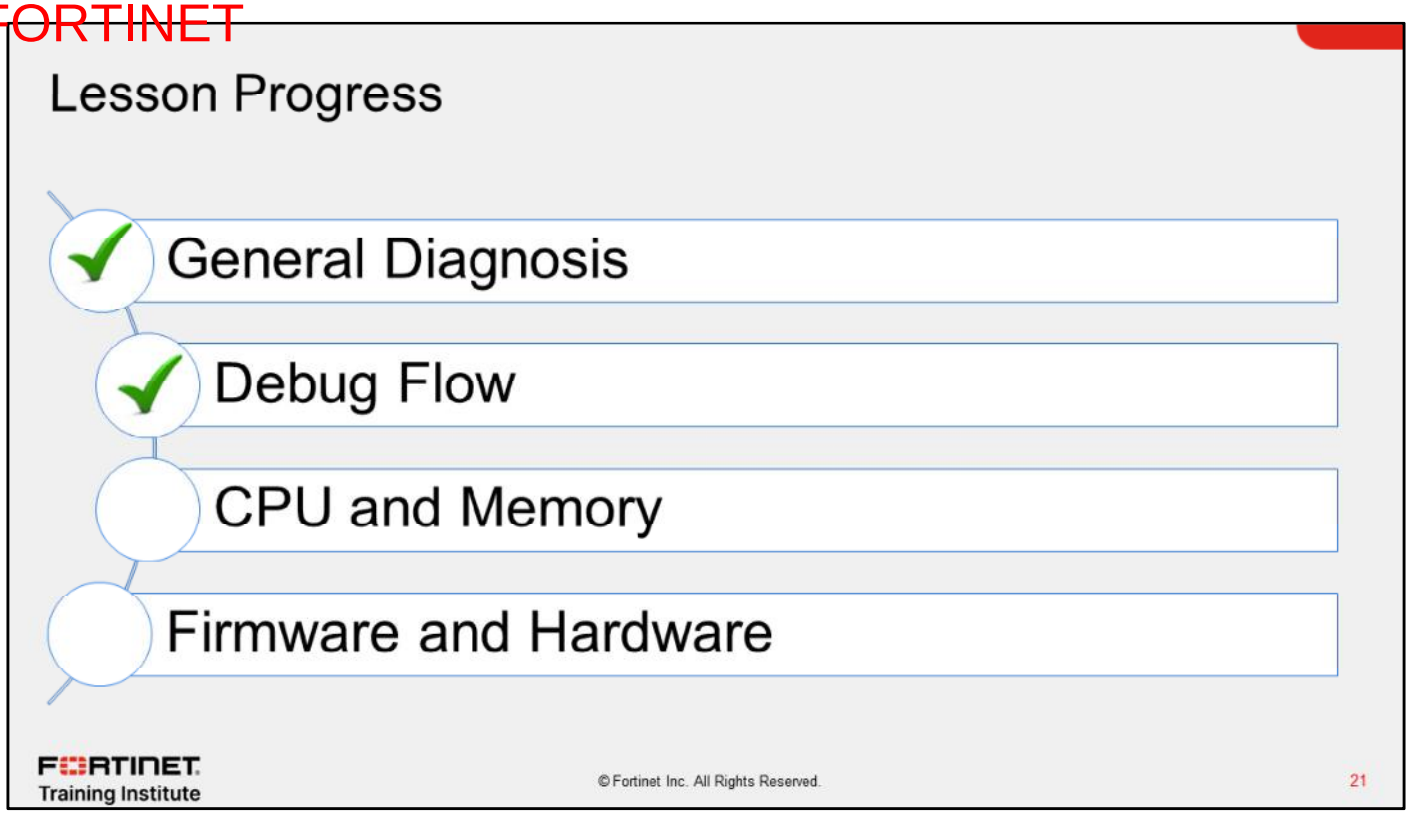

Good job! You now understand debug flow.

Now, you will learn about FortiGate CPU and memory diagnosis.

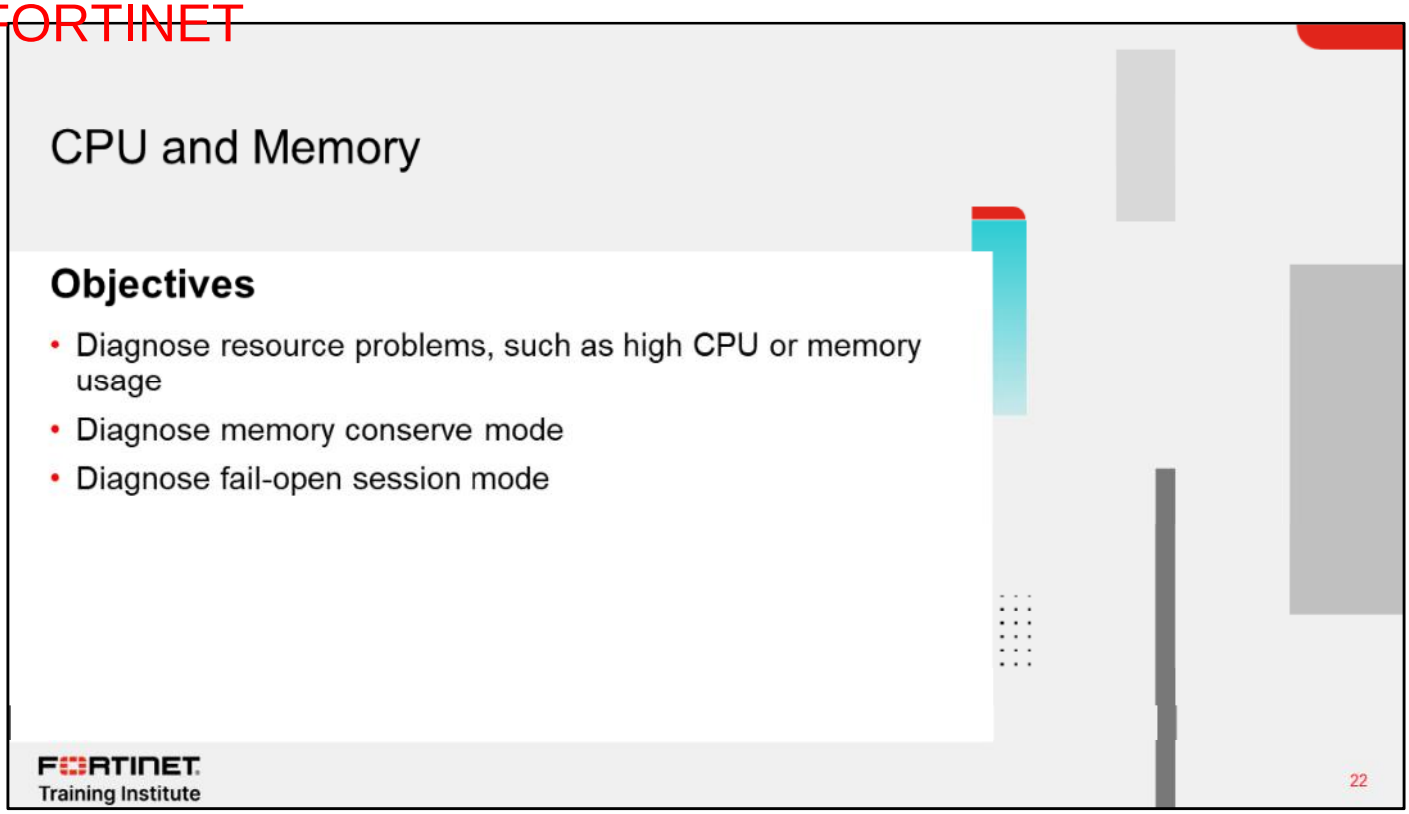

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in CPU and memory, you will be able to diagnose the most common CPU and memory problems.

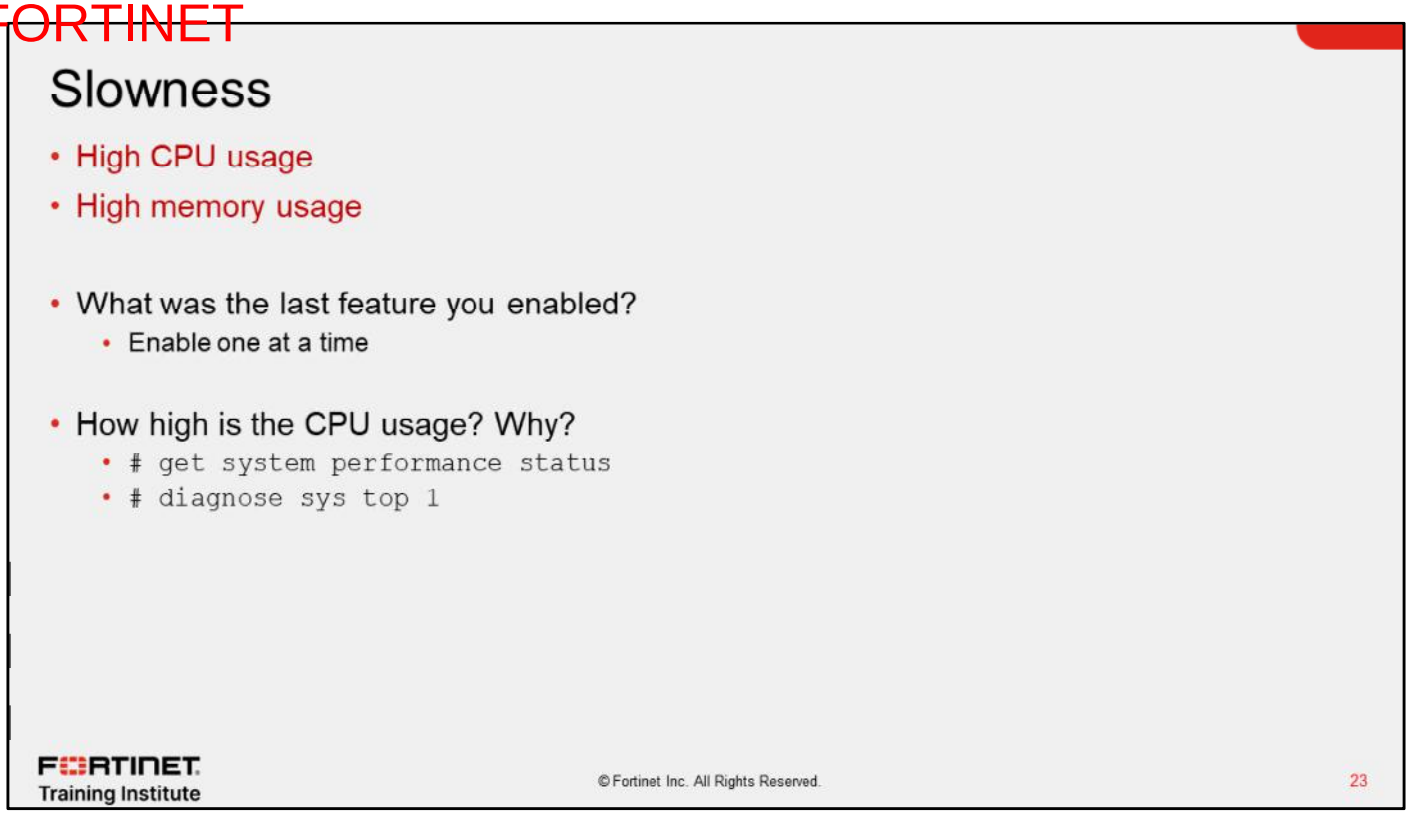

Not all problems are network connectivity failures. Sometimes, there are resource problems in the devices.

What else could cause latency? After you have eliminated problems with the physical media and bandwidth usage, you should check the FortiGate resources usage: CPU and memory.

If usage is high, there are tools that can identify which feature is consuming the most CPU. Additionally, you can troubleshoot faster if you know precisely which change (if any) corresponds with the time the problem began.

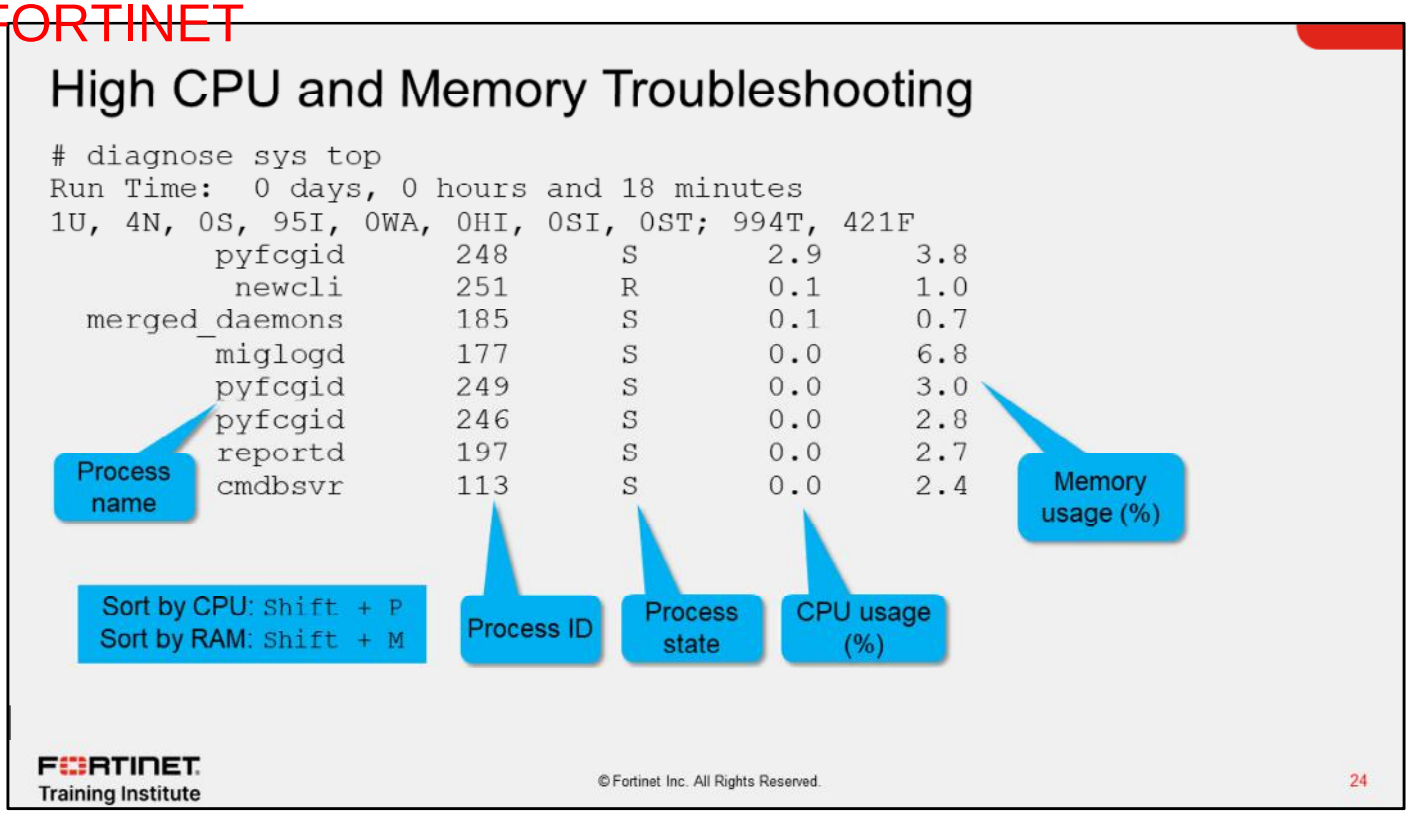

Next, examine the output for `diagnose sys top`. It lists processes that use the most CPU or memory. Some common processes include:

- `ipsengine`, `scanunitd`, and other inspection processes
- `reportd`
- `fgfmd` for FortiGuard and FortiManager connections
- `forticron` for scheduling
- Management processes (`newcli`, `miglogd`, `cmdb`, `sshd`, and `httpsd`)

To sort the list by highest CPU usage, press `Shift+P`. To sort by highest RAM usage, press `Shift+M`.

Memory Conserve Mode

- FortiOS protects itself when memory usage is high
  - It prevents using so much memory that FortiGate becomes unresponsive
- Three configurable thresholds:

Threshold	Definition	Default (% of total RAM)
Green	Threshold at which FortiGate exits conserve mode	82%
Red	Threshold at which FortiGate enters conserve mode	88%
Extreme	Threshold at which new sessions are dropped	95%

```
config system global
    set memory-use-threshold-red <percentage>
    set memory-use-threshold-extreme <percentage>
    set memory-use-threshold-green <percentage>
end
```

FORTINET Training Institute

25

If memory usage becomes too high, FortiGate may enter into memory conserve mode. While FortiGate is in memory conserve mode, it must take action to prevent memory usage from increasing, which could cause the system to become unstable and inaccessible.

Memory conserve mode is never a desirable state because it impacts the user traffic.

Three different configurable thresholds define when FortiGate enters and exits conserve mode. If memory usage goes above the percentage of total RAM defined as the red threshold, FortiGate enters conserve mode. The actions that the device takes depend on the device configuration.

If memory usage keeps increasing, it might exceed the extreme threshold. While memory usage is above this highest threshold, all new sessions are dropped.

The third configuration setting is the green threshold. If memory usage goes below this threshold, FortiGate exits conserve mode.

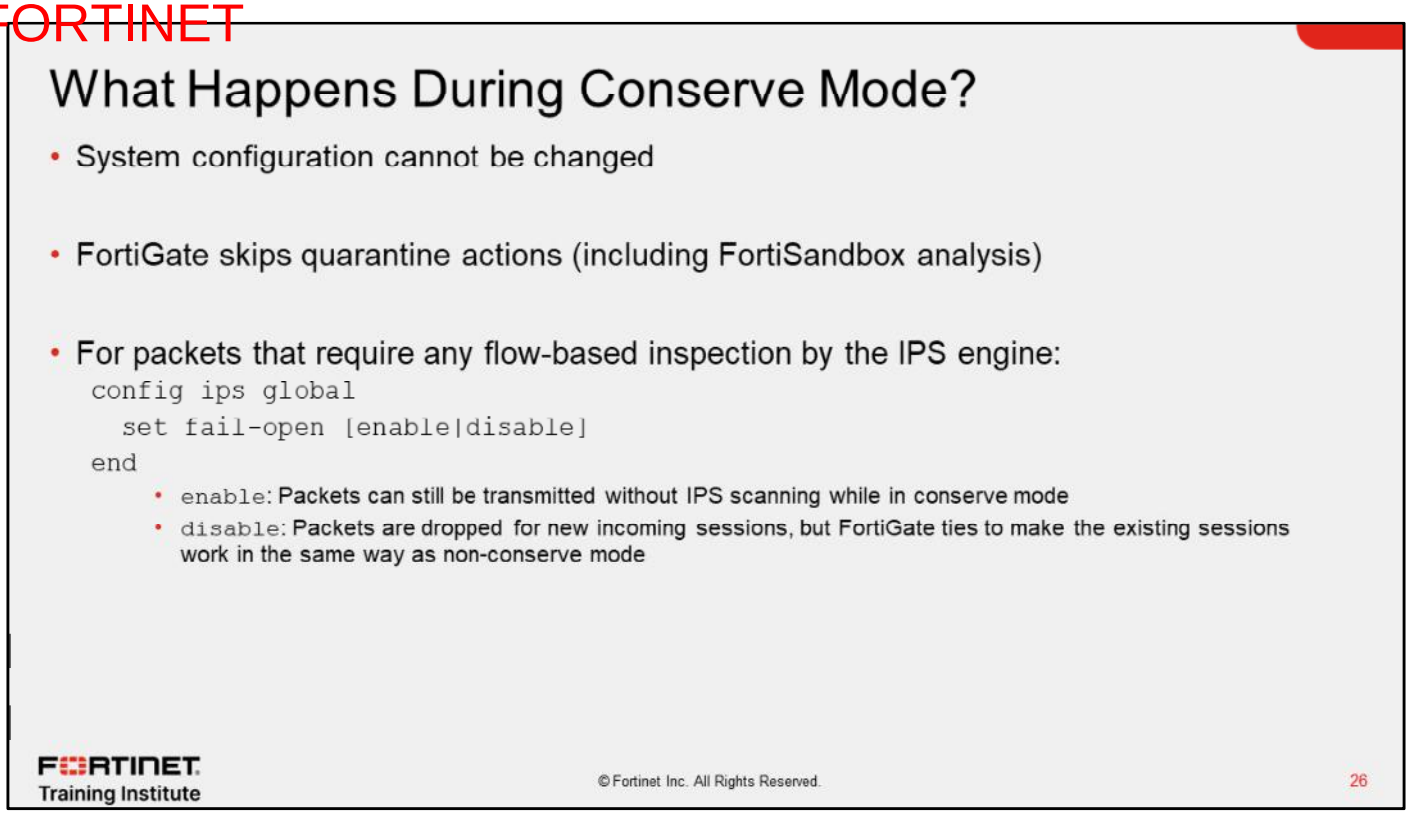

What actions does FortiGate take to preserve memory while in conserve mode?

- FortiGate does not accept configuration changes, because they might increase memory usage.
- FortiGate does not run any quarantine action, including forwarding suspicious files to FortiSandbox.
- You can configure the `fail-open` setting under `config ips global` to control how the IPS engine behaves when the IPS socket buffer is full.

If the IPS engine does not have enough memory to build more sessions, the `fail-open` setting determines whether the FortiGate should drop the sessions or bypass the sessions without inspection.

It is important to understand that the IPS `fail-open` setting is not just for conserve mode —it kicks in whenever IPS fails. Most failures are due to a high CPU issue or a high memory (conserve mode) issue. Enable the setting so that packets can still be transmitted while in conserve mode (or during any other IPS failure) but are not inspected by IPS. Disable the setting so that packets are dropped for new incoming sessions, but allow FortiOS to try to make the existing sessions work in the same way as non-conserve mode.

Remember that the IPS engine is used for all types of flow-based inspections. The IPS engine is also used when FortiGate must identify the network application, regardless of the destination TCP/UDP port (for example, for application control). Note that NTurbo doesn't support the `fail-open` setting. If fail open is triggered, new sessions that would typically be accelerated with NTurbo are dropped, even if the `fail-open` setting is enabled.

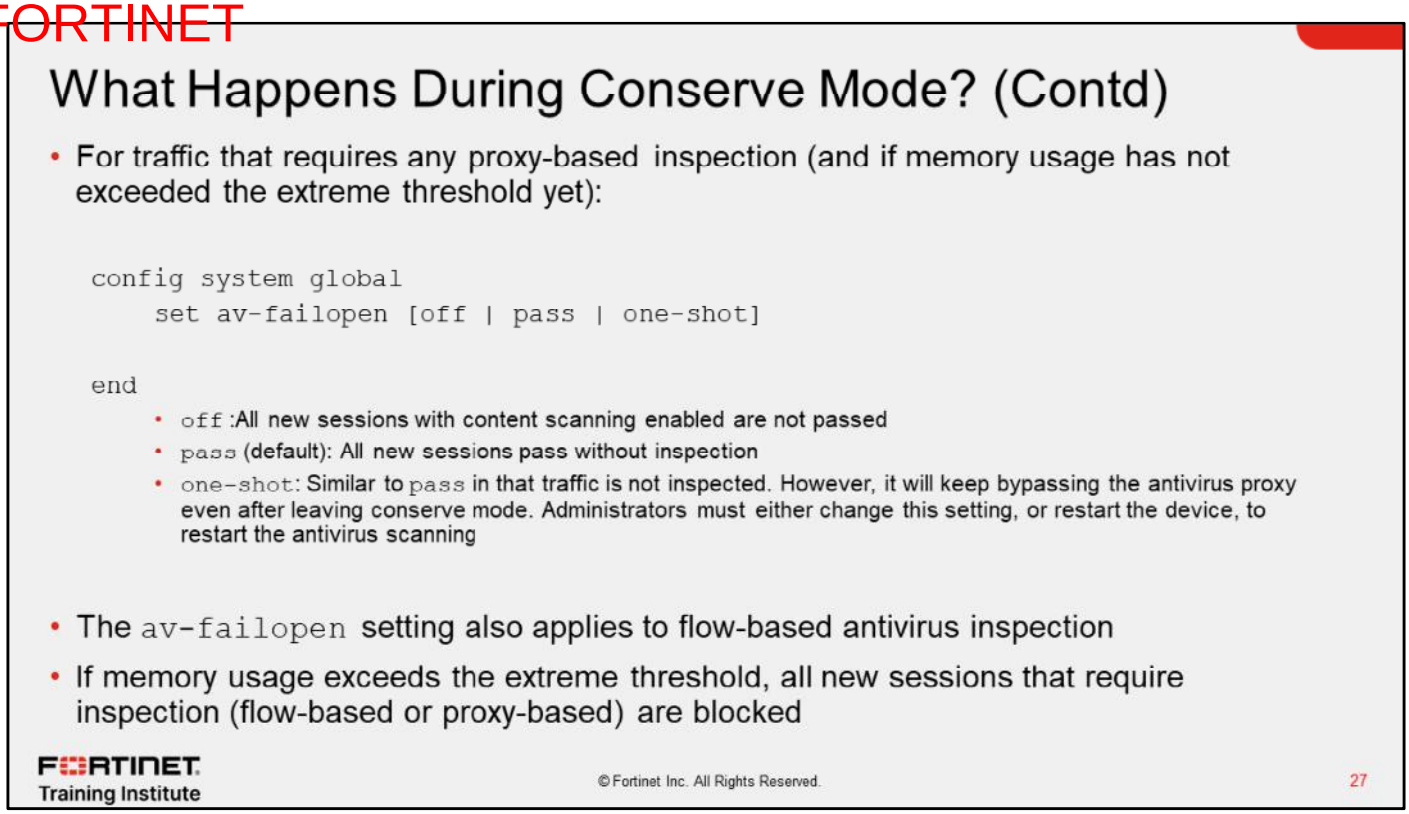

The `av-failopen` setting defines the action that is applied to any proxy-based inspected traffic, while the unit is in conserve mode (and as long as the memory usage does not exceed the extreme threshold). This setting also applies to flow-based antivirus inspection. Three different actions can be configured:

- `off`: All new sessions with content scanning enabled are not passed but FortiGate processes the current active sessions.
- `pass` (default): All new sessions pass without inspection until FortiGate switches back to non-conserve mode.
- `one-shot`: Similar to `pass` in that traffic passes without inspection. However, it will keep bypassing the antivirus proxy even after it leaves conserve mode. Administrators must either change this setting, or restart the unit to restart the antivirus scanning

However, if the memory usage exceeds the extreme threshold, new sessions are always dropped, regardless of the FortiGate configuration.

System Memory Conserve Mode Diagnostics

```
# diagnose hardware sysinfo conserve
memory conserve mode:                          on
total RAM:                                   3040 MB
memory used:                                 2706 MB   89% of total RAM
memory freeable:                              334 MB   11% of total RAM
memory used + freeable threshold extreme:    2887 MB   95% of total RAM
memory used threshold red:                   2675 MB   88% of total RAM
memory used threshold green:                 2492 MB   82% of total RAM
```

Off = no conserve mode
on = conserve mode

FORTINET Training Institute

28

The `diagnose hardware sysinfo conserve` command is used to identify if a FortiGate device is currently in memory conserve mode.

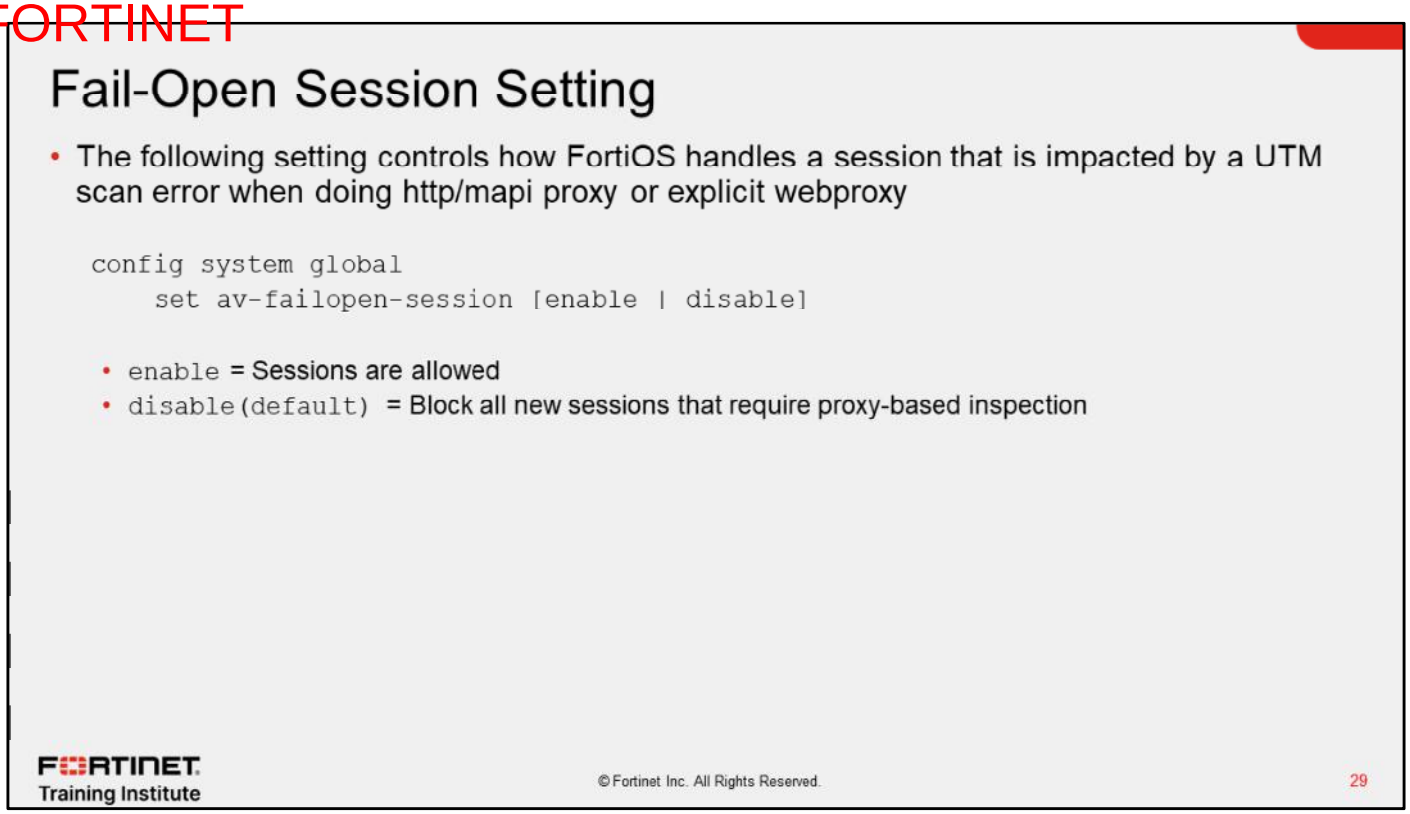

Another undesirable state for FortiGate is the fail-open session mode. This mode kicks in, not during a high-memory situation, but when a proxy on FortiGate runs out of available sockets to process more proxy-based inspected traffic.

If `av-failopen-session` is enabled, FortiGate allows all the sessions. Otherwise, by default, it blocks new sessions that require proxy-based inspection until new sockets become available.

Knowledge Check

1. Which action does FortiGate take during memory conserve mode?
   - ✓ A. Configuration changes are not allowed.
   - B. Administrative access is denied.

2. Which threshold is used to determine when FortiGate enters conserve mode?
   - A. Green
   - ✓ B. Red

FORTINET Training Institute

30

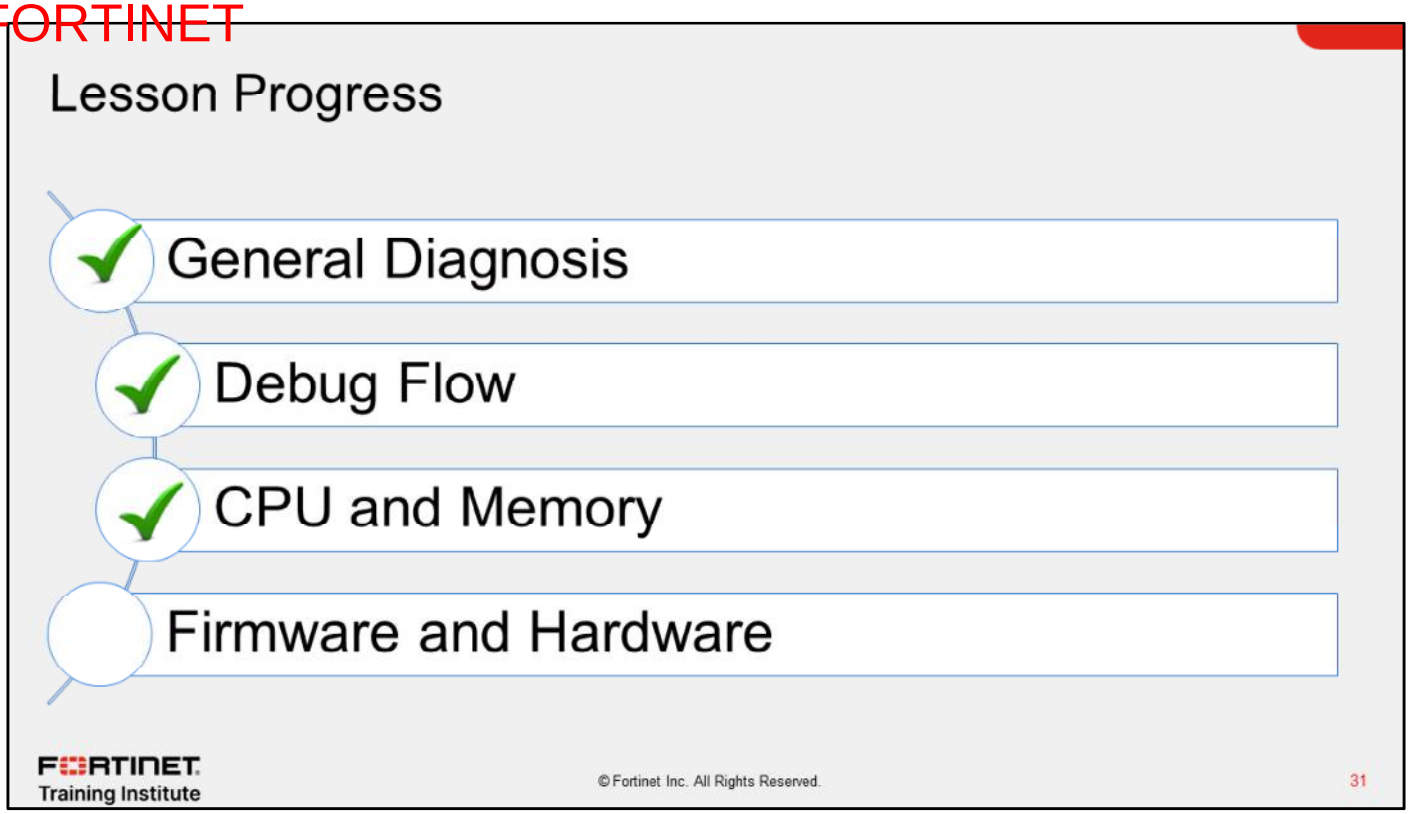

Good job! You now understand FortiGate CPU and memory diagnosis.

Now, you will learn about FortiGate firmware and hardware diagnosis.

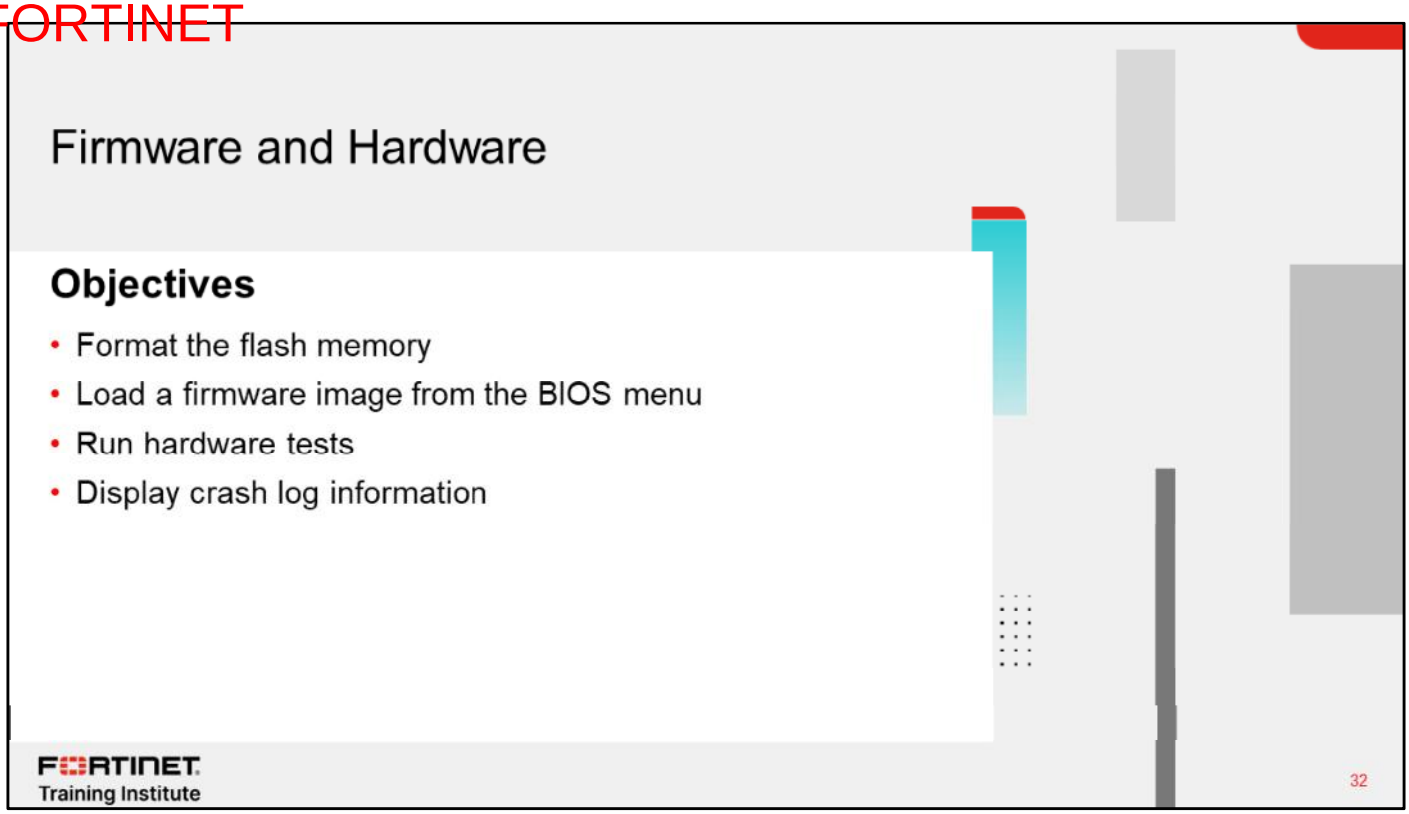

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in firmware and hardware, you will be able to diagnose the most common firmware and hardware problems.

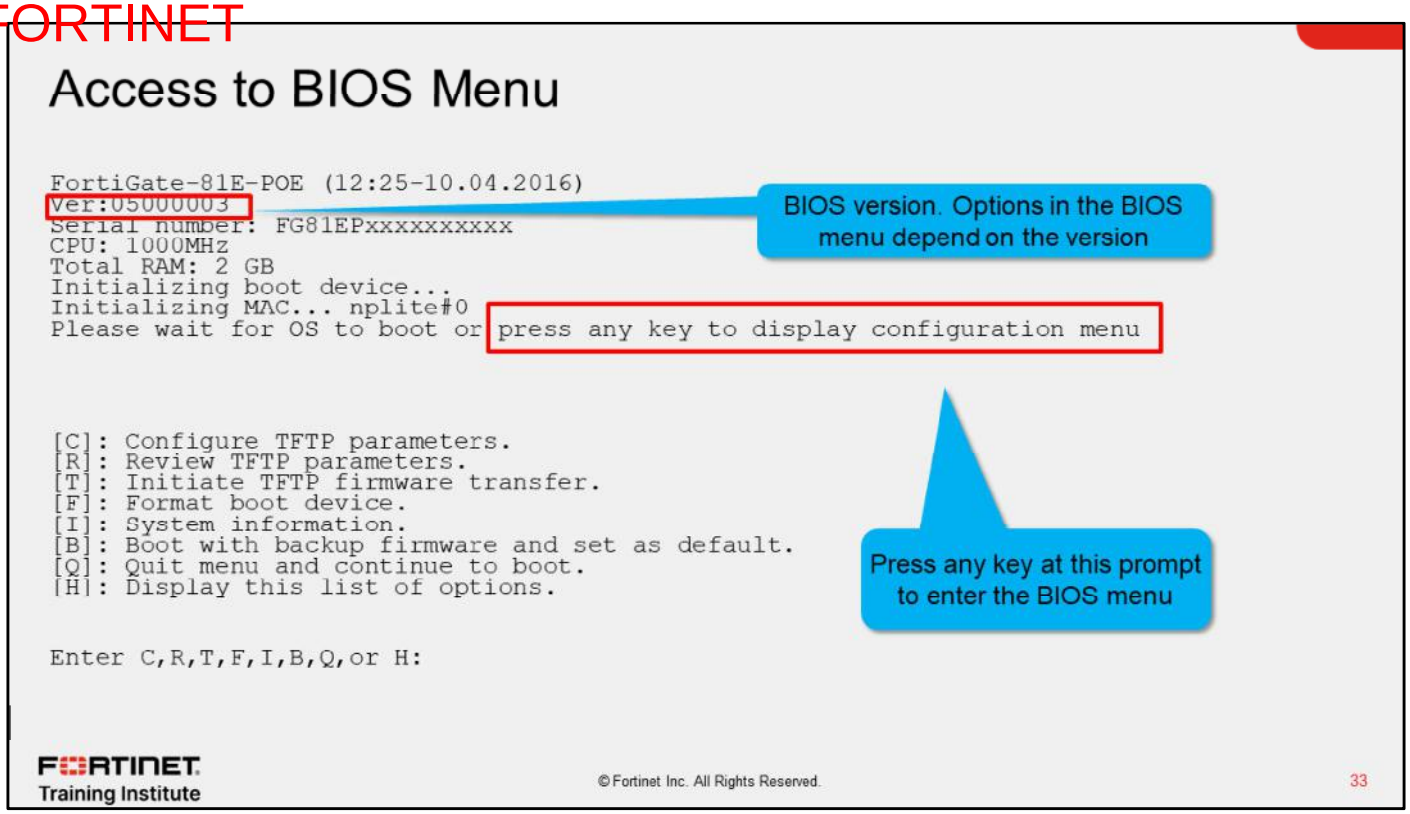

On the FortiGate BIOS, administrators can run some operations over the flash memory and firmware images. To access the BIOS menu, you must reboot the device while connected to the console port. The booting process, at one point, shows the following message:

```
Press any key to display configuration menu
```

While this prompt is displayed, press any key to interrupt the booting process and display the BIOS menu. In the BIOS menu, you can see the options shown on this slide.

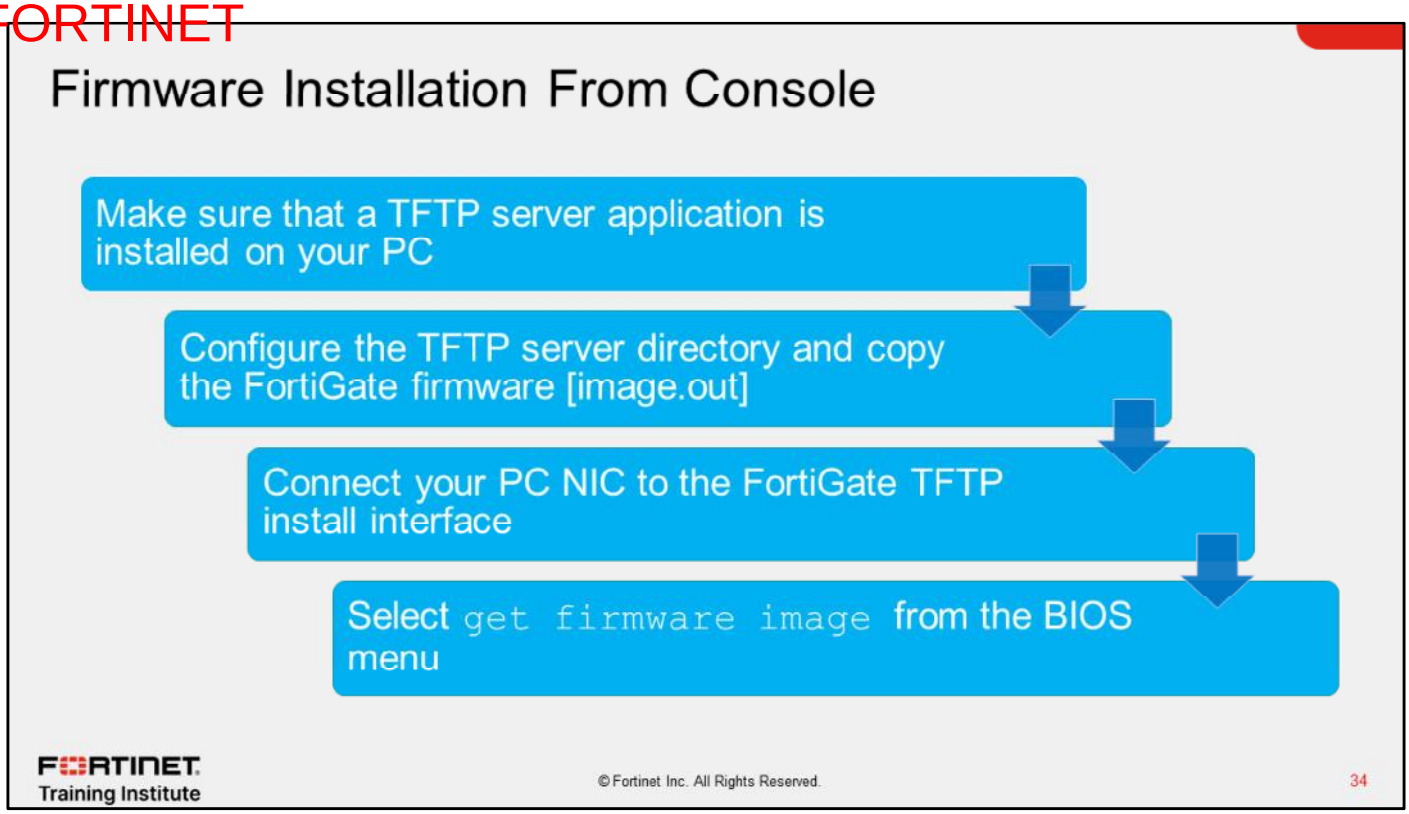

After reformatting the flash memory, you must install the firmware image from the BIOS menu. Follow these steps:

1. Run a TFTP server.
2. Configure the TFTP server with the folder where the firmware image file is stored.
3. Connect the PC Ethernet port to the FortiGate TFTP installation interface.
4. Select `get firmware image` from the BIOS menu.

The interface assigned as the TFTP installation interface depends on the model. However, and in most cases, it is either the *port1* or *internal* interface.

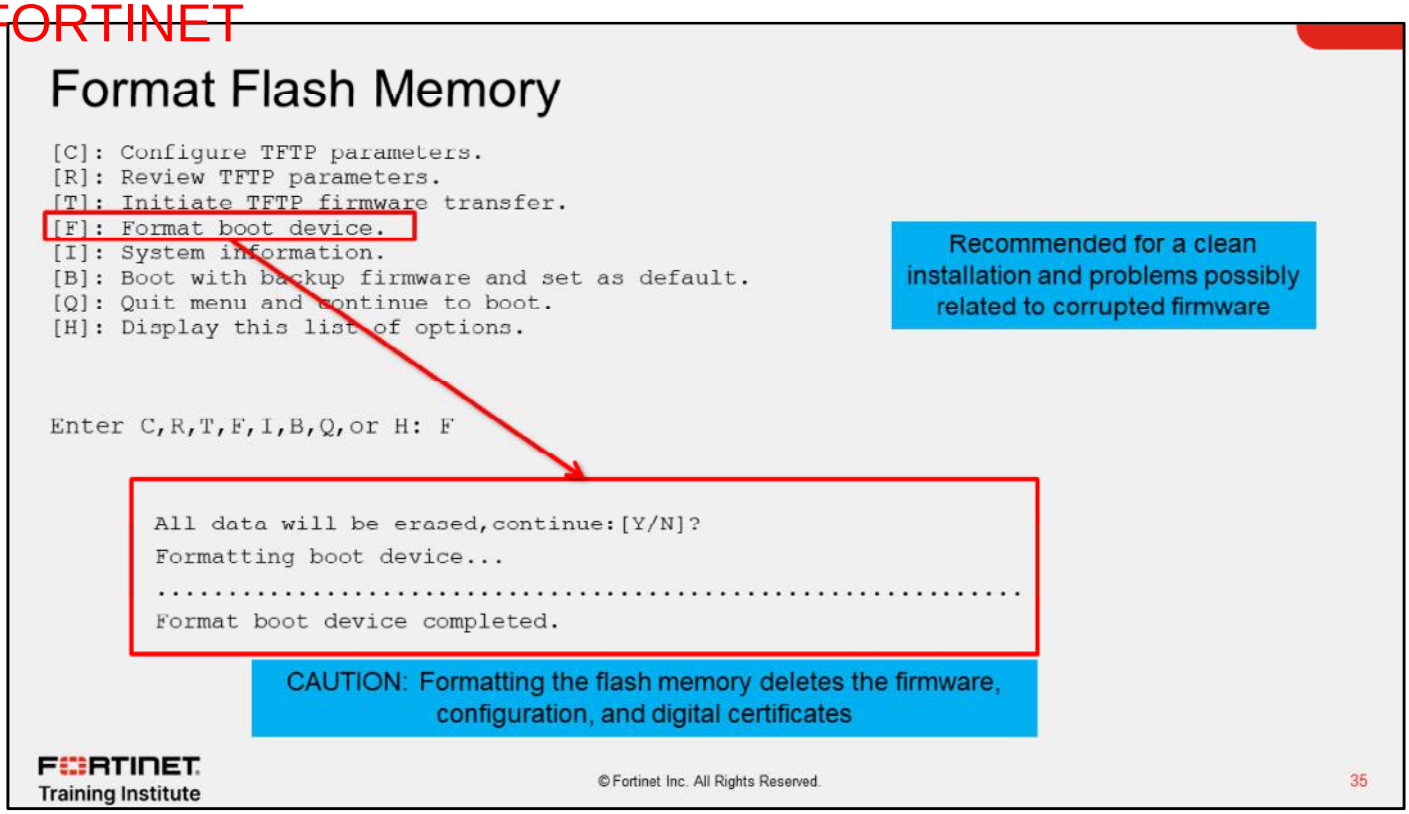

From the BIOS menu, select `F` to format the flash memory.

Doing this might be required if the firmware gets corrupted, or if the administrator wants to do a clean installation of new firmware. Keep in mind, though, that formatting the flash memory deletes any information stored on it, such as firmware images, configuration, and digital certificates.

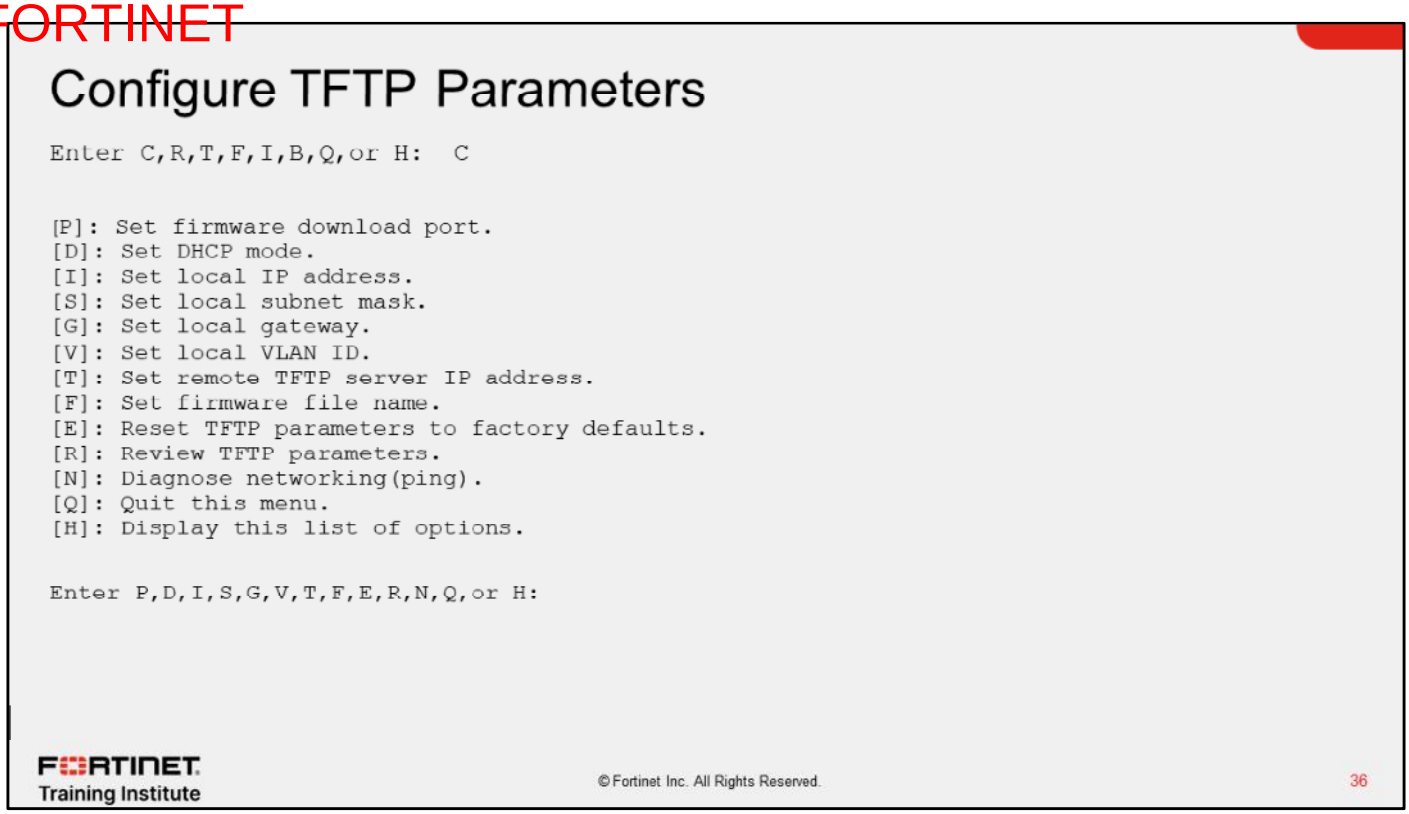

From the BIOS menu, select C to configure TFTP parameters. Use the menu options to configure parameters, such as local IP address, subnet mask , gateway address, and firmware file name.

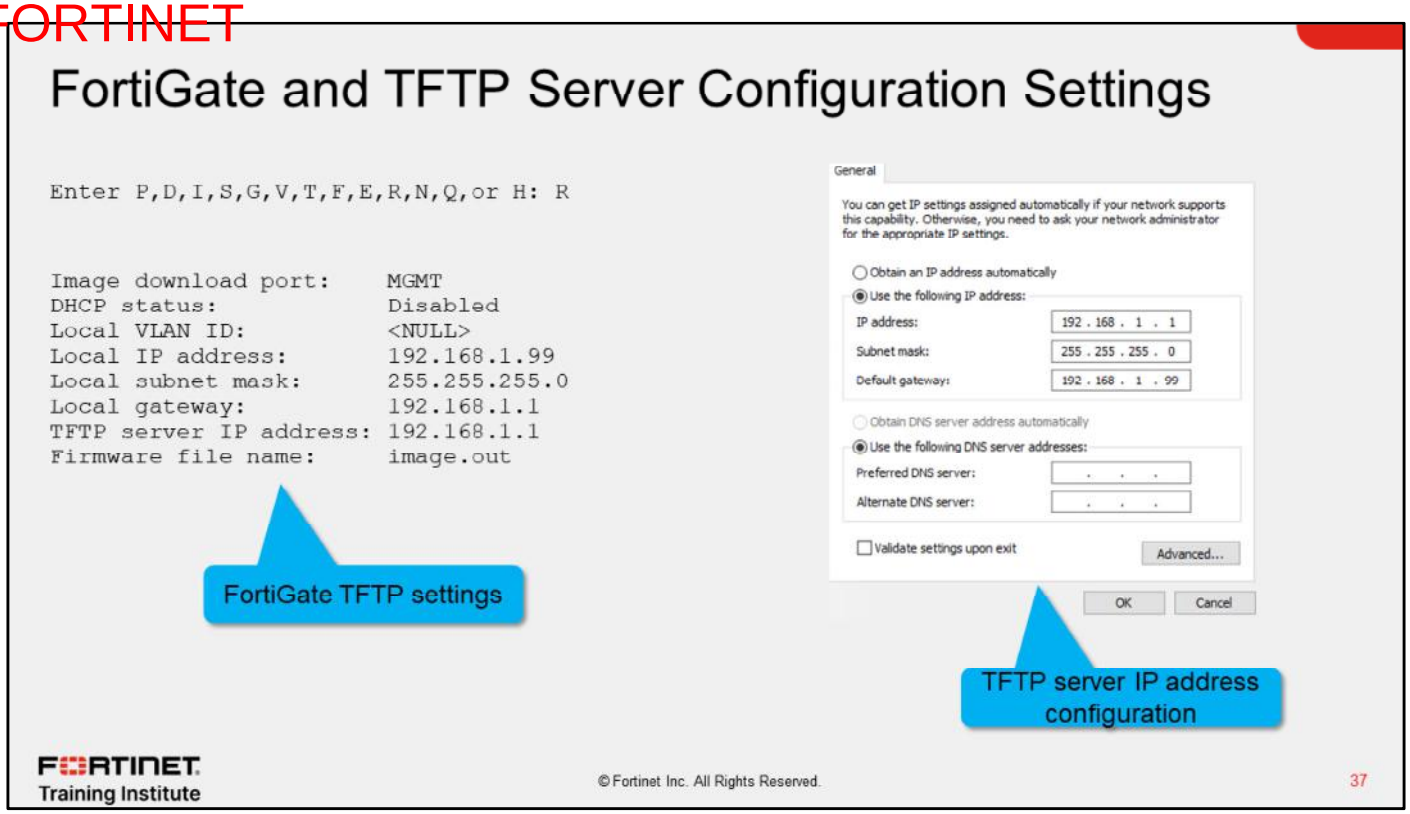

Press R to review the TFTP configuration settings.
After you have configured the TFTP parameters, press Q to return to the main configuration menu.

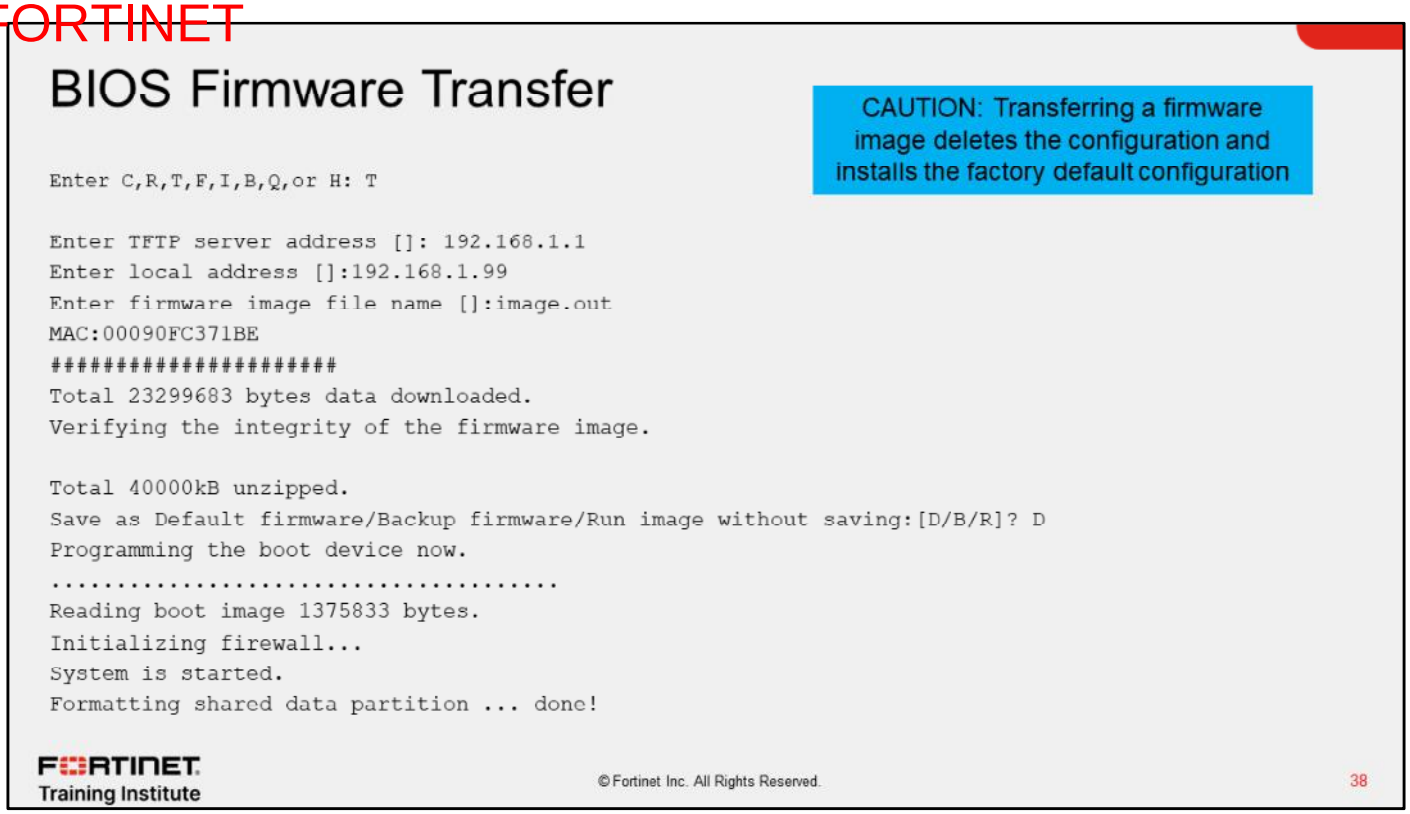

From the BIOS menu, press `T` to initiate the TFTP firmware transfer.

The BIOS requires you to enter:
- The IP address of the TFTP server
- The FortiGate IP address (it must be in the same class-C subnet as the TFTP server)
- The name of the firmware image

If everything is OK, you should see a series of pound signs, indicating that the device is downloading the image. The BIOS will then verify the integrity of the file and give you the following three options:

- `Save it as the default firmware`
- `Save it as the backup firmware`
- `Run the image without saving it`

If the firmware is going to be used in production, select the first option: `Save it as the default firmware`.

The last option (`Run the image without saving it`) allows you to run and test firmware without overwriting any existing firmware in the memory. After you have finished the tests and are ready to roll back the change, you must reboot the device, and the previously existing firmware will be used.

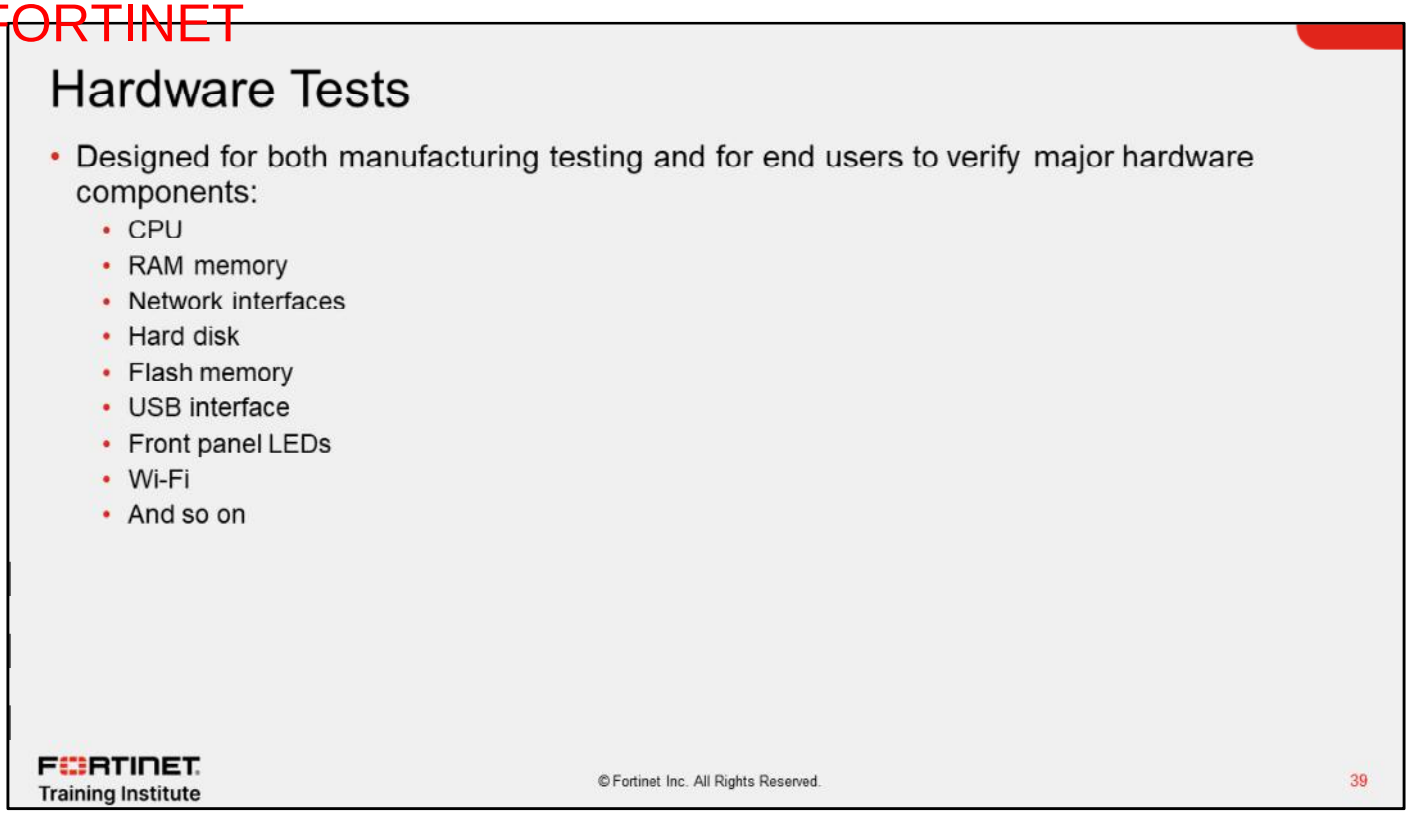

As with any other electronic device, damage to RAM can cause intermittent crashes.

If you suspect hardware failure, you can run hardware tests.

How do you run the hardware tests? It depends on the FortiGate model.

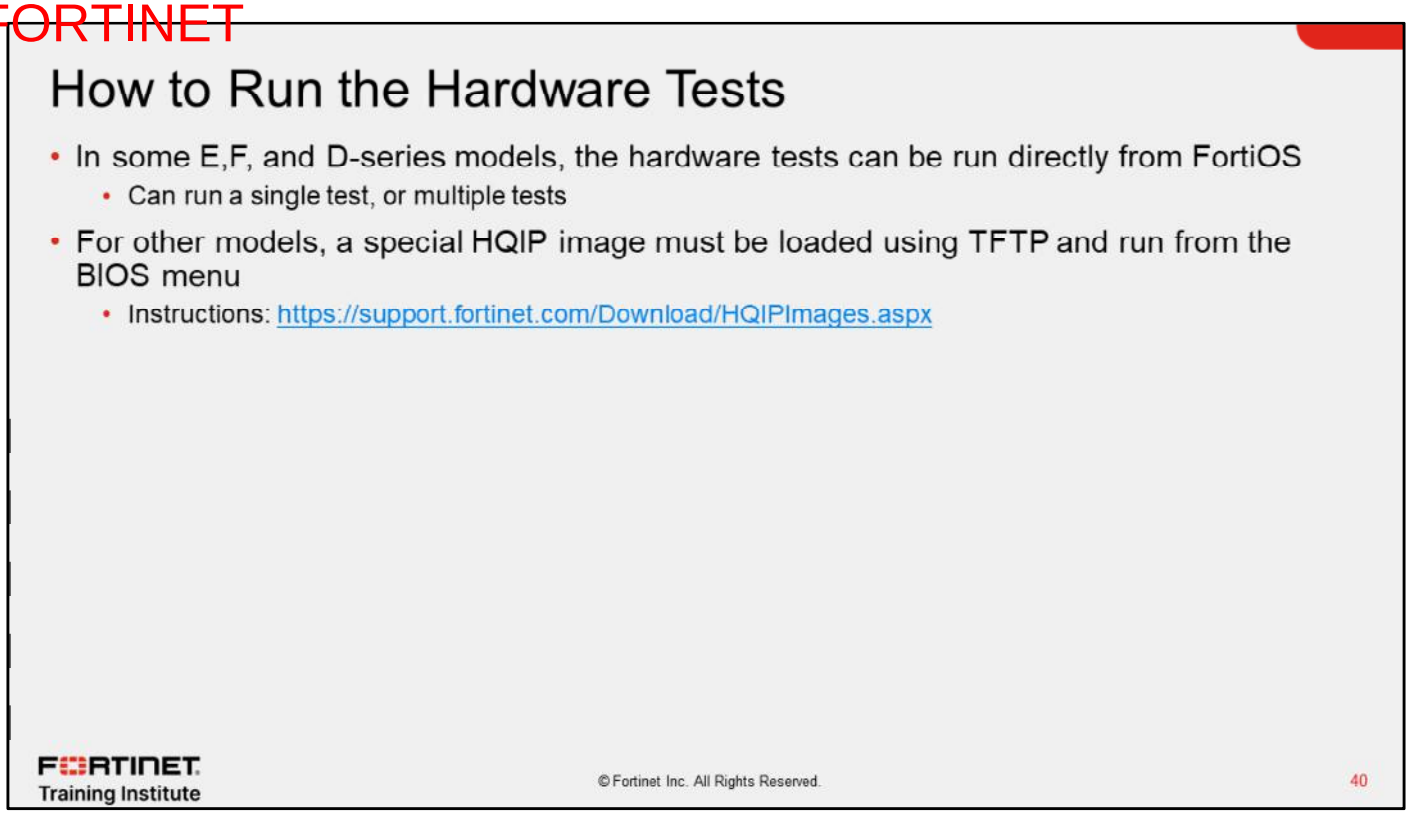

For some FortiGate E, F, and D-series models, you can run the hardware tests directly from the FortiOS CLI.

For other models, you must download special HQIP hardware testing images from the Fortinet Technical Support website.

The steps for uploading the hardware test image are the same as the ones used for uploading a firmware image. You can run the hardware test image without saving it in the flash memory, so any existing firmware image won't be overwritten.

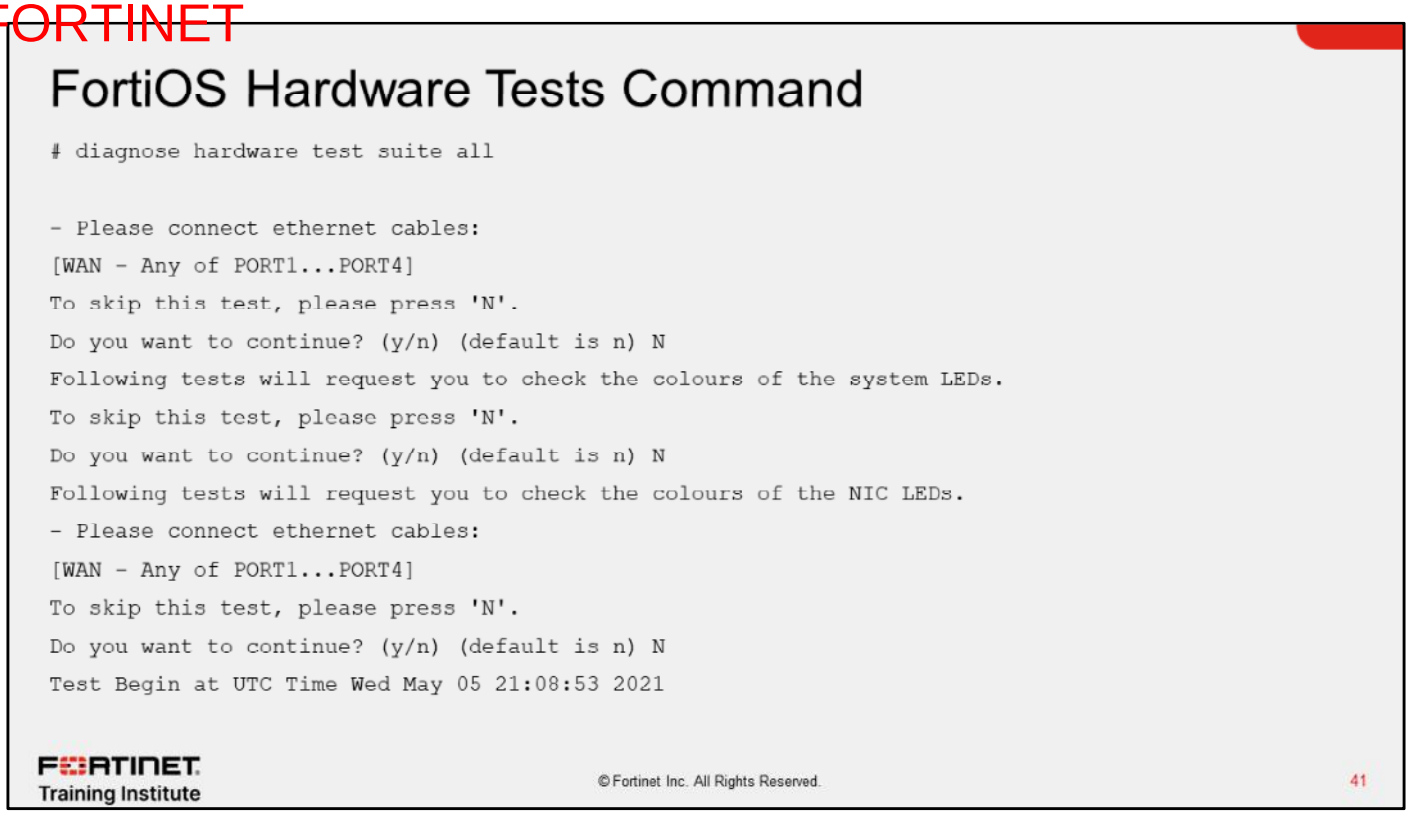

For some models, the command `diagnose hardware test suite all` runs the hardware tests from FortiOS. The hardware tests require user interaction while running. Users can skip some of the steps. Some tests require connecting external devices (such as USB flash drives) or network cables to FortiGate.

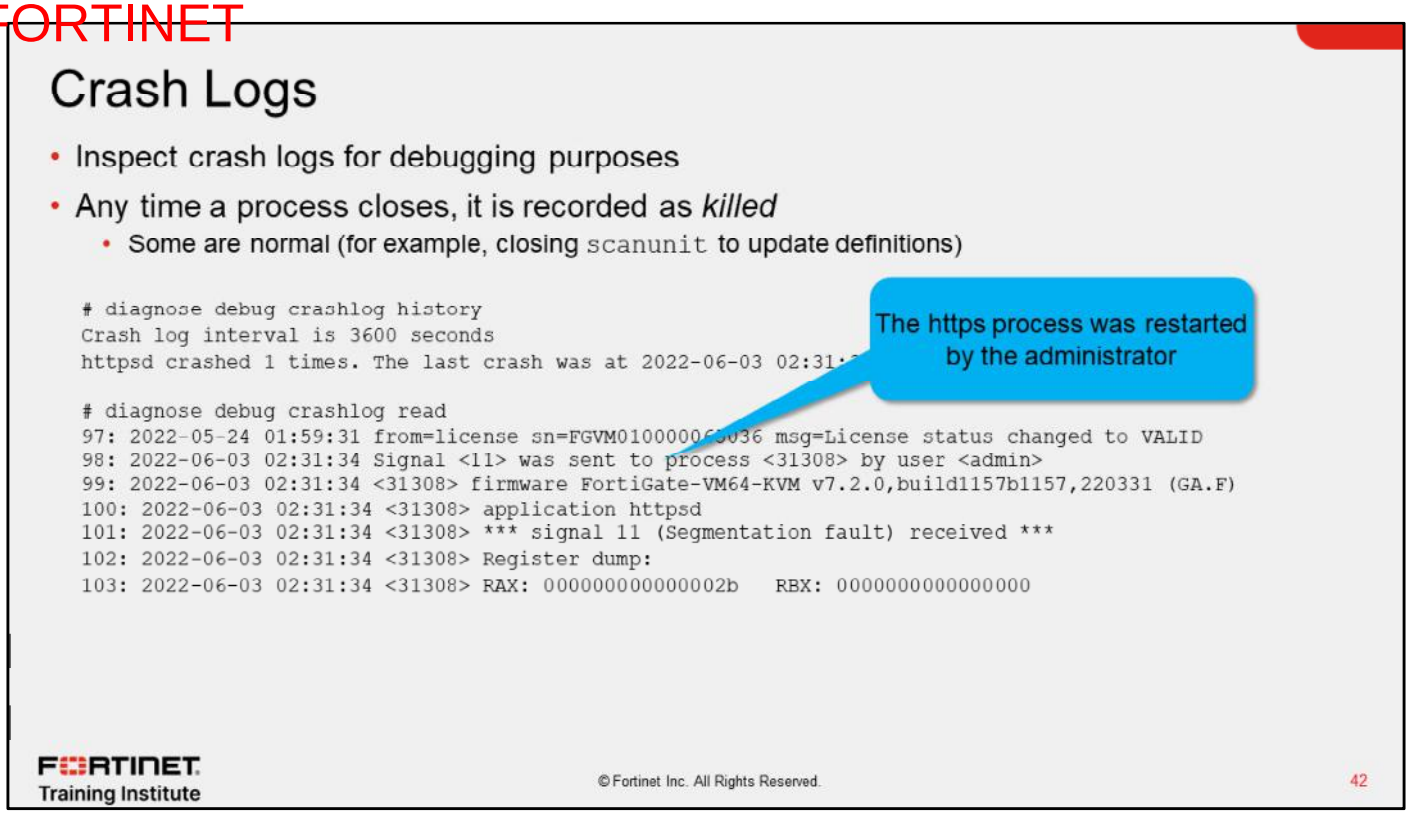

Another area you might want to monitor, purely for diagnostics, is the crash logs. Crash logs are available through the CLI.

Any time a process is closed for any reason, the crash log records this as a crash. Most of the logs in the crash log are normal. For example, any time the antivirus definitions package is updated, the `scanunit` process needs to close down in order to apply the new package. This is a normal shutdown. Some logs in the crash log shows they are initiated by a user, which indicates the administrator manually restarted a process.

Some logs in the crash log might indicate problems. For that reason, the crash logs are frequently requested by Fortinet Technical Support for troubleshooting purposes.

This slide shows the commands you have to use to get a crash log. The crashlog output shows the http process is restarted by the administrator.

Two commands can show information from the crash logs:

- `diagnose debug crashlog history` lists a summary of the processes that have crashed, how many crashes have happened, and the time of the last crash.
- `diagnose debug crashlog read` provides details about each crash, in addition to other system events, such as conserve mode entry and exit times.

Conserve Mode Events in Crash Logs

- The crash log also records conserve mode events
  - Entering:

```
12: 2021-04-06 14:10:16 logdesc="Kernel enters conserve mode" service=kernel
conserve=on free="127962
13: 2021-04-06 14:10:16 pages" red="128000 pages" msg="Kernel enters conserve
mode"
```

  - Exiting:

```
14: 2021-04-06 14:19:55 logdesc="Kernel leaves conserve mode" service=kernel
conserve=exit
15: 2021-04-06 14:19:55 free="192987 pages" green="192000 pages" msg="Kernel
leaves conserve mode"
```

FORTINET Training Institute

43

This slide shows the entries generated in the crash logs when FortiGate enters and exits memory conserve mode.

Knowledge Check

1. Which types of information are stored in the crash log?
   - ✓ A. Process crashes and conserve mode events
   - B. Traffic logs and security logs

2. Which protocol is used to upload new firmware from the console?
   - A. HTTP/HTTPS
   - ✓ B. TFTP

FORTINET Training Institute

44

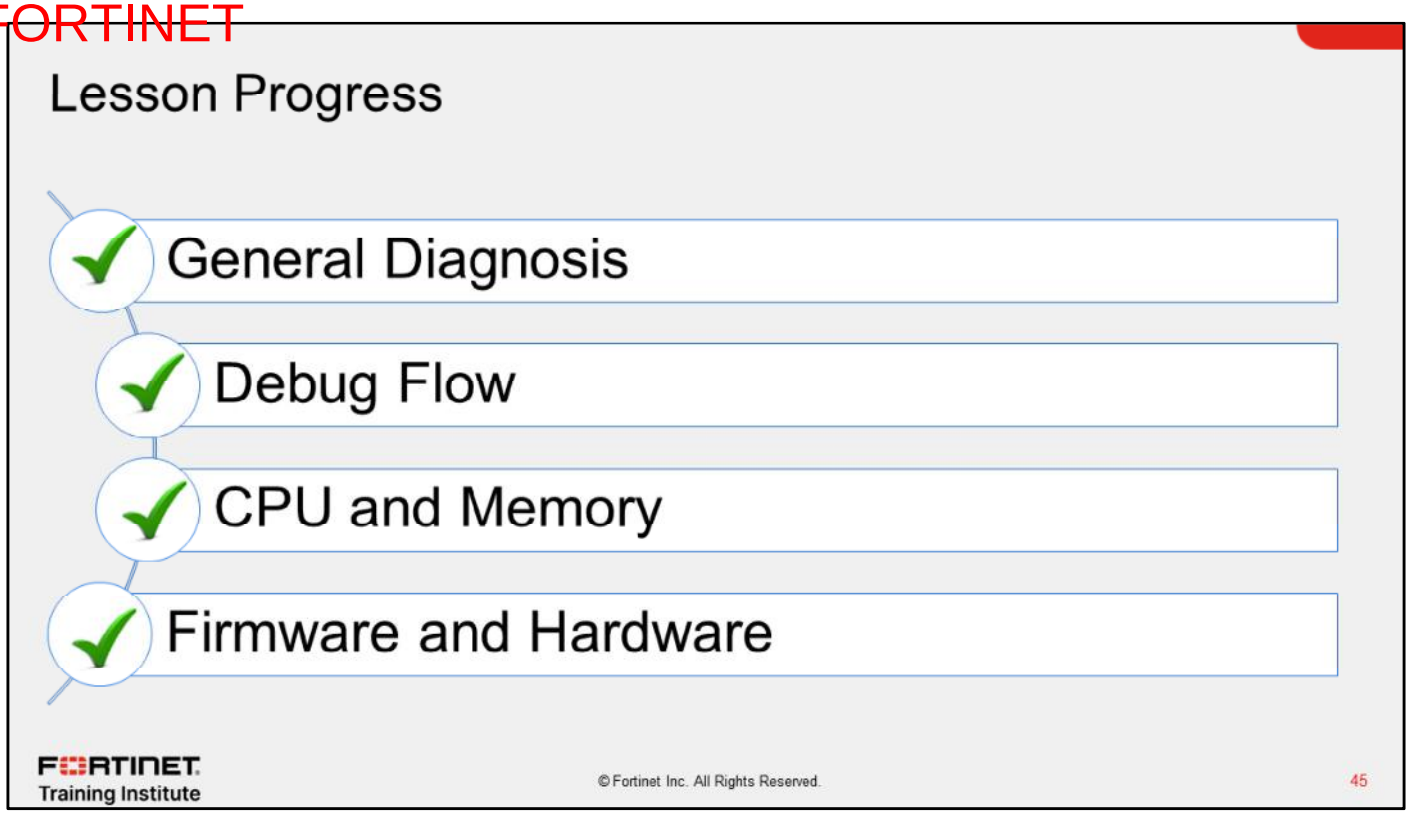

Congratulations! You have completed the lesson.

Now, you will review the objectives that you covered in this lesson.

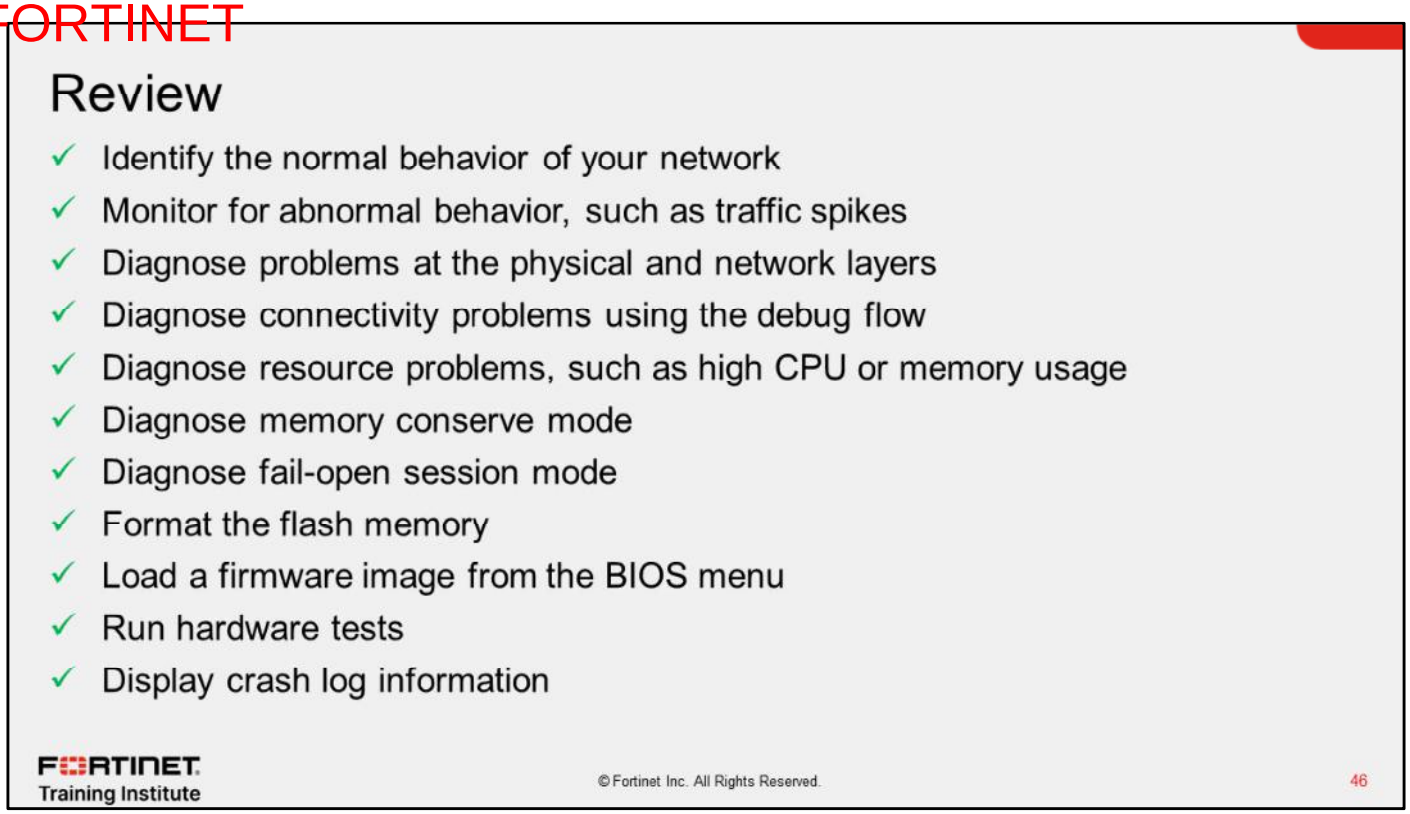

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use diagnostic commands and tools.