FortiGate Infrastructure Study Guide

for FortiOS 7.0

FORTINET NSE 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

1/24/2022

TABLE OF CONTENTS

Change Log

This table includes updates to the *FortiGate Infrastructure 7.0 Study Guide* dated 6/7/2021 to the updated document version dated 1/24/2022.

Change	Location
Various formatting fixes	Entire Guide
• Removed slides 15 • Fixed slide 27 (now slide 26) • Removed slides related to STP (slides 36-41)	Lesson 04
Added IKEv1 vs. IKEv2	Lesson 5: Slide 9
Updated HA A-A Load Balance slides	Lesson 7: Slides 29-34

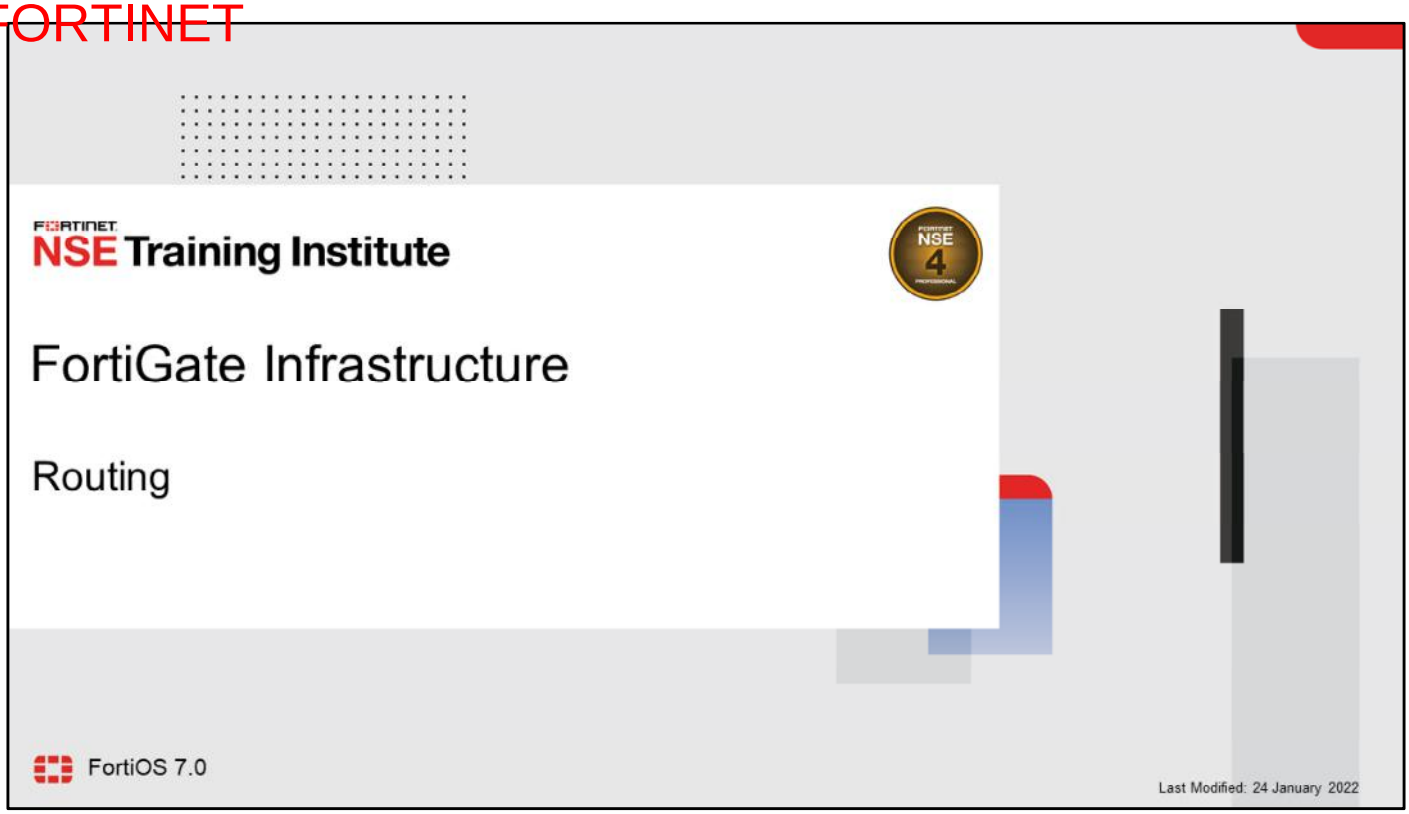

In this lesson, you will learn about the routing capabilities and features available on FortiGate.

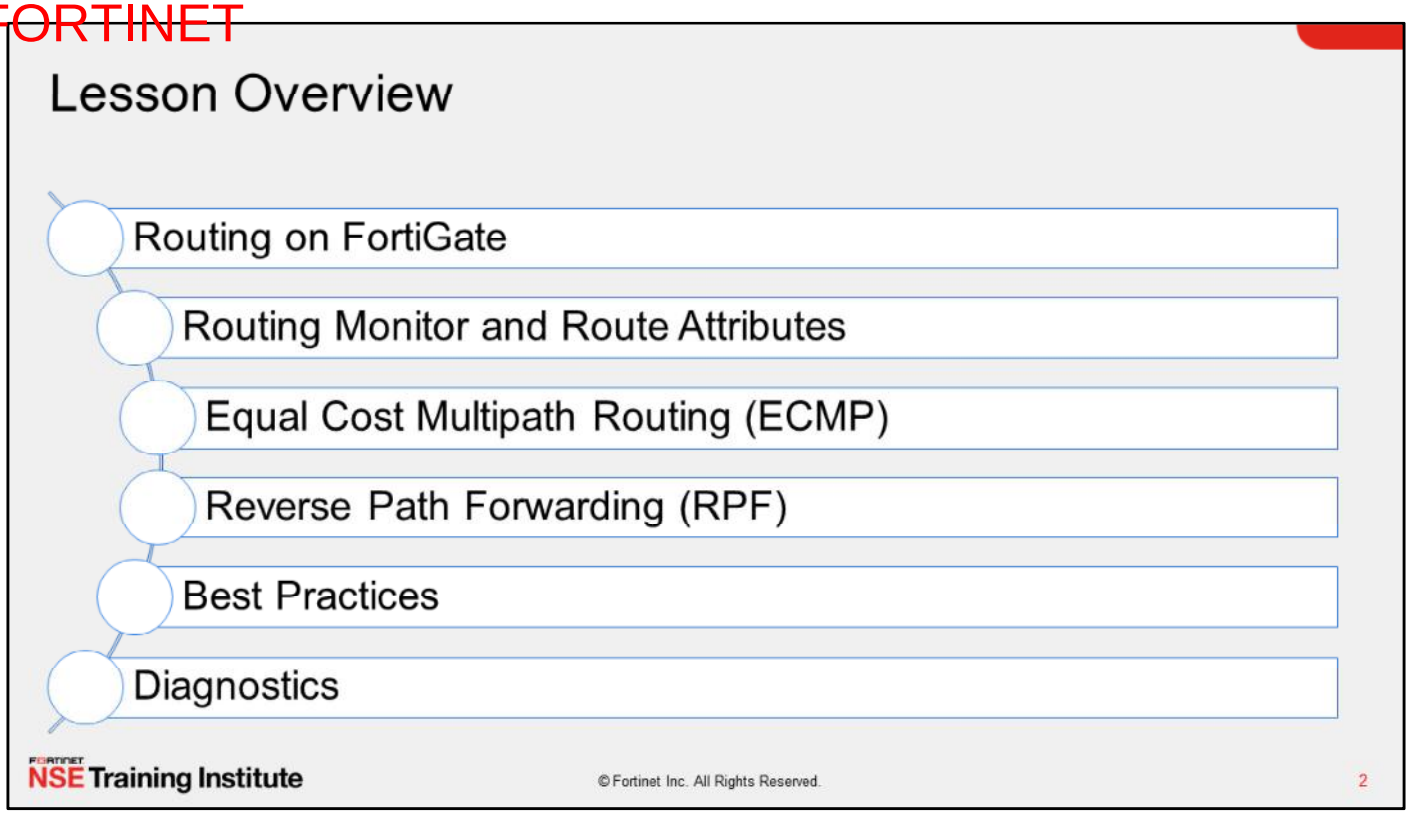

In this lesson, you will learn about the topics shown on this slide.

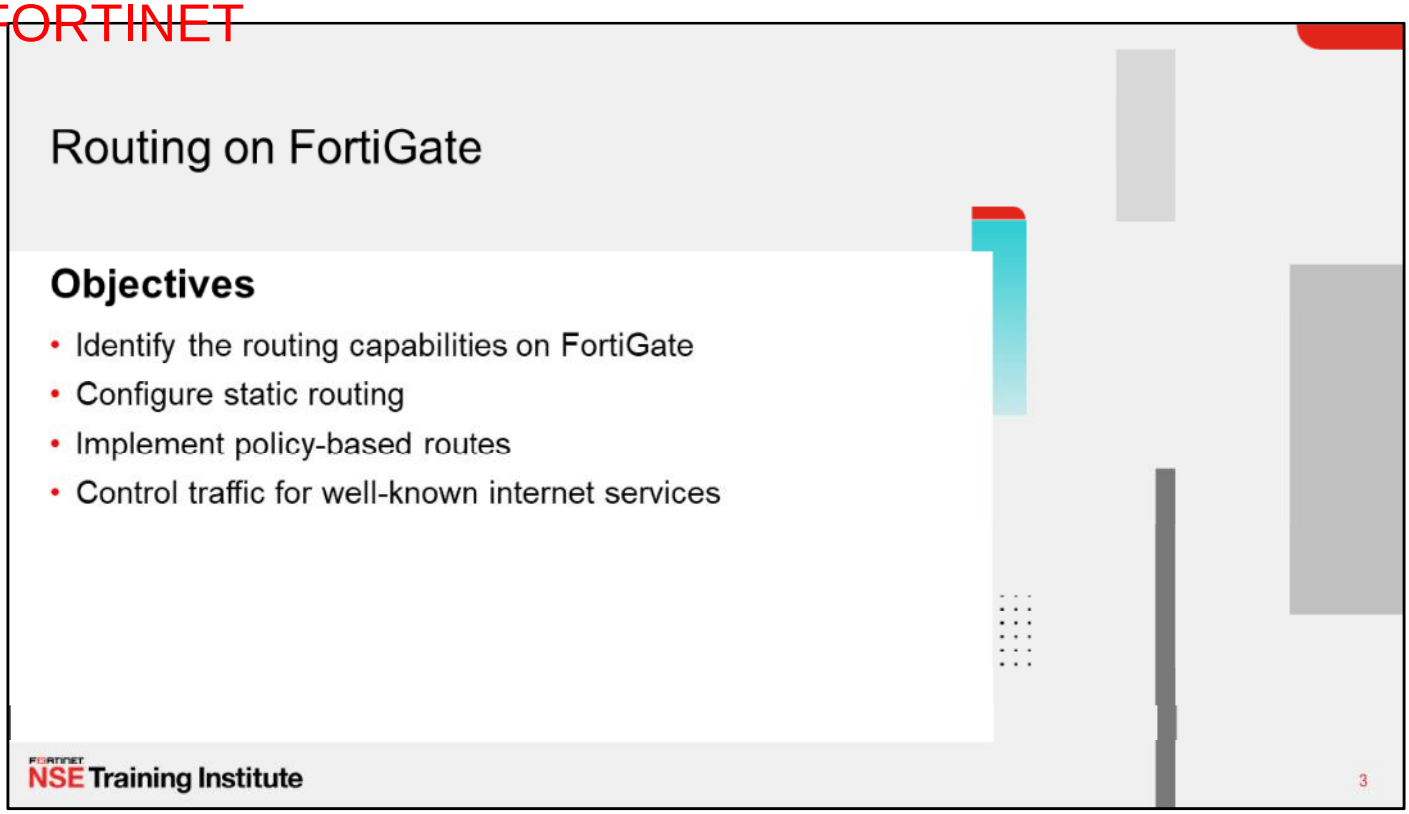

After completing this section, you should be able to achieve the objectives shown on this slide.

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

What Is IP-Layer Routing?

- Routing is how packets are sent along a path—from point-to-point on the network—from source to destination
  - If the destination is on a subnet that is not directly connected to the router, the packet is relayed to another router that is closer
  - Entries in the routing table can be configured manually, dynamically, or both
- A FortiGate in NAT mode, among other things, is an OSI Layer 3 router

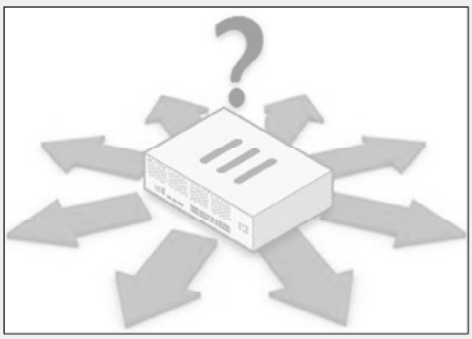

NSE Training Institute

4

What is routing?

Routing is how FortiGate in NAT mode decides where to send the packets that it receives and the packets that it generates.

All network devices that perform routing have a routing table. A routing table contains a series of rules. Each rule specifies the next *hop*, which may or may not be the final destination of the packet. Each routing *hop* in the routed path requires a routing table lookup to pass the packet along until it reaches the final destination.

When routing packets, FortiGate will first finds a matching route in its list of routes based on the packet's destination address. When performing this match, FortiGate evaluates the entire routing table to find the most specific match before selecting a route. If FortiGate finds multiple matches, it uses various route attributes to determine the best route.

Proper routing configuration is important. If routes are misconfigured, packets will not reach their destination and will be lost.

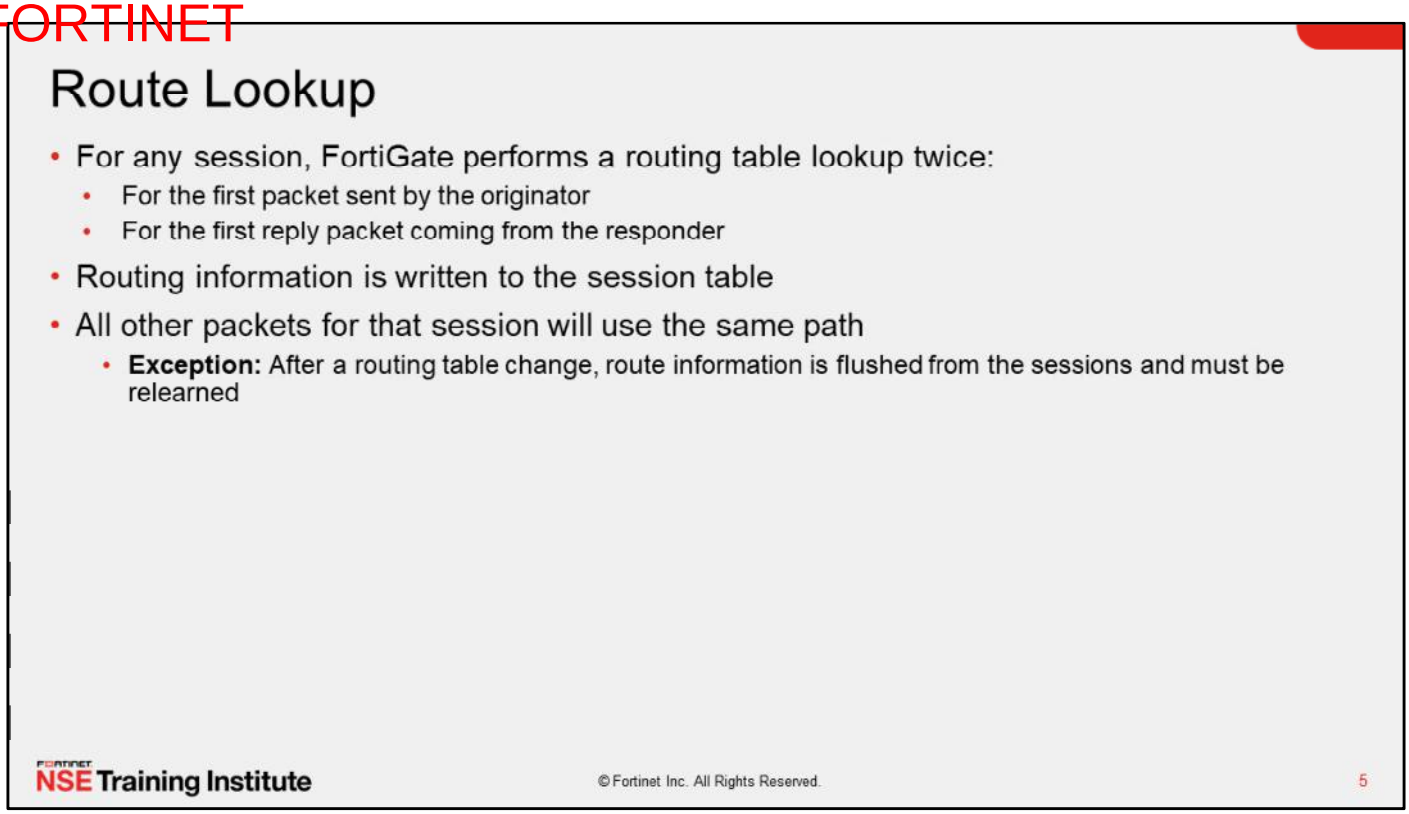

By default, many aspects of FortiGate are stateful. That is, FortiGate decides many things at the beginning of a session, when it receives the first packet.

For each session, FortiGate performs two routing 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, even after a change in the static routes. However, there is an exception to this rule: if there is a change in the routing table, FortiGate removes the route information for the session table, and then it makes additional routing table lookups to rebuild this information.

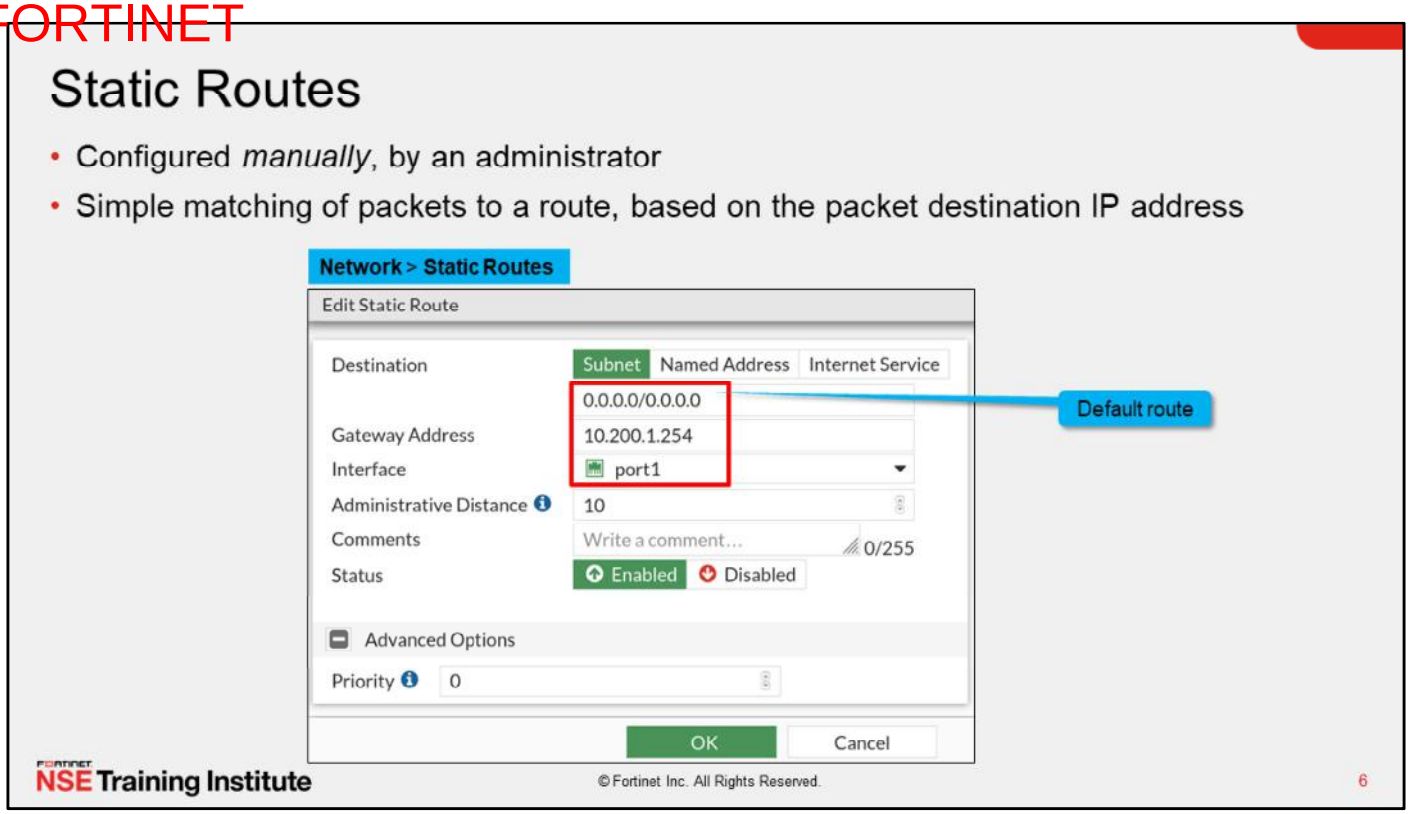

One type of manually configured route is called a static route. When you configure a static route, you are telling FortiGate, “When you see a packet whose destination is within a specific range, send it through a specific network interface, towards a specific router.” You can also configure the distance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

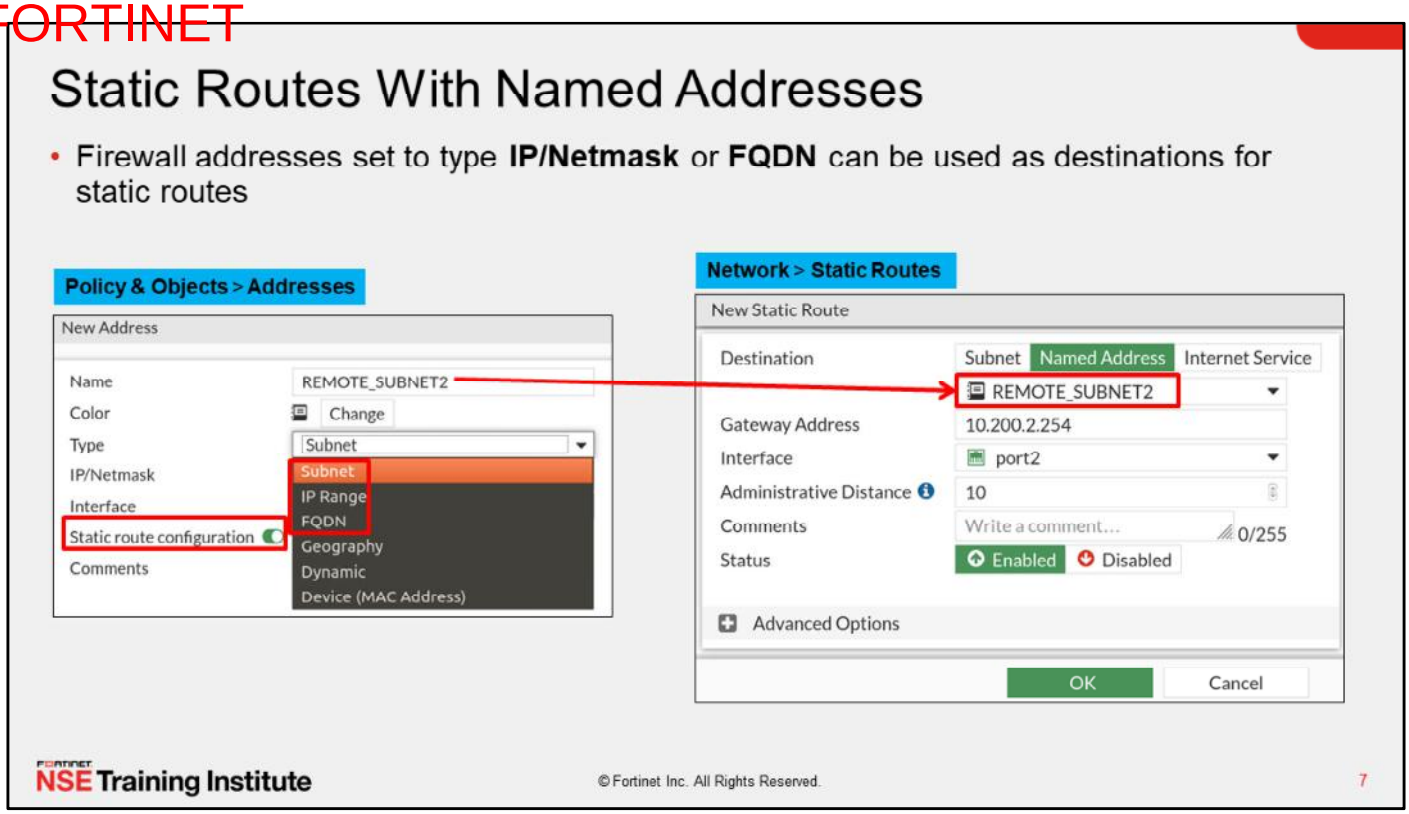

If you create a firewall address object with the type **IP/Netmask** or **FQDN**, you can use that firewall address as the destination of one or more static routes. First, enable **Static route configuration** in the firewall address configuration. After you enable it, the firewall address object becomes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

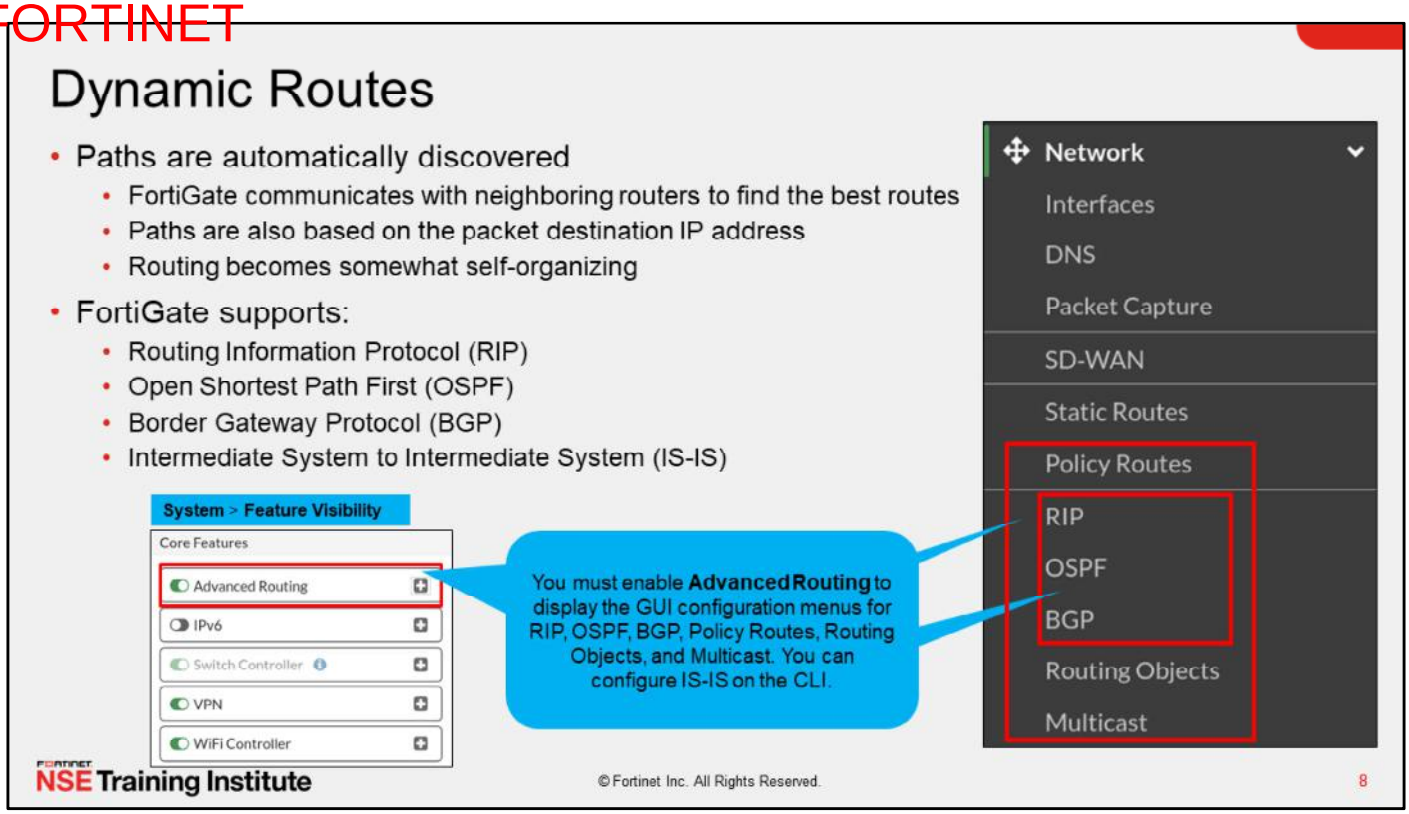

For large networks, manually configuring hundreds of static routes may not be practical. Your FortiGate can help, by learning routes automatically. FortiGate supports several dynamic routing protocols: RIP, OSPF, BGP, and IS-IS.

In dynamic routing, FortiGate communicates with nearby routers to discover their paths, and to advertise its own directly connected subnets. Discovered paths are automatically added to the FortiGate routing table. So verify that your neighbor routers are trusted and secured!

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

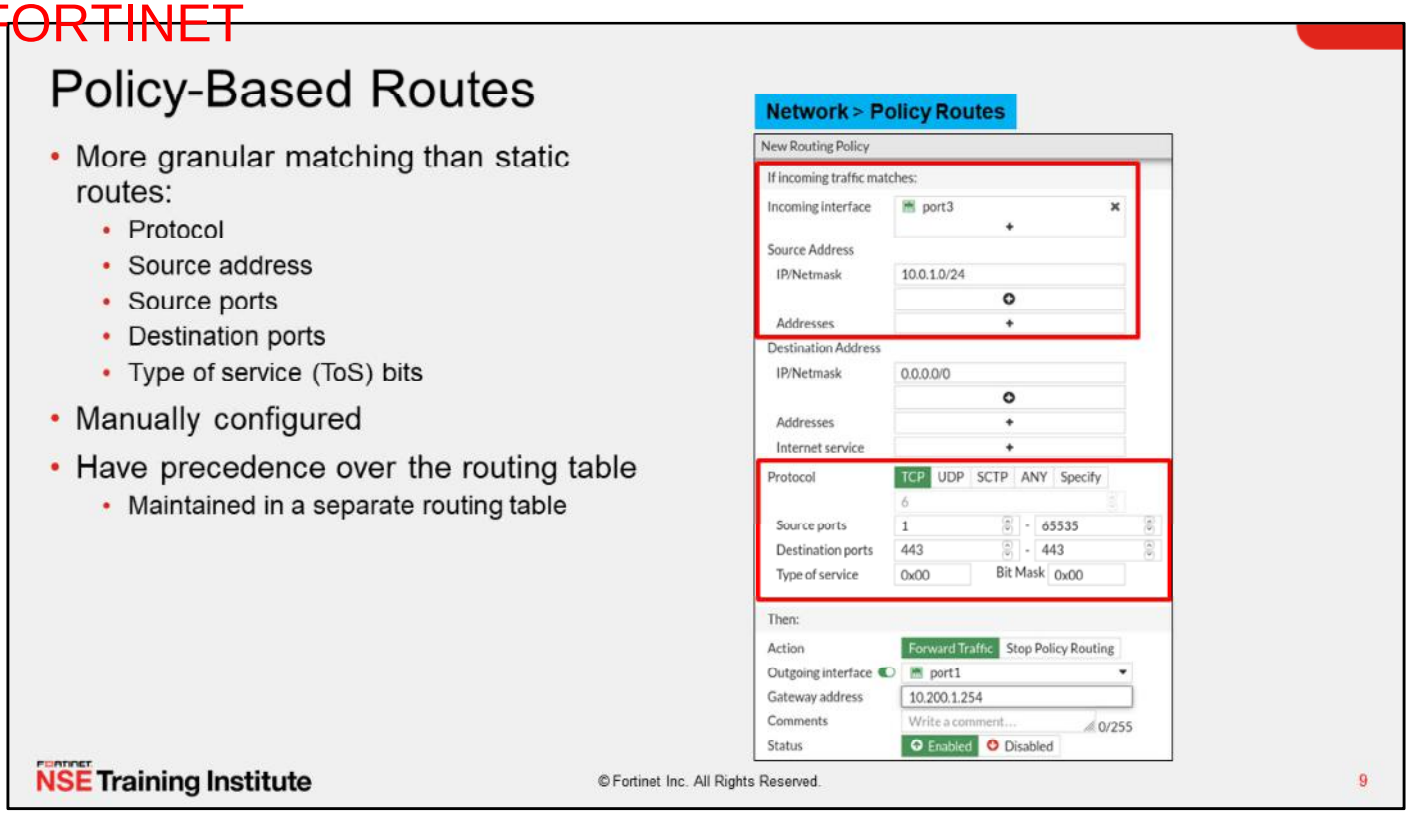

Static routes are simple and are often used in small networks. Policy-based routes, however, are more flexible. They can match more than just the destination IP address. For example, if you have two links—a slow one and a fast one—you can route packets from low-priority source IPs to the slow link.

Policy routes set to the action **Forward Traffic** have precedence over static and dynamic routes. So, if a packet matches the policy route, FortiGate bypasses any routing table lookup.

Like static routes, policy-based routes must be valid: a destination and gateway are required, and disconnected (or down) interfaces can't be used. For policy-based routes, packets must also match all specified subnets, ToS bits, and port number. So, if you don't want a setting to be included in the matching criteria, leave it blank.

Policy routes are maintained in a separate routing table by FortiGate, and have precedence over the regular routing table.

Policy-Based Routing Actions

- If traffic matches a policy-based route, FortiGate either:
  - Forwards traffic using the specified outgoing interface to the specified gateway
  - Stops policy routing and uses the routing table instead

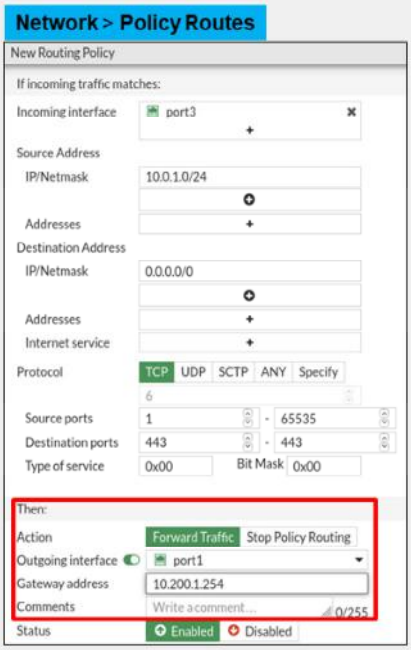

NSE Training Institute

10

When a packet matches a policy route, FortiGate takes one of two actions. Either it routes the packet to the configured interface and gateway, bypassing the routing table, or it stops checking the policy routes, so the packet is routed based on the routing table.

Remember, for a policy route to forward traffic out a specific interface, there should be an active route for that destination using that interface in the routing table. Otherwise, the policy route does not work.

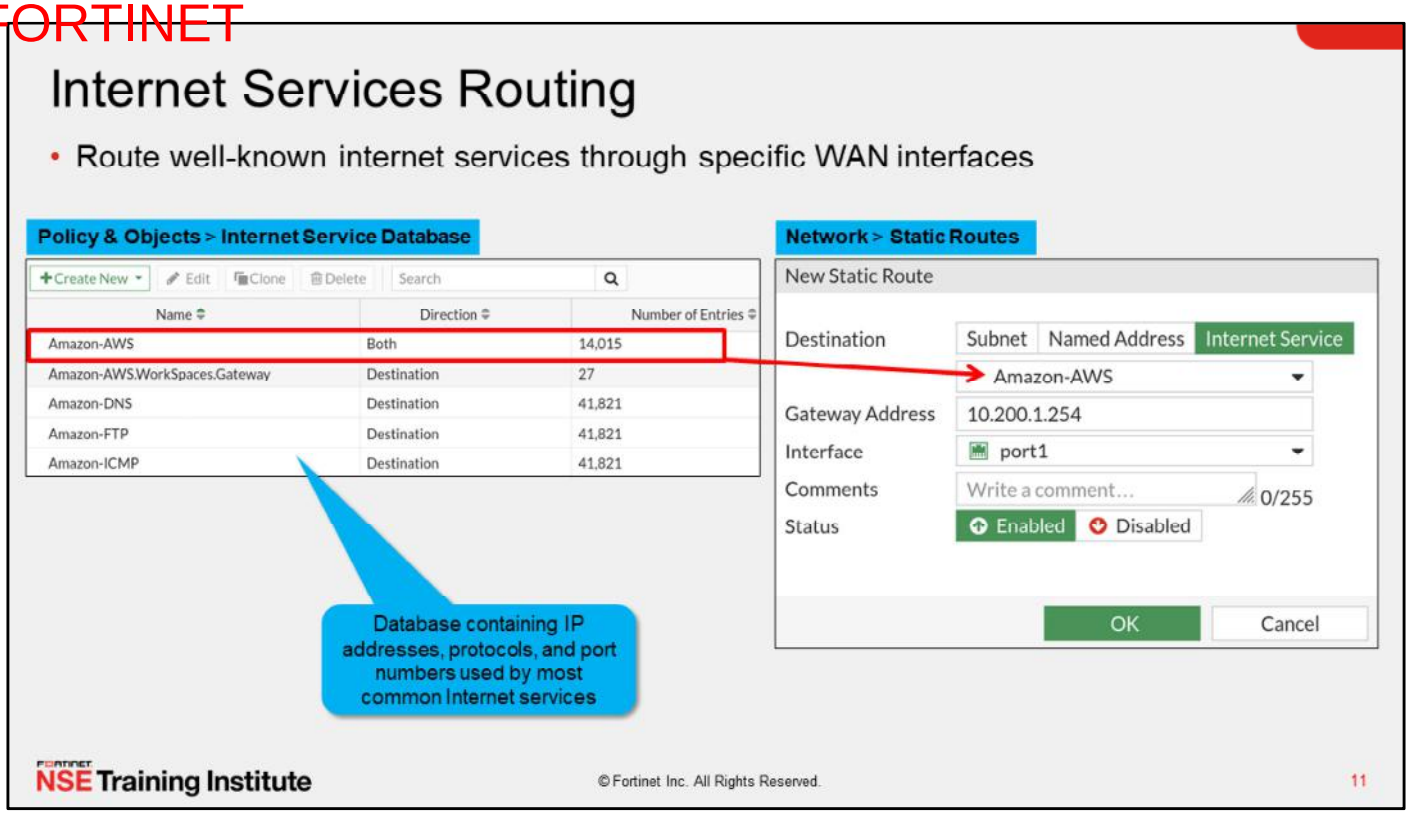

What happens if you need to route traffic to a public internet service (such as Amazon-AWS or Apple Store) through a specific WAN link? Say you have two ISPs and you want to route Netflix traffic through one ISP and all your other internet traffic though the other ISP. To achieve this goal, you need 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.

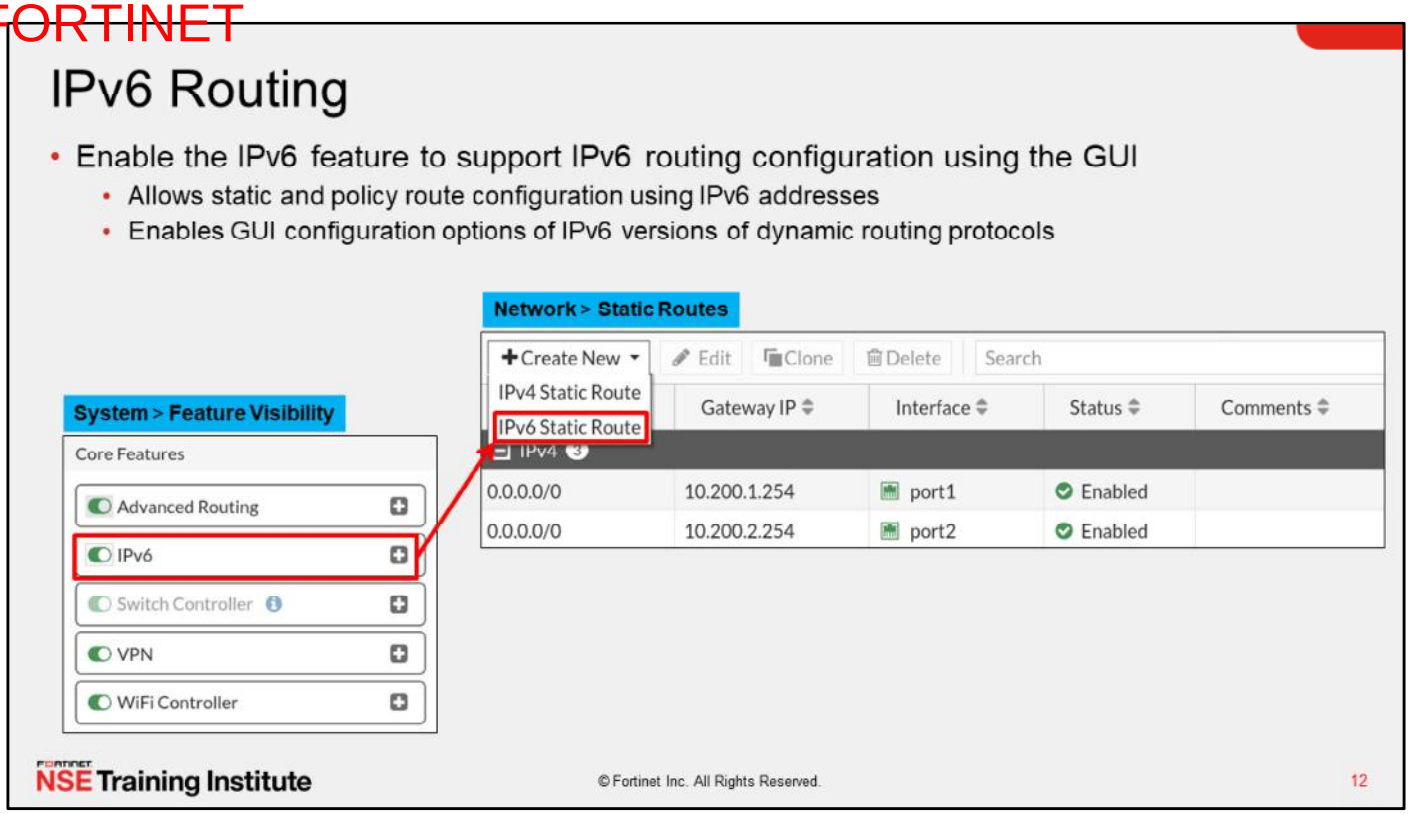

To enable routing configuration for IPv6 addresses using the GUI, you must enable **IPv6** in the **Feature Visibility** menu. Then, you can create static routes and policy routes with IPv6 addresses. Enabling the IPv6 feature also enables GUI configuration options for IPv6 versions of the dynamic routing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

NSE Training Institute

13

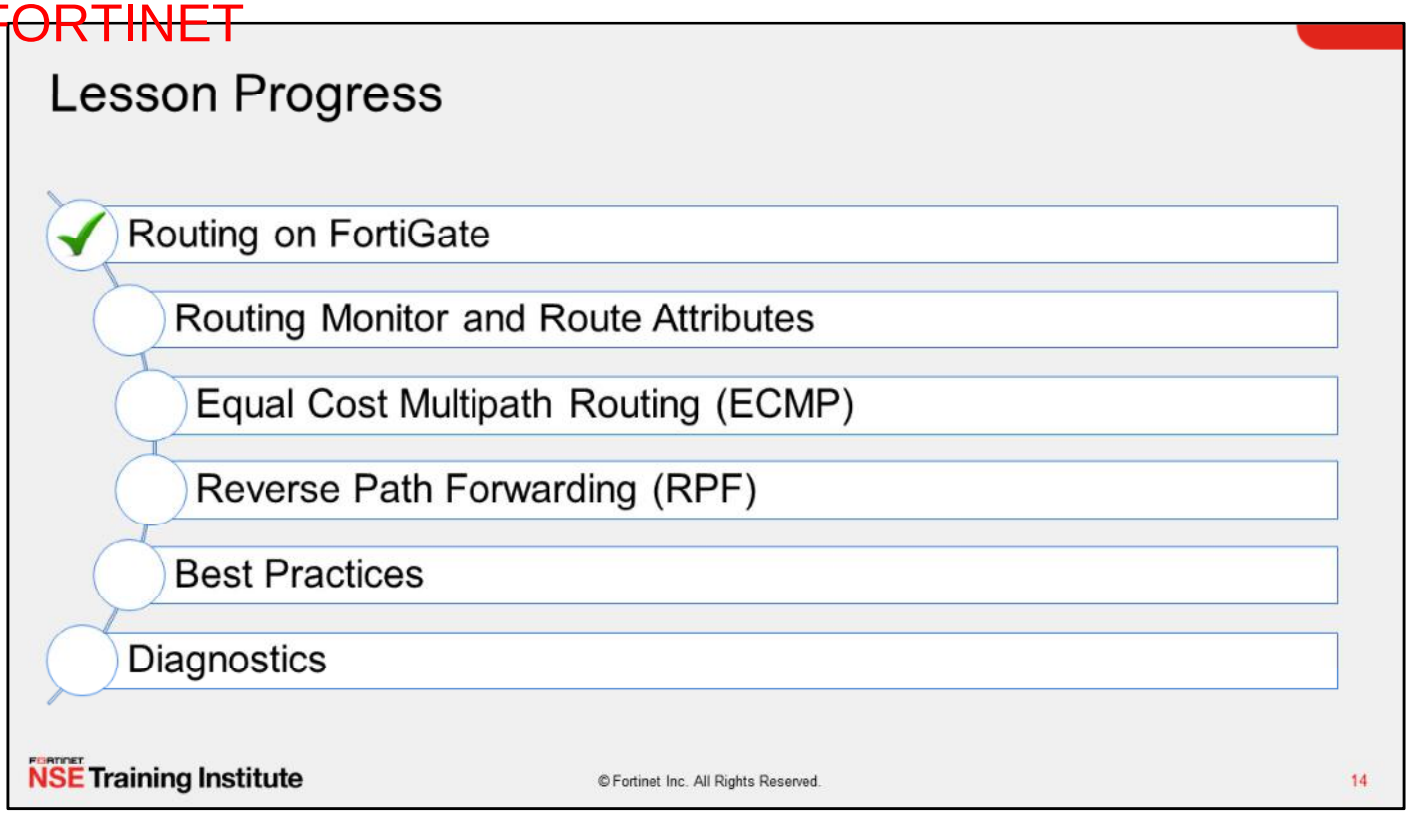

Good job! You now understand routing on FortiGate.

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

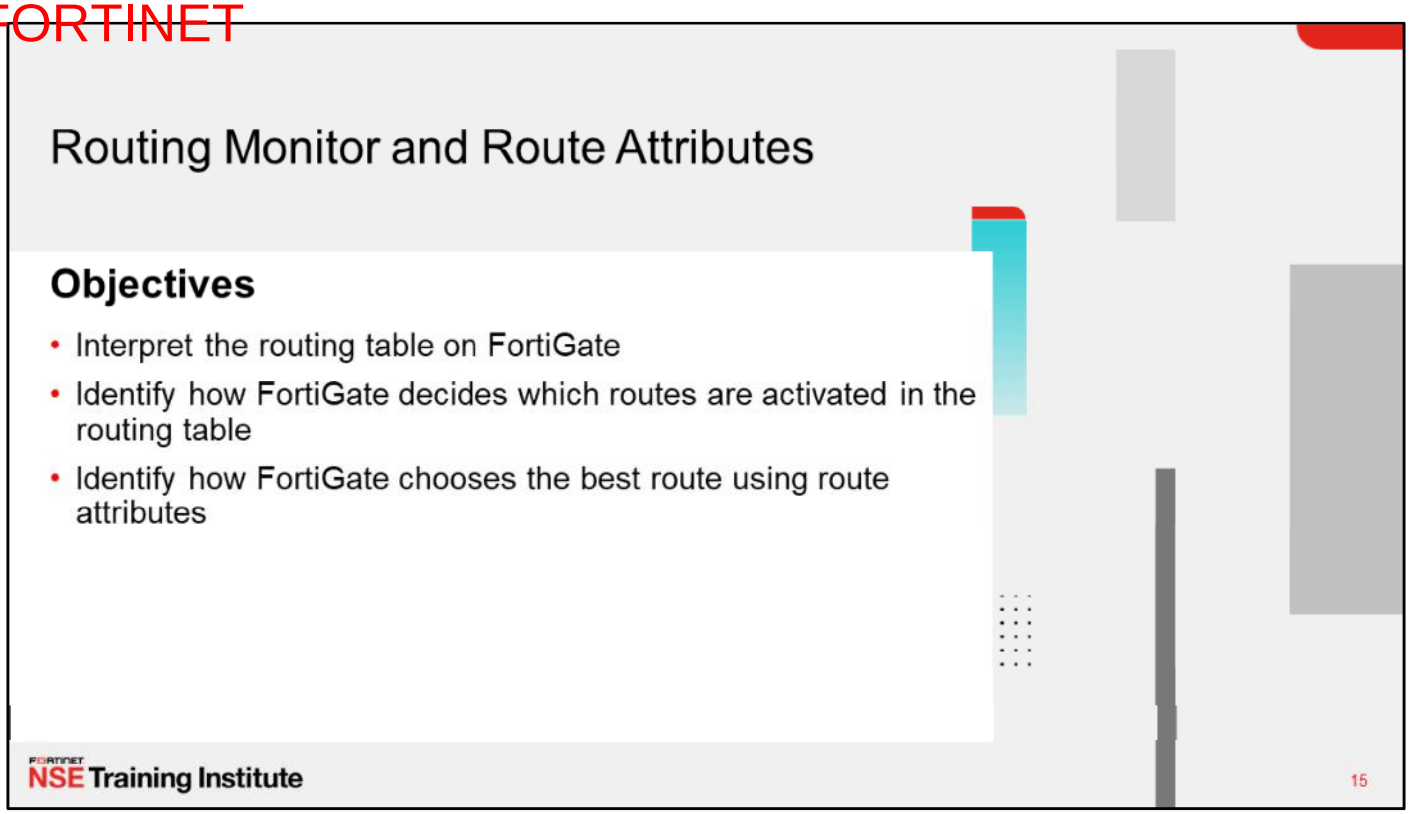

After completing this section, you should be able to achieve the objectives shown on this slide.

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

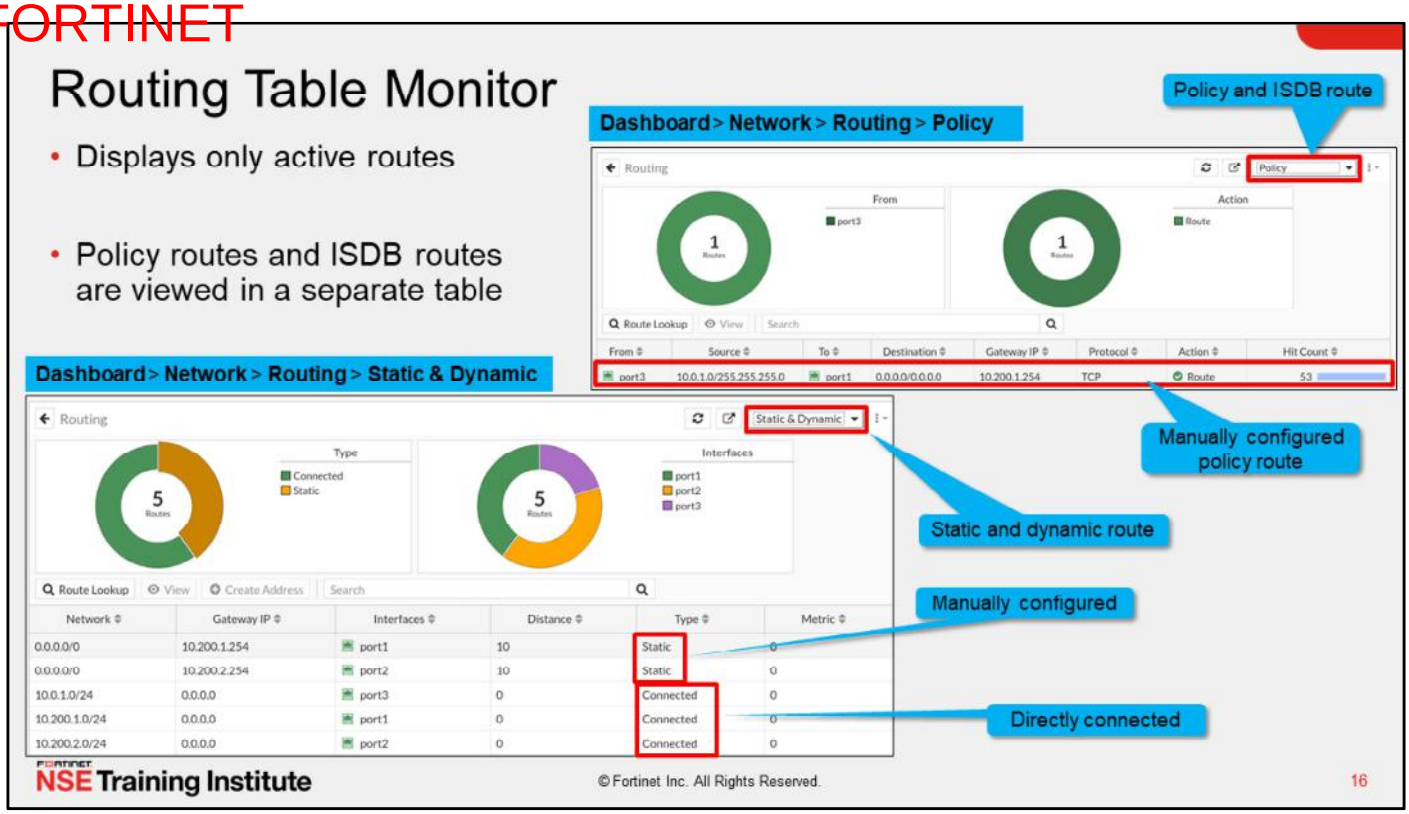

The routing table monitor on the FortiGate GUI shows the active routes.

Which routes, besides the static routes, are displayed here?

- Directly connected subnets: when a subnet is assigned to the FortiGate's interface, a route to the subnet is automatically added with **Connected** shown in the **Type** column. For connected routes to be displayed, the respective interface link must be up. This means that if an interface is down, or there is no link established, the route is not added.
- Dynamic routes: on larger networks, your FortiGate may receive routes from other routers, through protocols such as BGP or OSPF. FortiGate adds these routes to the routing table with the respective routing protocol's name under the **Type** column.

Which configured routes aren't displayed in the routing table monitor?

- Inactive routes
    - If an interface is administratively down, has its link down, or its gateway has been detected dead by the link monitor feature, then that route is considered inactive, and will not be added to the routing table.
- Standby routes. These are valid duplicate routes that have higher distance. For instance:
    - A second default static route with a higher distance
    - A dynamic route such as BGP, RIP, 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 are viewed in a separate table. ISDB routes are also added as policy routes in the policy route monitor.

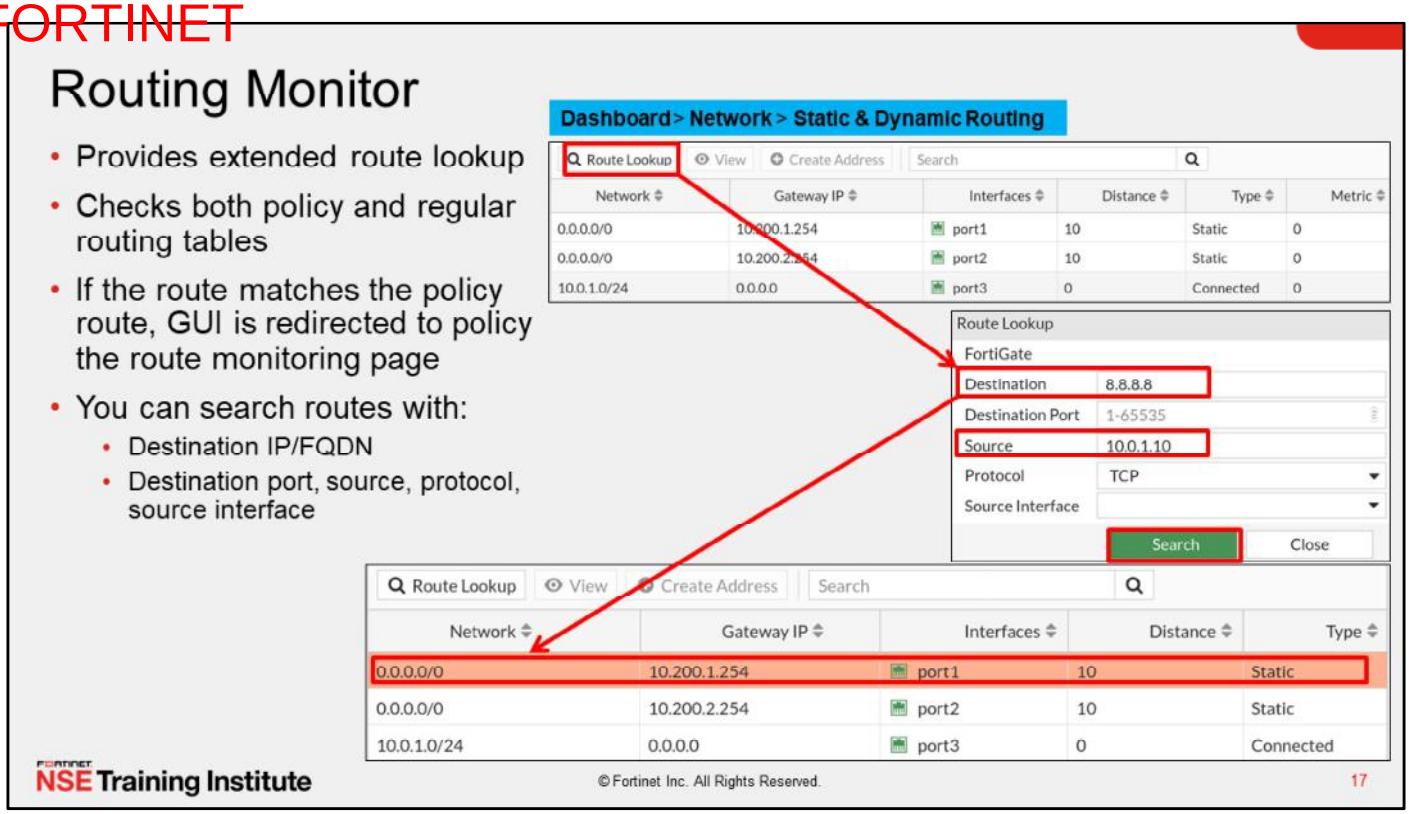

The **Routing Monitor** provides extended route lookup for the routing monitor. This feature checks both policy routes and the regular routing table. If the route matches the policy route, then you are redirected to the **Policy Route** page on the GUI.

In the **Route Lookup**, you can specify a destination address and optionally specify a destination port, source IP, protocol, and source interface to search a route based on these criteria.

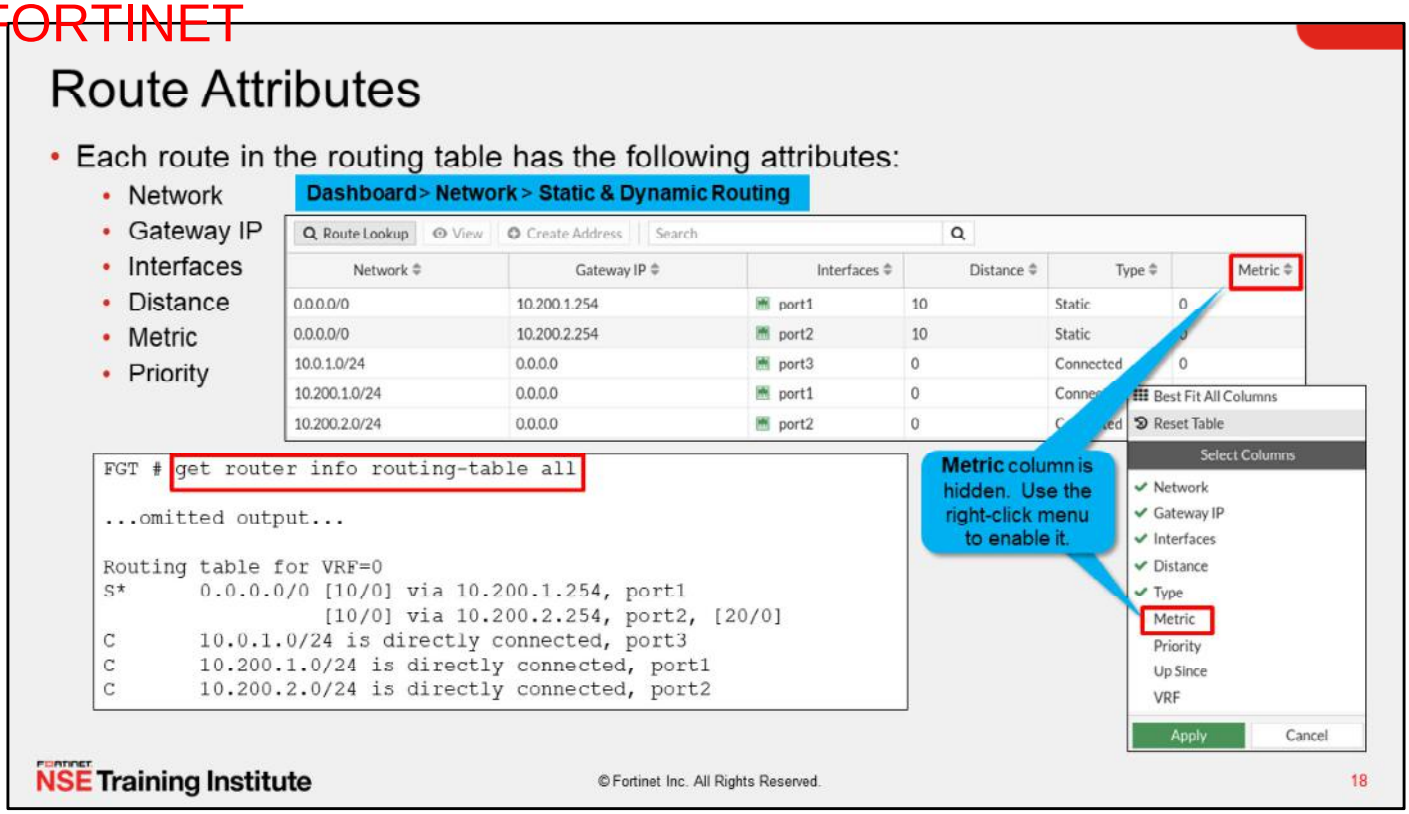

Each of the routes listed in the routing table includes several attributes with associated values.

The **Network** column lists the destination IP address and subnet mask that will be matched. The **Interface** column lists the interface that will be used 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 later in this lesson.

Distance

- Used to rank routes from most preferred (low distance value) to least preferred (high distance value)
- If multiple routes for the same destination exist, the one with the *lowest* distance is installed in the routing table (active), and the rest are not (standby)

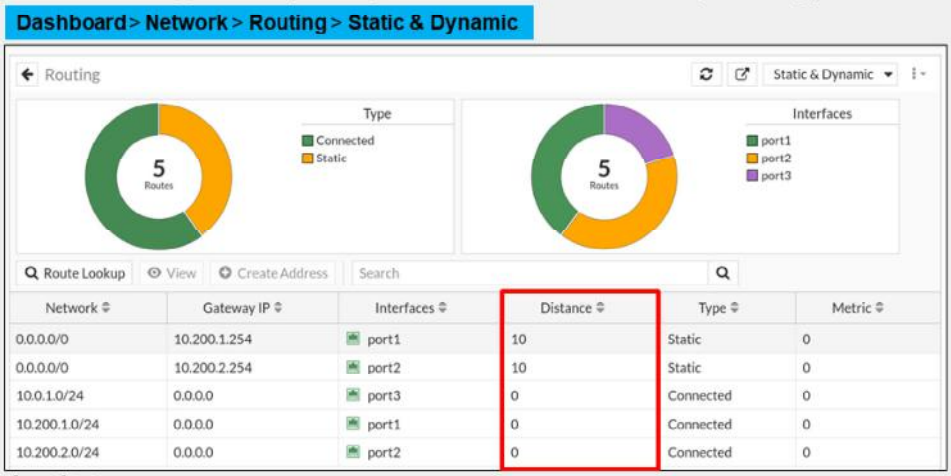

NSE Training Institute

19

Distance, or administrative distance, is a number that is used by routers to determine which route is preferred for a particular destination. If there are two routes to the same destination, the one with the smaller distance is considered *better* (active) and used for routing. The routes with higher distances (standby) are not installed in the routing table.

By default, routes learned through the RIP protocol have a higher distance value than routes learned through the OSPF protocol. OSPF is considered to be more accurate than RIP.

The following values are the default distances on FortiGate:
- 0 - directly connected
- 5 - DHCP gateway
- 20 - external BGP (EBGP) routes
- 200 - internal BGP (IBGP) routes
- 110 - OSPF routes
- 120 - RIP routes
- 10 - static routes

Metric

- Used by dynamic routing protocols to identify the best route to a destination
- If multiple dynamic routes have the same distance, then the metric is used to break the tie
  - The route with the lowest metric is chosen
- The calculation method differs among routing protocols

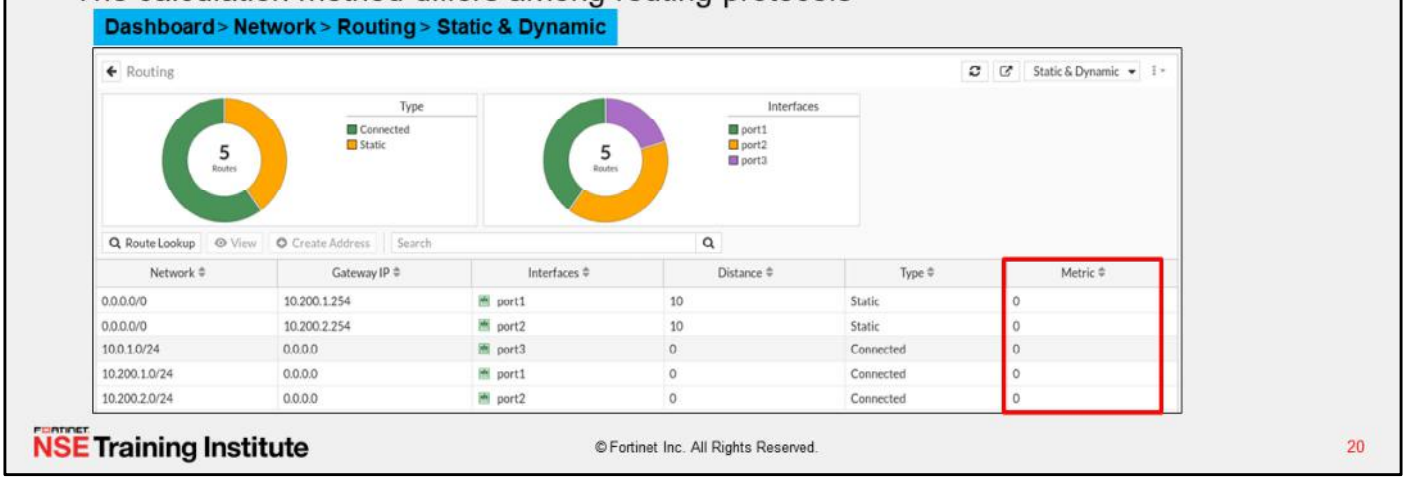

The metric attribute is used to determine the best route to a destination when dealing with routes learned through dynamic routing protocols. If two routes have the same distance, the metric value is used to break the tie. The route with the lowest metric is chosen for routing.

How the metric value is measured depends on the routing protocol. 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 how much bandwidth a link has.

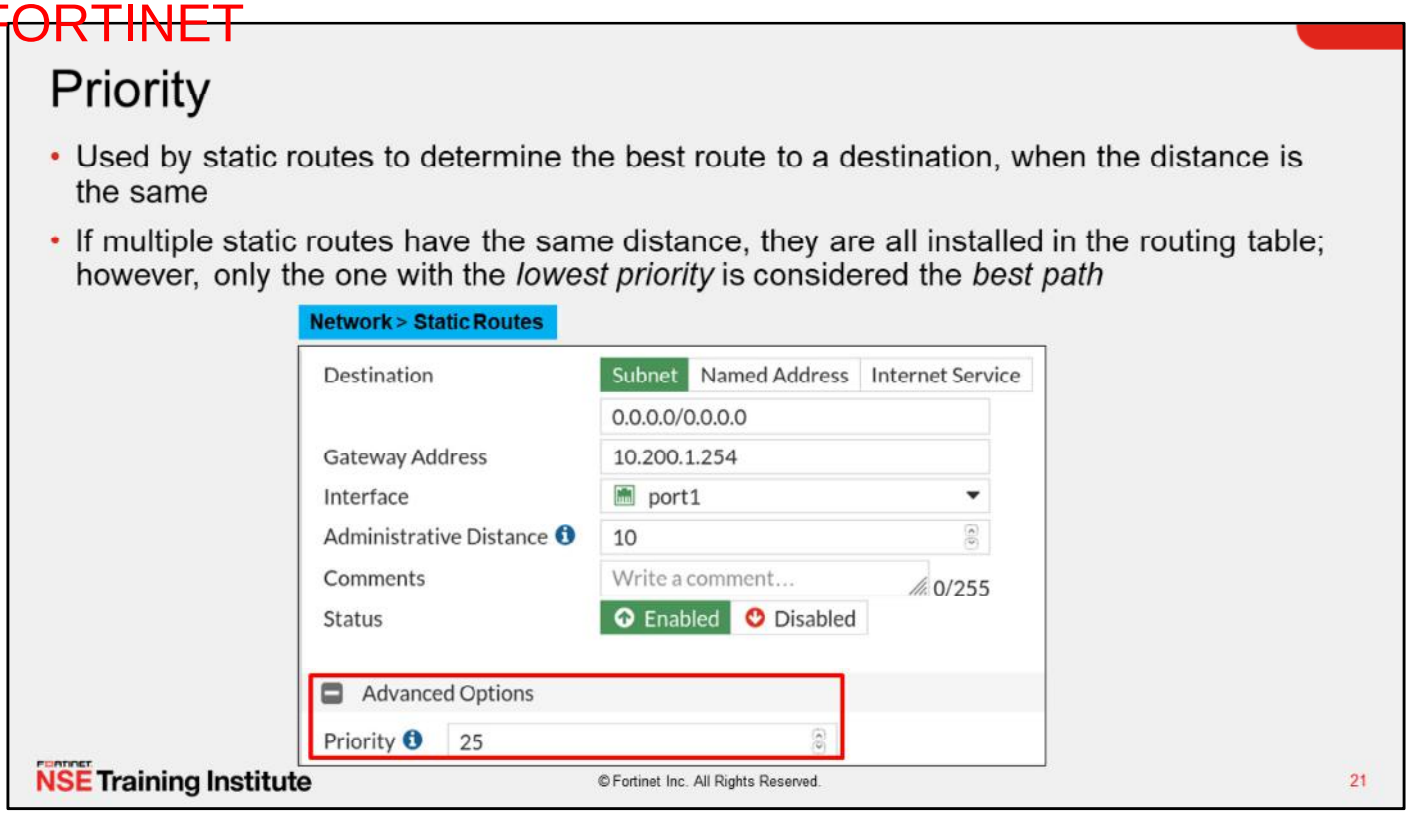

When multiple static routes have the same distance value, they are both installed in the routing table. So which route will be used to route matching packets? In the scenario shown on this slide, FortiGate uses the priority value as a tiebreaker to identify the best route. Routes with lower priority are always preferred.

The priority attribute is applicable only to static routes, and is configured under the **Advanced Options** on the GUI. By default, all static routes have a priority of `0`.

Priority values are viewed in the static route configuration, and on the routing table on the CLI, which you will learn about later in this lesson. They are not displayed on the GUI routing table.

Knowledge Check

1. The **Priority** attribute applies to which type of routes?
   - ✓ A. Static
   - B. Dynamic

2. Which attribute does FortiGate use to determine the *best* route for a packet, if it matches multiple dynamic routes that have the same **Distance**?
   - A. Priority
   - ✓ B. Metric

3. Which static route attribute does not appear on the GUI routing monitor?
   - A. Distance
   - ✓ B. Priority

NSE Training Institute

22

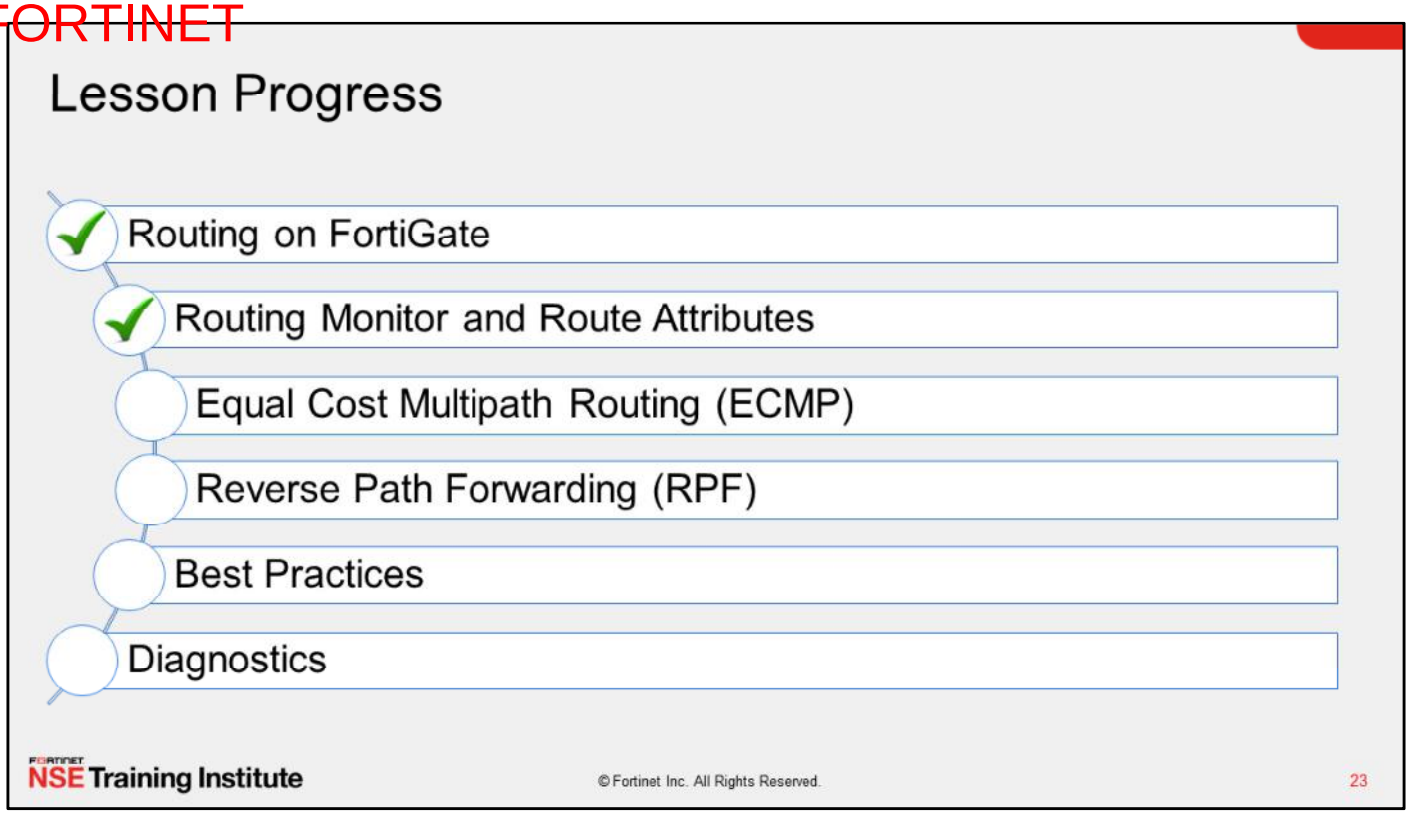

Good job! You now understand the routing monitor and route attributes.

Now, you will learn about ECMP routing.

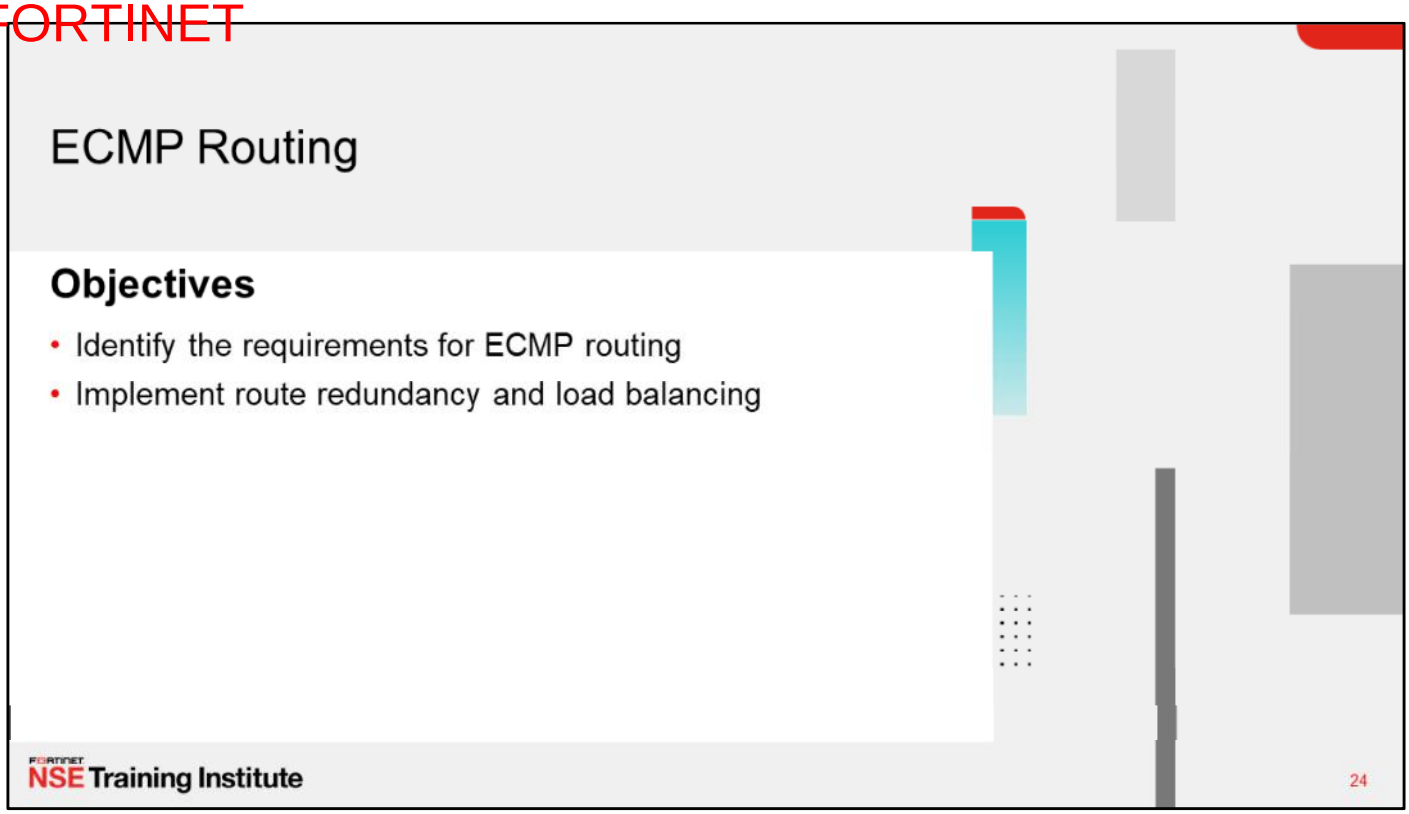

After completing this section, you should be able to achieve the objectives shown on this slide.

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

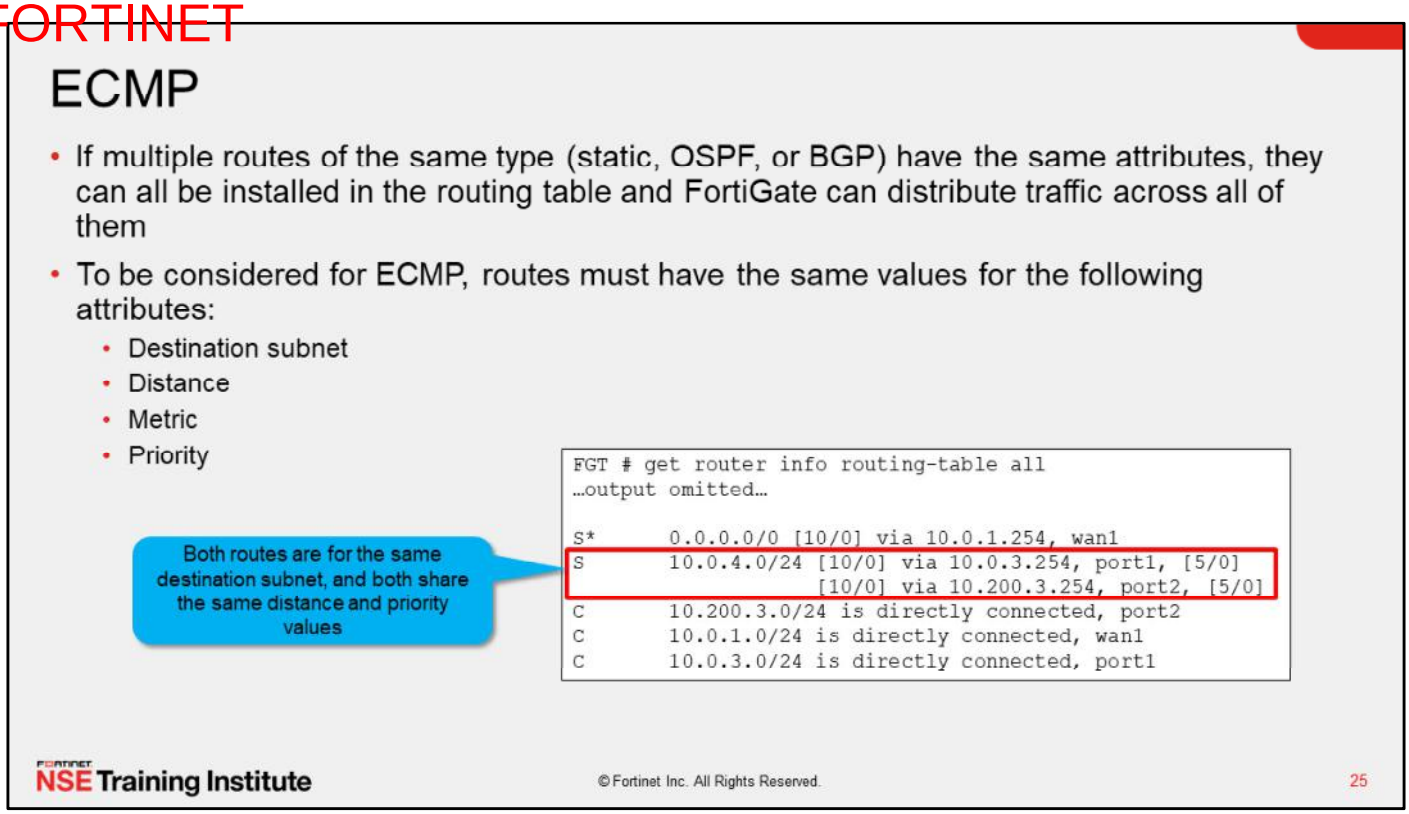

So far, you've learned about the different route attributes available for routers to identify the best route to a destination. So, what happens when two or more routes of the same type and to the same destination share the same values for all of the attributes? All routes are installed in the routing table and FortiGate will load balance traffic across all routes. This is called equal cost multi-path (ECMP).

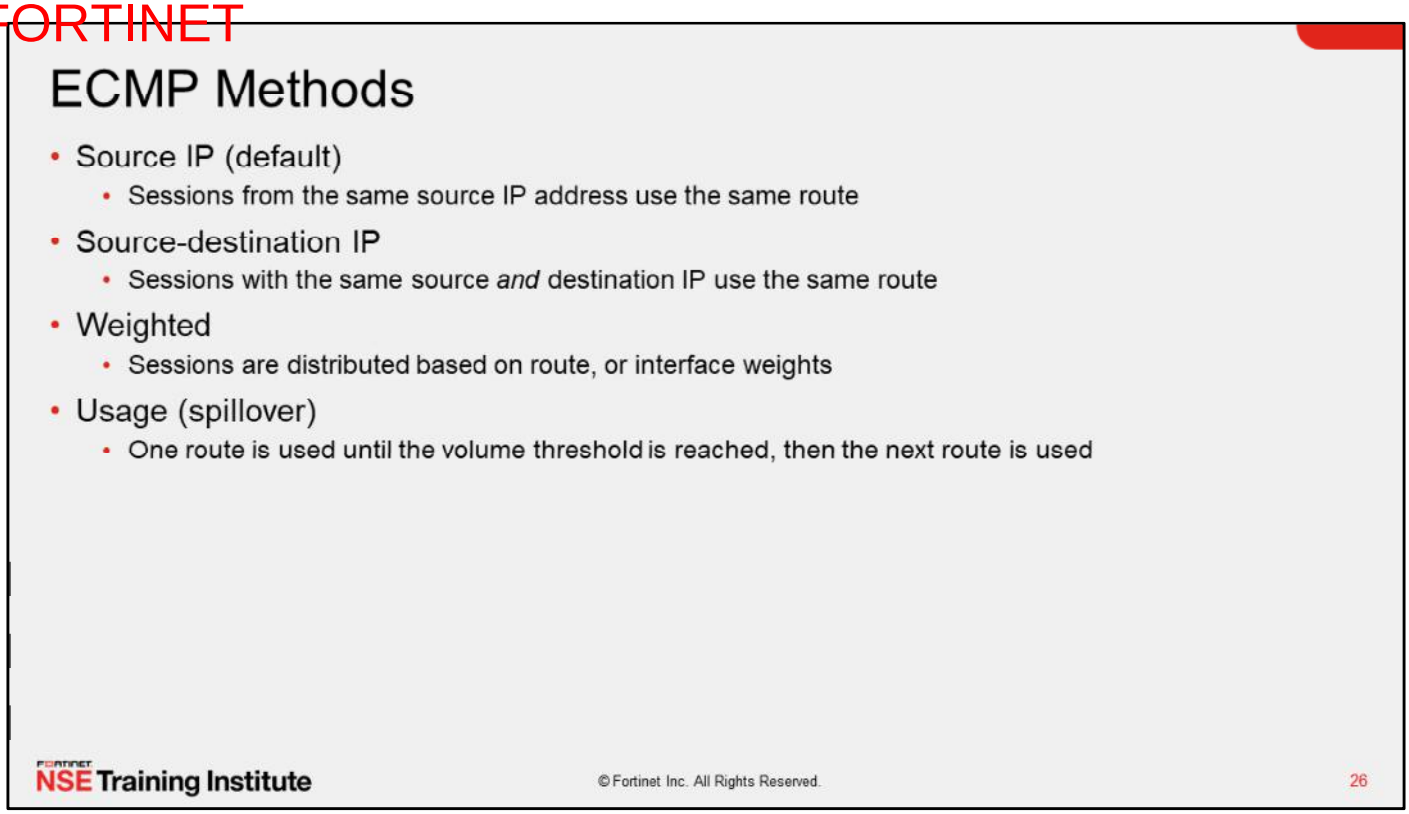

ECMP can load balance traffic using one of the following four methods: sessions can be balanced among equal routes depending on the source IP address, source and destination IP addresses, route or interface weights, or interface volume thresholds.

When using the source IP method, all traffic originating from the same source IP is expected to use the same path. The source-destination IP method works similarly, but it also factors in the destination IP. So, sessions from a specific source to a specific destination are expected to use the same path.

With the ECMP load balancing method set to weighted, FortiGate distributes sessions with different destination IPs by generating a random value to identify the route to select. The probability of selecting one route over another is based on the weight value of each route or interface. Higher weights are more likely to be selected.

There is an additional method called usage-based (or spillover). In usage-based load balancing, FortiGate uses a primary route until a traffic volume threshold is reached; after that, it uses the next available route.

If one of the ECMP routes fails and is removed from the routing table, the traffic is routed over the remaining routes. There is no specific configuration necessary for route failover.

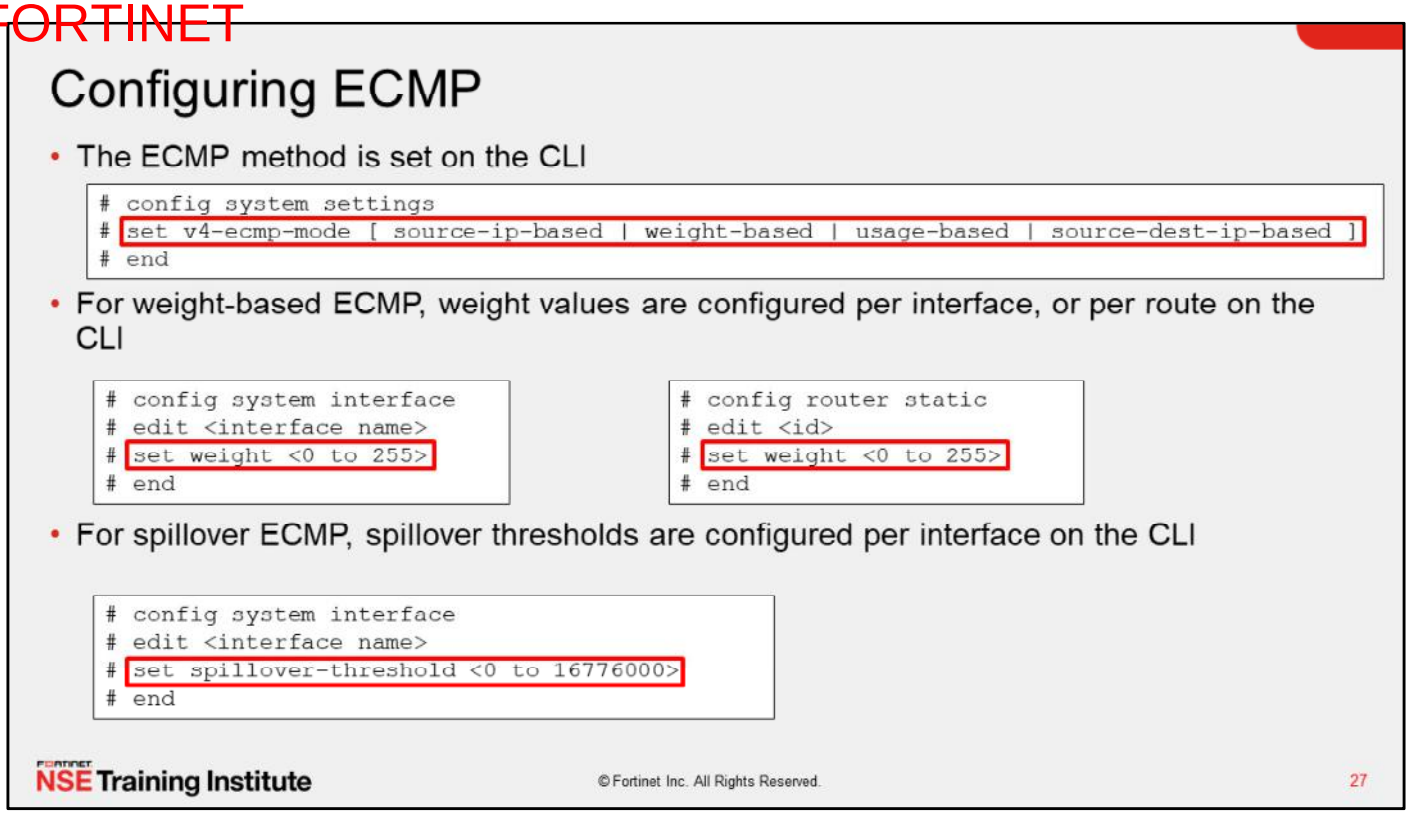

FortiGate uses `source-ip-based` as the default ECMP method. You can change this setting on the CLI using the commands shown on this slide.

For spillover-based ECMP, you must configure additional settings at the interface level. For weight-based ECMP, you must assign weight values to interfaces, or routes. You can do this on the CLI using the commands shown on this slide.

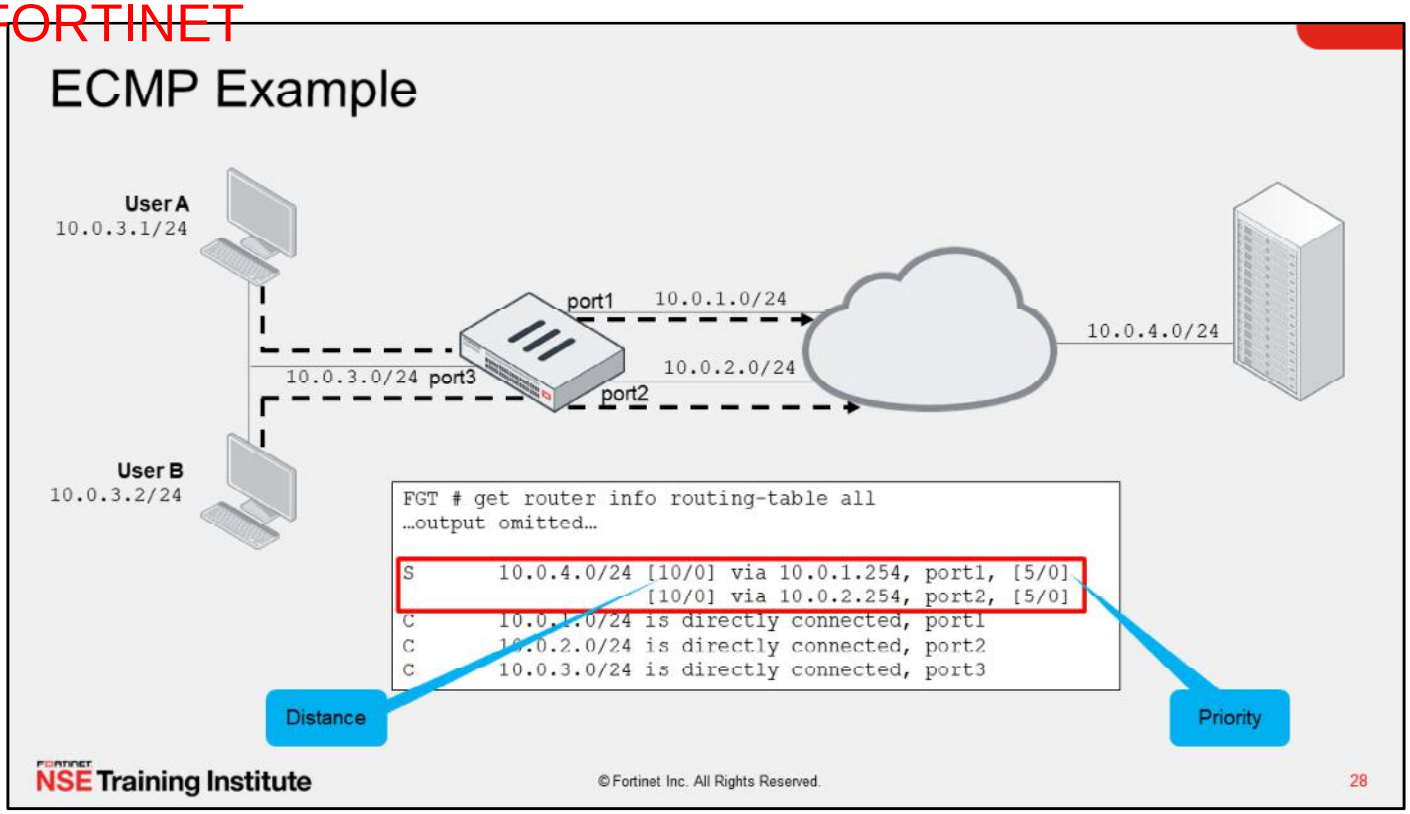

In the scenario shown on this slide, FortiGate has two equal candidate routes for the `10.0.4.0/24` subnet using port1 and port2 respectively. Using the default source-based ECMP method, FortiGate may use either route to deliver traffic destined for the `10.0.4.0/24` subnet from User A and User B. If port1 loses connectivity, FortiGate automatically uses port2 to deliver all traffic destined for the `10.0.4.0/24` subnet.

ECMP allows you to maintain multiple links for the same destination, as well as provide built-in failover. You can deploy this for any network resources that have high bandwidth demands, and are mission critical. Employing ECMP for these resources allows you to aggregate the available bandwidth of multiple links, and load balance traffic across those links.

When using ECMP, you must have the correct firewall policies in place to allow traffic to egress all interfaces participating in ECMP.

While you can use ECMP to maintain multiple internet (WAN) connections on FortiGate, it can be more efficient to use the software-defined WAN (SD-WAN) feature to accomplish this. You can still use ECMP for internal resources.

Knowledge Check

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

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

NSE Training Institute

29

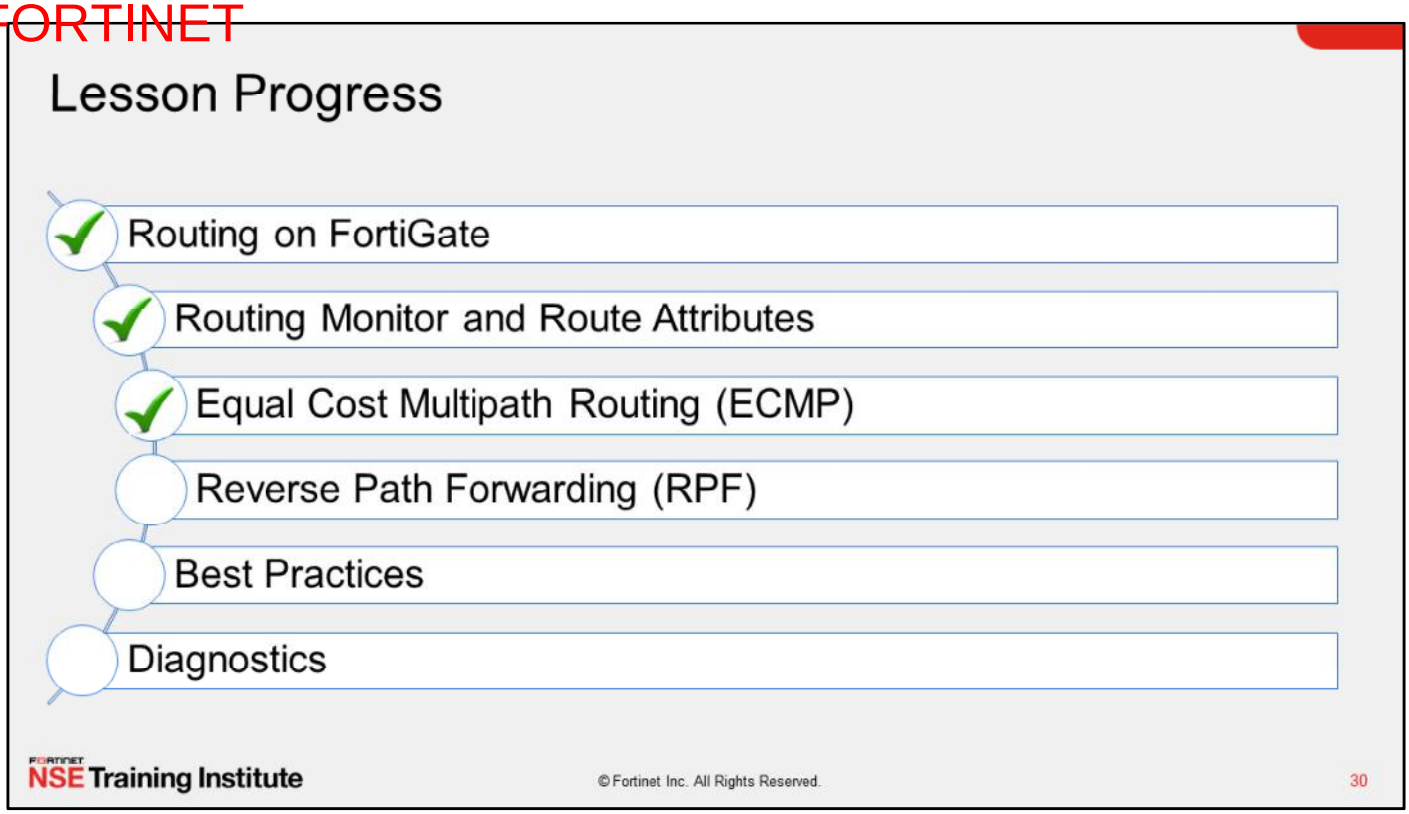

Good job! You now understand ECMP routing.

Now, you will learn about reverse path forwarding.

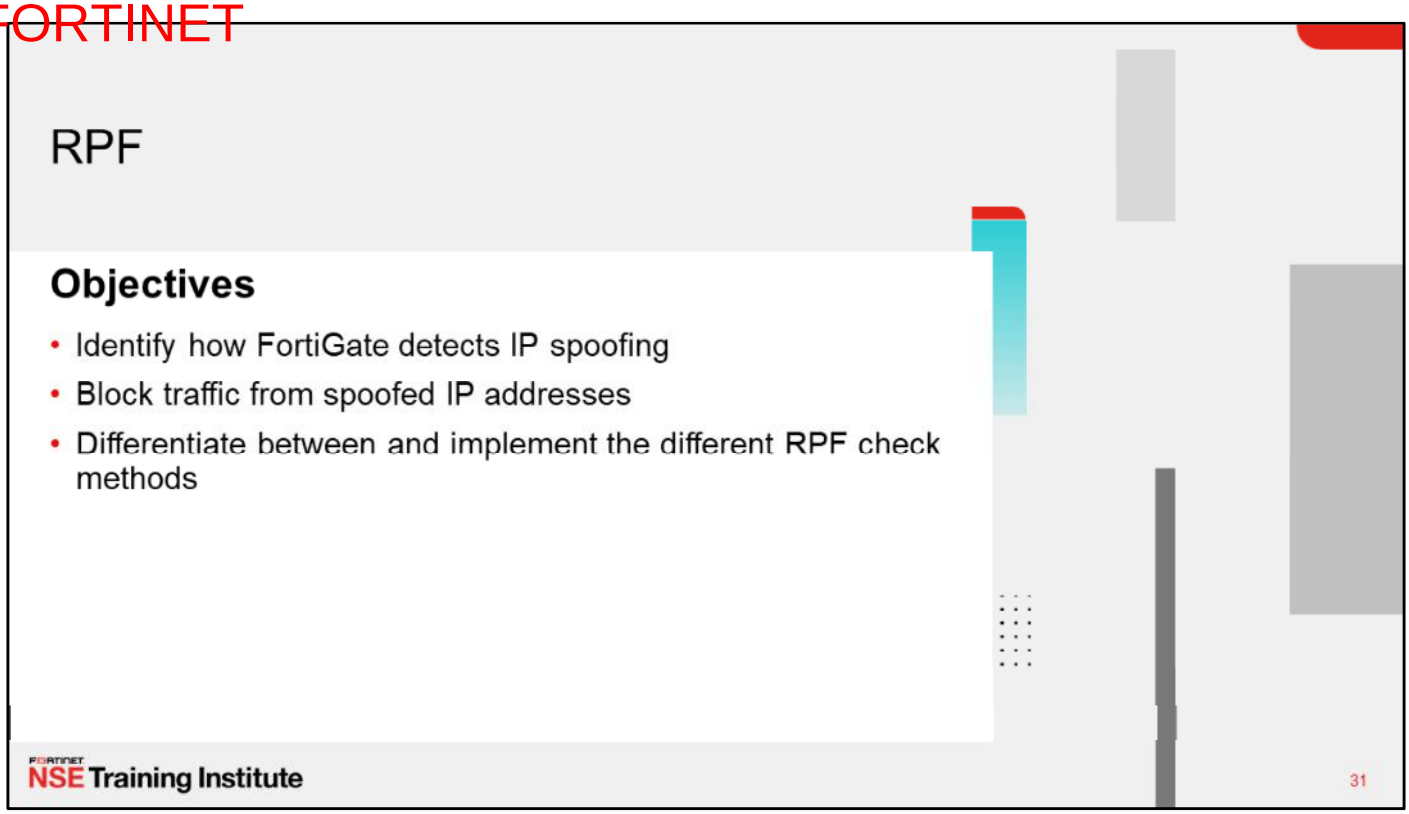

After completing this section, you should be able to achieve the objectives shown on this slide.

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

RPF

- Protects against IP spoofing attacks
- The source IP address is checked against the routing table for a return path
- RPF is only carried out on:
  - The first packet in the session, not on a reply
- Two methods:
  - Loose
  - Strict

NSE Training Institute

32

Packets are sometimes dropped for routing and security reasons. RPF is a mechanism that protects FortiGate and your network from IP spoofing attacks. It checks if there is a route back to the packet source. This check is run on the first packet of any new session.

There are two RPF methods: loose and strict.

RPF Checking

- RPF checks for an active route back to the source IP through the incoming interface
  - User A traffic is *accepted* because there is an active route (the default route) back to the source
  - User B and C packets are *denied* because there are no active routes back to those sources

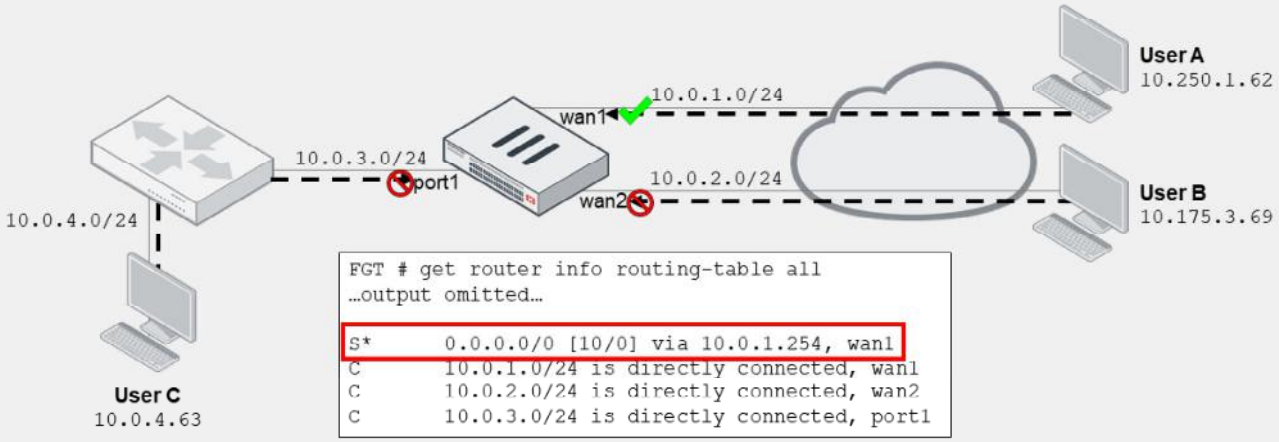

NSE Training Institute

33

In the scenario shown on this slide, incoming internet traffic arriving at `wan1` is accepted because the default route is a valid route back to the source.

However, there are two interfaces that do not route some incoming traffic: `port1` and `wan2`. `port1` do not route traffic because the subnet for user C is `10.0.4.0/24`. There is no active route for that subnet through `port1`. So, traffic coming from `10.0.4.0/24` to `port1` is dropped because it failed the RPF check.

The other interface that does not route traffic is `wan2`. While `wan2` is physically connected to the internet, the only IP addresses that are valid as sources or destinations for `wan2` are those in the `10.0.2.0/24` subnet. So, incoming traffic from any other source does not pass the RPF check and is dropped.

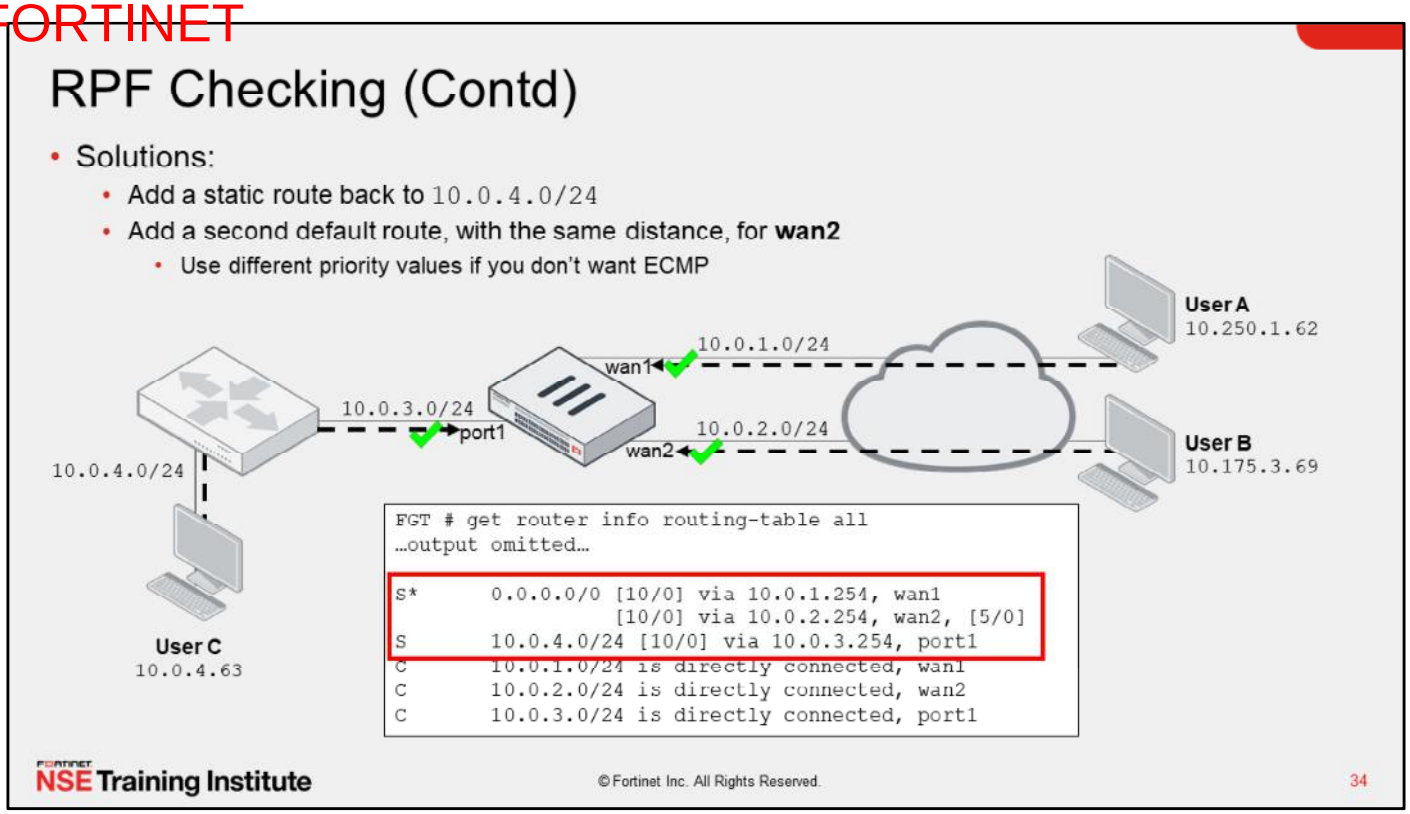

So, how can you fix this problem?

The first problem is fixed by adding a static route to `10.0.4.0/24 through port1`. Now, when FortiGate runs the RPF check for User C's packets, it finds an active route to that subnet through `port1` and the packet is accepted.

The second problem is also fixed by adding a static route. In this case, it is a default route for `wan2`. This second default route must have the same distance as the default route for `wan1`. This ensures that both routes are active in the routing table. They both can have different priorities, but they must have the same distance to be active.

This slide shows an example of when two routes with the same distance, but different priorities, are required. So, one route is the best (the one with the lowest priority), but both are active. The best route is used for outbound traffic, but both can receive incoming connections without failing the RPF check.

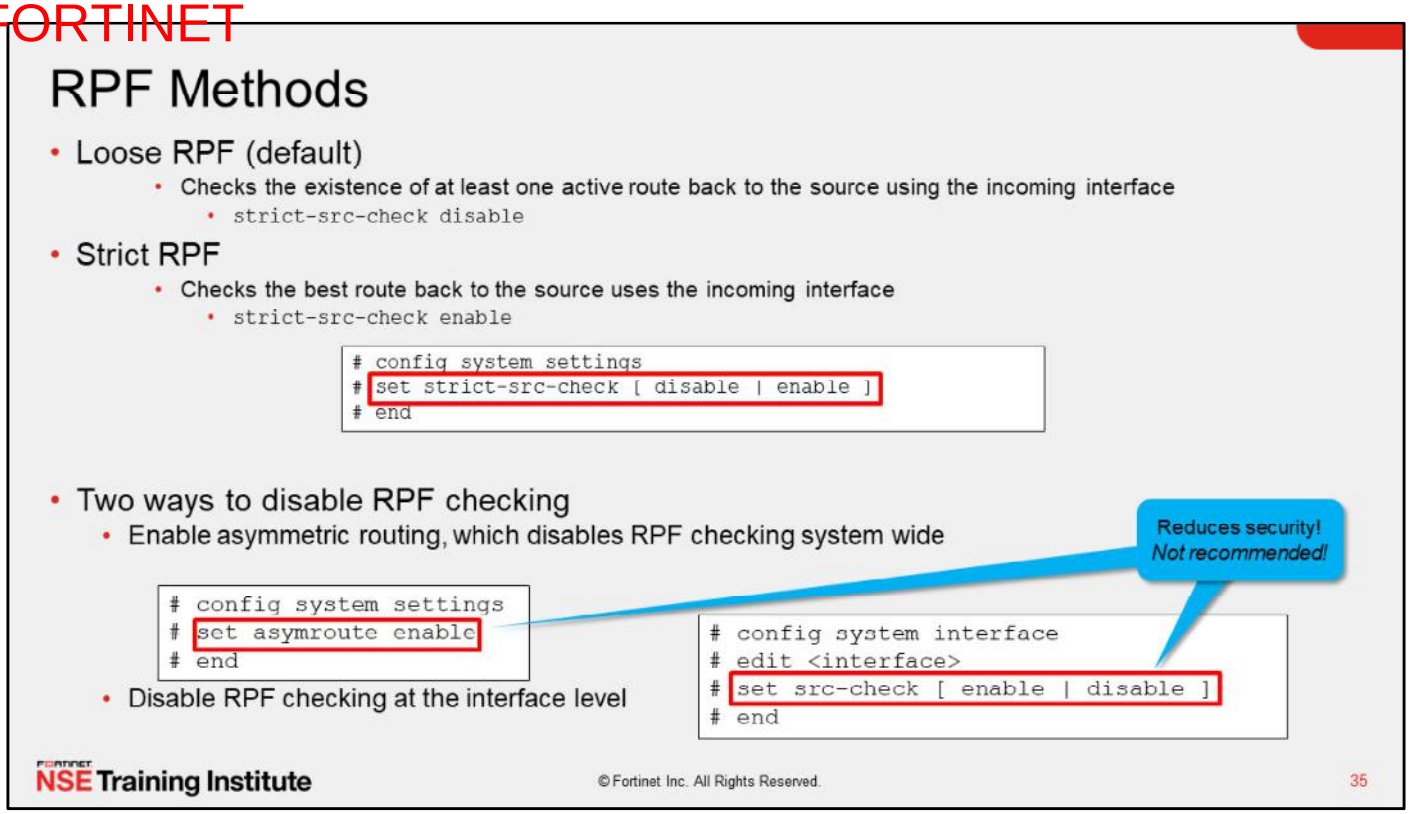

FortiGate can either strictly or loosely enforce RPF.

In loose mode, the packet is accepted as long as there is one active route to the source IP through the incoming interface. It does not have to be the best route, just an active one.

In strict mode, FortiGate checks that the best route to the source IP address is through the incoming interface. The route not only has to be active (as in the case of loose mode), but it also has to be the best.

You can disable RPF checking in two ways. If you enable asymmetric routing, it disables RPF checking system wide. However this reduces the security of your network. Features, such as antivirus and IPS become noneffective. So, if you need to disable RPF checking, you can do so at the interface level using the commands shown on this slide.

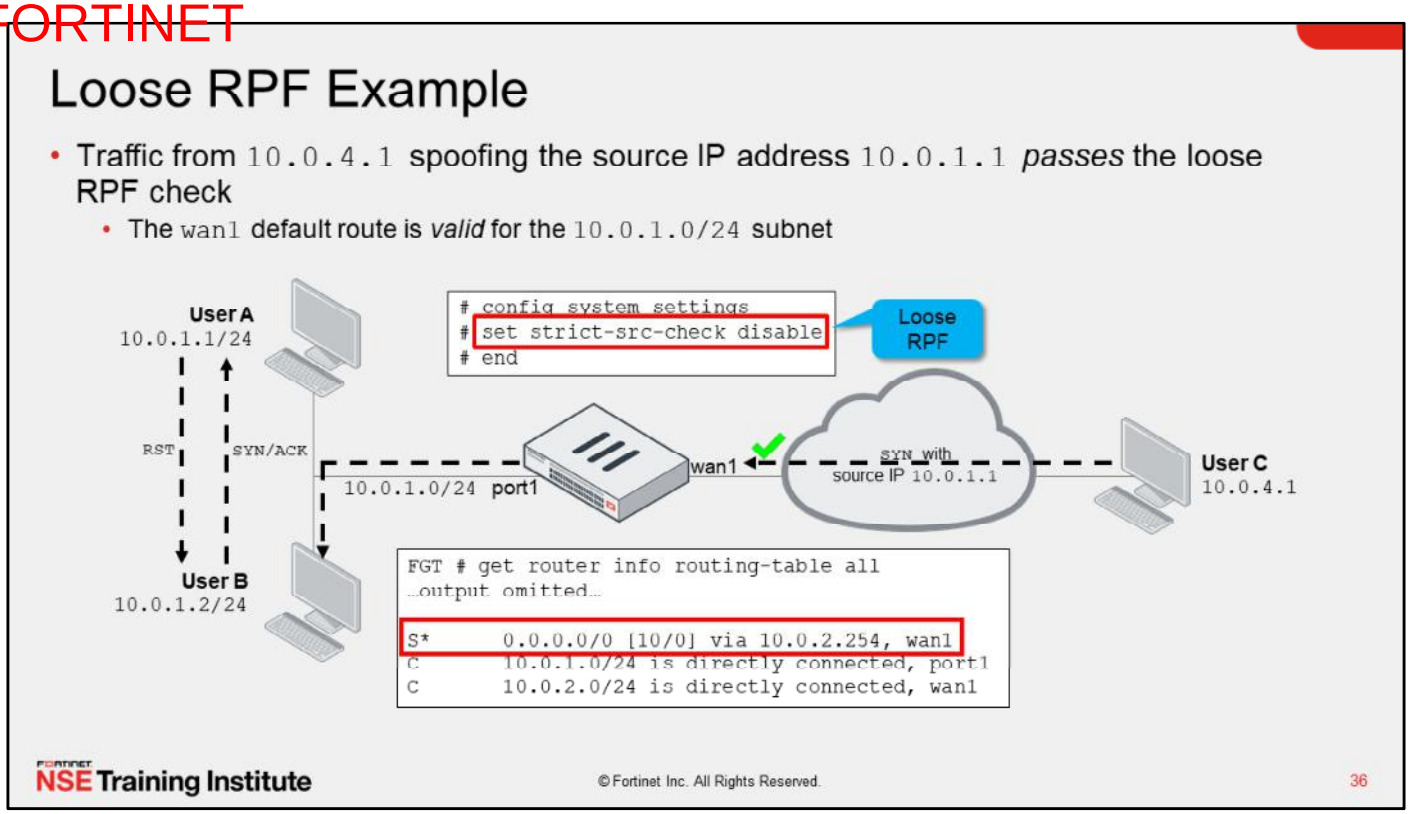

In the example shown on this slide, `10.0.4.1` sends a `SYN` packet to `10.0.1.2`, but spoofs a source IP of `10.0.1.1`. This makes the packet appear to be initiated from the internal network behind `port1`. Loose RPF *allows* this traffic because the active route on `wan1` is a default route (`0.0.0.0/0`).

Next, `10.0.1.2 (User B)` sends the `SYN/ACK` packet to the *real* device with the IP address of `10.0.1.1 (User A)`.

Since `10.0.1.1 (User A)` is not expecting any `SYN/ACK` packets (because it has not previously sent any `SYN` packet to `10.0.1.2`), it replies with an `RST` (reset) packet.

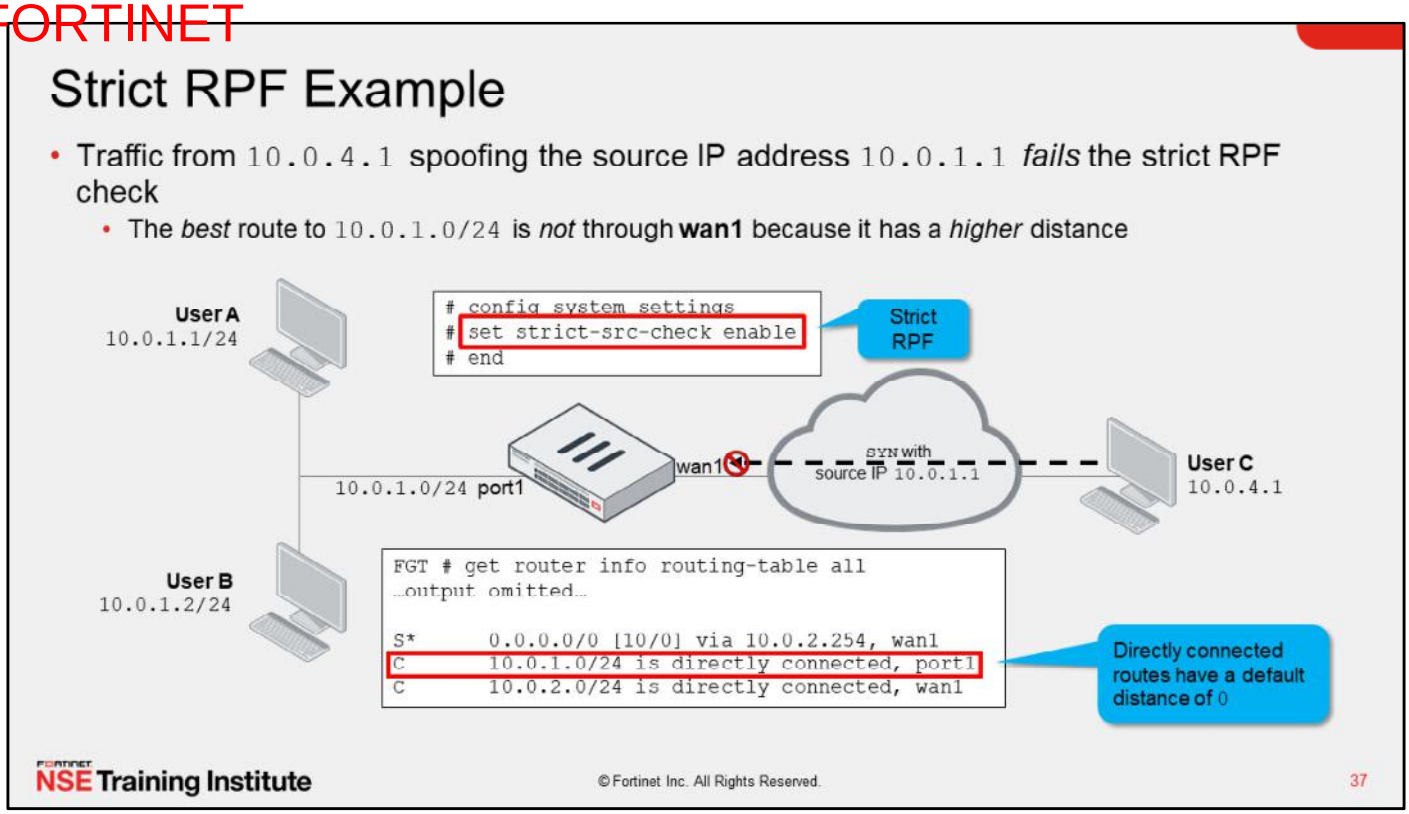

So, what happens in the same scenario when strict RPF is used?

Strict RPF drops the `SYN` packet. Even though the `wan1` default route is an active route for the `10.0.1.0/24` subnet, it's not the *best* route. The best route is through the `port1` interface because it has a *lower* distance value. Remember, the default distance value for directly connected routes is `0`, which is lower than the distance value of `10` of the default route.

Although strict RPF is more secure, it can cause false positives if you use dynamic routing. Dynamic routes can change quickly, and they could cause FortiGate to drop legitimate packets each time the preferred route changes. In general, it is recommended to use loose RPF in combination with firewall policies that block spoofed traffic, instead of using strict RPF for that purpose.

Knowledge Check

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

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

NSE Training Institute

38

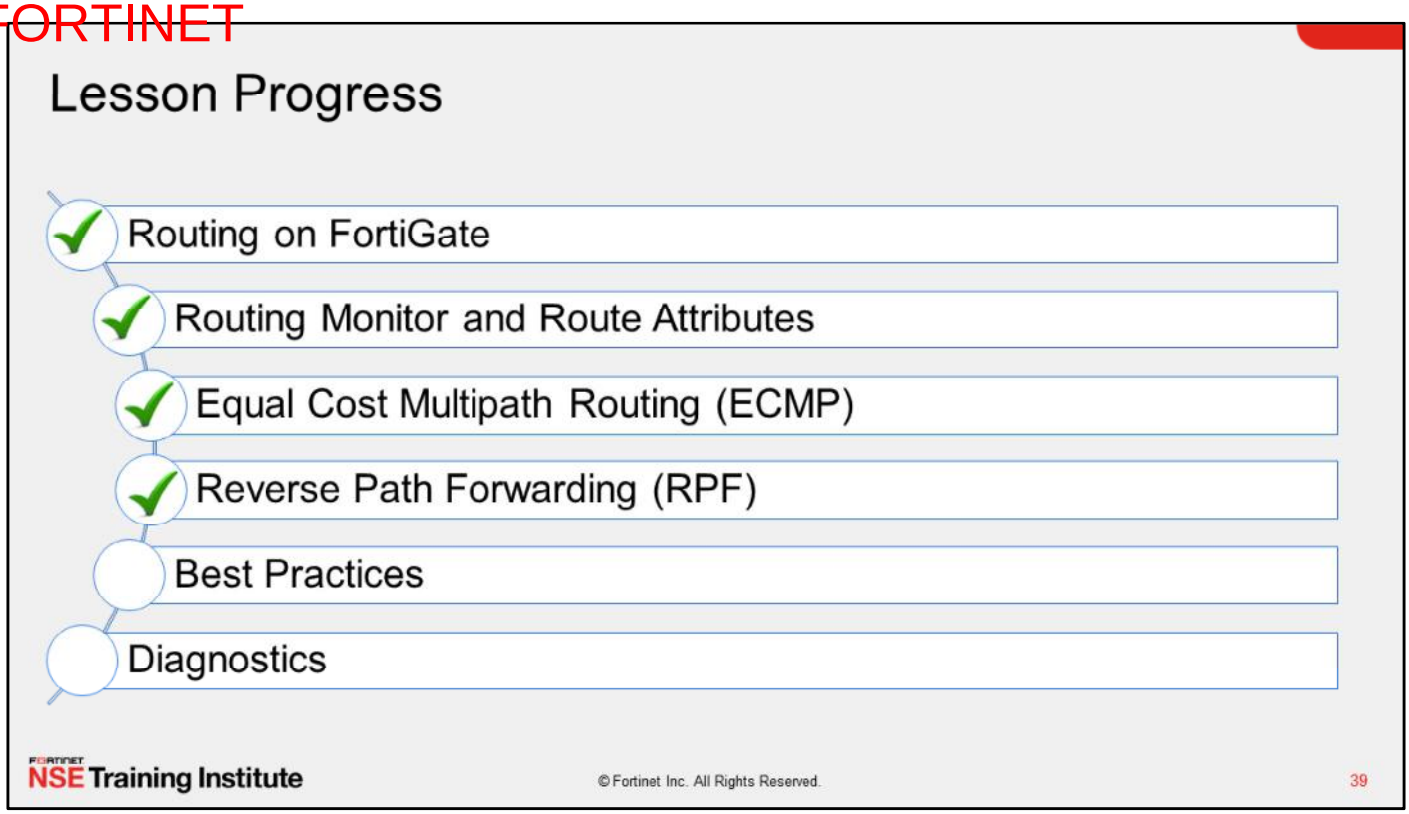

Good job! You now understand RPF.

Now, you will learn about routing best practices.

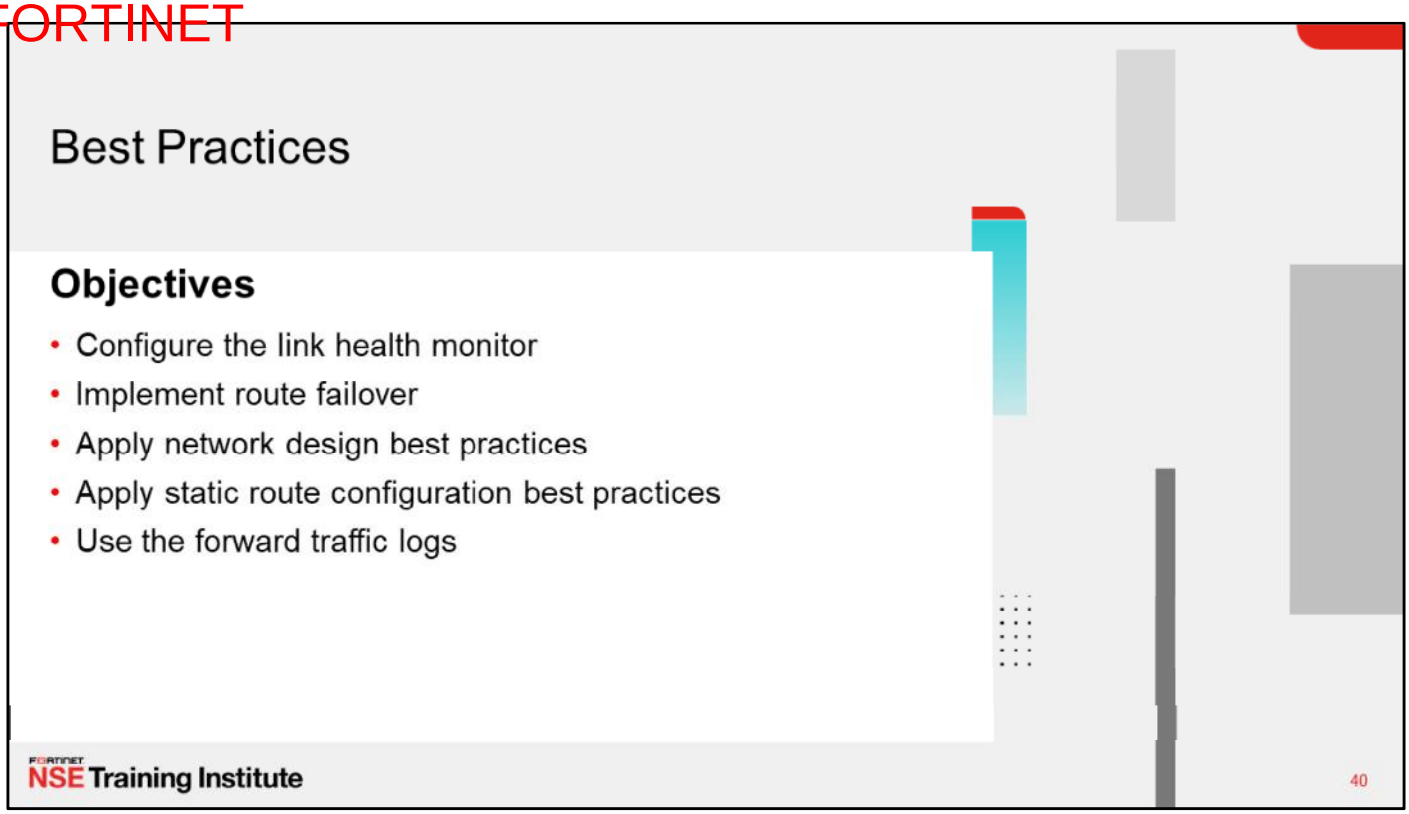

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in routing best practices, you should be able apply them in your own network and routing designs to maintain an effective and efficient routing configuration.

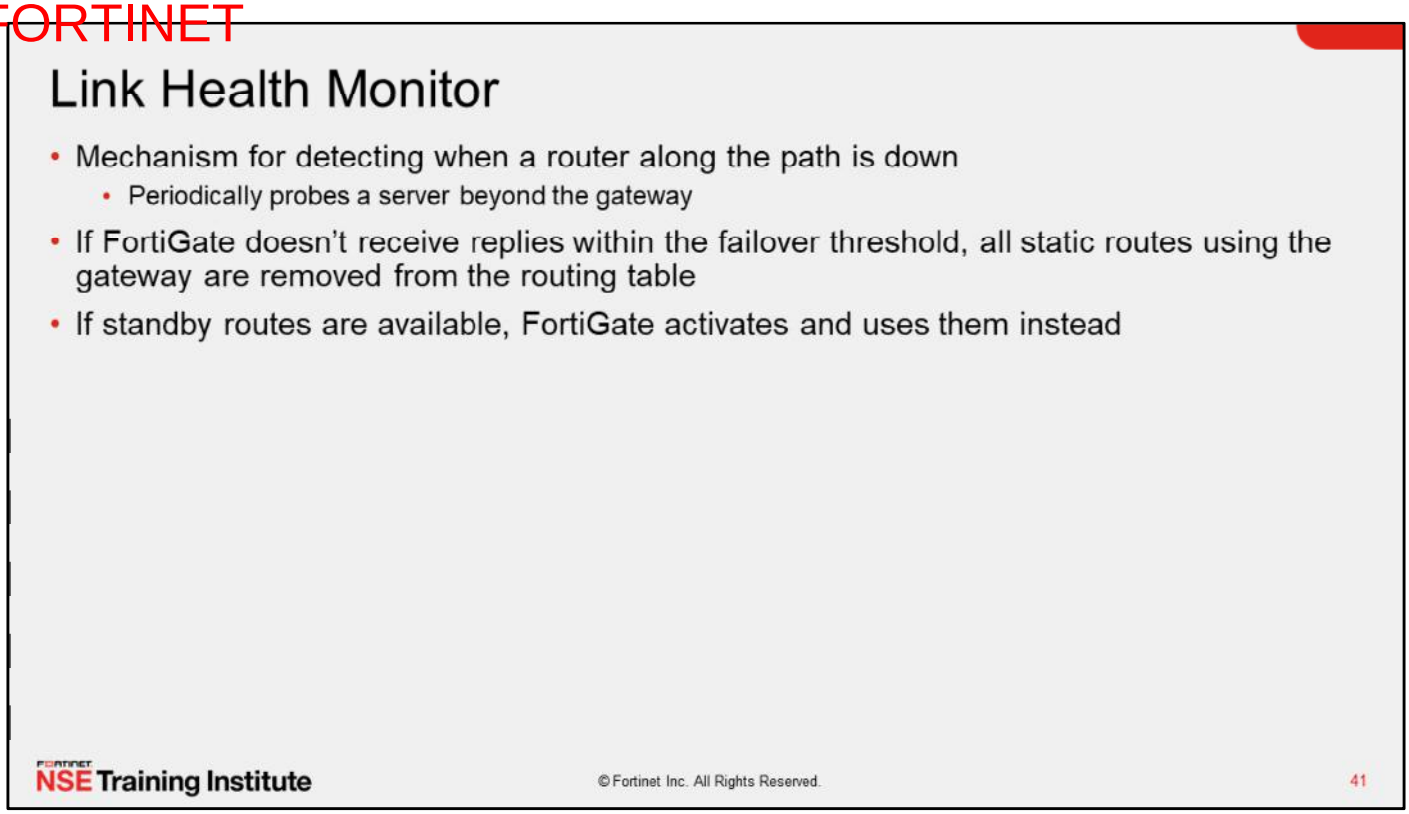

When using ECMP routing, you don't have to implement anything extra for route failover. Because of the design of ECMP, route failover happens automatically. How do you implement route failover when you're not using ECMP?

The link health monitor is a mechanism for detecting when a router along the path is down. It is often used where there are redundant routers onsite, such as for dual ISP links. When configured, FortiGate periodically sends probing signals through one of the gateways to a server that acts as a beacon. The server can be any host that should normally be reachable through that path. Usually, it's best to choose a stable server with robust infrastructure, and to choose the protocol to which the server would normally respond.

If FortiGate stops receiving a reply from the server, all the routes using that gateway are removed from the routing table. Alternatively, you can configure the device to administratively bring down an interface, so all routes using that interface are removed. While a monitored route is removed, FortiGate continues to send link health monitor signals. As soon as FortiGate receives a reply, it reactivates the associated routes.

It might be useful to choose a server that is indirectly attached, located one or two hops beyond the FortiGate's gateway. This does not exactly test availability of this one gateway, but rather the combination of gateways. That way, FortiGate will accurately indicates availability of services and subsequent hops.

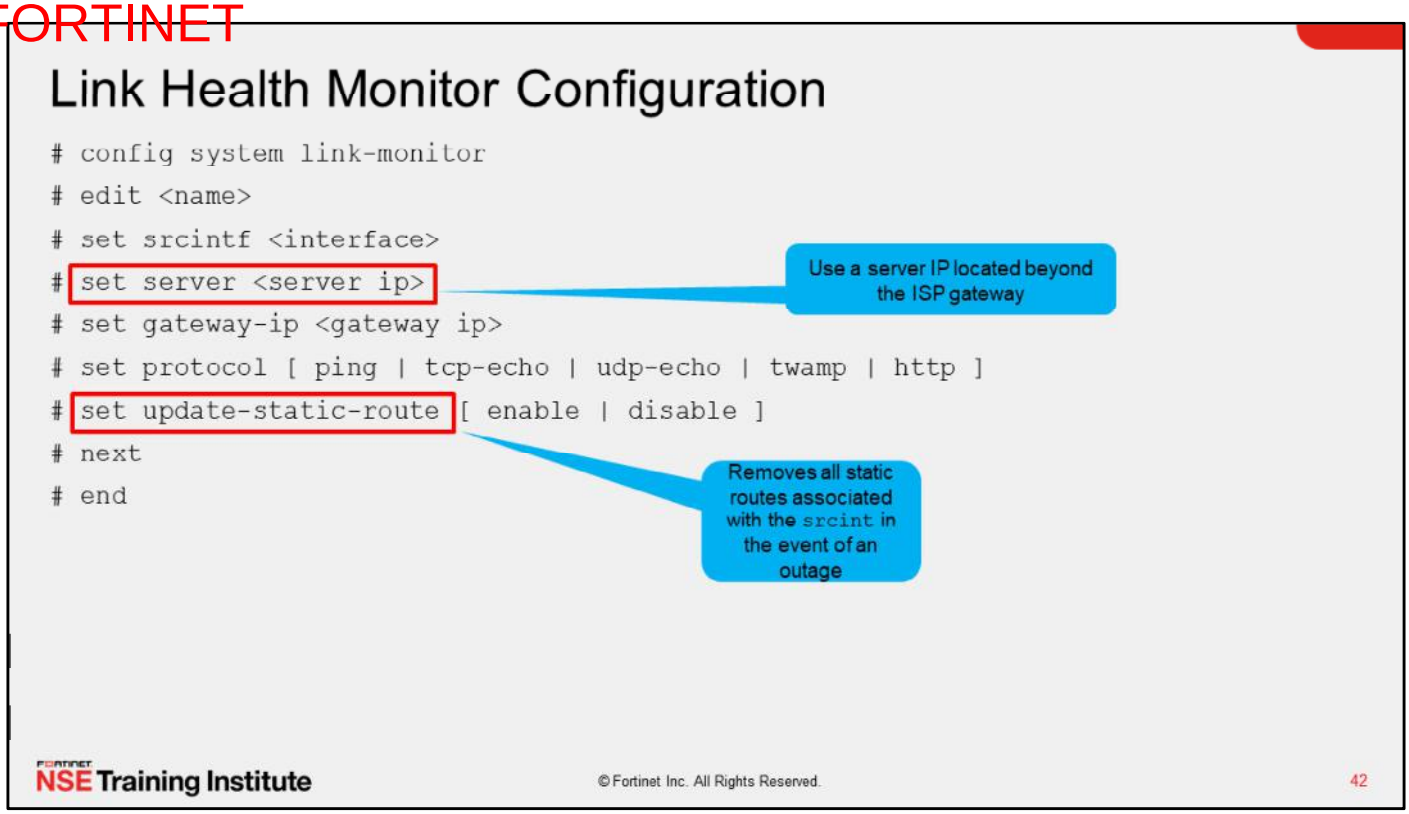

You configure the link health monitor on the CLI.

You must set the egress interface, the IP address of the gateway router, and the IP address and protocol (`http`, `ping`, `udp-echo`, `tcp-echo`, or `twamp`).

You must enable the `update-static-route` setting to ensure that FortiGate removes any matching static route, in the event the link health monitor detects an outage. This allows any secondary route configured with a higher distance to be activated.

You can configure multiple link health monitors, for example, one for each ISP.

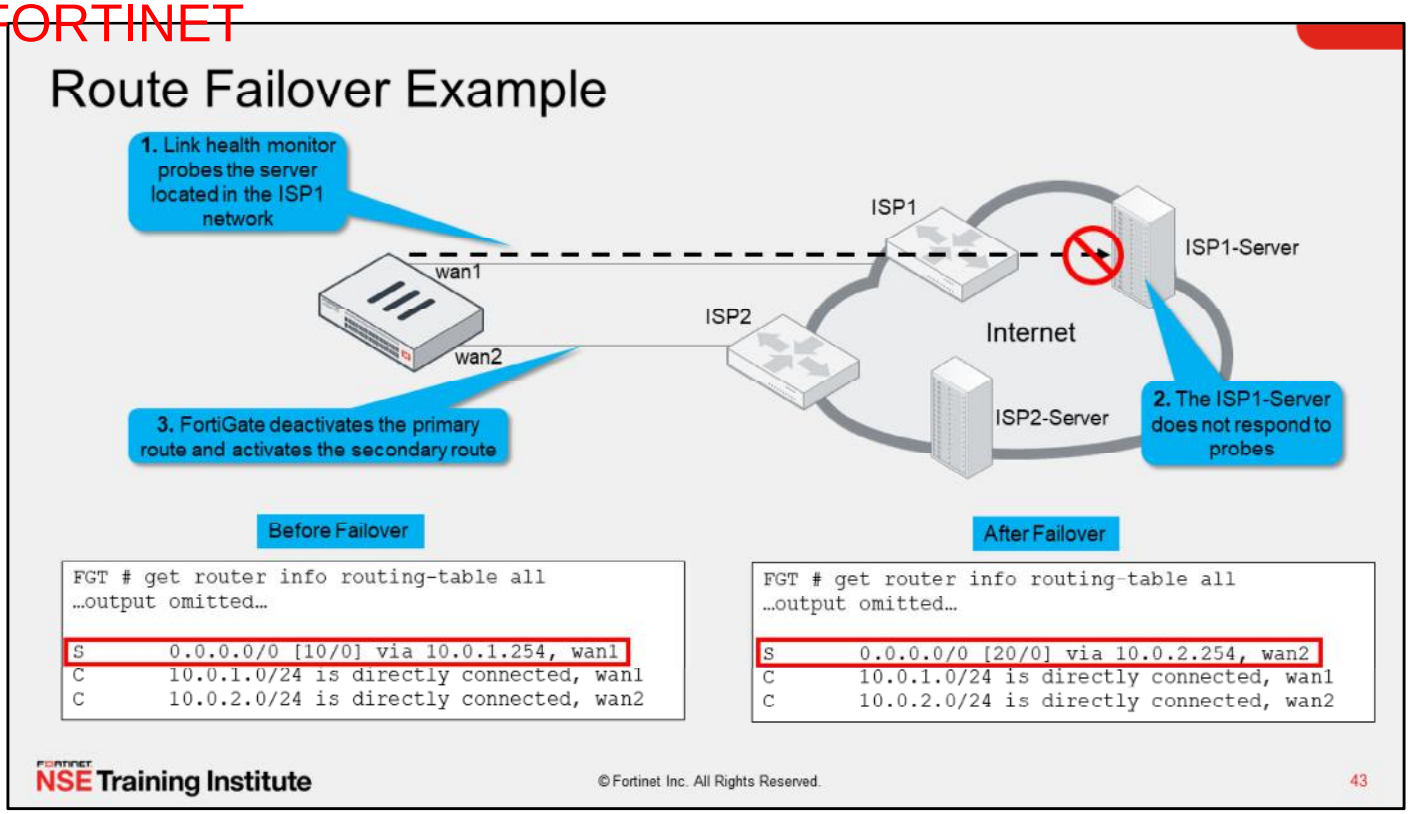

In the example shown on this slide, FortiGate has two ISP connections: ISP1 and ISP2. Within each ISP's network, there are servers that the FortiGate is probing.

Before failover, the `wan1` default route is active in the routing table because it has the lower distance. The `wan2` default route is configured with a higher distance and, therefore, standby. The link health monitor probes the ISP1-Server located within the ISP1 network through the `wan1` interface. When the ISP1-Server does not respond to the probing attempts, FortiGate removes the primary route from the routing table and, because there is a second default route through `wan2`, activates it to route traffic to the internet.

In this dual-ISP scenario, your first reaction might be to implement ECMP routing. After all, ECMP would allow you to utilize both ISP links at the same time, which increases the available bandwidth for internet traffic. However, it might not always be feasible to implement ECMP. Most of the time it is related to cost considerations—the secondary ISP might be charging based on bandwidth or data usage. This is where you can use the distance attribute, in conjunction with link health monitoring, to implement route failover.

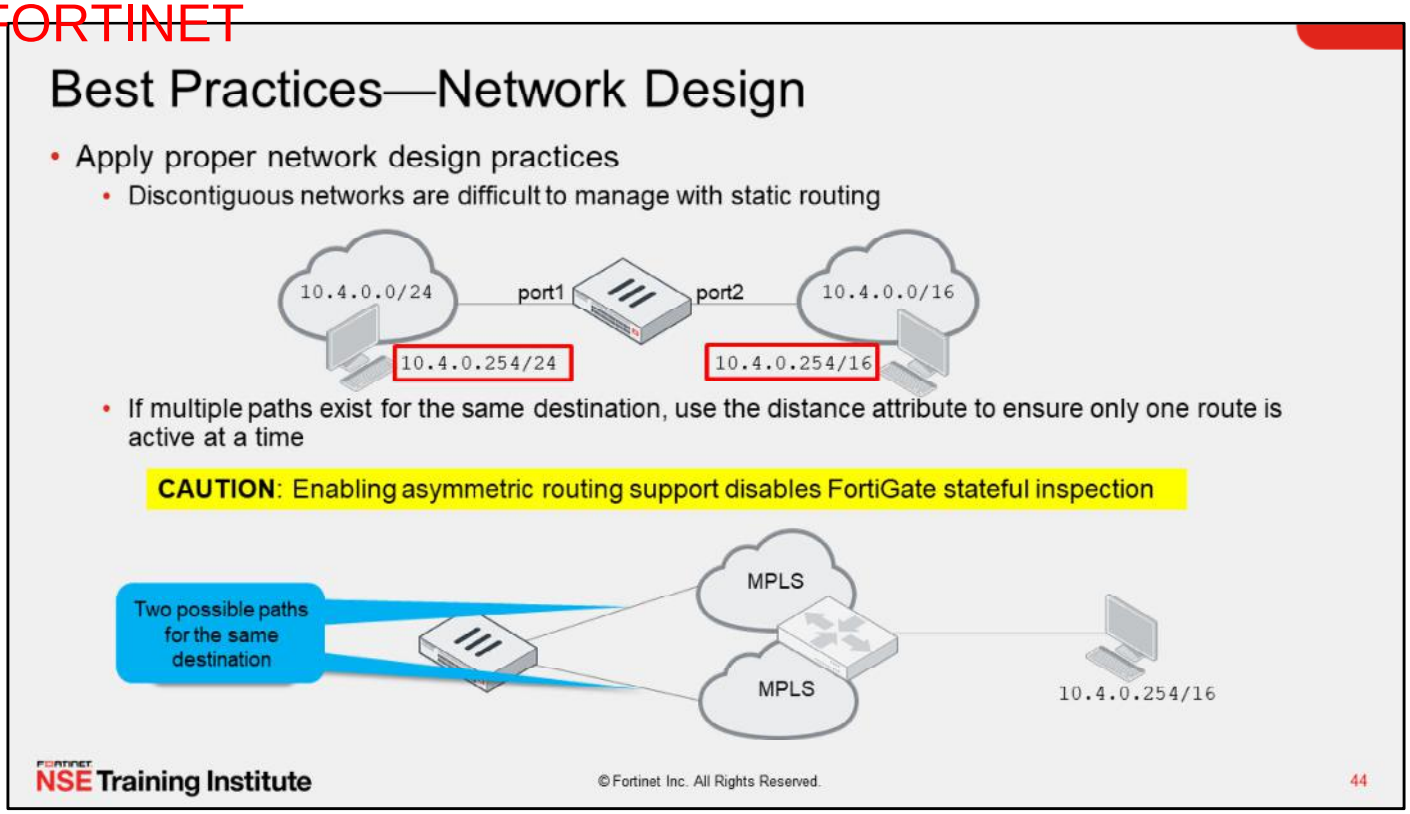

Routing best practices start at the network design phase. One of the biggest challenges with static routing is figuring out how to handle discontiguous networks. While dynamic routing protocols are better equipped to handle them, they still create issues with large routing tables and route summarization. In some situations, you cannot handle discontiguous networks with static routing alone, and may need to resort to NAT to allow traffic.

Another challenge to consider is asymmetric routing. Asymmetric routing is a situation where packets, in the same session, might flow through different routes to reach the destination. So, if multiple paths exist in your network for the same destination, consider using the distance attribute to ensure only one route is active at a time. You can also consider using ECMP; however, the effectiveness of this also depends on whether or not the remote side router is also capable of ECMP, or some form of session persistence. In other words, the remote side router also must send the reply packets back through the same path.

While you can configure the FortiGate to allow asymmetric routing, it is *highly* discouraged to do so. Enabling asymmetric routing disables the FortiGate stateful inspection capability. Antivirus and intrusion prevention will not be effective because FortiGate is unaware of sessions and treats each packet individually. Disable the RPF check at the interface level.

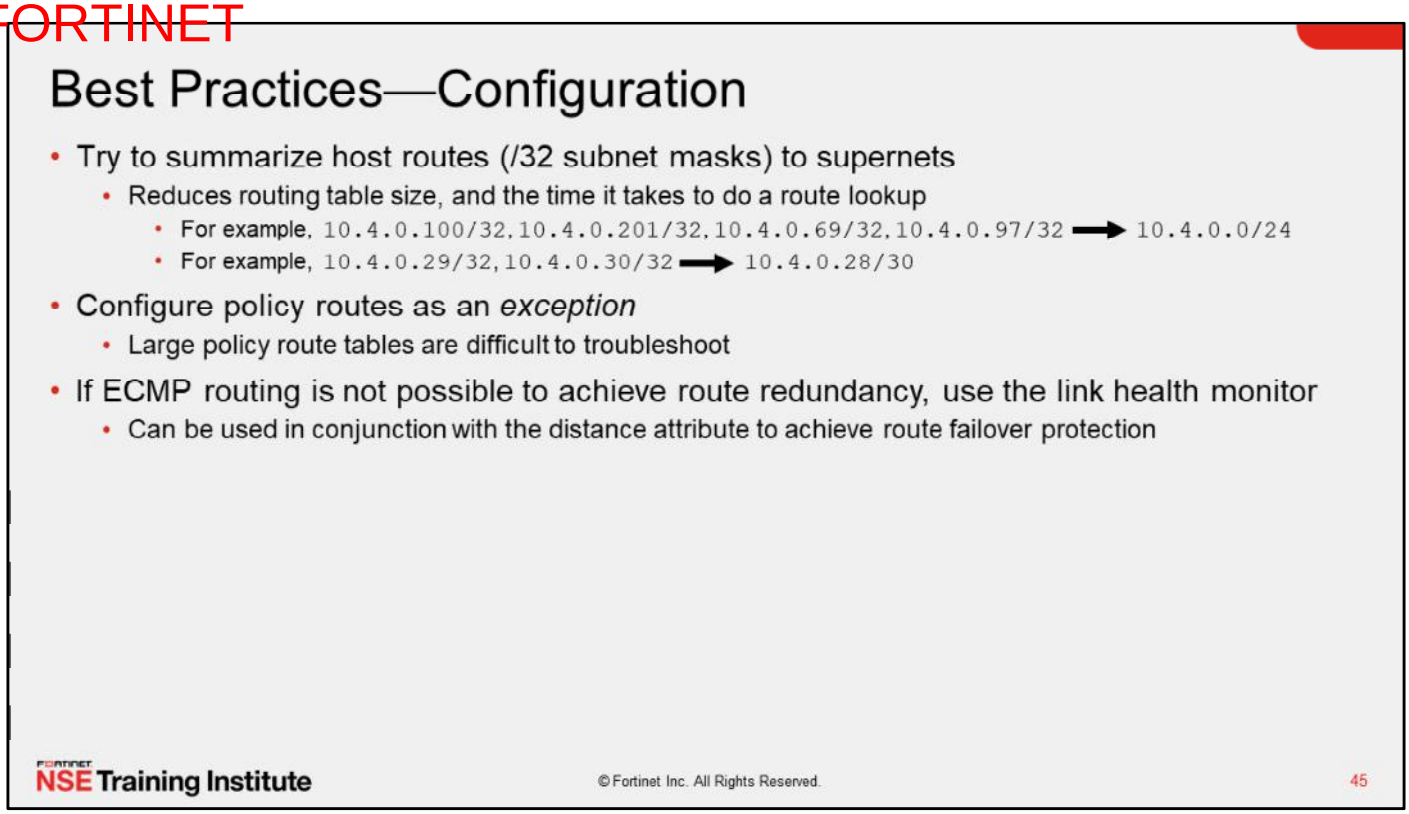

If you find yourself creating multiple host routes (/32 subnet mask), investigate whether you can summarize them in a supernet. Dynamic routing protocols are designed to summarize contiguous networks to keep routing tables small, reducing the size of routing updates and the time it takes to do a route lookup. You should employ the same methodology when creating static routes.

When configuring policy routes, treat them as an *exception* to the routing table. If you find yourself continuously having to create policy routes, you should re-evaluate your static route configuration to see if you can make adjustments there first. Remember, policy routes override the routing table. So, the only way to override a policy route is by configuring another one. If you don't plan your policy route configuration, it can quickly become an issue. Also, large policy route tables are difficult to troubleshoot.

Finally, if ECMP routing is not possible with multiple routes, use the link health monitor in conjunction with the distance attribute, to ensure you have route failover.

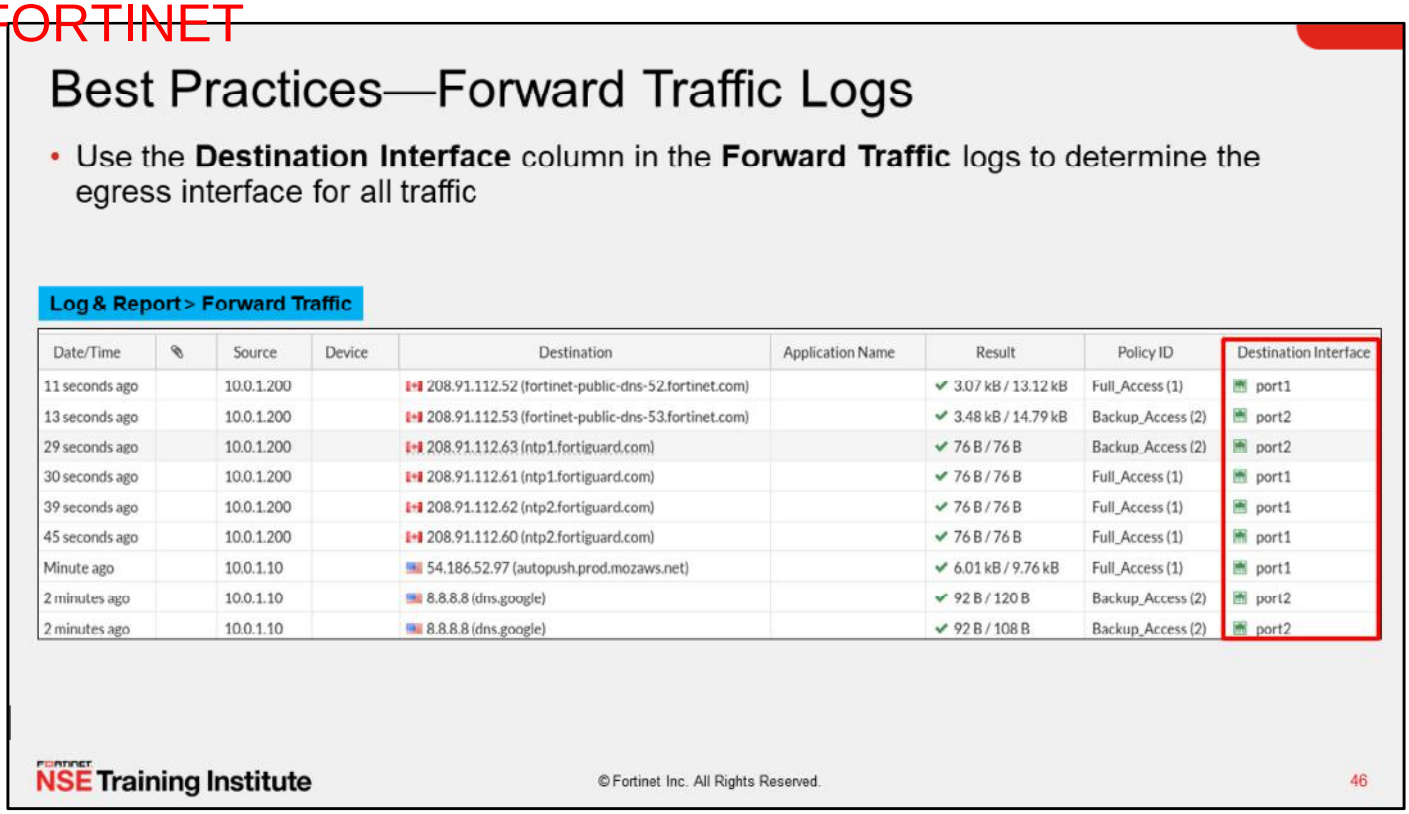

If you enable the **Destination Interface** column in the **Forward Traffic** logs, you can view the egress interface for traffic passing through your FortiGate device. You can use this information to determine which route is applied to which traffic stream, as well as identify any routing configuration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

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

NSE Training Institute

47

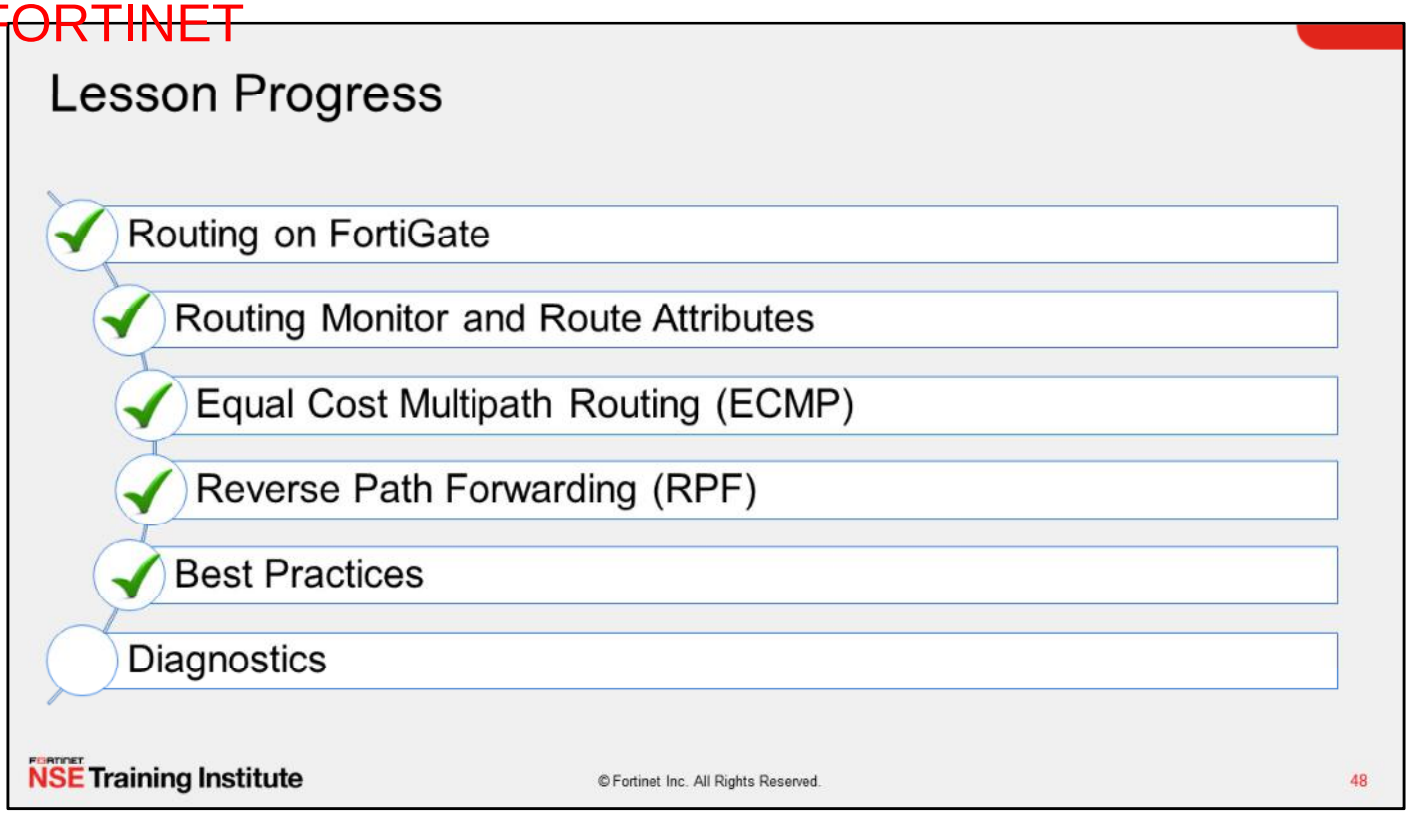

Good job! You now understand some routing best practices.

Now, you will learn about routing diagnostics.

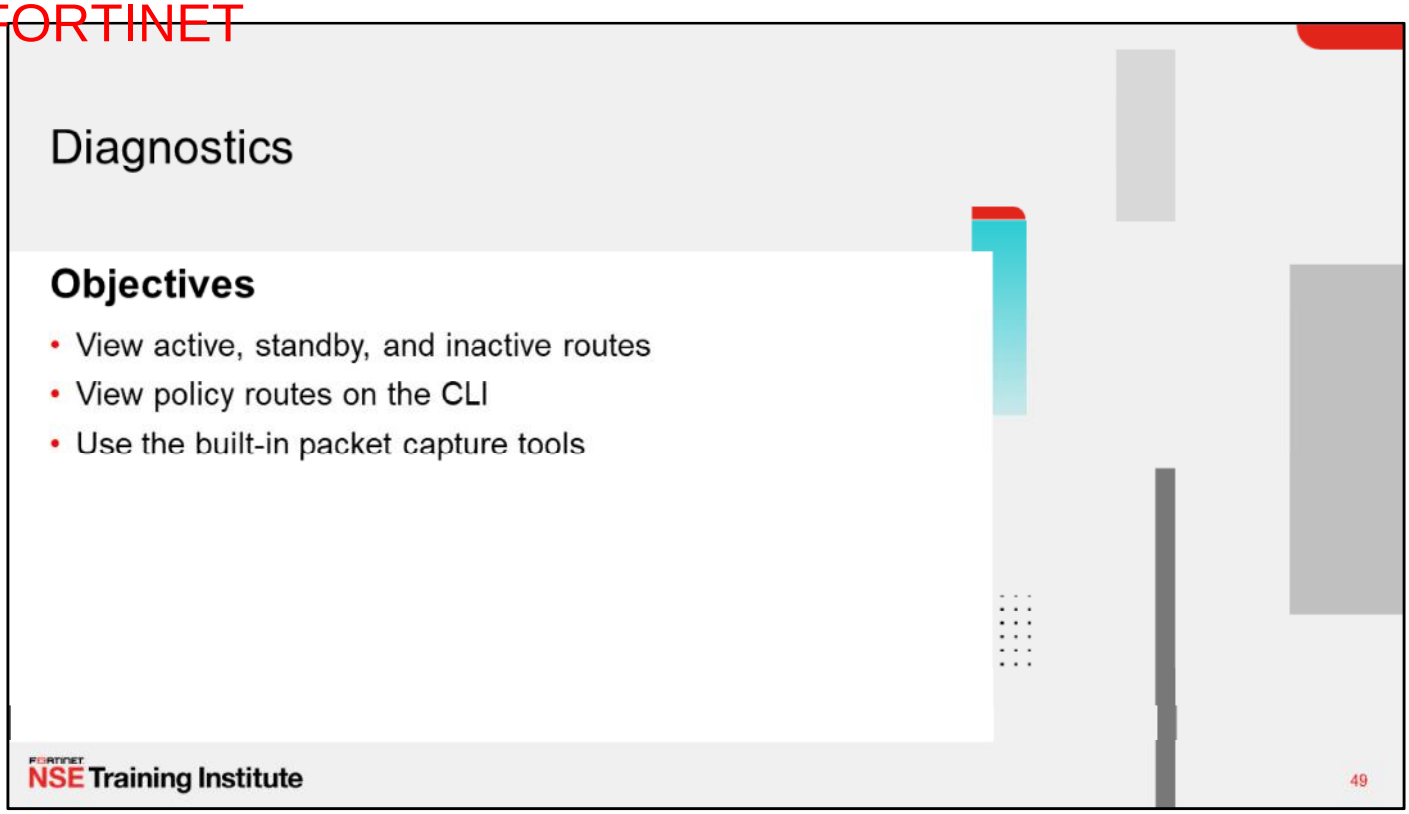

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in routing diagnostics, you should be able to determine the root causes of any routing failures in your network.

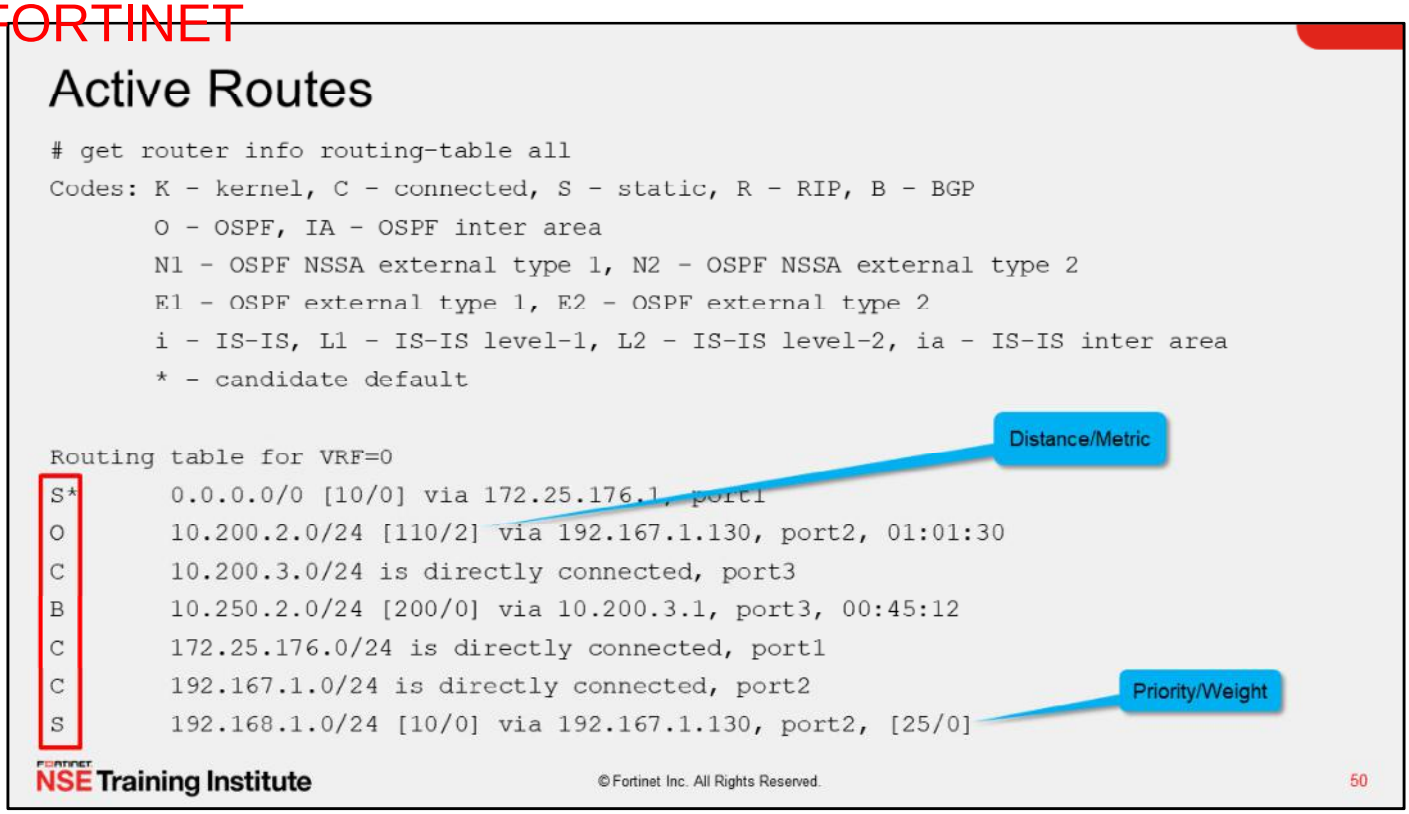

The CLI command shown on this slide displays all active routes in the routing table. The left-most column indicates the source for the route.

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 can also have priority and weight attributes, which are shown as the last pair of attributes for the respective route.

This command doesn't show standby or inactive routes. For example, when two static routes to the same destination subnet have different distances, the one with the lower distance is considered active and so is installed in the routing table. The one with the higher distance is considered standby and therefore is not installed in the routing table. So, this command displays only the route with the lowest distance (the active one).

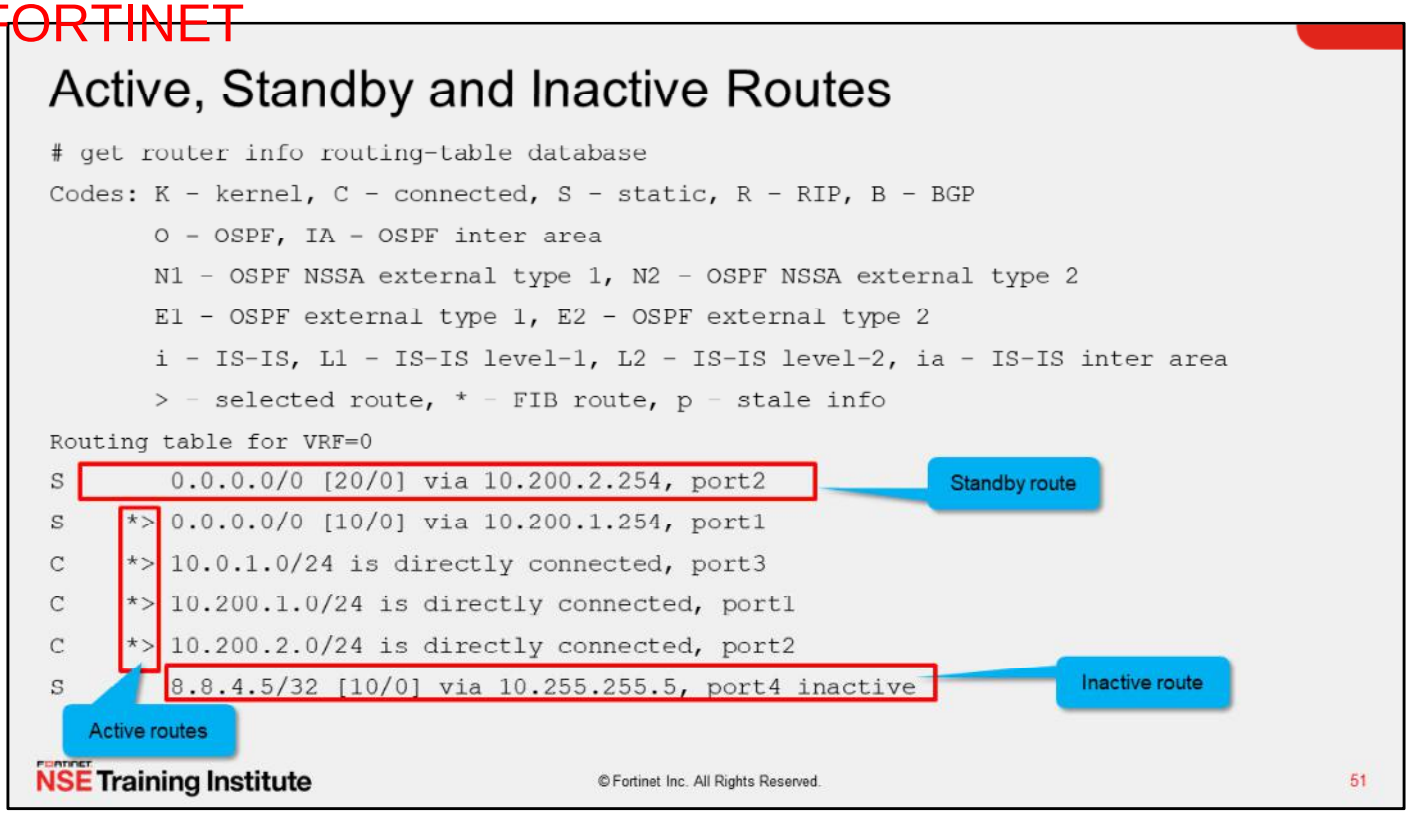

If you want to display active, standby, and inactive routes, use the CLI command shown on this slide.

In the example on this slide, the command shows one standby route. The route is the standby route because the default route over `port1` has a lower distance. The output also shows an inactive route. Routes are marked as inactive where the corresponding interface is administratively down, or has its link down, or when the gateway is detected dead by the link monitor.

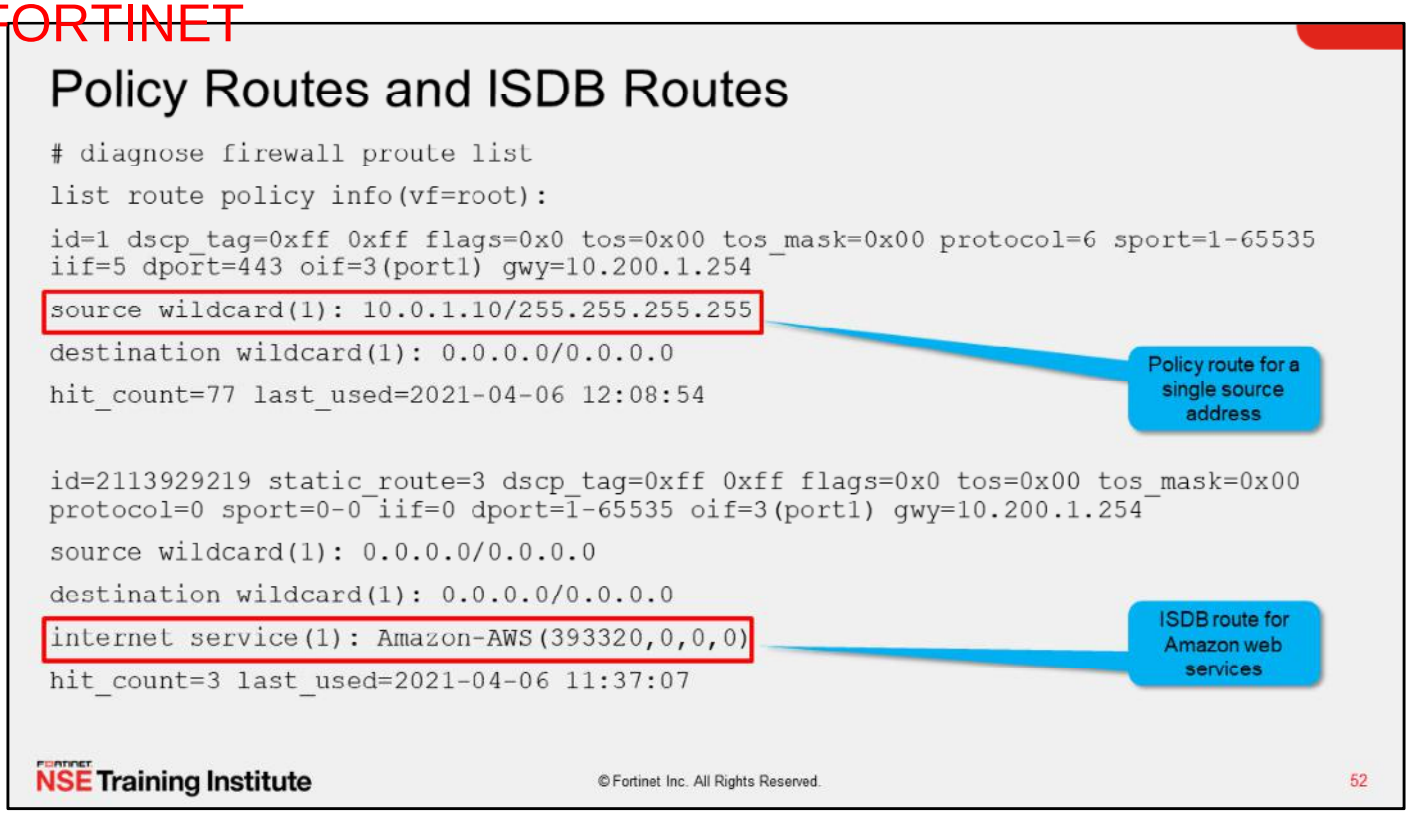

Policy routes, and static routes created using ISDB addresses are not added to the routing table; they are added to the policy routing table. You can display these routes using the CLI command shown on this slide. This command lists all the active routes in the policy routing table.

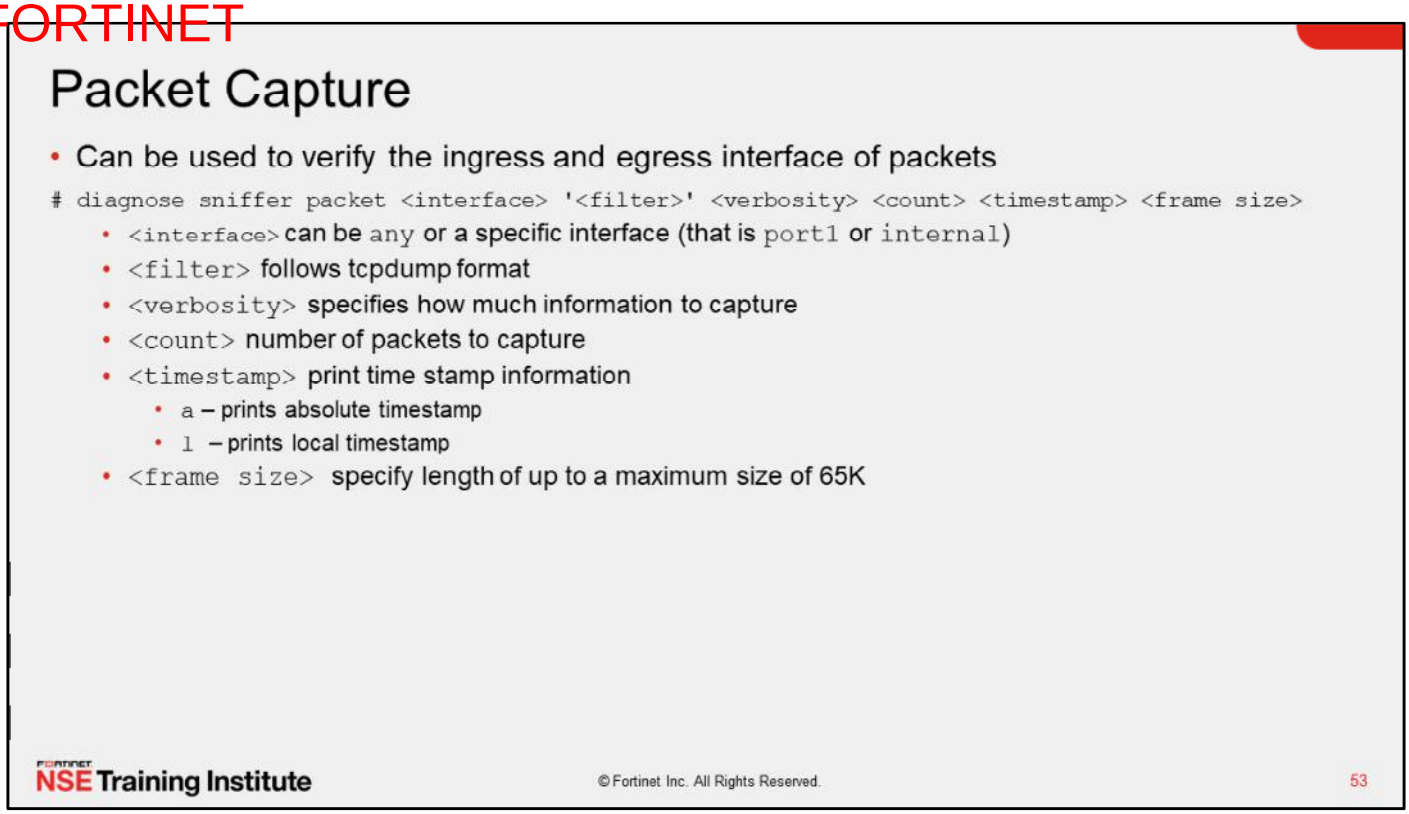

Packet captures, or *sniffers*, are one of the most useful sources of information for debugging routing problems. FortiGate includes a built-in traffic sniffer tool. You can use it to verify the ingress and egress interfaces of packets as they pass through. You can run the built-in sniffer from either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

The `<interface>` is the name of the physical or logical interface. If your account has the access profile super_admin, you can specify `any` to capture on all the interfaces. If you're using the `any` option, remember that the sniffer does not print any interface MAC addresses.

The filters are similar to `tcpdump` on Linux. 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 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 later in this course.

By default, the sniffer uses the MTU configured on the interface. Using the `<frame size>` argument, you can specify a length larger or smaller than the interface MTU. If you use the `any` interface, the sniffer will defaults to 1600 bytes.

Packet Capture Verbosity Level

Level	IP Headers	Packet Payload	Ethernet Headers	Interface Name
1	•			
2	•	•		
3	•	•	•	
4	•			•
5	•	•		•
6	•	•	•	•

- The most common levels are:
  - `4` – Prints the ingress and egress interfaces
    - You can verify how traffic is being routed, or if FortiGate is dropping packets
  - `3` or `6` – Prints the packet payload
    - You can convert this output to a packet capture (pcap) file that can be opened with a packet analyzer
  - If you don't specify a level, the sniffer uses level `1` by default

NSE Training Institute

54

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`).

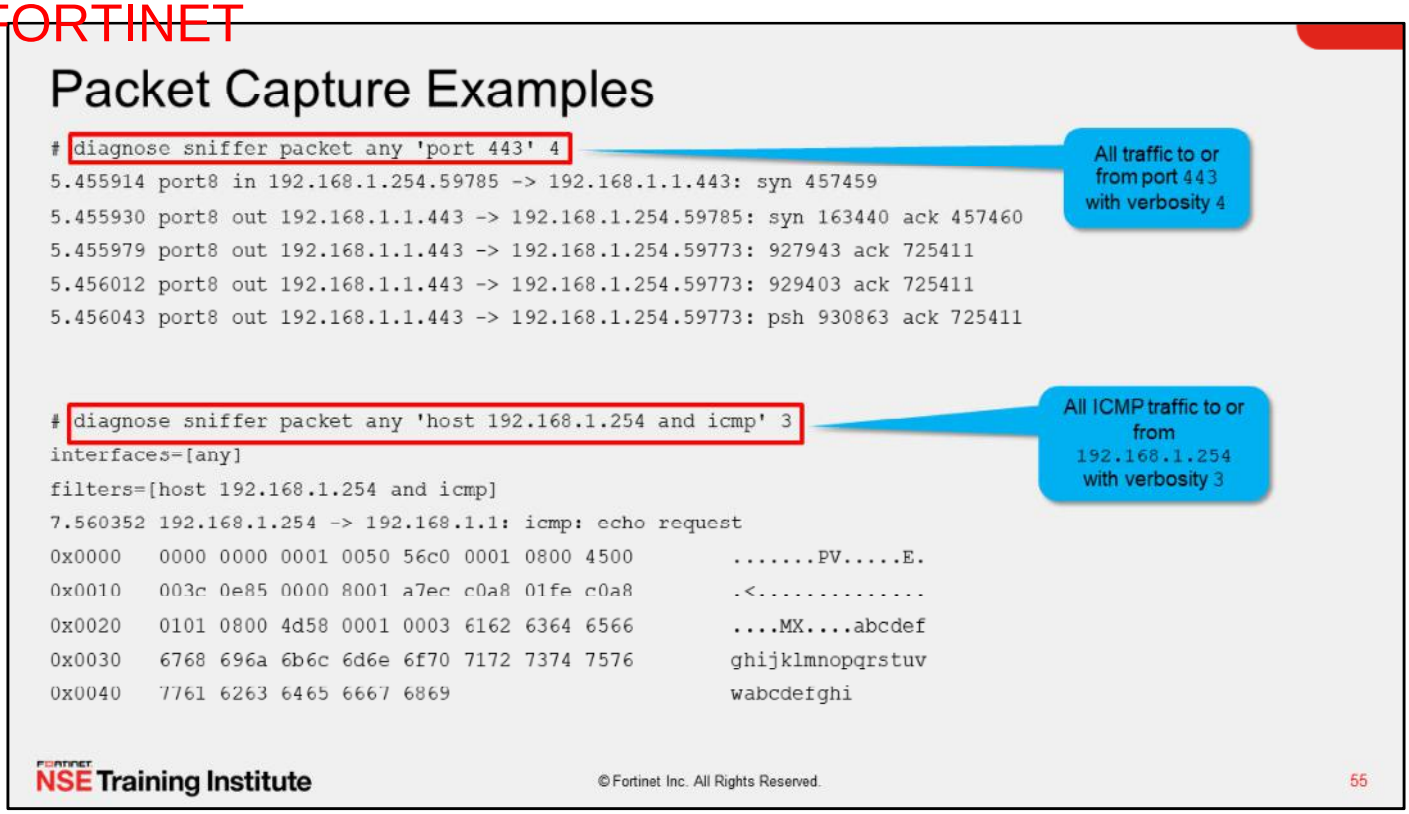

This slide shows two examples of packet capture outputs. The first example captures all traffic to and from port `443`. It uses verbosity `4`, so the information is easy to read. It displays one line per packet, containing the incoming and outgoing interface, IP addresses, port numbers, and type of packet (`SYN`, `SYN/ACK`, and so on).

The second example captures all ICMP traffic coming from or going to the host `192.168.1.254`. In this case, the verbosity is `3`, which is longer and more difficult to read as it includes the IP payload of the packets. However, this is one of the two verbosity levels to use (`6` being the other one) if you need to export the output to Wireshark.

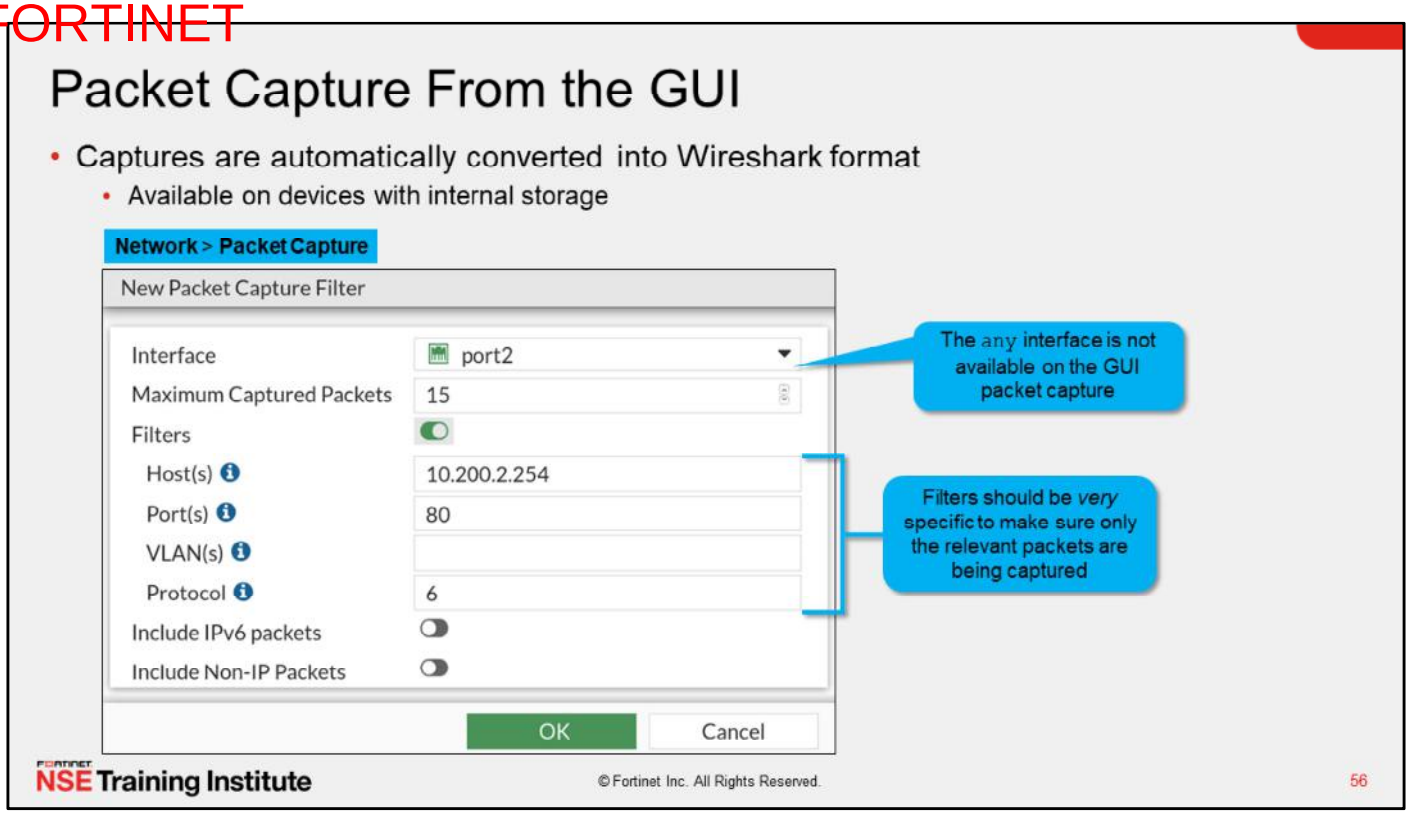

If your model of FortiGate has internal storage, you can capture packets on the GUI. The options are similar to those for the CLI. To run a trace, specify a source interface and a filter.

What is the main advantage of using the GUI over the CLI? You download the output in a file format (pcap) that is ready to be opened using Wireshark, without having to use a conversion script.

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.

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 commands 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

NSE Training Institute

57

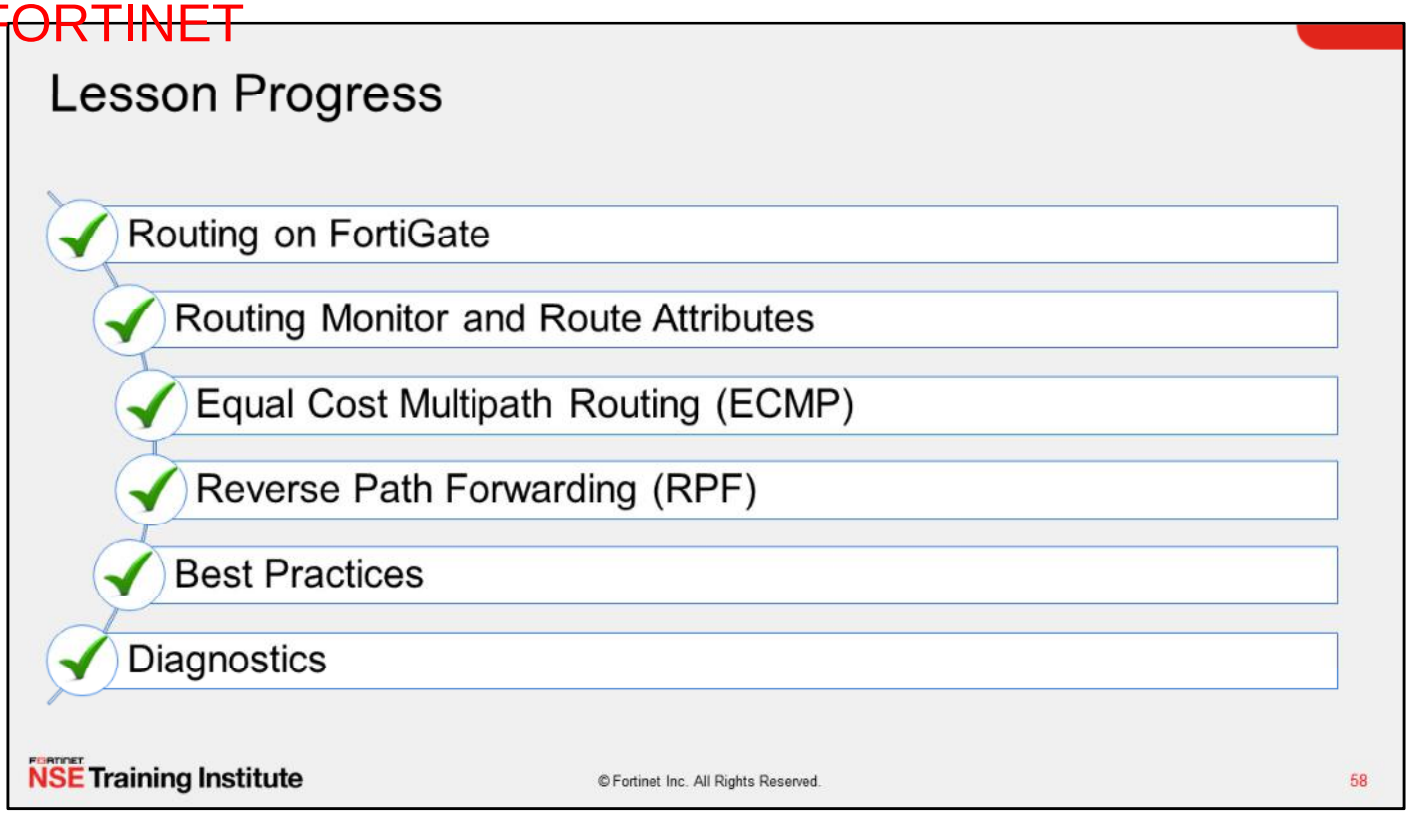

Congratulations! You have completed the lesson.

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

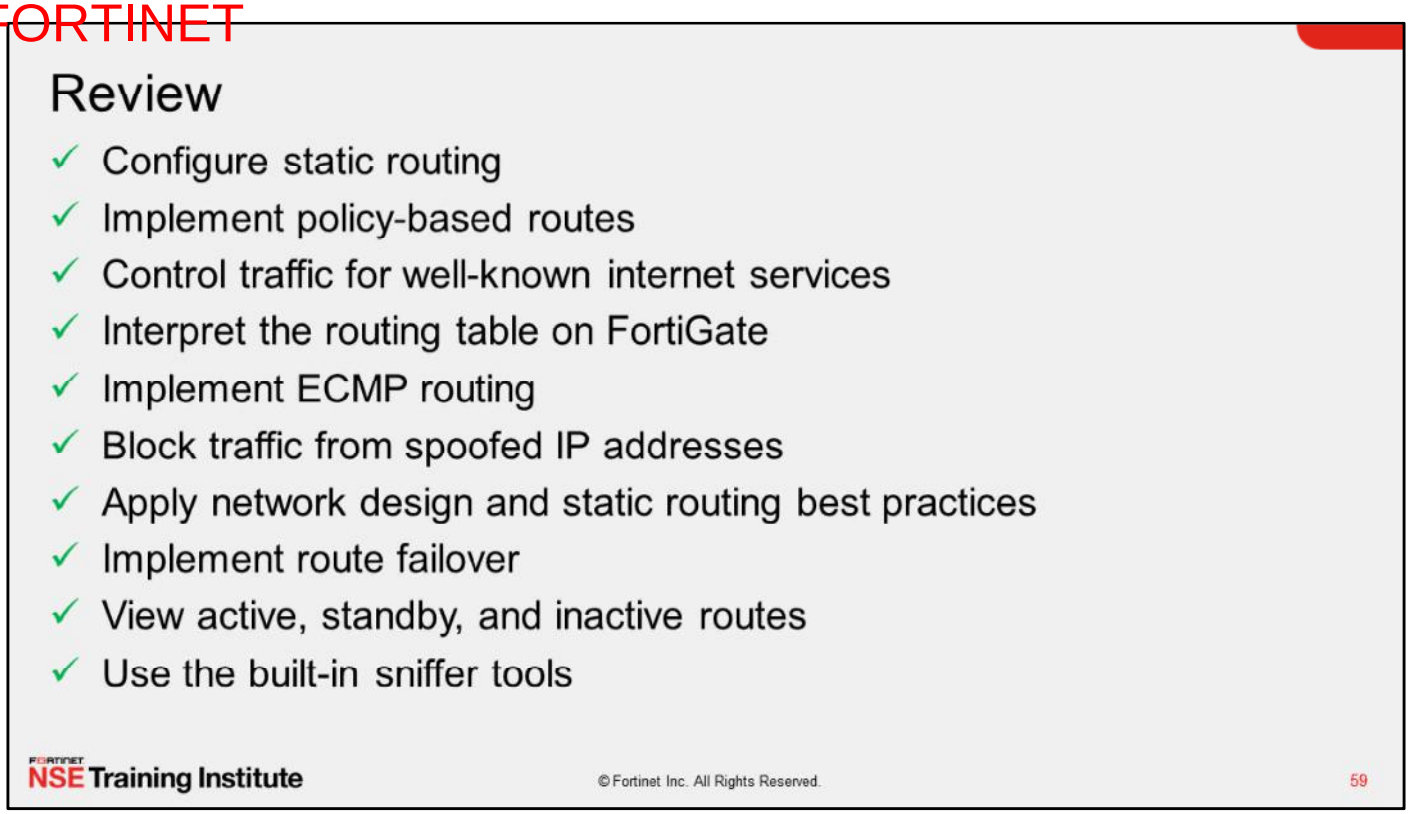

This slide shows the objectives that you covered in this lesson.

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

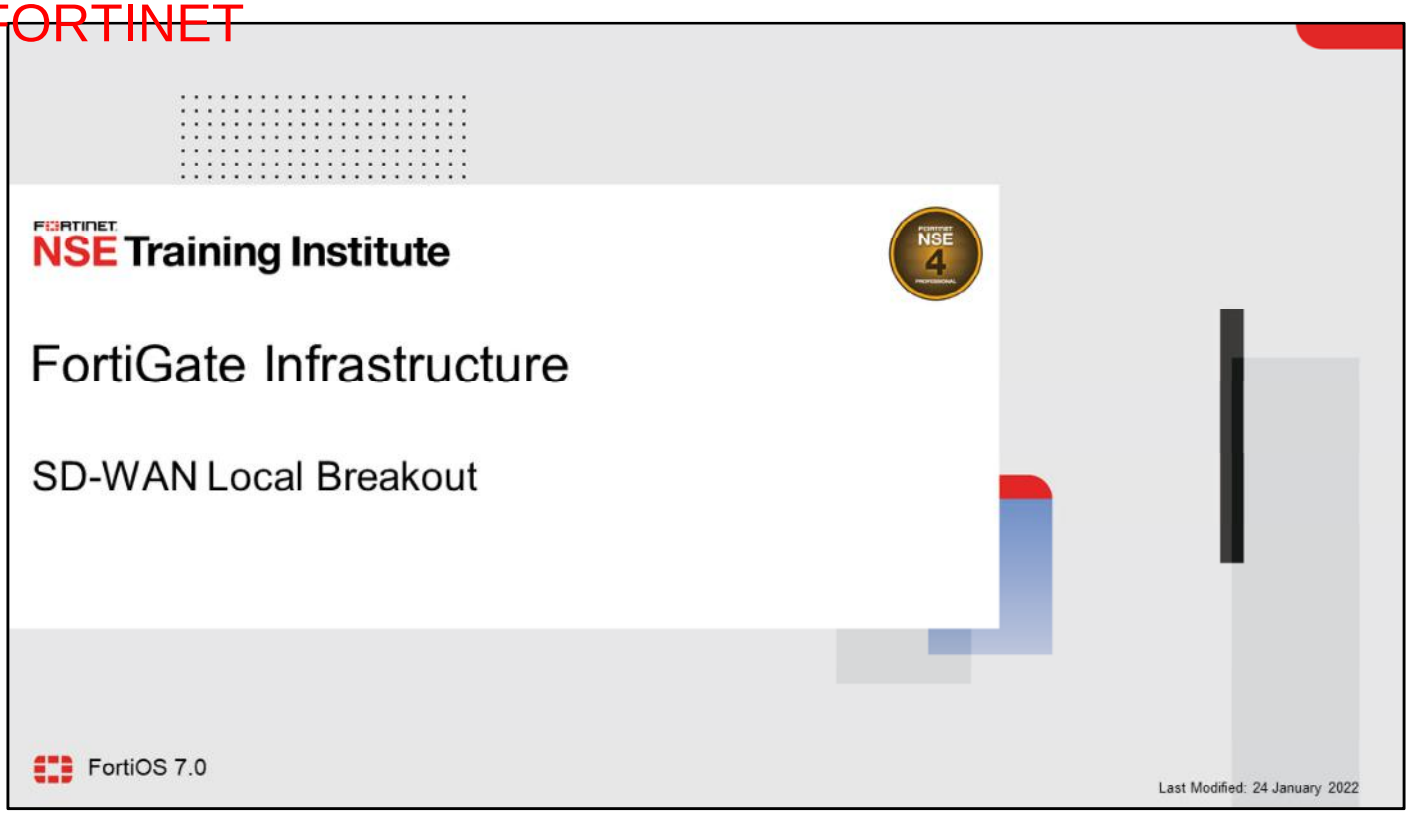

In this lesson, you will learn about the SD-WAN feature available on FortiGate.

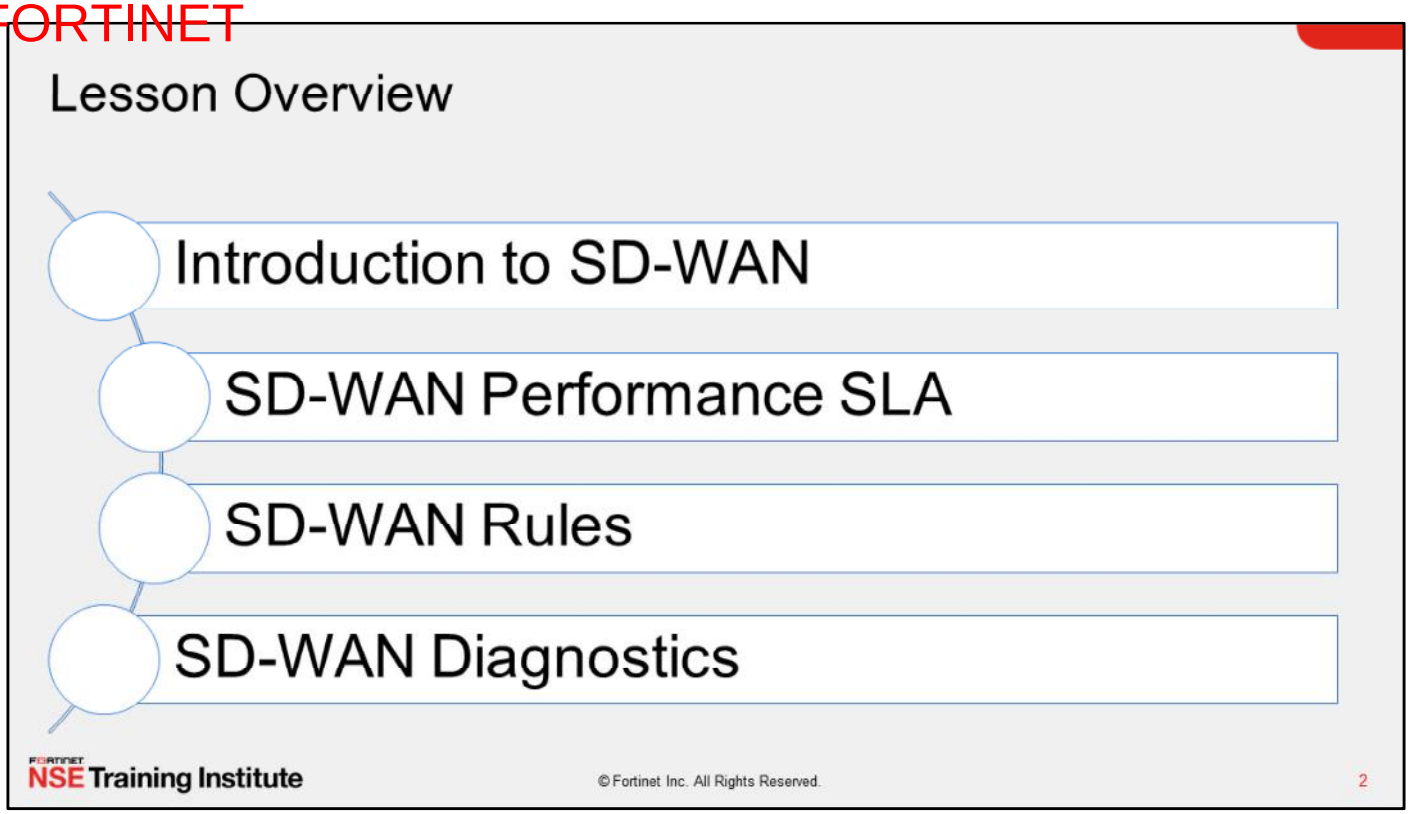

In this lesson, you will learn about the topics shown on this slide.

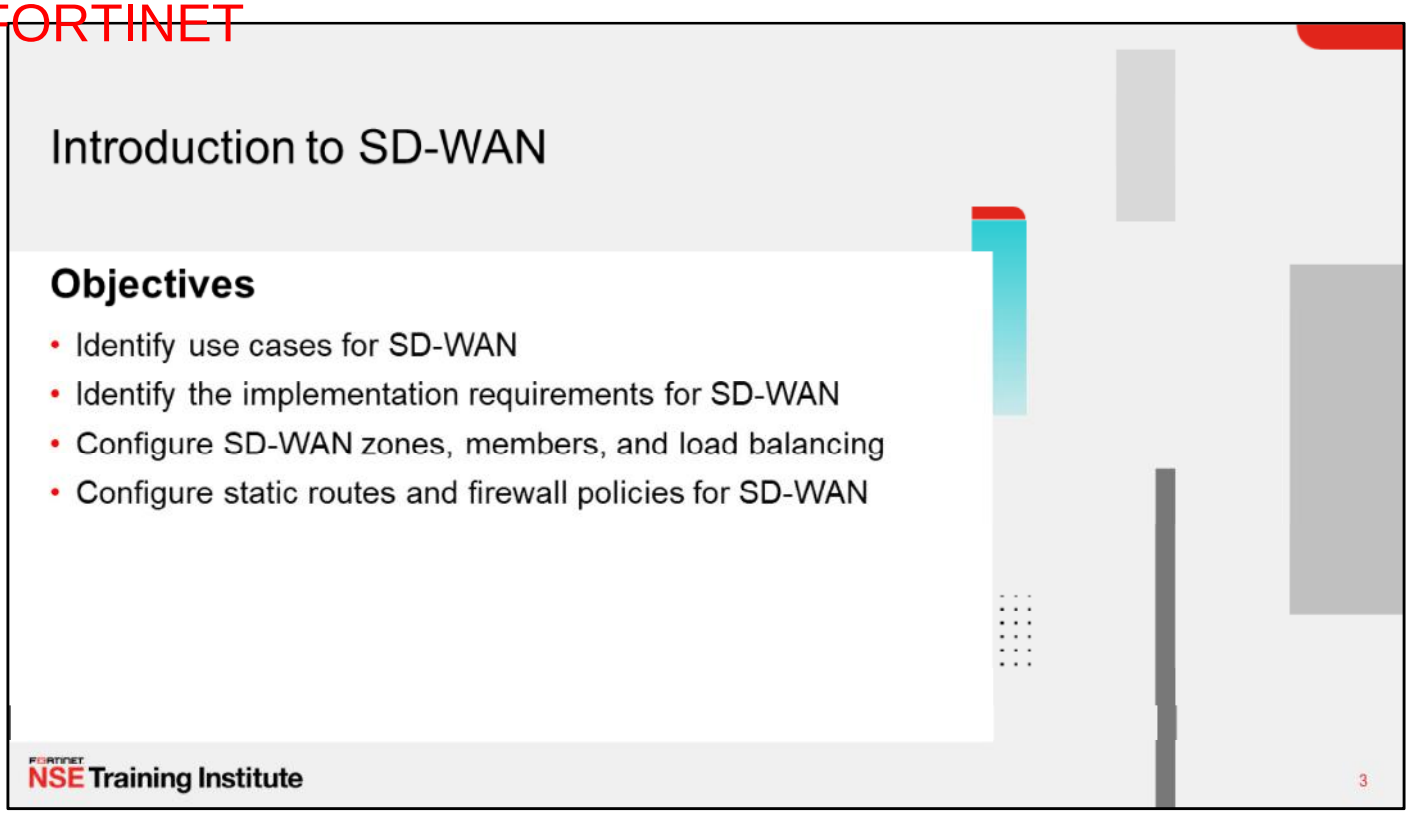

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in SD-WAN, you should be able to configure SD-WAN zones, members, and traffic load balancing to use multiple WAN links effectively on FortiGate.

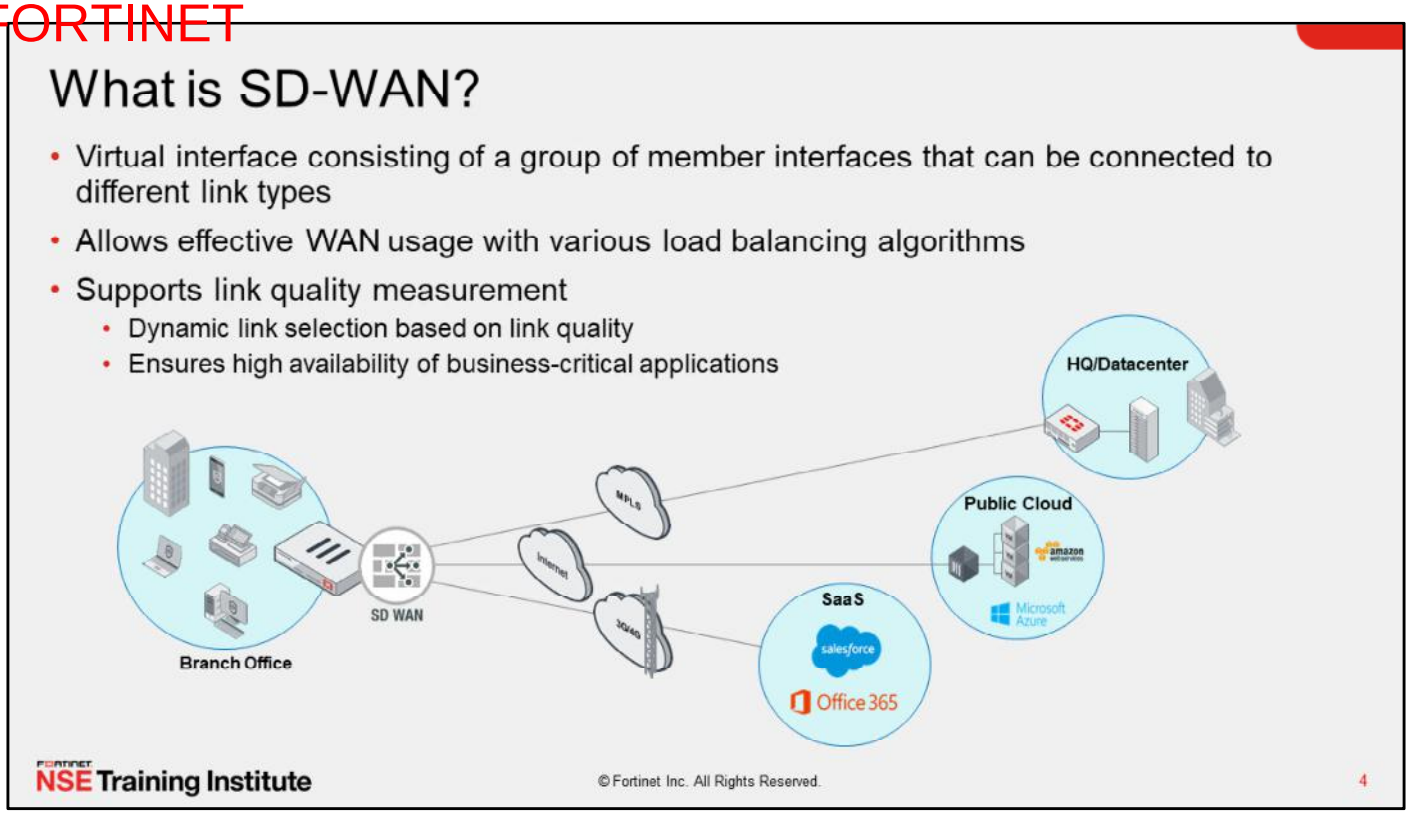

SD-WAN is a virtual interface consisting of a group of member interfaces that can be connected to different link types. FortiGate groups all the member interfaces into a single virtual interface: the SD-WAN interface. Using SD-WAN simplifies configuration because the administrator can configure a single set of routes and firewall policies and apply them to all member interfaces. There can be only one SD-WAN interface per VDOM.

One of the main motivators for deploying SD-WAN is effective WAN use, when you are using multiple WAN links. Effective WAN use is achieved using various load balancing algorithms, such as bandwidth usage, sessions, or application-aware routing. Another important feature of SD-WAN is link quality measurements. Using ping or HTTP echo, FortiGate can determine the latency, jitter, or packet loss percentage for each link, and dynamically select links based on these measurements. This ensures high availability (HA) for business-critical applications.

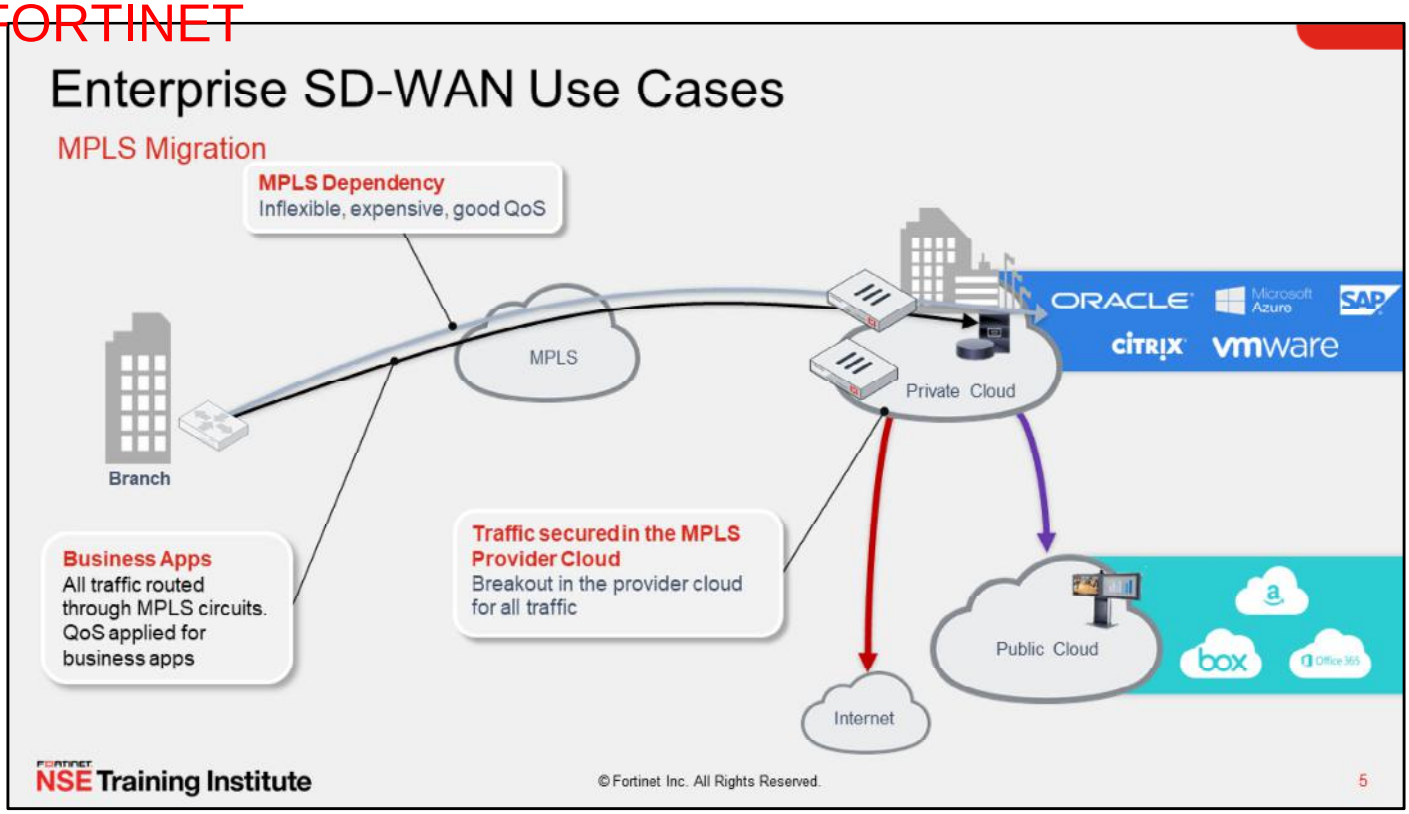

Now, you will learn about SD-WAN use cases. In the example shown on this slide, the customer depends heavily on expensive, inflexible MPLS. All the traffic is routed through the MPLS circuit to the provider cloud, then to the public cloud or internet, based on the applications. As mentioned earlier, there is no flexibility in this scenario, and yet it is an expensive solution for the customer. How can the customer add redundancy, flexibility, reliability, and, most importantly, security, without adding costly infrastructure? You will learn about the solution in this lesson.

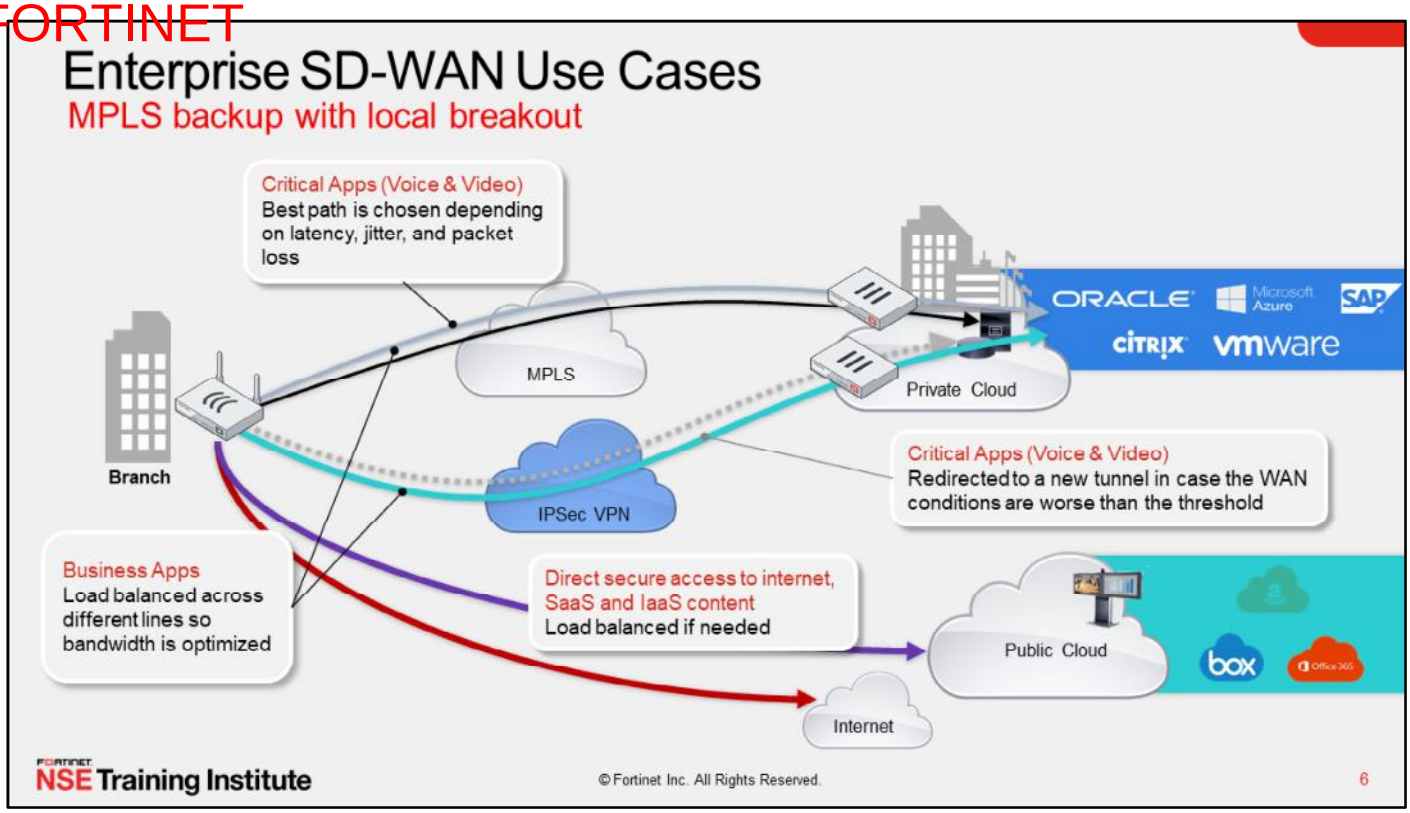

In the example shown on this slide, the customer would like to keep MPLS for business-critical applications, while adding flexibility and redundancy. MPLS is being used to send business-critical traffic (for example, voice and video) based on the best path with less delay, jitter, or packet loss. In case the current path degrades below policy threshold, business-critical traffic will be rerouted to a new tunnel. Also, non-critical traffic is load balanced across different lines to maximize bandwidth or minimize cost. At the same time, the branch can have direct secure access to the internet, which improves the cloud application performance, and can load balance SaaS and IaaS content if needed.

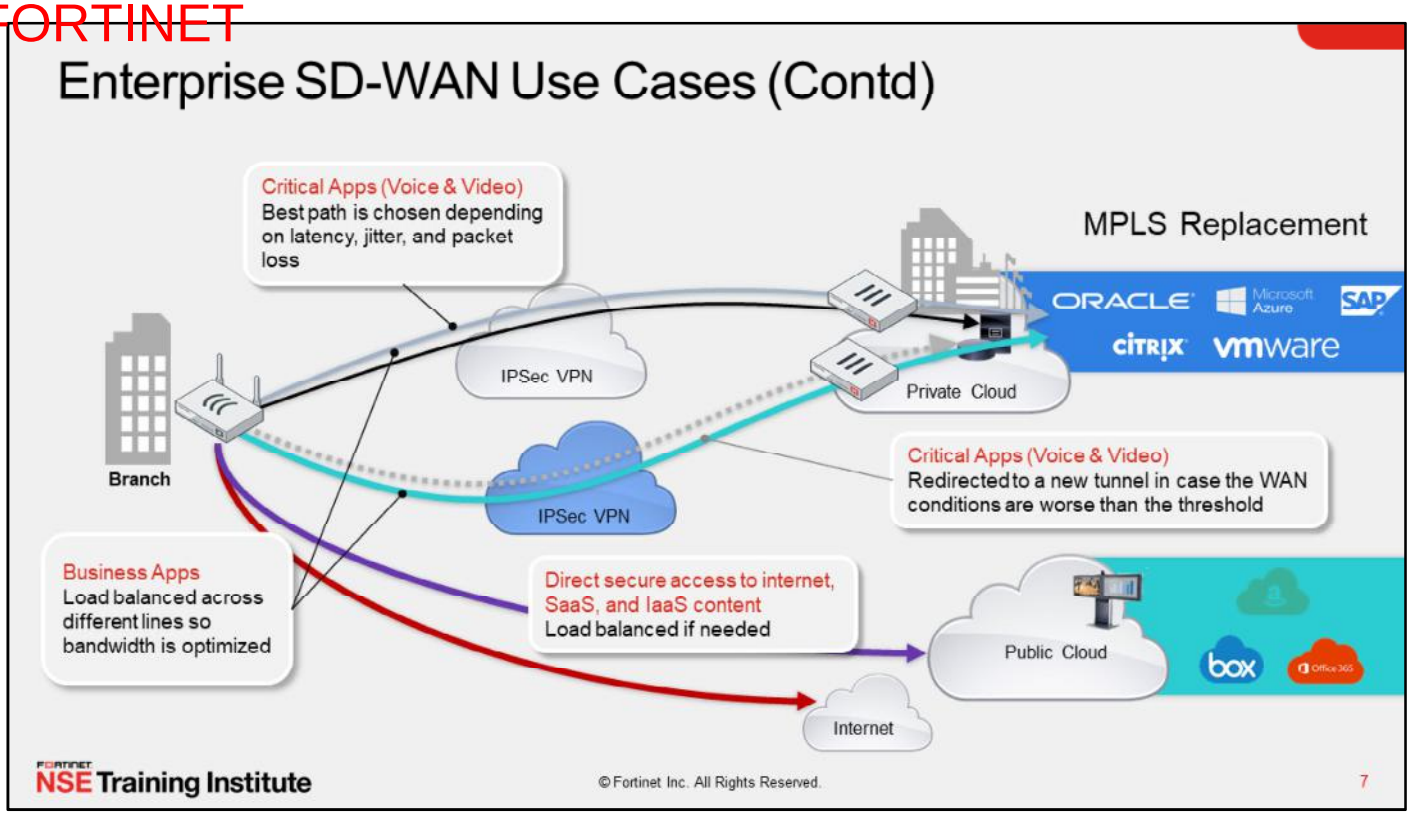

In the example shown on this slide, costly MPLS is replaced by two internet VPN tunnels, and yet gains robust resiliency and redundancy. By replacing MPLS, the customer can minimize cost while maximizing quality. The SD-WAN solution is a network application-aware solution and dynamically selects the best WAN to maintain higher SLA .

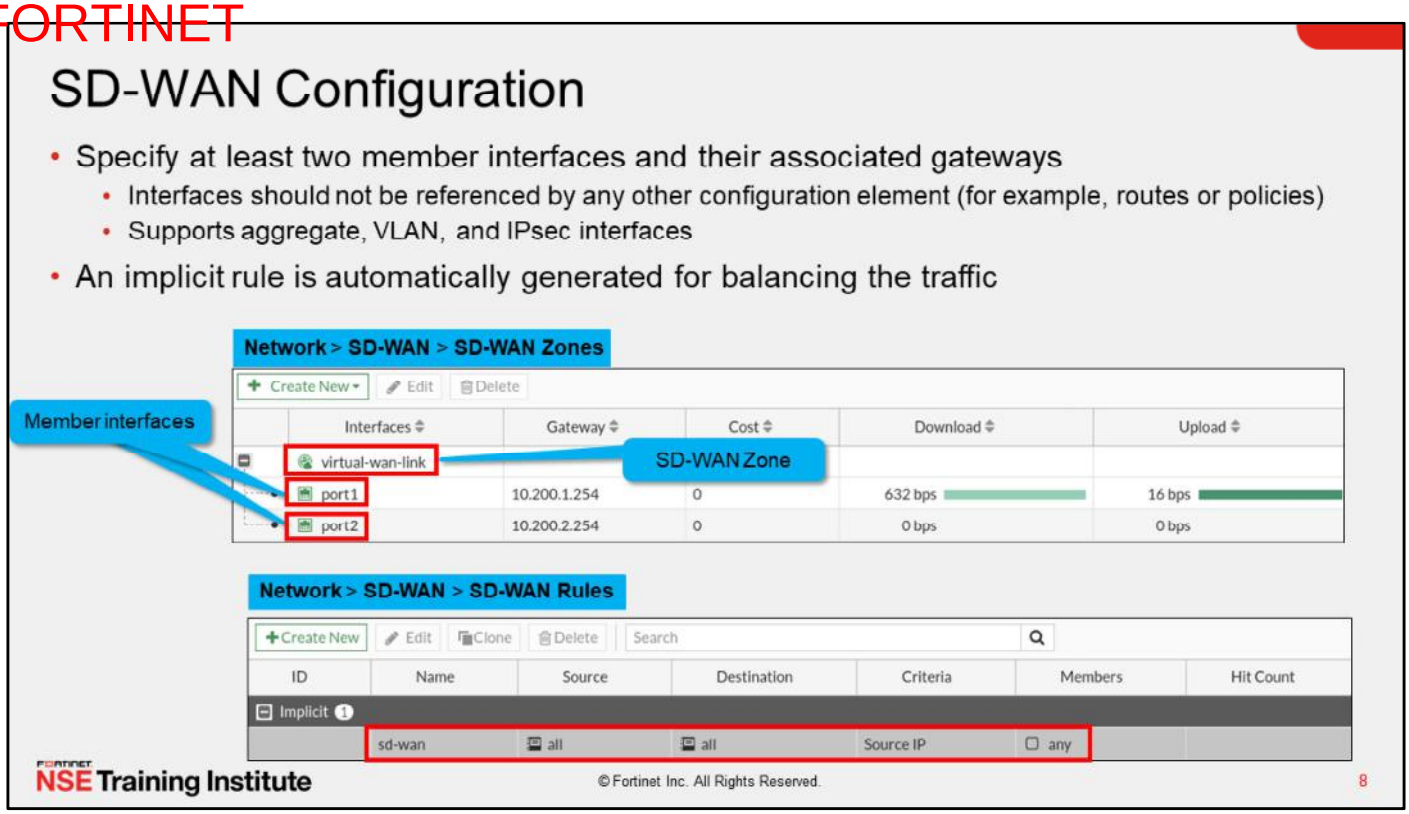

When you configure SD-WAN, you must specify at least two member interfaces, their associated gateways, and an SD-WAN zone. You should configure SD-WAN early during the initial setup of FortiGate because, if an interface is already referenced by a firewall policy or static route, you cannot use it as a member interface. If you intend to use an interface as an SD-WAN member, and that interface is being referenced by a firewall policy or static route, you must delete the associated firewall policy and static route before you can assign that interface as an SD-WAN member. SD-WAN supports physical interfaces as well as VLAN, aggregate, and IPsec interfaces.

You can also easily add another member interface at a later date, to add more bandwidth or QoS options.

FortiGate groups all the member interfaces into a single virtual interface for creating routes: the SD-WAN interface. Using SD-WAN simplifies configuration because the administrator can configure a single set of routes and apply them to all member interfaces. There can be only one SD-WAN interface per VDOM.

An implicit rule is automatically generated when you enable SD-WAN. If none of the conditions for any of the other rules are met, then the implicit rule is used. This implicit rule is designed to balance the traffic among all the available SD-WAN member links. You will learn about SD-WAN rules later in this lesson.

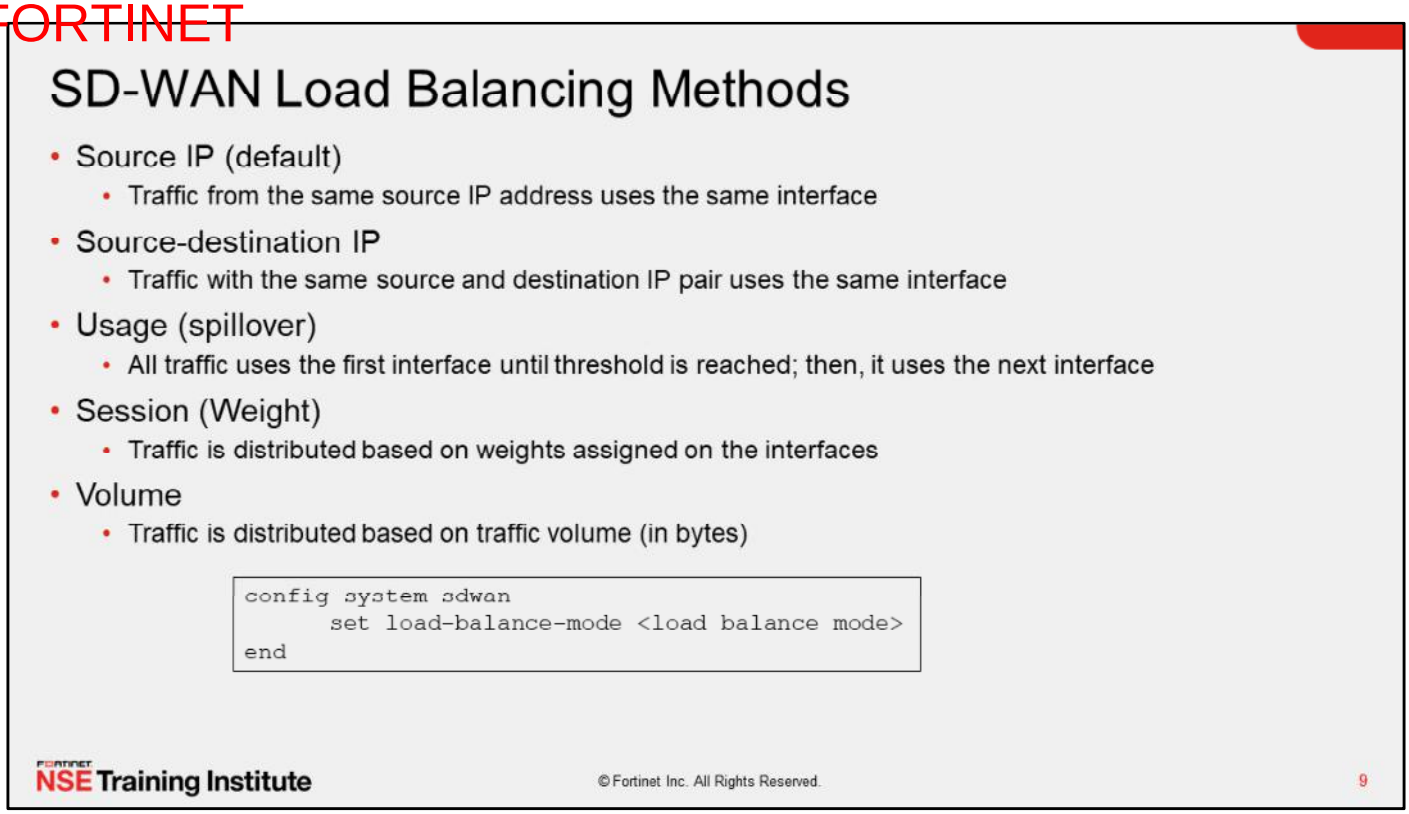

SD-WAN load balancing uses traffic distribution methods that are similar to those used by equal cost multipath (ECMP). However, SD-WAN link load balancing includes one more balancing method: volume.

By default, the `load-balance-mode` is set to `source-ip-based`. However, you can change the load balancing mode to any of the following:

- `source-ip-based`:
  - All traffic from a source IP is sent to the same interface.
- `weight-based`:
  - Interfaces with higher weights have higher priority and get more traffic.
- `usage-based`:
  - All traffic is sent to the first interface on the list. When the bandwidth on that interface exceeds the spillover limit, new traffic is sent to the next interface.
- `source-dest-ip-based`:
  - Source and destination IP load balancing. All traffic from a source IP to a destination IP is sent to the same interface.
- `measured-volume-based`:
  - Volume-based load balancing. Traffic is load balanced based on traffic volume (in bytes). More traffic is sent to interfaces with higher volume ratios.

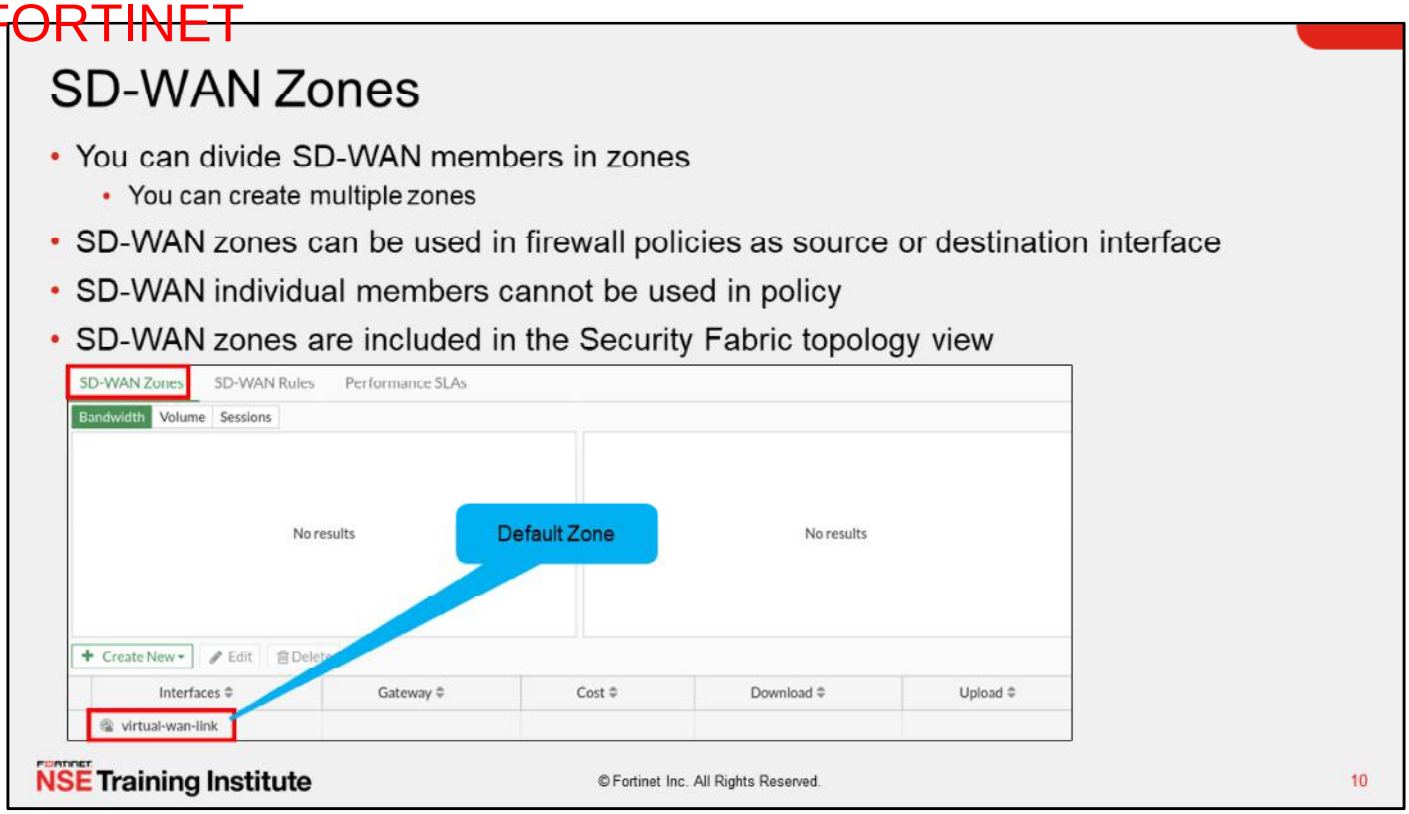

You can divide SD-WAN interfaces into smaller logical groups called SD-WAN zones. You can then use these SD-WAN zones in firewall policies to allow you to have more granular control over traffic being inspected and allowed. You can create multiple SD-WAN zones for SD-WAN members. However, SD-WAN members cannot be shared between multiple zones. By default, FortiGate has the **virtual-wan-link** zone created.

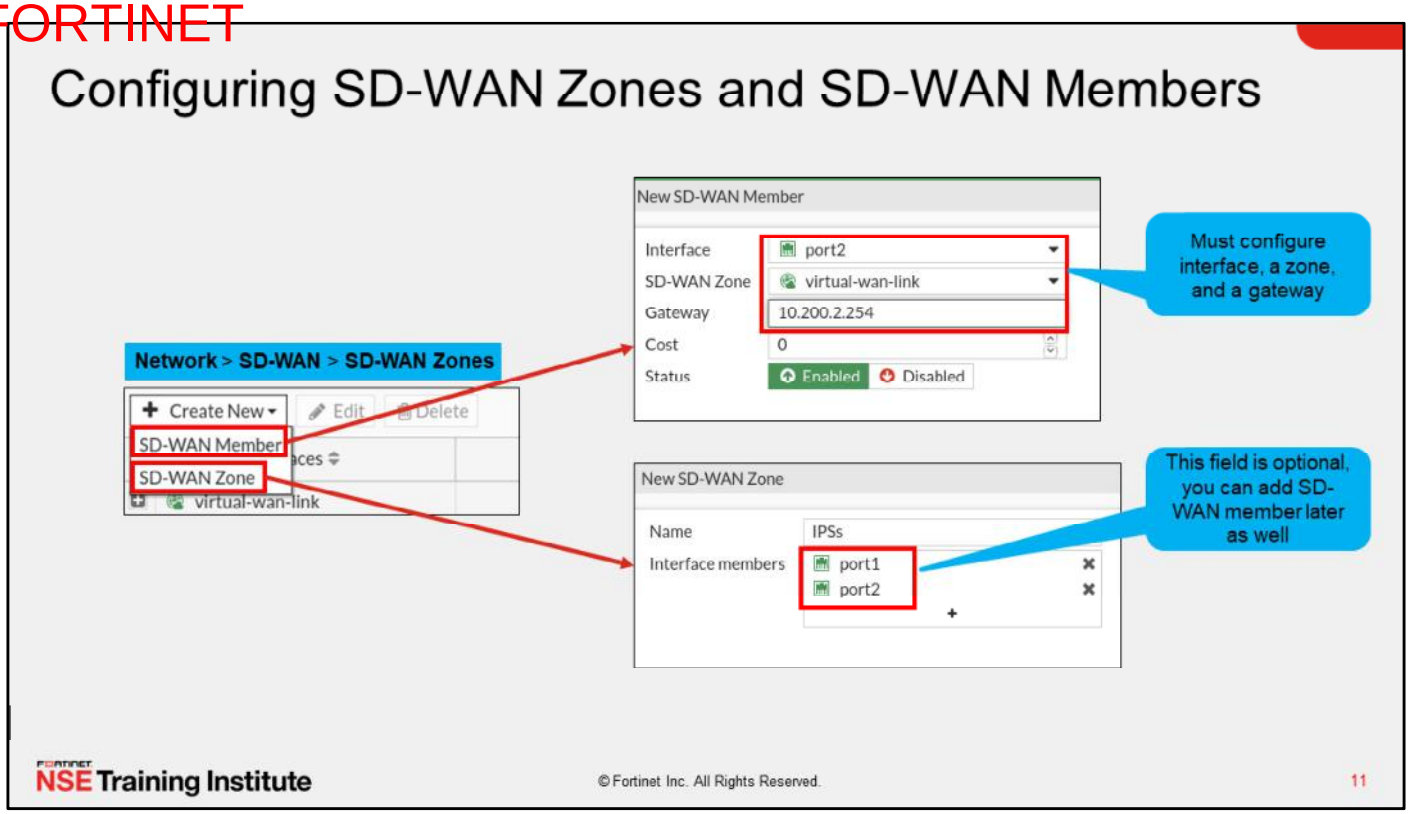

When you create an SD-WAN member, you must specify an interface, SD-WAN zone, and a gateway address. You can use the default SD-WAN zone created on FortiGate, or you can create a new SD-WAN zone. After you assign an SD-WAN interface to an SD-WAN zone, you have the option to change the zone.

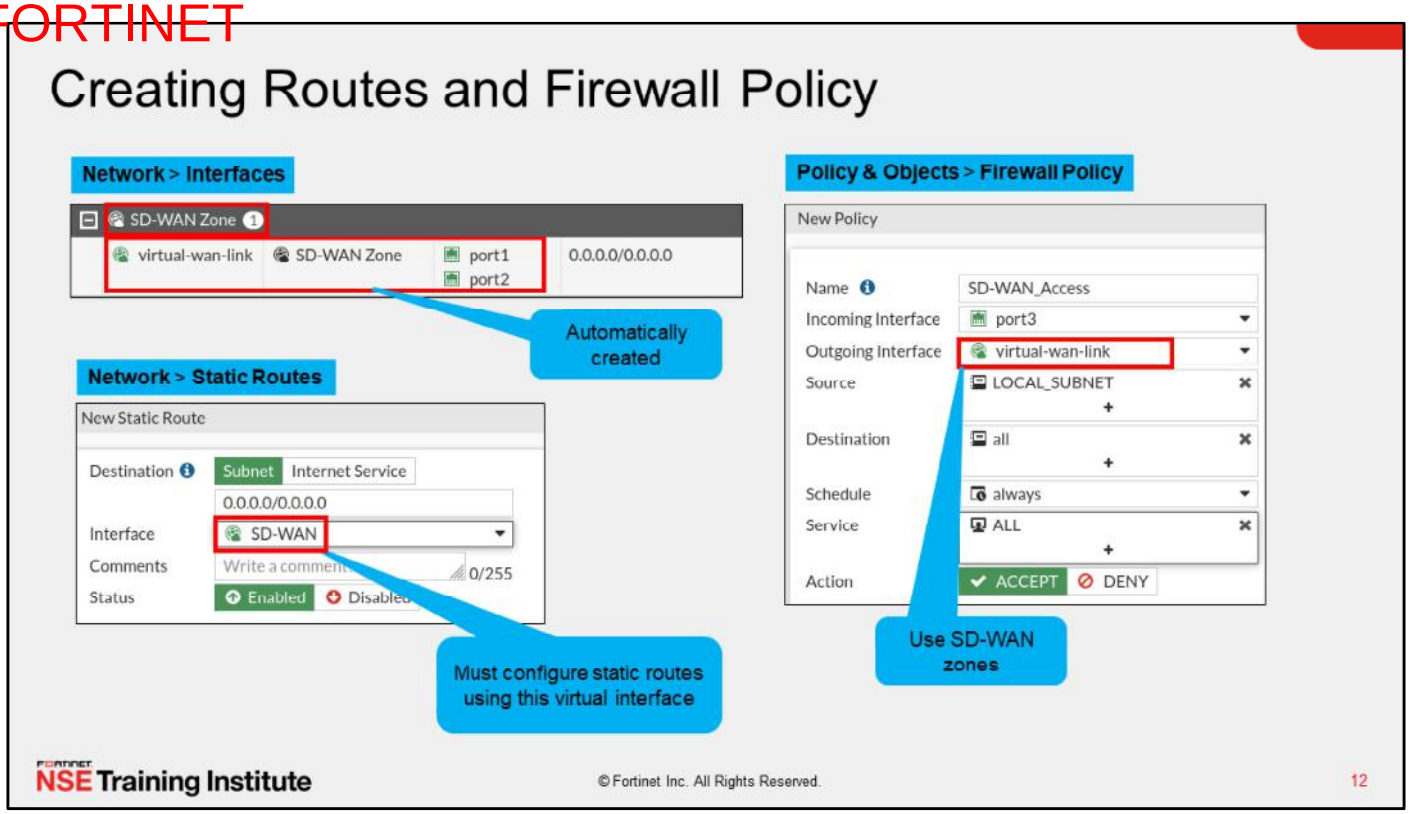

After you enable SD-WAN, and configure the member interfaces and the load balancing method, a logical interface with the name SD-WAN is automatically added to the interface list when you create a static route. Next, you must create the routes using this virtual interface.

For firewall policies, you must use the SD-WAN zones as a source interface or destination interface. You can't use individual members or an SD-WAN virtual interface in firewall policies. Firewall policies created with the SD-WAN interface allow traffic to be forwarded through any member interface.

You must still configure a default route when implementing SD-WAN. The default route configuration using the SD-WAN interface does not require a gateway address because FortiGate forwards packets to the appropriate gateway, based on the member interface gateway information.

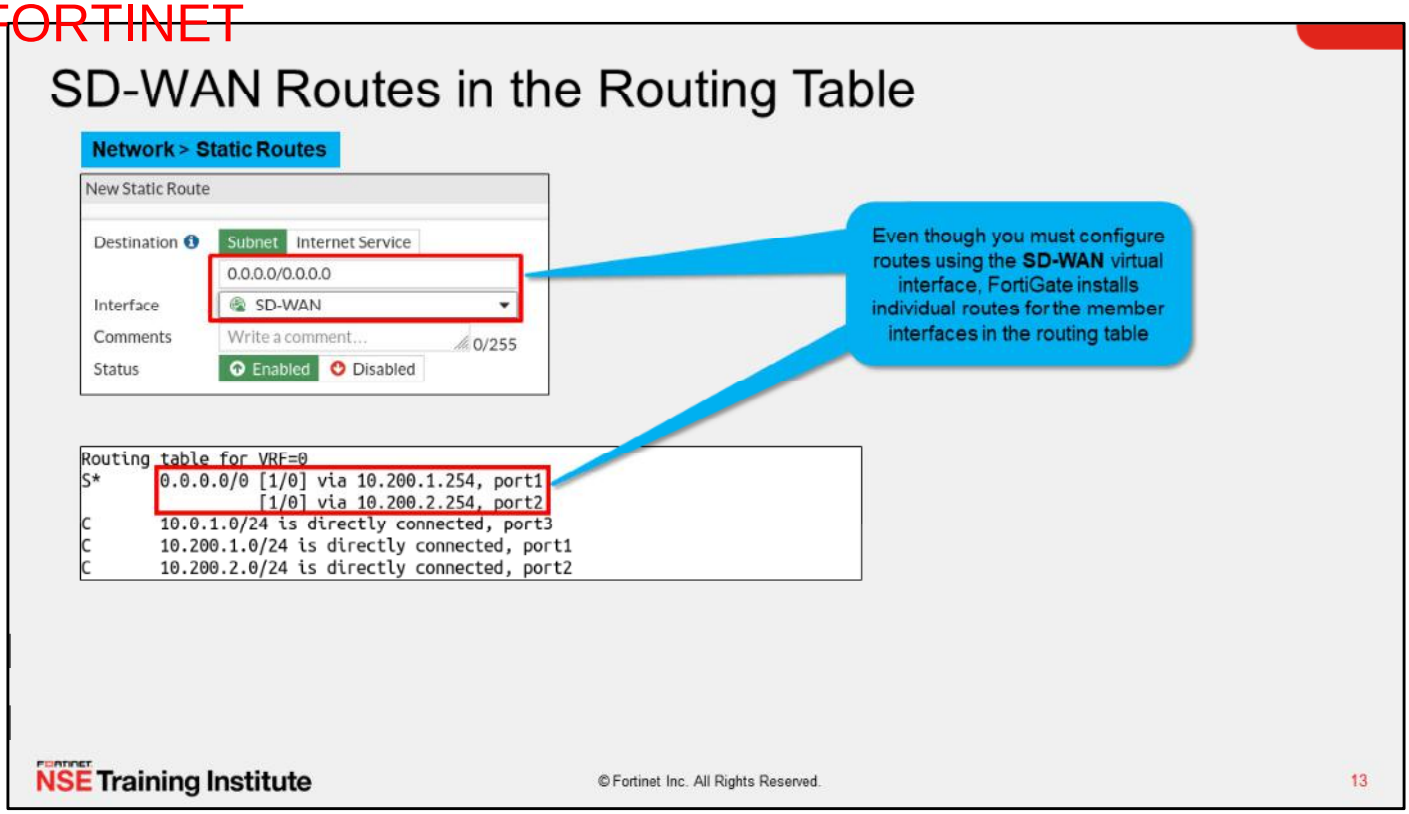

Even though you must configure routes using the **SD-WAN** virtual interface, FortiGate installs individual routes for the member interfaces in the routing table. These routes share the same attributes (destination address and subnet, distance, and priority) and are all active. This allows FortiGate to remove individual routes in the event of an interface outage, and redirect all traffic to the remaining member interfaces, without affecting the whole SD-WAN load balancing group.

Knowledge Check

1. Which method of load balancing is supported by SD-WAN but *not* supported by ECMP routing?
   - A. Sessions
   - ✓ B. Volume

2. Which configuration task is correct when implementing SD-WAN?
   - ✓ A. Configure a default route using the **SD-WAN** virtual interface.
   - B. Configure firewall policies for each individual member interface.

NSE Training Institute

14

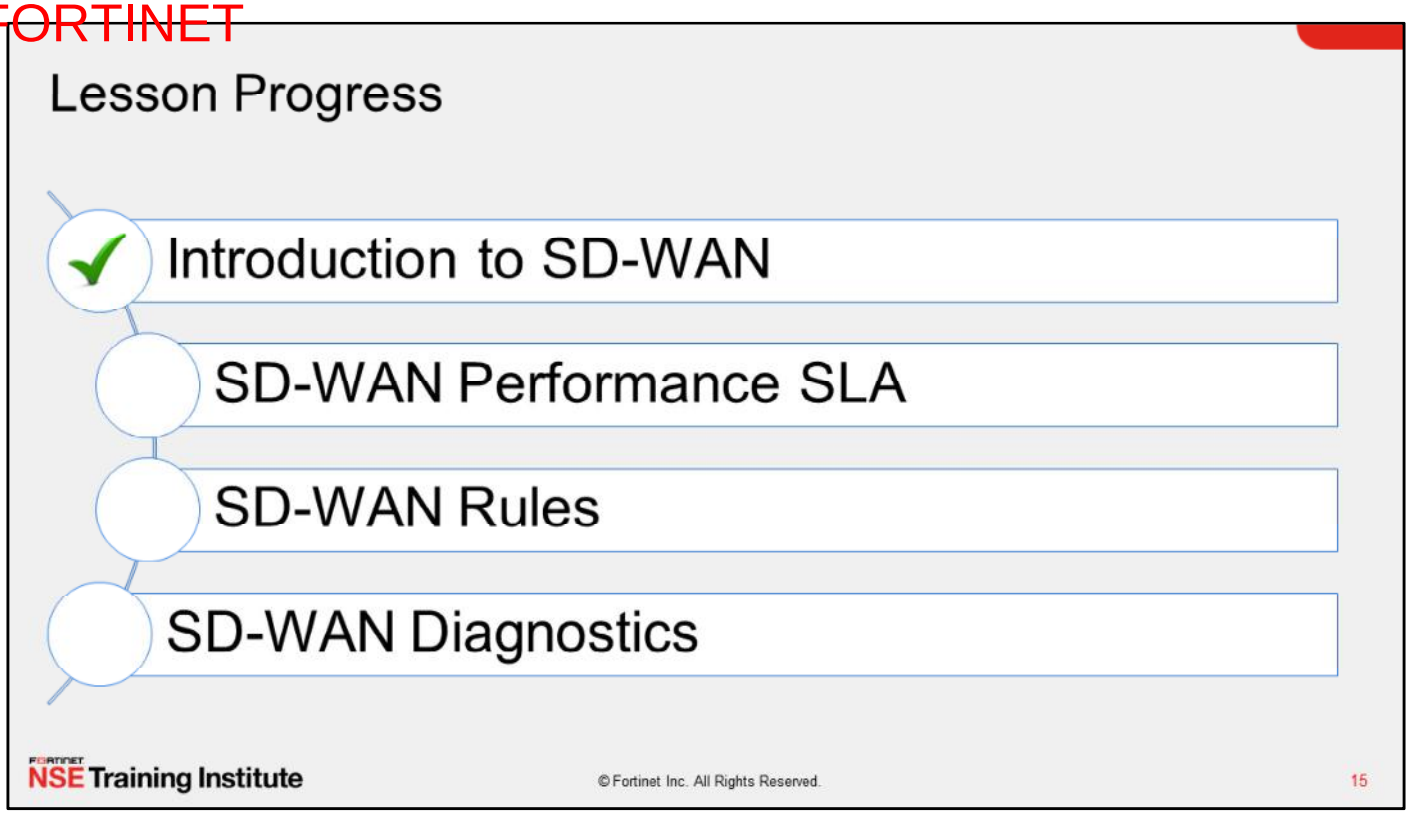

Good job! You now understand the SD-WAN feature on FortiGate.

Now, you will learn about the SD-WAN performance SLA, and link quality measurements.

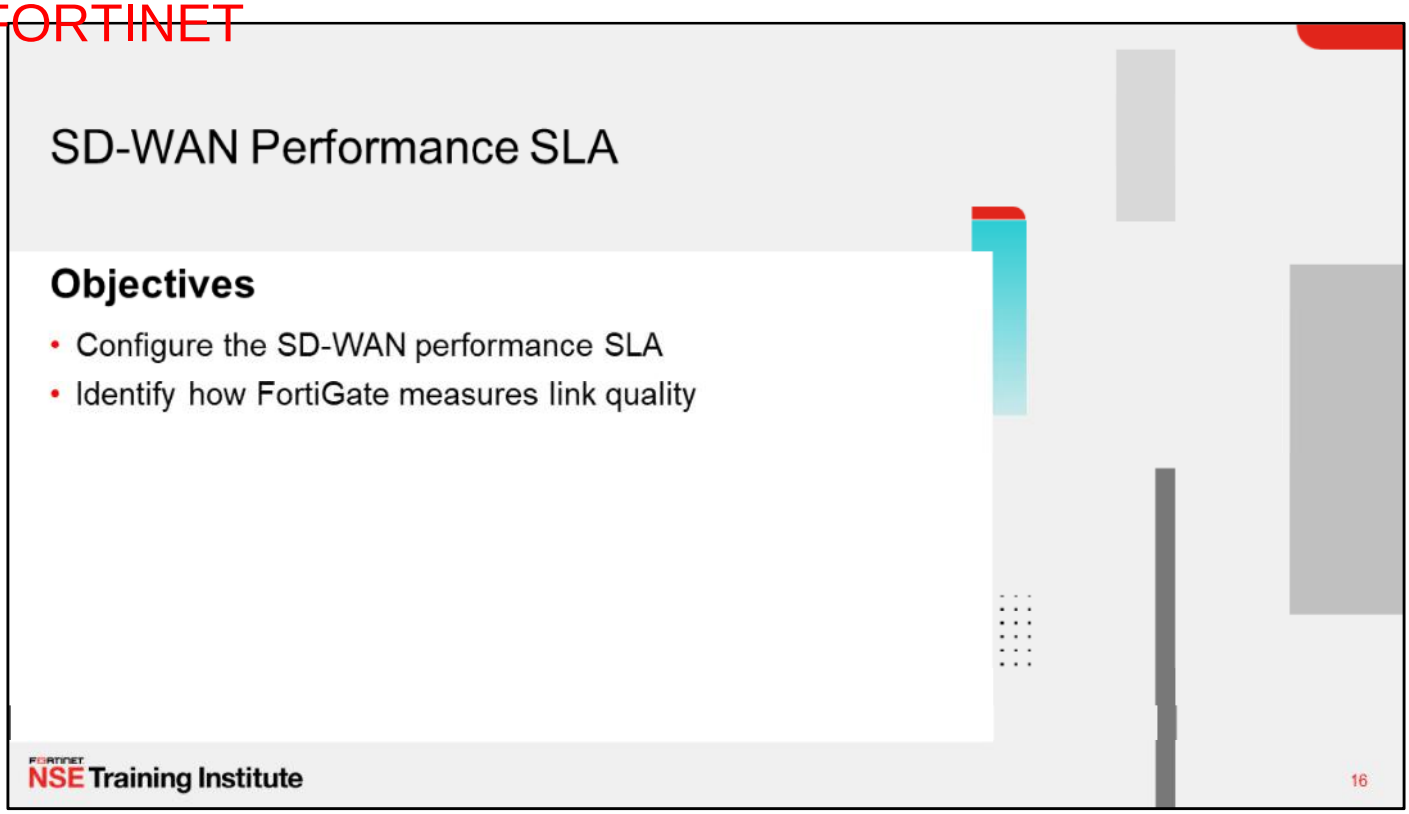

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the SD-WAN performance SLA, you should be able to configure the SD-WAN performance SLA, and identify how FortiGate measures link quality.

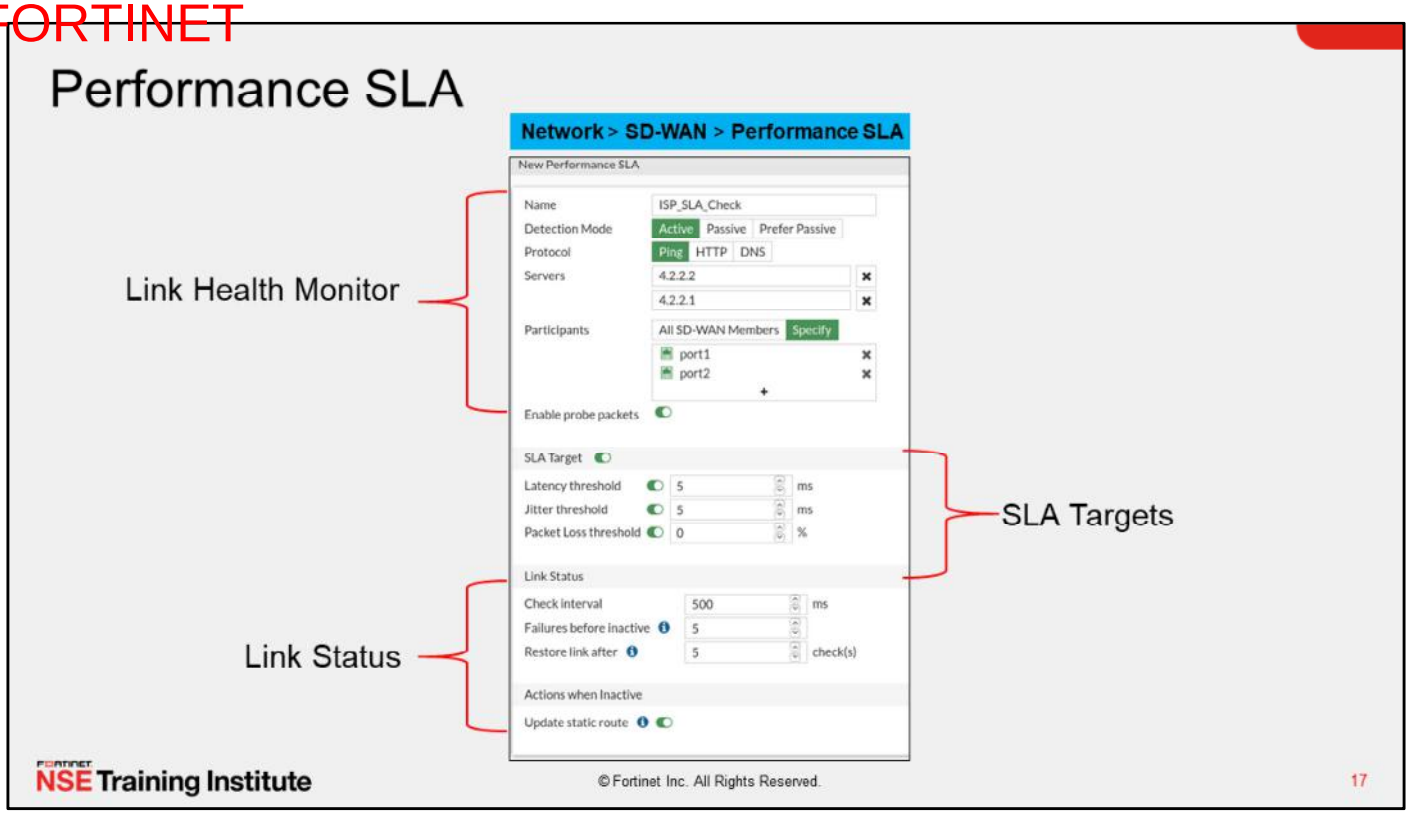

Now, you will learn about the two parts that make up the **Performance SLA** window: the link health monitor (or status check), and **SLA Targets**.

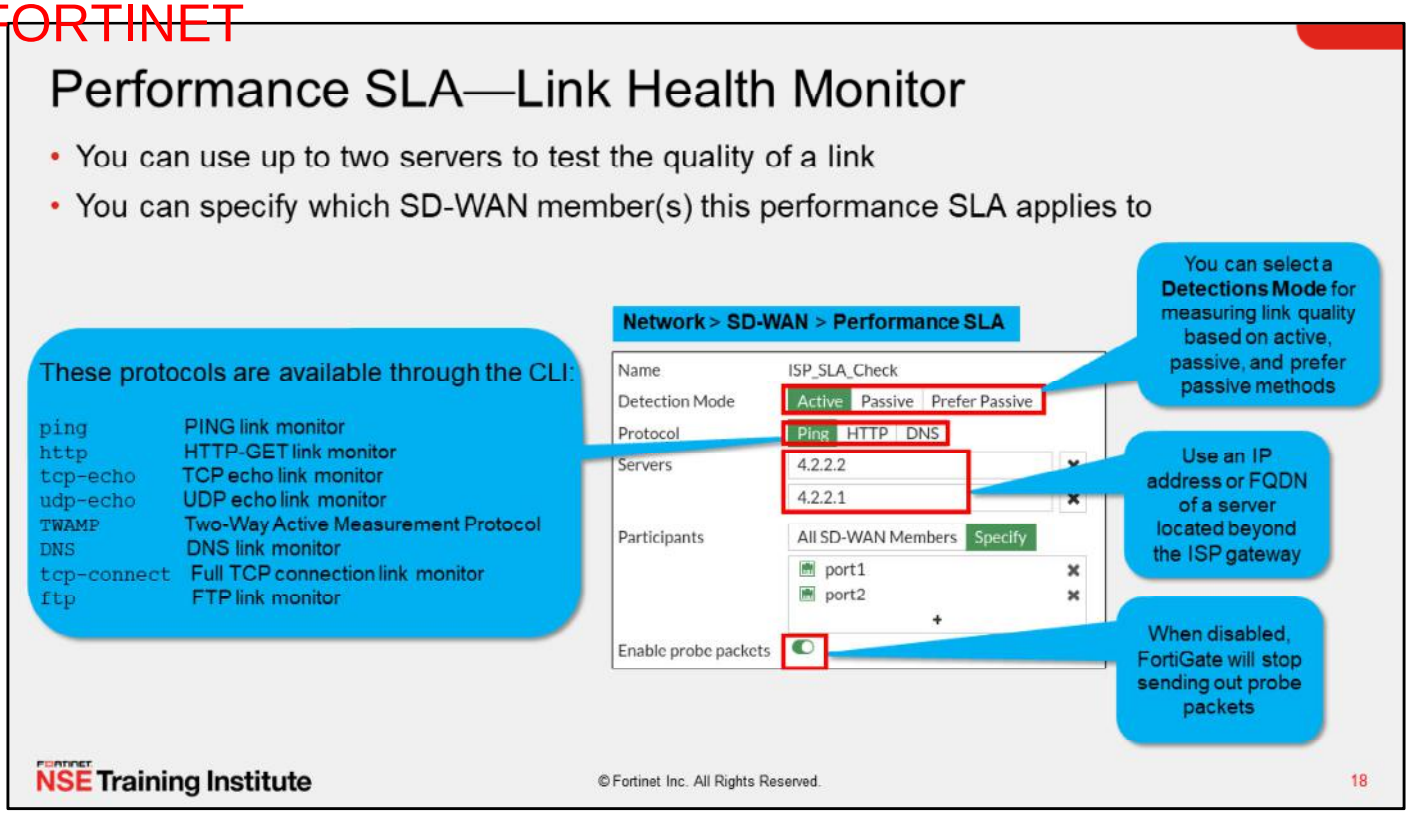

The link health monitor is a mechanism that detects when a router along the path is stopped or degraded.

FortiGate can check the status (or health) of each SD-WAN member interface participating in a performance SLA, based on the **Detection Mode** you have selected.

- **Active**: link health is measured by sending probe packets to the configured servers.
- **Passive**: link health is measured using session information that is captured on firewall policies that have `passive-wan-health-measurement` enabled.
- **Prefer Passive**: link health is measured using traffic passing through the SD-WAN members. When there is no traffic passing through the member interface, probe packets are sent to the configured servers to measure the link health.

You can specify up to two servers to act as your beacons. This guards against the server being at fault, and not the link.

On the GUI, you can either specify members or select all SD-WAN members as participants. On the CLI, you can now add member 0, which is equivalent to adding all members as participants for the particular performance SLA.

A FIB route entry is added in the kernel to reach the servers defined on the **Performance SLA** page over each participant interface. These kernel routes are flagged as `proto=17`. These kernel routes will act independently of the usual sources of routing.

The GUI provides three protocol options with which to perform the status check: **ping**, **HTTP**, and **DNS**, but on the CLI you have six options. Those options are: ping, HTTP, and DNS, just like on the GUI, but also TCP echo, UDP echo, and Two-Way Active Measurement Protocol (TWAMP).

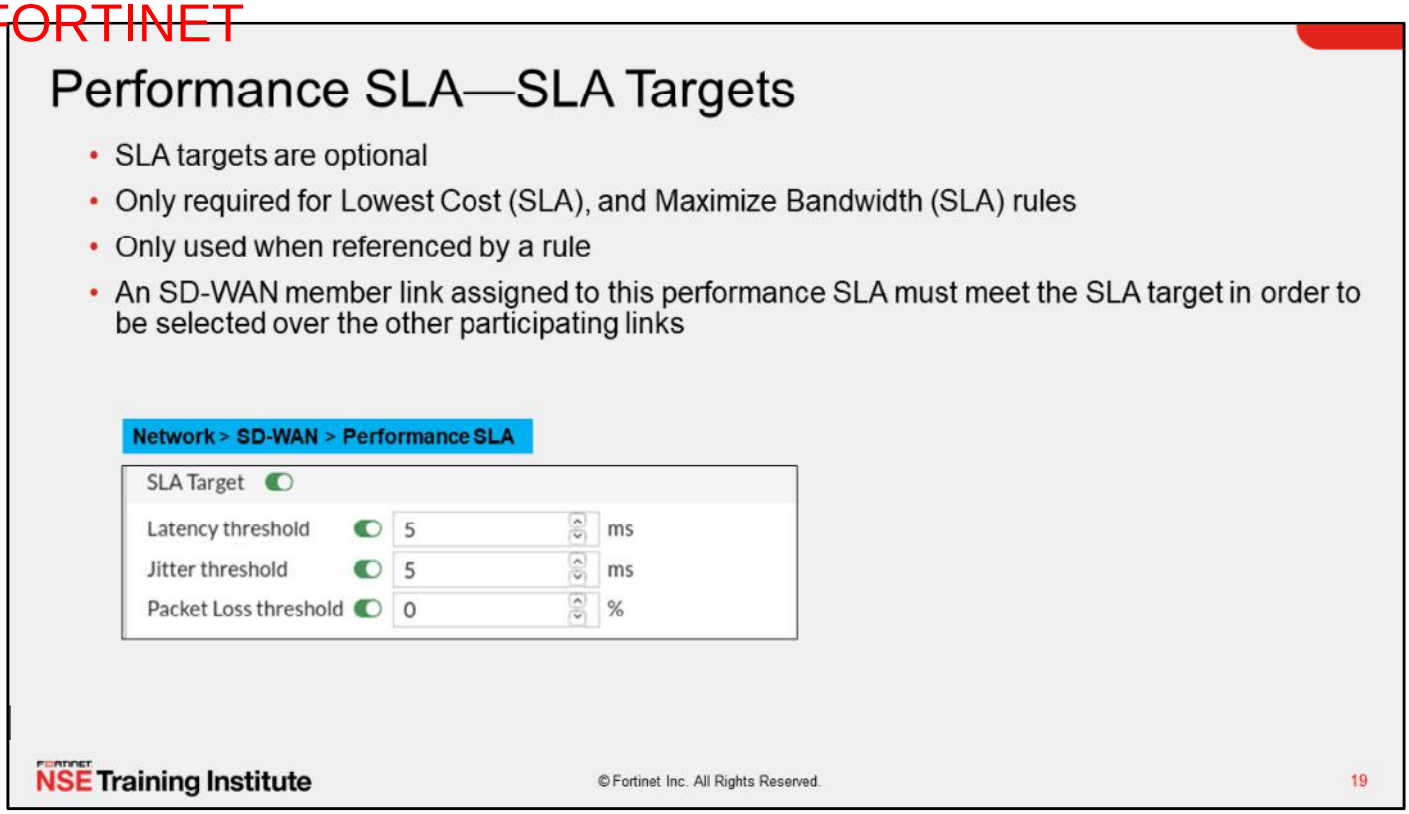

The quality of service for the traffic associated with this performance SLA is defined by the **SLA Targets.** An SD-WAN member link assigned to this performance SLA must meet the SLA target in order to be selected over the other participating links. You can configure the latency, jitter, and packet loss thresholds to meet your needs, and create granular SLA targets to fine-tune the SD-WAN for specific applications.

Although SLA targets are specified on the **Performance SLA** page, they are not actually used there. The values configured on that page, are used only when referenced by a rule. You can create multiple SLA targets per performance SLA, although there are few scenarios in which you would want to do that.

One scenario might be, if you are located in a branch office and use a few different applications that run on the same server headquarters. You could create one performance SLA to perform the health check on that server, but then have different SLA targets for the different applications. You could make the rules for some apps lenient, but more strict for others. If, however, the applications are running on different servers, then you would want to create different performance SLAs for each application, in order to have the health check go against the specific application server. And each performance SLA would require only one SLA target for that application.

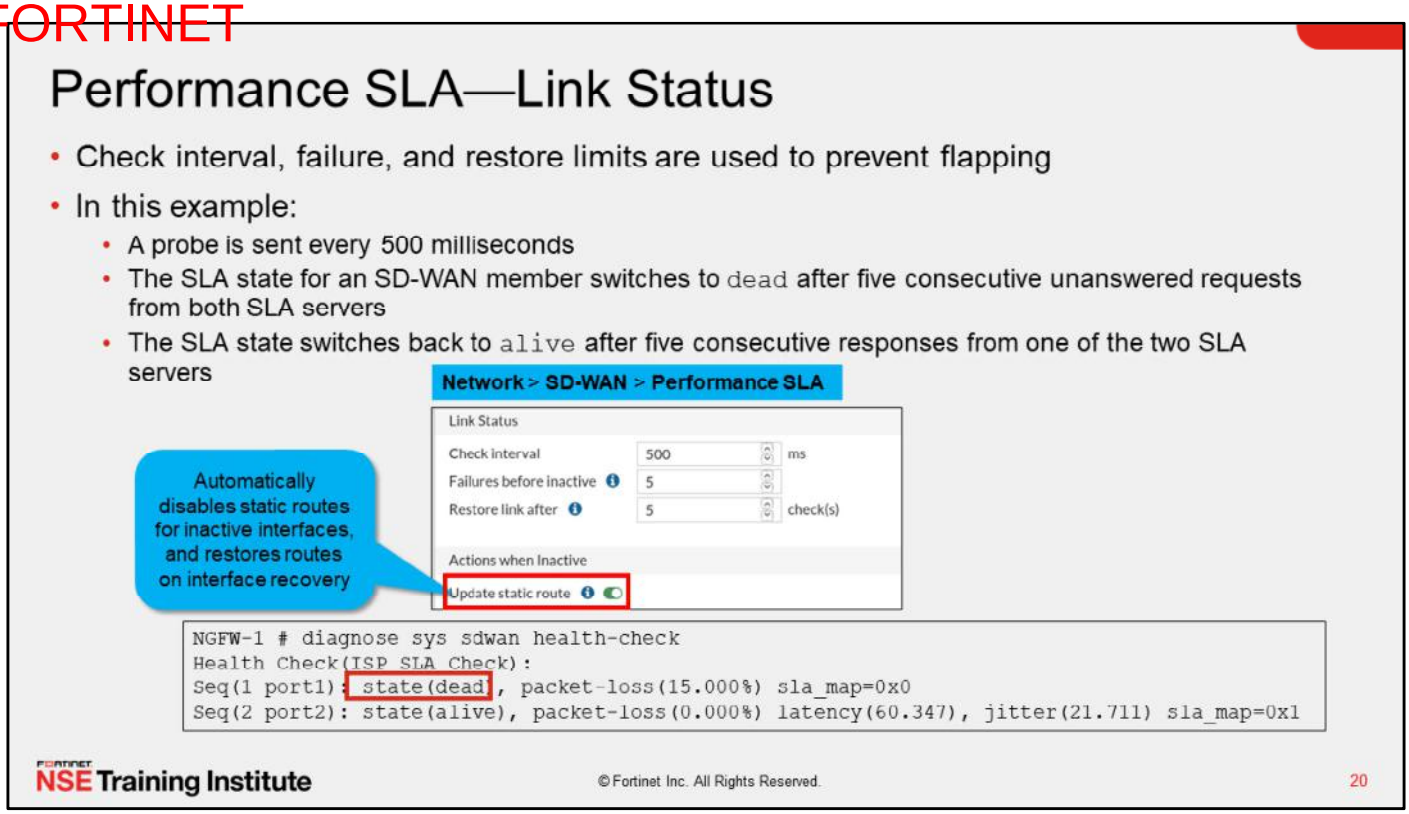

The **Link Status** contains settings that specify how often the system checks the link status to determine if it needs to transfer the traffic to another link. The **Failure before Inactive** and **Restore link after** settings prevent the system from continuously transferring traffic back and forth between links, a condition known as flapping.

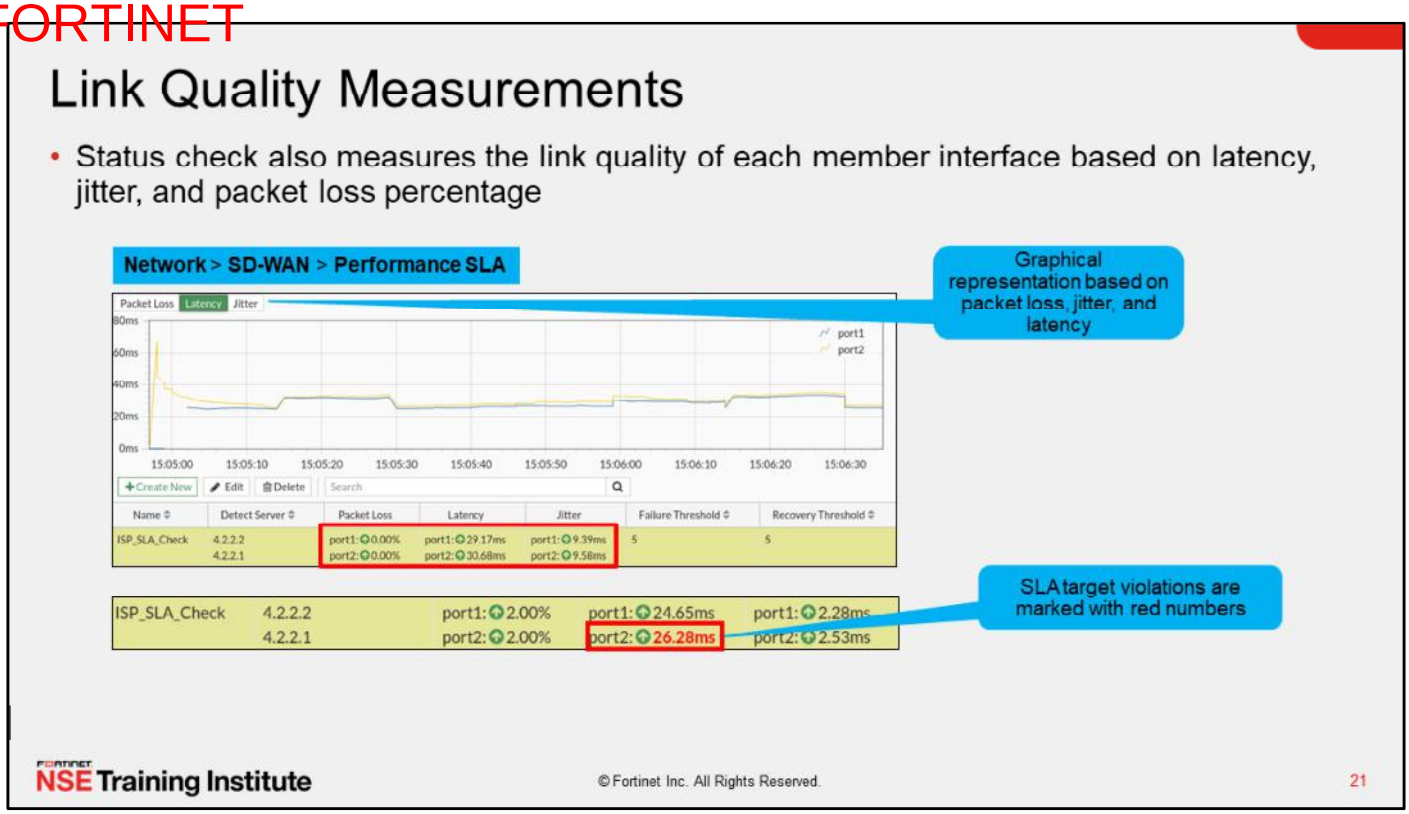

The **Performance SLA** (health checks) measures the quality of the links connected to the member interface participating in a performance SLA. Three different criteria are used for this measurement: latency, jitter, and packet loss percentage.

It's these values that are used against the SLA criteria within the rules that are used to route traffic based on the link quality of each member.

The **Packet loss, Latency,** and **Jitter** that are displayed are based on the replies (averaged over a short period) from the server that the performance SLA is using. The system starts with the first server. If that server becomes unavailable, then it switches to the second server. It stays with that second server until it becomes unavailable, at which point it goes back to the first server. If both servers are unavailable, then that performance SLA is deemed dead.

It is important to note that the green up arrows indicate only that the server is responding to the health check, regardless of the packet loss, latency, and jitter values. It is not an indication that any of the SLAs are being met.

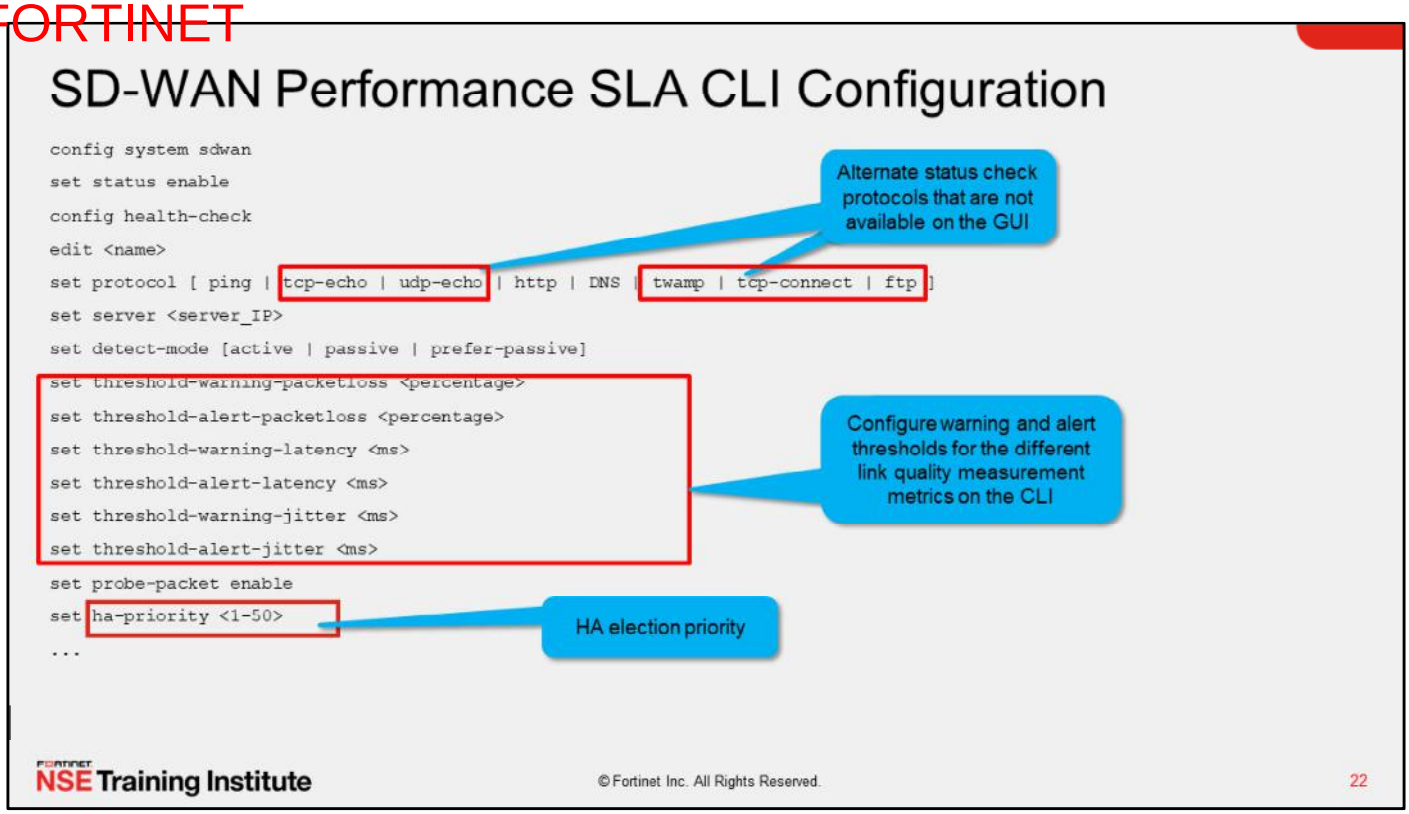

The CLI commands used to configure a performance SLA, provide more options. The `tcp-echo`, `udp-echo`, tcp-connect, ftp, and `twamp` options are available only on the CLI. These options provide different methods of measuring round-trip network performance between any two devices that support them. There are other CLI-only options that are available, based only on the performance SLA protocol you choose. For more information about these options, refer to the *CLI Reference Guide* on `docs.fortinet.com`.

You can configure the warning and alert thresholds for the latency, jitter, and packet loss quality checks. These are also not available on the GUI.

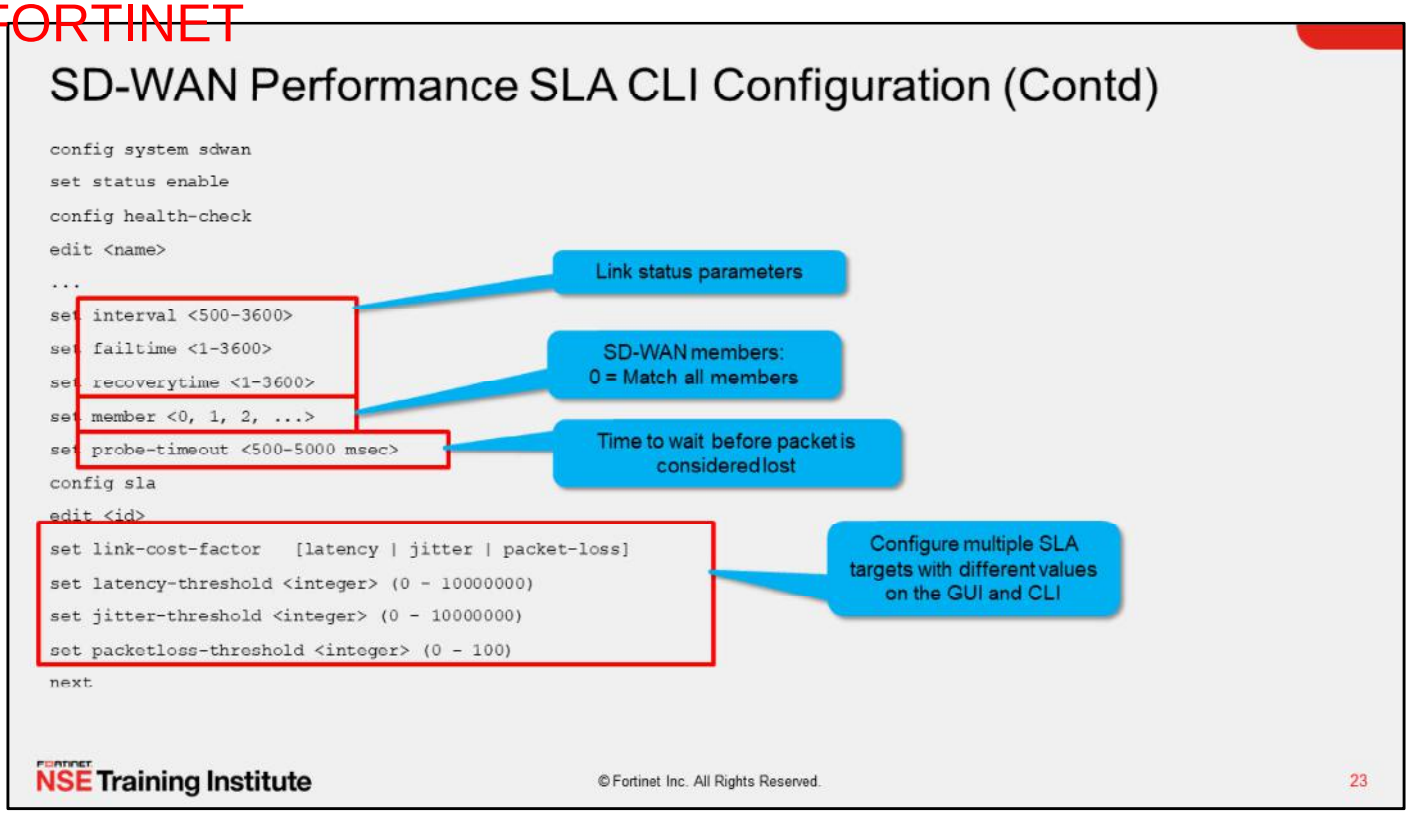

You can configure link status parameters using the commands shown on this slide. You can also configure multiple SLA targets with different values on both the GUI and CLI.

Knowledge Check

1. Which link attribute is used in SD-WAN link quality measurements?
   - A. Cost
   - ✓ B. Latency

2. Which status check protocol is available only on the CLI?
   - ✓ A. TCP-Echo
   - B. HTTP

NSE Training Institute

24

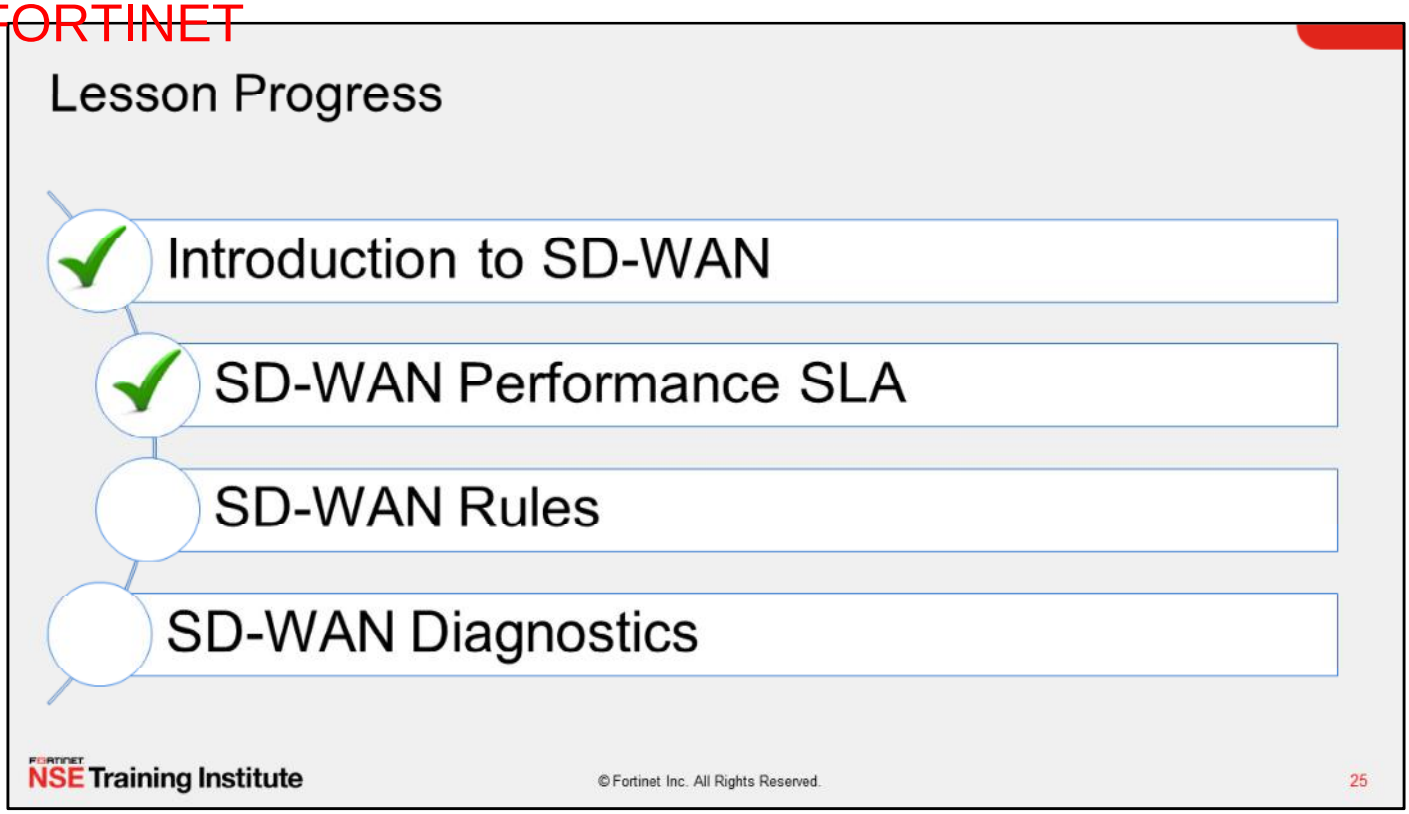

Good job! You now understand the SD-WAN performance SLA.

Now, you will learn about SD-WAN rules.

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

By demonstrating competence in SD-WAN rules, you should be able to configure dynamic link selection based on link quality to ensure high availability for business-critical applications.

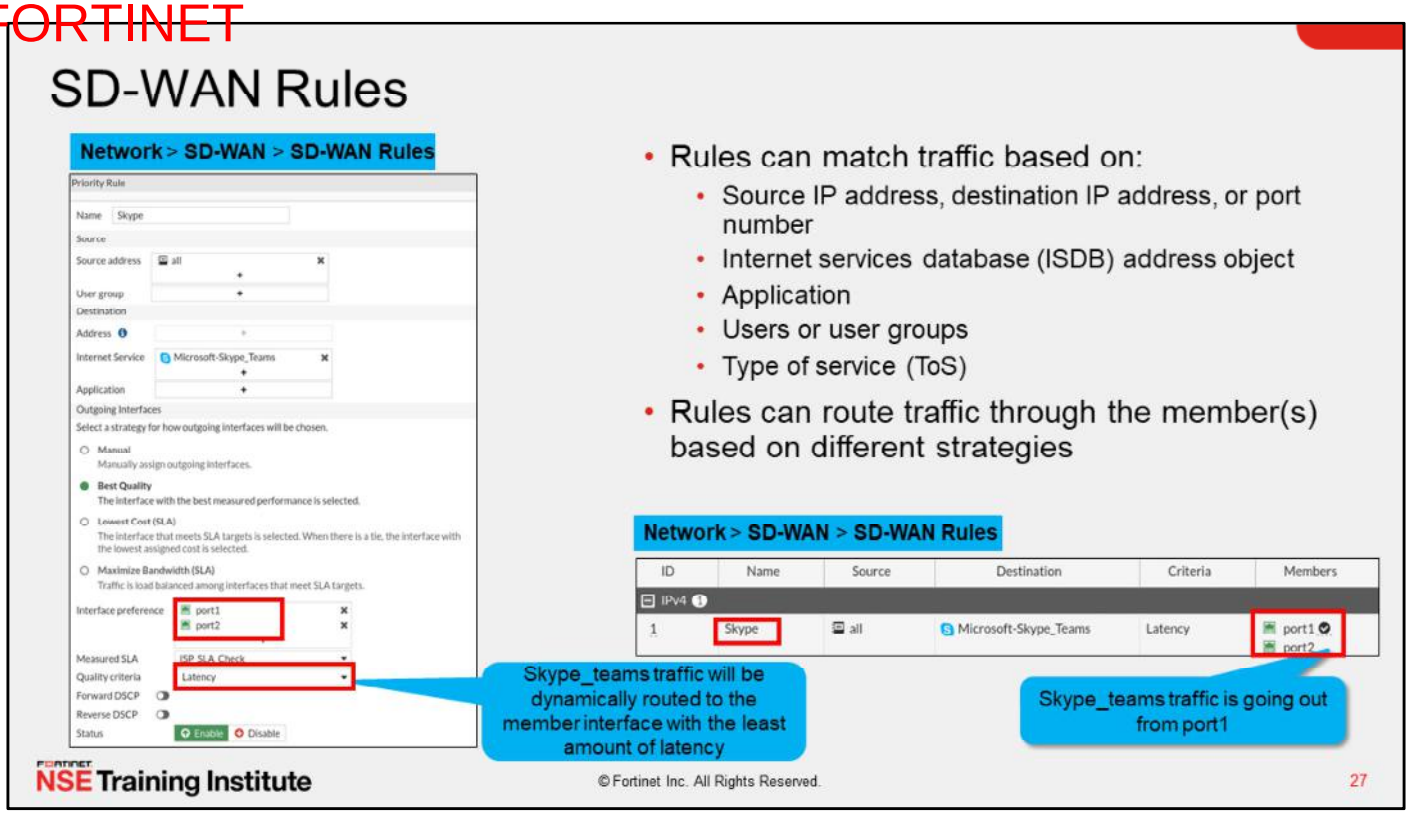

SD-WAN rules allow you to specify which traffic you want to route through which interface. You can configure the SD-WAN rules to choose the egress interface according to latency, jitter, or packet loss percentage, based on the settings you configured on the **SLA Targets** page of a link. The rules are evaluated in the same way as firewall policies: from top to bottom, using the first match. You can use the following parameters to match the traffic:

- Source IP address
- Destination IP address
- Destination port number
- ISDB address objects as destination
- Firewall application as destination
- Users or user groups
- Type of service (ToS)

SD-WAN rules offer great flexibility when matching traffic. For example, you can route Netflix traffic from specific authenticated users through one ISP, while routing the rest of your internet traffic through another ISP.

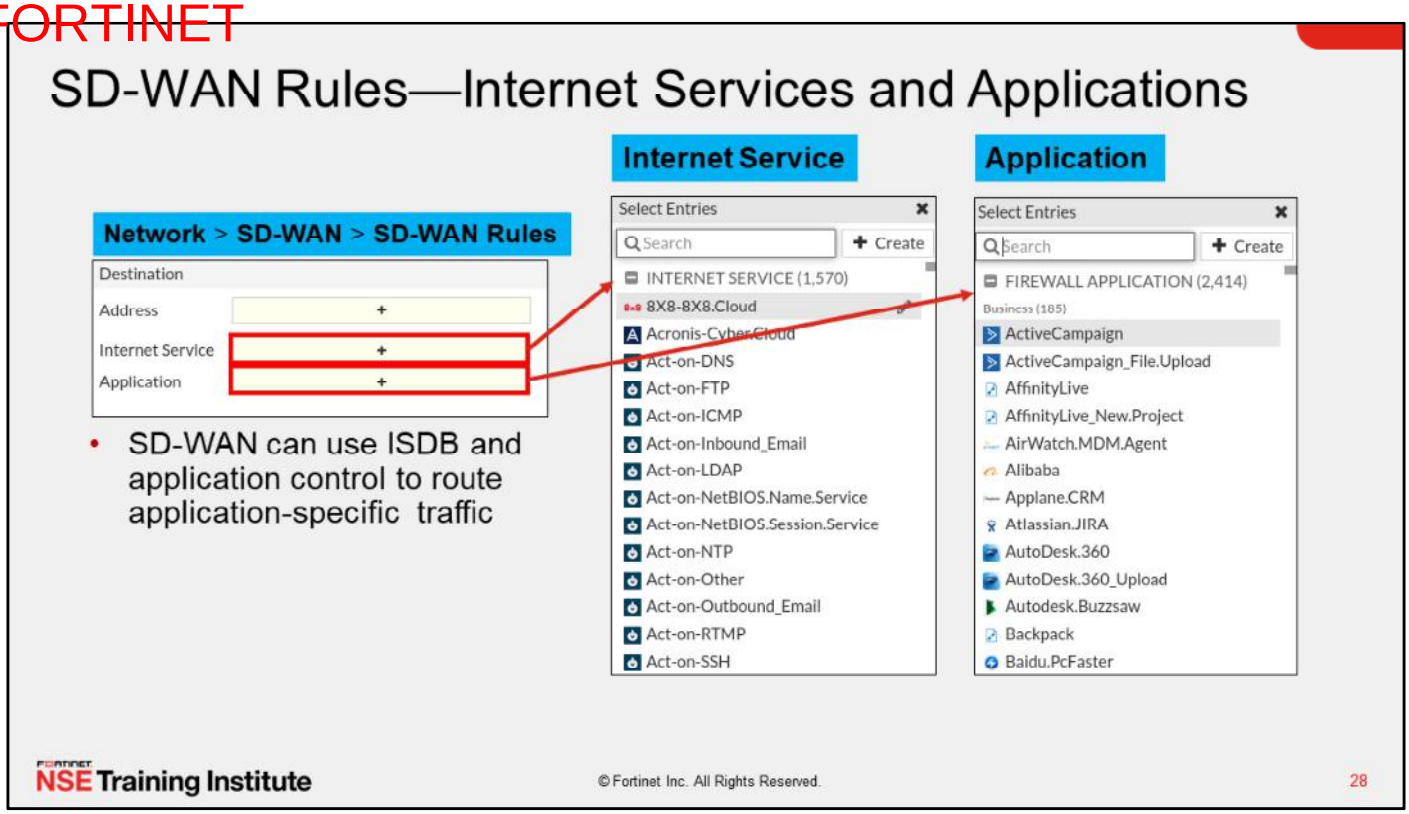

SD-WAN can use the Internet Services database, as well as the Application Control database to steer applications along a specific link.

FortiGuard maintains these databases, and FortiGate periodically gets an updated copy. When using the Application Control database, you should enable SSL inspection for the most accurate application identification.

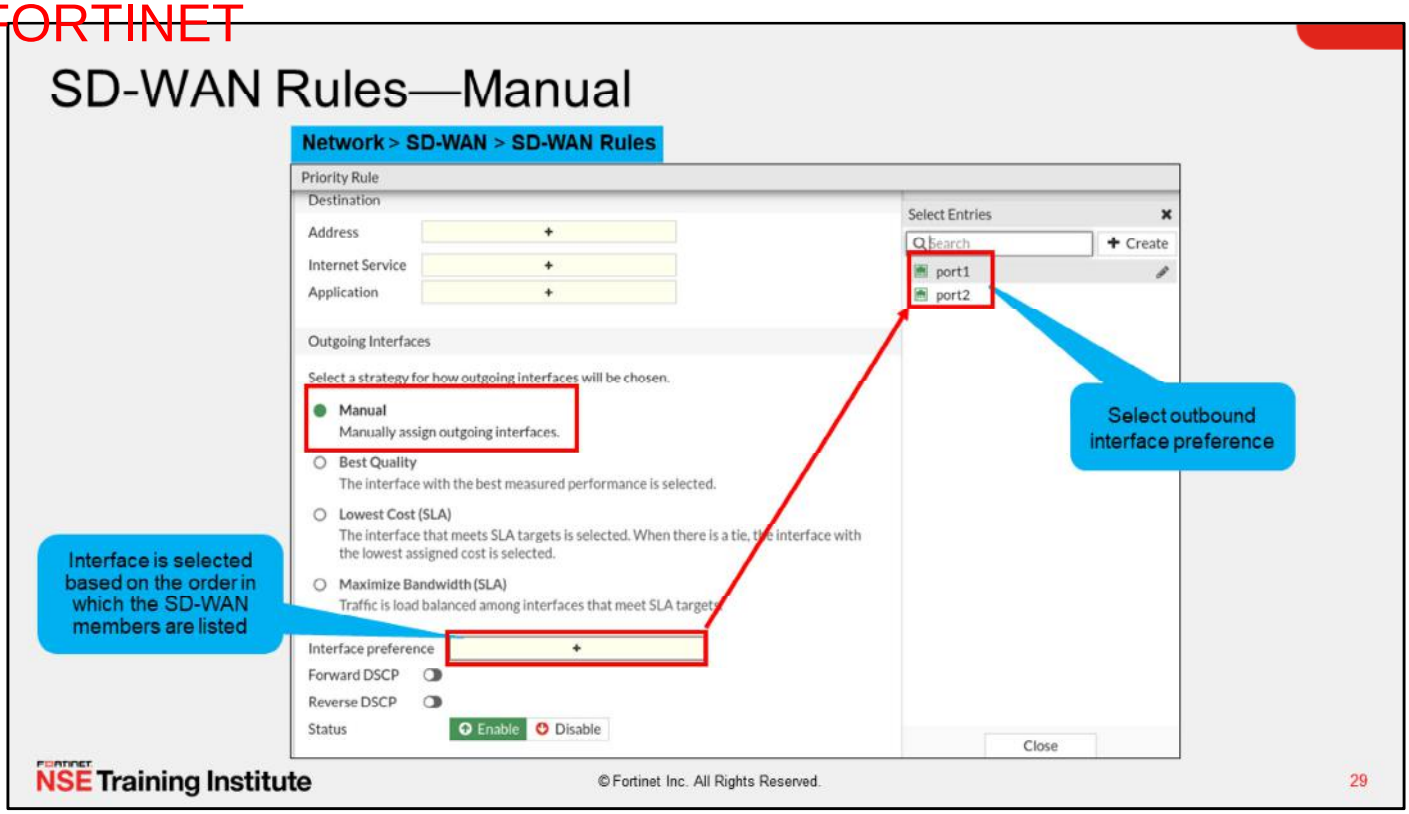

FortiGate SD-WAN offers four strategies for selecting outgoing interface(s): **Manual**, **Best Quality**, **Lowest Cost (SLA)**, and **Maximize Bandwidth (SLA)**.

If you select **Manual**, you can specify the interface priority you want to send traffic out from. If the traffic matches the rule criteria, the traffic will go out from the first available interface based on the interface preference. This strategy does not depend on performance SLA or SLA targets.

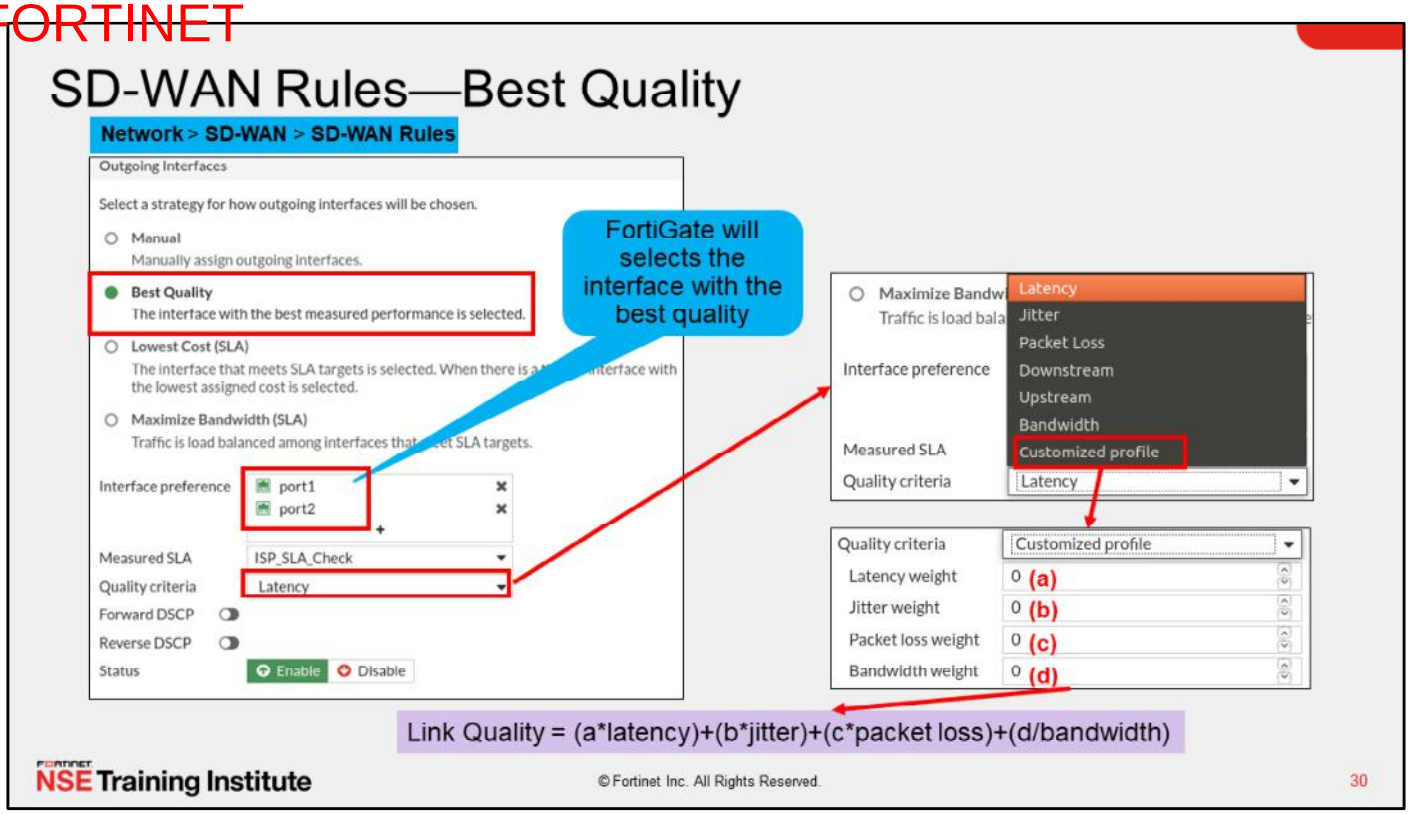

The best quality strategy is based on the performance of the network. In the example shown on this slide, port1 and port2 are included in the interface preference. So, port1 is used (because it is the first on the list) until the quality criteria of that network is 10% worse than that of port2, at which point port2 would take over. By default, the quality criteria is 10%, but you can change it on the CLI using the **`set link-cost-threshold`** command. Note that you don't use any of the SLAs here. The quality check on the performance SLA (DC_PBX_SLA) is using only the server information (health check) against the quality criteria. You can use the options of latency, jitter, and packet loss percentage. You can also use the bandwidth options (**Downstream** bandwidth, **Upstream** bandwidth, or **Bidirectional** bandwidth) so that FortiGate selects the link based on the available bandwidth of incoming, outgoing, or bidirectional traffic. This is useful because users may use some applications primarily for downloading, and other applications primarily for uploading.

The last option, **custom profile-1**, allows you to base the link selection on a combination of its criteria values. The link quality is determined by the equation.

The larger the value, the more weight that criteria will have in the selection. Leave the weight value at zero to exclude that criteria from the equation.

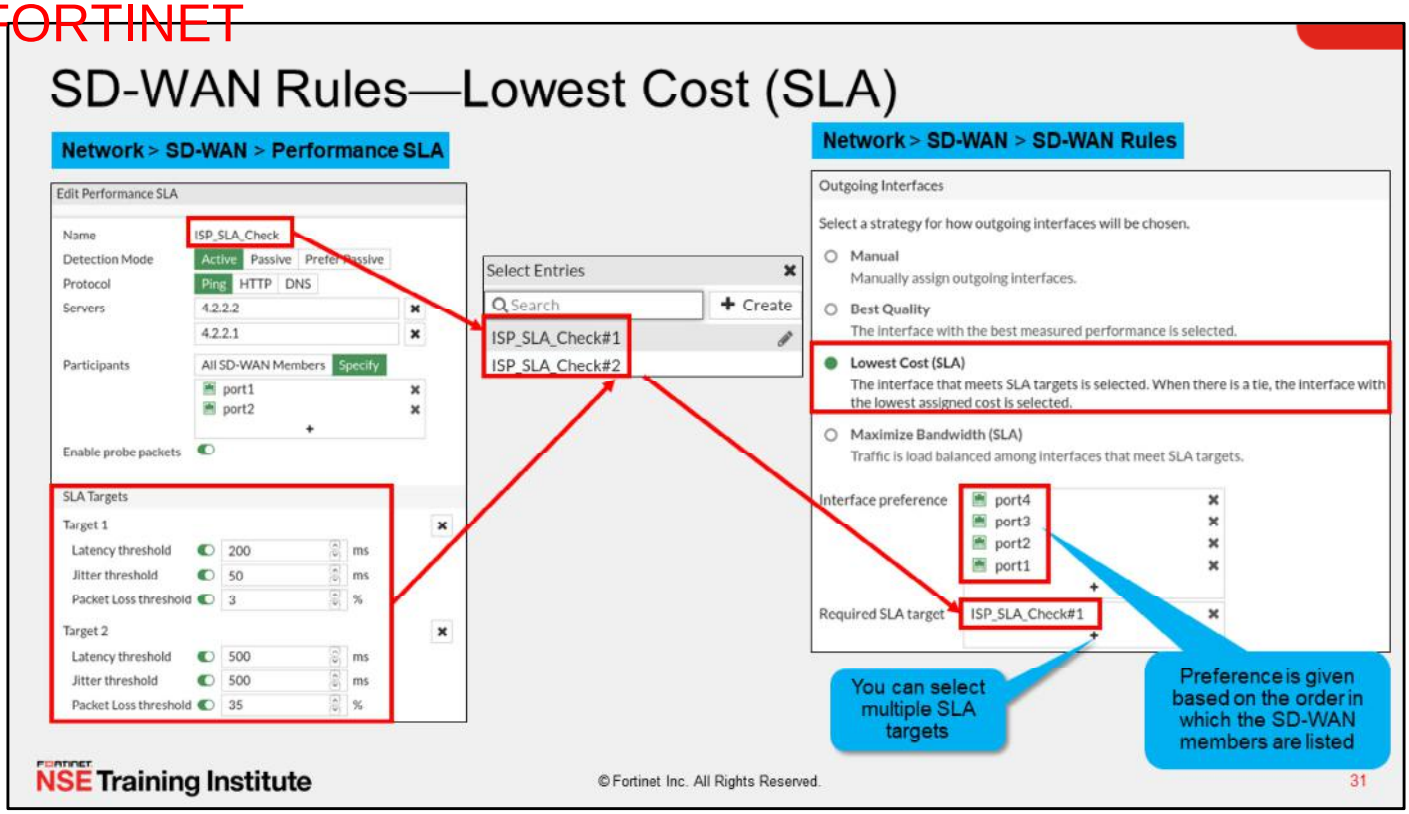

When you use the lowest cost (SLA) strategy, you select an SLA target from a performance SLA that you want to measure the traffic against. Note that even if a performance SLA has multiple SLA targets, you can select only one of the SLA targets from that particular performance SLA.

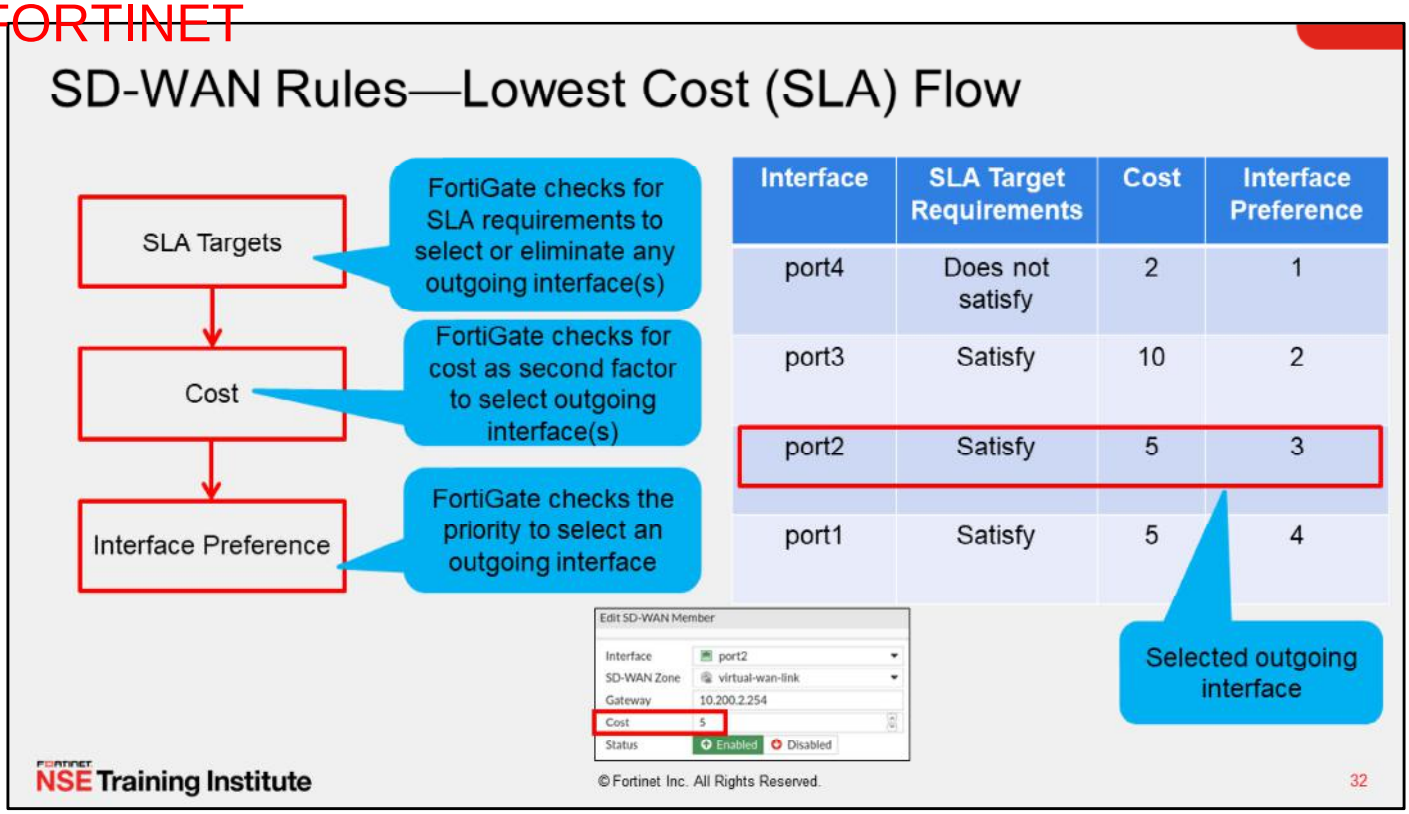

(Animation at the end of the slide)

FortiGate follows the flow shown on the slide to select an outgoing interface for **Lowest Cost (SLA)**.

For example, SD-WAN has four members: interface1, 2, 3, and 4. Refer to the table shown on the slide.

First, FortiGate considers the SLA and eliminates interface 4 from the outgoing interface selection. Based on the SLA, FortiGate will consider three interfaces.

Then, FortiGate considers the cost to eliminate any interface from consideration. You can configure the cost in **Network** > **SD-WAN** under **SD-WAN Interface Members**. FortiGate selects an interface with a lower cost. In the example shown on this slide, interfaces 1 and 2 have cost 5, and interface 3 has cost 10 set. In this case, FortiGate eliminates interface 3 from consideration and considers interface 1 and 2 for the outgoing interface.

Lastly, FortiGate checks the priority set for all the interfaces to select an outgoing interface. In the example shown on this slide, interface 2 has a higher priority than interface 1. In this case, FortiGate selects interface 2 as the outgoing interface for the traffic matching the rule. (click)

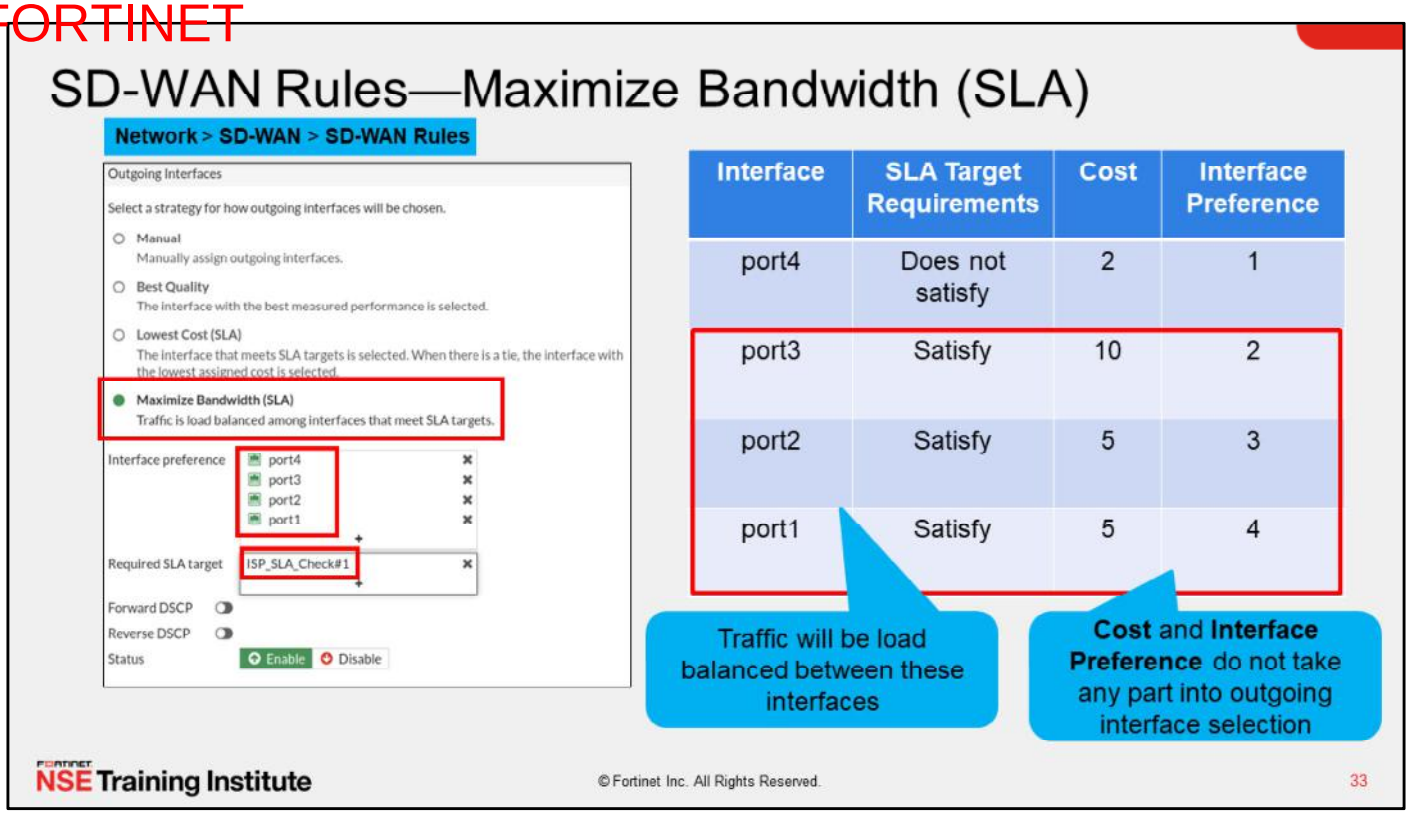

(Slide Contains animation)

This feature introduces a new load balance mode for the SD-WAN rule. If traffic matches the rule specifications, the traffic is load balanced among the selected members, which satisfies the SLA specification. If there are multiple SLA criteria, traffic is load balanced only to the members satisfying the most SLA criteria.

The example on this slide shows that interfaces 1, 2, and 3 satisfy the SLA requirements, and interface 4 does not. In this case, the traffic matching the rule is load balanced between interfaces 1, 2, and 3. (Click)

Using this method, FortiGate do not take cost or priority into consideration.

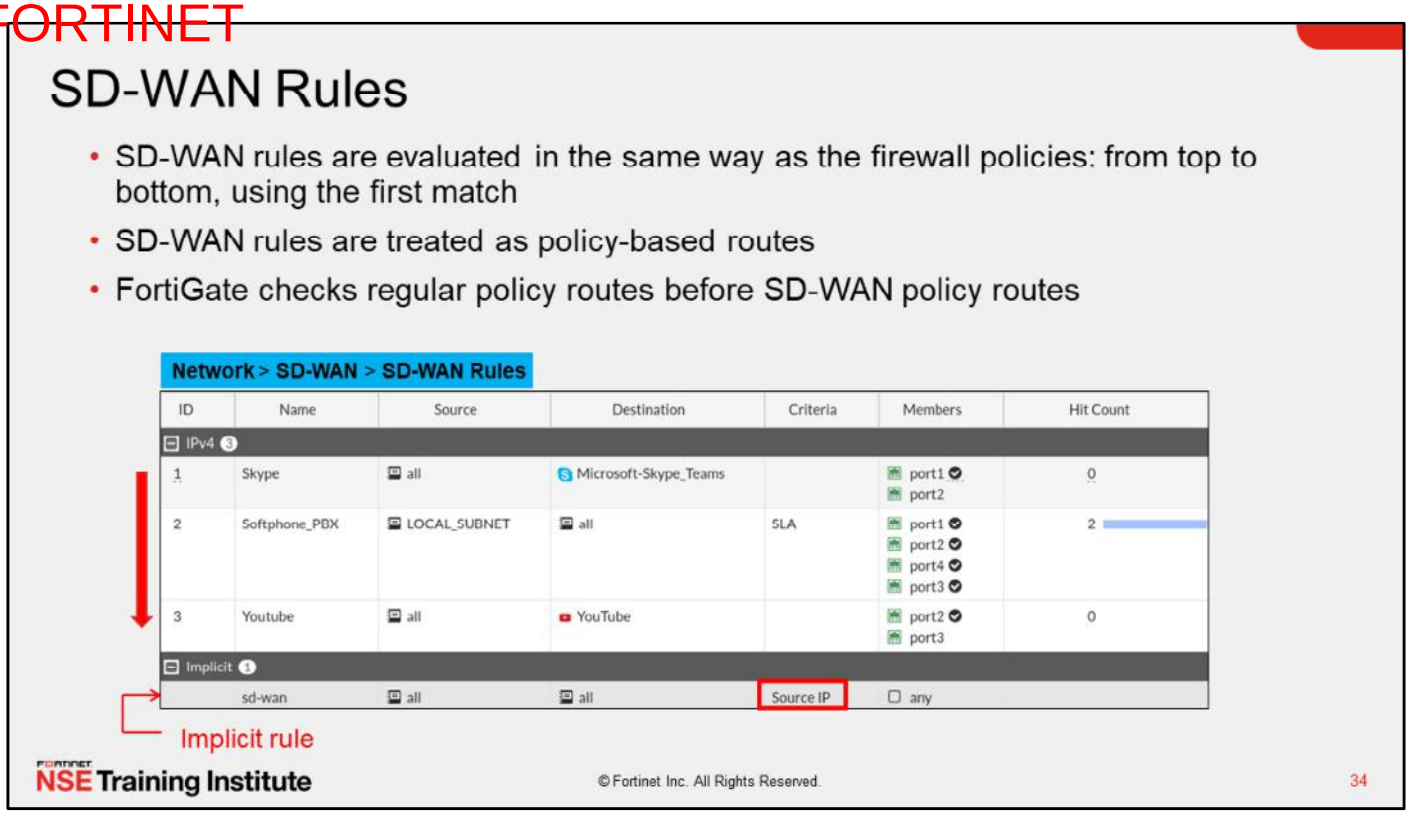

Application-specific rules are evaluated in the same way as the firewall policies: from top to bottom, using the first match.

An implicit rule is automatically generated when you enable SD-WAN. If none of the conditions of any of the other rules are met, then the implicit rule is used. This implicit rule is designed to balance the traffic among all the available SD-WAN member links.

Double-clicking the implicit rule displays the load balancing options.

Similar to ISDB routes, SD-WAN rules function as policy routes. They take precedence over any other routes in the routing table. When it comes to policy routing, FortiGate checks regular policy routes first, before checking SD-WAN policy routes.

Knowledge Check

1. What is an SD-WAN rule matching parameter for traffic sources?
   - ✓ A. User groups
   - B. IPS signatures

2. You can configure SD-WAN rules to choose the egress interface based on which parameter?
   - A. Weight
   - ✓ B. Latency

NSE Training Institute

35

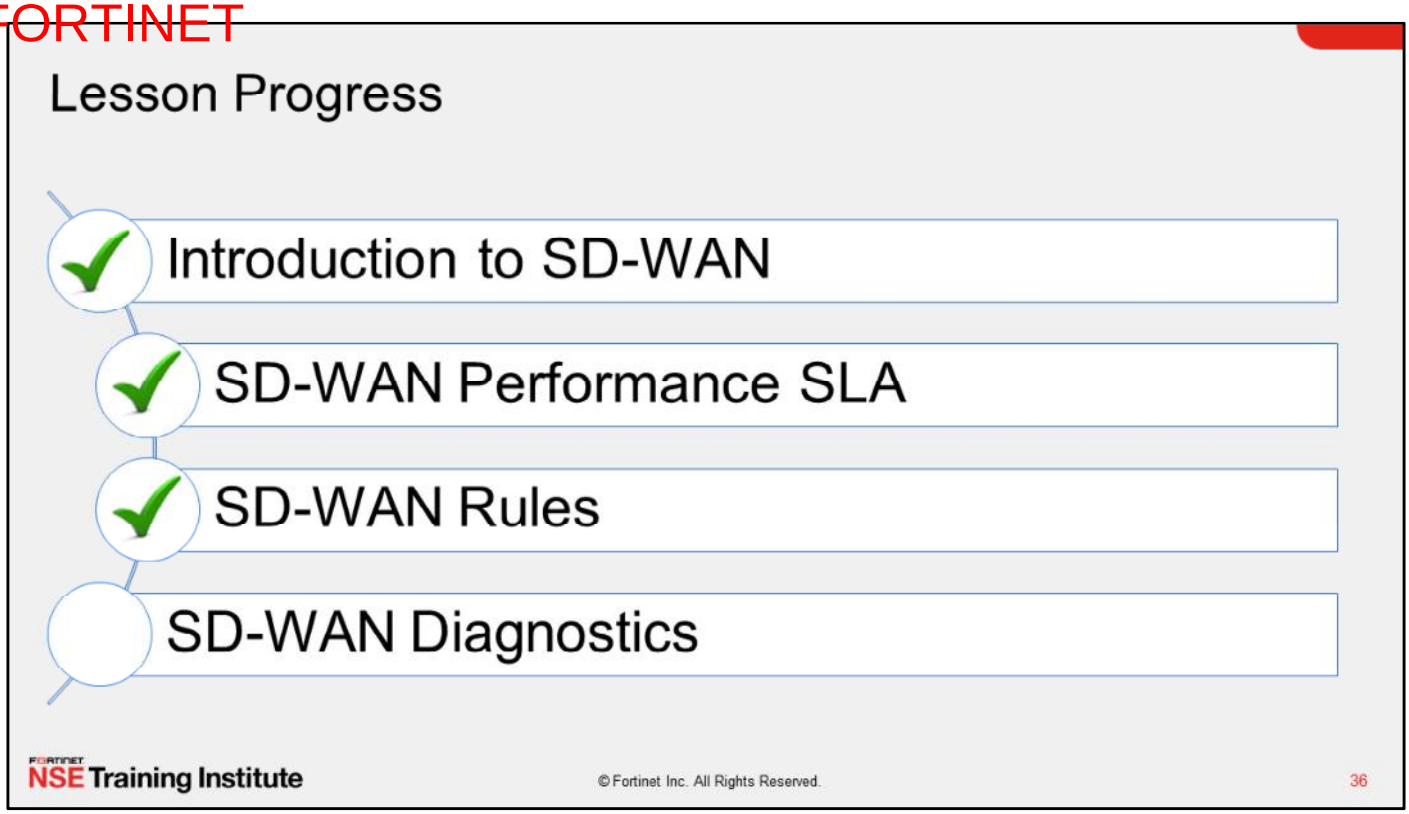

Good job! You now understand SD-WAN rules.

Now, you will learn about SD-WAN diagnostics.

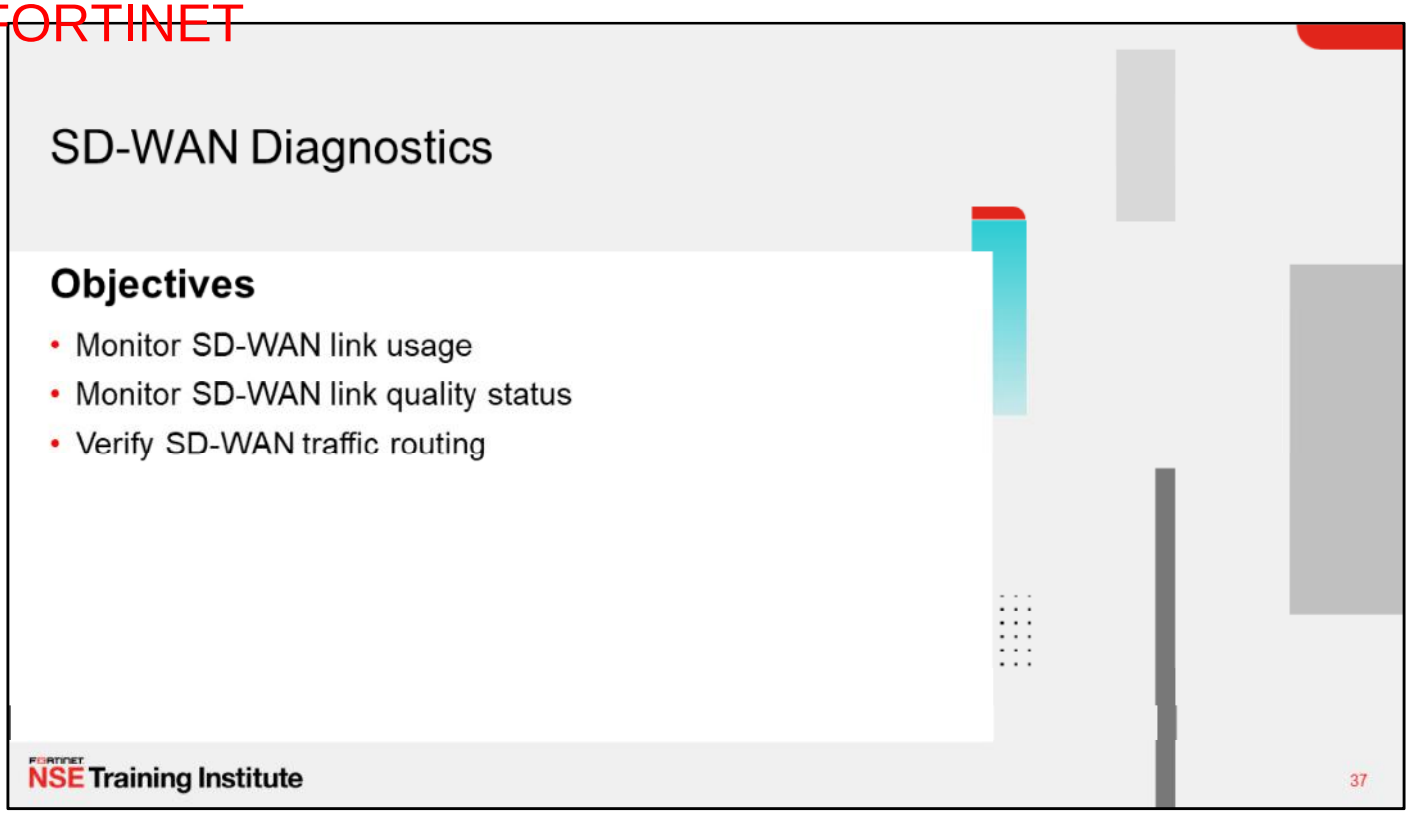

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in SD-WAN diagnostics, you should be able to maintain an efficient and effective SD-WAN solution.

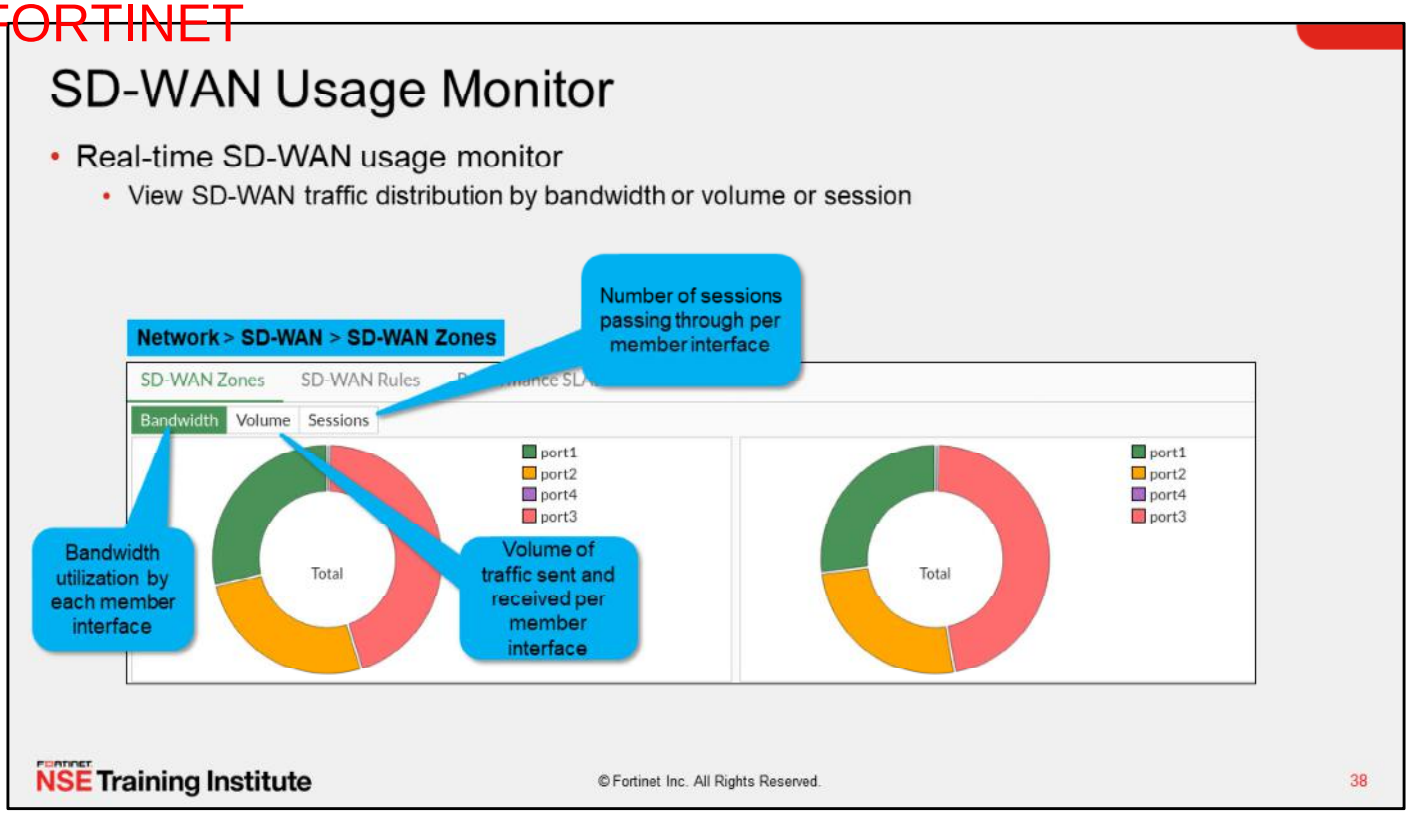

You can use the SD-WAN usage monitor to view traffic distribution between the member interfaces, based on bandwidth or volume.

The **Volume** view gives a better representation of the traffic sent and received across all the member interfaces; whereas the **Bandwidth** view shows you how much bandwidth each interface is using as a result of the sessions passing through them. The **Sessions** view shows the number of sessions passing through for each interface.

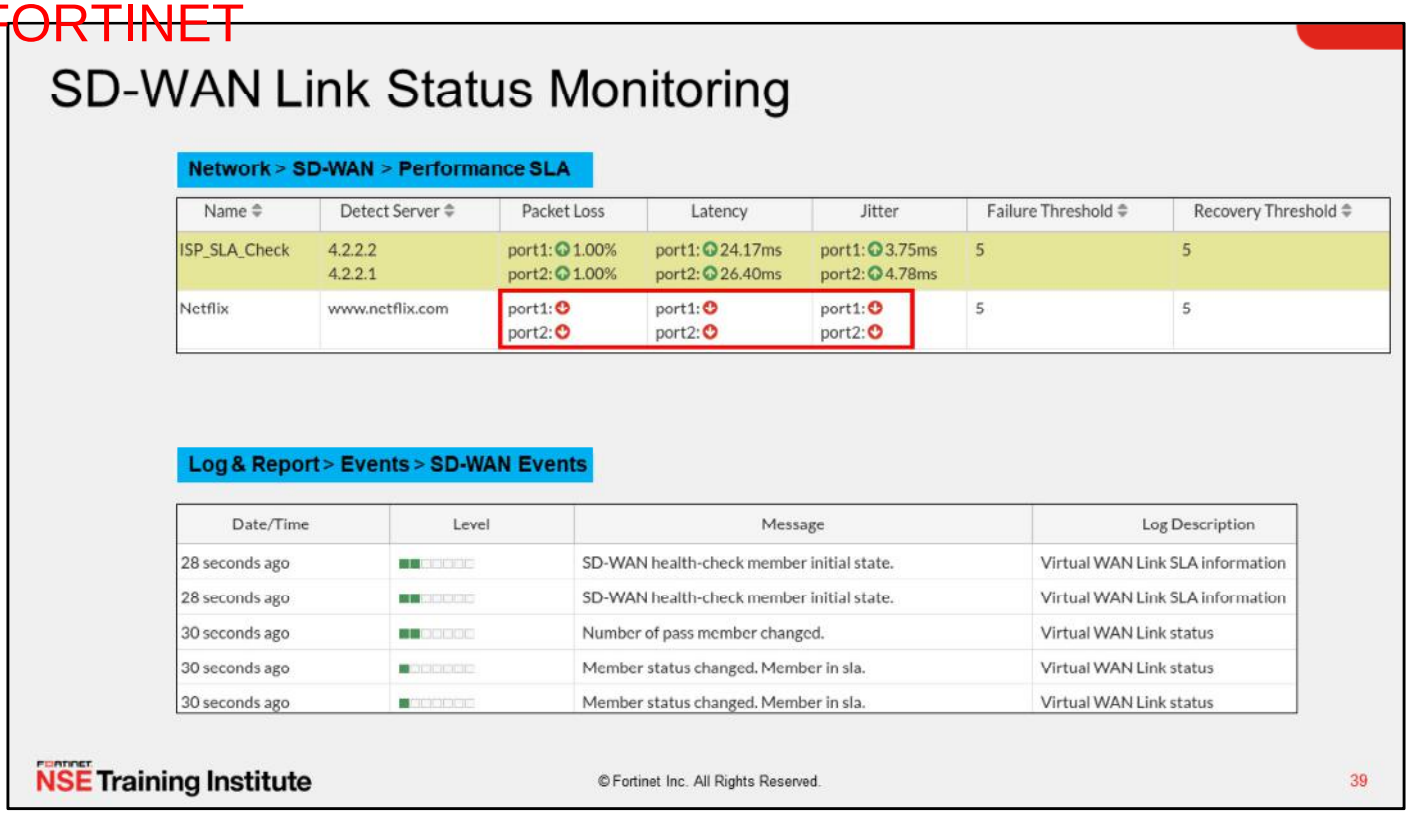

Because link quality plays a big role in link selection when using SD-WAN, monitoring the link quality status of the SD-WAN member interfaces is a good practice. You should investigate any prolonged issues with packet loss and latency to ensure your network traffic does not experience outage or degraded performance. Green arrows indicate interfaces are active in the SD-WAN group. Red arrows indicate that the interface is inactive for that specific status check.

FortiGate also generates system event logs when the route of an SD-WAN member interface is removed or added to the routing table. Use **Events** logs to review logs.

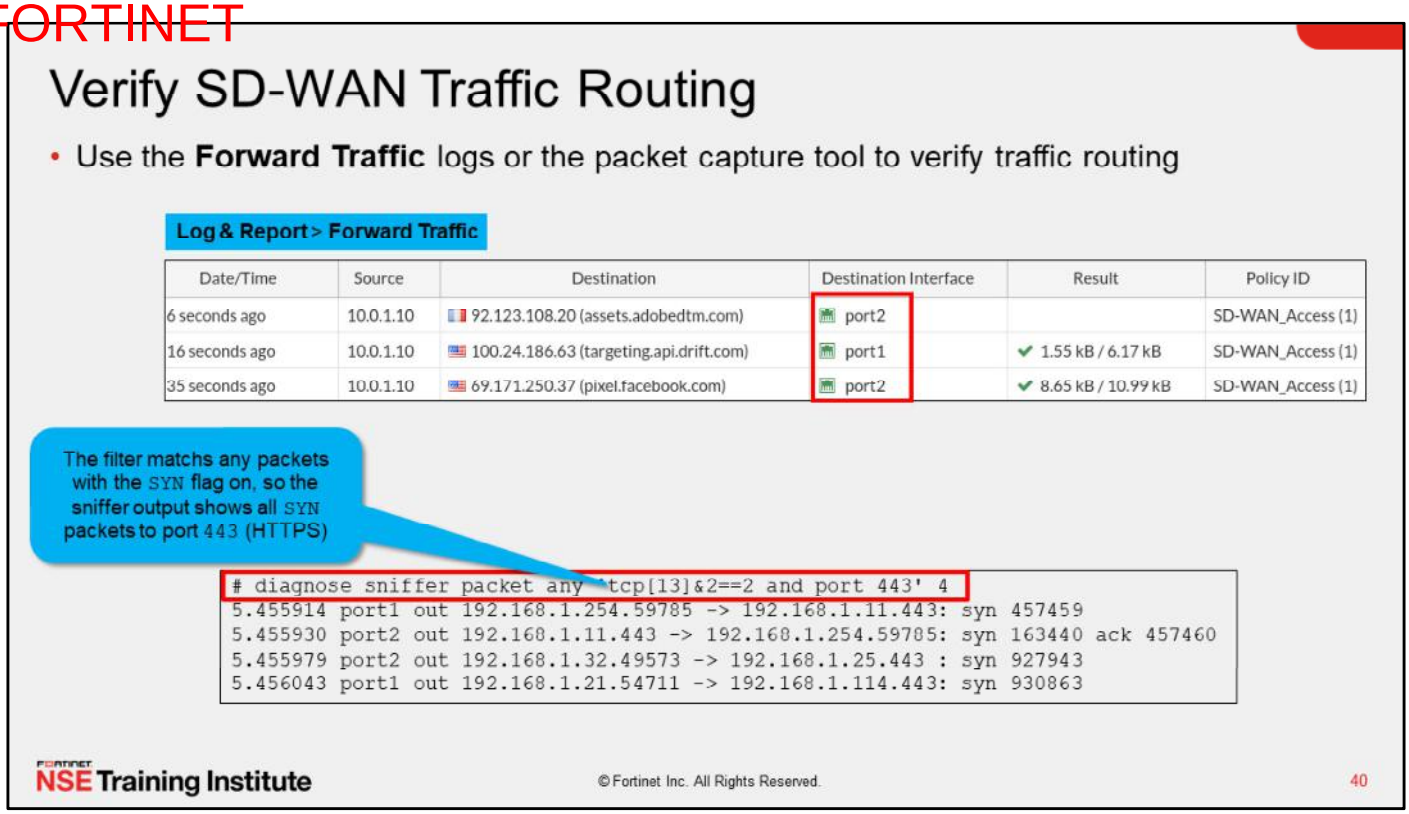

You can use the **Destination Interface** column in the **Forward Traffic** logs to verify that traffic is egressing the SD-WAN member interfaces. Alternatively, you can use verbosity levels `4` and `6` to view the egress interface using the CLI packet capture tool.

Knowledge Check

1. Which monitor should you use to monitor the session distribution across the SD-WAN member interfaces?
   - A. SD-WAN Link Status monitor
   - ✓ B. SD-WAN Usage monitor

2. When verifying SD-WAN traffic routing with the CLI packet capture tool, which verbosity level should you use?
   - A. 1
   - ✓ B. 4

NSE Training Institute

41

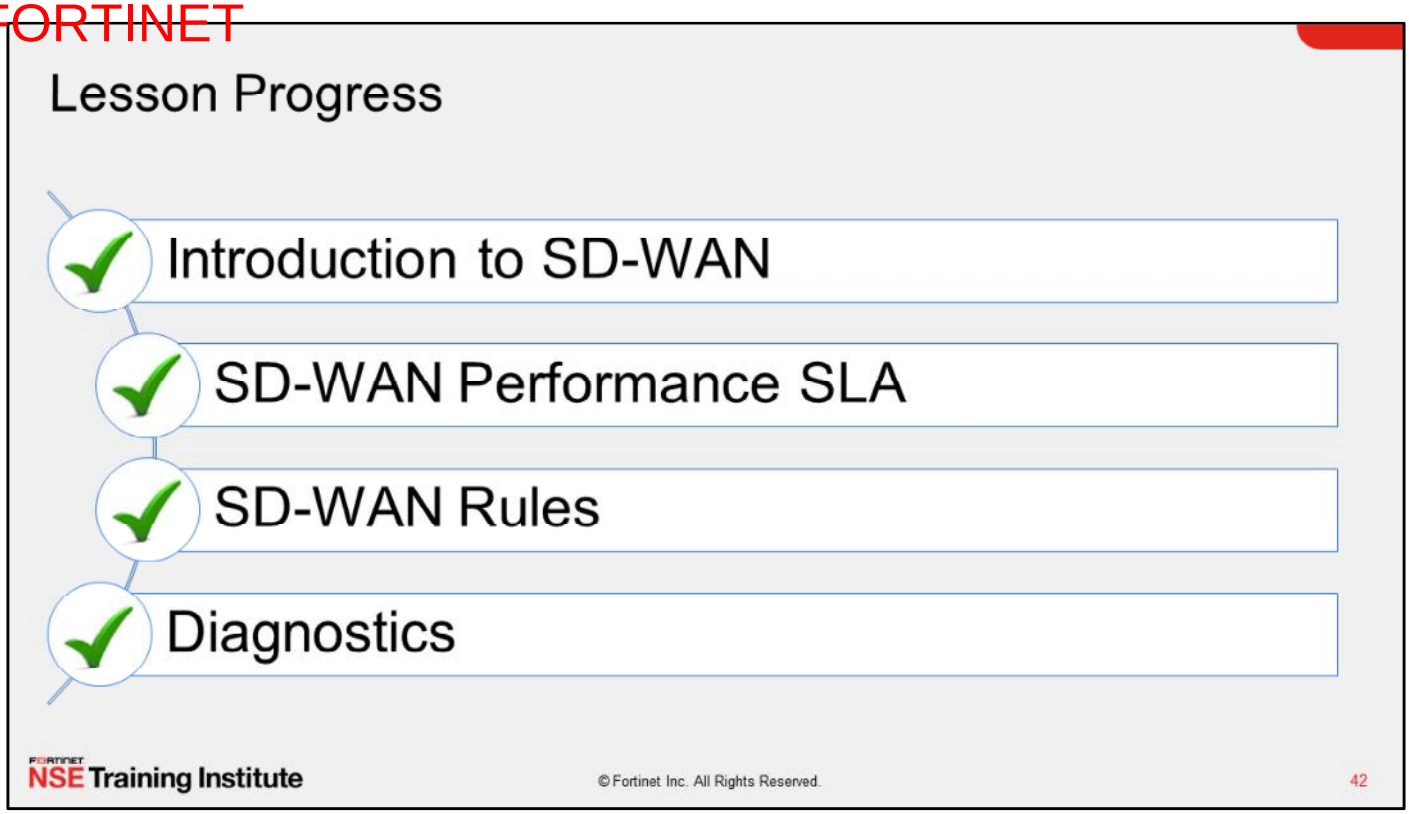

Congratulations! You have completed this lesson.

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

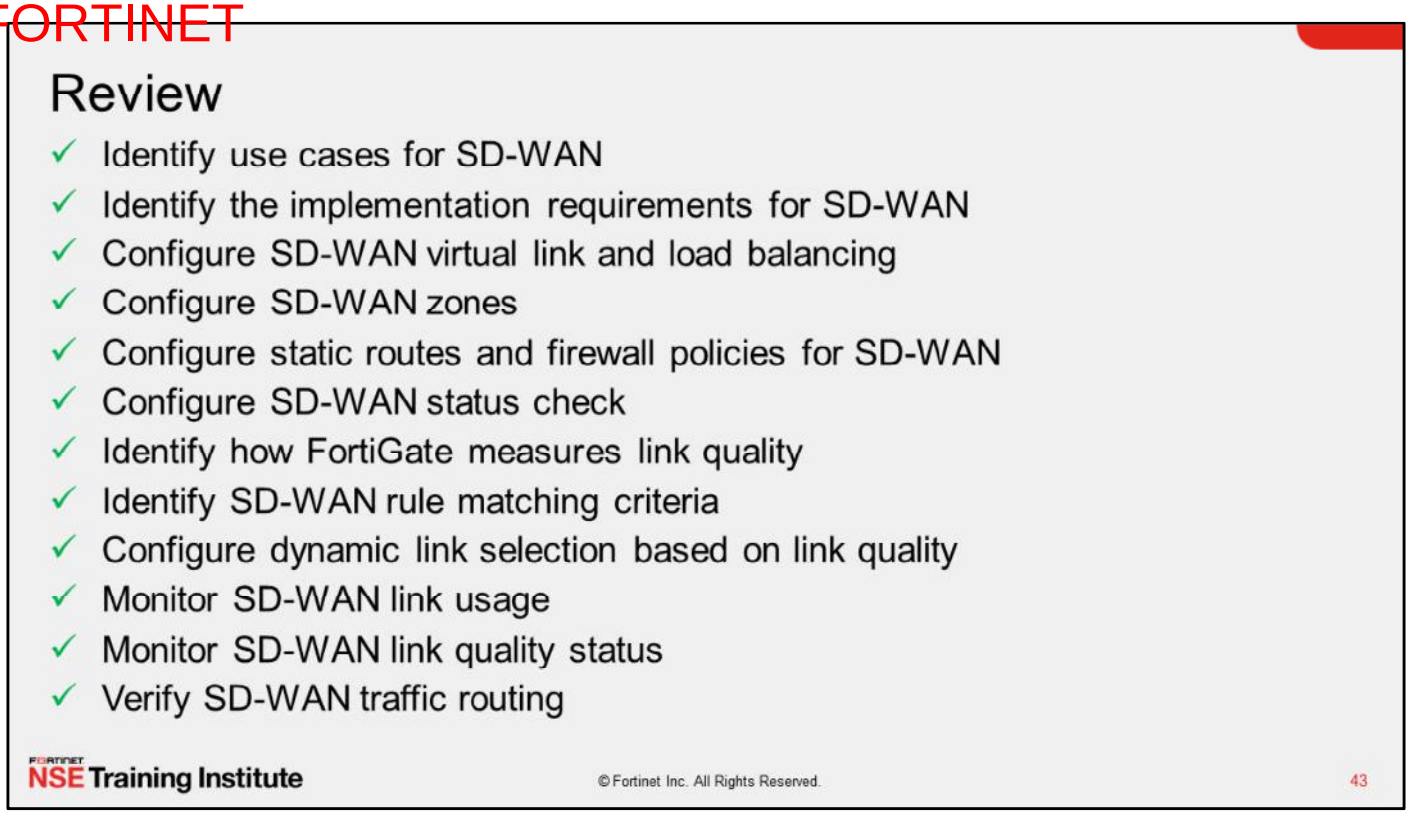

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure, maintain, and diagnose your FortiGate SD-WAN solution.

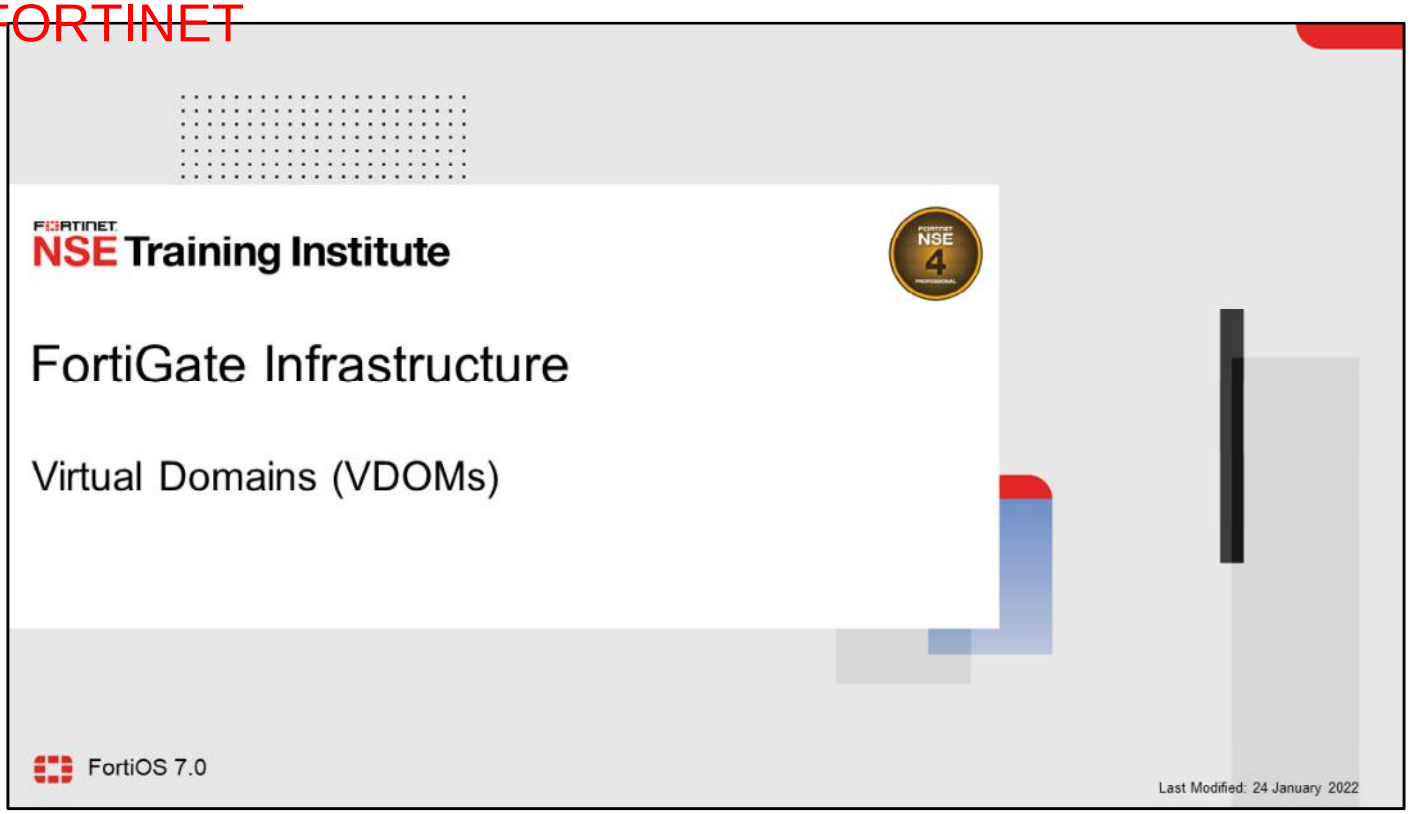

In this lesson, you will learn how to configure VDOMs, and examine examples of common use.

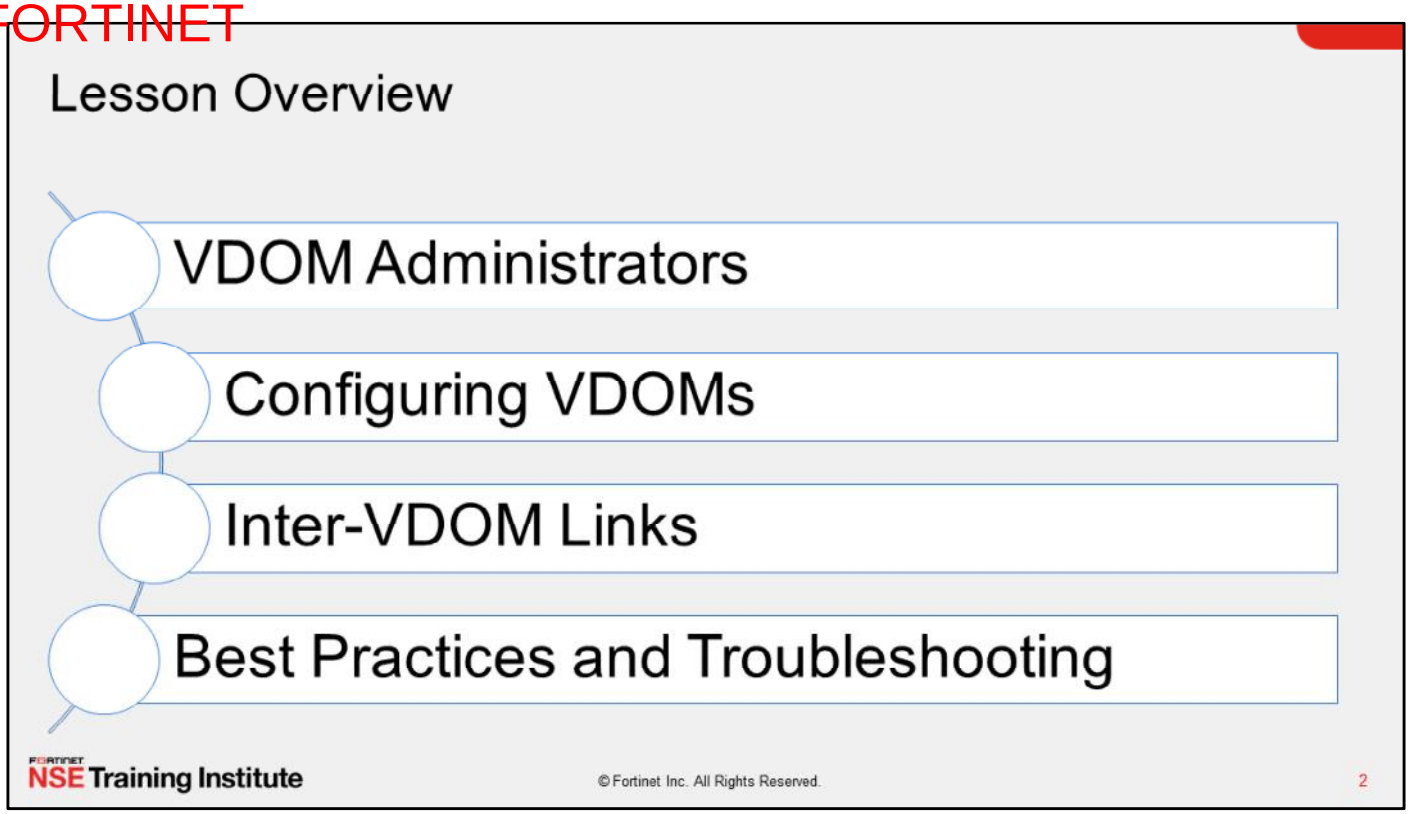

In this lesson, you will learn about the topics shown on this slide.

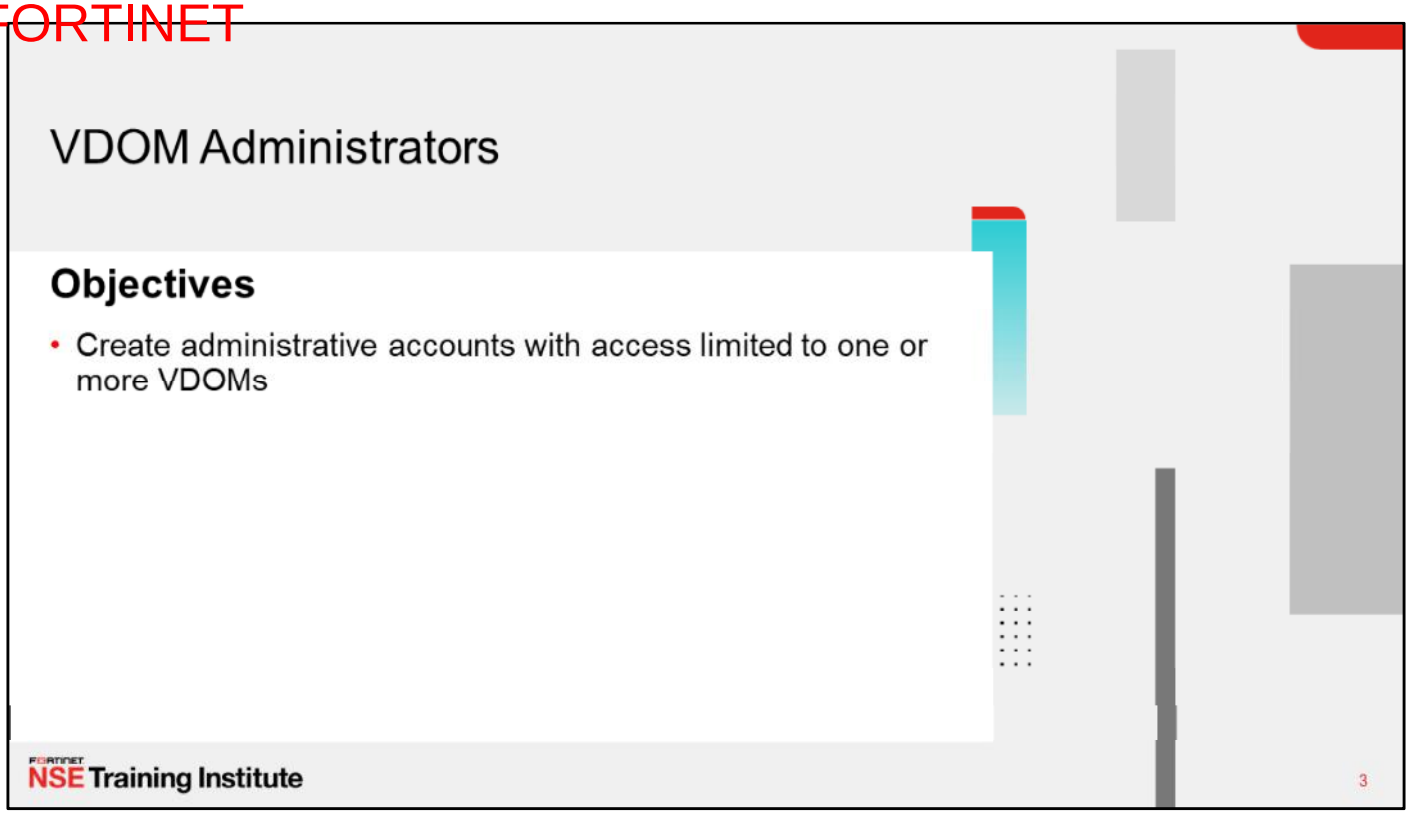

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in 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

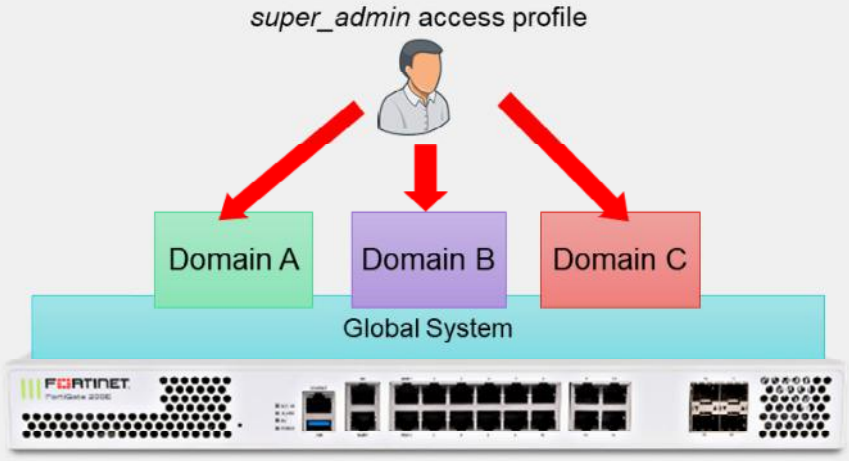

NSE Training Institute

4

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.

Per-VDOM Administration

- Other administrators can access only their *assigned* VDOMs
  - Cannot access the global settings

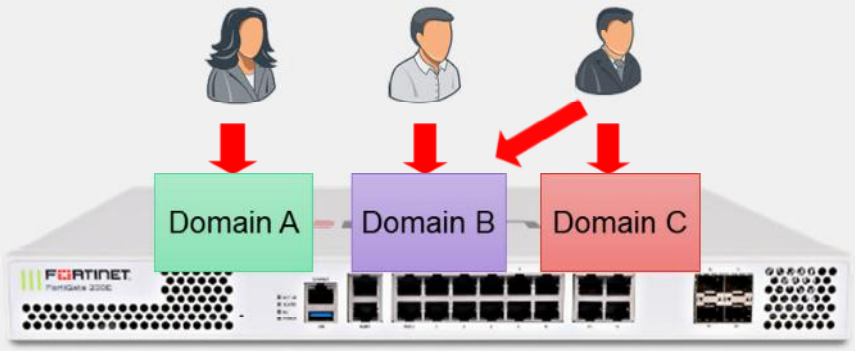

NSE Training Institute

5

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.

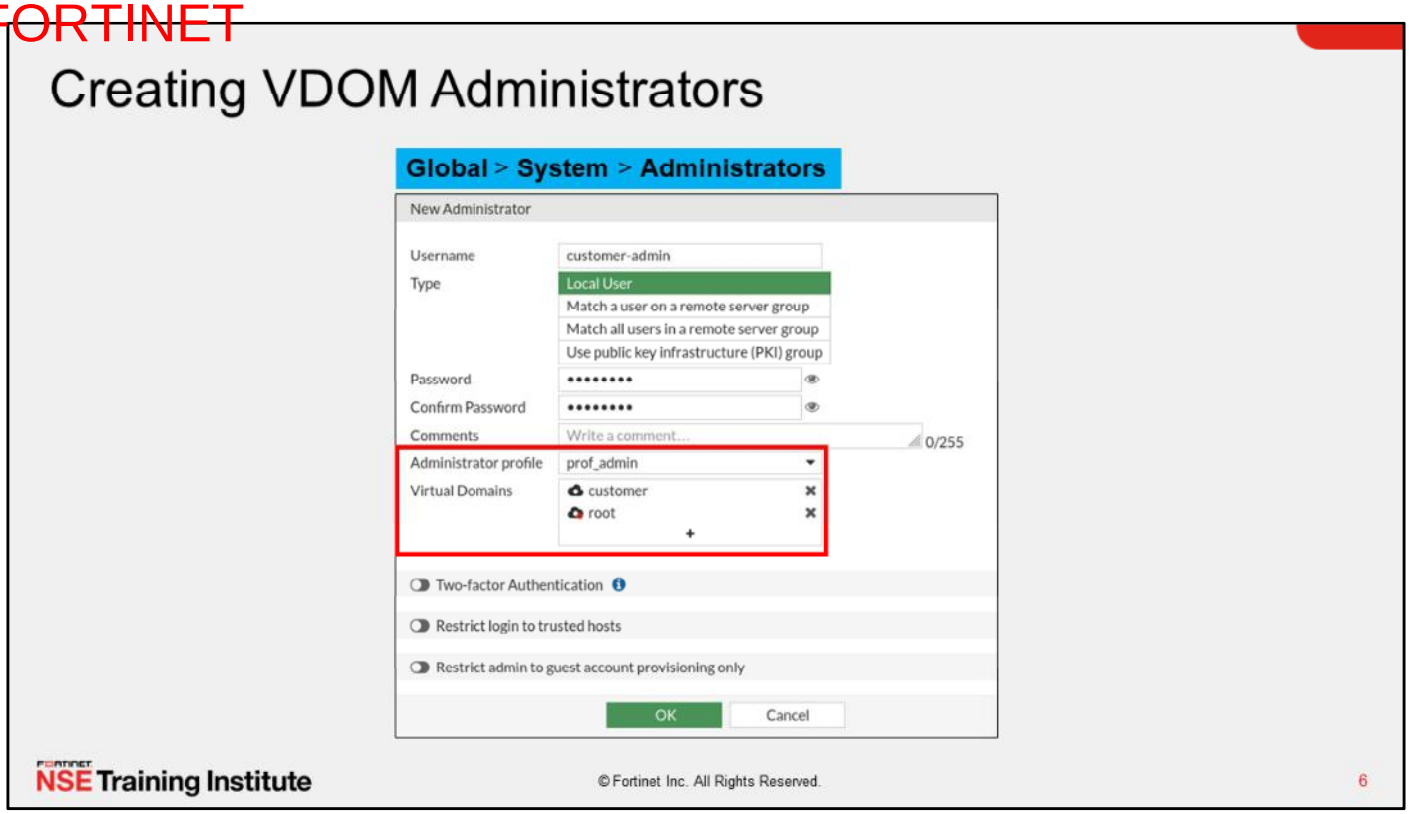

To create new administrator accounts and assign them to a VDOM, click **Global** > **System** > **Administrators**.

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.

NSE Training Institute

7

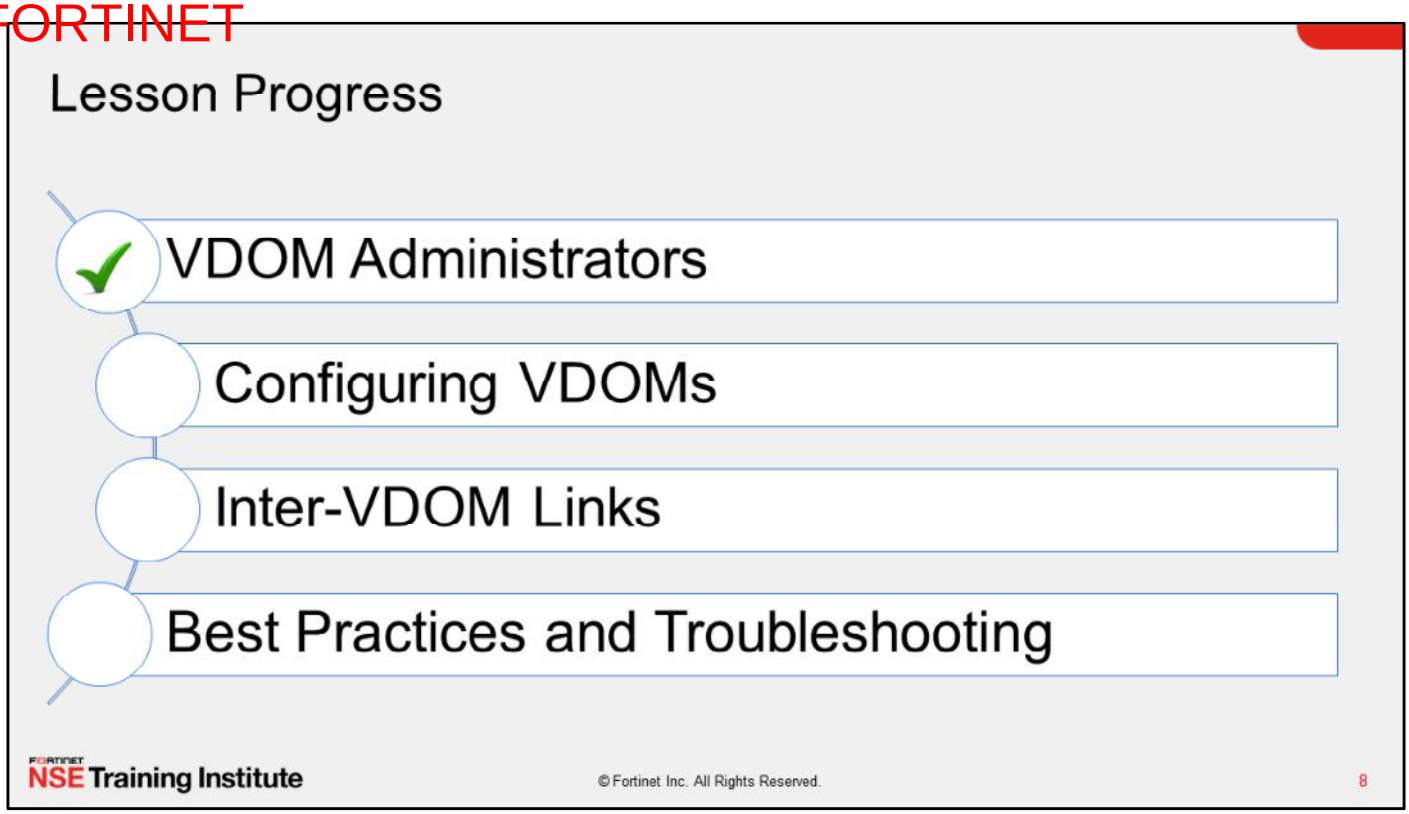

Good job! You now understand VDOM administrators.

Now, you'll learn how to configure VDOMs.

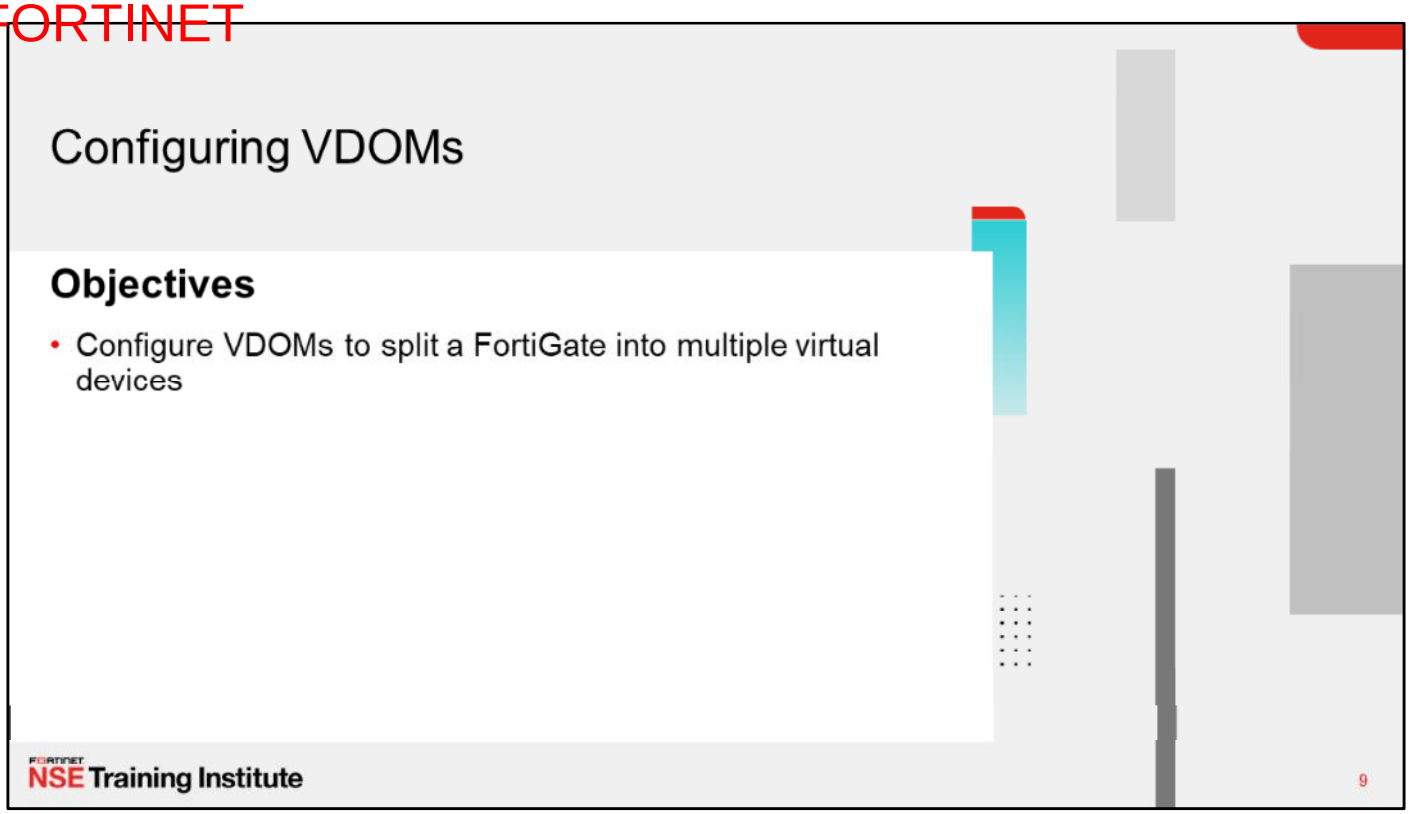

After completing this section, you will be able to achieve the objective shown on this slide.

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

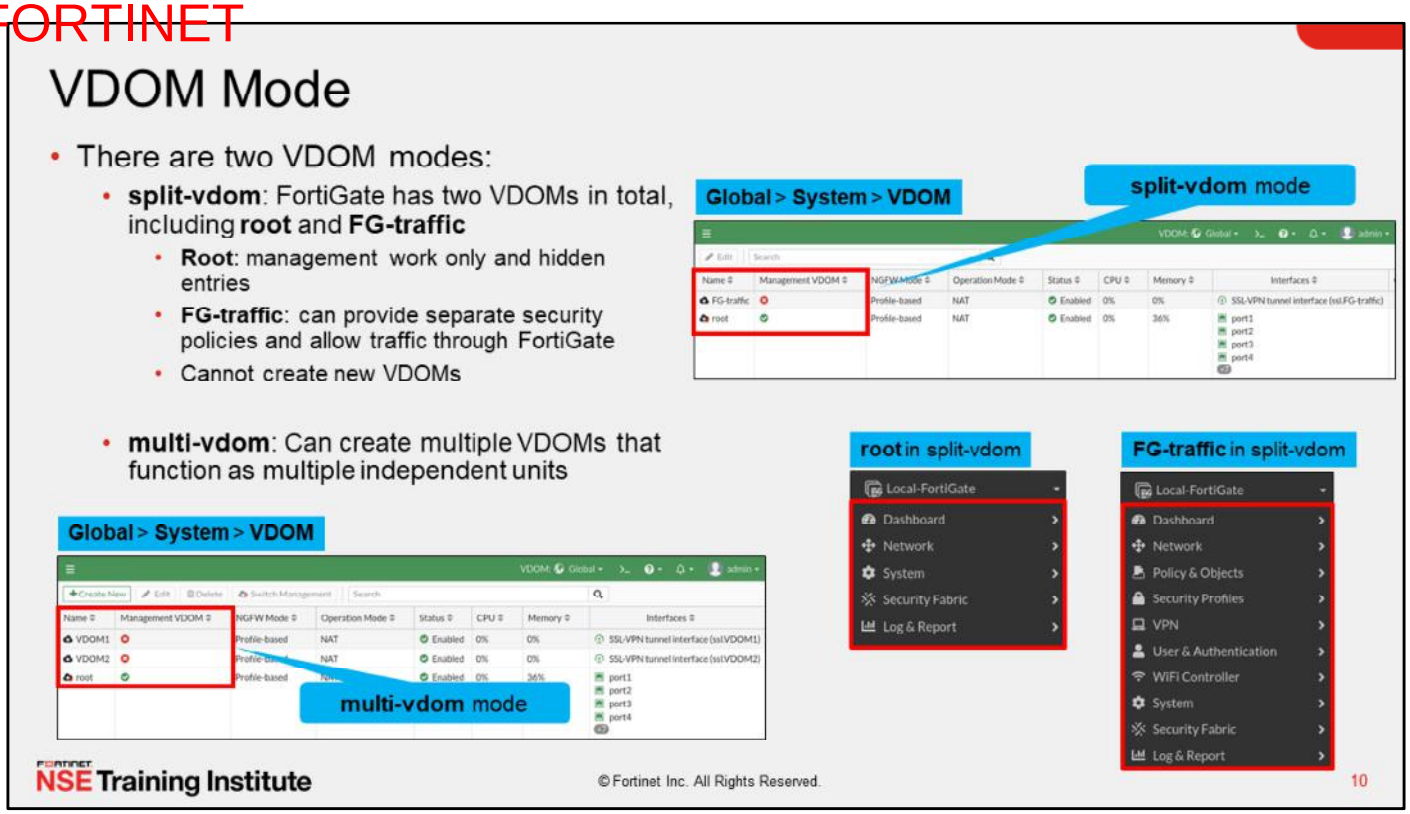

There are two VDOM modes: split-vdom and multi-vdom. In split-vdom, FortiGate has two VDOMs in total, including **root** and **FG-traffic** vdoms. You cannot add VDOMs in split-vdom mode.
1. **split-vdom** mode:

a) The **root** VDOM in split-vdom mode is the management VDOM and does only management work. The following navigation bar entries and pages are hidden in the **root** vdom:

- All **Policy & Object** entries
- **User & Device**, **Security Profiles**
- Traffic-related **FortiView** entries
- **VPN** entries
- **System** > **Fabric Connectors**, **Reputation**, **Feature Visibility**, **Object Tags** entries
- **Wan-Opt** entries
- Most route entries
- Most log event entries
- **Monitor** entries

b) The **FG-traffic** VDOM can provide separate security policies and allow traffic through FortiGate.

2. 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.

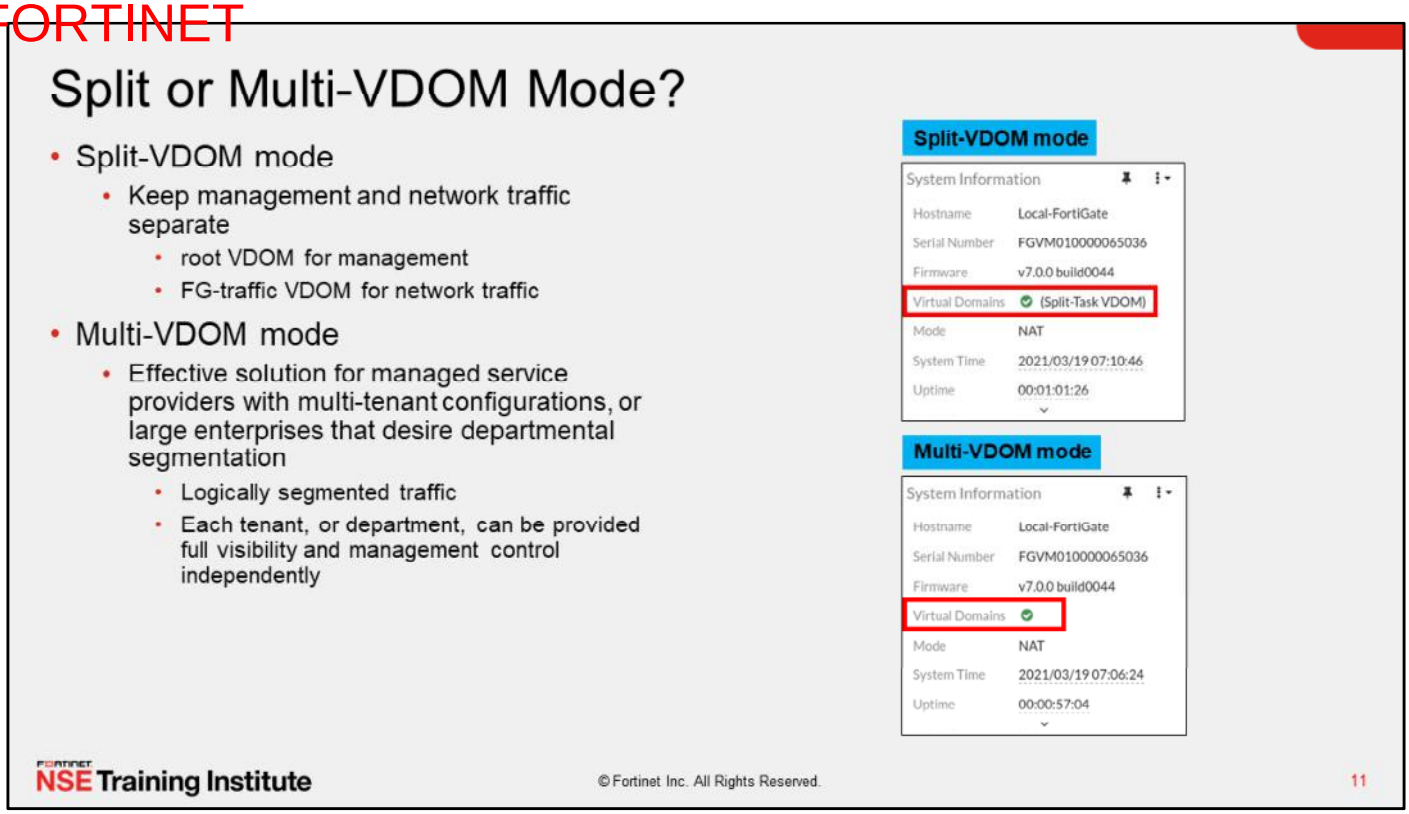

When would you select one VDOM mode over another? The answer depends on the environment, and the desired traffic segmentation methodology.

You can use split-VDOM mode when you do not want management traffic and network traffic to exist in the same VDOM. Use the root VDOM to isolate management traffic, while reserving the FG-traffic VDOM for network traffic.

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

Enabling VDOMs

- From the GUI:
  - Only available on specific models. If the option does not exist, use the CLI command

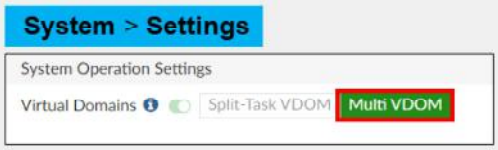

- From the CLI:

```
#config system global
    set vdom-mode no-vdom/split-vdom/multi-vdom
 end
```

NSE Training Institute

12

To enable VDOMs using the GUI, click **System** > **Settings** > **Virtual Domains**. Then, click **toggle vdom mode** in the **Virtual Domains** section. Note that on FortiGate-60 series and earlier models, you must enable VDOMs on the CLI only, by using following command:

```
config system global
    set vdom-mode   no-vdom/split-vdom/multi-vdom
end
```

Enabling VDOMs won't reboot your FortiGate device, but will log out all active administrator sessions. Enabling VDOMs restructures both the GUI and CLI, which you will see when you log in again. This also does not affect any traffic passing through FortiGate.

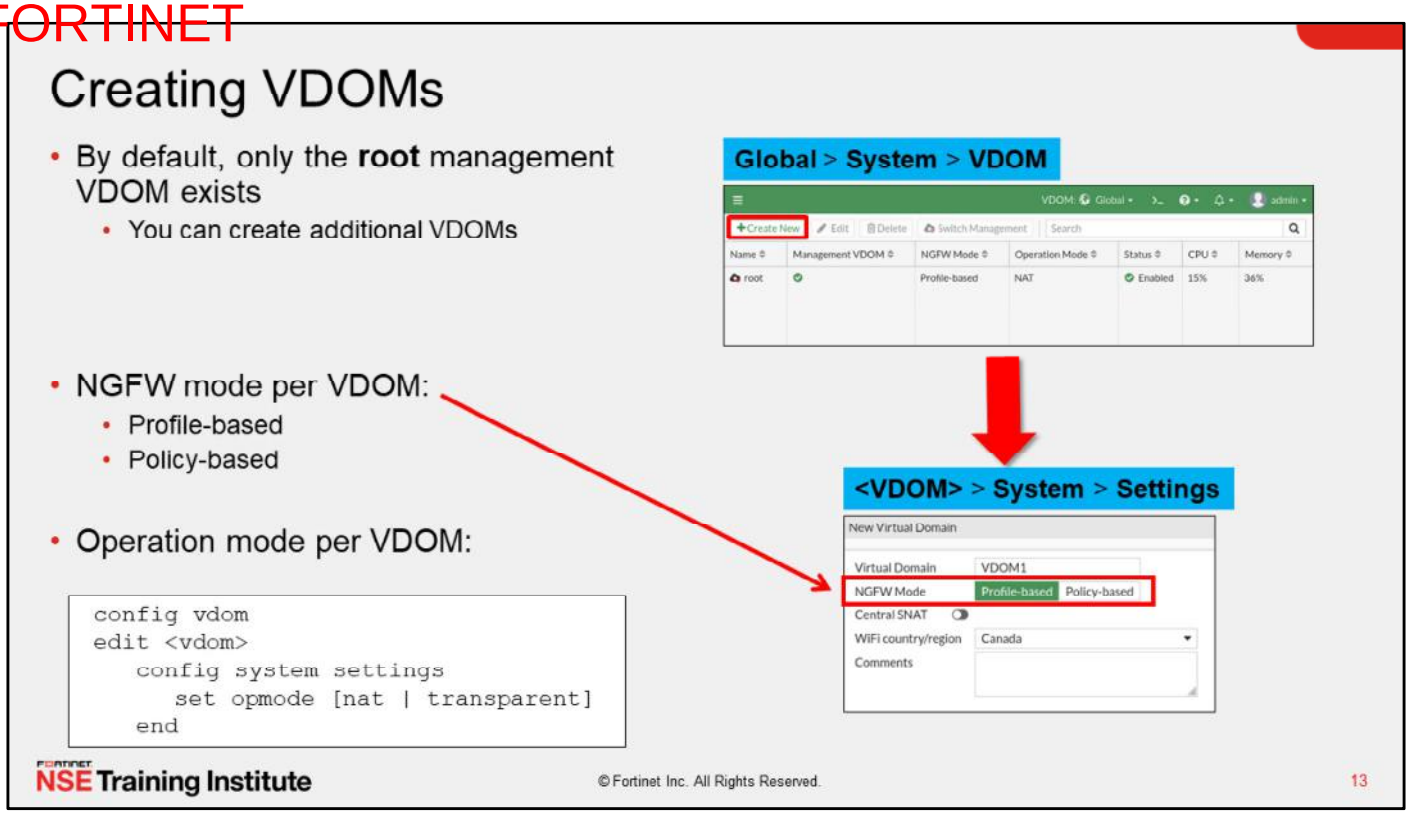

After enabling VDOMs in multi-vdom mode, by default, only one VDOM exists: the root VDOM. It's the default management VDOM.

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.

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

NSE Training Institute

14

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

- After you create VDOMs, you can assign interfaces to each VDOM

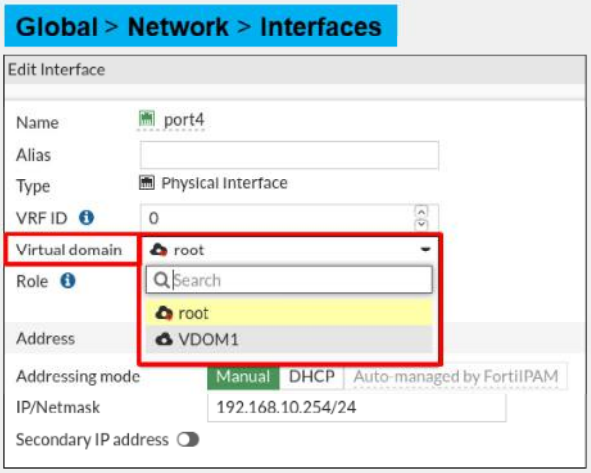

NSE Training Institute

15

After adding the additional VDOMs, you can specify which interfaces belong to each VDOM. Each interface (physical or VLAN) can belong to only one VDOM.

You can move interfaces from one VDOM to another, provided they have no references associated with them, for example, firewall policies.

Global Settings

- Affect all configured VDOMs:
  - Hostname
  - HA settings
  - FortiGuard settings
  - System time
  - Administrative accounts

NSE Training Institute

16

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.

Settings for Each VDOM

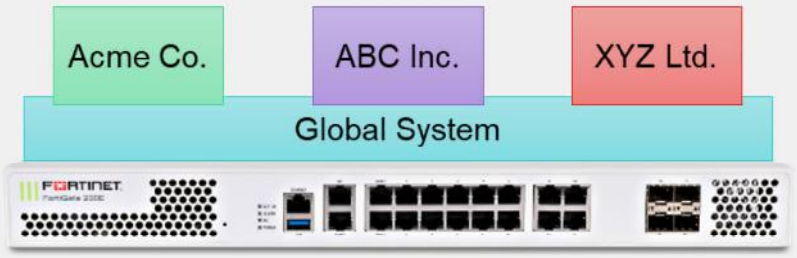

- Configured separately, in each VDOM:
  - Operating mode (transparent, NAT/route)
  - NGFW mode (profile-based, policy-based)
  - Routes and network interfaces
  - Firewall policies
  - Security profiles

NSE Training Institute

17

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

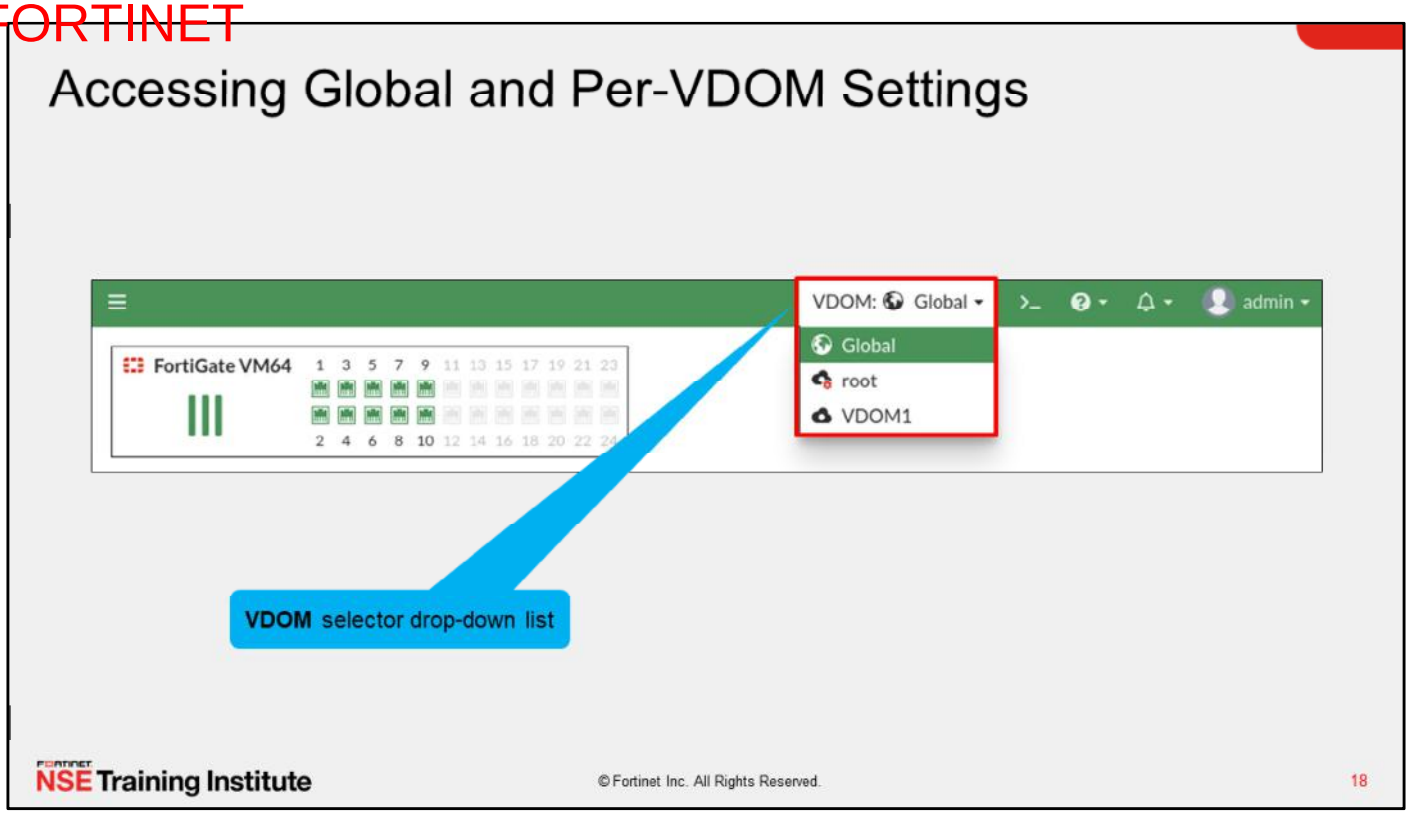

If you log in as most administrator accounts, you enter your VDOM automatically.

But, if you are logged in as the account named **admin**, you aren't assigned to any VDOM. As such, you have access to all VDOMs.

To enter a VDOM on the GUI, select the VDOM from the drop-down list at the top of the page.

Inside each VDOM, the submenu should be familiar; it is essentially the same navigation menu that you had before you enabled VDOMs. However, the global settings are moved to the **Global** menu.

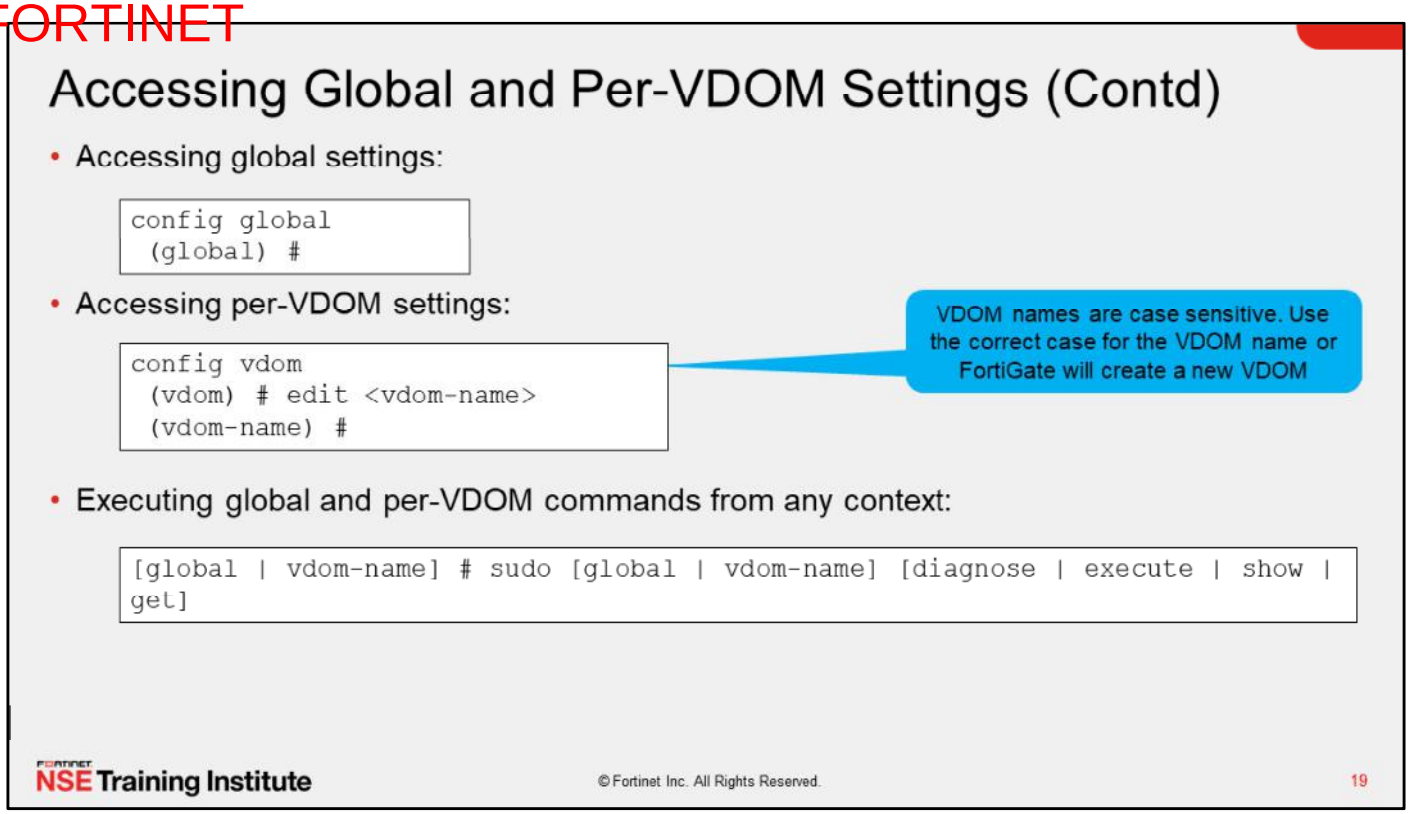

To access the global configuration settings on the CLI, you must enter `config global` to enter into the global context. After that, you can run global commands and change global configuration settings.

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 will create 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.

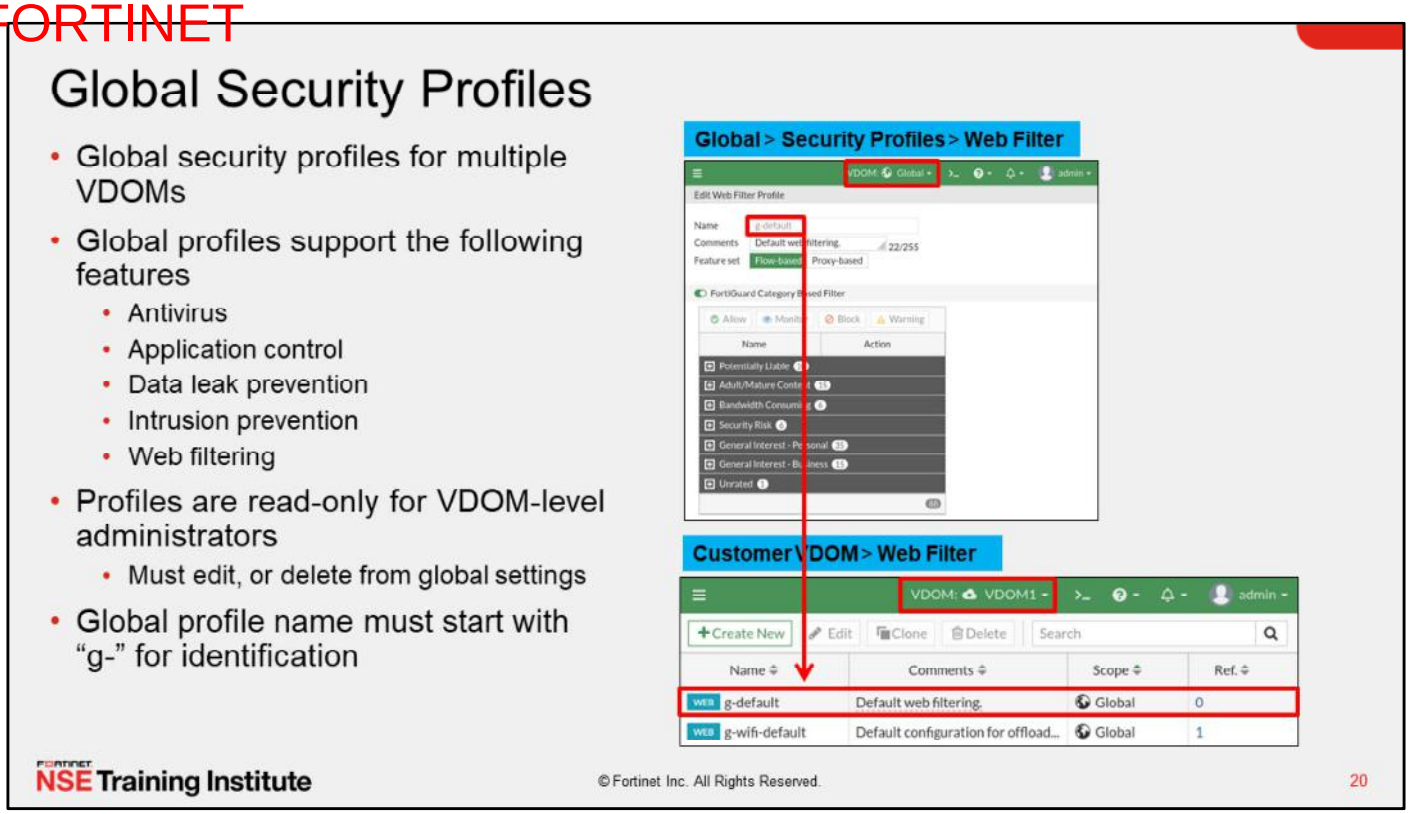

You can configure security profiles globally for use by multiple VDOMs, to avoid creating identical profiles for each VDOM separately. Global profiles are available for the following security features:

- Antivirus
- Application control
- Data leak prevention
- 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.

VDOMs in the Security Fabric

- Security fabric support
  - Physical topology displays downstream FortiGate devices individually
  - Logical topology displays interfaces grouped by VDOM
- Fabric connectivity through management VDOM

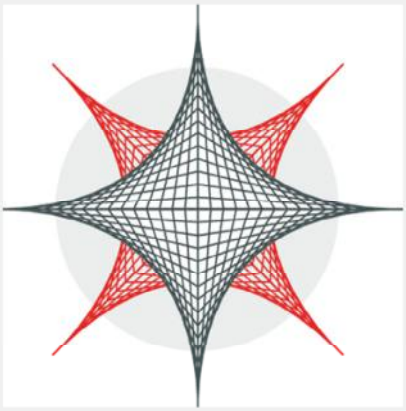

NSE Training Institute

21

You can include FortiGate devices configured with VDOMs in the Security Fabric. The physical Security Fabric topology represents each downstream FortiGate as a single FortiGate device. In the Security Fabric logical topology, each FortiGate displays ports grouped by VDOM. Downstream FortiGate devices must join the Security Fabric through the management VDOM of their upstream FortiGate device.

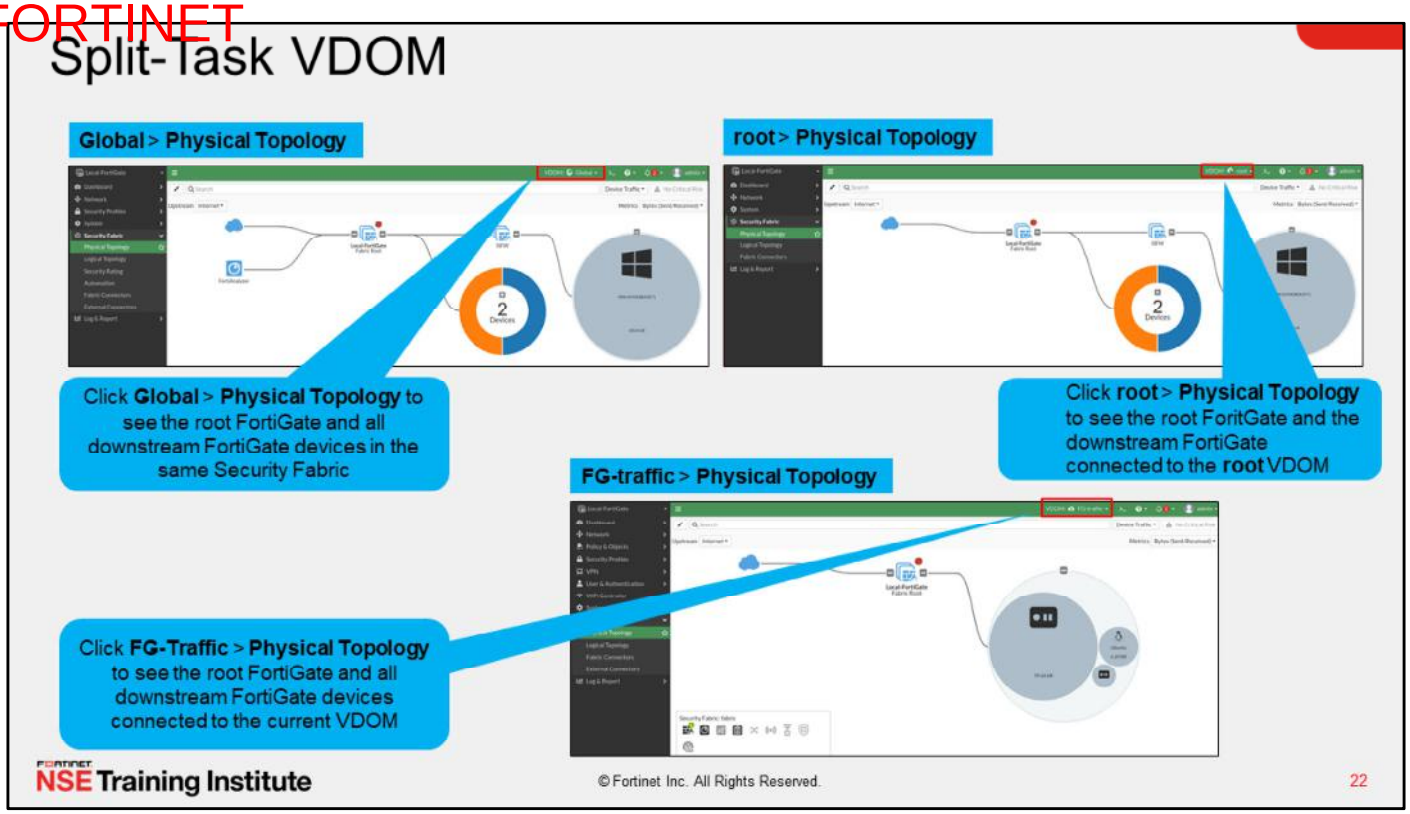

FortiGate Security Fabric connection settings are available on the **Security Fabric** > **Fabric Connectors** page. If the upstream FortiGate has VDOM mode enabled, it can allow downstream FortiGate devices to join the Security Fabric through one of the existing VDOMs. If the downstream FortiGate device has VDOM mode enabled, it can connect to the upstream FortiGate through the downstream FortiGate interface.

You can click **Global** > **Physical Topology** to see the root FortiGate device and *all* downstream FortiGate devices that are in the same Security Fabric as the root FortiGate device. You can click **root** > **Physical Topology** or **FG-Traffic** > **Physical Topology** to see the root FortiGate device and *only* the downstream FortiGate devices that are connected to the currently selected VDOM on the root FortiGate device.

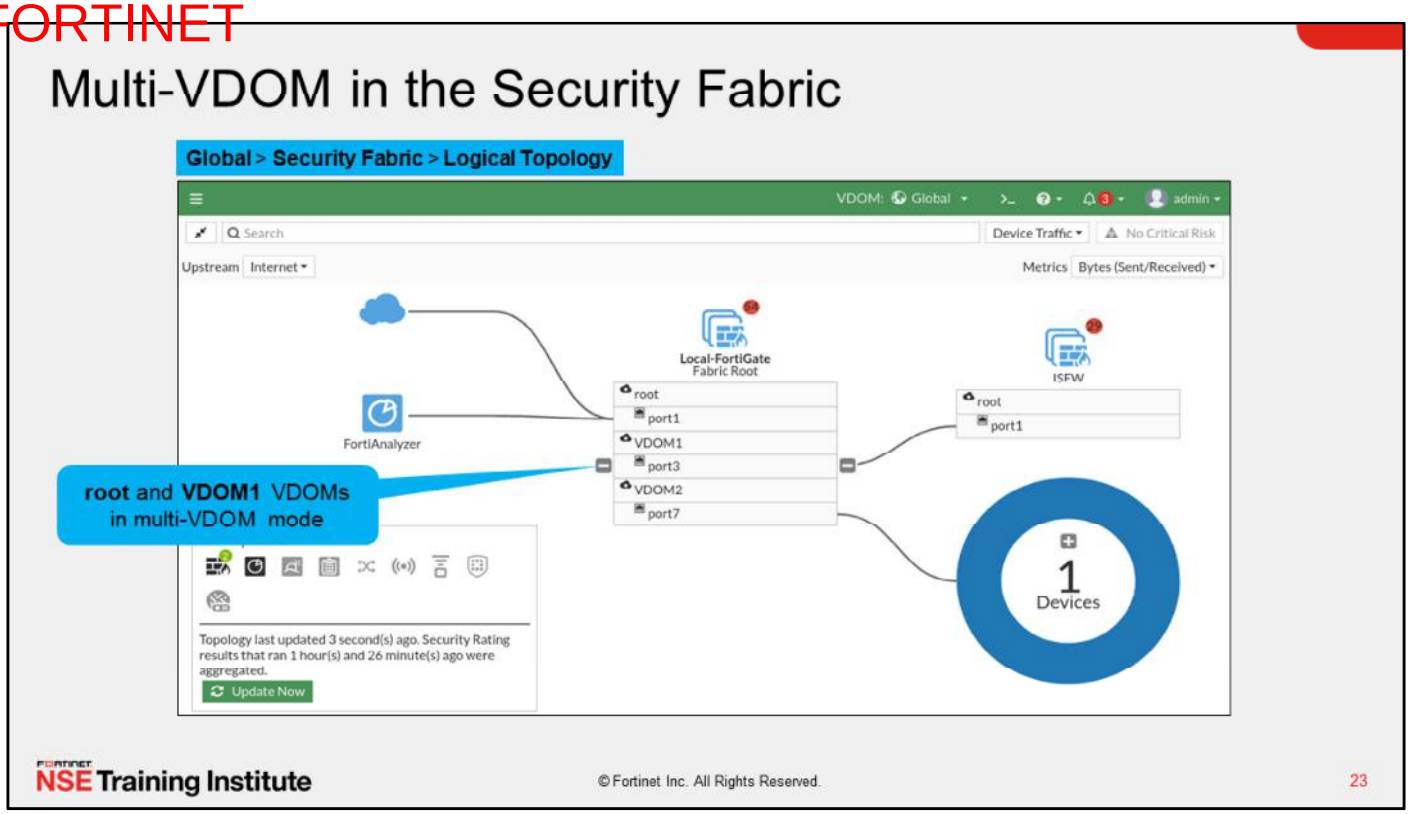

When you configure FortiGate devices in multi-vdom mode and add them to the security fabric, each VDOM with its assigned ports is displayed when one or more devices are detected. *Only* the ports with discovered and connected devices appear in the Security Fabric view and, because of this, you must enable **Device Detection** on ports you want to have displayed in the **Security Fabric**. VDOMs without ports with connected devices are not displayed. All VDOMs configured must be part of a single **Security Fabric**. In the example shown on this slide, the Local-FortiGate is configured in multi-VDOM mode, and has three VDOMs (root, VDOM1, and VDOM2), each with ports that have connected devices.

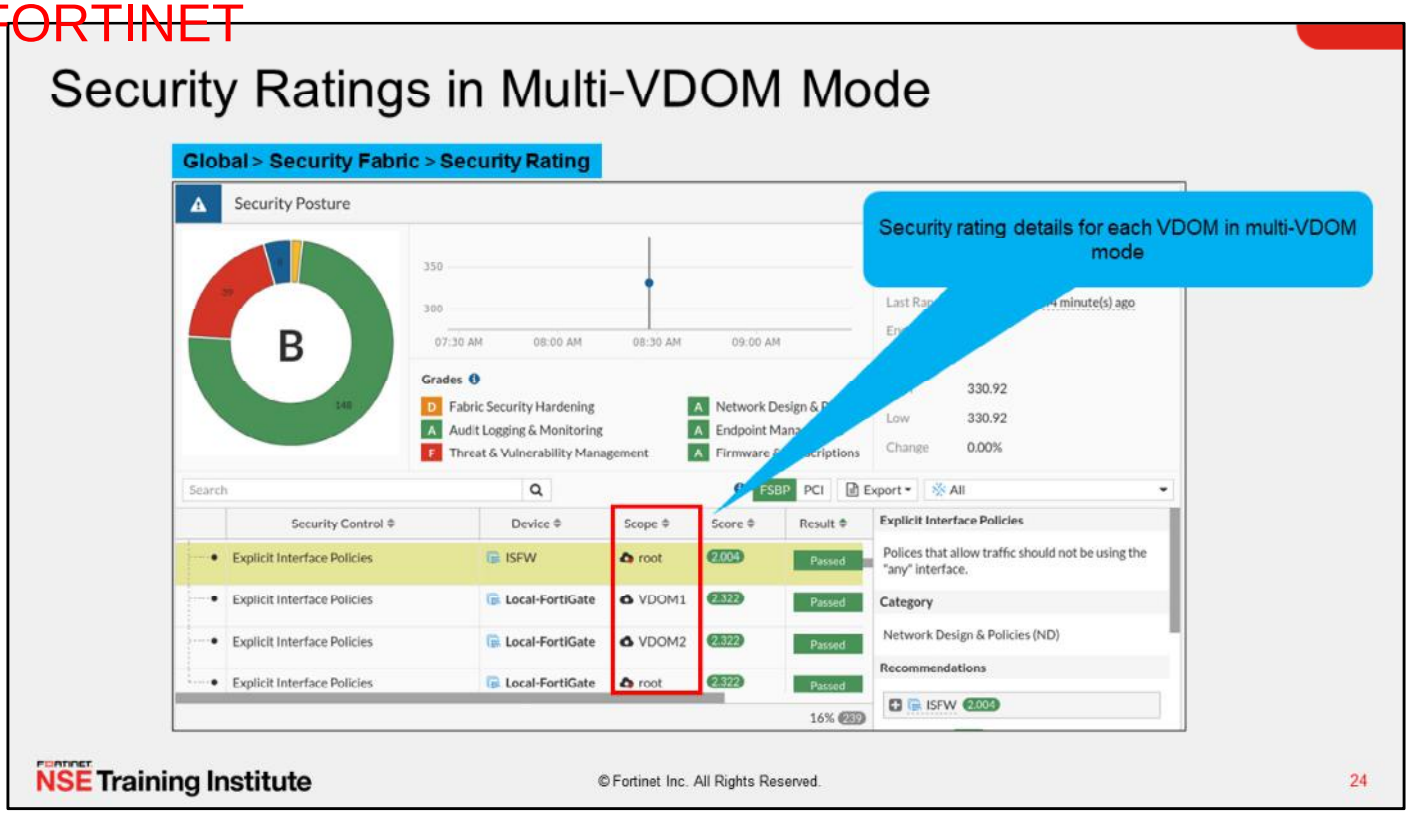

You can view security ratings at the global level. Within each of the three scorecards, each security control entry displays the associated VDOM information (if any) in the **Scope** column. This allows for a more granular breakdown of the Security Fabric deployment analysis.

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

25

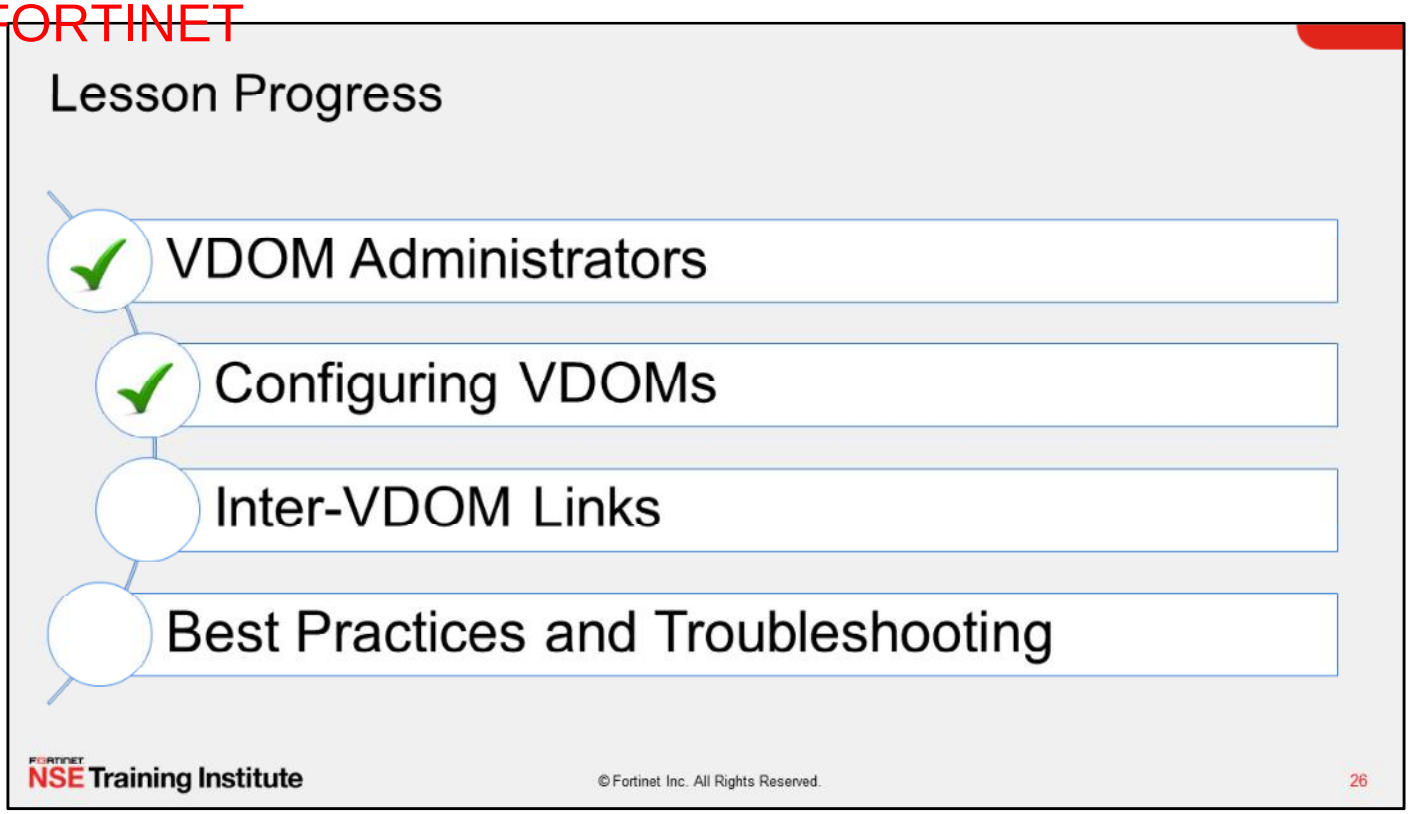

Good job! You now understand how to configure VDOMs.

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

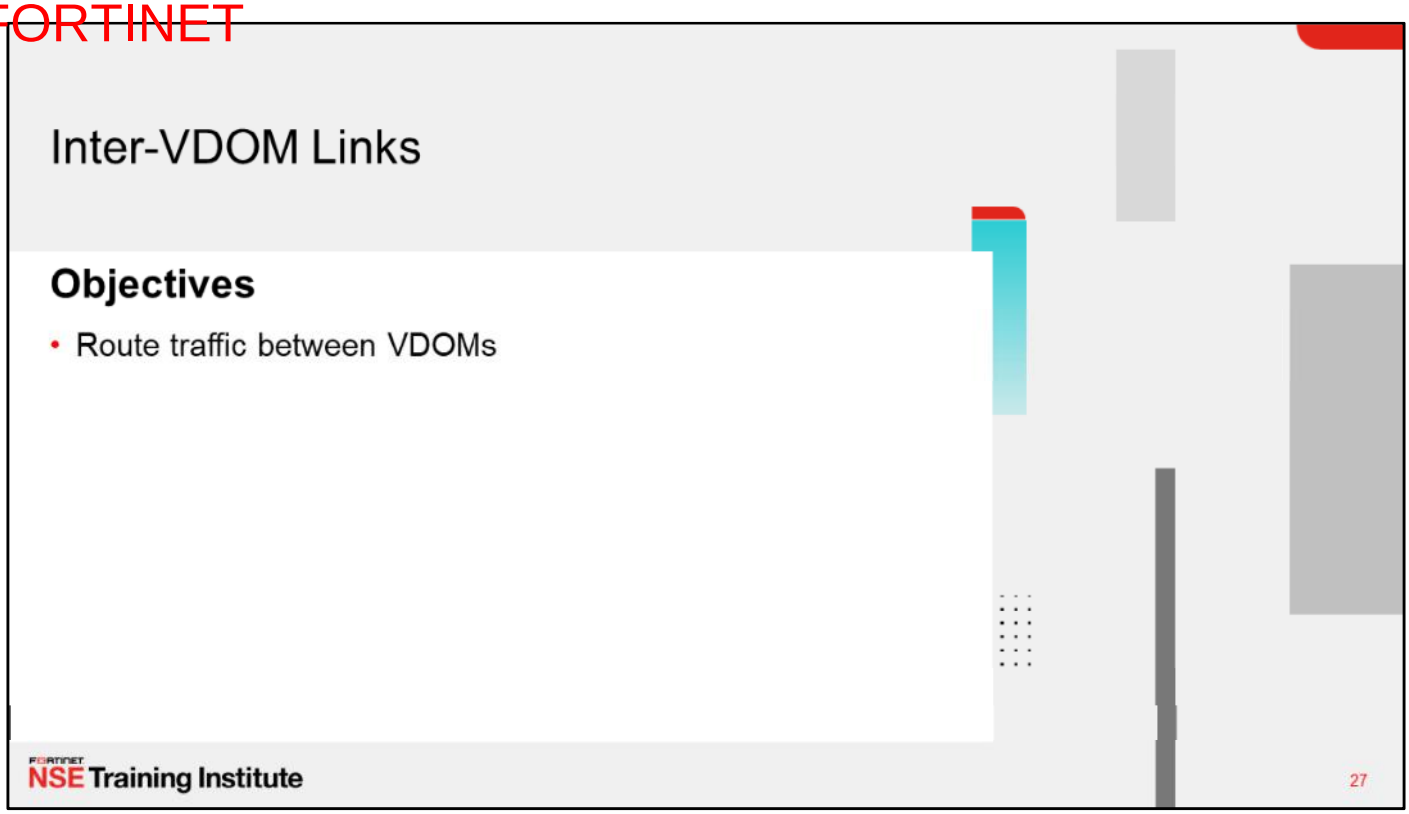

After completing this section, you should be able to achieve the objective shown on this slide.

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

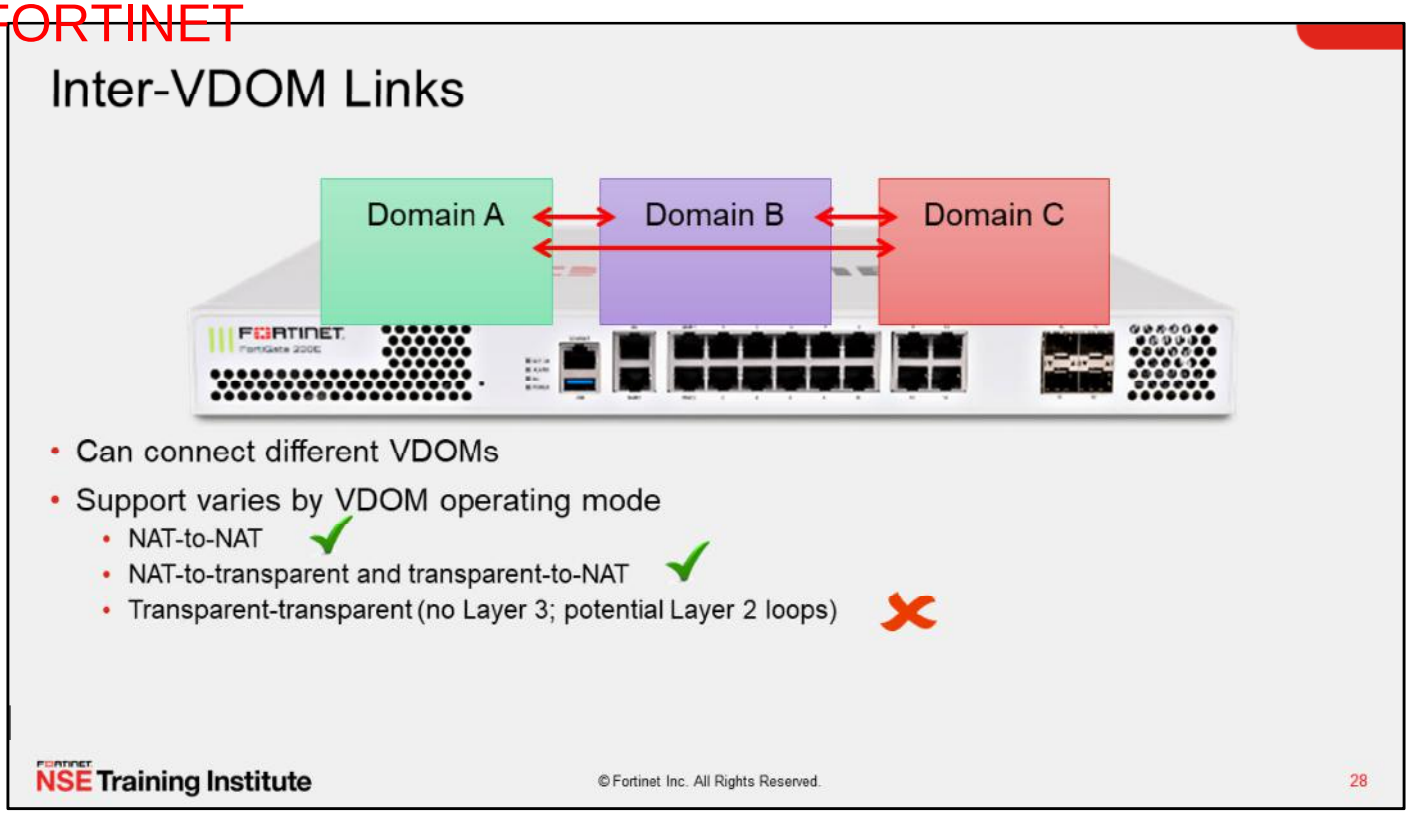

To review, each VDOM behaves like it is on a separate FortiGate device. With separate FortiGate devices, you would normally connect a network cable and configure routing and policies between them. But VDOMs are on the same FortiGate. So, how should you route traffic between them?

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 similar to 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.

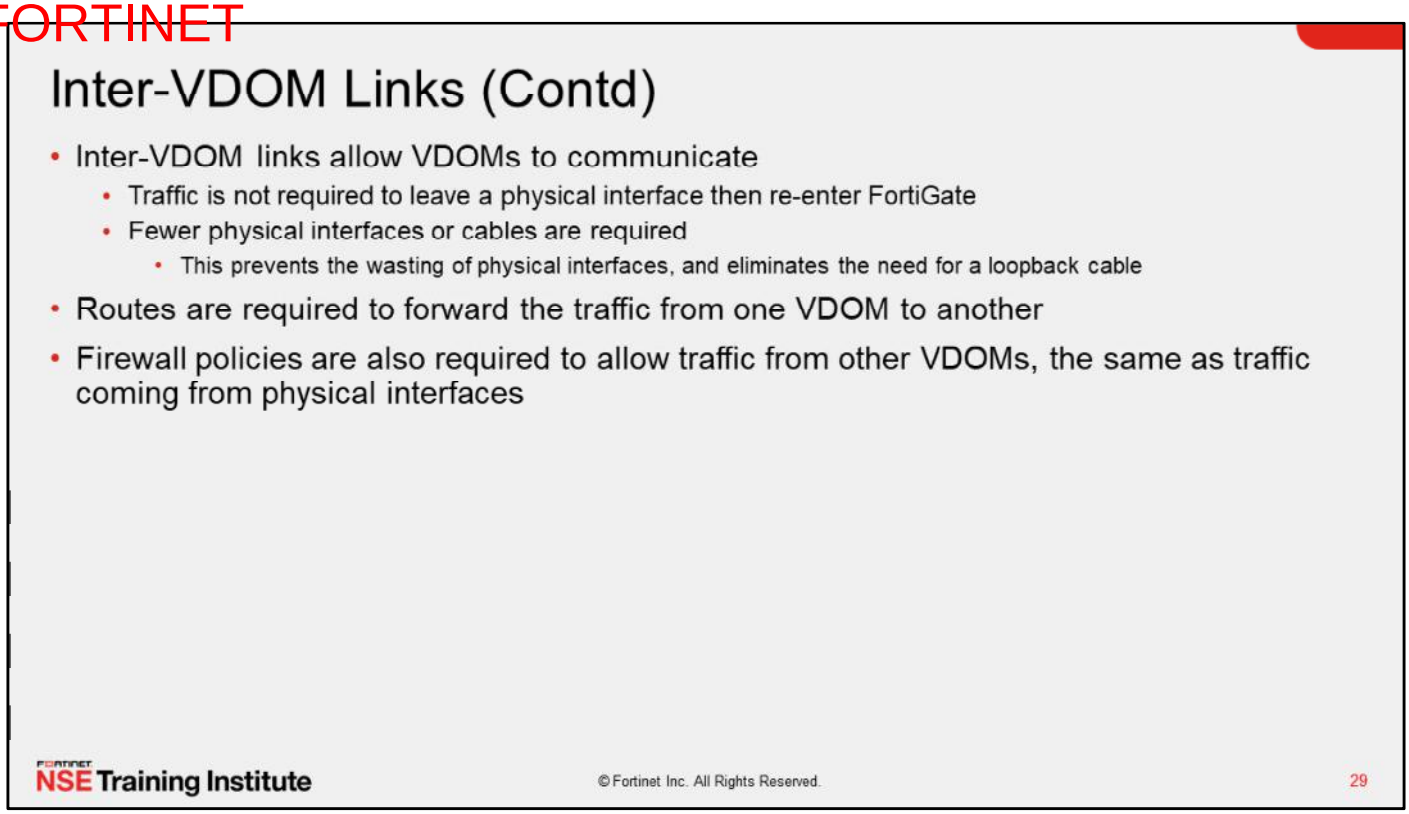

When creating inter-VDOM links, you must create the virtual interfaces. You must also create the appropriate firewall policies in each VDOM, just as you would if the traffic were arriving on a network cable, otherwise, FortiGate will block it.

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

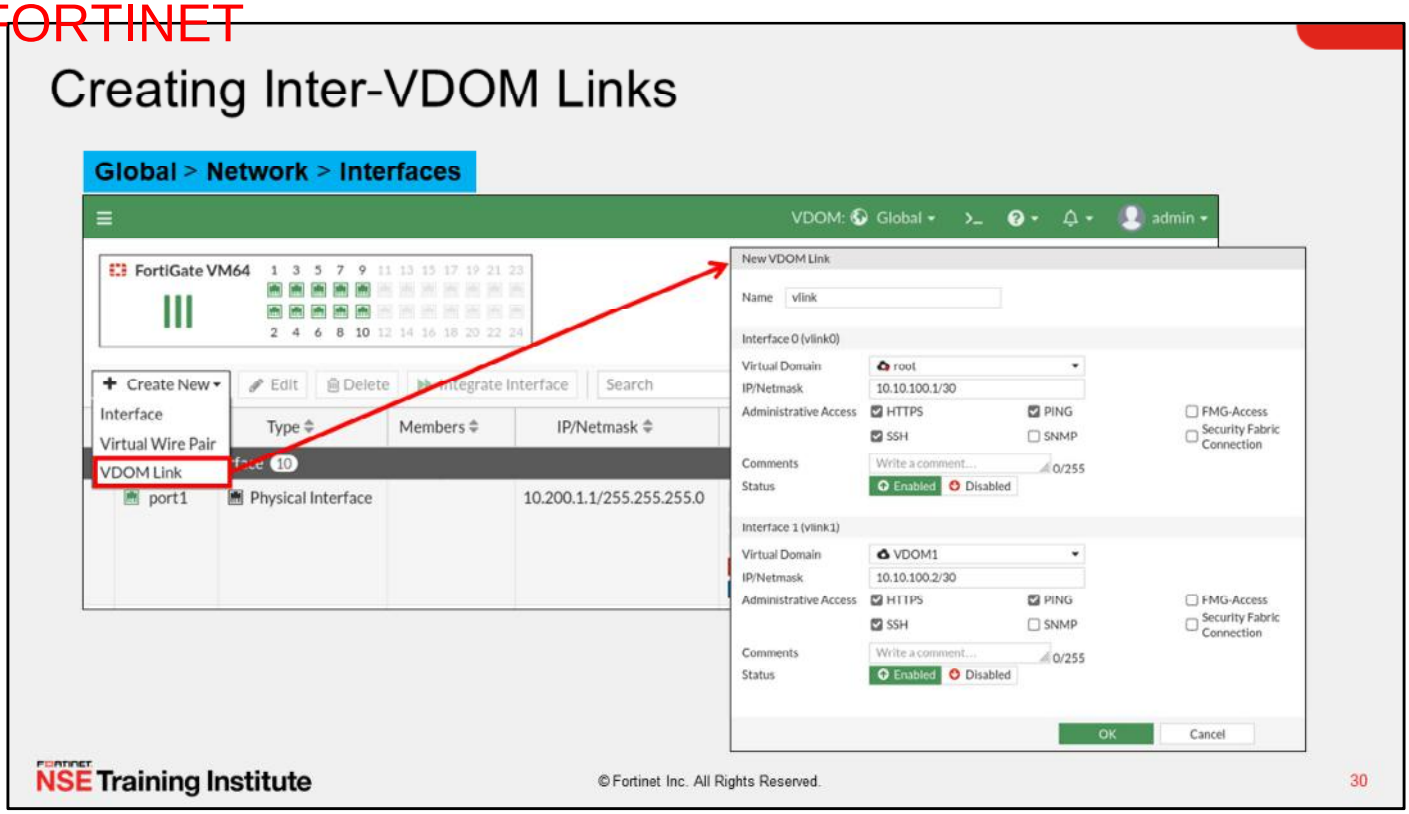

On the GUI, you create a network interface in the **Global** settings. To create the virtual interface, click **Create New**, and then select **VDOM Link**.

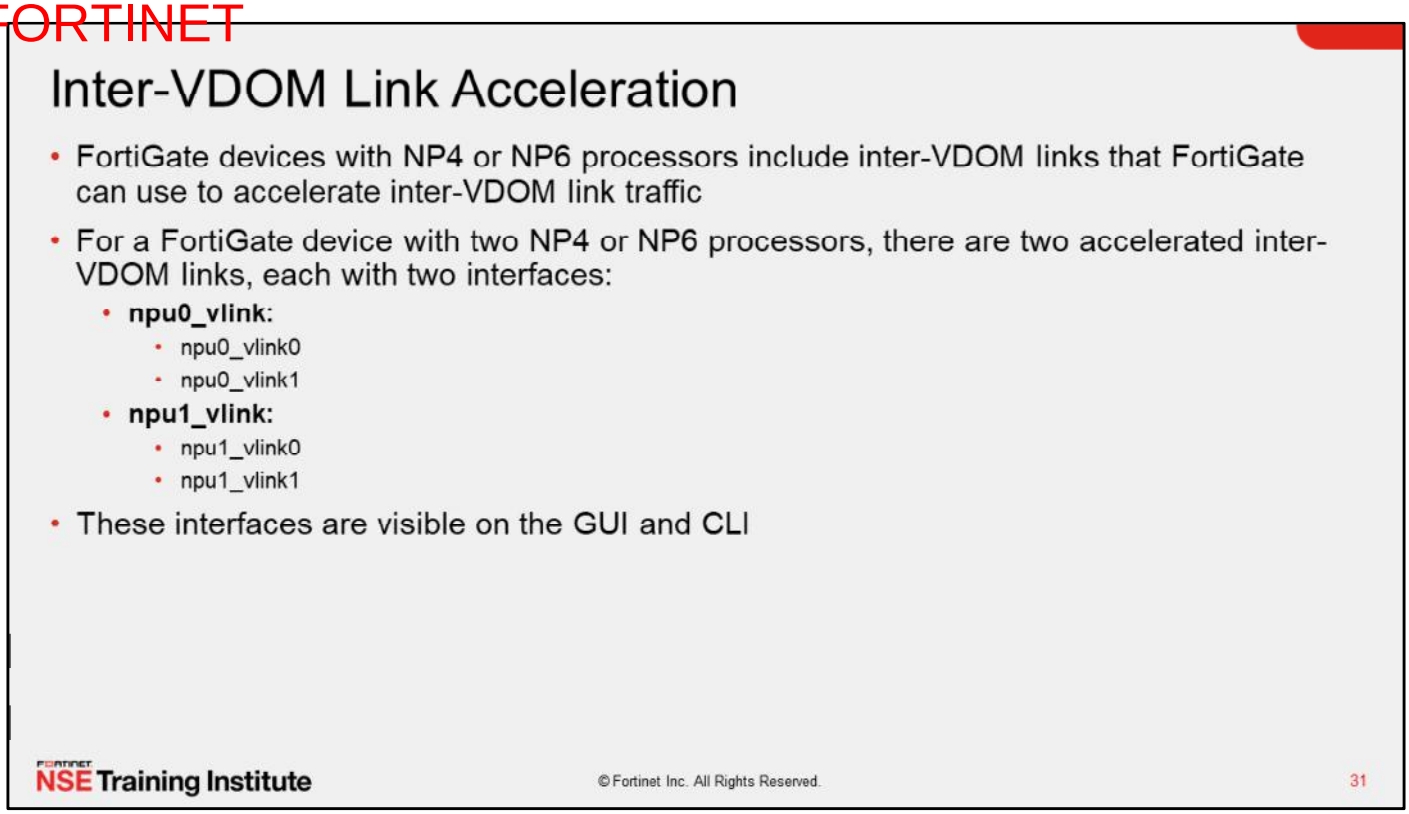

FortiGate devices with NP4 or NP6 processors include inter-VDOM links that FortiGate can use to accelerate inter-VDOM link traffic. For a FortiGate with two NP4 or NP6 processors, there are two accelerated inter-VDOM links, each with two interfaces:

- **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

NSE Training Institute

32

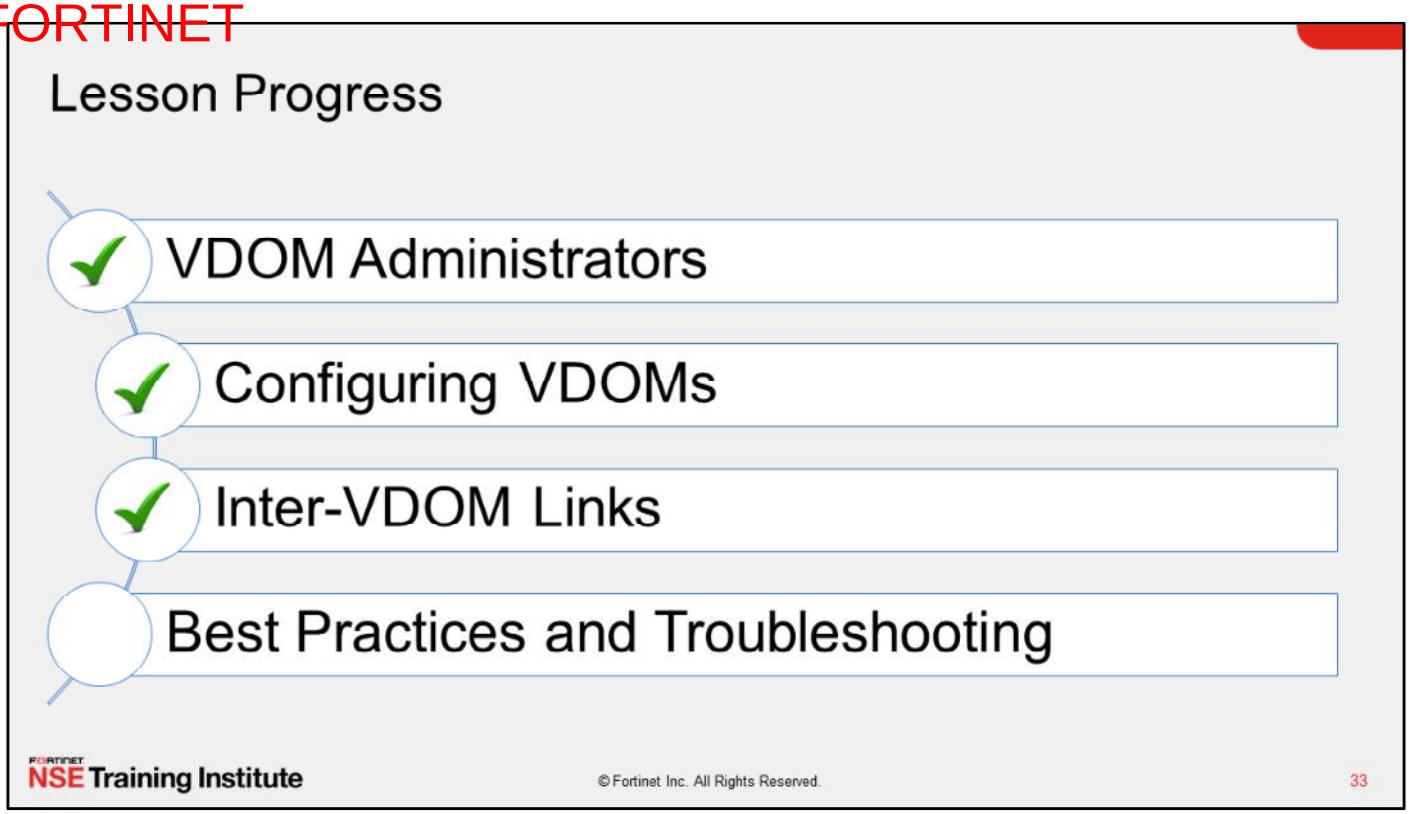

Good job! You now understand inter-VDOM Links.

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

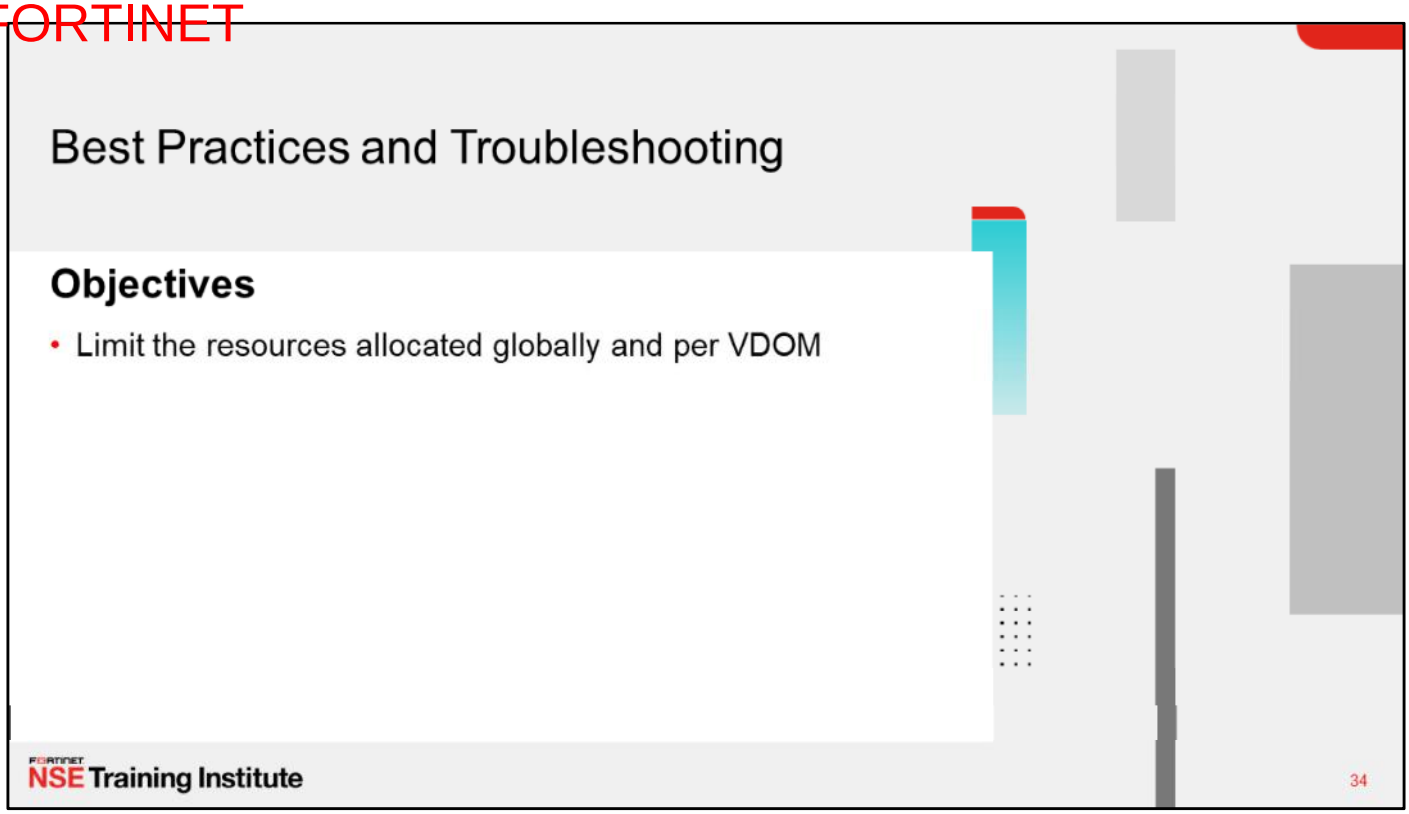

After completing this section, you should be able to achieve the objective shown on this slide.

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

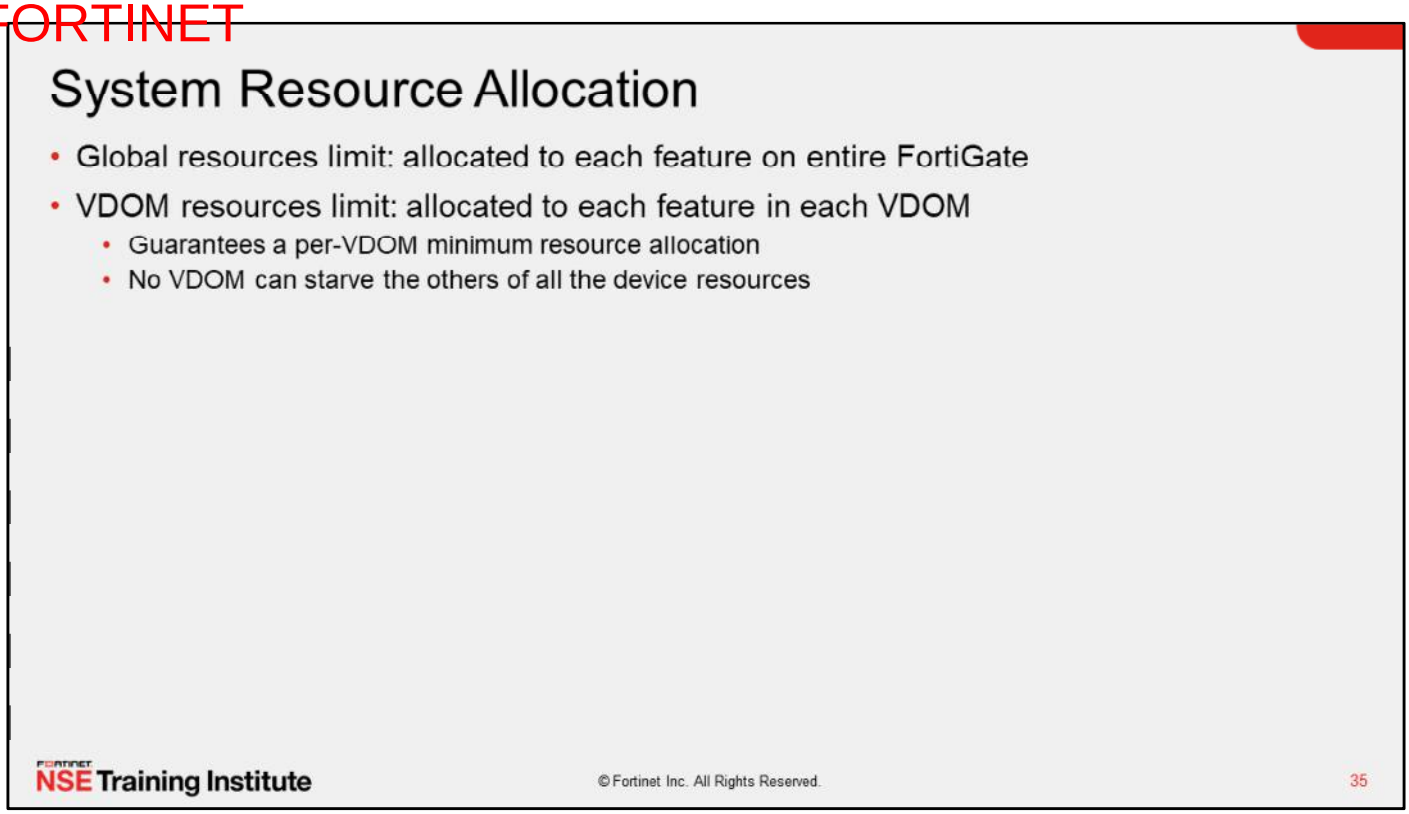

Remember, VDOMs are only a *logical* separation—each VDOM shares physical resources with the others.

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.

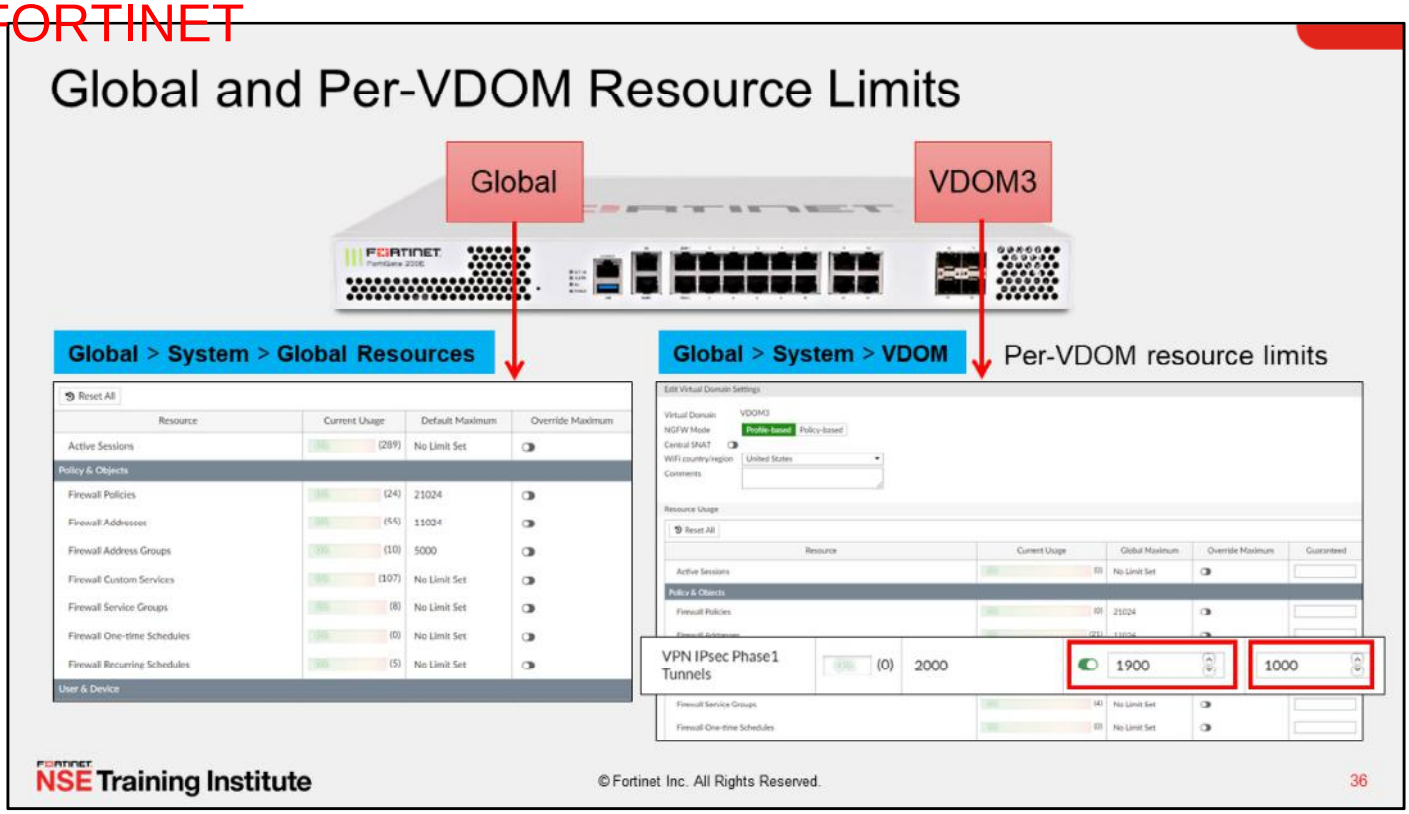

For example, a FortiGate with hardware powerful enough to handle up to 2000 IPsec VPN tunnels and configured with three VDOMs, could be configured as follows to meet specific criteria: VDOM1 and VDOM2 don't use IPsec VPN tunnels often. So, they are allowed to have up to 50 tunnels each. VDOM3, however,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

Monitoring VDOM Resources

- VDOM monitor displays:
  - CPU utilization
  - Memory utilization

Global > System > VDOM

+ Create New | Edit | Delete | Switch Management | Search

Name	Management VDOM	NGFW Mode	Operation Mode	Status	CPU	Memory	Interfaces
customer		Profile-based	NAT	Enabled	0%	7%	port3 SSL-VPN tunnel interface (ssl.customer)
root		Profile-based	NAT	Enabled	0%	38%	port1 port2 port4 port5 +2

NSE Training Institute

37

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

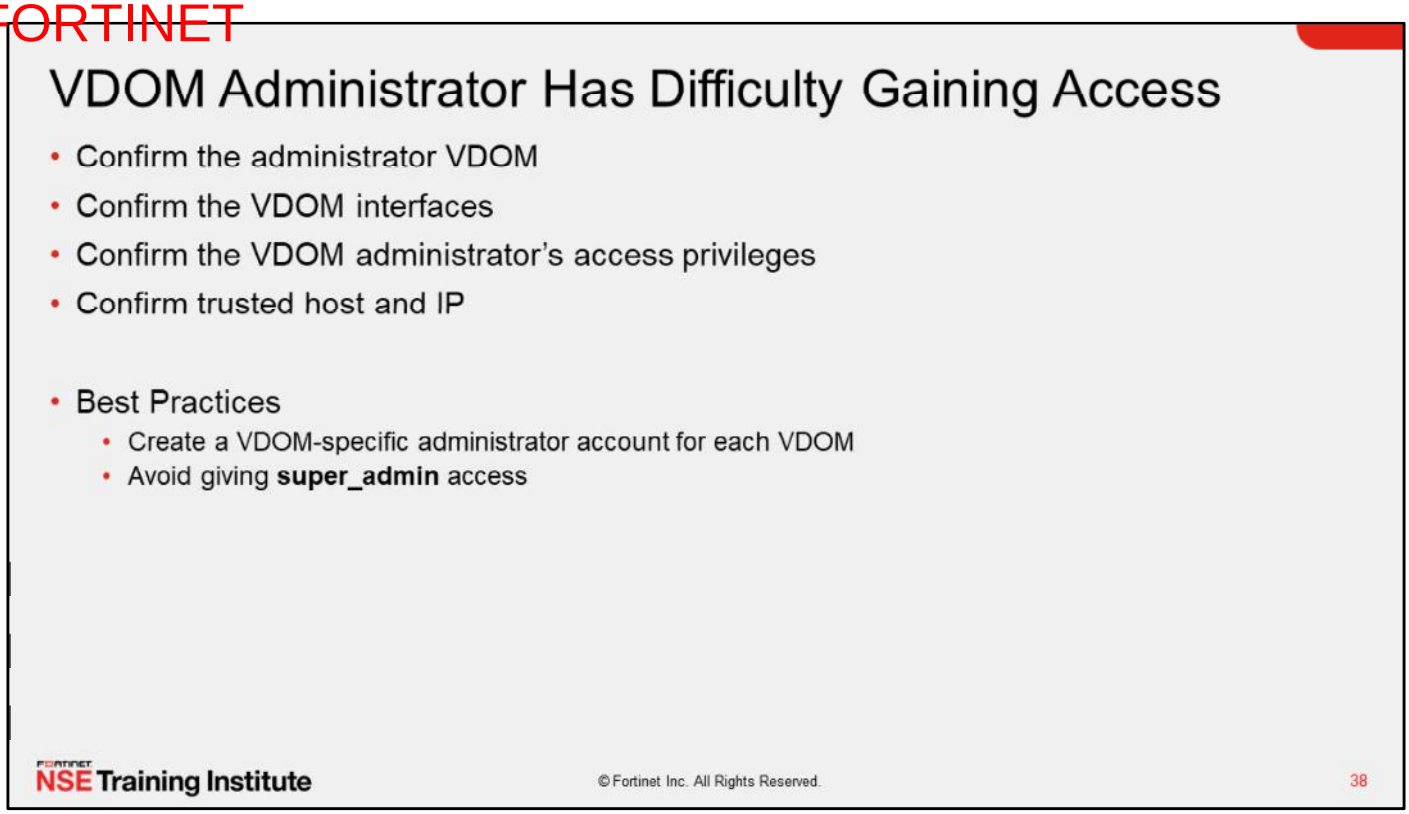

With VDOMs configured, administrators have an extra layer of permissions and may have problems accessing the desired information.

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.

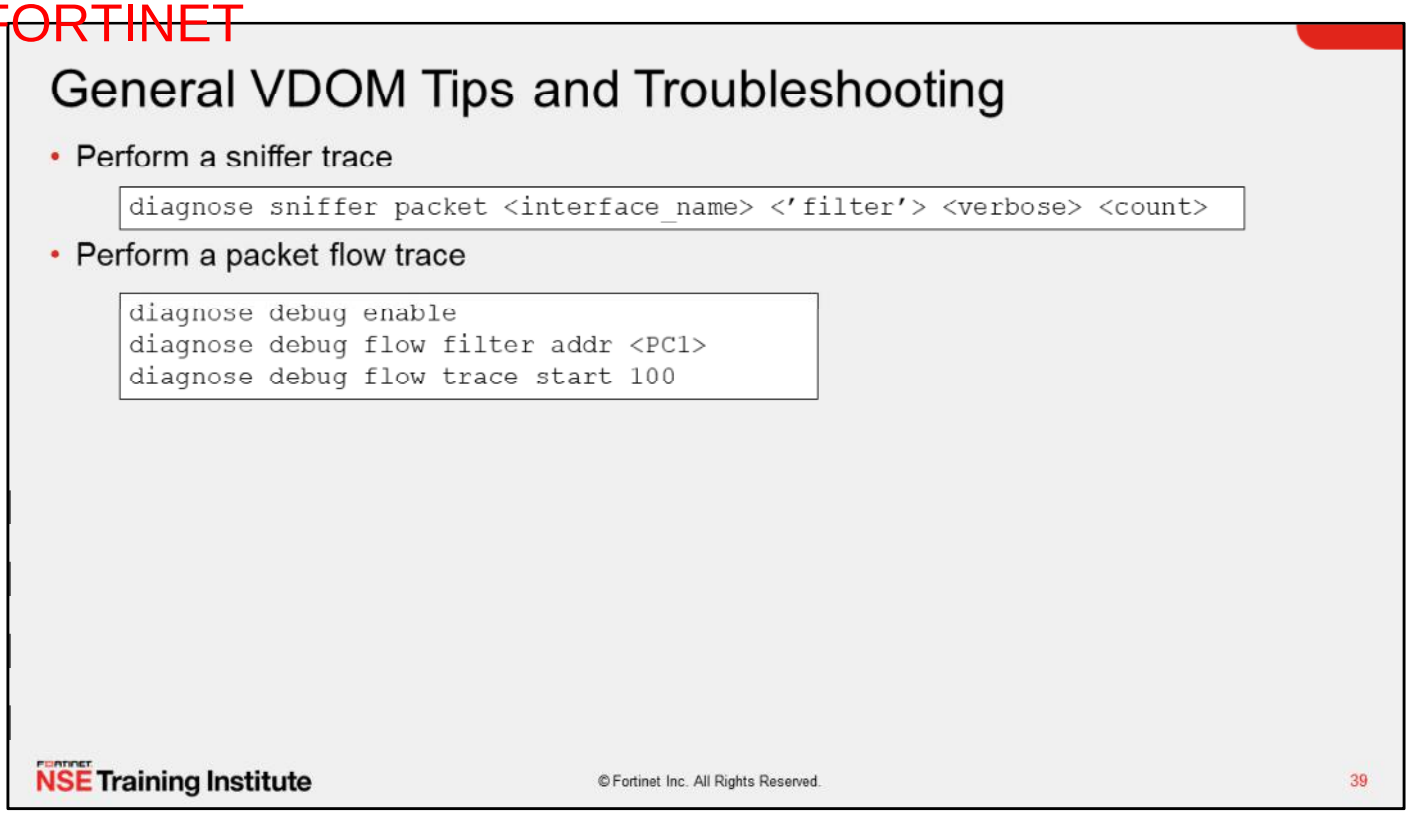

Besides ping and traceroute, there are additional tools for troubleshooting your VDOM configurations. The primary tools for VDOM troubleshooting include packet sniffing and debugging the packet flow.

- 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

NSE Training Institute

40

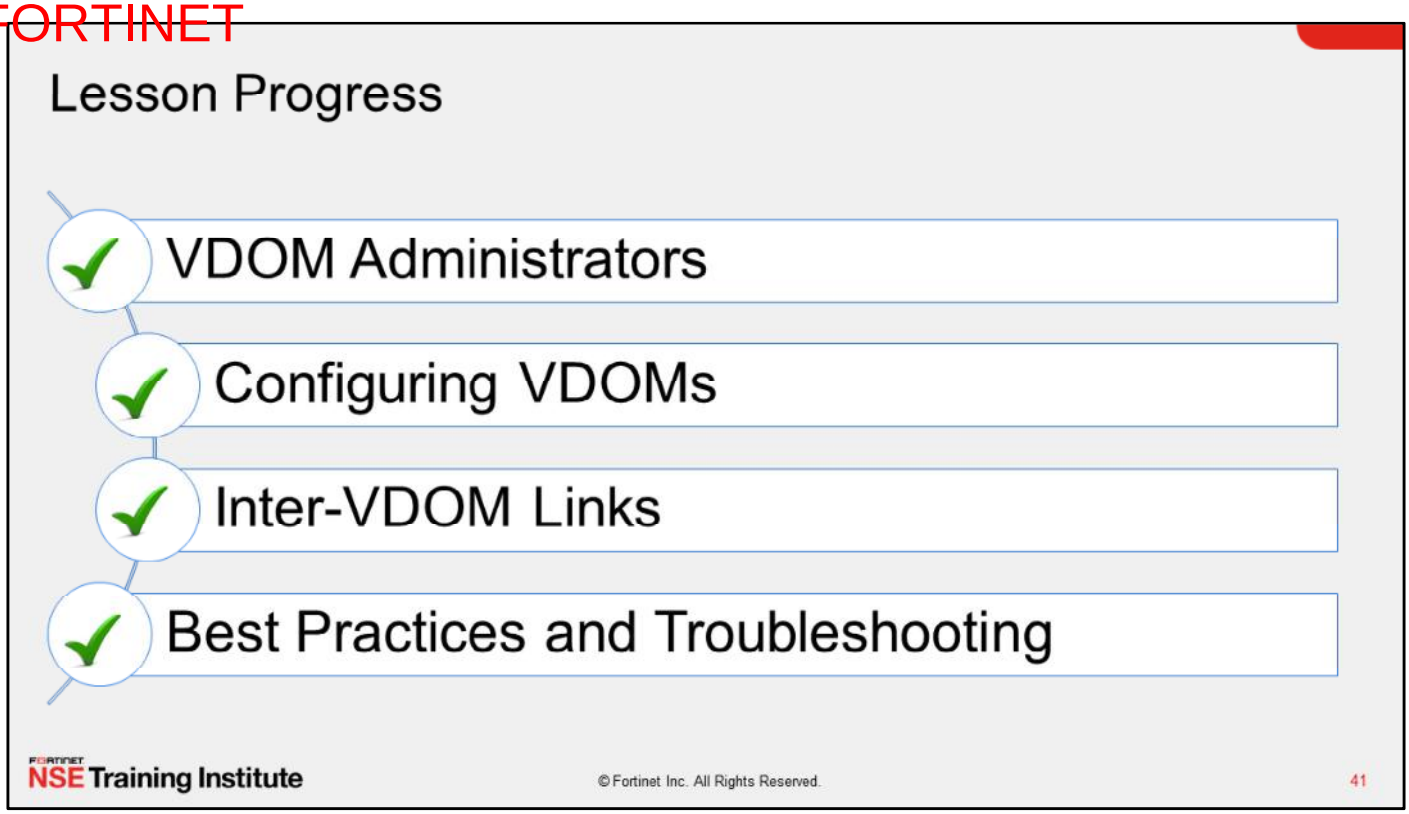

Congratulations! You have completed this lesson.

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

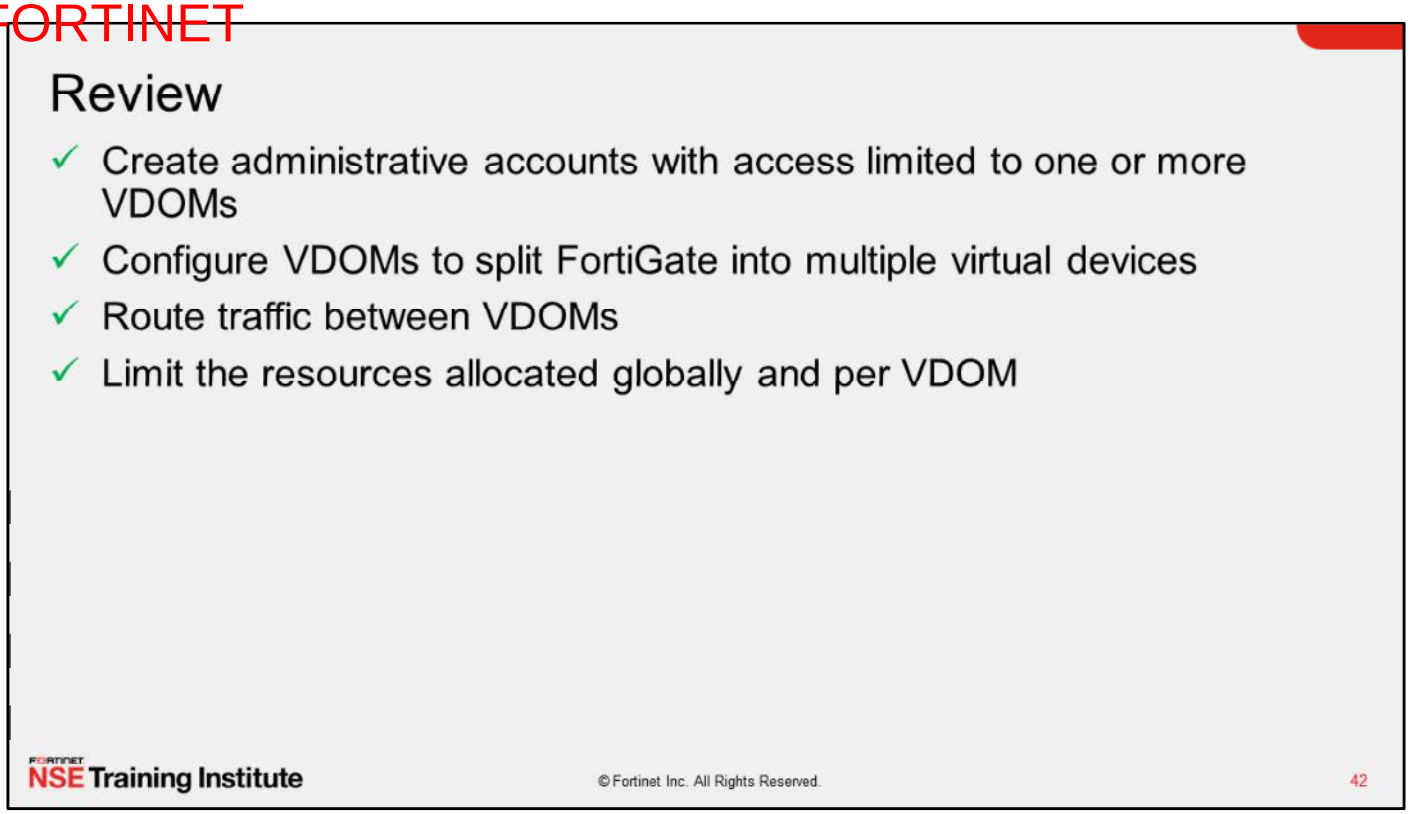

This slide shows the objectives that you covered in this lesson.

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

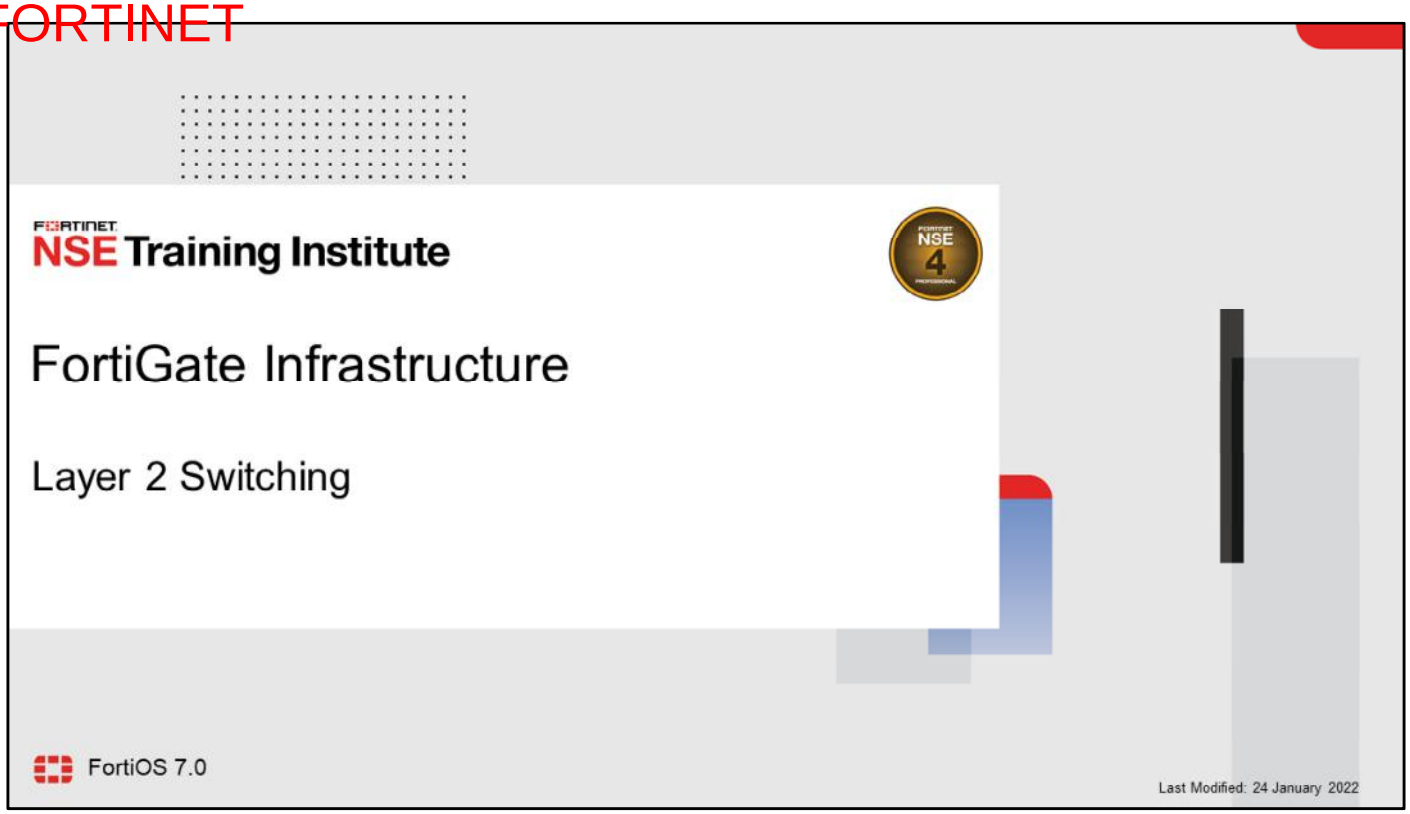

In this lesson, you will learn how to use transparent operation mode and Layer 2 switching on FortiGate.

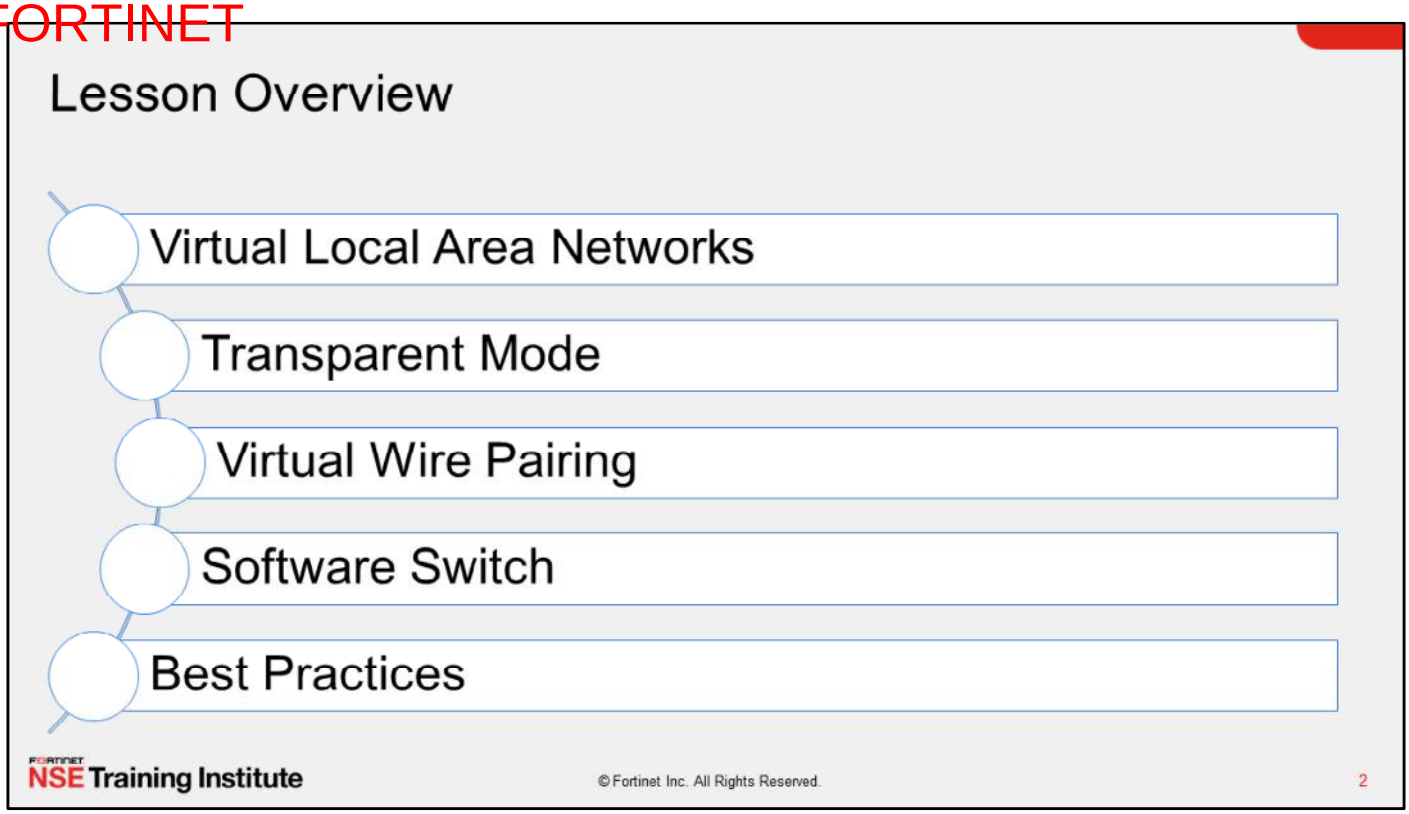

In this lesson, you will learn about the topics shown on this slide.

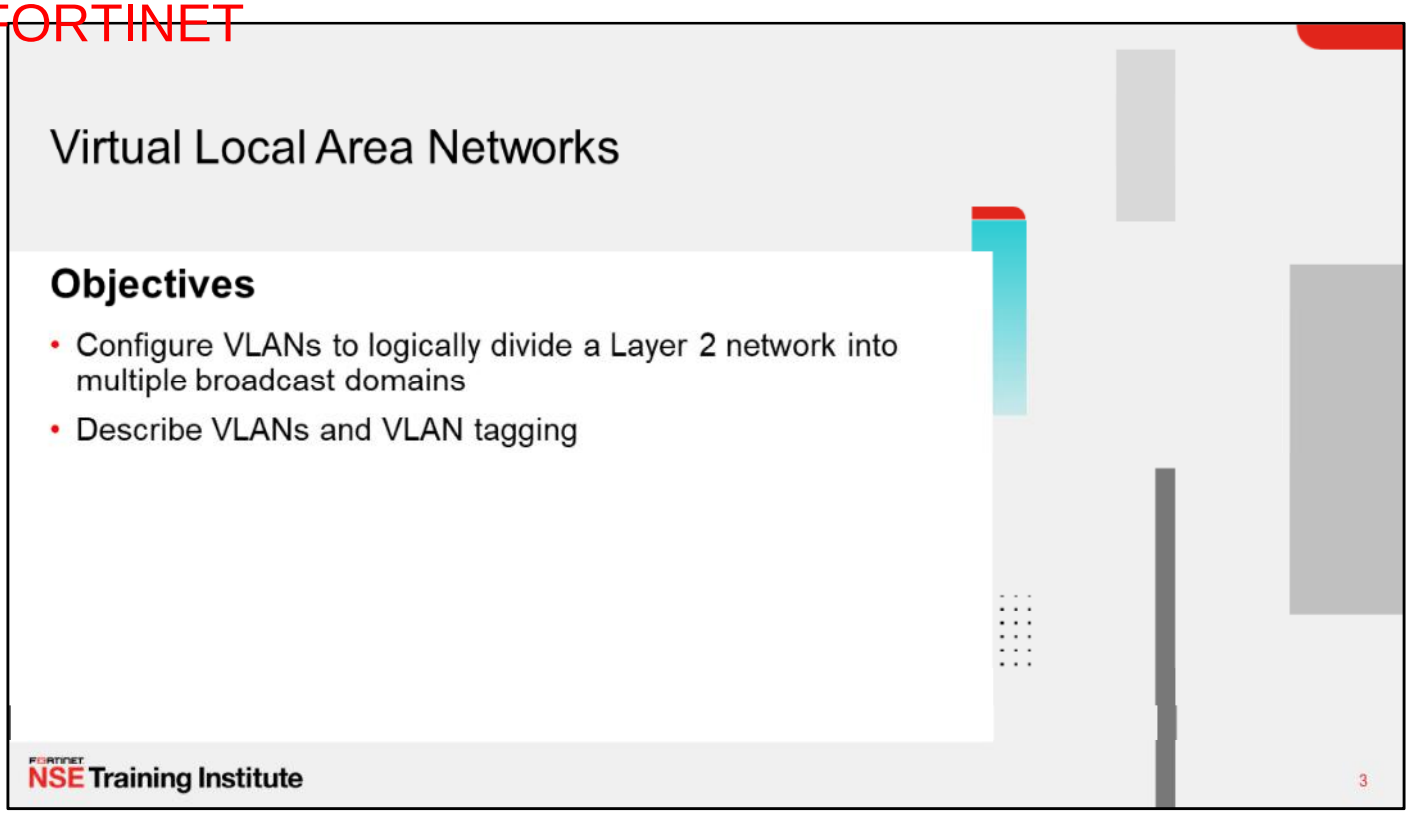

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding and configuring VLANs, you will be able to effectively divide your network into smaller, logical segments.

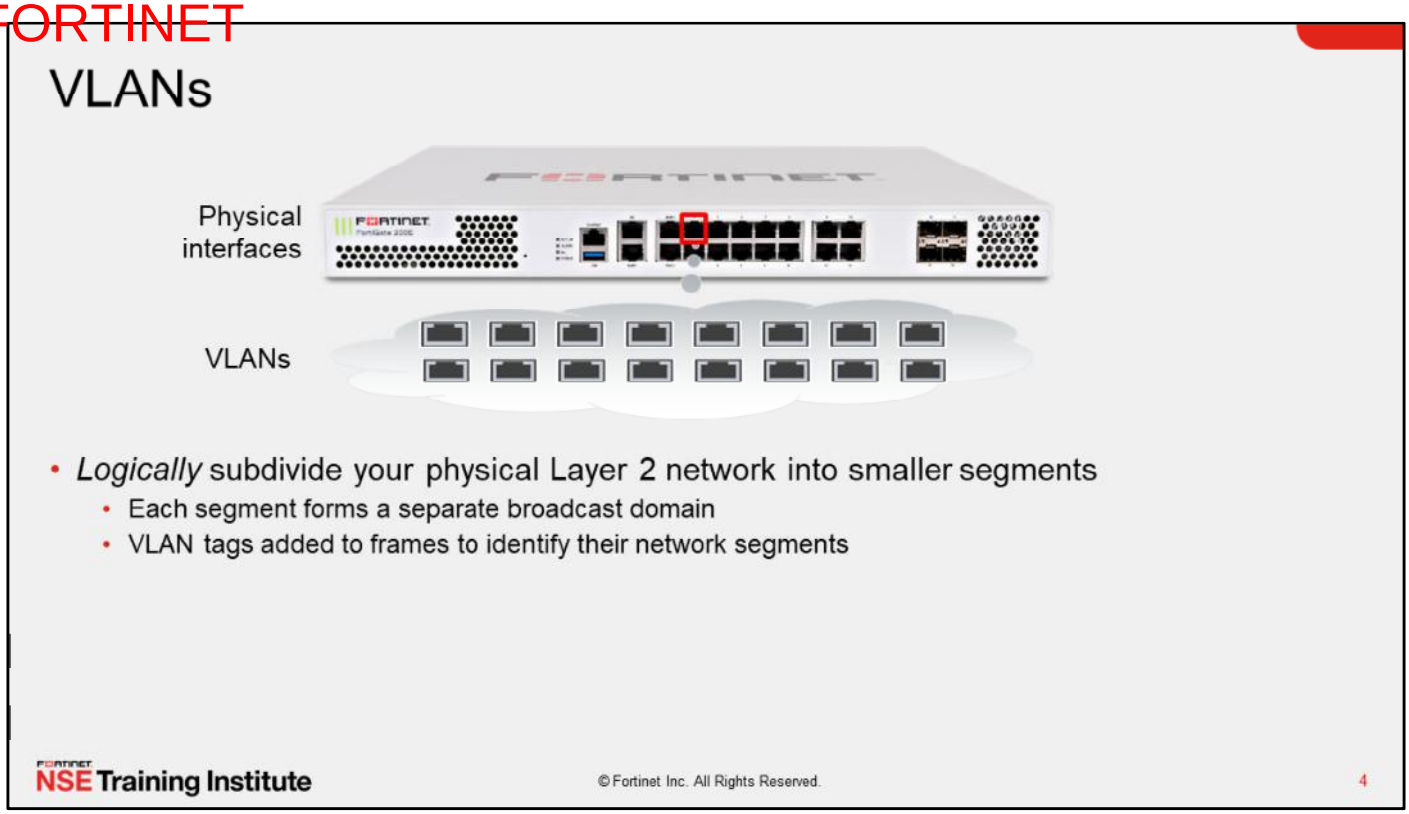

VLANs split your physical LAN into multiple, logical LANs. In NAT operation mode, each VLAN forms a separate broadcast domain. Multiple VLANs can coexist in the same physical interface, provided they have different VLAN IDs. In this way, a physical interface is split into two or more logical interfaces. A tag is added to each Ethernet frame to identify the VLAN to which it belongs.

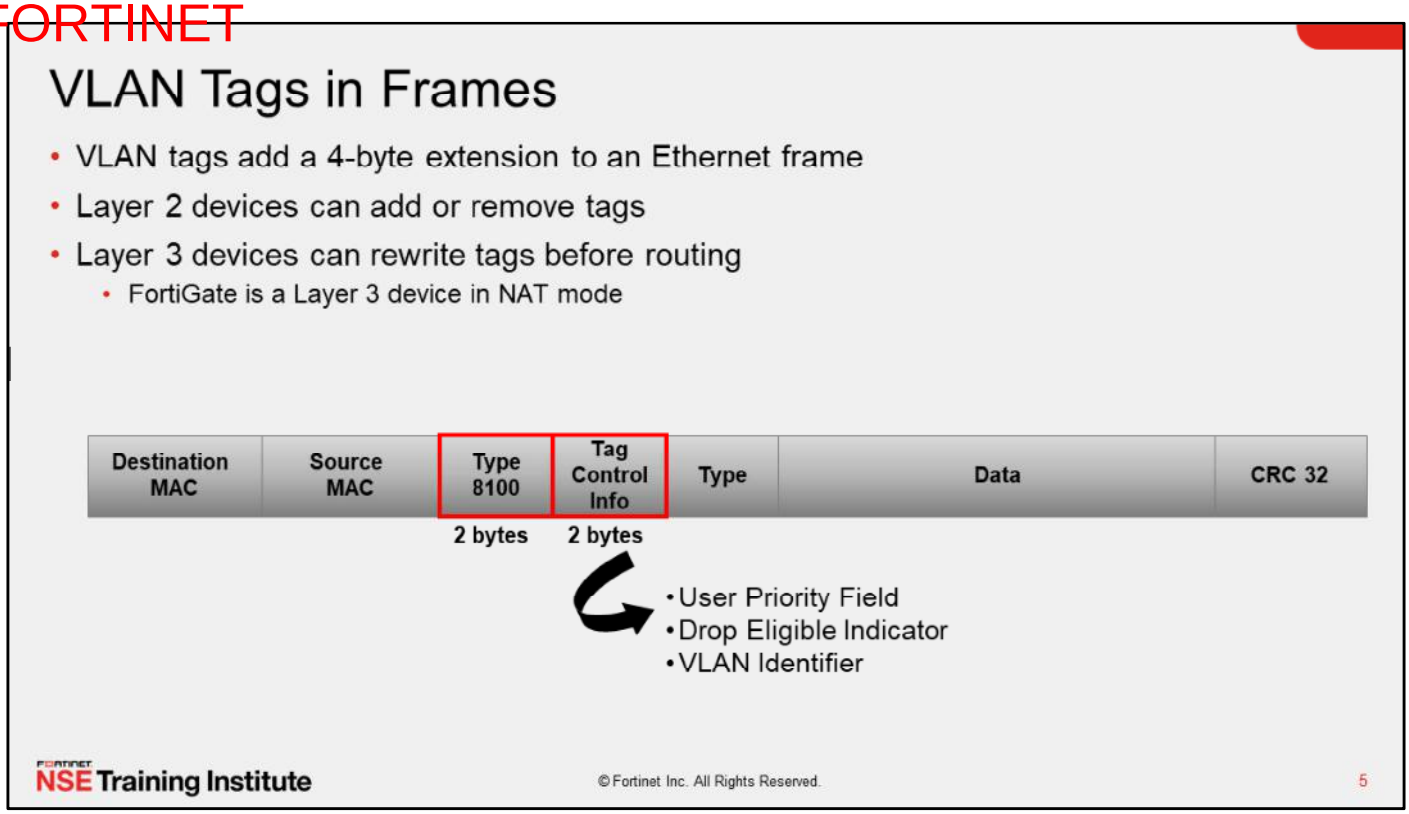

This slide shows an Ethernet frame. The frame contains the destination and source MAC addresses, the type, the data payload, and a CRC code, to confirm that it is not corrupted.

In the case of Ethernet frames with VLAN tagging, according to the 802.1q standard, four more bytes are inserted after the MAC addresses. They contain an ID number that identifies the VLAN.

An OSI Layer 2 device, such as a switch, can add or remove these tags from Ethernet frames, but it cannot change them.

A Layer 3 device, such as a router or FortiGate device, can change the VLAN tag before proceeding to route the packet. In this way, they can route traffic between VLANs.

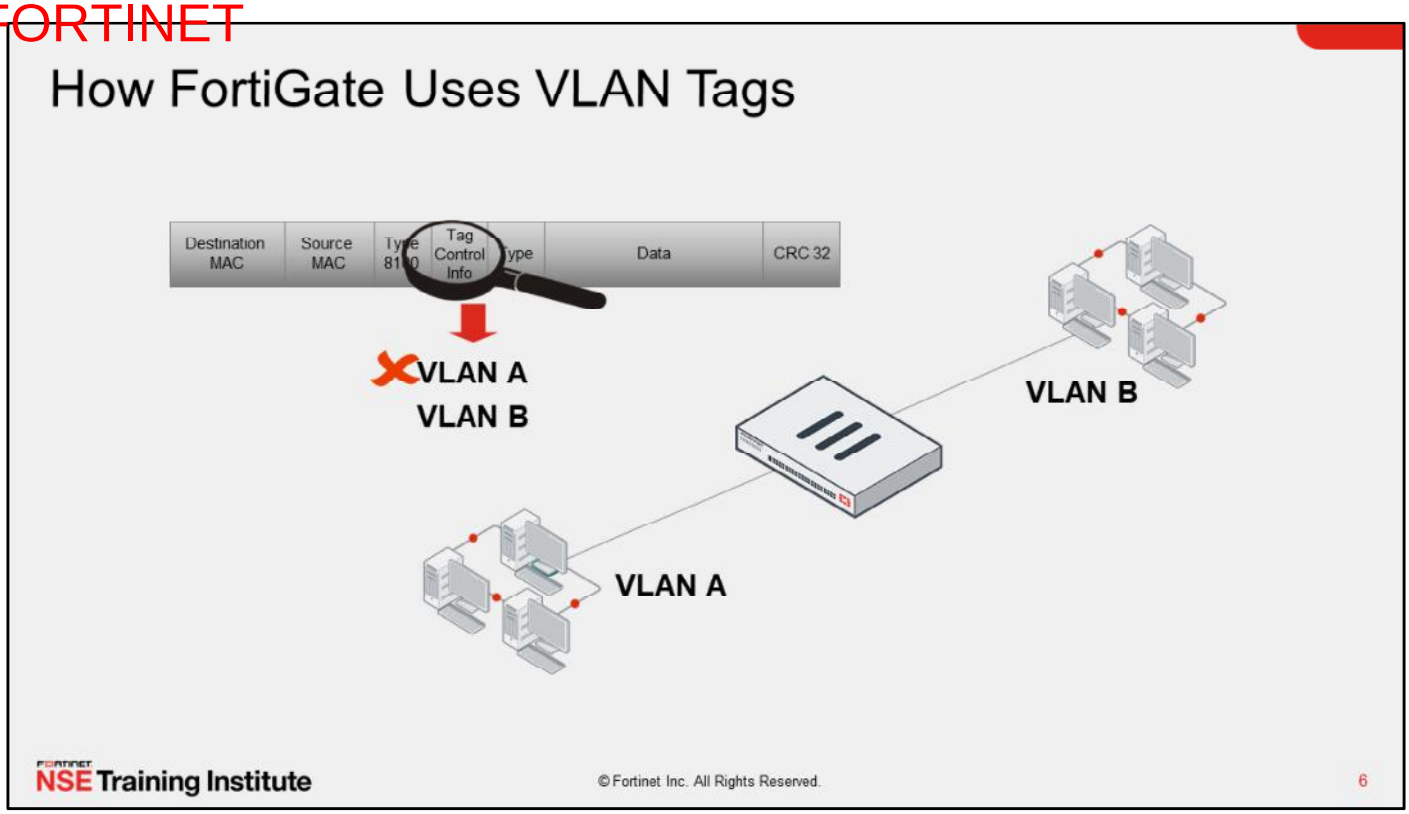

When operating in NAT mode, FortiGate operates as an OSI Layer 3 router in its most basic configuration. In this mode, a VLAN is an interface on the device. VLAN tags may be added on egress, removed on ingress, or rewritten based on a routing decision. FortiGate does not add VLAN tags on ingress (this is the responsibility of a previous device).

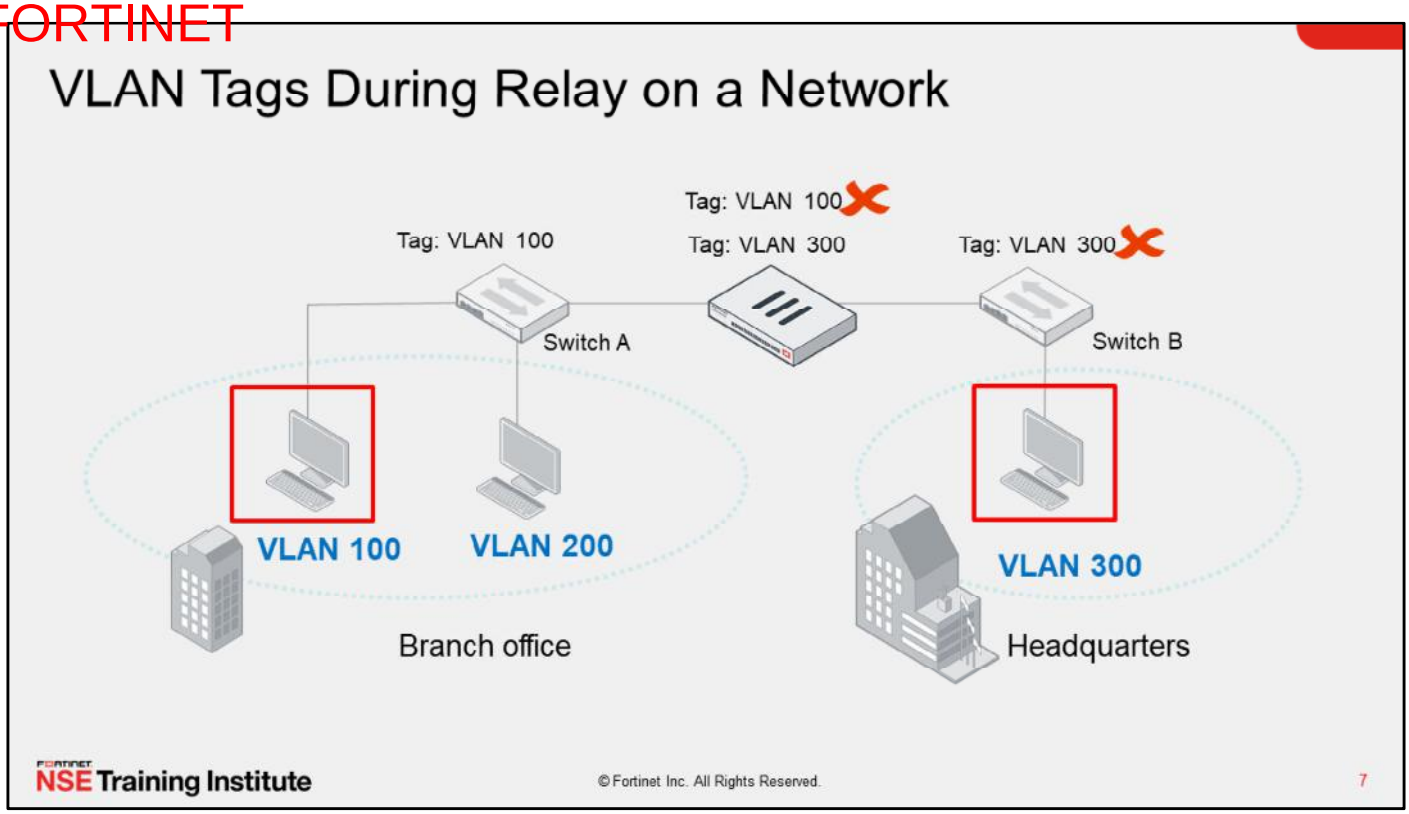

In the example shown on this slide of NAT operation mode, a host on VLAN 100 sends a frame to a host on VLAN 300. Switch A receives the frame on the untagged VLAN 100 interface. After that, it adds the VLAN 100 tag on the tagged trunk link between switch A and FortiGate.

FortiGate receives the frame on the VLAN 100 interface. Then, it routes the traffic from VLAN 100 to VLAN 300, rewriting the VLAN ID to VLAN 300 in the process.

Switch B receives the frame on the VLAN trunk interface and removes the VLAN tag before forwarding the frame to its destination on the untagged VLAN 300 interface.

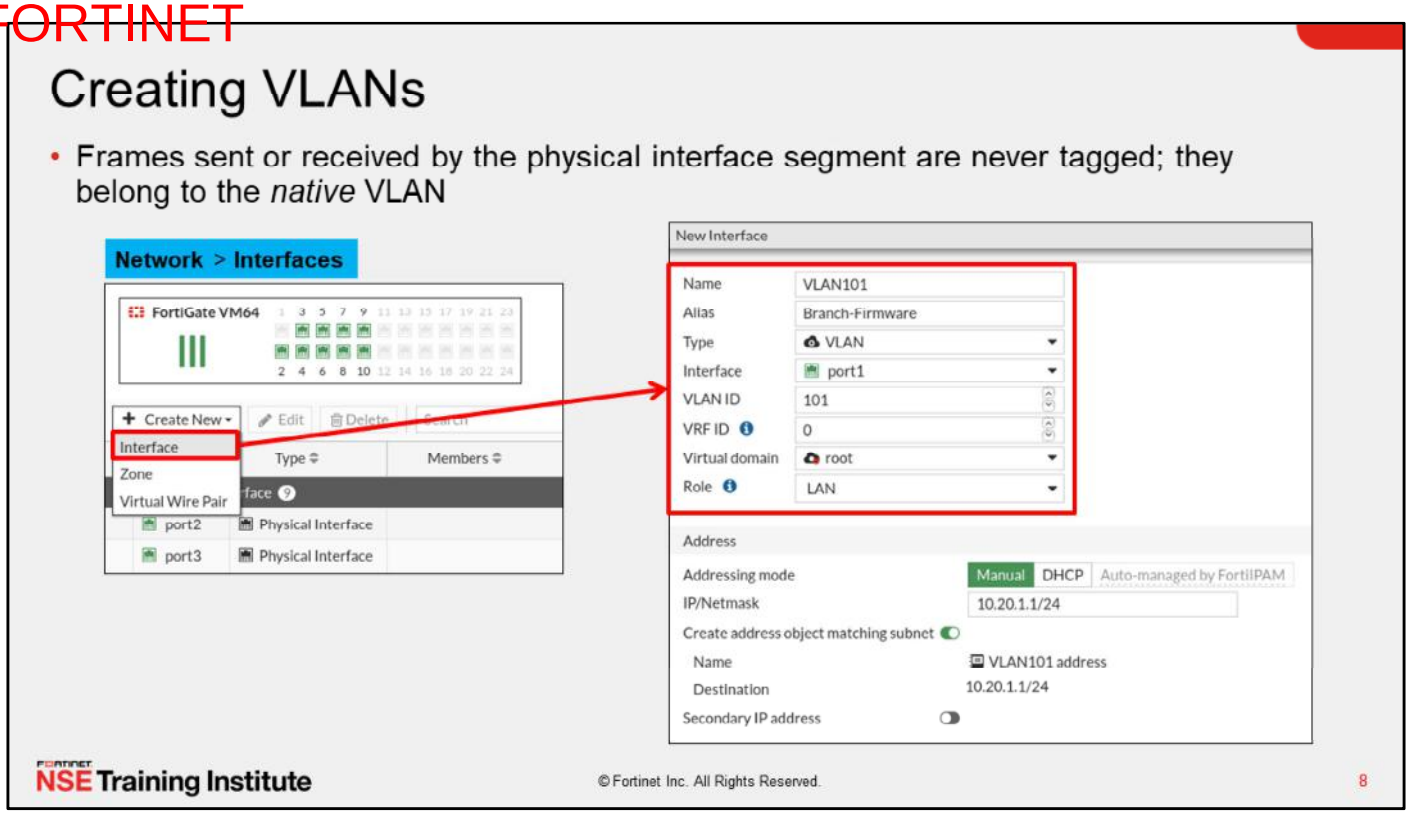

To create a VLAN using the GUI, click **Create New,** select **Interface,** and then, in the **Type** drop-down list, select **VLAN**. You must specify the VLAN ID and the physical interface to which the VLAN will be bound. Frames that belong to interfaces of that type are always tagged. On the other hand, frames sent or received by the physical interface segment are never tagged. They belong to what is called the *native* VLAN (VLAN ID 0).

Note that in a multi-VDOM environment, the physical interface and its VLAN sub-interface can be in separate VDOMs.

Knowledge Check

1. Which mode must the FortiGate VDOM be operating in, to route traffic between VLANs?
   - A. Transparent mode
   - ✓B. NAT mode

9

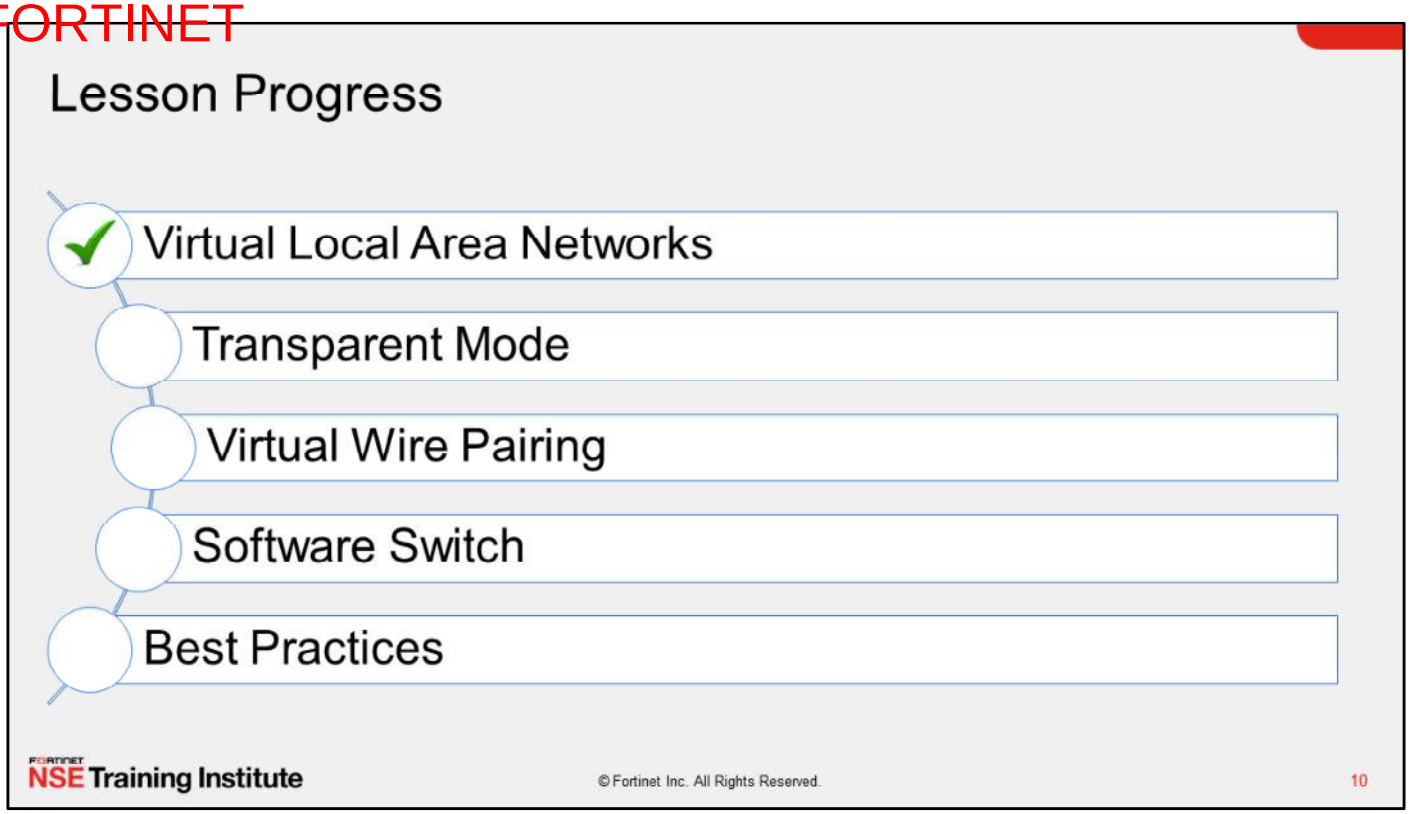

Good job! You now understand VLANs.

Now, you will learn about transparent mode.

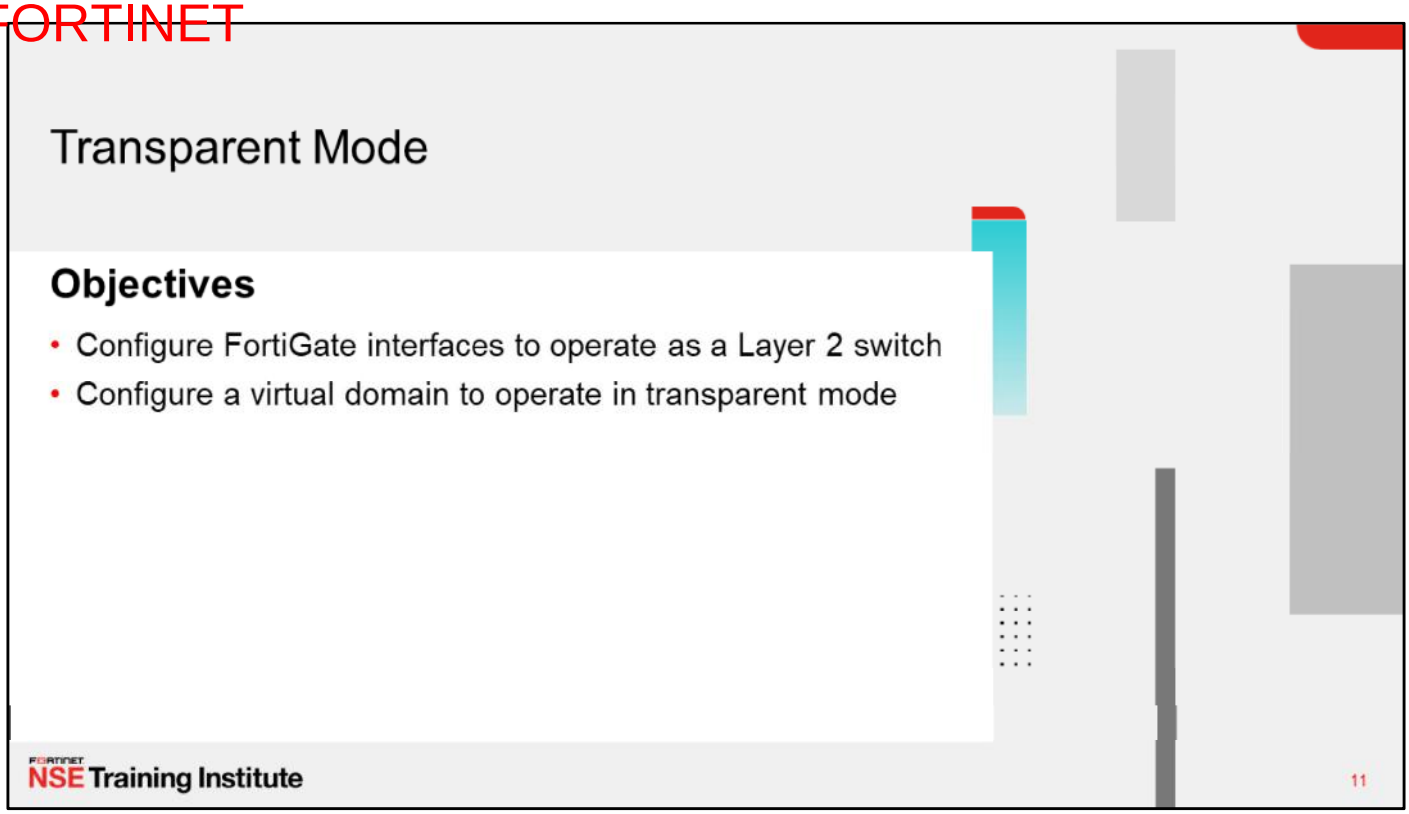

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding and configuring transparent operation mode, you will understand a key component of implementing Layer 2 switching on FortiGate.

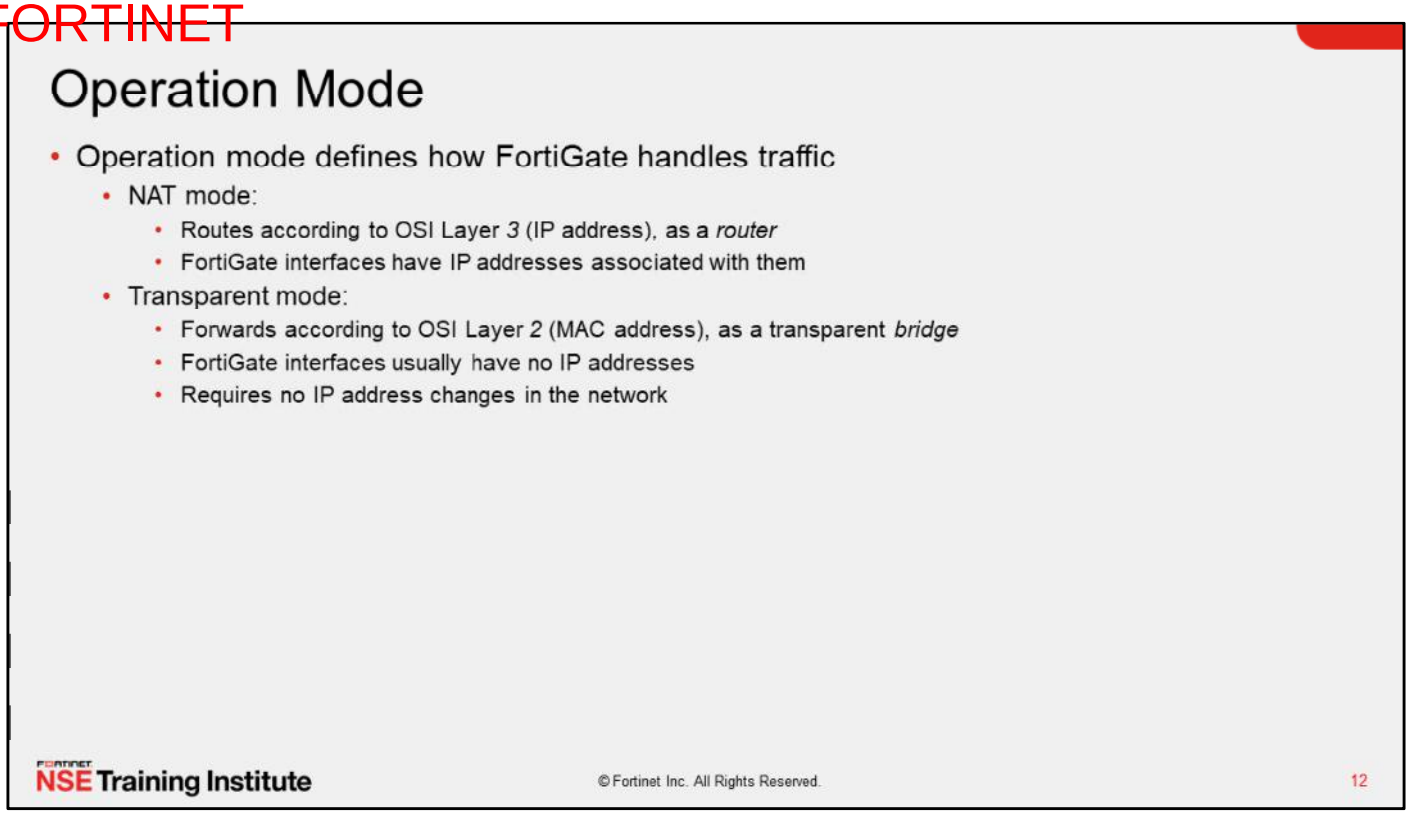

Traditional IPv4 firewalls and NAT mode FortiGate devices handle traffic the same way that routers do. Each interface must be in a different subnet and each subnet forms a different broadcast domain. FortiGate routes IP packets based on the IP header information, overwriting the source MAC address. So, 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.

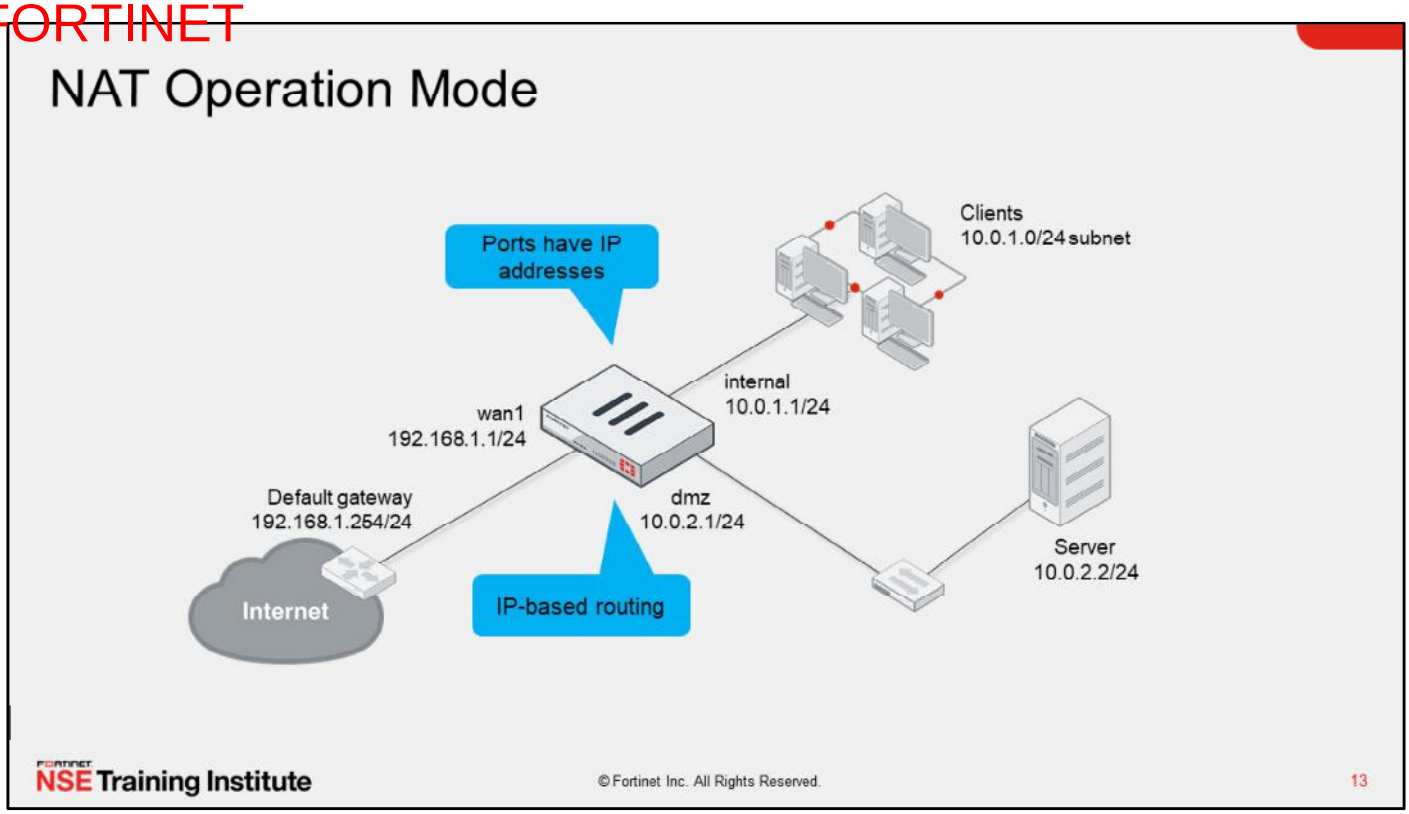

This slide shows an example of NAT operation mode.

FortiGate has three connected ports, each with separate IP subnets. All interfaces on FortiGate have IP addresses, and, in this case, NAT translates between networks. Firewall policies allow traffic to flow between networks.

FortiGate handles packets according to their routes. In most cases, routes are based on the destination IP address (at Layer 3 of the OSI model).

Clients on each subnet send frames that are destined for a FortiGate MAC address—not the real MAC address of the server.

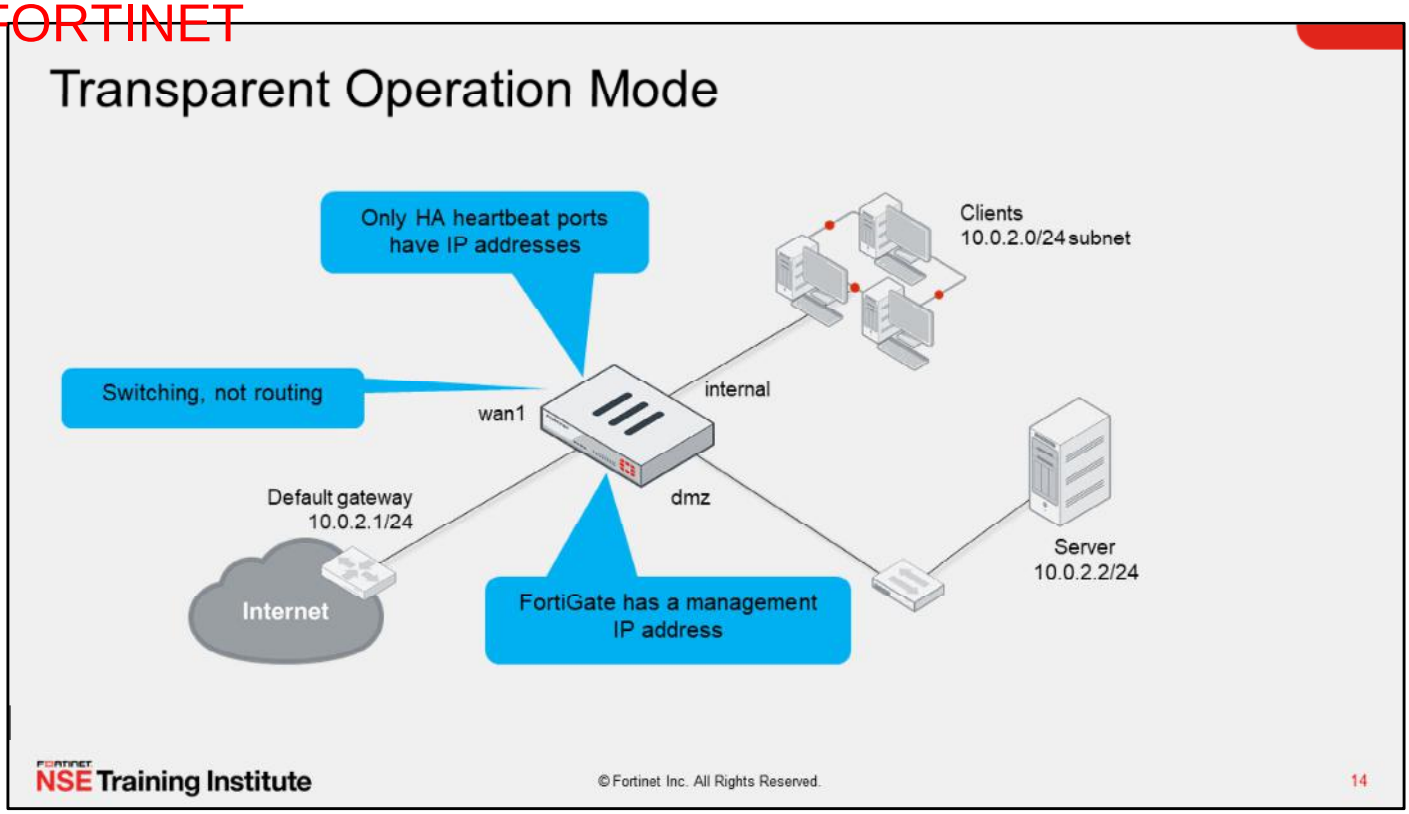

This slide shows an example of transparent operation mode. Firewall policies scan, then allow or block traffic. But there are differences.

Notice that the physical interfaces on FortiGate have no IP addresses. Therefore, FortiGate won't respond to ARP requests. However, there are some exceptions. For example, when changing to transparent operation mode, you must specify a management IP address in order to receive connections from your network administrators and send log messages, SNMP traps, alert email, and so on. This IP address is not assigned to a specific interface. It is assigned to the VDOM settings. The management IP address has no effect on traffic passing through FortiGate.

By default, a transparent mode FortiGate device won't perform NAT. Also, clients send frames destined directly for the real router or server MAC address.

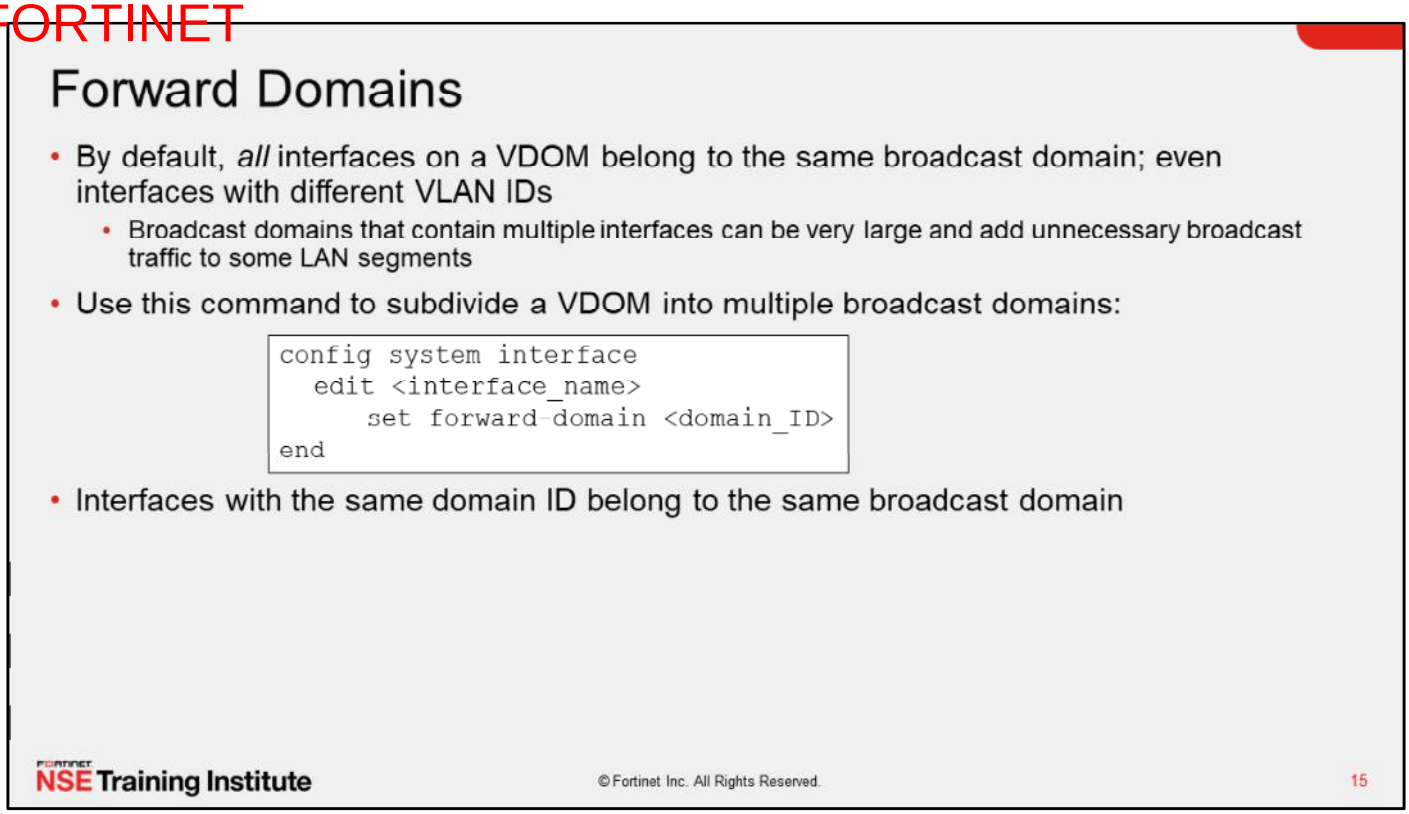

By default, in transparent operation mode, each VDOM forms a separate forward domain; however, interfaces do not. How does this affect the network?

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.

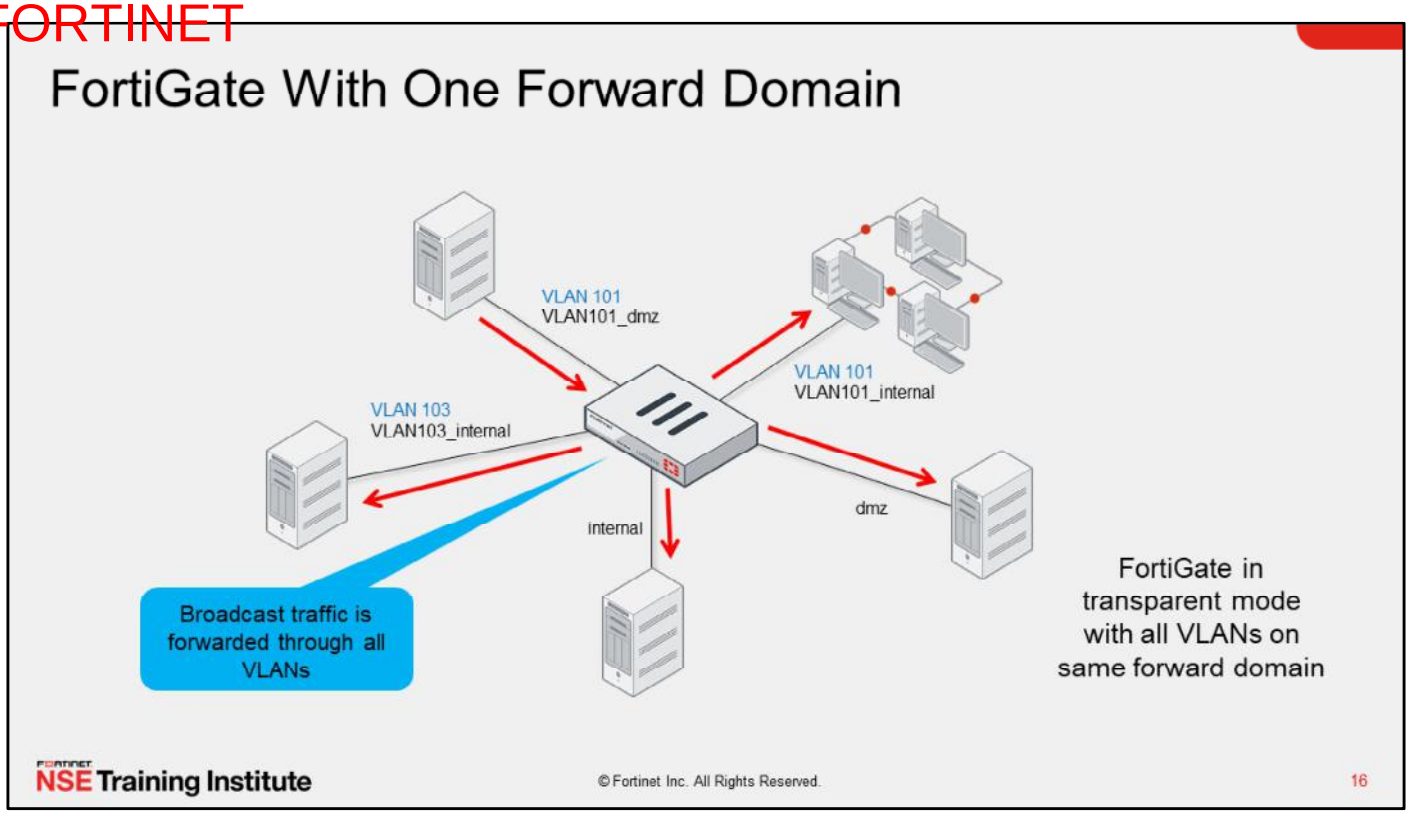

This slide illustrates a problem—a broadcast with all the interfaces on the forward domain 0 (default). One device sends an ARP request. It reaches FortiGate through one of the interfaces in the VDOM.

Because all interfaces belong to the same forward domain, FortiGate rebroadcasts to all the other interfaces, even to interfaces that belong to different VLANs. This generates unnecessary traffic. After that, the ARP reply still arrives on only one interface, and FortiGate learns that the MAC is on that interface.

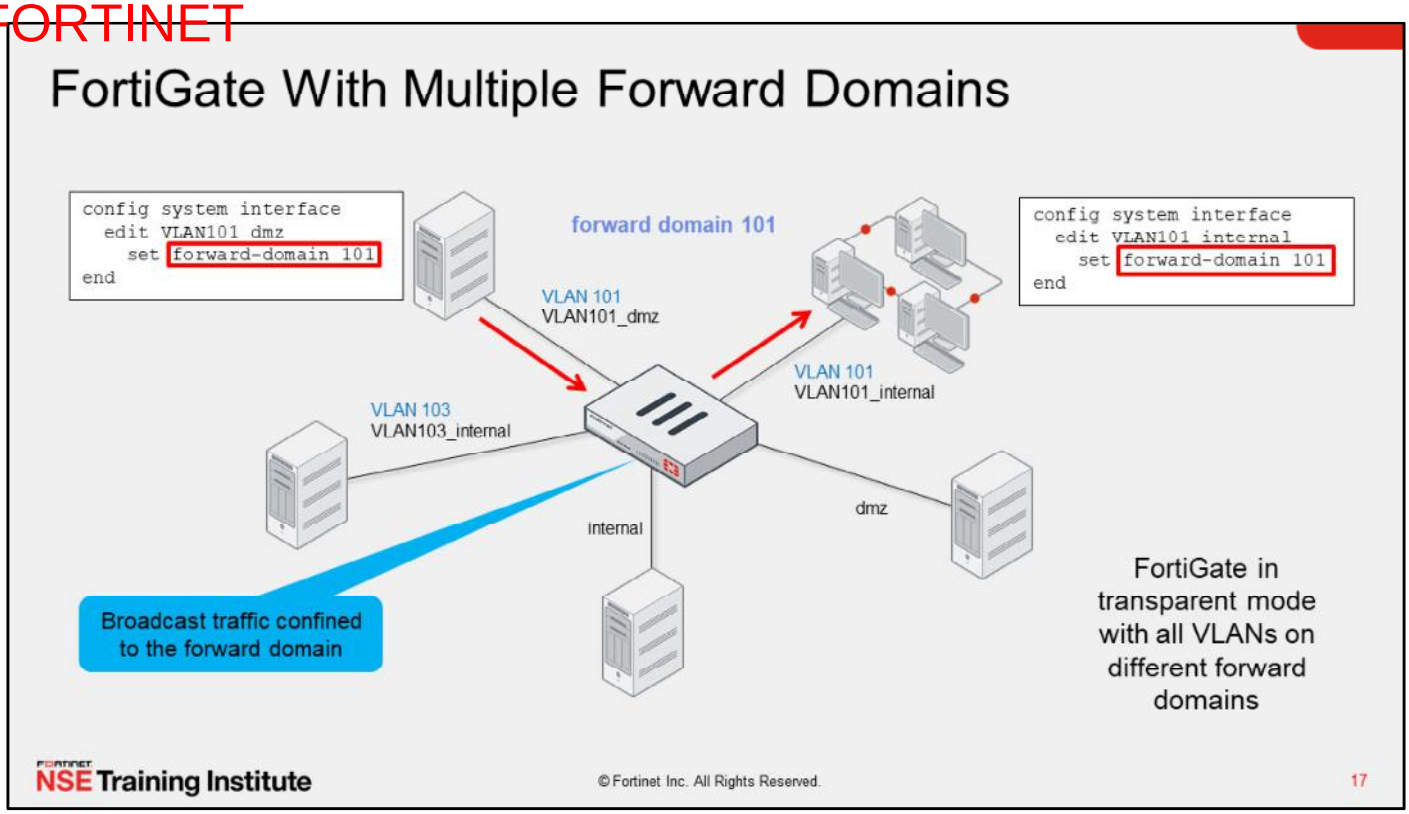

As you learned earlier in this lesson, forward domains are like broadcast domains.

The example on this slide shows the same network as before, but different forward domain IDs are assigned to each VLAN.

Traffic arriving on one interface is broadcast only to interfaces that are in the same forward domain ID.

Transparent Mode MAC Table

```
# diagnose netlink brctl name host <vdom name>.b
```

```
show bridge control interface inspect.b host.
fdb: size=2048, used=5, num=5, depth=1
Bridge inspect.b host table
port no device  devname mac addr                ttl     attributes
  2     22      vlink1  1e:44:d1:3a:00:15       0        Hit(0)
  1     3       port1   00:0c:29:b7:1d:ed       144      Hit(144)
  1     3       port1   00:0c:29:2e:e0:4e       0       Local Static
  1     3       port1   00:0c:29:8c:36:cc       0        Hit(0)
  2     22      vlink1  7e:73:da:d2:00:16       0       Local Static
```


18

This debug command lists the MAC address table in a VDOM operating in transparent mode. The table shown on this slide contains the outbound interfaces to reach each learned MAC address.

Knowledge Check

1. Which statement about FortiGate operating in transparent mode is true?
   - ✓A. It has a management IP address.
   - B. Each interface has its own IP address.

2. How can an administrator configure FortiGate to have four interfaces in the same broadcast domain?
   - A. Create a firewall policy on each of the four interfaces
   - ✓B. Configure the operation mode as transparent and use the same forward domain ID

NSE Training Institute

19

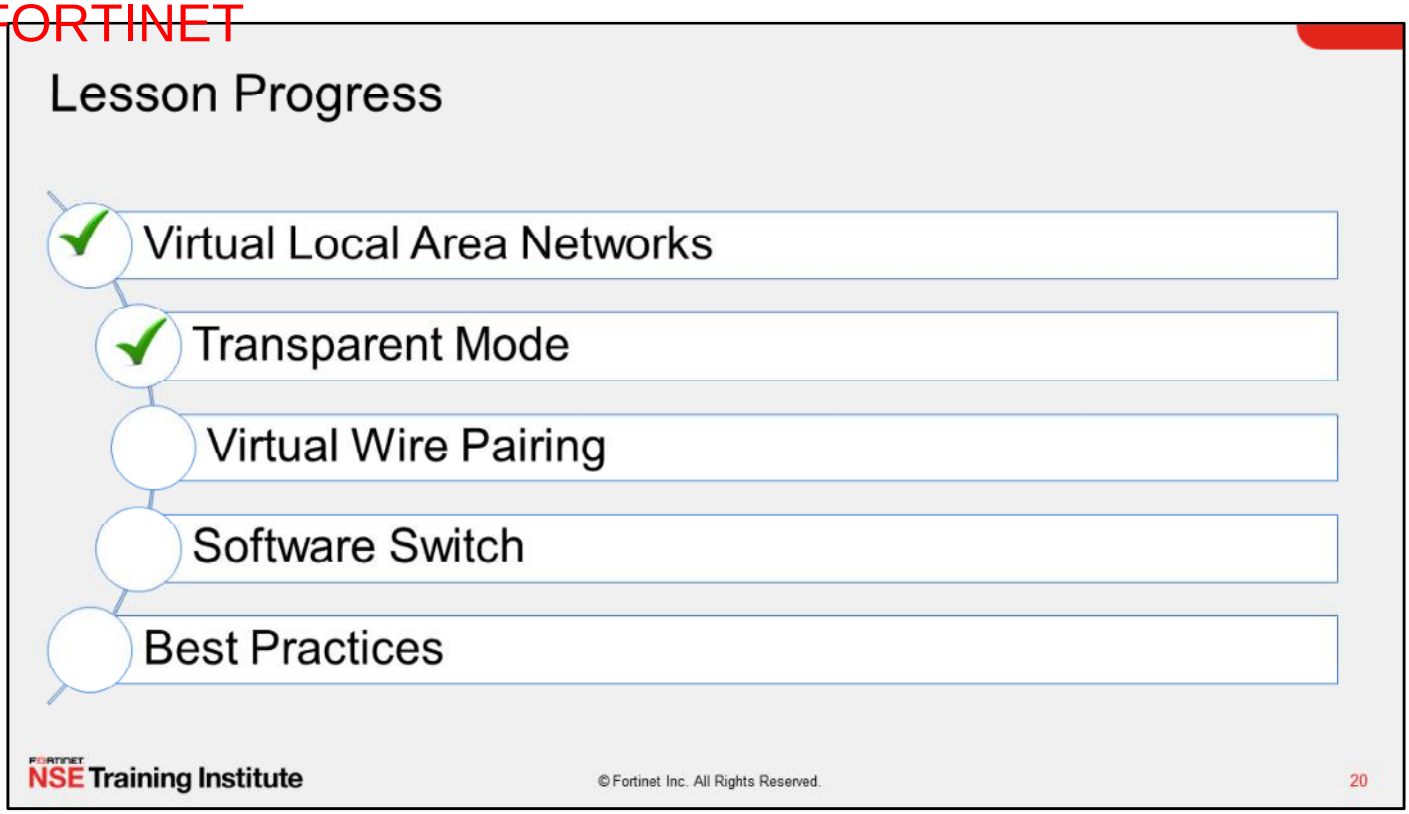

Good job! You now understand transparent mode.

Now, you will learn about virtual wire pairing.

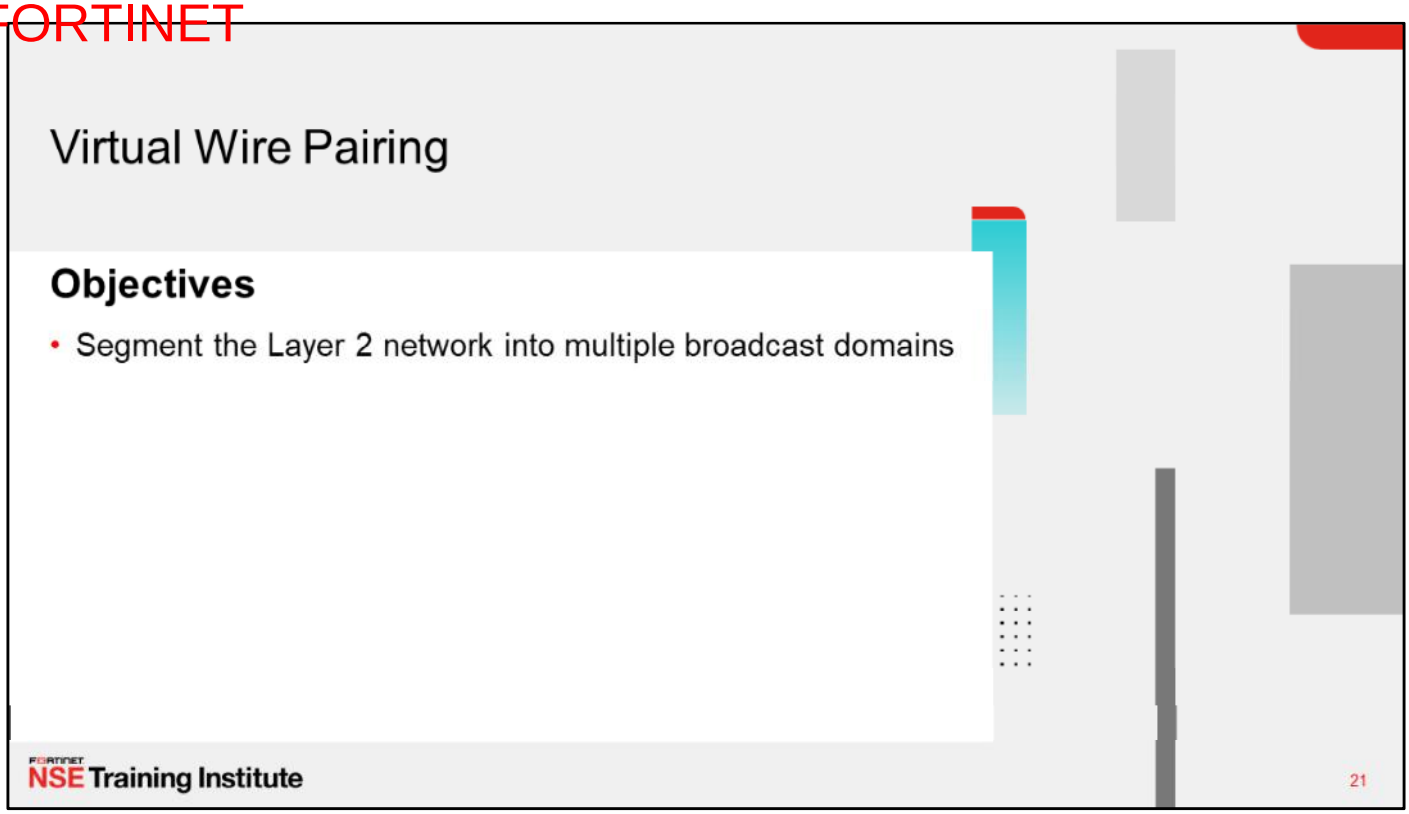

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in using virtual wire pairing, you will understand another way to create broadcast domains and be able to create interface pairs operating like transparent mode in a NAT VDOM.

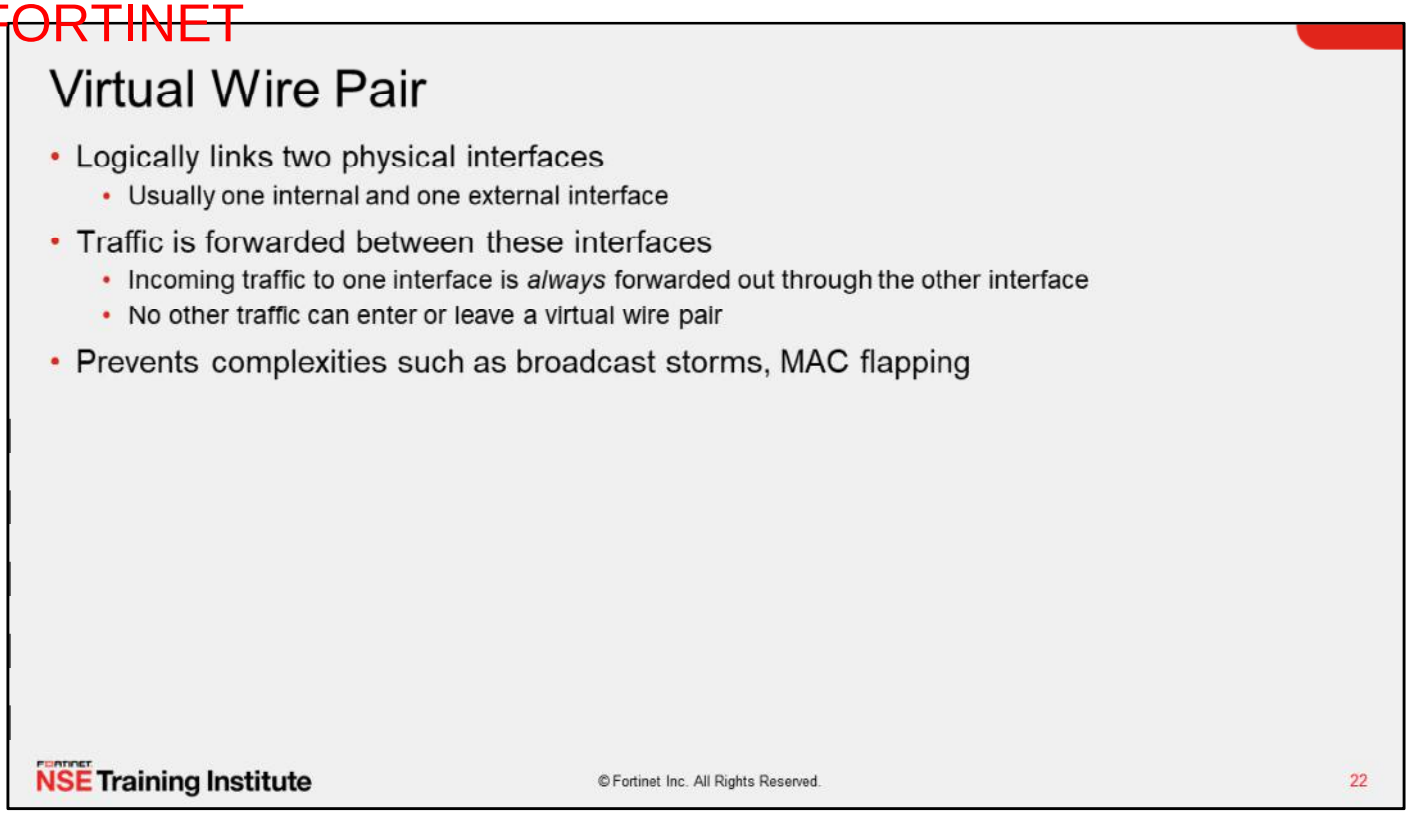

You can use virtual wire pairing when only two physical interfaces need to be connected to the same broadcast domain. The most frequently seen example of this is FortiGate connected between the internal network and the ISP router.

When you configure virtual wire pairing, two ports are logically bound or linked, acting like a filtered cable or pipe. All the traffic that arrived at one port is forwarded to the other port. This prevents issues related to broadcast storms or MAC address flapping.

You can create more than one virtual wire pair on FortiGate.

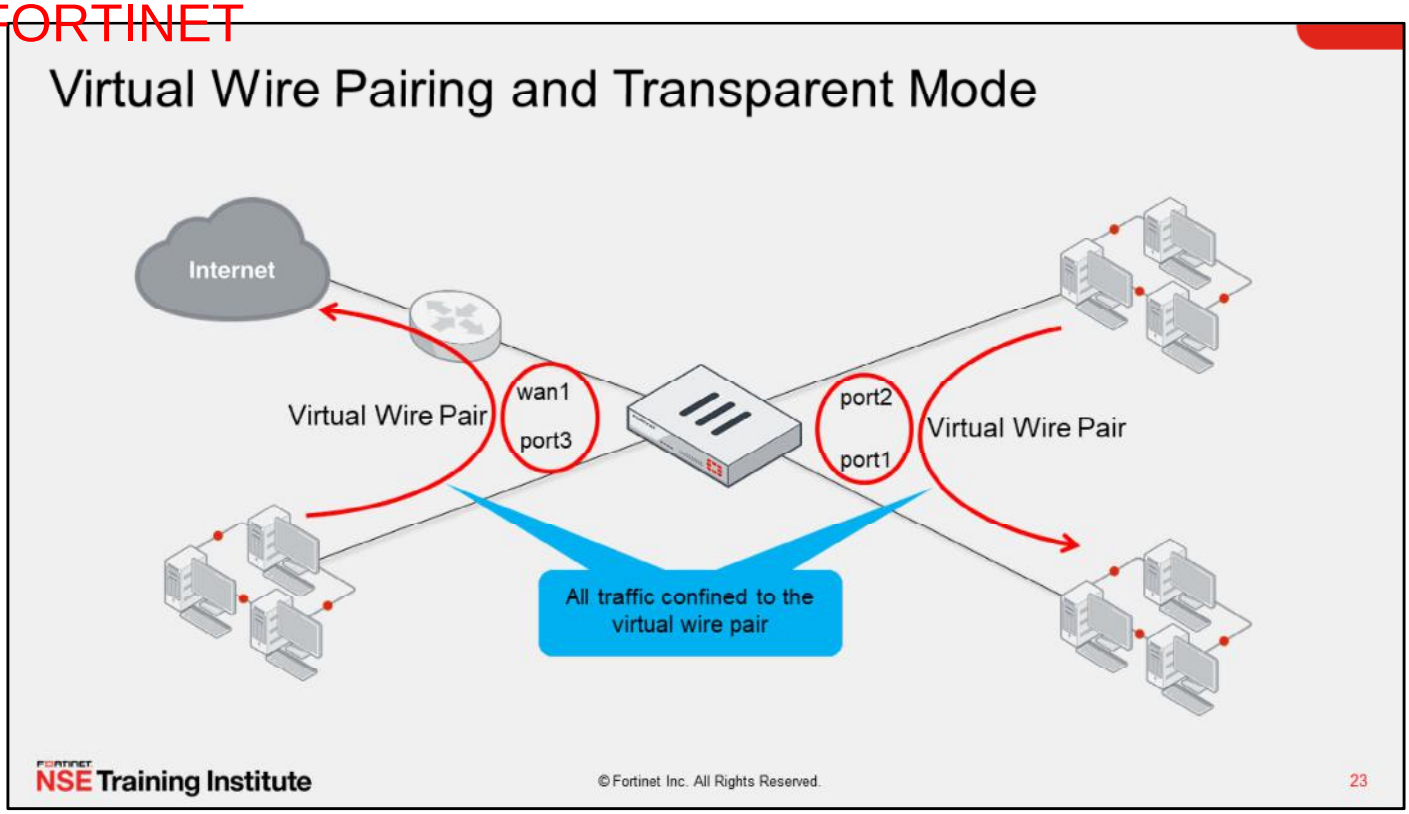

This slide shows an example of two virtual wire pairs used on FortiGate in transparent mode.

This FortiGate has four ports, each connected to different physical locations. But traffic is not allowed to flow between all four locations. Virtual wire pairing allows traffic only between ports in the same pair: between **port1** and **port2**, and between **port3** and **wan1**.

So, in this example, the network on **port3** can reach the internet through **wan1**. However, the networks on **port2** and **port1** can't reach the internet. They can reach only each other.

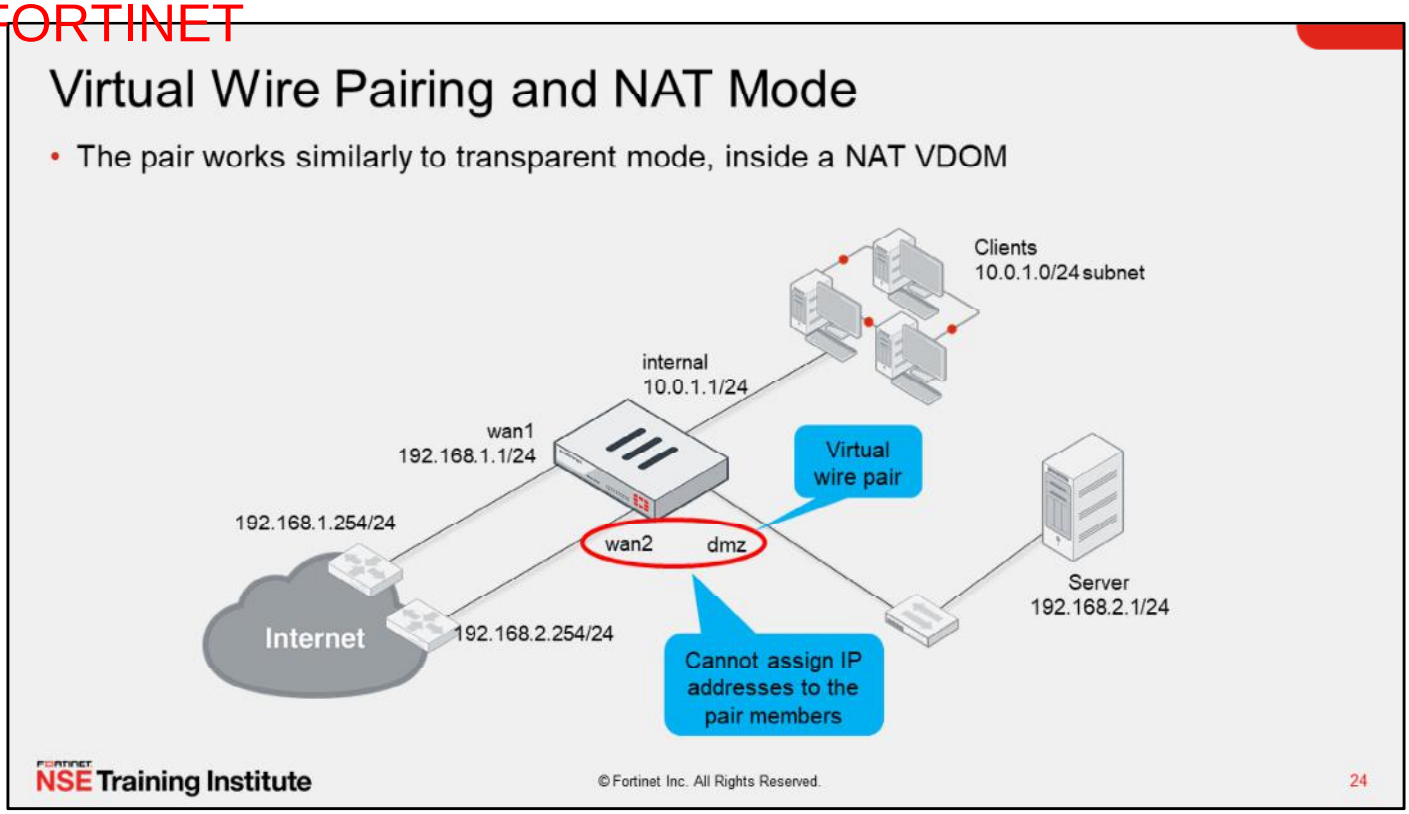

This slide shows an example of a virtual wire pair on FortiGate operating in NAT mode. In this example, IP packets ingressing interfaces **wan1** and **internal** are routed using the IP header information. Those two interfaces have different IP addresses, and each one forms a separate broadcast domain.

The case of interfaces **wan2** and **dmz** are different. Because these interfaces are configured as a virtual port pair, they don't have assigned IP addresses, and they form one single broadcast domain. Observe the IP addresses for the server and the router connected to **wan2**. They must both belong to the same subnet.

So, virtual wire pairing offers a way to mix NAT mode functionalities with some transparent mode functionalities in the same VDOM.

This scenario is most commonly used as a segmentation firewall. This configuration allows for integrating FortiGate into an existing network where the web server is on a public IP address. It prevents the need to use a virtual IP, and provides isolation from the rest of network. It also allows immediate integration of FortiGate into an established network and provides a migration path to the FortiGate infrastructure.

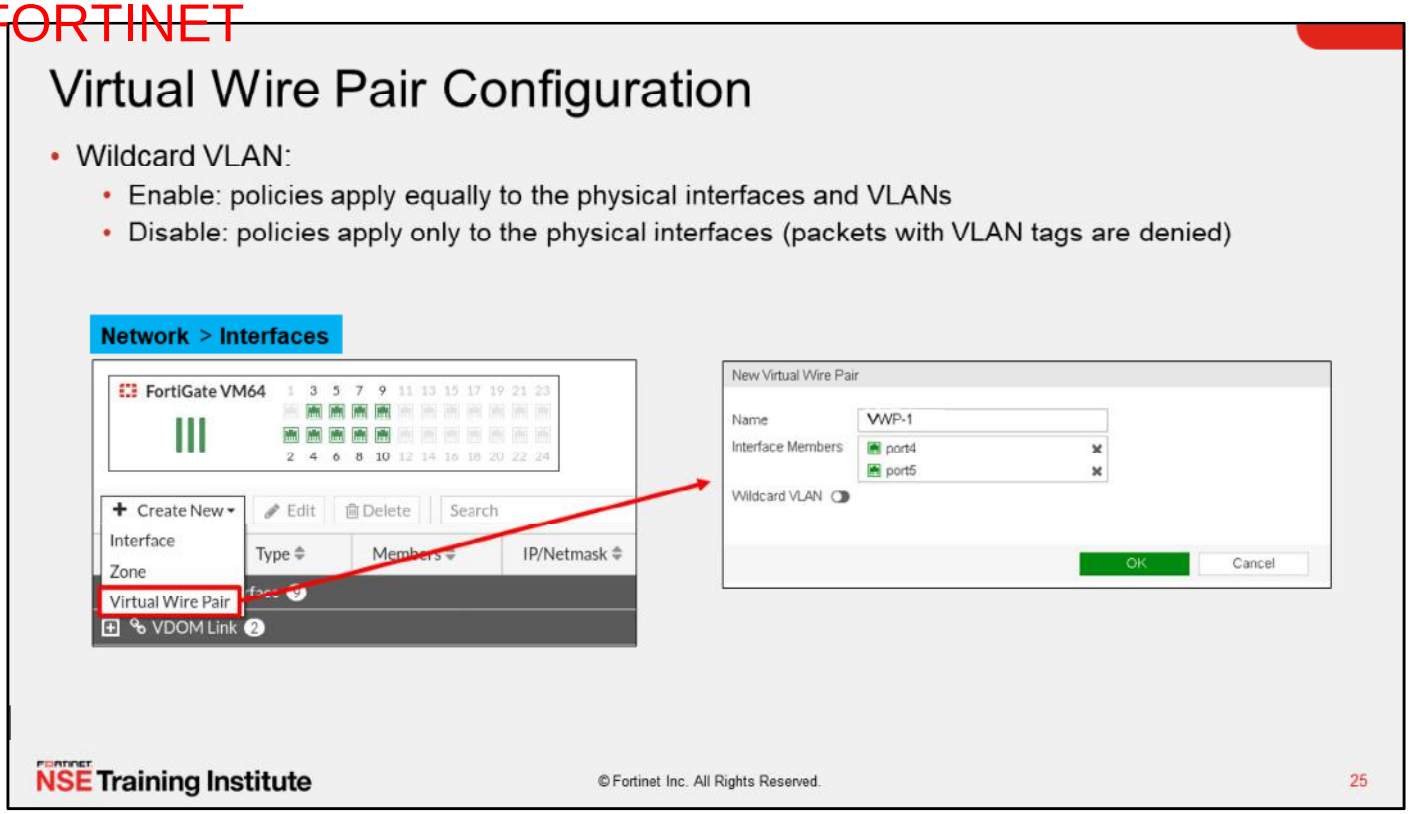

When you create a virtual wire pair, you must select two physical interfaces—no more, no less.

After selecting the two interfaces, you create the virtual wire pair policies that inspect the traffic crossing the virtual wire pair. The **Wildcard VLAN** setting specifies how those policies are applied to the different VLANs whose traffic flows between the pair:

- If you enable **Wildcard VLAN**, the virtual wire pair policies are applied equally to the physical interfaces and VLANs.
- If you disable **Wildcard VLAN**, the virtual wire pair policies are applied only to the physical interfaces. Traffic with any VLAN tag is denied.

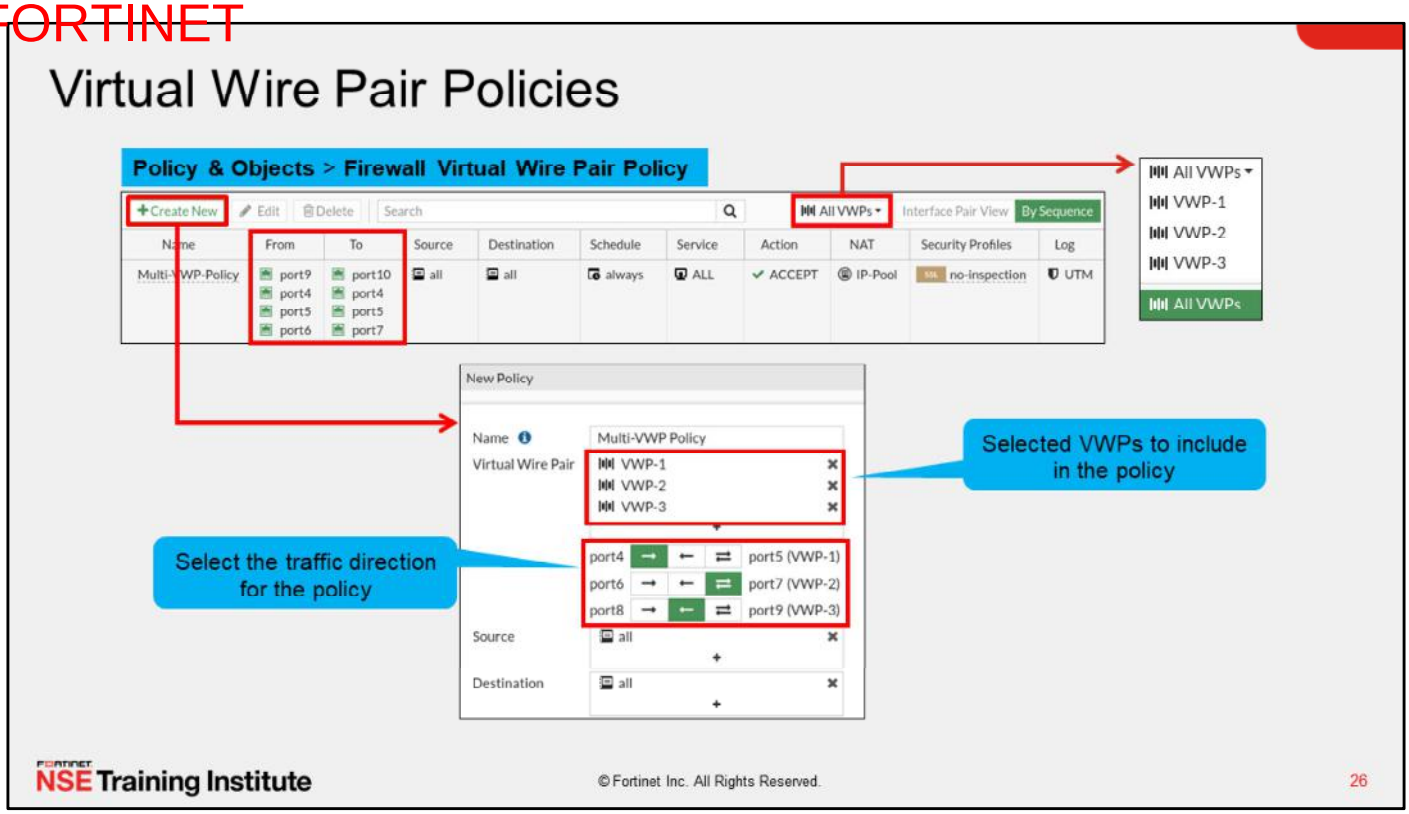

For profile-based configurations, you configure the firewall policies for virtual wire pairings under the **Policy & Objects** > **Firewall Virtual Wire Pair Policy** menu, which is displayed when at least one virtual wire pair has been crated.

Firewall virtual wire pair policies can include more than a single virtual wire pair. This capability can streamline the policy management process by eliminating the need to create multiple, similar policies for each virtual wire pair. When creating or modifying a policy, you can select the traffic direction for each VWP included in the policy.

You can use the drop-down list at the top of the view to filter the view to display only policies associated with specific VWPs. Policies that include multiple VWPs appear in the list of policies for all included VWPs.

For policy-based configurations, both the **Security Virtual Wire Pair Policy** and **Virtual Wire Pair SSL Inspection & Authentication** pages include the VWP drop-down list at the top of the view.

Knowledge Check

1. Which configuration setting must be enabled to allow VLAN-tagged traffic through a virtual wire pair?
   - A. Transparent bridging
   - ✓B. Wildcard VLAN

2. How is traffic handled in a virtual wire pair?
   - ✓A. Incoming traffic to one interface is always forwarded out through the other interface.
   - B. Traffic is forwarded based on the destination MAC address.

NSE Training Institute

27

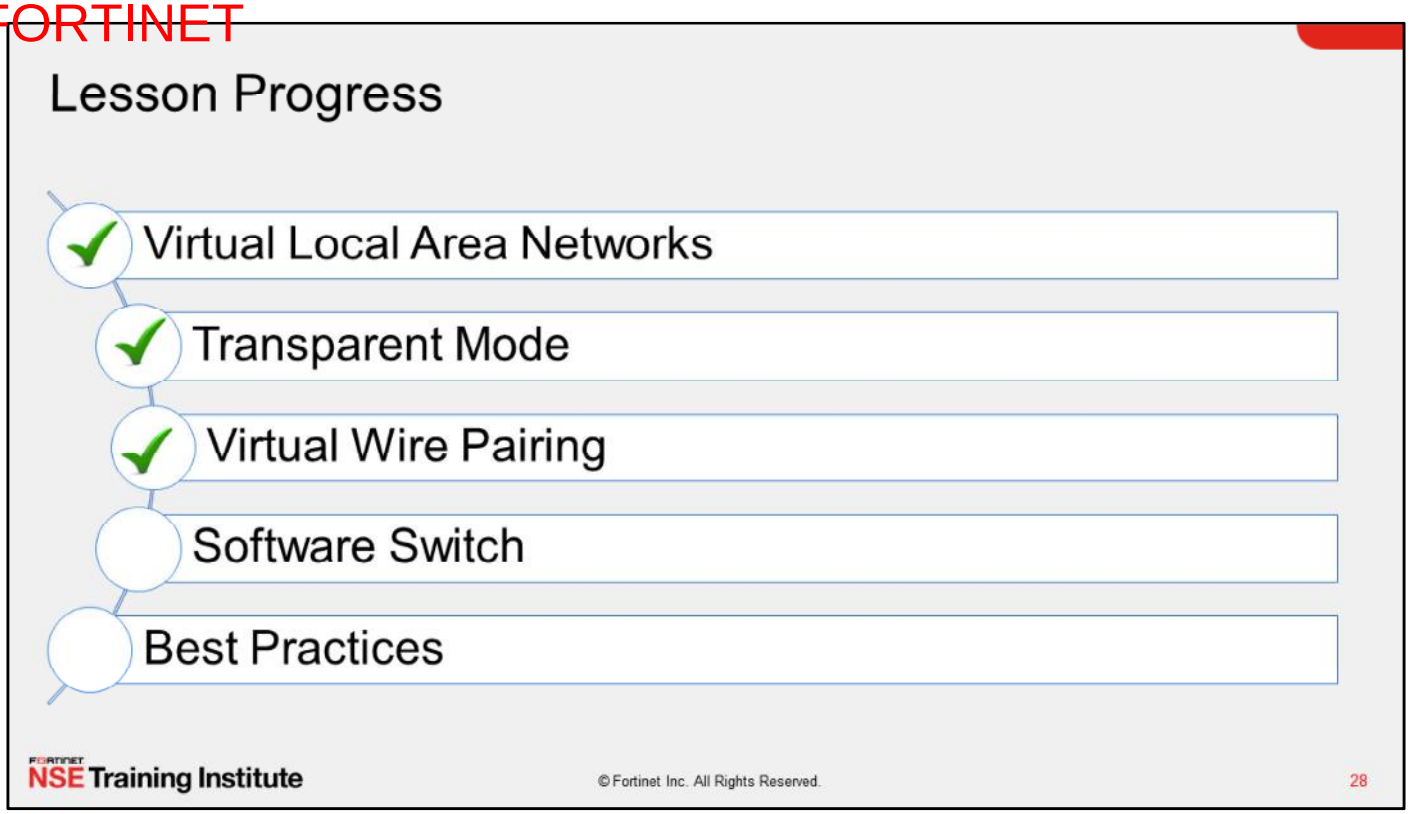

Good job! You now understand virtual wire pairing.

Now, you will learn about the software switch function.

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

By demonstrating competence in using the software switch feature, you will understand how to group multiple physical and wireless interfaces into a single virtual interface.

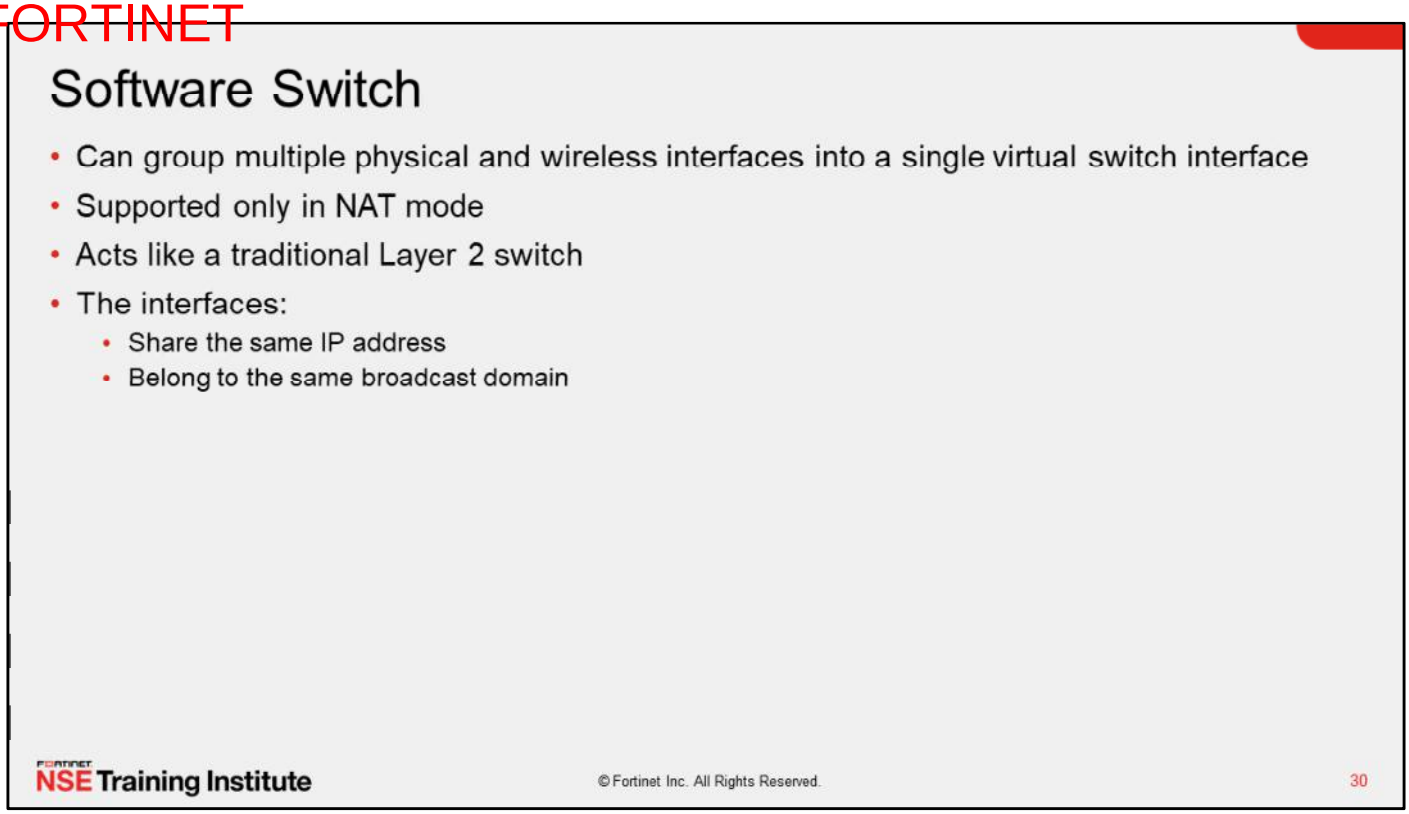

A software switch groups multiple interfaces to form a virtual switch, which acts as a traditional Layer 2 switch. This means that all switch interfaces are part of the same broadcast domain.

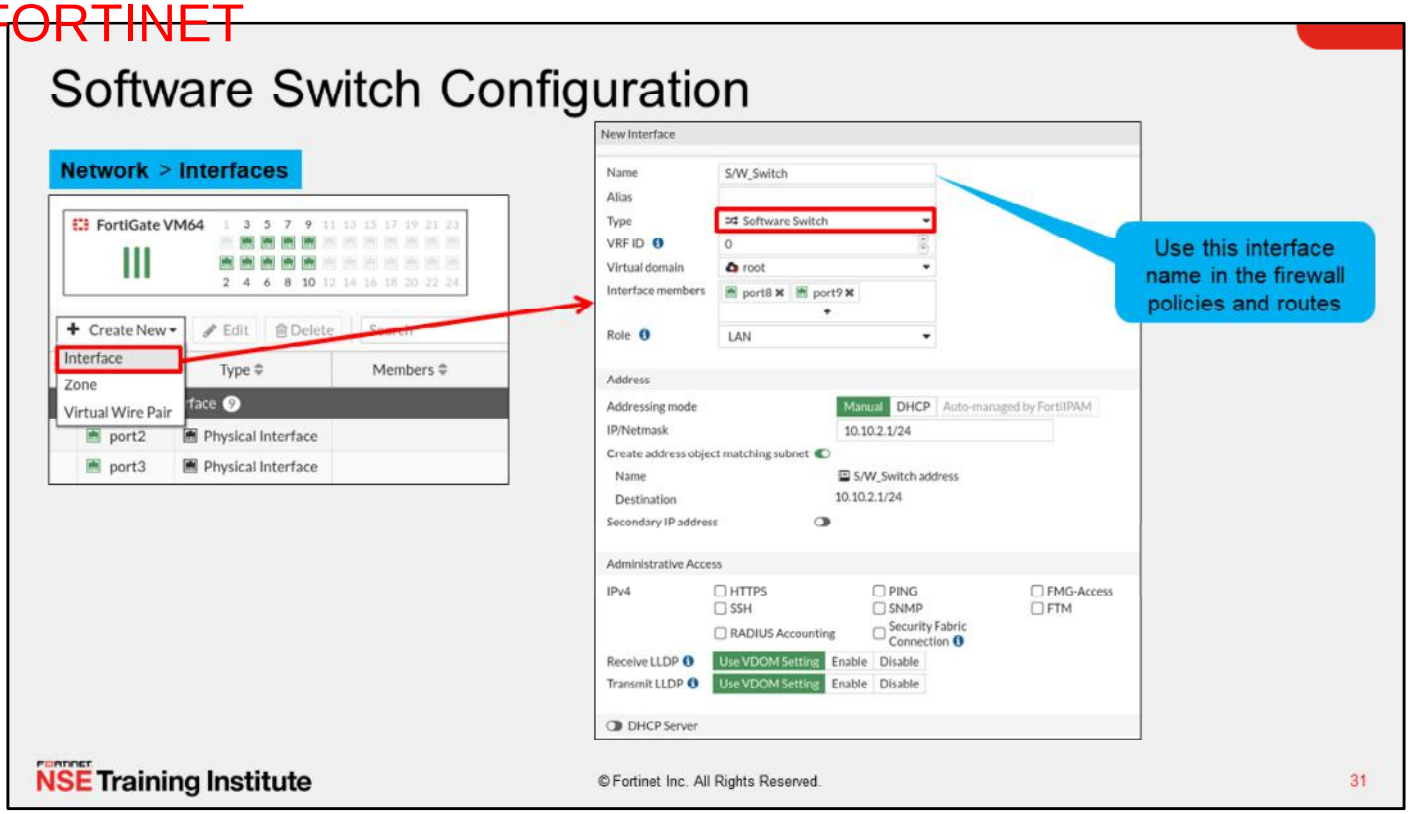

Each software switch has a virtual interface associated with it. Its IP address is shared by all the physical switch interfaces and member SSIDs. You use this virtual interface in the firewall policies and routing configuration.

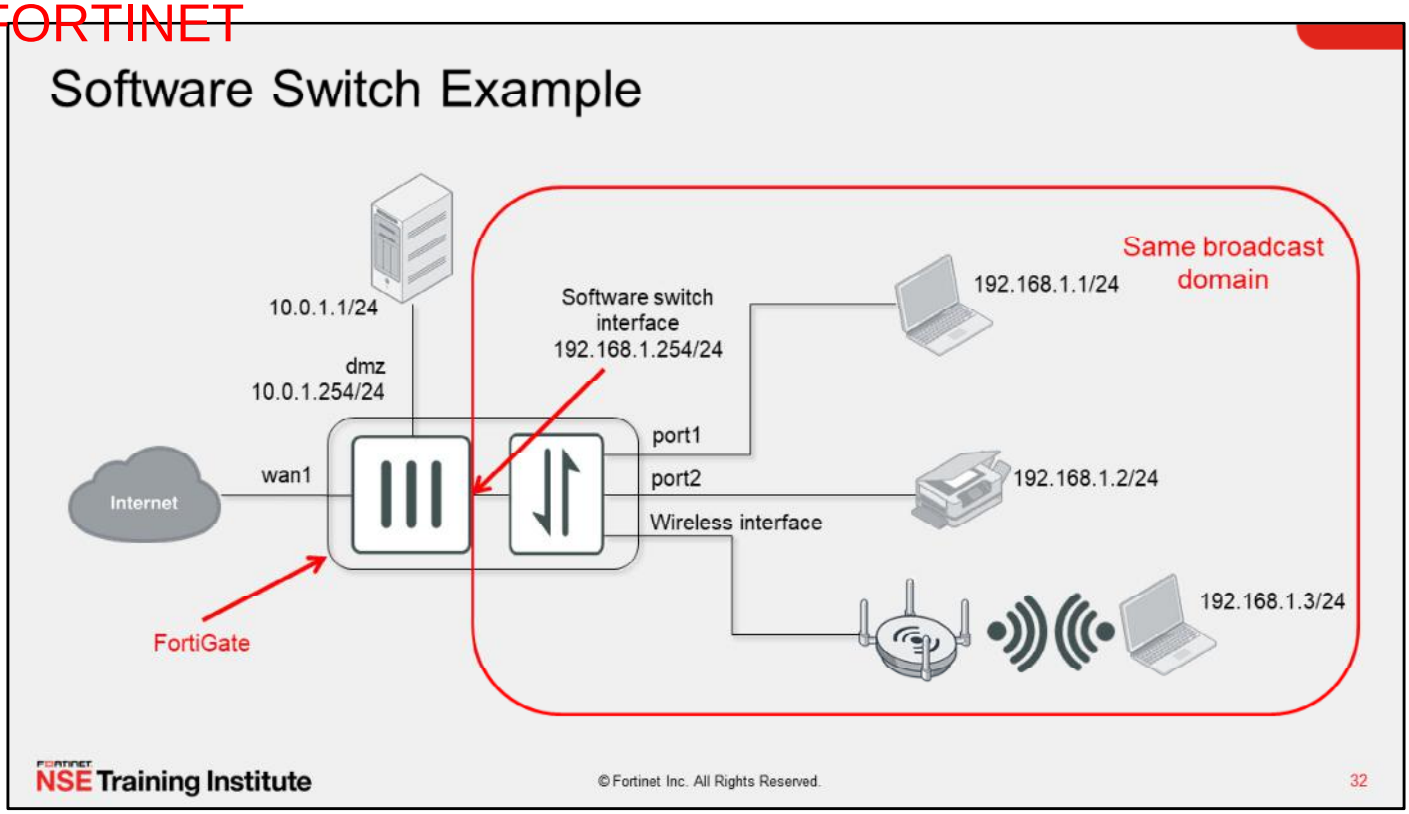

In the example shown on this slide, the administrator grouped a wireless interface with **port1** and **port2** to form a software switch. These three interfaces are part of the same broadcast domain. All the devices connected to the switch interfaces belong to the same IP subnet: `192.168.1.0/24`. This allows FortiGate to forward broadcast traffic from the wireless clients to **port1** and **port2**.

The software switch interface itself has an IP address, which is also in the same subnet: `192.168.1.0/24`. This is the default gateway IP address for all the devices connected to the software switch.

The server `10.0.1.1` is connected to an interface (**dmz**) that is not part of the software switch. So, it belongs to a different broadcast domain and IP subnet.

Knowledge Check

1. In which operating mode is the software switch function supported?
   A. Transparent mode
   ✓B. NAT mode

2. Which interface can be a member of a software switch?
   A. VLAN interface
   ✓B. Wireless interface

NSE Training Institute

33

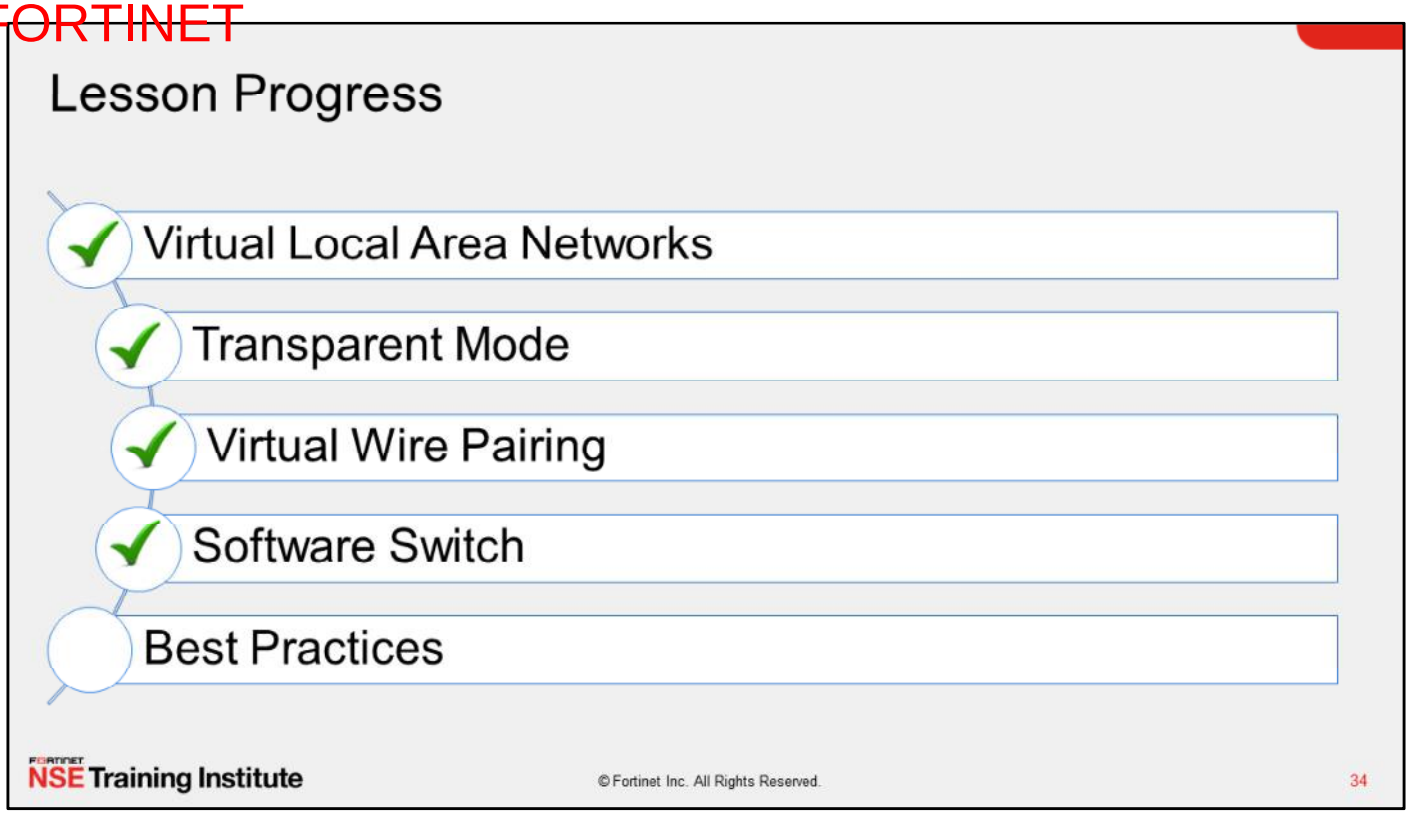

Good job! You now understand the software switch function.

Now, you will learn about spanning tree protocol (STP).

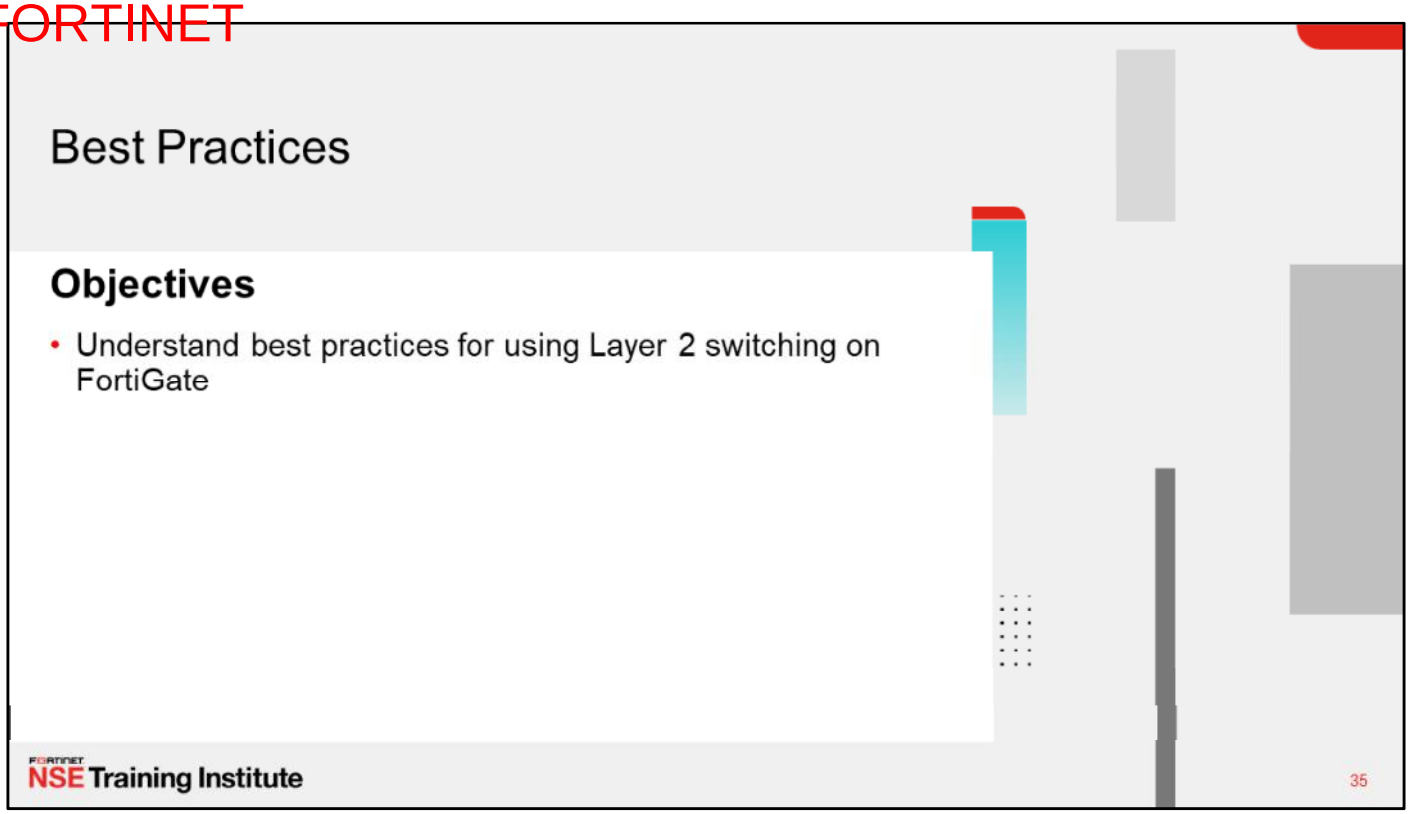

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in using best practices for Layer 2 switching on FortiGate, you will understand how to identify, resolve, and prevent common Layer 2 issues.

Best Practices

- Create forwarding domains when VLANs are used and set `vlanforward` to `disable` on all relevant physical interfaces
- The forward-domain ID can be different from the VLAN ID, but it is recommended for troubleshooting and readability to keep them the same
- When using forwarding domains, a router is required to move traffic between the forwarding domains
- Only interfaces from the same forwarding domains can have firewall policies between each other
- Because STP BPDUs are not forwarded by default, use caution when inserting FortiGate (or any other forwarding device) because this could break the spanning tree and lead to Layer 2 loops

NSE Training Institute

36

When you implement a FortiGate device that has Layer 2 features, this slide shows some best practices you should follow to avoid preventable issues within your network.

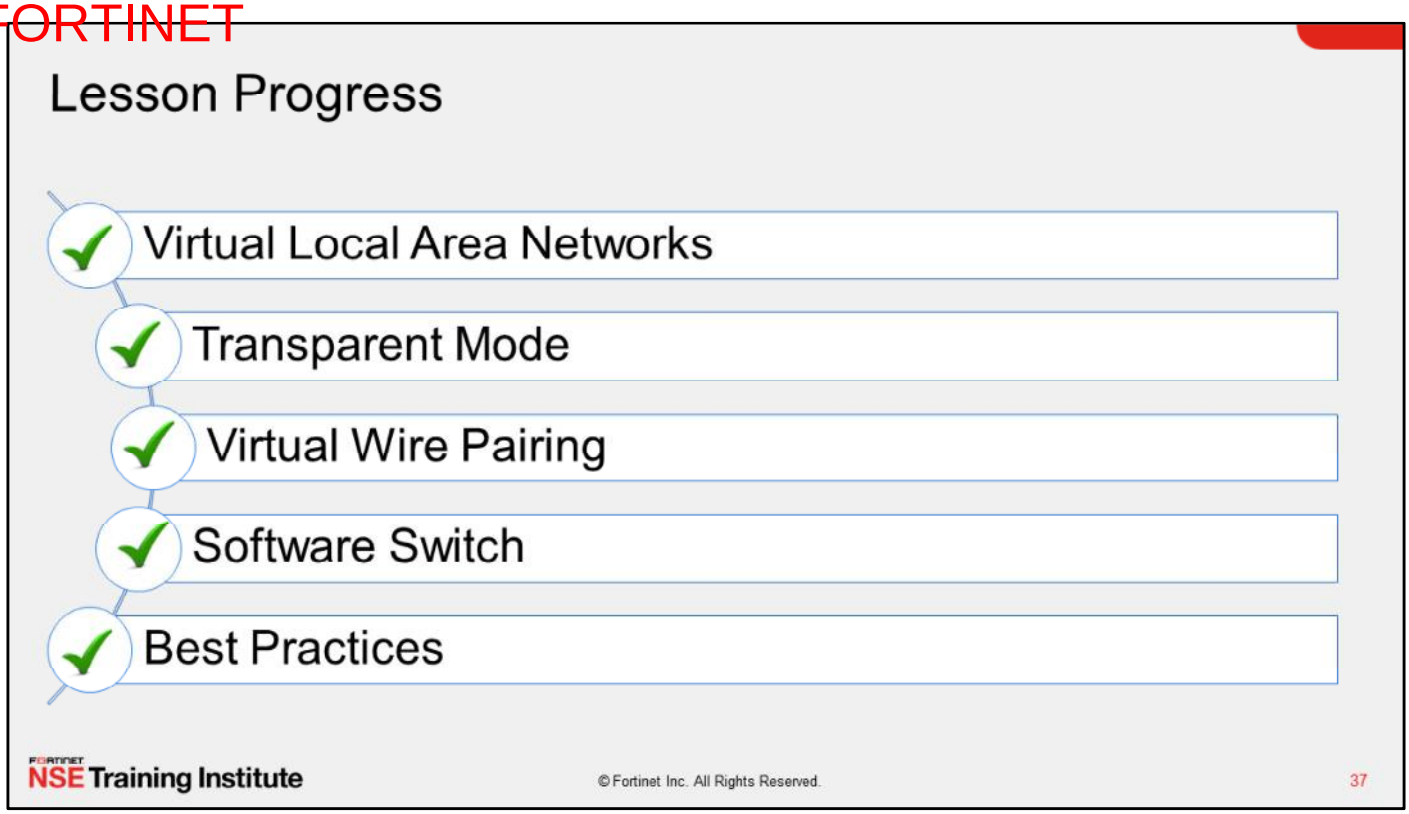

Congratulations! You have completed this lesson.

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

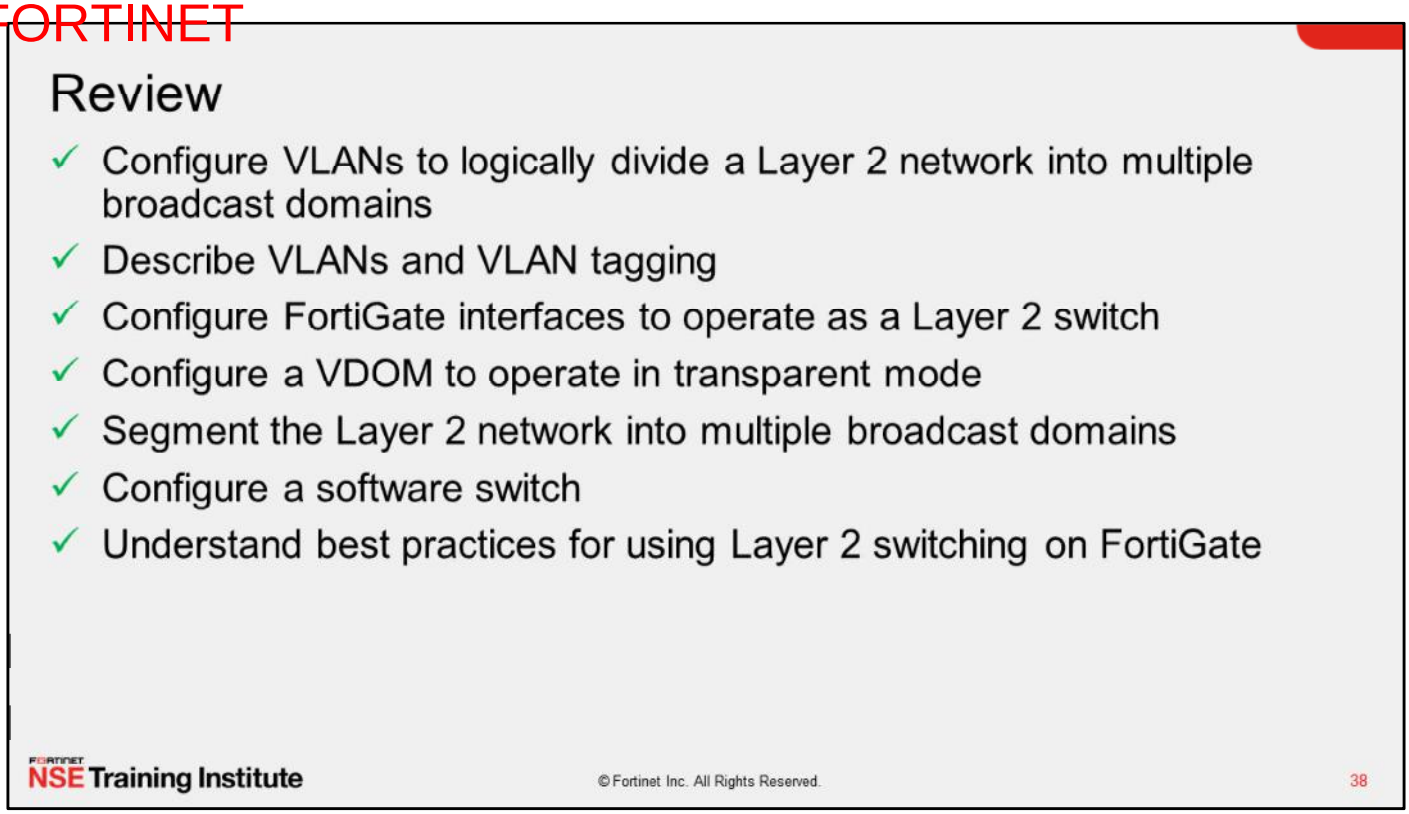

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use transparent operation mode and Layer 2 switching on FortiGate.

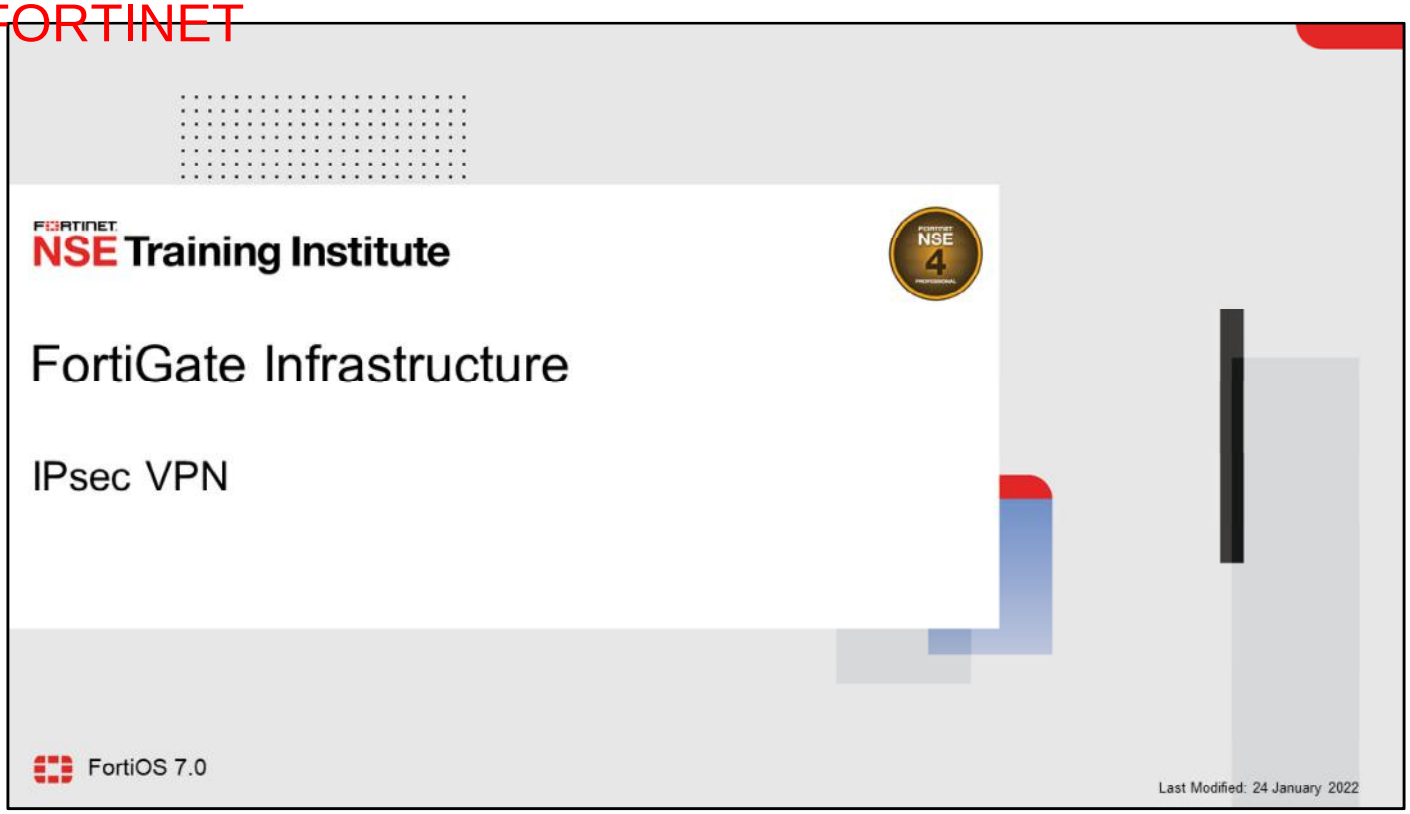

In this lesson, you will learn about the architectural components of IPsec VPN and how to configure them.

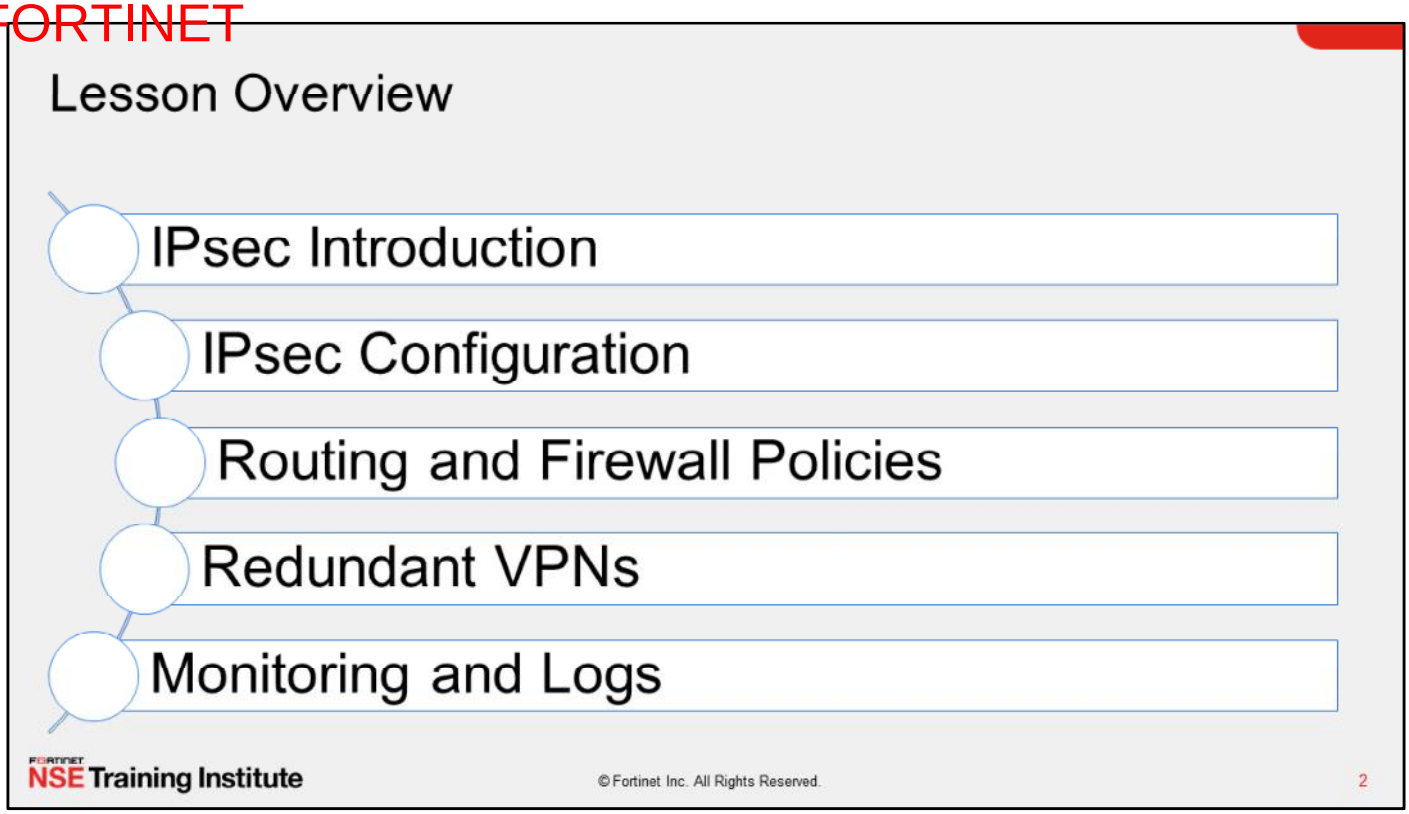

In this lesson, you will learn about the topics shown on this slide.

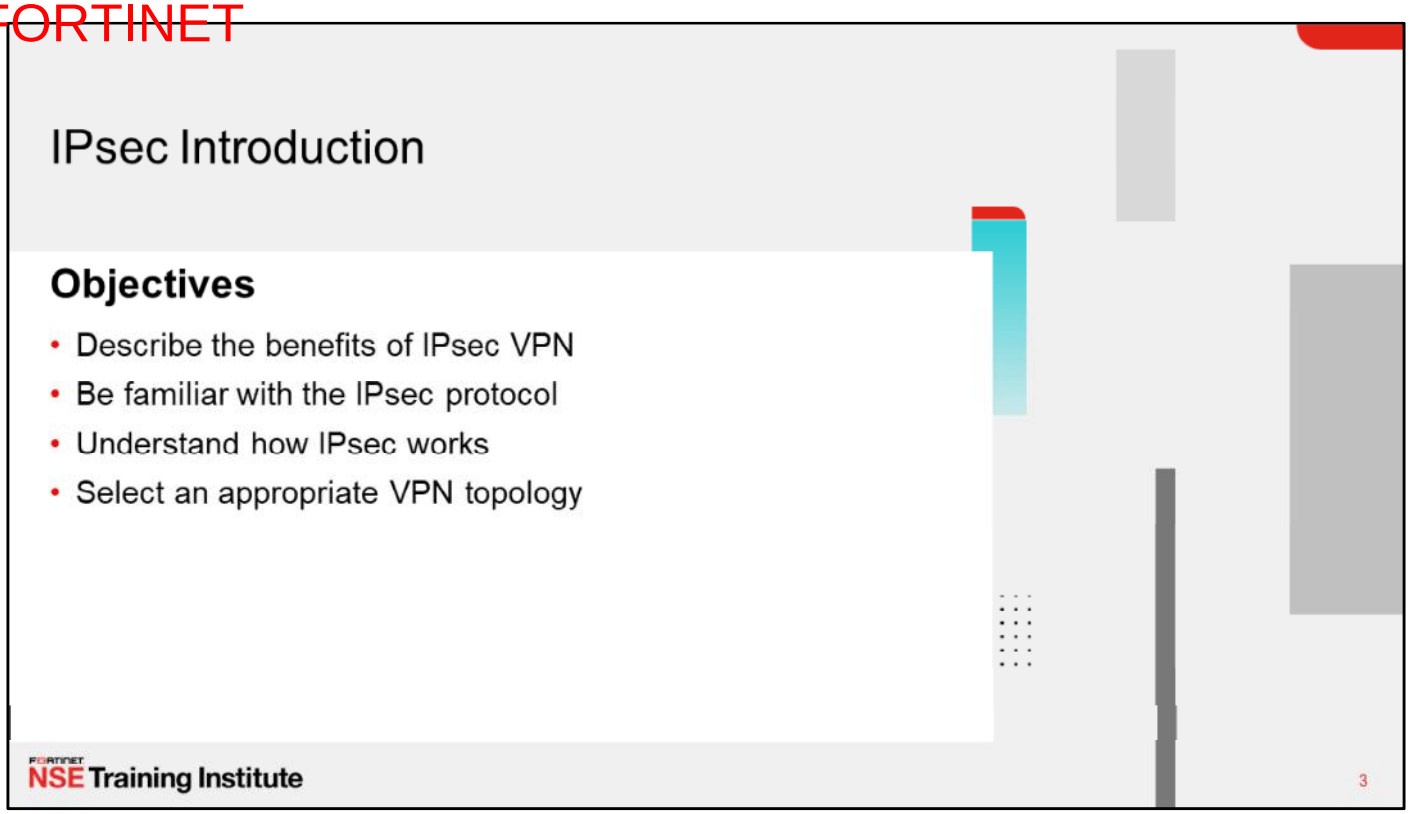

After completing this lesson, you should be able to achieve the objectives shown on this slide.

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

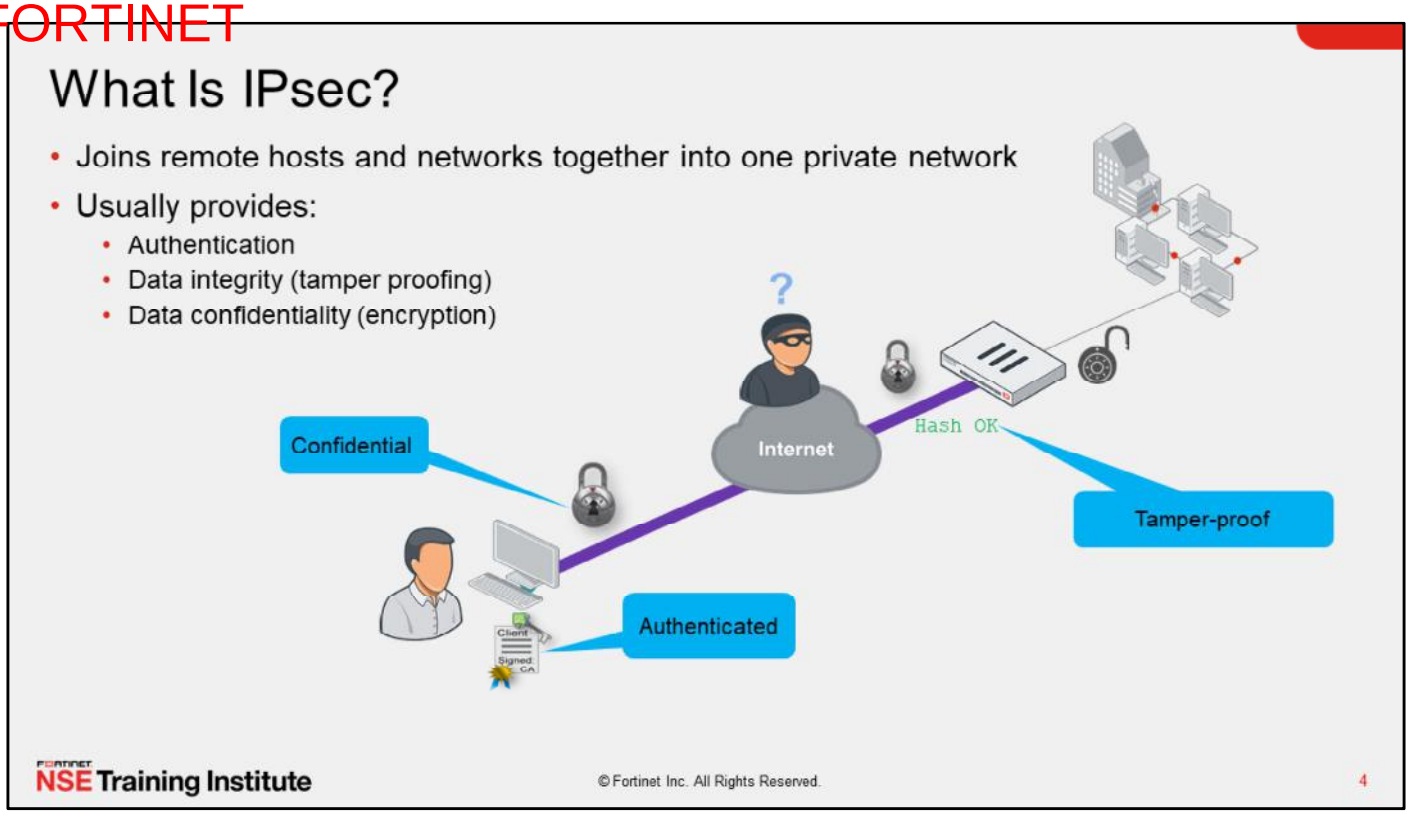

What is IPsec? When should you use it?

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

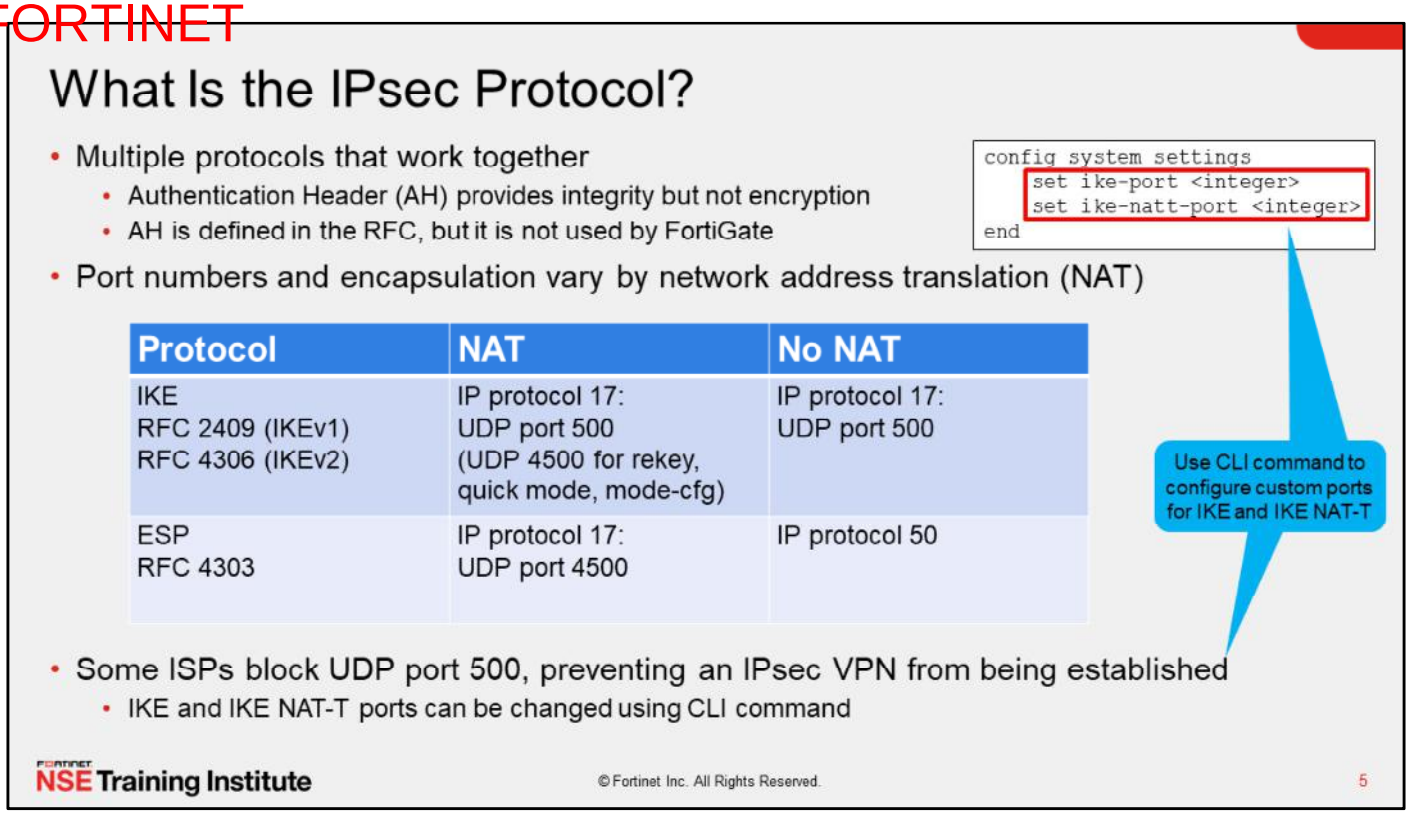

If you're passing your VPN through firewalls, it helps to know which protocols to allow.

IPsec is a suite of separate protocols. It includes:
- Internet Key Exchange (IKE): IKE is 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.
- Encapsulation Security Payload (ESP): ESP is the encapsulated security payload—the encrypted payload, essentially, the *data channel.*

So, if you need to pass IPsec traffic through a firewall, remember: allowing just one protocol or port number is usually not enough.

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

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

Some ISPs block UDP port 500, preventing an IPsec VPN from being established. You can use the CLI command shown on the slide to configure custom ports for IKE and IKE NAT-T.

The default UDP port for IKE traffic is 500, but you can select a custom port from port range 1024 to 65535. The default port for IKE traffic in NAT-T mode is 4500, but you can change it to a custom port from port range 1024 to 65535.

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 like SSL/TLS
  - Authentication
  - Handshake to exchange keys, settings

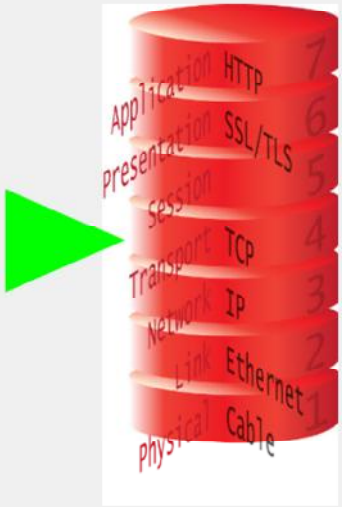

NSE Training Institute

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.

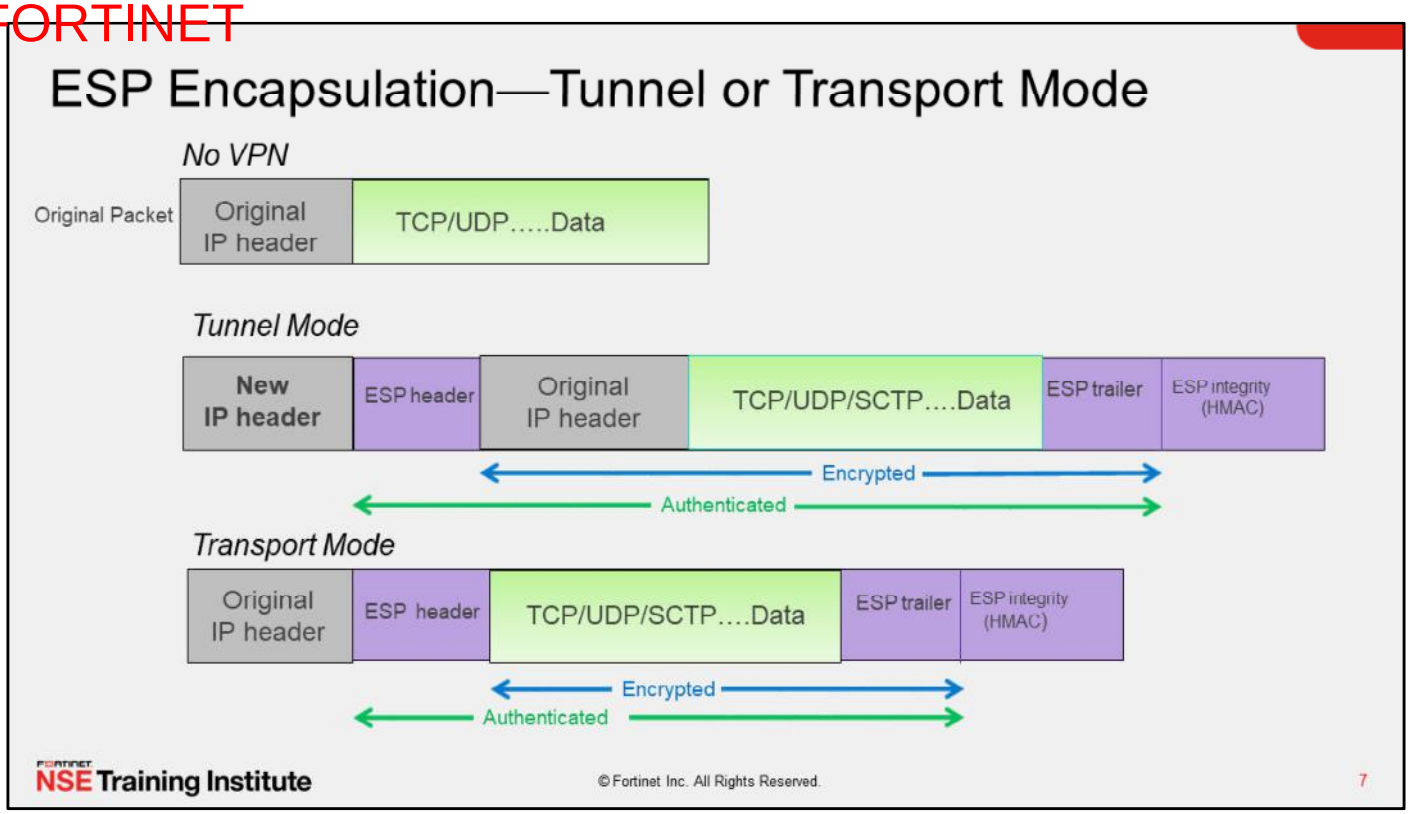

What's encapsulated? It depends on the encapsulation mode being used. IPsec can operate in two modes: transport mode and tunnel mode.

- 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.

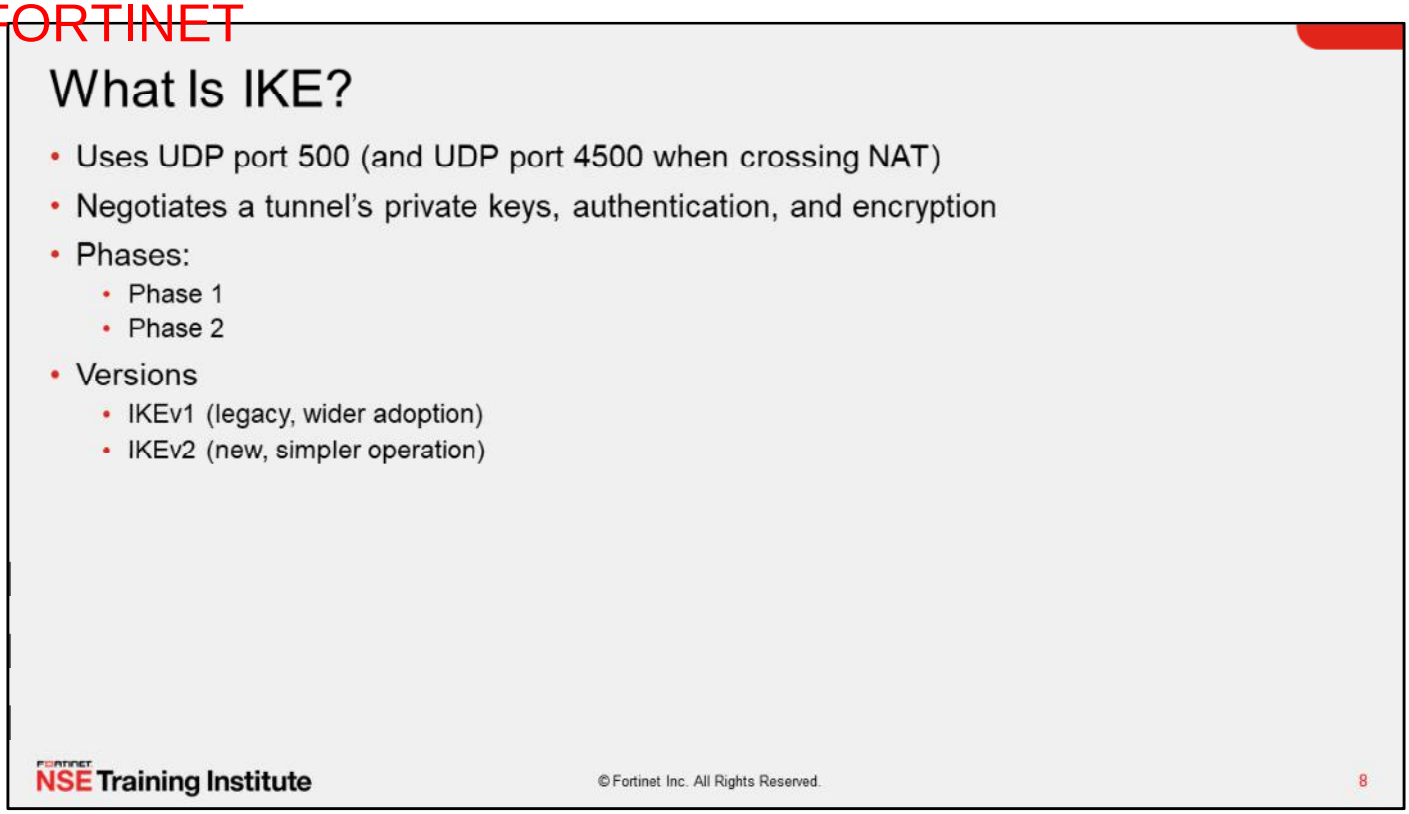

IKE uses UDP port 500. If NAT-T is enabled in a NAT scenario, IKE uses UDP port 4500.

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.

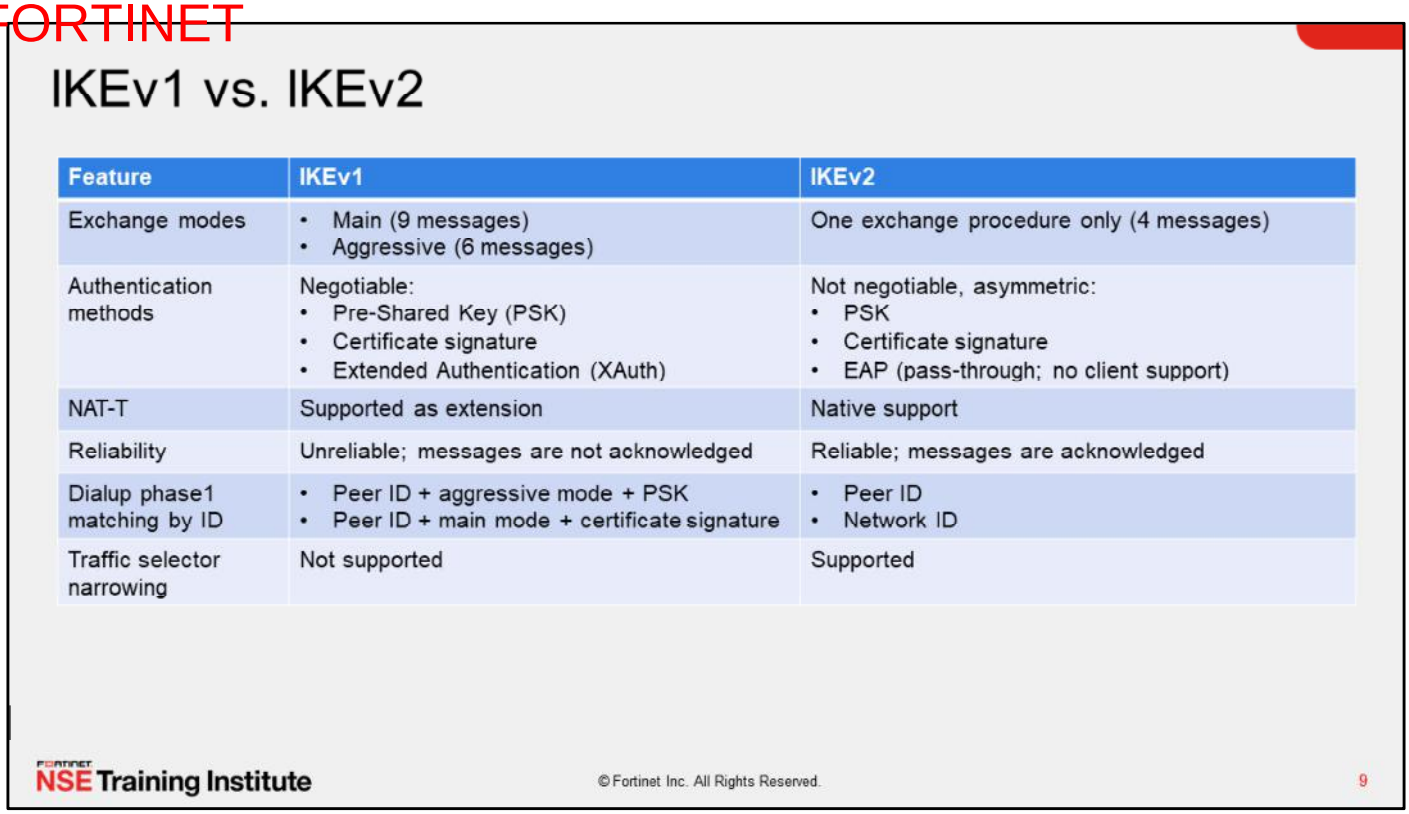

This slide shows a table comparing some of the IKEv1 and IKEv2 features supported by FortiOS. IKEv2 provides a simpler operation, which is the result of using a single exchange mode and requiring less messages to bring up the tunnel.

In terms of authentication methods, IKEv1 enables you to deny access to VPN peers without having to use certificate signature. This is possible by using IKEv1 and XAuth. With XAuth, you deny access to a peer by removing or disabling the peer username from the backend server, and then flushing the tunnel. The equivalent of XAuth on IKEv2 is EAP. However, the IKEv2 EAP implementation in FortiOS is pass-through only. That is, FortiOS doesn't support EAP as client, which means that you cannot revoke access to peers using IKEv2 unless you use certificate signature.

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

When you configure multiple dialup 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 dialup server regardless of the authentication mode in use. However, with IKEv1, you can use 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 proposed by the initiator. 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 proposed by the remote peer.

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.

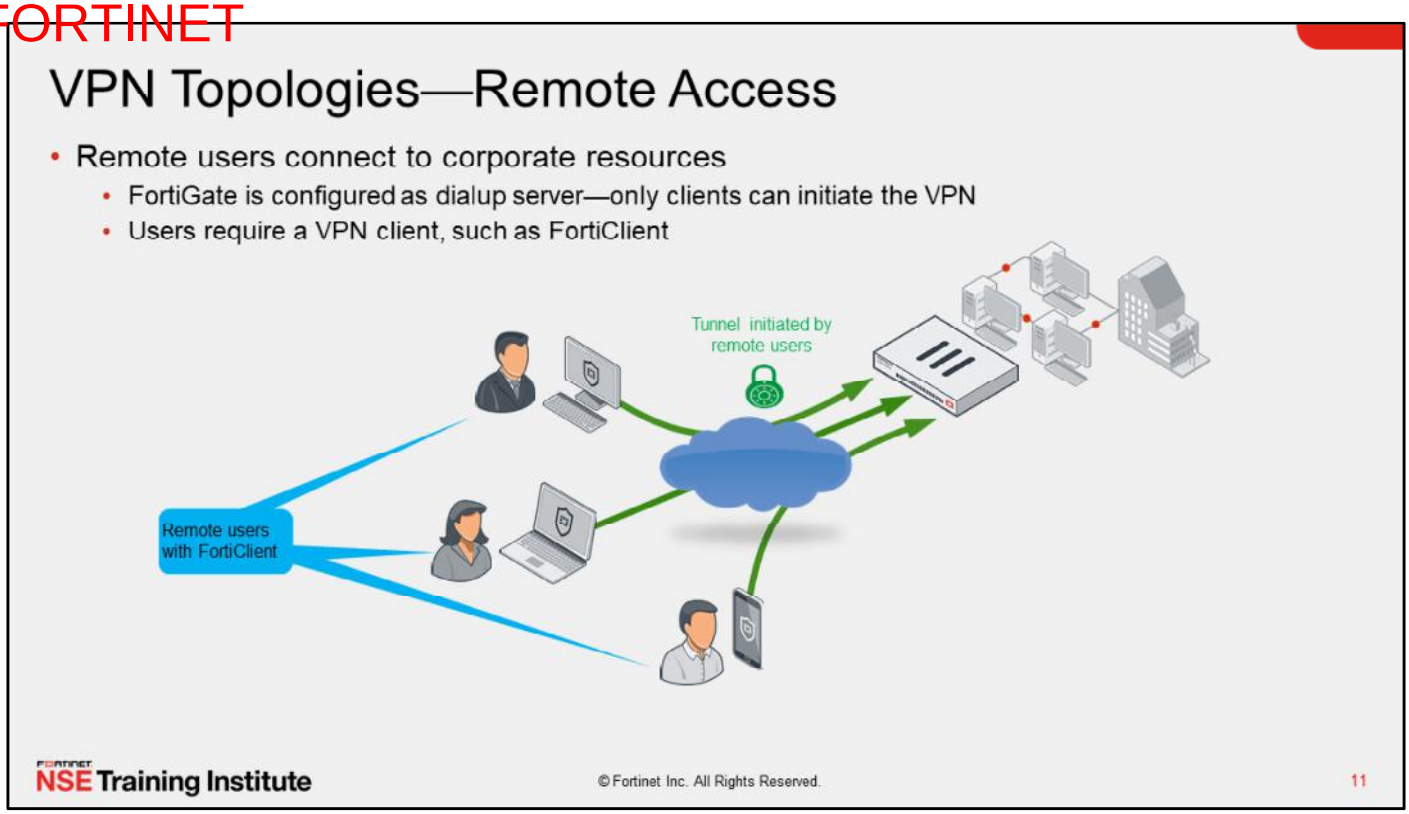

Use remote access VPNs when remote internet users need to securely connect to the office to access corporate resources. The remote user connects to a VPN server located on the corporate premises, such as FortiGate, to establish a secure tunnel. After the user is authenticated, FortiGate provides access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

In a remote access VPN, FortiGate is usually configured as a dialup server. You will learn more about dialup 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.

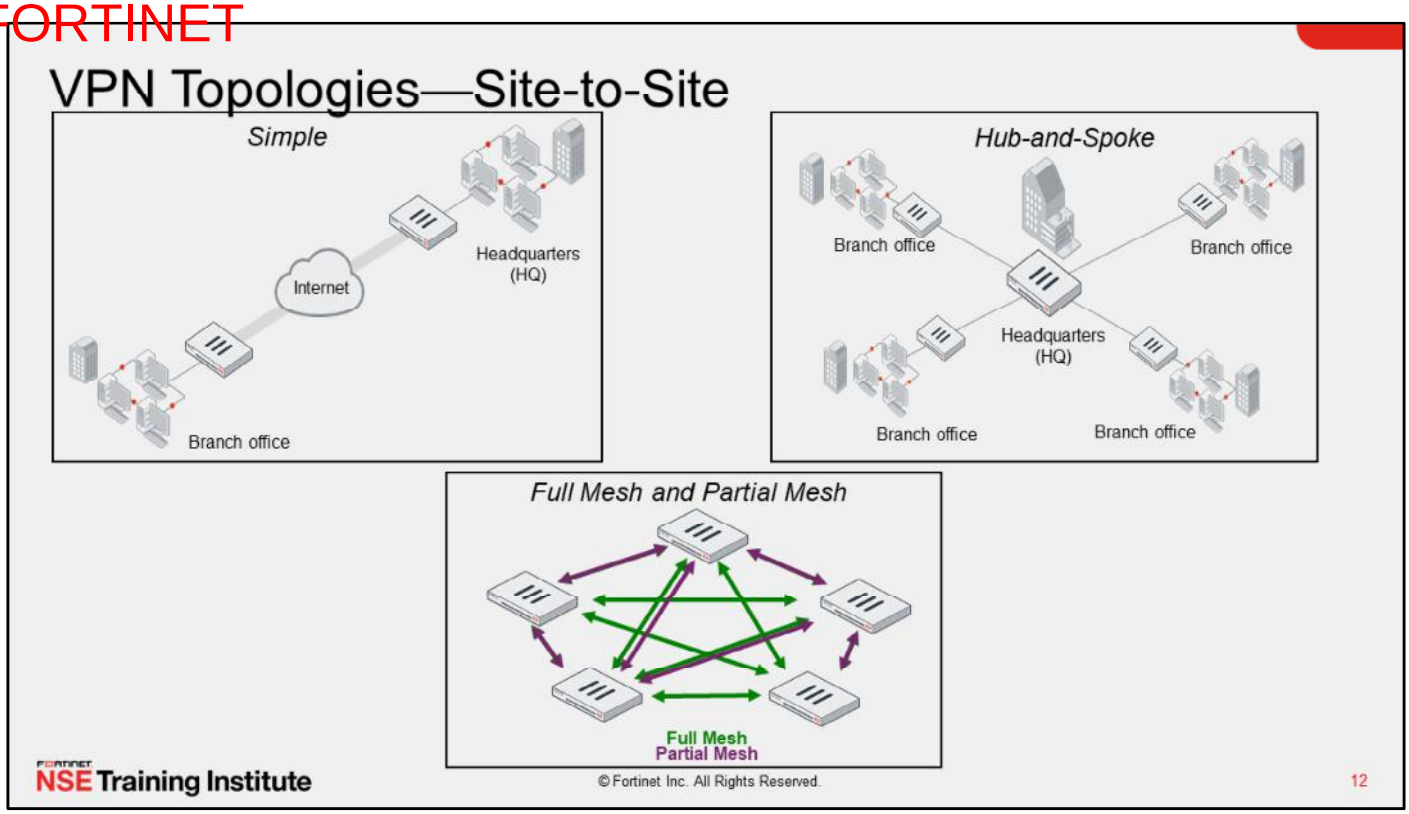

Site-to-site VPN is also known as LAN-to-LAN VPN. A *simple* site-to-site deployment involves two peers communicating directly to connect two networks located at different offices.

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

NSE 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 (ADVPN)

- 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
  - Requires the use of routing protocol for spokes to learn the routes to other spokes

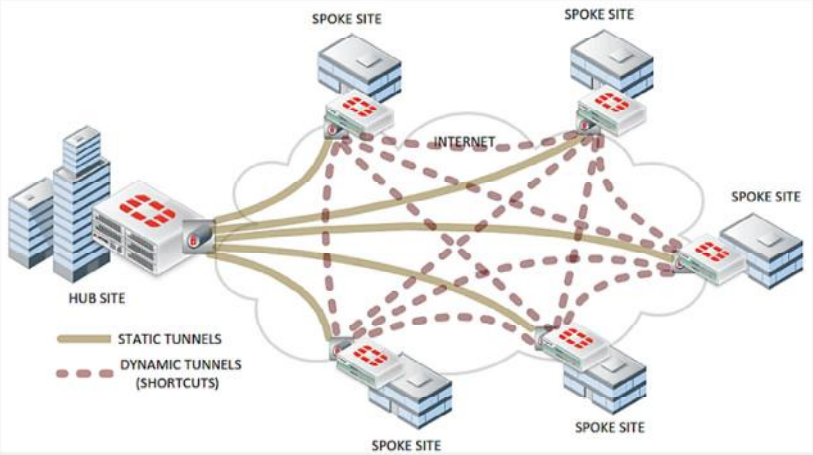

NSE Training Institute

14

Each VPN topology has its advantages and disadvantages.

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.

ADVPN requires a dynamic routing protocol running over the IPsec tunnels so that spokes can learn the routes to other spokes, after the dynamic VPNs negotiate.

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

NSE Training Institute

15

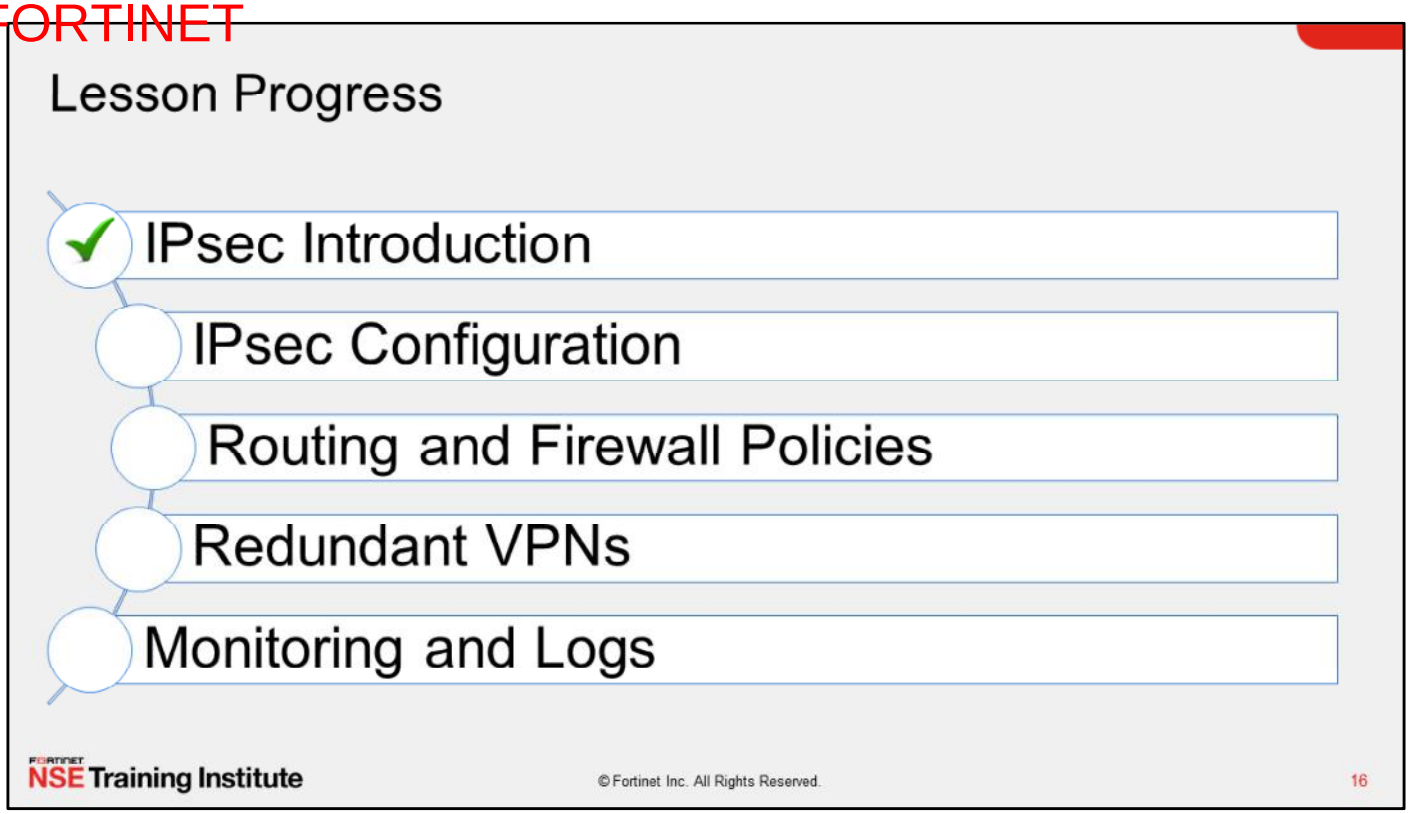

Good job! You have now been introduced to IPsec.

Now, you will learn about IPsec configuration.

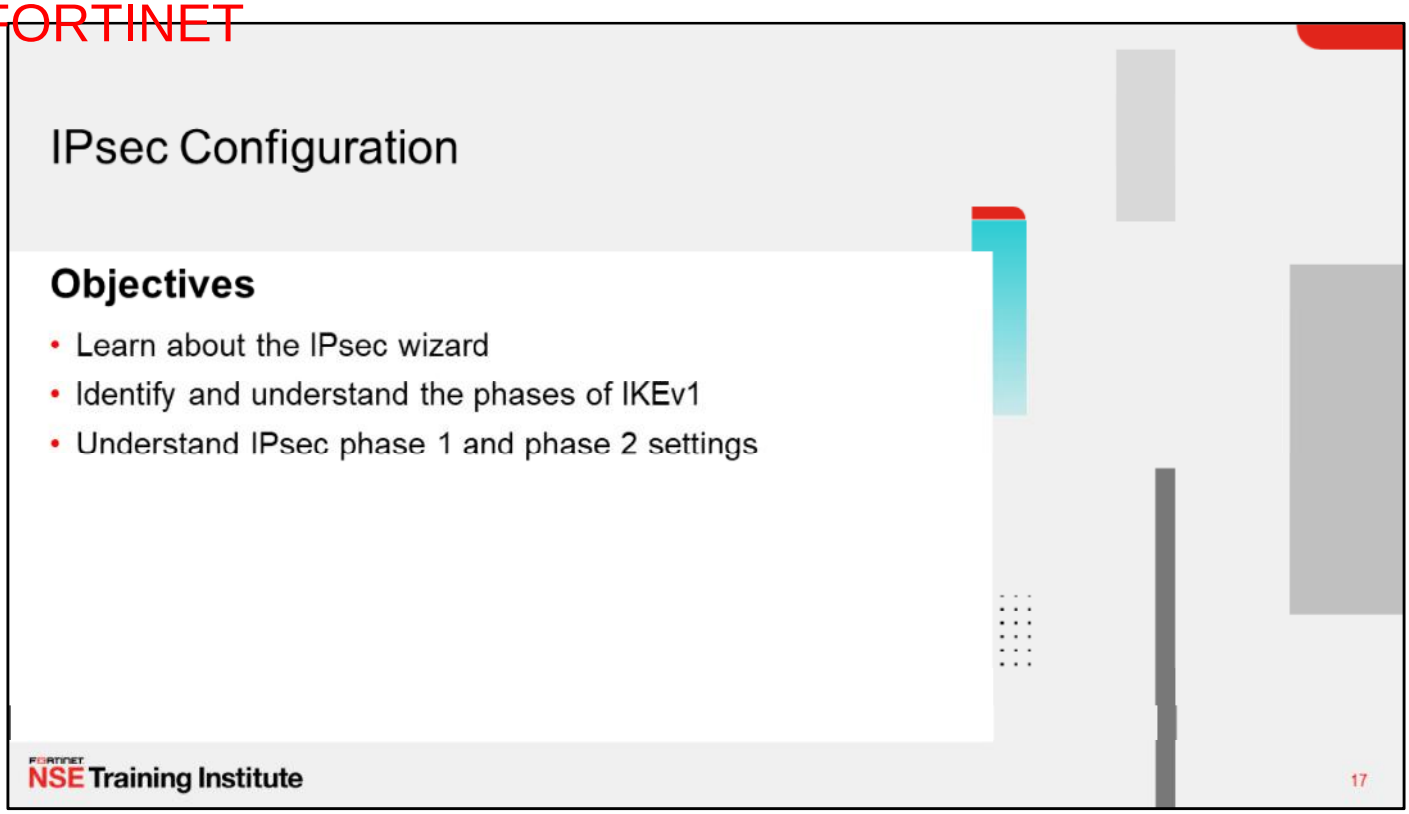

After completing this section, you should be able to achieve the objectives shown on this slide.

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

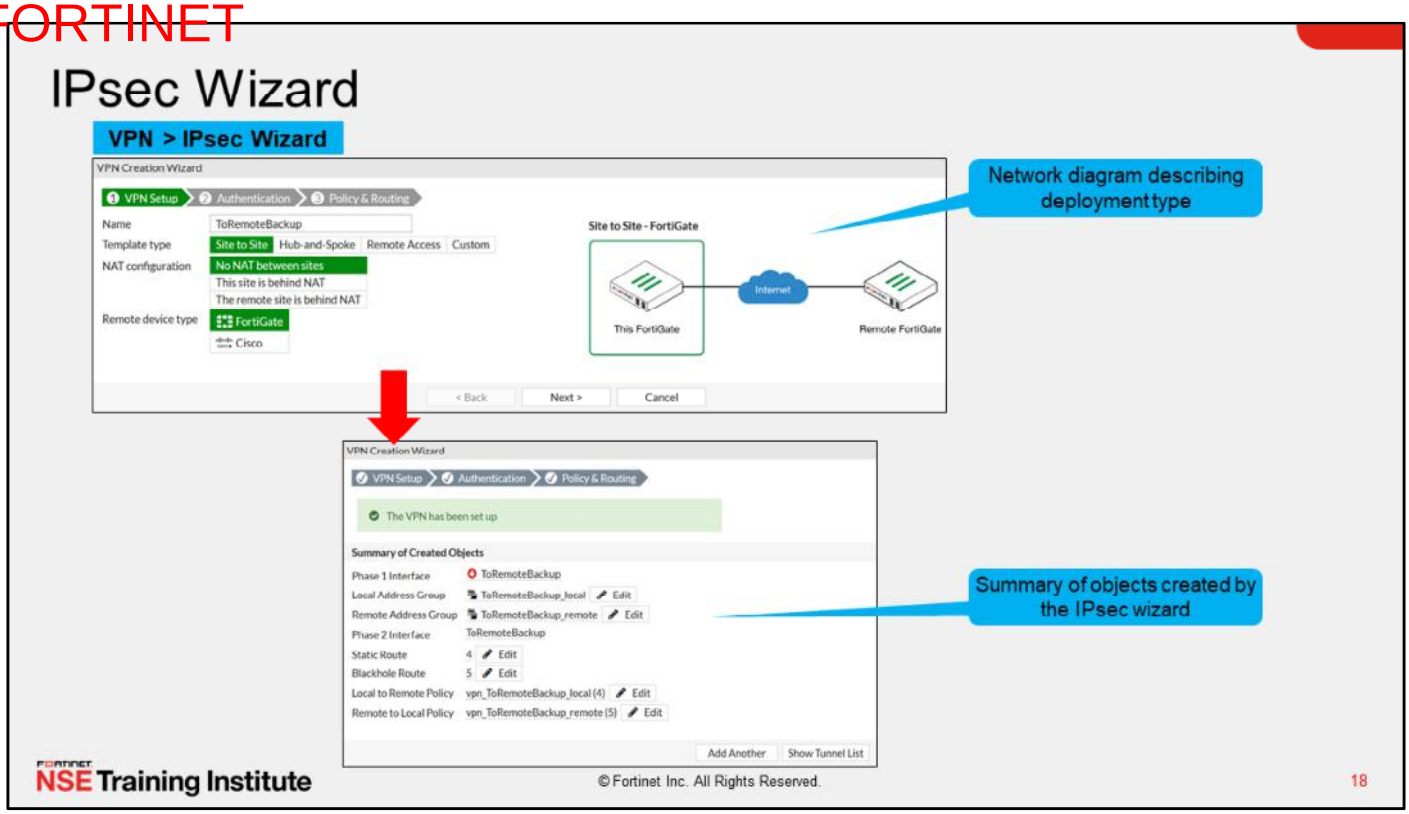

When you create an IPsec tunnel on the GUI, FortiGate redirects you to the **IPsec Wizard**. The wizard simplifies the creation of the new VPN by walking you through a three to four-step process. The first step is to select a template type. If you want to manually configure your VPN, you can select **Custom** 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.

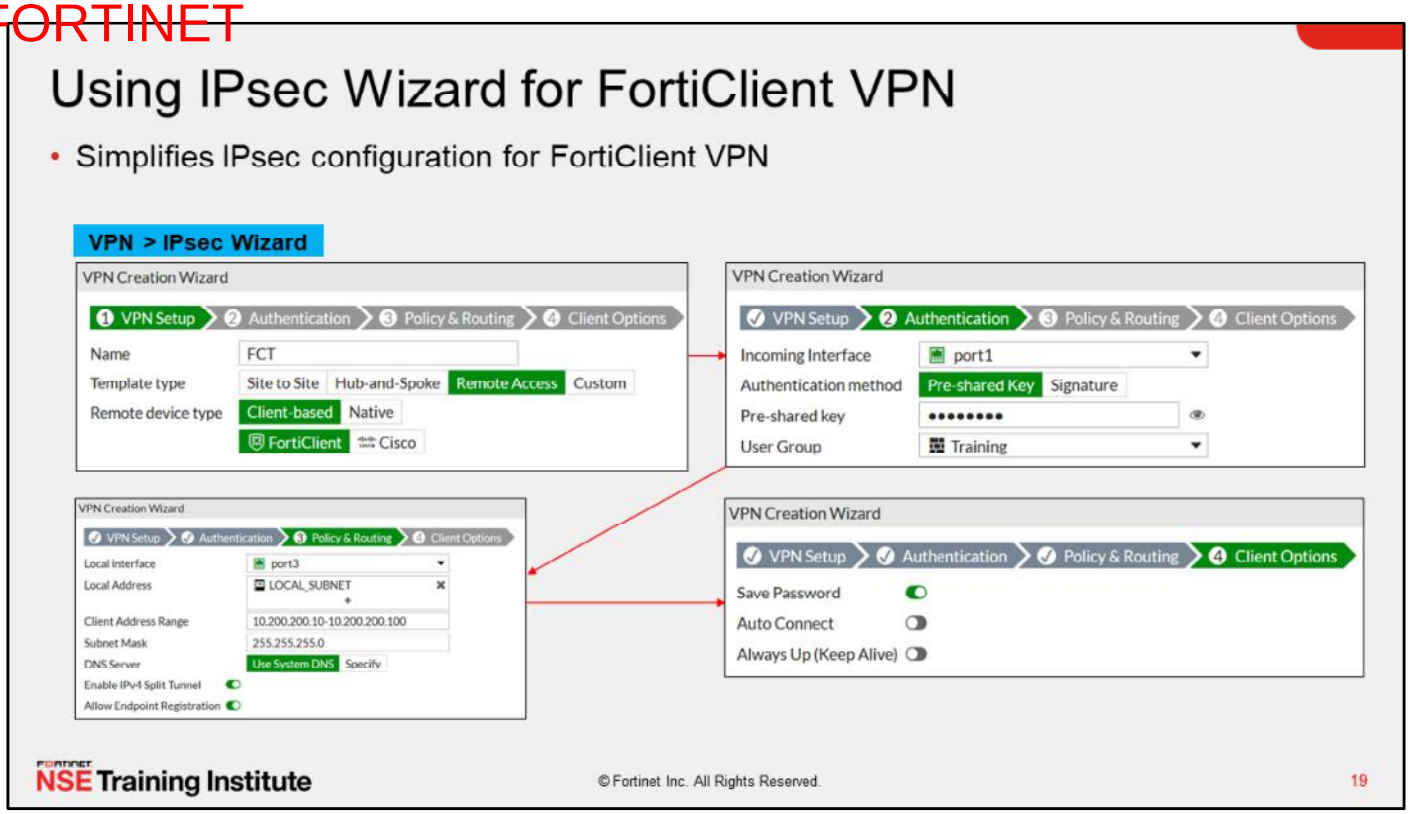

A common use of the IPsec wizard is for configuring a remote access VPN for FortiClient users. The wizard enables IKE mode config, XAuth, and other appropriate settings for FortiClient users. You will learn more about IKE mode config and XAuth in this lesson.

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

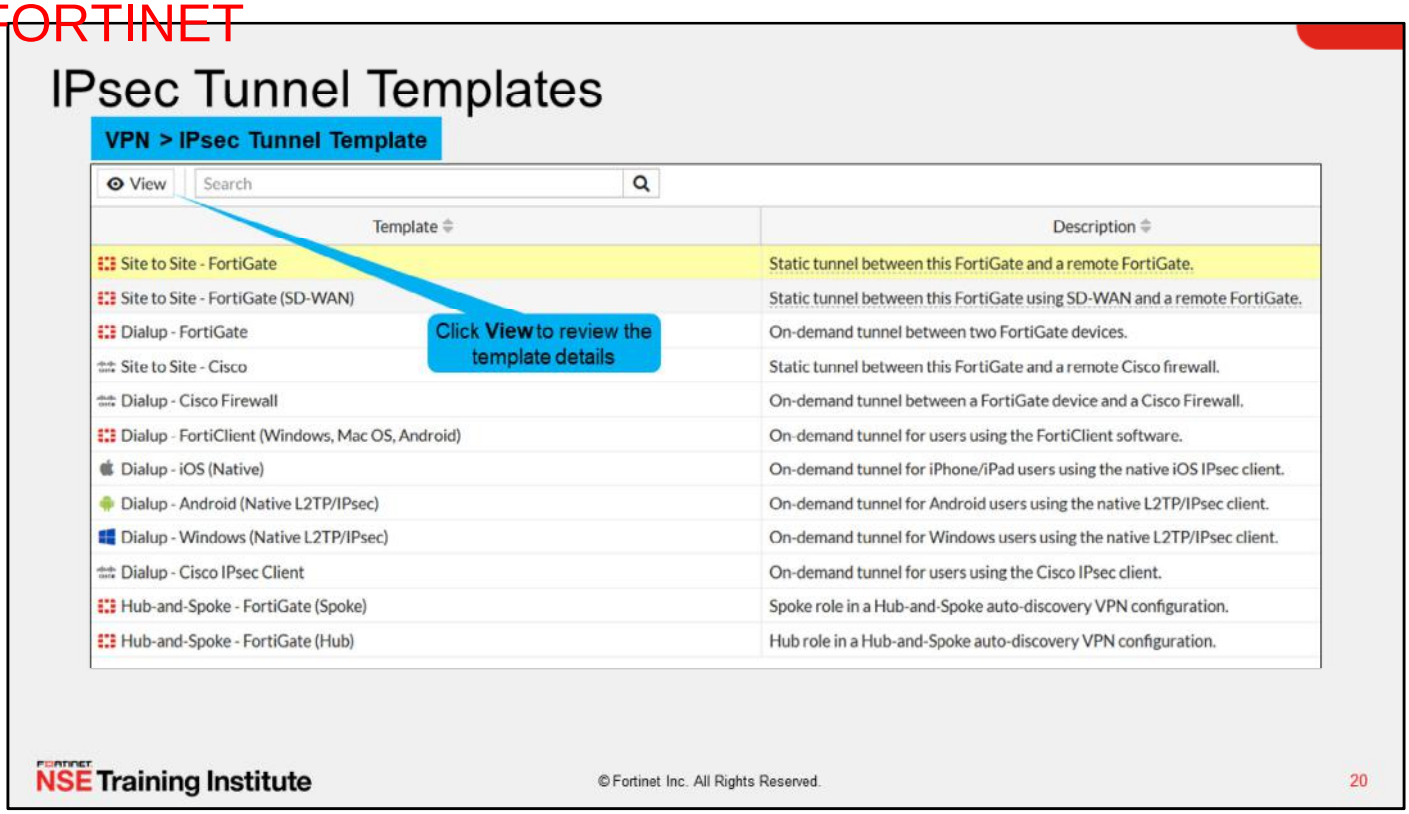

The IPsec wizard uses one of the templates shown on this slide when applying the configuration for the new IPsec tunnel. You can review the settings of a template by selecting the template, and then clicking **View**. You cannot change the template settings.

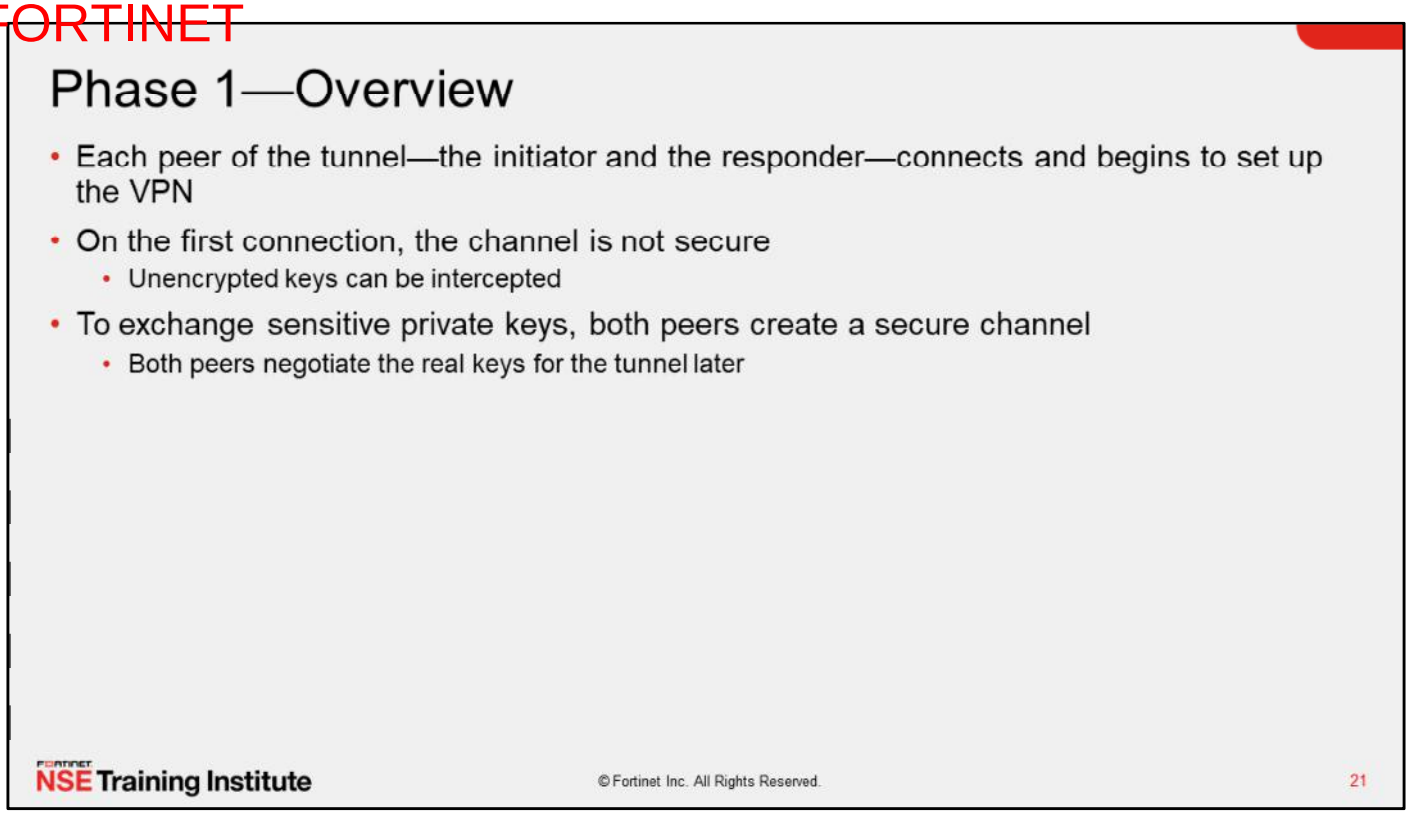

Phase 1 takes place when each peer of the tunnel—the initiator and the responder—connects and begins to set up the VPN. The initiator is the peer that starts the phase 1 negotiation, while the responder is the peer that responds to the initiator request.

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's identity, so how can they exchange sensitive private keys? They can't. First, both peers create a secure tunnel. They use this secure tunnel to protect strong authentication and negotiate the real keys for the tunnel later.

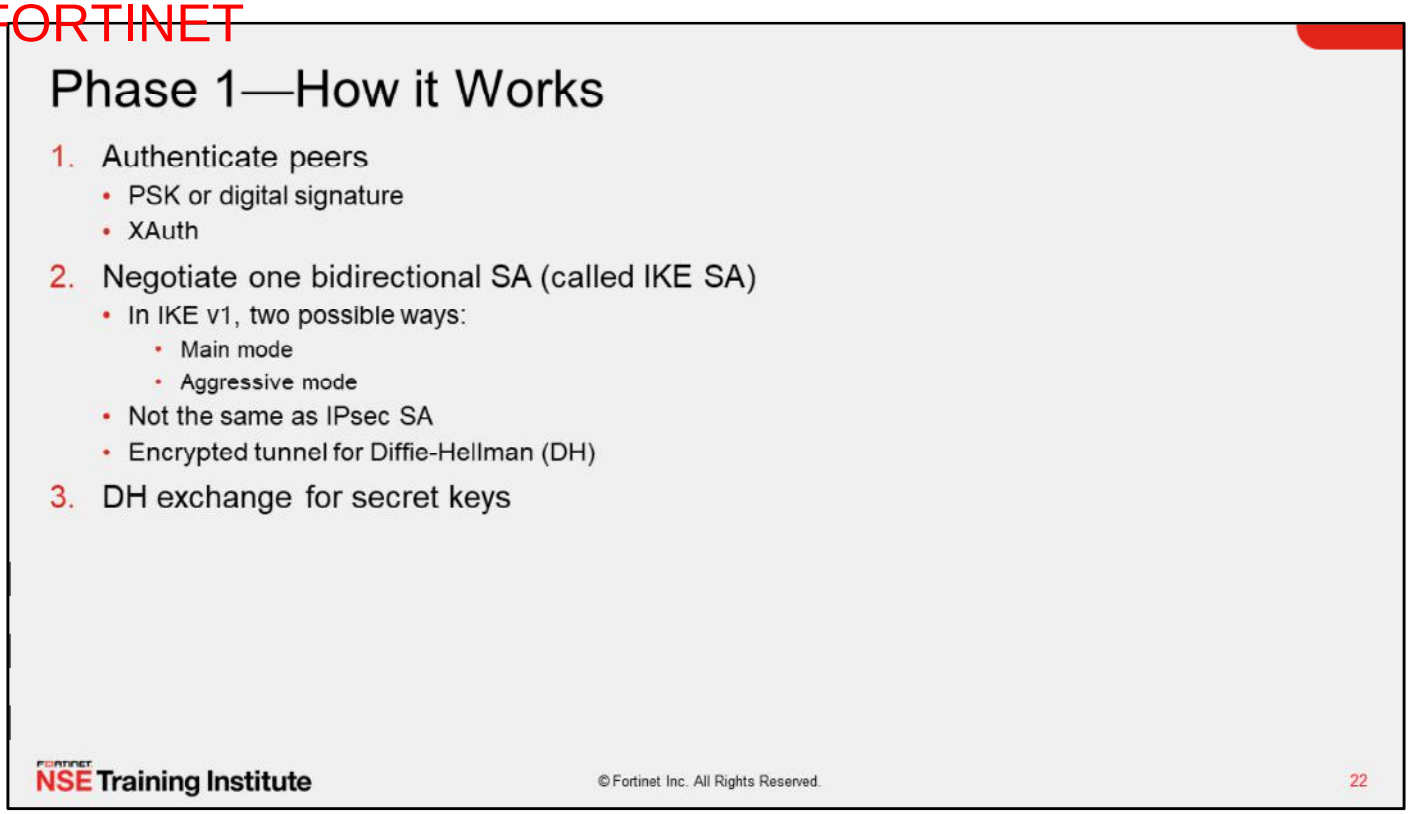

Now you'll examine how phase 1 works.

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.

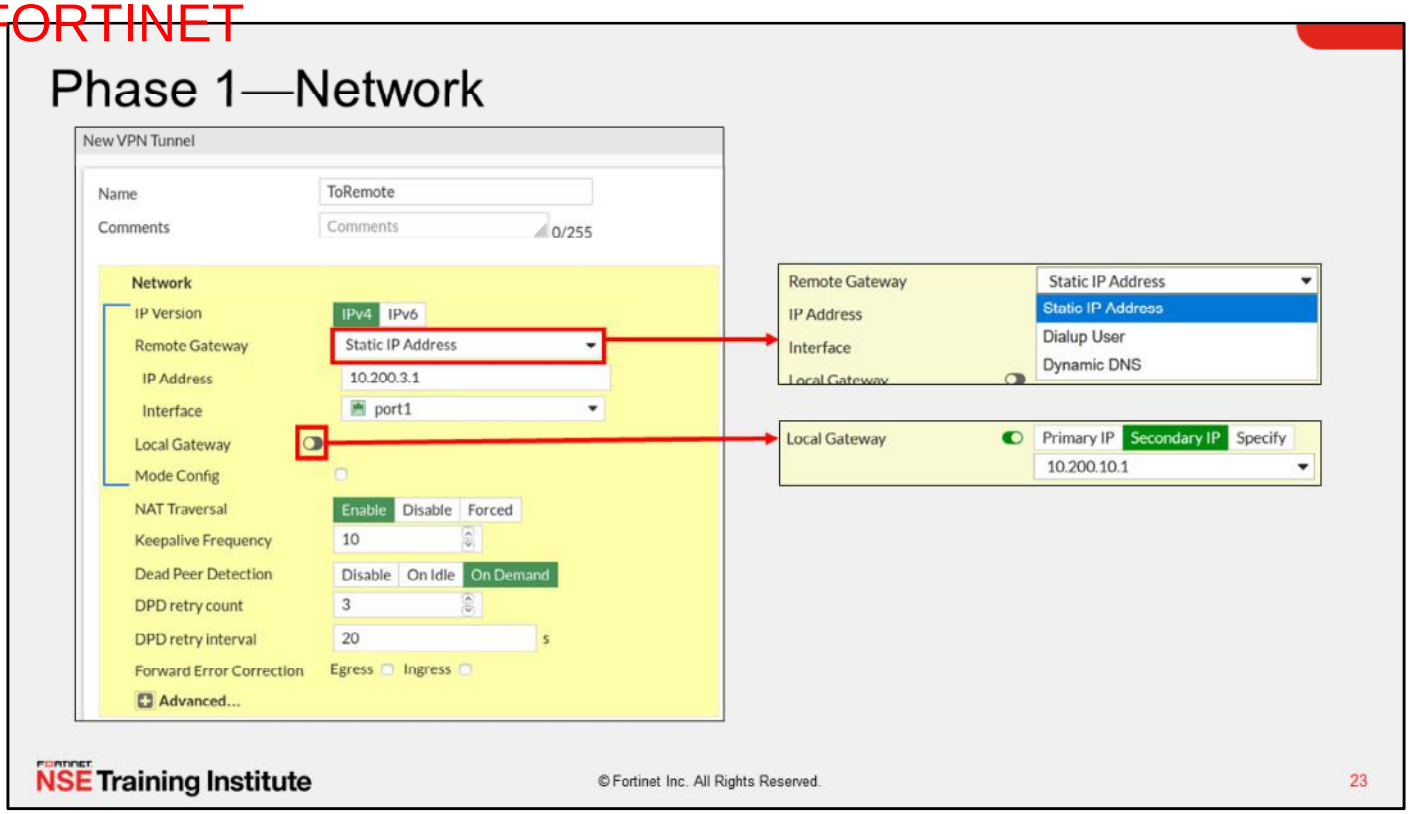

Phase 1 configuration is broken down on the GUI into four sections: **Network**, **Authentication**, **Phase 1 Proposal**, and **XAUTH**. You will learn about the settings available on each section. You will learn about some of these settings in more detail on separate slides.

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.

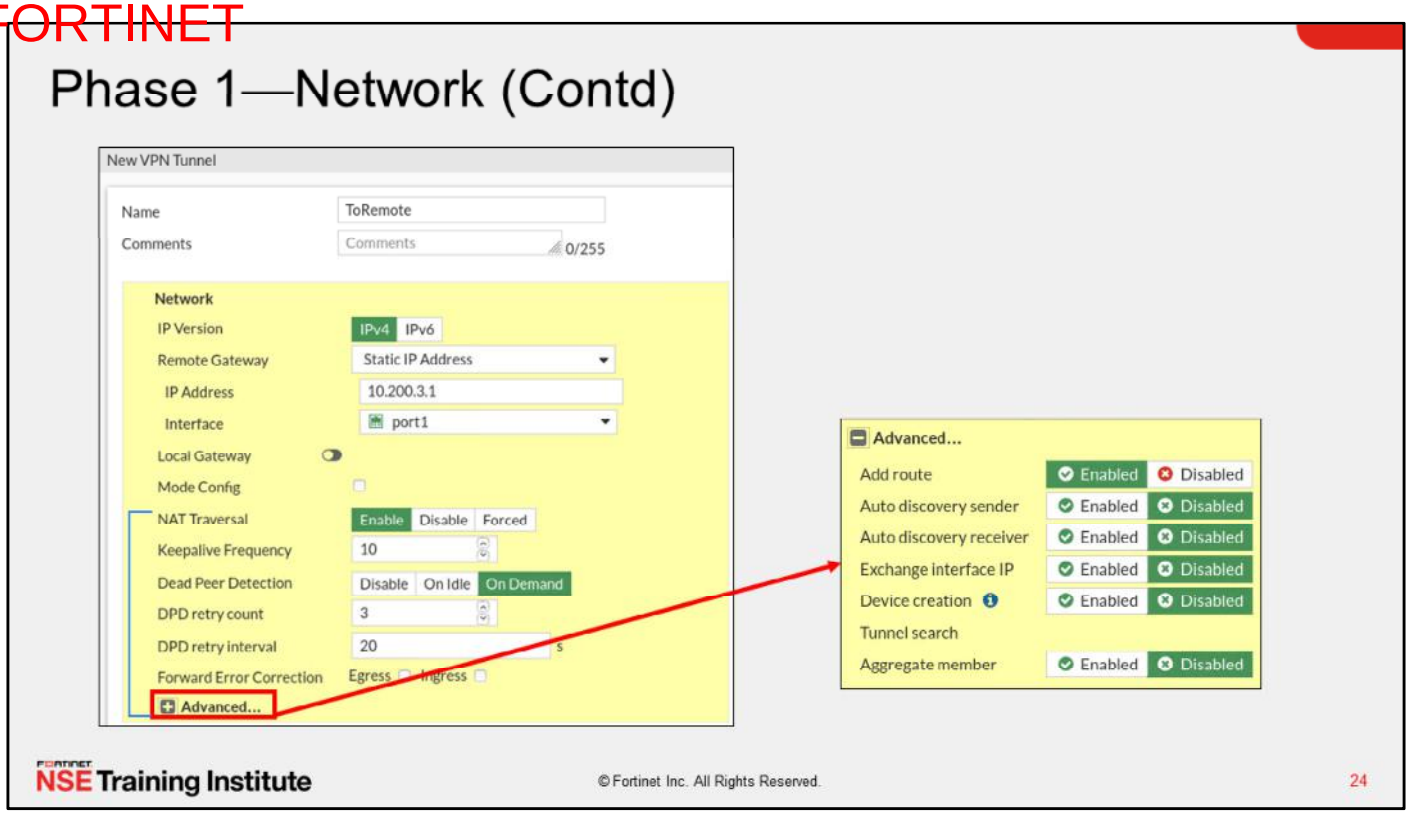

Below are the other options available on the GUI for the **Network** section:

- **NAT Traversal**. 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 unreliable 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 add-route if you are using a dynamic routing protocol over IPsec and do not want FortiGate to automatically add static routes.
  - **Auto discovery sender**. You must enable set auto-discovery-sender if you want ADVPN in the hub. This setting indicates that when IPsec traffic transits the hub, it should send a shortcut offer to the initiator of the traffic to indicate that it could perhaps establish a more direct connection (shortcut).
  - **Auto discovery receiver**. Enable ADVPN in the spoke. Use this command to indicate that this IPsec tunnel wants to participate in an autodiscovery VPN.
  - **Exchange interface IP**. In dial-up or ADVPN, this option allows the exchange of IPsec interface IP addresses. This allows a point-to-multipoint connection to the hub FortiGate device.
  - **Device creation**. When you enable device creation, the kernel creates an interface for every dialup client. For deployments with a large number of dialup clients, disable device creation to achieve higher performance.
  - **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.

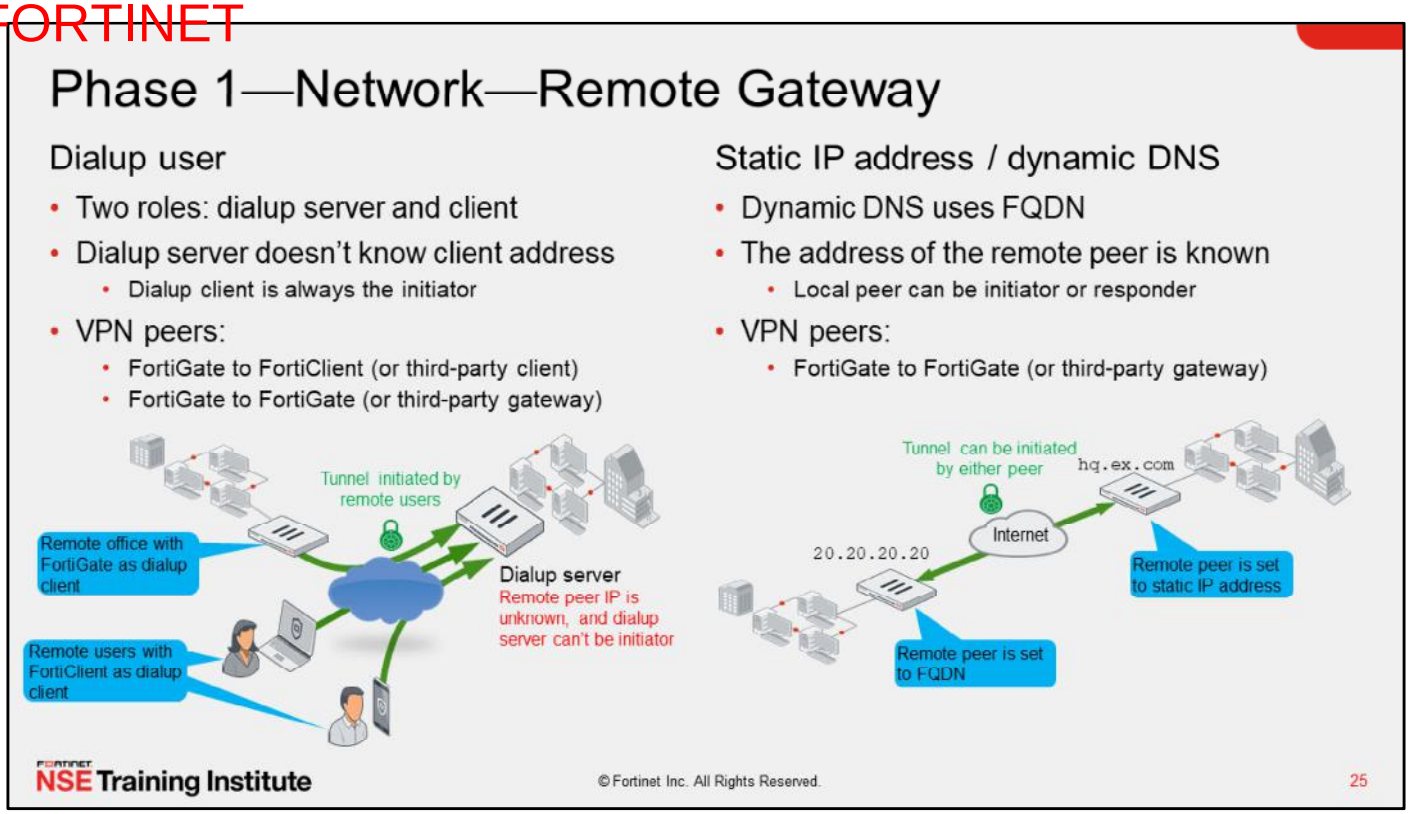

You have three options when configuring the remote gateway type of your VPN: **Dialup User**, **Static IP Address**, and **Dynamic DNS**.

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

Usually, dialup clients are remote and mobile employees with FortiClient on their computer or handheld devices. You can also have a FortiGate device acting as a dialup client for a remote office. One dialup 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 dialup setup, the dialup 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.

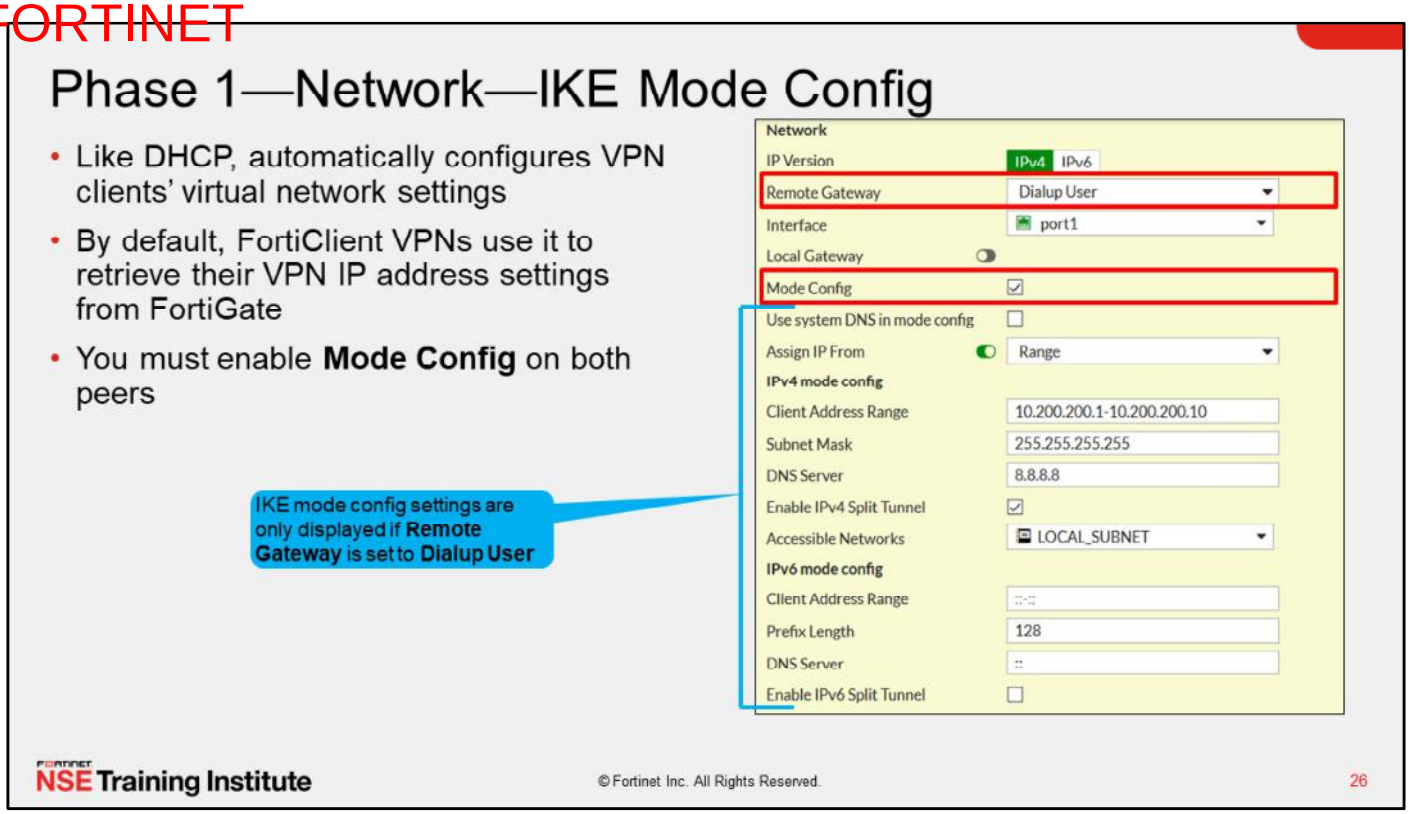

IKE **Mode Config** is similar to DHCP because a server assigns network settings such as IP address, netmask, and DNS servers, to clients. This assignment takes place over IKE messages.

When you enable **Mode Config** on a FortiGate device acting as dialup server, it pushes network settings to dialup clients. The dialup 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 dialup client, you can select **Mode Config** on the GUI, but the additional settings are not displayed.

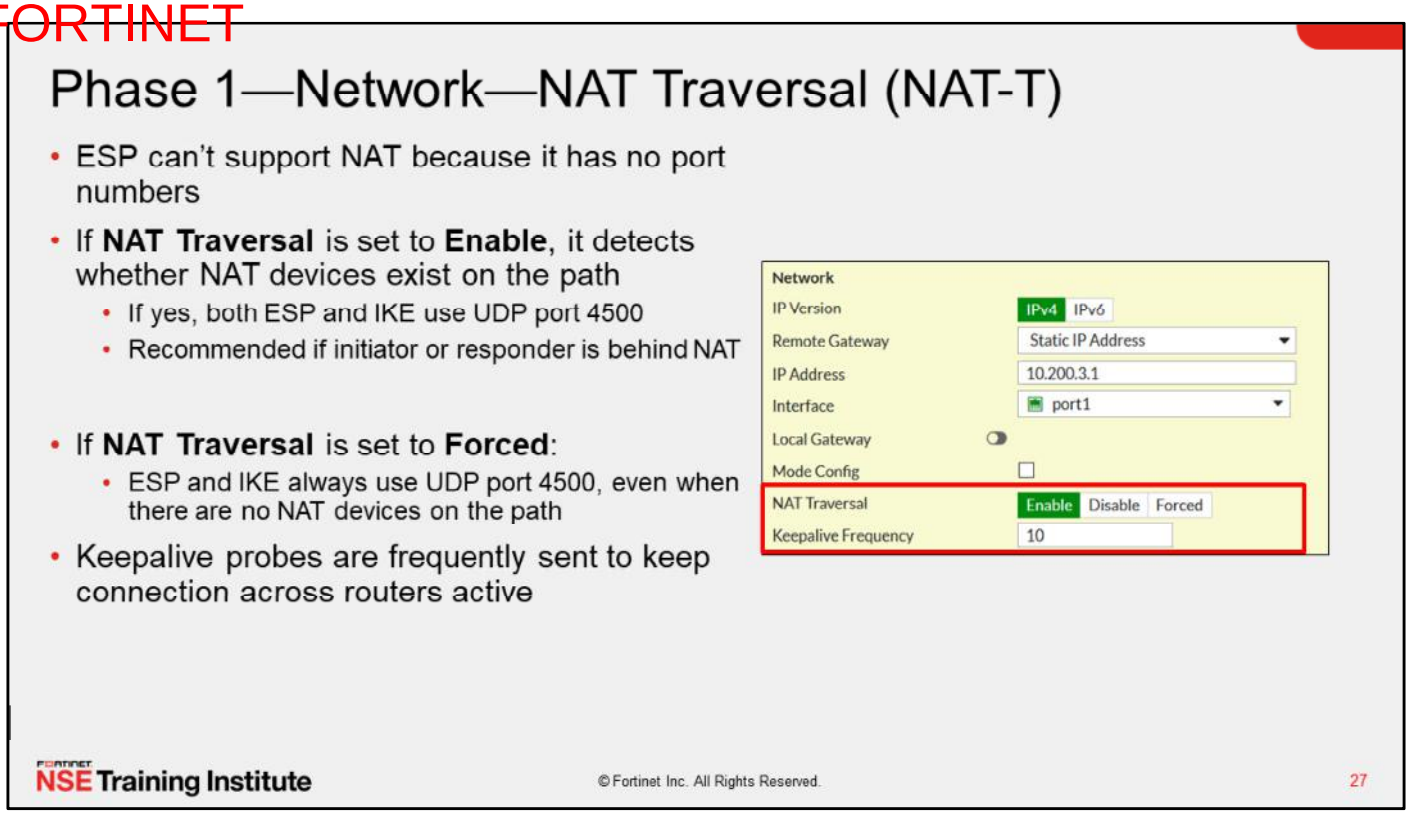

The ESP protocol usually has problems crossing devices that are performing NAT. One of the reasons is that ESP does not use port numbers, like TCP and UDP do, to differentiate one tunnel from another.

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 has 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

NSE 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 failover 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**.

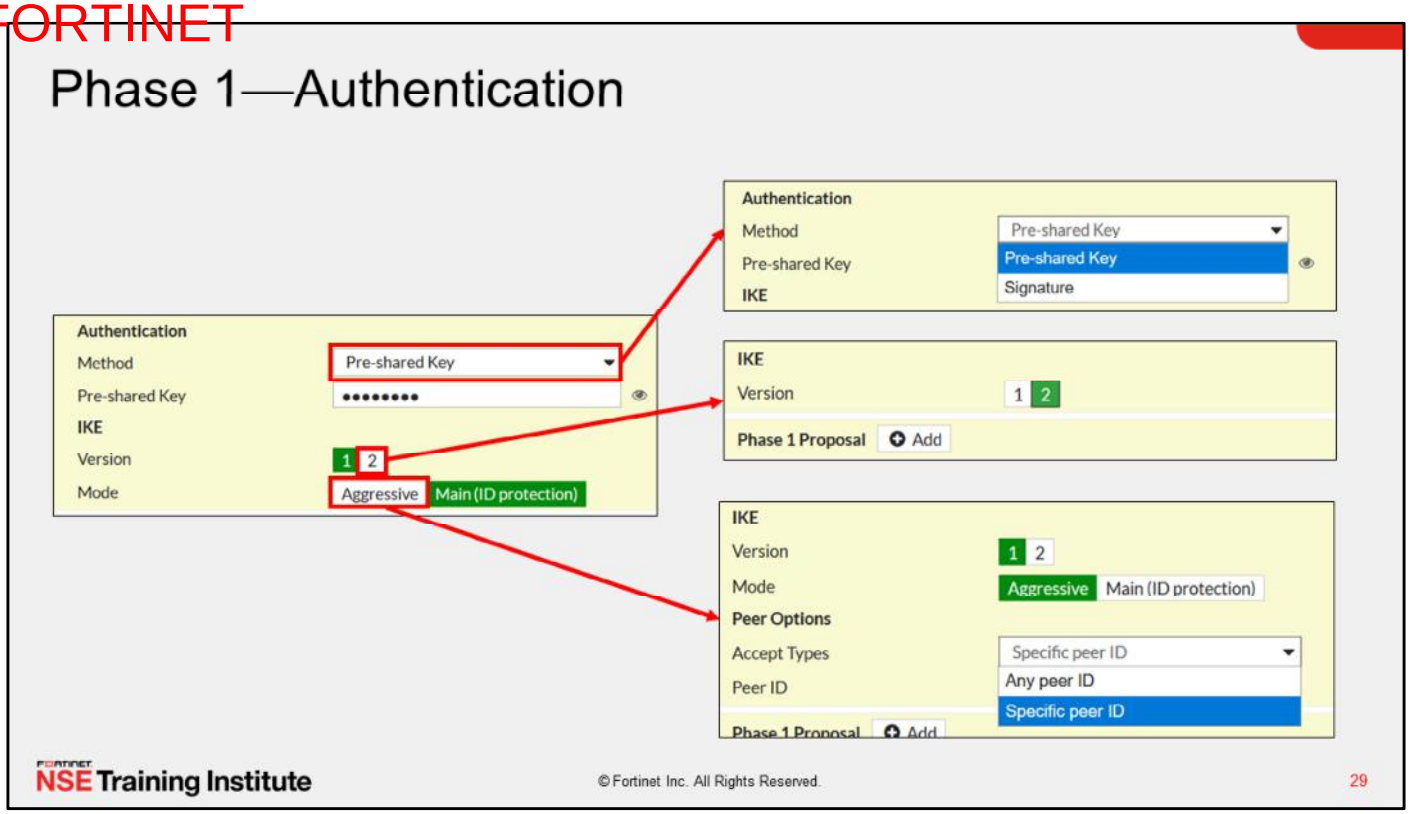

Now, you will learn about the **Authentication** section in phase 1 configuration:

- **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.

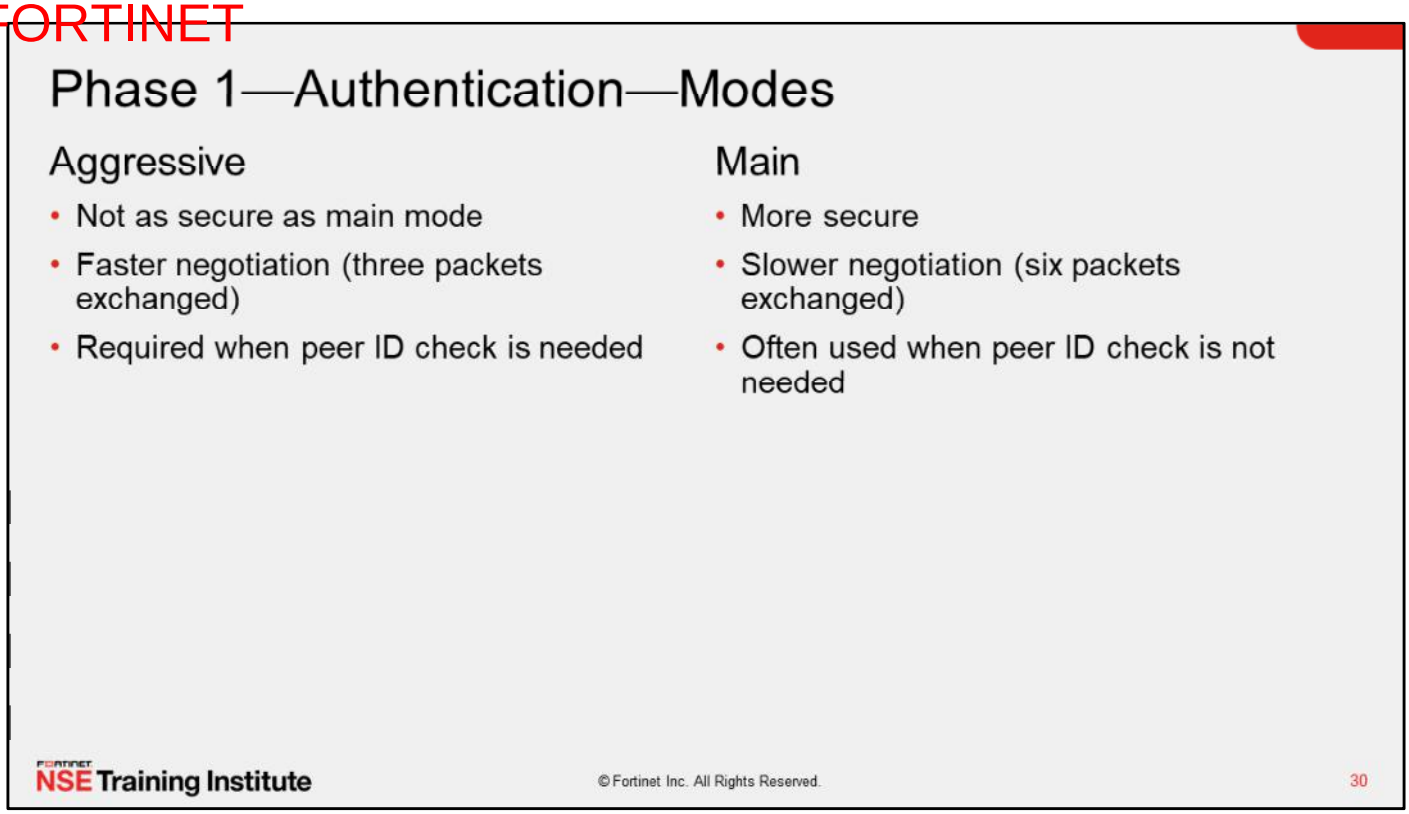

IKE supports two different negotiation modes: main and aggressive. Which one should you use?

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 dialup 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 dialup 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.

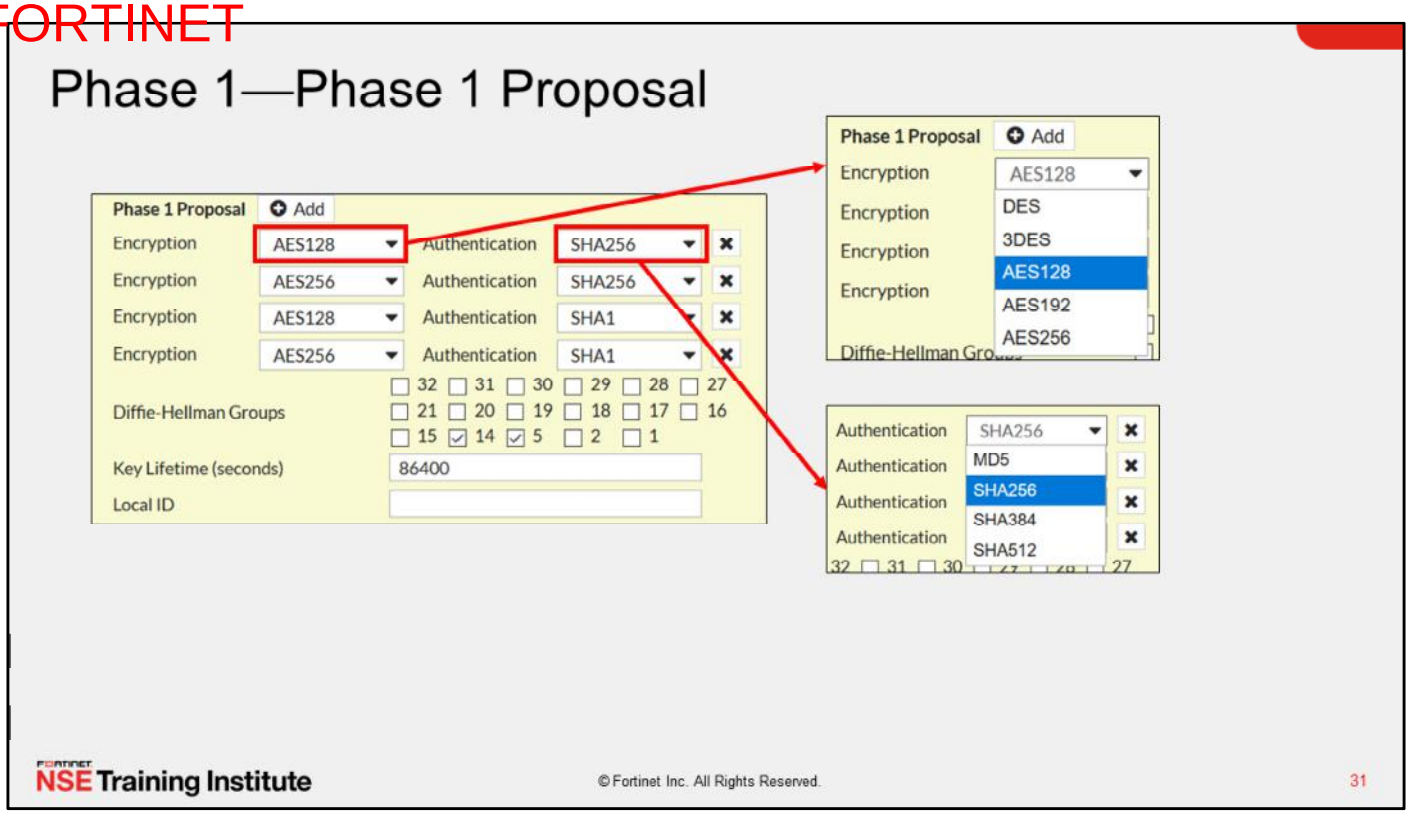

Now, you will learn about the **Phase 1 Proposal** section of phase 1 configuration. This section allows you to enable the different proposals that FortiGate supports when negotiating the IKE SA (or phase 1 SA). You can combine different parameters to suit your security needs. You must at least configure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

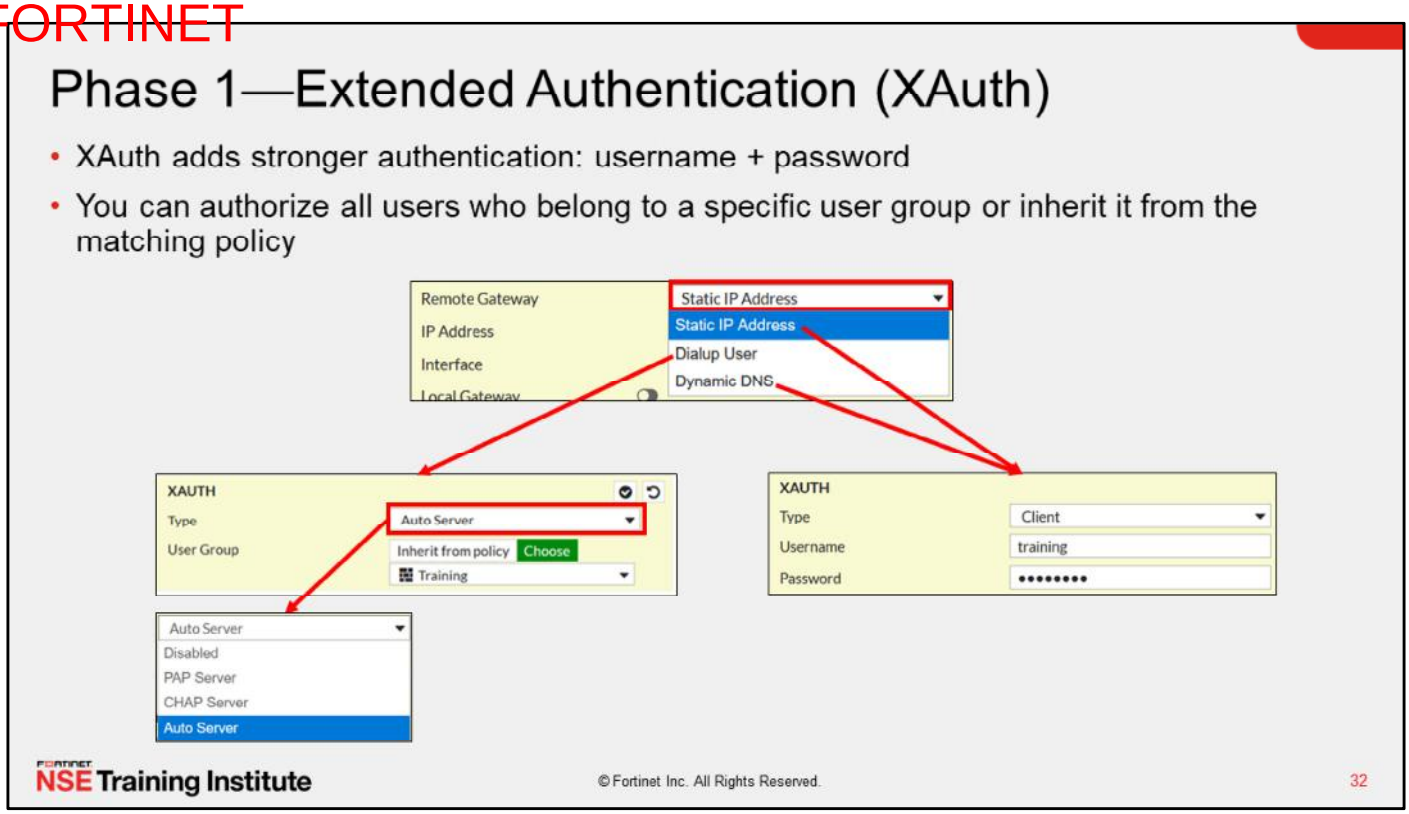

Phase 1 supports two types of authentication: pre-shared keys and digital signatures. The XAuth extension, sometimes called phase 1.5, forces remote users to authenticate additionally with their credentials (username and password). So, additional authentication packets are exchanged if you enable it.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

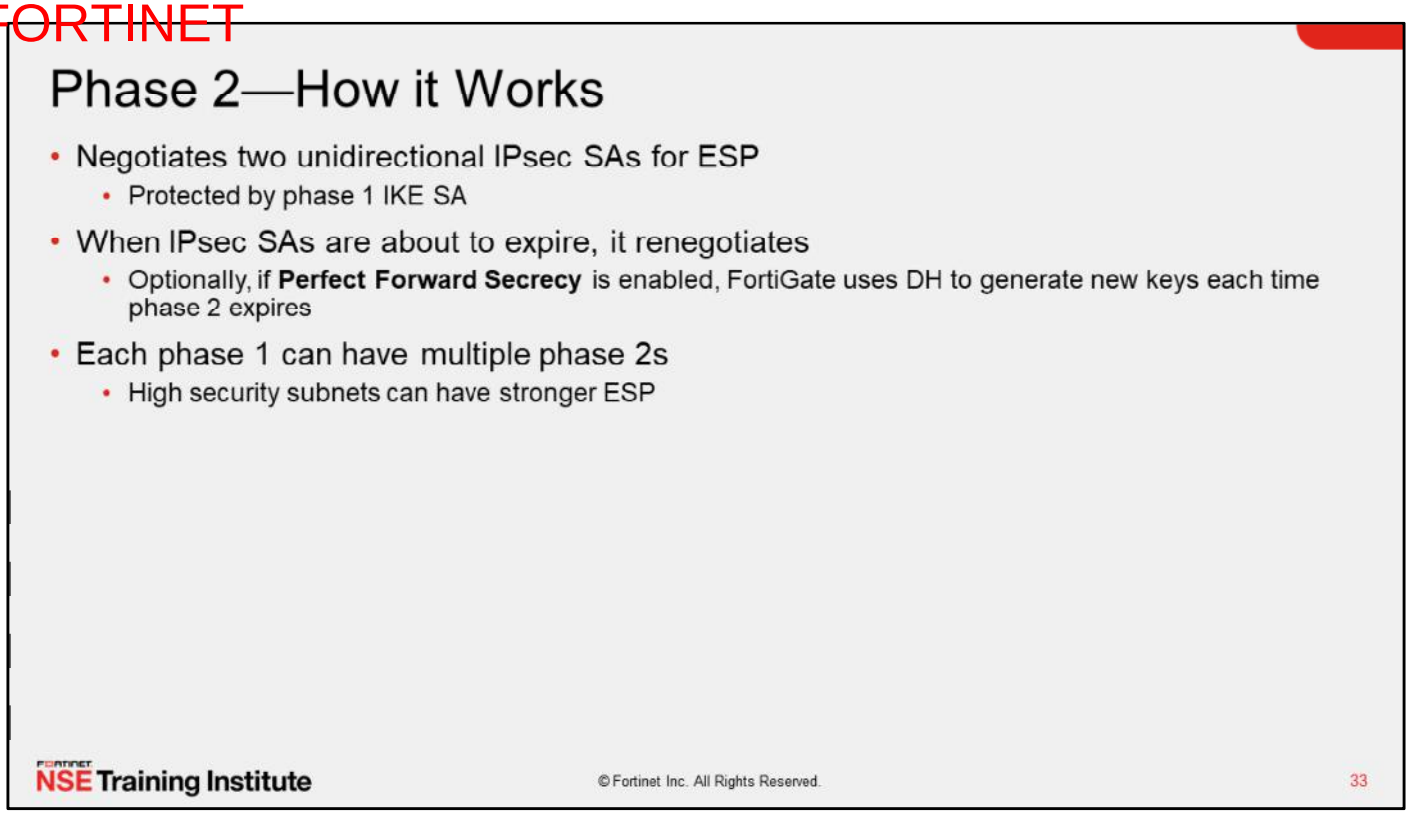

After phase 1 has established a secure channel to exchange data, phase 2 begins.

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.

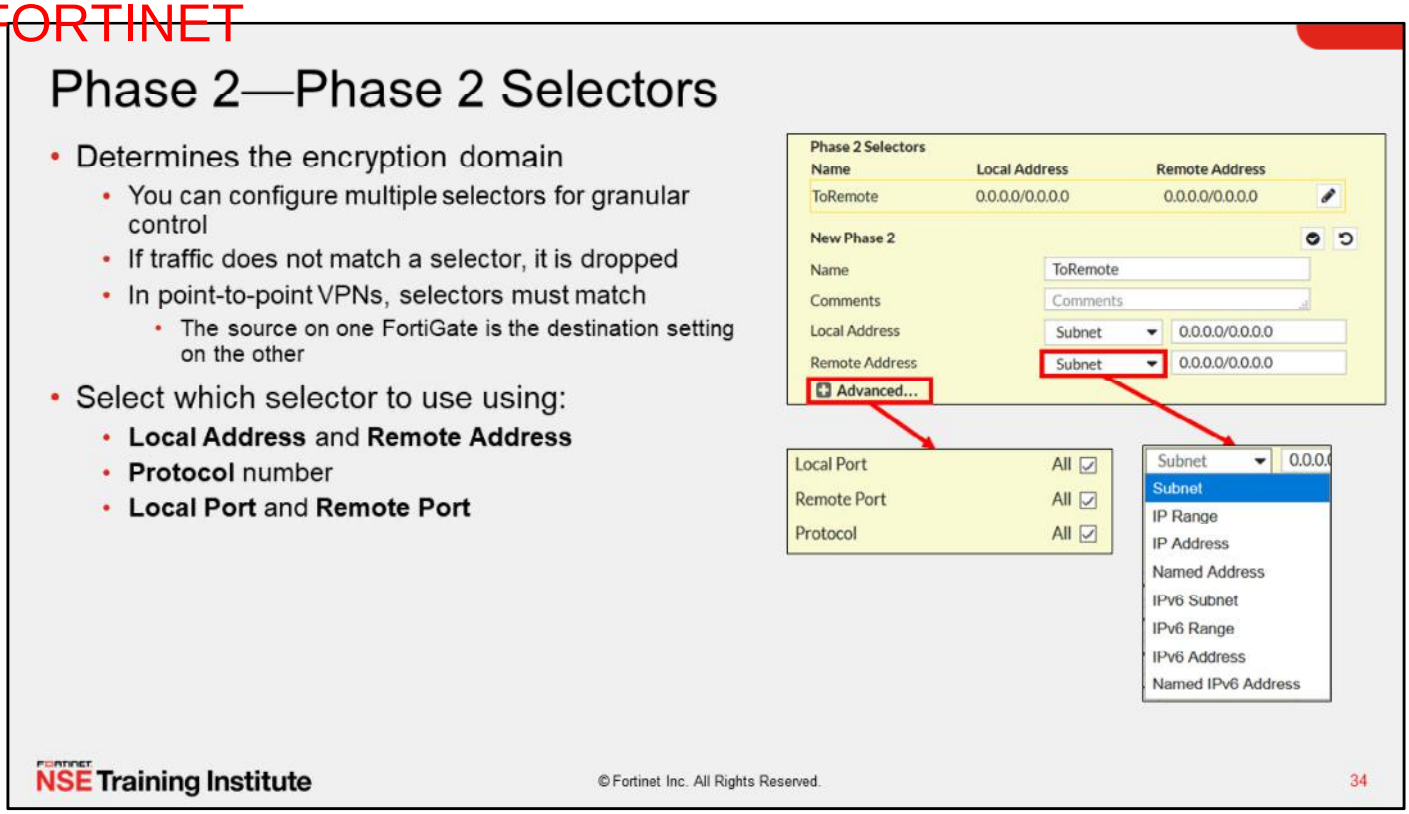

In phase 2, you must define the encryption domain (or interesting traffic) of your IPsec tunnel. The encryption domain refers to the traffic that you want to protect with IPsec, and it is determined by your phase 2 selector configuration.

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, 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 dialup user.

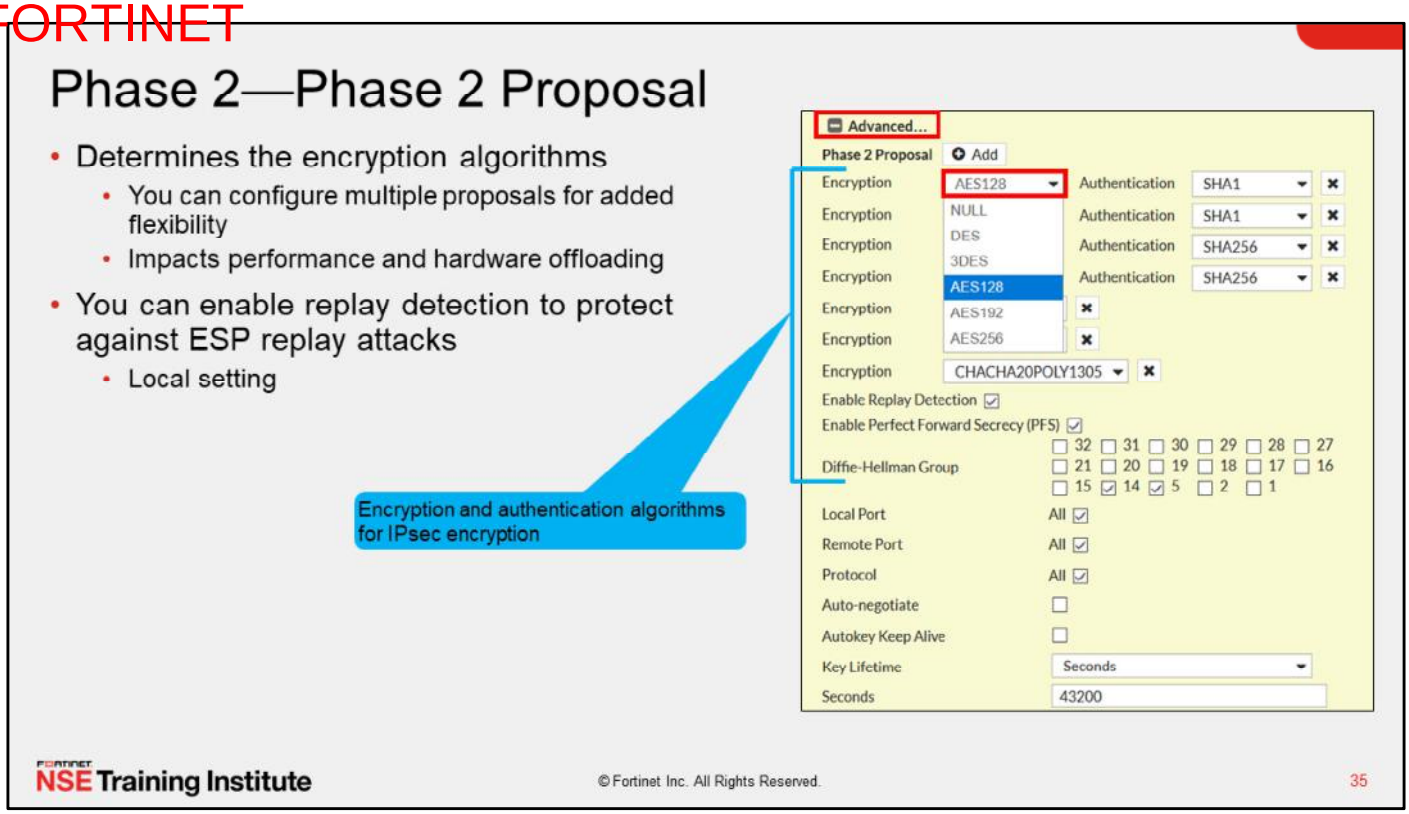

For every phase 2 selector, you need to configure one or more phase 2 proposals. A phase 2 proposal defines the algorithms supported by the peer for encrypting and decrypting the data over the tunnel. You can configure multiple proposals to offer more options to the remote peer when negotiating the IPsec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

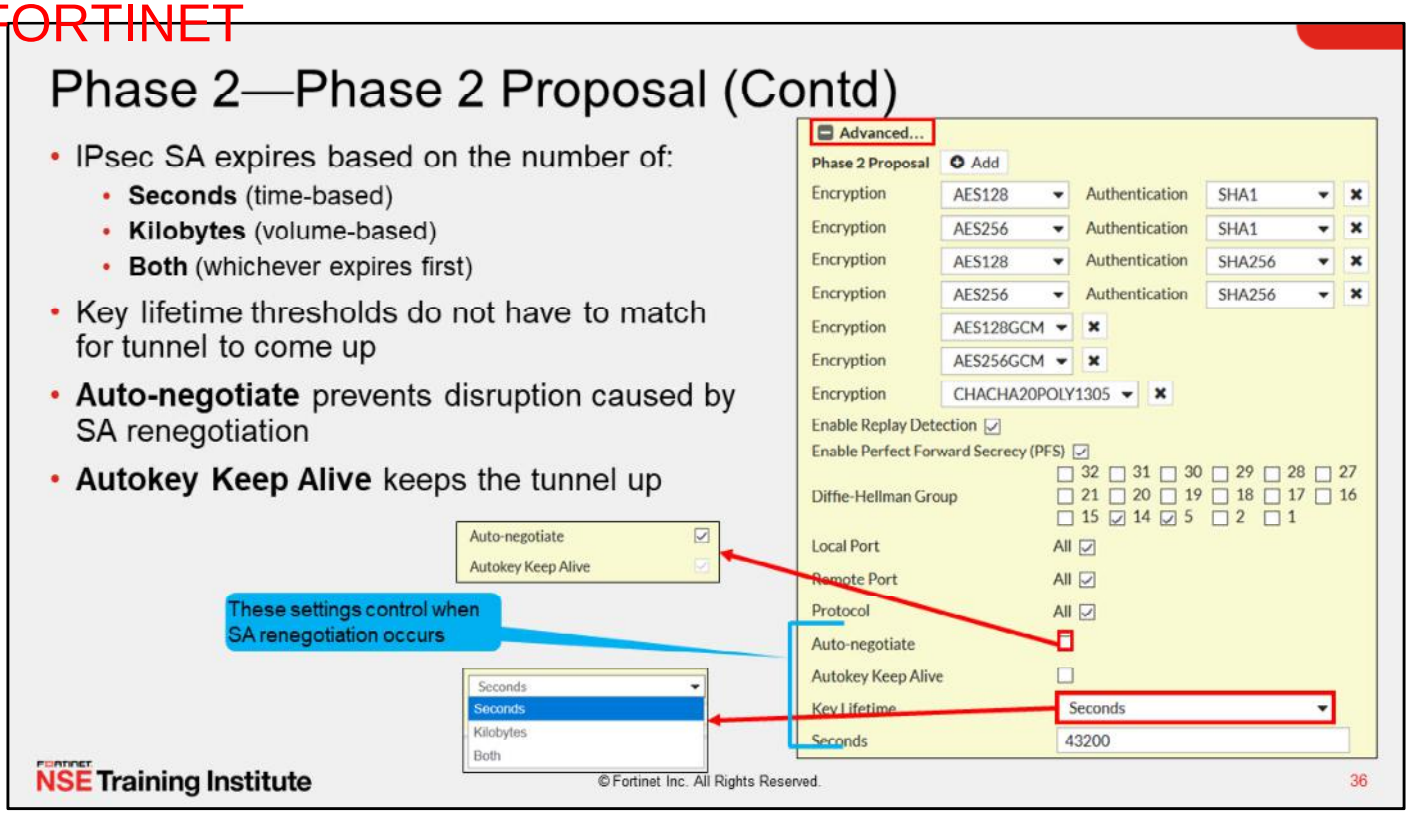

IPsec SAs are periodically renegotiated to improve security, but when does that happen? It depends on the key lifetime settings configured on the phase 2 proposal.

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.

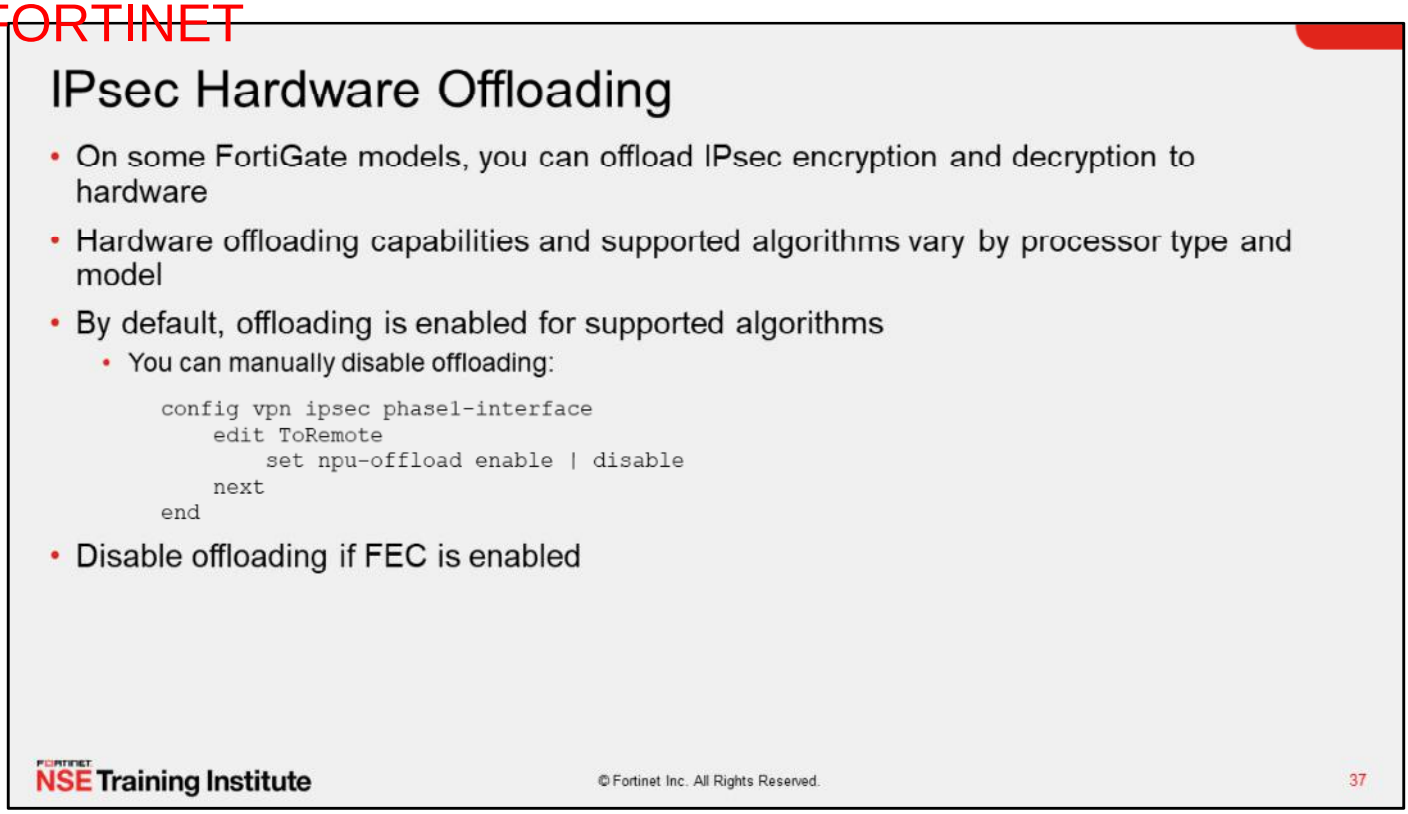

On some FortiGate models, you can offload the encryption and decryption of IPsec traffic to hardware. The supported algorithms depend on the NP unit model present on the FortiGate device. For a list of supported encryption algorithms for IPsec hardware offloading, refer to `https://docs.fortinet.com`.

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.

Finally, if you want to use FEC for IPsec tunnels, you must disable IPsec offload for the feature to work.

Knowledge Check

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

NSE Training Institute

38

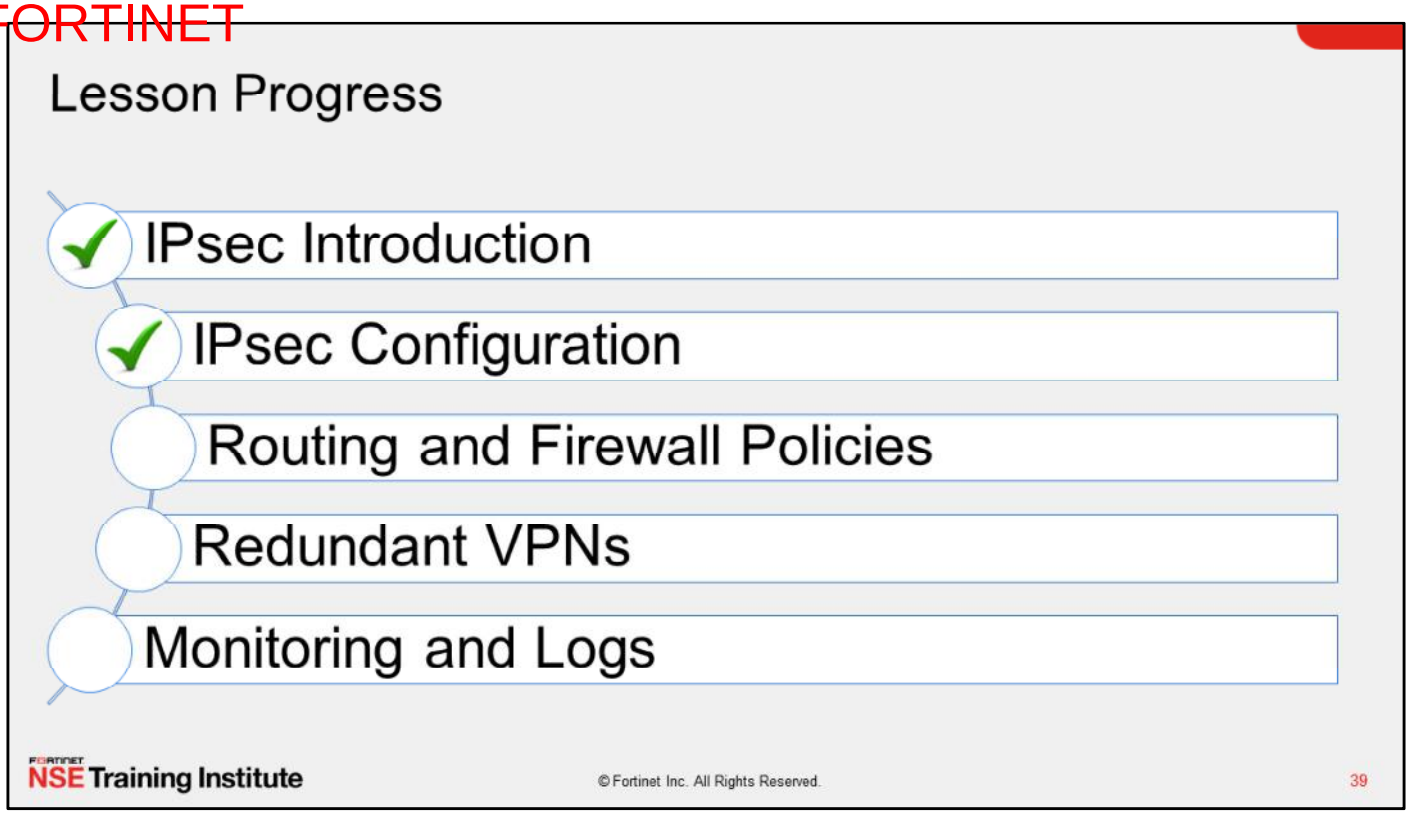

Good job! You now understand IPsec configuration.

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

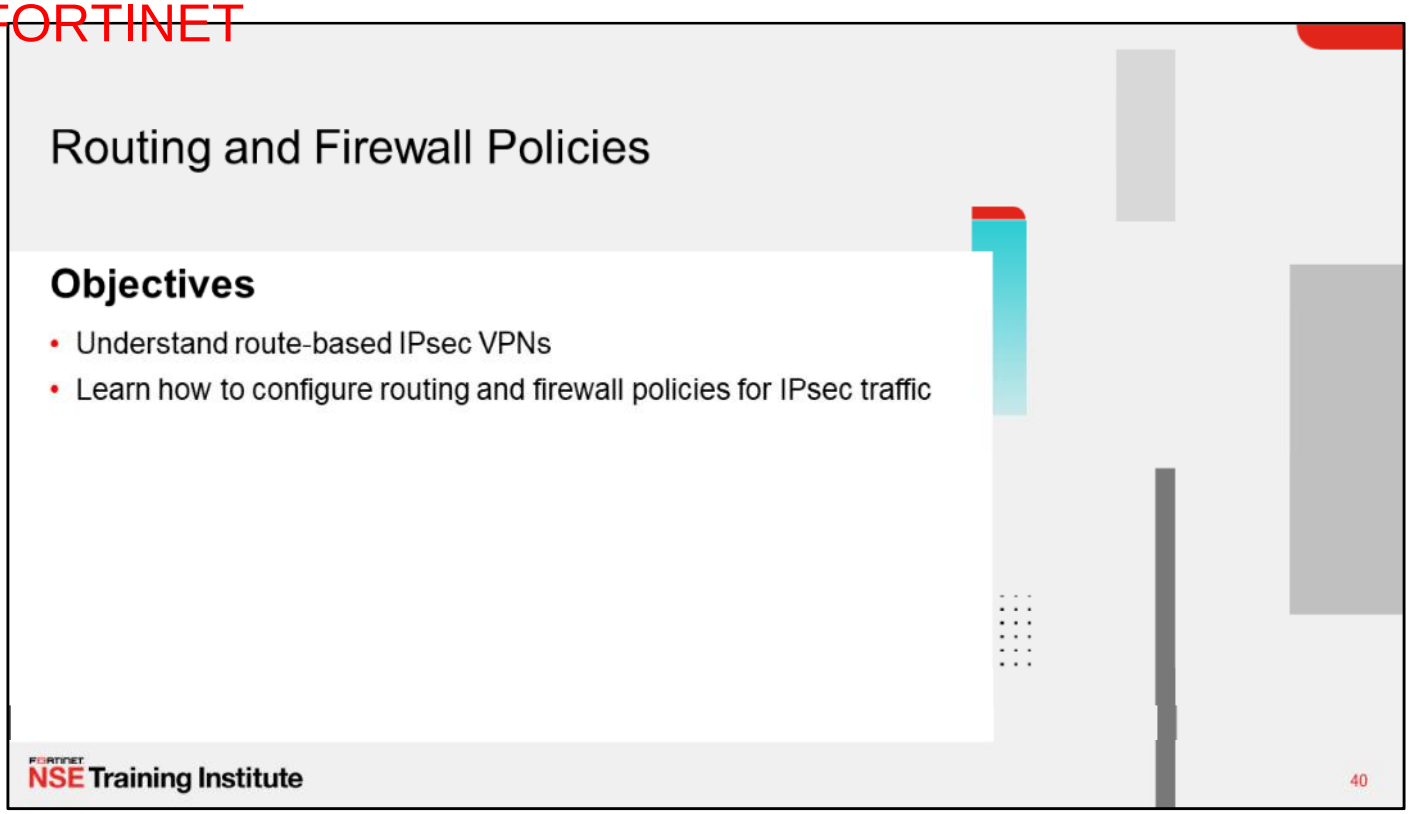

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

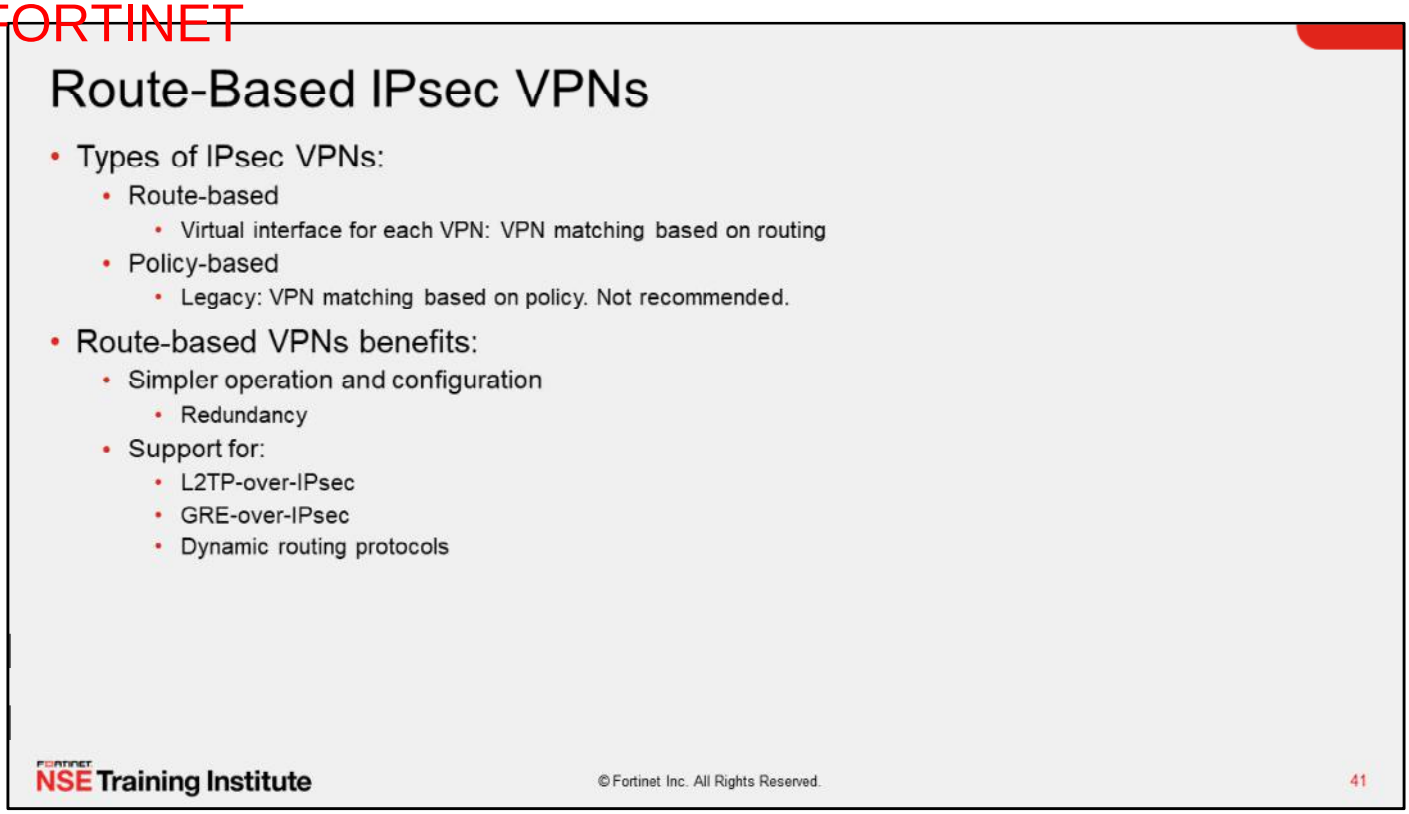

FortiGate supports two types of IPsec VPNs: route-based and policy-based. Policy-based is a legacy IPsec VPN that is supported only for backward compatibility reasons, and its use *is not recommended* for new deployments. Unless otherwise stated, all IPsec VPN references in this lesson are for route-based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

Dialup 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 dialup 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

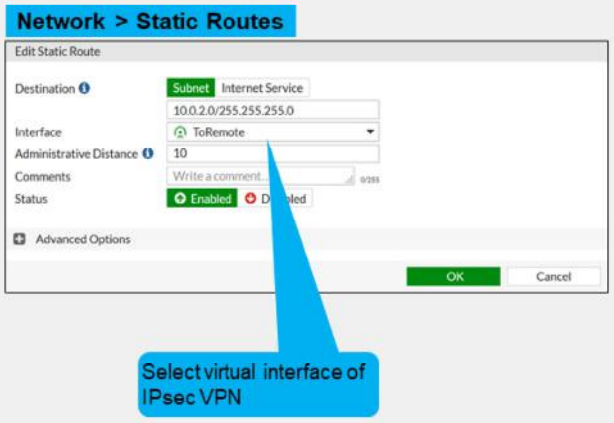

NSE Training Institute

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.

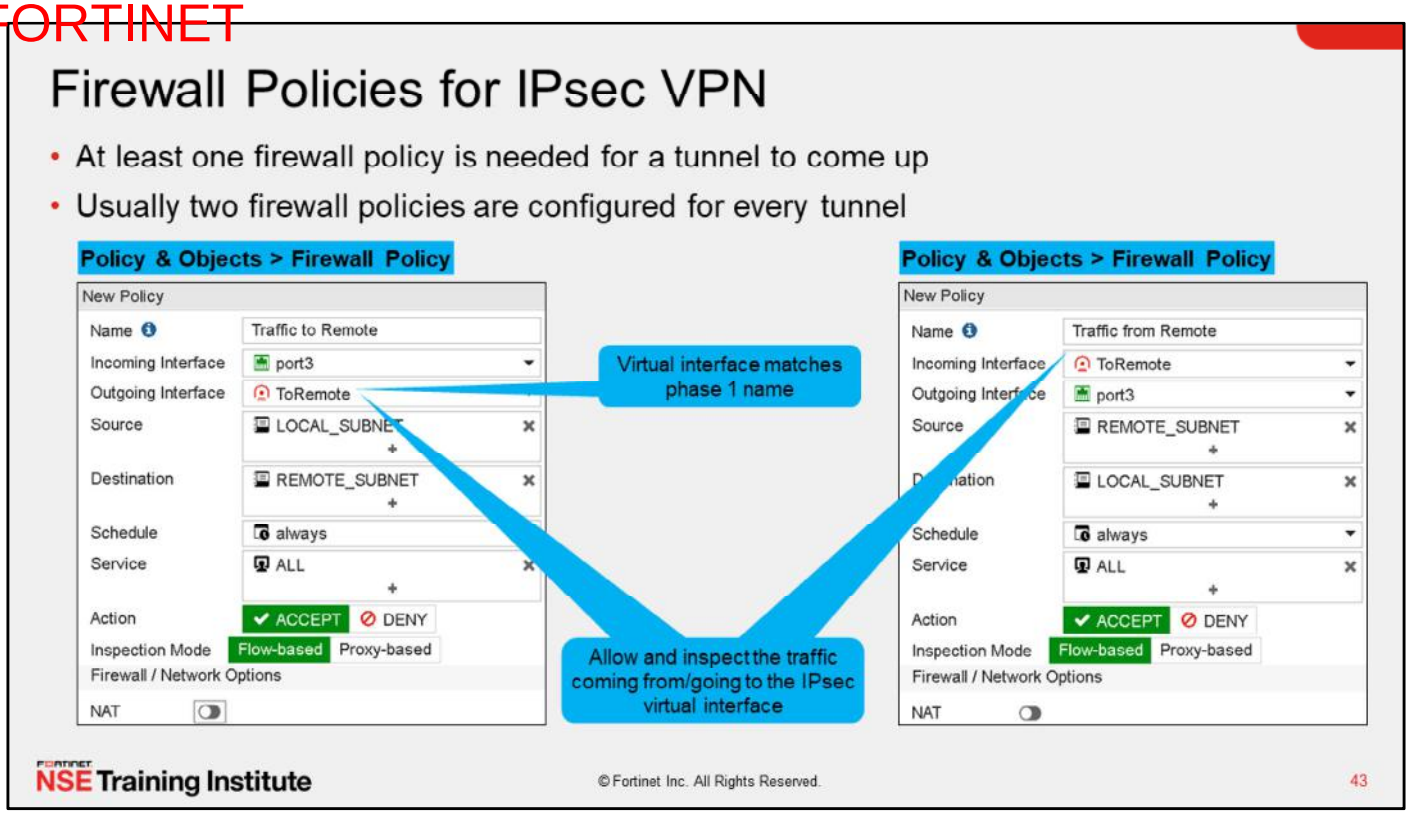

You must configure at least one firewall policy that accepts traffic on your IPsec tunnel. Otherwise, the tunnel will not come up.

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 VPN
   - ✓ B. Policy-based 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

NSE Training Institute

44

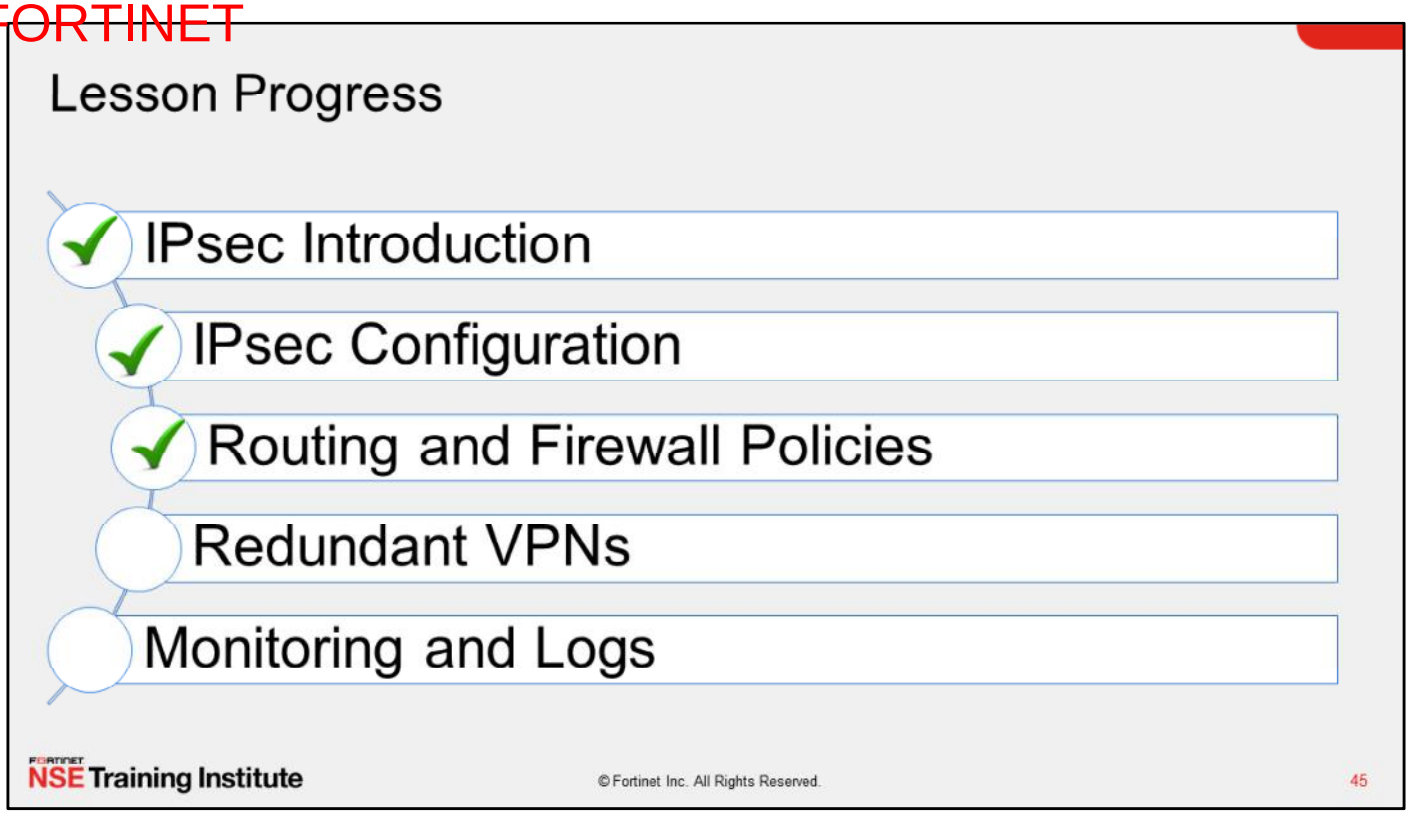

Good job! You now understand routing and firewall policies for IPsec traffic.

Now, you'll learn about redundant VPNs.

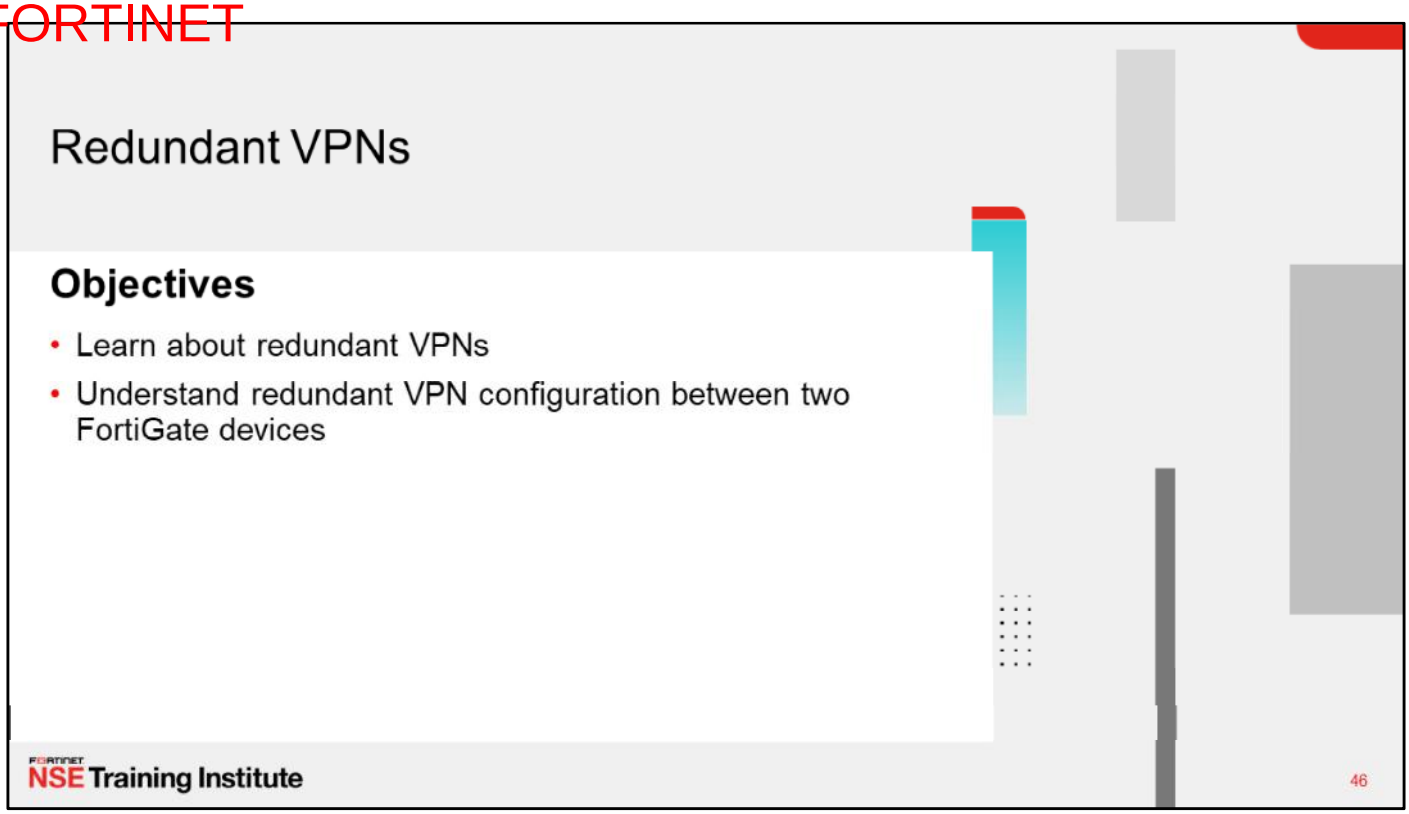

After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in redundant VPNs, you will able to add redundancy to your IPsec VPN deployment.

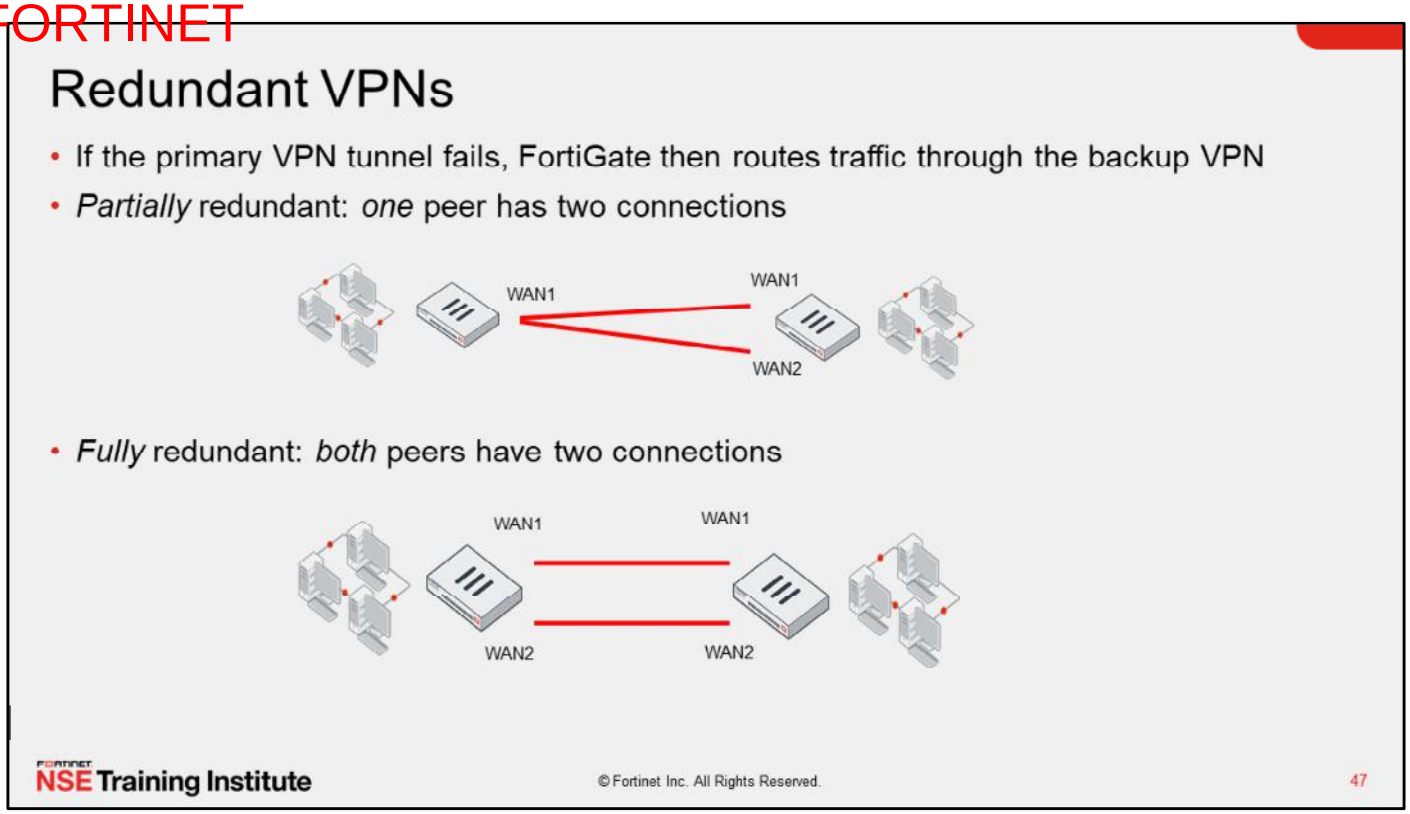

How can you make your IPsec VPN deployment more resilient? Provide a second ISP connection to your site and configure two IPsec VPNs. If the primary IPsec VPN fails, another tunnel can be used instead.

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

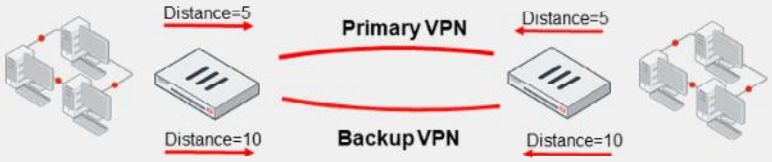

NSE Training Institute

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

NSE Training Institute

49

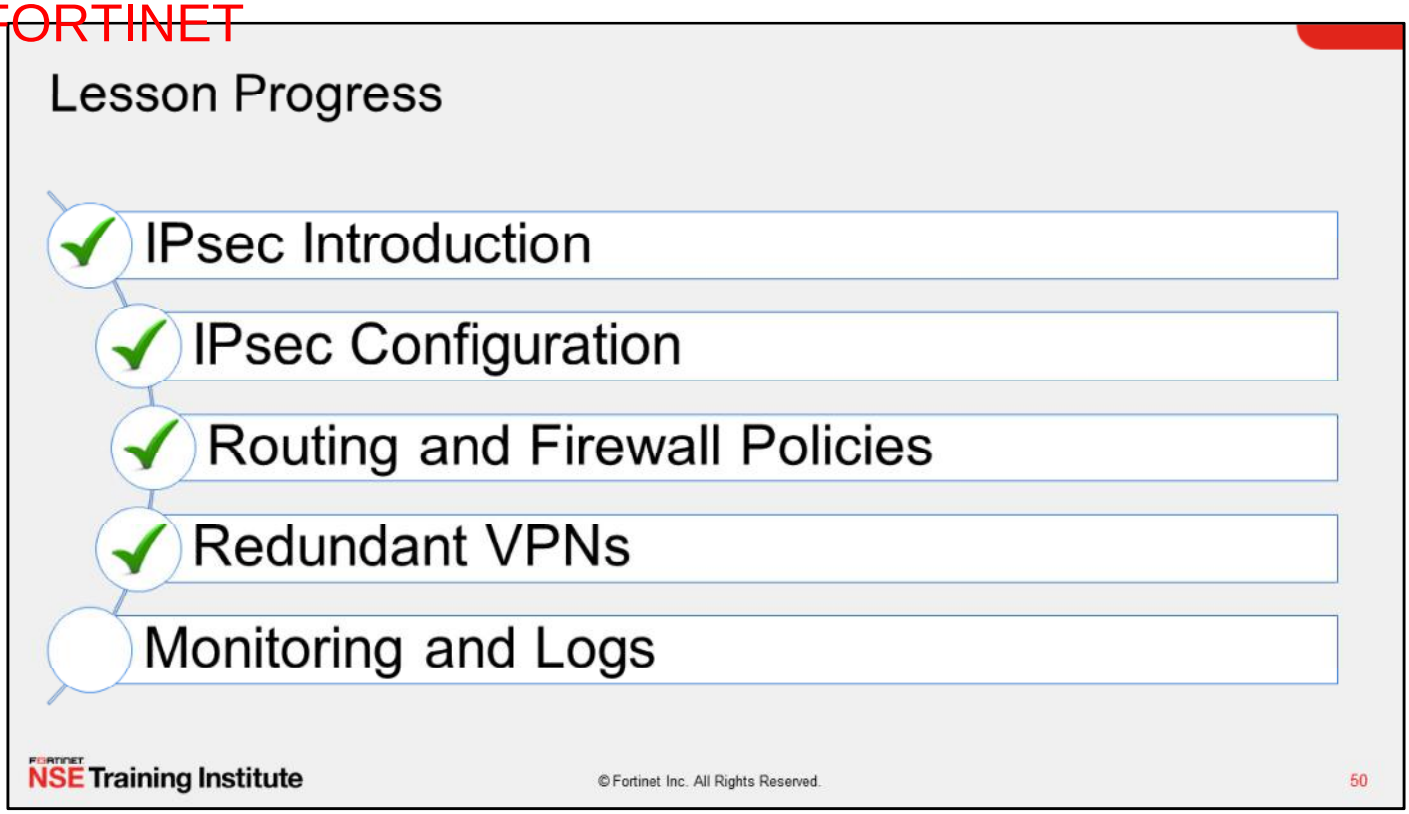

Good job! You now understand redundant VPNs.

Now, you'll learn about monitoring IPsec VPNs and reviewing their logs.

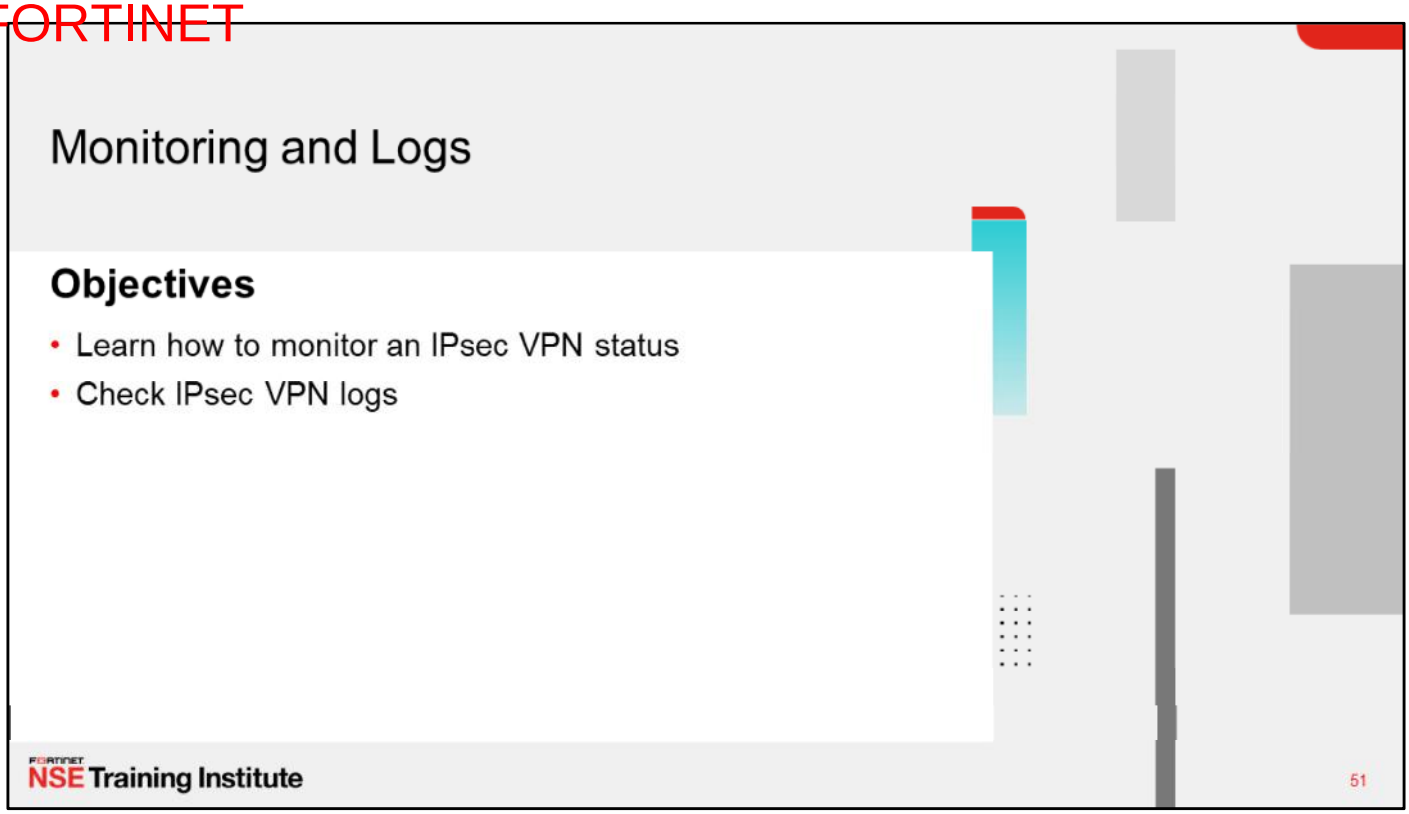

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in monitoring and logs, you will be able to monitor IPsec VPN and review past events.

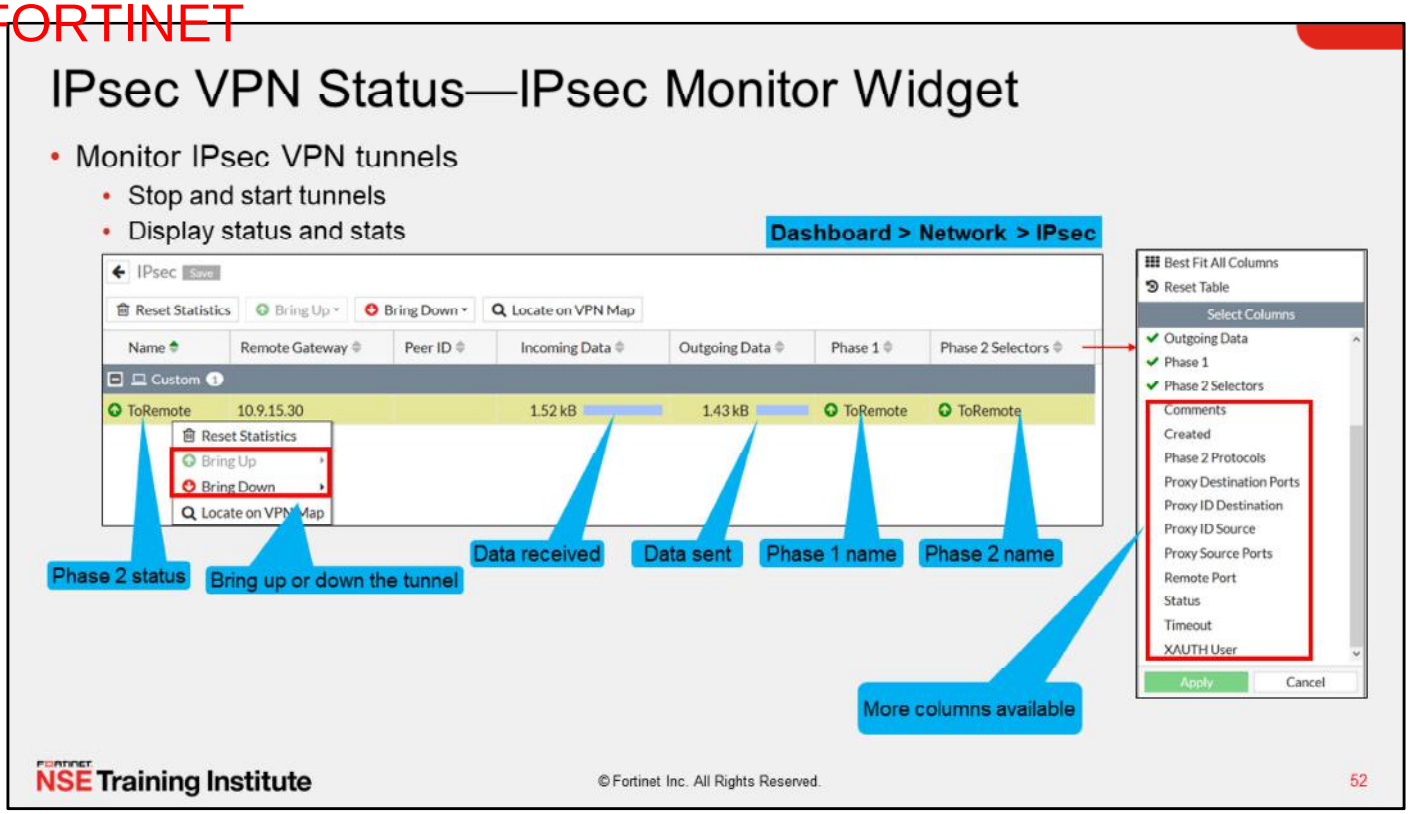

On the GUI dashboard, you can use the IPsec widget to monitor the status of your IPsec VPNs. The widget shows the phase 1 and phase 2 status of an IPsec VPN.

You can also bring up or down individual VPNs, and get additional details. When you bring up or down IPsec VPNs using the IPsec widget, you are affecting only the phase 2 status of the tunnel, *not* its phase 1 status.

If the tunnel is up, meaning phase 2 is up, a green up arrow is displayed next to its name. If it is down, then a red down arrow is displayed.

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.

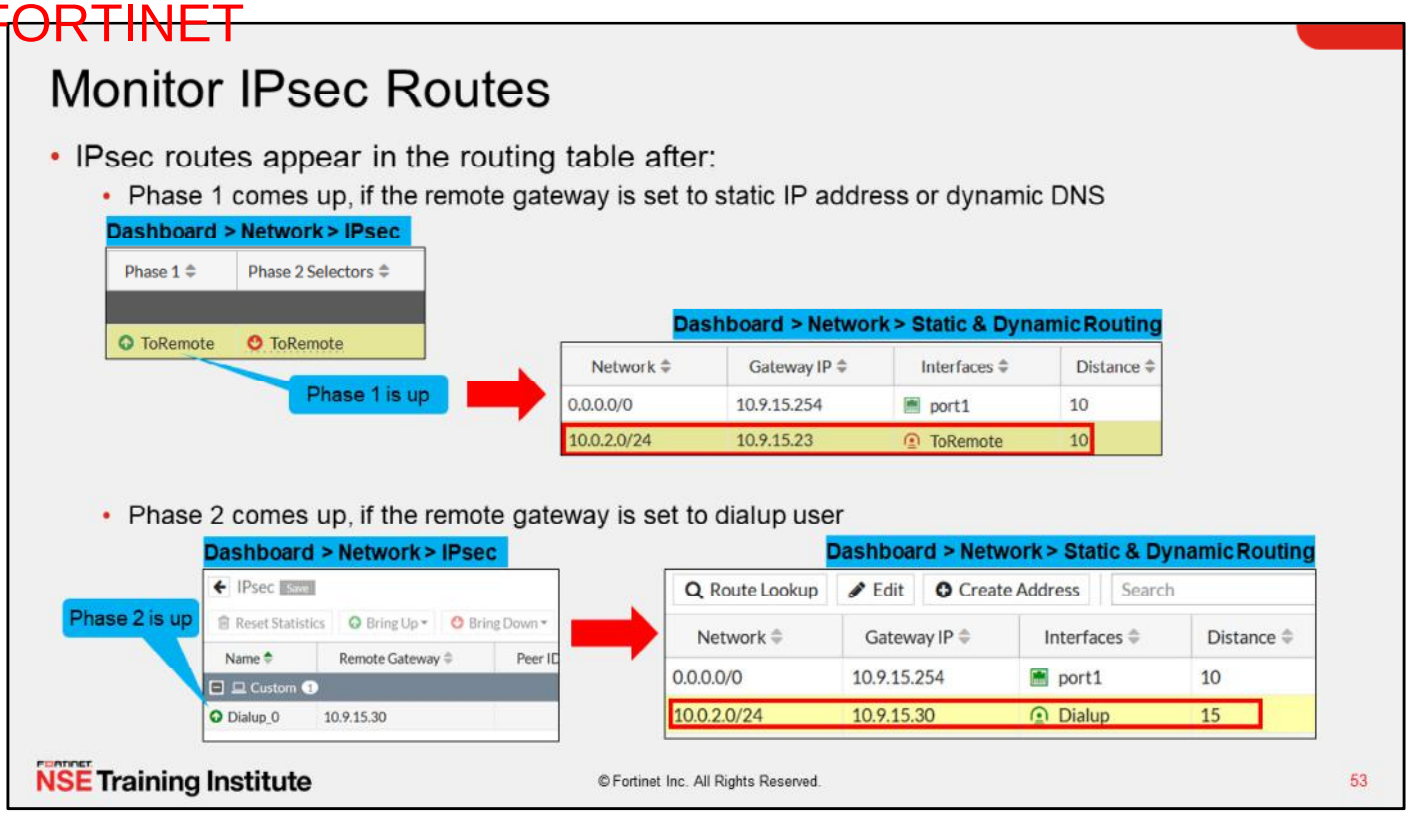

If you set the remote gateway to **Static IP Address** or **Dynamic DNS**, the static routes for these tunnels become active in the routing table after phase 1 comes up. Phase 1 negotiation is started automatically because automatic negotiation is enabled on phase 1 by default. This behavior allows FortiGate 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.

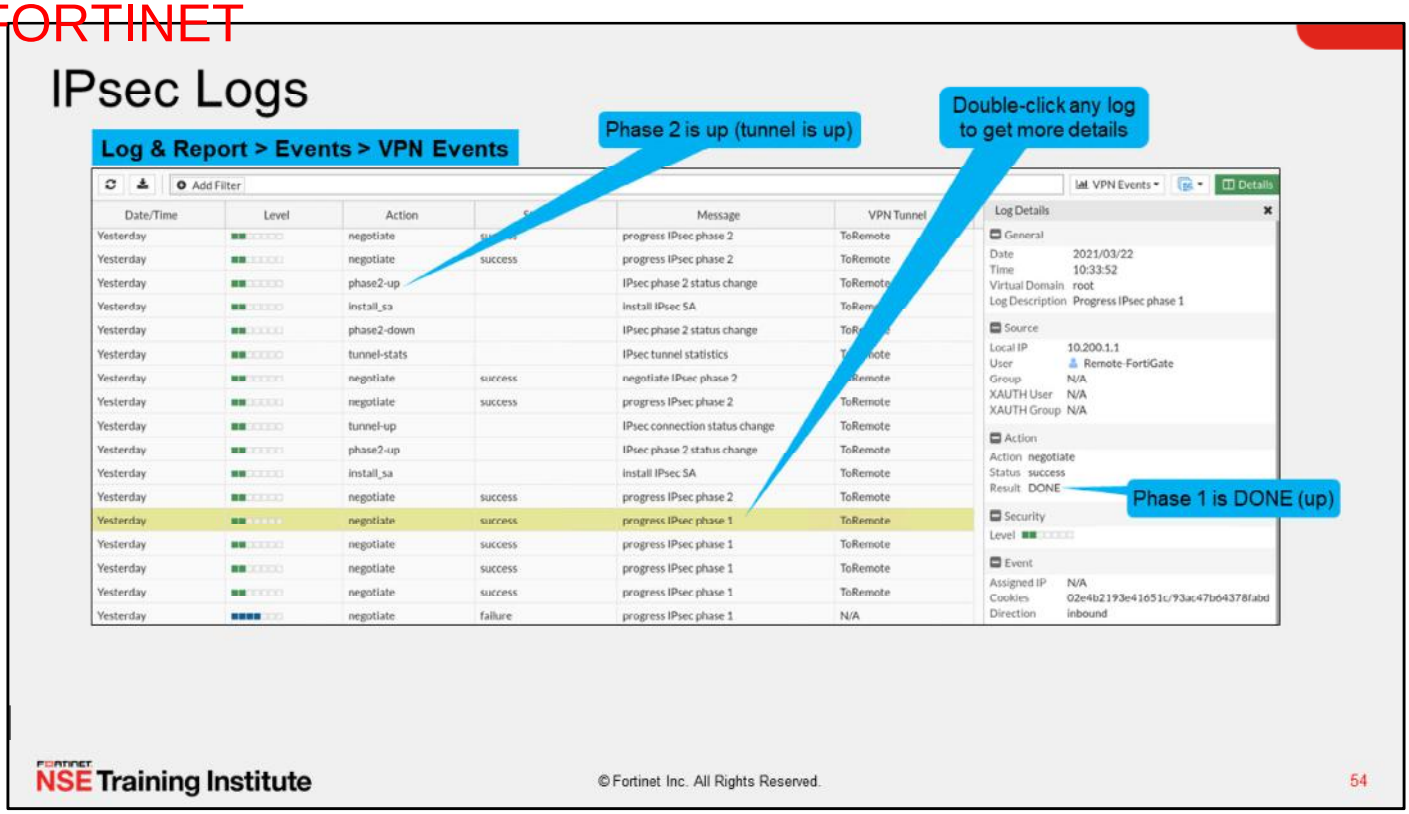

FortiGate logs IPsec VPN events by default. To view IPsec VPN event logs, click **Log & Report** > **Events** > **VPN Events**.

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 _____ status of an IPsec VPN.
   - A. Phase 1
   - ✓ B. Phase 2

2. When the remote gateway is set to dialup 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

NSE Training Institute

55

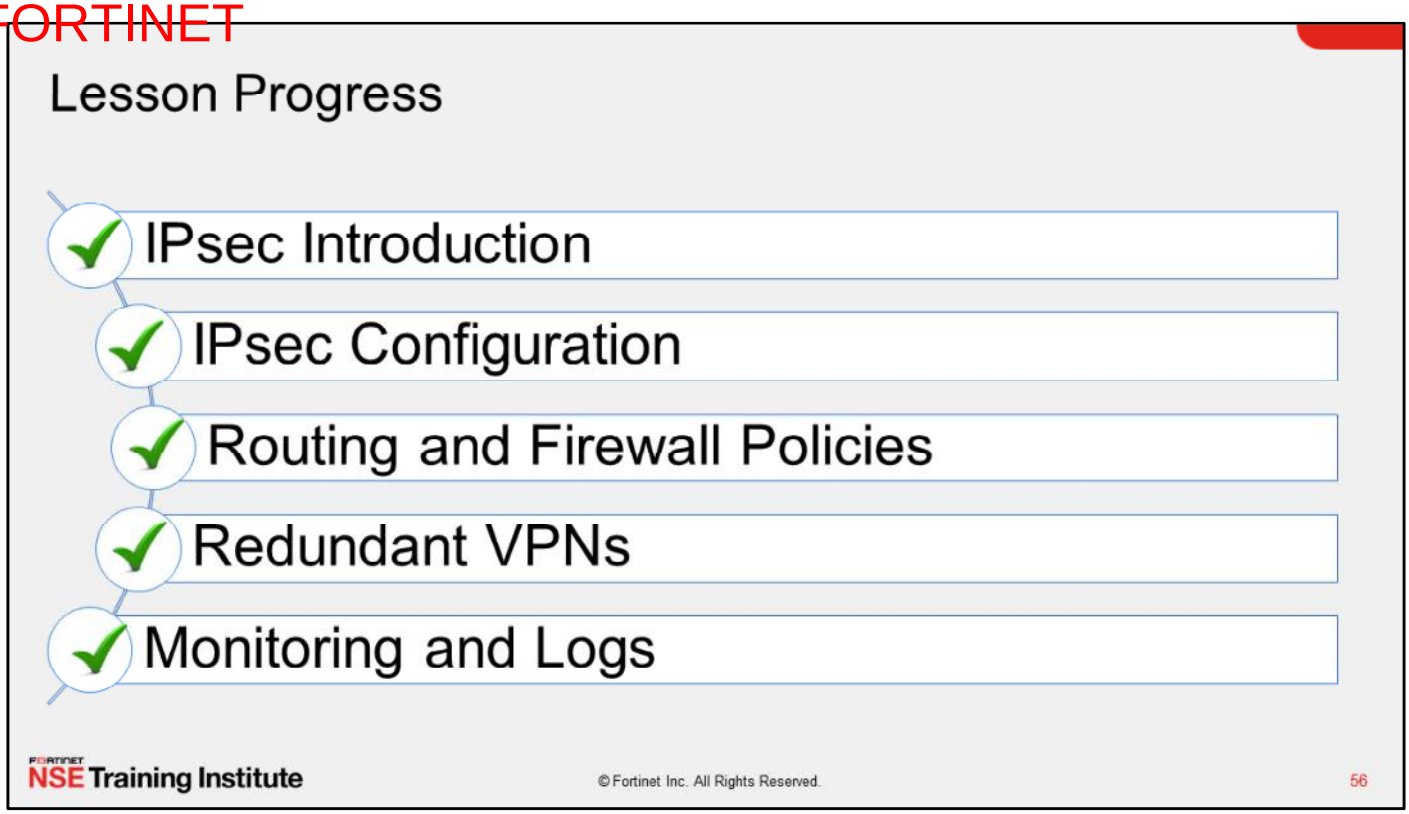

Congratulations! You have completed this lesson.

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

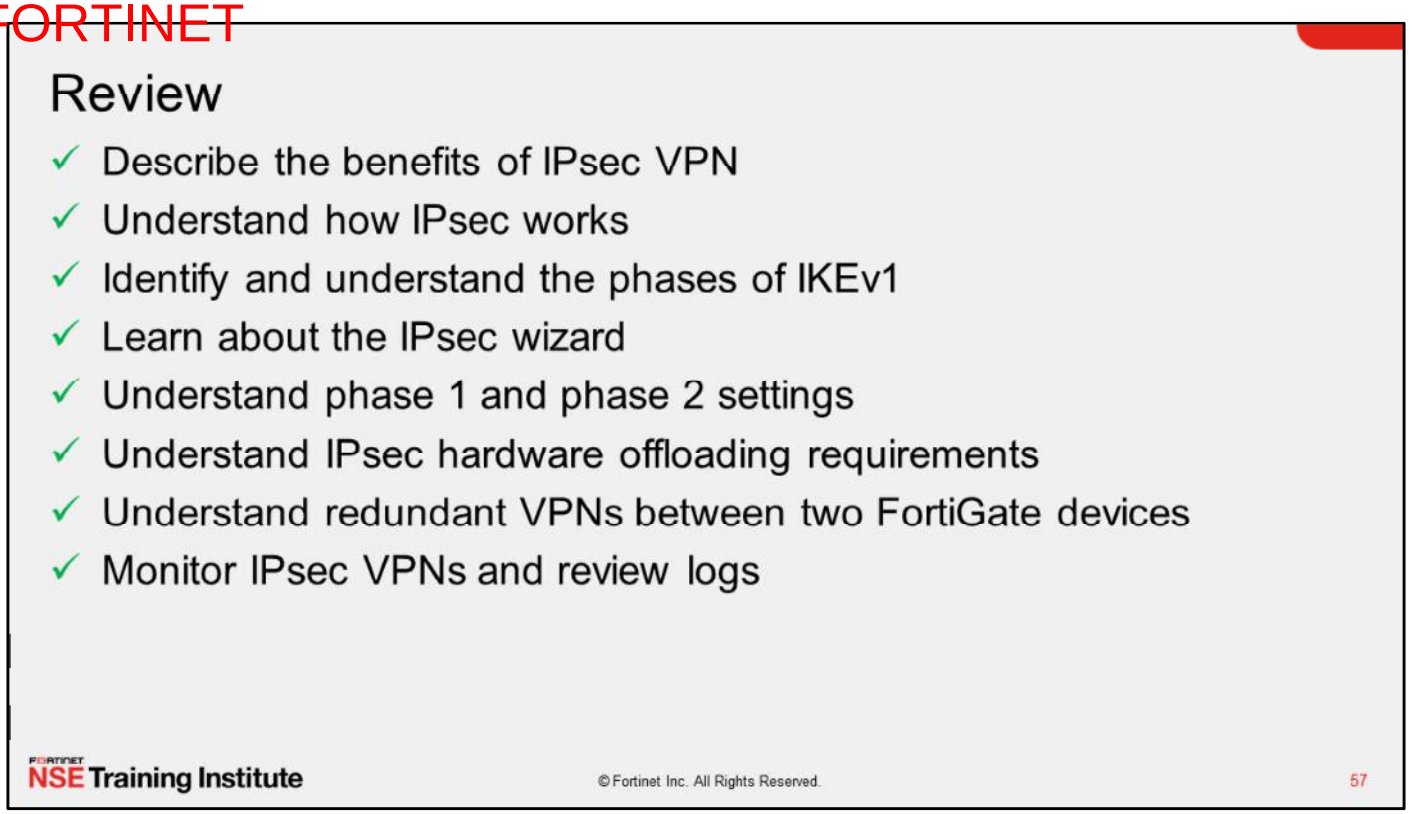

This slide shows the objectives that you covered in this lesson.

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.

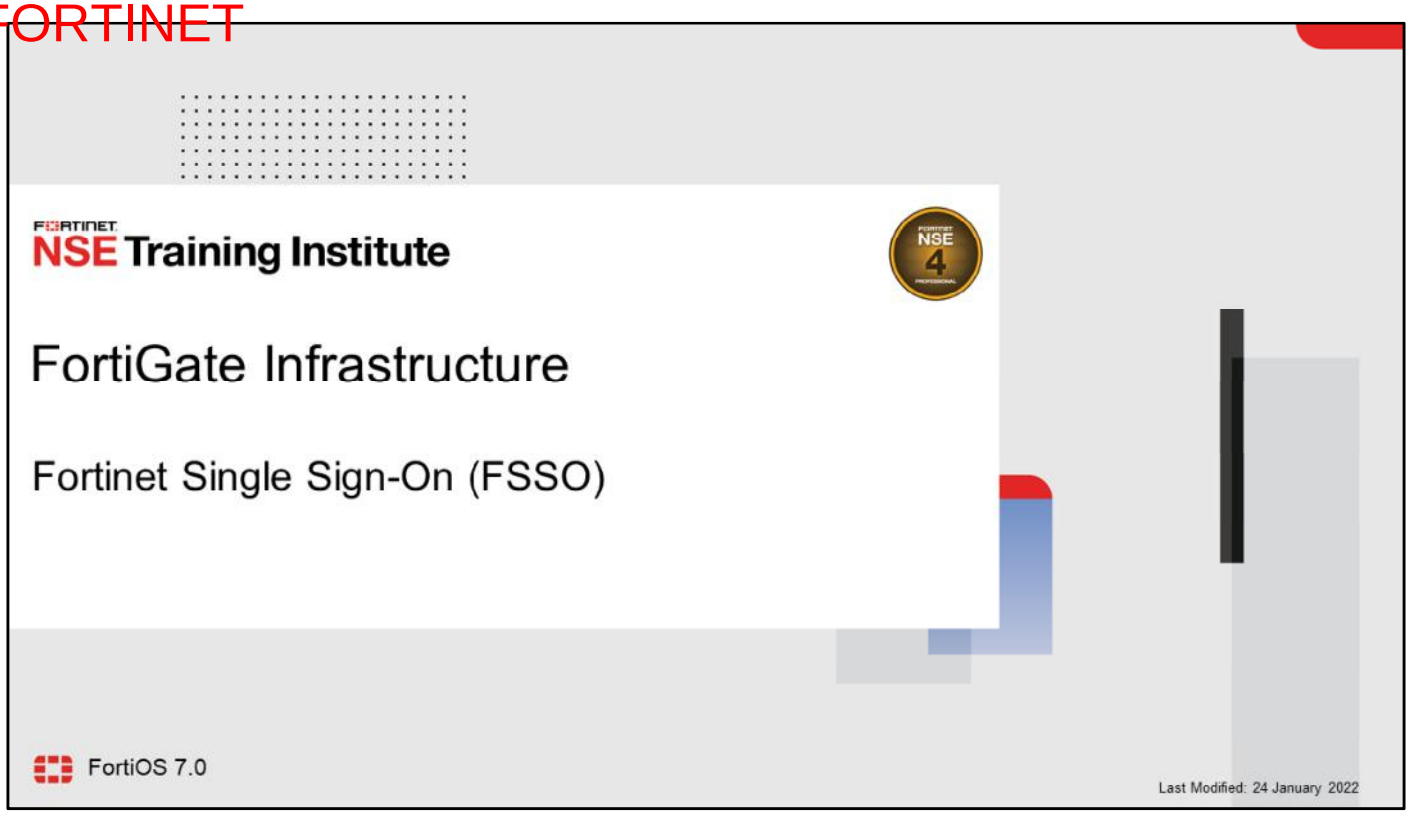

In this lesson, you will learn about Fortinet single sign-on (FSSO). When you use this feature, your users don't need to log on each time they access a different network resource.

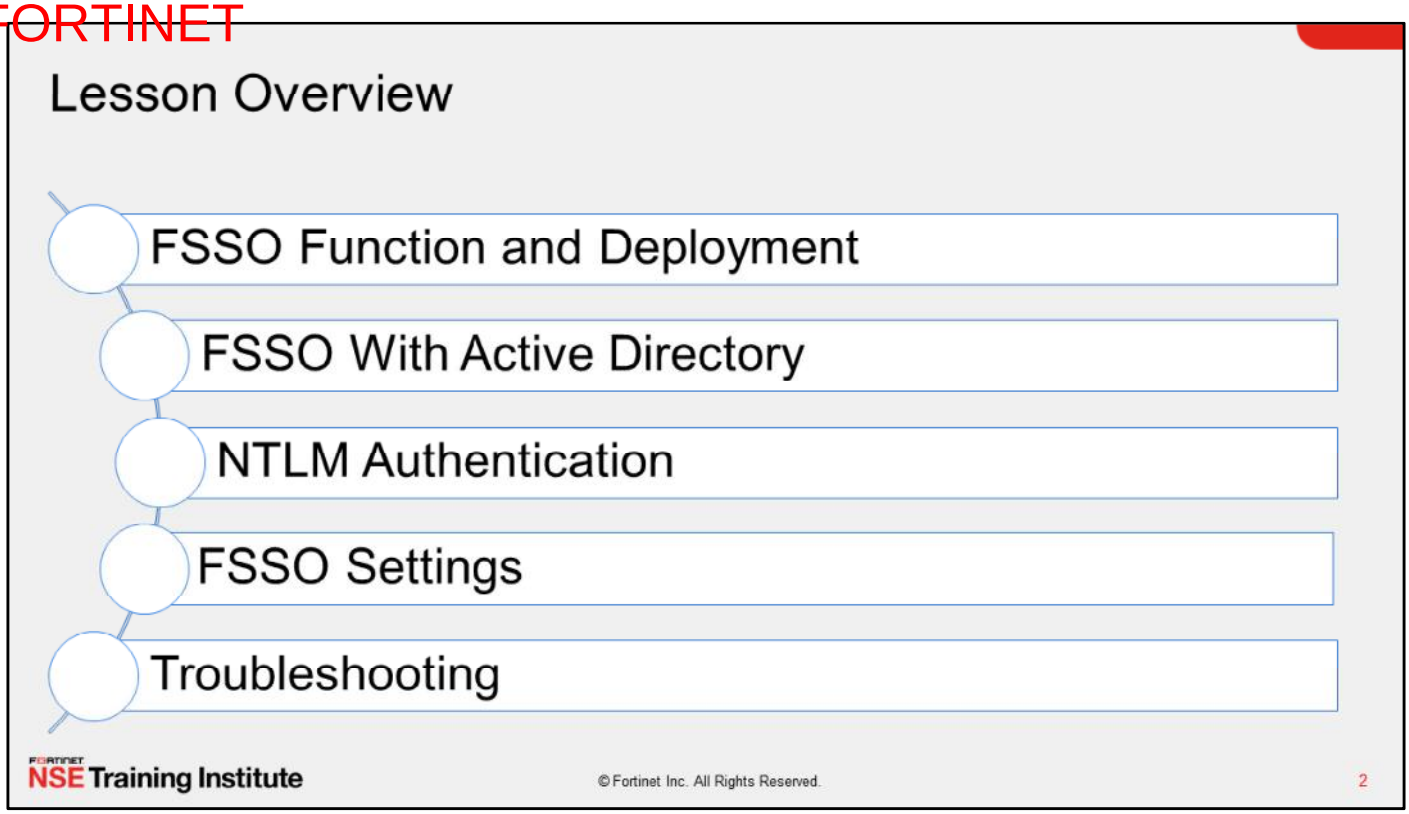

In this lesson, you will learn about the topics shown on this slide.

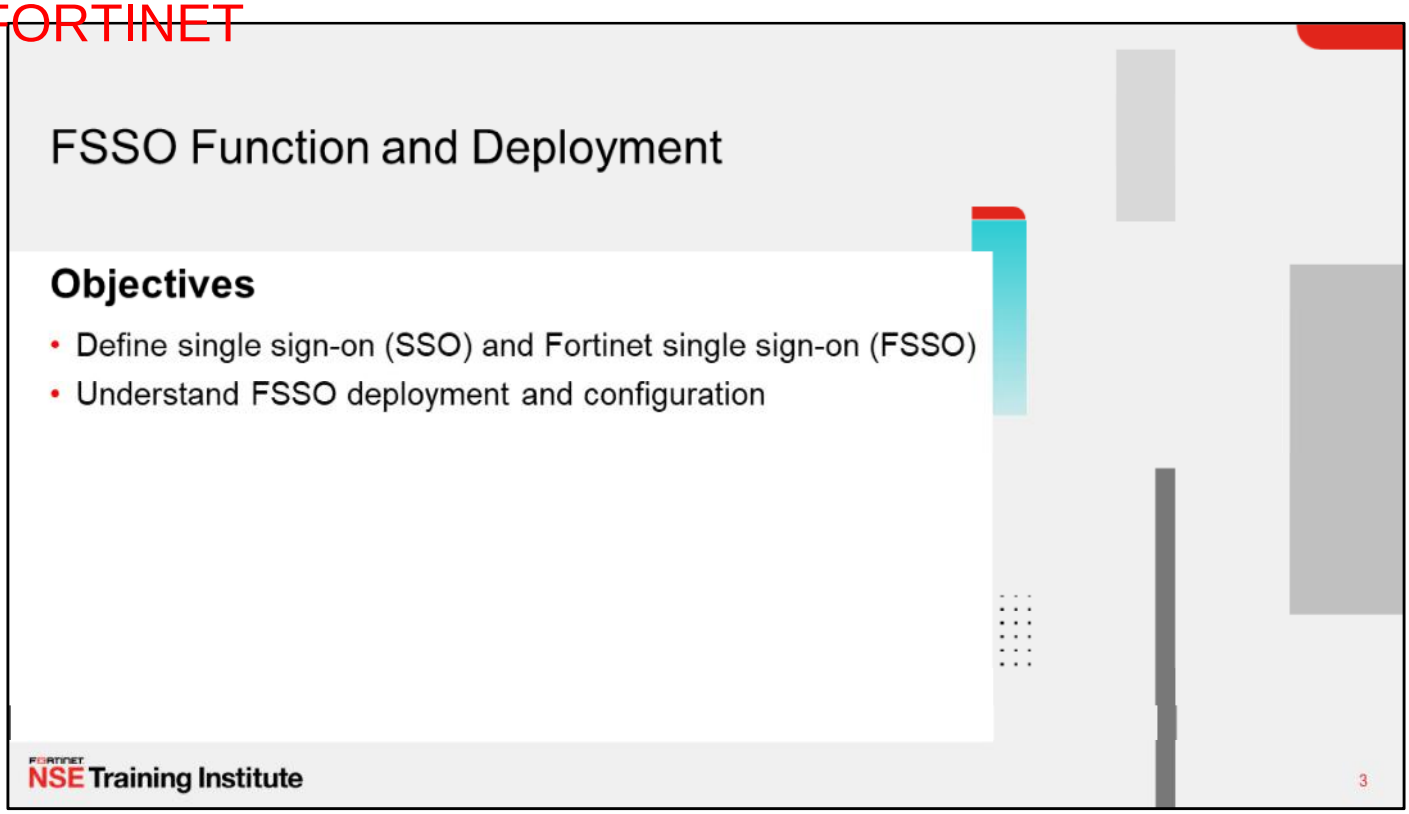

After completing this section, you should be able to achieve the objectives shown on this slide.

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

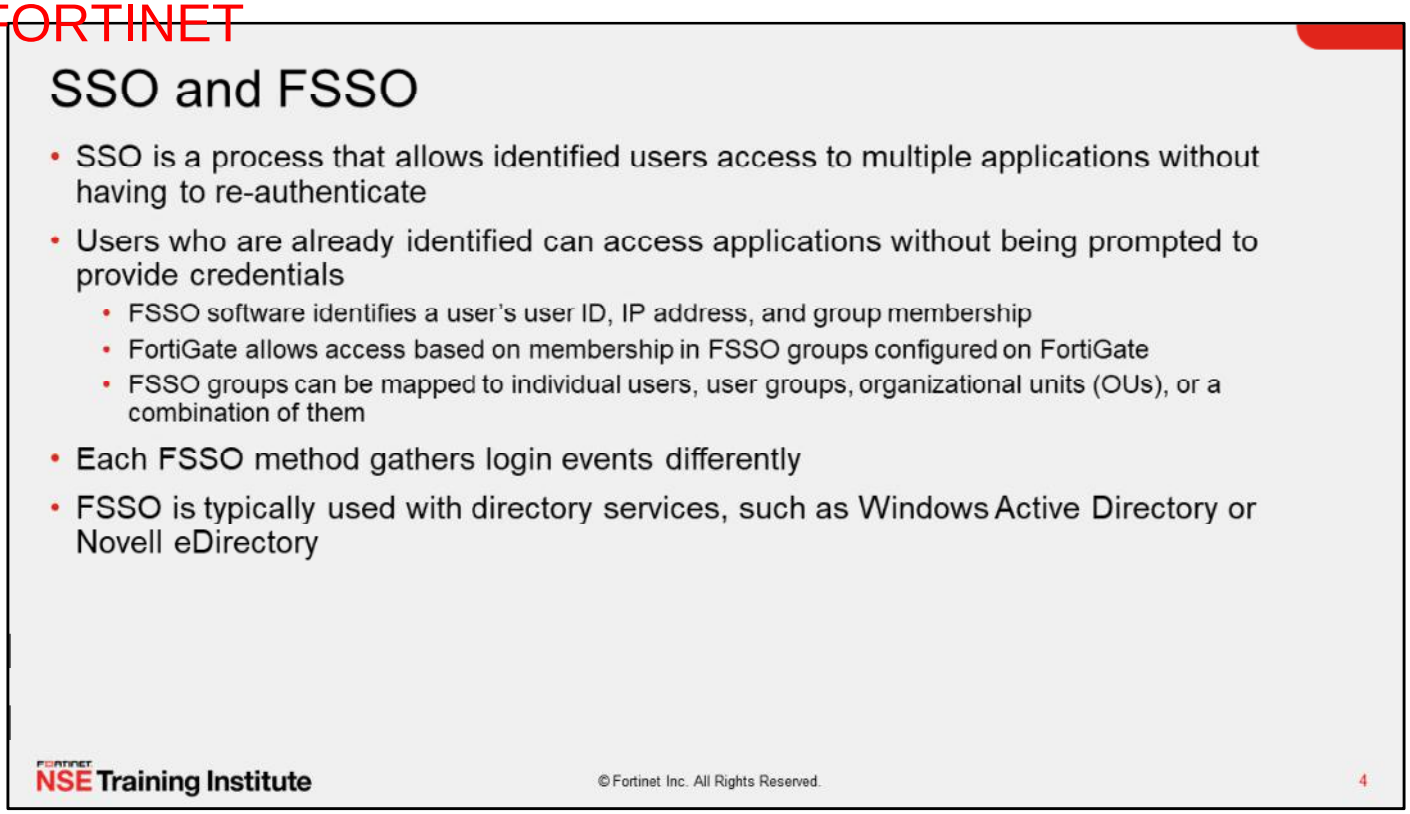

SSO is a process that allows users to be automatically logged in to every application after being identified, regardless of platform, technology, and domain.

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.

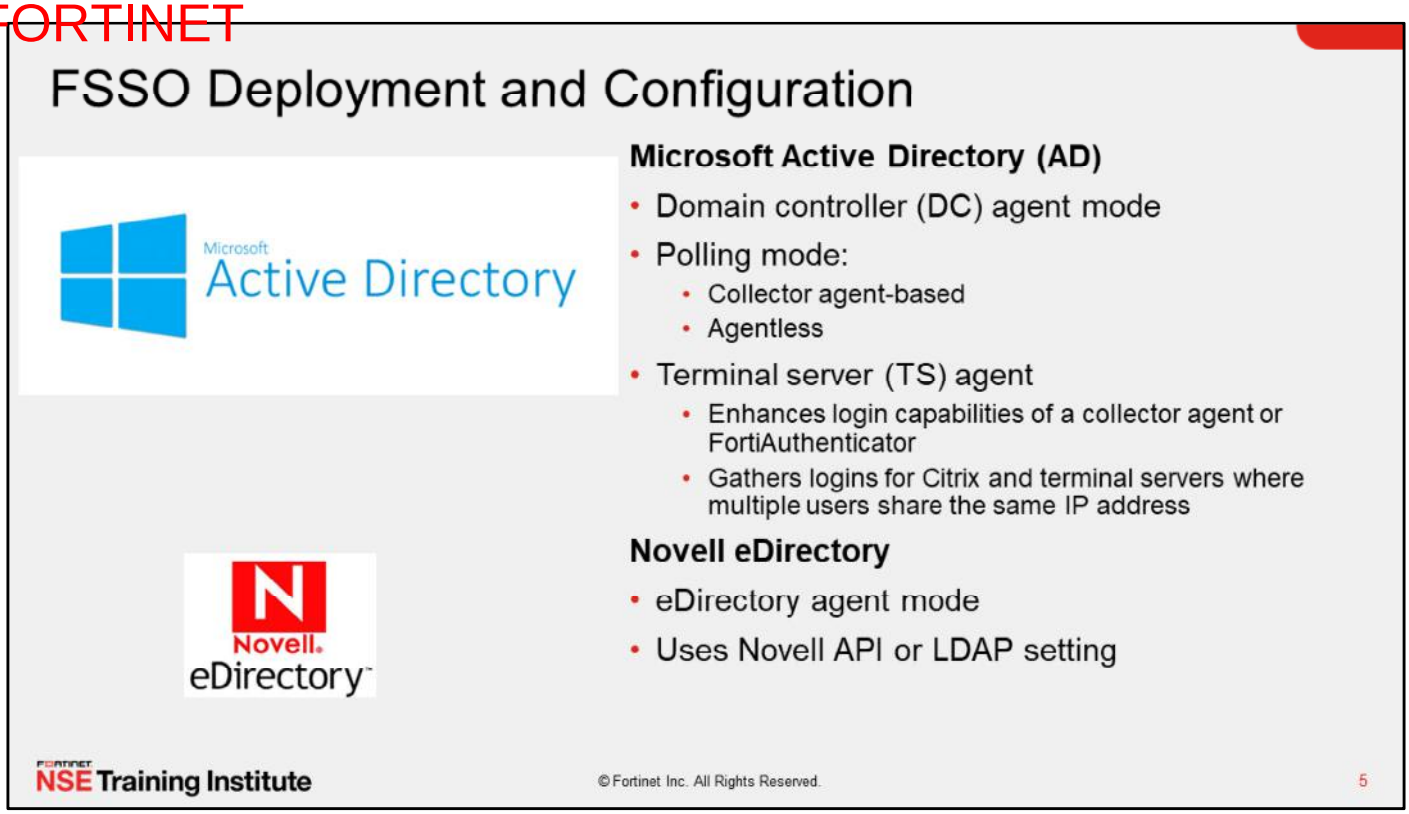

How you deploy and configure FSSO depends on the server that provides your directory services.

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.

Knowledge Check

1. In FSSO, FortiGate allows network access based on _____.
   - A. Active user authentication with username and password
   - ✓ B. Passive user identification by user ID, IP address, and group membership

2. Which working mode is used for monitoring user sign-on activities in Windows AD?
   - ✓ A. Polling mode (collector agent-based or agentless)
   - B. eDirectory agent mode

NSE Training Institute

6

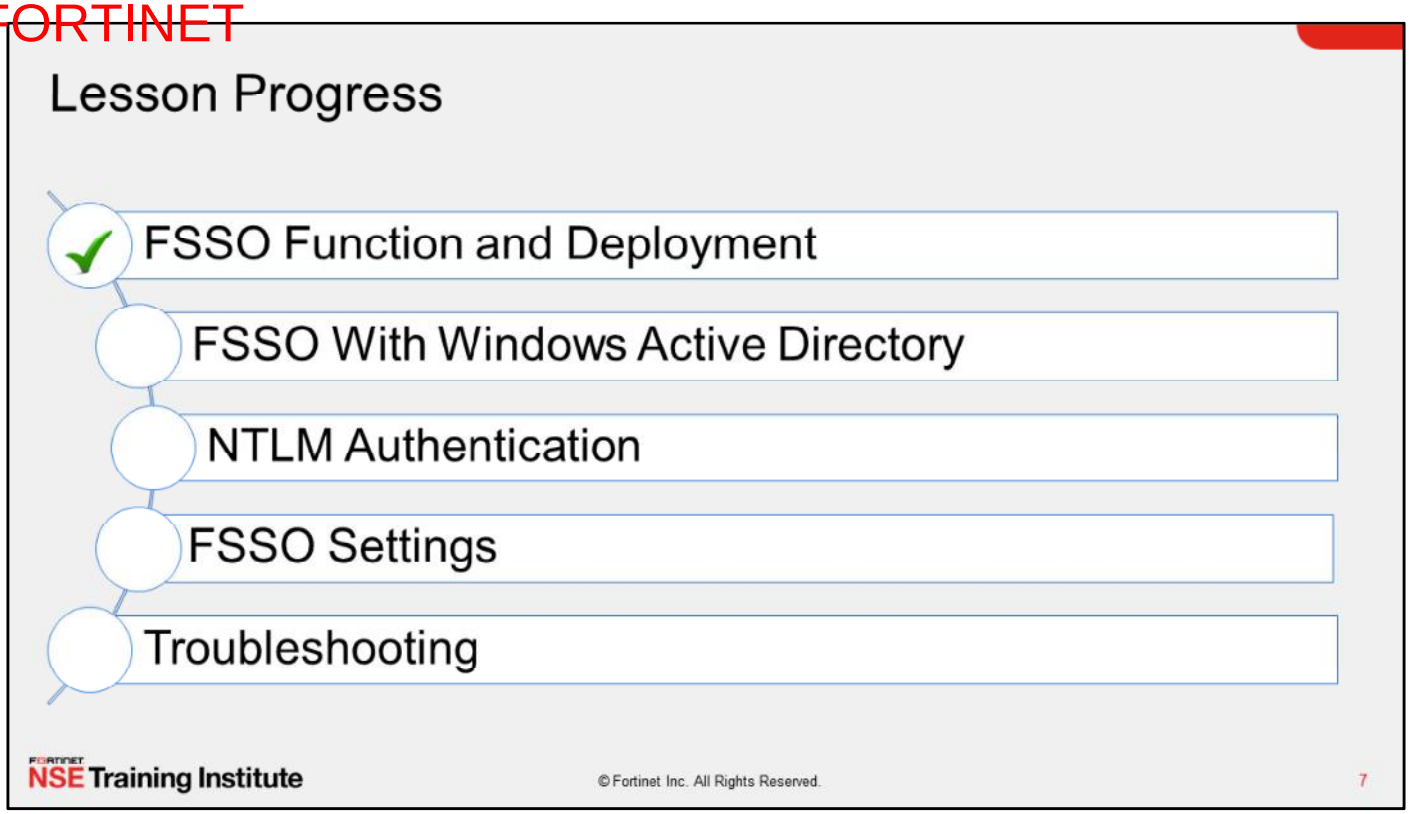

Good job! You now understand basic concepts about the function of FSSO and how it is deployed.

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.

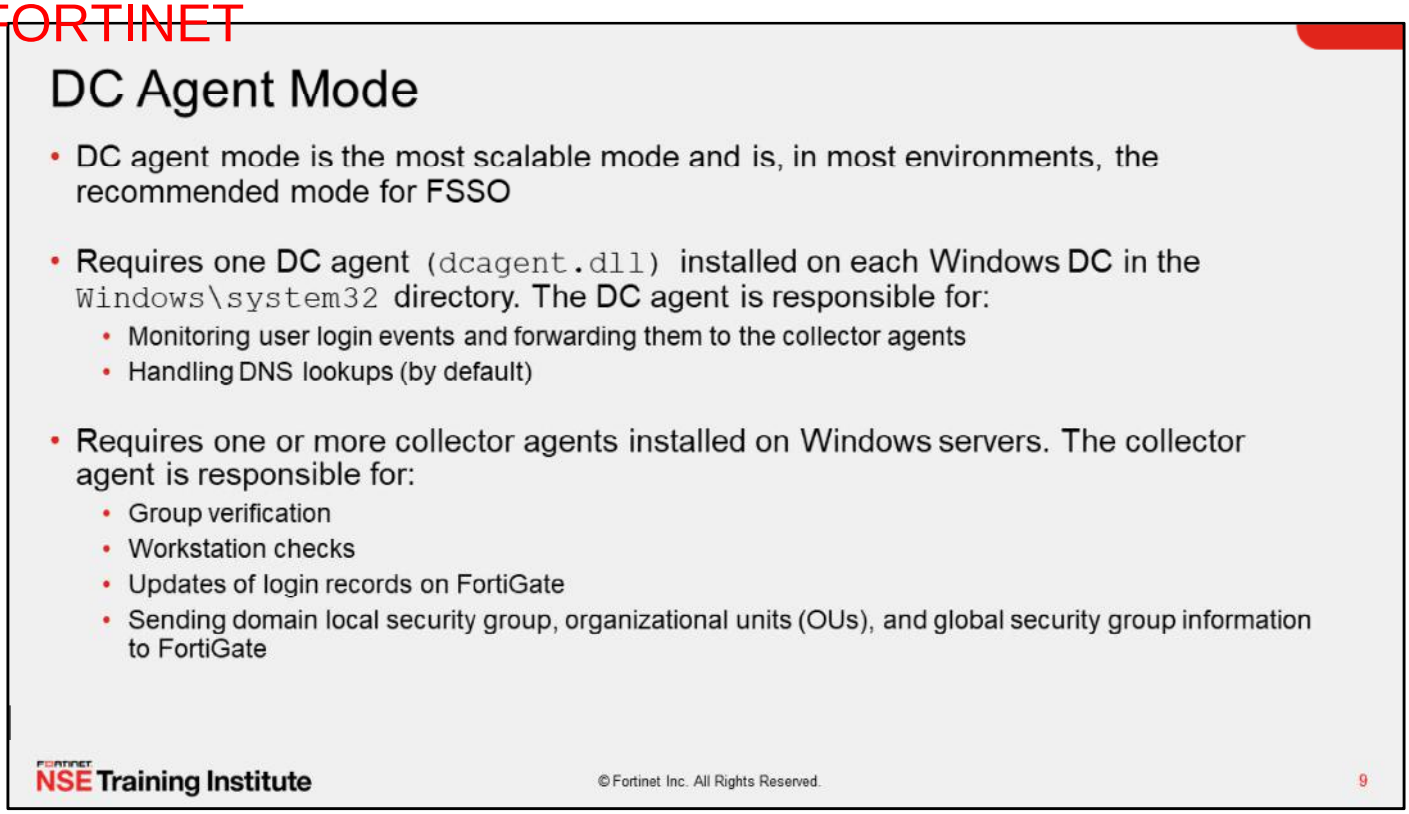

DC agent mode is considered the recommended mode for FSSO.

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`

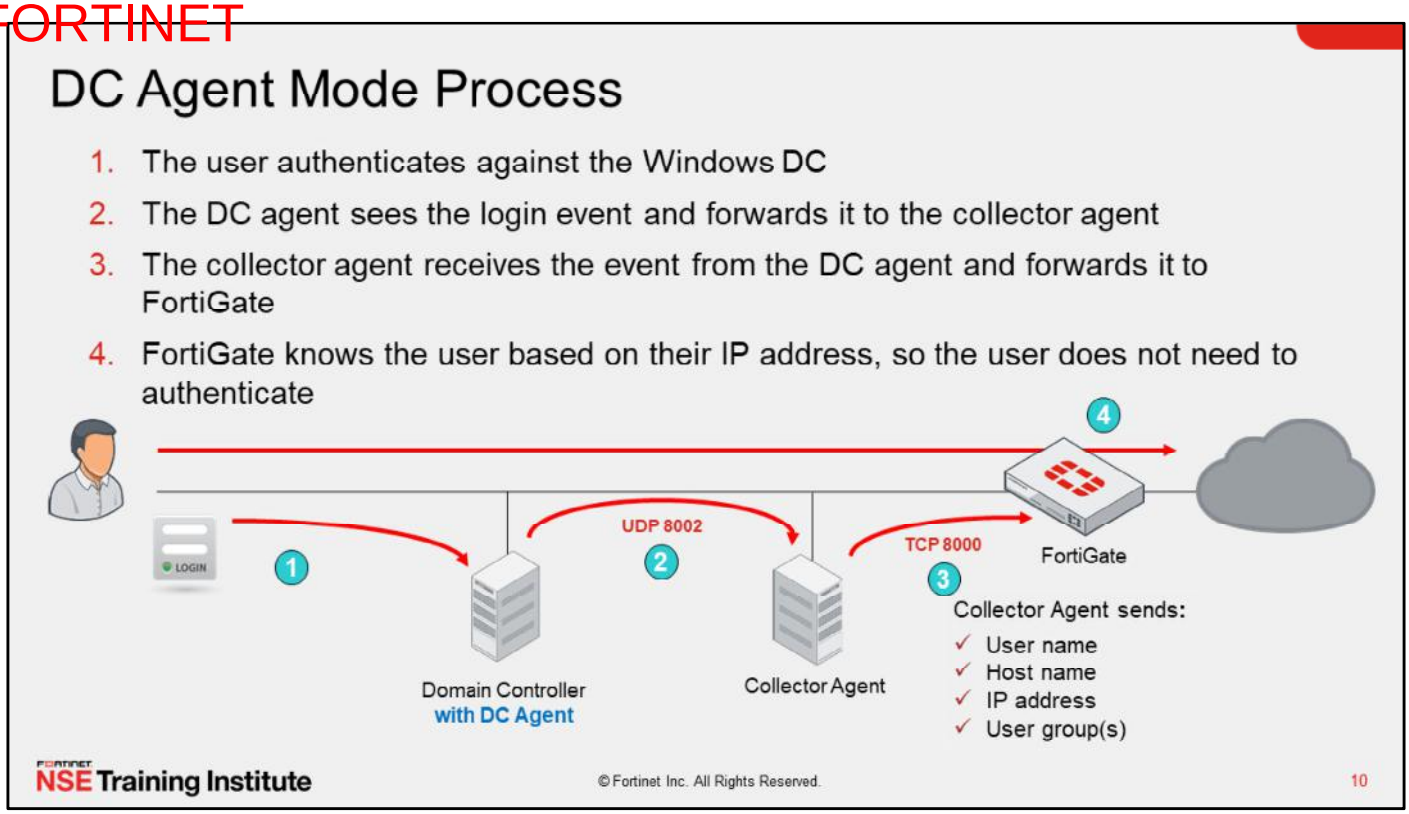

This slide shows the process of information passing between DC agents, the collector agent, and a FortiGate configured for FSSO authentication.

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.

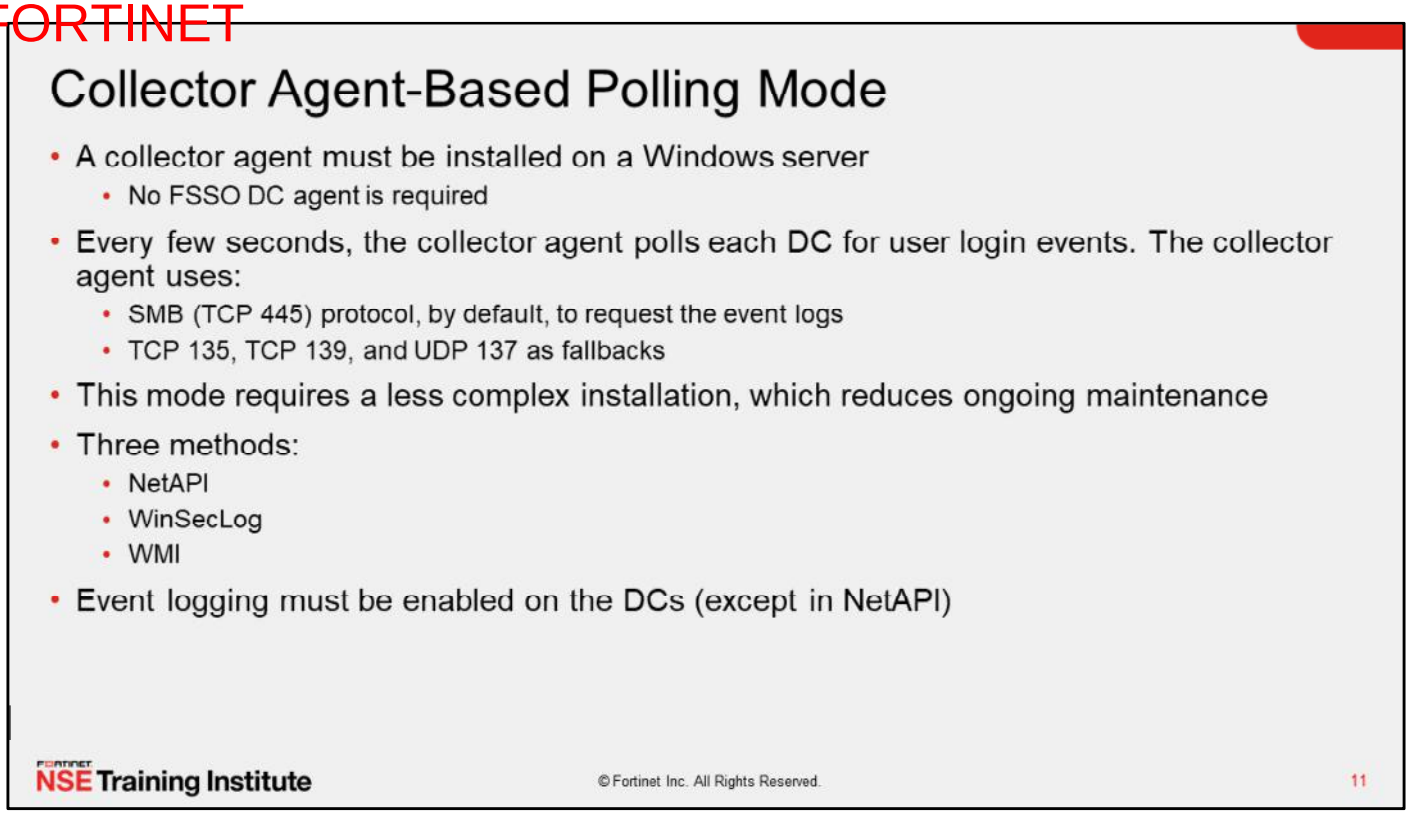

Polling mode can be collector agent-based or agentless.

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.

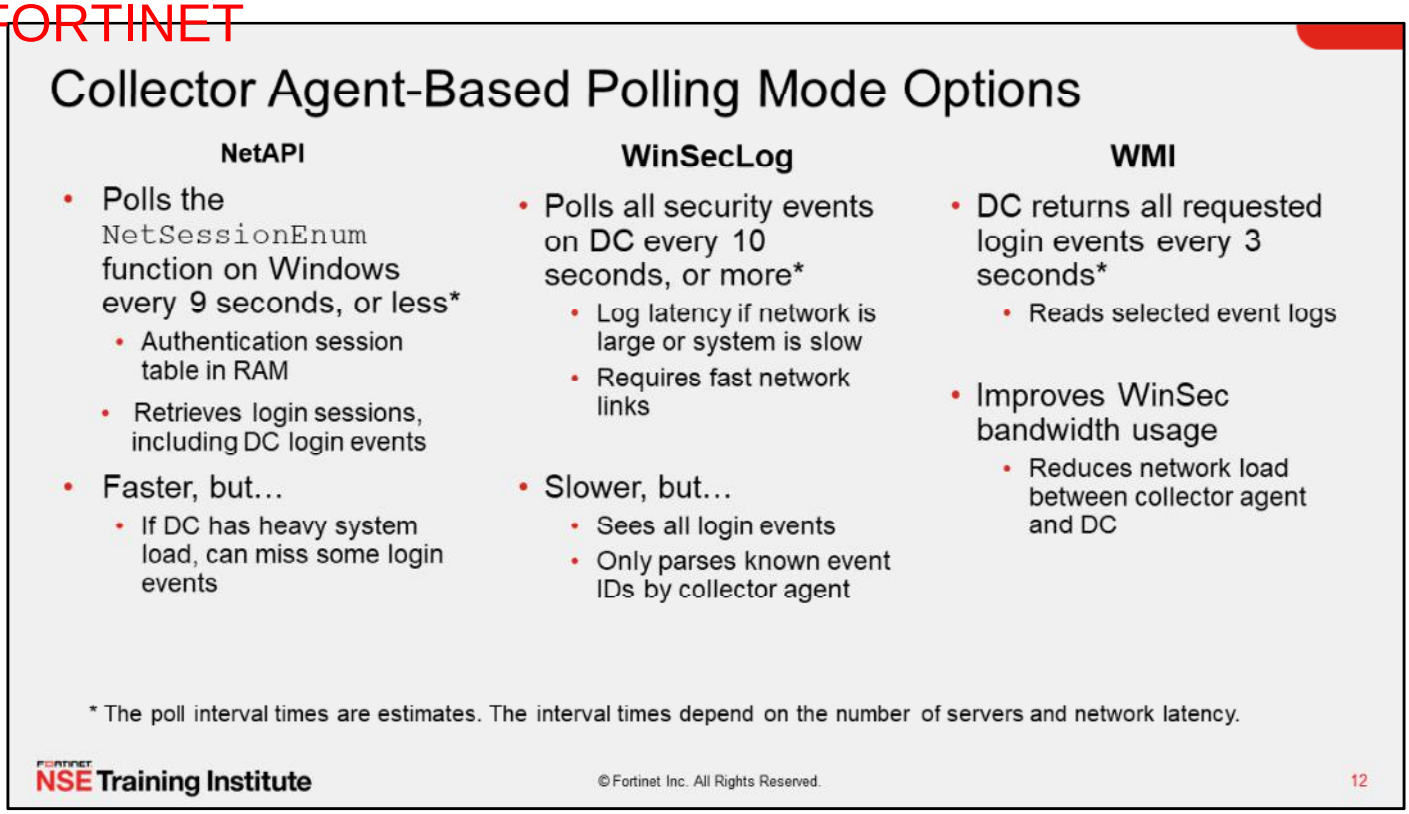

As previously stated, collector agent-based polling mode has three methods (or options) for collecting login information:

- **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.

- **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).

- **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.

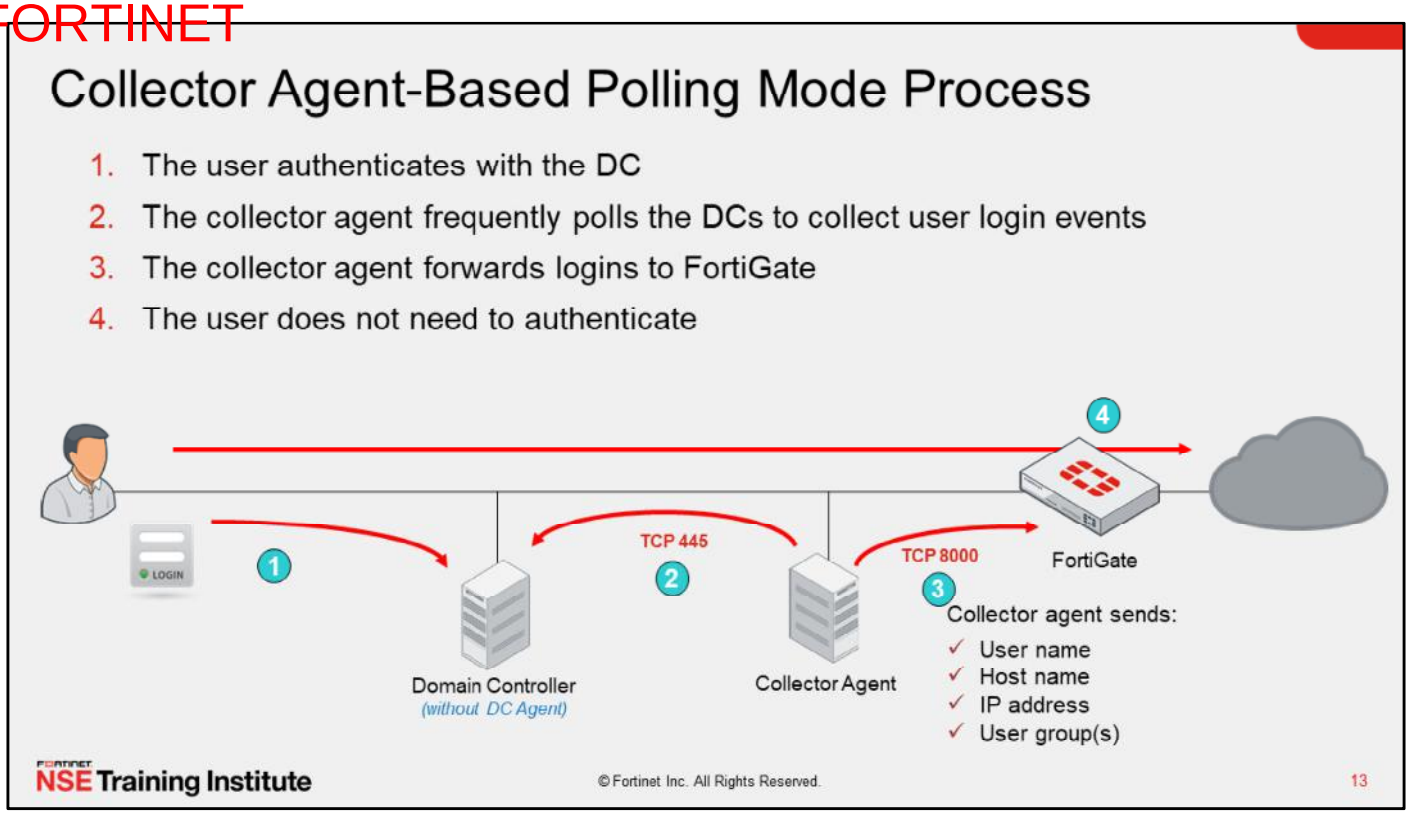

This slide shows an example of FSSO using the collector agent-based polling mode. This example includes a DC, a collector agent, and FortiGate, but the DC doesn't have the dcagent (or, alternatively, dcagent.dll) installed.

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.

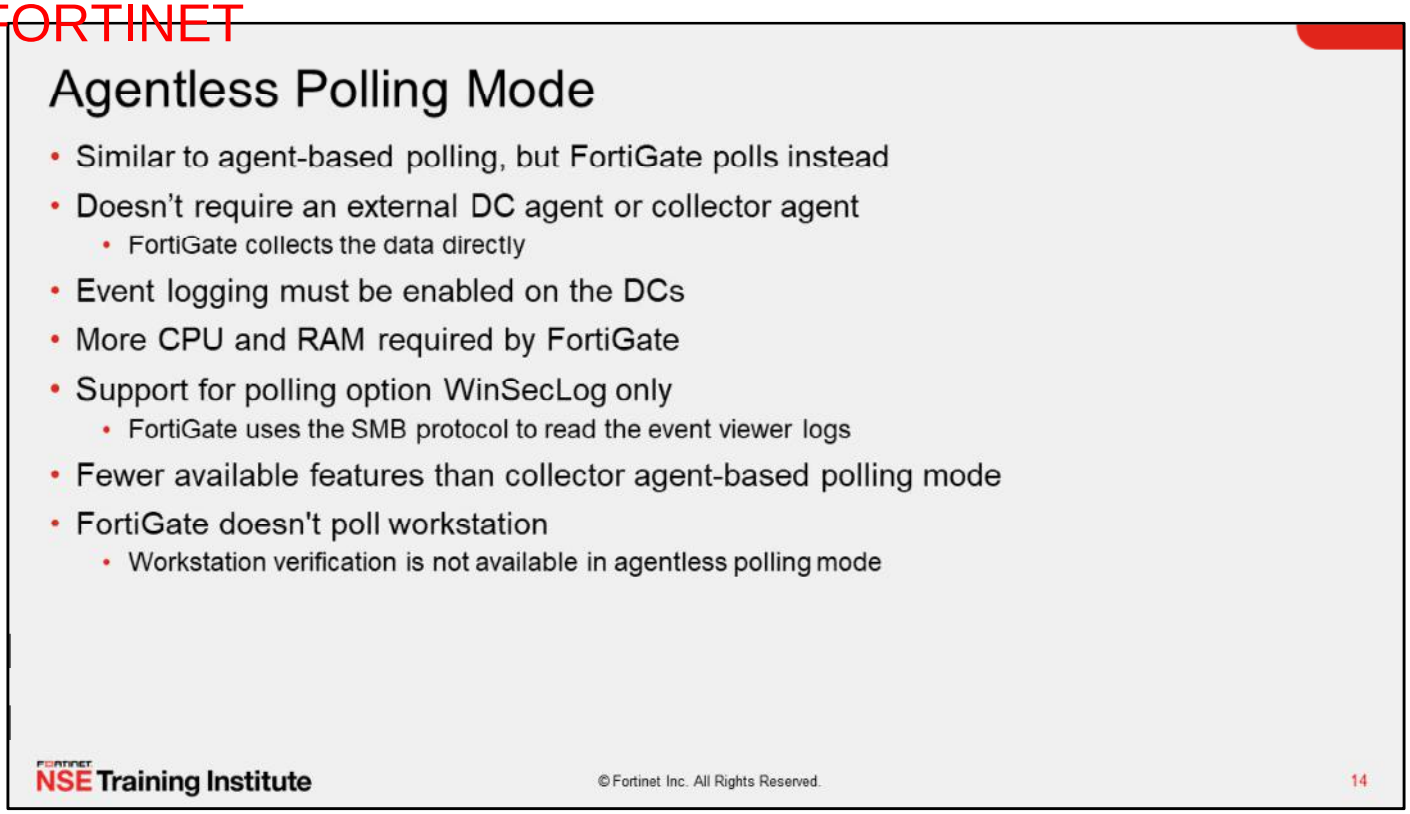

You can deploy FSSO without installing an agent. FortiGate polls the DCs directly, instead of receiving login information indirectly from a collector agent.

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.

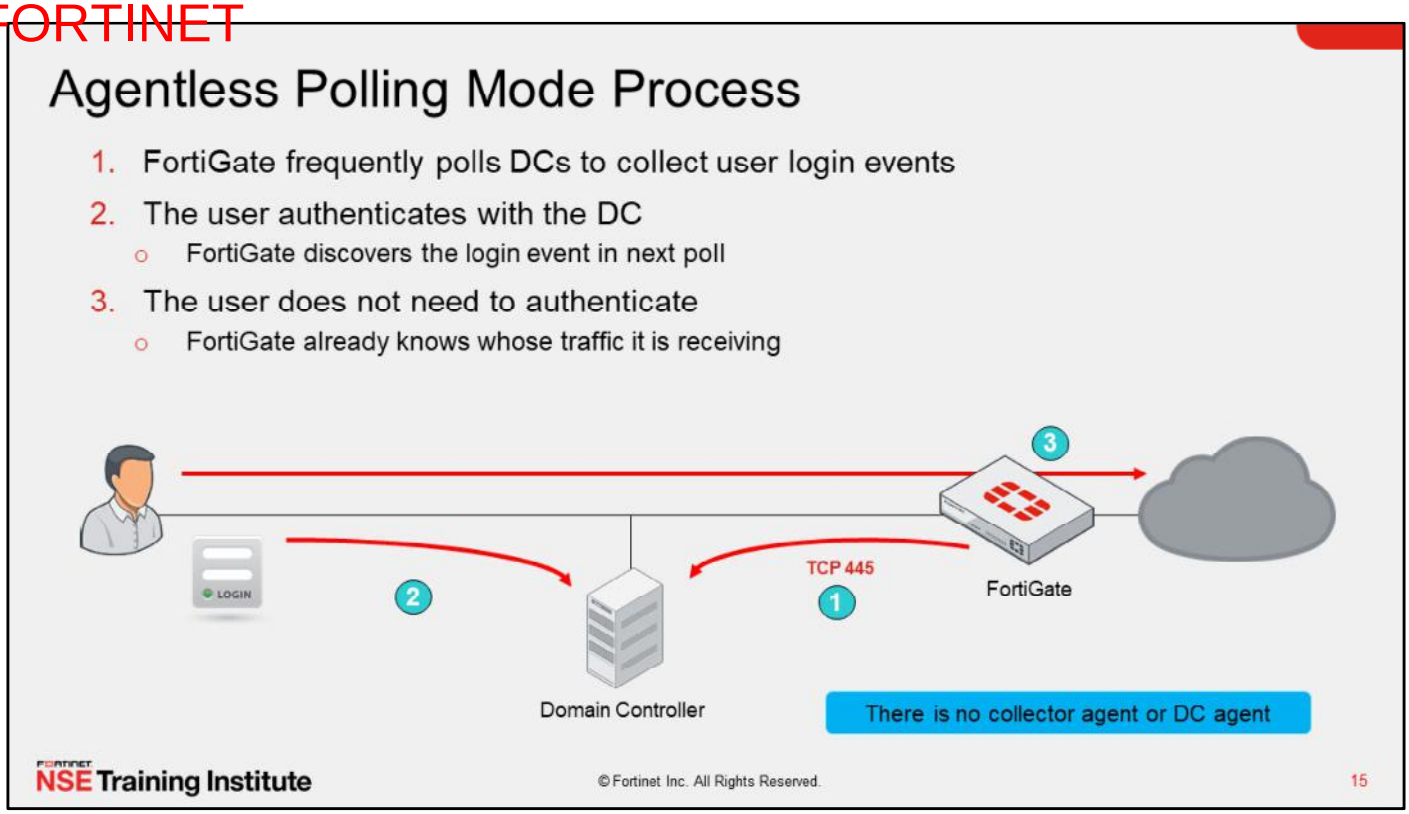

This slide shows how communication is processed without agents. (There is no collector agent or DC agent.)

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.

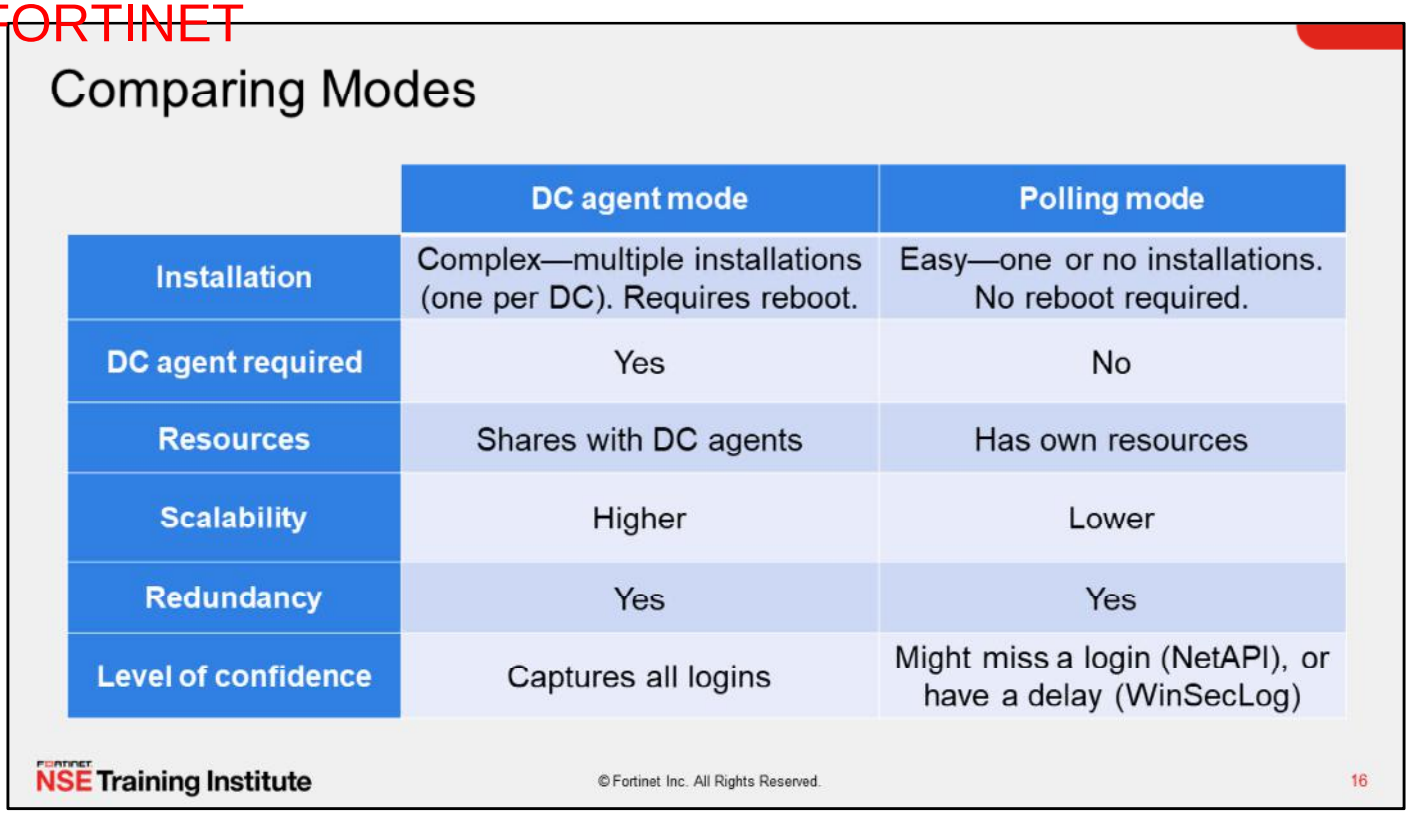

This table summarizes the main differences between DC agent mode and polling mode.

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.

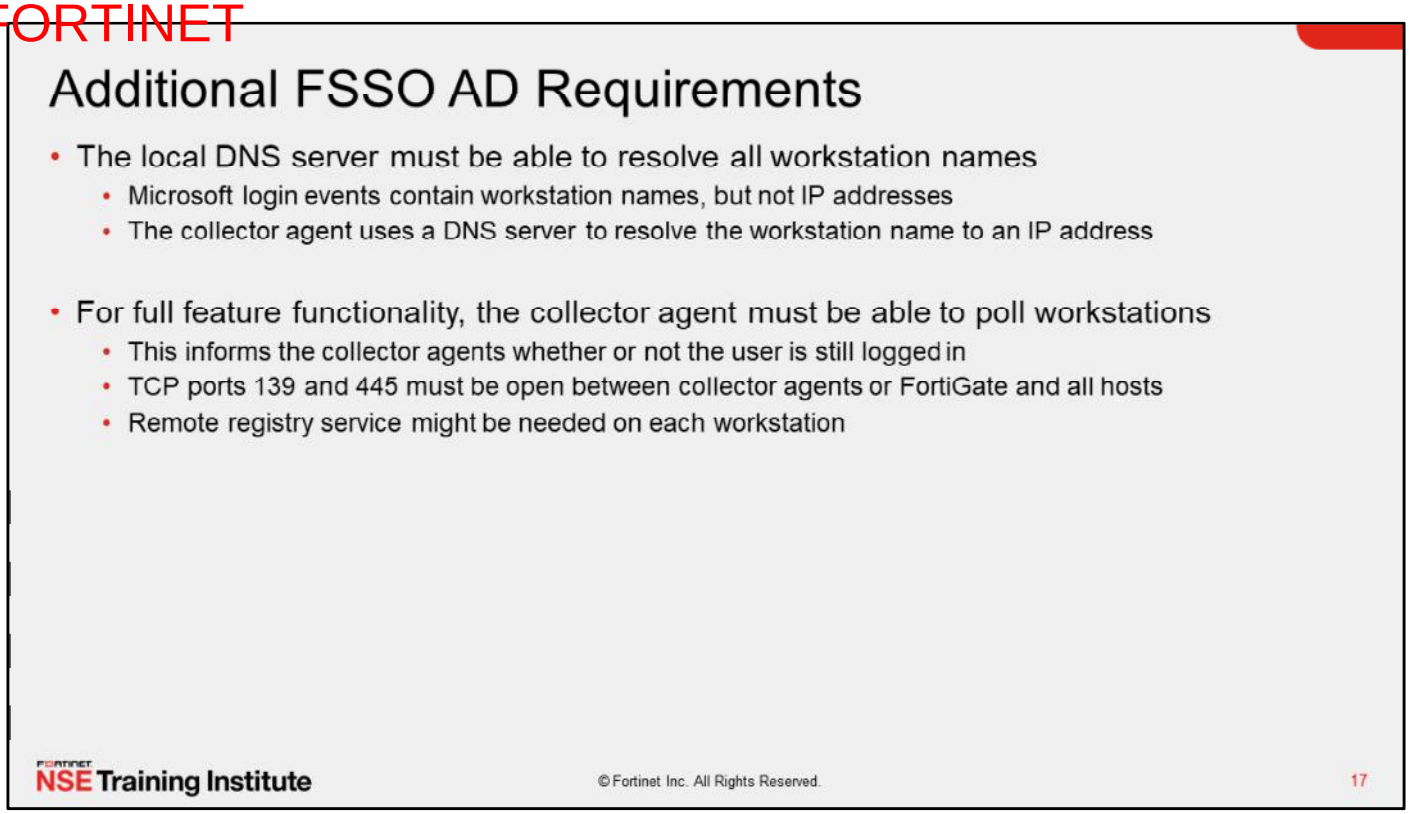

Regardless of the collector method you choose, some FSSO requirements for your AD network are the same:

- 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

NSE Training Institute

18

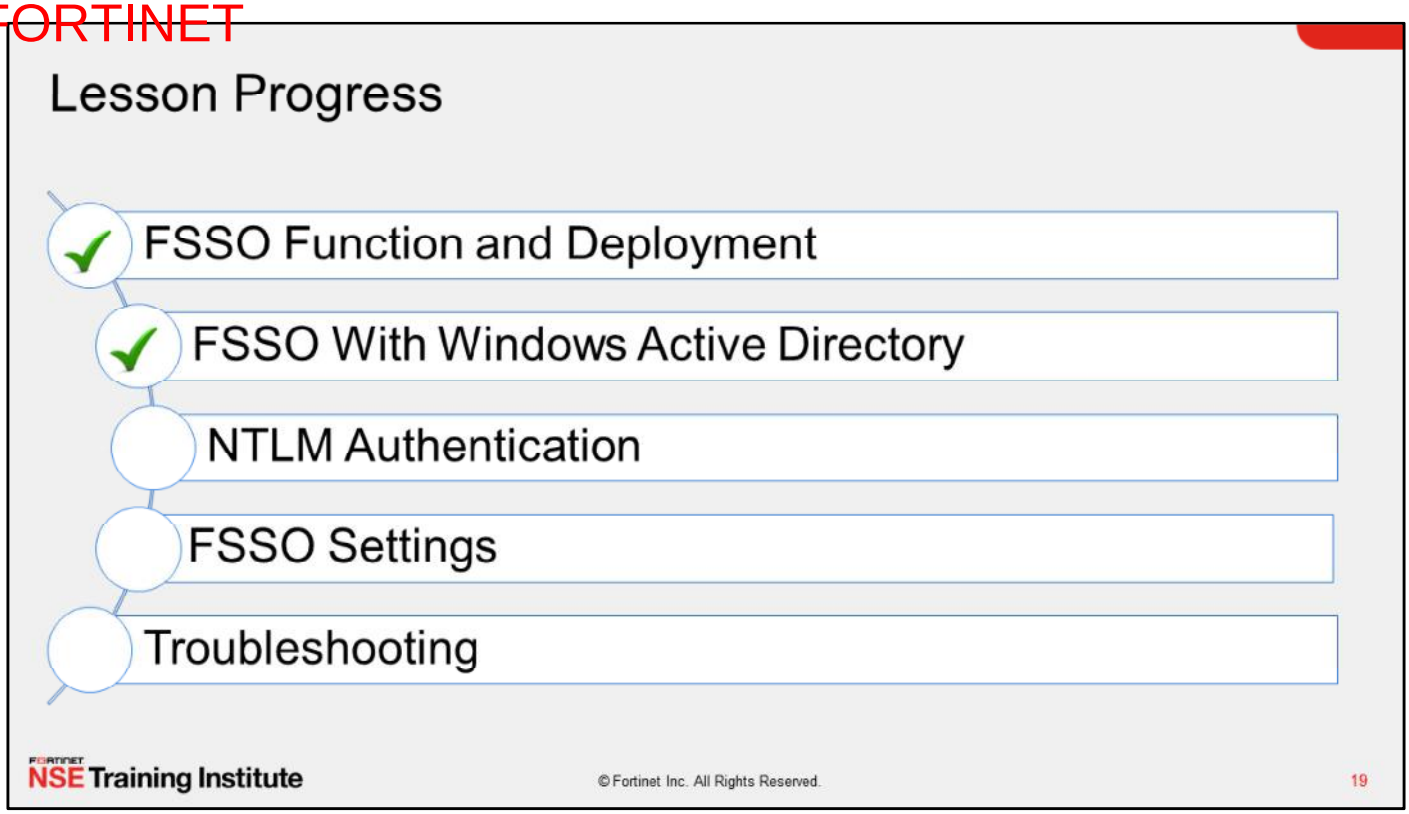

Good job! You now understand how FortiGate detects login events in Windows Active Directory (AD) using FSSO.

Now, you'll learn how NT LAN manager (NTLM) works and interacts with FSSO for a web-initiated login.

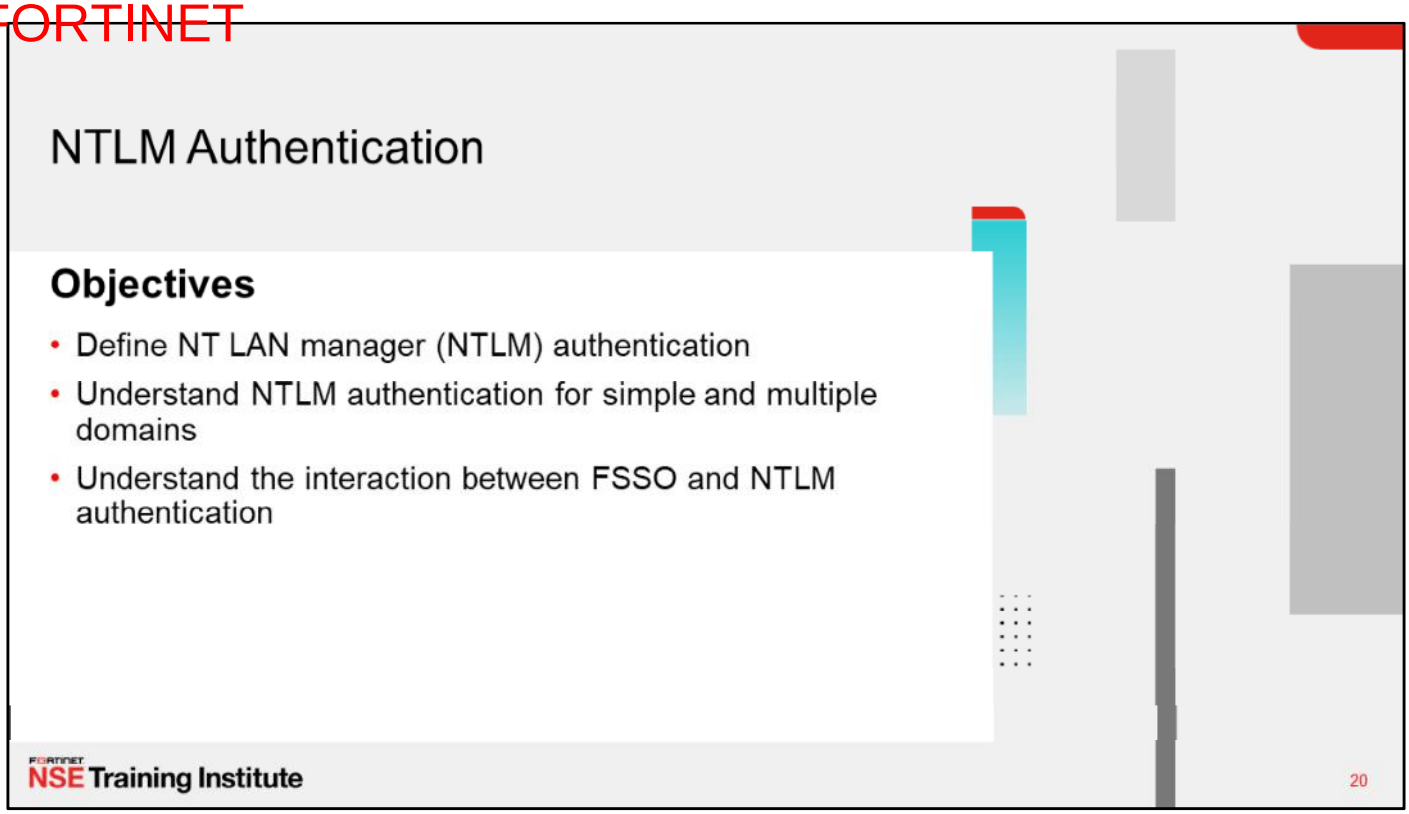

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding NTLM authentication and its interaction with FSSO, you will be able to configure a transparent web-initiated login session with NTLM authentication.

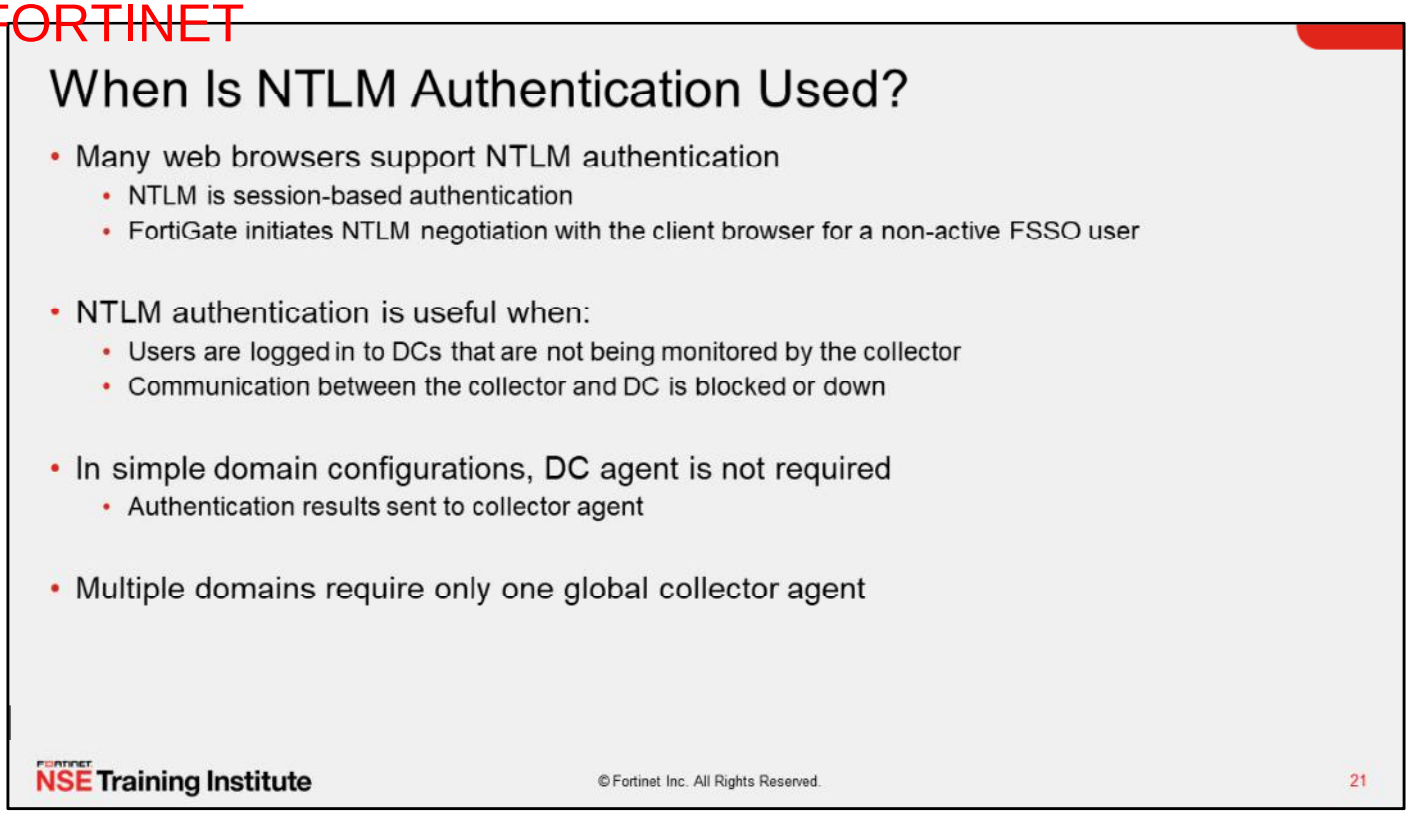

In an AD environment, FSSO can also work with NTLM, which is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users.

NTLM authentication does not require DC agents, but it is not fully invisible to users; users must enter their credentials during NTLM negotiation. NTLM authentication is a Microsoft-proprietary solution, so it can be implemented only in a Windows network.

NTLM is most often used when users authenticate against DCs that, for some reason, can’t be monitored by the collector agent, or when there are communication problems between the collector agent and one or more of the DC agents. In other words, NTLM authentication is best used as a backup to FSSO.

Notice that FortiGate cannot do NTLM authentication on its own. FortiGate needs to pass the user-entered credentials back to a collector agent for verification. The collector agent, in turn, responds to FortiGate with the appropriate user's groups if the authentication is successful.

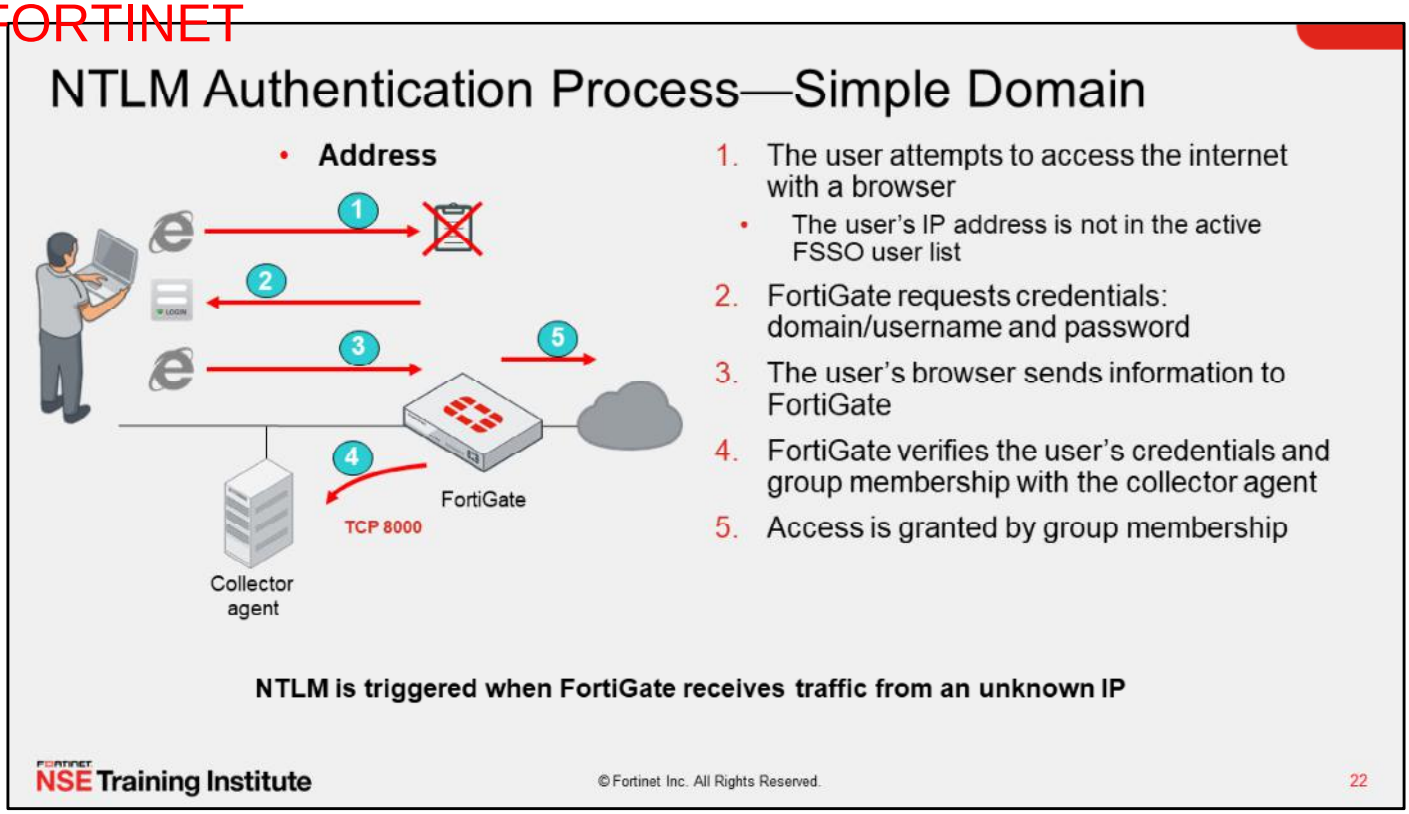

This slide shows how messages are processed during NTLM authentication in a simple domain configuration.

1. When both FSSO and NTLM are enabled, NTLM is used as a fallback for FSSO. When FortiGate receives traffic from an IP address that doesn't exist in the FSSO user list, NTLM is triggered.
2. FortiGate replies with an NTLM challenge, requesting credentials.
3. The user's browser sends the requested credentials.
4. FortiGate receives the user's credentials, then authenticates them with the collector agent over TCP port 8000. FortiGate also receives the names of the groups that the user belongs to.
5. If the credentials are correct, FortiGate authorizes access for the user.

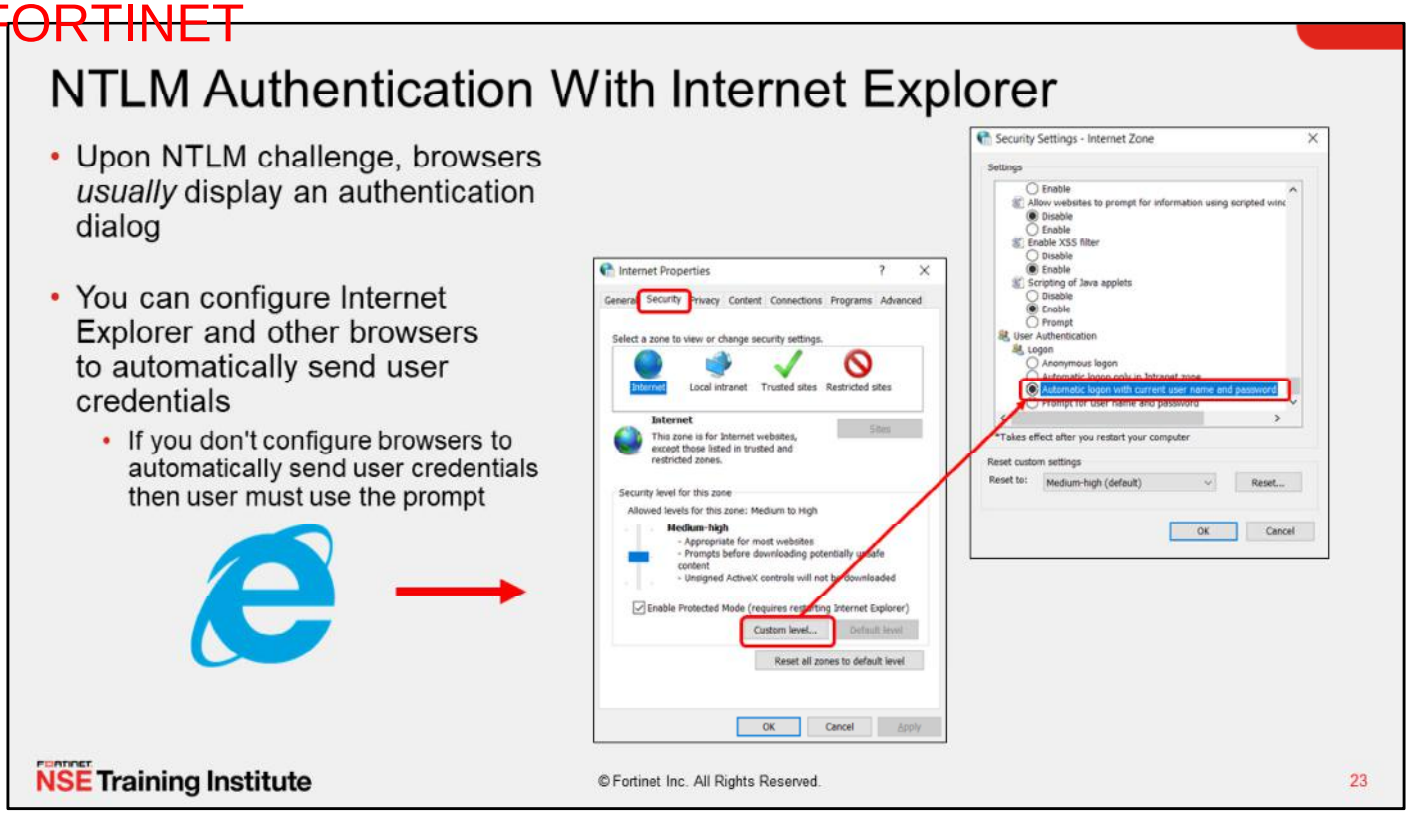

Unlike full FSSO, NTLM authentication is not transparent to users. In most browsers and, by default, in Internet Explorer, users must enter their credentials whenever the browser receives an NTLM authentication challenge.

However, you can configure Internet Explorer to automatically send the user's credentials each time it receives an NTLM challenge. To do this, in the **Internet Options** dialog, click **Custom level**. Then, in the **Settings** dialog, scroll to **User Authentication login**, and then select **Automatic login with current user name and password**.

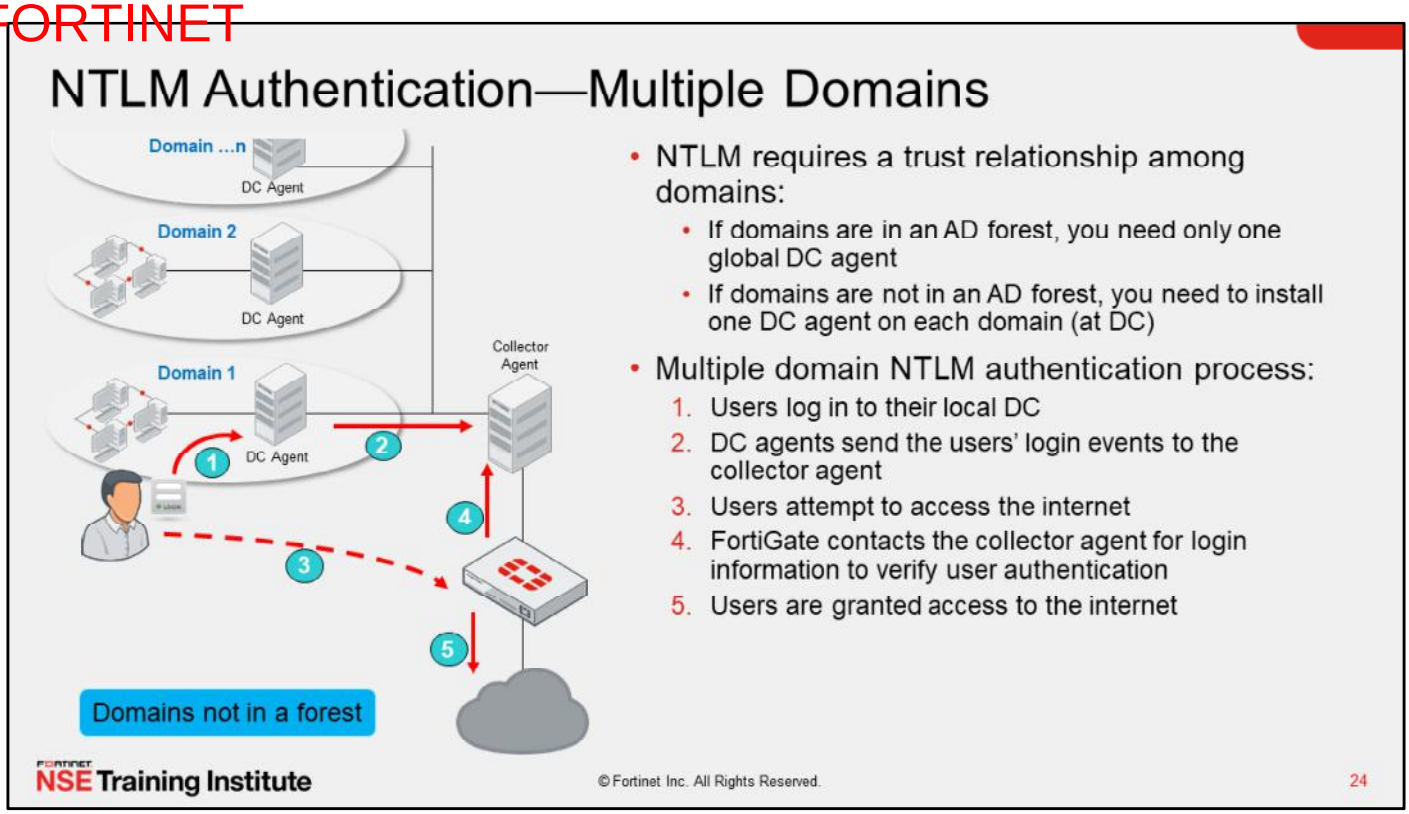

In a multiple domain environment for NTLM, it's important to have a trust relationship between the domains. When multiple domains exist in an AD forest, a trust relationship is automatically created, so only one DC agent is required on one of the domain controllers. But, when multiple domains are not in an AD forest, you have two options:

- Create a trust relationship between the domains through AD settings
- Install one DC agent on each domain, then use security policies for network access

If you decide to install one DC agent on each domain, the DC agent sends login information to the collector agent. This process works as follows:

1. The user logs in to their local DC.
2. The DC agent sends the user login event information to the collector agent.
3. The user attempts to access the internet.
4. FortiGate verifies that the user is authenticated by contacting the collector agent for the login information.
5. If the user is correctly authenticated, FortiGate allows them to access the internet.

Knowledge Check

1. When performing NTLM authentication, what information does the web browser supply to FortiGate?
   - ✓ A. The user's credentials (username and password)
   - B. The user's user ID, IP address, and group membership

2. What may cause an NTLM authentication to occur?
   - A. Traffic coming from an IP on the FSSO user list
   - ✓ B. Traffic coming from an IP not on the FSSO user list

NSE Training Institute

25

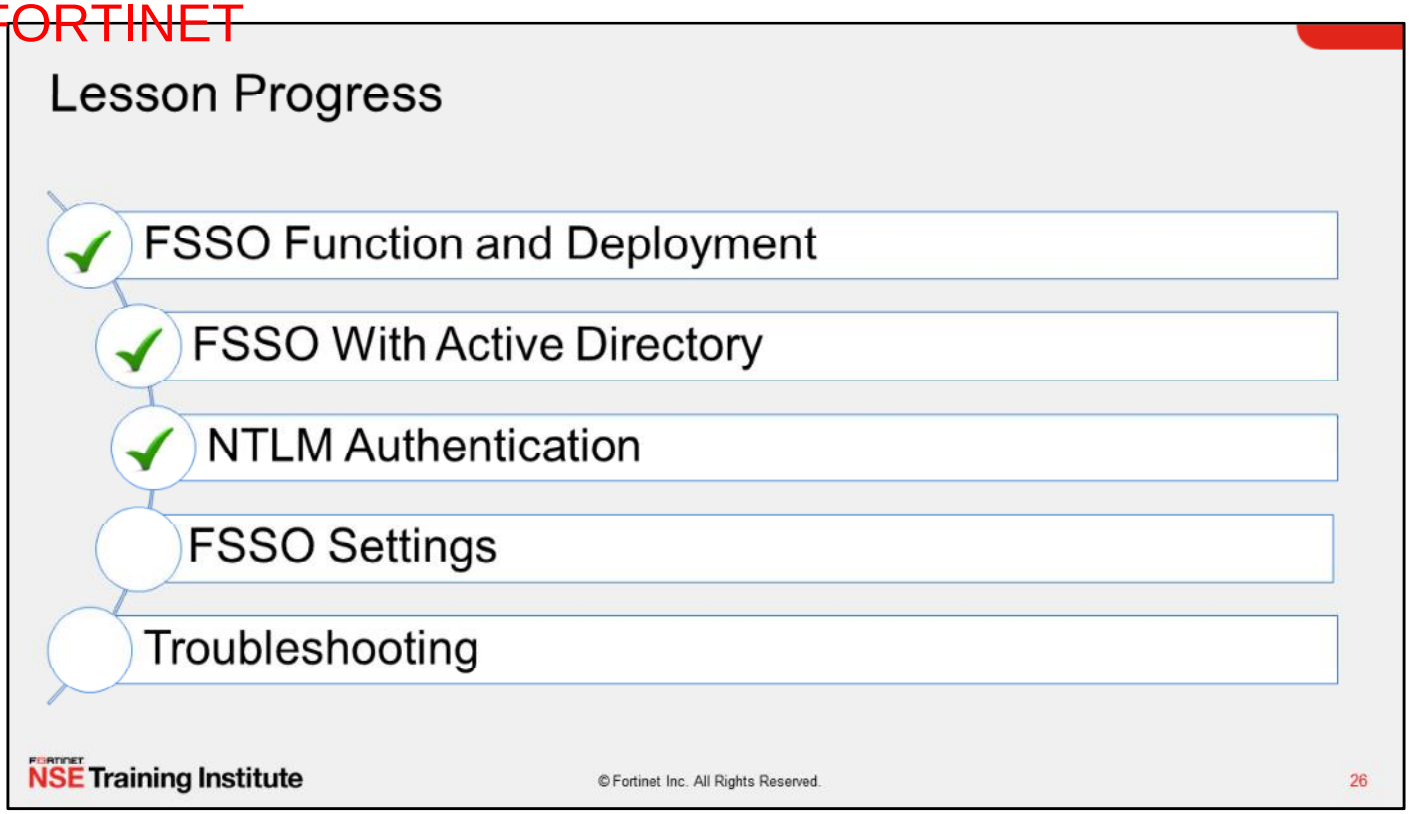

Good job! You now understand how to use FSSO with NTLM authentication.

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

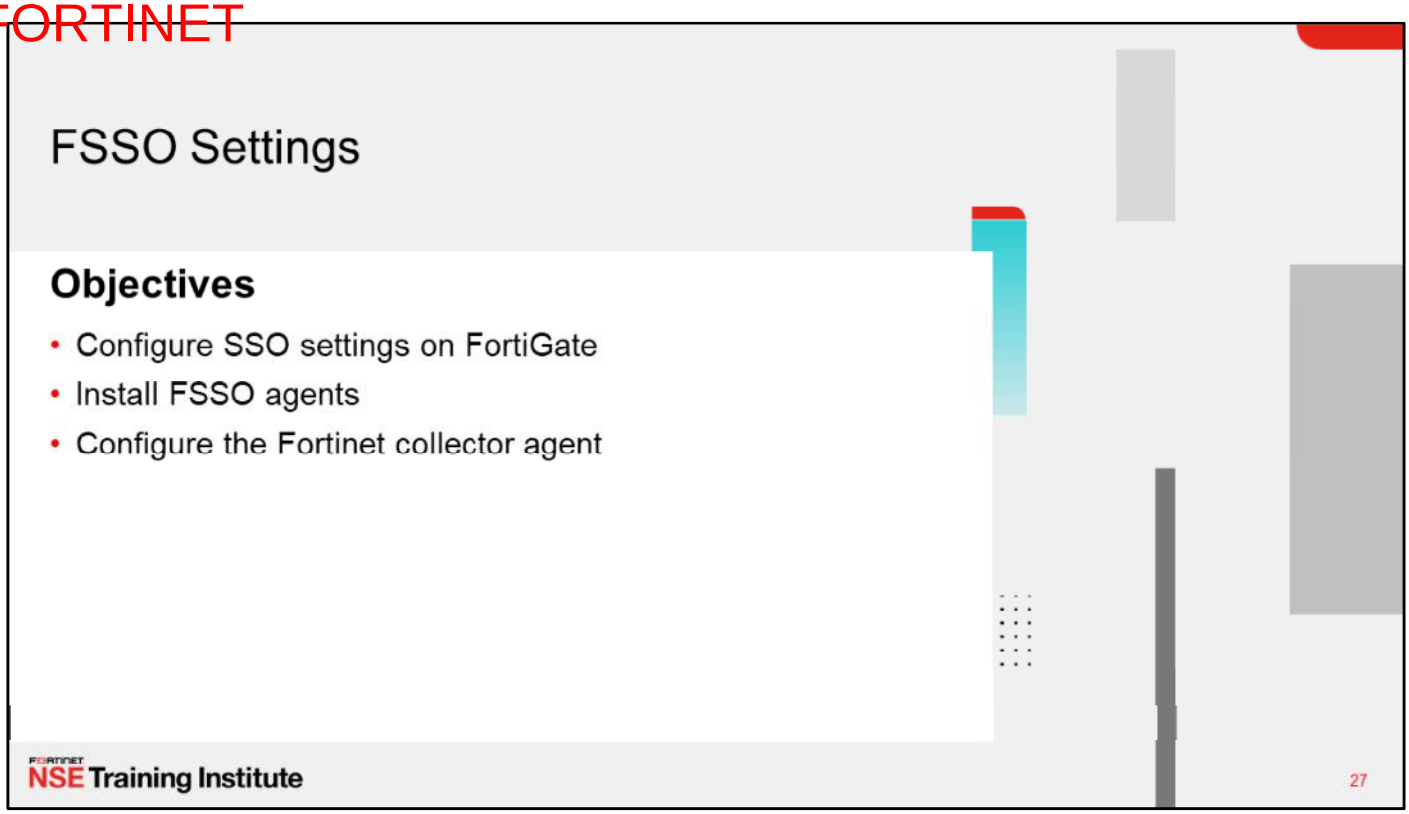

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

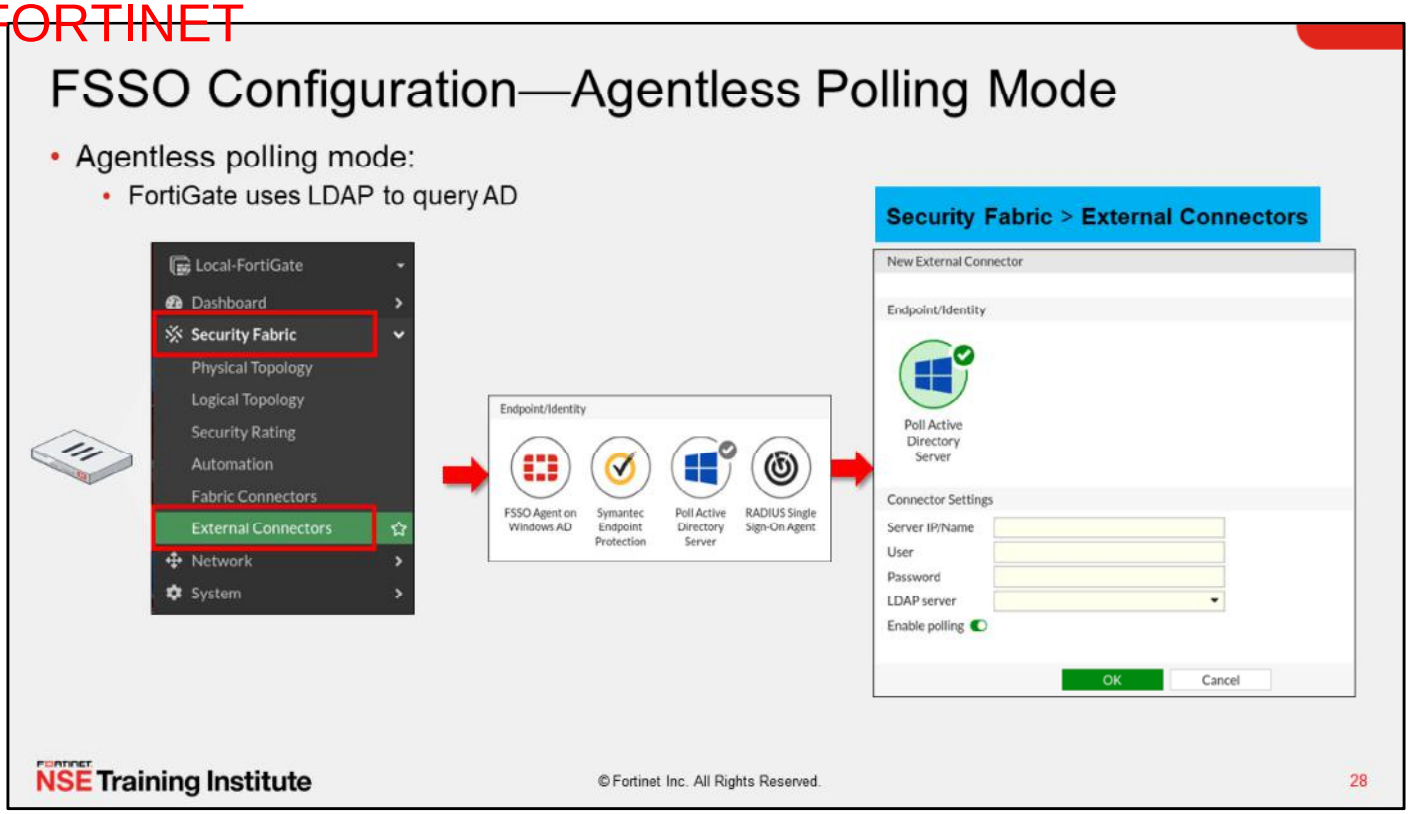

FortiGate FSSO configuration is straightforward.

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.

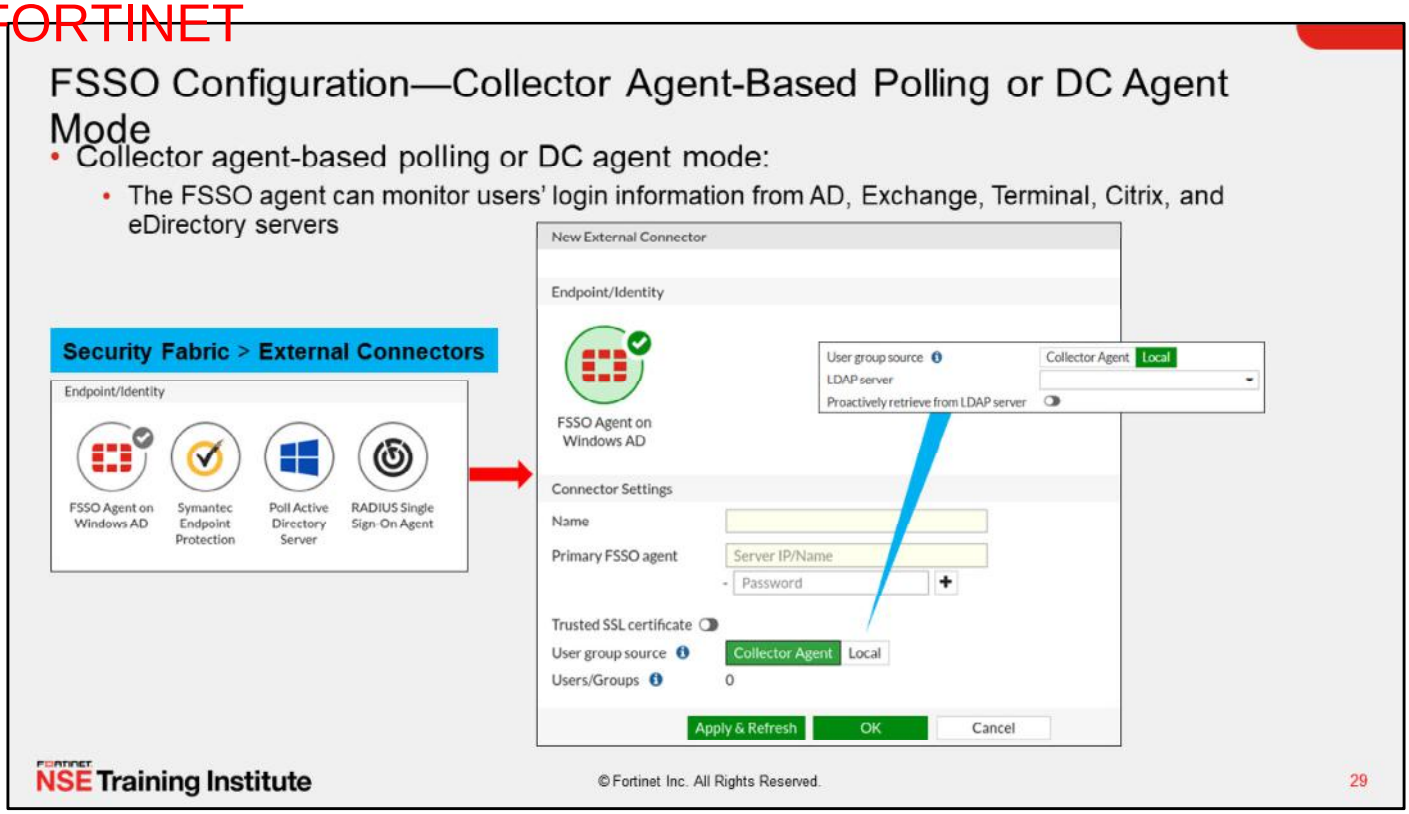

If you have collector agents, using either the DC agent mode or the collector agent-based polling mode, you must select **Fortinet Single-Sign-On Agent** and configure the IP address and password for each collector agent.

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** > **7.0** > **7.0.0** > **FSSO**

Example image below:

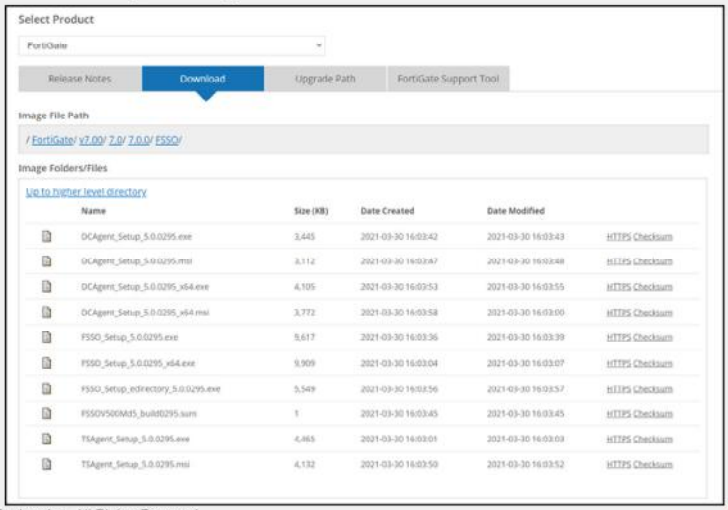

NSE Training Institute

30

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.

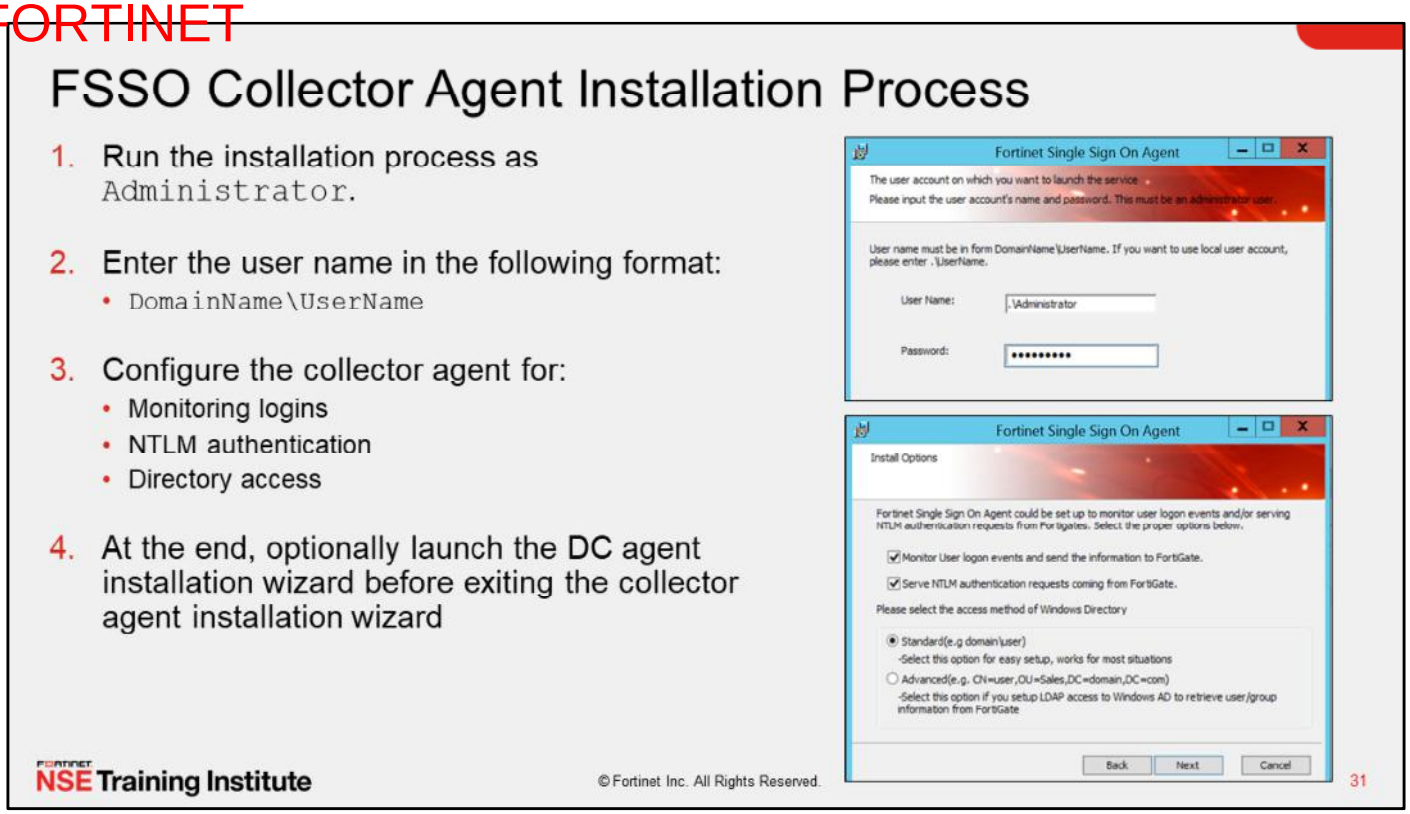

After you've downloaded the collector agent, run the installation process as Administrator and follow these steps in the installation wizard:

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 later 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.

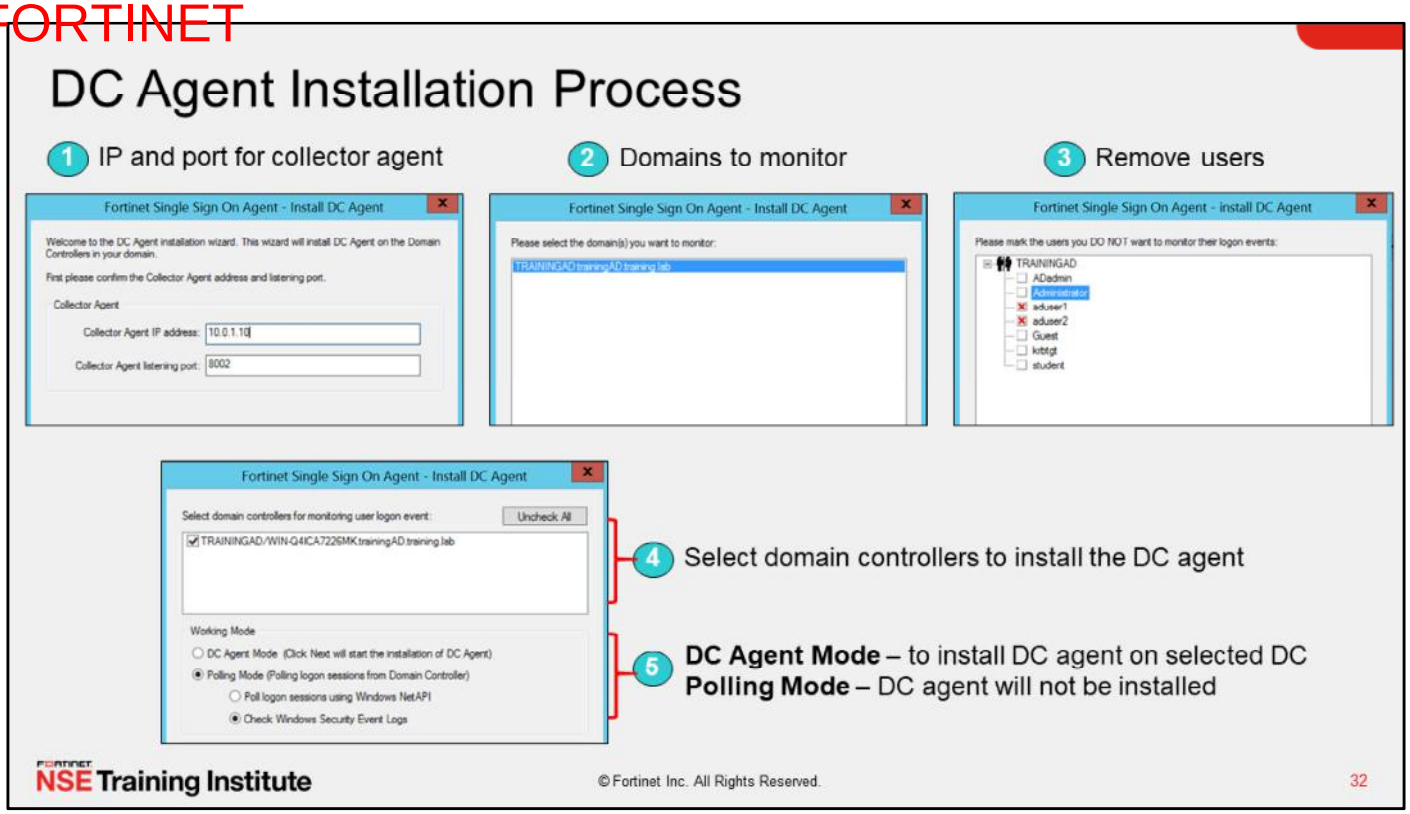

If you have just installed the collector agent and you selected **Launch DC Agent Install Wizard,** the installation process for domain controller agent automatically starts.

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.

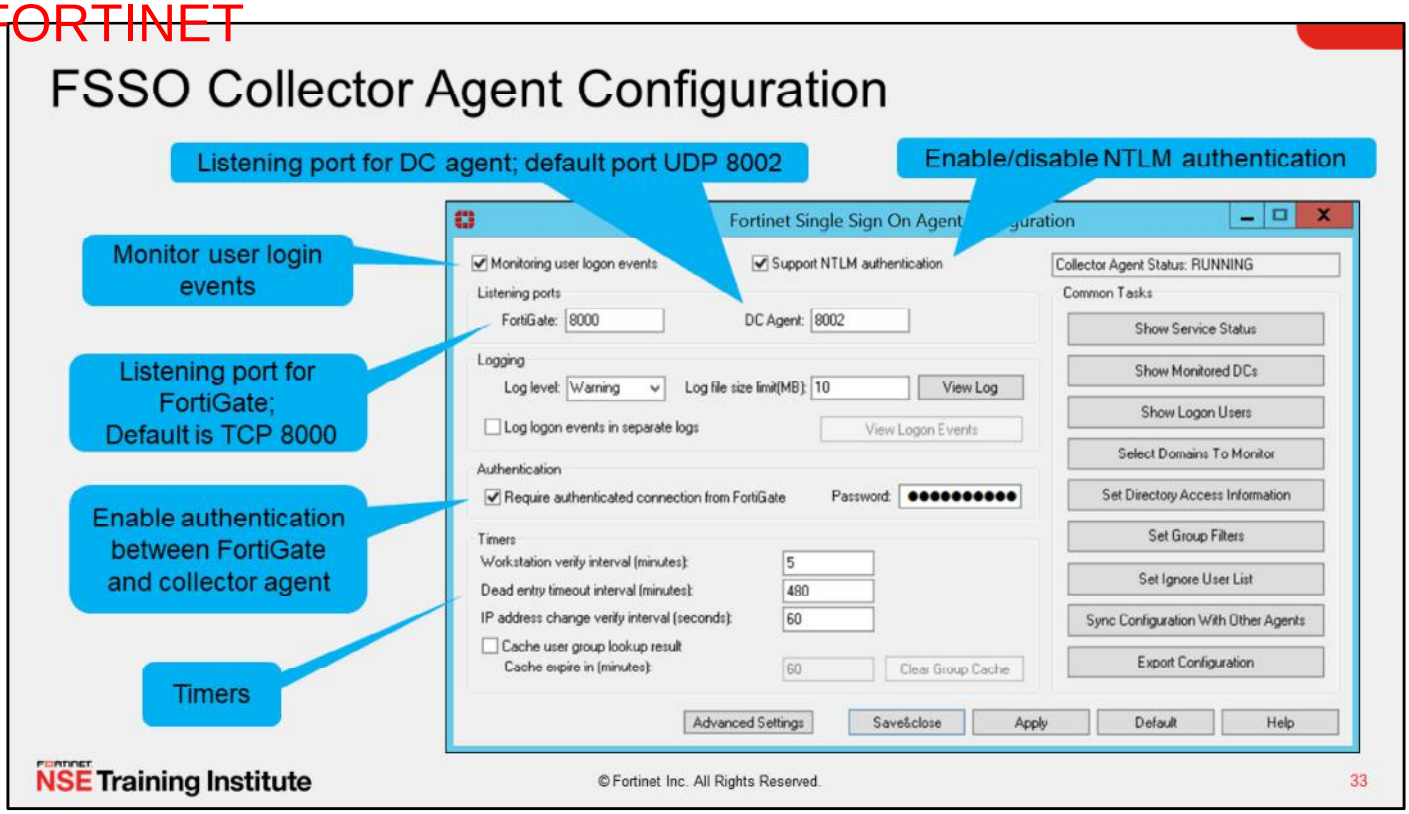

On the FSSO agent configuration GUI, you can configure settings such as:

- 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

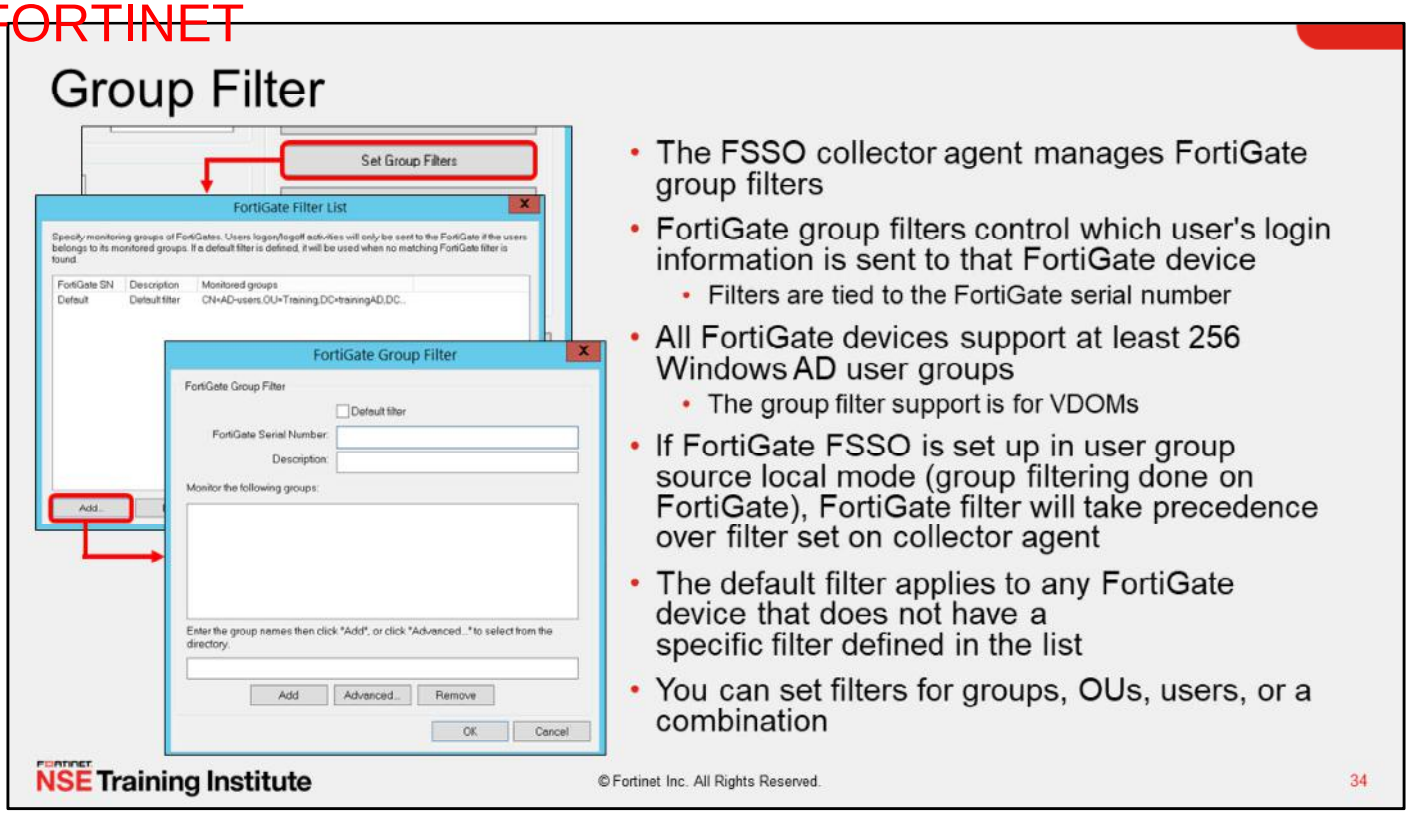

The FSSO collector agent allows you to configure a FortiGate group filter, which actively controls what user login information is sent to each FortiGate device. So, you can define which groups the collector agent passes to individual FortiGate devices.

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.

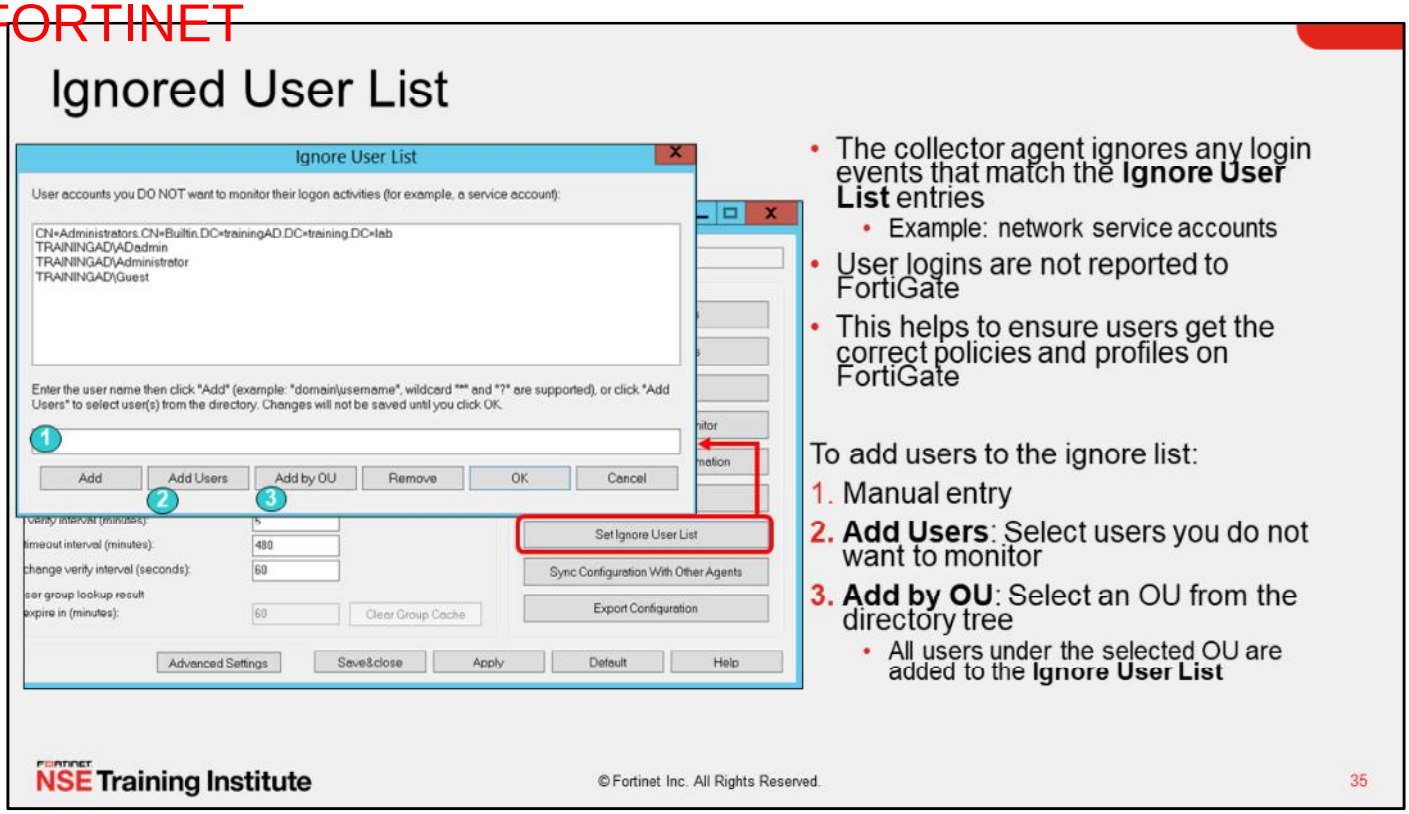

The FSSO collector agent ignores any login events that match the **Ignore User List** entries. Therefore, these login events are not recorded by the collector agent, nor are they reported to FortiGate.

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.

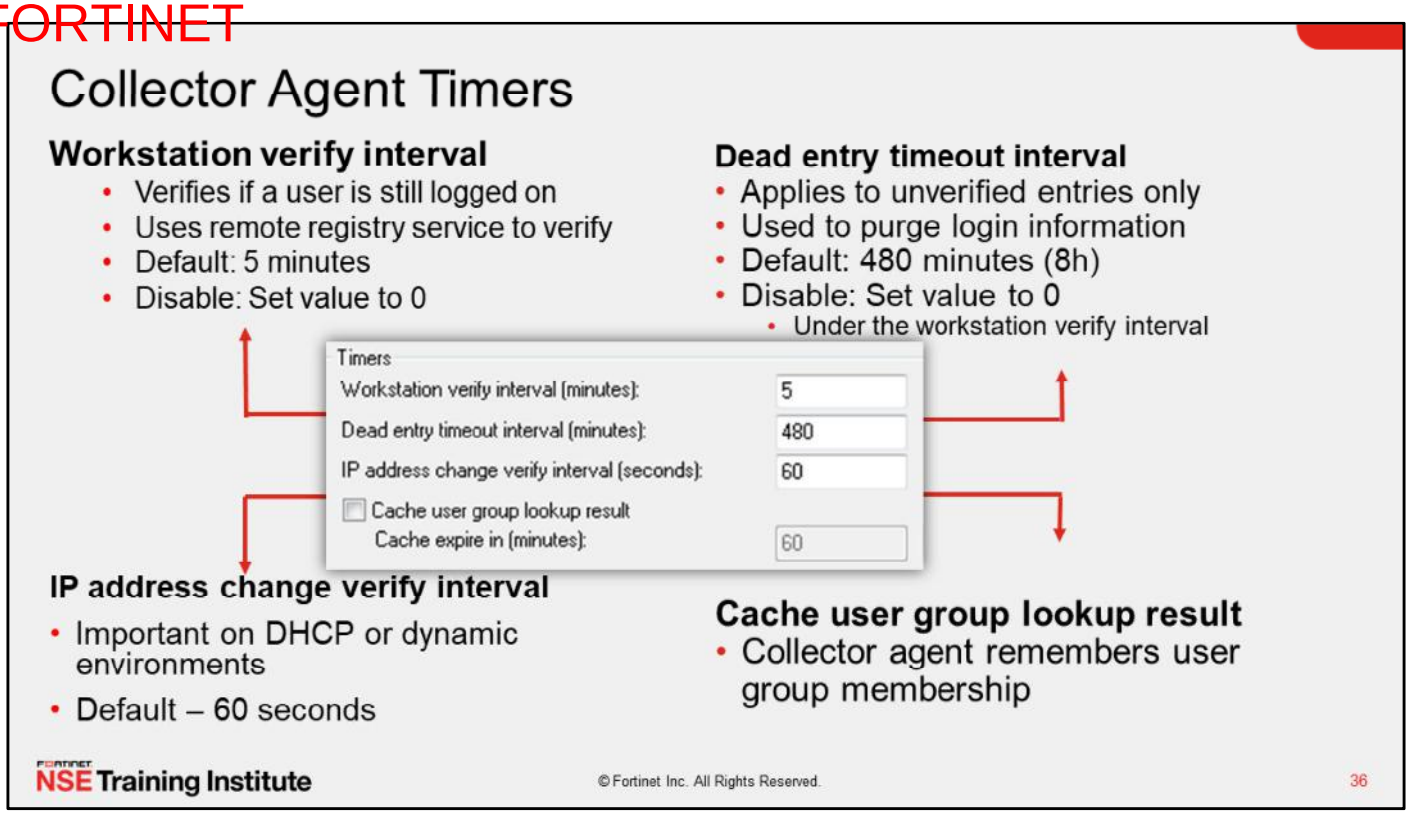

The FSSO collector agent timers play an important role in ensuring the correct operation of FSSO.

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.

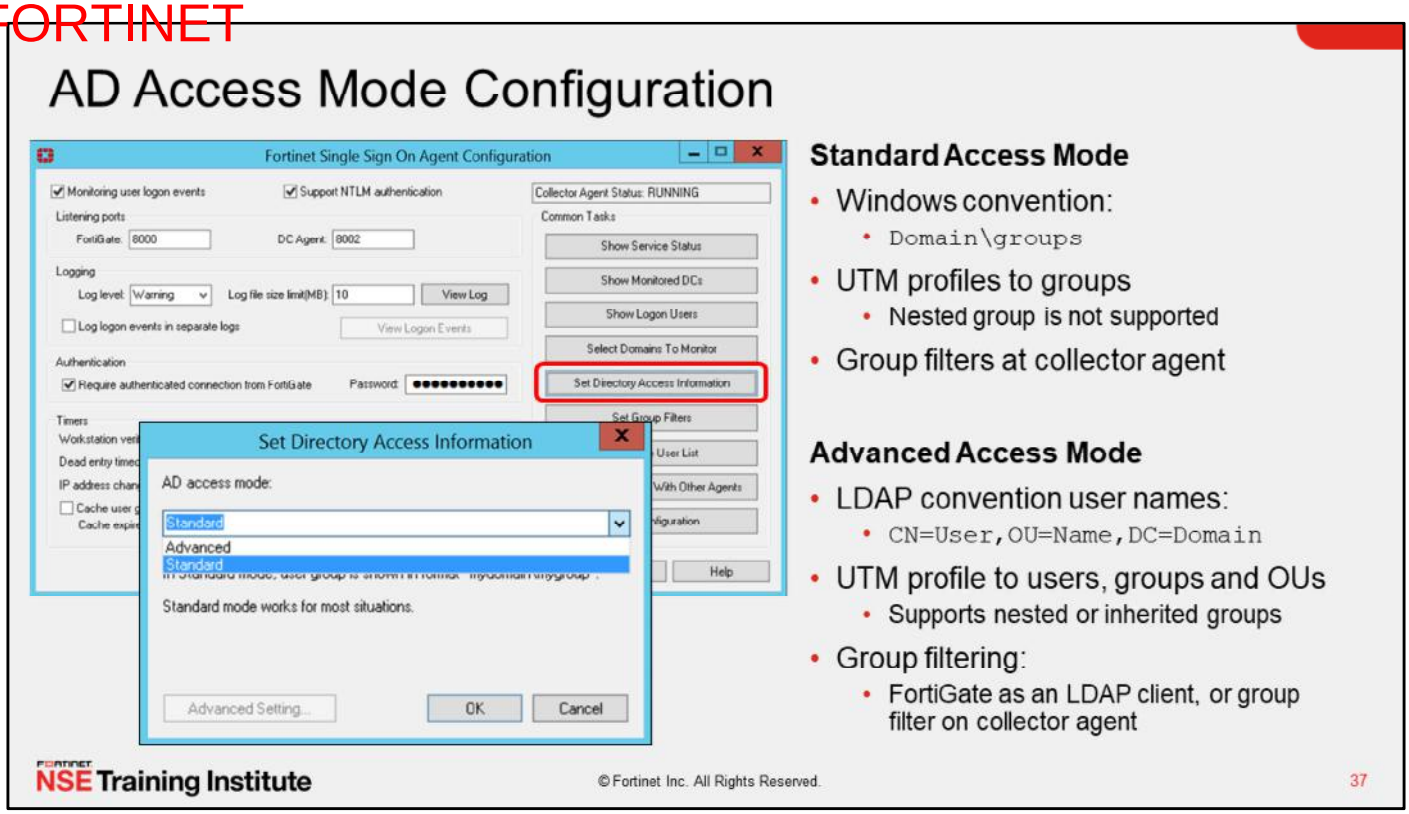

Another important FSSO setting is the AD access mode. You can set the AD access mode by clicking **Set Directory Access Information**. The AD access mode specifies how the collector agent accesses and collects the user and user group information. There are two modes that you can use to access AD user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

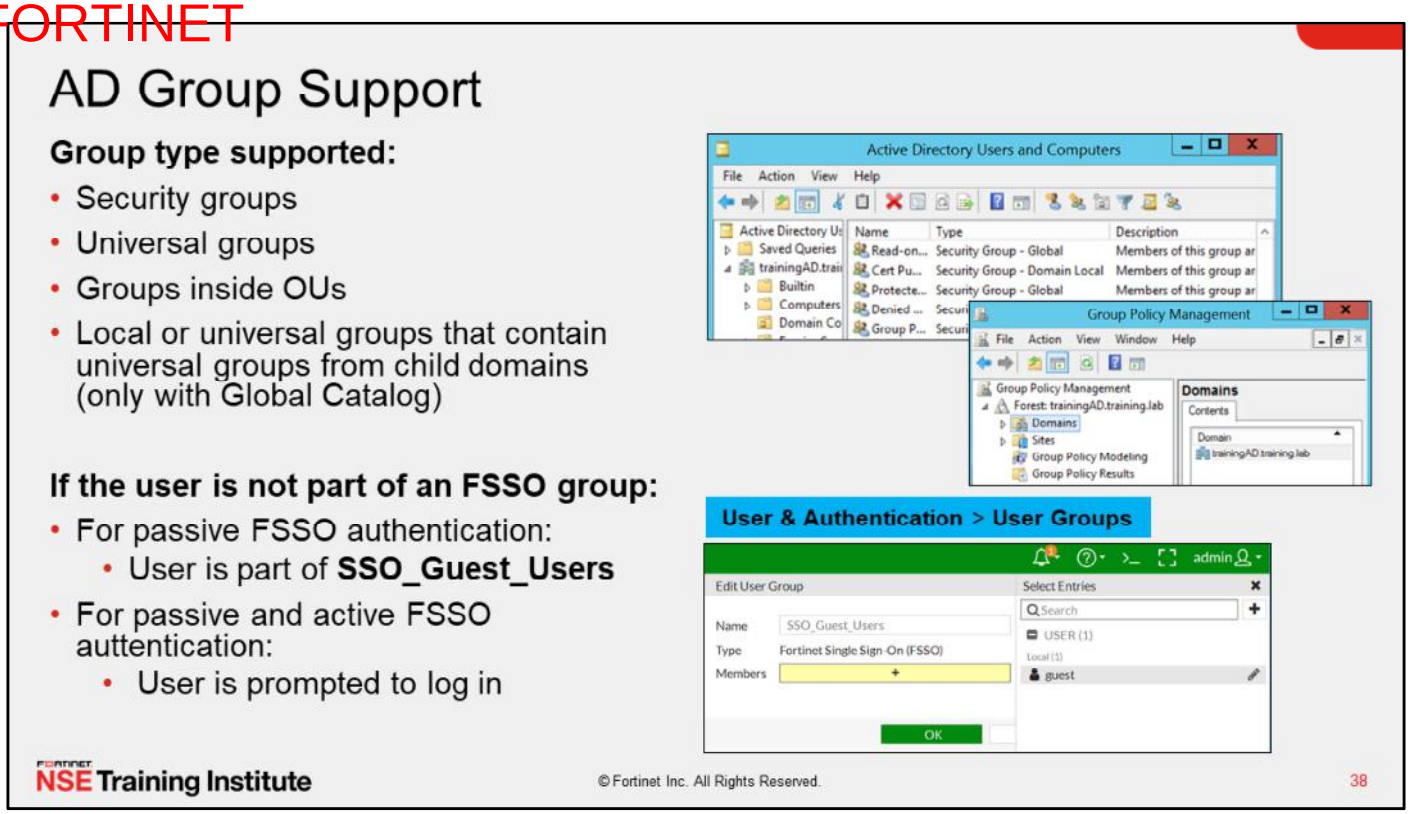

In AD settings, not all group types are supported. AD settings supports filtering groups only from:

- 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.

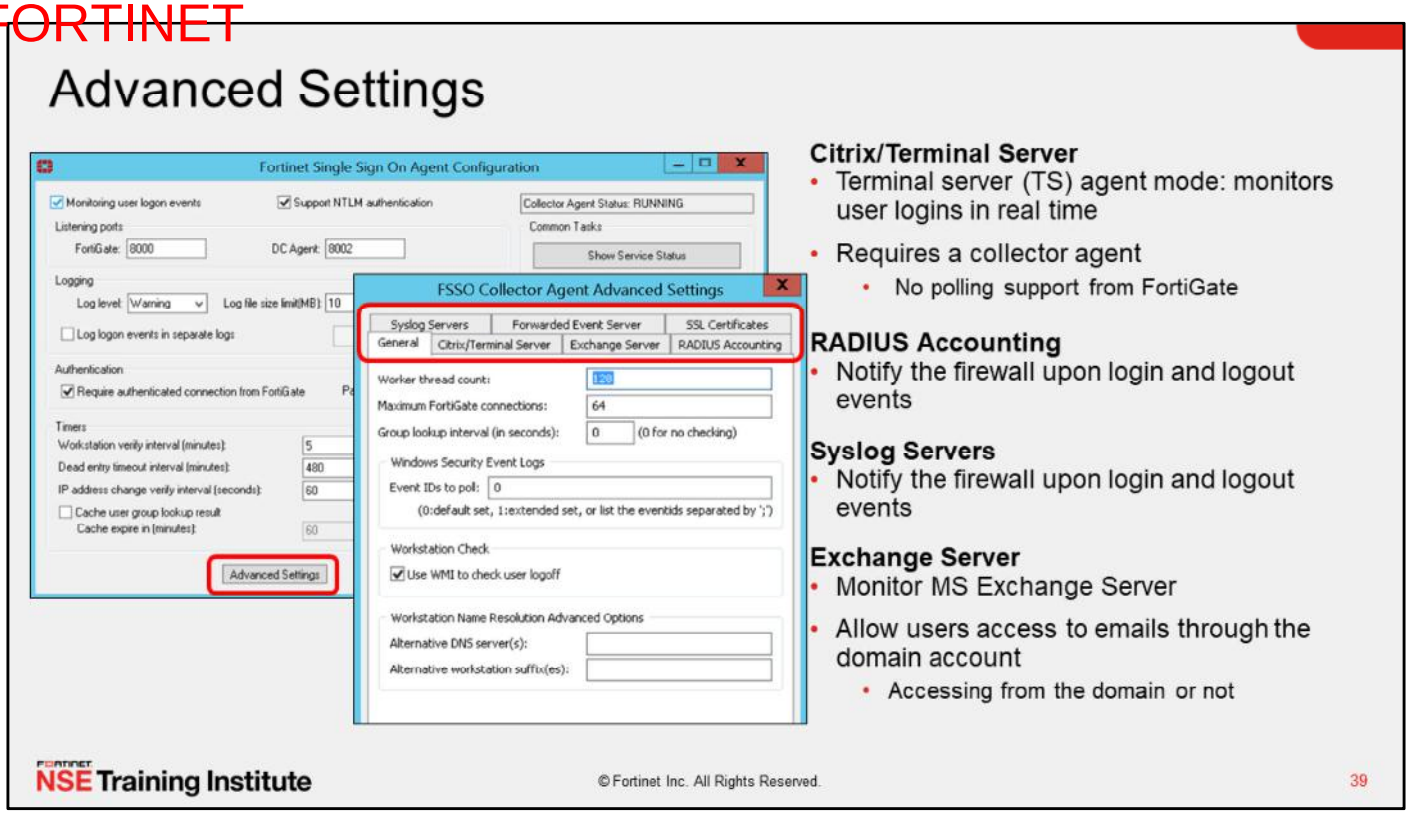

Depending on your network, you might need to configure advanced settings in your FSSO collector agent.

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`

NSE Training Institute

40

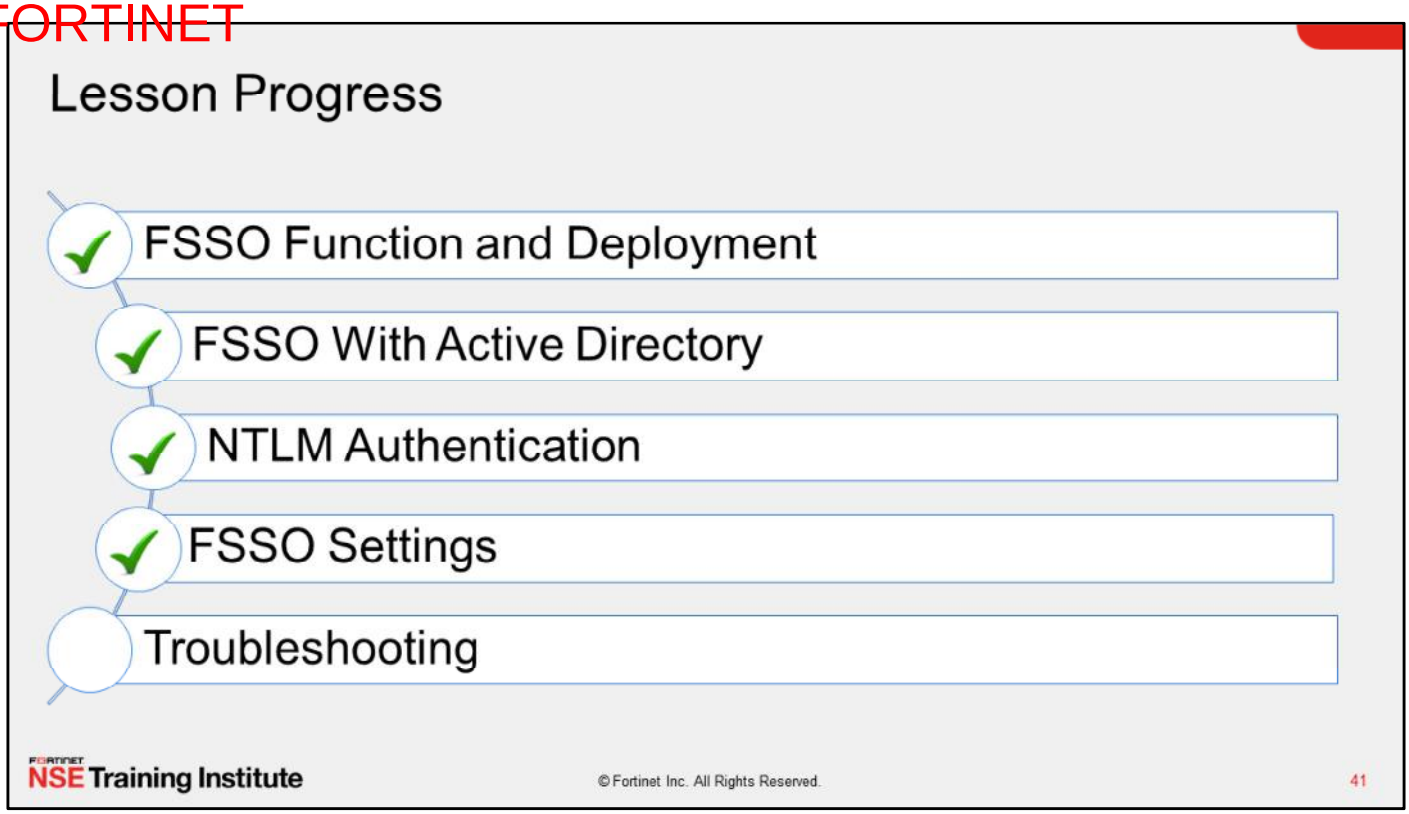

Good job! You now understand how to configure the SSO settings on FortiGate and the FSSO collector agent.

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.

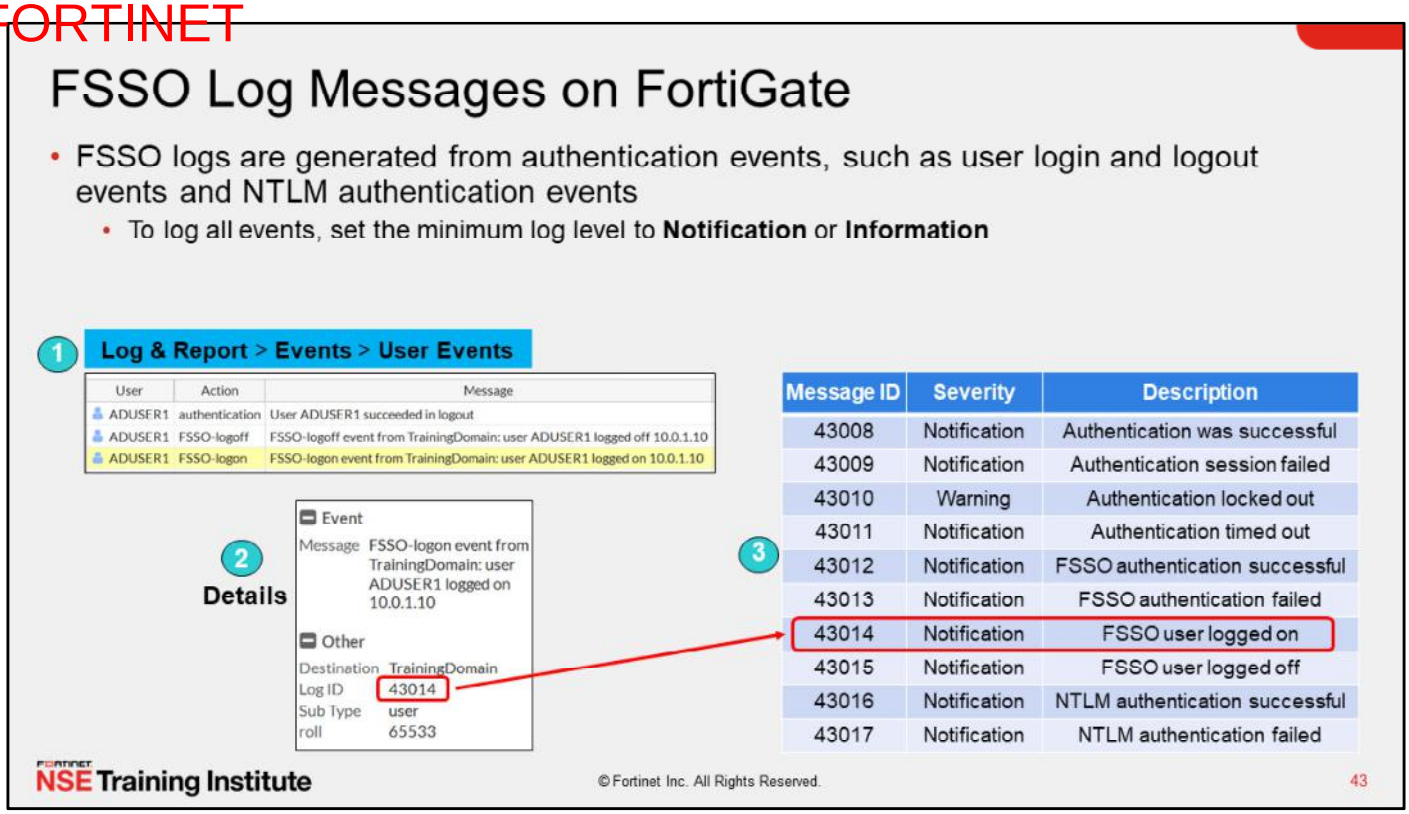

FSSO-related log messages are generated from authentication events. These include user login and logout events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues.

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.

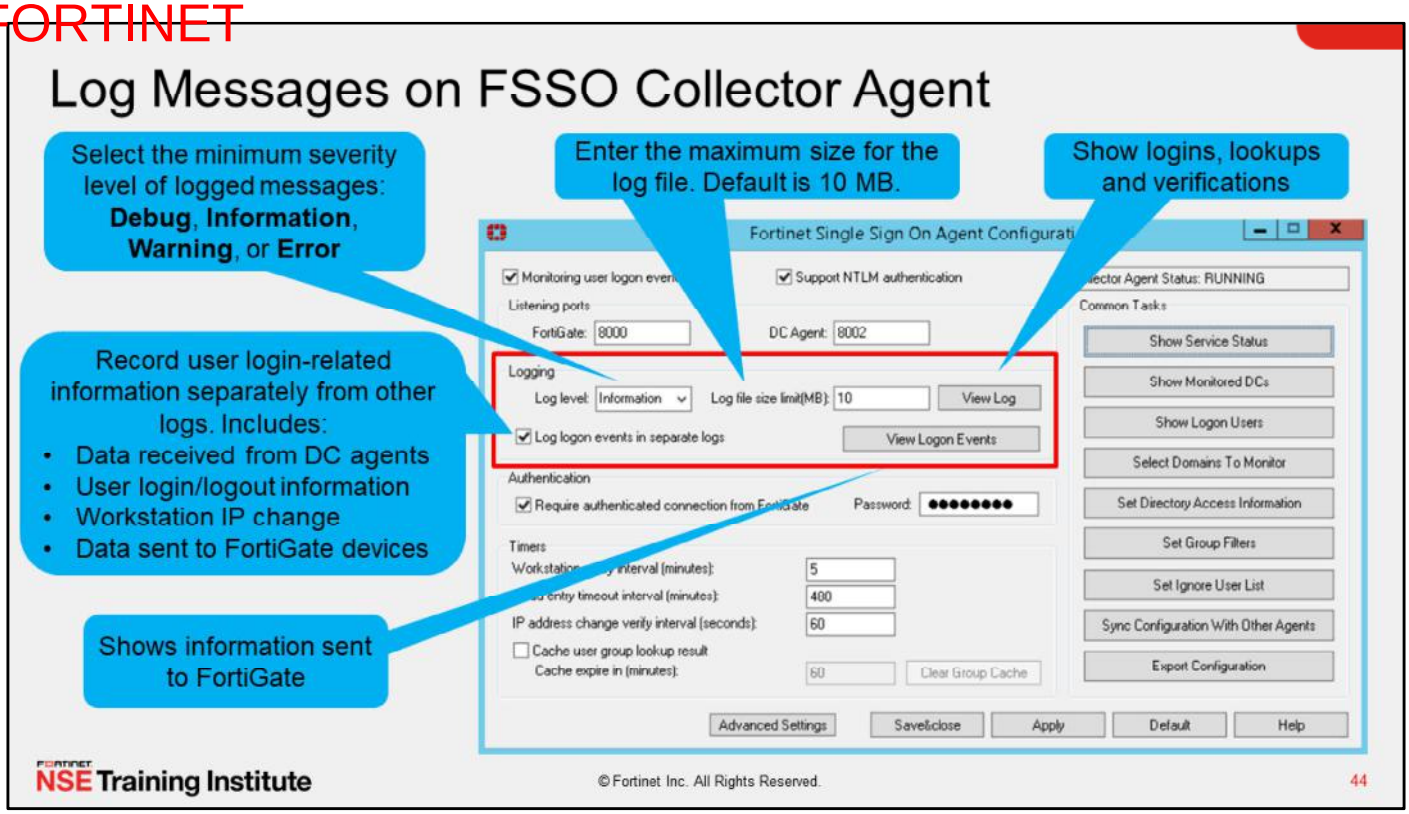

When troubleshooting FSSO agent-based deployments, you might want to look at the log messages generated directly on the FSSO collector agent.

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.

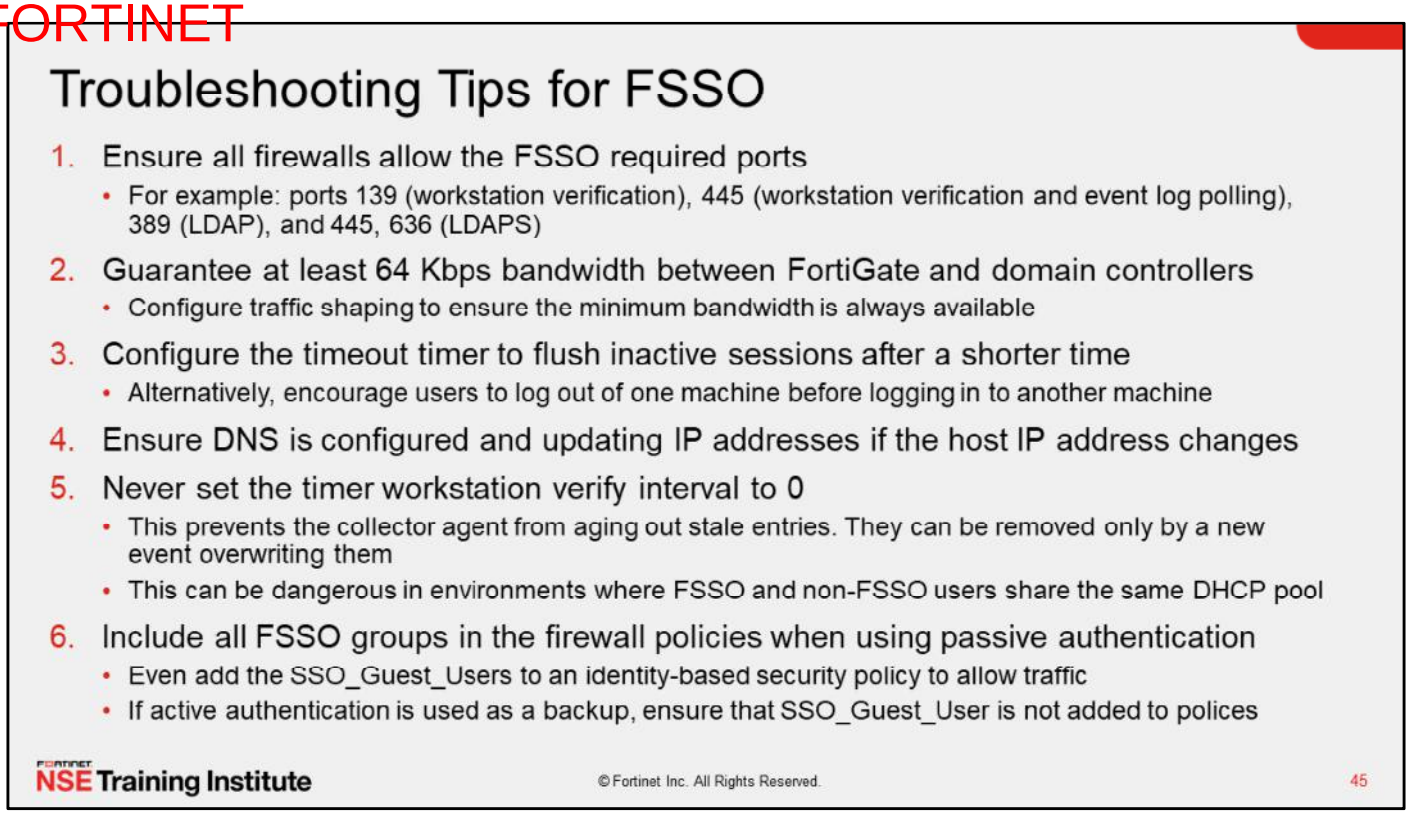

Begin with the following tips, which are useful in many FSSO troubleshooting situations:

- 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.

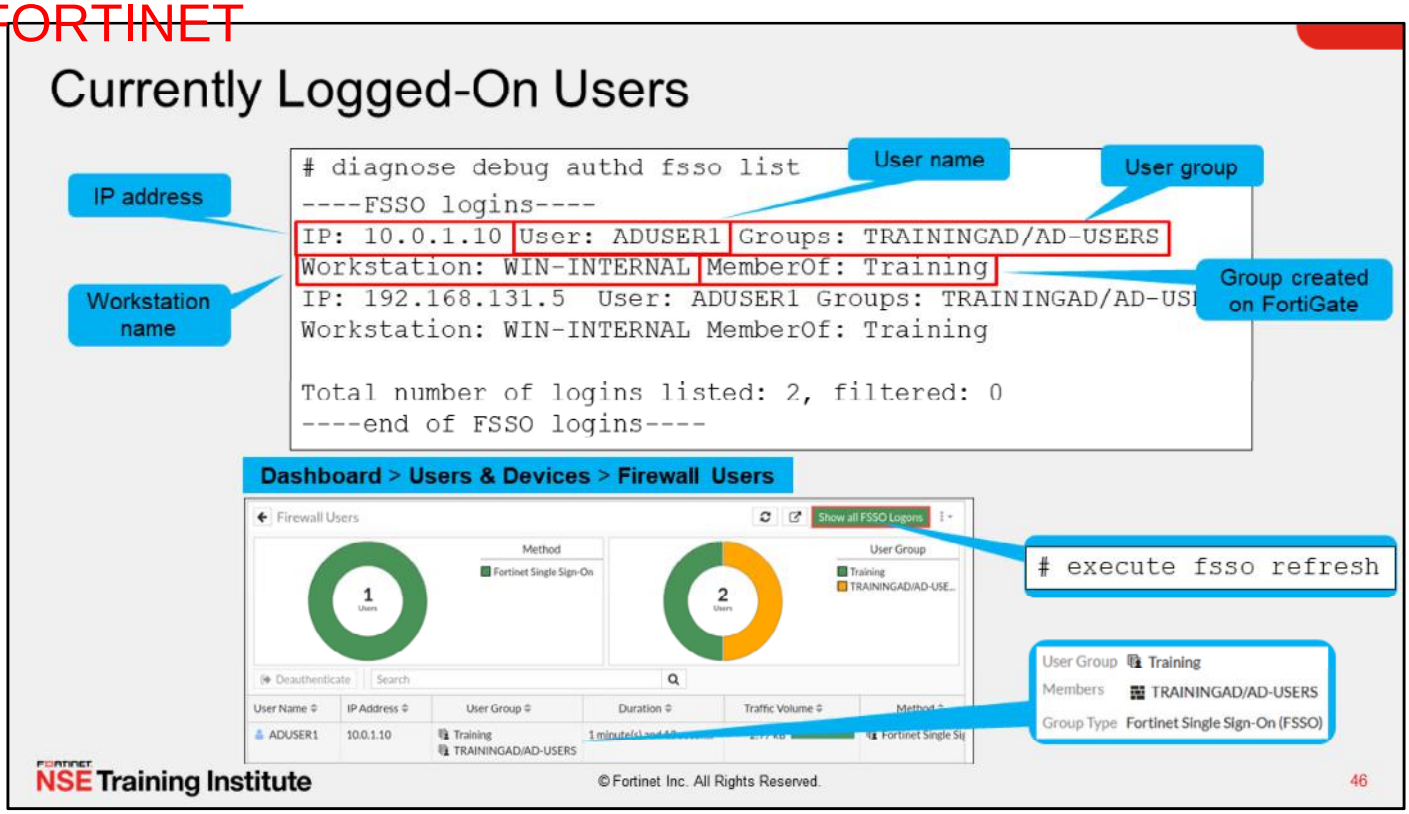

If applying the tips from the previous slide didn't solve your FSSO issues, you may need to apply some `debug` commands.

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.

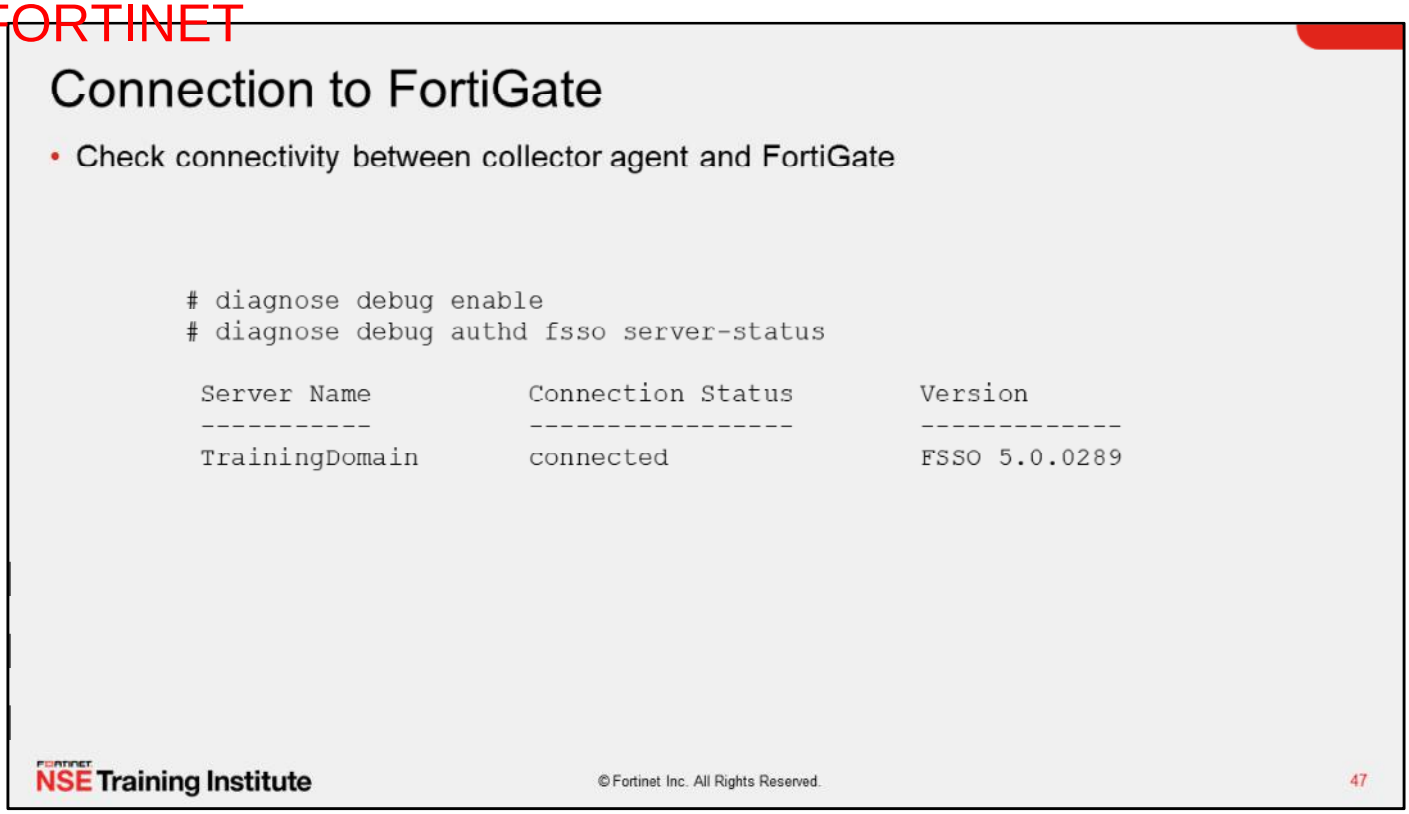

To show the status of communication between FortiGate and each collector agent, you can use the CLI command `diagnose debug authd fsso server-status`.

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

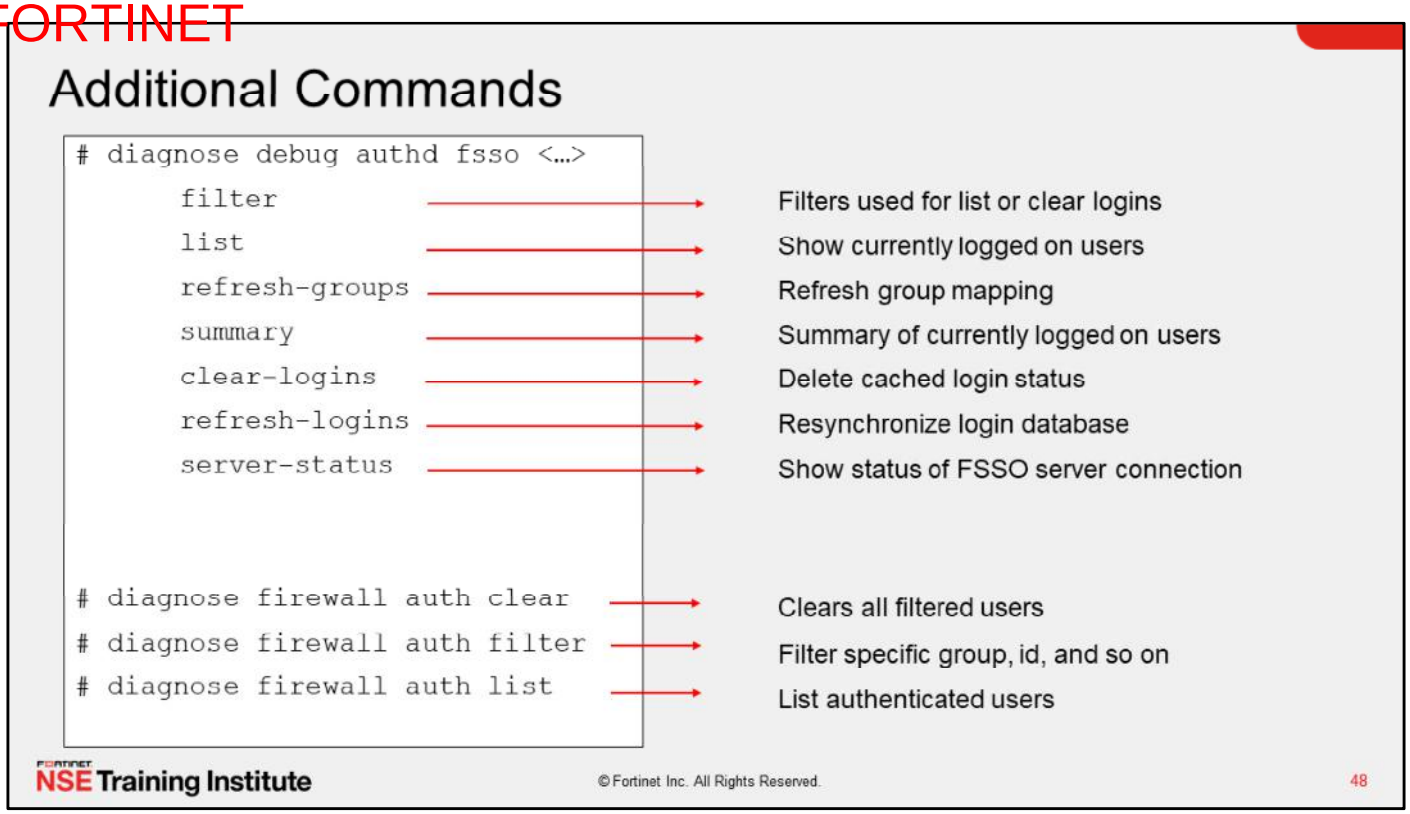

Also, available under `diagnose debug authd fsso` are commands for clearing the FortiGate cache of all currently logged in users, filtering the display of the list of logged in users, and refreshing the login and user group information.

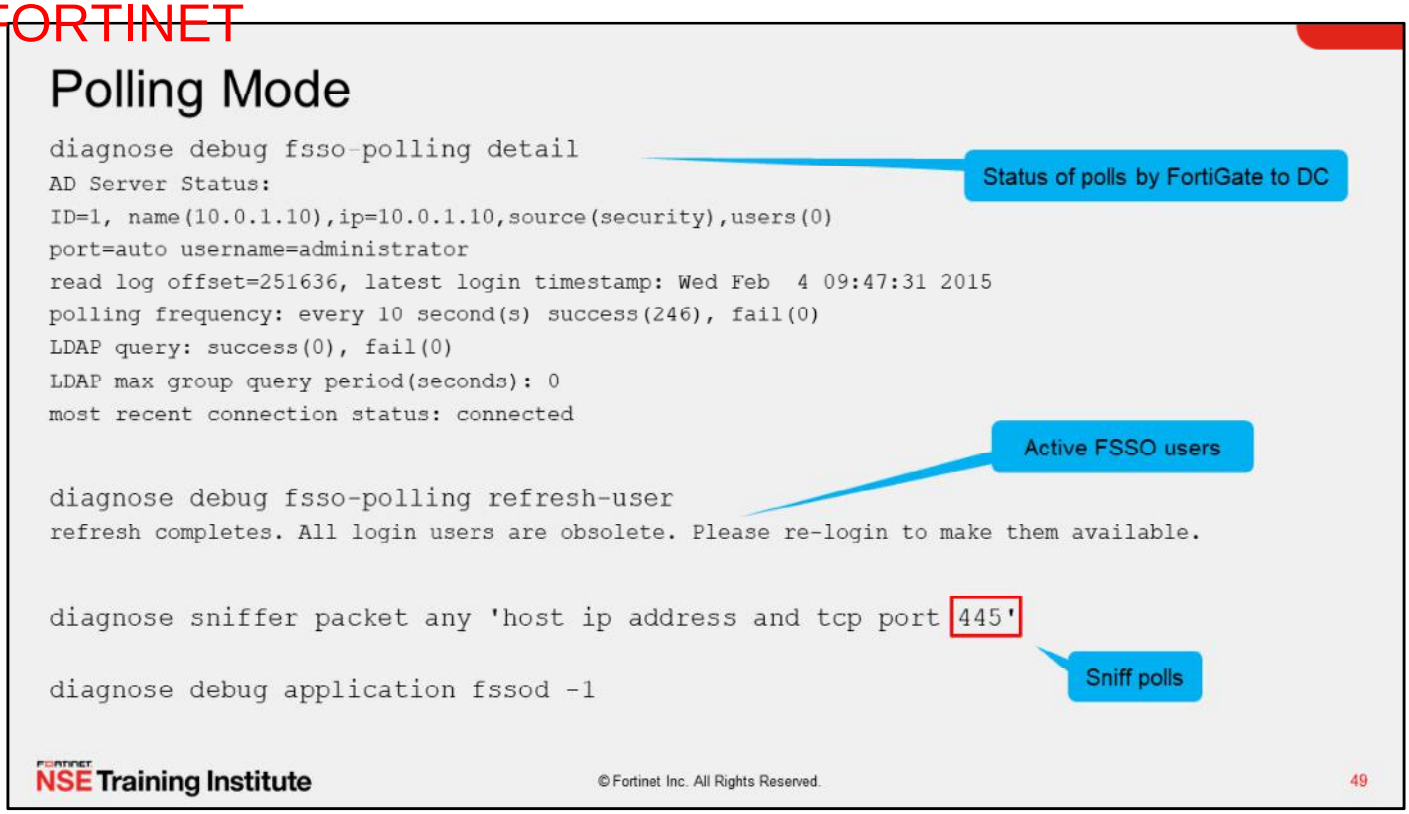

The command `diagnose debug fsso-polling detail` displays status information and some statistics related to the polls done by FortiGate on each DC in agentless polling. If the `read log offset` is incrementing, FortiGate is connecting to and reading the logs on the domain controller. If the `read log 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

NSE Training Institute

50

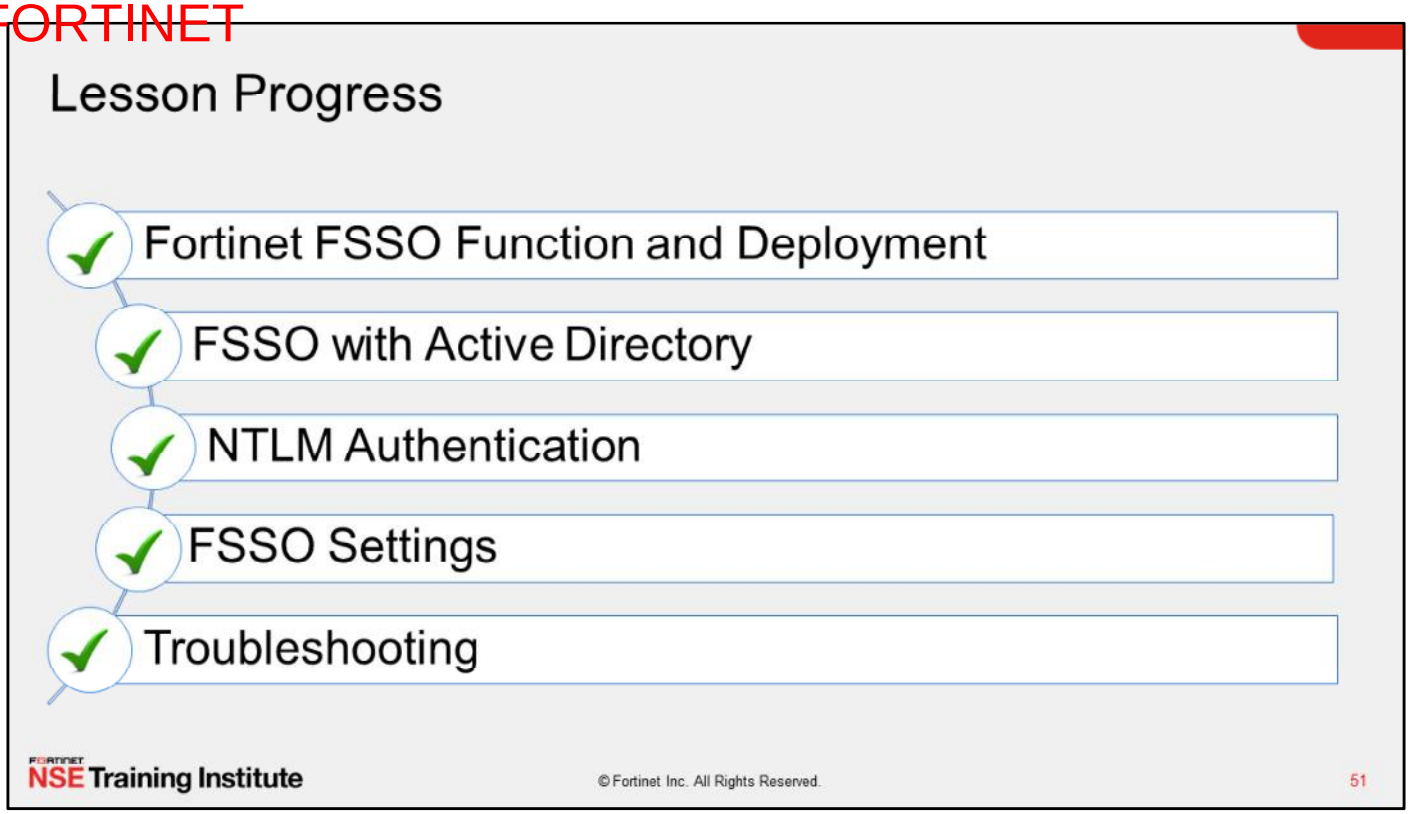

Congratulations! You have completed this lesson.

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

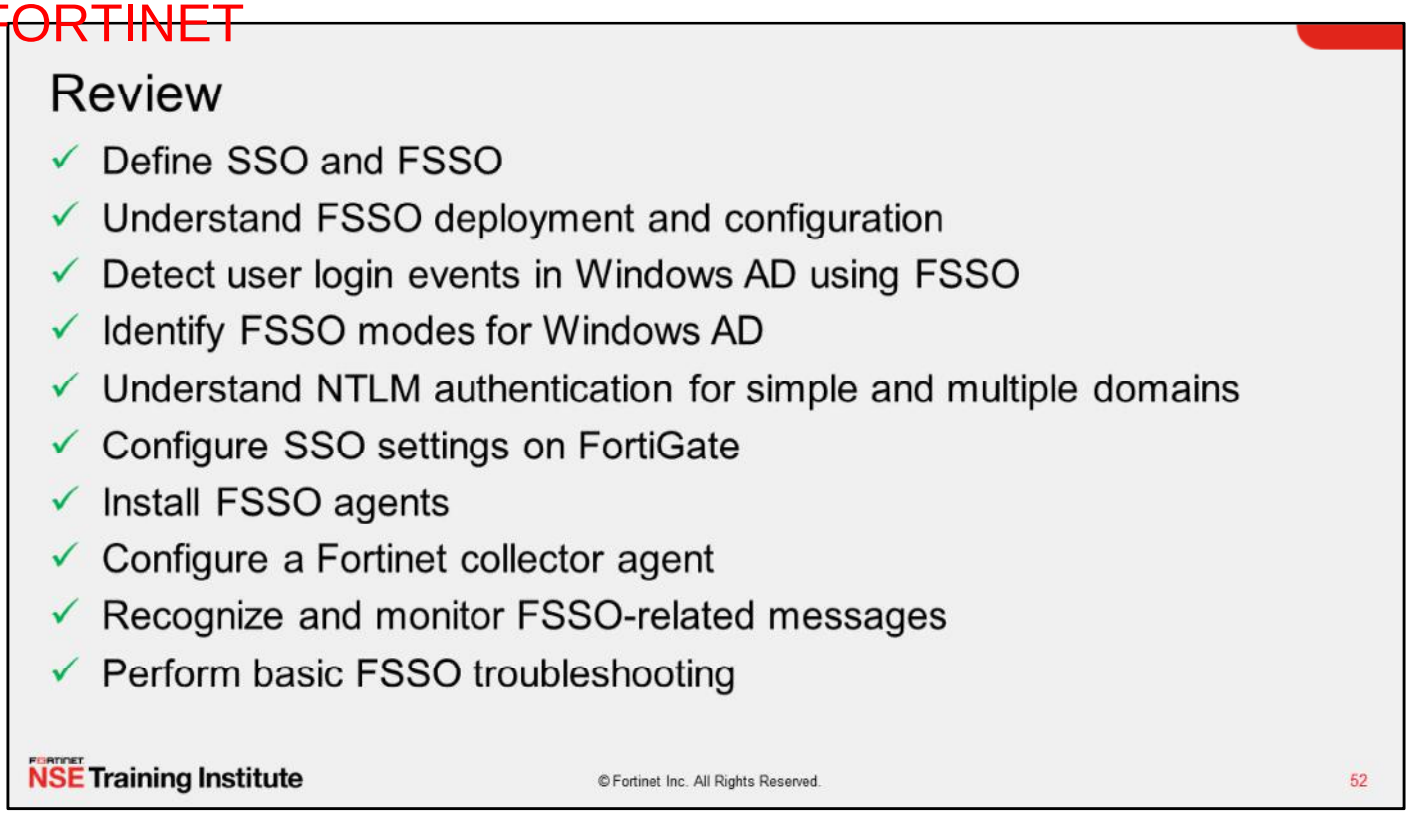

This slide shows the objectives that you covered in this lesson.

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.

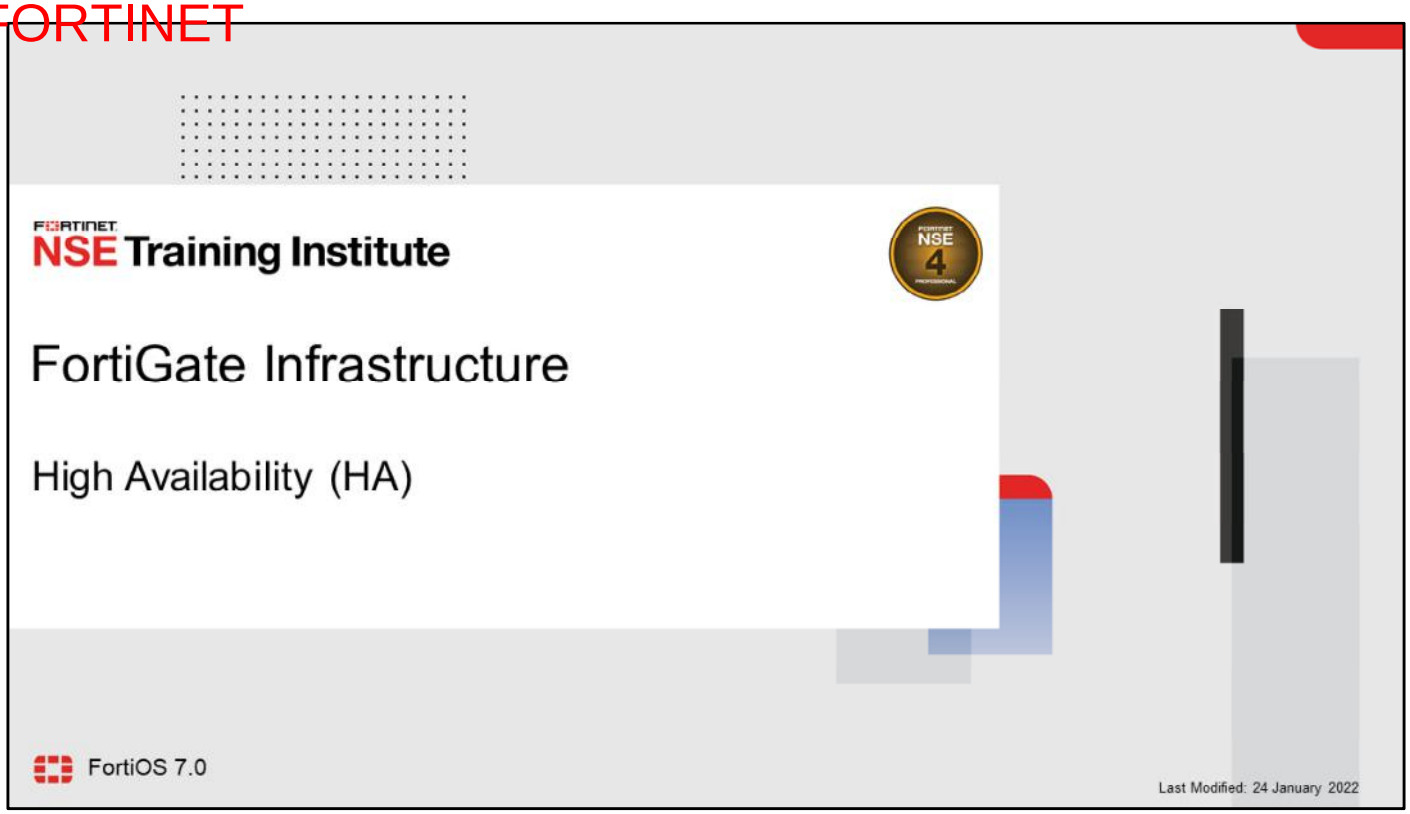

In this lesson, you will learn about the fundamentals of FortiGate high availability (HA) and how to configure it. FortiGate HA provides a solution for enhanced reliability and increased performance.

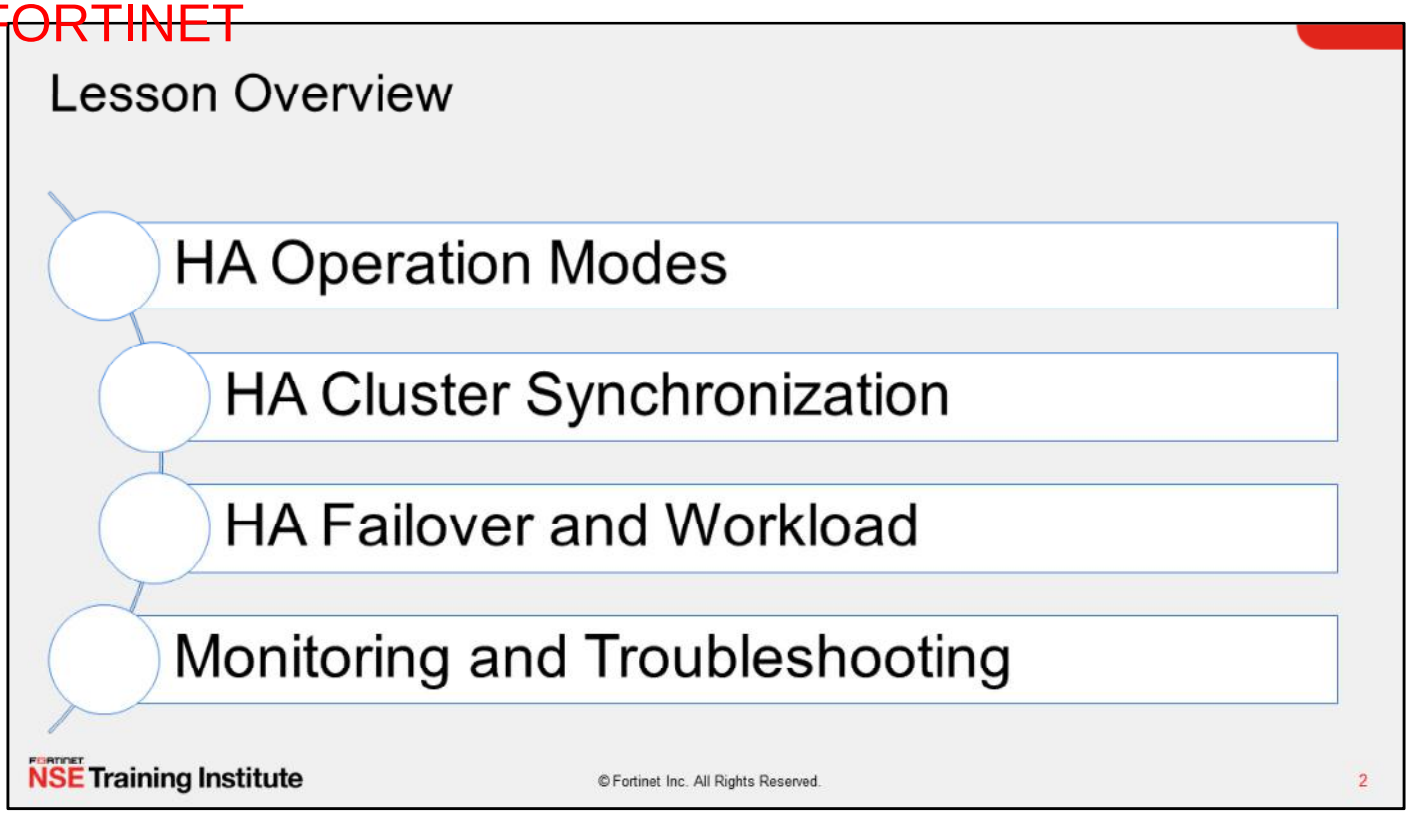

In this lesson, you will learn about the topics shown on this slide.

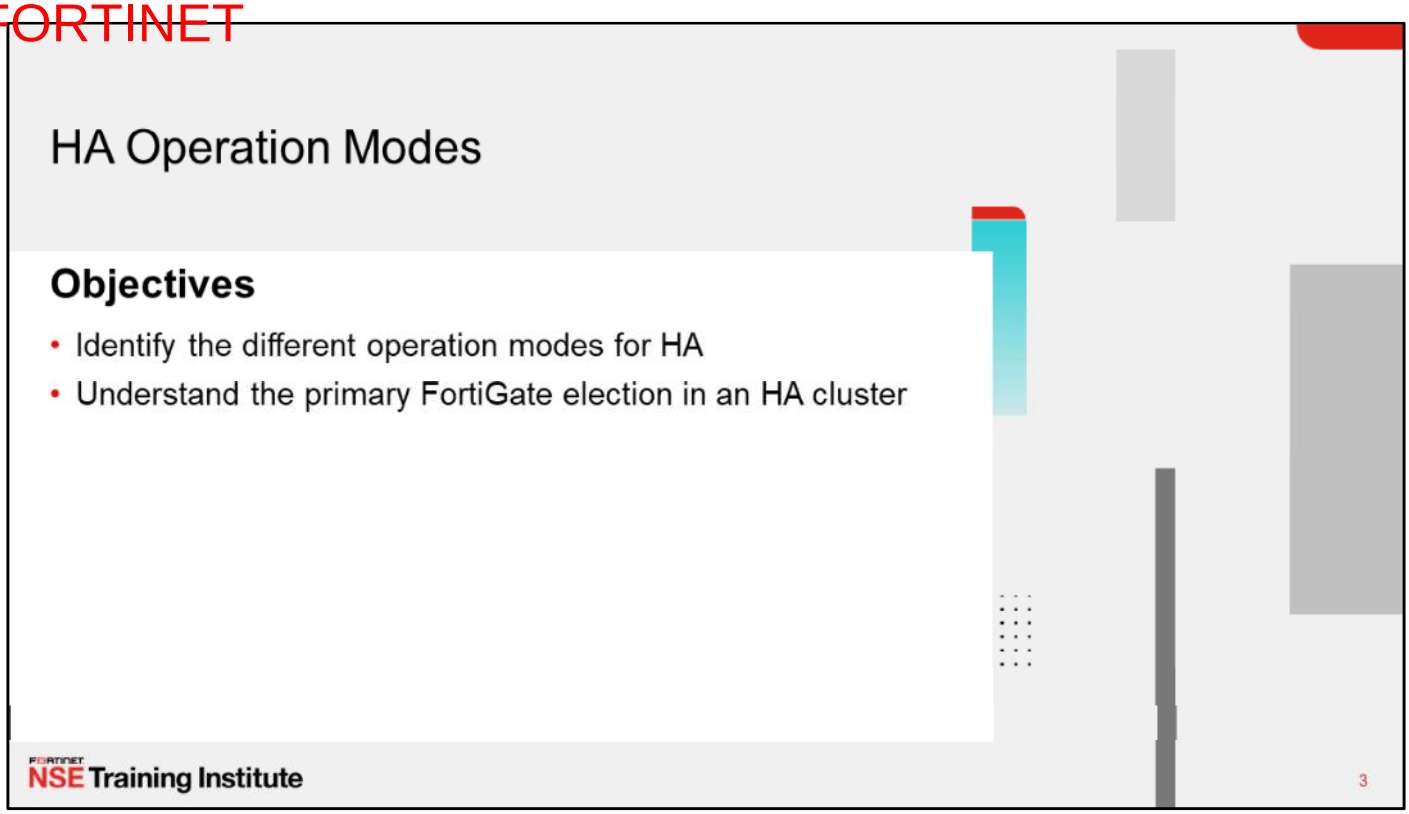

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

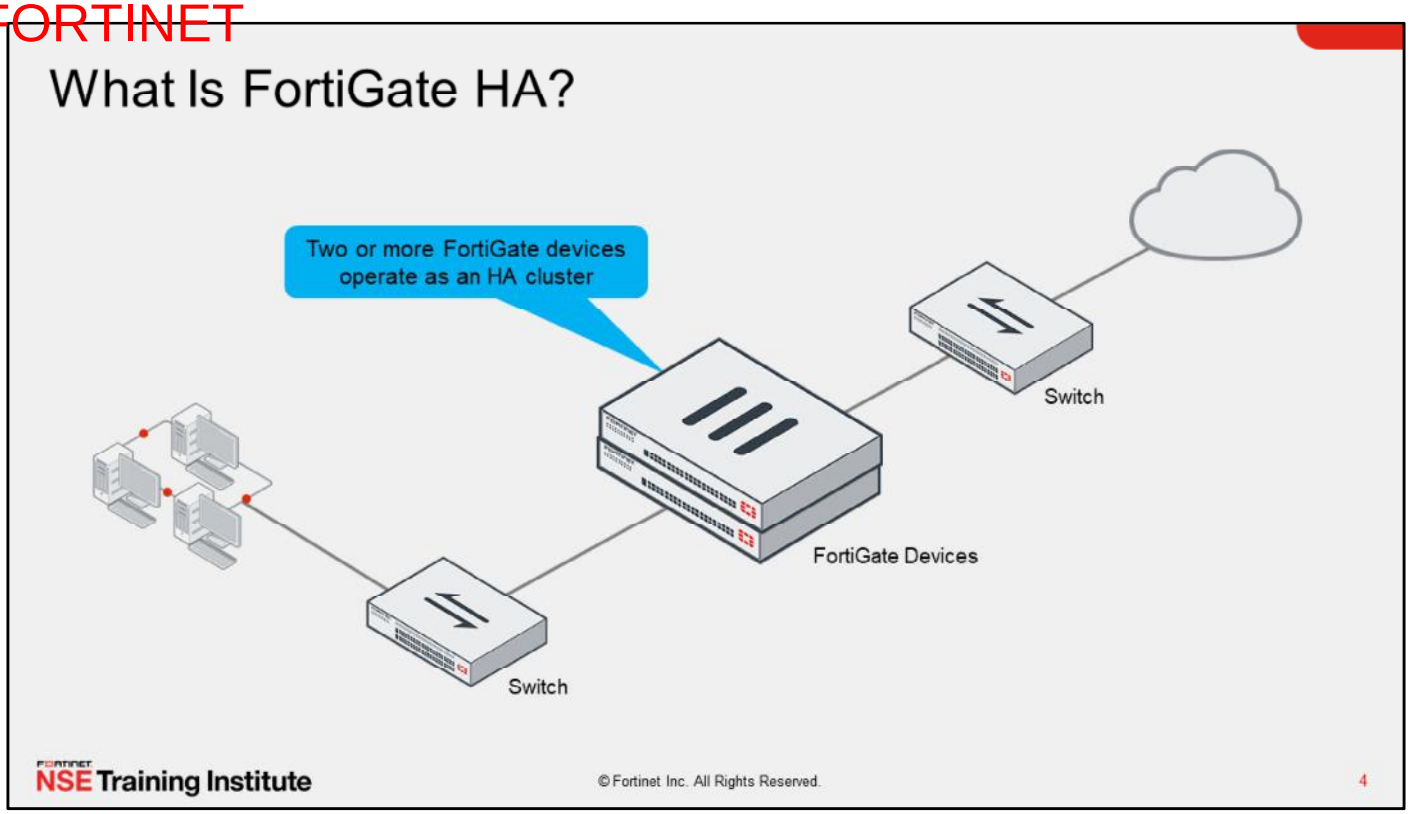

The idea of HA is simple. HA links and synchronizes two or more devices.

In FortiGate HA, one FortiGate device acts as the *primary* device (also called the *active* FortiGate). It synchronizes its configuration to the other devices. The other FortiGate devices are called *secondary* or *standby* devices.

A heartbeat link between all the appliances is used to detect unresponsive devices.

What is synchronized between the devices? Are all FortiGate devices processing traffic? Does HA improve availability, or does it improve throughput?

The answers vary, depending on the HA mode. There are currently two HA modes available: active-active and active-passive. Now, you will examine the differences.

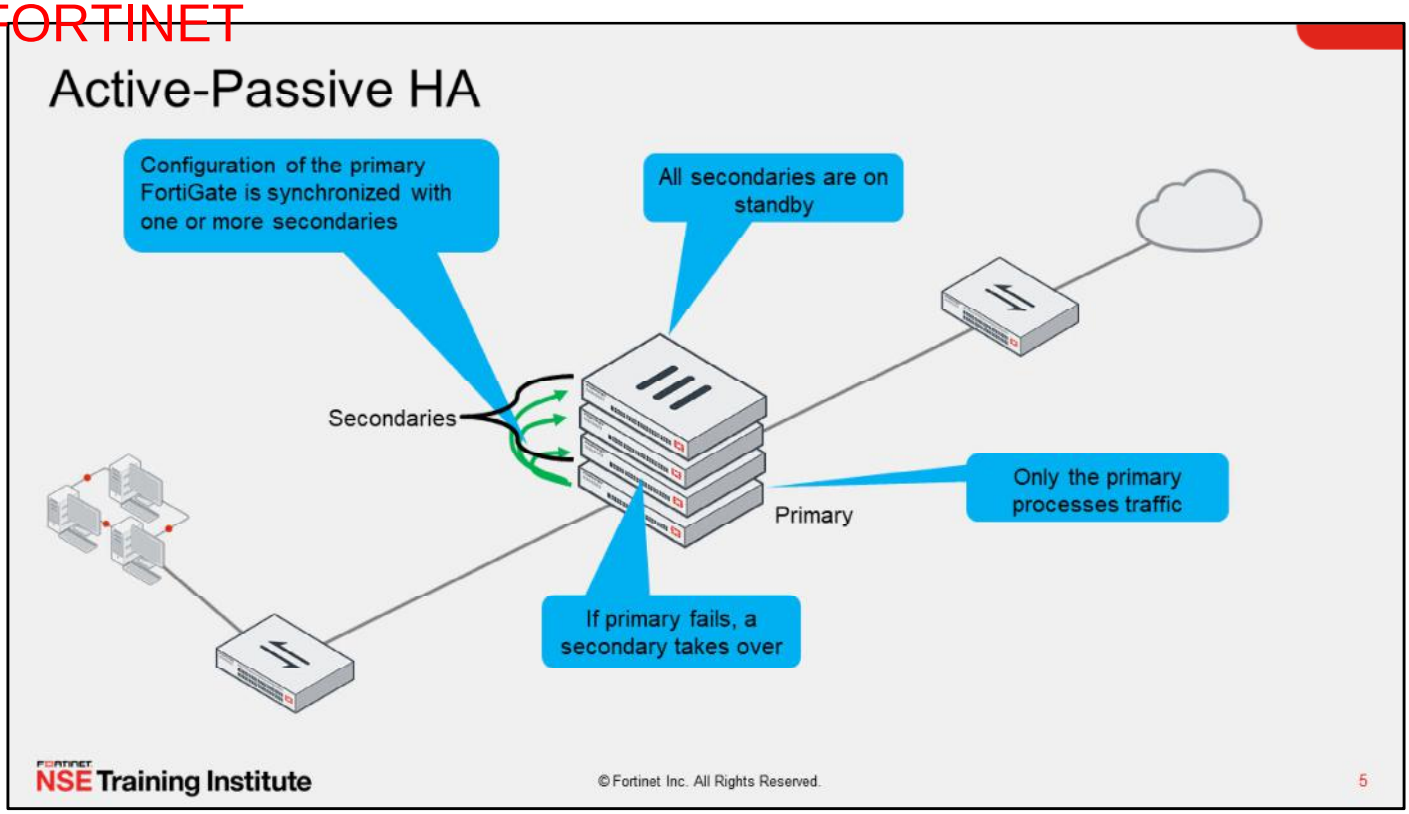

(slide contains animation)

First, take a look at active-passive mode. In either of the two HA operation modes, the configuration of the secondary FortiGate devices is synchronized with the configuration of the primary device.
(click)

In active-passive mode, the primary FortiGate is the only FortiGate device that actively processes traffic. Secondary FortiGate devices remain in passive mode, monitoring the status of the primary device.
(click)

If a problem is detected on the primary FortiGate, one of the secondary devices takes over the primary role. This event is called *HA failover*.

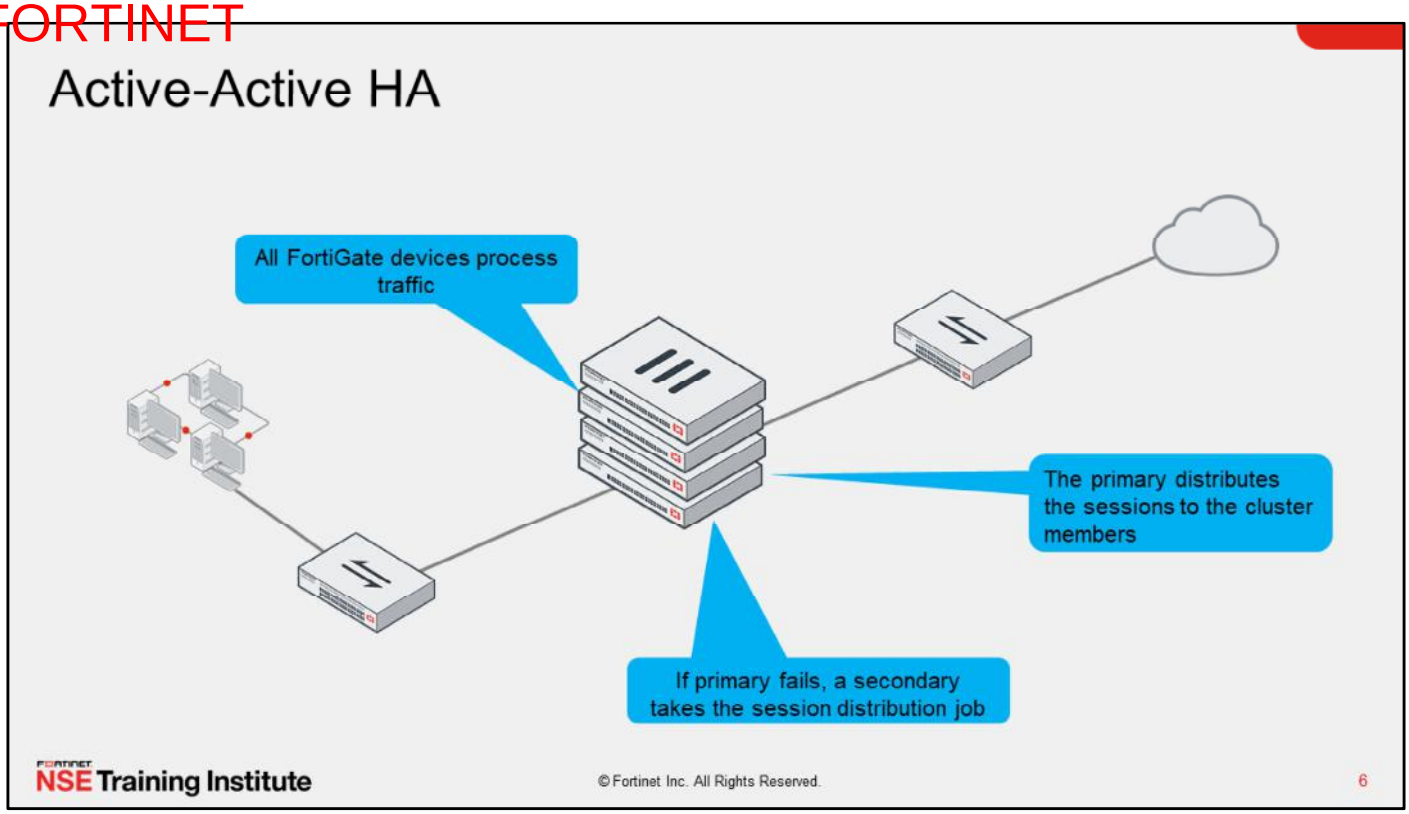

The other HA mode is active-active.

Like active-passive HA, in active-active HA, all FortiGate configurations are synchronized. Also, if a problem is detected on the primary device, one of the secondaries takes over the role of the primary, to process the traffic.

However, one of the main differences in active-passive mode is that in active-active mode, all the FortiGate devices are processing traffic. One of the tasks of a primary FortiGate in active-active mode is to balance some of the traffic among all the secondary devices.

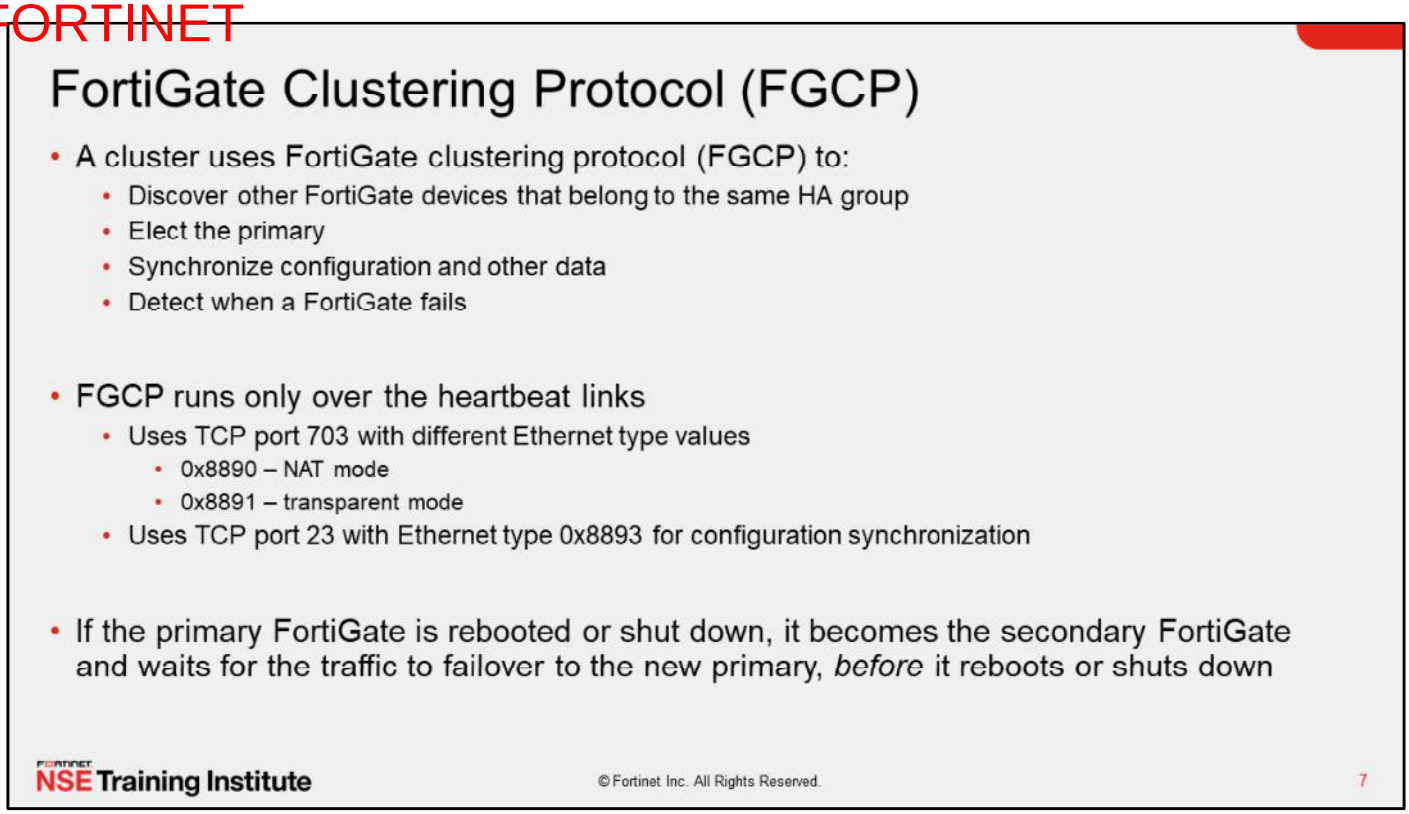

So how do the FortiGate devices in an HA cluster communicate?

FortiGate HA uses FGCP for HA-related communications. FGCP travels between the clustered FortiGate devices over the links that you have designated as the heartbeats.

You should create a heartbeat link between two FortiGate devices using a regular RJ45 or crossover cable. If you have another device between the two FortiGate devices, such as a switch, ensure that it is dedicated and isolated from the rest of your network. This way, critical FGCP traffic does not need to compete with the other traffic for bandwidth.

NAT mode clusters and transparent mode clusters use different Ethernet type values to discover and verify the status of other FortiGate devices in an operating cluster.

FortiGate devices in a cluster use Telnet sessions over TCP port 23, with Ethernet type 0x8893 over heartbeat links, to synchronize the cluster configuration and to connect to the CLI of another FortiGate in a cluster.

When you manually restart or shut down the primary FortiGate, before the primary FortiGate actually shuts down, it becomes the secondary device in an HA cluster, and waits for the traffic to failover to the new primary, before it shuts down or reboots.

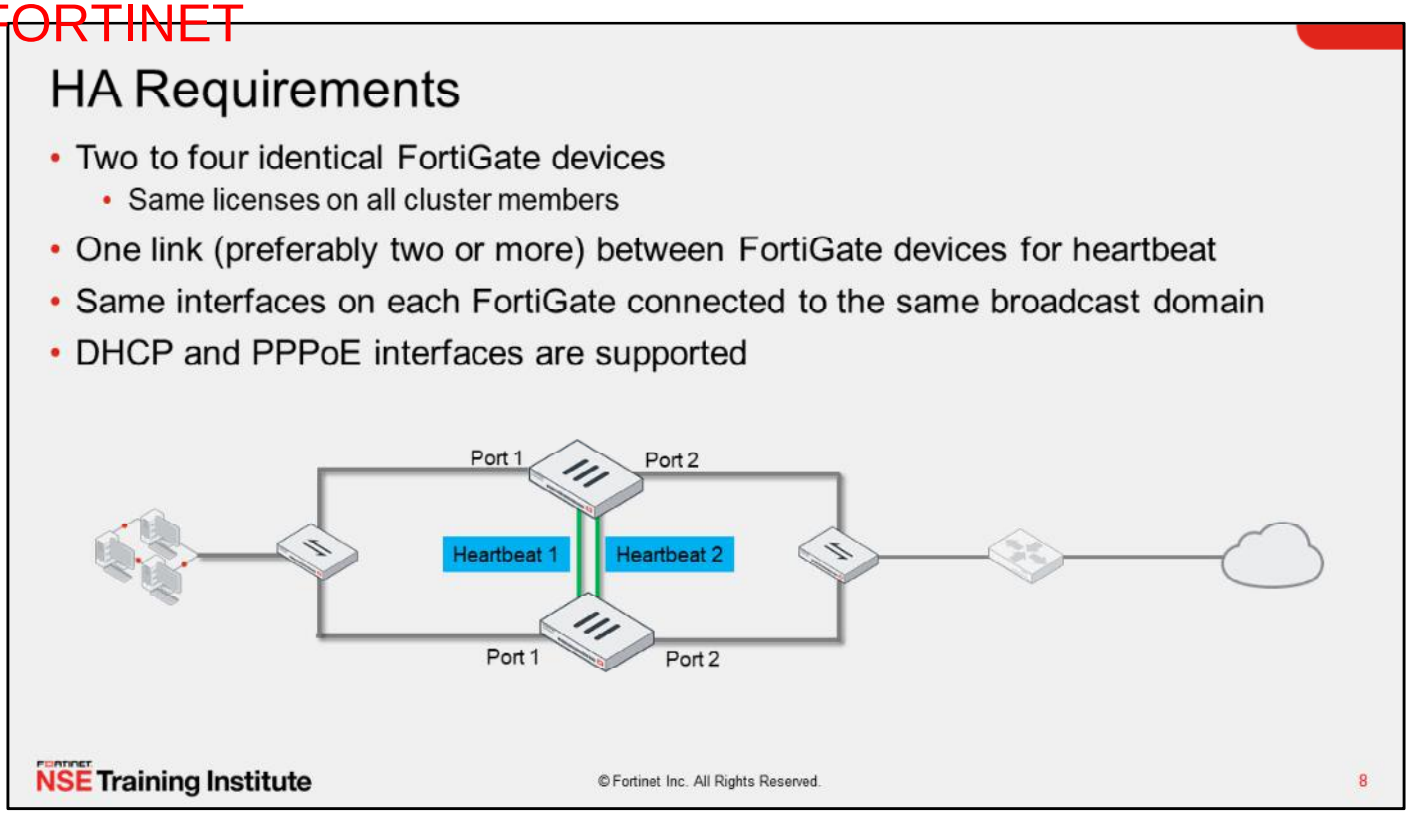

FortiGate HA configuration requires a specific setup and devices. First, the configuration requires at least two, but up to four, FortiGate devices with the same:

- Firmware
- Hardware model and VM license
- FortiGuard, FortiCloud, and FortiClient licenses
- Hard drive capacity and partitions
- Operating mode (transparent or NAT)

What if one of the FortiGate devices has a lower level of licensing than other FortiGate devices in the cluster? All of the FortiGate devices in the cluster revert to that lower licensing level. For example, if you purchase FortiGuard Web Filtering for only one of the FortiGate devices in a cluster, when the cluster is operating, none of the other cluster members support FortiGuard Web Filtering.

Second, the configuration requires at least one link between the FortiGate devices for HA communication. HA communication is called *heartbeat traffic*. For redundancy, you can create up to eight heartbeat interfaces. If one link fails, HA uses the next one, as indicated by priority and position in the heartbeat interface list.

Third, you must connect the same interfaces on each FortiGate device to the same switch or LAN segment. Note that in the example shown on the slide, the FortiGate devices are redundant to mitigate failure. But, the switches and their links are still a single point of failure. As you will see later, you can also have redundancy in the network switches and links.
As a best practice (and Fortinet recommendation), configure the FortiGate interfaces with static IP addresses when forming an HA cluster. Once an HA is formed, you can configure the DHCP or PPPoE addressing for an interface. If an interface is configured for DHCP or PPPoE, enabling HA may result in the interface receiving an incorrect address, or not being able to connect to the DHCP or PPPoE server correctly.

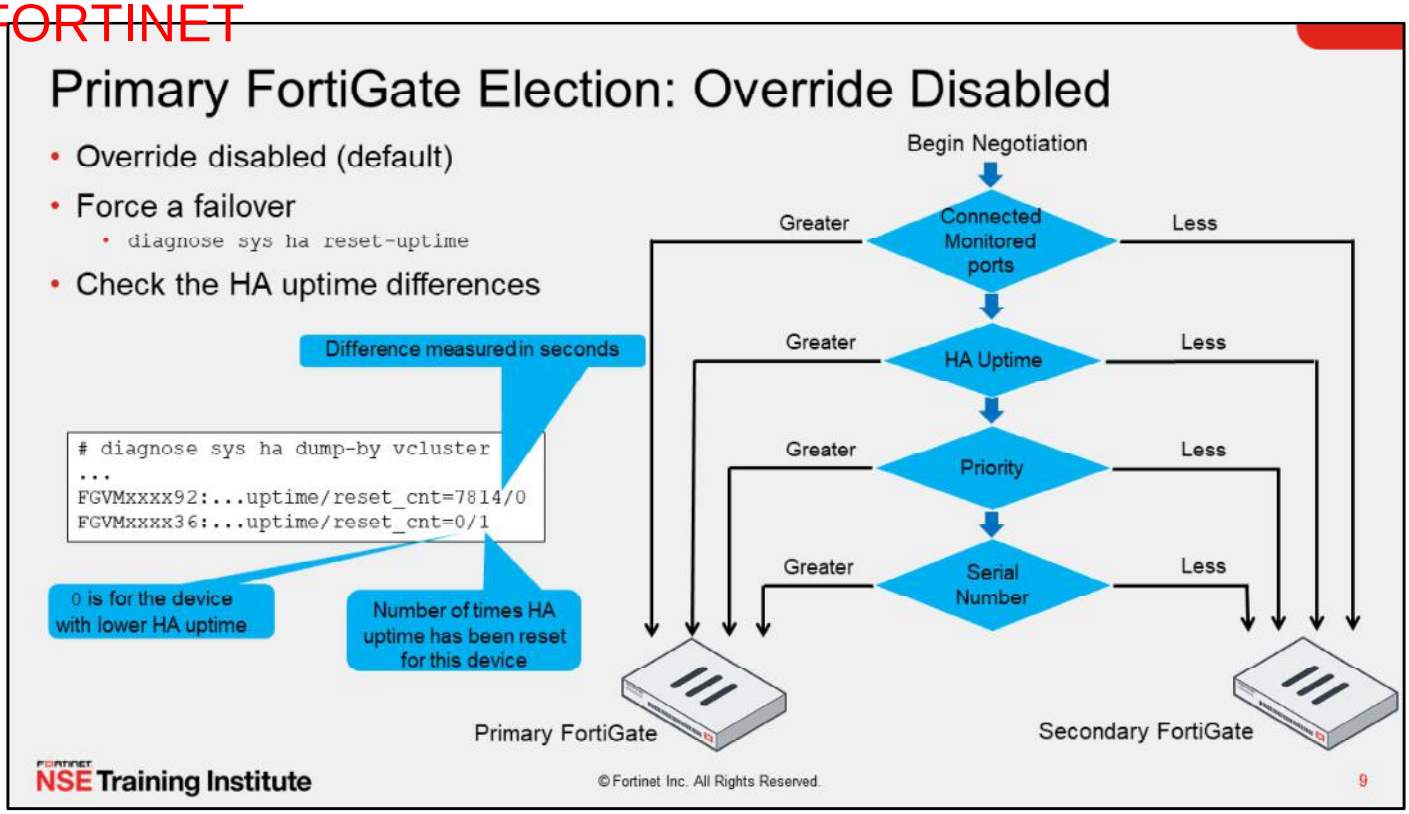

The process for electing the primary FortiGate depends on an HA setting called HA override. This slide shows the process and selection criteria that a cluster uses to elect the primary FortiGate when the HA override setting is disabled, which is the default behavior. Note that the selection process stops at the first matching criteria that successfully selects a primary FortiGate in a cluster.

1. The cluster first compares the number of monitored interfaces whose statuses are up. The FortiGate device with the most available monitored interfaces becomes the primary.
2. The cluster compares the HA uptimes of the individual devices. If the HA uptime of a device is at least five minutes more than the HA uptimes of the other FortiGate devices, it becomes the primary.
3. The FortiGate with the configured highest priority becomes the primary.
4. The cluster chooses the primary by comparing the serial numbers.

When HA override is disabled, the HA uptime has precedence over the priority setting. If, for any reason, you need to change which device is the current primary, you can manually force a failover event. When the override setting is disabled, the easiest way of doing this is by running the CLI command `diagnose sys ha reset-uptime` on the primary FortiGate.

Note that the `reset-uptime` command resets the HA uptime internally and does not affect the system up time displayed on the dashboard of a FortiGate. Also, if a monitored interface fails, or a FortiGate in a cluster reboots, the HA uptime for that FortiGate is reset to 0.
Note that you can view the HA uptime difference between the cluster members. The device with `0` in the `uptime` column indicates the device with lower uptime. In this example, the device ending with serial number `92` has an HA uptime `7814` seconds greater than the other device in the HA cluster. The `reset _cnt` column indicates the number of times HA uptime has been reset for that device.

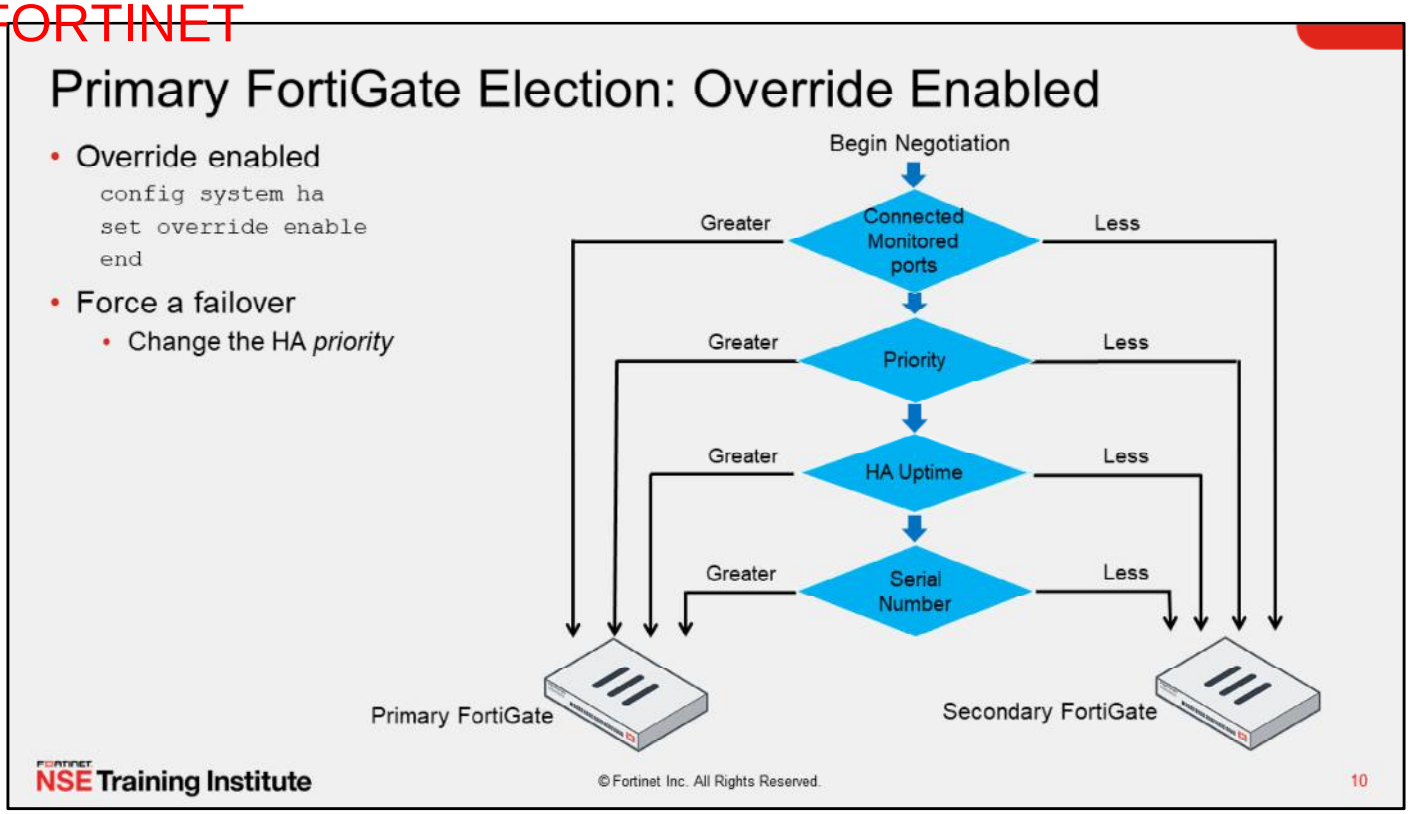

You can alter the order of the selection criteria that clusters consider when electing the primary FortiGate.

If the HA override setting is enabled, 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. When a primary becomes available again, it takes back its primary role from the secondary FortiGate that temporarily replaced it.

Note that the selection process stops at the first matching criteria that successfully selects a primary FortiGate in a cluster.

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 in one of the secondaries, or decrease the priority in the primary.

The override setting and device priority values are not synchronized to all cluster members. You must enable override and adjust device priority manually and separately for each cluster member.

Knowledge Check

1. To form an HA cluster, *all* FortiGate devices that will be included in the cluster must have which of the following?
   - A. The same FortiGate hostname
   - ✓B. The same firmware

2. What is the default criteria (override disabled) for selecting the HA primary device in an HA cluster?
   - ✓A. Connected monitored ports > HA uptime > priority > serial number
   - B. Priority > HA uptime > connected monitored ports > serial number

NSE Training Institute

11

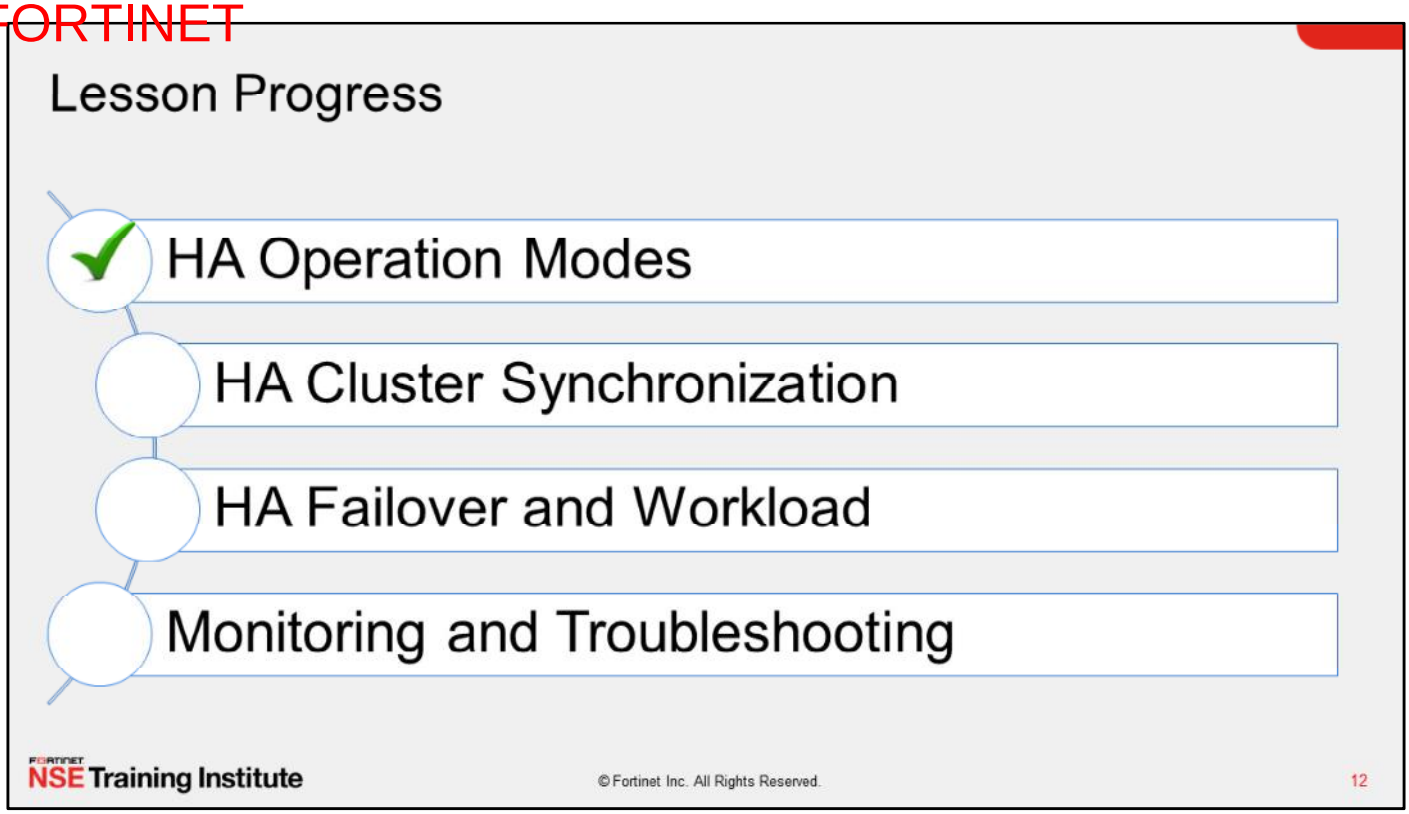

Good job! You now understand HA operation modes and the election of the primary FortiGate in an HA cluster.

Now, you will learn about HA cluster synchronization.

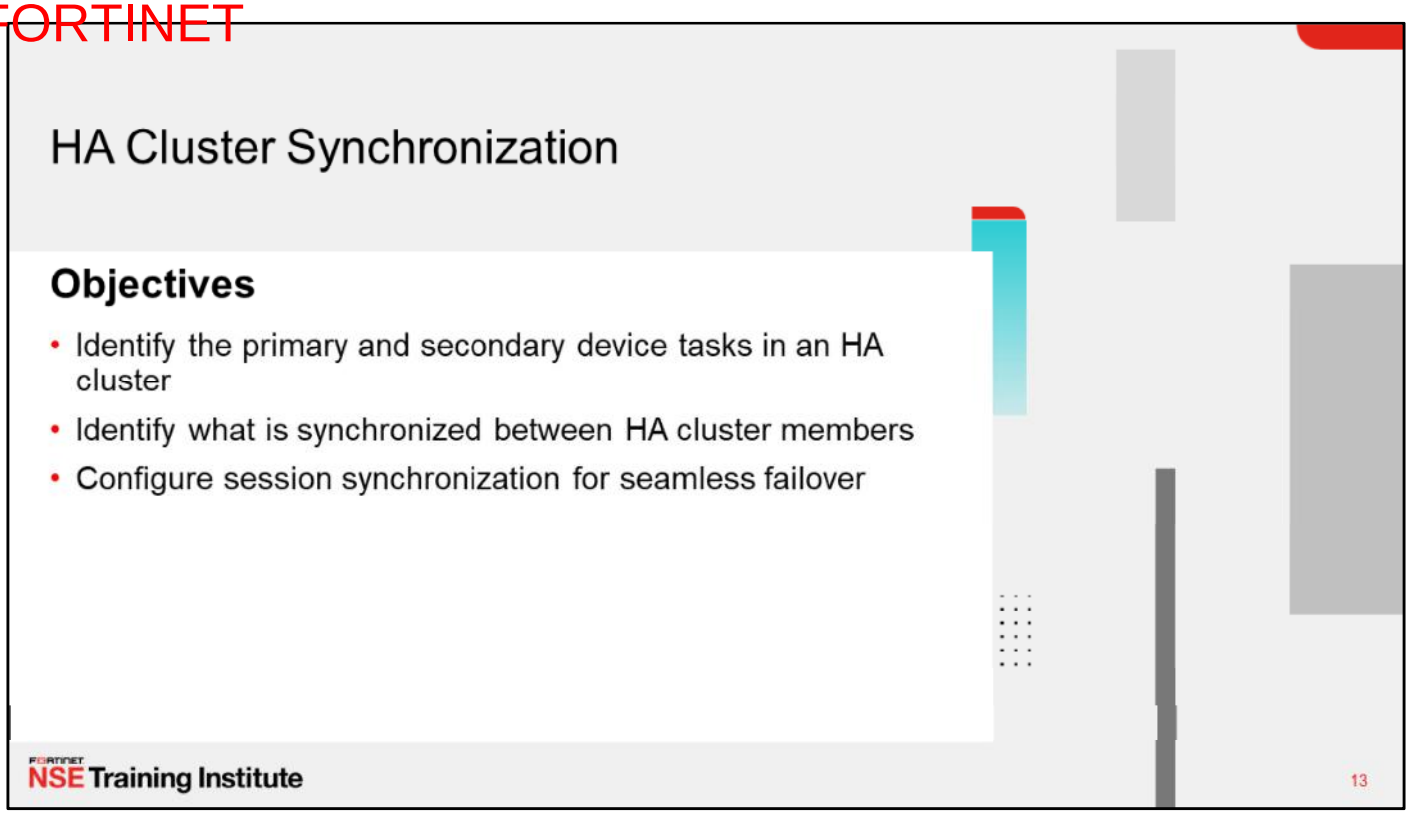

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in cluster synchronization, you will be able to identify the tasks of FortiGate devices and what is synchronized between the cluster members. You will also learn how to configure and use session synchronization for specific types of traffic for seamless failover.

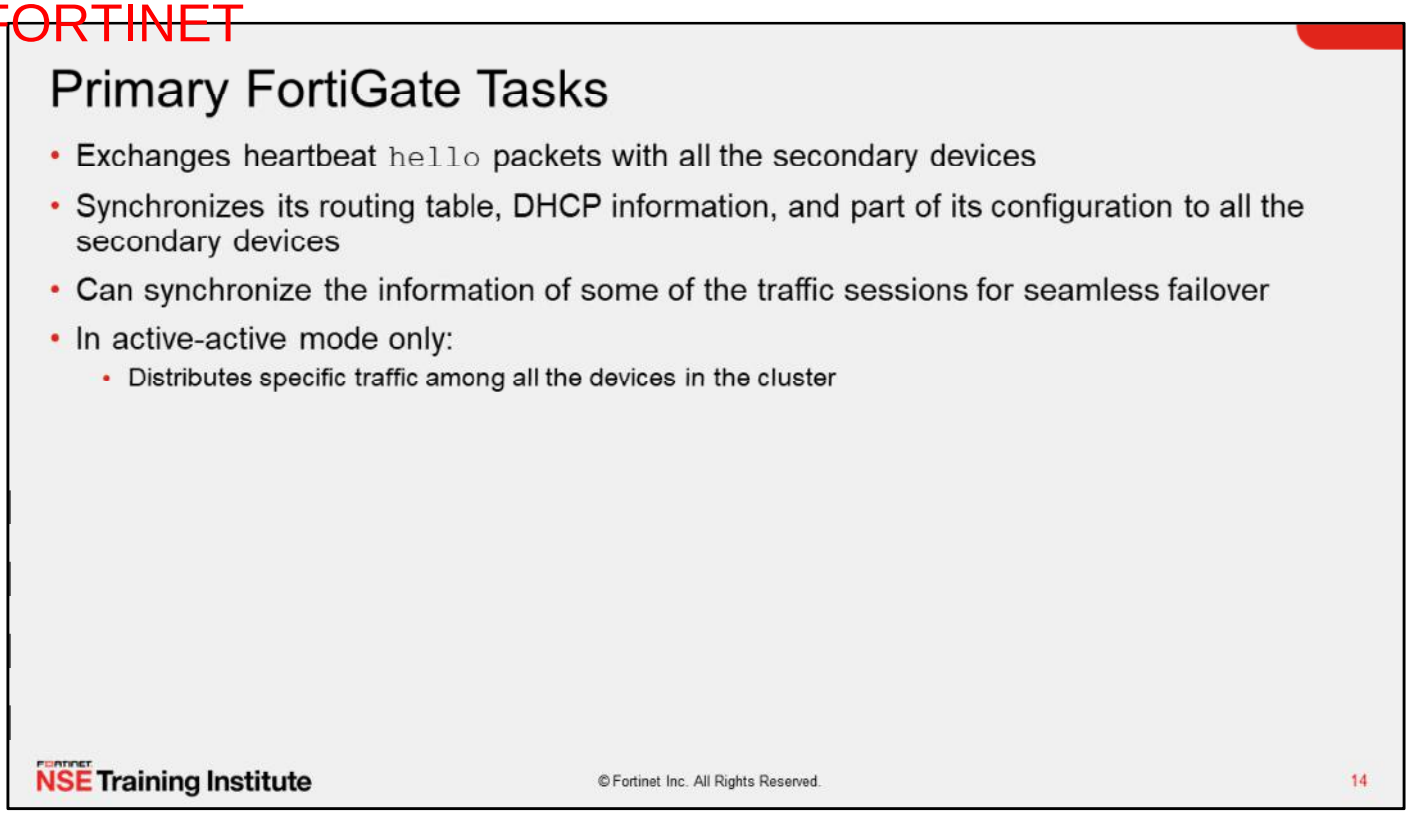

So, what are the tasks of a primary FortiGate?

It monitors the cluster by sending `hello` signals and listening for replies, to identify if other FortiGate devices are alive and available. It also synchronizes its routing table, DHCP information, and part of its configuration with the other devices.

Optionally, you can configure the primary FortiGate to synchronize some of the traffic session information to all the secondary devices. This allows a faster, seamless failover for some sessions. Some applications do not need to reestablish their sessions after a failure of a primary FortiGate. You will learn which session information you can synchronize later in the lesson.

In active-active mode only, a primary FortiGate also distributes specific traffic among all the available devices in the cluster.

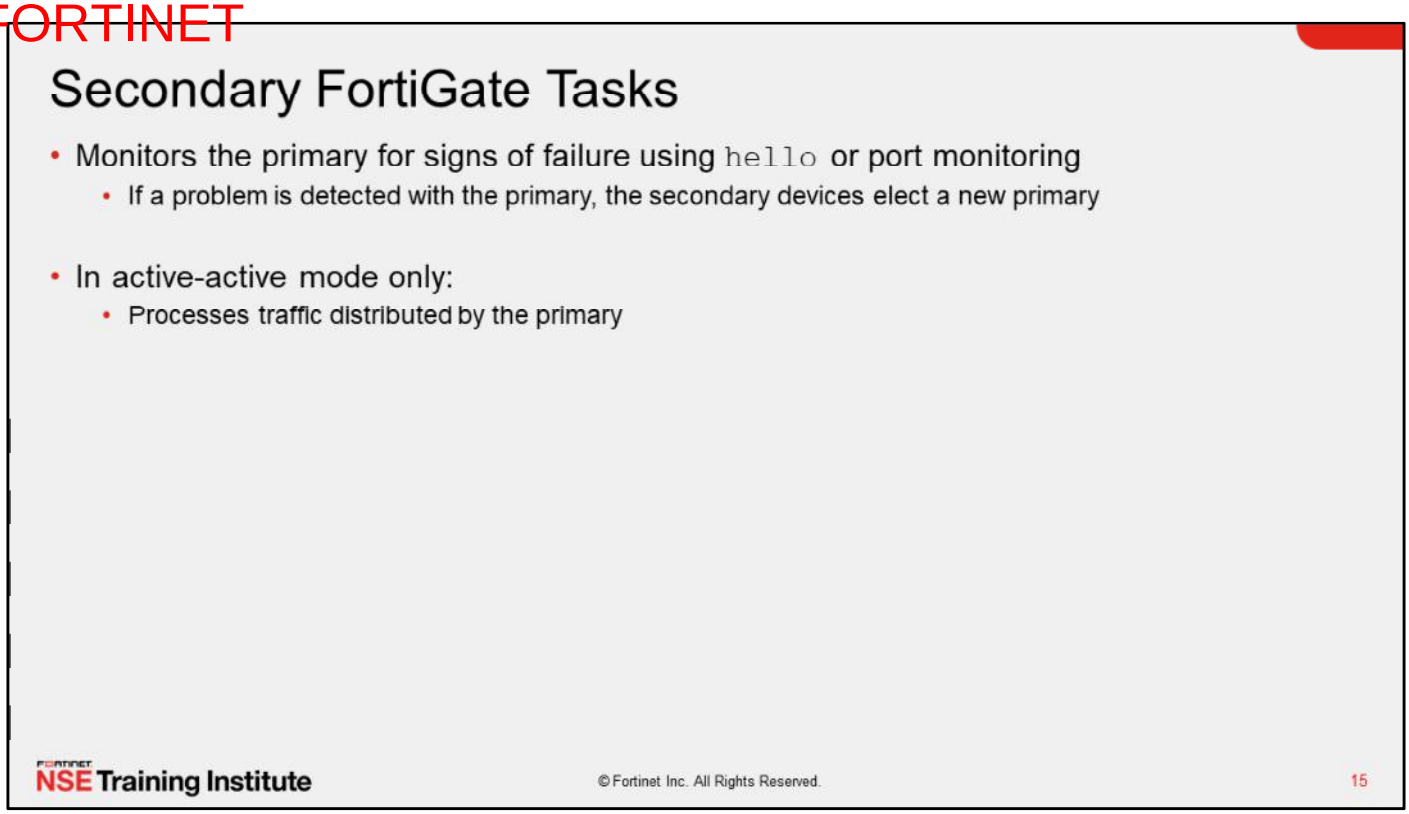

Now, take a look at the tasks of secondary FortiGate devices.

If the mode is active-passive, the secondaries simply wait, receiving synchronization data but not actually processing any traffic. If the primary FortiGate fails, the secondaries elect a new primary.

In active-active mode, the secondaries don’t wait passively. They process all traffic assigned to them by the primary device.

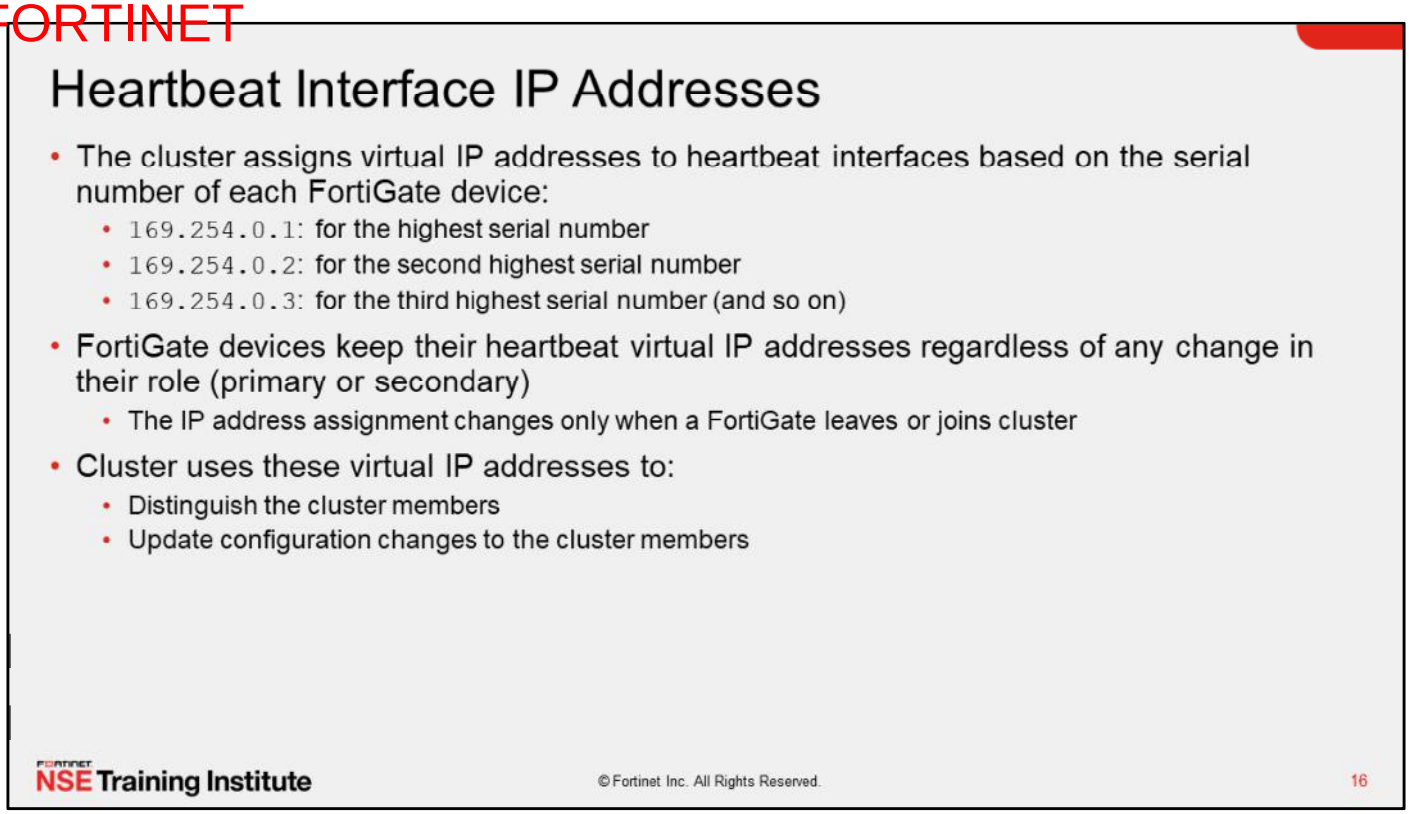

What about the heartbeat interfaces?

You don’t need to configure heartbeat interfaces. The FortiGate clustering protocol automatically negotiates 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 virtual IP address remains the same.

A change in the heartbeat IP addresses might 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 these virtual IP addresses to distinguish the cluster members and update configuration changes to the cluster members.

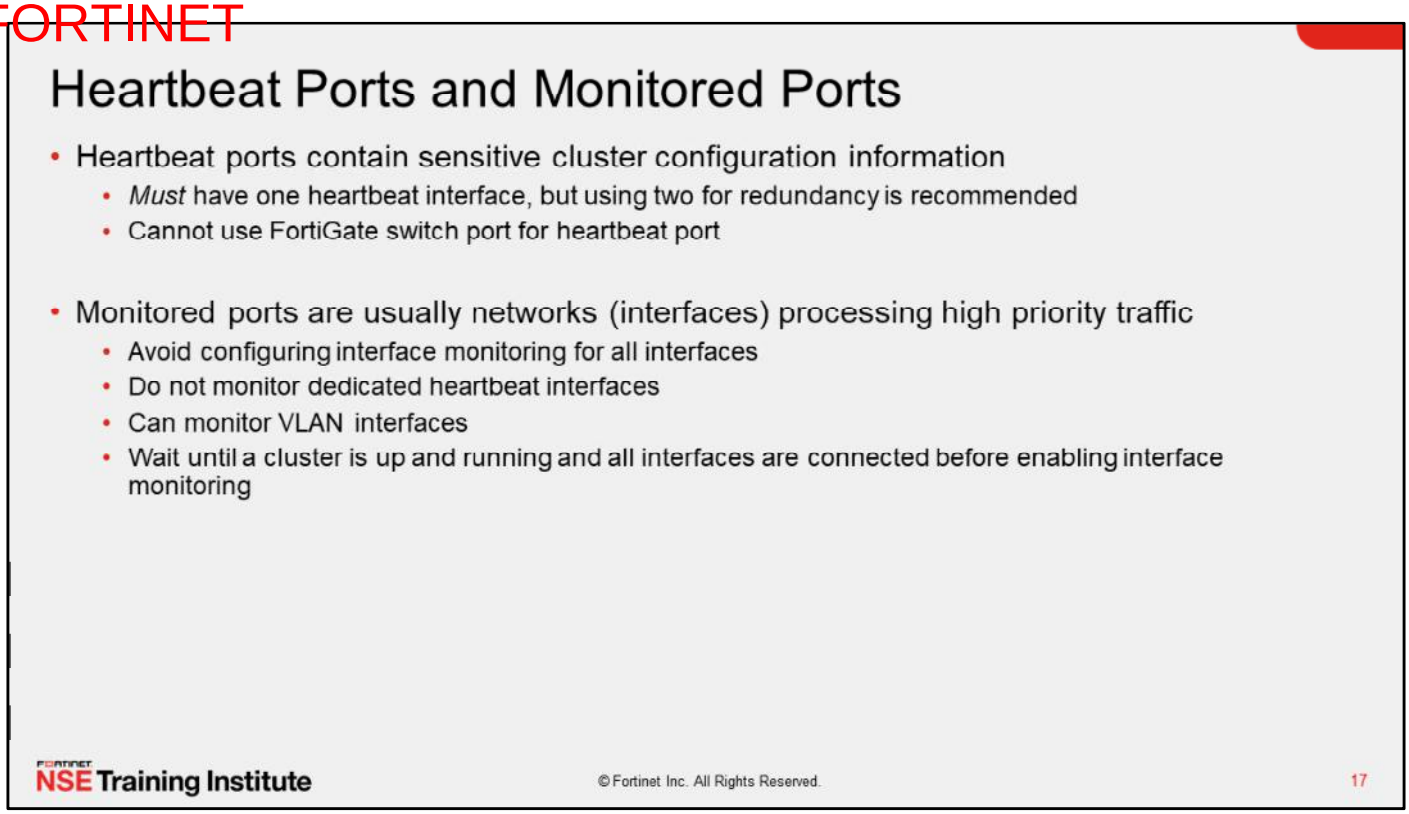

There are a few items that you must consider when connecting heartbeat interfaces and configuring interface monitoring:

- Heartbeat ports contain sensitive information about cluster configuration and require a fair amount of bandwidth to make sure cluster configurations are in a synchronized state at all times. You must have at least one port for the heartbeat traffic, and preferably two. As a best practice, configure an alias for the heartbeat interfaces. It helps to identify what these interfaces are being used for in an HA cluster.

Note that you can enable heartbeat communication for physical interfaces, but not for VLAN subinterfaces, IPsec VPN interfaces, redundant interfaces, 802.3ad aggregate interfaces, or FortiGate switch ports.

- You should configure interface monitoring only for those ports whose failure should trigger a device failover (for example, high-priority traffic ports). You should not configure port monitoring for dedicated heartbeat ports.

As a best practice, wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.

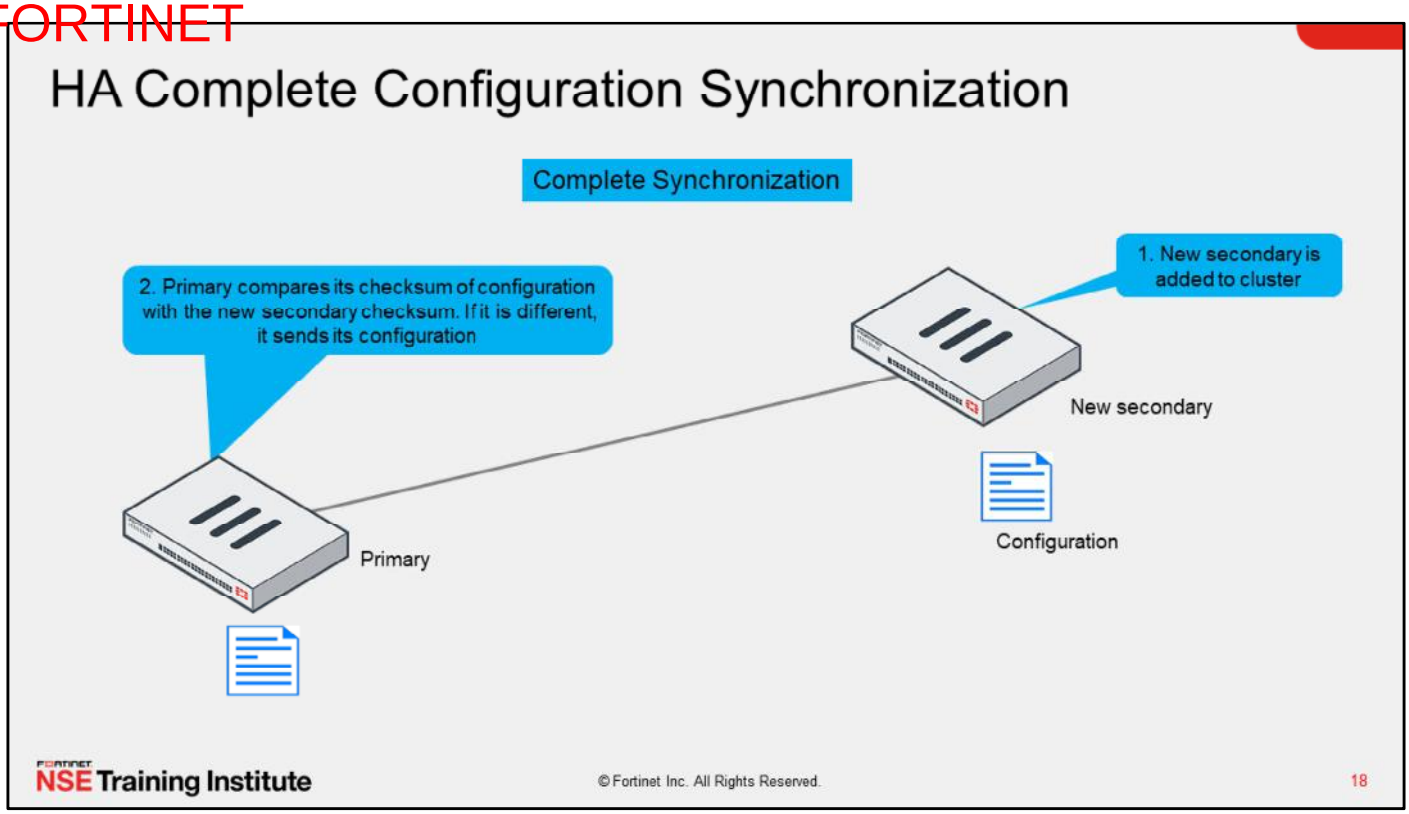

To prepare for a failover, an HA cluster keeps its configurations in sync. You will explore that now.

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.

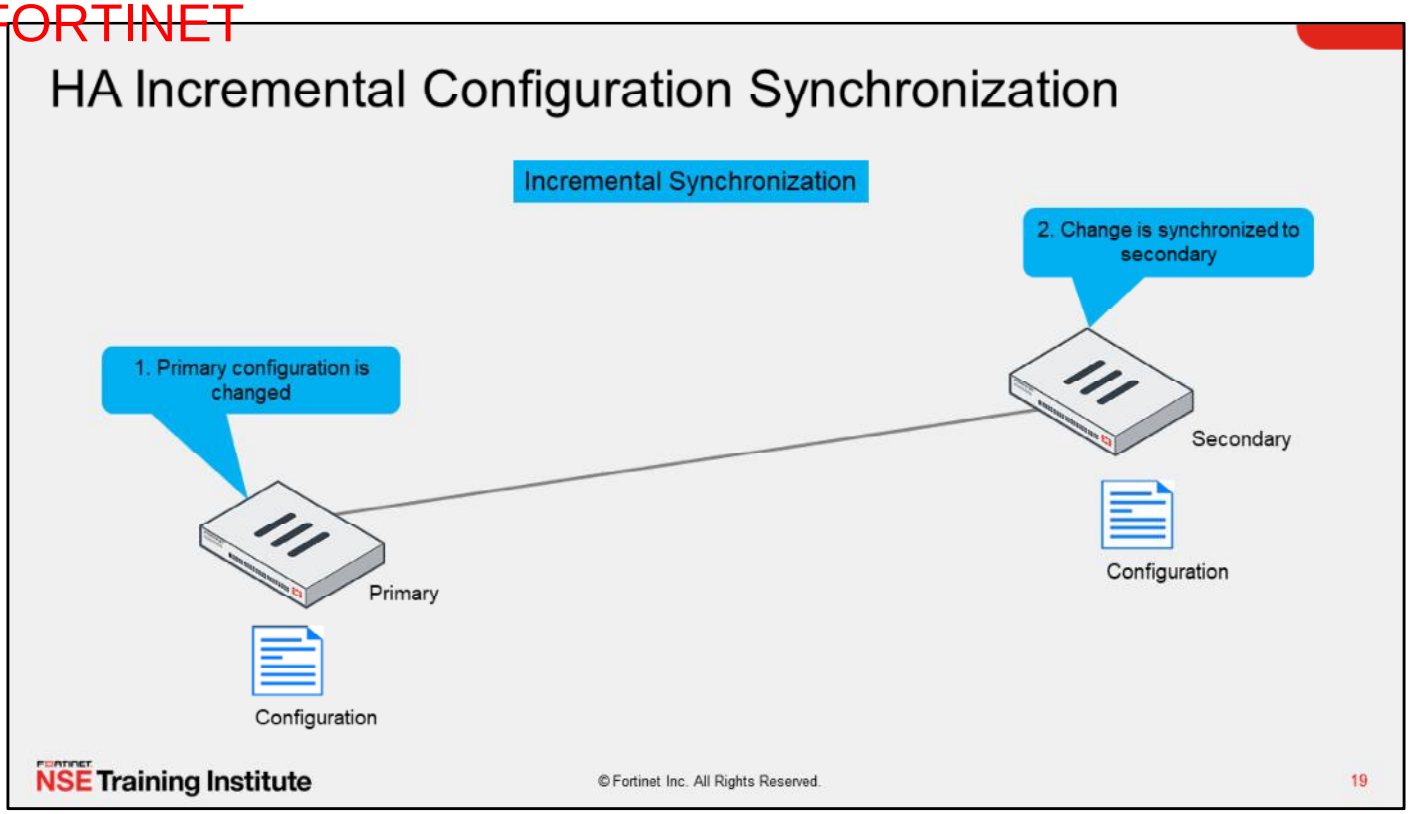

After the initial synchronization is complete, the primary sends any further configuration changes made by an administrator to all the secondaries. For example, if you create a firewall address object, the primary doesn't resend its complete configuration, it sends only the new object.

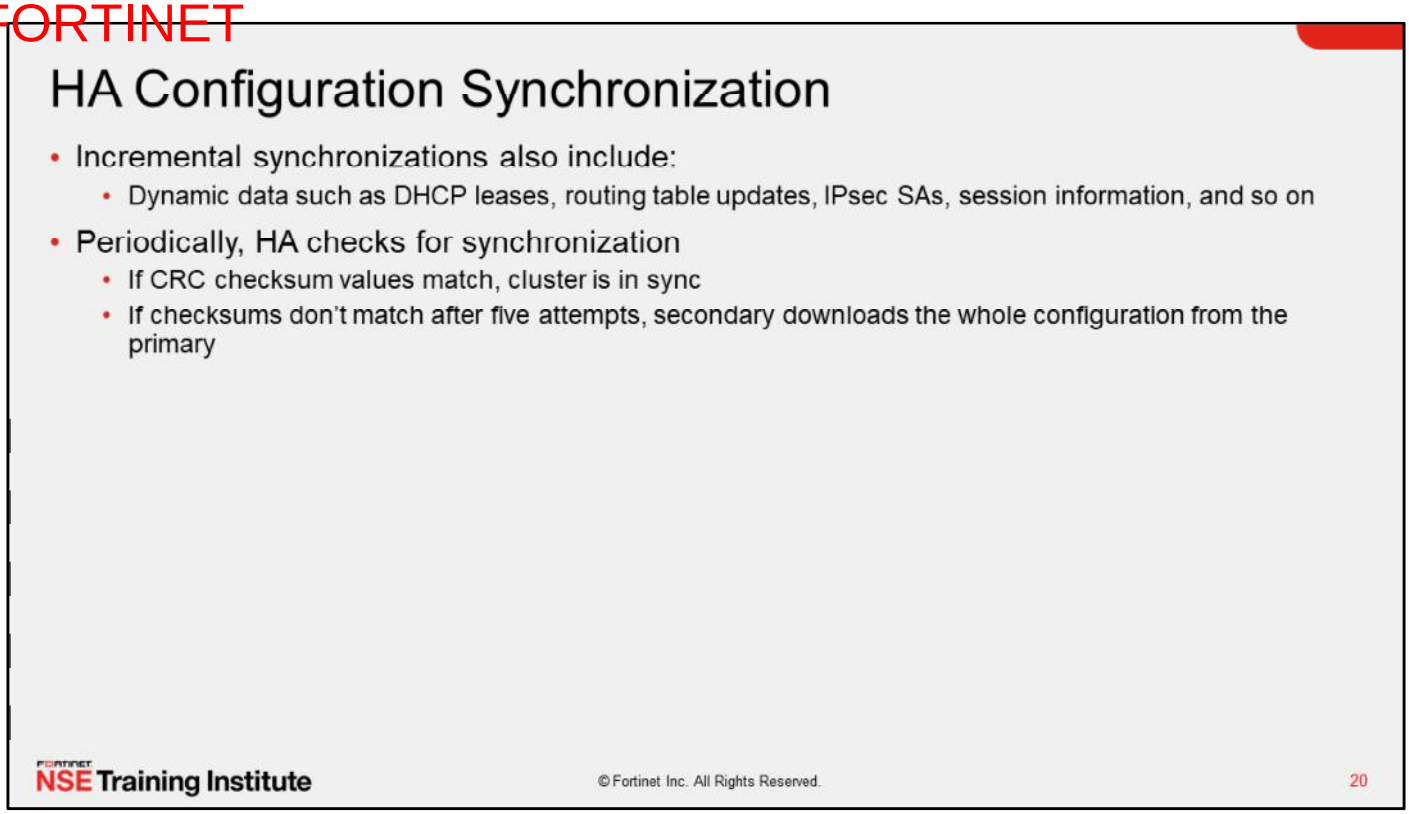

HA propagates more than just configuration details. Some runtime data, such as DHCP leases and routing tables, are also synchronized.

By default, the cluster checks every 60 seconds to ensure that all devices are synchronized. If any secondary is out of sync, the checksum of secondary devices is then checked every 15 seconds. If checksums don't match for five consecutive checks, a complete resynchronization is done.

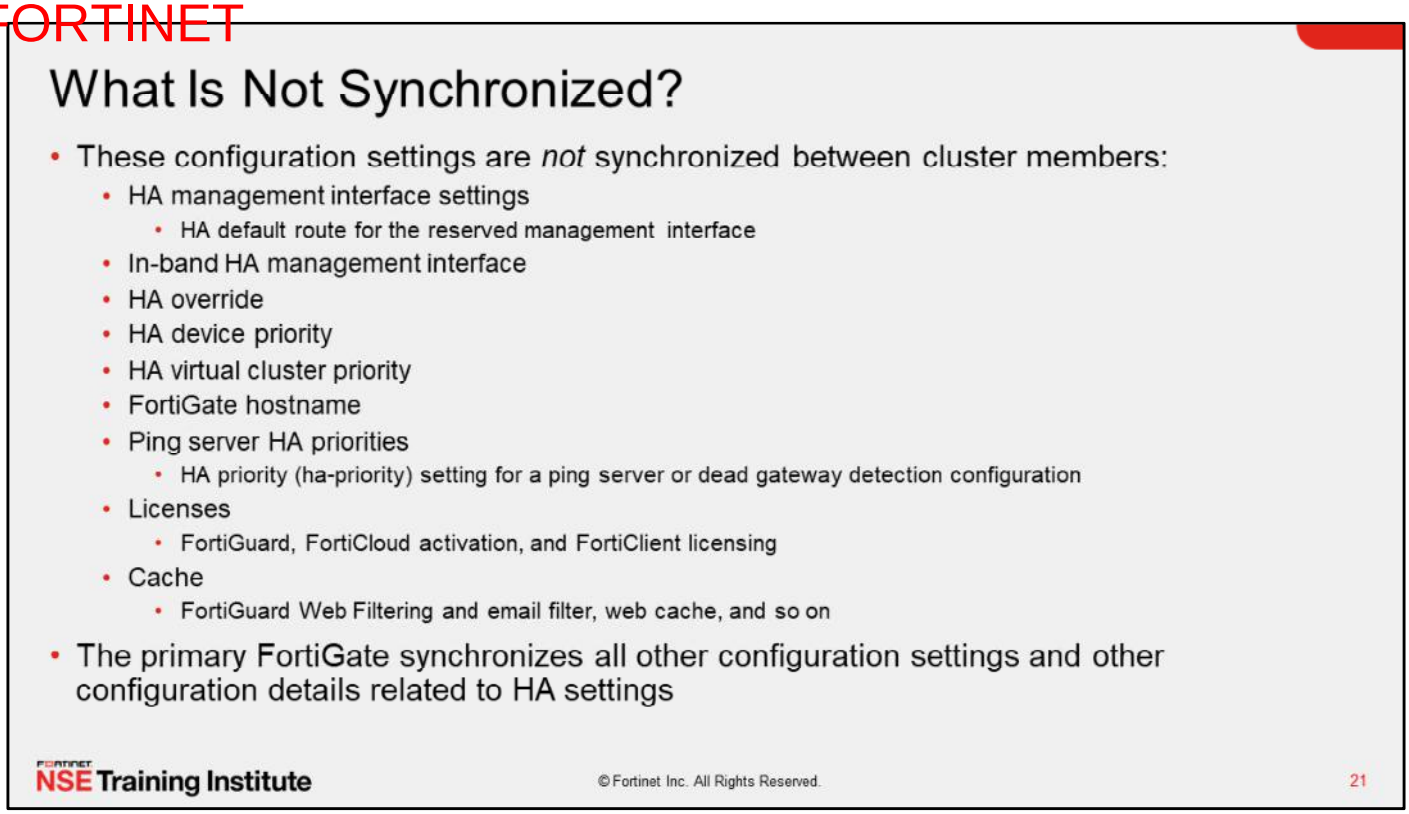

Not all the configuration settings are synchronized. There are a few that are not, such as:

- The 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
- The virtual cluster priority
- The FortiGate host name
- The HA priority setting for a ping server (or dead gateway detection) configuration
- Licenses
- Caches

The primary FortiGate synchronizes all other configuration settings, including other configurations related to HA settings.

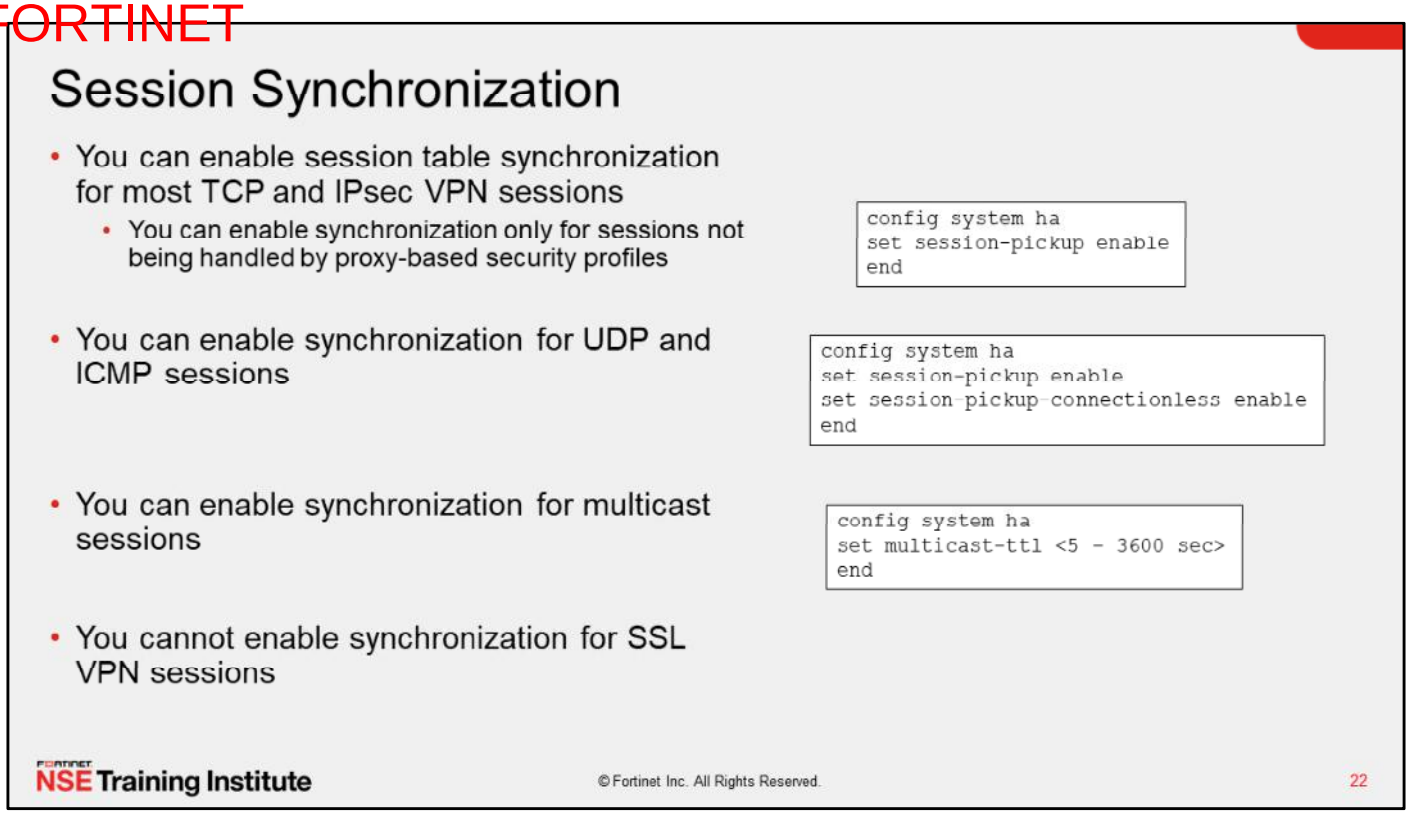

Session synchronization enables seamless failover for some traffic. The information of some sessions is synchronized, so when the primary fails, the new primary can take over those sessions where they were left and keep them open. Traffic might be interrupted for a few seconds, but the network applications don't need to reconnect the sessions again.

Once you enable session synchronization, the device synchronizes TCP and IPsec VPN sessions that comply with one requirement: they are not handled by proxy-based security profiles. However, sessions using flow-based security profiles are supported, but failed over sessions are no longer inspected by security profile functions.

Note that if both flow-based and proxy-based security profile features are applied to a TCP session, that session will not resume after a failover.

You can optionally enable the synchronization of UDP and ICMP sessions. Although both protocols are sessionless, entries are created in the FortiGate session table for each UDP and ICMP traffic flow. Usually, this synchronization is not required, because most of the network applications based on UDP or ICMP are able to keep the communication even when their session information is lost.

You can also enable synchronization of multicast sessions. The multicast time to live (TTL) timer controls how long to keep synchronized multicast routes on the secondary devices in the HA cluster so they are present on the secondary devices when it becomes the new primary device after a failover.

The synchronization of SSL VPN sessions is not supported.

Knowledge Check

1. Which information is synchronized between two FortiGate devices that belong to the same HA cluster?
   - ✓ A. Firewall policies and objects
   - B. FortiGate hostname

2. Which one of the following session types can be synchronized in an HA cluster?
   - A. SSL VPN sessions
   - ✓ B. IPsec VPN sessions

NSE Training Institute

23

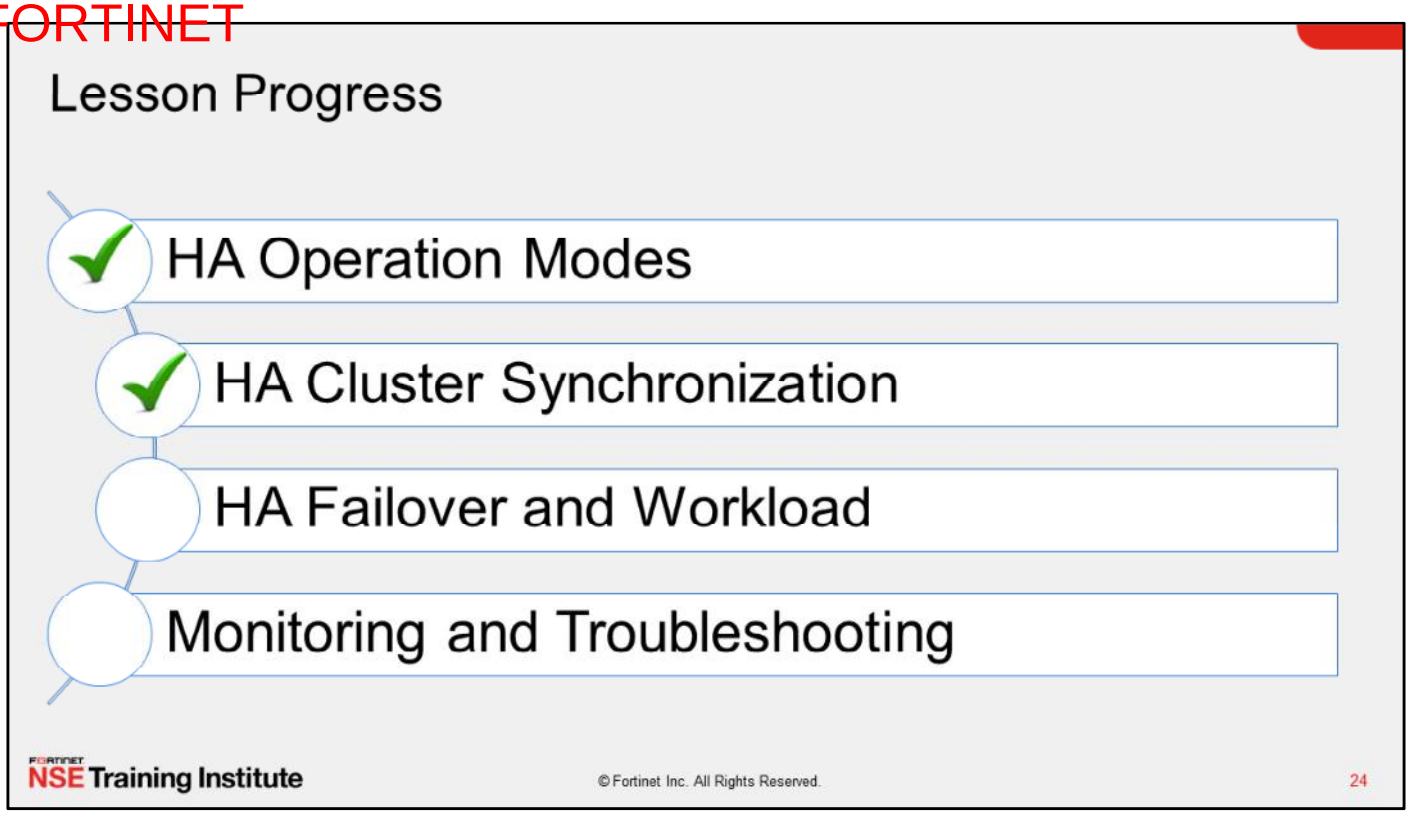

Good job! You now understand HA cluster synchronization.

Now, you will learn about HA cluster failover protection types and workload for primary and secondary FortiGate devices in an HA cluster.

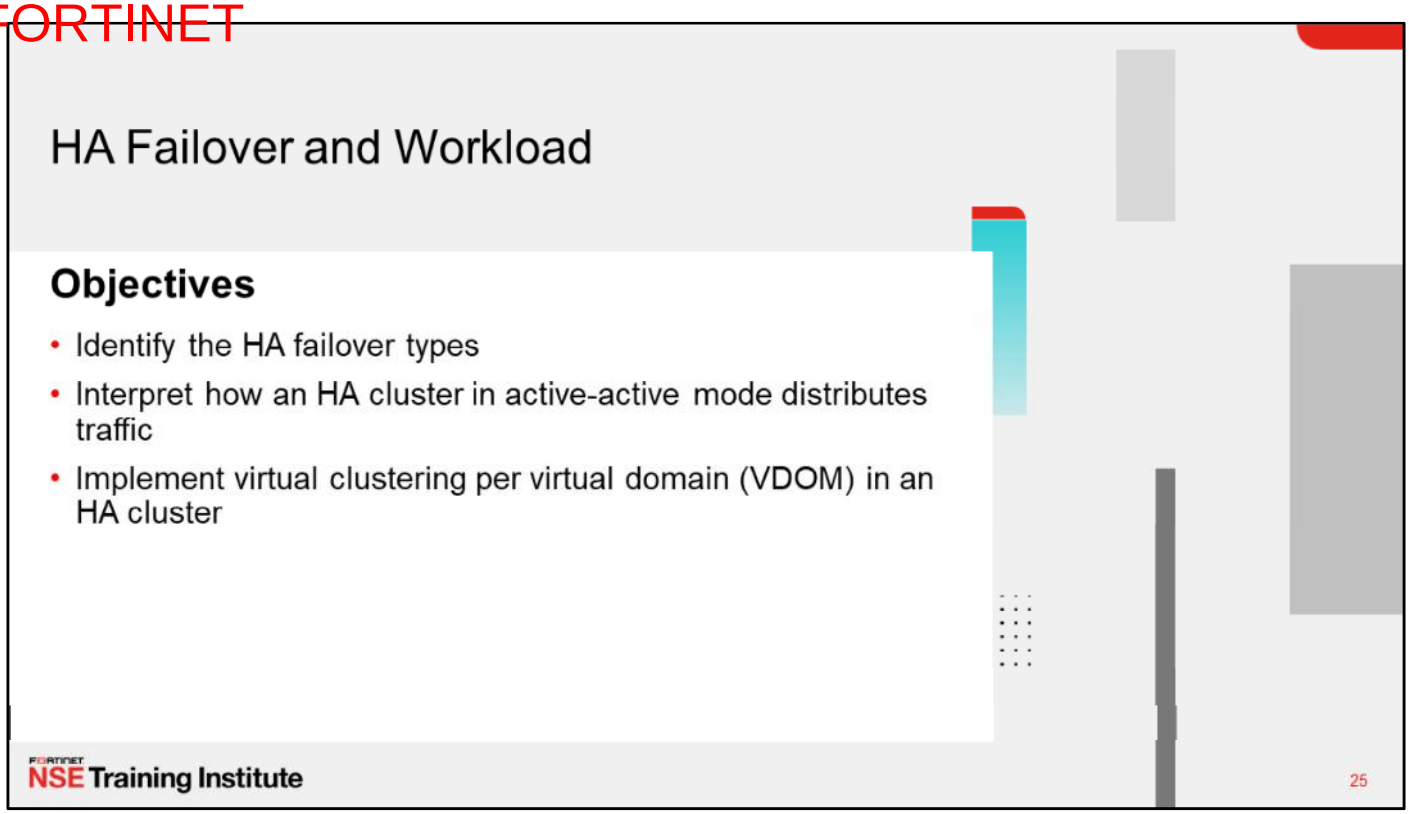

After completing this section, you should be able to achieve the objectives shown on this slide.

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 distribution of traffic in an active-active cluster and distributing traffic using virtual clustering.

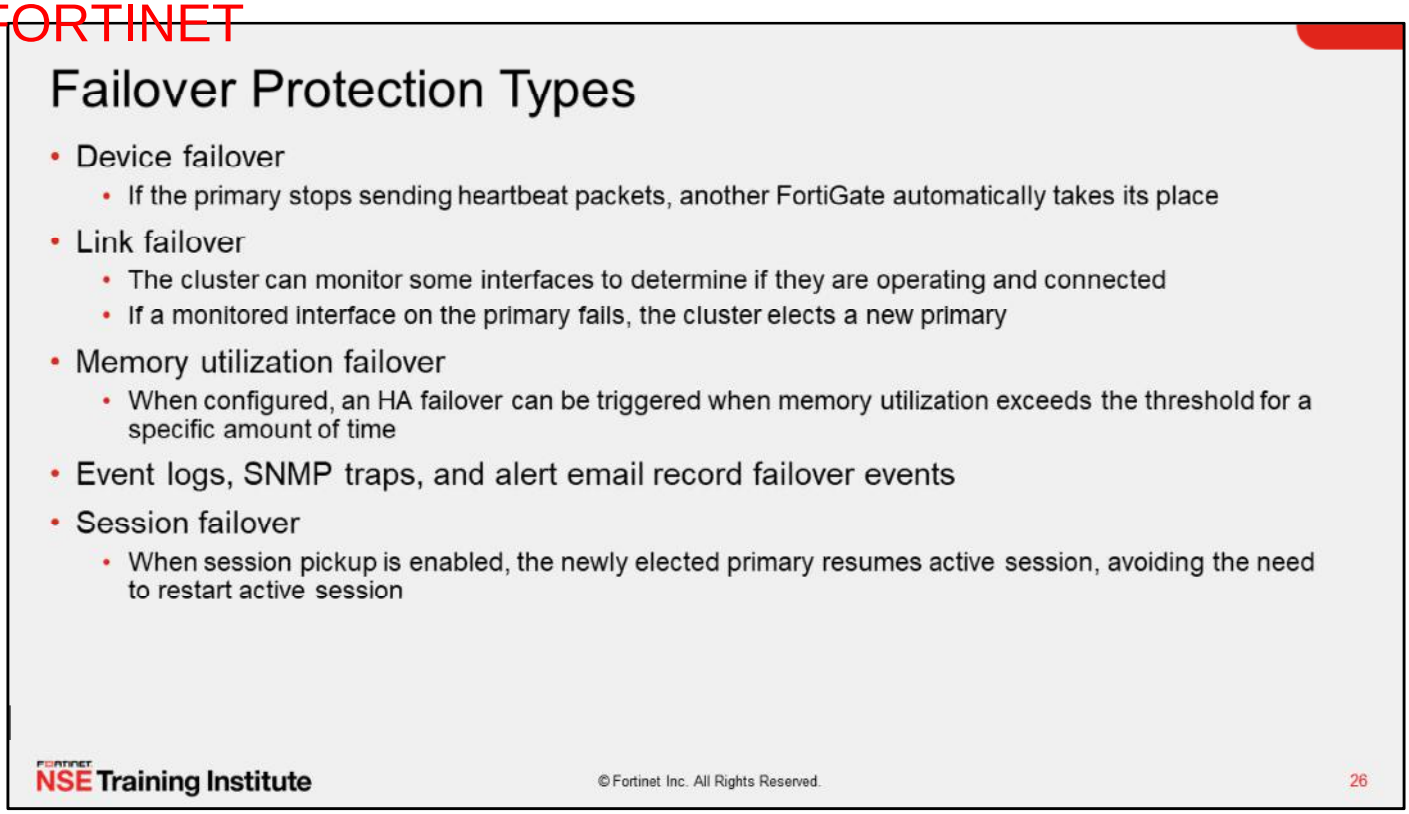

The most common types of failovers are device failovers and link failovers.

A device failover is triggered when the primary FortiGate stops sending heartbeat traffic. When this happens, the secondaries renegotiate a new 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 the link status of some interfaces. If a monitored interface on the primary FortiGate is unplugged, or its link status goes down, a new primary FortiGate is elected.

If you enable the memory-based failover, an HA failover can be triggered when memory utilization exceeds the threshold for a specific amount of time. You can enable the memory-based failover through `config system ha` CLI command.

There are multiple events that might trigger an HA failover, such as hardware or software failure in the primary FortiGate or an issue in one of the interfaces on the primary. When a failover occurs, an event log is generated. Optionally, the device can also generate an SNMP trap and an alert email.

If session pickup is enabled for the type of traffic you want to synchronize among cluster members, the sessions are resumed in the event of device failover or link failover.

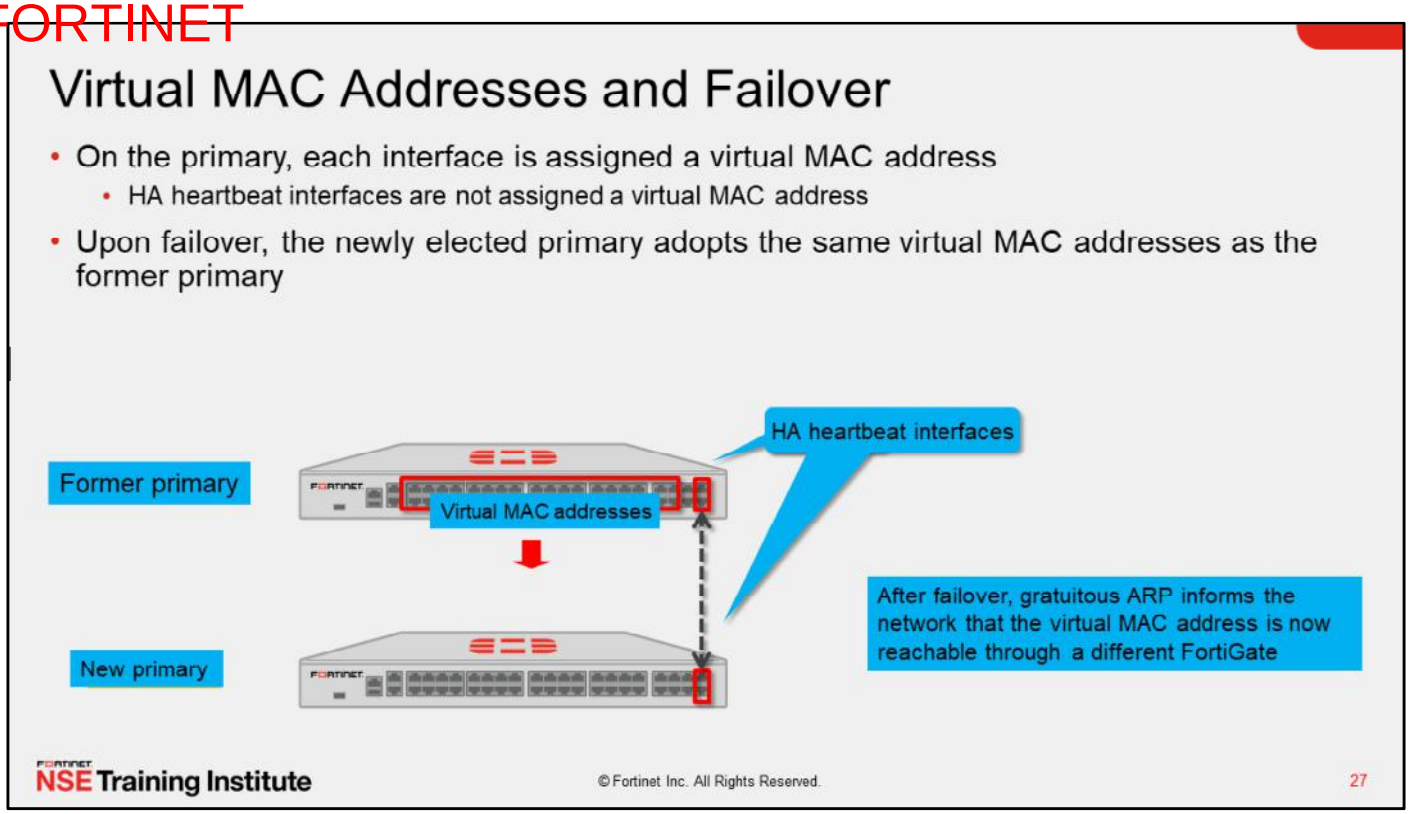

To forward traffic correctly, a FortiGate HA solution uses virtual MAC addresses. When a primary joins an HA cluster, each interface is assigned a virtual MAC address. The HA group ID is used in the creation of virtual MAC addresses assigned to each interface. So, if you have two or more HA clusters 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 to assign different HA group IDs to each cluster.

Through the heartbeats, the primary informs all secondaries 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.

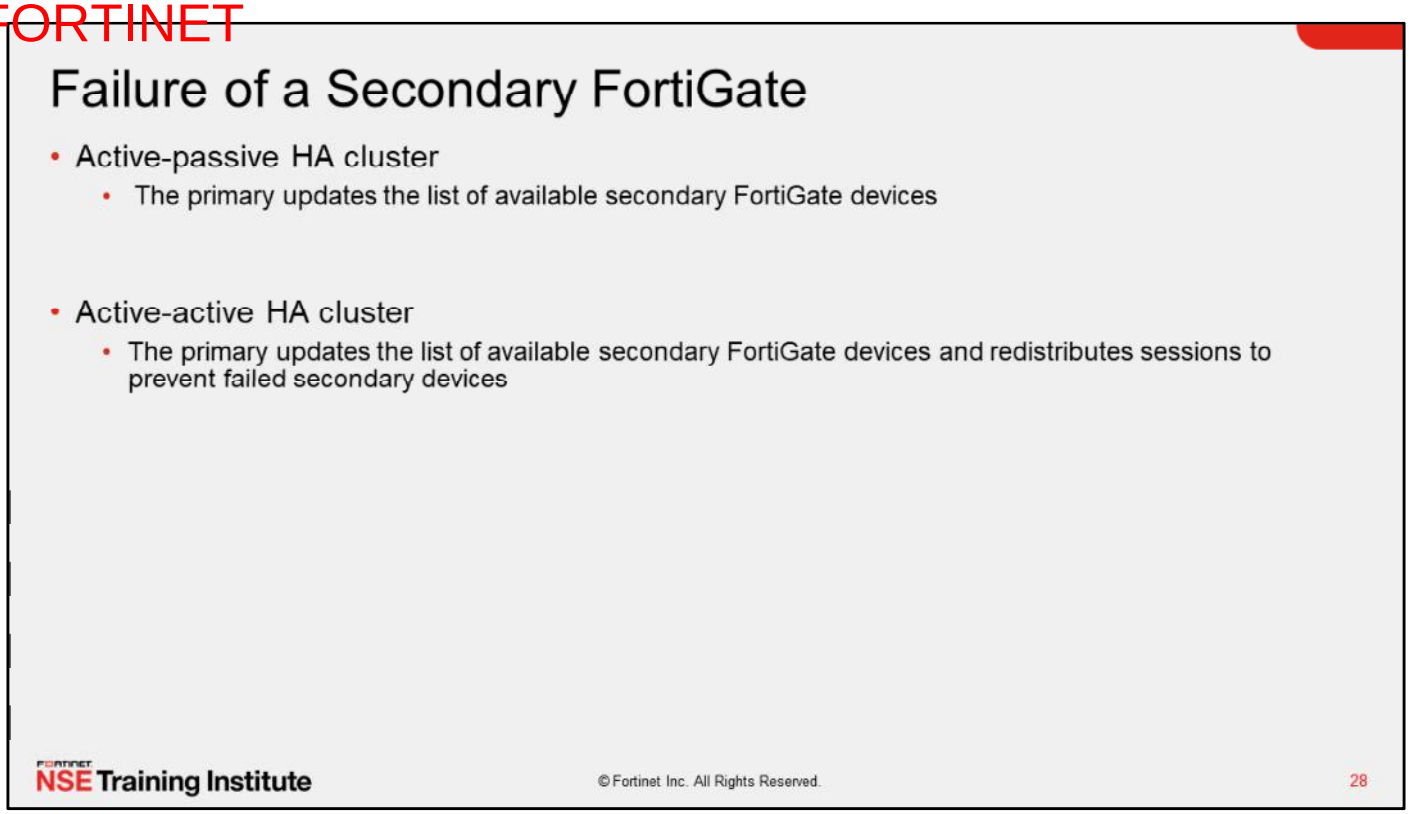

As you learned earlier in this lesson, if a primary fails, a new primary is elected. But what happens if a secondary FortiGate device fails? It depends on the HA mode.

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, all secondaries are handling traffic. So, the primary (which tracks and assigns sessions to each secondary) must not only update its list of available secondary FortiGate devices, it must also reassign sessions from the failed FortiGate to a different secondary FortiGate.

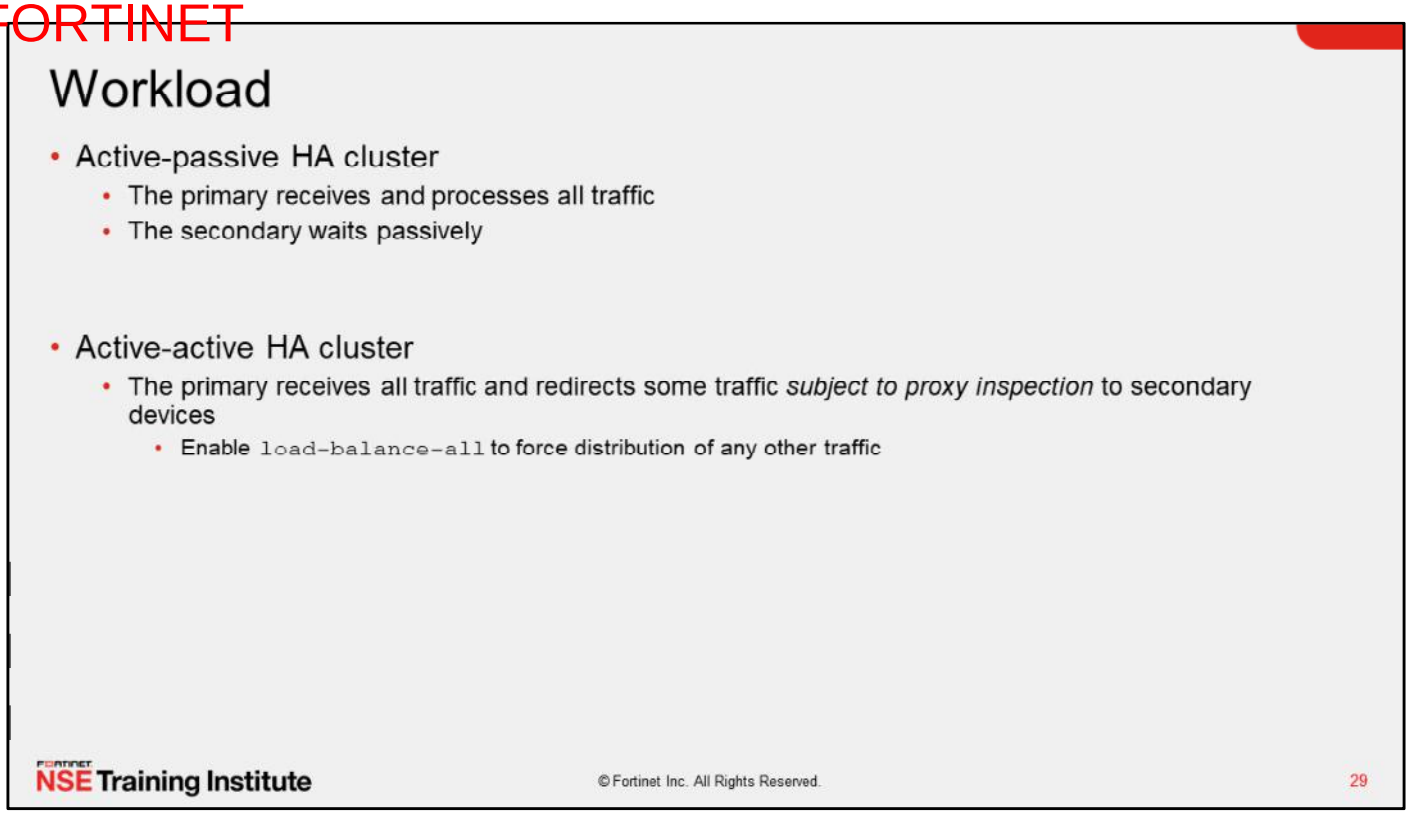

This is how the workload is distributed between roles, depending on the HA mode.

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 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—the setting is disabled by default.

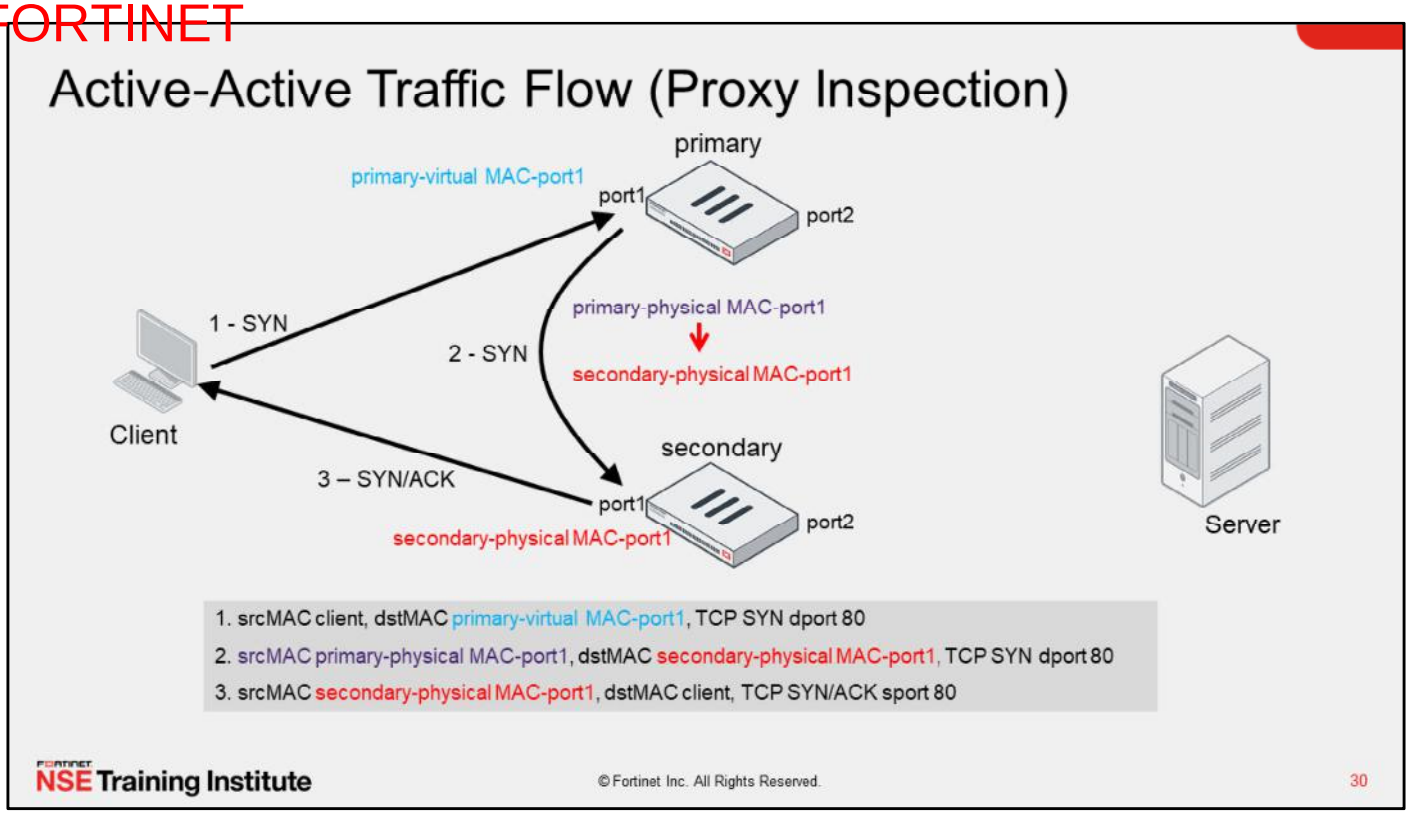

In active-active mode, the following always occur:

- 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, then the traffic is always 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 always sent by the secondary. The traffic is sourced from the physical MAC address of the egressing port on the secondary.

Now look at how an HA cluster in active-active mode distributes traffic when traffic is subject to proxy inspection. First, the client side sends a SYN packet, which is forwarded to port1 on the primary. The packet destination MAC address is the virtual MAC address on port1.

If the primary decides to distribute the session to a secondary FortiGate, the primary forwards the SYN packet to the selected secondary. In this example, the source MAC address of the forwarded 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.

The secondary responds with a SYN/ACK to the client, for which the source MAC address is the physical MAC address of port1 on the secondary, and the destination MAC address is the MAC address of the client.

Note that FortiGate has not contacted the server yet. As shown on the next slides, FortiGate contacts the server only after it finishes the 3-way handshake to the client. The same behavior—server is contacted only after the 3-way handshake to client is completed—is seen when FortiGate operates in standalone mode and performs proxy-based inspection. You will learn more about proxy-based inspection in another lesson.

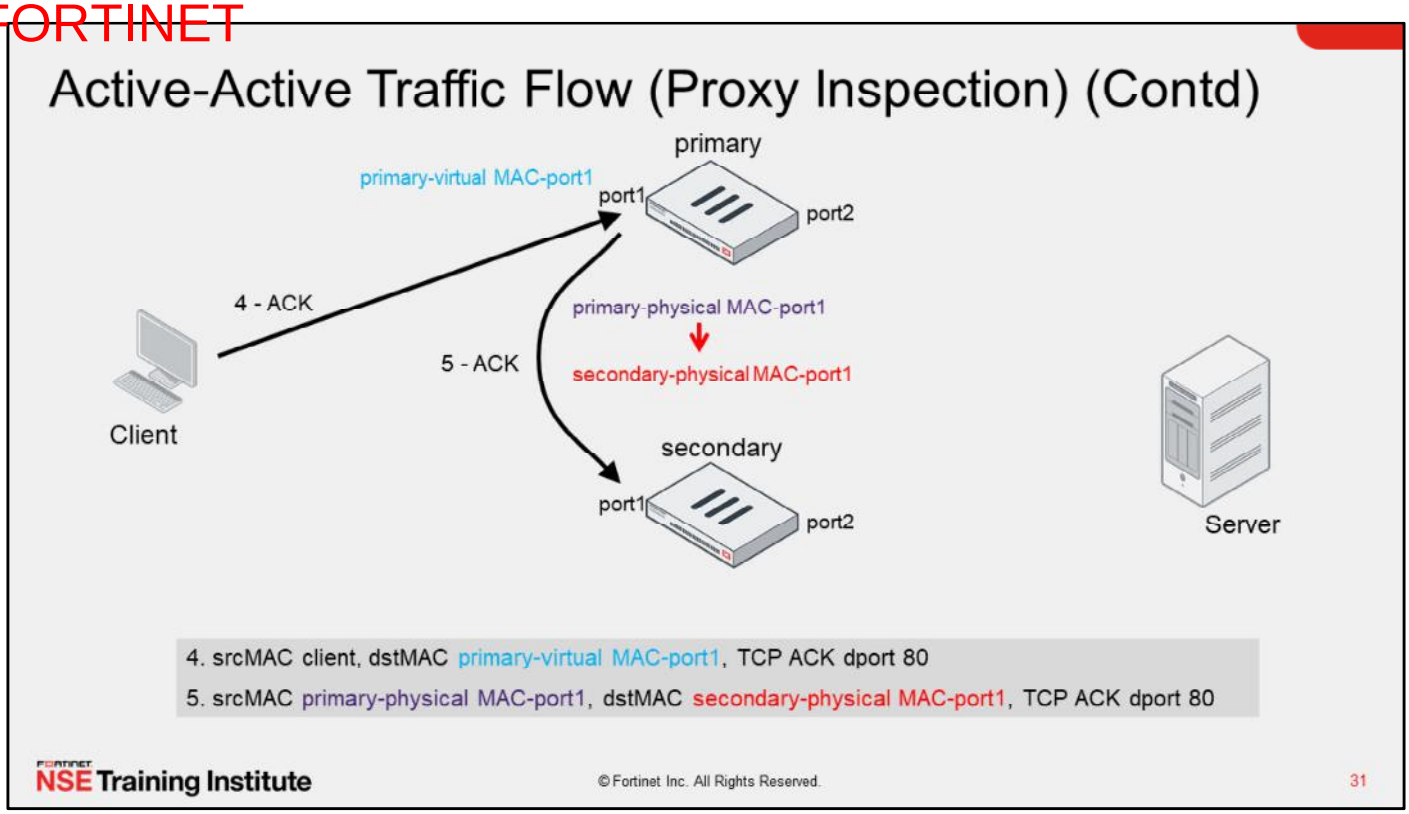

Next, the client acknowledges the SYN/ACK by sending an ACK to the cluster. The ACK packet is destined to port1 on the primary.

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 3-way handshake on the client side is complete.

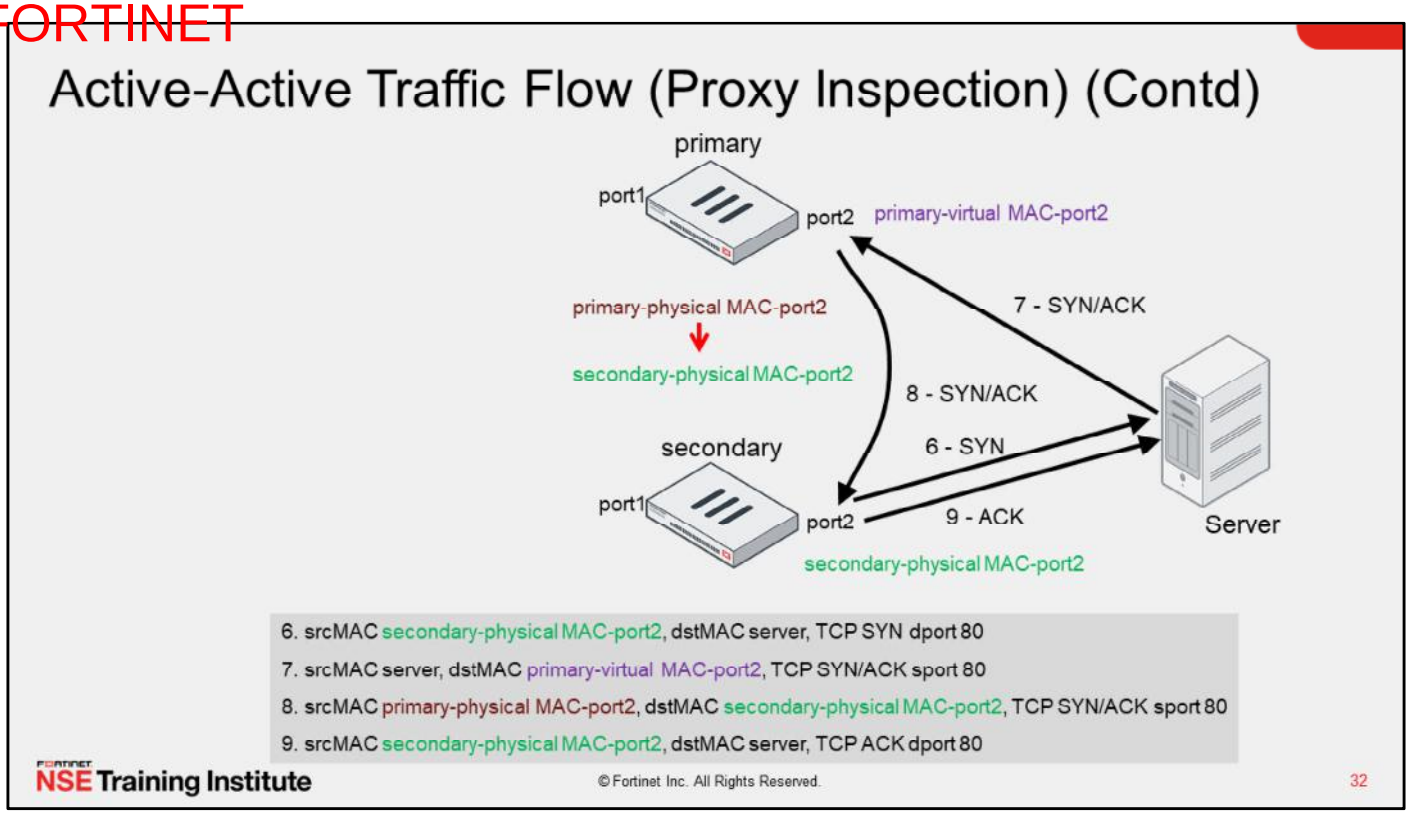

The secondary starts the connection with the server by directly sending a SYN packet using the physical MAC address of port2 as source.

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.

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 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.

The secondary responds with an ACK to the server, for which the source MAC address is the physical MAC address of port2 on the secondary, and the destination MAC address is the MAC address of the server.

The 3-way handshake on the server side is also complete. From now on, packets sent by the client will follow the same flow. For example, an HTTP GET request packet sent by the client will first be received by the primary, which will then forward it to the secondary. The secondary will perform proxy-based inspection on the packet. If the packet is allowed, then the secondary forwards the packet to the server. Any server response packets to the HTTP GET request are sent to the primary, which then forwards them 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 traffic from endpoints is always sent to the primary, expect to see more traffic on the primary than on any secondary FortiGate.

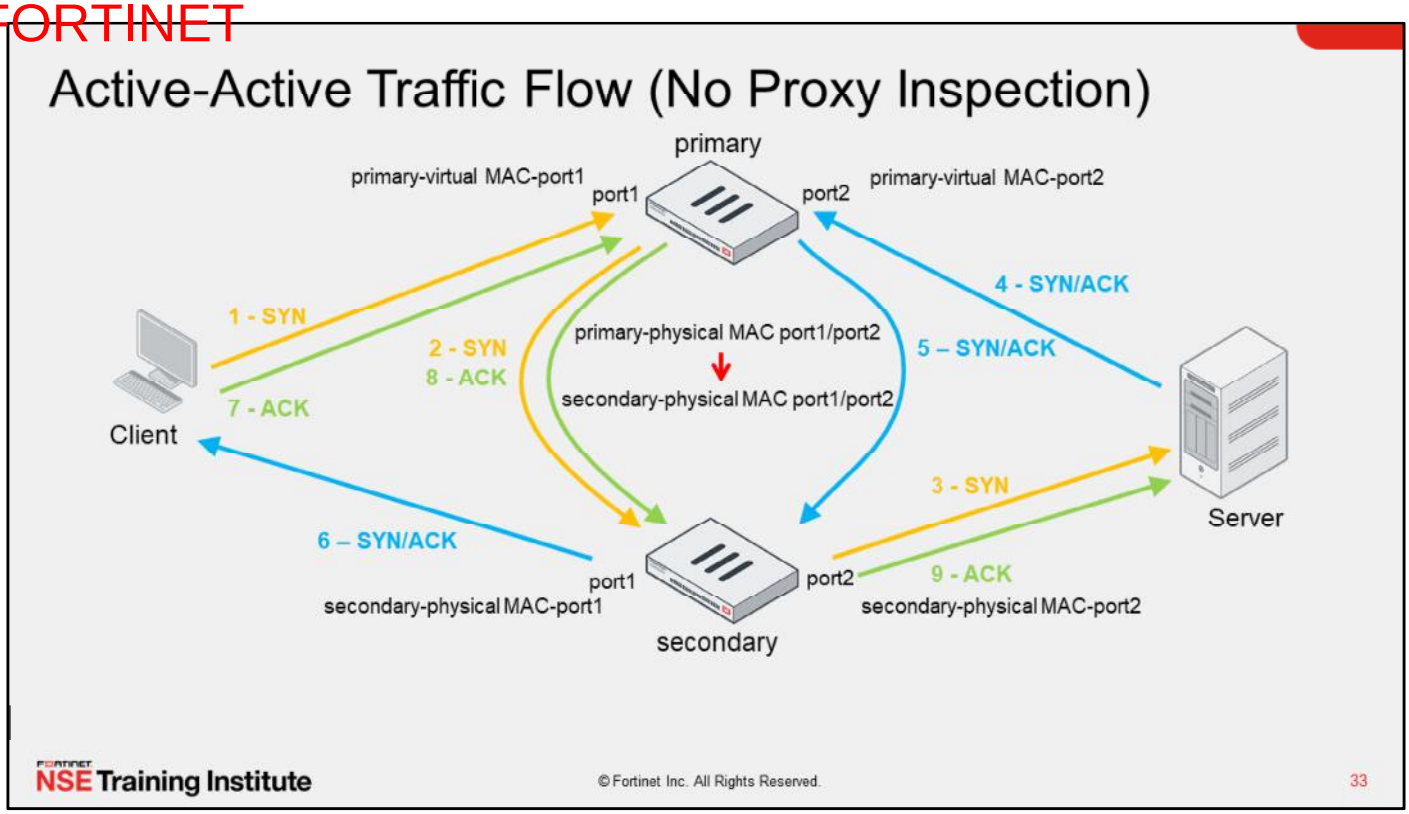

When there is no proxy inspection, that is, when traffic is either subject to flow inspection or no inspection at all, sessions are distributed to the secondary FortiGate only if you enable the `load-balance-all` setting (which is disabled by default) under HA configuration. In addition, as in proxy inspection, you will see the same basic behavior below:

1. Traffic sourced from the client or server and destined to the FortiGate cluster is always 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.
2. The primary FortiGate may, in turn, forward the traffic to the secondary if the session is to be load balanced.
3. When distributing the traffic to the secondary, FortiGate uses as source and destination MAC addresses the physical MAC addresses of the primary and secondary interfaces, respectively. This is also known as MAC address re-write.
4. If traffic has been load balanced to the secondary FortiGate, any traffic sourced from the cluster and destined to the endpoint is always sourced from the secondary FortiGate device. 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 the server. For example, instead of replying to the SYN packet sent by the client, FortiGate forwards the packet to the server through the secondary. Similarly, FortiGate forwards packets sent by the server to the client through the secondary.

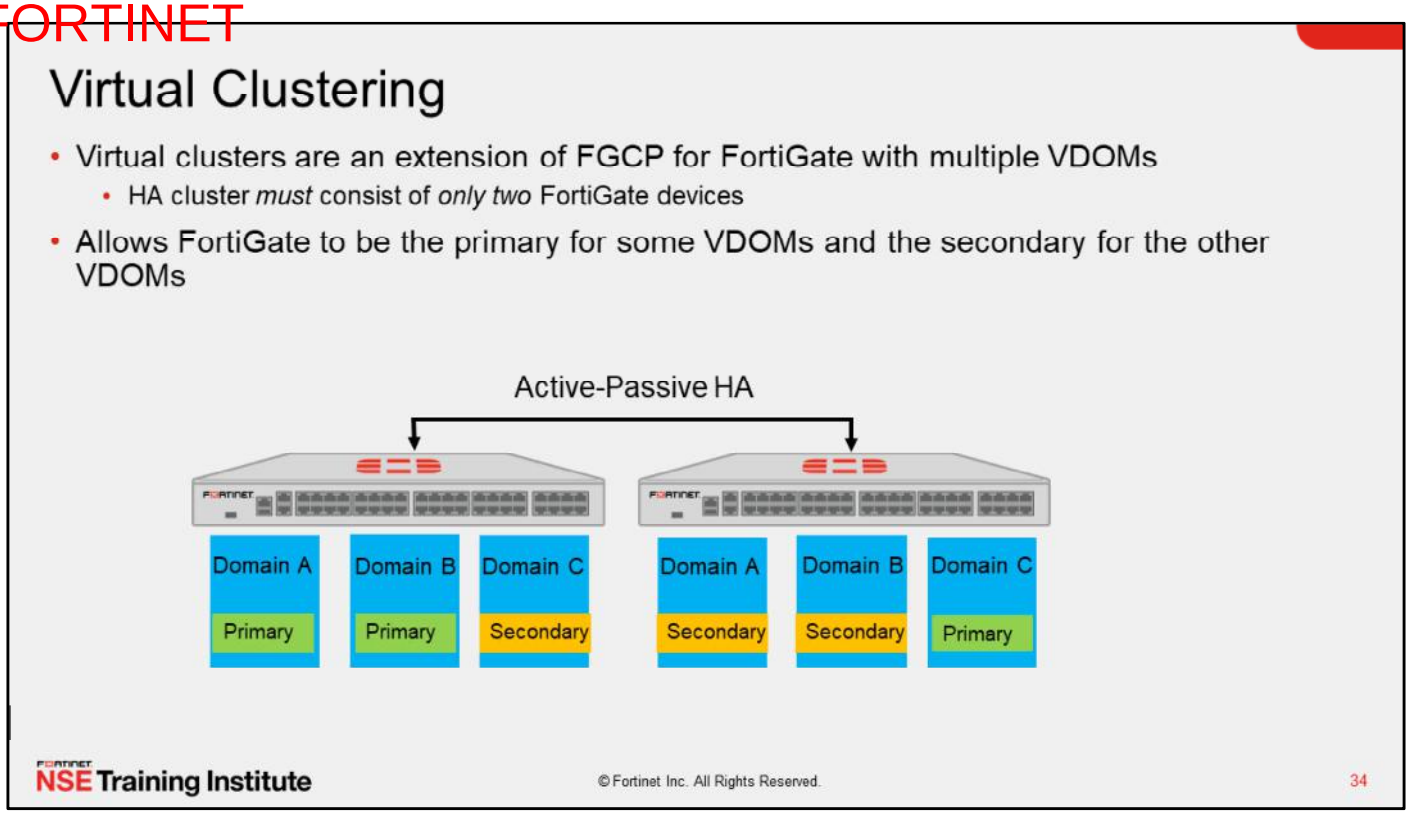

So far, you've learned about HA clustering where each FortiGate device acts as a whole security domain.

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.

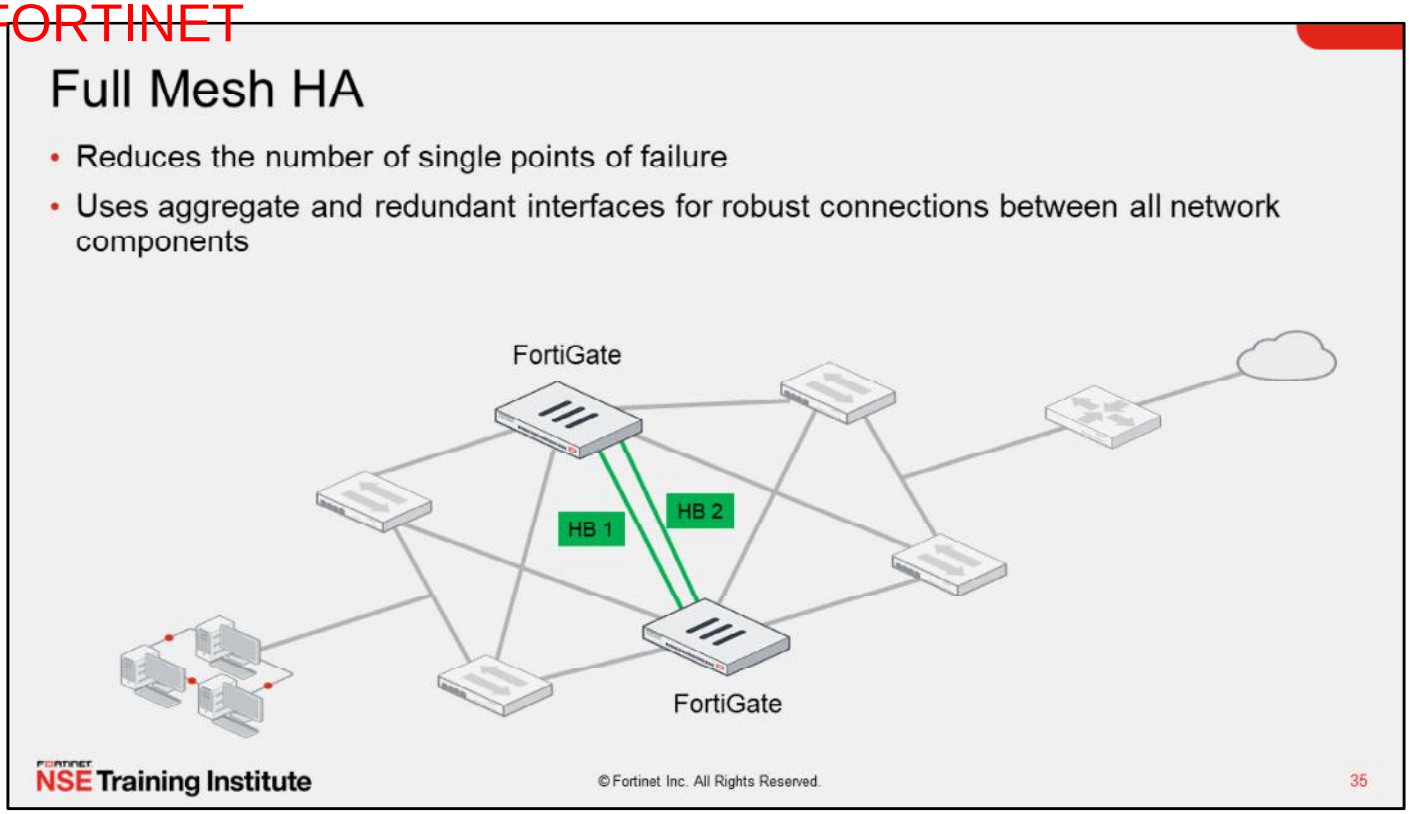

At the beginning of this lesson, you reviewed a simple HA topology. Now, take a look at a more robust topology. It is called *full mesh HA*.

The idea is to prevent any single point of failure, not only on the FortiGate devices, but also on the network switches and interfaces.

As you can see on this slide, you have two FortiGate devices for redundancy, and each FortiGate is connected to two redundant switches, using two different interfaces.

A full mesh HA is more complicated to assemble and administer, but it can provide the availability required by critical installations.

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. You can configure virtual clustering between only ____ FortiGate devices with multiple VDOMs in an active-passive HA cluster.
   - ✓ A. Two
   - B. Four

NSE Training Institute

36

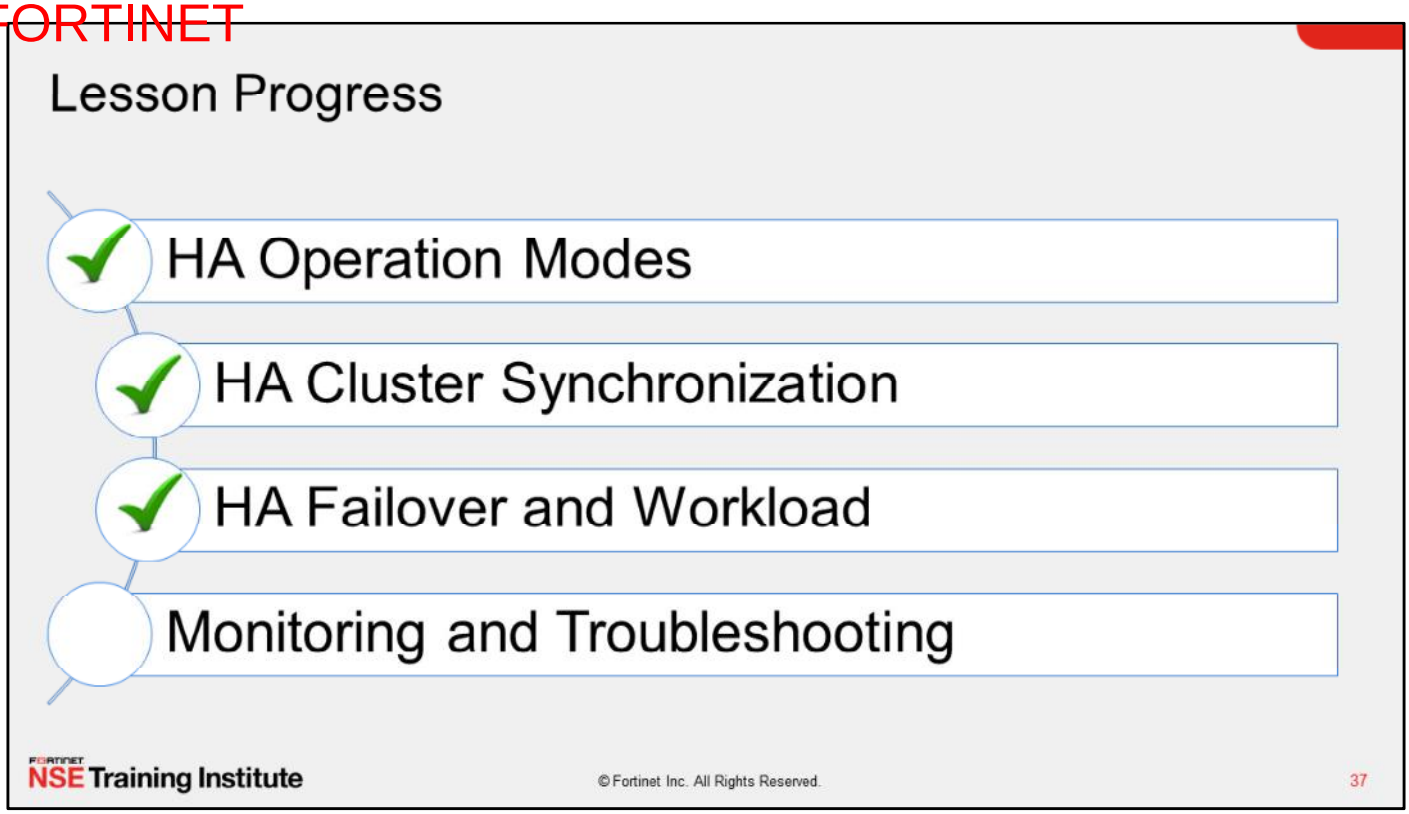

Good job! You now understand HA failover and workload.

Now, you will learn about monitoring and troubleshooting an HA cluster.

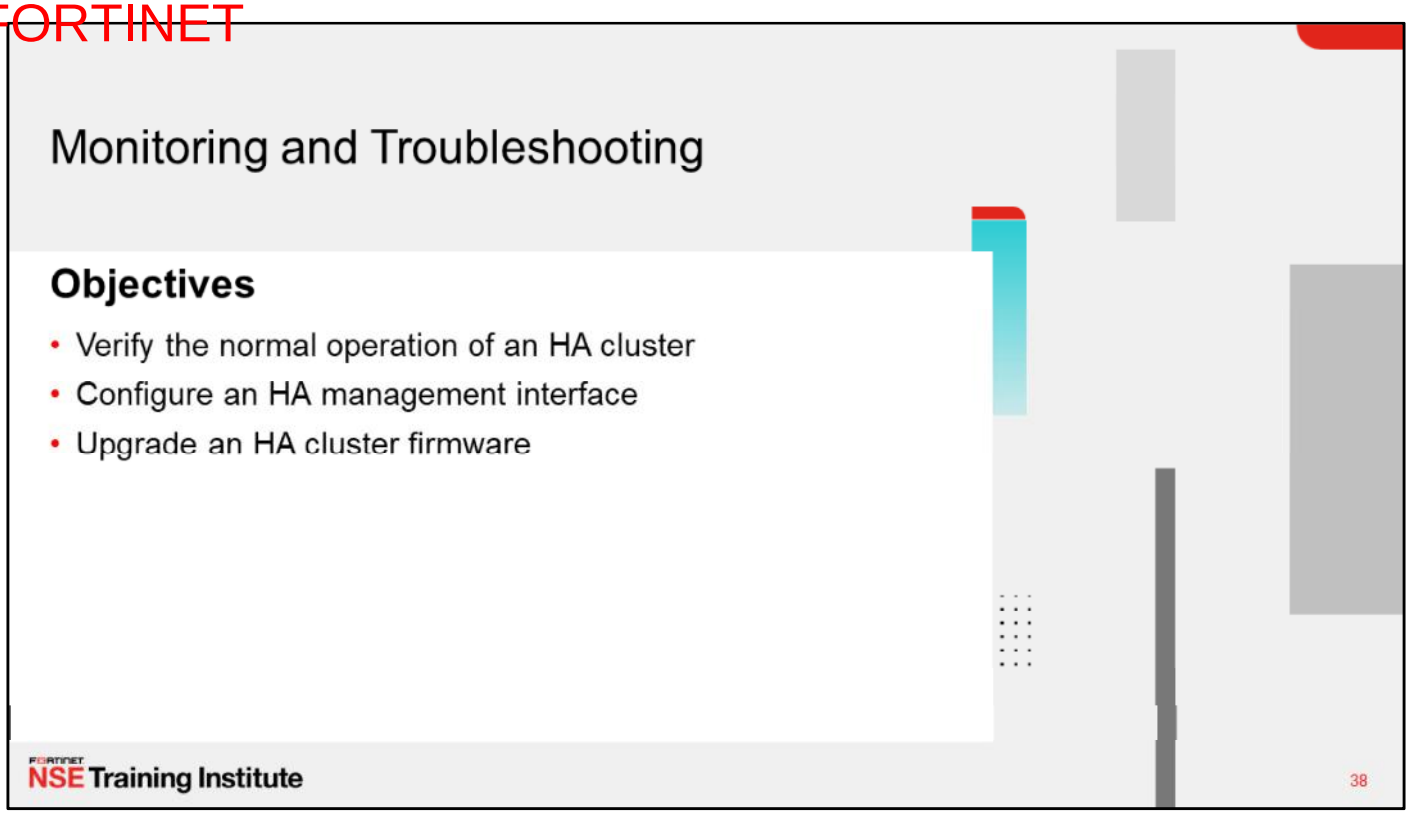

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

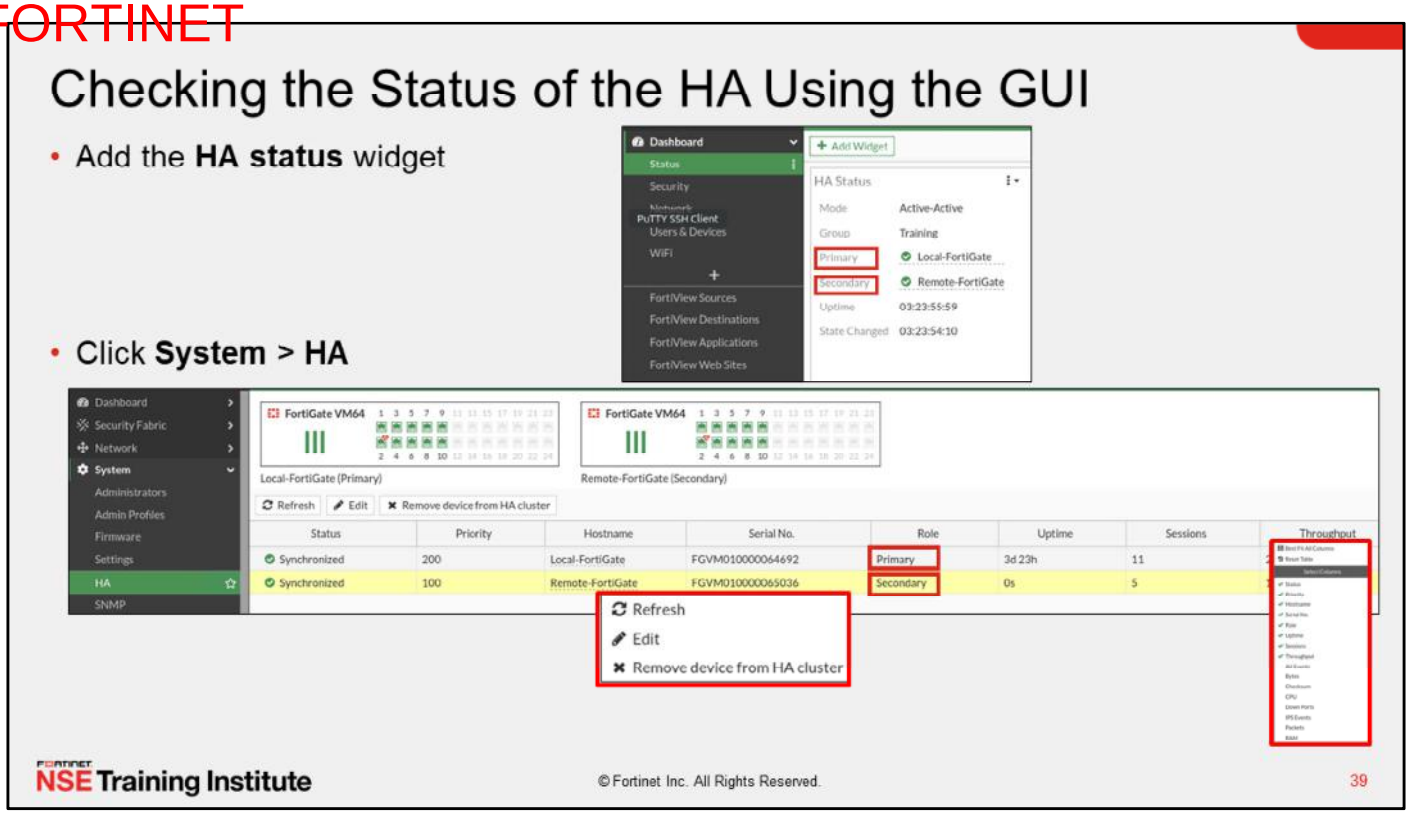

If the HA cluster has formed successfully, the GUI displays all the FortiGate devices in the cluster. It shows the synchronized status, hostnames, serial numbers, roles, priorities, uptime, and active sessions of the cluster members.

You can view the HA status by adding the **HA Status** dashboard widget. You can also view HA status information by clicking **System** > **HA**, where you can see more details by adding columns such as checksum, CPU, and memory, to name a few.

You can also disconnect a cluster member from the cluster and edit the HA configurations.

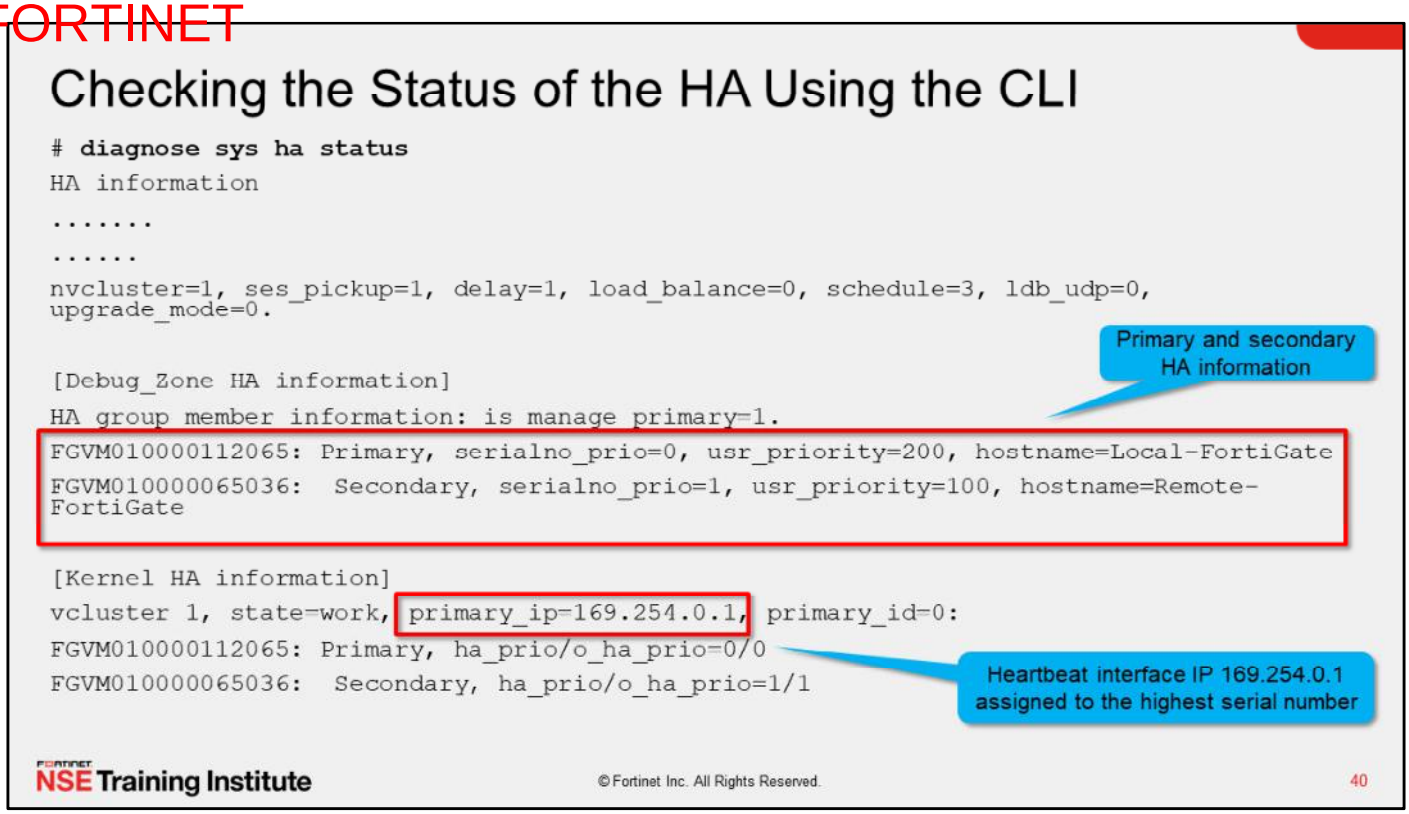

You can get more information about the status of the HA from the CLI. For example, the command `diagnose sys ha status` displays heartbeat traffic statistics, as well as the serial number and HA priority of each FortiGate device. This command also shows the heartbeat interface IP address automatically assigned to the FortiGate device with the highest serial number.

Remember, the heartbeat IP address assignment changes only when a FortiGate device leaves or joins the cluster.

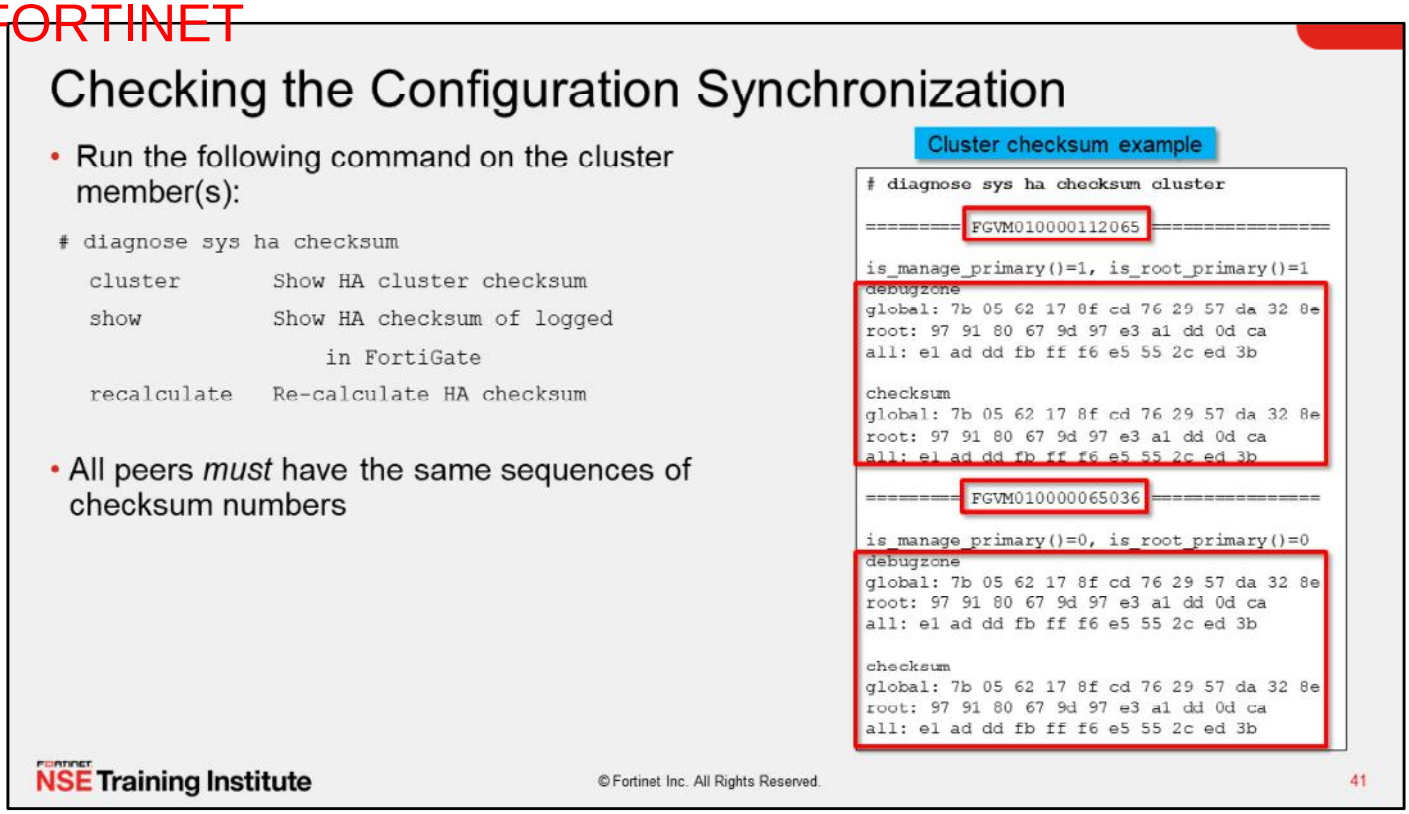

Another indication of the health of an HA cluster is the status of the configuration synchronization. The `diagnose sys ha checksum` command tree provides many options that you can use to check or recalculate the HA checksum.

To check that all the secondary configurations are synchronized with the primary configuration:
- Run the `diagnose sys ha checksum cluster` command to view the checksums of all cluster members from any FortiGate device in a cluster.
- The `diagnose sys ha checksum show` command shows the checksum of the individual FortiGate from which this command is run.
- You can also run the `diagnose sys ha checksum recalculate` command from any cluster member to recalculate the HA checksums.

If a secondary FortiGate displays exactly the same sequence of numbers as the primary, its configuration is well synchronized with the primary FortiGate in the cluster. In the example shown on this slide, the `diagnose sys ha checksum cluster` command is run to view the checksums of all cluster members.
- `global` represents the checksum of the global configuration, such as administrators, administrator profiles, global logging settings, and FortiGuard settings, and so on.
- `root` is the checksum for the root VDOM. If you have configured multiple VDOMs, you will see checksums of all configured VDOMs.
- `all` is the checksum of the global configuration, plus the checksums of all the VDOMs.

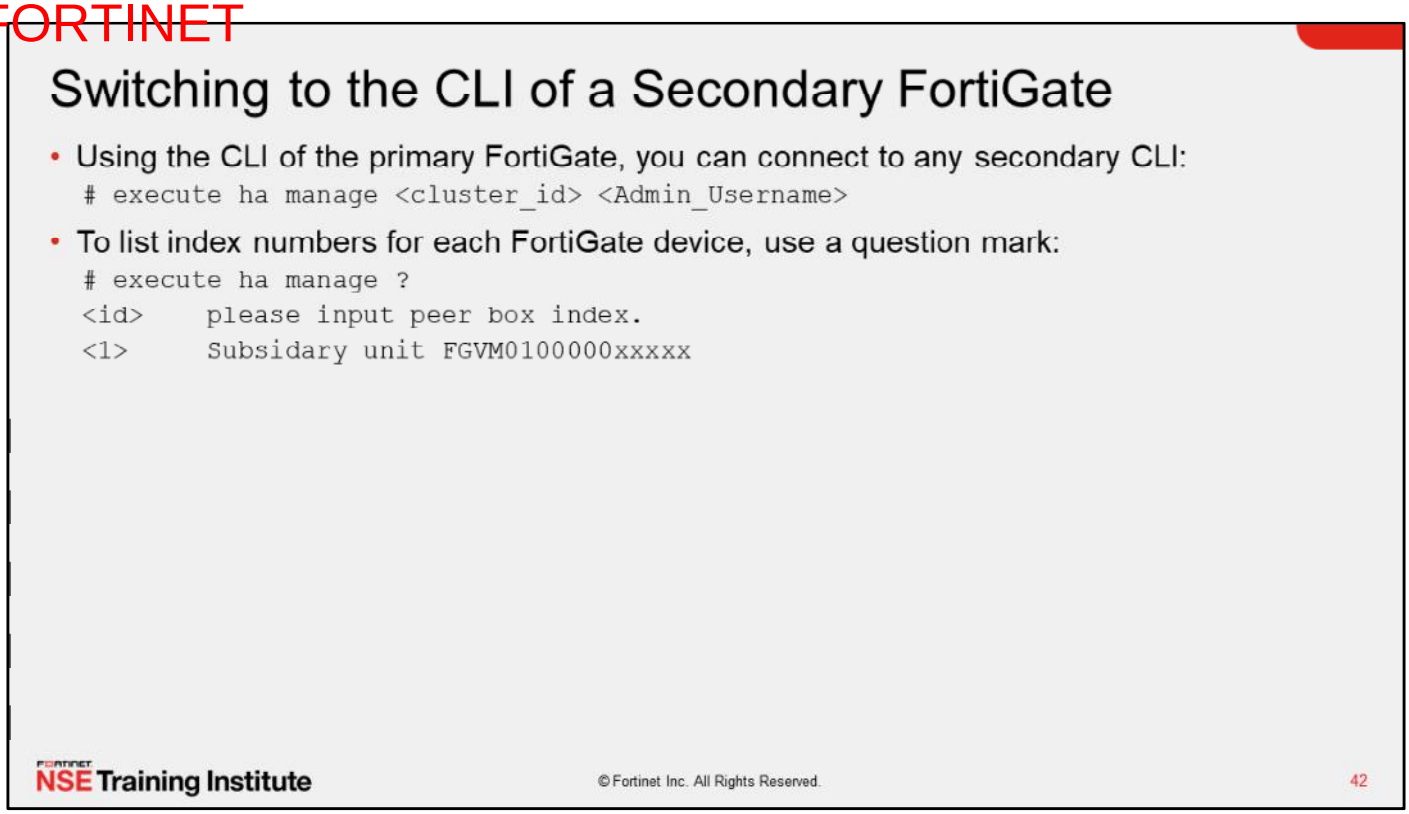

When troubleshooting a problem in an HA cluster, it is useful to know that you can connect to the CLI of any secondary FortiGate device from the CLI of the primary FortiGate device. You have to use the command `execute ha manage` with the secondary HA index for that purpose.

To get the list of secondary FortiGate devices with their HA indexes, you can add a question mark to the end of the `execute ha manage` command: `execute ha manage ?`.

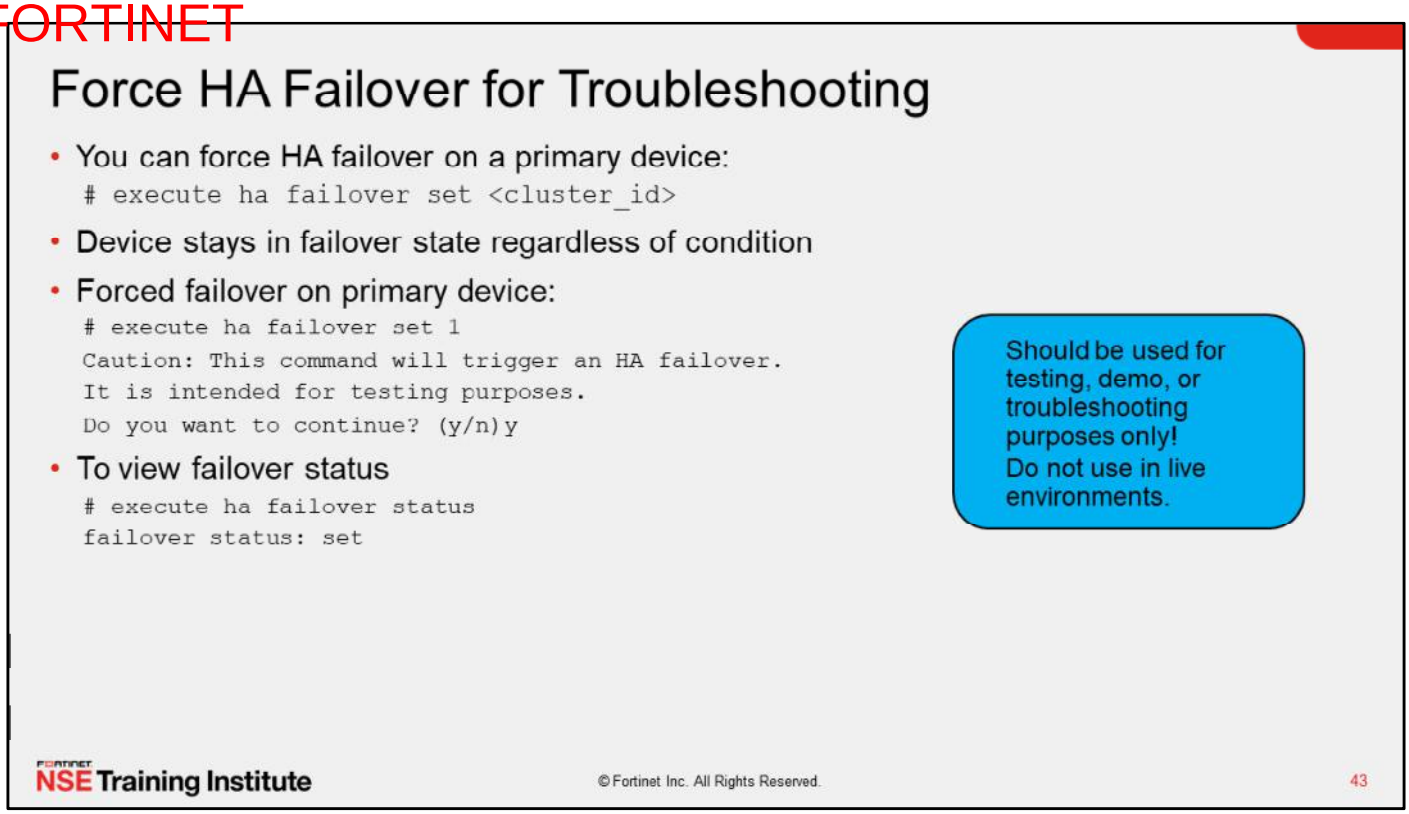

You can force HA failover on an HA primary device. The device stays in a failover state regardless of the conditions. The only way to remove the failover status is by manually turning it off.

Note that you should use this command only for testing, troubleshooting, and demonstrations. Do not use it in a production environment.

Force HA Failover for Troubleshooting (Contd)

- To view the system status of a device in forced HA failover:

```
# get system ha status
HA Health Status: OK
.........
Primary selected using:
    <2020/04/19 10:16:54> FGVM010000064692 is selected as the primary because it has EXE_FAIL_OVER
flag set.
    <2020/04/19 10:07:29> FGVM010000065036 is selected as the primary because it has the largest value
of override priority.
```

Forced failover was used to select primary

- To stop the failover status:

```
# execute ha failover unset 1
```

- To view the system status of a device after forced HA failover is disabled:

```
# get system ha status
.........
Primary selected using:
    <2020/04/19 10:38:28> FGVM010000065036 is selected as the primary because it has the
largest value of override priority.
    <2020/04/19 10:16:54> FGVM010000064692 is selected as the primary because it has
EXE_FAIL_OVER flag set.
```

Primary is selected based on device with highest priority

NSE Training Institute

44

Use the `get system ha status` command to confirm how the current primary was selected.

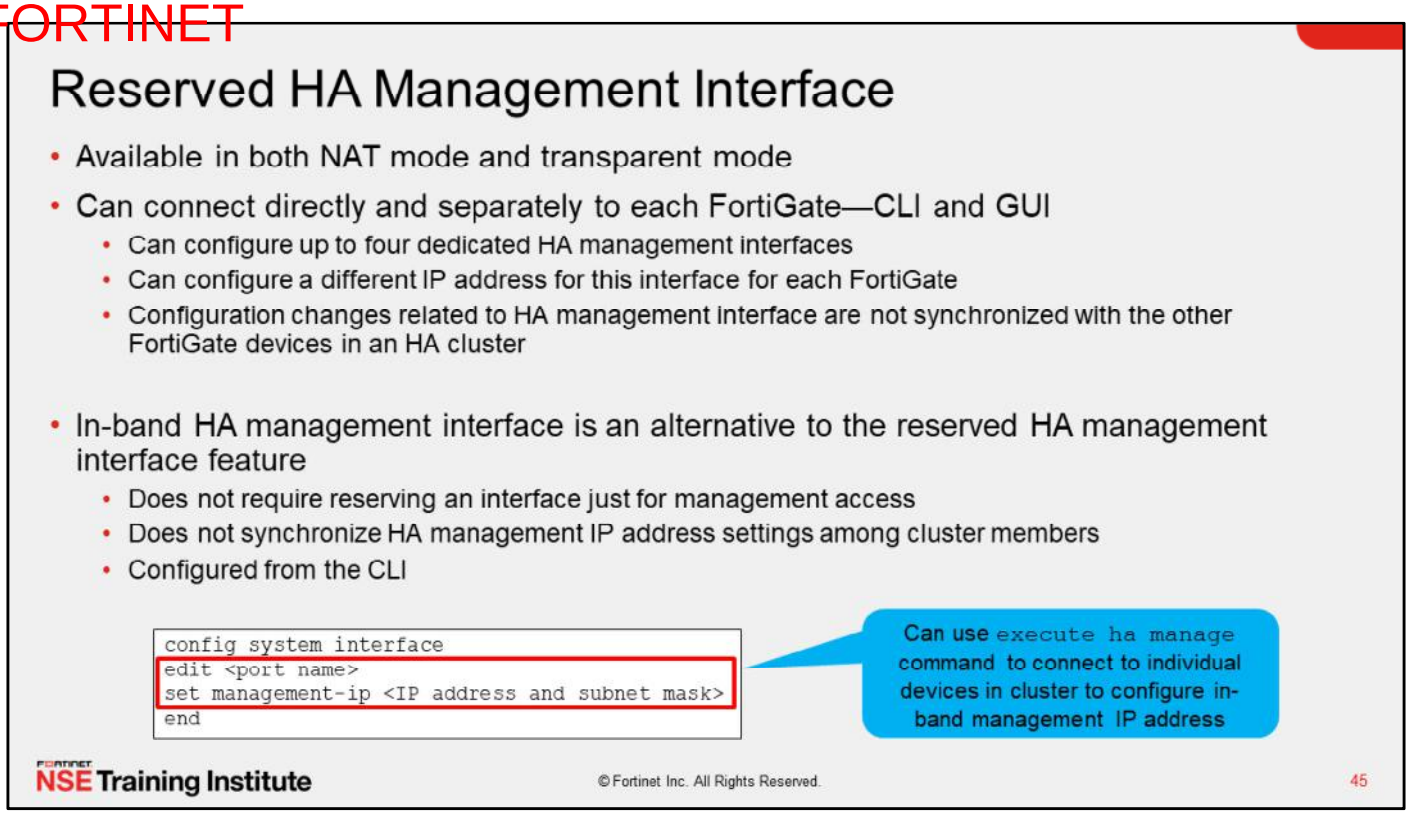

If you want to be able to connect to each device directly, you can reserve an interface for HA management. The FGCP cluster supports reserved HA management interfaces in both NAT and transparent mode. You can configure up to four dedicated management interfaces. The configuration of a reserved HA management interface is not synchronized between HA cluster members, and each device can have different management IP addresses. Each device can also use the HA reserved management interface to send SNMP traffic and logs independently.

What if you don't have a free interface available to reserve it as a dedicated HA management interface? You can configure a management IP address from the CLI on any interface that is connected to a network and processing traffic. This doesn't require reserving an interface just for management access, and is an alternative to reserving a dedicated HA management interface.

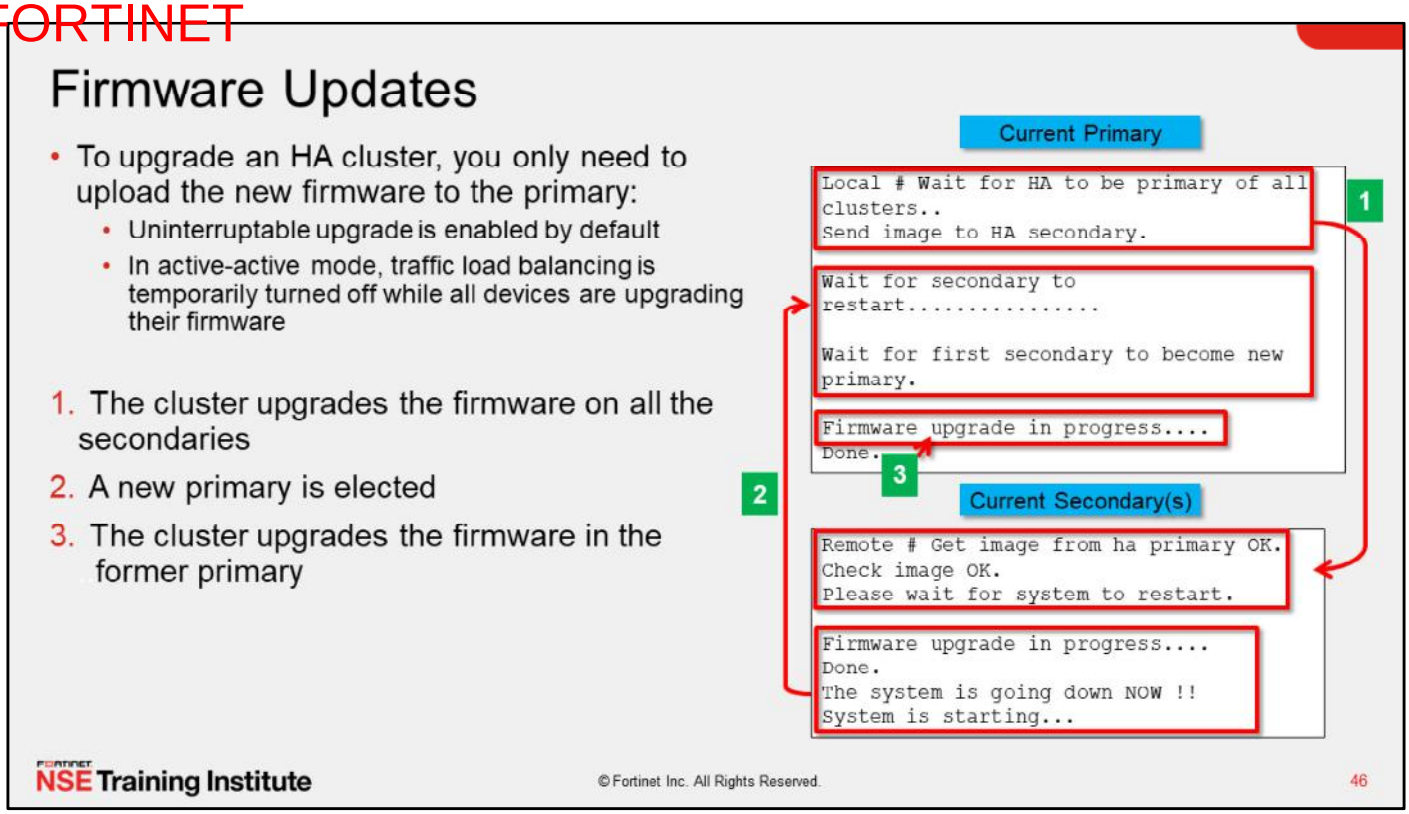

As with a standalone device, when upgrading an HA cluster, each updating FortiGate device must reboot. As the uninterruptable upgrade is enabled by default, the cluster upgrades the secondary FortiGate devices first. Once all the secondary FortiGate devices are running the new firmware, a new primary is elected and the firmware in the original primary device is upgraded.

If the cluster is operating in active-active mode, traffic load balancing is temporally disabled while all devices are upgrading their firmware.

You can change the firmware upgrade process by disabling the uninterruptable upgrade on the CLI using `config system ha`. This results in all FortiGate devices in a cluster being upgraded at the same time. This takes less time, but interrupts the traffic flow.

Knowledge Check

1. The heartbeat interface IP address `169.254.0.1` is assigned to which FortiGate in an HA cluster?
   - ✓ A. The FortiGate with the highest serial number
   - B. The FortiGate with the highest priority

2. Which statement about the firmware upgrade process on an HA cluster is true?
   - ✓ A. You need to upload the new firmware only to the primary FortiGate to upgrade an HA cluster.
   - B. The cluster members are not rebooted.

NSE Training Institute

47

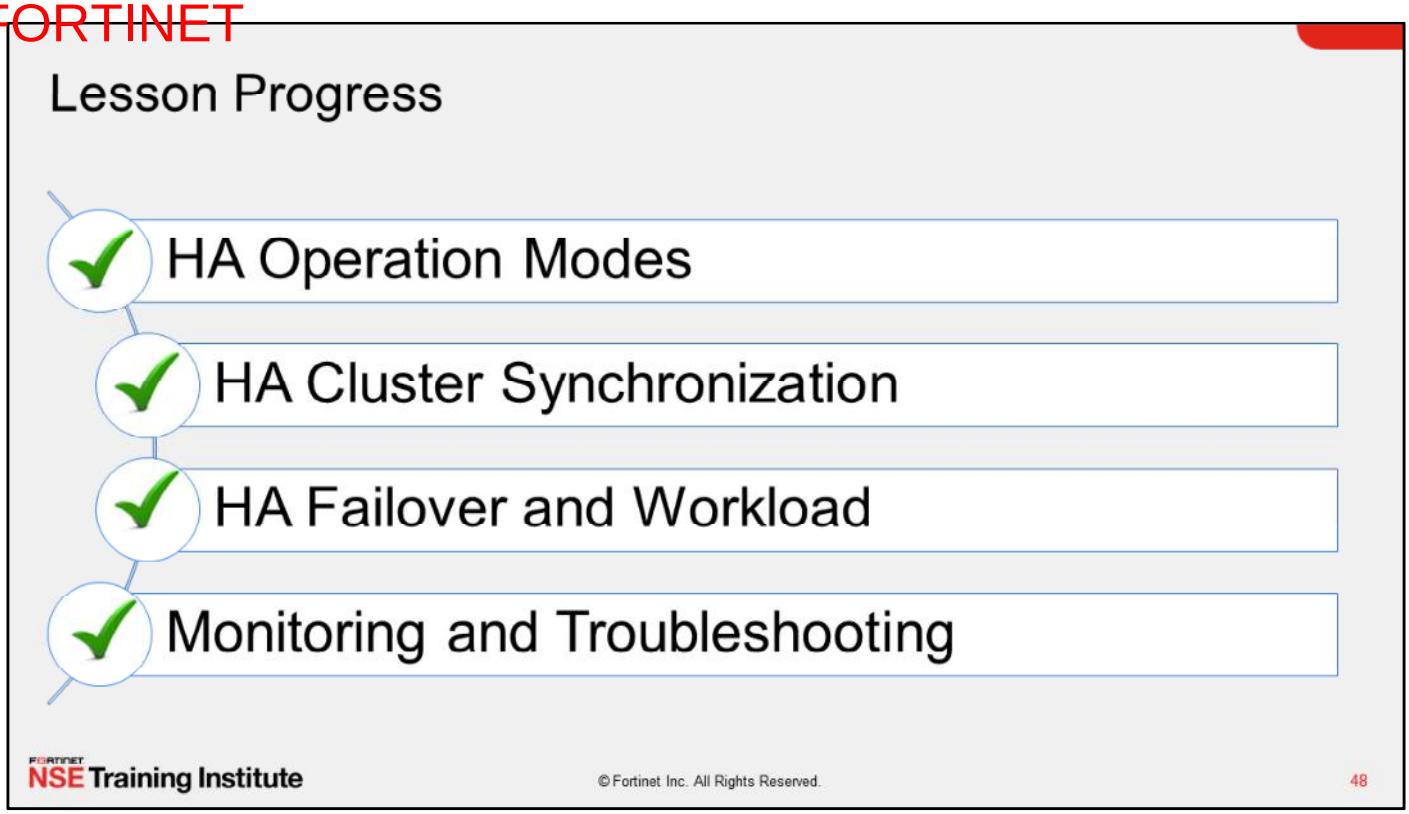

Congratulations! You have completed this lesson.

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

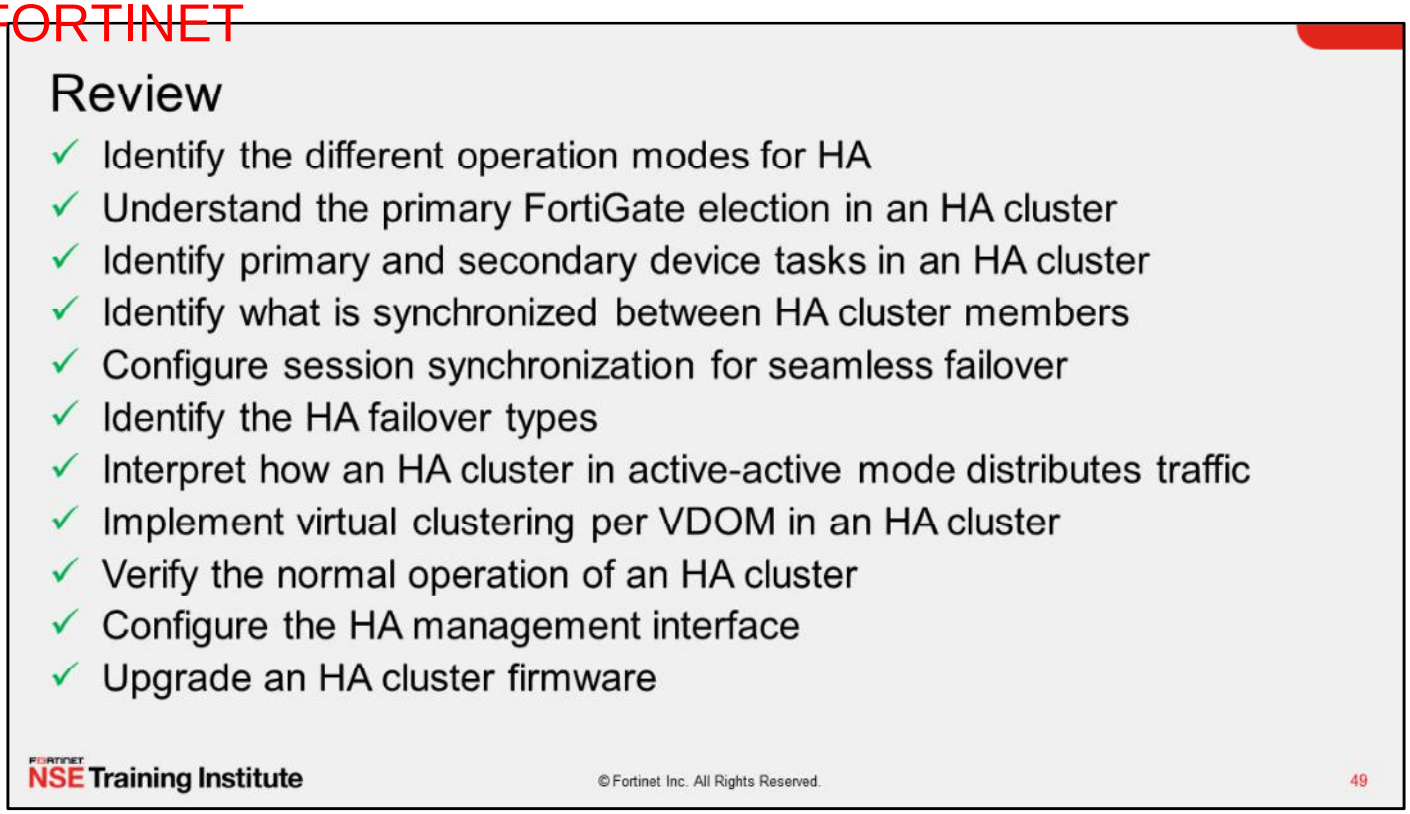

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned about the fundamentals of FortiGate high availability (HA) and how to configure it.

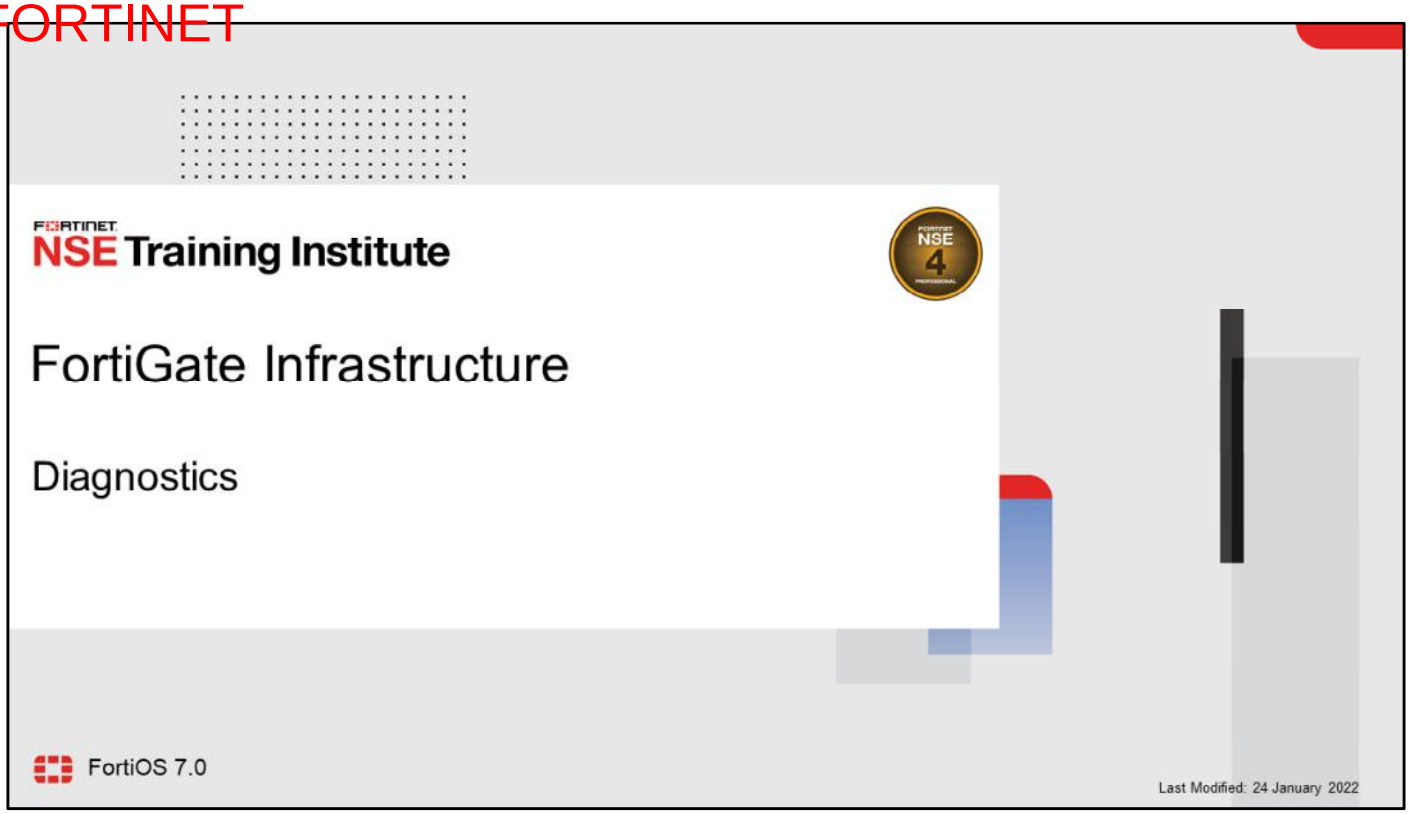

In this lesson, you will learn about using diagnostic commands and tools.

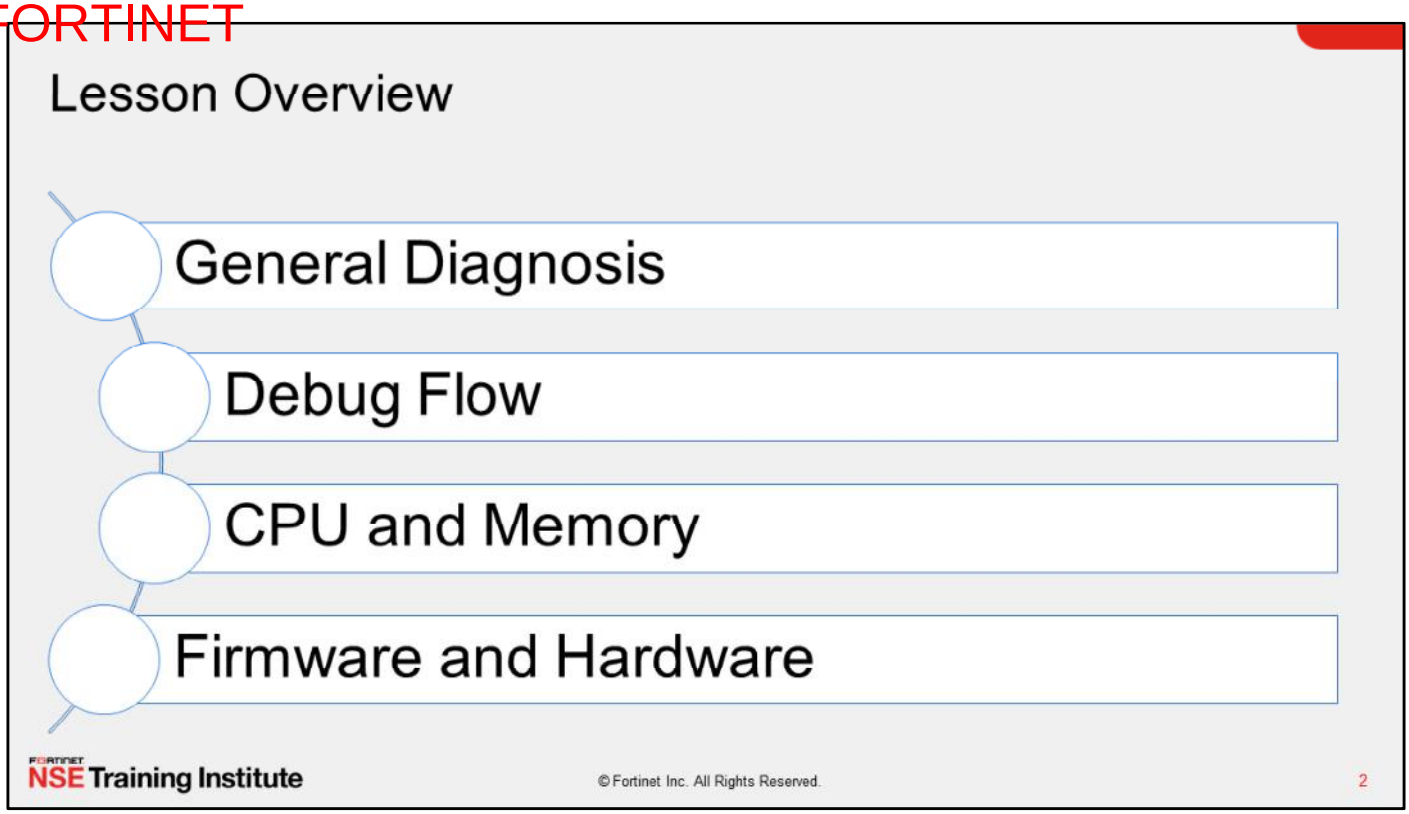

In this lesson, you will learn about the topics shown on this slide.

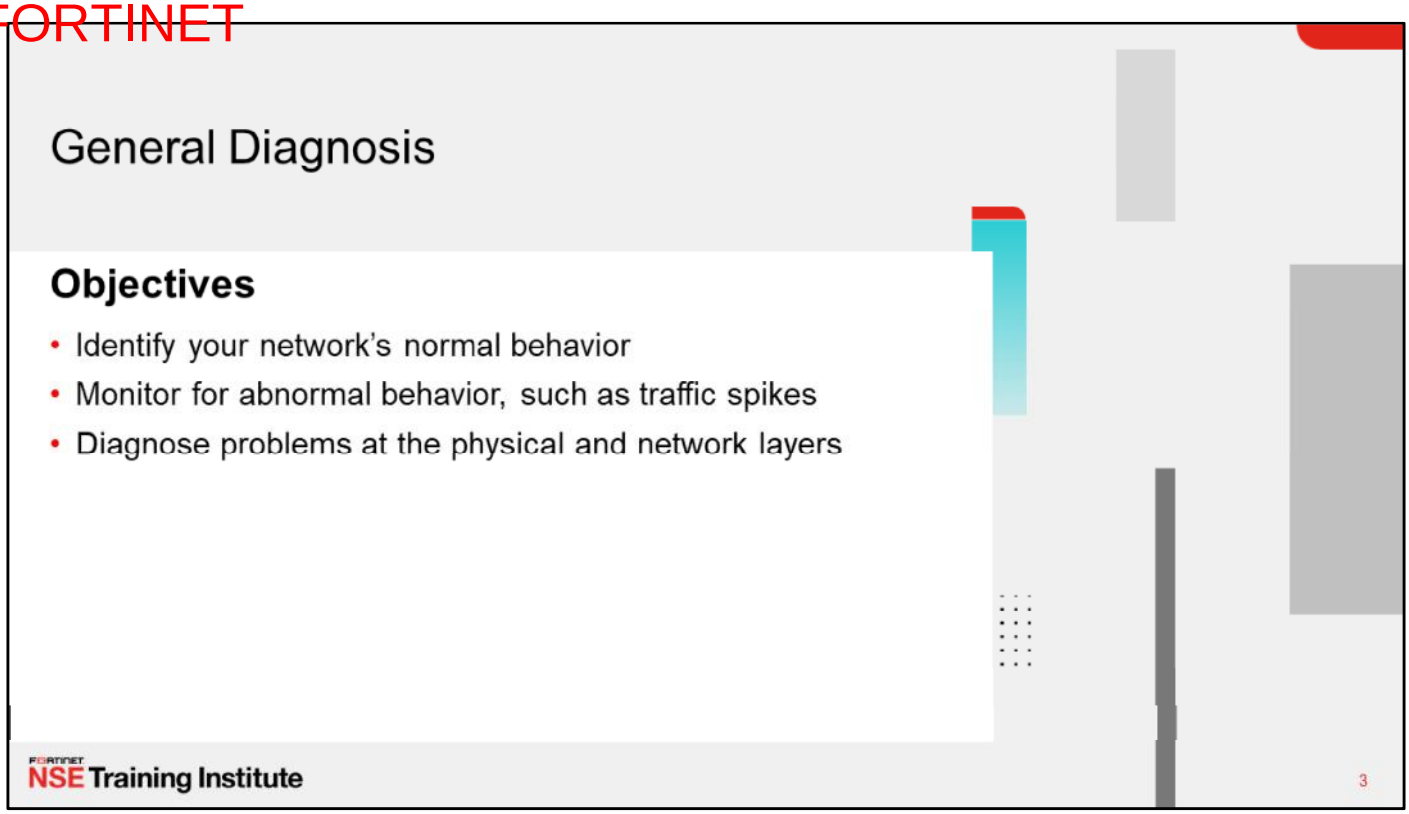

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in general diagnosis, you will be able to discover general information about the status of FortiGate.

Before a Problem Occurs

- Know what normal is (baseline):
  - CPU usage
  - Memory usage
  - Traffic volume
  - Traffic directions
  - Protocols and port numbers
  - Traffic pattern and distribution
- Why?
  - Abnormal behavior is difficult to identify, *unless* you know, relatively, what normal is

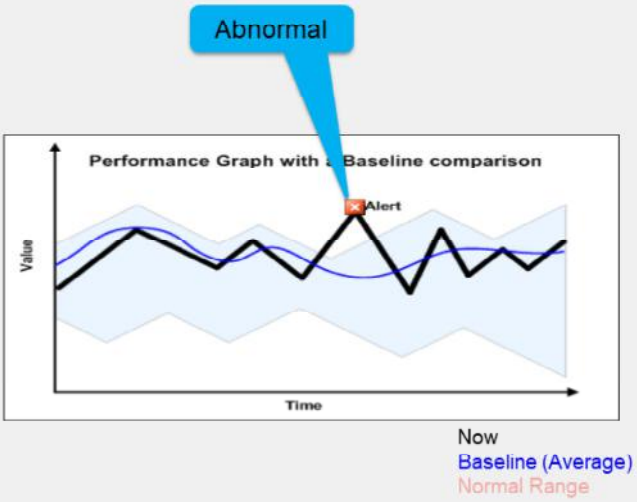

NSE Training Institute

4

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

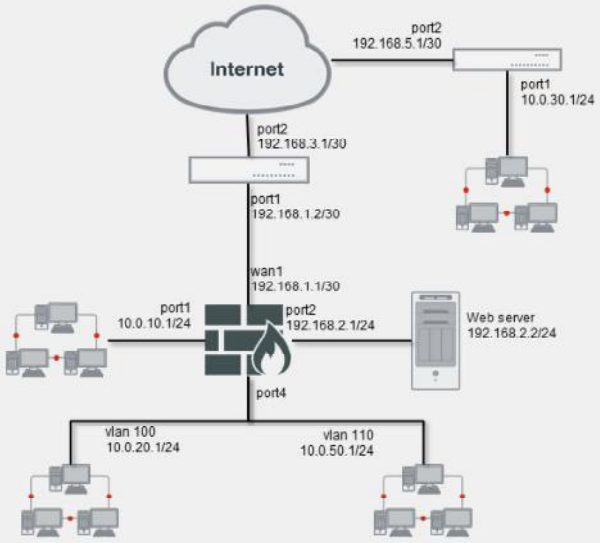

NSE Training Institute

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.

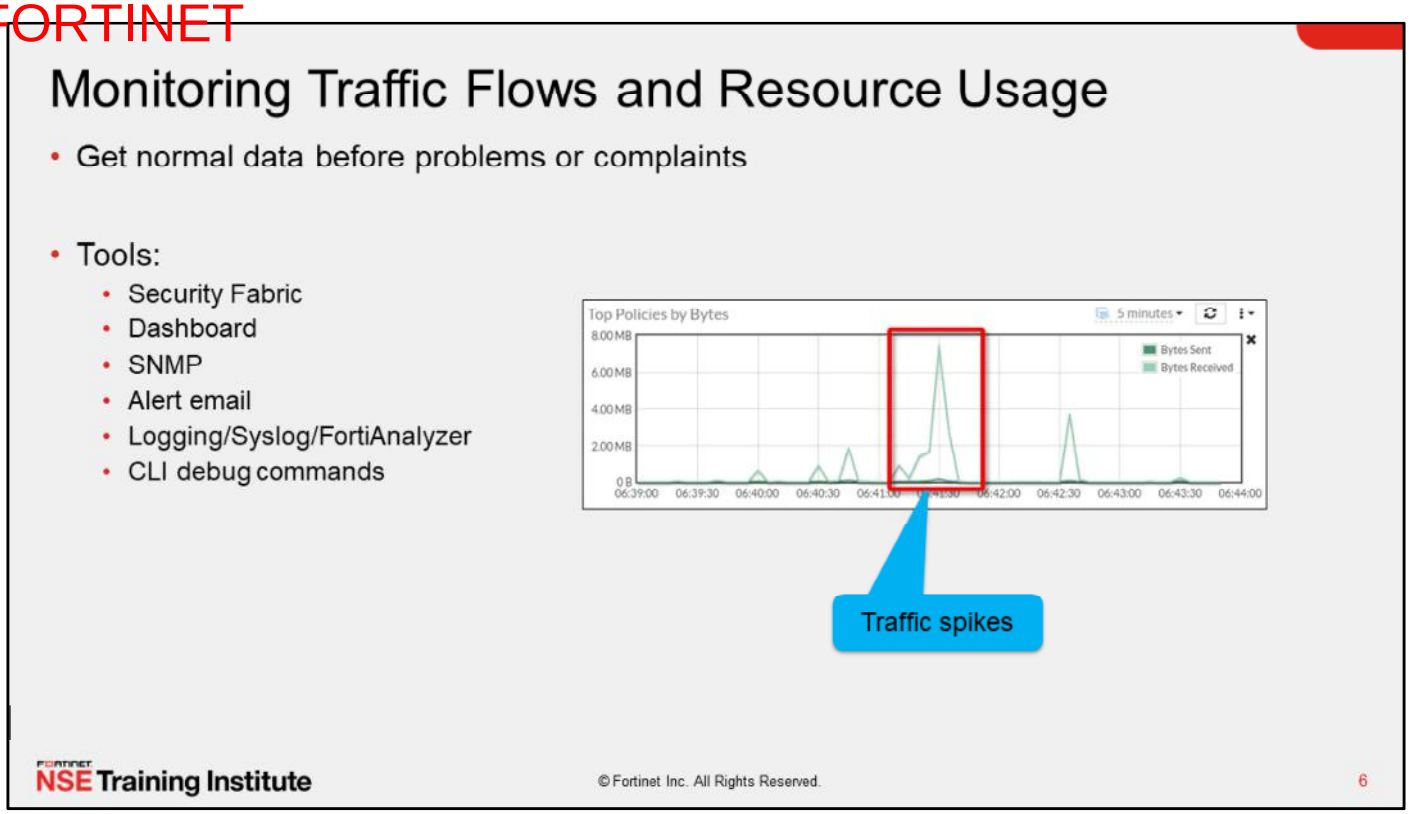

Another way to define normal is to know the average performance range. On an ongoing basis, collect data that shows normal usage.

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?

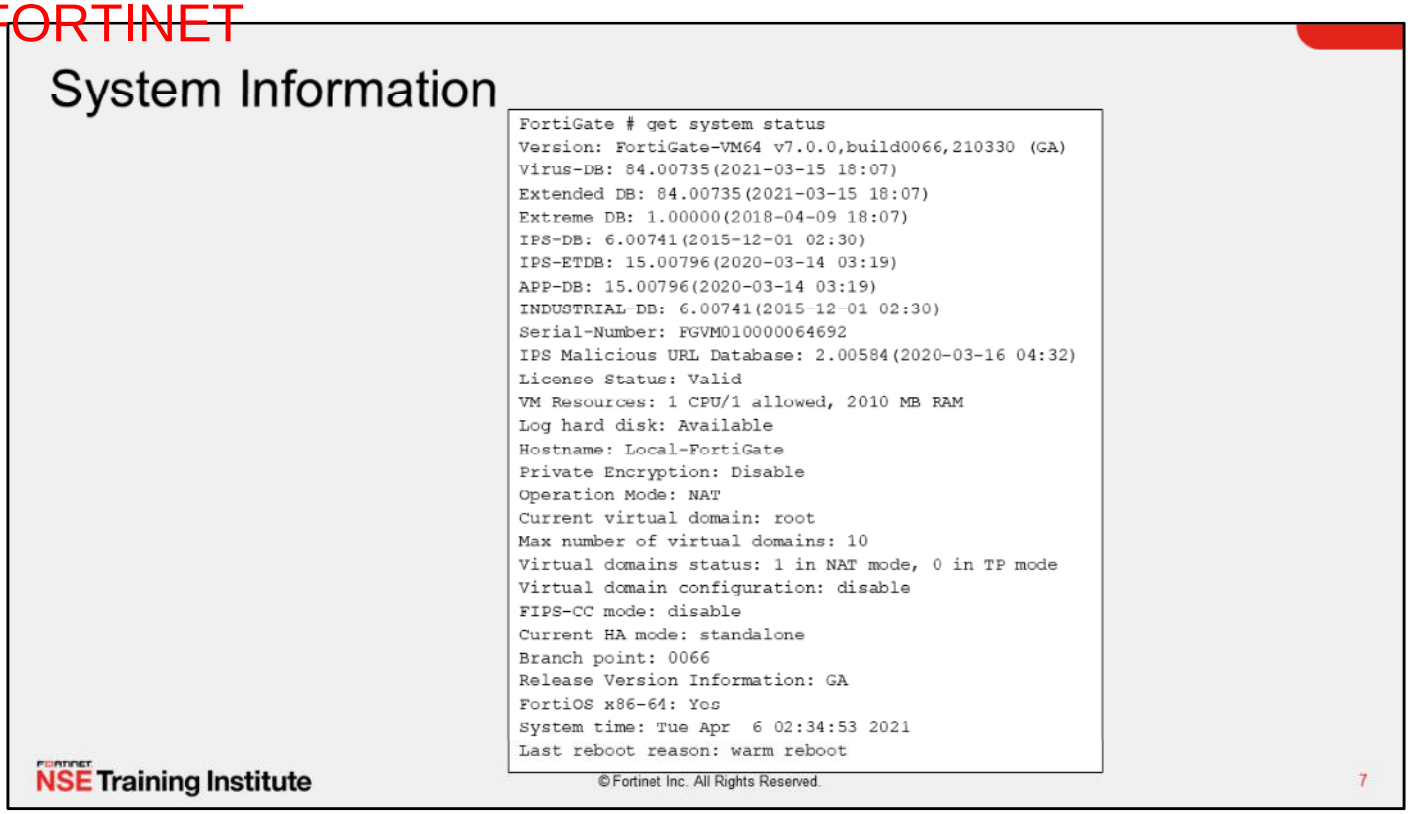

How can we get information about the current status? First, look at CLI commands; you can use them through a local console, even if network issues make GUI access slow or impossible.

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

System Information (Contd)

```
FortiGate # get hardware nic <interface_name>

Name:              port1
Driver:            virtio_net
Version:           1.0.0
Bus:               0000:00:03.0
Hwaddr:            02:09:0f:00:02:01
Permanent Hwaddr:02:09:0f:00:02:01
State:             up
Link:              up
Mtu:               1500
Supported:         1000full 10000full
Advertised:
Speed:             10000full
Auto:              disabled
RX Ring:                   256
TX Ring:                   256
Rx packets:                670785
Rx bytes:                  949908714
Rx compressed:             0
Rx dropped:                0
...
```

```
...
Rx errors:                 0
  Rx Length err:           0
  Rx Buf overflow:         0
  Rx Crc err:              0
  Rx Frame err:            0
  Rx Fifo overrun:         0
  Rx Missed packets:       0
Tx packets:                57752
Tx bytes:                  4993066
Tx compressed:             0
Tx dropped:                0
Tx errors:                 0
  Tx Aborted err:          0
  Tx Carrier err:          0
  Tx Fifo overrun:         0
  Tx Heartbeat err:        0
  Tx Window err:           0
Multicasts:                0
Collisions:                0
```

NSE Training Institute

8

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 output might vary depending on the model and NIC driver version. In all cases, the output shows the physical MAC address, administrative status, and link status.

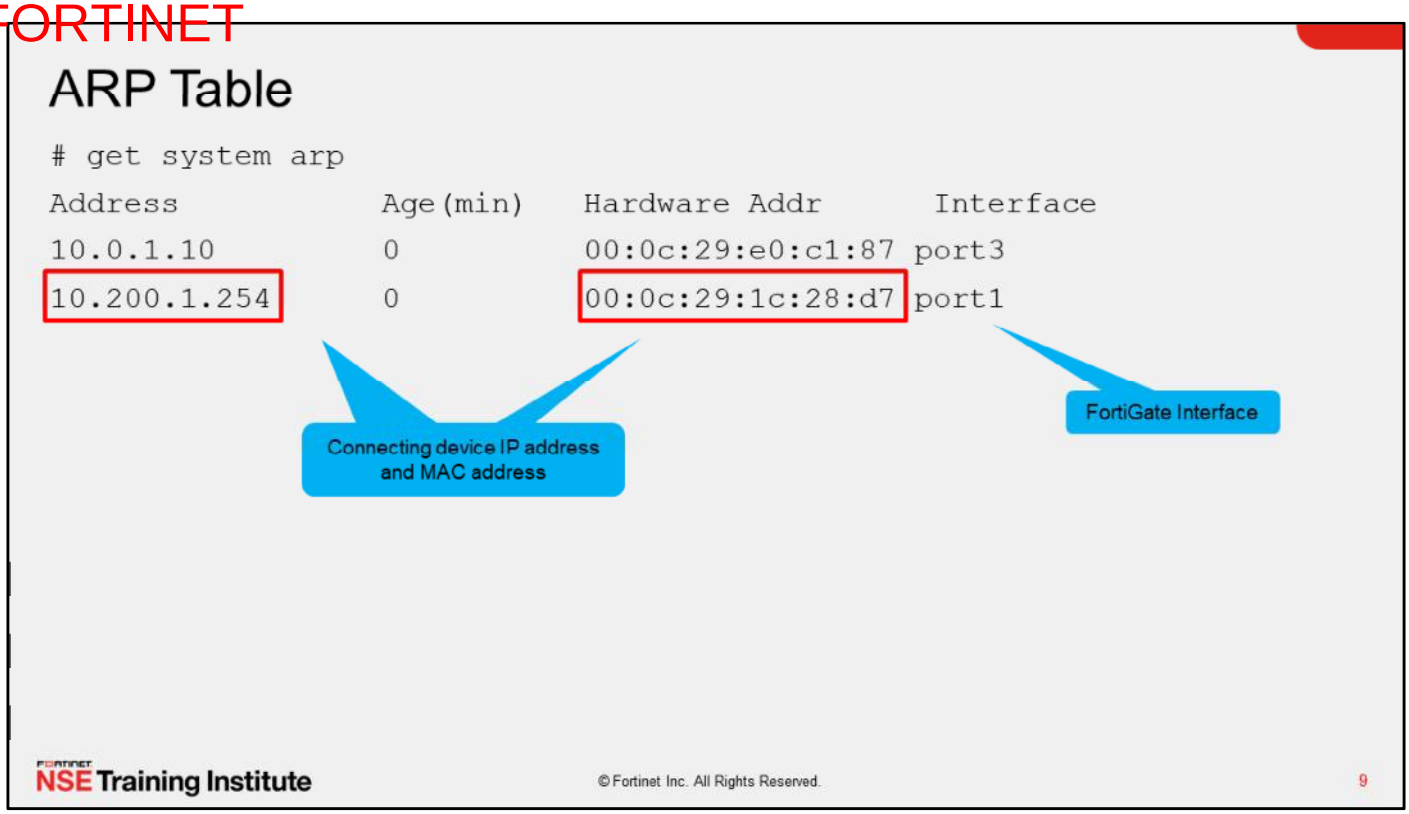

If you suspect that there is an IP address conflict, or that an IP has been assigned to the wrong device, you may need to look at the ARP table. The `get system arp` command is used for that purpose. It shows the FortiGate interface, IP address, and associated MAC address. This command lists the information 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
```

NSE Training Institute 10

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`

NSE Training Institute

11

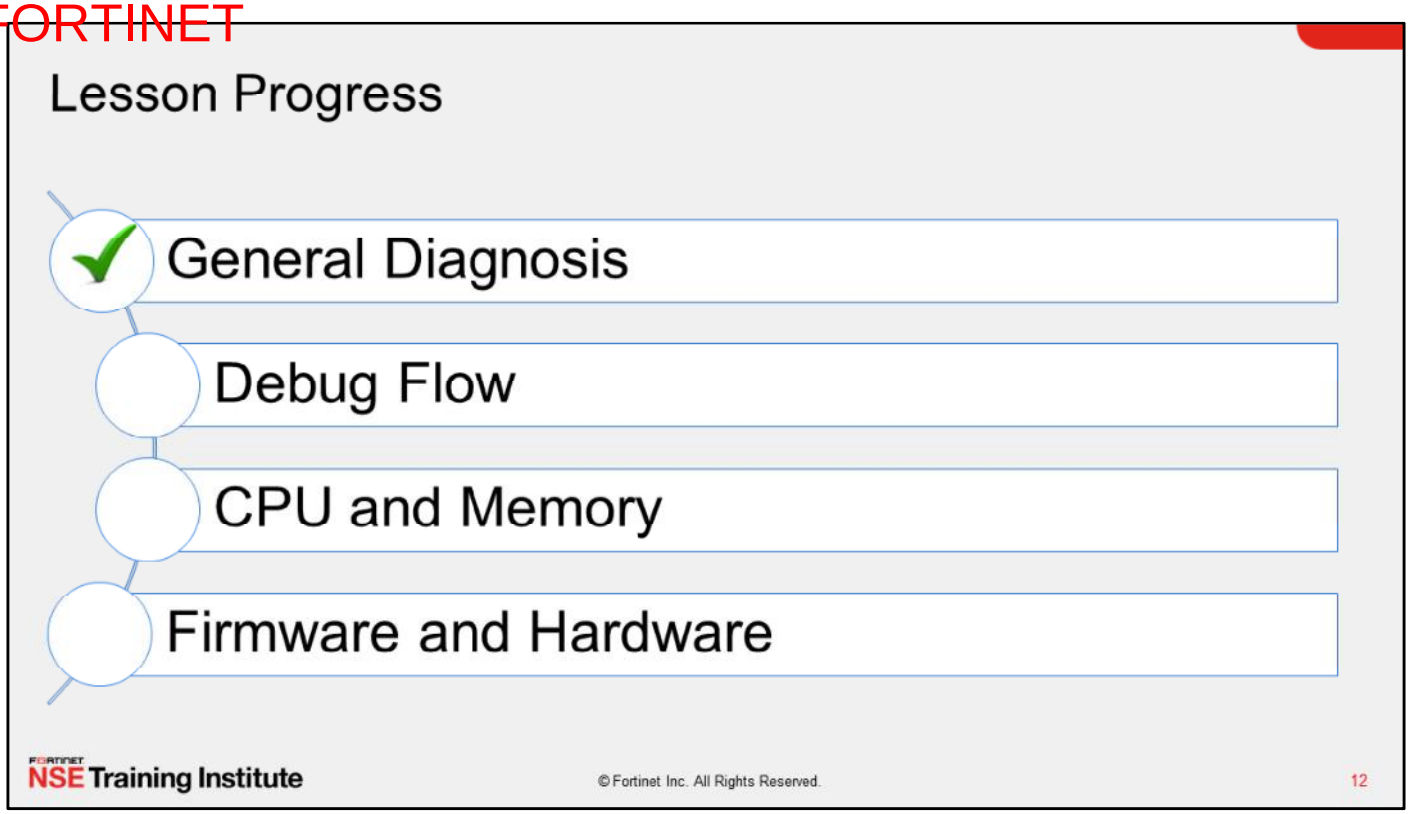

Good job! You now understand general diagnostics.

Now, you will learn about debug flow.

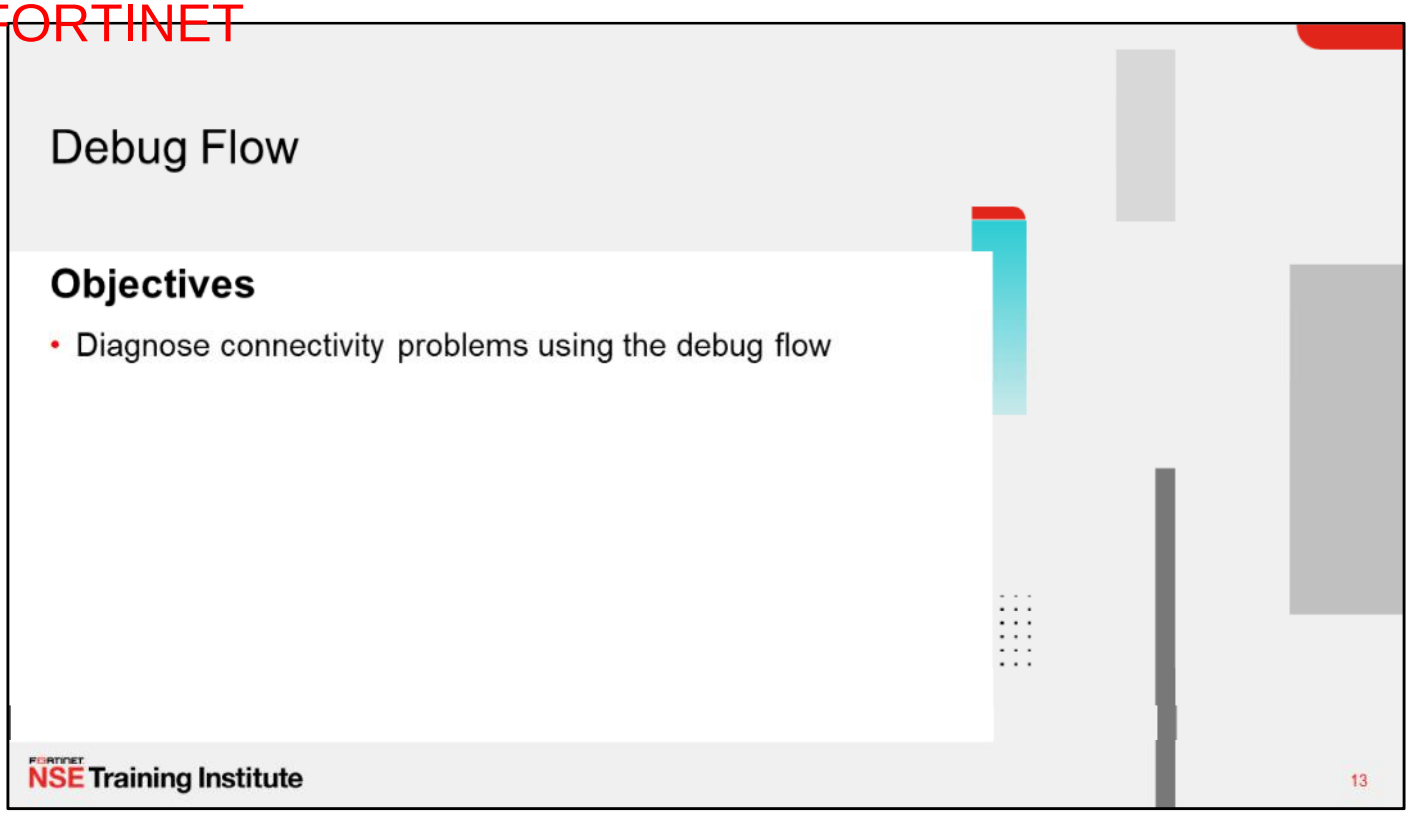

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the debug flow, you will be able to diagnose connectivity problems.

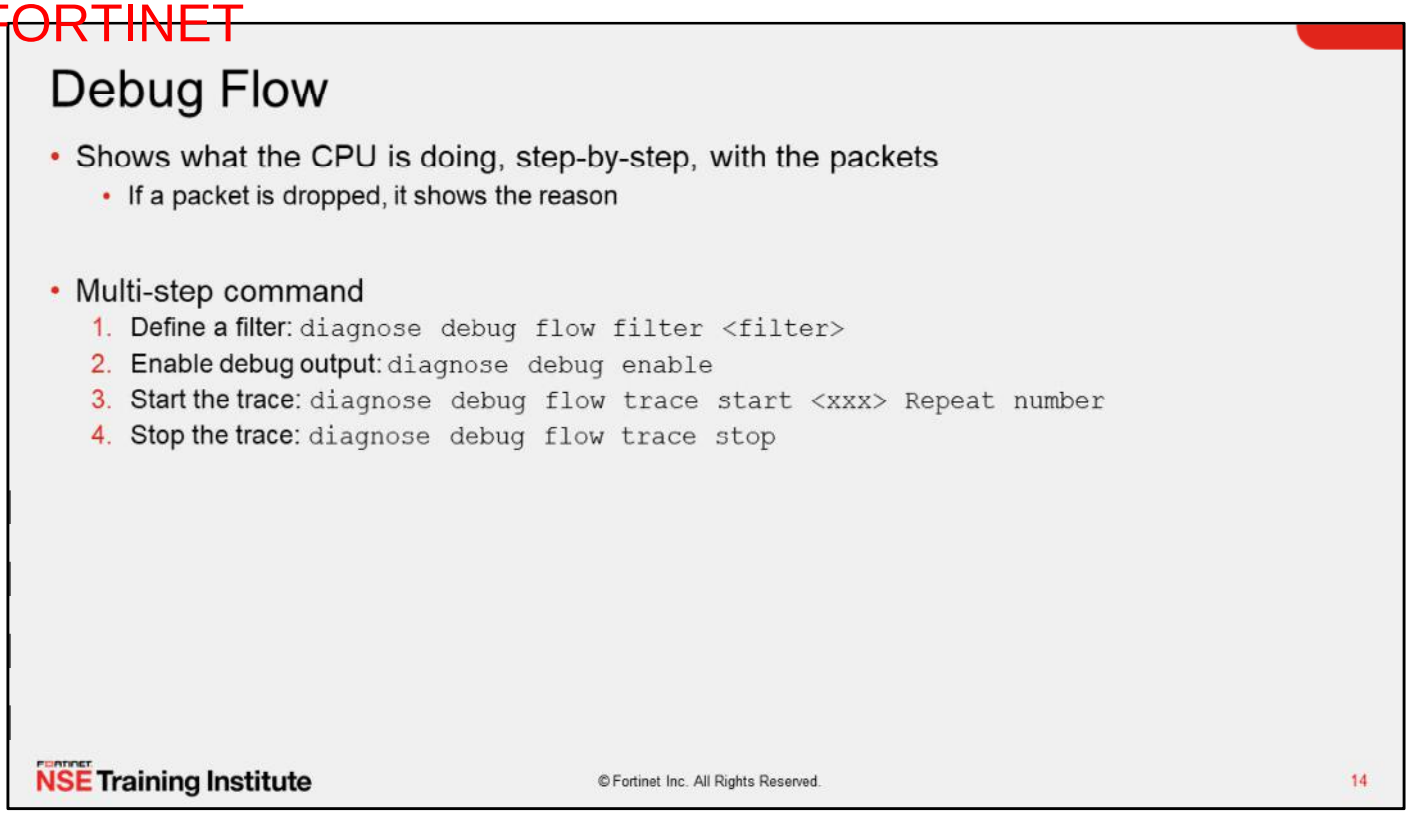

If FortiGate is dropping packets, can a packet capture (sniffer) be used to identify the reason? To find the cause, you should use the debug (packet) flow.

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.

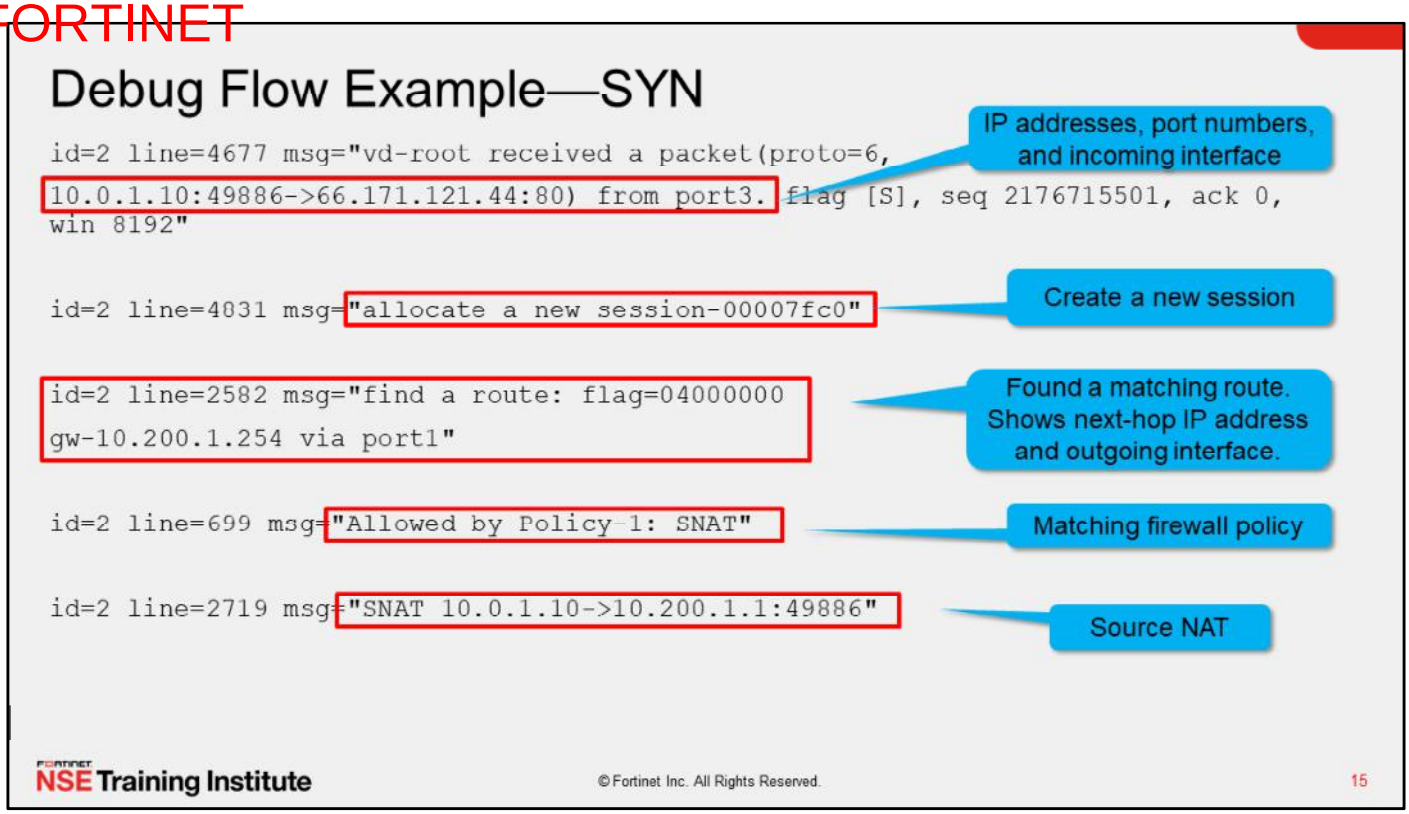

This slide shows an example of a debug flow output, which captures the first packet of a TCP three-way handshake, the SYN packet. It shows:

- 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

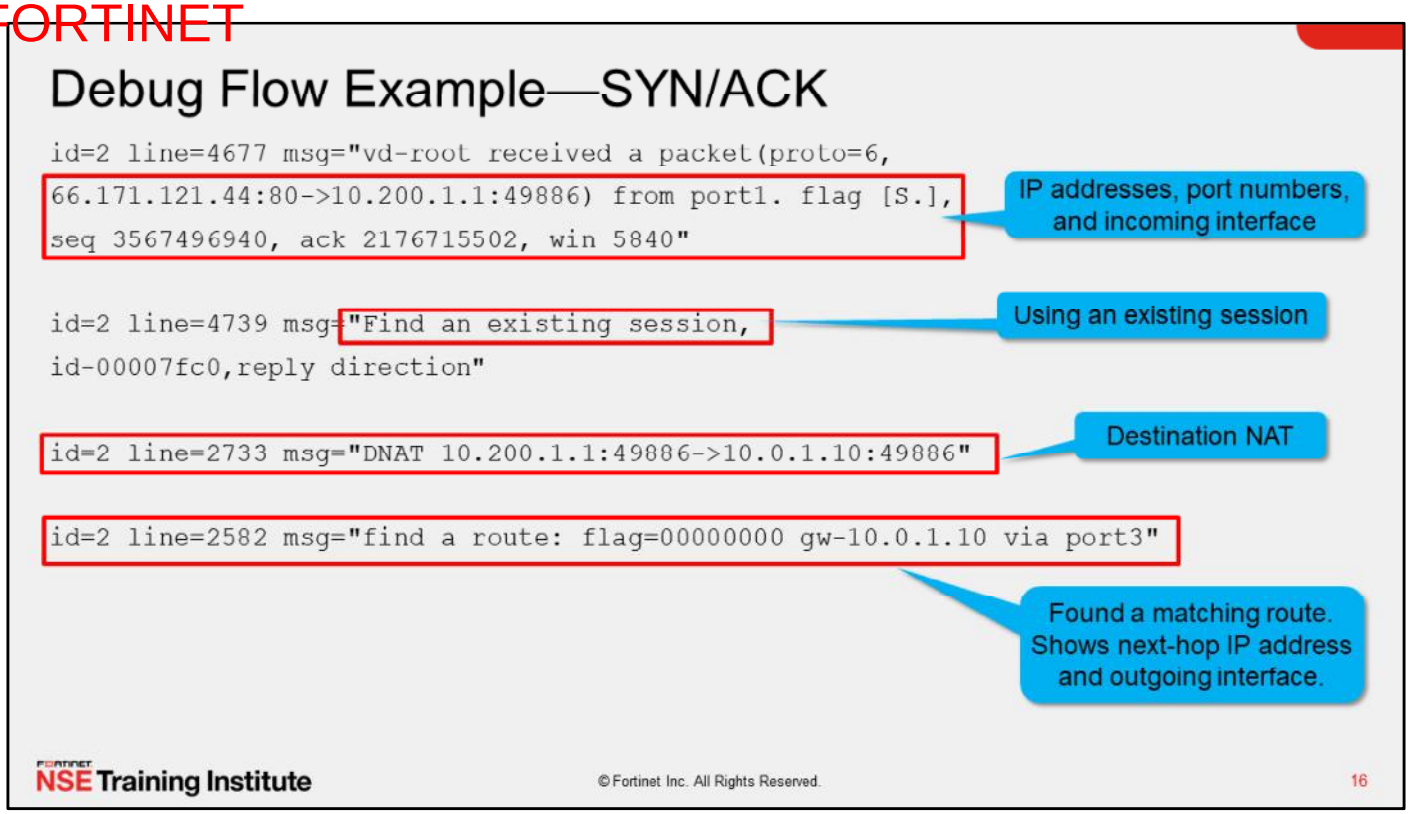

This slide shows the output for the SYN/ACK packet. It shows:

- 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.
- 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.

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

NSE Training Institute

17

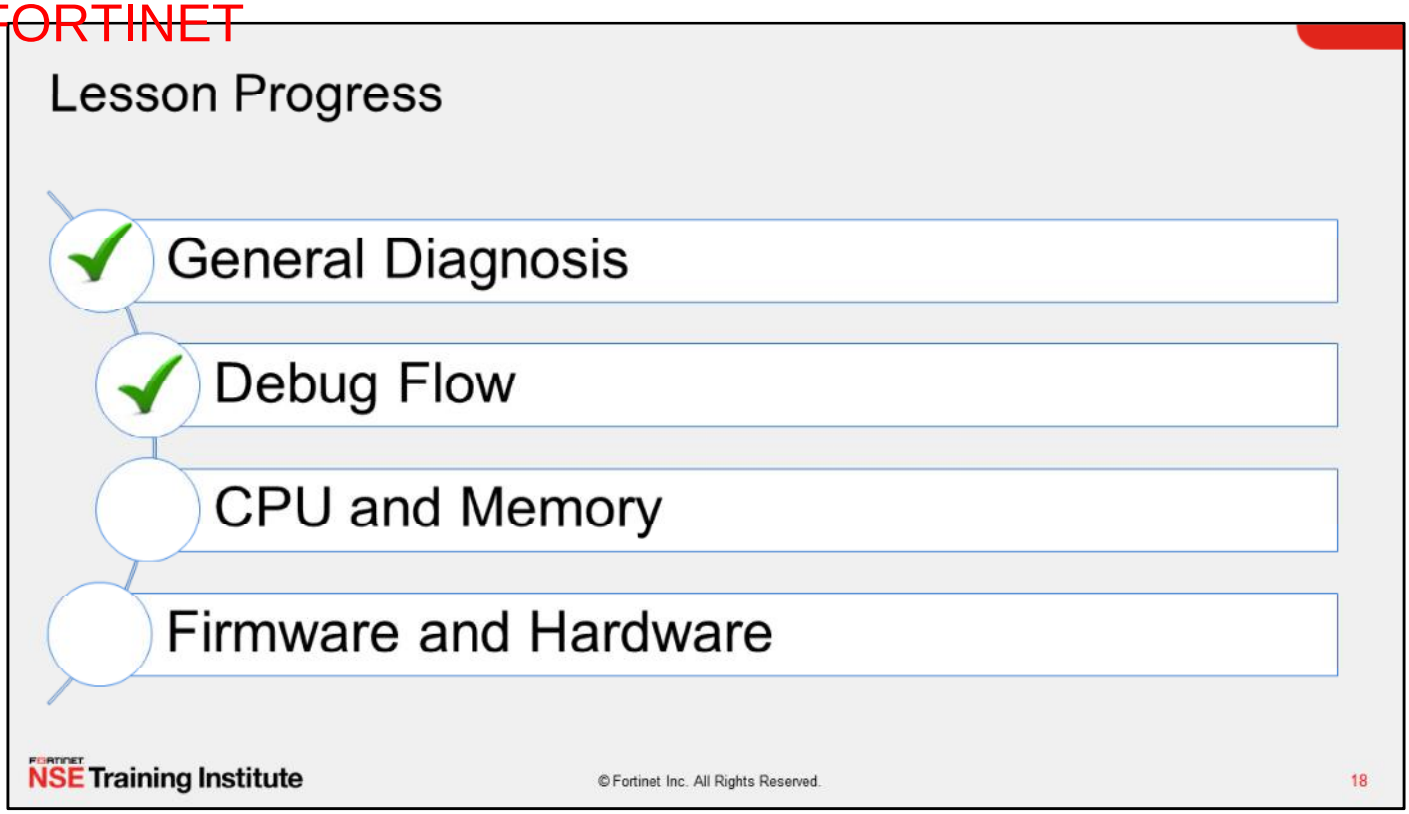

Good job! You now understand debug flow.

Now, you will learn about FortiGate CPU and memory diagnosis.

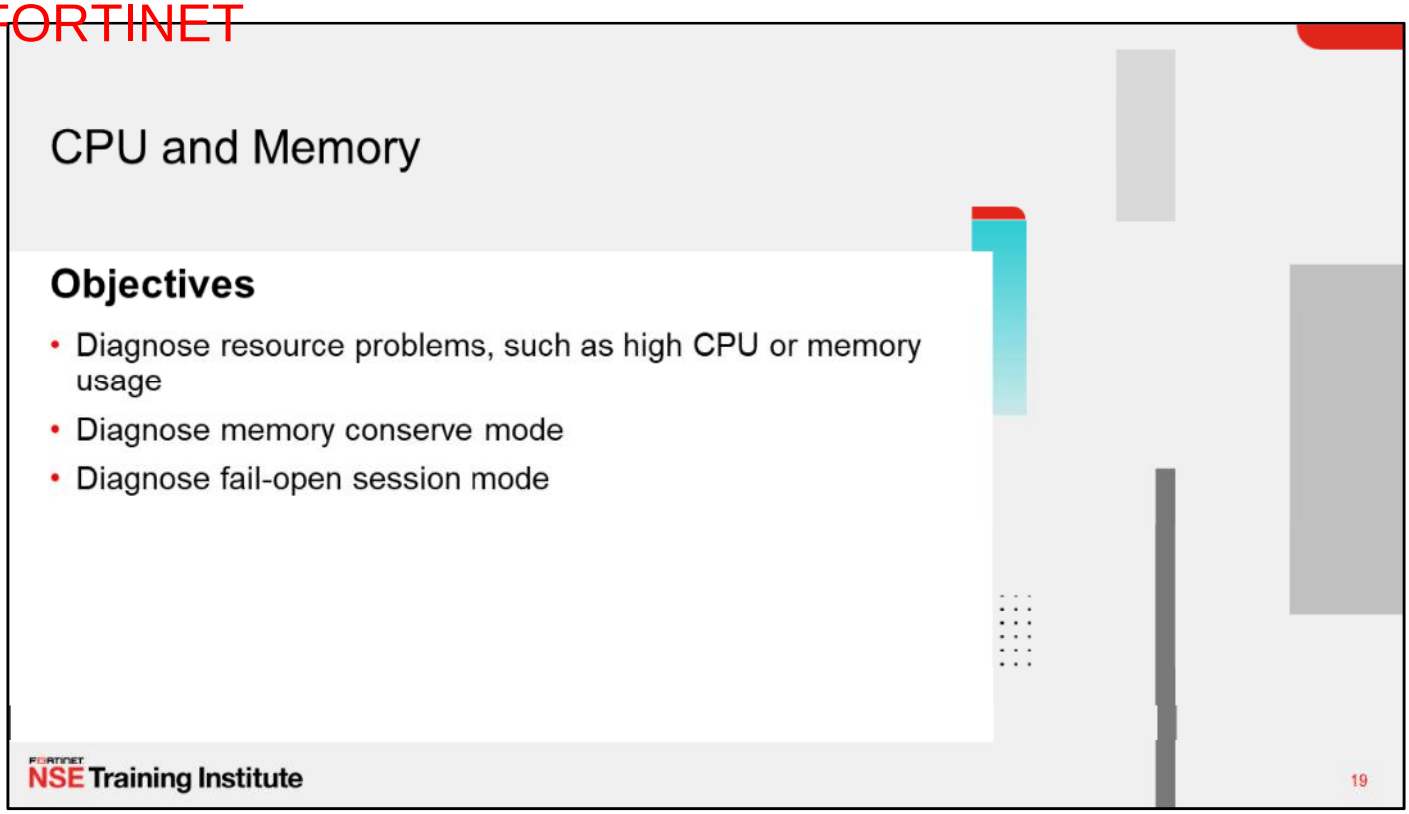

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in CPU and memory, you will be able to diagnose the most common CPU and memory problems.

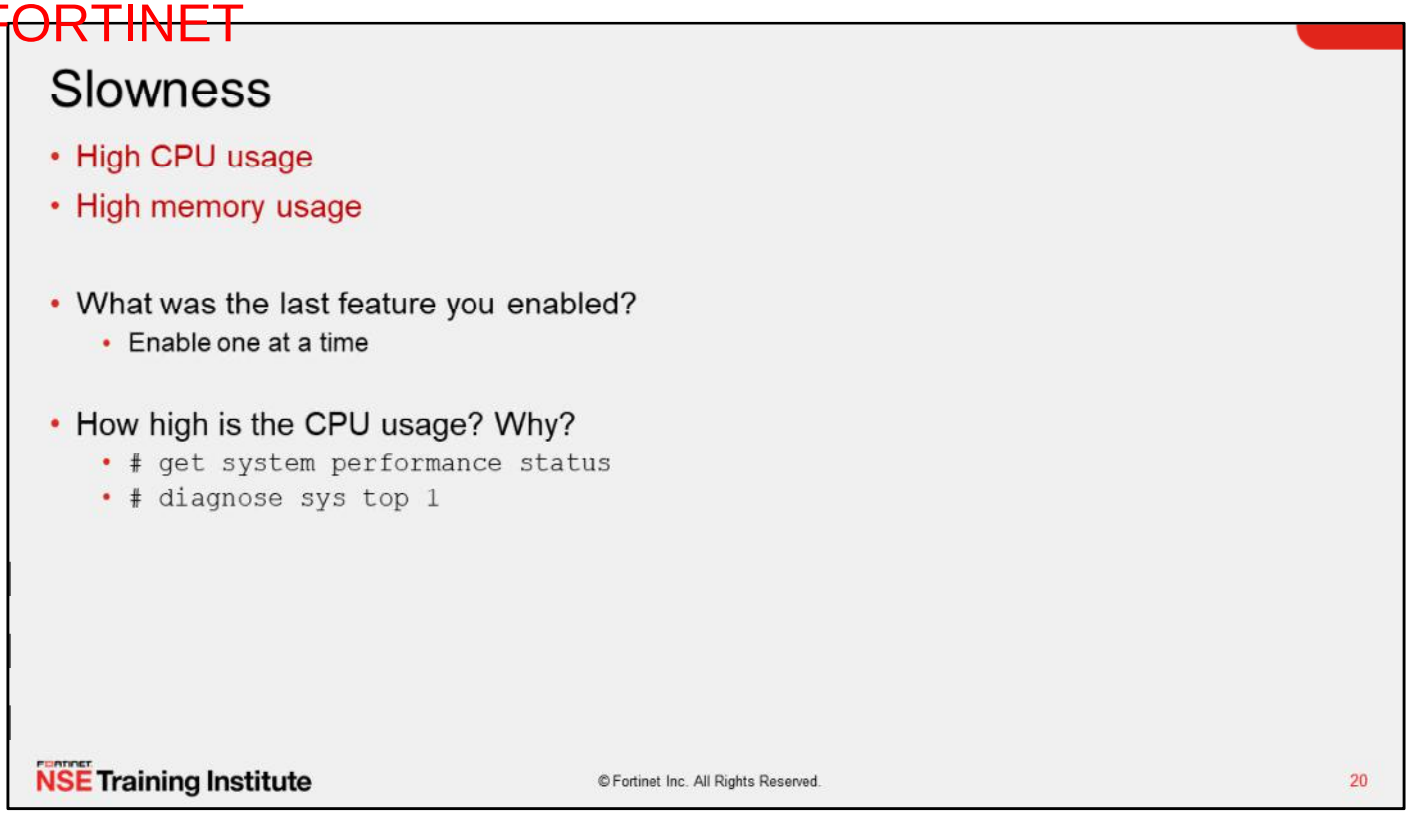

Not all problems are network connectivity failures. Sometimes, there are resource problems in the devices.

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.

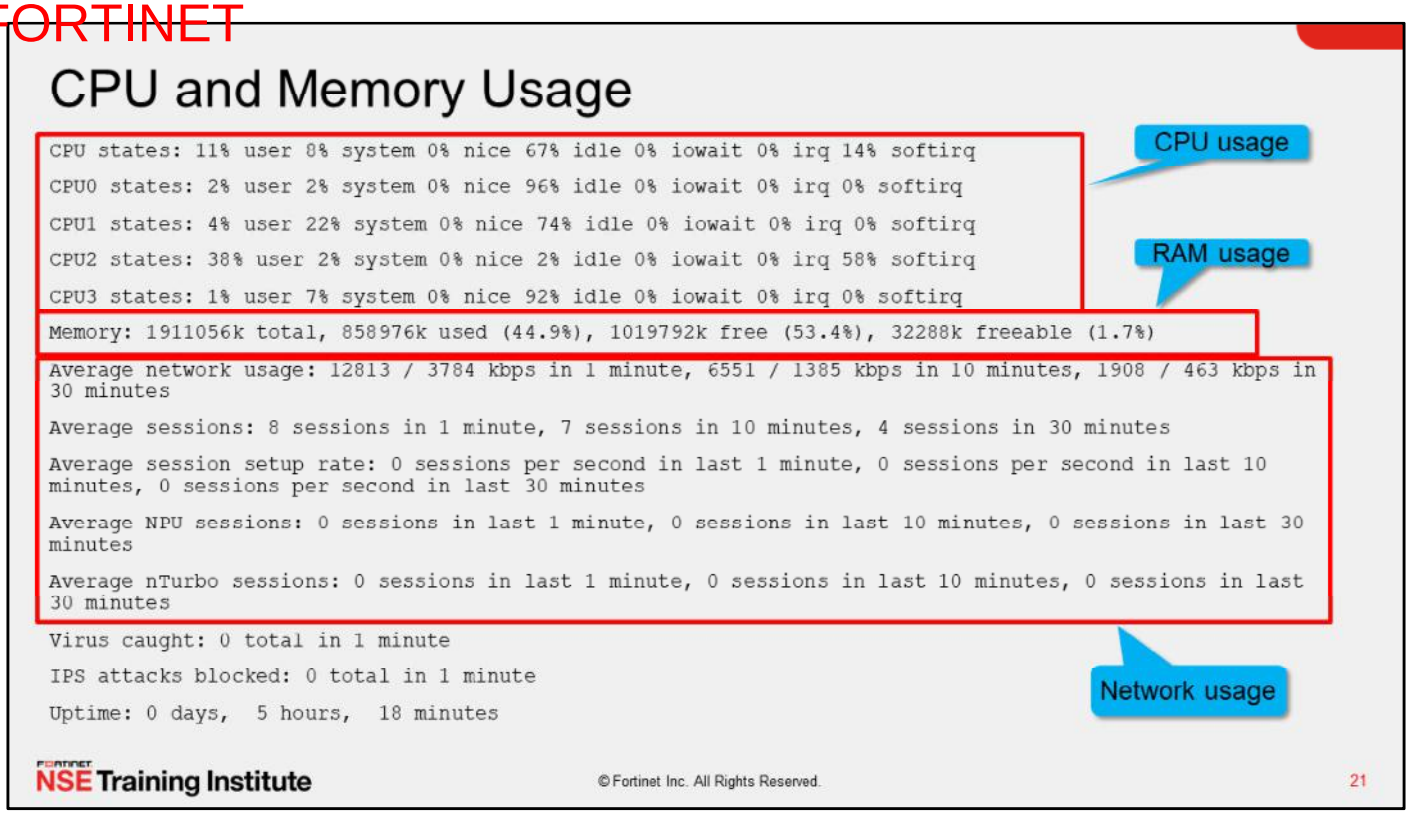

Begin by looking at the `get system performance status` command.

At the top of the example on this slide, the output shows that this FortiGate model has four CPUs: CPU0, CPU1, CPU2, and CPU3. This is followed by the RAM usage.

At the bottom of the example on this slide, the output shows the network traffic.

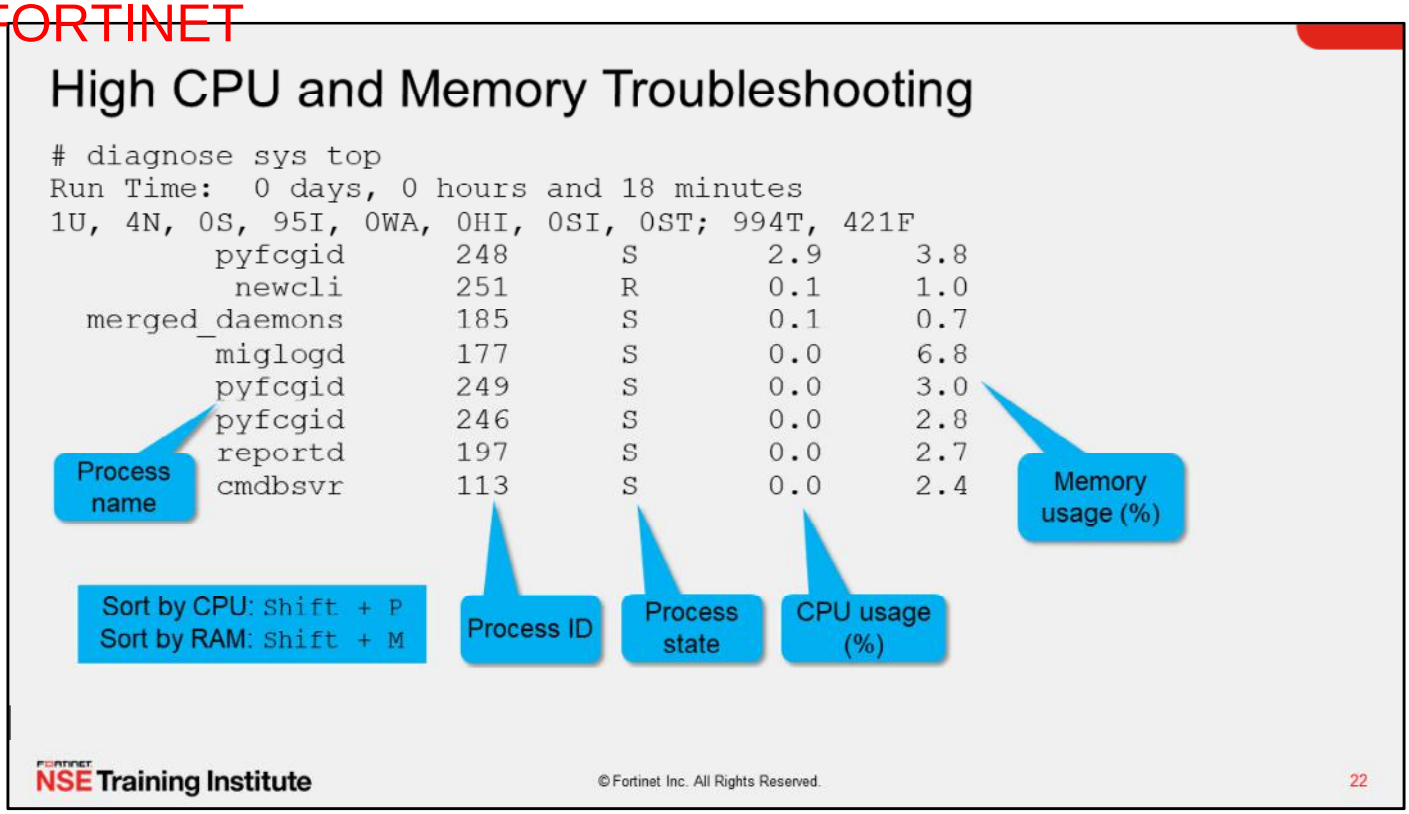

Next, examine the output for `diagnose sys top`. It lists processes that use the most CPU or memory. Some common processes include:

- `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
```

NSE Training Institute

23

If memory usage goes 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.

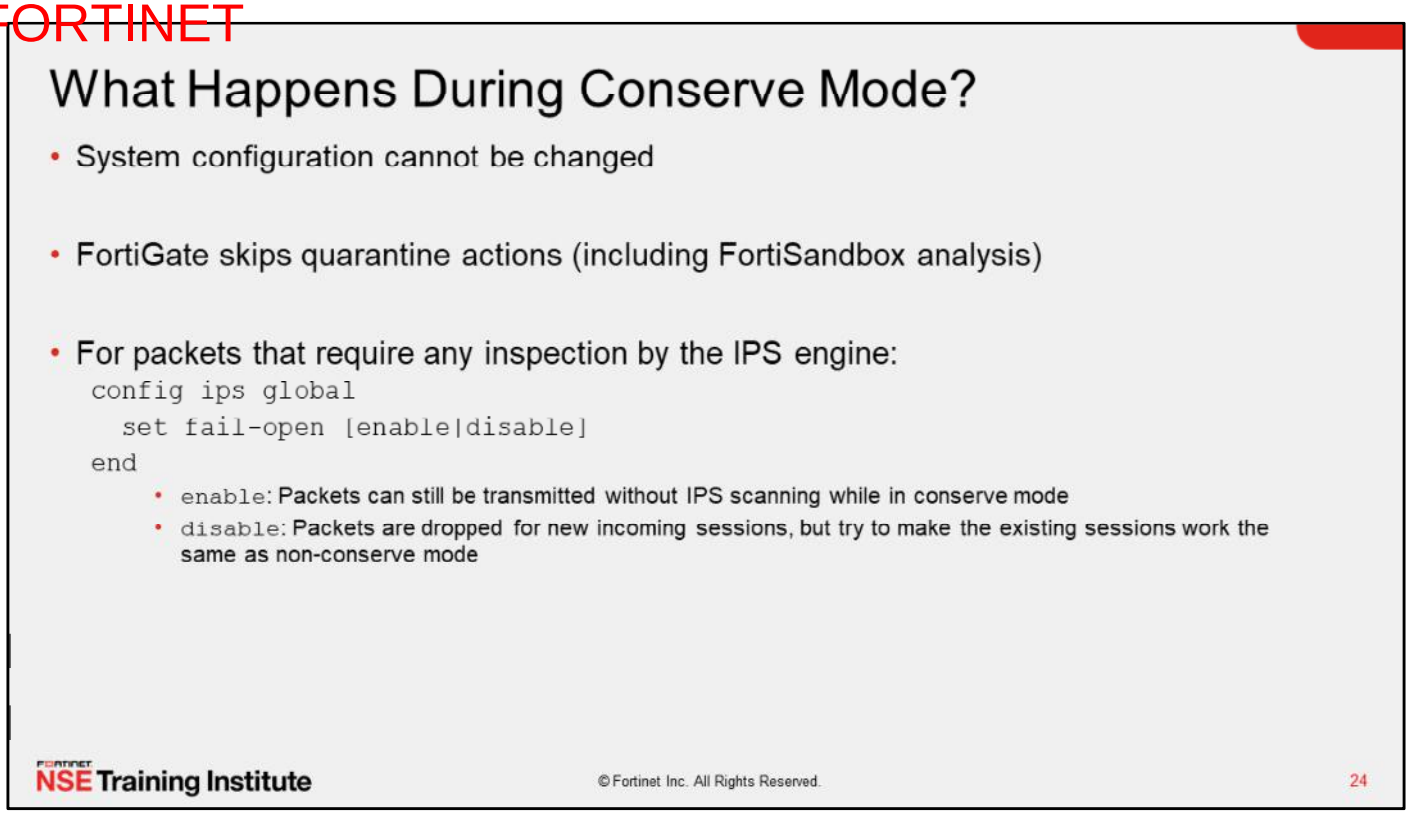

What actions does FortiGate take to preserve memory while in conserve mode?

- 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.
- The `fail-open` setting under `config ips global` controls how the IPS engine behaves during the IPS failures. It is important to understand that IPS fail open is not just for conserve mode. It kicks in whenever IPS fails. The most common failure could be due to a high CPU issue, or it could be due to a high memory (conserve mode) issue as well. If the setting is enabled, packets can still be transmitted while in conserve mode (or during any other IPS failure) but not inspected by IPS. If the setting is disabled, packets are dropped for new incoming sessions, but FortiOS will try to make the existing sessions work the same 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.

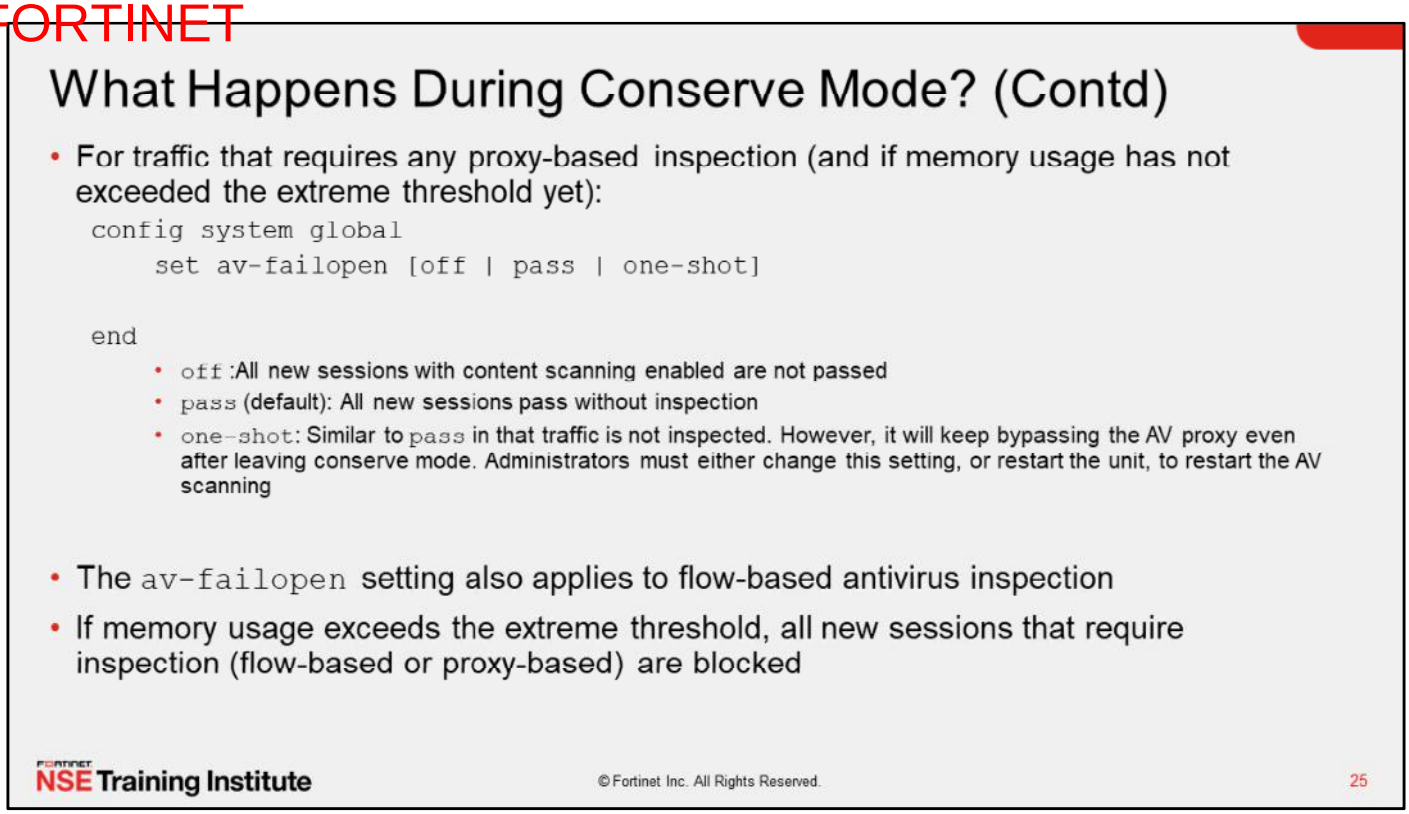

The `av-failopen` setting defines the action that is applied to any proxy-based inspected traffic, while the unit is in conserve mode (and as long as the memory usage does not exceed the extreme threshold). This setting also applies to flow-based AV inspection. Three different actions can be configured:

- `off`: All new sessions with content scanning enabled are not passed
- `pass` (default): All new sessions pass without inspection
- `one-shot`: Similar to `pass` in that traffic passes without inspection. However, it will keep bypassing the AV proxy even after it leaves conserve mode. Administrators must either change this setting, or restart the unit, to restart the AV 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

NSE Training Institute

26

The `diagnose hardware sysinfo conserve` command is used to identify if a FortiGate device is currently in memory conserve mode.

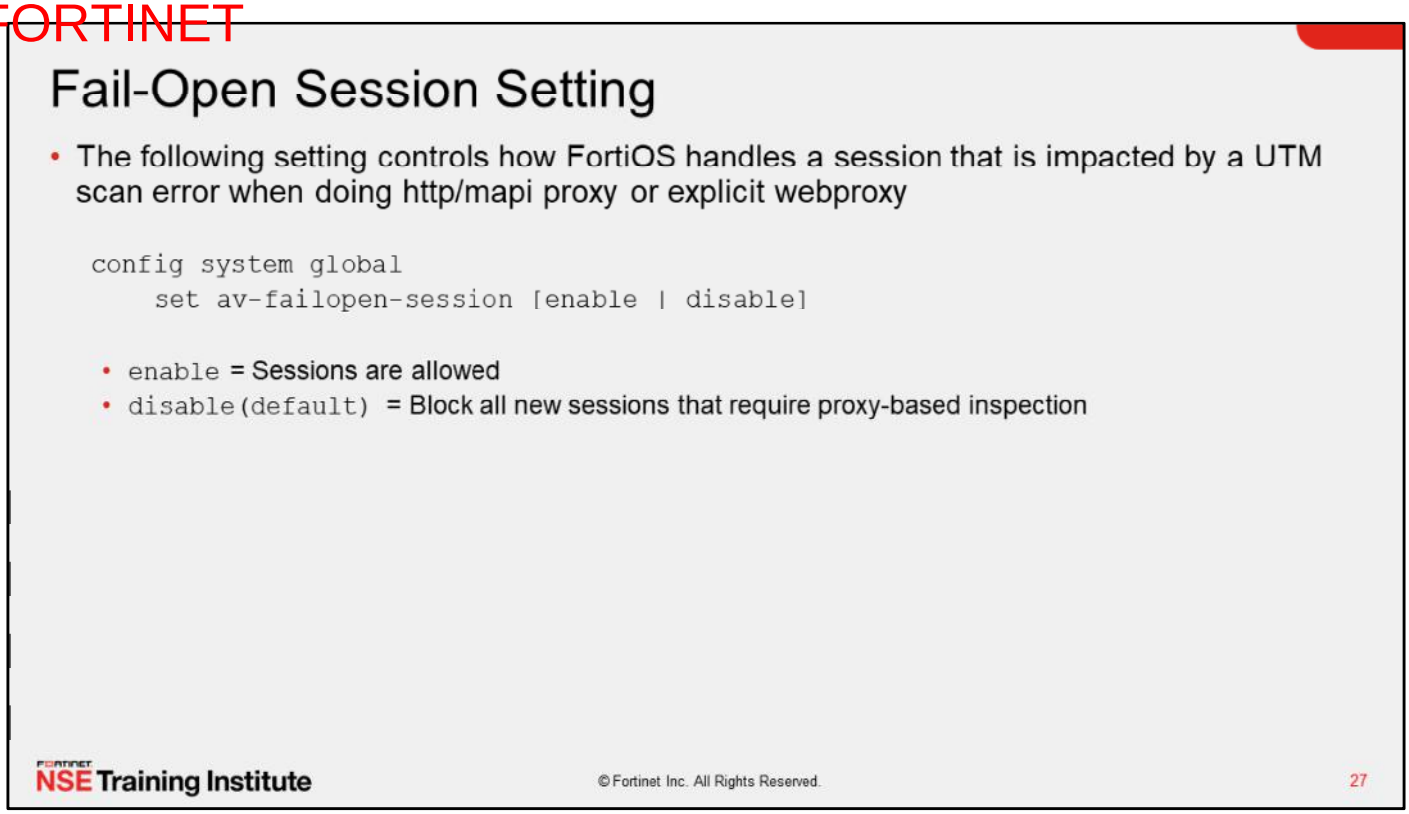

Another undesirable state for FortiGate is the fail-open session mode. This mode kicks in, not during a high-memory situation, but when a proxy on FortiGate runs out of available sockets to process more proxy-based inspected traffic.

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

NSE Training Institute

28

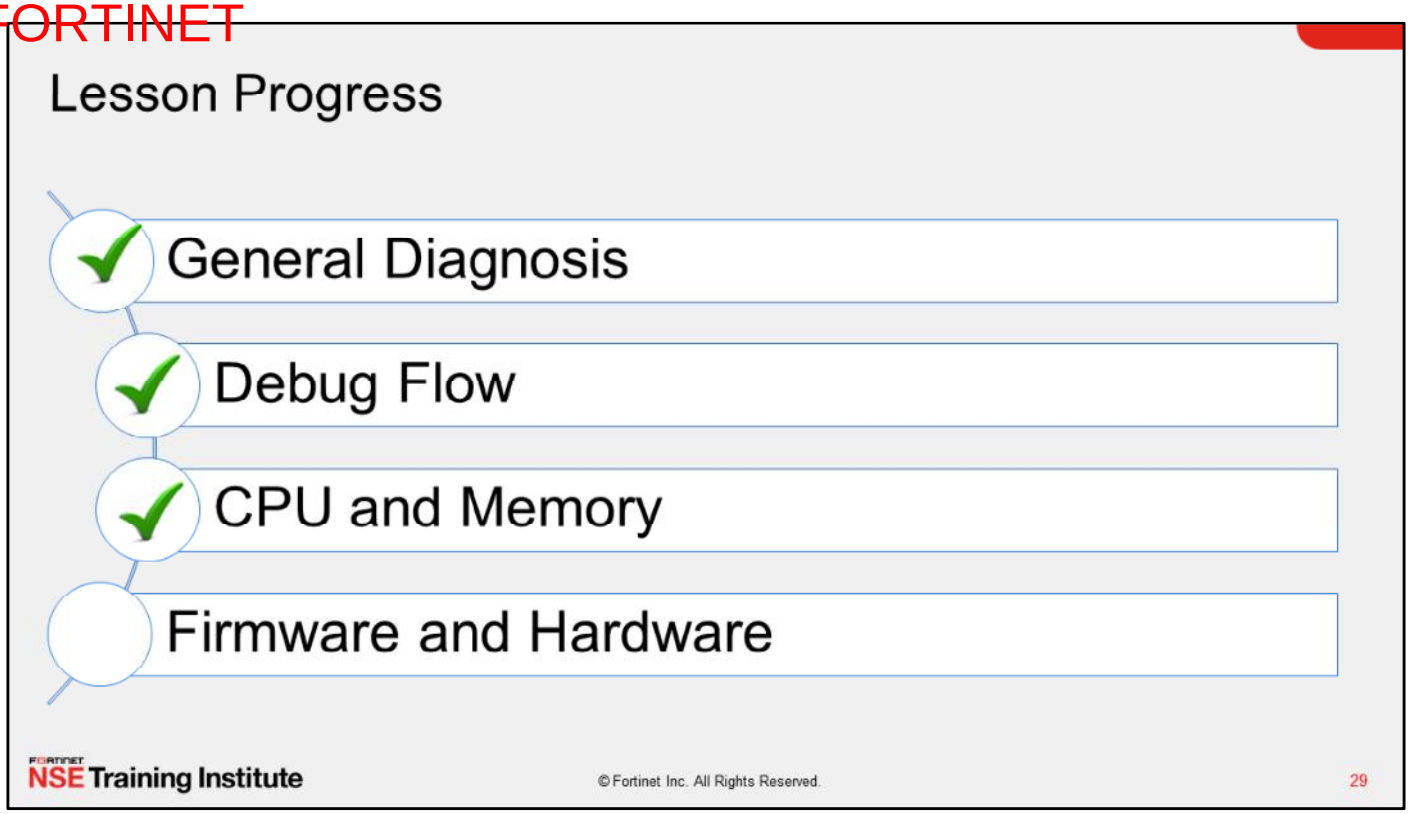

Good job! You now understand FortiGate CPU and memory diagnosis.

Now, you will learn about FortiGate firmware and hardware diagnosis.

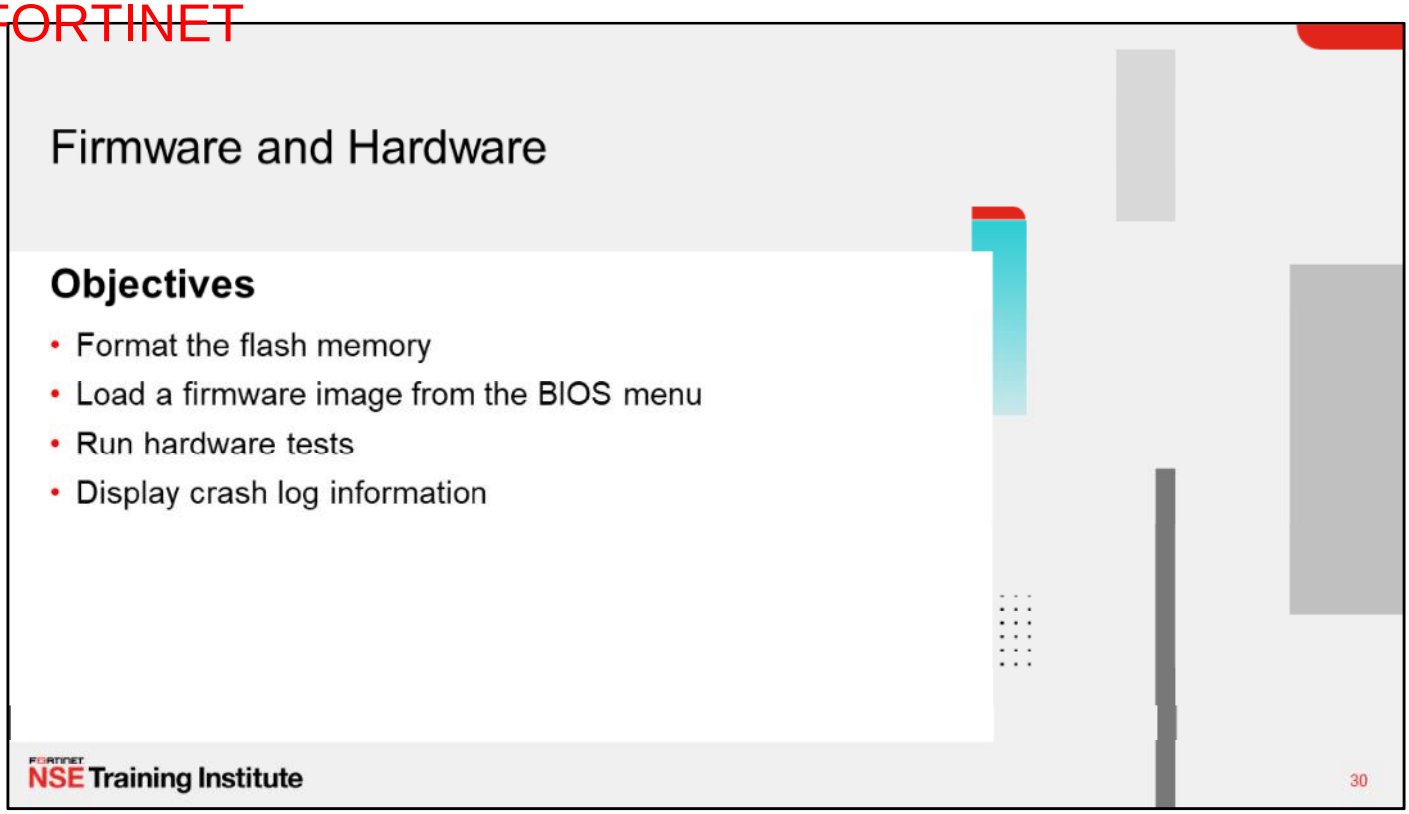

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in firmware and hardware, you will be able to diagnose the most common firmware and hardware problems.

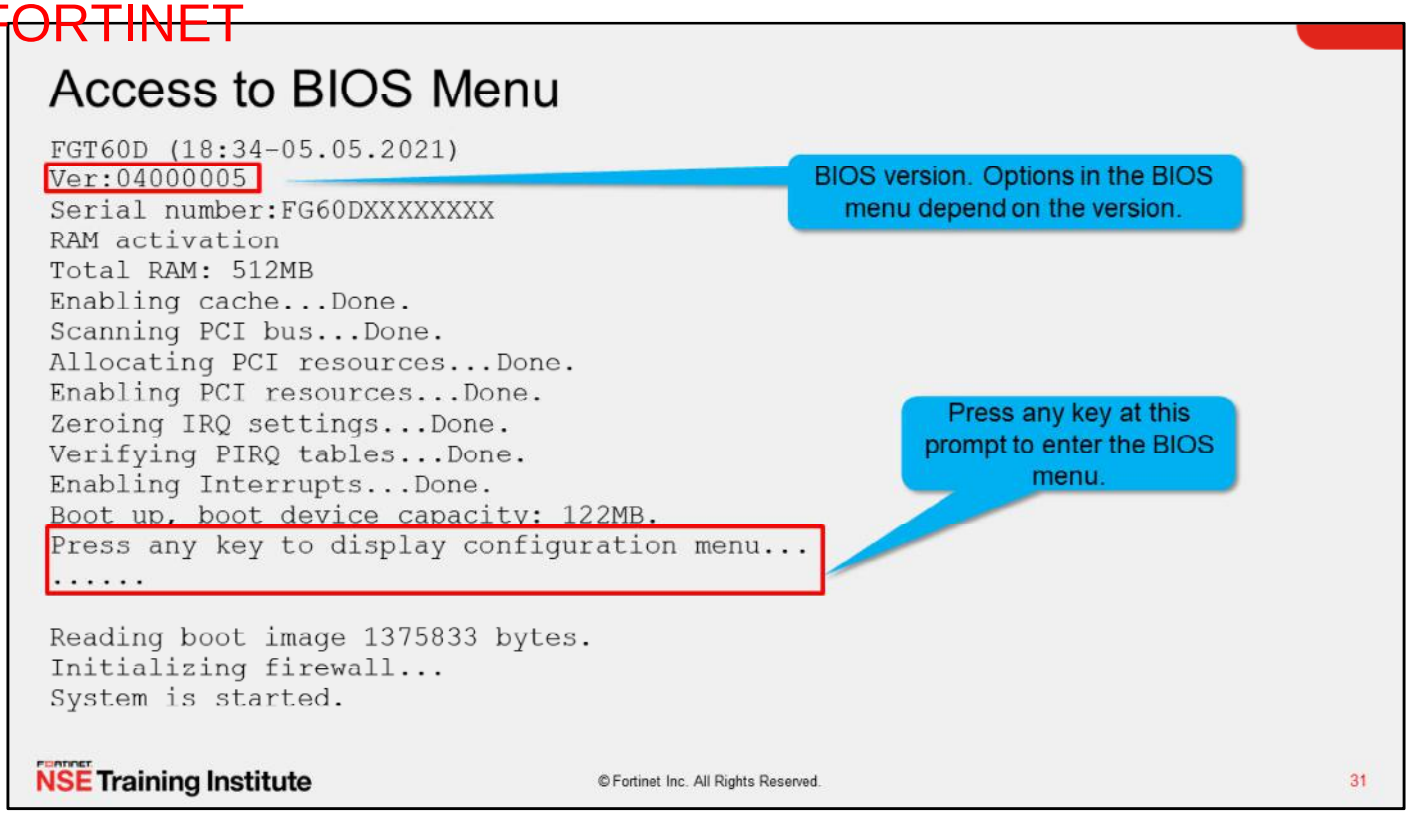

On the FortiGate BIOS, administrators can run some operations over the flash memory and firmware images. To access the BIOS menu, you must reboot the device while connected to the console port. The booting process, at one point, shows the following message:

```
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.

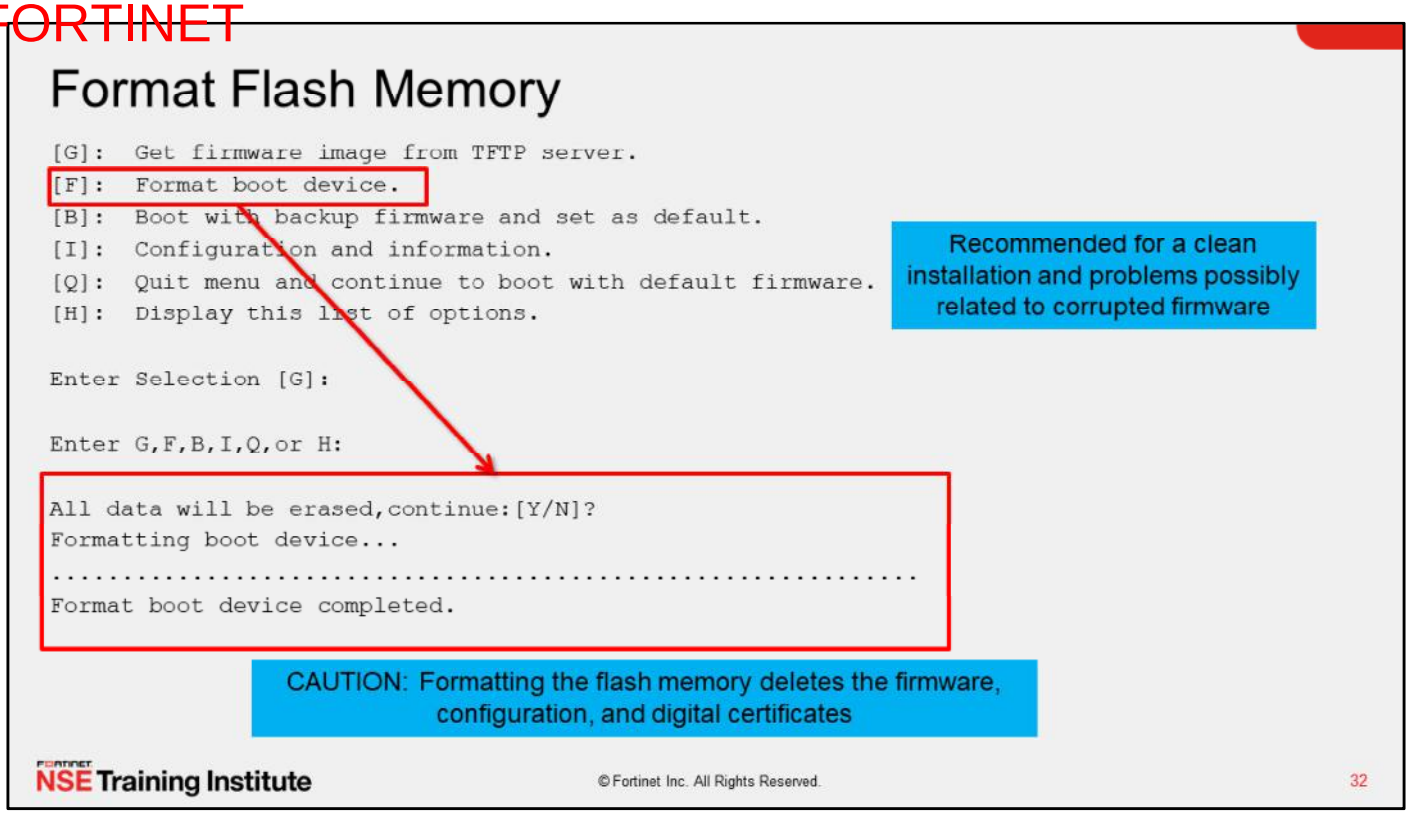

From the BIOS menu, select `F` to format the flash memory.

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.

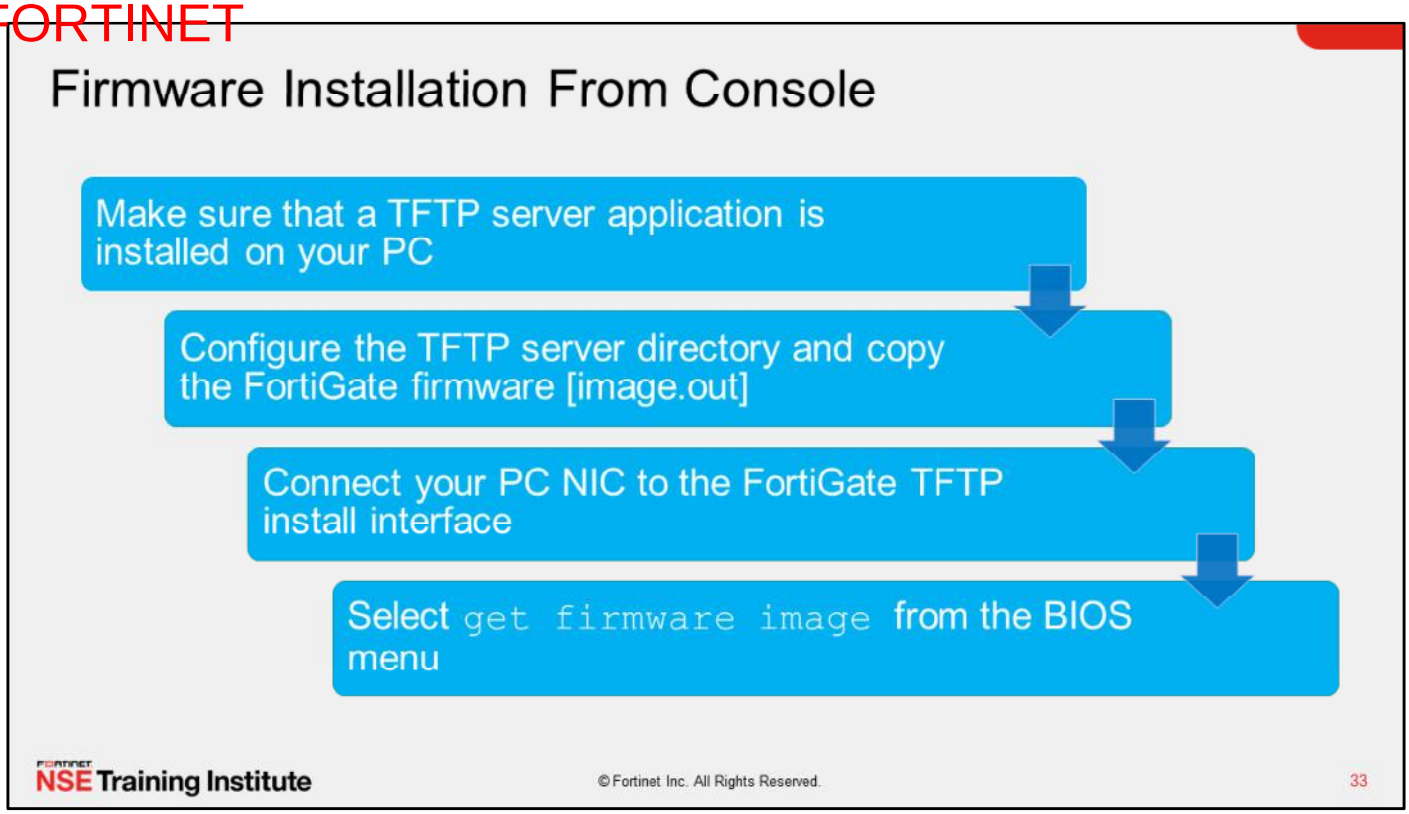

After reformatting the flash memory, you must install the firmware image from the BIOS. Follow these steps:

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 install interface.
4. Select `get firmware image` from the BIOS menu.

The interface assigned as the TFTP install interface depends on the model. However, and in most cases, it is either the *port1* or *internal* interface.

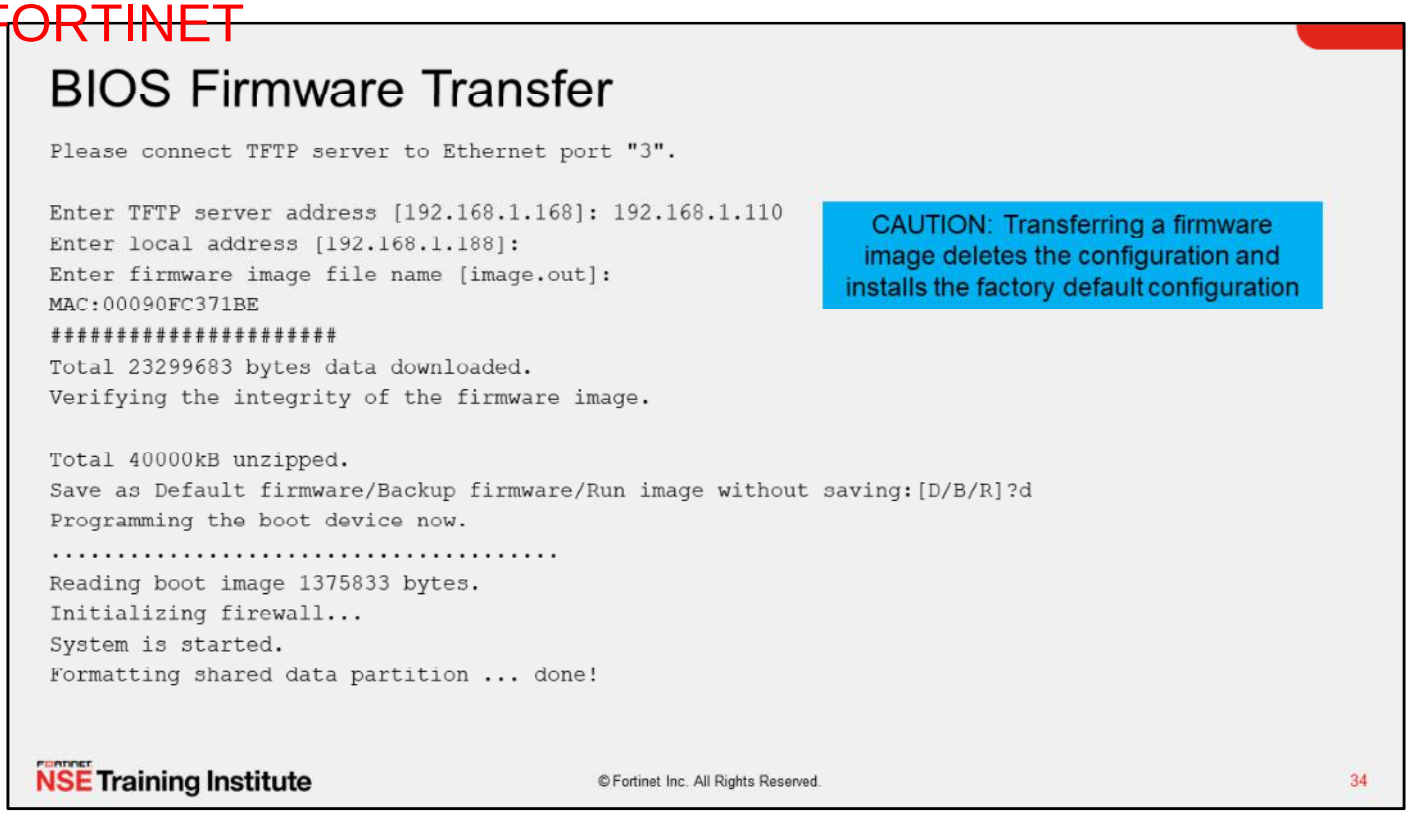

From the BIOS menu, select option `G` to install a new firmware.

The BIOS will ask for:
- 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 flash. Once 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.

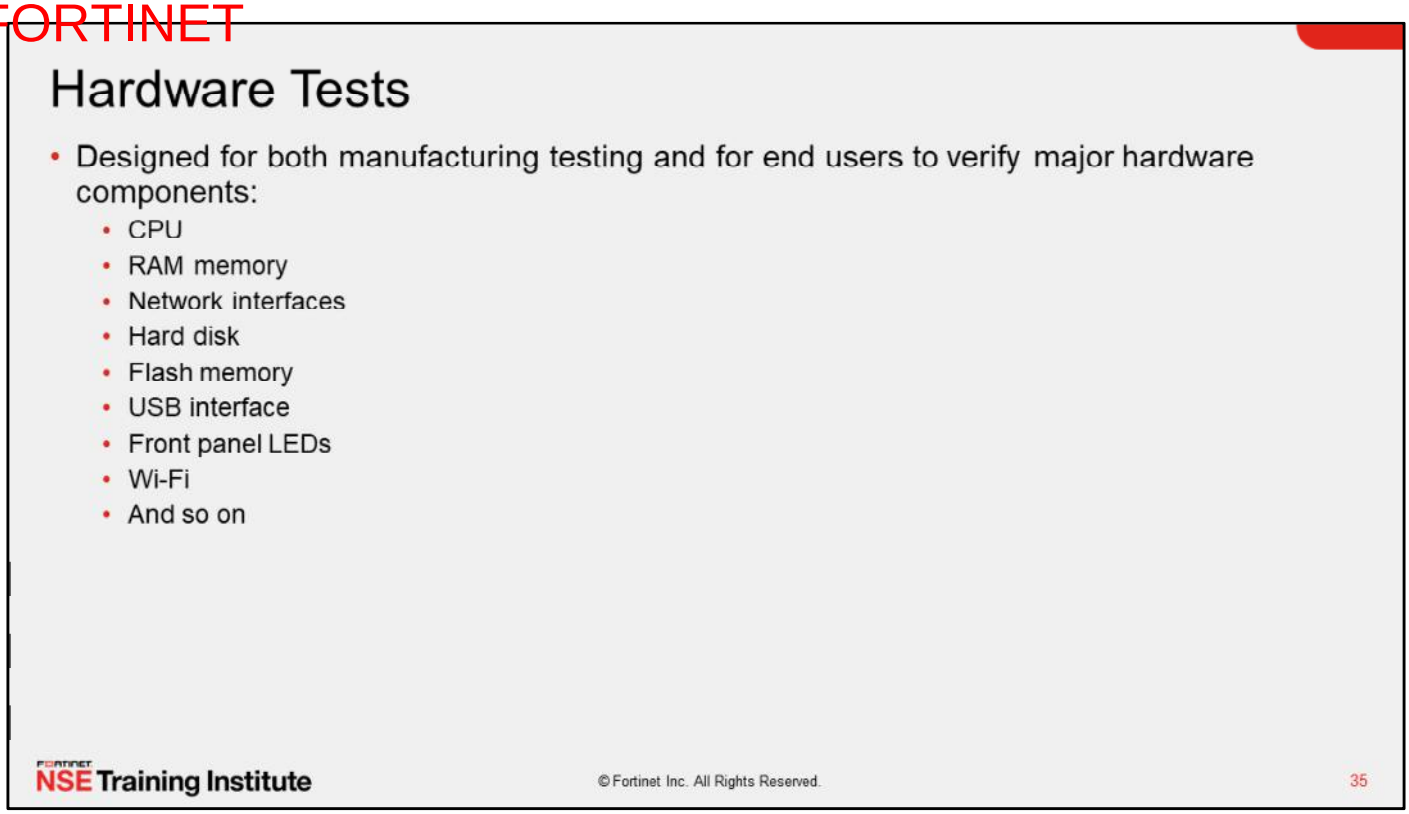

As with any other electronic device, damage to RAM can cause intermittent crashes.

If you suspect hardware failure, you can run hardware tests.

How do you run the hardware tests? It depends on the FortiGate model.

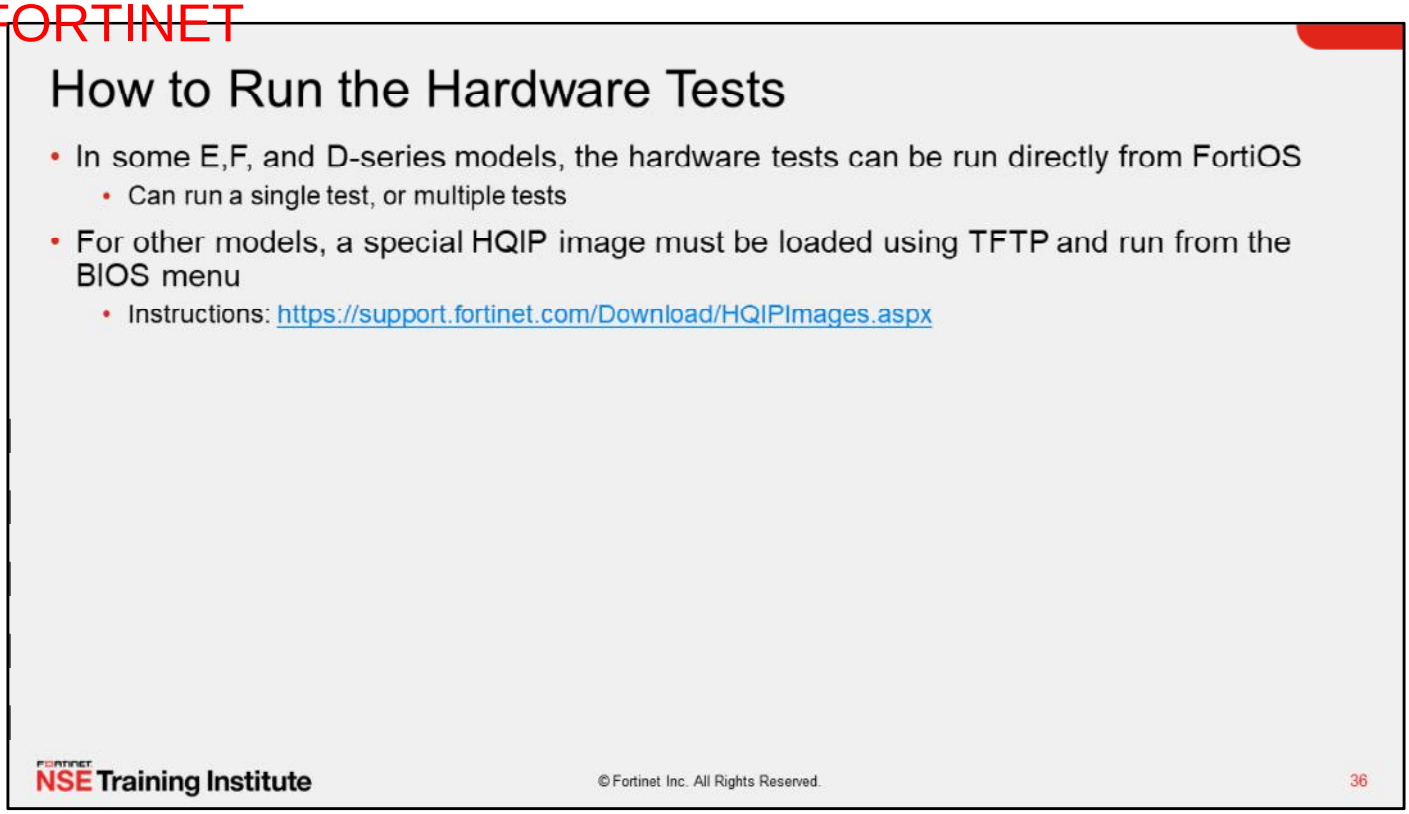

For some FortiGate E, F, and D-series models, you can run the hardware tests directly from the FortiOS CLI.

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.

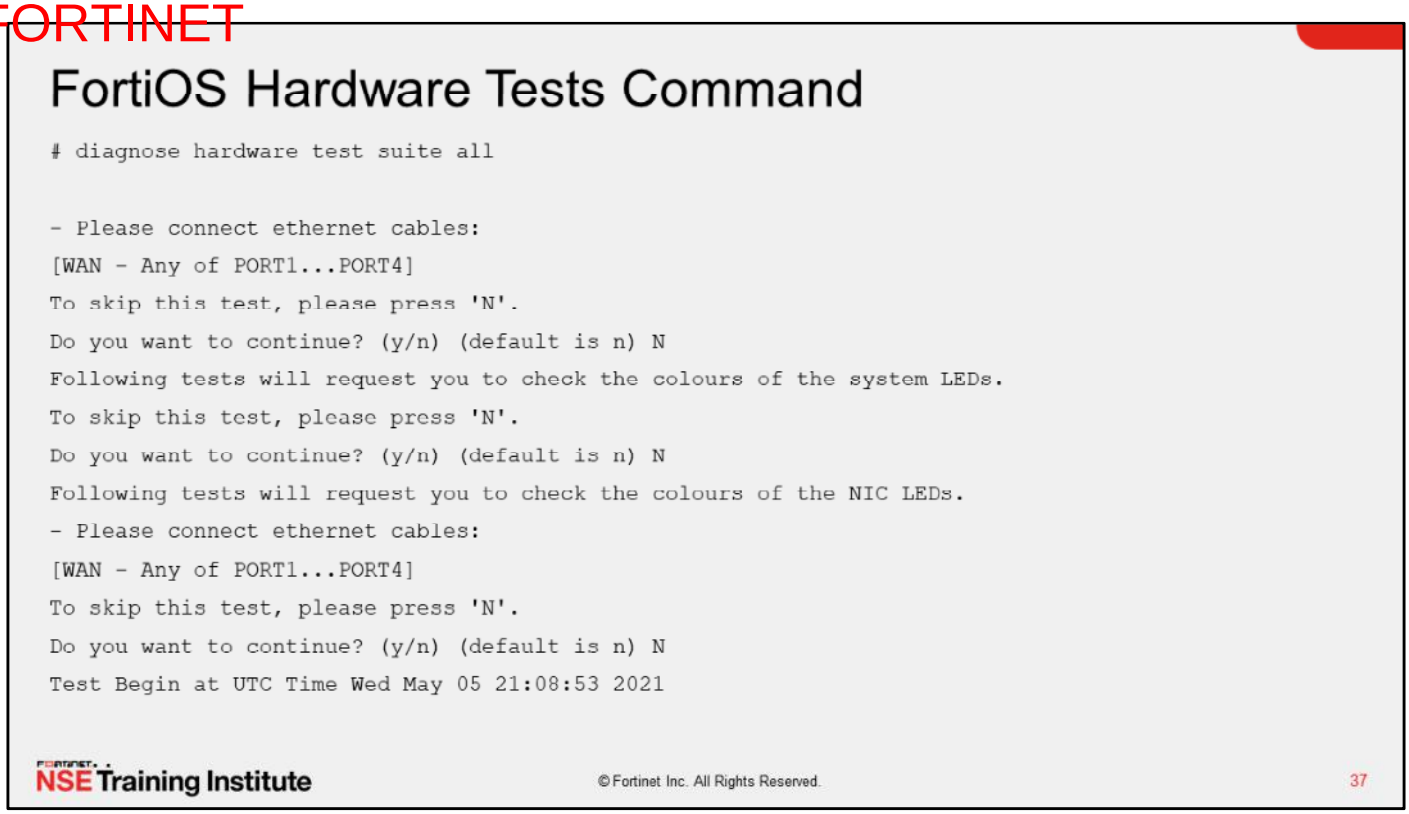

For some models, the command `diagnose hardware test suite all` runs the hardware tests from FortiOS. The hardware tests require user interaction while running. Users can skip some of the steps. Some tests require connecting external devices (such as USB flash drives) or network cables to FortiGate.

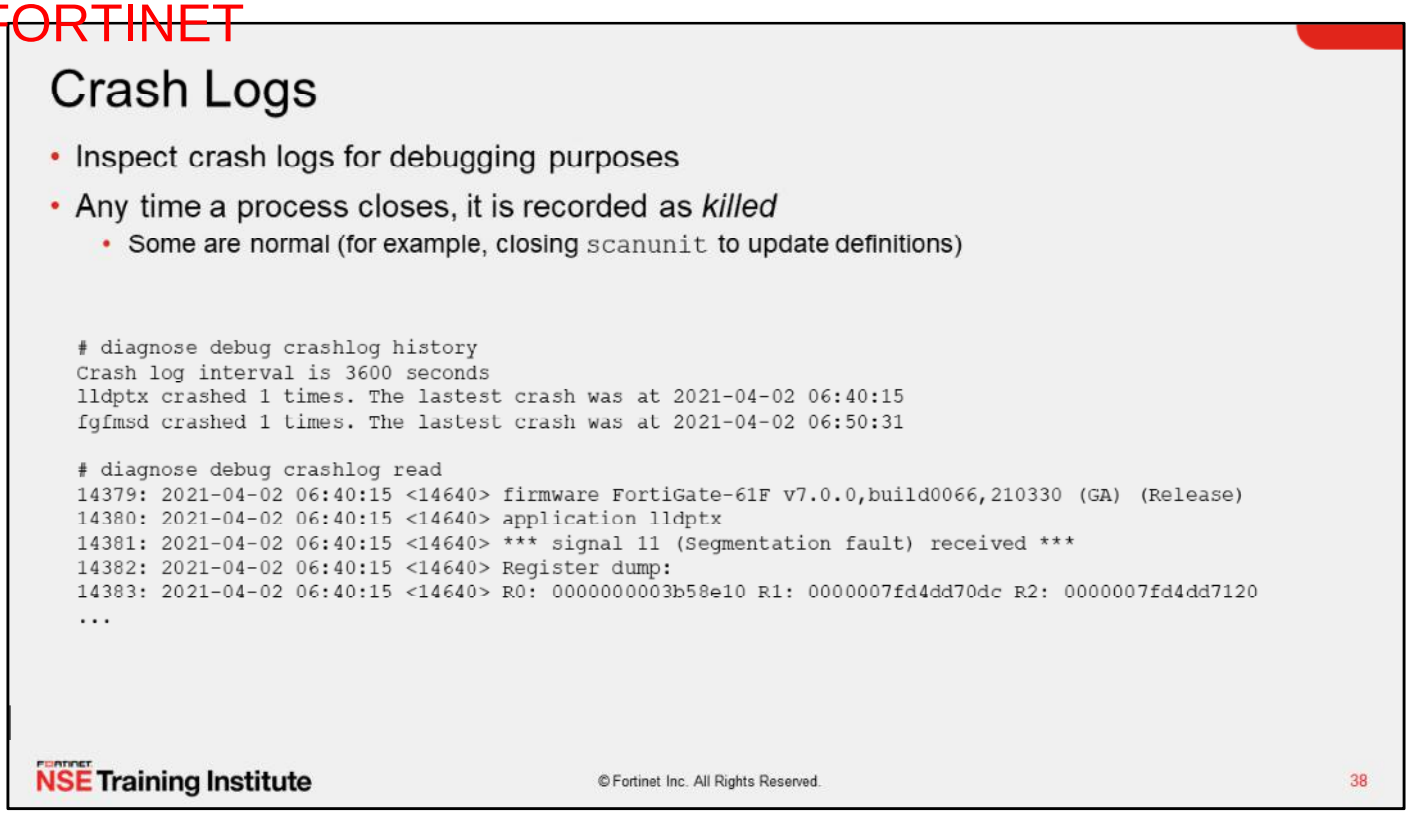

Another area you might want to monitor, purely for diagnostics, is the crash logs. Crash logs are available through the CLI. Any time a process is closed for any reason, the crash log records this as a crash. Most of the logs in the crash log are normal. For example, any time the antivirus definitions 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 might indicate problems. For that reason, crash logs are frequently requested by Fortinet Technical Support for troubleshooting purposes. This slide shows the command you have to use to get a crash log.

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"
```

NSE Training Institute 39

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

NSE Training Institute

40

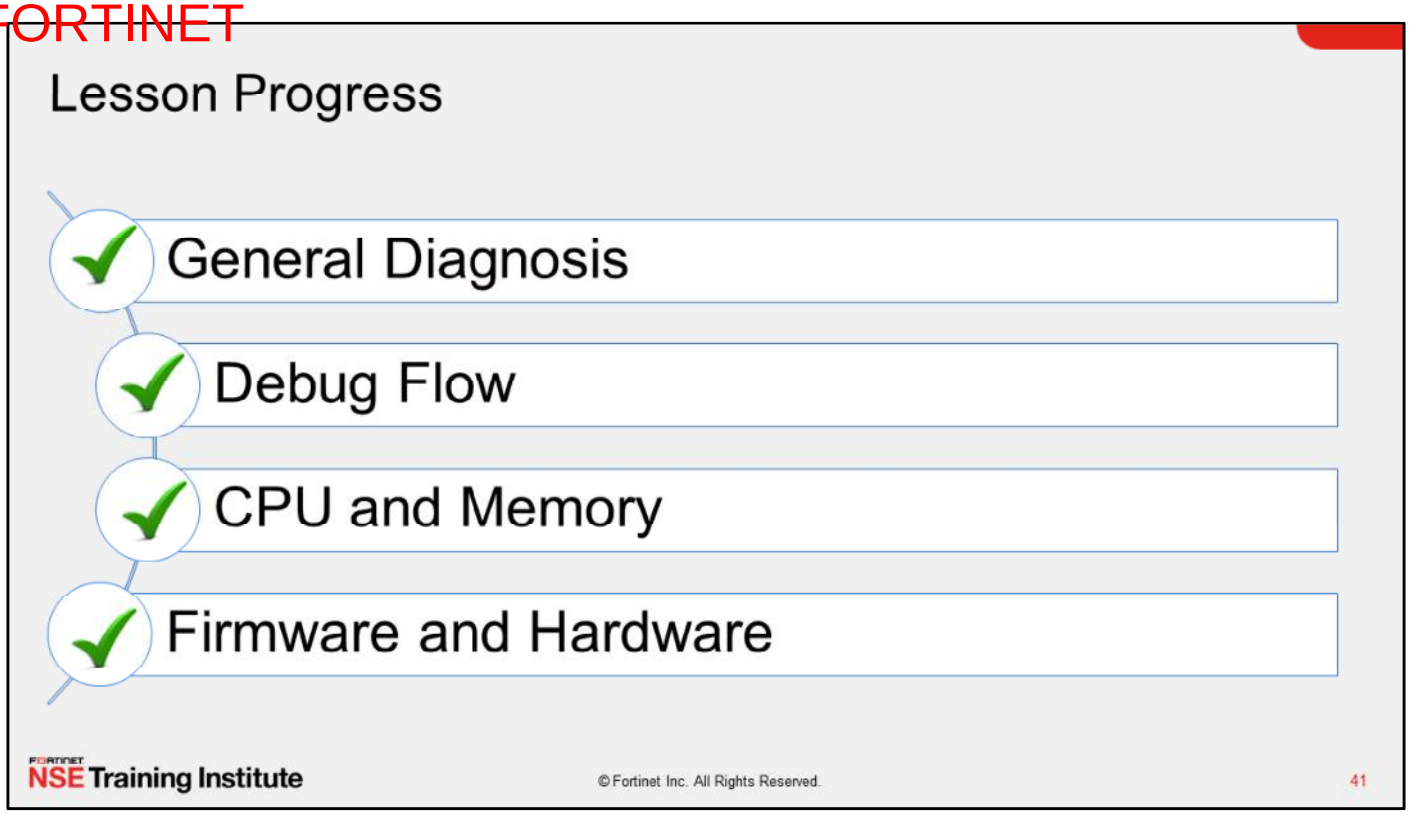

Congratulations! You have completed the lesson.

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

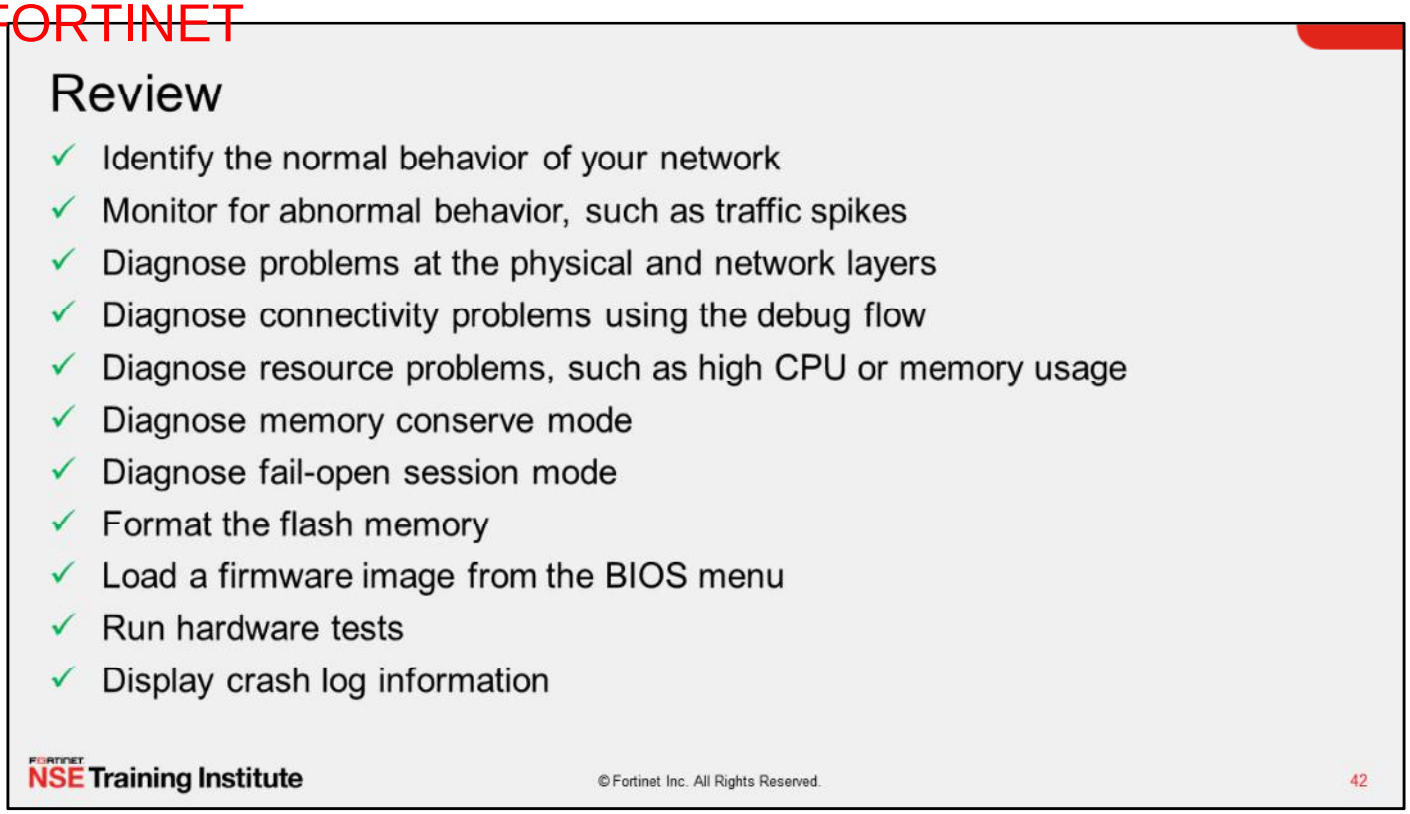

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use diagnostic commands and tools.