FortiGate Security Study Guide

for FortiOS 7.2

FORTINET®
Training Institute

Fortinet Training

https://training.fortinet.com

Fortinet Document Library

https://docs.fortinet.com

Fortinet Knowledge Base

https://kb.fortinet.com

Fortinet Fuse User Community

https://fusecommunity.fortinet.com/home

Fortinet Forums

https://forum.fortinet.com

Fortinet Support

https://support.fortinet.com

FortiGuard Labs

https://www.fortiguard.com

Fortinet Network Security Expert Program (NSE)

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

Fortinet | Pearson VUE

https://home.pearsonvue.com/fortinet

Feedback

Email: askcourseware@fortinet.com

6/13/2022

TABLE OF CONTENTS

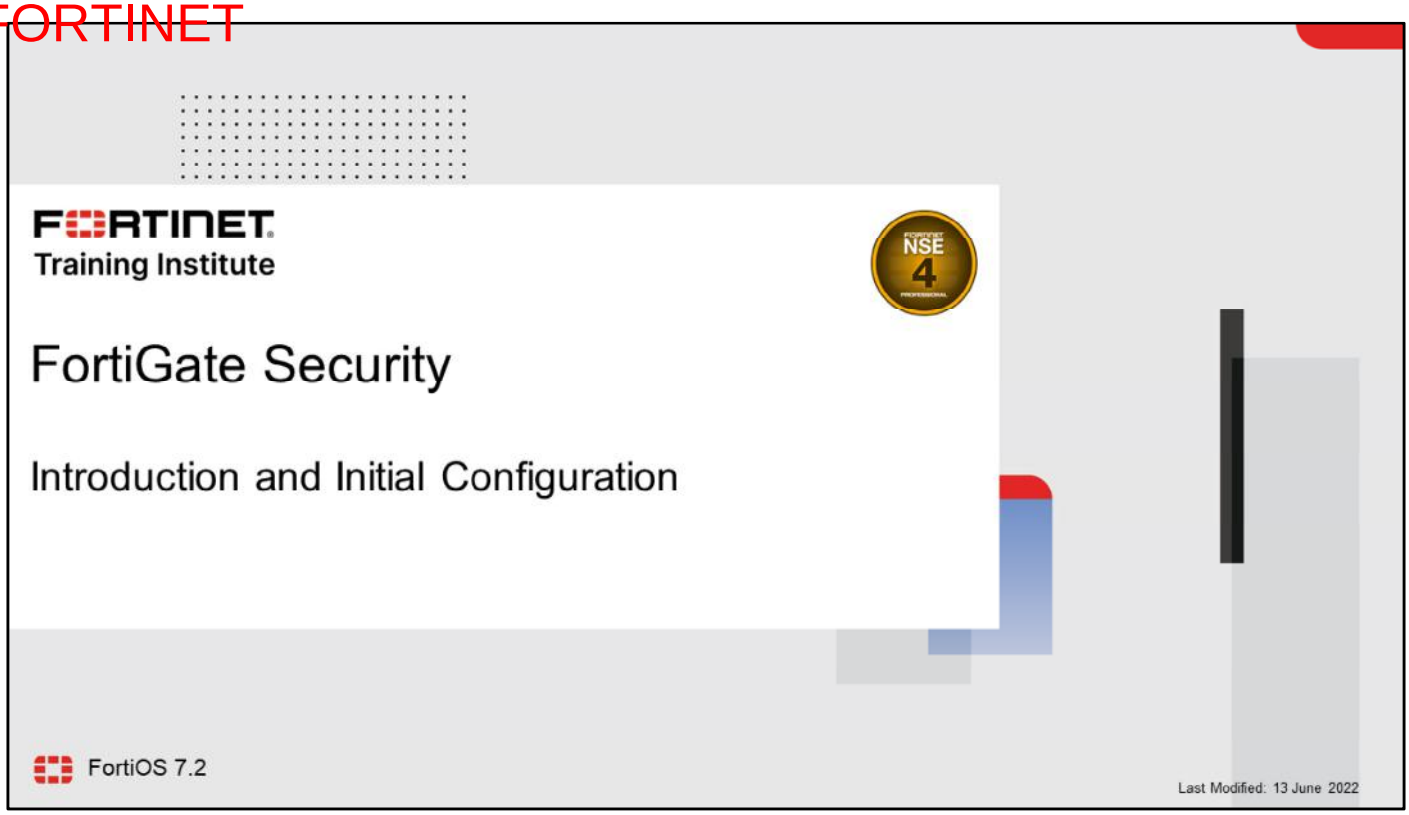

In this lesson, you will learn about FortiGate administration basics and the components within FortiGate that you can enable to extend functionality. This lesson also includes details about how and where FortiGate fits into your existing network architecture.

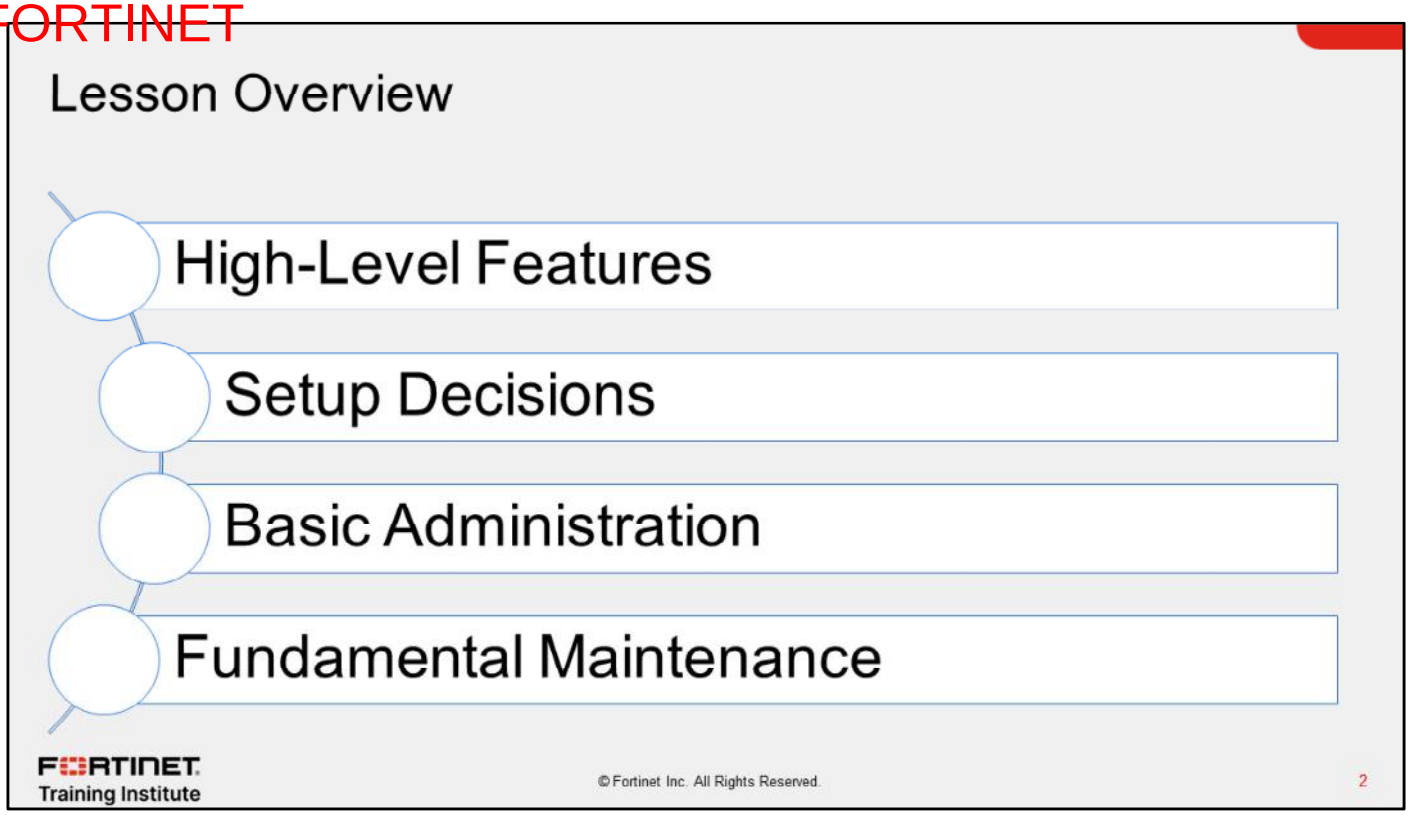

In this lesson, you will explore the topics shown on this slide.

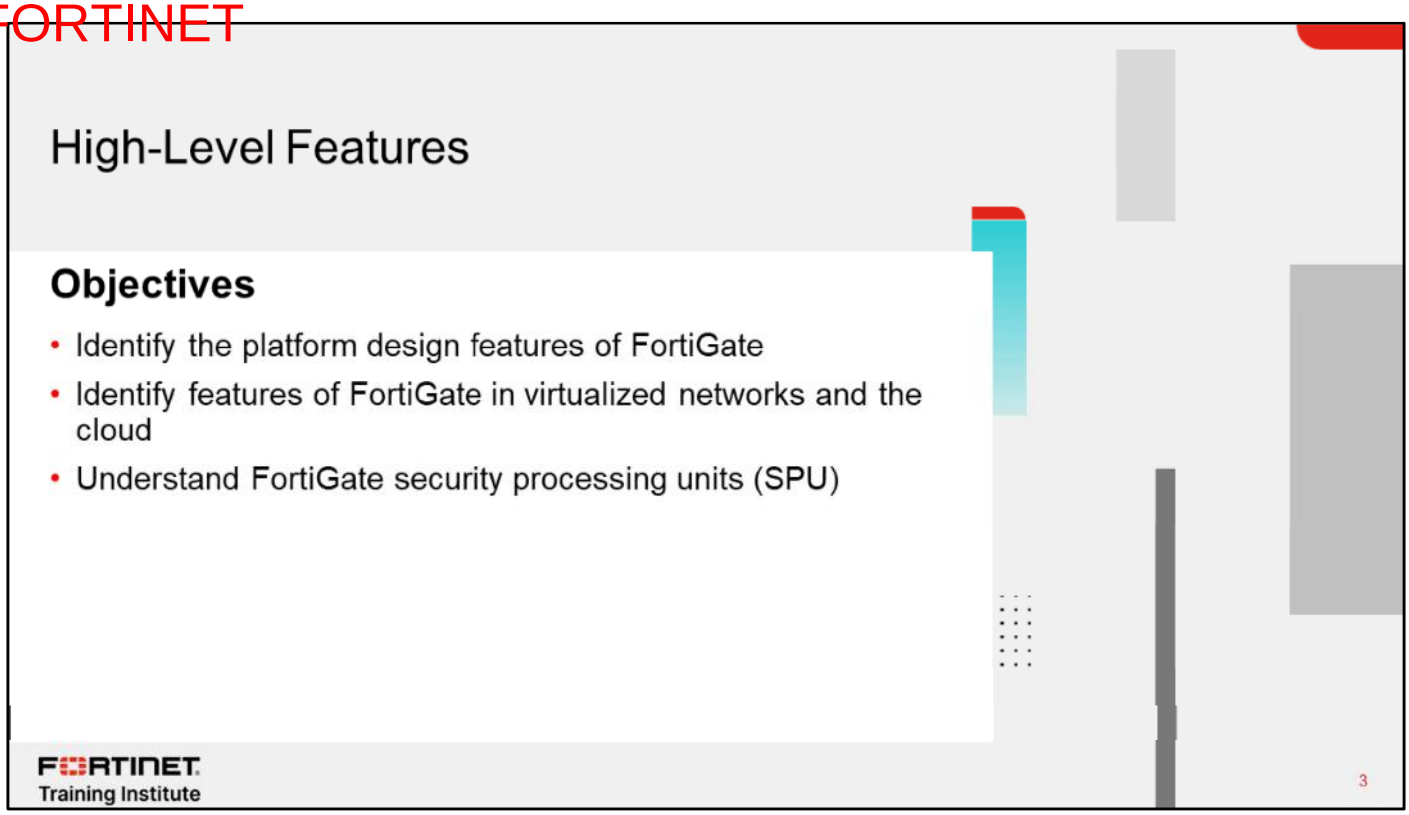

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in identifying the platform design features of FortiGate, FortiGate features in virtualized networks and the cloud, as well as the FortiGate security processing units, you will be able to describe the fundamental components of FortiGate and explain the types of tasks that FortiGate can perform.

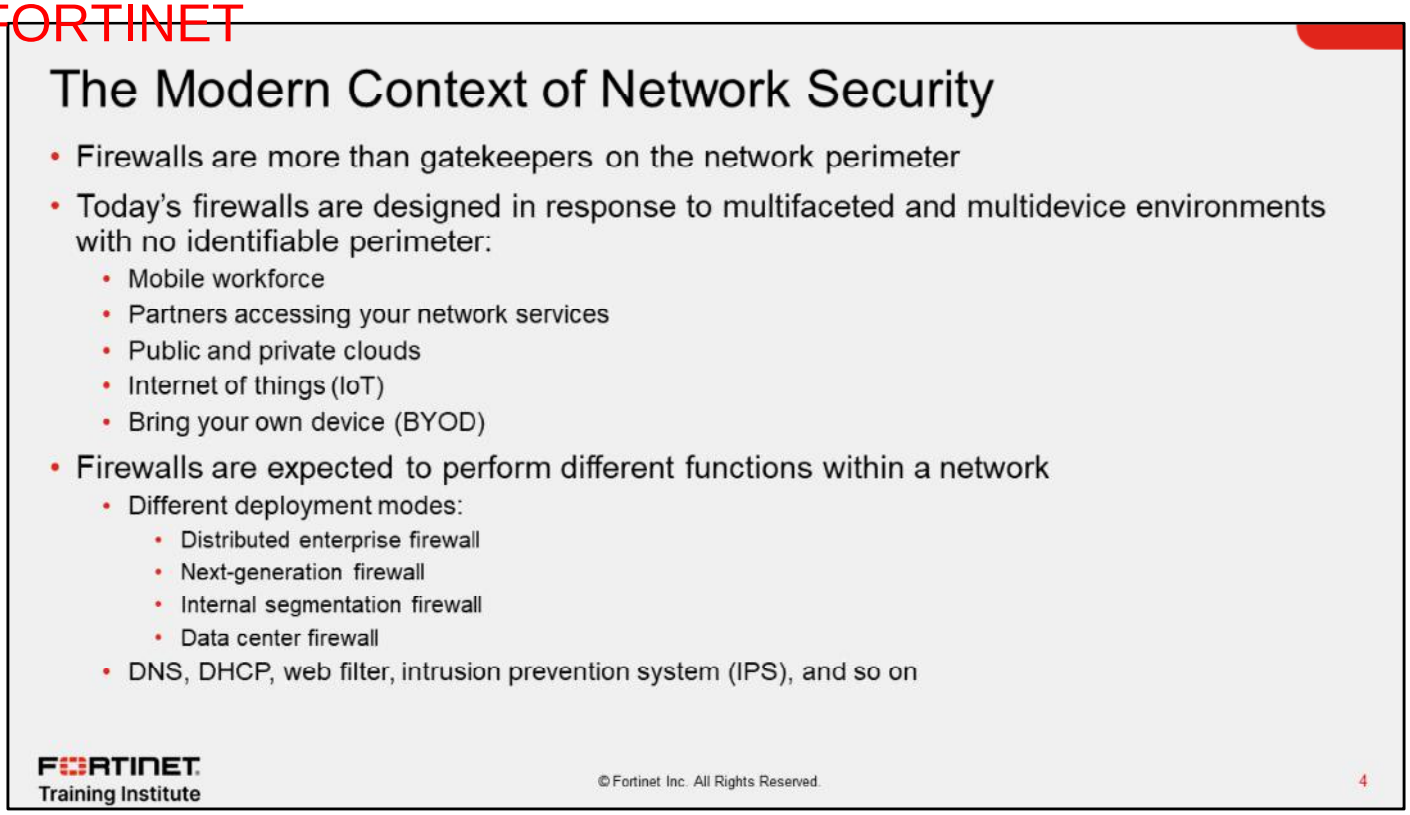

In the past, the common way of protecting a network was securing the perimeter and installing a firewall at the entry point. Network administrators used to trust everything and everyone inside the perimeter.

Now, malware can easily bypass any entry-point firewall and get inside the network. This could happen through an infected USB stick, or an employee's compromised personal device being connected to the corporate network. Additionally, because attacks can come from inside the network, network administrators can no longer inherently trust internal users and devices.

What's more, today's networks are highly complex environments whose borders are constantly changing. Networks run vertically from the LAN to the internet, and horizontally from the physical network to a private virtual network and to the cloud. A mobile and diverse workforce (employees, partners, and customers) accessing network resources, public and private clouds, the IoT, and BYOD programs all conspire to increase the number of attack vectors against your network.

In response to this highly complex environment, firewalls have become robust multifunctional devices that counter an array of threats to your network. Thus, FortiGate can act in different modes or roles to address different requirements. For example, FortiGate can be deployed as a data center firewall whose function is to monitor inbound requests to servers and to protect them without increasing latency for the requester. Or, FortiGate can be deployed as an internal segmentation firewall as a means to contain a network breach.

FortiGate can also function as DNS and DHCP servers, and be configured to provide web filter, antivirus, and IPS services.

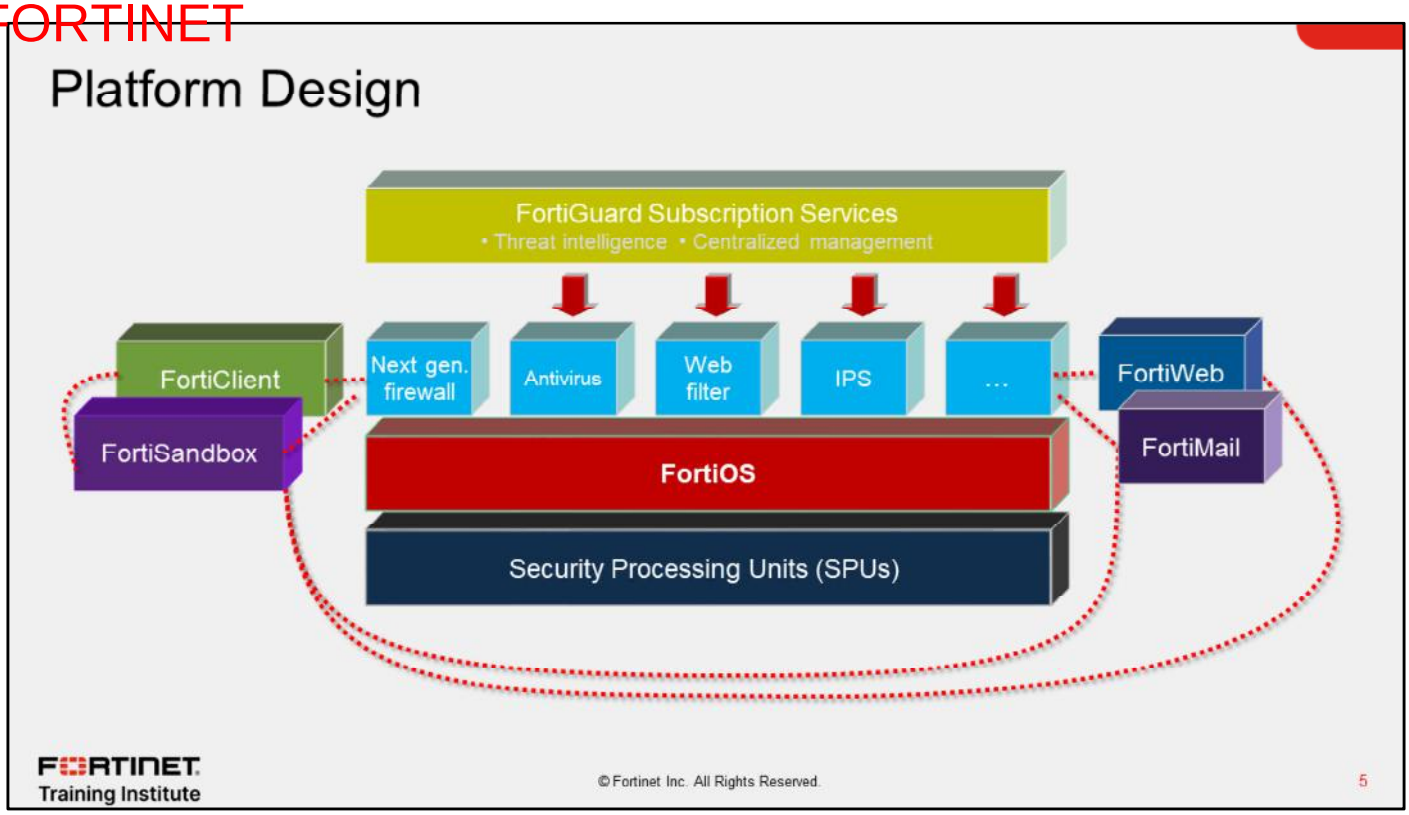

In the architecture diagram shown on this slide, you can see how FortiGate platforms add strength, without compromising flexibility. Like separate, dedicated security devices, FortiGate is still *internally* modular. Plus:

- **Devices add duplication.** Sometimes, dedication *doesn't* mean efficiency. If it's overloaded, can one device borrow free RAM from nine others? Do you want to configure policies, logging, and routing on 10 separate devices? Does 10 times the duplication bring you 10 times the benefit, or is it a hassle? For smaller to midsize businesses or enterprise branch offices, unified threat management (UTM) is often a superior solution, compared to separate dedicated appliances.
- **FortiGate hardware isn't just off-the-shelf.** It's carrier-grade. Most FortiGate models have one or more specialized circuits, called ASICs, that are engineered by Fortinet. For example, a CP or NP chip handles cryptography and packet forwarding more efficiently. Compared to a single-purpose device with only a CPU, FortiGate can have dramatically better performance. This is especially critical for data centers and carriers where throughput is business critical.
  (The exception? Virtualization platforms—VMware, Citrix Xen, Microsoft, or Oracle Virtual Box—have general-purpose vCPUs. But, virtualization might be worthwhile because of other benefits, such as distributed computing and cloud-based security.)
- **FortiGate is flexible.** If all you need is fast firewalling and antivirus, FortiGate won't require you to waste CPU, RAM, and electricity on other features. In each firewall policy, you can enable or disable UTM and next-generation firewall modules. Also, you won't pay more to add VPN seat licenses later.
- **FortiGate cooperates.** A preference for open standards instead of proprietary protocols means less vendor lock-in and more choice for system integrators. And, as your network grows, FortiGate can leverage other Fortinet products, such as FortiSandbox and FortiWeb, to distribute processing for deeper security and optimal performance—a total Security Fabric approach.

Topology in the Cloud

- **Deploy FortiGate in virtualized networks**
  - FortiGate VM – Same features as physical appliance *except* SPUs
- VMs or physical appliances
  - Configuration is essentially the same

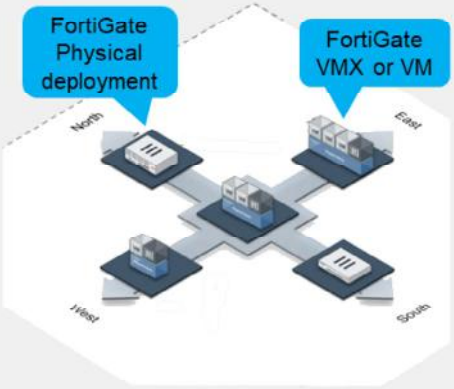

FortiGate VM Specifications	
Licenses	Max. 1 / 2 / 4 / 8 vCPU
Hypervisor	VMware, Hyper-V, KVM, Citrix Xen Server, Open Source Xen, Azure, Amazon AWS BYOL & on-demand
Memory	Max. 1/4/8/12 GB
10/100/1000 Interfaces	2-4 virtual NICs
Storage Capacity	40+ GB

FORTINET Training Institute

6

FortiGate VMs have the same features as physical FortiGate devices**,** ***except*** for hardware acceleration**.** Why? First, the hardware abstraction layer software for hypervisors is made by VMware, Xen, and other hypervisor manufacturers, *not* by Fortinet. Those other manufacturers don't make the Fortinet proprietary SPU chips. But there is another reason, too. The purpose of generic virtual CPUs and other virtual chips for hypervisors is to abstract the hardware details. That way, all VM guest OSs can run on a common platform, no matter the different hardware on which the hypervisors are installed. Unlike vCPUs or vGPUs that use generic, *non-optimal* RAM and vCPUs for abstraction, SPU chips are specialized *optimized* circuits. Therefore, a virtualized ASIC chip would not have the same performance benefits as a physical SPU chip.

If performance on equivalent hardware is less, you may wonder why anyone would use a FortiGate VM. In large-scale networks that change rapidly and may have many tenants, equivalent processing power and distribution may be achievable using larger amounts of cheaper, general purpose hardware. Also, trading some performance for other benefits may be worth it. You can benefit from faster network and appliance deployment and teardown.

Either VMs or physical appliances (low or high-end models), the configuration of the security instances is essentially identical, using same FortiOS version and FortiGuard real-time threat intelligence.

Knowledge Check

1. Which is a more accurate description of a modern firewall?
   - A. A device that inspects network traffic at an entry point to the internet and within a simple, easily defined network perimeter
   - ✓ B. A multifunctional device that inspects network traffic from the perimeter or internally, within a network that has many different entry points

2. Which solution specific to Fortinet enhances performance and reduces latency for specific features and traffic?
   - ✓ A. Acceleration hardware, called SPUs
   - B. Increased RAM and CPU power

FORTINET Training Institute

7

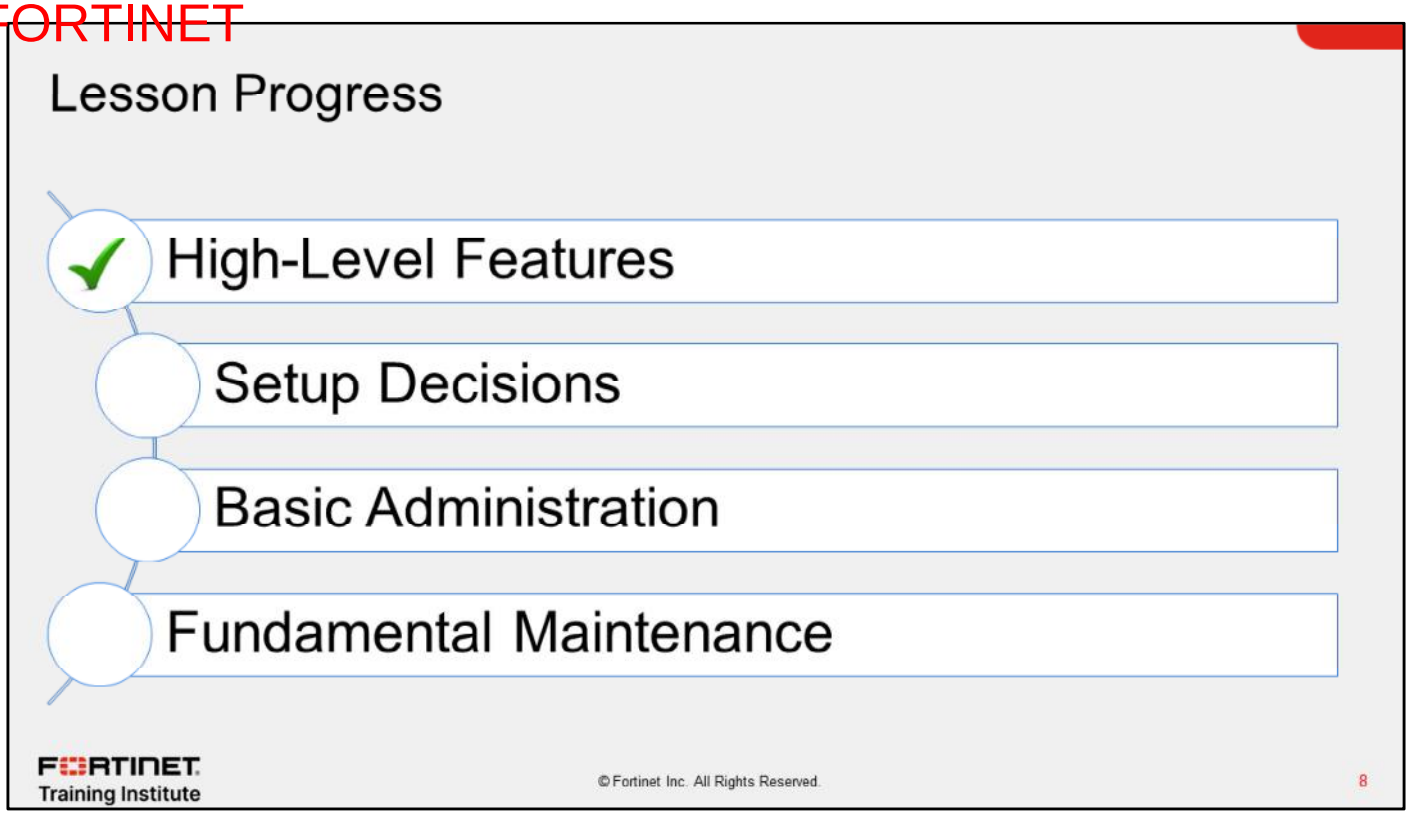

Good job! You now understand some of the high-level features of FortiGate.

Now, you will learn how to perform the initial setup of FortiGate and learn about why you might decide to use one configuration over another.

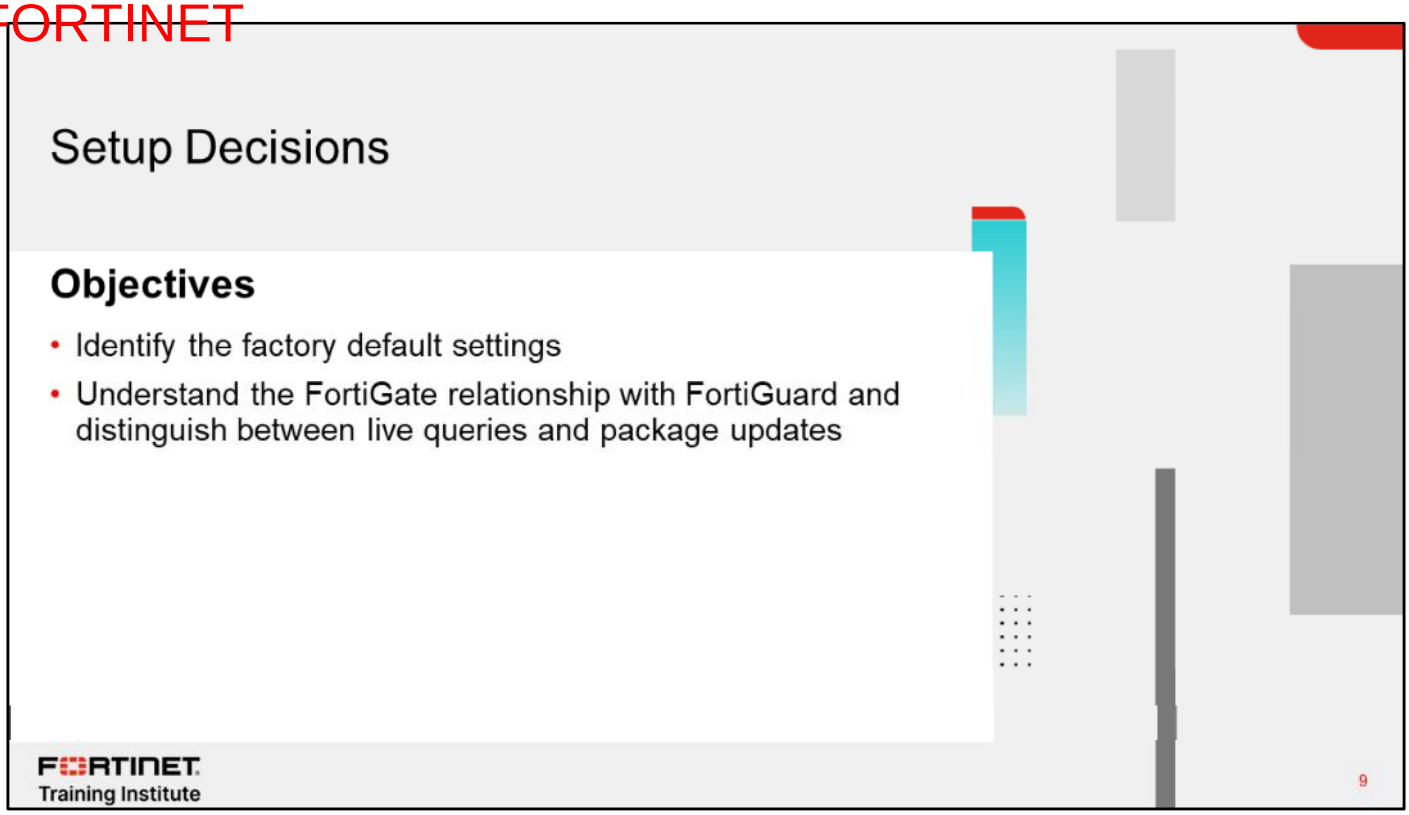

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in setting up FortiGate, you will be able to use the device effectively in your own network.

Factory Default Settings

- IP: `192.168.1.99/24`
  - MGMT interface on high-end and mid-range models
  - Port1 or internal interface on entry-level models
- PING, HTTPS, and SSH protocol management enabled
- Built-in DHCP server is enabled on port1 or internal interface
  - Only on entry-level models that support DHCP server
- Default login:
  User: `admin`
  Password: (blank)
  - Both are case sensitive
  - Modify the default (blank) password
- Can access FortiGate on the CLI
  - Console: without network
  - CLI console widget and terminal emulator, such as PuTTY or Tera Term

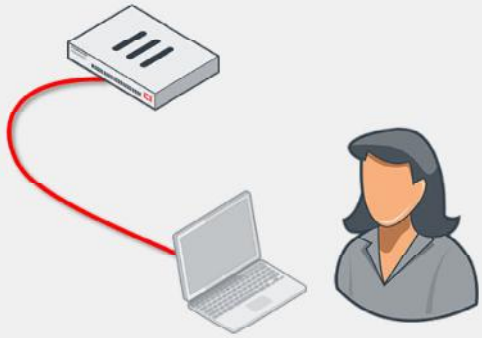

FORTINET Training Institute

10

Network address translation (NAT) mode is the default operation mode. What are the other factory default settings? After you've removed FortiGate from its box, what do you do next?

Now you'll take a look at how you set up FortiGate.

Attach your computer network cable to port1 or the internal switch ports (on the entry-level model). For high-end and mid-range models, connect to the MGMT interface. In most entry-level models, there is a DHCP server on that interface, so, if your computer's network settings have DHCP enabled, your computer should automatically get an IP, and you can begin setup.

To access the GUI on FortiGate or FortiWifi, open a web browser and visit `https://192.168.1.99`.

The default login information is public knowledge. Never leave the default password blank. Your network is only as secure as your FortiGate `admin` account. Once you logged in with default login details, you'll see a message to change the default blank password for the admin user password. Before you connect FortiGate to your network, you should set a complex password. You'll also be asked to apply additional configuration such as hostname, dashboard setup, register with FortiCare, and so on.

All FortiGate models have a console port and/or USB management port. The port provides CLI access without a network. You can access the CLI using the CLI console widget on the GUI, or from a terminal emulator, such as PuTTY or Tera Term.

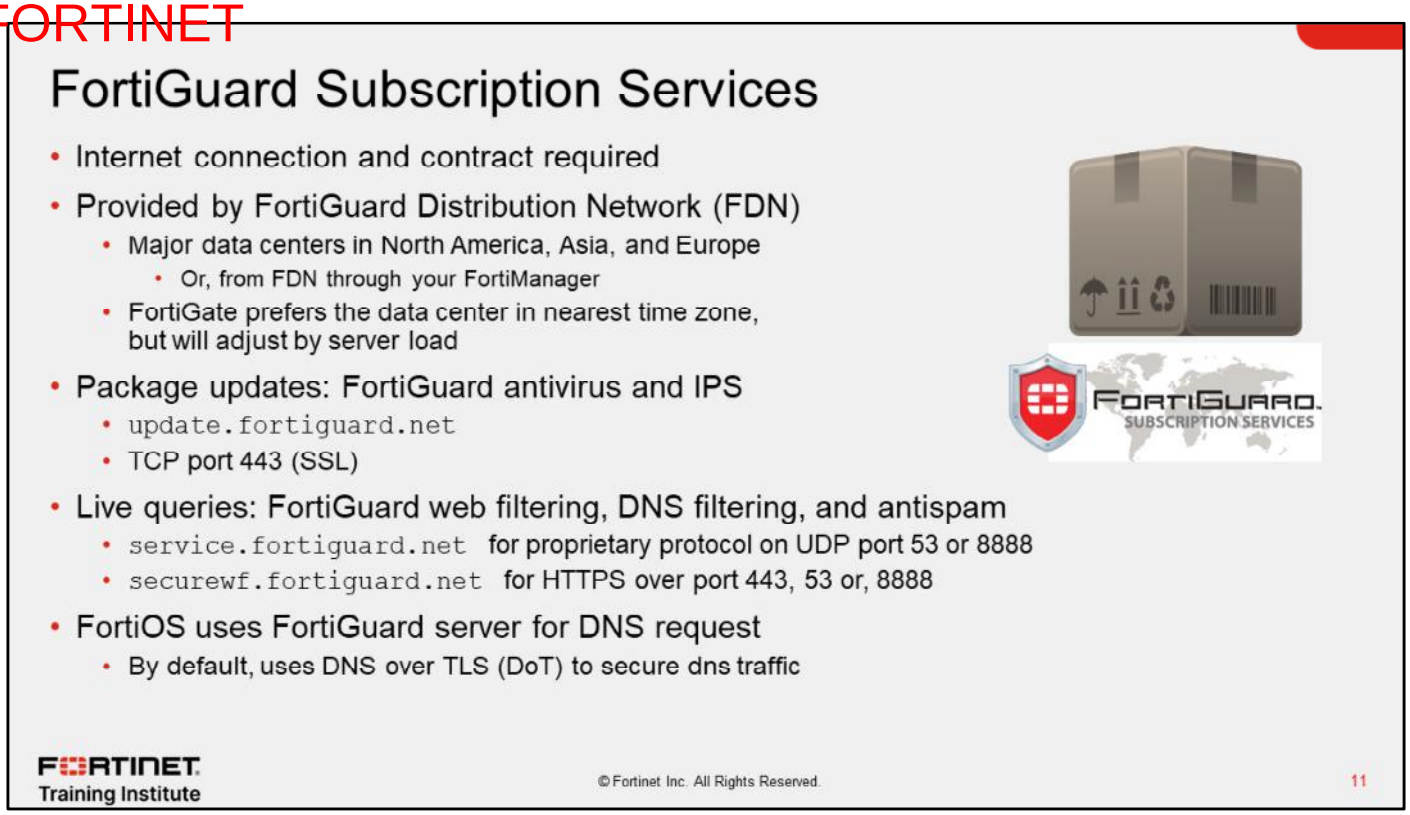

Some FortiGate services connect to other servers, such as FortiGuard, in order to work. FortiGuard Subscription Services provide FortiGate with up-to-date threat intelligence. FortiGate uses FortiGuard by:

- Periodically requesting packages that contain a new engine and signatures
- Querying the FDN on an individual URL or host name

By default, the FortiGuard server location is set to anywhere FortiGate selects a server based on server load, from any part of the world. However, you have the option to change the FortiGuard server location to USA. In this case, FortiGate selects a USA-based FortiGuard server.

Queries are real-time; that is, FortiGate asks the FDN every time it scans for spam or filtered websites. FortiGate queries, instead of downloading the database, because of the size and frequency of changes that occur to the database. Also, you can select queries to use UDP or HTTPs for transport; the protocols are not designed for fault tolerance, but for speed. So, queries require that your FortiGate device has a reliable internet connection.

Packages, like antivirus and IPS, are smaller and don't change as frequently, so they are downloaded (in many cases) only once a day. They are downloaded using TCP for reliable transport. After the database is downloaded, their associated FortiGate features continue to function, even if FortiGate does not have reliable internet connectivity. However, you should still try to avoid interruptions during downloads—if your FortiGate device must try repeatedly to download updates, it can’t detect new threats during that time.

When using FortiGuard servers for DNS, FortiOS uses DNS over TLS (DoT) by default to secure the DNS traffic. New FortiGuard DNS servers have been added as primary and secondary servers.

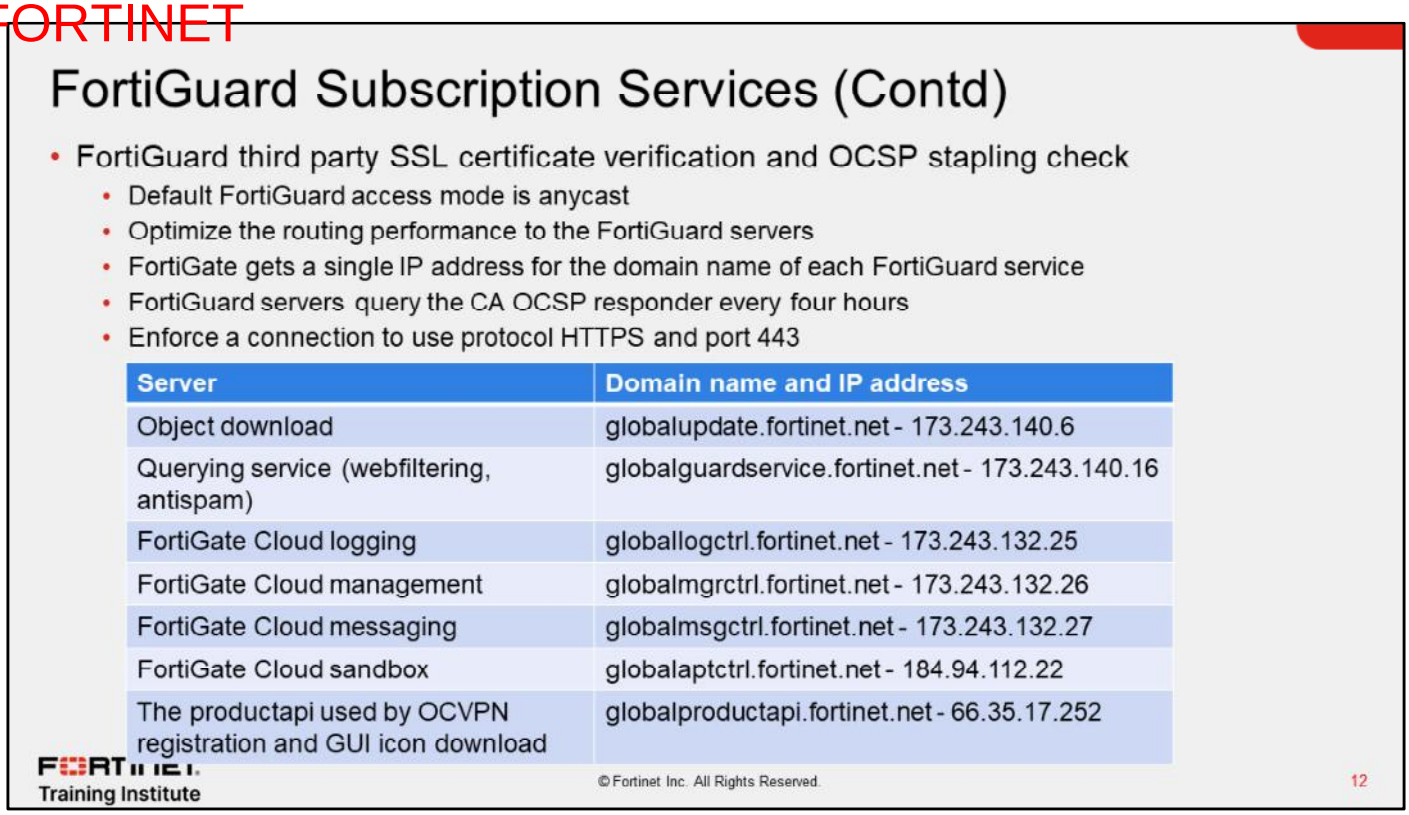

Now, third-party SSL certificate verification and OCSP stapling check has been implemented for all FortiGuard servers. By default, the FortiGuard access mode is *anycast* on FortiGate, to optimize the routing performance to the FortiGuard servers. The FortiGuard server has one IP address to match its domain name. FortiGate connects with a single server address, regardless of where the FortiGate device is located.

The domain name of each FortiGuard service is the common name in the certificate of that service. The certificate is signed by a third-party intermediate CA. The FortiGuard server uses the Online Certificate Status Protocol (OCSP) stapling technique, so that FortiGate can always validate the FortiGuard server certificate efficiently. FortiGate will complete the TLS handshake only with a FortiGuard server that provides a *good* OCSP status for its certificate. Any other status results in a failed SSL connection.

The FortiGuard servers query the OCSP responder of the CA every four hours and update its OCSP status. If FortiGuard is unable to reach the OCSP responder, it keeps the last known OCSP status for seven days.

FortiGate aborts the connection to the FortiGuard server if:
- The CN in the server certificate does not match the domain name resolved from the DNS.
- The OCSP status is not good.
- The issuer-CA is revoked by the root-CA.

The FortiGuard access mode *anycast* setting forces the rating process to use protocol HTTPS, and port 443. The table on this slide shows a list of some of the FortiGuard servers and their domain names and IP addresses.

Knowledge Check

1. Which protocol does FortiGate use to download antivirus and IPS packages?
   - A. UDP
   - ✓ B. TCP

2. How does FortiGate check content for spam or malicious websites?
   - ✓ A. Live queries to FortiGuard over UDP or HTTPS
   - B. Local verification using a downloaded web filter database locally on FortiGate

FORTINET Training Institute

13

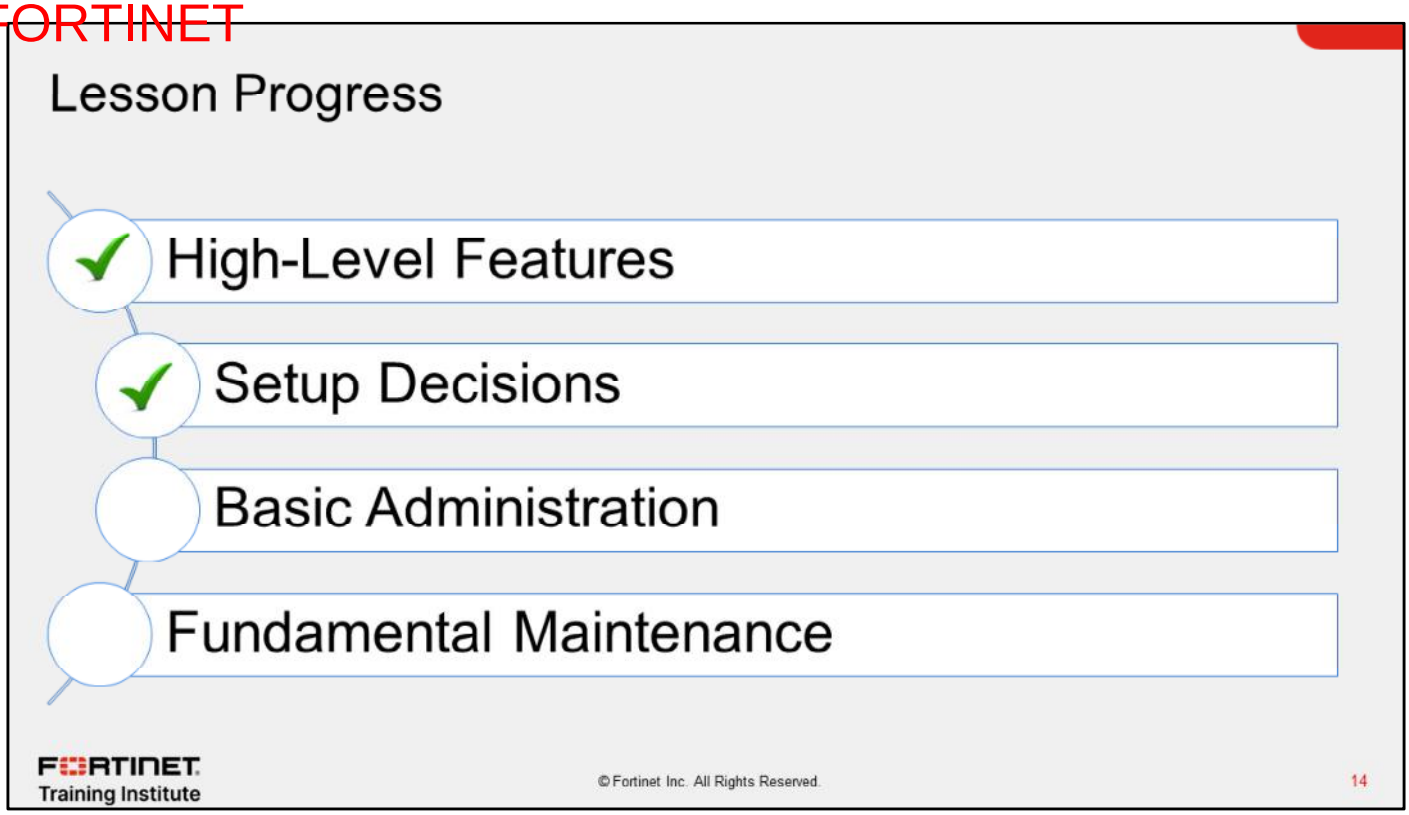

Good job! You now understand how to perform the initial setup of FortiGate and why you might decide to use one configuration over another. Now, you will learn about basic administration.

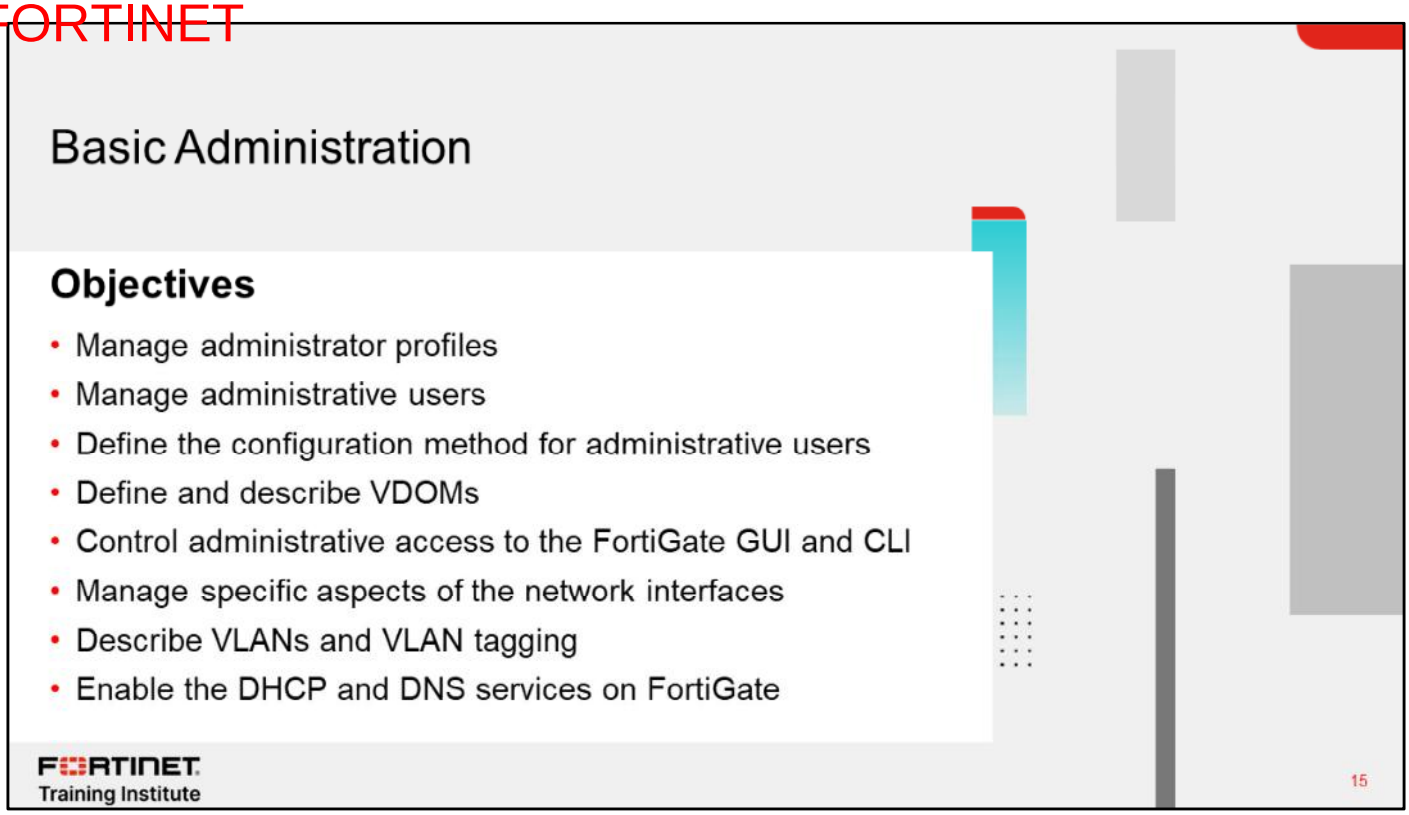

After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in basic administration, you will be able to better manage administrative users and implement stronger security practices around administrative access.

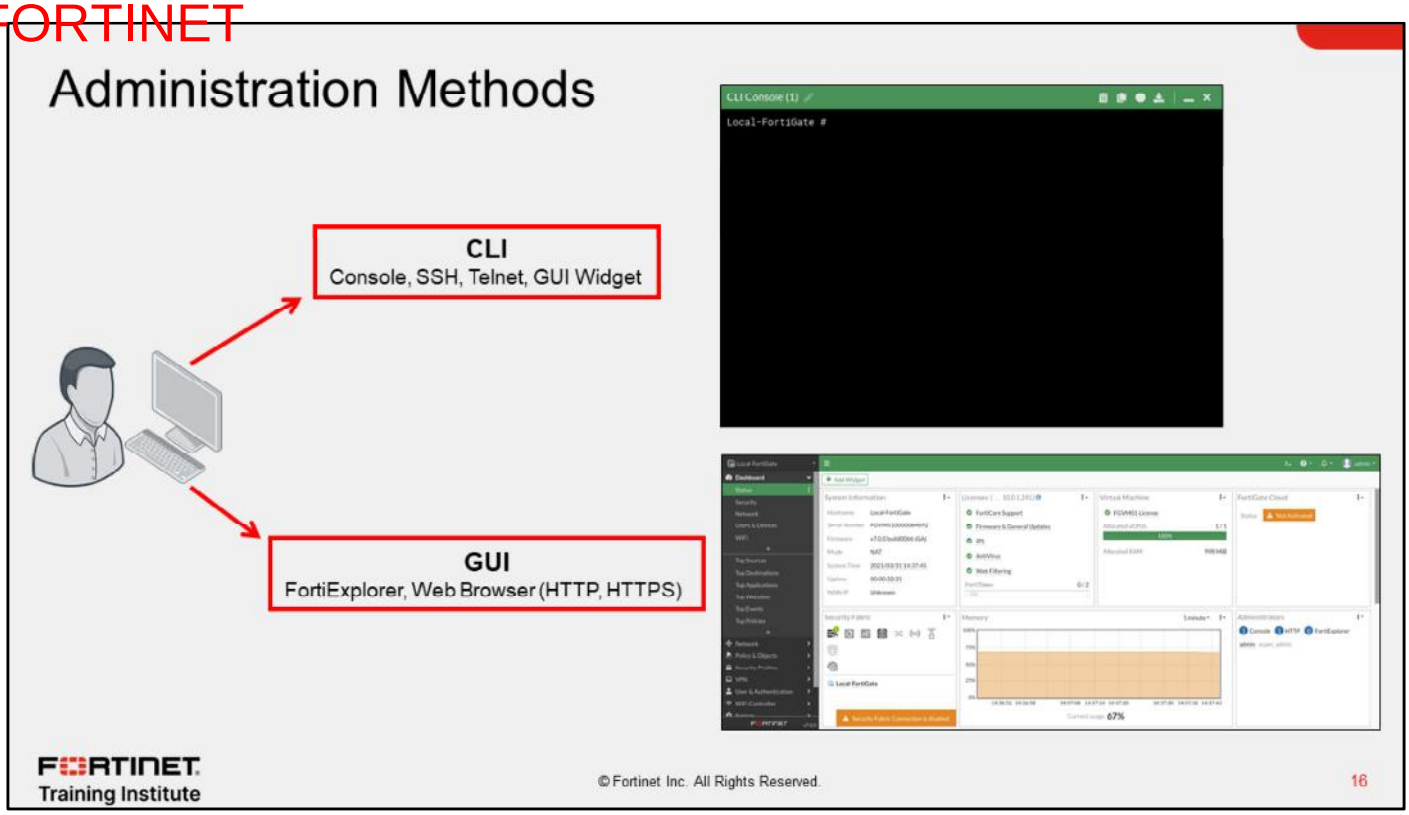

Most features are available on both the GUI and CLI, but there are a few exceptions. You can't view reports on the CLI. Also, advanced settings and diagnostic commands for super users are usually not available on the GUI.

As you become more familiar with FortiGate, and especially if you want to script its configuration, you might want to use the CLI in addition to the GUI. You can access the CLI through either the JavaScript widget on the GUI named **CLI Console**, or through a terminal emulator such as Tera Term (http://ttssh2.sourceforge.jp/index.html.en) or PuTTY (http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html). Your terminal emulator can connect through the network—SSH or telnet—or the local console port.

SNMP and some other administrative protocols are also supported, but they are read-only. You can't use them for basic setup.

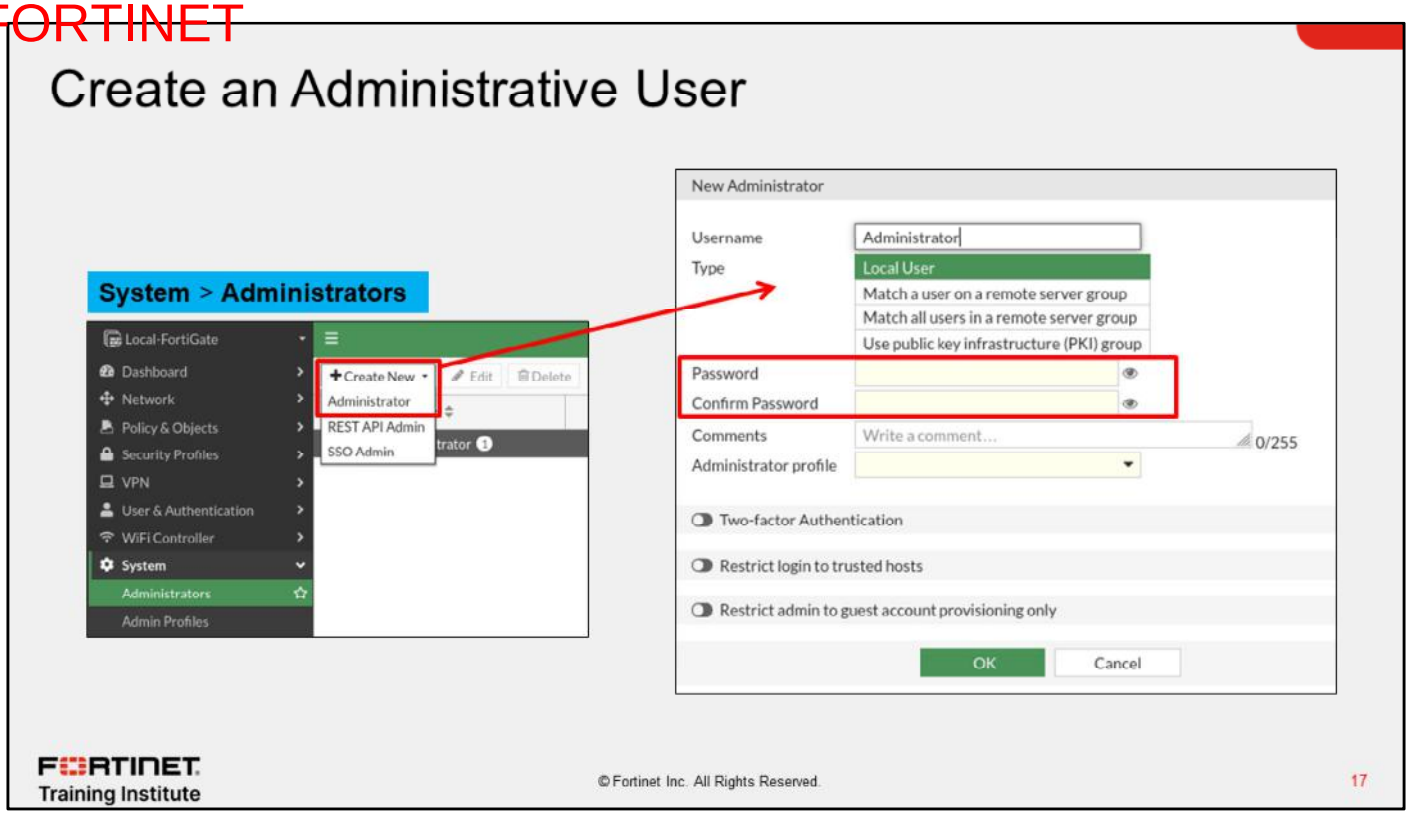

Whichever method you use, start by logging in as admin. Begin by creating separate accounts for other administrators. For security and tracking purposes, it is a best practice for each administrator to have their own account.

In the **Create New** drop-down list, you can select either **Administrator** or **REST API Admin**. Typically, you will select **Administrator** and then assign an **Administrator Profile**, which specifies that user's administrative permissions. You could select **REST API Admin** to add an administrative user who would use a custom application to access FortiGate with a REST API. The application would allow you to log in to FortiGate and perform any task that your assigned **Administrator Profile** permits.

Other options not shown here, include:

- Instead of creating accounts on FortiGate itself, you could configure FortiGate to query a remote authentication server.
- In place of passwords, your administrators could authenticate using digital certificates that are issued by your internal certification authority server.

If you do use passwords, ensure that they are strong and complex. For example, you could use multiple interleaved words with varying capitalization, and randomly insert numbers and punctuation. Do not use short passwords, or passwords that contain names, dates, or words that exist in any dictionary. These are susceptible to brute force attack. To audit the strength of your passwords, use tools such as L0phtcrack (http://www.l0phtcrack.com/) or John the Ripper (http://www.openwall.com/john/). Risk of a brute force attack is increased if you connect the management port to the internet.

In order to restrict access to specific features, you can assign permissions.

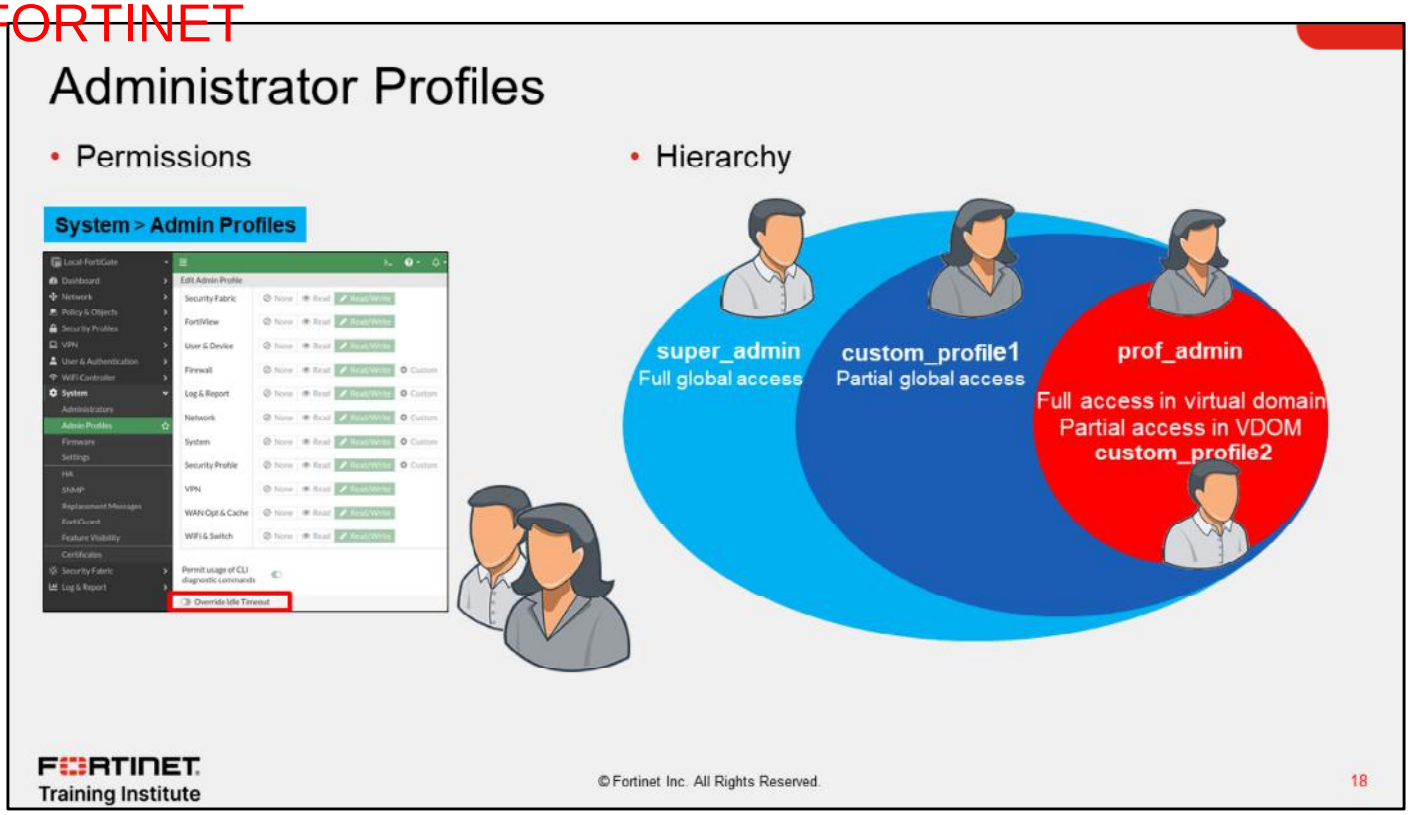

When assigning permissions to an administrator profile, you can specify read-and-write, read-only, or none to each area.

By default, there is a special profile named **super_admin**, which is used by the account named **admin**. You can't change it. It provides full access to everything, making the *admin* account similar to a root **superuser** account.The **prof_admin** is another default profile. It also provides full access, but unlike **super_admin**, it applies only to its virtual domain—not the global settings of FortiGate. Also, you can change its permissions.

You aren't required to use a default profile. You could, for example, create a profile named **auditor_access** with read-only permissions. Restricting a person's permissions to those necessary for his or her job is a best practice, because even if that account is compromised, the compromise to your FortiGate device (or network) is not total. To do this, create administrator profiles, then select the appropriate profile when configuring an account.

The **Override Idle Timeout** feature allows the admintimeout value, under `config system accprofile`, to be overridden per access profile. You can configure administrator profiles to increase inactivity timeout and facilitate use of the GUI for central monitoring. Note that you can do this on a per-profile basis, to prevent the option from being unintentionally set globally. So, what are the effects of administrator profiles?

It's actually more than just read or write access. Depending on the type of administrator profile that you assign, an administrator may not be able to access the entire FortiGate device. For example, you could configure an account that can view only log messages. Administrators may not be able to access global settings outside their assigned virtual domain either. Virtual domains (VDOMs) are a way of subdividing the resources and configurations on a single FortiGate. Administrators with a smaller scope of permissions cannot create, or even view, accounts with more permissions.

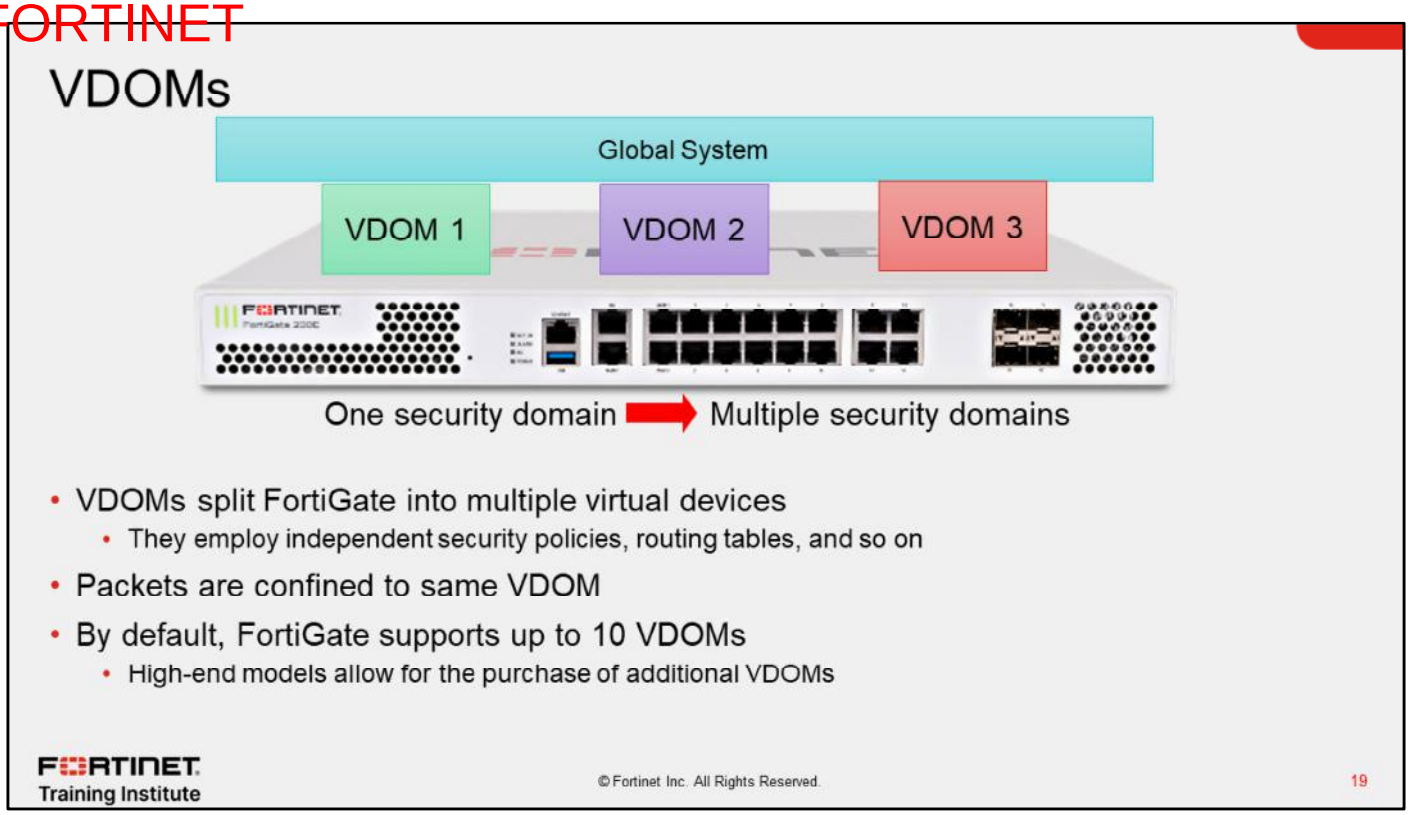

What if, more than segmenting your network, you want to subdivide policies and administrators into multiple security domains?

In that case, you can enable FortiGate VDOMs, which split your FortiGate into multiple logical devices. Each VDOM has independent security policies and routing tables. Also, and by default, traffic from one VDOM cannot go to a different VDOM. This means that two interfaces in different VDOMs can share the same IP address, without any overlapping subnet problems.

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

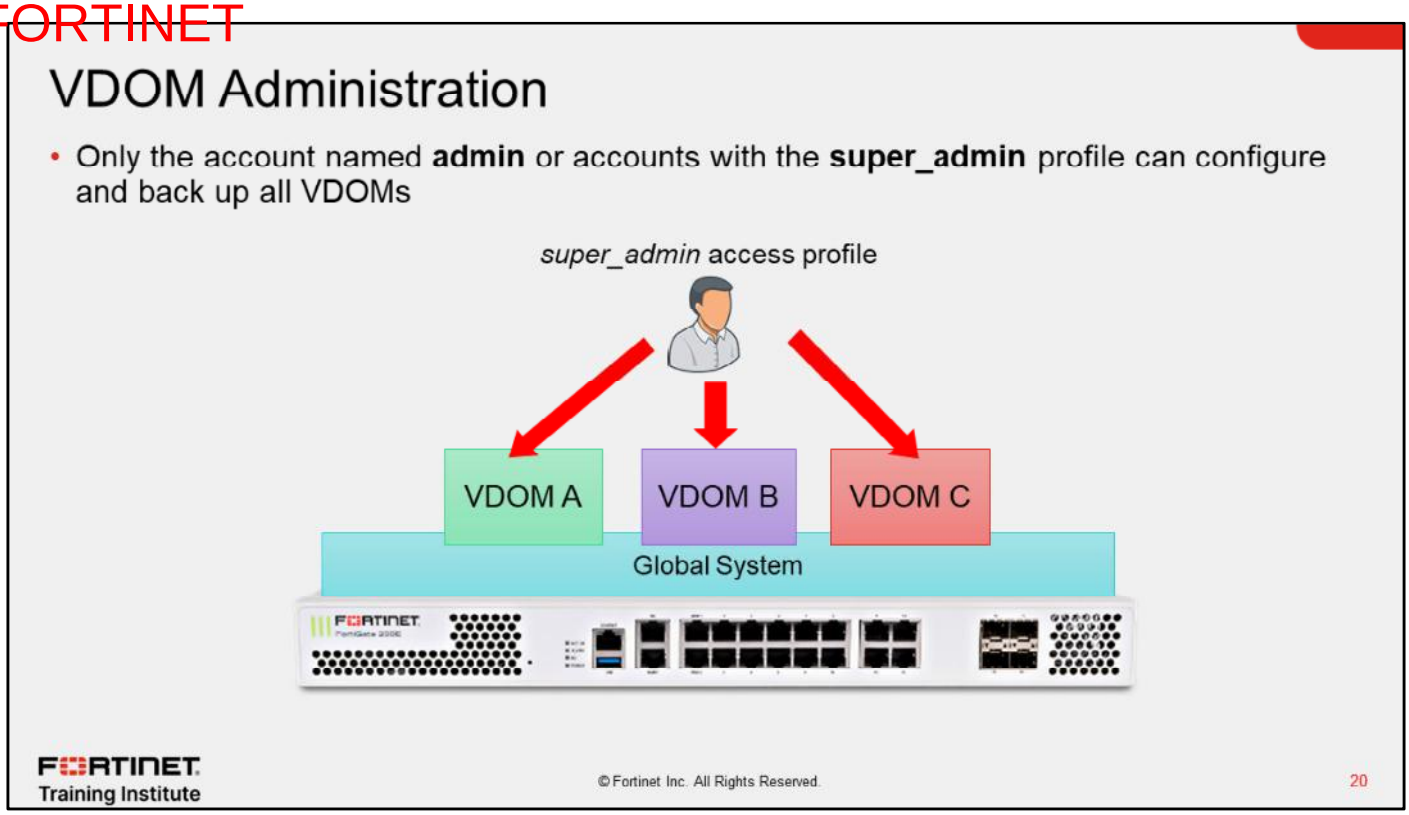

If you want to grant access to all VDOMs and global settings, select **super_admin** as the access profile when configuring the administrator account. Similar to the account named **admin**, this account can configure all VDOMs.

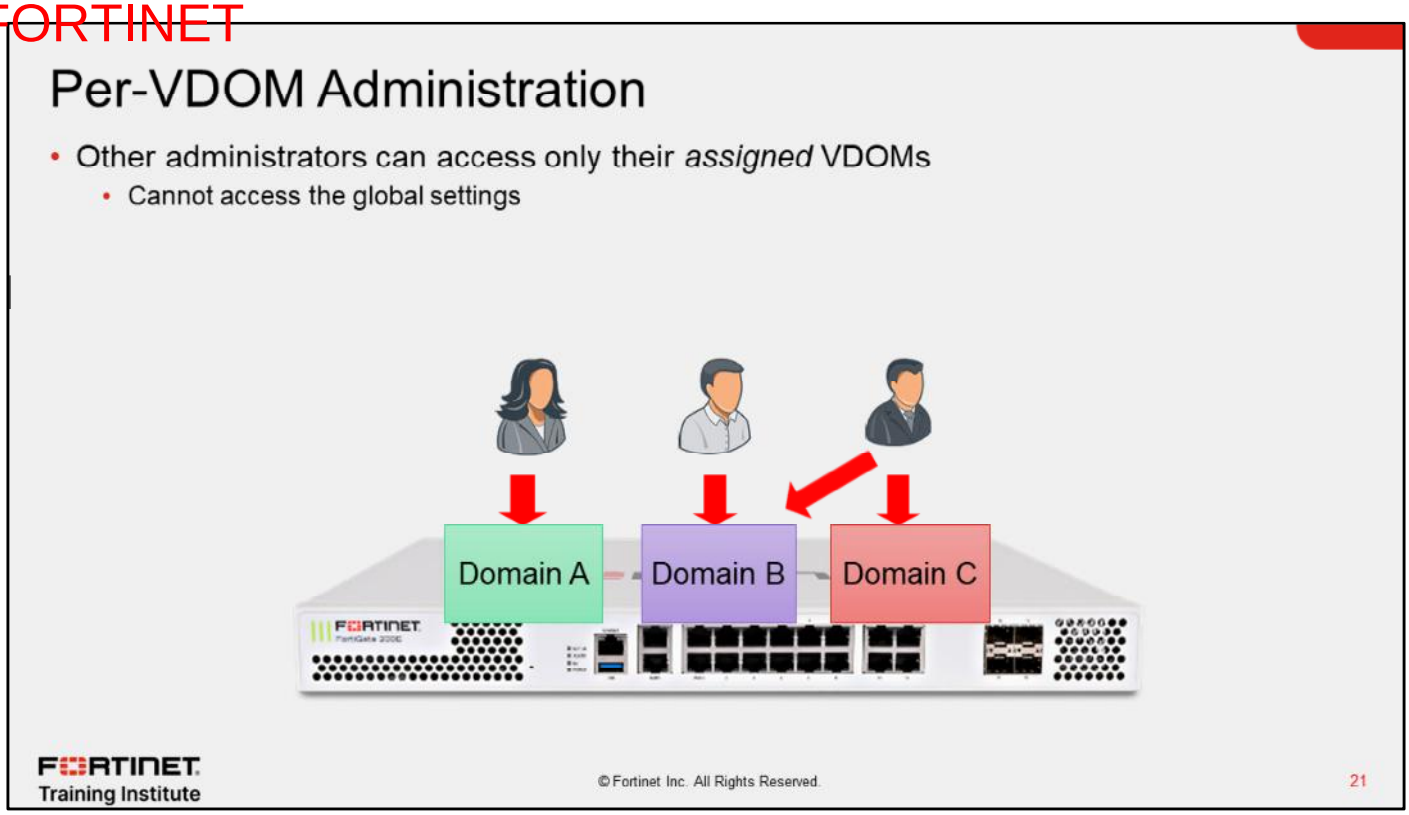

In most cases, you start by creating one administrator account per VDOM. That administrator is chiefly responsible for that domain, including the configuration backups of that VDOM. In larger organizations, you may need to make multiple VDOM administrators. You can assign multiple administrators to each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

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

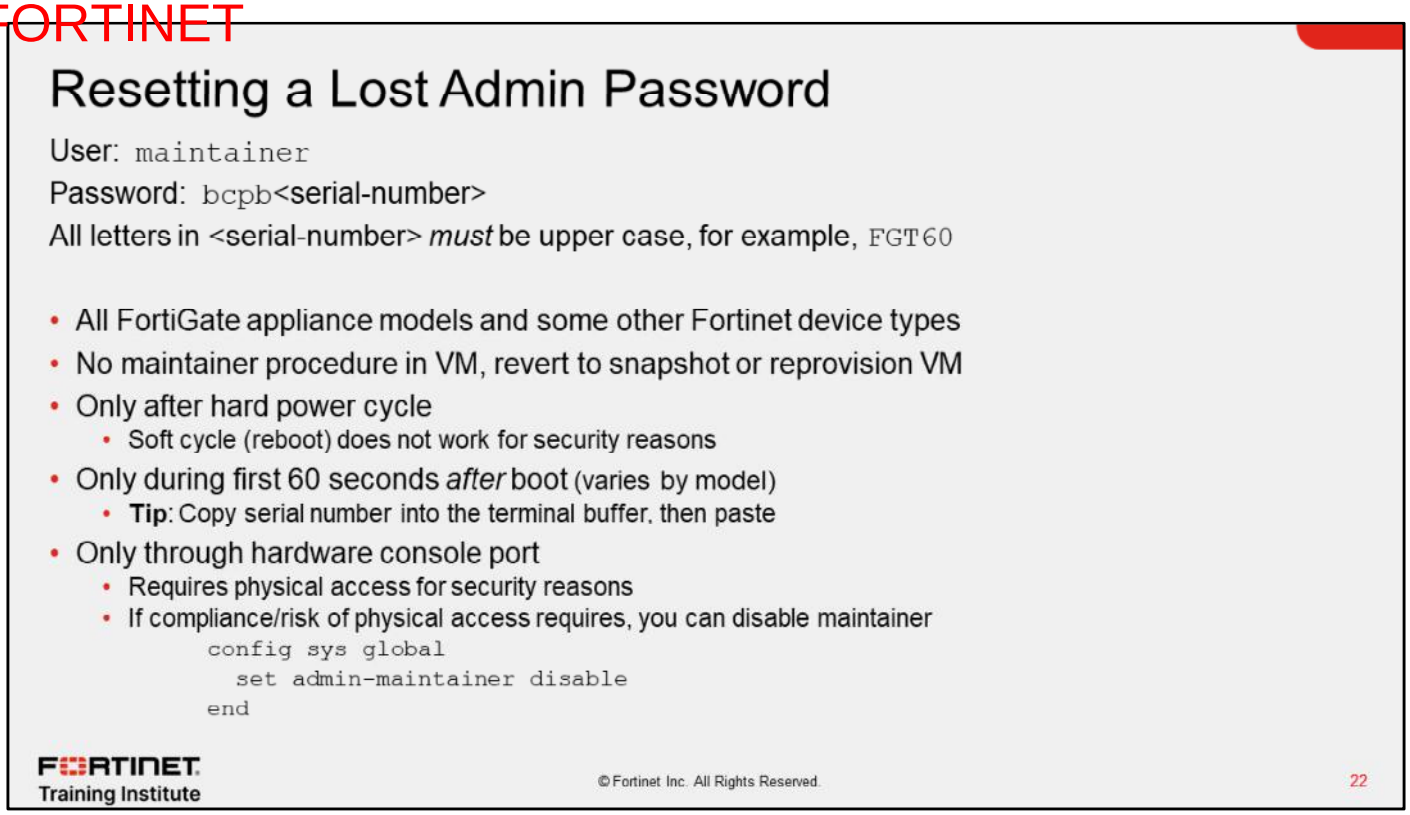

What happens if you forget the password for your `admin` account, or a malicious employee changes it?

This recovery method is available on all FortiGate devices and even some non-FortiGate devices, like FortiMail. There is no maintainer procedure in the VM. The administrator must revert to a snapshot or reprovision the VM and restore the configuration. It's a *temporary* account, only available through the local console port, and only after a hard reboot—disrupting power by unplugging or turning off the power, then restoring it. You must physically shut off FortiGate, then turn it back on, not reboot it through the CLI.

The `maintainer` login is available for login only for about 60 seconds after the restart completes (or less time on older models).

If you cannot ensure physical security, or have compliance requirements, you can disable the `maintainer` account. Use caution if you disable `maintainer` and then lose your `admin` password, because you cannot recover access to your FortiGate device. In order to regain access in this scenario, you will need to reload the device. This will reset to the device to its factory default settings.

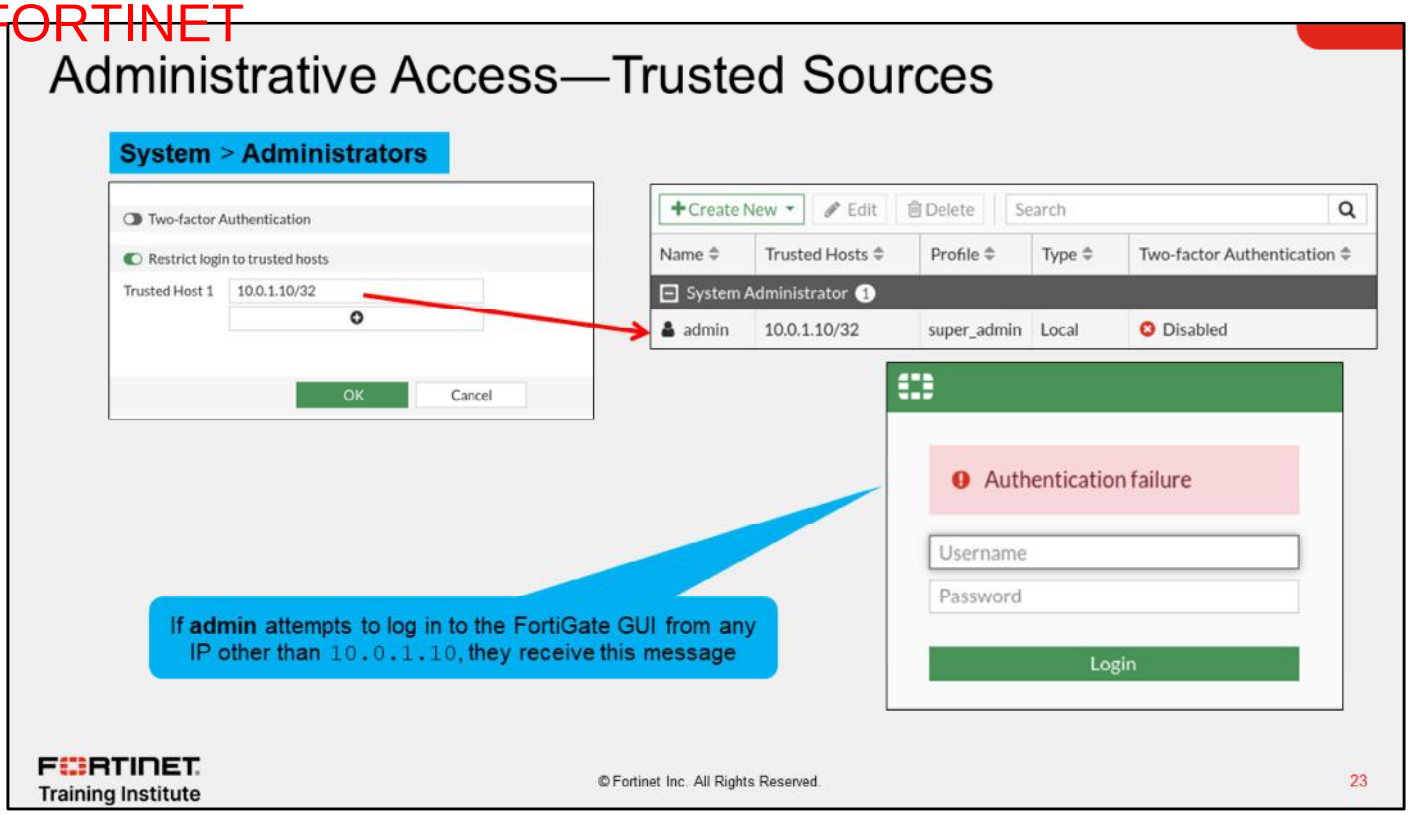

Another way to secure FortiGate is to define the hosts or subnets that are trusted sources from which to log in.

In this example, we have configured `10.0.1.10` as the only trusted IP for **admin** from which **admin** logs in. If **admin** attempts to log in from a machine with any other IP, they will receive an authentication failure message.

Note that If trusted hosts are configured on all administrators and an administrator is trying to log in from an IP address that is not set on any of the trusted hosts for any administrators, then the administrator will not get the login page but rather will receive the message: “Unable to contact server”.

If you leave any IPv4 address as `0.0.0.0/0`, it means that connections from any source IP will be allowed. By default, `0.0.0.0/0` is the configuration for the administrator, although you may want to change this.

Notice that each account can define its management host or subnet differently. This is especially useful if you are setting up VDOMs on FortiGate, where the VDOM administrators may not even belong to the same organization. Be aware of any NAT that occurs between the desired device and FortiGate. You can easily prevent an administrator from logging in from the desired IP address if it is later NATed to another address before reaching FortiGate, thus defeating the purpose of the trusted hosts.

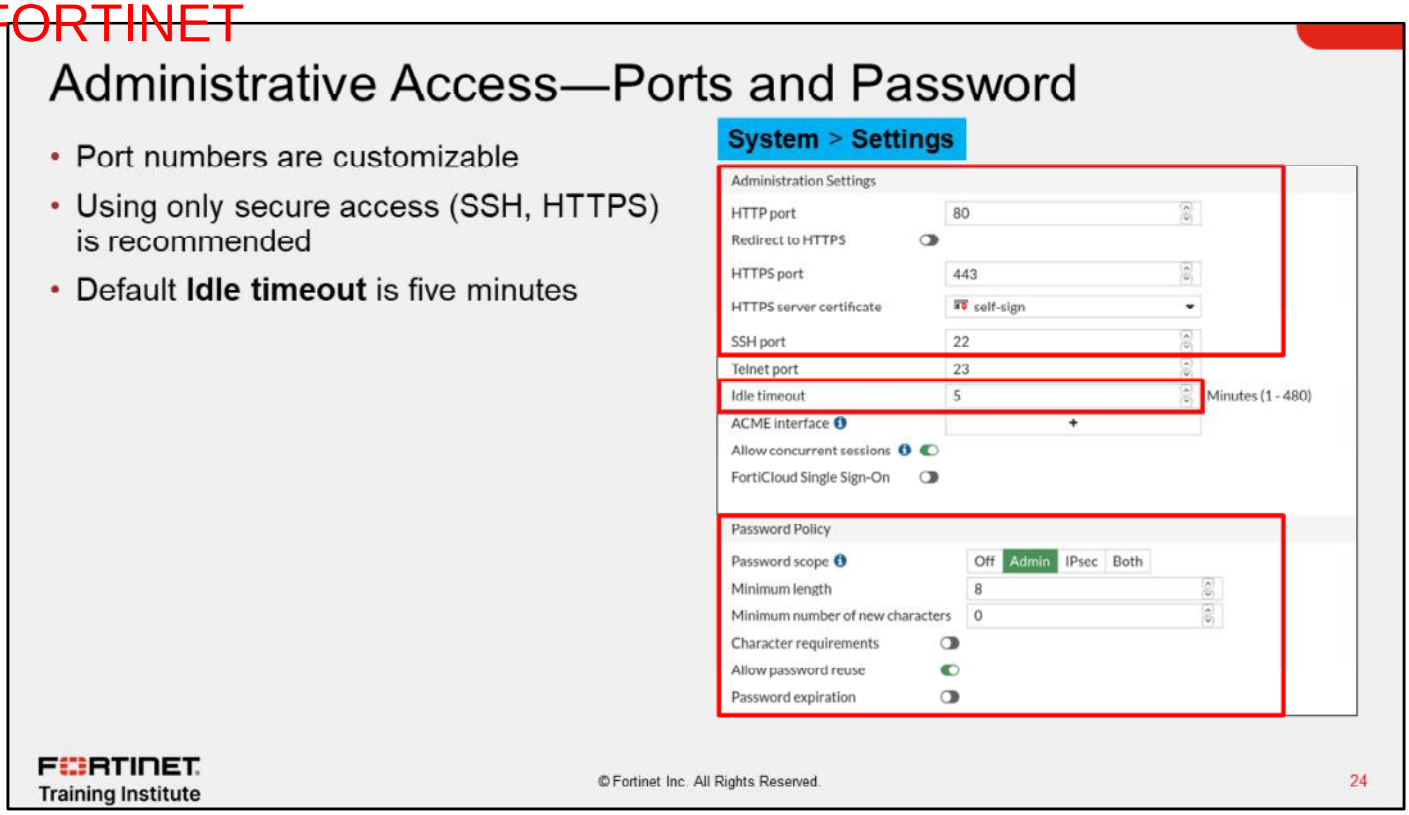

You may also want to customize the administrative protocols port numbers.

You can choose whether to allow concurrent sessions. You can use concurrent sessions to avoid accidentally overwriting settings, if you usually keep multiple browser tabs open, or accidentally leave a CLI session open without saving the settings, then begin a GUI session and accidentally edit the same settings differently.

For better security, use only secure protocols, and enforce password complexity and changes.

The **Idle timeout** settings specifies the number of minutes before an inactive administrator session times out (default is five minutes). A shorter idle timeout is more secure, but increasing the timer can help reduce the chance of administrators being logged out while testing changes.

You can override the idle timeout setting per administrator profile using the **Override Idle Timeout** setting.

You can configure an administrator profile to increase inactivity timeout and facilitate use of the GUI for central monitoring. The **Override Idle Timeout** setting allows the **admintimeout** value, under **config system accprofile**, to be overridden per access profile.

Note that you can do this on a per profile basis, to avoid the option from being unintentionally set globally.

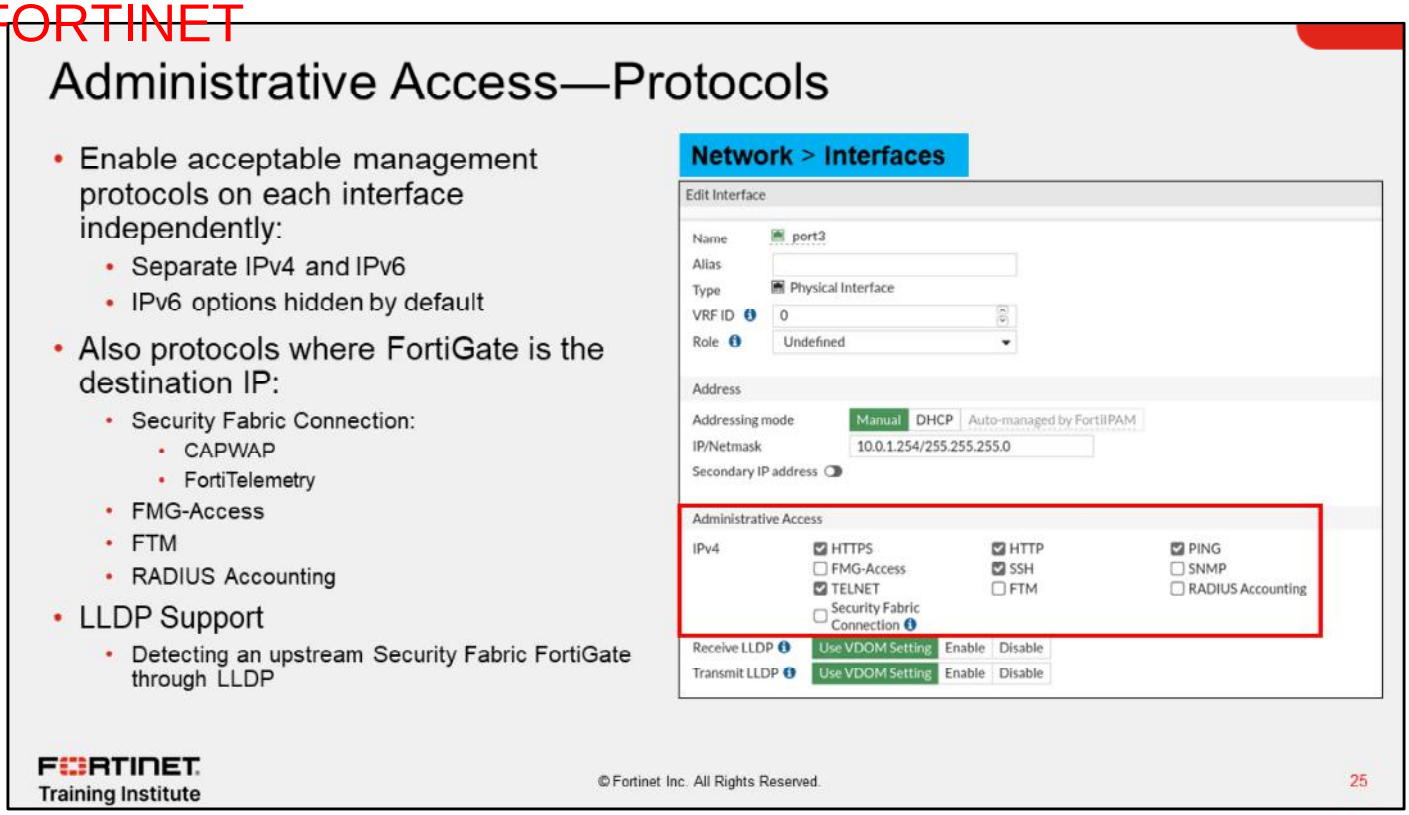

You've defined the management subnet—that is, the trusted hosts—for each administrator account. How do you enable or disable management protocols?
This is specific to each interface. For example, if your administrators connect to FortiGate only from port3, then you should disable administrative access on all other ports. This prevents brute force attempts and also insecure access. Your management protocols are HTTPS, HTTP, PING, and SSH. By default, the HTTP and TELNET option is not visible on the GUI.

Consider the location of the interface on your network. Enabling PING on an internal interface is useful for troubleshooting. However, if it's an external interface (in other words, exposed to the internet), then the PING protocol could expose FortiGate to a DoS attack. You should disable protocols that do not encrypt data flow, such as HTTP and TELNET. IPv4 and IPv6 protocols are separate. It's possible to have both IPv4 and IPv6 addresses on an interface, but only respond to pings on IPv6.

Security Fabric connection includes CAPWAP and FortiTelemetry. Protocols like FortiTelemetry are *not* for administrative access, but, like GUI and CLI access, they are protocols where the packets have FortiGate as a destination IP. Use the FortiTelemetry protocol specifically for managing FortiClient and the Security Fabric. Use the CAPWAP protocol for FortiAP, FortiSwitch, and FortiExtender when they are managed by FortiGate. Use the FMG-Access protocol specifically for communicating with FortiManager when that server is managing multiple FortiGate devices. Use the RADIUS accounting protocol when FortiGate needs to listen for and process RADIUS accounting packets for single sign-on authentication. FTM, or FortiToken Mobile push, supports second-factor authentication requests from a FortiToken mobile app.
When you assign the interface roles LAN or WAN to the appropriate interfaces, your FortiGate uses the Link Layer Discovery Protocol (LLDP) to detect if there's an upstream FortiGate in your network. If FortiGate discovers an upstream FortiGate, you're prompted to configure the upstream FortiGate device to join the Security Fabric.

Interface IPs

- In NAT mode, you can't use interfaces until they have an IP address:
  - Manually assigned
  - Automatic
    - DHCP
    - PPPoE

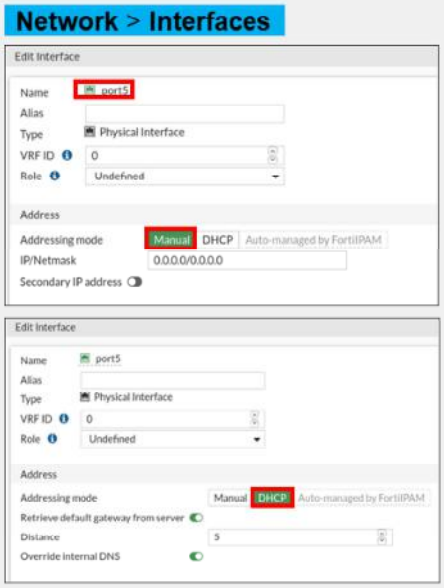

FORTINET Training Institute

26

When FortiGate is operating in NAT mode, every interface that handles traffic must have an IP address. When in NAT mode, FortiGate can use the IP address to source the traffic, if it needs to start or reply to a session, and as a destination address for devices trying to contact FortiGate or route traffic through it. There are multiple ways to get an IP address:

- Manually
- Automatically, using either DHCP or PPPoE (available on the CLI)

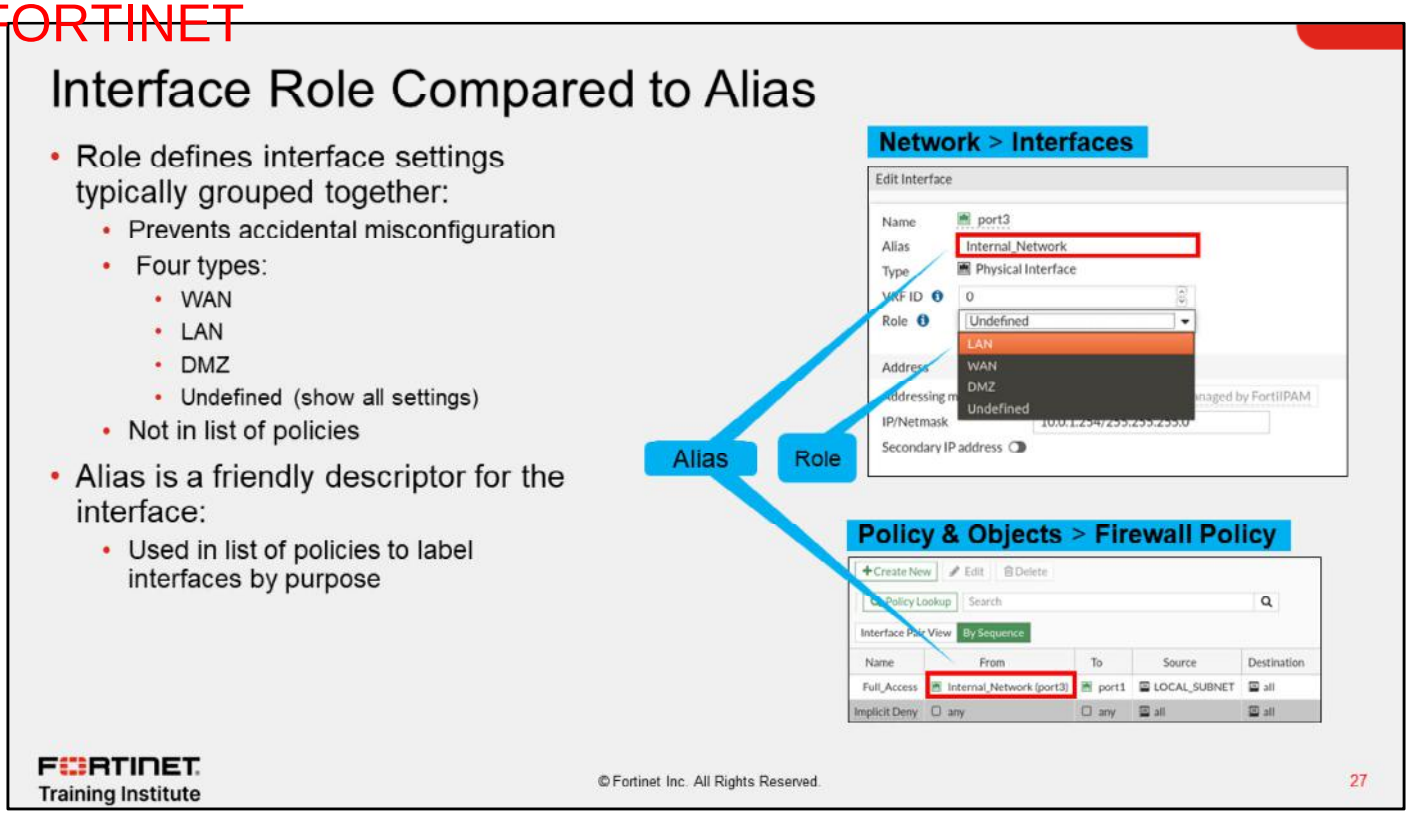

How many times have you seen network issues caused by a DHCP server—not client—enabled on the WAN interface?

You can configure the interface role. The roles shown on the GUI are the usual interface settings for that part of a topology. Settings that do not apply to the current role are hidden on the GUI. (All settings are always available on the CLI regardless of the role.) This prevents accidental misconfiguration.

For example, when the role is configured as **WAN**, there is no DHCP server and device detection configuration available. Device detection is usually used to detect devices internally on your LAN.

If there is an unusual case, and you need to use an option that's hidden by the current role, you can always switch the role to **Undefined**. This displays all options.

To help you remember the use of each interface, you can give them aliases. For example, you could call port3 `internal_network`. This can help to make your list of policies easier to comprehend.

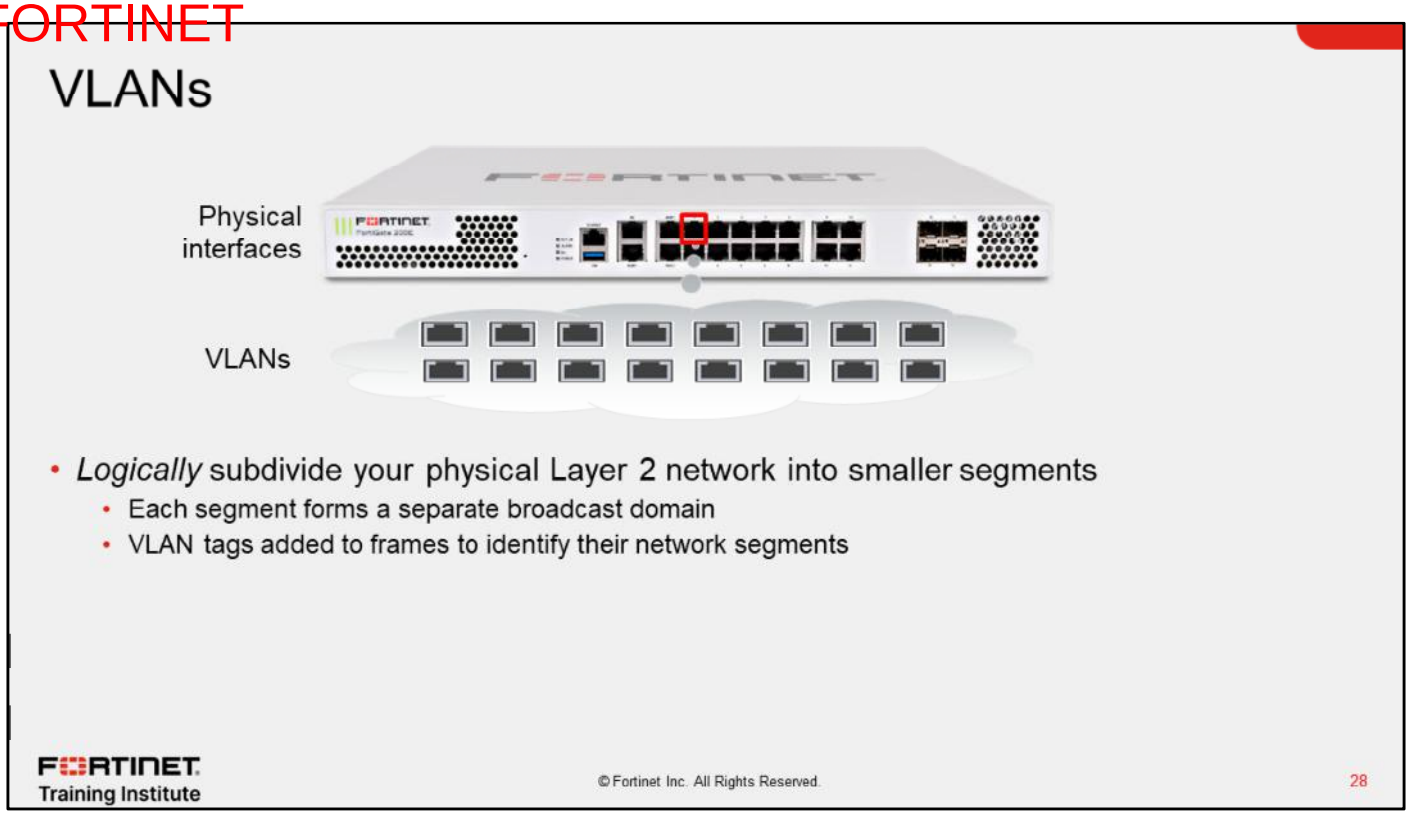

VLANs split your physical LAN into multiple, logical LANs. In NAT operation mode, each VLAN forms a separate broadcast domain. Multiple VLANs can coexist in the same physical interface, provided they have different VLAN IDs. In this way, a physical interface is split into two or more logical interfaces.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

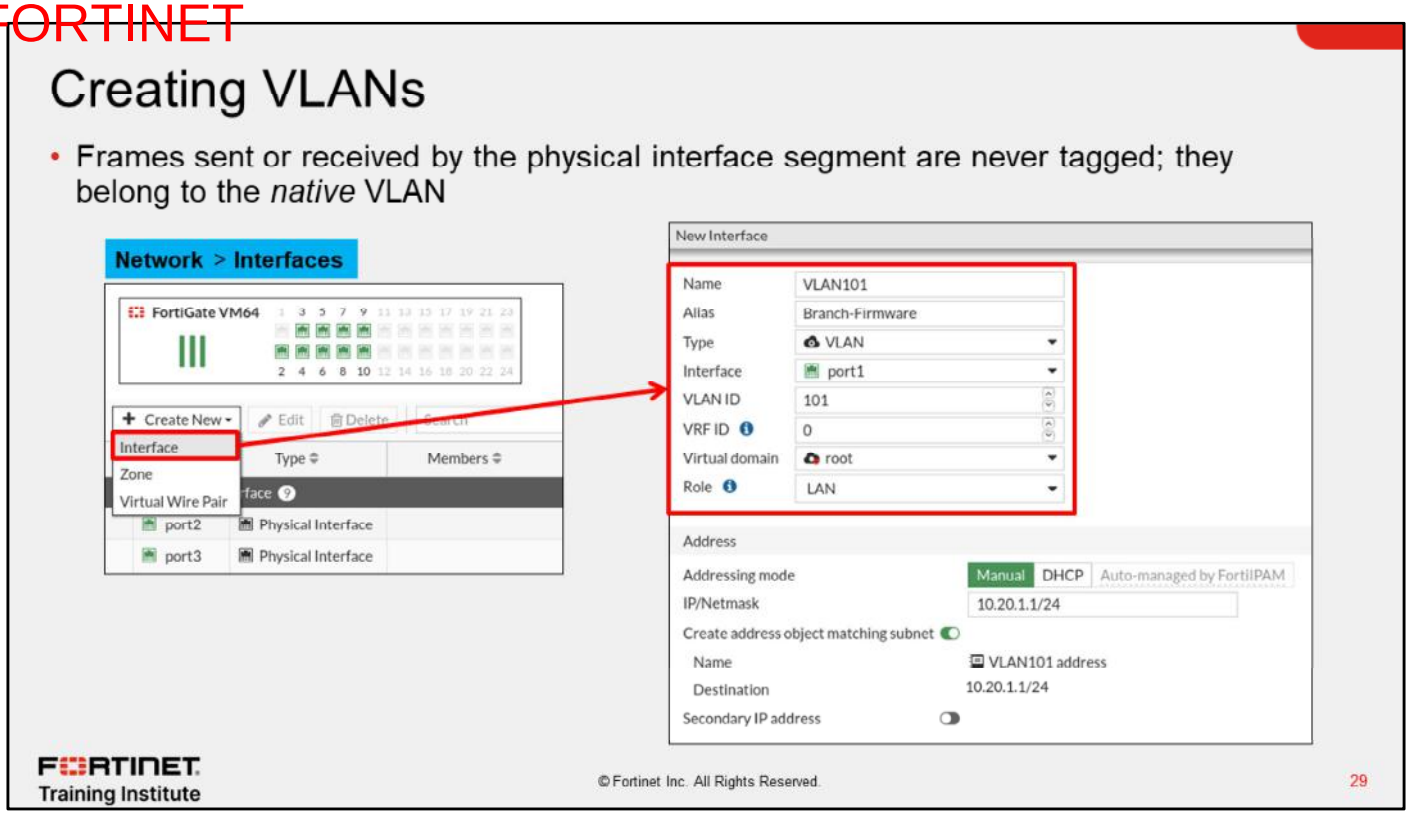

To create a VLAN using the GUI, click **Create New,** select **Interface,** and then, in the **Type** drop-down list, select **VLAN**. You must specify the VLAN ID and the physical interface to which the VLAN will be bound. Frames that belong to interfaces of that type are always tagged. On the other 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.

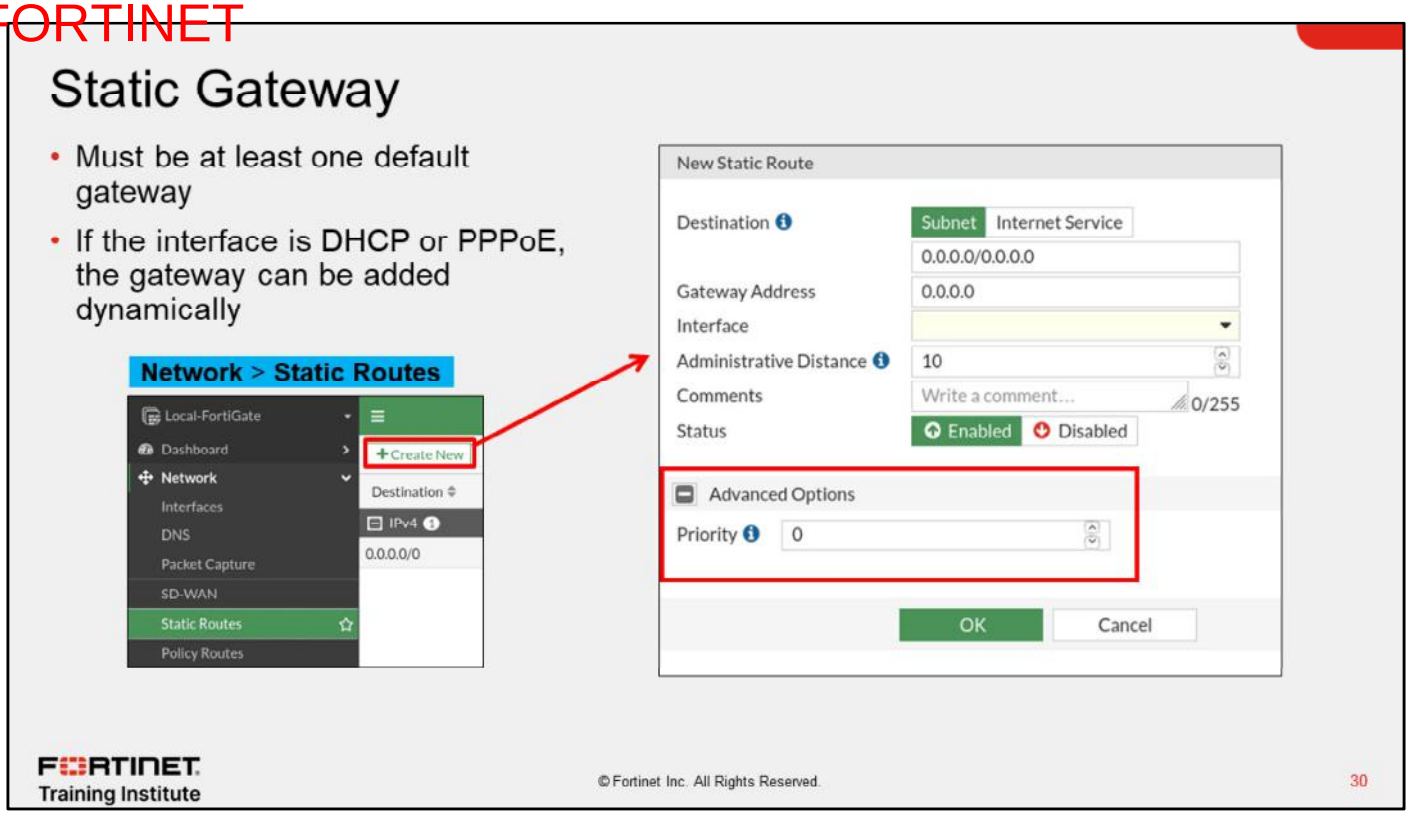

Before you integrate FortiGate into your network, you should configure a default gateway.

If FortiGate gets its IP address through a dynamic method such as DHCP or PPPoE, then it should also retrieve the default gateway.

Otherwise, you must configure a static route. Without this, FortiGate will not be able to respond to packets outside the subnets directly attached to its own interfaces. It probably also will not be able to connect to FortiGuard for updates, and may not correctly route traffic.

You should make sure that FortiGate has a route that matches all packets (destination is `0.0.0.0/0`), known as a default route, and forwards them through the network interface that is connected to the internet, to the IP address of the next router.

Routing completes the basic network settings that are required before you can configure firewall policies.

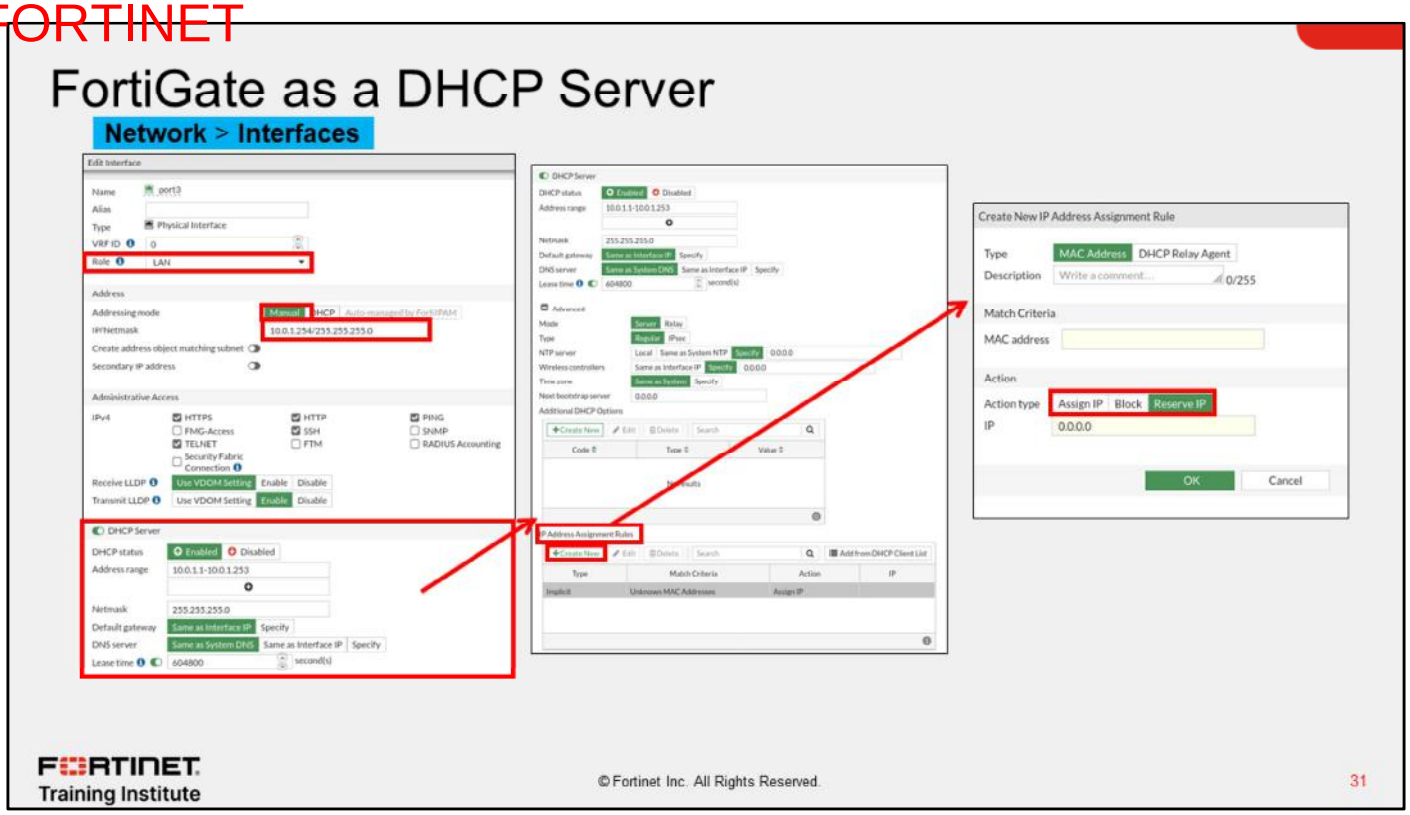

Wireless clients are not the only ones that can use FortiGate as their DHCP server.

For an interface (such as port3), select the **Manual** option, enter a static IP, and then enable the **DHCP Server** option. Options for the built-in DHCP server appear, including provisioning features, such as DHCP options and IP address assignment rules. You can also block specific MAC addresses from receiving an IP address.

Note that the screenshot on the middle of the slide shows that you can create IP address assignment rules in the **IP Address Assignment Rule** section. The **IP Address Assignment Rule** section allows you to assign, block or reserve the IP address to the host. It also allows you to select actions for unknown MAC addresses. The default action is **Assign IP**; however, you can change the default action type to **Assign IP** or **Block**.

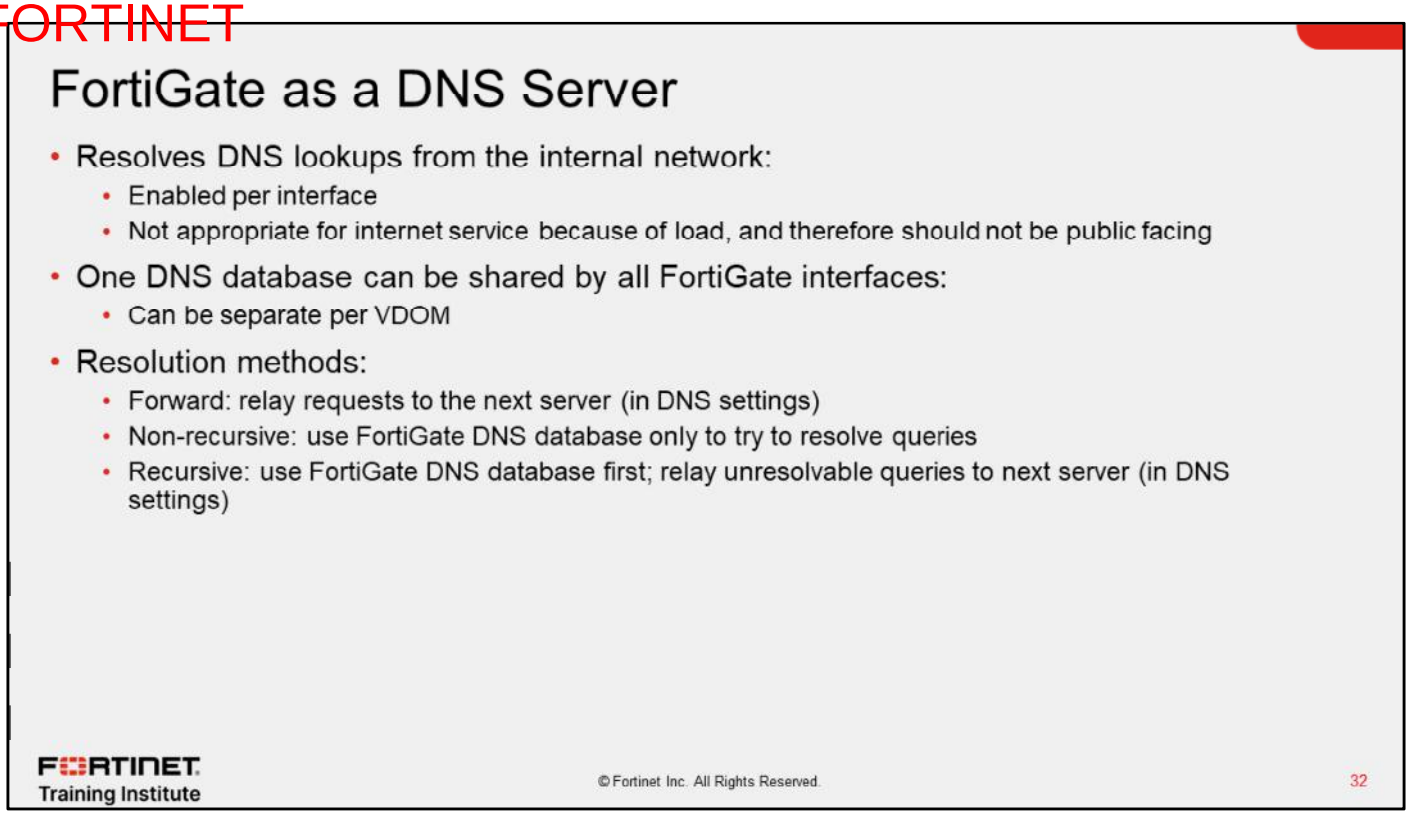

You can configure FortiGate to act as your local DNS server. You can enable and configure DNS separately on each interface.

A local DNS server can improve performance for your FortiMail device or other devices that use DNS queries frequently. If your FortiGate device offers DHCP to your local network, you can use DHCP to configure those hosts to use FortiGate as both the gateway and DNS server.

FortiGate can answer DNS queries in one of three ways:

- Forward: relays all queries to a separate DNS server (that you have configured in **Network** > **DNS**); that is, it acts as a DNS relay instead of a DNS server.
- Non-Recursive: replies to queries for items in the FortiGate DNS databases and does not forward unresolvable queries.
- Recursive: replies to queries for items in the FortiGate DNS databases and forwards all other queries to a separate DNS server for resolution.

You can configure all modes on the GUI or CLI.

Knowledge Check

1. How do you restrict logins to FortiGate from only specific IP addresses?
   - A. Change the FortiGate management interface IP address.
   - ✓ B. Configure a trusted host.

2. As a best security practice when configuring administrative access to FortiGate, which protocol should you disable?
   - ✓ A. Telnet
   - B. SSH

3. When configuring FortiGate as a DNS server, which resolution method uses only the FortiGate DNS database to try to resolve queries?
   - ✓ A. Non-recursive
   - B. Recursive

FORTINET Training Institute

33

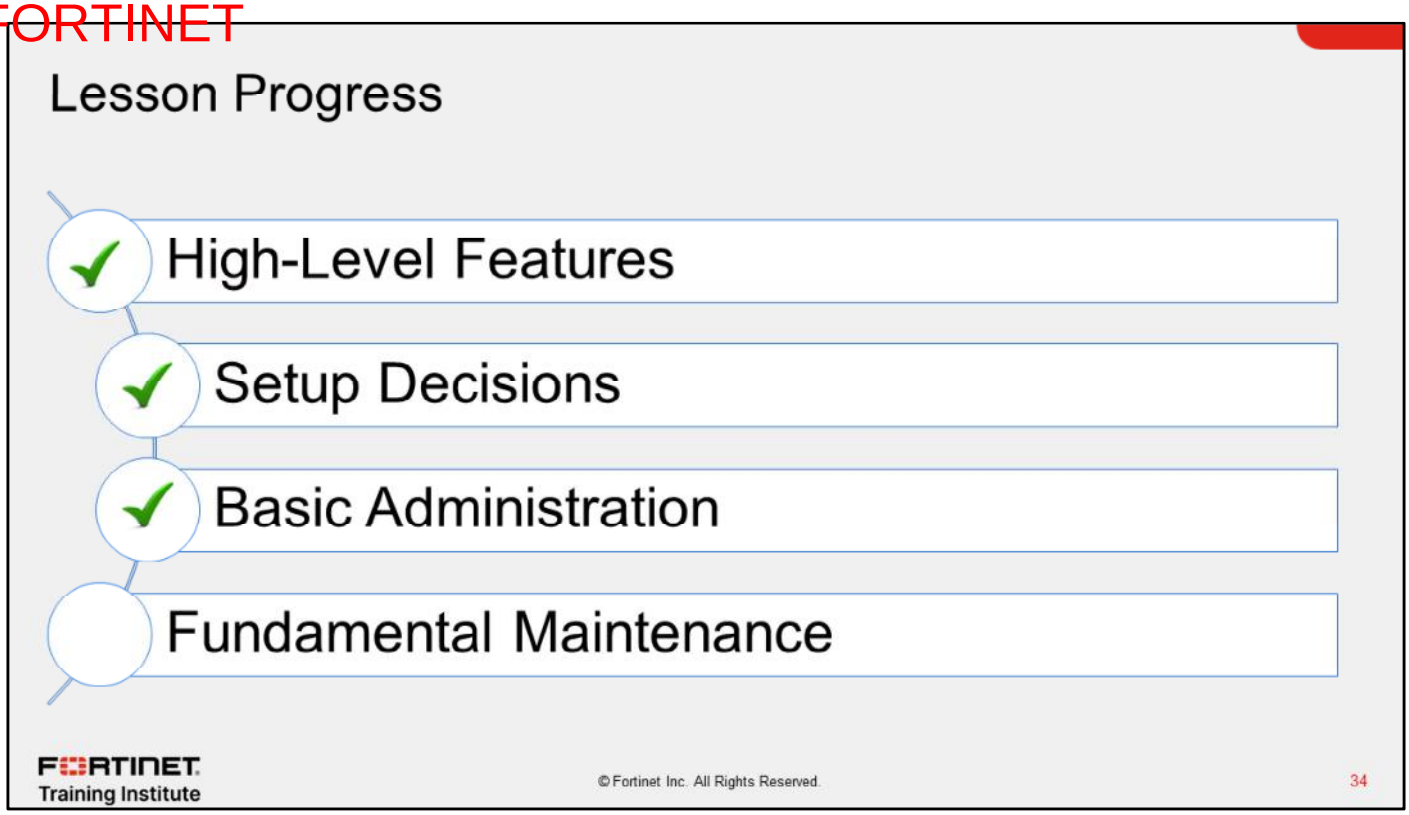

Good job! You now have the knowledge needed to carry out some basic administrative tasks. You also know how to enable DHCP and DNS services on FortiGate.

Now, you will learn about fundamental maintenance.

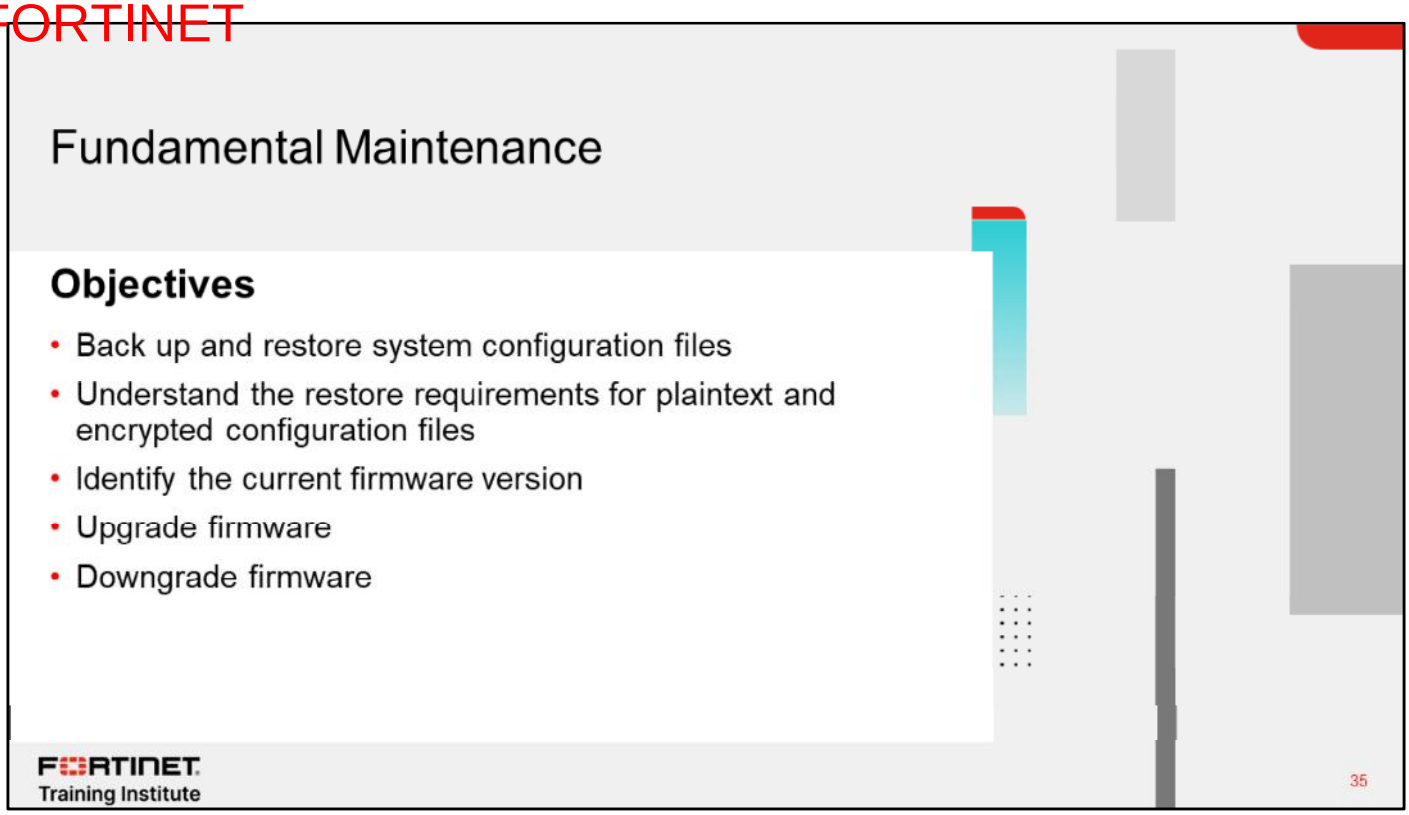

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the basic maintenance of FortiGate, you will be able to perform the vital activities of backing up and restoring configurations, upgrading and downgrading firmware, and ensuring that FortiGate remains reliably in service throughout its lifecycle.

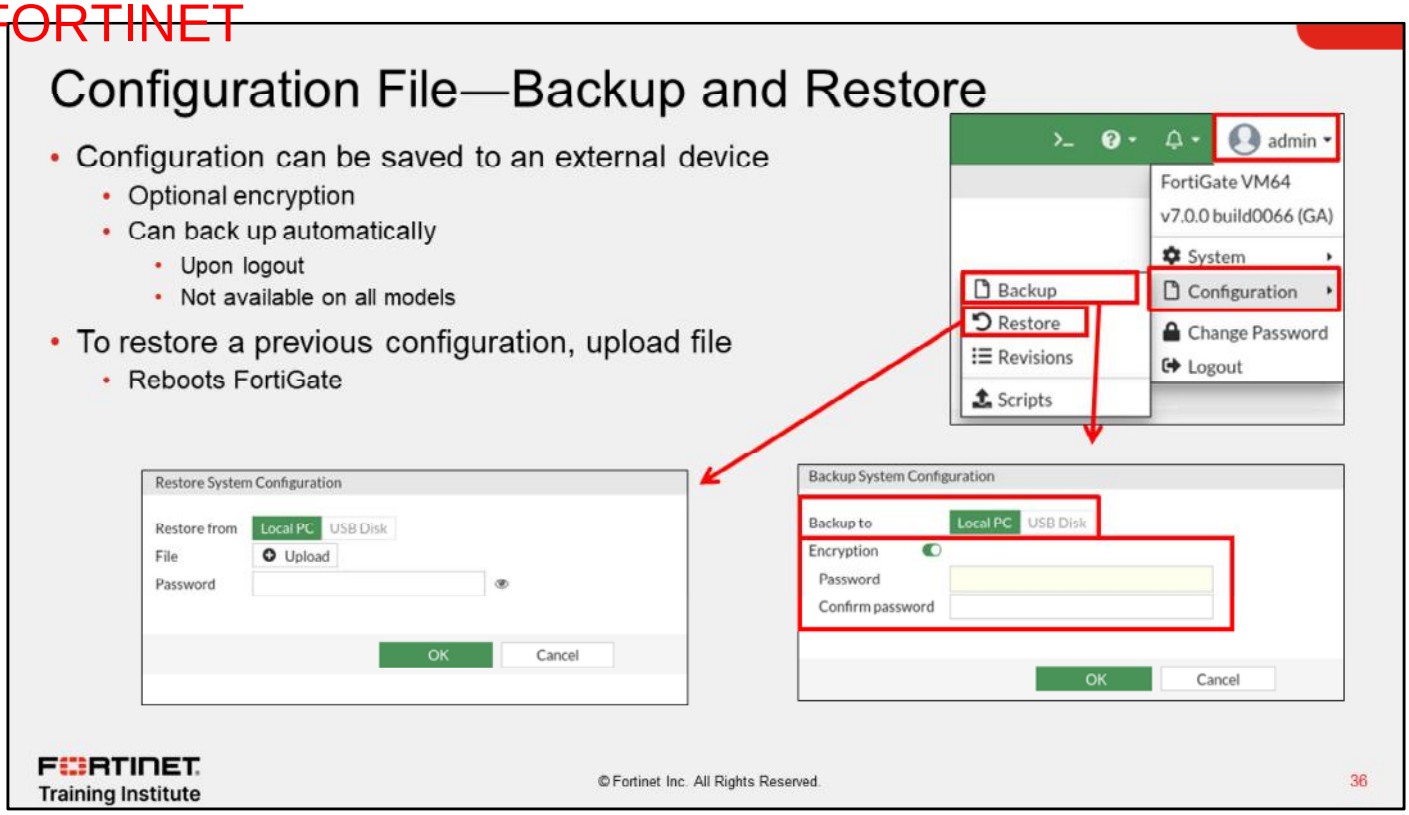

Now that FortiGate has basic network settings and administrative accounts, you will learn how to back up the configuration. In addition to selecting the destination of the backup file, you can choose to encrypt or not to encrypt the backup file. Even if you choose not to encrypt the file, which is the default, the passwords stored in the file are hashed, and, therefore, obfuscated. The passwords that are stored in the configuration file would include passwords for the administrative users and local users, and preshared keys for your IPSec VPNs. It may also include passwords for the FSSO and LDAP servers.

The other option is to encrypt the configuration file with a password. Besides securing the privacy of your configuration, it also has some effects you may not expect. After encryption, the configuration file cannot be decrypted without the password and a FortiGate of the same model and firmware. This means that if you send an encrypted configuration file to Fortinet technical support, even if you give them the password, they cannot load your configuration until they get access to the same model of FortiGate. This can cause unnecessary delays when resolving your ticket.

If you enable virtual domains (VDOMs), subdividing the resources and configuration of your FortiGate device, each VDOM administrator can back up and restore their own configurations. You don't have to back up the entire FortiGate configuration, however, it is still recommended.

Backups are needed to help speed up the return to production in the event of an unforeseen disaster that damages FortiGate. Having to recreate hundreds of policies and objects from scratch takes a significant amount of time, while loading a configuration file on a new device takes much less.

Restoring a configuration file is very similar to backing one up and restarts FortiGate.

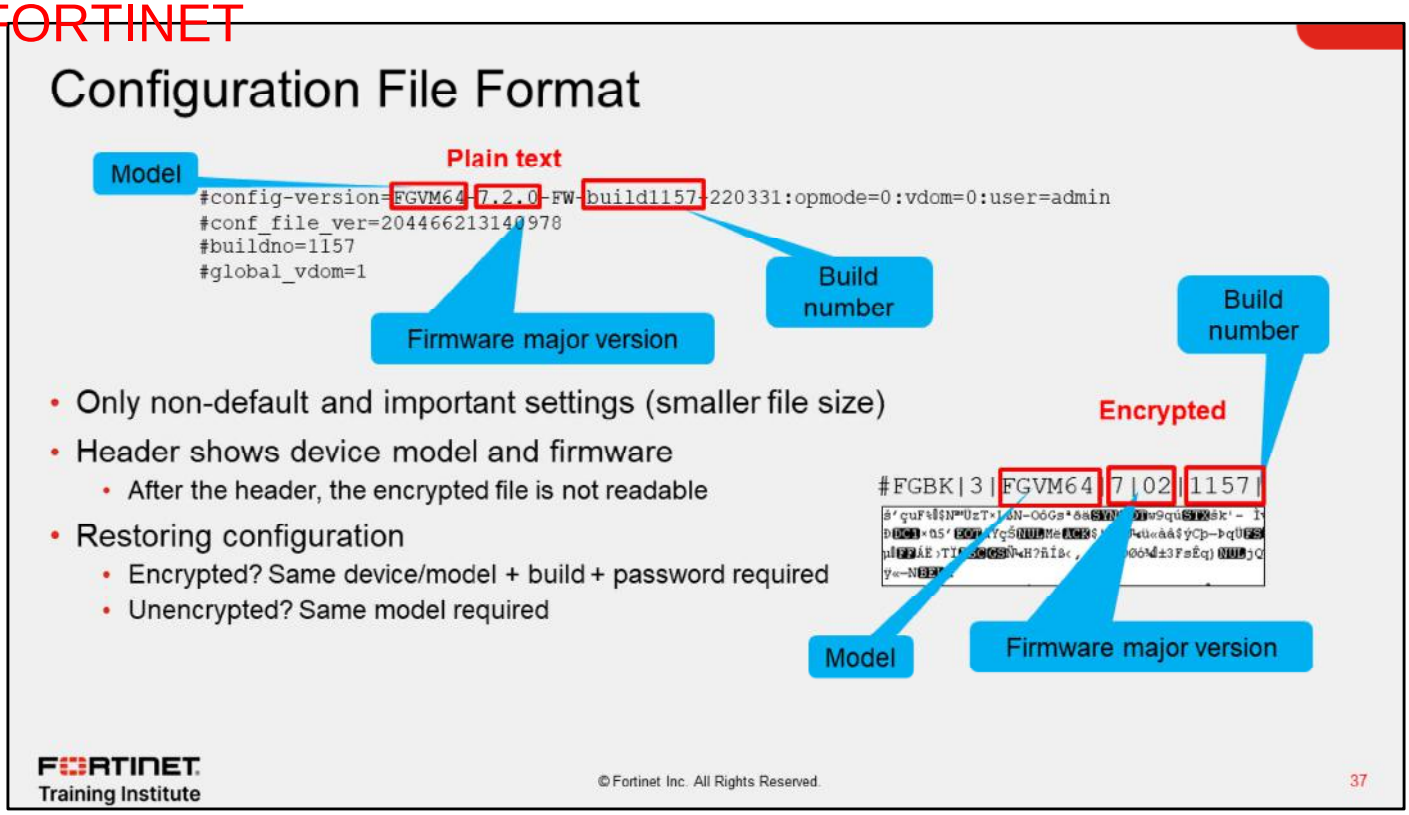

If you open the configuration file in a text editor, you'll see that both encrypted and unencrypted configuration files contain a cleartext header that contains some basic information about the device. The example on this slide shows what information is included. To restore an encrypted configuration, you must upload it to a FortiGate device of the same model and firmware, then provide the password.

To restore an unencrypted configuration file, you are required to match only the FortiGate model. If the firmware is different, FortiGate will attempt to upgrade the configuration. This is similar to how it uses upgrade scripts on the existing configuration when upgrading firmware. However, it is still recommended to match the firmware on FortiGate to the firmware listed in the configuration file.

Usually, the configuration file contains only non-default settings, plus few default, yet crucial, settings. This minimizes the size of the backup, which could otherwise be several megabytes in size.

Configuration File Format (Contd)

- Support YAML
- Configuration can be backedup and restored by CLI only
  - # execute backup yaml-config {ftp | tftp} <filename> server [username] [password]
  - # execute restore yaml-config {ftp | tftp} <filename> server [username] [password]

Default Format

```
config system global
  set admintimeout 480
  set alias "FortiGate-100F"
end
config system settings
  set default-voip-alg-mode kernel-helper-based
  set gui-dynamic-routing enable
end
config system interface
     edit "port1"
          set vdom "root"
          set ip 204.126.10.3 255.255.254.0
          set allowaccess ping
          config secondaryip
              edit 1
                set ip 204.126.10.2 255.255.255.0
                set allowaccess ping
              end
```

YAML Format

```
config_system_global:
    admintimeout:480
    alias:FortiGate-100F
config_system_settings:
    default-voip-alg-mode: kernel-helper-based
    gui-dynamic-routing: enable
config_system_interface:
    - port1:
        vdom: root
        ip: "204.126.10.3 255.255.254.0"
        allowaccess: ping
        secondaryip:
            - 0:
                ip: "204.126.10.2 255.255.255.0"
                allowaccess: ping
```

FORTINET Training Institute

38

YAML format becomes more and more popular often use to create configuration files. FortiOS now supports YAML format, you can take a backup as well as restore YAML configuration file using CLI commands. You must provide server type: ftp or tftp, filename, server IP address, and user credential's to backup or restore configuration file in YAML format.

This slide shows the sample configuration to understand the difference between the default file format and YAML format.

Upgrade Firmware

- You can view the current firmware version on the dashboard or in **System** > **Fabric Management** (or on the CLI: `get system status`)
- If there is an updated firmware version, you are notified
- Firmware can be updated by clicking **Upgrade** and then selecting the **All Upgrades** or **File Upload** option
- Make sure you read the *Release Notes* to verify the upgrade path and other details

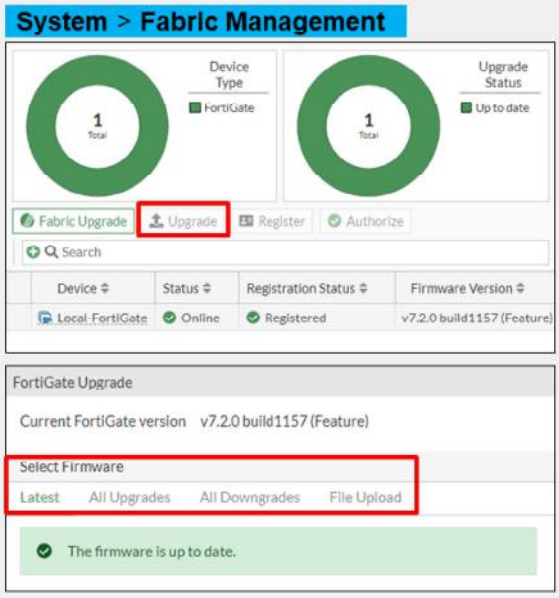

FORTINET Training Institute

39

You can view the current firmware version in multiple places on the FortiGate GUI. When you first log in to FortiGate, the landing page is the dashboard. You can see the firmware version in the **System** widget. This information is also found at **System** > **Fabric Management**. And, of course, you can retrieve the information on the CLI using the command `get system status`.

If a new version of the firmware is available, you are notified on the dashboard and on the **Fabric Management** page. The **Fabric Management** page allows administrators to manage the firmware running on each FortiGate, FortiAP, and FortiSwitch in the Security Fabric, and to authorize and register these Fabric devices.

You can use **Upgrade** option to upgrade firmware of the selected device. The **Fabric Upgrade** option upgrades firmware for the root FortiGate as well as Fabric devices. You can also use this option to upgrade firmware for a non-Security Fabric FortiGate with managed FortiSwitch and FortiAP devices.
The **Fabric Upgrade** option uses released firmware images from FortiGuard.

You can also use the **Register** option to register a selected device to FortiCare and an **Authorize** option to authorize a selected device for use in security fabric.

Remember to read the *Release Notes* to make sure that you understand the supported upgrade path. The *Release Notes* also provide pertinent information that may affect the upgrade.

Knowledge Check

1. When restoring an encrypted system configuration file, in addition to needing the FortiGate model and firmware version from the time the configuration file was produced, what must you also provide?
   - ✓A. The password to decrypt the file
   - B. The private decryption key to decrypt the file

2. Which document should you consult to increase the chances of success before upgrading or downgrading firmware?
   - A. Cookbook
   - ✓B. Release Notes

FORTINET Training Institute

40

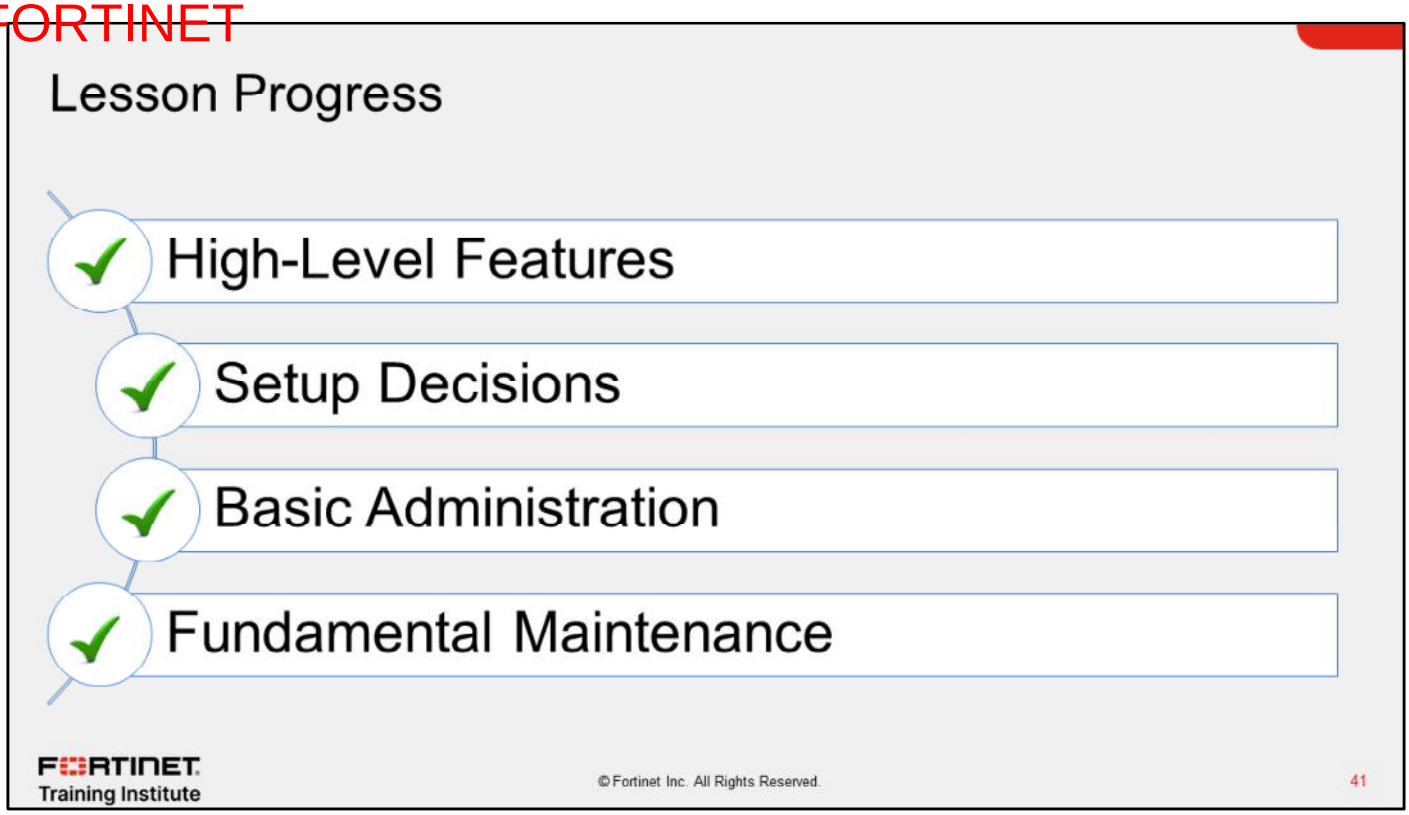

Congratulations! You have completed this lesson.

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

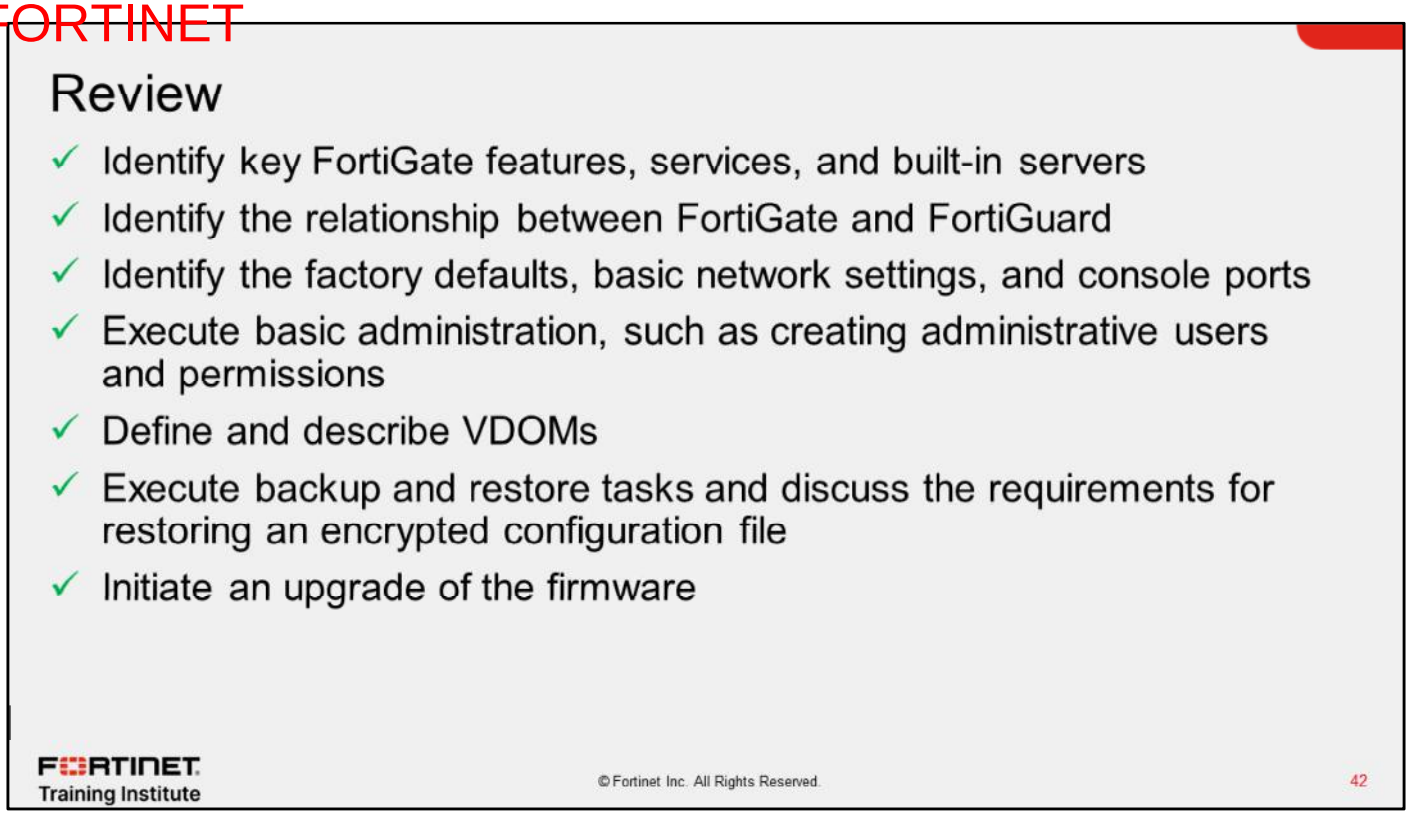

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how and where FortiGate fits into your network and how to perform basic FortiGate administration.

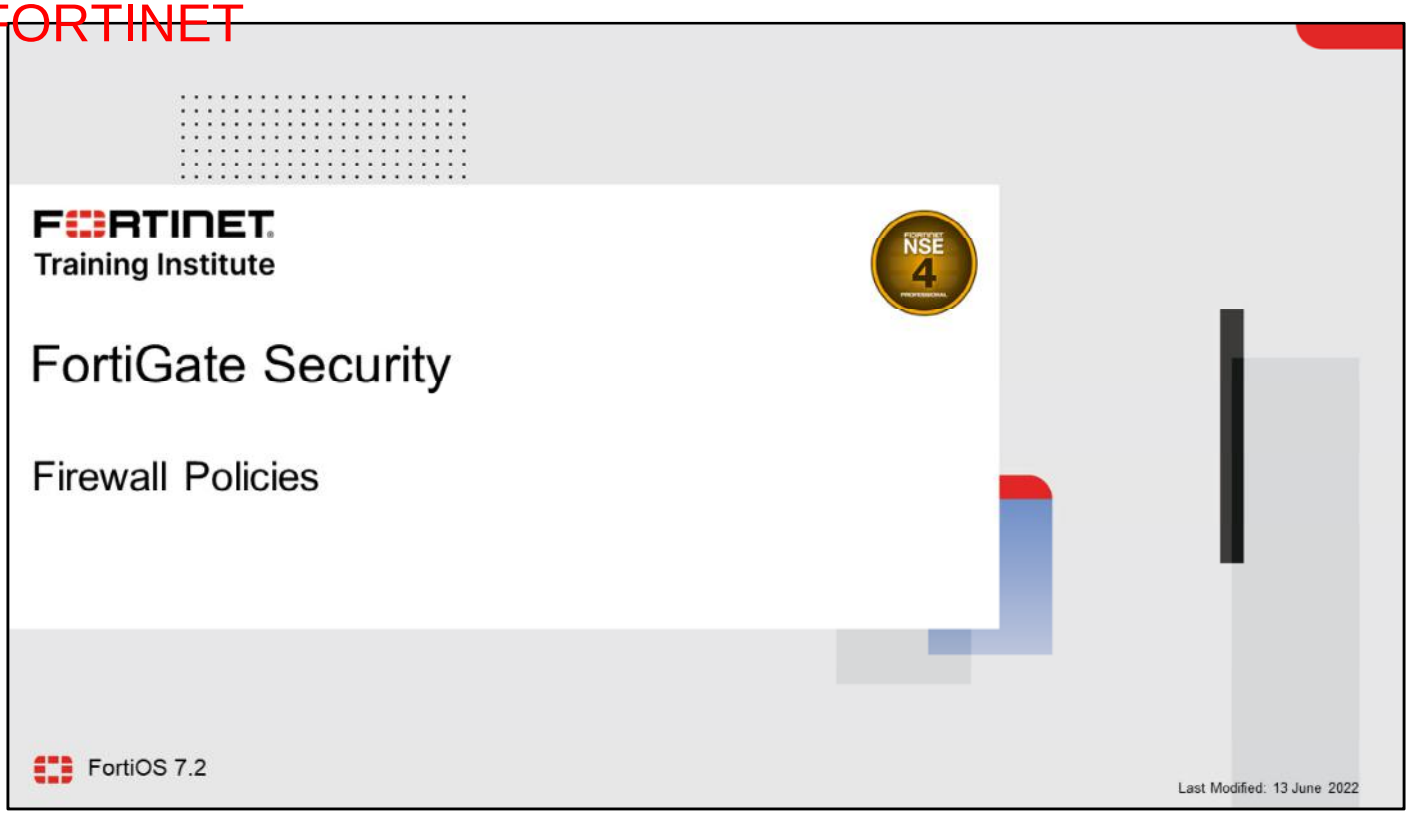

In this lesson, you will learn about firewall policies and how to apply them to allow and deny traffic passing through FortiGate. At its core, FortiGate is a firewall, so almost everything that it does to your traffic is linked to your firewall policies.

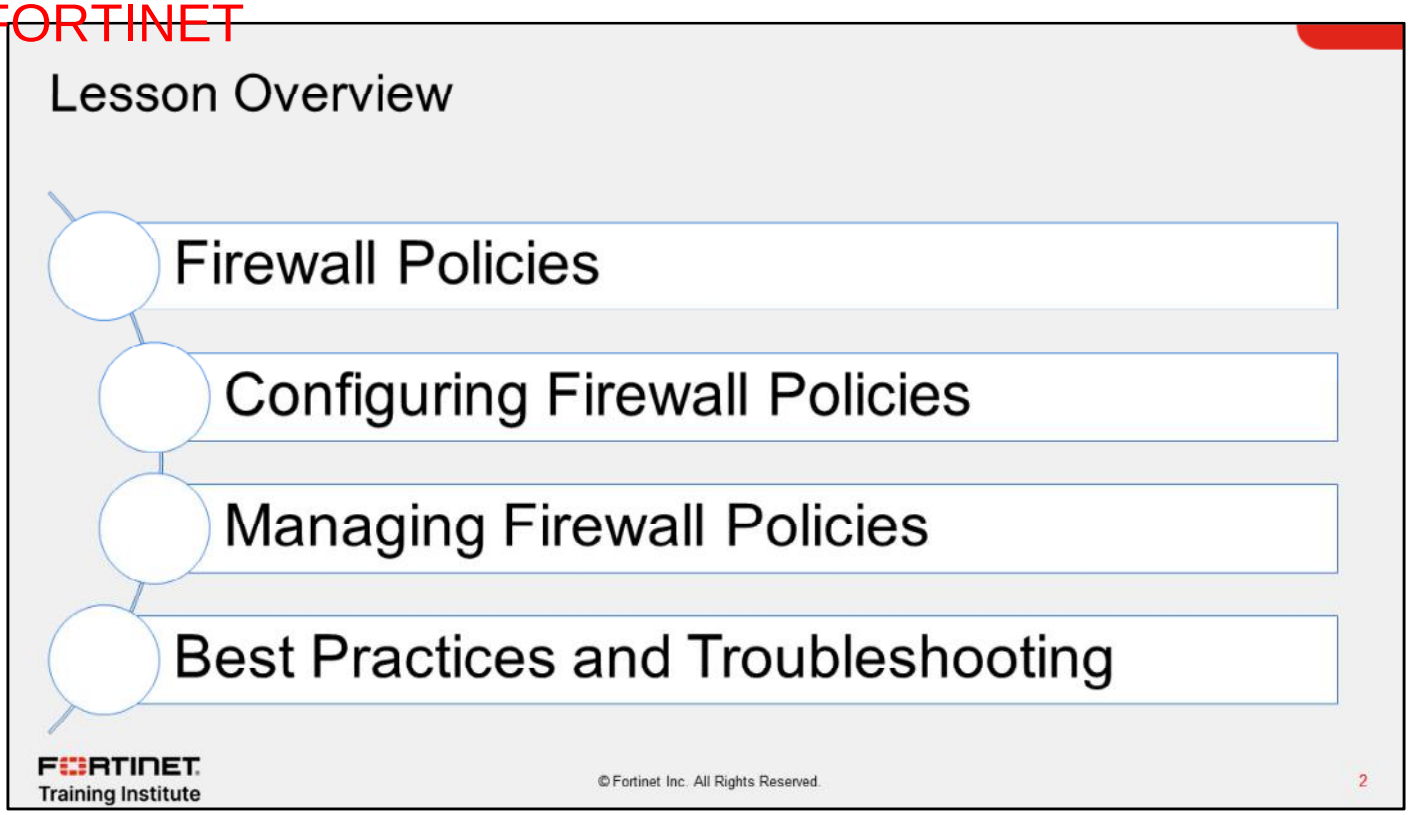

In this lesson, you will learn about the topics shown on this slide.

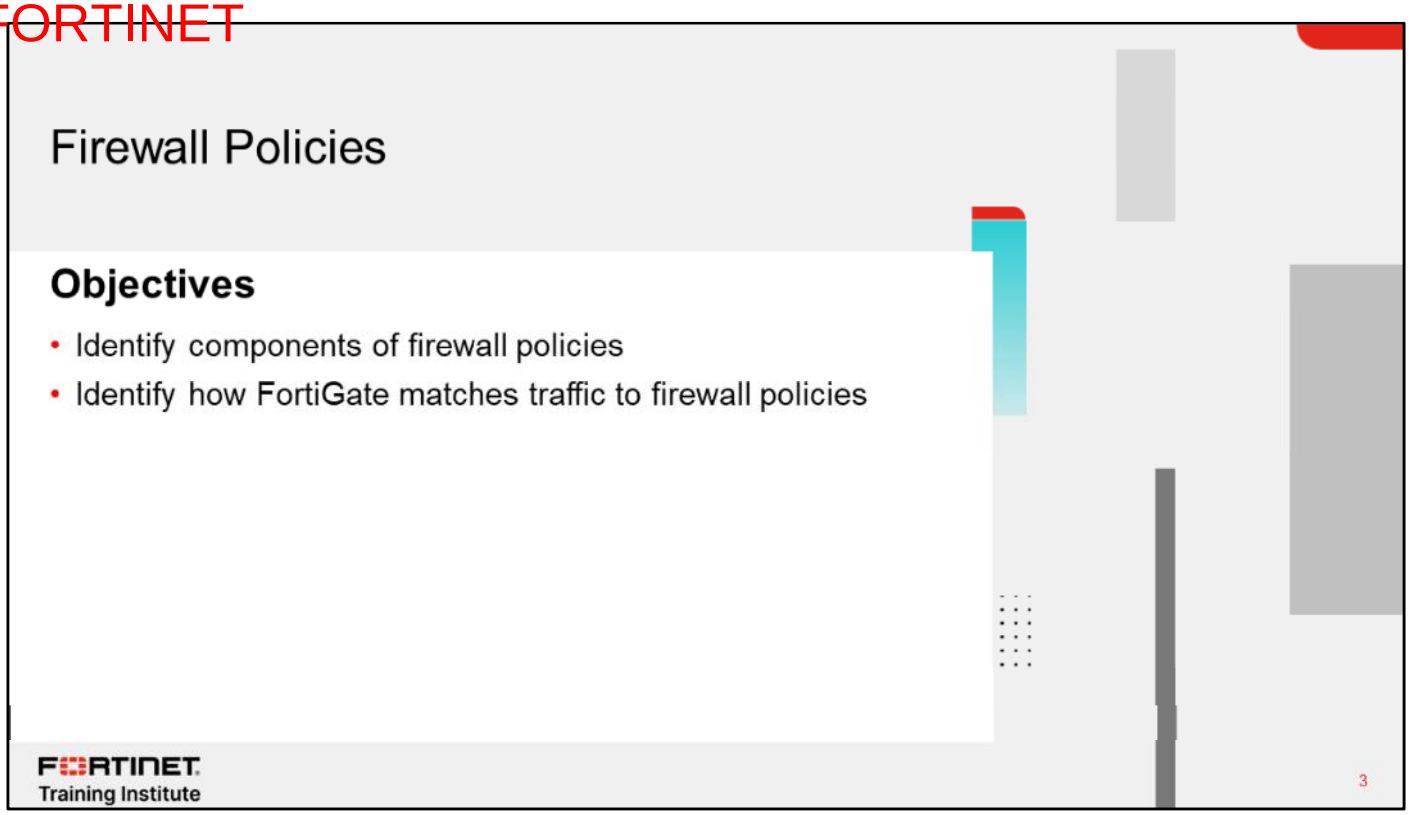

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in identifying the different components of firewall policies, and recognizing how FortiGate matches traffic with firewall policies and takes appropriate action, you will have a better understanding of how firewall policies interact with network traffic.

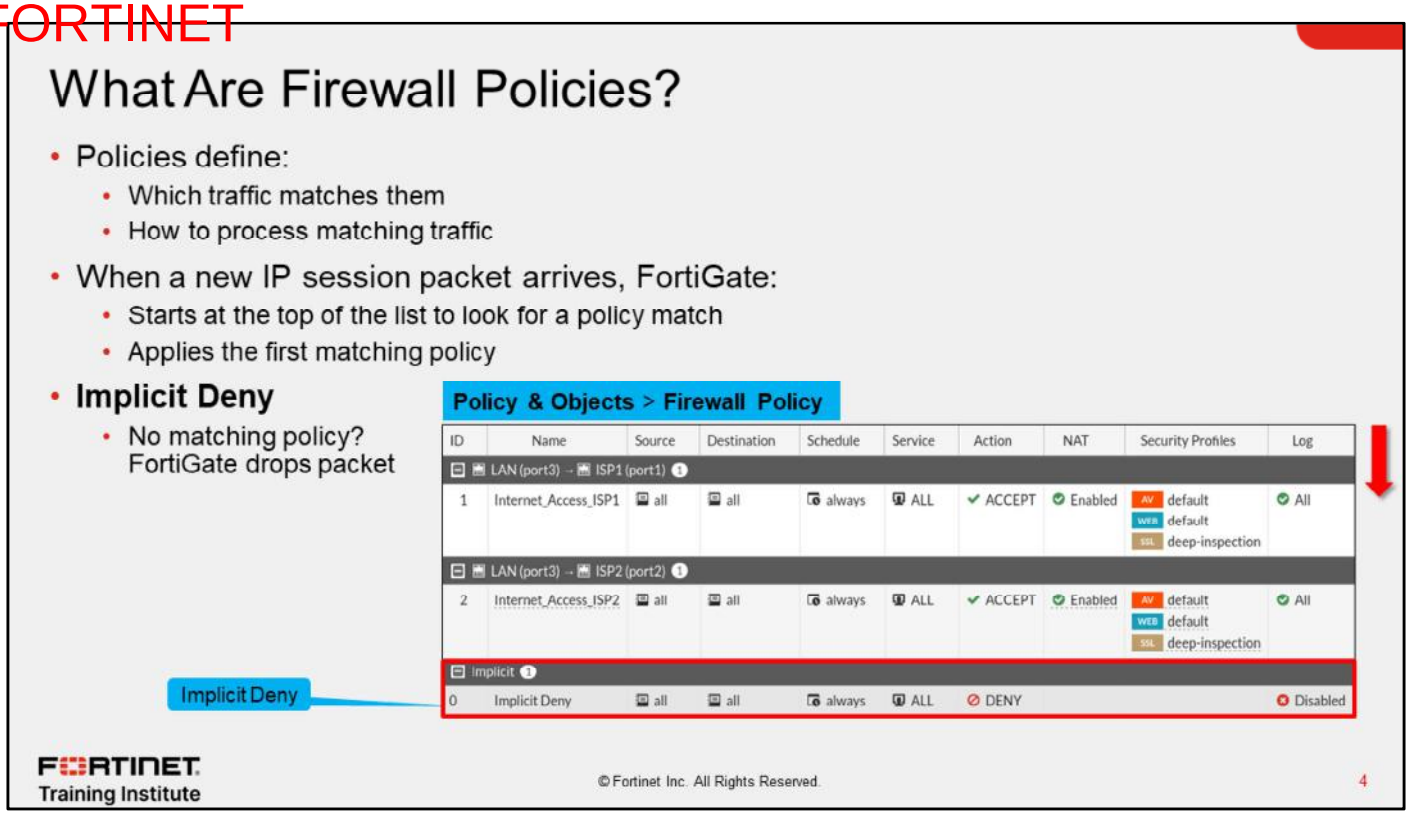

To begin, you will learn what firewall policies are.

Any traffic passing through a FortiGate must be associated with a firewall policy. A policy is a set of instructions that controls traffic flow through the FortiGate. These instructions determine where the traffic goes, how it's handled, and whether it's allowed to pass through the FortiGate. In summary, firewall policies are sets of rules that specify which traffic is allowed through the FortiGate and what FortiGate should do when traffic matches a policy.

Should the traffic be allowed? FortiGate bases this decision on simple criteria. FortiGate analyzes the source of the traffic, the destination IP address, and the service. If the policy does not block the traffic, FortiGate begins a more computationally expensive security profile inspection—often known as Unified Threat Management (UTM)—such as antivirus, application control, and web filtering, if you've chosen it in the policy. These inspections block the traffic if there is a security risk, for example, if the traffic contains a virus. Otherwise, the traffic is allowed.

Will network address translation (NAT) be applied? Is authentication required? Firewall policies also determine the answers to these questions. After processing is finished, FortiGate forwards the packet toward its destination.

FortiGate looks for the matching firewall policy from *top to bottom* and, if a match is found, the traffic is processed based on the firewall policy. If no match is found, the traffic is dropped by the default **Implicit Deny** firewall policy.

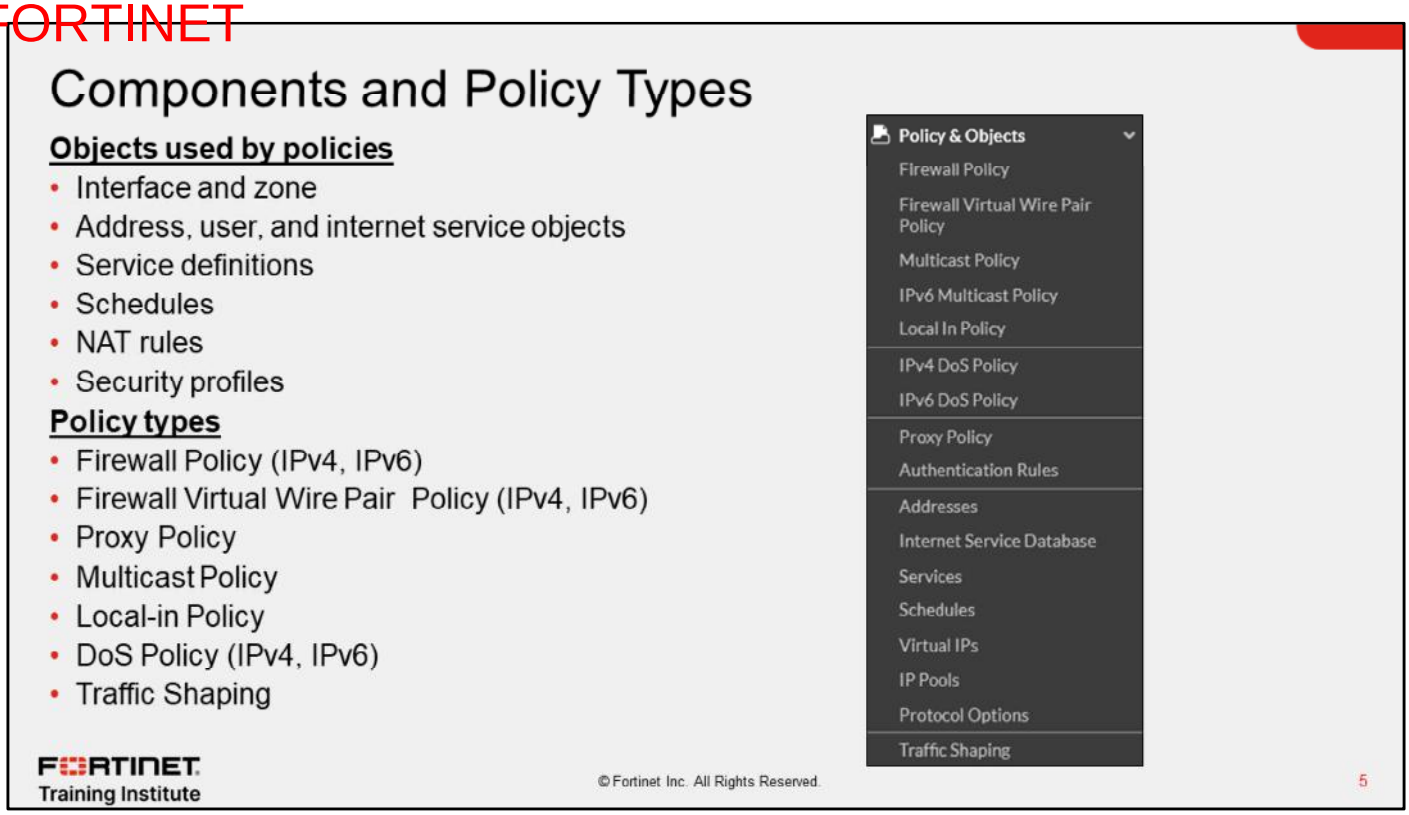

Each policy matches traffic and applies security by referring to the objects that you've defined, such as addresses and profiles.

Common policy types are:

- Firewall Policy: A firewall policy consists of set of rules that control traffic flow through FortiGate.
- Firewall Virtual Wire Pair Policy: A virtual wire pair policy is used to control the traffic between the interfaces in a virtual wire pair.
- Multicast Policy: A multicast policy allows multicast packets to pass from one interface to another.
- Local In Policy: A local-in policy controls the traffic to a FortiGate interface and can be used to restrict administrative access.
- DoS Policy: A denial-of-service (DoS) policy checks for the anomalous patterns in the network traffic that arrives at a FortiGate interface.

By default, only **Firewall Policy** is visible under **Policy and Object**. Other policies are available based on the interface configurations and advanced features enabled through **Feature Visibility**.
In this lesson, you will learn about IPv4 firewall policies, because they are the most commonly used policies.

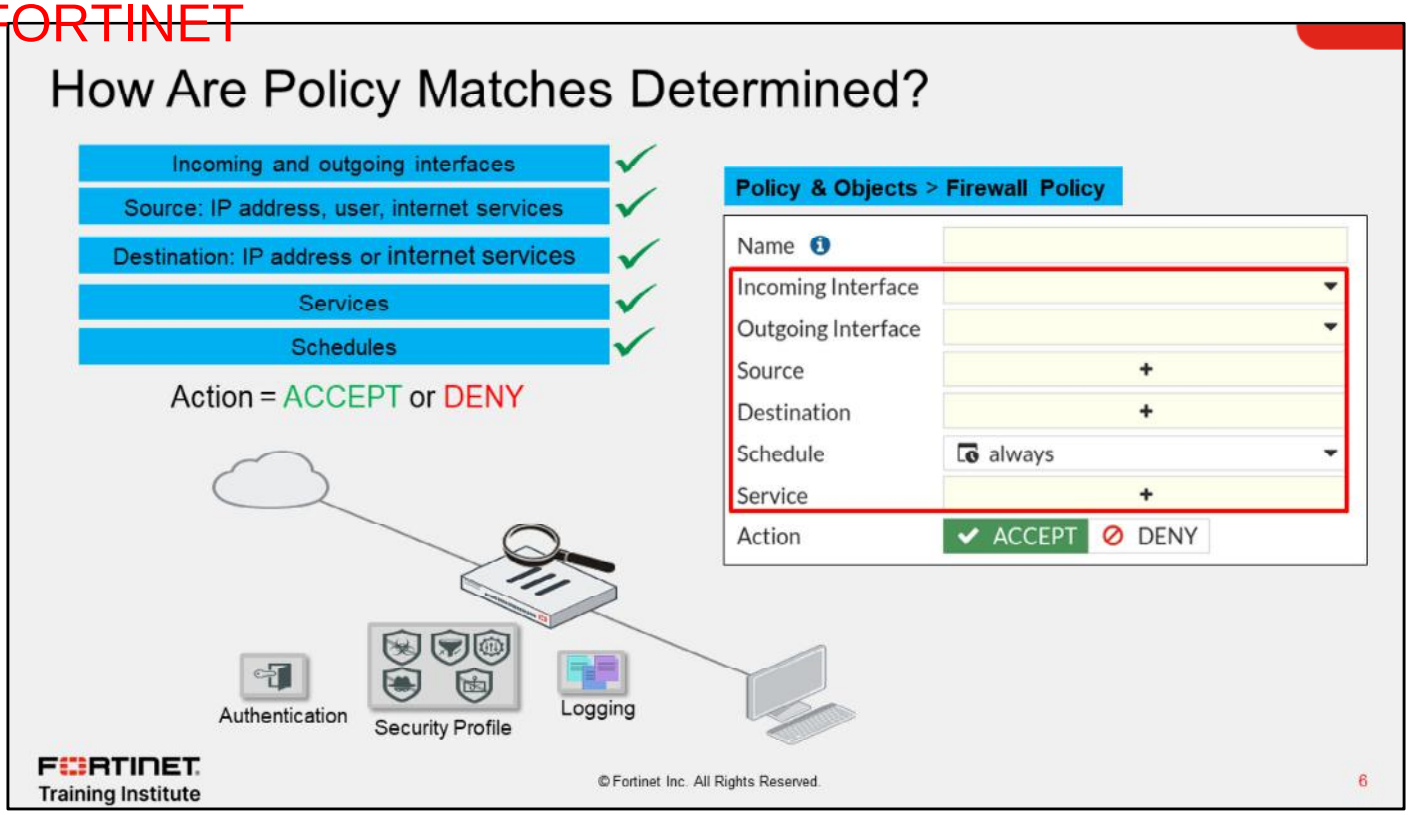

When a packet arrives, how does FortiGate find a matching policy? Each policy has match criteria, which you can define using the following objects:

- **Incoming Interface**
- **Outgoing Interface**
- **Source**: IP address, user, internet services
- **Destination**: IP address or internet services
- **Service**: IP protocol and port number
- **Schedule**: Specific times to apply policy

If the traffic matches a firewall policy, FortiGate applies the action configured in the firewall policy:

- If the **Action** is set to **DENY**, FortiGate drops the session.
- If the **Action** is set to **ACCEPT**, FortiGate allows the session and applies other configured settings for packet processing, such as user authentication, source NAT, antivirus scanning, web filtering, and so on.

When FortiGate receives traffic, it evaluates the packet's source IP address, destination IP address, and the requested service (protocol and port number). It also checks the incoming interface and the outgoing interface it needs to use. Based on this information, FortiGate identifies the firewall policy and evaluates the traffic. If the traffic matches the policy, then FortiGate applies the action (Accept/Deny) defined in the policy.
For example, to block incoming FTP traffic to all but a few FTP servers, define the addresses of the FTP servers as the destination, and select FTP as the service. You probably *wouldn't* specify a source (often any location on the internet is allowed) or schedule (FTP servers are usually always available, day or night). Finally, set the **Action** setting to **ACCEPT**.

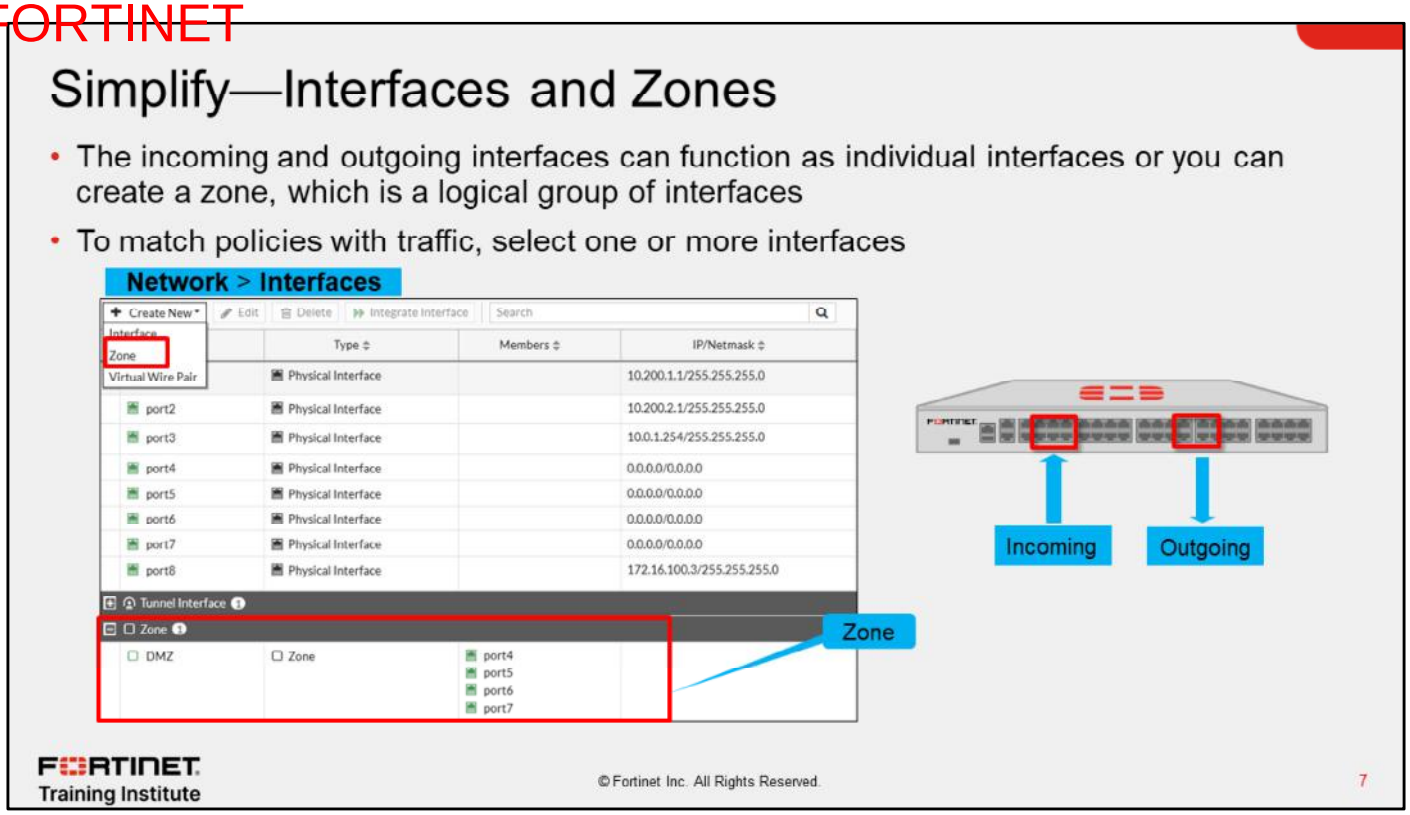

To begin describing how FortiGate finds a policy for each packet, let's start with the interfaces.

Packets arrive on an incoming, or ingress, interface. Routing determines the outgoing, or egress, interface. In each policy, you *must* set a source and destination interface; even if one or both are set to **any**. Both interfaces must match the policy's interface criteria in order to be a successful match.

For example, if you configure policies between port3 (LAN) ingress and port1 (WAN) egress and a packet arrives on port2, the packet will *not* match your policies and, therefore, would be dropped because of the implicit deny policy at the end of the list. Even if the policy is from port3 (LAN) ingress to any egress, the packet would still be dropped because it did not match the incoming interface.

To simplify policy configuration, you can group interfaces into logical zones. For example, you could group port4 to port7 as a DMZ. You can create zones on the **Interfaces** page. However, you should note that you cannot reference an interface in a zone individually, and, if you need to add the interface to the zone, you must remove all references to that interface (for example, firewall policies, firewall addresses, and so on). If you think you might need to reference interfaces individually, you should set multiple source and destination interfaces in the firewall policy, instead of using zones.

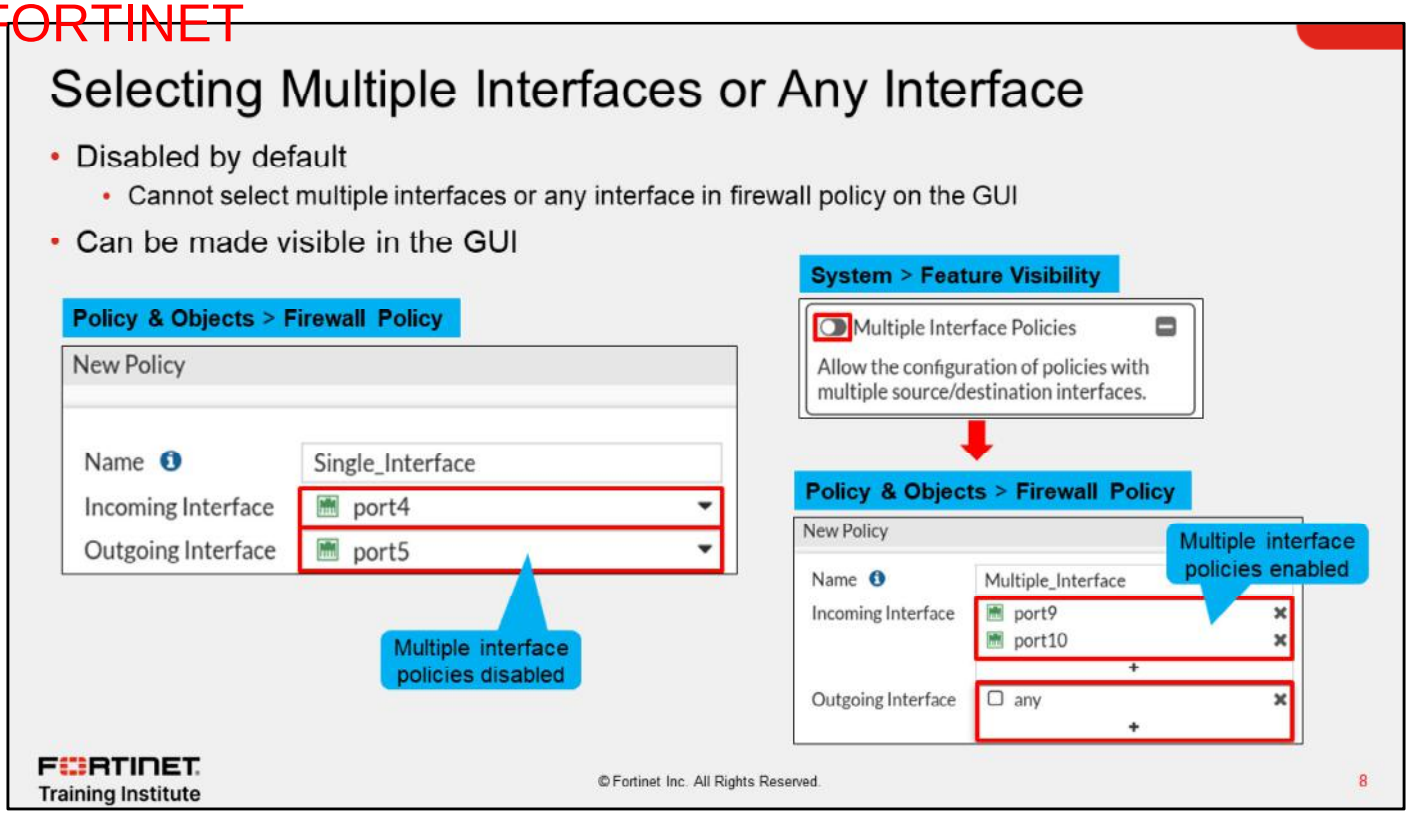

By default, you can select only a single interface as the incoming interface and a single interface as the outgoing interface. This is because the option to select multiple interfaces, or **any** interface in a firewall policy, is disabled on the GUI. However, you can enable the **Multiple Interface Policies** option on the **Feature Visibility** page to disable the single interface restriction.

You can also specify multiple interfaces, or use the `any` option, if you configure a firewall policy on the CLI, regardless of the default GUI setting.

It is also worth mentioning that when you choose the **any** interface option, you cannot select multiple interfaces for that interface. In the example shown on this slide, because **any** is selected as the outgoing interface, you cannot add any additional interfaces, because **any** interface implies that all interfaces have already been selected.

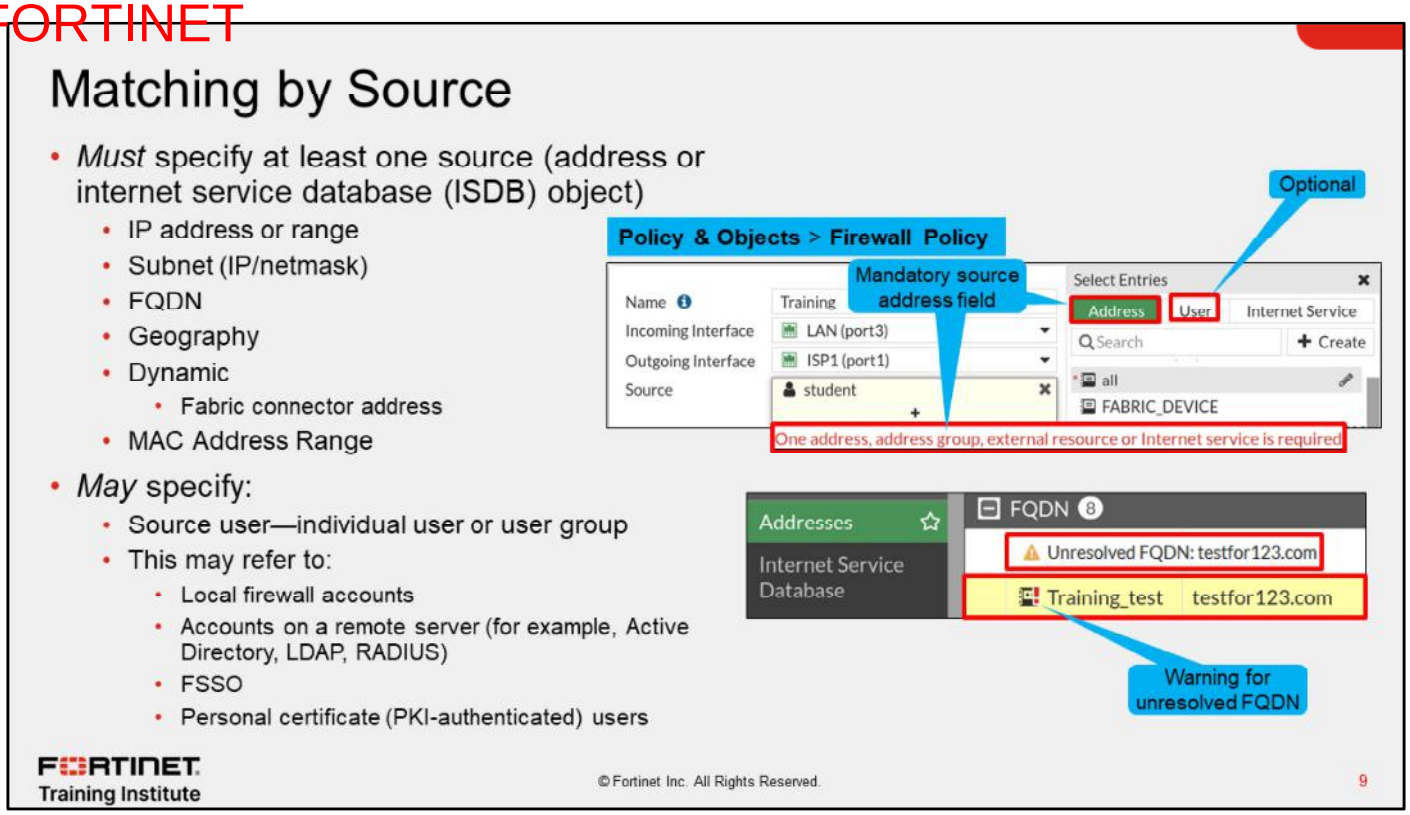

The next match criteria that FortiGate considers is the packet's source.

In each firewall policy, you *must* select a source address object. Optionally, you can refine your definition of the source address by *also* selecting a user, or a user group, which provides a much more granular match, for increased security. You can also select ISDB objects as the source in the firewall policy, which you will learn about later in this lesson.

When selecting a fully qualified domain name (FQDN) as the source address, it must be resolved by DNS and cached in FortiGate. Make sure FortiGate is configured properly for DNS settings. If FortiGate is not able to resolve an FQDN address, it will present a warning message, and a firewall policy configured with that FQDN may not function properly.

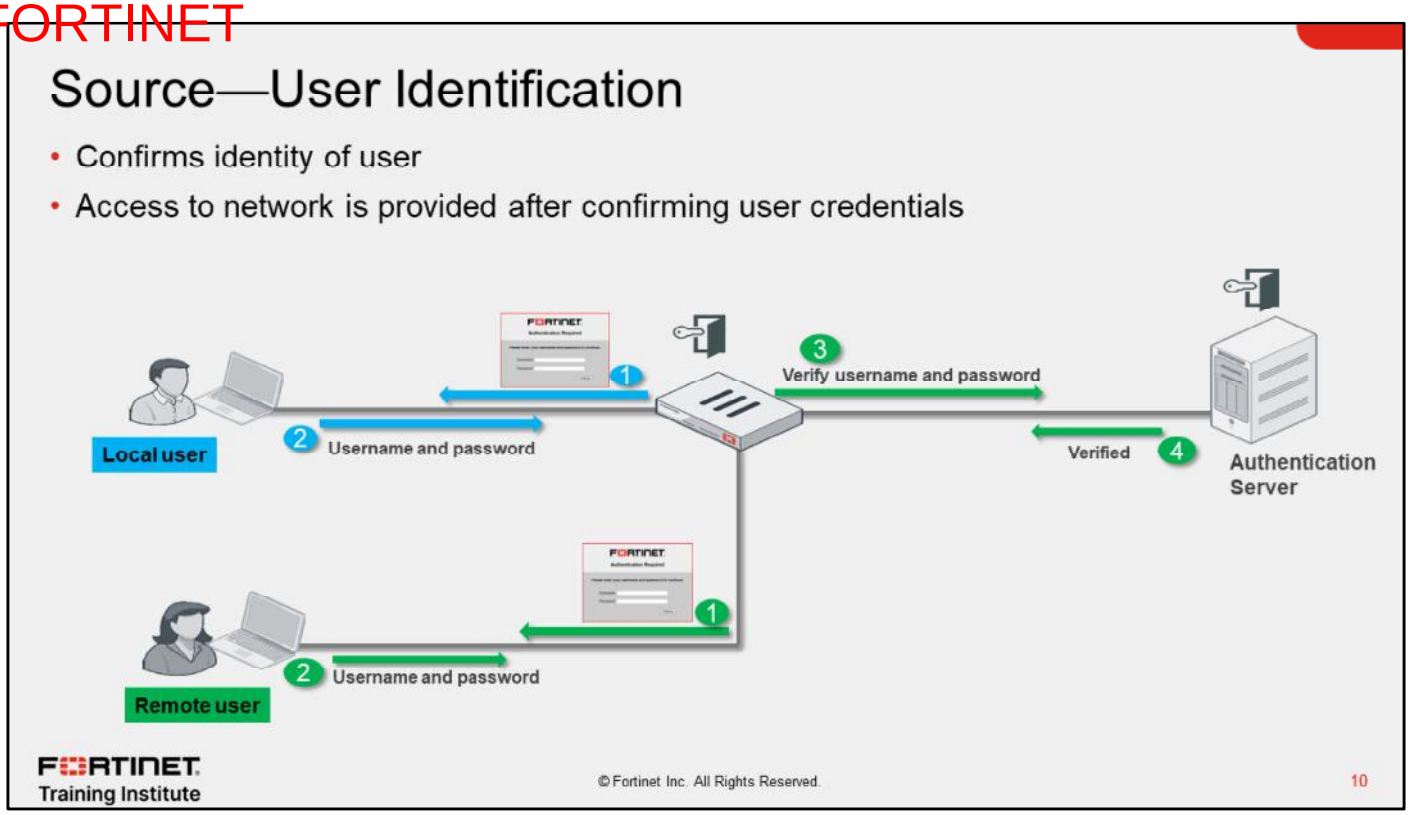

If a user is added as part of the source, FortiGate must verify the user before allowing or denying access based on the firewall policy. There are different ways that a user can authenticate.

For local users, the username and password is configured *locally* on FortiGate. When a local user authenticates, the credentials that they enter must match the username and password configured locally on FortiGate.

For a remote user (for example, LDAP or RADIUS), FortiGate receives the username and password from the remote user and passes this information to the authentication server. The authentication server verifies the user login credentials and updates FortiGate. After FortiGate receives that information, it grants access to the network based on the firewall policy.

A Fortinet single sign-on (FSSO) user's information is retrieved from the domain controller. Access is granted based on the group information on FortiGate.

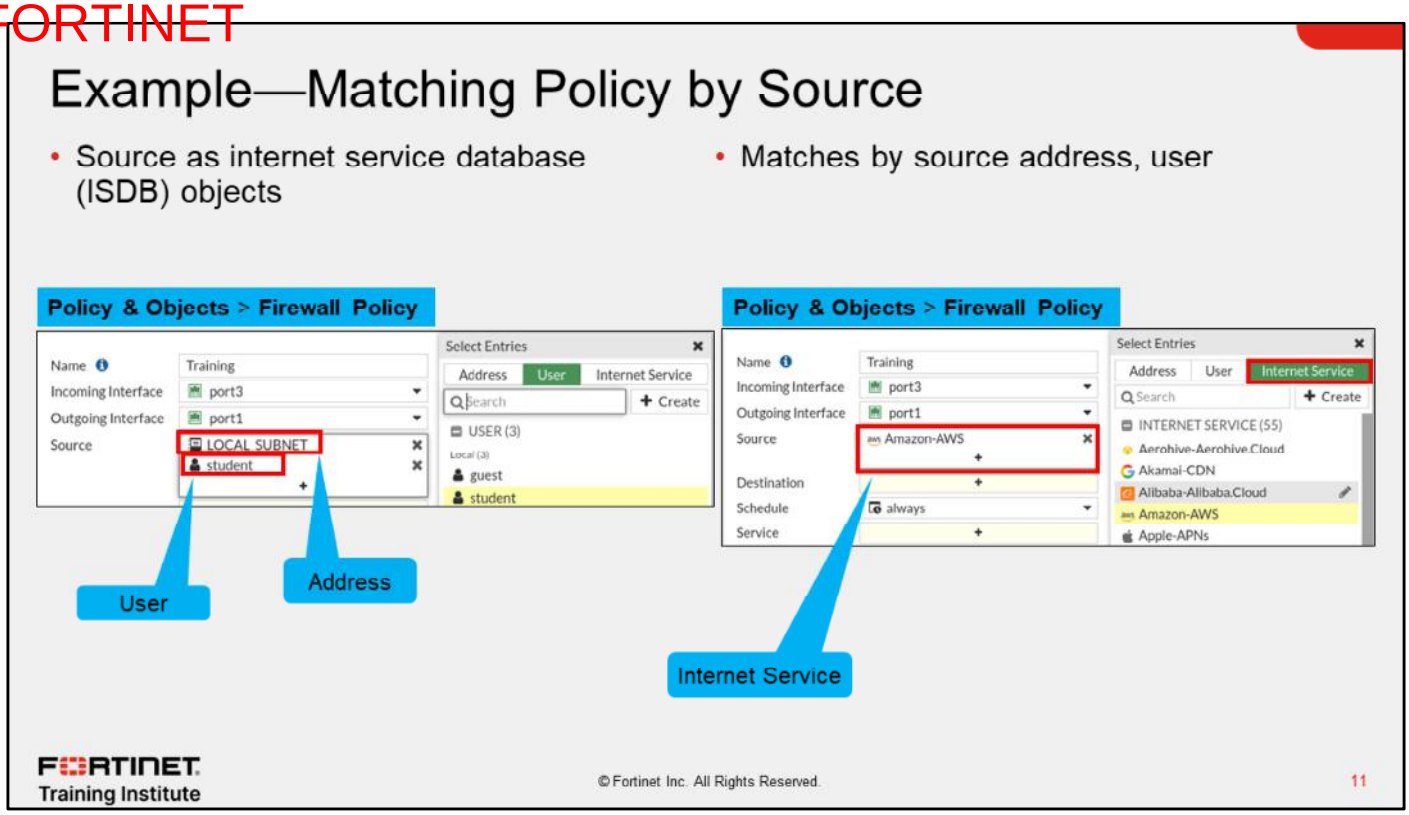

In the example shown on this slide, source selectors identify the specific subnet and user group. Remember, user is an optional object. The user object is used here to make the policy more specific. If you wanted the policy to match more traffic, you would leave the user object undefined.

You can also use internet service (ISDB) objects as a source in the firewall policy. There is an either/or relationship between internet service objects and source address objects in firewall policies. This means that you can select either a source address or an internet service, but not both.

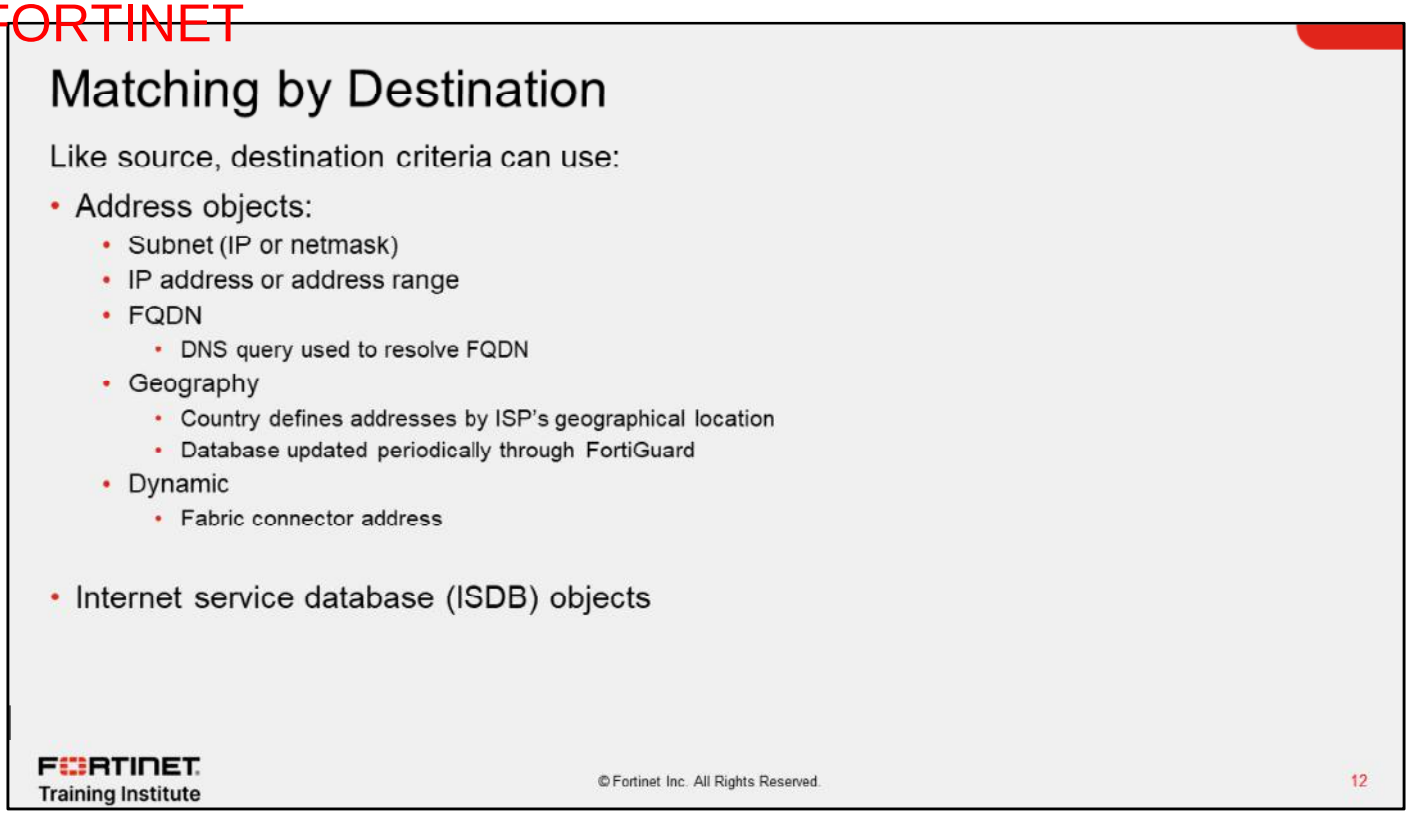

Like the packet's source, FortiGate also checks the destination address for a match.

You can use address objects or ISDB objects as destinations in the firewall policy. The address object may be a host name, IP subnet, or range. If you enter an FQDN as the address object, make sure that you've configured your FortiGate device with DNS servers. FortiGate uses DNS to resolve those FQDN host names to IP addresses, that appear in the IP header.

You can also choose geographic addresses, which are groups or ranges of addresses that are assigned to a country. FortiGuard is used to update these objects.

Why is there is no option to select a user? The user identification is determined at the ingress interface, and packets are forwarded only to the egress interface after the user is successfully authenticated.

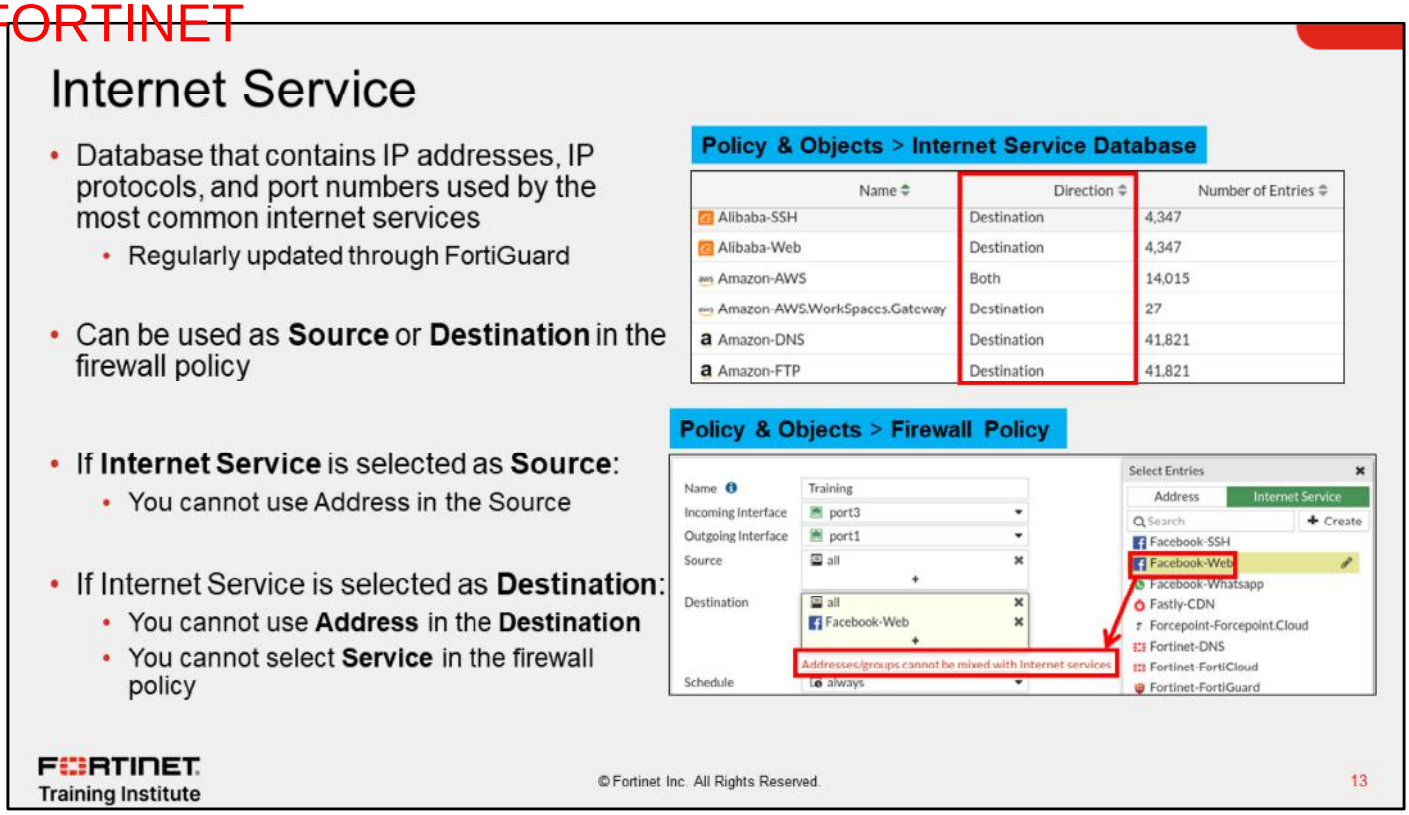

Internet Service is a database that contains a list of IP addresses, IP protocols, and port numbers used by the most common internet services. FortiGate periodically downloads the newest version of this database from FortiGuard. You can select these as **Source** or **Destination** in the firewall policy.

What happens if you need to allow traffic to only a few well-known public internet destinations, such as Dropbox or Facebook?

When configuring your firewall policy, you can use **Internet Service** as the destination in a firewall policy, which contains all the IP addresses, ports, and protocols used by that service. For the same reason, you cannot mix regular address objects with ISDB objects, and you cannot select services on a firewall policy. The ISDB objects already have services information, which is hardcoded.

Compared with address objects, which you need to check frequently to make sure that none of the IP addresses have changed or appropriate ports are allowed, internet services helps make this type of deployment easier and simpler.

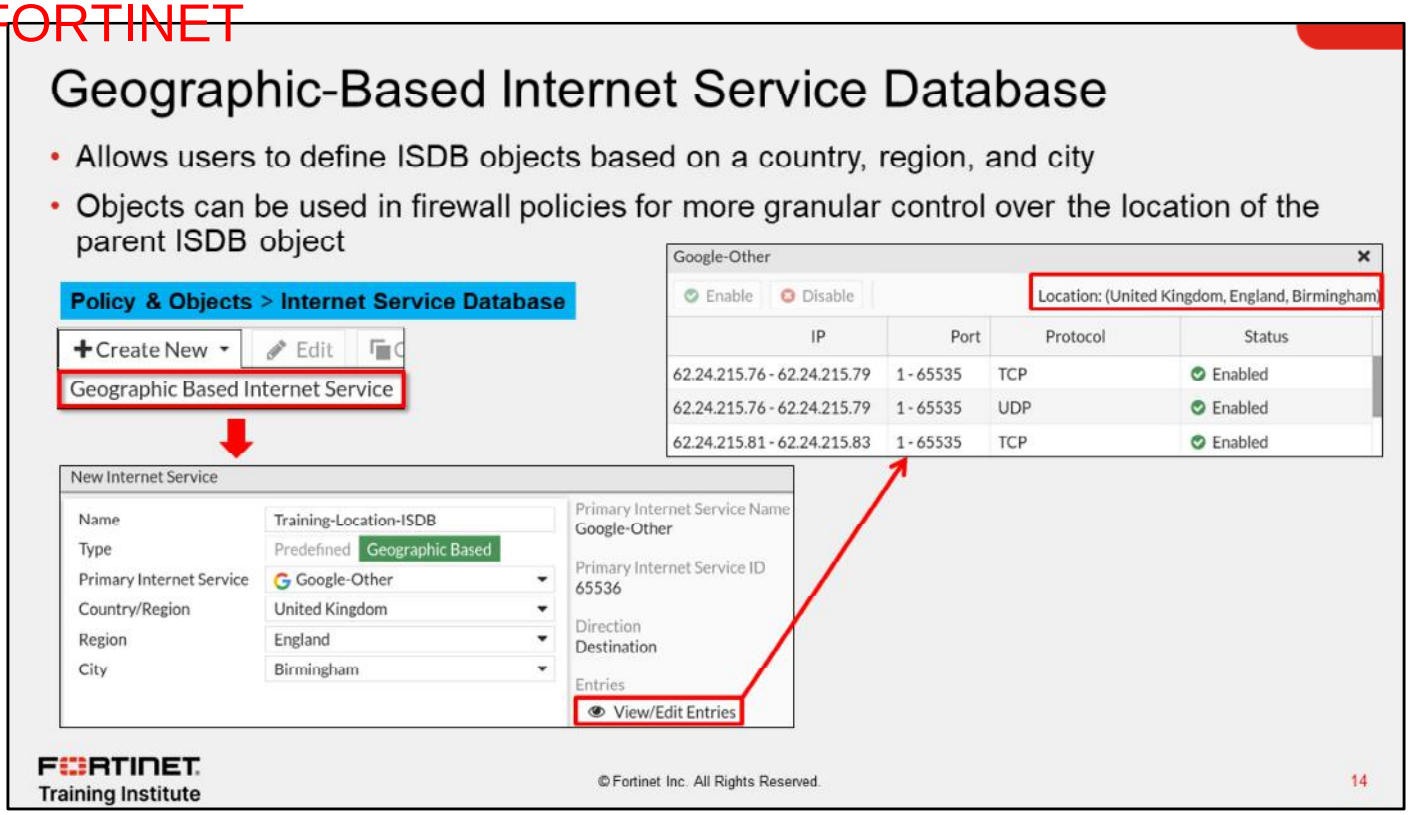

Geographic-based ISDB objects allow users to define a country, region, and city. These objects can be used in firewall policies for more granular control over the location of the parent ISDB object.

ISDB objects are referenced in policies by name, instead of by ID.

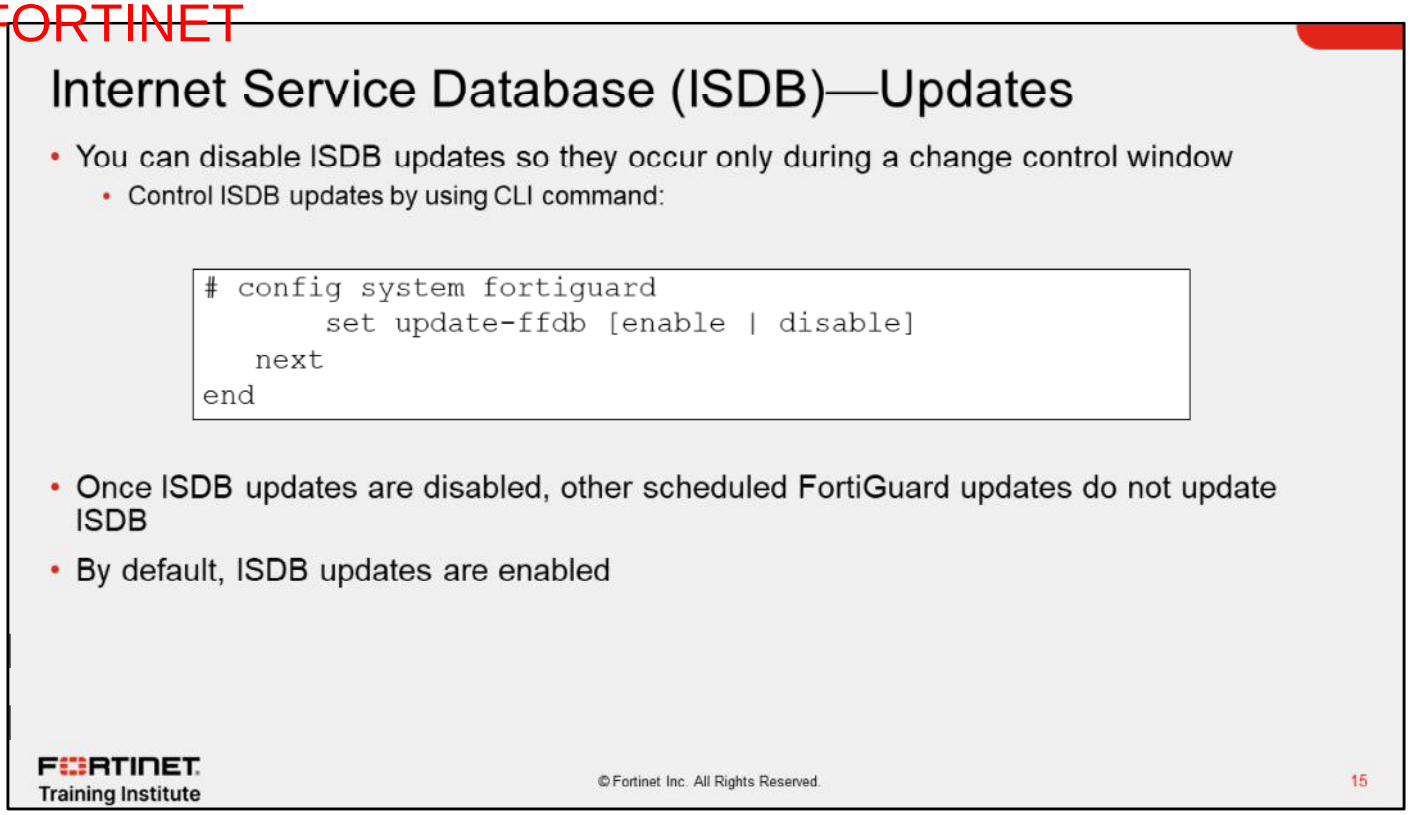

You can disable ISDB updates, so they occur only during a change control window. Once ISDB updates are disabled, other scheduled FortiGuard updates for IPS, AV, and so on, do not update ISDB. By default, ISDB updates are enabled.

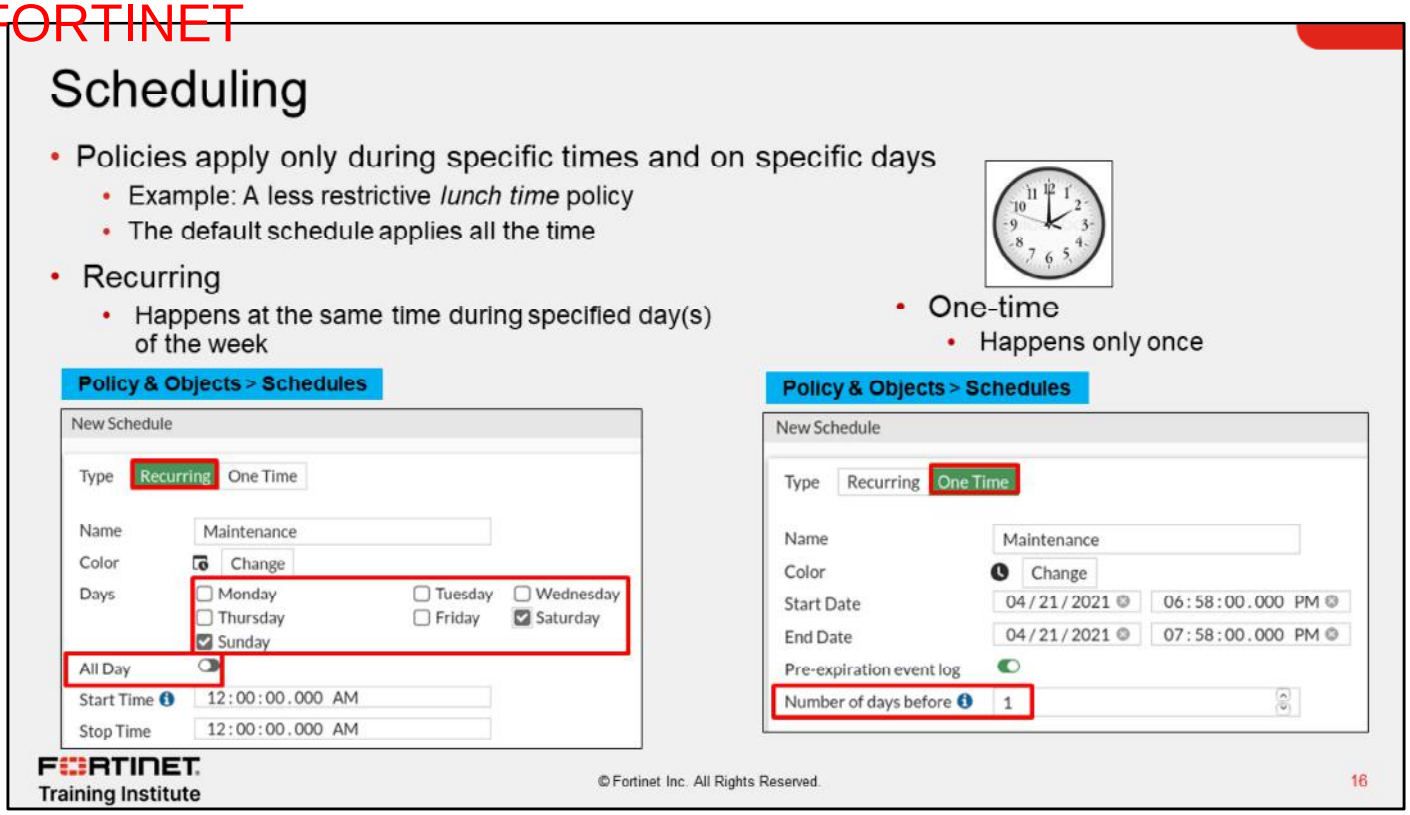

Schedules add a time element to the policy. For example, you might use a policy to allow backup software to activate at night, or create a test window for a remote address that is allowed for testing purposes.

Schedules can be configured and use a 24-hour time clock. There are a few configuration settings worth mentioning:

- **Recurring**: If you enable **All Day**, traffic will be allowed for 24 hours for the days selected. When configuring recurring schedules, if you set the stop time earlier than the start time, the stop time will occur the next day. For example, if you select Sunday as the day, 10:00 as the start time, and 09:00 as the stop time, the schedule will stop on Monday at 09:00. If the start and stop time are identical, the schedule will run for 24 hours.
- **One-time**: The start date and time must be earlier than the stop date and time. You can also enable **Pre-expiration event log**, which will generate an event log N number of days before the schedule expires, where N can be from 1 to 100 days.

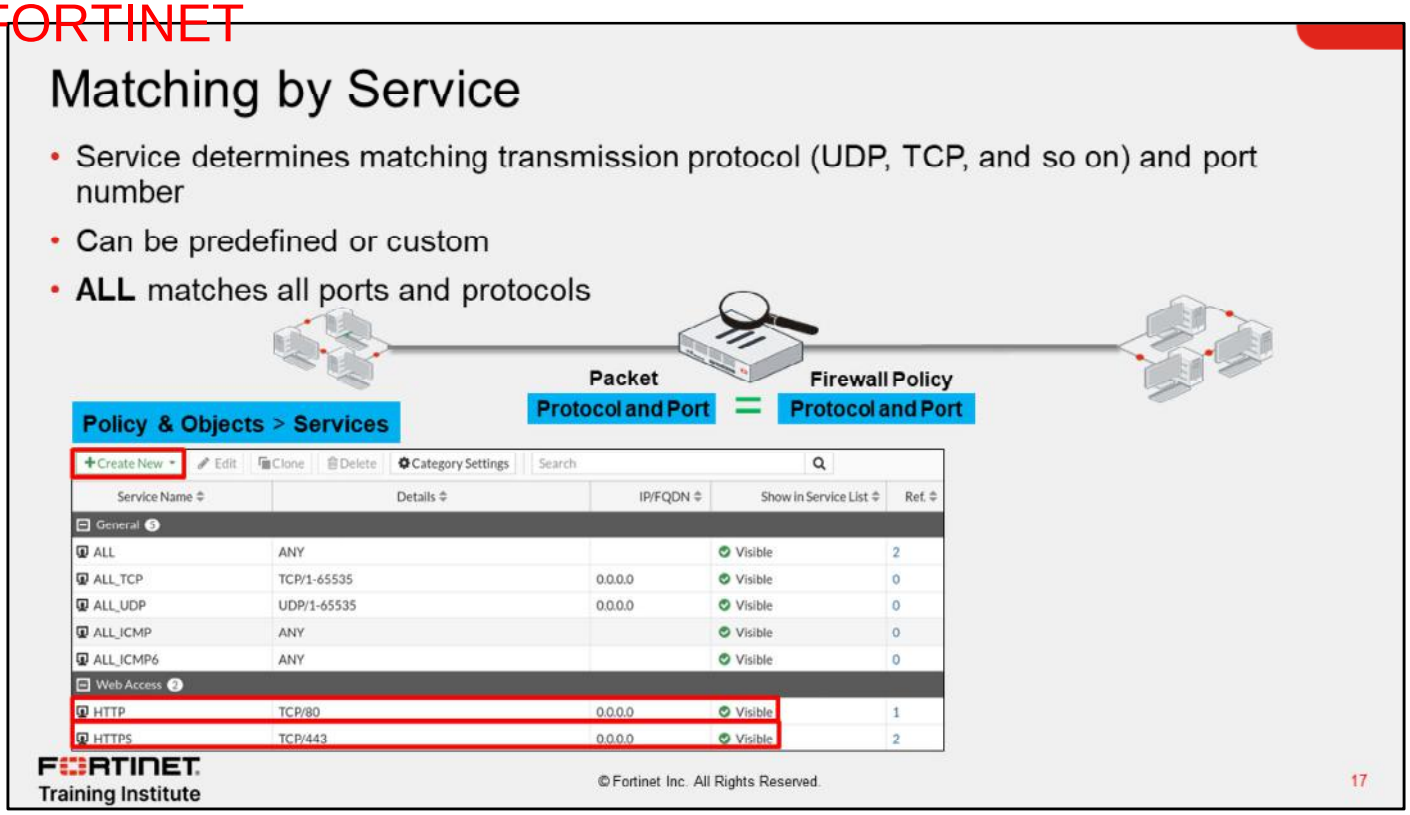

Another criterion that FortiGate uses to match policies is the packet's service.

At the IP layer, protocol numbers (for example, TCP, UDP, SCTP, and so on) together with source and destination ports, define each network service. Generally, only a destination port (that is, the server's listening port) is defined. Some legacy applications may use a specific source port, but in most modern applications, the source port is randomly identified at transmission time, and therefore is not a reliable way to define the service.

For example, the predefined service object named HTTP is TCP destination port 80, and the predefined service object named HTTPS is TCP destination port 443. However, the source ports last for only a short time and, therefore, are not defined.

By default, services are grouped together to simplify administration by categories. If the predefined services don't meet your organizational needs, you can create one or more new services, service groups, and categories.

Knowledge Check

1. What criteria does FortiGate use to match traffic to a firewall policy?
   - ✓A. Source and destination interfaces
   - B. Security profiles

2. What must be selected in the **Source** field of a firewall policy?
   - ✓A. At least one address object or ISDB
   - B. At least one source user and one source address object

FORTINET Training Institute

18

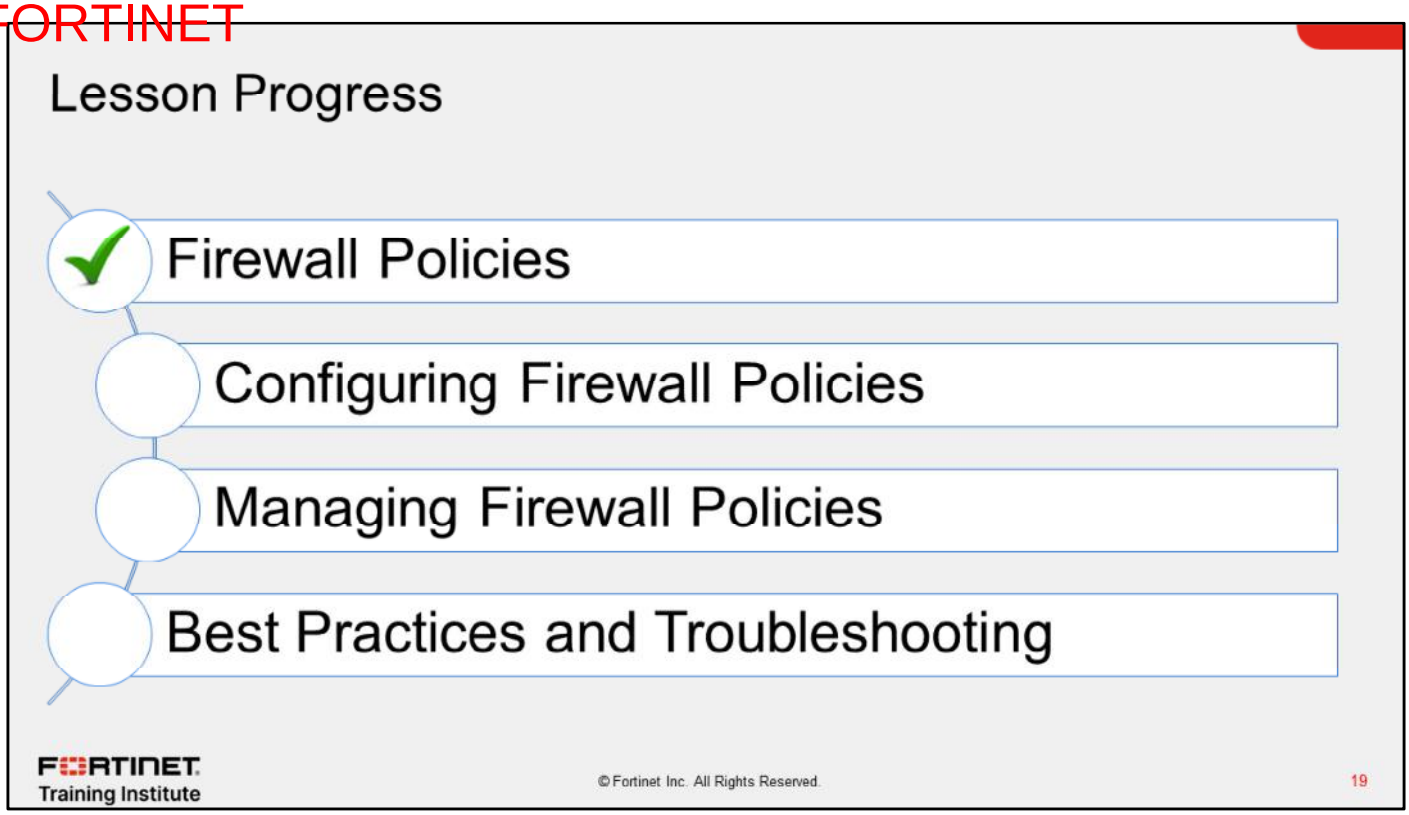

Good job! You now understand the components used in firewall policies and matching criteria used by FortiGate.

Now, you'll learn how to configure firewall policies.

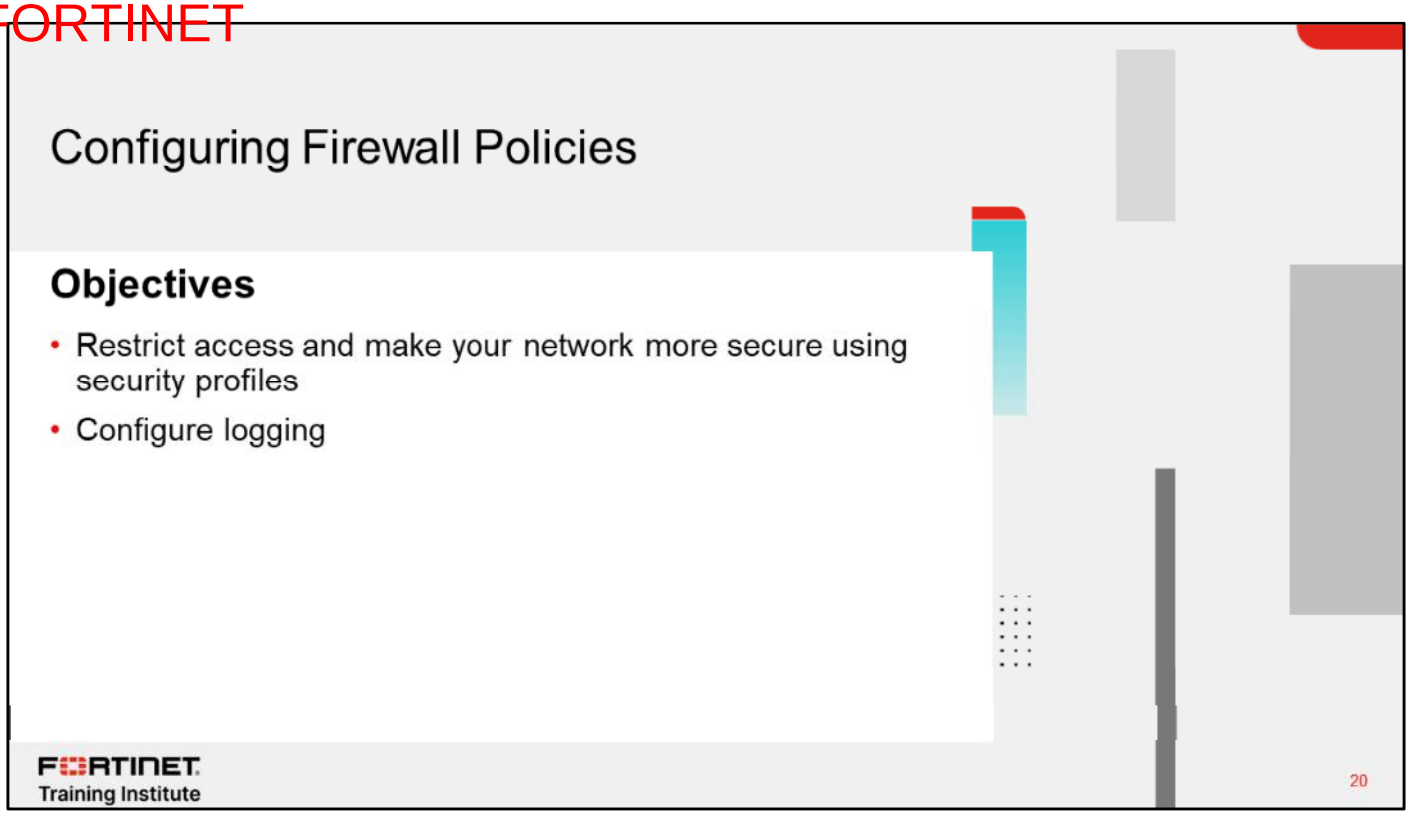

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in configuring firewall policies, you will be able to apply the correct settings, such as security profiles, logging, and traffic shaping, to firewall policies on FortiGate, and make your network more secure.

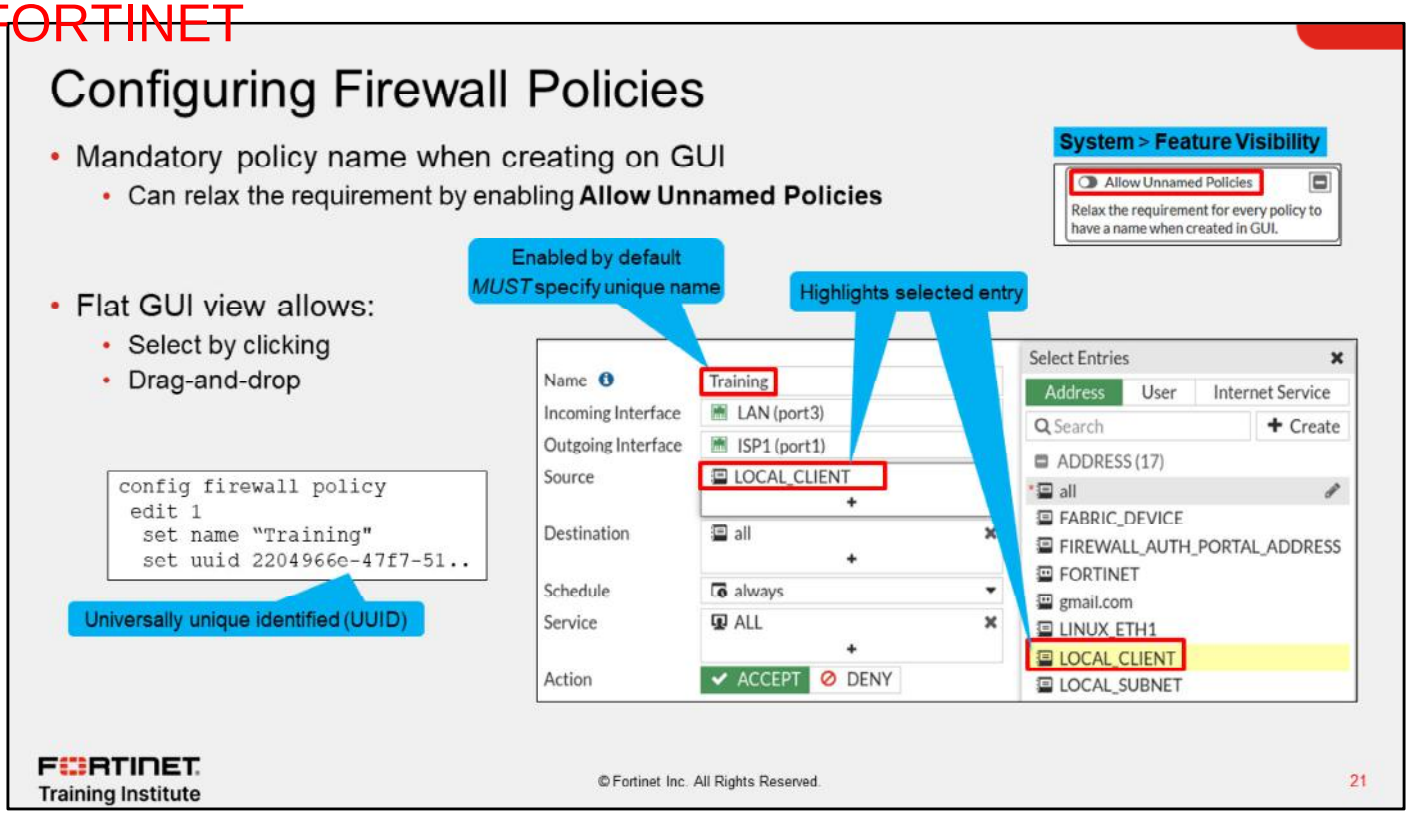

When you configure a new firewall policy on the GUI, you *must* specify a unique name for the firewall policy because it is enabled by default, while it is optional on the CLI. This helps the administrator to quickly identify the policy that they are looking for. However, you can make this feature optional on the GUI on the **Feature Visibility** page by enabling **Allow Unnamed Policies**.

Note that if a policy is configured without a policy name on the CLI, and you modify that existing policy on the GUI, you *must* specify a unique name. The FortiGate flat GUI view allows you to select interfaces and other objects by clicking or dragging and dropping from the list populated on the right side.

You can select **Internet Service** as the source. **Internet Service** is a combination of one or more addresses and one or more services associated with a service found on the internet, such as an update service for software.

There are many other options that you can configure in the firewall policy, such as firewall and network options, security profiles, logging options, and enabling or disabling a policy.

When creating firewall objects or policies, a universally unique identifier (UUID) attribute is added so that logs can record these UUIDs and improve functionality when integrating with FortiManager or FortiAnalyzer.

When creating firewall policies, remember that FortiGate is a stateful firewall. As a result, you need to create only one firewall policy that matches the direction of the traffic that initiates the session. FortiGate will automatically remember the source-destination pair and allow replies.

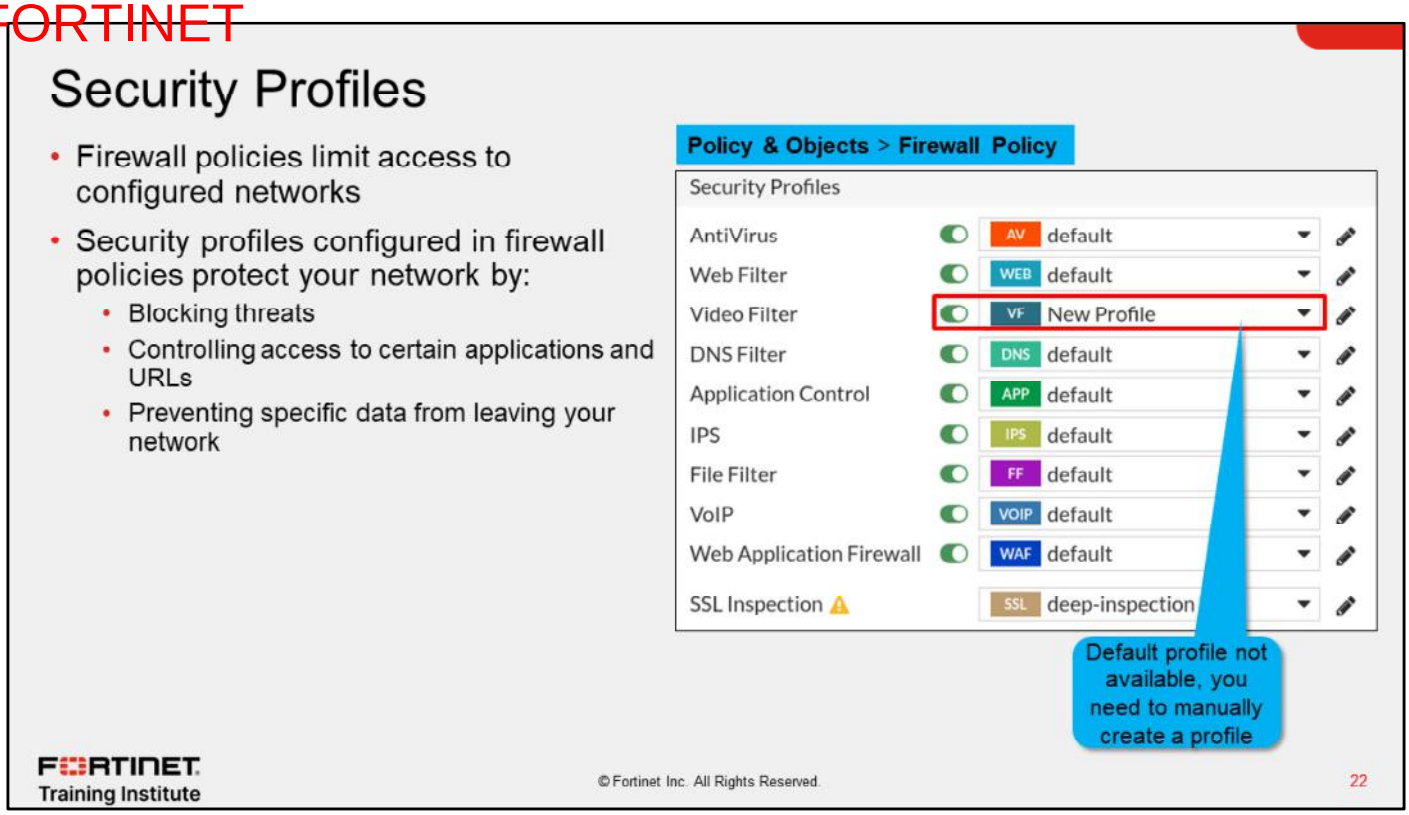

One of the most important features that a firewall policy can apply is security profiles, such as IPS and antivirus. A security profile inspects each packet in the traffic flow, where the session has already been conditionally accepted by the firewall policy.

When inspecting traffic, FortiGate can use one of two methods: flow-based inspection or proxy-based inspection. Different security features are supported by each inspection type.

Note that by default, the **Video Filter**, **VOIP**, **and Web Application Firewall** security profile option is not visible in the policy page on the GUI. You need to enable them on the **Feature Visibility** page.

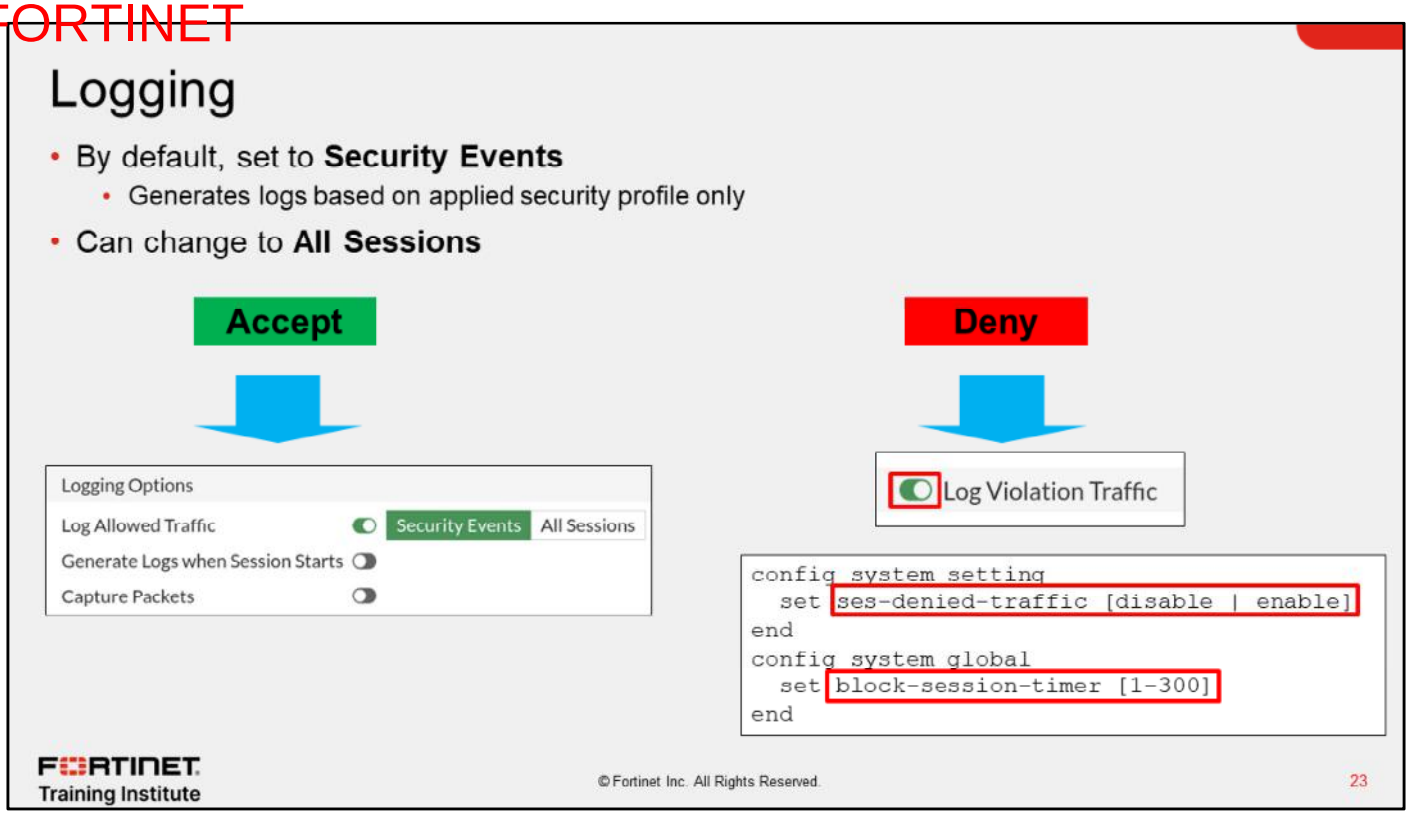

If you have enabled logging in the policy, FortiGate generates traffic logs after a firewall policy closes an IP session.

By default, **Log Allowed Traffic** is enabled and set to **Security Events** and generates logs for only the applied security profiles in the firewall policy. However, you can change the setting to **All Sessions**, which generates logs for all sessions.

If you enable **Generate Logs when Session Starts**, FortiGate creates a traffic log when the session begins. FortiGate also generates a second log for the same session when it is closed. But remember that increasing logging decreases performance, so use it only when necessary.

During the session, if a security profile detects a violation, FortiGate records the attack log immediately. To reduce the number of log messages generated and improve performance, you can enable a session table entry of dropped traffic. This creates the denied session in the session table and, if the session is denied, all packets of that session are also denied. This ensures that FortiGate does not have to do a policy lookup for each new packet matching the denied session, which reduces CPU usage and log generation.

This option is in the CLI, and is called `ses-denied-traffic`. You can also set the duration for block sessions. This determines how long a session will be kept in the session table by setting `block-session-timer` in the CLI. By default, it is set to 30 seconds.
If the GUI option **Generate Logs when Session Starts** is not displayed, this means that your FortiGate device does not have internal storage. This option is on the CLI, regardless of internal storage, and is called `set logtraffic-start enable`.

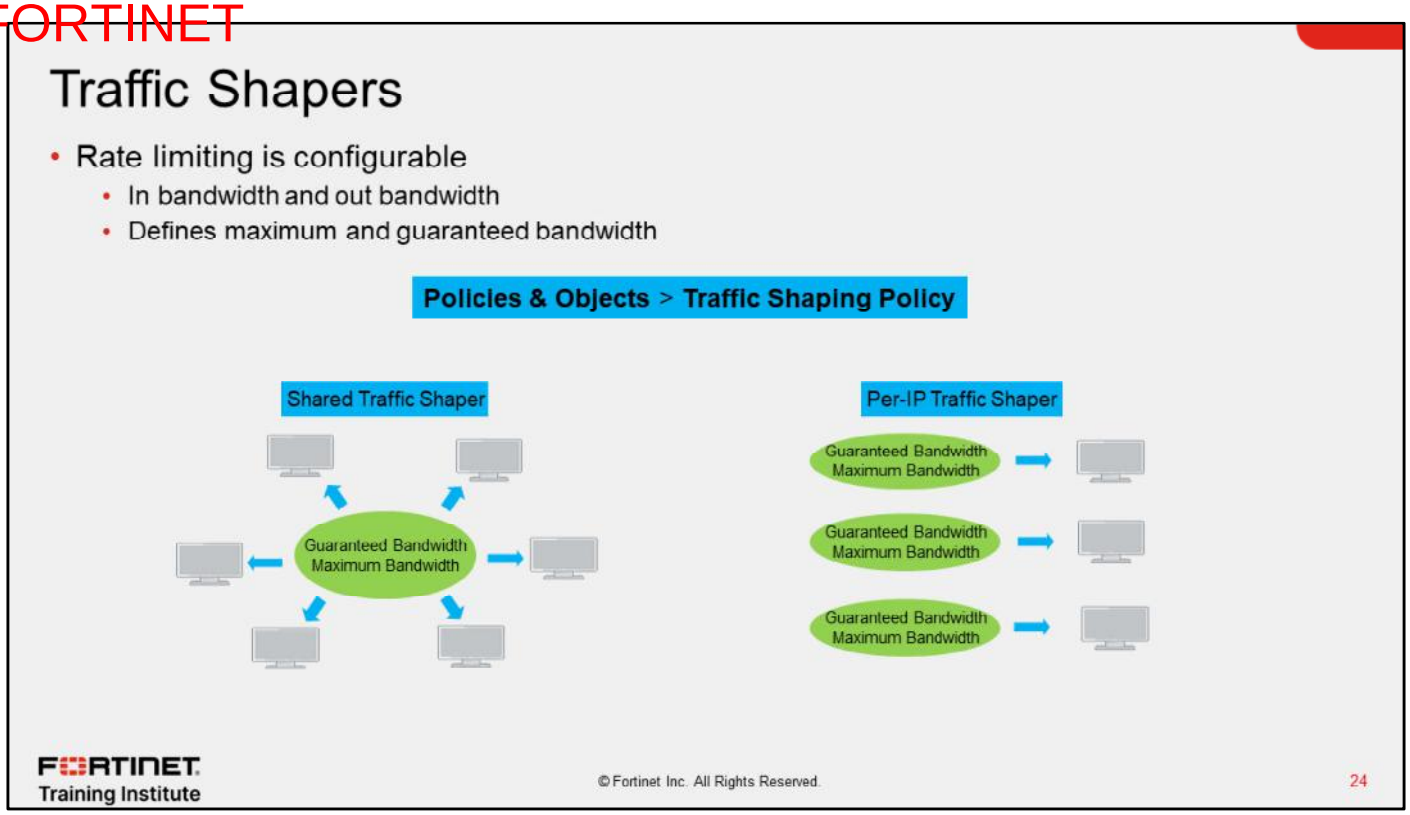

You can configure two types of traffic shapers: shared and per IP.

A shared shaper applies a total bandwidth to all traffic using that shaper. The scope can be per policy or for all policies referencing that shaper. FortiGate can count the packet rates of ingress and egress to police traffic.

FortiGate allows you to create three types of traffic shaping policies:

- Shared policy shaping: bandwidth management of security policies
- Per-IP shaping: bandwidth management of user IP addresses
- Application control shaping: bandwidth management by application

When creating traffic shaping policies, you must ensure that the matching criteria is the same as the firewall policies you want to apply shaping to. Note that these apply equally to TCP and UDP, and UDP protocols may not recover as gracefully from packet loss.

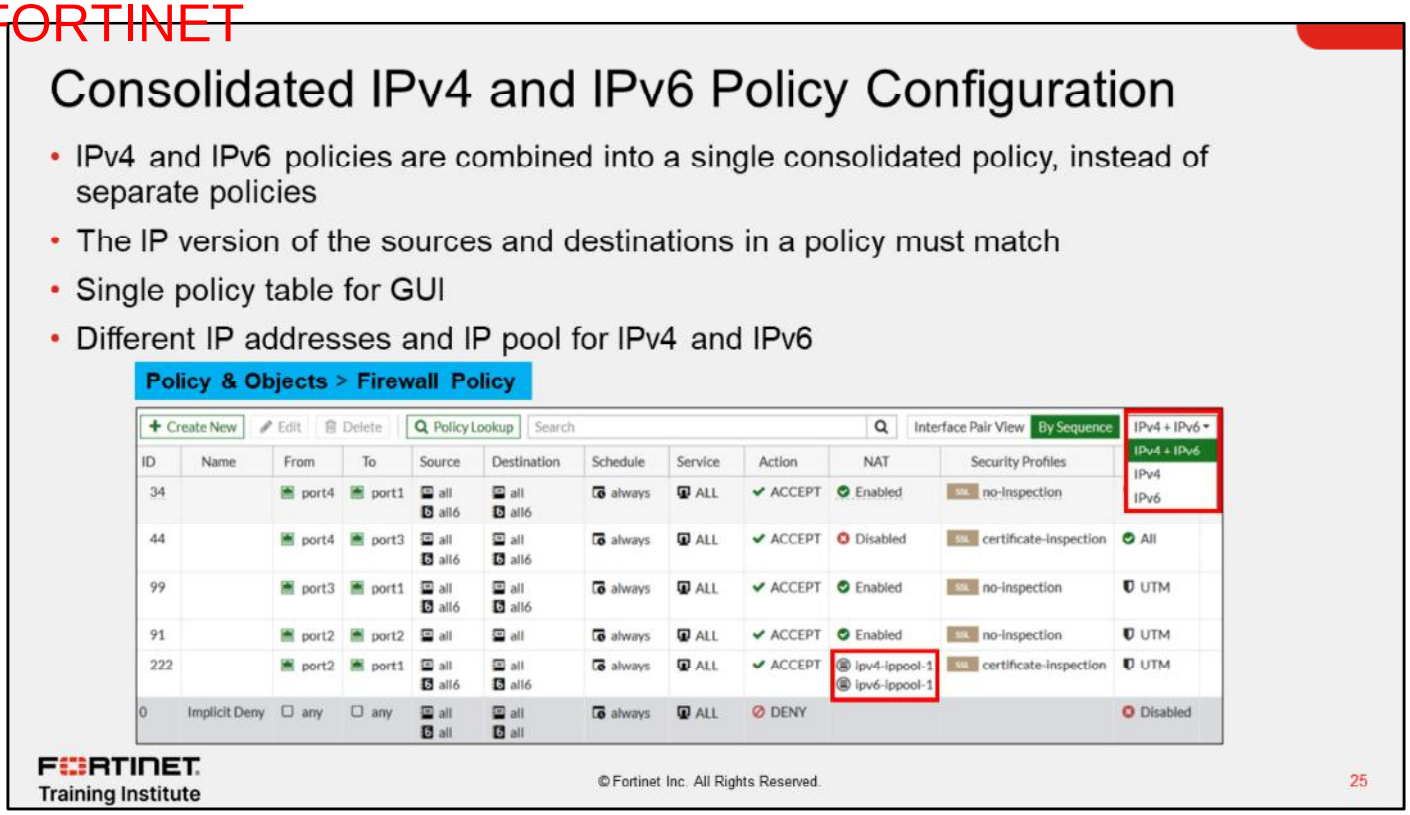

By default, IPv4 and IPv6 policies are combined into a single consolidated policy, rather than creating and maintaining two different policy sets for IPv4 and IPv6.

You can share the **Incoming Interface**, **Outgoing Interface**, **Schedule**, and **Service** fields with both IPv4 and IPv6. For source addresses, destination addresses, and IP pool, you must select addresses for both IPv4 and IPv6.
While configuring a consolidated firewall policy, you can configure a policy with IPv4 source addresses, IPv4 destination addresses, and an IPv4 IP pool, without specifying any IPv6 references. You can also configure the policy with the same behavior for IPv6. However, if you want to combine IPv4 and IPv6, you must select both IPv4 addresses and IPv6 addresses in the **Source** and **Destination** address fields in the firewall policy.The IP version of the sources and destinations in a policy must match. For example, a policy cannot have only an IPv4 source and an IPv6 destination. The policy table in the GUI can be filtered to show policies with IPv4, IPv6, or IPv4 and IPv6 sources and destinations.

Note that, by default, the **IPv6** option is not visible in the policy table on the GUI. You must enable **IPv6** on the **Feature Visibility** page.

Knowledge Check

1. To configure a firewall policy, you must include a firewall policy name when configuring using the _____.
   - A. CLI
   - ✓ B. GUI

2. What is the purpose of applying security profiles to a firewall policy?
   - A. To allow access to specific subnets
   - ✓ B. To protect your network from threats, and control access to specific applications and URLs

FORTINET Training Institute

26

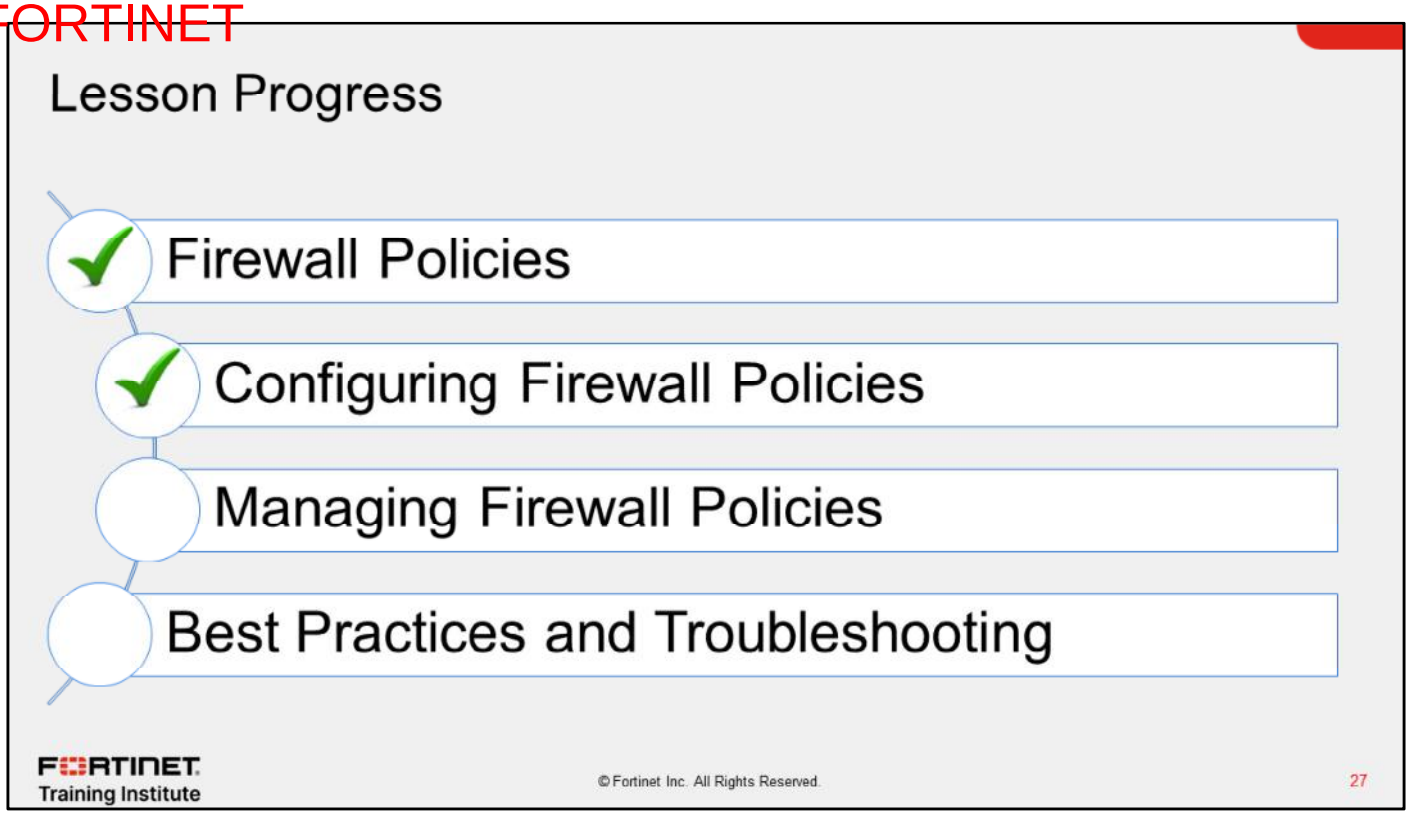

Good job! You now understand how to configure firewall policies on FortiGate.

Next, you'll learn how to manage and fine-tune settings for firewall policies.

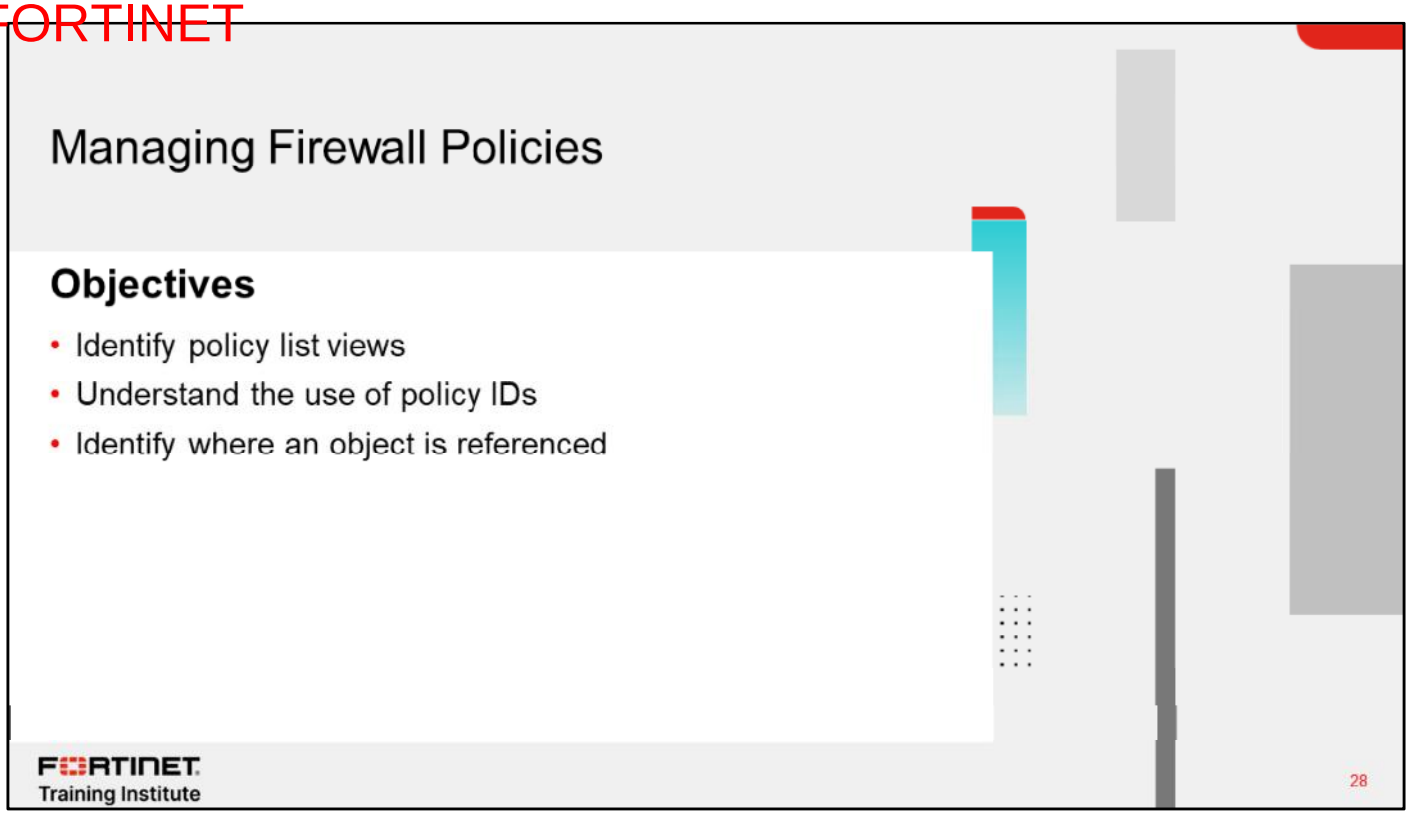

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in managing firewall policies, you will be able to understand the use of the policy ID of a firewall policy. Also, you will be able to pinpoint object usage, and simplify policies using object groups.

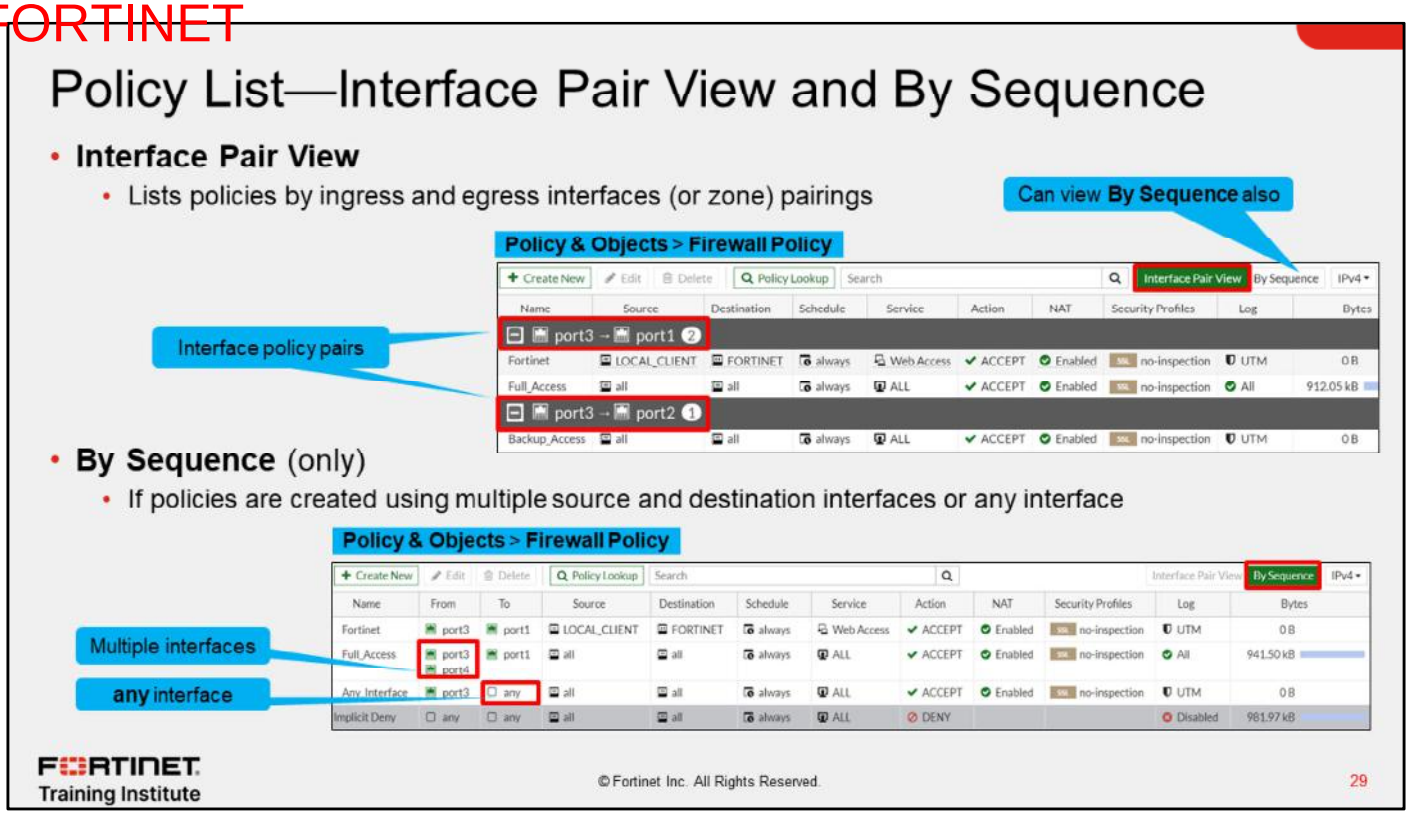

Firewall policies appear in an organized list. The list is organized either in **Interface Pair View** or **By Sequence**.

By default, the policy list appears in **Interface Pair View**. Each section contains policies in the order that they are evaluated for matching traffic and are arranged by ingress-egress interface pair. Alternatively, you can view your policies as a single, comprehensive list by selecting **By Sequence** at the top of the page. In this view the policies are also listed in the order in which they are evaluated for traffic matching, but they are not grouped.

In some cases, you cannot choose the view. For example, if you use multiple source or destination interfaces, or the **any** interface in a firewall policy, you cannot separate policies into sections by interface pairs—some would be triplets or more. In this case, policies always appear in a single list (**By Sequence**).

To help you remember the use of each interface, you add aliases by editing the interface on the **Network** page. For example, you could call port1 *ISP1*. This can help to make your list of policies easier to understand.

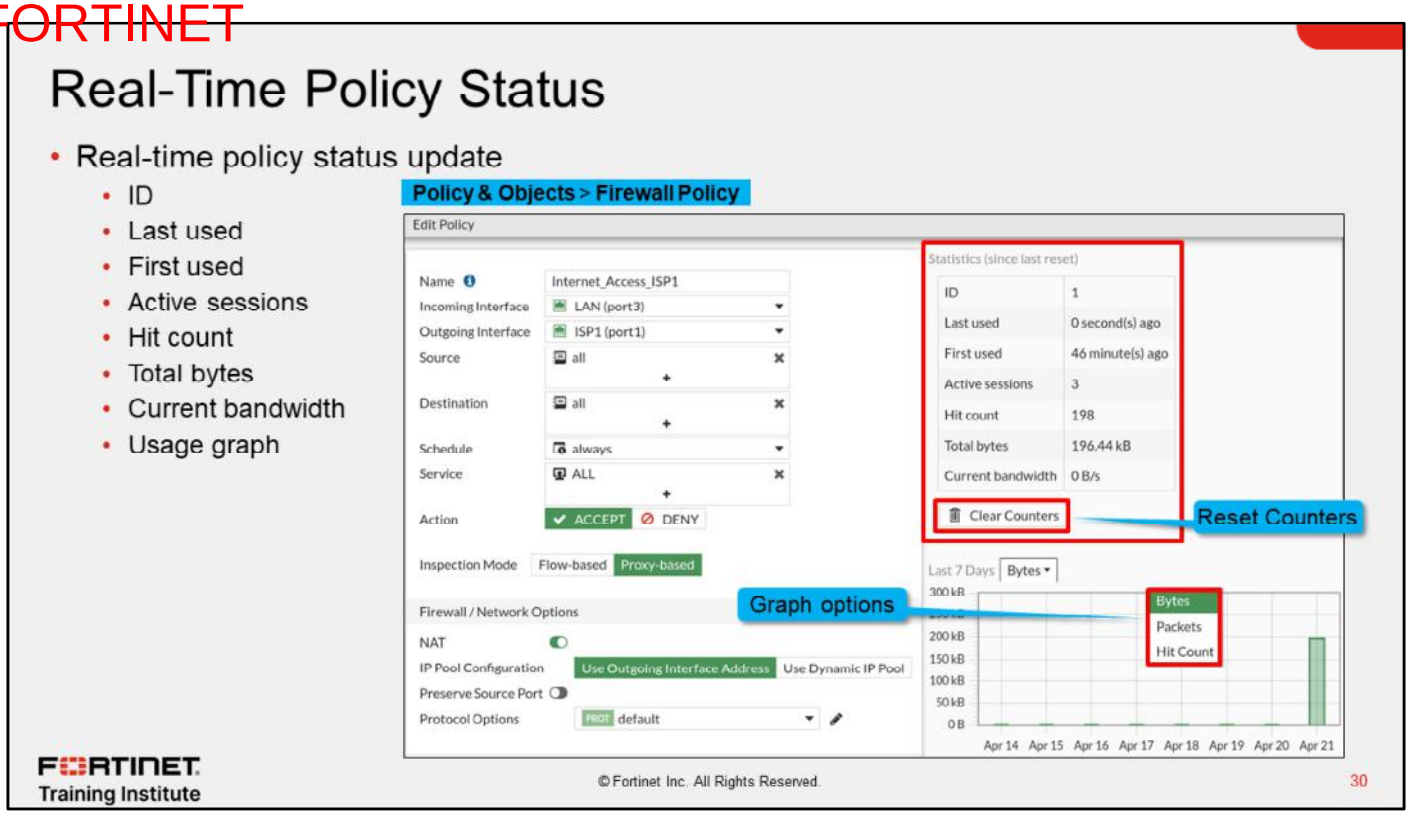

When you edit the policy, policy information will be visible.

This feature is very useful if an administrator wanted to check the policy usage, such as last used, first used, hit count, active sessions, and so on.

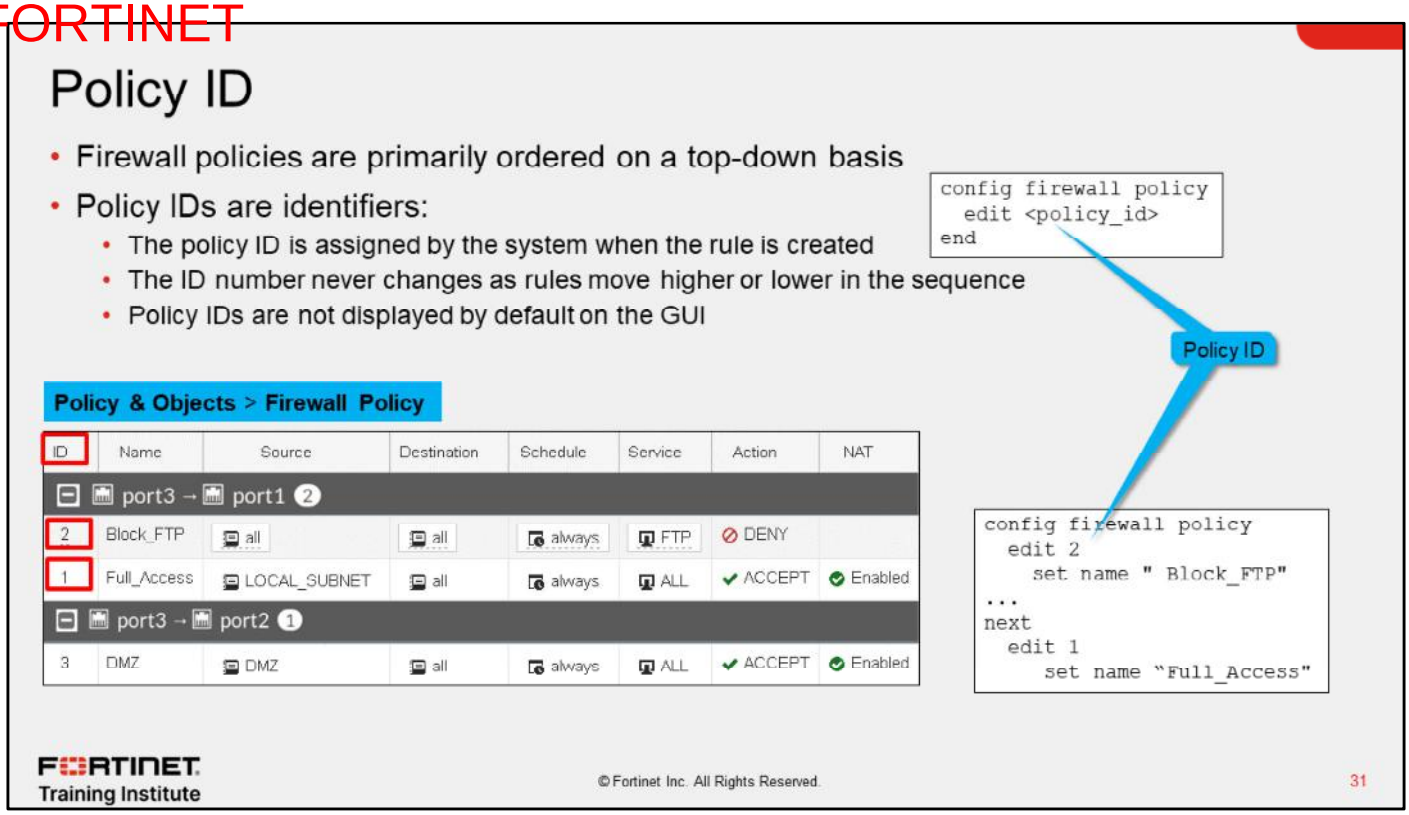

An important concept to understand about how firewall policies work is the precedence of order, or, if you prefer a more recognizable term, first come, first served.

Policy IDs are identifiers. By default, policy IDs are not displayed on the policy list GUI. You can add a policy **ID** column using the **Configure Table** settings icon.

FortiGate automatically assigns a policy ID when you create a new firewall policy on the GUI. The policy ID never changes, even if you move the rule higher or lower in the sequence.

If you enable **Policy Advanced Options,** then you can manually assign a policy ID, while creating a new policy. If a duplicate entry is found, the system produces an error, so you can assign a different available policy ID number.

Policy Advanced Options is not available on the GUI by default, you must enable it on the **Feature Visibility** page.

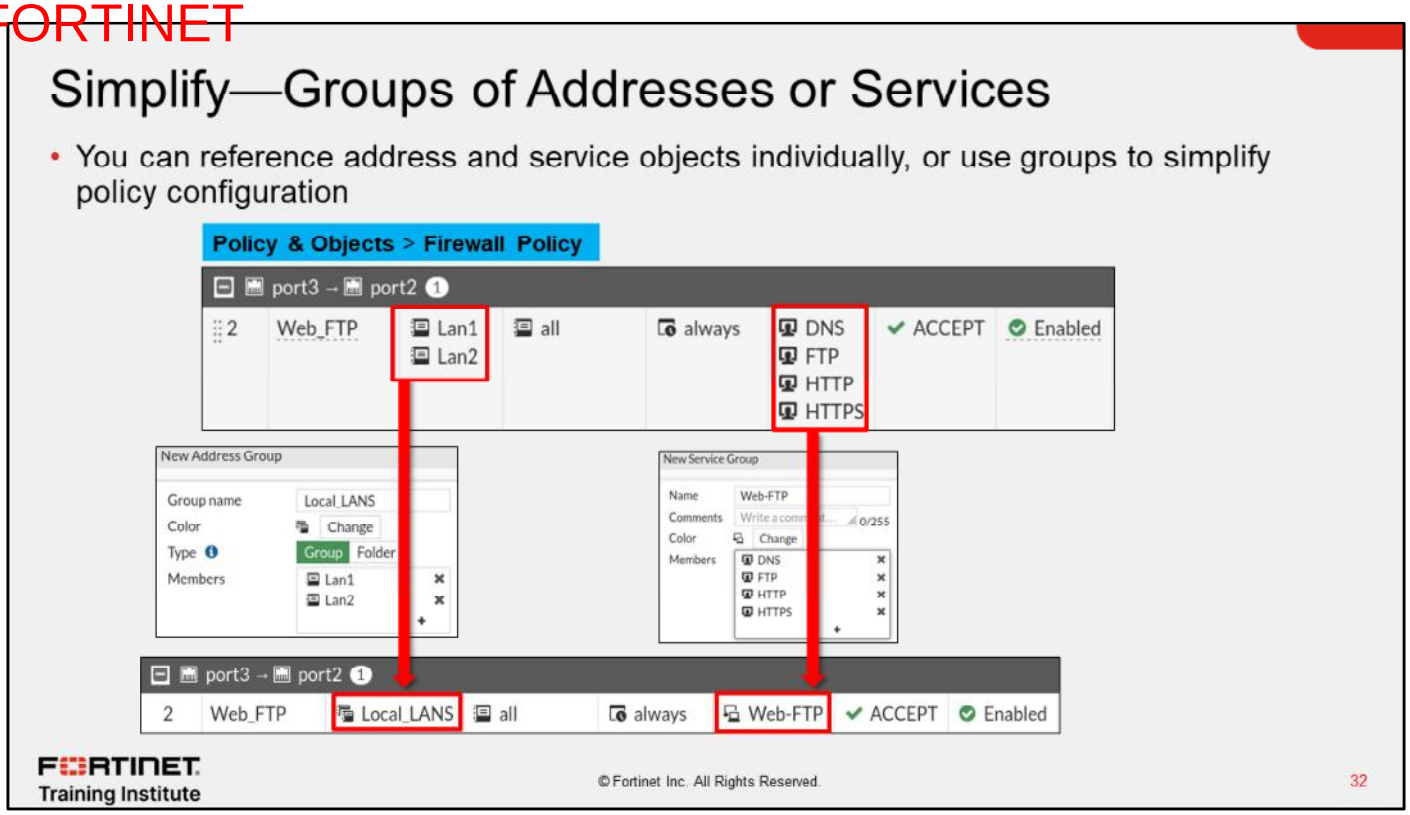

To simplify administration, you can group service and address objects. Then, you can reference that group in the firewall policy, instead of selecting multiple objects each time, or making multiple policies.

This slide shows that four services are used to configure the policy: HTTP, HTTPS, FTP, and DNS. DNS is used by browsers to resolve URLs to IP addresses because people remember domain names for websites instead of IP addresses. If you need to make many policies for web and FTP traffic, then it makes sense to create a service object named **Web-FTP**. That way, you don't have to manually select all four services each time you make a policy. Policies can reference the **Web-FTP** service group instead.

Also, you can consolidate source addresses in source groups.

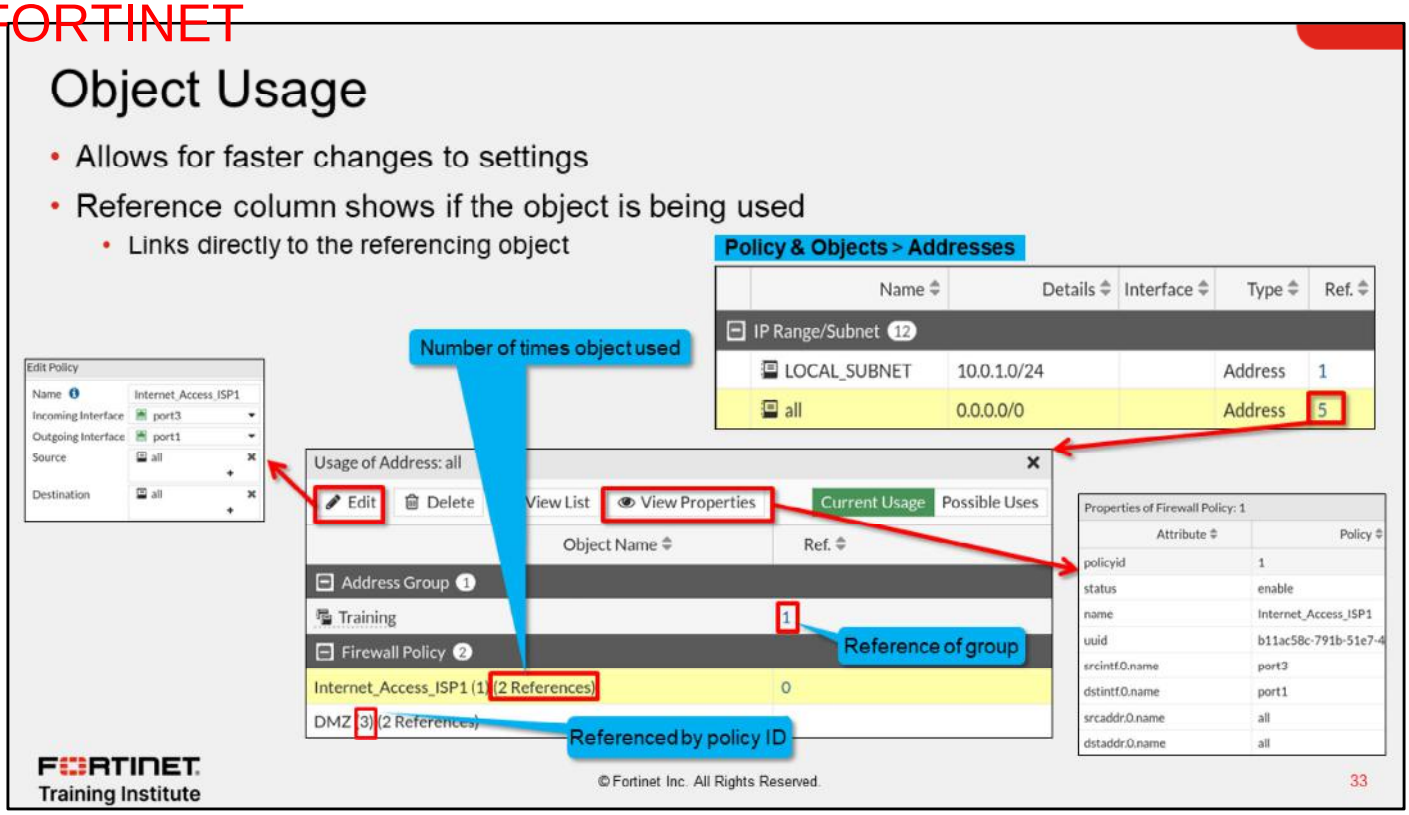

You've just seen several component objects that can be reused as you make policies. What if you want to delete an object?

If an object is being used, you can't delete it. First, you *must* reconfigure the objects that are currently using it. The GUI provides a simple way to find out where in the FortiGate's configuration an object is being referenced. Take a look at the numbers in the **Ref** column. They are the number of places where that object is being used. The number is actually a link, so if you click it, you can see which objects are using it.

In the example shown on this slide, the **all** address object is being used by the **Training** address group and three firewall policies. If you select a firewall policy, you can use the **Edit**, **View List**, and **View Properties** tabs.

- **Edit**: allows you to edit the selected object. In this example, it shows the edit page for the firewall policy ID 1.
- **View List**: allows you to view selected objects in its category. In this example, it will show you the list of all the firewall policies.
- **View Properties**: shows where the object is used in that configuration. In this example, address object **all** is being used in the destination address and source address of that firewall policy.

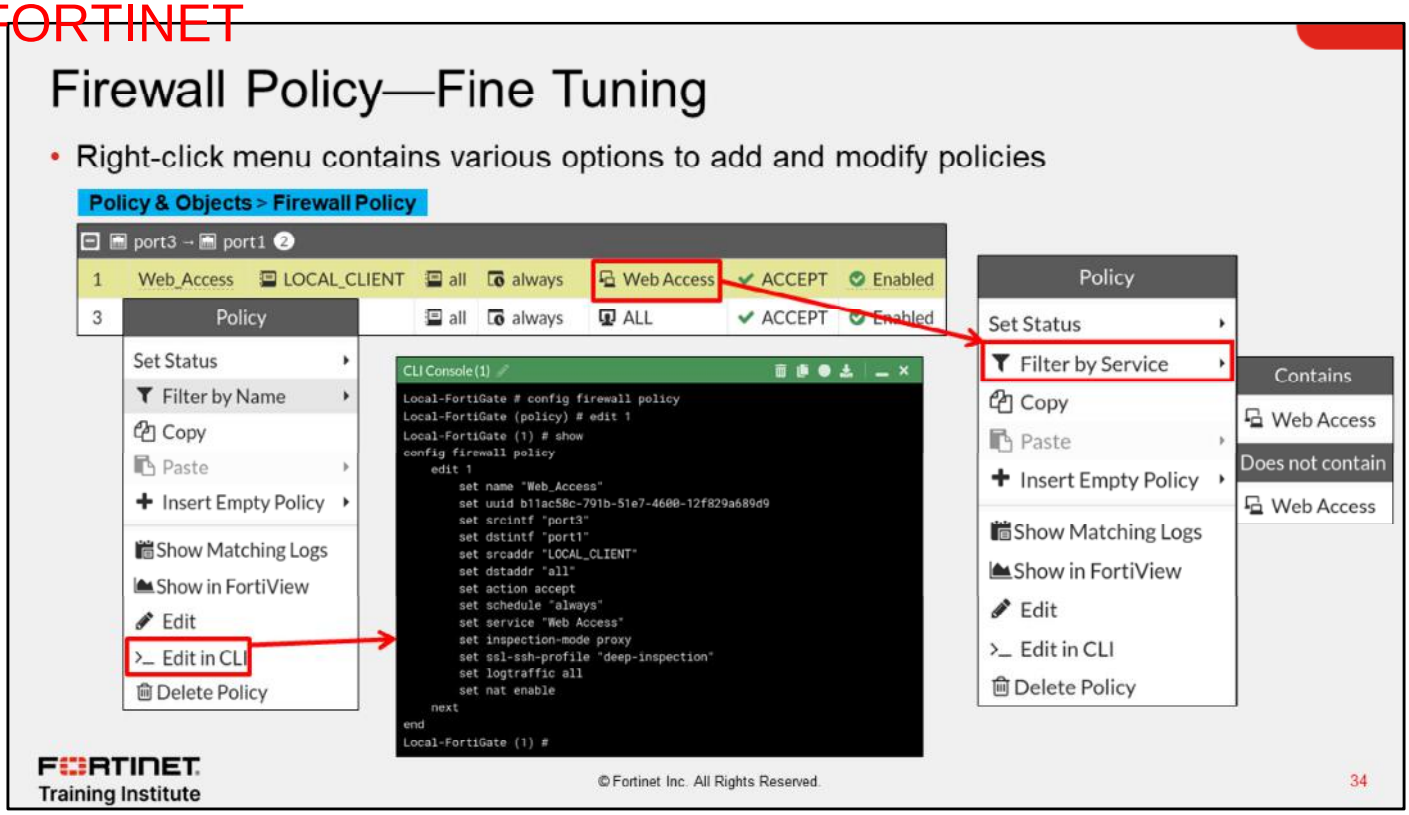

You can right-click any firewall policy to see different menu options to edit or modify the policy. The options include enabling or disabling a firewall policy, inserting firewall policies (above or below), copying and pasting policies, and cloning reverse (only if NAT is disabled on that policy).

Clicking **Edit in CLI** opens the CLI console for the selected firewall policy or object. It shows the configured settings on the CLI and can modify the selected firewall policy or object directly on the **CLI Console**.

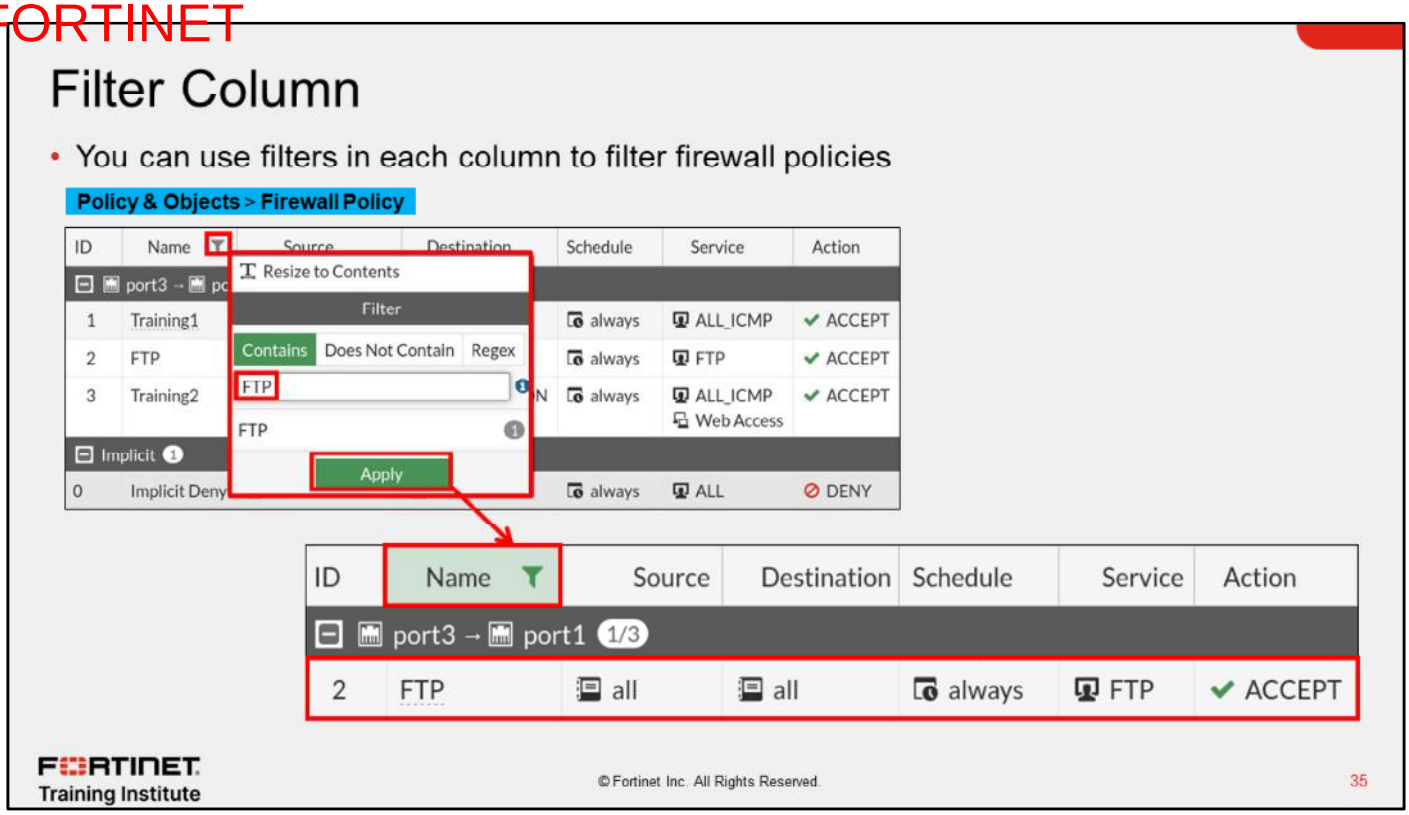

You can filter firewall policies on the GUI using filters in each column. You can add the **ID** column and then click the **ID** column filter icon to filter and search policies based on policy id numbers. You can click the **Name** filter icon to search policies based on policy name, and so on.

Knowledge Check

1. If you configure a firewall policy with the **any** interface, you can view the firewall policy list only in which view?_____ .
   - ✓ A. The **By Sequence View**
   - B. The **Interface Pair View**

FORTINET Training Institute

36

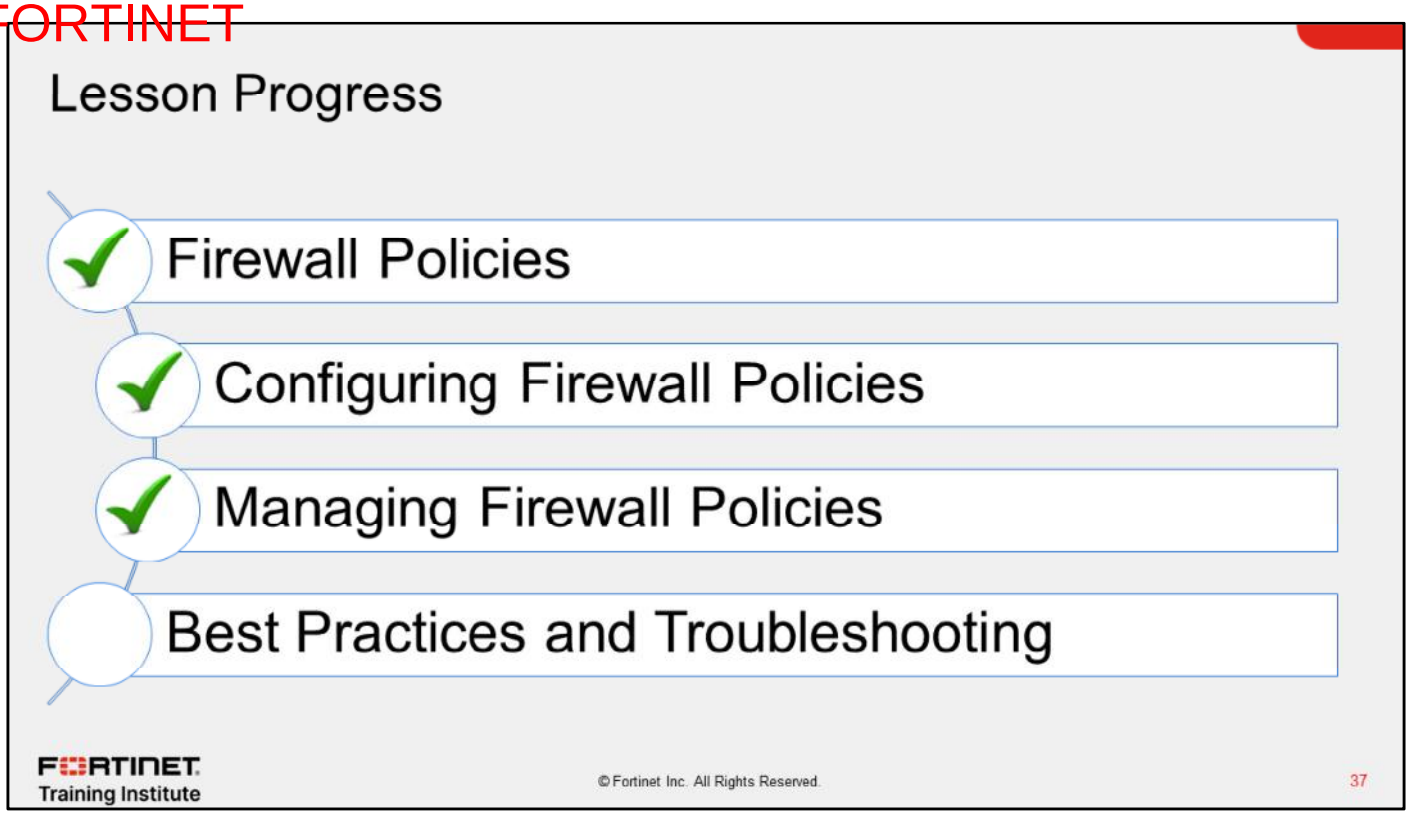

Good job! You now understand how to manage firewall policies on FortiGate.

Now, you'll learn about best practices and troubleshooting related to firewall policies.

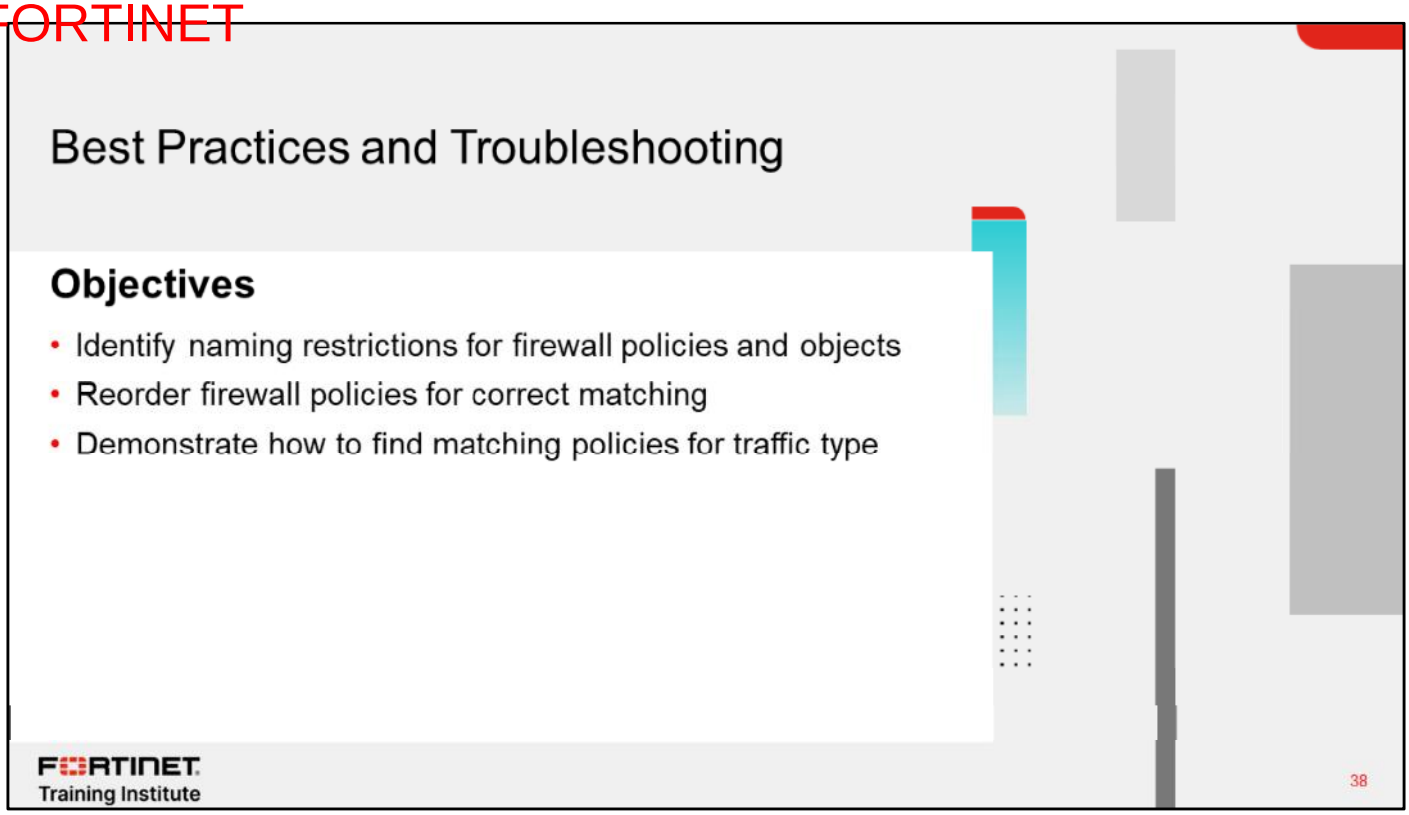

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in knowing firewall policy restrictions and using policy matching techniques, you will be able to apply best practices and basic troubleshooting techniques when working with firewall policies.

Naming Rules and Restrictions

- Most firewall object name fields accept up to 35 characters
- Supported characters in a firewall object name:
  - Numbers: 0 to 9
  - Letters: A to Z (uppercase and lower case)
  - Special characters: hyphen - and underscore _
  - Spaces
    - Avoid using spaces in general
- Some special characters are supported in passwords, comments, replacement messages, and so on
  - < > ( ) # " " ' '

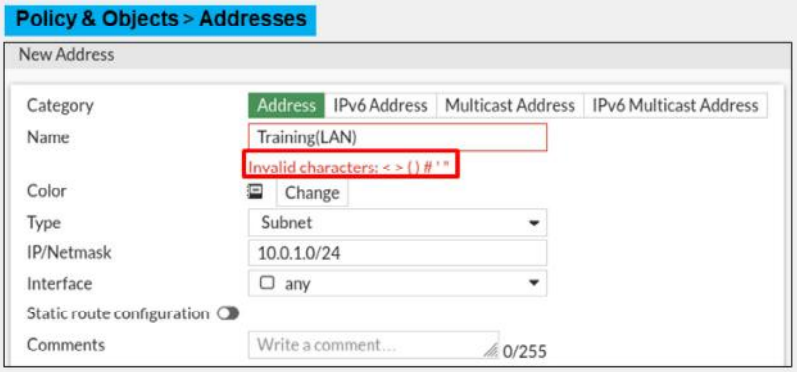

FORTINET Training Institute

39

When configuring names for firewall objects, only specific characters are supported. For example, `Training(LAN)` is not a valid name for an address object because it includes special characters that are not supported. Although spaces are supported in the names, as a best practice, avoid using spaces in names. Instead, use a hyphen or underscore. Using spaces can cause issues when trying to modify on the CLI, or troubleshooting.

However, many special characters are supported in passwords, comments, replacement messages, and so on.

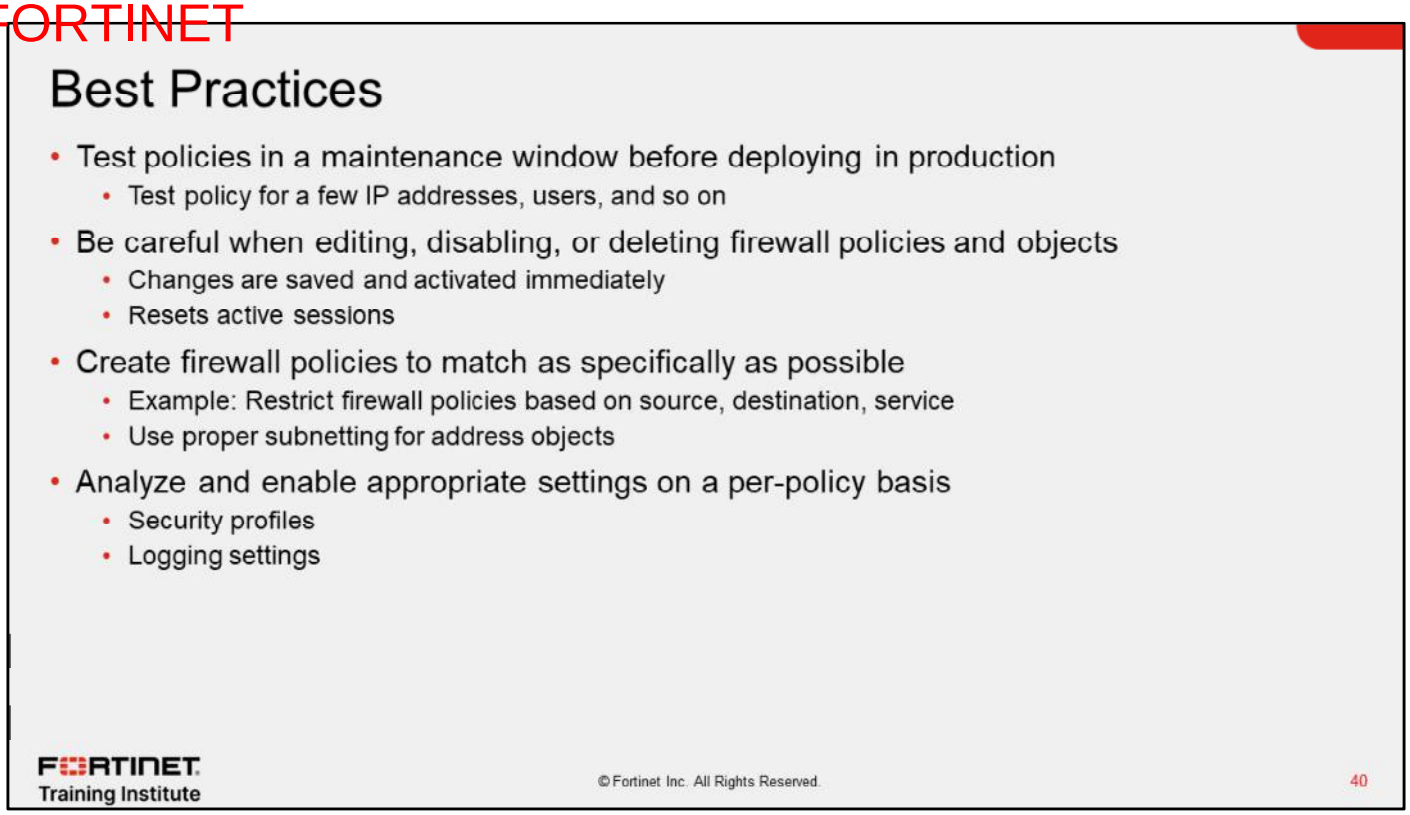

Always plan a maintenance window and create a test case for a few IP addresses and users, before implementing configuration changes in the production network. Any configuration changes made using the GUI or CLI take effect immediately, and can interrupt service.

As a best practice, try to configure firewall policies as specifically as possible. This helps to restrict access to only those resources. For example, use correct subnets when configuring address objects.

Another setting worth mentioning is security profiles. Security profiles help to provide appropriate security for your network. Proper logging configuration can also help you to analyze, diagnose, and resolve common network issues.

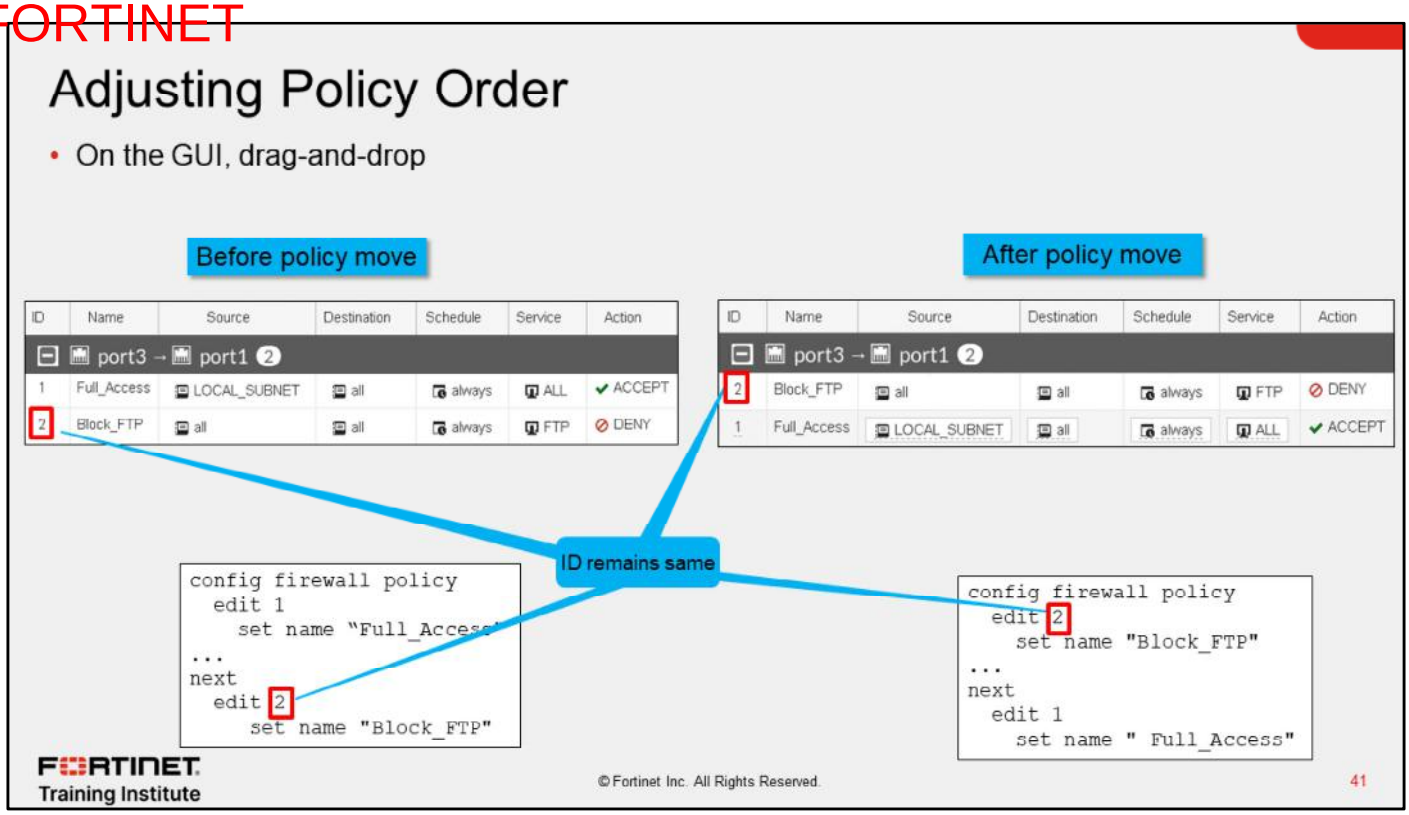

Remember you learned that only the first matching policy applies? Arranging your policies in the *correct position* is important. It affects which traffic is blocked or allowed. In the section of the applicable interface pair, FortiGate looks for a matching policy, beginning at the top. So, you should put more specific policies at the top; otherwise, more general policies will match the traffic first, and more granular policies will never be applied.

In the example shown on this slide, you're moving the **Block_FTP** policy (ID 2) that matches only FTP traffic, to a position above a more general **Full_Access** (accept everything from everywhere*)* policy. Otherwise, FortiGate would always apply the first matching policy in the applicable interface pairs—**Full_Access**–and never reach the **Block_FTP** policy.

When moving the policies across the policy list, policy IDs remain unchanged.

Note that FortiGate assigns the next highest available ID number as policies are created.

Note that policy IDs are identifiers and are not displayed by default on the policy list GUI. You can add a policy **ID** column using the **Configure Table** settings icon.

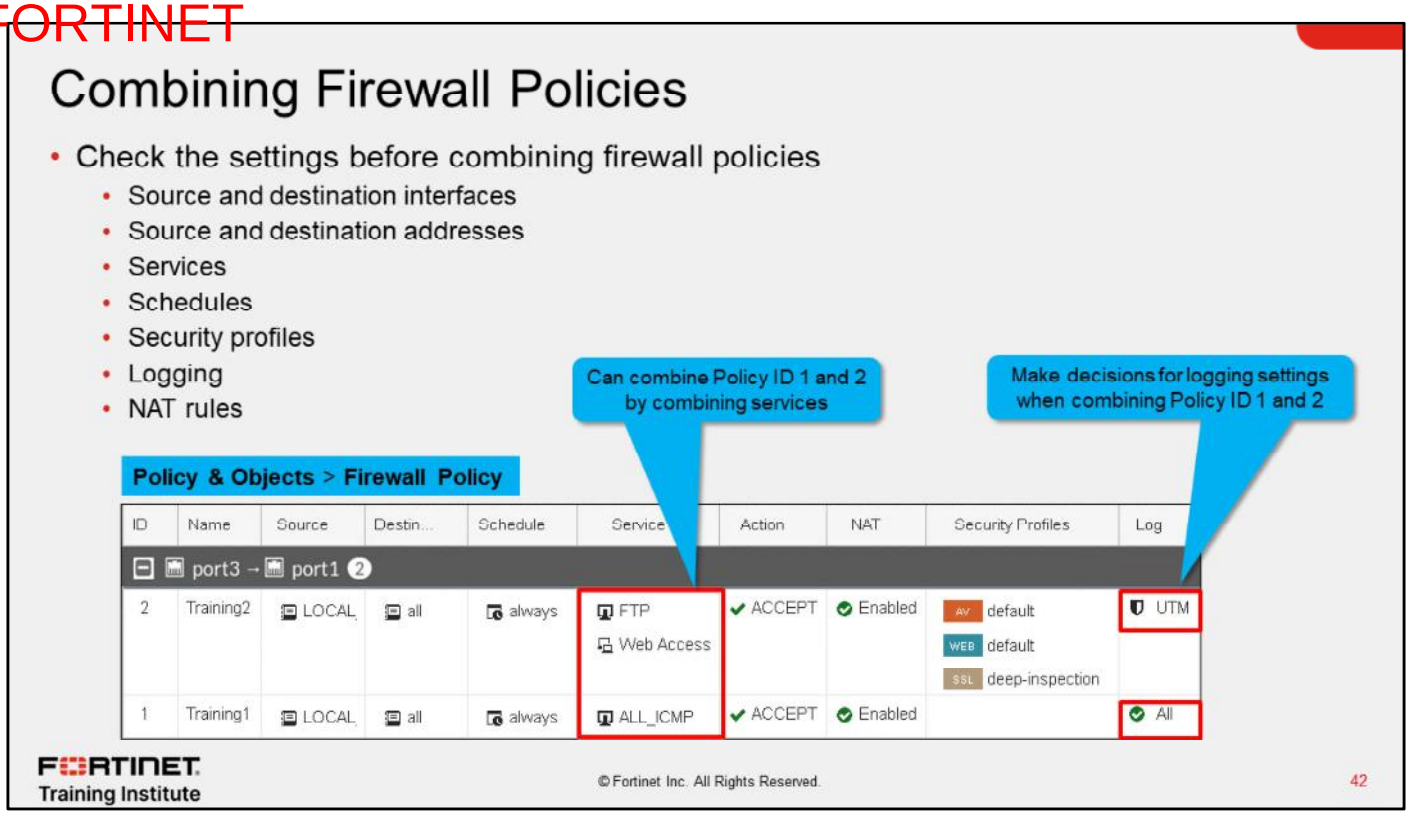

In order to optimize and consolidate firewall policies, always check all configured settings. In the example shown on this slide, the two firewall policies have differences in terms of services, security profiles, and logging settings. You can consolidate these two firewall policies by combining services and choosing appropriate logging settings.

If you select **Security Events** (UTM) for the logging settings, traffic logs will not be generated for **ALL_ICMP** traffic.

Note that the **ALL_ICMP** service is not subject to web filter and antivirus scans, which means that applying these security profiles to the ICMP traffic will result in the traffic passing through without being inspected.

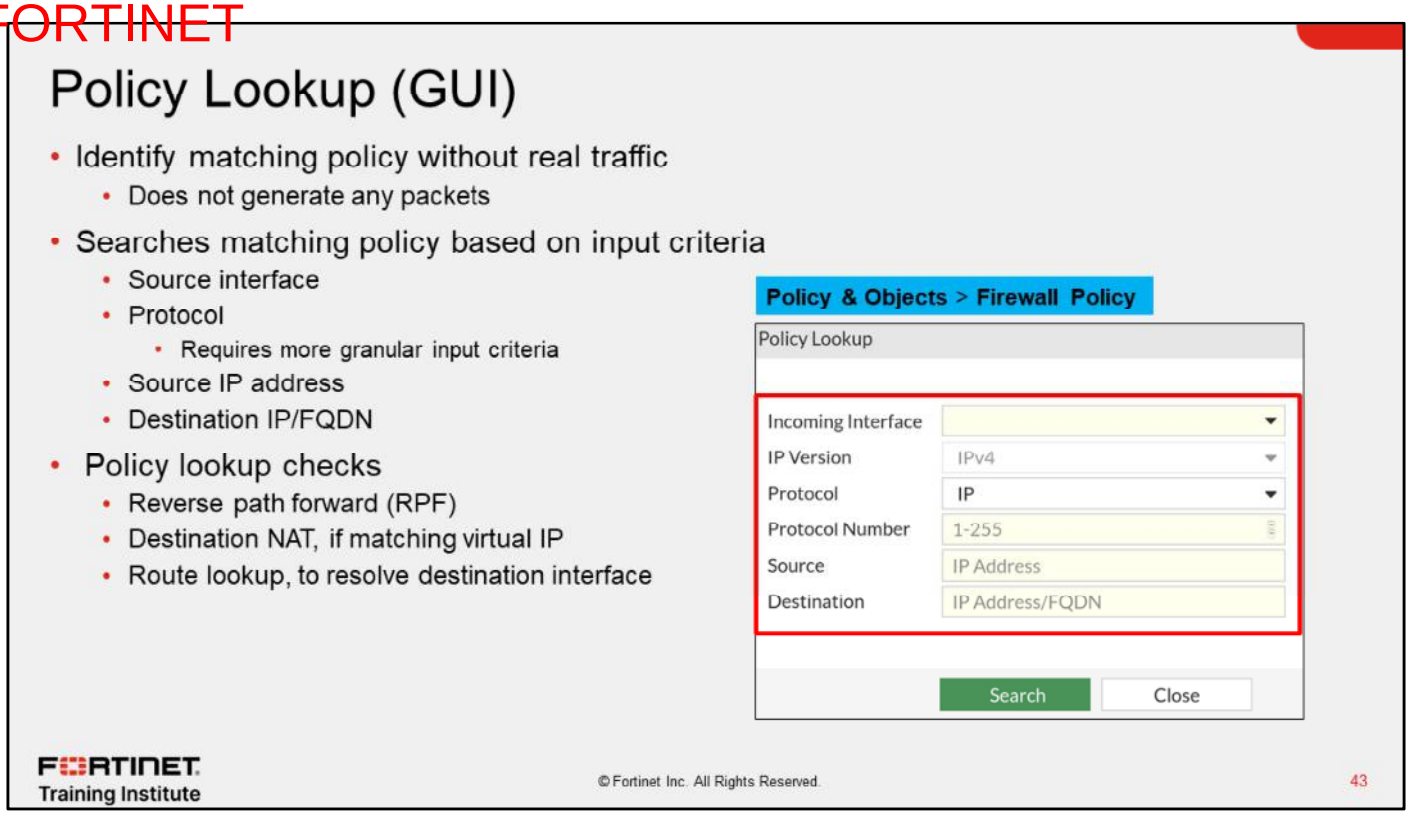

You can find a matching firewall policy based on the policy lookup input criteria. Policy lookup creates a packet flow over FortiGate without real traffic. From this, policy lookup can extract a policy ID from the flow trace and highlight it on the GUI policy configuration page.

Depending on the protocol you select (for example, TCP, UDP, IP, ICMP, and so on), you need to define other input criteria. For example, when you select TCP as the protocol, you need to define the source address, source port (optional), destination port, and destination address. When you select ICMP as the protocol, you need to define the ICMP type/code, source address, and destination address.

When FortiGate is performing policy lookup, it performs a series of checks on ingress, stateful inspection, and egress, for the matching firewall policy, from top to bottom, before providing results for the matching policy.

Note that if the firewall policy status is set to **disable**, the policy lookup skips the disabled policy and checks for the next matching policy in the list.

When FortiGate is in Transparent mode, it does not support the policy lookup function.

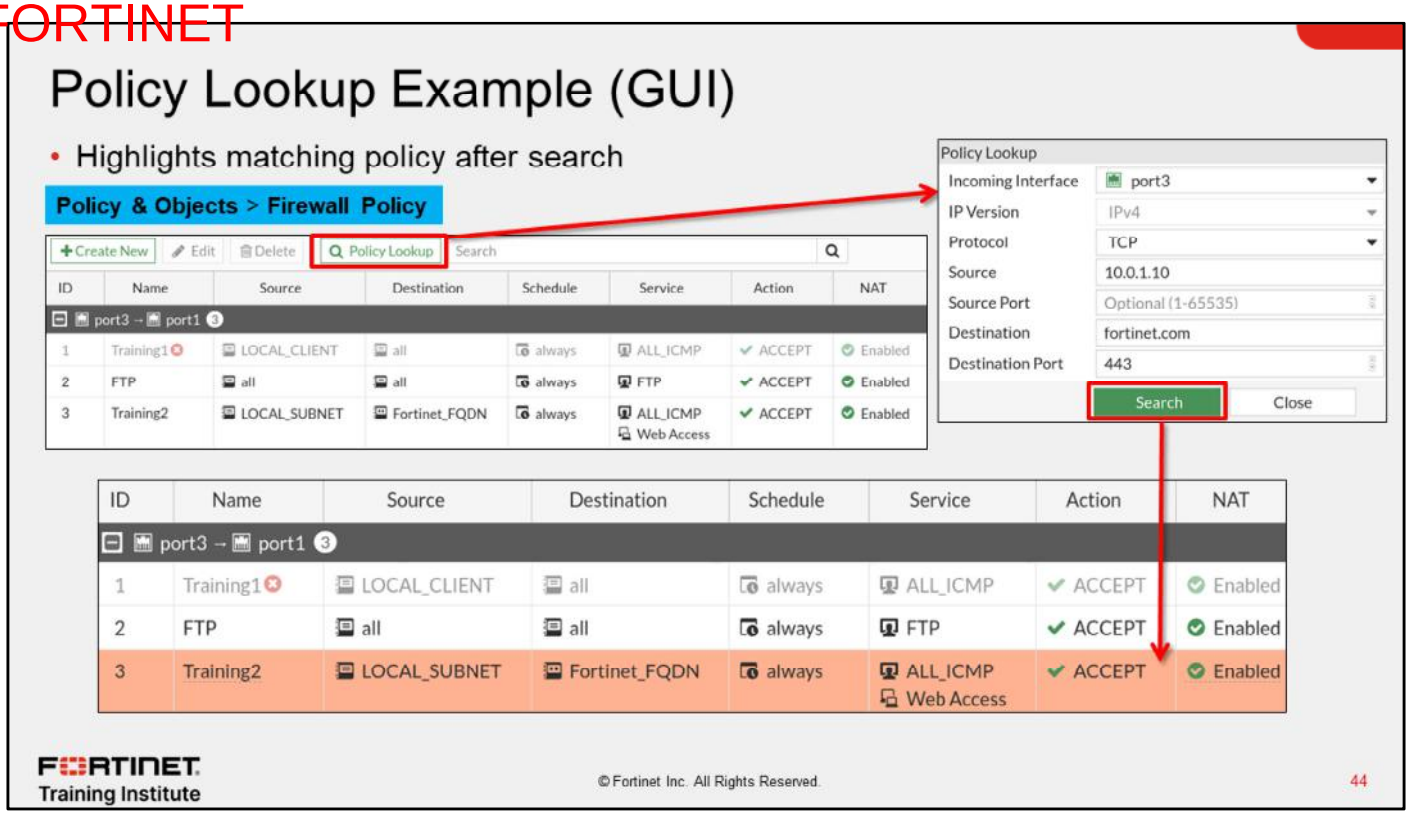

Based on the input criteria, after clicking **Search**, the trace result is selected and highlighted on the **Firewall Policy** page.

Why didn't policy **ID #1** or **ID #2** match the input criteria?

Because policy **ID #1** status is set to **disable**, policy lookup skips the disabled policy. For firewall policy **ID #2** , it doesn't match the destination port specified in the policy lookup matching criteria.

Knowledge Check

1. Which of the following naming formats is correct when configuring a name for a firewall address object?
   - ✓ A. Good_Training
   - B. Good(Training)

2. What is the purpose of the policy lookup feature on FortiGate?
   - ✓ A. To find a matching policy based on input criteria
   - B. To block traffic based on input criteria

FORTINET Training Institute

45

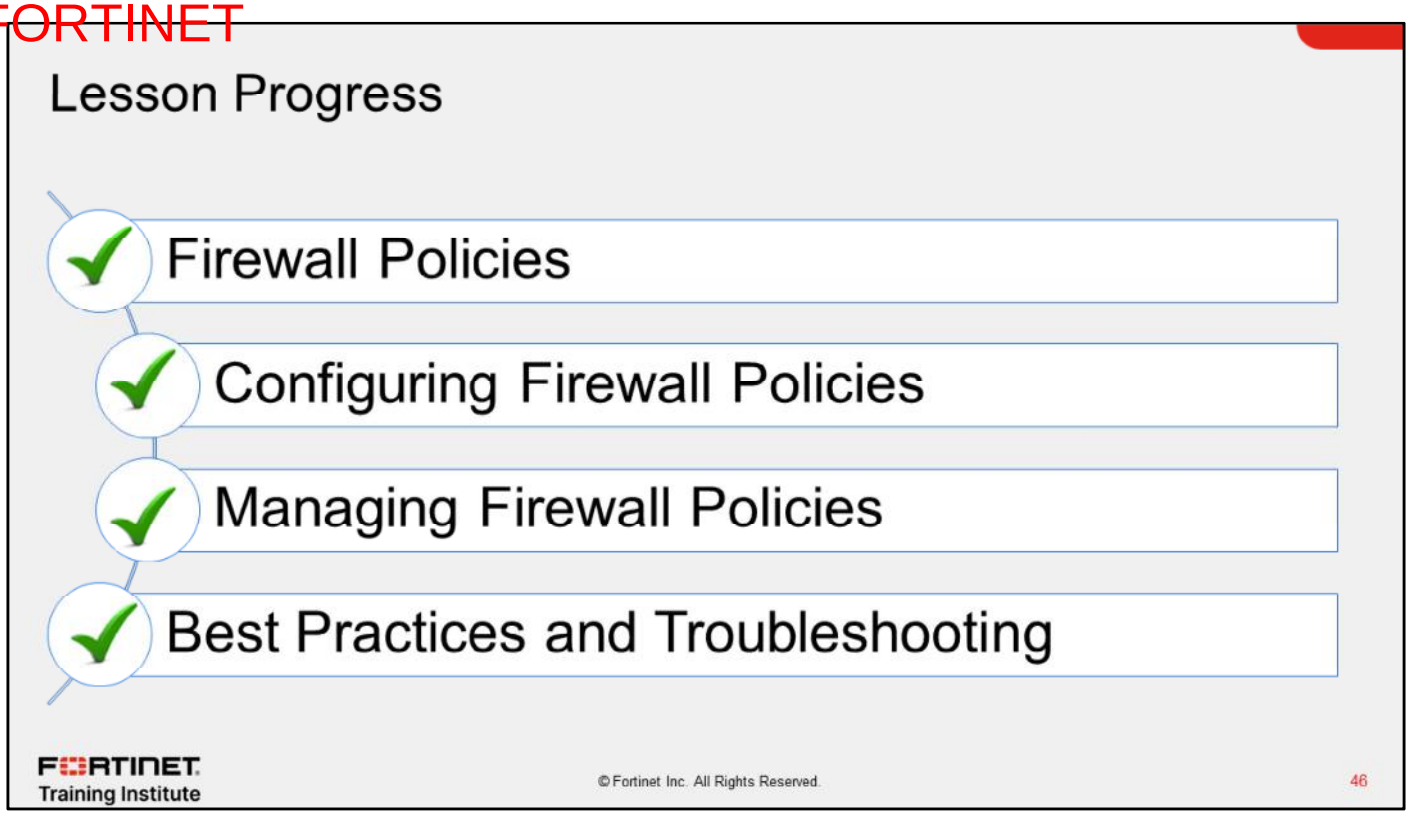

Congratulations! You have completed this lesson.

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

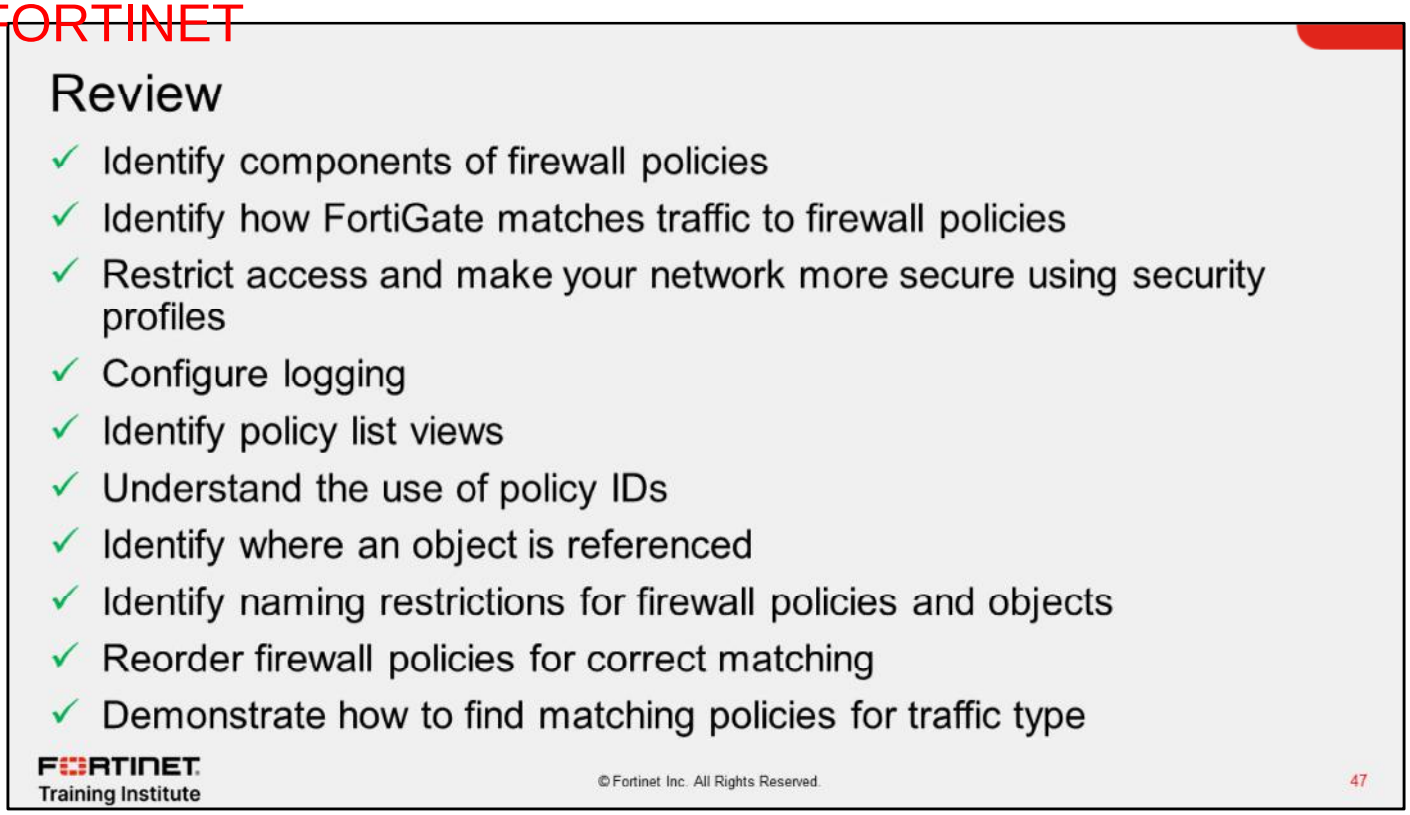

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure, use, and manage firewall policies.

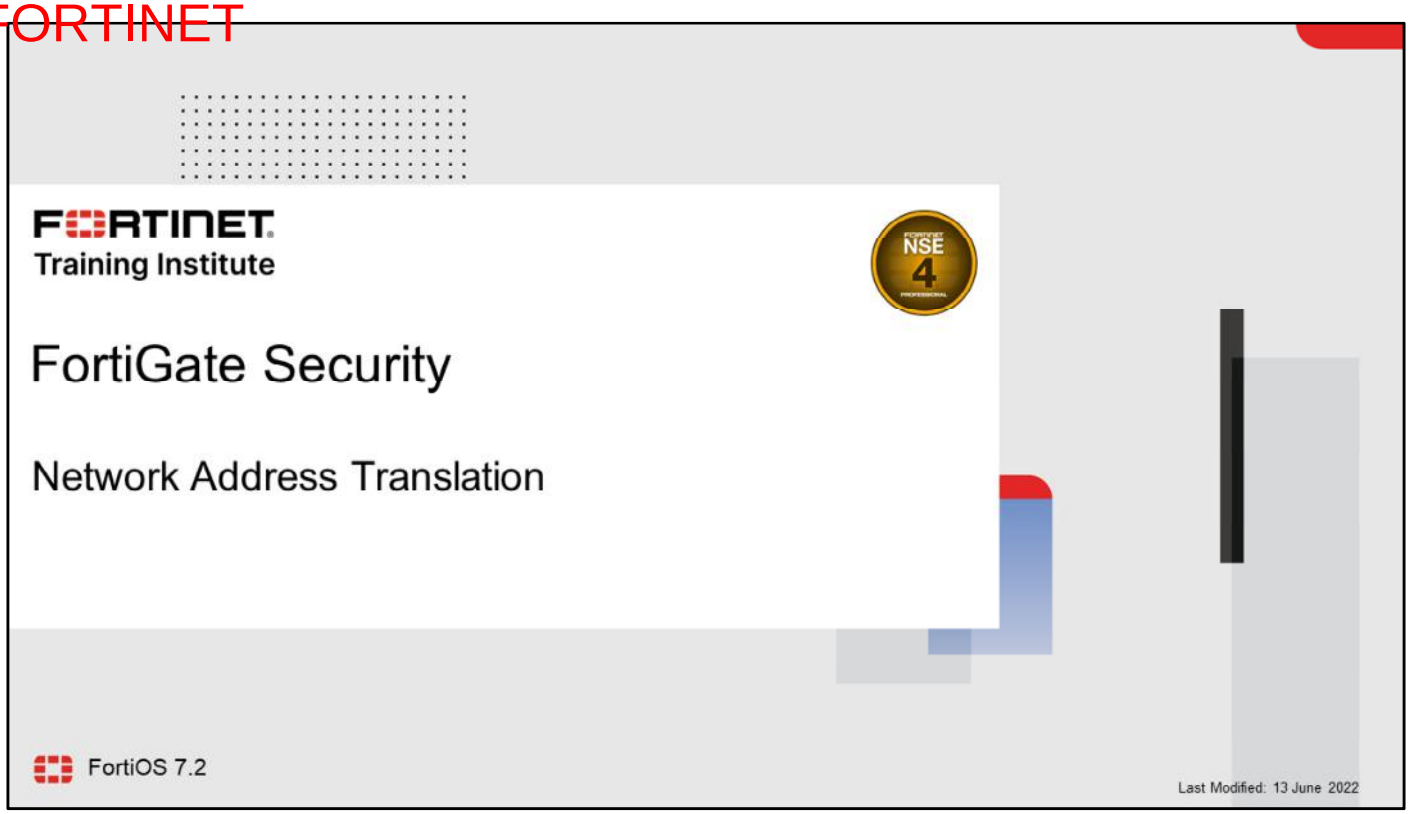

In this lesson, you will learn how to configure network address translation (NAT) and use it to implement source NAT (SNAT) and destination NAT (DNAT) for the traffic passing through FortiGate.

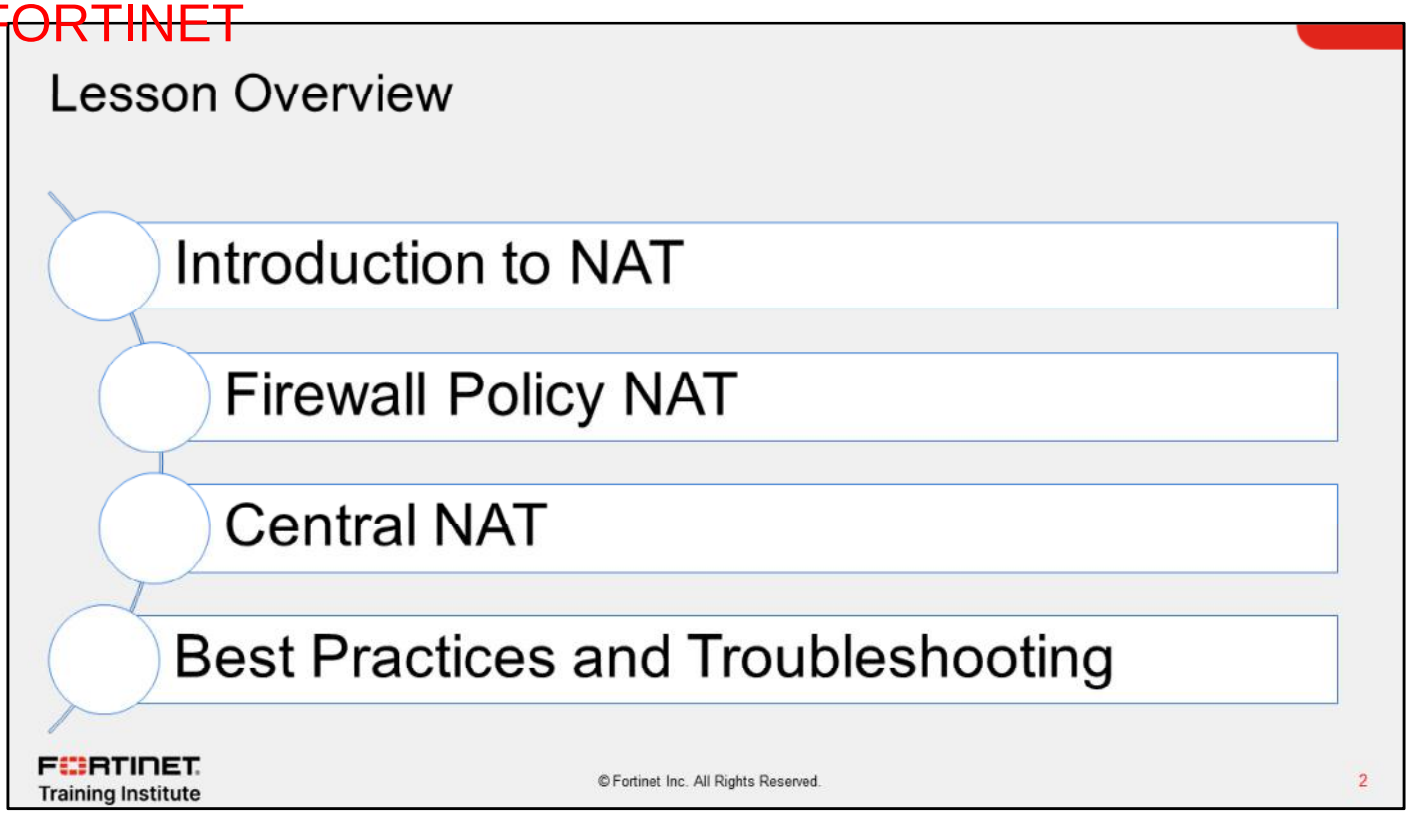

In this lesson, you will learn about the topics shown on this slide.

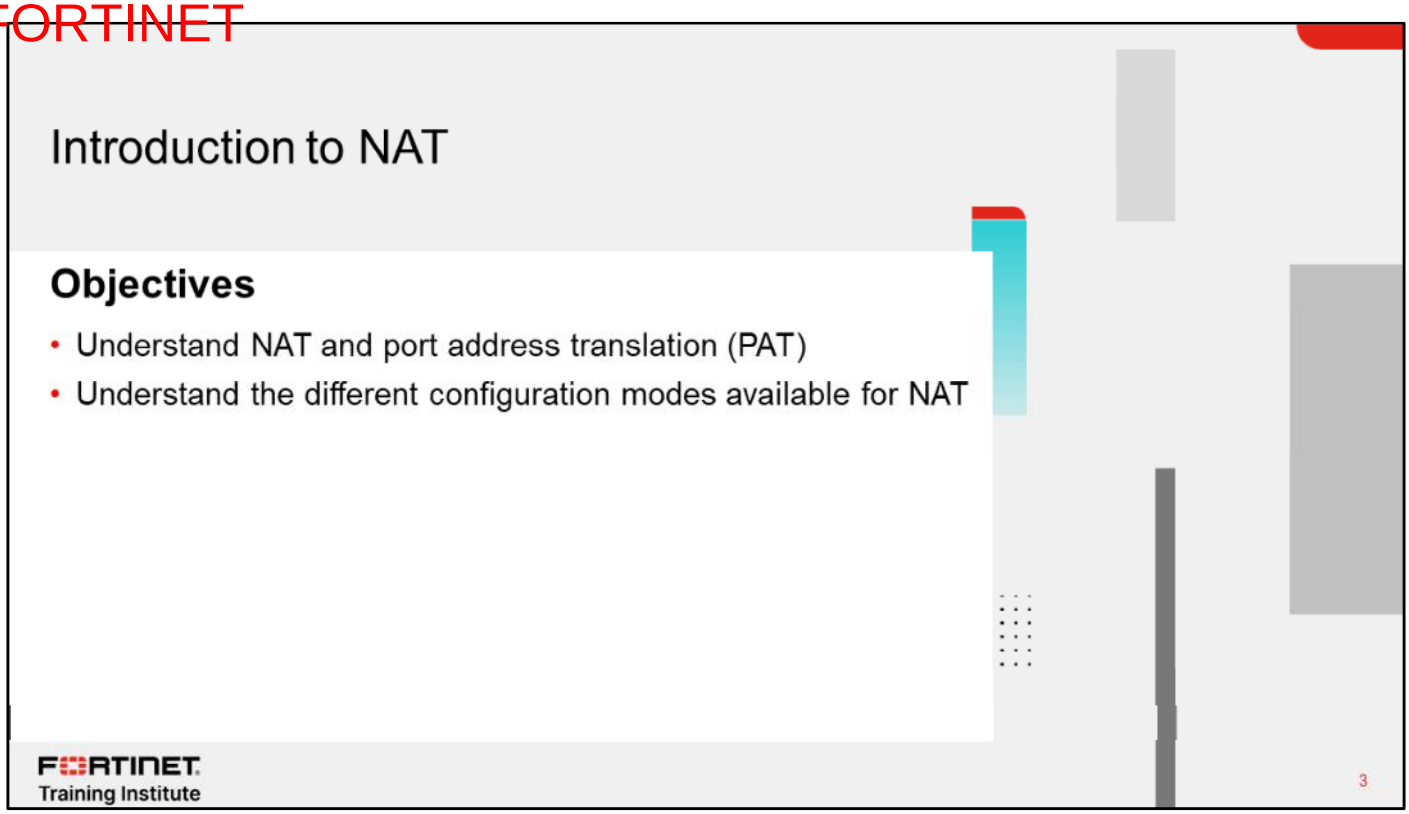

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding how NAT and PAT work, and the available NAT configuration modes, you will be well-positioned to plan the implementation of NAT in your network.

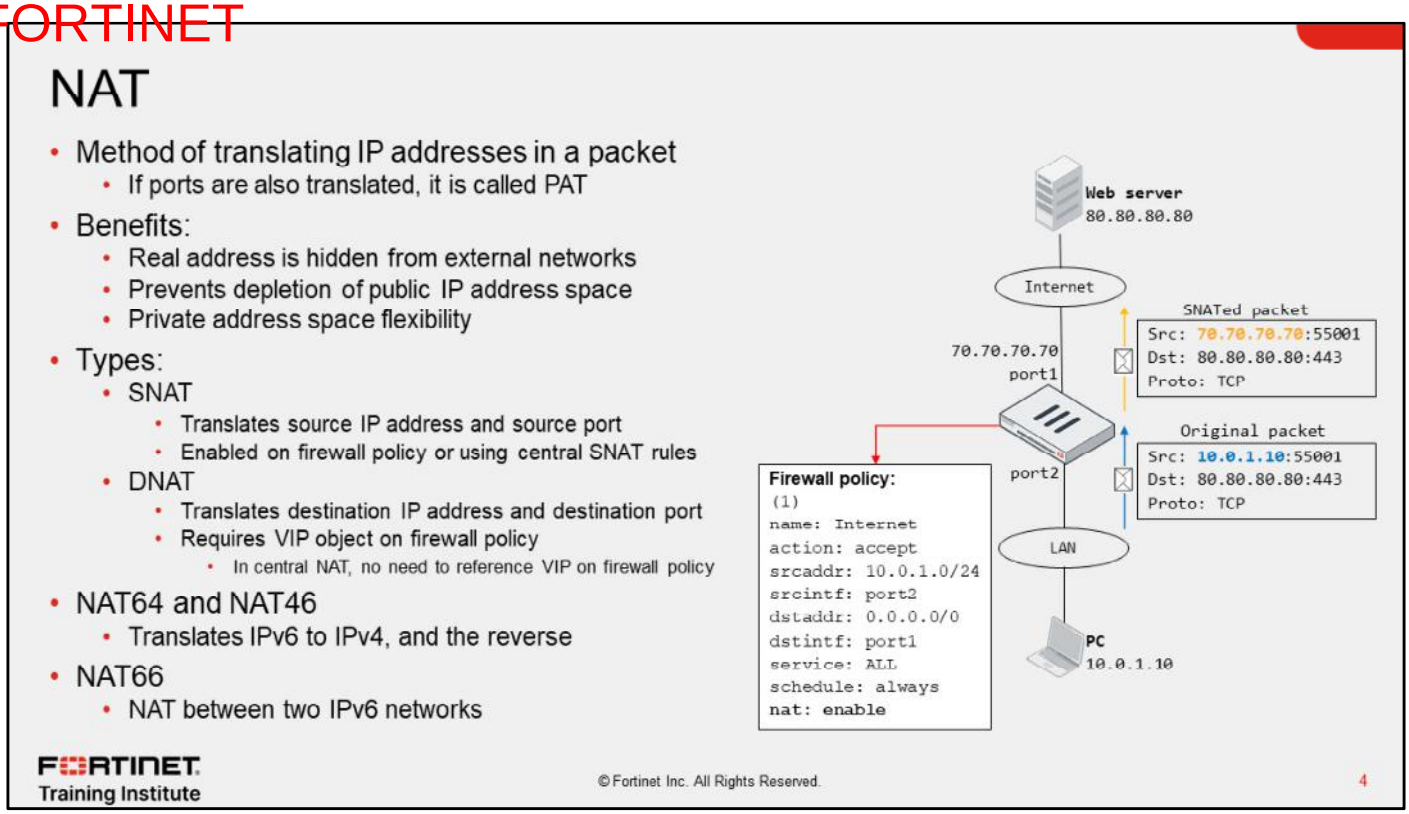

NAT is a method that enables a NAT device such as a firewall or router, to translate (or map) the IP address in a packet to another IP address, usually for connectivity purposes. If the port information in the packet is also translated, then the translation method is called PAT. NAT provides the following benefits:

- Security: The real address of a device is hidden from external networks.
- Public address depletion prevention: Hundreds of computers can share the same public IPv4 address.
- Private address flexibility: The addresses can stay the same, even if ISPs change. You can reuse private addresses in multiple networks.

There are two types of NAT: SNAT and DNAT. In SNAT, a NAT device translates the source IP address and source port in a packet. In DNAT, a NAT device translates the destination IP address and destination port. You can configure FortiGate to perform SNAT and DNAT as follows:

- For SNAT, you enable NAT on the matching firewall policy. Alternatively, you can enable central NAT to configure central SNAT rules for the VDOM.
- For DNAT, you configure virtual IPs (VIPs) and then reference them on the matching firewall policy. If you enable central NAT, you configure central DNAT rules and VIP objects for DNAT.

The example on this slide shows the most common use case for NAT: SNAT. FortiGate, acting as a NAT device, translates the private IP address assigned to the PC to the public address assigned by your ISP. The private-to-public source address translation is needed for the PC to access the internet web server.

NAT64 and NAT46 refer to the methods that translate an IPv6 address to an IPv4 address and the reverse, respectively. They enable you to communicate IPv6 networks with IPv4 networks, and the reverse. NAT66 consists of translating addresses between two IPv6 networks.

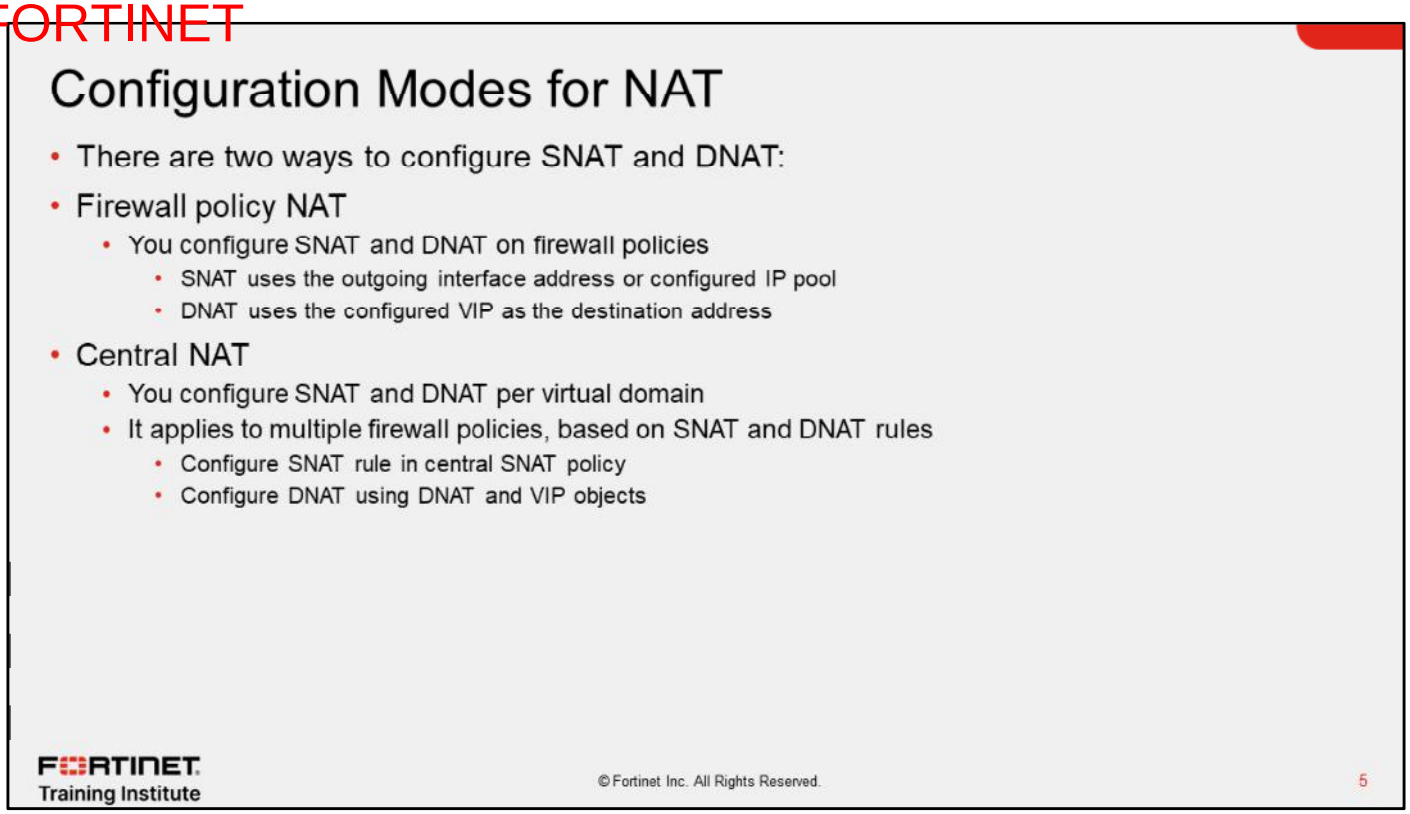

You configure NAT using firewall policy NAT mode or central NAT.

When you use firewall policy NAT mode, you must configure SNAT and DNAT for each firewall policy.

When you use central NAT, you configure NAT per virtual domain by configuring SNAT and DNAT rules. The result is that SNAT and DNAT settings automatically apply to multiple firewall policies, as opposed to each firewall policy in firewall policy NAT.

As a best practice, when you use central NAT, you should configure specific SNAT and DNAT rules so that they match only the desired firewall policies in your configuration.

Both firewall policy NAT and central NAT produce the same results; however, some deployment scenarios are best suited to firewall policy NAT and some are best suited to central NAT.

Firewall policy NAT is suggested for deployments that include relatively few NAT IP addresses and where each NAT IP address would have separate policies and security profiles. Central NAT is suggested for more complex scenarios where multiple NAT IP addresses have identical policies and security profiles, or in next generation firewall (NGFW) policy mode, where the appropriate policy may not be determined at the first packet.

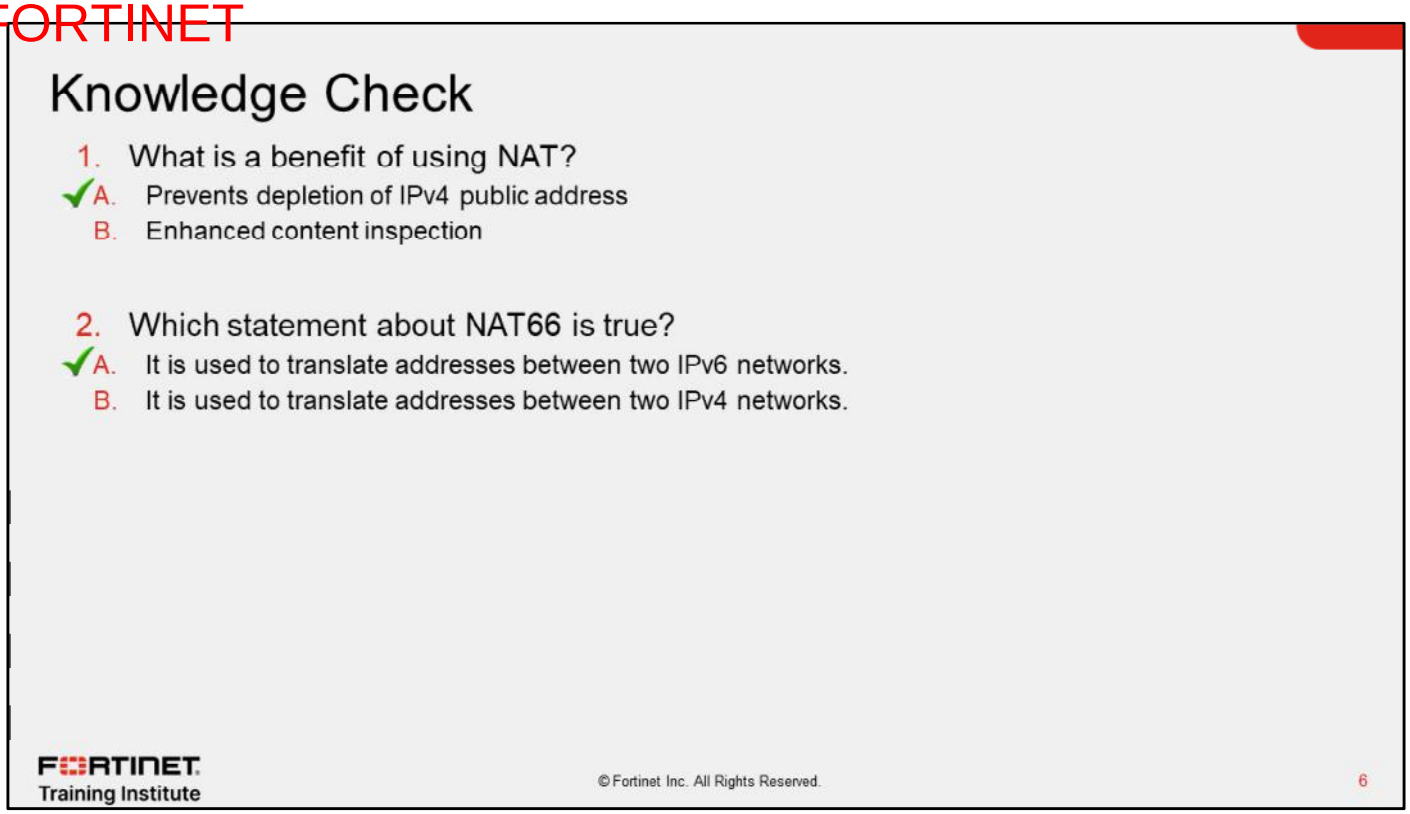

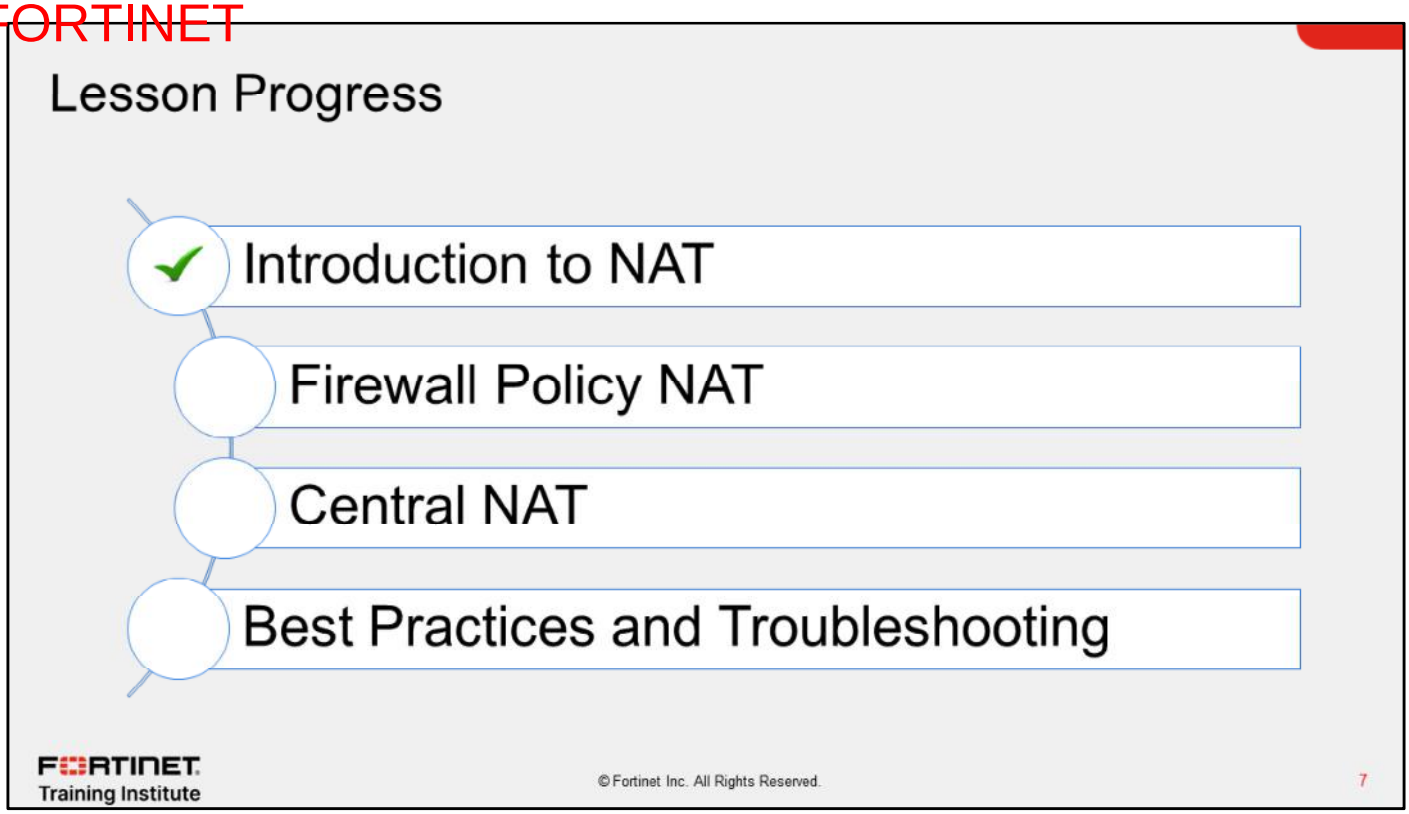

Good job! You now know about NAT.

Now, you'll learn about firewall policy NAT.

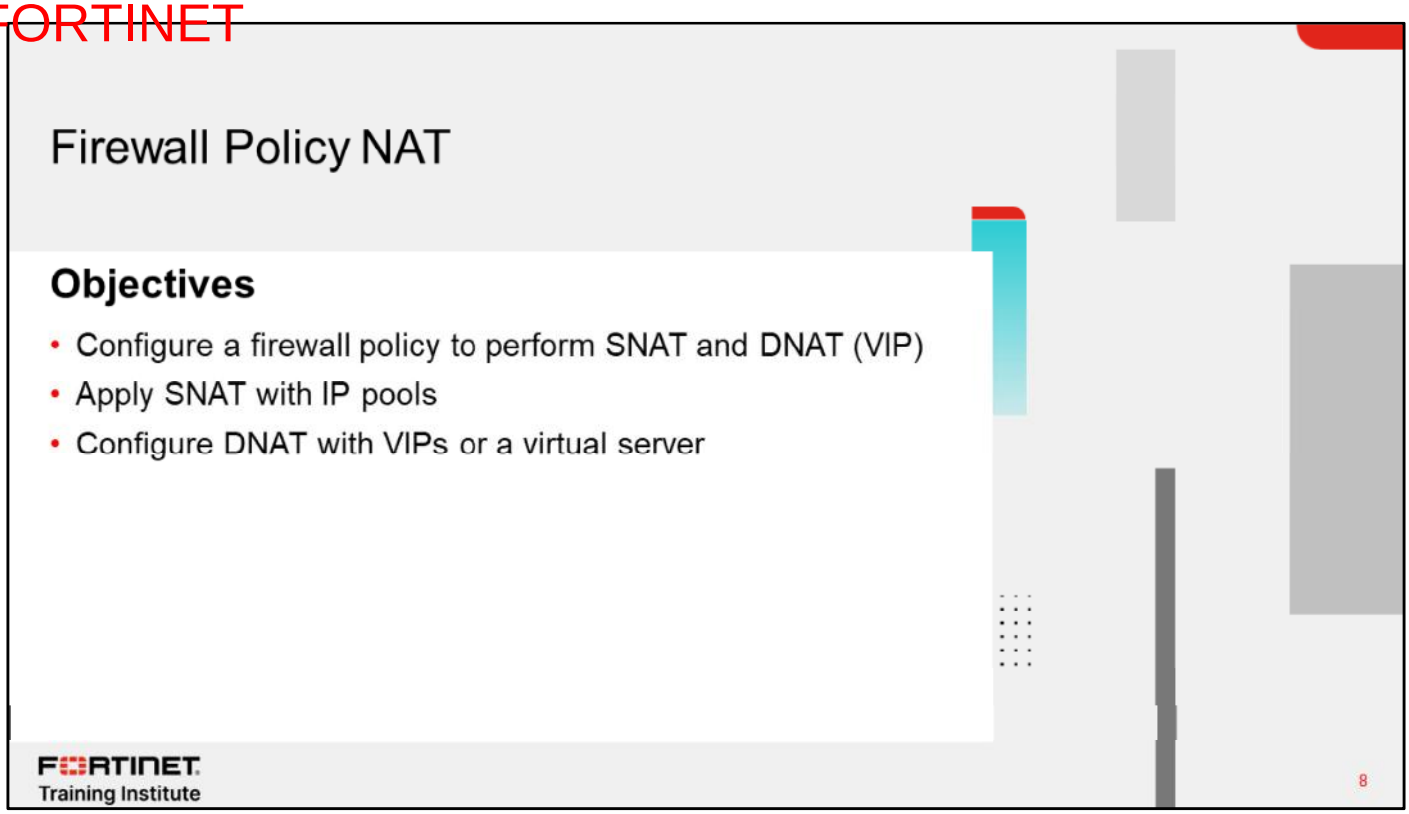

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in these areas, you will be able to configure firewall policies and apply appropriate SNAT and DNAT, and understand how it is applied to the traffic traversing through FortiGate.

There are two ways to configure firewall policy SNAT:

- Use the outgoing interface address.
- Use the dynamic IP pool.

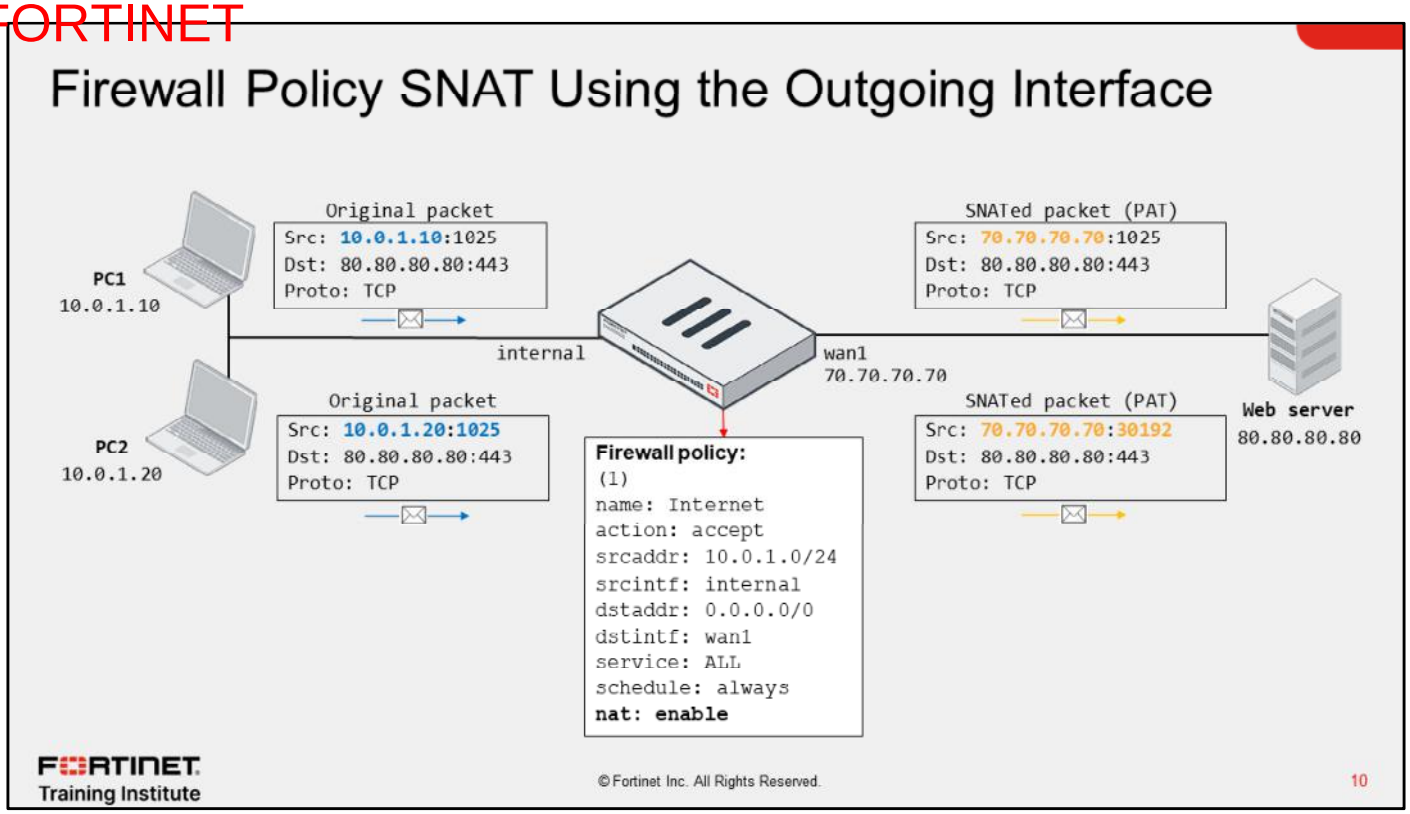

When you select **Use Outgoing Interface Address** on the matching firewall policy, FortiGate uses the egress interface address as the NAT IP for performing SNAT.

If there are multiple devices behind FortiGate, FortiGate performs many-to-one NAT. This is also known as PAT. FortiGate assigns to each connection sharing the egress interface address a port number from a pool of available ports. The assignment of a port enables FortiGate to identify packets associated with the connection and then perform the corresponding translation. This is the same behavior as the overload IP pool type, which you will also learn about.

Optionally, you may select a fixed port, in which case the source port translation is disabled. With a fixed port, if two or more connections require the same source port for a single IP address, only one connection is established.

The example on this slide shows two PCs behind FortiGate that share the same public IP address (`70.70.70.70`) to access the internet web server `80.80.80.80`. Because **Use Outgoing Interface Address** is enabled on the firewall policy—`set nat enable` on the CLI—the source IP address of the PCs is translated to the egress interface address. The source port, however, is not always translated. It depends on the available ports and the connection 5-tuple. In the example shown on this slide, FortiGate translates the source port of the connection from PC2 only. Otherwise, the two connections would have the same information on the session table for the reply traffic, which would result in a session clash.

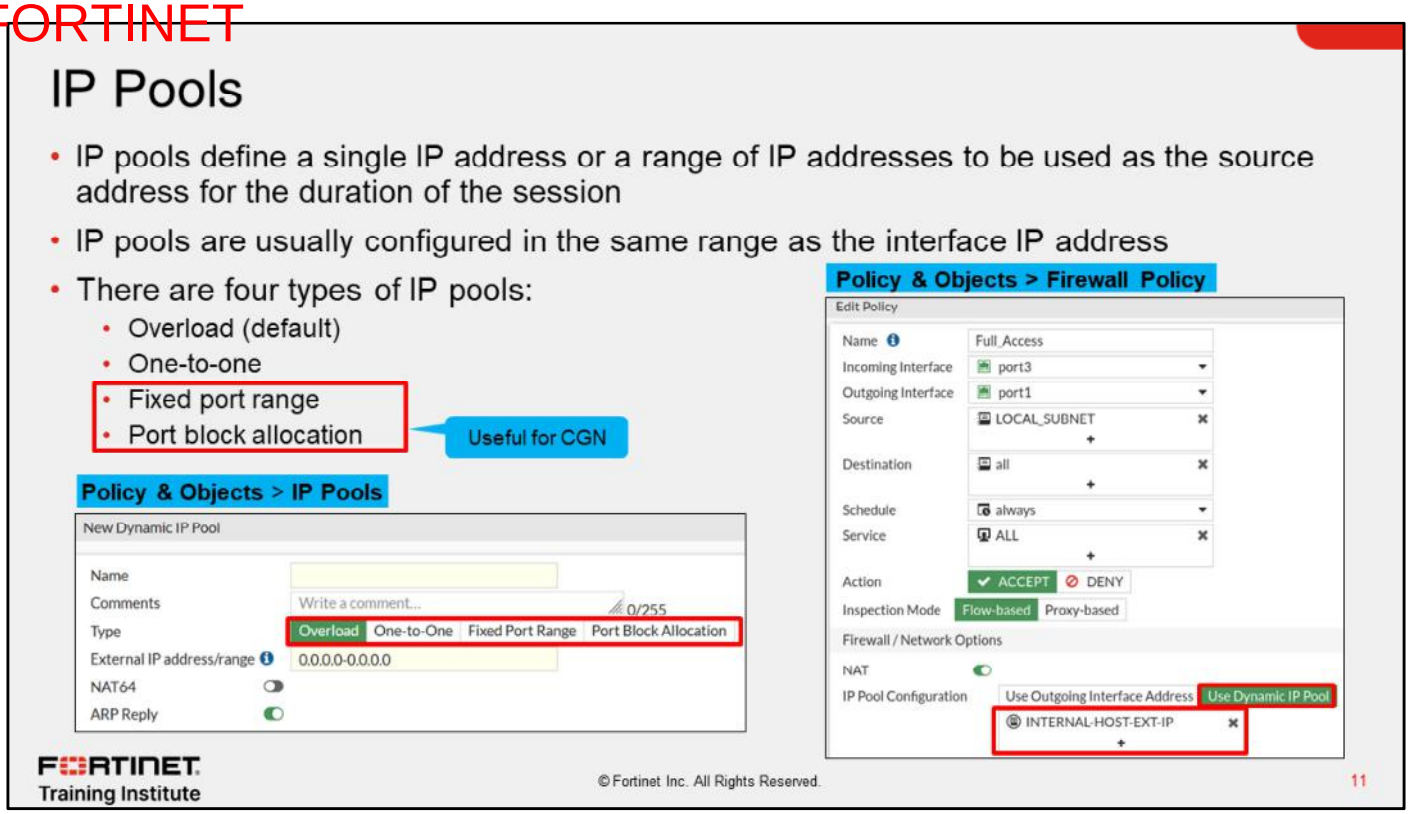

IP pools are a mechanism that allow sessions leaving the FortiGate firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses are used instead of the IP address assigned to that FortiGate interface.

IP pools are usually configured in the same range as the interface IP address.

When you configure the IP pools that will be used for NAT, there is a limitation that you must take into account. If the IP addresses in the IP pool are different from the IP addresses that are assigned to the interface(s), communications based on those IP addresses *may* fail *if the routing is not properly configured.* For example, if the IP address assigned to an interface is `172.16.100.1/24`, you cannot choose `10.10.10.1` to `10.10.10.50` for the IP pool unless appropriate routing is configured.

There are four types of IP pools that you can configure on the FortiGate firewall:
- Overload
- One-to-one
- Fixed port range
- Port block allocation

The fixed port range and port block allocation types are more common carrier-grade NAT (CGN) deployments.

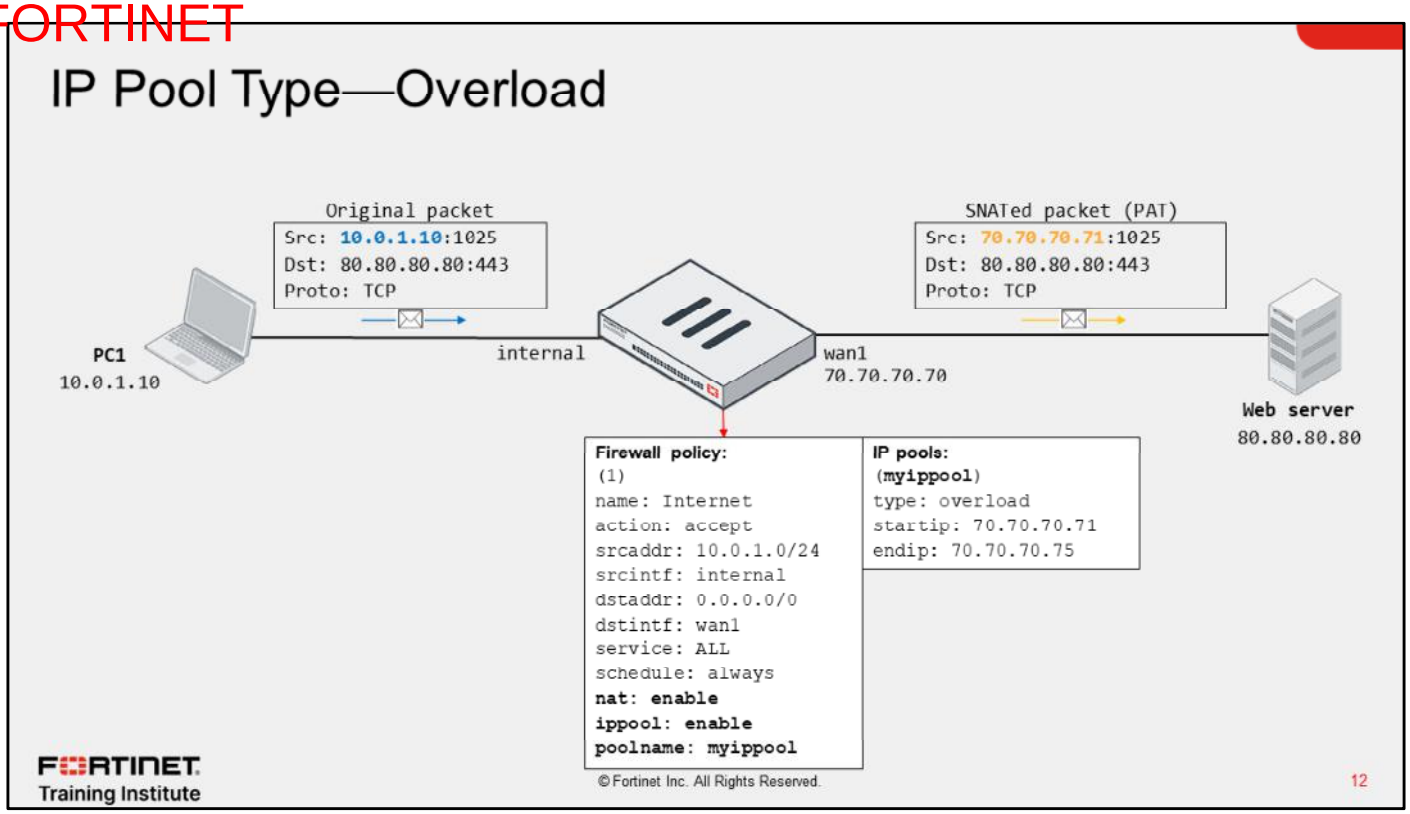

If you use an IP pool, the source address is translated to an address from that pool, rather than the egress interface address. The larger the number of addresses in the pool, the greater the number of connections that the pool can support.

The default IP pool type is overload. In the overload IP pool type, a many-to-one or many-to-few relationship and port translation is used.

In the example shown on this slide, source IP `10.0.1.10` is translated to the address `70.70.70.71`, which is one of the addresses defined in the IP pool (`70.70.70.71 – 70.70.70.75`).

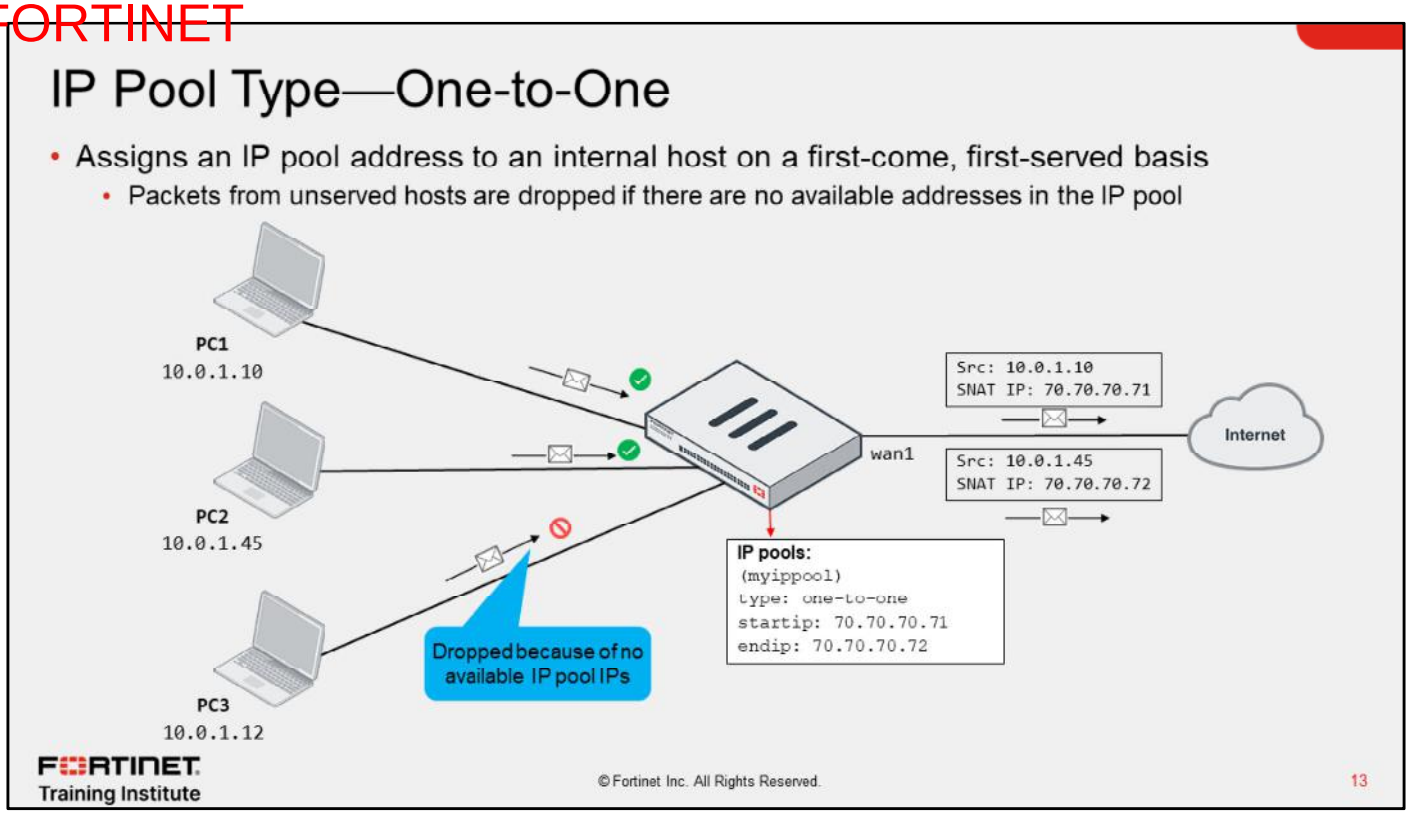

In the one-to-one pool type, FortiGate assigns an IP pool address to an internal host on a first-come, first-served basis.

There is a single mapping of an internal address to an external address. That is, an IP pool address is not shared with any other internal host, thus the name one-to-one. If there are no more addresses available in the IP pool, FortiGate drops packets from unserved hosts.

The example on this slide shows three internal hosts accessing the internet. PC1 and PC2 packets are received first by FortiGate and, therefore, served with addresses `70.70.70.71` and `70.70.70.72`, respectively. However, FortiGate drops packets sourced from PC3 because they arrived last, which is when there are no more available addresses in the IP pool to choose from.

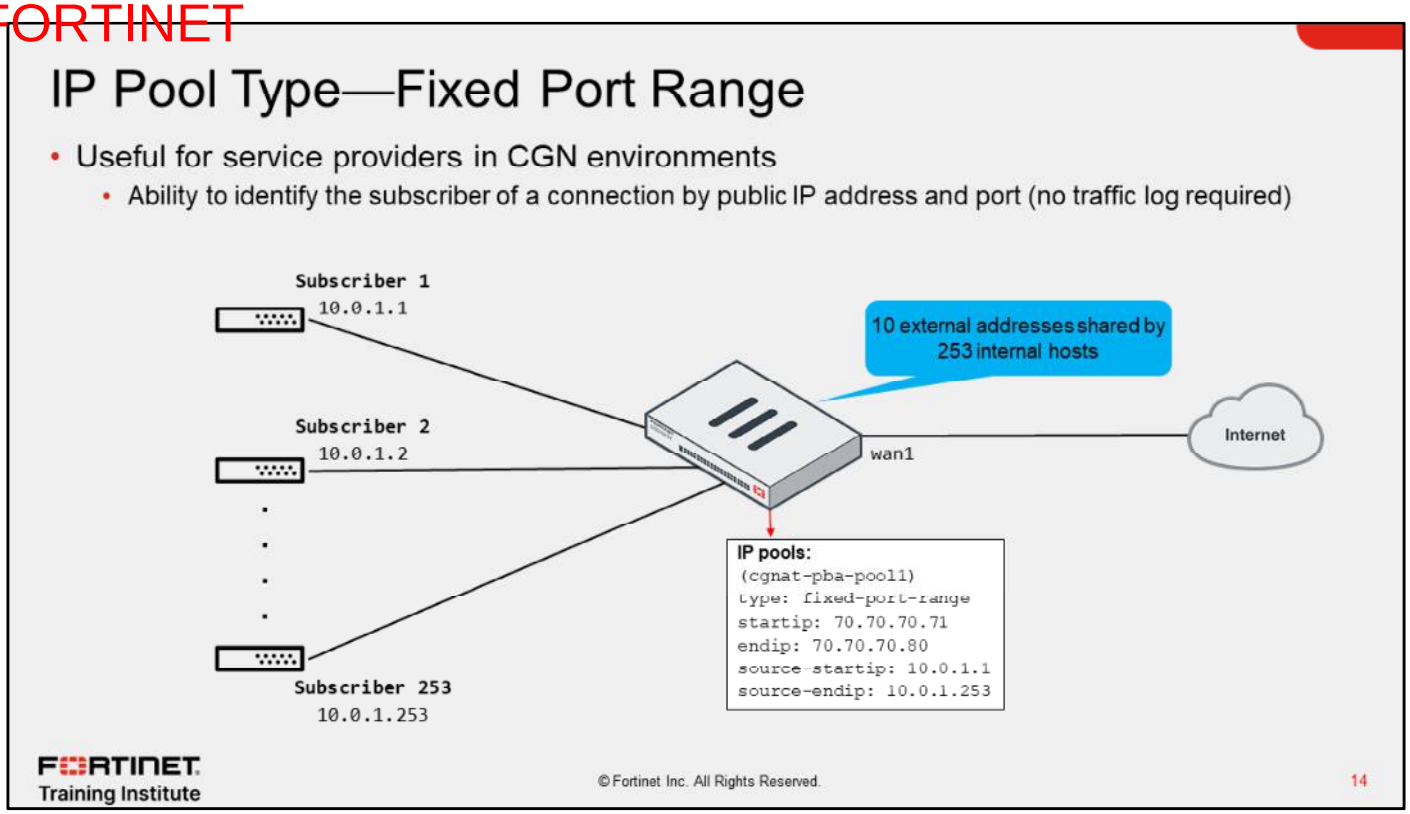

ISPs must be able to identify the subscriber responsible for a given connection should authorities require it. If the ISP performs NAT to subscriber traffic, then the traffic will share one or more public addresses. One way to track the traffic and, therefore, the NAT details for each connection, is by logging them. However, this can result in a huge number of resources that the ISP needs to dedicate for logging purposes only.

Another option is to deploy a CGN-focused feature such as a fixed port range IP pool. Fixed port range IP pools enable administrators to track connections by public address and port without having to log every session. When you configure a fixed port range IP pool, you indicate a range of external IP addresses that FortiGate uses to perform NAT on traffic sourced from a range of internal IP addresses. It is called fixed port range because FortiGate calculates the port block size and the number of available port blocks for the IP pool based on the number of configured internal and external IP addresses. FortiGate then allocates one or more port blocks to internal hosts when performing NAT, which is what enables the administrator to track connections without having to log them.

The example on this slide shows a fixed port range IP pool. The internal address range `10.0.1.1` to `10.0.1.253` maps to the external address range `70.70.70.71` to `70.70.70.80`. That is, FortiGate shares ten external addresses with 253 internal addresses.

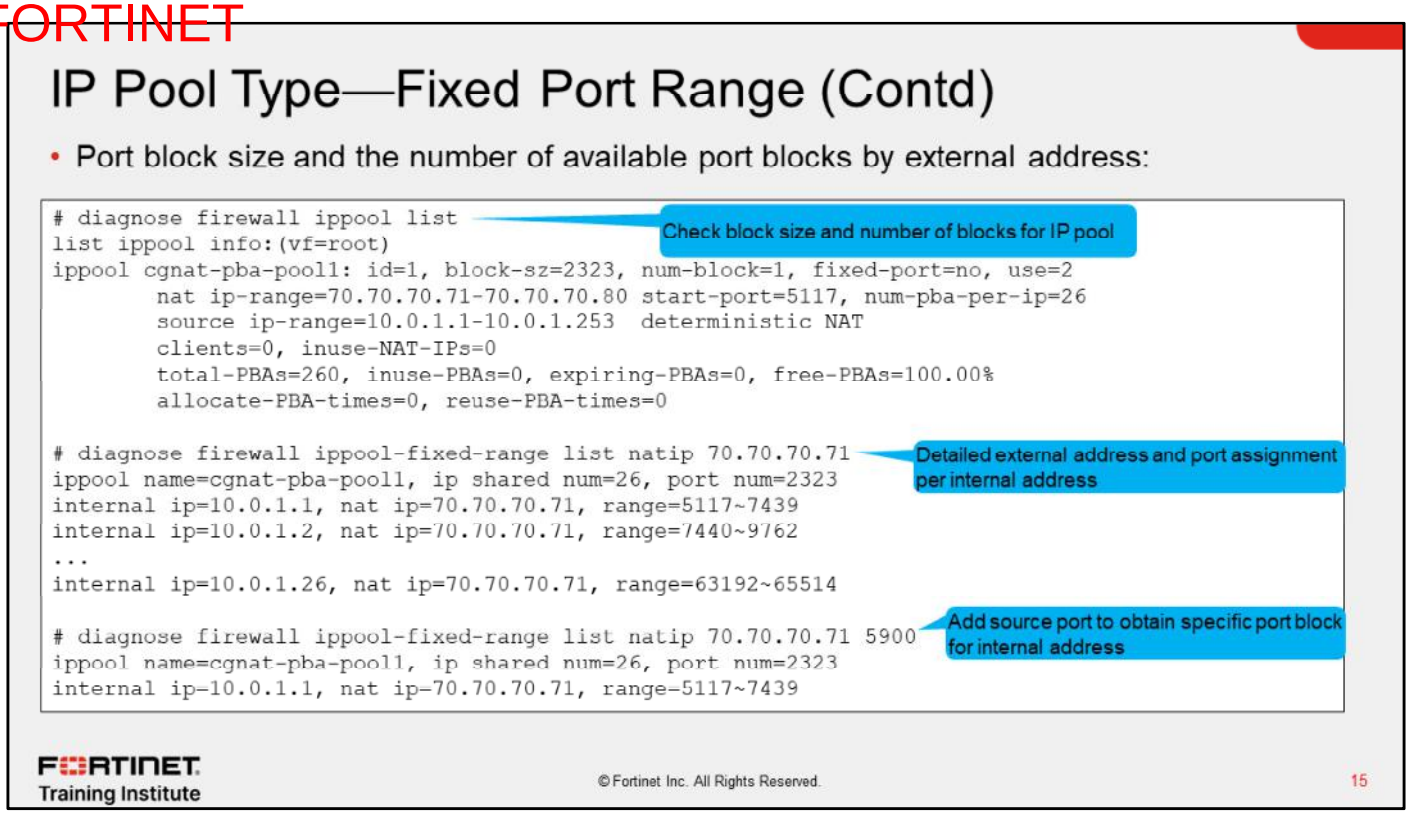

You can use the `diagnose firewall ippool list` command to identify the block size and number of blocks assigned to each external address in the fixed port range IP pool.

You can also use the `diagnose firewall ippool-fixed-range list natip` command to view detailed external address and port assignment information per internal address, as shown on this slide. The result is that you can identify subscribers by providing the public address and port of a connection.

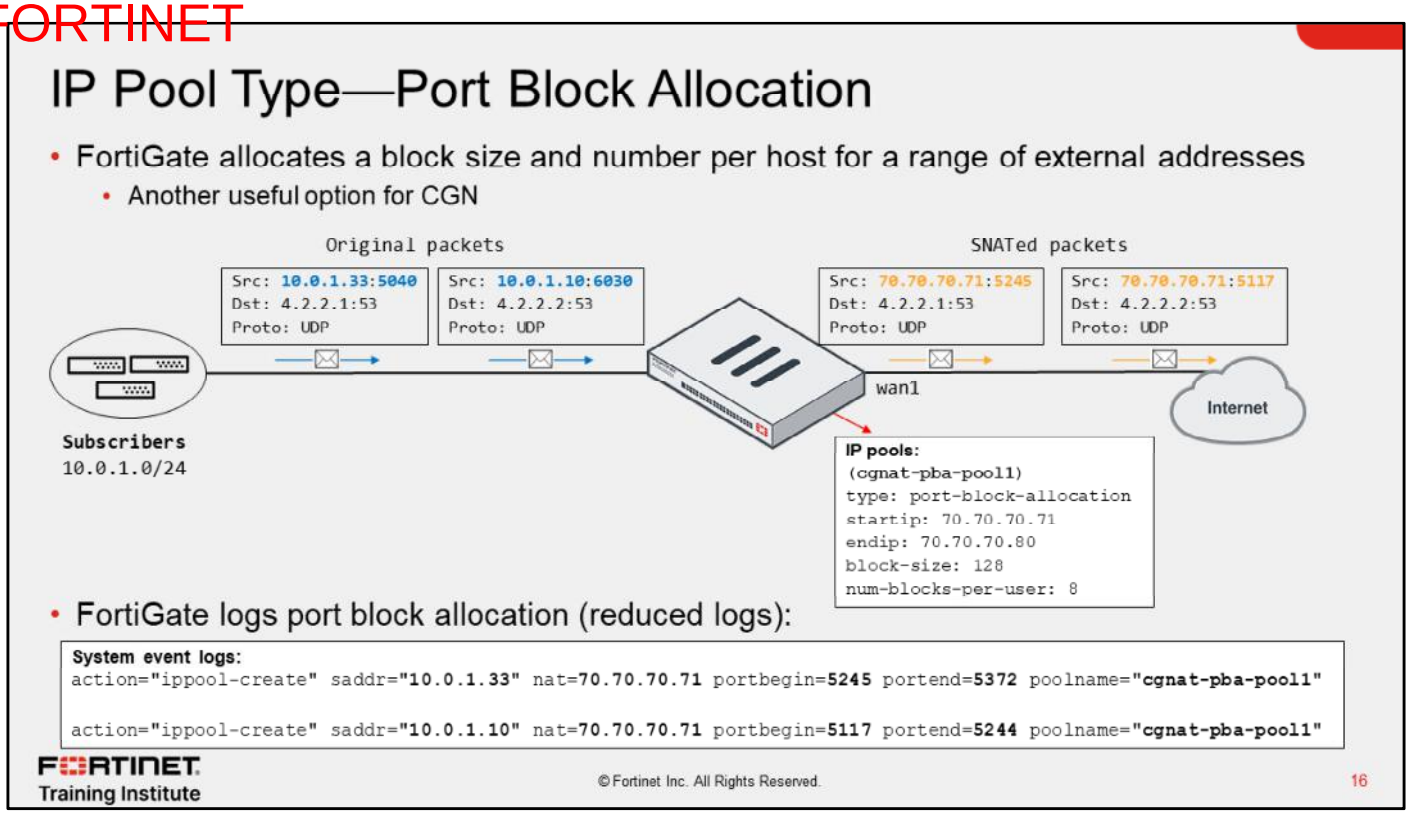

The port block allocation IP pool is also a useful option for CGN. It give administrators a more flexible way to control user port allocation for NAT. Unlike the fixed port range IP pool, which requires you to define internal and external IP address ranges, with port block allocation, you define the external IP address range only. You must also indicate the block port size and the number of blocks that FortiGate allocates to each host (or source IP address). The result is that each source IP address is limited to the number of blocks and ports configured in the IP pool, thus preventing port exhaustion caused by a few hosts.

For logging purposes, when FortiGate allocates a port block to a host, it generates a system event log to inform the administrator. The administrator can then look at the system event logs to identify internet connections made by a device should the authorities require such information. That is, like the fixed port block case, the administrator doesn't have to log the traffic for connection identification purposes.

The example on this slide shows how port block allocation assignment takes place. FortiGate allocates port blocks on a first-come, first-served basis. The port block allocation is made when FortiGate receives a packet from unserved hosts. In the example, `10.0.1.10` and `10.0.1.33` are unserved hosts that try to access the internet. FortiGate then allocates the port blocks to each host and performs the respective SNAT on traffic. Upon allocation, FortiGate also generates system event logs with the port block allocation details to inform the administrator.

Note that the system event logs shown on this slide have been cut to fit the slide.

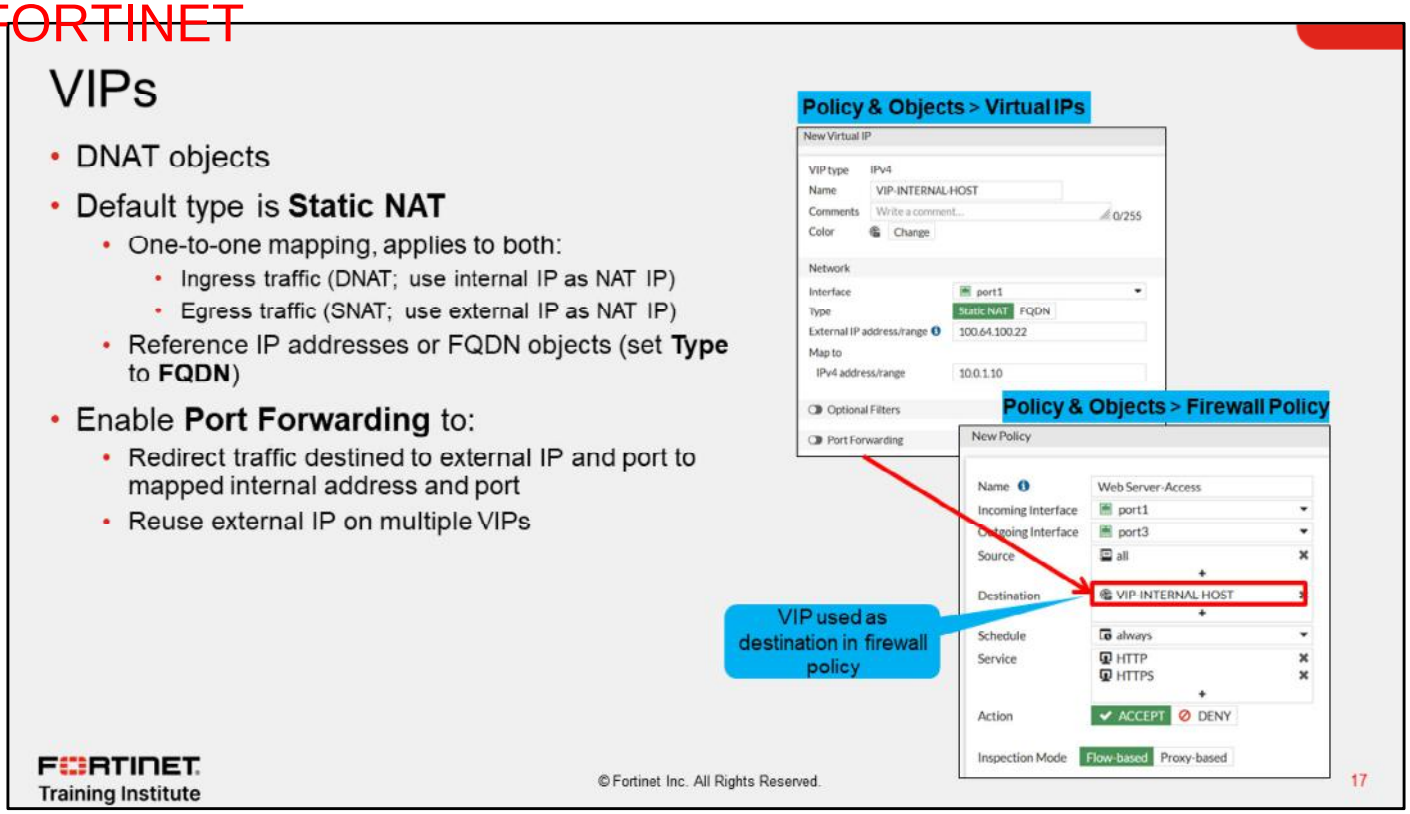

VIPs are DNAT objects. For sessions matching a VIP, the destination address is translated; usually a public internet address is translated to the private network address of a server. VIPs are selected in the firewall policy **Destination** field.

The default VIP type is **Static NAT**. This is a one-to-one mapping. This means that:

1. FortiGate performs DNAT on ingress traffic destined to the external IP address defined in the VIP, regardless of the protocol and port of the connection, provided the matching firewall policy references the VIP as **Destination**.
2. FortiGate uses as NAT IP the external IP address defined in the VIP when performing SNAT on all egress traffic sourced from the mapped address in the VIP, provided the matching firewall policy has NAT enabled. That is, FortiGate doesn't use the egress interface address as NAT IP.

Note that you can override the behavior described in step 2 by using an IP pool. You can also select **FQDN** as **Type**. When you select **FQDN**, you can configure FQDN address objects as external and internal IP addresses. This enables FortiGate to automatically update the external and internal IP addresses used by the VIP in case the FQDN resolved address change.

Optionally, you can enable **Port Forwarding** on the VIP to instruct FortiGate to redirect the traffic matching the external address and port in the VIP to the mapped internal address and port. When you enable port forwarding, FortiGate no longer performs one-to-one mapping. This means that you can reuse the same external address and map it to different internal addresses and ports provided the external port is unique. For example, you can configure a VIP so connections to the external IP 70.70.70.70 on port 8080 map to the internal IP 192.168.0.70 on port 80. You can then configure another VIP so connections to the external IP 70.70.70.70 on port 8081 map to the internal IP 192.168.0.71 on port 80.

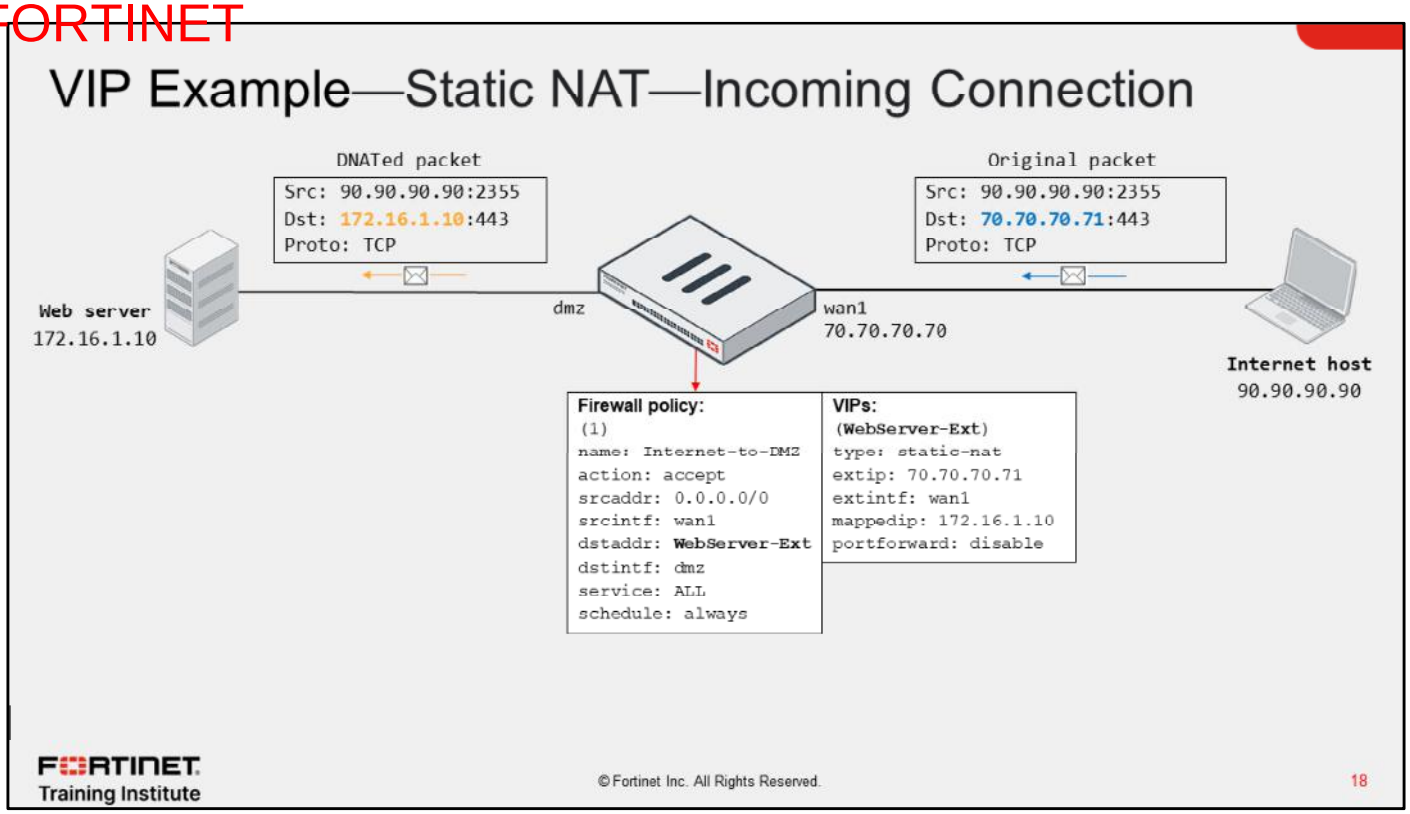

In the example shown on this slide, the internet host initiates a connection to `70.70.70.71` on TCP port 443. On FortiGate, the traffic matches the firewall policy ID 1, which references the `WebServer-Ext` VIP as destination. Because the VIP is configured as static NAT and has port forwarding disabled, then FortiGate translates the destination address of the packet to `172.16.1.10` from `70.70.70.71`. Note that the destination port doesn't change because port forwarding is disabled.

Also note that the external interface address is different from the external address configured in the VIP. This is not a problem as long as the upstream network has its routing properly set. You can also enable ARP reply on the VPN (enabled by default) to facilitate routing on the upstream network. You will learn more about ARP reply in this lesson.

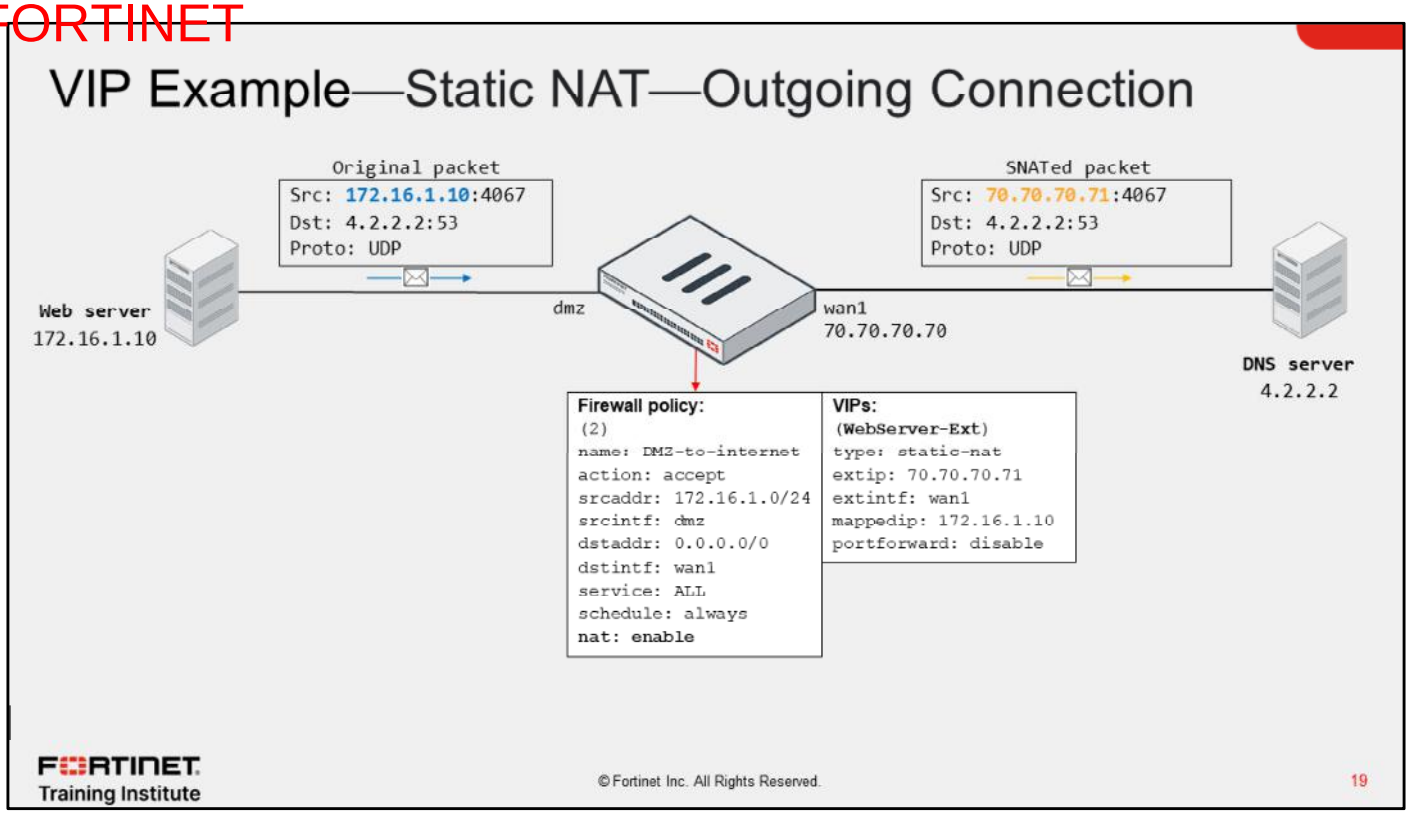

Now, suppose that the internal web server (`172.16.1.10`) initiates a DNS connection to the internet DNS server (`4.2.2.2`). On FortiGate, the traffic matches the firewall policy ID 2, which has `nat` enabled. Because the source address matches the internal address of the VIP, and because the VIP is configured as static NAT with port forwarding disabled, FortiGate translates the source address of the packet to `70.70.70.71` from `172.16.1.10`. Note that FortiGate doesn't have to perform PAT because the static NAT VIP equals one-to-one mapping. That is, the external IP is used by the web server only for SNAT.

Also note that FortiGate uses the VIP external address for SNAT if the VIP is referenced in an incoming firewall policy. That is, if you don't configure firewall policy ID 1, which is shown on the previous slide, or if you disable the firewall policy, then FortiGate doesn't automatically use the external IP for translating the source address of the web server. Instead, FortiGate uses the egress interface address (`70.70.70.70`).

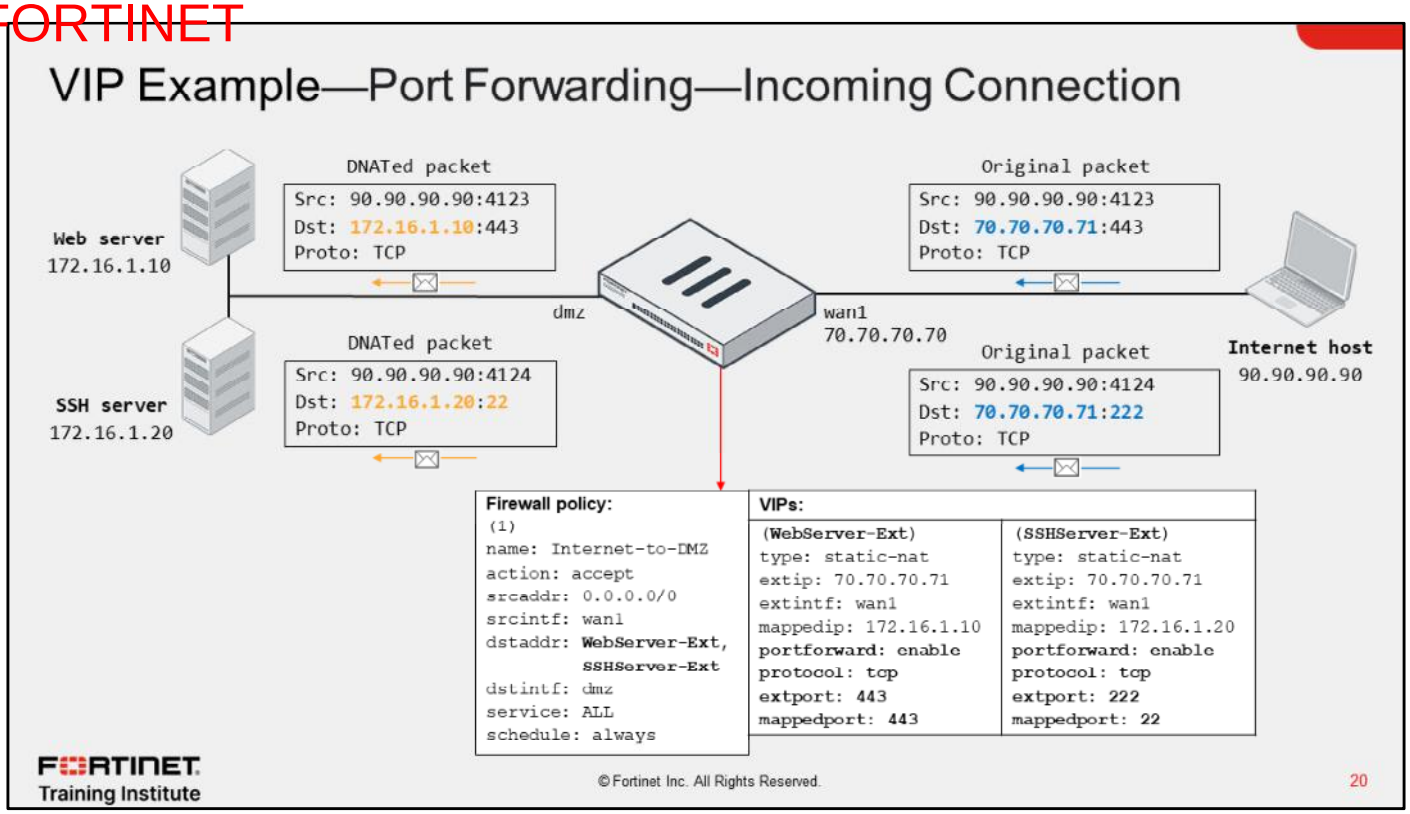

The example on this slide shows how FortiGate handles two incoming connections to the same external address, but on different ports. FortiGate forwards each connection to a different internal host based on the VIP mapping settings. This is possible because port forwarding is enabled on the VIPs, which enables FortiGate to redirect the external traffic to the corresponding internal address and port, while using the same external address.

Both connections match the firewall policy ID, which references two VIPs as destination. The HTTPS connection matches the `WebServer-Ext` VIP, and the SSH connection matches the `SSHServer-Ext` VIP. Note that for the SSH connection, FortiGate also translates the destination port to 22 from 222.

Although not shown on this slide, outgoing connections sourced from the web and SSH server would result in FortiGate using as NAT IP the egress interface address for SNAT, providing there is a matching firewall policy with `nat` enabled.

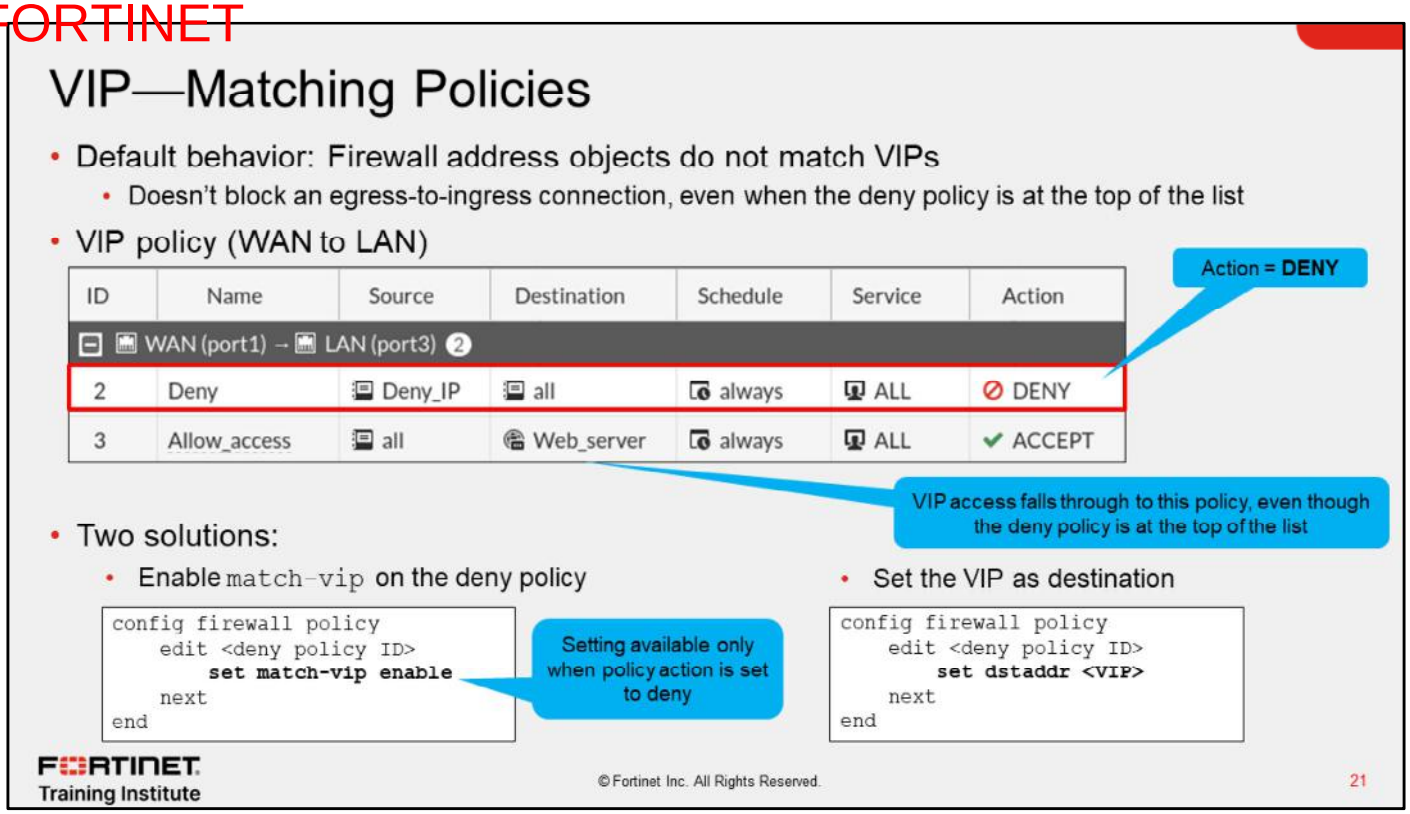

In FortiOS, VIPs and firewall address objects are completely different. They are stored separately with no overlap. This means that, by default, firewall address objects do not match VIPs.

In the example shown on this slide, the destination of the first firewall policy is set to **all**. Even though this means all destination addresses (`0.0.0.0/0`), by default, this doesn't include the external addresses defined on the VIPs. The result is that traffic destined to the external address defined on the **Web_server** VIP skips the first policy and matches the second policy instead.

But what if you want the first policy to block all incoming traffic to all destinations, including the traffic destined to any VIPs?. This is useful if your network is under attack, and you want to temporarily block all incoming external traffic. You can do this by enabling `match-vip` on the first firewall policy. Enabling `match-vip` instructs FortiGate to also check for VIPs during policy evaluation. Note that the `match-vip` setting is available only when the firewall policy action is set to **DENY**.

In case you want to block only traffic destined to one or more VIPs, you can reference the VIPs as the destination address on the deny firewall policy.

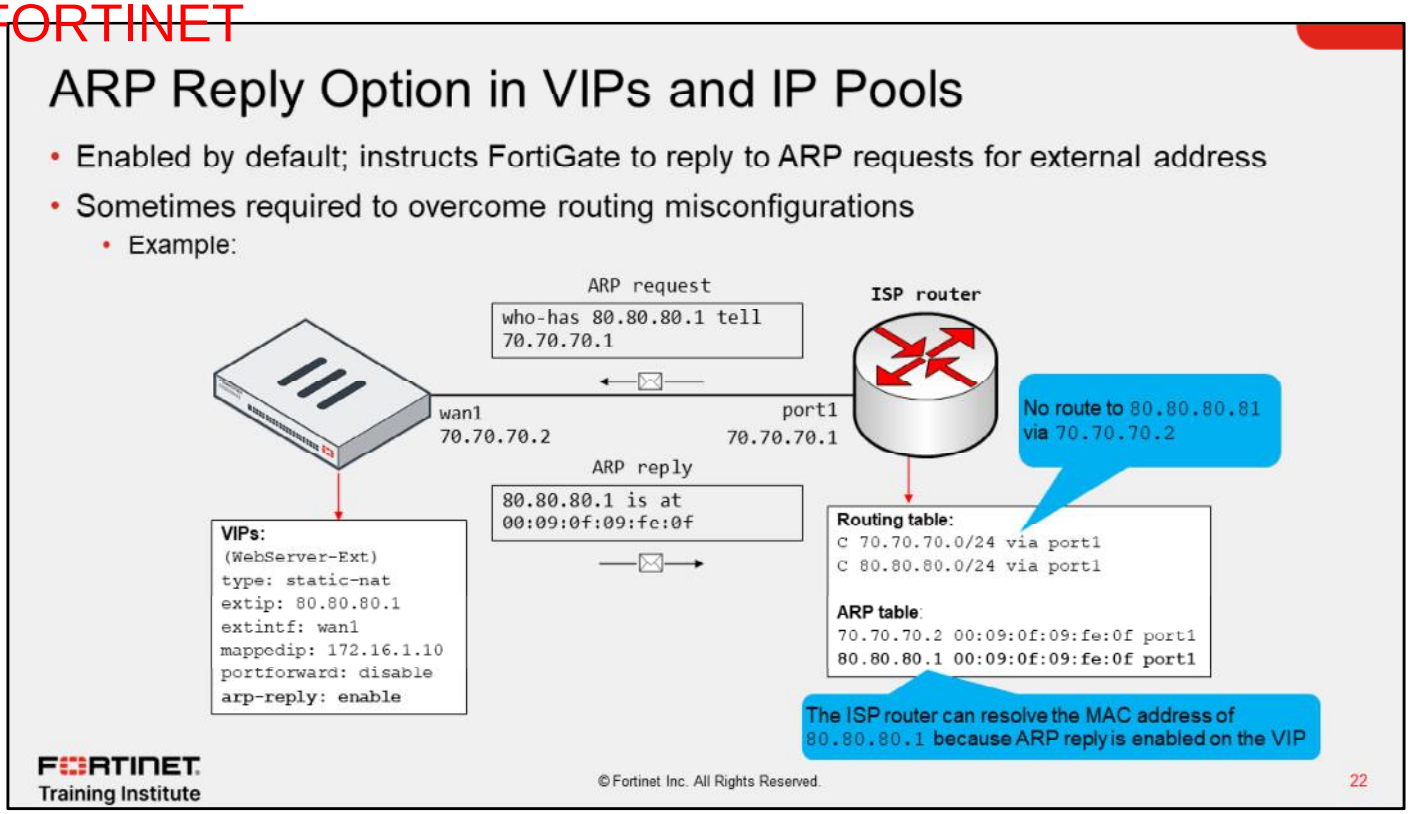

When you configure a VIP or an IP pool, ARP reply is enabled by default. When ARP reply is enabled, FortiGate replies to incoming ARP requests for the external address configured in the VIP and IP pools.

Enabling ARP reply is usually not required in most networks because the routing tables on the adjacent devices contain the correct next hop information, so the networks are reachable. However, sometimes the routing configuration is not fully correct, and having ARP reply enabled can solve the issue for you. For this reason, it's a best practice to keep ARP reply enabled.

Consider the example shown on this slide, which shows an internet connection between FortiGate and an ISP router. The example also shows a simplified version of the ISP router routing table and ARP table.

The ISP assigns the FortiGate administrator the public subnet `80.80.80.0/24` to deploy internet-facing services. The administrator configured the VIP shown on this slide to provide internet users with access to the company web server. While testing, the administrator confirms that internet users can reach the web server at `80.80.80.1`.

However, the administrator is likely unaware that having ARP reply enabled was key for a successful connectivity. The reason is that the ISP router doesn't have a route in its routing table to access the `80.80.80.0/24` subnet through the `70.70.70.2` gateway. Instead, the routing table contains a connected route for the subnet through port1. The result is that the ISP router generates ARP requests out of port1 to resolve the MAC address of any of the addresses in the `80.80.80.0/24` subnet. Nonetheless, because FortiGate responds to ARP requests for the external address in the VIP, the ISP router is able to resolve the MAC address successfully.

Knowledge Check

1. What is the default IP pool type?
   A. One-to-one
   ✓B. Overload

2. Which of the following is the default VIP type?
   ✓A. `static-nat`
   B. `load-balance`

FORTINET Training Institute

23

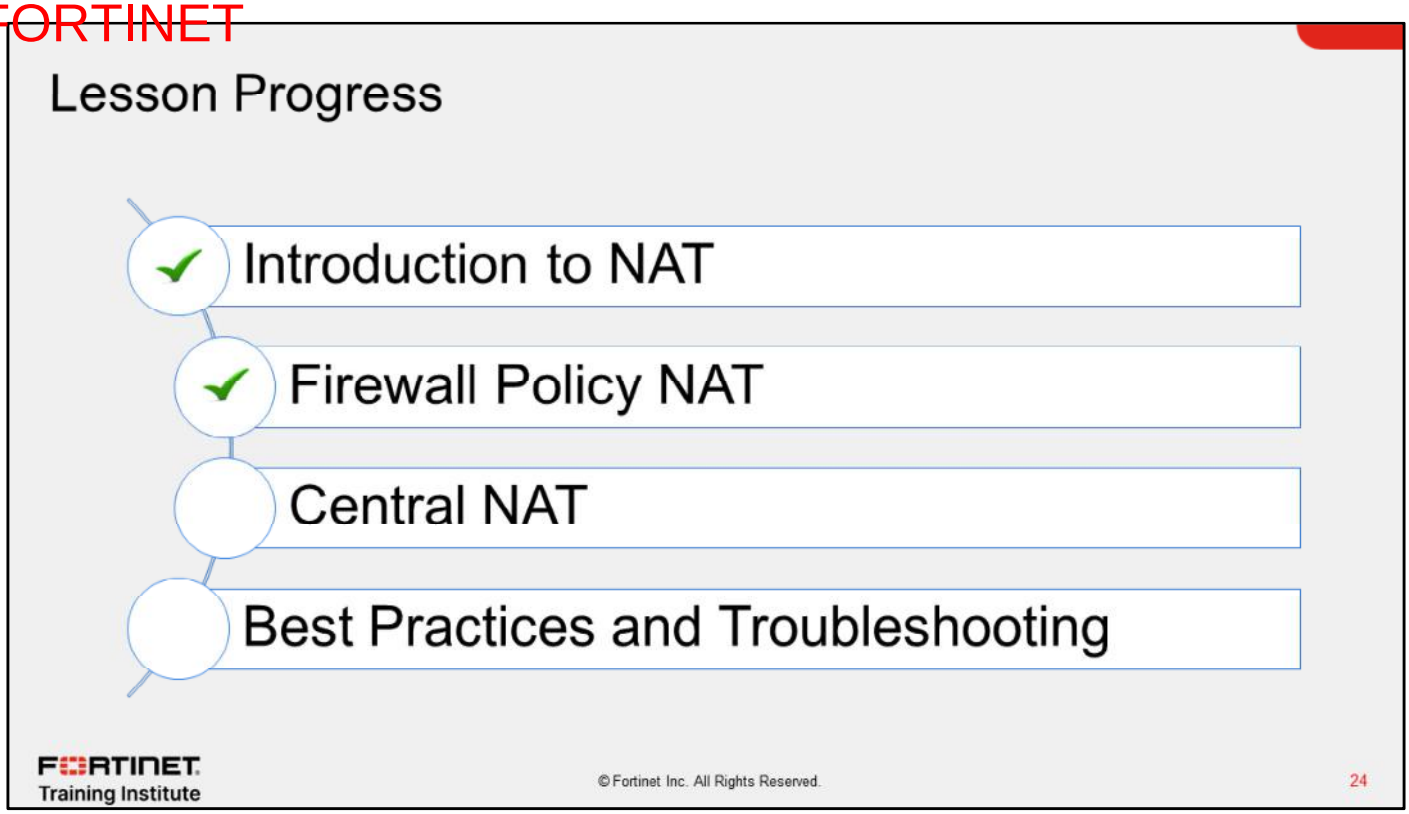

Good job! You now understand firewall policy NAT.

Now, you'll learn about central NAT.

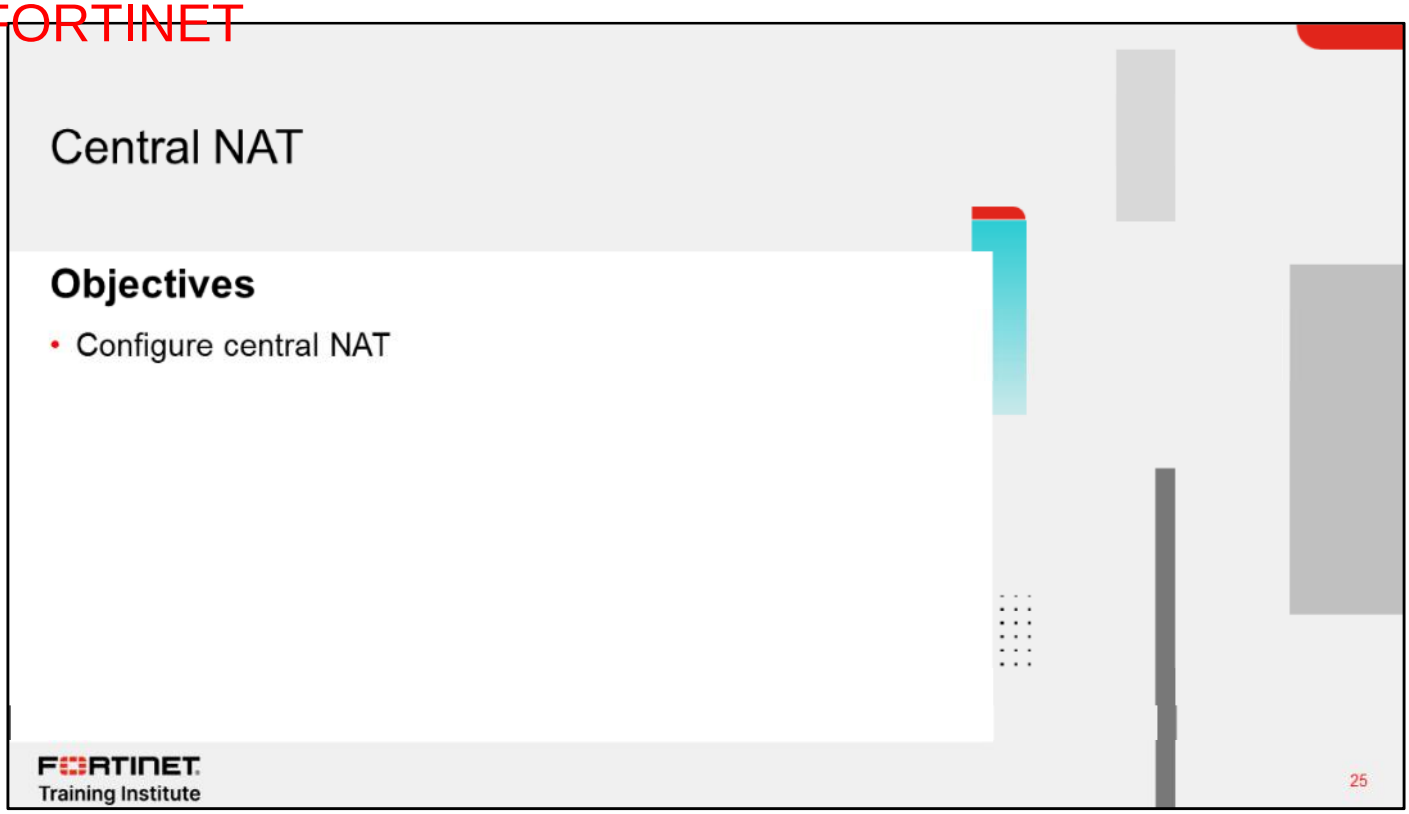

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in configuring central NAT to perform SNAT and DNAT, you will be able to use NAT on a more granular level to control IP address, protocol, and port translation.

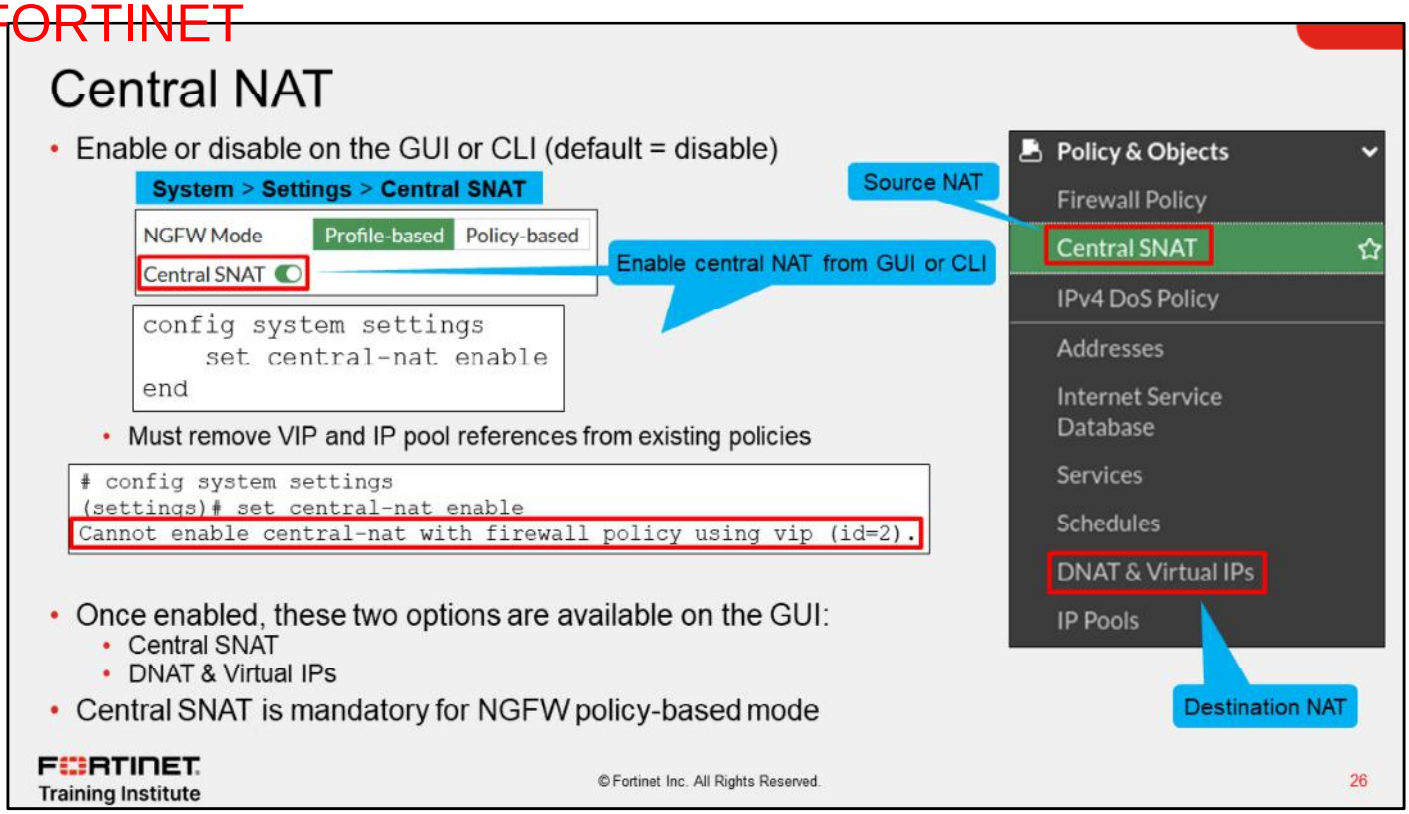

By default, central NAT is disabled. You can enable it on the CLI or the GUI. After central NAT is enabled, the following two options are available to be configured on the GUI:

- **Central SNAT**
- **DNAT & Virtual IPs**

What happens if you try to enable central NAT, but there are still IP pools or VIPs configured in firewall policies?

The CLI does not allow this and presents a message referencing the firewall policy ID with the VIP or IP pool. You *must* remove VIP or IP pool references from existing firewall policies in order to enable central NAT.

Central SNAT is mandatory for the new NGFW policy-based mode. This means SNAT behaves only according to the NAT settings found by clicking **Policy & Objects** > **Central SNAT**.

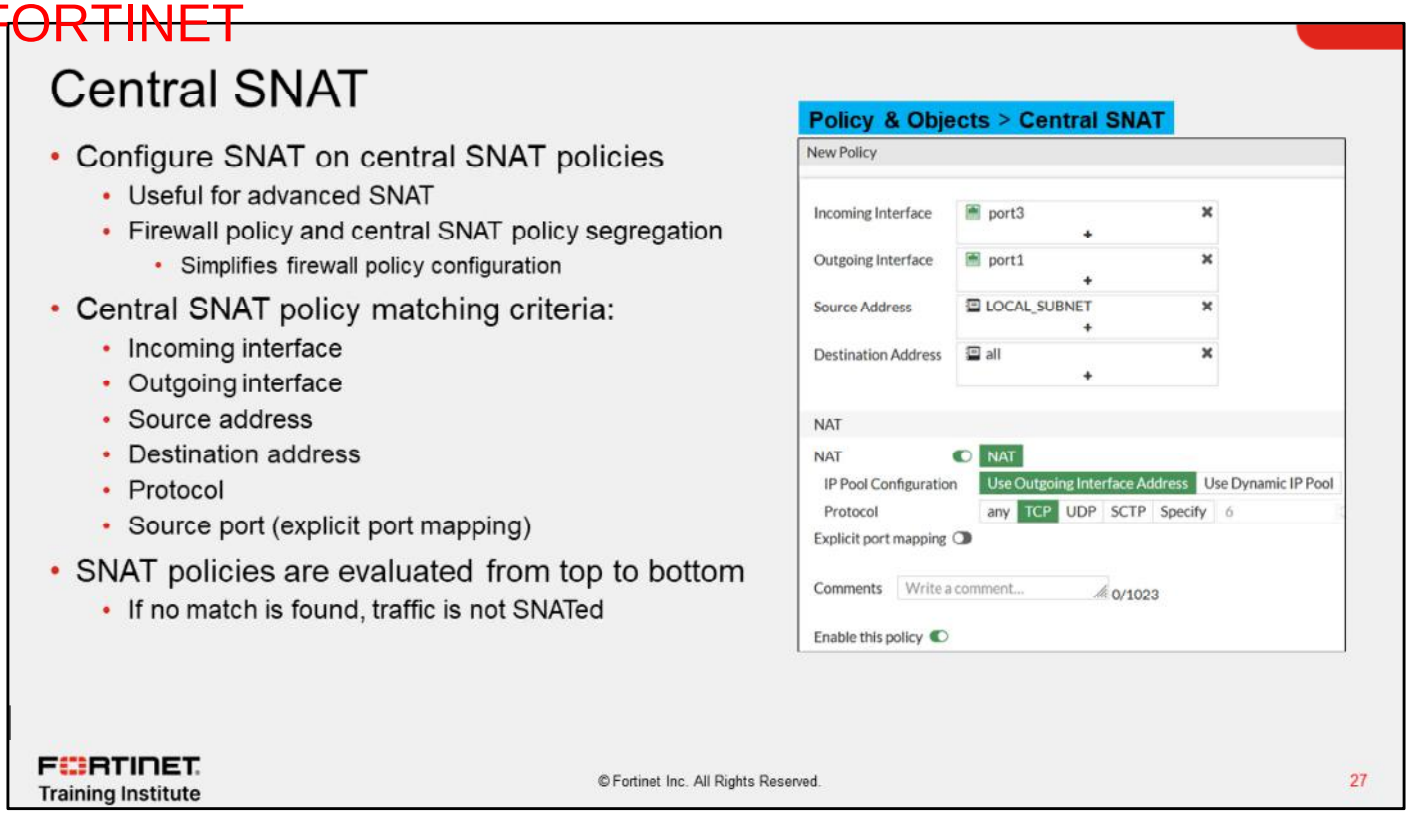

When you enable central NAT, you configure SNAT on the central SNAT page on the FortiGate GUI.

The main benefit of using central NAT for SNAT is firewall policy and central SNAT policy segregation. This is particularly useful for advanced SNAT configurations comprising multiple networks and IP pools. Instead of enabling NAT and selecting IP pools on firewall policies, you configure SNAT policies for all the accepted traffic by the firewall policies. This way, you focus your firewall policy configuration on what kind of traffic to accept, and your SNAT policies on what portion of the accepted traffic to translate and the SNAT mapping to follow. The result is that you simplify your firewall policy configuration by removing the SNAT settings from it.

When you configure SNAT policies, you can configure the following matching criteria:

- Incoming interface
- Outgoing interface
- Source address
- Destination address
- Protocol
- Source port (explicit port mapping)

You must also indicate whether you want to perform SNAT using the outgoing interface address or an IP pool. Note that if you enable central NAT mode, FortiGate doesn't perform SNAT on traffic unless you configure the corresponding matching central SNAT policy. Similarly, if the traffic doesn't match any of the configured SNAT policies, FortiGate doesn't perform SNAT on the traffic either.

Like firewall policies, SNAT policies are processed from *top to bottom* and, if a match is found, the source address and source port are translated based on the central SNAT policy mapping settings.

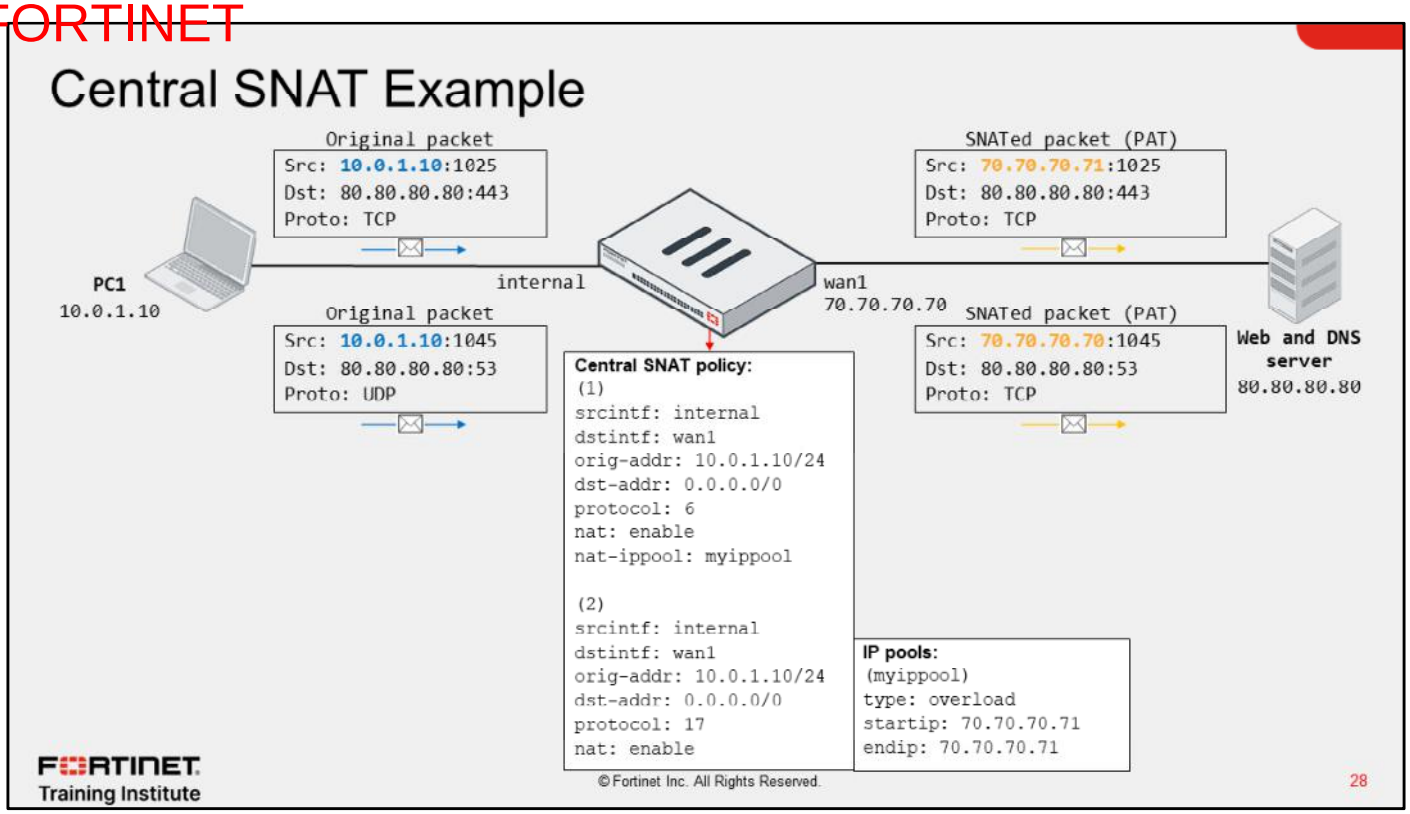

In the example shown on this slide, PC1 (`10.0.1.10`) initiates two connections to the external server (`80.80.80.80`). The HTTPS connection matches central SNAT policy ID 1 and, therefore, the source address is translated to the IP pool address (`70.70.70.71`). The DNS connection matches central SNAT policy ID 1, which doesn't reference an IP pool. The result is that the source address of the DNS connection is translated to the external interface address (`70.70.70.70`).

Although not shown on this slide, there are firewall policies configured that accept both connections.

Now, what if PC1 initiates an ICMP connection to the server? Because there is no matching central SNAT policy, then FortiGate wouldn't perform SNAT for the ICMP connection.

Central DNAT and VIPs

- Kernel has DNAT rules based on configured VIPs
  - You no longer reference VIPs in firewall policies
- Firewall policy
  - Destination address must match the VIP mapped address
    - DNAT takes place before firewall policy lookup

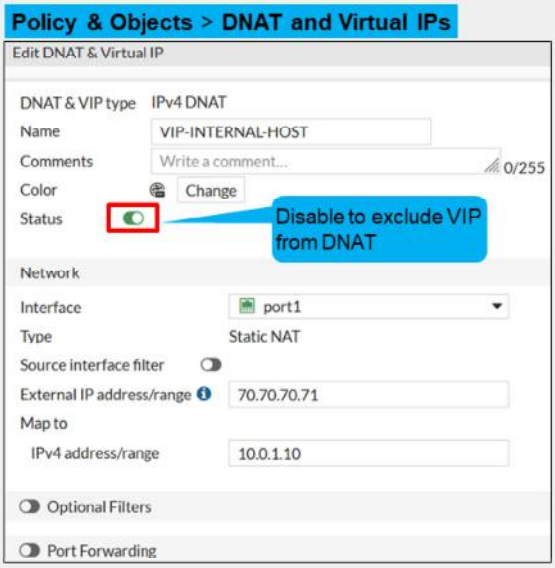

FORTINET Training Institute

29

When you enable central NAT, you no longer reference VIPs on firewall policies. Instead, FortiGate automatically creates a rule in the kernel to perform DNAT for the matching traffic based on the configured VIPs. You configure the VIPs on the **DNAT and Virtual IPs** page.

Like in the central SNAT case, you must also have a matching firewall policy that accepts the traffic you want to DNAT. However, instead of referencing the VIP, you reference the mapped internal address as destination in the firewall, and *not* the external address. This is because for ingress traffic, DNAT takes place before the firewall policy lookup. That is, FortiGate considers the translated destination address during the firewall policy lookup process.

In central NAT mode, VIPs take effect right after you create them. In case you want to exclude a VIP from DNAT, you can disable the object on the FortiGate GUI by using the **Status** button.

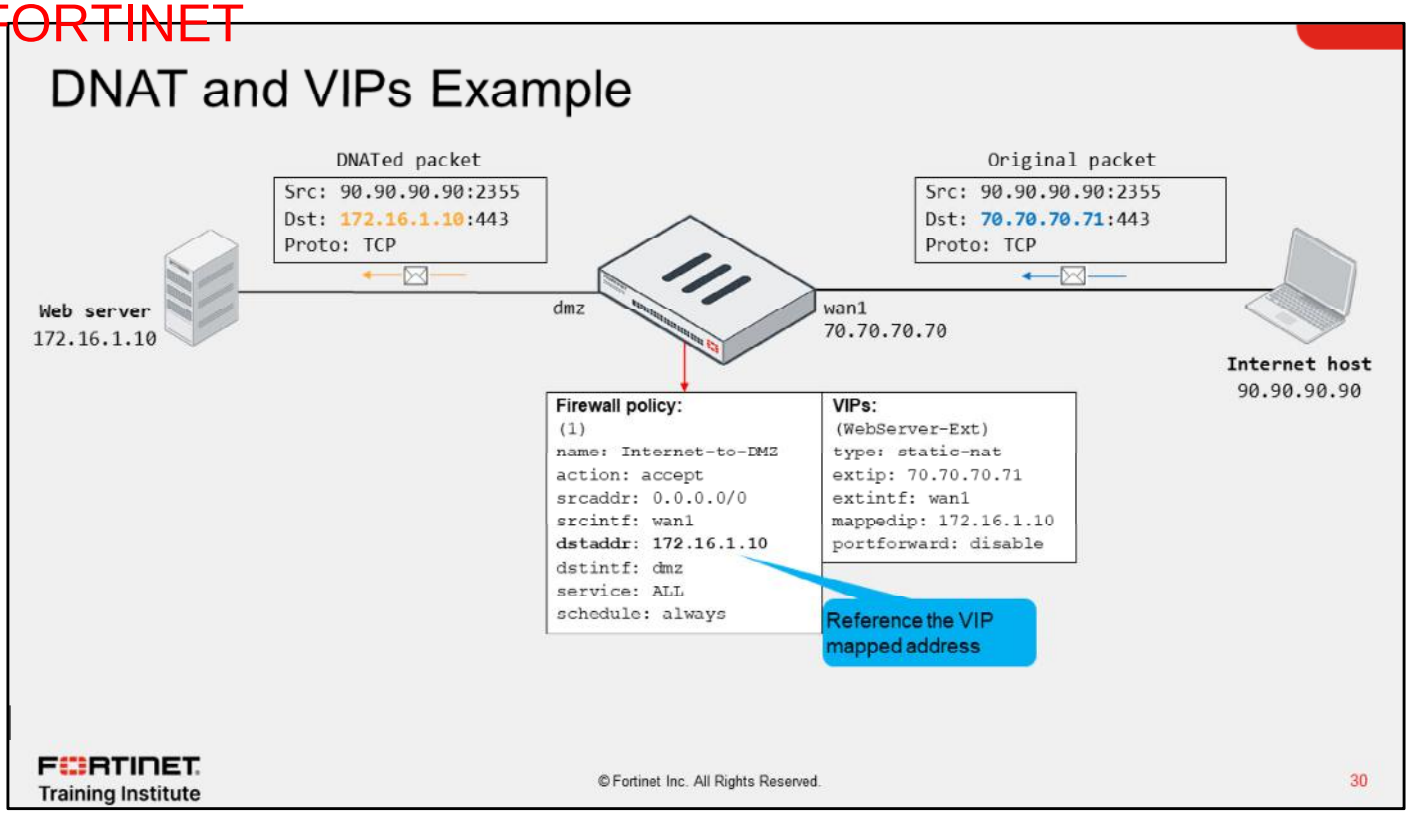

In the example shown on this slide, the internet host initiates a connection to `70.70.70.71` on TCP port 443. On FortiGate, the traffic matches the firewall policy ID 1, which references the web server internal address (`172.16.1.10`) as the destination. Because the VIP is configured as static NAT and has port forwarding disabled, then FortiGate translates the destination address of the packet to `172.16.1.10` from `70.70.70.71`.

Note that you configure the firewall policy to match the VIP mapped address as the destination, and *not* the VIP external address.

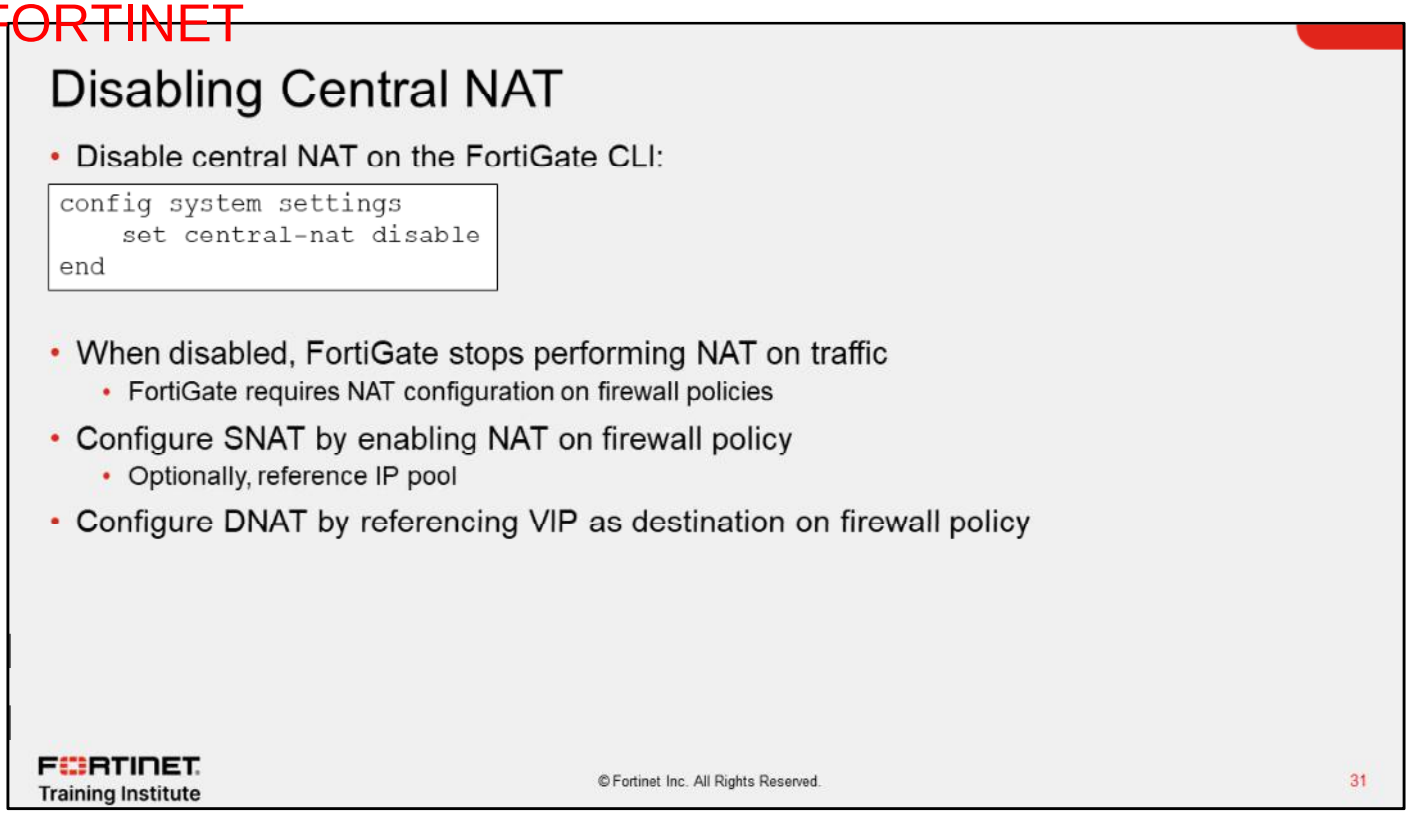

You can disable central NAT on the FortiGate CLI by disabling `central-nat` under `config system settings`.

However, note that when you disable central NAT, FortiGate stops performing NAT on traffic because it now requires the NAT configuration to be applied on the corresponding firewall policies. For FortiGate to perform SNAT, you must enable NAT on the respective firewall policy and, optionally, reference the IP pool. For DNAT, you must reference the VIP object as the destination on the corresponding firewall policies.

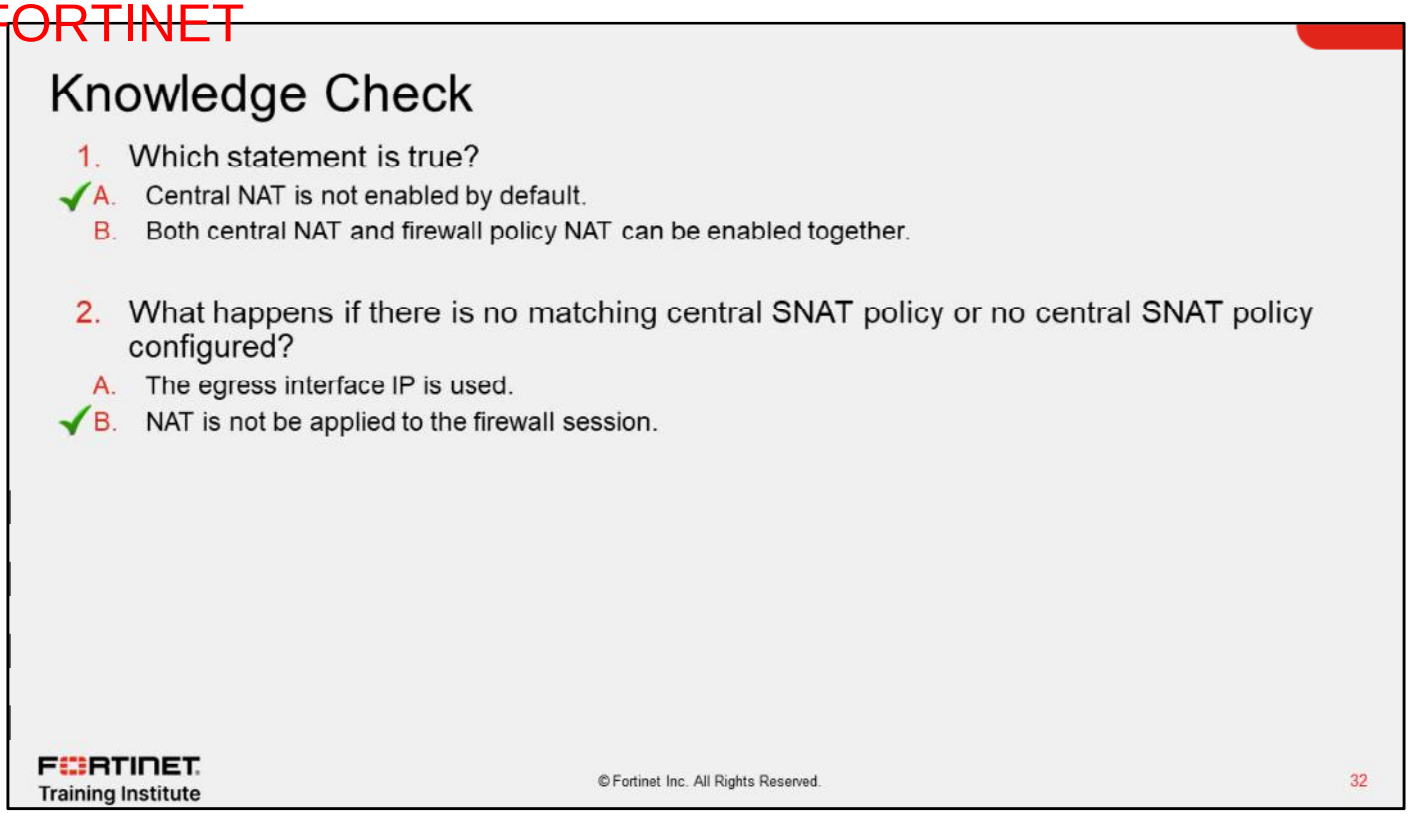

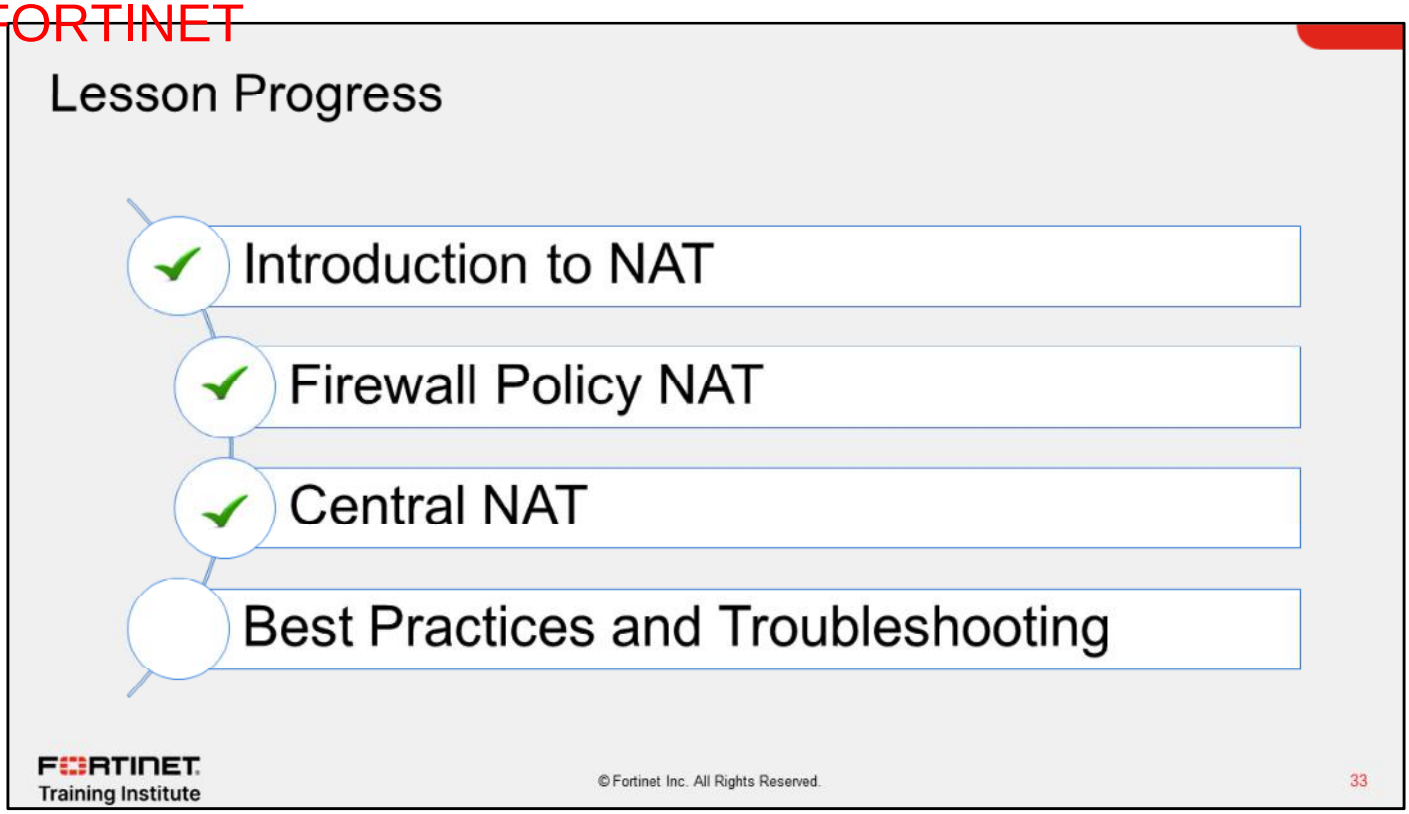

Good job! You now understand central NAT.

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

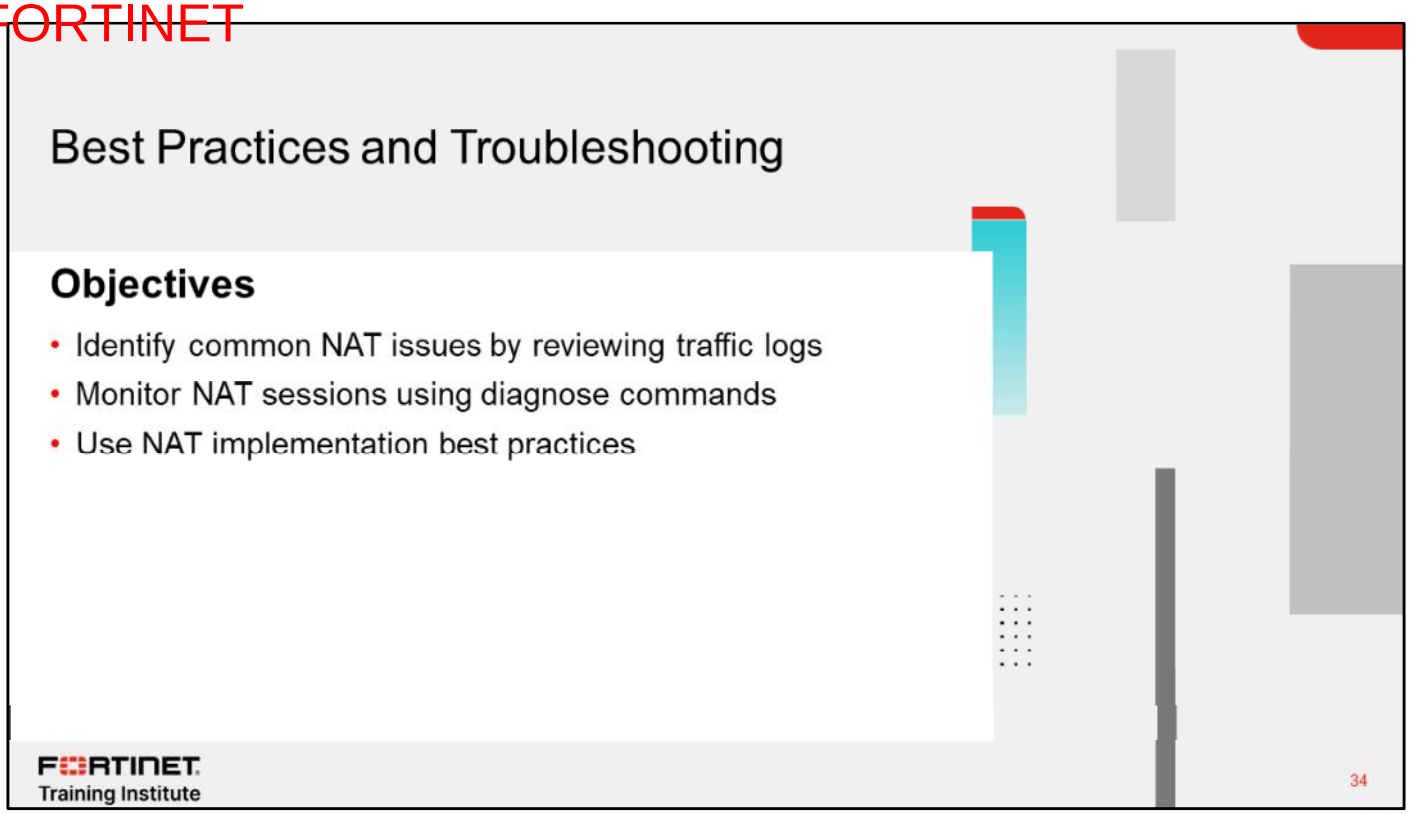

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in using traffic logs, diagnose commands, and best practices for NAT implementation, you should be able to monitor and troubleshoot common NAT issues, and successfully implement NAT in your network.

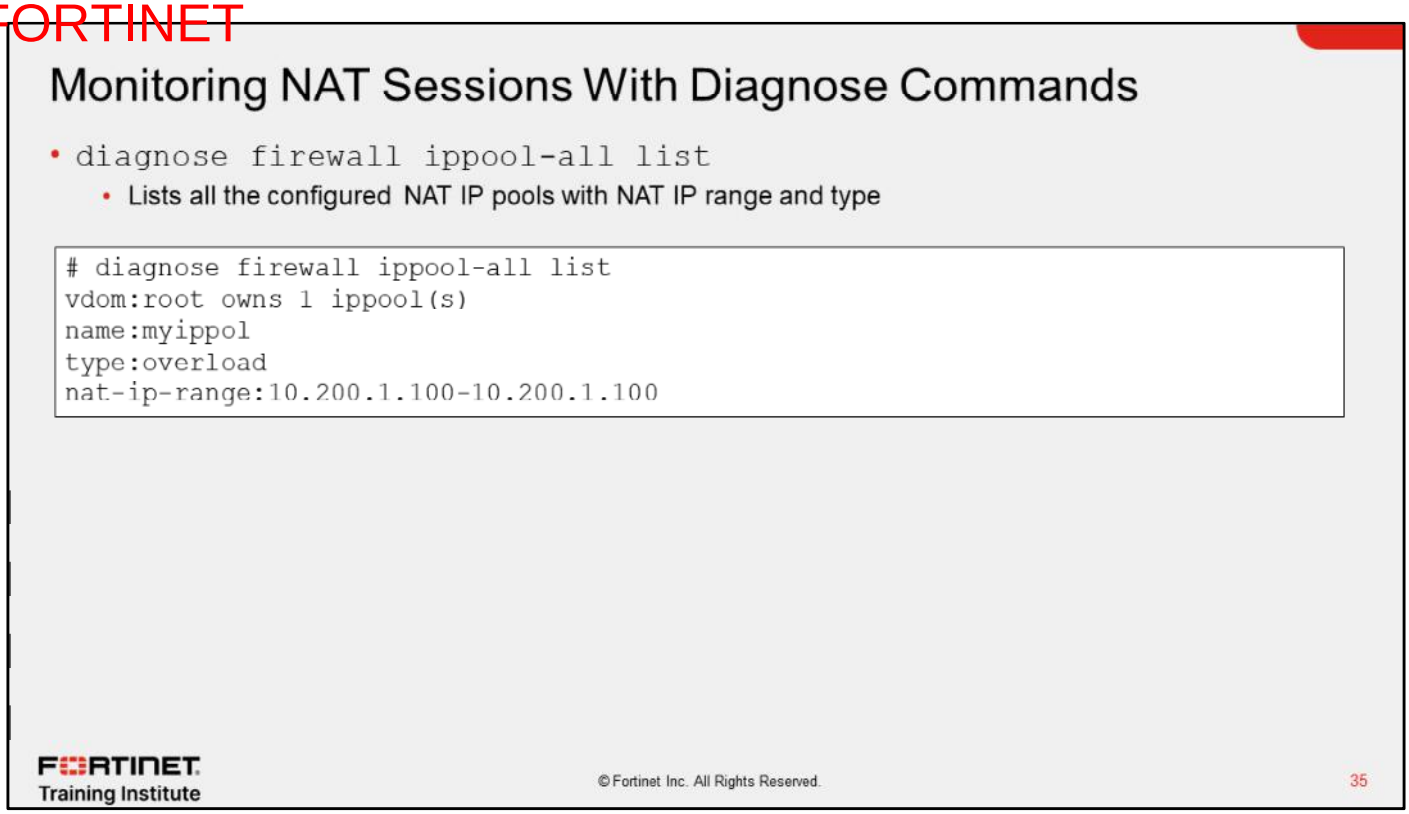

You can run the `diagnose firewall ippool-all list` command to display the configured IP pools and their settings.

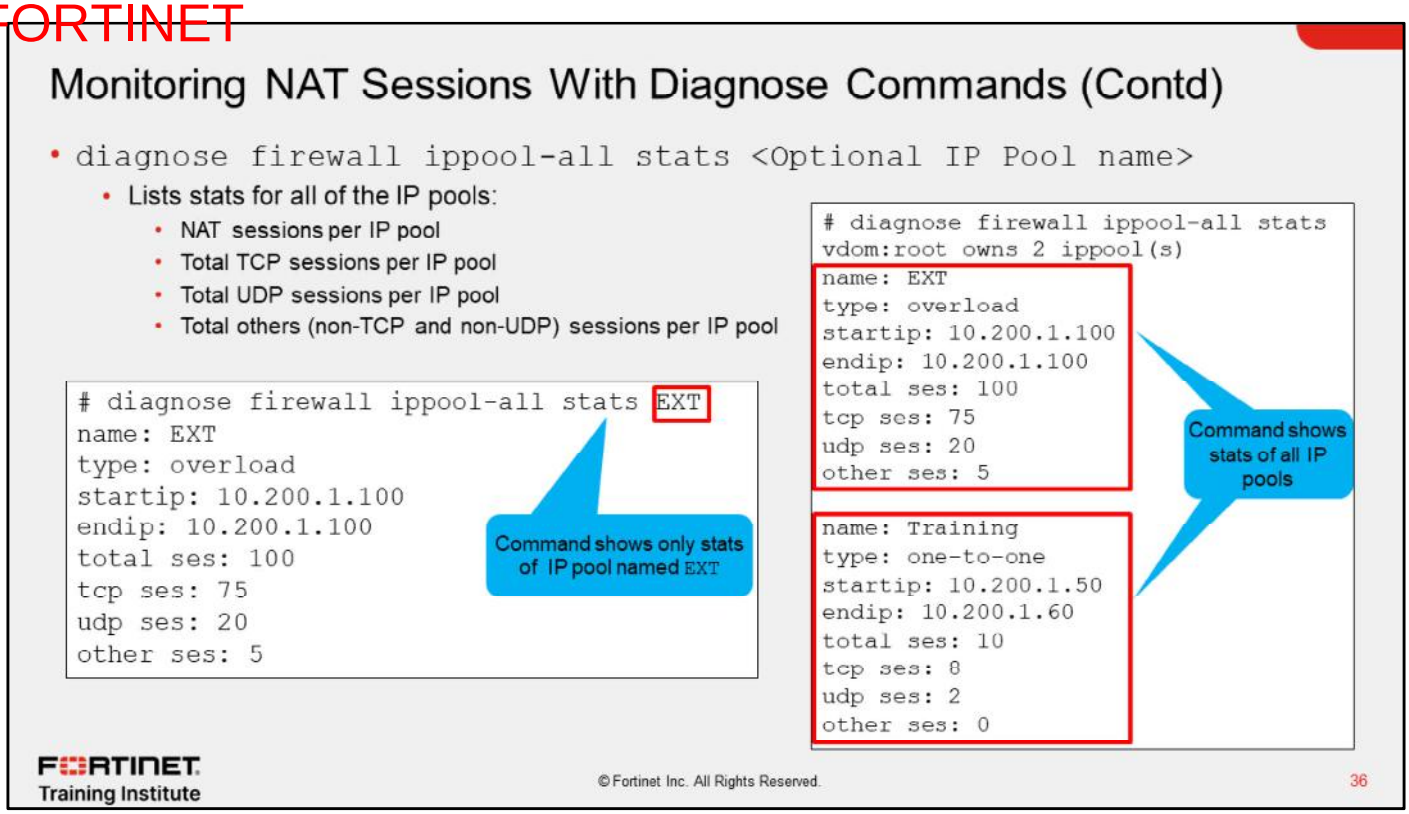

The `diagnose firewall ippool-all stats` command shows the stats for all IP pools.

The `stats` command provides the following data and information:

- NAT sessions per IP pool
- Total TCP sessions per IP pool
- Total UDP sessions per IP pool
- Total others (non-TCP and non-UDP) sessions per IP pool

Optionally, you can filter the output for a specific IP pool by using the name of the IP pool.

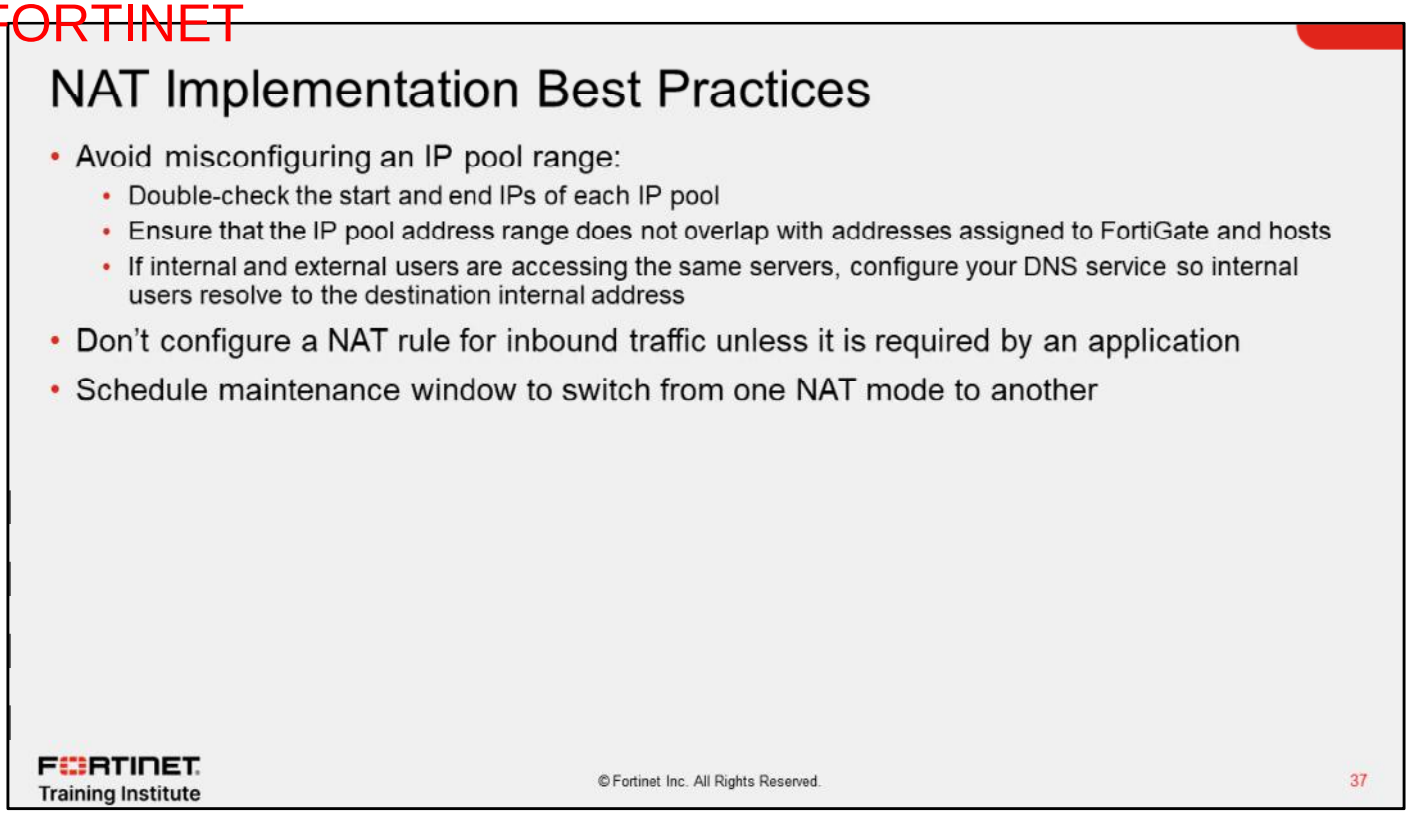

Use the following best practices when implementing NAT:

- Avoid misconfiguring an IP pool range:
  - Double-check the start and end IPs of each IP pool.
  - Ensure that the IP pool address range does not overlap with addresses assigned to FortiGate interfaces or to any hosts on directly connected networks.
  - If you have internal and external users accessing the same servers, configure your DNS services so internal users resolve to use the destination internal address instead of its external address defined in the VIP.
- Don't configure a NAT rule for inbound traffic unless it is required by an application. For example, if there is a matching NAT rule for inbound SMTP traffic, the SMTP server might act as an open relay.
- You must schedule a maintenance window to switch from central NAT mode to firewall policy NAT mode, or from firewall policy NAT mode to central NAT mode. Switching between NAT modes can create a network outage.

Knowledge Check

1. An administrator wants to check the total number of TCP sessions for an IP pool named `INTERNAL`. Which CLI command should the administrator use?

✓ A. `diagnose firewall ippool-all stats INTERNAL`
B. `diagnose firewall ippool-all list INTERNAL`

FORTINET Training Institute

38

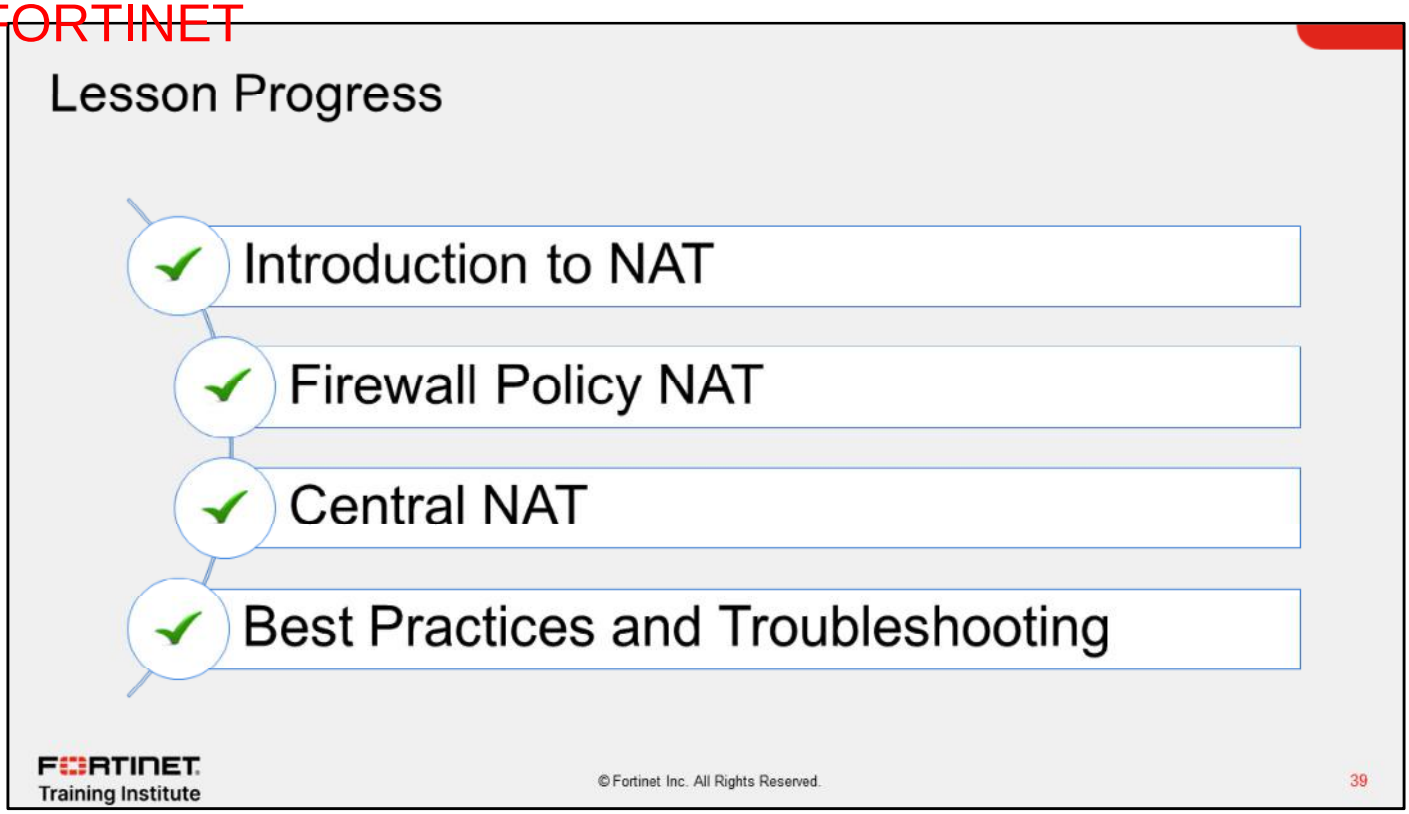

Congratulations! You have completed this lesson.

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

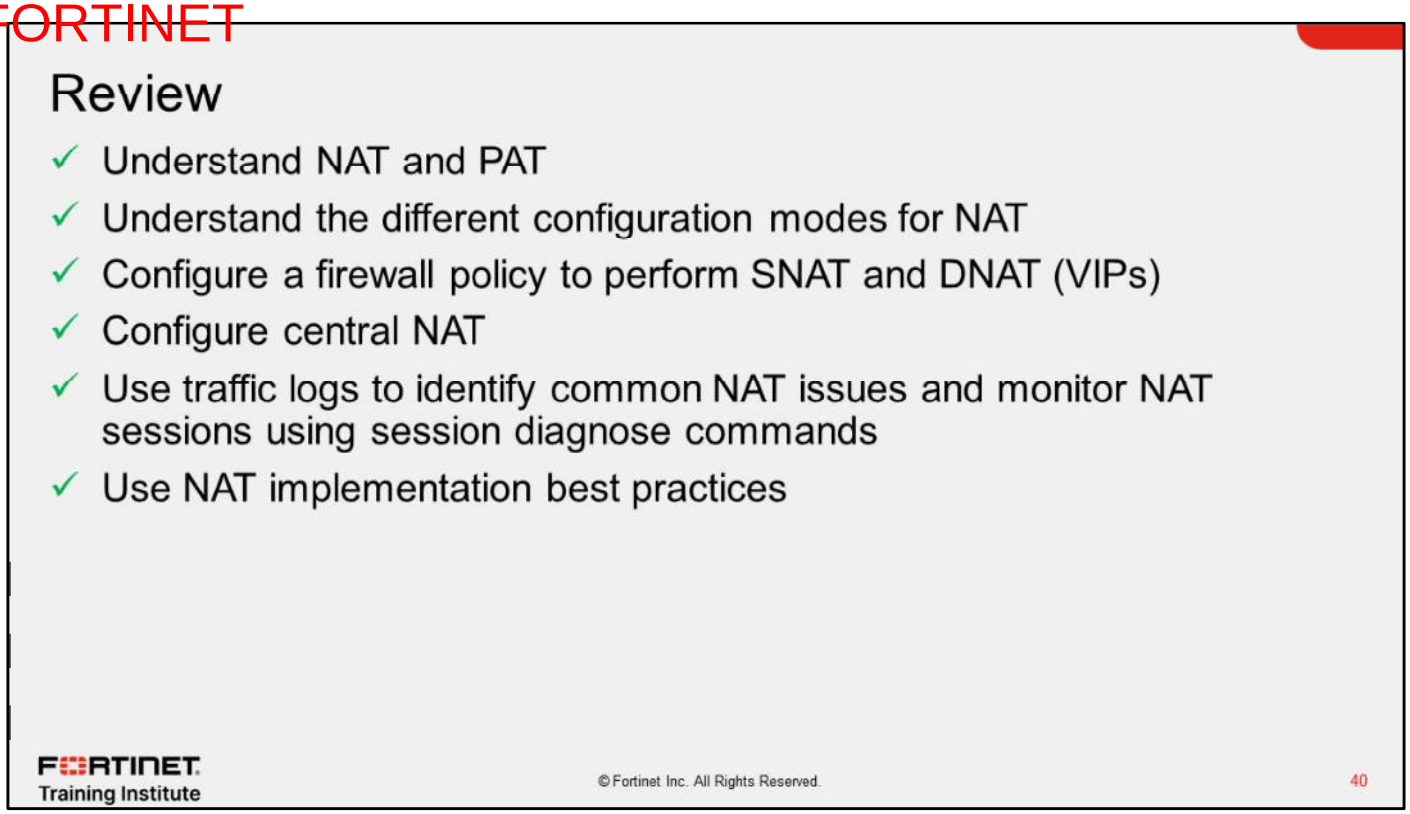

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to understand and configure NAT so that you can use it in your network.

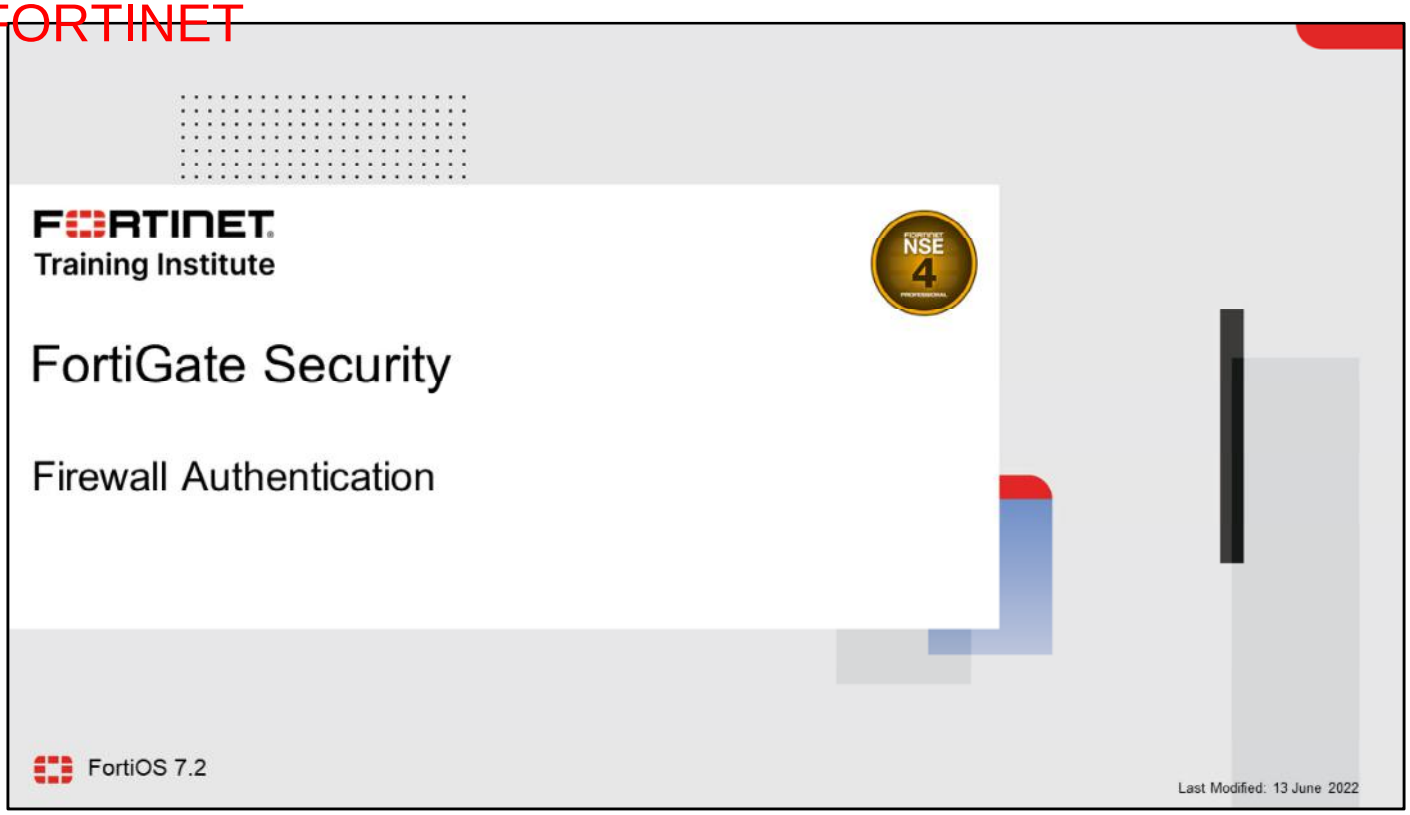

In this lesson, you will learn about using authentication on the firewall policies of FortiGate.

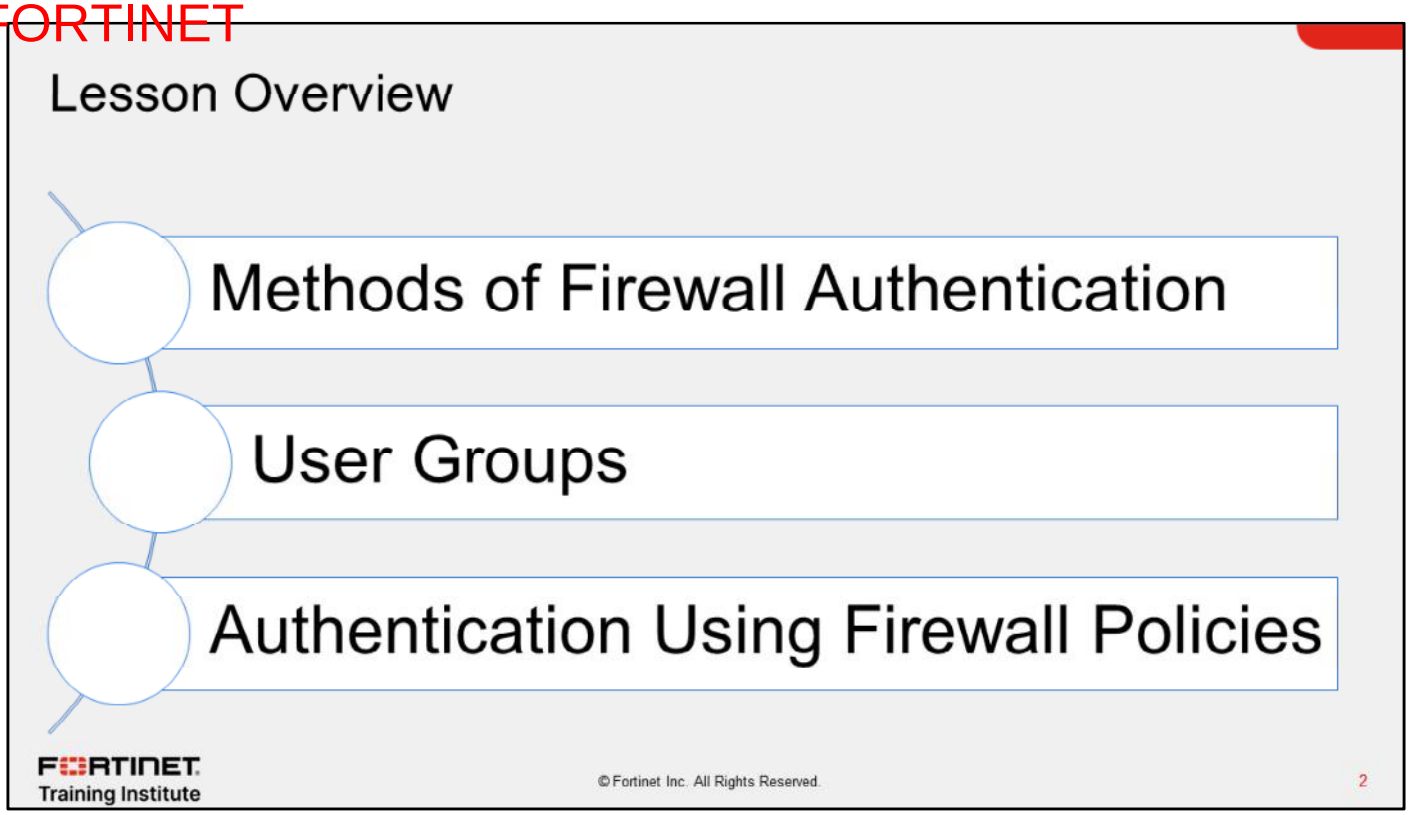

In this lesson, you will learn about the topics shown on this slide.

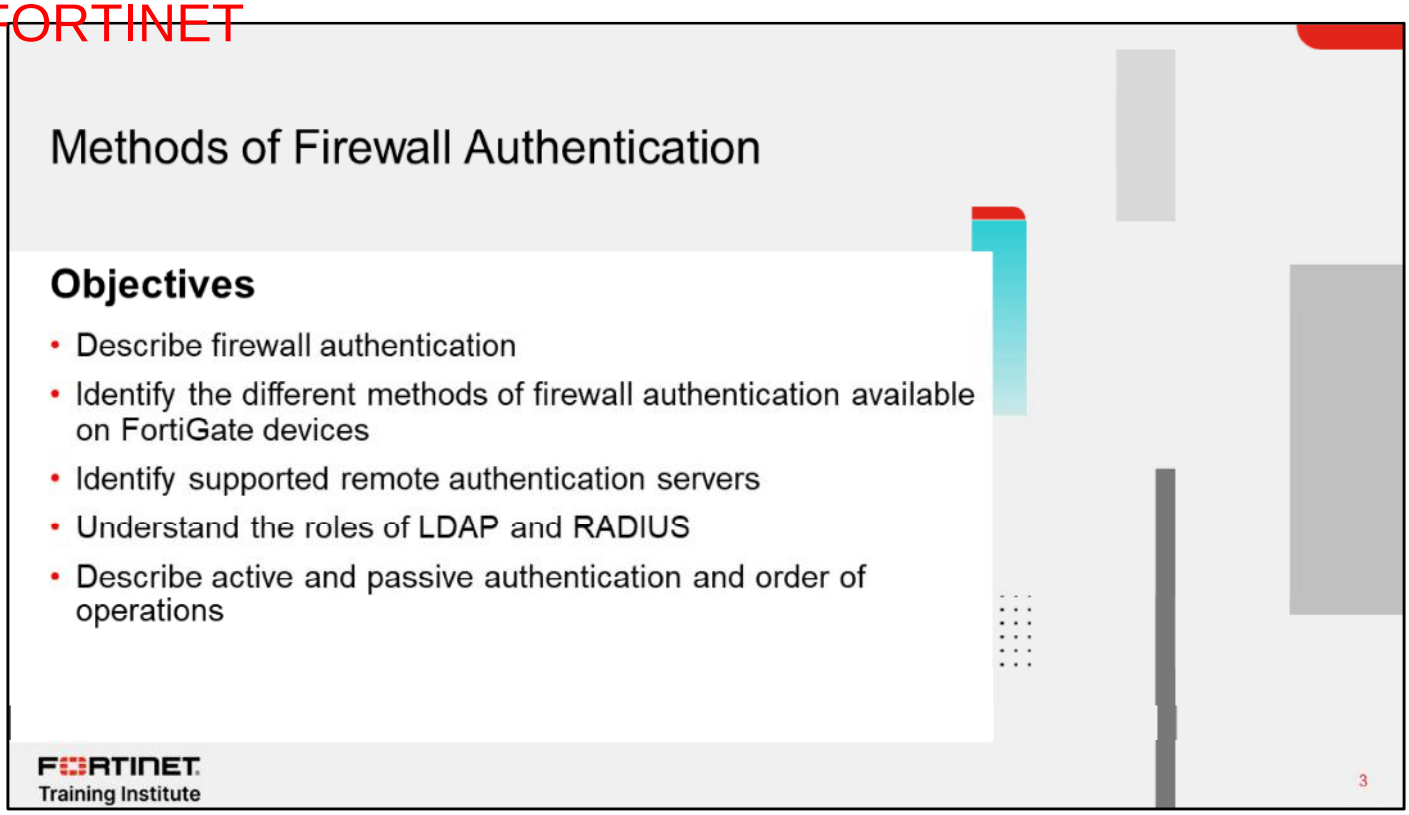

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in methods of firewall authentication, you will be able to describe and identify the supported methods of firewall authentication available on FortiGate.

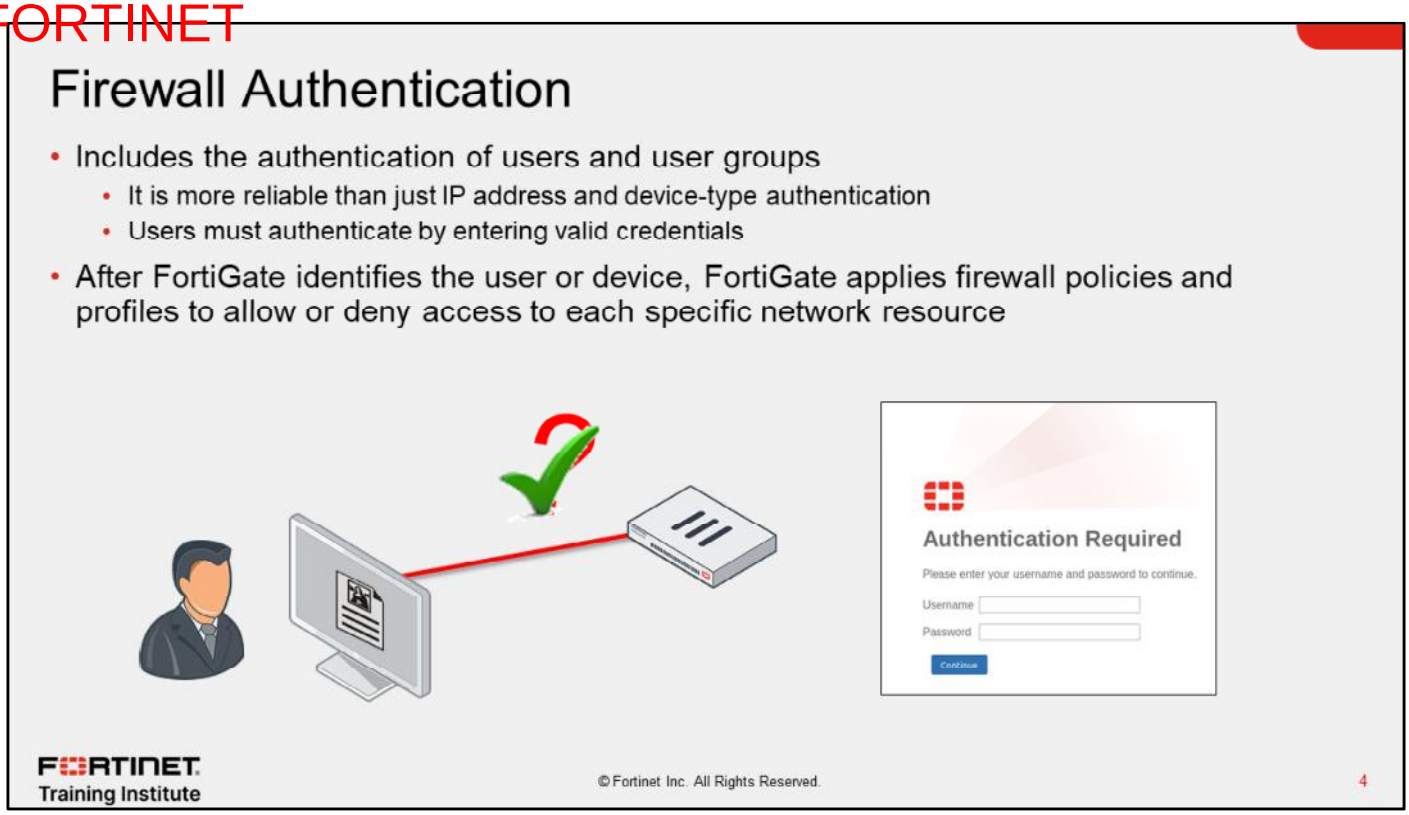

Traditional firewalling grants network access by verifying the source IP address and device. This is inadequate and can pose a security risk because the firewall cannot determine who is using the device to which it is granting access.

FortiGate includes authentication of users and user groups. As a result, you can follow individuals across multiple devices.

Where access is controlled by a user or user group, users must authenticate by entering valid credentials (such as username and password). After FortiGate validates the user, FortiGate applies firewall policies and profiles to allow or deny access to specific network resources.

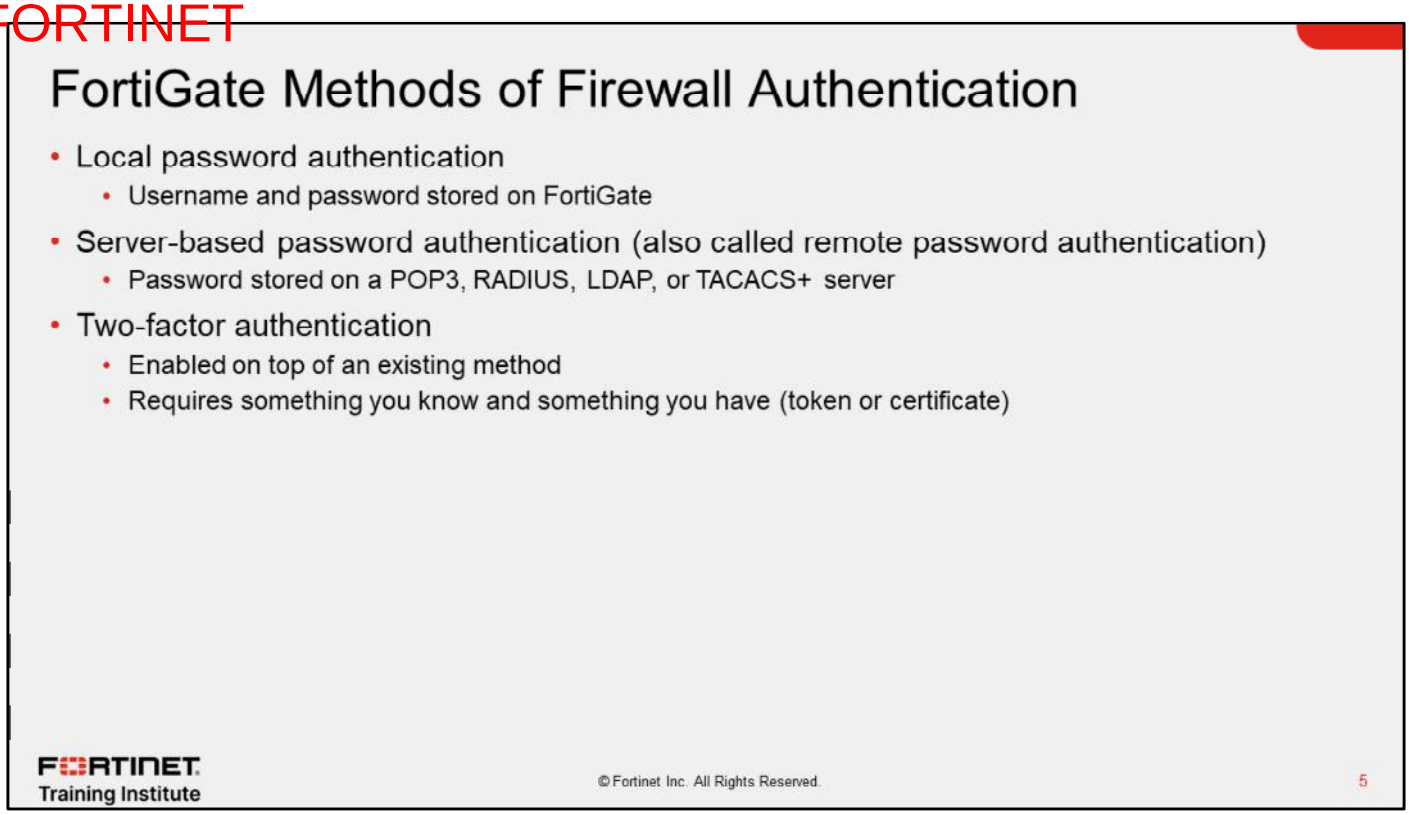

FortiGate supports multiple methods of firewall authentication:

- Local password authentication
- Server-based password authentication (also called remote password authentication)
- Two-factor authentication
  This is a system of authentication that is enabled on top of an existing method—it cannot be enabled without first configuring one of the other methods. It requires something you know, such as a password, and something you have, such as a token or certificate.

During this lesson, you will learn about each method of firewall authentication in detail.

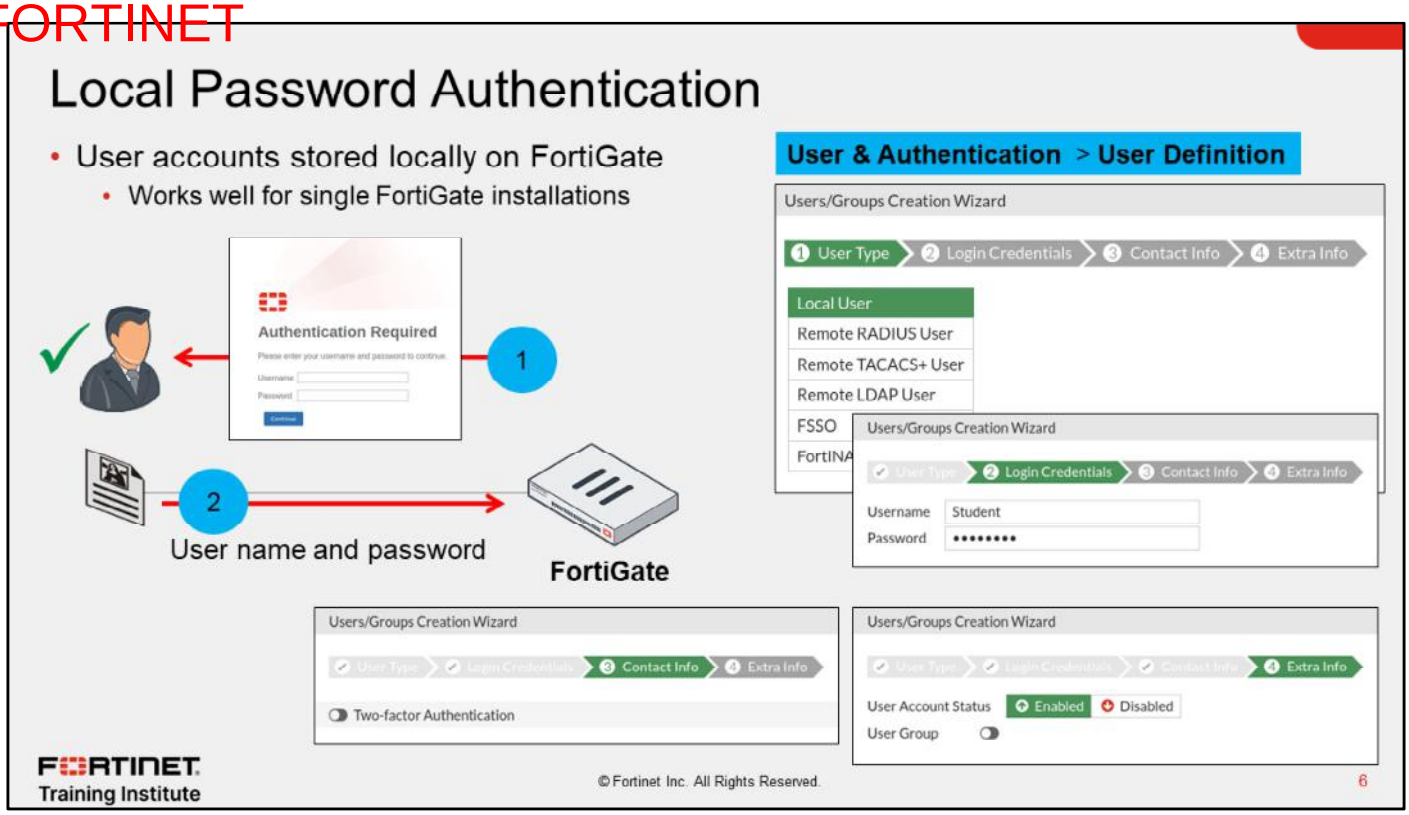

The simplest method of authentication is local password authentication. User account information (username and password) is stored locally on the FortiGate device. This method works well for a single FortiGate installation.

Local accounts are created on the **User Definition** page where a wizard takes you through the process. For local password authentication, select **Local User** as the user type and create a username and password. If desired, you can also add email and SMS information to the account, enable two-factor authentication, and add the user to a preconfigured user group.

After you create the user, you can add the user—or any preconfigured user group in which the user is a member—to a firewall policy, in order to authenticate. You will learn about user groups and firewall policies in this lesson.

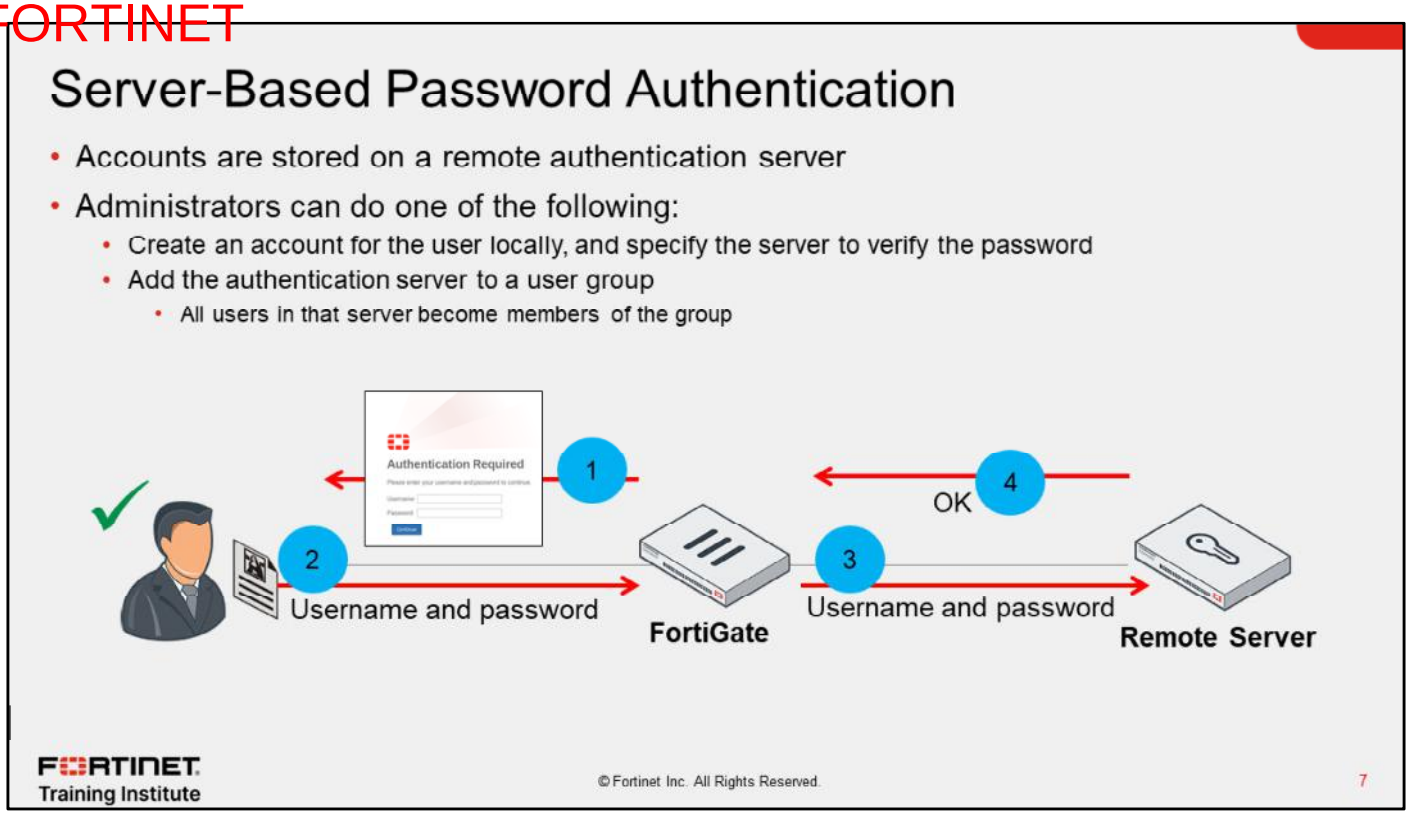

When server-based password authentication is used, a remote authentication server authenticates users. This method is desirable when multiple FortiGate devices need to authenticate the same users or user groups, or when adding FortiGate to a network that already contains an authentication server.

When you use a remote authentication server to authenticate users, FortiGate sends the user's entered credentials to the remote authentication server. The remote authentication server responds by indicating whether the credentials are valid or not. If valid, FortiGate consults its configuration to deal with the traffic. Note that it is the remote authentication server—not FortiGate—that evaluates the user credentials.

When the server-based password authentication method is used, FortiGate does not store all (or, in the case of some configurations, any) of the user information locally.

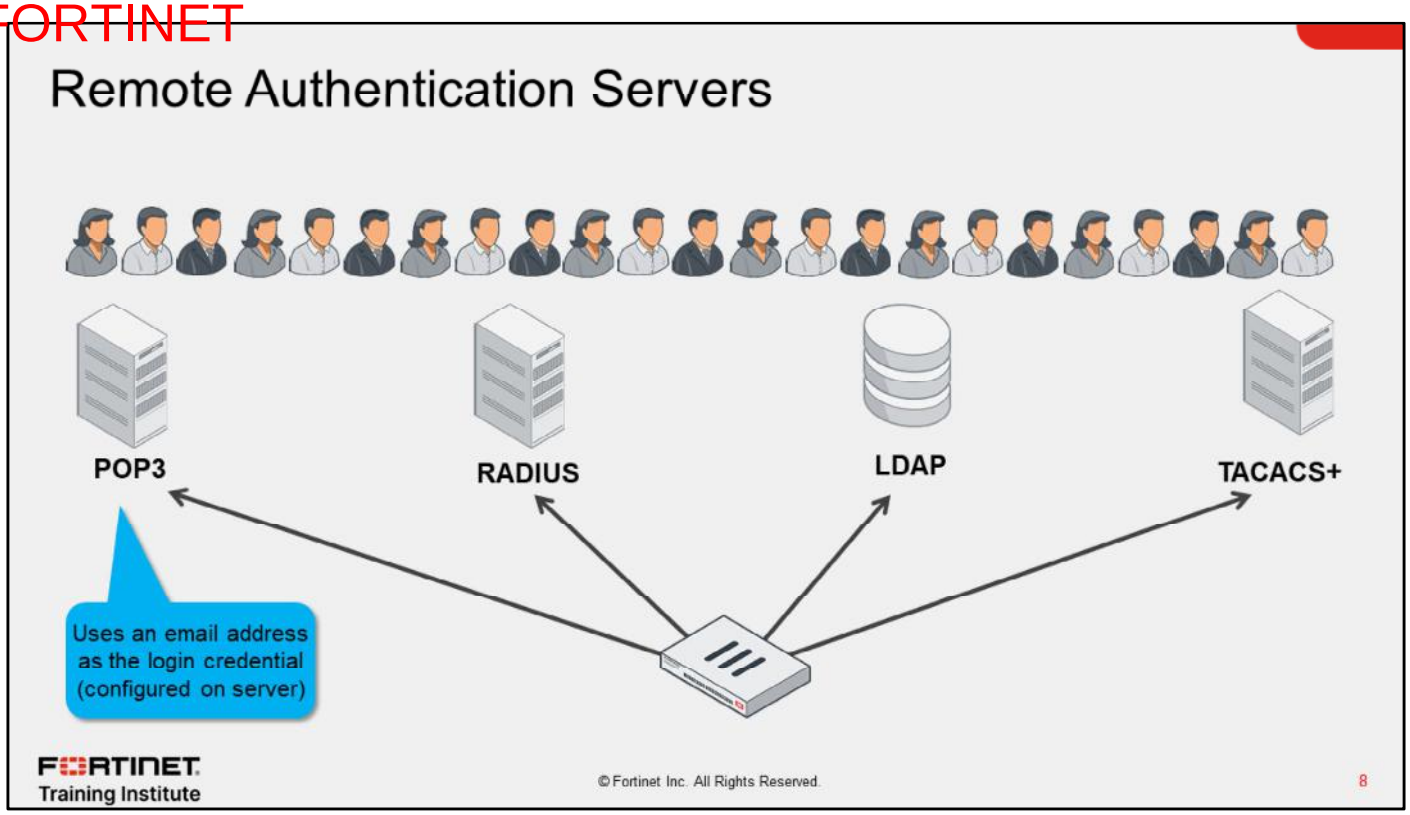

FortiGate provides support for many remote authentication servers, including POP3, RADIUS, LDAP, and TACACS+.

POP3 is the only server that requires an email address as the login credential. All other remote authentication servers use the user name. Some POP3 servers require the full email with domain (user@example.com), others require the suffix only, while still others accept both formats. This requirement is determined by the configuration of the server and is not a setting on FortiGate. You can configure POP3 authentication only though the CLI. Note that you can configure LDAP to validate with email, rather than the user name.

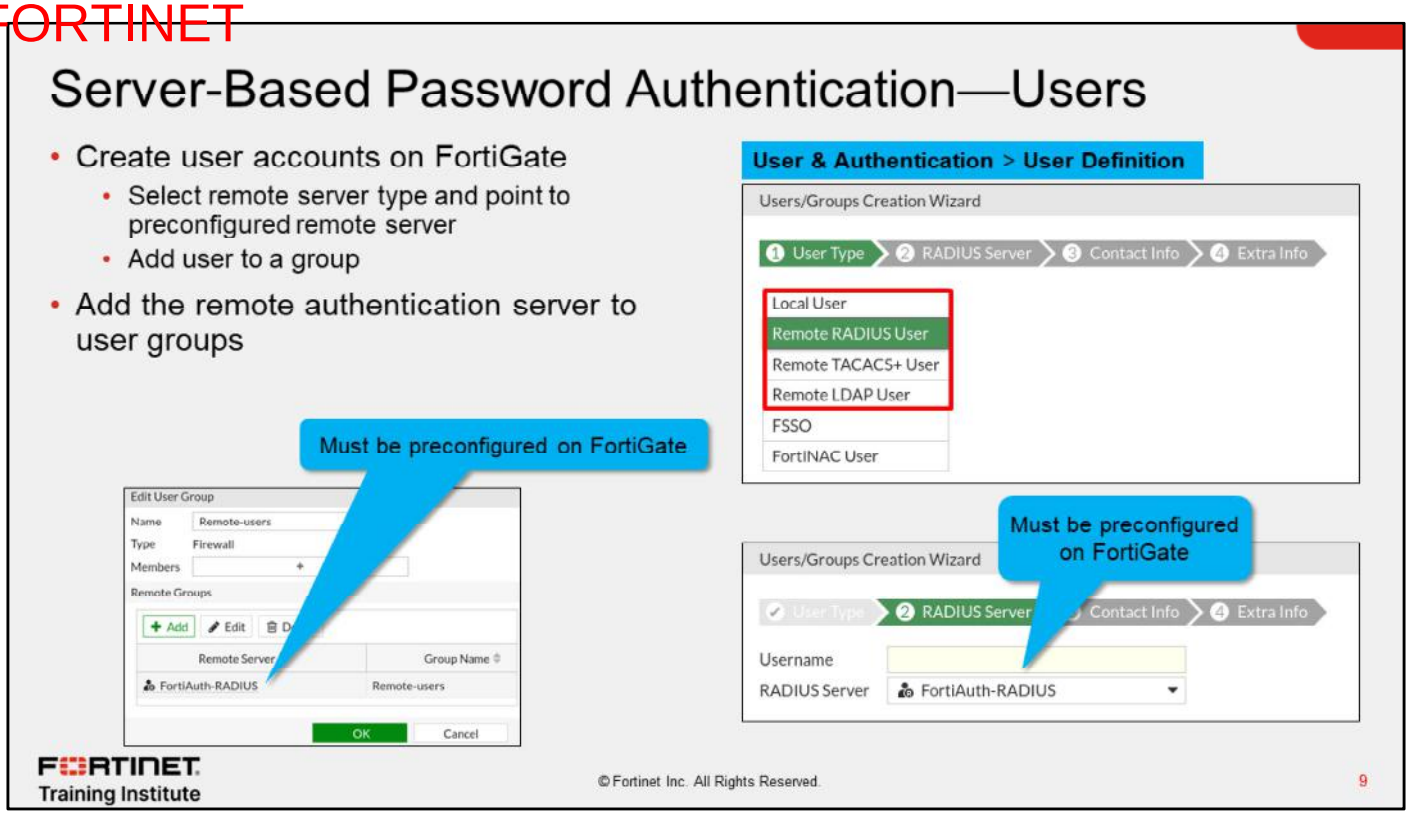

You can configure FortiGate to use external authentication servers in the following two ways:

- Create user accounts on FortiGate. With this method, you must select the remote authentication server type (RADIUS, TACACS+, or LDAP), point FortiGate to your preconfigured remote authentication server, and add the user to an appropriate group. This is usually done when you want to add two-factor authentication to your remote users. Remember, POP3 is only configurable through the CLI.
- Add the remote authentication server to user groups. With this method, you must create a user group and add the preconfigured remote server to the group. Accordingly, any user who has an account on the remote authentication server can authenticate. If you are using other types of remote servers, such as an LDAP server, as the remote authentication server, you can control access to specific LDAP groups, as defined on the LDAP server.

Similar to local password authentication, you must then add the preconfigured user group (in which the user is a member) to a firewall policy in order to authenticate. You will learn about user groups and firewall policies later in this lesson.

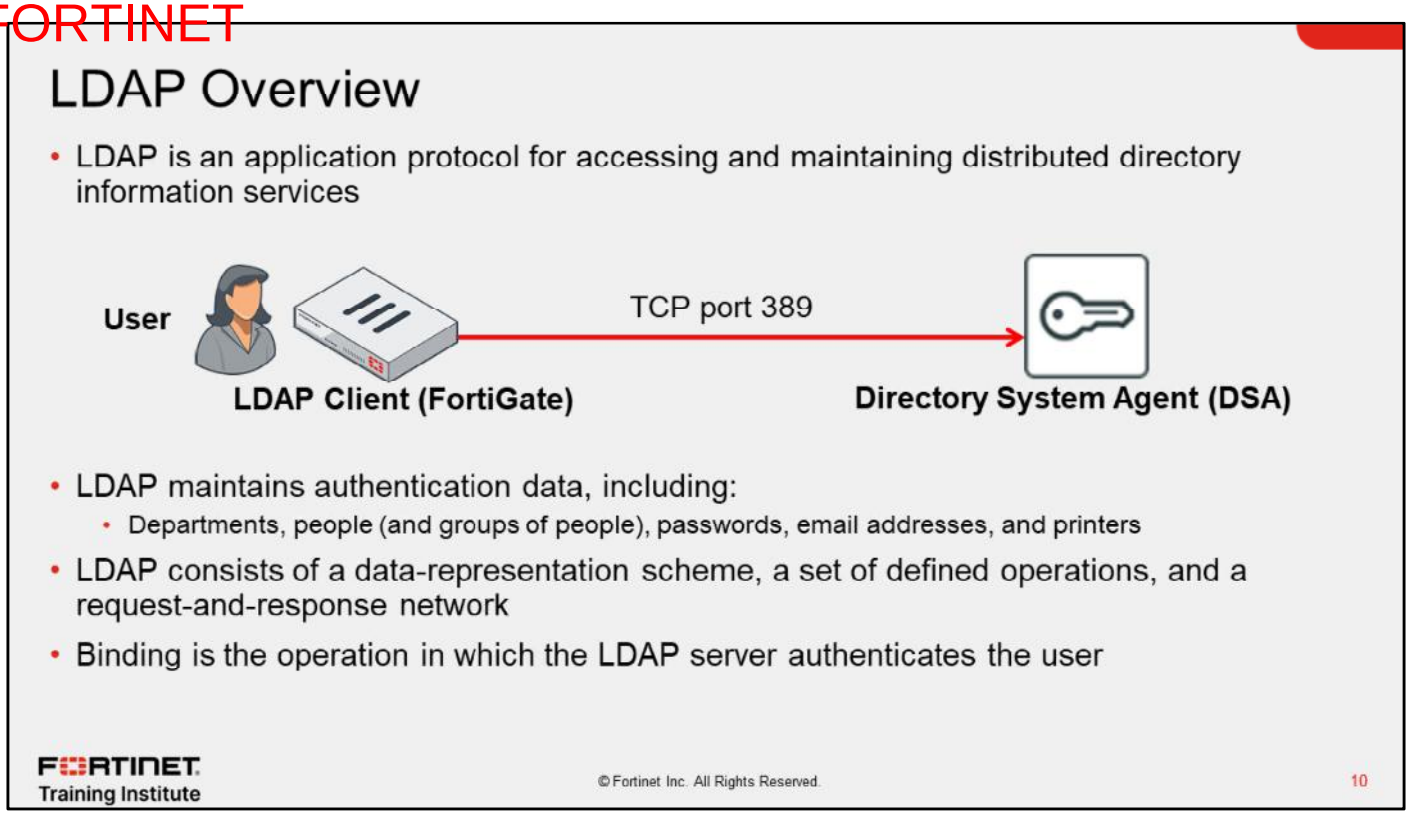

Lightweight Directory Access Protocol (LDAP) is an application protocol used for accessing and maintaining distributed directory information services.

The LDAP protocol is used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, and printers. LDAP consists of a data-representation scheme, a set of defined operations, and a request-and-response network.

The LDAP protocol includes a number of operations that a client can request, such as search, compare, and add or delete an entry. Binding is the operation in which the LDAP server authenticates the user. If the user is successfully authenticated, binding allows the user access to the LDAP server, based on that user's permissions.

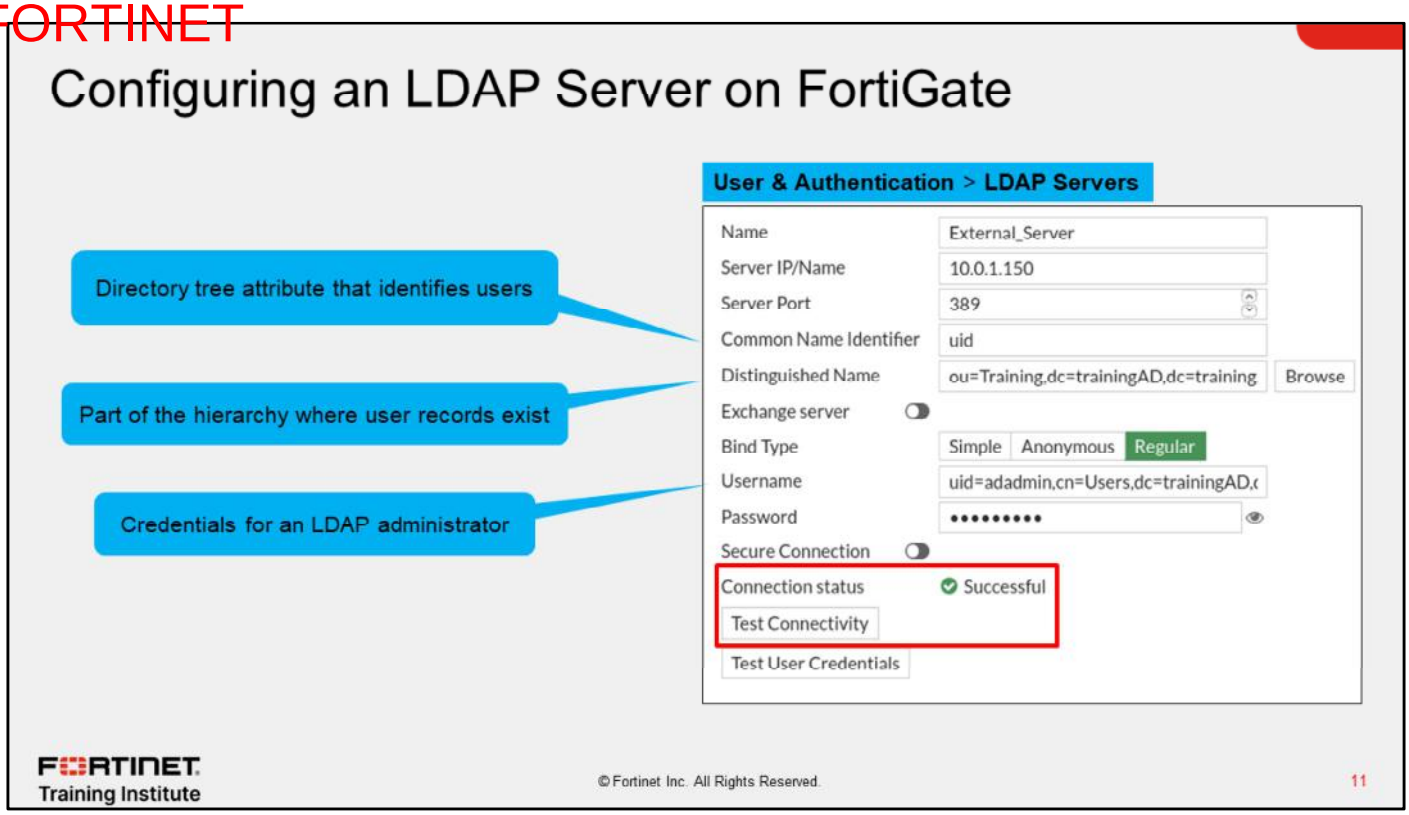

On the **LDAP Servers** page, you can configure FortiGate to point to an LDAP server for server-based password authentication. The configuration depends heavily on the server's schema and security settings. Windows Active Directory (AD) is very common.

The **Common Name Identifier** setting is the attribute name you use to find the user name. Some schemas allow you to use the attribute userid. AD most commonly uses `sAMAccountName` or `cn`, but can use others as well.

The **Distinguished Name** setting identifies the top of the tree where the users are located, which is generally the dc value; however, it can be a specific container or ou. You must use the correct X.500 or LDAP format.

The **Bind Type** setting depends on the security settings of the LDAP server. You must use the setting **Regular** (to specify a regular bind) if you are searching across multiple domains and require the credentials of a user that is authorized to perform LDAP queries (for example, an LDAP administrator).

If you want to have a secure connection between FortiGate and the remote LDAP server, enable **Secure Connection** and include the LDAP server protocol (LDAPS or STARTTLS) as well as the CA certificate that verifies the server certificate. LDAPS uses port 636 for communication.

The **Test Connectivity** button tests only whether the connection to the LDAP server is successful or not. To test whether a user's credentials can successfully authenticate, you can use the **Test User Credentials** button or use the CLI.

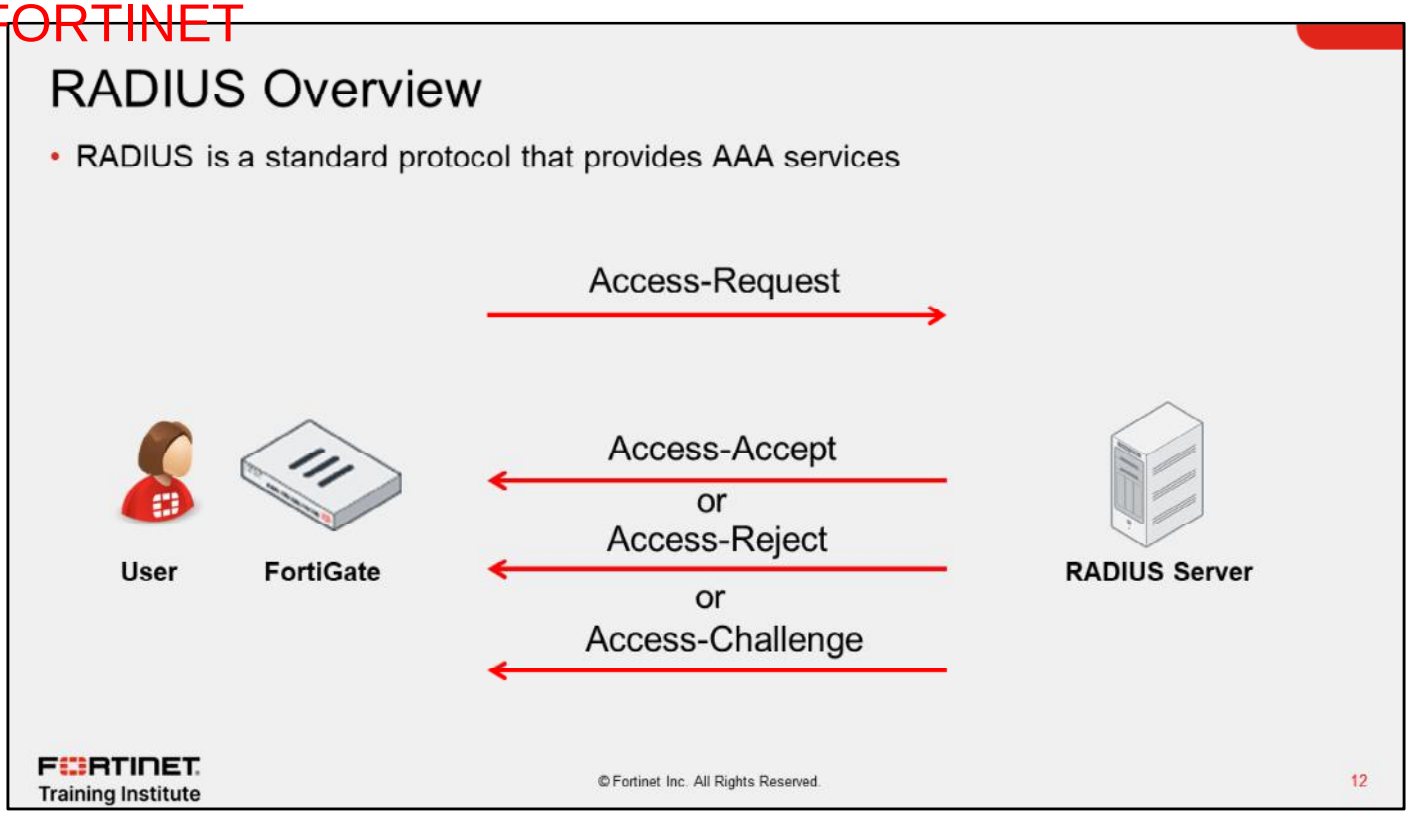

RADIUS is much different from LDAP, because there is no directory tree structure to consider. RADIUS is a standard protocol that provides authentication, authorization, and accounting (AAA) services.

When a user is authenticating, the client (FortiGate) sends an `ACCESS-REQUEST` packet to the RADIUS server. The reply from the server is one of the following:

- `ACCESS-ACCEPT`, which means that the user credentials are ok
- `ACCESS-REJECT`, which means that the credentials are wrong
- `ACCESS-CHALLENGE`, which means that the server is requesting a secondary password ID, token, or certificate. This is typically the reply from the server when using two-factor authentication.

Not all RADIUS clients support the RADIUS challenge method.

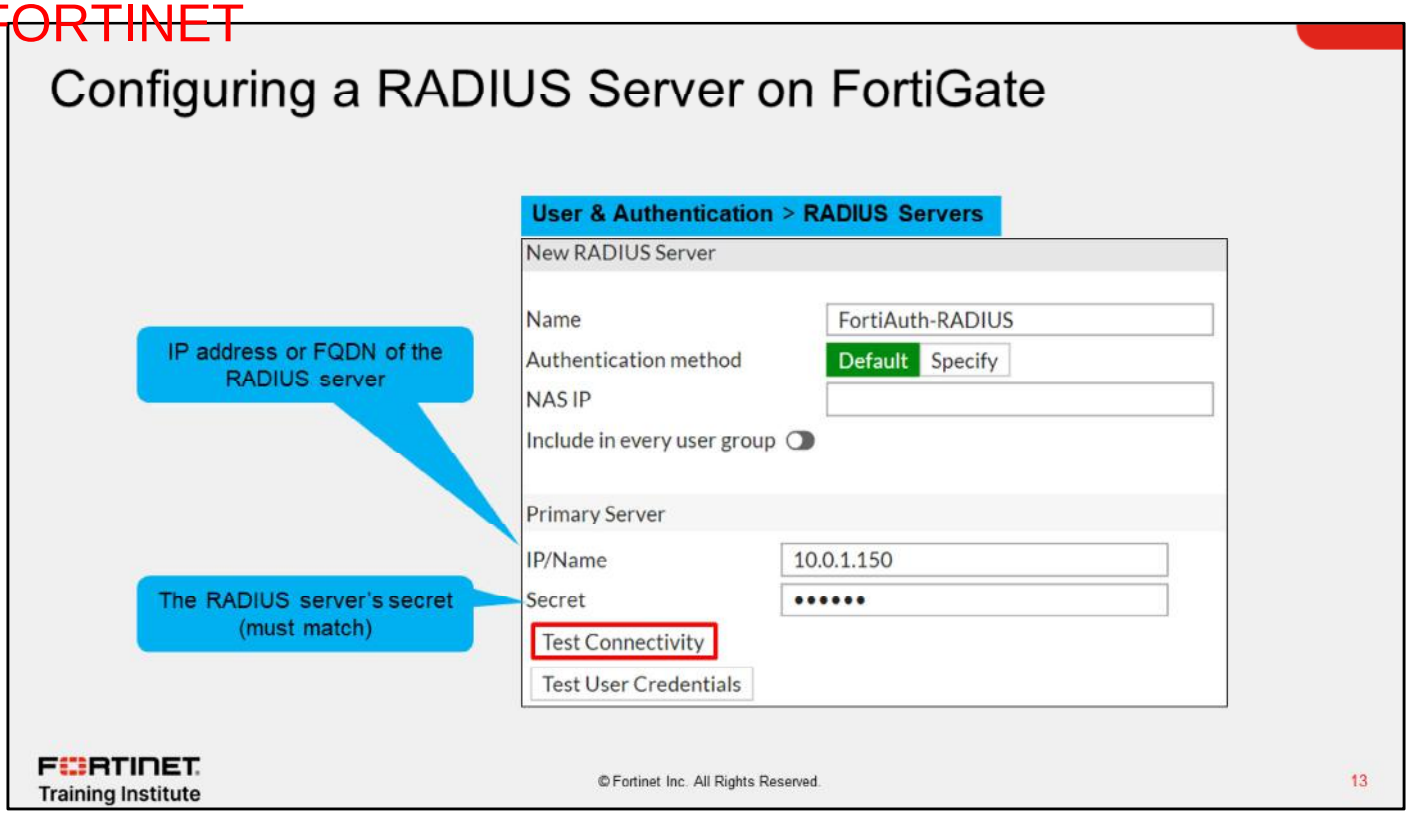

You can configure FortiGate to point to a RADIUS server for server-based password authentication through the **RADIUS Servers** page.

The **Primary Server IP/Name** setting is the IP address or FQDN of the RADIUS server.

The **Primary Server Secret** setting is the secret that was set up on the RADIUS server in order to allow remote queries from this client. Backup servers (with separate secrets) can be defined in case the primary server fails. Note that FortiGate must be listed on the RADIUS server as a client of that RADIUS server or else the server will not reply to queries done by FortiGate.

The **Authentication Method** setting refers to the authentication protocol that the RADIUS server supports. Options include chap, pap, mschap, and mschap2. If you select **Default**, FortiGate will use pap, mschap2, and chap (in that order).

Unlike LDAP configurations, the **Test Connectivity** button used in the example shown on this slide can test actual user credentials, but, like LDAP, you can also test this using the CLI.

The **Include in every User Group** option adds the RADIUS server and all users that can authenticate against it, to every user group created on FortiGate. So, you should enable this option only in very specific scenarios (for example, when only administrators can authenticate against the RADIUS server and policies are ordered from least restrictive to most restrictive).

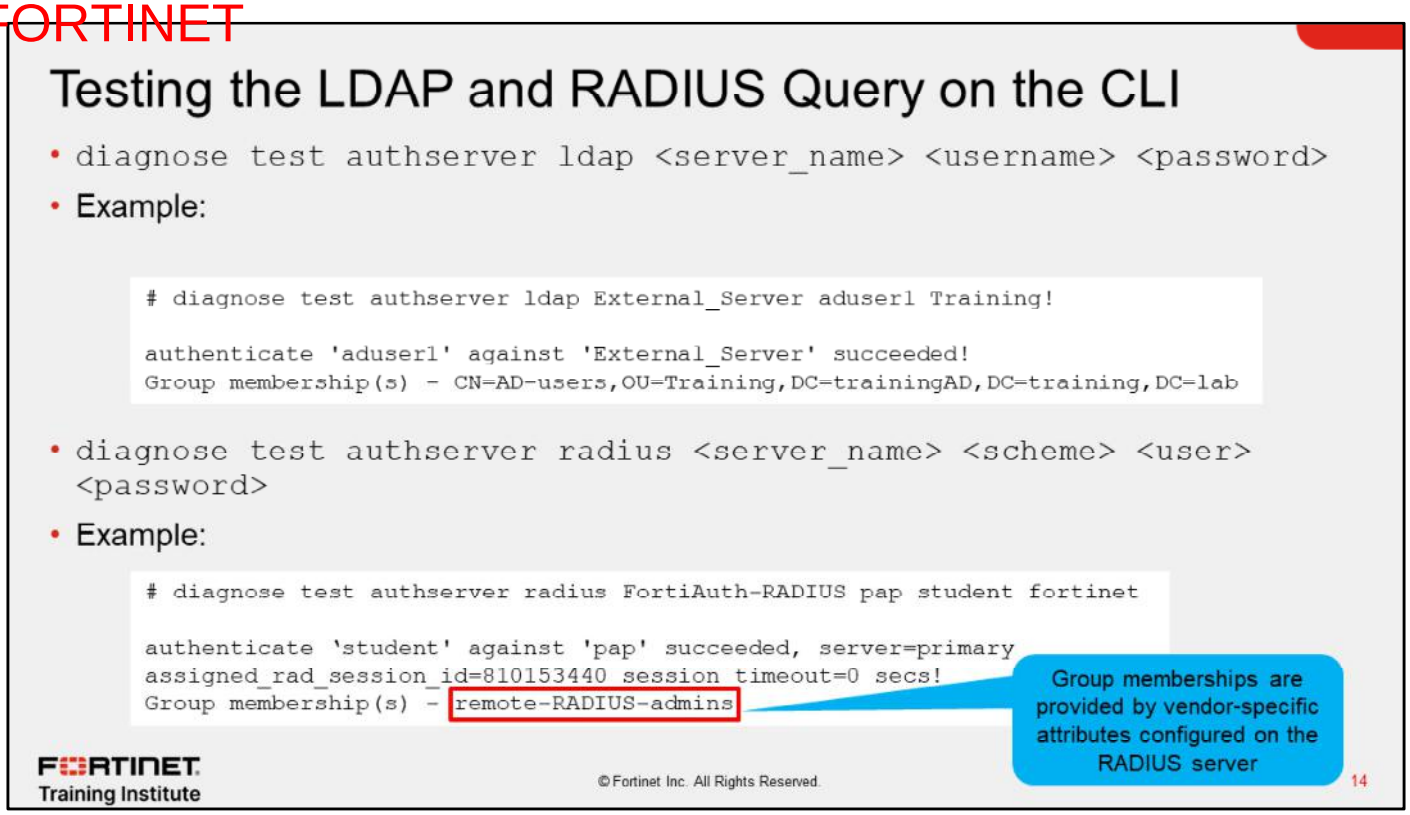

Use the `diagnose test authserver` command on the CLI to test whether a user's credentials can successfully authenticate. You want to ensure that authentication is successful, before implementing it on any of your firewall policies.

The response from the server reports success, failure, and group membership details.

Testing RADIUS is much the same as testing LDAP. Use the `diagnose test authserver` command on the CLI to test whether a user's credentials can successfully authenticate. Again, you should do this to ensure authentication is successful before implementing it on any of your firewall policies.

Like LDAP, it reports success, failure, and group membership details, depending on the server's response. Deeper troubleshooting usually requires RADIUS server access.

Note that Fortinet has a vendor-specific attributes (VSA) dictionary to identify the Fortinet-proprietary RADIUS attributes. This capability allows you to extend the basic functionality of RADIUS. You can obtain the Fortinet VSA dictionary from the Fortinet Knowledge Base (`kb.fortinet.com`).

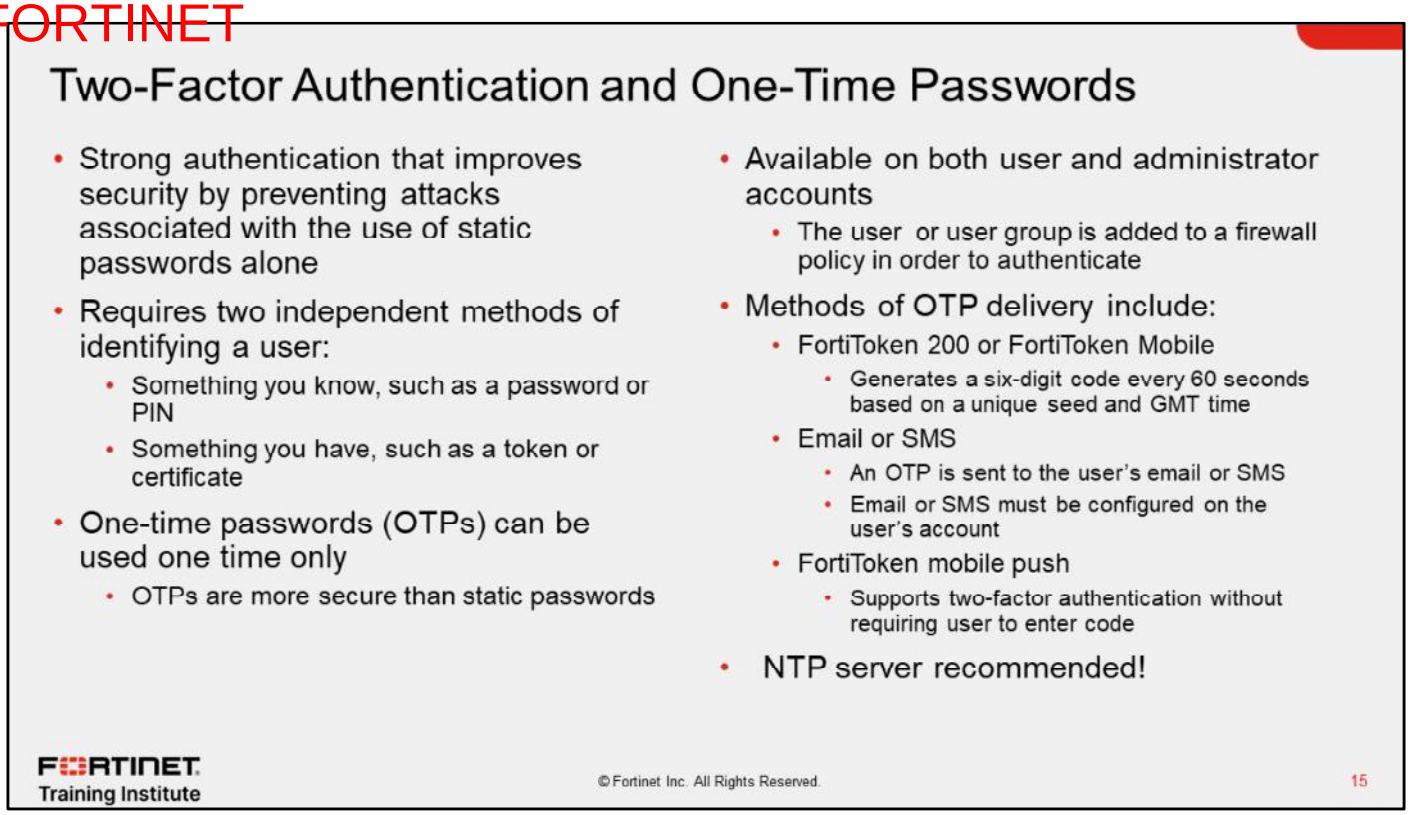

Traditional user authentication requires your user name plus something you know, such as a password. The weakness in this traditional method of authentication is that if someone obtains your username, they need only your password to compromise your account. Furthermore, since people tend to use the same password across multiple accounts (some sites with more security vulnerabilities than others), accounts are vulnerable to attack, regardless of password strength.

Two-factor authentication, on the other hand, requires something you know, such as a password, and something you have, such as a token or certificate. Because this method places less importance on often vulnerable passwords, it makes compromising the account more complex for an attacker. You can use two-factor authentication on FortiGate with both user and administrator accounts. The user (or user group to which the user belongs) is added to a firewall policy in order to authenticate. Note that you cannot use two-factor authentication with explicit proxies.

You can use one-time passwords (OTPs) as your second factor. OTPs are more secure than static passwords because the passcode changes at regular intervals and is valid for only a short amount of time. Once you use the OTP, you can't use it again. So, even if it is intercepted, it is useless. FortiGate can deliver OTPs through tokens, such as FortiToken 200 (hardware token) and FortiToken Mobile (software token), as well as through email or SMS. To deliver an OTP over email or SMS, the user account must contain user contact information.

FortiTokens and OTPs delivered through email and SMS are time based. FortiTokens, for example, generate a new, six-digit password every 60 seconds (by default). An NTP server is highly recommended to ensure the OTPs remain in sync. FortiToken Mobile Push allows users to accept the authorization request from their FortiToken mobile app, without the need to enter an additional code.

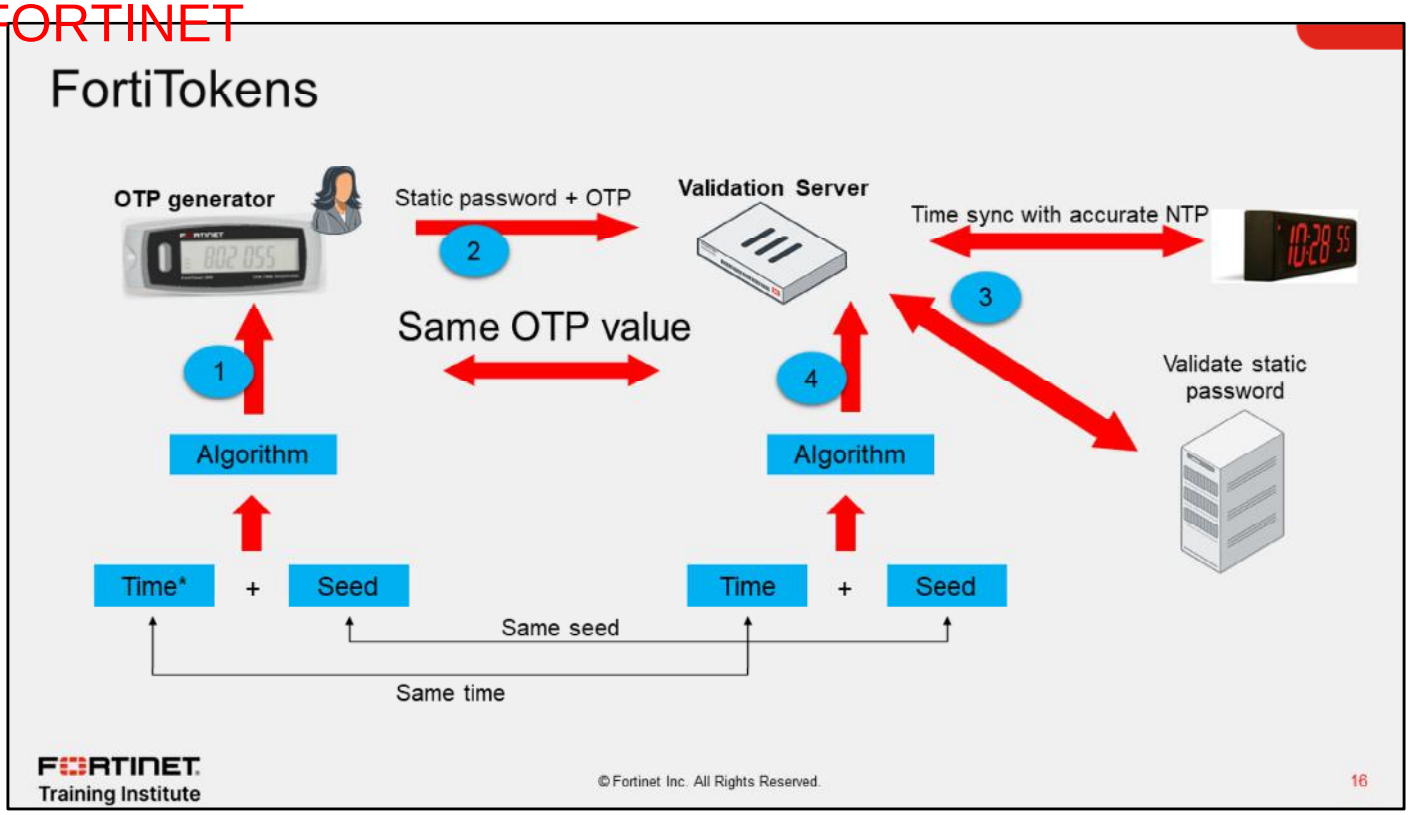

Tokens use a specific algorithm to generate an OTP. The algorithm consists of:

- A seed: a unique, randomly-generated number that does not change over time
- The time: obtained from an accurate internal clock

Both seed and time go through an algorithm that generates an OTP (or passcode) on the token. The passcode has a short life span, usually measured in seconds (60 seconds for FortiToken 200, possibly more or less for other RSA key generators). Once the life span ends, a new passcode generates.

When using two-factor authentication using a token, the user must first log in with a static password followed by the passcode generated by the token. A validation server (FortiGate) receives the user's credentials and validates the static password first. The validation server then proceeds to validate the passcode. It does so by regenerating the same passcode using the seed and system time (which is synchronized with the one on the token) and comparing it with the one received from the user. If the static password is valid, and the OTP matches, the user is successfully authenticated. Again, both the token and the validation server must use the same seed and have synchronized system clocks. As such, it is crucial that you configure the date and time correctly on FortiGate, or link it to an NTP server (which is recommended).

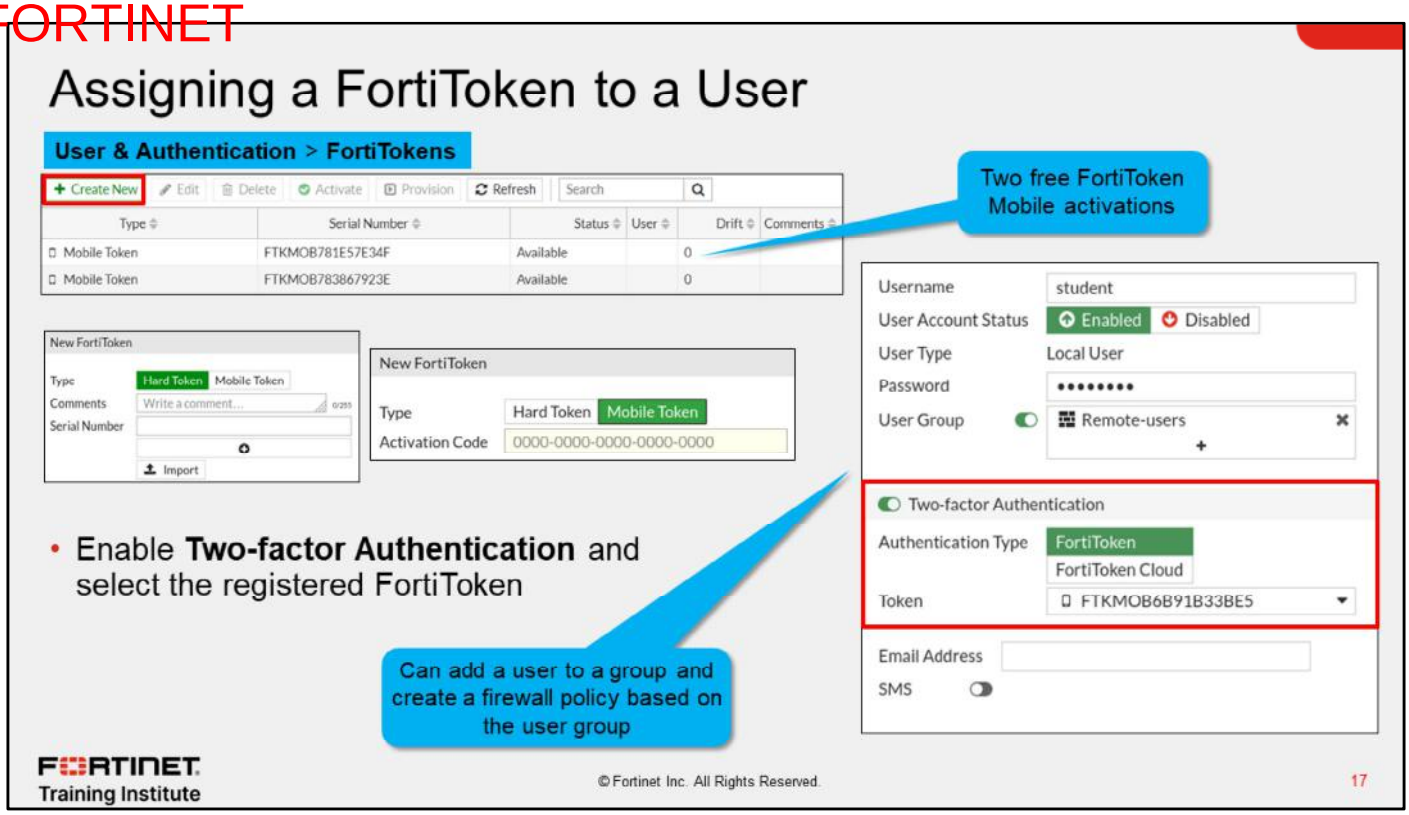

You can add a FortiToken 200 or FortiToken Mobile to FortiGate on the **FortiTokens** page.

A hard token has a serial number that provides FortiGate with details on the initial seed value. If you have several hard tokens to add, you can import a text file, where one serial number is listed per line.

A soft token requires an activation code. Note that each FortiGate (and FortiGate VM) provides two free FortiToken Mobile activations. You must purchase any additional tokens from Fortinet.

You cannot register the same FortiToken on more than one FortiGate. If you want to use the same FortiToken for authentication on multiple FortiGate devices, you must use a central validation server, such as FortiAuthenticator. In that case, FortiTokens are registered and assigned to users on FortiAuthenticator, and FortiGate uses FortiAuthenticator as its validation server.

After you have registered the FortiToken devices with FortiGate, you can assign them to users to use as their second-factor authentication method. To assign a token, edit (or create) the user account and select **Enable Two-factor Authentication**. On the **Token** drop-down list, select the registered token you want to assign.

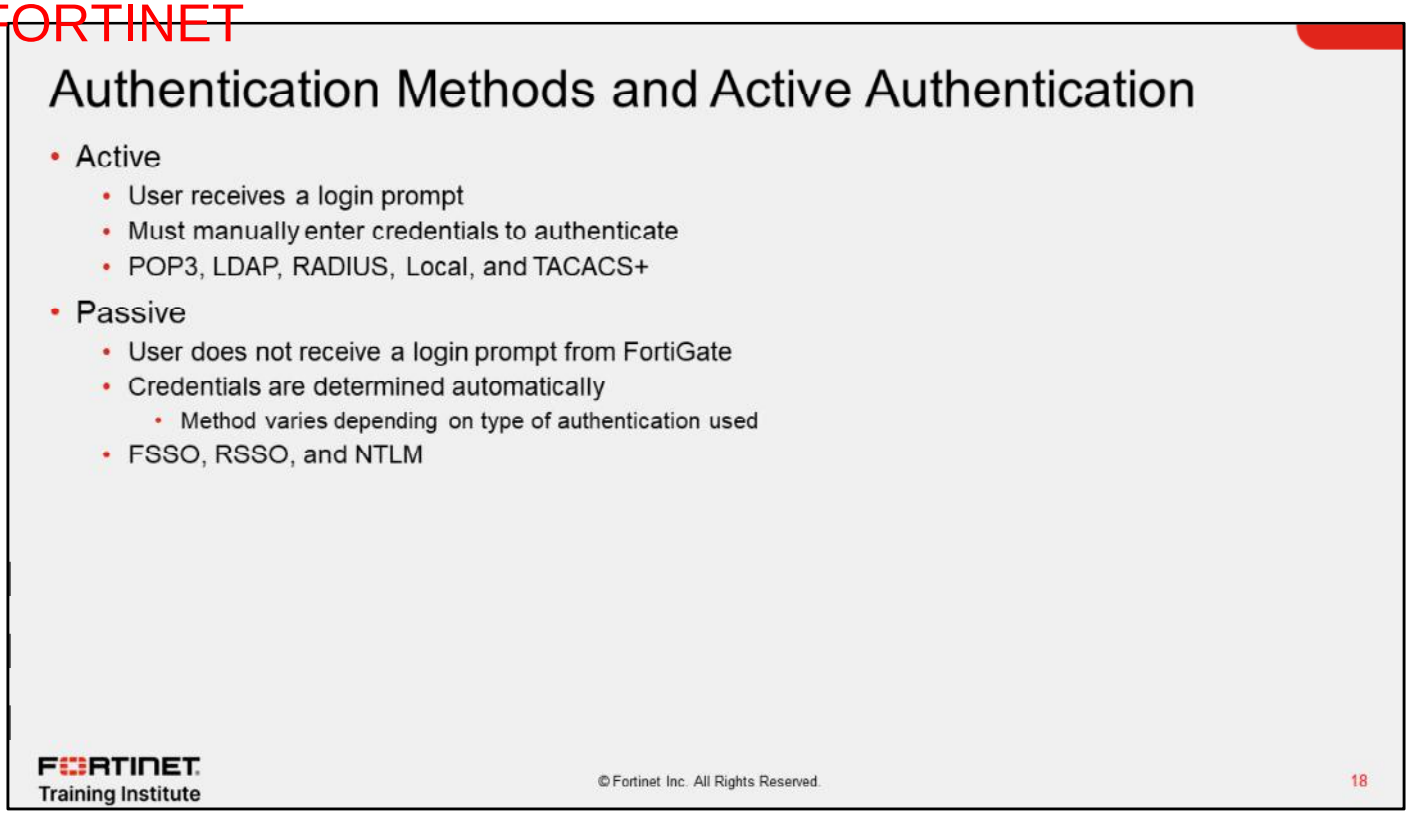

All the authentication methods you've learned about—local password authentication, server-based authentication, and two-factor authentication—use active authentication. Active authentication means that users are prompted to manually enter their login credentials before being granted access.

But not all users authenticate the same way. Some users can be granted access transparently, because user information is determined without asking the user to enter their login credentials. This is known as passive authentication. Passive authentication occurs with the single sign-on method for server-based password authentication: FSSO, RSSO, and NTLM.

Knowledge Check

1. Which firewall authentication method does FortiGate support?
   - ✓ A. Local password authentication
   - B. Biometric authentication

2. A remote LDAP user is trying to authenticate with a username and password. How does FortiGate verify the login credentials?
   - A. FortiGate queries its own database for user credentials.
   - ✓ B. FortiGate sends the user-entered credentials to the remote server for verification.

3. When FortiGate uses a RADIUS server for remote authentication, which statement about RADIUS is true?
   - A. FortiGate must query the remote RADIUS server using the distinguished name (dn).
   - ✓ B. RADIUS group memberships are provided by vendor-specific attributes (VSAs) configured on the RADIUS server

FORTINET Training Institute

19

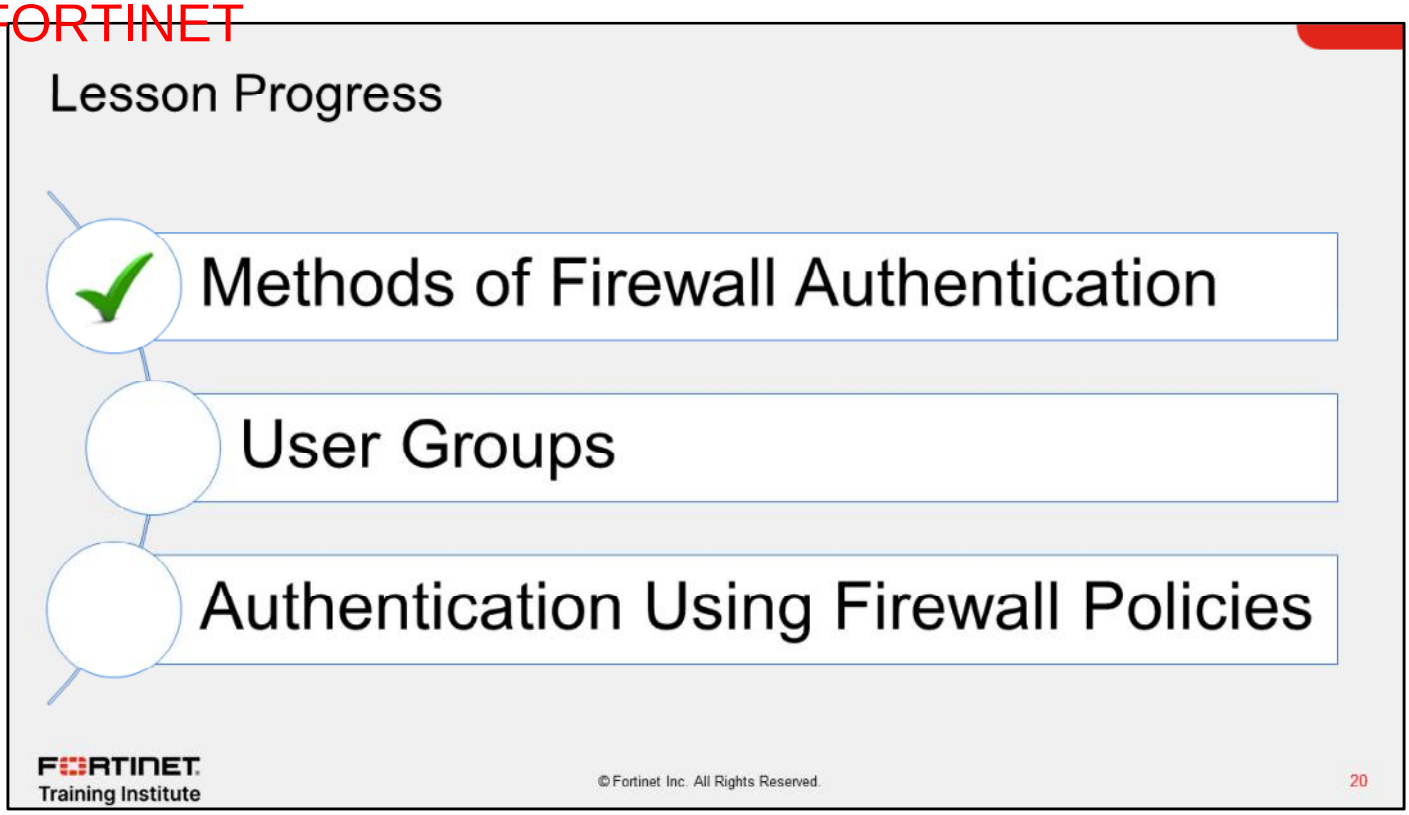

Good job! You now understand the basics of firewall authentication.

Now, you will learn about user groups.

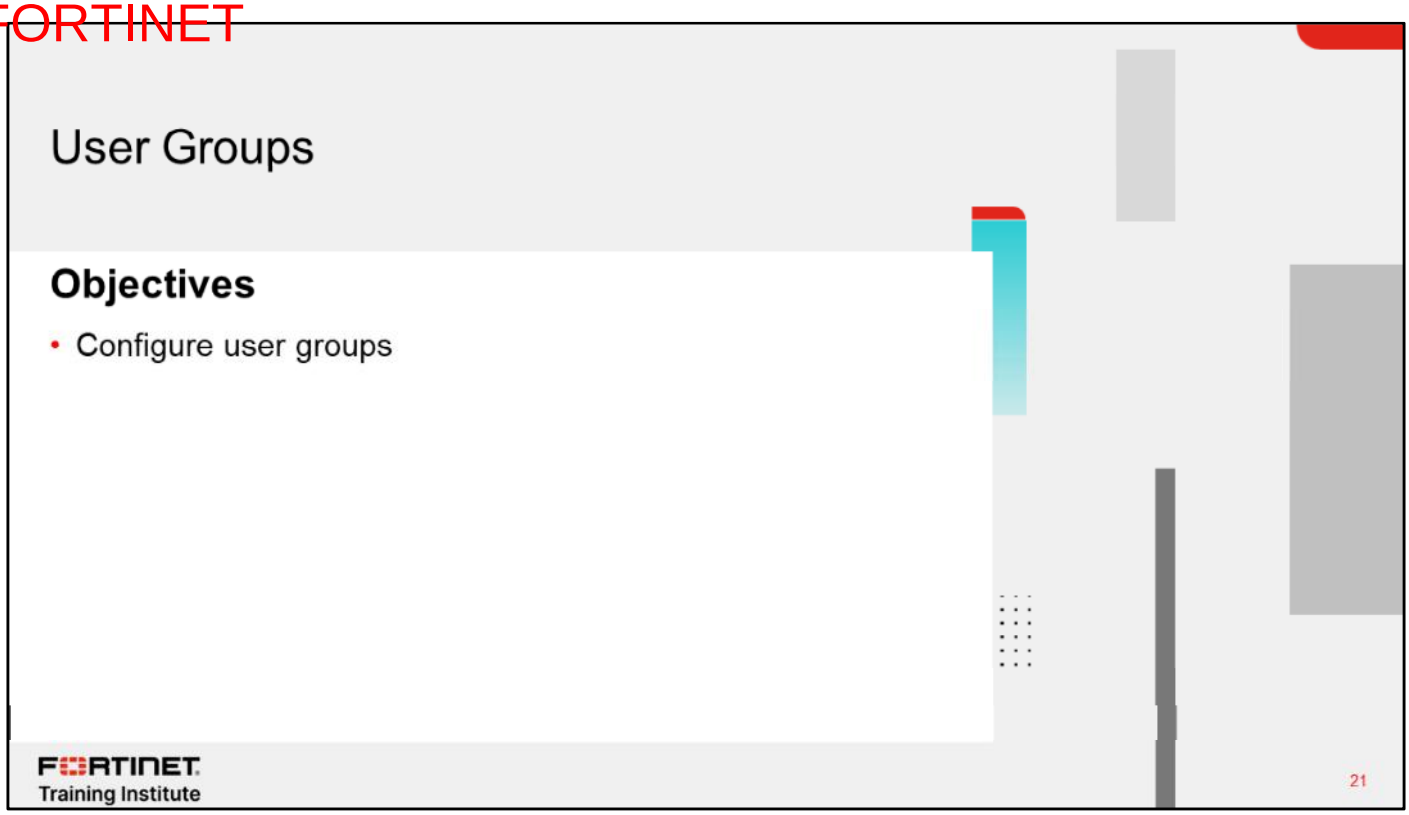

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in user groups, you will be able to configure user groups to efficiently manage firewall policies.

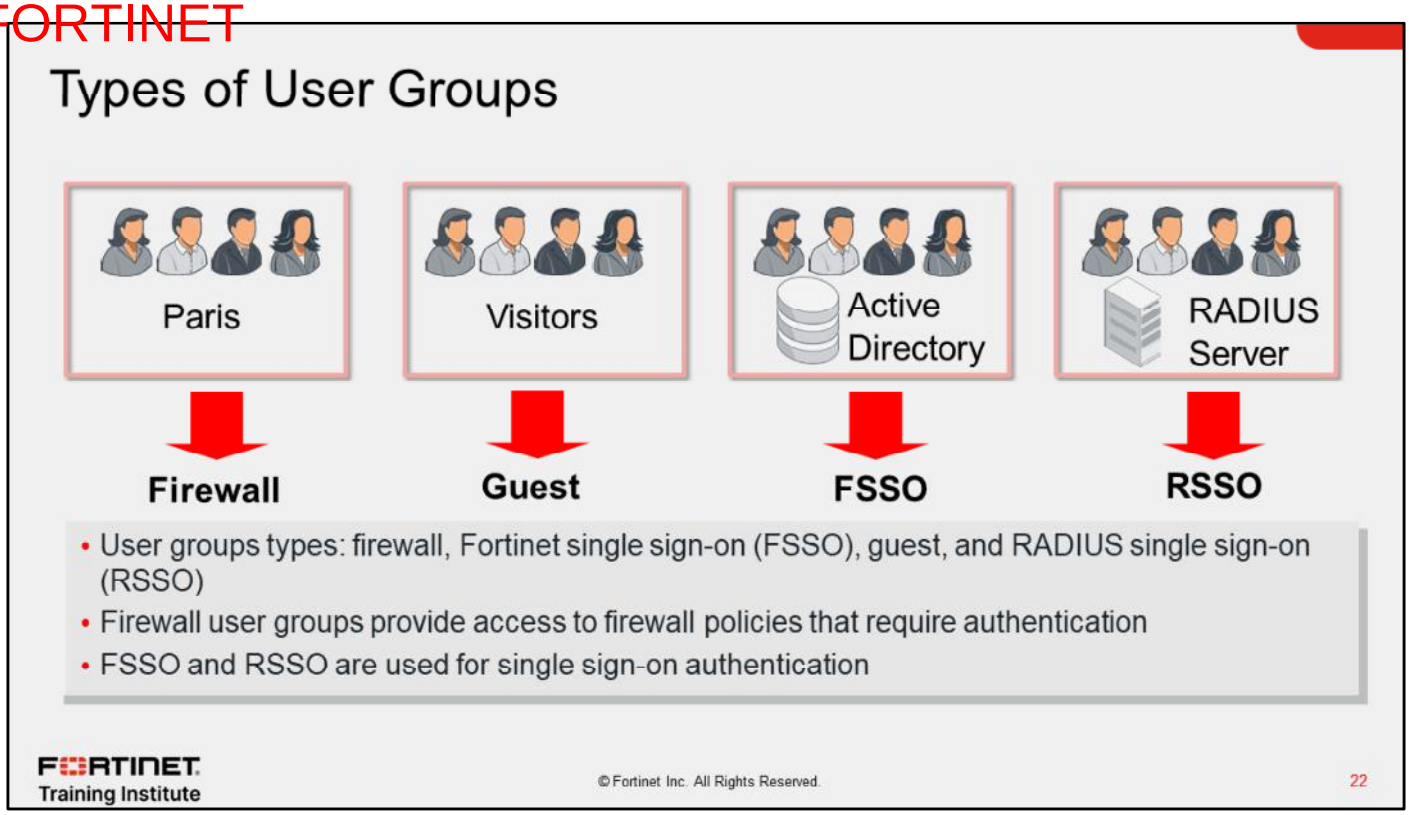

FortiGate allows administrators to assign users to groups. Usually, groups are used to more effectively manage individuals that have some kind of shared relationship. You might want to group employees by business area, such as finance or HR, or by employee type, such as contractors or guests.

After you create user groups, you can add them to firewall policies. This allows you to control access to network resources because policy decisions are made on the group as a whole. You can define both local and remote user groups on a FortiGate device. There are four user group types:

- Firewall
- Guest
- Fortinet single sign-on (FSSO)
- RADIUS single sign-on (RSSO)

The firewall user groups on FortiGate do not need to match any type of group that may already exist on an external server, such as an LDAP server. The firewall user groups exist solely to make configuration of firewall policies easier.

Most authentication types have the option to make decisions based on the individual user, rather than just user groups.

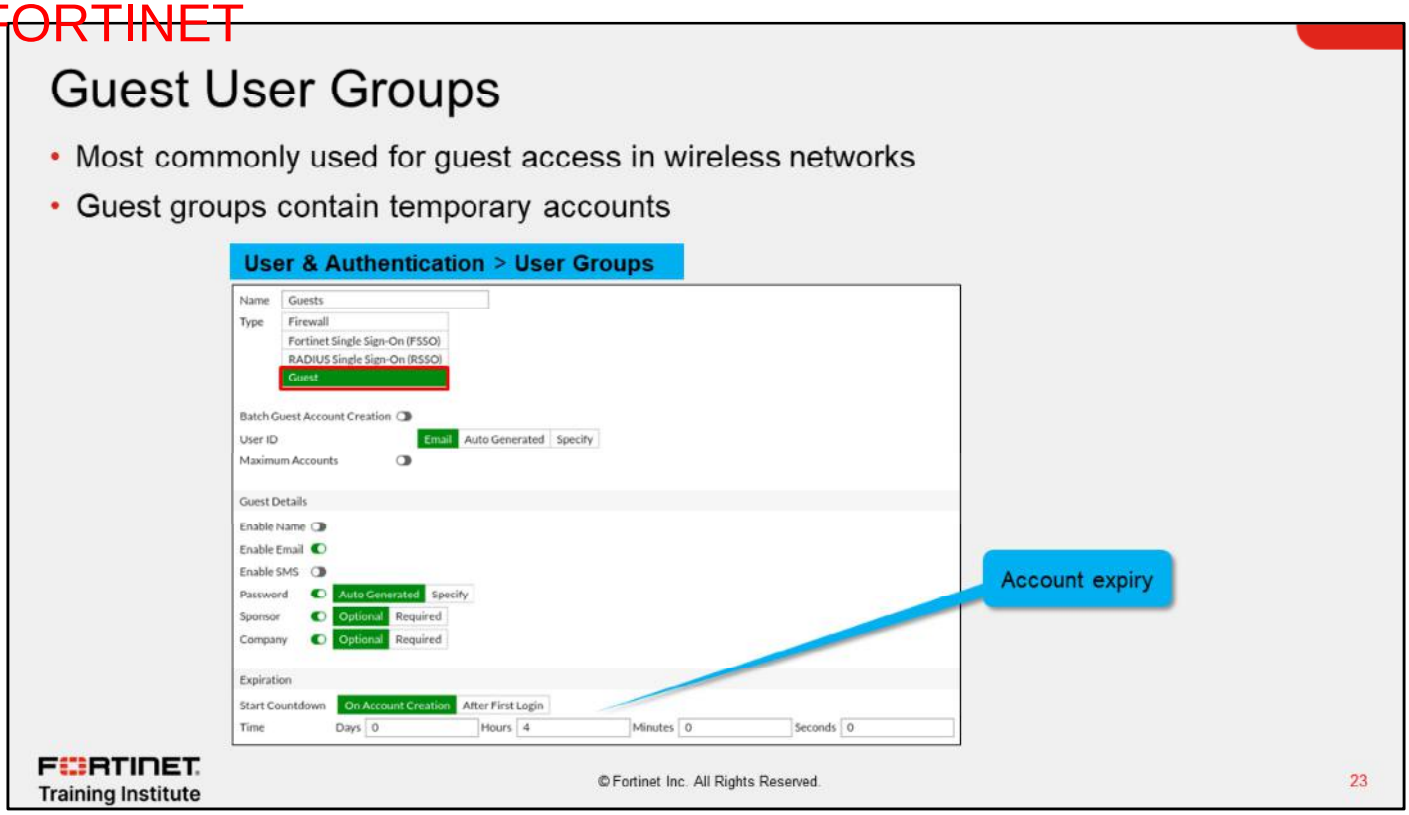

Guest user groups are different from firewall user groups because they contain exclusively temporary guest user accounts (the whole account, not just the password). Guest user groups are most commonly used in wireless networks. Guest accounts expire after a predetermined amount of time.

Administrators can manually create guest accounts or create many guest accounts at once using randomly generated user IDs and passwords. This reduces administrator workload for large events. Once created, you can add accounts to the guest user group and associate the group with a firewall policy.

You can create guest management administrators who have access only to create and manage guest user accounts.

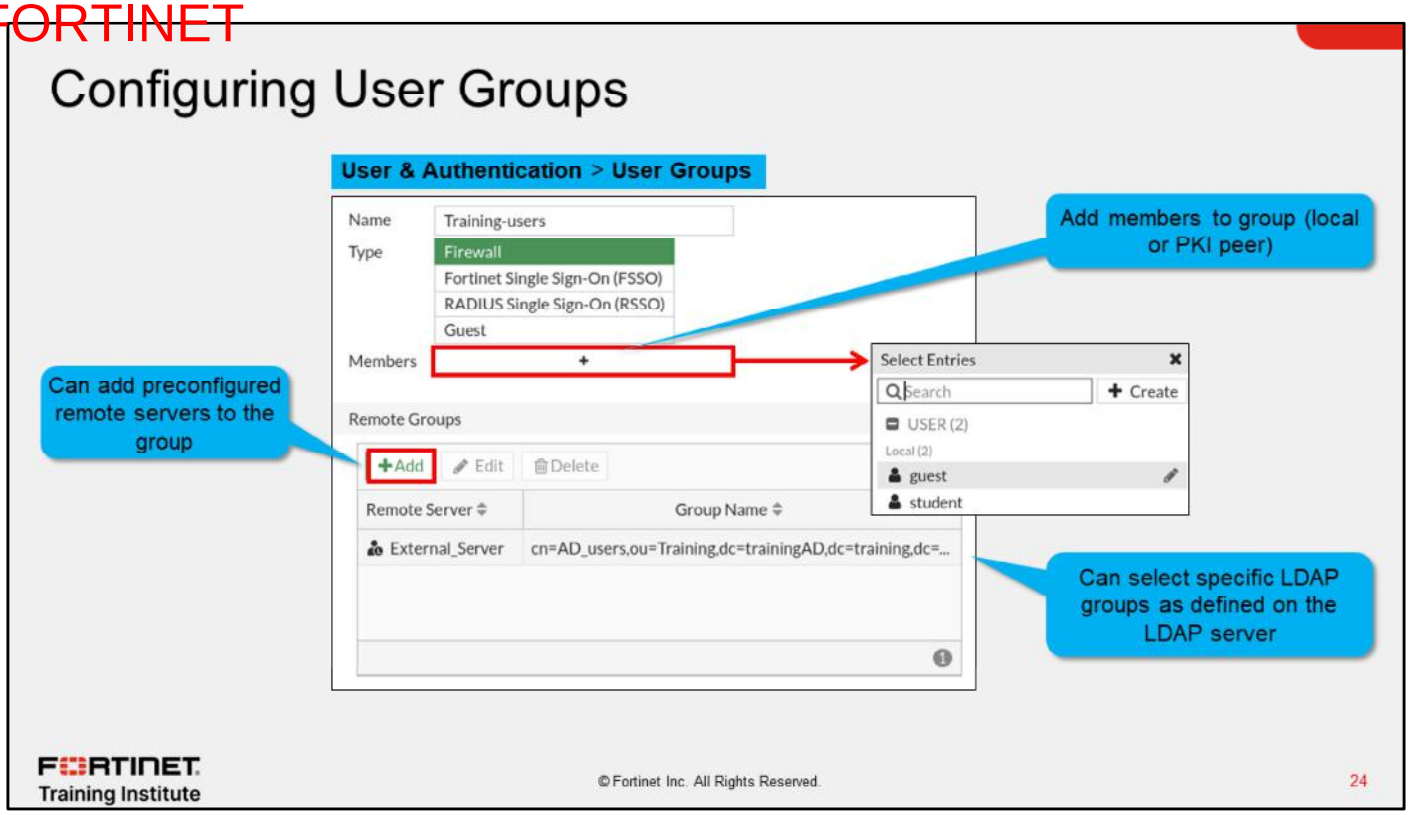

You can configure user groups on the **User Groups** page. You must specify the user group type and add users to the group. Depending on the group you create, you require different configurations. For the firewall user group, for example, members can consist of local users, PKI peer users, and users from one or more remote authentication servers. If your remote authentication server is an LDAP server, you can select specific LDAP groups to add to your user group, as defined on the LDAP server. Note that you can also select RADIUS groups, but this requires additional configuration on your RADIUS server and FortiGate (see the Fortinet Knowledge Base at kb.fortinet.com).

User groups simplify your configuration if you want to treat specific users in the same way, for example, if you want to provide the entire training department with access to the same network resources. If you want to treat all users differently, you need to add all users to firewall policies separately.

Knowledge Check

1. Which statement about guest user groups is true?
   - ✓ A. Guest user group accounts are temporary.
   - B. Guest user group account passwords are temporary.

2. Guest accounts are most commonly used for which purposes?
   - A. To provide temporary visitor access to corporate network resources
   - ✓ B. To provide temporary visitor access to wireless networks

FORTINET Training Institute

25

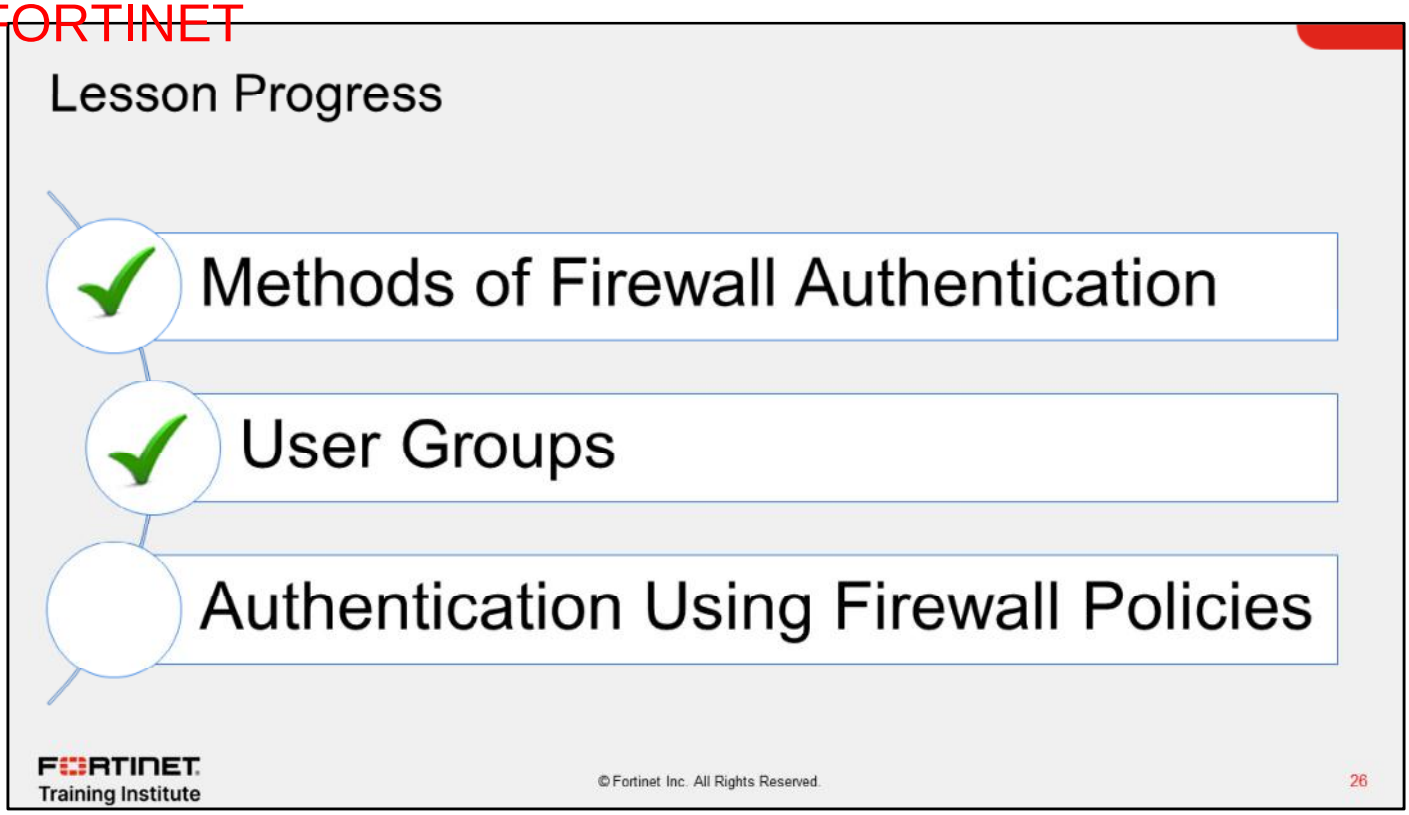

Good job! You now understand the basics of user groups.

Now, you will learn about using firewall policies for authentication.

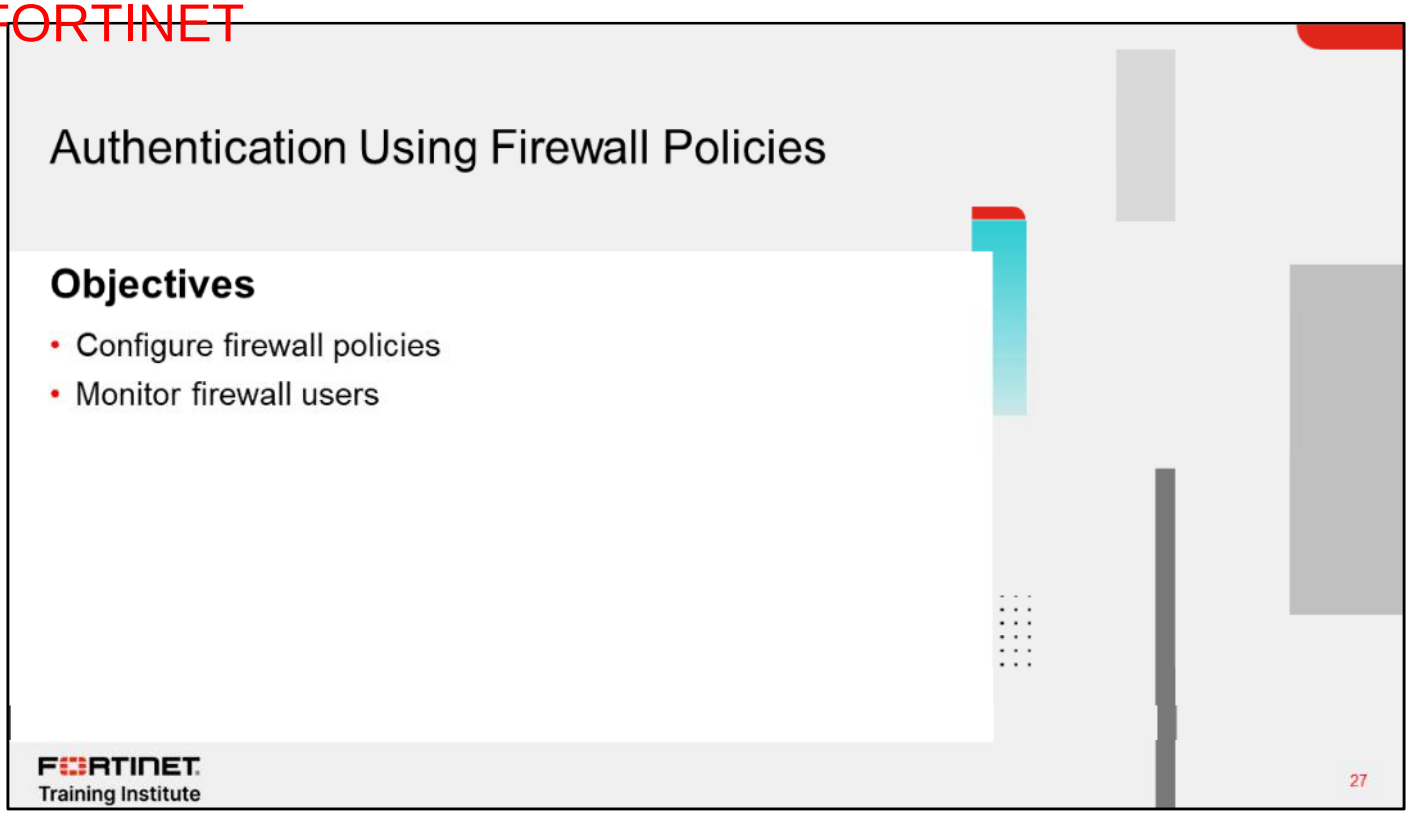

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in firewall policies, you will be able to configure firewall policies to enforce authentication on specific users and user groups.

Firewall Policy—Source

- Firewall policies can use user and user group objects to define the source. The objects include:
  - Local firewall accounts
  - External (remote) server accounts
  - PKI (certificate) users
  - FSSO users
- Anyone who belongs to the group and provides correct information will have a successful authentication

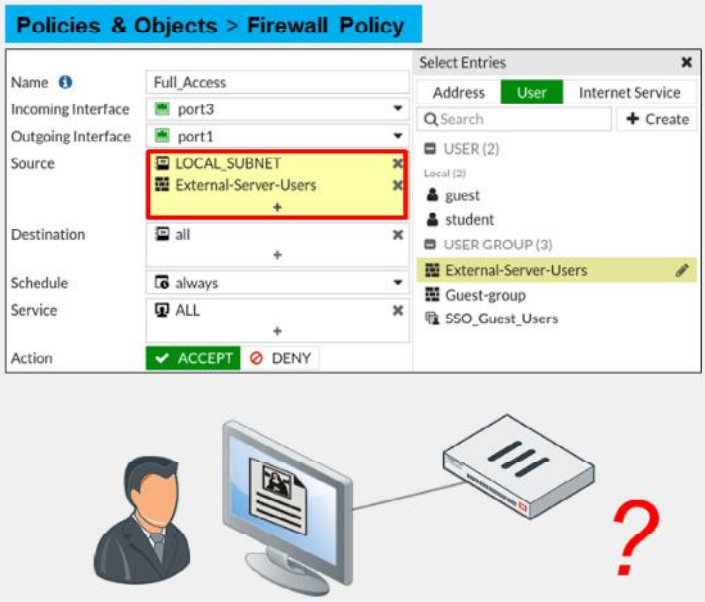

FORTINET Training Institute

28

A firewall policy consists of access and inspection rules (compartmentalized sets of instructions) that tell FortiGate how to handle traffic on the interface whose traffic they filter. After the user makes an initial connection attempt, FortiGate checks the firewall policies to determine whether to accept or deny the communication session. However, a firewall policy also includes a number of other instructions, such as those dealing with authentication. You can use the source of a firewall policy for this purpose. The source of a firewall policy must include the source address (IP address), but you can also include the user and user group. In this way, any user, or user group that is included in the source definition for the firewall policy can successfully authenticate.

User and user group objects can consist of local firewall accounts, external server accounts, PKI users, and FSSO users.

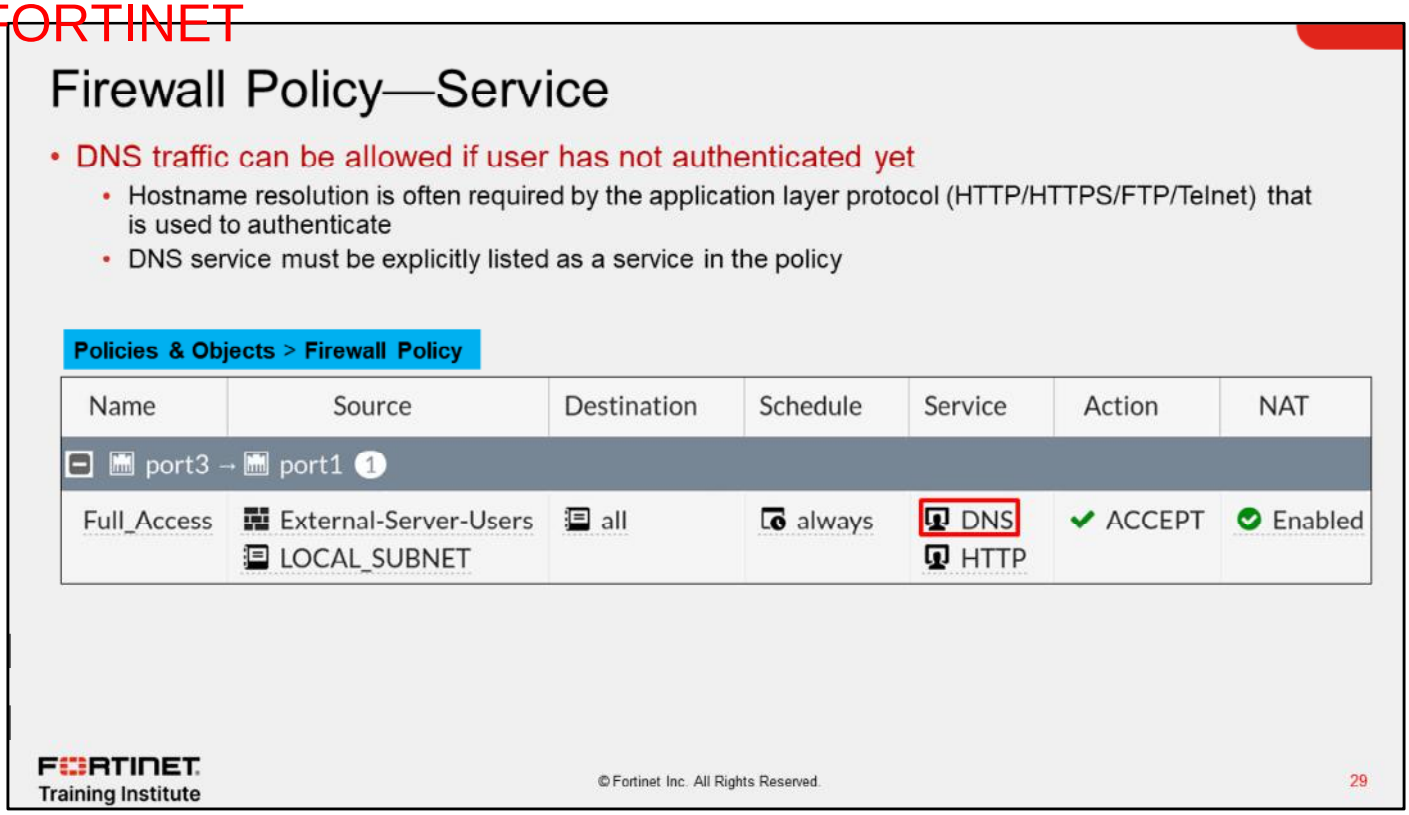

A firewall policy also checks the service in order to transport the named protocols or group of protocols. No service (with the exception of DNS) is allowed through the firewall policy before successful user authentication. DNS is usually used by HTTP so that people can use domain names for websites, instead of their IP address. DNS is allowed because it is a base protocol and will most likely be required to initially see proper authentication protocol traffic. Hostname resolution is almost always a requirement for any protocol. However, the DNS service must still be defined in the policy as allowed, in order for it to pass.

In the example shown on this slide, policy sequence 1 (Full_Access) allows users to use external DNS servers in order to resolve host names, before successful authentication. DNS is also allowed if authentication is unsuccessful because users need to be able to try to authenticate again. Any service that includes DNS would function the same way, like the default ALL service.

HTTP service is TCP port 80 and does not include DNS (UDP port 53).

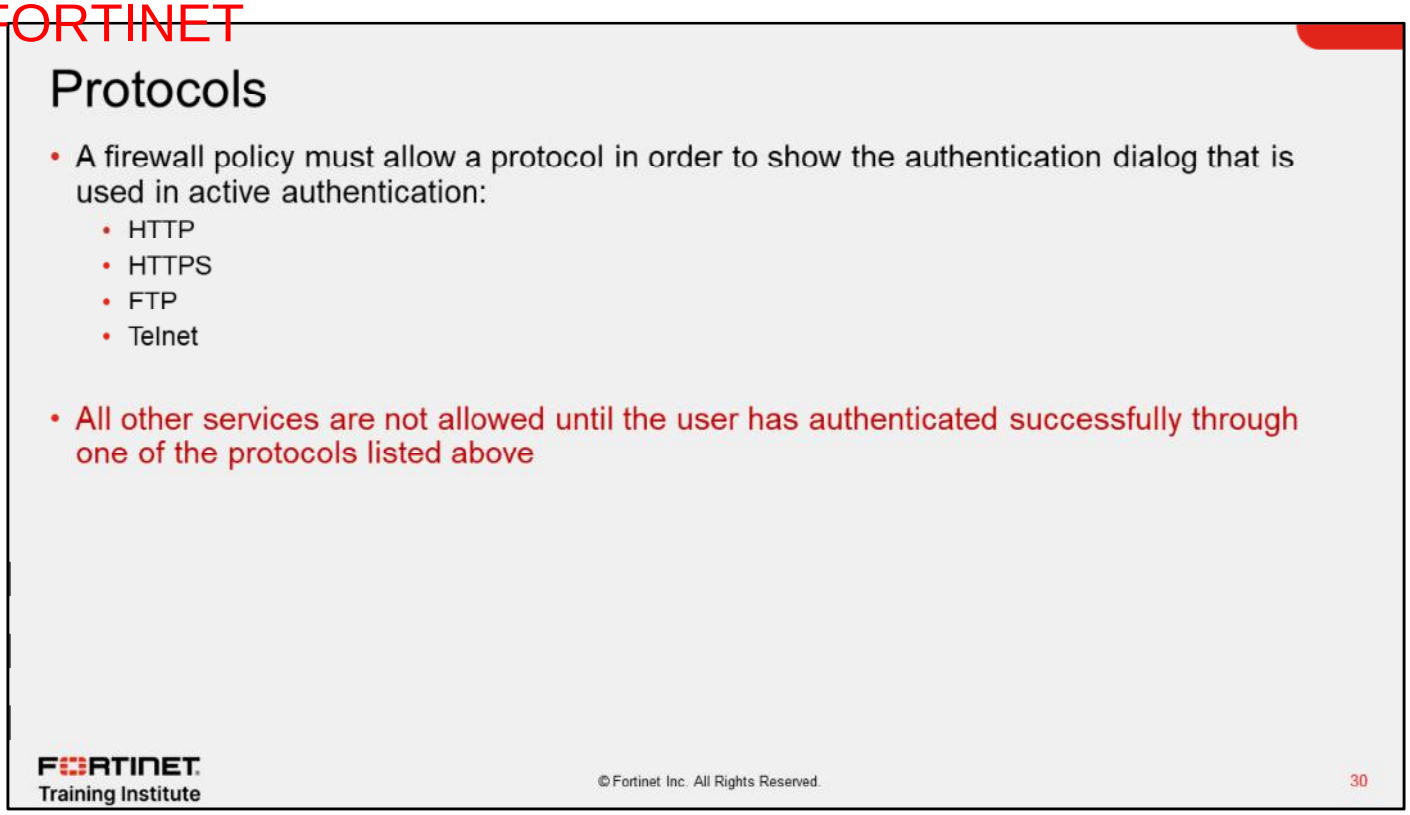

As well as the DNS service, the firewall policy must specify the allowed protocols, such as HTTP, HTTPS, FTP, and Telnet. If the firewall policy that has authentication enabled does not allow at least one of the supported protocols used for obtaining user credentials, the user will not be able to authenticate.

Protocols are required for all authentication methods that use active authentication (local password authentication, server-based password authentication, and two-factor authentication). Active authentication prompts the user for user credentials based on the following:

- The protocol of the traffic
- The firewall policy

Passive authentication, on the other hand, determines the user identity behind the scenes, and does not require any specific services to be allowed within the policy.

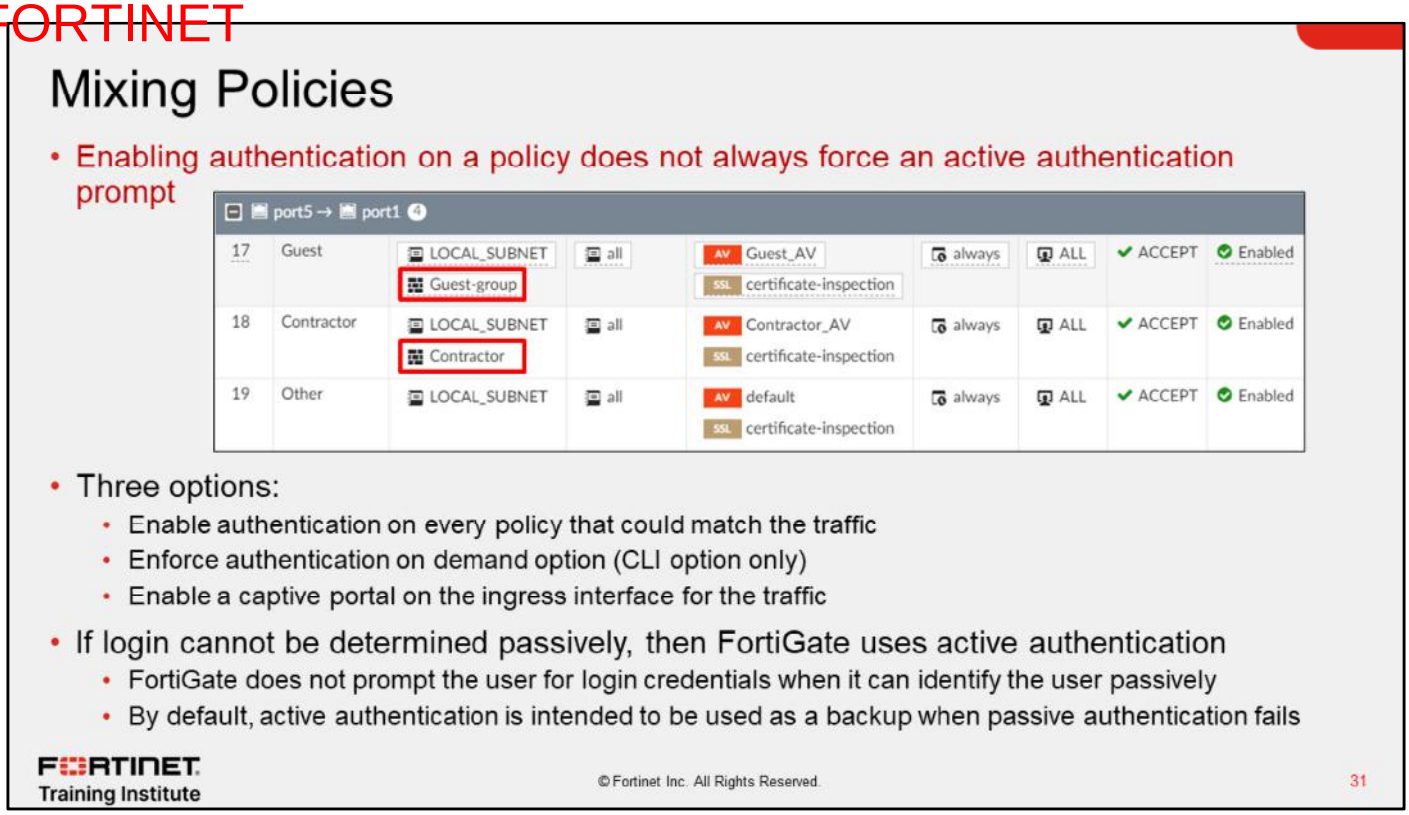

In the example shown on this slide, assuming active authentication is used, any initial traffic from LOCAL_SUBNET will not match policy sequence 17 (Guest). Policy sequence 17 looks for both IP and user, and user group information (LOCAL_SUBNET and Guest-group respectively), and since the user has not yet authenticated, the user group aspect of the traffic does not match. Since the policy match is not complete, FortiGate continues its search down the sequence list, to see if there is a complete match.

Next, FortiGate evaluates policy sequence 18 to see if the traffic matches. It will not for the same reason it did not match 17.

Finally, FortiGate evaluates policy sequence 19 to see if the traffic matches. It matches all criteria, so traffic is allowed with no need to authenticate.

When you use only active authentication, if all possible policies that could match the source IP have authentication enabled, then the user will receive a login prompt (assuming they use an acceptable login protocol). In other words, if policy sequence 19 also had authentication enabled, the users would receive login prompts.

If you use passive authentication and it can successfully obtain user details, then traffic from LOCAL_SUBNET with users that belong to Guest-group will apply to policy sequence 17, even though policy sequence 19 does not have authentication enabled.

If you use both active and passive authentication, and FortiGate can identify a user's credentials through passive authentication, the user never receives a login prompt, regardless of the order of any firewall policies. This is because there is no need for FortiGate to prompt the user for login credentials when it can identify who the user is passively. When you combine active and passive authentication methods, active authentication is

intended to be used as a backup, to be used only when passive authentication fails.

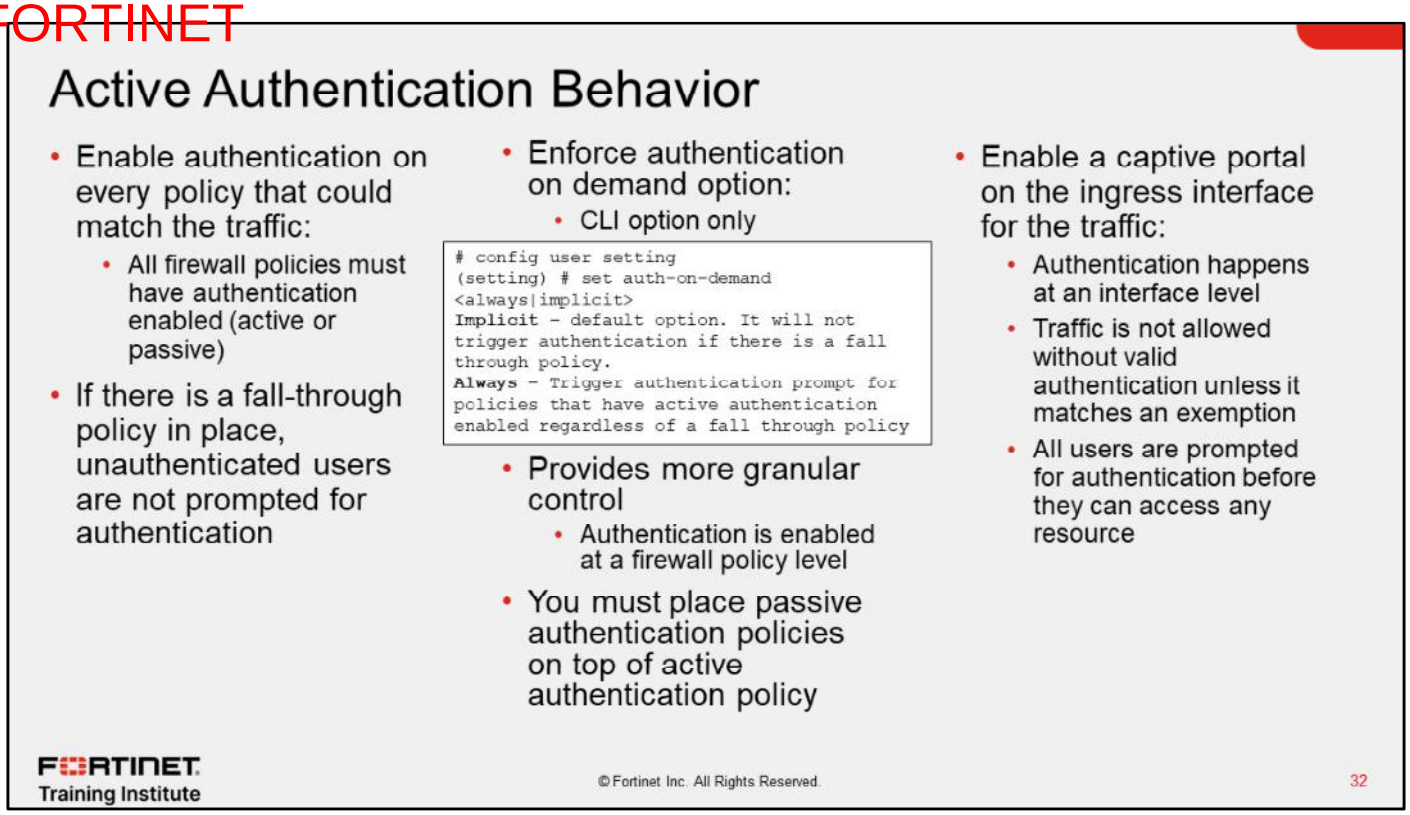

As mentioned earlier, there are three different ways you can alter active authentication behavior. If you have an active authentication firewall policy followed by a fall-through policy that does not have authentication enabled on it, then all traffic will use the fall-through policy. This means that users are not asked to authenticate. By default, all traffic passes through the catch-all policy without being authenticated. You can alter this behavior by enabling authentication on all firewall policies. When you enable authentication, all the systems must authenticate before traffic is placed on the egress interface.

Alternatively, only on the CLI, you can change the `auth-on-demand` option to `always`. This instructs FortiGate to trigger an authentication request, if there is a firewall policy with active authentication enabled. In this case, the traffic is allowed until authentication is successful.

If you want to have all users connect to a specific interface, then it is better to enable captive portal authentication at the interface level. This way, all devices must authenticate before they are allowed to access any resources.

Authentication Timeout

```
#config user setting
   set auth-timeout-type [idle-timeout|hard-timeout|new-session]
 end
```

- Timeout specifies how long a user can remain idle before the user must authenticate again
  - Default is five minutes
- Three options for behavior:
  - Idle (default): no traffic for that amount of time
  - Hard: authentication expires after that amount of time, regardless of activity
  - New session: authentication expires if no new session is created in that amount of time

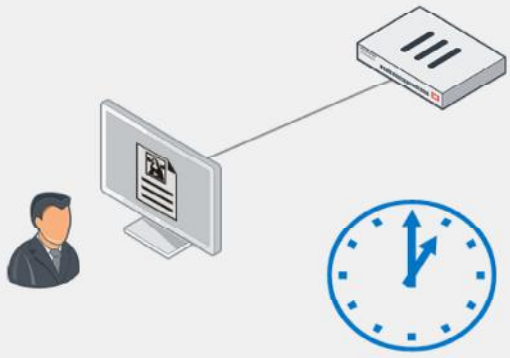

FORTINET Training Institute

33

An authentication timeout is useful for security purposes. It minimizes the risk of someone using the IP of the legitimate authenticated user. It also ensures users do not authenticate and then stay in memory indefinitely. If users stayed in memory forever, it would eventually lead to memory exhaustion.

There are three options for timeout behavior:

- **Idle**: looks at the packets from the host IP. If there are no packets generated by the host device in the configured timeframe, then the user is logged out.
- **Hard**: time is an absolute value. Regardless of the user's behavior, the timer starts as soon as the user authenticates and expires after the configured value.
- **New session**: even if traffic is being generated on existing communications channels, the authentication expires if no new sessions are created through the firewall from the host device within the configured timeout value.

Choose the type of timeout that best suits the authentication needs of your environment.

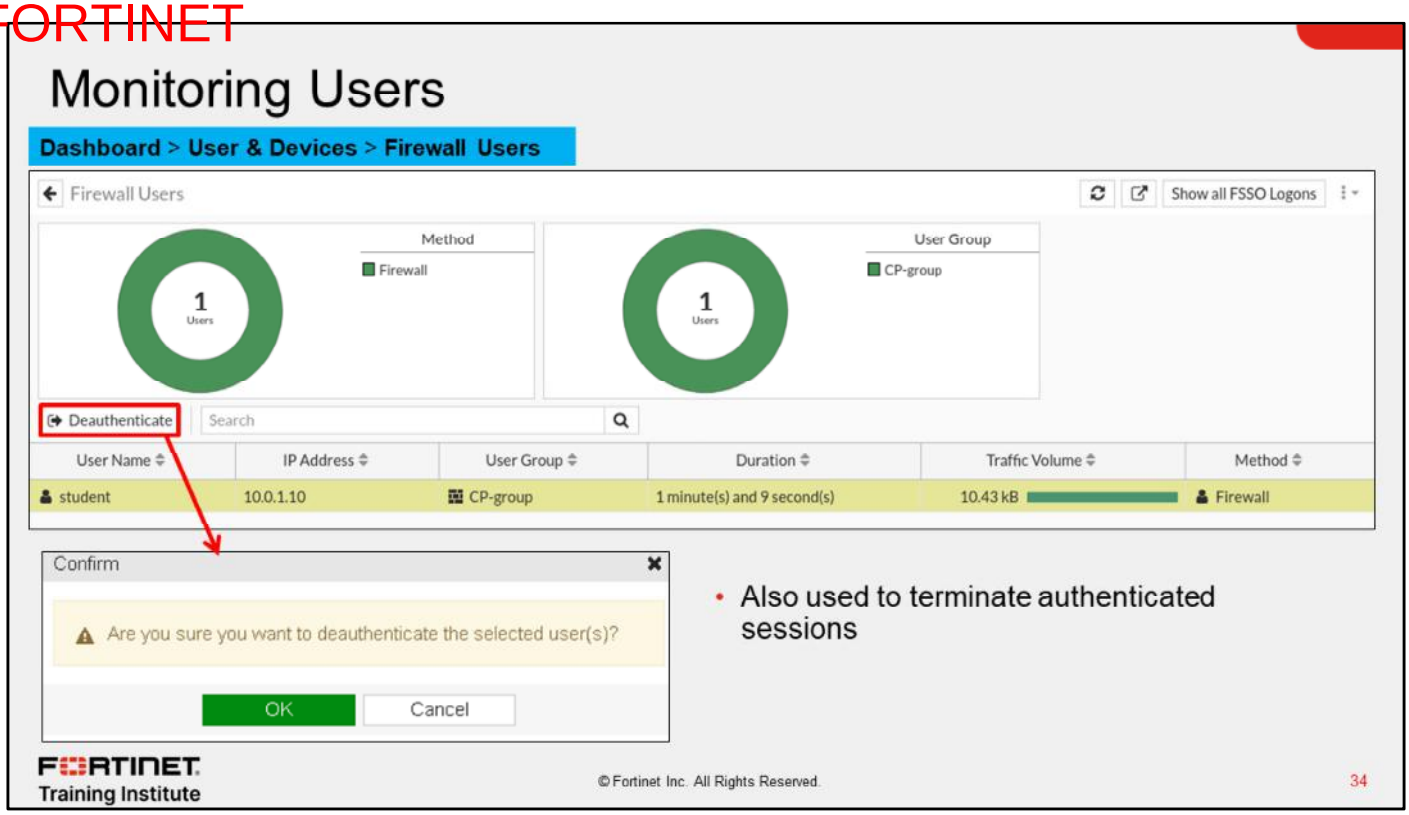

You can monitor users who authenticate through your firewall policies using the **Dashboard** > **User & Devices** > **Firewall Users** page. It displays the user, user group, duration, IP address, traffic volume, and authentication method.

It does not include administrators, because they are not authenticating through firewall policies that allow traffic. They are logging in directly on FortiGate.

This page also allows you to disconnect a user, or multiple users, at the same time.

Knowledge Check

1. Firewall policies dictate whether a user or device can or cannot authenticate on a network. Which statement about firewall authentication is true?
   - ✓ A. Firewall policies can be configured to authenticate certificate users.
   - B. The order of the firewall policies always determines whether a user's credentials are determined actively or passively.
2. Which statement about active authentication is true?
   - A. Active authentication is always used before passive authentication.
   - ✓ B. The firewall policy must allow the HTTP, HTTPS, FTP, and/or Telnet protocols in order for the user to be prompted for credentials.
3. Which statement best describes the authentication idle timeout feature on FortiGate?
   - A. The length of time FortiGate waits for the user to enter their authentication credentials
   - ✓ B. The length of time an authenticated user is allowed to remain authenticated without any packets being generated by the host device

FORTINET Training Institute

35

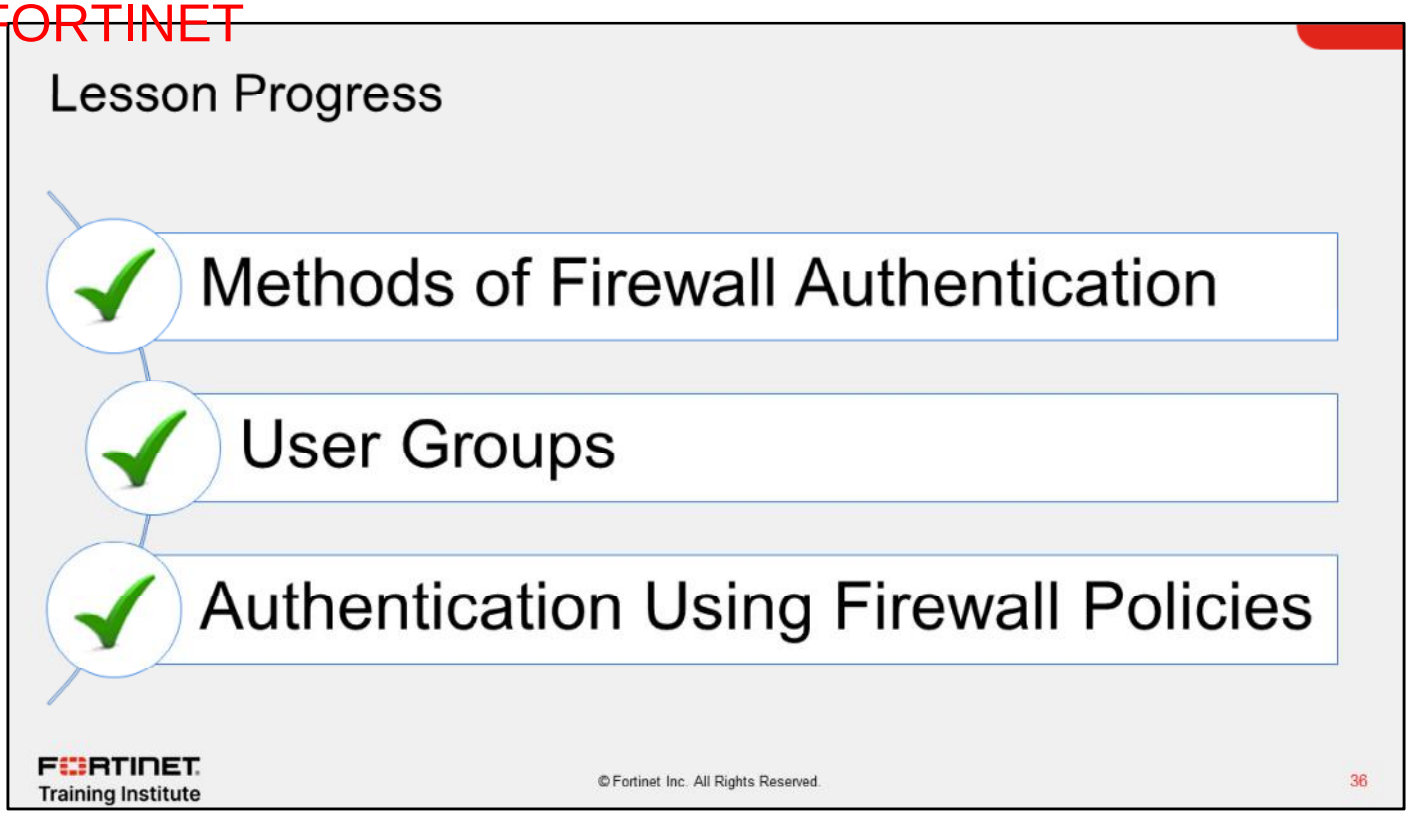

Congratulations! You have completed this lesson.

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

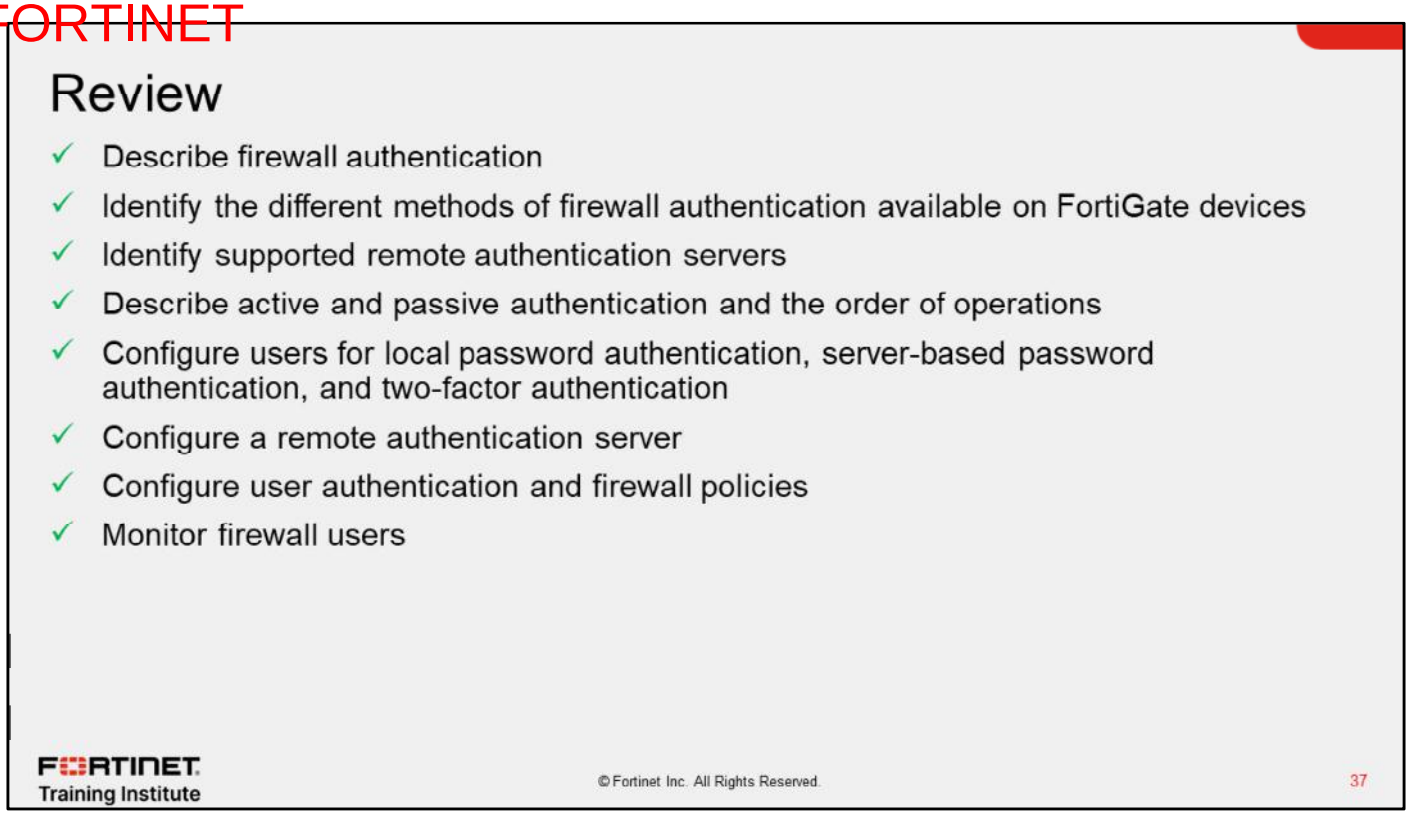

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use authentication on the firewall policies of FortiGate.

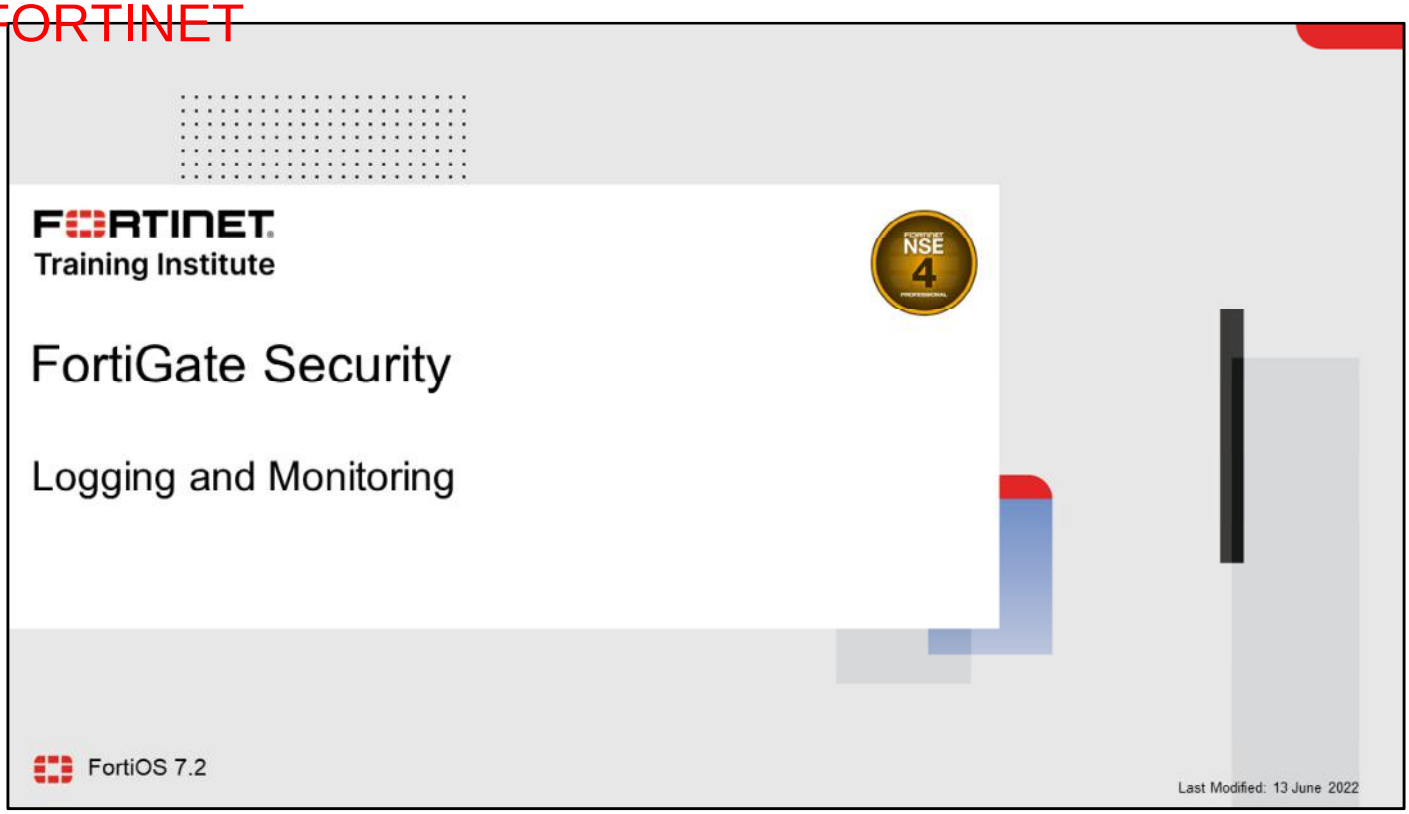

In this lesson, you will learn how to configure local and remote logging on FortiGate; view, search, and monitor logs; and protect your log data.

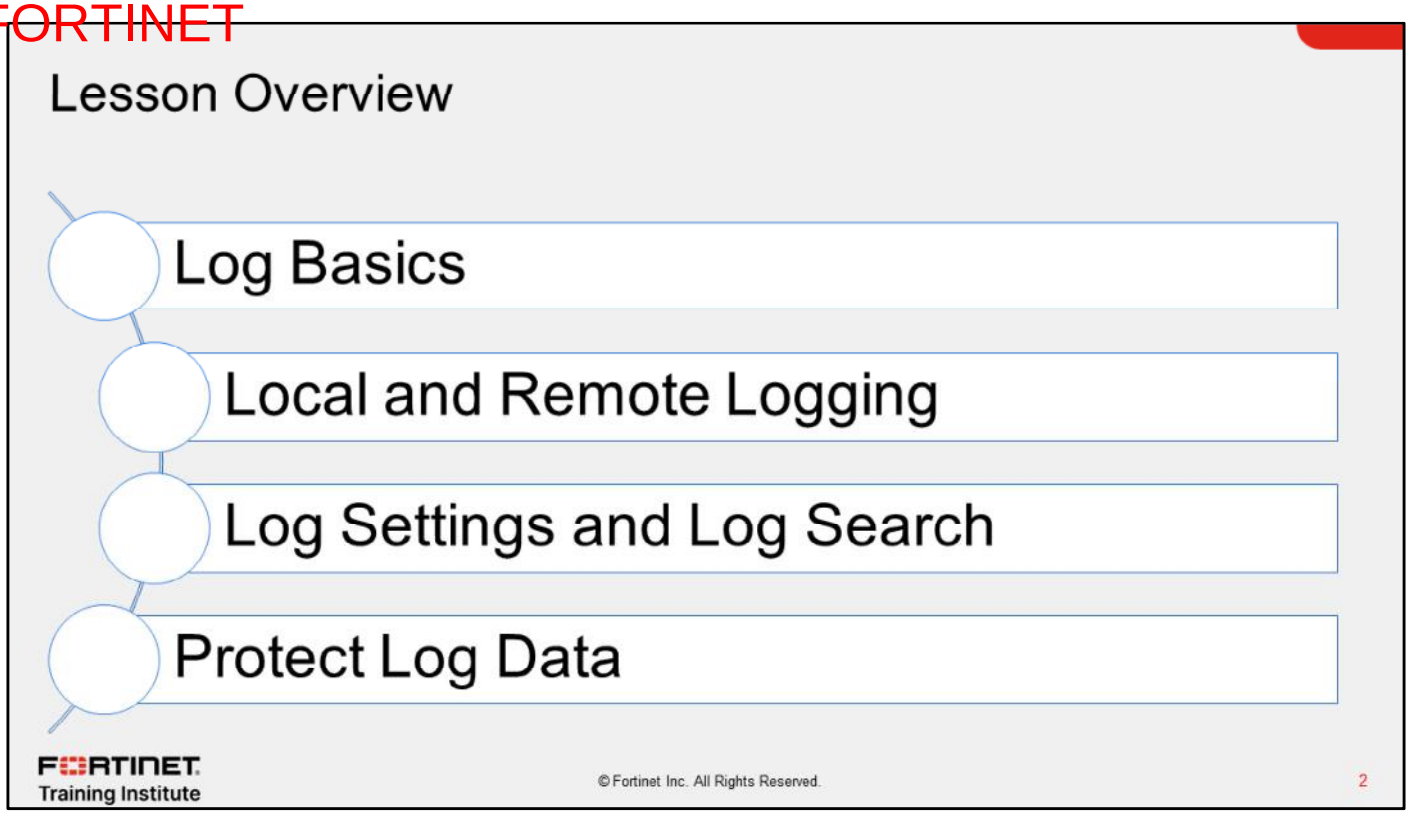

In this lesson, you will learn about the topics shown on this slide.

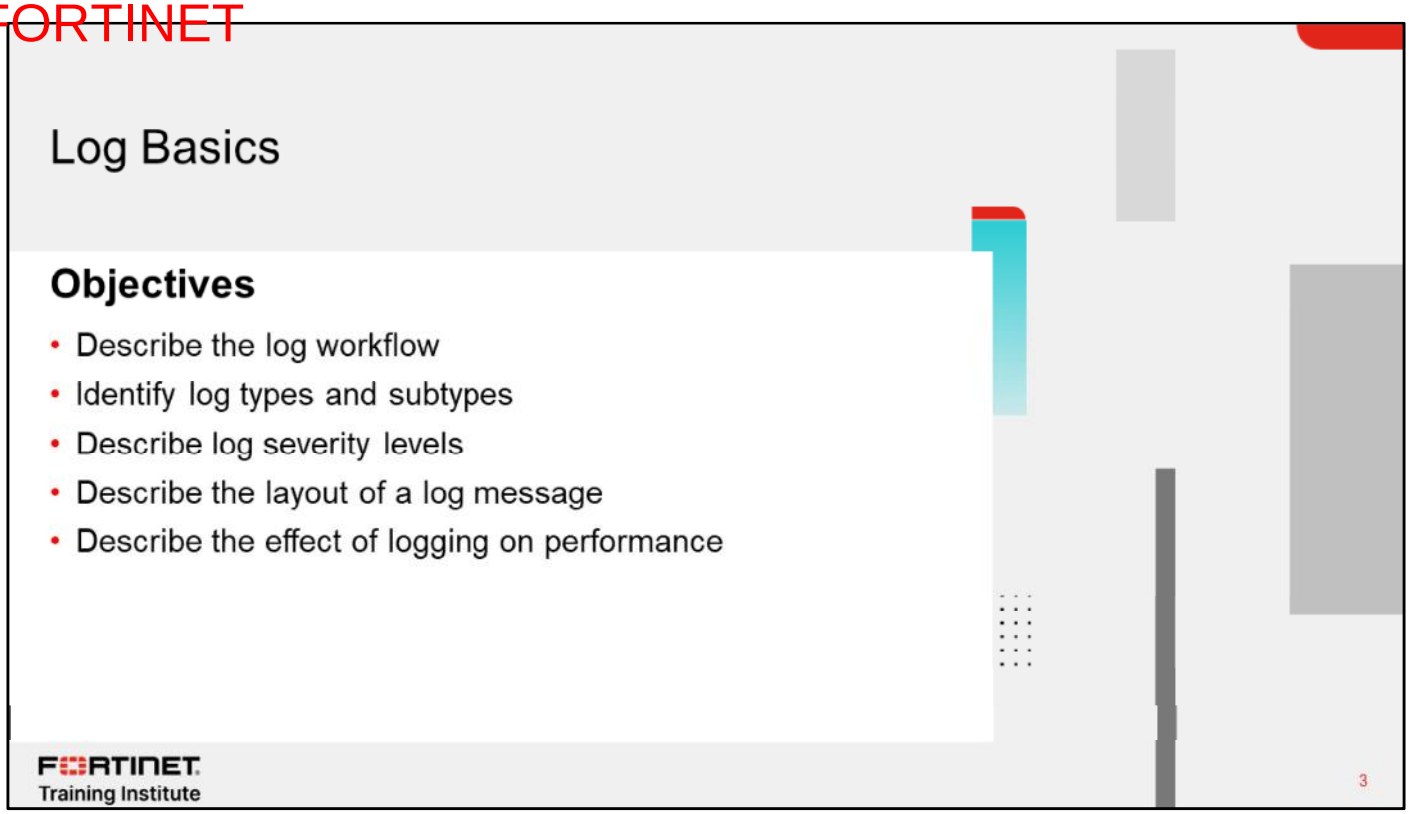

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in log basics, you will be able to more effectively analyze log data from your database.

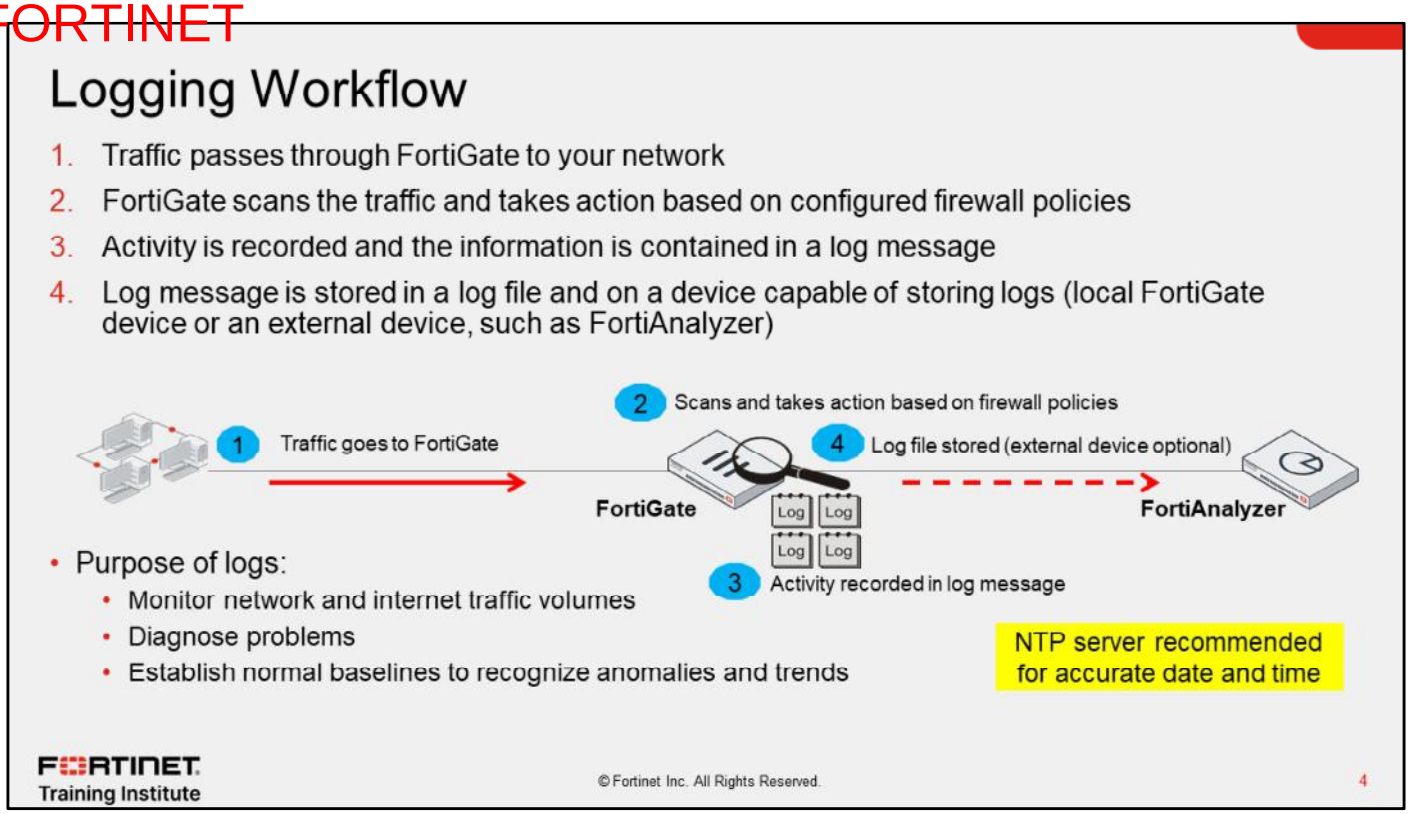

When traffic passes through FortiGate to your network, FortiGate scans the traffic, and then takes action based on the firewall policies in place. This activity is recorded, and the information is contained in a log message. The log message is stored in a log file. The log file is then stored on a device capable of storing logs. FortiGate can store logs locally on its own disk space, or can send logs to an external storage device, such as FortiAnalyzer.

The purpose of logs is to help you monitor your network traffic, locate problems, establish baselines, and more. Logs provide you with a greater perspective of your network, allowing you to adjust your network security settings if necessary.

Some organizations have legal requirements when it comes to logging, so it is important to be aware of your organization's policies during configuration.

For effective logging, your FortiGate system date and time should be accurate. You can either manually set the system date and time, or configure FortiGate to keep its time correct automatically by synchronizing with a Network Time Protocol (NTP) server. An NTP server is highly recommended.

Log Types and Subtypes

- *Traffic* logs record traffic flow information, such as an HTTP/HTTPS request and its response (if any)
- *Event* logs record system and administrative events, such as adding or modifying a setting, or daemon activities
- *Security* logs record security events, such as virus attacks and intrusion attempts, based on the security profile type (log type = utm)

Traffic	Event	Security
Forward	Endpoint	Application Control
Local	High Availability	Antivirus
Sniffer	General System	DNS Query
	User	File Filter
	Router	Web Filter
	VPN	Intrusion Prevention
	SD-WAN	Anomaly
	WiFi	SSL
	CIFS	SSH
	Security Ratings	
	SDN Connector	

FORTINET Training Institute

5

To FortiGate, there are three different types of logs: traffic logs, event logs, and security logs. Each type is further divided into subtypes.

Traffic logs record traffic flow information, such as an HTTP/HTTPS request and its response, if any. It contains subtypes named forward, local, and sniffer.

- Forward traffic logs contain information about traffic that FortiGate either accepted or rejected according to a firewall policy.
- Local traffic logs contain information about traffic directly to and from the FortiGate management IP addresses. They also include connections to the GUI and FortiGuard queries.
- Sniffer logs contain information related to traffic seen by the one-arm sniffer.

Event logs record system and administrative events, such as adding or modifying a setting, or daemon activities. It contains the subtypes listed on the slide.

- System event logs contain information related to operations, such as automatic FortiGuard updates and GUI logins.
- User logs contain login and logout events for firewall policies with user authentication.
- Router, VPN, and wireless subtypes include logs for those features. For example, VPN contains IPsec and SSL VPN log entries.

Finally, security logs record security events, such as virus attacks and intrusion attempts. They contain log entries based on the security profile type (log type = utm), including the subtype listed on the slide.

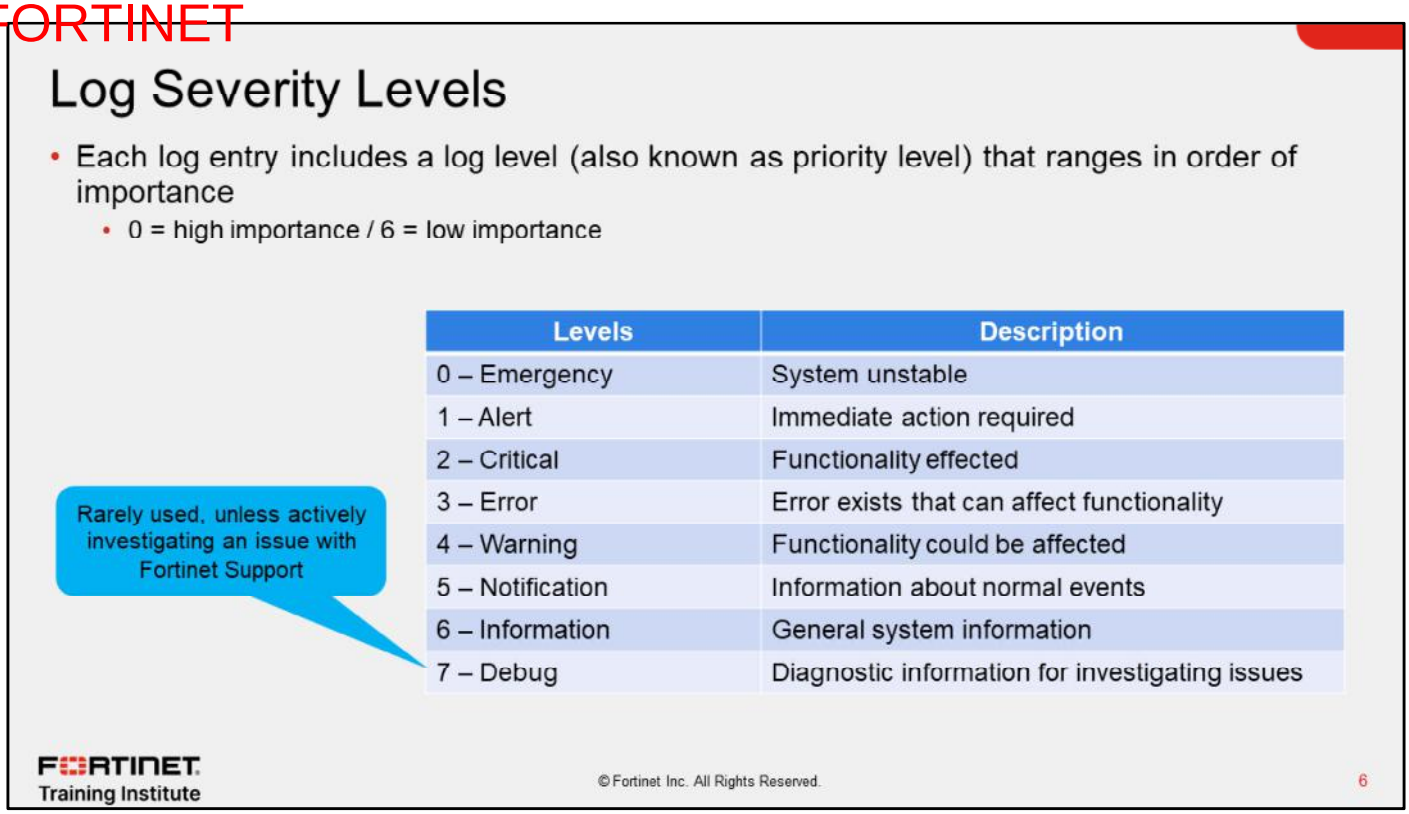

Each log entry includes a log level (or priority level) that ranges in order of importance from emergency to information.

There is also a debug level that includes diagnostic information into the event log. The debug level is rarely used, unless you are actively investigating an issue with Fortinet Support. Generally, the lowest level you want to use is information, but even this level generates many logs and can cause premature hard disk failure. Depending on the type of log and the needs of your organization, you may want to log only notification levels or higher.

You and your organization’s policies dictate what must be logged.

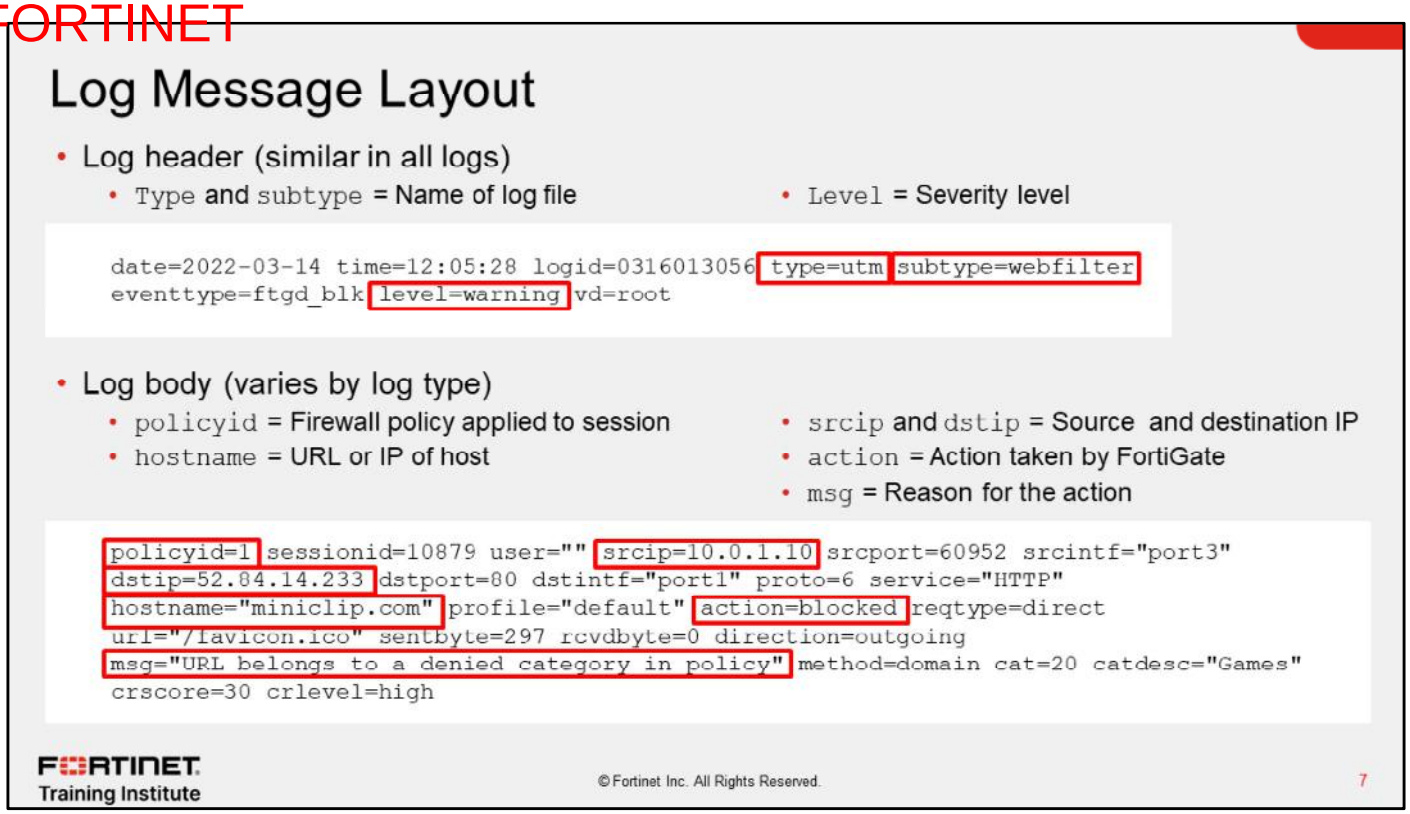

Every log message has a standard layout comprising two sections: a header and a body.

The header contains fields that are common to all log types, such as originating date and time, log identifier, log category, severity level, and virtual domain (VDOM). The value of each field, however, is specific to the log message. In the raw log entry example shown on this slide, the log type is `UTM`, the subtype is `webfilter`, and the level is `warning`. The type and subtype of logs determine what fields appear in the log body.

The body, therefore, describes the reason why the log was created, and the actions taken by FortiGate. These fields vary by log type. In the example shown on this slide, the fields are as follows:

- The `policyid` field indicates which firewall rule matched the traffic
- The `srcip` field indicates the source IP address
- The `dstip` field indicates the destination IP address
- The `hostname` field indicates the URL or IP of the host
- The `action` field indicates what FortiGate did when it found a policy that matched the traffic
- The `msg` field indicates the reason for the action taken. In this example, the action is `blocked`, which means that FortiGate prevented this IP packet from passing, and the reason is because it belonged to a denied category in the firewall policy.

If you log to a third-party device, such as a syslog server, knowing the log structure is crucial to integration. For information on log structures and associated meanings, visit http://docs.fortinet.com.

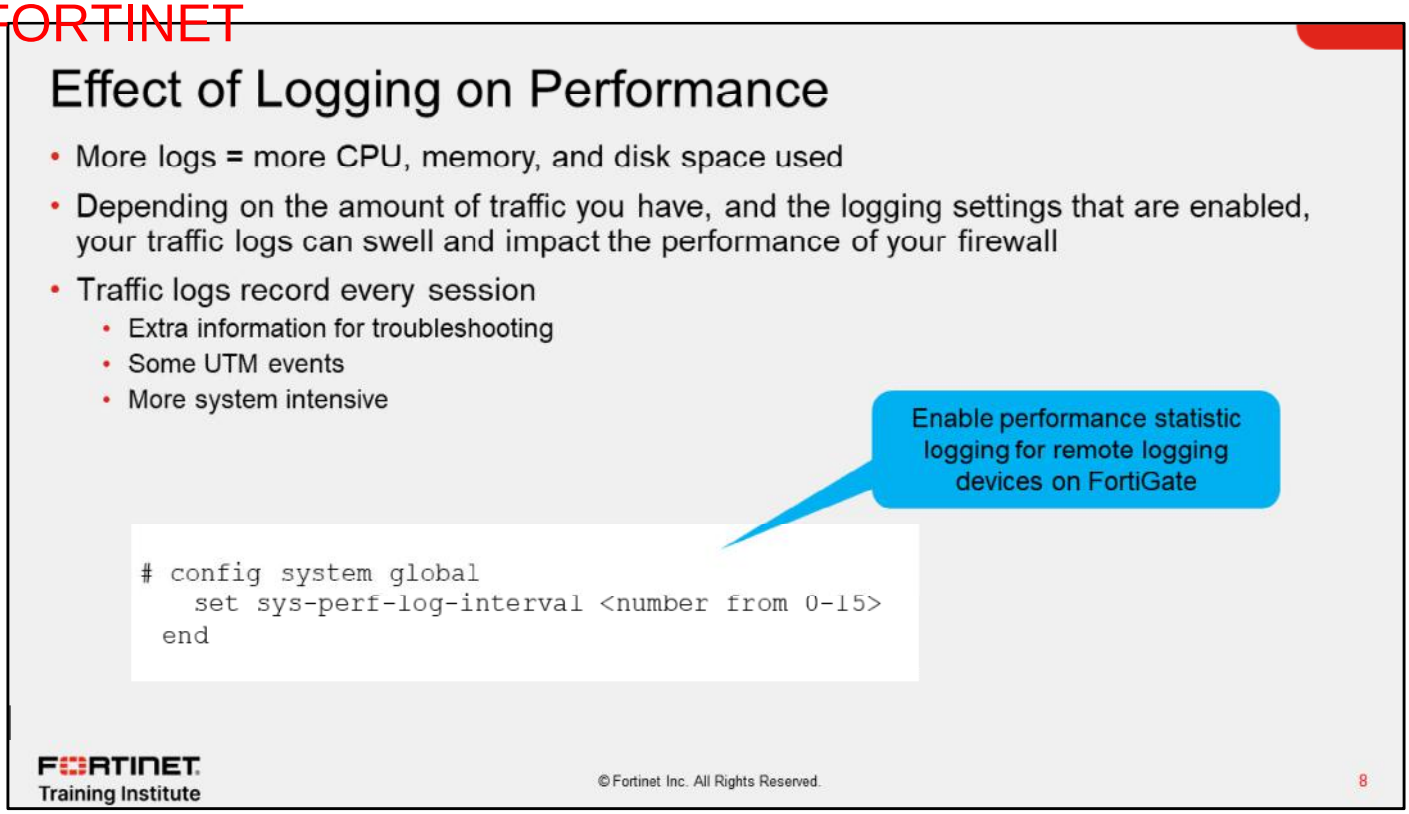

It is important to remember that the more logs that get generated, the heavier the toll on your CPU, memory, and disk resources. Storing logs for a period of time also requires disk space, as does accessing them. So, before configuring logging, make sure it is worth the extra resources and that your system can handle the influx.

Also important to note is the logging behavior with security profiles. Security profiles can, depending on the logging settings, create log events when a traffic matching the profile is detected. Depending on the amount of traffic you have, and the logging settings that are enabled, your traffic logs can swell and, ultimately, impact the performance of your firewall.

When using remote logging devices, such as FortiAnalyzer and syslog, you can enable performance statistic logging to occur every 1-15 minutes (0 to disable). This is not available for local disk logging or FortiCloud.

Knowledge Check

1. Which type of logs are application control, web filter, and antivirus?
   A. Event
   ✓B. Security

2. The log _____ contains fields that are common to all log types, such as originating date and time, log identifier, log category, and VDOM.
   ✓A. header
   B. body

FORTINET Training Institute

9

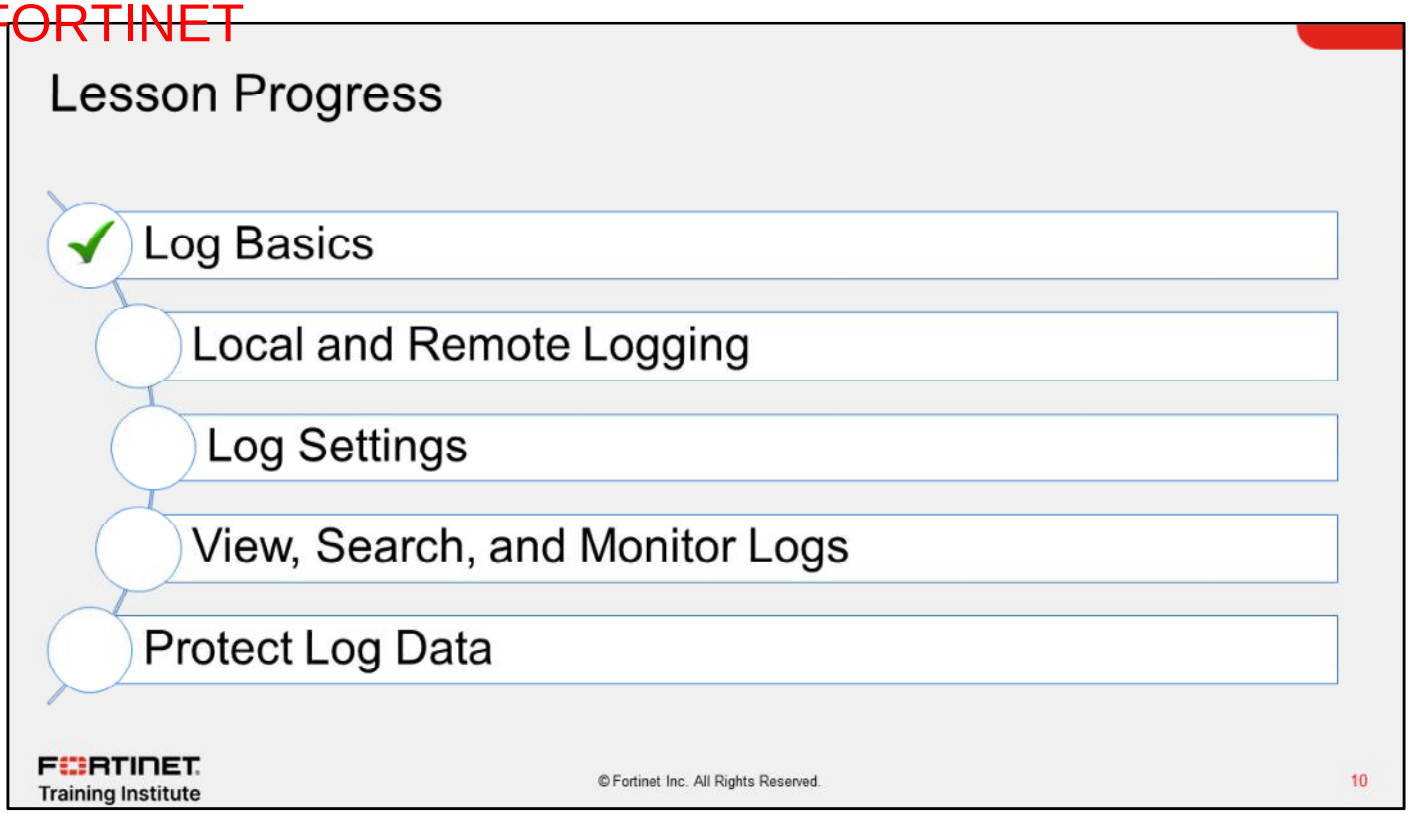

Good job! You now understand log basics.

Now, you will learn about local logging.

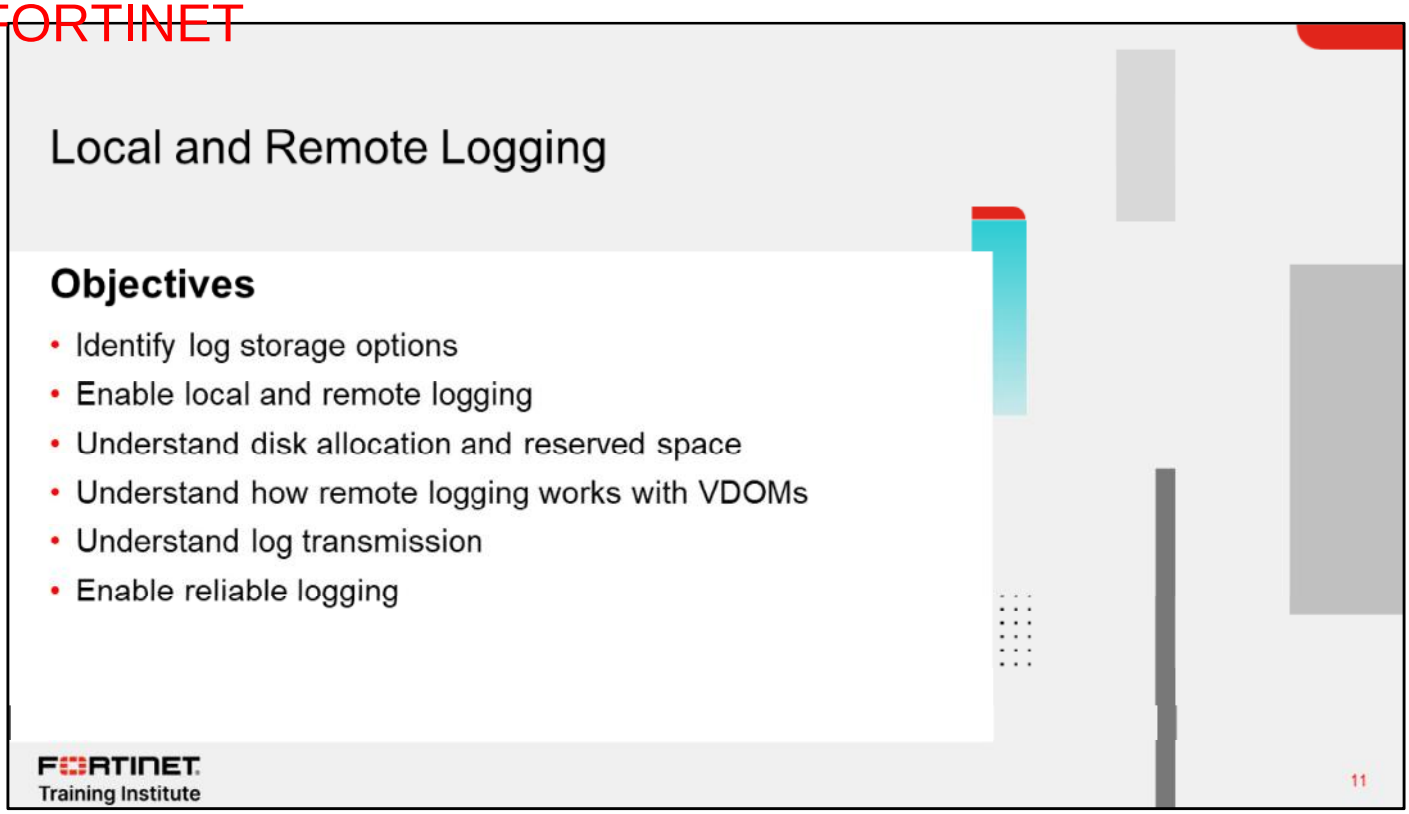

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in local logging, you will be able to successfully store logs to local disk and retain those logs, based on your requirements.

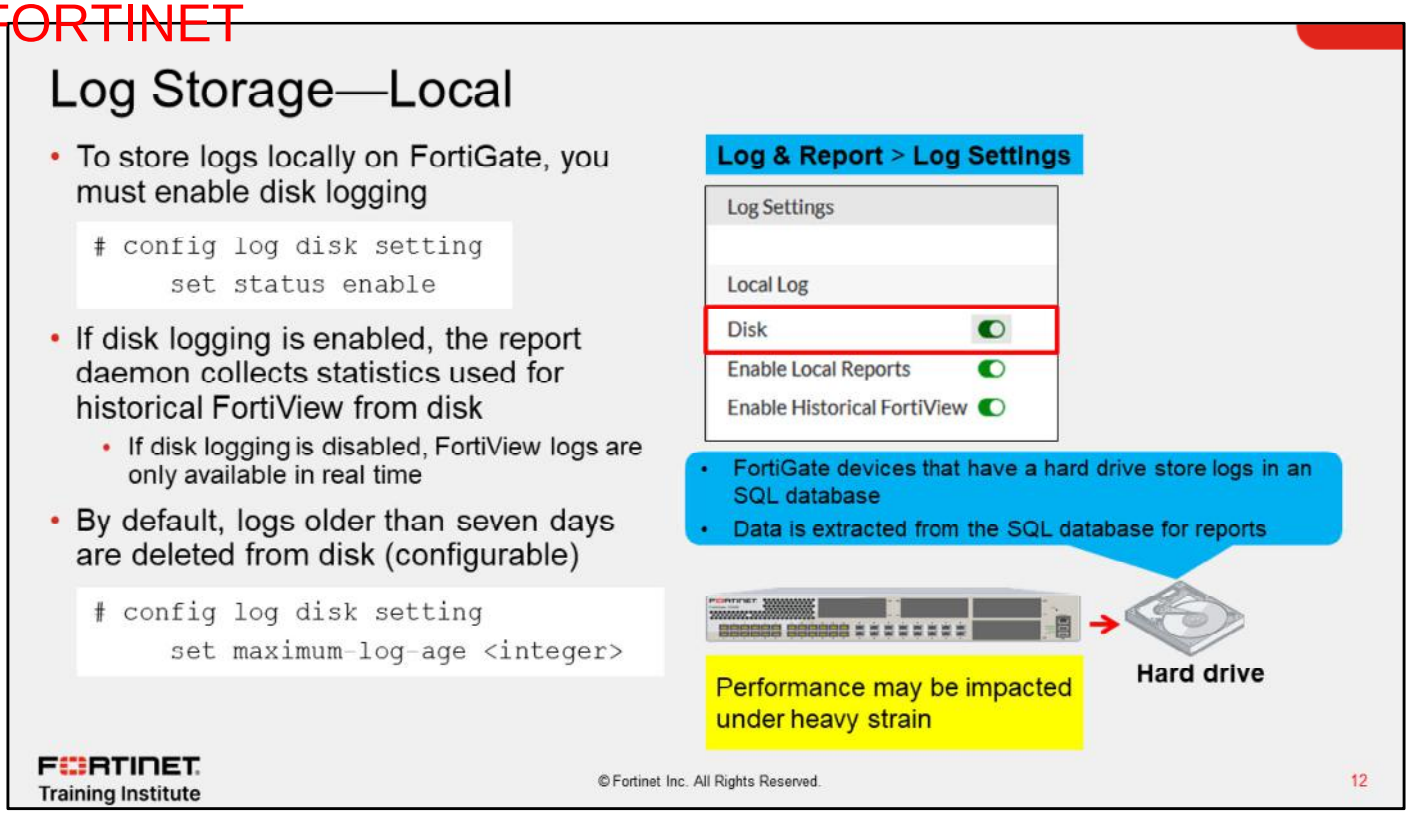

Typically, mid-level to high-end FortiGate models have a hard drive. FortiGate can store logs on its hard drive. This operation is known as local logging or disk logging. Depending on the model series, disk logging may be enabled by default.

FortiGate can store all log types, including log archives and traffic logs, locally. Traffic logs and log archives are larger files, and need a lot of room when being logged by FortiGate.

Under heavy log usage, disk logging will result in a performance impact.

If you are using the local hard disk on a device for WAN optimization, you cannot also log to disk, unless your device has two separate disks. If your device has two separate disks, you can use one for WAN optimization and the other for logging. If you are using the local hard disk for WAN optimization, and only one disk is available, you can log to remote FortiAnalyzer devices or syslog servers.

If you want to store logs locally on FortiGate, you must enable disk logging on the **Log Settings** page. Only some FortiGate models support disk logging. If your FortiGate does not support disk logging, you can log to an external device instead.

You must enable disk logging in order for information to appear on the FortiView dashboards. If disabled, logs display in real-time only. You can also enable this setting using the CLI `config log disk setting` command.

By default, logs older than seven days are deleted from the disk. This value is configurable.

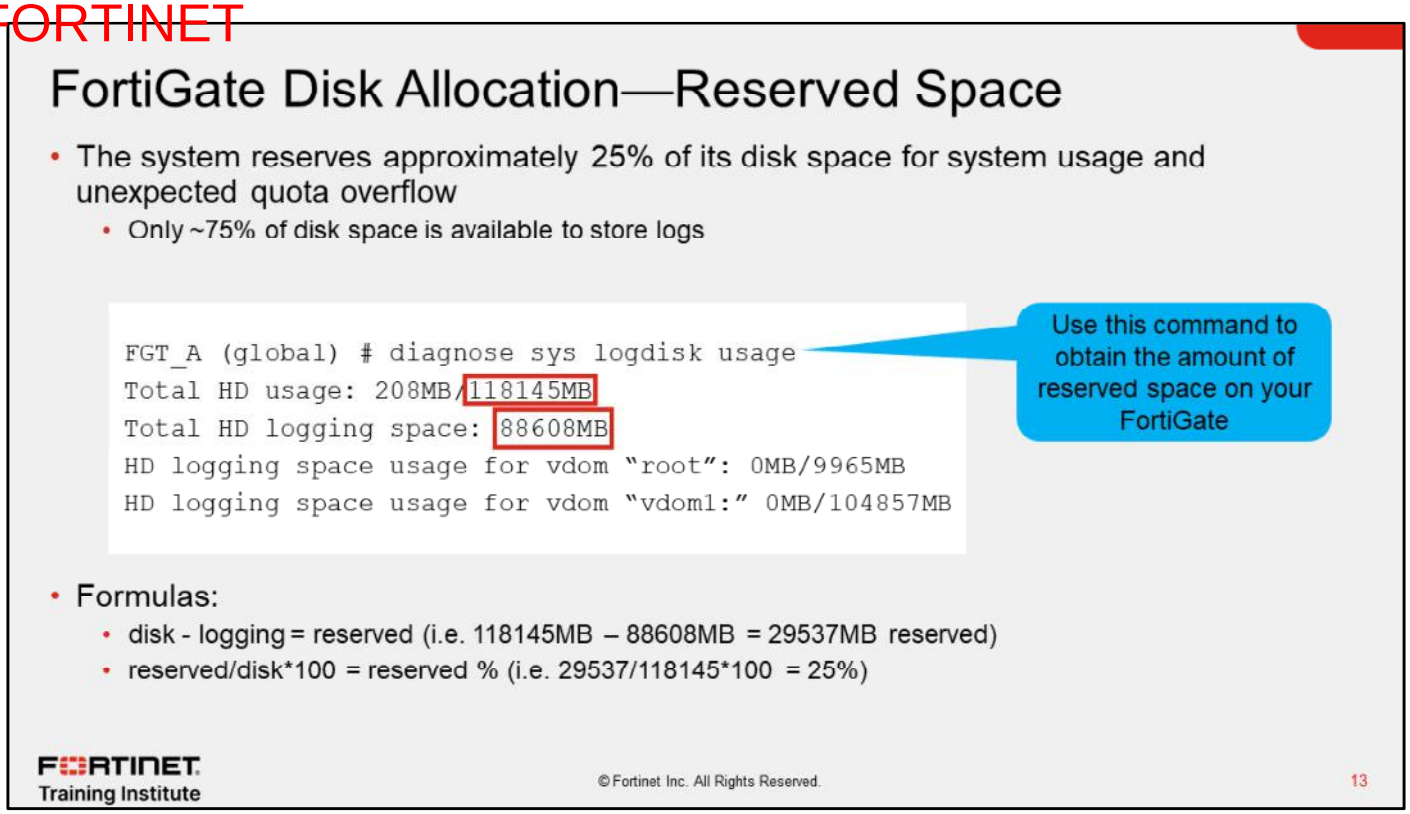

If you decide to log locally on FortiGate, be aware that the entire disk space is not available to store logs. The FortiGate system reserves approximately 25% of its disk space for system usage and unexpected quota overflow.

To determine the amount of reserved space on your FortiGate, use the CLI command `diagnose sys logdisk usage`. Subtract the total logging space from the total disk space to calculate the reserved space.

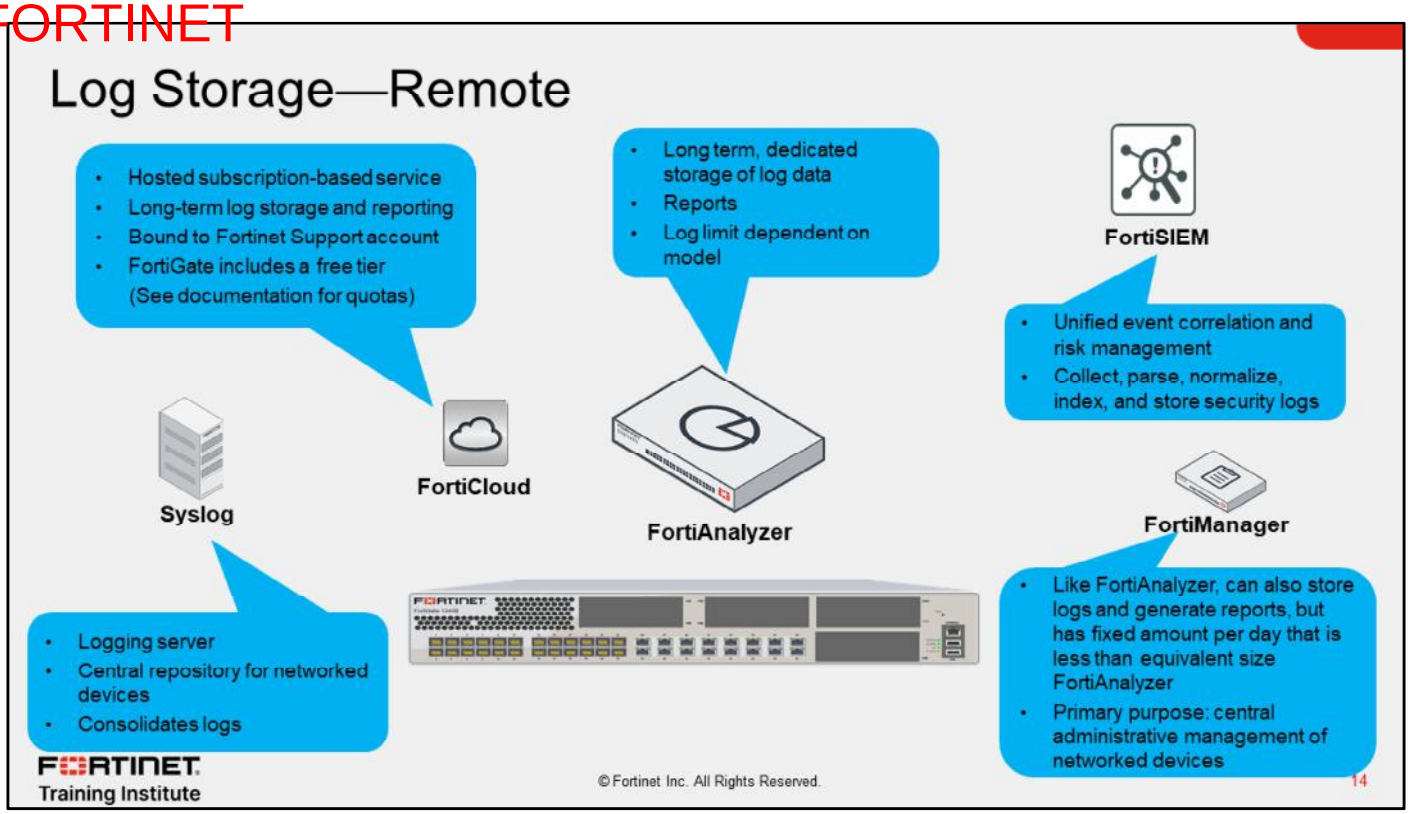

You can configure FortiGate to store logs on syslog servers, FortiCloud, FortiSIEM, FortiAnalyzer, or FortiManager. These logging devices can also be used as a backup solution. Whenever possible, it is preferred to store logs externally.

Syslog is a logging server that is used as a central repository for networked devices.

FortiCloud is a Fortinet subscription-based, hosted security management and log retention service that offers long-term storage of logs with reporting. If you have a smaller network, FortiCloud is usually more feasible than buying a dedicated logging device. Note that every FortiGate offers a free tier and will keep logs for seven days. You must upgrade to the paid service to retain logs for one year.

FortiSIEM provides unified event correlation and risk management that can collect, parse, normalize, index, and store security logs.

FortiAnalyzer and FortiManager are external logging devices with which FortiGate can communicate. You can place FortiAnalyzer or FortiManager in the same network as FortiGate, or outside of it. While FortiAnalyzer and FortiManager share a common hardware and software platform and can both take log entries, FortiAnalyzer and FortiManager actually have different capabilities that are worth noting. The primary purpose of FortiManager is to centrally manage multiple FortiGate devices. As such, log volumes are limited to a fixed amount per day, which are less than the equivalent size FortiAnalyzer. On the other hand, the primary purpose of FortiAnalyzer is to store and analyze logs, so the log limit is much higher (though the limit is model dependent). Note that local logging is not required for you to configure logging to FortiAnalyzer or FortiManager.

FortiAnalyzer and FortiManager Log Storage

- FortiGate can send logs to both FortiAnalyzer and FortiManager (FortiGate must be a registered device)

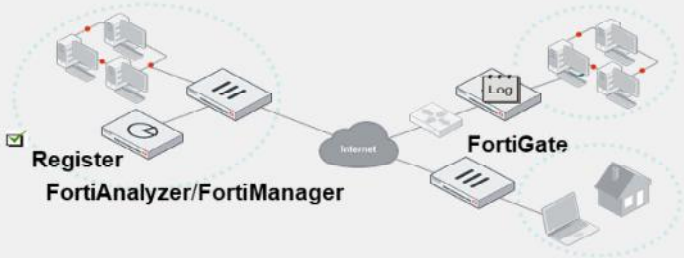

Log & Report > Log Settings

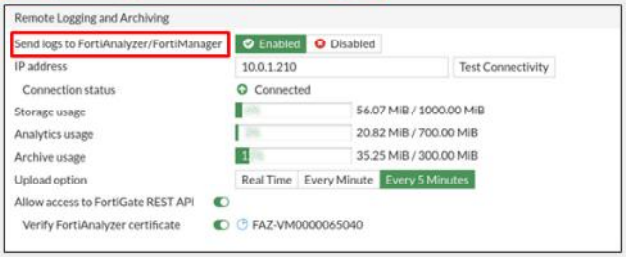

- Can configure up to three separate FortiAnalyzer and FortiManager devices or one cloud FortiAnalyzer instance using the CLI
  - Multiple devices may be needed for redundancy
  - Generating and sending logs requires resources—be aware!

```
# config log [fortianalyzer | fortianalyzer-
cloud|fortianalyzer2|fortianalyzer3] setting
  set status enable
  set server <server_IP>
  end
```

Commands *not* cumulative

FORTINET Training Institute

15

The process to configure FortiGate to send logs to FortiAnalyzer or FortiManager is identical. For FortiGate to send logs to either device, you must register FortiGate with FortiAnalyzer or FortiManager. After it is registered, FortiAnalyzer or FortiManager can begin to accept incoming logs from FortiGate.

You can configure remote logging to FortiAnalyzer or FortiManager using both the GUI and CLI.

Note that the **Test Connectivity** function on the GUI will report as failing until FortiGate is registered on FortiAnalyzer or FortiManager, because it is not yet authorized to send logs.

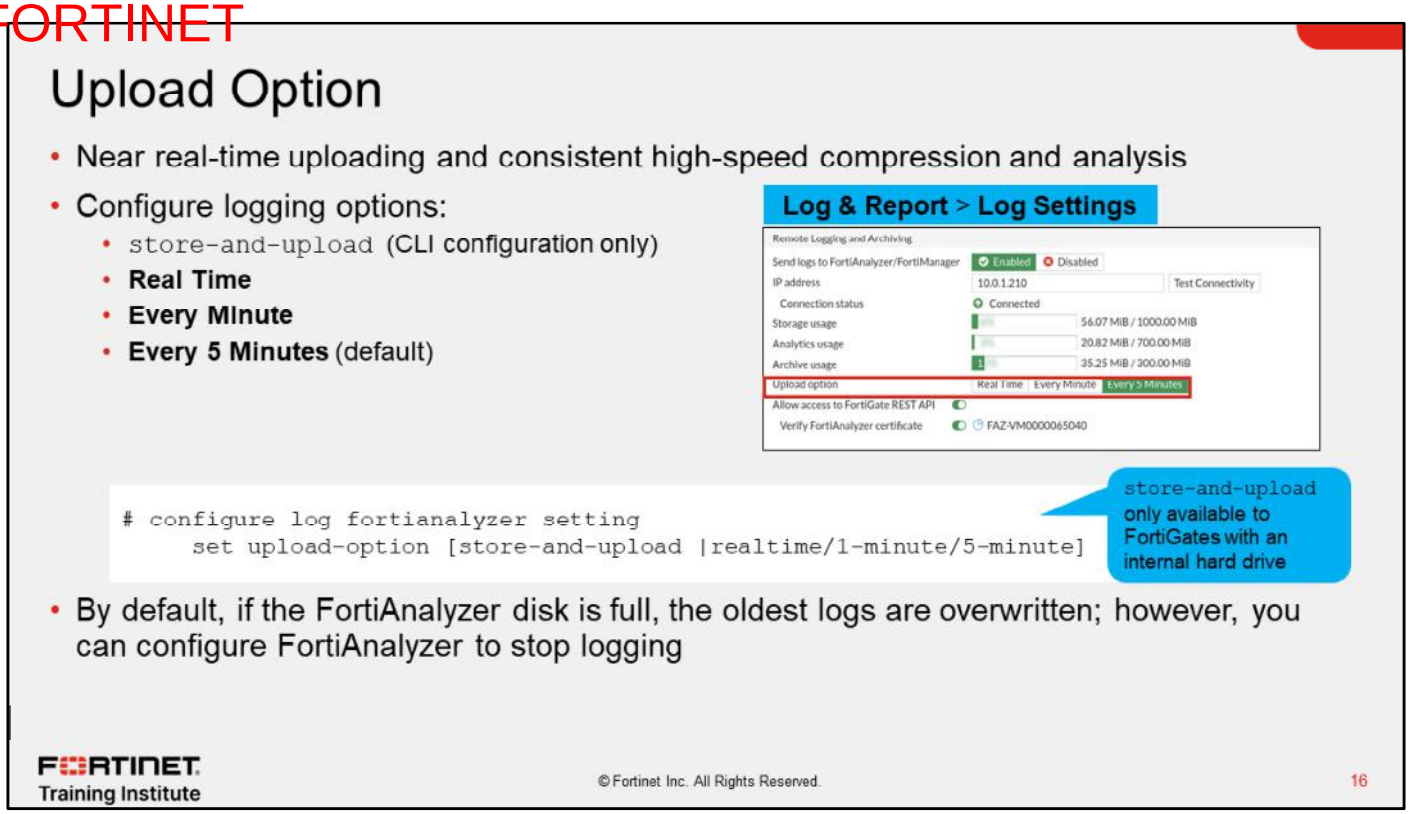

FortiGate allows near real-time uploading and consistent high-speed compression and analysis to FortiAnalyzer and FortiManager.

On the GUI, upload options include **Real Time**, **Every Minute**, and **Every 5 Minutes** (default).

If your FortiGate model includes an internal hard drive, you also have the `store-and-upload` option. This allows you to store logs to disk and then upload to FortiAnalyzer or FortiManager at a scheduled time (usually a low bandwidth time). You can configure the `store-and-upload` option, as well as a schedule, on the CLI only.

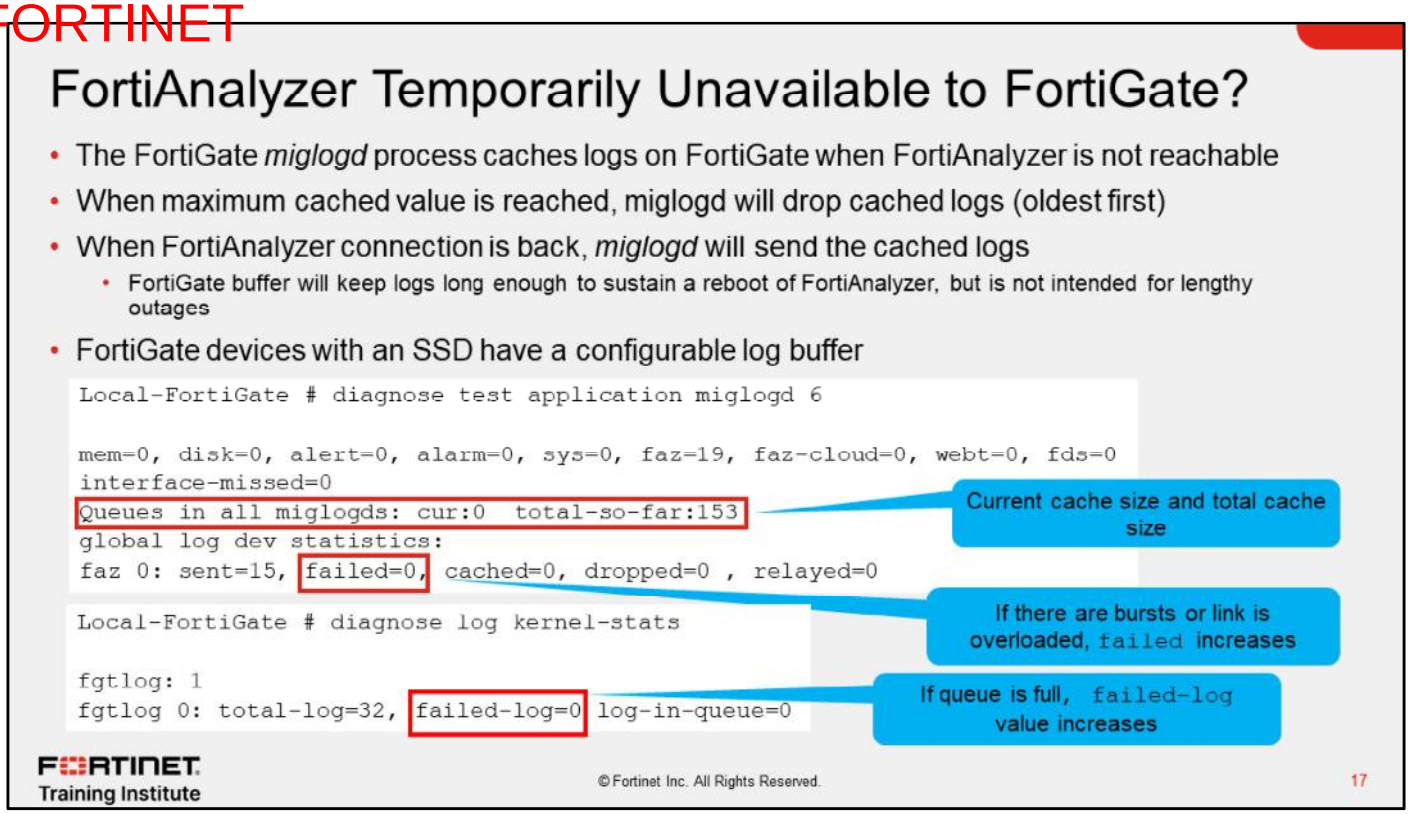

If FortiAnalyzer becomes unavailable to FortiGate for any reason, FortiGate uses its *miglogd* process to cache the logs. There is a maximum value to the cache size, and the *miglogd* process will begin dropping cached logs (oldest first) once this value is reached. When the connection between the two devices is restored, the miglogd process begins to send the cached logs to FortiAnalyzer. Therefore, the FortiGate buffer keeps logs long enough to sustain a reboot of your FortiAnalyzer (if you are upgrading the firmware, for example), but it is not intended for a lengthy FortiAnalyzer outage.

On FortiGate, the CLI command `diagnose test application miglogd 6` displays statistics for the miglogd process, including the total cache size and current cache size.

The CLI command `diagnose log kernel-stats` will show an increase in `failed-log` if the cache is full and needs to drop logs.

FortiGate devices with an SSD disk have a configurable log buffer. When the connection to FortiAnalyzer is unreachable, FortiGate can buffer logs on disk if the memory log buffer is full. The logs queued on the disk buffer can be sent successfully after the connection to FortiAnalyzer is restored.

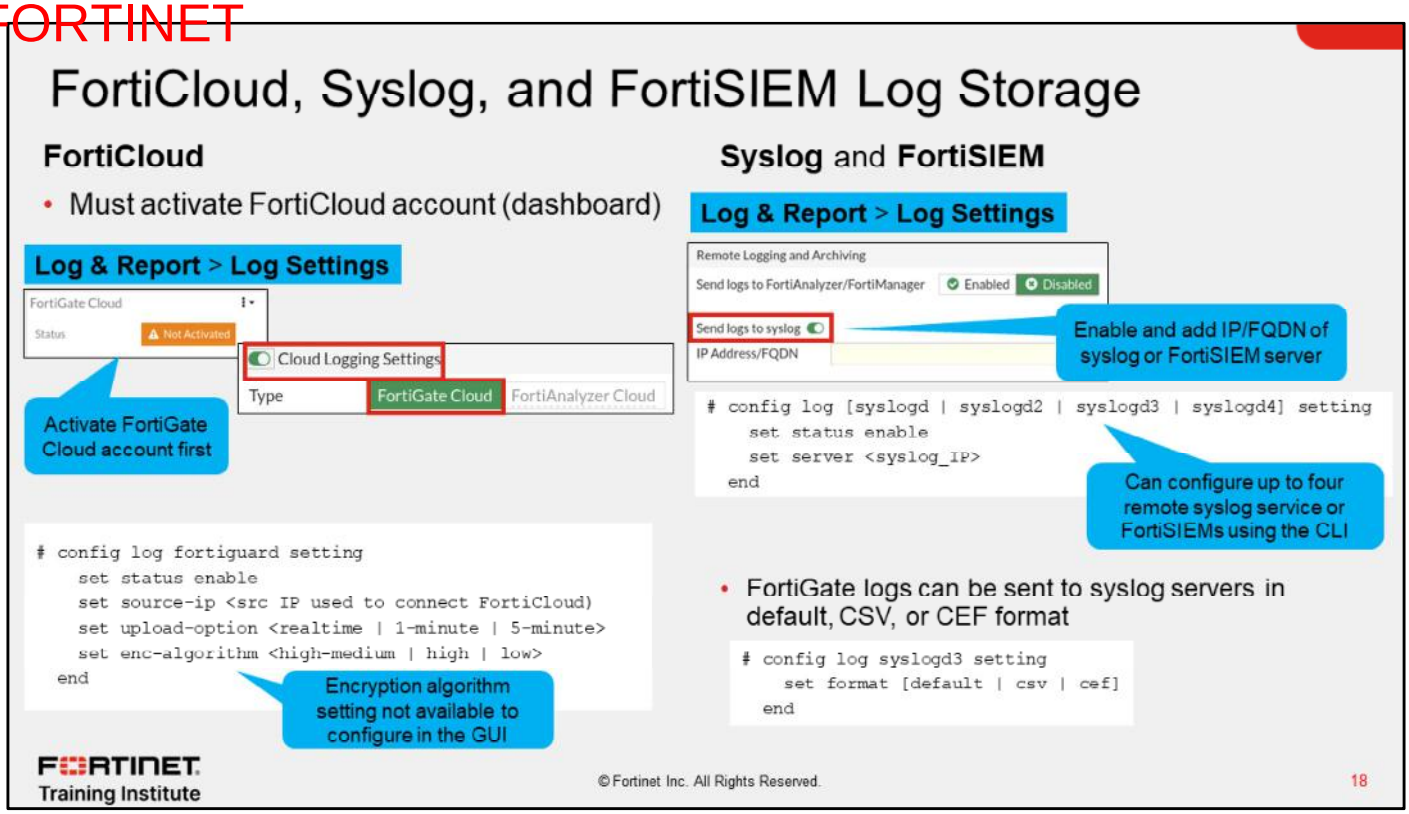

Similar to FortiAnalyzer and FortiManager, you can configure remote logging to FortiCloud on the **Log Settings** page or the CLI. However, you must first activate your FortiCloud account, so FortiGate can communicate with your FortiCloud account. Once complete, you can enable FortiCloud logging and set the upload option. If you want to store your logs to disk first and then upload to FortiCloud, you must specify a schedule. When disk usage is set to WAN optimization (`wanopt`), the store and upload option for logging to FortiCloud is removed.

You can also configure remote logging to syslog and FortiSIEM on the **Log Settings** page or the CLI. You can configure FortiGate to send logs to up to four syslog servers or FortiSIEM devices using the `config log syslogd` CLI command.

FortiGate supports sending logs to syslog in CSV and CEF format, which is an open log management standard that provides interoperability of security-related information between different network devices and applications. CEF data can be collected and aggregated for analysis by enterprise management or Security Information and Event Management (SIEM) systems, such as FortiSIEM. You can configure each syslog server separately to send log messages in CEF or CSV format.

You can configure an individual syslog to use CSV and CEF format using the CLI. The example shown on this slide is for syslogd3. All other syslog settings can be configured as required independently of the log message format, including the server address and transport (UDP or TCP) protocol.

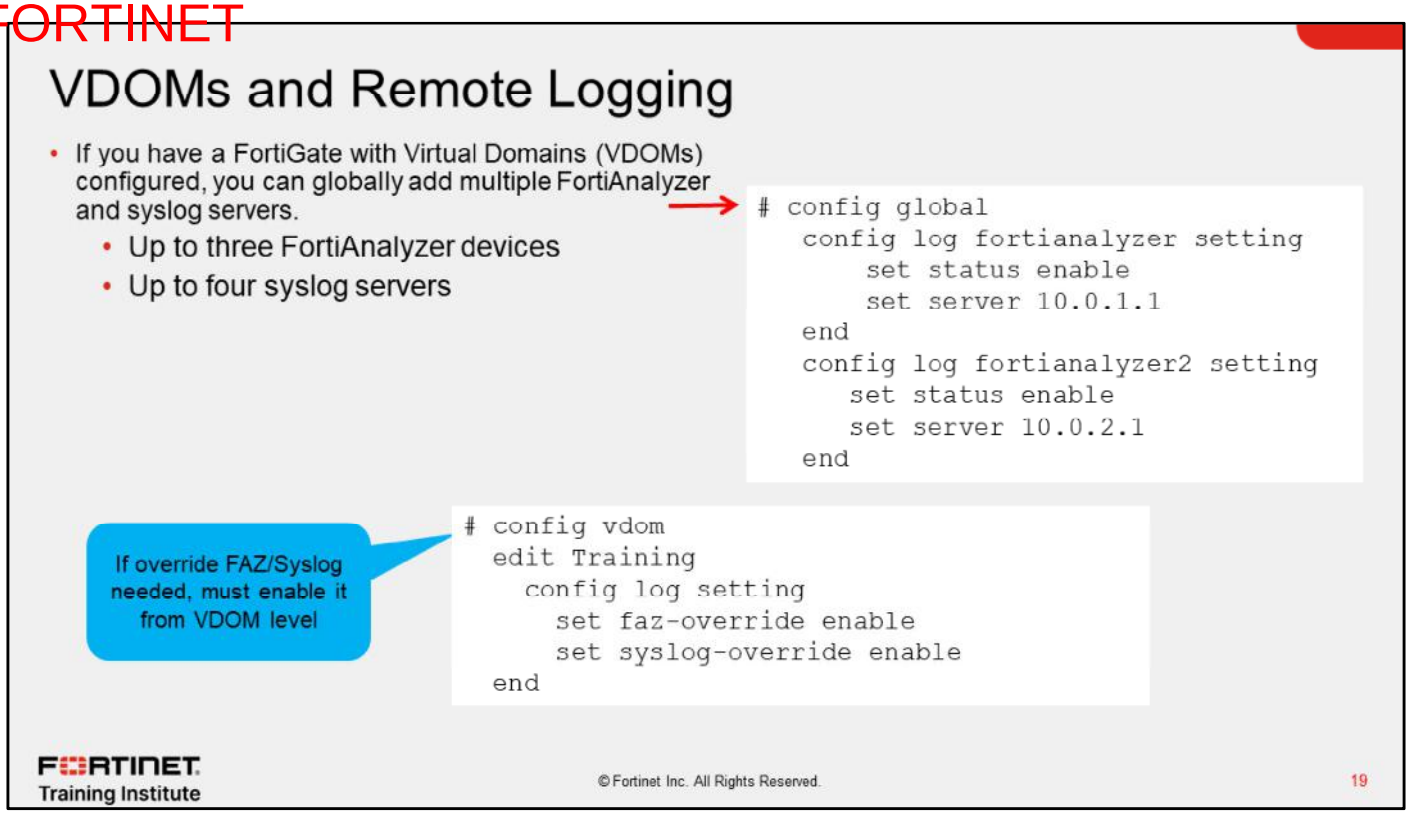

If you have a FortiGate with virtual domains (VDOMs) configured, you can globally add multiple FortiAnalyzer and syslog servers. You can configure up to three FortiAnalyzer devices and up to four syslog servers under global settings.

Log Transmission

- FortiGate uses UDP 514 for log transmission by default

```
config log fortianalyzer setting
    set status enable
    set server "10.0.1.210"
    set serial "FAZ-VM0000065040"
    set enc-algorithm high-medium
    set upload-option realtime
end
```

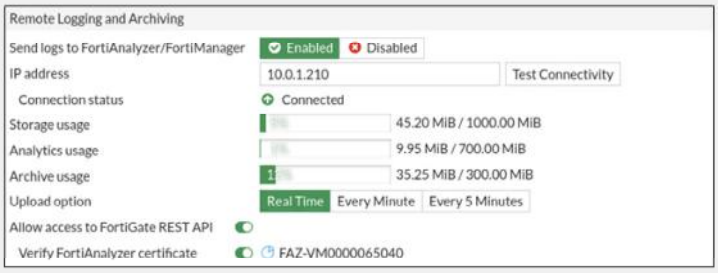

- Log messages are stored on disk and transmitted to FortiAnalyzer as plain text in LZ4 compressed format
  - Reduces disk log size and reduces log transmission time and bandwidth usage

FORTINET Training Institute

20

FortiGate uses UDP port 514 for log transmission by default .

Log messages are stored on disk and transmitted to FortiAnalyzer as plain text in LZ4 compressed format. This reduces disk log size and reduces log transmission time and bandwidth usage.

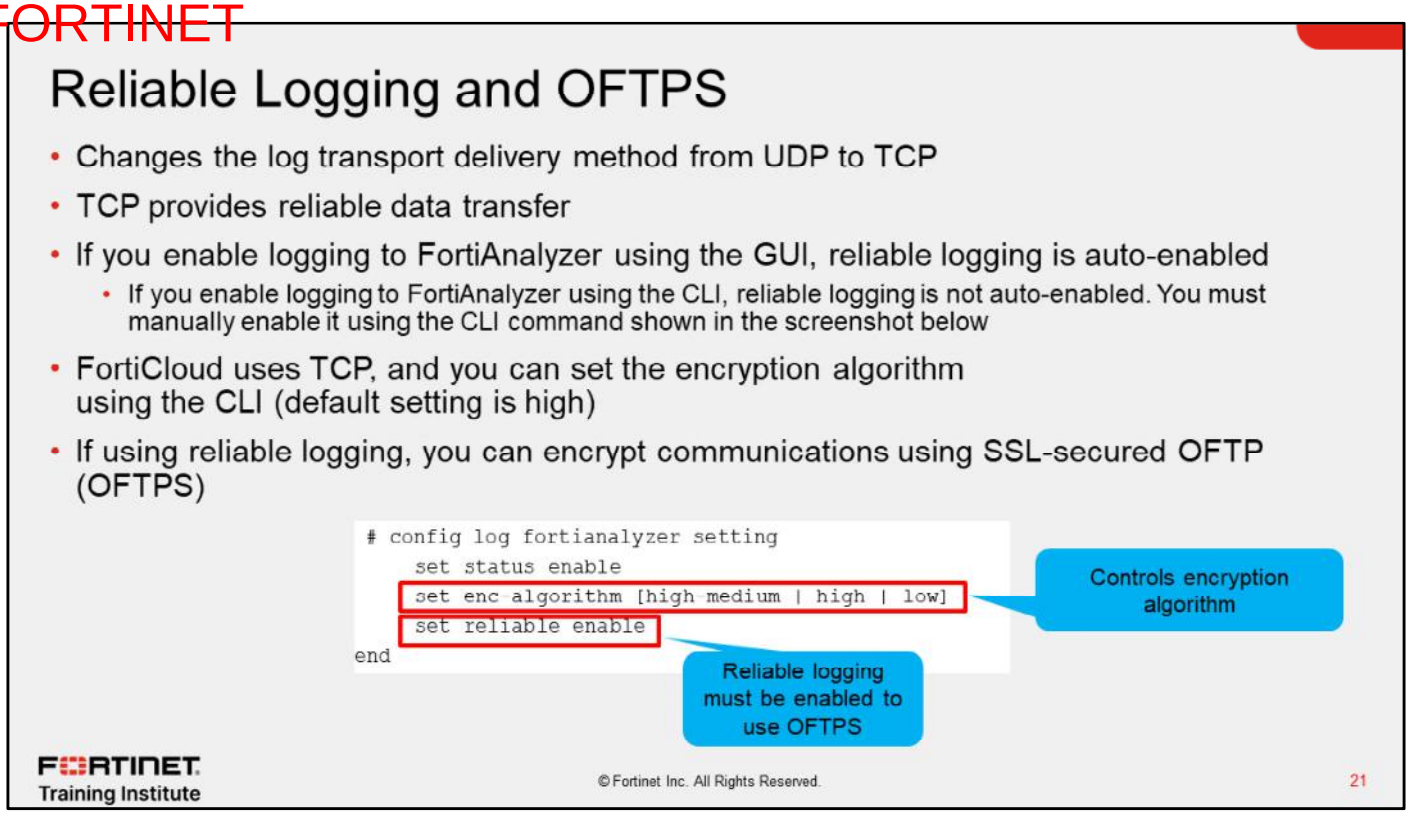

When you enable reliable logging on FortiGate, the log transport delivery method changes from User Datagram Protocol (UDP) to Transmission Control Protocol (TCP). TCP provides reliable data transfer, guaranteeing that the transferred data remains intact and arrives in the same order in which it was sent.

If you enable logging to FortiAnalyzer or FortiManager using the GUI, reliable logging is automatically enabled. If you enable logging using the CLI, you must enable reliable logging using the CLI command shown on this slide.

Logging to FortiCloud uses TCP, and you can set the encryption algorithm using the CLI. The default encryption setting is high.

Optionally, if using reliable logging, you can encrypt communications using SSL-encrypted OFTP traffic, so when a log message is generated, it is safely transmitted across an unsecure network. You can choose the level of SSL protection used by configuring the `enc-algorithm` setting on the CLI.

When both FortiGate and FortiAnalyzer are running version 7.2 or later, and reliable logging is configured, FortiGate keeps logs in a *confirm queue* until it verifies those logs were received by FortiAnalyzer. This is achieved by using sequence numbers (`seq_no`) to track the logs received. FortiOS periodically queries FortiAnalyzer for the latest `seq_no` of the last log received, and clears logs from the *confirm queue* up to the `seq_no`.

Knowledge Check

1. Which storage type is preferred for logging?
   - ✓ A. Remote logging
   - B. Hard drive

2. Which protocol does FortiGate use to send encrypted logs to FortiAnalyzer?
   - ✓ A. OFTPS
   - B. SSL

3. If you enable reliable logging, which transport protocol will FortiGate use?
   - A. UDP
   - ✓ B. TCP

FORTINET Training Institute

22

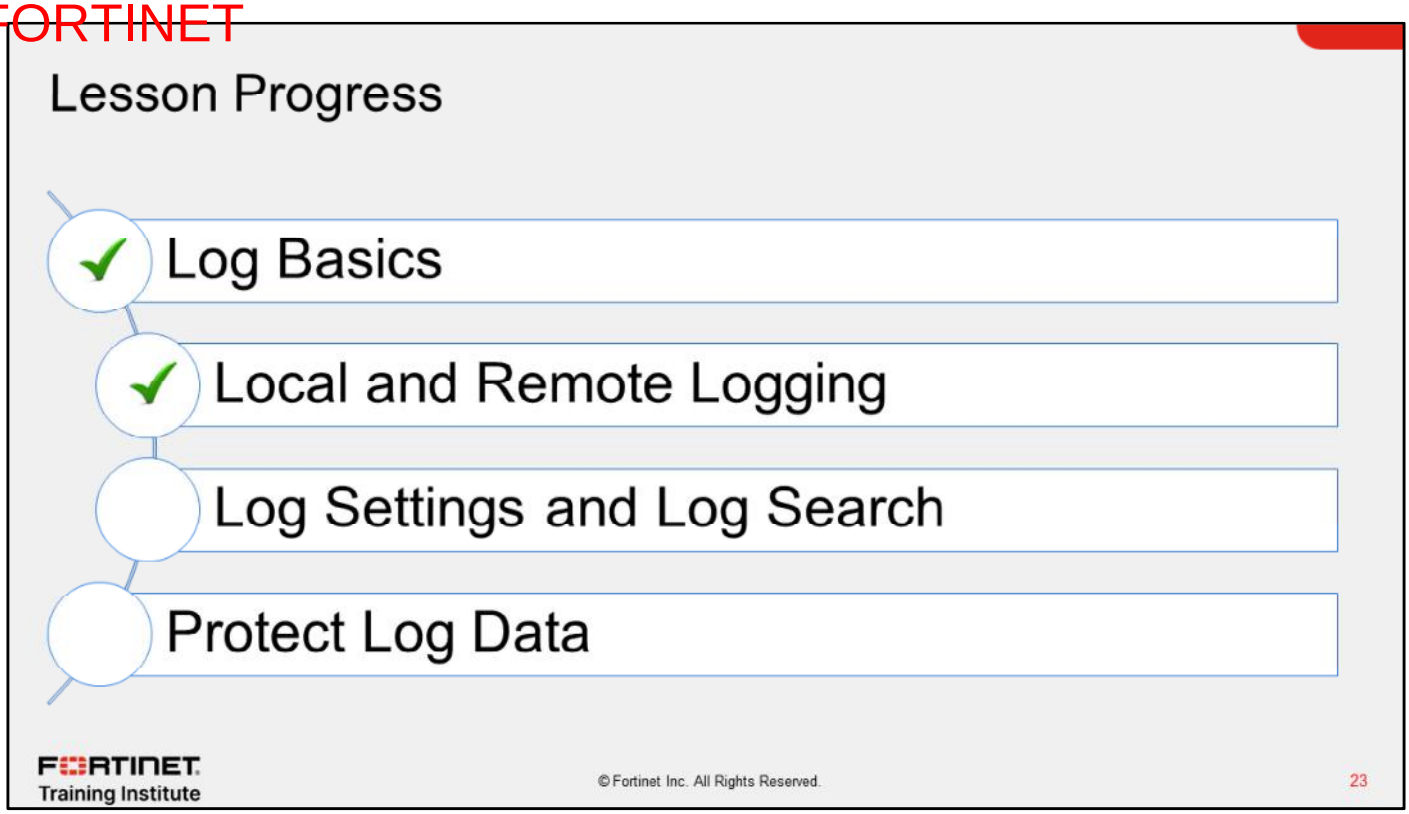

Good job! You now understand remote logging.

Now, you will learn about log settings.

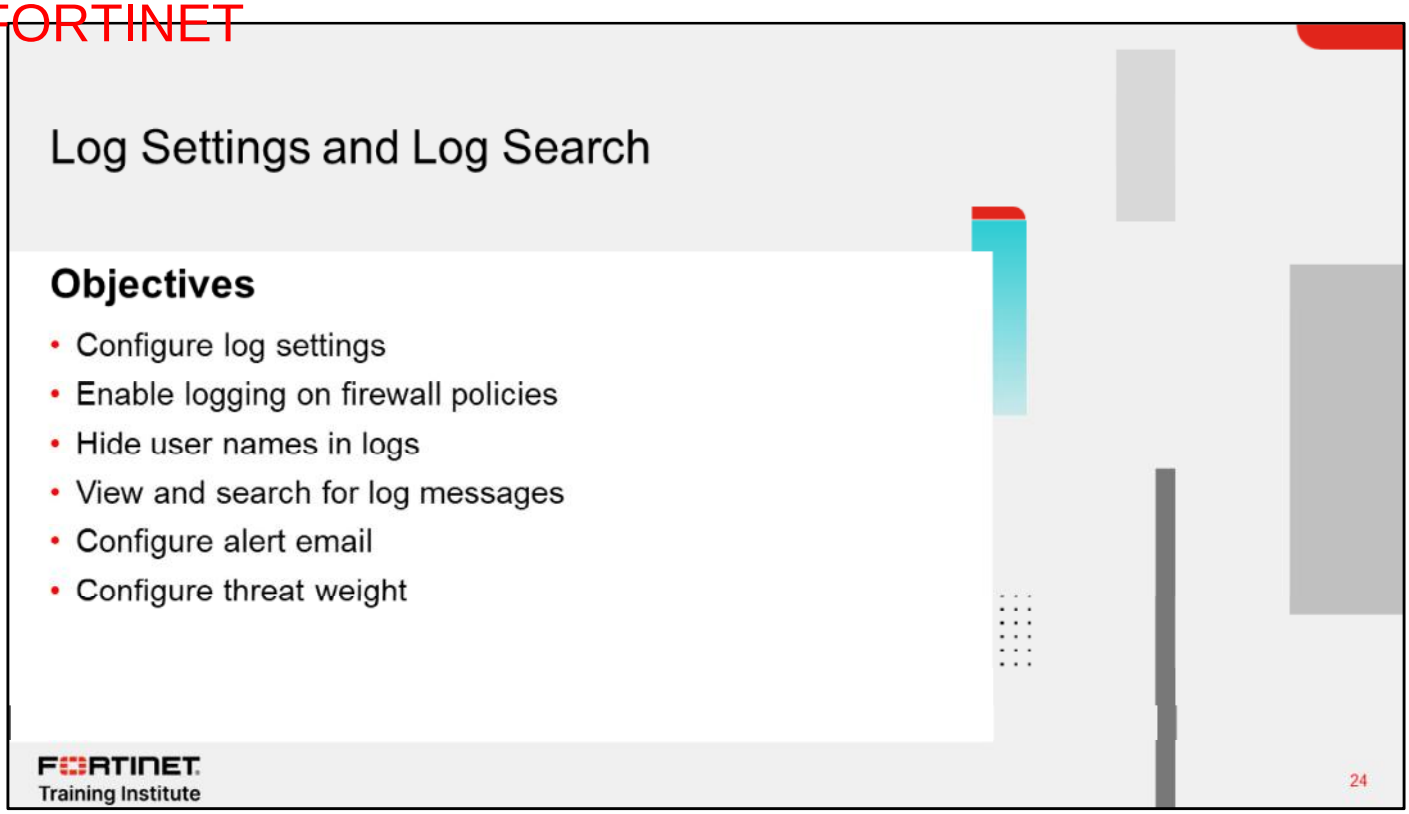

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in log settings, you will be able to successfully enable logging on your FortiGate, and ensure logs are generated on traffic caused by traffic passing through your firewall policies.

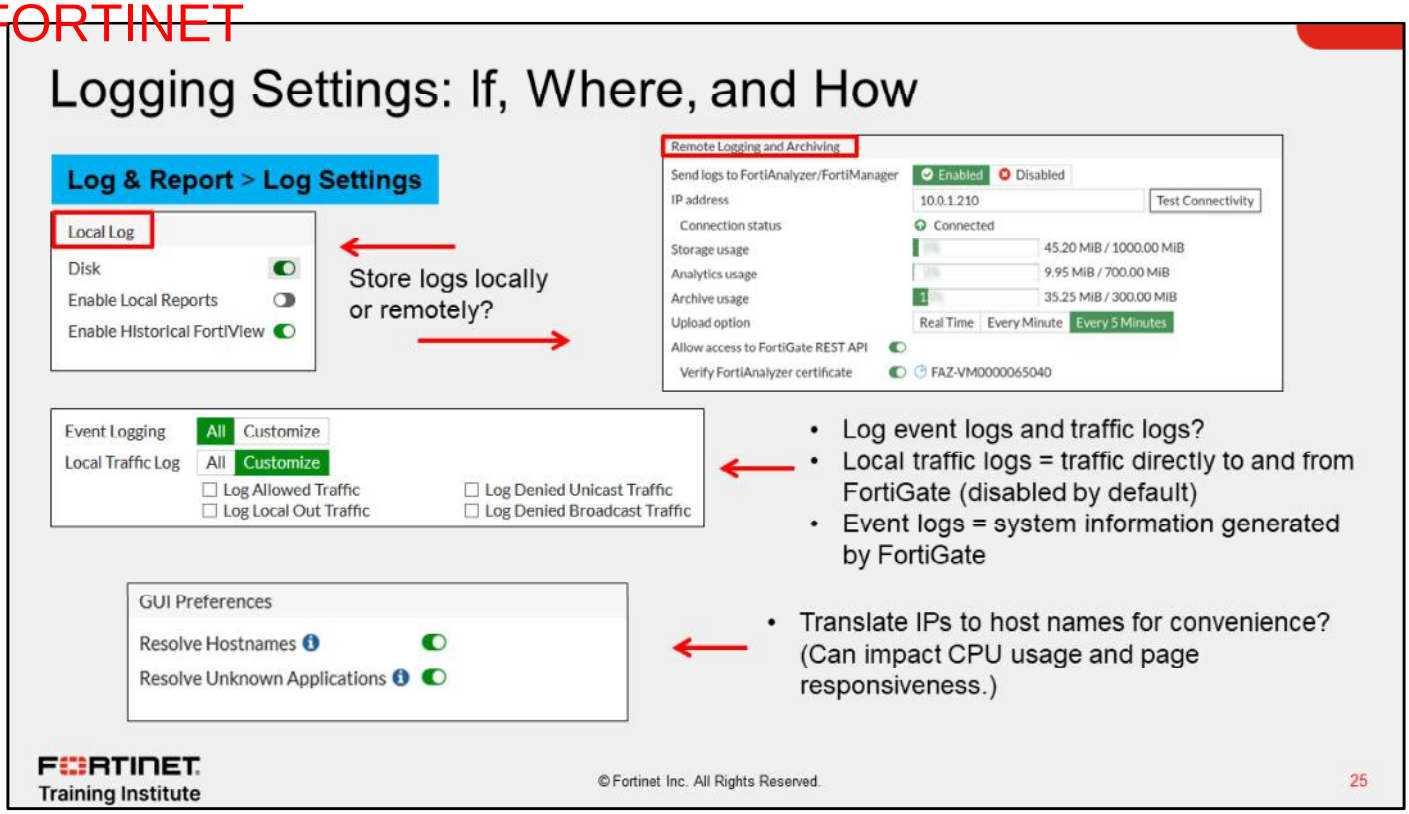

The **Log Settings** page allows you to decide if, where, and how a log is stored.

As previously discussed, you must configure whether to store logs locally on your FortiGate disk, or remotely to an external device, such as FortiAnalyzer.

You must also configure what event logs and local traffic logs to capture. By default, this option is disabled because of the large number of logs they can generate.

Event logs provide all of the system information generated by FortiGate, such as administrator logins, configuration changes made by administrators, user activity, and daily operations of the device—they are not directly caused by traffic passing through firewall policies. The event logs you choose to enable depend on what features you are implementing and what information you need to get from the logs.

The **Resolve Hostnames** feature resolves IP addresses to host names. This requires FortiGate to perform reverse DNS lookups for all IP addresses. If your DNS server is not available, or is slow to reply, it can impact your ability to look through the logs, because the requests will time out.

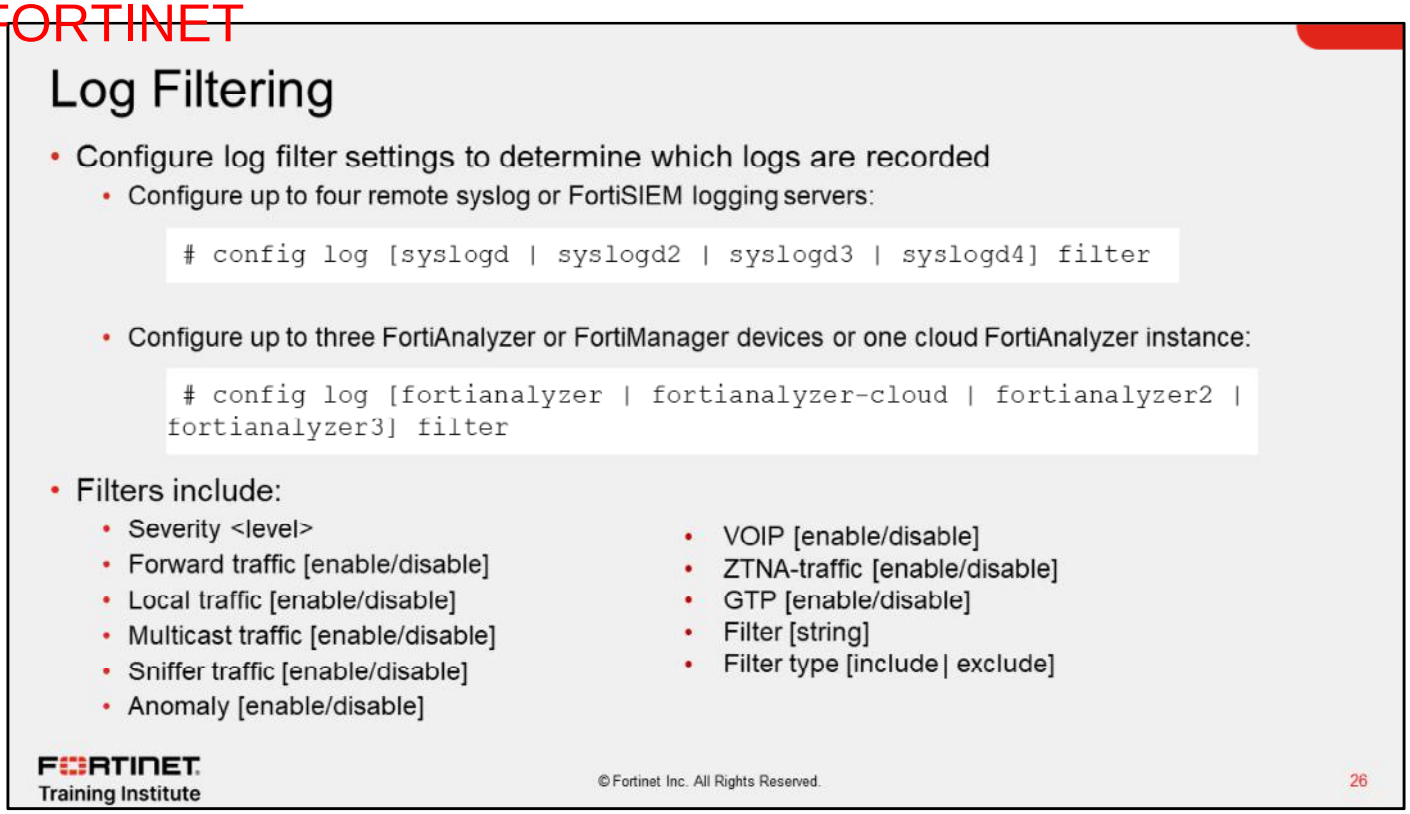

While you use the log settings on the GUI to configure which event logs and local traffic logs to capture, you can set more granular options using the CLI.

You can configure FortiGate to send logs to external servers. You can control which logs are sent to each of these devices separately, using the command `config log syslogd filter` for remote syslog or FortiSIEM, and the command `config log fortianalyzer filter` for FortiAnalyzer or FortiManager devices.

In this way, you can set devices to different logging levels and/or send only certain types of logs to one device and other types (or all logs) to others. For example, you can send all logs at information level and above to `fortianalyzer`, alert level and above to `fortianalyzer2`, and only traffic logs to `fortianalyzer3`.

For example, the following commands configure the log filter for the first syslog server to include only logs related to traffic directly to and from the FortiGate management IP addresses, with a severity level of *critical* or higher*:*

```
#config log syslogd filter
#(filter) set severity critical
#(filter) set local-traffic enable
```

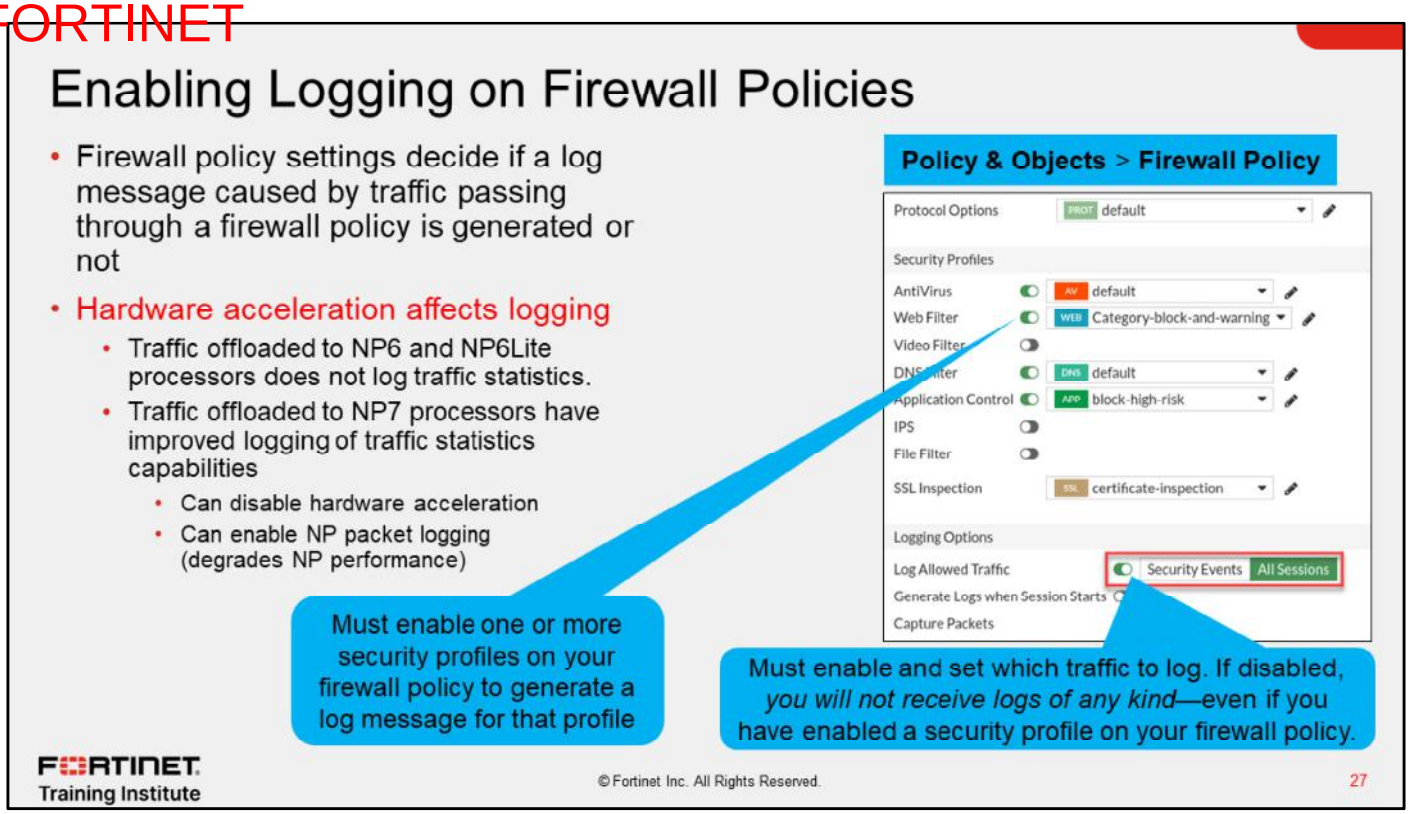

After you configure all logging settings, you can enable logging on your firewall policies. Only when enabled on a firewall policy can a log message—caused by traffic passing through that firewall policy—generate.

Generally, if you configure FortiGate to inspect traffic, you should also enable logging for that security feature to help you track and debug your traffic flow. Except for violations that you consider to be low in severity, you'll want to know if FortiGate is blocking attacks. Most attacks don't result in a security breach on the first try. A proactive approach, when you notice a persistent attacker whose methods seem to be evolving, can avoid a security breach. To get early warnings like this, enable logging for your security profiles.

To enable logging on traffic passing through a firewall policy, you must do the following:

1. Enable the desired security profile(s) on your firewall policy.
2. Enable **Log Allowed Traffic** on that firewall policy. This setting is vital. If disabled, you will not receive logs of any kind—even if you have enabled a security profile on your firewall policy. You can choose to log only security events, or log all sessions:
    - **Security Events**: If enabled (along with one or more security profiles), security log events appear in the forward traffic log and security log. A forward traffic log generates for packets causing a security event.
    - **All Sessions**: If enabled, a forward traffic log generates for every single session. If one or more security profiles are also enabled, security log events appear in the forward traffic log and security log.

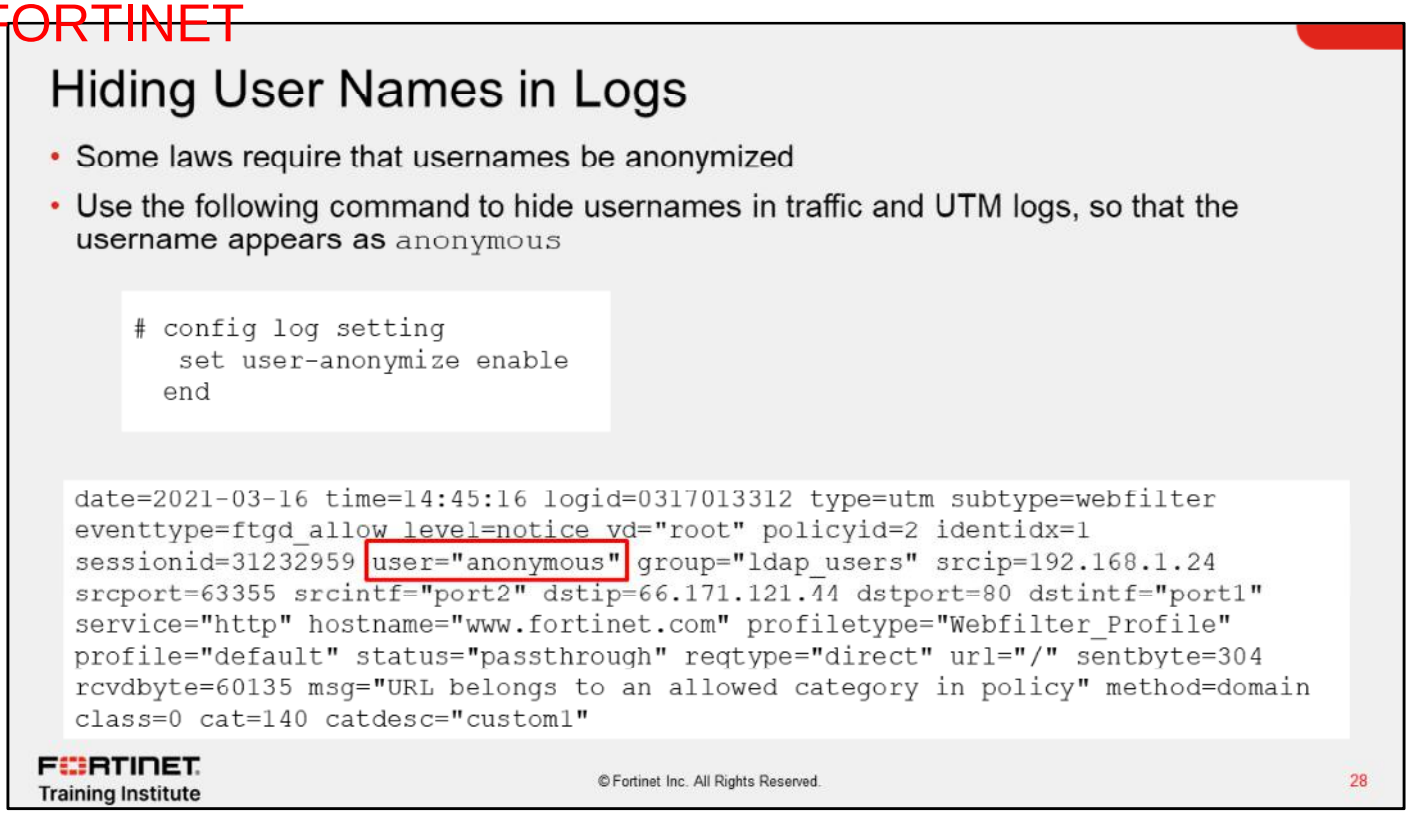

On FortiGate, you can hide usernames in traffic logs and UTM logs, so that the username appears as `anonymous`. This is useful, because some countries do not permit non-anonymized logging.

To anonymize usernames, use the `set user-anonymize enable` CLI command.

It is assumed that logging is enabled in firewall policies and security profiles, and that identity-based policies are configured on FortiGate.

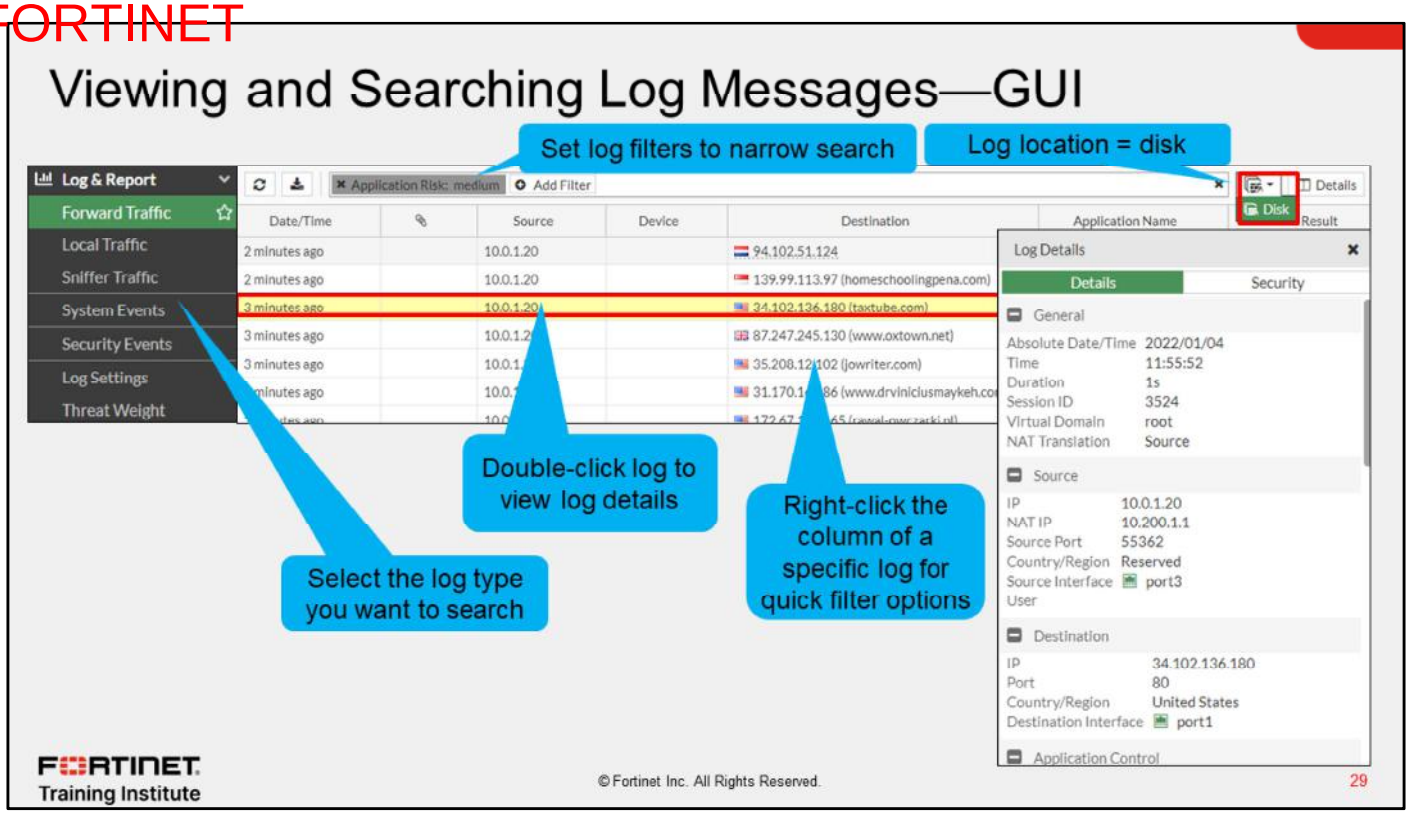

You can access your logs on the GUI in the **Log & Report** menu.

Select the type of log you want to view, such as **Forward Traffic**. Logs on the GUI appear in a formatted table view. The formatted view is easier to read than the raw view, and it enables you to filter information when viewing log messages. To view the log details, select the log in the table. The log details then appear in the **Log Details** pane on the right side of the window.

If archiving is enabled on security profiles that support it (such as DLP), archived information appears in the **Log Details** pane in the **Archived Data** section. Archived logs are also recorded when using FortiAnalyzer or FortiCloud.

If you configure FortiGate to log to multiple locations, you can change the log display location in this section. In the example shown on this slide, the log location is set to **Disk**. If logging to a syslog, you must view logs on the syslog instead.

To navigate the logs more efficiently, you can set up log filters. The more information you specify in the filter, the easier it is to find the precise log entry. Filters are configurable for each column of log data on the display. Click **Add Filter** to select the filter in the drop-down list that appears. If you see data that you want to filter on in a log that is already in the table, you can right-click that data to select the quick filter option. For example, if you see an antivirus log in the table with a specific botnet name, right-click the botnet name in the table, and a quick filter option opens, allowing you to filter on all logs with that botnet name.

By default, the most common columns are shown and less common columns are hidden. To add columns, right-click any column field and, in the pop-up menu that opens, select the column in the **Available Columns** section.

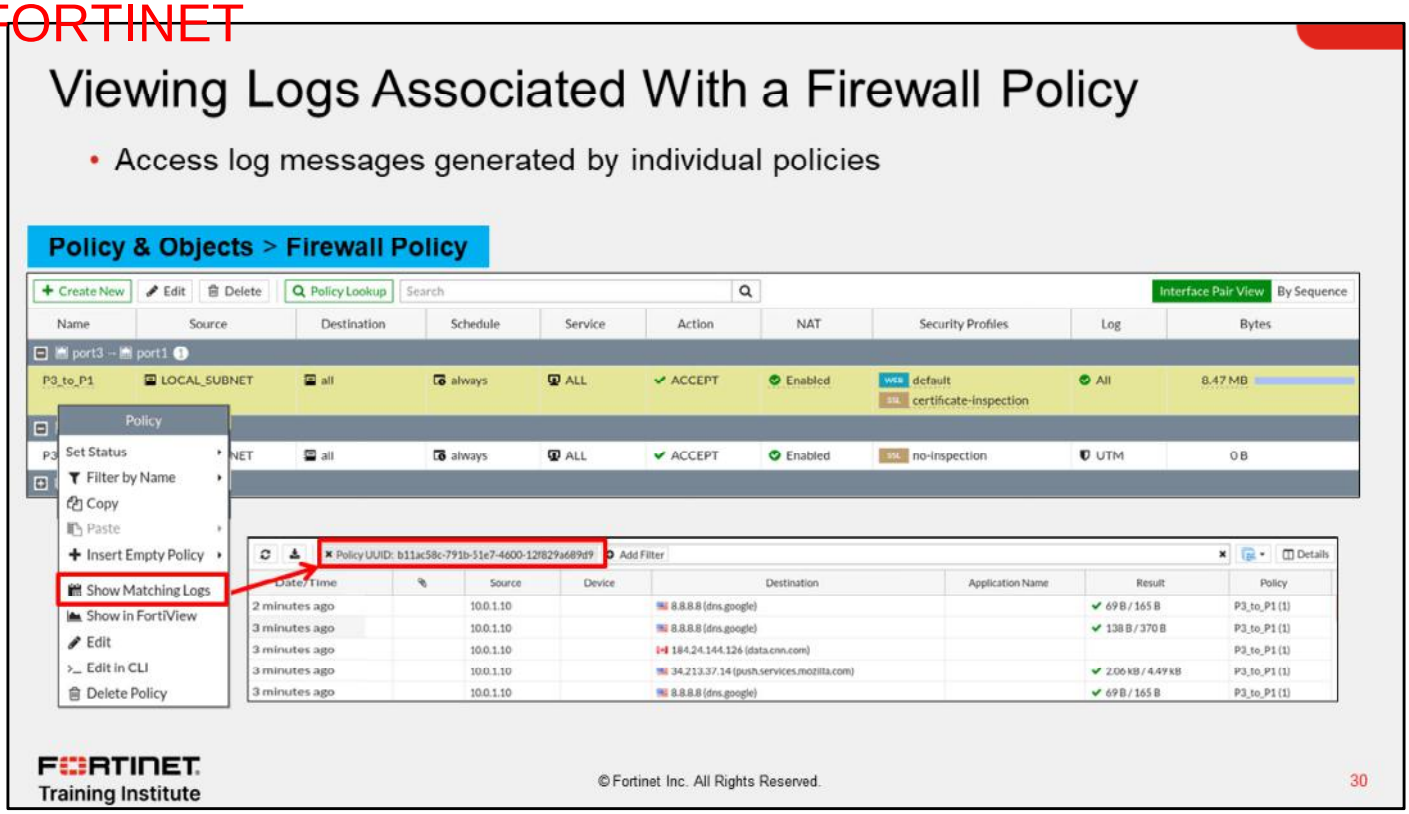

You can also access log messages generated by individual policies. Right-click the policy for which you want to view all associated logs and, in the pop-up menu, select **Show Matching Logs**. FortiGate takes you to the **Forward Traffic** page where a filter is automatically set based on the policy UUID.

You are not restricted from viewing log messages on the GUI. You can also view log messages on the CLI, using the `execute log display` command. This command allows you to see specific log messages that you already configured within the `execute log filter` command. The `execute log filter` command configures what log messages you will see, how many log messages you can view at one time (a maximum of 1000 lines of log messages), and the type of log messages you can view.

Logs appear in the raw format view. The raw format displays logs as they appear within the log file.

Similar to the GUI, if you have configured either a syslog or SIEM server, you will not be able to view log messages on the CLI.

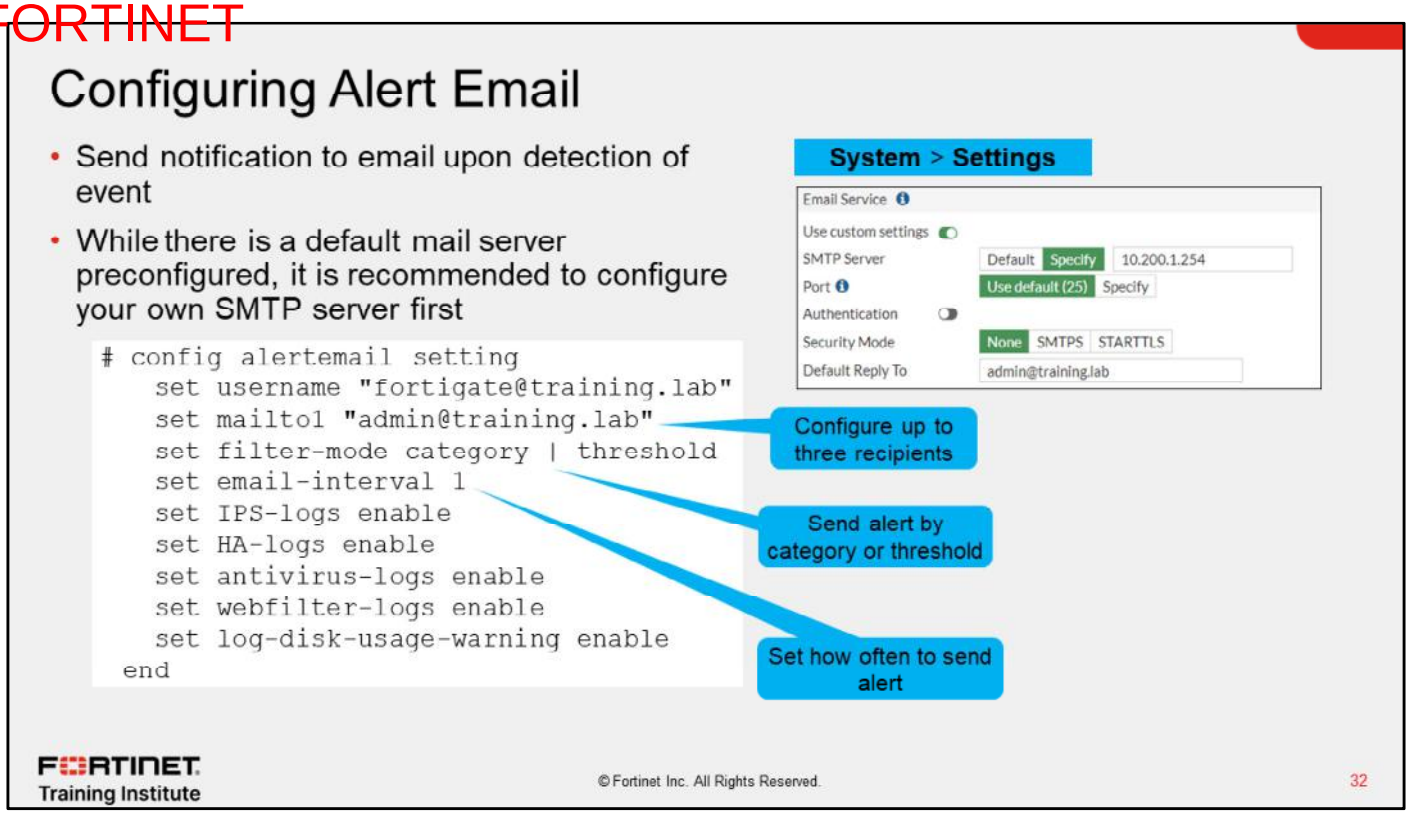

Because you can't always be physically watching the logs on the device, you can monitor events by setting up alert email. Alert emails provide an efficient and direct method of notifying an administrator of events.

Before you configure alert email, you should configure your own SMTP server on your FortiGate first. The FortiGate has an SMTP server preconfigured, but it is recommended that you use your internal email server if you have one.

You can configure alert emails using the CLI. You can trigger alert emails based on event (such as any time an intrusion is detected or the web filter blocked traffic), or on minimum log severity level (such as all logs at the Alert level or above). You can configure up to three recipients.

Configuring Threat Weight

- Prioritize solving the most relevant issues by configuring severity levels for IPS signatures, web categories, and applications with a threat weight
- Set risk level values for low, medium, high, and critical

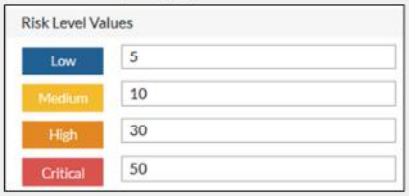

- View detected threats from **Dashboard > Security**

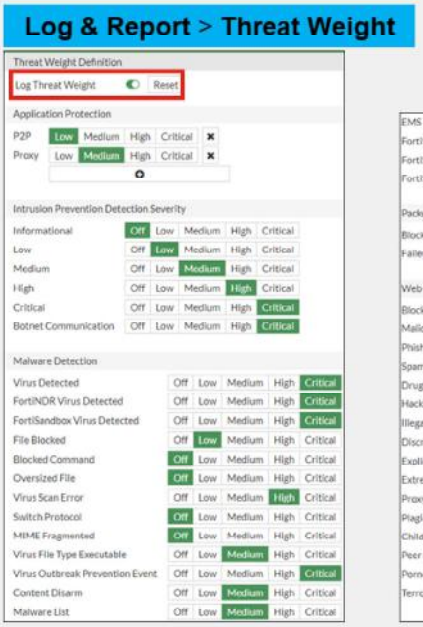

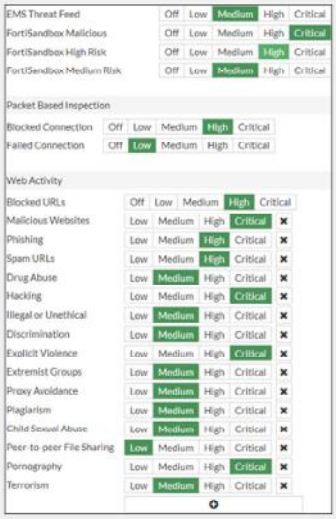

FORTINET Training Institute

33

In order to prioritize solving the most relevant issues easily, you can configure severity levels for IPS signatures, web categories, and applications that are associated with a threat weight (or score).

On the **Threat Weight** page, you can apply a risk value of either low, medium, high, or critical to each category-based item. Each of these levels includes a threat weight. By default, low = 5, medium = 10, high = 30, and critical = 50. You can adjust these threat weights based on your organizational requirements.

After threat weight is configured, you can view all detected threats on the **Security** page. You can also search for logs by filtering based on their threat score.

Note that threat weight is for informational purposes only. FortiGate will not take any action based on threat weight.

Knowledge Check

1. In your firewall policy, which setting must you enable to generate logs on traffic sent through that firewall policy?
   - ✓ A. Log Allowed Traffic
   - B. Event Logging

2. With email alerts, you can trigger alert emails based on _____ or log severity level.
   - ✓ A. event
   - B. threat weight

FORTINET Training Institute

34

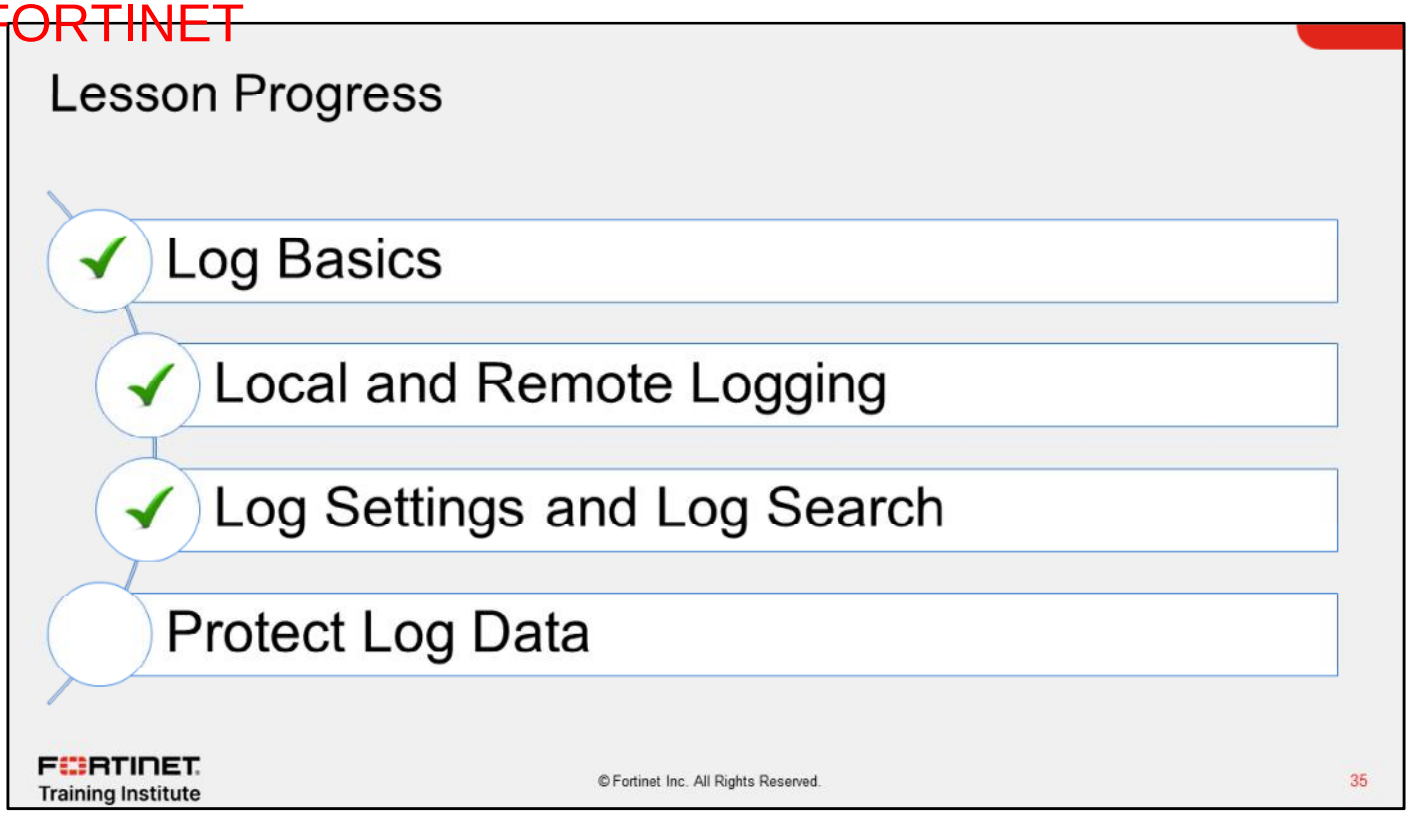

Good job! You now understand how to troubleshoot communication issues.

Now, you will learn how you can protect your log data.

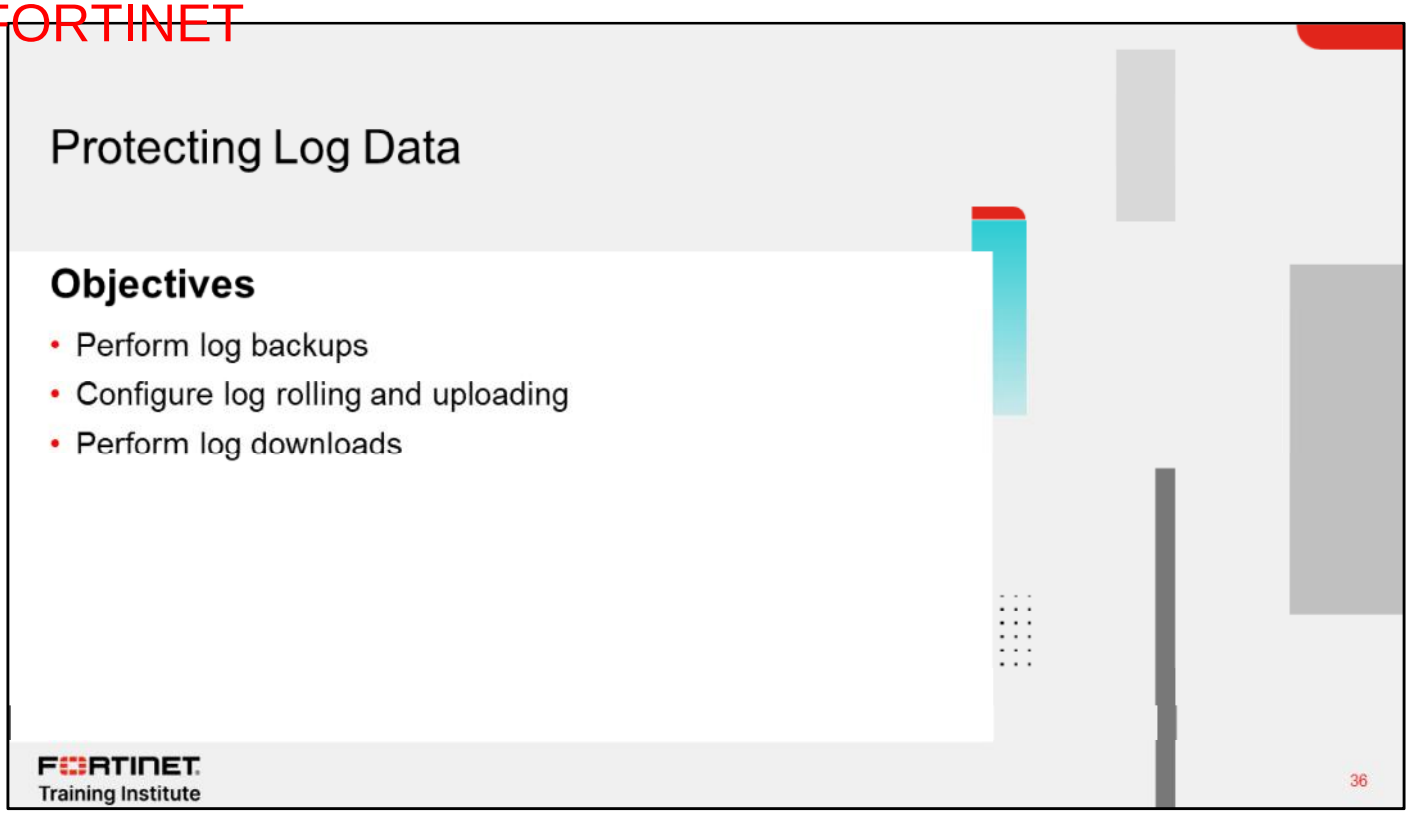

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in using various methods to protect your logs, you will be able to meet organizational or legal requirements for logs.

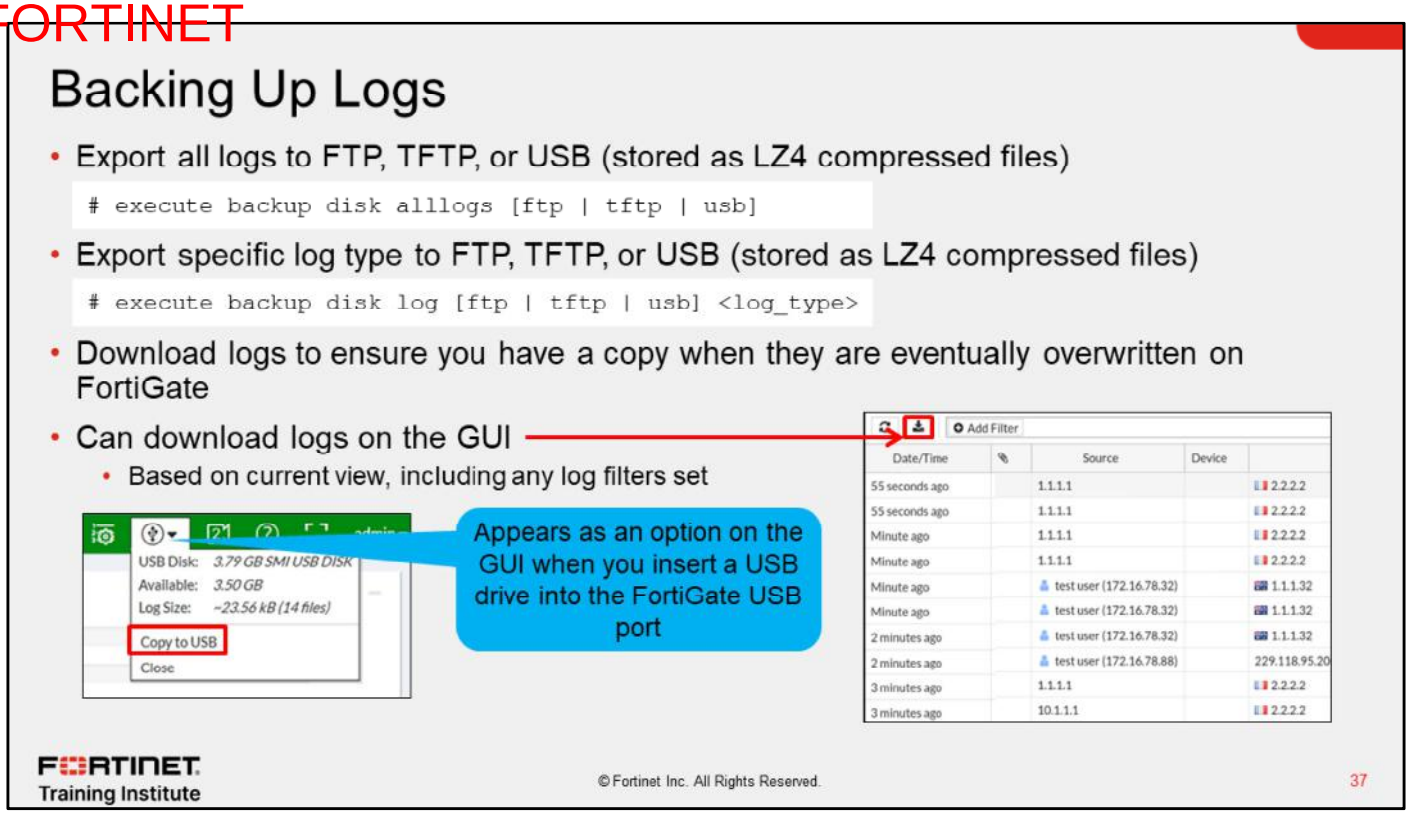

You can also protect your log data by performing log backups. A backup operation copies log files from the database to a specified location.

The `execute backup disk alllogs` command backs up all logs to FTP, TFTP, or USB, while the `execute backup disk log <log type>` command backs up specific log types (such as web filter or IPS) to FTP, TFTP, or USB. These logs are stored in LZ4 format.

You can also use the GUI to back up logs to a USB drive, or to your computer disk. This ensures that you still have a copy when the originals are eventually overwritten on FortiGate.

You can download logs by clicking the download icon on the associated log type page (for example, **Forward Traffic** or **Web Filter**). This downloads only the logs in the results table—not all logs on disk. As such, you can add log filters if you want to download only a subset of logs. When you download log messages from the GUI, you are downloading log messages in raw format.

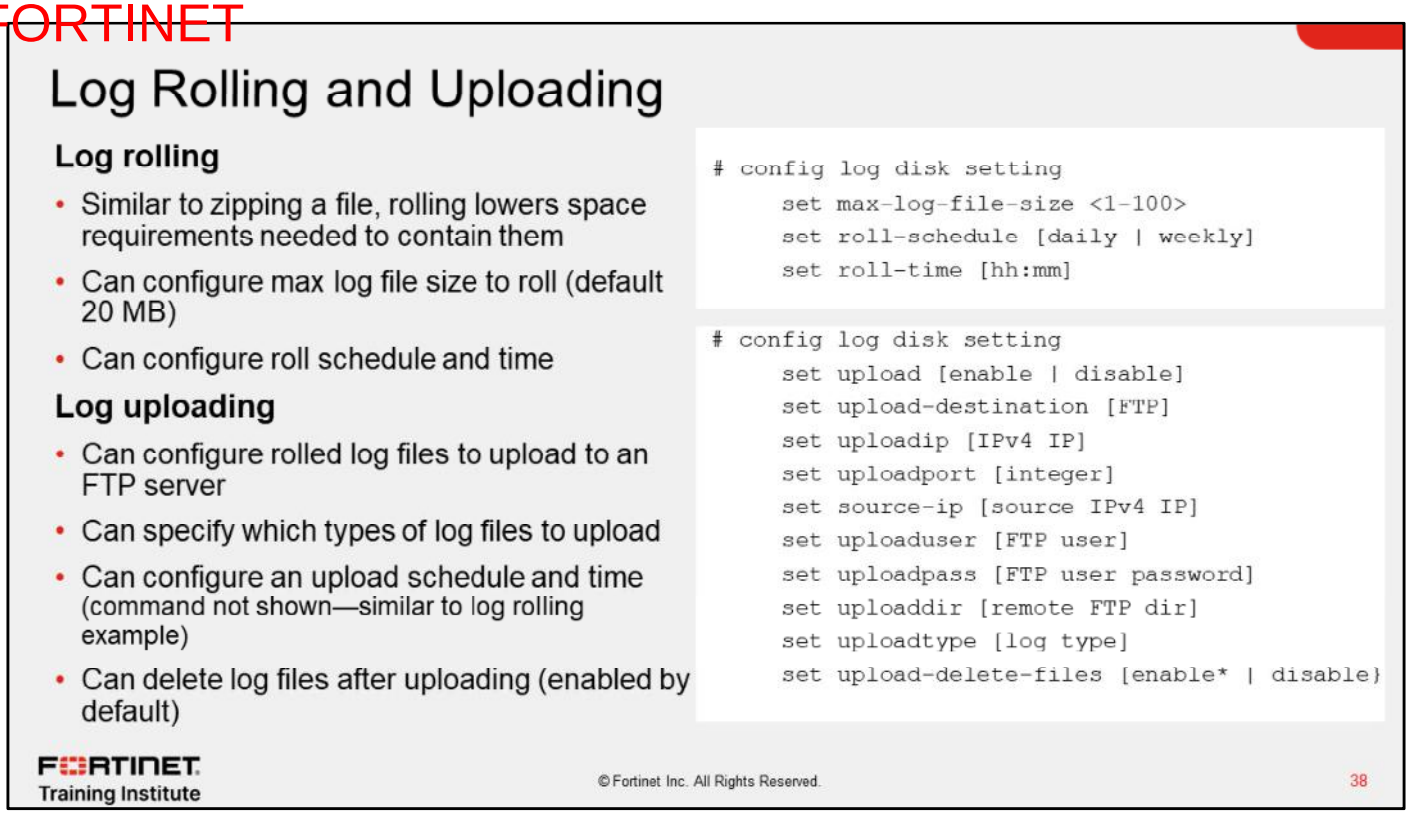

Using the `config log disk setting` command, you can configure logs to roll (which is similar to zipping a file) to lower the space requirements needed to contain them so they don't get overwritten. By default, logs roll when they reach 20 MB in size. You can also configure a roll schedule and time.

Using the same CLI command, you can also configure rolled logs to upload to an FTP server to save disk space. You can configure which types of log files to upload, when, and whether to delete files after uploading.

Knowledge Check

1. What happens when logs roll?
   - ✓A. It lowers the space requirements needed to contain those logs.
   - B. They are uploaded to an FTP server.

2. When you download logs on the GUI, _____
   - A. all logs in the SQL database are downloaded.
   - ✓B. only your current view, including any filters set, are downloaded.

FORTINET Training Institute

39

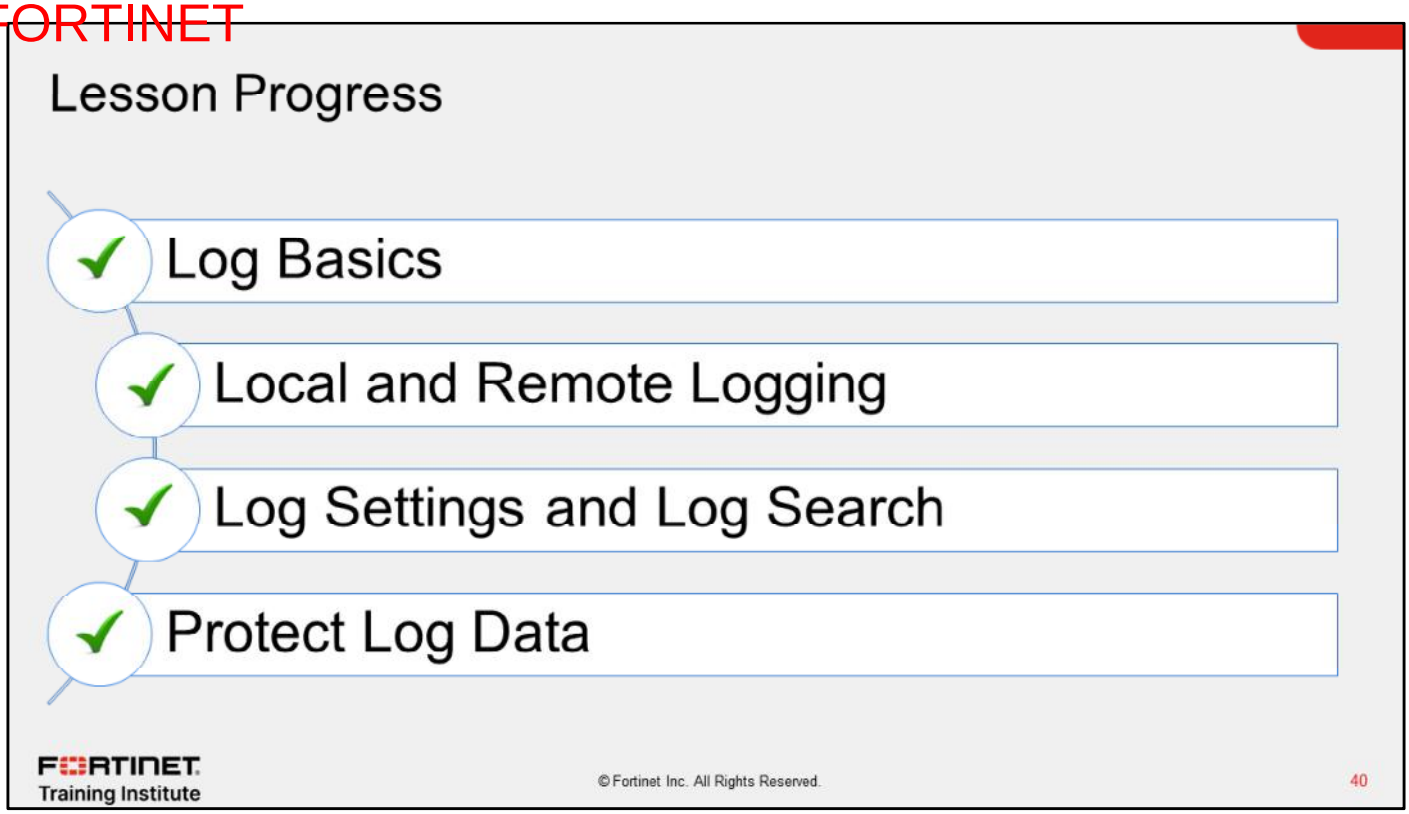

Congratulations! You have completed this lesson. Now, you will review the topics that you covered in this lesson.

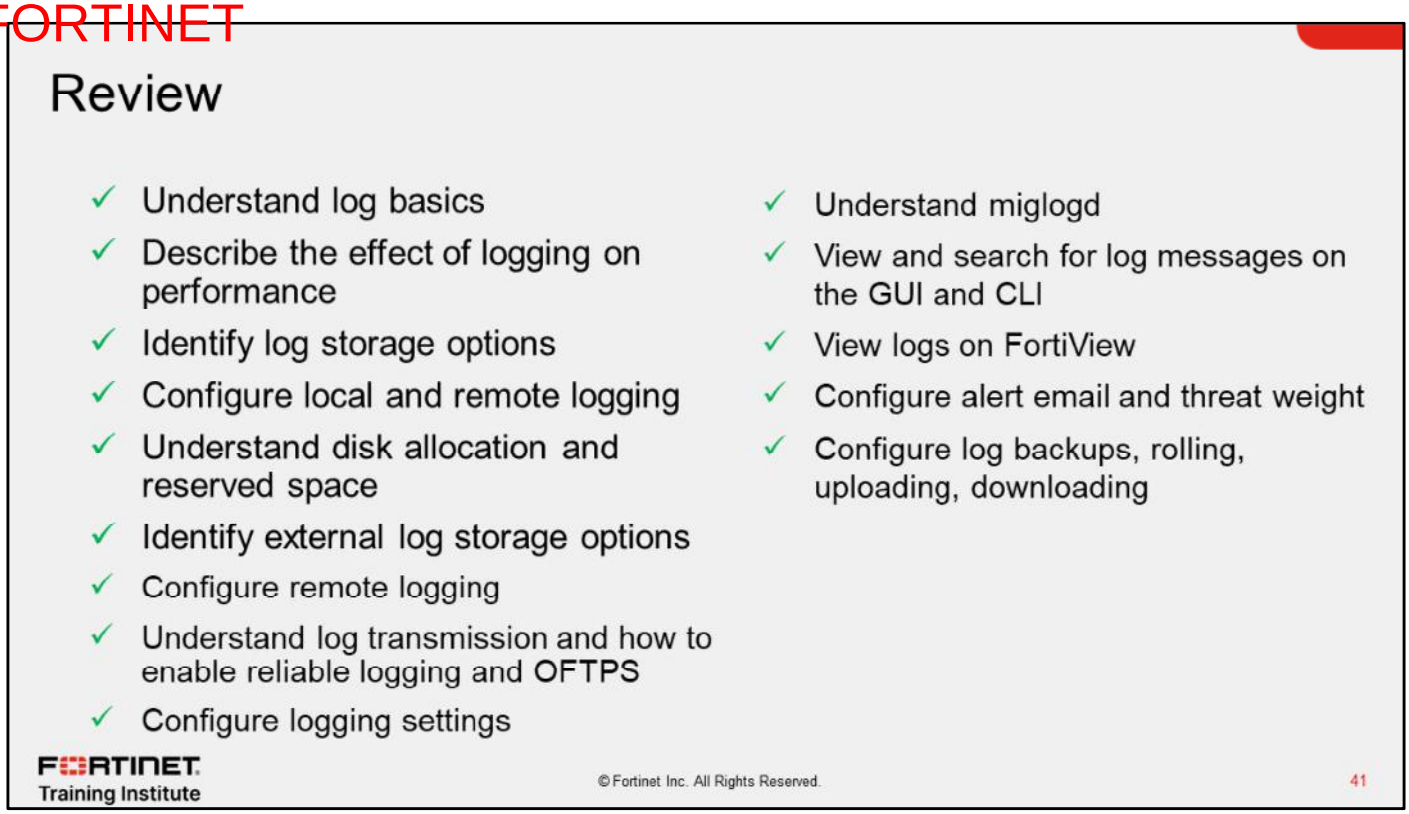

This slide shows the topics that you covered in this lesson.

By mastering the topics covered in this lesson, you learned to configure local and remote logging, view logs, search logs, and protect your log data.

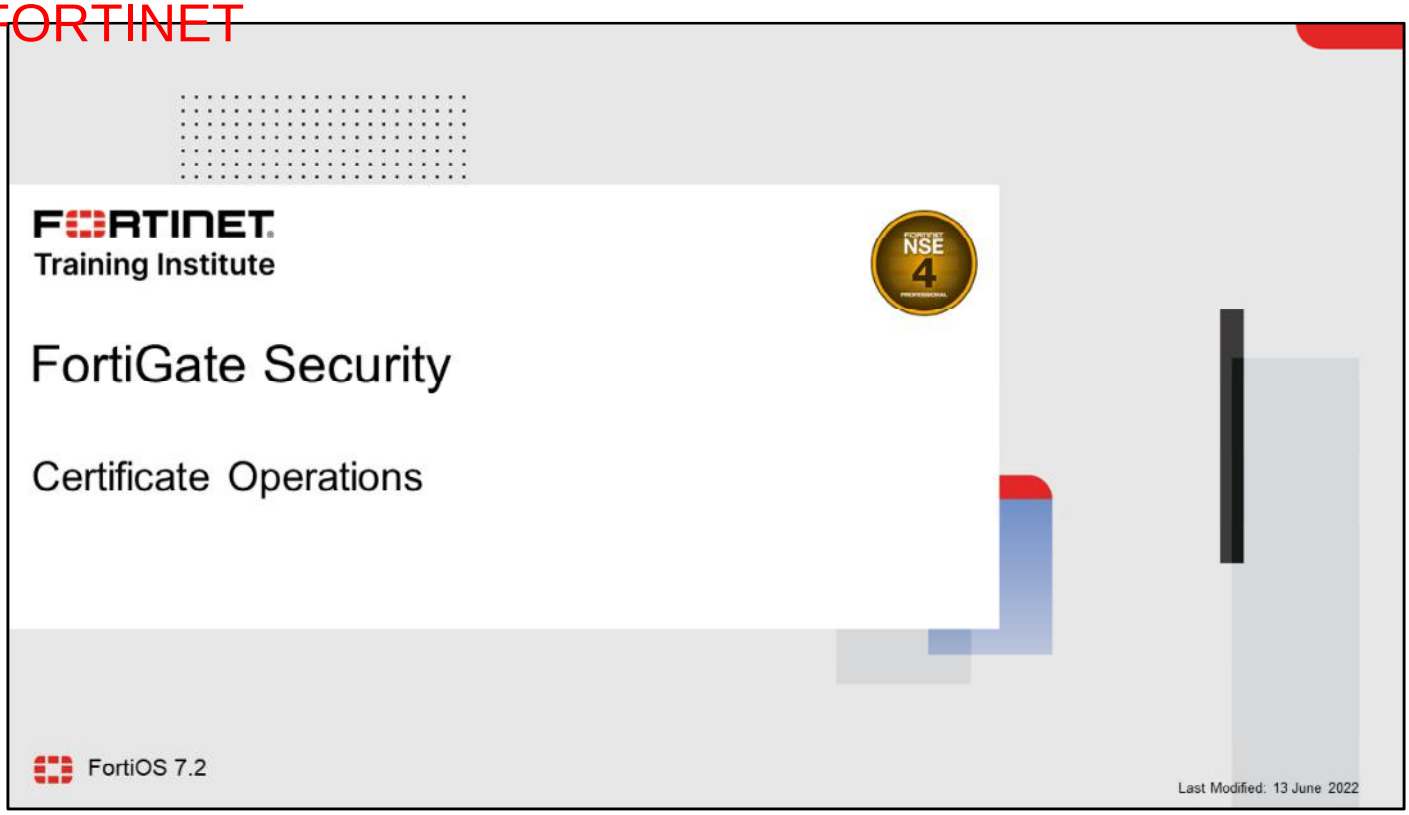

In this lesson, you will learn why FortiGate uses digital certificates and how to configure FortiGate to use certificates (including to inspect the contents of encrypted traffic).

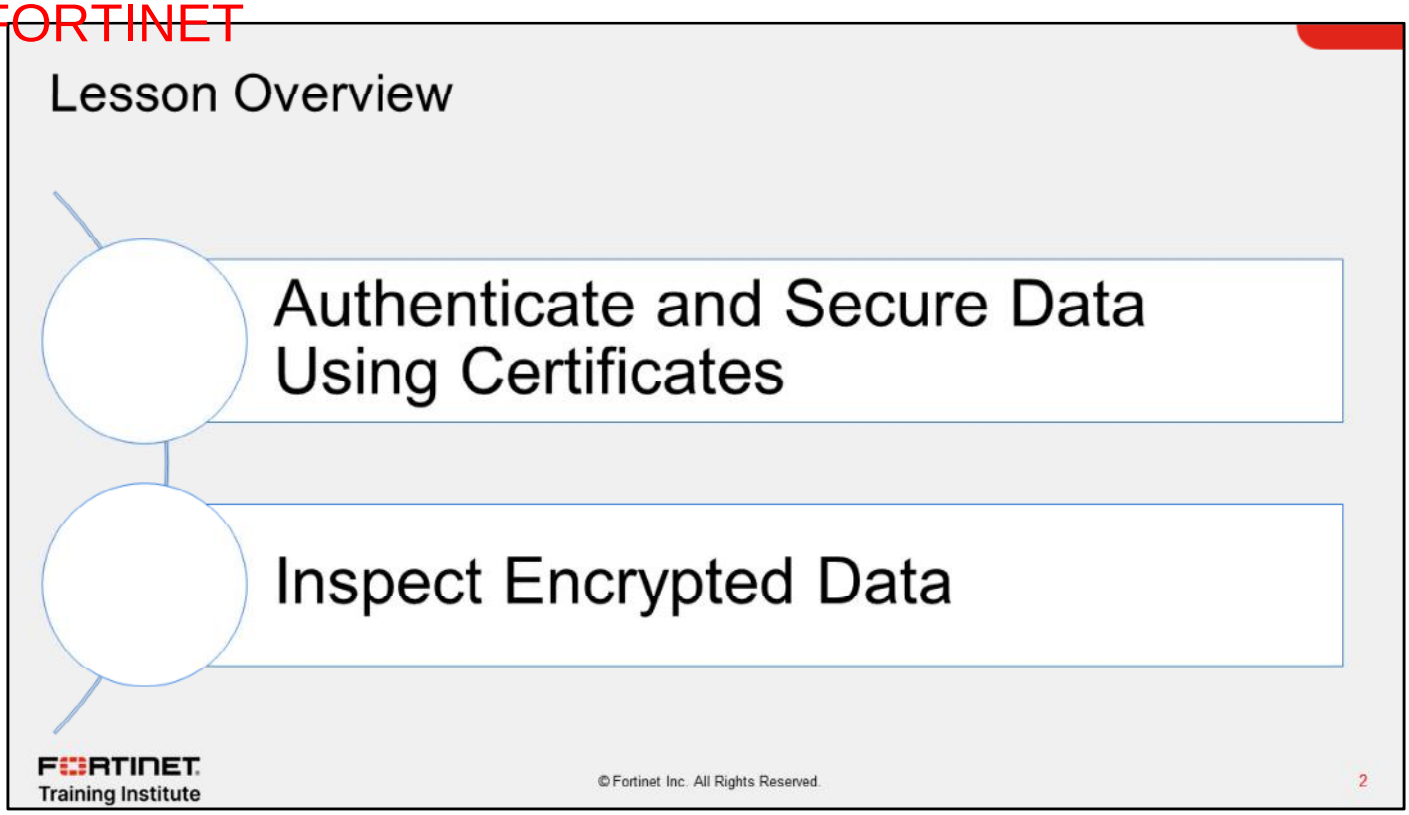

In this lesson, you will learn about the topics shown on this slide.

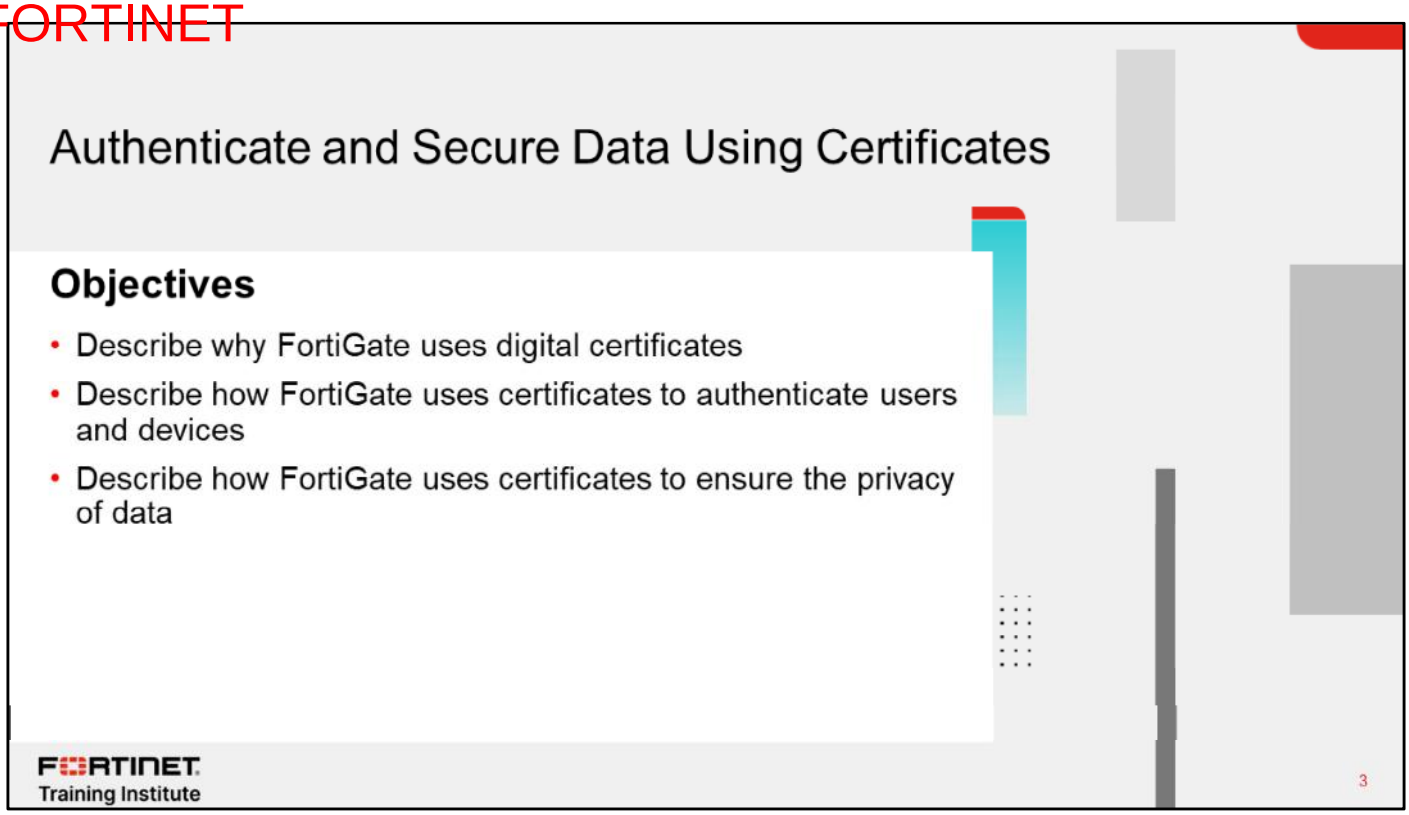

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating an understanding of how FortiGate uses certificates, you will be better able to judge how and when certificates could be used in your own networks.

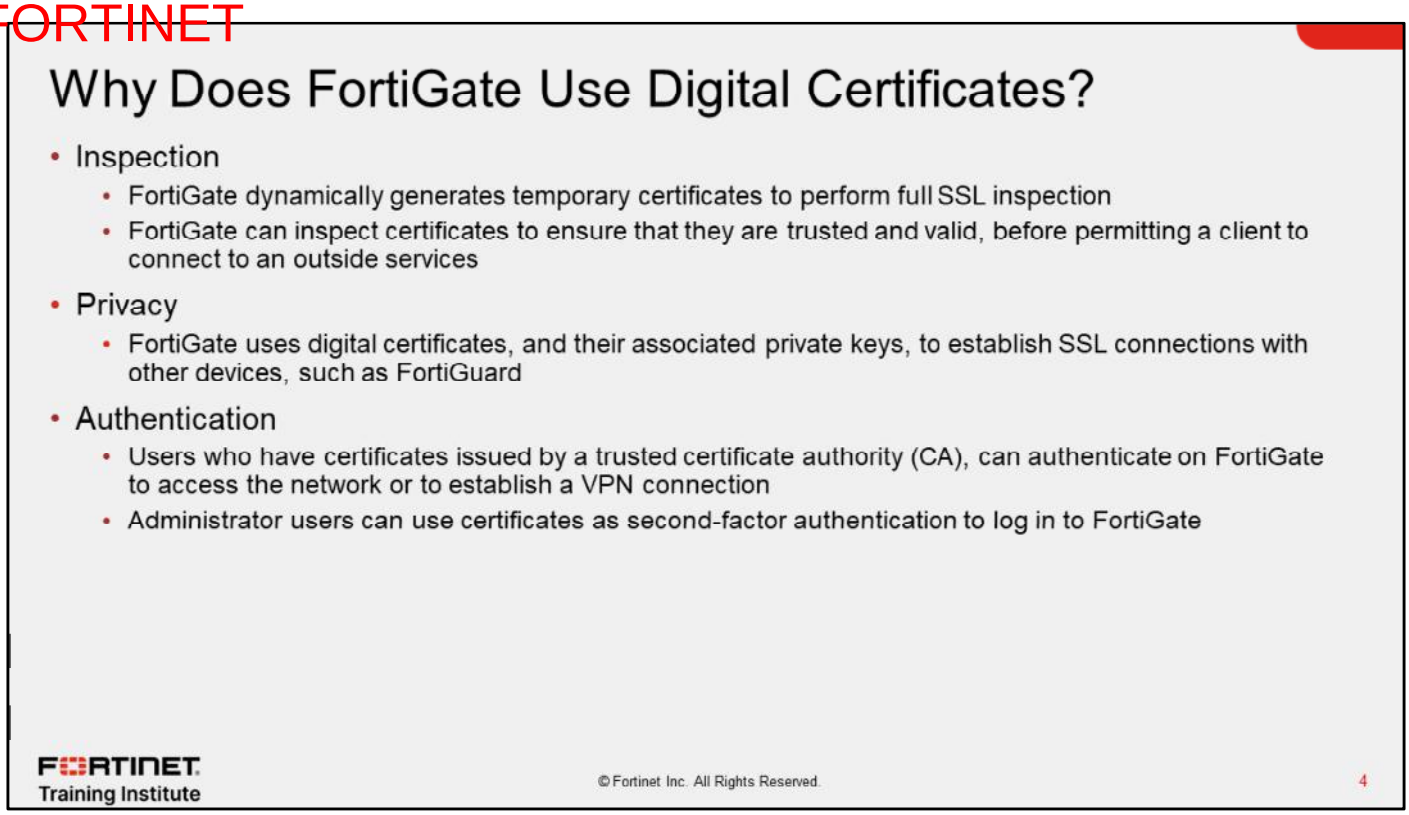

FortiGate uses digital certificates to enhance security.

FortiGate uses digital certificates for inspection. The device can generate certificates on demand for the purpose of inspecting encrypted data that is transferred between two devices; essentially, a man-in-the-middle (MITM) attack. FortiGate can also inspect certificates to identify people and devices (in the network and on the internet), before it permits a person or device to make a full connection to the entity that it is protecting. If FortiGate trusts the certificate, it permits the connection. But if FortiGate does not trust the certificate, it can prevent the connection. How you configure FortiGate determines the behavior; however, other policies that are being used may also affect whether connection attempts are accepted or rejected.

FortiGate uses digital certificates to enforce privacy. Certificates, and their associated private keys, ensure that FortiGate can establish a private SSL connection to another services, such as FortiGuard, a web browser, or a web server.

FortiGate also uses certificates for authentication. Users who have certificates issued by a known and trusted CA can authenticate on FortiGate to access the network or to establish a VPN connection. Administrator users can use certificates as a second-factor authentication to log in to FortiGate.

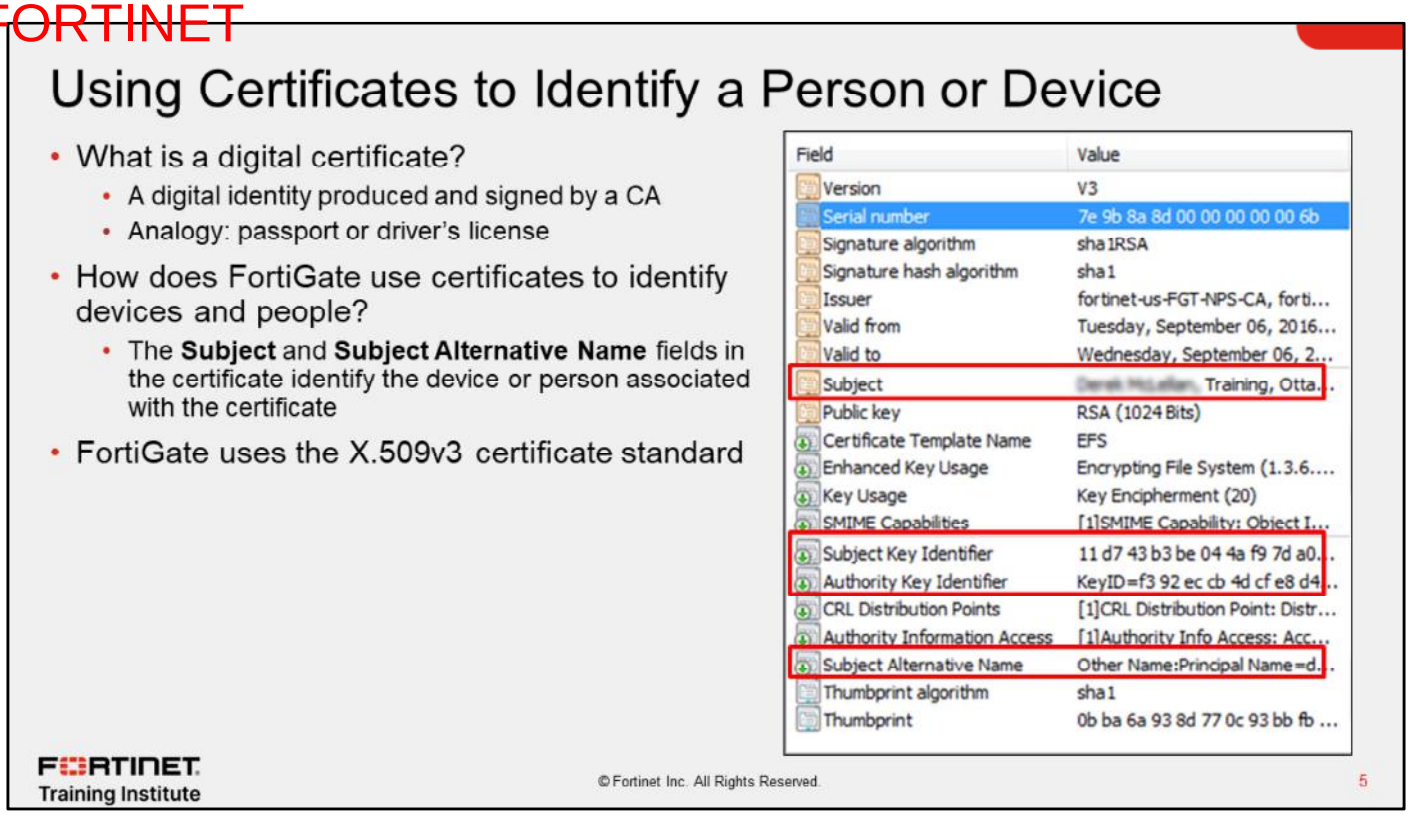

What is a digital certificate?

A digital certificate is a digital document produced and signed by a CA. It identifies an end entity, such as a person (example, Joe Bloggins), a device (example, webserver.acme.com), or thing (example, a certificate revocation list). FortiGate identifies the device or person by reading the value in the **Subject** field, which is expressed as a distinguished name (DN). FortiGate could also use alternate identifiers, shown in the **Subject Alternative Name** field, whose values could be a network ID or an email address, for example. FortiGate can use the **Subject Key Identifier** and **Authority Key Identifier** values to determine the relationship between the issuer of the certificate (identified in the **Issuer** field) and the certificate. FortiGate supports the X.509v3 certificate standard, which is the most common standard for certificates.

How Does FortiGate Trust Certificates?

- FortiGate does the following checks against a certificate before trusting it and using it:
  - Revocation check
    - You must download the relevant certificate revocation lists (CRLs) to FortiGate or configure FortiGate to use OCSP
    - Certificates are identified by a serial number on the CRL
  - CA certificate possession
    - FortiGate uses the **Issuer** value to determine if FortiGate possesses the corresponding CA certificate
    - Without the corresponding CA certificate, FortiGate cannot trust the certificate
  - Validity dates
  - Digital signature validation
    - The verification of the digital signature on the certificate must pass

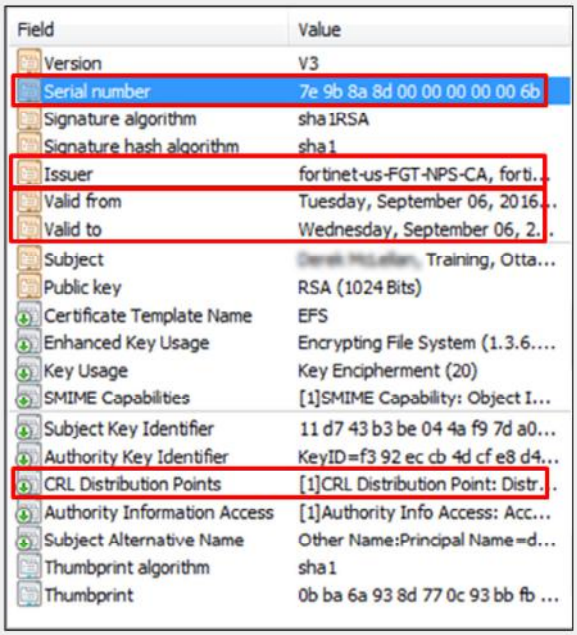

FORTINET Training Institute

6

FortiGate runs the following checks before it trusts the certificate:

- Checks the CRLs locally (on FortiGate) to verify if the certificate has been revoked by the CA. If the serial number of the certificate is listed on the CRL, then the certificate has been revoked and it is no longer trusted. FortiGate also supports Online Certificate Status Protocol (OCSP), where FortiAuthenticator acts as the OCSP responder.
- Reads the value in the **Issuer** field to determine if it has the corresponding CA certificate. Without the CA certificate, FortiGate does not trust the certificate. FortiOS uses the Mozilla CA certificate store. You can view the list by clicking **Security Profiles** > **SSL Inspection** > **View Trusted CA List** > **Factory Bundles**.
- Verifies that the current date is between the **Valid From** and **Valid To** values. If it is not, the certificate is rendered invalid.
- Validates the signature on the certificate. The signature must be successfully validated. Because a valid signature is a critical requirement for trusting a certificate, it may be useful to review how FortiGate verifies digital signatures.

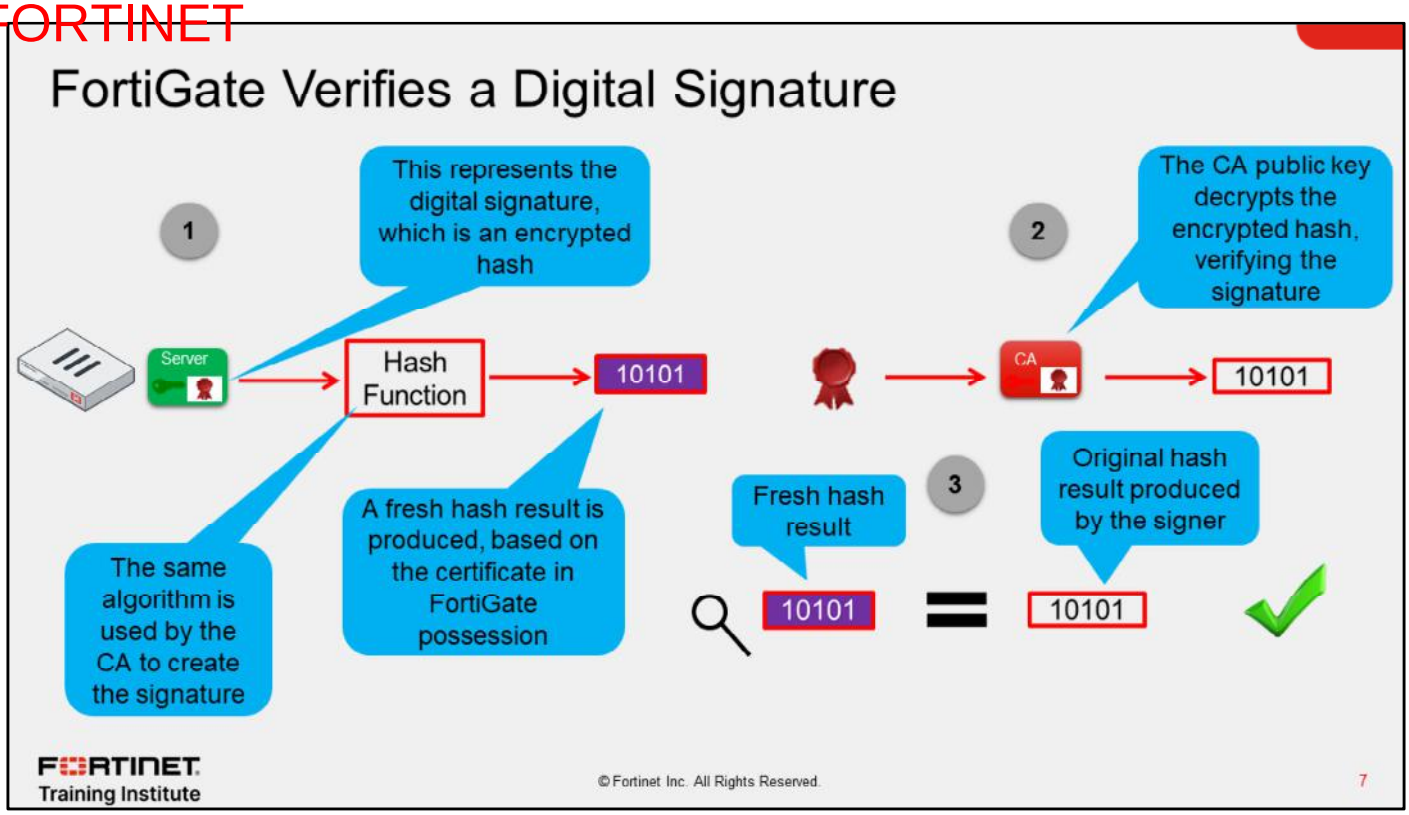

Before it generates a digital signature, the CA runs the content of the certificate through a hash function, which produces a hash result. The hash result, which is a mathematical representation of the data, is referred to as the *original hash result*. The CA encrypts the original hash result using its private key. The encrypted hash result is the digital signature.

When FortiGate verifies the digital signature, it runs the certificate through a hash function, producing a fresh hash result. FortiGate must use the same hash function, or hashing algorithm, that the CA used to create the digital signature. The hashing algorithm is identified in the certificate.

In the second part of the verification process, FortiGate decrypts the encrypted hash result (or digital signature) using the CA public key, and applying the same algorithm that the CA used to encrypt the hash result. This process verifies the signature. If the key cannot restore the encrypted hash result to its original value, then the signature verification fails.

In the third, and final, part of the verification process, FortiGate compares the fresh hash result to the original hash result. If the two values are identical, then the integrity of the certificate is confirmed. If the two hash results are different, then the version of the certificate that FortiGate has is not the same as the one that the CA signed, and data integrity fails.

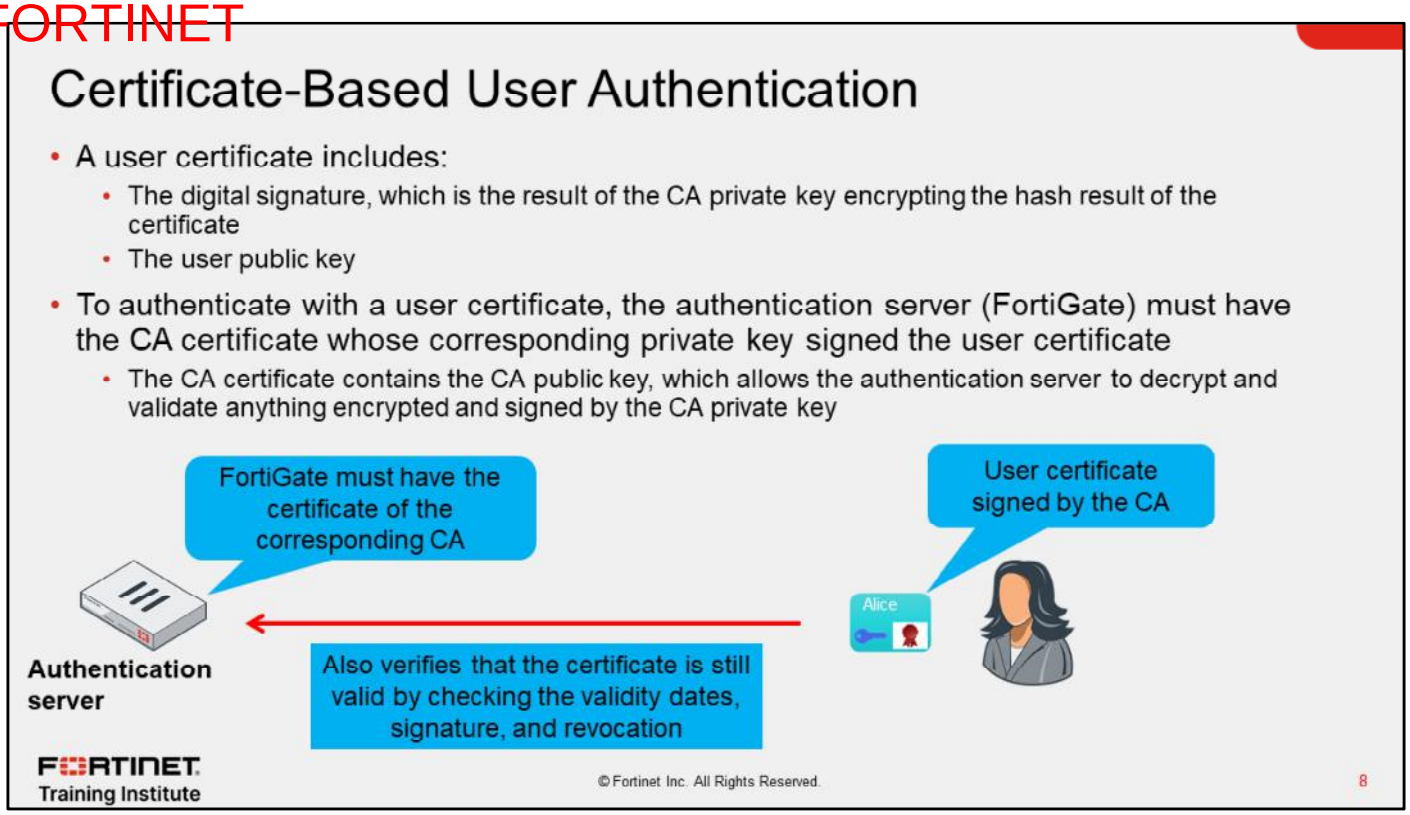

Certificate-based user authentication uses an end-entity certificate to identify the user. This certificate contains the user public key and the signature of the CA that issued the certificate. The authentication server (for example, FortiGate) must have the CA certificate whose private key signed the user certificate. FortiGate verifies that the certificate signature is valid, that the certificate has not expired, and that the certificate hasn't been revoked. If any of these verifications fail, the certificate-based user authentication fails.

You can configure FortiGate to require that administrators use certificates for second-factor authentication. The process for verifying administrator certificates is the same.

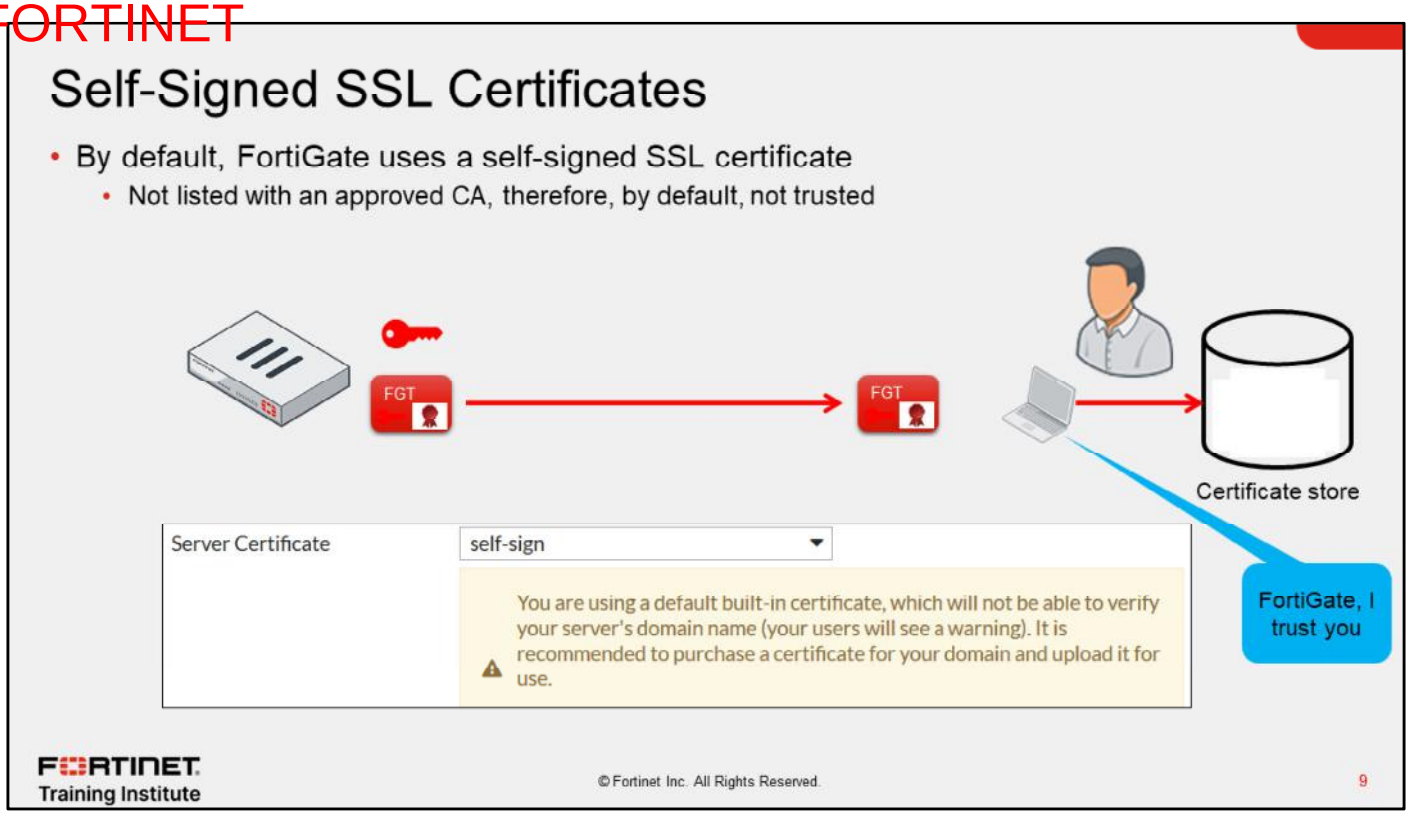

As you can see in the example shown on this slide, trust in the web model is determined by whether or not your certificate store possesses the CA certificate that is required to verify the signature on the SSL certificate. Certificate stores come prepopulated with root and subordinate CA certificates. You can choose to add or remove the certificates, which will affect which websites you trust.

By default, FortiGate uses a self-signed certificate to authenticate itself to HTTPS clients.

You can configure self-signed certificates to establish SSL sessions, just like those certificates issued by Verisign, Entrust Datacard, and other certificate vendors. But, because self-signed certificates do not come prepopulated in client certificate stores, your end users get a security warning. You can choose to add the self-signed certificate to clients, or to purchase an SSL certificate from an approved CA vendor for your FortiGate device.

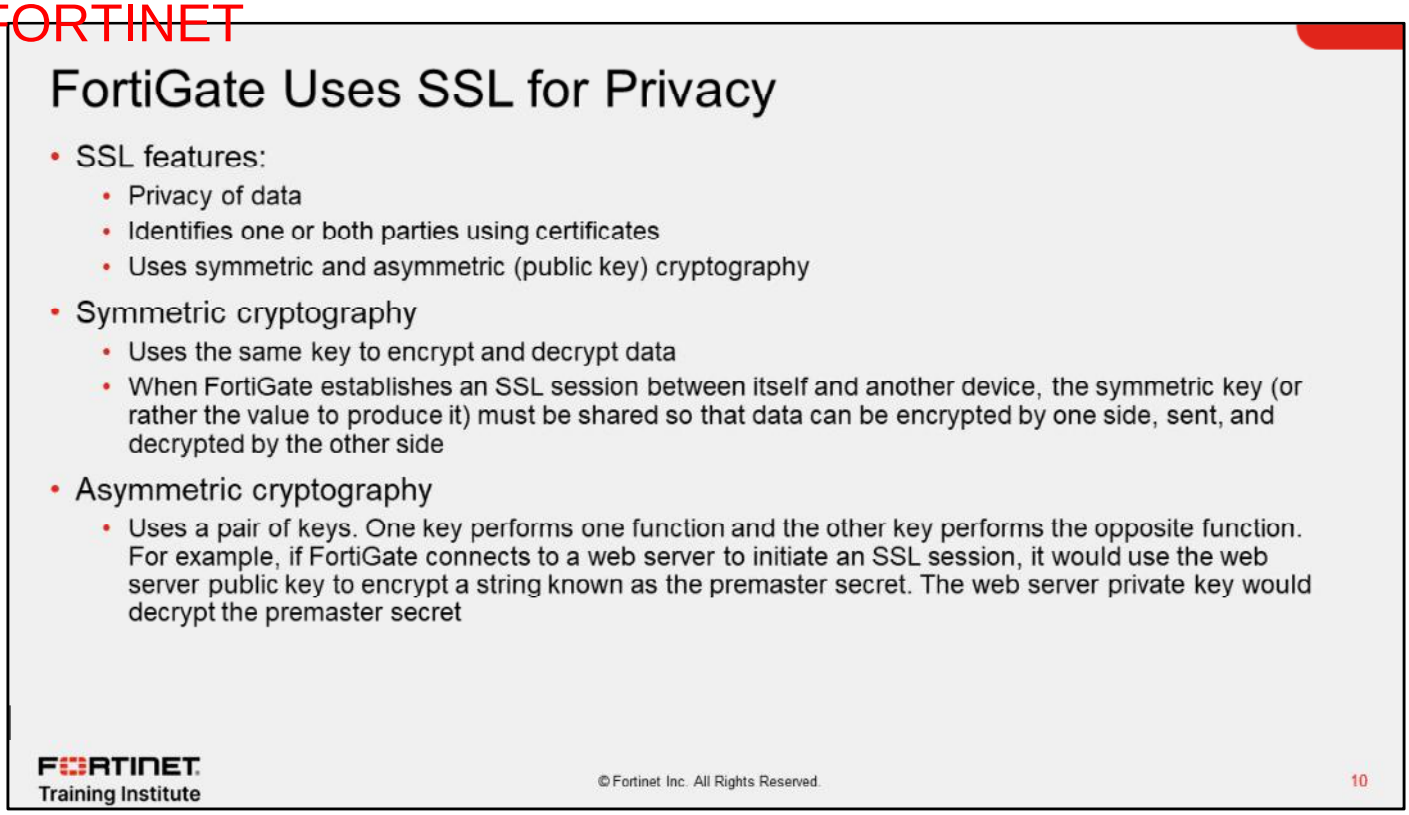

FortiGate uses SSL to ensure that data remains private when connecting with servers, such as FortiGuard, and with clients, such as a web browser. Another feature of SSL is that FortiGate can use it to identify one or both parties using certificates. SSL uses symmetric and asymmetric cryptography to establish a secure session between two points.

It is beneficial to understand the high-level process of an SSL handshake, in order to understand how FortiGate secures private sessions.

An important attribute of symmetric cryptography is that the same key is used to encrypt and decrypt data. When FortiGate establishes an SSL session between itself and another device it must share, the symmetric key (or rather the value required to produce it), so that data can be encrypted by one side, sent, and decrypted by the other side.

Asymmetric cryptography uses a pair of keys: one key performs one function and the other key performs the opposite function. When FortiGate connects to a web server, for example, it uses the web server public key to encrypt a string known as the premaster secret. The web server private key decrypts the premaster secret.

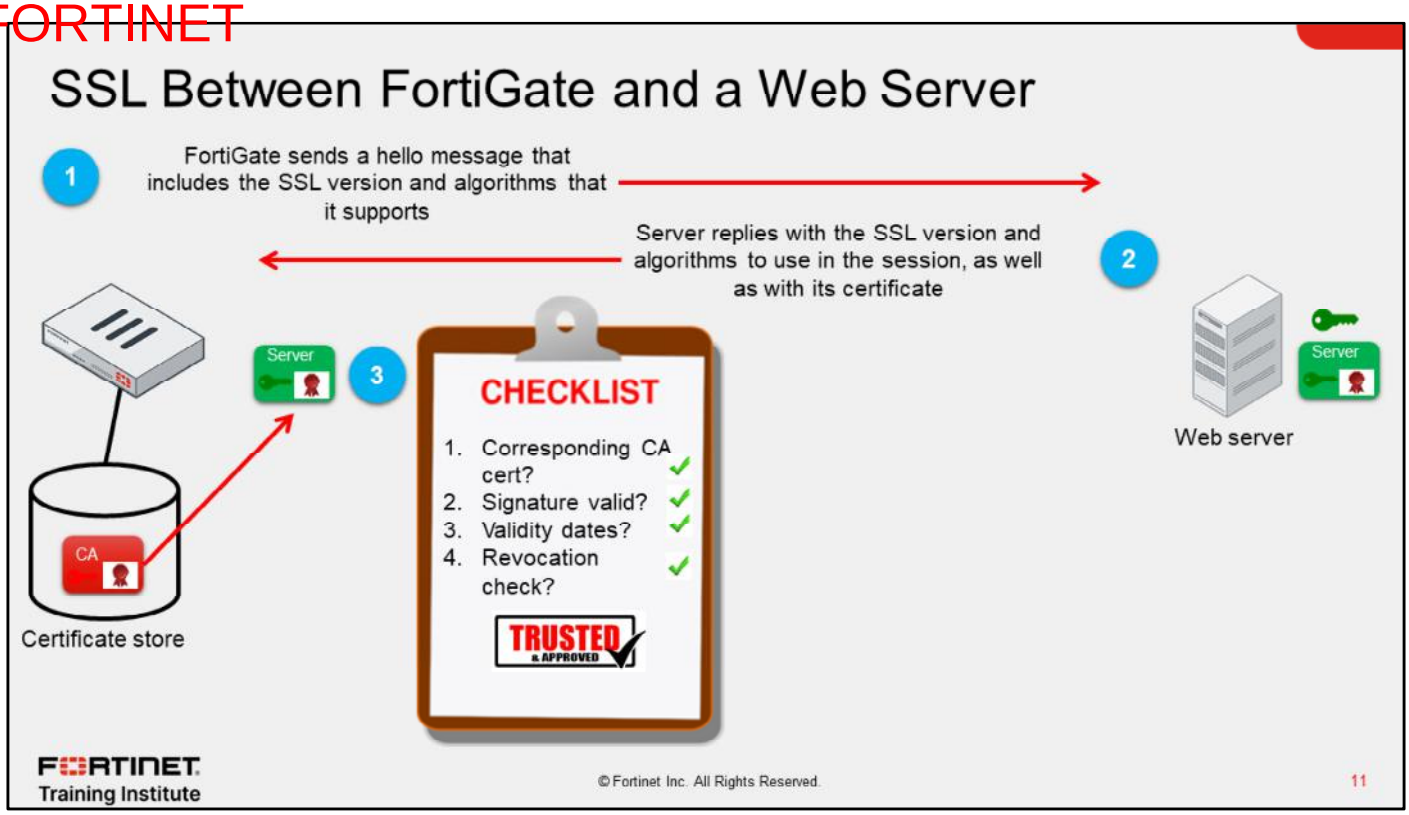

Now, you will learn more about the process of establishing an SSL session.

In the first step of the example shown on this slide, FortiGate connects to a web server that is configured for SSL. In the initial hello message, the browser provides critical information that is needed to communicate with the web server. This information includes the SSL version number and the names of the cryptographic algorithms that it supports.

In the second step, the web server receives the message from FortiGate and chooses the first suite of cryptographic algorithms included in the message, and verifies that it is also supported by the web server. The web server replies with the chosen SSL version and cipher suite, and then sends its certificate to FortiGate. Note that the certificate information is passed as cleartext over the public network. The information contained in a certificate is typically public, so this is not a security concern.

In the third step, FortiGate validates the web server certificate. The checklist shown on this slide represents the checks that FortiGate performs on the certificate to ensure that it can be trusted. If FortiGate determines that the certificate can be trusted, then the SSL handshake continues.

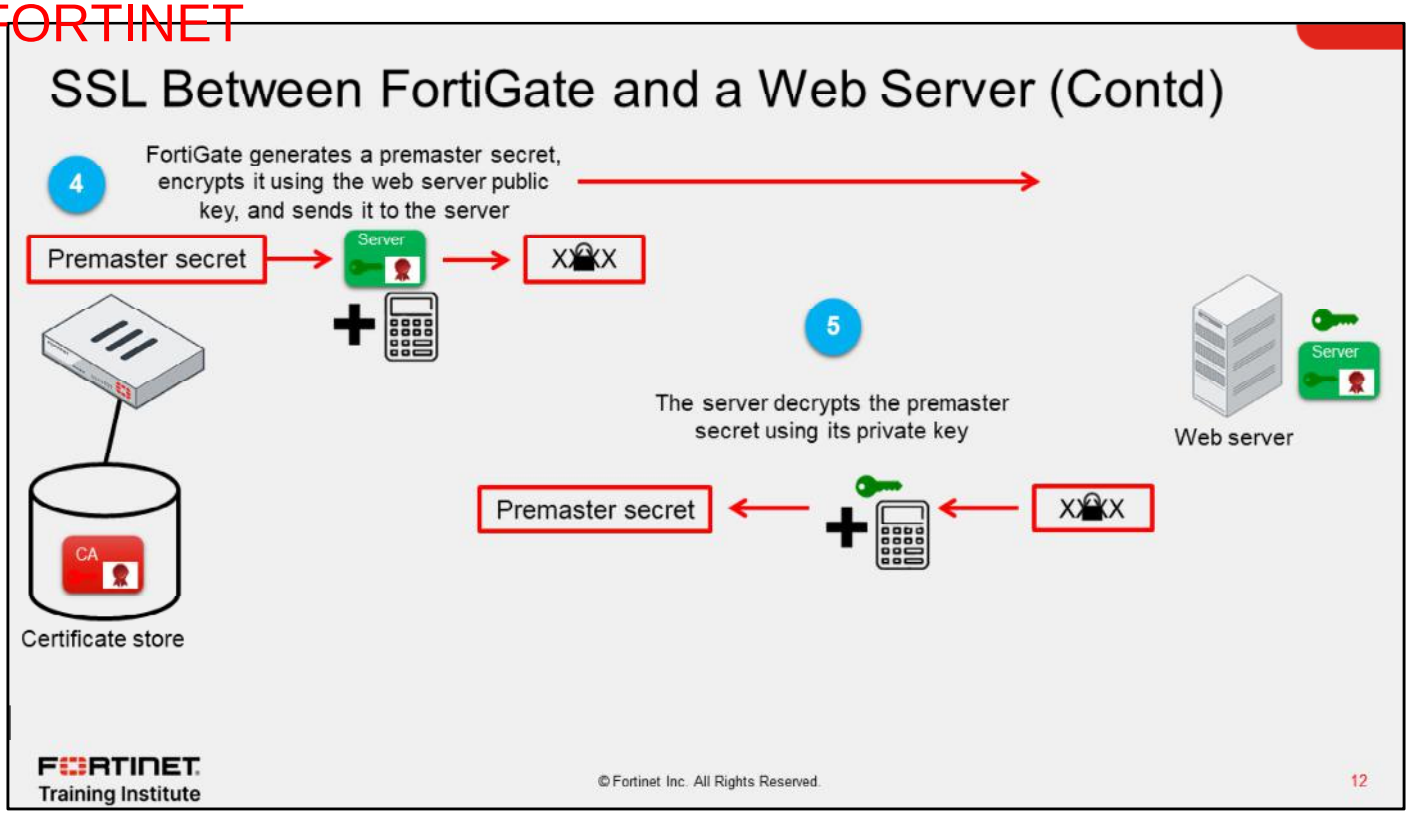

In the fourth step, FortiGate generates a value known as the premaster secret. FortiGate uses the server public key, which is in the certificate, to encrypt the premaster secret. FortiGate then sends the encrypted premaster secret to the web server. If a third-party intercepted the premaster secret, they would be unable to read it, because they do not have the private key.

In the fifth step, the web server uses its private key to decrypt the premaster secret. Now, both FortiGate and the web server share a secret value that is known by only these two devices.

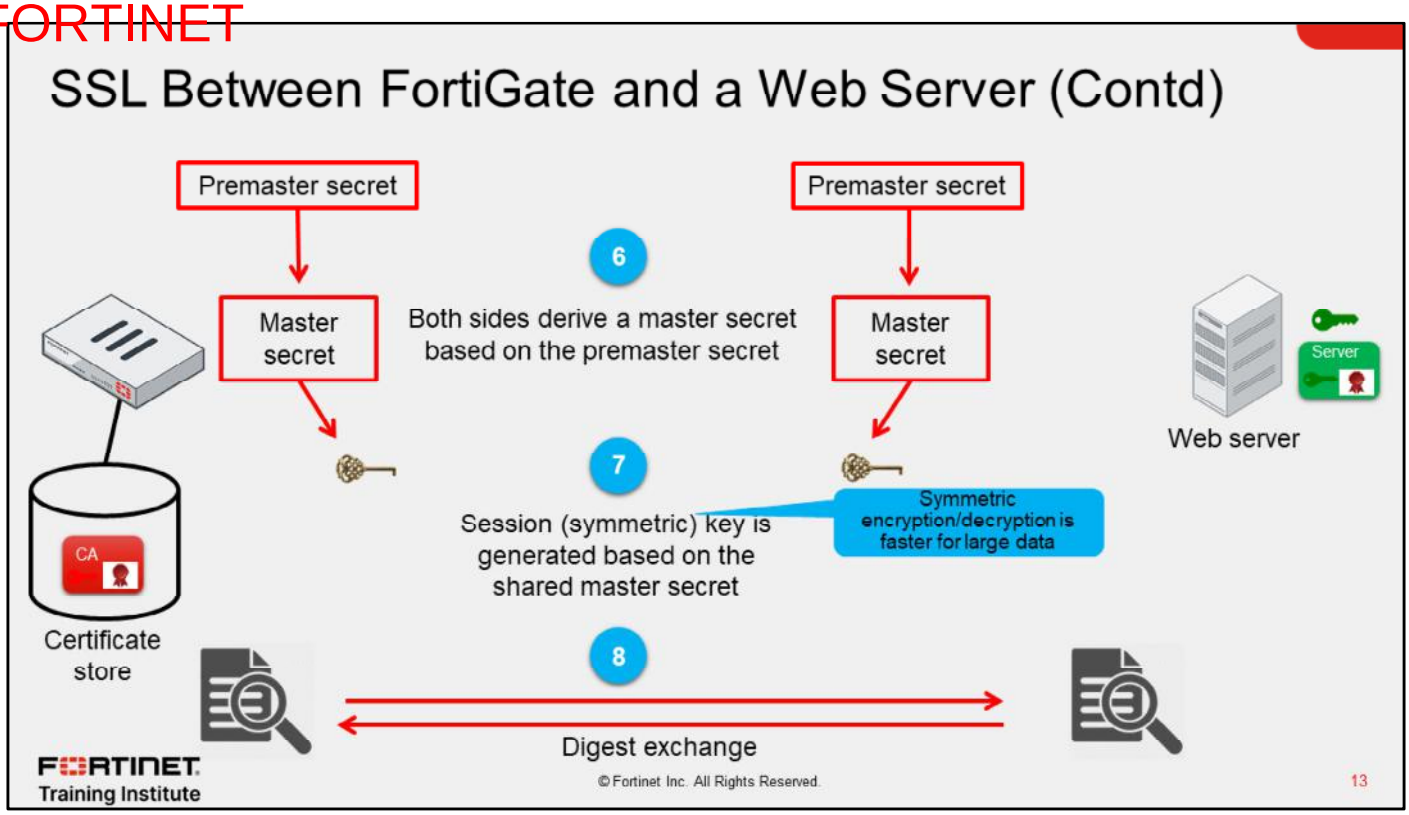

In the sixth step, both FortiGate and the web server derive the master secret based on the premaster secret.

In the seventh step, based on the master secret value, FortiGate and the web server generate the session key. The session key is a symmetric key. The main advantage of symmetric key over asymmetric keys is that it is fast and efficient for large amounts of data. It is required to encrypt and decrypt the data. Because both sides have the session key, both sides can encrypt and decrypt data for each other.

In the eighth and final step before these two entities establish the secure connection, both FortiGate and the web server send each other a summary (or digest) of the messages sent so far. The digests are encrypted with the session key. The digests ensure that none of the messages exchanged during the creation of the session have been intercepted or replaced. If the digests match, the secure communication channel is established.

The SSL handshake is now complete. Both FortiGate and the web server are ready to communicate securely, using the session keys to encrypt and decrypt the data they send over the network or internet.

Knowledge Check

1. Which attribute or extension identifies the owner of a certificate?
   - ✓ A. The subject name in the certificate
   - B. The unique serial number in the certificate

2. How does FortiGate determine if a certificate has been revoked?
   - ✓ A. It checks the CRL that resides on FortiGate.
   - B. It retrieves the CRL from a directory server.

FORTINET Training Institute

14

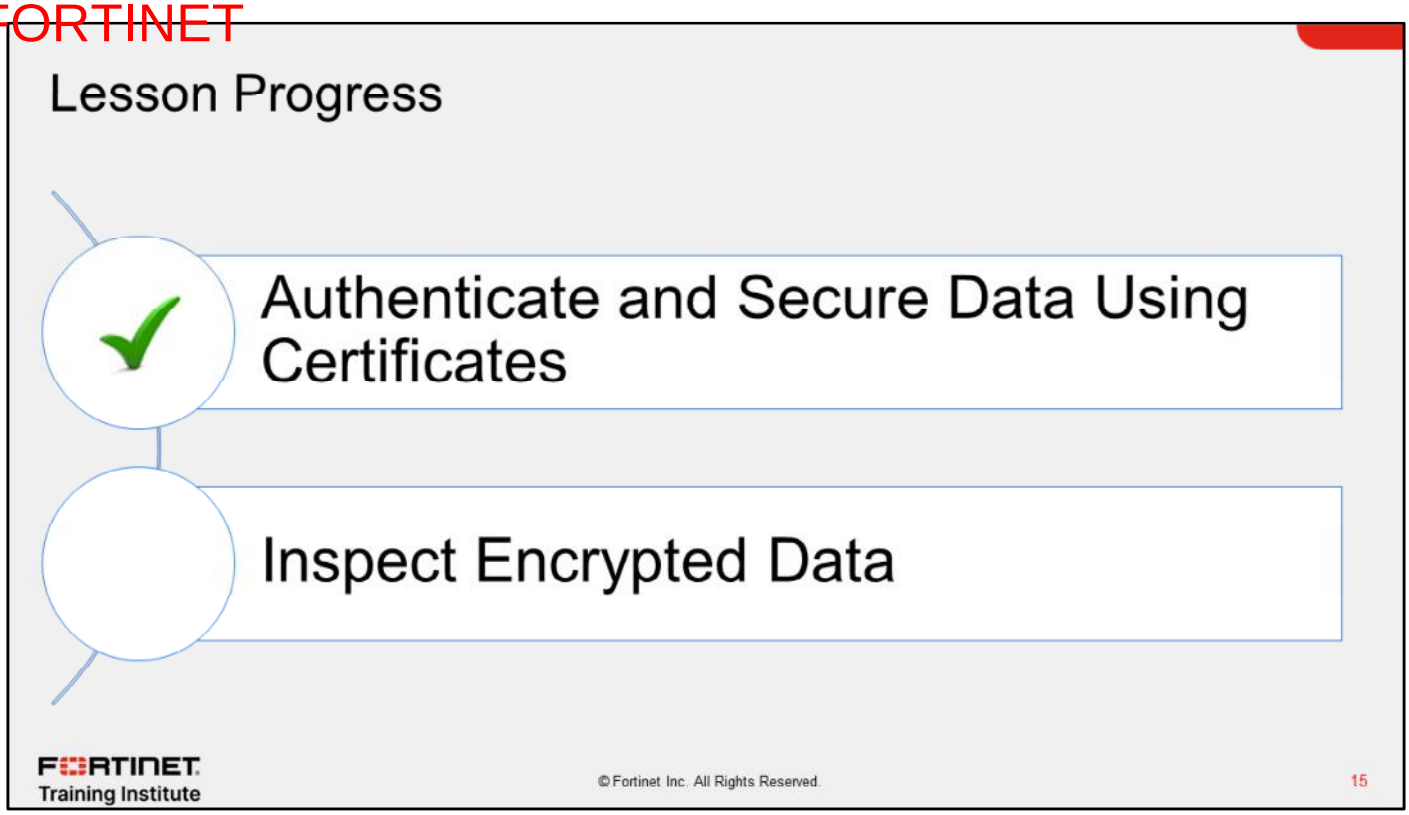

Good job! You now understand why and how FortiGate uses certificates to authenticate devices and people. You also understand how FortiGate uses certificates to ensure the privacy of data as it flows from FortiGate to another device, or from another device to FortiGate.

Now, you will learn how to inspect encrypted data.

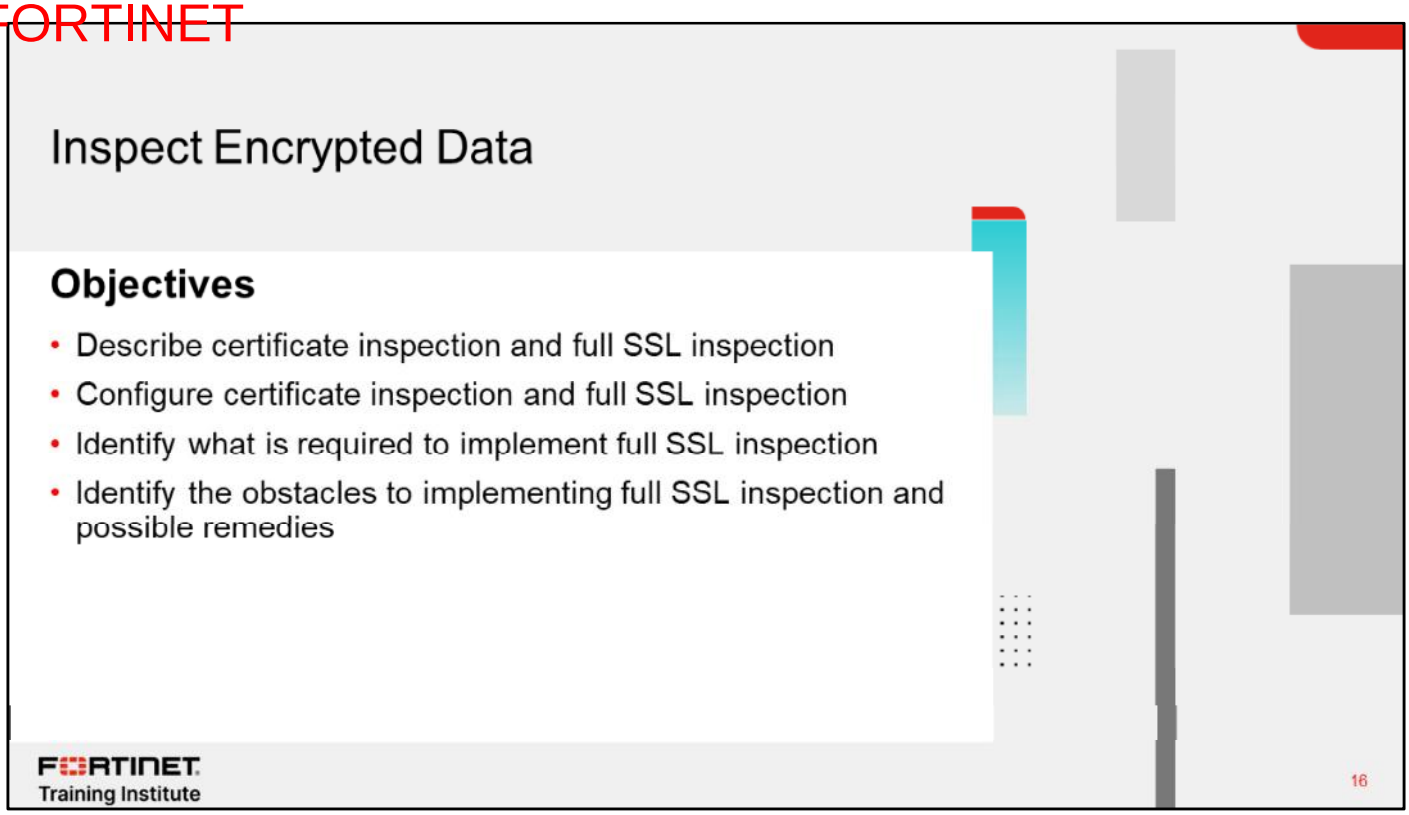

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding and configuring full SSL inspection and certificate inspection, you will be able to implement one of these SSL inspection solutions in your network.

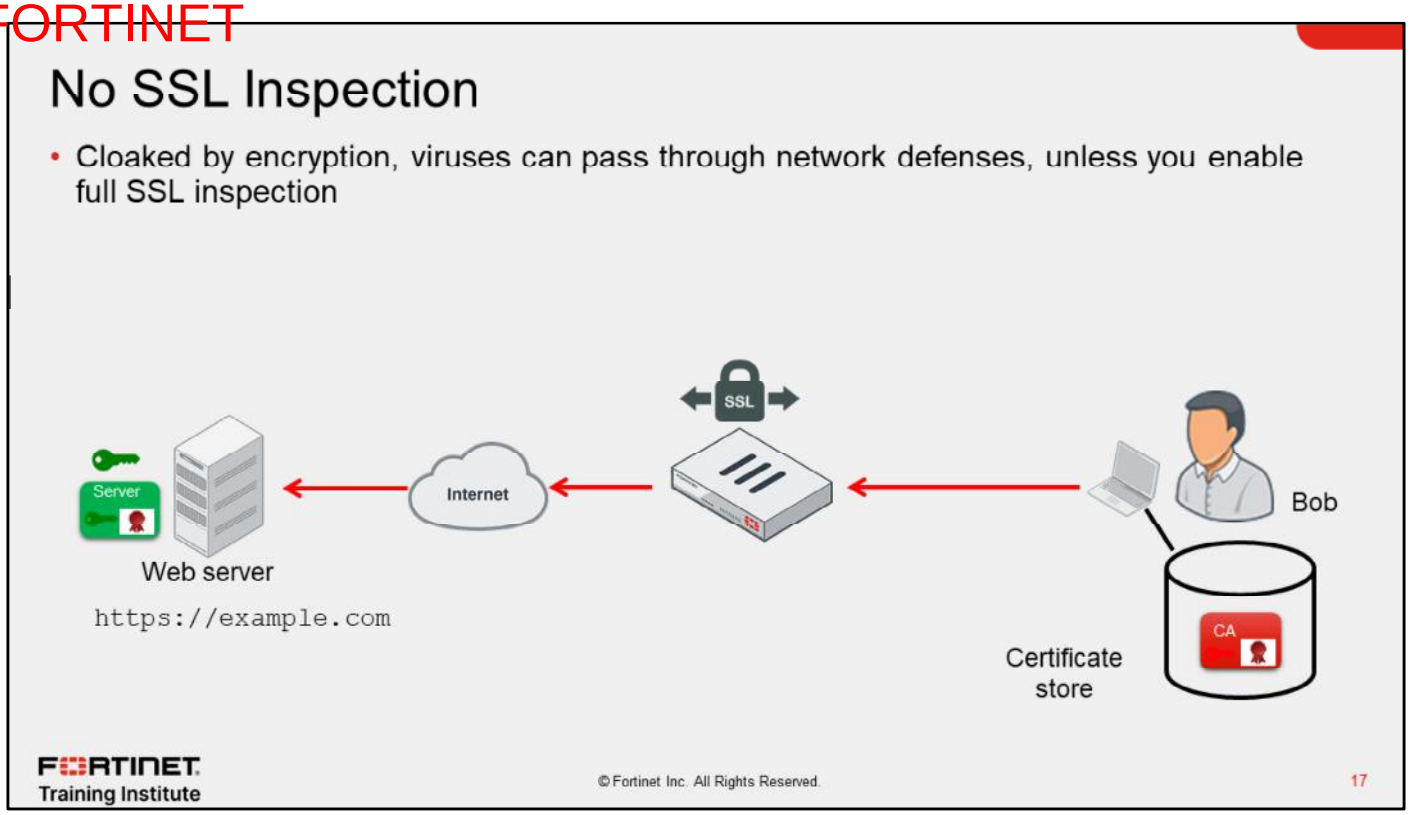

(slide contains animation)

While there are benefits to using HTTPS, there are risks associated with its use as well, because encrypted traffic can be used to get around normal defenses. For example, if a session is encrypted when you download a file containing a virus, the virus might get past your network security measures.

In the example shown on this slide, Bob connects to a site with a certificate issued by a legitimate CA. Because the CA is an approved CA, the CA verification certificate is in Bob's certificate store, and Bob's browser is able to establish an SSL session with the `example.com` site. However, unknown to Bob, the example.com site has been infected with a virus. The virus, cloaked by encryption, passes through FortiGate undetected and enters Bob's computer. The virus is able to breach security because full SSL inspection is not enabled.

You can use full SSL inspection, also known as deep inspection, to inspect encrypted sessions.

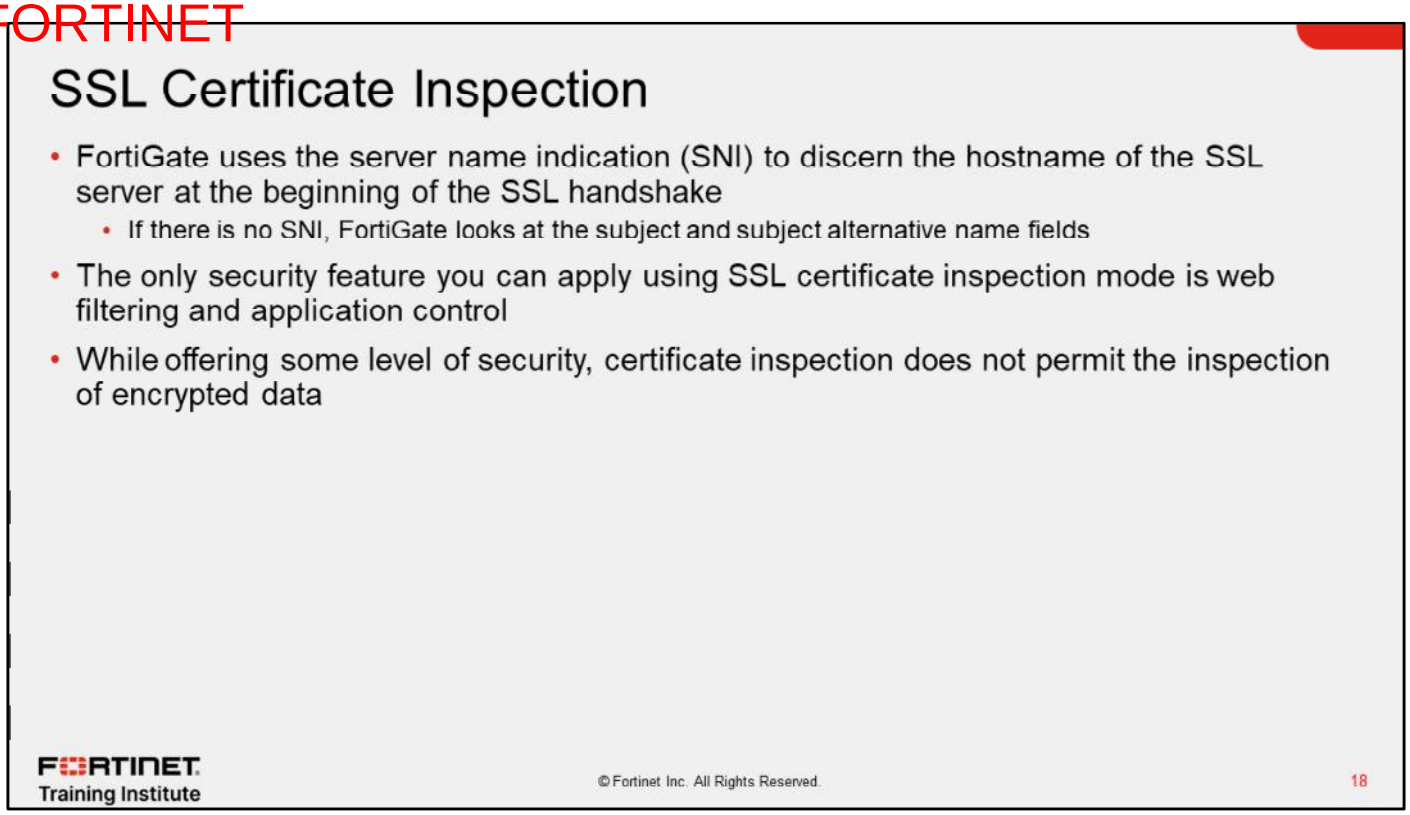

During the exchange of hello messages at the beginning of an SSL handshake, FortiGate parses server name indication (SNI) from client Hello, which is an extension of the TLS protocol. The SNI tells FortiGate the hostname of the SSL server, which is validated against the DNS name before receipt of the server certificate. If there is no SNI exchanged, then FortiGate identifies the server by the value in the **Subject** field or **SAN** (subject alternative name) field in the server certificate.

When you use certificate inspection, FortiGate inspects only the header information of the packets. You use certificate inspection to verify the identity of web servers. You can also use it to make sure that the HTTPS protocol isn't used as a workaround to access sites you have blocked using web filtering.

The only security feature that you can apply using SSL certificate inspection mode is web filtering and application control. However, since only the packet is inspected, this method does not introduce certificate errors and can be a useful alternative to full SSL inspection when you use web filtering.

Certificate inspection offers some level of security, but it does *not* allow FortiGate to inspect the flow of encrypted data between the outside server and the internal client.

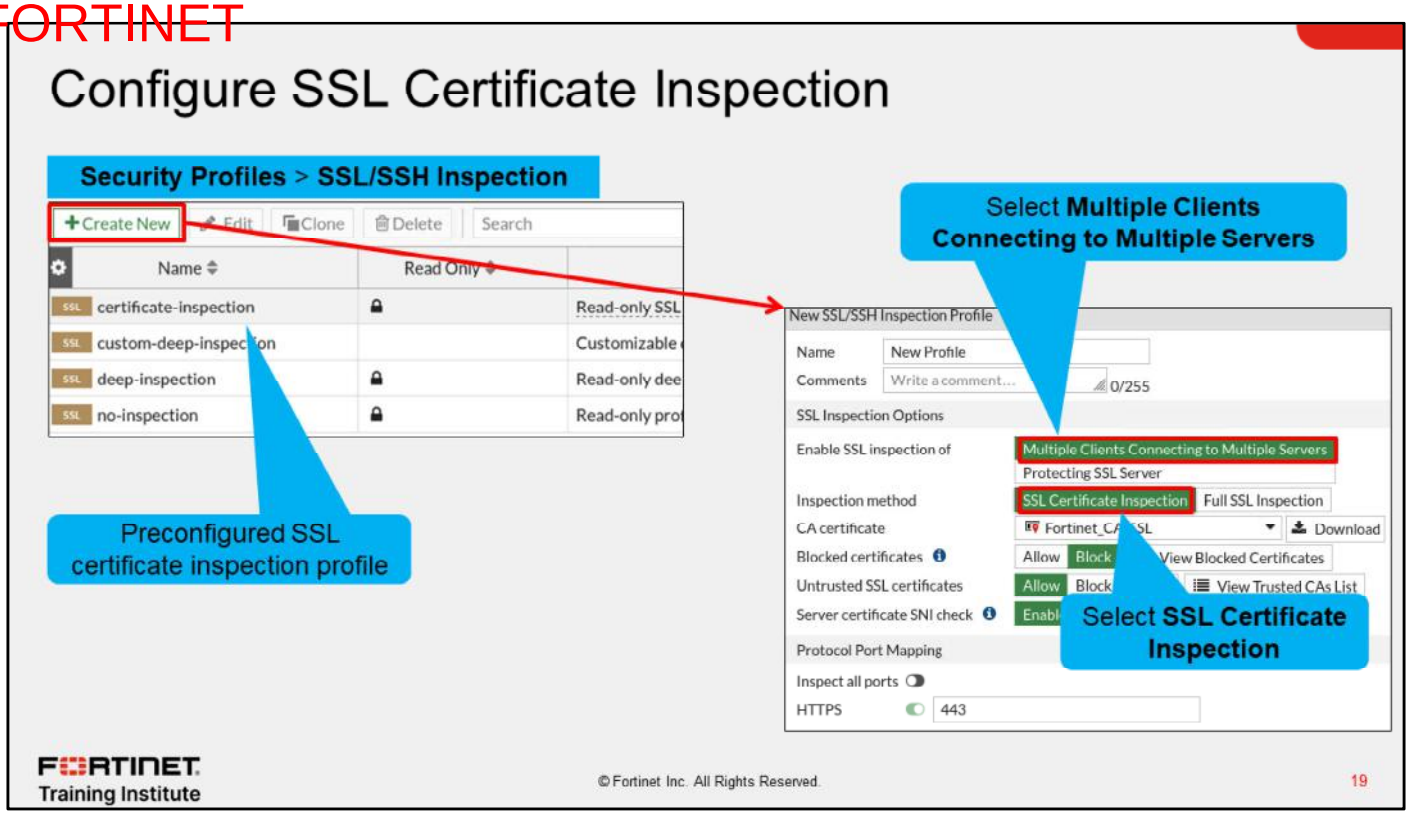

FortiGate has a read-only preconfigured profile for SSL certificate inspection named **certificate-inspection**. If you want to enable SSL certificate inspection, select this profile when configuring a firewall policy.

Alternatively, you can create your own profile for SSL certificate inspection by following the steps below:

1. On the FortiGate GUI, click **Security Profiles** > **SSL/SSH Inspection**.
2. Click **Create New** to create a new SSL/SSH inspection profile.
3. Select **Multiple Clients Connecting to Multiple Servers**, and click **SSL Certificate Inspection**.

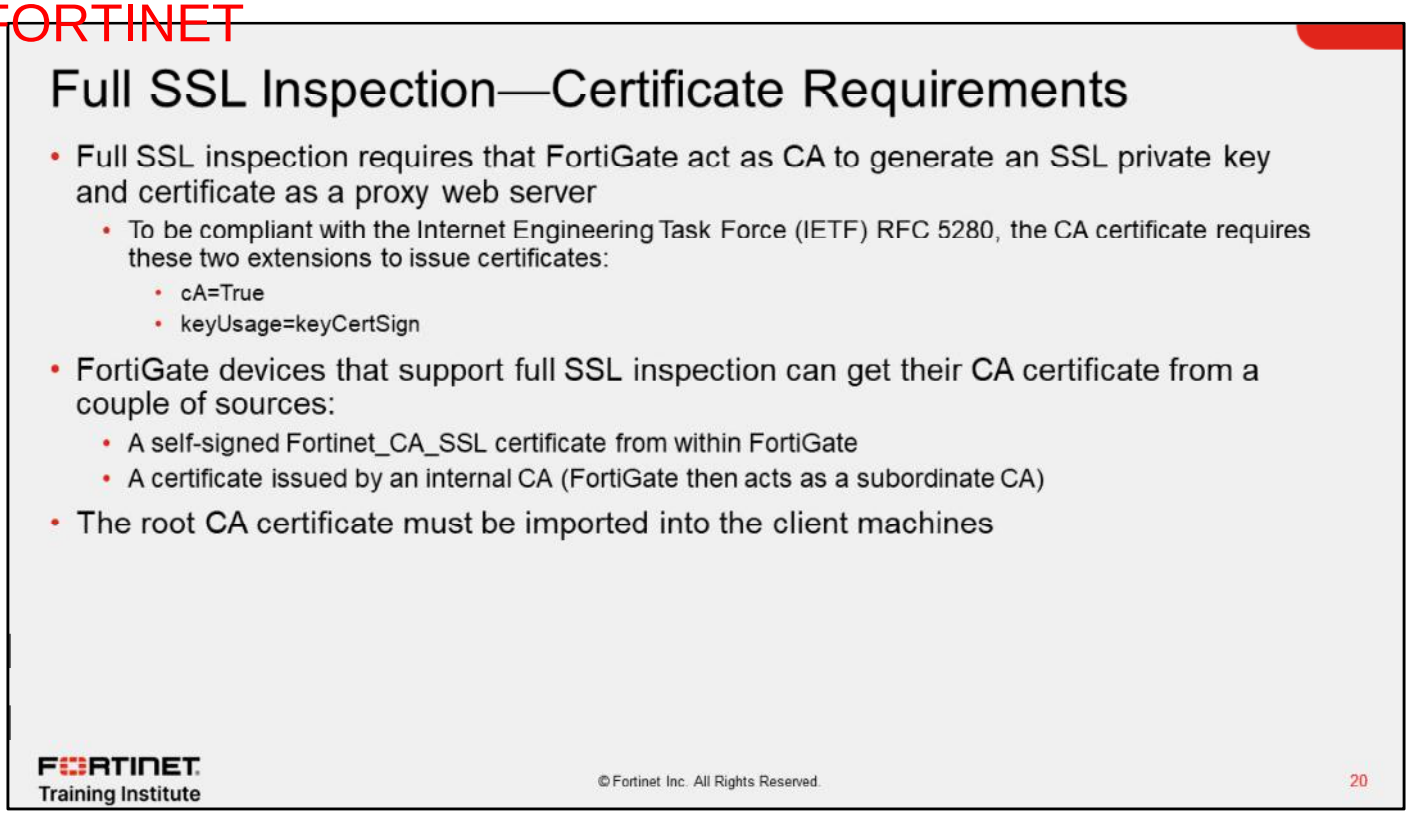

FortiGate performs web proxy and must act as a CA in order for it to perform full SSL inspection. The internal CA must generate an SSL private key and certificate each time an internal user connects to an external SSL server. The key pair and certificate are generated *immediately* so the user connection with the web server is not delayed.

Although it appears as though the user browser is connected to the web server, the browser is connected to FortiGate. FortiGate is acting as a proxy web server. In order for FortiGate to act in these roles, its CA certificate must have the basic constraints extension set to **cA=True** and the value of the **keyUsage** extension set to **keyCertSign**.

The **cA=True** value identifies the certificate as a CA certificate. The **keyUsage=keyCertSign** value indicates that the certificate corresponding private key is permitted to sign certificates. For more information, see *RFC 5280 Section 4.2.1.9 Basic Constraints*.

All FortiGate devices that support full SSL inspection can use the self-signed Fortinet_CA_SSL certificate that is provided with FortiGate, or an internal CA, to issue FortiGate a CA certificate. When FortiGate uses an internal CA, FortiGate acts as a subordinate CA. Note that your client machines and devices must import the root CA certificate, in order to trust FortiGate and accept an SSL session. You must install the chain of CA certificates on FortiGate. FortiGate sends the chain of certificates to the client, so that the client can validate the signatures and build a chain of trust.

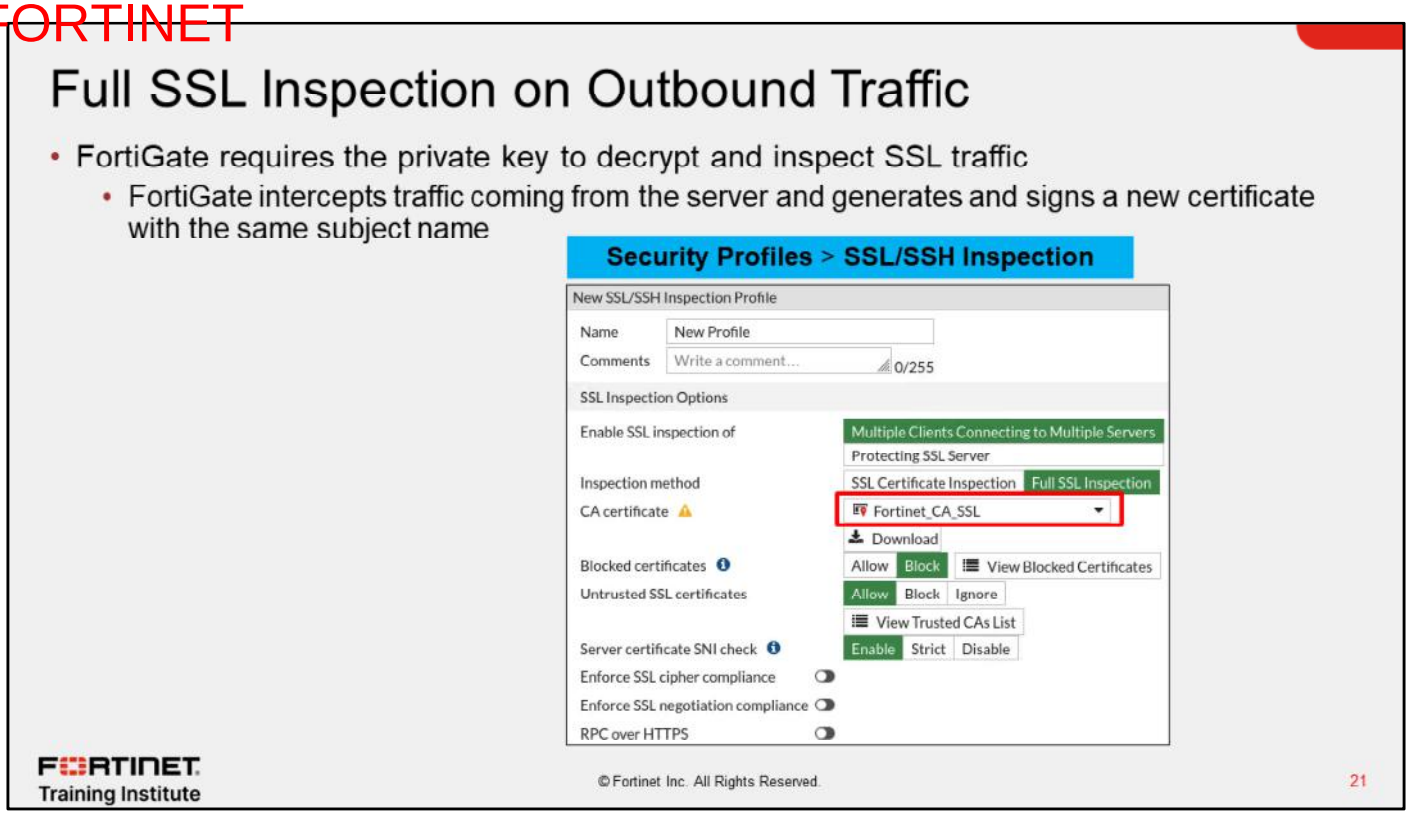

Some FortiGate devices offer a mechanism to inspect encrypted data that flows between external SSL servers and internal clients. Without full SSL inspection, FortiGate cannot inspect encrypted traffic, because the firewall does not have the SSL key that is required to decrypt the data, and that was negotiated between client and server during the SSL handshake.

There are two possible configurations for full SSL inspection: one for outbound traffic and one for inbound traffic.

If the connection request is outbound (initiated by an internal client to an external server), you must select the option, **Multiple Clients Connecting to Multiple Servers**. Then, you must select the CA certificate that will be used to sign the new certificates. In the example shown on this slide, it is the built-in **FortiGate_CA_SSL** certificate, which is available on FortiGate devices that support SSL inspection. You will also learn about configuring full SSL inspection for inbound traffic in this lesson.

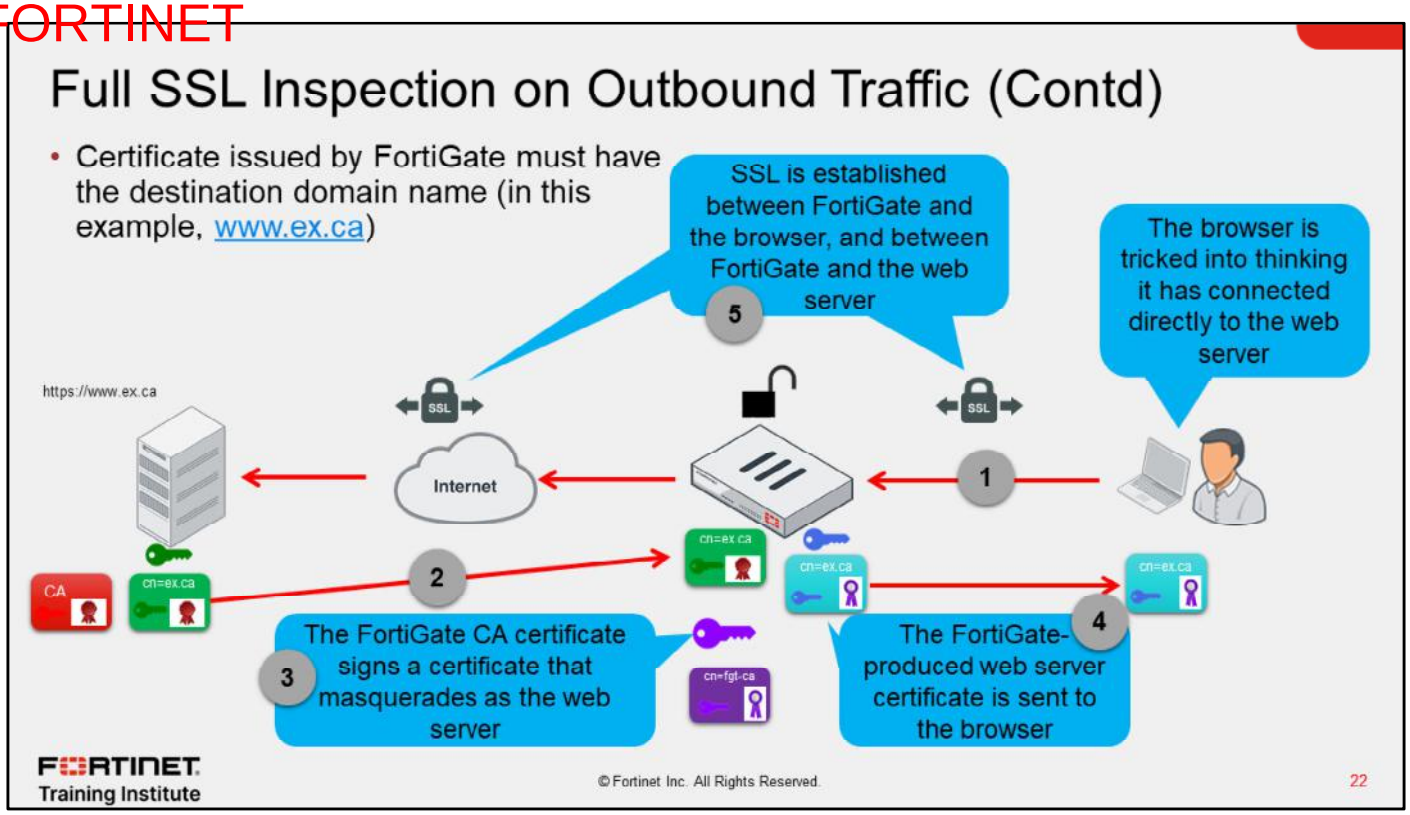

In step 1, an internal web browser connects to an SSL-enabled web server. Normally, when a browser connects to a secure site, the web server sends its certificate to the browser. However, in step 2, FortiGate intercepts the web server certificate. In step 3, the FortiGate internal CA generates a new key pair and certificate. The new certificate subject name must be the DNS name of the website (for example, `ex.ca`). In steps 4 and 5, the new key pair and certificate are used to establish a secure connection between FortiGate and the web browser. A new temporary key pair and certificate are generated each time a client requests a connection with an external SSL server.

Outward facing and included in step 5, FortiGate uses the web server certificate to initiate a secure session with the web server. In this configuration, FortiGate can decrypt the data from both the web server and the browser, in order to scan the data for threats before re-encrypting it and sending it to its destination. This scenario is, essentially, an MITM attack.

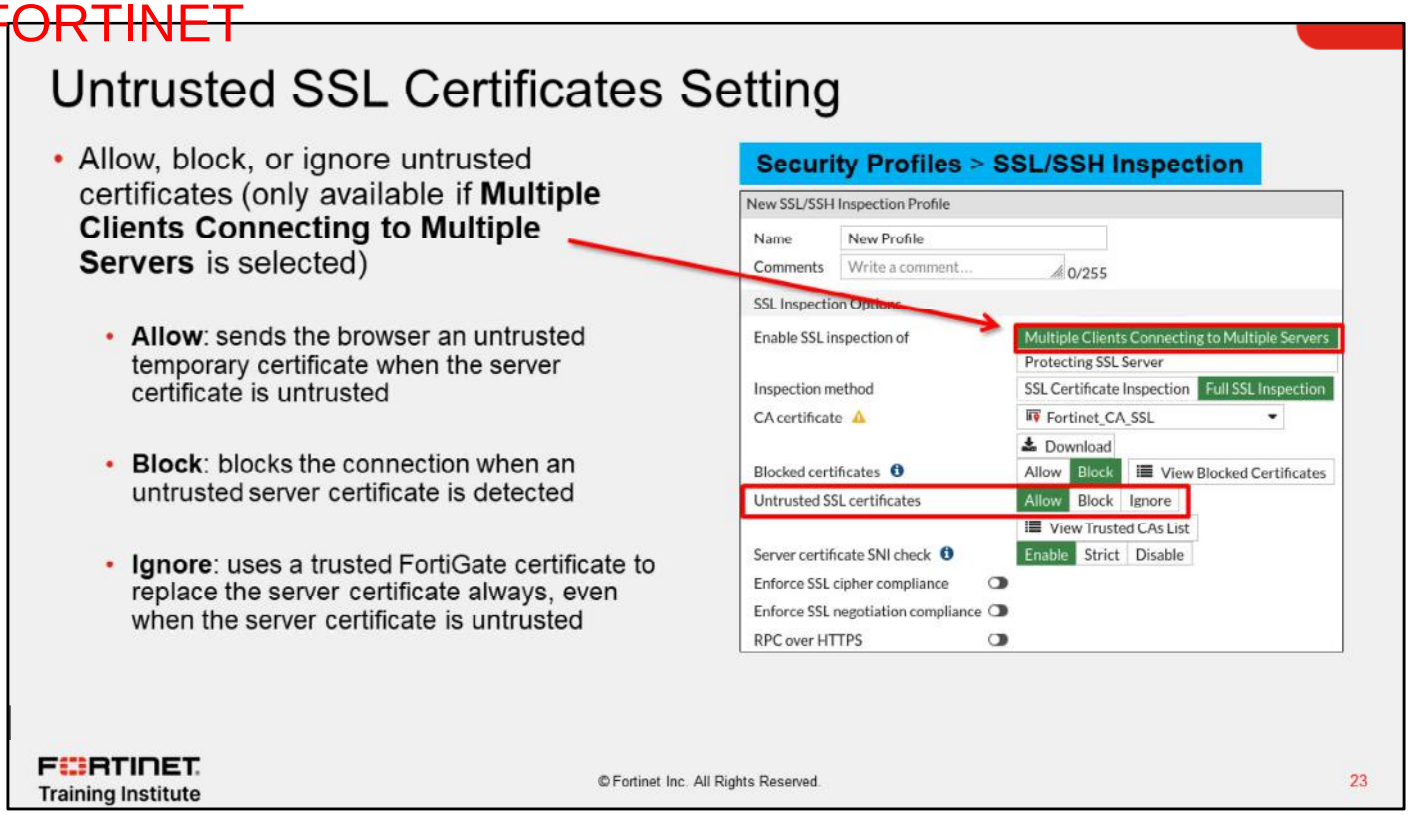

The browser presents a certificate warning when you attempt to access an HTTPS site that uses an untrusted certificate. Untrusted certificates include self-signed SSL certificates, unless the certificate is imported into the browser-trusted certificate store. FortiGate has its own configuration setting on the **SSL/SSH Inspection** page, which includes options to **Allow**, **Block**, or **Ignore** untrusted SSL certificates.

When you set the **Untrusted SSL certificates** setting to **Allow** and FortiGate detects an untrusted SSL certificate, FortiGate generates a temporary certificate signed by the built-in Fortinet_CA_Untrusted certificate. FortiGate then sends the temporary certificate to the browser, which presents a warning to the user indicating that the site is untrusted. If FortiGate receives a trusted SSL certificate, then it generates a temporary certificate signed by the built-in Fortinet_CA_SSL certificate and sends it to the browser. If the browser trusts the Fortinet_CA_SSL certificate, the browser completes the SSL handshake. Otherwise, the browser also presents a warning message informing the user that the site is untrusted. In other words, for this function to work as intended, you must import the Fortinet_CA_SSL certificate into the trusted root CA certificate store of your browser. The Fortinet_CA_Untrusted certificate must not be imported.

When the setting is set to **Block** and FortiGate receives an untrusted SSL certificate, FortiGate blocks the connection outright, and the user cannot proceed.

When the setting is set to **Ignore**, FortiGate sends the browser a temporary certificate signed by the Fortinet_CA_SSL certificate, regardless of the SSL certificate status—trusted or untrusted. FortiGate then proceeds to establish SSL sessions.

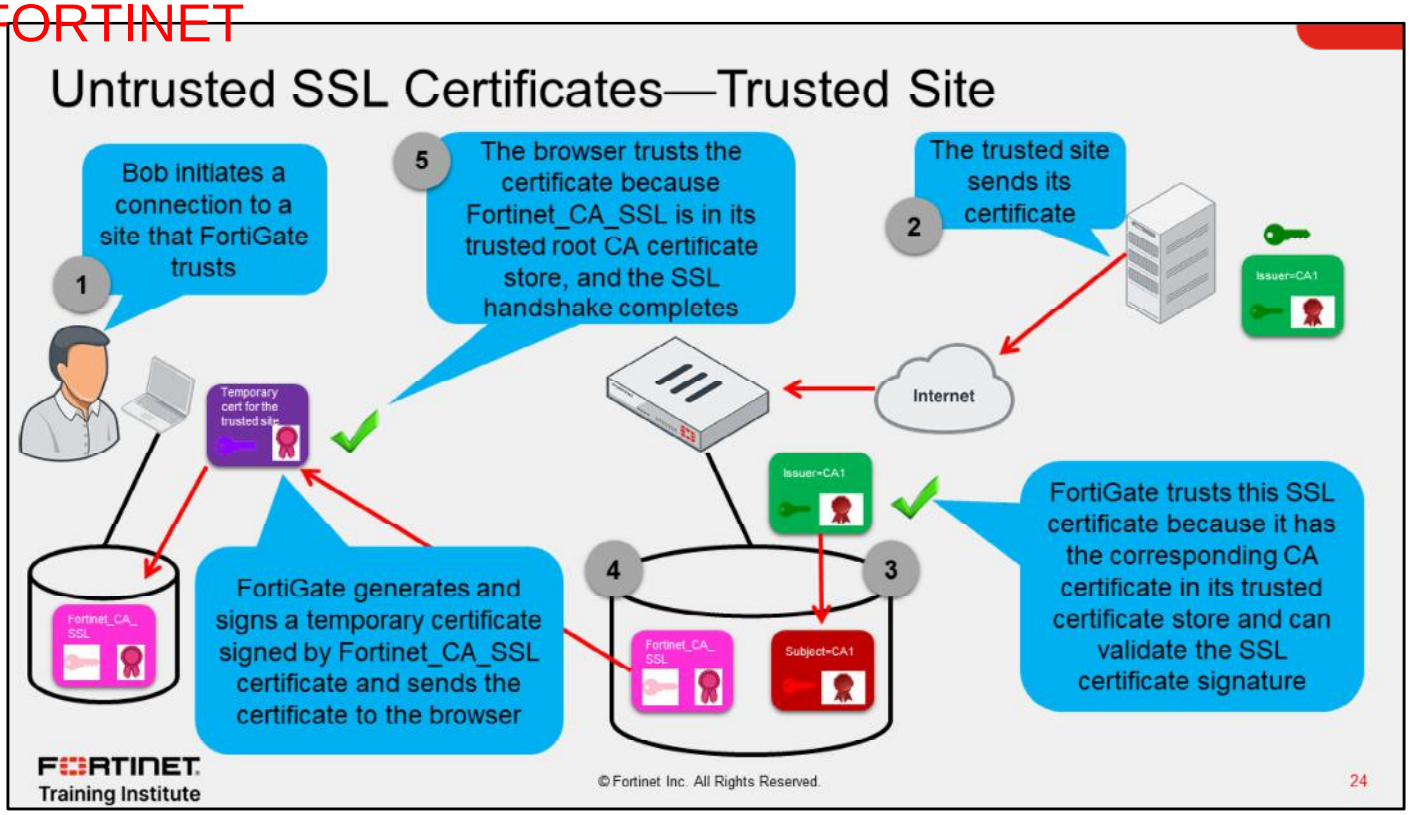

The scenario shown on this slide describes how FortiGate handles a trusted external site regardless of the **Untrusted SSL Certificate** setting.

In step 1, the browser initiates a connection with an external site that is trusted by FortiGate. In step 2, the trusted server sends its SSL certificate to FortiGate. In step 3, FortiGate trusts the certificate because it has the corresponding CA certificate in its trusted certificate store. FortiGate can validate the signature on the SSL certificate. In step 4, because FortiGate trusts the SSL certificate, it generates a temporary certificate signed by the Fortinet_CA_SSL certificate. FortiGate sends the temporary certificate to the browser. Finally, in step 5, the browser trusts the temporary certificate because the Fortinet_CA_SSL certificate is in its trusted root CA store. After the browser finishes validating the certificate, it completes the SSL handshake with FortiGate. Next, FortiGate continues the SSL handshake with the trusted server.

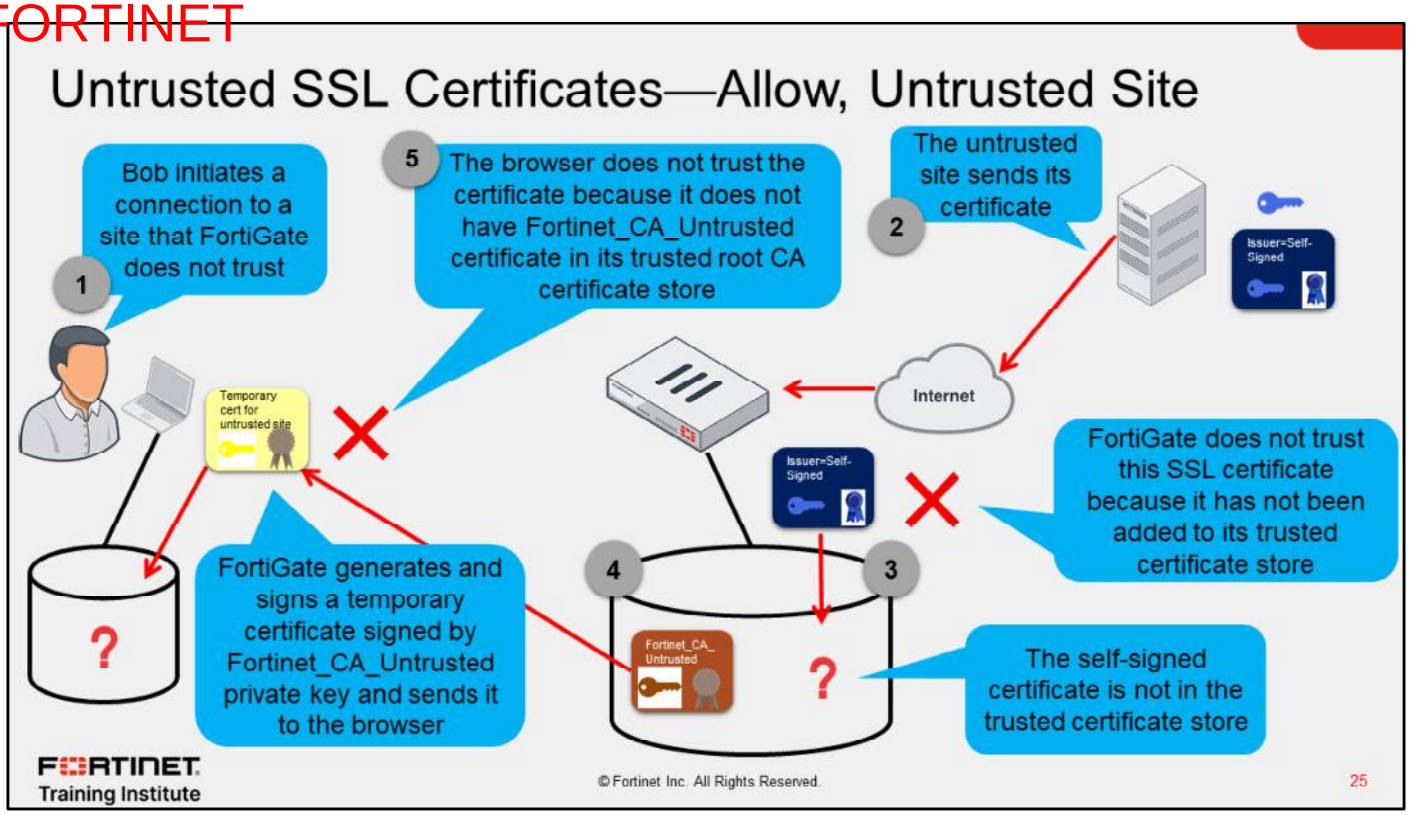

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Allow**.

In step 1, the browser initiates a connection with an external site that is *not* trusted by FortiGate. In step 2, the untrusted server sends its self-signed SSL certificate to FortiGate. In step 3, FortiGate does not find a copy of the certificate in its trusted certificate store and, therefore, does not trust the SSL certificate. In step 4, because FortiGate does not trust the SSL certificate, it generates a temporary certificate signed by the Fortinet_CA_Untrusted certificate. This temporary certificate is sent to the browser. In step 5, the browser does not trust the temporary certificate because it does not have the Fortinet_CA_Untrusted certificate in its trusted root CA store. The browser displays a warning alerting the user that the certificate is untrusted. If the user decides to ignore the warning and proceed, the browser completes the SSL handshake with FortiGate. Next, FortiGate continues the SSL handshake with the untrusted server.

The user may have the option to write this temporary certificate to the browser trusted certificate store. However, this has no impact in the future. The next time the user connects to the same untrusted site, a new temporary certificate is produced for the session.

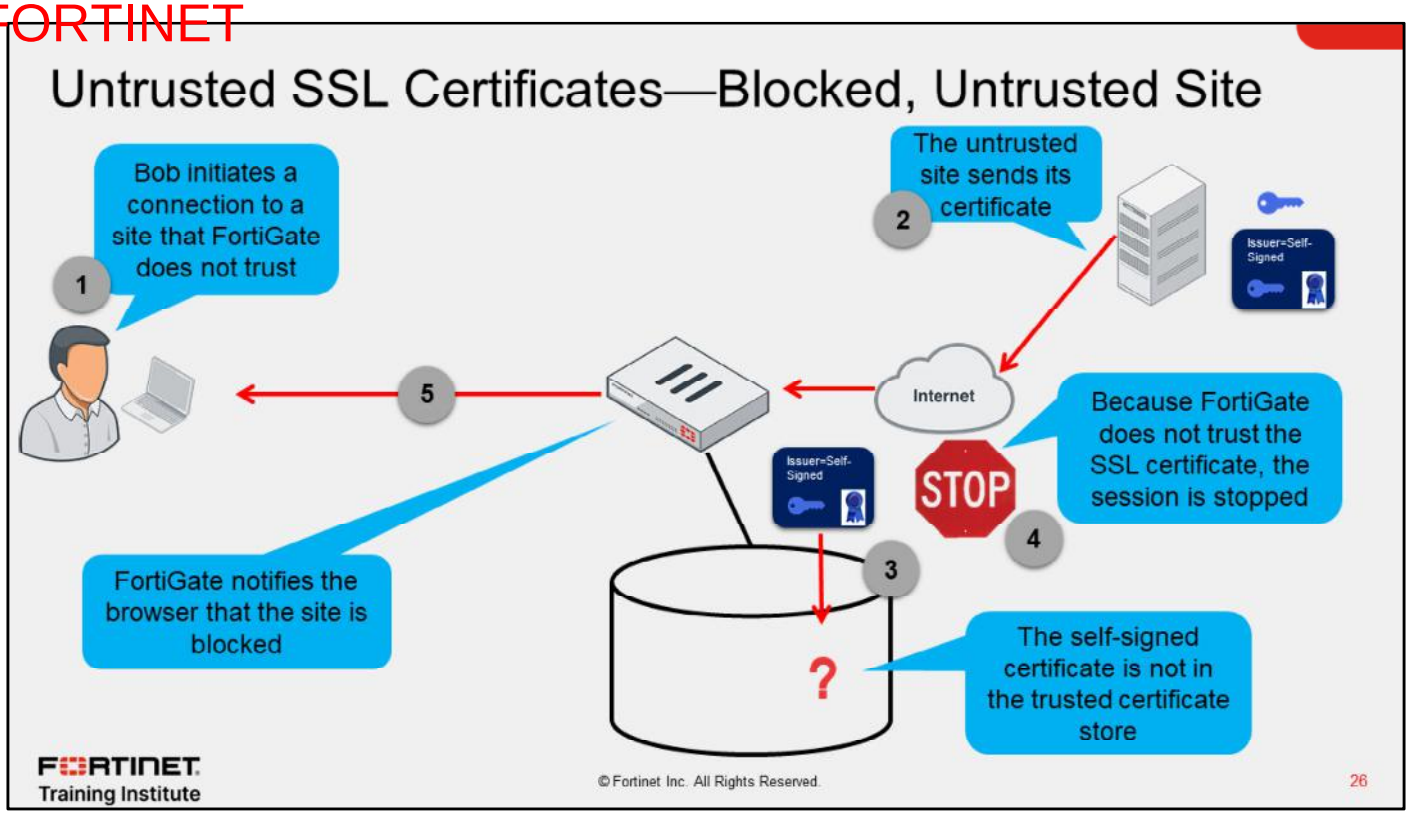

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Block**.

In step 1, the browser initiates a connection with an external site that is *not* trusted by FortiGate. In step 2, the untrusted server sends its self-signed SSL certificate to FortiGate. In step 3, FortiGate does not find the certificate in its trusted certificate store and, therefore, does not trust the SSL certificate. In step 4, because FortiGate does not trust the SSL certificate, it stops the session. In step 5, FortiGate notifies the browser that the site is blocked.

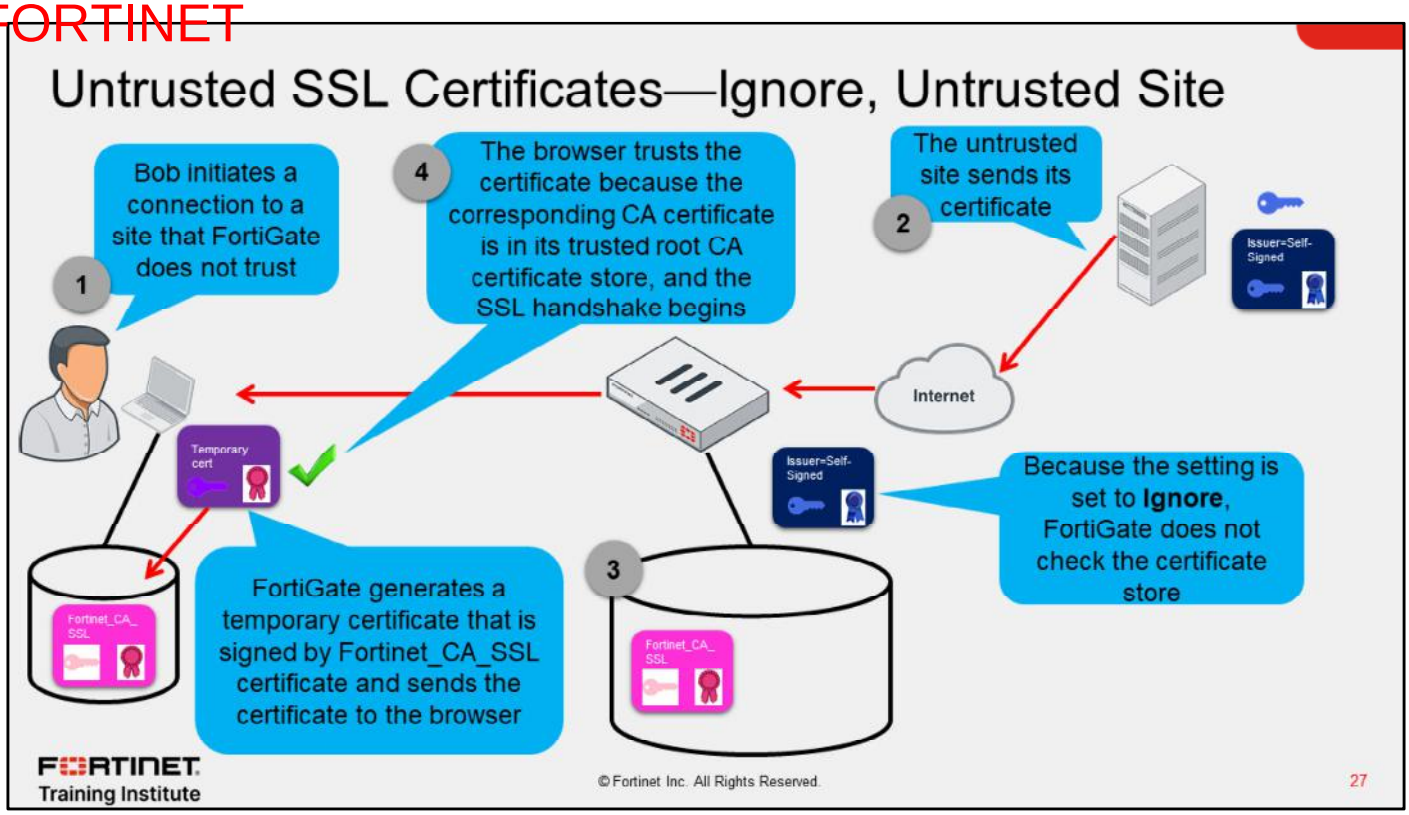

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Ignore**.

In step 1, the browser initiates a connection with an external site that is *not* trusted by FortiGate. In step 2, the untrusted server sends its self-signed SSL certificate to FortiGate. Because the setting is set to **Ignore**, FortiGate does not check the certificate store. In step 3, FortiGate generates a temporary certificate signed by Fortinet_CA_SSL certificate, and sends the certificate to the browser. In step 4, the browser trusts the certificate because Fortinet_CA_SSL certificate is in its trusted root CA store. After the browser finishes checking the certificate, it completes the SSL handshake with FortiGate. Next, FortiGate continues the SSL handshake with the trusted server.

A connection to a trusted site is handled the same way.

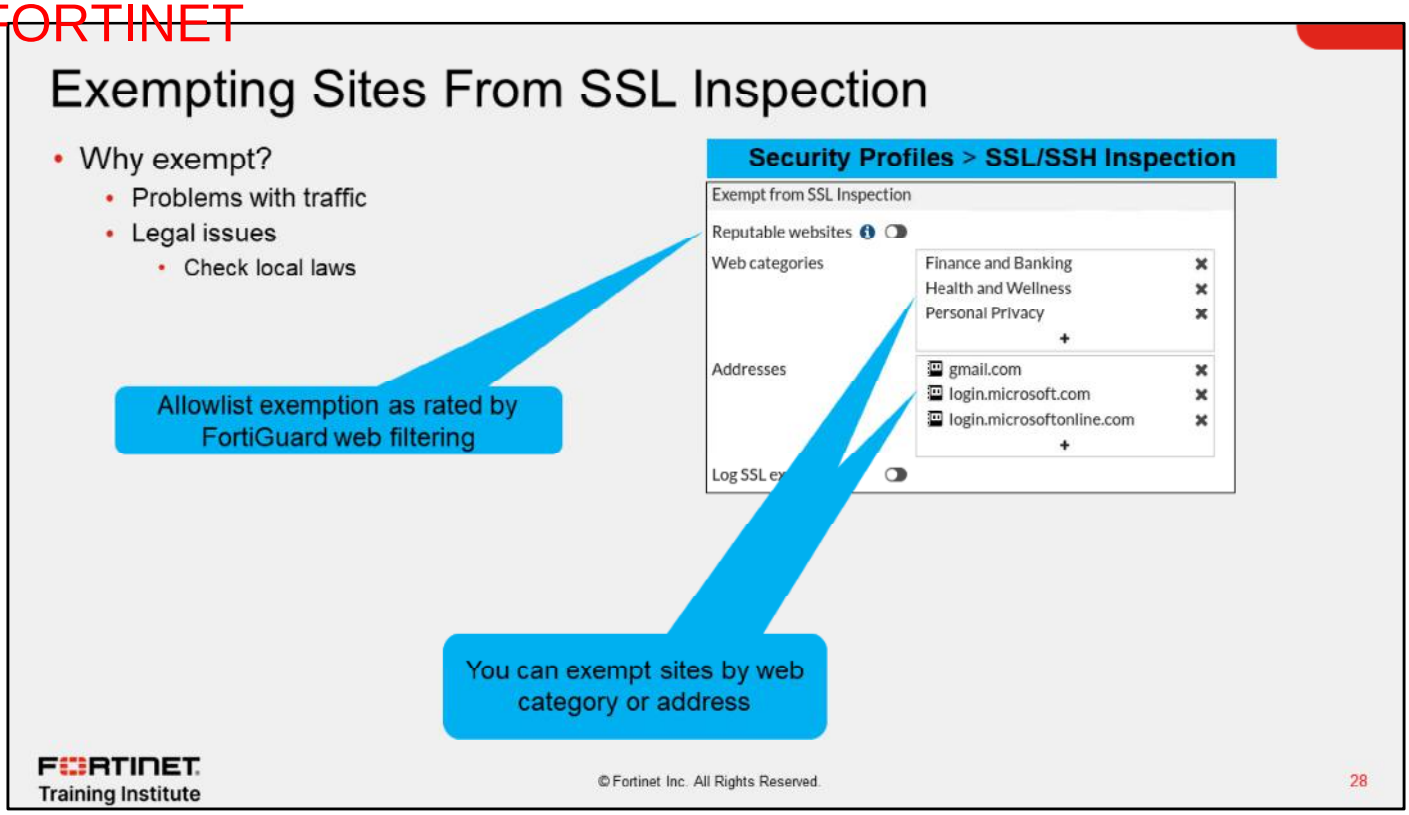

Within the full SSL inspection profile, you can also specify which SSL sites, if any, you want to exempt from SSL inspection. You may need to exempt traffic from SSL inspection if it is causing problems with traffic, or for legal reasons.

Performing SSL inspection on a site that is enabled with HTTP public key pinning (HPKP), for example, can cause problems with traffic. Remember, the only way for FortiGate to inspect encrypted traffic is to intercept the certificate coming from the server, and generate a temporary one. After FortiGate presents the temporary SSL certificate, browsers that use HPKP refuse to proceed. The SSL inspection profile, therefore, allows you to exempt specific traffic.

Laws protecting privacy might be another reason to bypass SSL inspection. For example, in some countries, it is illegal to inspect SSL bank-related traffic. Configuring an exemption for sites is simpler than setting up firewall policies for each individual bank. You can exempt sites based on their web category, such as finance or banking, or you can exempt them based on their address. Alternatively, you can enable **Reputable websites**, which excludes an allowlist of reputable domain names maintained by FortiGuard from full SSL inspection. This list is periodically updated and downloaded to FortiGate devices through FortiGuard.

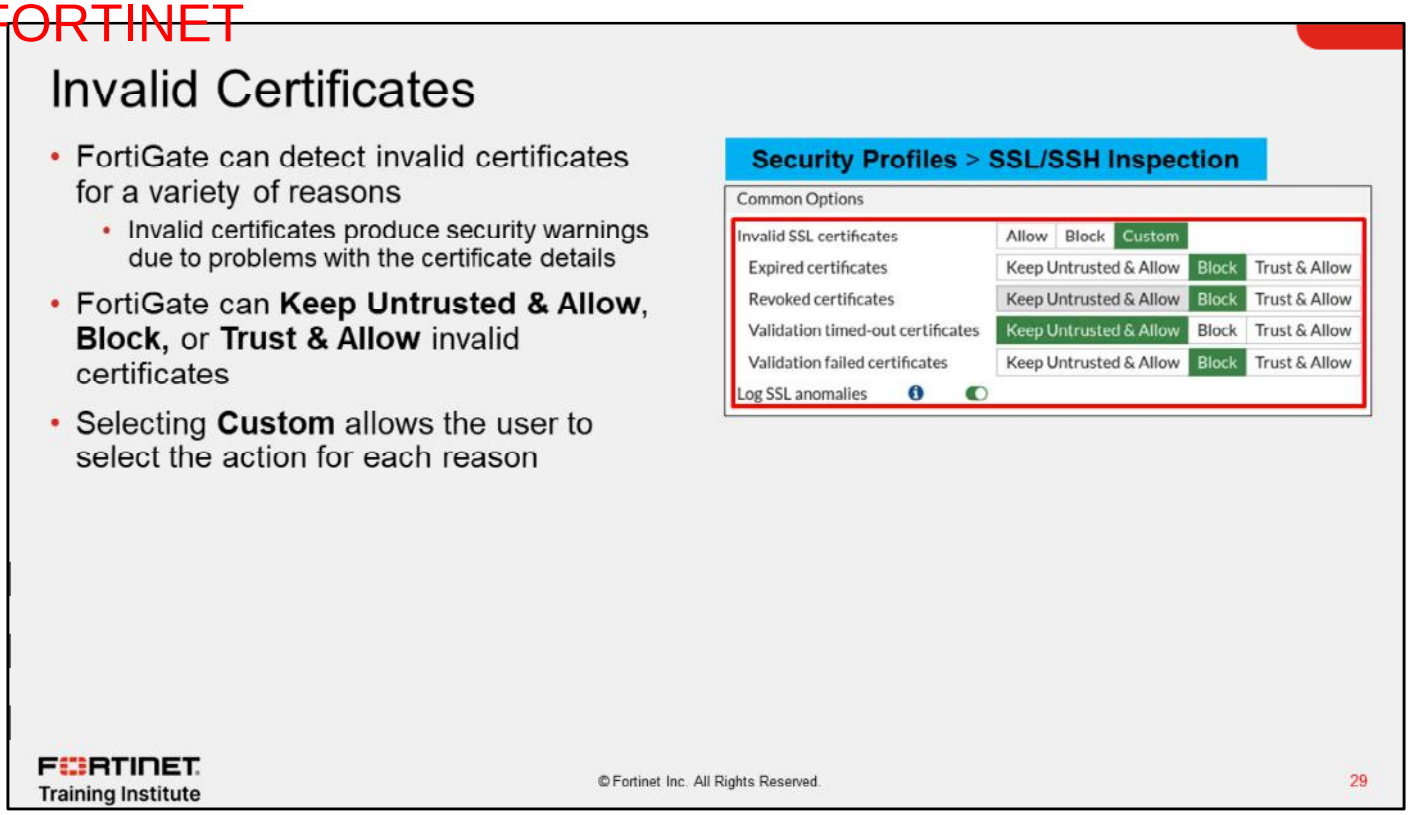

FortiGate can detect certificates that are invalid for the following reasons:

- Expired: The certificate is expired.
- Revoked: The certificate has been revoked based on CRL or OCSP information.
- Validation timeout: The certificate could not be validated because of a communication timeout.
- Validation failed: The certificate could not be validated because of a communication error.

When a certificate fails for any of the reasons above, you can configure any of the following actions:

- **Keep untrusted & Allow**: FortiGate allows the website and lets the browser decide the action to take. FortiGate takes the certificate as *trusted*.
- **Block**: FortiGate blocks the content of the site.
- **Trust & Allow**: FortiGate allows the website and takes the certificate as *trusted*.

The certificate check feature can be broken down into two major checks, which are done in parallel:

- FortiGate checks if the certificate is invalid because of the four reasons described on this slide.
- FortiGate performs certificate chain validation based on the CA certificates installed locally and the certificates presented by the SSL server. This is described in this lesson.

Based on the actions configured and the check results, FortiGate presents the certificate as either trusted (signed by Fortinet_CA_SSL) or untrusted (signed by Fortinet_CA_Untrusted), and either allows the content or blocks it. You can also track certificate anomalies by enabling the **Log SSL anomalies** option.

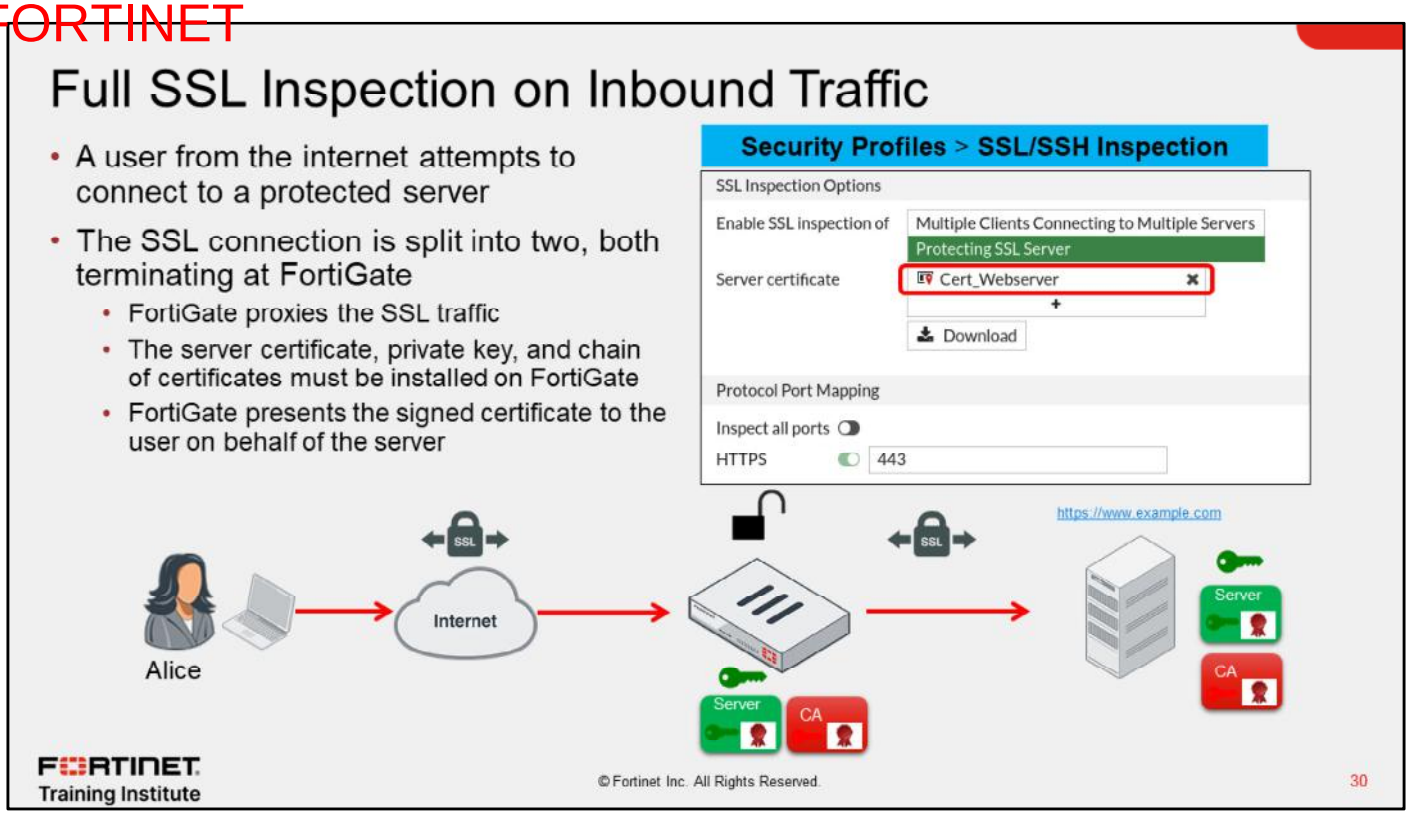

In the example shown on this slide, FortiGate is protecting a web server. This is the second configuration option for full SSL inspection. When configuring the SSL inspection profile for this server, you must select **Protecting SSL Server**, import the server key pair to FortiGate, and then select the certificate from the **Server Certificate** drop-down list.

When Alice attempts to connect to the protected server, FortiGate becomes a surrogate web server by establishing the secure connection with the client using the server key pair. FortiGate also establishes a secure connection with the server, but acting as a client. This configuration allows FortiGate to decrypt the data from either direction, scan it, and if it is clean, re-encrypt it and send it to the intended recipient.

You must install the server certificate and private key plus the chain of certificates required to build the chain of trust. FortiGate sends the chain of certificates to the browser for this purpose.

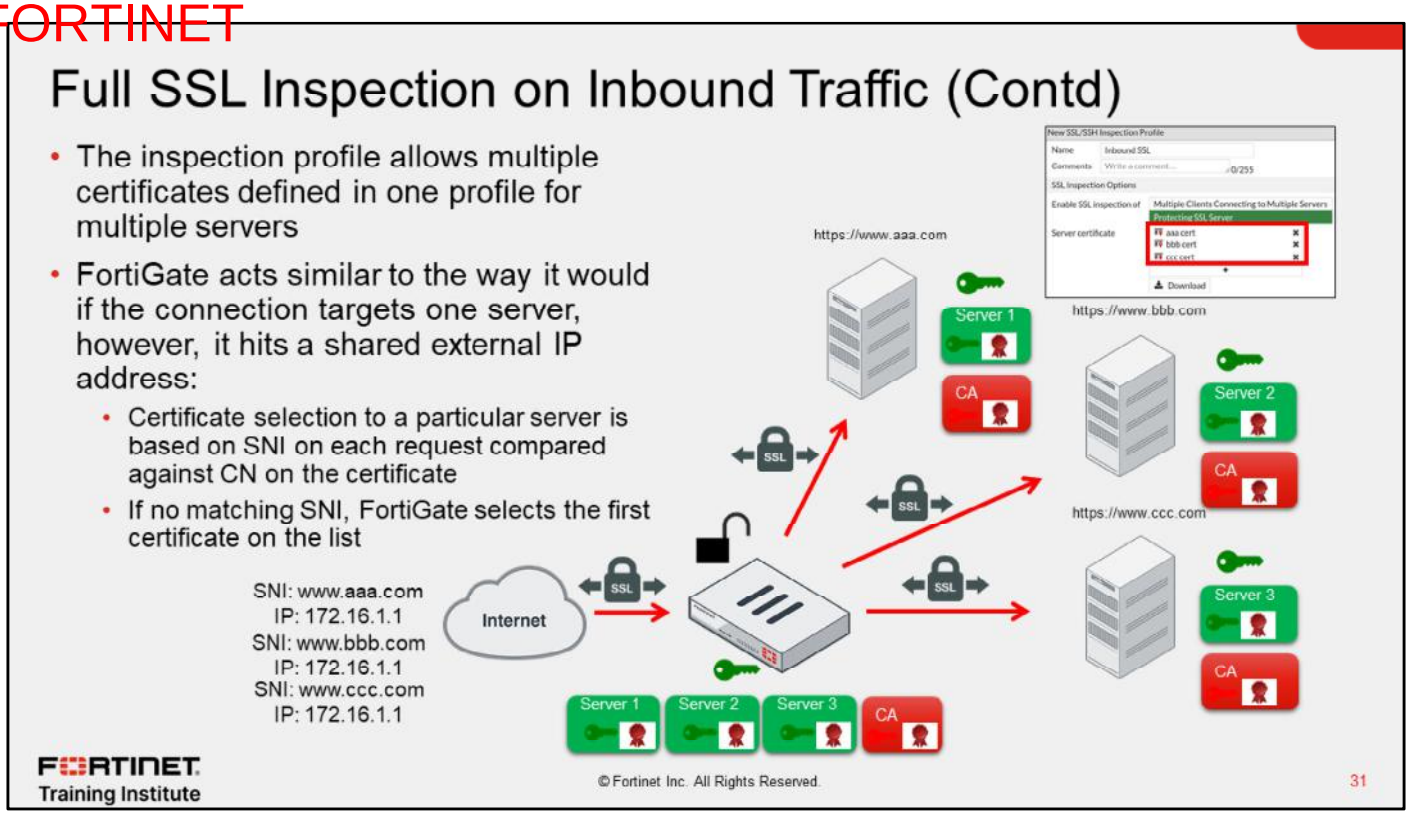

By creating a full SSL inspection profile on inbound traffic, you can configure the profile to use multiple web sites if they are approachable by the same external IP address. When FortiGate receives client and server hello messages, it selects the certificate to perform the full SSL inspection based on server name indication (SNI) value against the common name (CN) on the certificate part of the inspection profile. If a certificate CN matches the SNI on the request, FortiGate then selects this certificate to replace the original certificate and uses it to inspect the traffic.

If the SNI does not match the CN in the certificate list in the SSL profile, then FortiGate selects the first server certificate in the list.

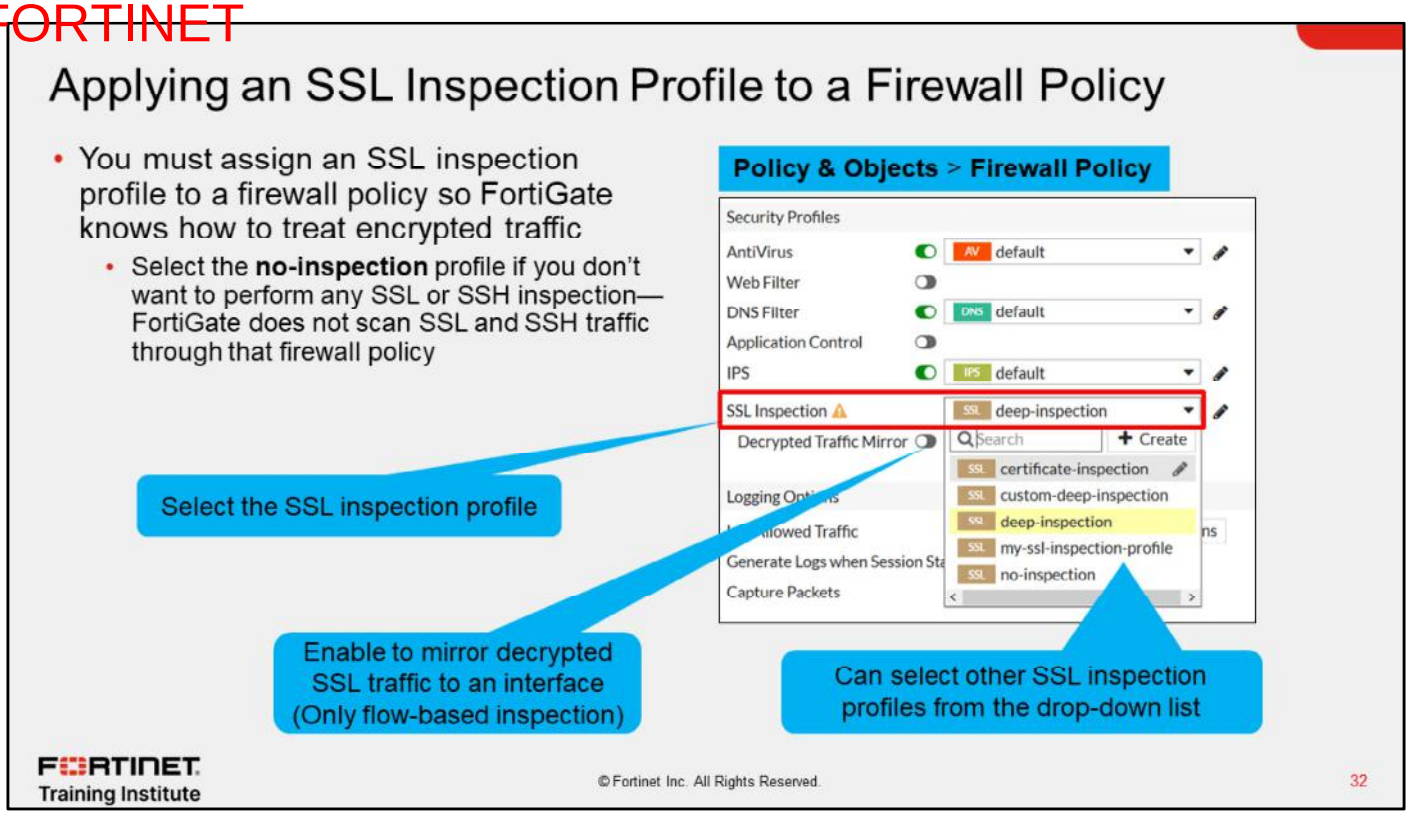

After you create and configure an SSL inspection profile, you must assign it to a firewall policy so FortiGate knows how to inspect encrypted traffic. Most of the internet traffic is being encrypted nowadays. For this reason, you usually want to enable SSL inspection to protect your network from security threats transported over encrypted traffic. If you don't want to enable SSL or SSH inspection, select the **no-inspection** profile from the drop-down list. If SSL inspection is not enabled in a policy, FortiGate will not scan SSL or SSH encrypted traffic matching that policy.

If you select a profile with full SSL inspection enabled, the option **Decrypted Traffic Mirror** appears. Enable this option if you want FortiGate to send a copy of the decrypted SSL traffic to an interface, it only works with flow-based inspection. When you enable **Decrypted Traffic Mirror**, FortiGate displays a window with the terms of use for this feature. The user must agree with the terms before they can use the feature.

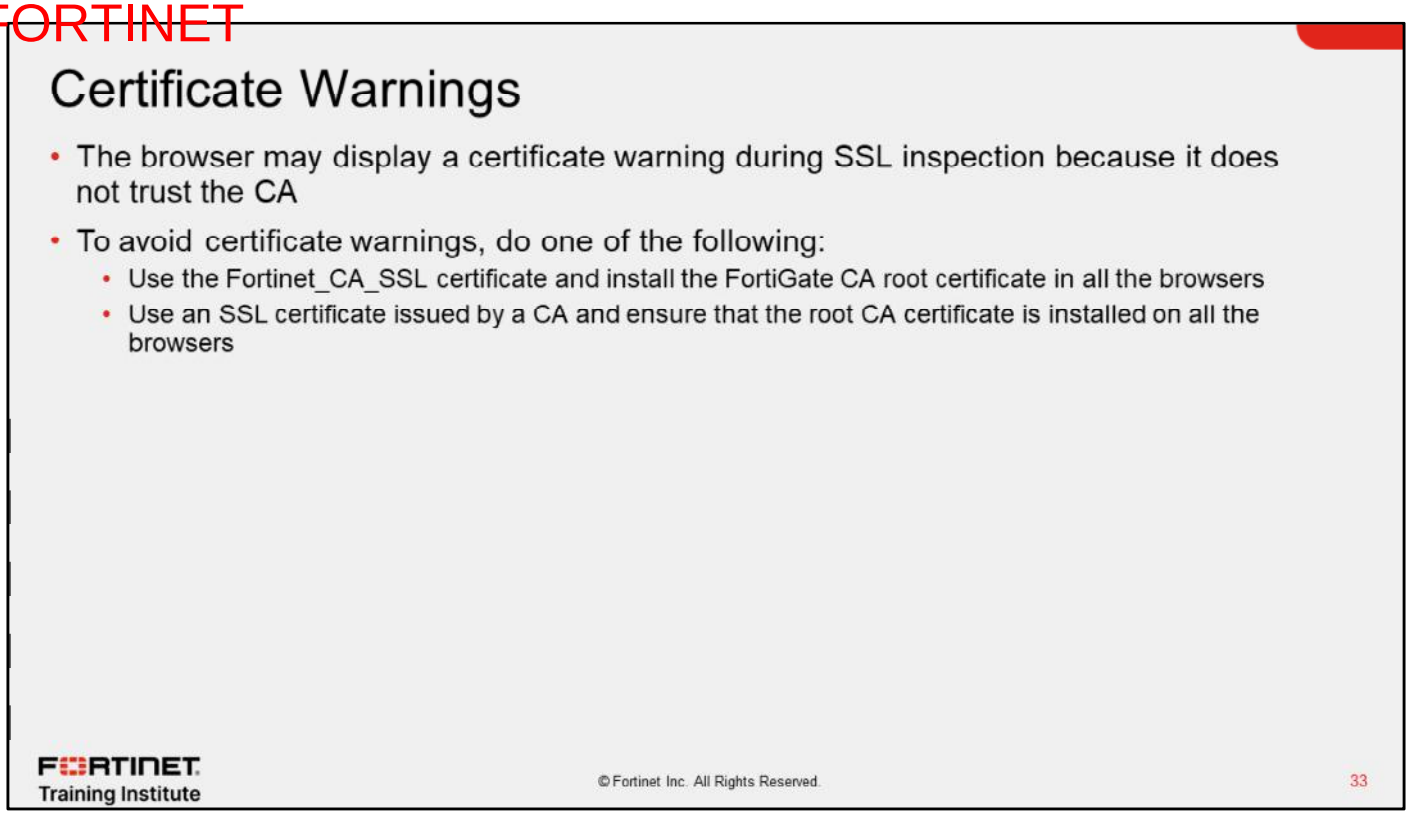

When doing full SSL inspection using the FortiGate self-signed CA, your browser displays a certificate warning each time you connect to an HTTPS site. This is because the browser is receiving certificates signed by FortiGate, which is a CA it does not know and trust. The browser also displays a certificate warning when performing SSL certificate inspection and an HTTPS website is blocked by FortiGate. Because FortiGate needs to present a replacement message to the browser, FortiGate performs MITM and signs the certificate with its self-signed CA as well.

You can avoid this warning by doing one of the following:

- Download the Fortinet_CA_SSL certificate and install it on all the workstations as a trusted root authority.
- Use an SSL certificate issued by a CA and ensure the certificate is installed in the necessary browsers.

You must install the SSL certificate on FortiGate and configure the device to use that certificate for SSL inspection. If the SSL certificate is signed by a subordinate CA, ensure that the entire chain of certificates—from the SSL certificate to the root CA certificate—is installed on FortiGate. Verify that the root CA is installed on all client browsers. This is required for trust purposes. Because FortiGate sends the chain of certificates to the browser during the SSL handshake, you do not have to install the intermediate CA certificates on the browsers.

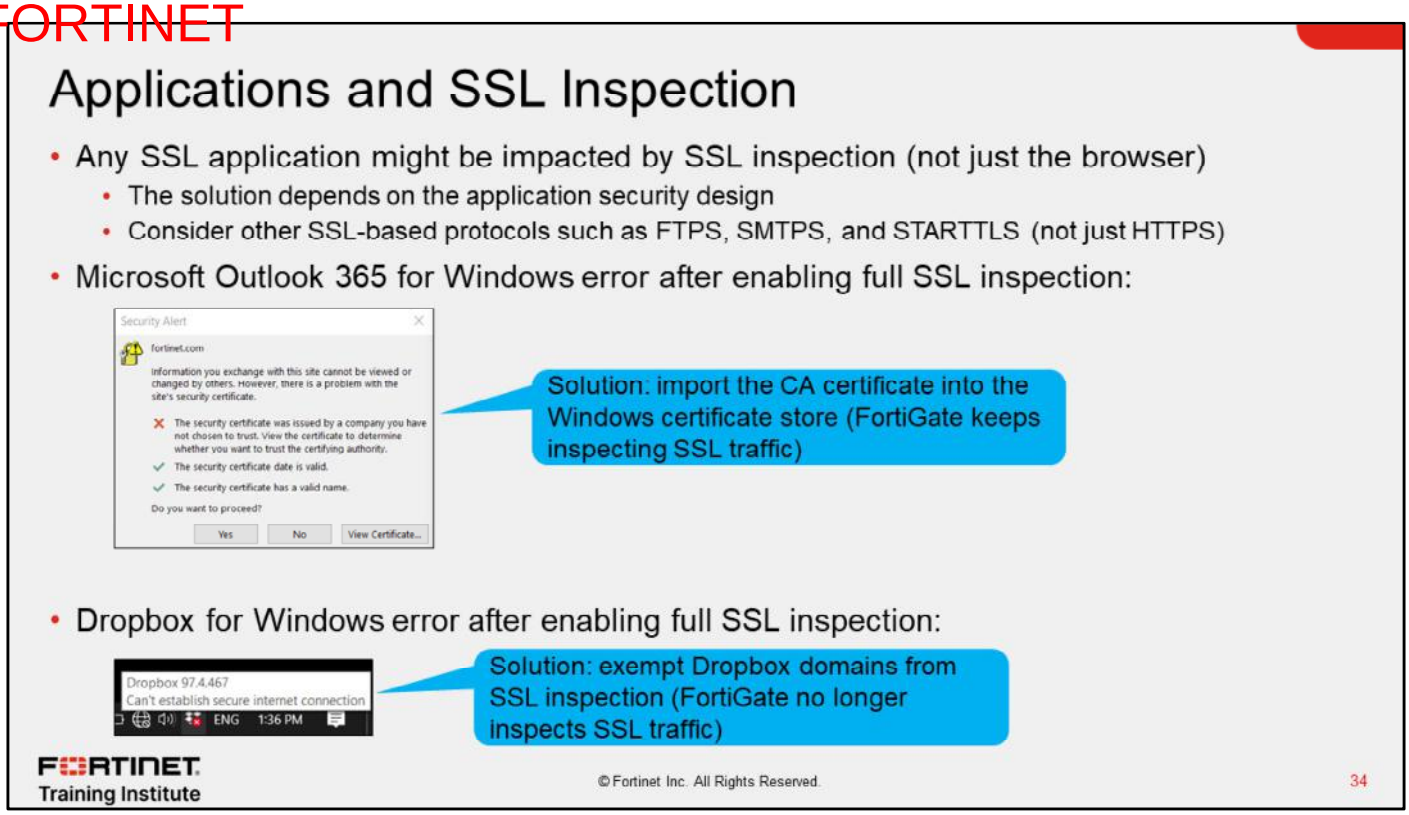

More and more applications are using SSL to securely exchange data over the internet. While most of the content in this lesson centers around the operation and impact of SSL inspection on browsers, the same applies to other applications using SSL as well. After all, the browser is just another application using SSL on your device.

For this reason, when you enable SSL inspection on FortiGate, you need to consider the potential impact on your SSL-based applications. For example, Microsoft Outlook 365 for Windows reports a certificate error when you enable full SSL inspection because the CA certificate used by FortiGate is not trusted. To solve this issue, you can import the CA certificate into your Windows certificate store as a trusted root certificate authority. Because Microsoft Outlook 365 trusts the certificates in the Windows certificate store, then the application won't report the certificate error anymore. Another option is to exempt your Microsoft Exchange server addresses from SSL inspection. While this prevents the certificate error, you are no longer performing SSL inspection on email traffic.

There are other applications that have built-in extra security checks that prevent MITM attacks, such as HPKP or certificate pinning. For example, Dropbox uses certificate pinning to ensure that no SSL inspection is possible on user traffic. As a result, when you enable full SSL inspection on FortiGate, your Dropbox client stops working and reports that it can't establish a secure connection. In the case of Dropbox, the only way to solve the connection error is by exempting the domains Dropbox connects to from SSL inspection.

In addition, remember that SSL is leveraged by different protocols, not just HTTP. For example, there are other SSL-based protocols such as FTPS, POP3S, SMTPS, STARTTLS, LDAPS, and SIP TLS. If you have an application using any of these SSL-based protocols, and you have turned on SSL inspection along with a security profile that inspects those protocols, then the applications may report an SSL or certificate error. The solution depends on the security measures adopted by the application.

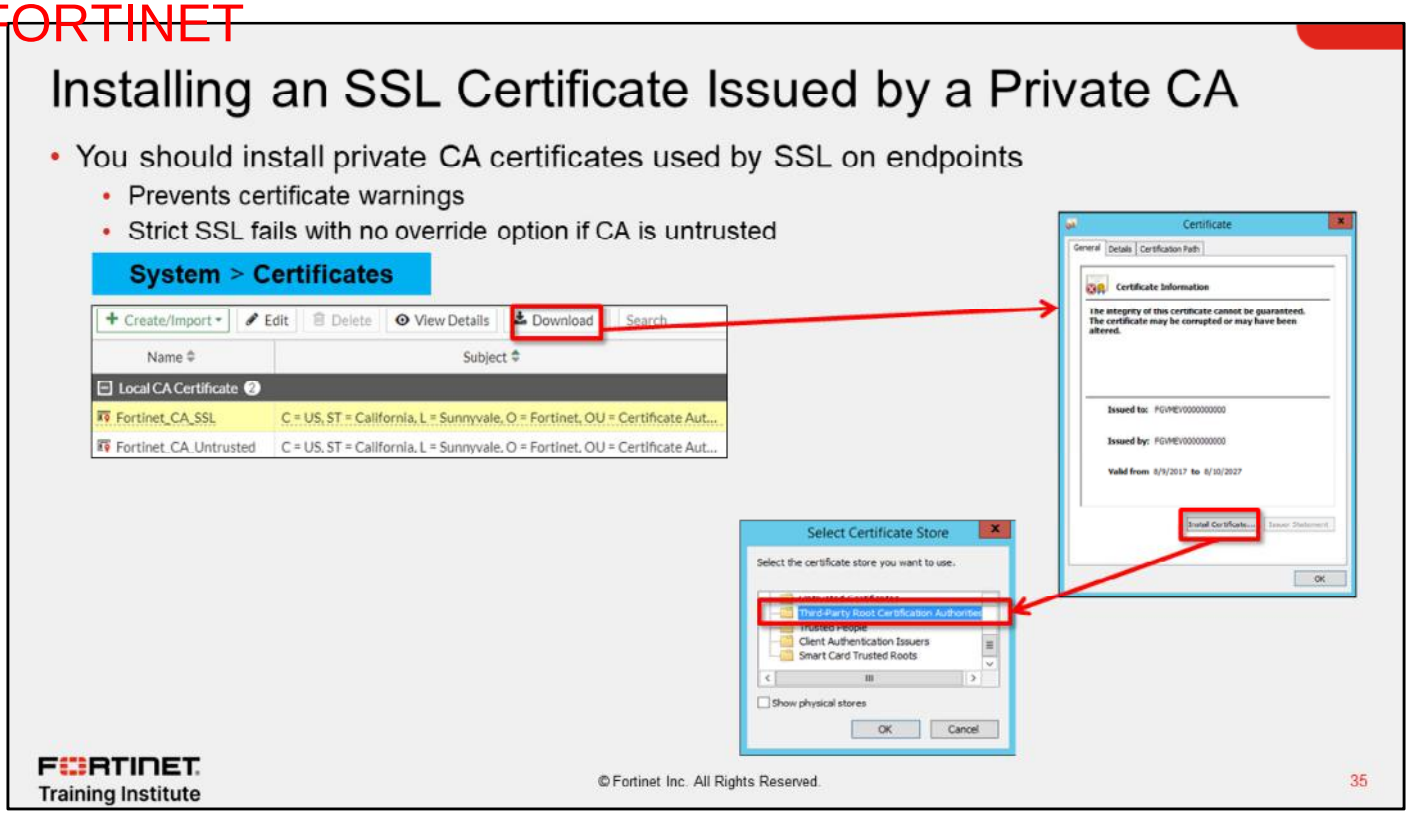

If you are using an SSL certificate issued by a private CA, you must install the CA certificate in the list of trusted CAs. If you fail to do this, a warning message appears in your web browser any time you access an HTTPS website. Encrypted communications might also fail, simply because the CA that issued and signed the certificate is untrusted.

After you download the SSL certificate from FortiGate, you can install it on any web browser or operating system. Not all browsers use the same certificate repository. For example, Firefox uses its own repository, while Internet Explorer and Chrome store certificates in a system-wide repository. In order to prevent certificate warnings, you must install the SSL certificate as a trusted root CA.

When you install the certificate, make sure that you save it to the certificate store for root authorities.

Knowledge Check

1. Which certificate extension and value is required in the FortiGate CA certificate in order to enable full SSL inspection?
   - A. CRL DP=ca_arl.arl
   - ✓ B. cA=True

2. Which configuration requires FortiGate to act as a CA for full SSL inspection?
   - ✓ A. Multiple clients connecting to multiple servers
   - B. Protecting the SSL server

FORTINET Training Institute

36

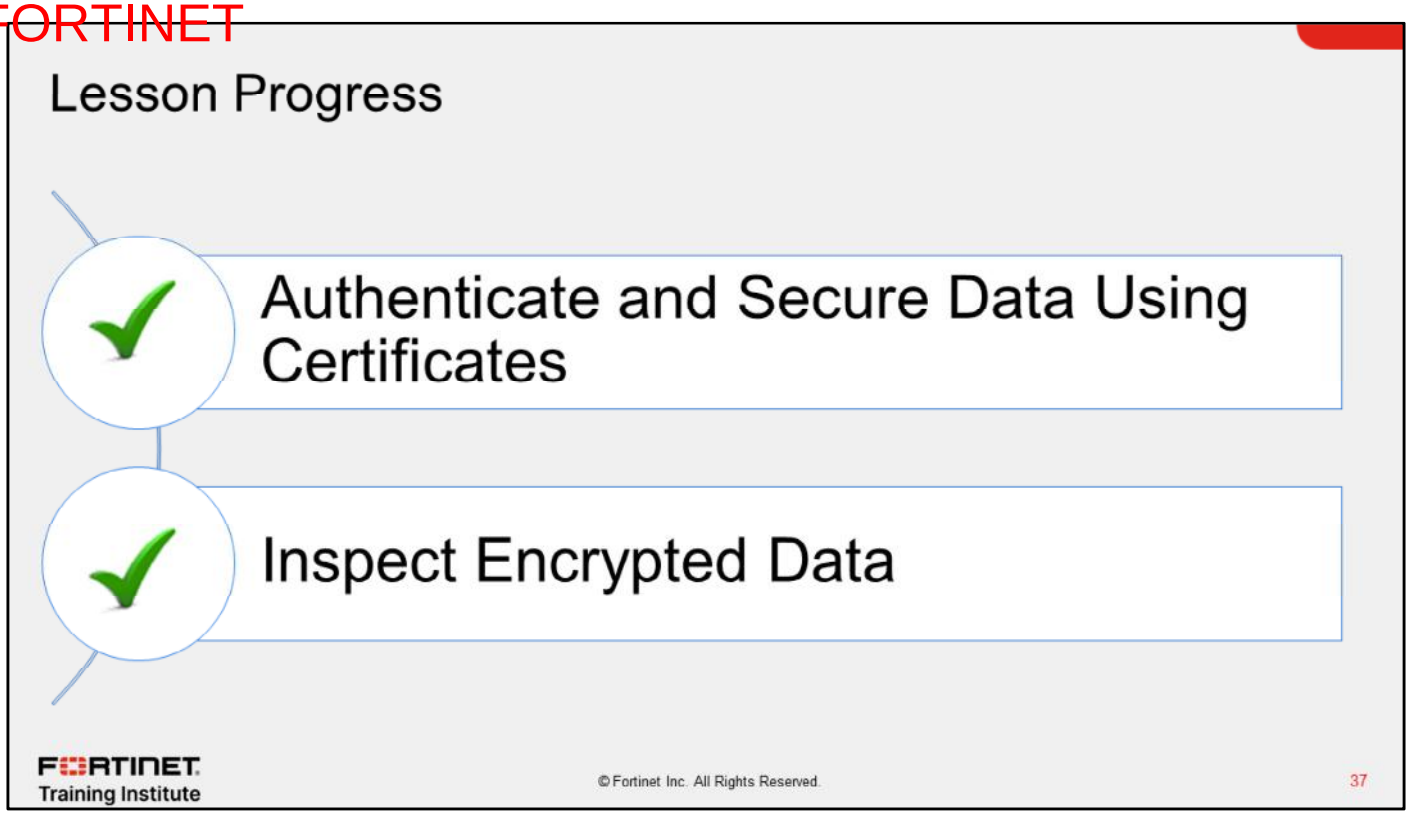

Congratulations! You have completed this lesson.

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

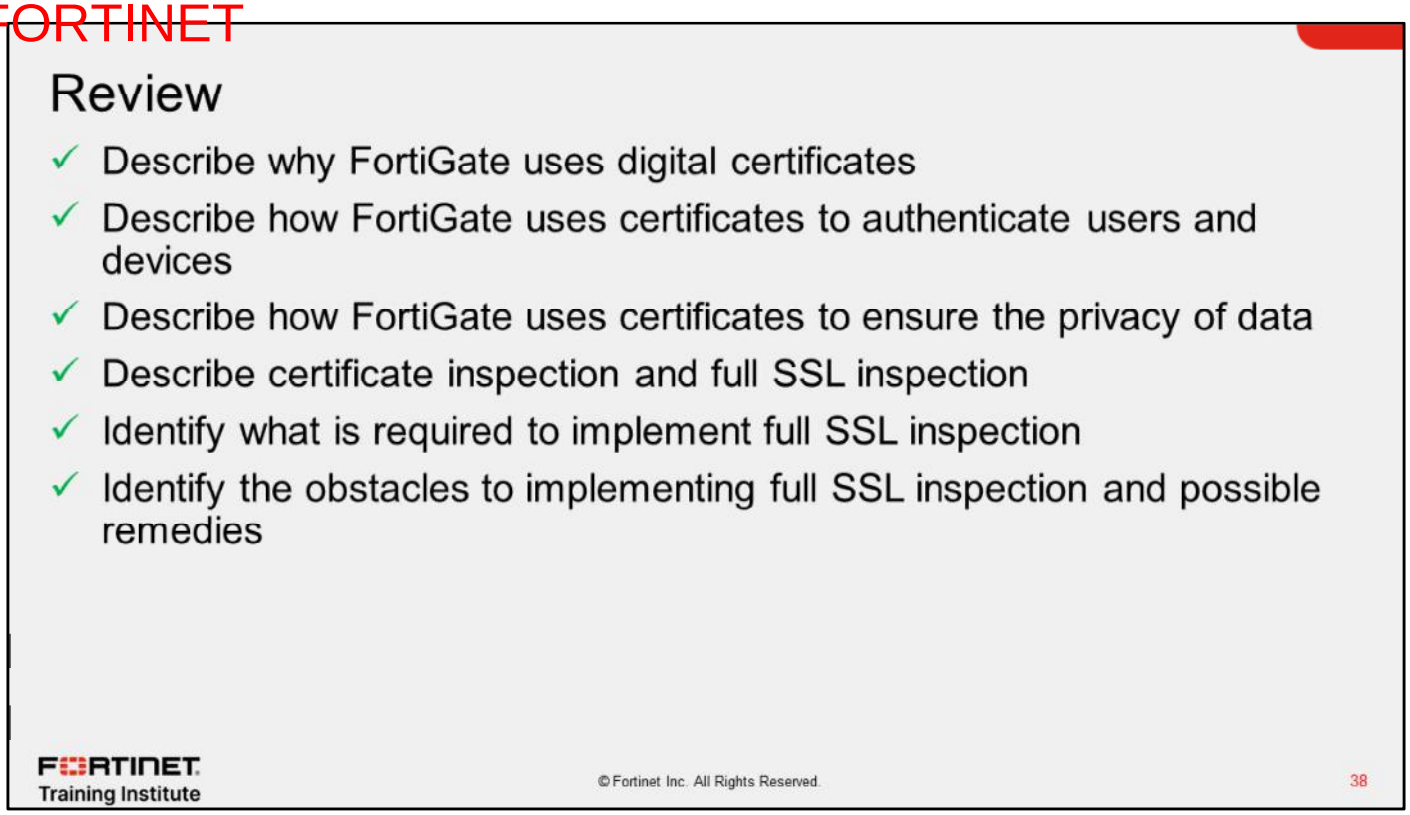

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how FortiGate uses certificates, and how to manage and work with certificates in your network.

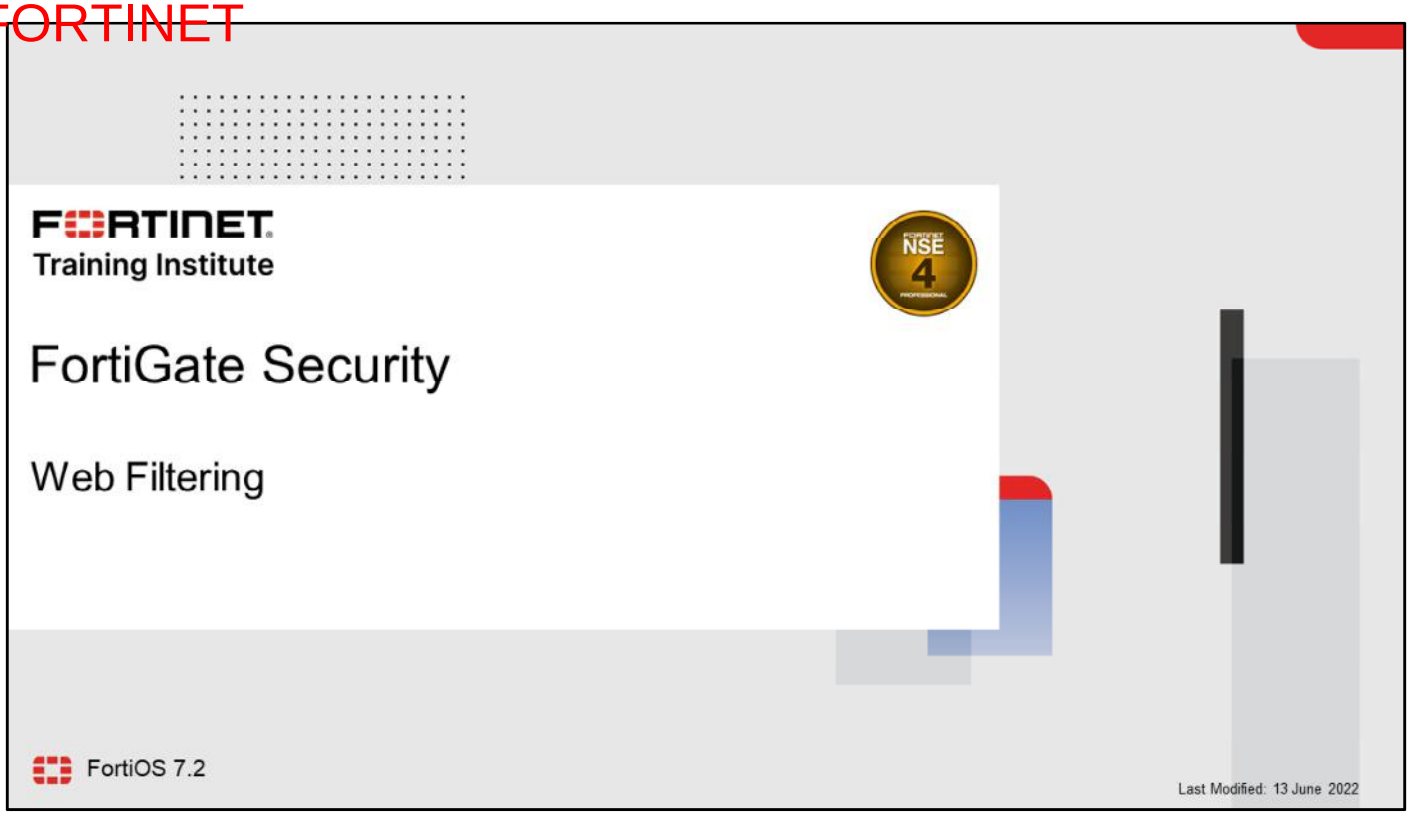

In this lesson, you will learn how to configure web filtering on FortiGate to control web traffic in your network.

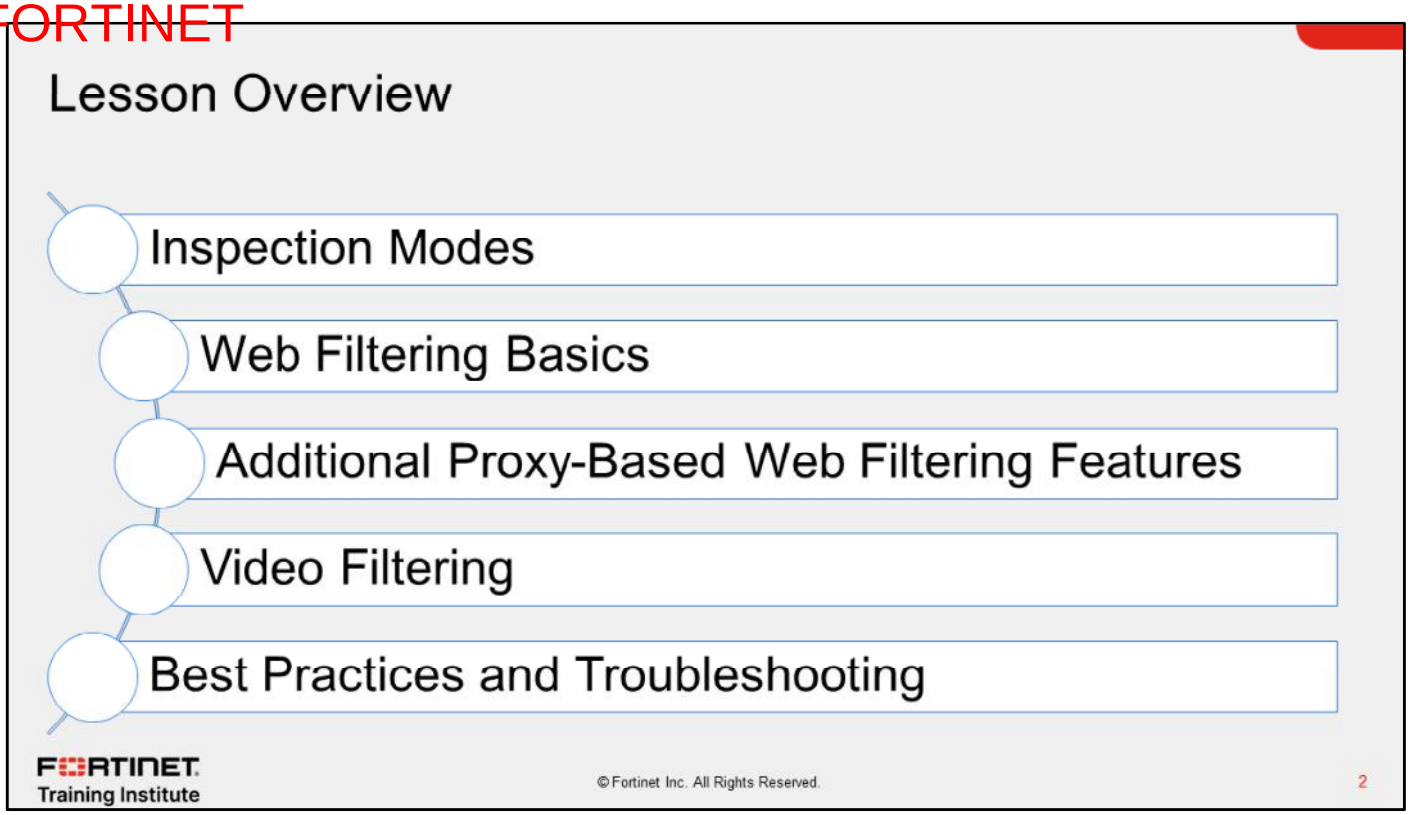

In this lesson, you will learn about the topics shown on this slide.

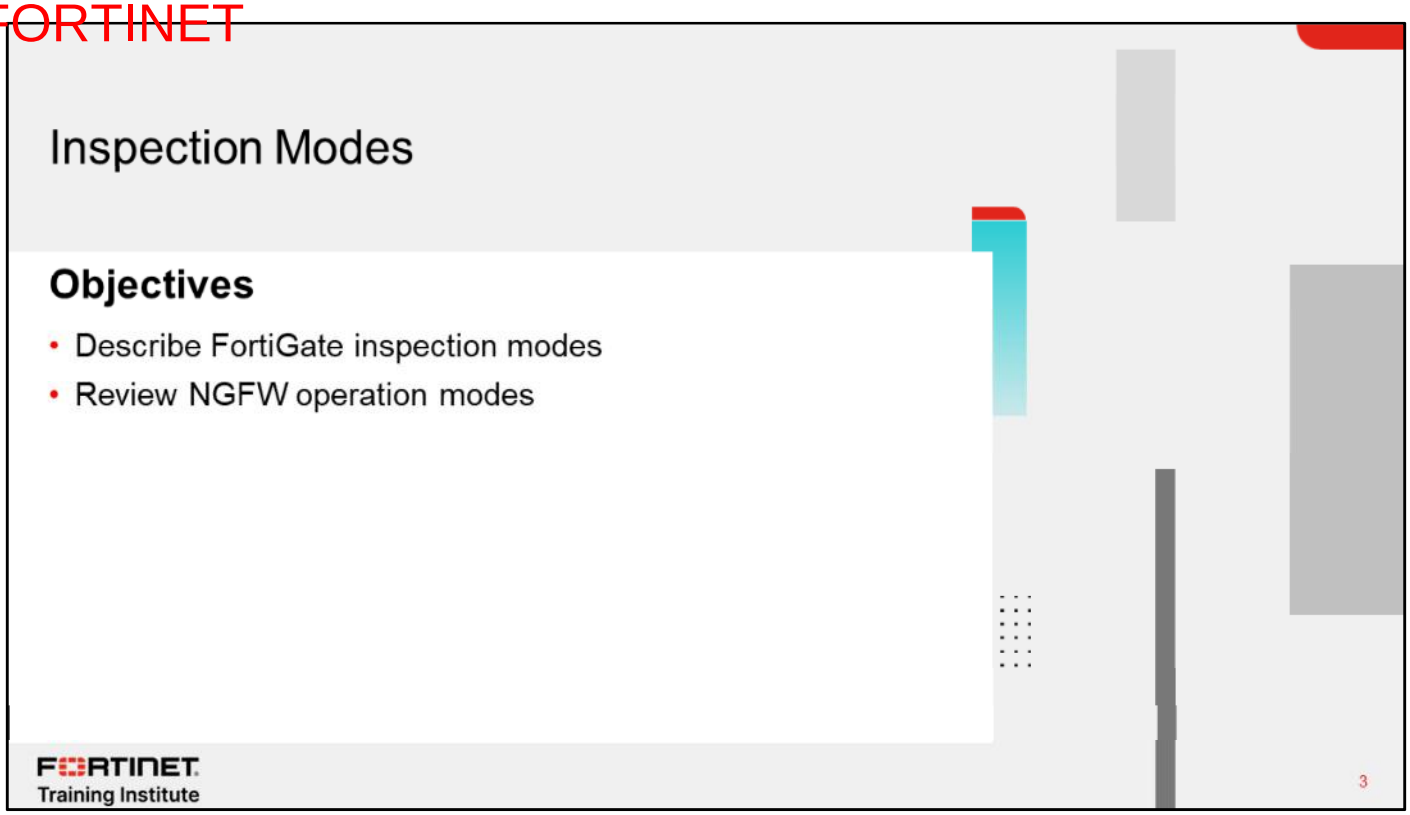

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding inspection modes, you will be able to implement the appropriate inspection modes to support the desired security profiles.

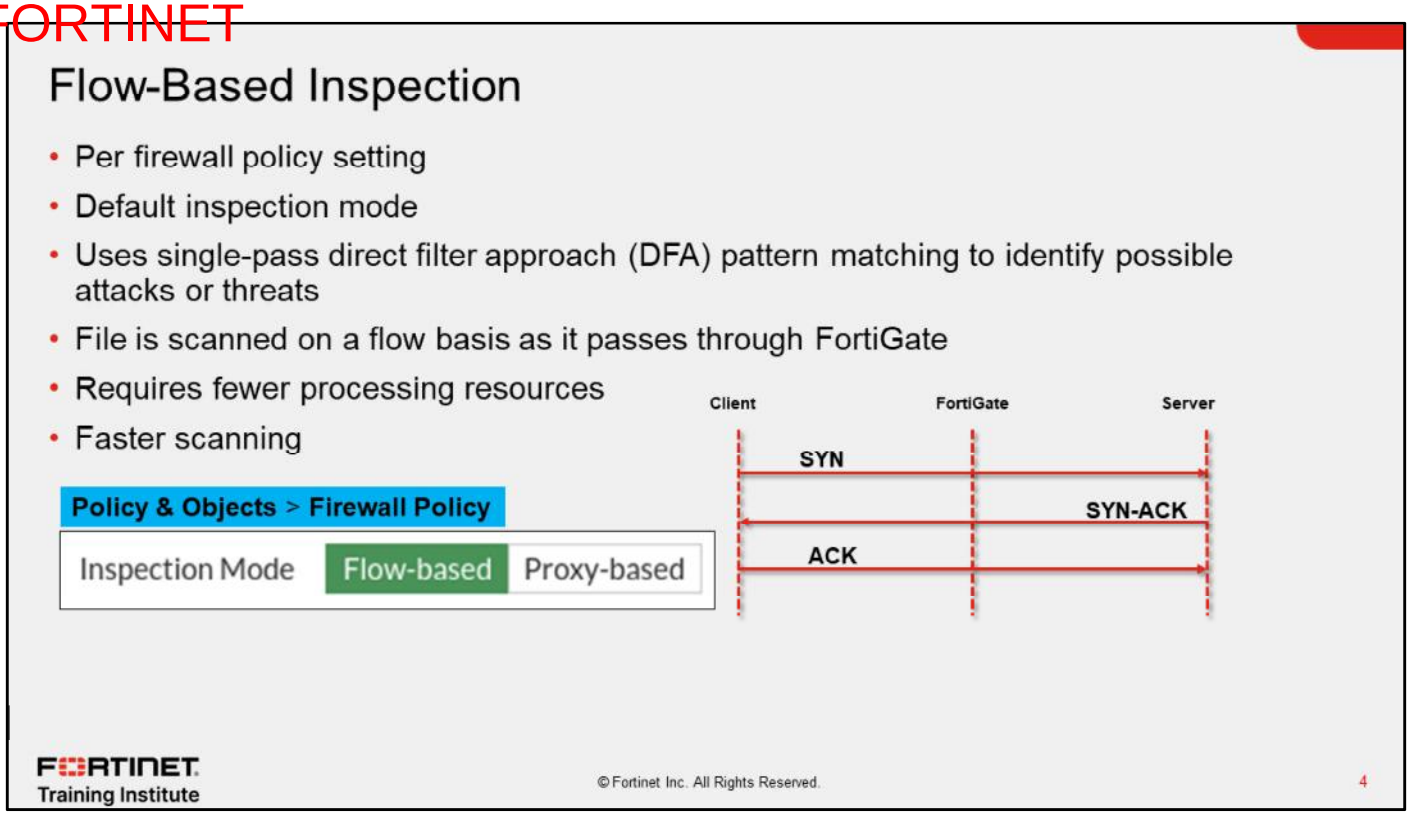

Flow-based inspection mode examines the file as it passes through FortiGate, without any buffering. As each packet arrives, it is processed and forwarded without waiting for the complete file or web page. If you are familiar with the TCP flow analysis of Wireshark, then that is essentially what the flow engine sees. Packets are analyzed and forwarded as they are received. Original traffic is not altered. Therefore, advanced features that modify content, such as safe search enforcement, are not supported.

The advantages of flow-based mode are:
- The user sees a faster response time for HTTP requests compared to proxy based
- There is less chance of a time-out error because of the server at the other end responding slowly

The disadvantages of flow-based mode are:
- A number of security features that are available in proxy-based mode are not available in flow-based mode
- Fewer actions are available based on the categorization of the website by FortiGuard services

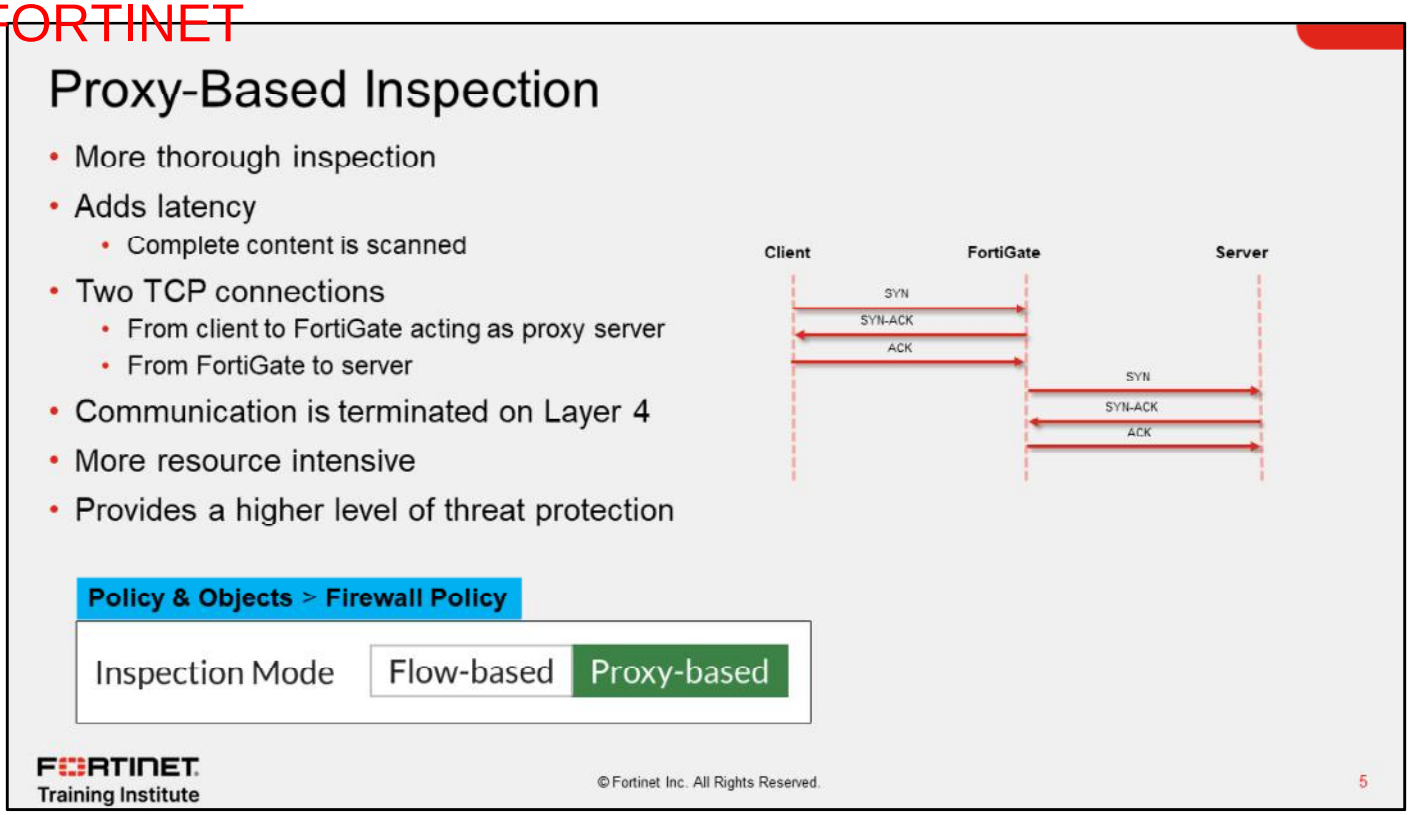

Proxy-based scanning refers to transparent proxy. It's called transparent because, at the IP layer, FortiGate is not the destination address, but FortiGate *does* intercept the traffic. When proxy-based inspection is enabled, FortiGate buffers traffic and examines it *as a whole*, before determining an action. Because FortiGate examines the data as a whole, it can examine more points of data than it does when using flow-based inspection.

In TCP connections, the FortiGate proxy generates the SYN-ACK to the client, and completes the three-way handshake with the client, before creating a second, new connection to the server. If the payload is less than the oversize limit, the proxy buffers transmitted files or emails for inspection, before continuing transmission. The proxy analyzes the headers and may change the headers, such as HTTP host and URL, for web filtering. If a security profile decides to block the connection, the proxy can send a replacement message to the client. This adds latency to the overall transmission speed.

Proxy-based inspection is more thorough than flow-based inspection, yielding fewer false positives and negative results.

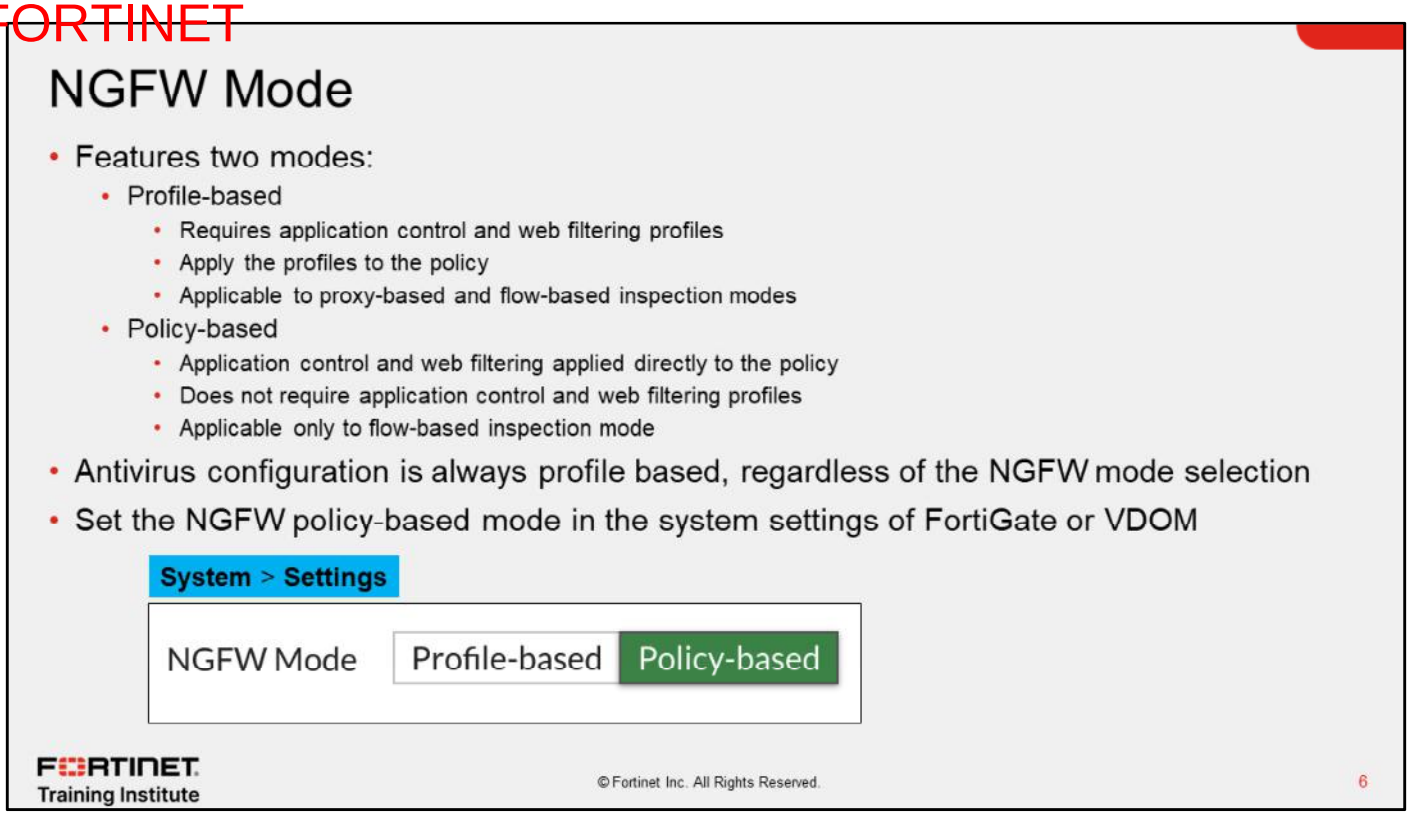

FortiGate, or the individual VDOM, has two next-generation firewall (NGFW) modes available:

1. Profile-based mode: Requires administrators to create and use application control and web filter profiles and apply them to a firewall policy. Profile-based mode is applicable to use flow-based or proxy-based inspection mode as per the policy.
2. Policy-based mode: Administrators can apply application control and web filter configuration directly to a security policy. Flow-based inspection mode is the only applicable process available in policy-based NGFW mode.

Antivirus scanning is available as a security profile that you can apply in a profile-based NGFW mode firewall policy or policy-based NGFW mode security policy.

You can change NGFW mode in the system settings of FortiGate or the individual VDOM. Note that the change will require you to remove all existing policies in either mode.

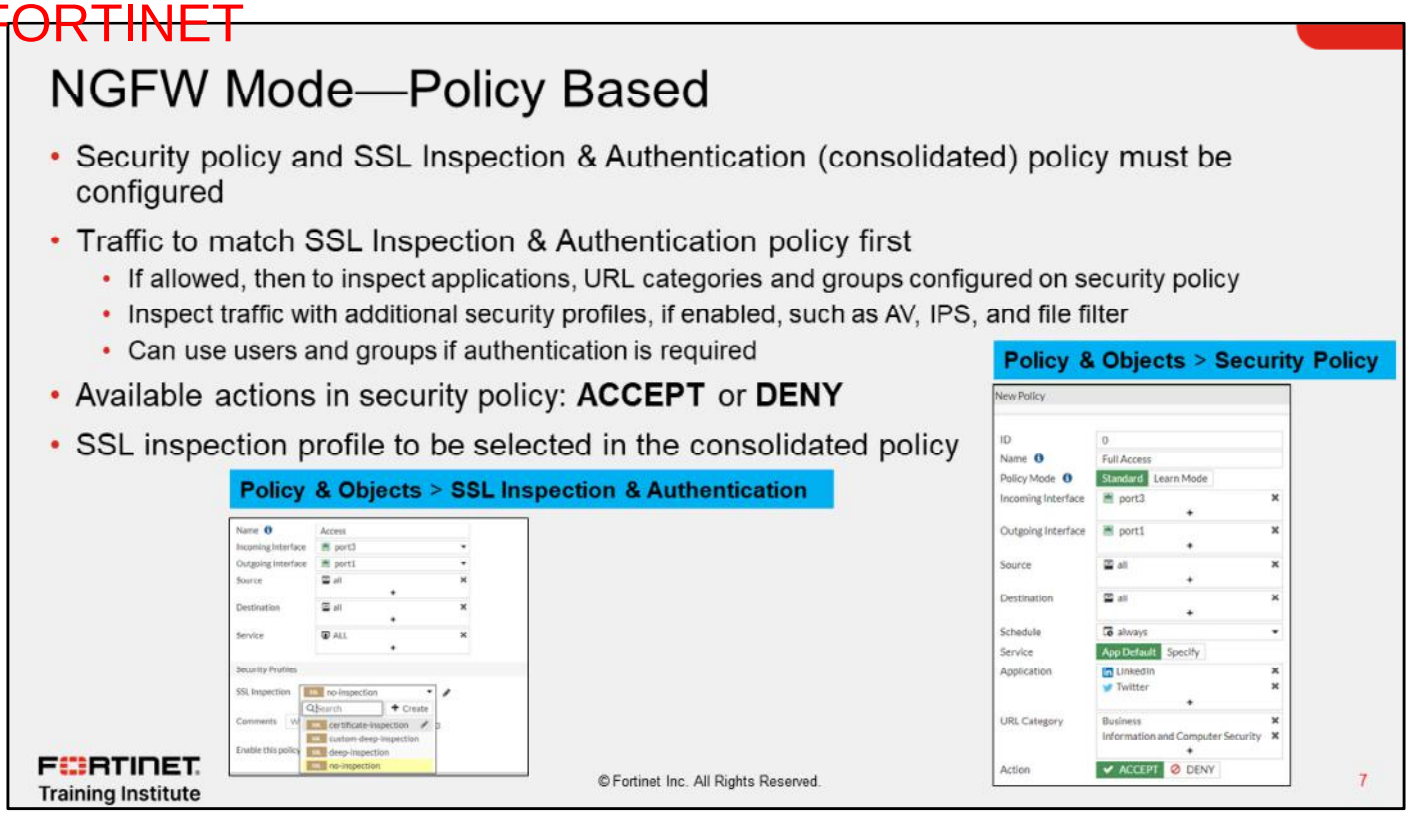

If you configured FortiGate to use NGFW policy-based mode or created a VDOM specifically to provide NGFW policy-based mode, you must configure a few policies to allow traffic.

SSL Inspection & Authentication (consolidated) policy: This policy allows traffic from a specific user or user group to match the criteria specified within the consolidated policy, and inspect SSL traffic using the SSL inspection profile selected. FortiGate can either accept or deny the traffic.

Security policy: If the traffic is allowed according to the consolidated policy, FortiGate then processes it based on the security policy to analyze additional criteria, such as URL categories, groups for web filtering, and application control. Also, if enabled, the security policy further inspects traffic using security profiles such as AV, IPS, and file filter.

Knowledge Check

1. How does NGFW policy-based mode differ from profile-based mode?
   - A. Policy-based flow inspection supports web profile overrides.
   - ✓ B. Policy-based flow inspection defines URL filters directly in the firewall policy.

2. Which statement about proxy-based web filtering is true?
   - A. It requires more resources than flow-based
   - ✓ B. It transparently analyzes the TCP flow of the traffic

FORTINET Training Institute

8

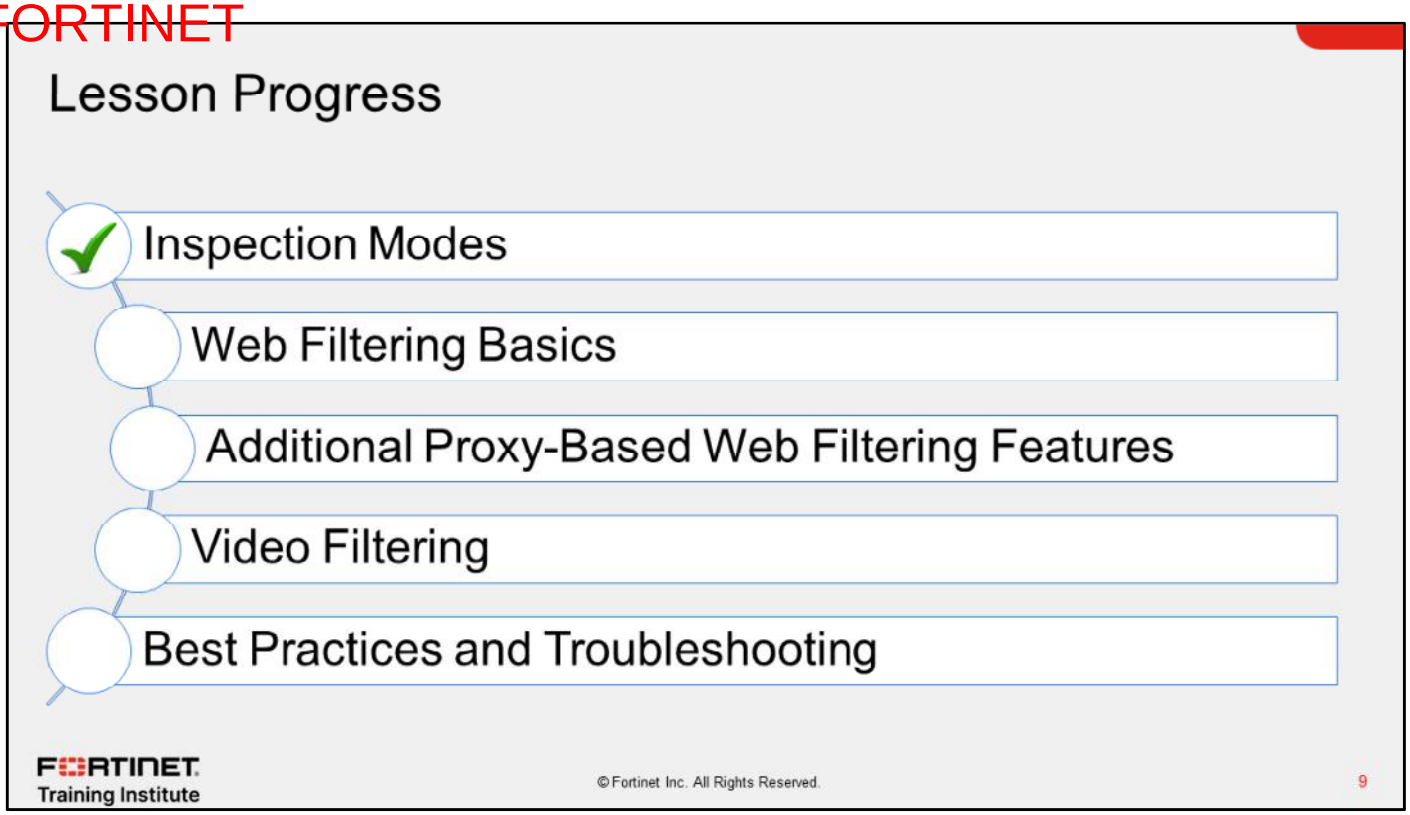

Good job! You now understand inspection modes.

Now, you will learn about web filtering basics.

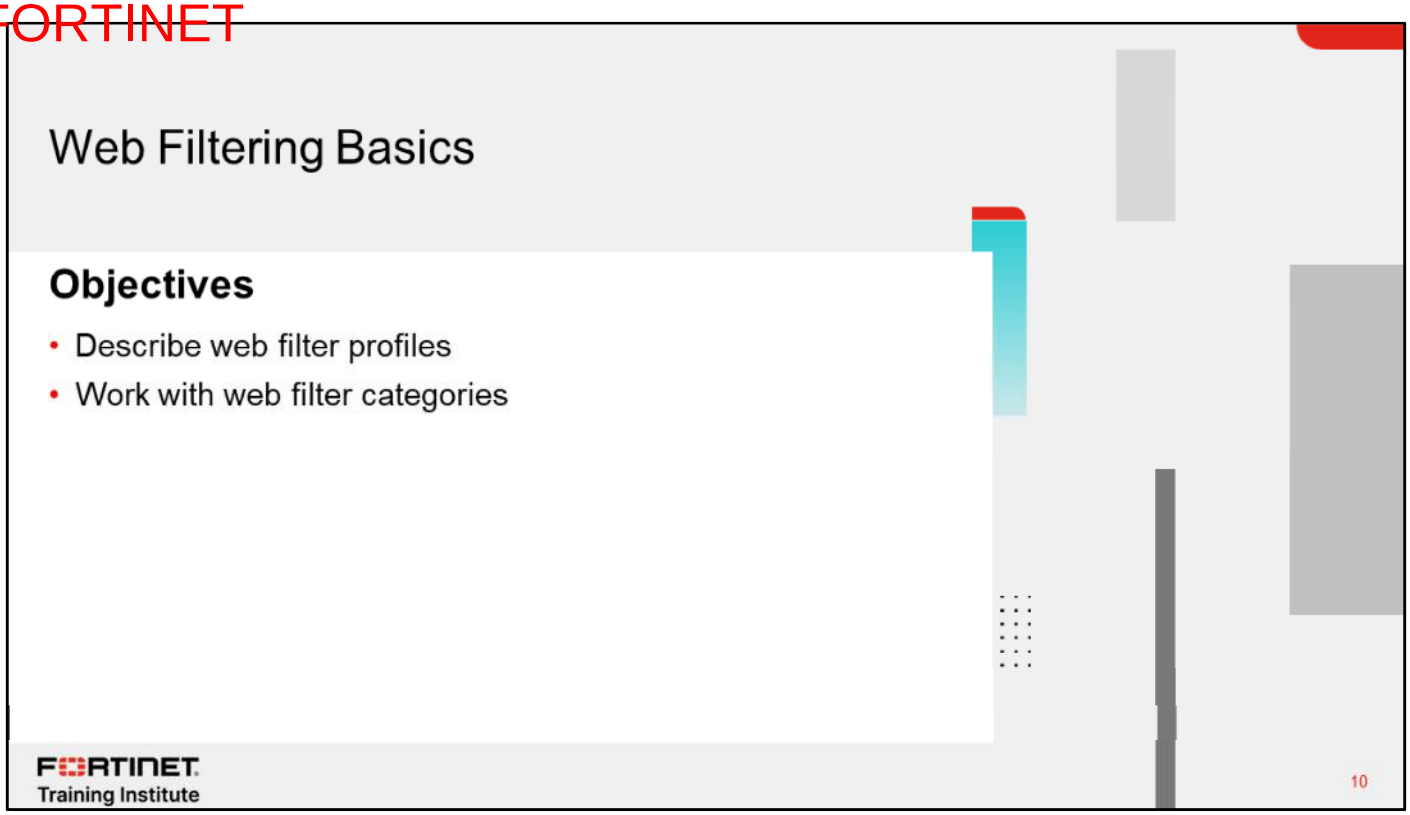

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in web filtering basics, you will be able to describe web filter profiles and use FortiGuard web filter profiles.

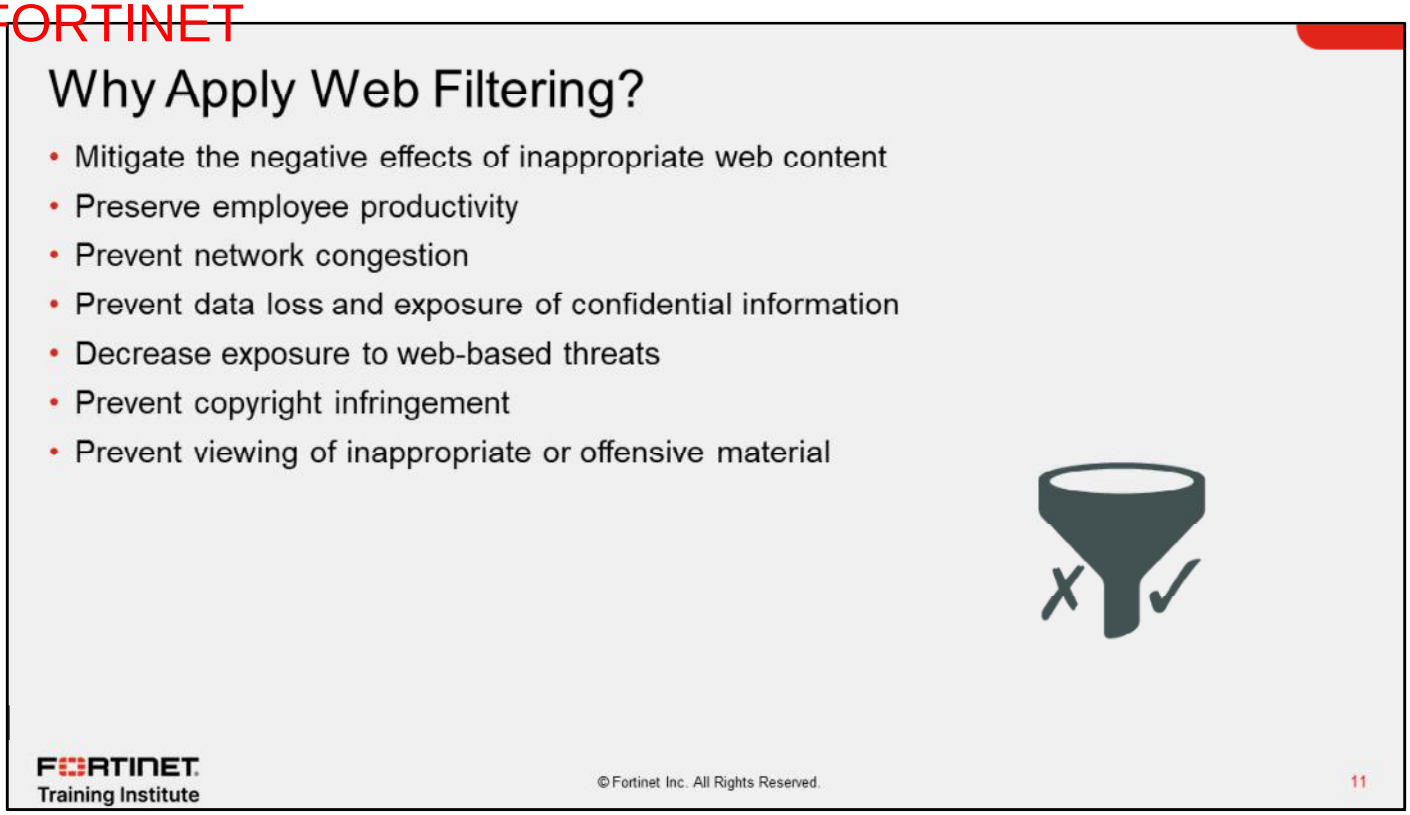

Web filtering helps to control, or track, the websites that people visit. There are many reasons why network administrators apply web filtering, including to:

- Preserve employee productivity
- Prevent network congestion, where valuable bandwidth is used for non-business purposes
- Prevent loss or exposure of confidential information
- Decrease exposure to web-based threats
- Limit legal liability when employees access or download inappropriate or offensive material
- Prevent copyright infringement caused by employees downloading or distributing copyrighted materials
- Prevent children from viewing inappropriate material

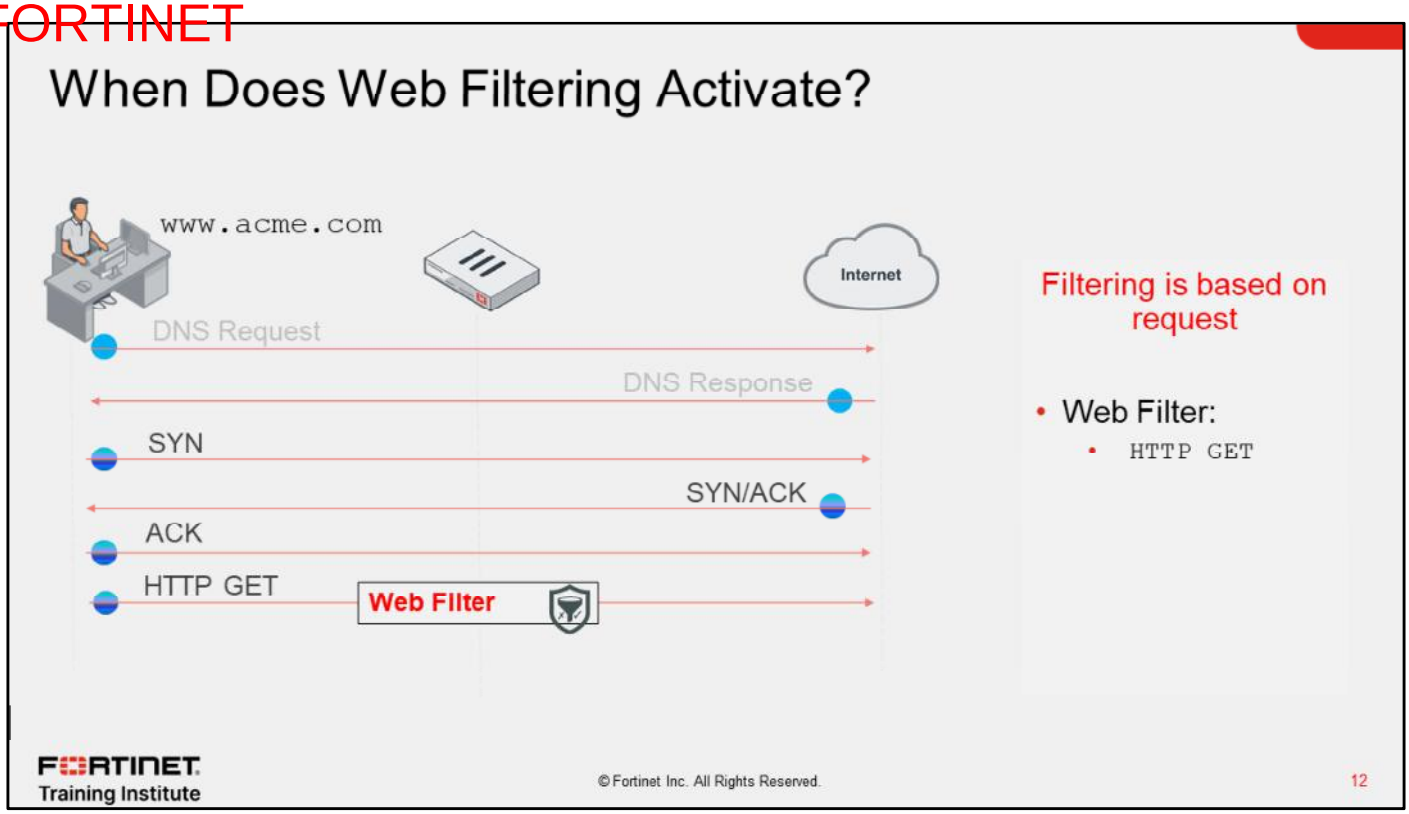

(slide contains animation)

The example on this slide shows the flow of an HTTP filter process.

FortiGate looks for the `HTTP GET` request to collect URL information and perform web filtering.

So, as shown, in HTTP the domain name and URL are separate pieces. The domain name might look like the following in the header: `Host: www.acme.com`, and the URL might look like the following in the header: `/index.php?login=true`.

If you filter by domain, sometimes it blocks too much. For example, the blogs on `tumblr.com` are considered different content, because of all the different authors. In that case, you can be more specific, and block by the URL part, `tumblr.com/hacking`, for example.

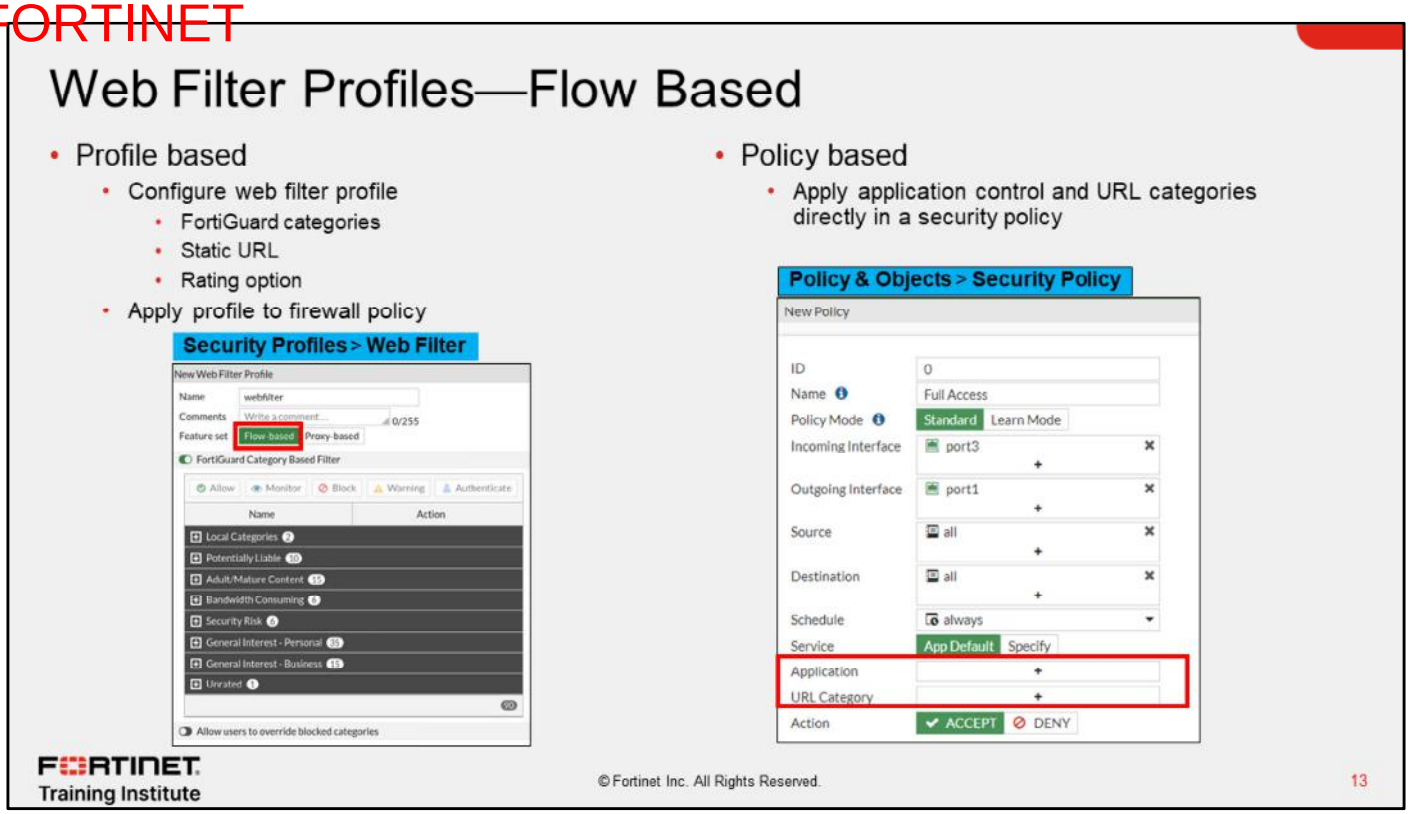

Now, you will look at the web filter profile.

You can configure this security profile to use a feature set for proxy-based or flow-based inspection modes. However, depending on the mode you select, the available settings are different. Flow-based inspection has fewer available options.

In the examples shown on this slide, the web filter profile has a FortiGuard category-based filter that categorizes the websites based on categories and subcategories by FortiGuard. FortiGate offers two NGFW options:

- **Profile-Based** (default)
  - Web filters are defined as security profiles and applied to the firewall policy
- **Policy-Based**
  - URL categories are defined directly under the firewall policy

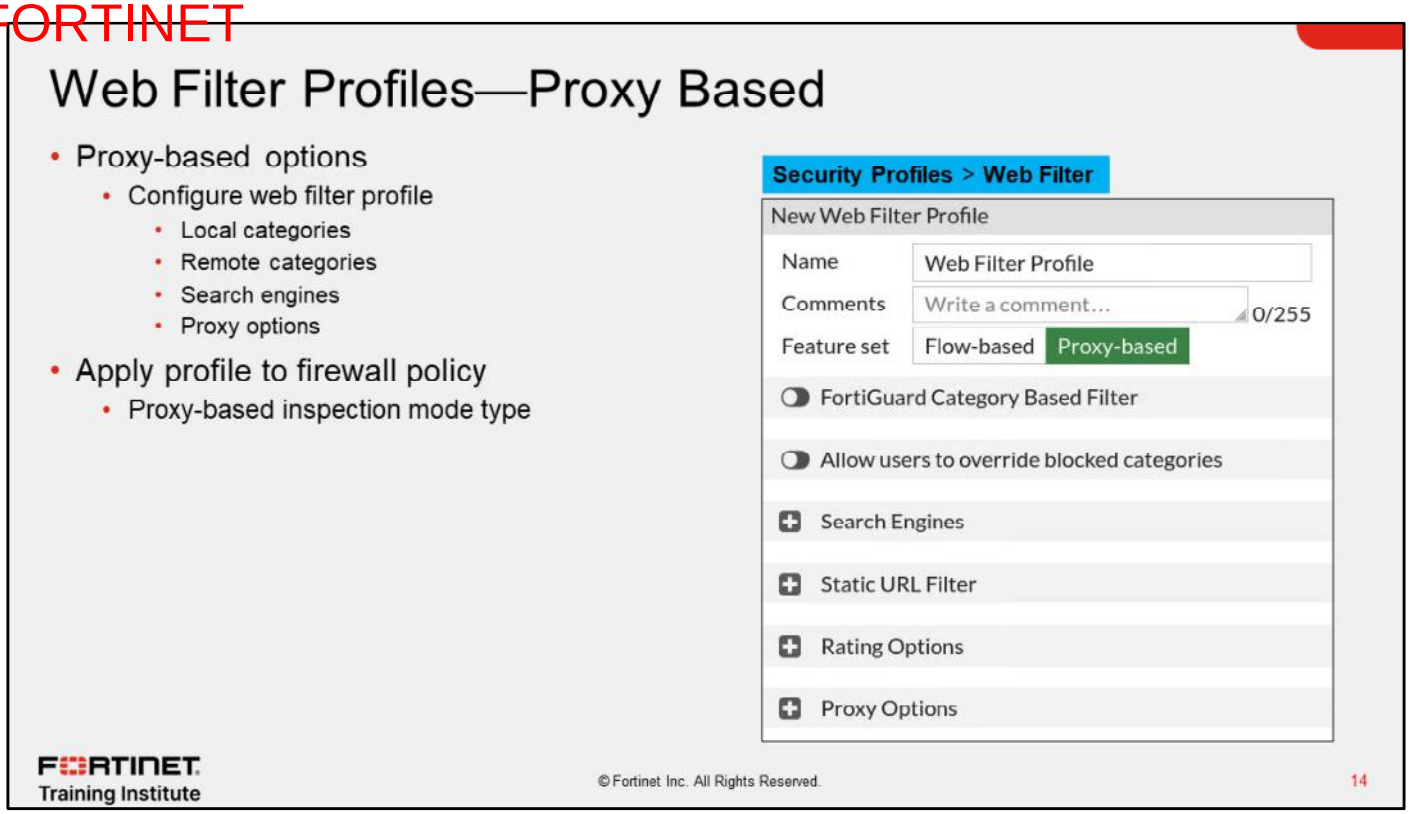

In the example shown on this slide, the security profile is configured to use a proxy-based feature set. The profile is available to a firewall policy configured to use proxy-based inspection mode. Other local options include:

- **Search Engines**
- **Static URL Filter**
- **Rating Options**
- **Proxy Options**

After you configure your web filter profile, apply this profile to your firewall policy so the filtering is applied to your web traffic.

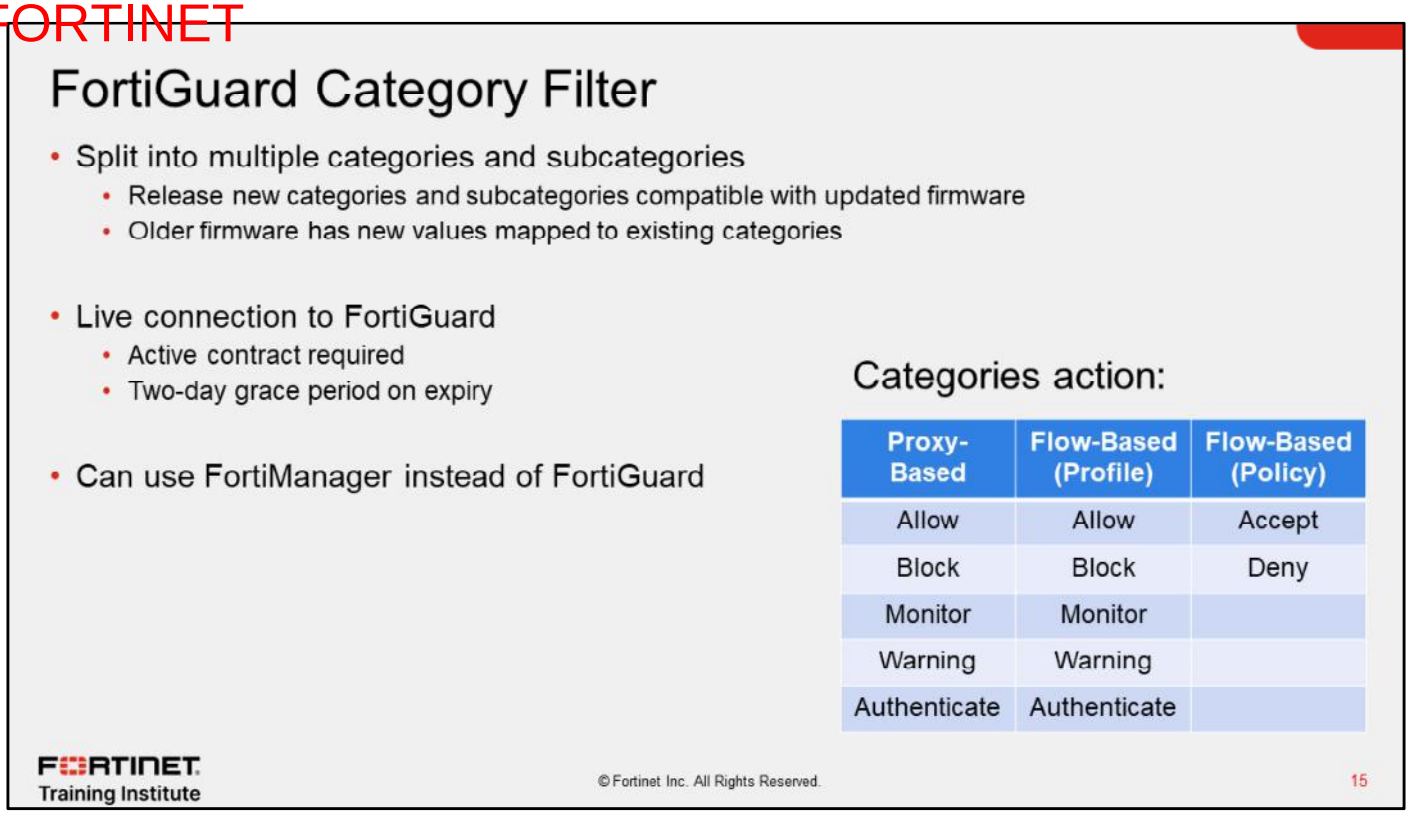

Rather than block or allow websites individually, FortiGuard category filtering looks at the category that a website has been rated with. Then, FortiGate takes action based on that category, not based on the URL.

FortiGuard category filtering is a live service that requires an active contract. The contract validates connections to the FortiGuard network. If the contract expires, there is a two-day grace period during which you can renew the contract before the service cuts off. If you do not renew, after the two-day grace period, FortiGate reports a rating error for every rating request made. In addition, by default, FortiGate blocks web pages that return a rating error. You can change this behavior by enabling the **Allow websites when a rating error occurs** setting. You will learn more about this setting in this lesson.

You can configure FortiManager to act as a local FortiGuard server. To do this, you must download the databases to FortiManager, and configure FortiGate to validate the categories against FortiManager, instead of FortiGuard.

You can enable the FortiGuard category filtering on the web filter. Categories and subcategories are listed, and you can customize the actions to perform individually.

The actions available depend on the mode of inspection:
- Proxy: Allow, Block, Monitor, Warning, and Authenticate
- Flow-based, profile-based: Allow, Block, Monitor, Warning, and Authenticate
- Flow-based, policy-based: Action defined in a security policy (accept or deny)

To review the complete list of categories and subcategories, visit `www.fortiguard.com/webfilter/categories`.

Web Filter FortiGuard Category Action—Warning

- Category Action =

- Exclusive for web filtering
  - Proxy mode
  - Flow mode (profile-based only)
  - Not available in:
    - Static URL filtering feature
- FortiGuard warning page
  - Customizable warning interval

FORTINET Training Institute

16

The warning action informs users that the requested website is not allowed by the internet policies. However, the action gives the user the option to proceed to the requested website, or return to the previous website.

You can customize the warning interval, so you can present this warning page at specific times, according to the configured period.

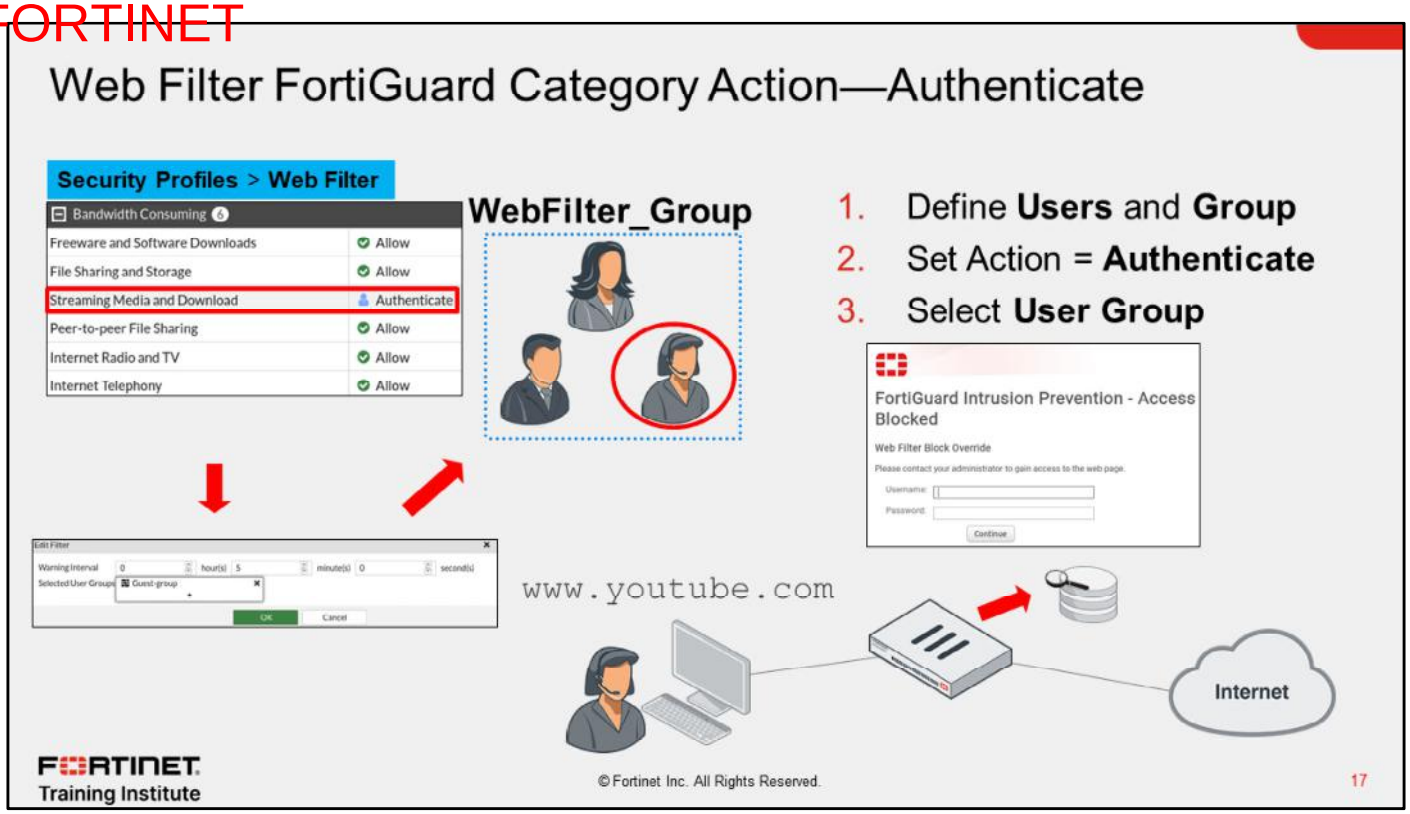

The authenticate action blocks the requested websites, unless the user enters a successful username and password. Local authentication and remote authentication using LDAP, radius and so on are supported for web filtering authentication.

You can customize the interval of time to allow access. Users are not prompted to authenticate again if they access other websites in the same category until the timer expires.

Choosing this action prompts you to define user groups that are allowed to override the block.

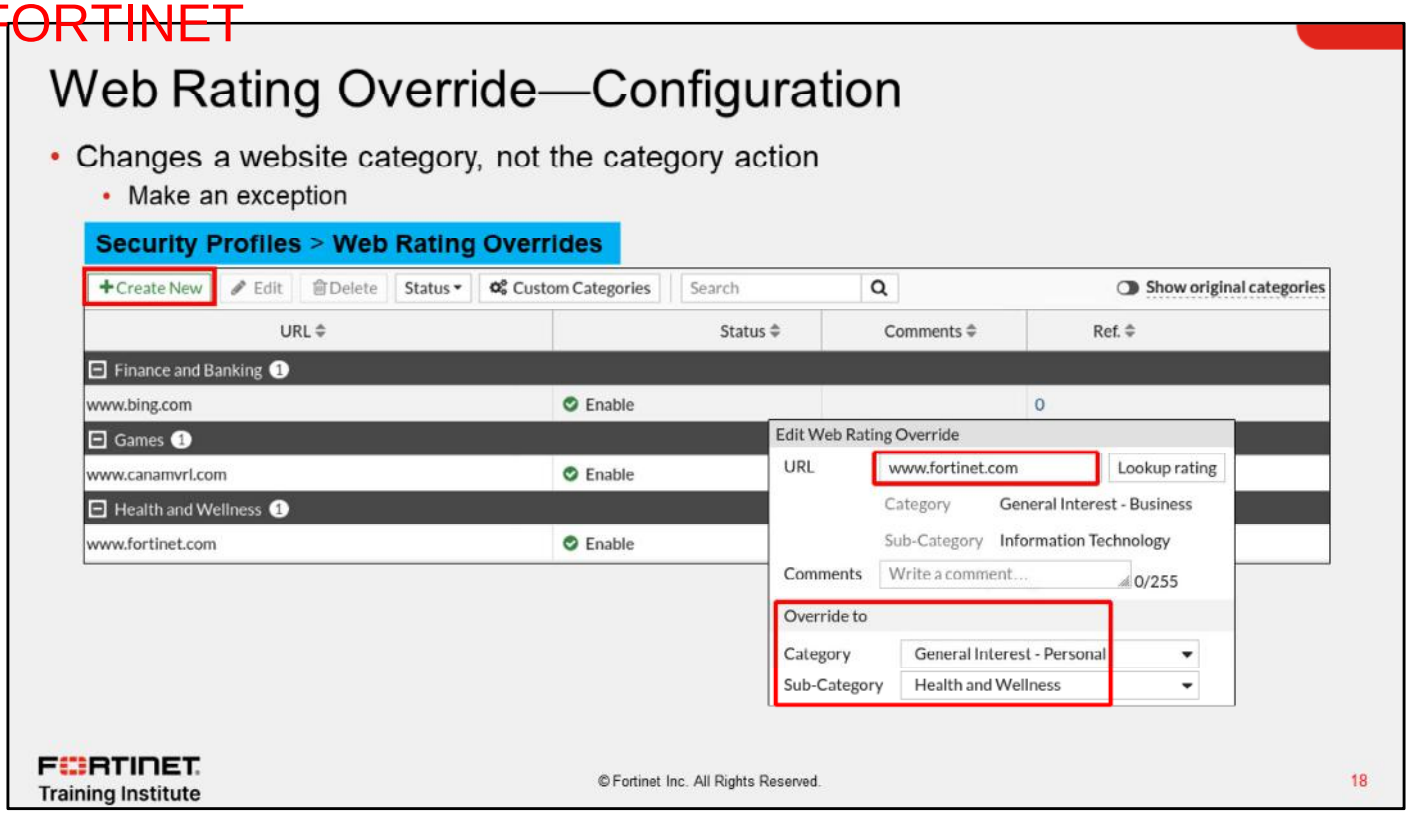

If you want to make an exception, for example, rather than unblock access to a potentially unwanted category, change the website to an allowed category. You can also do the reverse. You can block a website that belongs to an allowed category.

Remember that changing categories does not automatically result in a different action for the website. This depends on the settings within the web filter profile.

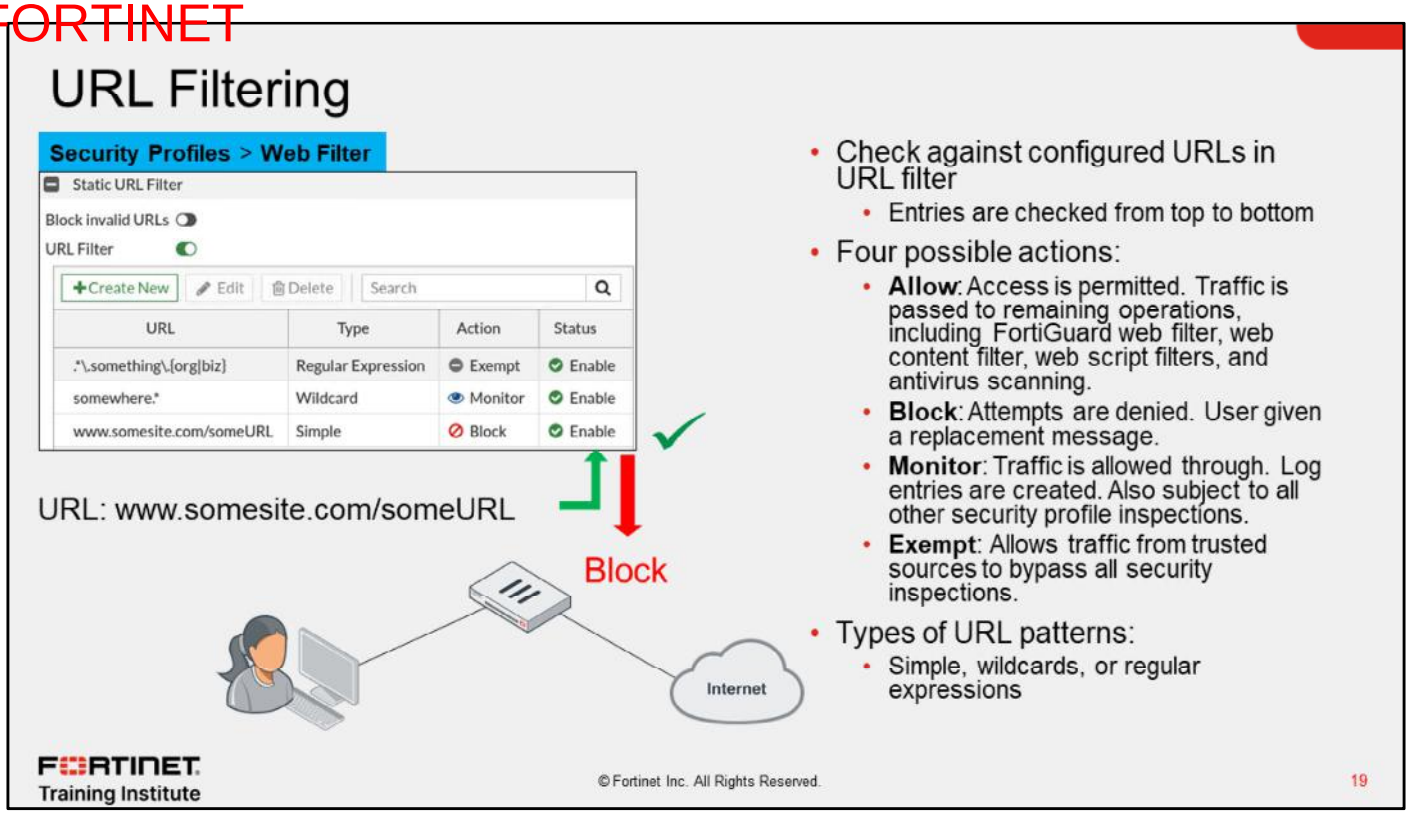

Static URL filtering is another web filter feature. Configured URLs in the URL filter are checked against the visited websites. If a match is found, the configured action is taken. URL filtering has the same patterns as static domain filtering: simple, regular expressions, and wildcard.

Take a look at how it works.

When a user visits a website, FortiGate looks at the URL list for a matching entry. In the example shown on this slide, the website matches the third entry in the table, which is set as type **Simple**. This type means that the match must be exact—there is no option for a partial match with this pattern. Also, the action is set to **Block**, so FortiGate displays a block page message.

Knowledge Check

1. Which of the following can act as a local FortiGuard server?
   - ✓ A. FortiManager
   - B. FortiAnalyzer

2. Which action in URL filtering will bypass all security profiles?
   - ✓ A. Exempt
   - B. Allow

FORTINET Training Institute

20

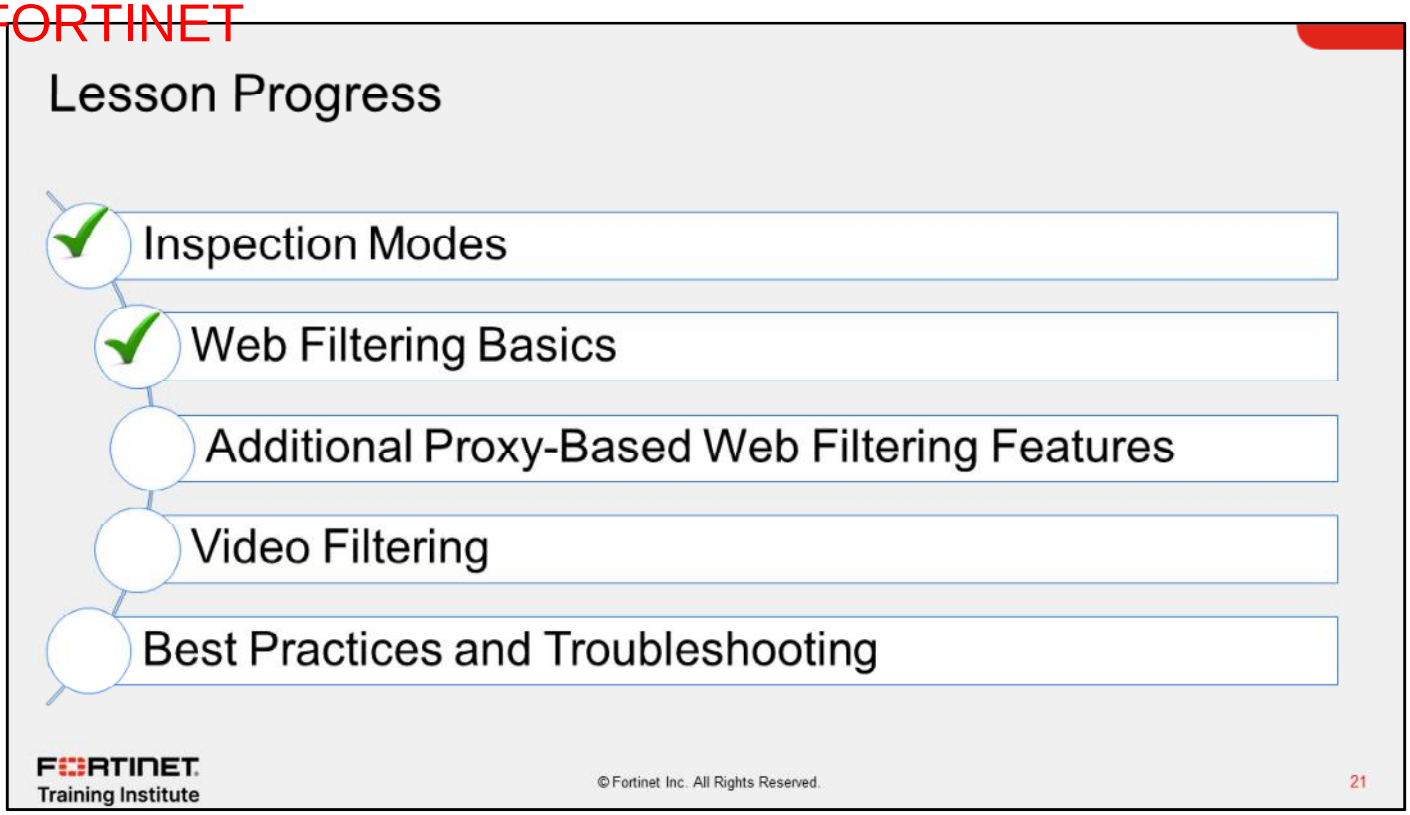

Good job! You now understand the basics of web filtering.

Now, you will learn about additional proxy-based web filtering features.

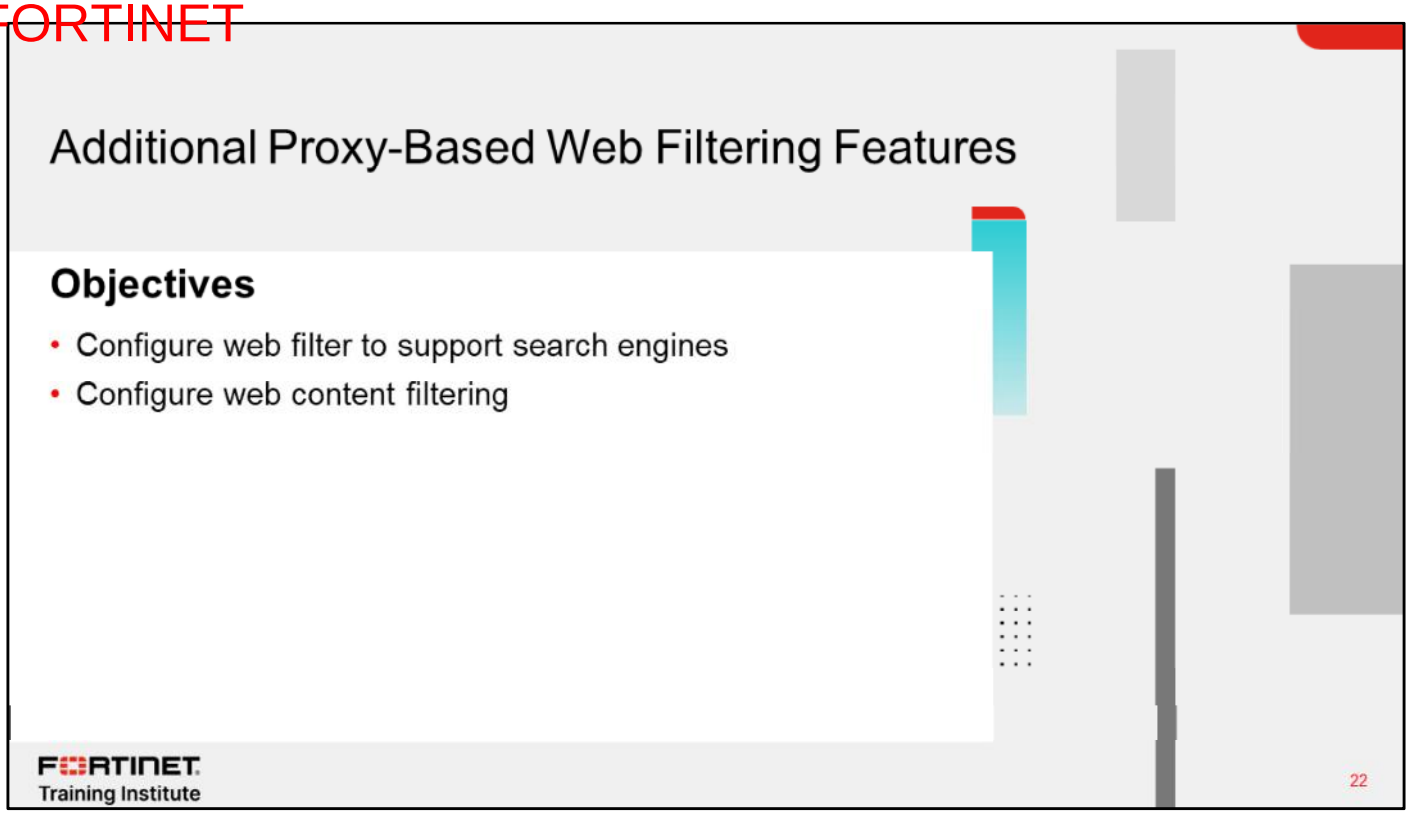

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in additional proxy-based web filtering features, you will be able to search engine filters, and web content filtering.

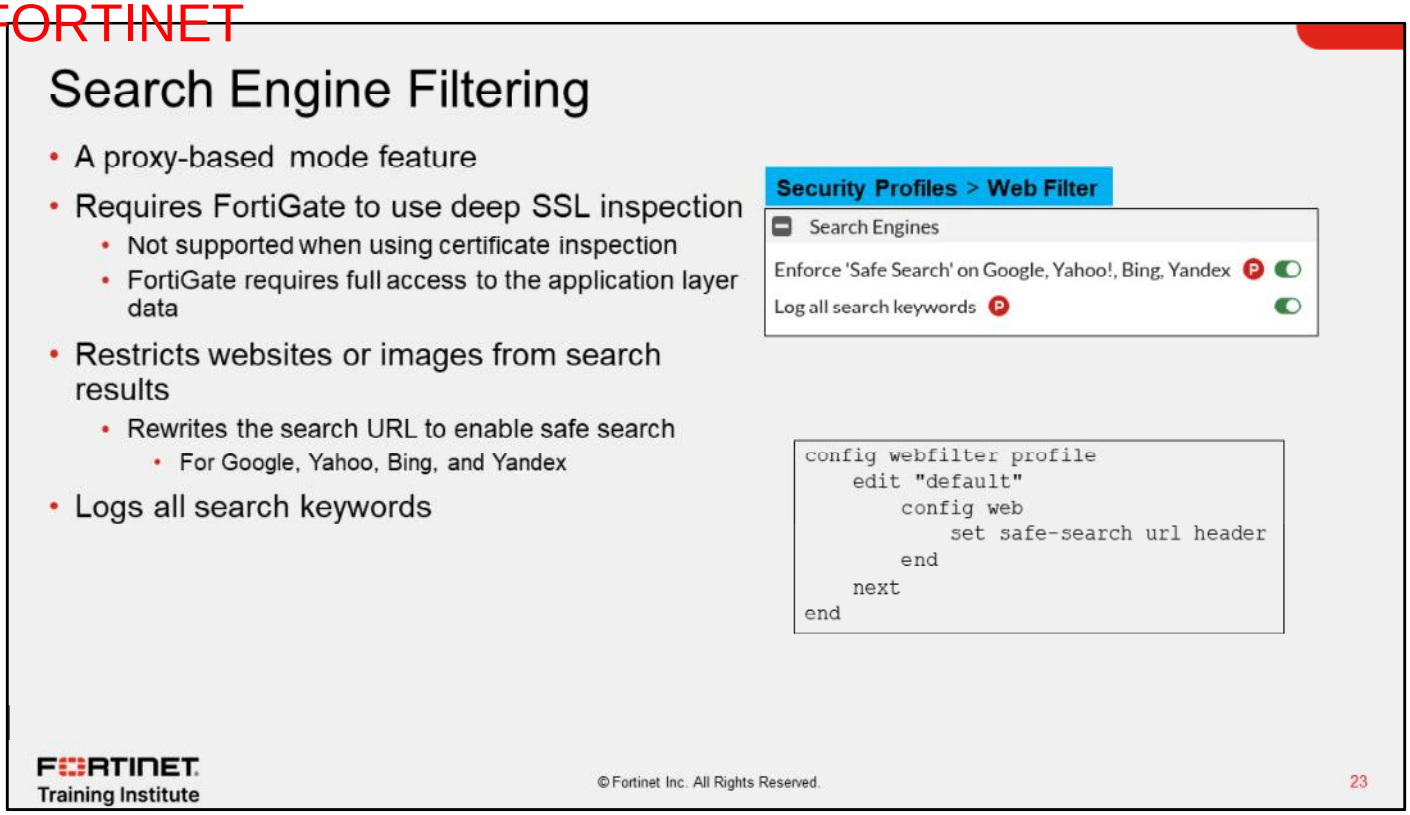

Search engine filtering is available when you configure a web filter profile while setting the feature set to proxy-based.

Safe search is an option that some browsers support. It applies internal filters to the search results. When you enable safe search for the supported search sites, FortiGate appends code to the URL to enforce the use of safe search. For example, on a Google search, FortiGate adds the string `&safe=active` to the URL in the search. So, even if it is not locally enabled in the browser, FortiGate applies safe search to the requests when they pass through. Safe search is supported for Google, Yahoo, Bing, and Yandex.

As a proxy-based web filter feature, search engine filtering is supported only when using full SSL inspection because FortiGate requires access to the full header.

Web Content Filtering

- Requires FortiGate to use SSL deep inspection
- Controls access to web pages containing specific patterns
- Scans the content of every website accepted by security policies
- Matches content from wildcards or Perl regular expressions
- The maximum number of web content patterns in a list is 5000
- Actions:
  - Exempt
  - Block

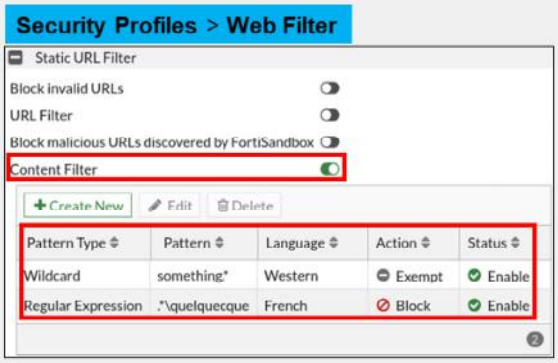

FORTINET Training Institute

24

You can also control web content in the web filter profile by blocking access to websites containing specific words or patterns. This helps to prevent access to sites with questionable material.

You can add words, phrases, patterns, wildcards, and Perl regular expressions to match content on websites. You configure this feature on a per-web-filter-profile basis, not at the global level. So, it is possible to add multiple web content filter lists and then select the best list for each web filter profile.

The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the scores of banned words and phrases on the page. If the sum is higher than the threshold set in the web filter profile, FortiGate blocks the site.

The maximum number of web content patterns in a list is 5000.

Like search engine filtering, web content filtering requires that FortiGate uses deep SSL inspection because FortiGate requires full access to the packet headers.

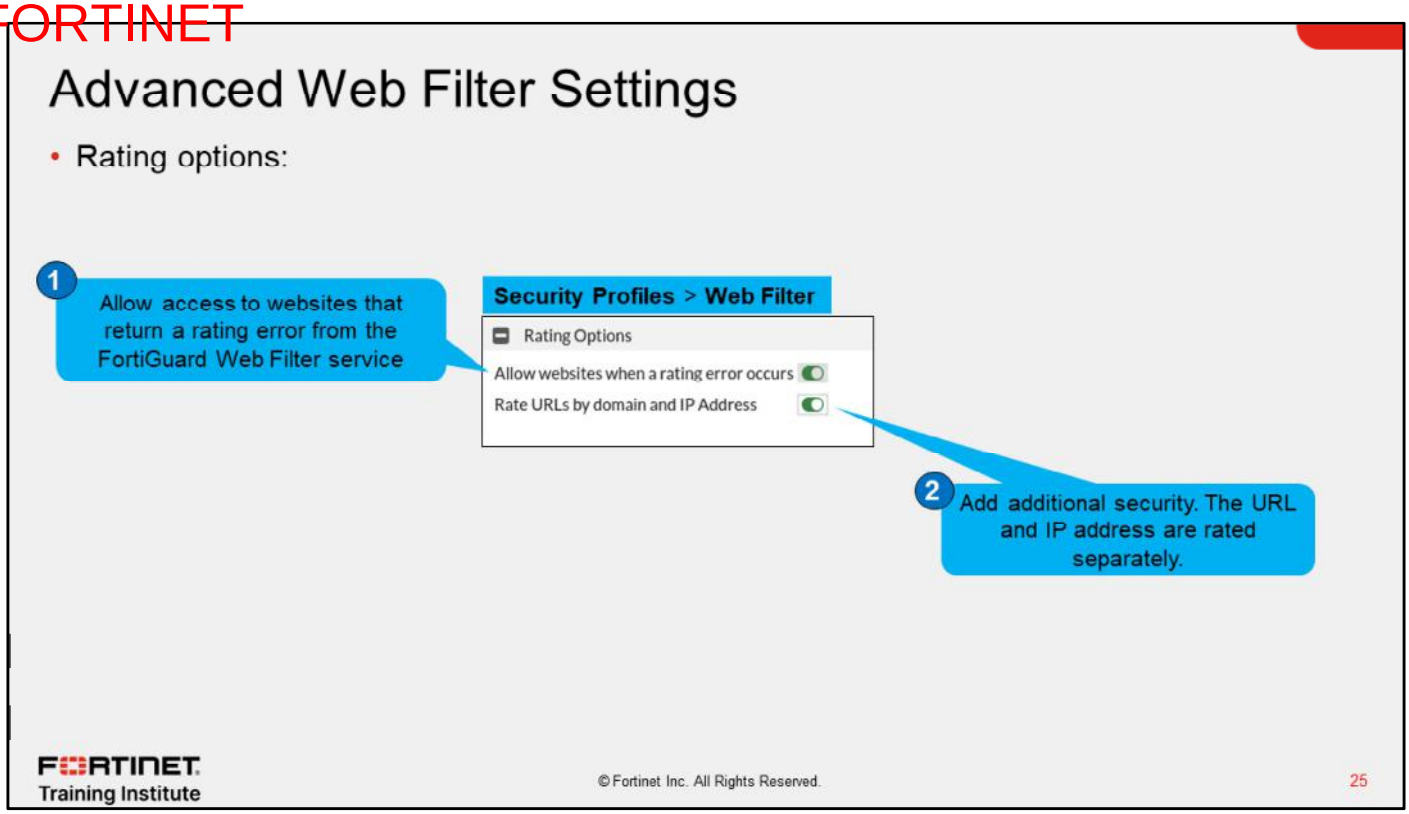

You can use advanced web filtering settings to improve the web filter.

The rating options are as follows:

1. **Allow websites when a rating error occurs**. If a rating error occurs from the FortiGuard web filter service, users have full unfiltered access to all websites.
2. **Rate URLs by domain and IP Address**. This option sends both the URL and the IP address of the requested site for checking, providing additional security against attempts to bypass the FortiGuard system.

Knowledge Check

1. Which of the following is used for matching content in Web Content Filtering ?
   - ✓ A. Perl regular expressions
   - B. Boolean operators

2. Which feature can be used for restricting websites or images from search results ?
   - A. Web Content Filtering
   - ✓ B. Search Engine Filtering

FORTINET Training Institute

26

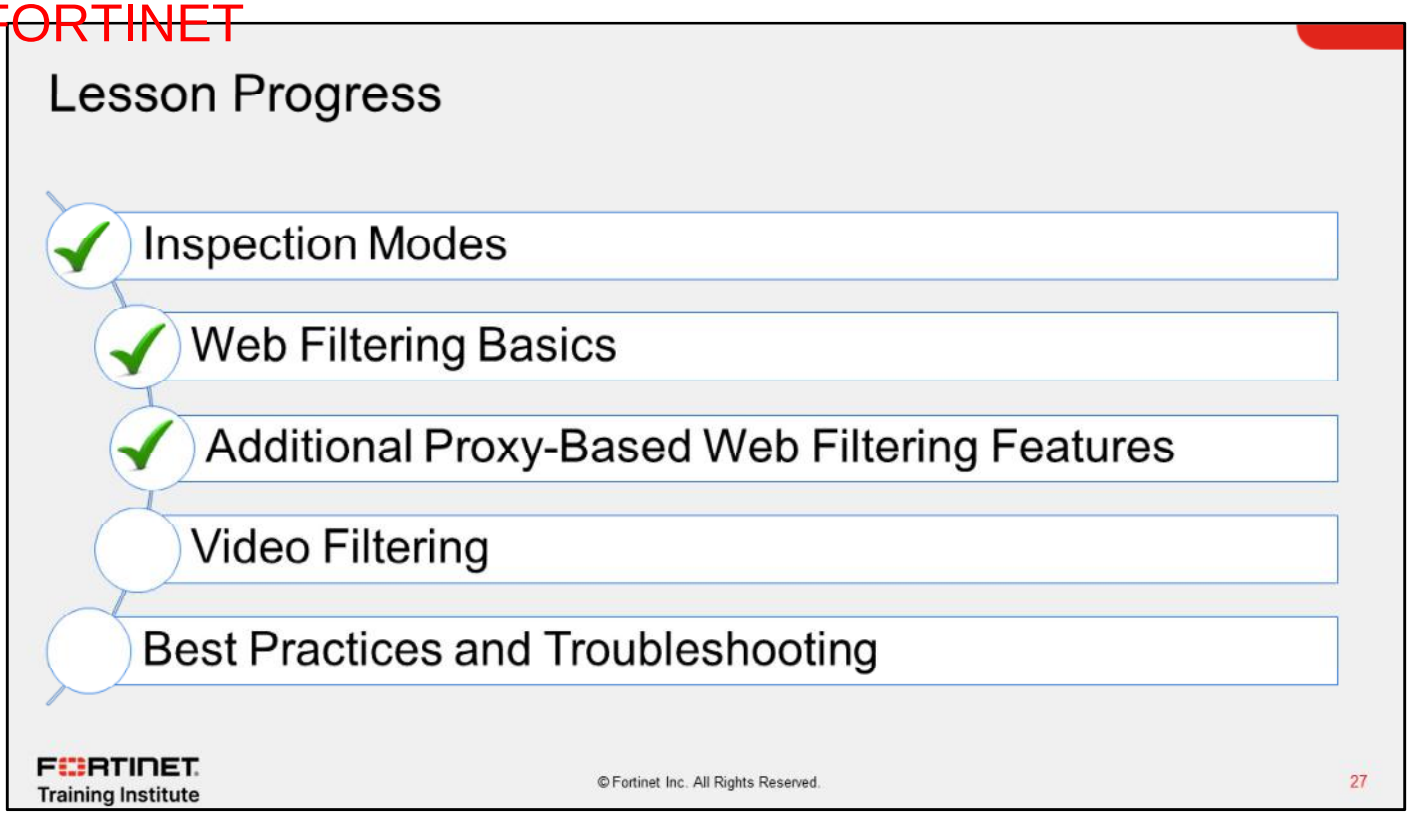

Good job! You now understand additional proxy-based web filtering features.

Now, you will learn about video filtering.

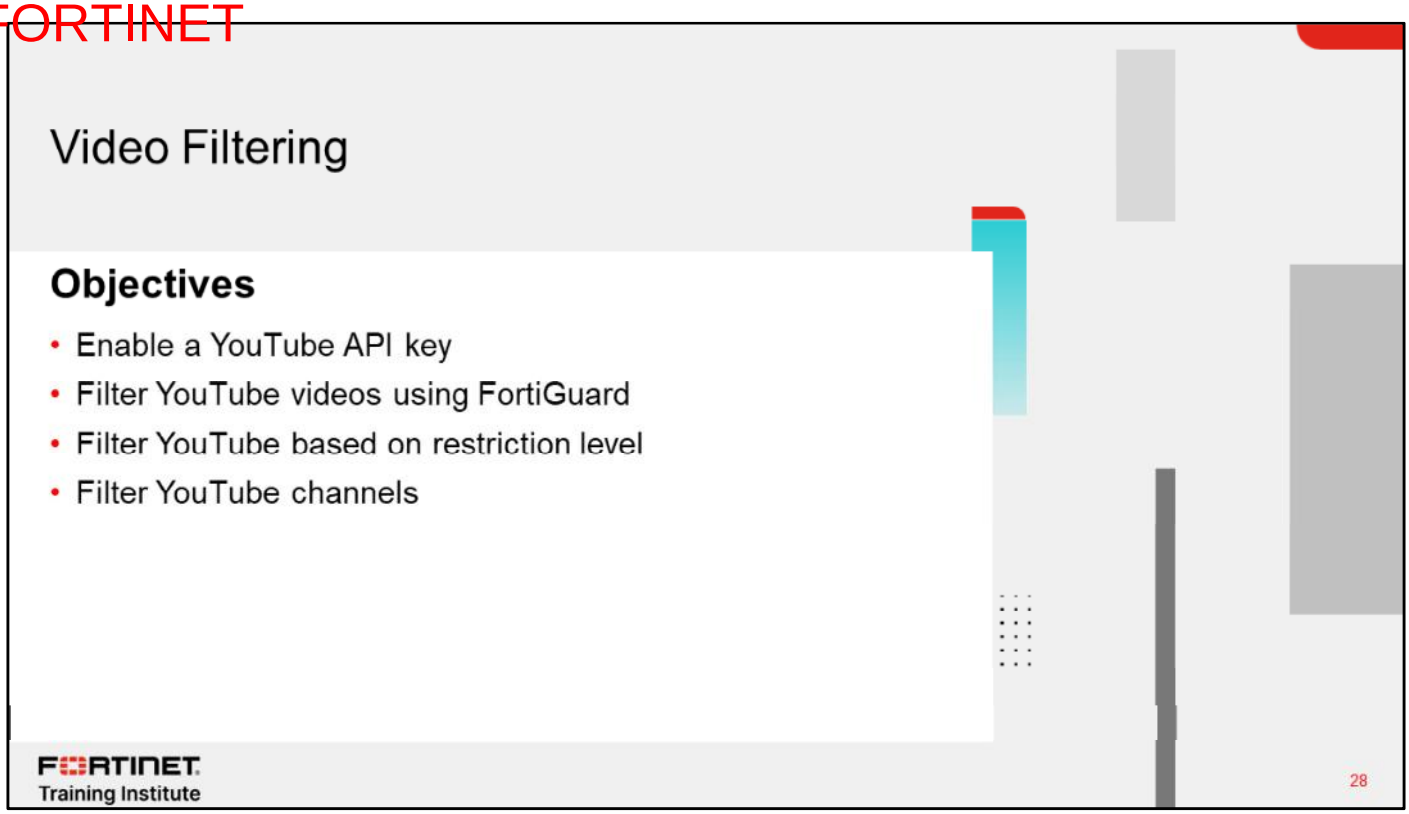

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in video filtering, you will be able to control access to YouTube using FortiGuard categories and YouTube static IDs.

Video Filter Profile

- Controls YouTube content access:
  - To allow, monitor, or block based on category
  - To allow, monitor, or block access to channels
  - To set restriction levels
- Separate FortiGuard license for video filtering
- Supported only on proxy-based firewall policy
- Requires full SSL inspection
- Requires YouTube API key
- Enables YouTube key on CLI
  - You can add multiple YouTube API keys
- Filters videos using two methods:
  - FortiGuard categories
  - Channel IDs

```
config videofilter youtube-key
    edit 1
        set key "youtube_api_key"
    next
end
```

FORTINET Training Institute

29

Video filtering allows you to control access to YouTube content using parameters that are associated with the video channel, video categories, or the video itself. It is part of the FortiGuard service, which requires a separate license bundled with the other FortiGuard security services.

To apply the video filter profile, proxy-based firewall polices currently allow you to enable the video filter profile. You must enable full SSL inspection on the firewall policy.

You must obtain a YouTube API key to use the video filter feature. The API key allows FortiGate to match parameters identified when users access YouTube content, and match the parameters with the local categories defined on the video filter.

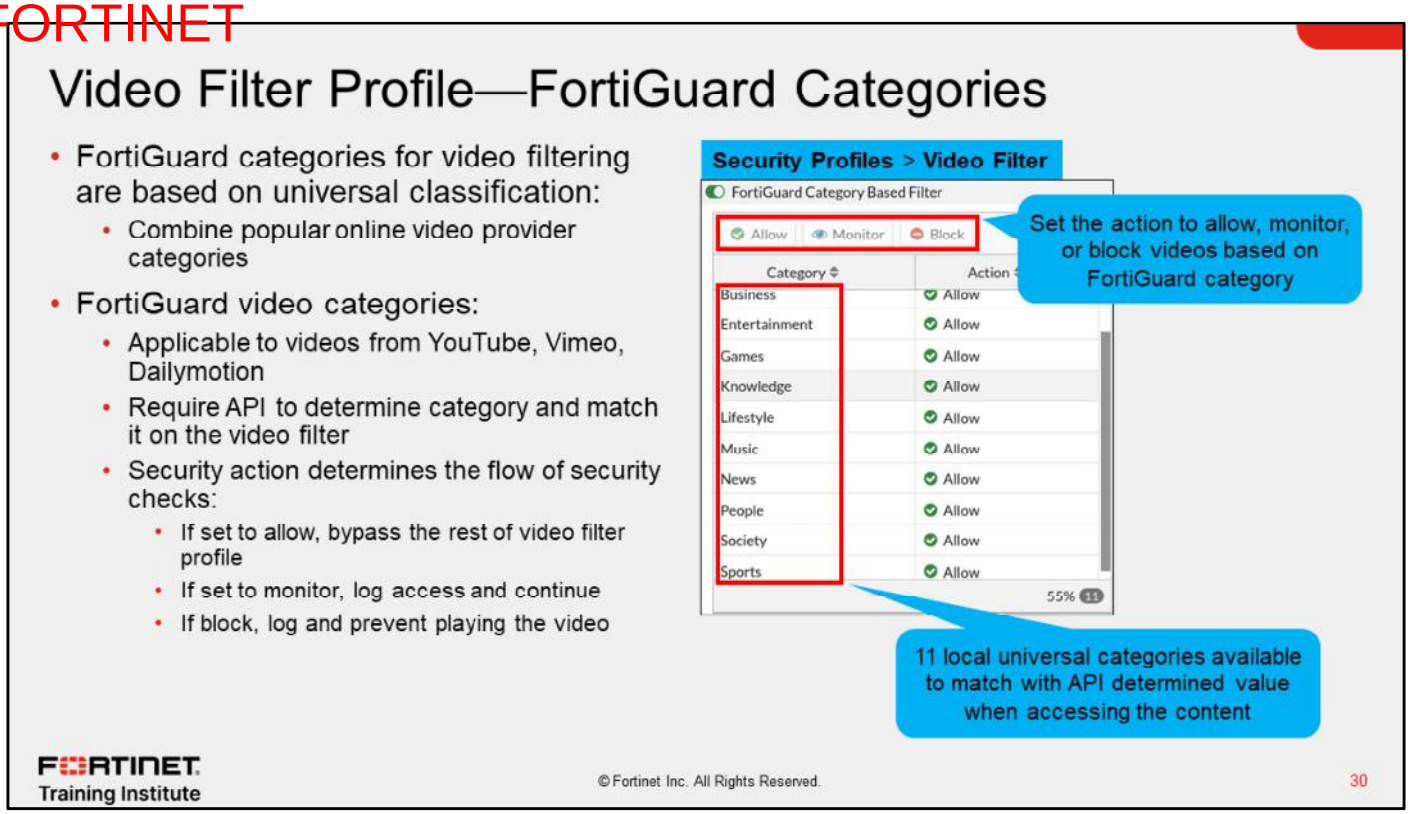

The video filter can identify videos using universal categories used by major online video content providers, such as YouTube. The generic classification combines multiple categories by these providers into one category. For example, the FortiGuard video category **Entertainment** includes YouTube categories, such as entertainment, comedy, movies, shows, and trailers.

The FortiGuard video categories are universal, to cover the common classifications used in the categories of online video content providers. Currently, it is applicable to content hosted by YouTube, Vimeo, and Dailymotion. Some of these providers offer API queries that enable FortiGate to identify the content and match it to local FortiGuard video categories.

In a video filter profile, if a FortiGuard category is allowed, the video content bypasses the rest of the security checks configured on the video profile, such as channel override and YouTube restriction level. If the action is set to monitor or block, then the video content undergoes further security checks configured on the video filter profile.

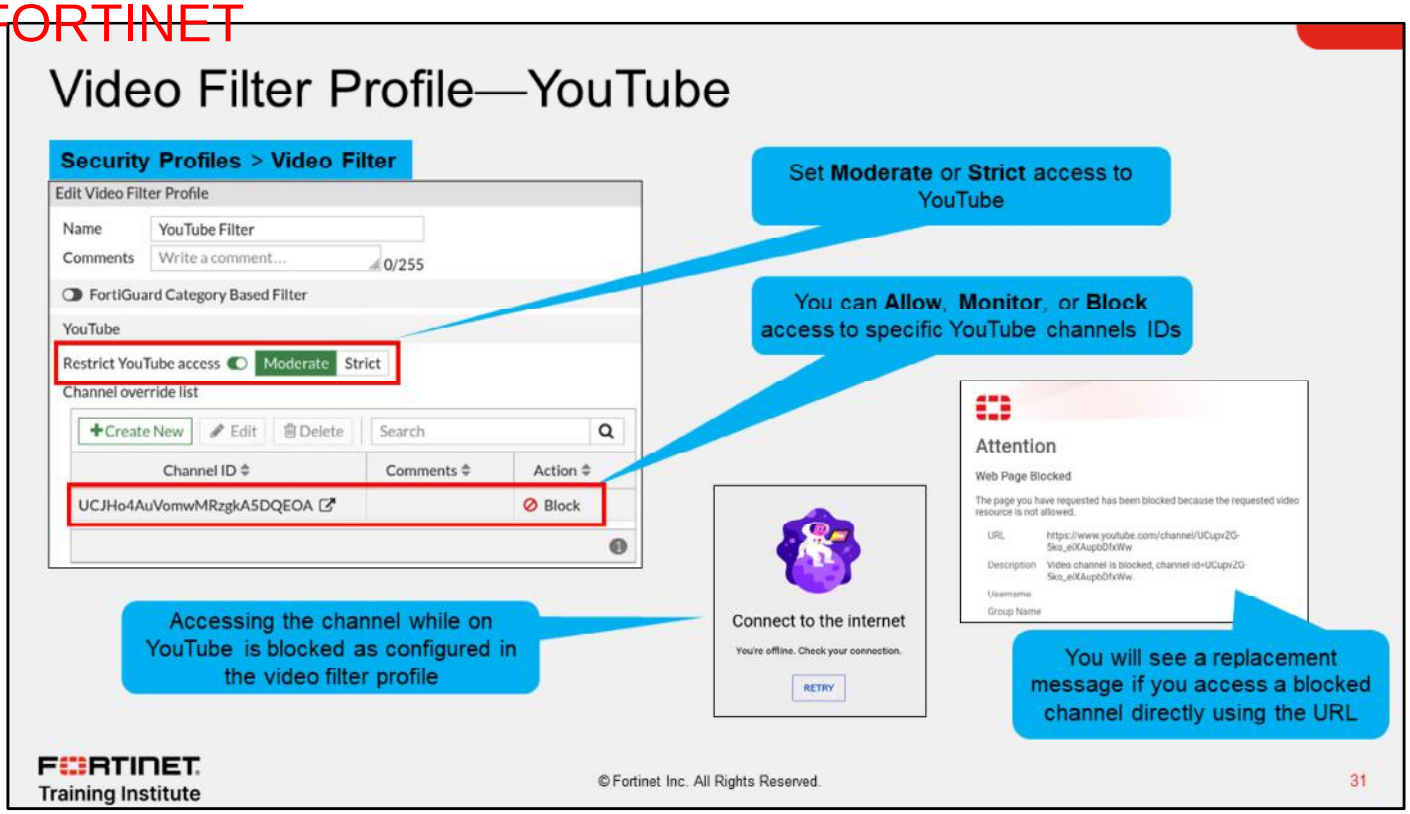

You can restrict YouTube access on a video filter by setting the restriction level to **Moderate** or **Strict**. When users access YouTube content using the firewall policy with the video filter profile applied, the users are given only content that is screened according to a filter applied by Google. Moderate restricted access is similar to strict but makes more videos available.

The YouTube channel ID is used to identify YouTube channels. It allows FortiGate to apply actions to access related content on the channel. These actions can allow, monitor, or block access to the channel. If a video filter has a channel override to block a specific YouTube channel, access to this channel is stopped only to this particular channel. If a user attempts to access the channel while surfing YouTube content, an error message appears telling the user that they must connect to the internet. If the user accesses the channel using the URL, a blocked replacement message shows up to confirm the reason why access is blocked.

Knowledge Check

1. Which is required by FortiGate to configure YouTube video filtering?
   - ✓ A. YouTube API key
   - B. Username

2. Which action in video filtering will prevent the videos from playing?
   - A. Deny
   - ✓ B. Block

FORTINET Training Institute

32

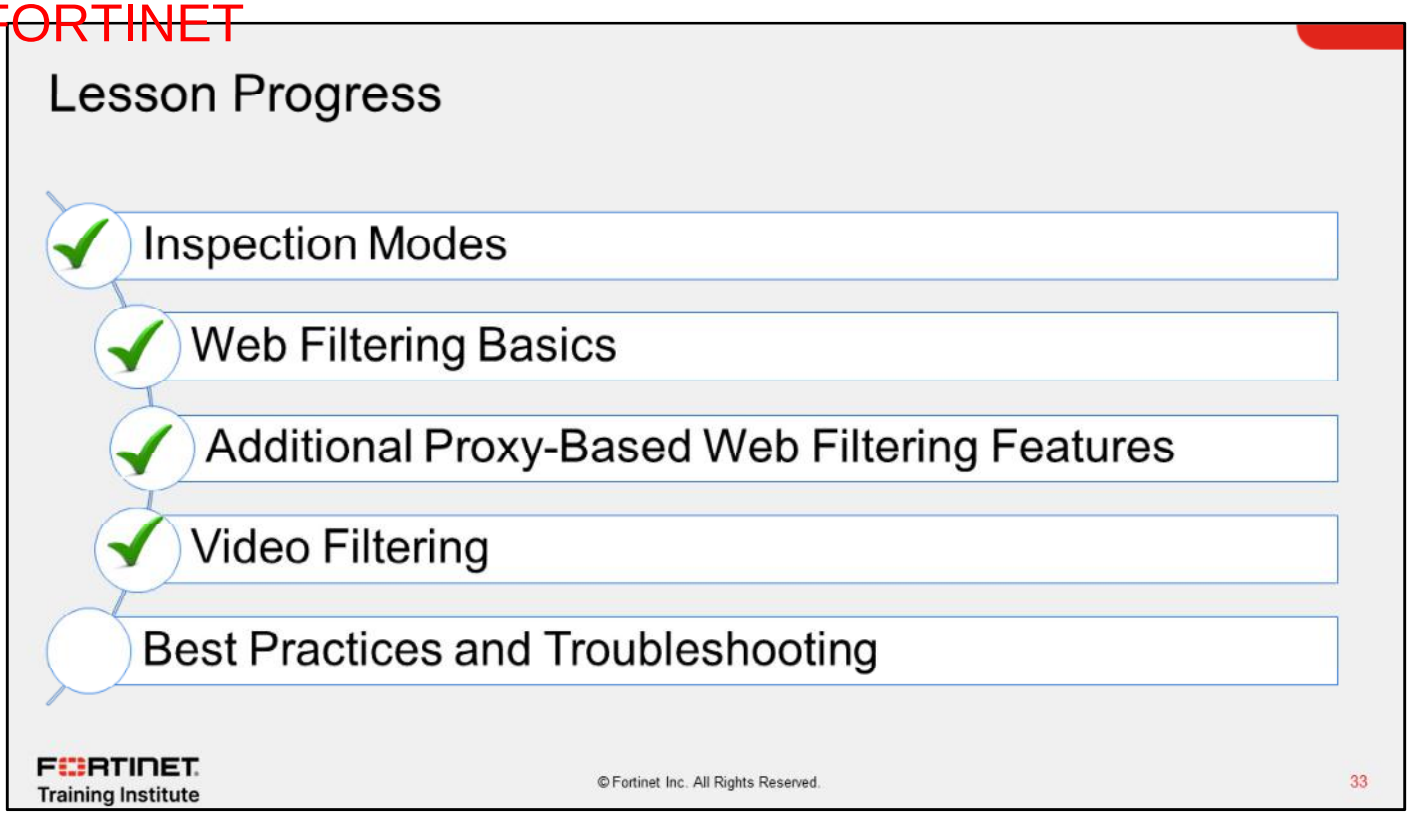

Good job! You now understand the video filtering feature.

Now, you will learn about best practices and troubleshooting.

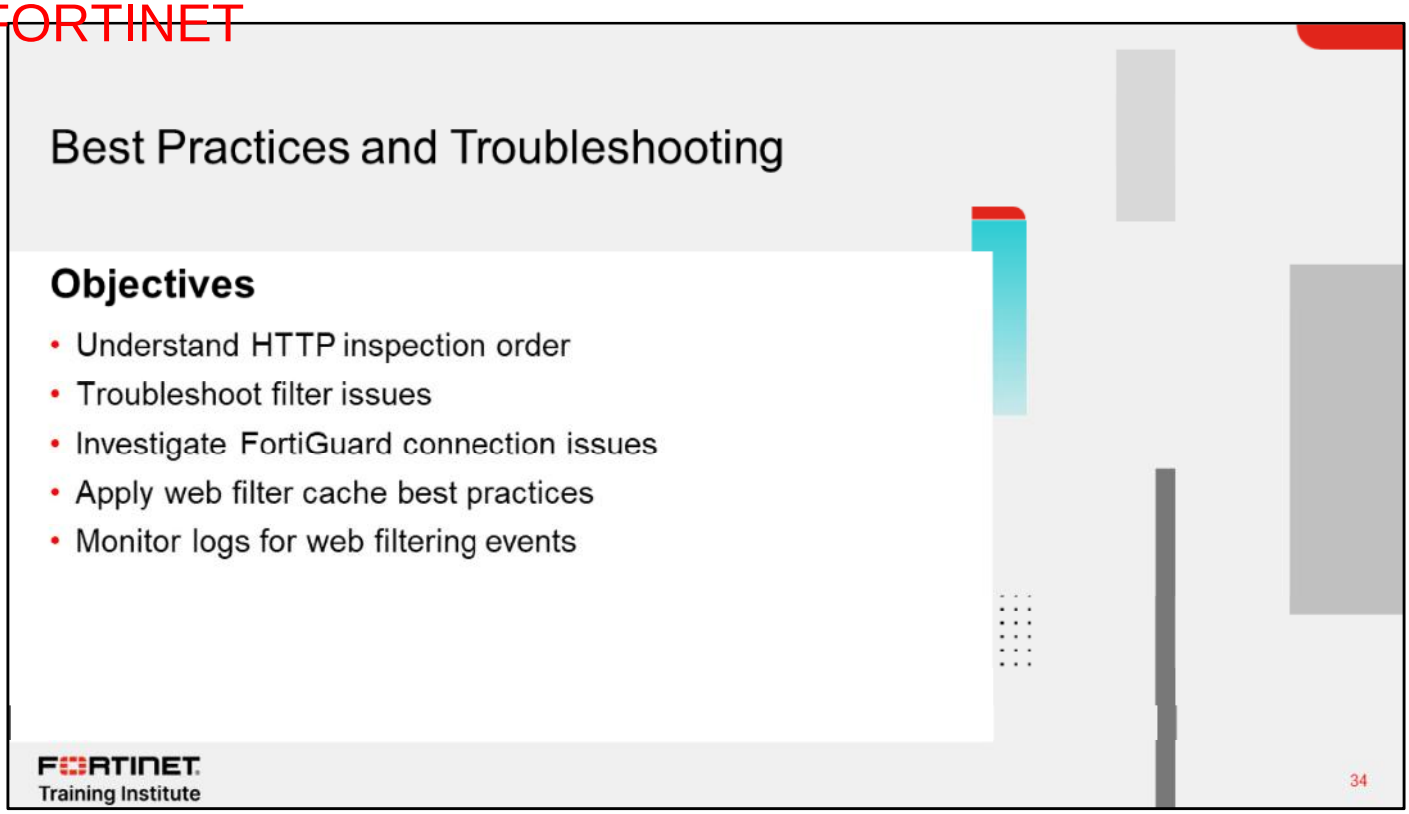

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in best practices and troubleshooting, you will be able to apply various best practices and troubleshooting techniques to avoid and investigate common issues.

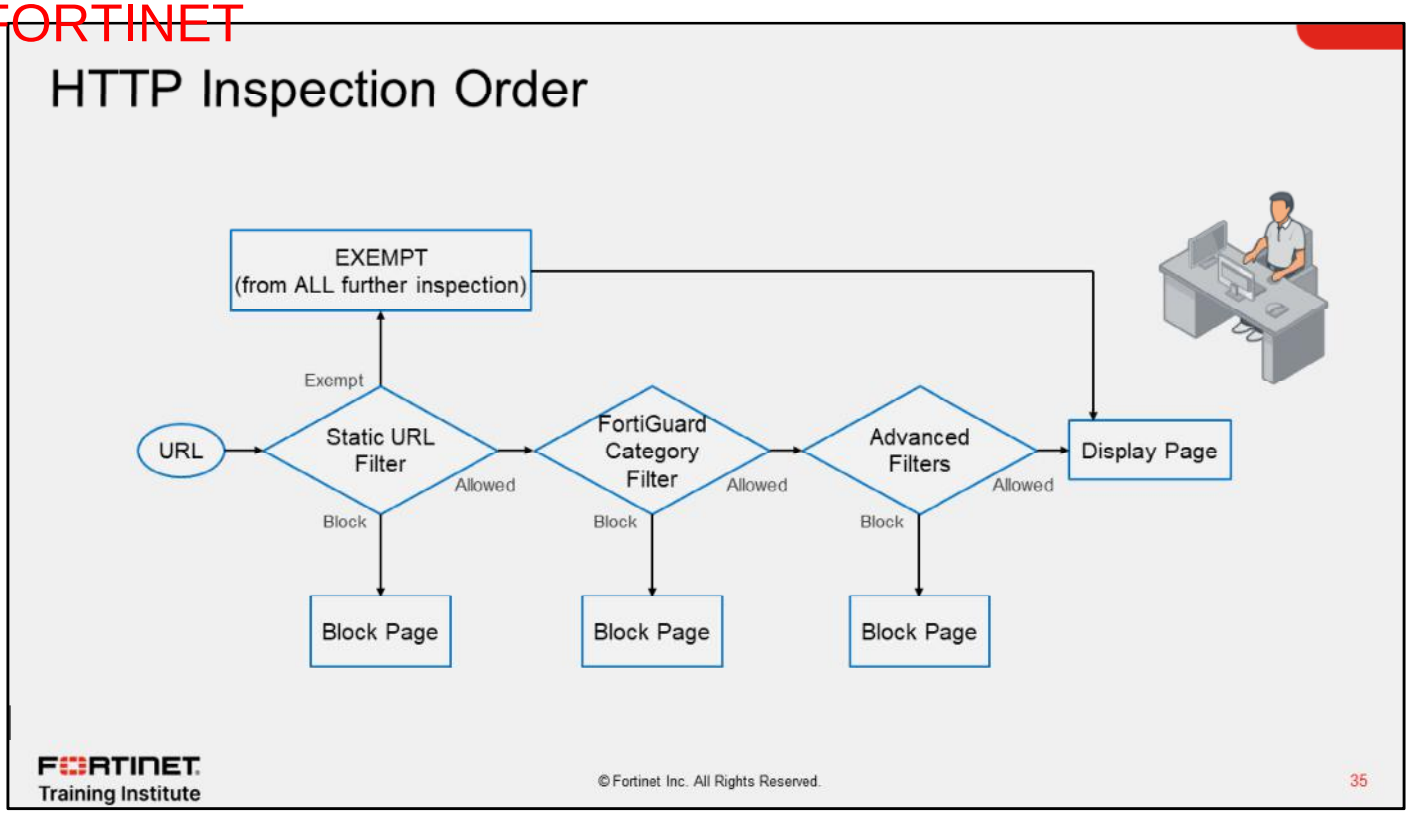

Remember that the web filtering profile has several features. So, if you have enabled many of them, the inspection order flows as follows:

1. The local static URL filter
2. FortiGuard category filtering (to determine a rating)
3. Advanced filters (such as safe search or removing Active X components)

For each step, if there is no match, FortiGate moves on to the next check enabled.

Apply the Filters

- It's not working. Why?
  - Did you apply the security profiles to the firewall policies?
  - Did you apply the SSL inspection profile, if needed?


```
config firewall policy
 edit 1
 set webfilter-profile <profile>
 next
end

config firewall profile-group
 edit <group name>
    set webfilter-profile <profile>
 next
end
```

FORTINET Training Institute

36

You have configured your security profiles, but they are not performing web inspection. Why?

Check to see if you have applied the security profiles to your firewall policies. Also, make sure that the SSL inspection profile is applied as needed.

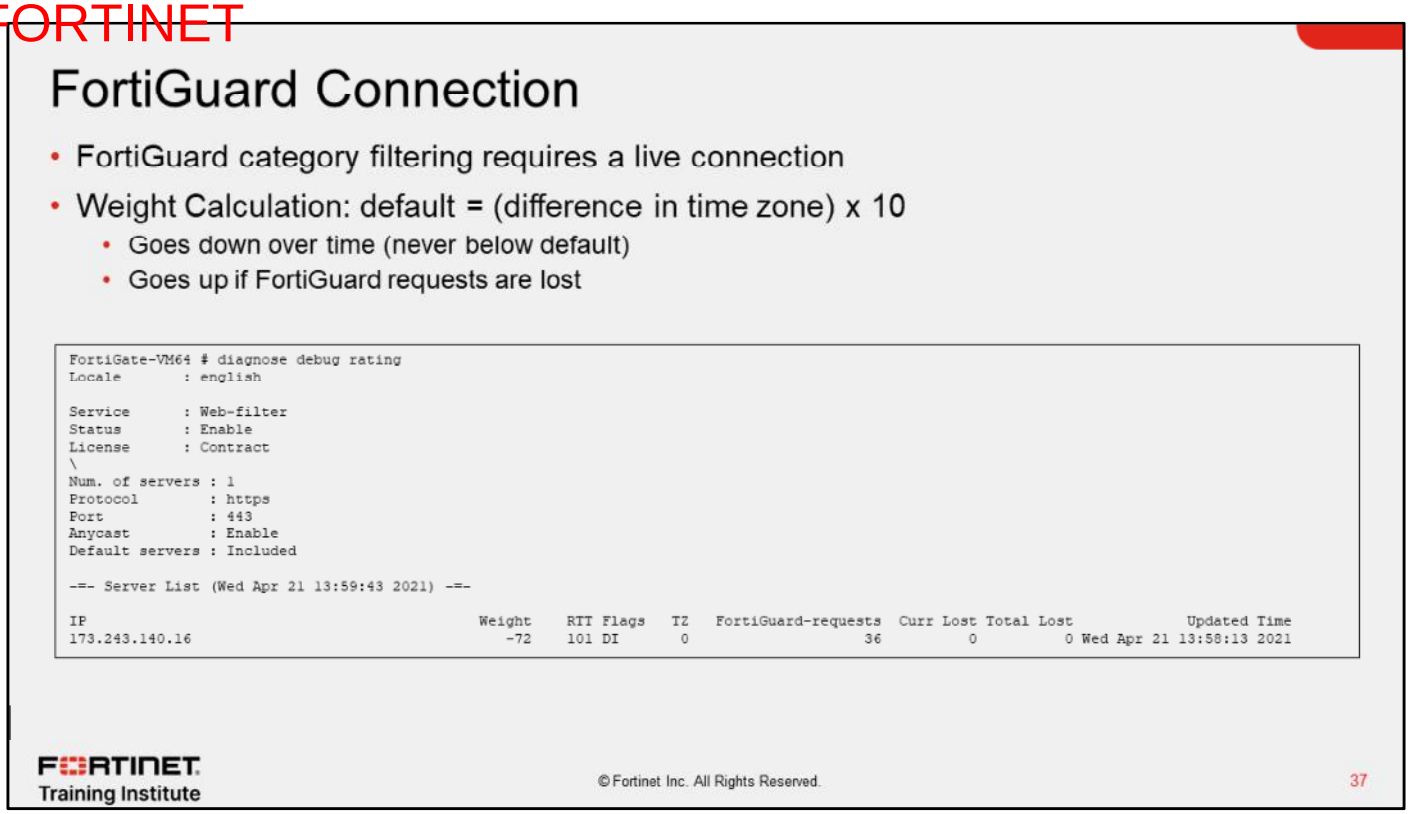

Category-based filtering requires a live connection to FortiGuard.

You can verify the connection to FortiGuard servers by running the `diagnose debug rating` CLI command. This command displays a list of FortiGuard servers you can connect to, as well as the following information:

- `Weight`: Based on the difference in time zone between FortiGate and this server (modified by traffic)
- `RTT`: Return trip time
- `Flags`: `D` (IP returned from DNS), `I` (Contract server contacted), `T` (being timed), `F` (failed)
- `TZ`: Server time zone
- `FortiGuard-requests`: The number of requests sent by FortiGate to FortiGuard
- `Curr Lost`: Current number of consecutive lost FortiGuard requests (in a row, resets to `0` when one packet succeeds)
- `Total Lost`: Total number of lost FortiGuard requests

The list is of variable length depending on the FortiGuard Distribution Network.

Web Filter Cache

- Improves performance by reducing requests to FortiGuard
- Cache is checked before sending a request to the FortiGuard server
  - FortiGate remembers response of visited websites
  - TTL settings control the number of seconds the query results are cached
  - Request is considered a rating error after timeout (15 seconds as default)
- HTTPS port 443 enforced by default FortiGuard or FortiManager communications
  - Disable FortiGuard anycast setting on CLI to use UDP ports 443, 53, or 8888
- Enabled by default—default TTL is 60 minutes (3600 seconds)

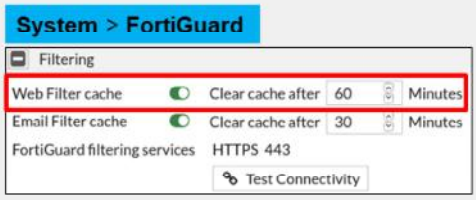


```
config system fortiguard
    set fortiguard-anycast {enable|disable}
    set protocol {udp|https}
    set port {8888|53|443}
    set webfilter-timeout {<1> - <30>}
end
```

FORTINET Training Institute

38

FortiGate can maintain a list of recent website rating responses in memory. So, if the URL is already known, FortiGate doesn't send back a rating request.

By default, FortiGate is configured to enforce the use of HTTPS port 443 to perform live filtering with FortiGuard or FortiManager. Other ports and protocols are available by disabling the FortiGuard anycast setting on the CLI. These ports and protocols to query the servers (FortiGuard or FortiManager) HTTPS port 53 and port 8888, UDP port 443, port 53, and port 8888. If you are using UDP port 53, any kind of inspection reveals that this traffic is not DNS and prevents the service from working. In this case, you can switch to the alternate UDP port 443 or port 8888, or change the protocol to HTTPS, but these ports are not guaranteed to be open in all networks, so you must check beforehand.

Caching responses reduces the amount of time it takes to establish a rating for a website. Also, memory lookup is much quicker than packets travelling on the internet.

The timeout defaults to 15 seconds, but you can set it as high as 30 seconds, if necessary.

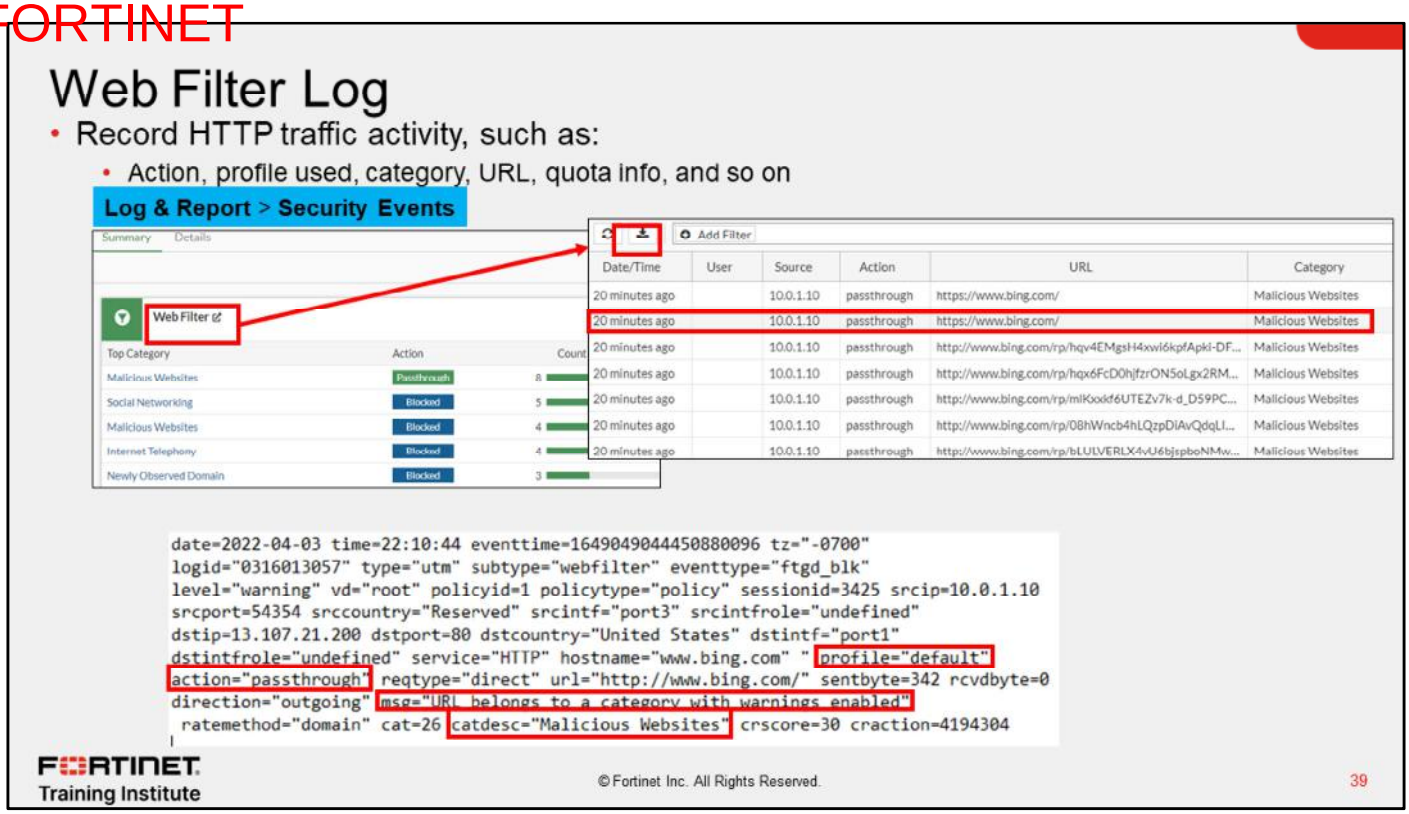

Now, take a look at the web filter log and report feature.

This slide shows an example of a log message. Access details include information about the FortiGuard quota and category (if those are enabled), which web filter profile was used to inspect the traffic, the URL, and more details about the event.

You can also view the raw log data by clicking the download icon at the top of the GUI. The file downloaded is a plain text file in a syslog format.

Knowledge Check

1. You have configured your security profiles, but they are not performing web or DNS inspection. Why?
   - A. The certificate is not installed correctly.
   - ✓ B. The profile is not associated with the correct firewall policy.

FORTINET Training Institute

40

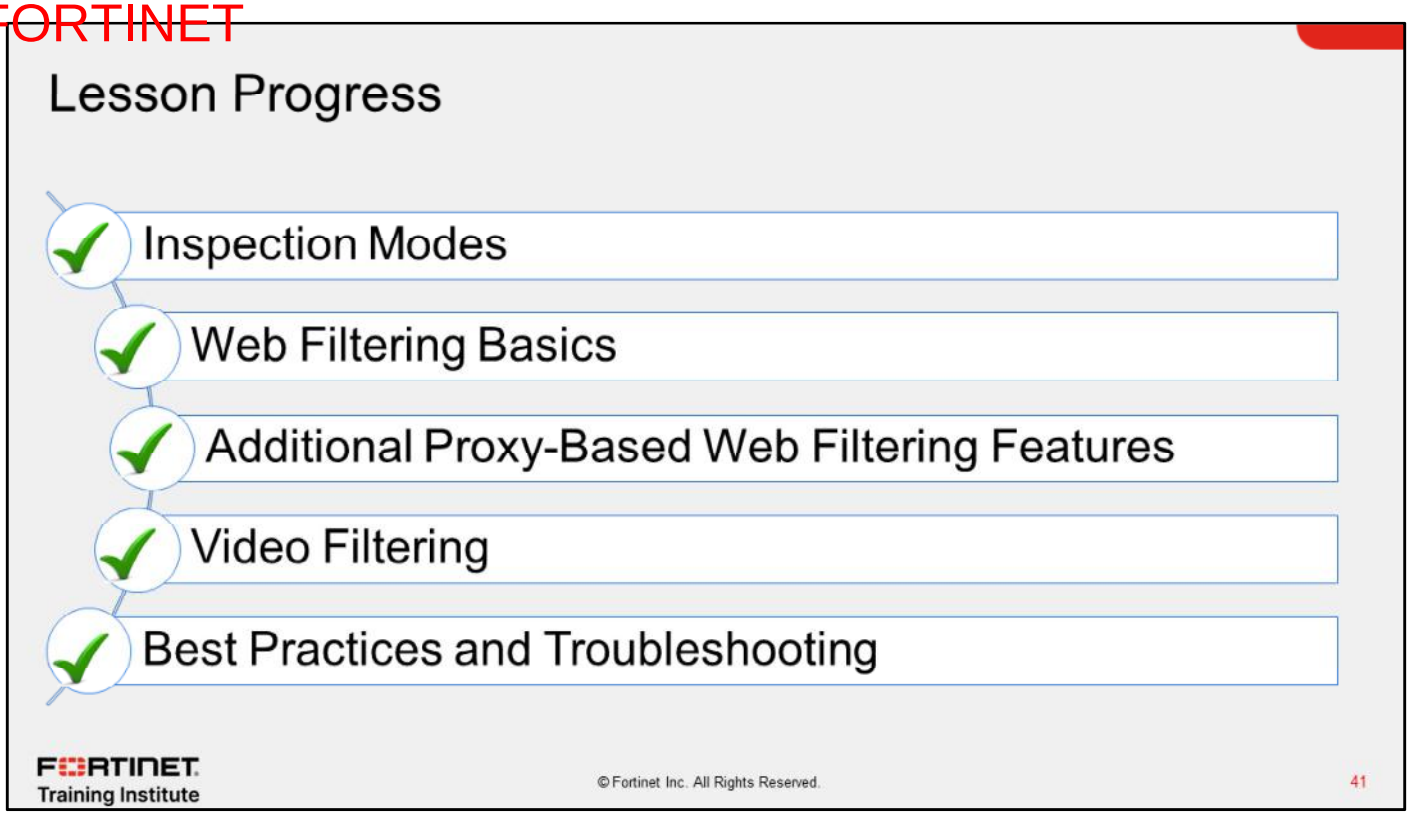

Congratulations! You have completed this lesson.

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

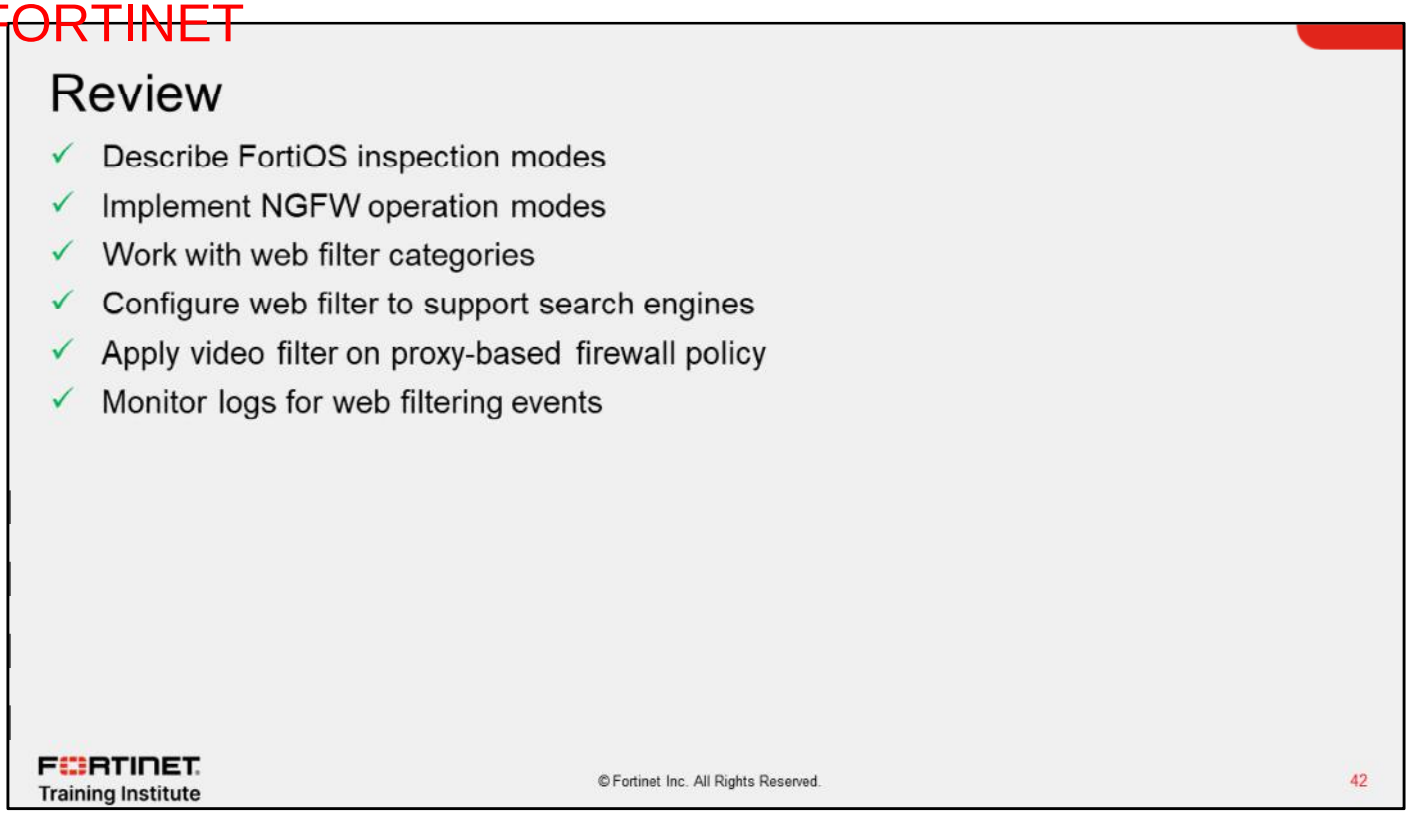

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure web filtering on FortiGate to control web traffic in your network.

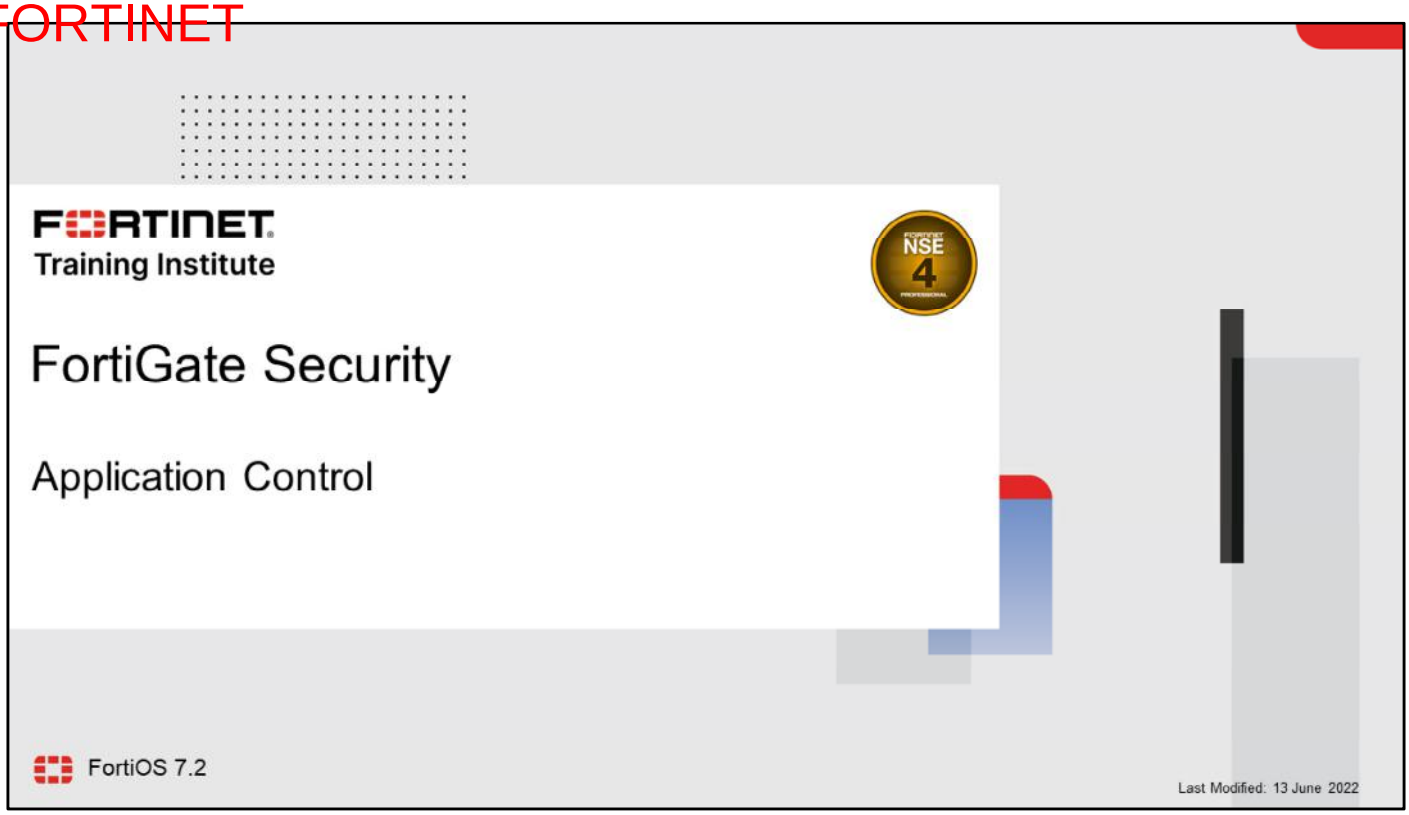

In this lesson, you will learn how to monitor and control network applications that may use standard or non-standard protocols and ports—beyond simply blocking or allowing a protocol, port number, or IP address.

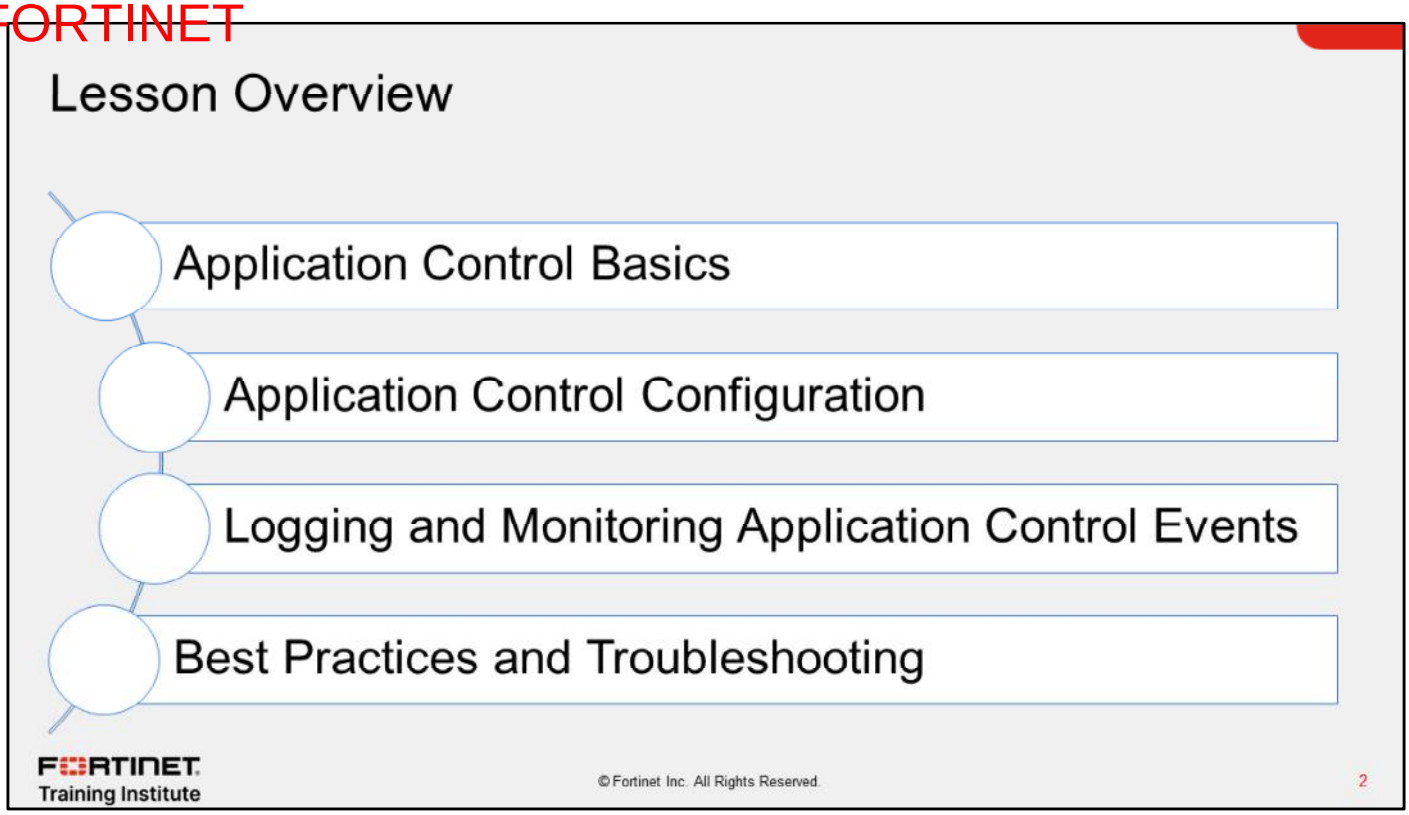

In this lesson, you will learn about the topics shown on this slide.

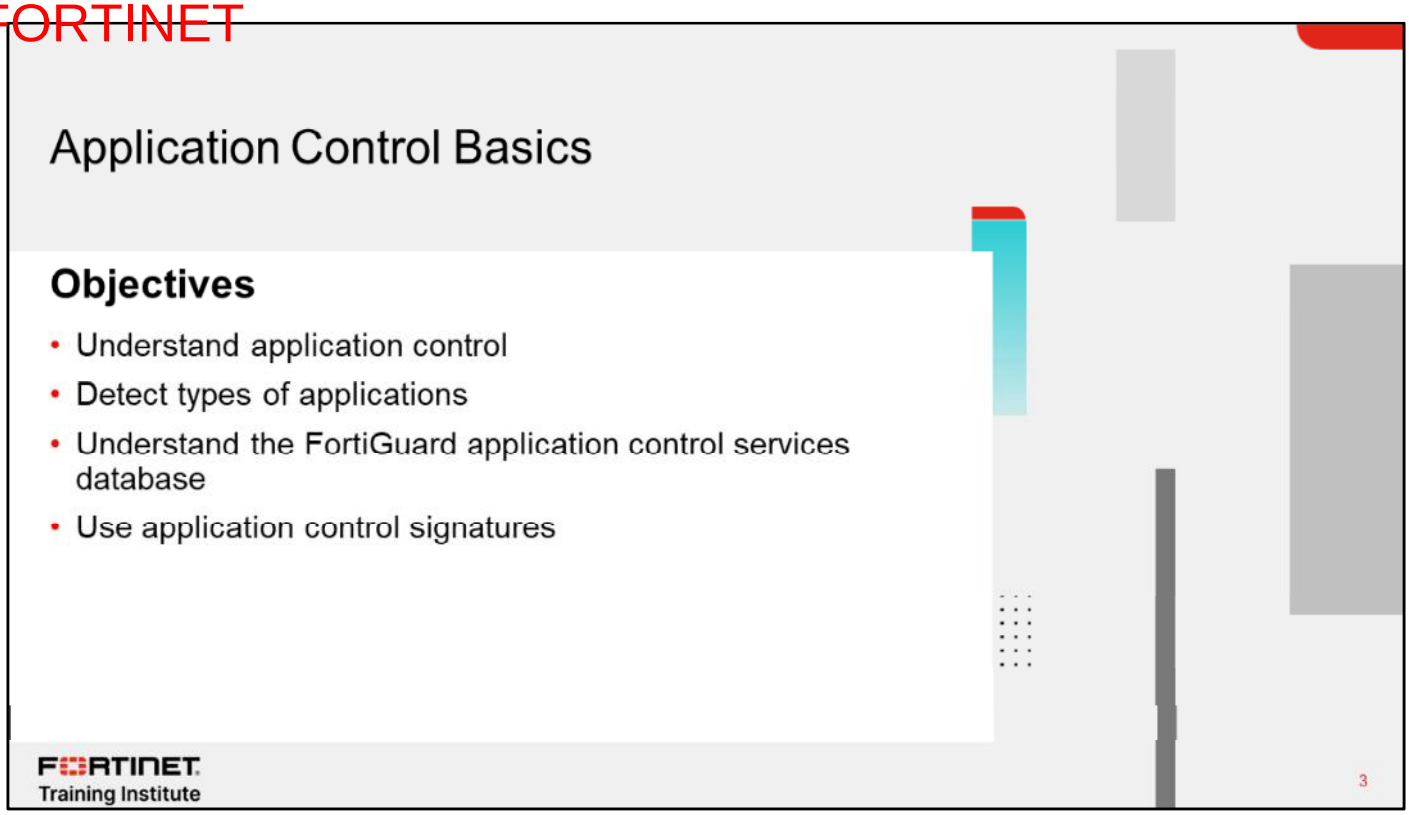

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control basics, you will be able to understand how application control works on FortiGate.

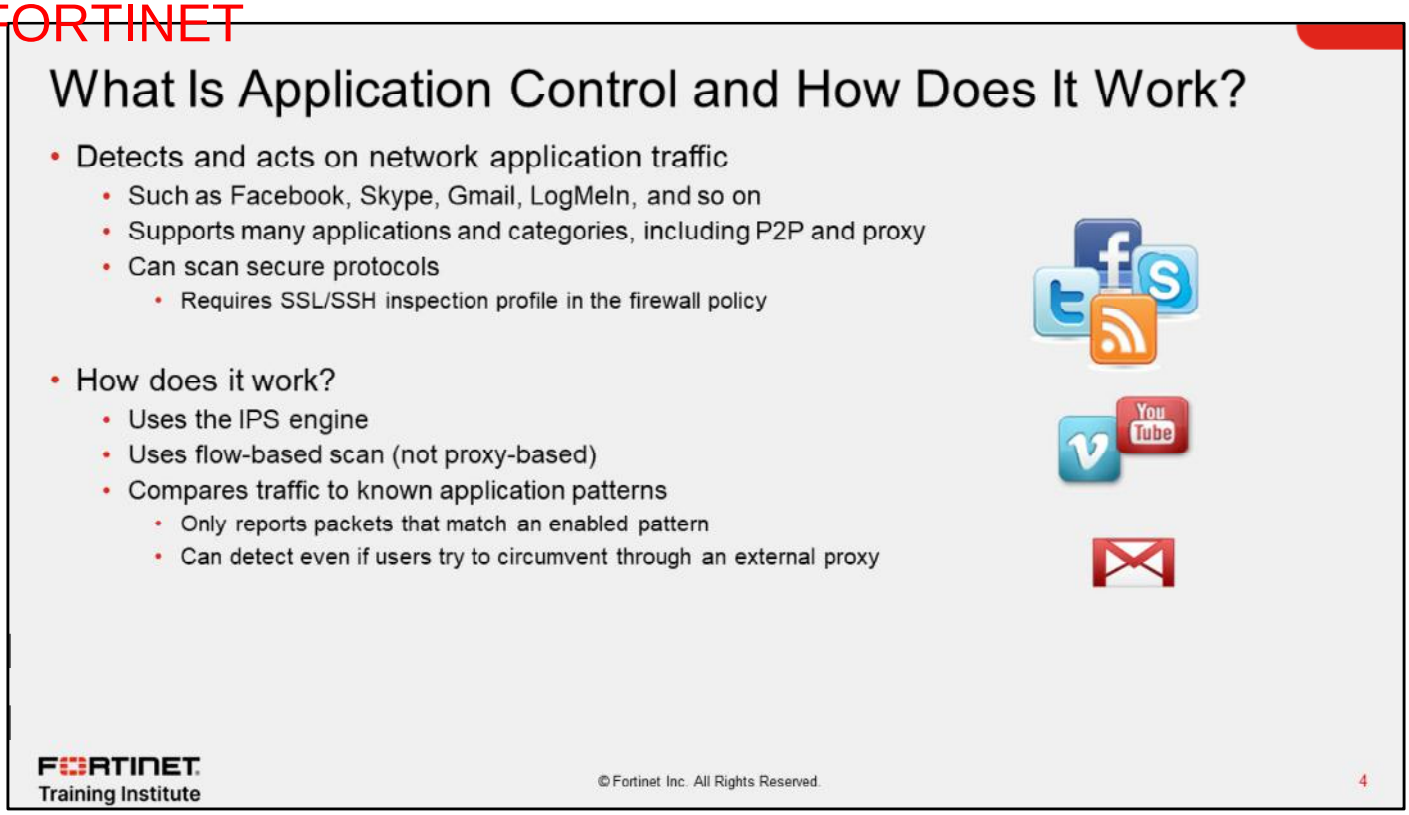

Application control detects applications—often applications that consume a lot of bandwidth—and allows you to take appropriate action related to application traffic, such as monitoring, blocking, or applying traffic shaping.

Application control identifies applications, such as Google Talk, by matching known patterns to the application's transmission patterns. Therefore, an application can be accurately identified, only if its transmission pattern is unique. However, not every application behaves in a unique way. Many applications reuse pre-existing, standard protocols and communication methods. For example, many video games, such as *World of Warcraft,* use the BitTorrent protocol to distribute game patches.

Application control can be configured in proxy-based and flow-based firewall policies. However, because application control uses the IPS engine, which uses flow-based inspection, inspection is always flow-based. By comparison, when applying web filtering and antivirus through an HTTP proxy, the proxy first parses HTTP and removes the protocol, and then scans only the payload inside.

Why does FortiGate use a flow-based scan for application control?

Unlike other forms of security profiles, such as web filtering or antivirus, application control is not applied by a proxy. It uses an IPS engine to analyze network traffic and detect application traffic, even if the application is using standard or non-standard protocols and ports. It doesn't operate using built-in protocol states. It matches patterns in the entire byte stream of the packet, and then looks for patterns.

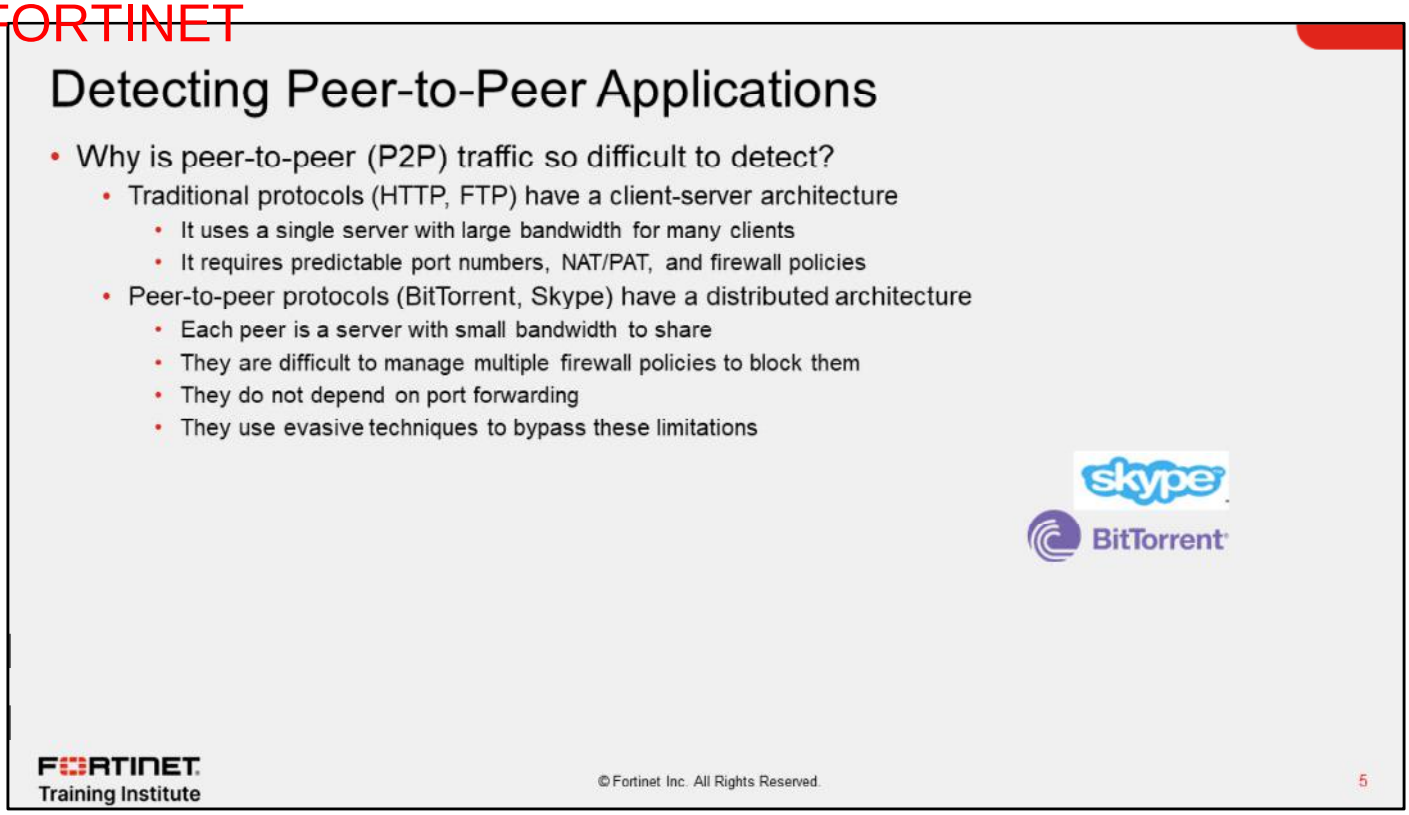

When HTTP and other protocols were designed, they were designed to be easy to trace. Because of that, administrators could easily give access to single servers behind NAT devices, such as routers and, later, firewalls.

But when P2P applications were designed, they had to be able to work without assistance—or cooperation—from network administrators. In order to achieve this, the designers made P2P applications able to bypass firewalls and incredibly hard to detect. Port randomization, pinholes, and changing encryption patterns are some of the techniques that P2P protocols use.

These techniques make P2P applications difficult to block using a firewall policy, and also make them difficult to detect by proxy-based inspection.

Flow-based inspection using the IPS engine can analyze packets for pattern matching, and then look for patterns to detect P2P applications.

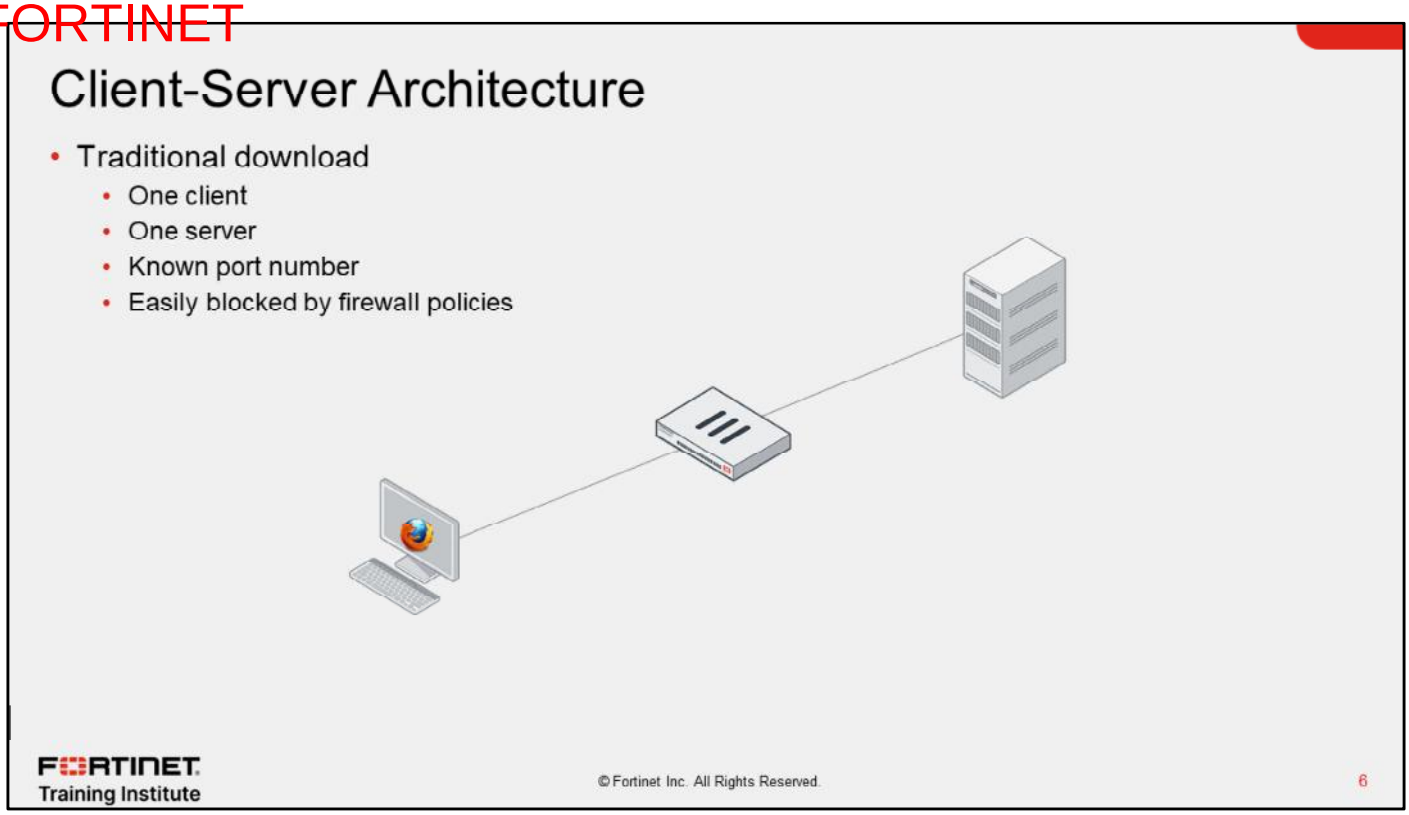

This slide shows a traditional, client-server architecture. There may be many clients of popular sites, but often, such as with an office file server, it's just one client and one server.

Traditional downloads use a defined protocol over a standard port number. Whether it's from a web or FTP site, the download is from a single IP address, to a single IP address. So, blocking this kind of traffic is easy: you only need one firewall policy.

But, it's more difficult to block traffic from peer-to-peer downloads. Why?

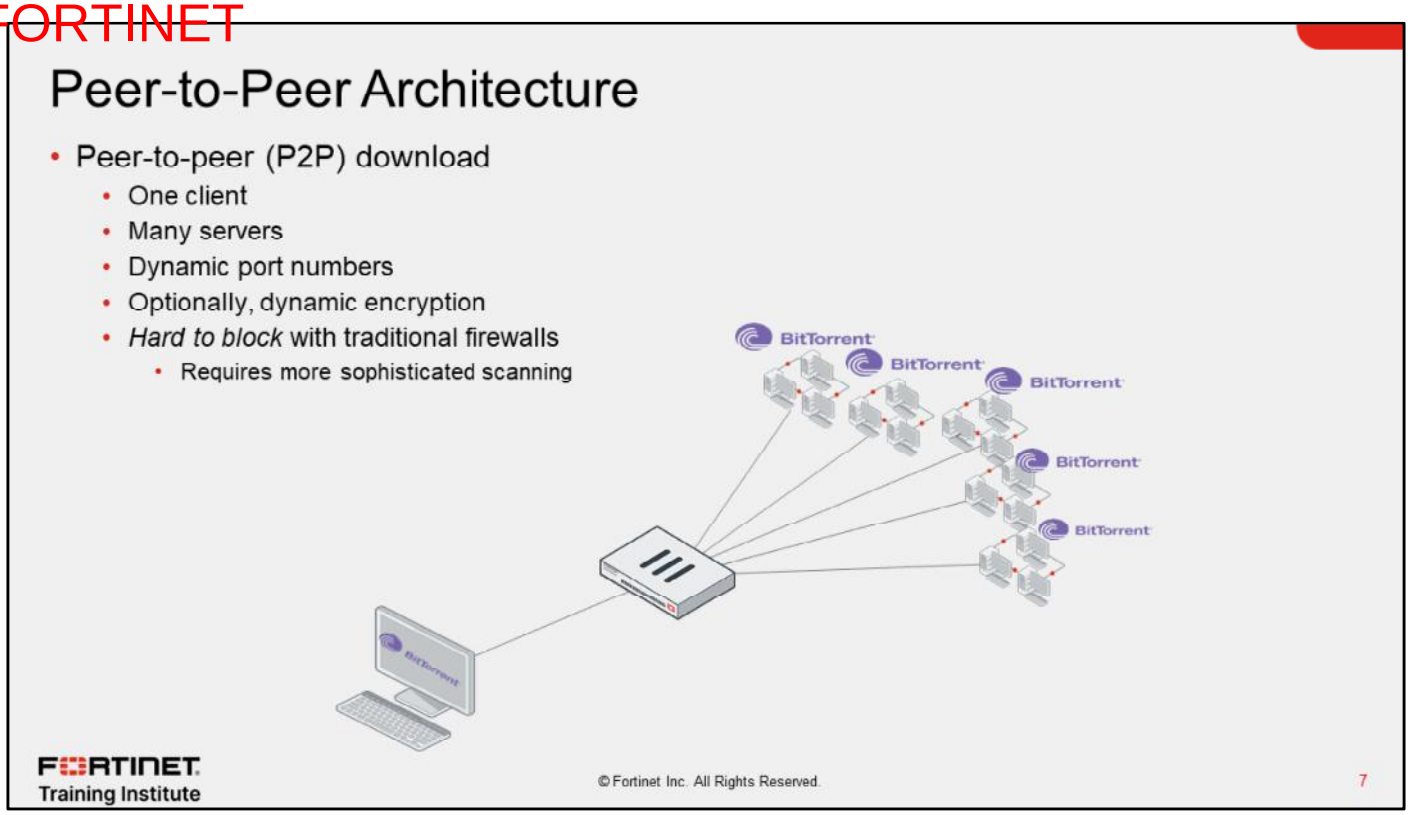

Peer-to-peer (P2P) downloads divide each file among multiple (theoretically unlimited) peers. Each peer delivers part of the file. While having many clients is a disadvantage in client-server architectures, it is an advantage for P2P architecture because, as the number of peers increases to n, the file is delivered n times faster.

Because popularity increases the speed of delivery—unlike traditional client-server architecture where popularity could effectively cause a denial of service (DoS) attack on the server—some software, such as BitTorrent distributions of Linux, and games distributing new patches, leverage this advantage. Even if each client has little bandwidth, together they can offer more bandwidth for the download than many powerful servers.

Consequently, in order to download the file, the requesting peer can consume much more bandwidth per second than it would from only a single server. Even if there is only one peer in your network, it can consume unusually large amounts of bandwidth. Because the protocols are usually evasive, and there will be many sessions to many peers, they are difficult to completely block.

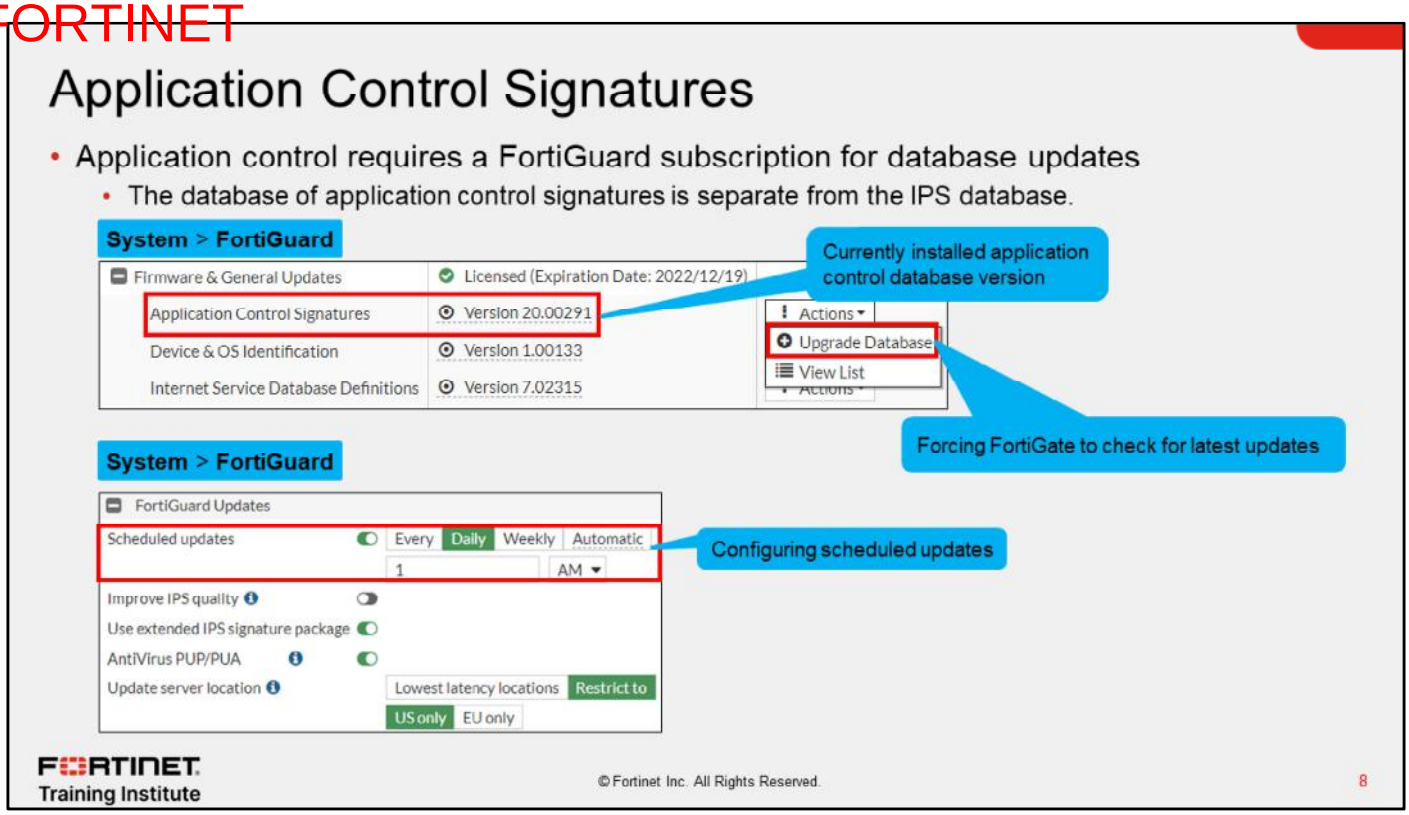

Before you try to control applications, it's important to understand the signatures used by application control.

How does application control detect the newest applications and changes to application protocols?

Application control updates come as part of the standard FortiCare support contract, but it requires a subscription for database updates. The database for application control signatures is separate from the intrusion prevention system (IPS) database. You can configure FortiGate to automatically update its application control signature database on the FortiGuard page. The application control signature database information is also displayed on the FortiGuard page.

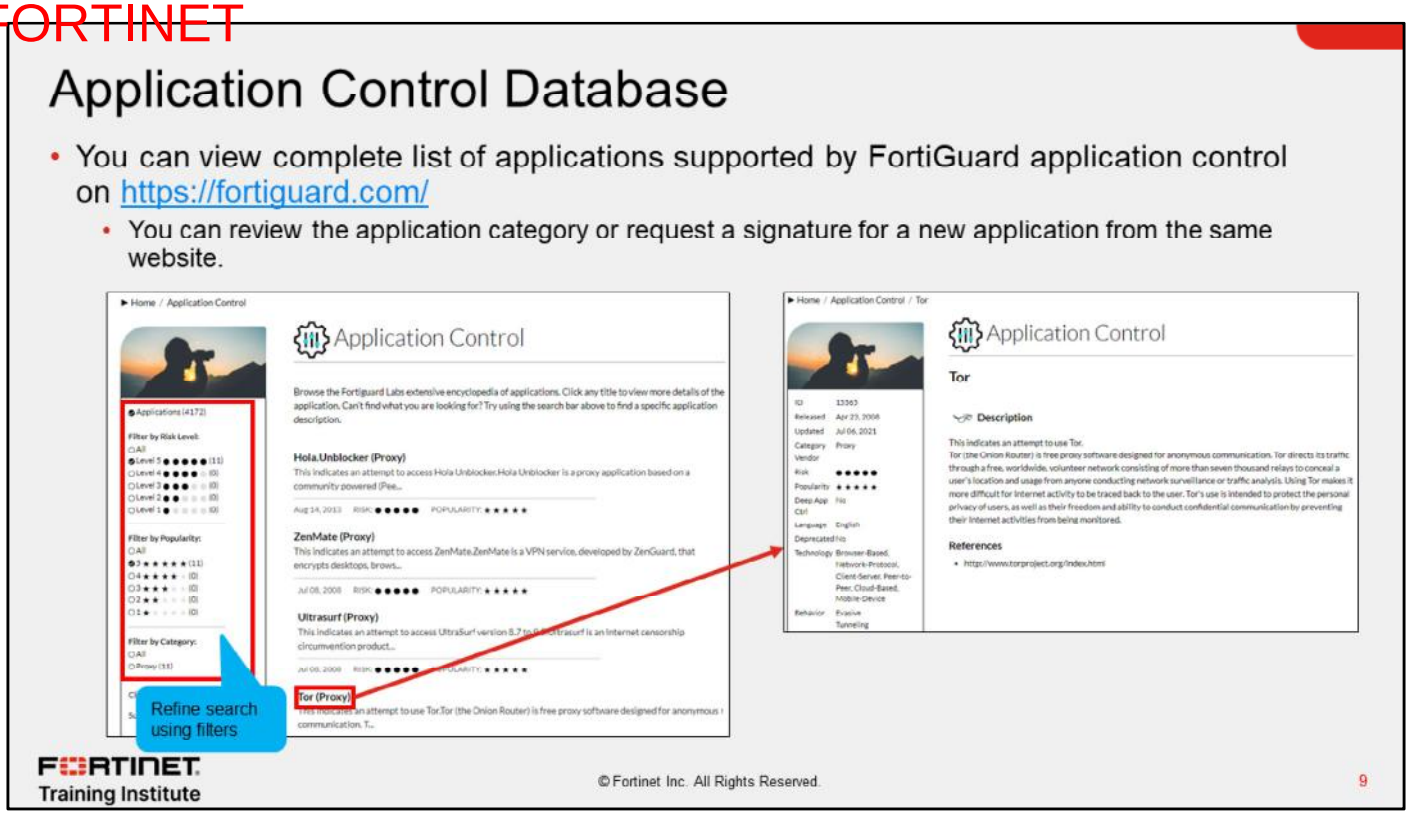

You can view the latest version of the application control database on the FortiGuard website, or by clicking an individual application signature in the application control profile.

The application control database provides details about application control signatures based on category, popularity, and risk, to name a few.

When building an application control signature, the FortiGuard security research team evaluates the application and assigns a risk level based on the type of security risk. The rating is Fortinet-specific, and not related to the common vulnerability scoring system (CVSS) or other external systems. The rating can help you decide whether or not to block an application.

On the FortiGuard website, you can read details about each signature's related application.

This slide shows an example for an application called Tor. Tor is a web proxy, so it belongs in the proxy category. A best practice is to create test policies that you can use to observe policy behavior.

If the most recent FortiGuard update does not include a definition for an application that you need to control, you can submit a request on the FortiGuard website to have the application added. You can also submit a request to re-evaluate an application category, if you believe an application should belong to a different category.

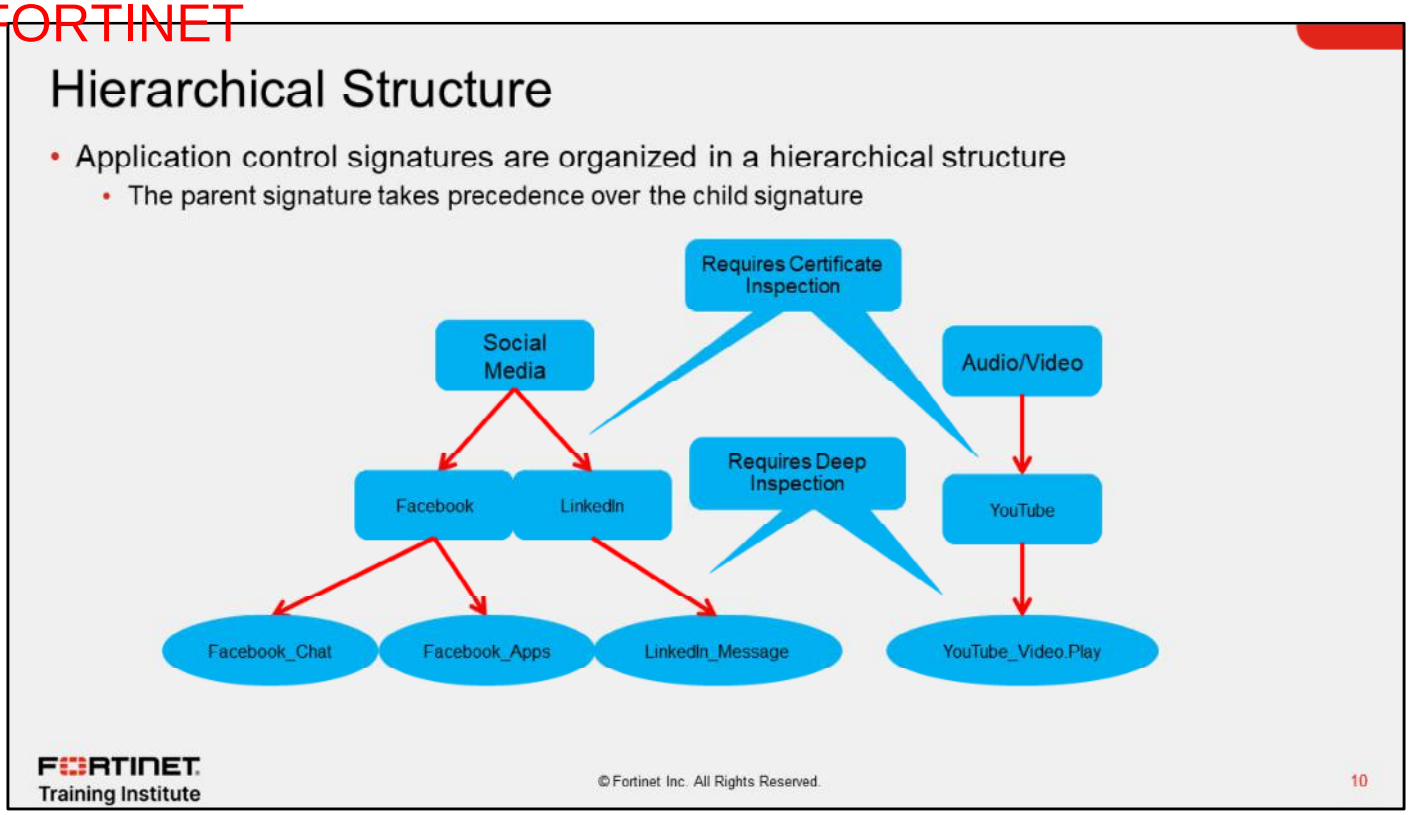

Many web applications offer functionality that can be embedded in third-party websites or applications. For example, you can embed a Facebook **Like** button at the end of an article, or reference a YouTube video on an educational website. FortiOS gives administrators all the tools they need to inspect sub-application traffic. The FortiGuard application control signature database is organized in a hierarchical structure. This gives you the ability to inspect the traffic with more granularity. You can block Facebook applications while allowing users to collaborate using Facebook chat.

Knowledge Check

1. Which statement about application control is true?
   - ✓ A. Application control uses the IPS engine to scan traffic for application patterns.
   - B. Application control is unable to scan P2P architecture traffic.

2. Which statement about the application control database is true?
   - ✓ A. The application control database is separate from the IPS database.
   - B. The application control database must be updated manually.

FORTINET Training Institute

11

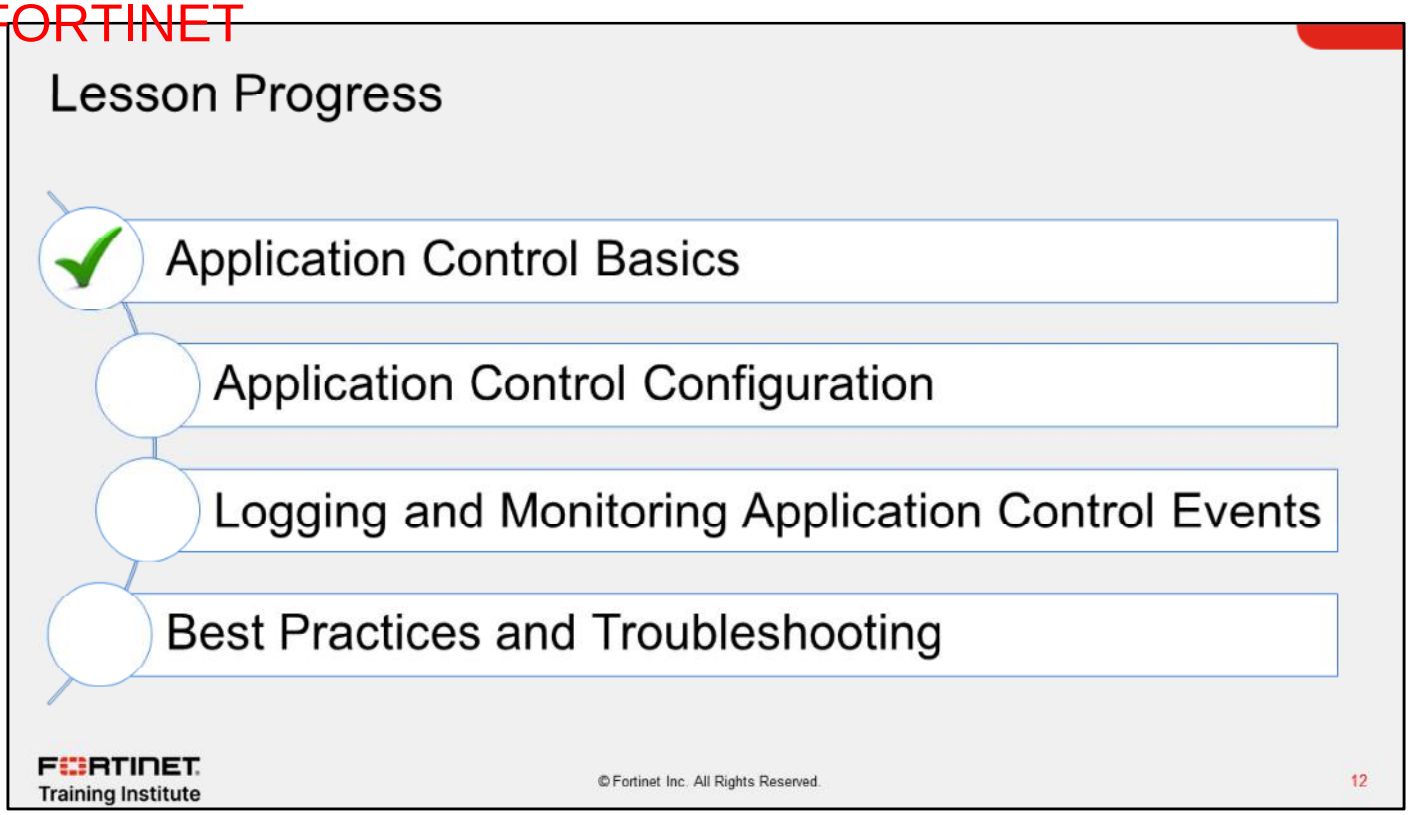

Good job! You now understand basic application control functionality.

Now, you will learn about application control configuration.

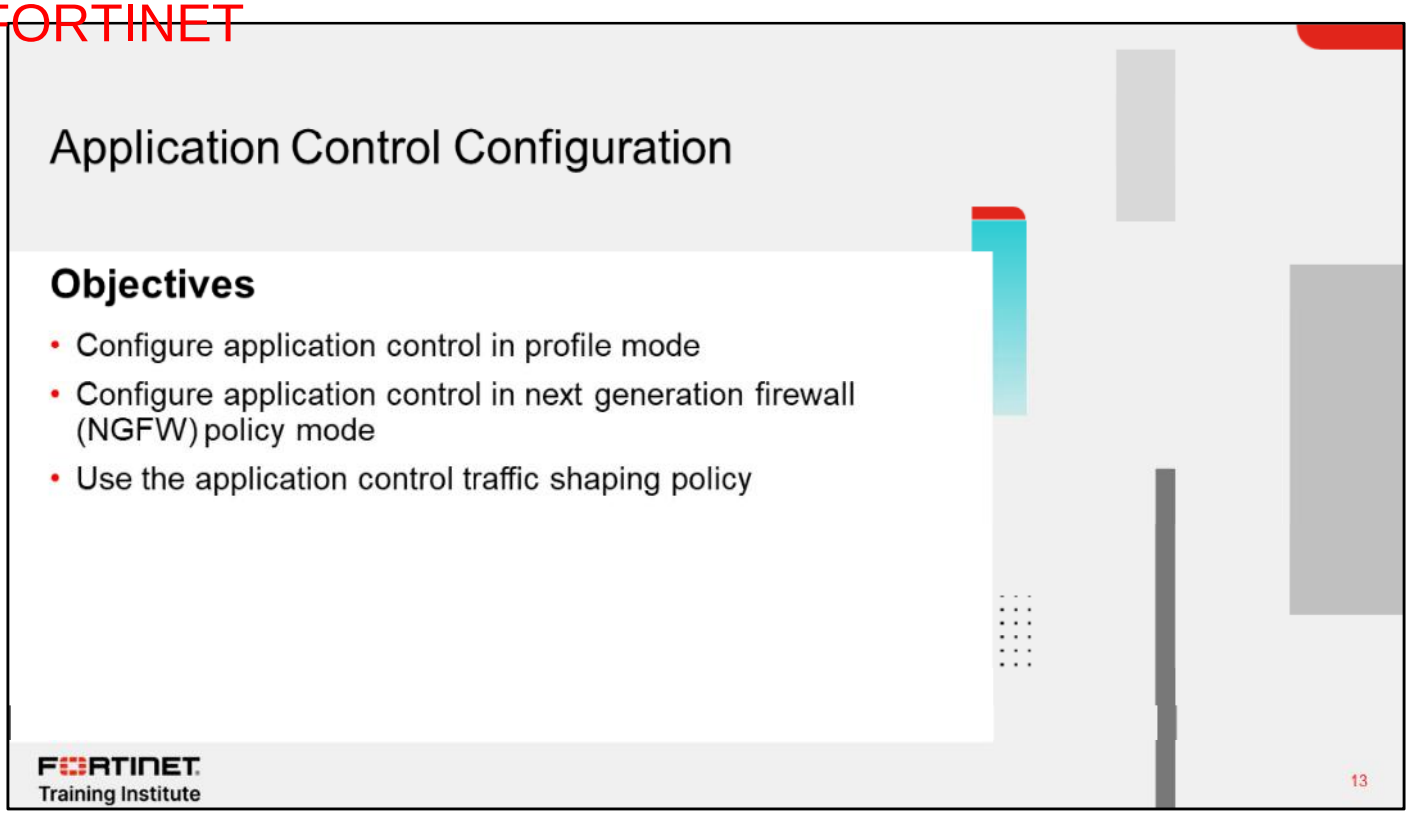

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in configuring the application control operation modes that are available on FortiOS, you will be able to use application control effectively in both profile mode and NGFW policy mode.

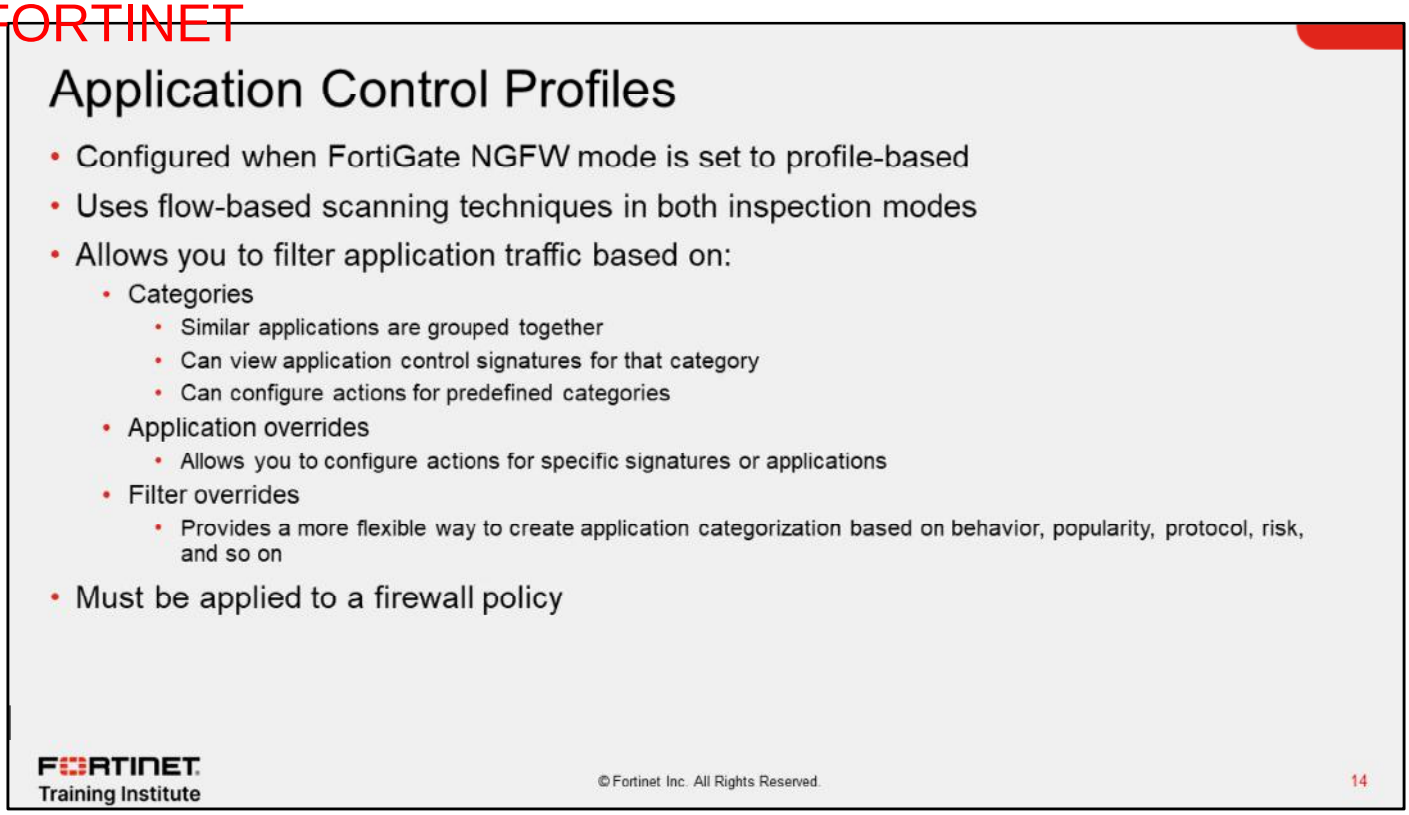

When FortiGate or a VDOM is operating in flow-based (NGFW mode set to profile-based, policy set to flow-based) inspection mode or policy set to proxy-based inspection mode, to configure application control, administrators must create an application control *profile* and apply that profile to a firewall policy.

It is important to note that the application control profile uses flow-based scanning techniques, regardless of which inspection mode is used on the policy.

The application control profile consists of three different types of filters:

- Categories: Groups applications based on similarity. For example, all applications that are capable of providing remote access are grouped in the **Remote Access** category. You can view the signatures of all applications in a category or apply an action to a category as a whole.
- Application overrides: Provides the flexibility to control specific signatures and applications.
- Filter overrides: Useful when a predefined category does not meet your requirements and you want to block all applications based on criteria that is not available in categories. You can configure the categorization of applications based on behavior, popularity, protocol, risk, vendor, or the technology used by the applications, and take action based on that.

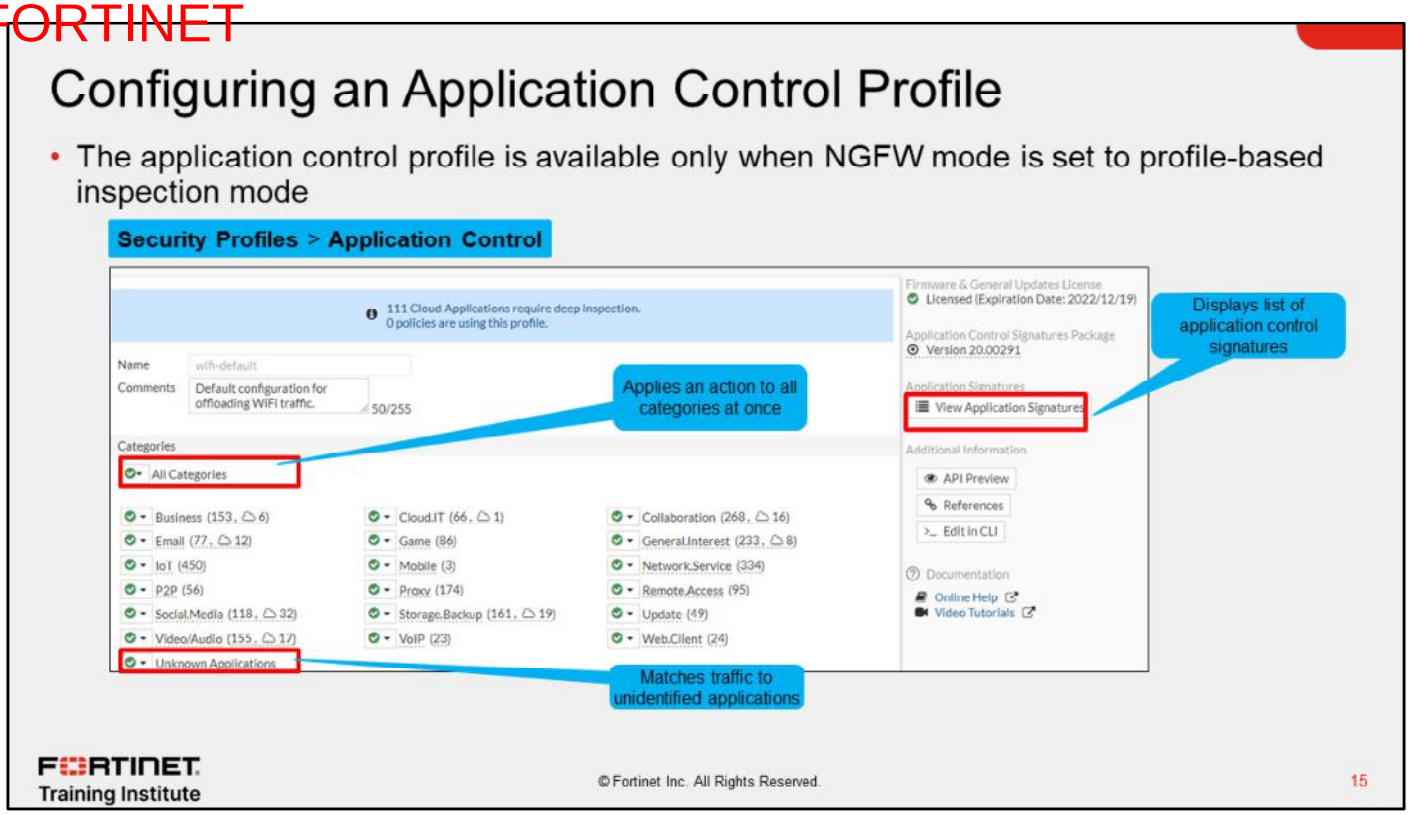

The application control profile is configured on the **Application Control** page. You can configure actions based on categories, application overrides, and filter overrides. You can also view the list of application control signatures by clicking **View Application Signatures**.

At the top of the **Application Control** profile page, you will see a summary of how many cloud applications require deep inspection. Cloud applications that use SSL encryption cannot be scanned without a deep inspection profile. FortiGate must decrypt the traffic in order to perform inspection and control application traffic.

The **Unknown Applications** setting matches traffic that can't be matched to any application control signature and identifies the traffic as `unknown application` in the logs. Factors that contribute to traffic being identified as `unknown application` include:

- How many rare applications your users are using
- Which IPS database version you are using

Identifying traffic as unknown can cause frequent log entries. Frequent log entries decrease performance.

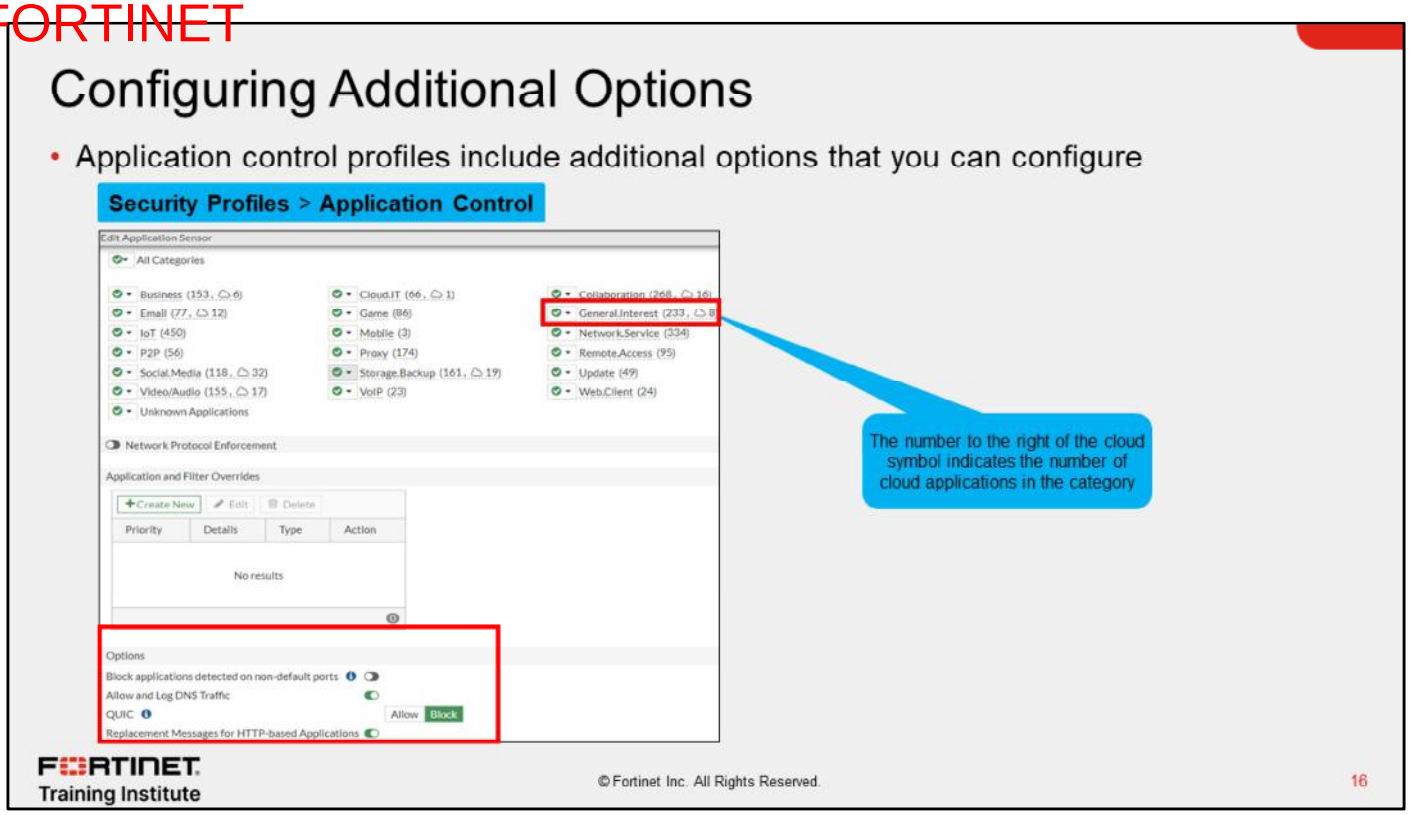

The number listed to the right of the cloud symbol indicates the number of cloud applications in the category.

Allow and Log DNS Traffic: Enable this option to allow DNS traffic for the application sensor. Depending on the application and how often it queries DNS servers, enabling this setting can use significant system resources.

QUIC: QUIC is a protocol from Google that uses UDP Instead of the standard TCP connections for web access. UDP is not scanned by web filtering. Allowing QUIC instructs FortiGate to inspect Google Chrome packets for a QUIC header and to generate logs as QUIC messages. Blocking QUIC forces Google Chrome to use HTTP2/TLS1.2 and FortiGate to log QUIC as blocked. The default action for QUIC is **Block**.

Replacement Messages for HTTP-based Applications: This setting allows you to replace blocked content from HTTP/HTTPS applications with an explanation for the user's benefit. For non-HTTP/HTTPS applications, FortiGate only drops the packets or resets the TCP connection.

After you've configured the application control profile, select the profile in the firewall policy. Like any other security profile, the settings you configure in the application control profile are not applied globally. FortiGate applies the application control profile settings only to traffic governed by the firewall policy in which you've selected the application control profile. This allows granular control.

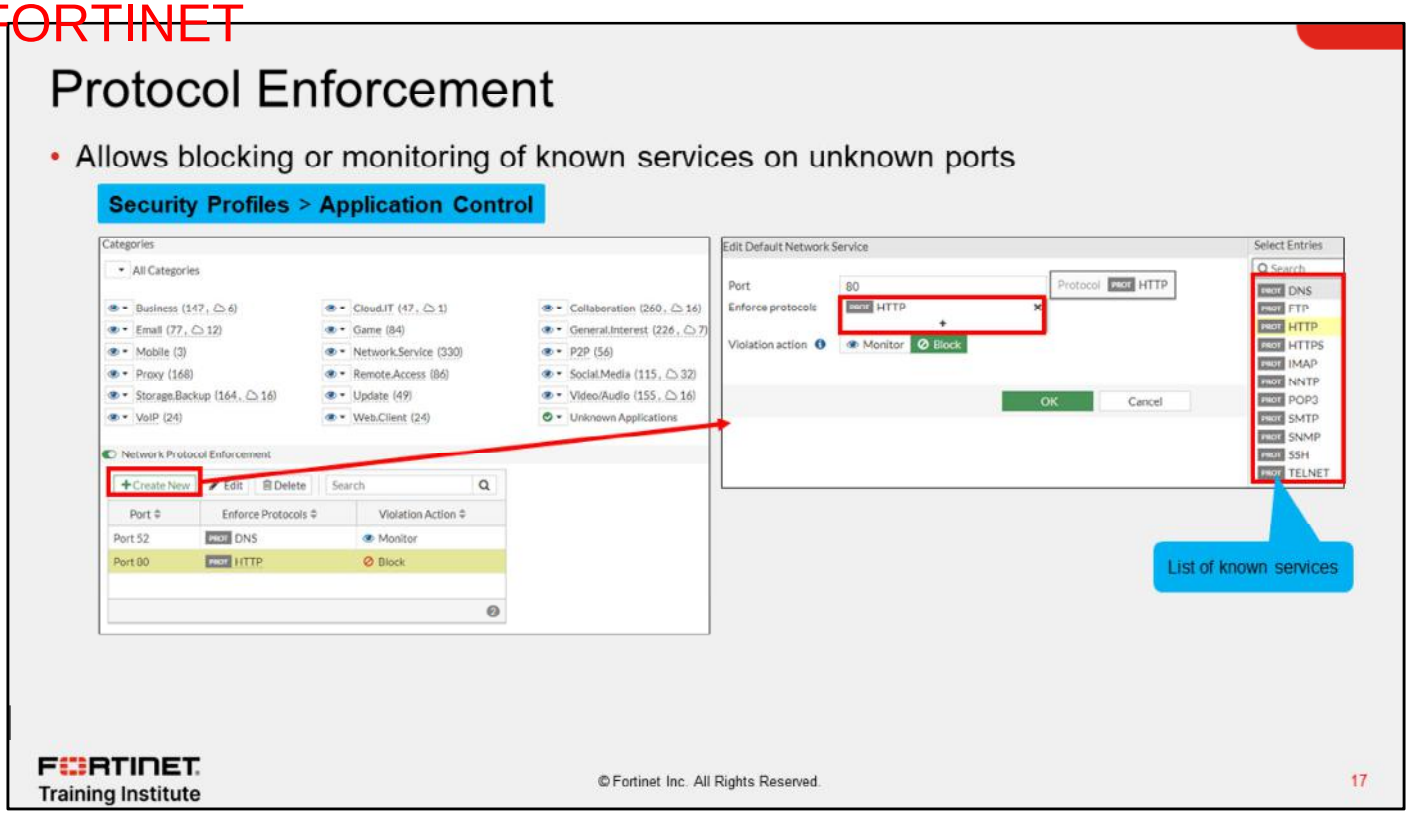

Protocol enforcement is added to the application control profile, allowing the administrator to configure network services (for example, FTP, HTTP, and HTTPS) on known ports (for example, 21, 80, and 443), while blocking those services on other ports.

The feature takes action in the following scenarios:

- When one protocol dissector confirms the service of network traffic, protocol enforcement can check whether the confirmed service is whitelisted under the server port. If it is not, then the traffic is considered a violation and IPS can take the action (for example, block) specified in the configuration.
- There is no confirmed service for network traffic. It would be considered a service violation if IPS dissectors rule out all the services enforced under its server port. For example, if port 21 is configured for FTP and IPS Dissector could not decide on the exact service but is sure it is not FTP. If the port of the non-ftp traffic is 21, it will be a violation.

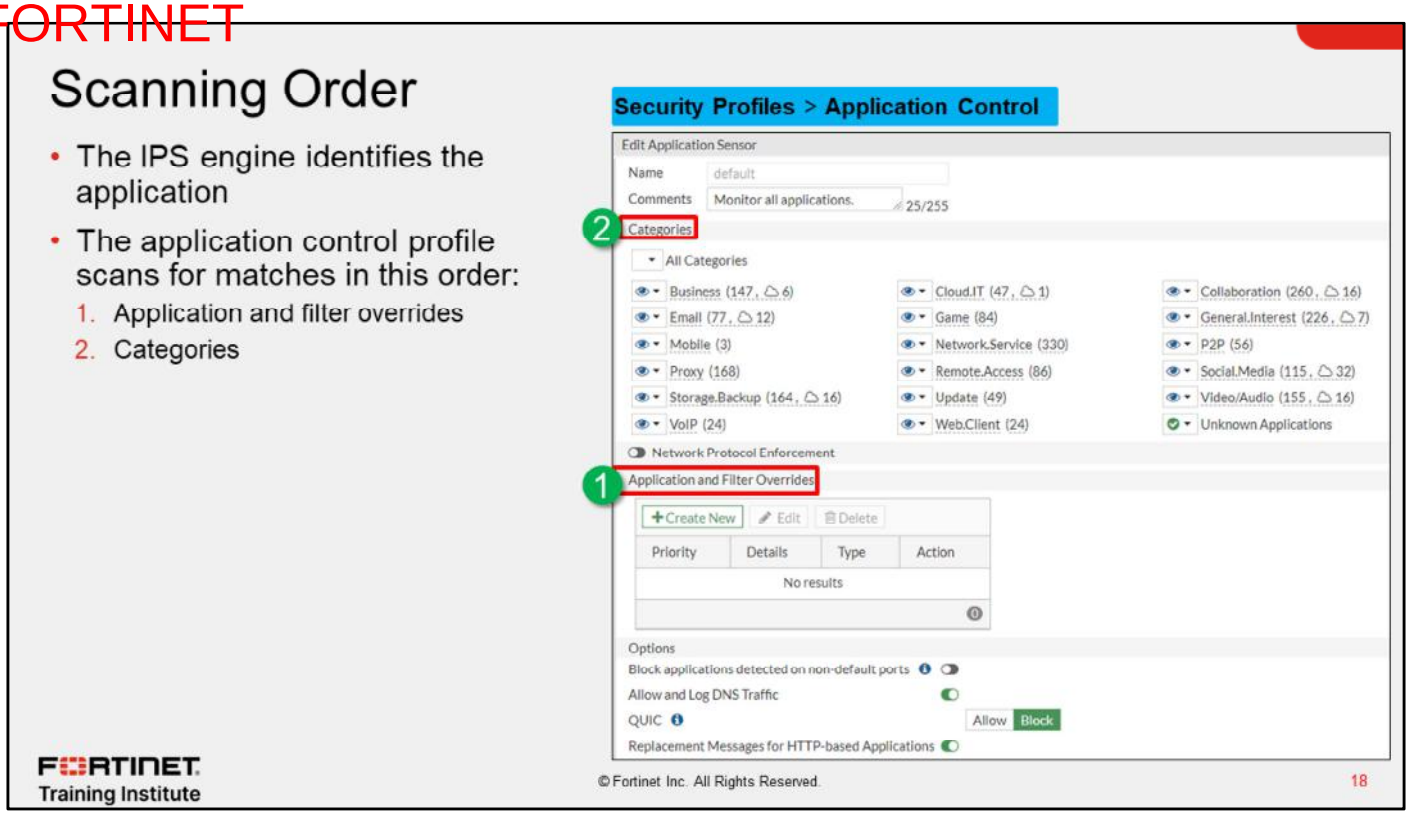

The IPS engine examines the traffic stream for a signature match.

Then, FortiGate scans packets for matches, in this order, for the application control profile:

1. Application and filter overrides: If you have configured any application overrides or filter overrides, the application control profile considers those first. It looks for a matching override starting at the top of the list, like firewall policies.
2. Categories: Finally, the application control profile applies the action that you've configured for applications in your selected categories.

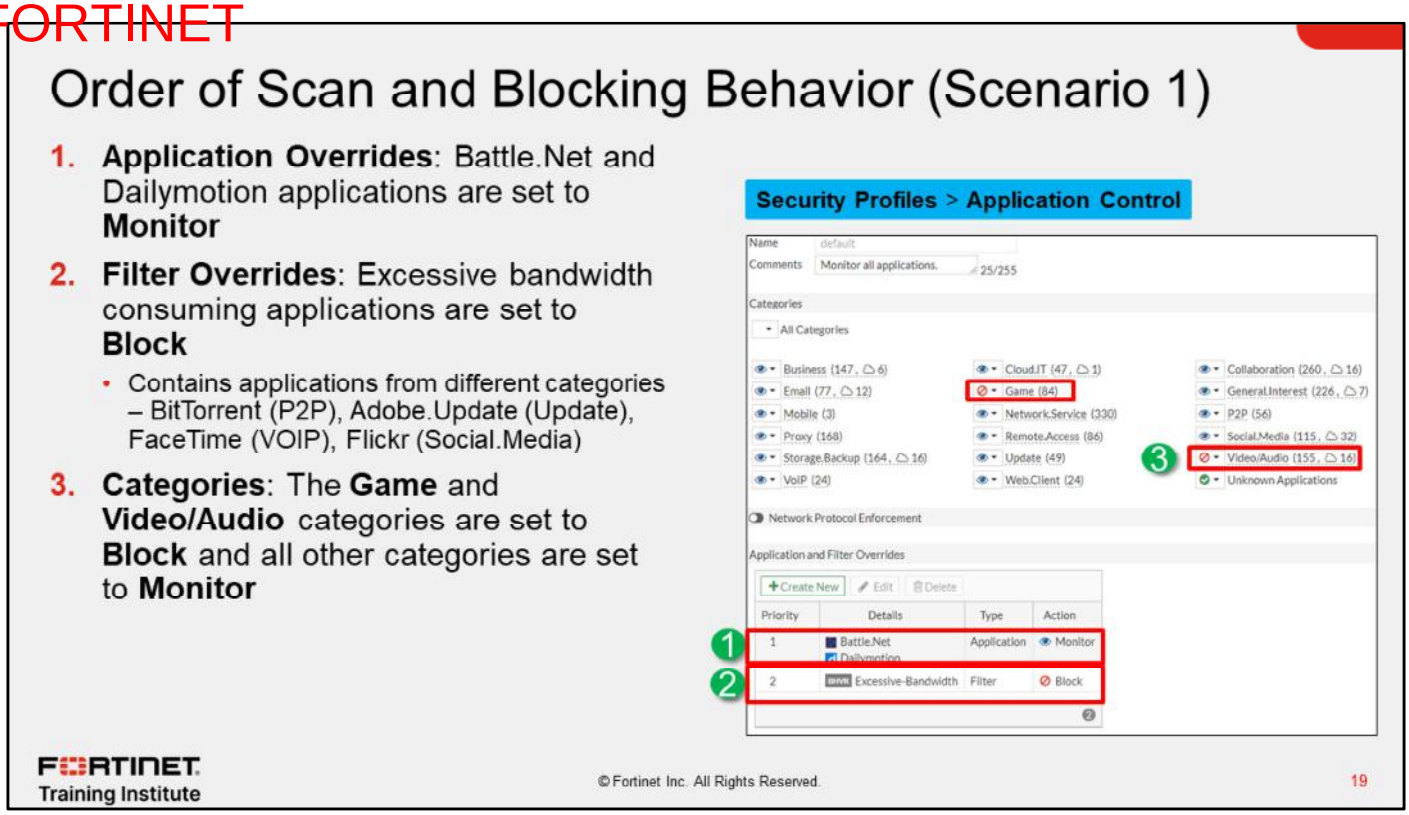

In the example profile shown on this slide, the application control profile blocks the **Game** and **Video/Audio** categories. For applications in these categories, FortiGate responds with the application control HTTP block message. (It is slightly different from the web filtering HTTP block message.) All other categories are set to **Monitor**, except **Unknown Applications**, and are allowed to pass traffic.

In the **Application and Filter Overrides** section, you can see that some exceptions are specified. Instead of being set to **Block**, **Battle.Net** (**Game**) and **Dailymotion** (**Video/Audio**) are set to **Monitor**. Because application overrides are applied first in the scan, these two applications are allowed, and generate logs.

Next, the scan checks for **Application and Filter Overrides**. Because a filter override is configured to block applications that use excessive bandwidth, it blocks all applications using excessive bandwidth, regardless of categories that allow these applications.

This slide shows an example of how several security profile features could work together, overlap, or work as substitutes, on the same traffic.

After the application control profile scan is done, FortiGate begins other scans, such as web filtering. The web filtering scan could block Battle.Net and Dailymotion, but it would use its own block message. Also, web filtering doesn't check the list of application control overrides. *So, even if an application control override allows an application, web filtering could still block it.*

Similarly, static URL filtering has its own exempt action, which bypasses all subsequent security checks. However, application control occurs before web filtering, so that the web filtering exemption *cannot* bypass application control.

Order of Scan and Blocking Behavior (Scenario 2)

1. **Filter Overrides**: Excessive bandwidth consuming applications are set to **Block**
   - Contains applications from different categories – BitTorrent (P2P), Adobe.Update (Update), FaceTime (VOIP), Flickr (Social.Media)
2. **Application Overrides**: Battle.Net and Dailymotion applications are set to **Monitor**
3. **Categories**: The **Game** and **Video/Audio** categories are set to **Block** and all other categories set to **Monitor**

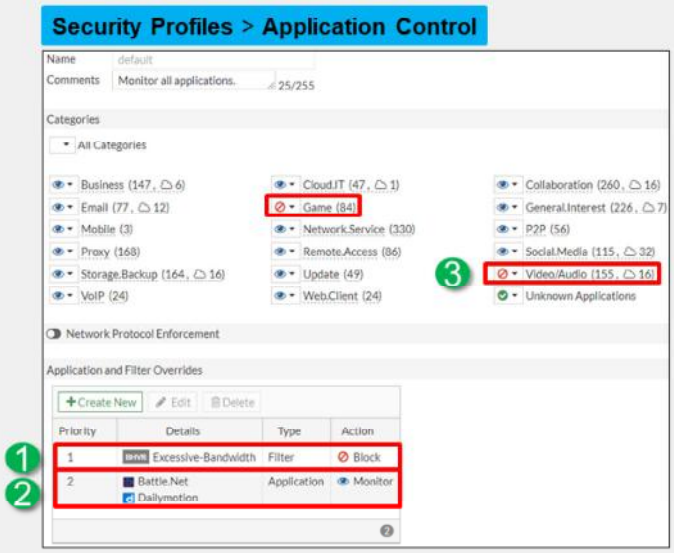

FORTINET Training Institute

20

In the example profile shown on this slide, the filter override has been moved above the application override. In this scenario, the filter override (**Excessive-Bandwidth**) is blocked and, since **Dailymotion** falls under the excessive bandwidth category, Dailymotion is blocked even though it is set to **Monitor** under the **Application and Filter Overrides** section.

The priority in which application and filter overrides are placed takes precedence.

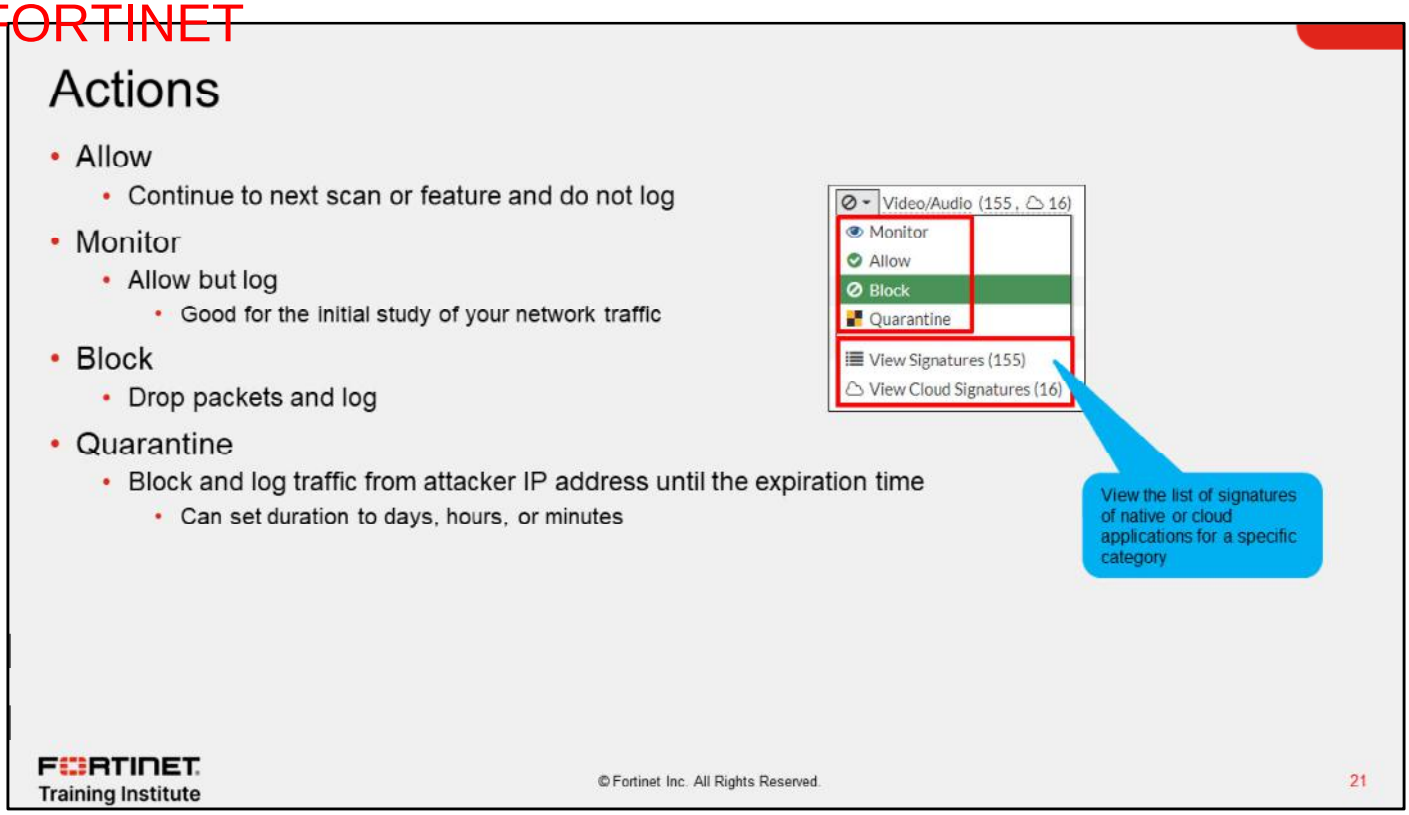

For each filter in the application control profile, you must indicate an action—what FortiGate does when traffic matches. Actions include the following:

- **Allow**: Passes the traffic and does not generate a log
- **Monitor**: Passes the traffic, but also generates a log message
- **Block**: Drops the detected traffic and generates a log message
- **Quarantine**: Blocks the traffic from an attacker IP until the expiration time is reached and generates a log message

The **View Signature** action allows you to view signatures from a particular category only and is *not* a configurable action. The **View Cloud Signatures** action allows you to view application signatures for cloud applications from a particular category.

Which is the correct action to choose?

If you're not sure which action to choose, **Monitor** can be useful initially, while you study your network. Later, after you have studied your network traffic, you can fine-tune your filter selection by choosing the most appropriate action. The action you choose also depends on the application. If an application requires feedback to prevent instability or other unwanted behavior, then you might choose **Quarantine** instead of **Block**. Otherwise, the most efficient use of FortiGate resources is to block.

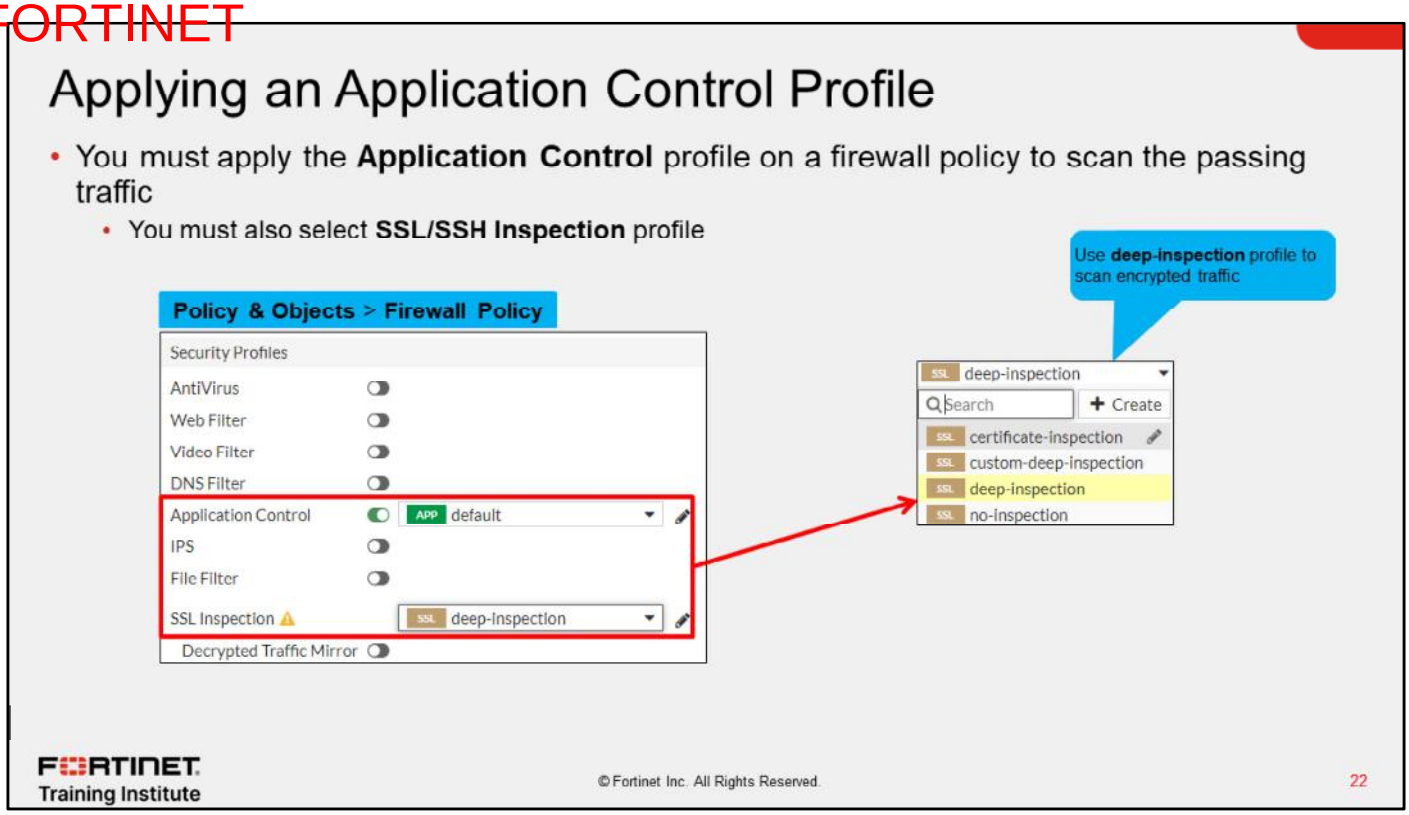

After you configure an application control profile, you must apply it to a firewall policy. This instructs FortiGate to start scanning application traffic that is subject to the firewall policy.

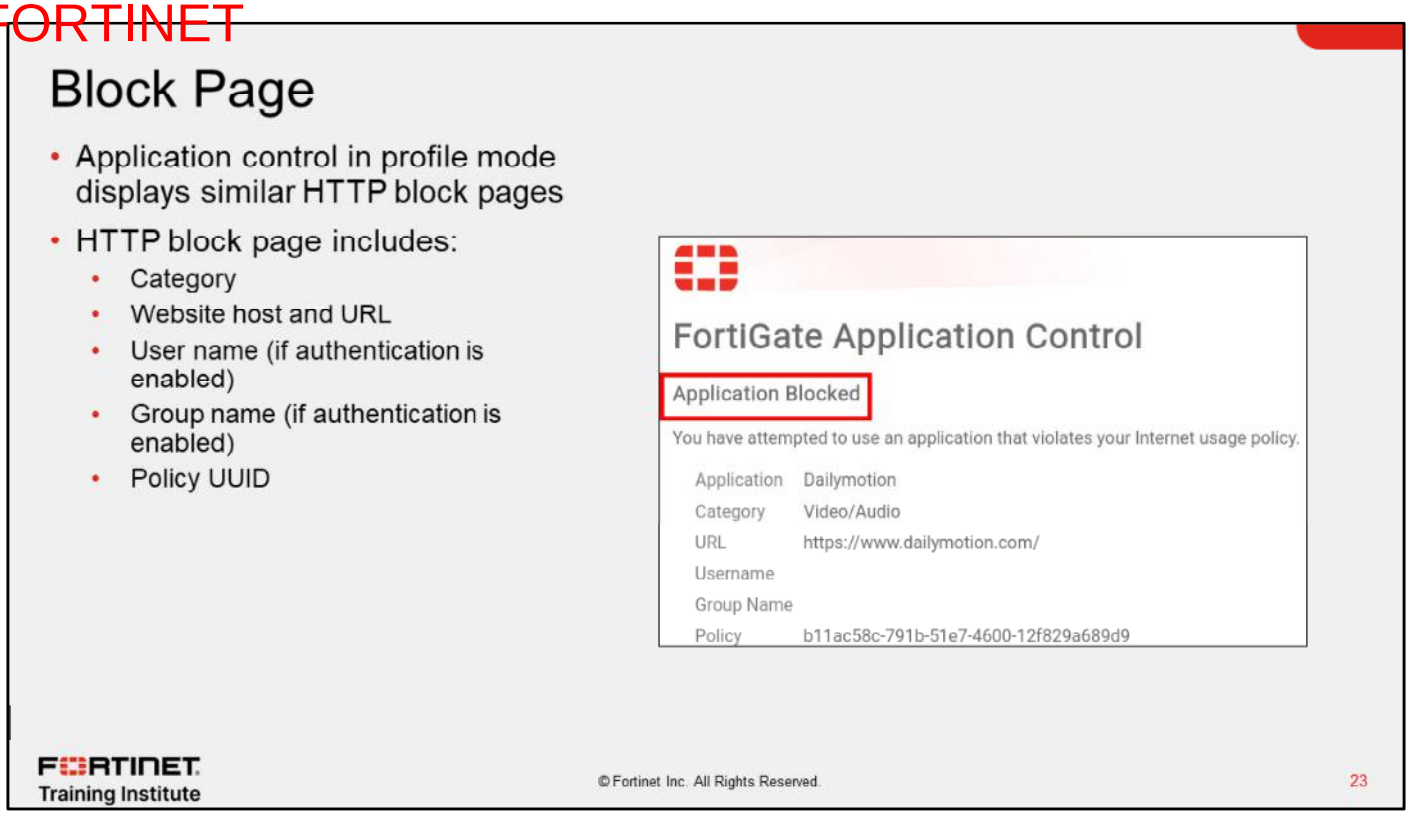

For HTTP-based applications, application control can provide feedback to the user about why their application was blocked. This is called a block page, and it is similar to the one you can configure for URLs that you block using FortiGuard web filtering.

It is also worth mentioning that, if deep inspection is enabled in the firewall policy, all HTTPS-based applications provide this block page.

The block page contains the following information:
- Signature that detected the application (in this case, Dailymotion)
- Signature's category (Video/Audio)
- URL that was specifically blocked (in this case, the index page of `www.dailymotion.com`), since a web page can be assembled from multiple URLs
- User name (if authentication is enabled)
- Group name (if authentication is enabled)
- UUID of the policy governing the traffic

The last item in this list can help you to identify which policy on FortiGate blocked the page, even if you have a large number of policies with many FortiGate devices securing different segments.

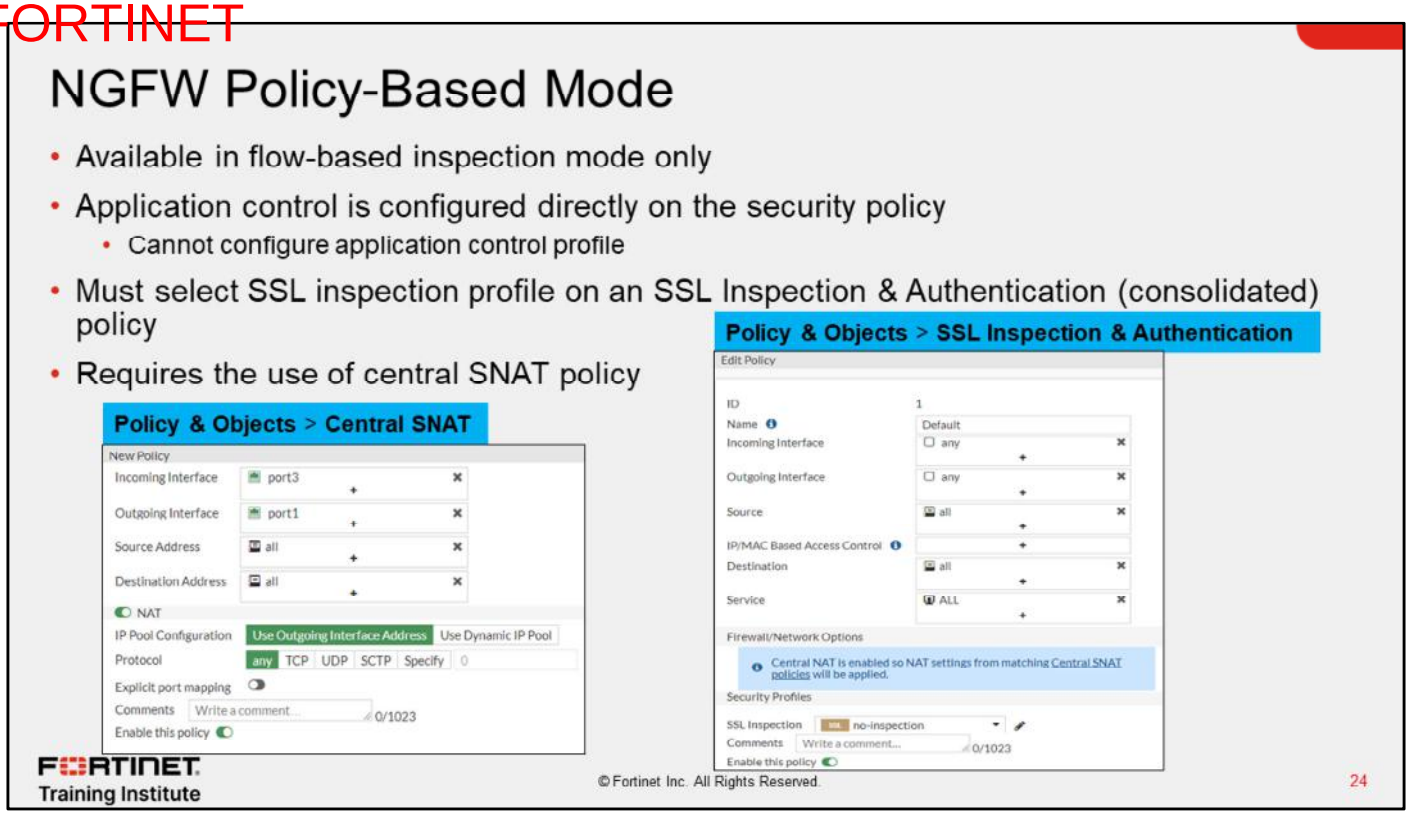

When FortiGate is operating in NGFW policy-based mode, administrators can apply application control to a security policy directly, instead of having to create an application control profile first, and then apply that to a firewall policy. Eliminating the need to use an application control profile makes it easier for the administrator to select the applications or application categories they want to allow or deny in the firewall policy.

It is important to note that all security policies in an NGFW policy-based mode VDOM or FortiGate must specify an SSL/SSH inspection profile on a consolidated policy. NGFW policy-based mode also requires the use of central source NAT (SNAT), instead of NAT settings applied within the firewall policy.

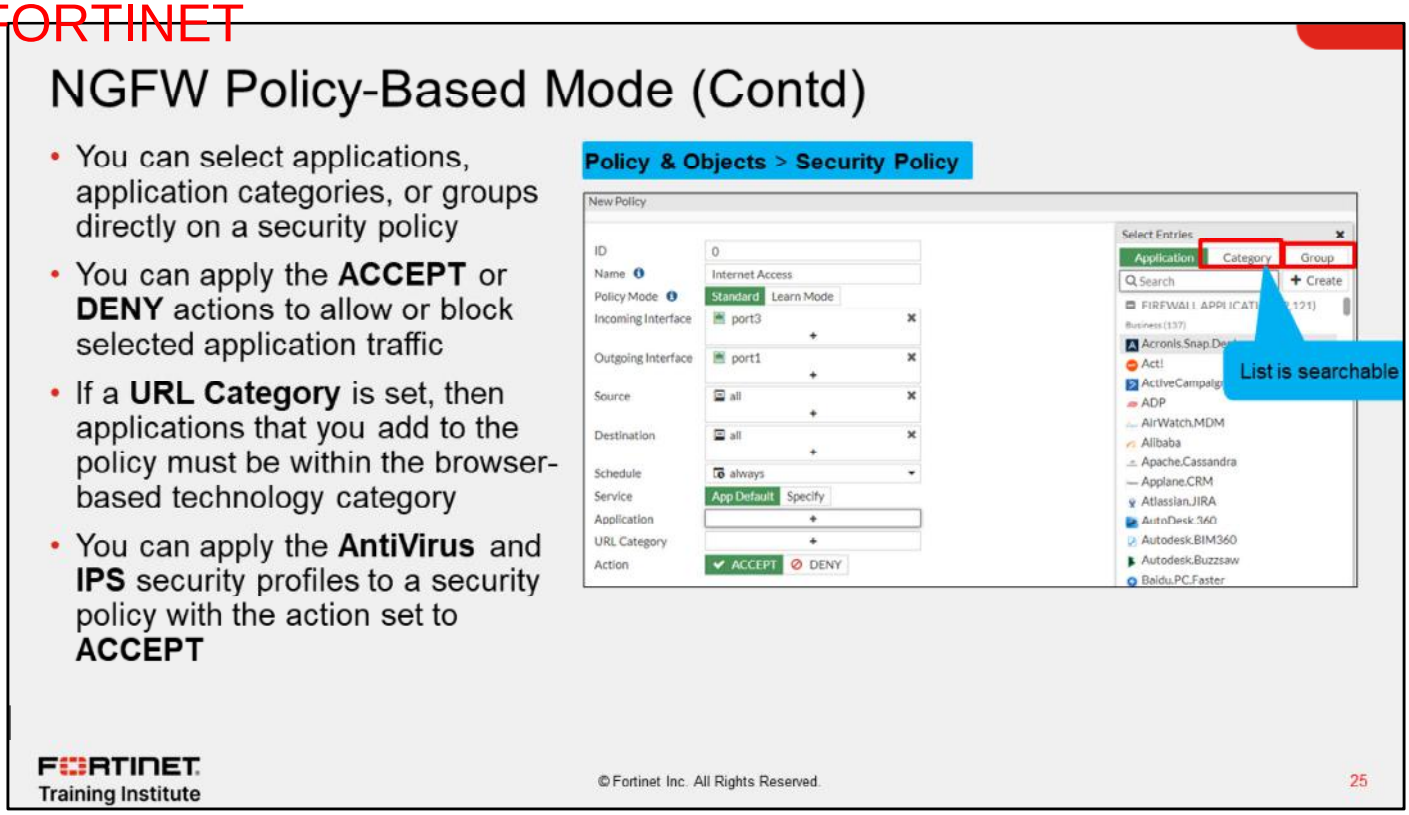

You can select one or more applications, application groups, and application categories on a security policy in the **Application** section. After you click the **+** icon for an application**,** a pop-up window opens. In that window, you can search for and select one or more application signatures, application groups, or application categories. Based on the applications, groups, and application categories applied to the policy, FortiOS applies the security action to the application traffic.

You can configure the **URL Category** within the same security policy; however, adding a URL filter causes application control to scan applications in only the browser-based technology category, for example, Facebook Messenger on the Facebook website.

You can also configure the **Group** with multiple applications and application categories. This allows the administrator to mix multiple applications and categories.

In addition to applying a URL category filter, you can also apply **AntiVirus** and **IPS** security profiles to application traffic that is allowed to pass through.

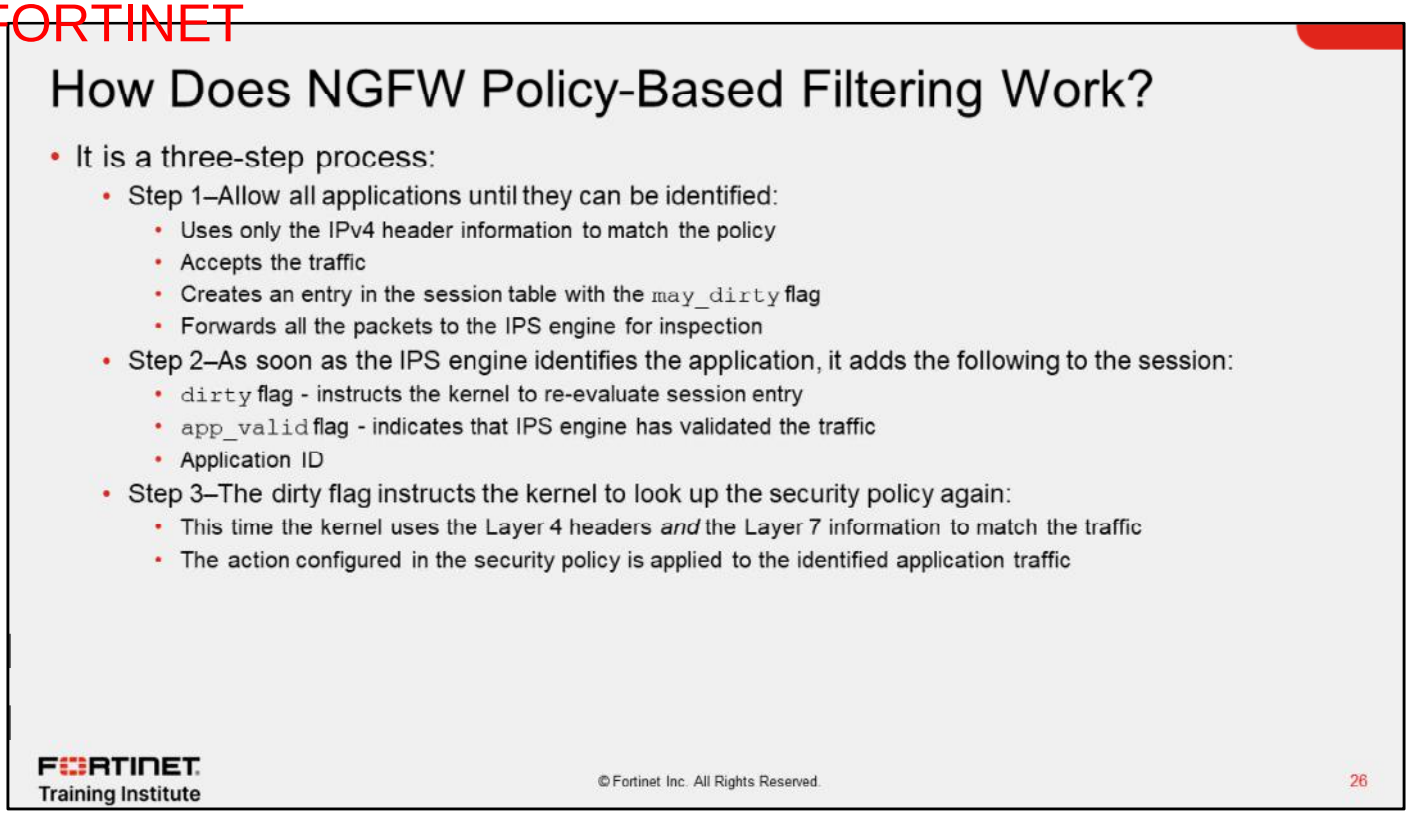

FortiOS uses a three-step process to perform NGFW policy-based application filtering. Here is a brief overview of what happens at each step.

In step 1, FortiOS allows all traffic while forwarding packets to the IPS engine for inspection and identification of the traffic. At the same time, FortiOS creates an entry in the session table allowing the traffic to pass and it adds a `may_dirty` flag to it.

In step 2, as soon as the IPS engine identifies the application, it updates the session entry with the following information: `dirty` flag, `app_valid` flag, and an application ID.

In step 3, the FortiOS kernel performs a security policy lookup again, to see if the identified application ID is listed in any of the existing security policies. This time the kernel uses both Layer 4 and Layer 7 information for policy matching. After the criteria matches a firewall policy rule, the FortiOS kernel applies the action configured on the security policy to the application traffic.

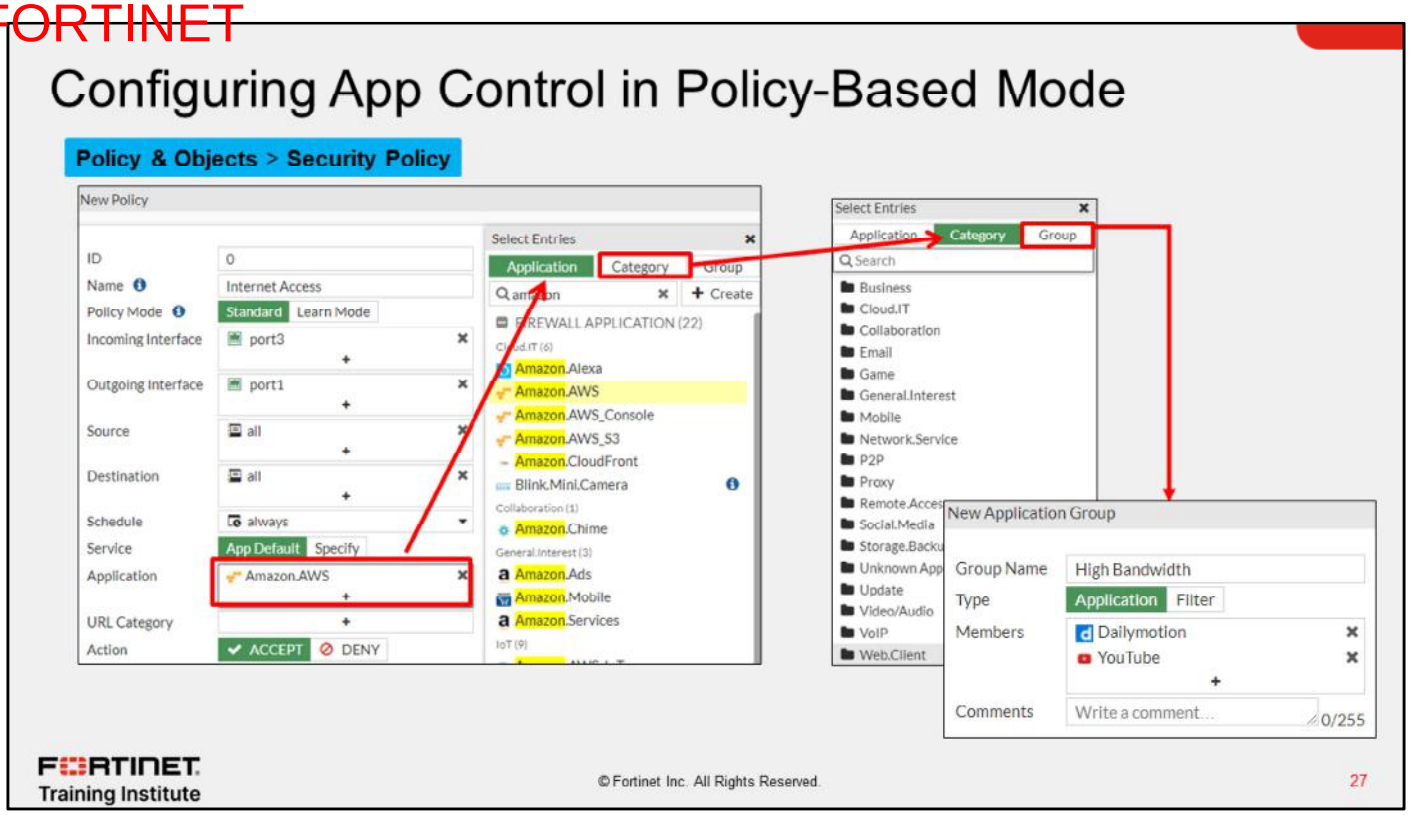

Configuring application control in NGFW policy-based mode is simple. You can create a new security policy or edit an existing security policy. In the **Application** section, select the applications, categories, or groups that you want to allow or deny, and change the security policy **Action** accordingly. For applications that you selected to allow, you can further enhance network security by enabling antivirus scanning and IPS control. You can also enable the logging of **Security Events** or **All Sessions** to ensure that all application control events are logged.

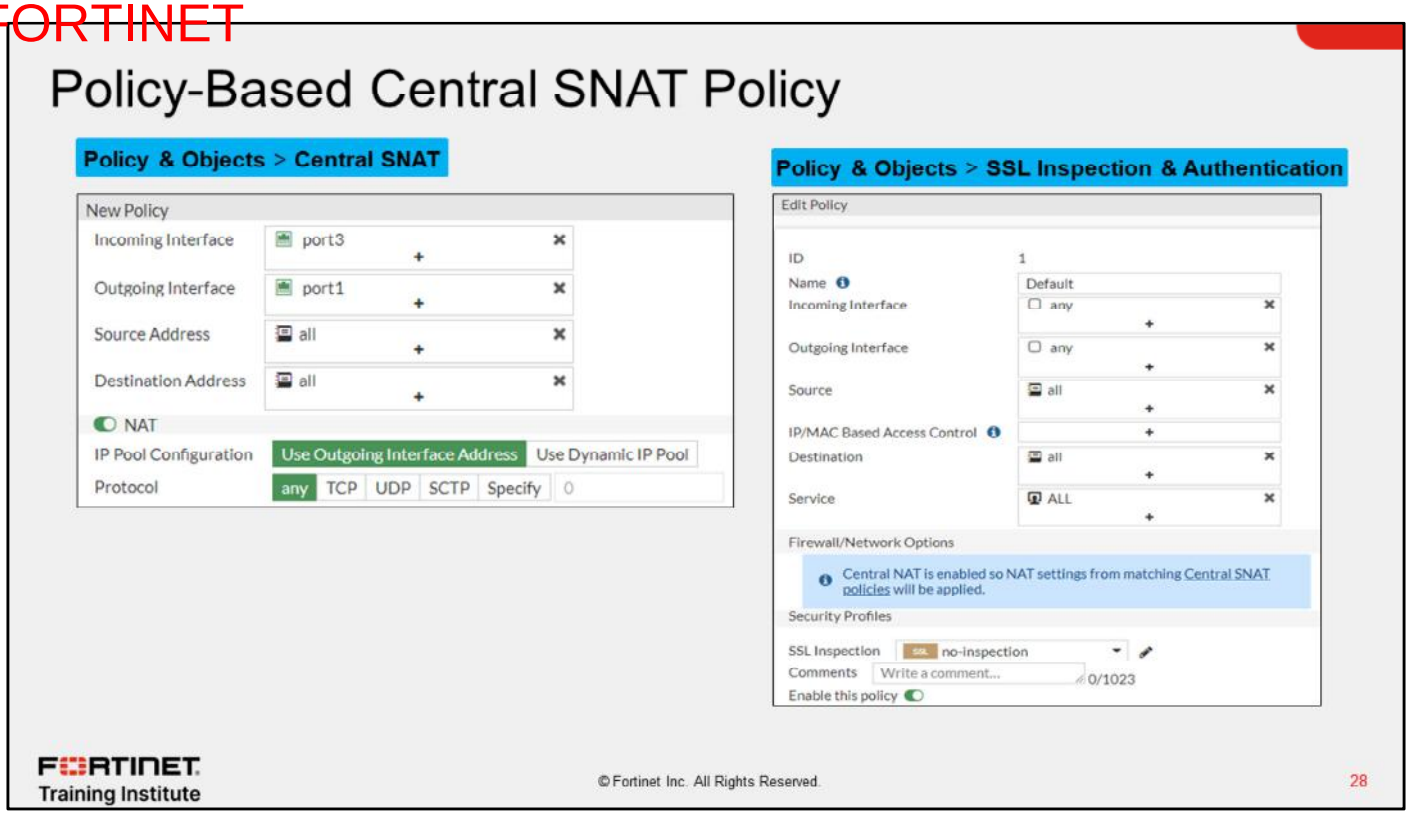

You must have a matching central SNAT policy in NGFW policy-based mode to be able to pass traffic. FortiGate applies NAT on the traffic based on the criteria defined in the central SNAT policy.

It is extremely important to arrange security policies in **Policy & Objects**, so that the more specific policies are located at the top to ensure proper use of application control.

A default **SSL Inspection & Authentication** policy inspects traffic accepted by any of the security firewalls, and by using the **certificate-inspection** SSL inspection profile.

NGFW Policy Matching

- Based on the configuration shown in the screenshot:
  - Facebook, Flickr, Instagram, and Pinterest application traffic is blocked by policy ID 1
  - All other Social.Media (for example, LinkedIn) application traffic is allowed by policy ID 2
  - All applications that belong to the P2P application category are blocked by policy ID 3
  - All other traffic and applications are allowed by policy ID 4

Policy & Objects > Security Policy

ID	Name	Source	Destination	Schedule	Service	Applications	Action	Security Profiles	Log
port3 → port1 4									
1	Blocking apps	all	all	always	App Default	Facebook Flickr Instagram Pinterest	DENY		All
2	Allow social media	all	all	always	App Default	Social.Media	ACCEPT	default	All
3	Blocking P2P Apps	all	all	always	App Default	P2P	DENY		Disabled
4	Allow all	all	all	always	App Default		ACCEPT	default	UTM

FORTINET Training Institute

29

NGFW policy matching works using a top-to-bottom approach. You must have a specific policy above a more broad or open policy. For example, if you would like to block Facebook but allow the **Social.Media** category, you must place the policy blocking Facebook traffic above the policy allowing the **Social.Media** category.

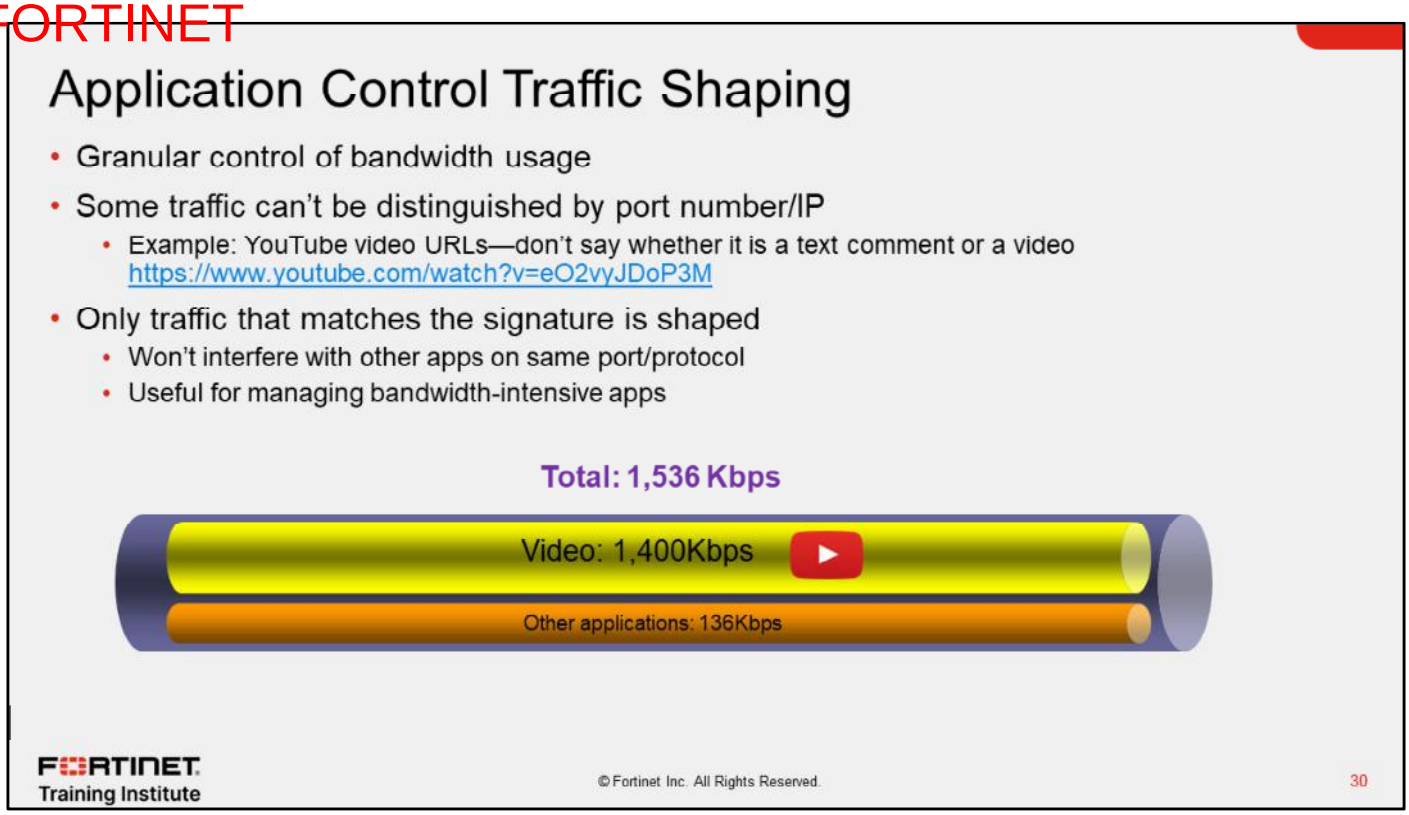

If an application is necessary, but you must prevent it from impacting bandwidth then, instead of blocking it entirely, you can apply a rate limit to the application. For example, you can rate limit applications used for storage or backup leaving enough bandwidth for more sensitive streaming applications, such as video conferencing.

Applying traffic shaping to applications is very useful when you're trying to limit traffic that uses the same TCP or UDP port numbers as mission-critical applications. Some high-traffic web sites, such as YouTube, can be throttled in this way.

Examine the details of how throttling works. Not all URL requests to `www.youtube.com` are for video. Your browser makes several HTTPS requests for:

- The web page itself
- Images
- Scripts and style sheets
- Video

All of these items have separate URLs. If you analyze a site like YouTube, the web pages themselves don't use much bandwidth; it is the video content that uses the most bandwidth. But, since all content is transported using the same protocol (HTTPS), and the URLs contain dynamically generated alphanumeric strings, traditional firewall policies can't block or throttle the traffic by port number or protocol because they are the same. Using application control, you can rate limit only videos. Doing this prevents users from saturating your network bandwidth, while still allowing them to access the other content on the site, such as for comments or sharing links.

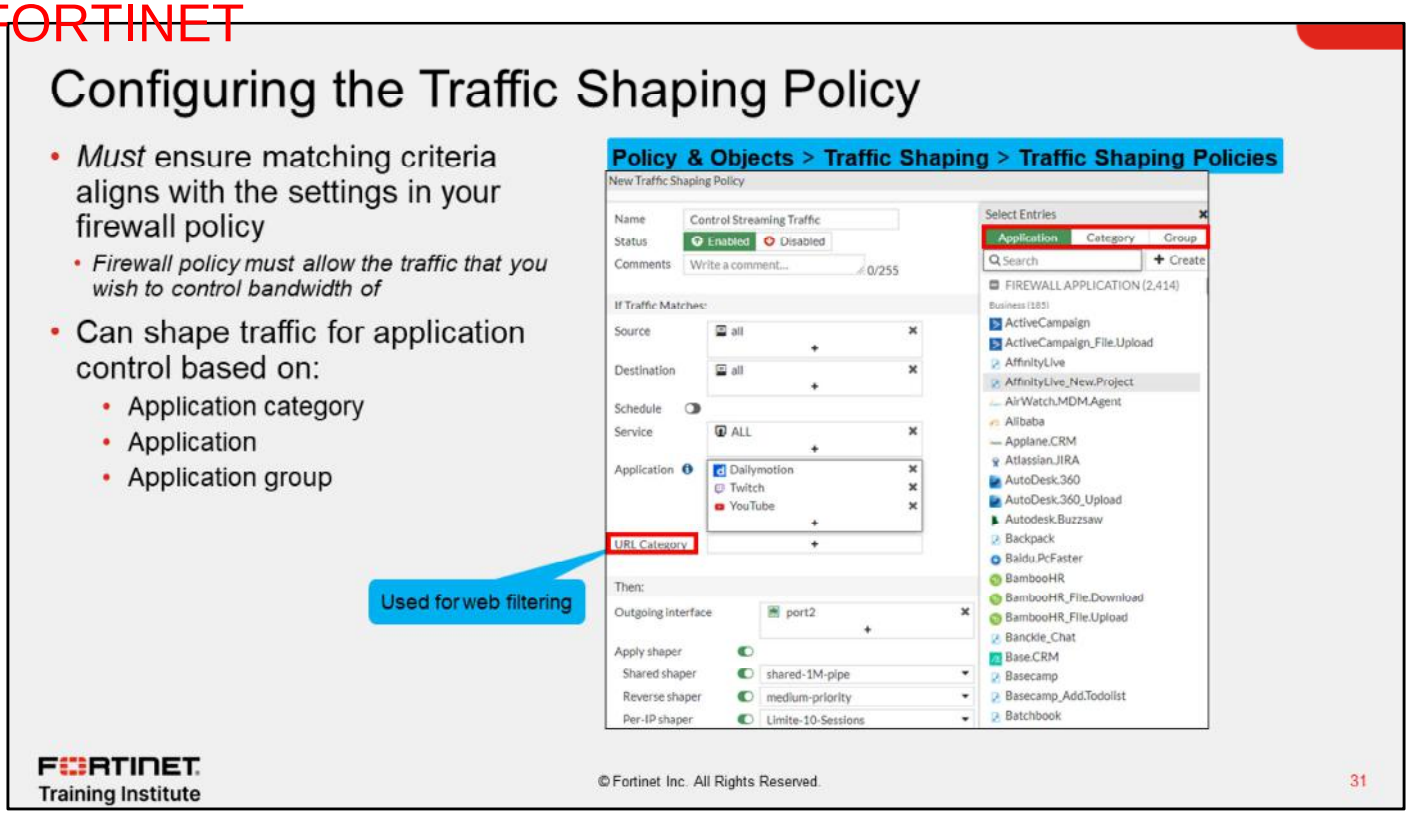

You can limit the bandwidth of an application category, application group, or specific application by configuring a traffic shaping policy. You can also apply traffic shaping to FortiGuard web filter categories and to the application group.

You must ensure that the matching criteria aligns with the firewall policy or policies to which you want to apply shaping. It does not have to match outright. For example, if the source in the firewall policy is set to **all** (`0.0.0.0/0.0.0.0`), you can set the source in the traffic shaping policy to any source that is included in **all**, for example, **LOCAL_SUBNET** (`10.0.1.0/24`).

If the traffic shaping policy is not visible in the GUI, you can enable it on the **Feature Visibility** page.

There are two types of shapers that you can configure on the **Traffic Shaping Policy** page, and you can apply them in the traffic shaping policy:

- **Shared shaper**: applies a total bandwidth to all traffic using that shaper. The scope can be per policy or for all policies referencing that shaper.
- **Per-IP shaper**: applies traffic shaping to all source IP addresses in the security policy. Bandwidth is equally divided among the group.

Note that the outgoing interface is usually the egress interface (WAN). The **Shared shaper** setting is applied to ingress-to-egress traffic, which is useful for restricting bandwidth for uploading. The **Reverse Shaper** setting is also a shared shaper, but it is applied to traffic in the reverse direction (egress-to-ingress traffic). This is useful for restricting bandwidth for downloading or streaming, because it limits the bandwidth from the external interface to the internal interface.

Knowledge Check

1. Which statement about application control in an NGFW policy-based configuration is true?

✓ A. Applications are applied directly to the security policies.
B. The application control profile must be applied to firewall policies.

2. Which statement about the HTTP block page for application control is true?

✓ A. It can be used only for web applications.
B. It works for all types of applications.

FORTINET Training Institute

32

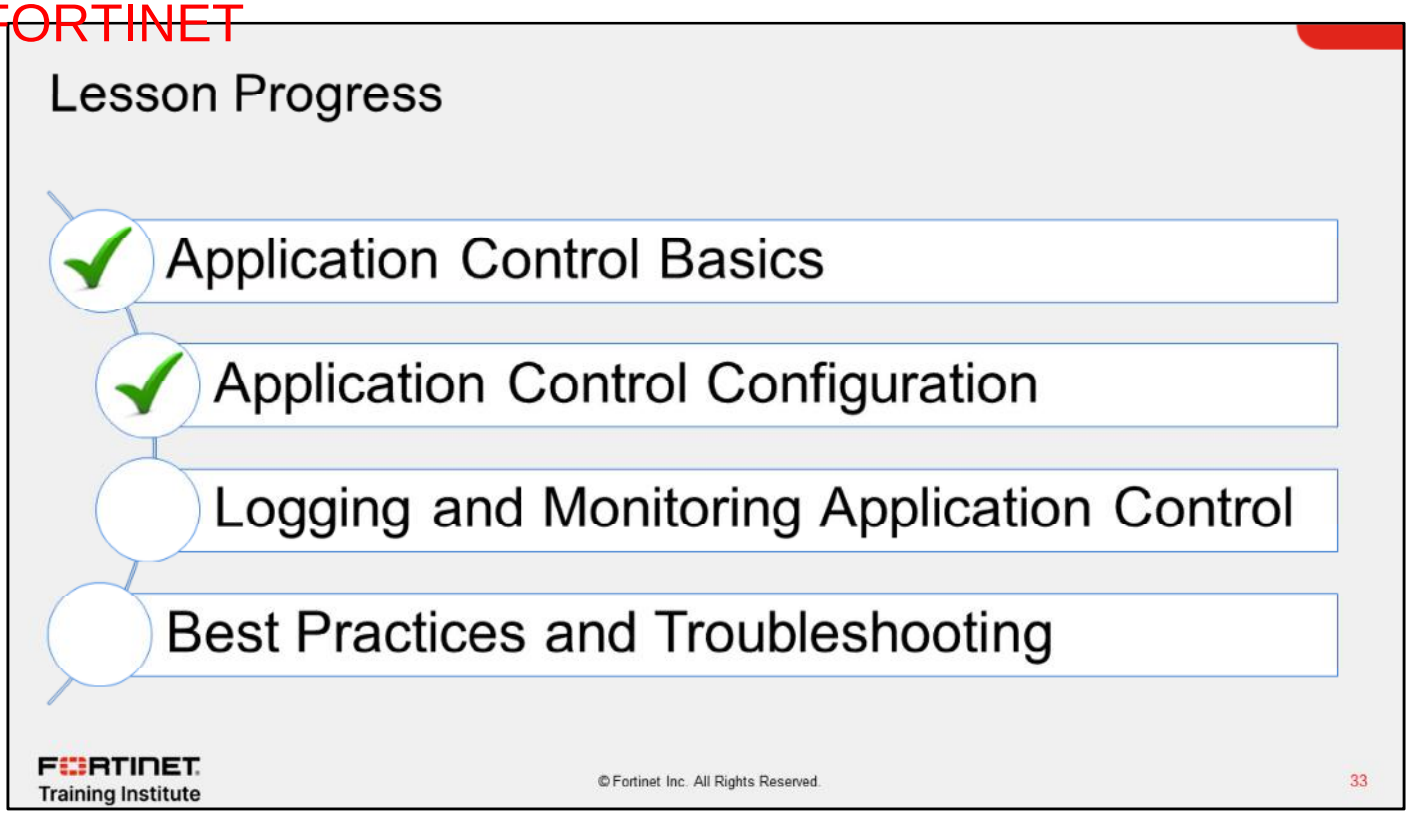

Good job! You now understand application control configuration.

Now, you will learn about logging and monitoring application control events.

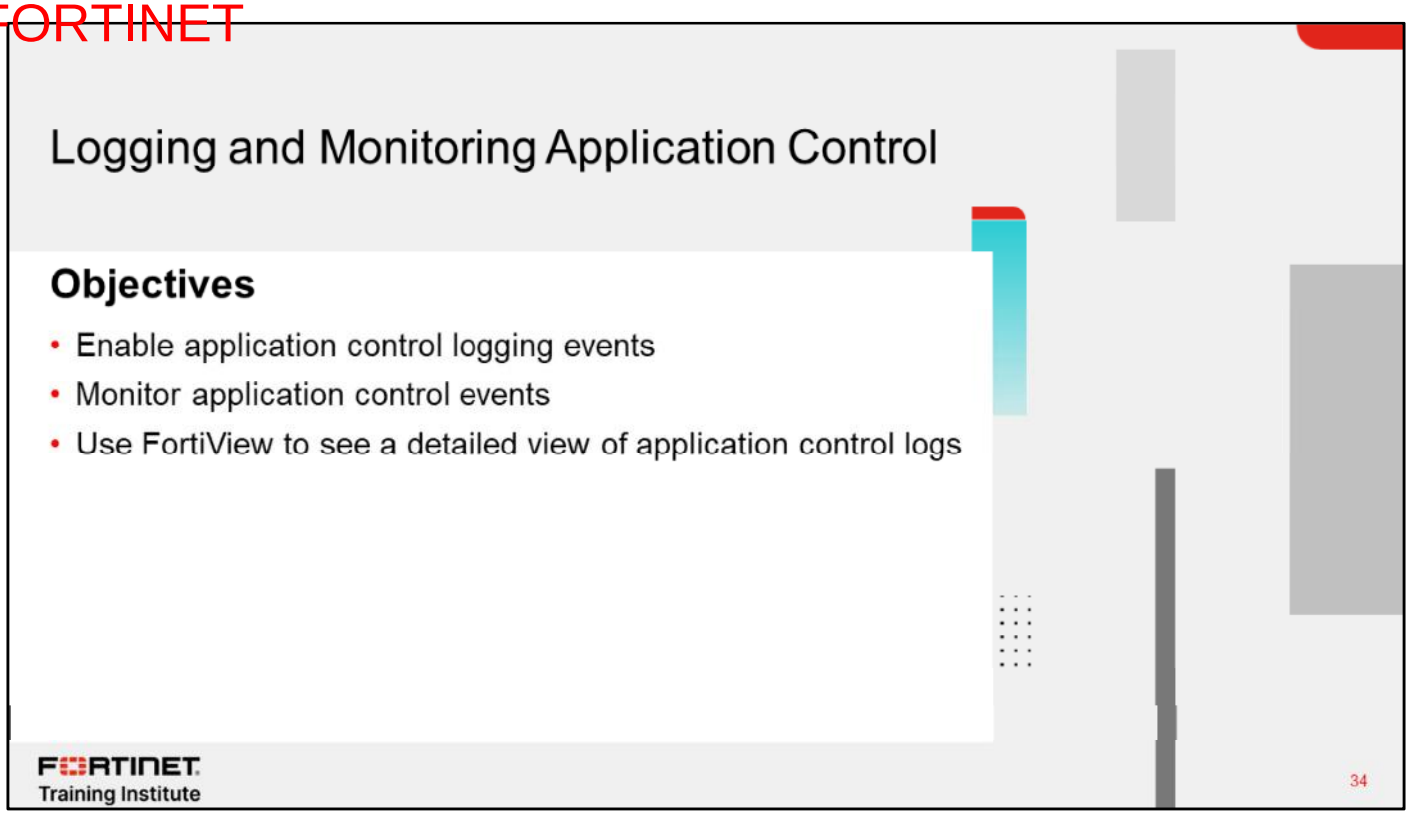

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control configuration, including reviewing application control logs, you will be able to effectively use and monitor application control events.

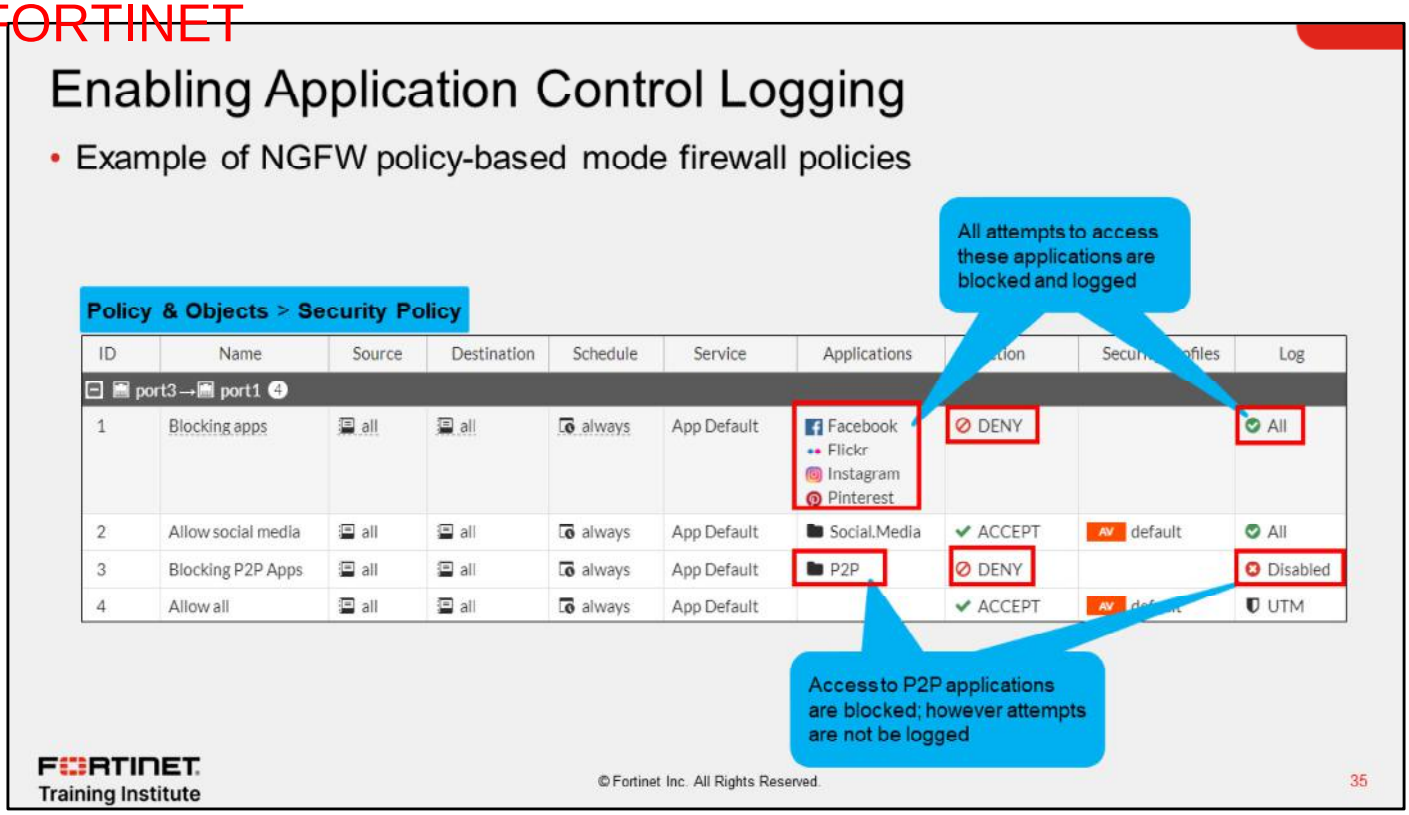

Regardless of which operation mode application control is configured in, you must enable logging on the security or firewall policy. When you enable the logging of security events or all sessions on a security or firewall policy, application control events are also logged. You must apply application control to the security or firewall policy to enable application control event logging.

When the **Deny** action is selected on a security or firewall policy, you must enable the **Log Violations** option to generate application control events for blocked traffic.

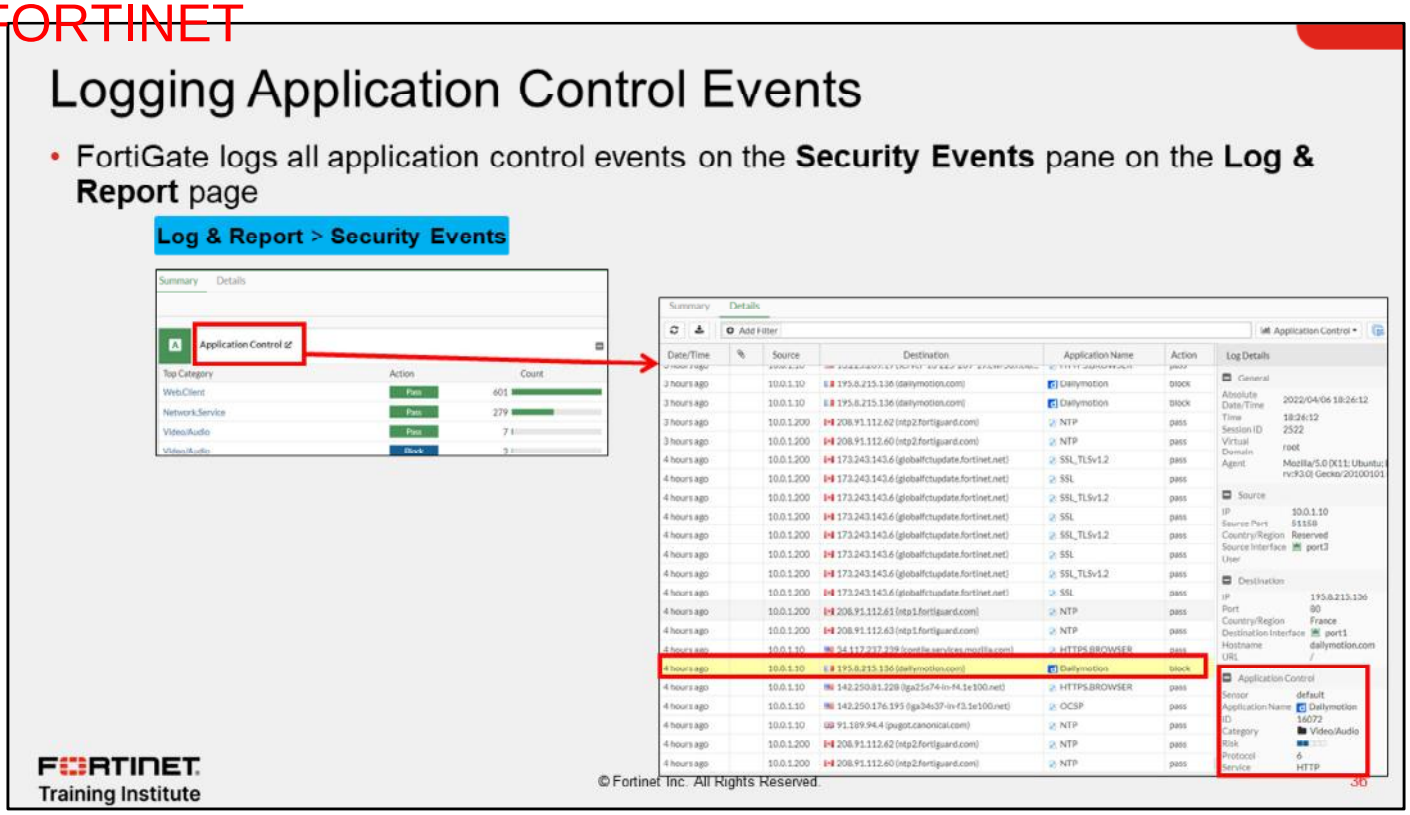

FortiGate logs all application control events on the **Security Events** pane on the **Log & Report** page. You can view the logs by clicking on **Application Control**.

In the example shown on this slide, the default application control profile blocks access to **Dailymotion**. You can view this information in the **Log Details** section, as well as information about the log source, destination, application, and action.

Note that application control generates this log message using a profile-based configuration. The log message for an NGFW policy-based configuration, does not include information that does not apply, such as application sensor name. The remainder of the information and structure of the log message is the same for each log, regardless of which inspection mode FortiGate is using.

You can also view the details on the **Forward Traffic** logs pane, where firewall policies record activity. You can also find a summary of the traffic to which FortiGate applied application control. Again, this is because application control is applied by a firewall policy. To find out which policy applied application control, you can review either the **Policy ID** or the **Policy UUID** fields of the log message.

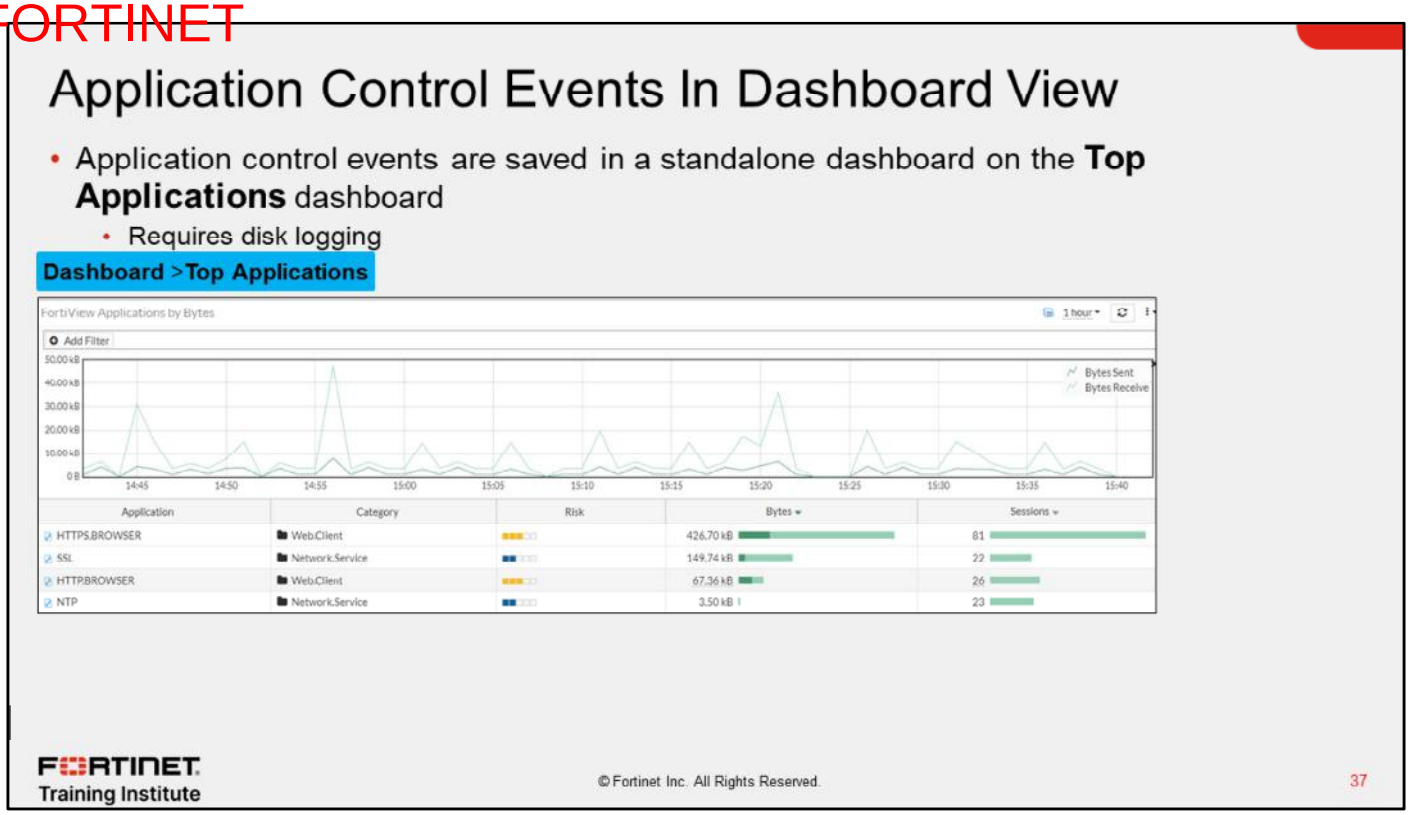

On the **Dashboard** menu, the **Top Applications** standalone page provides details about each application, such as the application name, category, and bandwidth. You can drill down further to see more granular details by double-clicking an individual log entry. The detailed view provides information about the source, destination, policies, or sessions for the selected application.

Knowledge Check

1. Where do you enable logging of application control events?
   - ✓ A. Application control logs are enabled in the firewall policy configuration.
   - B. Application control logs are enabled on the **FortiView Applications** page of FortiGate.

2. Which piece of information is not included in the application event log when using NGFW policy-based mode?
   - ✓ A. Application control profile name
   - B. Application name

FORTINET Training Institute

38

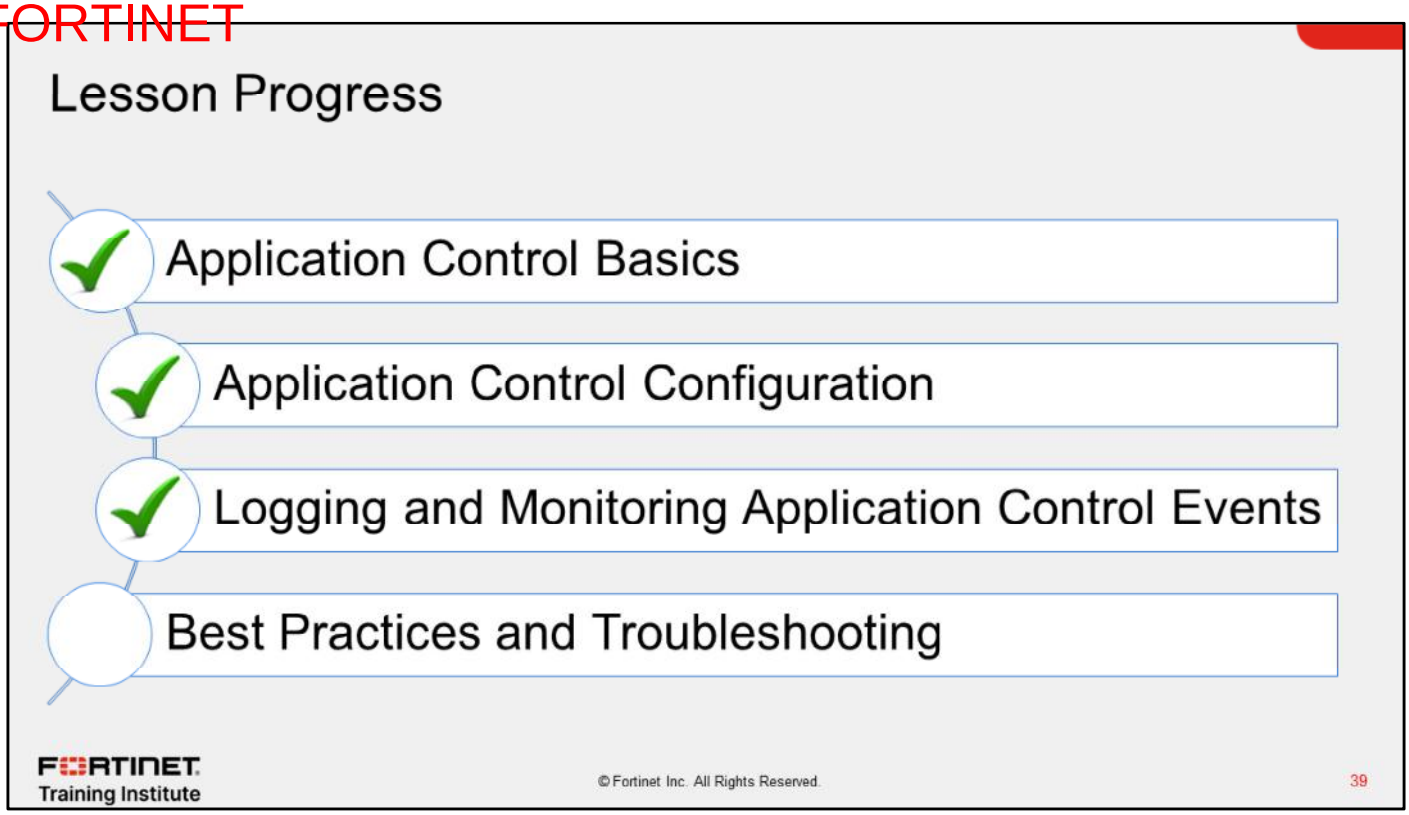

Good job! You now understand application control logging and monitoring.

Now, you will learn about application control best practices and troubleshooting.

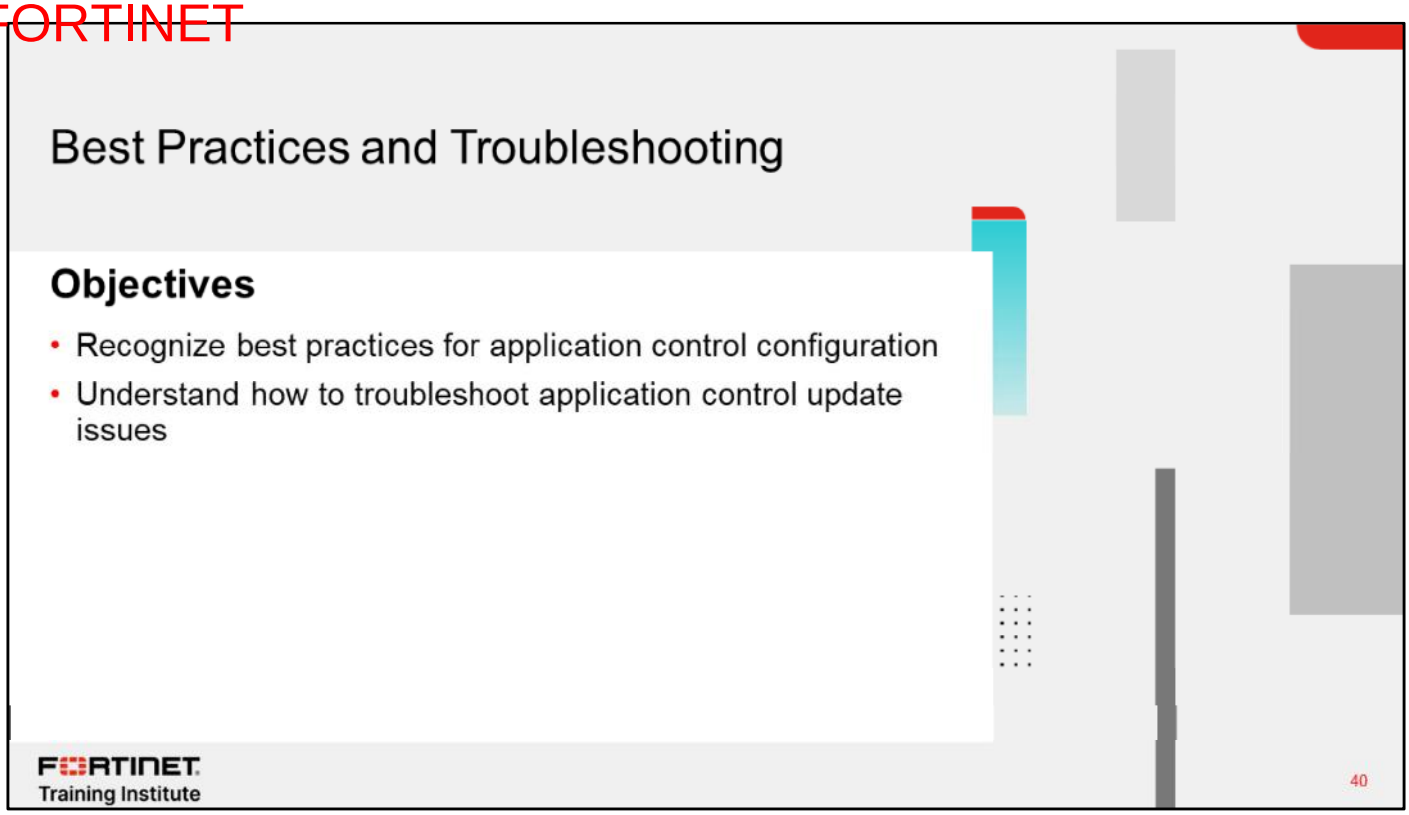

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control best practices and troubleshooting, you will be able to configure and maintain an effective application control solution.

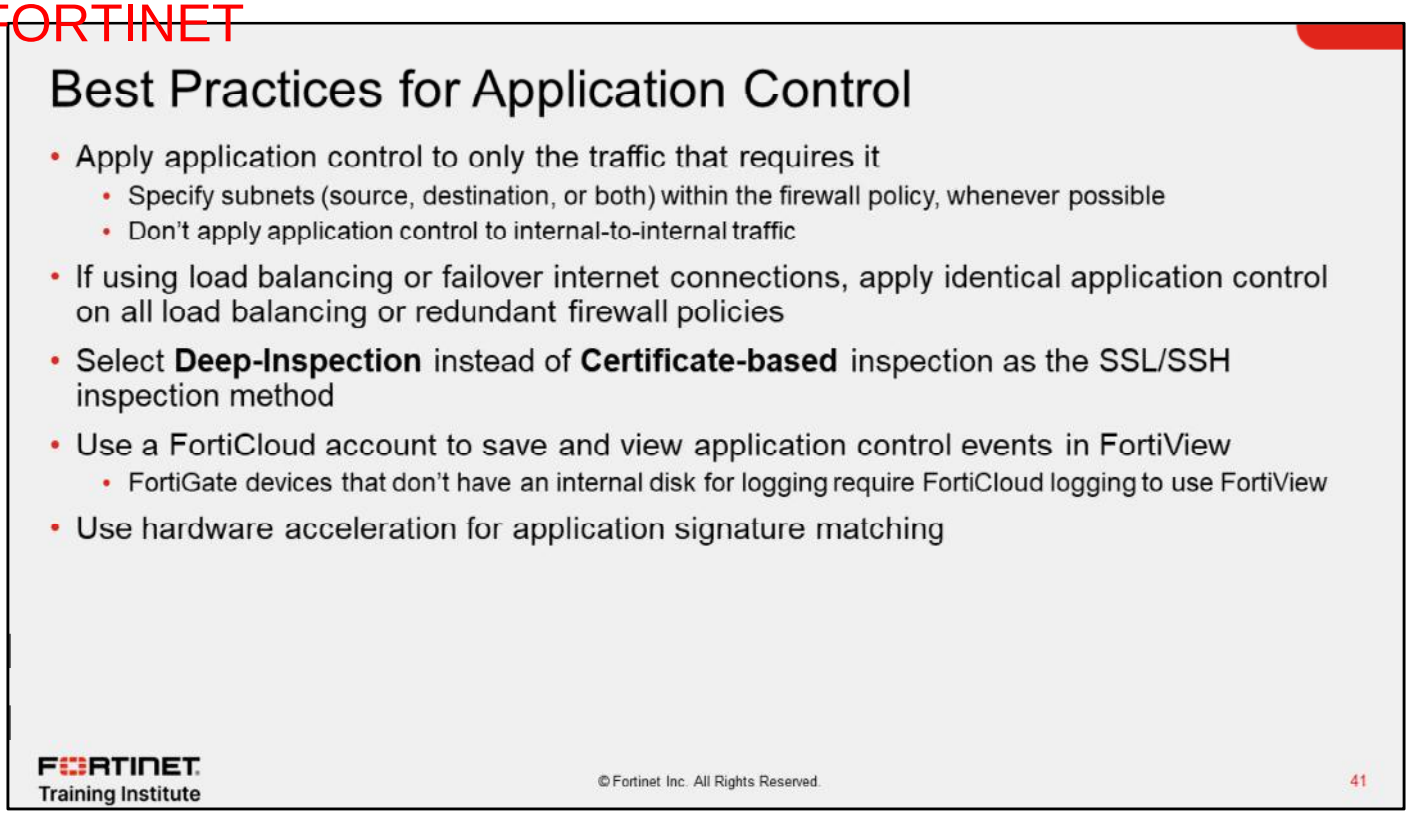

This slide lists some best practices to keep in mind when implementing application control on FortiGate.

Not all traffic requires an application control scan. Don't apply application control to internal-only traffic.

To minimize resource use on FortiGate, be as specific as possible when creating firewall policies. This reduces resource use, and also helps you build a more secure firewall configuration.

Create identical firewall policies for all redundant internet connections, to ensure that the same inspection is performed on failover traffic. Select **Deep-Inspection** instead of **Certificate-based** inspection for the SSL/SSH inspection mode, to ensure content inspection is performed on encryption protocols.

FortiGate models that feature specialized chips, such as network processors and content processors, can offload and accelerate application signature matching for enhanced performance.

You can use a FortiCloud account to save and view application control logs in FortiView, on FortiGate devices that do not have a log disk.

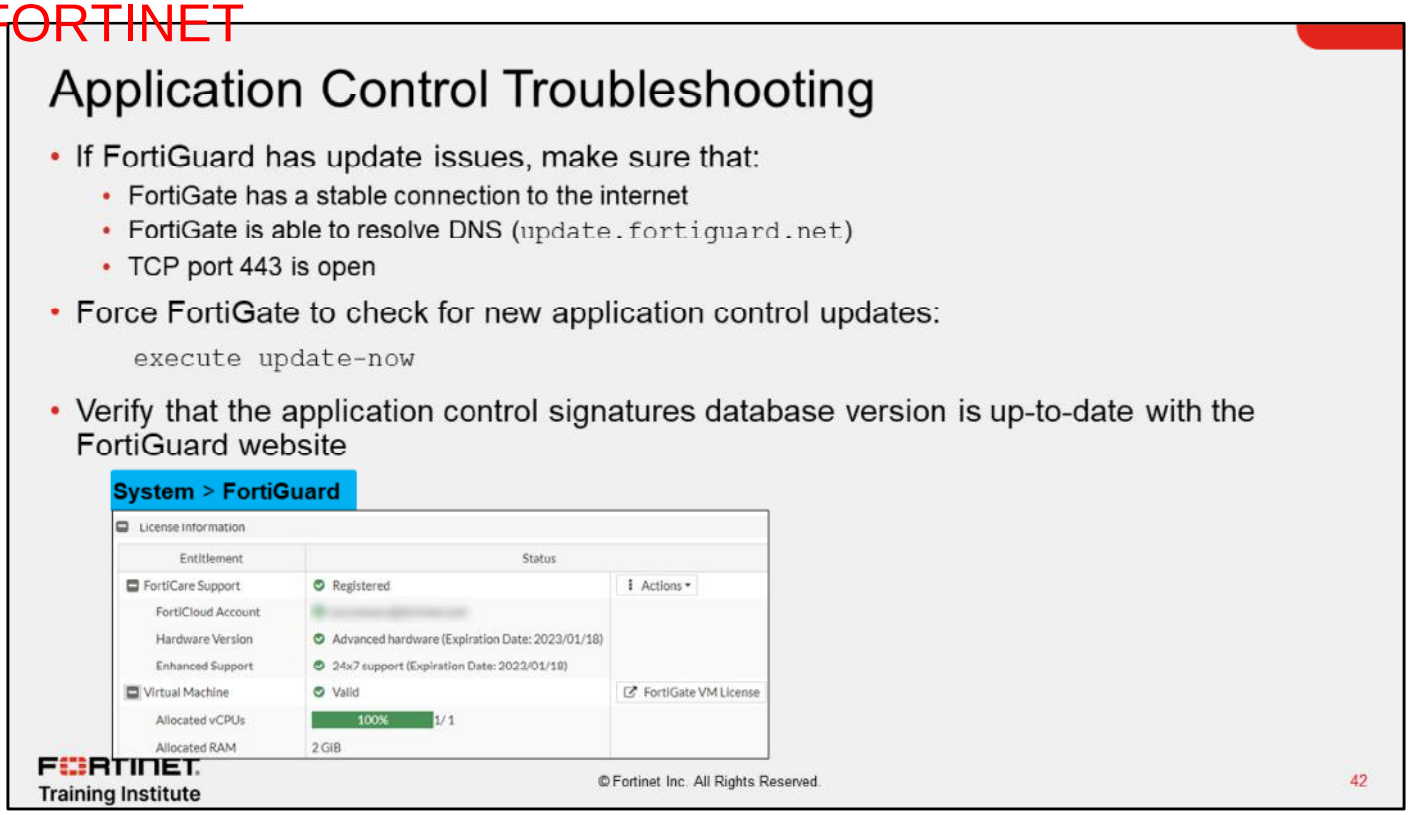

If you are experiencing issues with a FortiGuard application control update, start troubleshooting the issue with the most basic steps:

- Make sure that FortiGate has a stable connection to the internet or FortiManager (if FortiGate is configured to receive updates from FortiManager)
- If the internet connection is stable, check DNS resolution on FortiGate
- If FortiGate is installed behind a network firewall, make sure that port443 is being allowed from FortiGate

You can check the FortiGuard website for the latest version of the application control database. If your locally installed database is out-of-date, try forcing FortiGate to check for the latest updates by running the `execute update-now` command.

Knowledge Check

1. Which protocol does FortiGate use with FortiGuard to receive updates for application control?
   - A. UDP
   - ✓ B. TCP

2. Which SSL/SSH inspection method is recommended for use with application control scanning to improve application detection?
   - A. Certificate-based inspection profile
   - ✓ B. Deep-inspection profile

FORTINET Training Institute

43

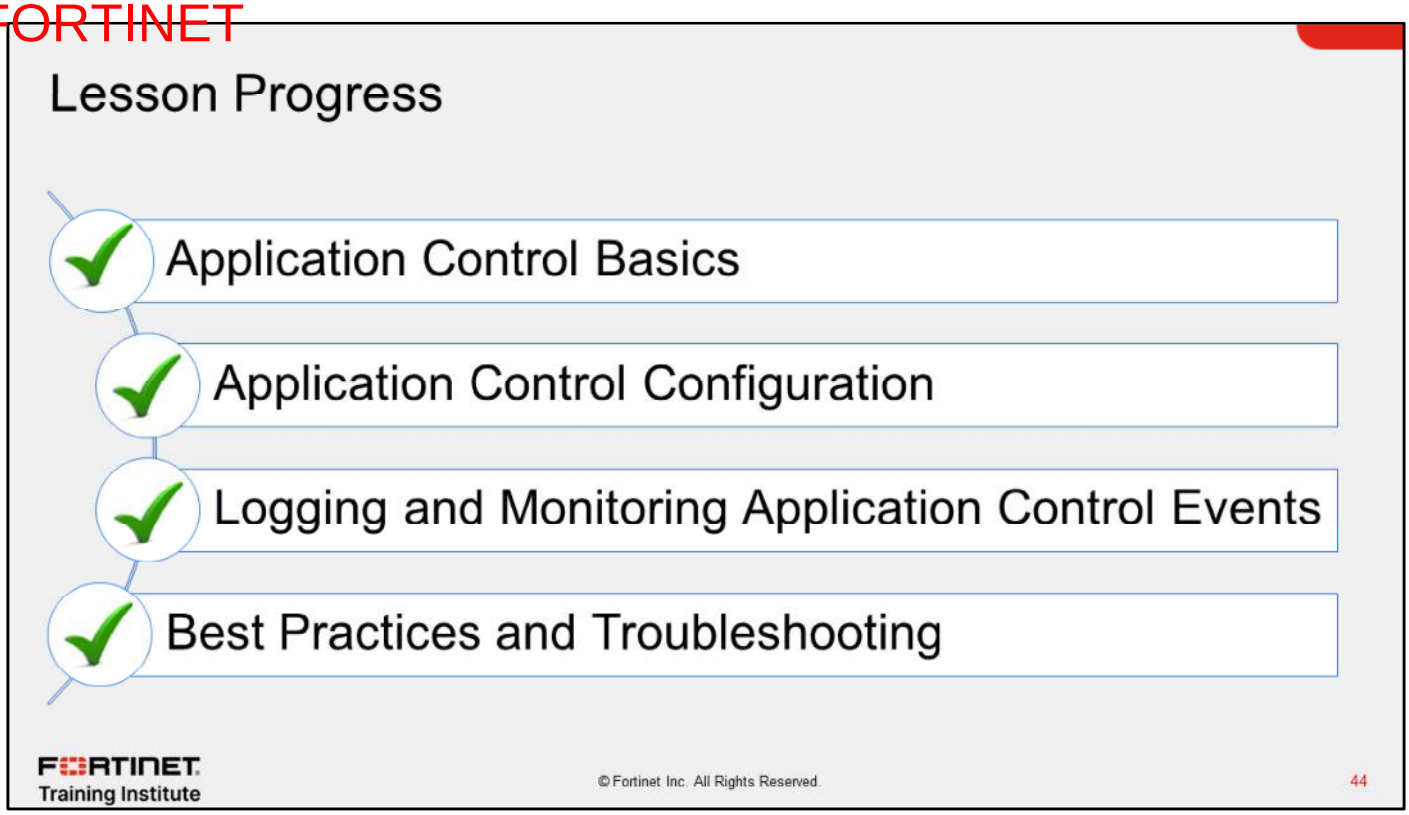

Congratulations! You have completed this lesson.

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

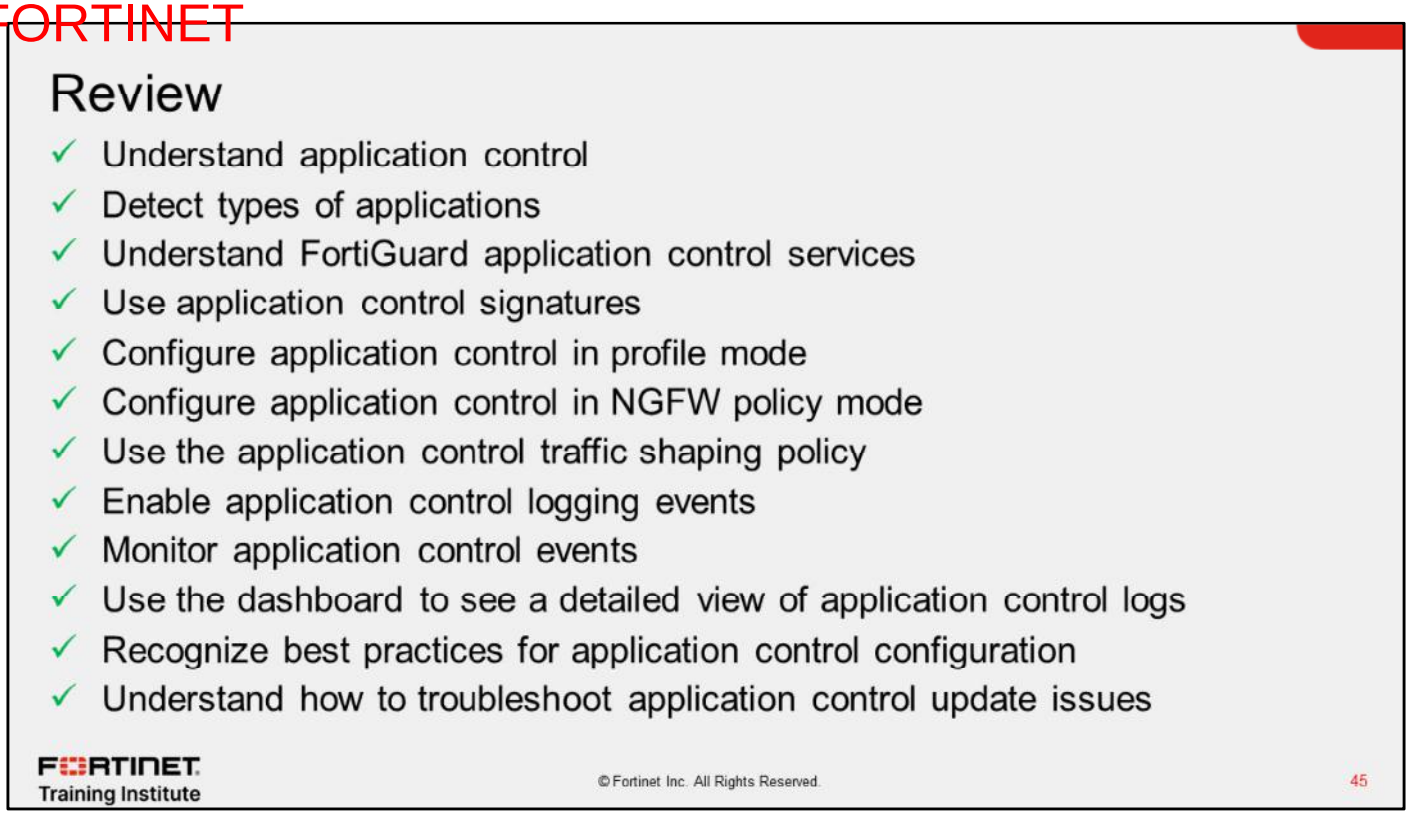

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use methods beyond simply blocking protocols, port numbers, or IP addresses, to monitor and control both standard and non-standard network applications.

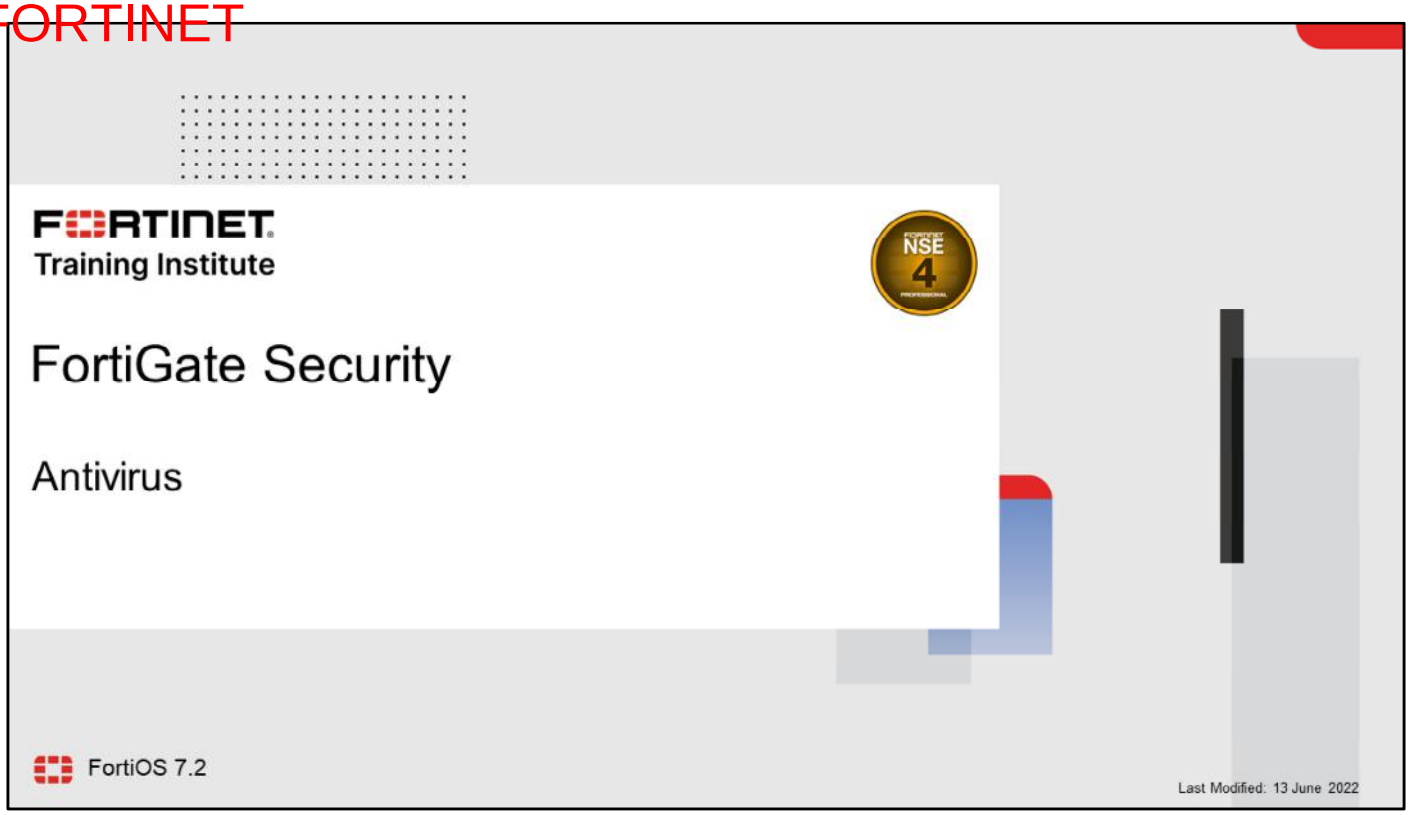

In this lesson, you will learn how to use FortiGate to protect your network against viruses.

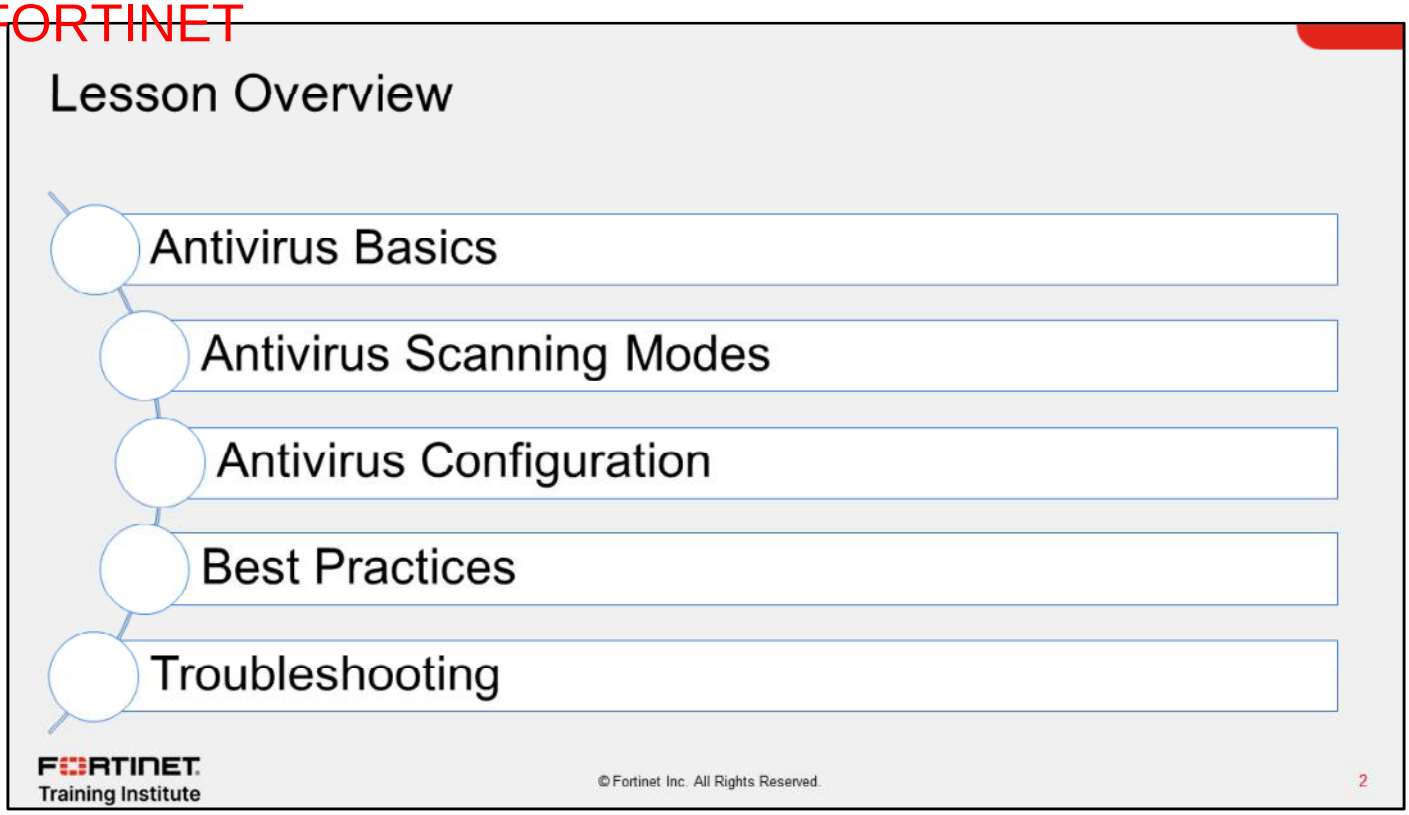

In this lesson, you will learn about the topics shown on this slide.

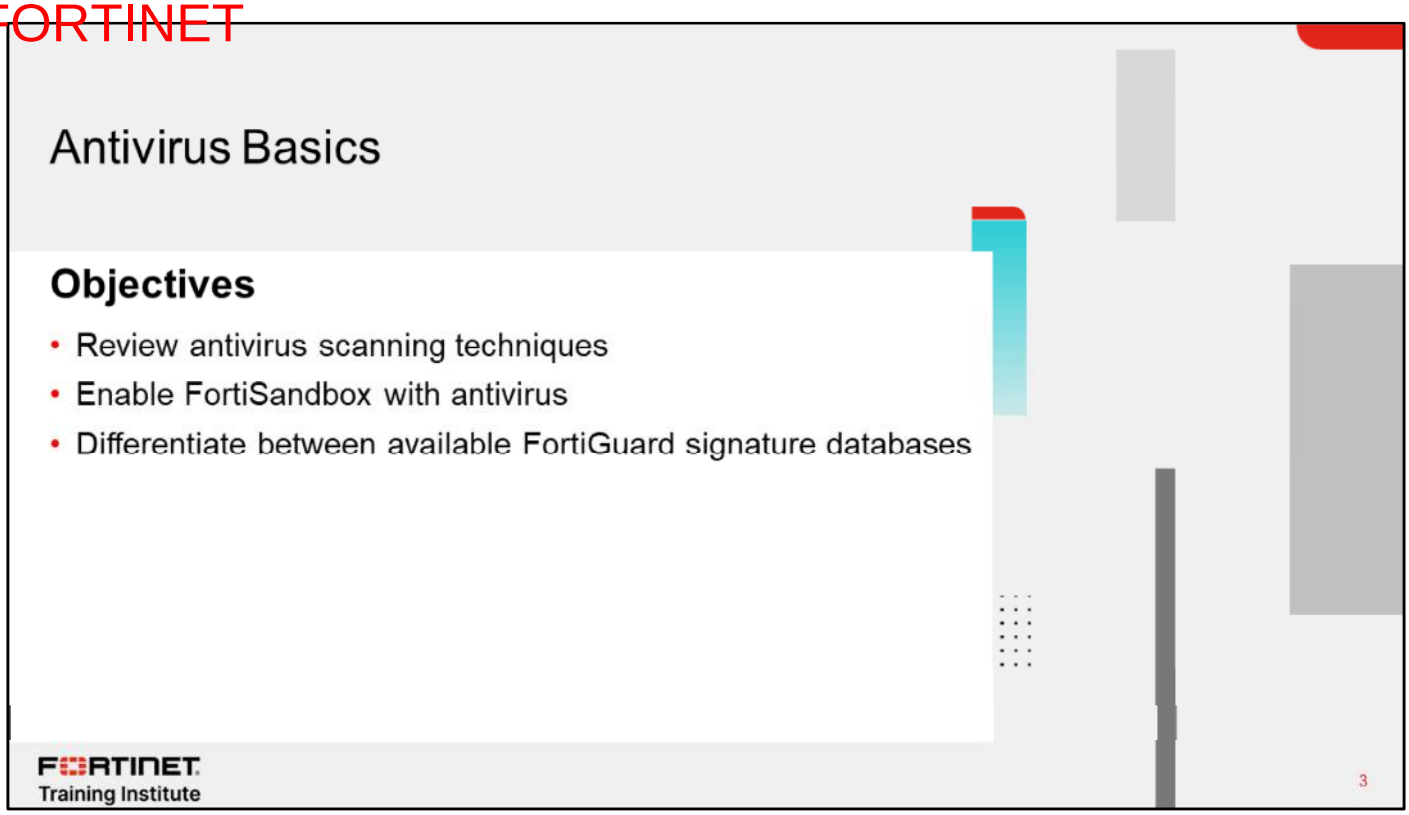

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus basics, you will be able to understand and apply antivirus on FortiGate.

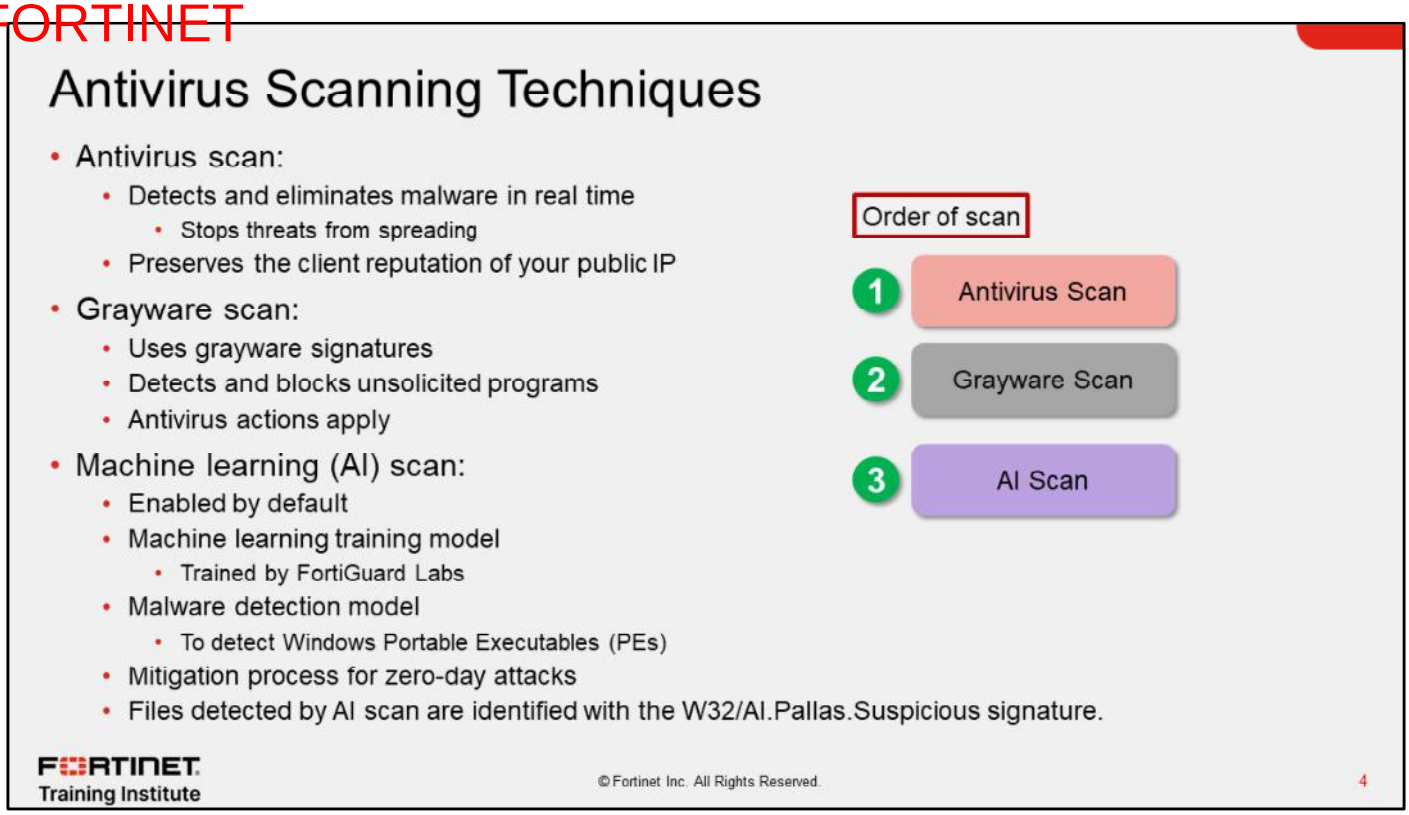

Like viruses, which use many methods to avoid detection, FortiGate uses many techniques to detect viruses. These detection techniques include:

- Antivirus scan: This is the first, fastest, simplest way to detect malware. It detects viruses that are an exact match for a signature in the antivirus database.
- Grayware scan: This scan detects unsolicited programs, known as grayware, that have been installed without the user's knowledge or consent. Grayware is not technically a virus. It is often bundled with innocuous software, but *does* have unwanted side effects, so it is categorized as malware. Often, grayware can be detected with a simple FortiGuard grayware signature.
- Machine learning (AI) scan: These scans are based on probability, so they increase the possibility of false positives, but they also detect zero-day attacks. Zero-day attacks are malwares that are new, unknown, and, therefore, have no existing associated signature. If your network is a frequent target, enabling an AI scan may be worth the performance cost because it can help you to detect a virus before the outbreak begins. Files detected by AI scan are identified with the W32/AI.Pallas.Suspicious signature.

If all antivirus features are enabled, FortiGate applies the following scanning order: antivirus scan, followed by grayware scan, followed by AI scan.

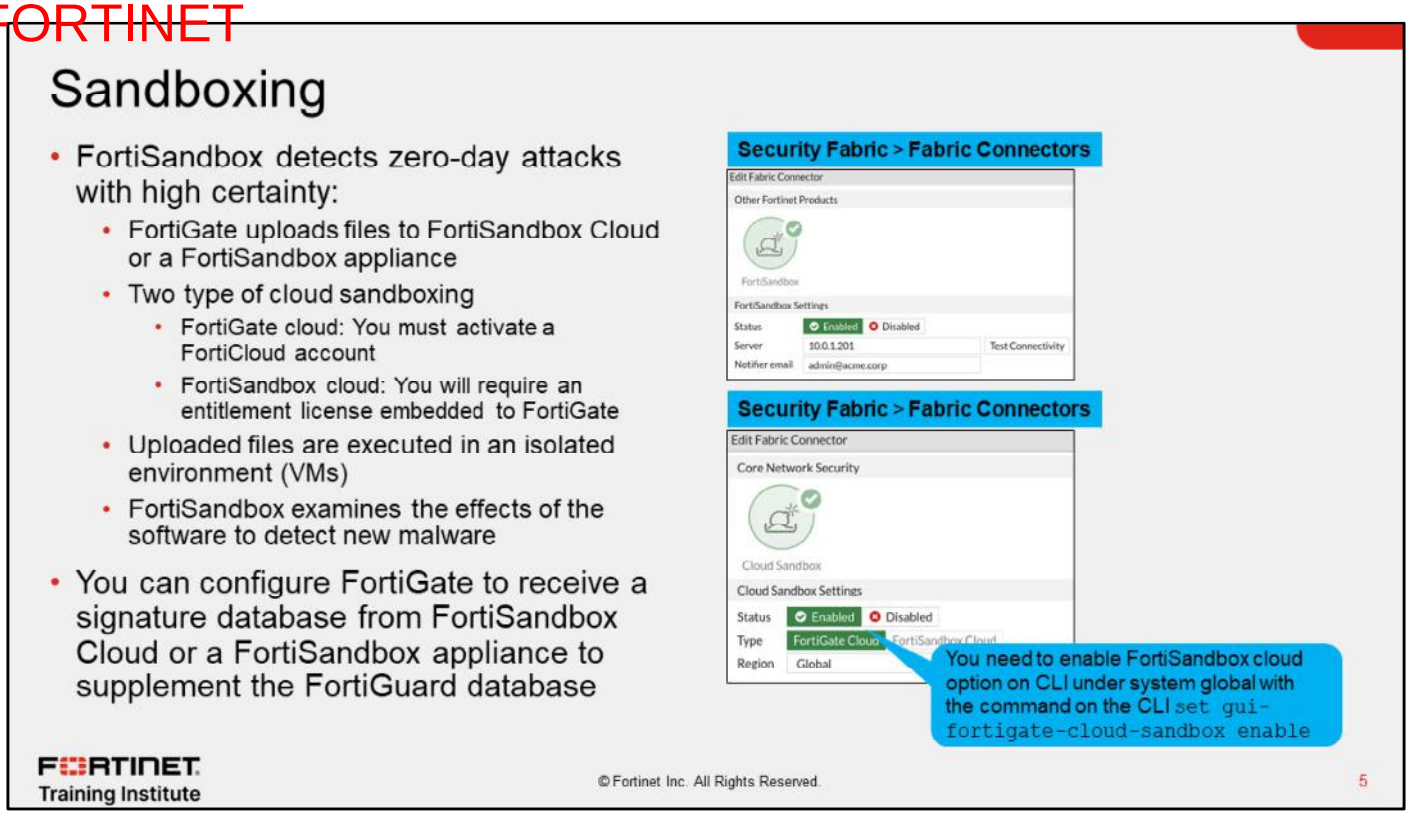

What if AI scans are too uncertain? What if you need a more sophisticated, more certain way to detect malware and find zero-day viruses?

You can integrate your antivirus scans with either FortiSandbox Cloud or a FortiSandbox appliance. Note you will need to enable cloud sandboxing on the CLI under system global settings for configuration options to appear on GUI . For environments that require more certainty, FortiSandbox executes the file within a protected environment (VMs), then examines the effects of the software to see if it is dangerous.

For example, let's say you have two files. Both alter the system registry and are, therefore, suspicious. One is a driver installation—its behavior is normal—but the second file installs a virus that connects to a botnet command and control server. Sandboxing would reveal the difference.

FortiGate can be configured to receive a supplementary signature database from FortiSandbox based on the sandboxed results.

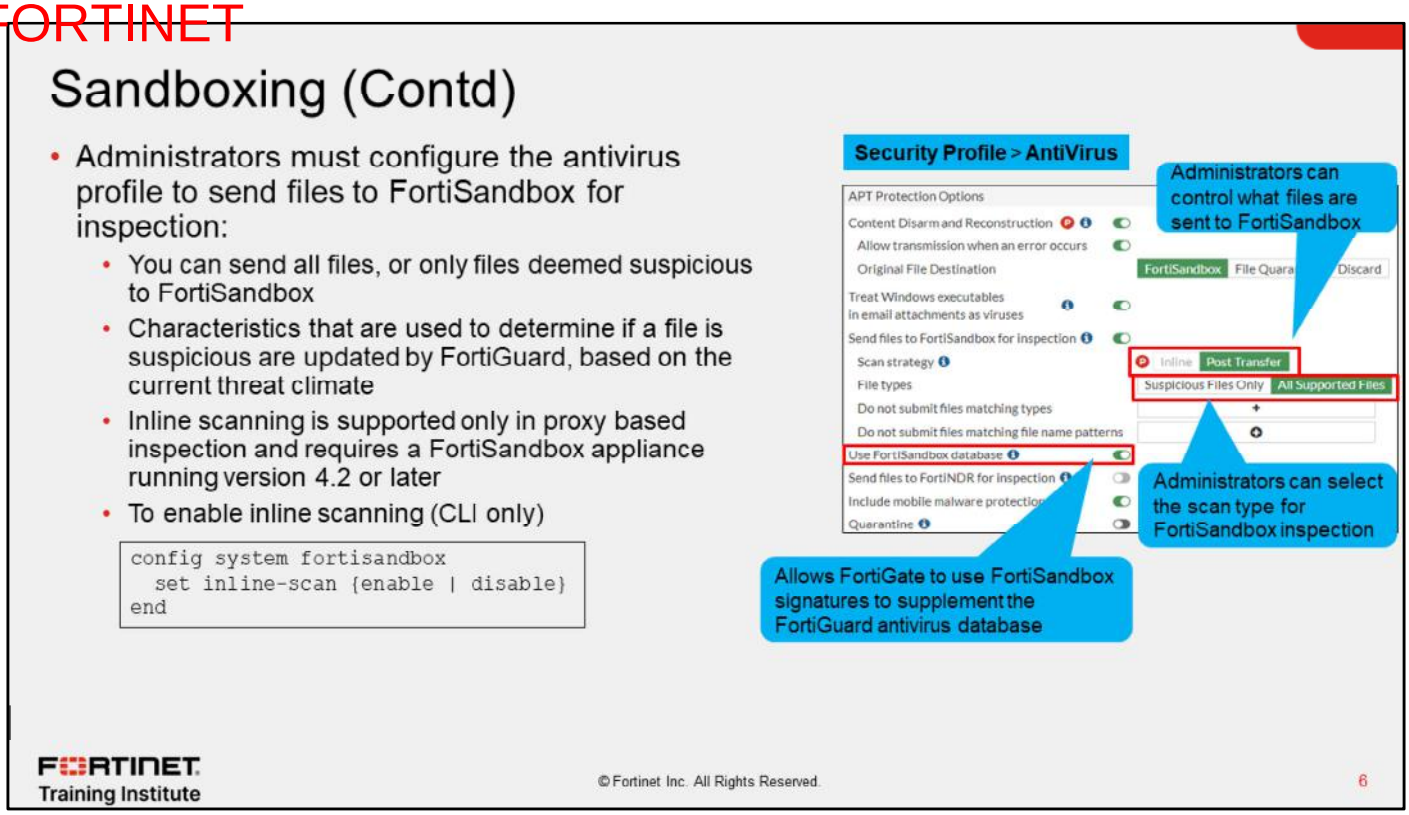

FortiOS is smart when it comes to determining what files are sent to FortiSandbox. One feature FortiOS uses for this is content disarm and reconstruction (CDR), a proxy-based feature that you will learn more about in this lesson. When CDR processes files, the original documents can be saved to FortiSandbox.

FortiGuard provides FortiGate with information based on the current threat climate, that is used to determine if a file should be deemed suspicious or not. FortiGate provides the administrator with granular control when it comes to determining what type of files are sent to FortiSandbox for further investigation. Administrators also have the option to use the FortiSandbox database, in conjunction with the FortiGuard antivirus database, to enhance their network security.

FortiSandbox inline scanning is supported only in proxy inspection mode. You will need to enable inline scanning under system fortisandbox settings and then select **Inline** in the antivirus profile. When the setting is enabled, the client's file is held by FortiSandbox for inspection, and an appropriate configured action is applied once a verdict is returned. Inline scanning is not supported on FortiSandbox Cloud or FortiGate Cloud Sandbox.

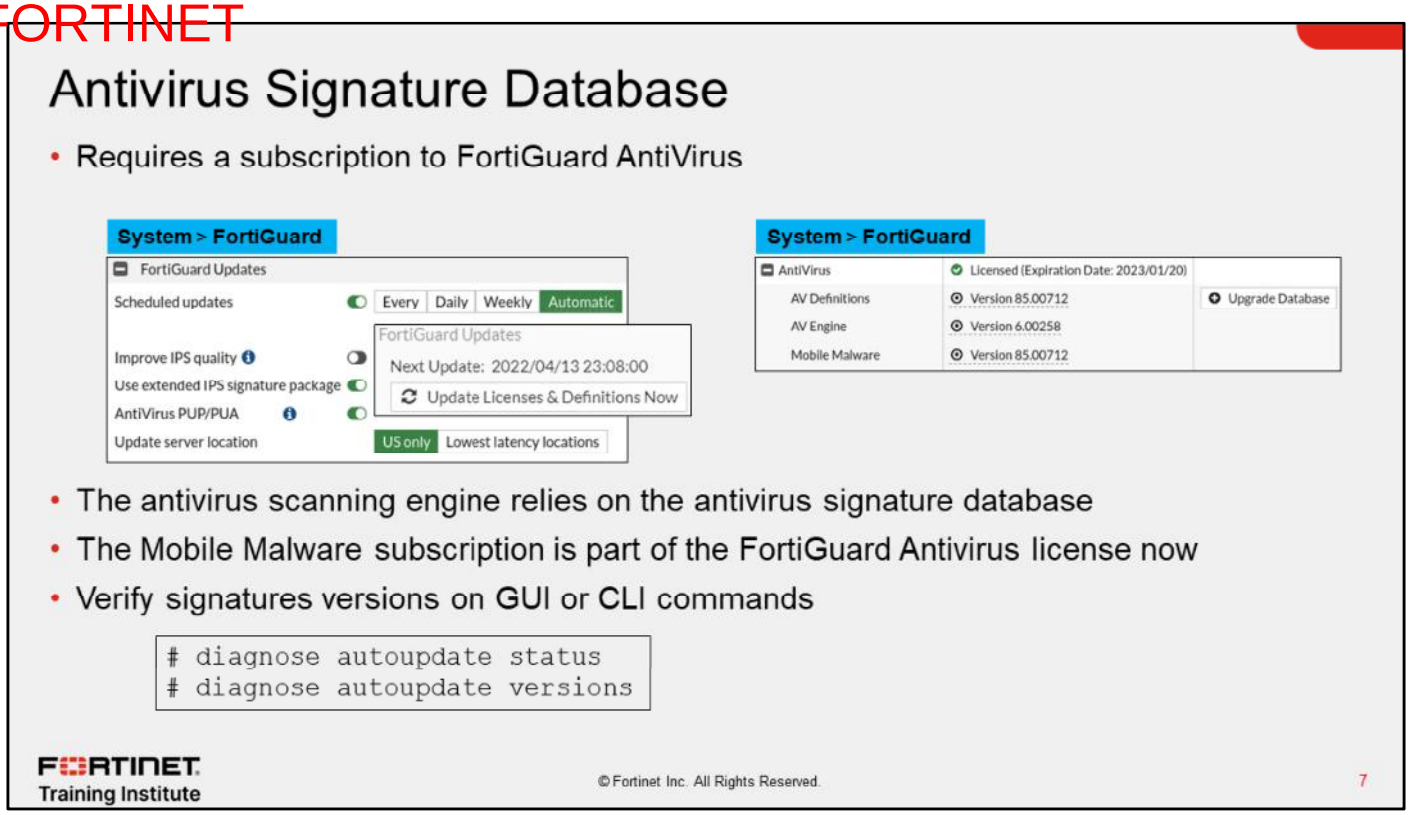

Scheduled updates allow you to configure scheduled updates at regular intervals, such as hourly, daily, weekly, or automatically within every hour. You can also enable **AntiVirus PUP/PUA**, which allows antivirus grayware checks for potentially unwanted programs and applications.

Regardless of which method you select, you *must* enable virus scanning in at least one firewall policy. Otherwise, FortiGate will not download any updates. Alternatively, you can download packages from the Fortinet customer service and support website (requires subscription), and then manually upload them to your FortiGate. You can verify the update status and signature versions from the **FortiGuard** page on the GUI or using the CLI console.

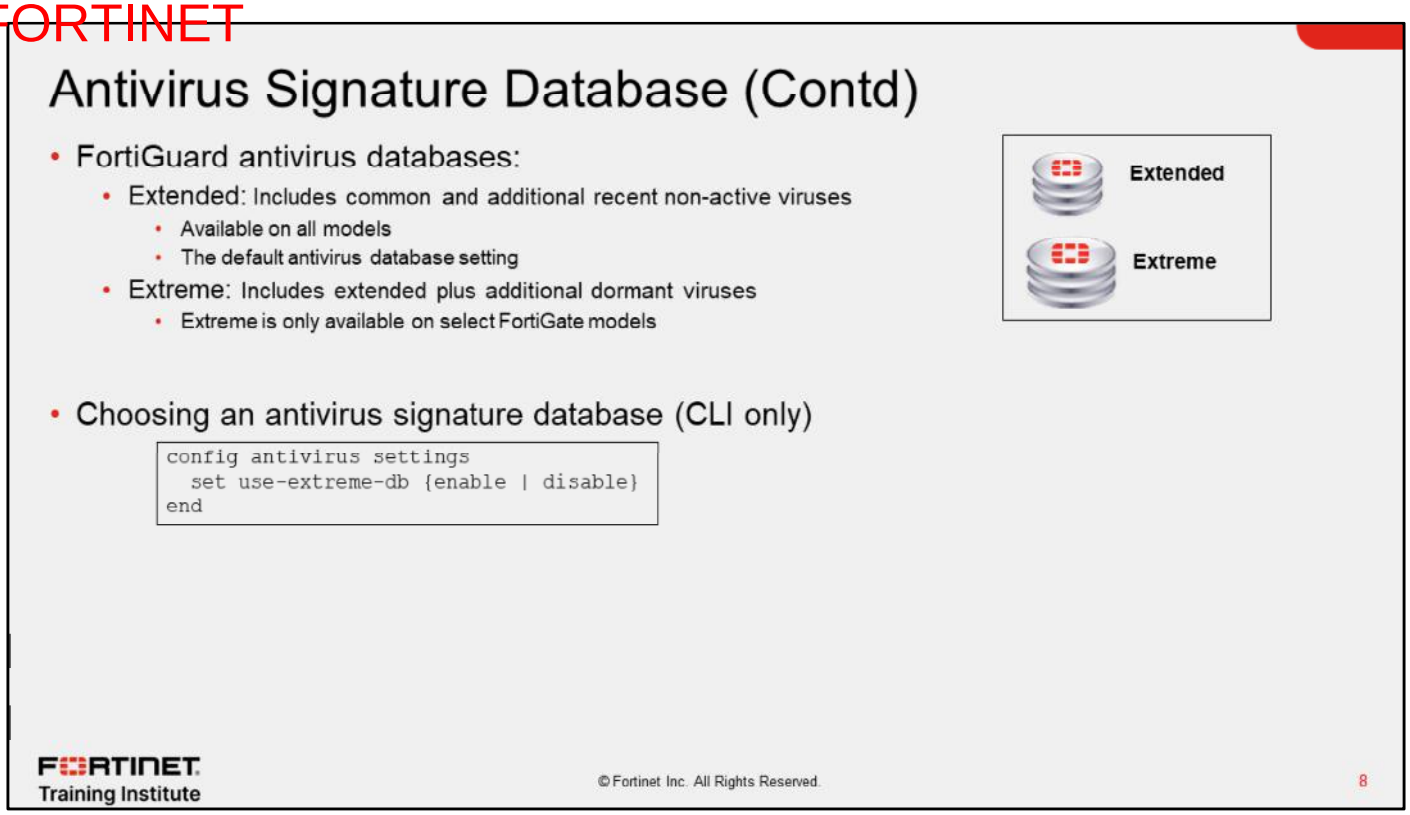

Multiple FortiGuard antivirus databases exist, which you can configure using CLI commands. Support for each database type varies by FortiGate model.

All FortiGate devices include the extended database. The extended database contains signatures for viruses that have been detected in recent months, as identified by the FortiGuard Global Security Research Team. The extended database also detects viruses that are no longer active.

The extreme database is intended for use in high-security environments. The extreme database detects all known viruses, including viruses targeted at legacy operating systems that are no longer widely used. Most FortiGate models support the extreme database.

FortiGuard Protection Services

- CDR
  - CDR removes exploitable content and replaces it with content that's known to be safe
- Virus outbreak prevention
  - Additional layer of protection that keeps your network safe from newly emerging malware
  - Quick virus outbreaks can infect a network before signatures can be developed to stop them
  - Outbreak protection stops these virus outbreaks until signatures become available in FortiGuard
- Malware block list
  - Manual external malware signatures to support antivirus database
  - The block list can be in the form of MD5, SHA1, and SHA256 hashes
  - Defined as a Security Fabric connector

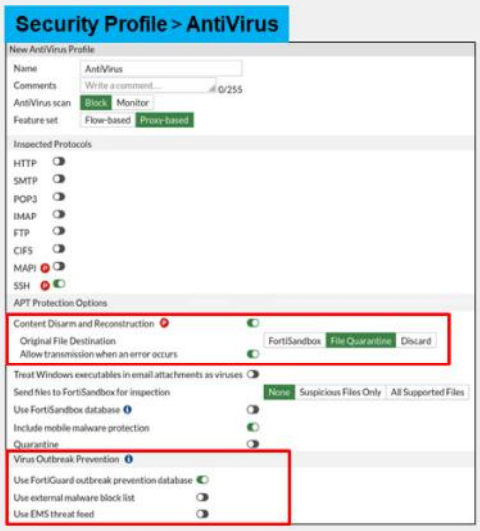

FORTINET Training Institute

9

CDR: The CDR removes exploitable content and replaces it with content that's known to be safe. As files are processed through an enabled antivirus profile, content that's found to be malicious or unsafe is replaced with content that allows the traffic to continue, but doesn't put the recipient at risk. Content that can be scanned includes PDF and Microsoft Office files leaving the network on CDR-supported protocols (such as HTTP, SMTP, IMAP, and POP3—MAPI isn't supported). When the client tries to download the file, FortiGate removes all exploitable content in real-time, and then sends the original file to FortiSandbox for inspection. The client can download the original file by logging in to FortiSandbox.

Virus outbreak prevention: This is an additional layer of protection that keeps your network safe from newly emerging malware. Quick virus outbreaks can infect a network before signatures can be developed to stop them. Outbreak protection stops these virus outbreaks until signatures become available in FortiGuard. FortiGate must have a zero-hour virus outbreak (ZHVO) license. FortiGate adds hash-based virus detection for new threats that are not yet detected by the antivirus signatures. When the file is sent to the scanunit deamon, buffers are hashed and a request is sent to the urlfilter deamon. After checking against its request cache for known signatures, the urlfilter deamon sends an antivirus request to FortiGuard with the remaining signatures. FortiGuard returns a rating that is used to determine if the scanunit deamon should report the file as harmful or not. Jobs remain suspended in the scanunit deamon until the client receives a response, or the request times out.

Malware block list: FortiGate can enhance the antivirus database by linking a dynamic external malware block list to FortiGate. The list is hosted on a web server and is available through HTTP/HTTPS URL defined within the Security Fabric malware hash list. The hash list can be in the form of MD5, SHA1, and SHA256 hashes, and is written on separate lines on a plaintext file. The malware block list can be defined as a Security Fabric connector and configured to pull the list dynamically, by setting the refresh rate.

Knowledge Check

1. If antivirus, grayware, and AI scans are enabled, in what order are they performed?
   - A. AI scan, followed by grayware scan, followed by antivirus scan
   - ✓ B. Antivirus scan, followed by grayware scan, followed by AI scan

2. Which databases can be manually selected for use in antivirus scanning?
   - ✓ A. Extended and Extreme
   - B. Quick, Normal, and Extreme

FORTINET Training Institute

10

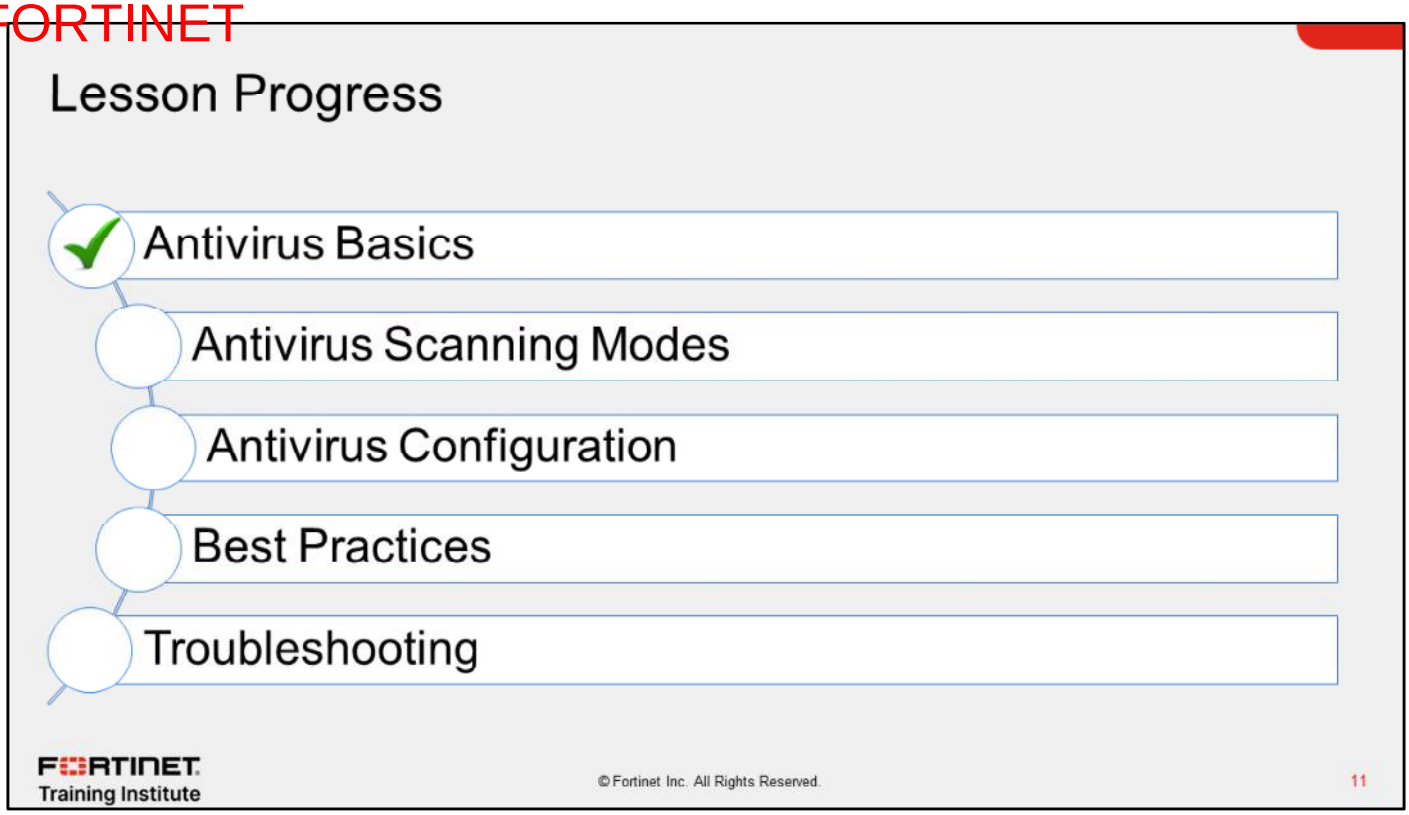

Good job! You now understand the basics of antivirus functionality.

Now, you will learn about antivirus scanning modes.

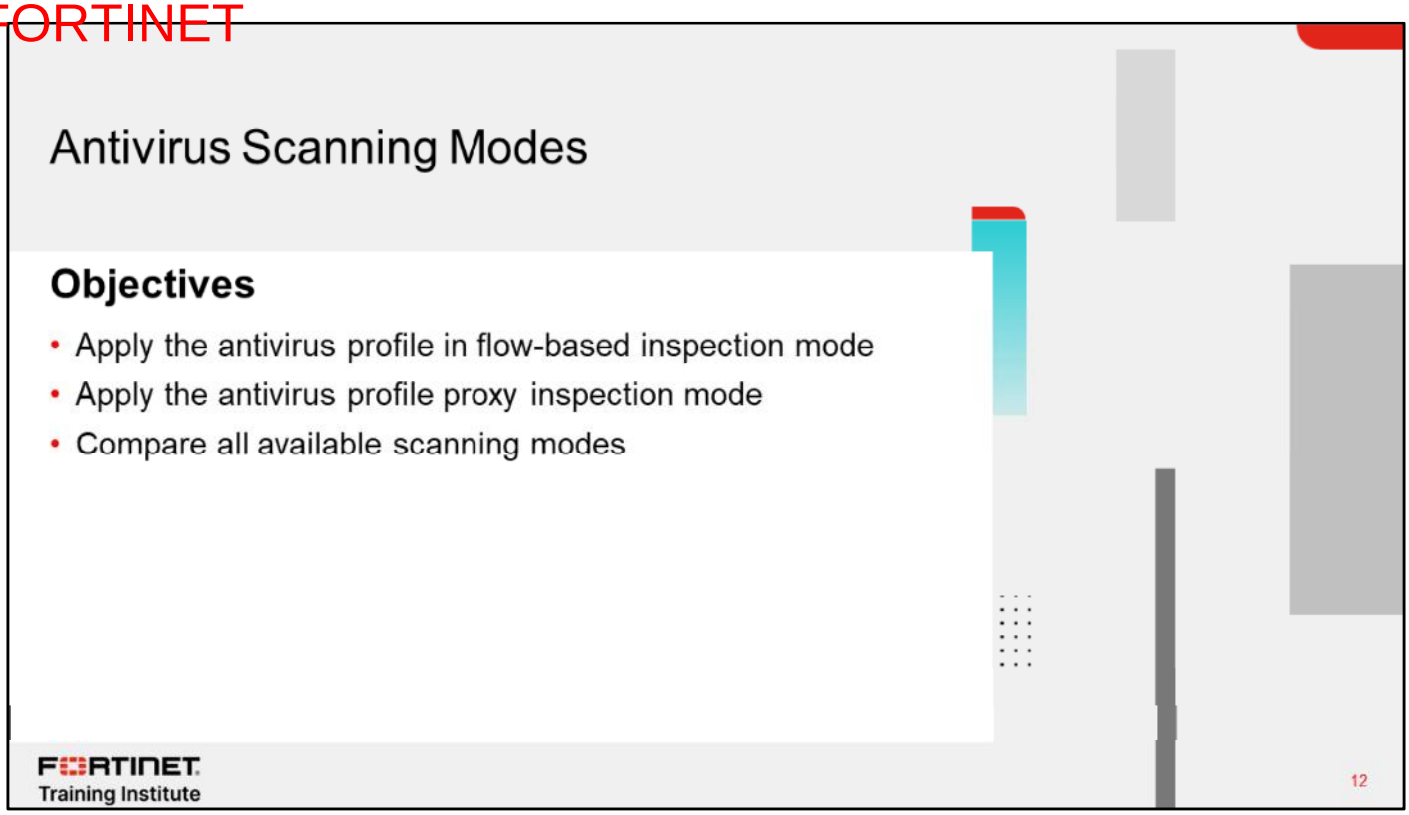

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in all antivirus scanning modes available in FortiOS, you will be able to use the antivirus profile in an effective manner.

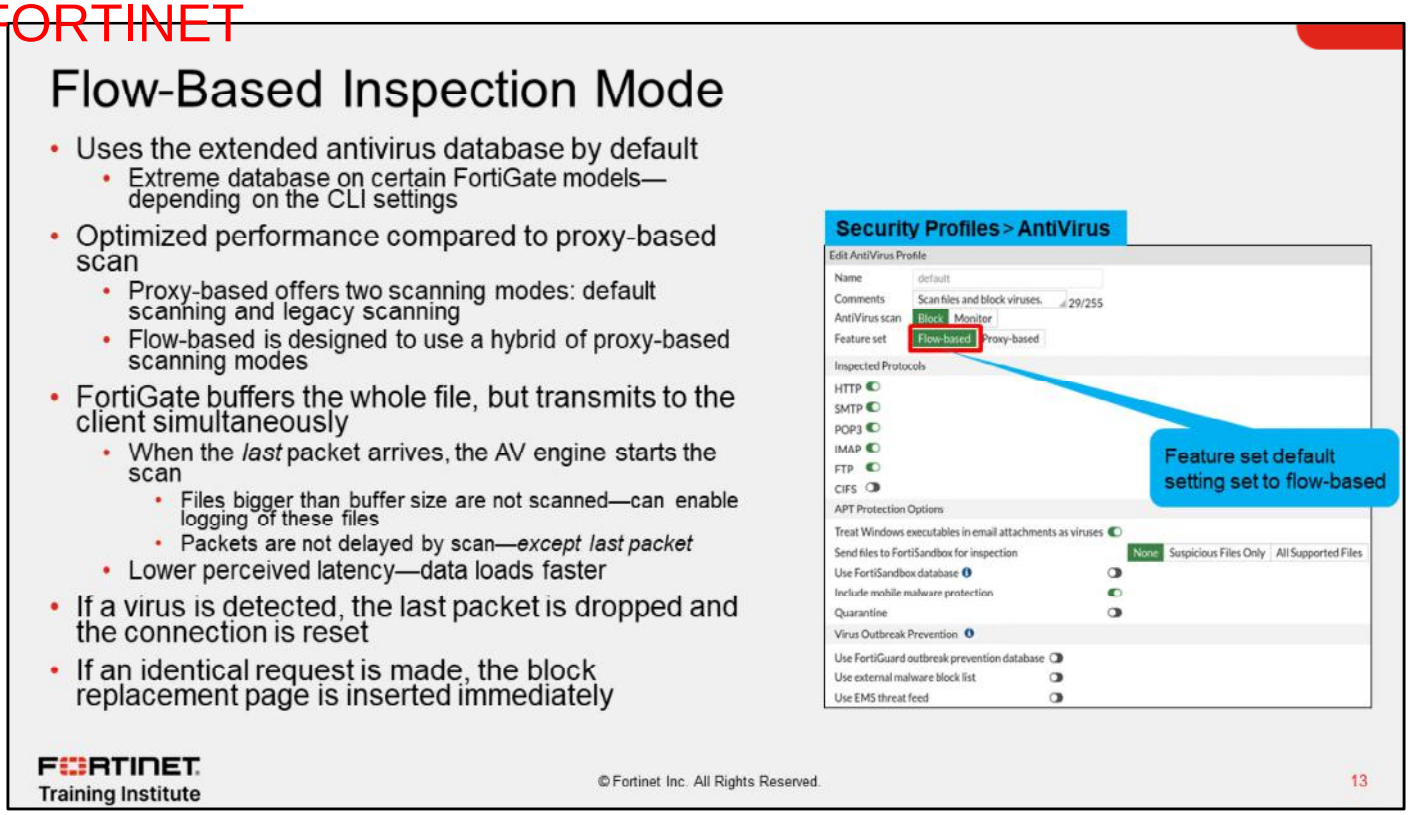

AV can operate in flow-based or proxy-based inspection mode, both of which use the full AV database (extended or extreme–depending on the CLI settings).

Flow-based inspection mode uses a hybrid of the scanning modes available in proxy-based inspection: the default scanning mode and the legacy scanning mode. The default mode enhances the scanning of nested archive files without buffering the container archive file. The legacy mode buffers the full container, and then scans it.

In flow-based inspection mode, the IPS engine reads the payload of each packet, caches a local copy, and forwards the packet to the receiver at the same time. Because the file is transmitted simultaneously, flow-based mode consumes more CPU cycles than proxy-based. However, depending on the FortiGate model, some operations can be offloaded to SPUs to improve performance. When FortiGate receives the last packet of the file, it puts the packet on hold and sends a copy to the IPS engine. The IPS engine extracts the payload and assembles the whole file, and then sends the whole file to the AV engine for scanning.

Two possible scenarios can occur when a virus is detected:

- When a virus is detected on a TCP session where some packets have been already forwarded to the receiver, FortiGate resets the connection and does not send the last piece of the file. Although the receiver got most of the file content, the file has been truncated and therefore, can't be opened. The IPS engine also caches the URL of the infected file, so that if a second attempt to transmit the file is made, the IPS engine will then send a block replacement message to the client instead of scanning the file again.
- If the virus is detected at the start of the connection, the IPS engine sends the block replacement message immediately.

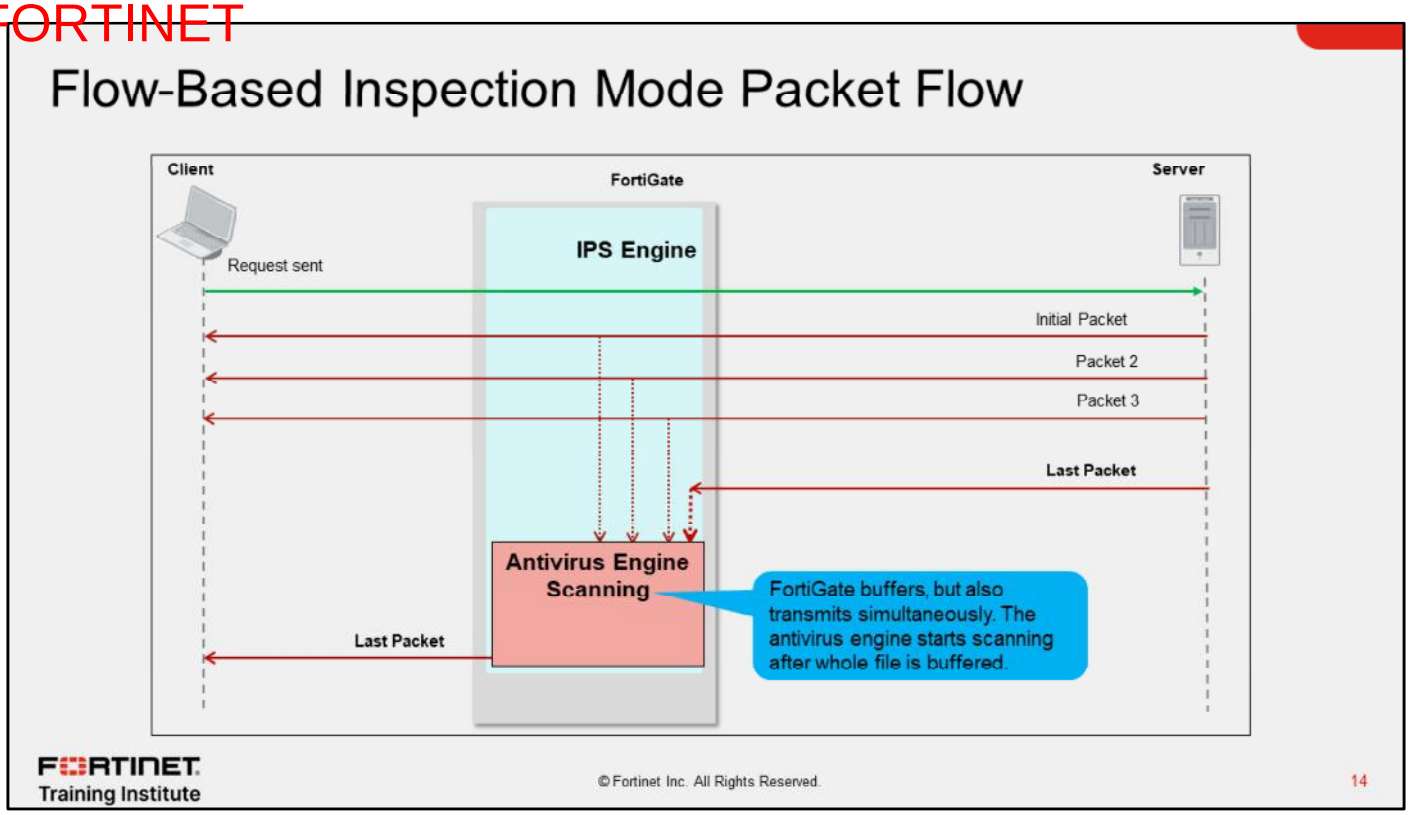

As you can see on this slide, the client sends a request and starts receiving packets immediately, but FortiGate also caches those packets at the same time. When the last packet arrives, FortiGate caches it and puts it on hold. Then, the IPS engine extracts the payload of the last packet, assembles the whole file, and sends it to the antivirus engine for scanning. If the antivirus scan does not detect any viruses, and the result comes back clean, the last cached packet is regenerated and delivered to the client. However, if a virus is found, the last packet is dropped. Even if the client has received most of the file, the file will be truncated and the client will be not able to open a truncated file.

Regardless of which mode you use, the scan techniques give similar detection rates. How can you choose between the scan engines? If performance is your top priority, then flow inspection mode is more appropriate. If security is your priority, proxy inspection mode—with client comforting disabled—is more appropriate.

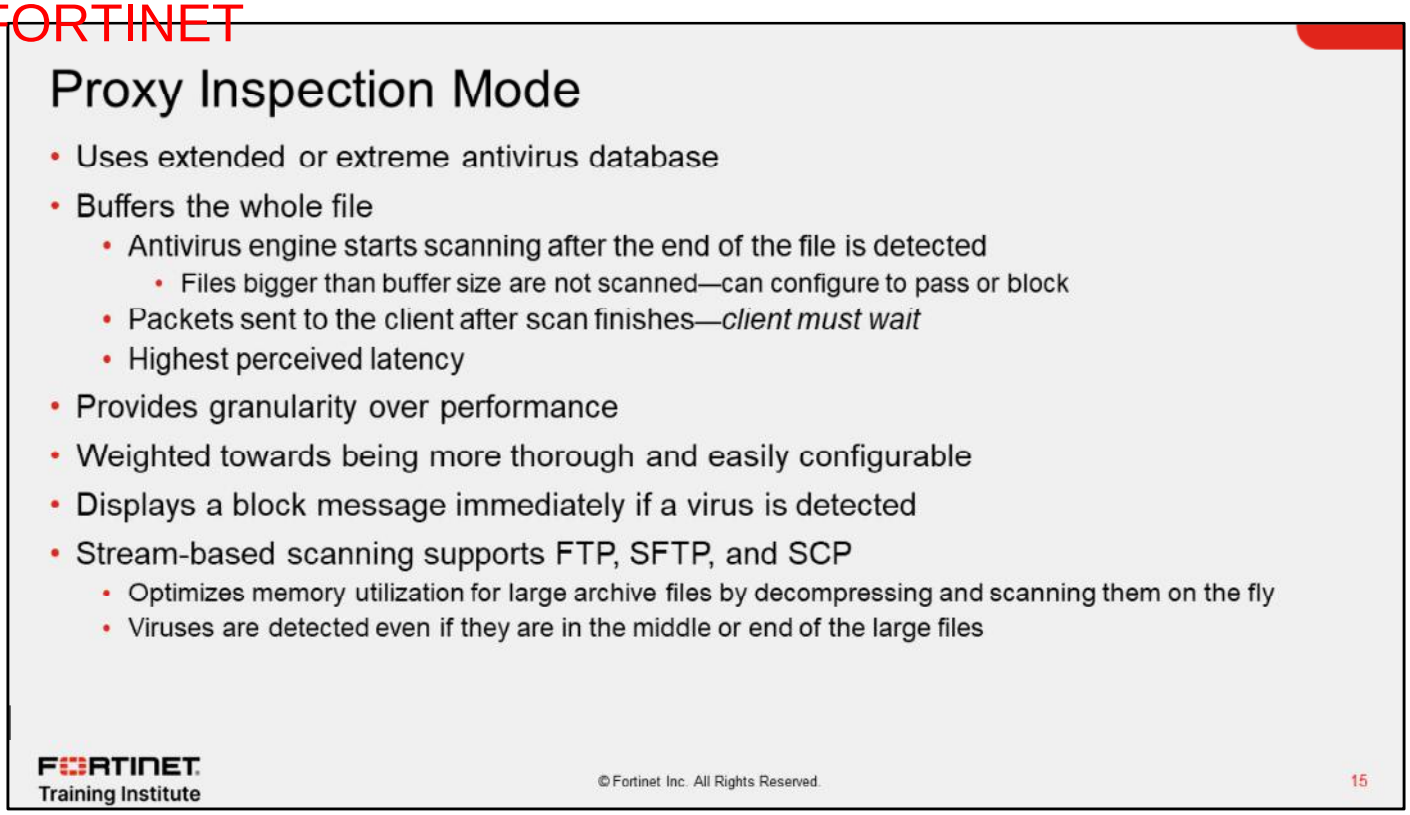

Each protocol's proxy picks up a connection and buffers the entire file first (or waits until the oversize limit is reached) before scanning. The client must wait for the scanning to finish. If a virus is detected, the block replacement page is displayed immediately. Because FortiGate has to buffer the whole file and then do the scanning, it takes a long time to scan. Also, from the client point of view, it has to wait for the scanning to finish and might terminate the connection due to lack of data.

You can configure client comforting for HTTP and FTP from the `config firewall profile-protocol-options` command tree. This allows the proxy to slowly transmit some data until it can complete the buffer and finish the scan. This prevents a connection or session timeout. No block replacement message appears in the first attempt, as FortiGate is transmitting the packets to the end client.

Using proxy inspection antivirus allow you to use the stream-based scanning, which is enabled by default. Stream-based scanning scans large archive files by decompressing the files and then scanning and extracting them at the same time. This process optimized memory utilization to conserve resources on FortiGate. Viruses are detected even if they are in the middle or towards the end of these large files.

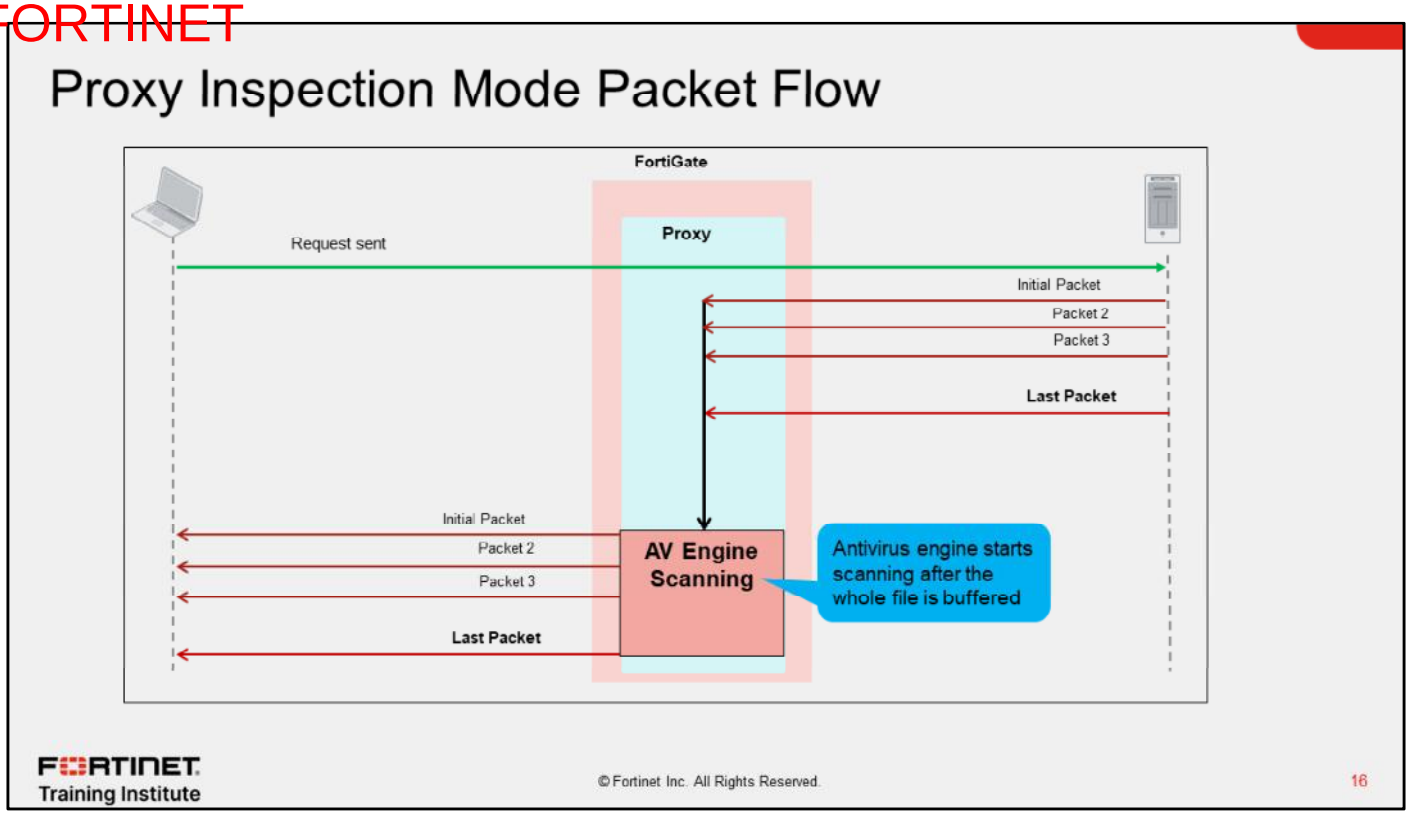

With a proxy inspection mode scan, the client sends a request and FortiGate starts buffering the whole file, then sends it to the antivirus engine for scanning. If the file is clean (without any viruses), FortiGate starts transmitting the file to the end client. If a virus is found, no packets are delivered to the end client and the proxy sends the replacement block message to the end client.

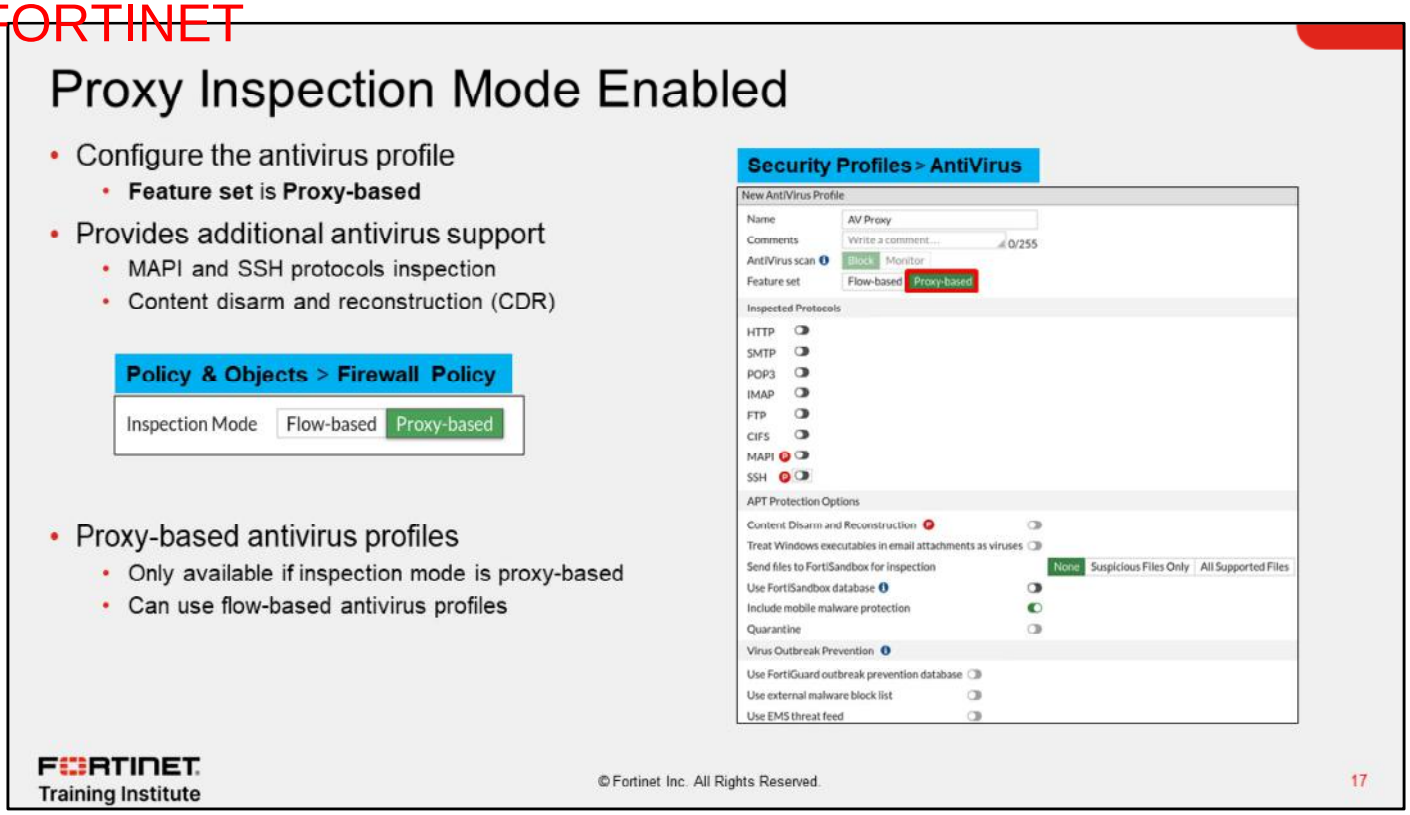

Applying a proxy-based antivirus profile requires two sections in FortiGate configuration to use non-default settings:

1. Antivirus profile
2. Firewall policy

Antivirus profile provides the option to select a proxy-based approach as the inspection mode within the profile. This allows the profile to inspect MAPI and SSH protocols traffic, as well as to sanitize Microsoft documents and PDF files using the content disarm and reconstruction (CDR) feature.

If the inspection mode on the antivirus profile is set to **Proxy-based**, it is only available when the firewall policy inspection mode is set to **Proxy-based**.

Knowledge Check

1. What three additional features of an antivirus profile are available in proxy-based inspection mode?
   - ✓ A. MAPI, SSH, and CDR
   - B. Full and quick

2. What antivirus database is limited to specific FortiGate models?
   - A. Extended
   - ✓ B. Extreme

FORTINET Training Institute

18

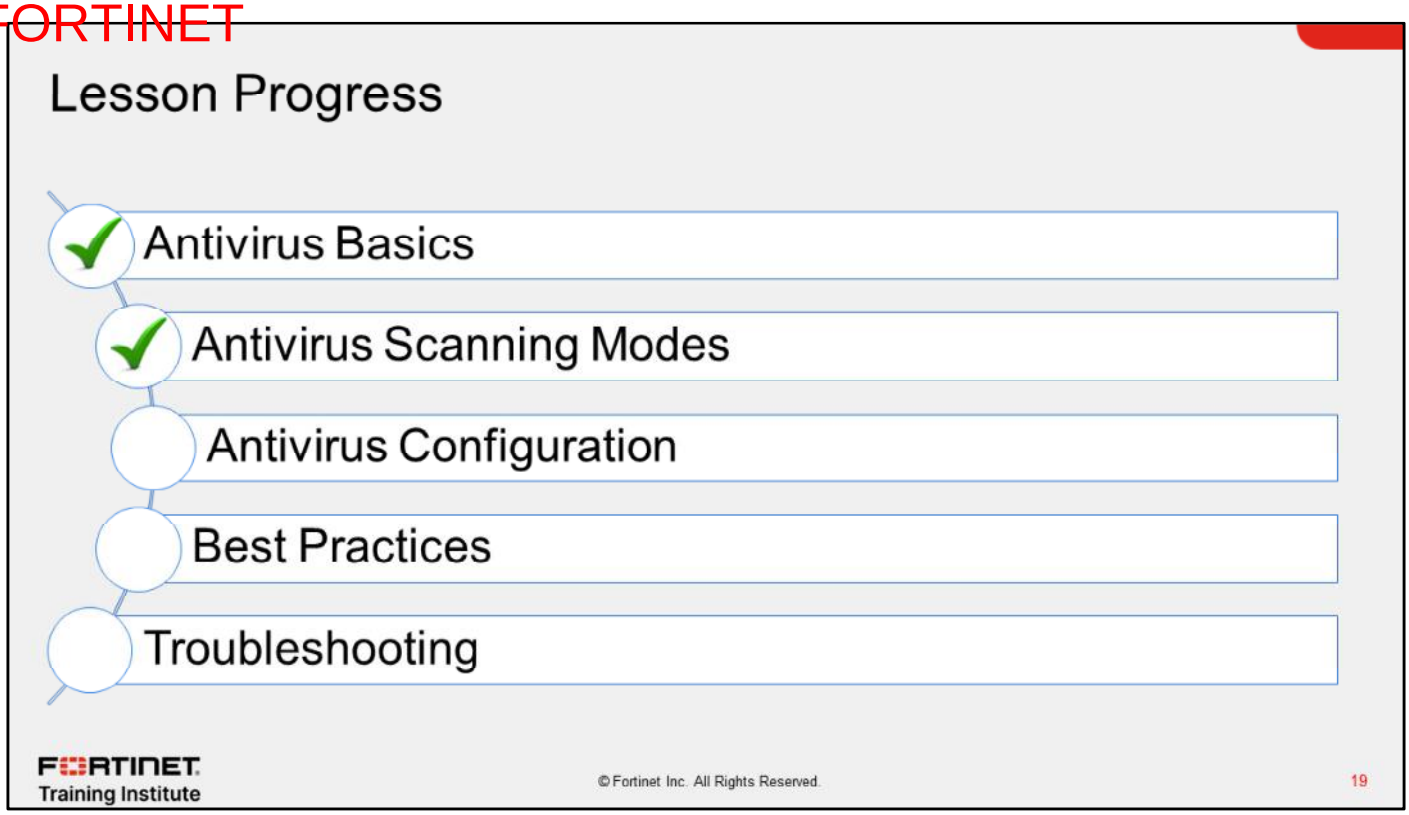

Good job! You now understand antivirus scanning modes.

Now, you will learn about antivirus configuration.

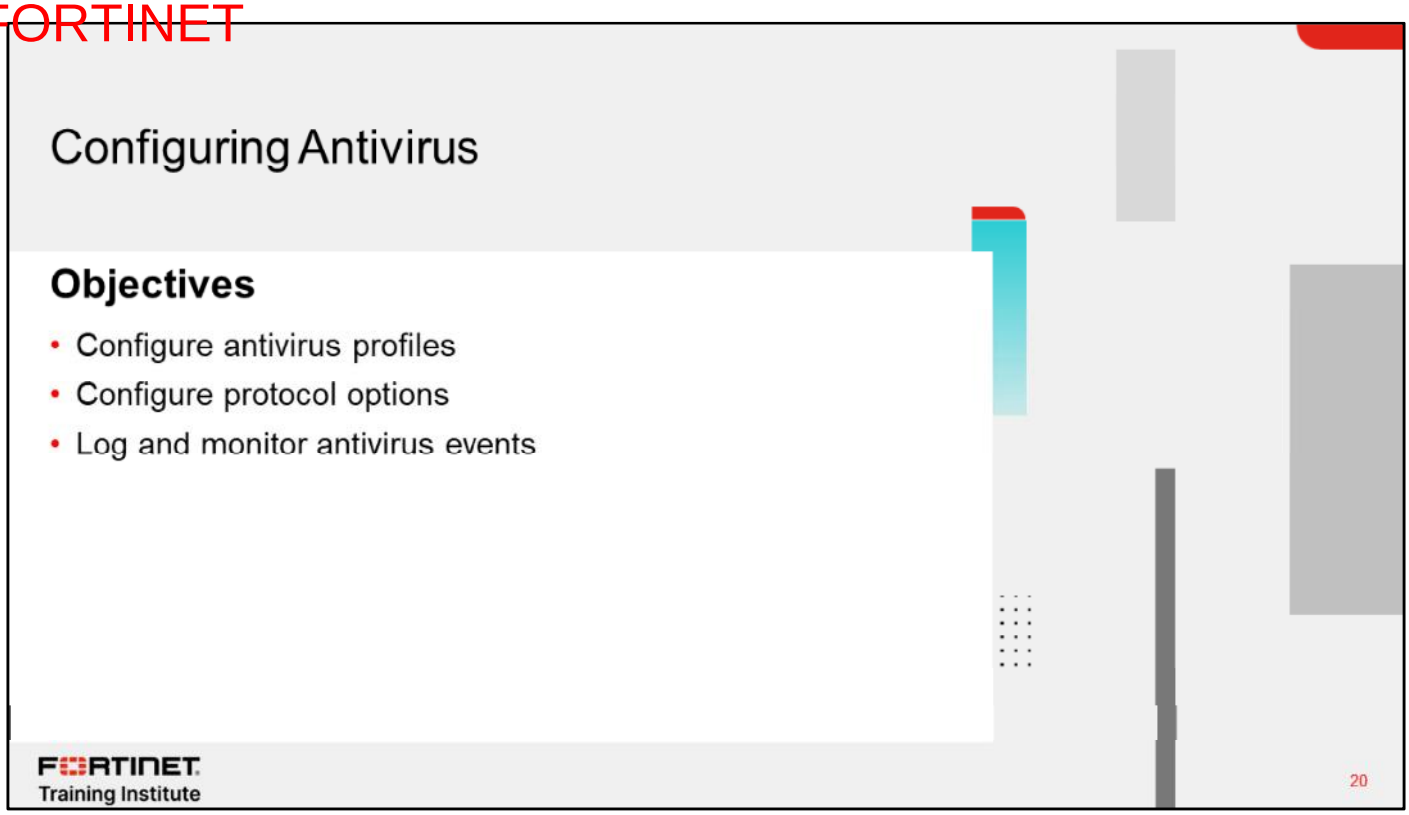

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus configuration, including reviewing antivirus logs, you will be able to use the antivirus profile in an effective manner.

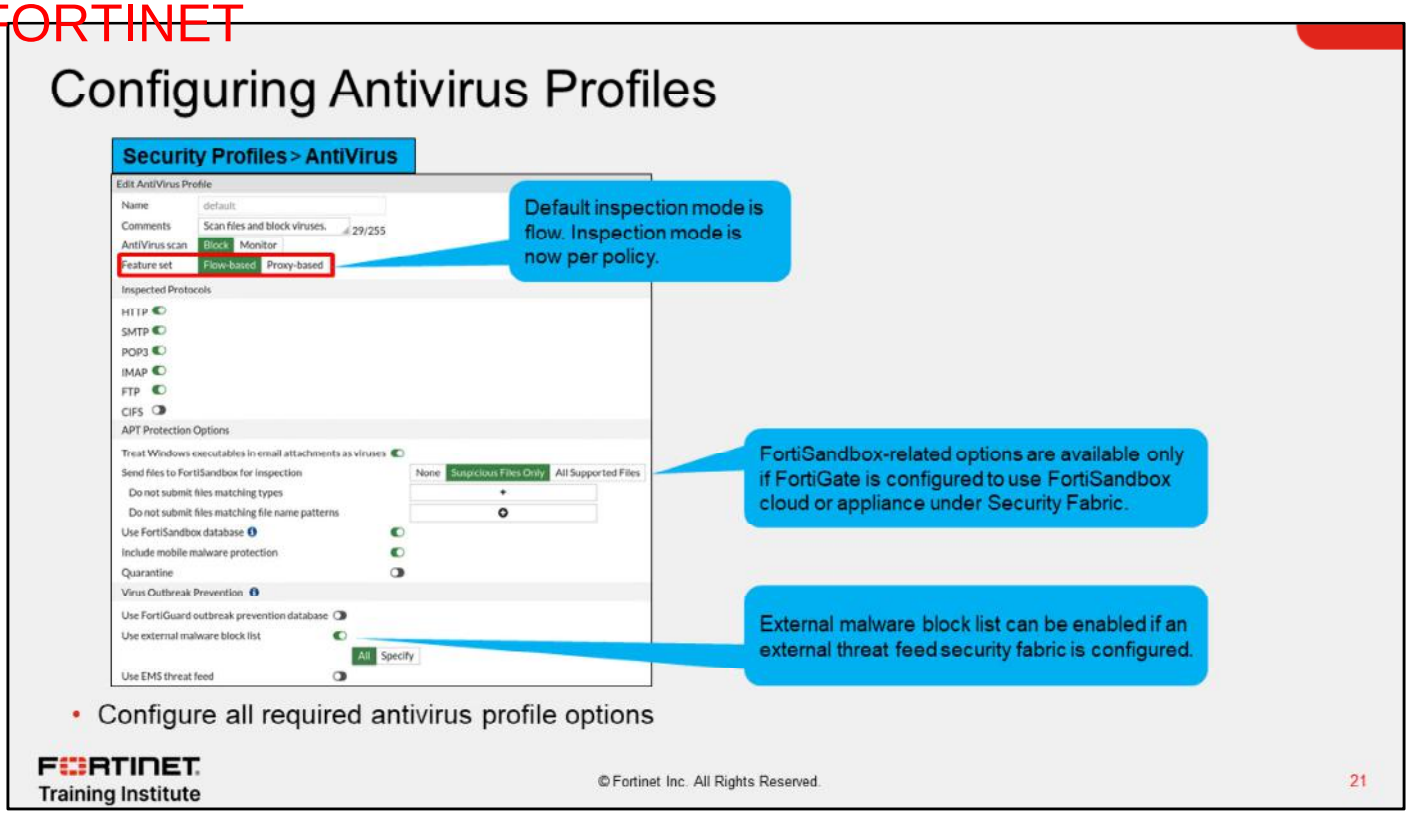

The antivirus profile can be configured on the **AntiVirus** page. Since the default inspection mode on a firewall policy is flow-based, **Feature set** is required to be set to **Flow-based**. If the inspection mode of the firewall policy is proxy-based, **Feature set** can be set to **Proxy-based,** which allows specific functions that are only available using proxy-based inspection mode firewall policy such as MAPI protocol and CDR.

Both feature sets provide the following options:
APT Protection Options:

- **Treat Windows executables in email attachment as viruses**: By default, this option is enabled and files (including compressed files) identified as Windows executables can be treated as viruses.
- **Send files to FortiSandbox for inspection:** If FortiSandbox cloud or appliance is configured, you can configure the antivirus profile to send malicious files to FortiSandbox for behaviour analysis. If tagged as malicious, any future files matching the same behavior will be blocked if **Use FortiSandbox database** is enabled.

Virus Outbreak Prevention:

- **Use FortiGuard Virus outbreak prevention database**: FortiGuard virus outbreak prevention is an additional layer of protection that keeps your network safe from newly emerging malware. Quick virus outbreaks can infect a network before signatures can be developed to stop them. Outbreak protection stops these virus outbreaks until signatures become available on FortiGuard.
- **Use external malware block List:** FortiGate can enhance the antivirus database by linking a dynamic external malware block list to FortiGate. Malware block list can be defined as a Security Fabric connector and configured to pull the list dynamically by setting the refresh rate.

In the antivirus profile, you can define what FortiGate should do if it detects an infected file. After you configure an antivirus profile, you must apply it in the firewall policy.

Configuring Protocol Options

- More granular control
- Allows configuration of:
  - Protocol port mappings
  - Common options
  - Web and email options
- Configure for both proxy-based and flow-based firewall policies
  - From the GUI, on the **Protocol Options** page
  - From the CLI, using the `config firewall profile-protocol-options` command

```
config firewall profile-protocol-options
edit <profile_name>
config <protocol_name>
```

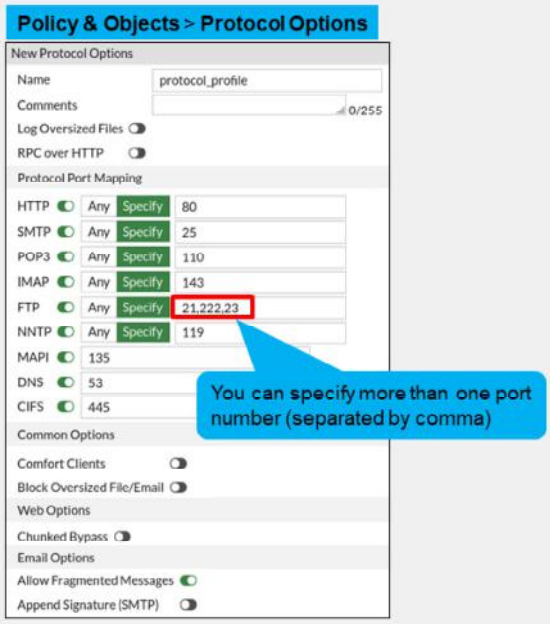

FORTINET Training Institute

22

Protocol options provide more granular control than antivirus profiles. You can configure protocol port mappings, common options, web options, and email options, to name a few.

You can configure protocol options on the **Protocol Options** page on the GUI or from the CLI. Protocol options are used by antivirus and other security profiles, such as web filtering, DNS filtering, and data loss prevention (DLP), to name a few.

Once protocol options are configured, they are applied in the firewall policy.

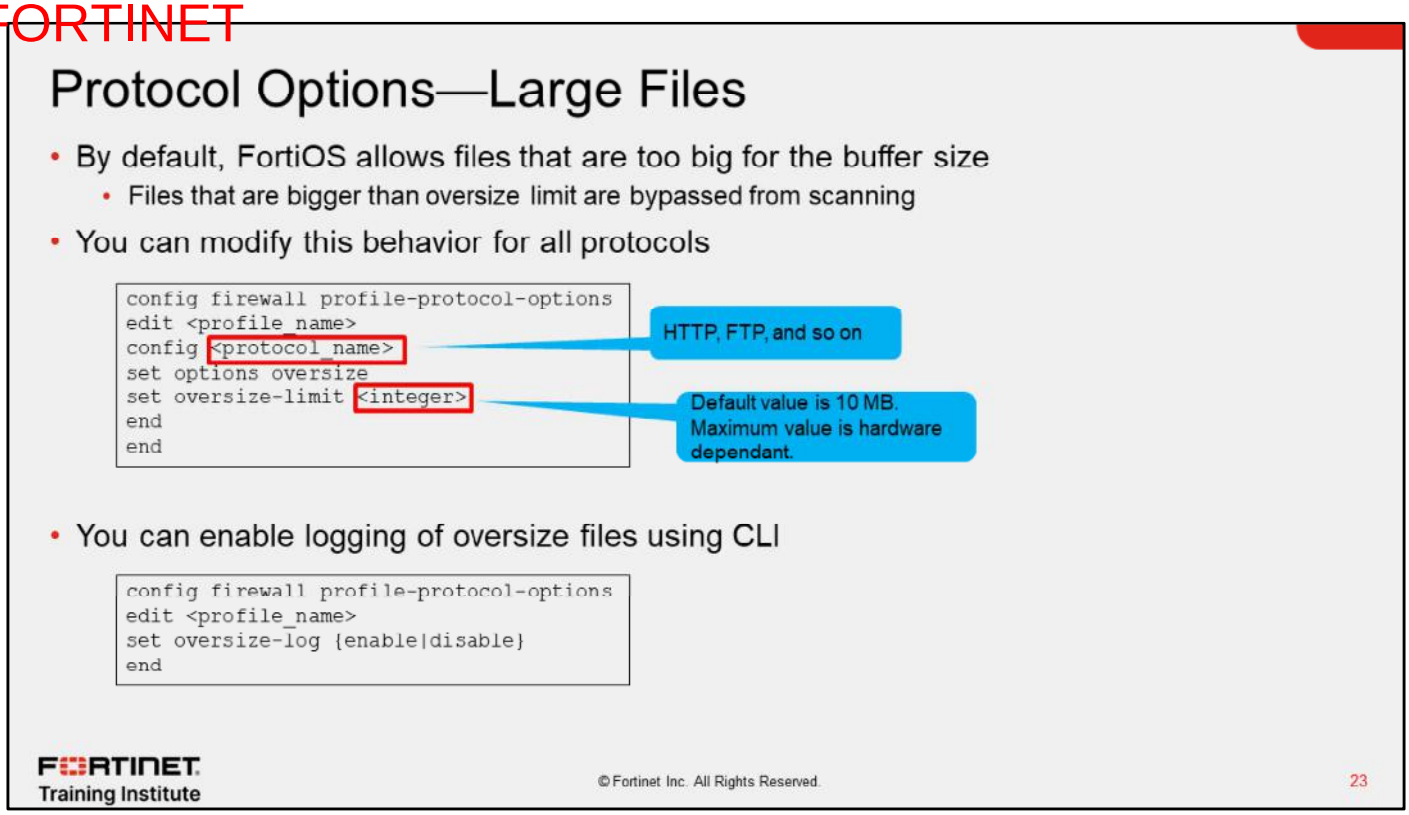

So what is the recommended buffer limit? It varies by model and configuration. You can adjust the `oversize-limit` for your network for optimal performance. A smaller buffer minimizes proxy latency (for both scanning modes) and RAM usage, but that may allow viruses to pass through undetected. When a buffer is too large, clients may notice transmission timeouts. You need to balance the two.

If you aren't sure about the value to set `oversize-limit` to, you can temporarily enable `oversize-log` to see if your FortiGate is scanning large files frequently. You can then adjust the value accordingly.

Files that are bigger than the oversize limit are bypassed from scanning. You can enable logging of oversize files by enabling the `oversize-log` option from the CLI.

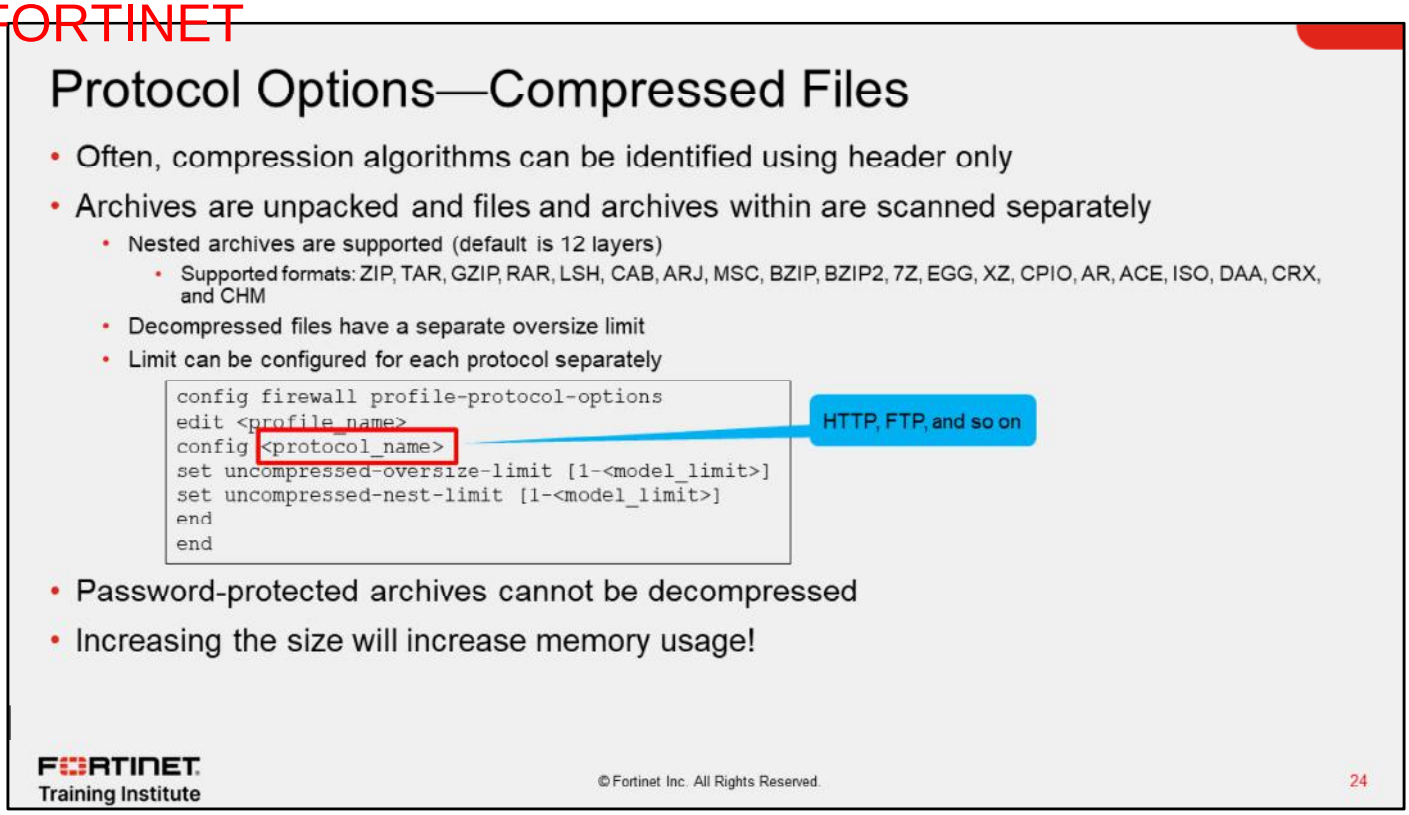

Large files are often compressed. When compressed files go through scanning, the compression acts like encryption: the signatures won't match. So, FortiGate must decompress the file in order to scan it.

Before decompressing a file, FortiGate must first identify the compression algorithm. Some archive types can be correctly identified using only the header. Also, FortiGate must check whether the file is password protected. If the archive is protected with a password, FortiGate can’t decompress it, and, therefore, can’t scan it.

FortiGate decompresses files into RAM. Just like other large files, the RAM buffer has a maximum size. Increasing this limit may decrease performance, but it allows you to scan larger compressed files.

If an archive is nested—for example, if an attacker is trying to circumvent your scans by putting a ZIP file inside the ZIP file—FortiGate will try to undo all layers of compression. By default, FortiGate will attempt to decompress and scan up to 12 layers deep, but you can configure it to scan up to the maximum number supported by your device (usually 100). Often, you shouldn’t increase this setting because it increases RAM usage.

Applying the Antivirus Profile

- Apply the antivirus profile and protocol options on the firewall policy, to scan traffic
- Ensure that **deep-inspection** is selected for the **SSL/SSH Inspection** setting—required to scan encrypted protocols

FORTINET Training Institute

25

Before FortiGate devices can start scanning traffic for malware, you need to apply the antivirus profile, the protocol options, and SSL/SSH inspection profiles on the firewall policy.

In full SSL inspection level, FortiGate terminates the SSL/TLS handshake at its own interface, before it reaches the server. When certificates and private keys are exchanged, it is with FortiGate and not the server. Next, FortiGate starts a second connection with the server.

Because traffic is unencrypted while passing between its interfaces, FortiGate can inspect the contents and look for matches with the antivirus signature database, before it re-encrypts the packet and forwards it.

For these reasons, full SSL inspection level is the only choice that allows antivirus to be effective.

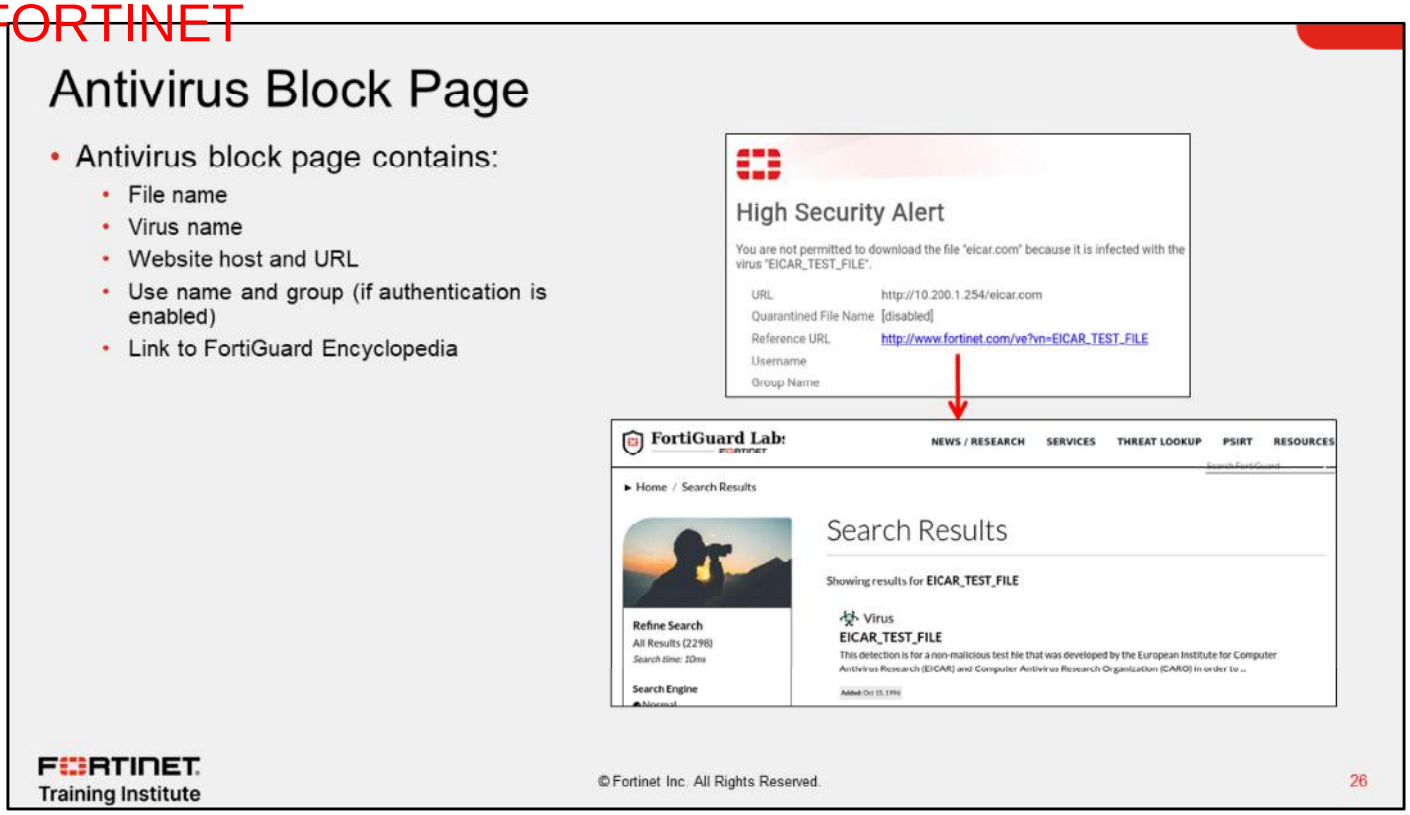

For antivirus scanning in proxy-based inspection mode (with client comforting disabled), the block replacement page is displayed *immediately* when a virus is detected.

For flow-based inspection mode scanning, if a virus is detected at the start of the stream, the block replacement page is displayed at the *first attempt*. If a virus is detected after a few packets have been transmitted, the block replacement page is *not* displayed. However, FortiGate caches the URL and can display the replacement page immediately, on the second attempt.

Note that if deep inspection is enabled, all HTTPS-based applications also display the block replacement message.

The block page includes the following:
- File name
- Virus name
- Website host and URL
- User name and group (if authentication is enabled)
- Link to FortiGuard Encyclopedia—which provides analysis, recommended actions (if any), and detection availability

You can go directly to the FortiGuard website to view information about other malware, and scan, submit, or do both, with a sample of a suspected malware.

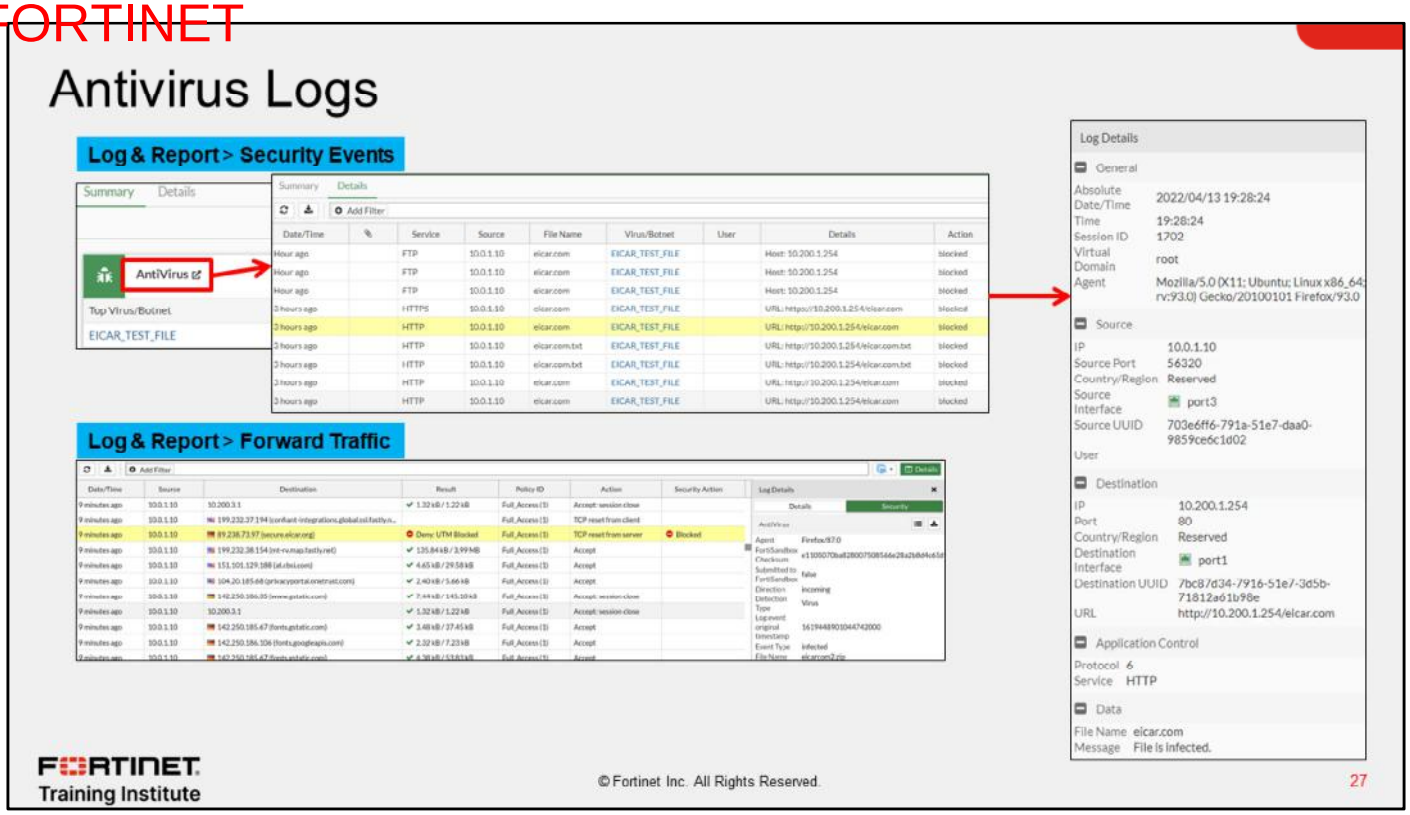

If you enable logging, you can find details on the **AntiVirus** log page under **Security Events**.

When the antivirus scan detects a virus, by default, it creates a log about what virus was detected, as well as the action, policy ID, antivirus profile name, and detection type. It also provides a link to more information on the FortiGuard website.

You can also view log details on the **Forward Traffic** log page, where firewall policies record traffic activity. You'll also find a summary of the traffic on which FortiGate applied an antivirus action.

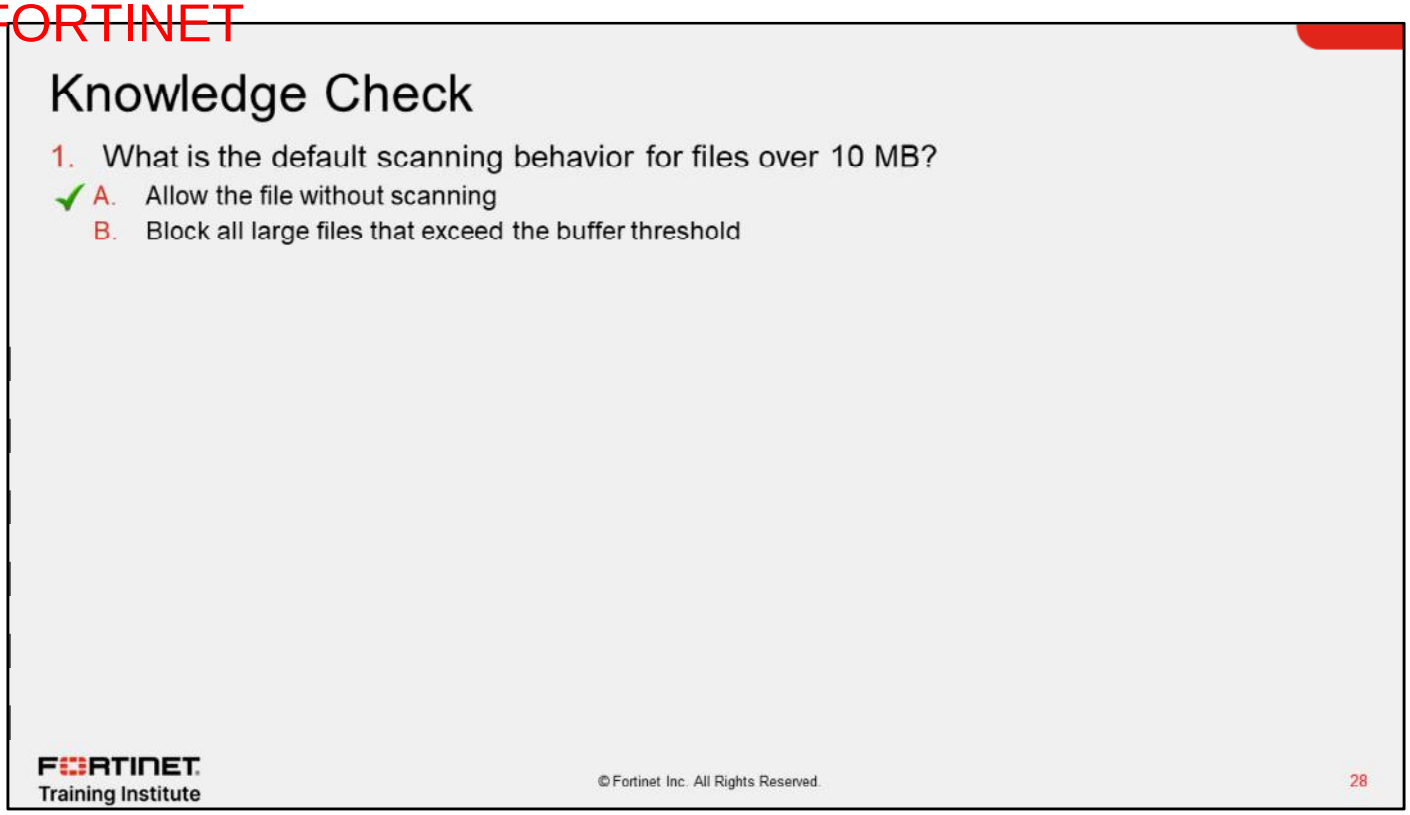

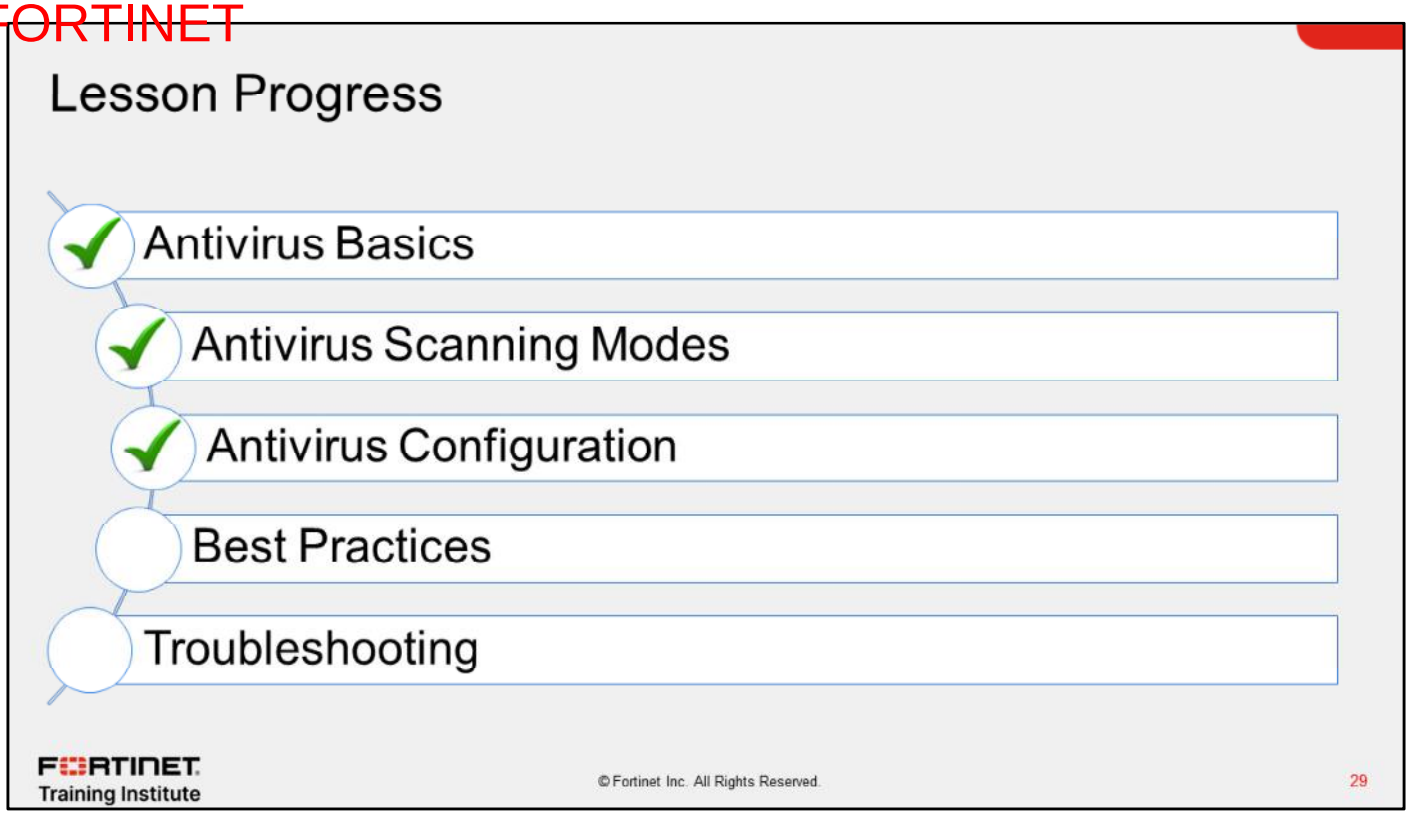

Good job! You now understand antivirus configuration.

Now, you will learn about some antivirus best practices.

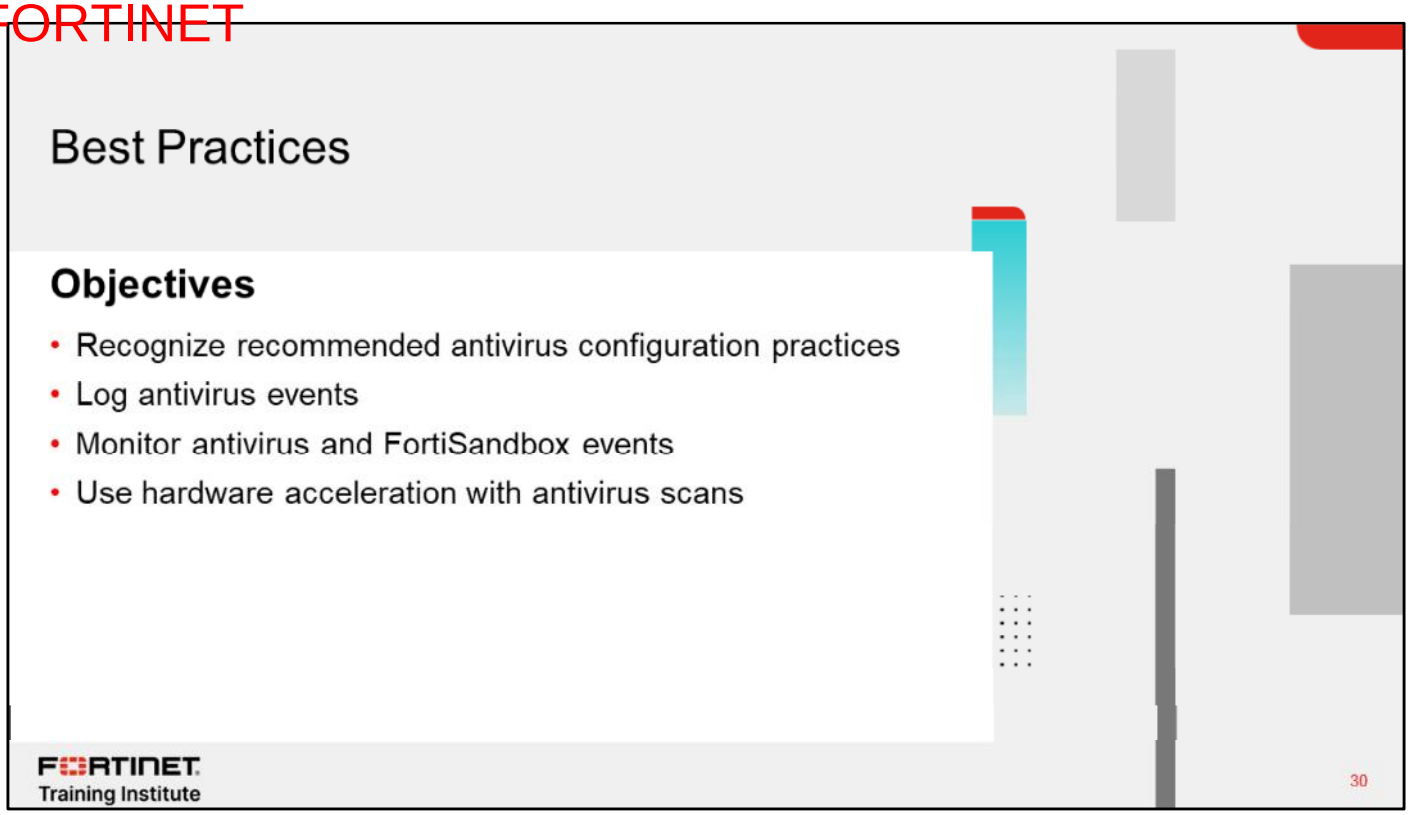

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus best practices, you will be able to configure an effective antivirus solution.

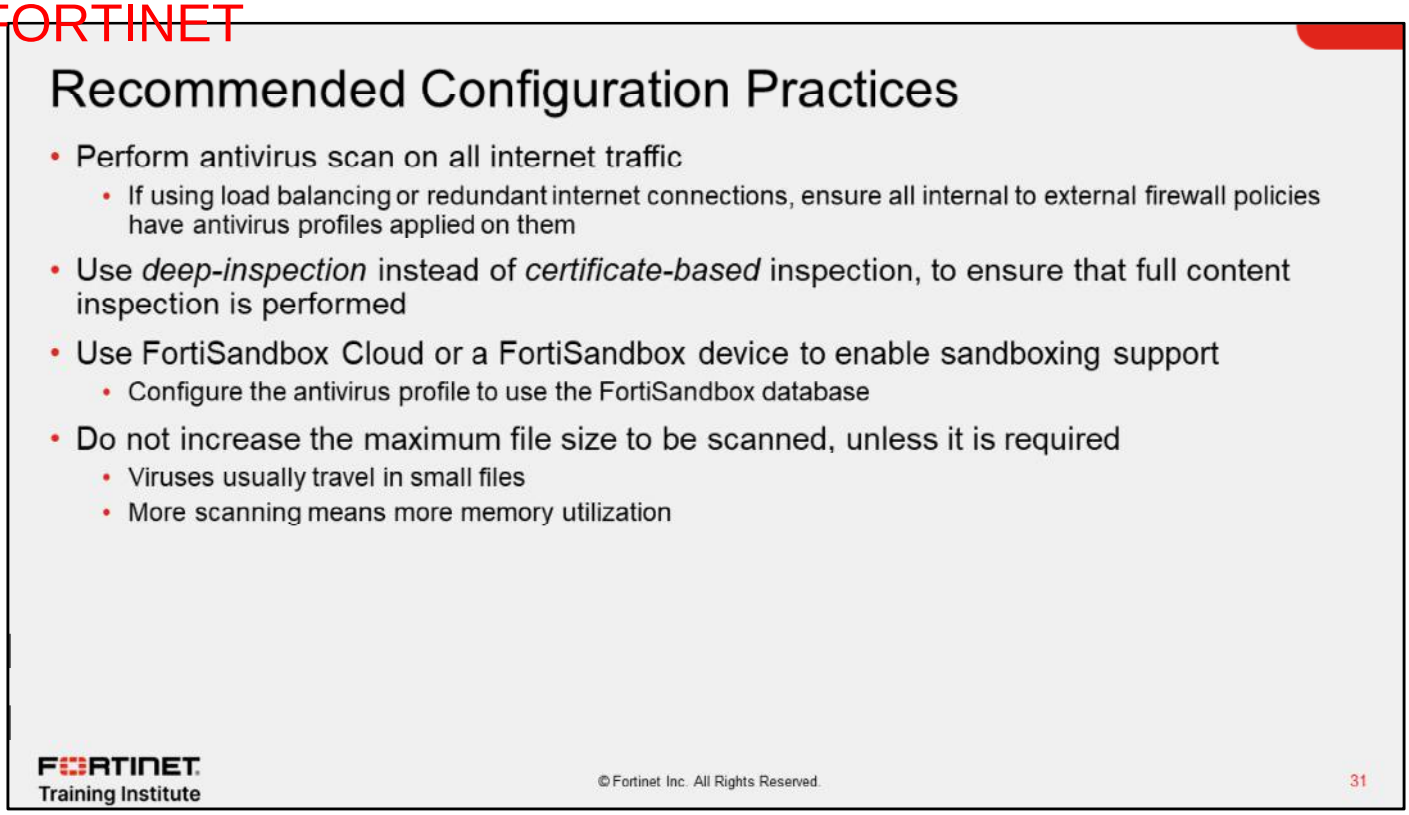

The following are some best practices to follow when configuring antivirus scanning for use on FortiOS:

- Enable antivirus scanning on all internet traffic. This includes internal to external firewall policies, and any VIP firewall policies.
- Use *deep-inspection* instead of *certificate-based* inspection, to ensure that full content inspection is performed.
- Use FortiSandbox for protection against new viruses.
- Do not increase the maximum file size to be scanned, unless there is good reason, or you need to do so in order to meet a network requirement.

Log Antivirus Events

- Enable logging of oversized files
  - This will ensure that files that are not *scanned* are *logged*
- Ensure that firewall policies with antivirus applied have security events logging enabled
- Use standalone dashboard to monitor threats to your network
  - Dashboard organizes threats based on network segments on the device

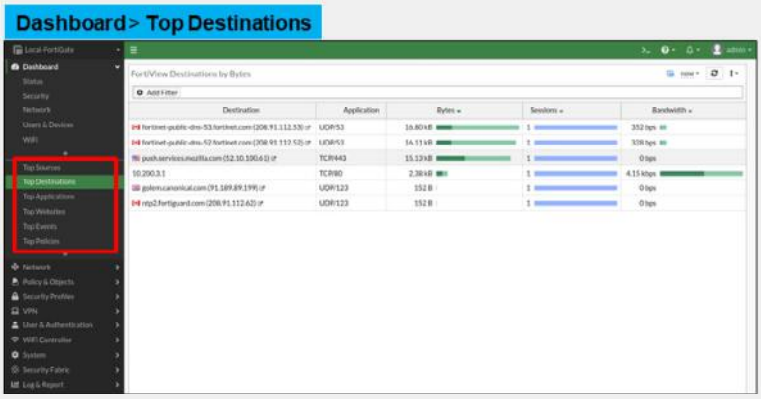

FORTINET Training Institute

32

Logging is an important part of managing a secure network. Enable logging for oversized files so that if there are files that are not scanned, you can be aware of it. Also, ensure that security events logging is enabled on all firewall policies using security profiles. Use the standalone dashboards to view relevant information regarding threats to your network. The standalone dashboard organizes information into network segments and breaks it down into various categories.

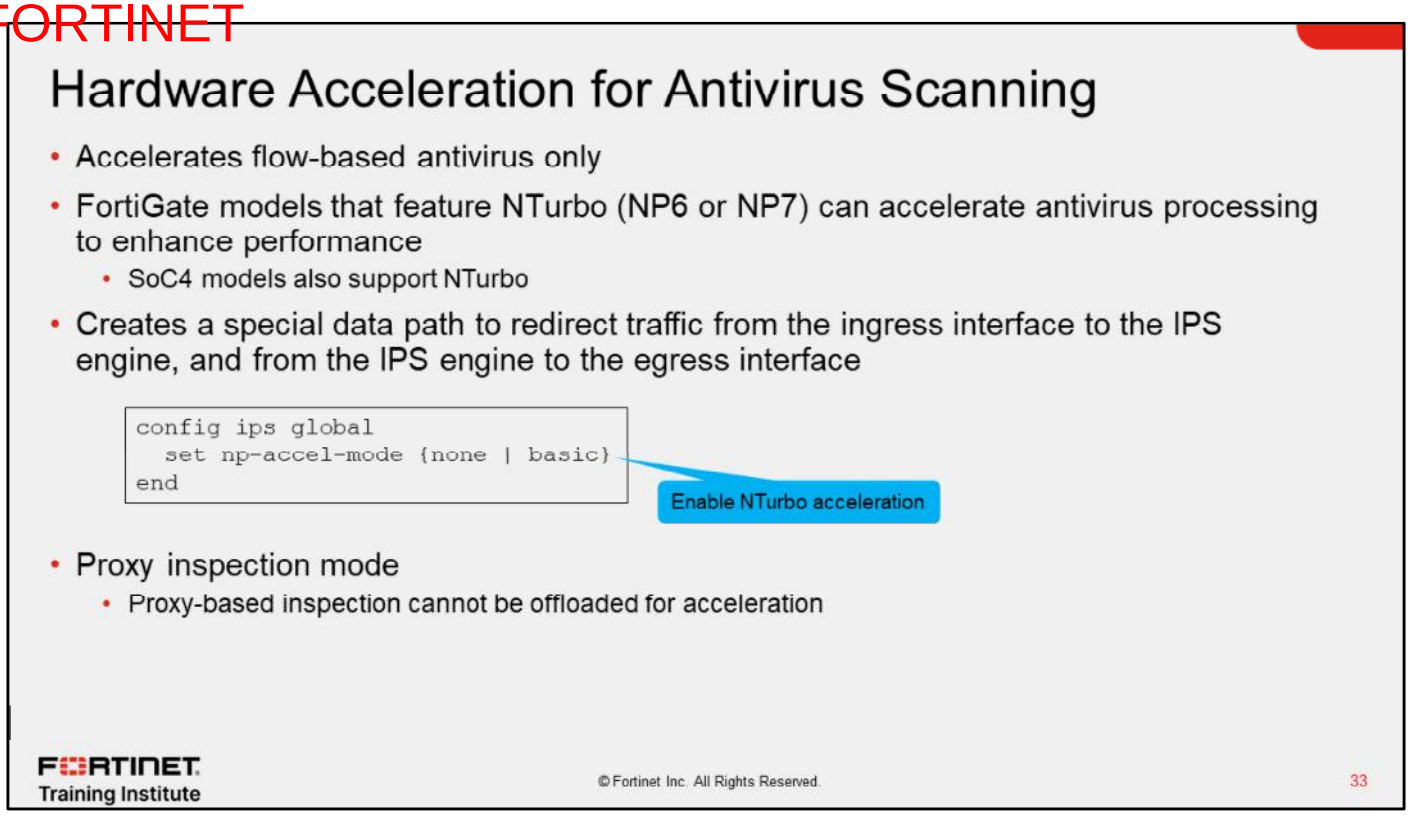

The FortiGate main CPU is responsible for performing UTM/NGFW inspection on the network traffic. FortiGate models that have specialized chips can offload inspection tasks to enhance performance while providing the same level of protection. FortiGate devices that support the NTurbo feature can offload UTM/NGFW sessions to network processors. NTurbo creates a special data path to redirect traffic from the ingress interface to the IPS engine, and from the IPS engine to the egress interface. This can improve performance by accelerating antivirus inspection, without sacrificing security.

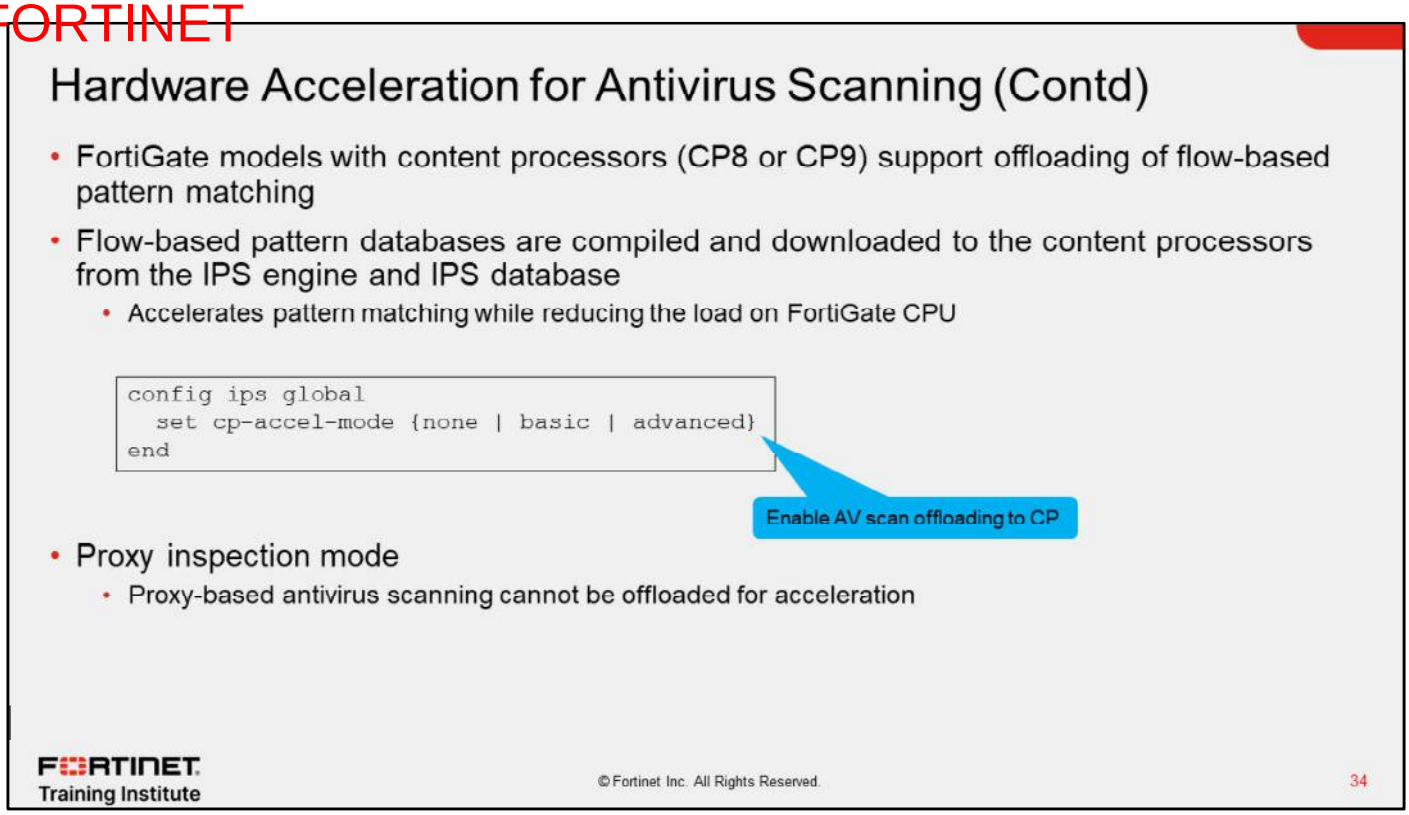

FortiGate models that have CP8 or CP9 content processors can offload flow-based pattern matching to CP8 or CP9 processors. When CP acceleration is enabled, flow-based pattern databases are compiled and downloaded to the content processors from the IPS engine and IPS database. This reduces load on the FortiGate CPU because flow-based pattern matching requests are redirected to the CP hardware. Before flow-based inspection is applied to the traffic, the IPS engine uses a series of decoders to determine the appropriate security modules that can be used, depending on the protocol of the packet and policy settings. In addition, if SSL inspection is configured, the IPS engine also decrypts SSL packets. SSL decryption is also offloaded and accelerated by CP8 or CP9 processors.

Knowledge Check

1. Which type of inspection mode can be offloaded using NTurbo hardware acceleration?
   - A. Proxy-based
   - ✓ B. Flow-based

2. What does the logging of oversized files option do?
   - ✓ A. Enables logging of all files that cannot be scanned because of oversize limit
   - B. Logs all files that are over 5 MB

FORTINET Training Institute

35

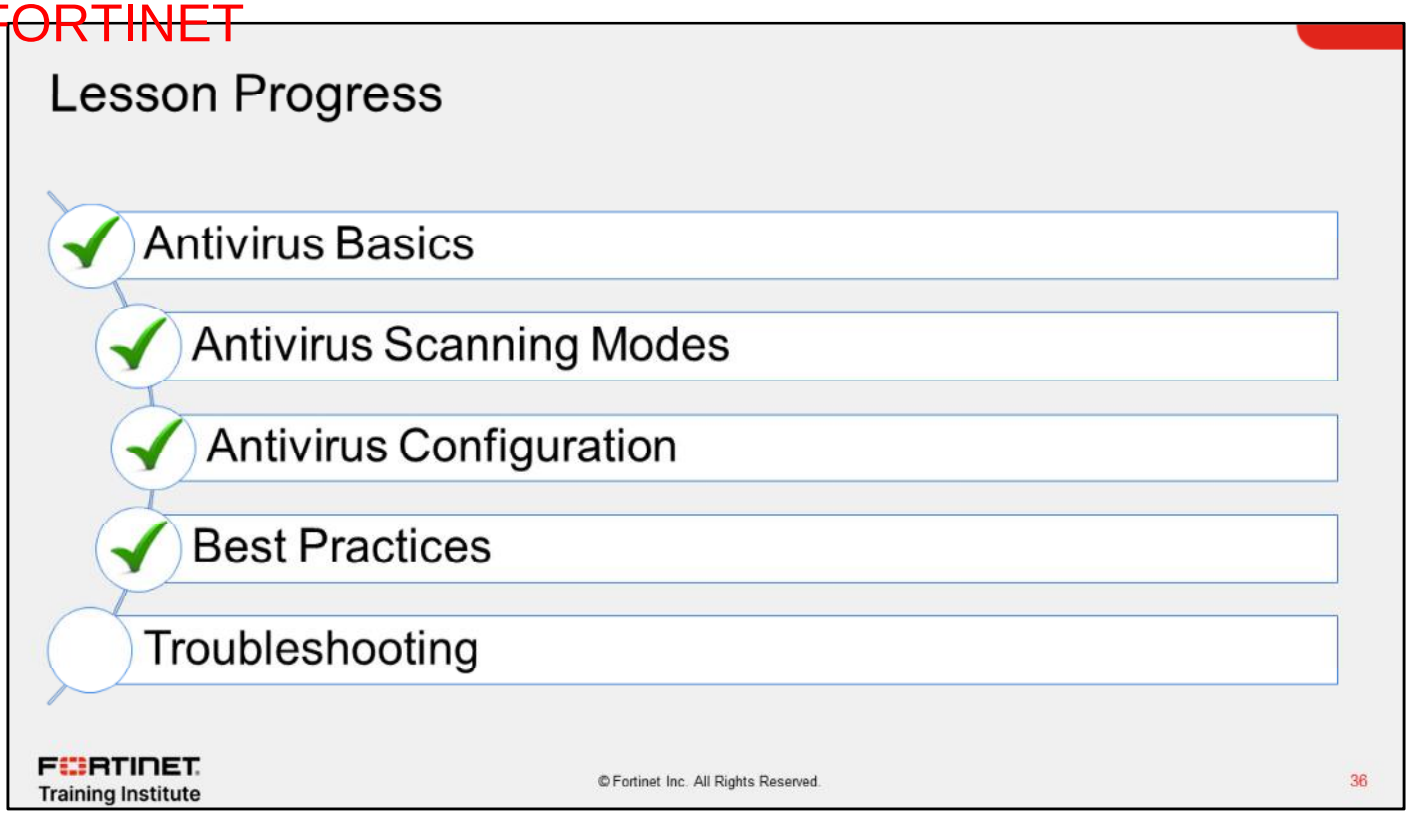

Good job! You now understand antivirus best practices.

Now, you will learn about antivirus troubleshooting.

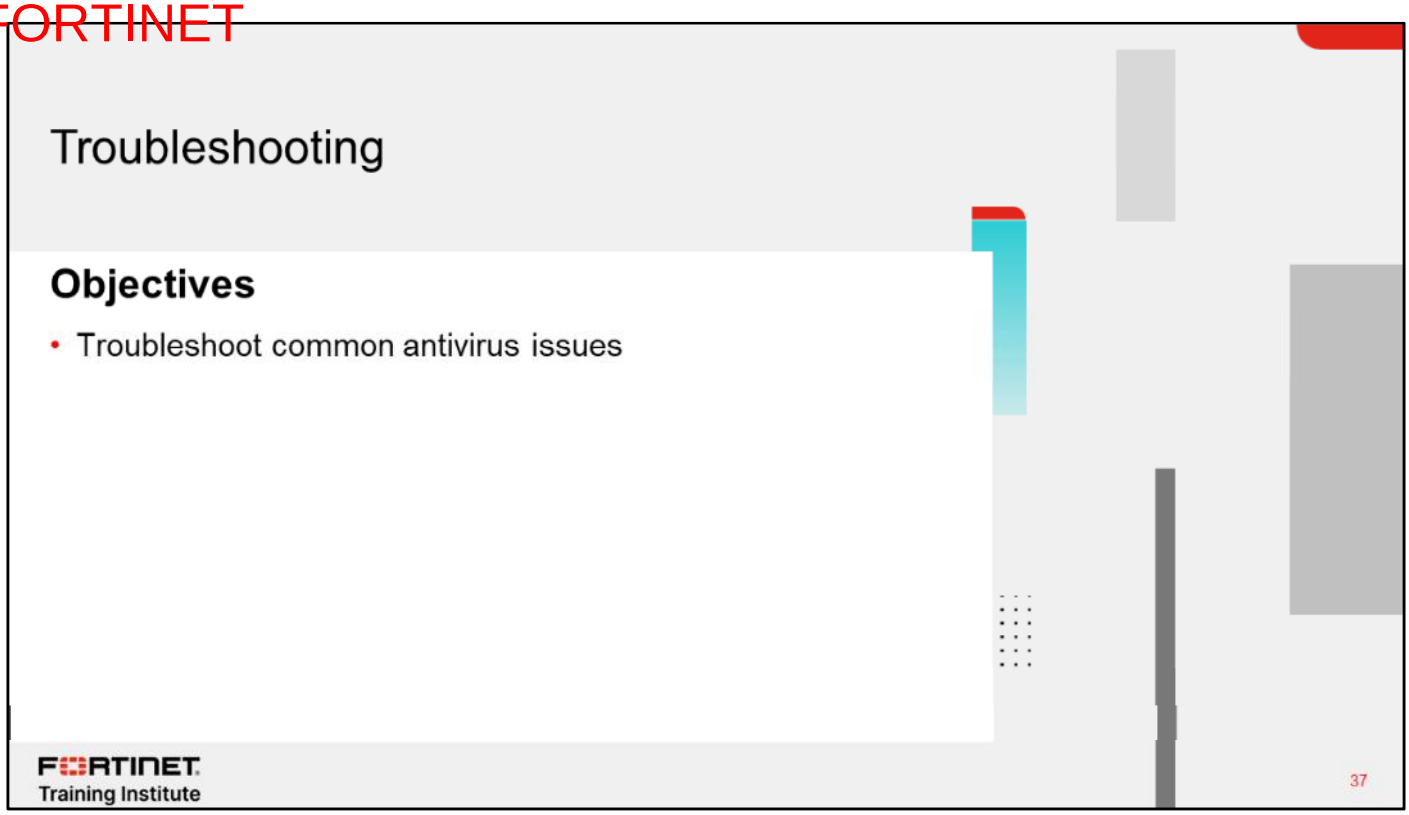

After completing this section, you should be able to troubleshoot common issues with antivirus.

By demonstrating competence in troubleshooting common antivirus issues, you will be able to configure and maintain an effective antivirus solution.

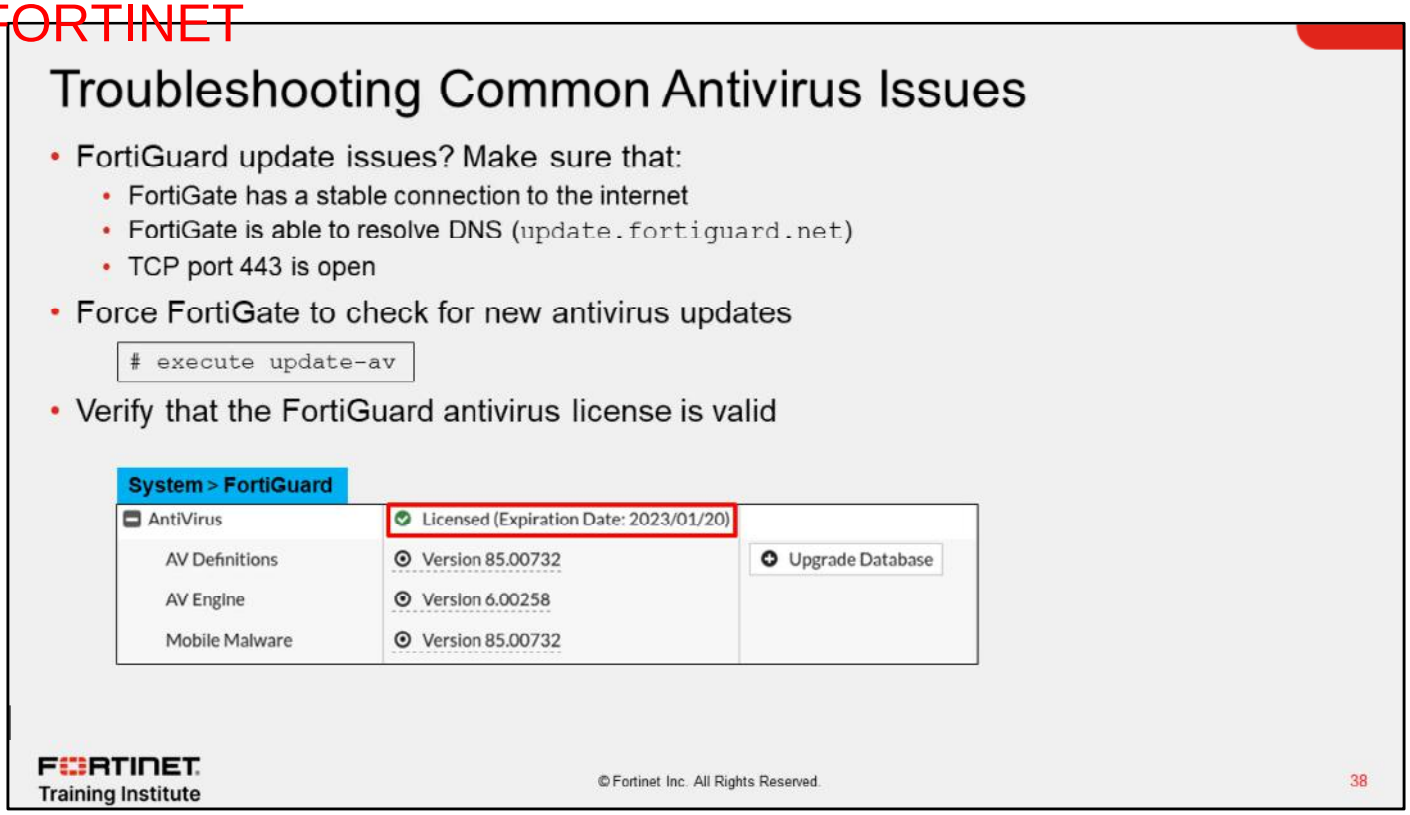

If you are having issues with the antivirus license or FortiGuard updates, start troubleshooting with basic connectivity tests. Most of the time, issues related to updates are caused by connectivity problems with FortiGuard servers. You can perform the following to handle common antivirus issues:

- Make sure that FortiGate has a stable internet connection and can resolve DNS (`update.fortinet.net`).
- If there is another firewall between FortiGate and the internet, make sure TCP port 443 is open and traffic is allowed from and to the FortiGate device.
- Force FortiGate to check for new virus updates using the CLI command: `execute update-av`.
- Verify that the FortiGate device is registered and has a valid antivirus service contract.

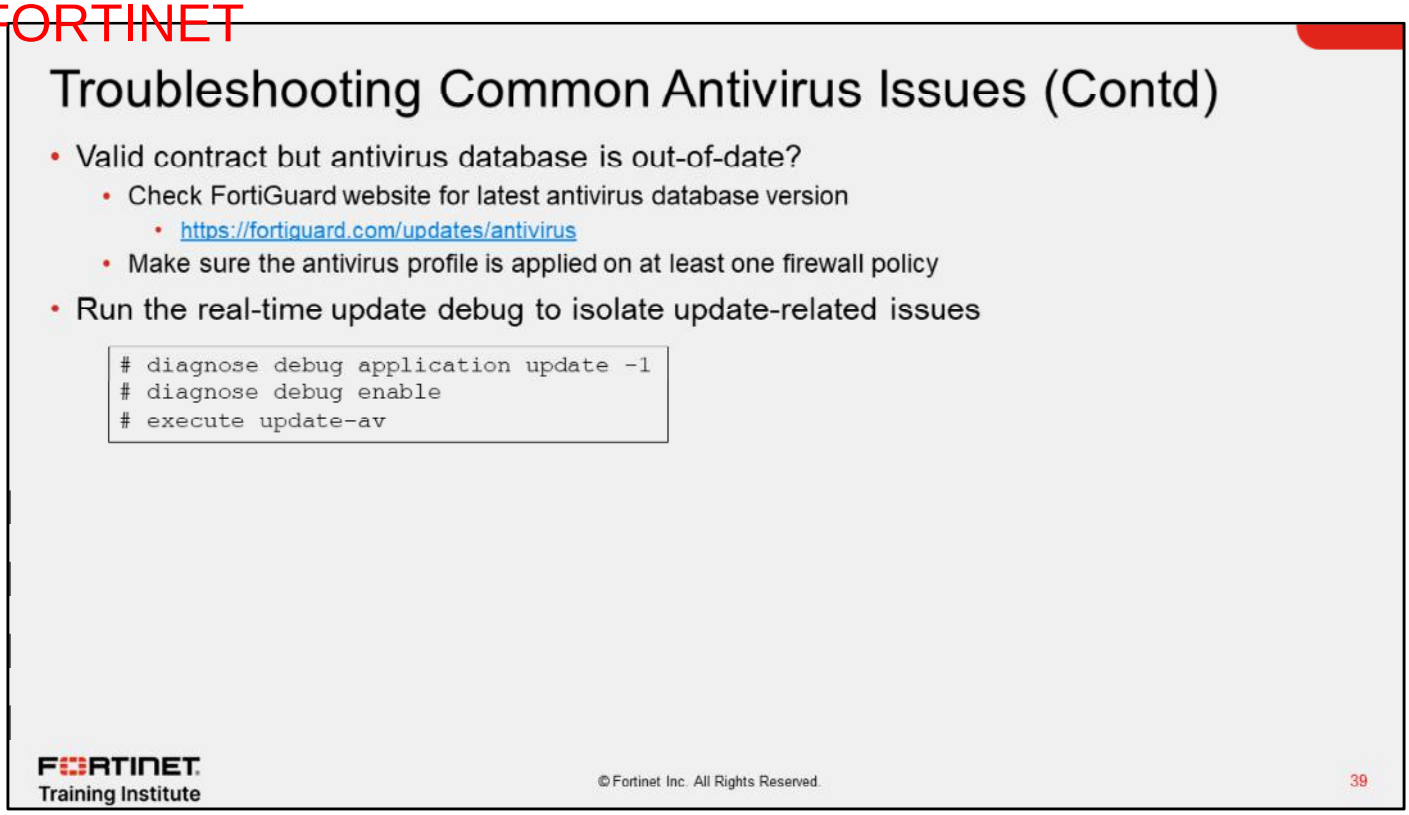

What if FortiGate shows a valid license but the antivirus database is out-of-date?

Check the current database version installed on your FortiGate and compare the version number with the current release on the FortiGuard website. FortiGate may not update the antivirus database if it is not being used (applied on a firewall policy). Make sure the antivirus profile is applied on at least one firewall policy. If you continue to see issues with the update, run the real-time debug command to identify the problem.

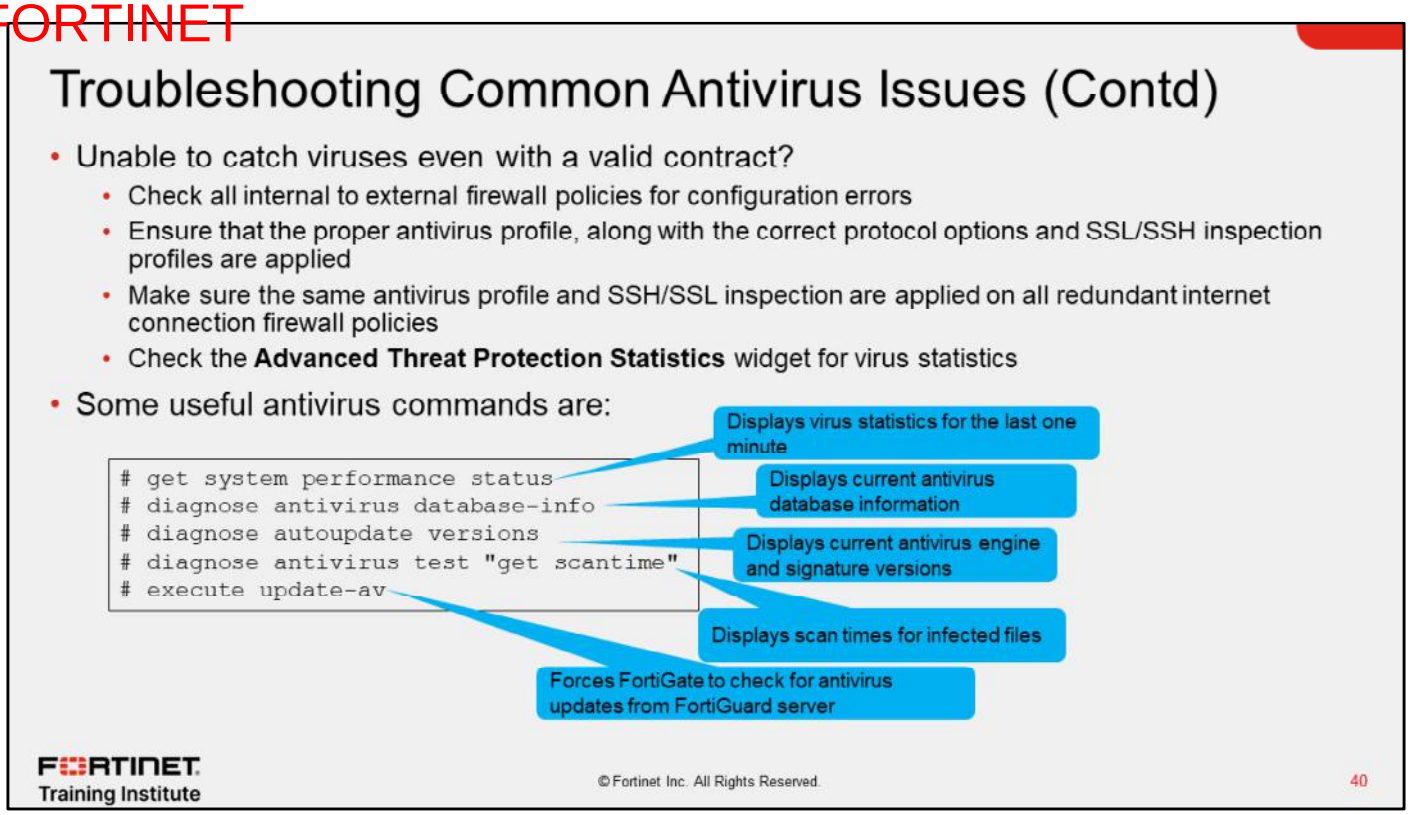

What if you have a valid contract and updated database, and you are still having issues catching viruses? Start troubleshooting for basic configuration errors. Most of the time, issues are caused by misconfiguration on the device. You can verify them as following:

- Make sure that the correct antivirus profile is applied on the right firewall policy.
- Make sure that you are using the same antivirus profile and SSL/SSH inspection on all internet connection firewall policies.
- Add and use advanced the threat protection statistics widget to get the latest virus statistics from the unit.

These are some of the commands that can be used to retrieve information and troubleshoot antivirus issues:

- `get system performance status`: Displays statistics for the last one minute.
- `diagnose antivirus database-info`: Displays current antivirus database information.
- `diagnose autoupdate versions`: Displays current antivirus engine and signature versions.
- `diagnose antivirus test "get scantime"`: Displays scan times for infected files.
- `execute update-av`: Forces FortiGate to check for antivirus updates from the FortiGuard server.

Knowledge Check

1. What command do you use to force FortiGate to check for new antivirus updates?
   A. `execute update antivirus`
   ✓ B. `execute update-av`

FORTINET Training Institute

41

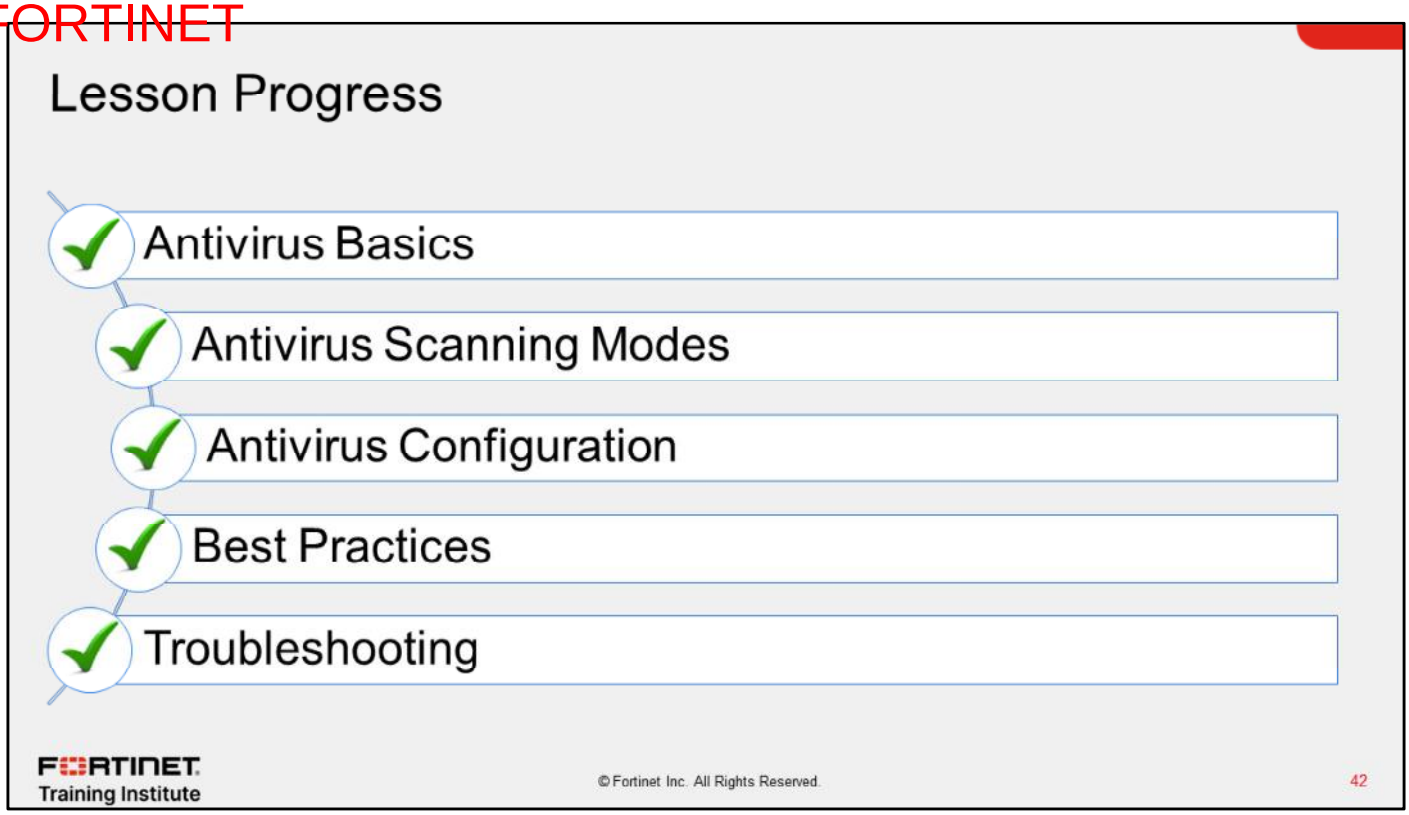

Congratulations! You have completed this lesson.

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

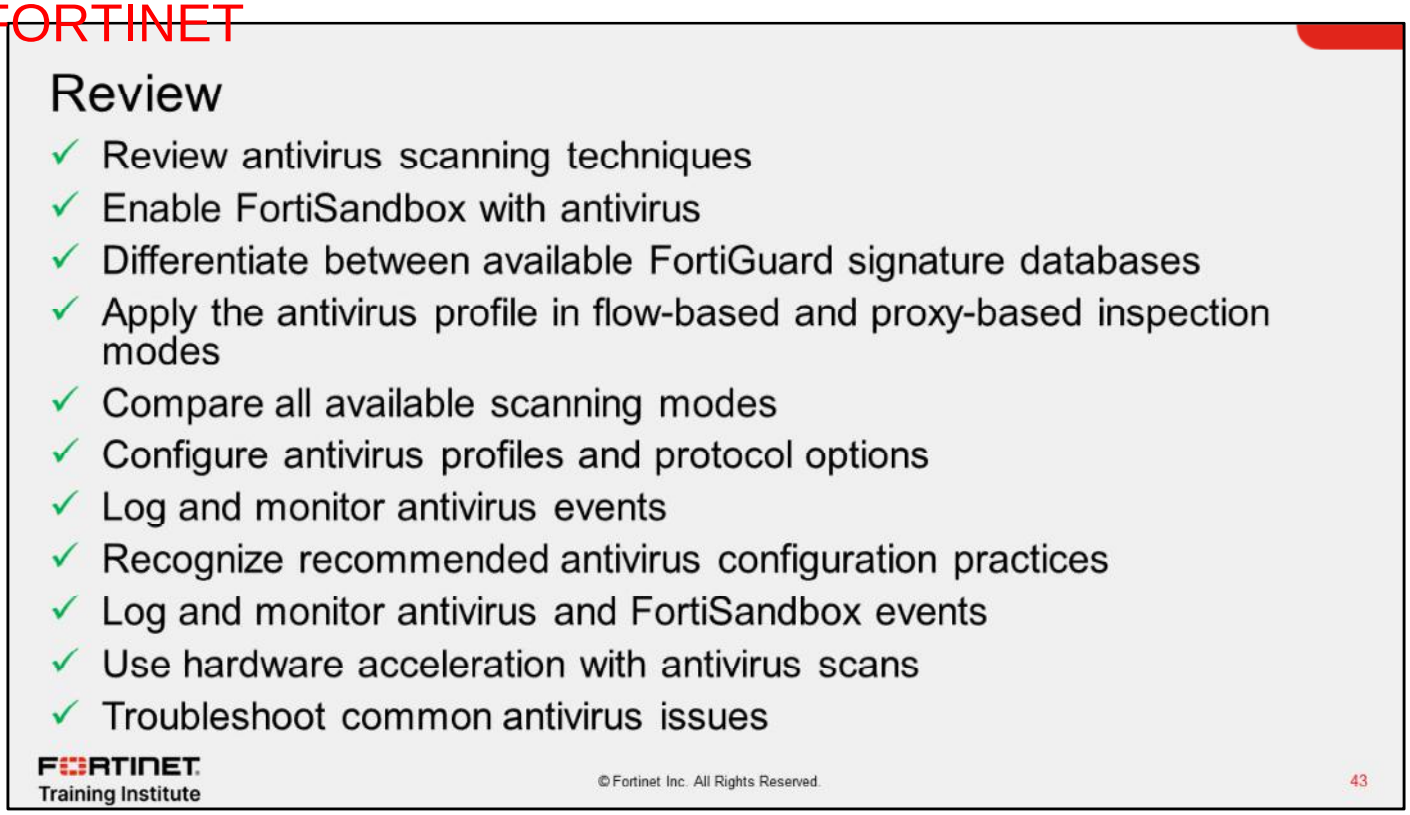

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to use FortiGate features and functions to protect your network against viruses.

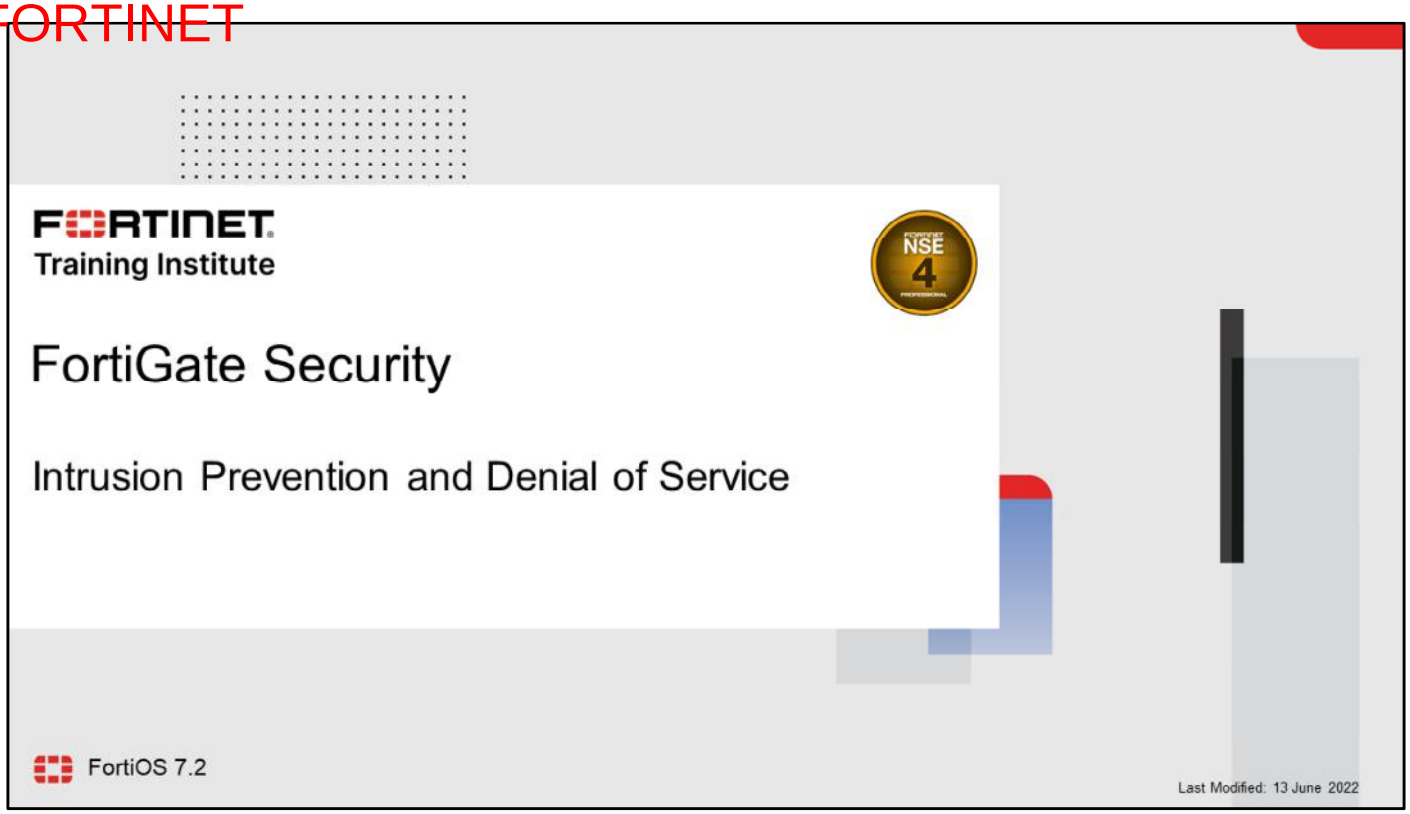

In this lesson, you will learn how to use FortiGate to protect your network against intrusions and denial of service (DoS) attacks.

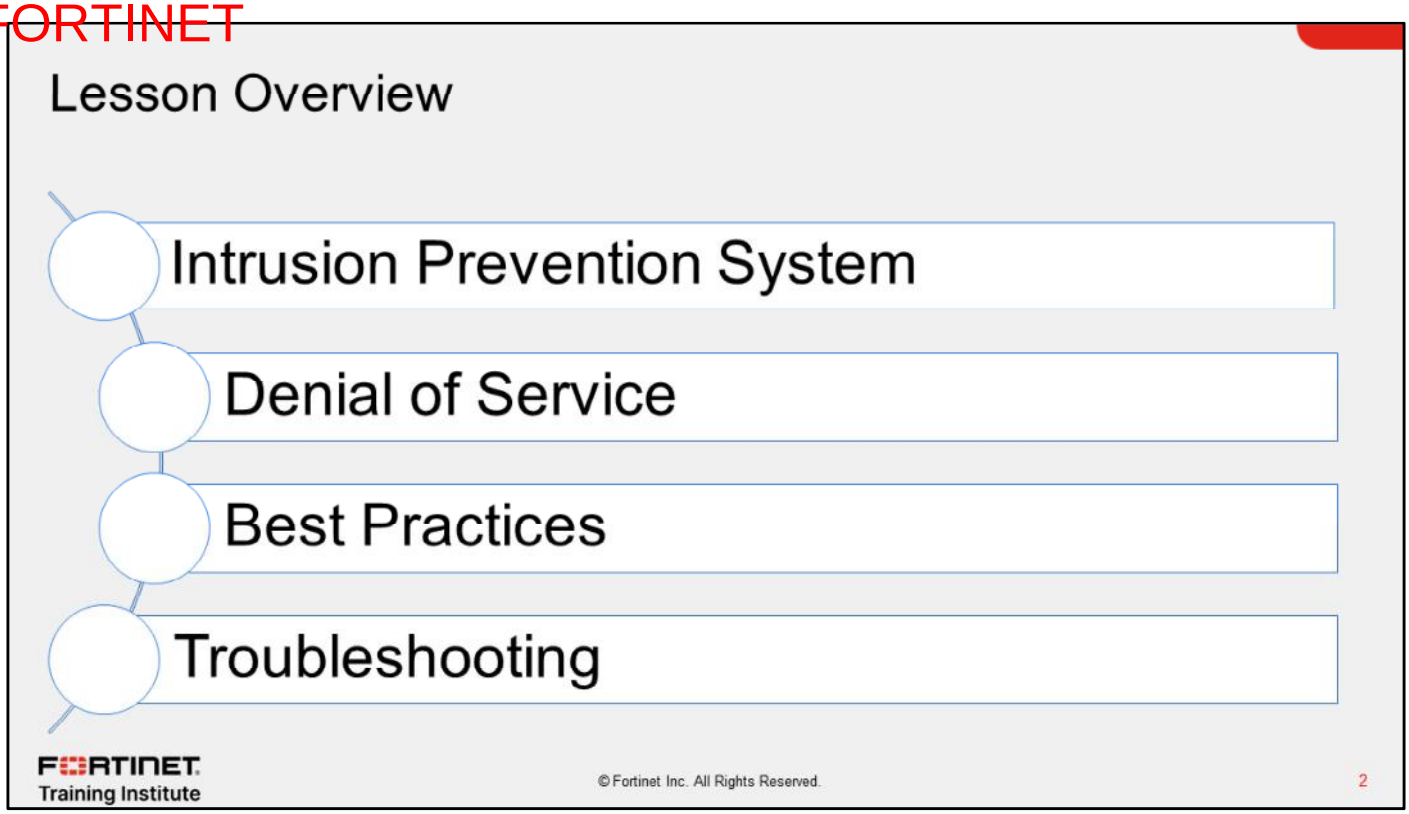

In this lesson, you will learn about the topics shown on this slide.

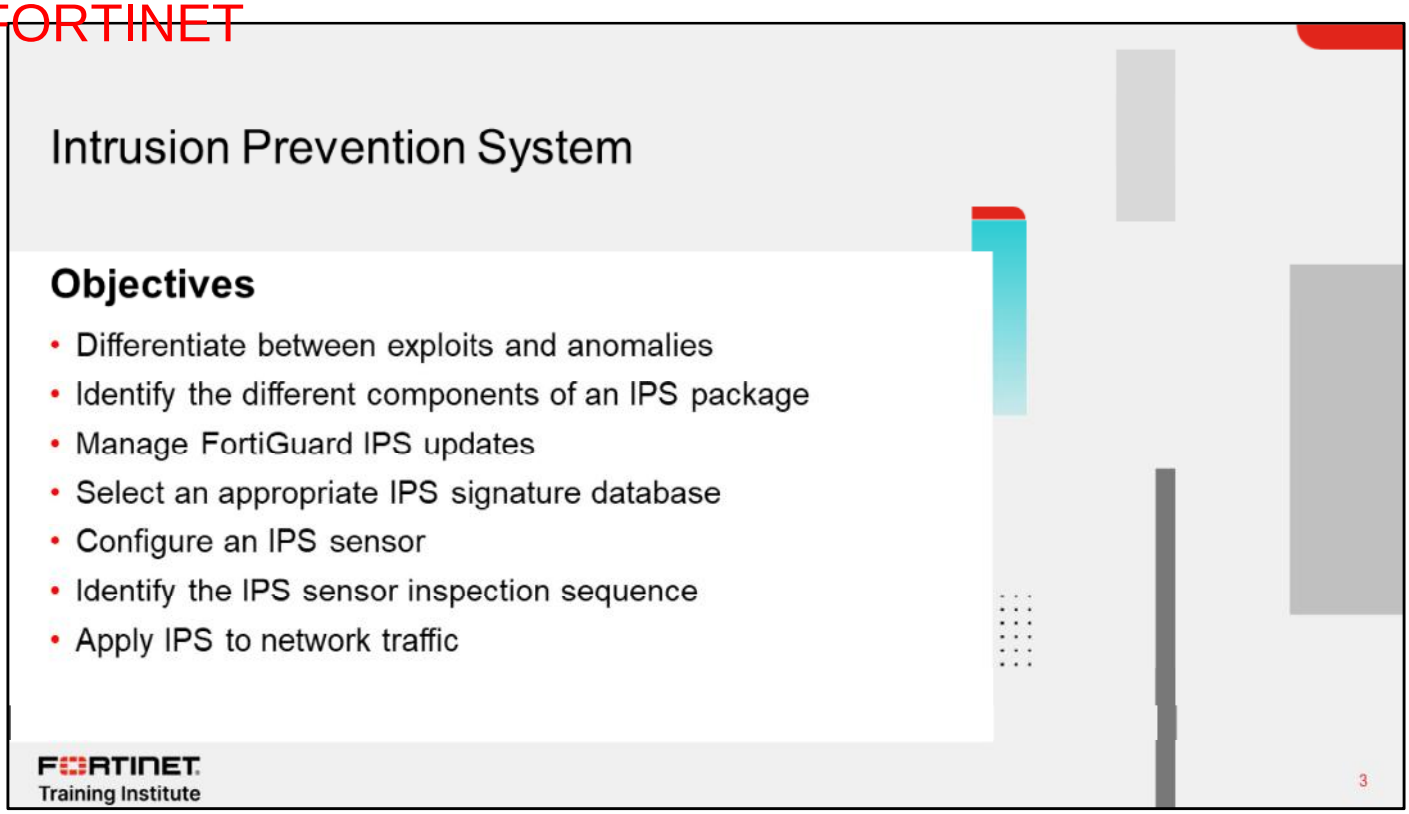

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in intrusion prevention systems (IPS), you should be able to implement an effective IPS solution to protect your network from intrusion.

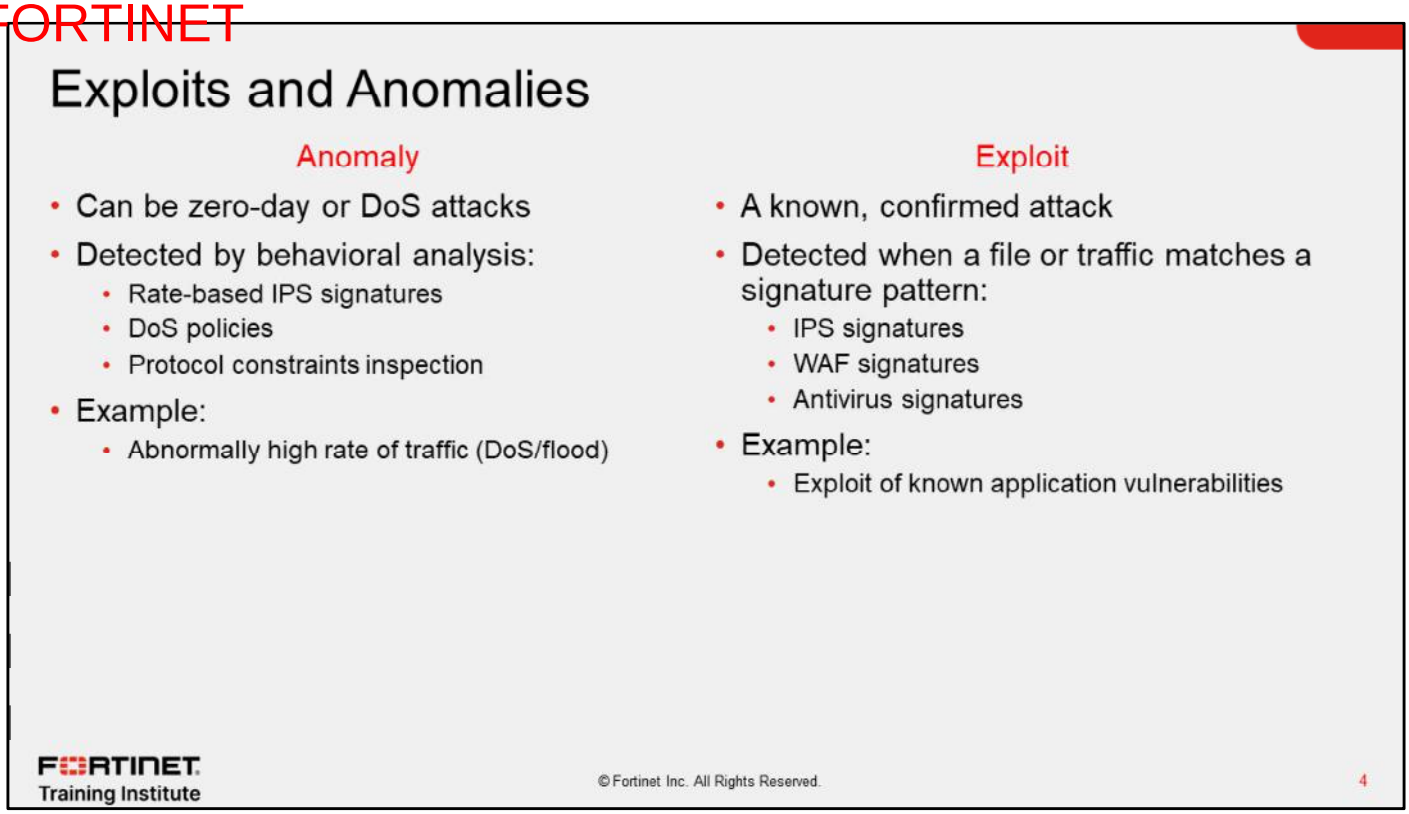

It's important to understand the difference between an anomaly and an exploit. It's also important to know which FortiGate features offer protection against each of these types of threats.

Exploits are known attacks, with known patterns that can be matched by IPS, web application firewall (WAF), or antivirus signatures.

Anomalies are unusual behaviors in the network, such as higher-than-usual CPU usage or network traffic. Anomalies must be detected and monitored (and, in some cases, blocked or mitigated) because they can be the symptoms of a new, never-seen-before attack. Anomalies are usually better detected by behavioral analysis, such as rate-based IPS signatures, DoS policies, and protocol constraints inspection.

IPS

- Flow-based detection and blocking
  - Known exploits that match signatures
  - Network errors and protocol anomalies
- IPS components
  - IPS signature databases
  - Protocol decoders
  - IPS engine
    - Application control
    - Antivirus (flow-based)
    - Web filter (flow-based)
    - Email filter (flow-based)

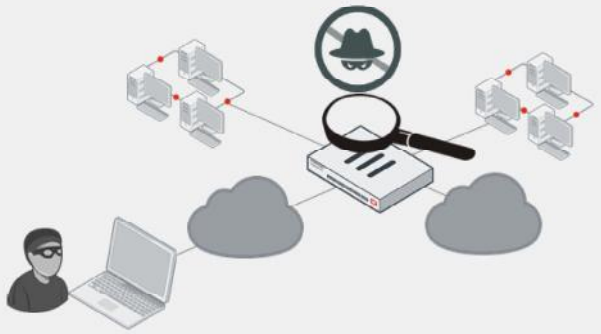

FORTINET Training Institute

5

IPS on FortiGate uses signature databases to detect known attacks. Protocol decoders can also detect network errors and protocol anomalies.

The IPS engine is responsible for most of the features shown in this lesson: IPS and protocol decoders. It's also responsible for application control, flow-based antivirus protection, web filtering, and email filtering.

What Are Protocol Decoders?

- Decoders parse protocols
- IPS signatures find parts of a protocol that don't conform
  - For example, too many HTTP headers, or a buffer overflow attempt
- Unlike proxy-based scans, IPS often does not require IANA standard ports
  - Automatically selects decoder for protocol at each OSI layer

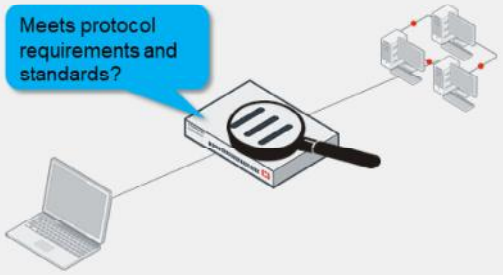

FORTINET Training Institute

6

How does the IPS engine determine if a packet contains an attack or anomaly?

Protocol decoders parse each packet according to the protocol specifications. Some protocol decoders require a port number specification (configured on the CLI), but usually, the protocol is automatically detected. If the traffic doesn't conform to the specification—if, for example, it sends malformed or invalid commands to your servers—then the protocol decoder detects the error.

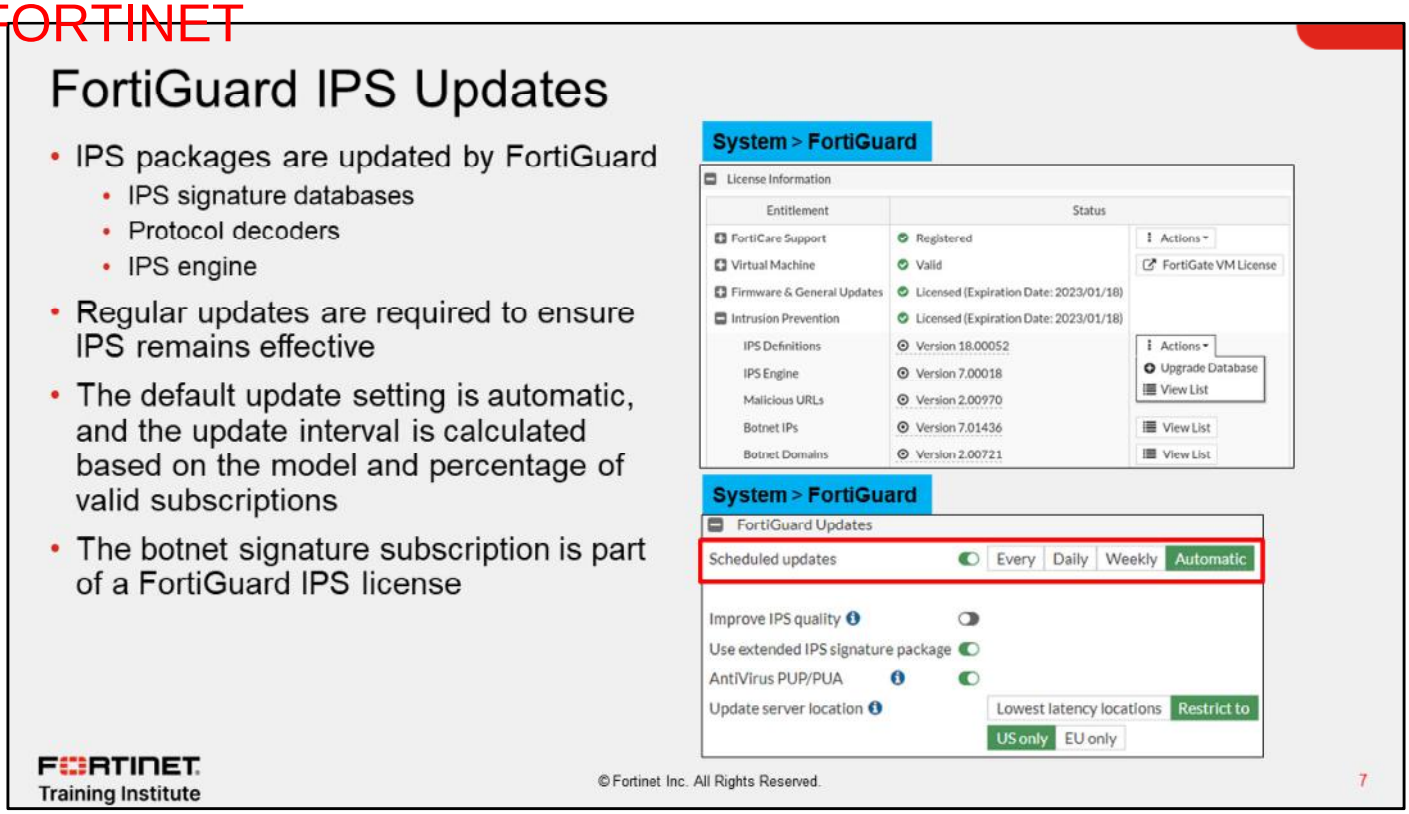

By default, an initial set of IPS signatures is included in each FortiGate firmware release. FortiGuard updates the IPS signature database with new signatures. That way, IPS remains effective against new exploits. Unless a protocol specification or RFC changes (which doesn't happen very often), protocol decoders are rarely updated. The IPS engine itself changes more frequently, but still not often.

The FortiGuard IPS service updates the IPS signatures most often. The FortiGuard research team identifies and builds new signatures, just like antivirus signatures. So, if your FortiGuard Services contract expires, you can still use IPS. However, just like antivirus scans, IPS scans become increasingly ineffective the longer the signatures are not updated—old signatures won't defend against new attacks.

The default auto-update schedule for FortiGuard packages has been updated. Previously, the frequency was a reoccurring random interval within two hours. Starting in FortiOS 7.0, the frequency is automatic, and the update interval is calculated based on the model and percentage of valid subscriptions. The update interval is within one hour.

For example, an FG-501E has 78% valid contracts. Based on this device model, FortiOS calculates the update schedule to be every 10 minutes. You can verify the system event logs, which are generated approximately every 10 minutes.

IPS is a FortiGuard subscription, and includes a botnet signature database. The botnet IP database is part of the ISDB updates. The botnet domains database is part of the AV updates, and only the botnet signatures require the FortiGuard IPS license subscription.

Choosing the Signature Database

- Regular
  - Common attacks with fast, certain identification (default action is block)
- Extended
  - Performance intensive

FORTINET Training Institute

8

The IPS signature database is divided into the regular and extended databases. The regular signature database contains signatures for common attacks whose signatures cause rare or no false positives. It's a smaller database, and its default action is to block the detected attack.

The extended signature database contains additional signatures for attacks that cause a significant performance impact, or don't support blocking because of their nature. In fact, because of its size, the extended database is not available for FortiGate models with a smaller disk or RAM. But, for high-security networks, you might be required to enable the extended signatures database.

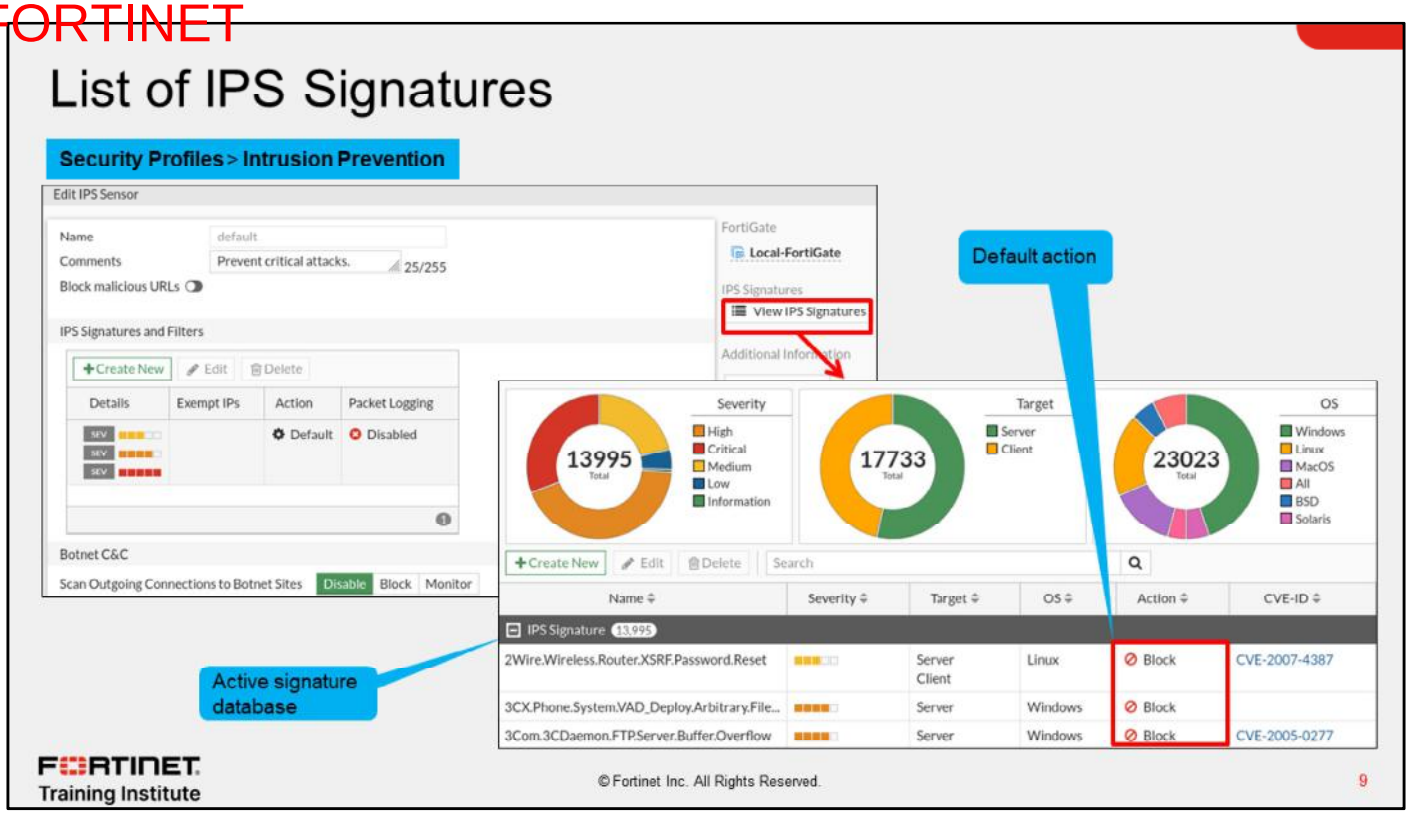

After FortiGate downloads a FortiGuard IPS package, new signatures appear in the signature list. When configuring FortiGate, you can change the **Action** setting for each sensor that uses a signature.

The default action setting is often correct, except in the following cases:

- Your software vendor releases a security patch. Continuing to scan for exploits wastes FortiGate resources.
- Your network has a custom application with traffic that inadvertently triggers an IPS signature. You can disable the setting until you notify Fortinet so that the FortiGuard team can modify the signature to avoid false positives.

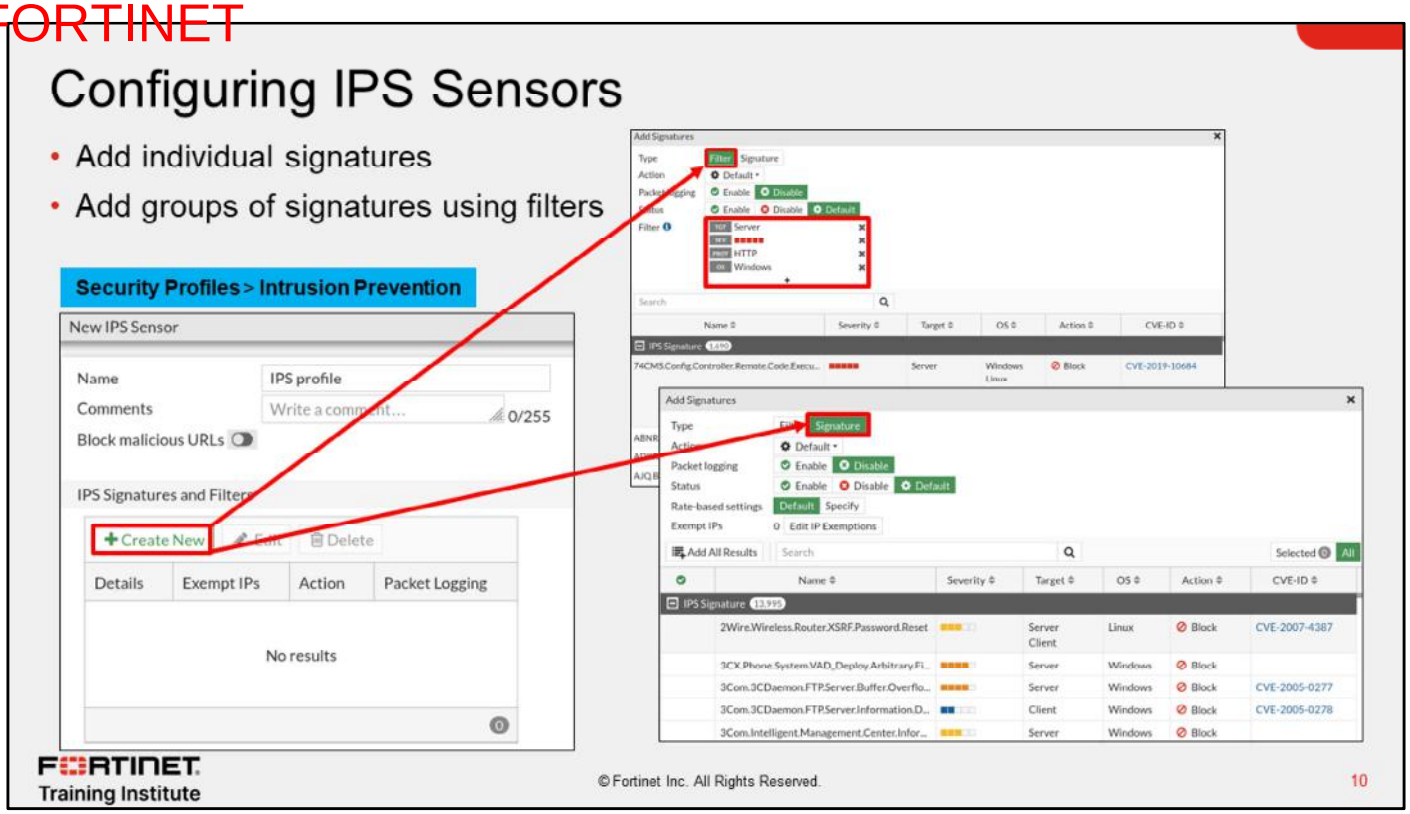

There are two ways to add predefined signatures to an IPS sensor. One way is to select the signatures individually. After you select a signature in the list, the signature is added to the sensor with its default action. Then, you can right-click the signature and change the action.

The second way to add a signature to a sensor is using filters. FortiGate adds all the signatures that match the filters.

The purpose of the IPS feature is to protect the inside of the network from outside threats.

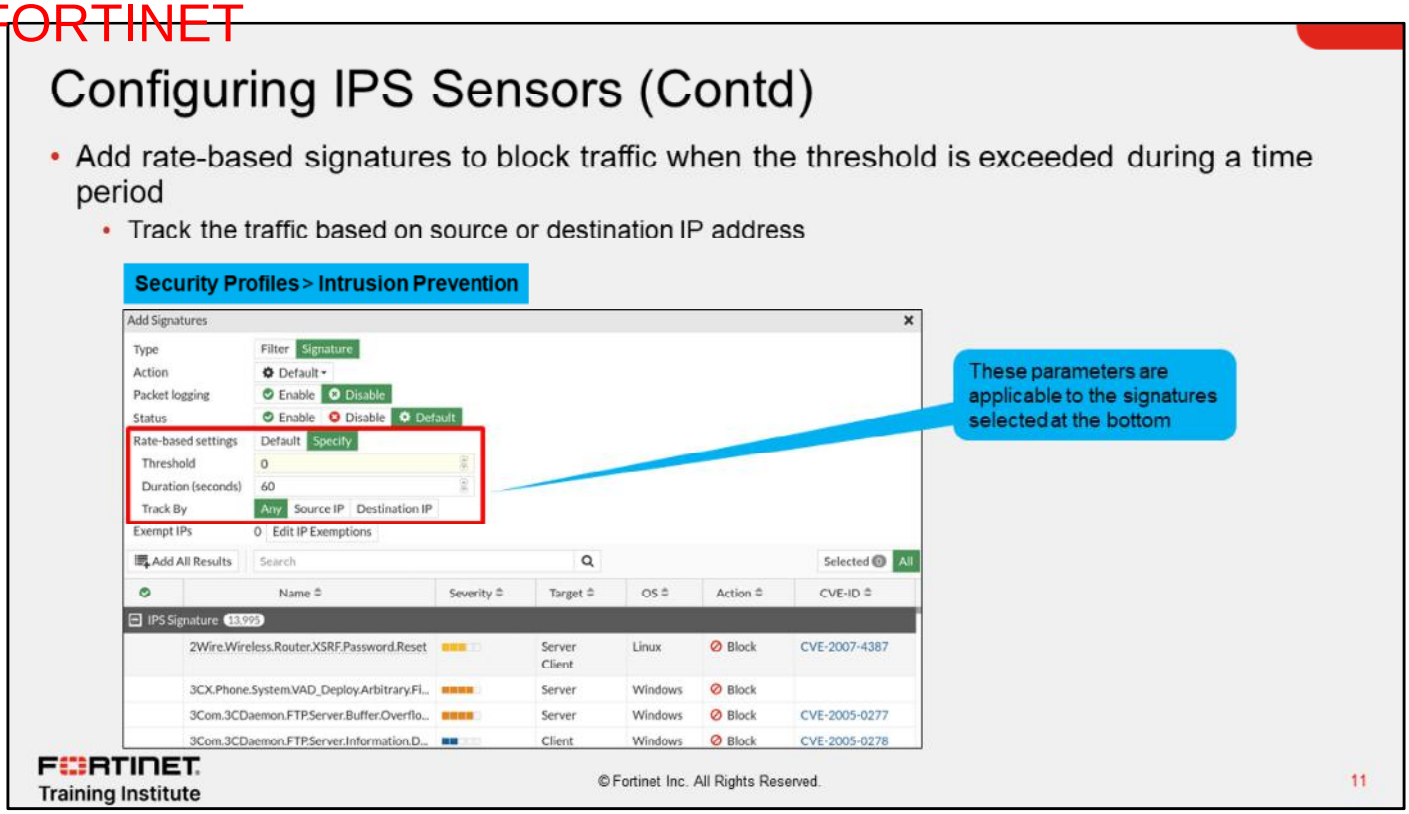

You can also add rate-based signatures to block specific traffic when the threshold is exceeded during the configured time period. You should apply rate-based signatures only to protocols you actually use. Then, configure **Duration** to block malicious clients for extended periods. This saves system resources and can discourage a repeat attack. FortiGate does not track statistics for that client while it is temporarily blocklisted.

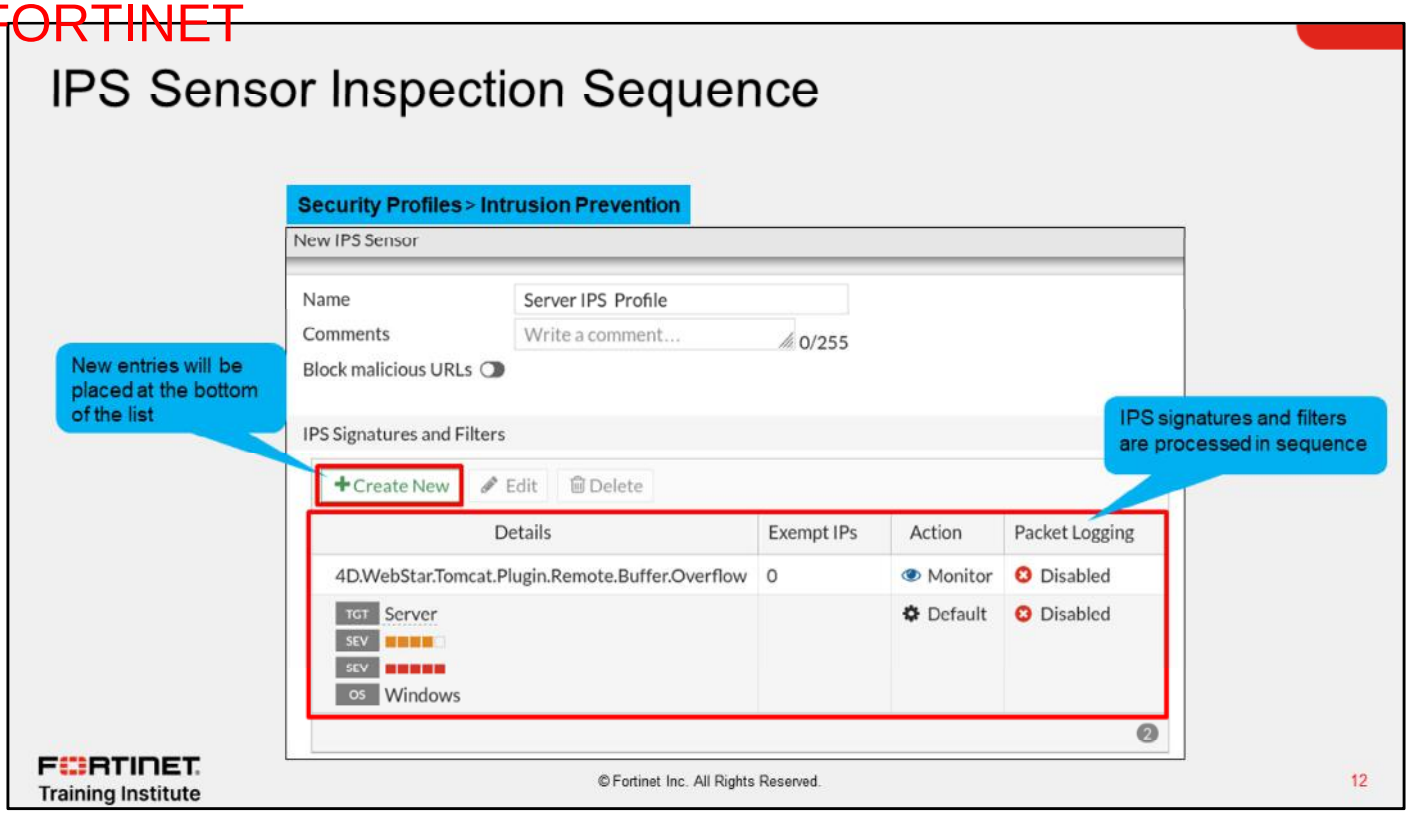

When the IPS engine compares traffic with the signatures in each filter, order matters. The rules are similar to firewall policy matching; the engine evaluates the filters and signatures at the top of the list first, and applies the first match. The engine skips subsequent filters.

So, position the most likely matching filters, or signatures, at the top of the list. Avoid making too many filters, because this increases evaluations and CPU usage. Also, avoid making very large signature groups in each filter, which increase RAM usage.

In the event of a false-positive outbreak, you can add the triggered signature as an individual signature and set the action to **Monitor**. This allows you to monitor the signature events using IPS logs, while investigating the false-positive issue.

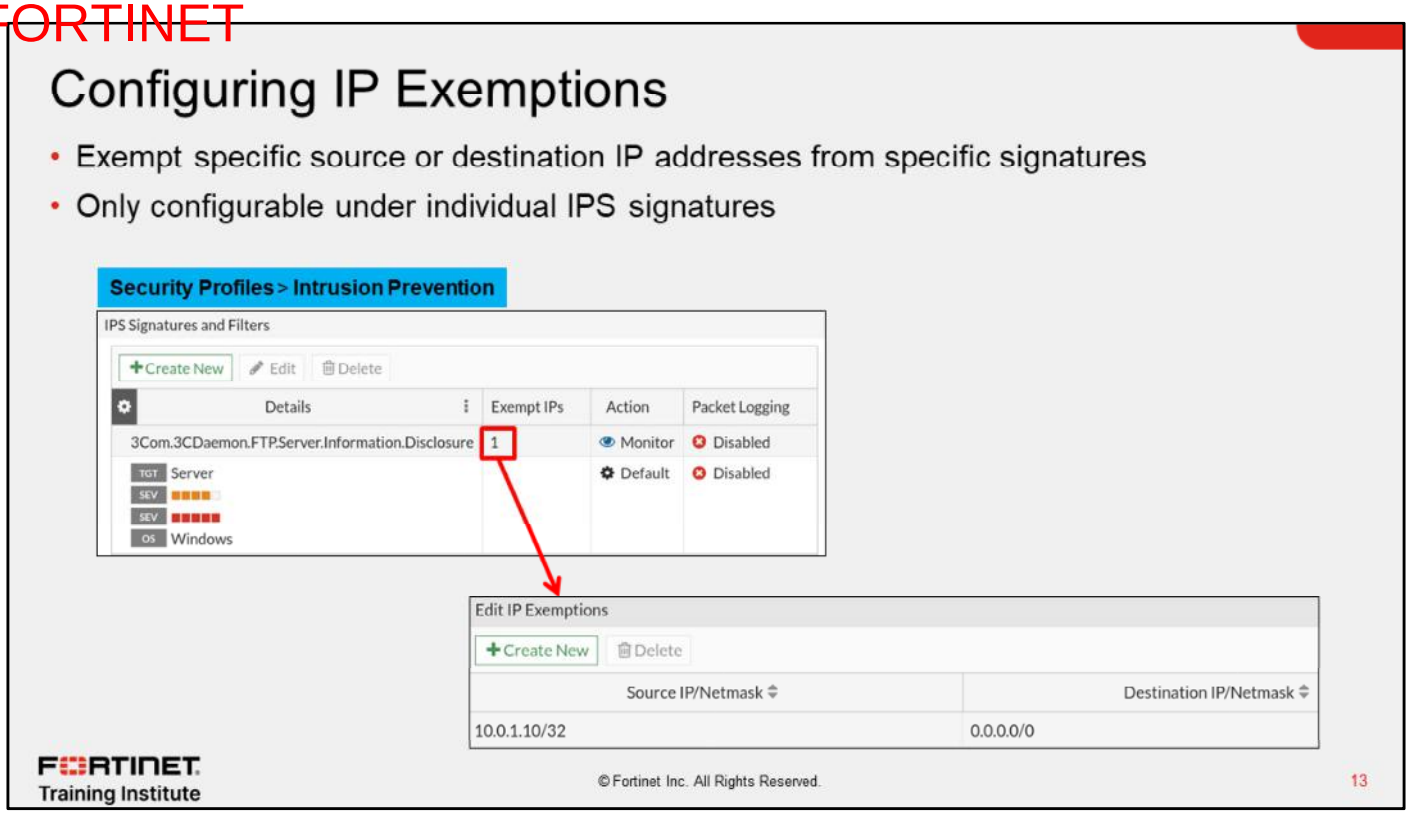

Sometimes it is necessary to exempt specific source or destination IP addresses from specific signatures. This feature is useful during false-positive outbreaks. You can temporarily bypass affected endpoints until you investigate and correct the false-positive issue.

You can configure IP exemptions on individual signatures only. Each signature can have multiple exemptions.

When you create a new entry to add signatures or filters, you can select the action by clicking **Action.**

Select **Allow** to allow traffic to continue to its destination. Select **Monitor** to allow traffic to continue to its destination and log the activity. Select **Block** to silently drop traffic matching any of the signatures included in the entry. Select **Reset** to generate a TCP RST packet whenever the signature is triggered. Select **Default** to use the default action of the signatures.

Quarantine allows you to quarantine the attacker's IP address for a set duration. You can set the quarantine duration to any number of days, hours, or minutes.

If you enable **Packet logging**, FortiGate saves a copy of the packet that matches the signature.

IPS Signature Filter Options—CVE Pattern

- IPS signature filter options include CVE pattern
  - Allows you to filter IPS signatures based on CVE IDs or with a CVE wildcard
  - For example, to configure CVE patterns for CVE-2010-0177

```
# config ips sensor
    edit "cve"
     set comment "cve"
      config entries
        edit 1
          set cve "cve-2010-0177"
          set status enable
          set log-packet enable
          set action block
        next
      end
   next
end
```

- For example, the CVE of the IPS signature `Mozilla.Firefox.PluginArray.NsMimeType.Code.Execution` is CVE-2010-0177
- This matches the CVE filter in the IPS sensor, so traffic is blocked and logged

```
date=2022-04-13 time=15:44:56 logid="0419016384"
type="utm" subtype="ips" eventtype="signature"
level="alert" vd="vd1" eventtime=1594593896666145871
tz="-0700" severity="critical" srcip=10.1.100.22
srccountry="Reserved" dstip=172.16.200.55
srcintf="port2" srcintfrole="undefined"
dstintf="port1" dstintfrole="undefined"
sessionid=1638 action="dropped" proto=6
service="HTTPS" policyid=1
attack="Mozilla.Firefox.PluginArray.NsMimeType.Code.E
xecution" srcport=58298 dstport=443
hostname="172.16.200.55" url="/Mozilla"
direction="incoming" attackid=20853 profile="sensor-
1" ref="http://www.fortinet.com/ids/VID20853"
incidentserialno=124780667 msg="web client:
Mozilla.Firefox.PluginArray.NsMimeType.Code.Execution
," crscore=50 craction=4096 crlevel="critical"
```

FORTINET Training Institute

15

IPS signature filter options include the CVE pattern. The CVE pattern option allows you to filter IPS signatures based on CVE IDs or with a CVE wildcard, ensuring that any signatures tagged with that CVE are automatically included.

Enabling Botnet Protection

- The botnet database:
  - Part of the IPS contract
  - Should be used with the IPS profile to maximize the protection of internal endpoints
- Can be enabled only on the IPS profile
- Administrators can set the action to **Block** or **Monitor**
- IPS logs are generated

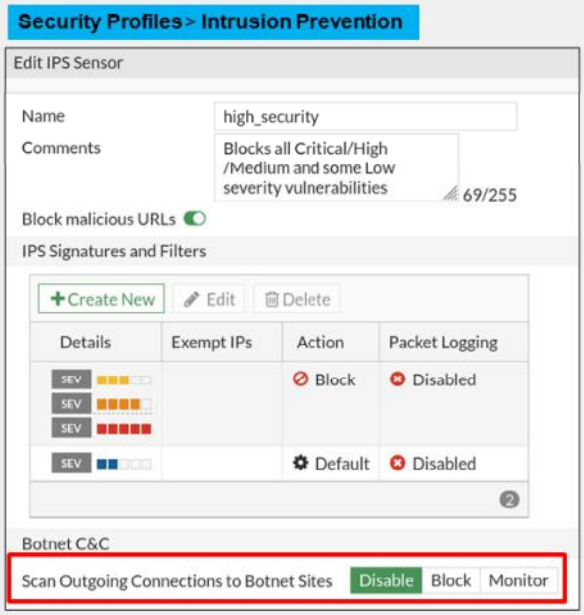

FORTINET Training Institute

16

Since the botnet database is part of the FortiGuard IPS contract, administrators can enable scanning of botnet connections to maximize their internal security. You enable botnet scanning on the IPS profile that you applied the firewall policy on. You can also enable scanning of botnet connections using the CLI.

There are three possible actions for botnet and C&C:
- **Disable**: Do not scan connections to botnet servers
- **Block**: Block connections to botnet servers
- **Monitor**: Log connections to botnet servers

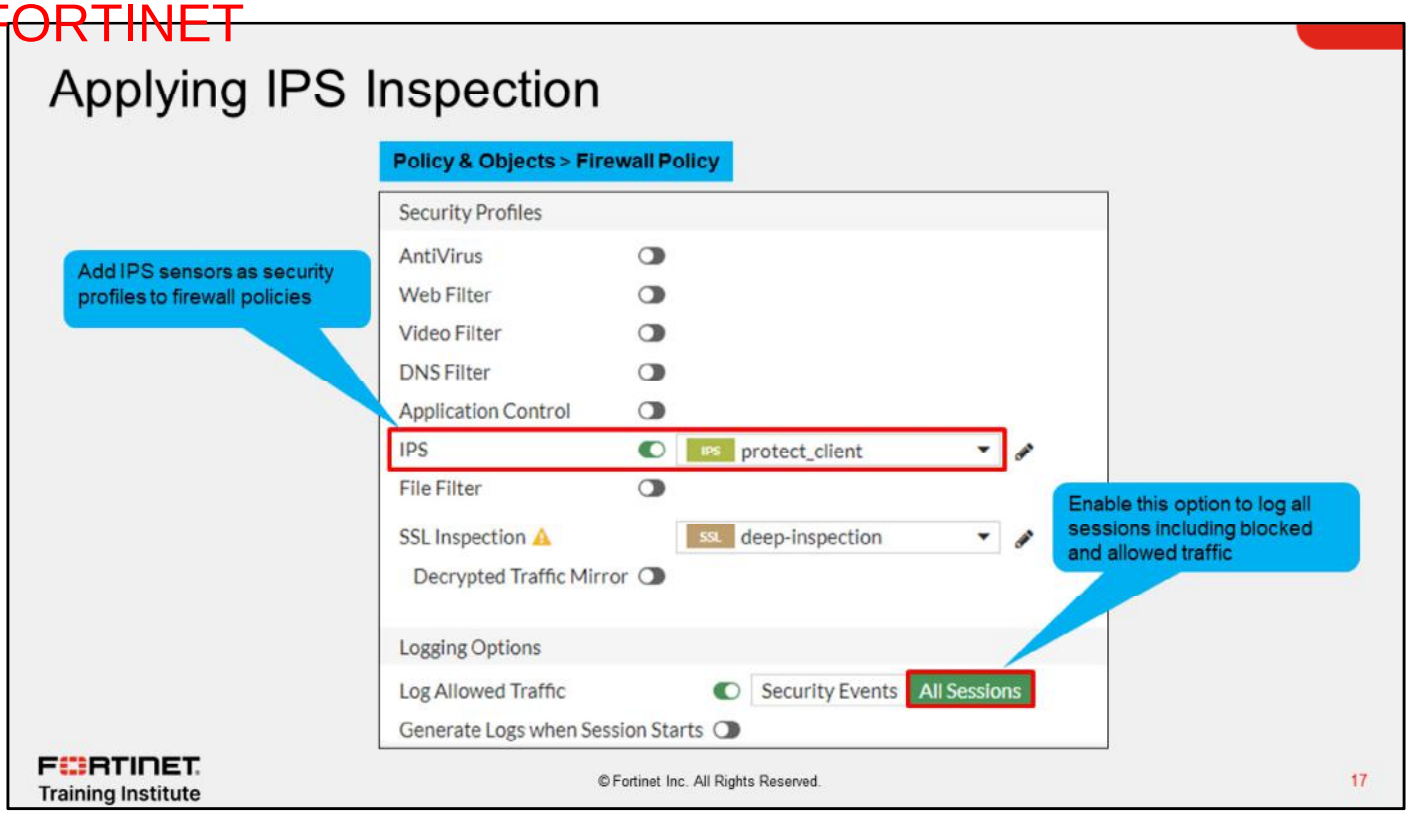

To apply an IPS sensor, you must enable **IPS** and then select the sensor in a firewall policy. By default, FortiGate logs all security events. This means you can see any traffic that is being blocked by IPS.

If you think some traffic should be blocked but is passing through the policy, you should change the **Log Allowed Traffic** method to **All Sessions**. This will log all traffic processed by that firewall policy, and not just the traffic that is blocked by the security profiles. This can help you in identifying false negative events.

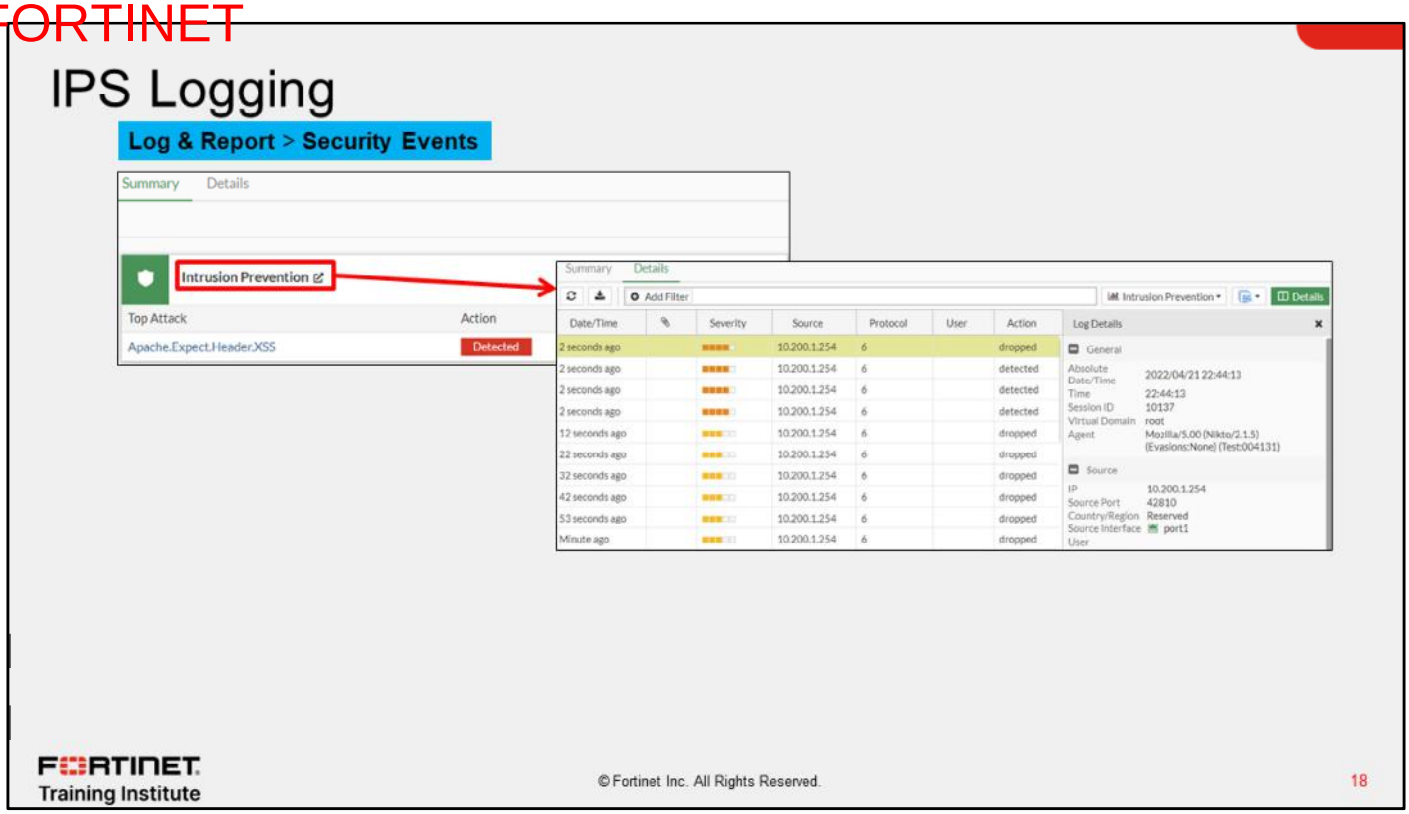

If you enabled security events logging in the firewall policies that apply IPS, you can view events are logged on the **Security Events** pane on the **Log & Report** page. You can view the logs by clicking on **Intrusion Prevention**.

You should review IPS logs frequently. The logs are an invaluable source of information about the kinds of attacks that are being targeted at your network. This helps you develop action plans and focus on specific events, for example, patching a critical vulnerability.

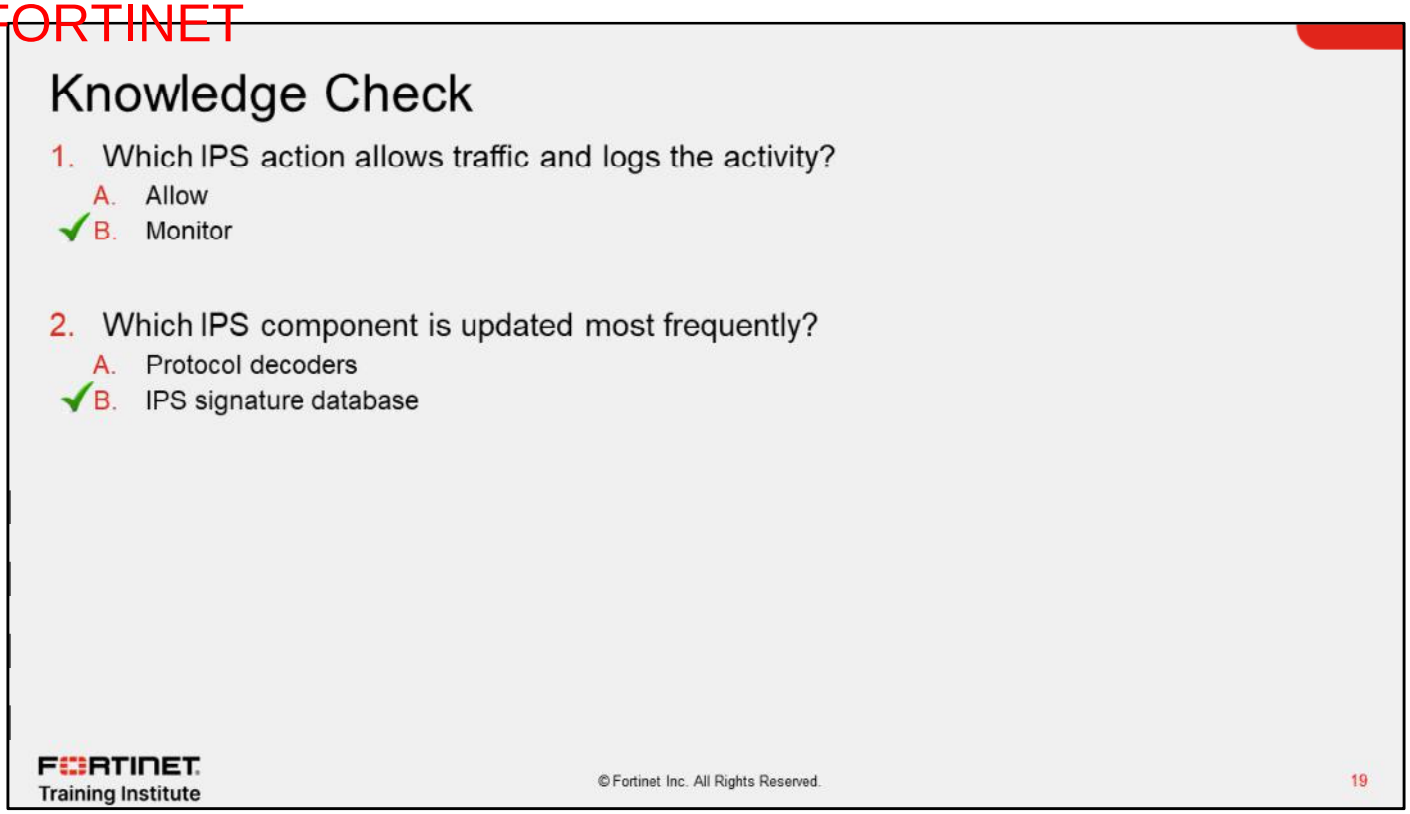

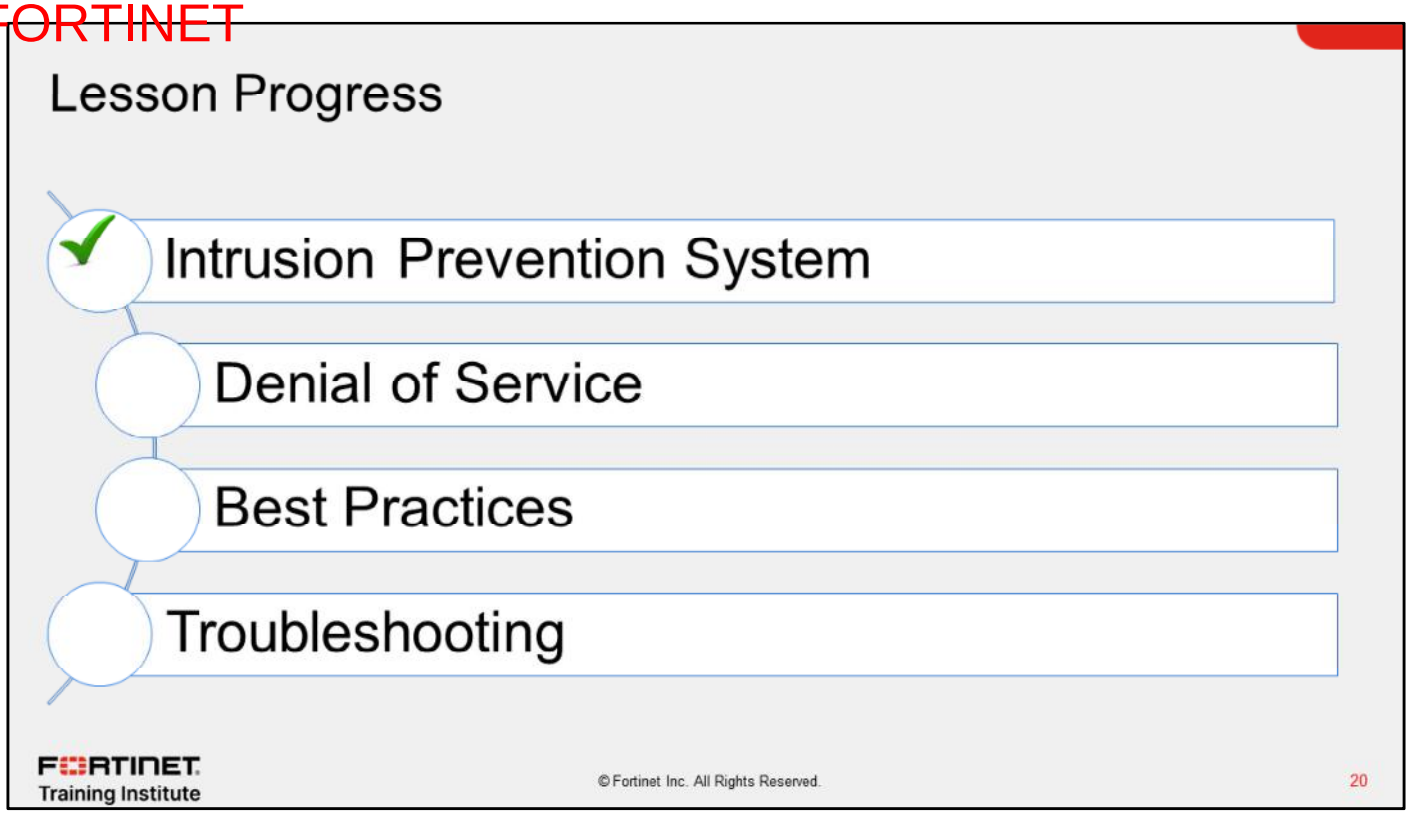

Good job! You now understand the IPS on FortiGate.

Now, you will learn about DoS.

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

By demonstrating competence in Denial of Service (DoS), you should be able to protect your network from common DoS attacks.

DoS Attacks

- Attacker sessions consume all resources—RAM, CPU, port numbers
- Slows down or disables the target until it can't serve legitimate requests

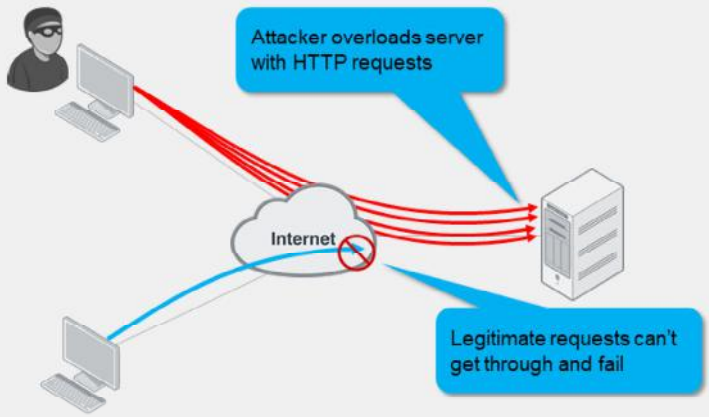

FORTINET Training Institute

22

So far, you have learned about signatures that match illegal commands and invalid protocol implementations. Those are easy to confirm as attacks.

What about attacks that function by exploiting asymmetric processing or bandwidth between clients and servers?

The goal of a DoS attack is to overwhelm the target—to consume resources until the target can't respond to legitimate traffic. There are many ways to accomplish this. High-bandwidth use is only one type of DoS attack. Many sophisticated DoS attacks, such as Slowloris, don't require high bandwidth.

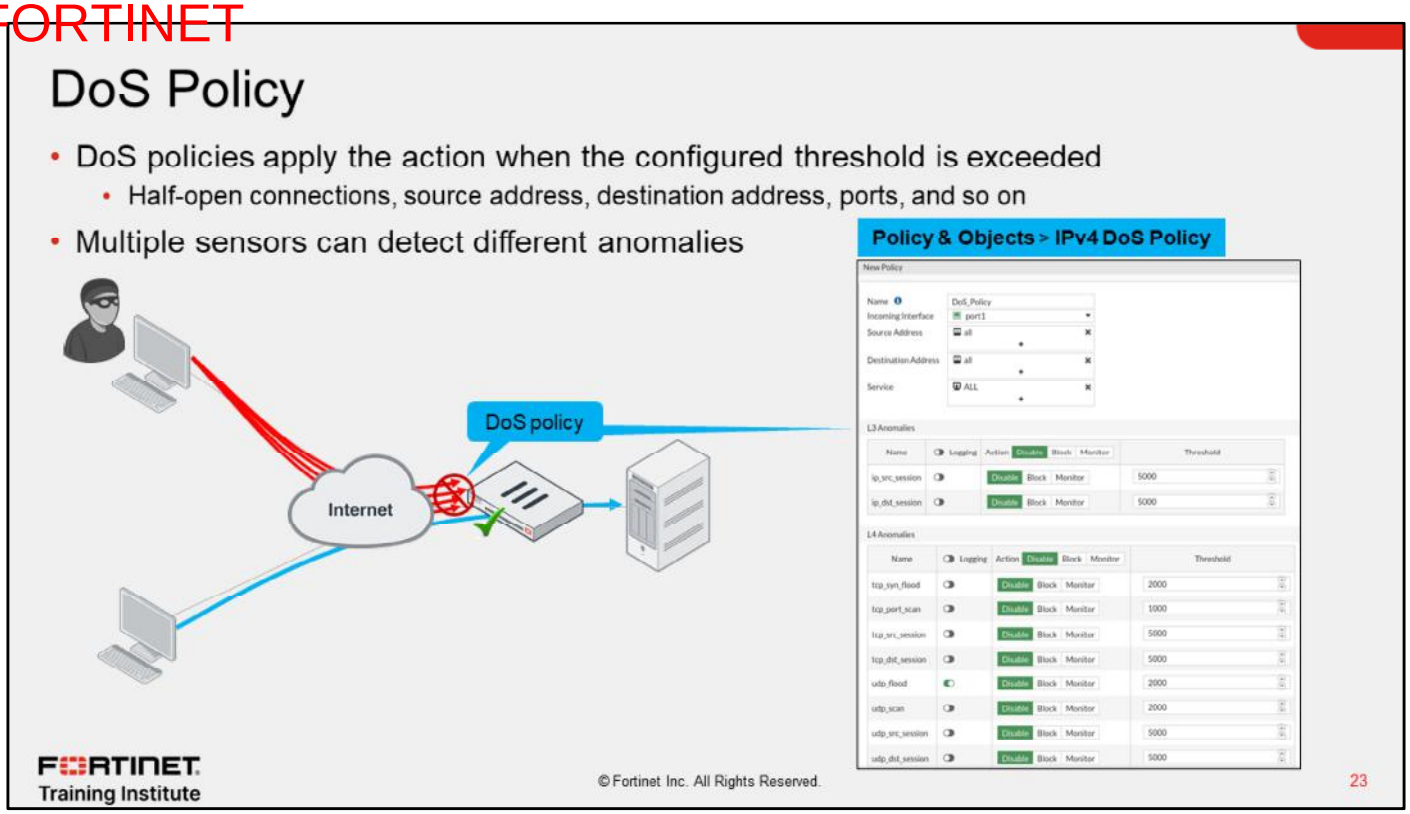

To block DoS attacks, apply a DoS policy on a FortiGate that is located between attackers and all the resources that you want to protect.

DoS filtering is done early in the packet handling process, which is handled by the kernel.

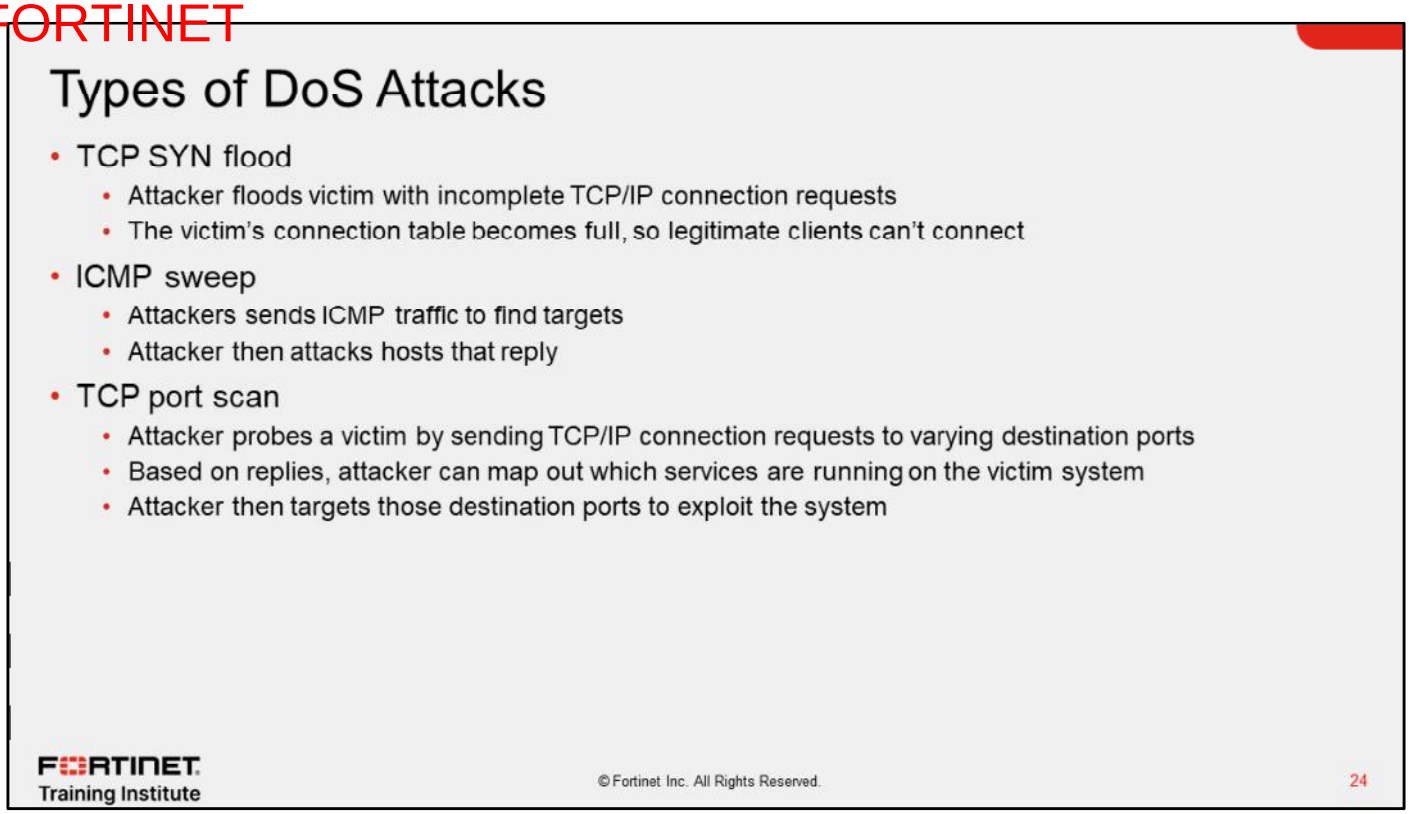

In TCP, the client sends a SYN packet to initiate a connection. The server must respond with a SYN/ACK packet, and save the connection information in RAM while it waits for the client to acknowledge with an ACK packet. Legitimate clients ACK quickly and begin to transmit data. But malicious clients continue to send more SYN packets, half-opening more connections, until the server's connection table is full. Once the server's table is full, it can't accept more connections and begins to ignore all new clients.

ICMP is used during troubleshooting: devices respond with success or error messages. However, attackers can use ICMP to probe a network for valid routes and responsive hosts. By doing an ICMP sweep, the attacker can gain information about your network before crafting more serious exploits.

Attackers use port scanning to determine which ports are active on a system. The attacker sends TCP SYN requests to varying destination ports. Based on the replies, the attacker can map out which services are running on the system, and then proceed to exploit those services.

Types of DoS Attacks (Contd)

- Distributed DoS
  - Many of the same characteristics of an individual DoS attack
  - However, attack originates from multiple sources

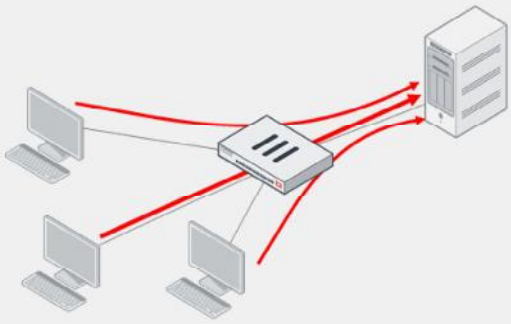

FORTINET Training Institute

25

An individual DoS attack is a flood of traffic coming from a single address. It can originate from the internet, or even from your internal network. Typically, a single device makes many connections or sessions, and possibly uses much bandwidth to connect to a single location. A variation of this is the distributed denial of service attack, or DDoS. It has many of the same characteristics as an individual DoS attack, but the main difference is that multiple devices are all attacking one destination at the same time.

DoS Policy Configuration

- Can apply multiple DoS policies to any physical or logical interface
- Types
  - Flood
    - Detects a large volume of the same type of traffic
  - Sweep/scan
    - Detects probing attempts
  - Source (SRC)
    - Detects a large volume of traffic from an individual IP
  - Destination (DST)
    - Detects a large volume of traffic destined for an individual IP

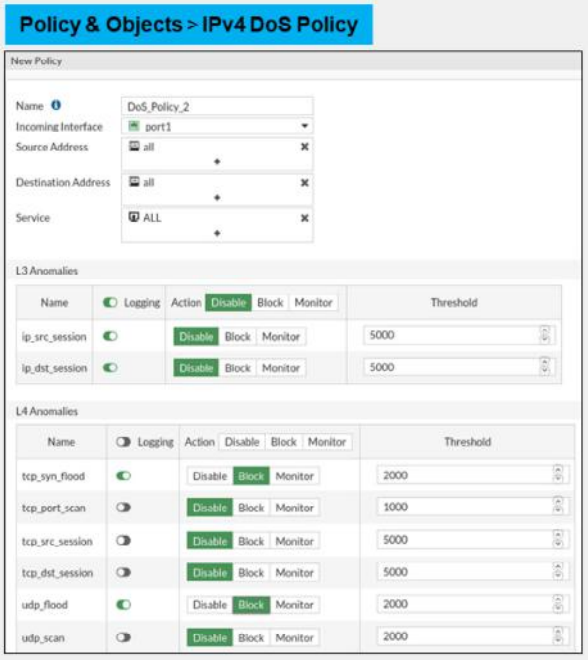

FORTINET Training Institute

26

You can apply DoS protection to four protocols: TCP, UDP, ICMP, and SCTP. And, you can apply four different types of anomaly detection protocols:

- A flood sensor detects a high volume of that specific protocol, or signal in the protocol.
- A sweep/scan detects probing attempts to map which of the host ports respond and, therefore, might be vulnerable.
- Source signatures look for large volumes of traffic originating from a single IP address.
- Destination signatures look for large volumes of traffic destined for a single IP address.

When you implement DoS for the first time, if you don't have an accurate baseline for your network, be careful not to completely block network services. To prevent this from happening, configure the DoS policy initially to log, but not block. Using the logs, you can analyze and identify normal and peak levels for each protocol. Then, adjust the thresholds to allow normal peaks, while applying appropriate filtering.

The threshold for flood, sweep, and scan sensors are defined as the maximum number of sessions or packets per second. The thresholds for source and destination sensors are defined as concurrent sessions. Thresholds that are too high can exhaust your resources before the DoS policies trigger. Thresholds that are too low will cause FortiGate to drop normal traffic.

Knowledge Check

1. Which DoS anomaly sensor can be used to detect and block the probing attempts of a port scanner?
   - A. tcp_syn_flood
   - ✓ B. tcp_port_scan

2. Which behavior is a characteristic of a DoS attack?
   - A. Attempts to exploit a known application vulnerability
   - ✓ B. Attempts to overload a server with TCP SYN packets

FORTINET Training Institute

27

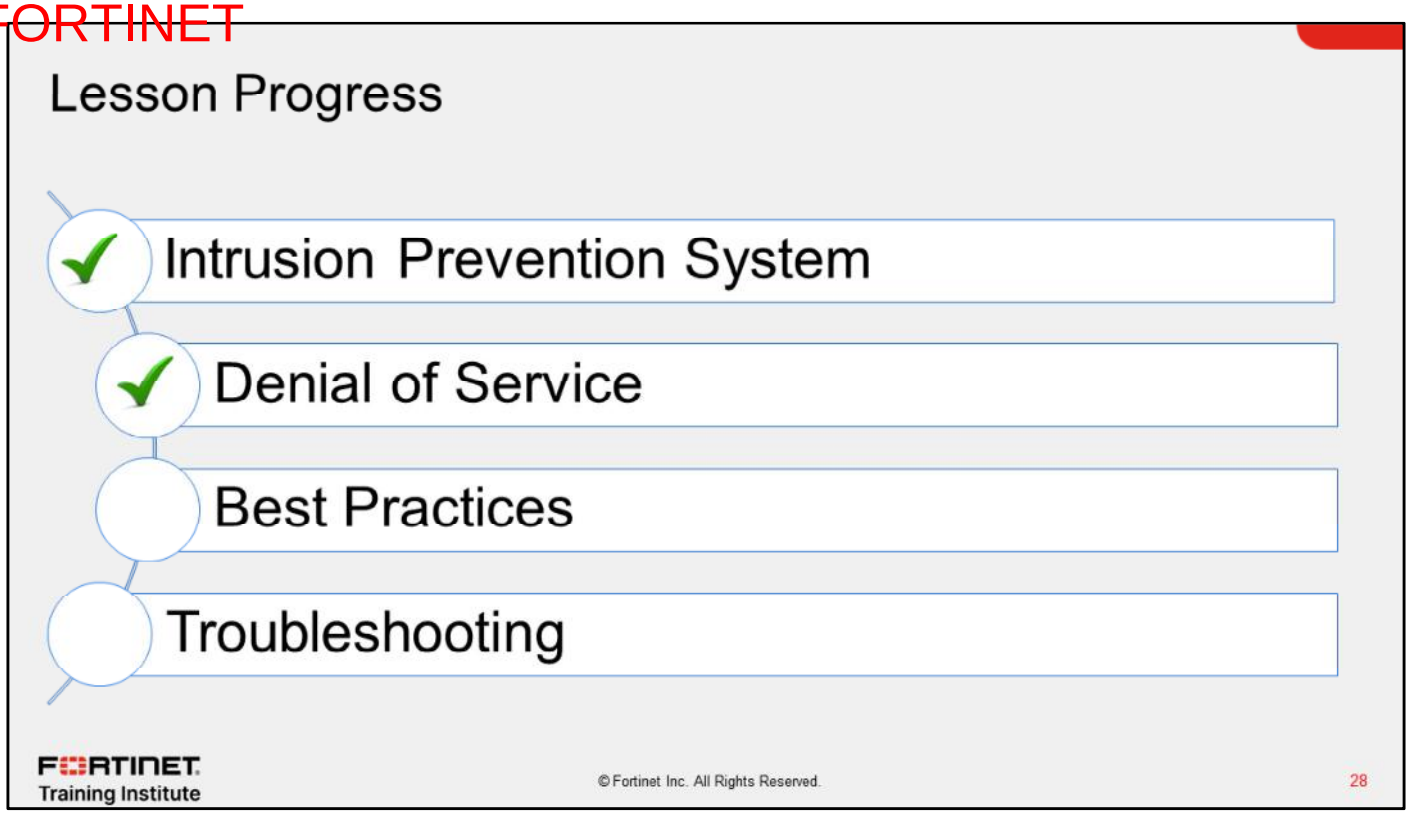

Good job! You now understand how to protect your network from DoS attacks on FortiGate.

Now, you will learn about IPS best practices.

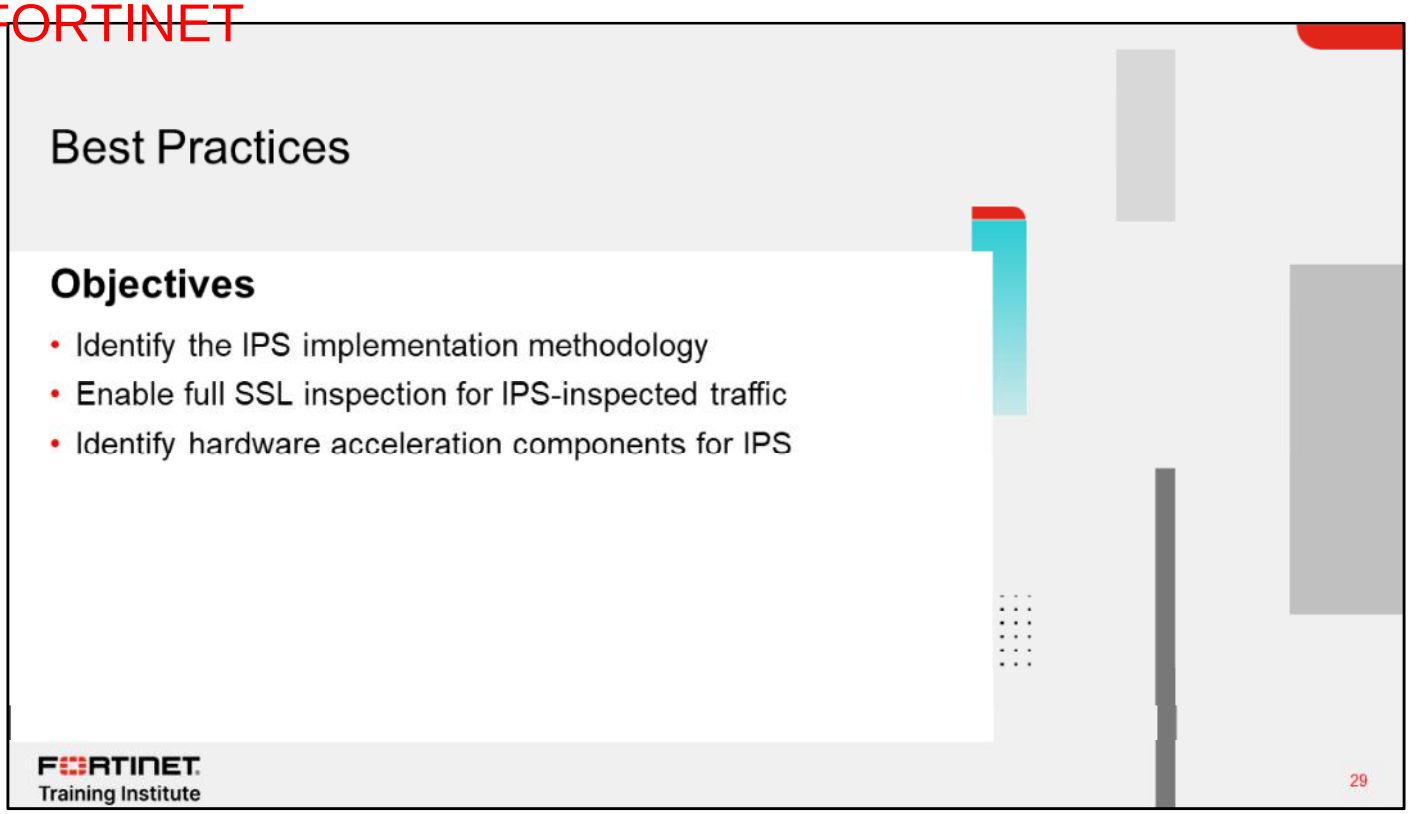

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in identifying IPS implementation best practices, you should be able to deploy an IPS solution on FortiGate that is efficient and effective. You should also be able to apply full SSL inspection for IPS-inspected traffic, as well as identify hardware acceleration components for IPS.

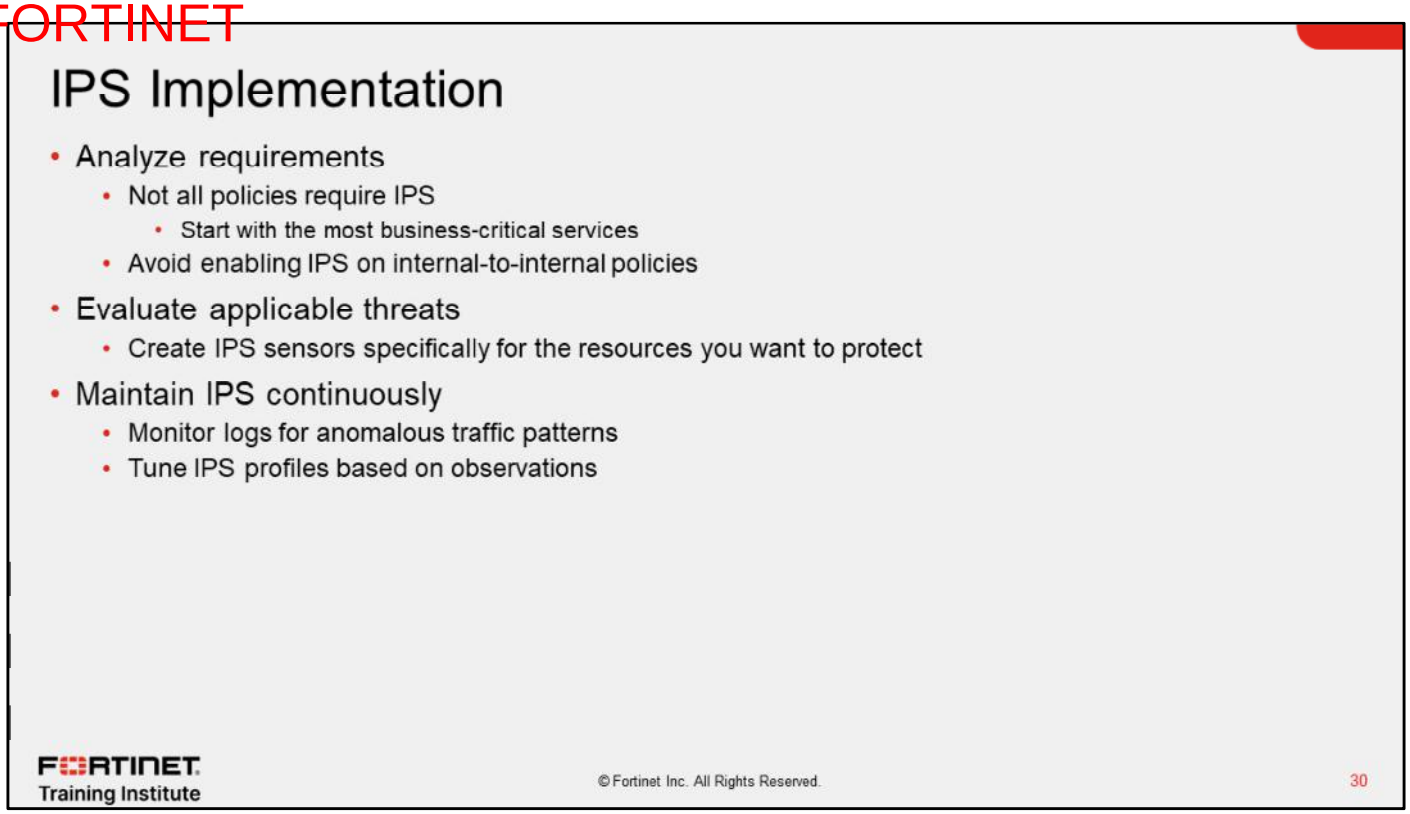

Before you implement IPS, you must analyze the needs of your network. Enabling the default profiles across all policies quickly causes issues, the least of which are false positives. Performing unnecessary inspections on all network traffic can cause high resource utilization, which can hamper the ability of FortiGate to process regular traffic.

You must also evaluate applicable threats. If your organization runs only Windows, there is no need to scan for Mac OS vulnerabilities. It is also important to consider the direction of the traffic. There are many IPS signatures that apply only to clients, and many signatures that apply only to servers. Create IPS sensors specific to the resources you want to protect. This makes sure that FortiGate is not scanning traffic with irrelevant signatures.

Lastly, IPS is not a *set-and-forget* implementation. You must monitor logs regularly for anomalous traffic patterns, and adjust your IPS profile configuration based on your observations. You should also audit your internal resources regularly to identify if certain vulnerabilities still apply to your organization.

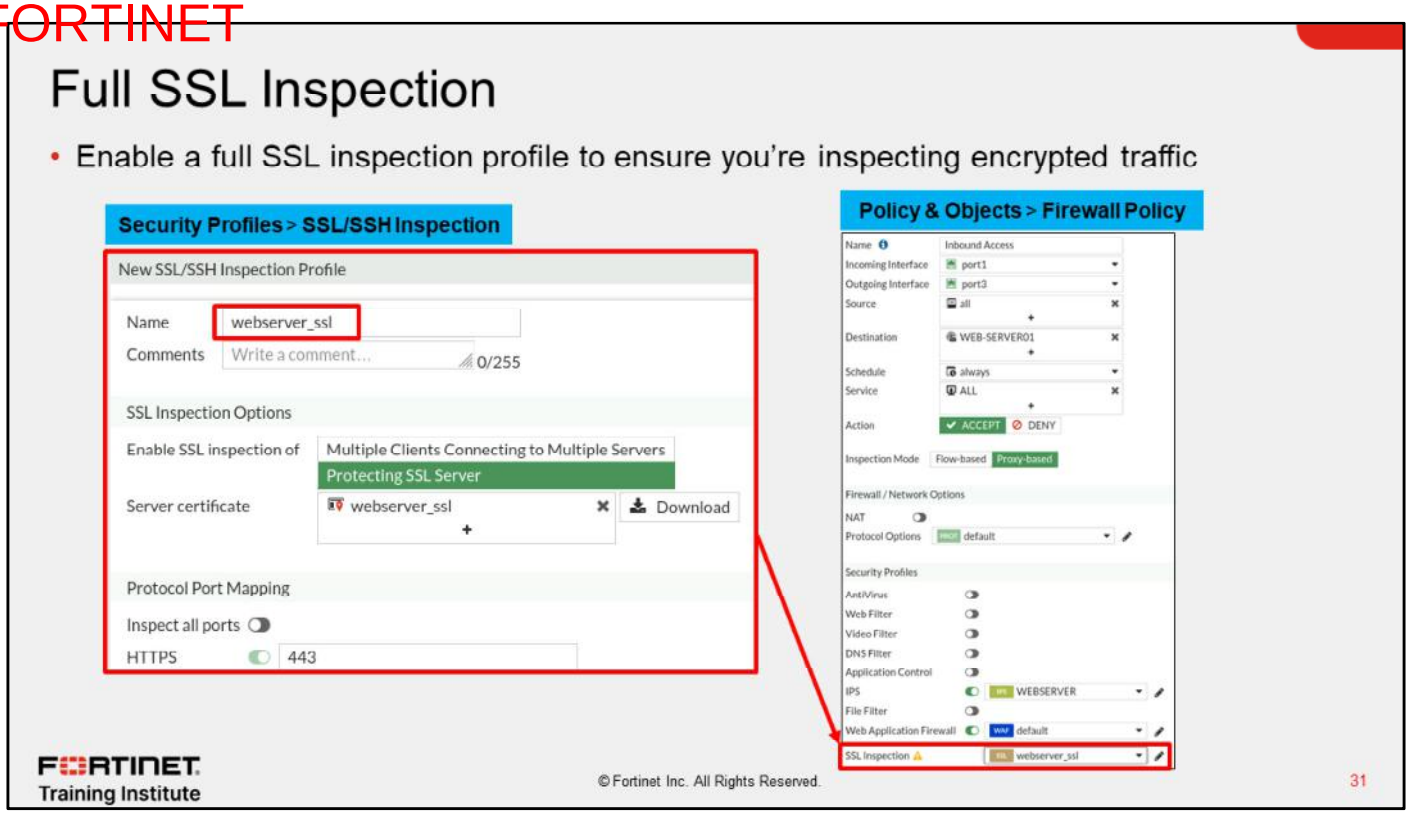

Certain vulnerabilities apply only to encrypted connections. In some of these cases, FortiGate can't identify the threat reliably if it can't parse the payload. For this reason, you must use an SSL inspection profile if you want to get the maximum benefit from your IPS and WAF features.

The example on this slide shows an SSL inspection profile configured to protect a server. This policy, when applied to inbound traffic, can apply IPS and WAF inspection on encrypted traffic reliably, because FortiGate can decrypt encrypted sessions and inspect all parts of the packet.

It's important to note that DoS policies do not have the ability to assign SSL inspection profiles. This is because DoS does not require SSL inspection to maximize its detection ability, because it does not inspect the packet payload. DoS inspects only specific session types and their associated volume.

Hardware Acceleration

- FortiGate models with NP6, NP7, and SoC4 can benefit from NTurbo acceleration (`np-accel-mode`)
- FortiGate models with CP8 or CP9 support offloading of IPS pattern matching to the content processor (`cp-accel-mode`)

```
fgt # get hardware status
Model name: FortiGate-300D
ASIC version: CP8
ASIC SRAM: 64M
CPU: Intel(R) Core(TM) i3-3220 CPU @ 3.30GHz
Number of CPUs: 4
RAM: 7996 MB
Compact Flash: 15331 MB /dev/sda
Hard disk: 114473 MB /dev/sdb
USB Flash: not available
Network Card chipset: Intel(R) Gigabit Ethernet
Network Driver (rev.0003)
Network Card chipset: FortiASIC NP6 Adapter (rev.)
```

```
# config ips global
# set np-accel-mode [ basic | none ]
# set cp-accel-mode [ basic | advanced | none ]
# end
```

`np-accel-mode`

- `basic`: offloads IPS processing to NP

`cp-accel-mode`

- `basic`: offloads basic IPS pattern matching to CP8 or CP9
- `advanced`: offloads more types of IPS pattern matching
  - Only available on devices with two or more CP8s or one or more CP9s

FORTINET Training Institute

32

Usually, traffic requiring inspection, such as antivirus or IPS, is handled by the CPU on FortiGate. However, there are specialized chips on specific FortiGate models that can offload these inspection tasks. This frees up CPU cycles to manage other tasks, and also accelerates sessions requiring security inspection.

FortiGate models that support a feature called NTurbo can offload IPS processing to NP6, NP7, or SoC4 processors. If the command `np-accel-mode` is available under `config system global`, the FortiGate model supports NTurbo.

Some FortiGate models also support offloading IPS pattern matching to CP8 or CP9 content processors. If the command `cp-accel-mode` is available under `config ips global`, the FortiGate model supports IPS pattern matching acceleration to its CP8 or CP9 processor.

Knowledge Check

1. Which chipset uses NTurbo to accelerate IPS sessions?
   - A. CP9
   - ✓ B. SoC4

2. Which feature requires full SSL inspection to maximize its detection capability?
   - ✓ A. WAF
   - B. DoS

FORTINET Training Institute

33

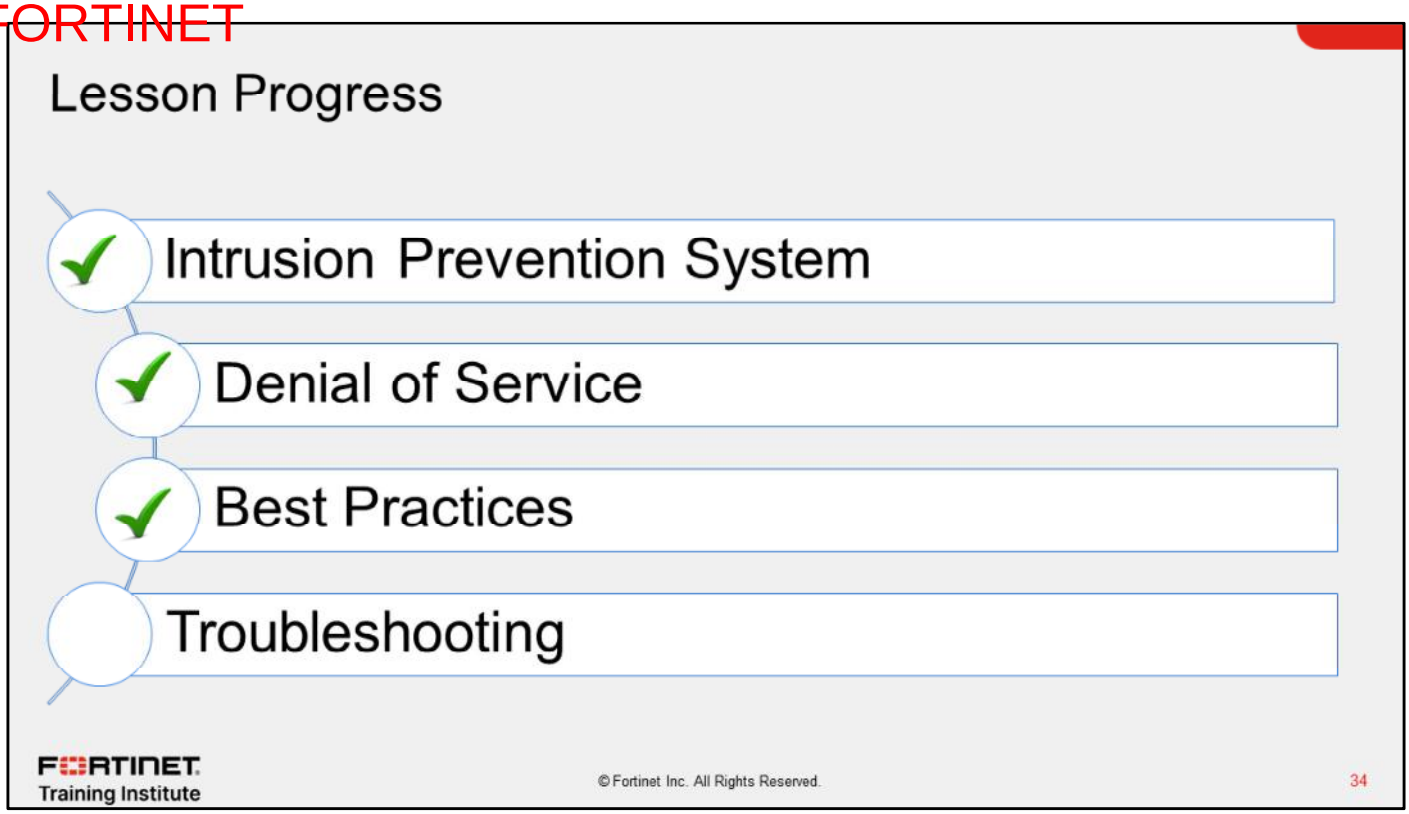

Good job! You now understand some best practices for IPS implementation on FortiGate.

Now, you will learn about IPS troubleshooting.

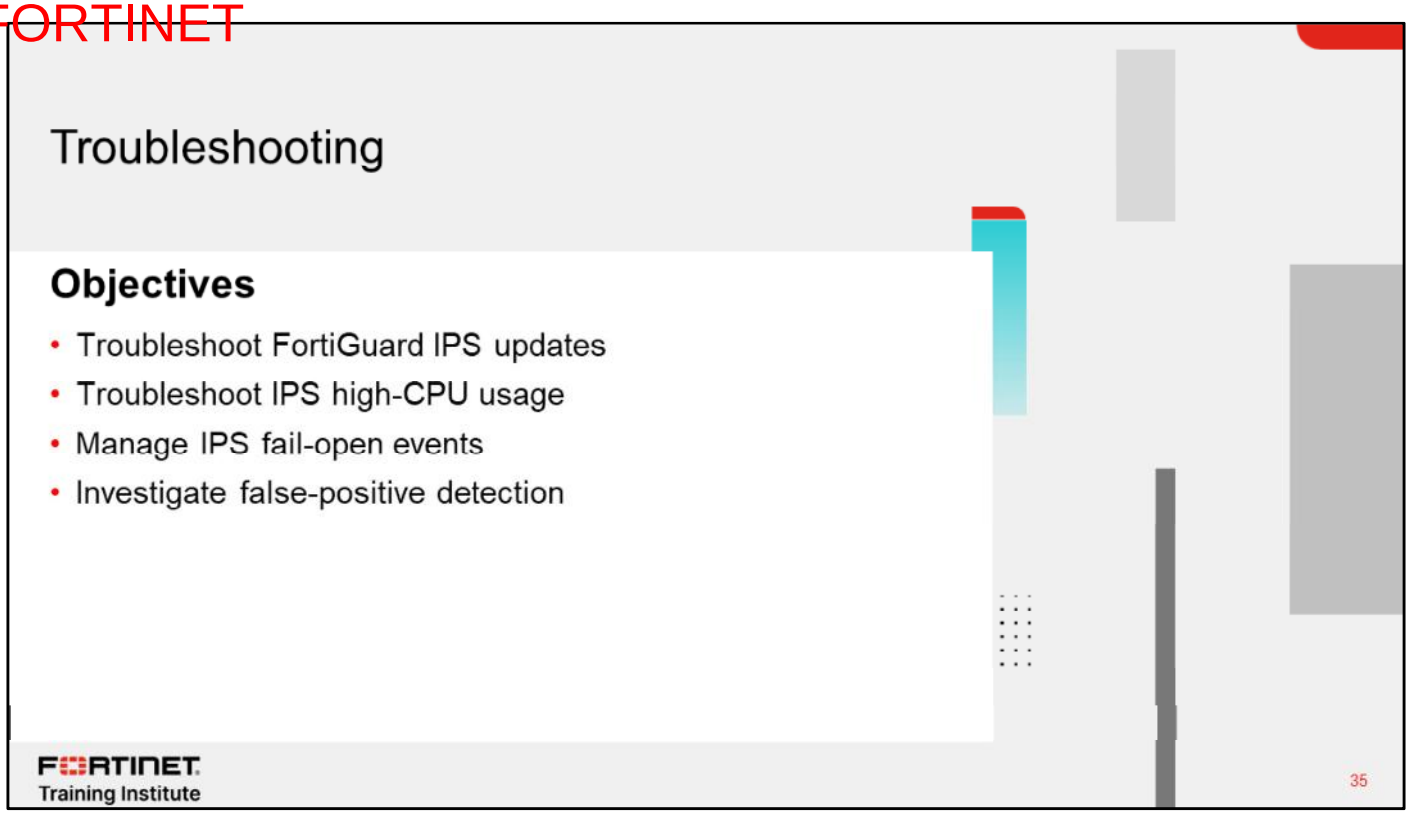

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in troubleshooting, you should be able to identify, investigate, and manage some common issues with IPS deployments on FortiGate.

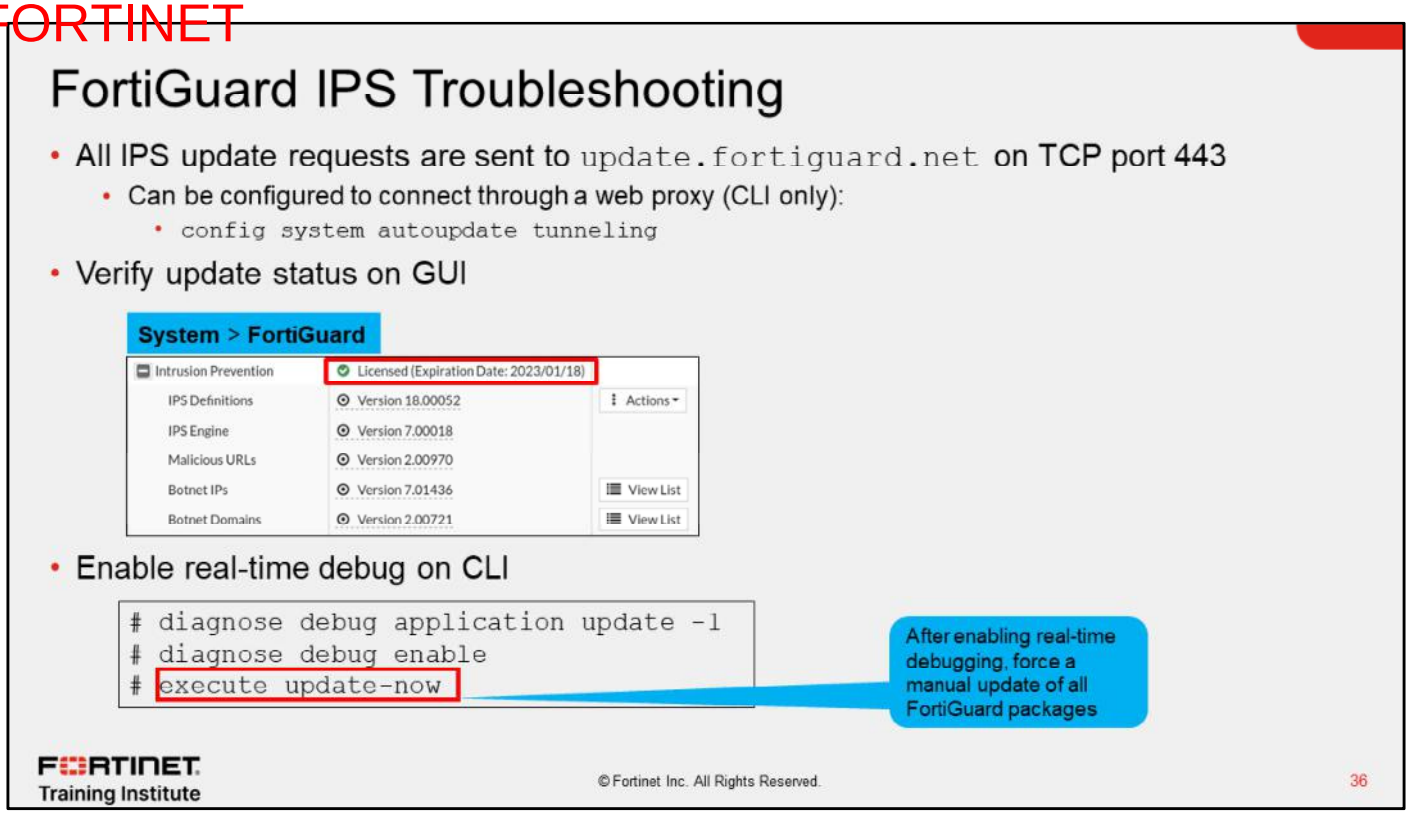

FortiGate IPS update requests are sent to `update.fortiguard.net` on TCP port 443. You can also configure FortiGate to connect through a web proxy for updates.

You should check the last update timestamp regularly. You can verify it on the GUI. If there is any indication that the IPS definitions are not updating, you should investigate. Always make sure FortiGate has proper DNS resolution for `update.fortiguard.net`. If, by chance, there are any intermediary devices between the FortiGate and the internet, make sure the correct firewall rules are in place to allow traffic on port443. Any intermediary devices performing SSL inspection on this traffic can also cause issues with updates.

Finally, you can use the FortiGuard update debug to monitor update events in real time.

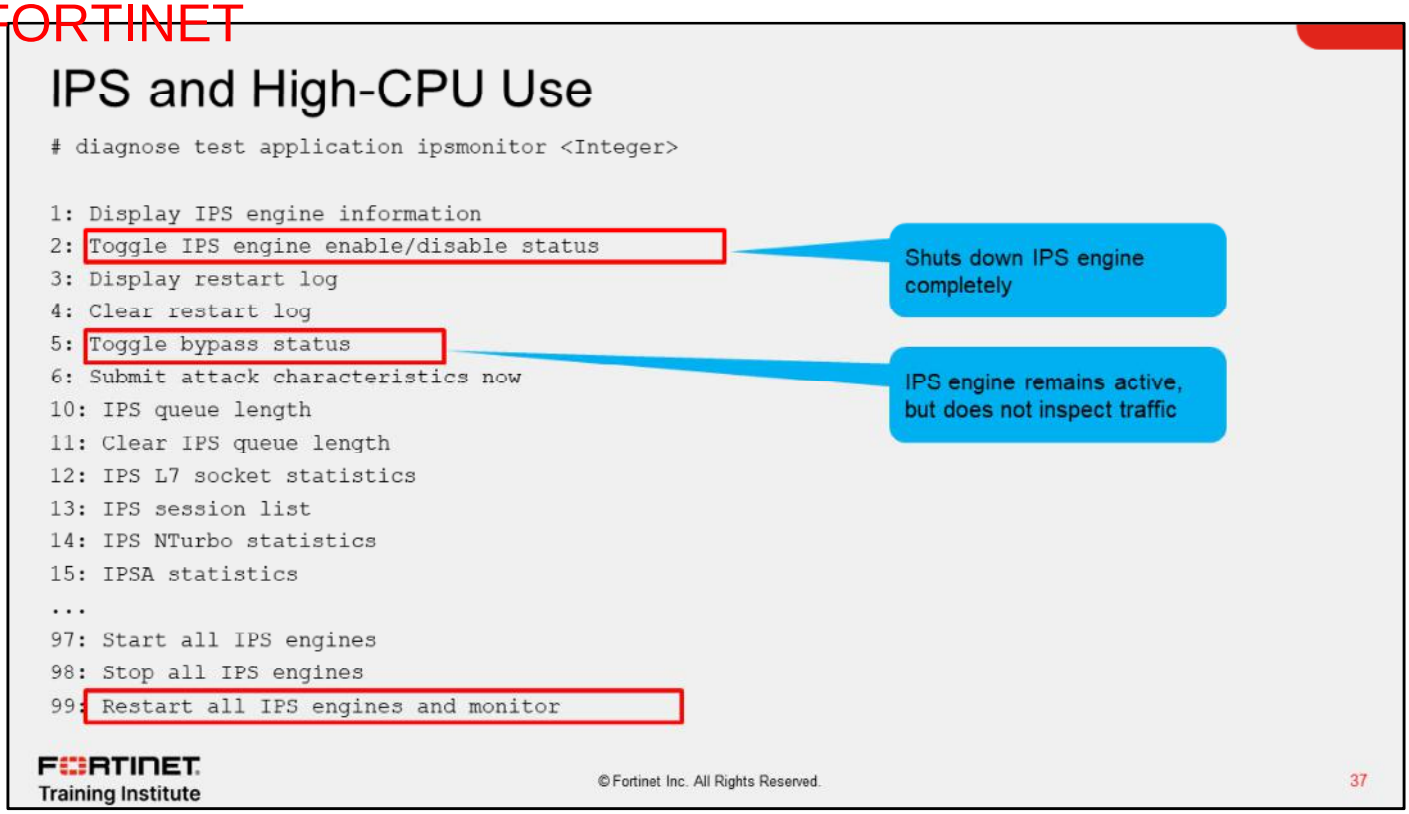

Short spikes in CPU usage by IPS processes can be caused by firewall policy or profile changes. These spikes are usually normal. Spikes might happen when FortiGate has hundreds of policies and profiles, or many virtual domains. Continuous high-CPU use by the IPS engines is not normal, and you should investigate it. You can use the command shown on this slide, along with displayed options, to troubleshoot these issues.

If there are high-CPU use problems caused by the IPS, you can use the `diagnose test application ipsmonitor` command with option `5` to isolate where the problem might be. Option `5` enables IPS bypass mode. In this mode, the IPS engine is still running, but it is not inspecting traffic. If the CPU use decreases after that, it usually indicates that the volume of traffic being inspected is too high for that FortiGate model.

If the CPU use remains high after enabling IPS bypass mode, it usually indicates a problem in the IPS engine, which you must report to Fortinet Support. You can disable the IPS engine completely using option `2`. If you want to restore IPS inspection of traffic after you finish troubleshooting, use option `5` again.

Another recommendation to keep in mind: if you need to restart the IPS, use option `99`, as shown on this slide. This guarantees that all the IPS-related processes restart properly.

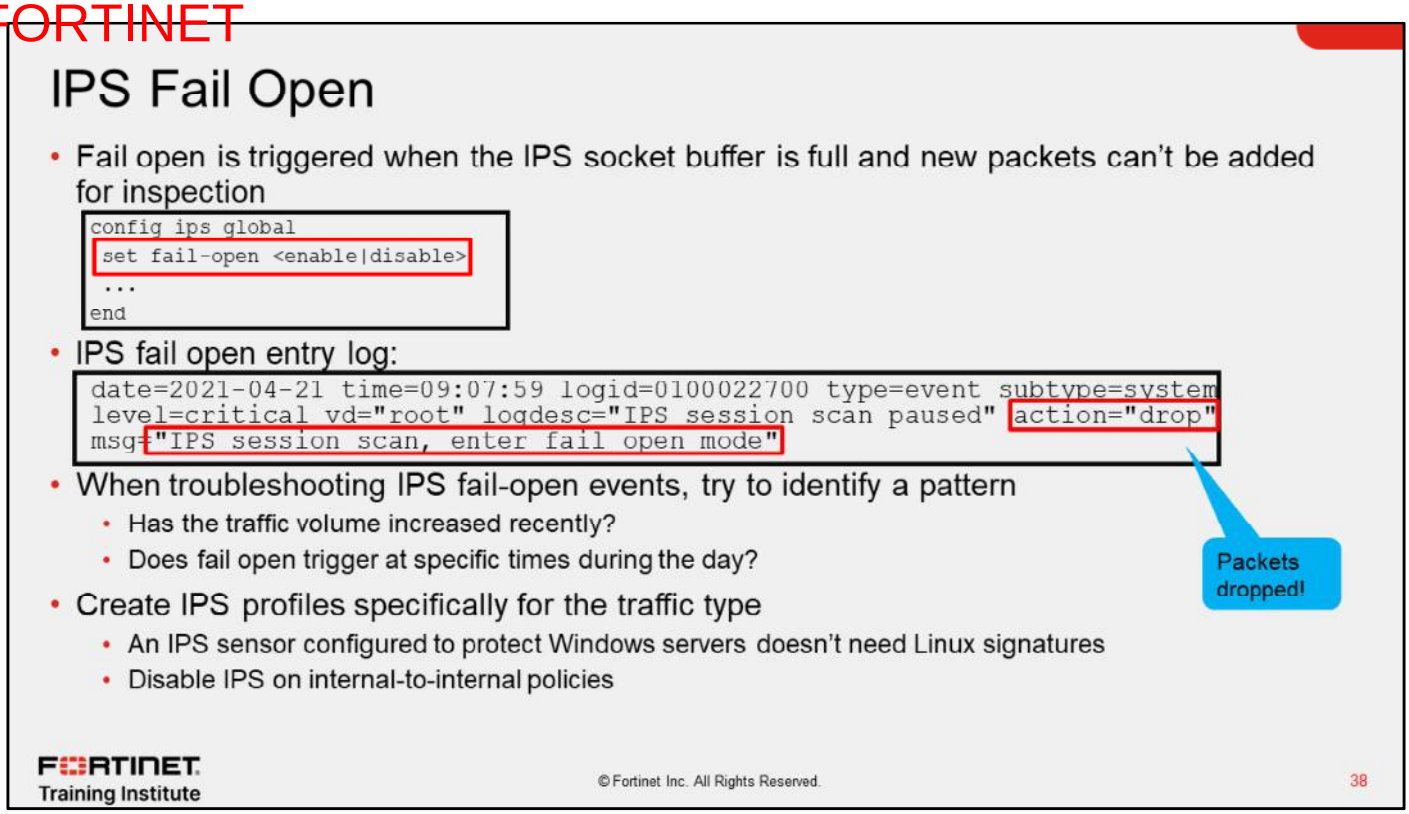

IPS goes into fail-open mode when there is not enough available memory in the IPS socket buffer for new packets. What happens during this state depends on the IPS configuration. If the `fail-open` setting is enabled, some new packets (depending on the system load) will pass through without being inspected. If it is disabled, new packets are dropped.

Frequent IPS fail open events usually indicate that IPS can't keep up with the traffic demands. So, try to identify patterns. Has the traffic volume increased recently? Have throughput demands increased? Does fail open trigger at specific times during the day?

Tune and optimize your IPS configuration. Create IPS profiles specific to the type of traffic being inspected, and disable IPS profiles on policies that don't need them.

Knowledge Check

1. Which FQDN does FortiGate use to obtain IPS updates?
   - ✓ A. update.fortiguard.net
   - B. service.fortiguard.com

2. When IPS fail open is triggered, what is the expected behavior, if the IPS fail-open option is set to enabled?
   - ✓ A. New packets pass through without inspection
   - B. New packets are dropped

FORTINET Training Institute

39

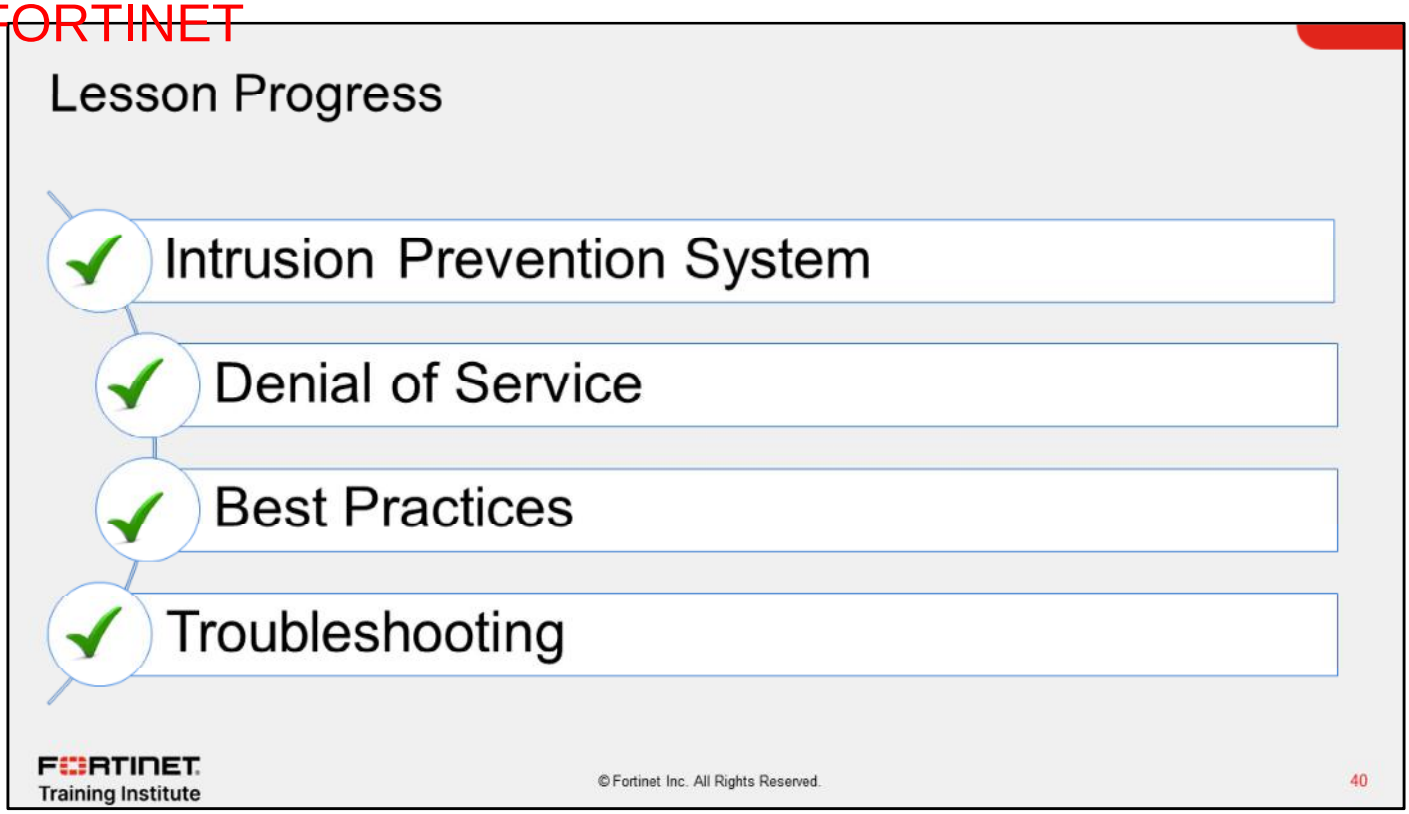

Congratulations! You have completed this lesson.

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

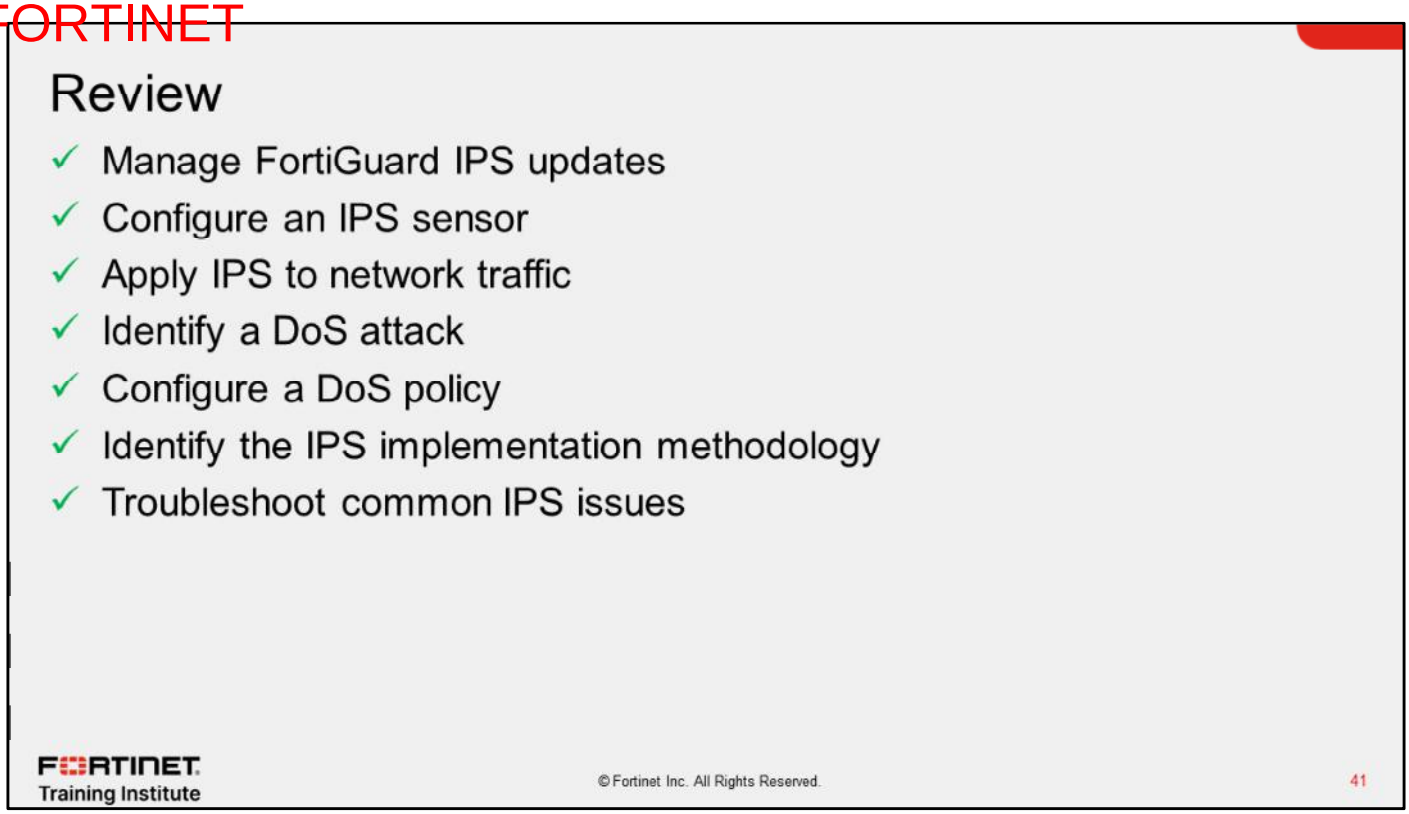

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you gained the skills and knowledge you need to configure, maintain, and troubleshoot the FortiGate IPS solution.

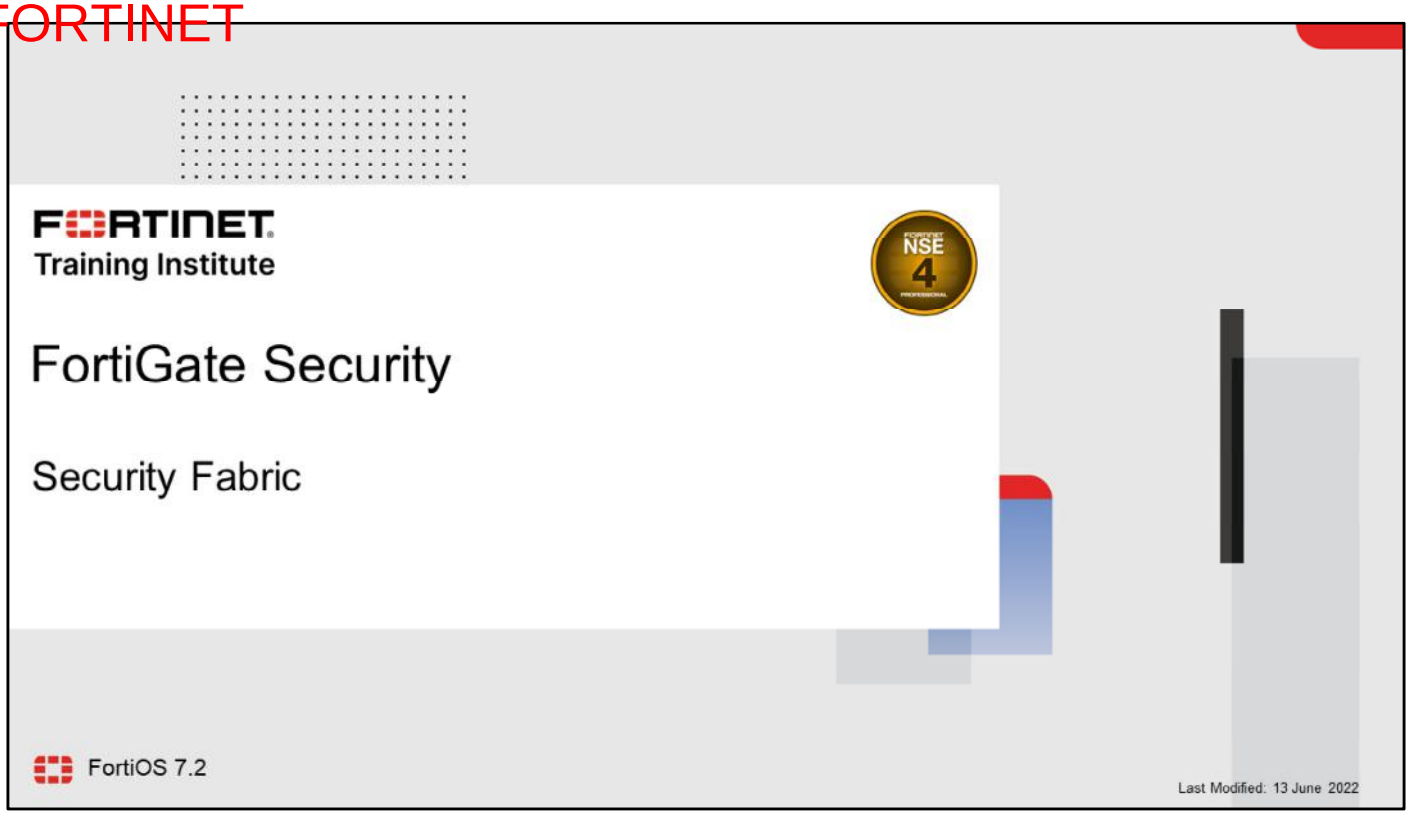

In this lesson, you will learn about the Fortinet Security Fabric.

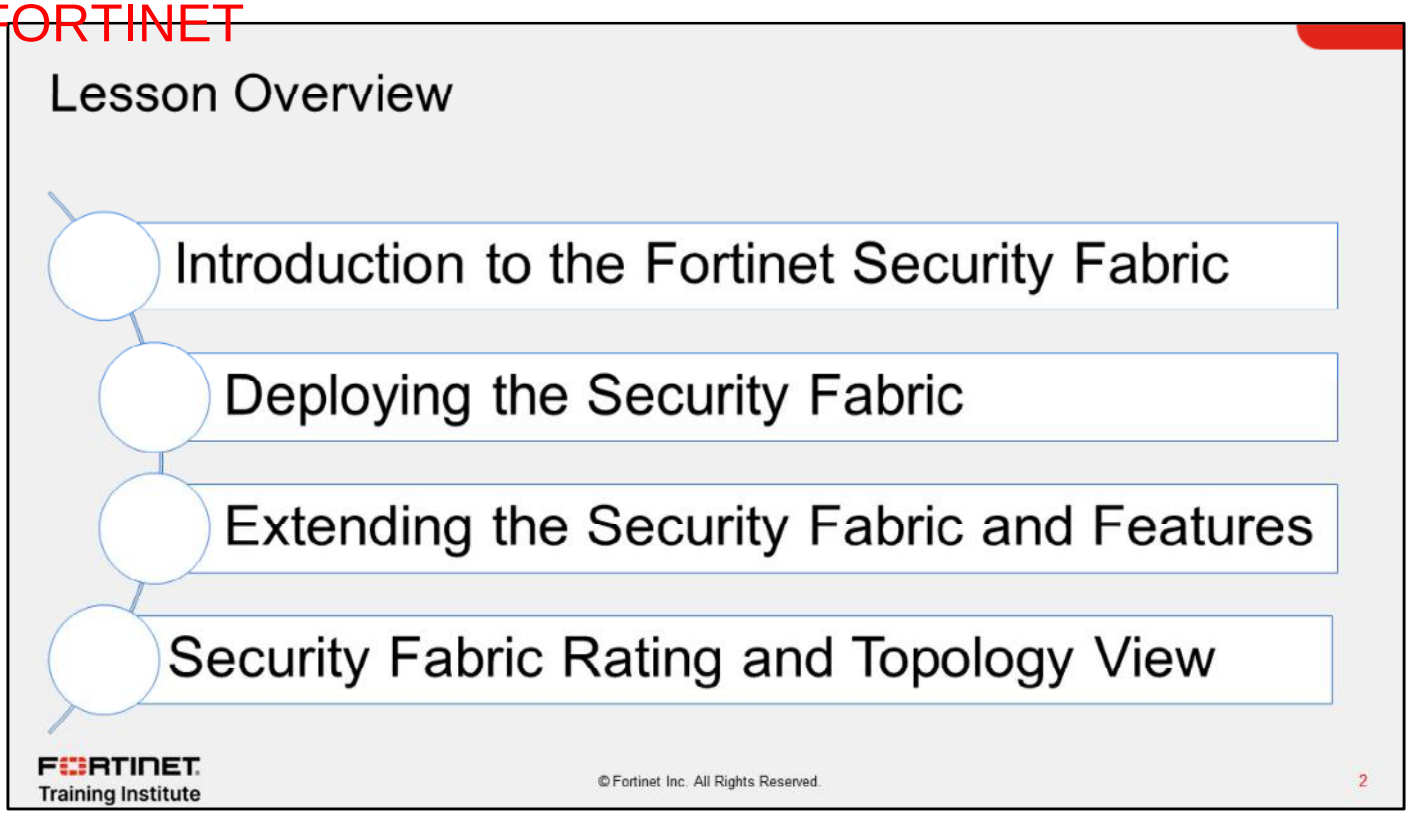

In this lesson, you will learn about the topics shown on this slide.

By demonstrating competence in deploying the Fortinet Security Fabric, using and extending the Security Fabric features, and understanding its topology, you will be able to use the Fortinet Security Fabric effectively in your network.

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

By demonstrating competence in understanding key concepts of the Fortinet Security Fabric, you will better understand the value of the Security Fabric, the servers that comprise it, and how to deploy it.

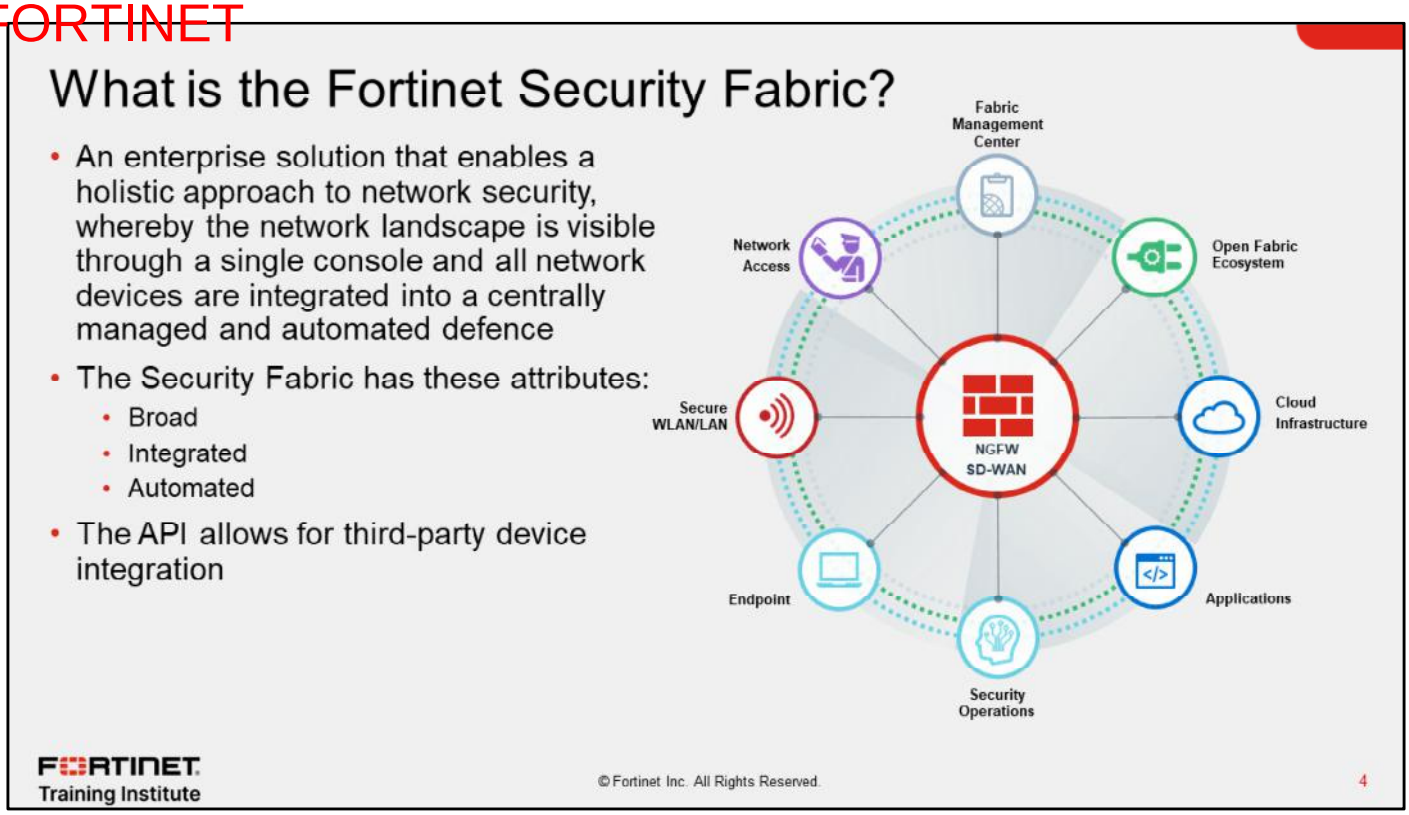

What is the Fortinet Security Fabric?

It is a Fortinet enterprise solution that enables a holistic approach to network security, whereby the network landscape is visible through a single console and all network devices are integrated into a centrally managed and automated defence.

The network devices include all components, from physical endpoints to virtual devices in the cloud. Because devices are centrally managed and are sharing threat intelligence with one another in real time, and are receiving updates from Fortinet at the macro level, your network can quickly identify, isolate, and neutralize threats as they appear.

The Security Fabric has the following attributes:
- **Broad**: It provides visibility of the entire digital attack surface to better manage risk
- **Integrated**: It provides a solution that reduces the complexity of supporting multiple point products
- **Automated**: Threat intelligence is exchanged between network components in real-time allowing for automated response to threats

A fourth attribute could be added to this description of the Security Fabric: *open*. The API and protocol are available for other vendors to join and for partner integration. This allows for communication between Fortinet and third-party devices.

Why a Security Fabric?

- Many administrators lack visibility of their network defences, making their networks more susceptible to undetected network infiltration
- Network complexity and sophisticated malware (soon to be augmented by AI), necessitates a centralized and holistic approach to security

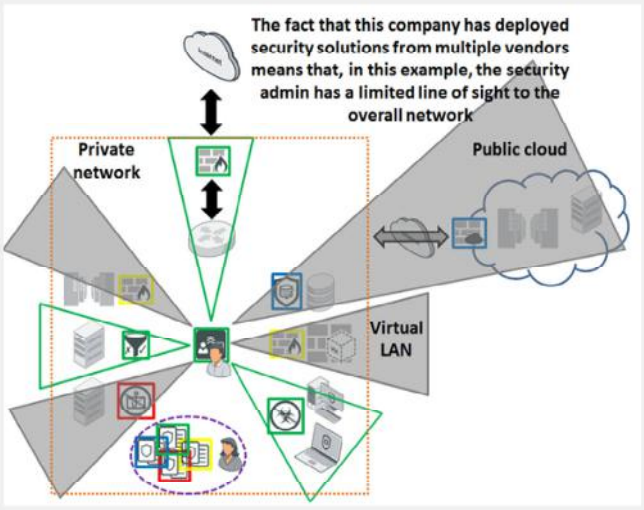

FORTINET Training Institute

5

Why has Fortinet deemed the Security Fabric an essential solution for a robust network defence?

As networks evolved and various new types of threats surfaced, point security products were deployed to address these emerging threats. Often, these piecemeal solutions were effective, but deploying products using different standards and protocols meant that defence assets could not be effectively coordinated.

The illustration on the right side of the slide tells a story of a network that has deployed security solutions from four different vendors. The administrator at the center, working from the security console, has visibility into only some of the security solutions. This lack of visibility of the entire network defence is a serious flaw, and could allow a foreign infiltrator to breach network defences undetected.

The sheer complexity of today's networks compounds this problem. In addition, increasingly sophisticated malware has an expanding attack surface on which to exploit, because networks have broken out of the confines of a traditional network perimeter and have expanded to virtualized networks and public clouds. Add to this mix, the ever growing numbers of unmanaged devices, as a result of BYOD programs, and you have the perfect security storm.

The most feasible solution is to build a centrally managed, holistic approach to security, whereby you have a clear line of sight to all potential infiltration points and can coordinate defences to contain and neutralize network breaches.

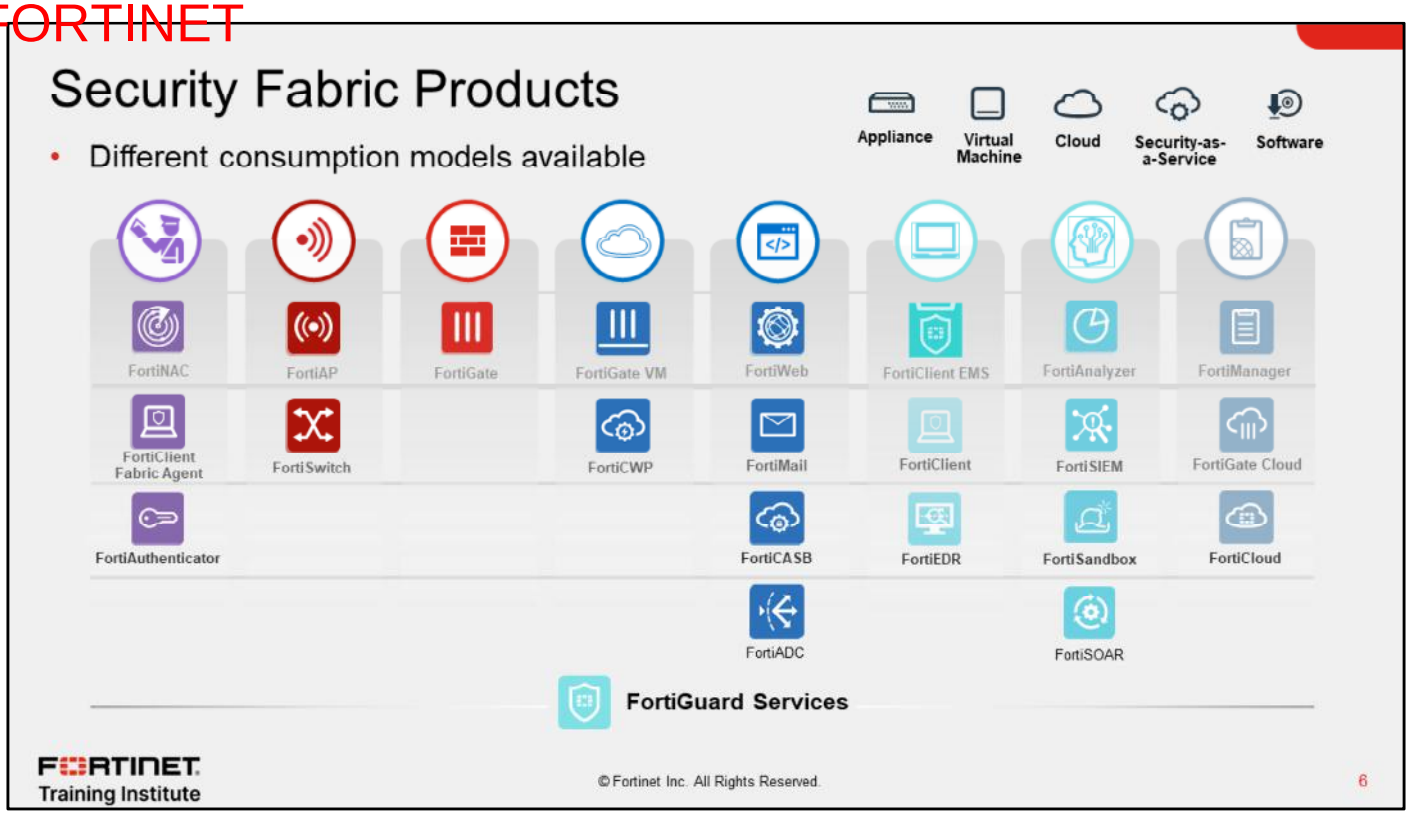

As shown on this slide, the Fortinet Security Fabric offers eight solutions: network access, security WLAN/LAN, public and private cloud infrastructure, applications, endpoint, security operations, open fabric ecosystem, and fabric management center. Each of these solutions is based on specific use cases and involve the integration of specific Fortinet products.

The Fortinet Security Fabric offers network security with FortiGate, IPS, VPN, SD-WAN. It also offers multi-cloud strategy across public clouds, private clouds, hybrid clouds, and software as a service (SaaS). It also offers quite a sophisticated endpoint offering ranging from the Fabric Agent all the way up to full endpoint protection, email security, web application security, secure access across distributed enterprises and SD-WAN environments, advanced threat protection, management and analytics, and security information and event management (SIEM).

All of these are underscored and supported by FortiGuard Services, which deliver AI-powered intelligence and protection across the Security Fabric.

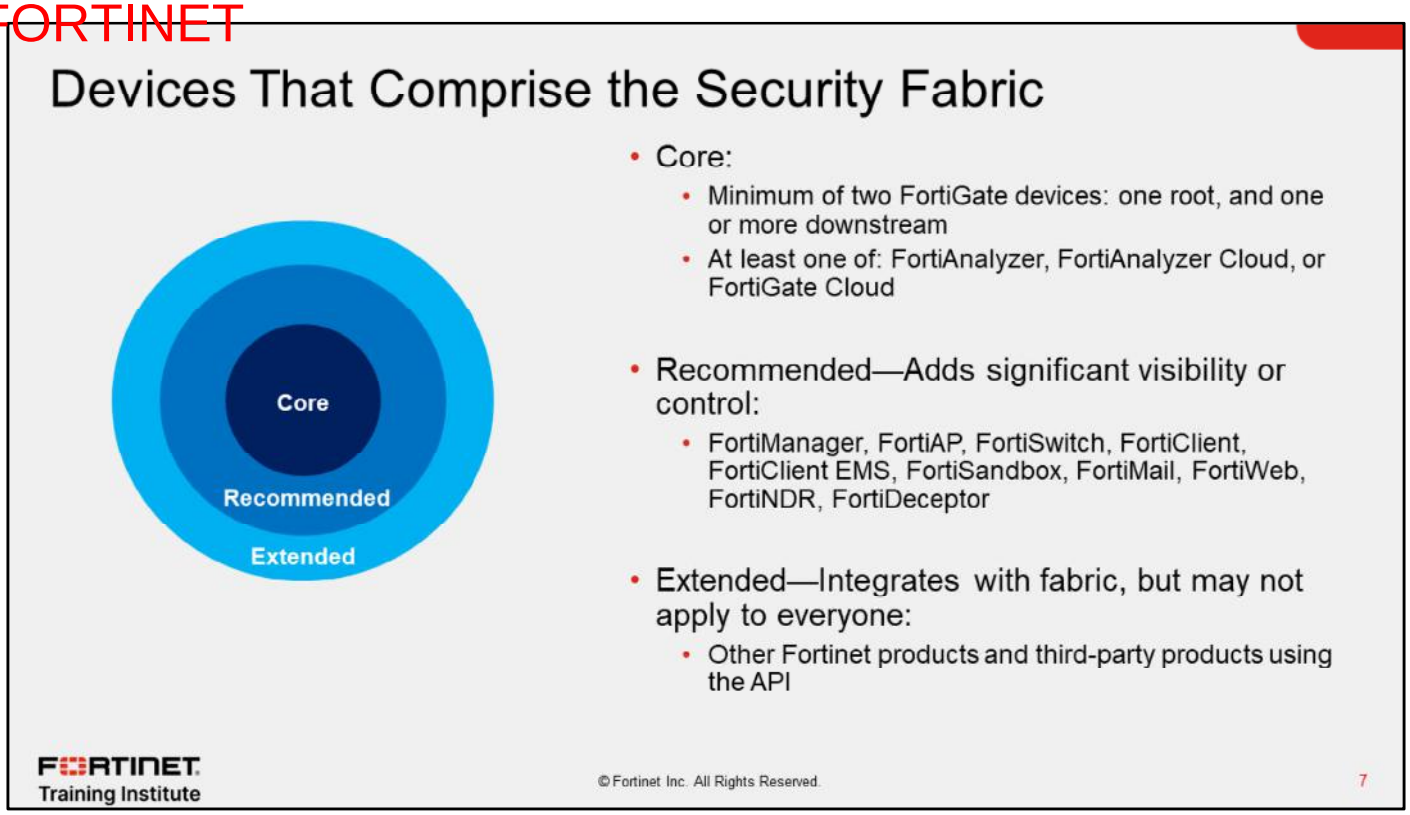

You must have a minimum of two FortiGate devices at the core of the Security Fabric, plus one FortiAnalyzer or cloud logging solution. FortiAnalyzer Cloud or FortiGate Cloud can act as the cloud logging solution. The FortiGate devices must be running in NAT mode.

To add more visibility and control, Fortinet recommends adding FortiManager, FortiAP, FortiClient, FortiClient EMS, FortiSandbox, FortiMail, FortiWeb, FortiNDR, FortiDeceptor, and FortiSwitch.

The solution can be extended by adding other network security devices, including several third-party products.

Knowledge Check

1. What is the Fortinet Security Fabric?
   - ✓ A. A Fortinet solution that enables communication and visibility among devices of your network
   - B. A device that can manage all your firewalls

2. Which combination of devices must participate in the Security Fabric?
   - ✓ A. A FortiAnalyzer and two or more FortiGate devices
   - B. A FortiMail and two or more FortiGate devices

FORTINET Training Institute

8

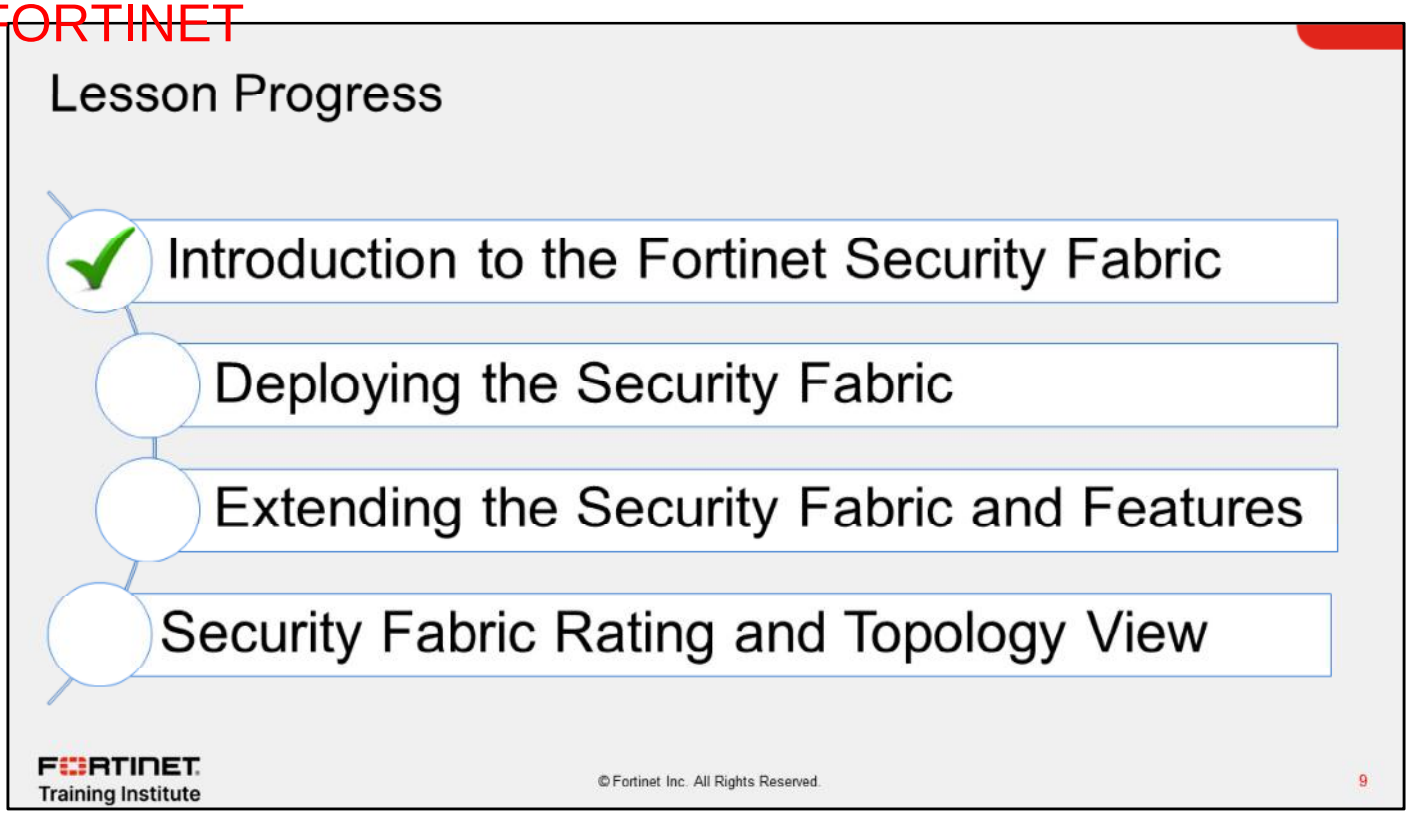

Good job! You now understand the basics of the Fortinet Security Fabric.

Next, you'll learn how to deploy the Security Fabric in your network environment.

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

By demonstrating competence in the deployment of the Fortinet Security Fabric, you will better understand the value of the Security Fabric and how it helps to manage all your network devices more efficiently.

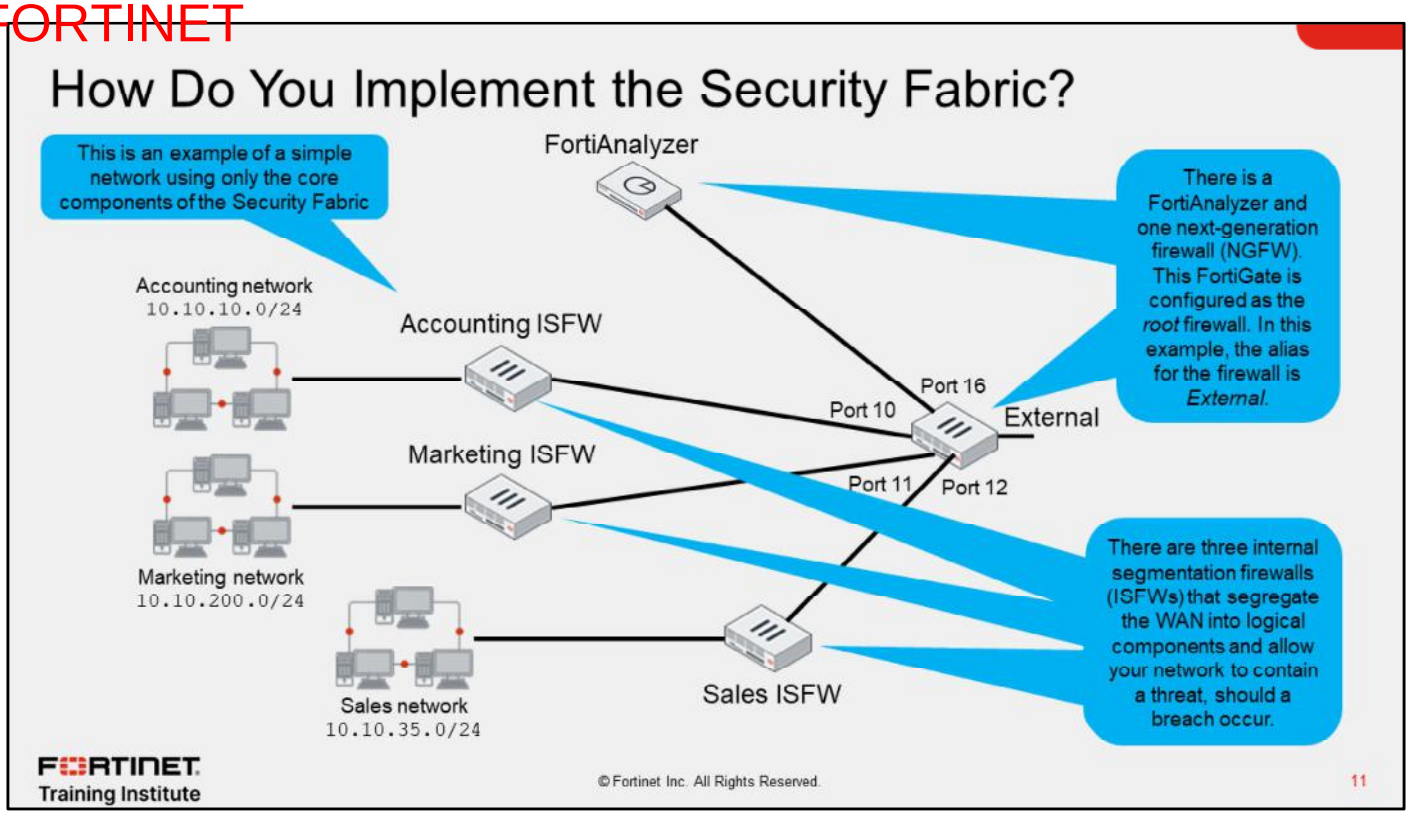

This simple network that comprises only the core devices of a Security Fabric includes one FortiAnalyzer and four next-generation firewall (NGFW) FortiGate devices.

The FortiGate device named External is acting as the edge firewall and is configured as the *root* firewall within the Security Fabric.

Downstream from the root firewall, three internal segmentation firewalls compartmentalize the WAN in order to contain breaches and to control access to various LANs. This example uses Accounting, Marketing, and Sales LANs.

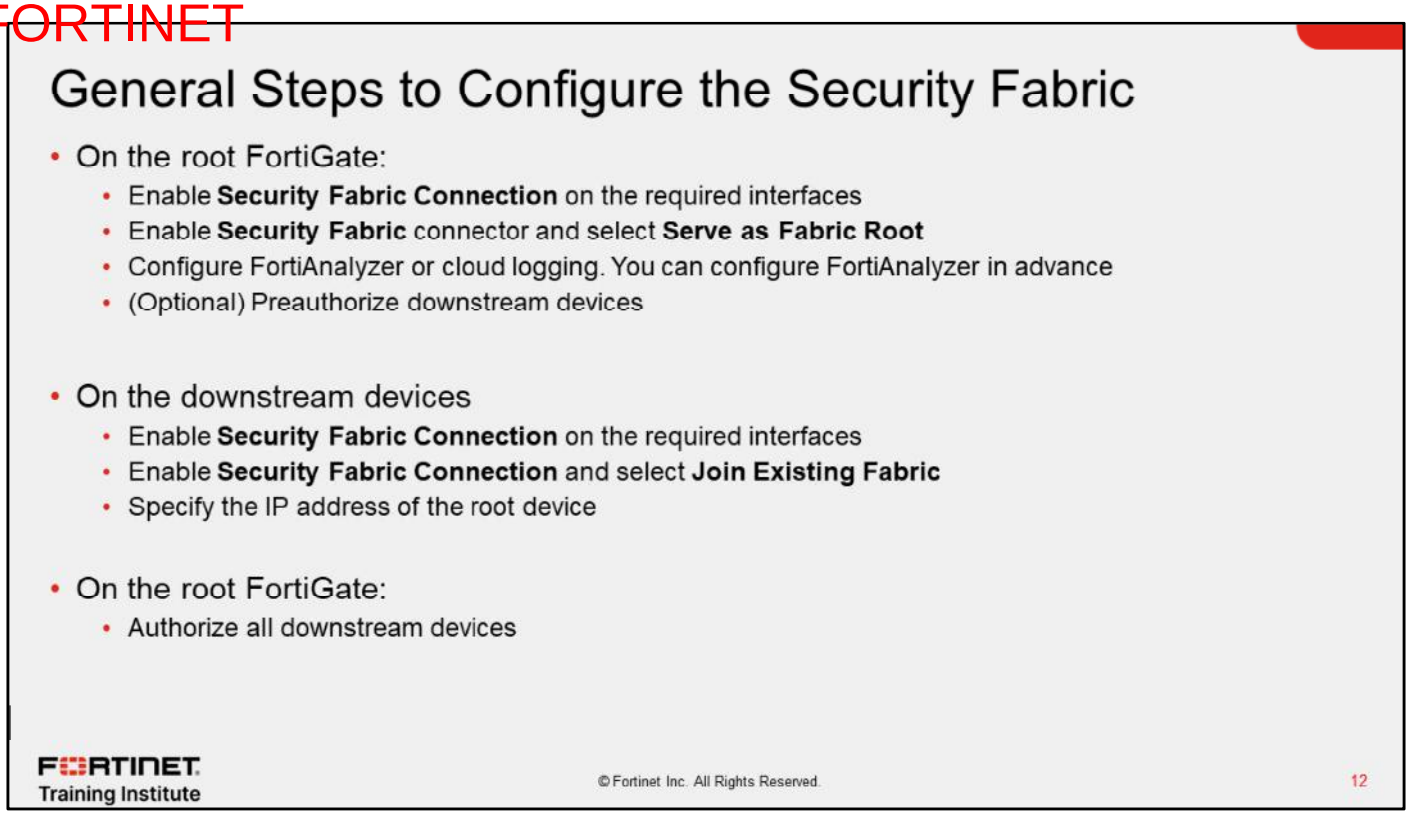

To configure a new security fabric, follow these general steps:

First, on the root FortiGate, you must enable **Security Fabric Connection** on the interfaces that face any downstream FortiGate. Then, enable the Security Fabric connector, and select **Serve as Fabric Root. Y**ou also need to configure FortiAnalyzer or a cloud logging solution. This logging configuration will be pushed to all the downstream FortiGate devices.

Optionally, you can preauthorize your downstream devices by adding their serial numbers. When you add the serial number of a Fortinet device to the trusted list on the root FortiGate, the device can join the Security Fabric as soon as it connects. After you authorize the new FortiGate, additional connected FortiAP and FortiSwitch devices automatically appear in the topology tree. On the topology tree, it's easier for you to authorize them with one click.

The second step in implementing the Security Fabric is configuring the downstream Fortinet devices. On the downstream FortiGate devices, you must enable **Security Fabric Connection** and **Device Detection** on the interfaces facing the downstream FortiGate devices. On the **Fabric Connectors** page, select **Join Existing Fabric** and add the root (upstream) FortiGate IP address.

The third step in implementing the Security Fabric is to authorize the downstream FortiGate devices on the root FortiGate.

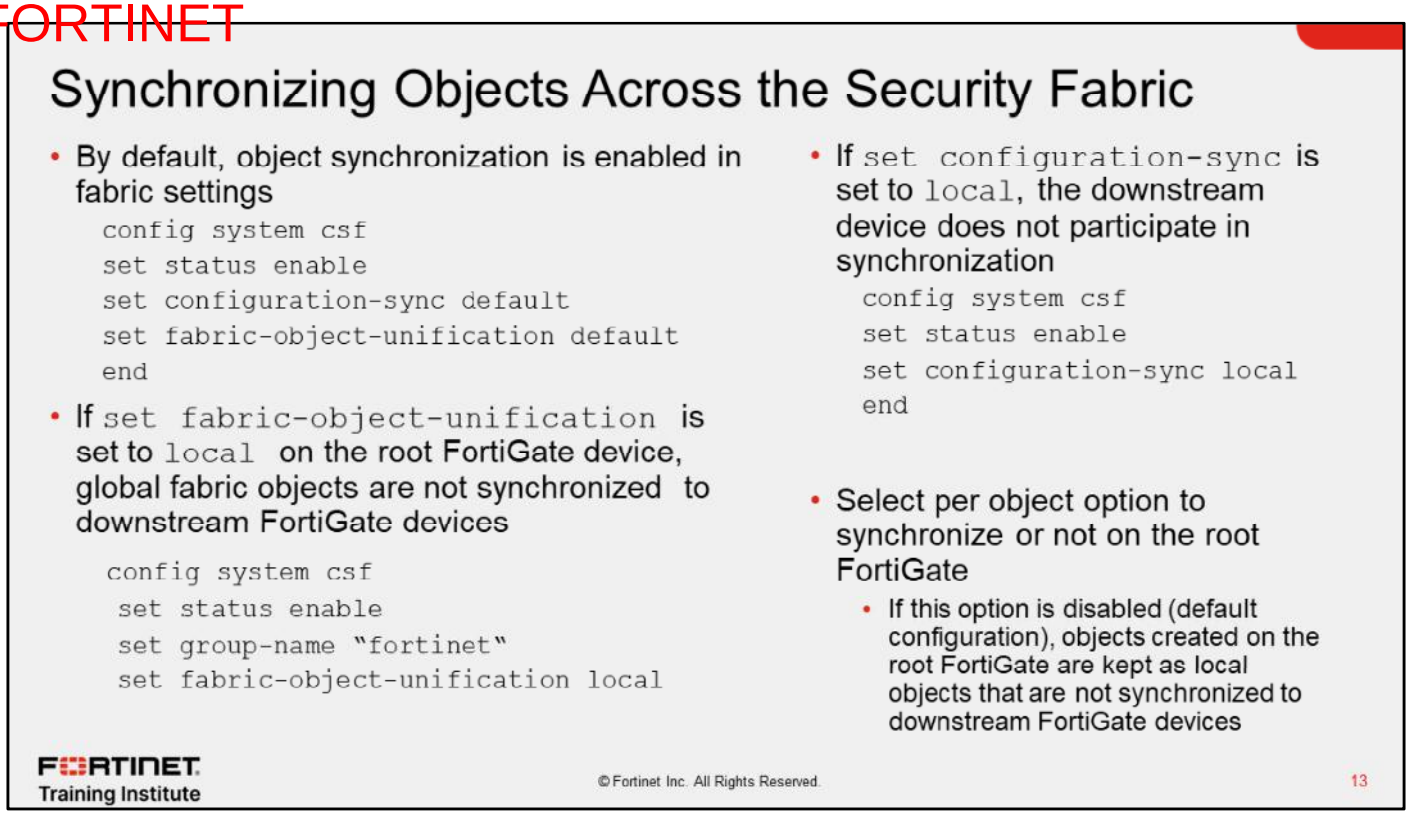

When the Security Fabric is enabled, settings to sync various objects, such as addresses, services, and schedules, from the upstream FortiGate to all downstream FortiGate devices is enabled by default. Synchronization always happens from the root FortiGate to downstream FortiGate devices. Any object that can be synced will be available on downstream FortiGate devices after synchronization.

The CLI command `set fabric-object-unification` is only available on the root FortiGate. When set to `local`, global objects will not be synchronized to downstream devices in the Security Fabric. The default value is `default`.

The CLI command `set configuration-sync local` is used when a downstream FortiGate doesn't need to participate in object synchronization. When set to `local` on a downstream FortiGate, the device does not synchronize objects from the root, but will still participate in sending the synchronized object downstream.

You can also enable or disable per object synchronization in the Security Fabric. This option is not available for objects you create on a downstream FortiGate. Fabric synchronization is disabled by default for supported fabric objects, and these fabric objects are kept as locally created objects on all the FortiGate devices in the Security Fabric. If object synchronization is disabled on the root FortiGate, using the command `set fabric-object disable`, firewall addresses and address groups will not be synchronized to downstream FortiGate devices.

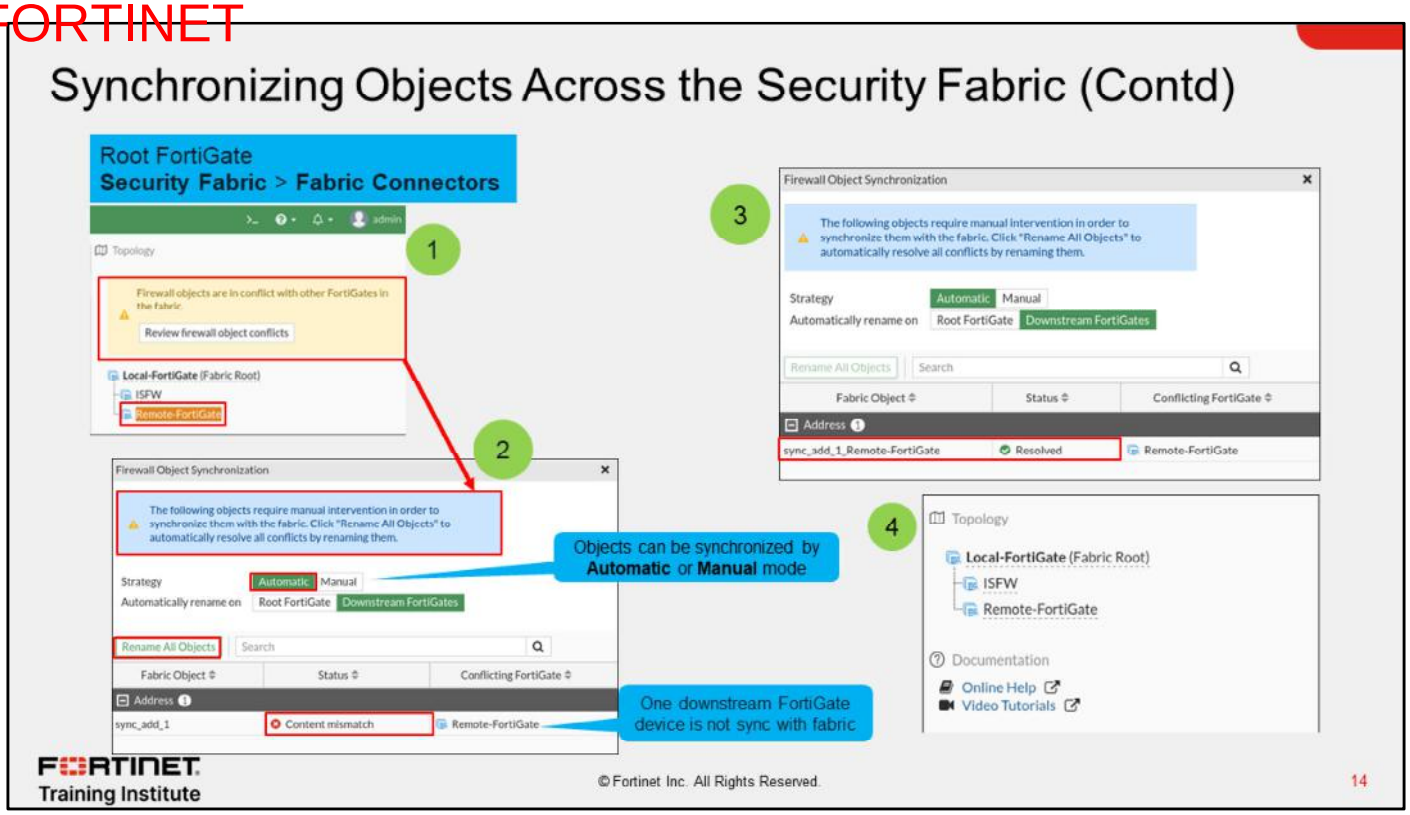

If an object conflict occurs during synchronization, you'll get a notification in the topology tree.

The process to resolve a syncing conflict is as follows:

1. The notification icon displays this message: **Firewall objects are in conflict with other FortiGates in the fabric**. **Remote-FortiGate** is highlighted in amber. Click **Review firewall object conflicts**.

2. On the **Firewall Object Synchronization** page, you can see that both the root FortiGate and downstream FortiGate devices contain the **synn_add_1** object (with a different IP address/subnet schema on each device), causing a status of **Content mismatch**. In the **Strategy** field, there are two options to resolve the conflict: **Automatic** and **Manual**. If you select **Automatic**, as shown in this example, you can then click **Rename All Objects**.

3. **Remote-FortiGate** is appended to the name of the downstream FortiGate device **sync_Add_1** address object and the status changes to **Resolved**.

4. In the topology tree, none of the FortiGate devices are highlighted because there is no conflict.

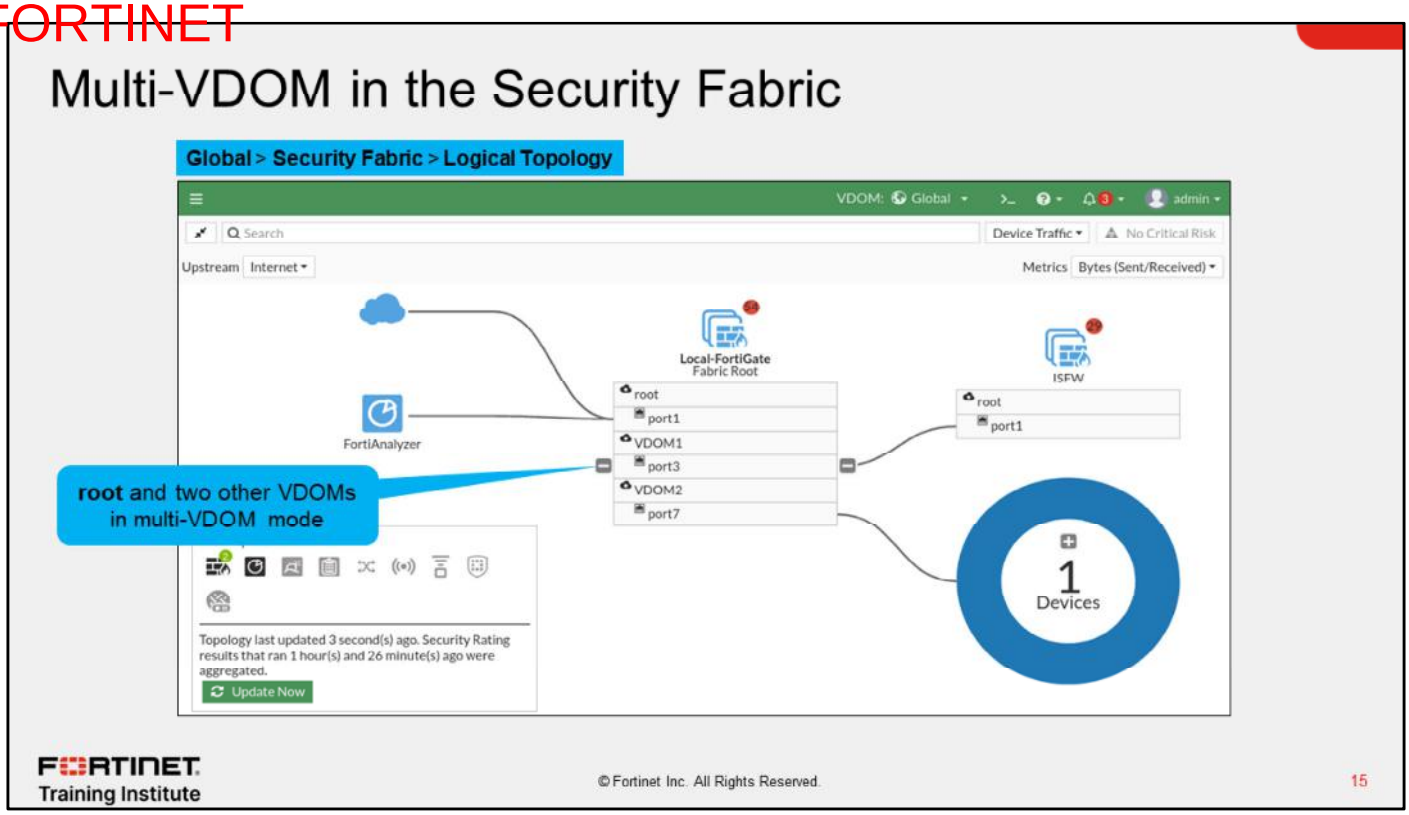

When you configure FortiGate devices in multi-vdom mode and add them to the Security Fabric, each VDOM with its assigned ports is displayed when one or more devices are detected. *Only* the ports with discovered and connected devices appear in the Security Fabric view and, because of this, you must enable **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.

Device Identification–Agentless vs. Agent

Agentless

- Useful feature for the Security Fabric topology view
- Requires direct connectivity to FortiGate
- Detection methods:
  - HTTP user agent
  - TCP fingerprinting
  - MAC address vendor codes
  - DHCP
  - Microsoft Windows browser service (MWBS)
  - SIP user agent
  - Link Layer Discovery Protocol (LLDP)
  - Simple Service Discovery Protocol (SSDP)
  - QUIC
  - FortiOS-VM detection
    - FortiOS-VM vendor ID in IKE messages
    - FortiOS-VM vendor ID in FortiGuard web filter and spam filter requests

Agent (FortiClient)

- Location and infrastructure independent

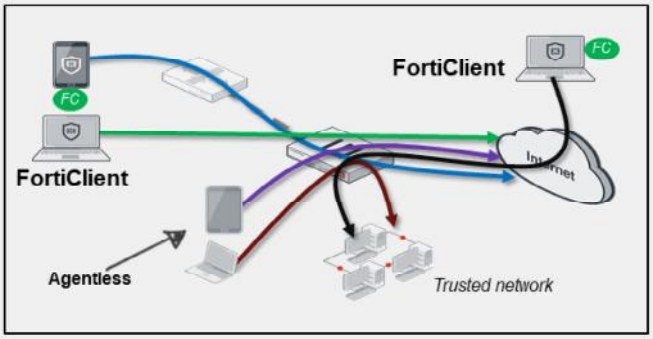

FORTINET Training Institute

16

Device identification is an important component in the Security Fabric. FortiGate detects most of the third-party devices in your network and added into the topology view in the Security Fabric. There are two device identification techniques: with an agent and without an agent (agentless).

Agentless identification uses traffic from the device. Devices are indexed by their MAC address and there are various ways to identify devices, such as HTTP user-Agent header, TCP fingerprint, MAC address OUI, and FortiOS-VM detection methods, to name a few. Agentless device identification is only effective if FortiGate and the workstations are directly connected network segments, where traffic is sent directly to FortiGate, and there is no intermediate router or Layer 3 device between FortiGate and the workstations.

Note that FortiGate uses a *first come, first served* approach to determine the device identity. For example, if a device is detected by the HTTP user agent, FortiGate updates its device table with the detected MAC address and scanning stops as soon as the type has been determined for that MAC address.

Agent-based device identification uses FortiClient. FortiClient sends information to FortiGate, and the device is tracked by its unique FortiClient user ID (UID).

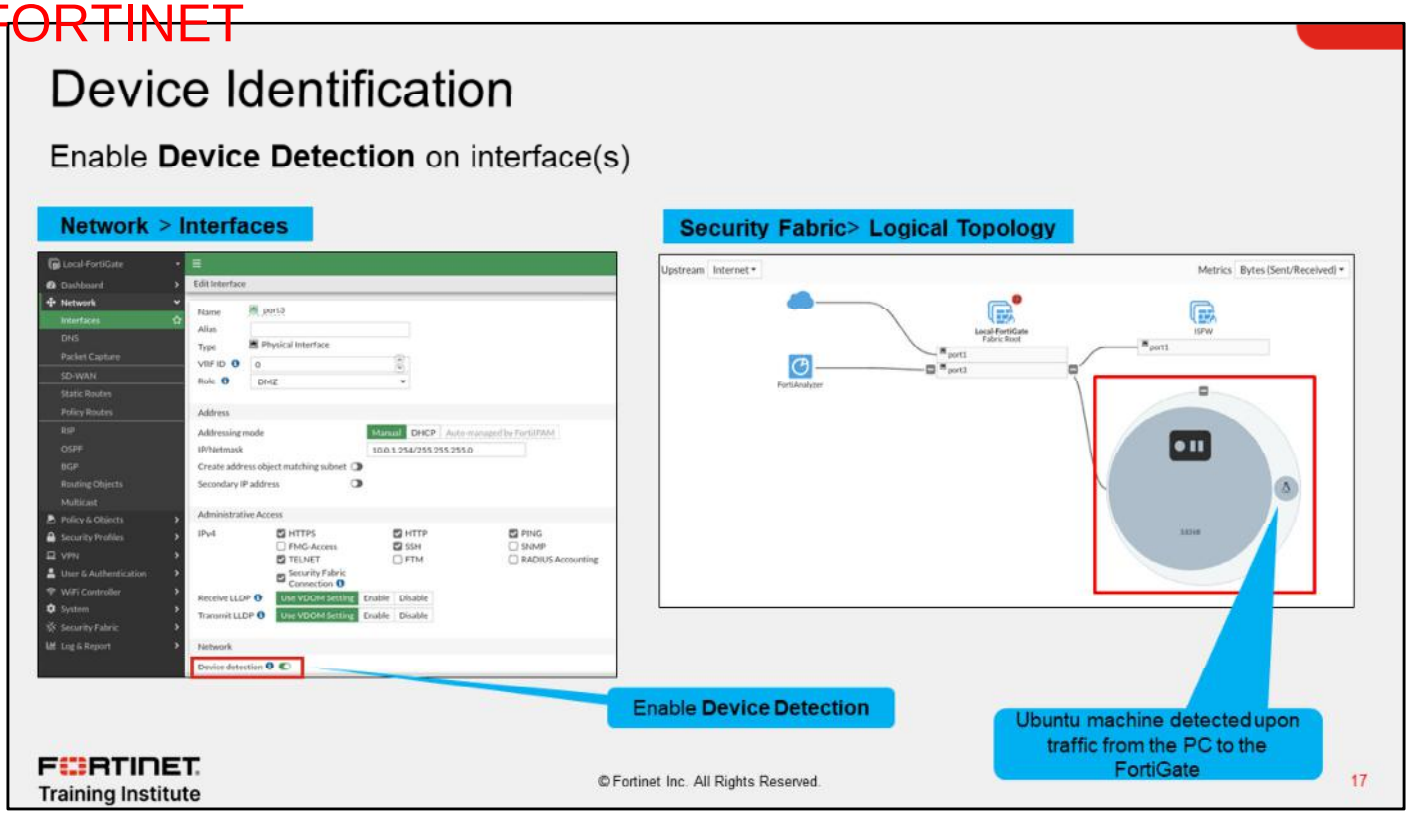

By default, FortiGate uses device detection (passive scanning), which runs scans based on the arrival of traffic.

Knowledge Check

1. What are the two mandatory settings of the Security Fabric configuration?
   - ✓ A. Fabric name and Security Fabric role
   - B. Fabric name and FortiManager IP address

2. From where do you authorize a device to participate in the Security Fabric?
   - A. From the downstream FortiGate
   - ✓ B. From the root FortiGate

FORTINET Training Institute

18

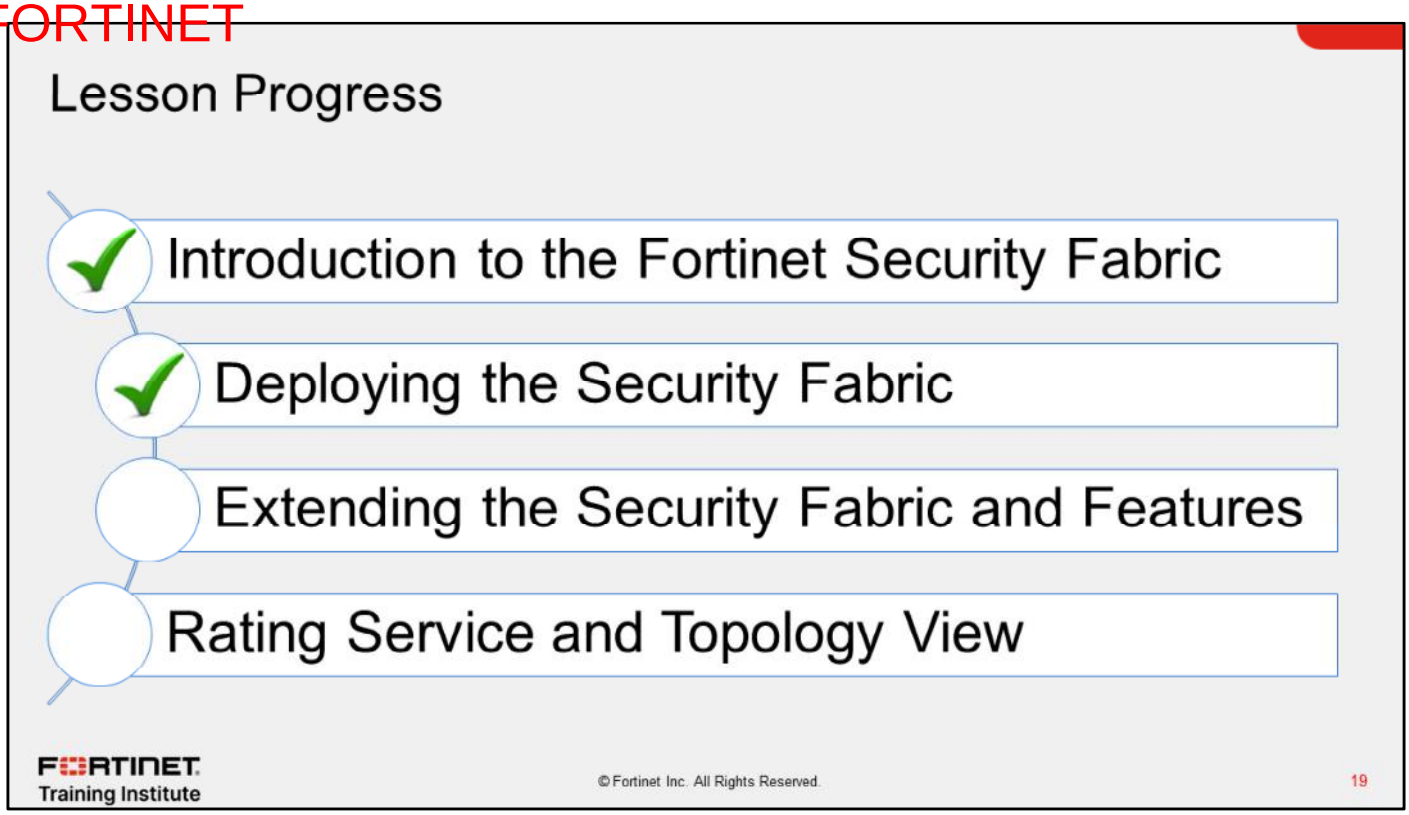

Good job! You now know how to deploy the Security Fabric.

Next, you'll learn about Security Fabric features and how to extend the Security Fabric in your network environment.

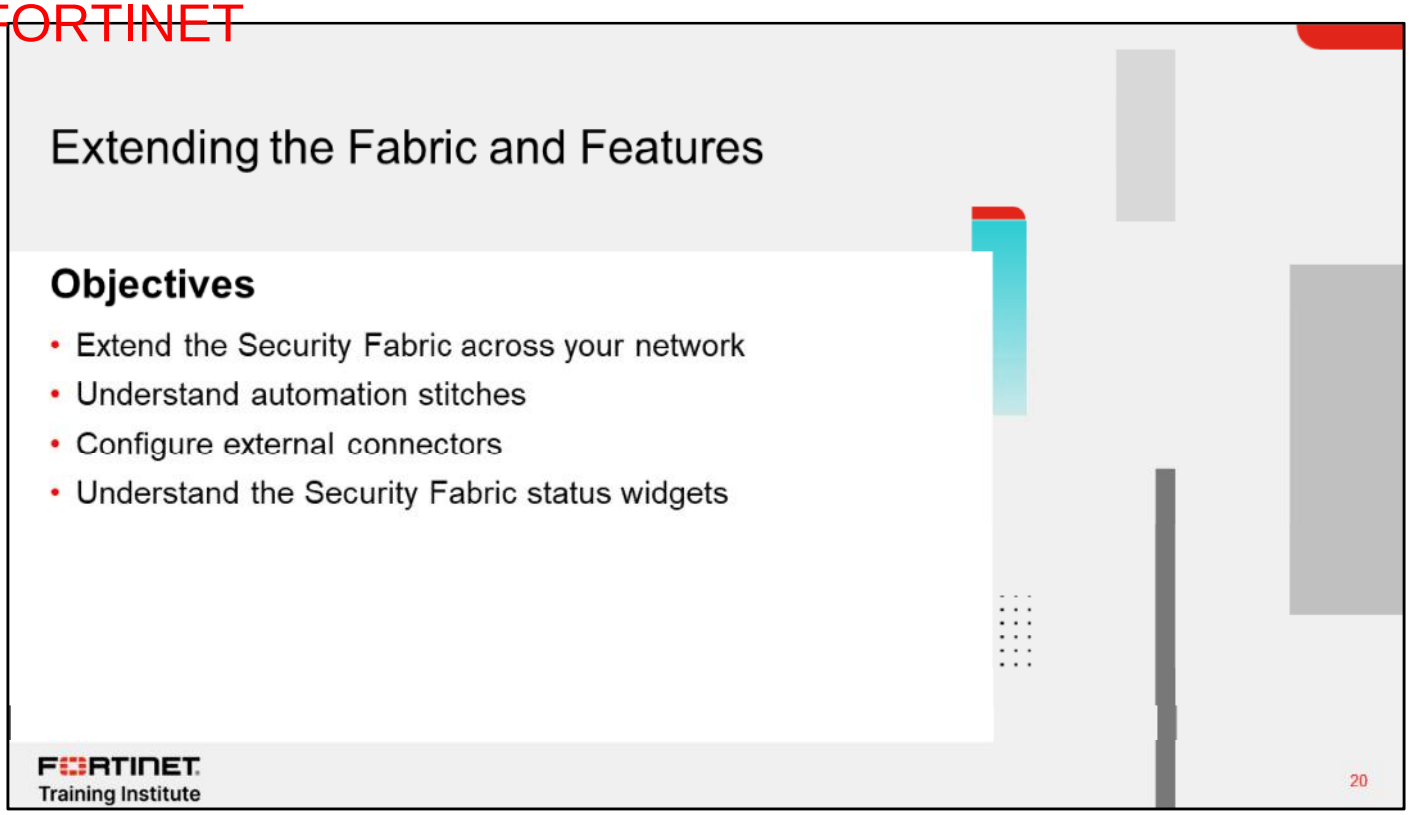

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the extending the Fortinet Security Fabric, you will better understand the value of the Security Fabric and how it helps to manage all your network devices from a single point of device.

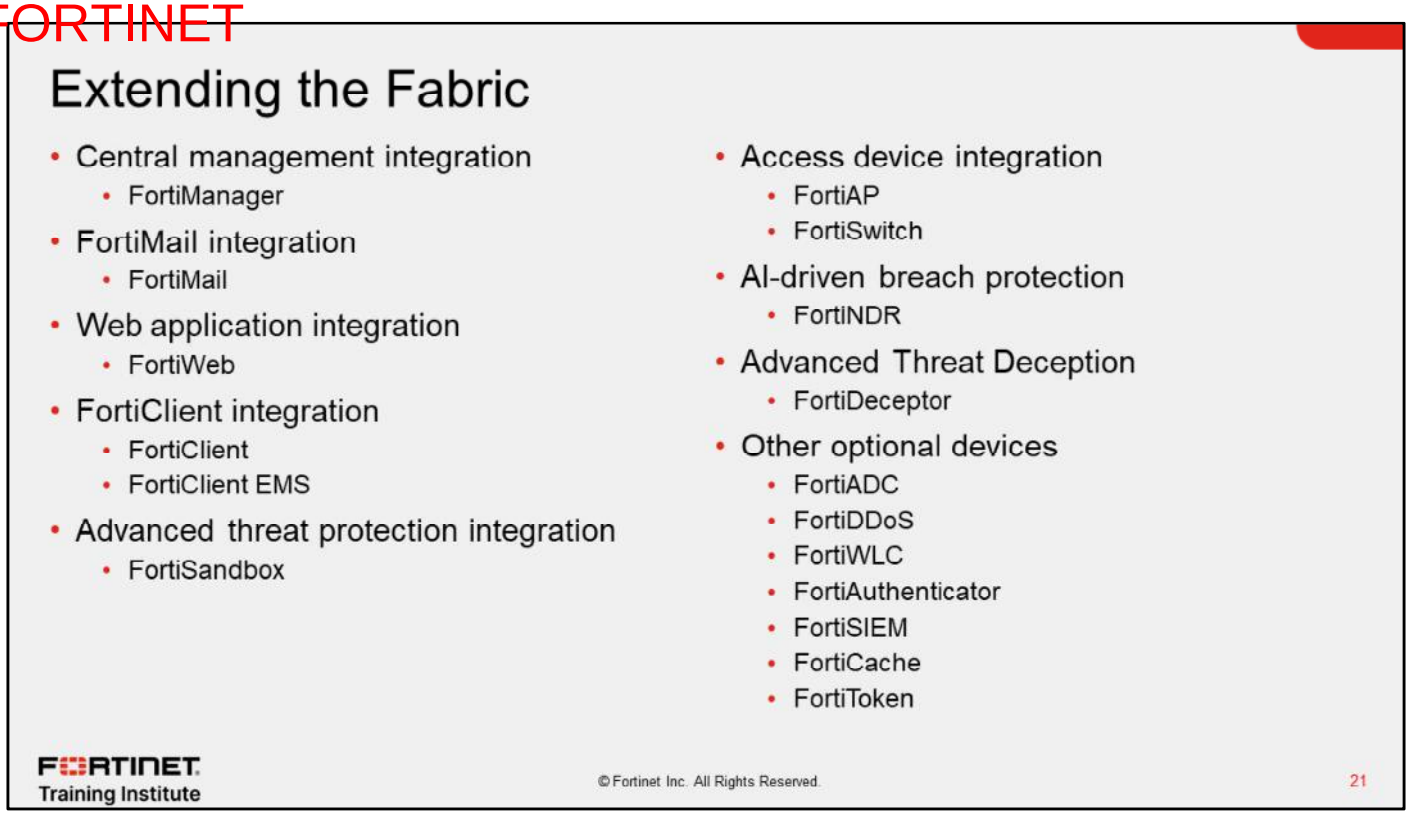

The slide shows the list of products that Fortinet recommends to extend the Security Fabric.

For example, Fortinet recommends using a FortiManager for centralized management of all FortiGate devices and to access devices in the Security Fabric. You can also extend the Security Fabric down to the access layer by integrating FortiSwitch and FortiAP devices.

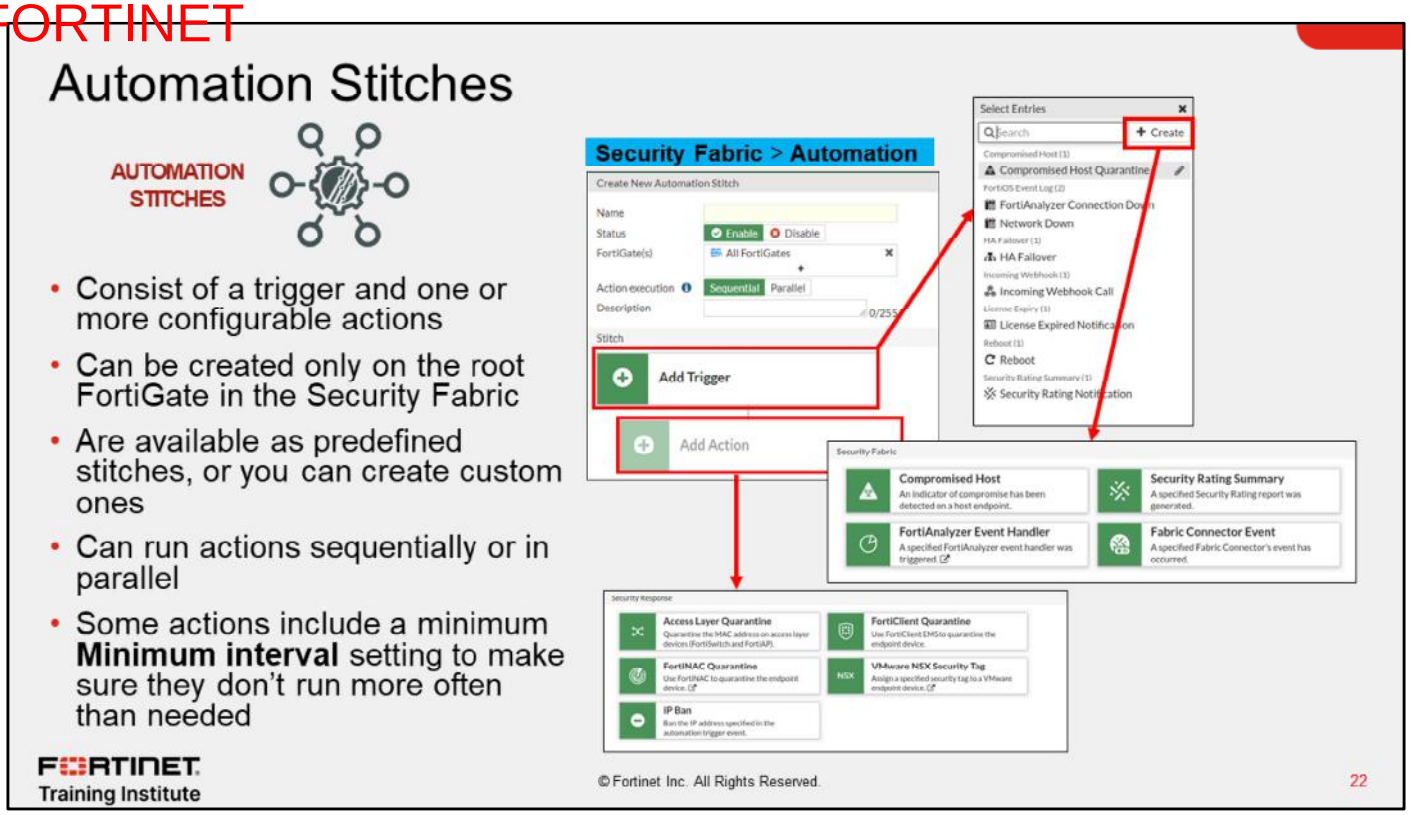

Administrator-defined automated workflows (called stitches) cause FortiOS to automatically respond to an event in a preprogrammed way. Because this workflow is part of the Security Fabric, you can set up automation stitches for any device in the Security Fabric. However, the Security Fabric is not required to use stitches.

Each automation stitch pairs a trigger and one or more actions. FortiOS has several predefined stitches, triggers and actions. However, you can create custom automation based on the available options.

Automation stitches allow you to monitor your network and take appropriate action when the Security Fabric detects a threat. You can use automation stitches to detect events from any source in the Security Fabric and apply actions to any destination.

You can configure the **Minimum internal (seconds)** setting on some of the available actions to make sure they don't run more often than needed.

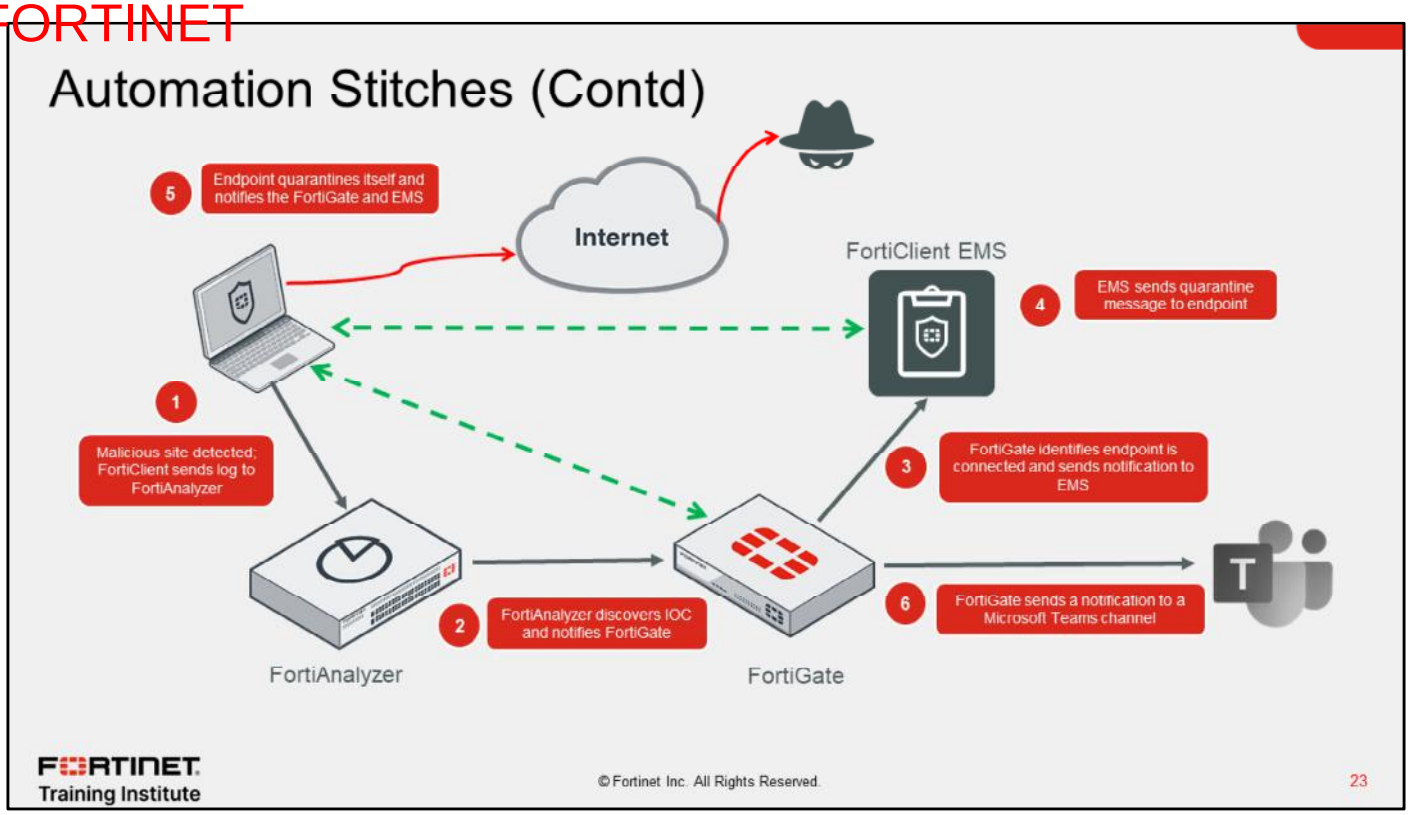

This slide shows an example of how automation stitches can be configured to work in the Security Fabric:

1. FortiClient sends logs to FortiAnalyzer.
2. FortiAnalyzer discovers IoCs in the logs and notifies FortiGate.
3. FortiGate identifies whether FortiClient is a connected endpoint, and whether it has the login credentials for the FortiClient EMS that FortiClient is connected to. With this information, FortiGate sends a notification to FortiClient EMS to quarantine the endpoint.
4. FortiClient EMS searches for the endpoint and sends a quarantine message to it.
5. The endpoint receives the quarantine message and quarantines itself, blocking all network traffic. The endpoint notifies FortiGate and EMS of the status change.
6. FortiGate sends a notification to a Microsoft Teams channel to alert the administrators about the event.

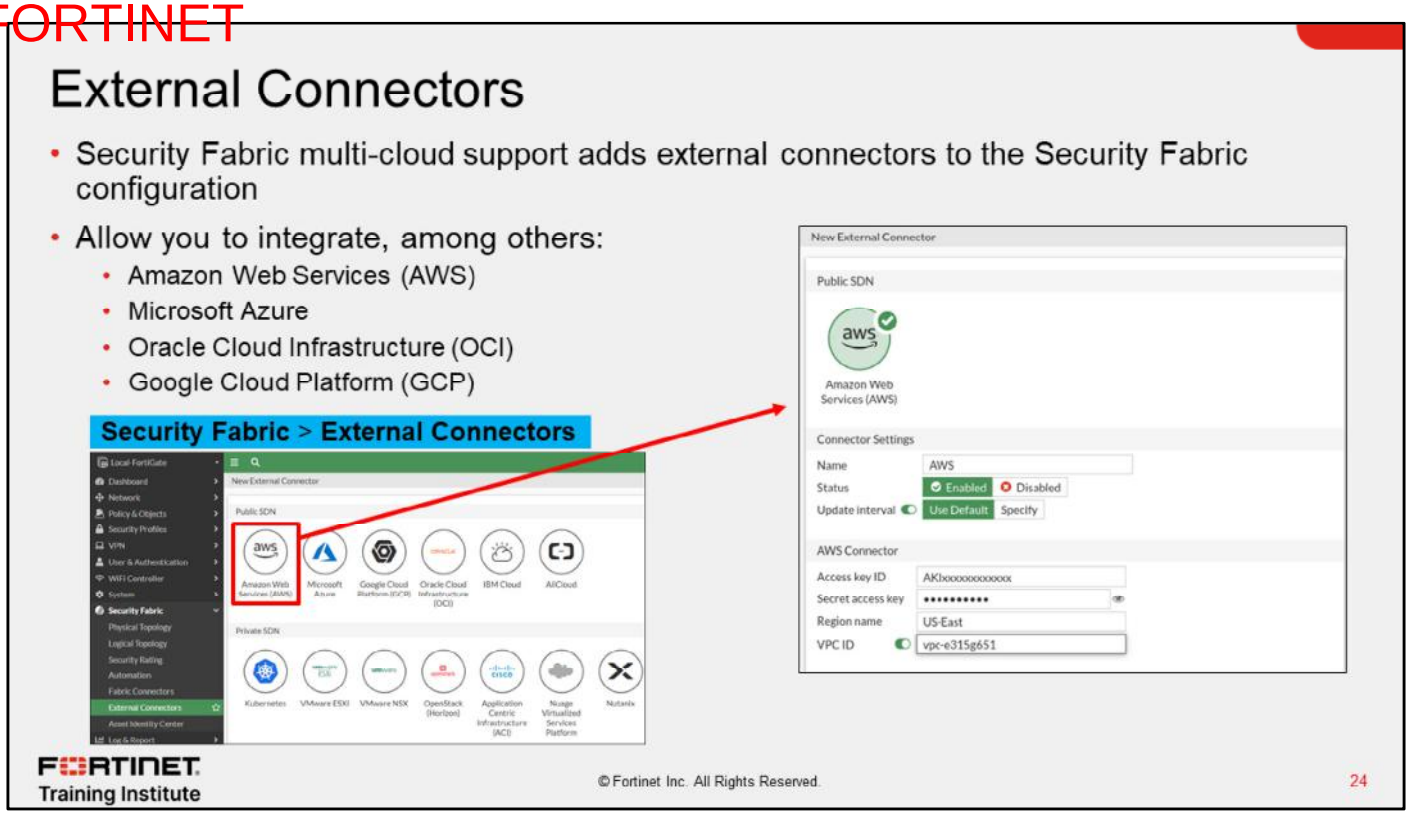

External connectors allow you to integrate multi-cloud support, such as Microsoft Azure and AWS, among others.

In an application-centric infrastructure (ACI), the SDN connector serves as a gateway bridging SDN controllers and FortiGate devices. For example, the SDN connector can register itself to APIC in the Cisco ACI fabric, polls objects of interest, and translates them into address objects. The translated address objects and associated endpoints populate on FortiGate.

The Security Fabric Status Widget

- The name of your Security Fabric
- Icons indicating the other devices in the Security Fabric
- The names of the FortiGate devices in the Security Fabric

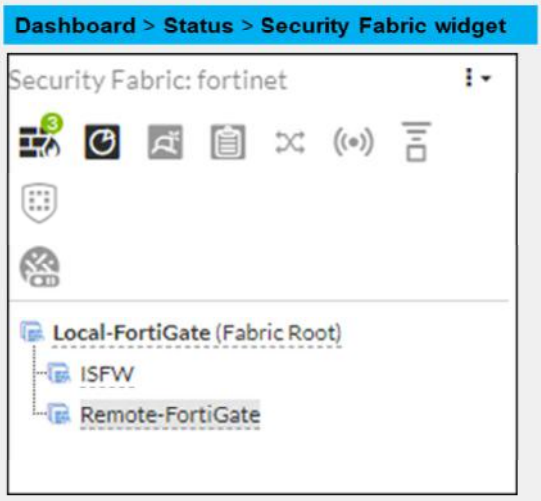

FORTINET Training Institute

25

The **Security Fabric Status** widget shows a visual summary of the devices in the Security Fabric.

You can hover over the icons at the top of the widget to display a quick view of their statuses. From here, you can click to authorize FortiAP and FortiSwitch devices that are connected to an authorized FortiGate.

Icons represent the other Fortinet devices that can be used in the Security Fabric:

- Devices in blue are connected in your network.
- Devices in gray are not configured, or not detected in your network.
- Devices in red are no longer connected, or not authorized in your network.

Knowledge Check

1. Why should an administrator extend the Security Fabric to other devices?
   - ✓ A. To provide a single pane of glass for management and reporting purposes
   - B. To eliminate the need to purchase licenses for FortiGate devices in the Security Fabric

2. What is the purpose of Security Fabric external connectors?
   - ✓ A. External connectors allow you to integrate multi-cloud support with the Security Fabric
   - B. External connectors allow you to connect the FortiGate command line interface (CLI)

FORTINET Training Institute

26

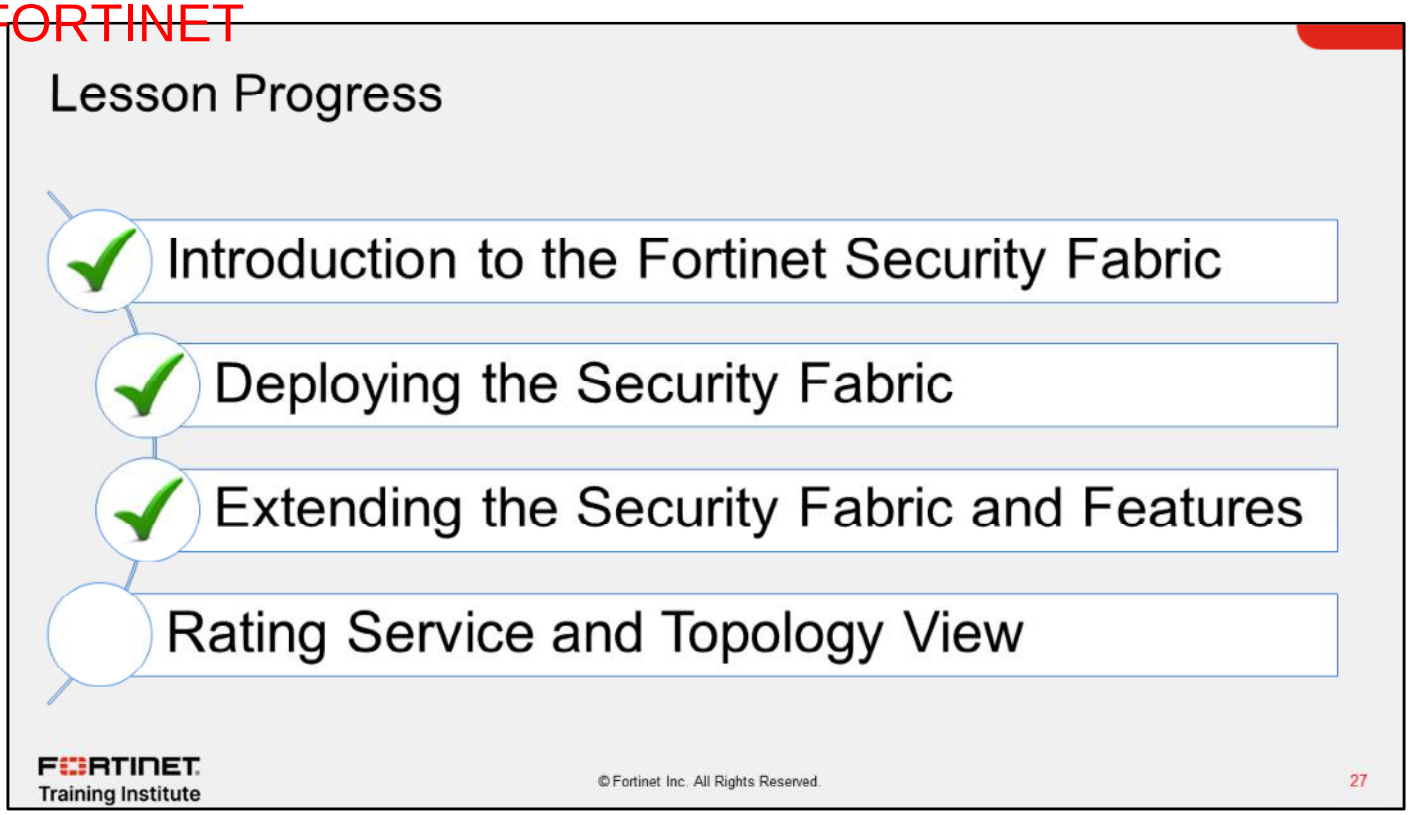

Good job! You now know how to extend the Security Fabric and its features.

Next, you'll learn about the Security Fabric Rating service and topology view.

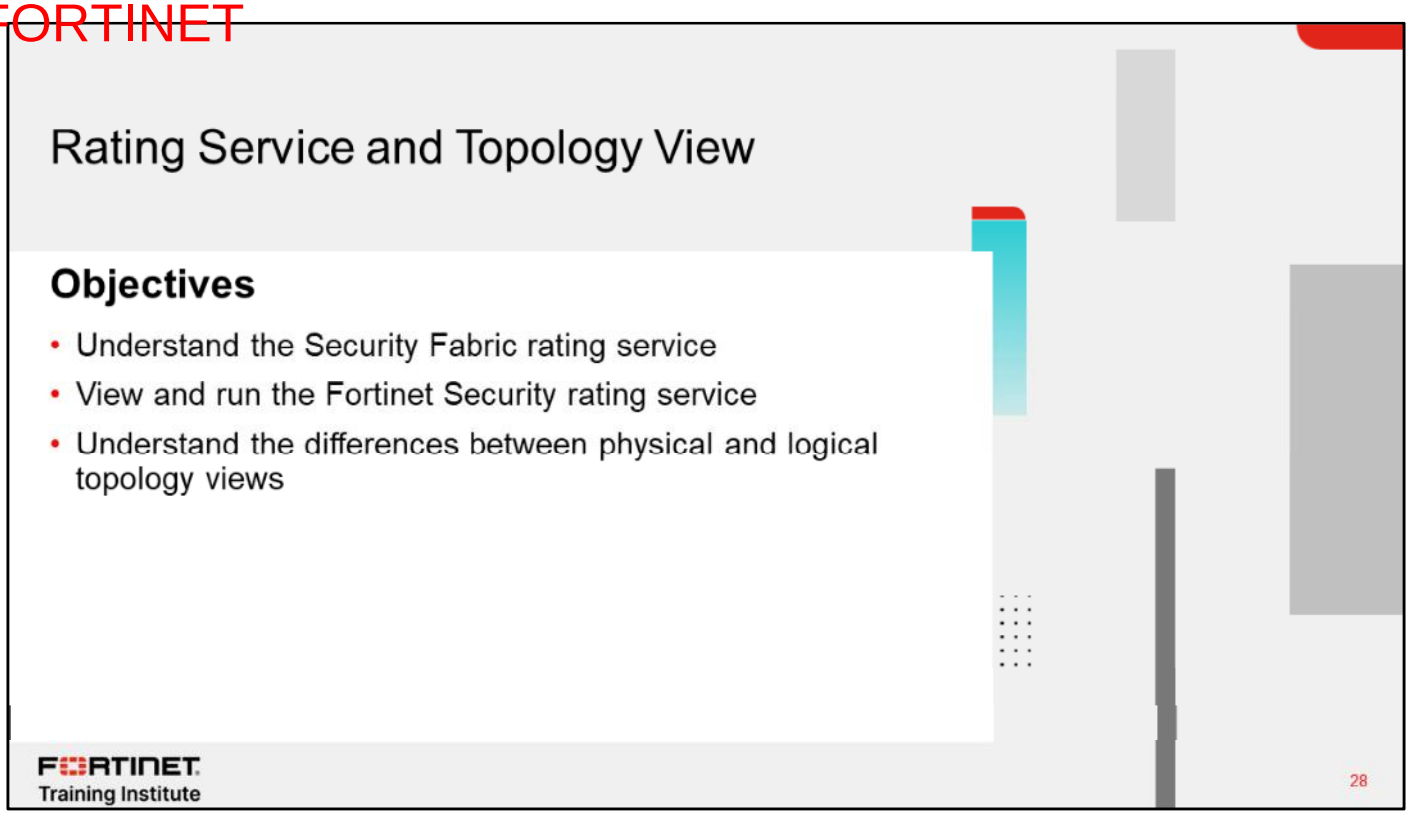

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the Fortinet Security rating service and topology views, you should be able to have clear visibility of your network devices.

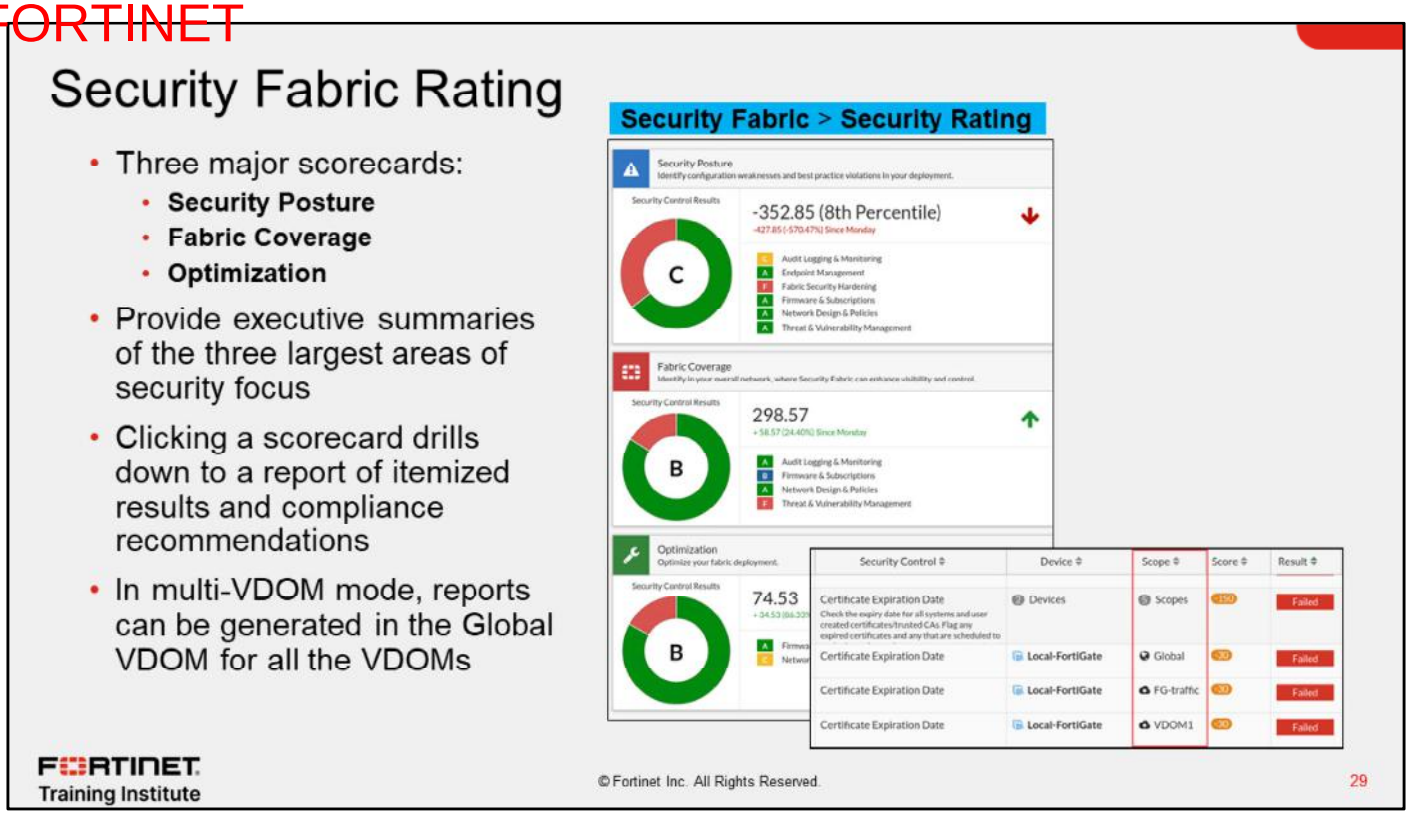

Security rating is a subscription service that requires a security rating license. This service provides the ability to perform many *best practices*, including password checks, to audit and strengthen your network security.

The **Security Rating** page is separated into three major scorecards: **Security Posture, Fabric Coverage**, and **Optimization**.

These scorecards provide executive summaries of the three largest areas of security focus in the Security Fabric.

The scorecards show an overall letter grade and breakdown of the performance in subcategories. Click a scorecard to drill down to a detailed report of itemized results and compliance recommendations. The point score represents all passed and failed items in that area. The report includes the security controls that were tested, linking them to specific FSBP or PCI compliance policies. You can click **FSBP** and **PCI** to reference the corresponding standard.

In multi-VDOM mode, administrators with read/write access can generate security rating reports in the Global VDOM for all the VDOMs on the device. Administrators with read-only access can view the report, but not generate it.

On the scorecards, the **Scope** column shows the VDOM or VDOMs that the security rating checked. On checks that support **Easy Apply**, you can run the remediation on all the associated VDOMs.

The security rating event log is available on the root VDOM.

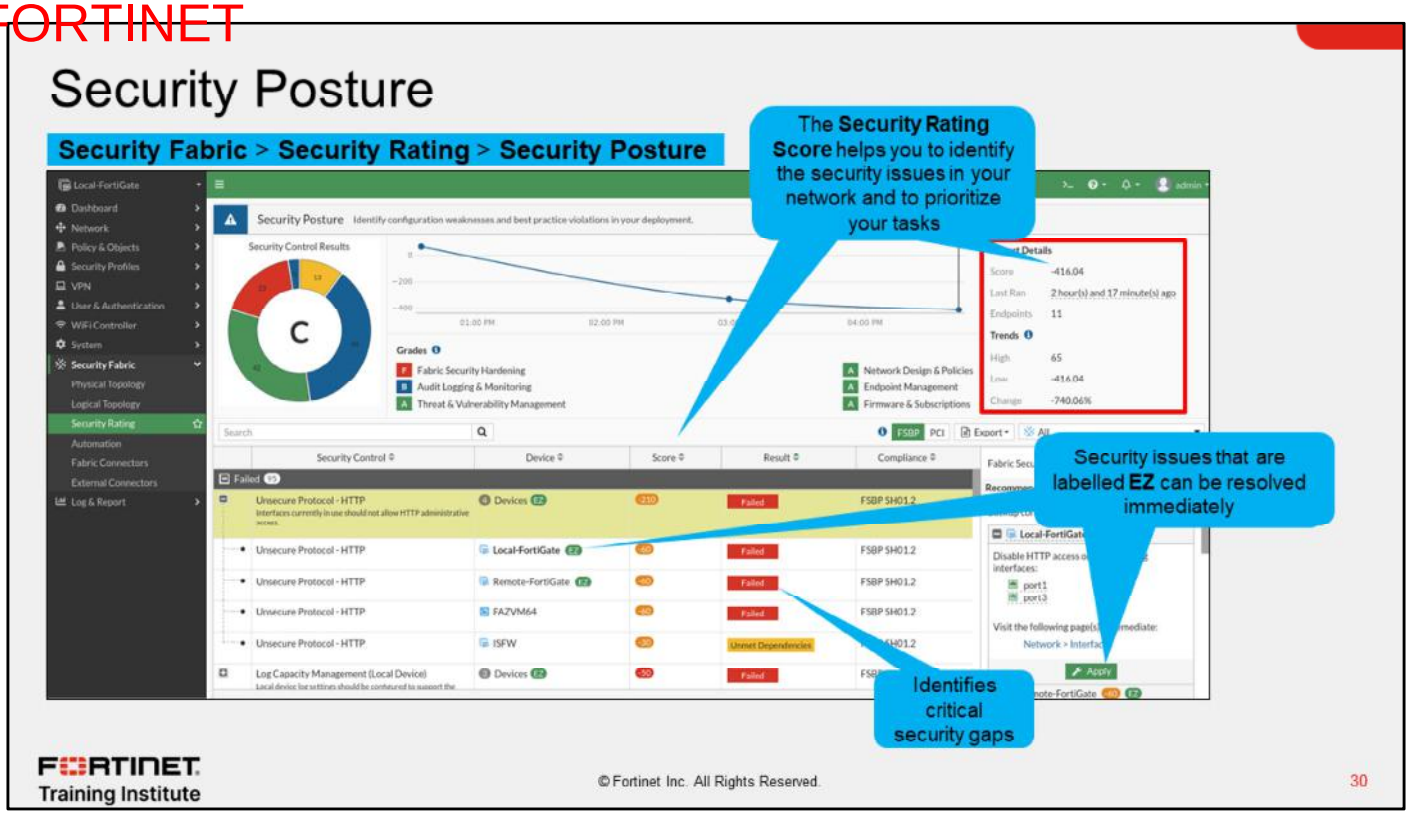

Click the **Security Posture** scorecard on the **Security Rating** page to expand the scorecard and see more details.

The security posture service now supports the following:

- Customer rankings by percentile using security audit (FortiGuard data): Security rating now supports sending results to FortiGuard, and receiving statistics from FortiGuard. Results are displayed to customer in the form of percentile.
- Security audits running in the background, not just on demand, when an administrator is logged in to the GUI. When you view the security audit page, the latest saved security audit data is loaded. From the GUI, you can run audits on demand and view results for different devices in the Security Fabric. You can also view all results or just failed test results.
- New security checks that can help you make improvements to your organization's network. These results include enforcing password security, applying recommended login attempt thresholds, encouraging two-factor authentication, and more.

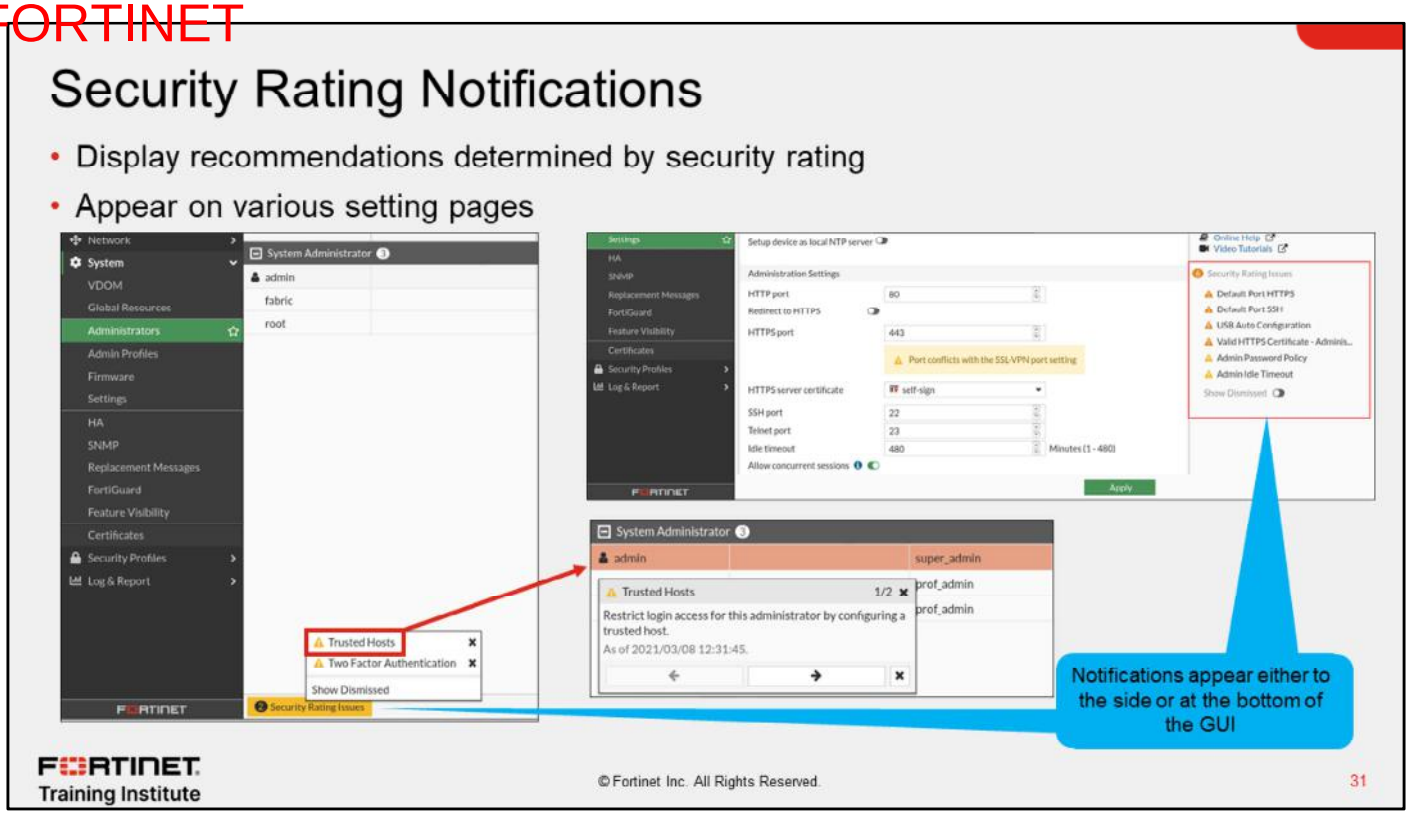

Security rating provides recommendations and highlights issues with the configuration of the FortiGate settings. These recommendations and issues appear as notifications on the **Settings** page.

Click a notification to display the page where the setting needs to be fixed. This prevents you from having to go back and forth between the **Security Fabric** > **Security Rating** page and the various settings pages.

Notifications appear either to the side or at the bottom of the GUI. You can also dismiss the notifications.

In the example shown on this slide, some of the issues found are that FortiGate is using the default HTTPS and SSH ports, and that the administrator password policy is not enabled. The security rating check also recommends that you configure trusted hosts and two-factor authentication.

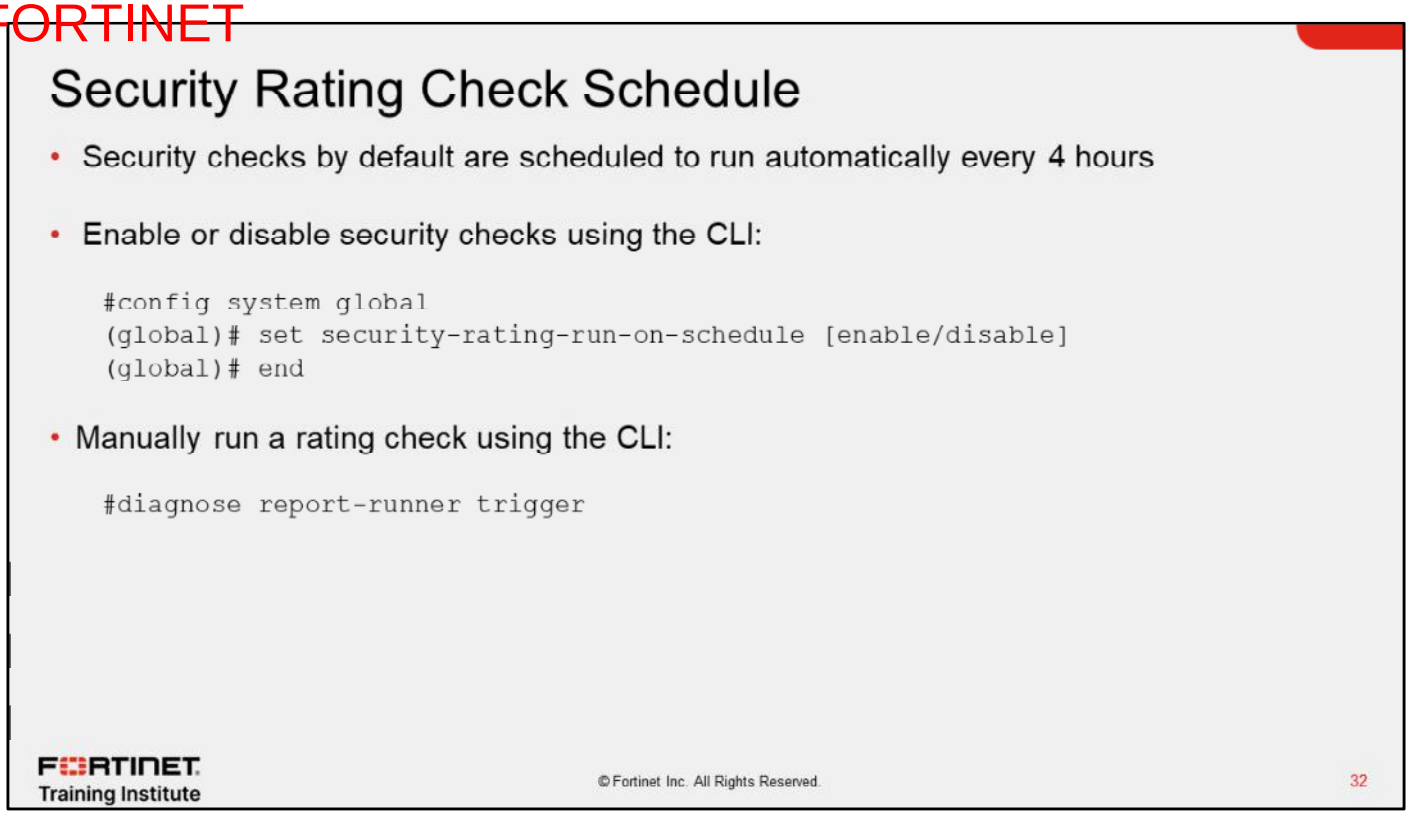

Security rating checks by default are scheduled to run automatically every four hours.

Use the following commands to enable or disable security checks using the CLI:

```
#config system global
(global)# set security-rating-run-on-schedule [enable/disable]
(global)# end
```

Use the following command to manually run a rating check using the CLI:

```
#diagnose report-runner trigger
```

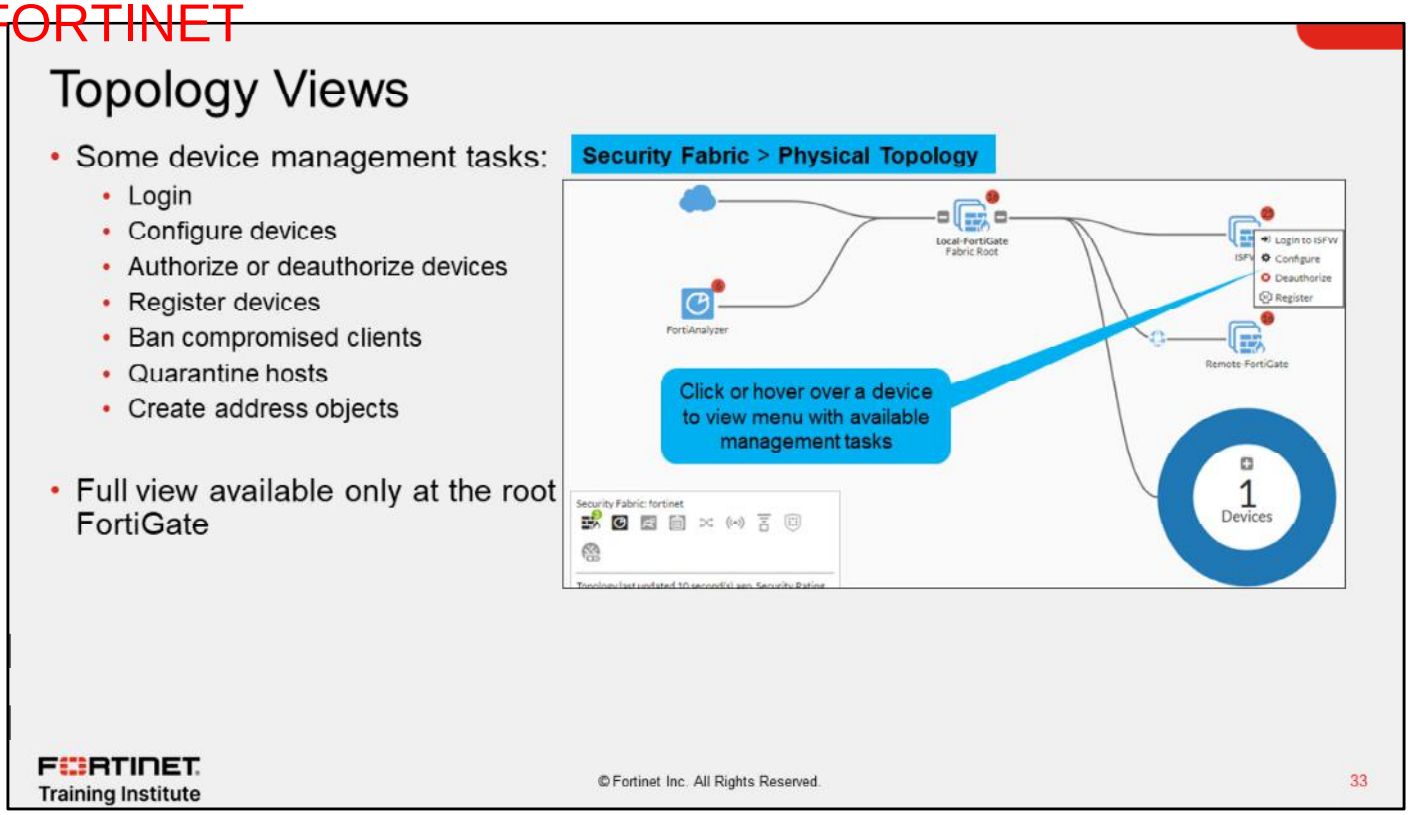

You can view the Security Fabric topology on the FortiGate GUI, from the **Security Fabric** menu. You can select the **Physical Topology** or **Logical Topology** view. To view the complete network, you must access the topology views on the root FortiGate in the Security Fabric.

The **Physical Topology** view displays your network as a bubble chart of interconnected devices. These devices are grouped based on the upstream device they are connected to. The bubbles appear smaller or larger, based on their traffic volume. You can double-click any bubble to resize it and view more information about the device.

The **Logical Topology** view is similar to the **Physical Topology** view, but it shows the network interfaces, logical or physical, that are used to connect devices in the Security Fabric.

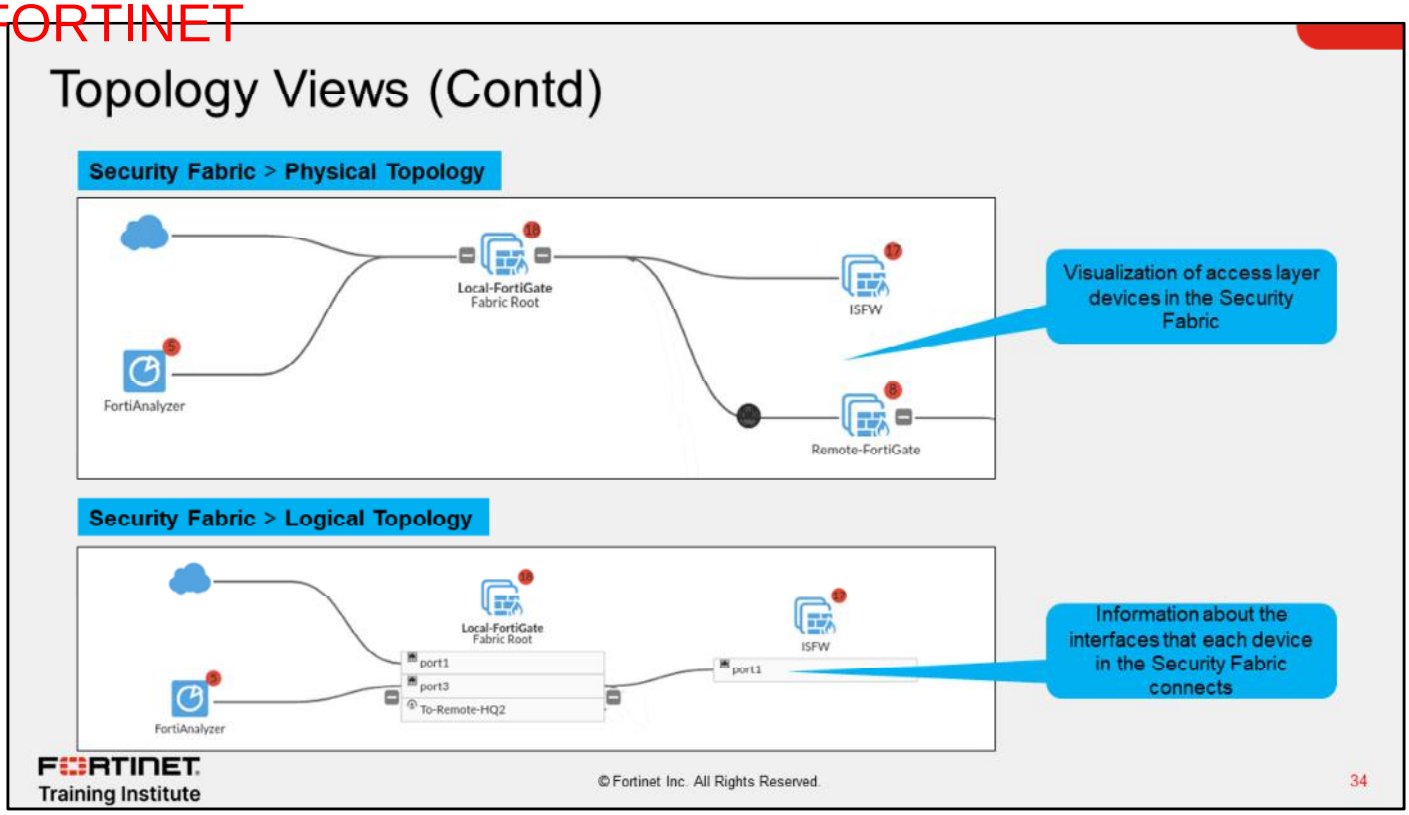

This slide shows the difference between the **Physical Topology** view and the **Logical Topology** view.

Knowledge Check

1. Which one is a part of the Security Rating scorecard?
   - A. Firewall Policy
   - ✓ B. Optimization

2. From which view can an administrator deauthorize a device from the Security Fabric?
   - ✓ A. From the physical topology view
   - B. From the FortiView

FORTINET Training Institute

35

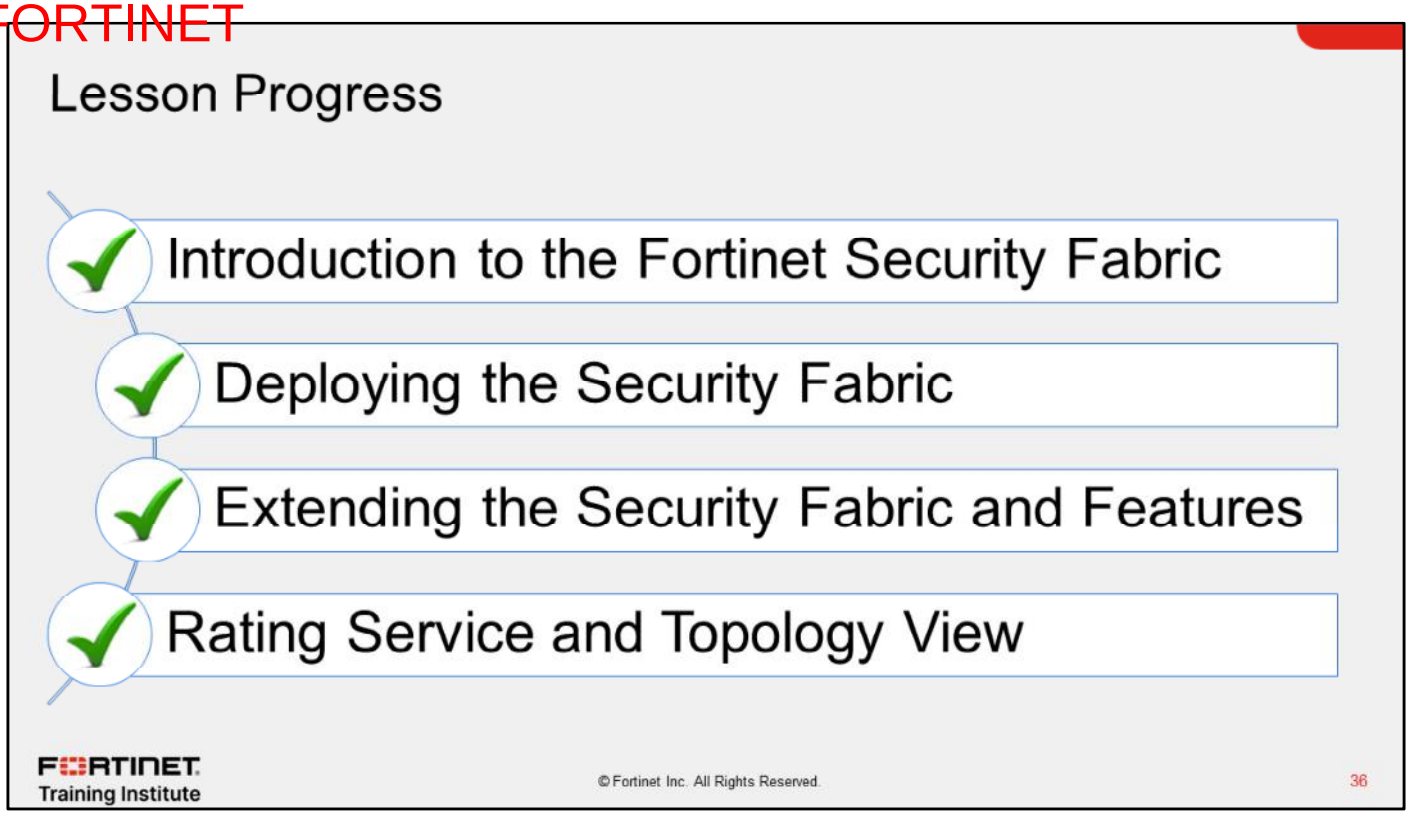

Congratulations! You have completed this lesson.

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

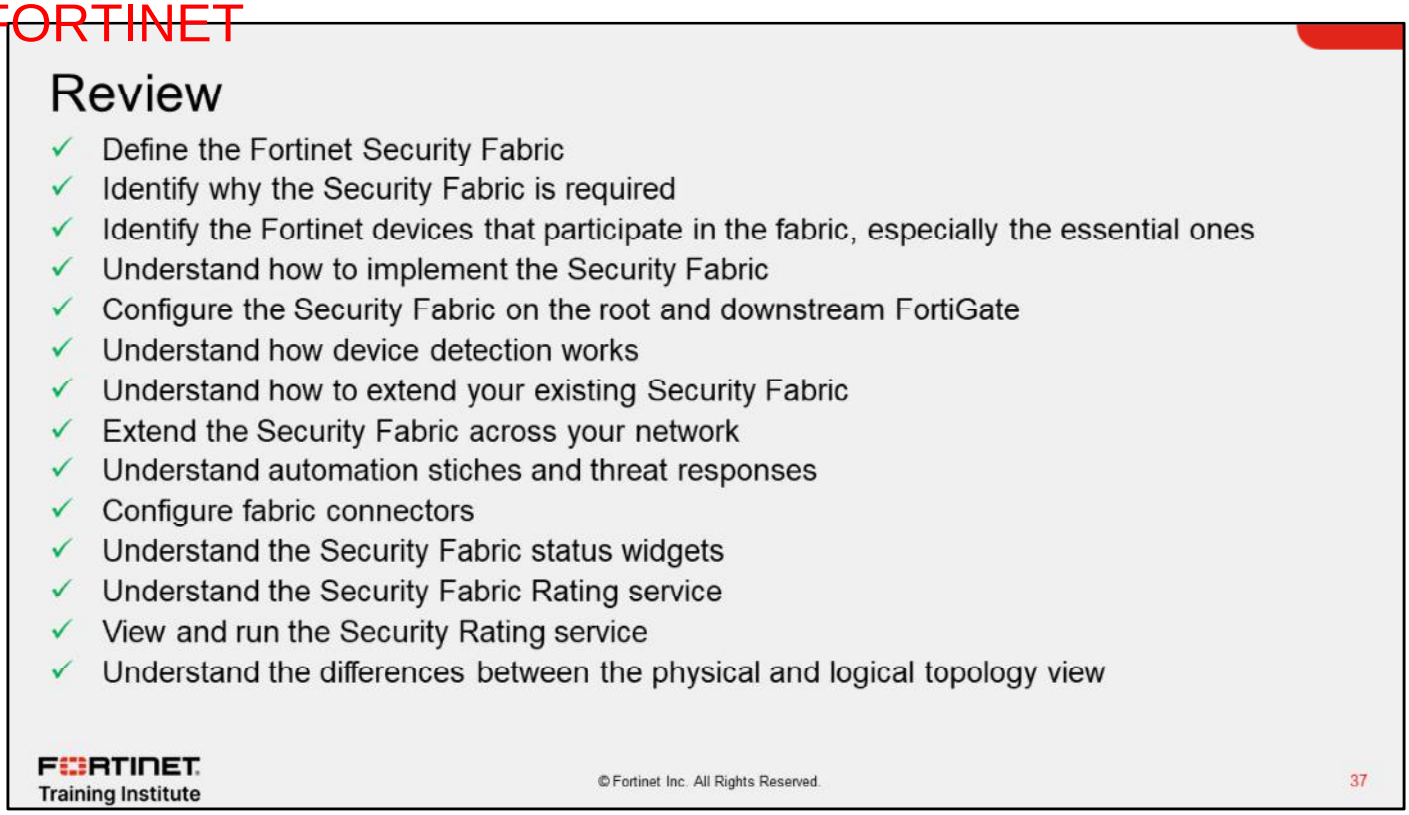

This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to configure and use the Fortinet Security Fabric.