FortiGate Security Study Guide

for FortiOS 7.0

FORTINET NSE Training Institute

Fortinet Training

https://training.fortinet.com

Fortinet Document Library

https://docs.fortinet.com

Fortinet Knowledge Base

https://kb.fortinet.com

Fortinet Fuse User Community

https://fusecommunity.fortinet.com/home

Fortinet Forums

https://forum.fortinet.com

Fortinet Support

https://support.fortinet.com

FortiGuard Labs

https://www.fortiguard.com

Fortinet Network Security Expert Program (NSE)

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

Fortinet | Pearson VUE

https://home.pearsonvue.com/fortinet

Feedback

Email: askcourseware@fortinet.com

1/26/2022

TABLE OF CONTENTS

Change Log

This table includes updates to the *FortiGate Security 7.0 Study Guide* dated 6/7/2021 to the updated document version dated 1/26/2022.

Change	Location
Various formatting fixes	Entire Guide
• Removed slides: Administration Methods, basic CLI commands, Two Factor Authentication, Link Aggregation • Added Slides : Related to VDOMs	Lesson 1
Updated CLI command `update-ffdb`	Lesson 3: Slide 15
Updated Mixing Policies slide example	Lesson 5: Slide 37
Fixed notes	Lesson 10: Slide 14
Fixed notes	Lesson 11: Slide 8
Trim down SSL VPN content and added ZTNA section	Lesson 12

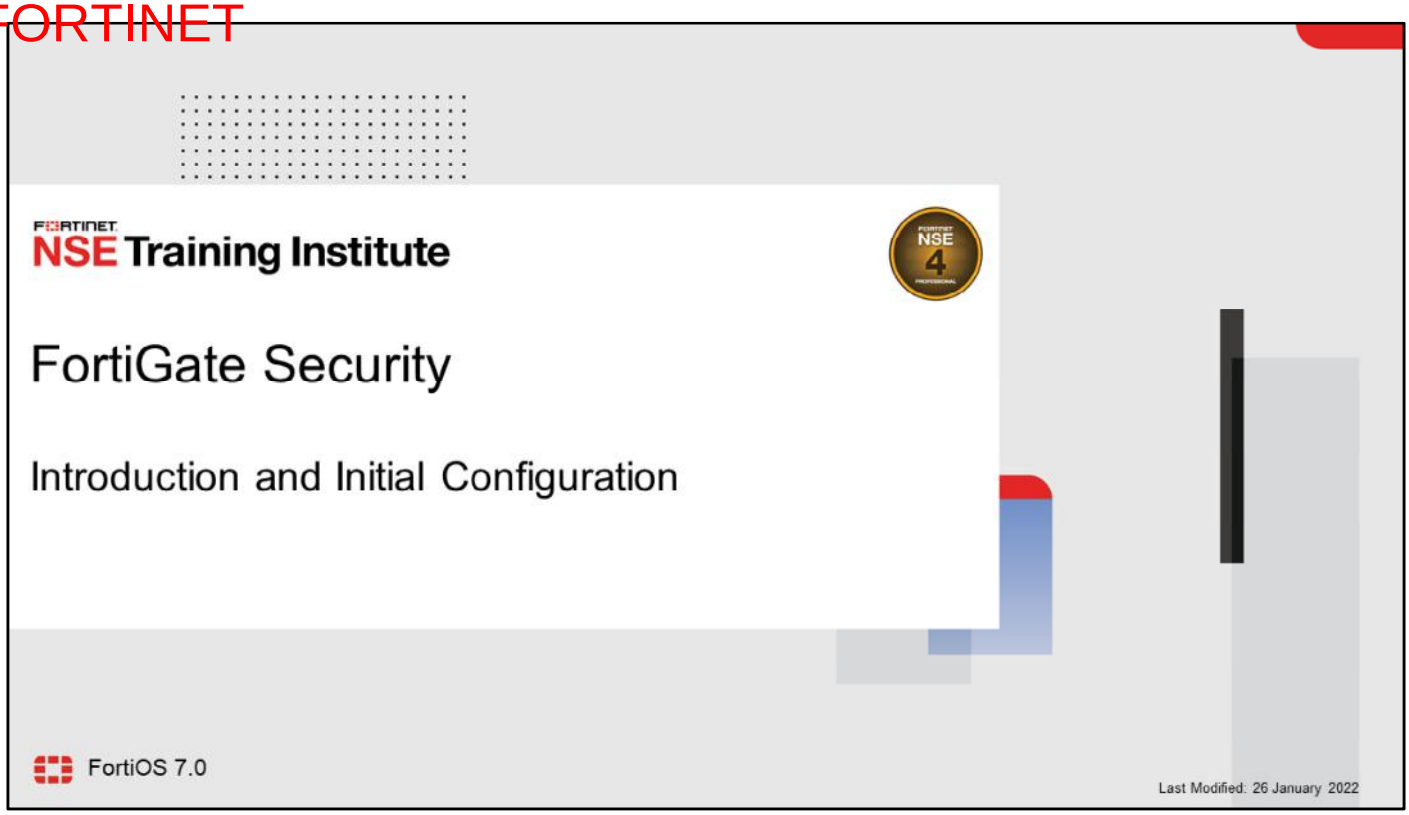

In this lesson, you will learn about FortiGate administration basics and the components within FortiGate that you can enable to extend functionality. This lesson also includes details about how and where FortiGate fits into your existing network architecture.

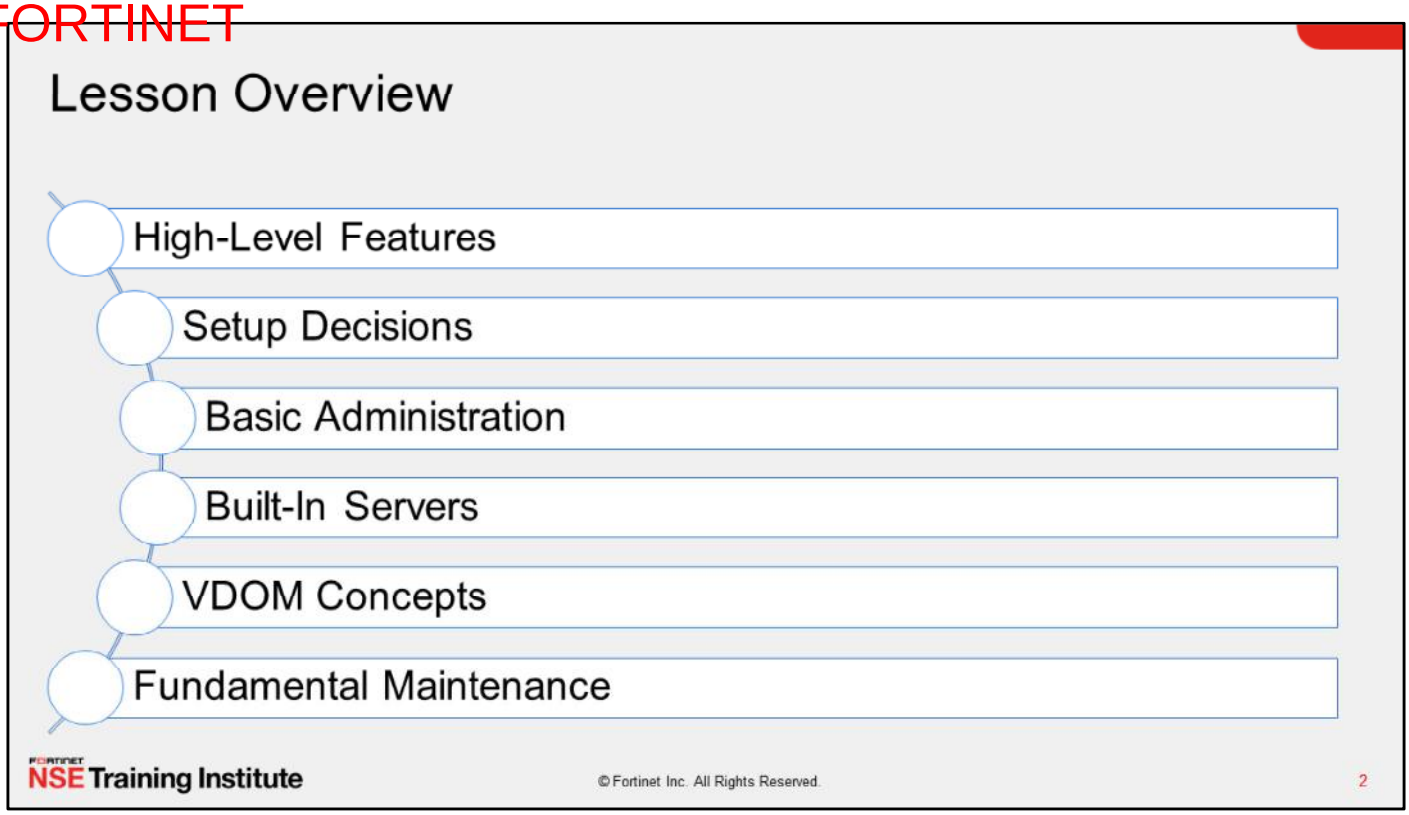

In this lesson, you will explore the topics shown on this slide.

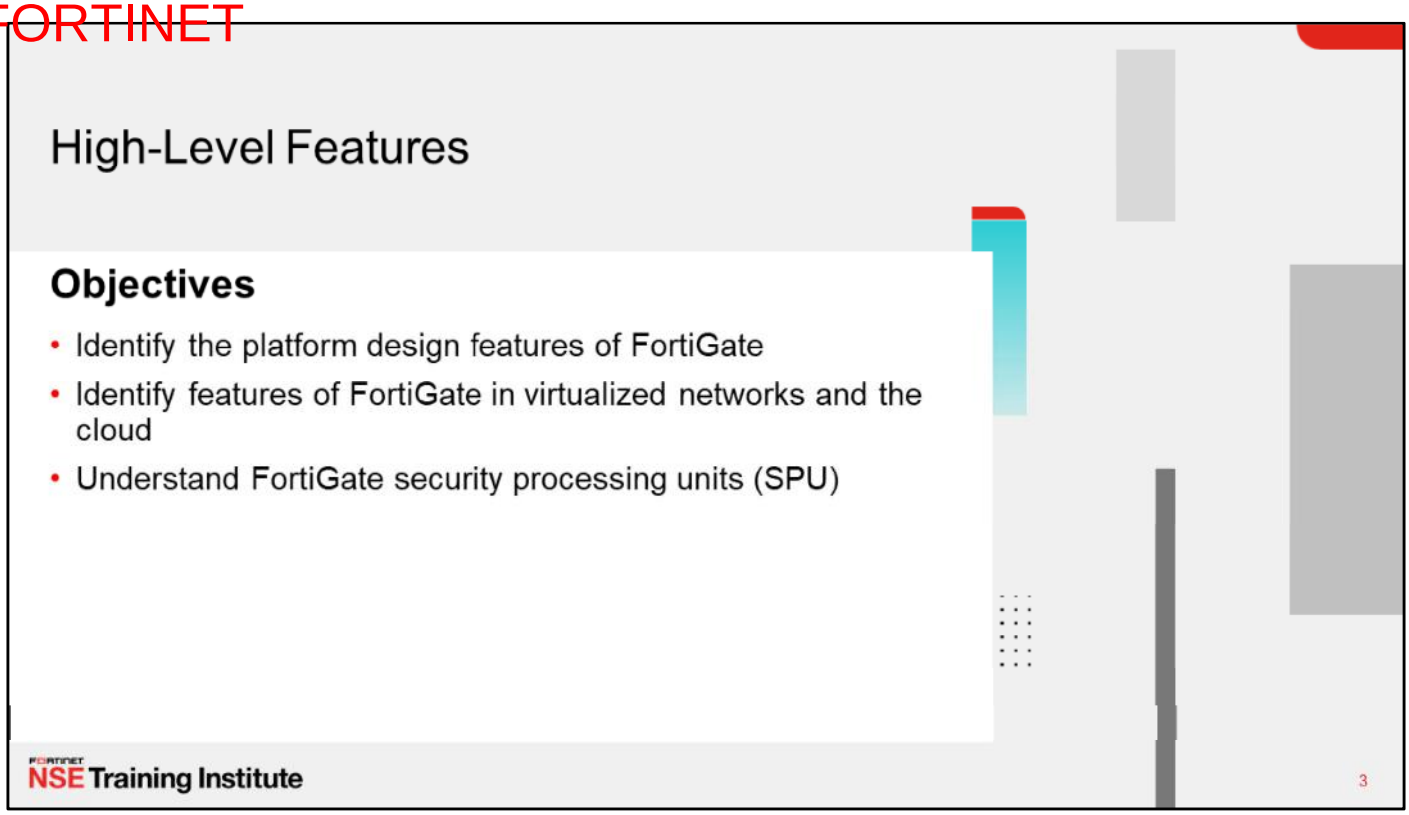

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

The Modern Context of Network Security

- Firewalls are more than gatekeepers on the network perimeter
- Today's firewalls are designed in response to multifaceted and multidevice environments with no identifiable perimeter:
  - Mobile workforce
  - Partners accessing your network services
  - Public and private clouds
  - Internet of things (IoT)
  - Bring your own device (BYOD)
- Firewalls are expected to perform different functions within a network
  - Different deployment modes:
    - Distributed enterprise firewall
    - Next-generation firewall
    - Internal segmentation firewall
    - Data center firewall
  - DNS, DHCP, web filter, intrusion prevention system (IPS), and so on

NSE Training Institute © Fortinet Inc. All Rights Reserved. 4

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.

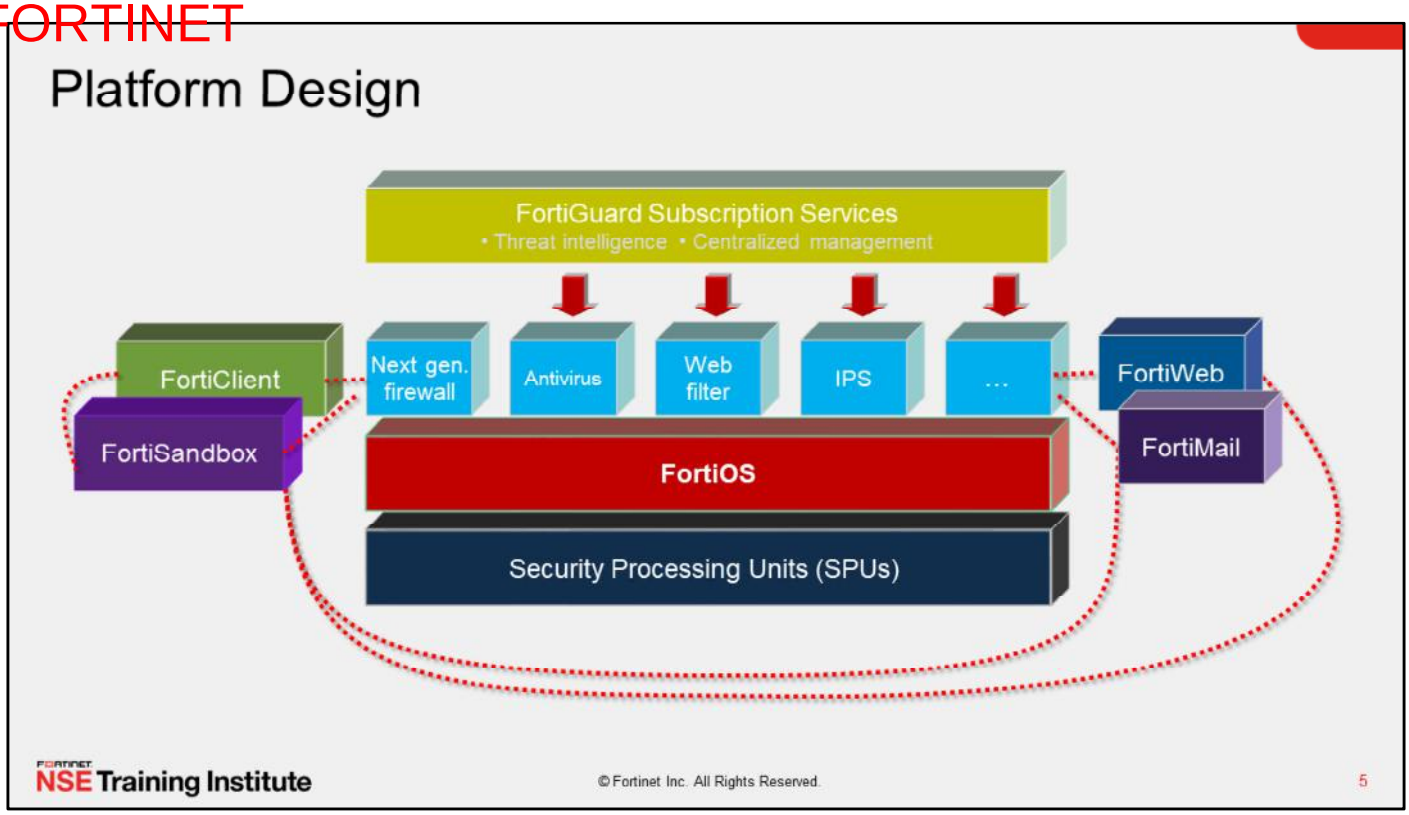

In the architecture diagram shown on this slide, you can see how FortiGate platforms add strength, without compromising flexibility. Like separate, dedicated security devices, FortiGate is still *internally* modular. Plus:

- **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
  - FortiGate VMX – Subset of features for VMware NSX-V (east-west) data flows
  - FortiGate Connector for Cisco ACI – Subset for Cisco ACI (north-south) data flows. Integrates physical *or* virtual appliance.
- **Faster setup and teardown: SDN + VMs**

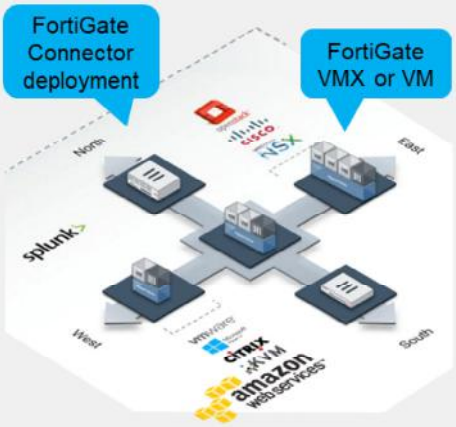

FortiGate VM Specifications	
Licenses	Max. 1 / 2 / 4 / 8 vCPU
Hypervisor	VMware, Hyper-V, KVM, Citrix Xen Server, Open Source Xen, Azure, Amazon AWS BYOL & on-demand
Memory	Max. 1/4/8/12 GB
10/100/1000 Interfaces	2-4 virtual NICs
Storage Capacity	40+ GB

NSE Training Institute

6

FortiGate virtual machines (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 Fortinet's 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 would anyone 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.

FortiGate VMX and the FortiGate Connector for Cisco ACI are specialized versions of FortiOS and an API that allow you to orchestrate rapid network changes through standards, such as OpenStack for software-defined networking (SDN).

- FortiGate VM is deployed as a guest VM on the hypervisor.
- FortiGate VMX is deployed inside the virtual networks of a hypervisor, *between* guest VMs.
- FortiGate Connector for Cisco ACI allows ACI to deploy physical *or* virtual FortiGate VMs for north-south traffic.

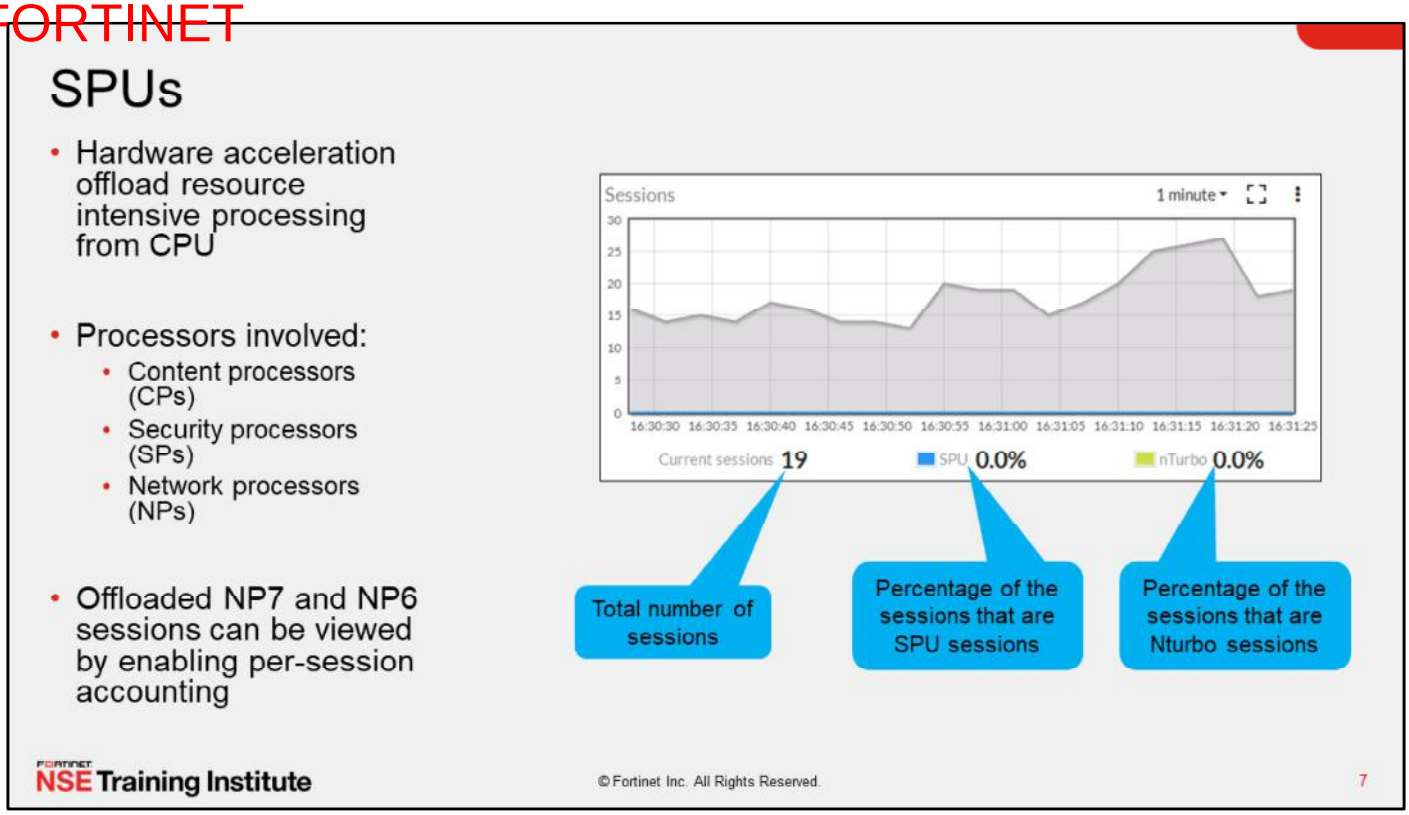

All Fortinet hardware acceleration hardware has been renamed security processing units (SPUs). This includes NPx and CPx processors.

Most FortiGate models have specialized acceleration hardware, called SPUs that can offload resource-intensive processing from main processing (CPU) resources. Most FortiGate devices include specialized content processors (CPs) that accelerate a wide range of important security processes, such as virus scanning, attack detection, encryption, and decryption. (Only selected entry-level FortiGate models do not include a CP processor.)

SPU and nTurbo data is now visible in a number of places on the GUI. For example, the **Active Sessions** column pop-up in the firewall policy list and the **Sessions** dashboard widget. Per-session accounting is a logging feature that allows FortiGate to report the correct bytes per packet numbers per session for sessions offloaded to an NP7, NP6 or NP6lite processor.

The following example shows the **Sessions** dashboard widget tracking SPU and nTurbo sessions. **Current sessions** shows the total number of sessions, **SPU** shows the percentage of these sessions that are SPU sessions, and **Nturbo** shows the percentage that are nTurbo sessions.

NTurbo offloads firewall sessions that include flow-based security profiles to NP6 or NP7 network processors. Without NTurbo, or with NTurbo disabled, all firewall sessions that include flow-based security profiles are processed by the FortiGate CPU.

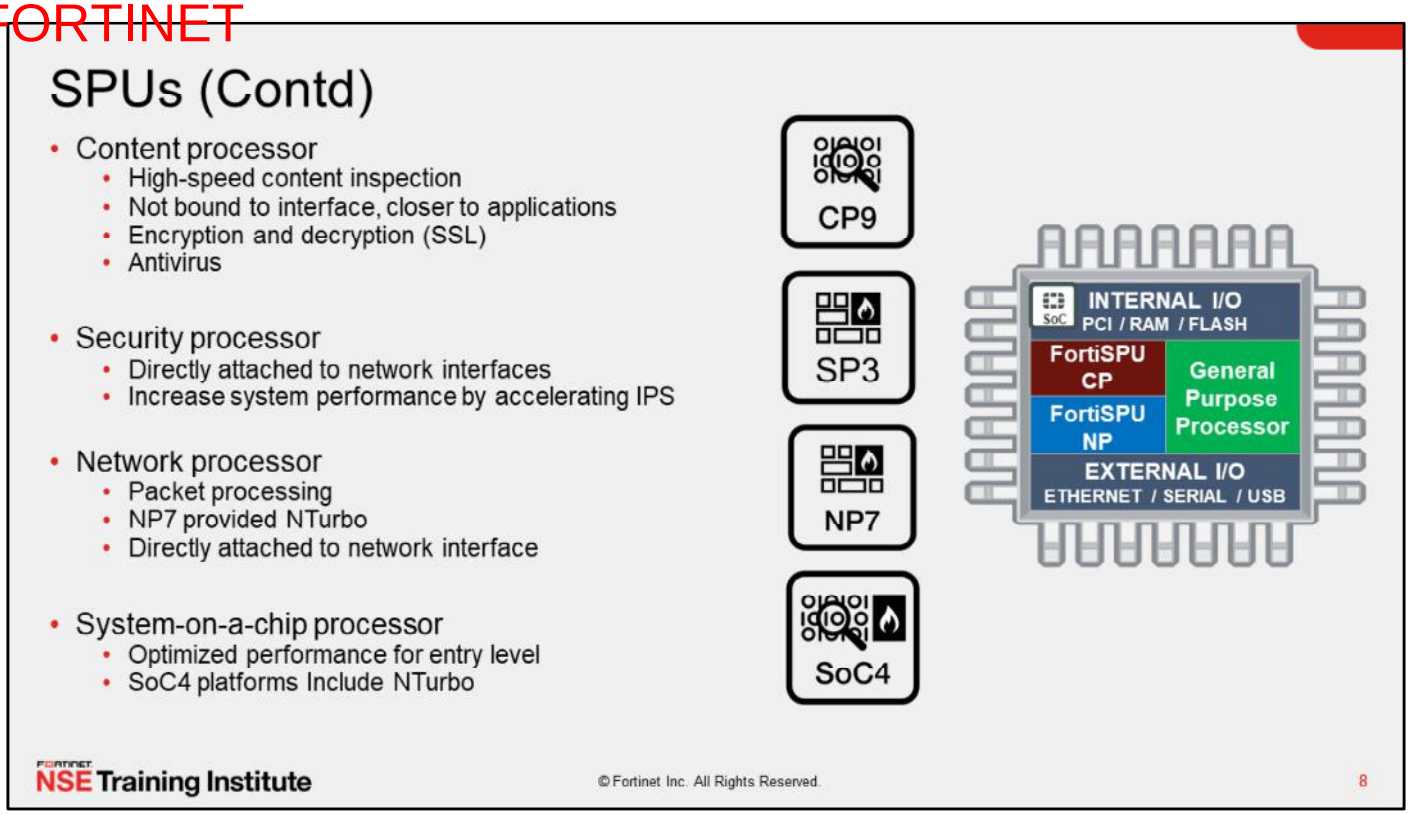

The Fortinet content processor (CP9) works outside of the direct flow of traffic, providing high-speed cryptography and content inspection services. This frees businesses to deploy advanced security whenever it is needed without impacting network functionality. CP8 and CP9 provide a fast path for traffic inspected by IPS, including sessions with flow-based inspection.

CP processors also accelerate intensive proxy-based tasks:
- Encryption and decryption (SSL)
- Antivirus

FortiSPU network processors work at the interface level to accelerate traffic by offloading traffic from the main CPU. Models that support FortiOS 6.4 or later contain NP6, NP6lite, and NP7 network processors.

Fortinet integrates content and network processors along with a RISC-based CPU into a single processor known as SoC4 for entry-level FortiGate security devices used for distributed enterprises. This simplifies device design and enables breakthrough performance without compromising on security.

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

NSE Training Institute

9

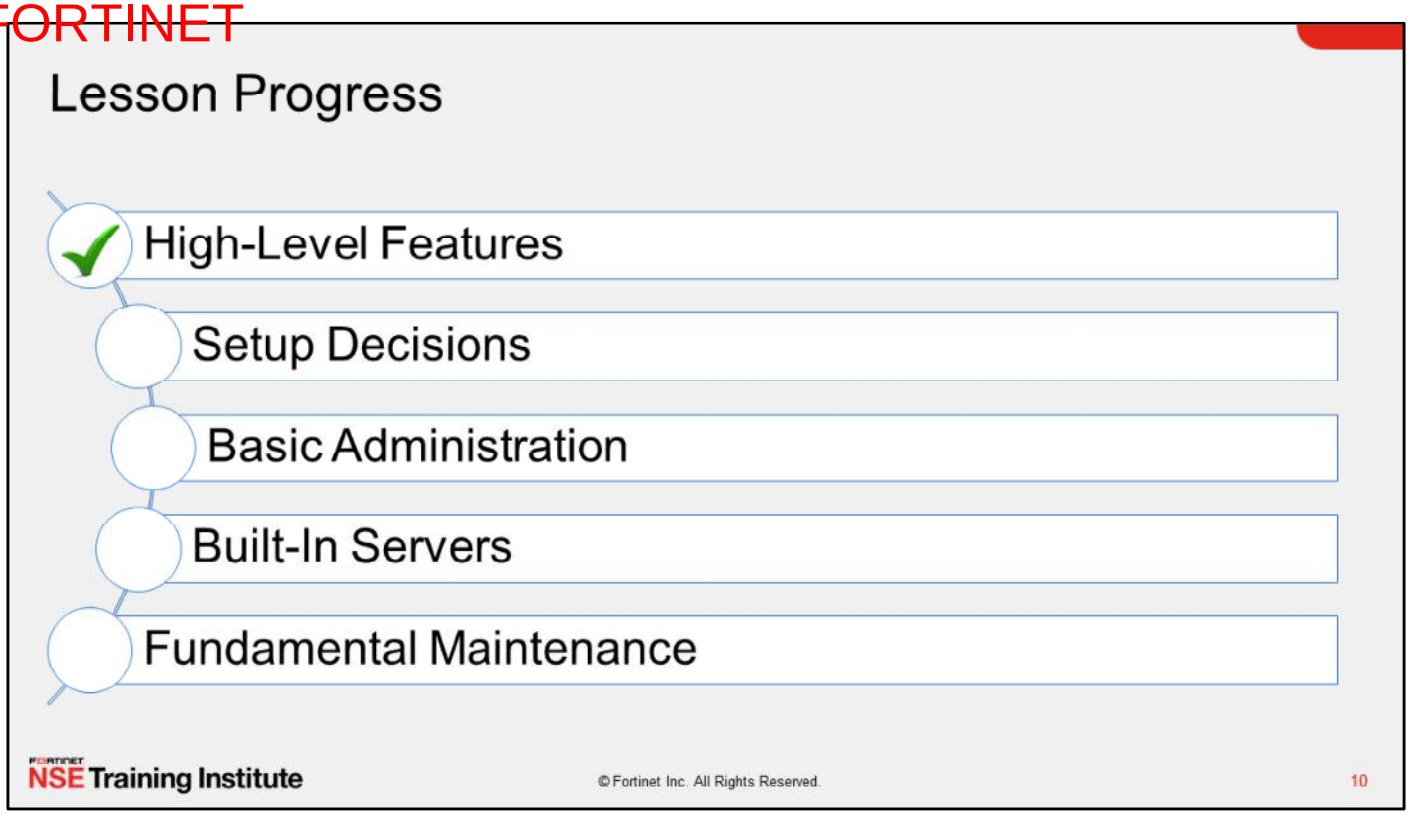

Good job! You now understand some of the high-level features of FortiGate.

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.

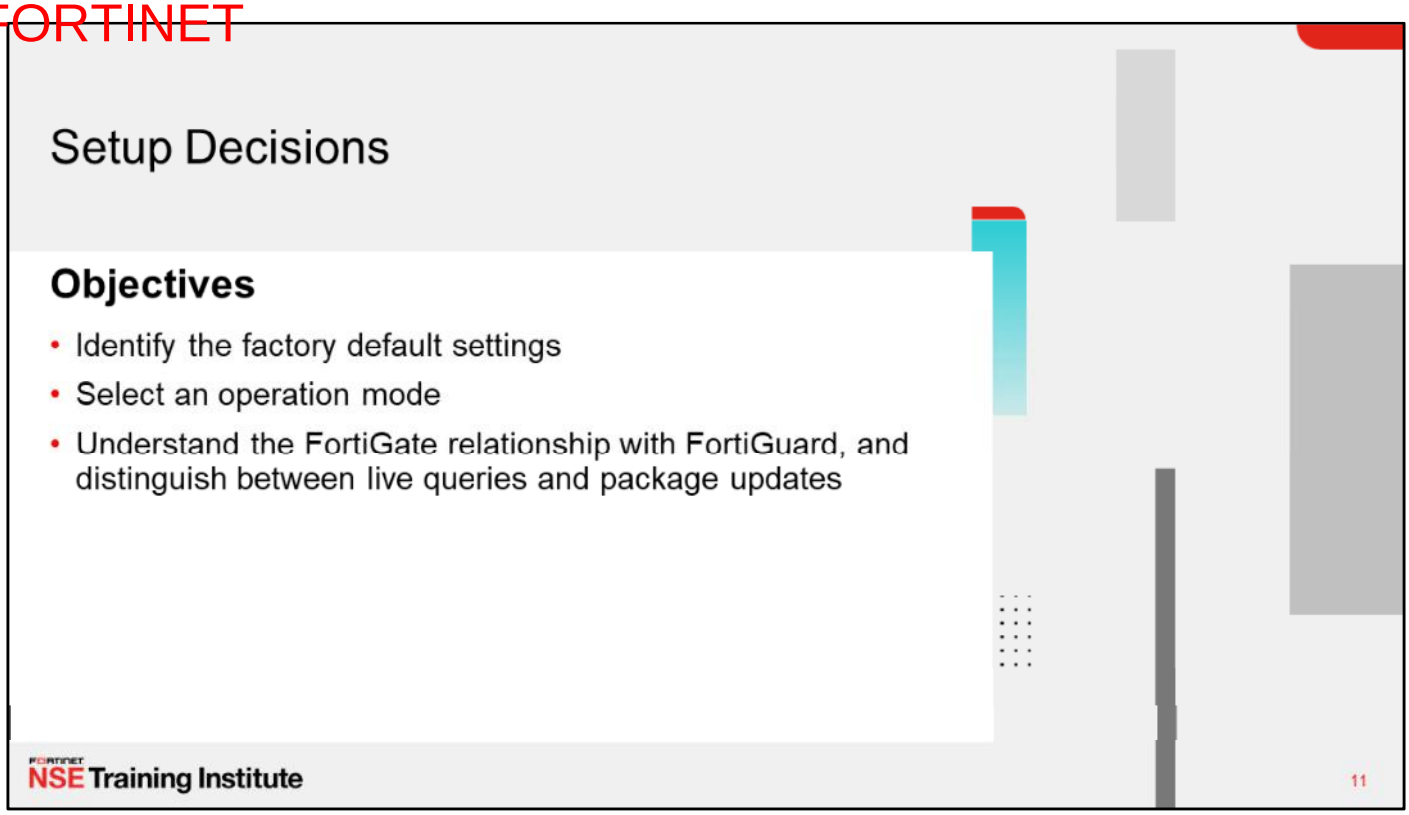

After completing this section, you should be able to achieve the objectives shown on this slide.

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

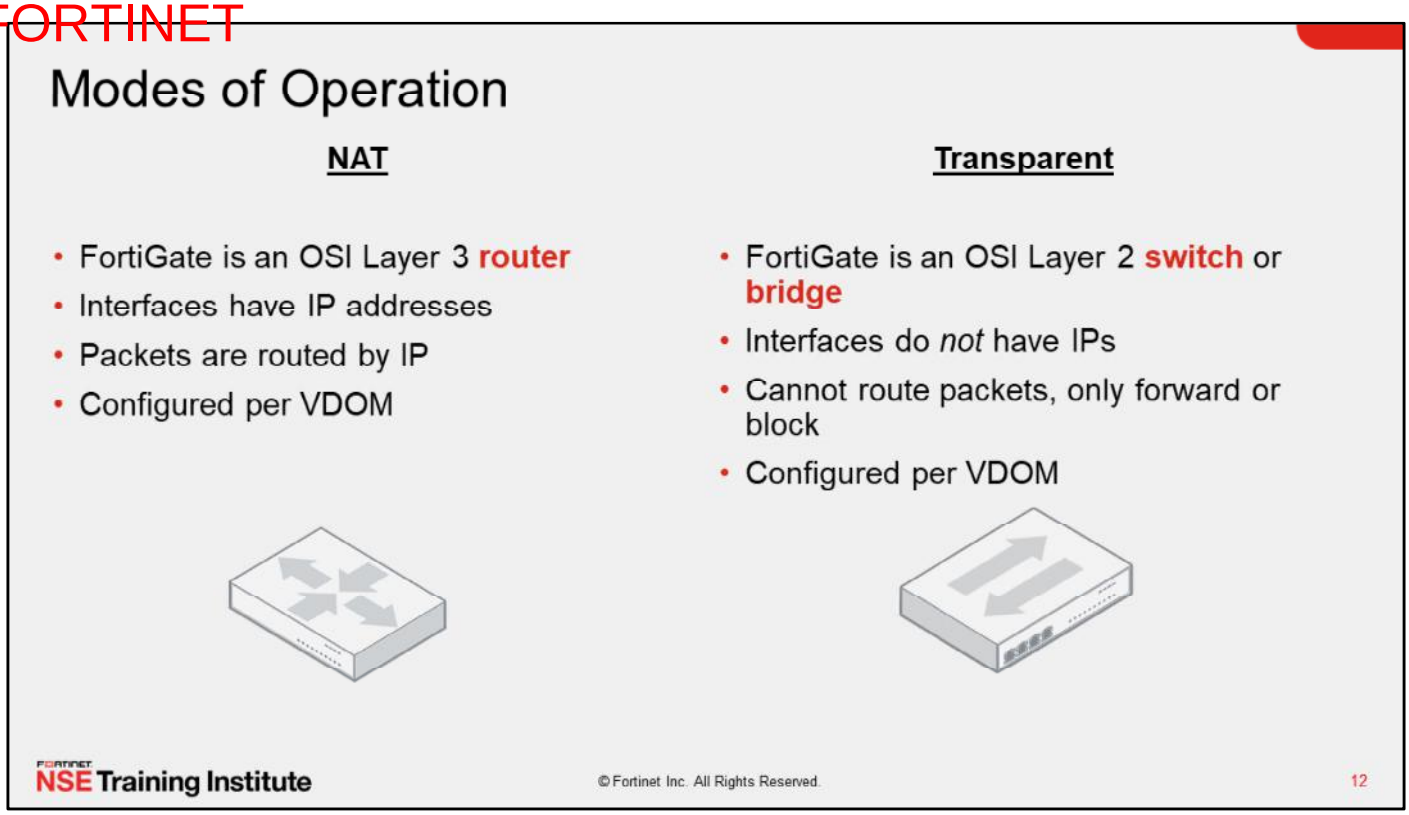

What about the network architecture? Where does FortiGate fit in?

When you deploy FortiGate, you can choose between two operating modes: NAT mode or transparent mode.

- In NAT mode, FortiGate routes packets based on Layer 3, like a router. Each of its logical network interfaces has an IP address and FortiGate determines the outgoing or egress interface based on the destination IP address and entries in its routing tables.
- In transparent mode, FortiGate forwards packets at Layer 2, like a switch. Its interfaces have no IP addresses and FortiGate identifies the outgoing or egress interface based on the destination MAC address. The device in transparent mode has an IP address used for management traffic.

Interfaces *can* be exceptions to the router versus switch operation mode, on an individual basis.

When you enable virtual domains (VDOMs) on FortiGate, you can configure each VDOM for NAT mode or transparent mode, regardless of the operation mode of other VDOMs on FortiGate. By default, VDOMs are disabled on the FortiGate device, but there is still one VDOM active: the *root* VDOM. It is always there in the background. When VDOMs are disabled, the NAT mode or transparent mode relates to the root VDOM.

VDOMs are a method of dividing a FortiGate device into two or more virtual devices that function as multiple independent devices. VDOMs can provide separate firewall policies and, in NAT mode, completely separate configurations for routing and VPN services for each connected network or organization. In transparent mode, VDOM applies security scanning to traffic and is installed between the internal network and the external network.

By default, a VDOM is in NAT mode when it is created. You can switch it to transparent mode, if required.

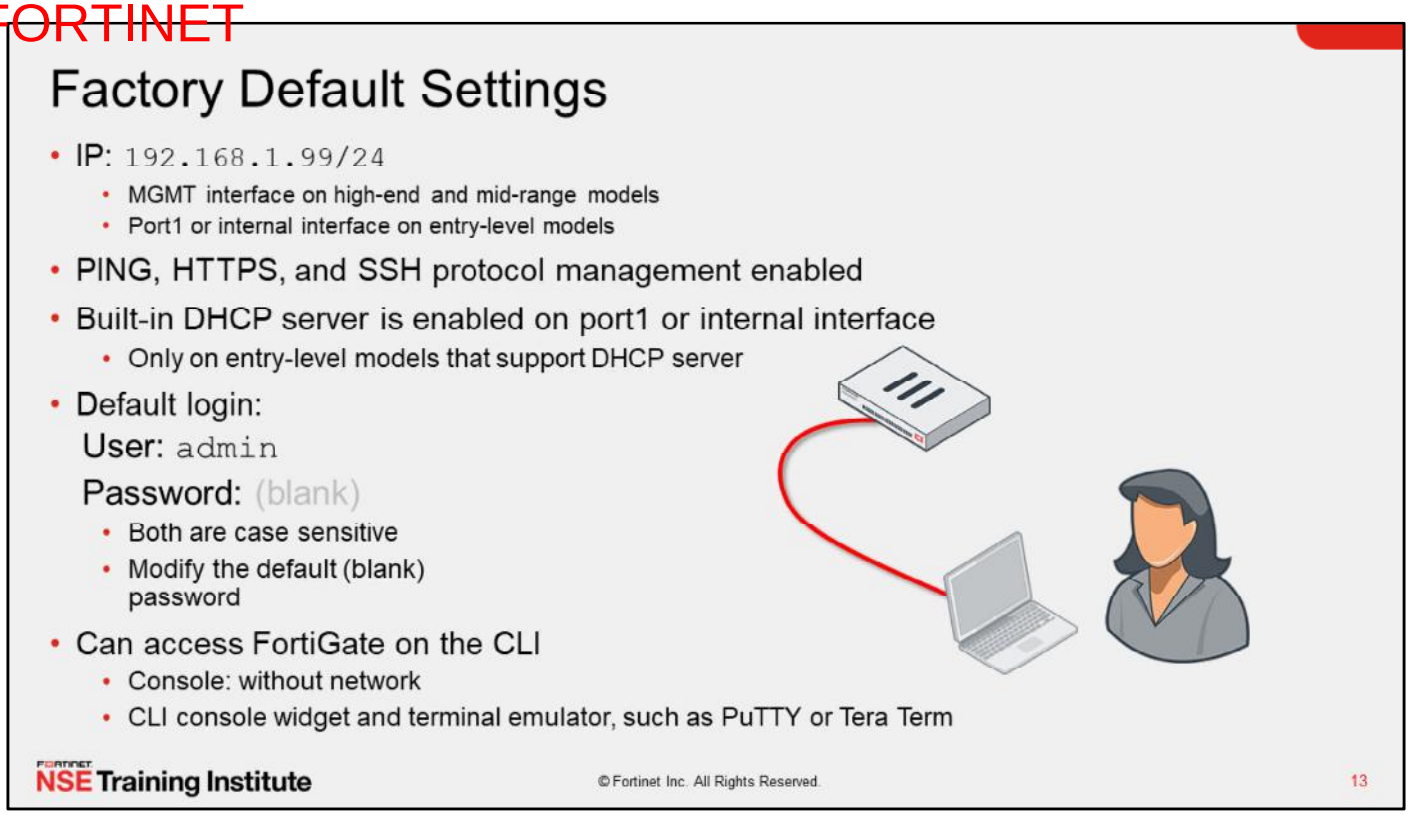

Network address translation (NAT) mode is the default operation mode. What are the other factory default settings? After you've removed FortiGate from its box, what do you do next?

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.

FortiGuard Subscription Services

- Internet connection and contract required
- Provided by FortiGuard Distribution Network (FDN)
  - Major data centers in North America, Asia, and Europe
    - Or, from FDN through your FortiManager
  - FortiGate prefers the data center in nearest time zone, but will adjust by server load
- Package updates: FortiGuard antivirus and IPS
  - `update.fortiguard.net`
  - TCP port 443 (SSL)
- Live queries: FortiGuard web filtering, DNS filtering, and antispam
  - `service.fortiguard.net` for proprietary protocol on UDP port 53 or 8888
  - `securewf.fortiguard.net` for HTTPS over port 443, 53 or, 8888

NSE Training Institute © Fortinet Inc. All Rights Reserved. 14

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.

FortiGuard Subscription Services (Contd)

- FortiGuard third party SSL certificate verification and OCSP stapling check
  - Default FortiGuard access mode is anycast
  - Optimize the routing performance to the FortiGuard servers
  - FortiGate gets a single IP address for the domain name of each FortiGuard service
  - FortiGuard servers query the CA OCSP responder every four hours
  - Enforce a connection to use protocol HTTPS and port 443

Server	Domain name and IP address
Object download	globalupdate.fortinet.net - 173.243.140.6
Querying service (webfiltering, antispam)	globalguardservice.fortinet.net - 173.243.140.16
FortiGate Cloud logging	globallogctrl.fortinet.net - 173.243.132.25
FortiGate Cloud management	globalmgrctrl.fortinet.net - 173.243.132.26
FortiGate Cloud messaging	globalmsgctrl.fortinet.net - 173.243.132.27
FortiGate Cloud sandbox	globalaptctrl.fortinet.net - 184.94.112.22
The productapi used by OCVPN registration and GUI icon download	globalproductapi.fortinet.net - 66.35.17.252

NSE Training Institute © Fortinet Inc. All Rights Reserved. 15

In FortiOS 6.4 or later, 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

NSE Training Institute

16

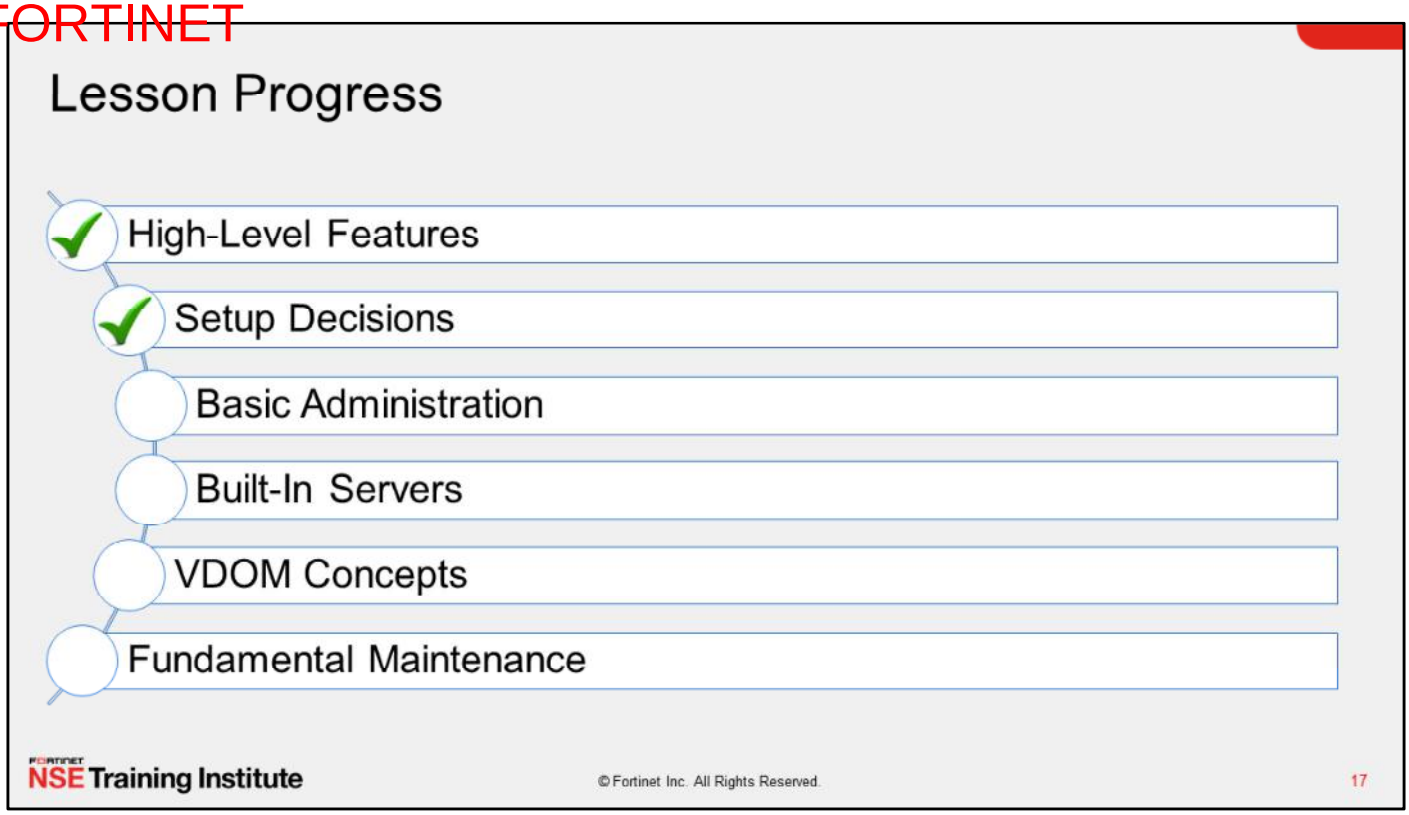

Good job! You now understand how to perform the initial setup of FortiGate and why you might decide to use one configuration over another. Now, you will learn about basic administration.

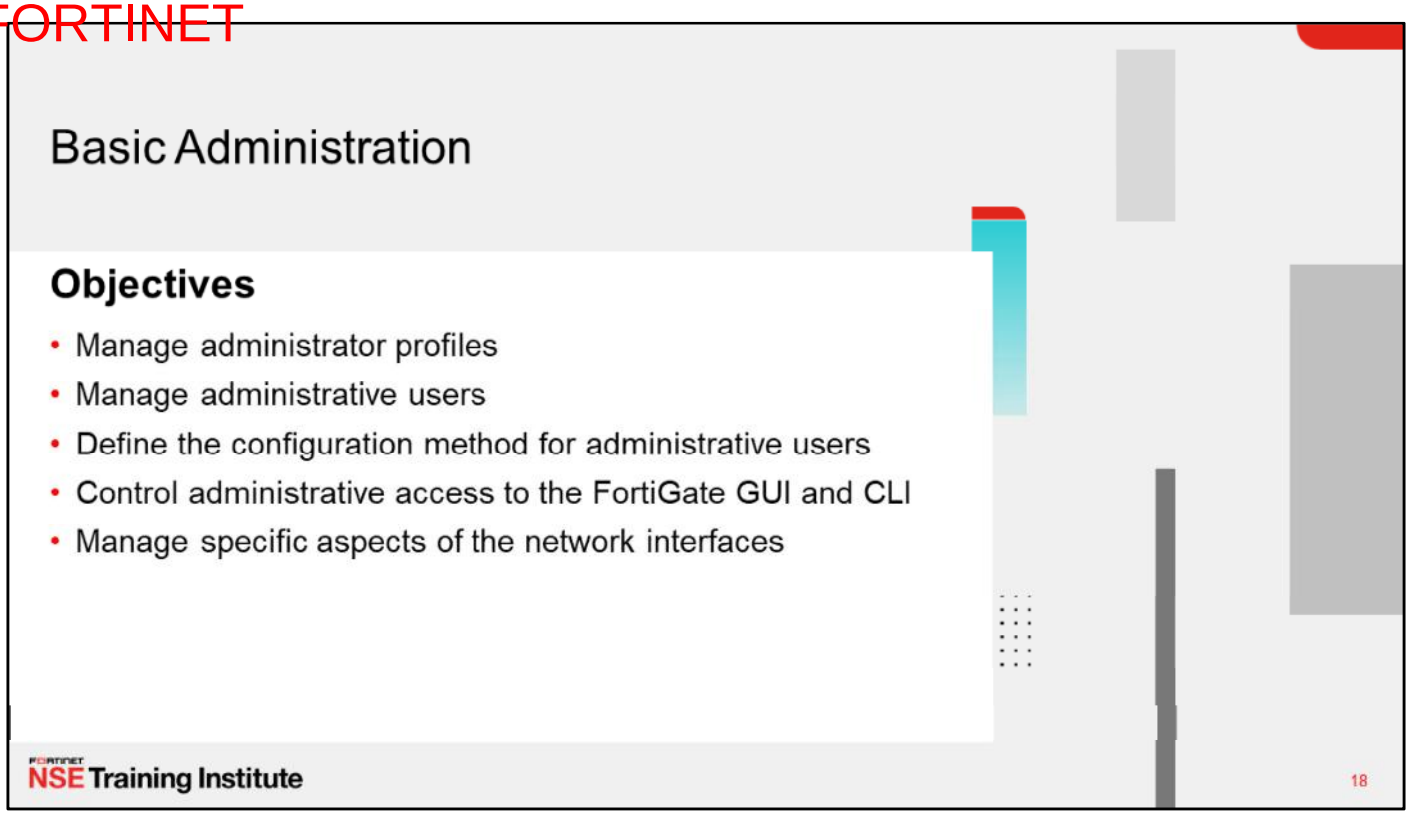

After completing this lesson, you should be able to achieve the objectives shown on this slide.

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

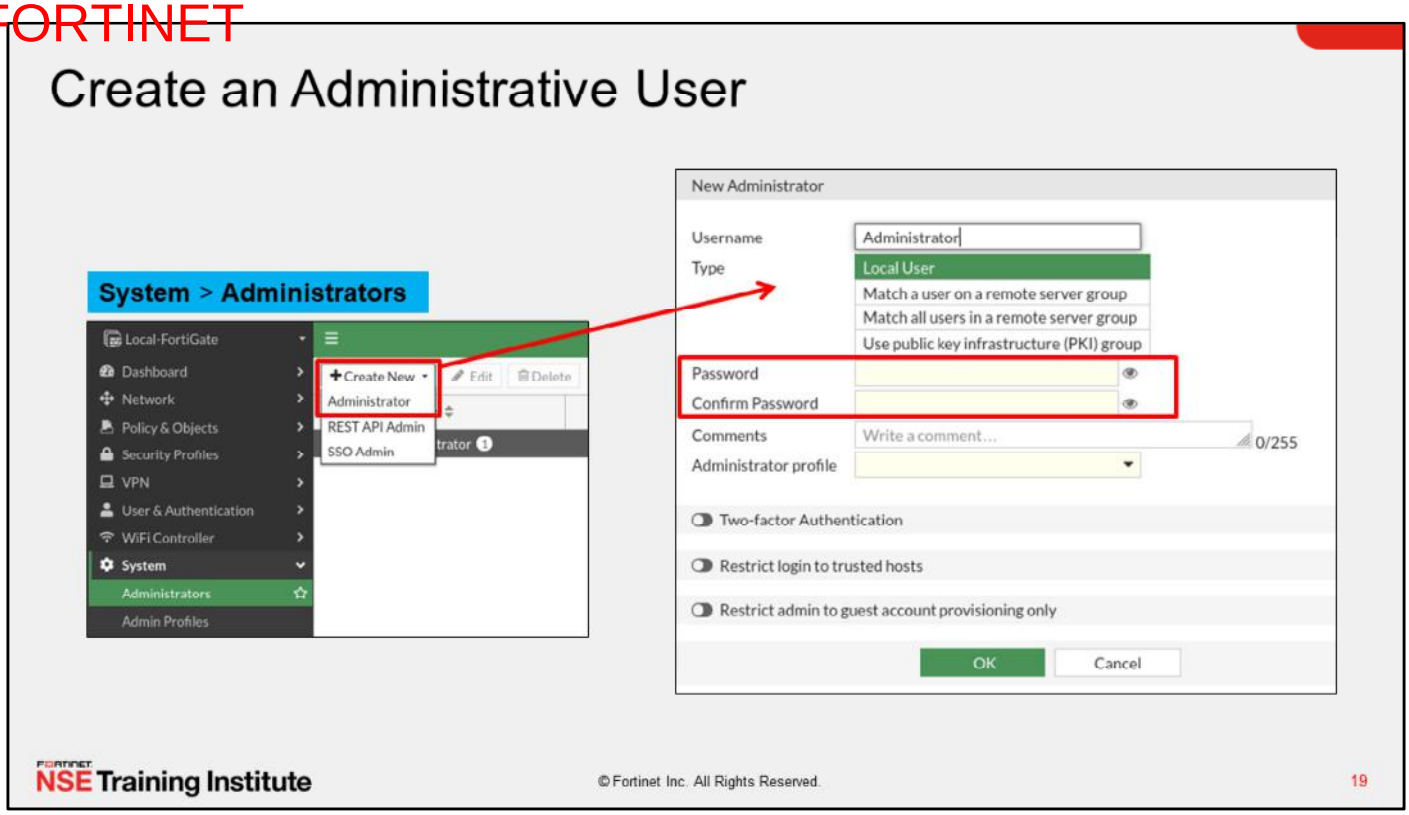

Whichever method you use, start by logging in as admin. Begin by creating separate accounts for other administrators. For security and tracking purposes, it is a best practice for each administrator to have their own account.

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.

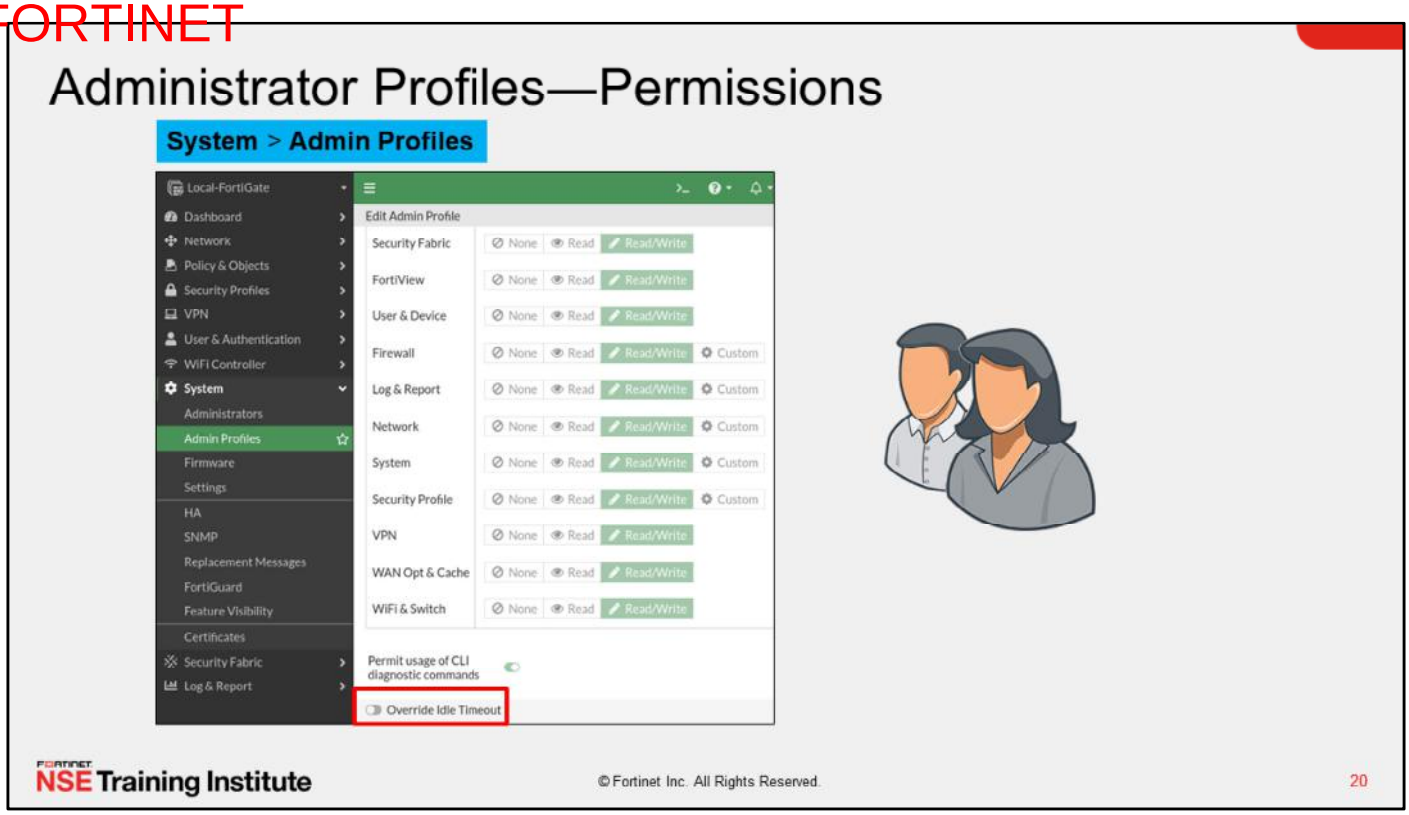

When assigning permissions to an administrator profile, you can specify read-and-write, read-only, or none to each area.

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.

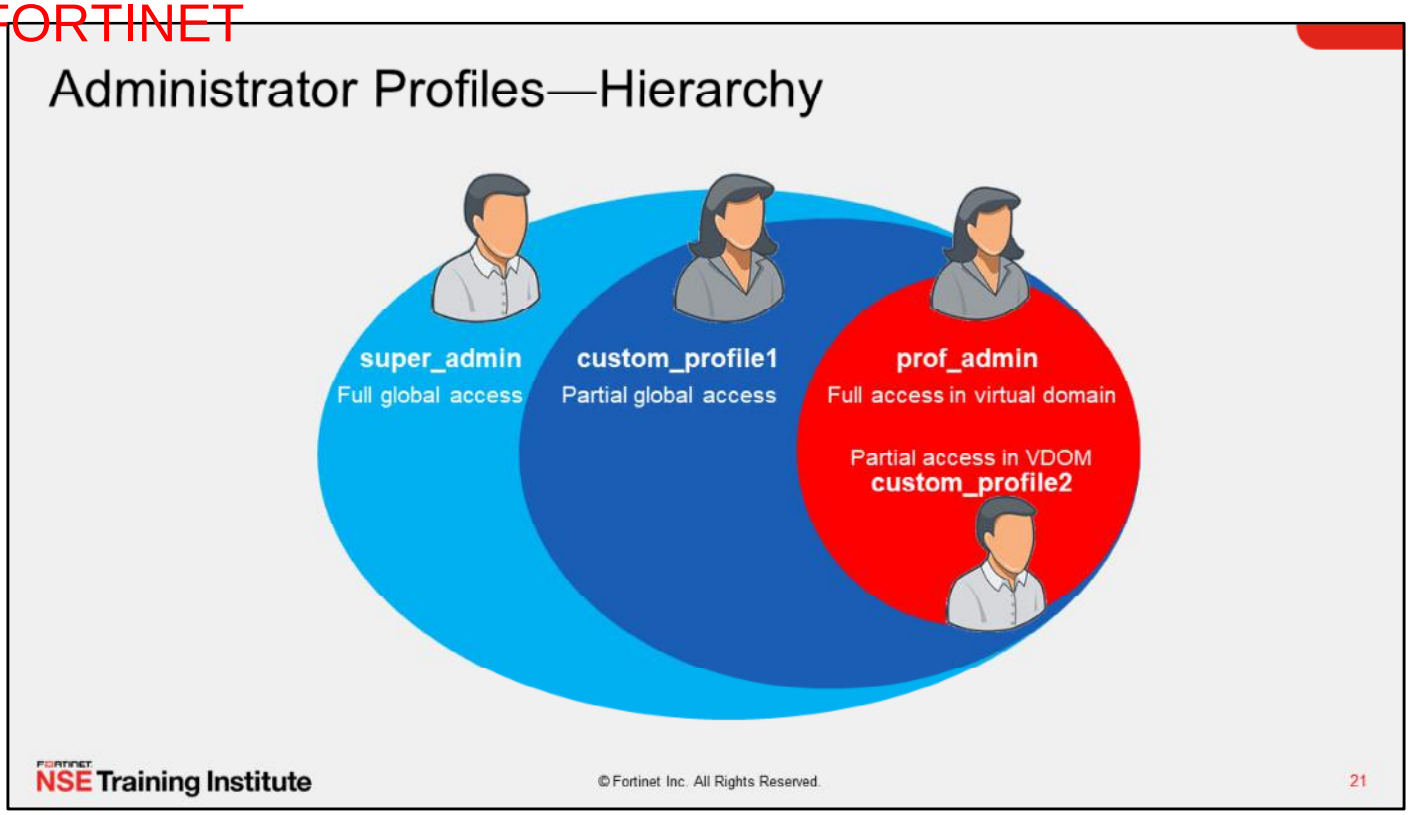

What are the effects of administrator profiles?

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. So, for example, an administrator using the **prof_admin** or a custom profile cannot see, or reset the password of accounts that use the **super_admin** profile.

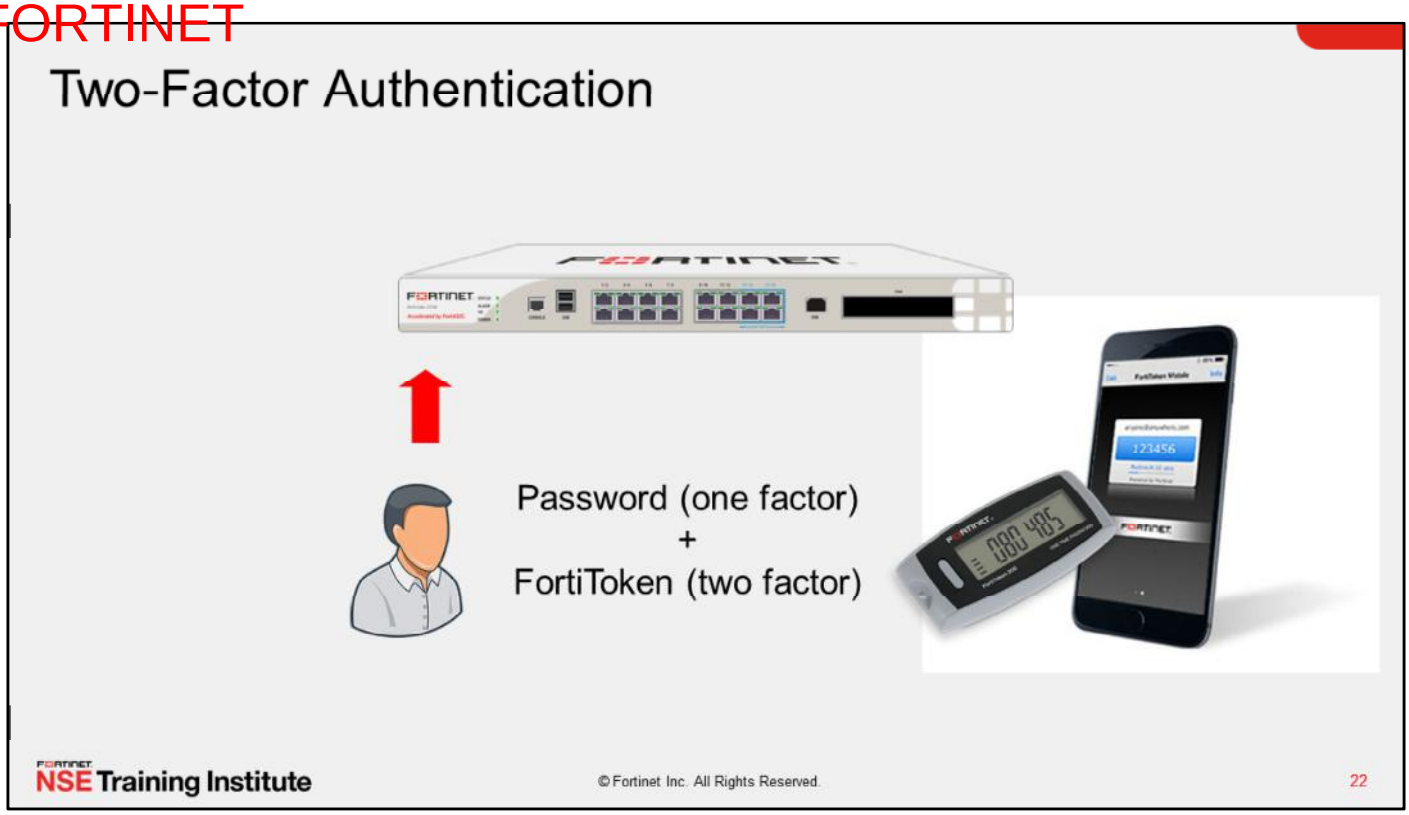

To further secure access to your network security, use two-factor authentication.

Two-factor authentication means that instead of using one method to verify your identity—typically a password or digital certificate—your identity is verified by two methods. In the example shown on this slide, two-factor authentication includes a password plus an RSA randomly generated number from a FortiToken that is synchronized with FortiGate.

Resetting a Lost Admin Password

User: `maintainer`

Password: `bcpb`<serial-number>

All letters in <serial-number> *must* be upper case, for example, `FGT60`

- All FortiGate appliance models and some other Fortinet device types
- No maintainer procedure in VM, revert to snapshot or reprovision VM
- Only after hard power cycle
  - Soft cycle (reboot) does not work for security reasons
- Only during first 60 seconds *after* boot (varies by model)
  - **Tip**: Copy serial number into the terminal buffer, then paste
- Only through hardware console port
  - Requires physical access for security reasons
  - If compliance/risk of physical access requires, you can disable maintainer

```
config sys global
  set admin-maintainer disable
end
```

NSE Training Institute © Fortinet Inc. All Rights Reserved. 23

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.

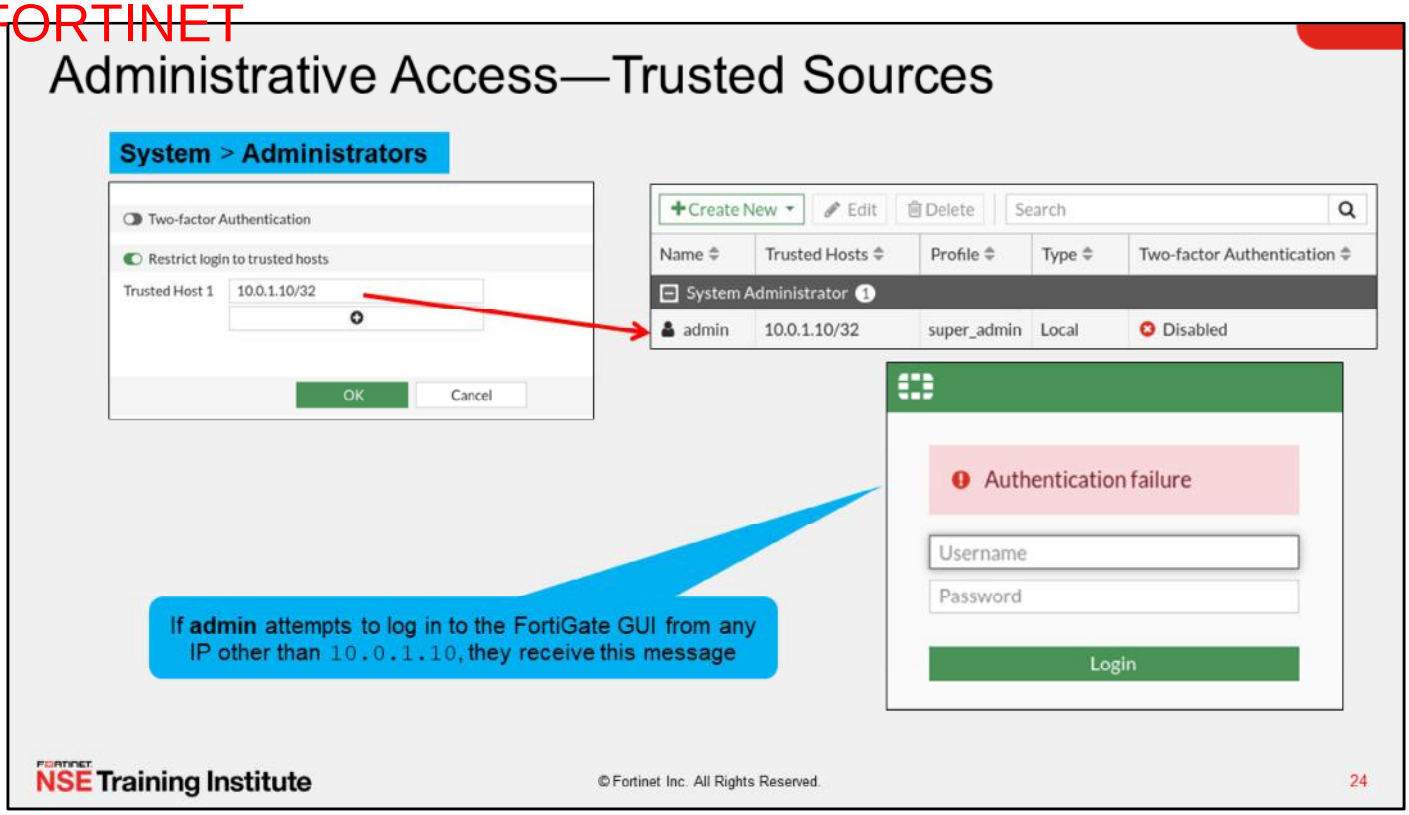

Another way to secure FortiGate is to define the hosts or subnets that are trusted sources from which to log in.

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.

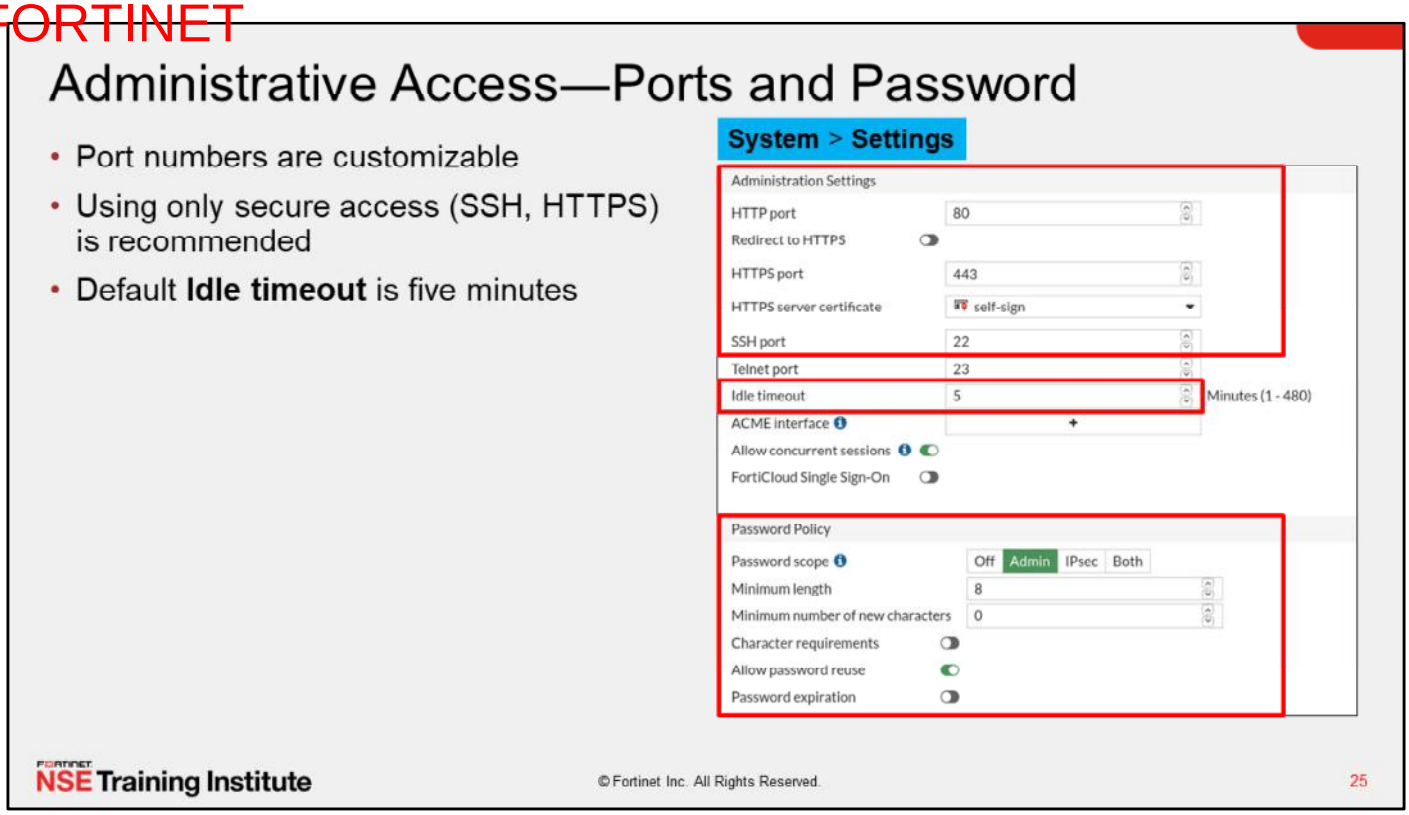

You may also want to customize the administrative protocols port numbers.

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.

Administrative Access—Protocols

- Enable acceptable management protocols on each interface independently:
  - Separate IPv4 and IPv6
  - IPv6 options hidden by default
- Also protocols where FortiGate is the destination IP:
  - Security Fabric Connection:
    - CAPWAP
    - FortiTelemetry
  - FMG-Access
  - FTM
  - RADIUS Accounting
- LLDP Support
  - Detecting an upstream Security Fabric FortiGate through LLDP

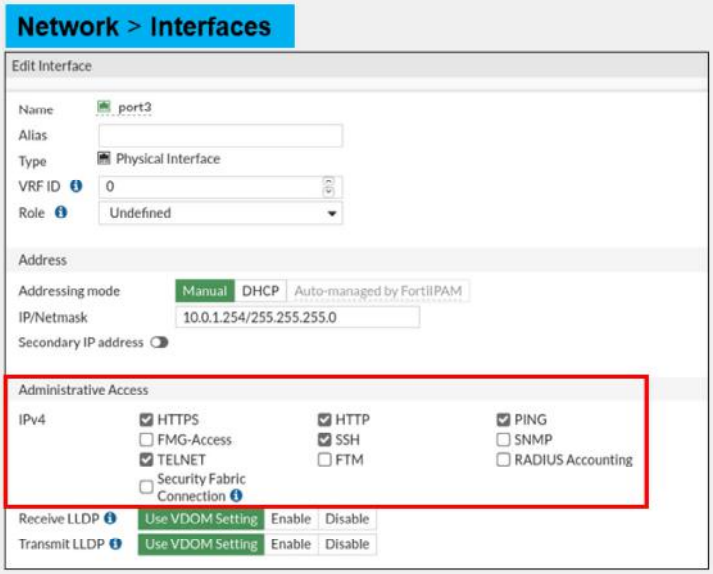

NSE Training Institute © Fortinet Inc. All Rights Reserved. 26

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.

Features Hidden by Default

- By default, some features like IPv6 are hidden on the GUI
  - Hidden features are not disabled
- In **Feature Visibility**, select whether to hide or show groups of features commonly used together

System > Feature Visibility

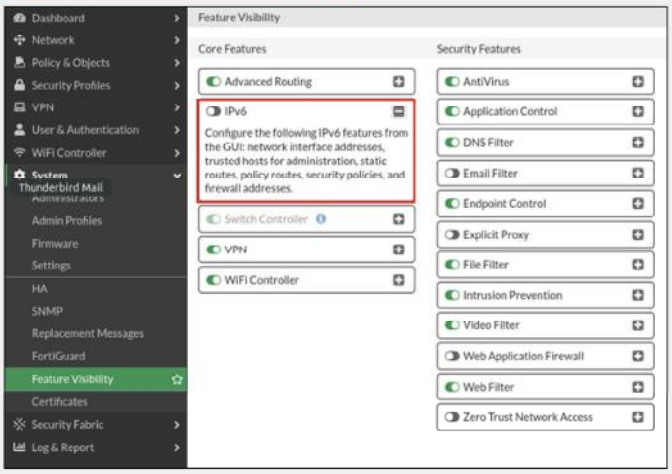

NSE Training Institute

27

FortiGate has hundreds of features. If you don't use all of them, hiding features that you don't use makes it easier to focus on your work.

Hiding a feature on the GUI does not disable it. It is still functional, and still can be configured using the CLI.

Some advanced or less commonly used features, such as IPv6, are hidden by default.

To show hidden features, click **System** > **Feature Visibility**.

Interface IPs

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

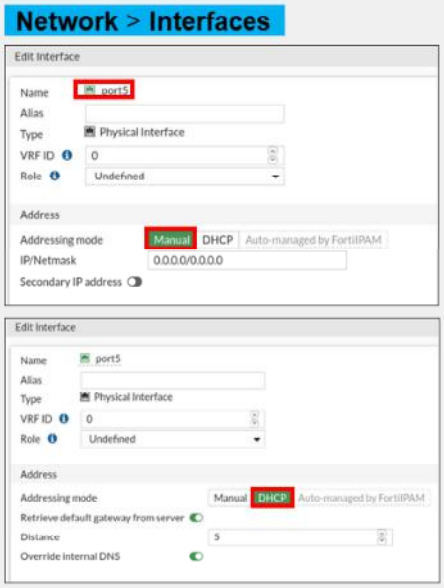

28

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)

Interface IPs (Contd)

- Select **Auto-managed by FortiIPAM** to better manage DHCP using FortiGuard
  - FortiIPAM is paid service
  - Provides IP address management solution
  - IP subnets do not overlap
- Exception: **One-Arm Sniffer**
  - Only available in CLI

```
config system interface
  edit port <number>
   set ips-sniffer-mode enable
  end
```

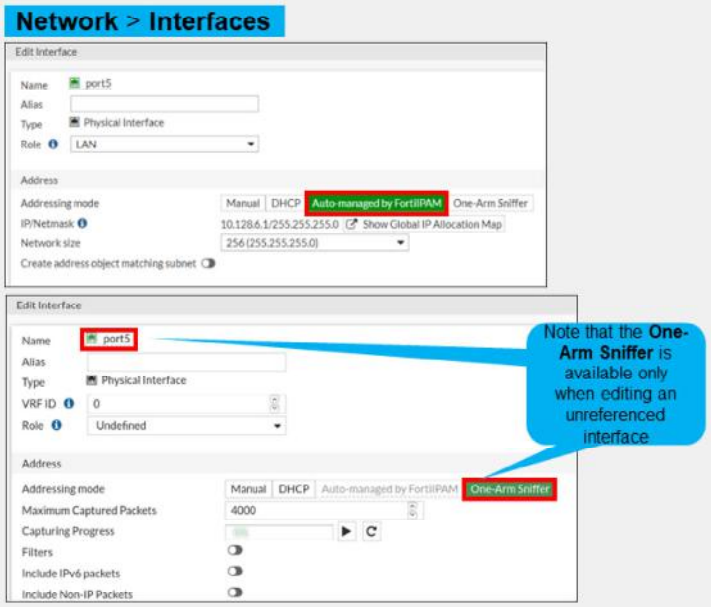

NSE Training Institute

29

FortiGate can use FortiIPAM to automatically assign IP addresses based on the configured network size for the FortiGate interface. FortiIPAM provides an on-premises IP address management solution when integrating network resources with FortiGate, and automatically assigns subnets to FortiGate to prevent duplicate IP addresses from overlapping within the same Security Fabric. Note that FortiIPAM is a paid service.

There is an exception to the IP address requirement: the **One-Arm Sniffer** interface type. This interfaces is *not* assigned an address.

When you select **One-Arm Sniffer** by enabling a sniffer on the CLI, the interface is not inline with the traffic flow. Rather, it is receiving a copy of the traffic from a mirrored port on a switch. The interface operates in promiscuous mode, scanning traffic that it sees, but is unable to make changes because the original packet has already been processed by the switch. As a result, one-arm sniffer mode is mostly used in proof of concept (POC), or in environments where corporate requirements state that traffic must not be changed, only logged. Once it is enabled, a **One-Arm Sniffer** option appears in the **Addressing mode** setting of a interface.

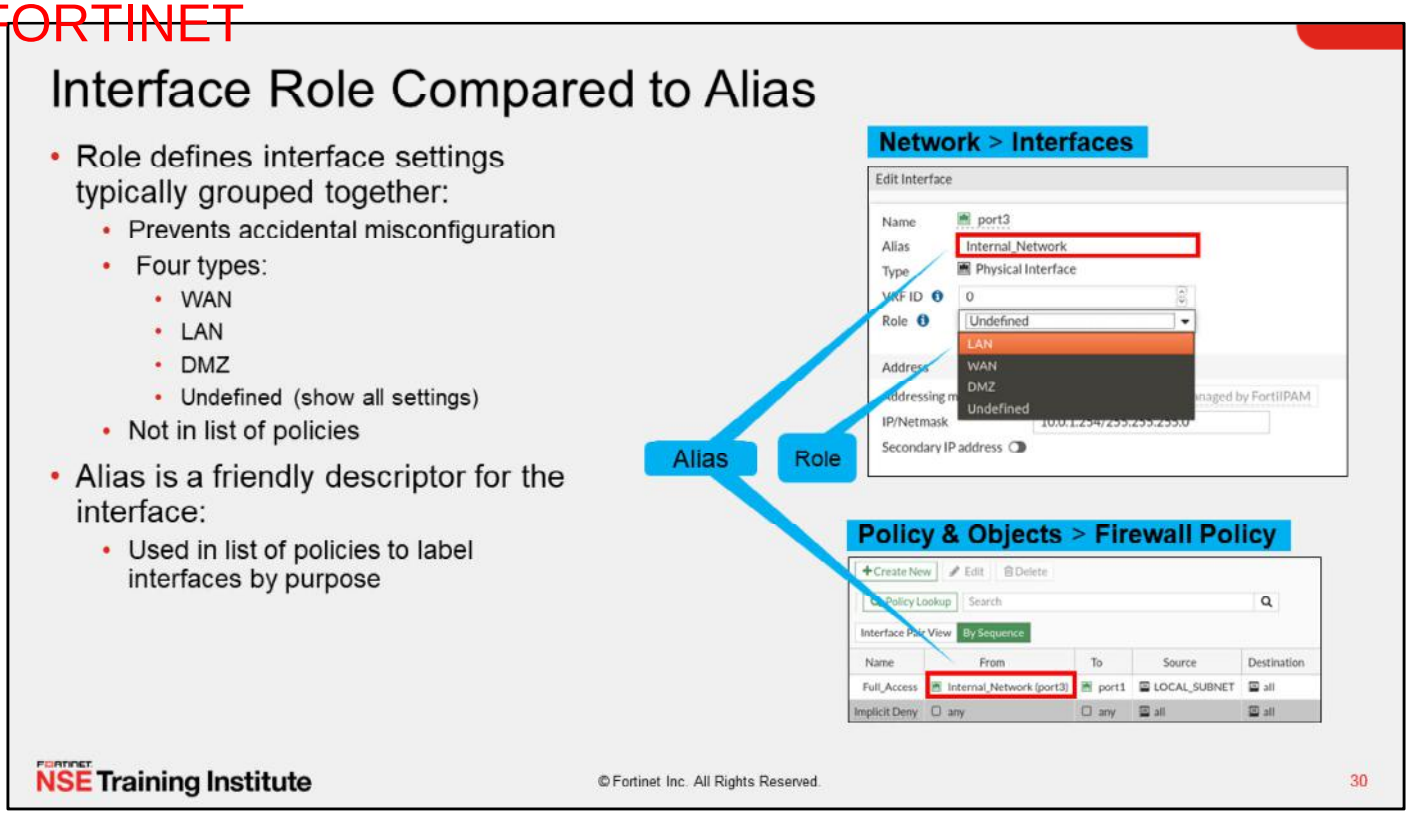

How many times have you seen network issues caused by a DHCP server—not client—enabled on the WAN interface?

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.

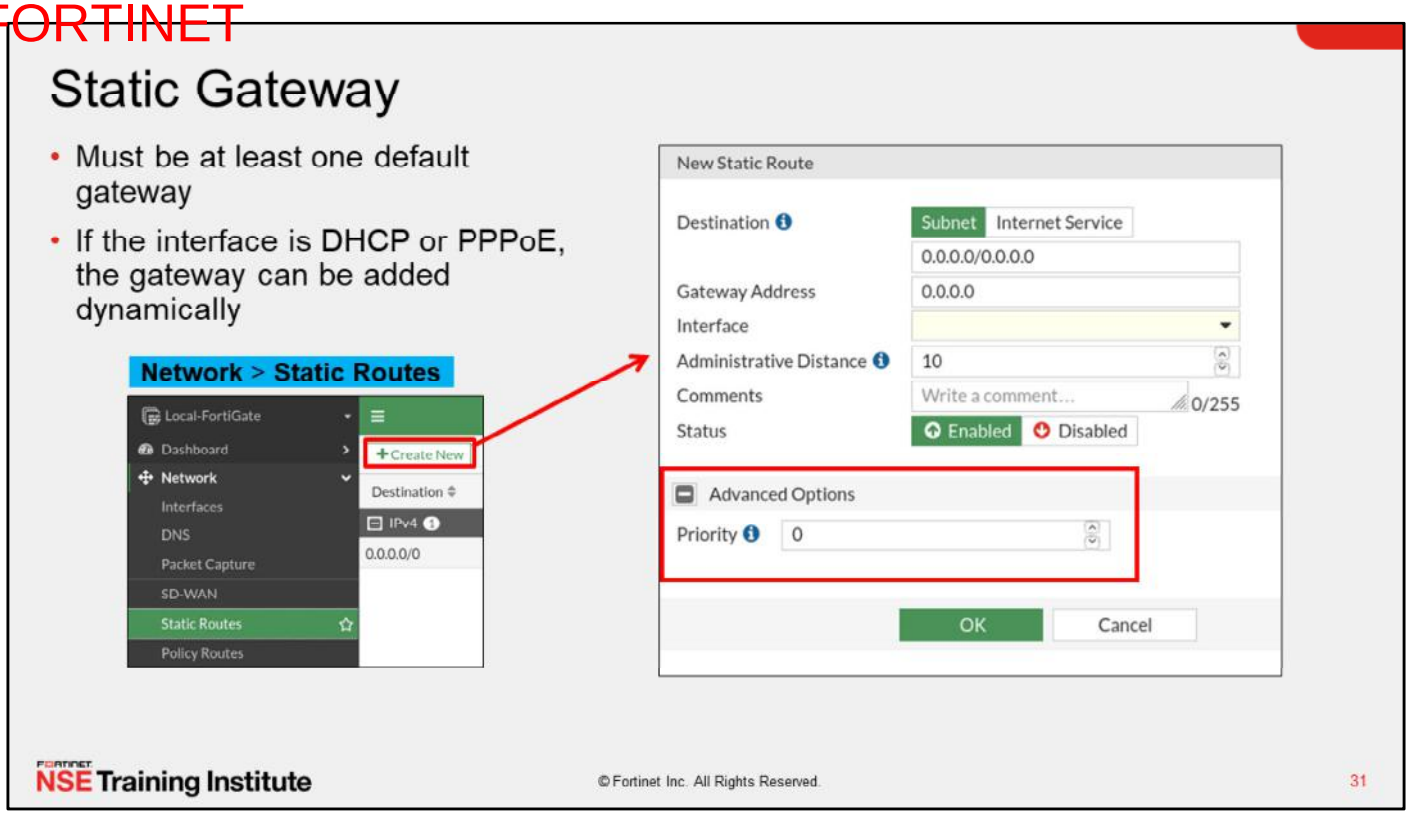

Before you integrate FortiGate into your network, you should configure a default gateway.

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.

Knowledge Check

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

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

NSE Training Institute

32

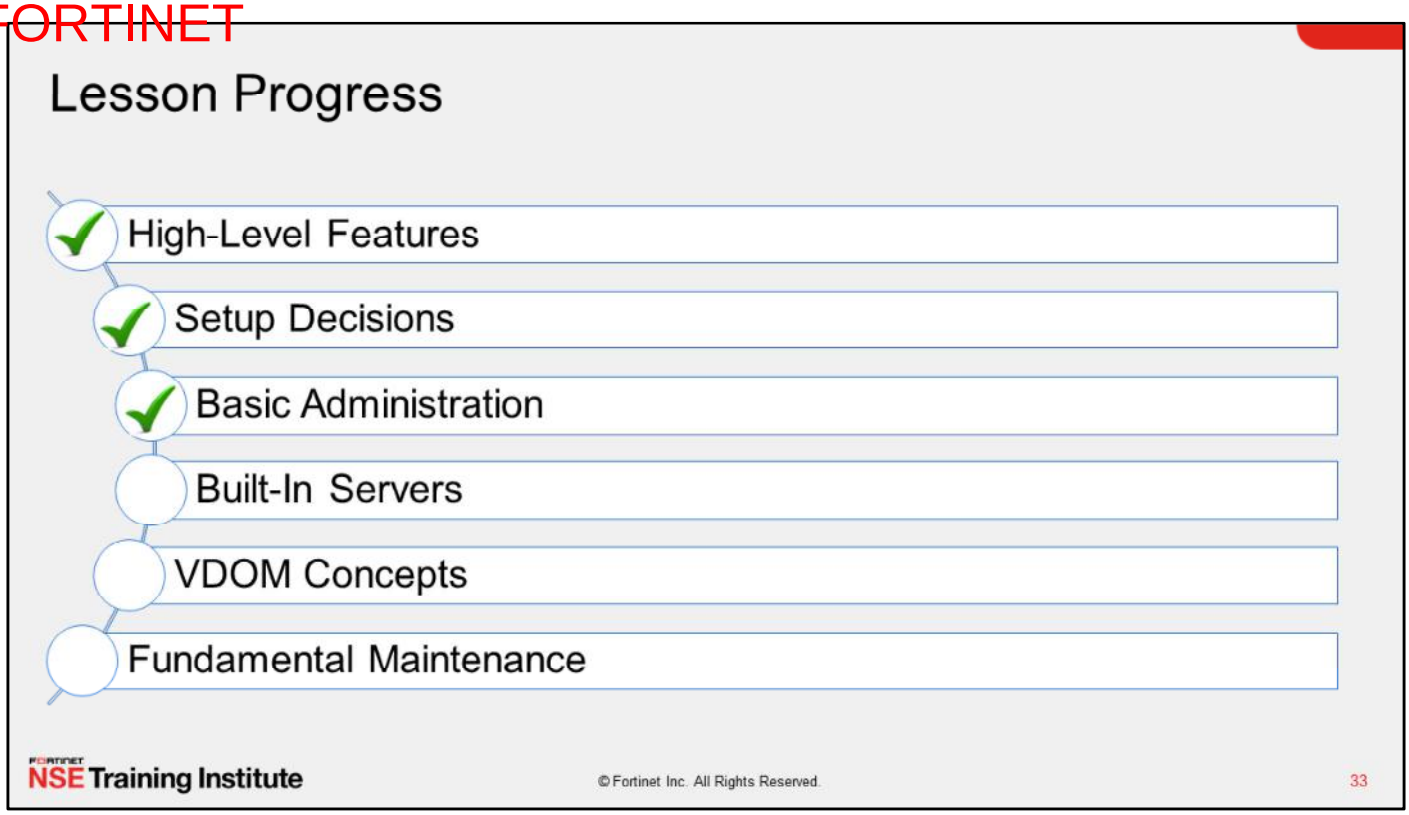

Good job! You now have the knowledge needed to carry out some basic administrative tasks. Now, you'll learn about built-in servers.

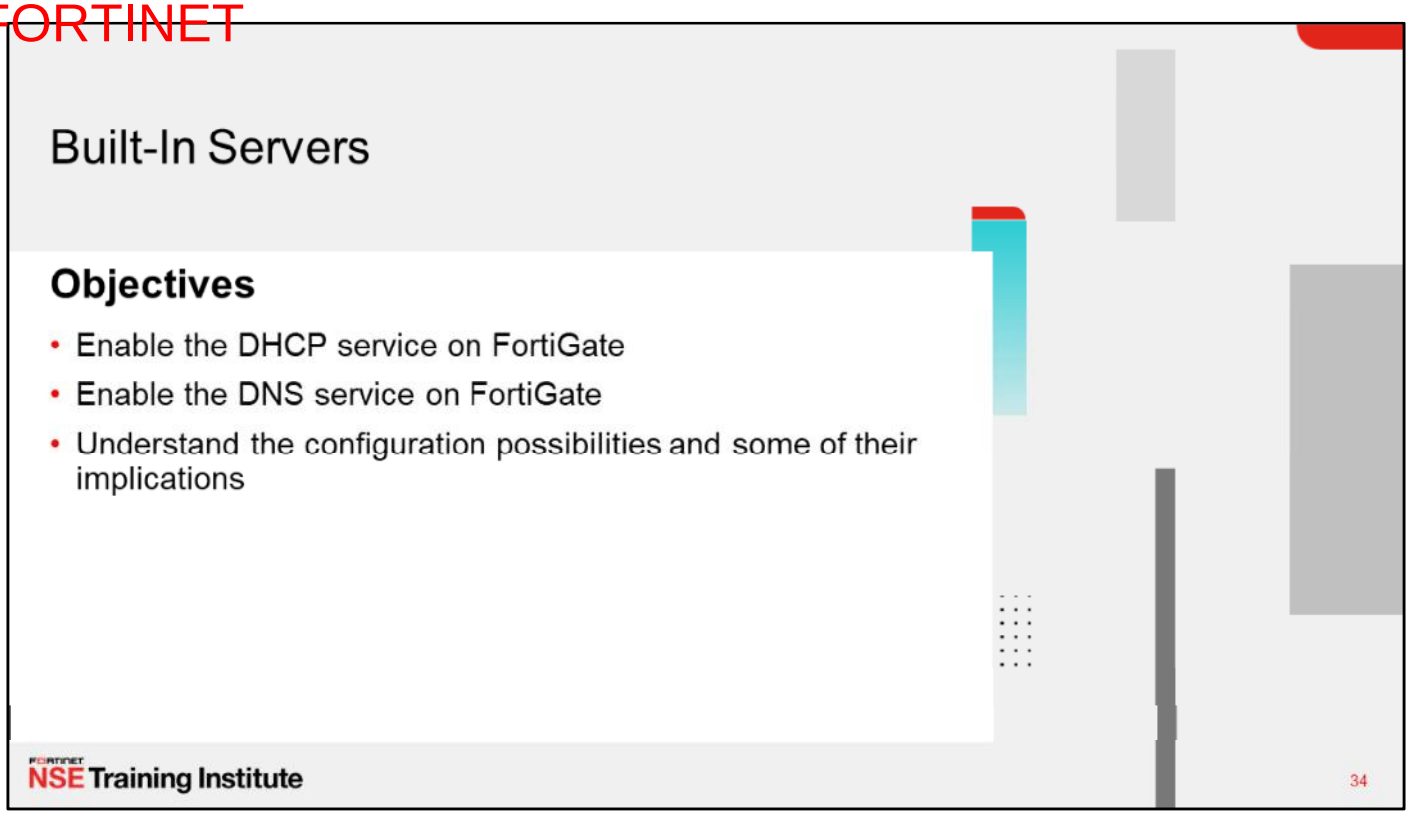

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in implementing the DHCP and DNS built-in servers, you will know how to provide these services through FortiGate.

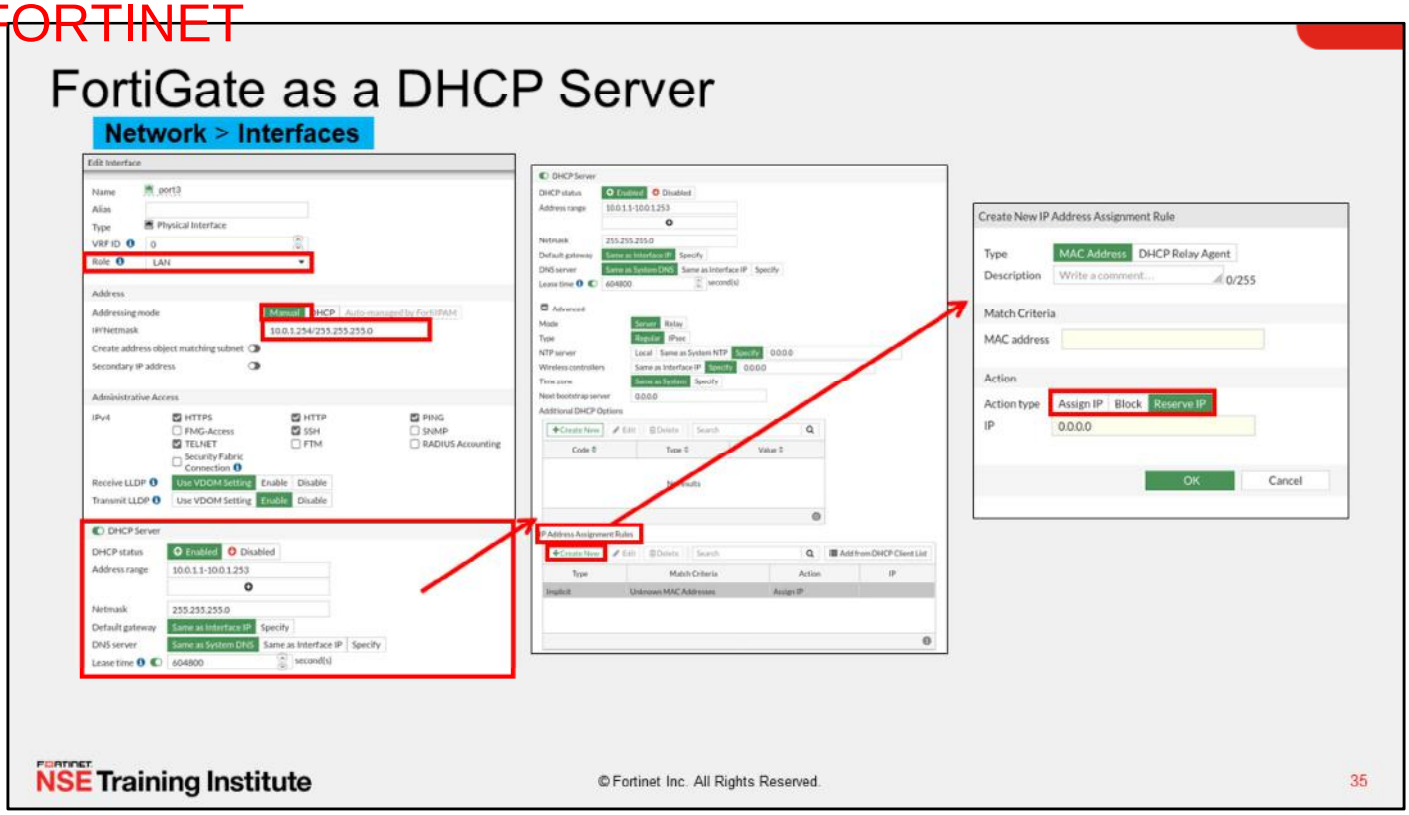

Wireless clients are not the only ones that can use FortiGate as their DHCP server.

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.

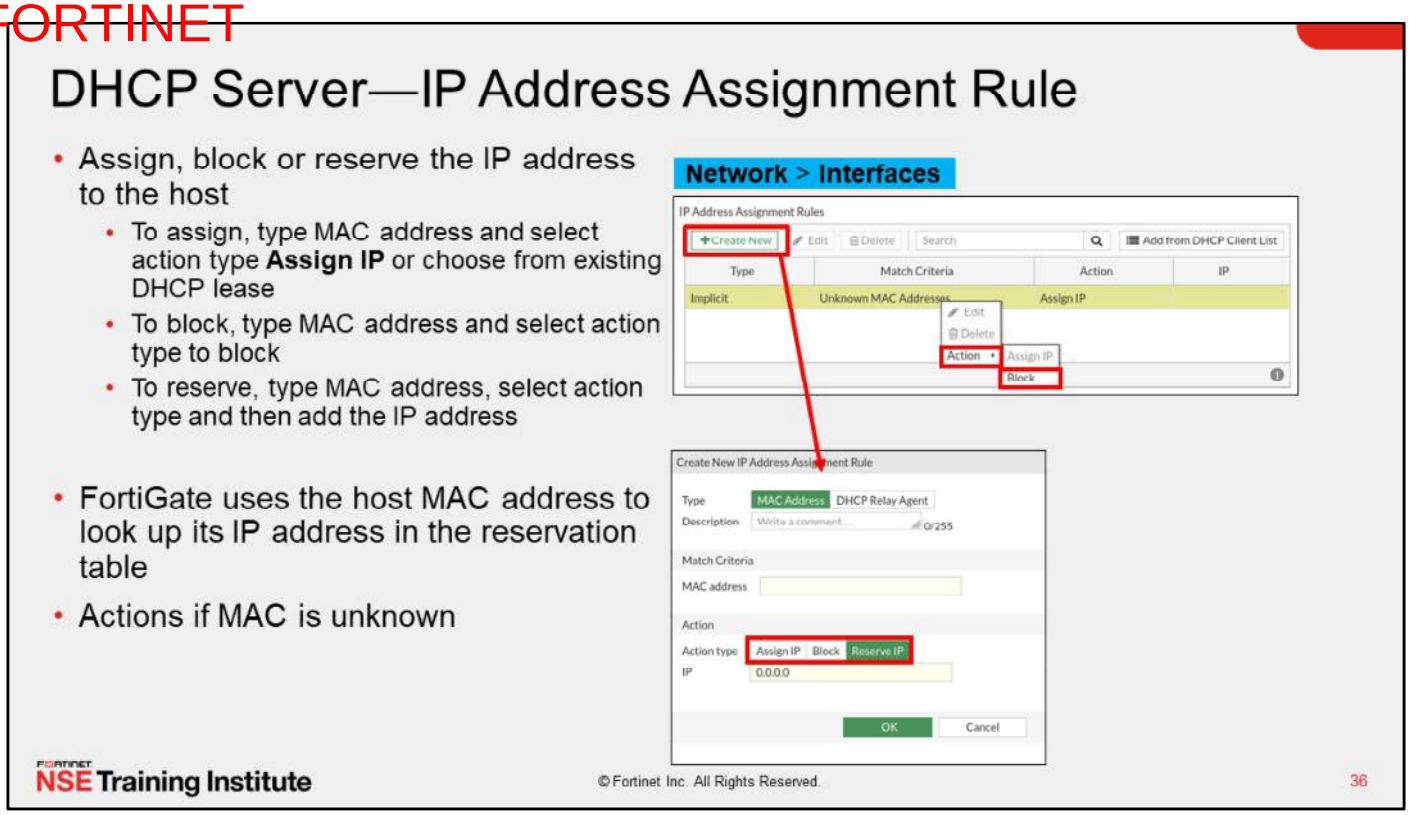

For the built-in DHCP server, you can reserve specific IP addresses for devices with specific MAC addresses.

The action selected for **Unknown MAC Addresses** defines what the FortiGate DHCP server does when it gets a request from a MAC address that is not explicitly listed. The default action is **Assign IP**; however, you can change the default action type to **Assign IP** or **Block**.

- **Assign IP**: permits the DHCP server to assign from its pool of addresses to the identified MAC address. A device receiving an IP address will always receive the same address provided that its lease has not expired.
- **Block**: is the computer with the identified MAC address and the **Block** option will not receive an IP address.
- **Reserve IP**: allows you to bind a specific IP to a MAC address.

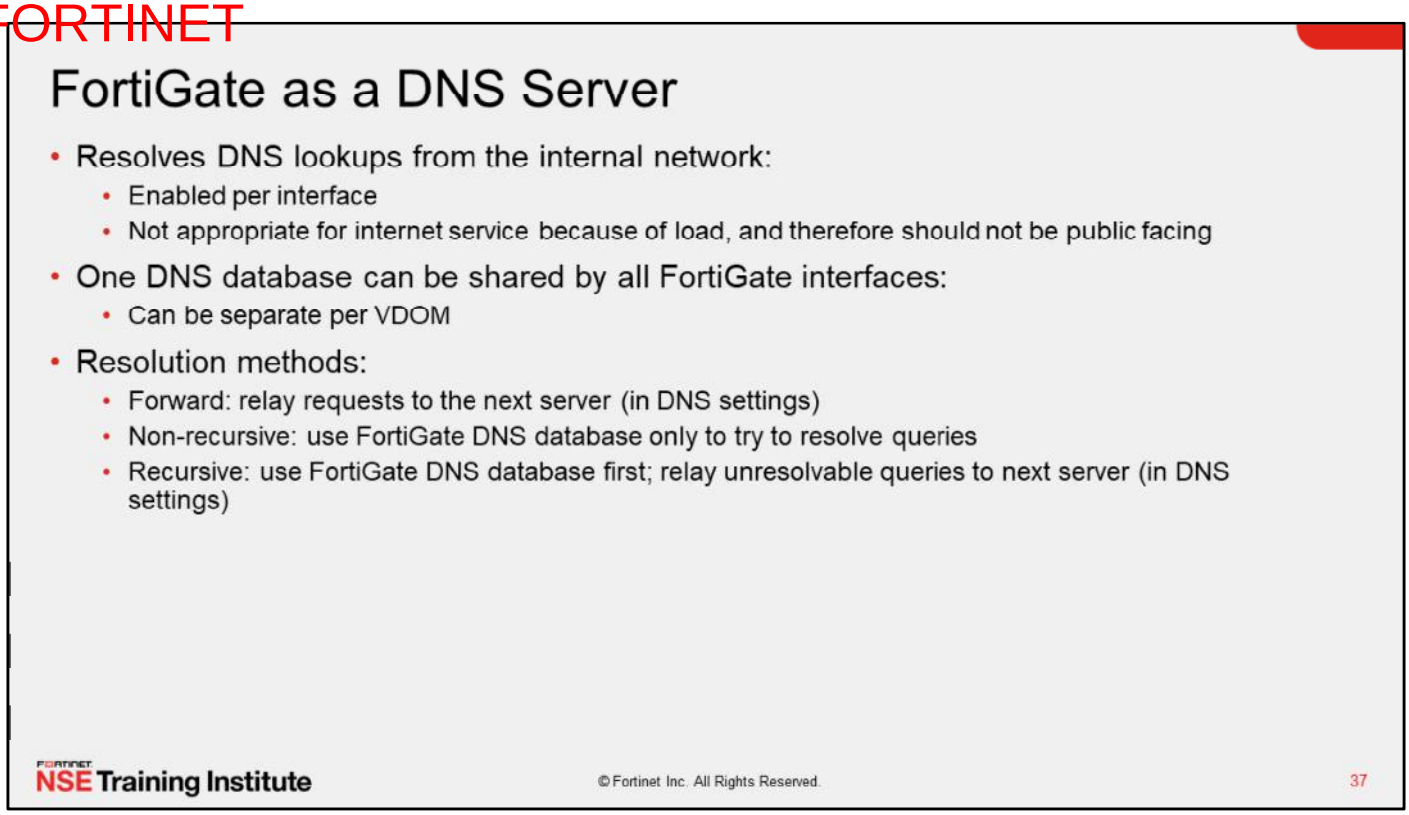

You can configure FortiGate to act as your local DNS server. You can enable and configure DNS separately on each interface.

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.

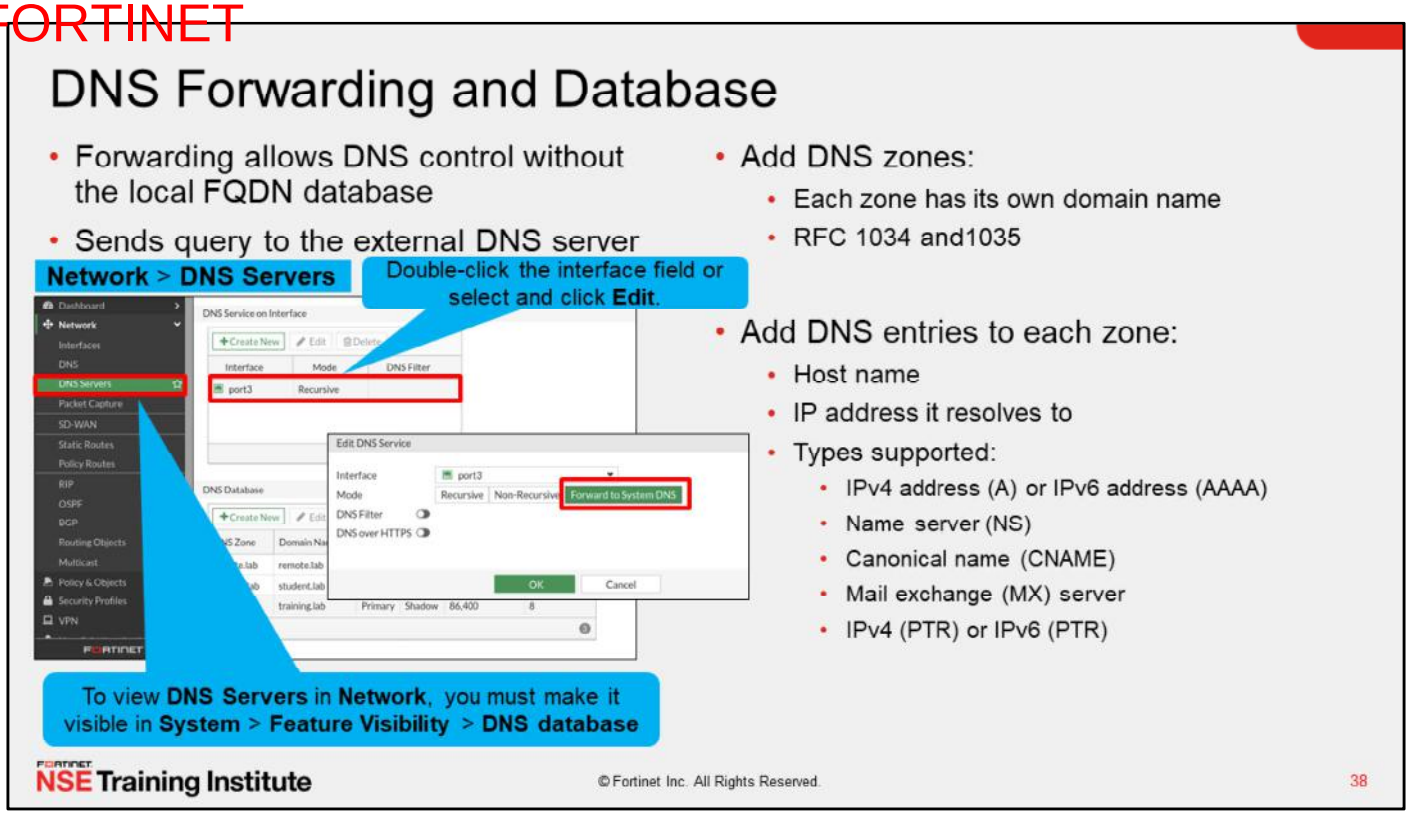

If you select **Recursive**, FortiGate queries its own database before forwarding unresolved requests to the external DNS servers.

If you select **Forward to System DNS**, you can control DNS queries within your own network, without having to enter any DNS names in the FortiGate DNS server.

If you choose to have your DNS server resolve queries, or you choose a split DNS, you must set up a DNS database on your FortiGate device.

This defines the host names that FortiGate resolves queries for. Note that FortiGate currently supports only the DNS record types listed on this slide.

Knowledge Check

1. When configuring FortiGate as a DHCP server, to restrict access by MAC address, what does the **Assign IP** option do?
   A. Assigns a specific IP address to a MAC address
   ✓ B. Dynamically assigns an IP to a MAC address

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

NSE Training Institute

39

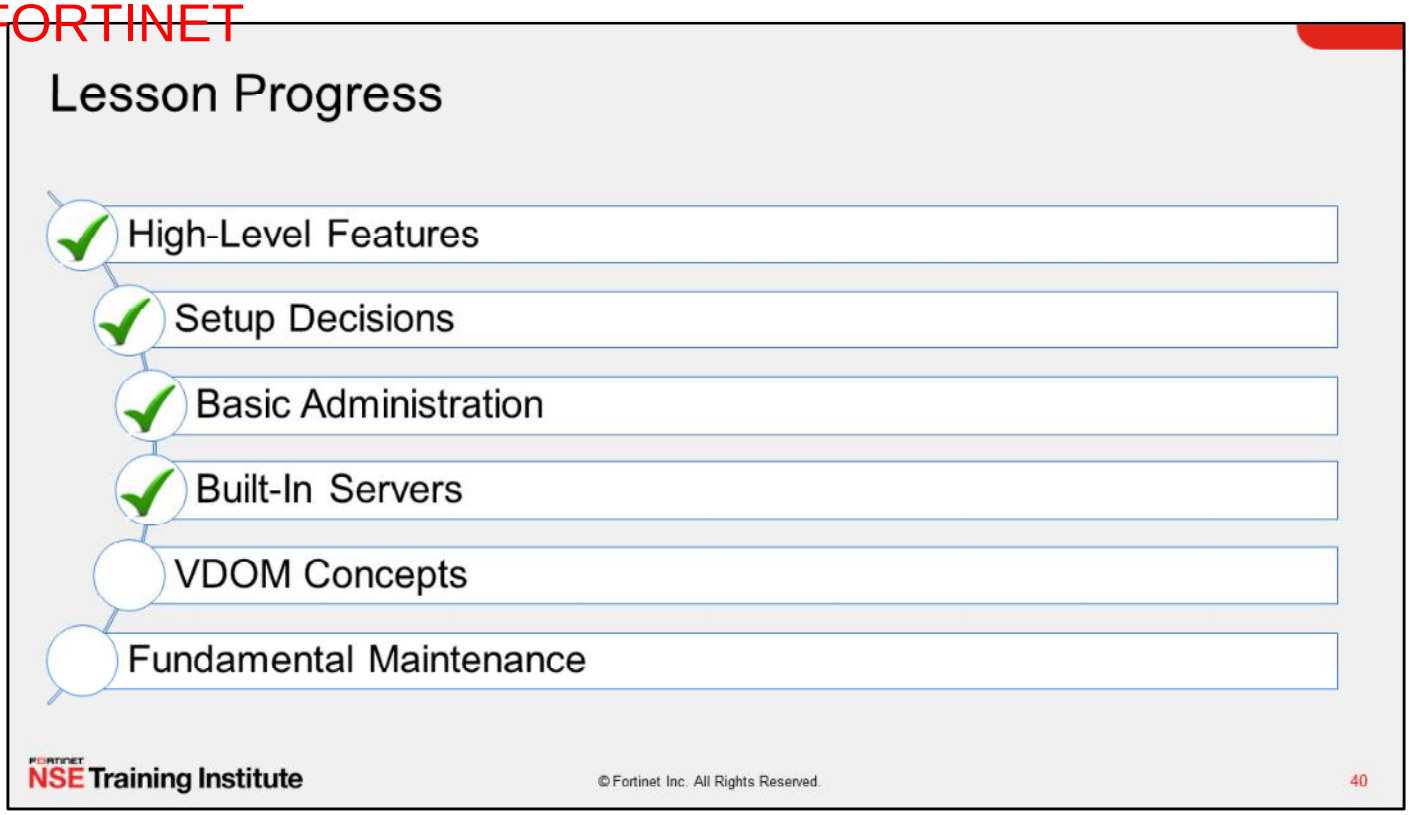

Good job! You now know how to enable DHCP and DNS services on FortiGate, and have some understanding of configuration possibilities. Now, you will learn about VDOM concept.

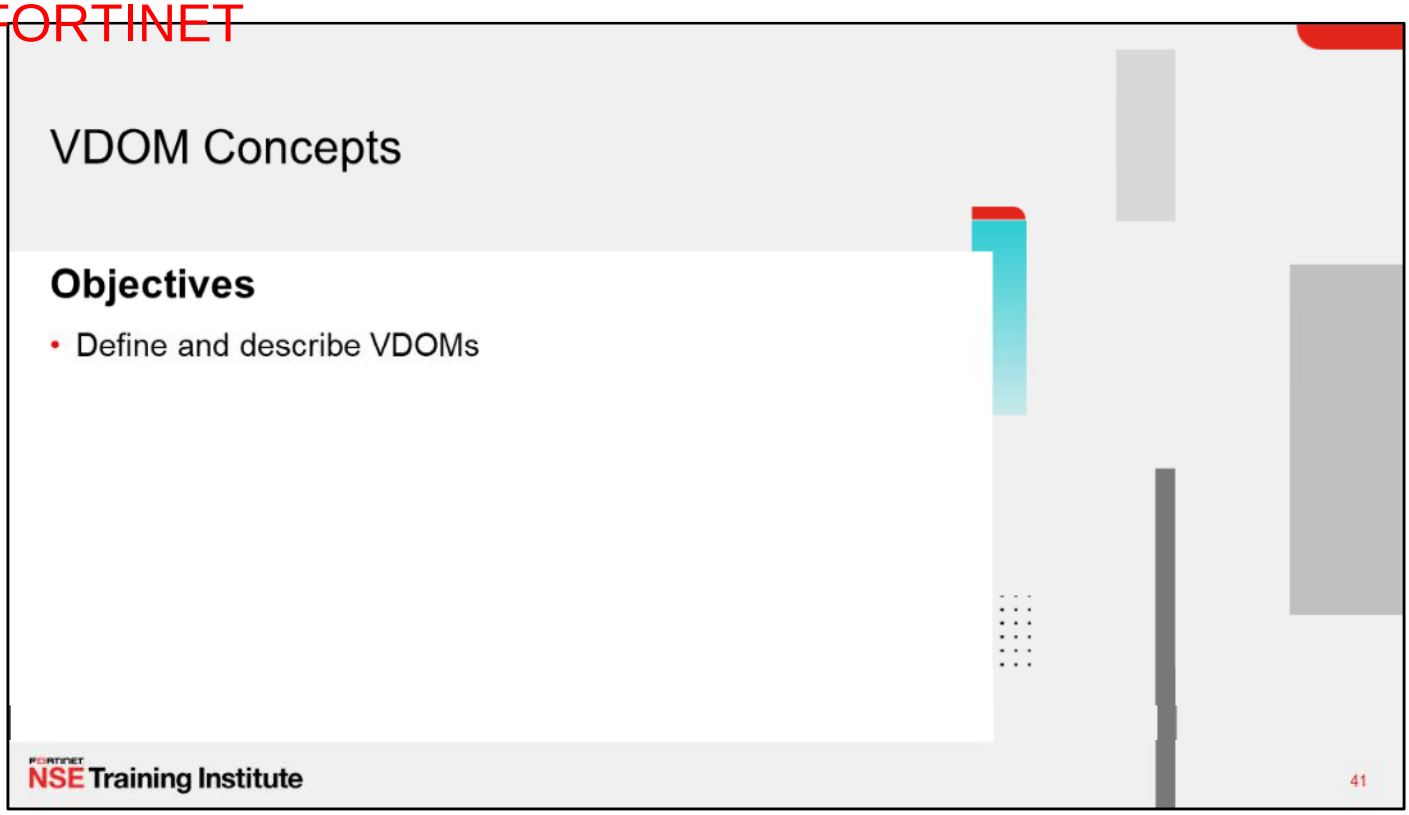

After completing this section, you should be able to achieve the objective shown on this slide.

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

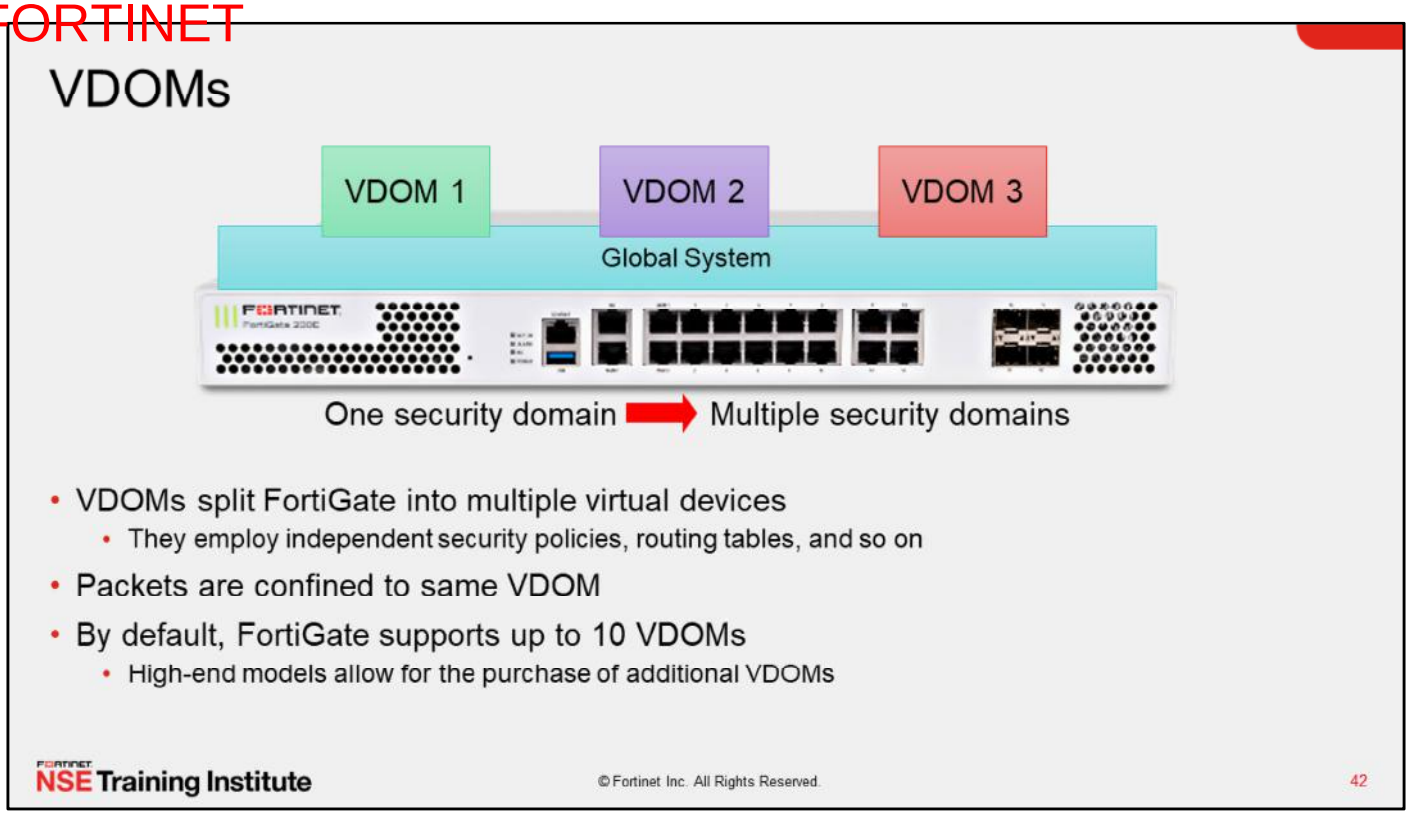

What if, more than segmenting your network, you want to subdivide policies and administrators into multiple security domains?

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.

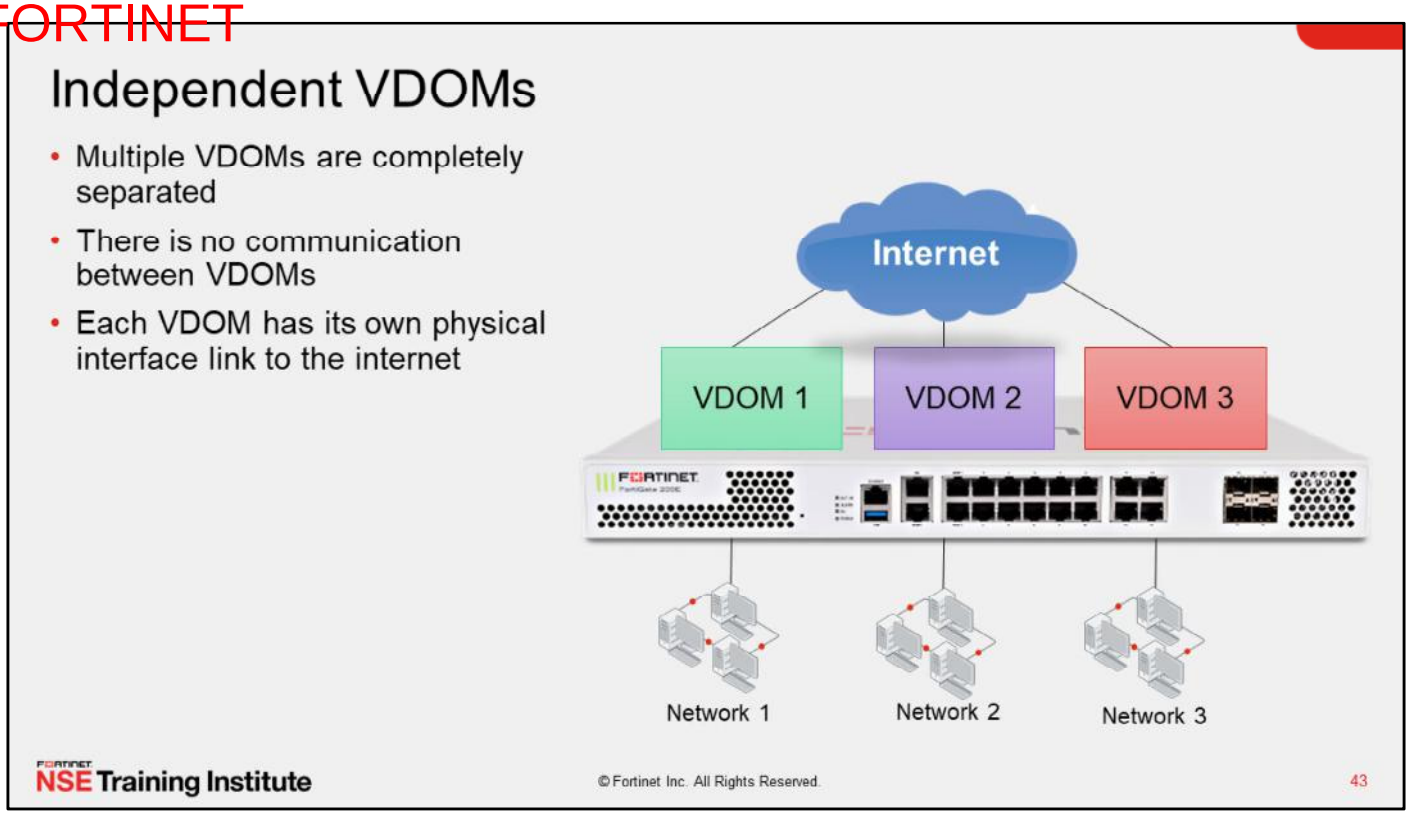

There are a few ways you can arrange your VDOMs. In the topology shown on this slide, each network accesses the internet through its own VDOM.

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

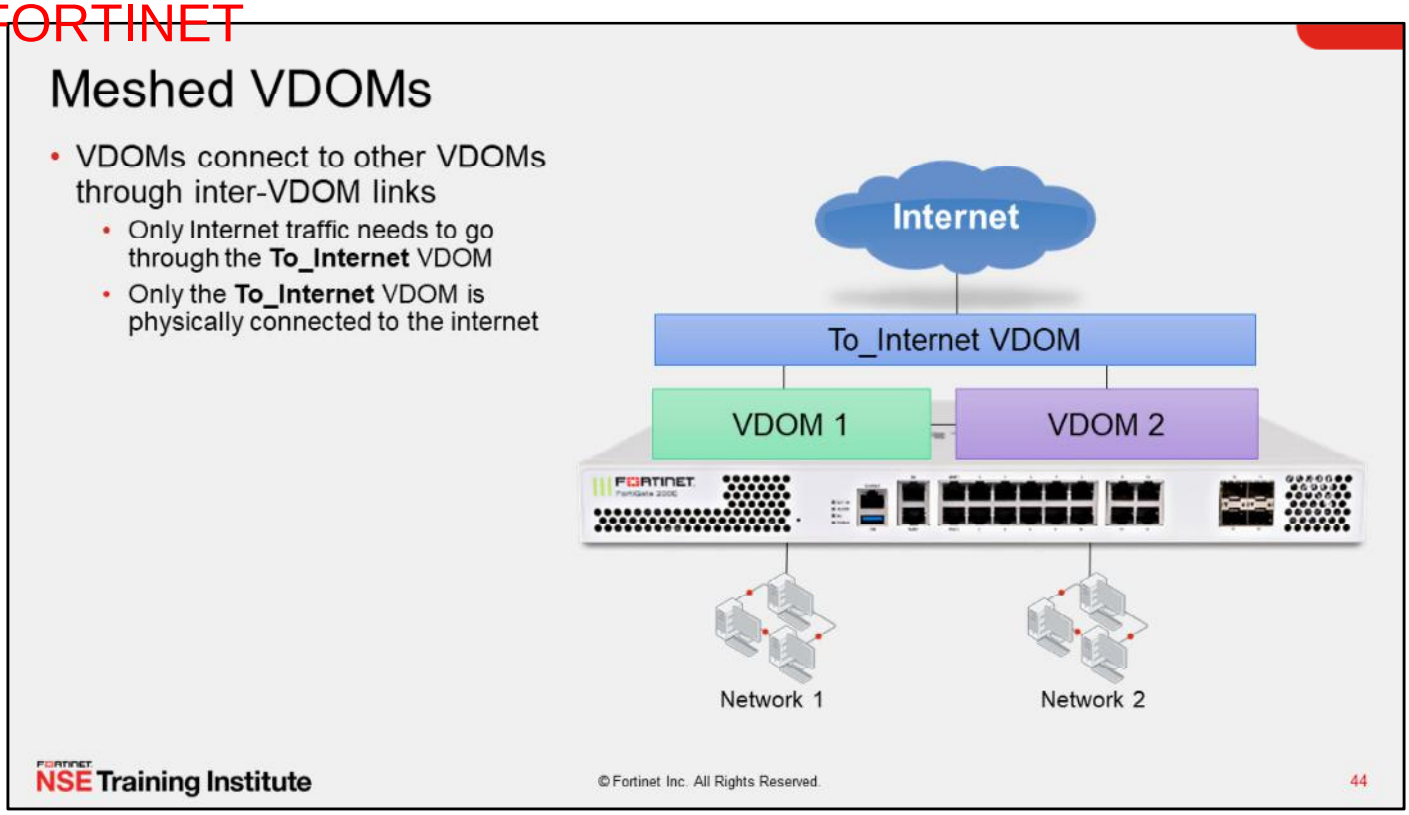

In the example topology shown on this slide, traffic again flows through a single pipe in the **To_Internet** VDOM toward the internet. Traffic between VDOMs doesn't need to leave FortiGate.

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

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

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

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

Management VDOM

- Where all the management traffic for FortiGate originates
- It *must* have access to all global services that FortiGate requires:
  - NTP
  - FortiGuard updates and queries
  - SNMP
  - DNS filtering
  - Logs—both FortiAnalyzer and syslog
  - As well as other FortiGate management-related services
- By default, the management VDOM is **root**
  - Can be reassigned to *any* VDOM in multi-vdom mode, but internet access is recommended because specific services, such as web filtering using the public FortiGuard servers, will not work without it

NSE Training Institute © Fortinet Inc. All Rights Reserved. 45

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

Some system daemons, such as NTP and FortiGuard updates, generate traffic coming from FortiGate. One, and only one, of the VDOMs on a FortiGate device is assigned the role of the management VDOM. Traffic coming from FortiGate to those global services originates from the management VDOM. By default, the root VDOM acts as the management VDOM, but you can manually reassign this task to a different VDOM in multi-vdom mode.

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

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

Knowledge Check

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

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

NSE Training Institute

46

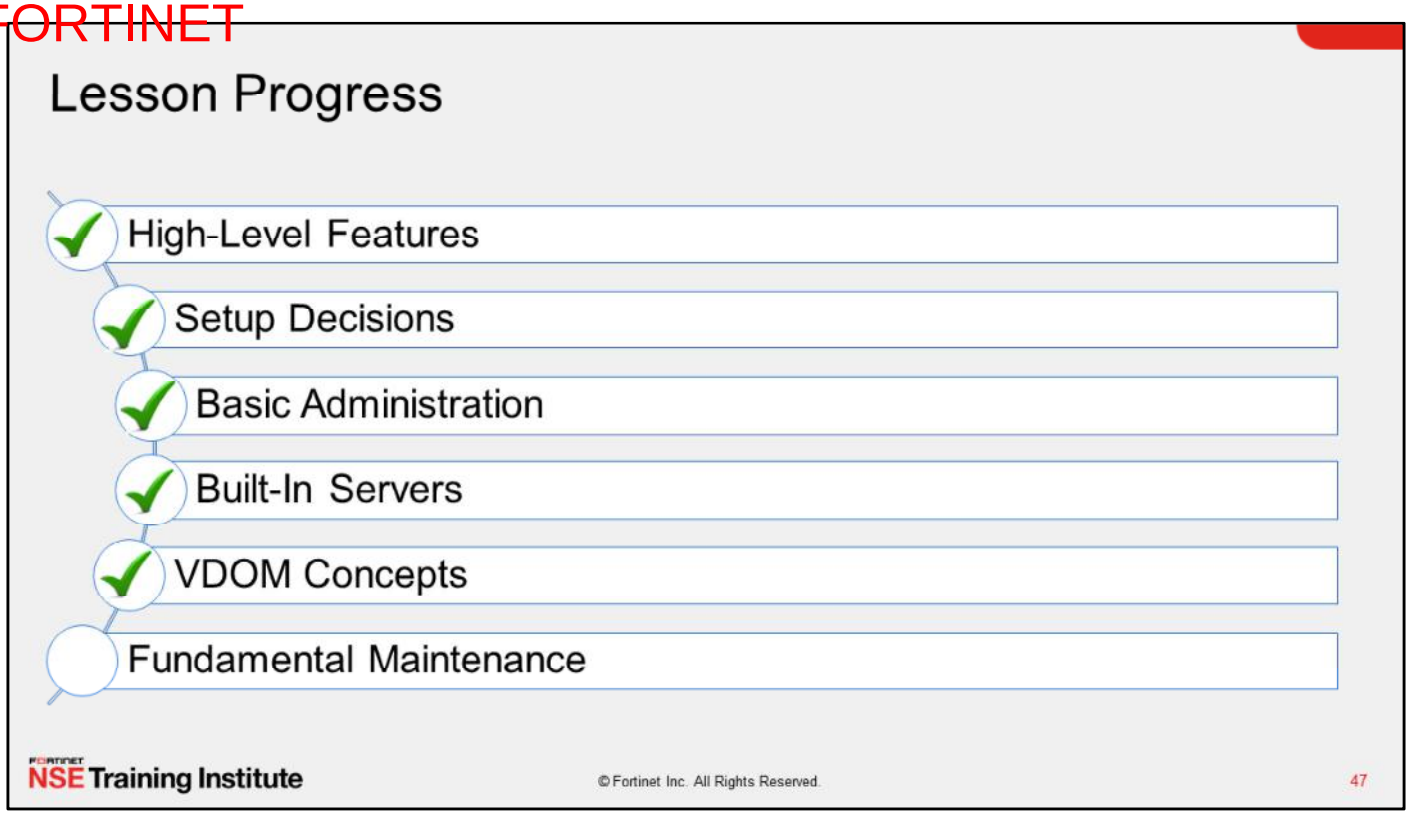

Good job! You now understand some basic concepts about VDOMs.

Now, you will learn about fundamental maintenance.

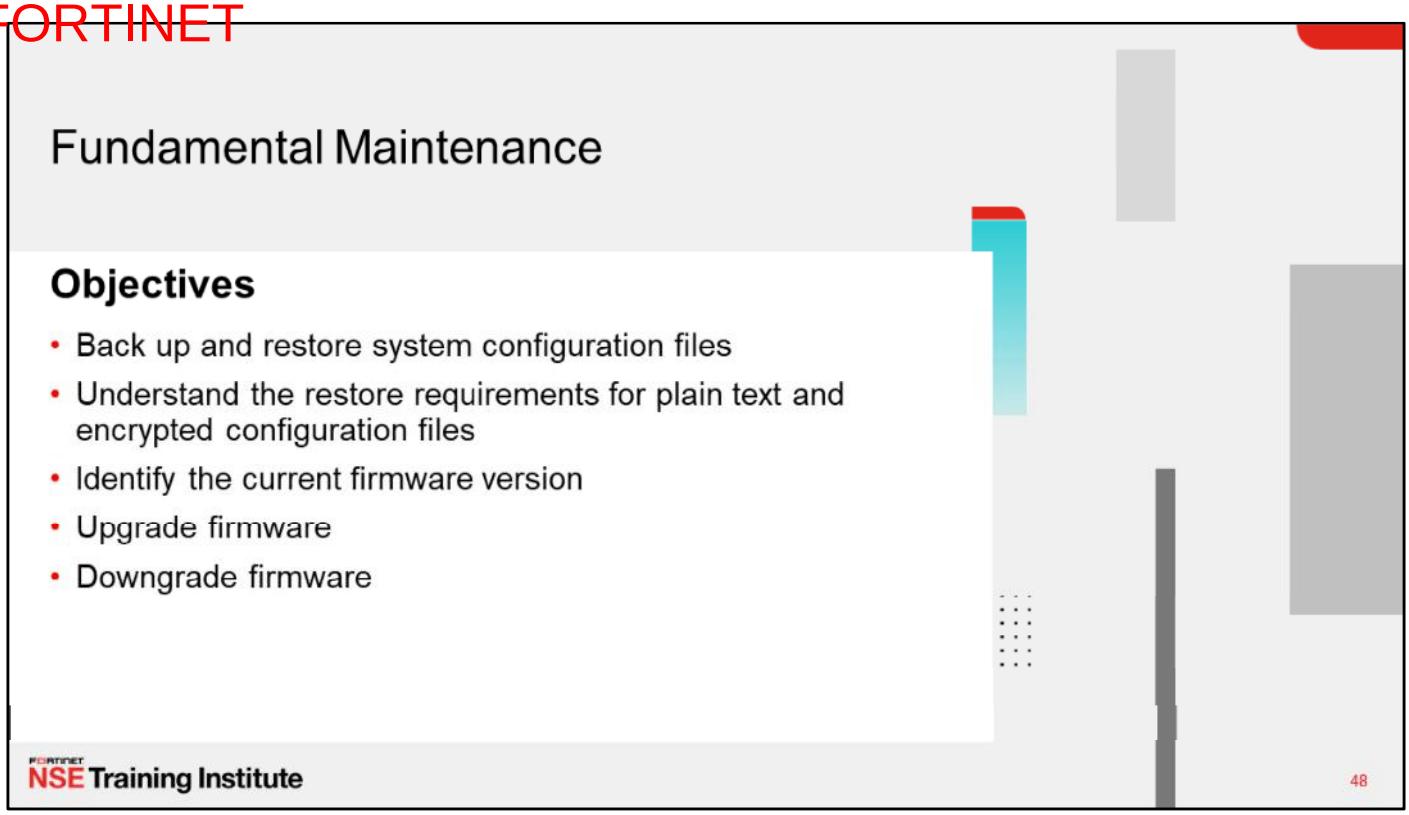

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

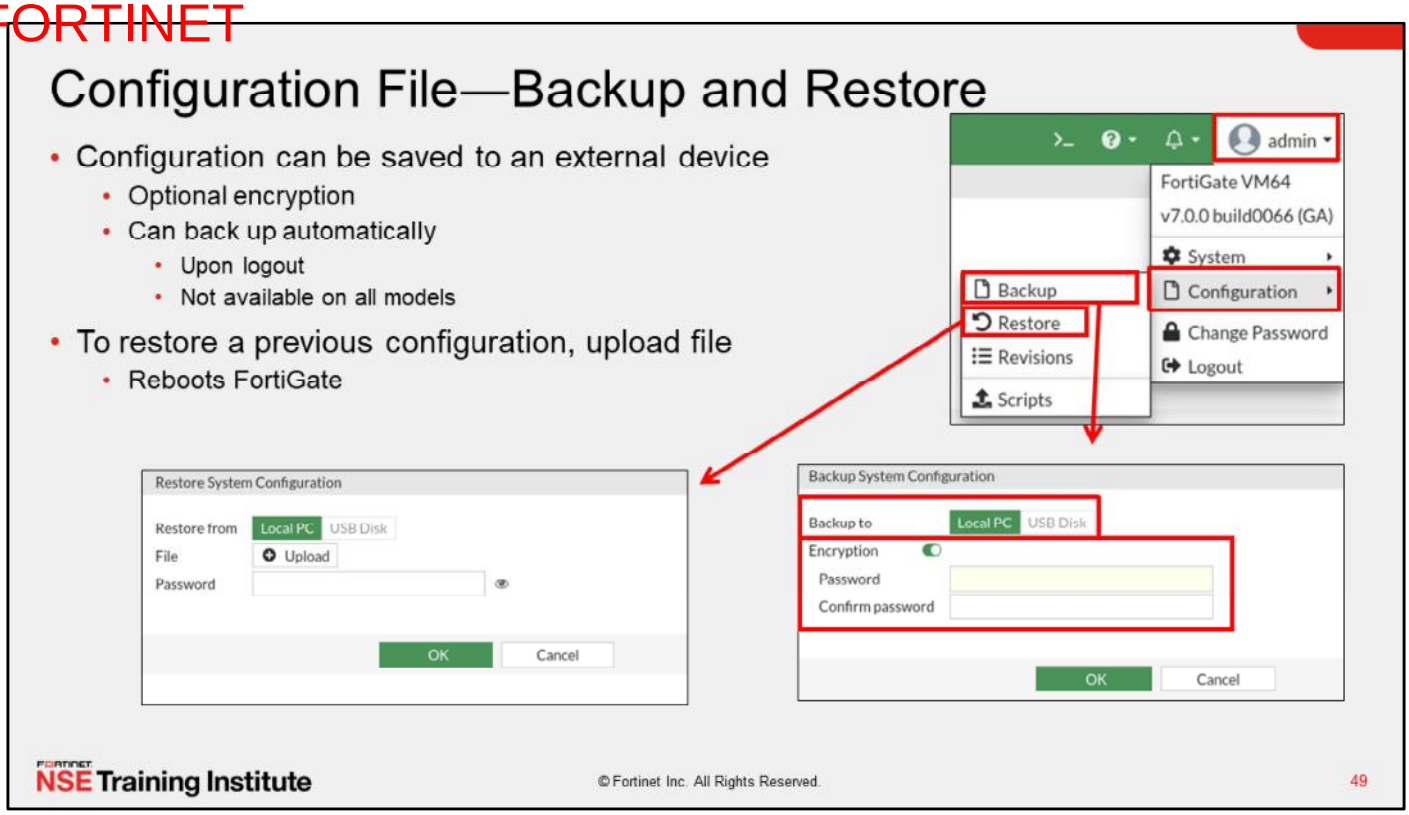

Now that FortiGate has basic network settings and administrative accounts, you will learn how to back up the configuration. In addition to selecting the destination of the backup file, you can choose to encrypt or not to encrypt the backup file. Even if you choose not to encrypt the file, which is the 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.

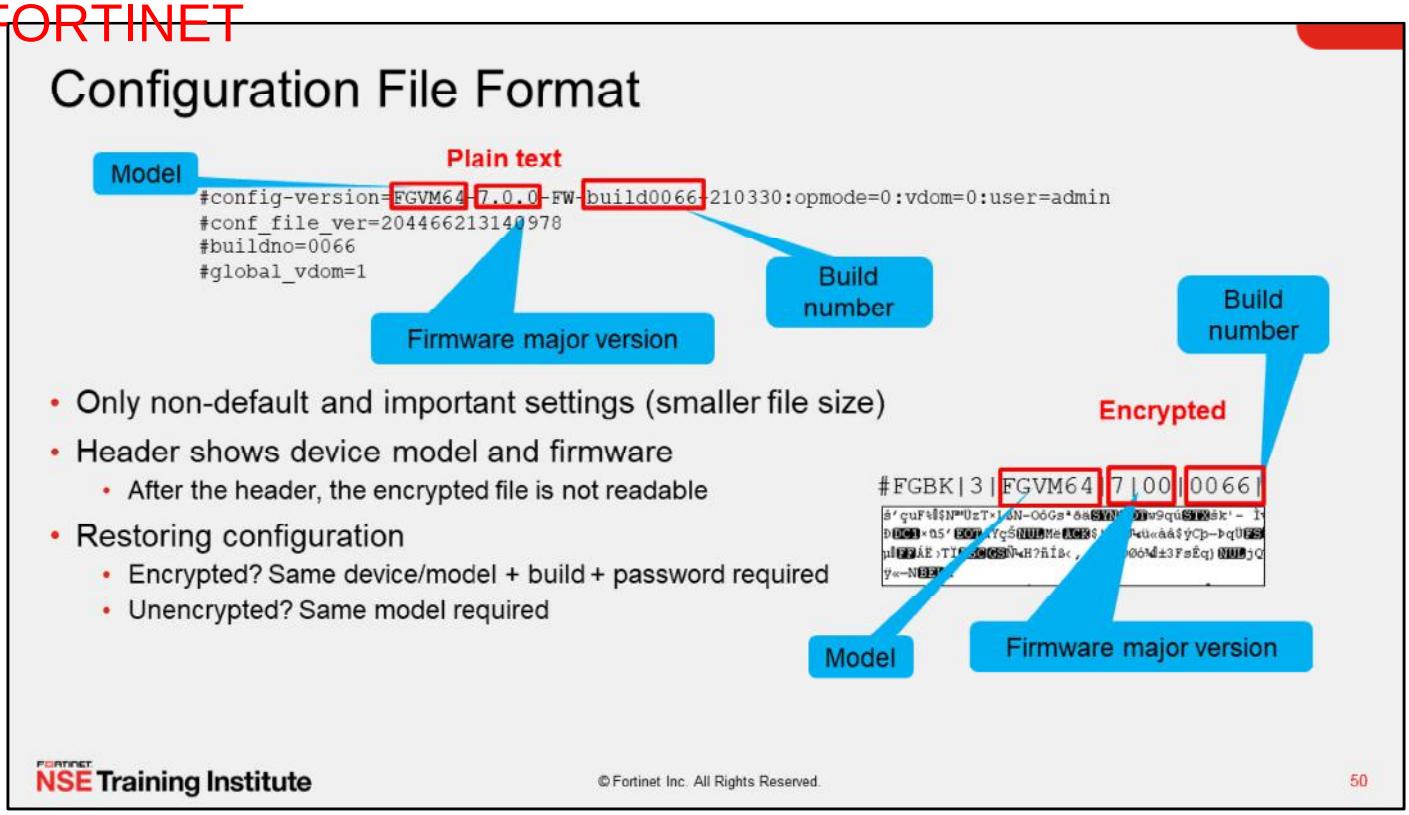

If you open the configuration file in a text editor, you'll see that both encrypted and unencrypted configuration files contain a cleartext header that contains some basic information about the device. The example on this slide shows what information is included. To restore an encrypted configuration, 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.

Upgrade Firmware

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

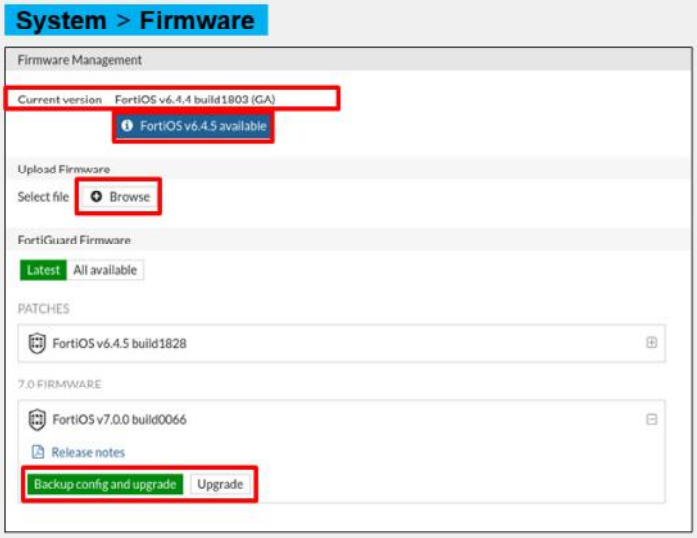

NSE Training Institute

51

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** > **Firmware**. 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 **Firmware** page.

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.

Upgrade Firmware Process

1. Back up the configuration (`full config backup` on GUI or CLI)
2. Download a copy of the current firmware, in case you need to revert
3. Have physical access, or a terminal server connected to local console, in case you need to revert
4. Read the *Release Notes*; they include the upgrade path and other useful information
5. Perform the upgrade

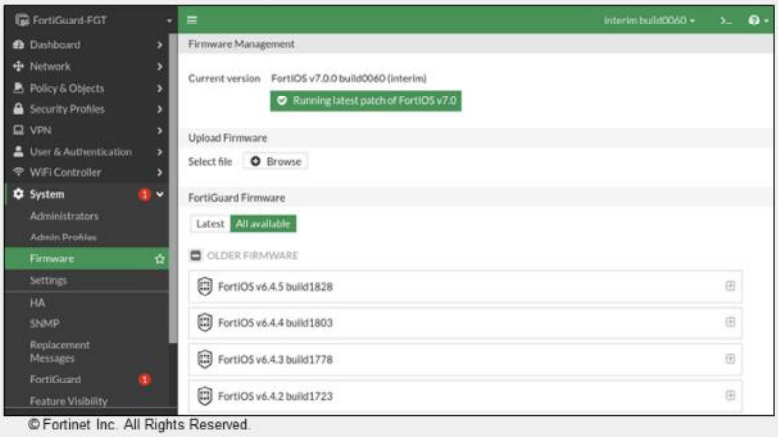

NSE Training Institute

52

Upgrading the firmware on FortiGate is simple. Click **System** > **Firmware**, and then browse to the firmware file that you have downloaded from `support.fortinet.com` or choose to upgrade online.

If you want to do a clean installation by overwriting both the existing firmware and its current configuration, you can do this using the local console CLI, within the boot loader menu, while FortiGate is rebooting. However, this is not the usual method.

Downgrade Firmware Process

1. Get the pre-upgrade configuration file
2. Download a copy of the current firmware, in case you need to revert
3. Have physical access, or a terminal server connected to the local console, in case you need to revert
4. Read the *Release Notes* (Does downgrade preserve configuration?)
5. Downgrade the firmware
6. If required, upload the configuration that matches the firmware version

NSE Training Institute © Fortinet Inc. All Rights Reserved. 53

You can also downgrade the firmware. Because settings change in each firmware version, you should have a configuration file in the syntax that is compatible with the firmware.

Remember to read the *Release Notes*. Sometimes a downgrade between firmware versions that preserves the configuration is not possible. In that situation, the only way to downgrade is to format the disk, then reinstall.

After you've confirmed that the downgrade is possible, verify everything again, then start the downgrade. After the downgrade completes, restore a configuration backup that is compatible with that version.

Why should you keep emergency firmware and physical access?

Earlier firmware versions do not know how to convert later configurations. Also, when upgrading through a path that is not supported by the configuration translation scripts, you *might* lose all settings except basic access settings, such as administrator accounts and network interface IP addresses. Another rare, but possible, scenario is that the firmware could be corrupted when you are uploading it. For all of those reasons, you should always have local console access during an upgrade. However, in practice, if you read the *Release Notes* and have a reliable connection to the GUI or CLI, it should not be necessary.

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

NSE Training Institute

54

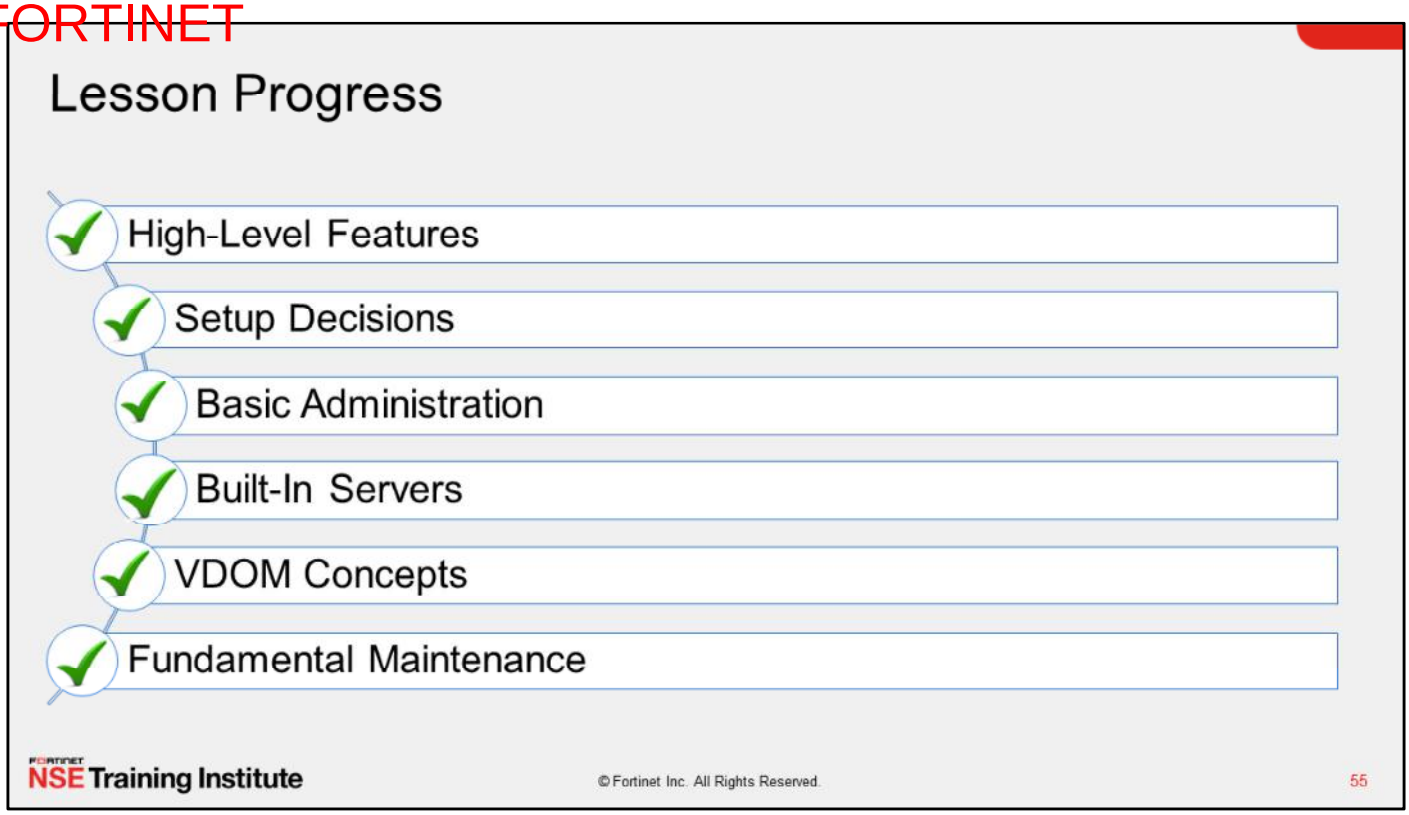

Congratulations! You have completed this lesson.

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

Review

- ✓ Identify key FortiGate features, services, and built-in servers
- ✓ Identify the differences between the two operating modes, and the relationship between FortiGate and FortiGuard
- ✓ Identify the factory defaults, basic network settings, and console ports
- ✓ Execute basic administration, such as creating administrative users and permissions
- ✓ Define and describe VDOMs
- ✓ Execute backup and restore tasks and discuss the requirements for restoring an encrypted configuration file
- ✓ Initiate an upgrade and downgrade of the firmware

NSE Training Institute

56

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.

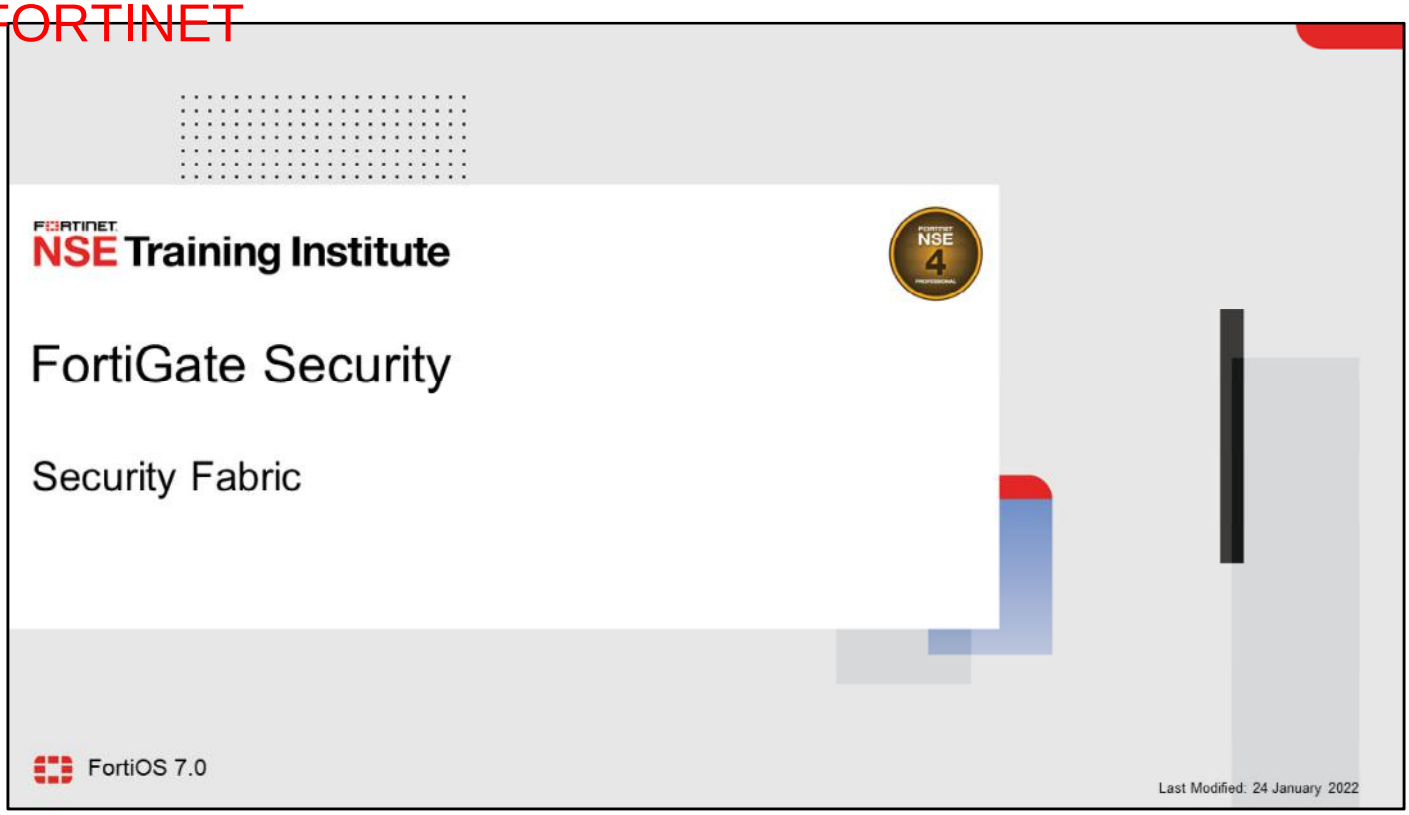

In this lesson, you will learn about the Fortinet Security Fabric.

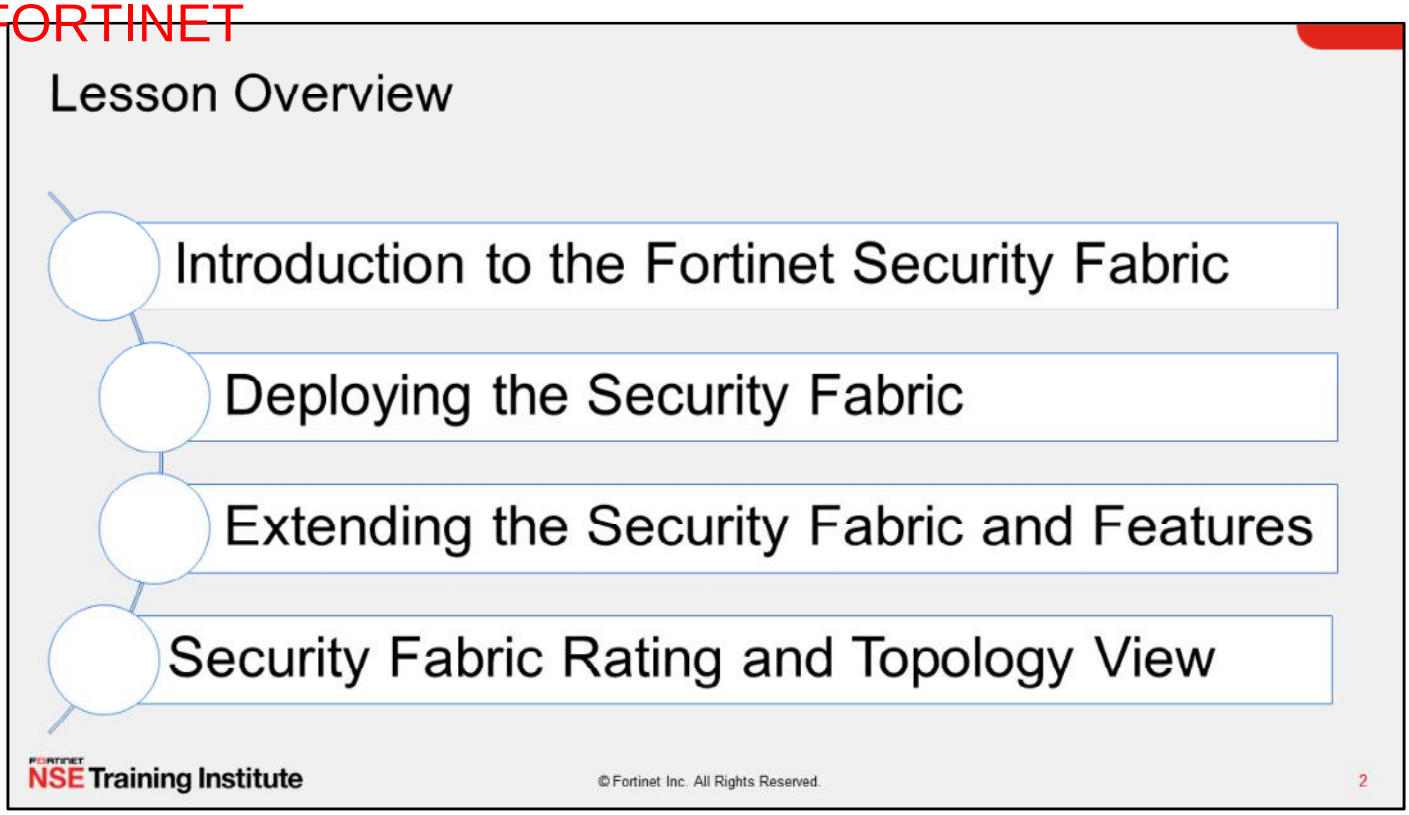

In this lesson, you will learn about the topics shown on this slide.

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.

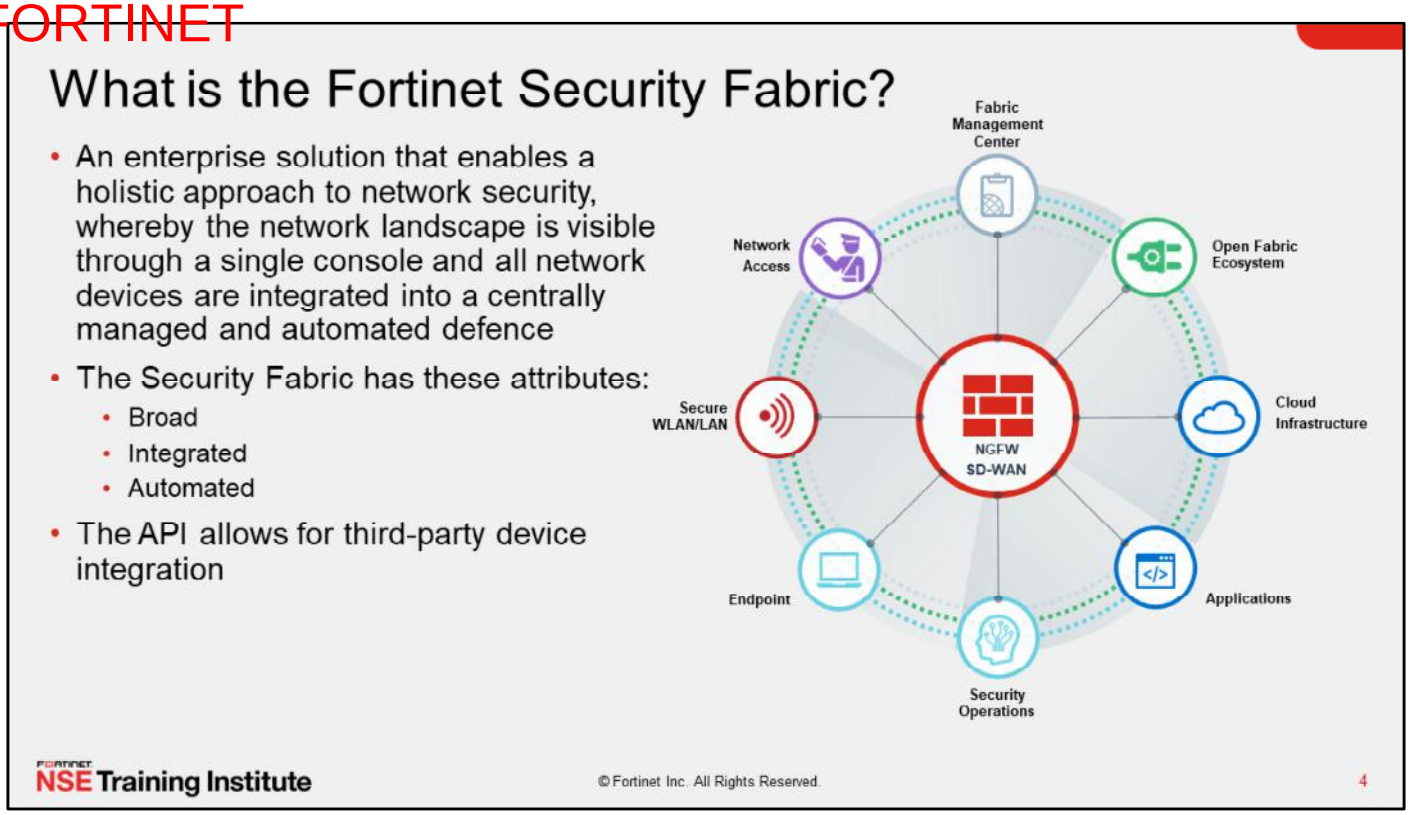

What is the Fortinet Security Fabric?

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

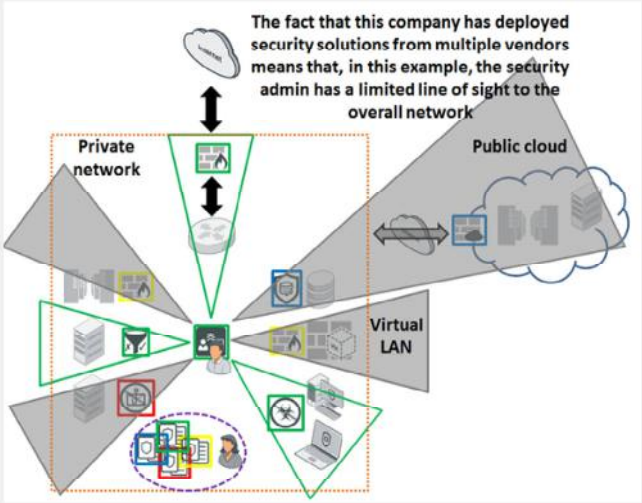

NSE Training Institute

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.

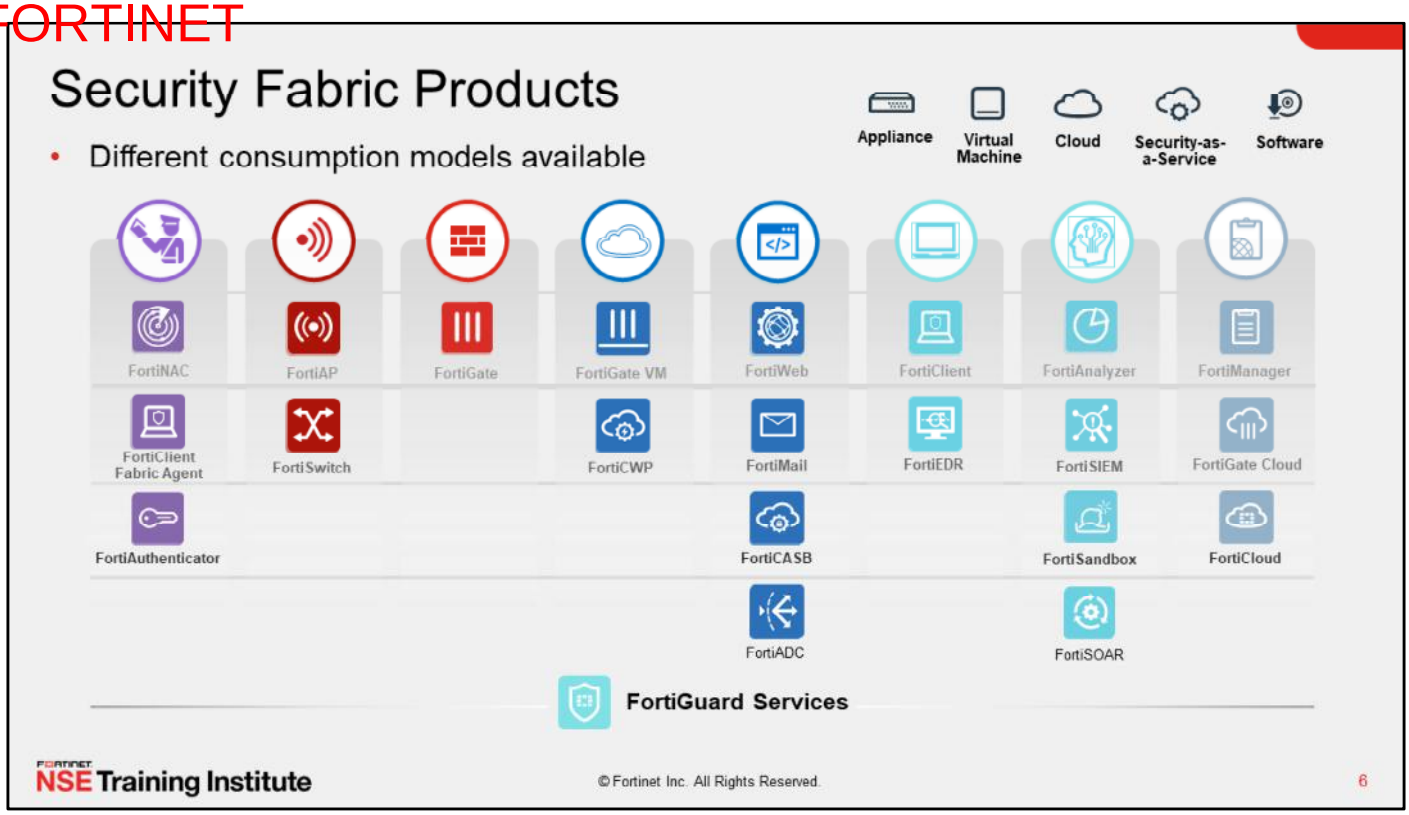

As shown on this slide, the Fortinet Security Fabric offers eight solutions: network access, security WLAN/LAN, public and private cloud infrastructure, applications, endpoint, security operations, open fabric ecosystem, and fabric management center. Each of these solutions is based on specific use cases 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.

Devices That Comprise the Security Fabric

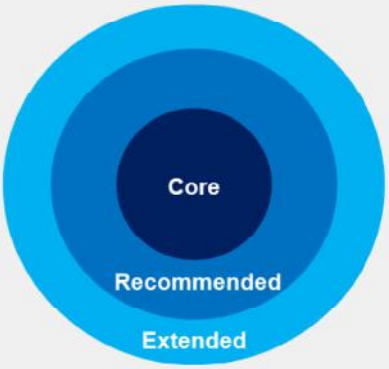

- Core:
  - Two FortiGate devices + FortiAnalyzer
- Recommended–Adds significant visibility or control:
  - FortiManager, FortiAP, FortiSwitch, FortiClient, FortiSandbox, FortiMail, FortiWeb, FortiAI
- Extended—Integrates with fabric, but may not apply to everyone:
  - Other Fortinet products and third-party products using the API

NSE Training Institute

7

FortiGate and FortiAnalyzer creates the core of the Security Fabric. To add more visibility and control, Fortinet recommends adding FortiManager, FortiAP, FortiClient, FortiSandbox, FortiMail, FortiWeb, FortiAI, and FortiSwitch. The solution can be extended by adding other network security devices.

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

NSE Training Institute

8

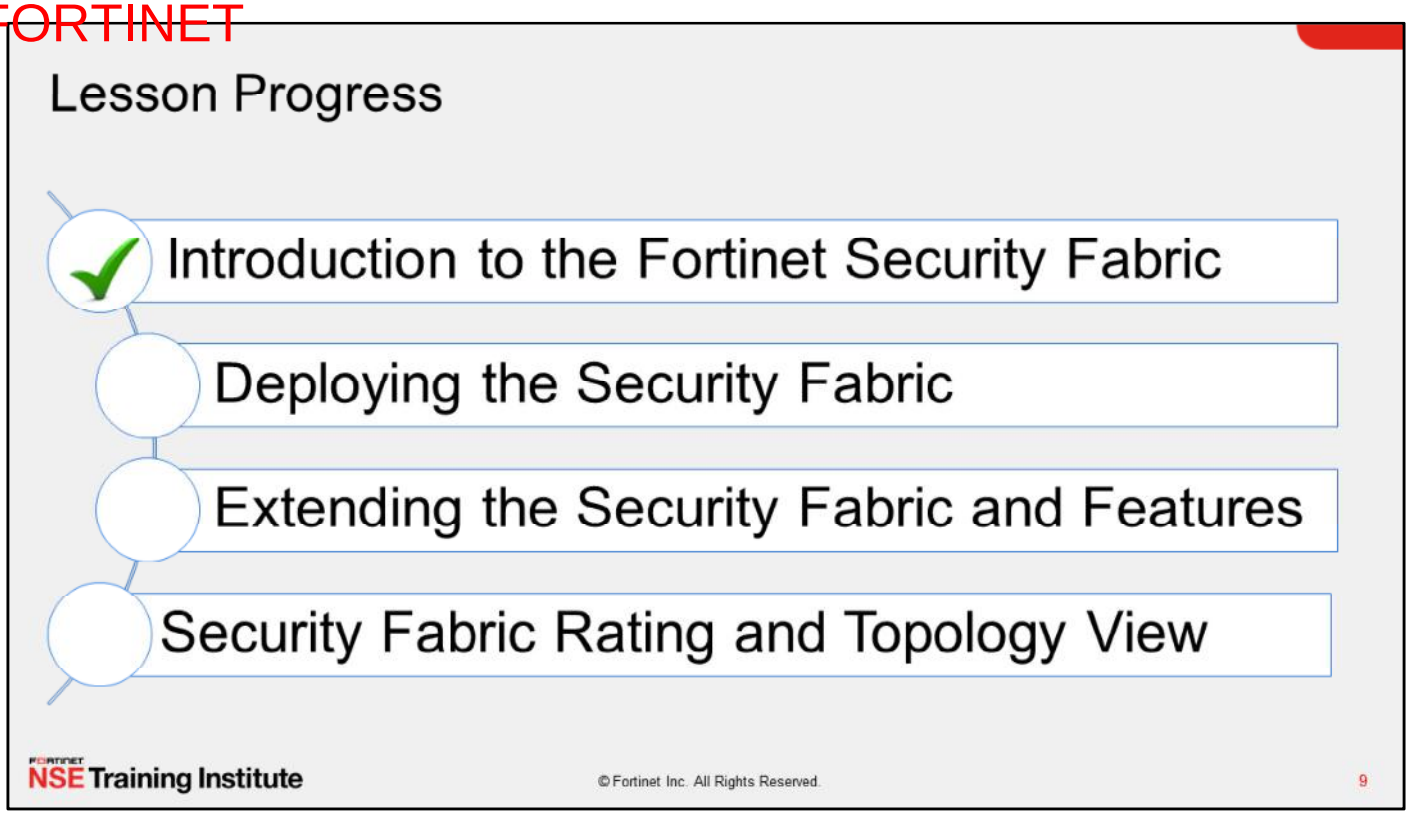

Good job! You now understand the basics of the Fortinet Security Fabric.

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.

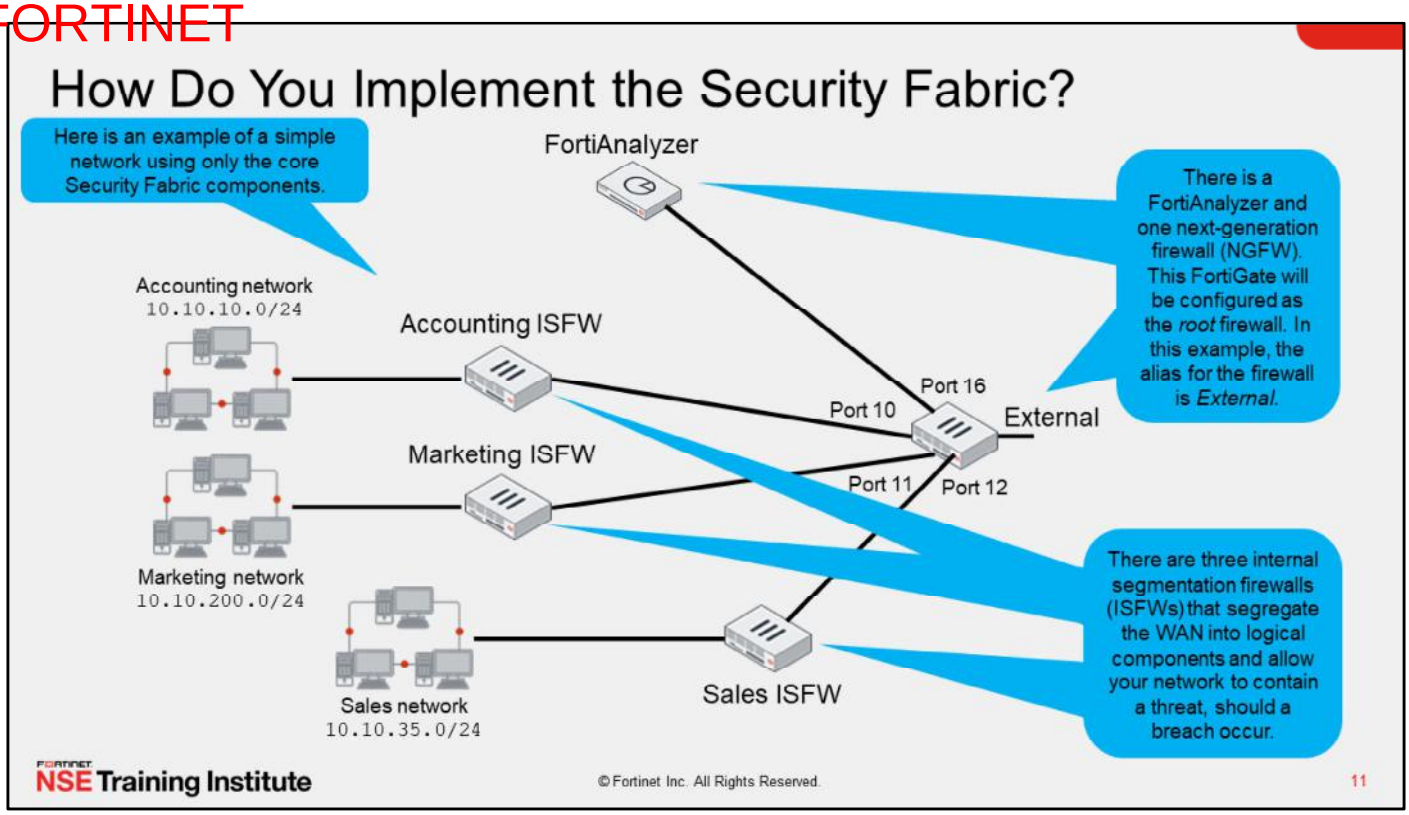

In this simple network that comprises only the core devices of a Security Fabric, there is one FortiAnalyzer and one next-generation firewall (NGFW) FortiGate. This implementation example is intended to be a high-level view only. For more detail, see `docs.fortinet.com`. The FortiGate device named External is acting as the edge firewall and will also be configured as the *root* firewall within the Security Fabric. Downstream from the root firewall there are three internal segmentation firewalls that compartmentalize the WAN in order to contain a breach and control access to various LANs. In this example, there are Accounting, Marketing, and Sales LANs.

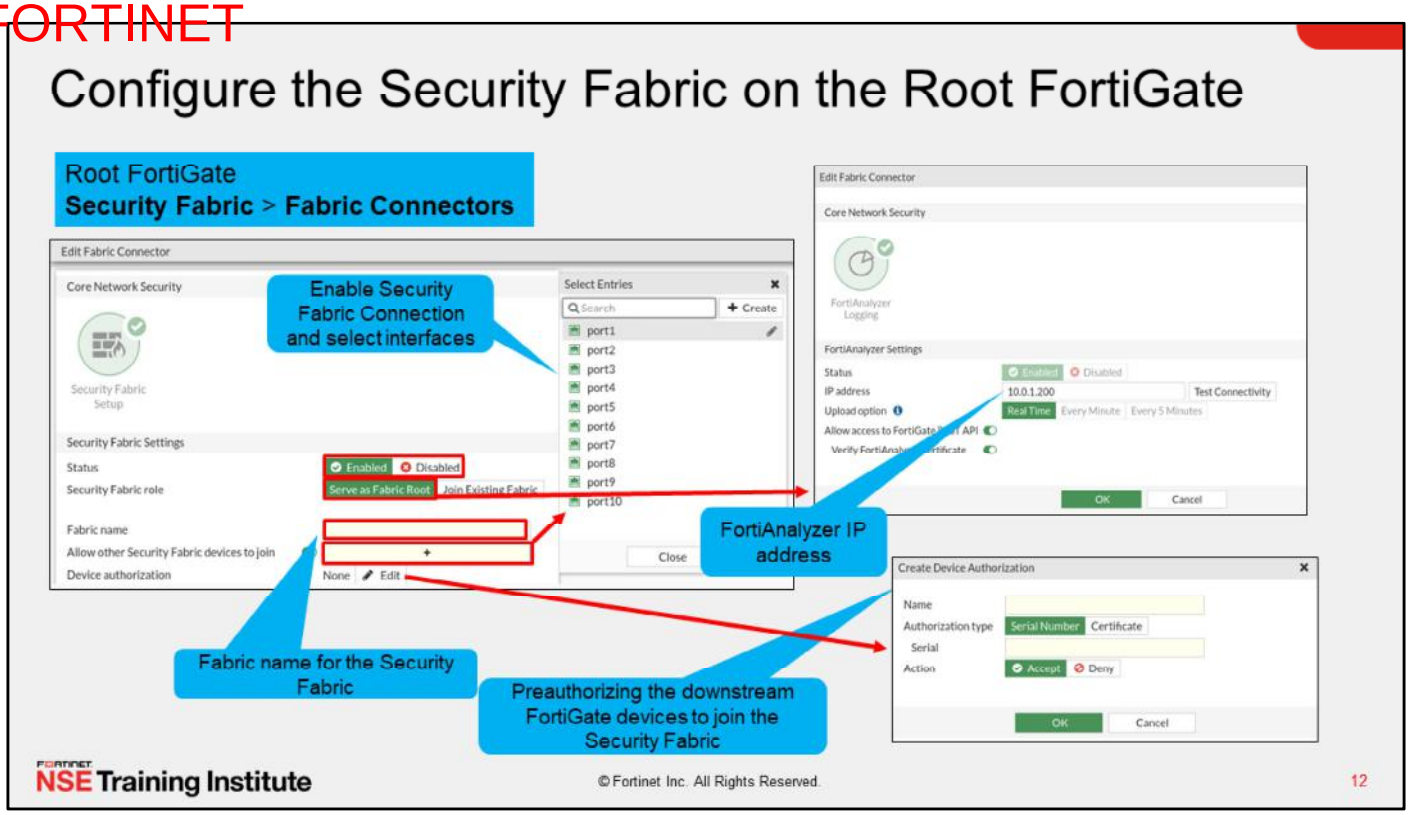

First, on the root FortiGate, you must enable **Security Fabric Connection** in the interfaces facing any downstream FortiGate. If you select **Serve as Fabric Root**, you also need to configure the FortiAnalyzer IP address. Then, you need to configure a fabric name for the Security Fabric. This FortiAnalyzer configuration will be pushed to all the downstream FortiGate devices. All downstream FortiGate devices send logs directly to FortiAnalyzer.

You can also preauthorize your downstream devices by adding the serial number of the device. 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.

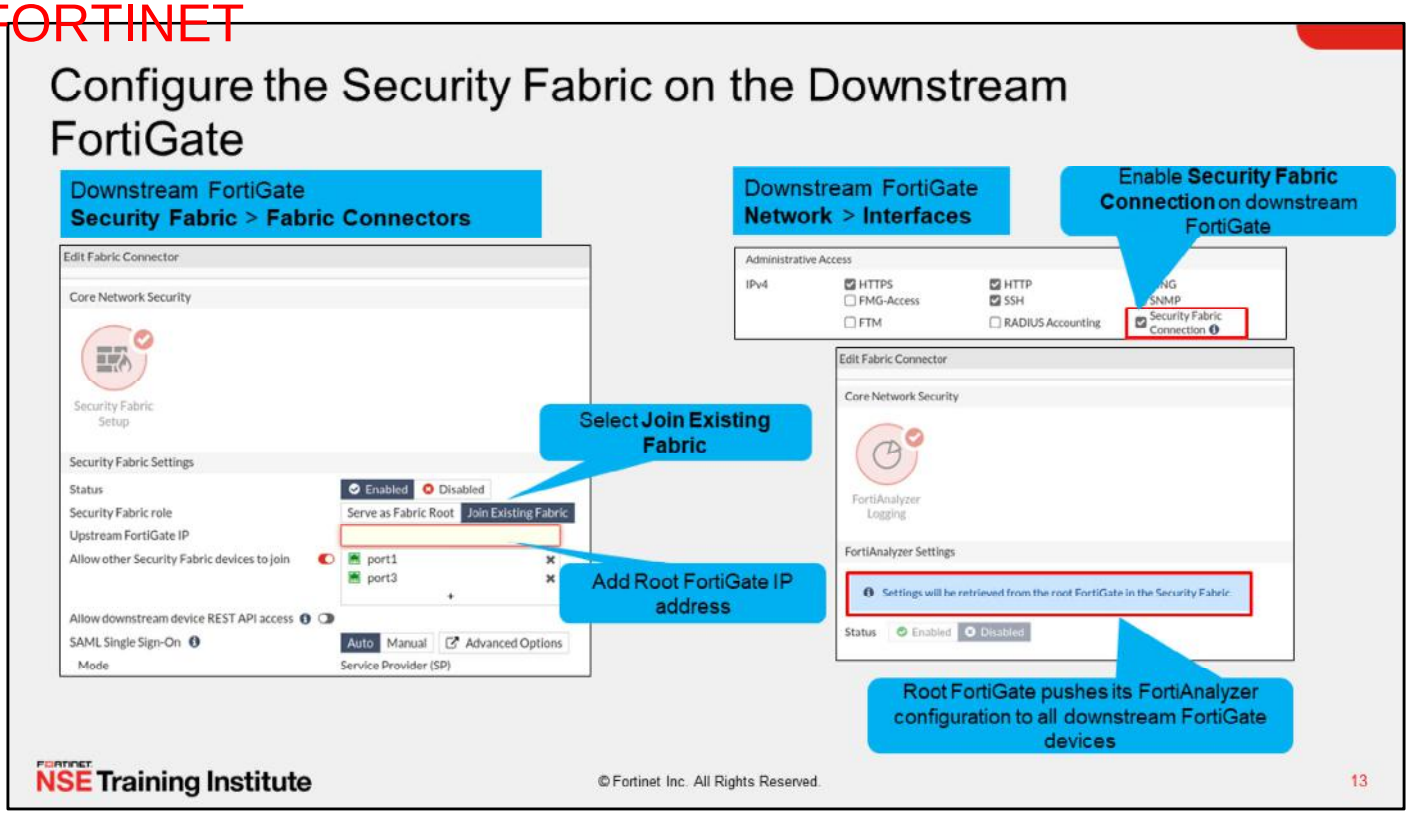

The second step in implementing the Security Fabric is configuring the downstream Fortinet devices. On the downstream FortiGate devices, you must enable **Security Fabric Connection** and **Device Detection** on the interfaces facing the downstream FortiGate devices. On the **Fabric Connectors** page, select **Join Existing Fabric** and add the root (upstream) FortiGate IP address. The root FortiGate pushes its FortiAnalyzer configuration to all downstream FortiGate devices.

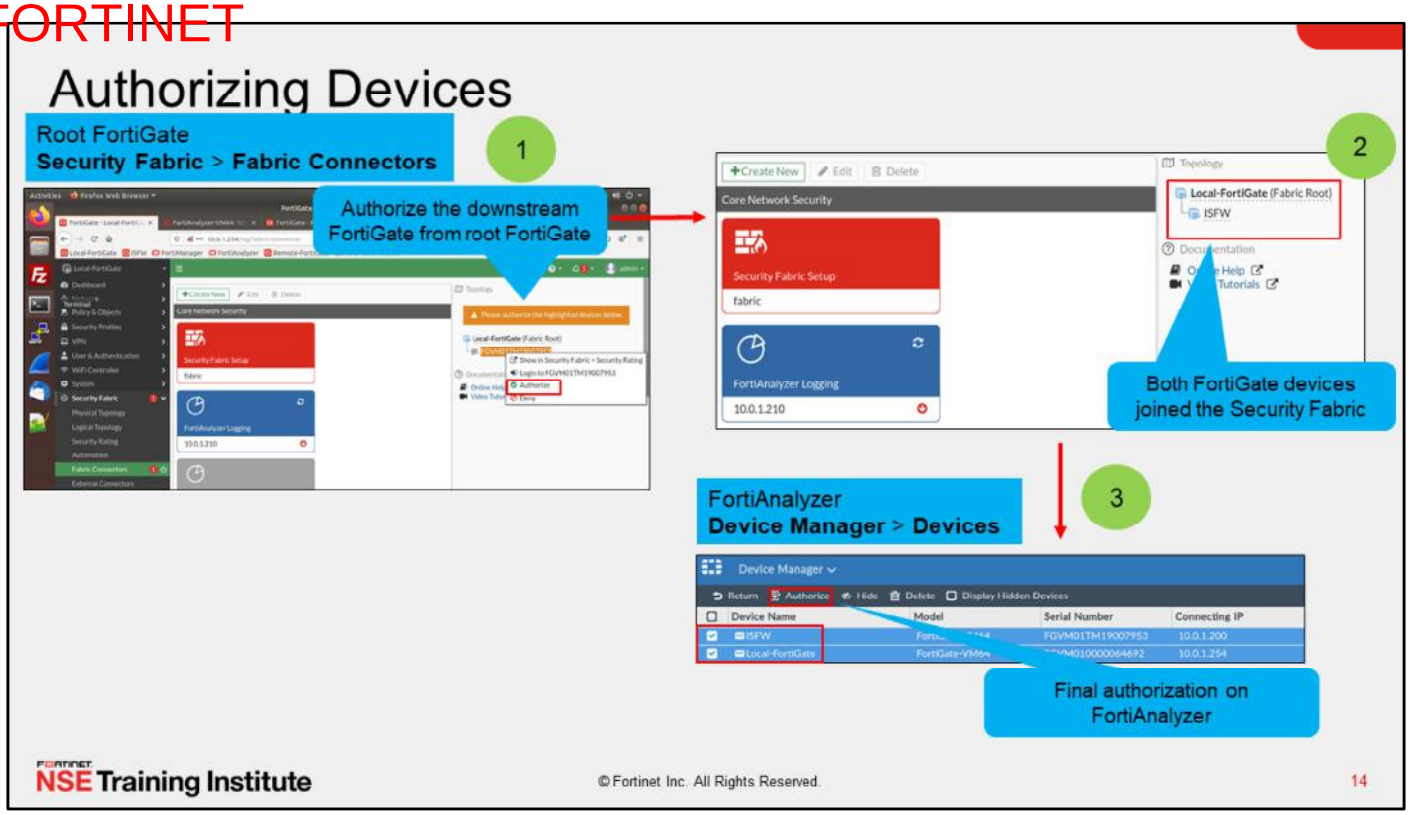

The third step in implementing the Security Fabric is to authorize the downstream FortiGate device on the both root FortiGate and the FortiAnalyzer. Click the serial number of the highlighted downstream FortiGate device and select **Authorize**. After few seconds, the downstream FortiGate will join the Security Fabric. In order to complete the full Security Fabric process, you will need to authorize all your devices on the FortiAnalyzer. From the FortiAnalyzer **Device Manager** section, select all your devices in the Security Fabric and click **Authorize**. After few seconds, you will notice all your authorized devices join the Security Fabric.

Synchronizing Objects Across the Security Fabric

- By default, object synchronization is enabled in fabric settings

```
config system csf
set status enable
set configuration-sync default
set fabric-object-unification default
end
```

- If `set fabric-object-unification` is set to `local` on the root FortiGate device, global fabric objects are not synchronized to downstream FortiGate devices

```
config system csf
 set status enable
 set group-name "fortinet"
 set fabric-object-unification local
```

- If `set configuration-sync` is set to `local`, the downstream device does not participate in synchronization

```
config system csf
set status enable
set configuration-sync local
end
```

- Select per object option to synchronize or not on the root FortiGate
  - By default, this option is disabled, and fabric objects are kept as locally created objects on FortiGate
  - If disabled on the root FortiGate, objects will not be synchronized to downstream FortiGate devices

NSE Training Institute

15

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 there is an object conflict during synchronization, you'll get a notification to resolve the conflict. In the topology tree, **Remote-FortiGate** is highlighted in amber because there is a conflict.

In the example shown on this slide, you will examine how to resolve a syncing conflict.
1. The notification icon displays this message: **Firewall objects are in conflict with other FortiGates in the fabric**. 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 has changed to **Resolved**.
4. In the topology tree, `none of the FortiGate devices are highlighted.`

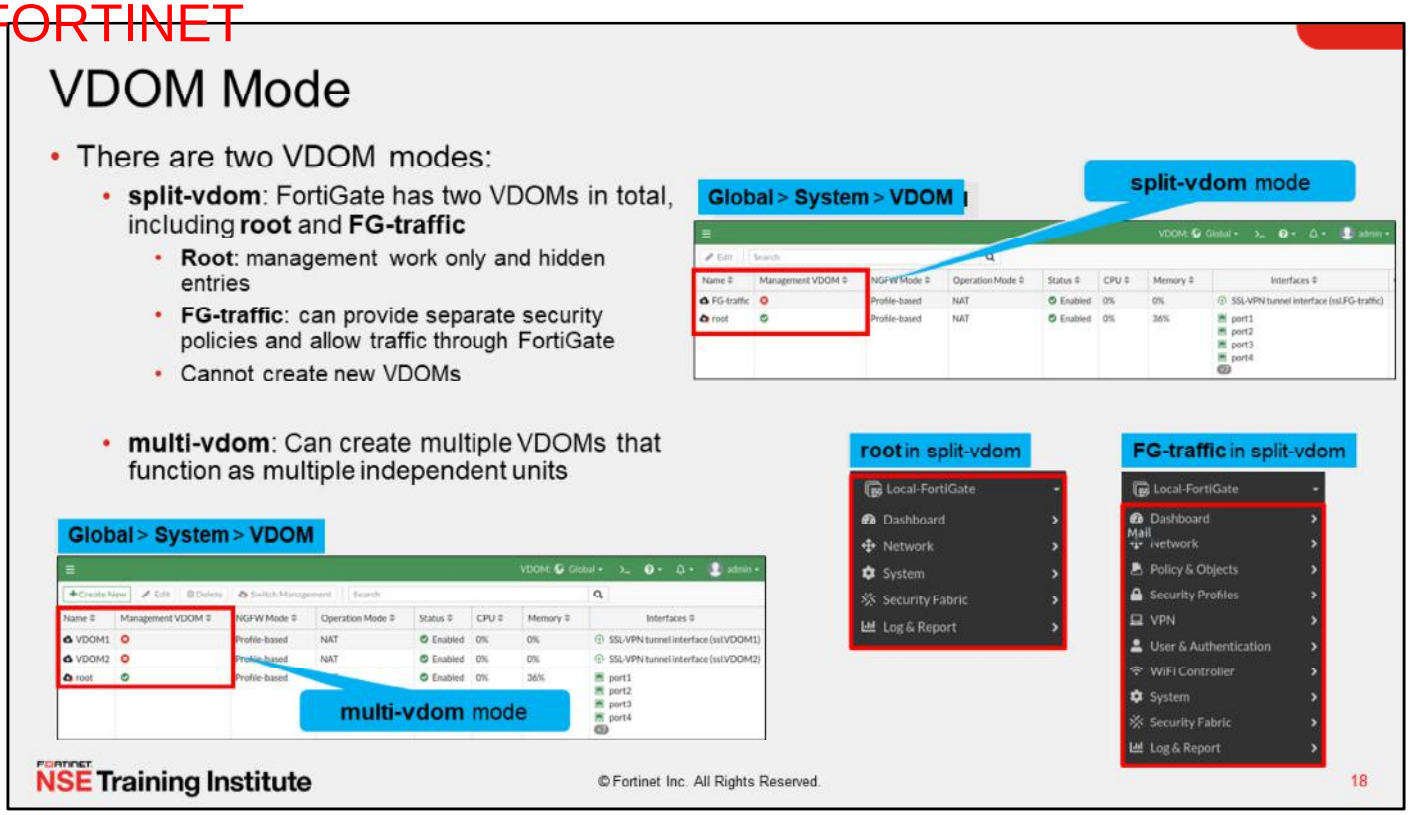

There are two VDOM modes: split-vdom and multi-vdom. In split-vdom mode, FortiGate has two VDOMs in total, including **root** and **FG-traffic** vdoms. You cannot add VDOMs in split-vdom mode.
1. **split-vdom** mode:

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

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

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

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

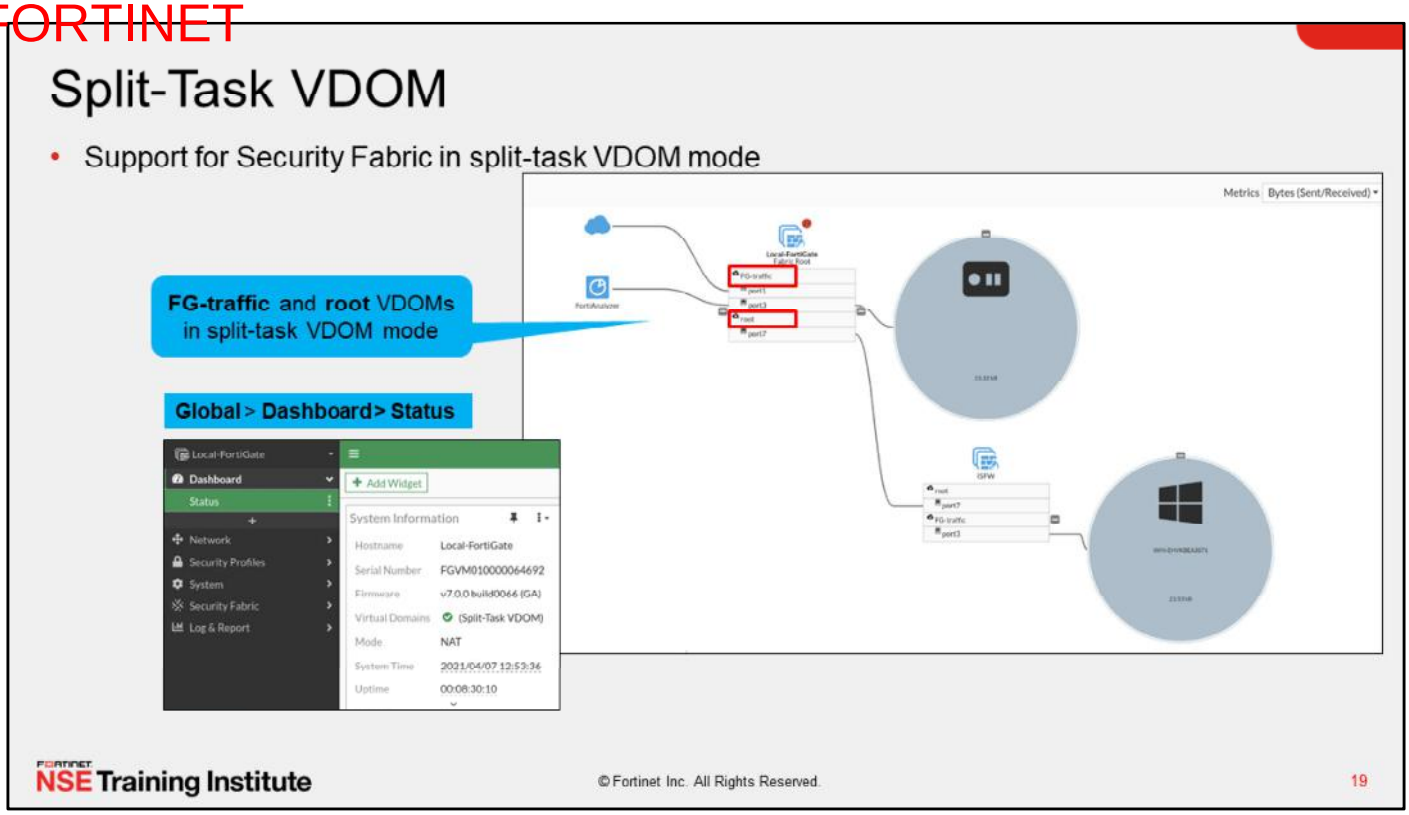

You can enable FortiGate Security Fabric in split-task VDOM mode. If you enable split-task VDOM mode on the upstream FortiGate device, it can allow downstream FortiGate devices to join the Security Fabric in the root and FG-traffic VDOMs. If split-task VDOM mode is enabled on the downstream FortiGate, it can connect to the upstream FortiGate only through the downstream FortiGate interface on the *root* VDOM.

Telemetry settings are shown in both global and VDOM contexts, but in the VDOM context, only the topology and FortiTelemetry-enabled interface fields are shown.

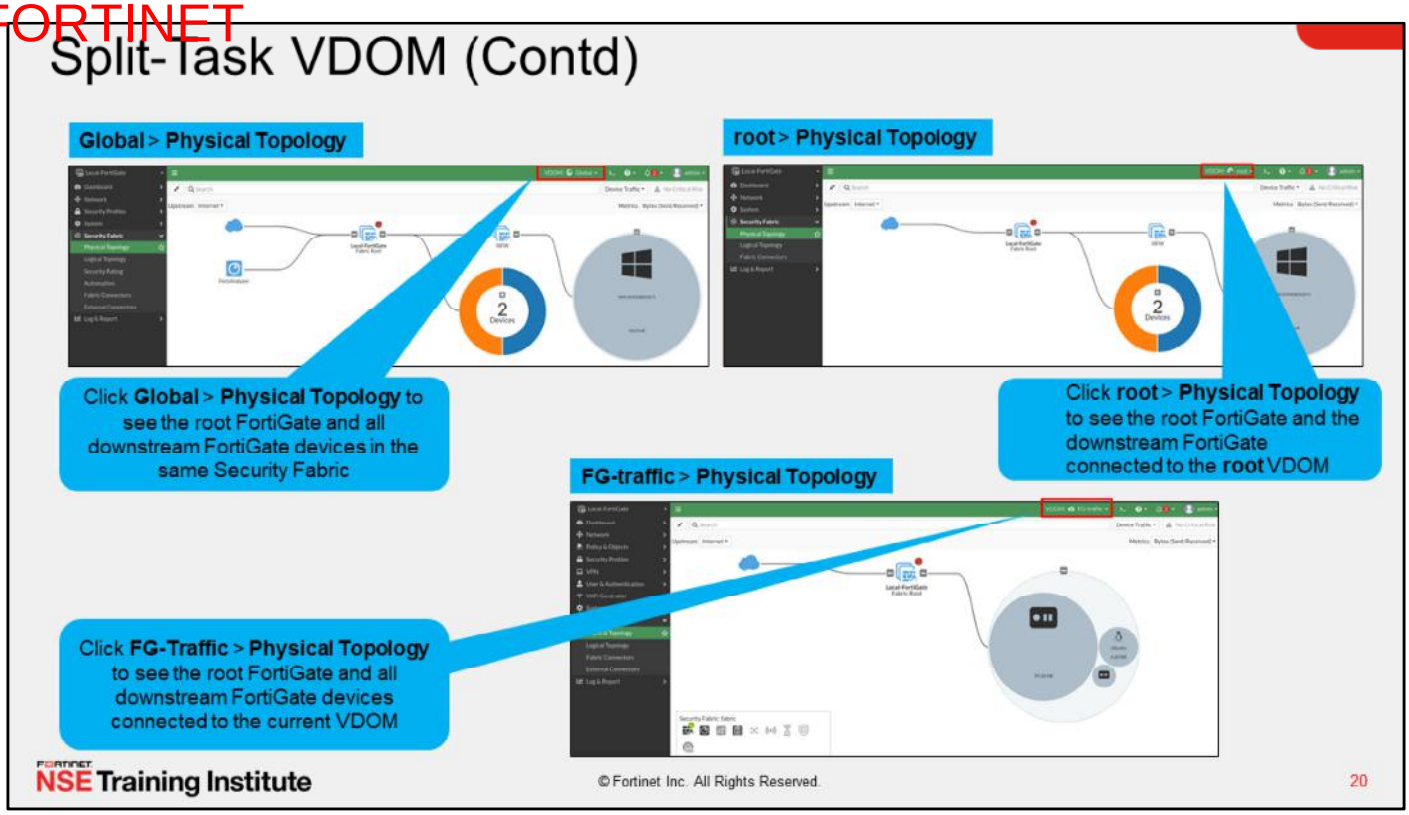

You can click **Global** > **Physical Topology** to see the root FortiGate and *all* downstream FortiGate devices that are in the same Security Fabric as the root FortiGate. You can click **root** > **Physical Topology** or **FG-Traffic** > **Physical Topology** to see the root FortiGate and *only* the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

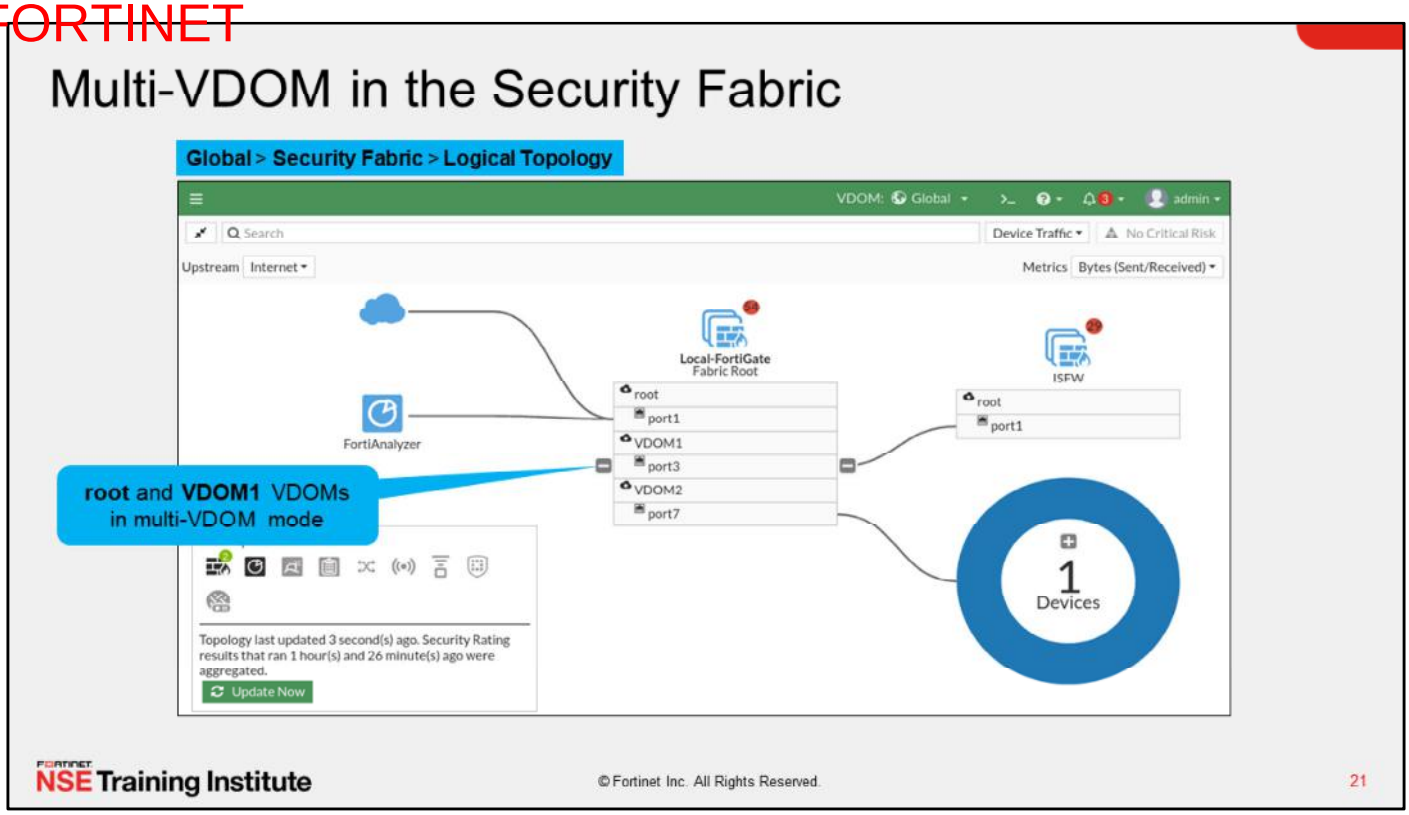

When you configure FortiGate devices in multi-vdom mode and add them to the Security Fabric, each VDOM with its assigned ports is displayed when one or more devices are detected. *Only* the ports with discovered and connected devices appear in the Security Fabric view and, because of this, you must enable **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

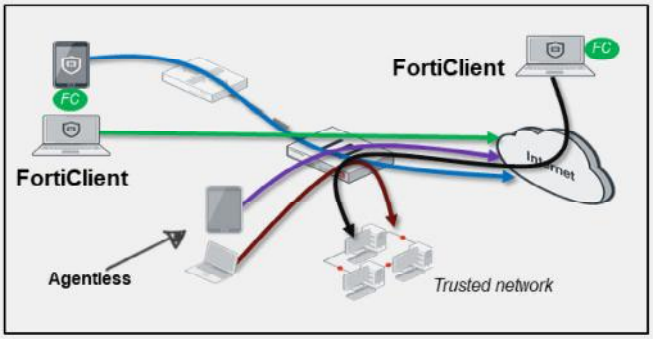

NSE Training Institute © Fortinet Inc. All Rights Reserved. 22

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).

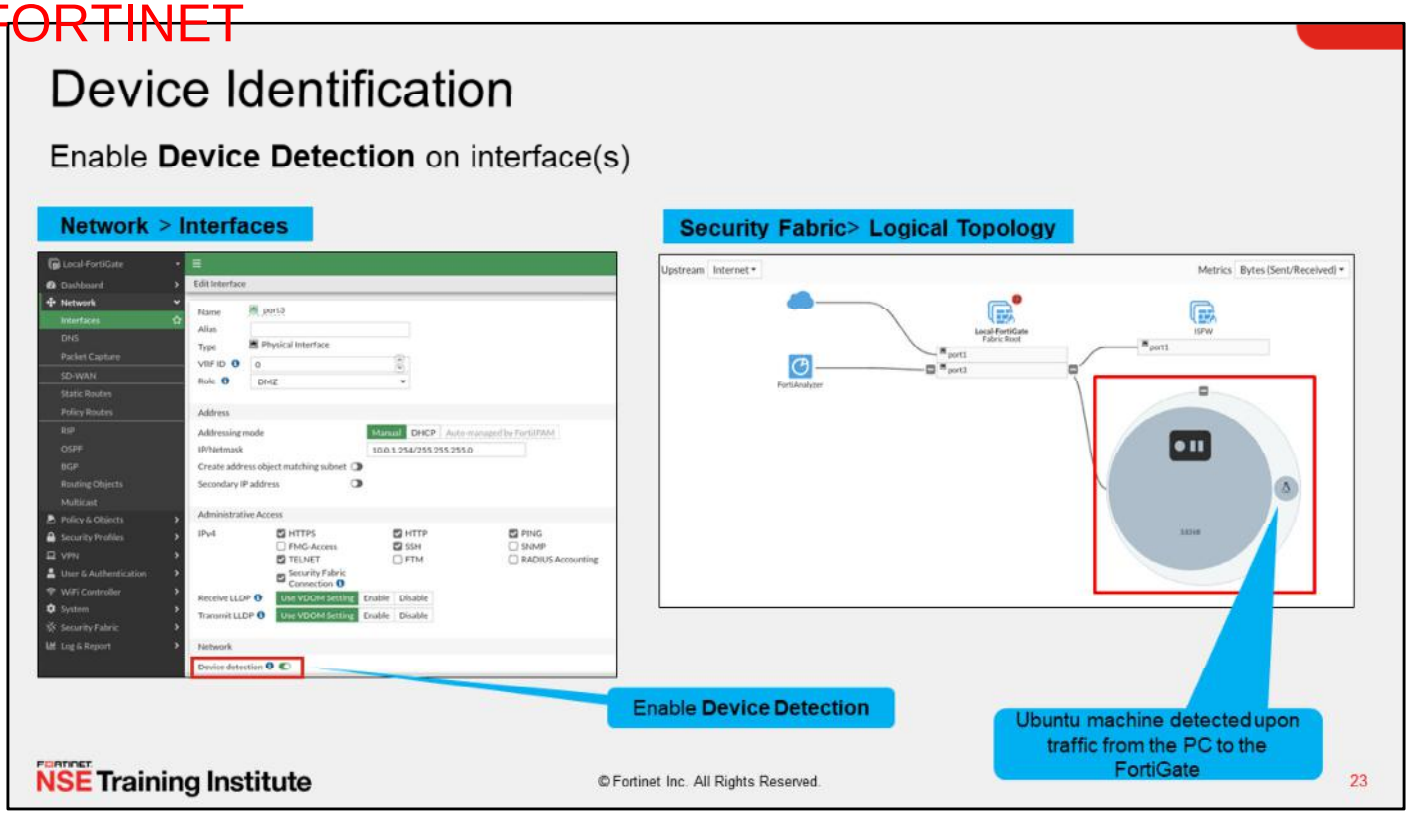

By default, FortiGate uses device detection (passive scanning), which runs scans based on the arrival of traffic.

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

NSE Training Institute

24

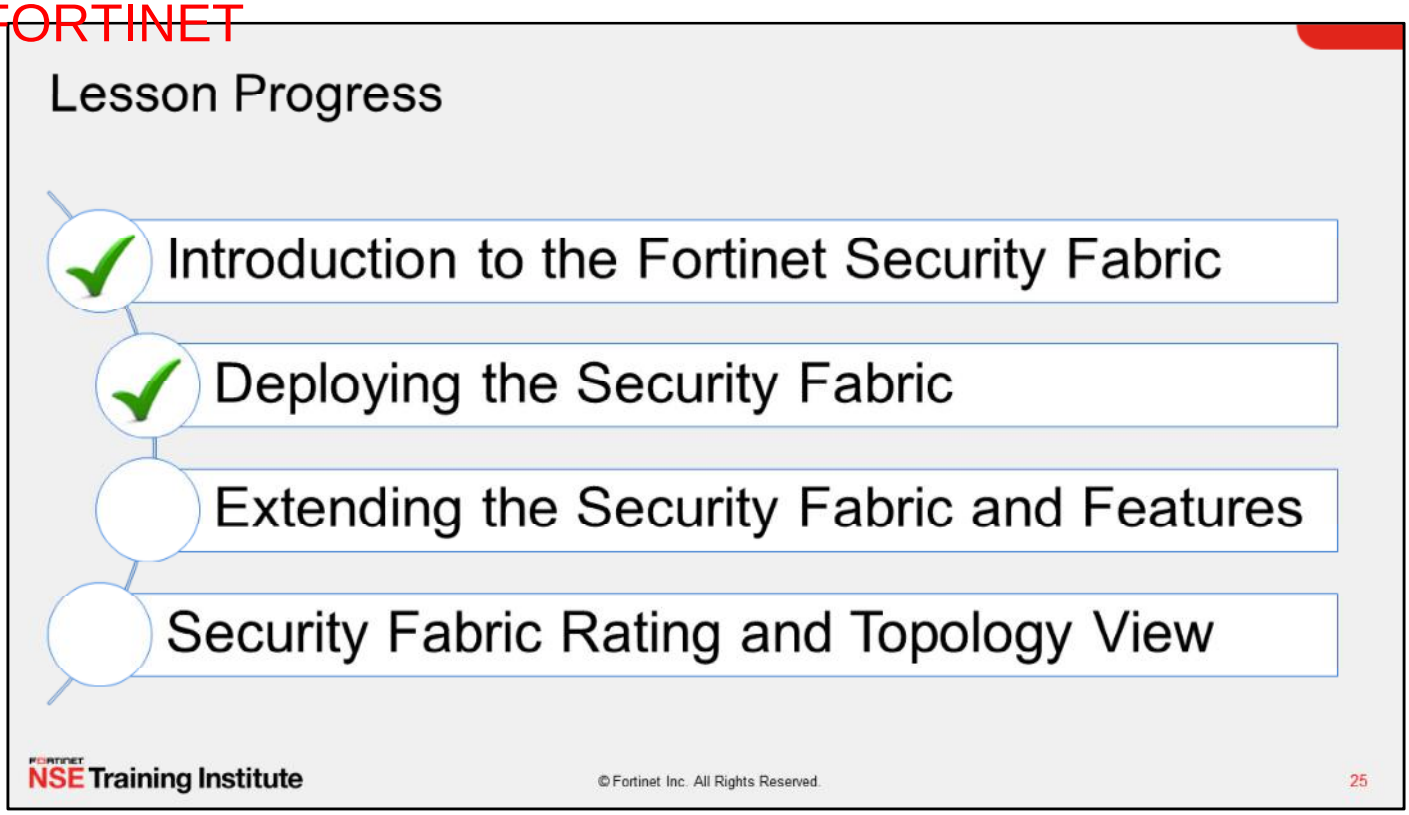

Good job! You now know how to deploy the Security Fabric.

Next, you'll learn about Security Fabric features and how to extend the Security Fabric in your network environment.

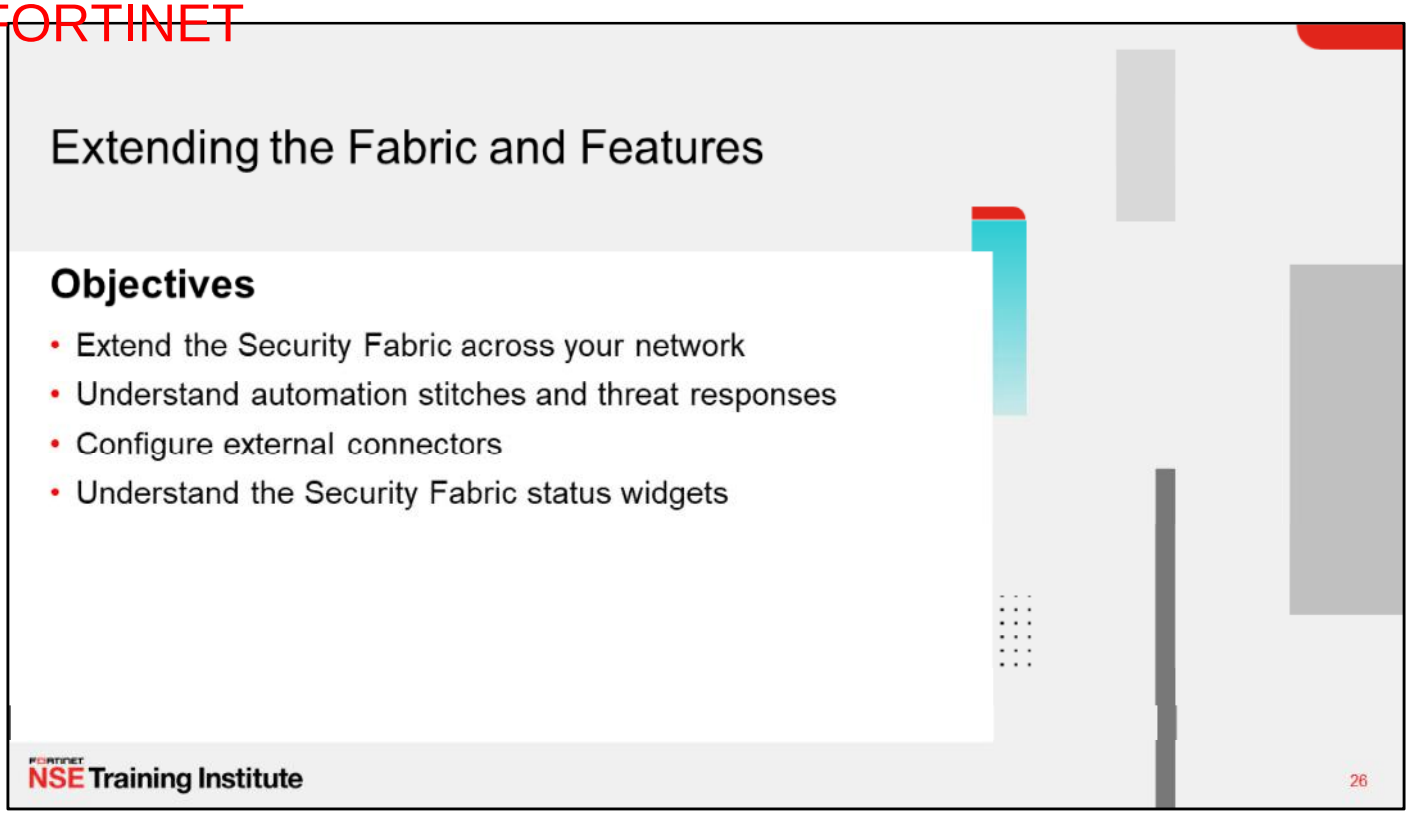

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

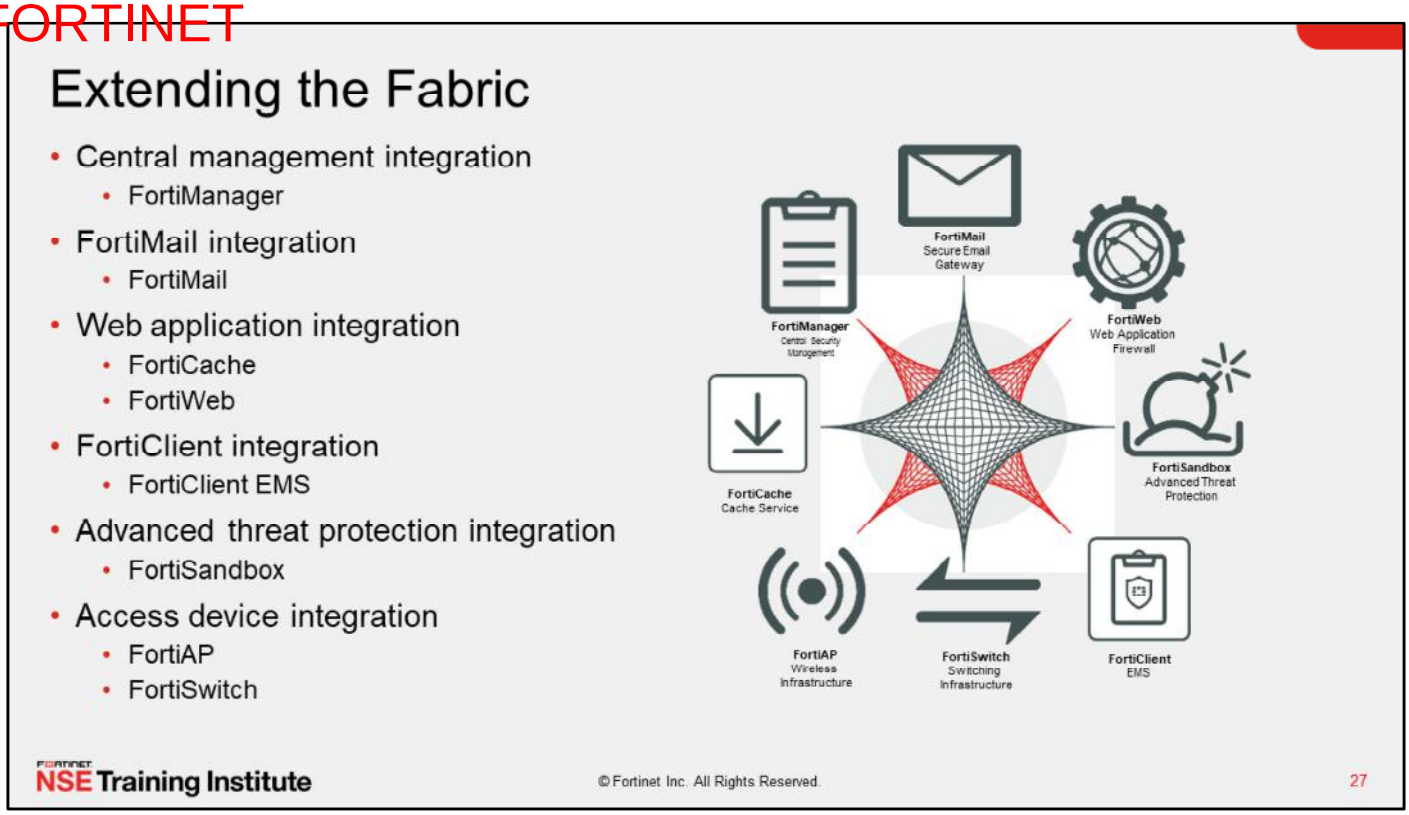

Fortinet recommends using a FortiManager for centralized management of all FortiGate devices and access devices in the Security Fabric. You can integrate FortiSwitch, and FortiAP devices to extend the Security Fabric down to the access layer. You can also extend the Security Fabric by integrating FortiMail, FortiWeb, FortiCache, FortiSandbox, and FortiClient EMS.

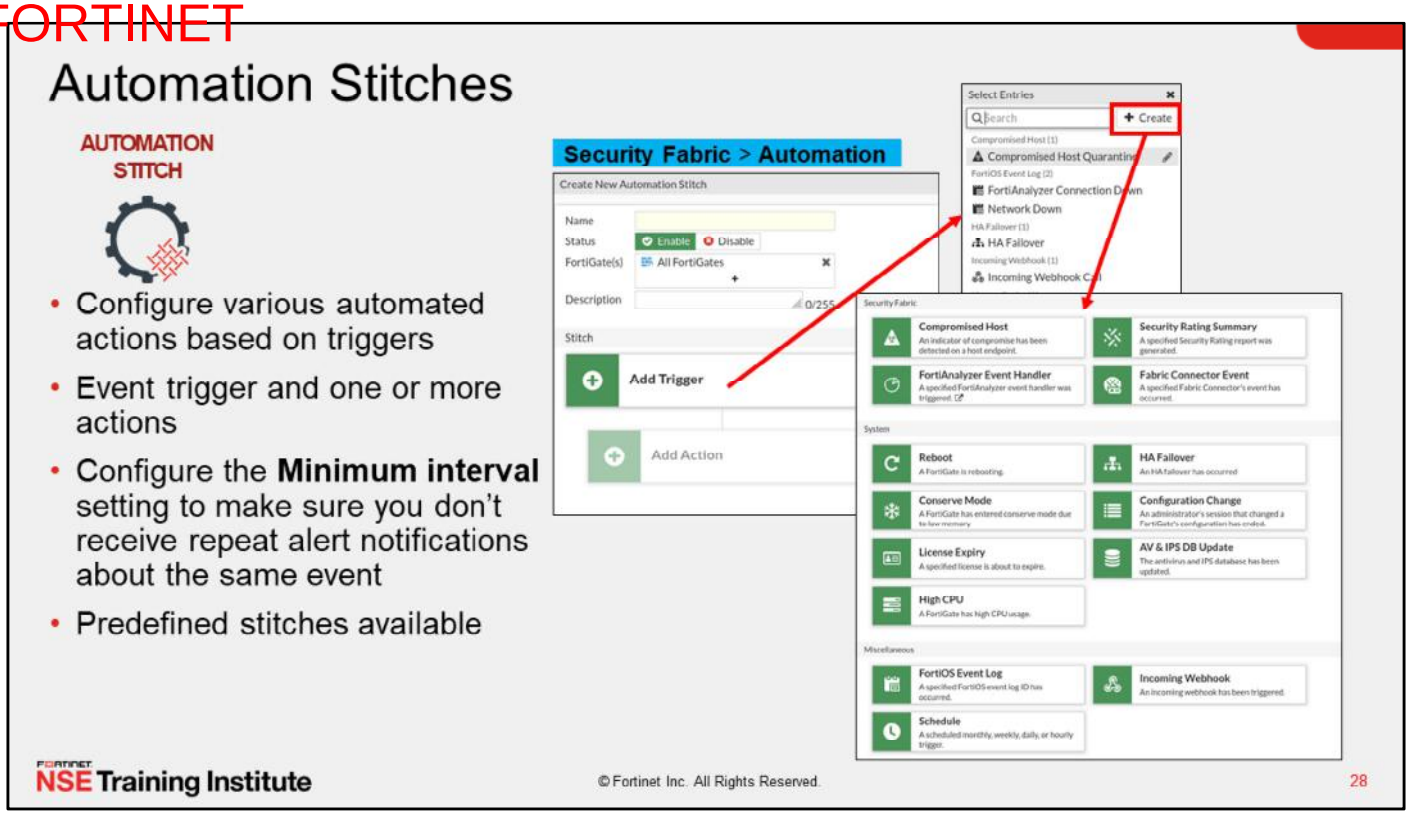

Administrator-defined automated work flows (called stitches) use if/then statements to cause the FortiOS to automatically respond to an event in a preprogrammed way. Because this workflow is part of the Security Fabric, you can set up if/then statements for any device in the Security Fabric. However, the Security Fabric is not required to use stiches.

Each automation stitch pairs an event trigger and one or more actions. 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 to make sure you don't receive repeat notifications about the same event. There are predefined stitches, triggers and actions available. However, you can create custom automation based on the available options.

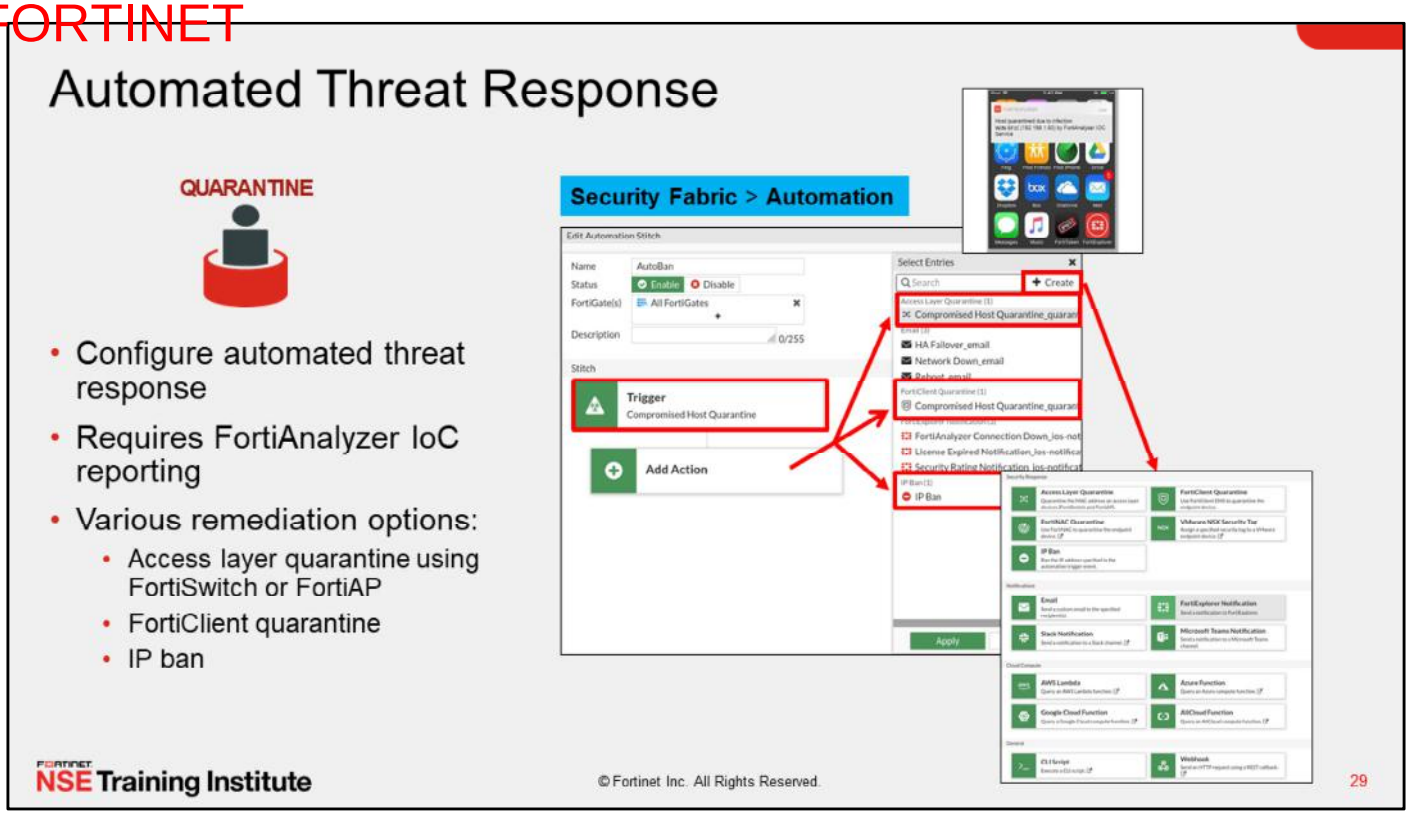

You can configure the **Compromised Host** trigger to create an automated threat response stitch. This trigger uses indicator of compromise (IoC) event reporting from FortiAnalyzer. Based on the **Threat level threshold** setting, you can configure the stitch to take different remediation steps:

- Quarantine the compromised host at the FortiSwitch or FortiAP
- Quarantine FortiClient on the compromised host using FortiClient EMS
- Ban the IP

You can also click **Monitor** > **Quarantine Monitor** to view quarantined and banned IP addresses. Quarantined addresses are automatically removed from quarantine after a configurable period of time. Banned IP addresses can be removed from the list only by administrator intervention.

Automated Threat Response (Contd)

NOTIFICATIONS

- Output notifications in various ways such as iOS Push or on the GUI dashboard
- Integrate with IFTTT and other cloud services

Security Fabric > Physical Topology

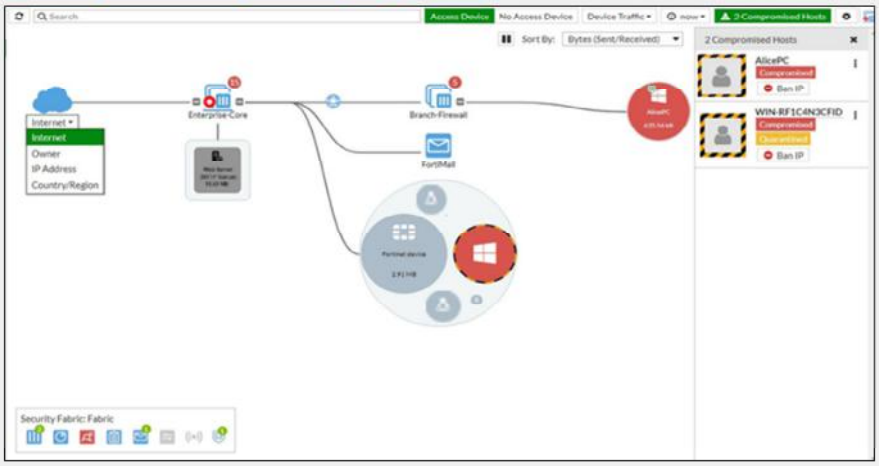

NSE Training Institute

30

You can also view compromised hosts on the FortiGate GUI and get output notifications in various ways such as iOS push. This feature is integrated with IFTT.

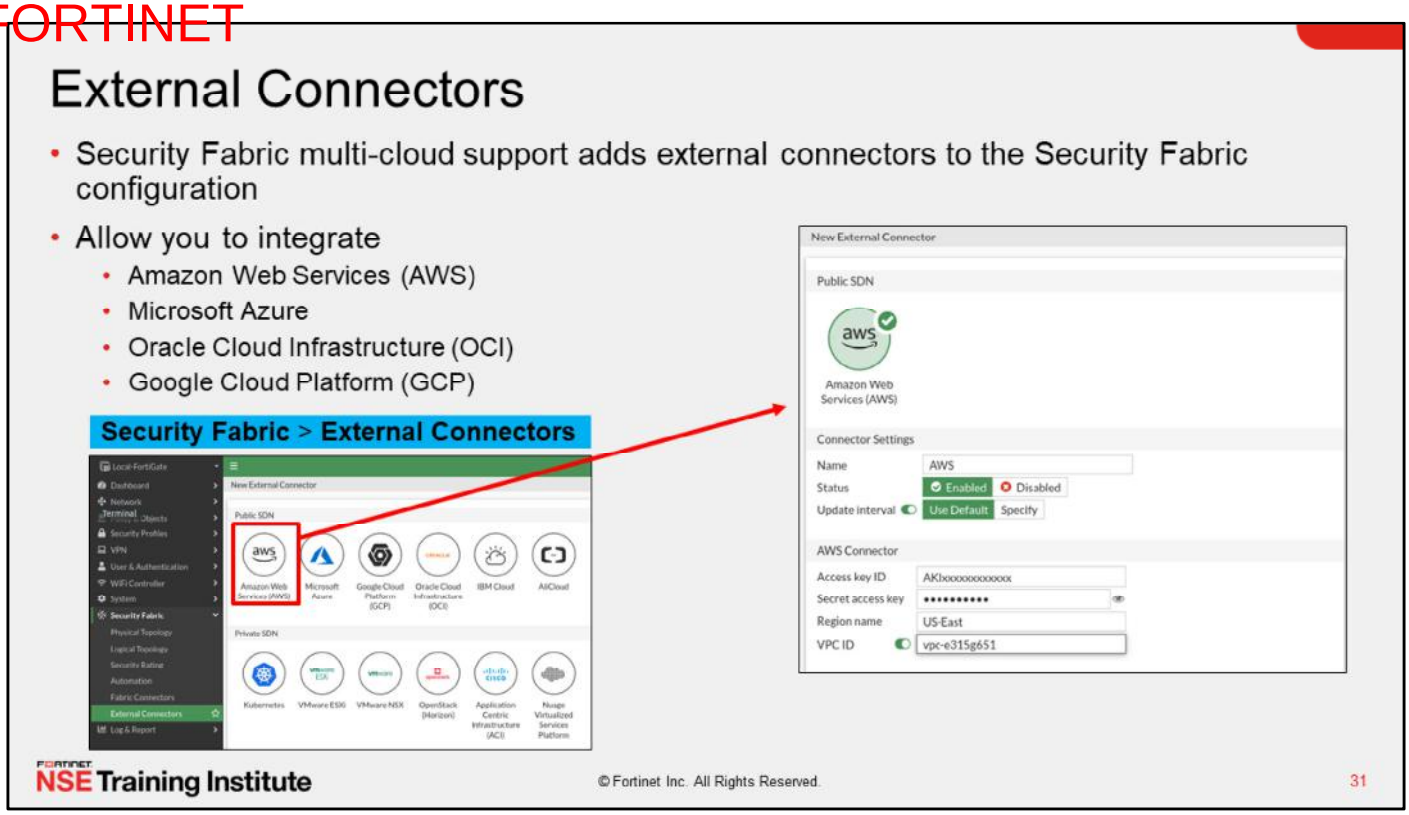

External connectors allow you to integrate multi-cloud support, such as ACI and AWS, to name a few.

In an application-centric infrastructure (ACI), the SDN connector serves as a gateway bridging SDN controllers and FortiGate devices. The SDN Connector registers itself to APIC in the Cisco ACI fabric, polls interested objects, and translates them into address objects. The translated address objects and associated endpoints populate on FortiGate.

FortiGate VM for Microsoft Azure also supports cloud-init and bootstrapping.

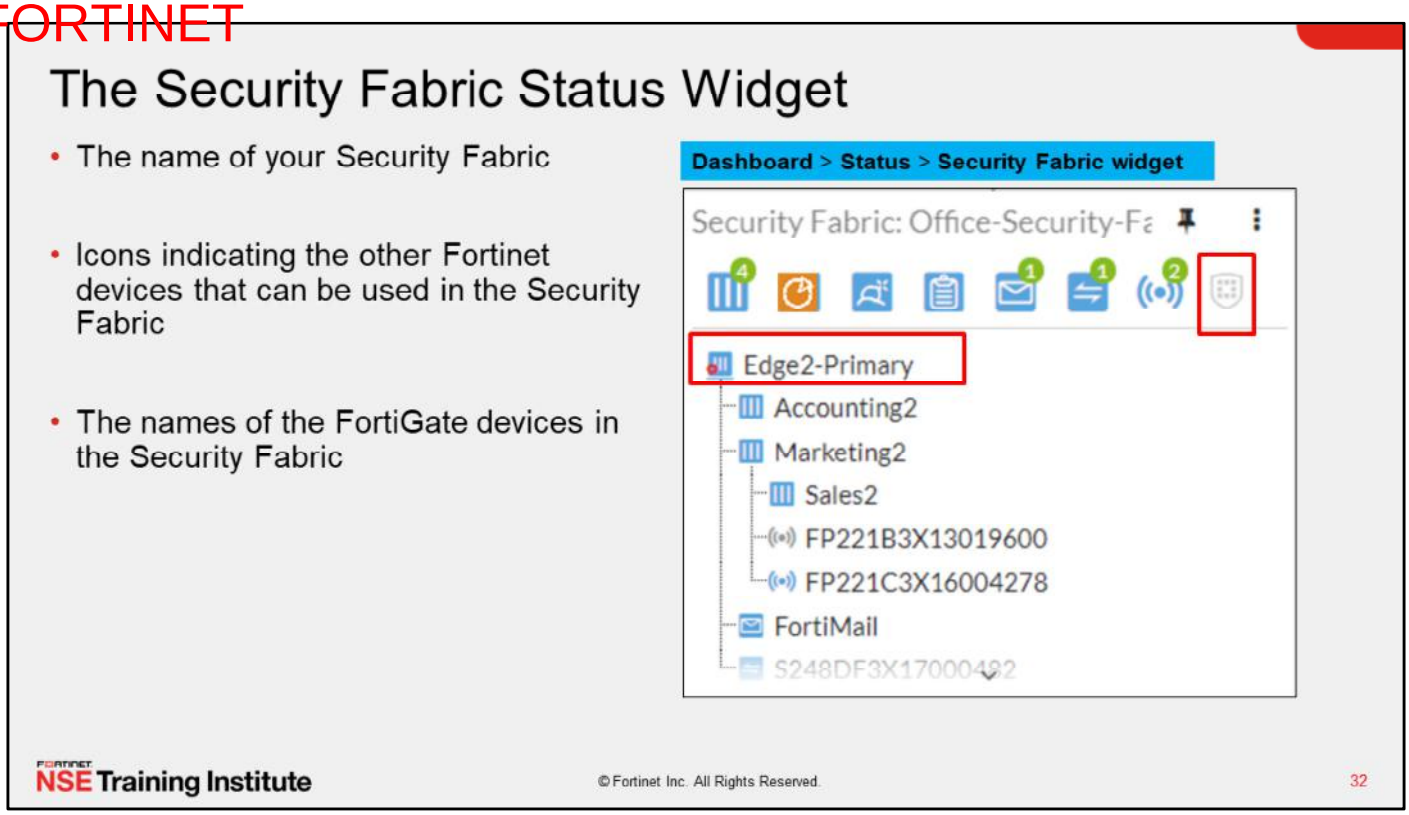

The **Security Fabric Status** widget shows a visual summary of many of the devices in the Security Fabric. You can hover over the icons at the top of the widget to get a quick view of the status of the Security Fabric, including the status of FortiTelemetry and devices in the Security Fabric. You can 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 unauthorized devices that are connected in your network.
- Devices in red are not detected in your network, but are recommended for the Security Fabric.
- An attention icon indicates a FortiGate or FortiWiFi waiting for authorization.

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)

NSE Training Institute

33

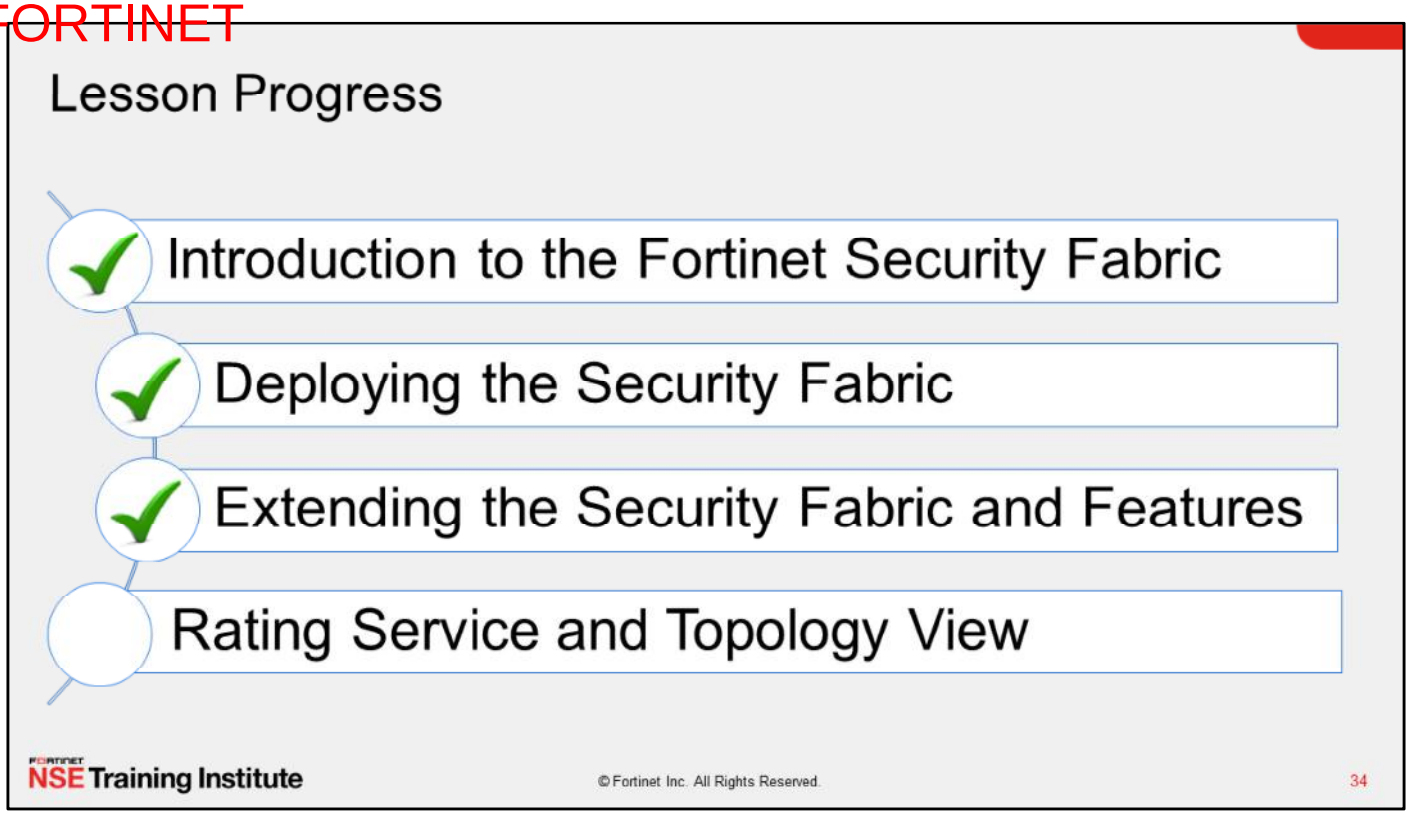

Good job! You now know how to extend the Security Fabric and its features.

Next, you'll learn about the Security Fabric Rating service and topology view.

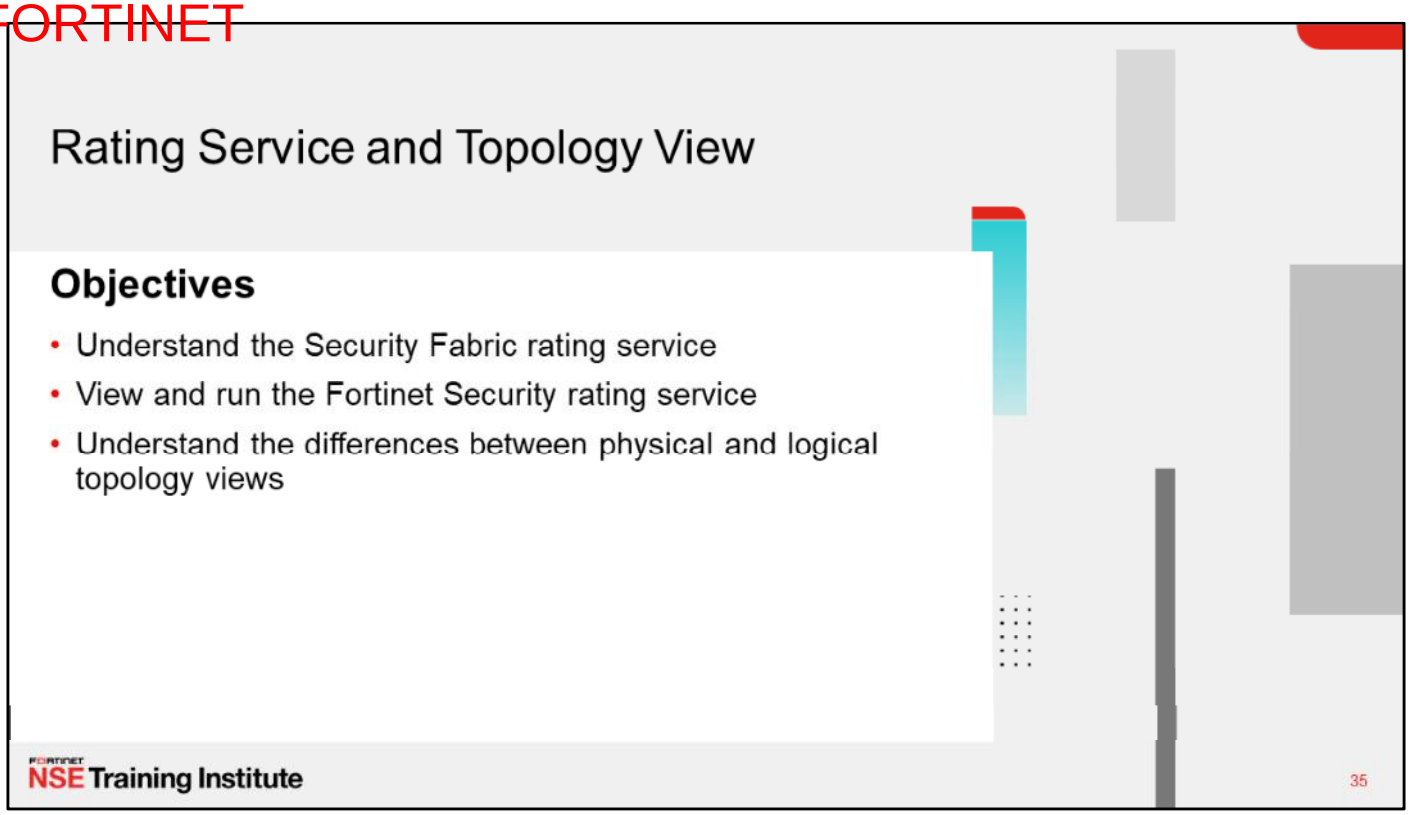

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in the Fortinet Security rating service and topology views, you should be able to have clear visibility of your network devices.

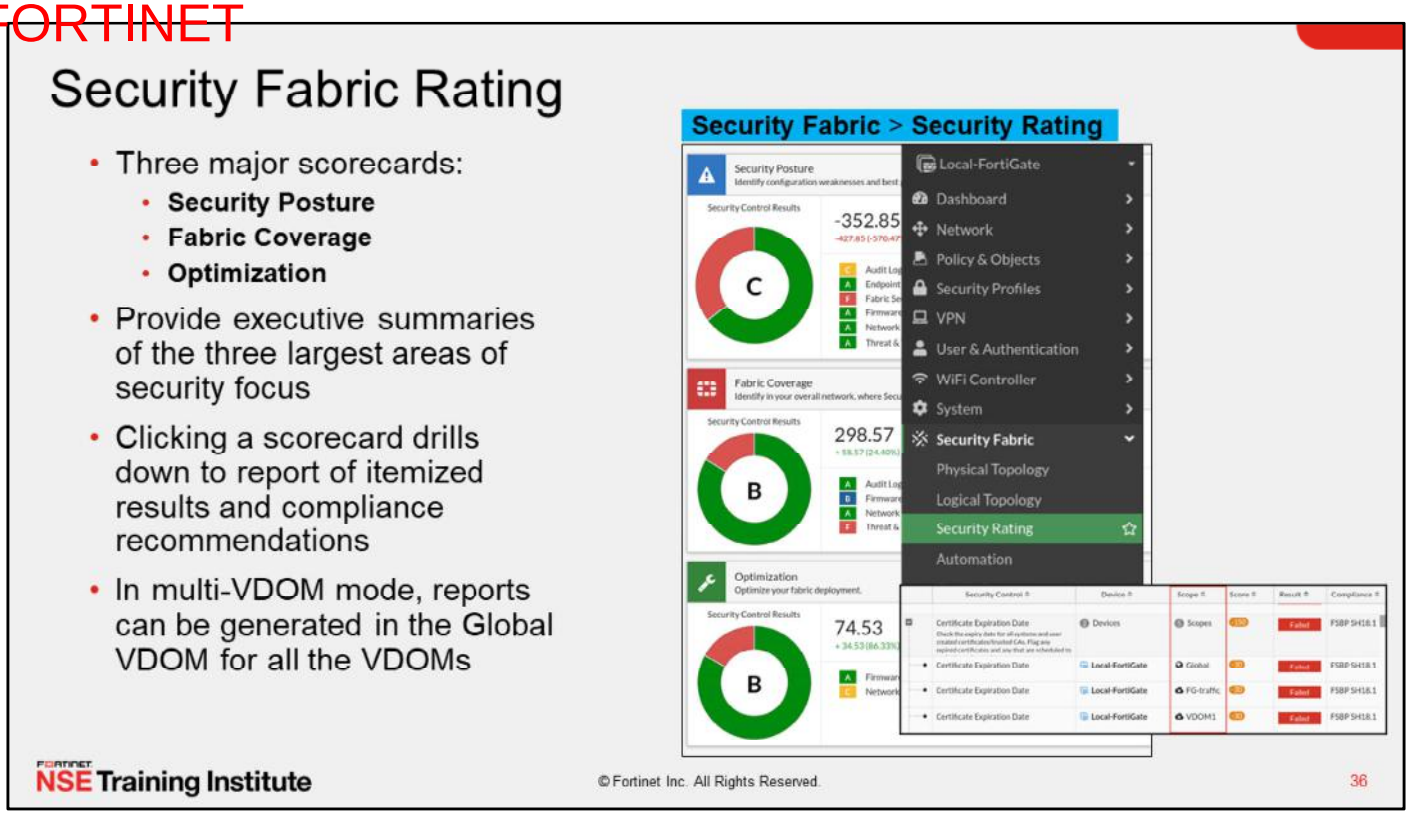

Security rating is a subscription service that requires a security rating license. This service now provides the ability to perform many *best practices*, including password checks, to audit and strengthen your network security.
The **Security Rating** page is separated into three major scorecards:

- **Security Posture**
- **Fabric Coverage**
- **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 sub-categories. Clicking a scorecard drills down to a detailed report of itemized results and compliance recommendations.
The point score represents the net score for all passed and failed items in that area. The report includes the security controls that were tested against, linking to specific FSBP or PCI compliance policies. You can click **FSBP** and **PCI** to reference the corresponding standard.

In multi-VDOM mode, security rating reports can be generated in the Global VDOM for all of the VDOMs on the device. Administrators with read/write access can run the security rating report in the Global VDOM. Administrators with read-only access can only view the report.

On the scorecards, the **Scope** column shows the VDOM or VDOMs that the check was run on. On checks that support **Easy Apply**, the remediation can be run on all of the associated VDOMs.

The security rating event log is available on the root VDOM.

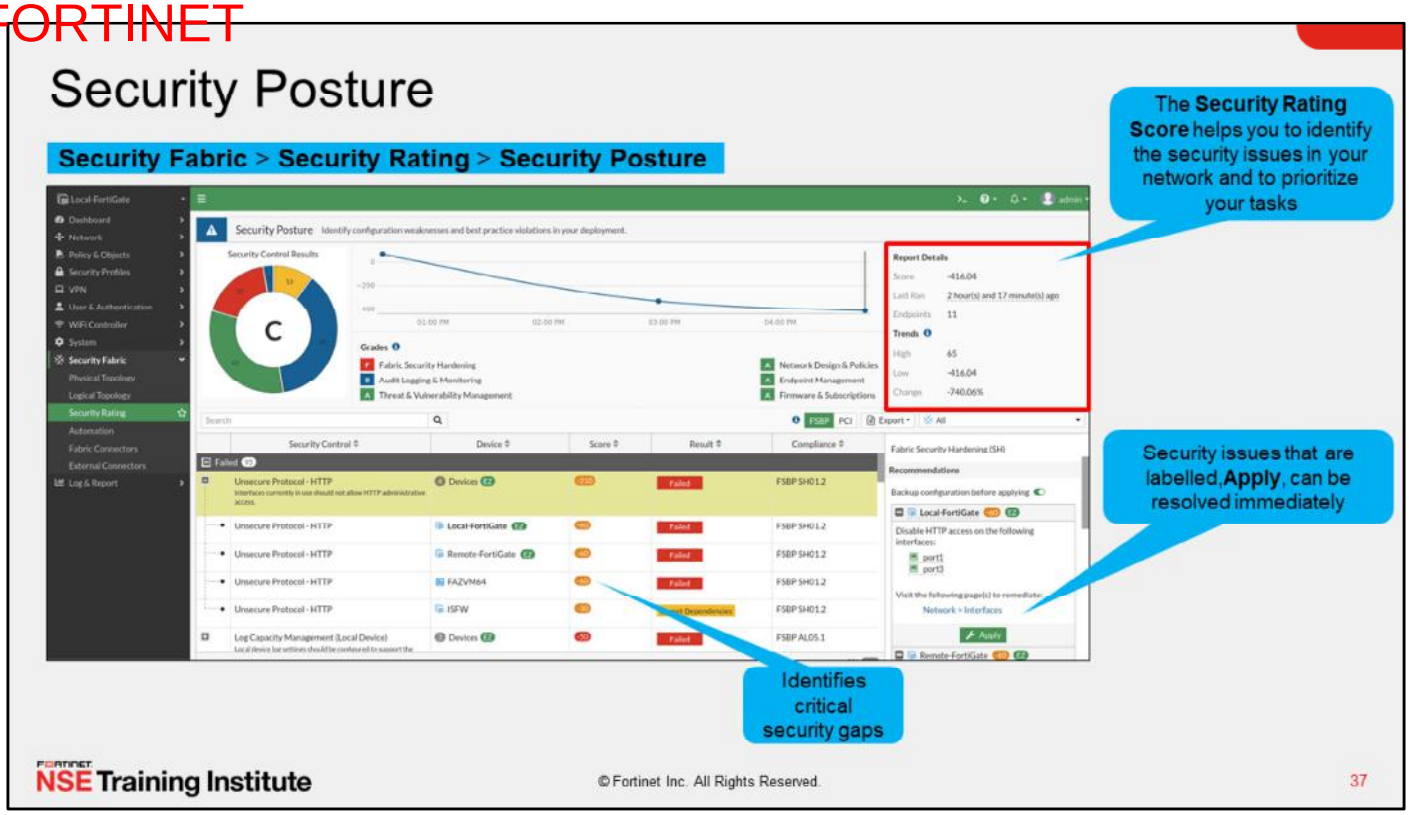

Click the **Security Posture** scorecard on the **Security Rating** page to expand the scorecard and see more details.

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.

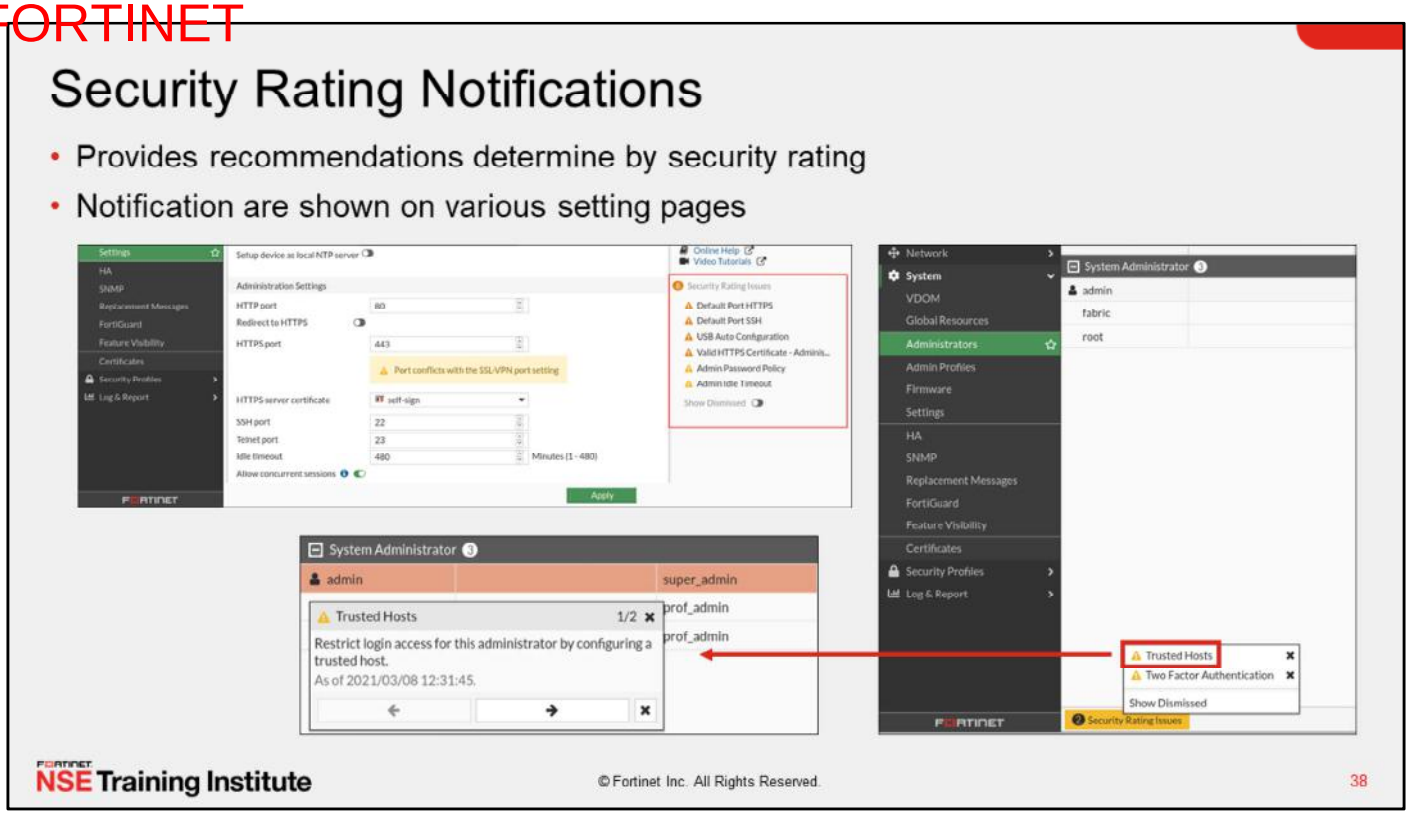

Security rating provides recommendation about FortiGate settings. These recommendations are shown as notifications on the settings page, which shows configuration issues as determined by security rating. An administrator can open the recommendation to see which configuration setting needs to be fixed. This helps administrator from going back and forth between the **Security Fabric** > **Security Rating** page and the various settings pages.

In the examples shown on this slide, FortiGate is using default HTTPS and SSH ports, and administrator password policy is not enabled. Another recommendation is to restrict login access by configuring a trusted host.

Notifications appear either in the gutter, the footer, or as a mutable. Notifications can also be dismissed.

FortiGuard Security Rating Service

- Enable or disable security checks using the CLI:

```
config system global
    set security-rating-run-on-schedule [enable/disable]
end
```

NSE Training Institute

39

The security rating service reports the security posture score per the Security Fabric group. FortiGuard Security Rating Service is a subscription-based service that takes the generated report and obtains analysis by FortiGuard. It compares security score results within the industry that the fabric group belongs to. All FortiGate devices in the group need to have FortiGuard Security Rating Service and the score can be obtained only on the Security Fabric root FortiGate. The score can be obtained after the security rating report is generated. The scores are presented as numbers and are based on the industry, the size of the organization, and the region.

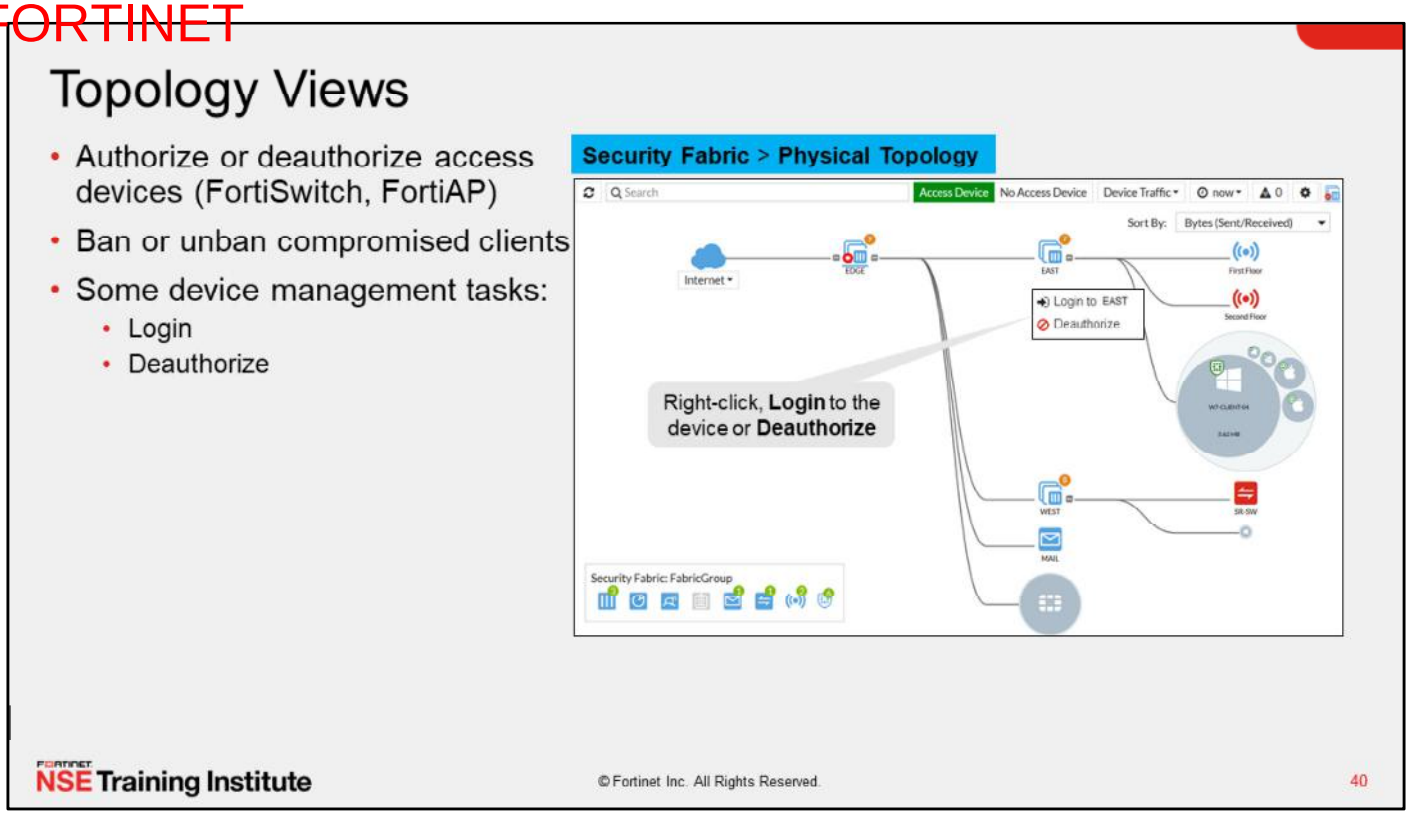

You can view the Security Fabric topology on the FortiGate GUI, from the **Security Fabric** menu. You can select the **Physical Topology** or **Logical Topology** view. To view the complete network, you must access the topology views on the root FortiGate in the Security Fabric.

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.

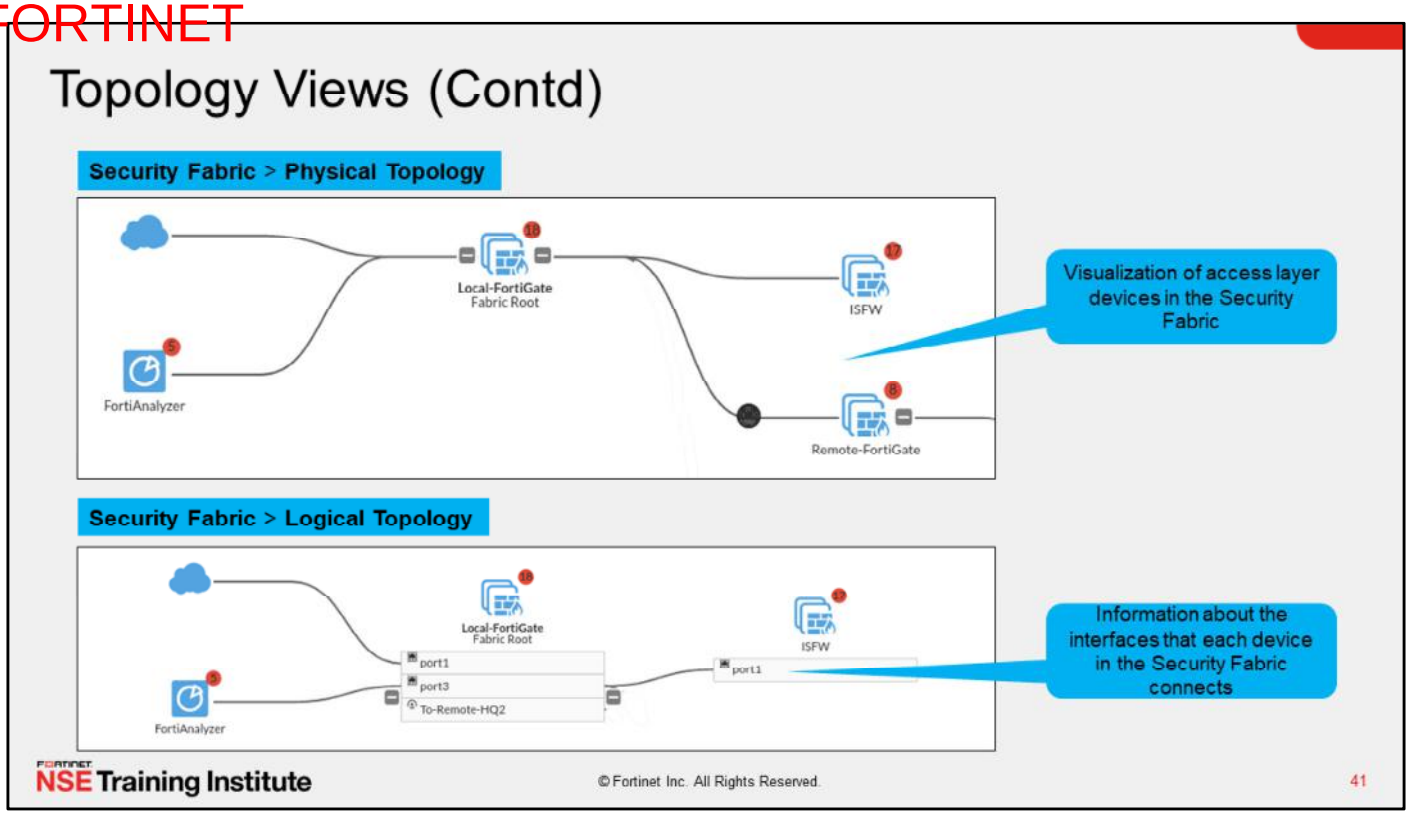

This slide shows the difference between the **Physical Topology** view and the **Logical Topology** view.

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

NSE Training Institute

42

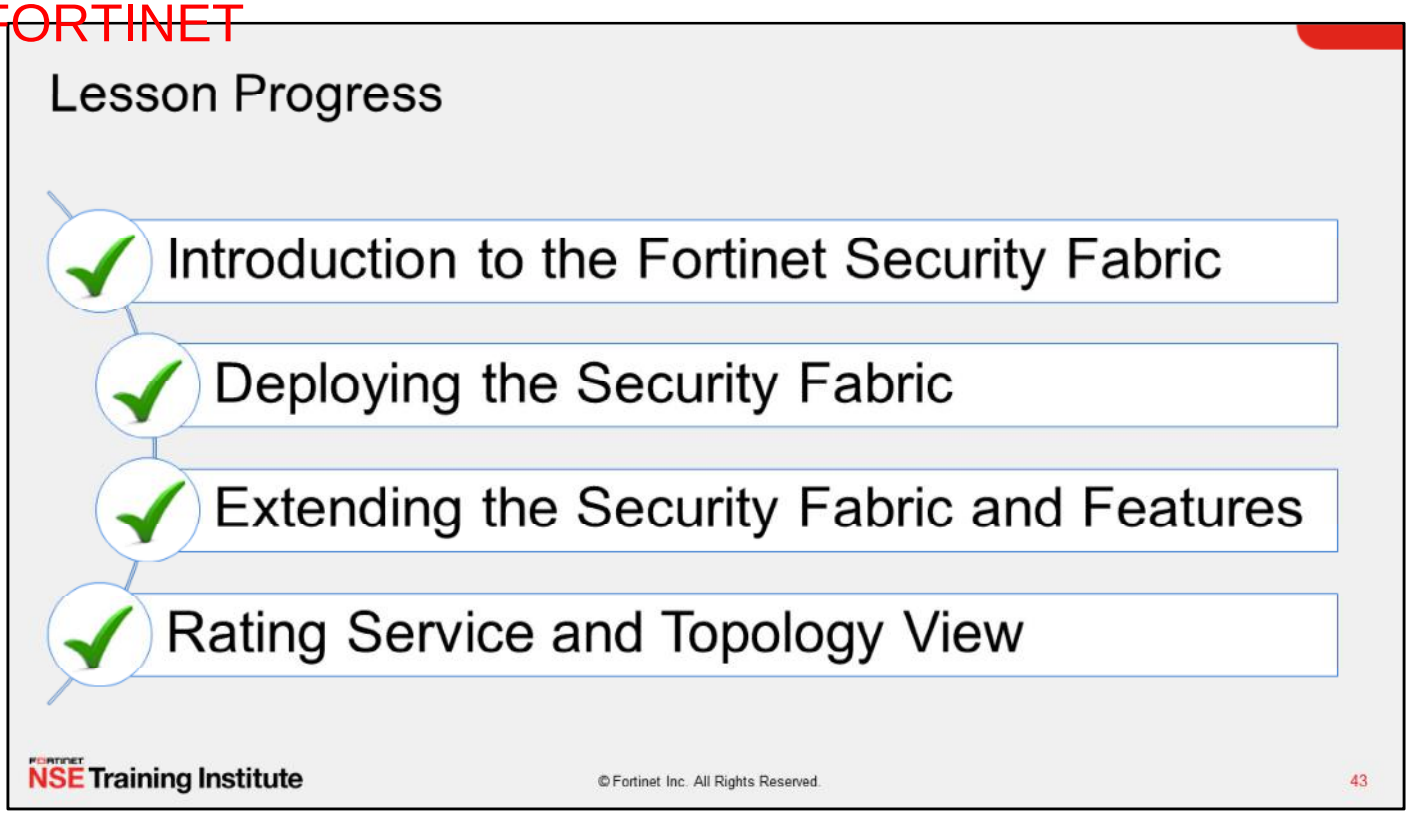

Congratulations! You have completed this lesson.

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

Review

- ✓ Define the Fortinet Security Fabric
- ✓ Identify why the Security Fabric is required
- ✓ Identify the Fortinet devices that participate in the fabric, especially the essential ones
- ✓ Understand how to implement the Security Fabric
- ✓ Configure the Security Fabric on the root and downstream FortiGate
- ✓ Understand how device detection works
- ✓ Understand how to extend your existing Security Fabric
- ✓ Extend the Security Fabric across your network
- ✓ Understand automation stiches and threat responses
- ✓ Configure fabric connectors
- ✓ Understand the Security Fabric status widgets
- ✓ Understand the Security Fabric Rating service
- ✓ View and run the Security Rating service
- ✓ Understand the differences between the physical and logical topology view

NSE Training Institute © Fortinet Inc. All Rights Reserved. 44

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.

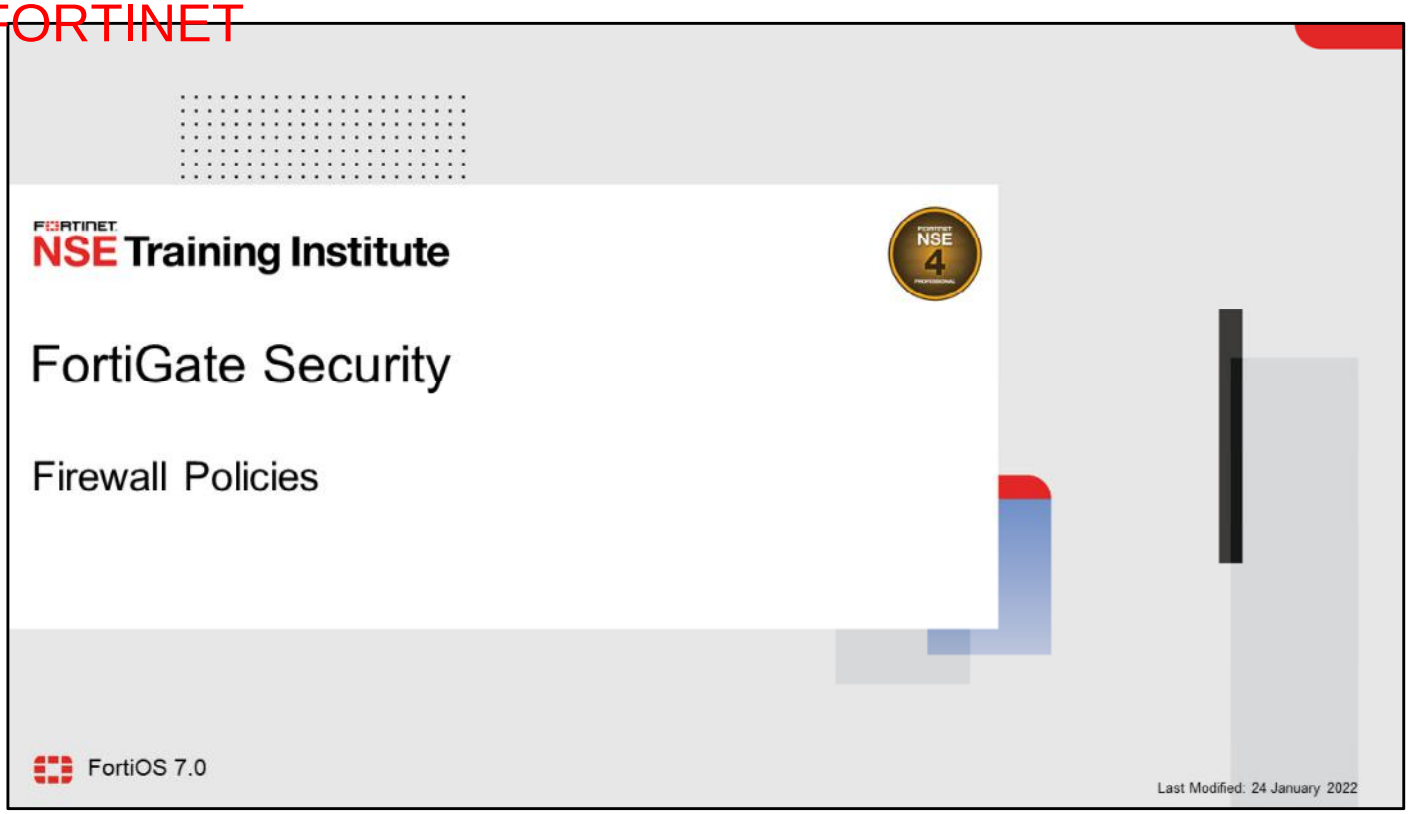

In this lesson, you will learn about firewall policies and how to apply them to allow and deny traffic passing through FortiGate. At its core, FortiGate is a firewall, so almost everything that it does to your traffic is linked to your firewall policies.

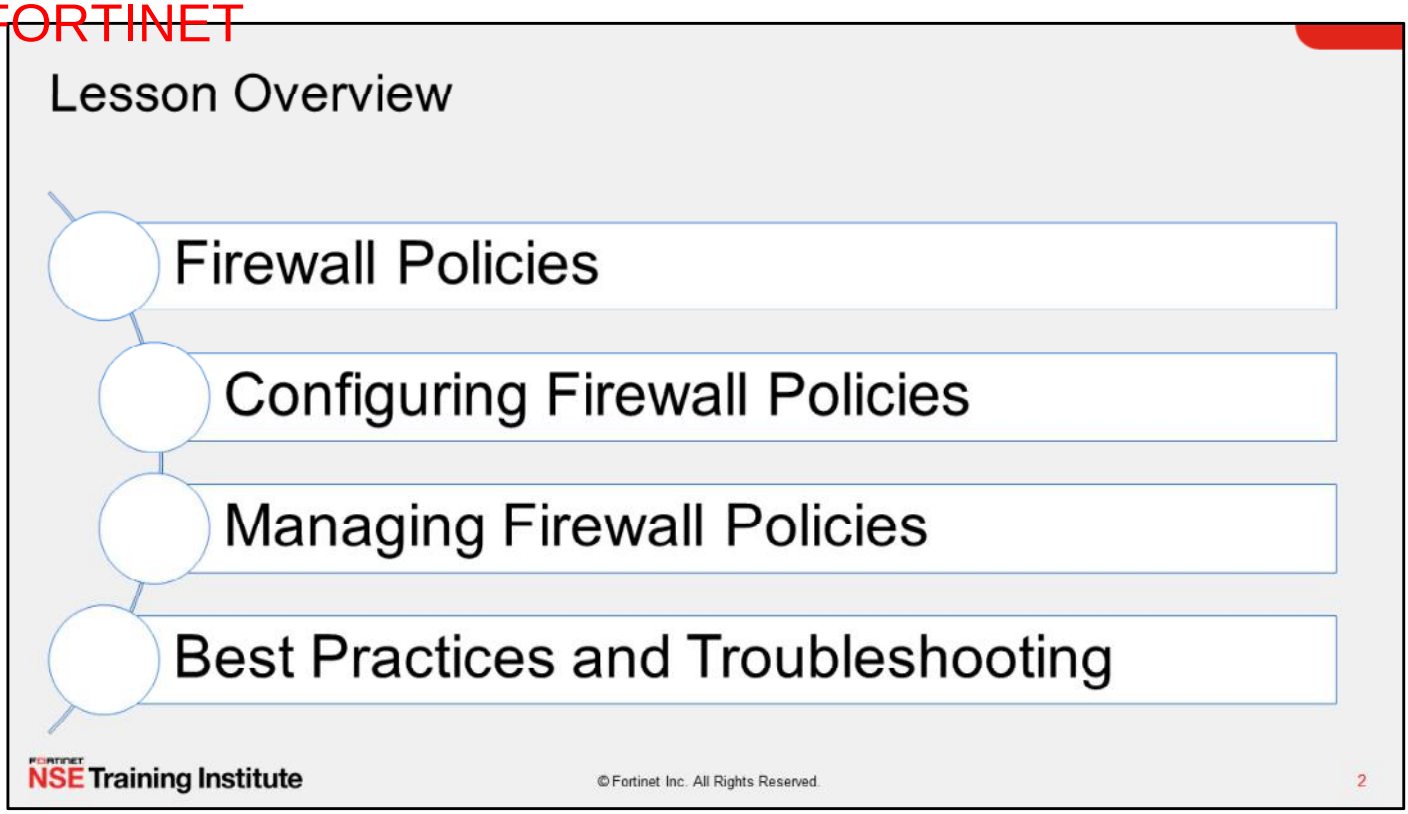

In this lesson, you will learn about the topics shown on this slide.

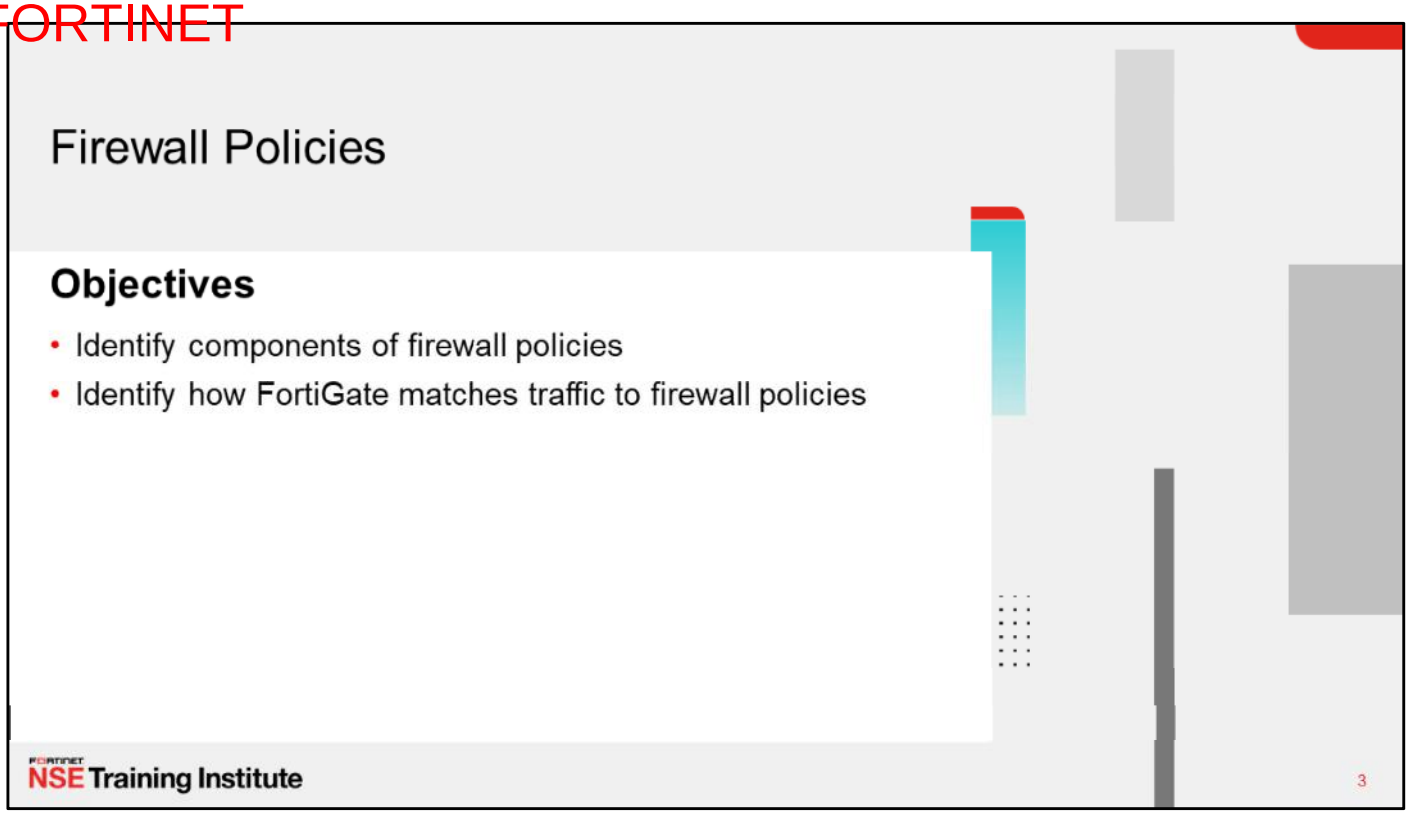

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

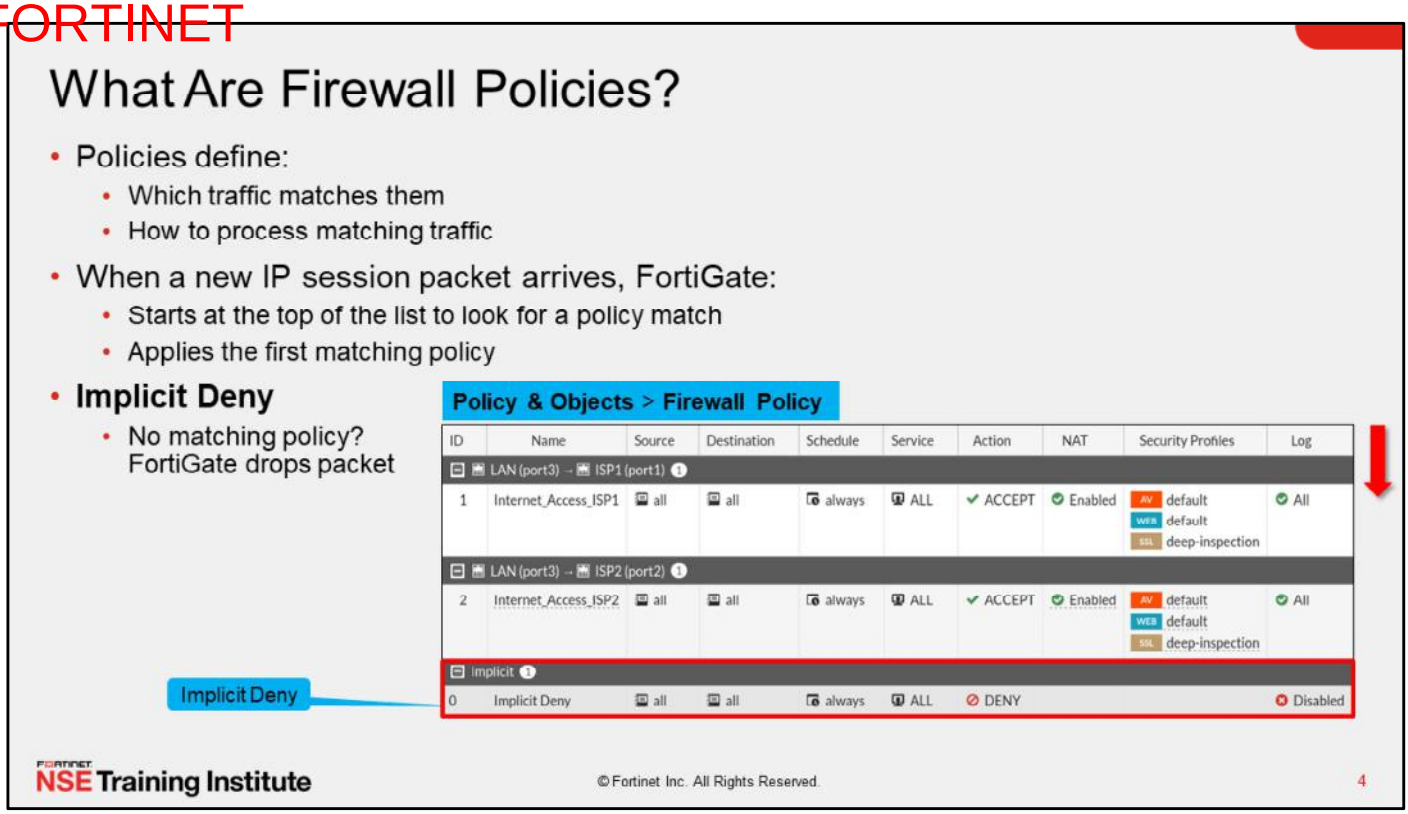

To begin, you will learn about what firewall policies are.

Firewall policies define which traffic matches them and what FortiGate does when traffic does match.

Should the traffic be allowed? Initially, FortiGate bases this decision on simple criteria, such as the source of the traffic. Then, 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. Those scans could block the traffic if, for example, it 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.

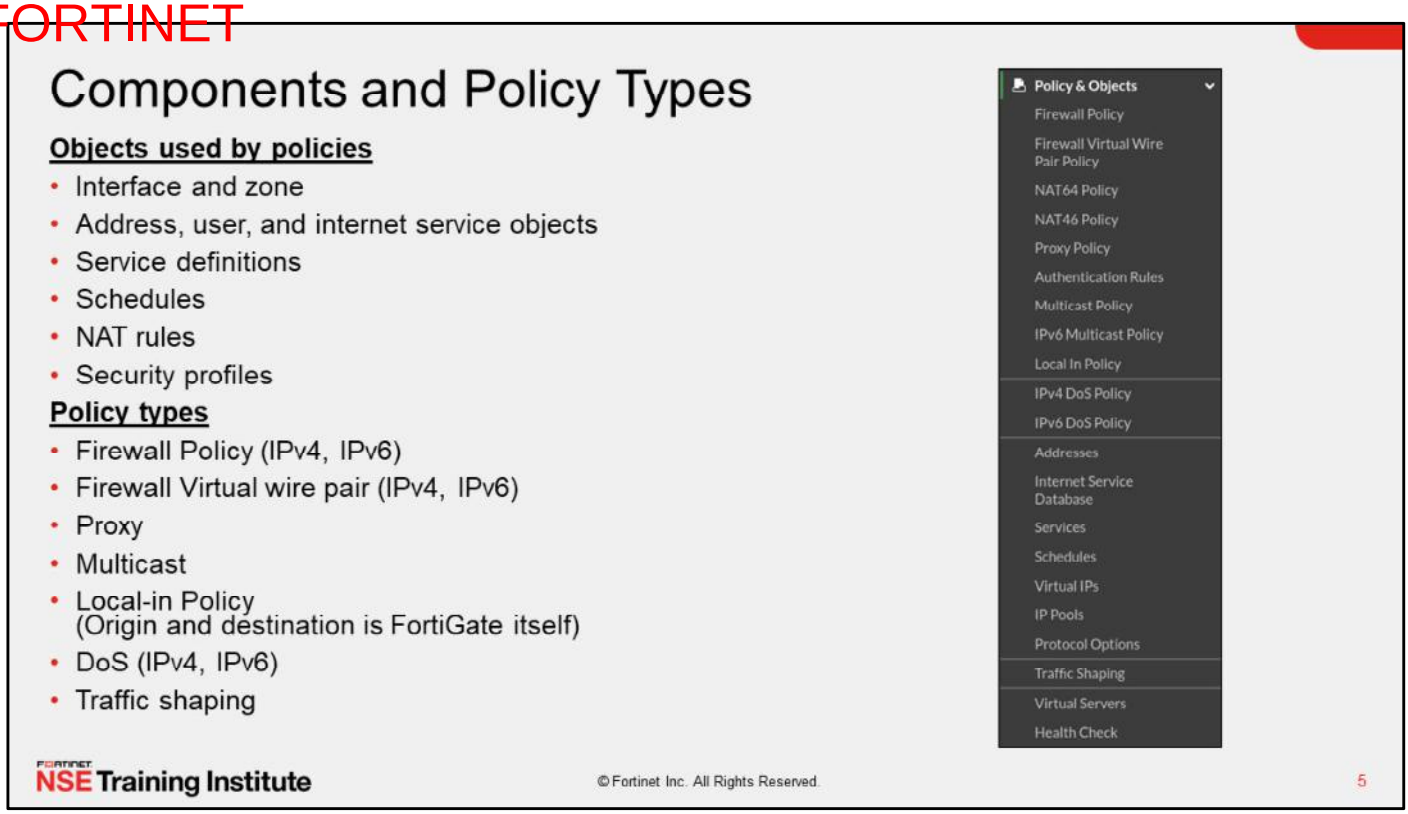

Each policy matches traffic and applies security by referring to the objects that you've defined, such as addresses and profiles.

What about other firewall policy types? Do IPv6 or virtual wire policies exist? Yes. These policies use slightly different objects that are relevant to their type. In this lesson, you will learn about IPv4 firewall policies, because they are the most common use case.

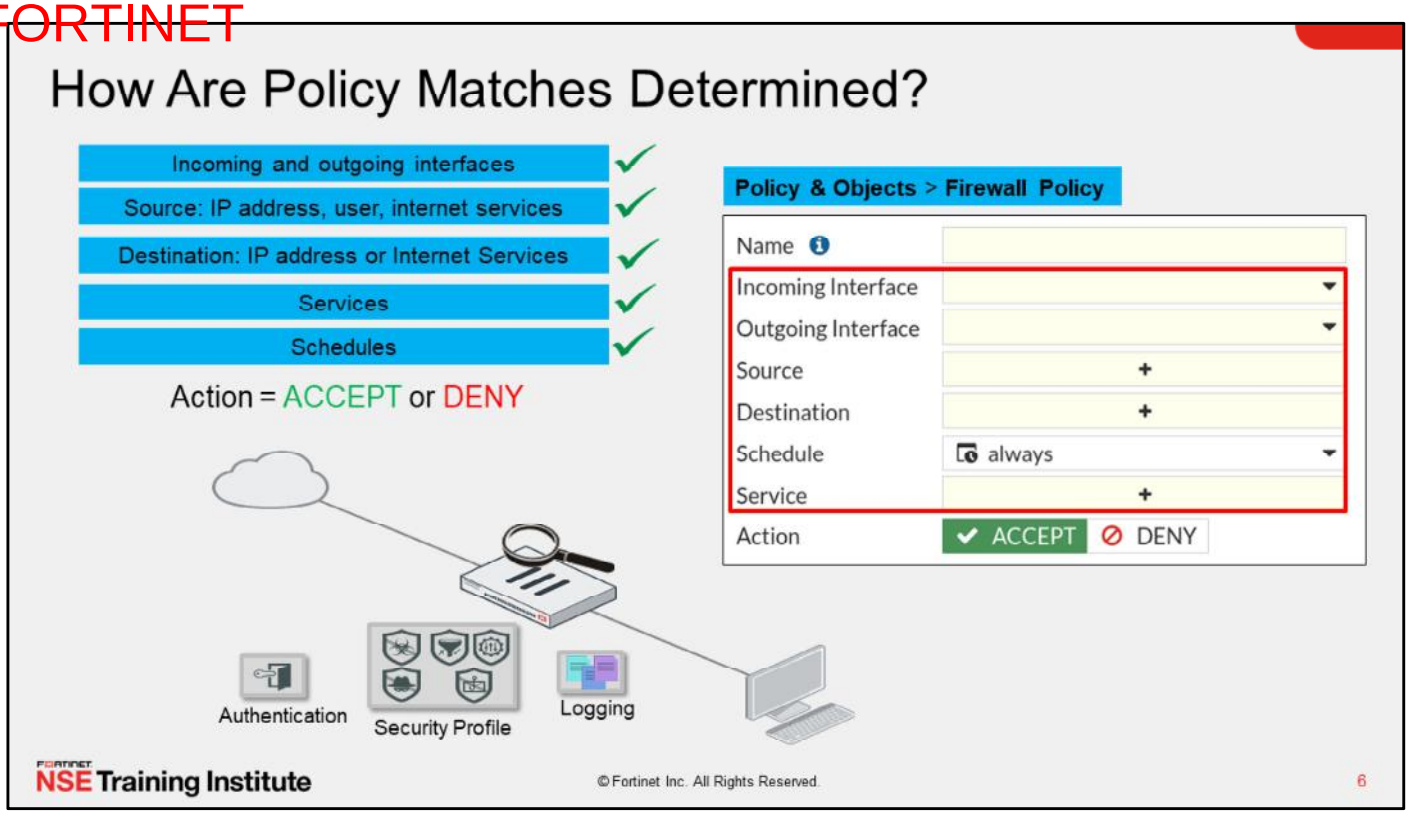

When a packet arrives, how does FortiGate find a matching policy? Each policy has match criteria, which you can define using the following objects:

- **Incoming Interface**
- **Outgoing Interface**
- **Source**: IP address, user, internet services
- **Destination**: IP address or internet services
- **Service**: IP protocol and port number
- **Schedule**: Applies during configured times

When 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 applies other configured settings for packet processing, such as antivirus scanning, web filtering, or source NAT.

For example, if you want to block incoming FTP to all but a few FTP servers, you would define the addresses of your FTP servers, select those 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 (usually FTP servers are always available, day or night). Finally, you would set the **Action** setting to **ACCEPT**.

This *might* be enough, but often you'll want more thorough security. Here, the policy also authenticates the user, scans for viruses, and logs blocked connection attempts.

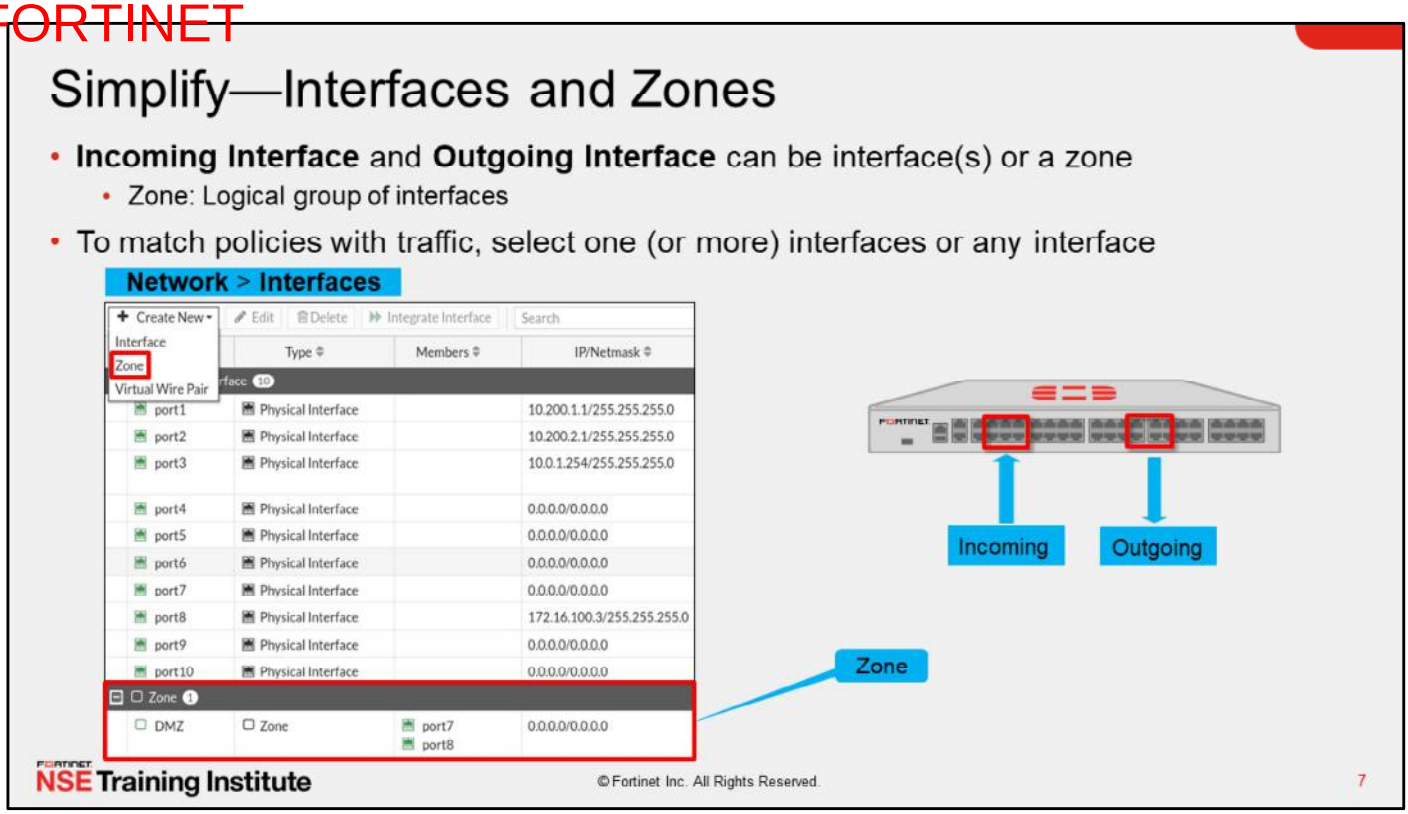

To begin describing how FortiGate finds a policy for each packet, let's start with the interface(s).

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 would *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 zone. 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.

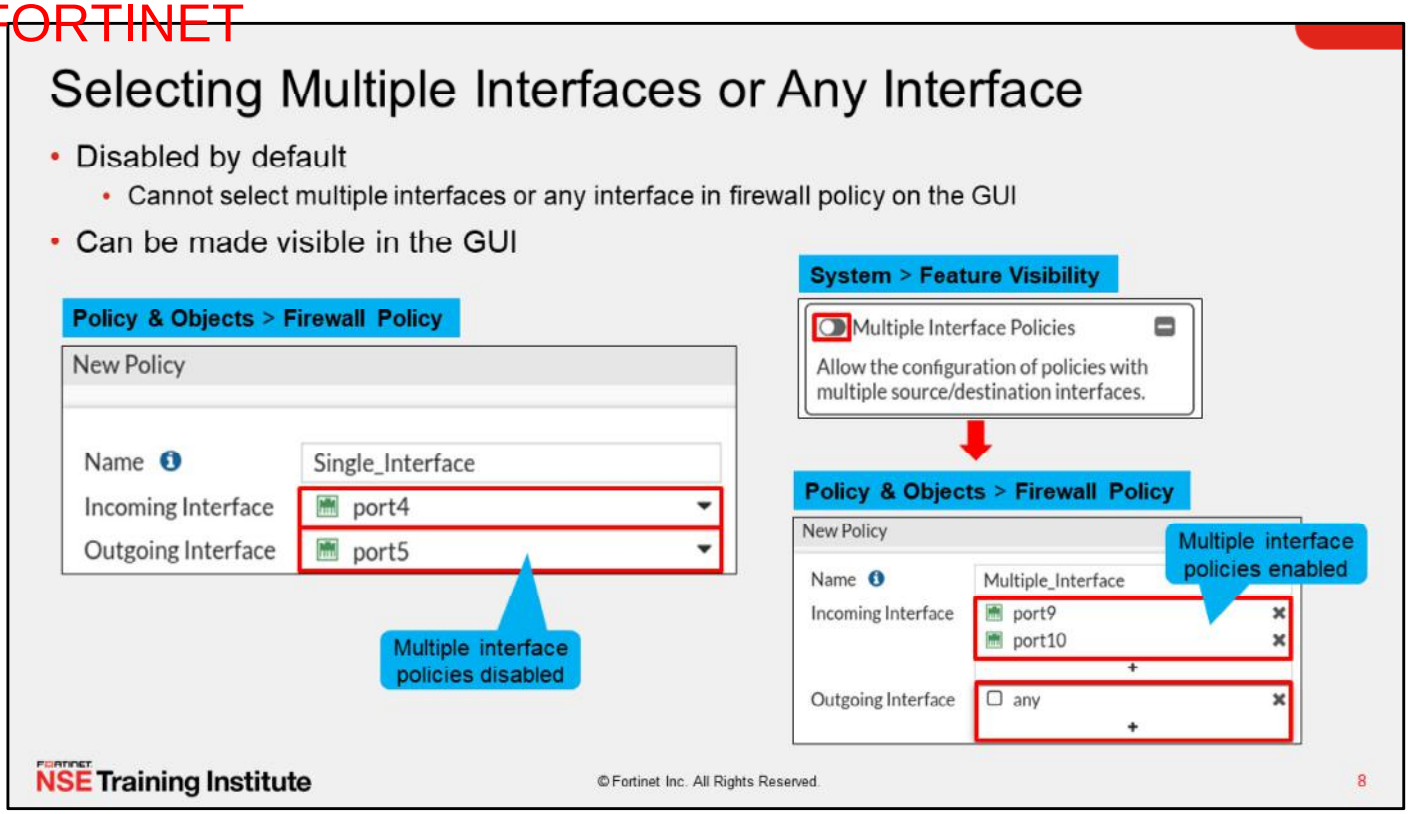

By default, you can select only a single interface as the incoming interface and a single interface as the outgoing interface. This is because the option to select multiple interfaces, or **any** interface in a firewall policy, is disabled on the GUI. However, you can enable the **Multiple Interface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

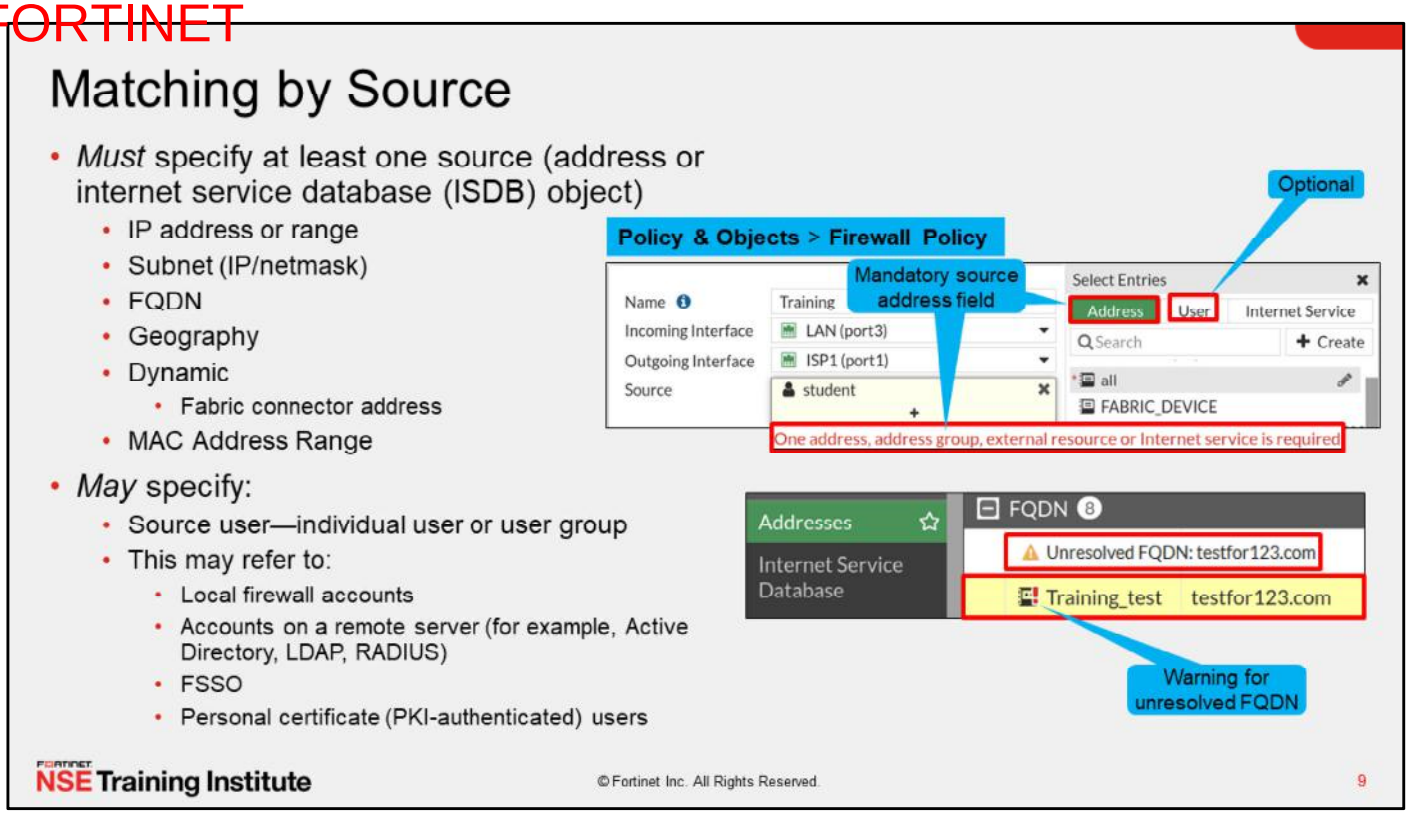

The next match criteria that FortiGate considers is the packet's source.

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.

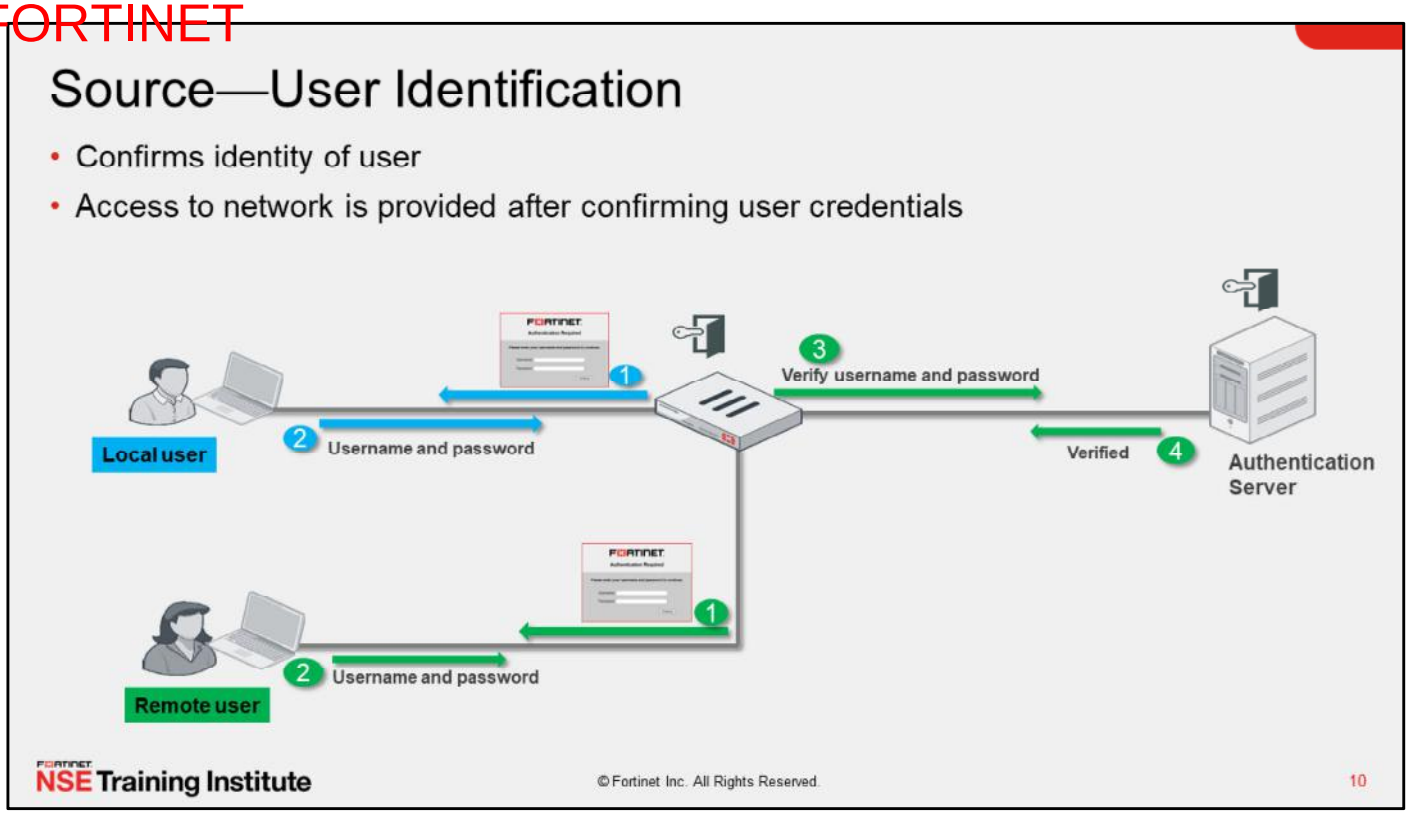

If a user is added as part of the source, FortiGate must verify the user before allowing or denying access based on the firewall policy. There are different ways that a user can authenticate.

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.

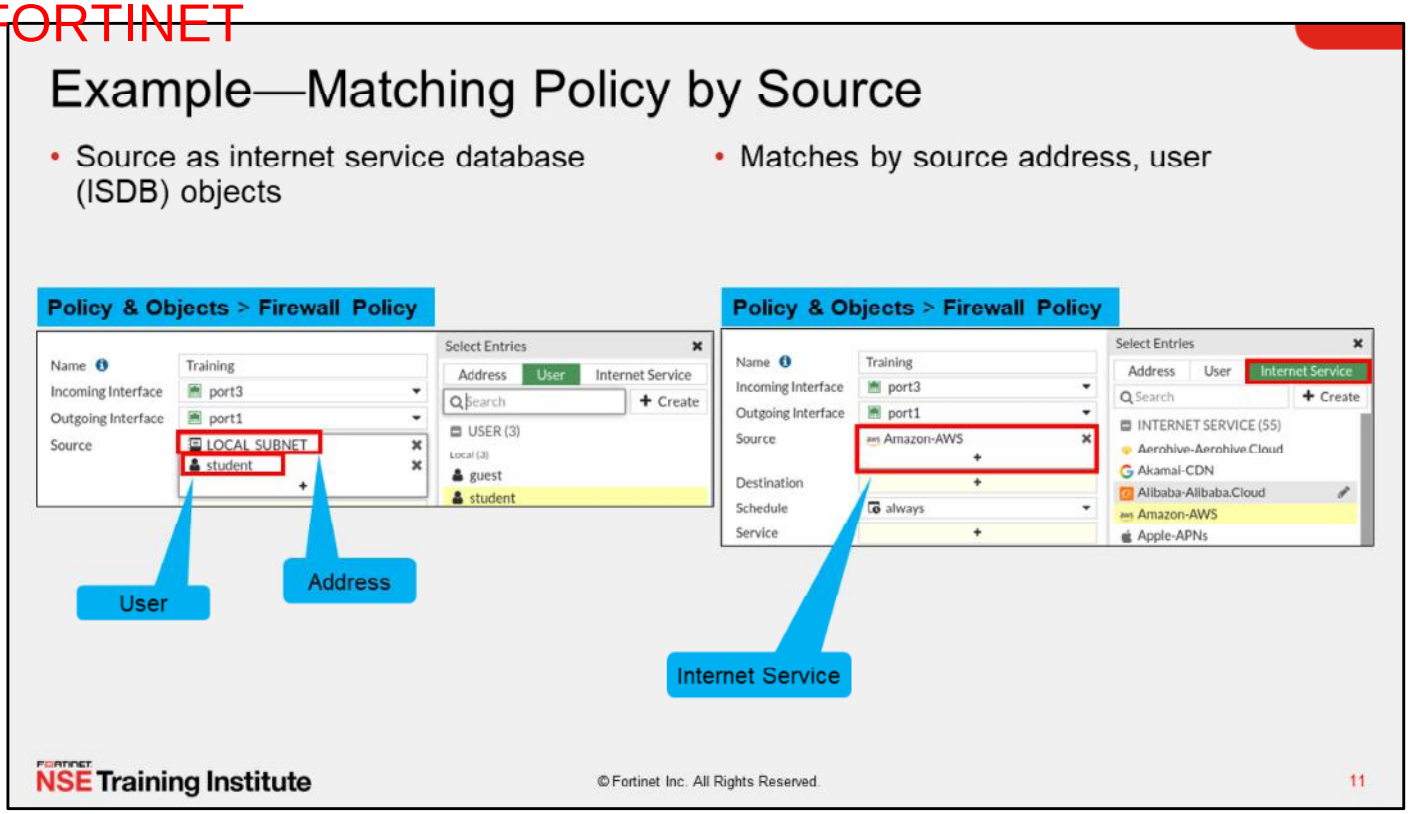

In the example shown on this slide, source selectors identify the specific subnet and user group. Remember, user is an optional object. The user object is used here to make the policy more specific. If you wanted the policy to match more traffic, you would leave the user object undefined.

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.

Matching by Destination

Like source, destination criteria can use:

- Address objects:
  - Subnet (IP or netmask)
  - IP address or address range
  - FQDN
    - DNS query used to resolve FQDN
  - Geography
    - Country defines addresses by ISP's geographical location
    - Database updated periodically through FortiGuard
  - Dynamic
    - Fabric connector address
- Internet service database (ISDB) objects

NSE Training Institute

12

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, which are what actually appear in the IP header.

You can select geographic addresses, which are groups or ranges of addresses allocated to a country, can be selected instead. You update these objects through FortiGuard.

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 user authentication is successful.

Internet Service

- Database that contains IP addresses, IP protocols, and port numbers used by the most common internet services
  - Regularly updated through FortiGuard
- Can be used as **Source** or **Destination** in the firewall policy
- If **Internet Service** is selected as **Source**:
  - You cannot use Address in the Source
- If Internet Service is selected as **Destination**:
  - You cannot use **Address** in the **Destination**
  - You cannot select **Service** in the firewall policy

Policy & Objects > Internet Service Database

Name	Direction	Number of Entries
Alibaba-SSH	Destination	4,347
Alibaba-Web	Destination	4,347
Amazon-AWS	Both	14,015
Amazon-AWS.WorkSpaces.Gateway	Destination	27
Amazon-DNS	Destination	41,821
Amazon-FTP	Destination	41,821

Policy & Objects > Firewall Policy

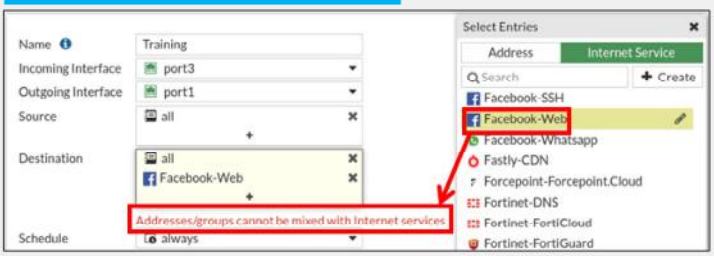

NSE Training Institute

13

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.

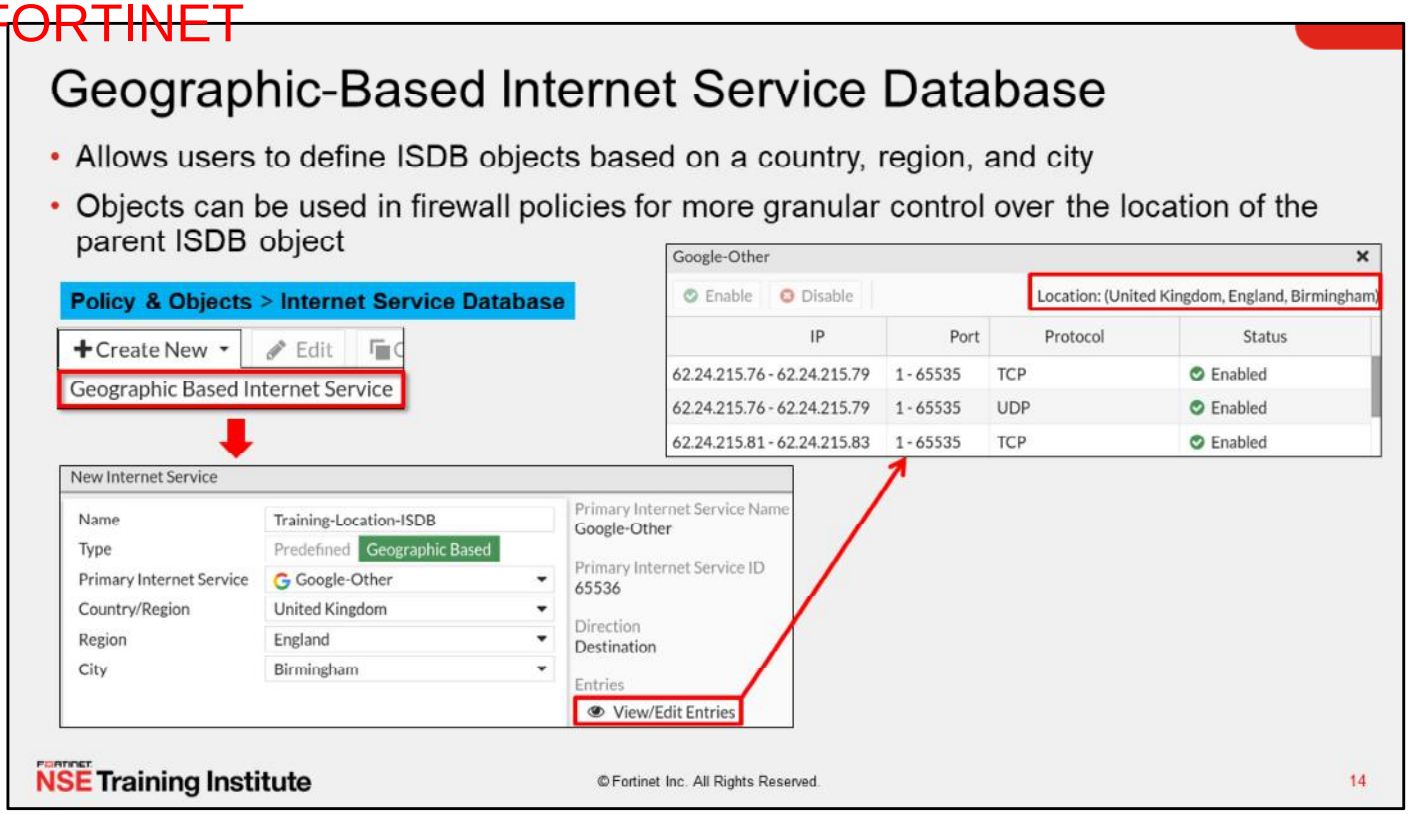

Geographic-based ISDB objects allow users to define a country, region, and city. These objects can be used in firewall policies for more granular control over the location of the parent ISDB object.

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

Internet Service Database (ISDB)—Updates

- You can disable ISDB updates so they occur only during a change control window
  - Control ISDB updates by using CLI command:

```
# config system fortiguard
        set update-ffdb [enable | disable]
   next
end
```

- Once ISDB updates are disabled, other scheduled FortiGuard updates do not update ISDB
- By default, ISDB updates are enabled

NSE Training Institute

15

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.

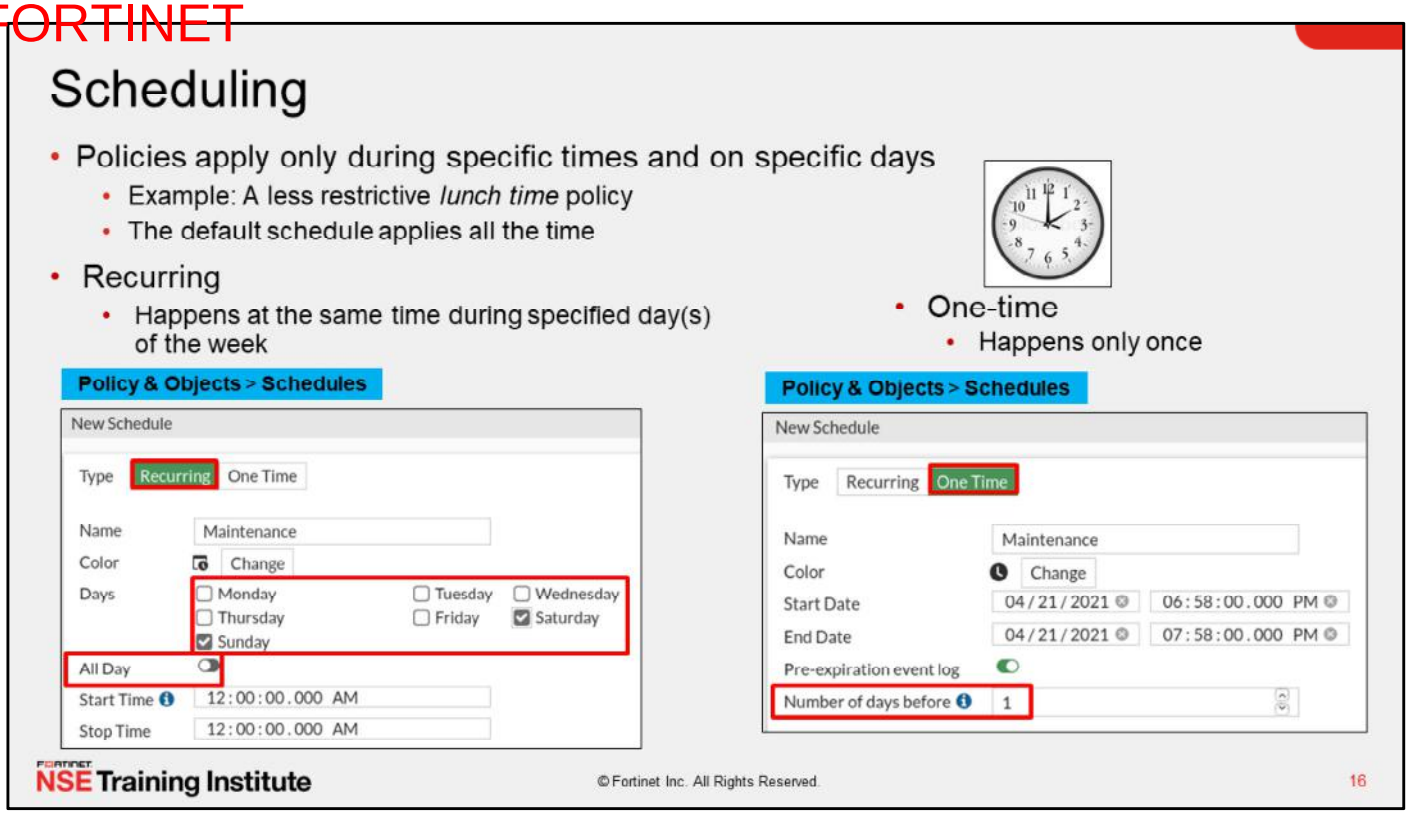

Schedules add a time element to the policy. For example, you might use a policy to allow backup software to activate at night, or create a test window for a remote address that is allowed for testing purposes.

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.

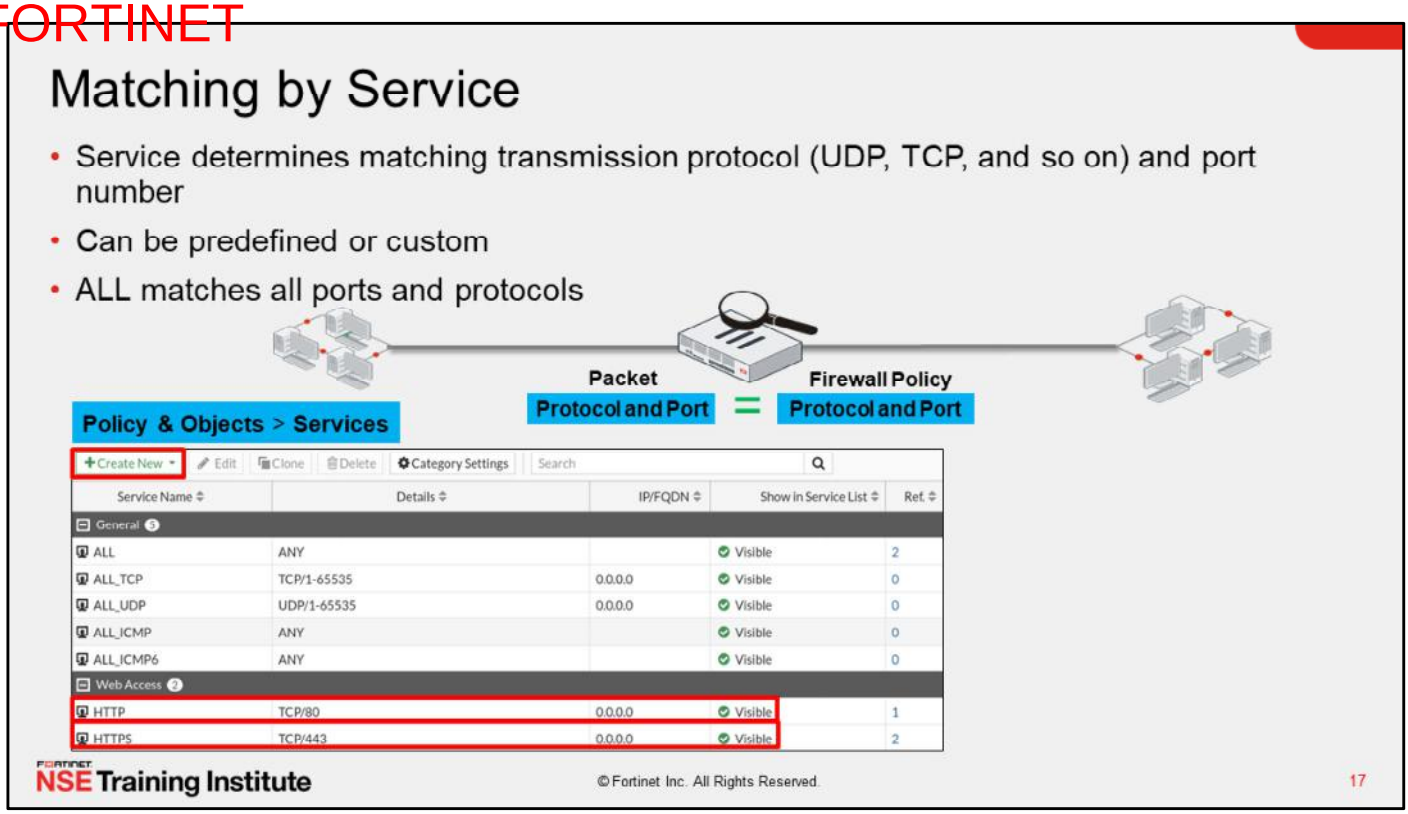

Another criterion that FortiGate uses to match policies is the packet's service.

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

NSE Training Institute

18

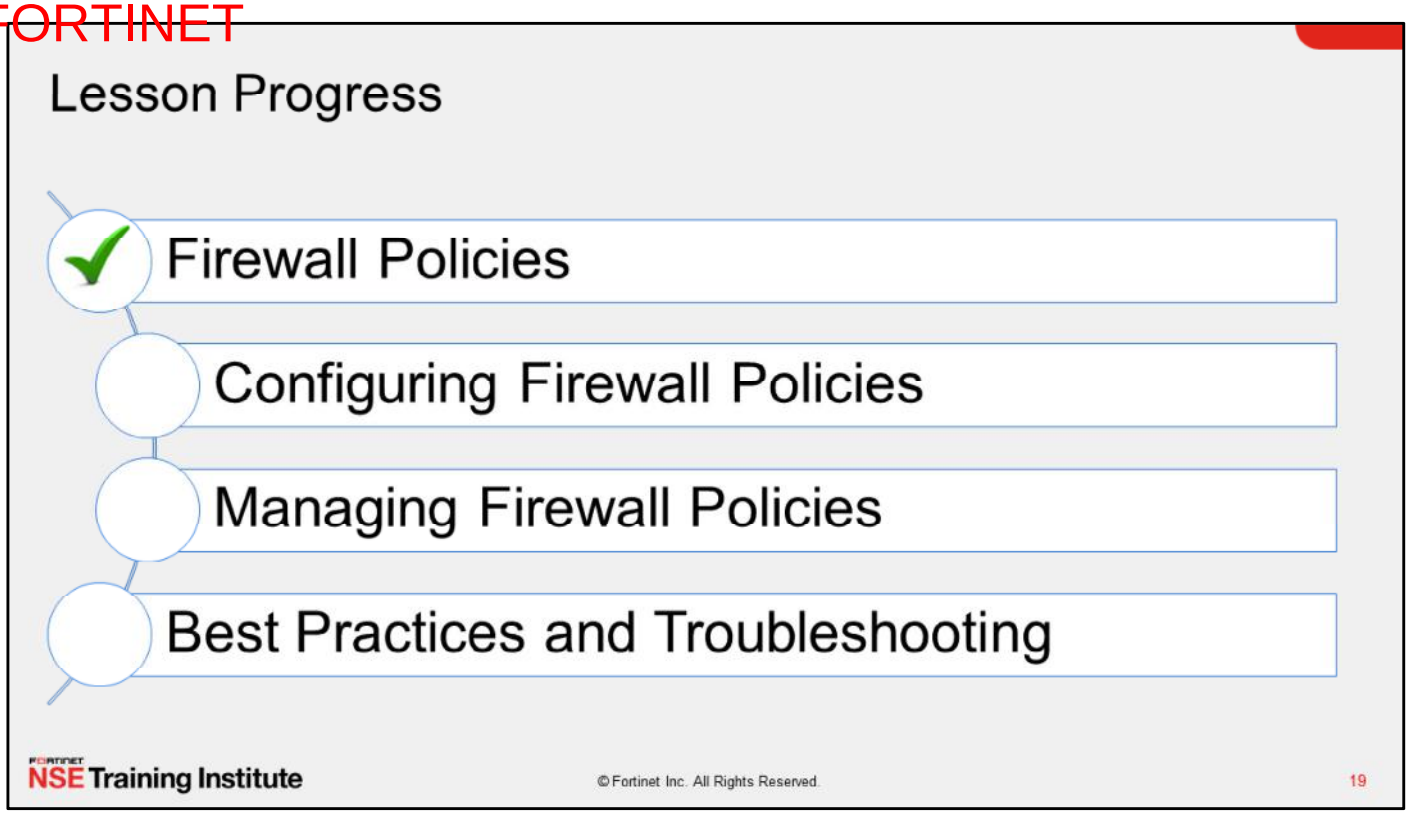

Good job! You now understand the components used in firewall policies and matching criteria used by FortiGate.

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

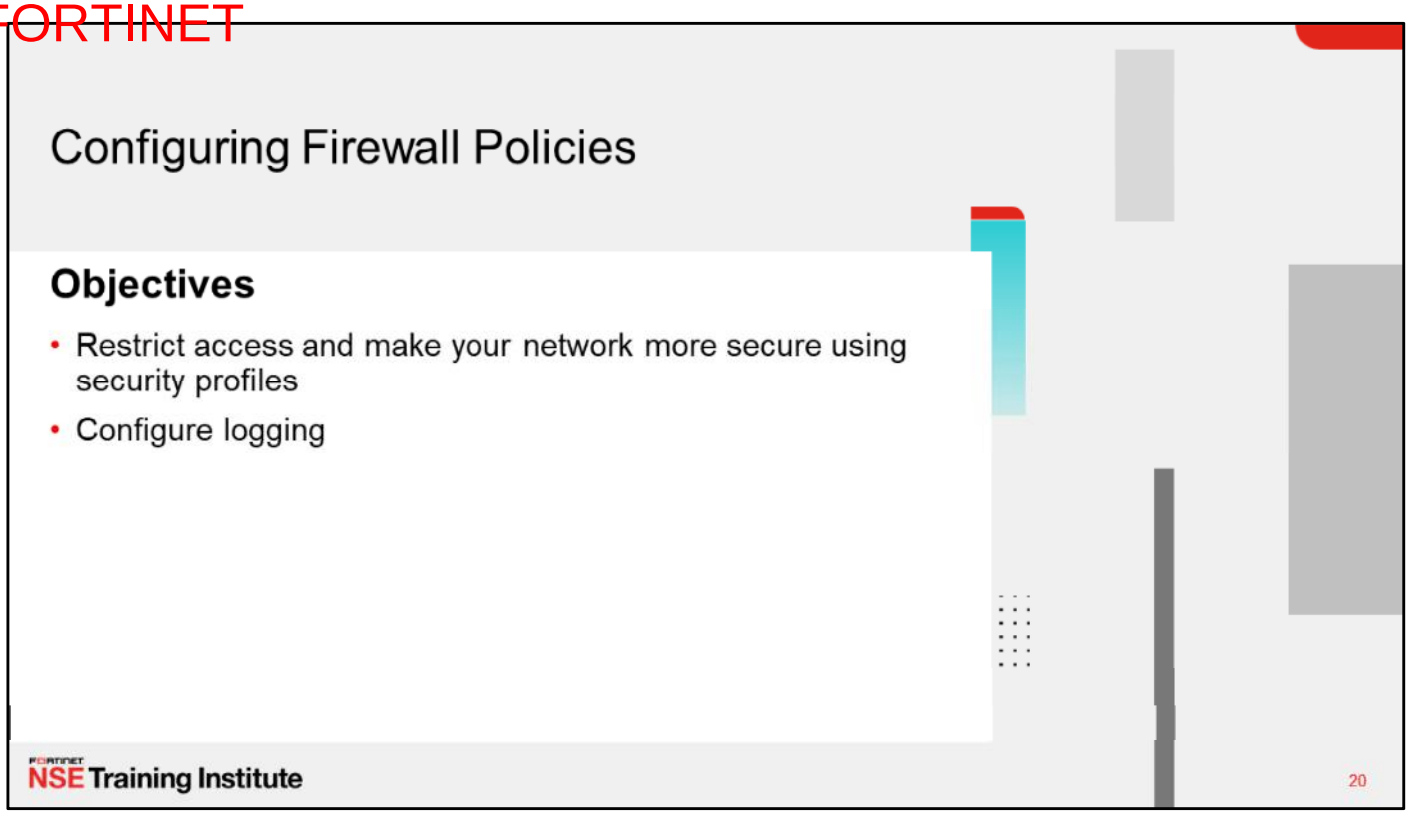

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

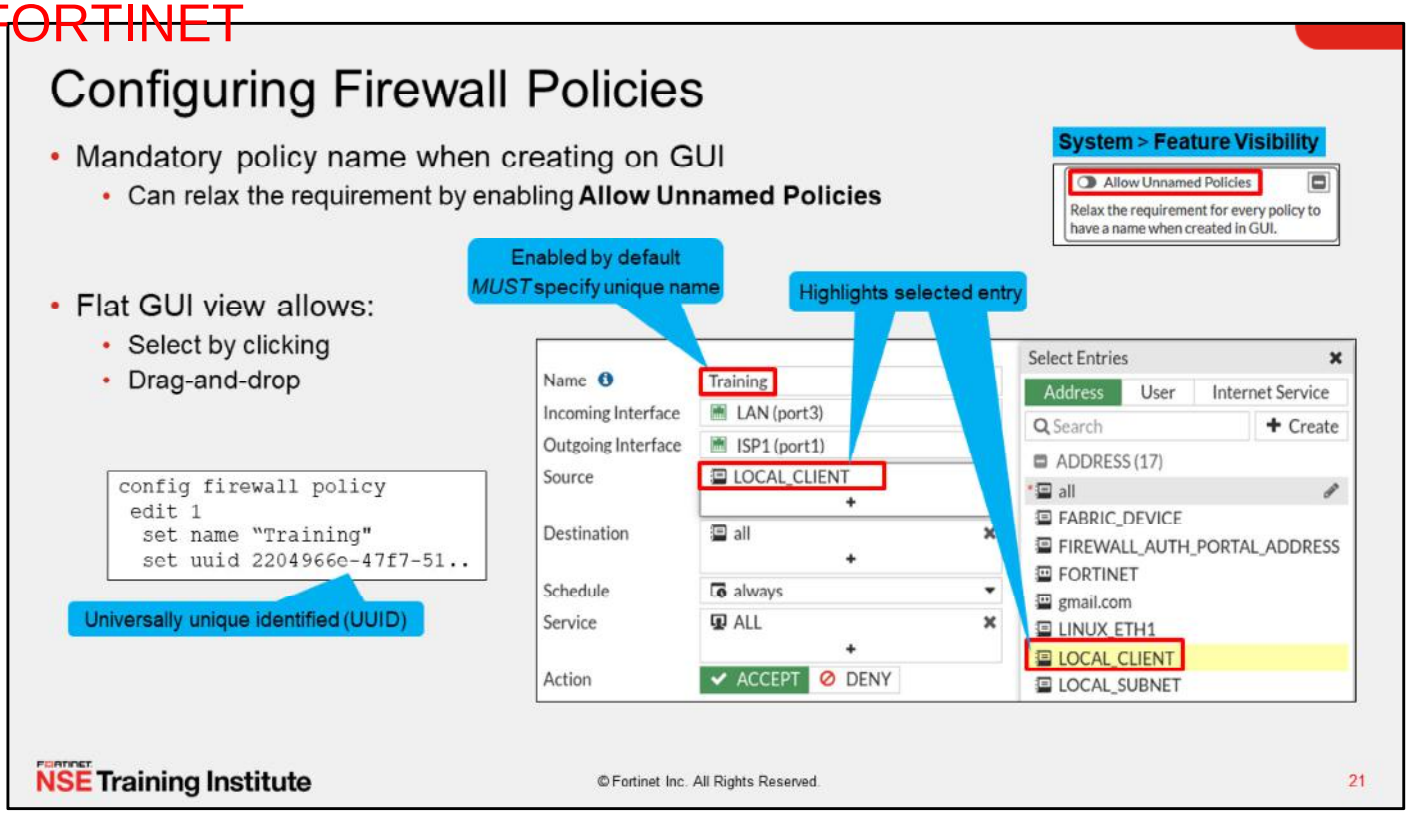

When you configure a new firewall policy on the GUI, you *must* specify a unique name for the firewall policy because it is enabled by default, while it is optional on the CLI. This helps the administrator to quickly identify the policy that they are looking for. However, you can make this feature optional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

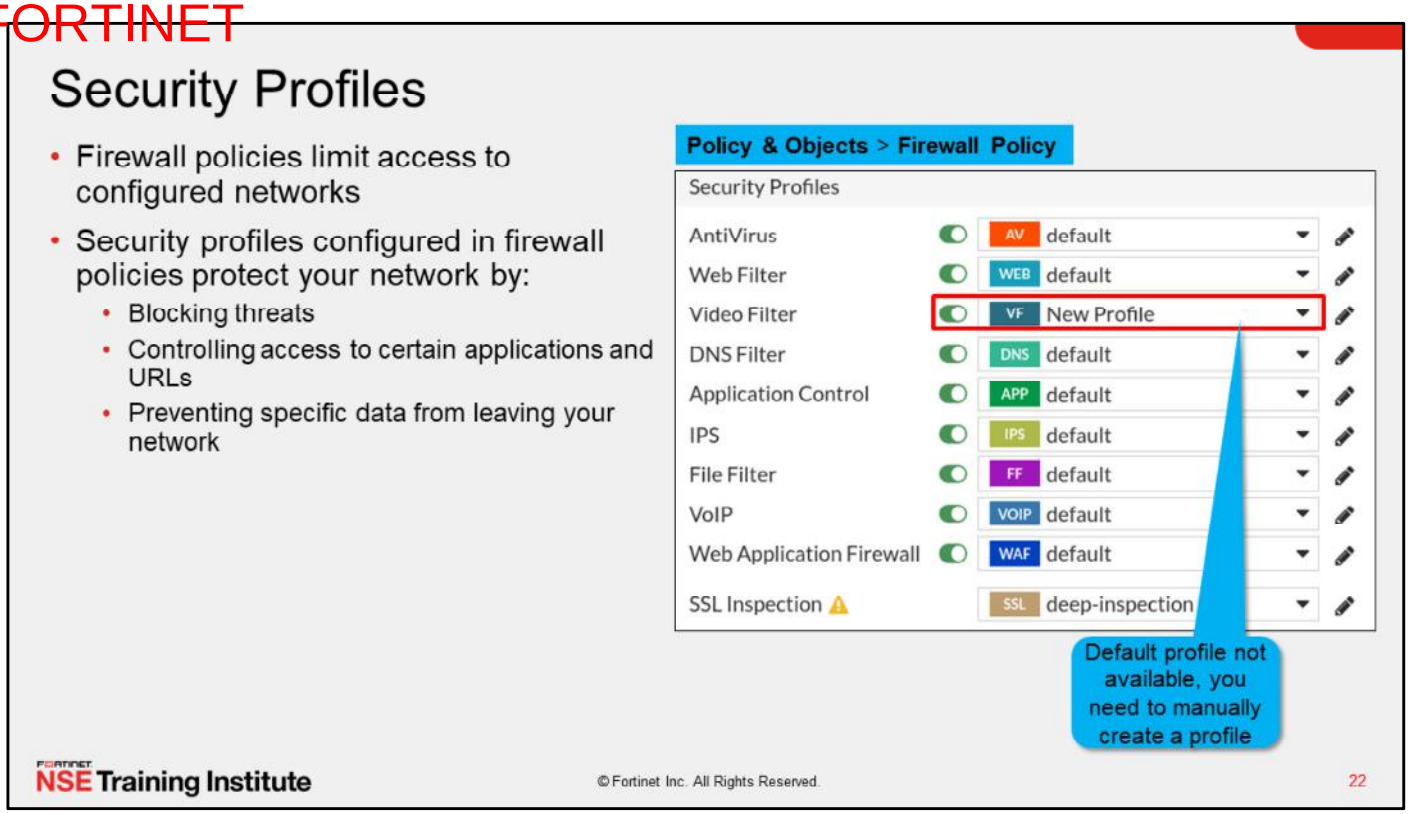

One of the most important features that a firewall policy can apply is security profiles, such as IPS and antivirus. A security profile inspects each packet in the traffic flow, where the session has already been conditionally accepted by the firewall policy.

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.

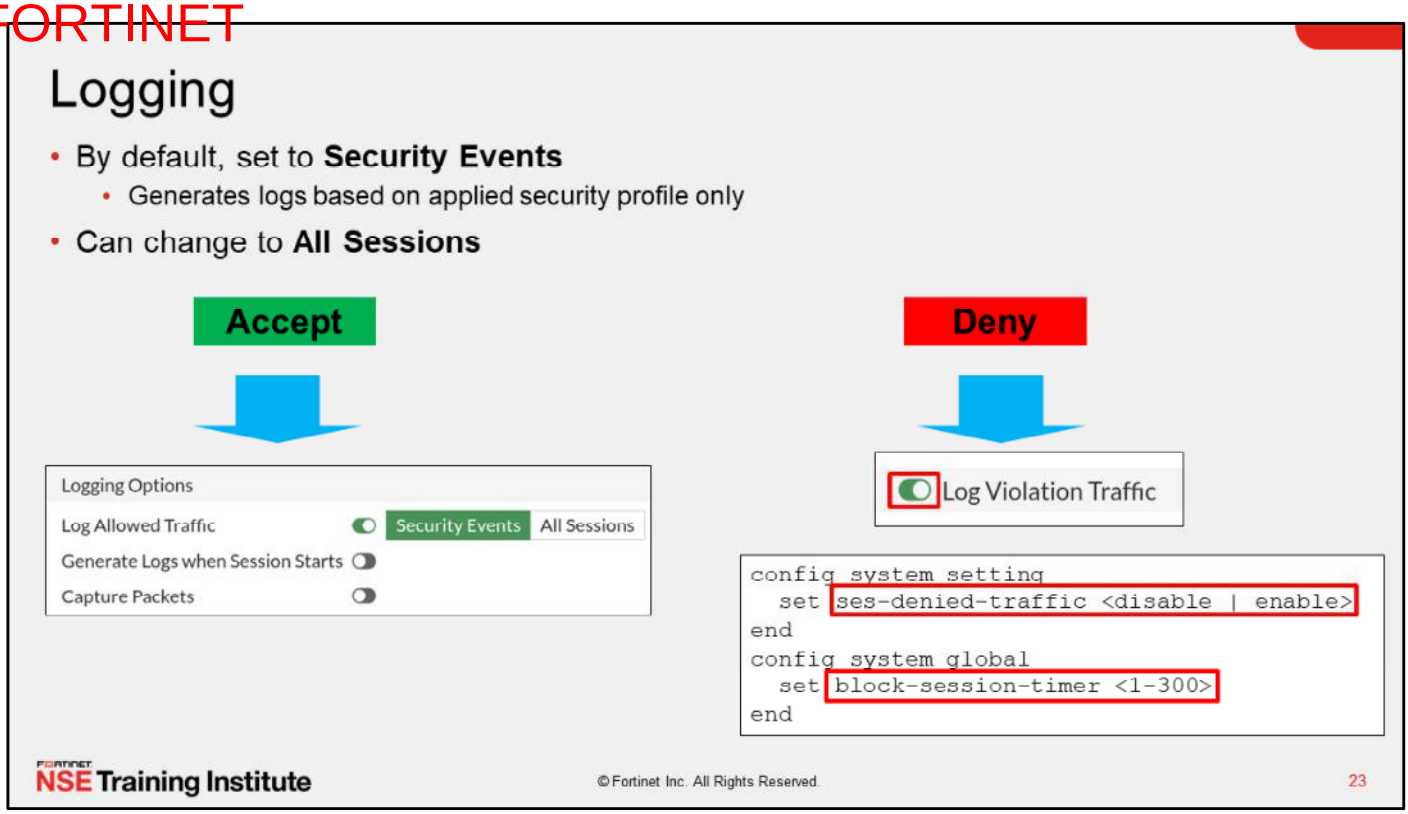

If you have enabled logging in the policy, FortiGate generates traffic logs after a firewall policy closes an IP session.

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`.

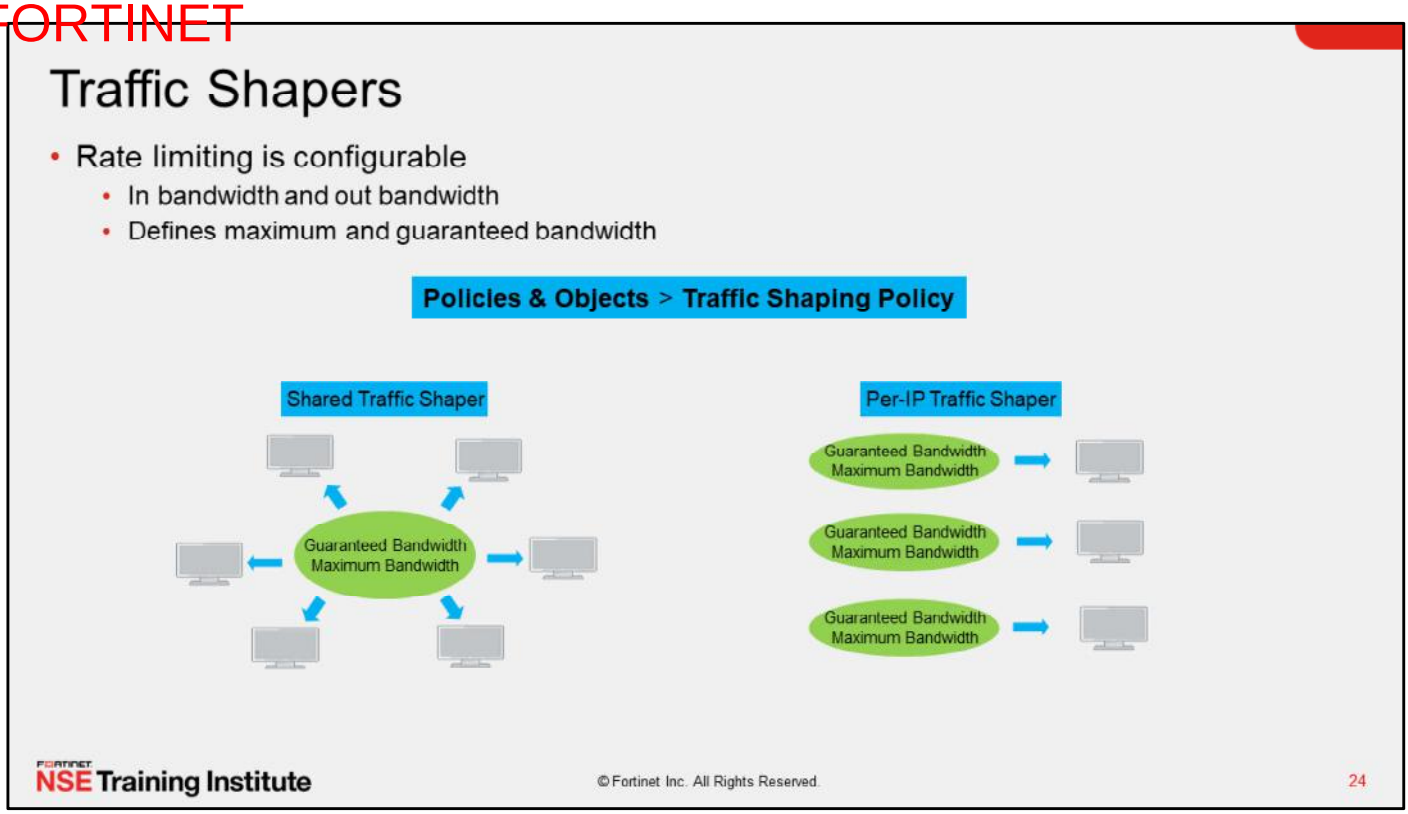

You can configure two types of traffic shapers: shared and per IP.

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.

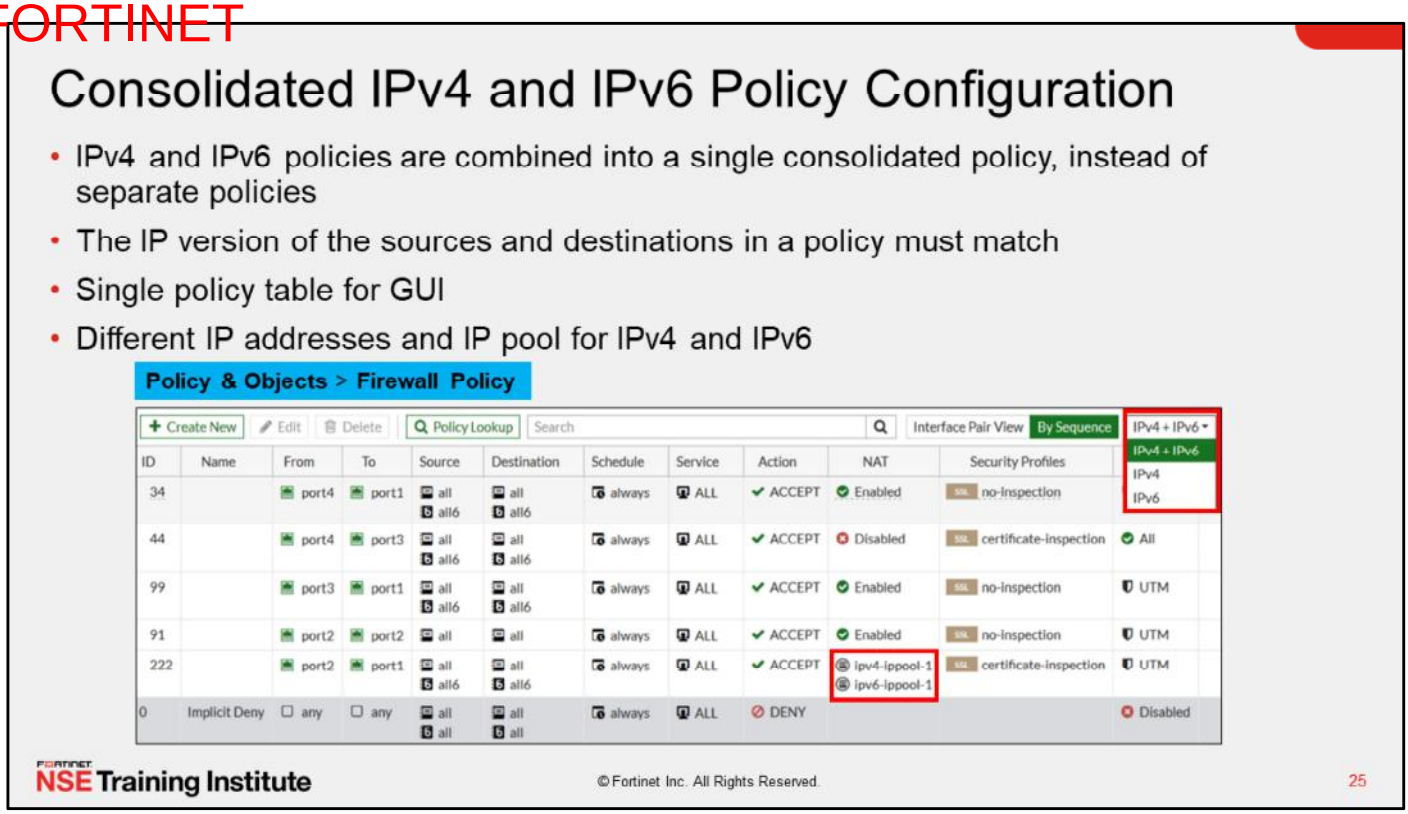

By default, IPv4 and IPv6 policies are combined into a single consolidated policy, rather than creating and maintaining two different policy sets for IPv4 and IPv6.

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

NSE Training Institute

26

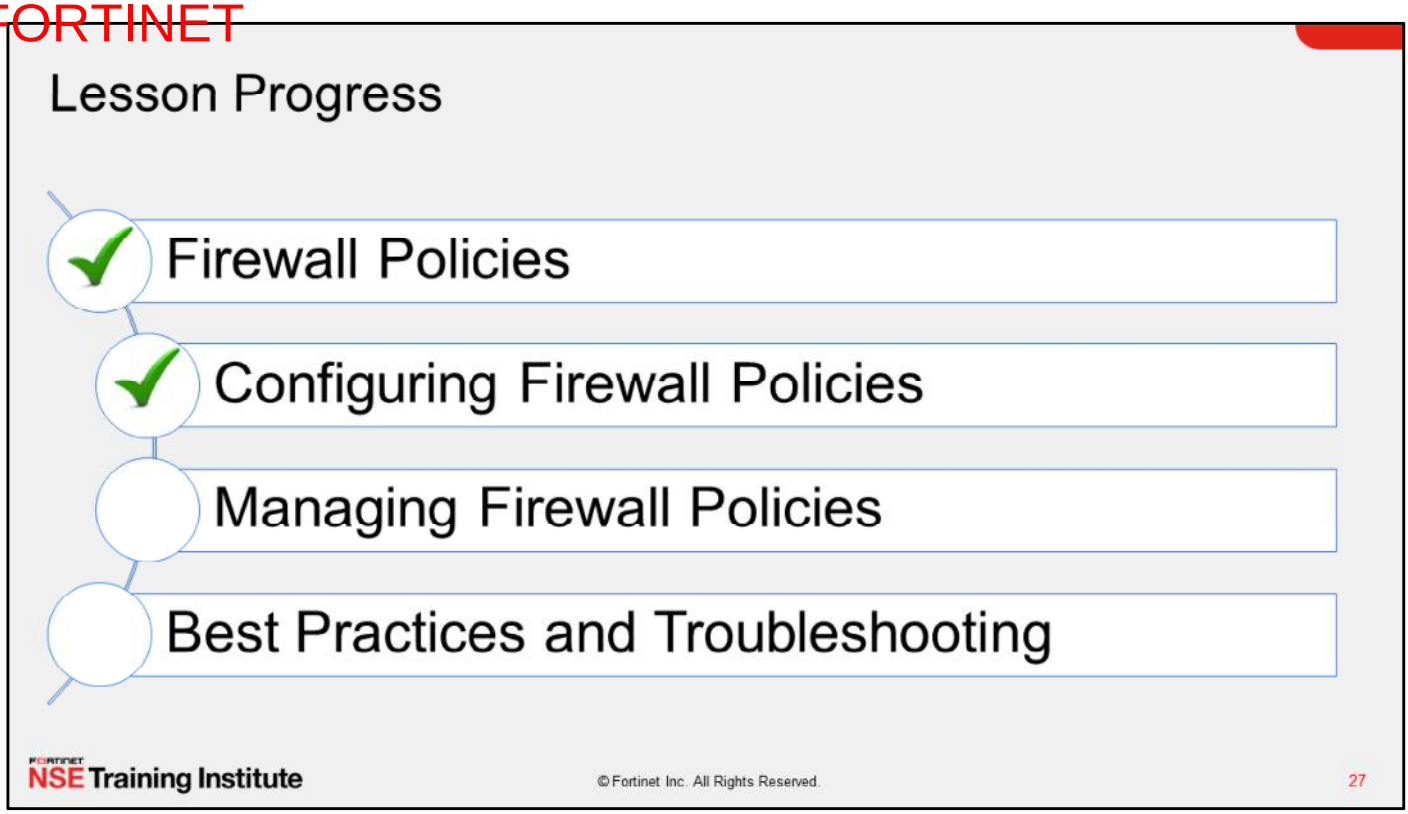

Good job! You now understand how to configure firewall policies on FortiGate.

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

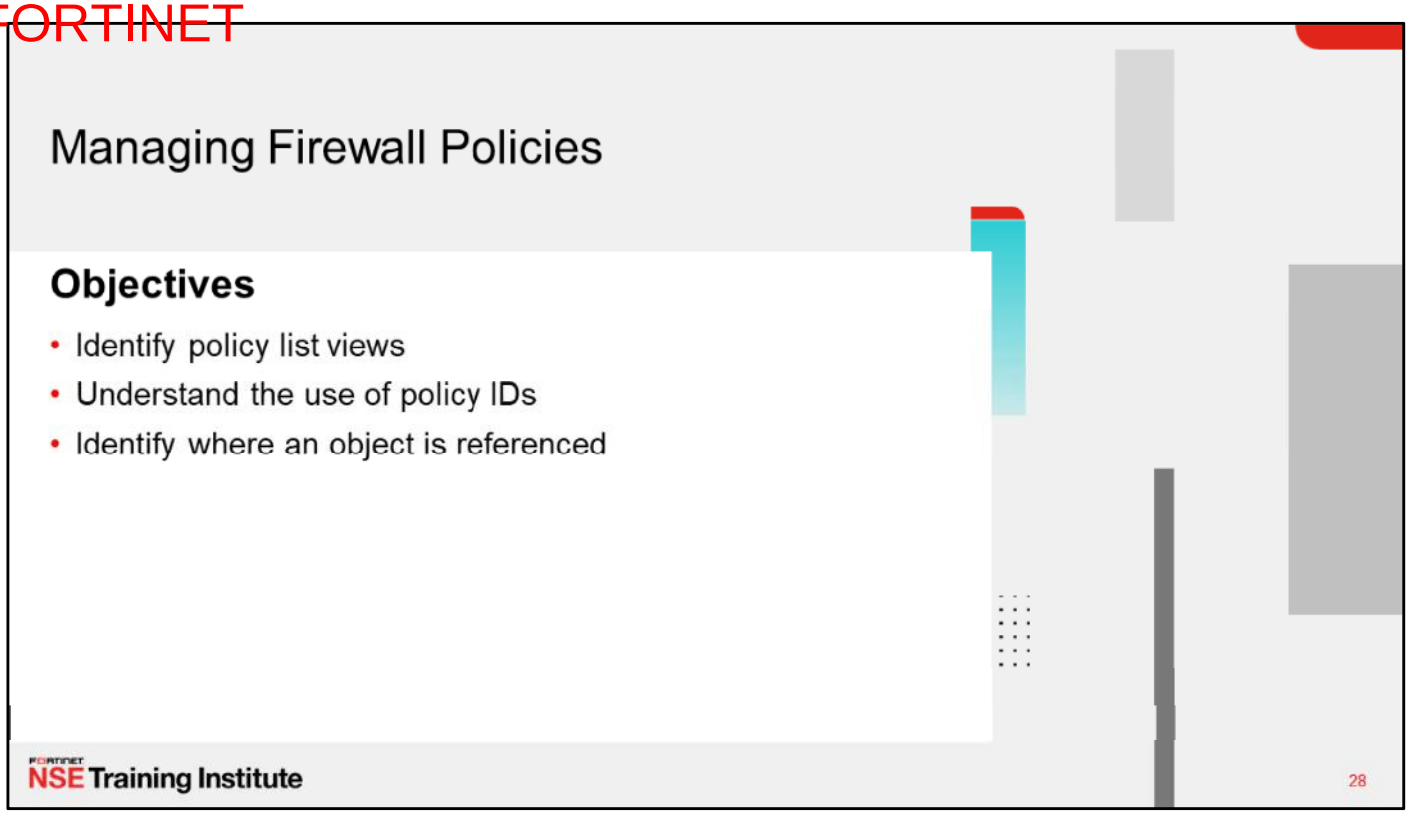

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

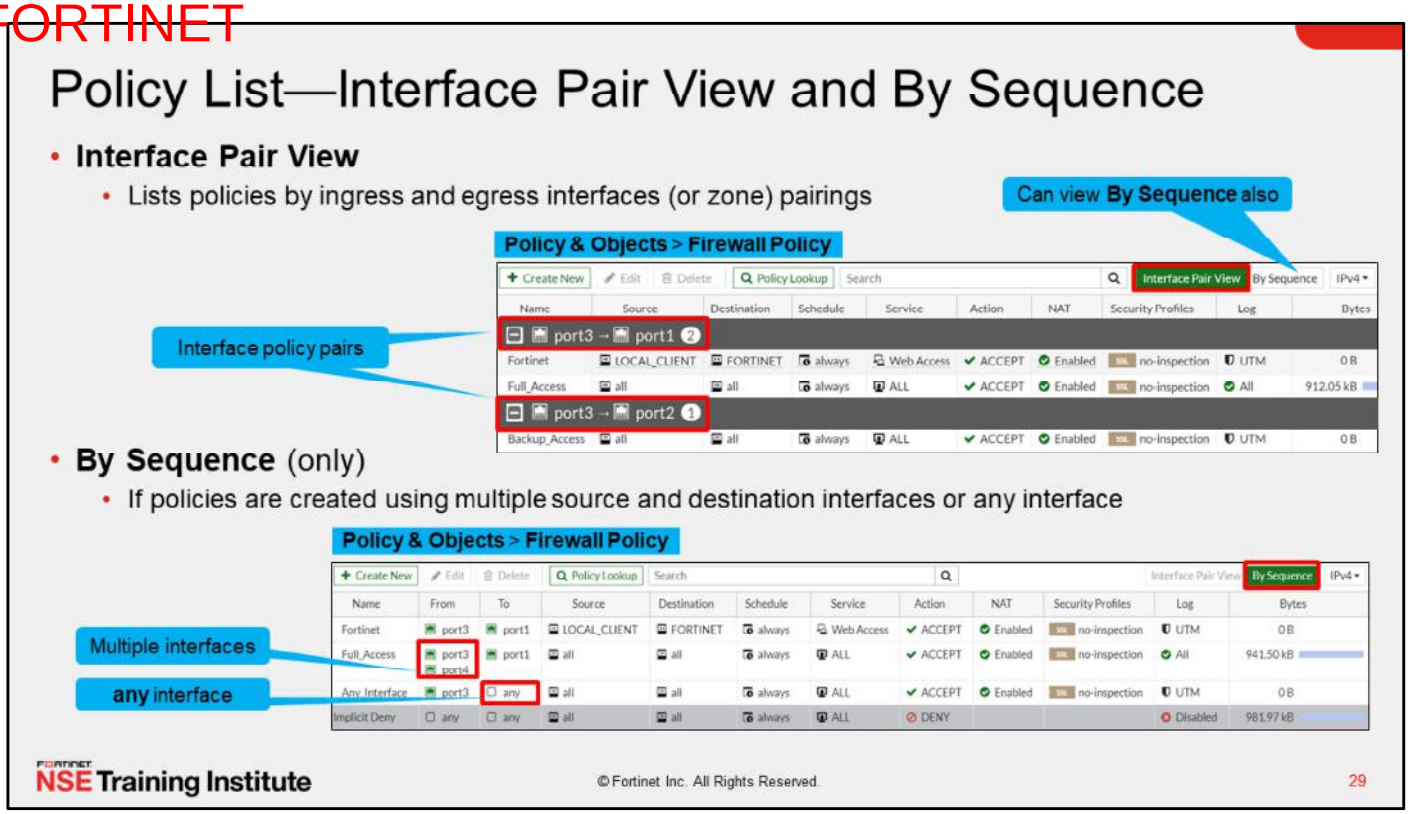

Firewall policies appear in an organized list. The list is organized either in **Interface Pair View** or **By Sequence**.

Usually, the list will appear in **Interface Pair View**. Each section contains policies for that ingress-egress pair. Alternatively, you can view your policies as a single, comprehensive list by selecting **By Sequence** at the top of the page.

In some cases, you won't have a choice of which view is used.

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. So instead, policies are then always displayed in a single list (**By Sequence**).

To help you remember the use of each interface, you can give them 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.

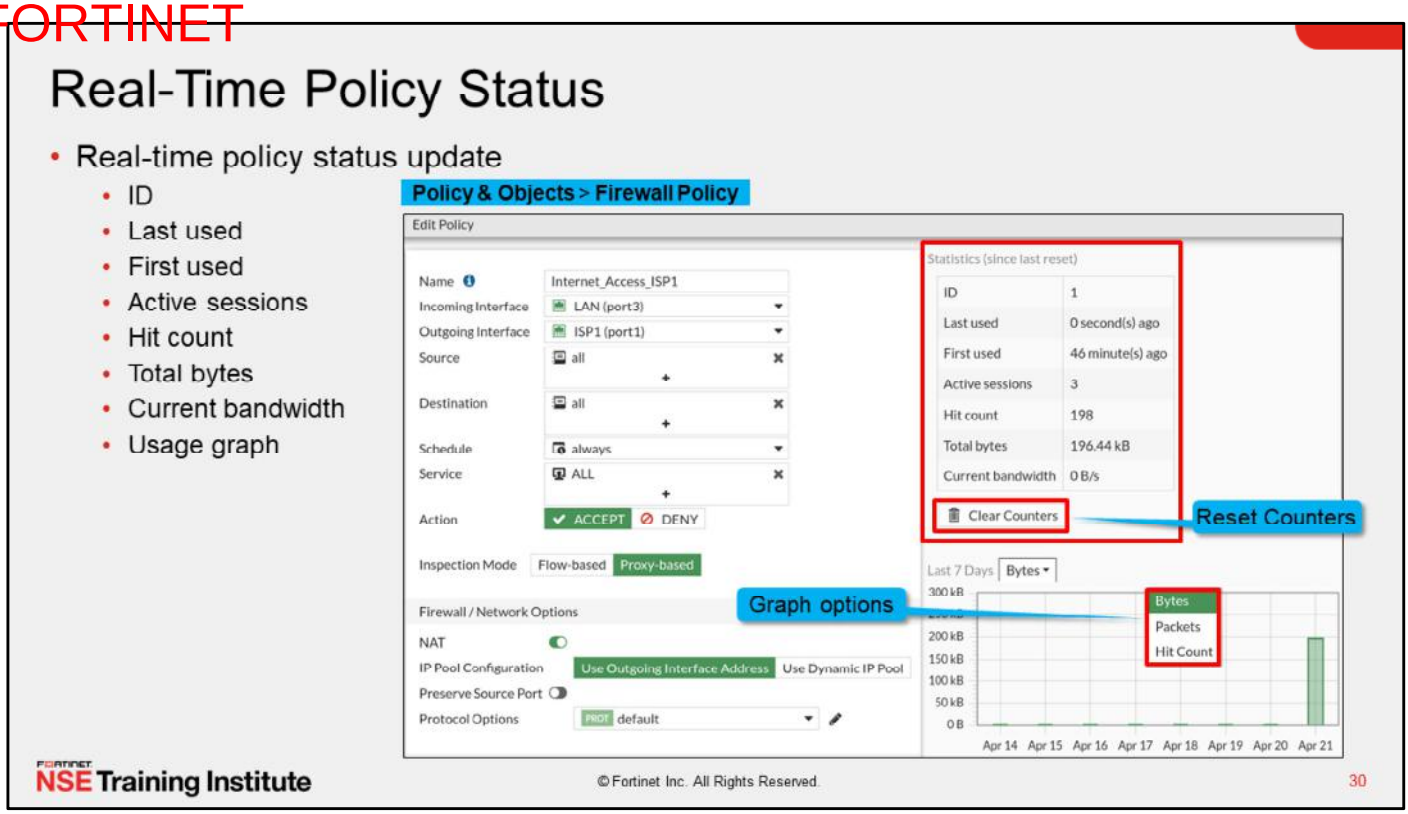

When you edit the policy, policy information will be visible.

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.

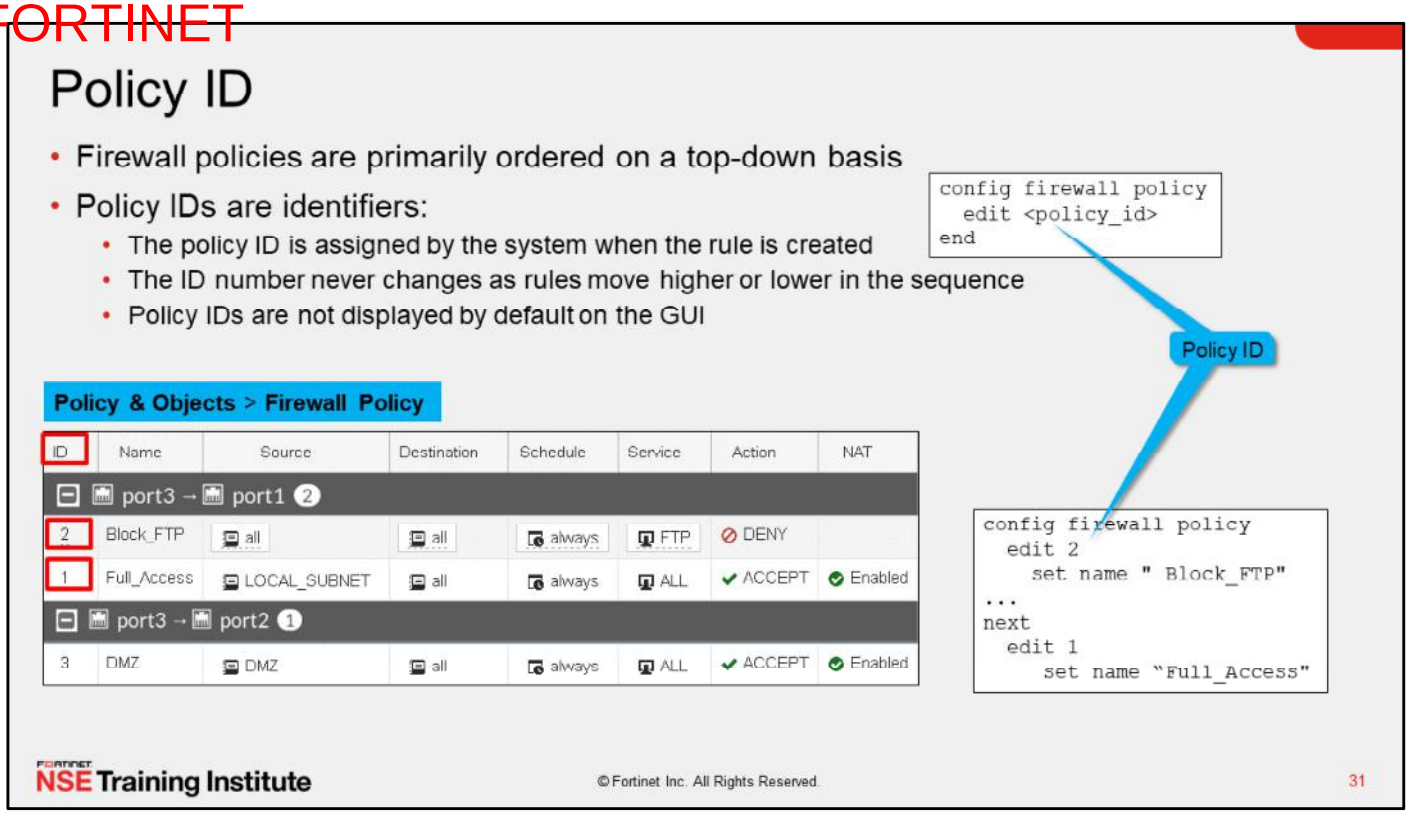

An important concept to understand about how firewall policies work is the precedence of order, or, if you prefer a more recognizable term, first come, first served.

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.

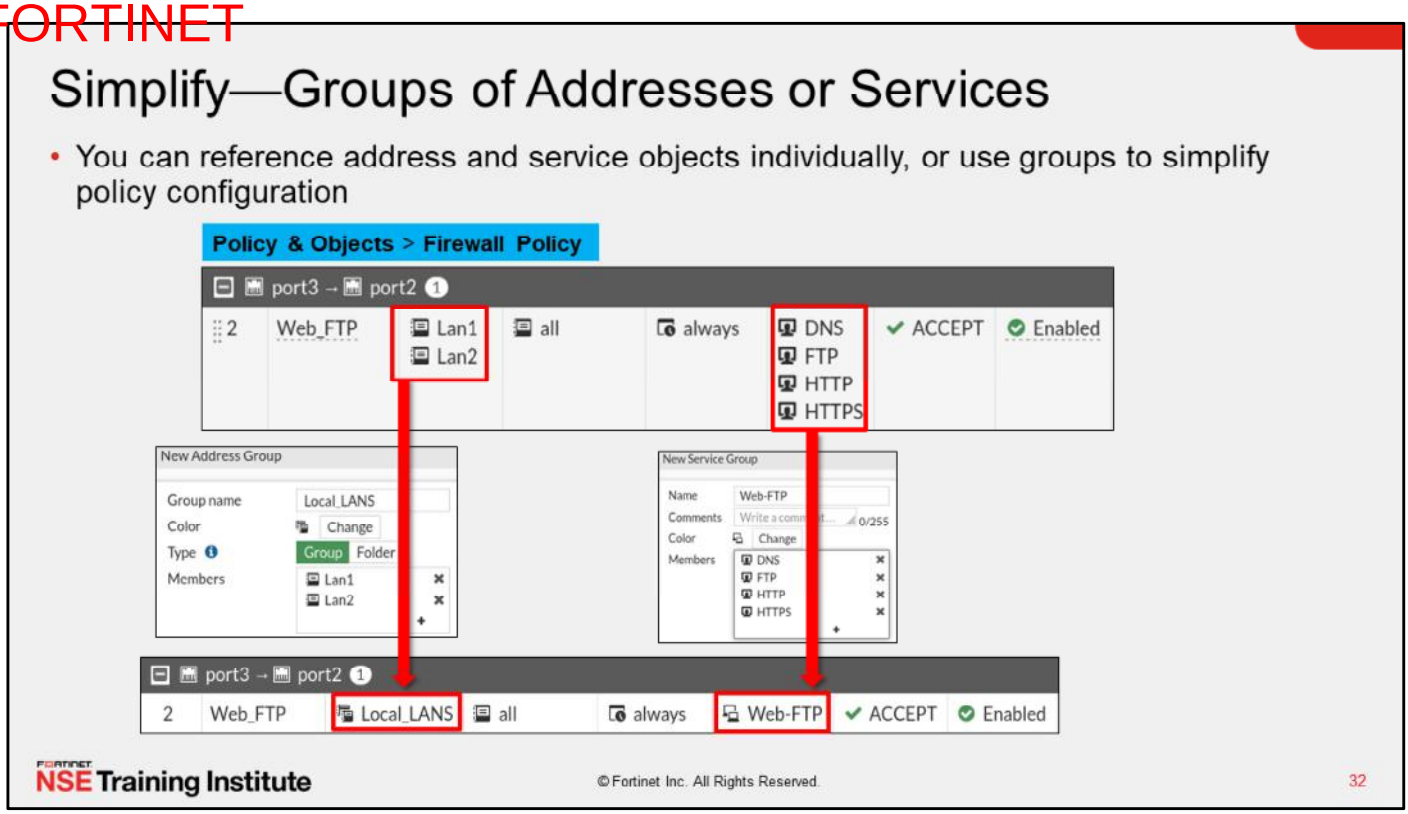

To simplify administration, you can group service and address objects. Then, you can reference that group in the firewall policy, instead of selecting multiple objects each time, or making multiple policies.

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.

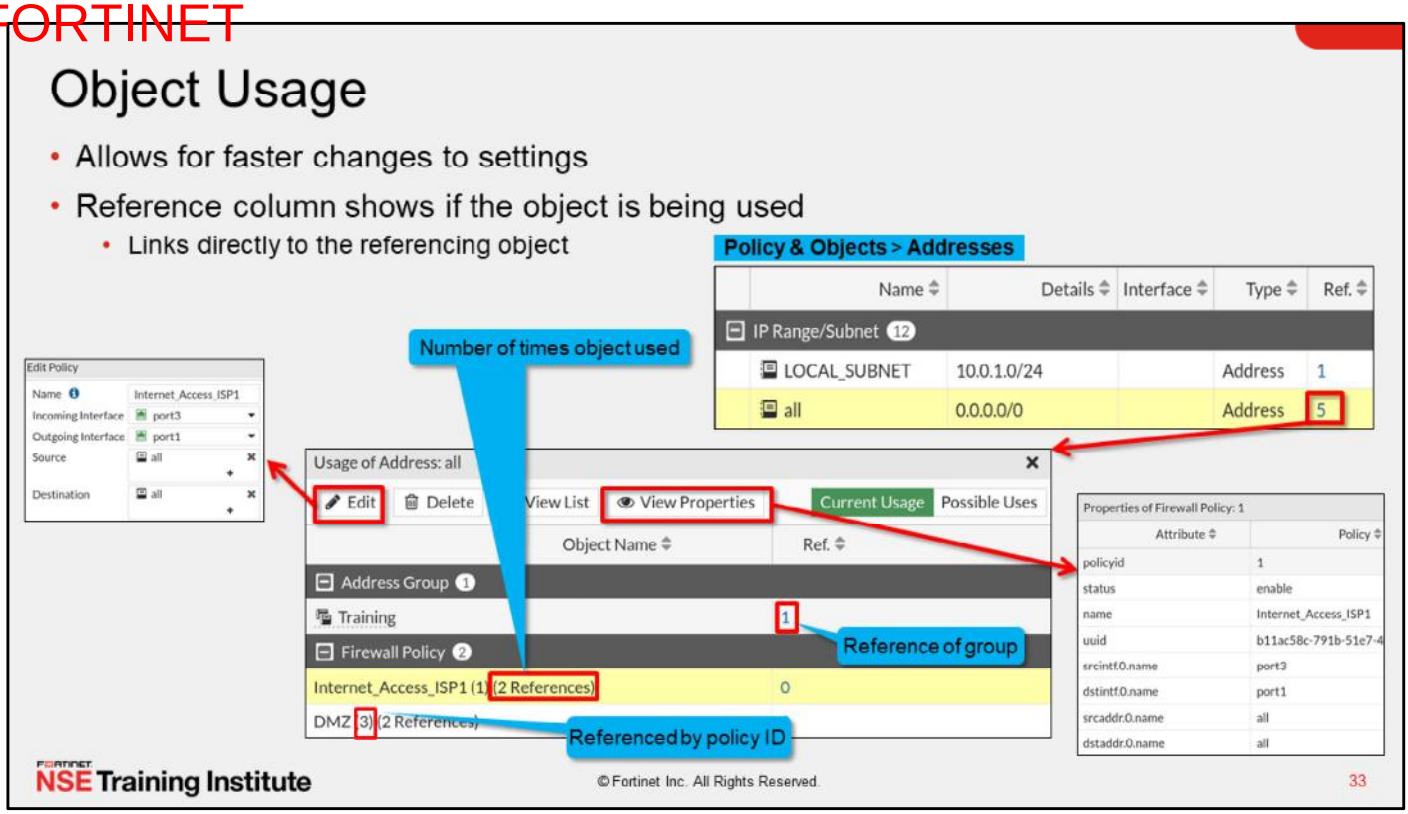

You've just seen several component objects that can be reused as you make policies. What if you want to delete an object?

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.

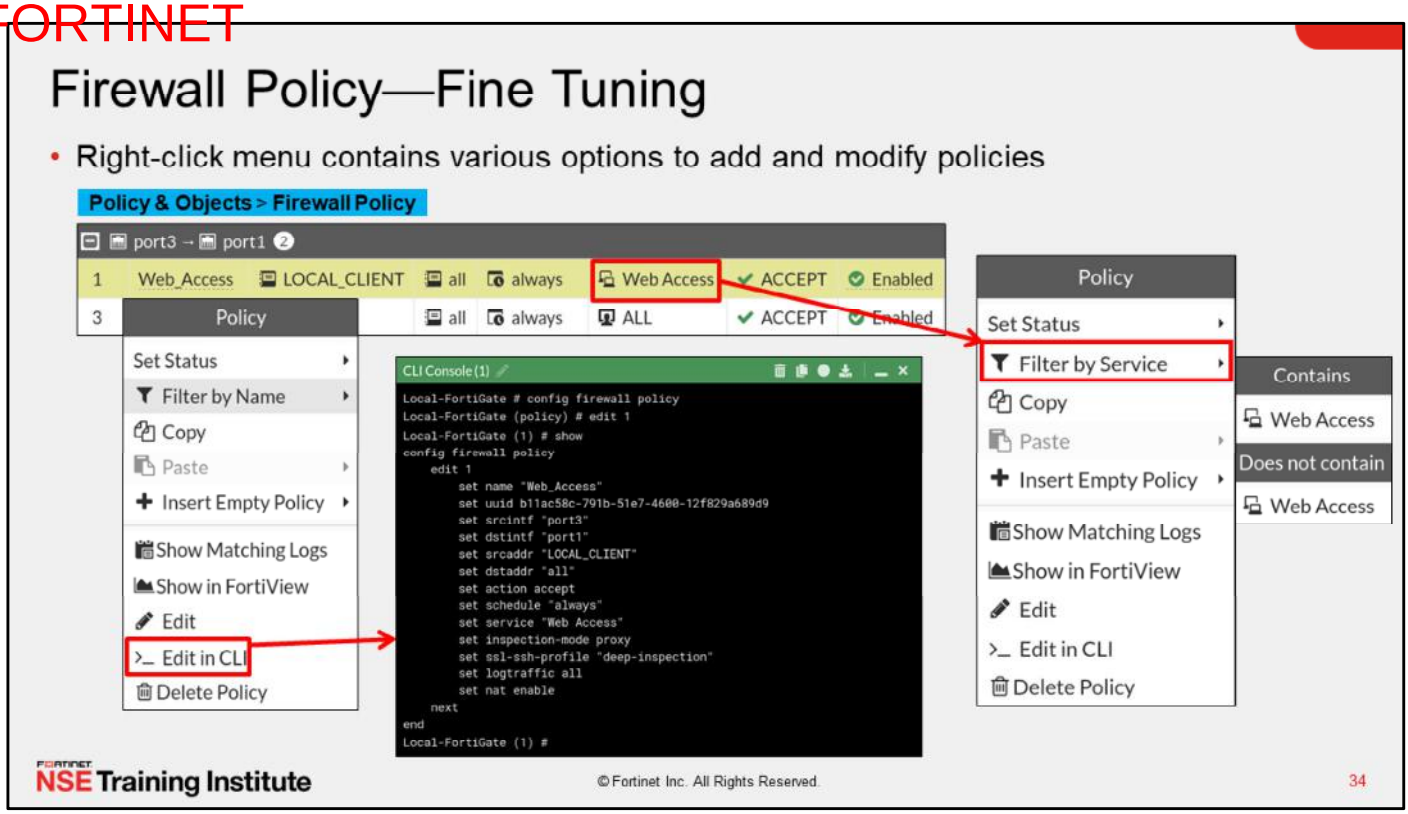

You can right-click any firewall policy to see different menu options to edit or modify the policy. The options include enabling or disabling a firewall policy, inserting firewall policies (above or below), copying and pasting policies, and cloning reverse (only if NAT is disabled on that policy).

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**.

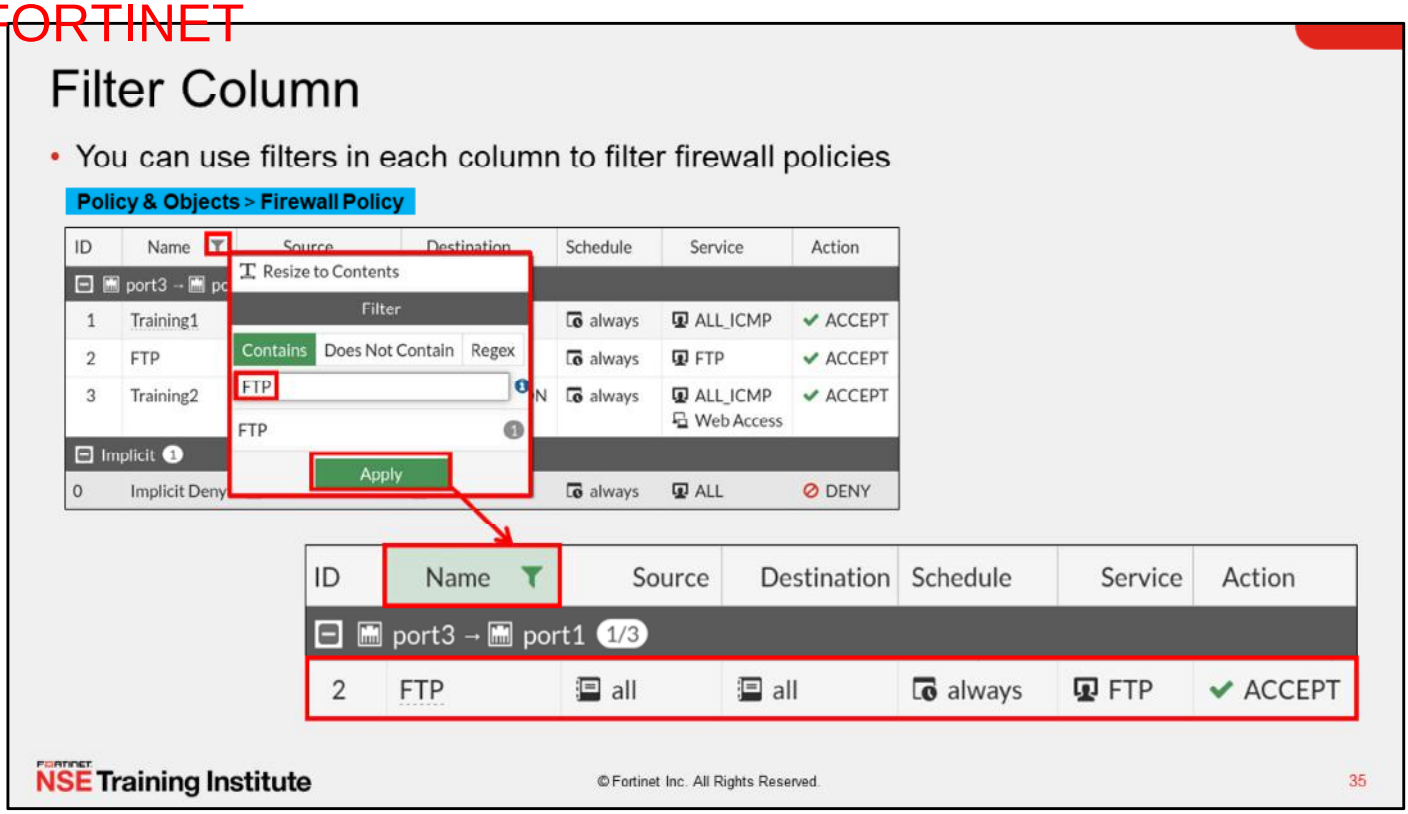

You can filter firewall policies on the GUI using filters in each column. You can add the **ID** column and then click the **ID** column filter icon to filter and search policies based on policy id numbers. You can click the **Name** filter icon to search policies based on policy name, and so on.

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**

NSE Training Institute

36

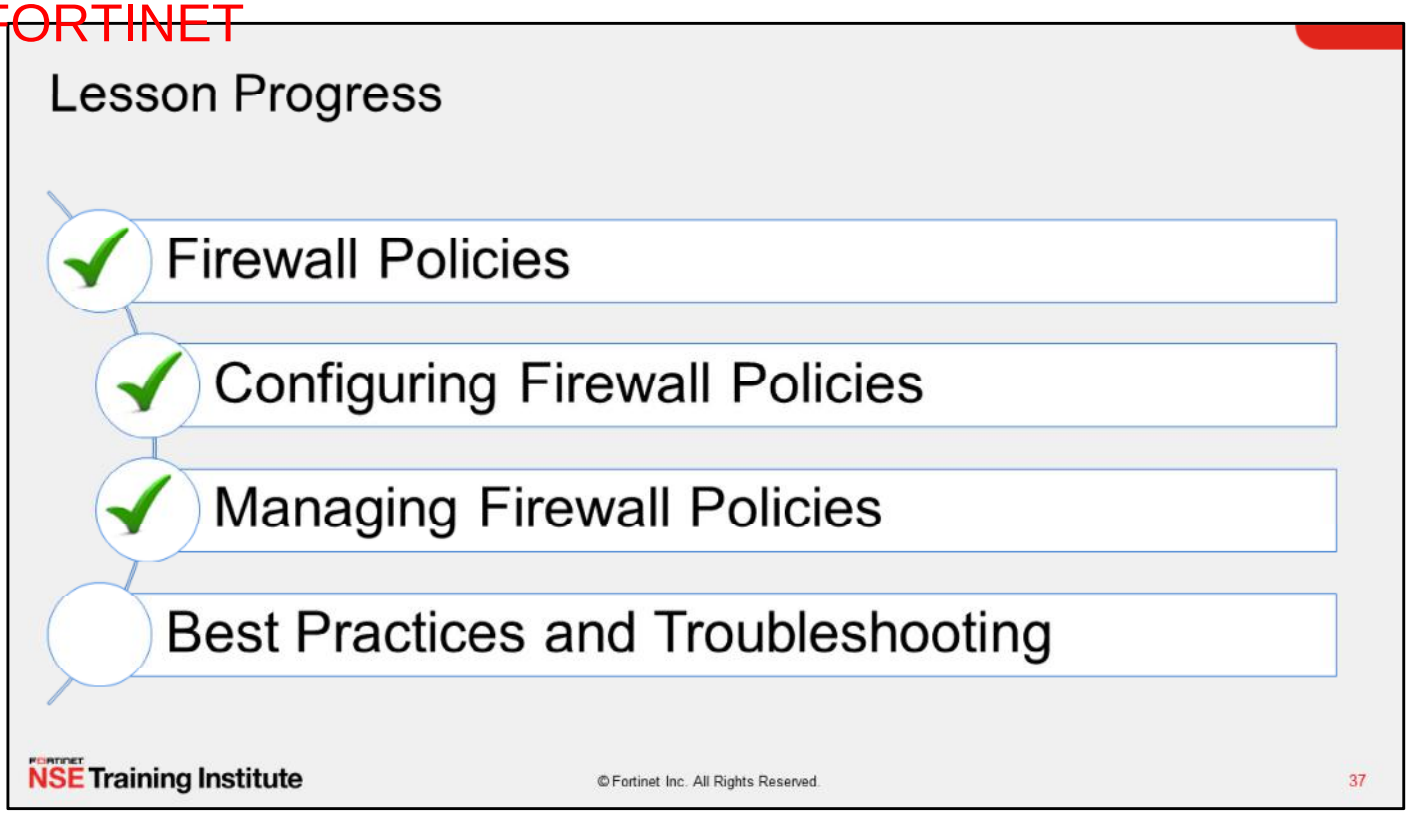

Good job! You now understand how to manage firewall policies on FortiGate.

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

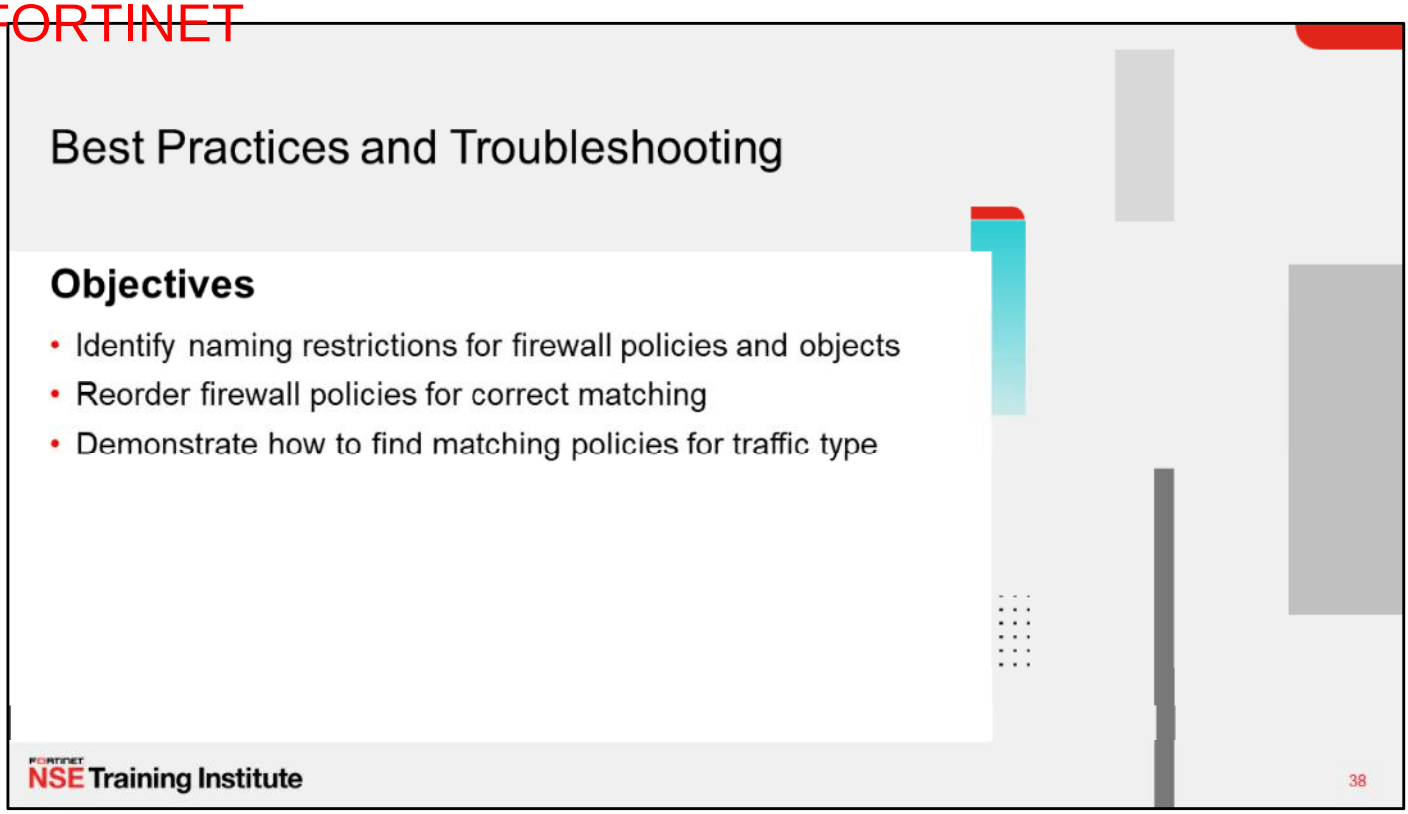

After completing this section, you should be able to achieve the objectives shown on this slide.

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
  - < > ( ) # " " ' '

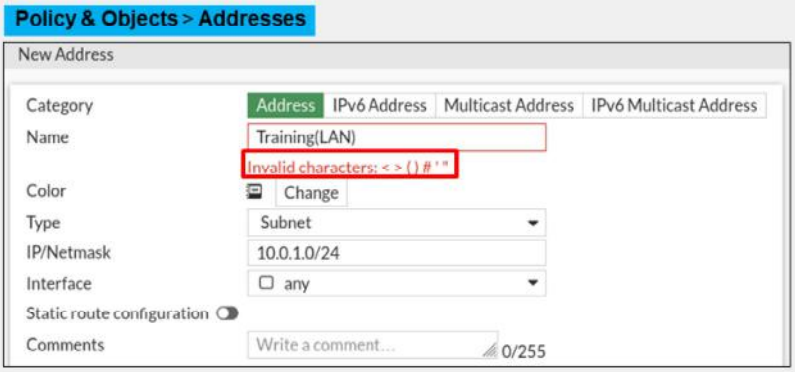

NSE Training Institute

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.

Best Practices

- Test policies in a maintenance window before deploying in production
  - Test policy for a few IP addresses, users, and so on
- Be careful when editing, disabling, or deleting firewall policies and objects
  - Changes are saved and activated immediately
  - Resets active sessions
- Create firewall policies to match as specifically as possible
  - Example: Restrict firewall policies based on source, destination, service
  - Use proper subnetting for address objects
- Analyze and enable appropriate settings on a per-policy basis
  - Security profiles
  - Logging settings

NSE Training Institute © Fortinet Inc. All Rights Reserved. 40

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.

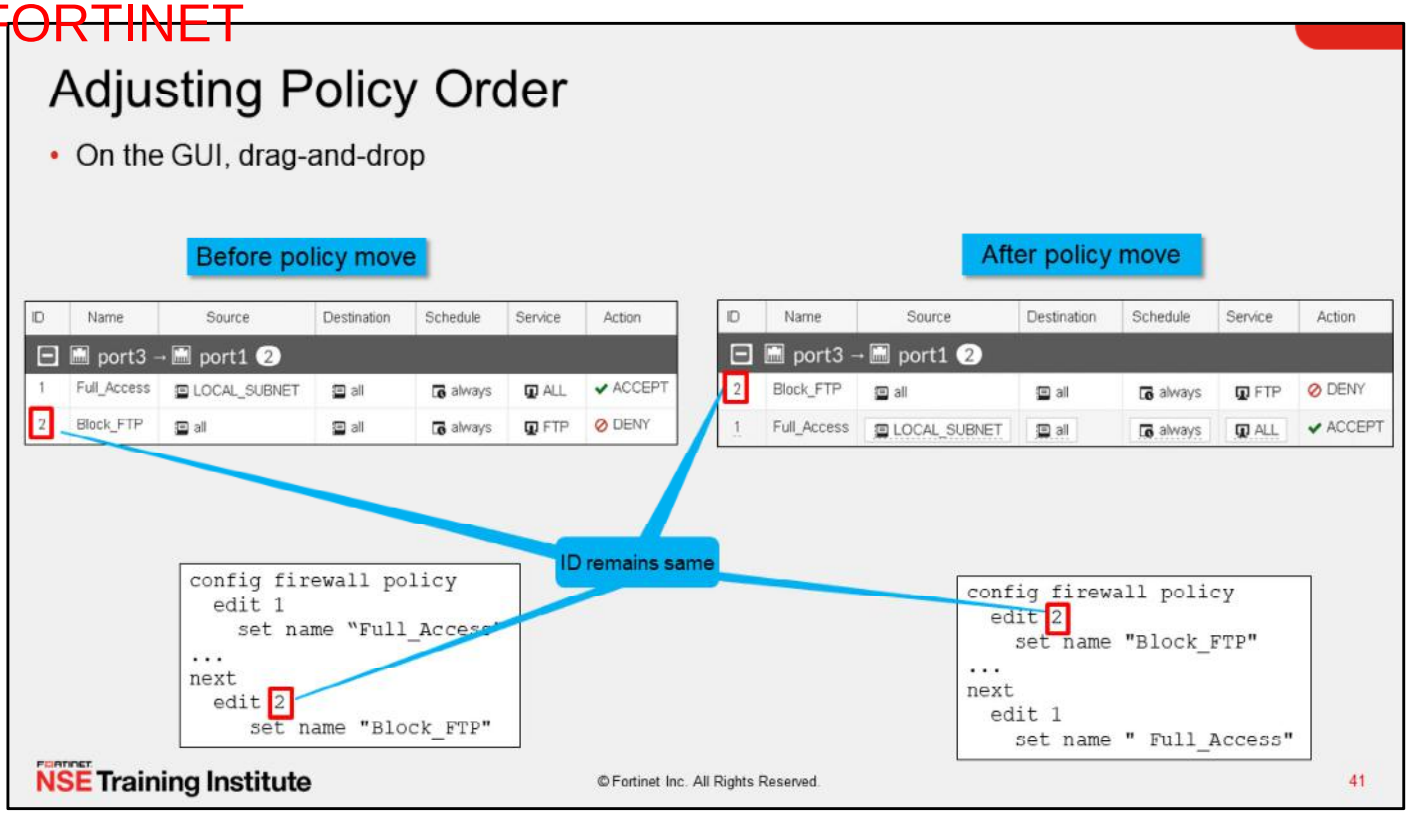

Remember you learned that only the first matching policy applies? Arranging your policies in the *correct position* is important. It affects which traffic is blocked or allowed. In the section of the applicable interface pair, FortiGate looks for a matching policy, beginning at the top. So, you should 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.

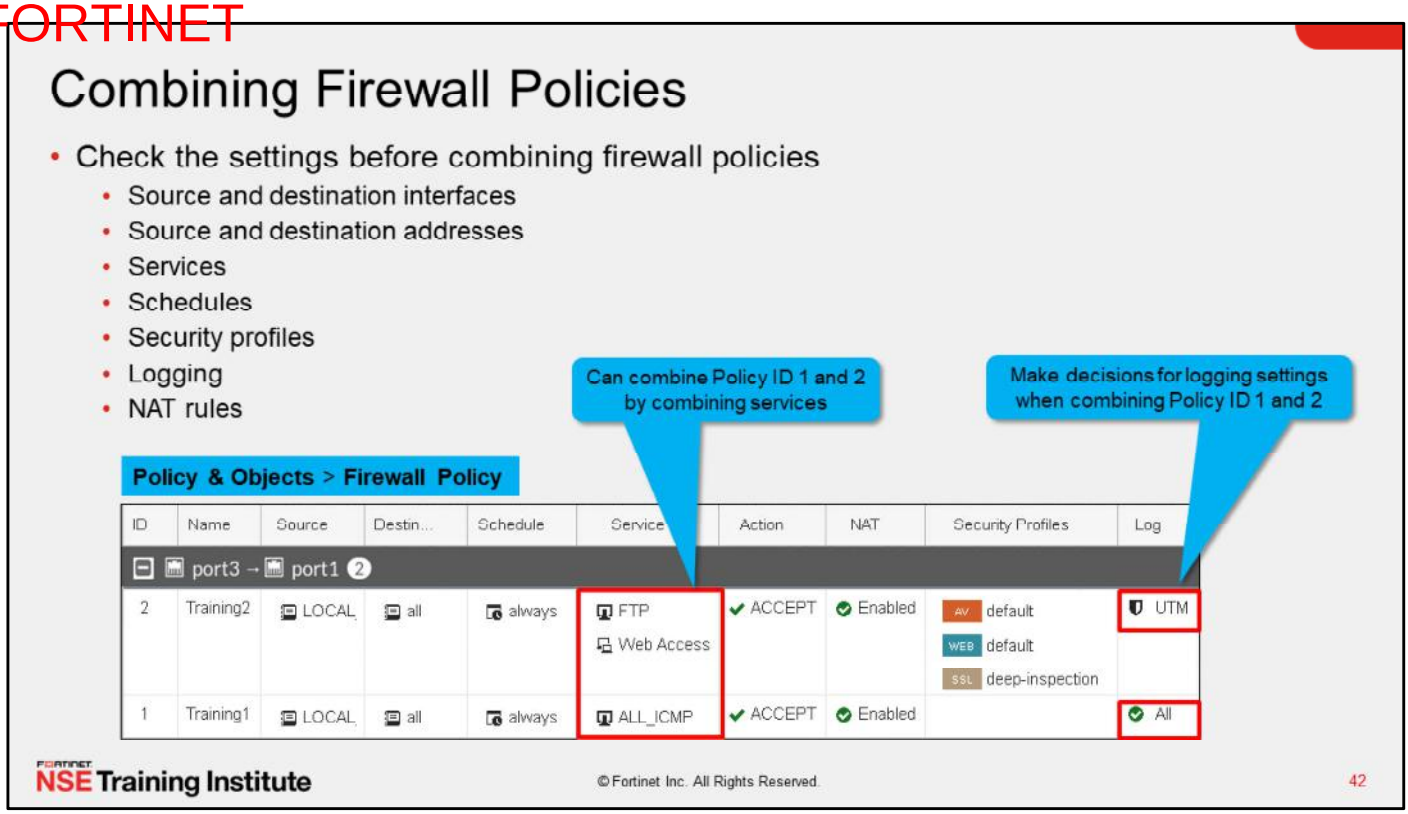

In order to optimize and consolidate firewall policies, always check all configured settings. In the example shown on this slide, the two firewall policies have differences in terms of services, security profiles, and logging settings. You can consolidate these two firewall policies by combining services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

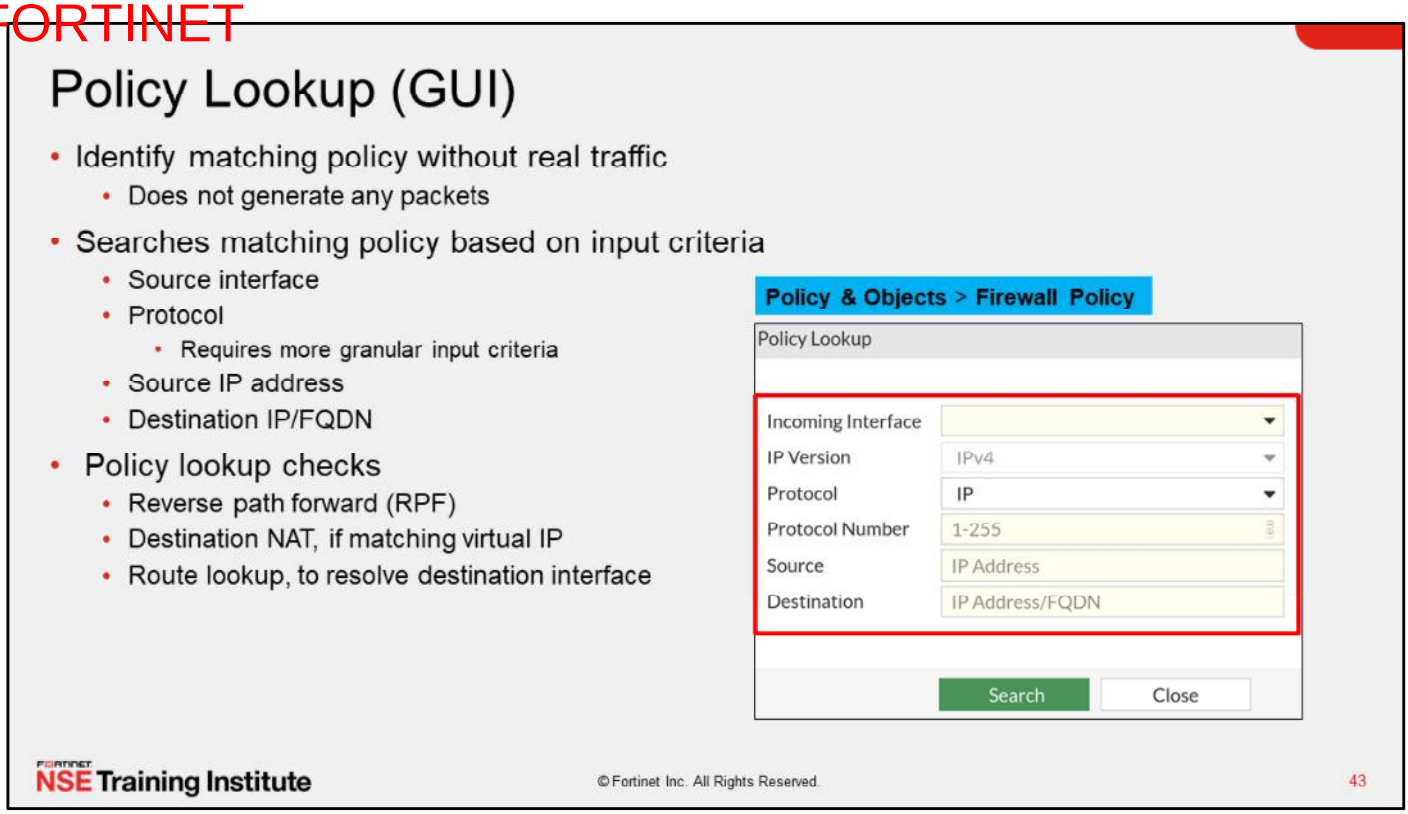

You can find a matching firewall policy based on the policy lookup input criteria. Policy lookup creates a packet flow over FortiGate without real traffic. From this, policy lookup can extract a policy ID from the flow trace and highlight it on the GUI policy configuration page.

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.

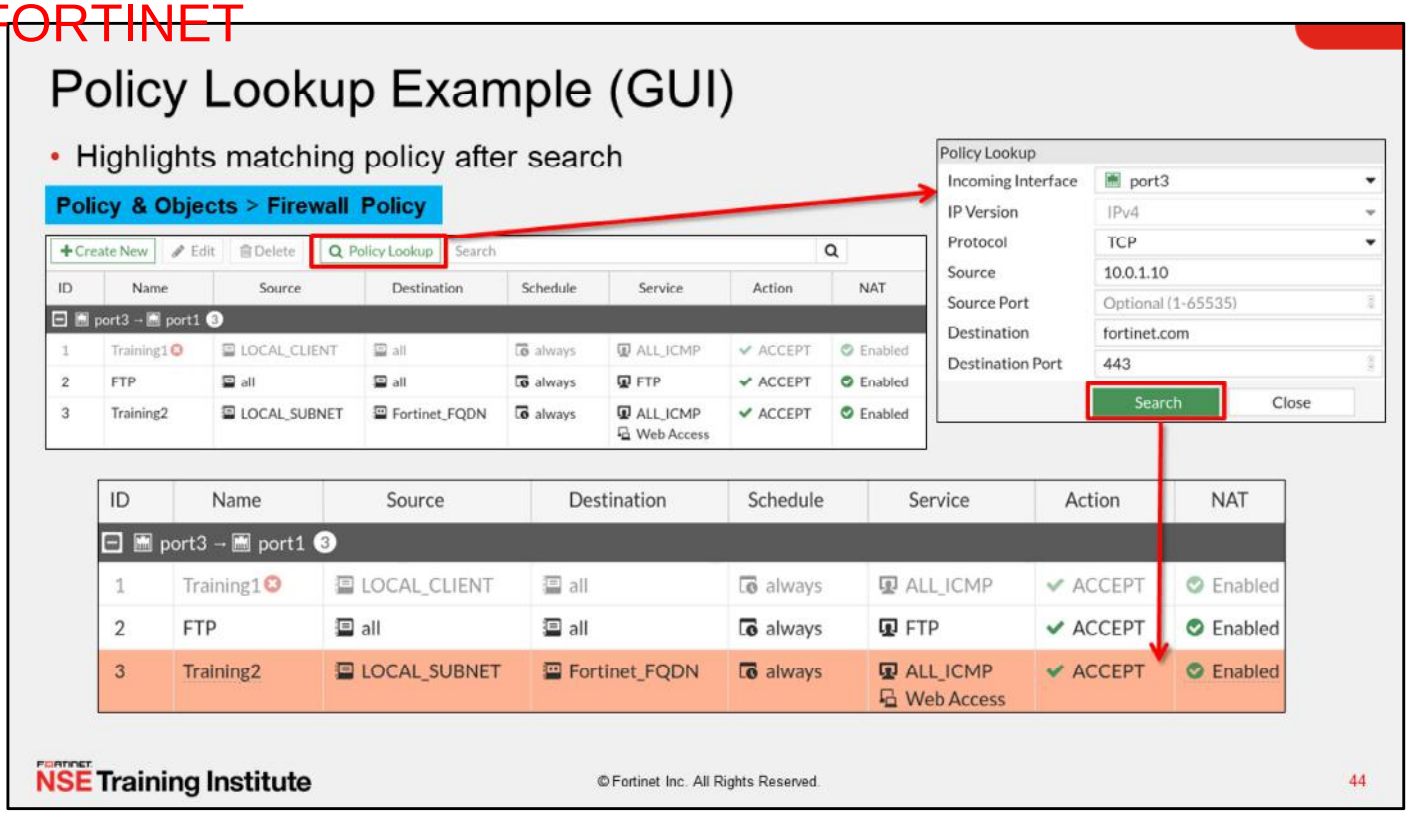

Based on the input criteria, after clicking **Search**, the trace result is selected and highlighted on the **Firewall Policy** page.

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

NSE Training Institute

45

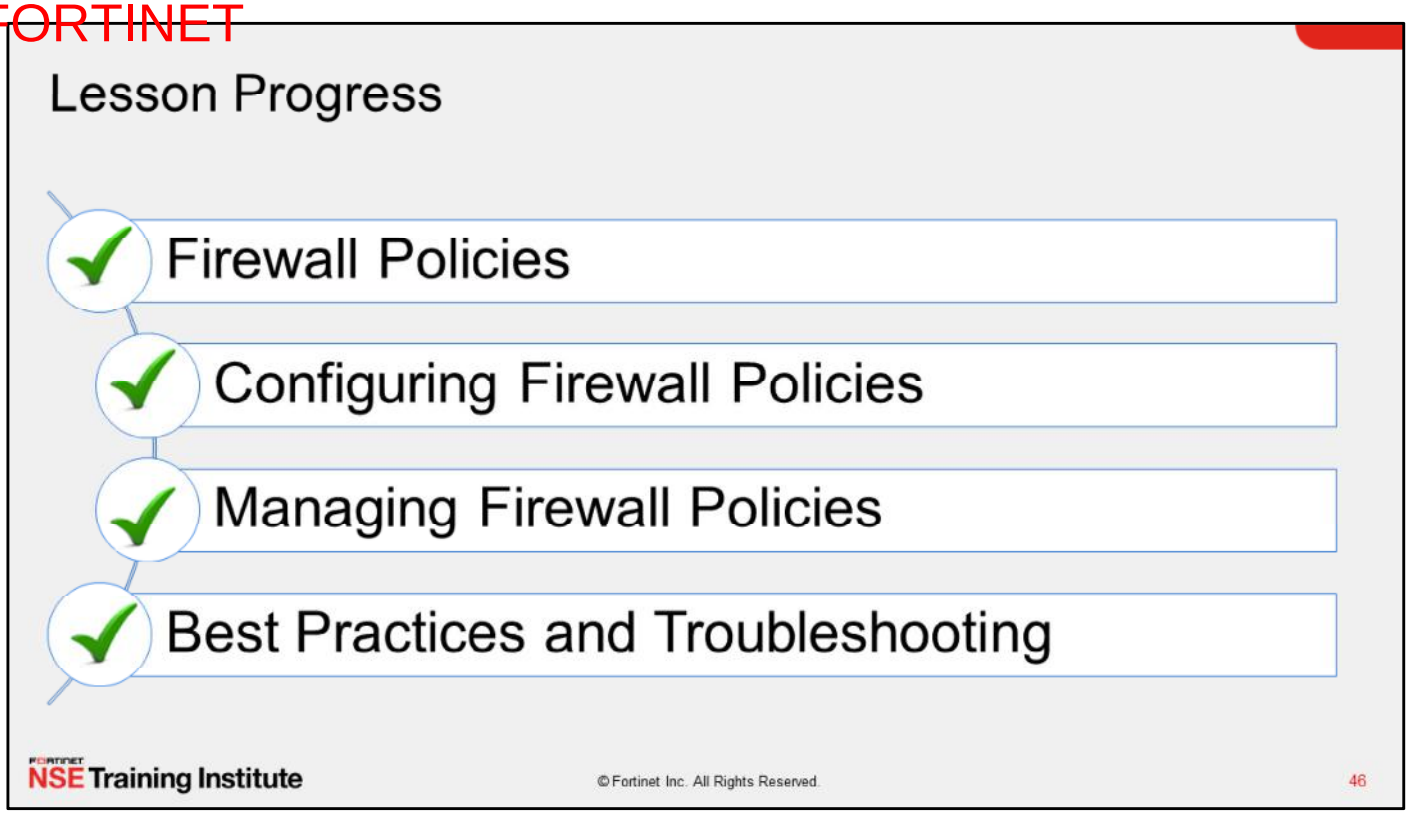

Congratulations! You have completed this lesson.

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

Review

- ✓ Identify components of firewall policies
- ✓ Identify how FortiGate matches traffic to firewall policies
- ✓ Restrict access and make your network more secure using security profiles
- ✓ Configure logging
- ✓ Identify policy list views
- ✓ Understand the use of policy IDs
- ✓ Identify where an object is referenced
- ✓ Identify naming restrictions for firewall policies and objects
- ✓ Reorder firewall policies for correct matching
- ✓ Demonstrate how to find matching policies for traffic type

NSE Training Institute © Fortinet Inc. All Rights Reserved. 47

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.

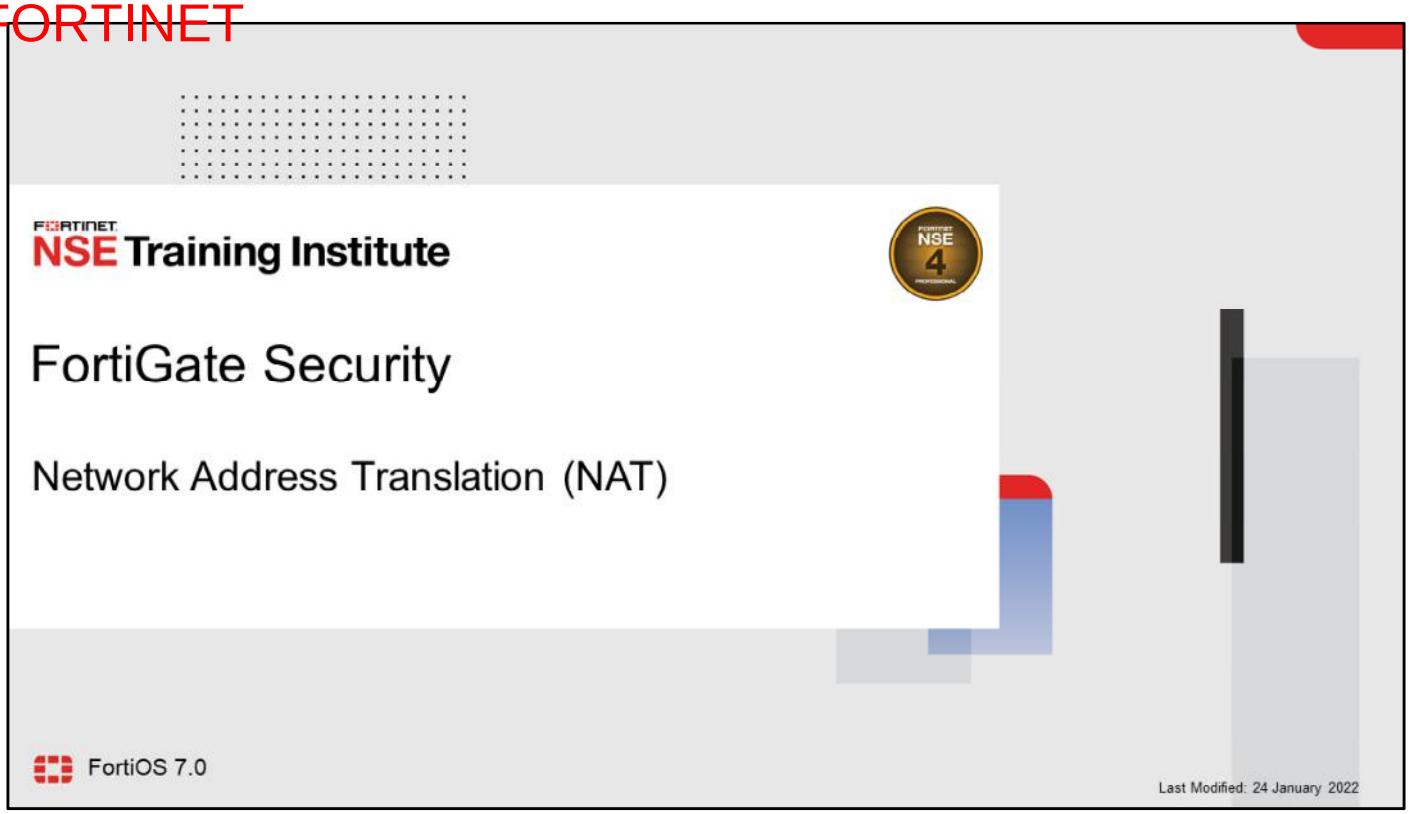

In this lesson, you will learn how to configure network address translation (NAT) and use it to implement source NAT and destination NAT for the traffic passing through FortiGate.

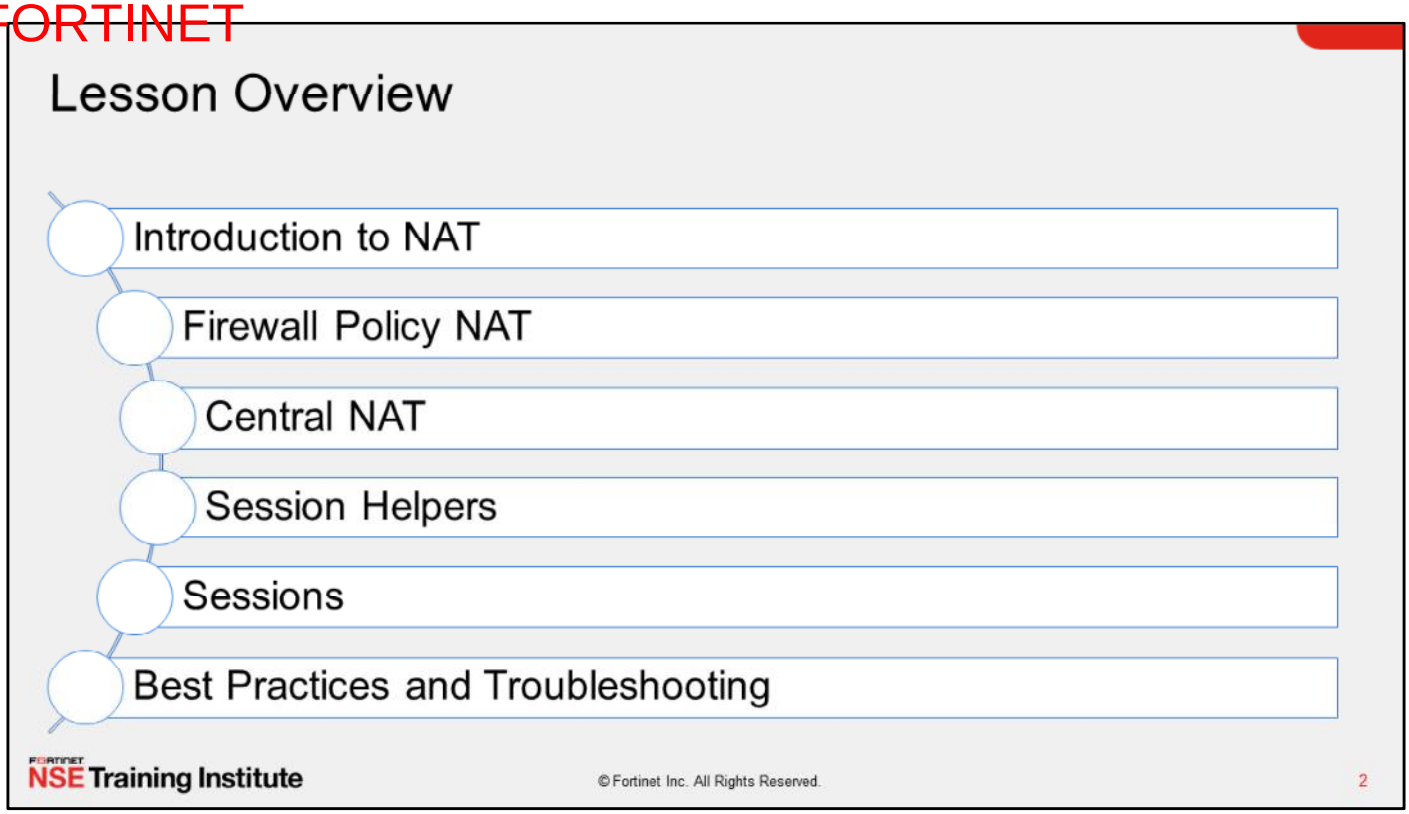

In this lesson, you will learn about the topics shown on this slide.

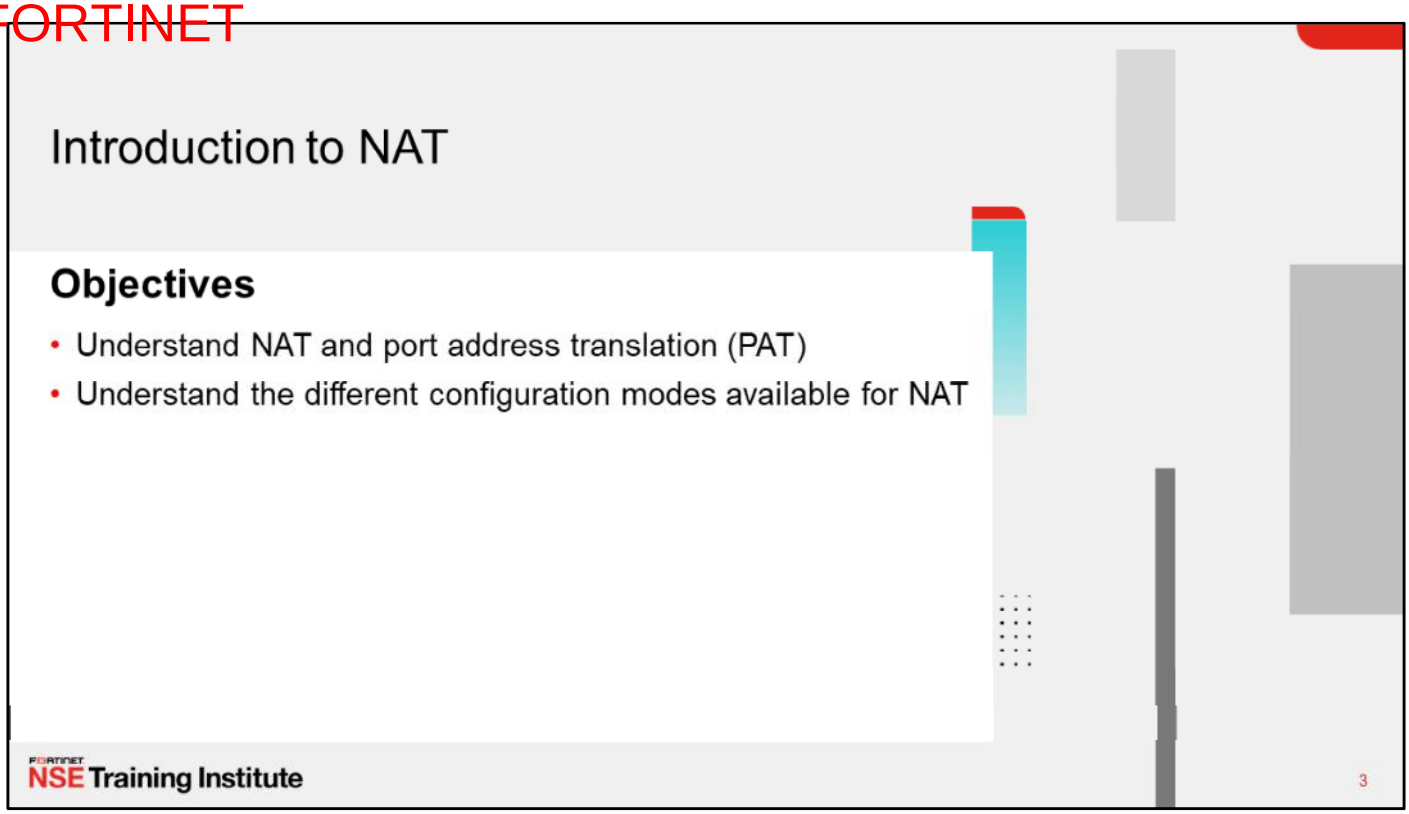

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding how NAT and PAT work, and the available NAT configuration modes, you will have a good start for planning the implementation of NAT in your network.

NAT and PAT

- NAT
  - Changes the IP layer address of a packet
    - Some protocols, like SIP, have addresses at the application layer, requiring session helpers or proxies
  - Source NAT (SNAT)
  - Destination NAT (DNAT)
- PAT (NAT overload)
  - Map multiple private IPv4 addresses to a single public IP address by using different source ports
- NAT64 and NAT46
  - A mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and the reverse
- NAT66
  - NAT between two IPv6 networks

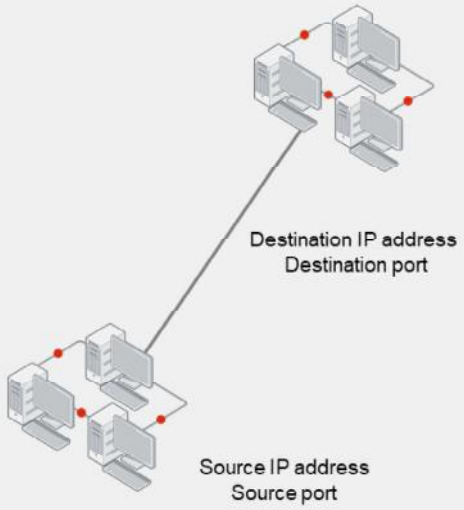

NSE Training Institute

4

NAT is the process that enables a single device, such as a firewall or router, to act as an agent between the internet, or public network, and a local, or private, network.

NAT is usually implemented for one, or a combination, of the following reasons:
- Improved security: The addresses behind the NAT device are virtually hidden.
- Amplification of addresses: Hundreds of computers can use as few as one public IP address.
- Internal address stability: The addresses can stay the same, even if internet service providers (ISPs) change.

NAT and PAT, also known as NAPT, translate internal, typically private, IP addresses to external, typically public or internet, IP addresses. In FortiOS, NAT and traffic forwarding apply to the same firewall policy. However, diagnostics clearly show NAT and forwarding as separate actions.

- For outgoing connections: you can use the NAT option in a central SNAT, IP pool, and central SNAT table, which is known as *source NAT*.
- For incoming connections: you can use virtual IPs (VIPs) and DNAT, which are known as *destination NAT*.

NAT64 and NAT46 are the terms used to refer to the mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and the reverse. Without this mechanism, an IPv6 node on a network, such as a corporate LAN, would not be able to communicate with a website that was in an IPv4-only environment, and IPv4 environments would not be able to connect to IPv6 networks.

NAT66 is NAT between two IPv6 networks.

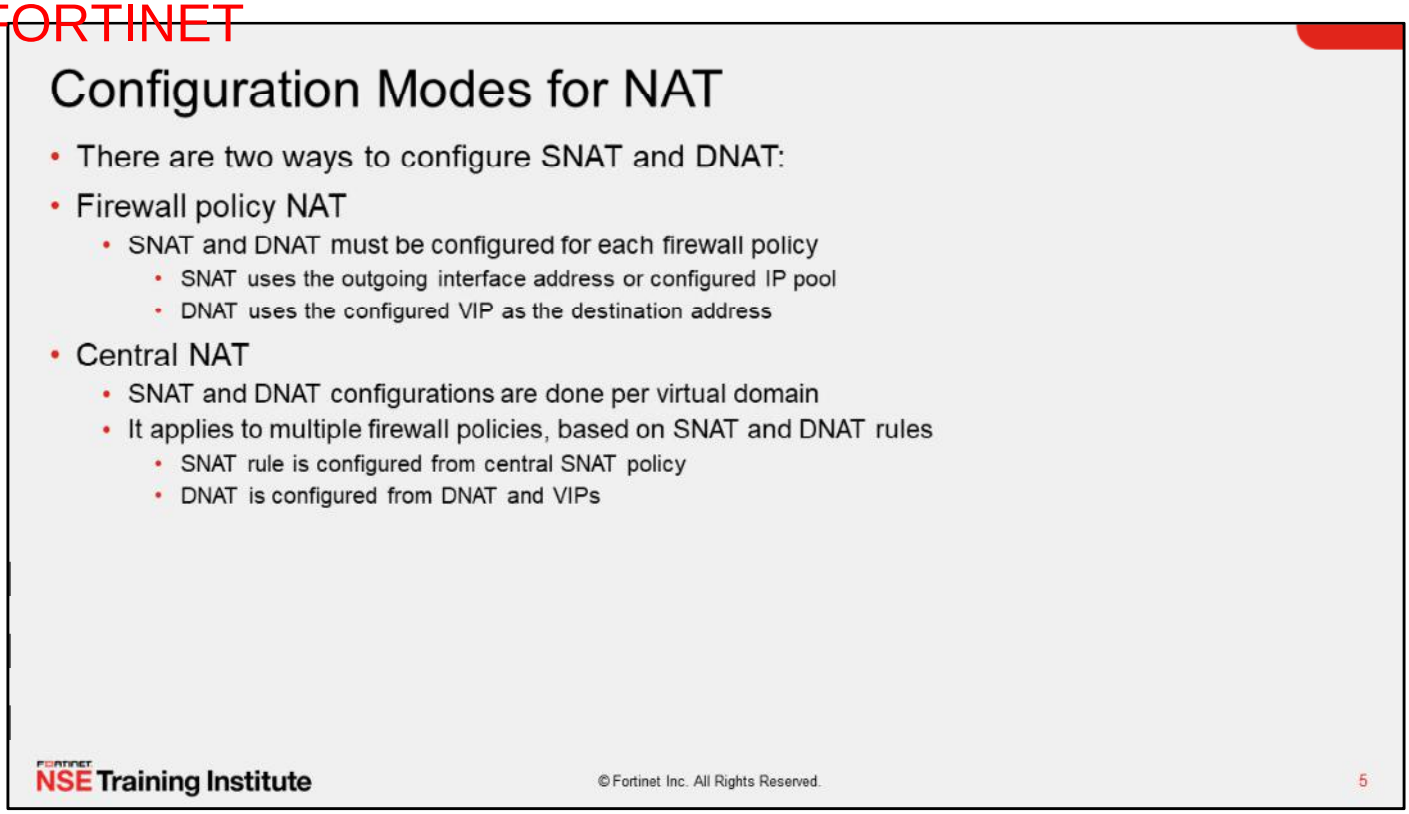

When you use firewall policy NAT mode, you must configure SNAT and DNAT for each firewall policy.

Central NAT configurations are done per virtual domain, which means SNAT and DNAT configurations automatically apply to multiple firewall policies. This is according to the SNAT and DNAT rules that you specify, 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.

Knowledge Check

1. What is NAT used for?
   - ✓ A. Preserving IP addresses
   - B. Traffic shaping

2. Which statement about NAT66 is true?
   - ✓ A. It is NAT between two IPv6 networks.
   - B. It is NAT between two IPv4 networks.

NSE Training Institute

6

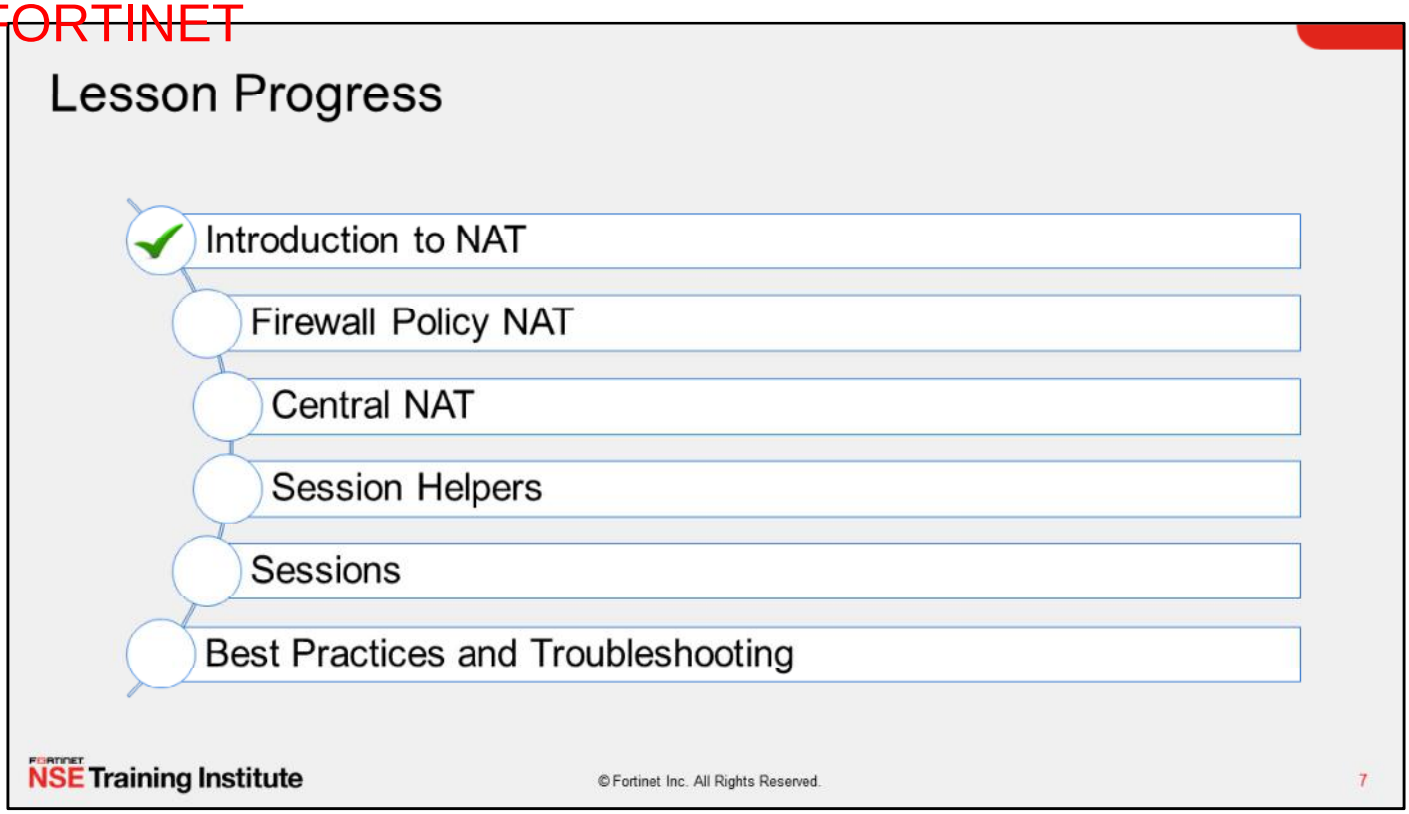

Good job! You now know about NAT.

Now, you'll learn about firewall policy NAT.

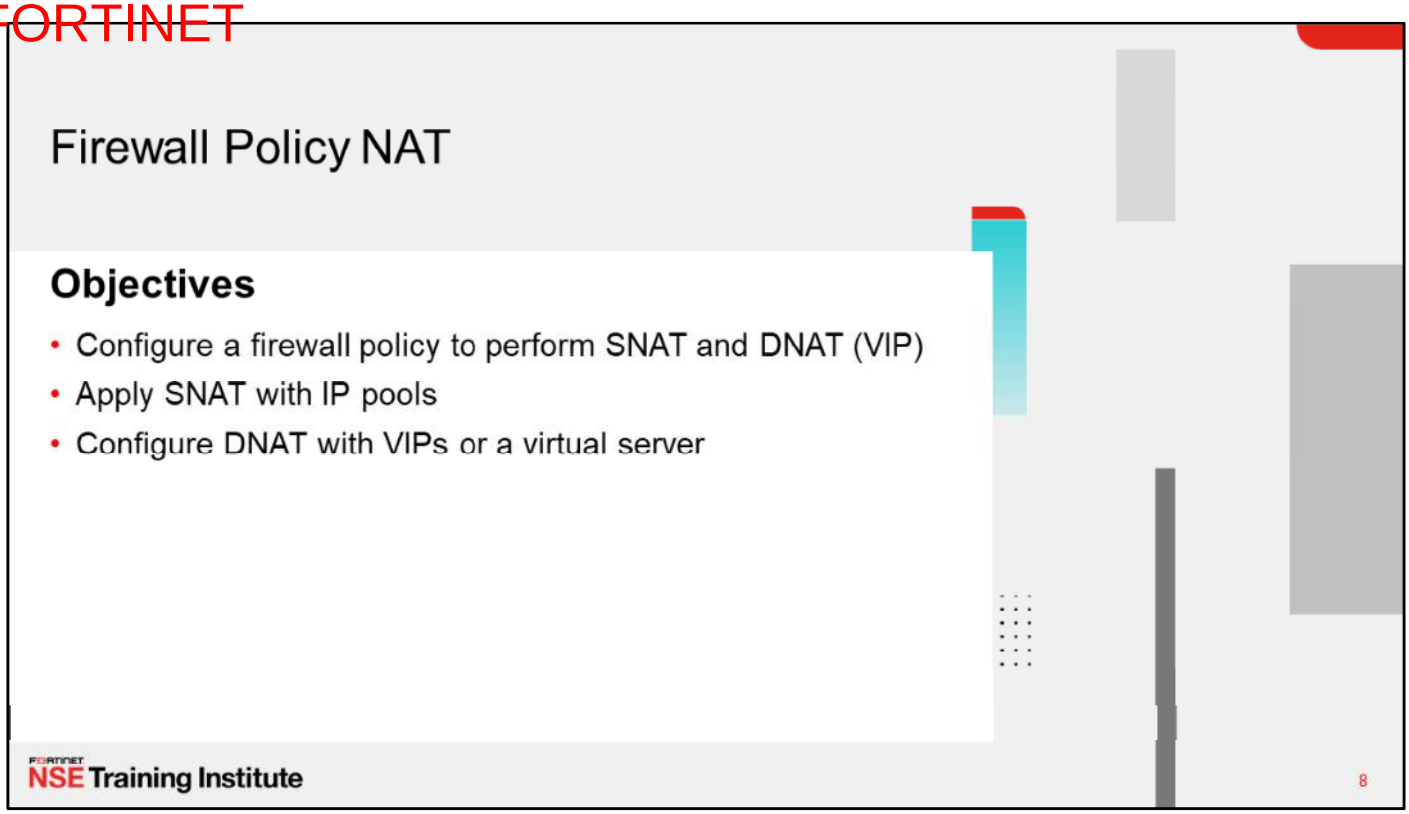

After completing this section, you should be able to achieve the objectives shown on this slide.

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 two ways to configure firewall policy SNAT:

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

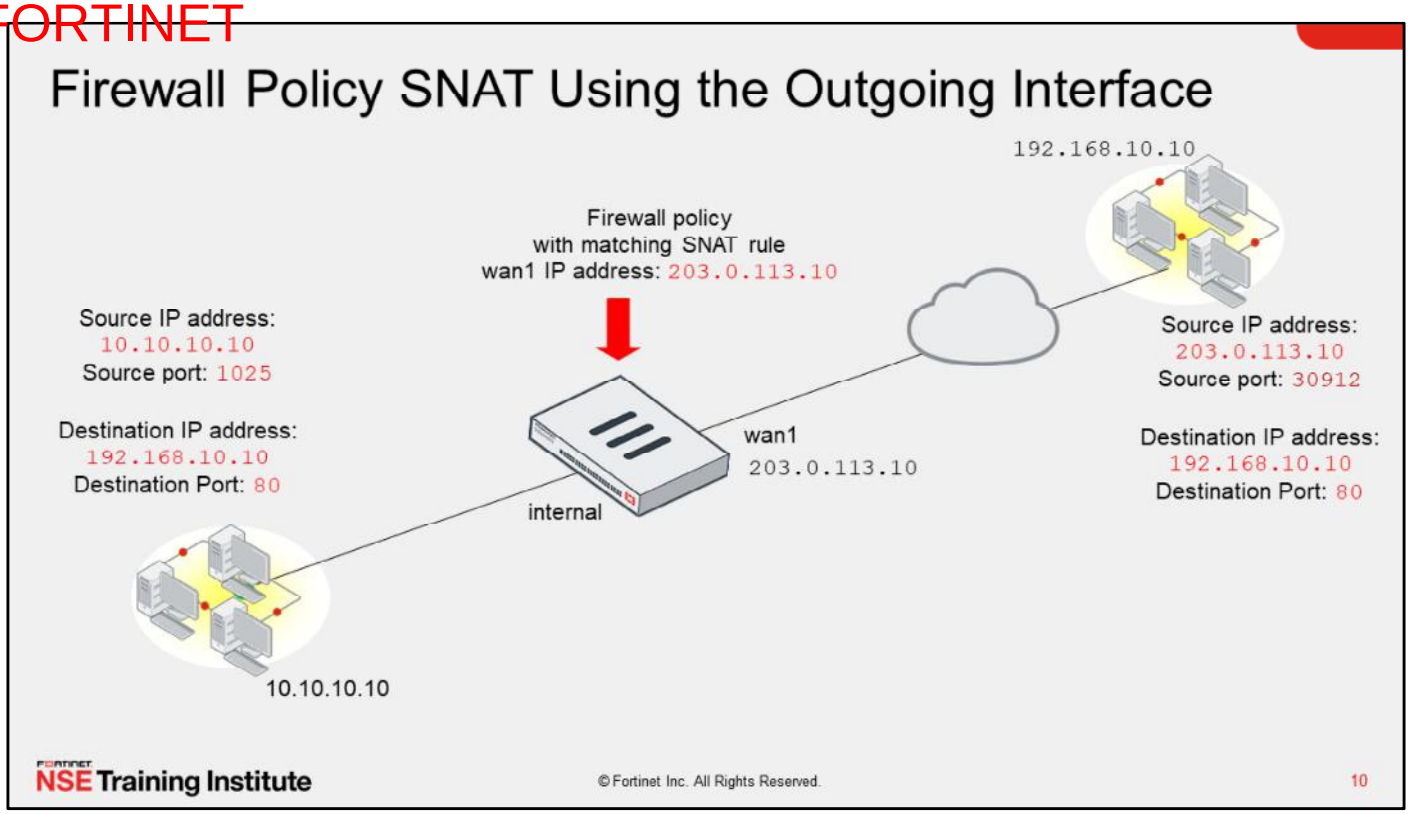

The source NAT option uses the egress interface address when NAT is enabled on the firewall policy. This is many-to-one NAT. In other words, PAT is used, and connections are tracked using the original source address and source port combinations, as well as the allocated source port. This is the same 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 can establish.

In the example shown on this slide, a firewall policy from internal to wan1 (IP address `203.0.113.10`) is created, and the user initiates traffic from source `10.10.10.10:1025` destined for `192.168.10.10:80`. Because NAT is enabled on the firewall policy, the source IP address is translated to the egress interface IP, with port translation.

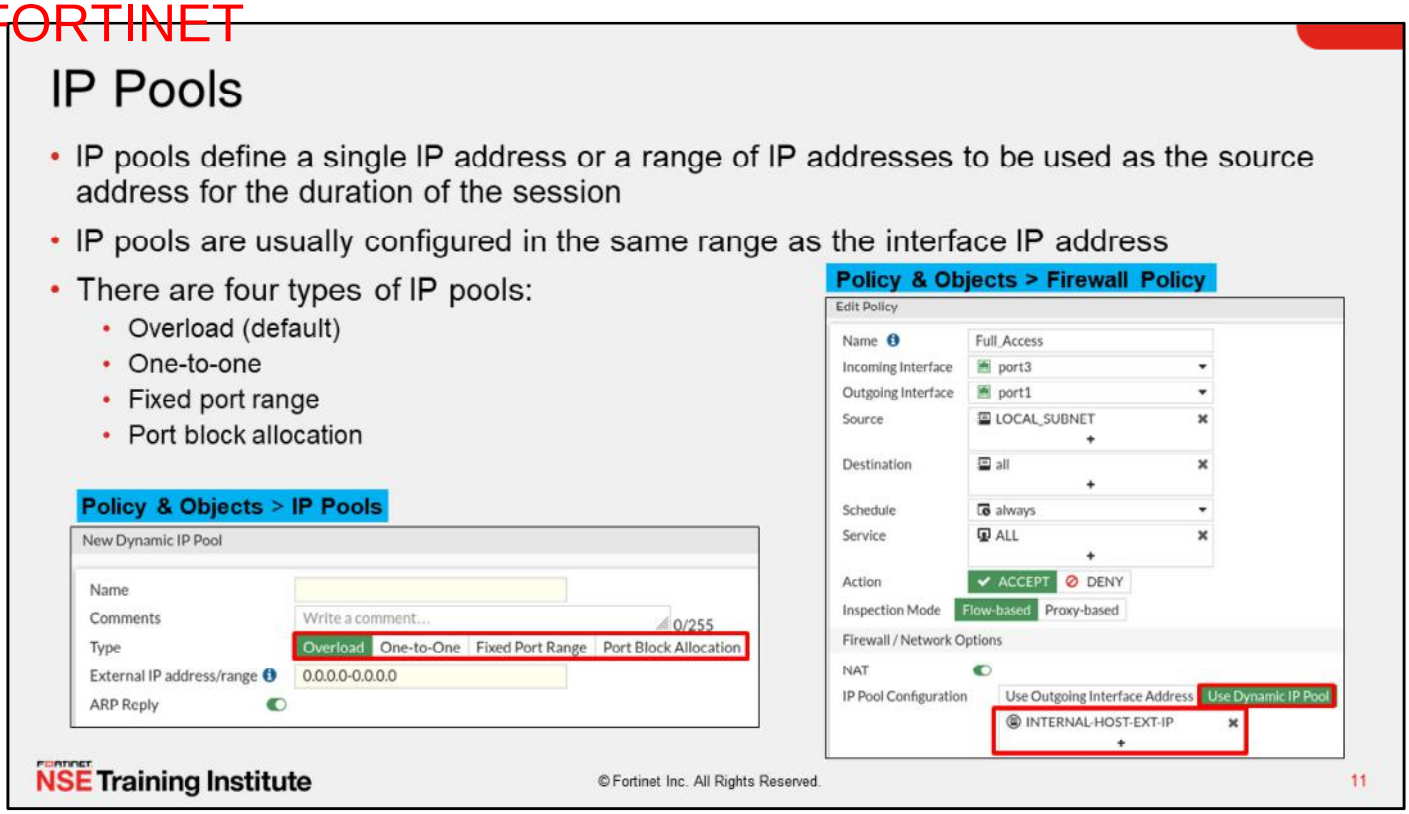

IP pools are a mechanism that allow sessions leaving the FortiGate firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses are used instead of the IP address assigned to that FortiGate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

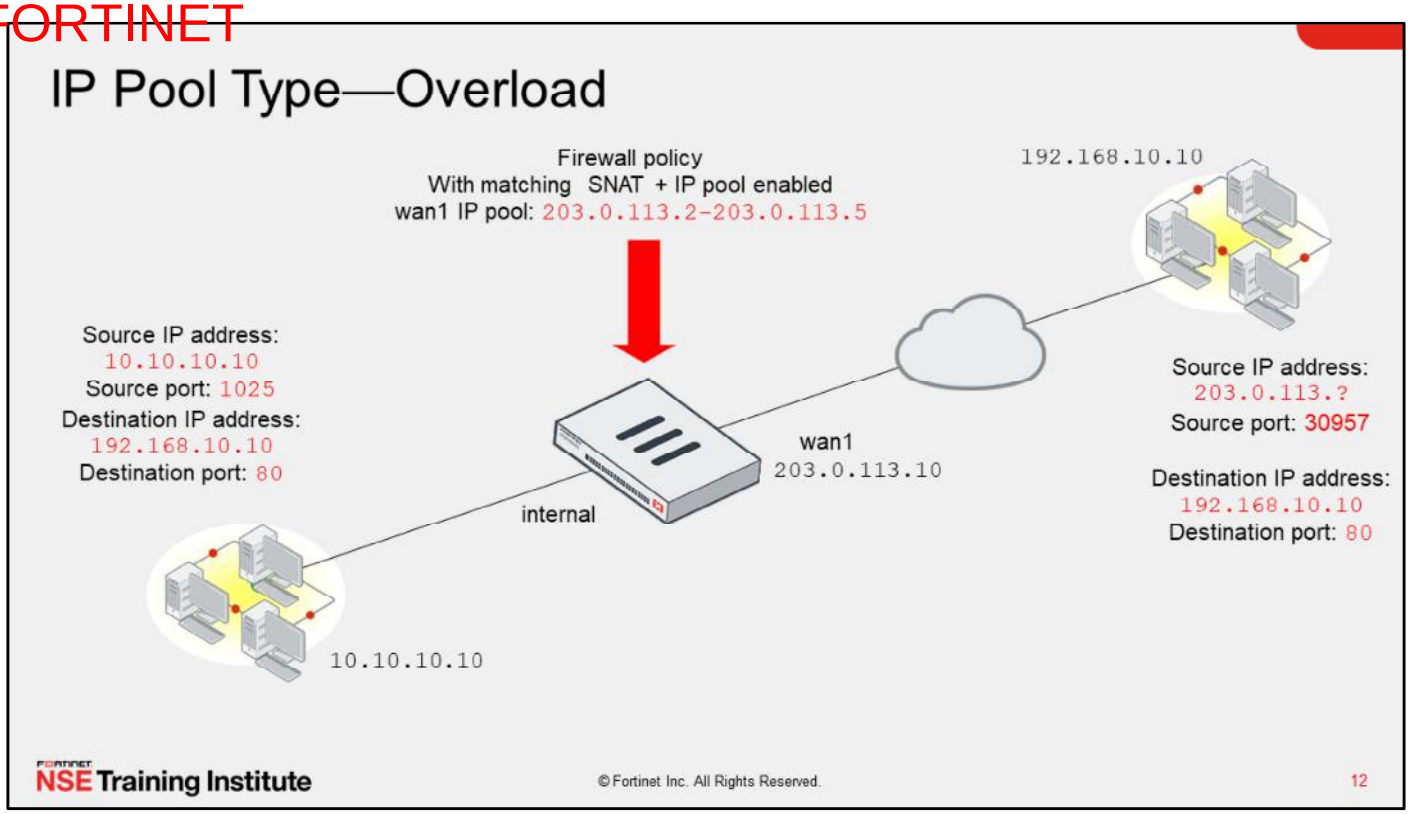

If you use an IP pool, the source address is translated to an address from that pool, rather than the egress interface address. The larger the number of addresses in the pool, the greater the number of connections that can be supported. For example, in an enterprise network where you require a greater number of connections, or in a network where you want one subnet to use one specific public IP over another to restrict access based on source IP address.

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 `10.10.10.10` is translated to an IP address from the IP pool (`203.0.113.2 - 203.0.113.5`).

IP Pool Type—One-to-One

- The IP pool type one-to-one associates an internal IP with a pool IP on a first-come, first-served basis
  - PAT is disabled

```
STUDENT # get system session list
PROTO   EXPIRE SOURCE          SOURCE-NAT       DESTINATION
DESTINATION-NAT
tcp     3598   10.0.1.10:2706  10.200.1.6:2706  10.200.1.254:80  -
tcp     3598   10.0.1.10:2704  10.200.1.6:2704  10.200.1.254:80  -
tcp     3596   10.0.1.10:2702  10.200.1.6:2702  10.200.1.254:80  -
tcp     3599   10.0.1.10:2700  10.200.1.6:2700  10.200.1.254:443 -
tcp     3599   10.0.1.10:2698  10.200.1.6:2698  10.200.1.254:80  -
tcp     3598   10.0.1.10:2696  10.200.1.6:2696  10.200.1.254:443 -
udp     174    10.0.1.10:2694  -                  10.0.1.254:53  -
udp     173    10.0.1.10:2690  -                  10.0.1.254:53  -
```

- Refuses the connection if there is no unallocated address

NSE Training Institute

13

In the one-to-one pool type, an internal IP address is mapped with an external address on a first-come, first-served basis.

There is a single mapping of an internal address to an external address. Mappings are not fixed and, if there are no more addresses available, a connection will be refused.

Also, in one-to-one, PAT is not required. In the example on this slide, you can see the same source port is shown for both the ingress and egress address.

IP Pool Type—Fixed Port Range

- The fixed port range IP pool type associates an internal IP range with an external IP range
  - A type of PAT

```
STUDENT # get system session list
PROTO   EXPIRE SOURCE           SOURCE-NAT        DESTINATION       DESTINATION-NAT
tcp     3574   10.0.1.11:60843  10.200.1.8:60843  216.23.154.83:80  -
tcp     3570   10.0.1.11:60809  10.200.1.8:60809  216.23.154.81:80  -
tcp     3590   10.0.1.11:60819  10.200.1.8:60819  216.23.154.74:80  -
tcp     3599   10.0.1.11:60817  10.200.1.8:60817  216.23.154.74:80  -
tcp     3586   10.0.1.11:60815  10.200.1.8:60815  216.23.154.81:80  -
tcp     3564   10.0.1.11:60807  10.200.1.8:60807  216.23.154.74:80  -
tcp     9      10.0.1.10:7112   10.200.1.7:7112   10.200.1.254:80   -
tcp     7      10.0.1.10:7110   10.200.1.7:7110   10.200.1.254:80   -
tcp     5      10.0.1.10:7108   10.200.1.7:7108   10.200.1.254:80   -
tcp     3      10.0.1.10:7106   10.200.1.7:7106   10.200.1.254:80   -
tcp     1      10.0.1.10:7104   10.200.1.7:7104   10.200.1.254:80   -
```

NSE Training Institute

14

For the overload and one-to-one IP pool types, you do not need to define the internal IP range. For the fixed port range type of IP pool, you can define both the internal IP range and external IP range.

Because each external IP address and the number of available port numbers is a specific number, if the number of internal IP addresses is also determined, you can calculate the port range for each address translation combination. This type of IP pool is called fixed port range and is a type of port address translation (PAT).

The fixed port range allows fixed mapping of the internal start IP or internal end IP range to the external start IP or external end IP range.

The example on this slide shows a fixed port range IP pool. The internal address range `10.0.1.10 to 10.0.1.11` maps to the external address range `10.200.1.7 to 10.200.1.8`.

IP Pool Type—Port Block Allocation

- The port block allocation IP pool type assigns a block size and number per host for a range of external IP addresses
  - Using a small 64-block size and one block
    - hping --faster –p 80 –S 10.200.1.254

```
STUDENT # diagnose sys session stat
misc info:          session_count=79 setup_rate=0 exp_count=0 clash=0
          memory_tension_drop=0 ephemeral=0/65536 removeable=0
delete=0, flush=0, dev_down=0/0 ses_walkers=0
TCP sessions:
           2 in ESTABLISHED state
          74 in SYN_SENT state
           1 in CLOSE_WAIT state
```

  - Using an overload type
    - hping --faster –p 80 –S 10.200.1.254

```
STUDENT # diagnose sys session stat
misc info:          session_count=10227 setup_rate=982 exp_count=0 clash=0
          memory_tension_drop=0 ephemeral=0/65536 removeable=0
delete=0, flush=0, dev_down=0/0 ses_walkers=0
TCP sessions:
          34 in ESTABLISHED state
       10117 in SYN_SENT state
           1  in SYN_RECV state
```

NSE Training Institute

15

IP Pool type port block allocation is also a type of PAT. It gives users a more flexible way to control the way external IPs and ports are allocated.

Users need to define **Block Size** and **Block Per User** and the external IP range. **Block Size** means how many ports each block contains. **Block per User** means how many blocks each host or (internal IP) can use.

The two CLI outputs shown on this slide illustrate the behavior difference between the port block allocation IP pool type and the default overload IP pool type.

Using `hping`, a rogue client generates many SYN packets per second. In the first example, the port block allocation type limits the client to 64 connections for that IP pool. Other users will not be impacted by the rogue client.

In the second example, the overload type imposes no limits, and the rogue client uses many more connections in the session table. Other users will be impacted.

The port block allocation timeout period is configurable on the FortiGate CLI.

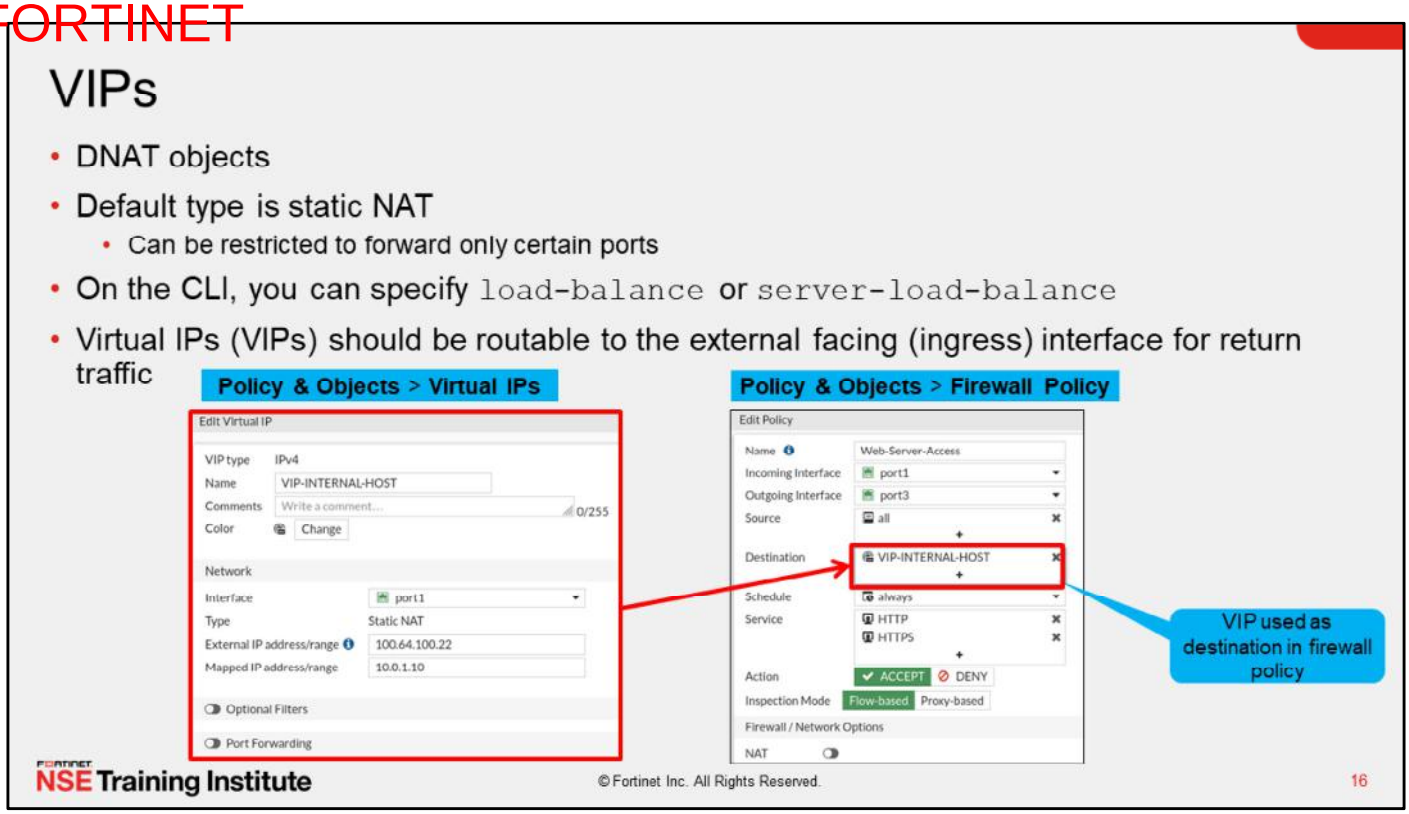

VIPs are DNAT objects. For sessions matching a VIP, the destination address is translated; usually a public internet address is translated to the private network address of a server. VIPs are selected in the firewall policy's **Destination** field.

The default VIP type is static NAT. This is a one-to-one mapping, which applies to incoming and outgoing connections; that is, an outgoing policy with NAT enabled would use the VIP address instead of the egress interface address. However, this behavior you can override using an IP pool.

You can restrict the static NAT VIP to forward only specific ports. For example, connections to the external IP on port 8080 map to the internal IP on port 80.

On the CLI, you can select the NAT type as `load-balance` and `server-load-balance`. Plain load balancing distributes connections from an external IP address to multiple internal addresses. Server load balancing builds on that mechanism, using a virtual server and real servers, and provides session persistence and server availability check mechanisms.

VIPs should be routable to the external facing (ingress) interface. FortiOS responds to ARP requests for VIP and IP pool objects. ARP responses are configurable.

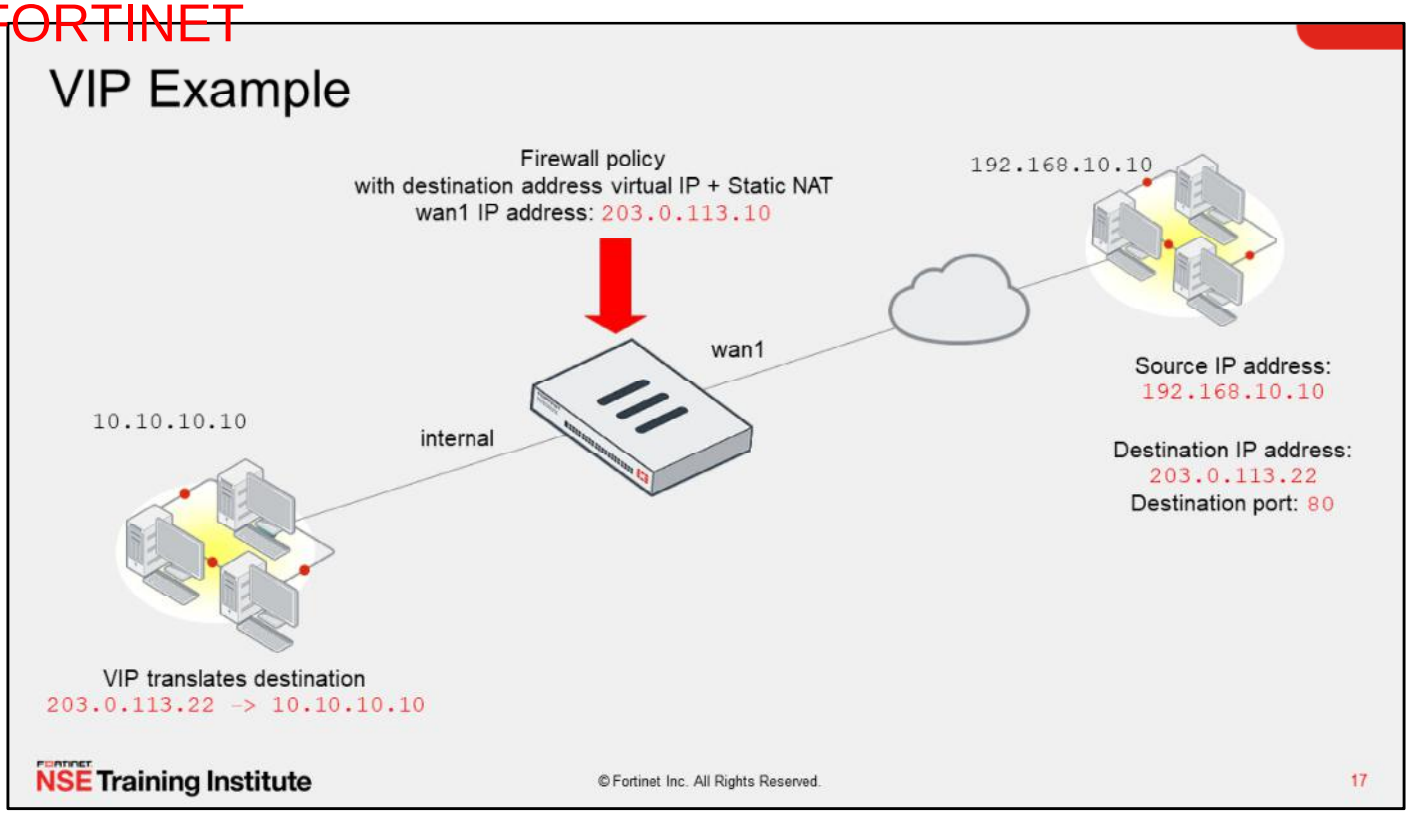

In the example shown on this slide, source IP address `192.168.10.10` is trying to access destination IP address `203.0.113.22` over port TCP `80`.

Connections to the VIP `203.0.113.22` are NATed to the internal host `10.10.10.10`.

Because this is a static NAT, all NATed outgoing connections from `10.10.10.10` will use VIP address `203.0.113.22` in the packet's IP source field and not the egress interface IP address `203.0.113.10`.

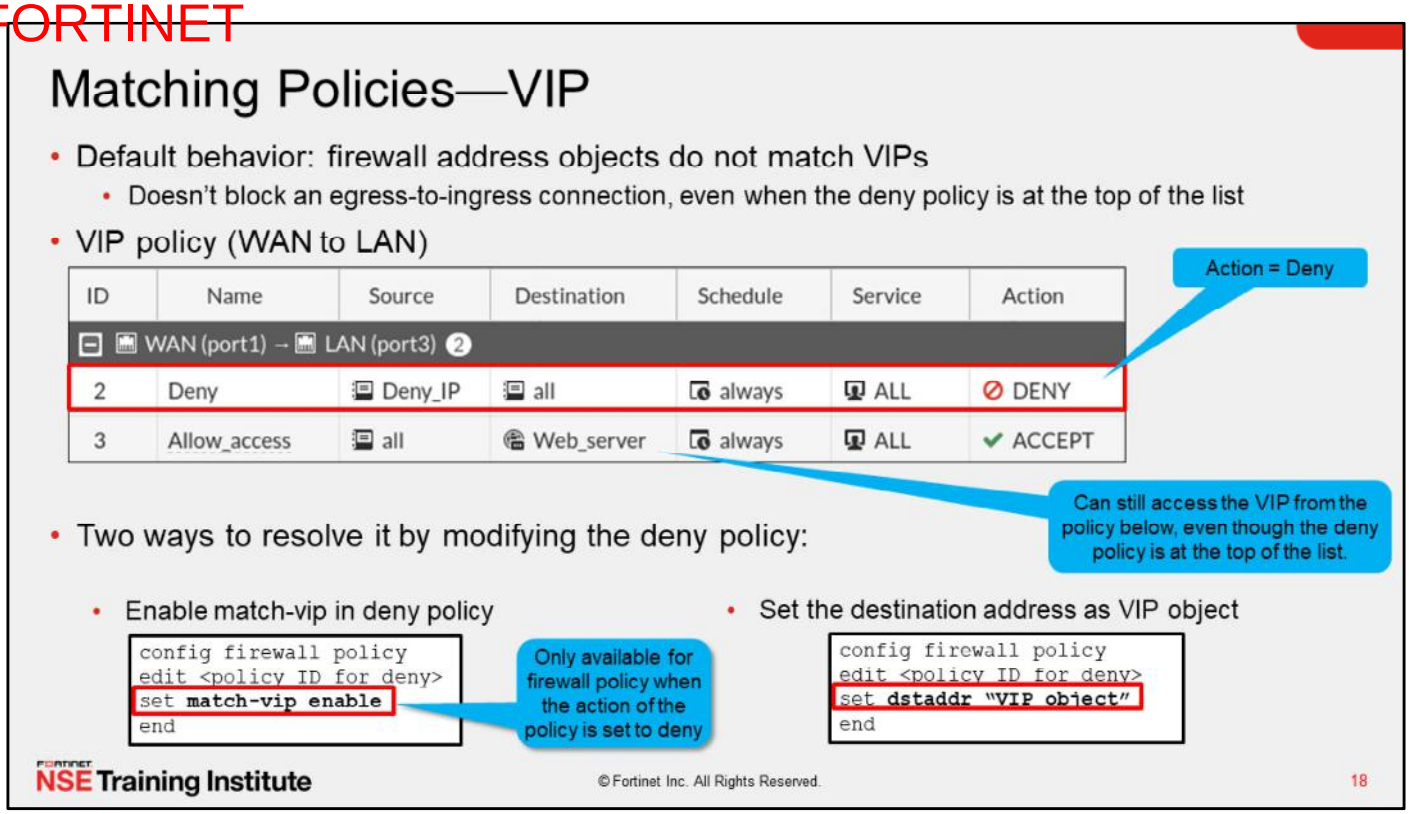

In FortiOS, VIPs and firewall address objects are completely different. They are stored separately with no overlap. By default, firewall address objects do not match VIPs. In the example shown on this slide, the **all** address object as a destination in the first policy does not include any VIPs, so traffic destined to the Webserver VIP skips the first policy and matches the second **Allow_access**. In order for the first policy to match the VIP, you need to either edit the policy on the CLI and `set match-vip enable`, which allows address objects to match the VIP address, or change the destination address of the first policy to be the VIP in question.

Traffic is permitted to fall through to the next policy; however, when you use VIP firewall policies, there can be some exceptions.

When VIP(s) are configured, for incoming (WAN to LAN) connections, it will be first matched against the VIP table.

In the example shown on this slide, a firewall policy from WAN to LAN is configured with a specific source and the action is **Deny**. There is a second firewall policy that is allowing access to the VIP (the destination address). Even though the deny firewall policy is at the top of the list, the denied source is still allowed by the second firewall policy to access the VIP.

In order to block traffic from the denied source, you must enable `set match-vip enable` in the deny firewall policy, which skips the VIP ID checking. Alternatively, you can configure the destination address as the virtual IP in the deny policy instead of **all**.

Note `set match-vip enable | disable` is only available for firewall policy when the action of the policy is set to deny.

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`

19

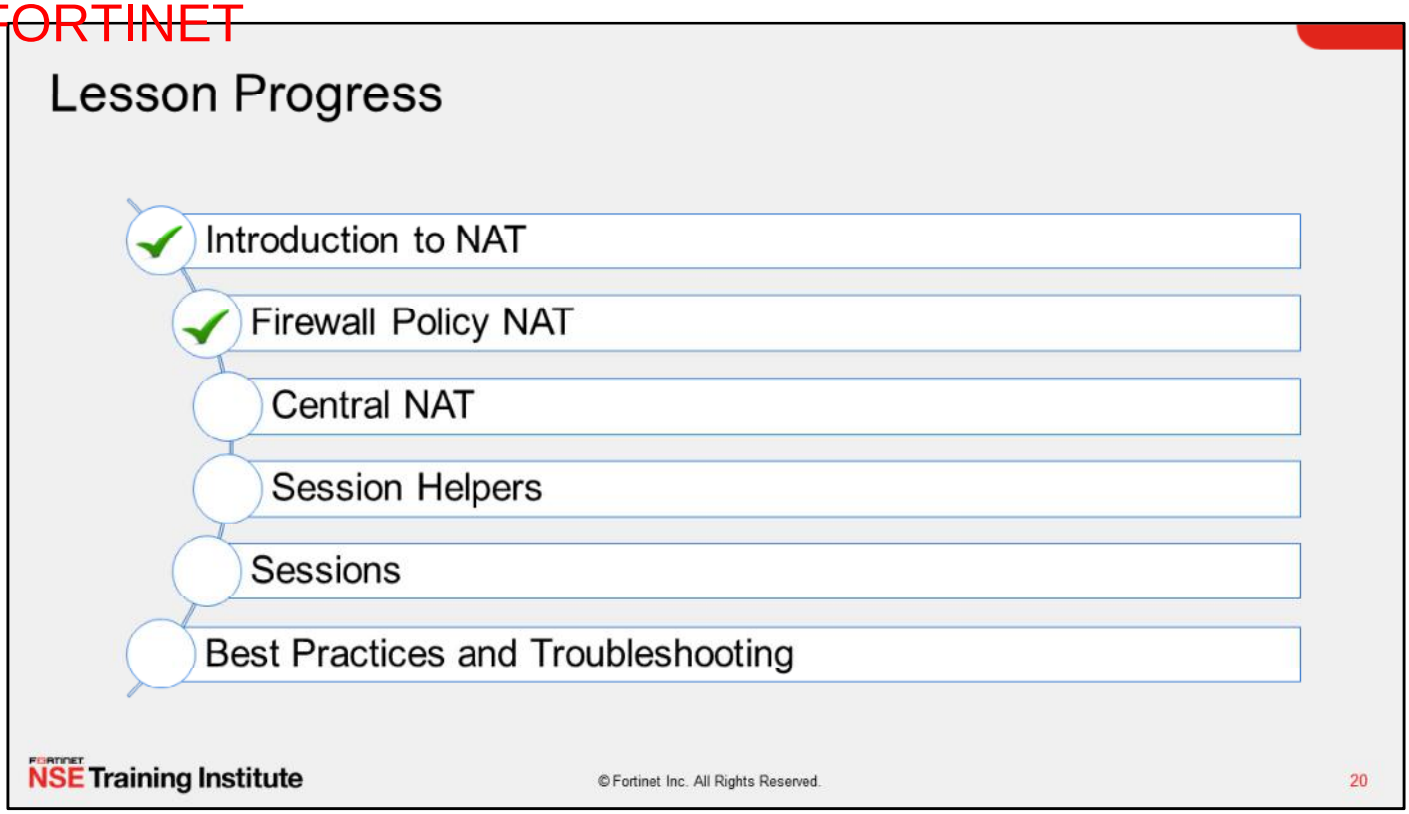

Good job! You now understand firewall policy NAT.

Now, you'll learn about central NAT.

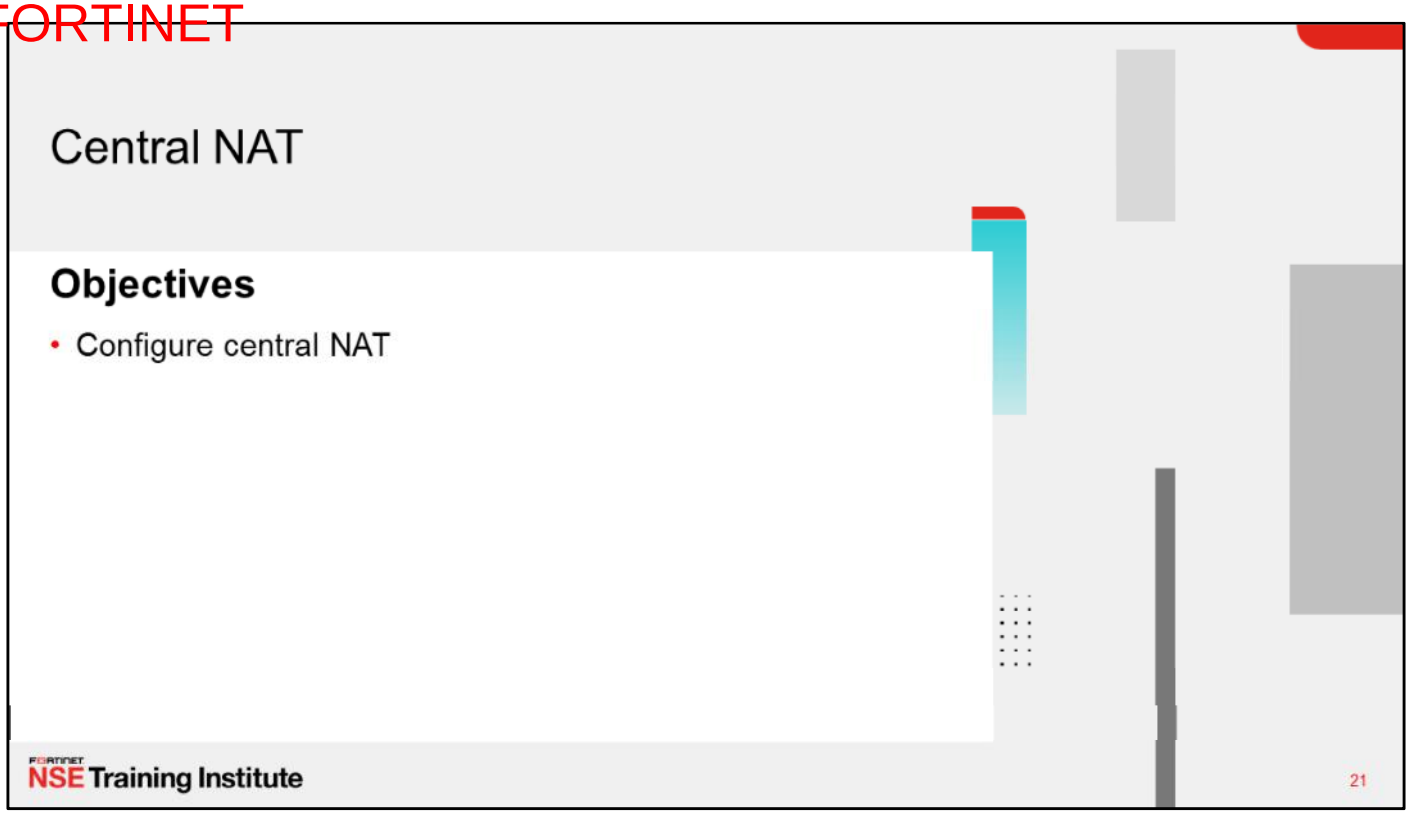

After completing this section, you should be able to achieve the objective shown on this slide.

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.

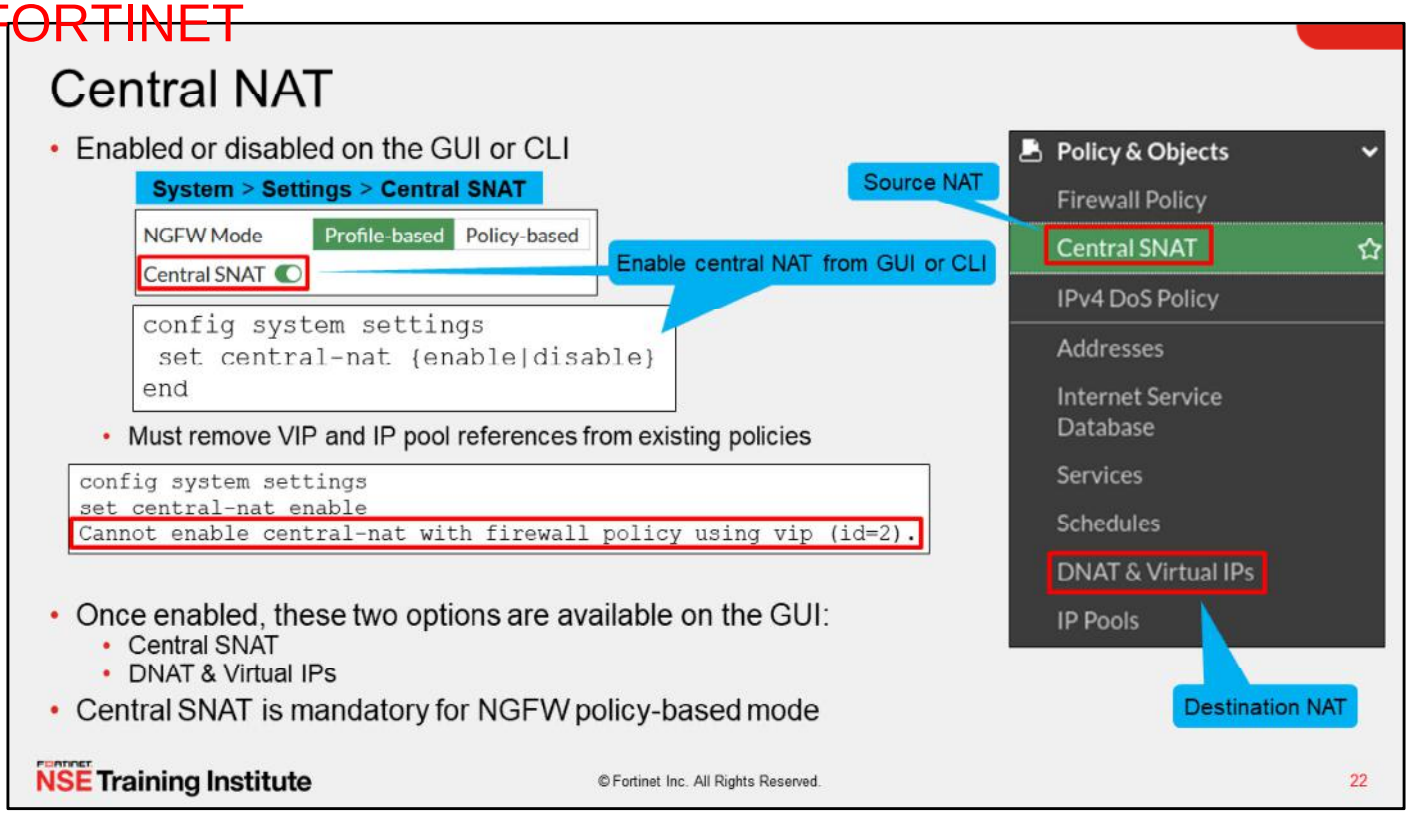

By default, central NAT is disabled. You can enable it on the CLI or the GUI. After central NAT is enabled, the following two options are available to be configured on the GUI:

- **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 will 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**.

Central SNAT

- SNAT configuration changes when you enable central NAT

Central NAT Enabled	Steps to Configure
SNAT	1. Define IP pool or use outgoing interface address 2. Configure central SNAT policy

- Central SNAT rules process from top to bottom
  - SNAT policy is selected according to the configuration of the firewall policy that matches the traffic
  - If no matching central SNAT rule exists, NAT will not be applied
- Matching criteria is based on:
  - Source interface
  - Destination interface
  - Source address
  - Destination address
  - Protocol
  - Source port
    - Most protocols don't need this

Policy & Objects > Central SNAT

ID	From	To	Source Address	Destination Address	Translated Address
IPv4 2					
1	LAN (port3)	WAN1 (port1)	all	all	SNAT-Pool
2	port4	WAN2 (port2)	LOCAL_SUBNET	REMOTE_SUBNET	INTERNAL-HOST-EXT-IP

NSE Training Institute

23

You can have more granular control, based on source and destination interfaces in the central SNAT policy, over traffic passing through firewall policies.

You can now define matching criteria in the central SNAT policy, based on:
- Source interface
- Destination interface
- Source address
- Destination address
- Protocol
- Source port

A matching central SNAT policy is mandatory for all firewall policies. If there is no matching SNAT policy, no NAT will be applied and a session is created using the original source IP address.

If the central SNAT policy criteria matches the traffic based on multiple firewall policies, the central SNAT policy is applied to those firewall policies.

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

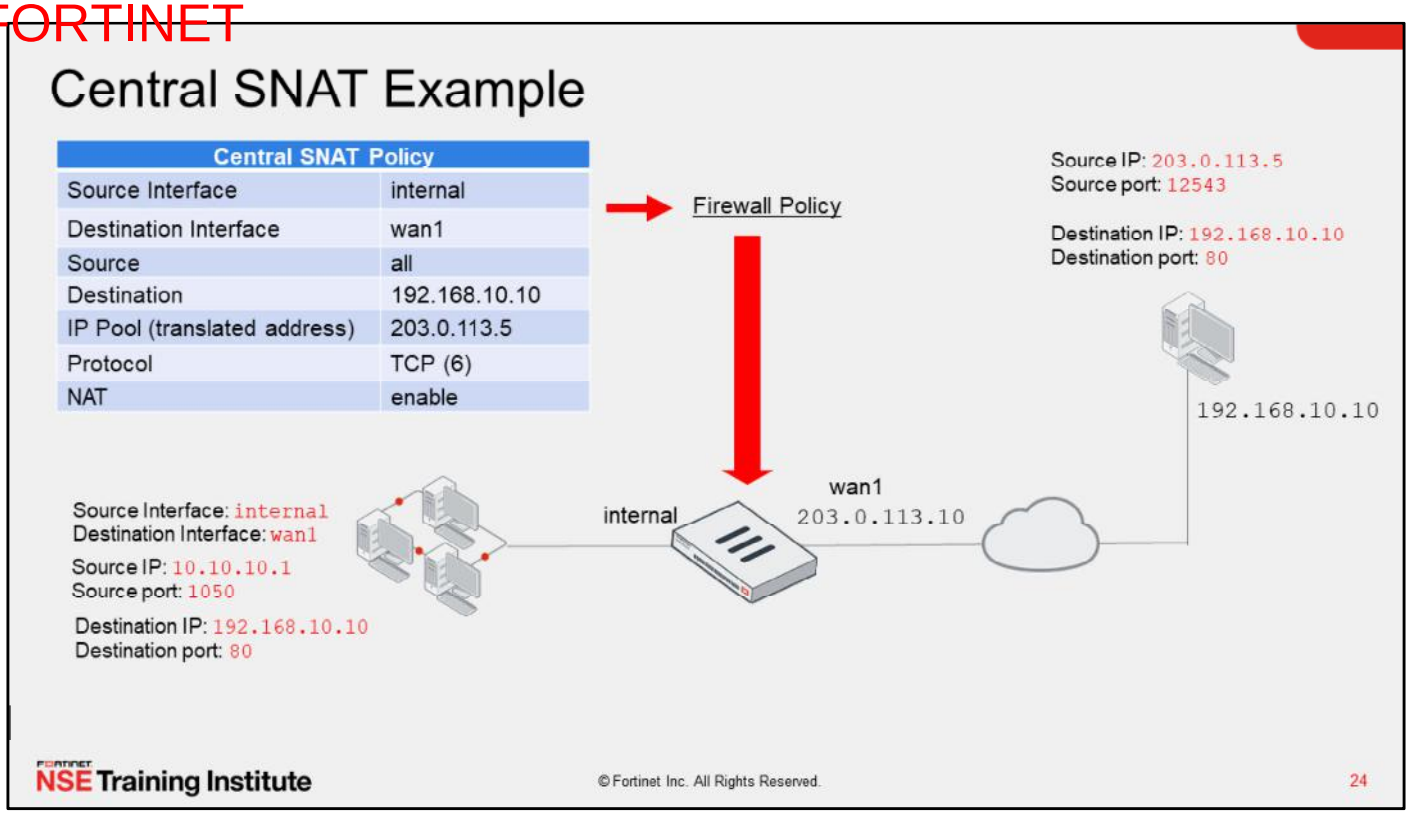

In the example shown on this slide, the central SNAT policy translates the source IP address to the defined IP pool address (`203.0.113.5`). However, the translation takes place only if the traffic matches all the variables defined in the central SNAT policy, that is, traffic from the source IP address through the source interface internal to FortiGate must be destined for destination IP address (`192.168.10.10`) through destination interface wan1 and the protocol must be TCP. For illustration purposes, only a single IP address is used for the destination, and the IP pool type is set to overload with a single IP address.

The firewall policy is created from internal to wan1. There is no NAT option available in firewall policy, and matching SNAT policy is mandatory to pass the traffic. If there is no matching central SNAT policy, no NAT will be applied and the session will be created using the original source IP address.

If the user tries any TCP-based sessions (for example HTTP, HTTPS) to the destination IP address `192.168.10.10`, the source IP address is translated to an IP pool address or addresses defined in the central NAT policy.

What if the user tries to send any ICMP or UDP-based traffic to `192.168.10.10`? Will the source address be translated to the IP pool defined in the central NAT policy?

Because the central SNAT policy does not match, FortiGate applies no NAT. What if the user tries TCP-based traffic to another destination IP address, `192.168.10.20`? Will the source address be translated to the IP pool defined in the central SNAT policy?

Again, the destination IP address of `192.168.10.20` does not match the central NAT policy, so FortiGate does not apply NAT and the firewall session is created using the original source IP.

Central DNAT and VIPs

- Enabling central NAT changes the DNAT configuration

Central NAT Enabled	Steps to Configure
Destination NAT (VIP)	Define DNAT & Virtual IPs (No additional configurations required)

- As soon as a VIP is created, a rule is created in the kernel to allow DNAT to occur
  - Firewall policy destination address—all or mapped IP of VIP
    - VIP cannot be selected in the firewall policy as the destination address
- You can exclude a VIP by changing the status of the VIP to disable from CLI, to easily manage VIP when central NAT is enabled

```
#config firewall vip
    edit "name of the VIP"
        set status disable
    next
end
```

NSE Training Institute

25

Traditionally on FortiGate, you select VIPs in the firewall policy as the destination address.

On FortiGate, you can configure DNAT and VIPs for DNAT. As soon as you configure a VIP, FortiGate automatically creates a rule in the kernel to allow DNAT to occur. You do not need to do any additional configuration.

Do you lose the granularity of being able to define a firewall policy for a specific VIP and services?

No, you don't. If you have several WAN-to-internal policies and multiple VIPs, and you want to allow specific services for specific VIPs, you can define each firewall policy with the destination address of the mapped IP of the VIP, and select the appropriate services to allow or deny.

VIP takes effect right after it is created when Central NAT is enabled, If you want to exclude a VIP, you can change the status of the VIP to disable on the CLI, to easily manage VIP when central NAT is enabled.

Note that if both central SNAT and central DNAT (VIP) are configured, the outgoing (internal-to-WAN) traffic will source NAT, based on the matching central SNAT policy configurations, and if there is no matching central SNAT policy no NAT is applied to outgoing traffic.

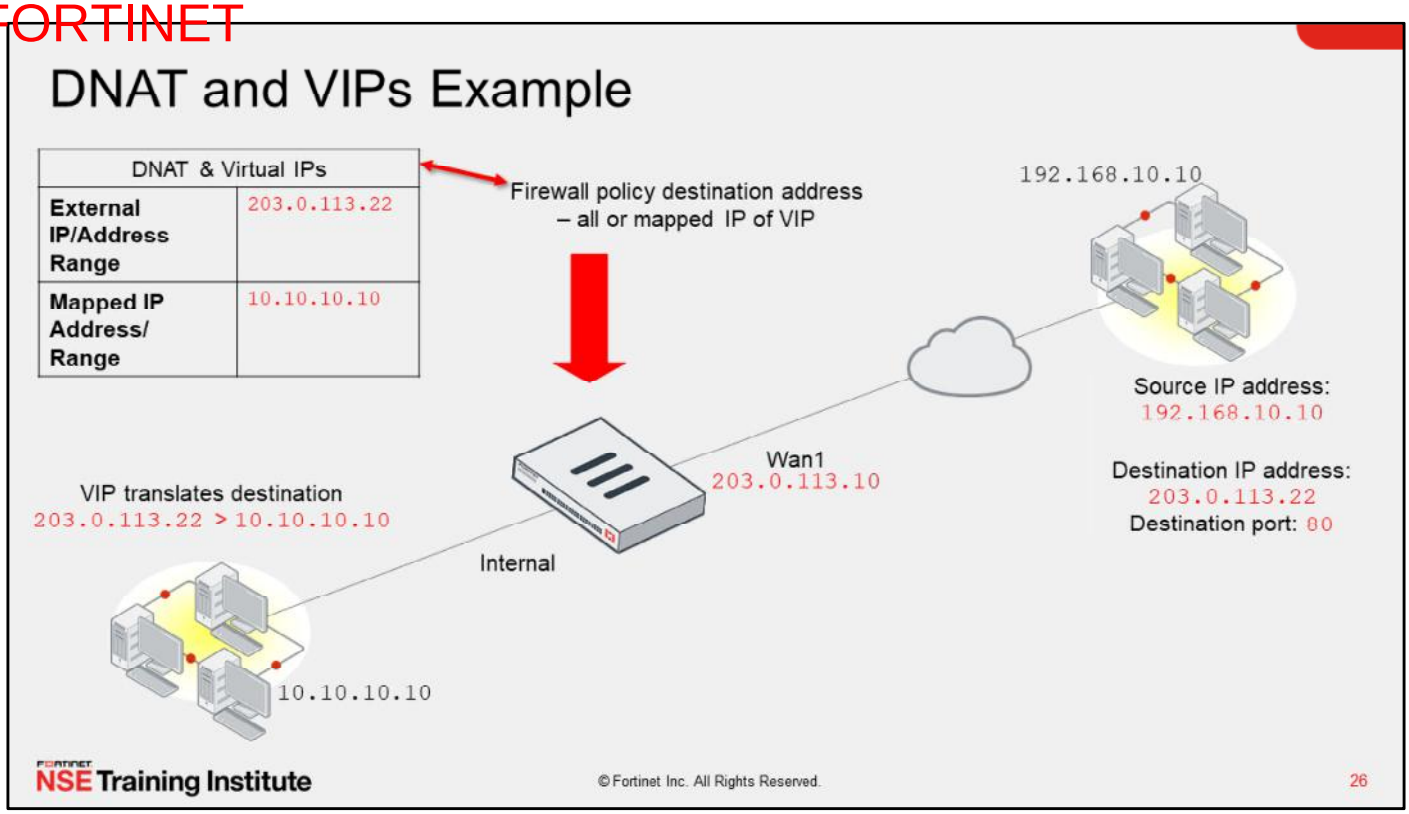

In the example shown on this slide, a DNAT and VIP rule is created to map external IP address `203.0.113.22` to internal IP address `10.10.10.10`. Remember, as soon as you create a VIP, a rule is created in the kernel to allow DNAT to occur.

The firewall policy from wan1 to internal is created with the destination address **all** or **Mapped IP Address/Range** (`10.10.10.10`) of the VIP.

The source IP address 192.168.10.10 is trying to access the destination IP address `203.0.113.22` over port TCP 80. Connections to the VIP `203.0.113.22` are NATed to the internal host `10.10.10.10`, without any additional configuration.

Disabling Central NAT

- If central NAT is enabled and configured for SNAT and DNAT, and then disabled, the following occurs:
  - Outgoing traffic SNAT is no longer performed
  - Incoming traffic that was previously configured with DNAT and VIPs stops working because there is no rule present in the kernel for DNAT

NSE Training Institute © Fortinet Inc. All Rights Reserved. 27

You can disable central NAT on the CLI by running `set central-nat disable` under the `config system setting`. What happens to firewall policies that are using central SNAT and DNAT rules, if central NAT is disabled?

For new firewall sessions, the incoming to outgoing firewall policies no longer perform SNAT. You must manually edit firewall policies to enable NAT and select appropriate IP pool addresses, which were previously tied to the central SNAT policy.

Egress-to-ingress firewall policies that use DNAT and VIP will no longer perform DNAT. In central NAT, the destination address in the firewall policy is simply an address object, not an actual VIP. Without the central NAT hook into the DNAT table, the address object causes a forward policy check failure—the traffic is denied by policy ID 0.

You must edit the egress-to-ingress firewall policies and select VIP as the destination address.

Knowledge Check

1. Which statement is true?
   - ✓ A. Central NAT is not enabled by default.
   - B. Both central NAT and firewall policy NAT can be enabled together.

2. What happens if there is no matching central SNAT policy or no central SNAT policy configured?
   - A. The egress interface IP will be used.
   - ✓ B. NAT will not be applied to the firewall session.

NSE Training Institute

28

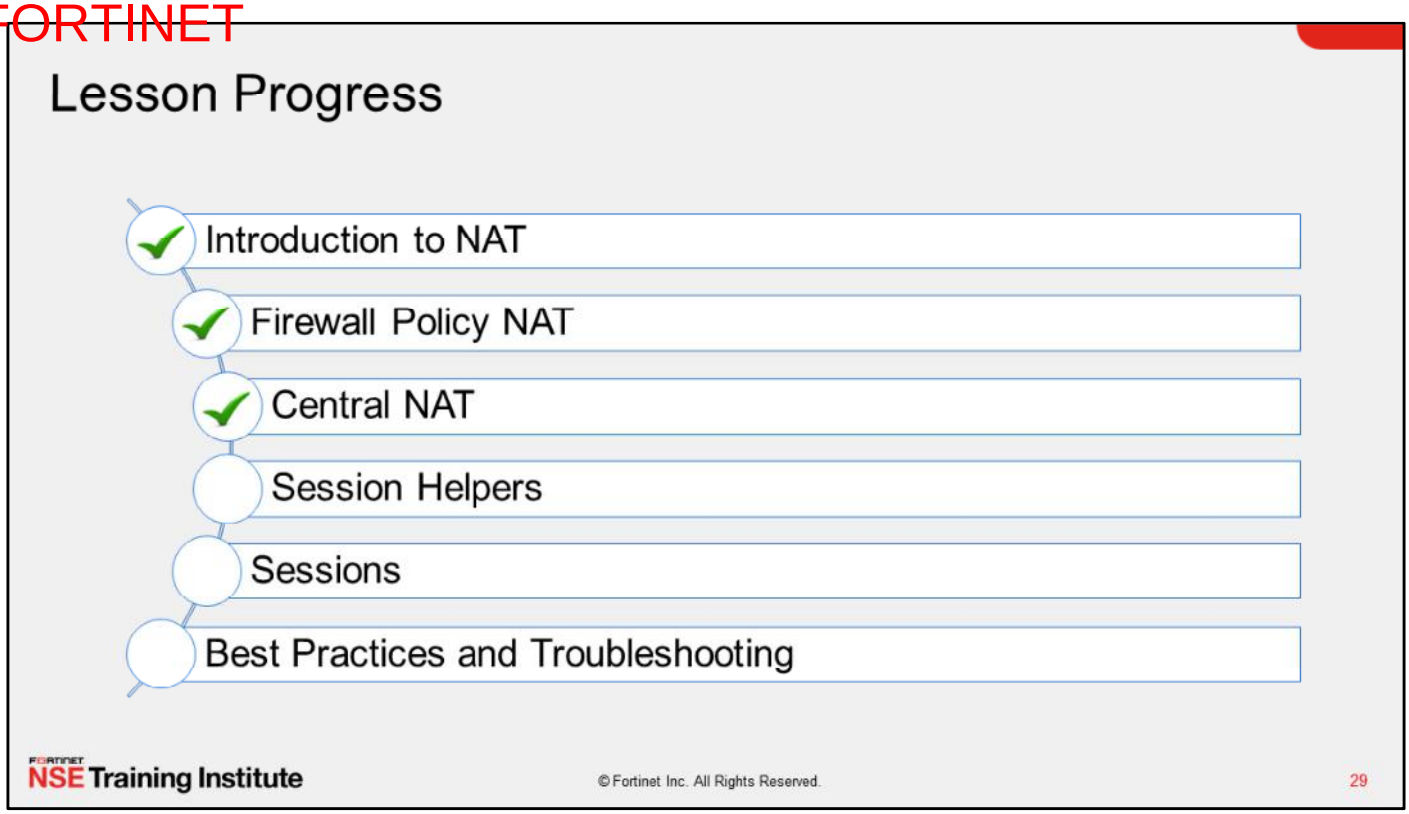

Good job! You now understand central NAT.

Now, you'll learn about session helpers.

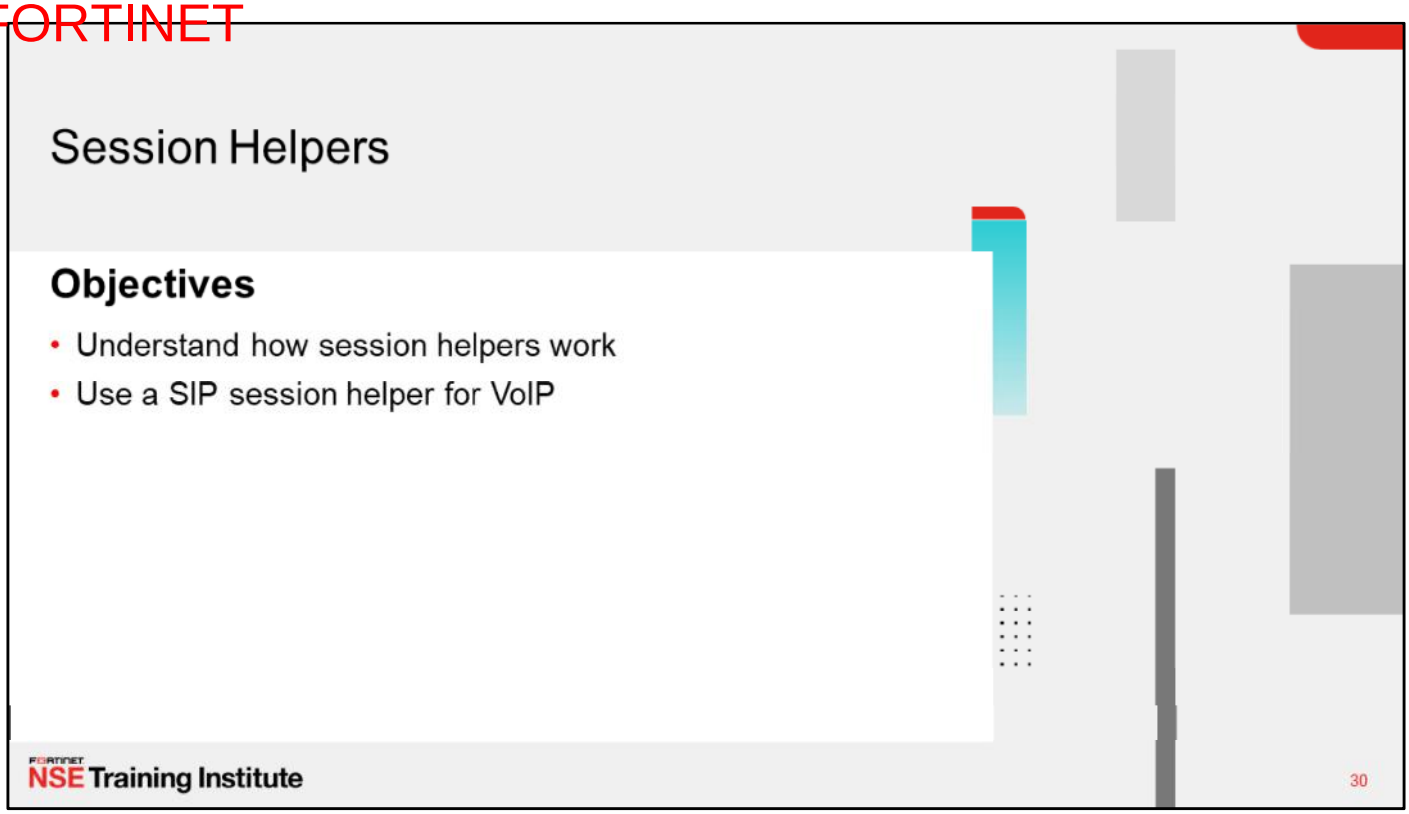

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding how session helpers work, you will be able to use session helpers to analyze data in the packets of some protocols, and allow those protocols to pass traffic through FortiGate.

Session Helpers

- Some traffic types require more packet modification for the application to work (configurable on the CLI), examples include:
  - The handling of FTP active mode connections—the control connection is separate from the data connection
  - Header rewrites in SIP SDP payloads required because of NAT actions
- To show configured session helpers, use this command:
  - `show system session-helper`
- Application layer gateway (ALG)
  - When more advanced application tracking and control is required, an ALG can be used—the VOIP profile is an example of an ALG.

NSE Training Institute © Fortinet Inc. All Rights Reserved. 31

Some application layer protocols are not fully independent of the lower layers, such as the network or transport layers. The addresses may be repeated in the application layer, for example. If the session helper detects a pattern like this, it may change the application headers, or create the required secondary connections.

A good example of this is an application that has both a control channel and a data or media channel, such as FTP. Firewalls will typically allow the control channel and rely on the session helpers to handle the dynamic data or media transmission connections.

When more advanced application tracking and control is required, you can use ALG. The VOIP profile is an example of an ALG.

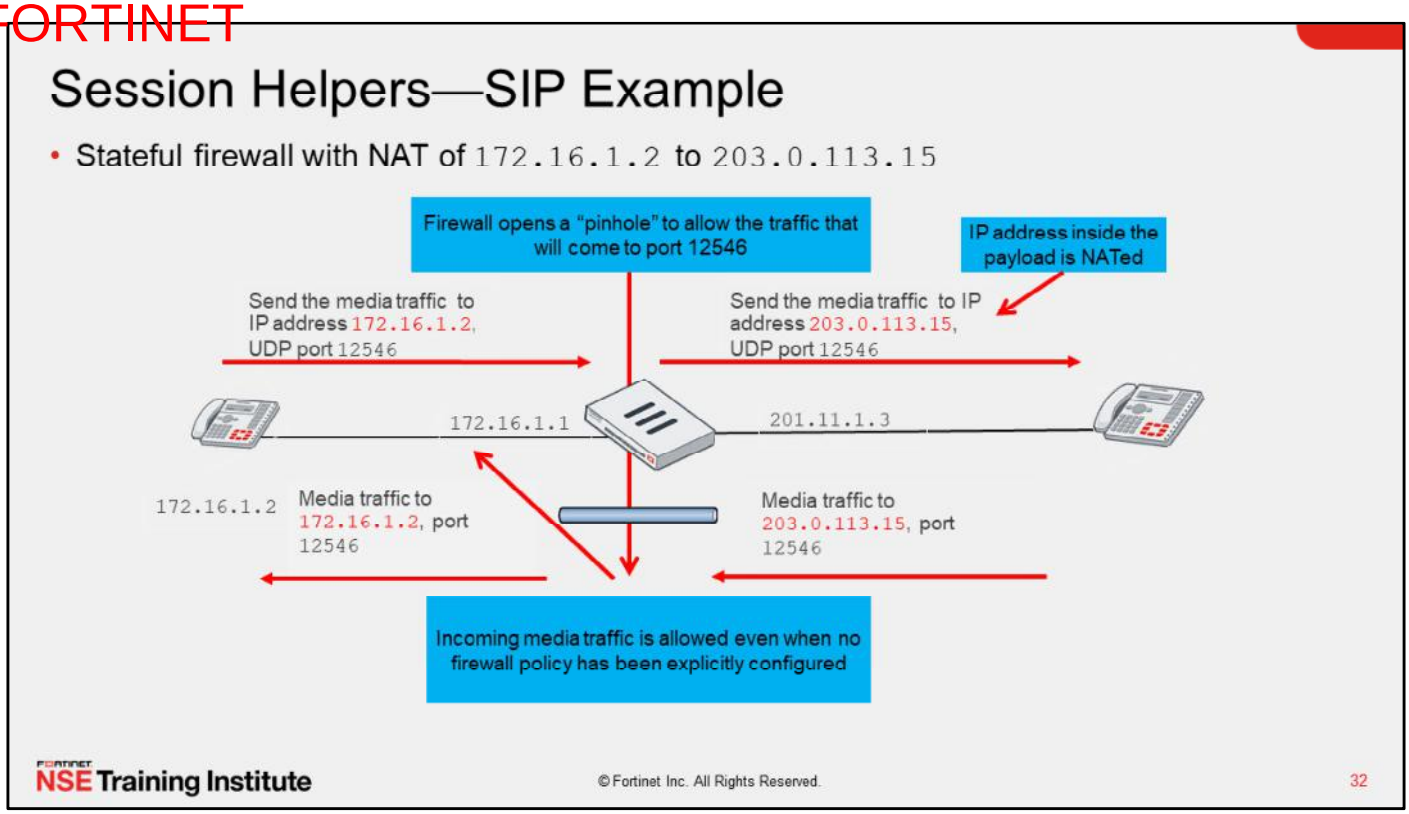

In the examples shown on this slide, the media recipient address in the SIP SDP payload is modified to reflect the translated IP address.

Notice how, because firewall policies are stateful, a pinhole is opened to allow reply traffic, even though you have not explicitly created a firewall policy to allow incoming traffic. This concept is used with some other protocols, such as NAT-T for IPsec.

Knowledge Check

1. Which method would you use for advanced application tracking and control?
   A. Session helper
   ✓ B. Application layer gateway

2. Which profile is an example of application layer gateway?
   A. WAF profile
   ✓ B. VOIP profile

NSE Training Institute

33

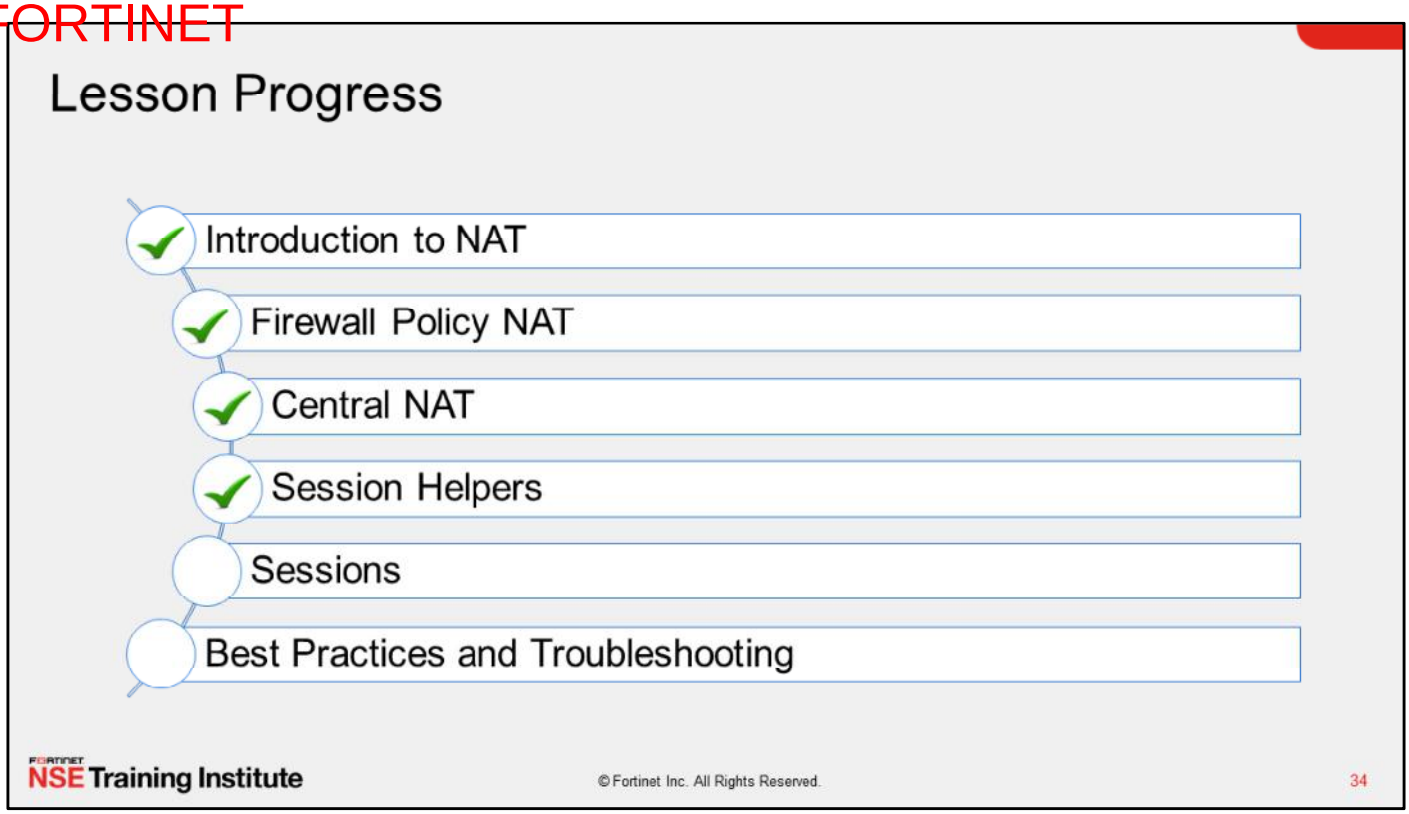

Good job! You now understand session helpers.

Now, you'll learn about sessions.

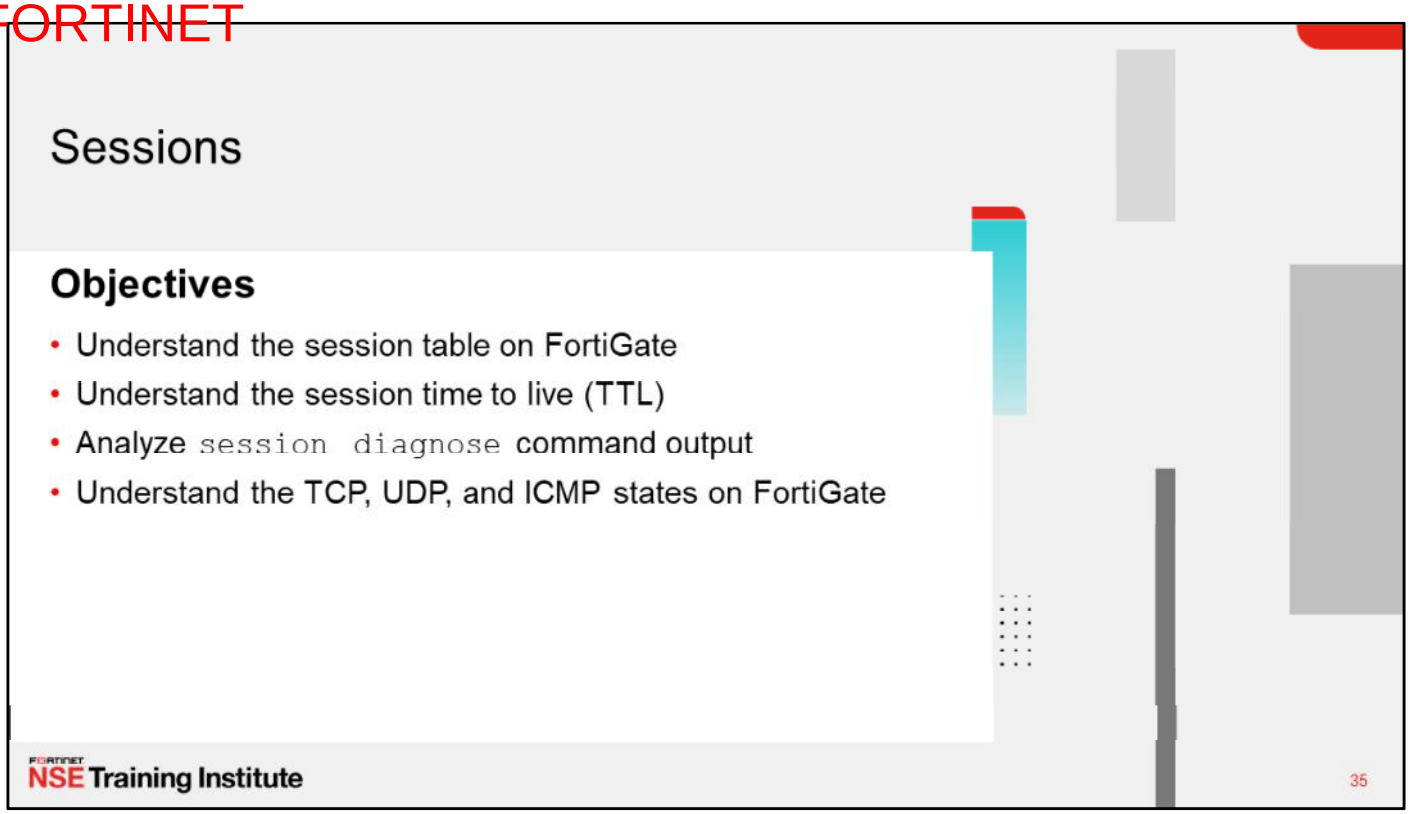

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding how a session table keeps track of the session information, you will be able to use that information effectively to understand the actions applied to traffic, such as SNAT, DNAT, and routing.

Session Table

- Accepted IP sessions are tracked in the kernel session table, but this can be affected by hardware acceleration
- The session table stores the following information about the session:
  - The source and destination addresses, port number pairs, state, and timeout
  - The source and destination interfaces
  - The source and destination NAT actions
- The session table stores the following performance metrics:
  - Maximum concurrent sessions
  - New sessions per second

Dashboard > FortiView Sessions

FortiView Sessions

Add Filter

Source	Device	Destination	Application	Protocol	Source Port	Destination Port	Bytes	Packets	Duration (seconds)
10.0.1.10		10.160.0.1	UDP/53	UDP	55061	53	378 B	6	1m 4s
10.0.1.10		8.8.8.8	UDP/53	UDP	60093	53	164 B	2	2m 15s
10.0.1.200		208.91.112.52	UDP/53	UDP	2022	53	197 B	3	26s
10.0.1.200		208.91.112.53	UDP/53	UDP	2022	53	628 B	10	23s

NSE Training Institute © Fortinet Inc. All Rights Reserved. 36

You can view the **Sessions** page on the GUI, but the CLI provides more information regarding sessions in the session table.

Firewall performance of connections for each session, and the maximum number of connections, are indicated by the session table. However, if your FortiGate device contains security processors designed to accelerate processing without loading the CPU, the session table information might not be completely accurate, because the session table reflects what is known to, and processed by, the CPU.

Session TTL

- When the session table is full, reducing timers may improve performance by closing sessions earlier; however, be careful not to close sessions *too* soon, because this can cause connection errors

TCP default TTL

```
config system session-ttl
set default 3600
end
```

Specific state timers

```
config system global
set tcp-halfclose-timer 120
set tcp-halfopen-timer 10
set tcp-timewait-timer 1
set udp-idle-timer 180
end
```

- Timers can be applied in policies and objects, and have precedence:
  - **Firewall Services** > **Firewall Policies** > **Global Sessions**

NSE Training Institute

37

Each session on FortiGate can idle for a finite time, which is defined by TTL. When the FortiGate detects the session is idle after some time of inactivity, and TTL is reached, the session is deleted from the session table.

Because the session table has a finite quantity of RAM that it can use on FortiGate, adjusting the session TTL can improve performance. There are global default timers, session state timers, and timers configurable in firewall objects.

Firewall Session Diagnostics

- `diagnose sys session`
  - The session table also indicates policy actions.
  - Clear any previous filter:
    - `diagnose sys session filter clear`
  - Set the filter:
    - `diagnose sys session filter ?`
    - `dport  Destination port`
    - `dst  Destination IP address`
    - `policy  Policy ID`
    - `sport  Source port`
    - `src  Source IP address`
  - List all entries matching the configured filter:
    - `diagnose sys session list`
  - Purge all entries matching the configured filter:
    - `diagnose sys session clear`

NSE Training Institute

38

The `diagnose sys session` command tree provides options to filter, clear, or show the list of sessions. You can also list brief information about sessions by running the `get system session list` command.

Before looking at the session table, first build a filter. To look at our test connection, you can filter on `dst 10.200.1.254` and `dport 80`.

Session Table—TCP Example

```
# diagnose sys session filter dst 10.200.1.254
# diag sys session filter dport 80
# diag sys session list

session info: proto=6 proto_state=05 duration=2 expire=78 timeout=3600 flags=00000000
socktype=0 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=may_dirty
statistic(bytes/packets/allow_err): org=538/6/1 reply=5407/6/0 tuples=2
tx speed(Bps/kbps): 5/0 rx speed(Bps/kbps): 2/0

orgin->sink: org pre->post, reply pre->post dev=5->3/3->5 gwy=10.200.1.254/10.0.1.10

hook=post dir=org act=snat 10.0.1.10:64624->10.200.1.254:80(10.200.1.1:64624)
hook=pre dir=reply act=dnat 10.200.1.254:80->10.200.1.1:64624(10.0.1.10:64624)
pos/(before,after) 0/(0,0), 0/(0,0)

misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=0

.............
```

TCP State

Session TTL

Routing operation

NAT operation

Policy ID

NSE Training Institute

39

In the example shown on this slide, you can see the session TTL, which reflects how long FortiGate can go without receiving any packets for this session, until it removes the session from its table.

Here you can see the routing and NAT actions that apply to the traffic. The firewall policy ID is also tracked.

The `proto_state` for TCP is taken from its state machine, which you'll learn about later in this lesson.

TCP States

- proto_state=05
  - First digit (from left to right): server-side state
    - 0 if no inspection, 1 if proxy or flow
  - Second digit (from left to right): client-side state

TCP State	Value	Expire Timer in sec (default)
NONE	0	10
ESTABLISHED	1	3600
SYN_SENT	2	120
SYN & SYN/ACK	3	60
FIN_WAIT	4	120
TIME_WAIT	5	1
CLOSE	6	10
CLOSE_WAIT	7	120
LAST_ACK	8	30
LISTEN	9	120

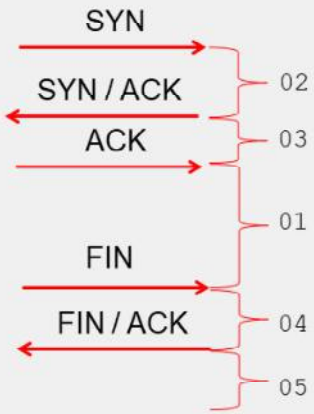

NSE Training Institute

40

Earlier in this lesson, you learned that the session table contains a number that indicates the current TCP state of the connection. These are the states of the TCP state machine. They are single-digit values, but `proto_state` is always shown as two digits. This is because FortiGate is a stateful firewall and keeps track of the original direction (client-side state) and the reply direction (server-side state). If there are too many connections in the SYN state for long periods of time, this indicates a SYN flood, which you can mitigate with DoS policies.

This table and flow graph correlates the second digit value with the different TCP session states. For example, when FortiGate receives the SYN packet, the second digit is `2`. It goes to `3` once the SYN/ACK is received. After the three-way handshake, the state value changes to `1`.

When a session is closed by both sides, FortiGate keeps it in the session table for a few seconds more, to allow any out-of-order packets that could arrive after the FIN/ACK packet. This is the state value `5`.

ICMP and UDP Protocol States

- Even though UDP is stateless, FortiGate still uses two session state values:

UDP State	Value
UDP traffic one way only	0
UDP traffic both ways	1

- ICMP has no state
  - `proto_state` is always 00

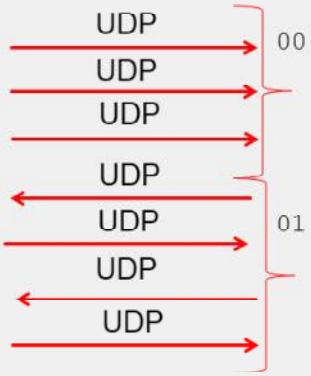

Although UDP is a message-oriented, stateless protocol, it doesn't inherently require confirmed bidirectional connections like TCP, so there is no connection state. However, the FortiGate session table does use the `proto_state=` field to track the unidirectional UDP as state `0`, and the bidirectional UDP as state `1`.

When FortiGate receives the first packet, it creates the entry and sets the state to `0`. If the destination replies, FortiGate updates the state flag to `1` for the remainder of the conversation.

Notably, ICMP, such as ping and traceroute, have no protocol state and it will always show `proto_state=00`.

Knowledge Check

1. If session diagnostic output indicates that a TCP protocol state is `proto_state=01`, which is true?

✓ A. The session is established.

B. The session is not established.

NSE Training Institute

42

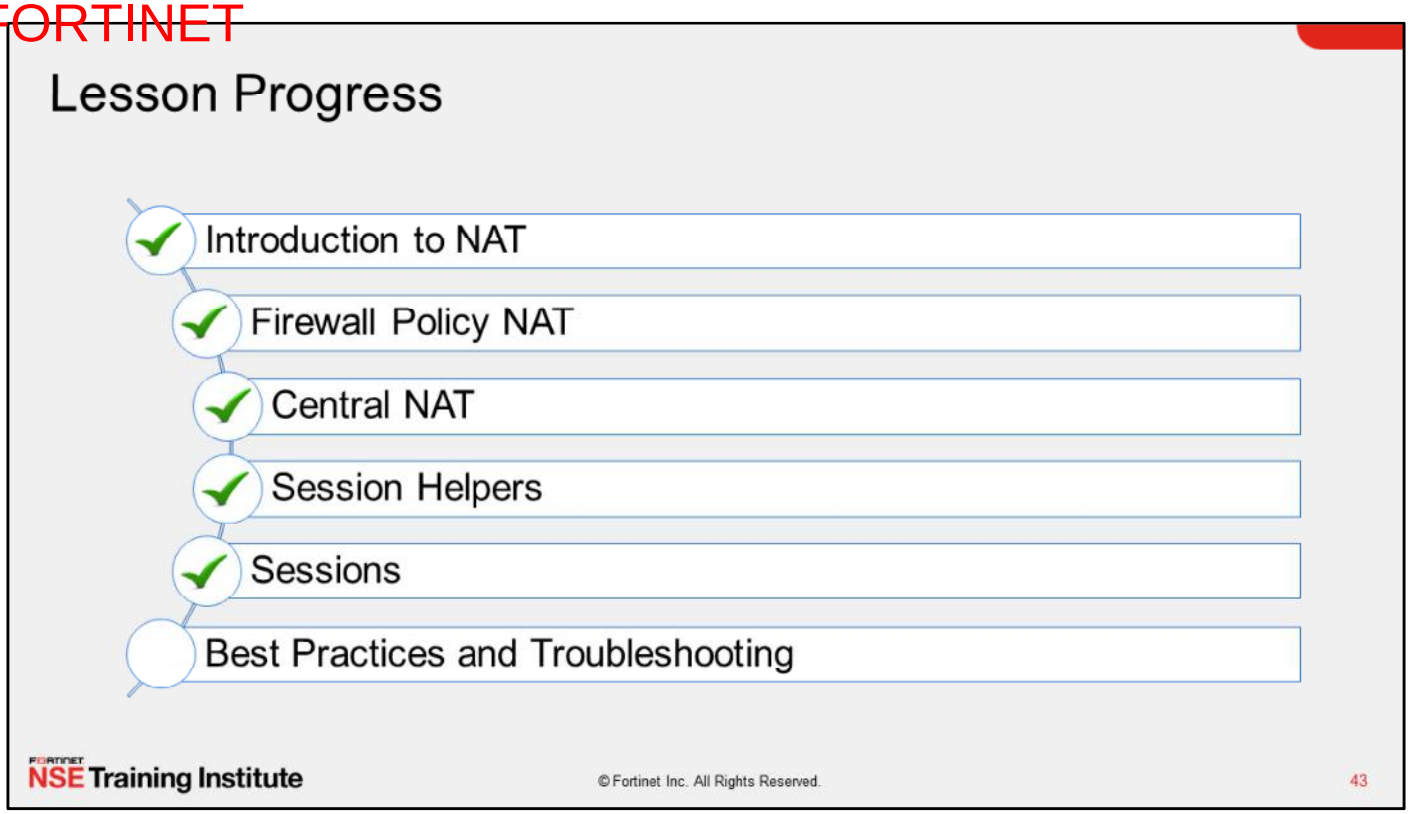

Good job! You now understand sessions.

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

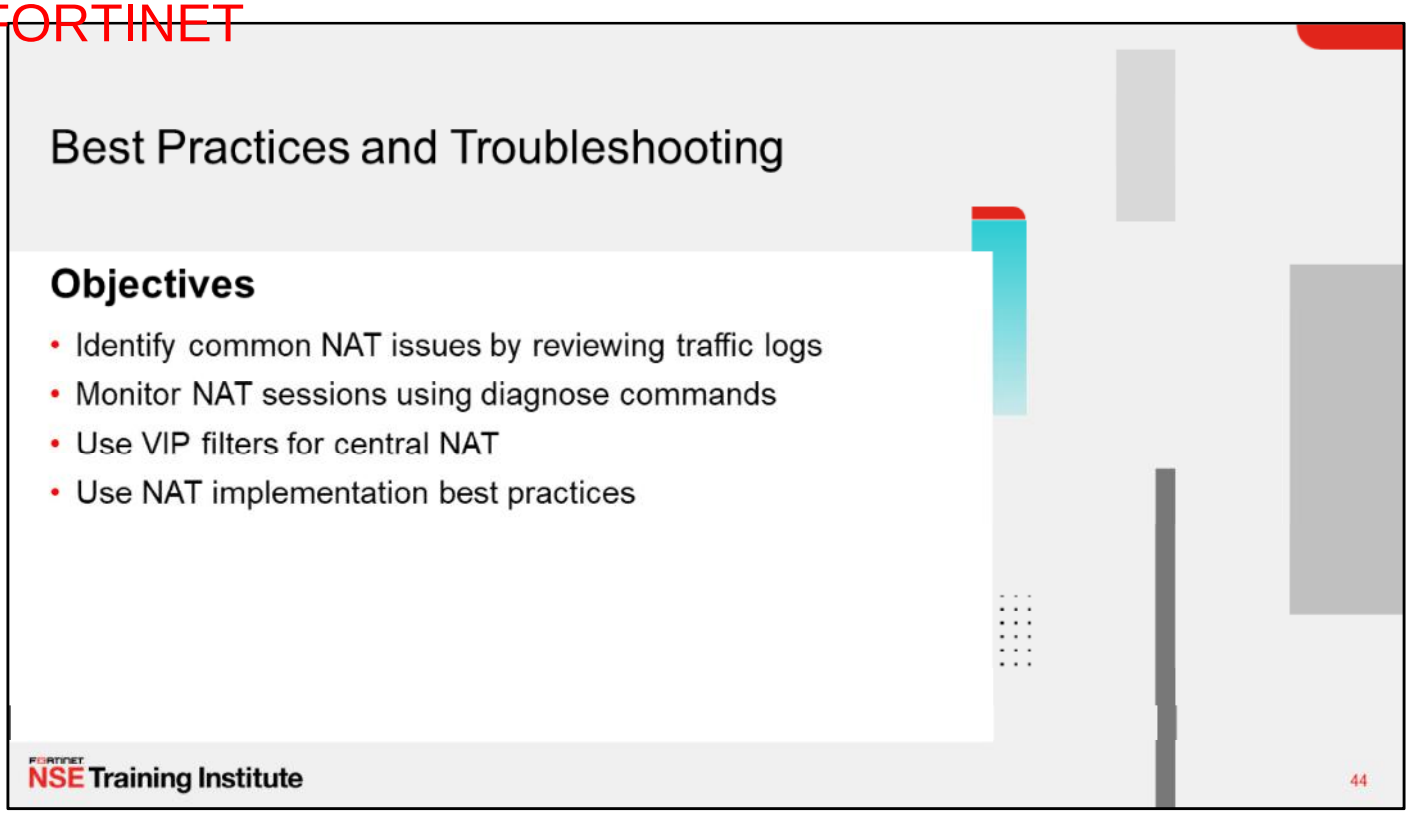

After completing this section, you should be able to achieve the objectives shown on this slide.

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

NAT Port Exhaustion

- If traffic log is enabled, the following log is displayed when the NAT ports are exhausted:

```
Message meets Alert condition date=2011-02-01 time=19:52:01 devname=master
device_id="" log_id=0100020007 type=event subtype=system pri=critical vd=root
service=kernel status=failure msg="NAT port is exhausted."
```

- NAT port exhaustion is also highlighted by a rise in the clash counter from the `diagnose system session stat` command:

```
# diagnose sys session stat
misc info: session_count=16 setup_rate=0 exp_count=0 clash=889
memory_tension_drop=0 ephemeral=1/16384 removeable=3
delete=0, flush=0, dev_down=16/69 ses_walkers=0
…
firewall error stat:
…
…
ids_recv=000fdc94
url_recv=00000000
av_recv=001fee47
fqdn_count=00000000
fqdn6_count=00000000
global: ses_limit=0 ses6_limit=0 rt_limit=0 rt6_limit=0
```

A number above 0 indicates that some sessions have been rejected because of NAT port exhaustion

NSE Training Institute © Fortinet Inc. All Rights Reserved. 45

NAT port exhaustion occurs when FortiGate is unable to allocate ports for performing NAT on new sessions because there are no available ports. When NAT port exhaustion occurs, FortiGate informs the administrator by displaying the log shown on this slide, with a severity of critical.

To address NAT port exhaustion, you must take one of the following actions:
- Create an IP pool that has more than one external IP tied to it (so it load balances across them)
- Reduce the number of sessions that require NAT

To receive important logs like this one, you must make sure that the necessary logging is enabled. On the FortiGate GUI, click **Log&Report** > **Log Settings**, to check that the default setting, logging to disk or memory, is activated.

NAT port exhaustion is also highlighted by a rise in the clash counter from the `diagnose system session stat` command.

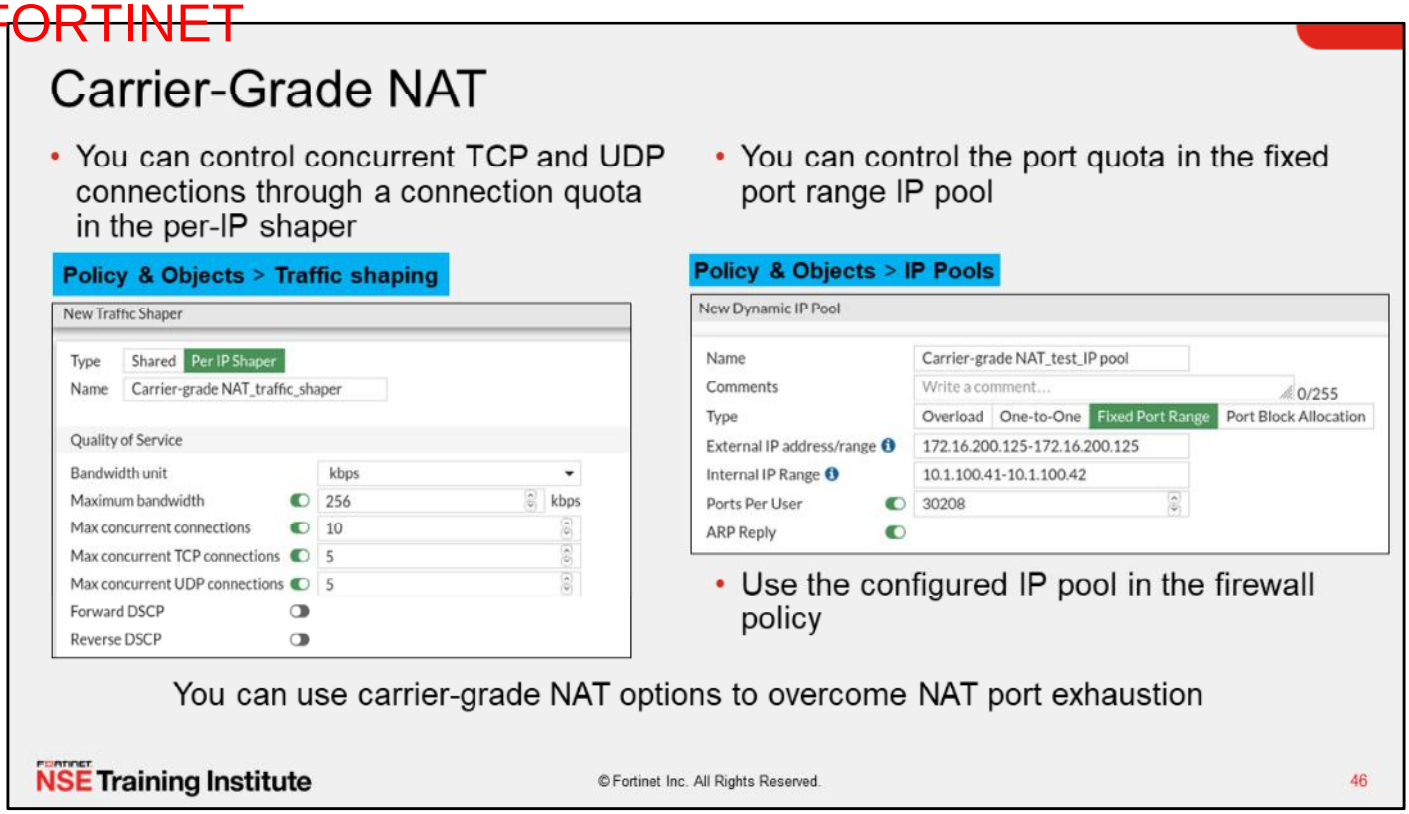

You can control concurrent TCP and UDP connections through a connection quota in the per-IP shaper. Maximum number of concurrent TCP and UDP sessions allowed by a shaper is from zero to 2097000, zero meaning no limit. You can select a value of zero only from the CLI configuration of the traffic shaper.

You can also control the port quota in the fixed port range IP pool. The number of ports for each user can be allocated from 32 to 60416, where 0 means default.

Monitoring NAT Sessions With Diagnose Commands

- `diagnose firewall ippool-all list`
  - Lists all the configured NAT IP pools with NAT IP range and type

```
Local-FortiGate # diagnose firewall ippool-all list
vdom:root owns 1 ippool(s)
name:INTERNAL-HOST-EXT-IP
type:overload
nat-ip-range:10.200.1.100-10.200.1.100
...........
...........
```

Command lists all configured IP Pools

NSE Training Institute

47

You can use the `diagnose firewall ippool-all list` command, which lists all of the configured NAT IP pools with their NAT IP range and type.

Monitoring NAT Sessions With Diagnose Commands (Contd)

- `diagnose firewall ippool-all stats <Optional IP Pool name>`
  - Lists stats for all of the IP pools:
    - 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

```
# diagnose firewall ippool-all stats EXT
name: EXT
type: overload
startip: 10.200.1.100
endip: 10.200.1.100
total ses: 100
tcp ses: 75
udp ses: 20
other ses: 5
```

Command show only stats of IP pool named EXT

```
# diagnose firewall ippool-all stats
vdom:root owns 2 ippool(s)
name: EXT
type: overload
startip: 10.200.1.100
endip: 10.200.1.100
total ses: 100
tcp ses: 75
udp ses: 20
other ses: 5

name: Training
type: one-to-one
startip: 10.200.1.50
endip: 10.200.1.60
total ses: 10
tcp ses: 8
udp ses: 2
other ses: 0
```

Command shows stats of all IP pools

NSE Training Institute

48

The `diagnose firewall ippool-all stats` shows the stats for all of the 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.

VIP Optional Filters

- Enabling the **Services** option:
  - Allows complex scenarios where multiple external sources of traffic use multiple services to connect to a single internal server
  - Avoids the requirement for numerous VIPs to be bundled into VIP groups

```
config firewall vip
    edit "WebServer-Services"
        set uuid dbb43cc8-cc25-51ea-748a-44d6a7c0f659
        set service "TCP_8090" "TCP_8091" "TCP_8092"
        set extip 203.0.113.10
        set extintf "any"
        set portforward enable
        set mappedip "10.0.1.10"
        set mappedport 80
    next
end
```

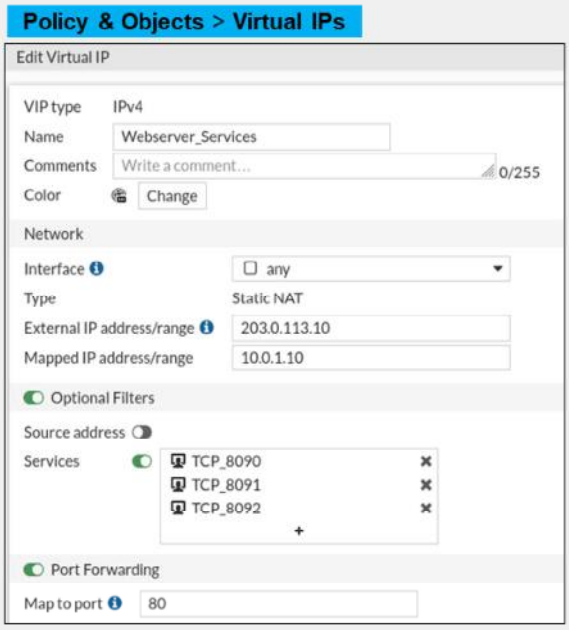

The **Services** option has been added to VIP objects. Virtual IP with services is a more flexible virtual IP mode. This mode allows users to define services to a single port number mapping.

This configuration was made possible to allow for complex scenarios where multiple sources of traffic are using multiple services to connect to a single computer, while requiring a combination of source and destination NAT, and not requiring numerous VIPs to be bundled into VIP groups.

In the example shown on this slide, TCP ports 8090, 8091, and 8092 are mapped to an internal webserver, TCP port 80. This allows remote connections to communicate with a server behind the firewall.

Once you apply the virtual IP configuration shown on the slide to the firewall policy, if a user accesses `203.0.113.10:8090` from external network, FortiGate maps it to `10.0.1.10:80` in the internal network. Similarly, if a user accesses `203.0.113.10:8091` or `203.0.113.10:8092` from an external network, FortiGate maps it to 10.0.1.10:80 in the internal network.

VIPs with different services are considered non-overlapping.

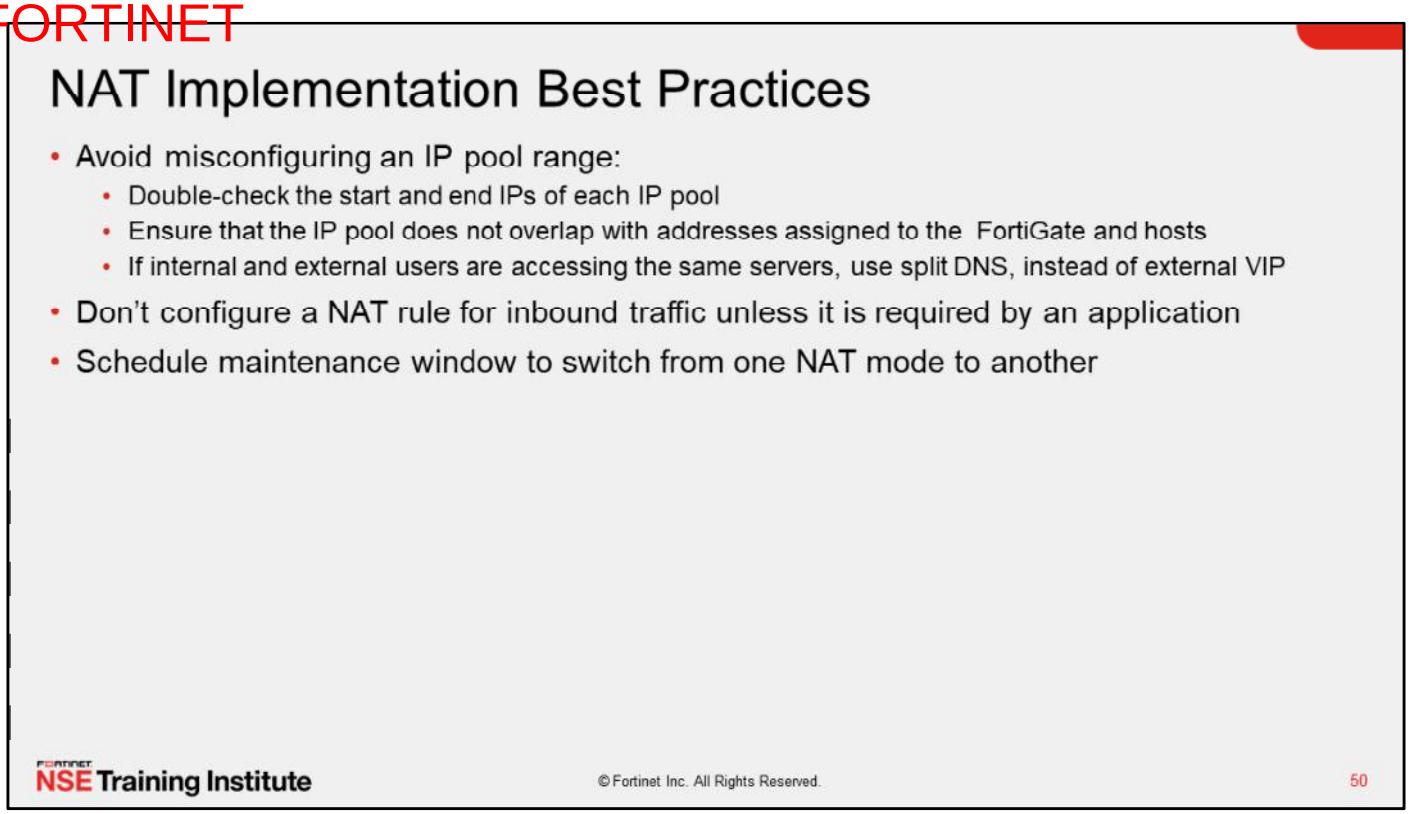

Use the following best practices when implementing NAT:

- Avoid misconfiguring an IP pool range:
    - Double-check the start and end IPs of each IP pool.
    - Ensure that the IP pool 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, use split DNS to offer an internal IP to internal users so that they don't have to use the external-facing 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`

51

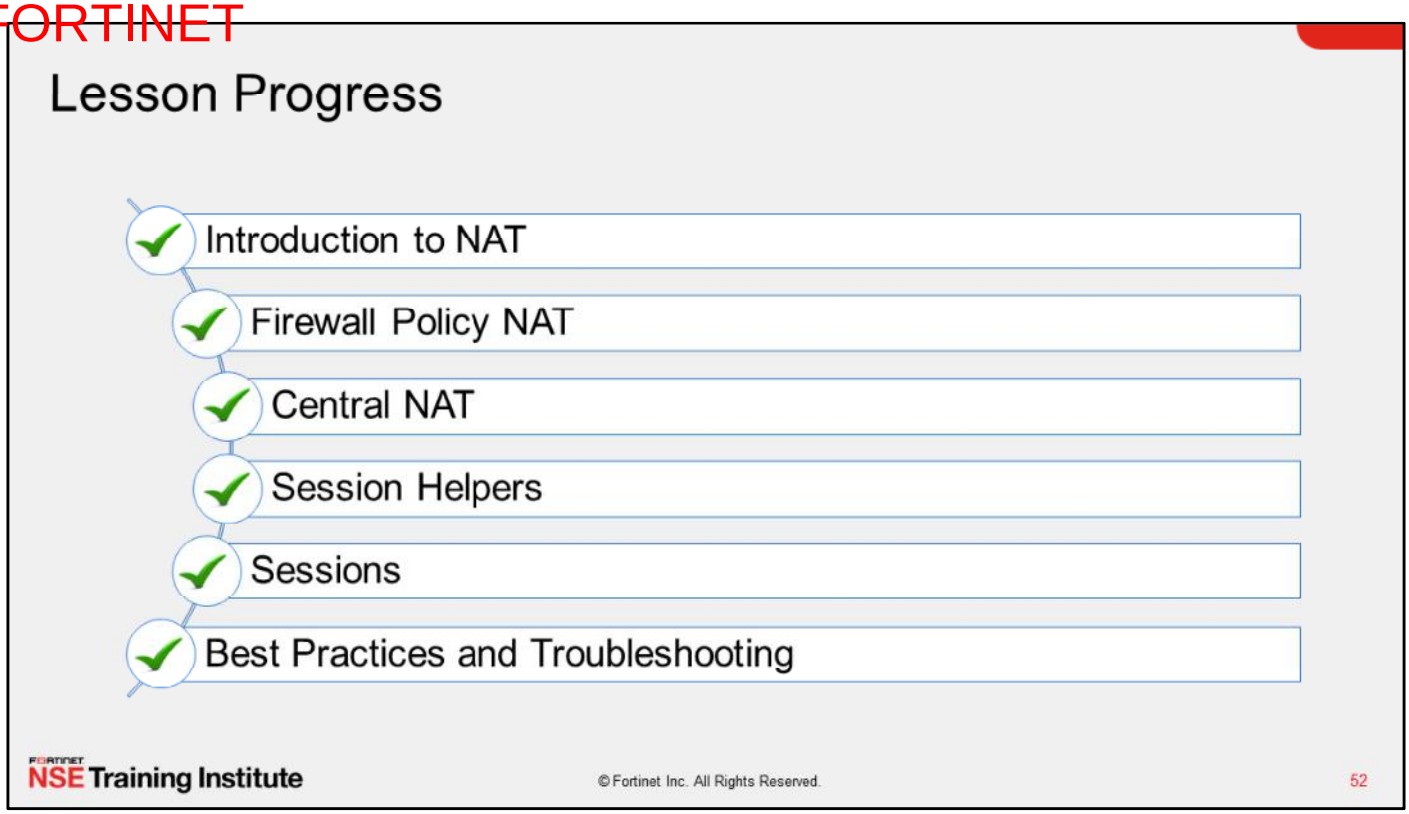

Congratulations! You have completed this lesson.

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

Review

- ✓ Understand NAT and PAT
- ✓ Understand the different configuration modes for NAT
- ✓ Configure a firewall policy to perform SNAT and DNAT (VIPs)
- ✓ Configure central NAT
- ✓ Understand session helpers and use a SIP session helper for VoIP
- ✓ Understand and interpret the session table
- ✓ Analyze the session diagnose command output
- ✓ Understand TCP, UDP, and ICMP states
- ✓ Use traffic logs to identify common NAT issues and monitor NAT sessions using session diagnose commands
- ✓ Use NAT implementation best practices

NSE Training Institute © Fortinet Inc. All Rights Reserved. 53

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.

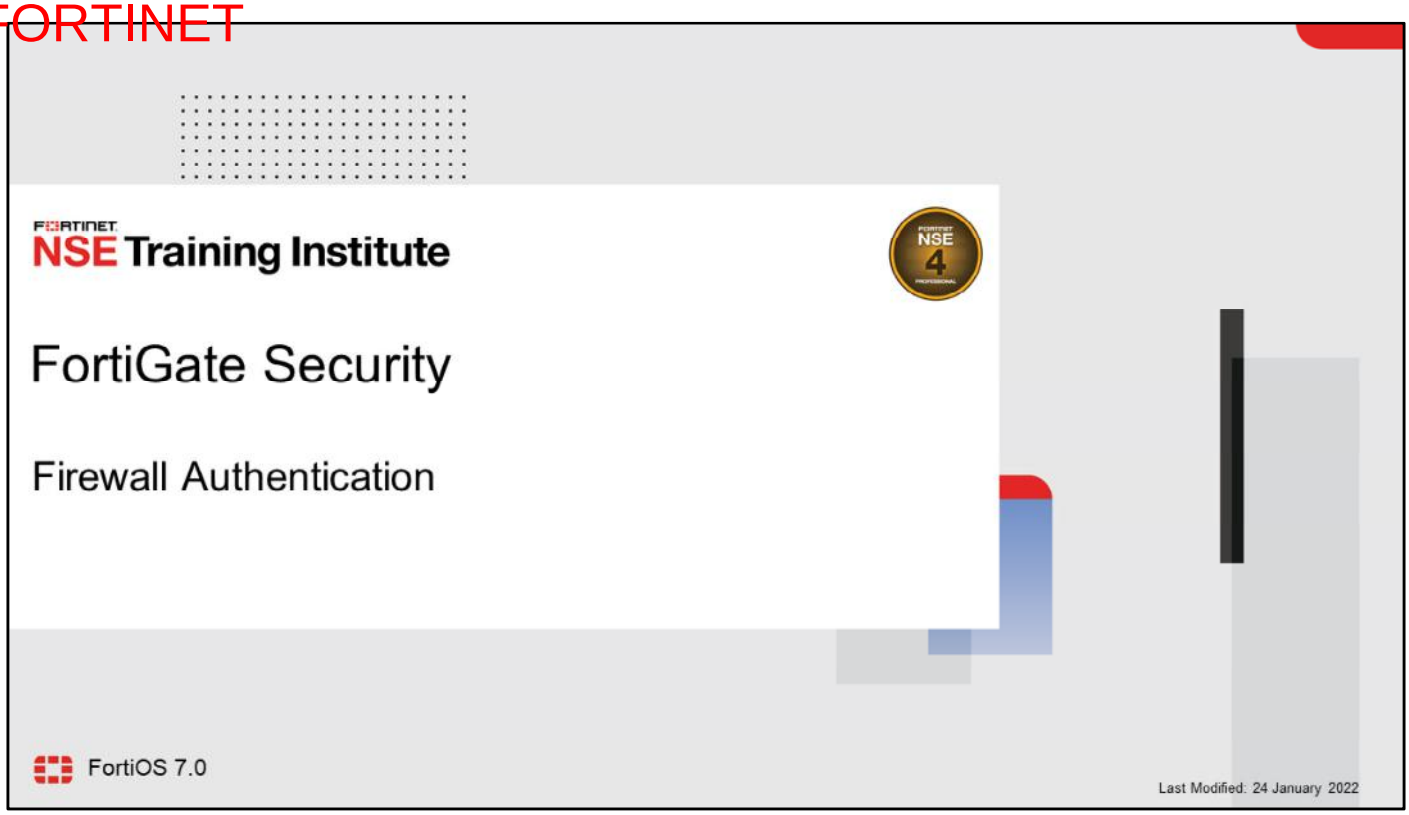

In this lesson, you will learn about using authentication on the firewall policies of FortiGate.

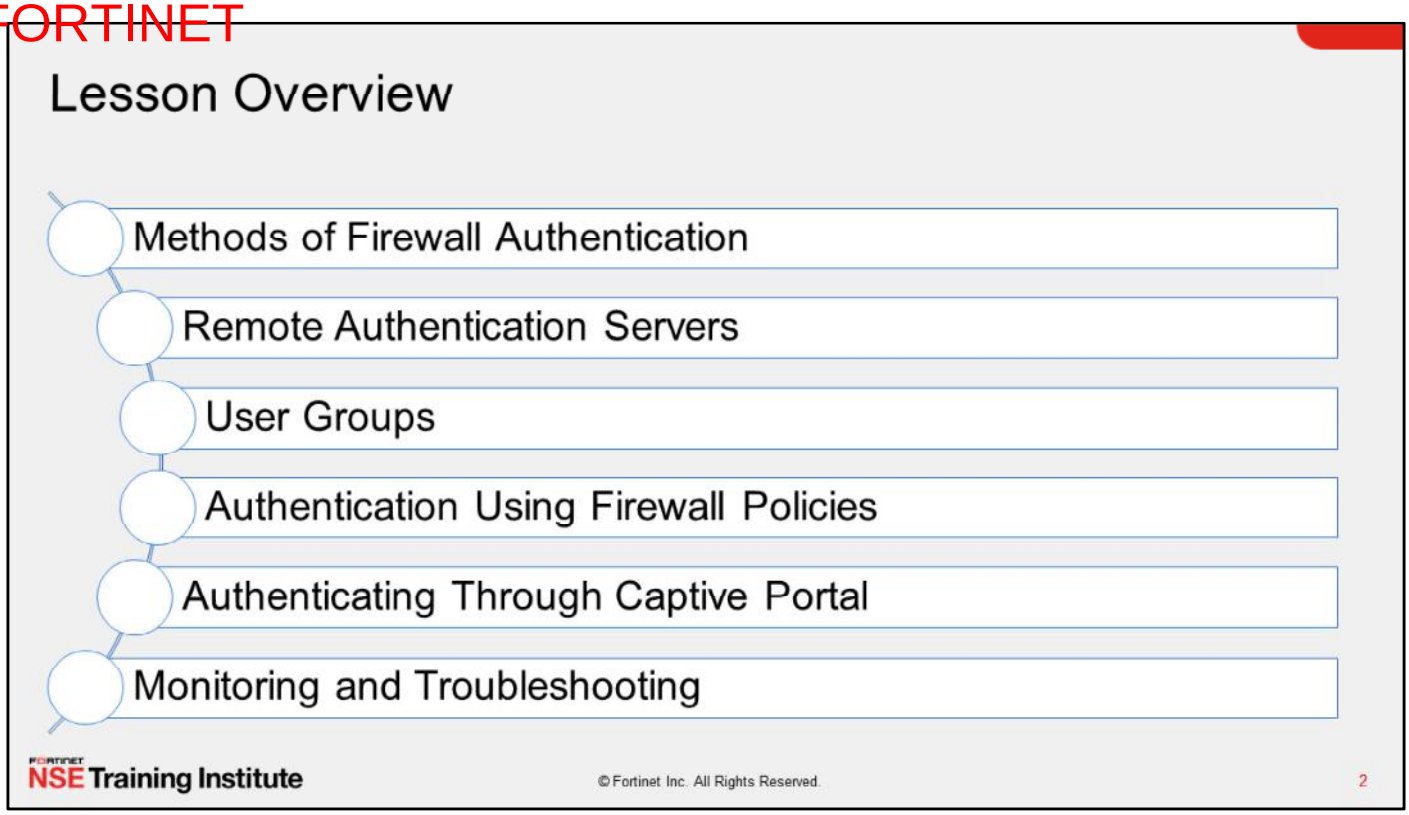

In this lesson, you will learn about the topics shown on this slide.

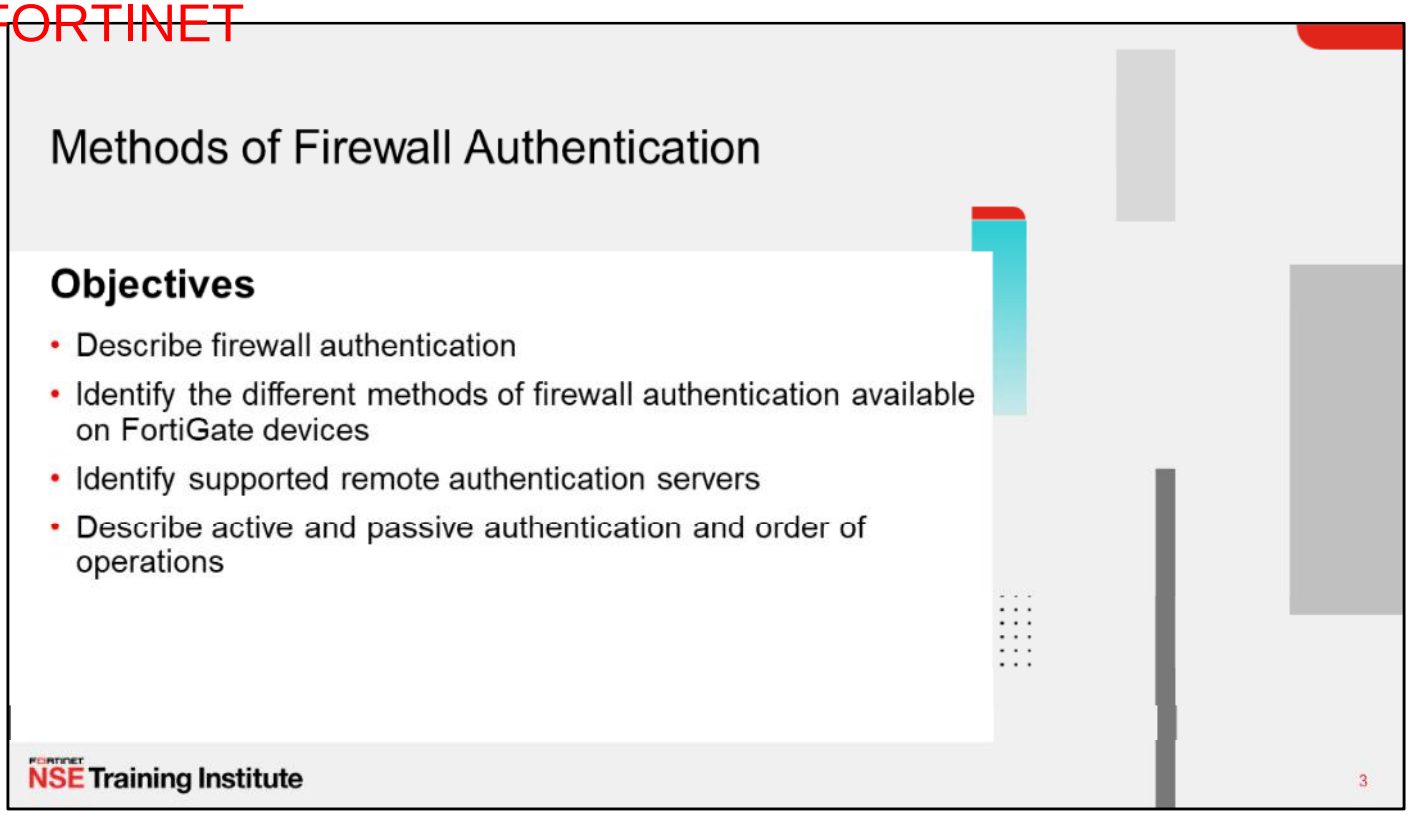

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

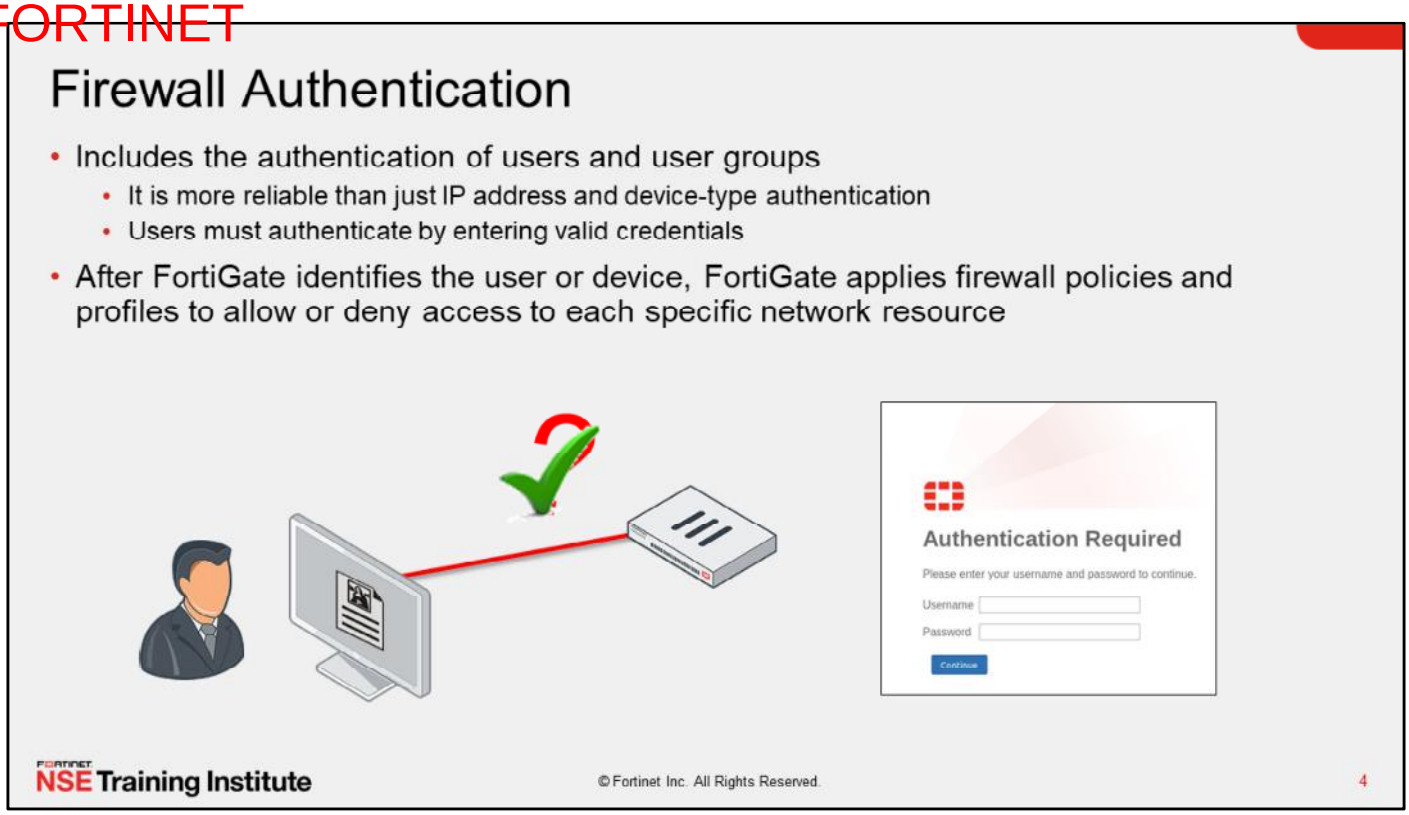

Traditional firewalling grants network access by verifying the source IP address and device. This is inadequate and can pose a security risk, because the firewall cannot determine who is using the device to which it is granting access.

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 Methods of Firewall Authentication

- Local password authentication
  - Username and password stored on FortiGate
- Server-based password authentication (also called remote password authentication)
  - Password stored on a POP3, RADIUS, LDAP, or TACACS+ server
- Two-factor authentication
  - Enabled on top of an existing method
  - Requires something you know and something you have (token or certificate)

NSE Training Institute

5

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.

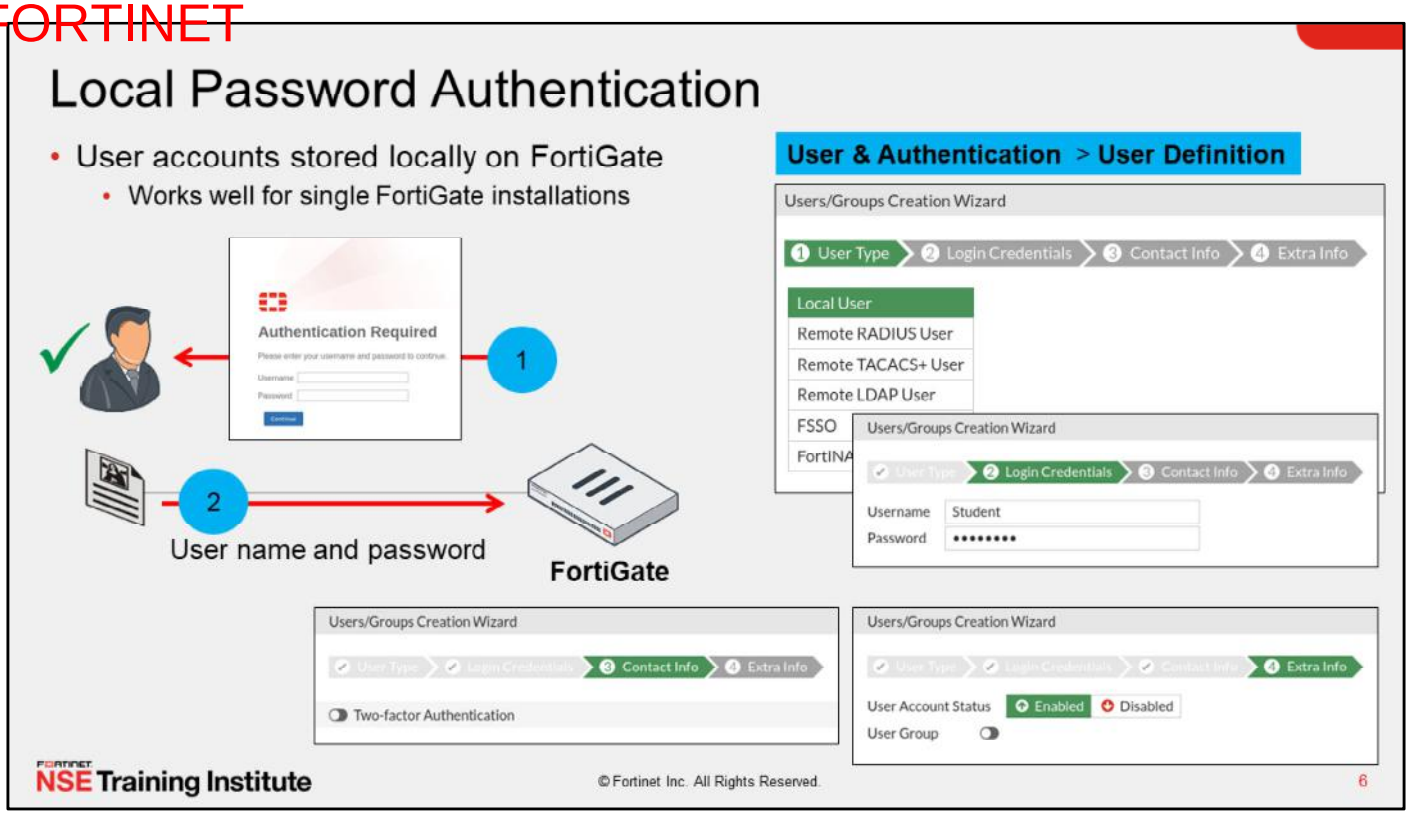

The simplest method of authentication is local password authentication. User account information (username and password) is stored locally on the FortiGate device. This method works well for a single FortiGate installation.

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.

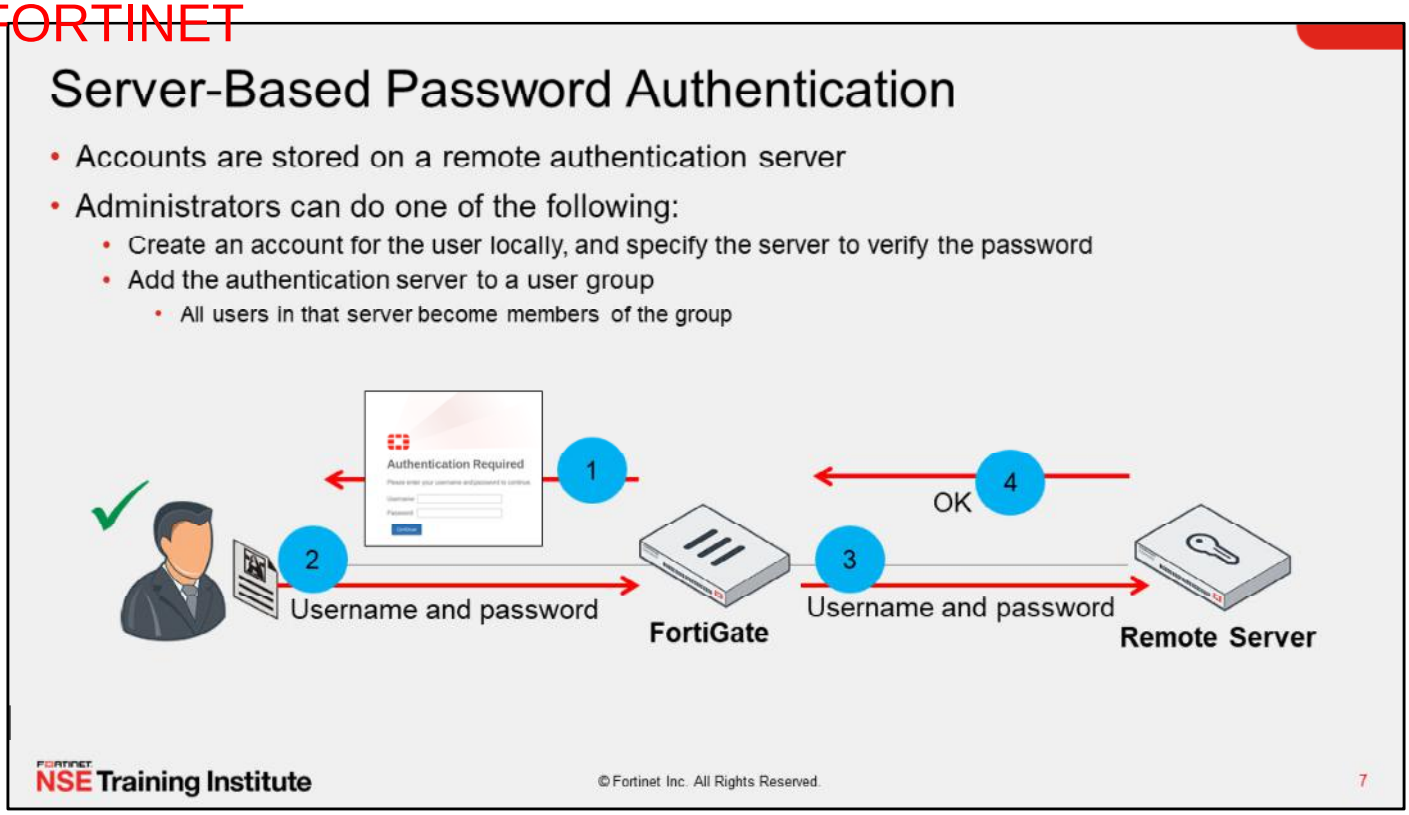

When server-based password authentication is used, a remote authentication server authenticates users. This method is desirable when multiple FortiGate devices need to authenticate the same users or user groups, or when adding FortiGate to a network that already contains an authentication server.

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.

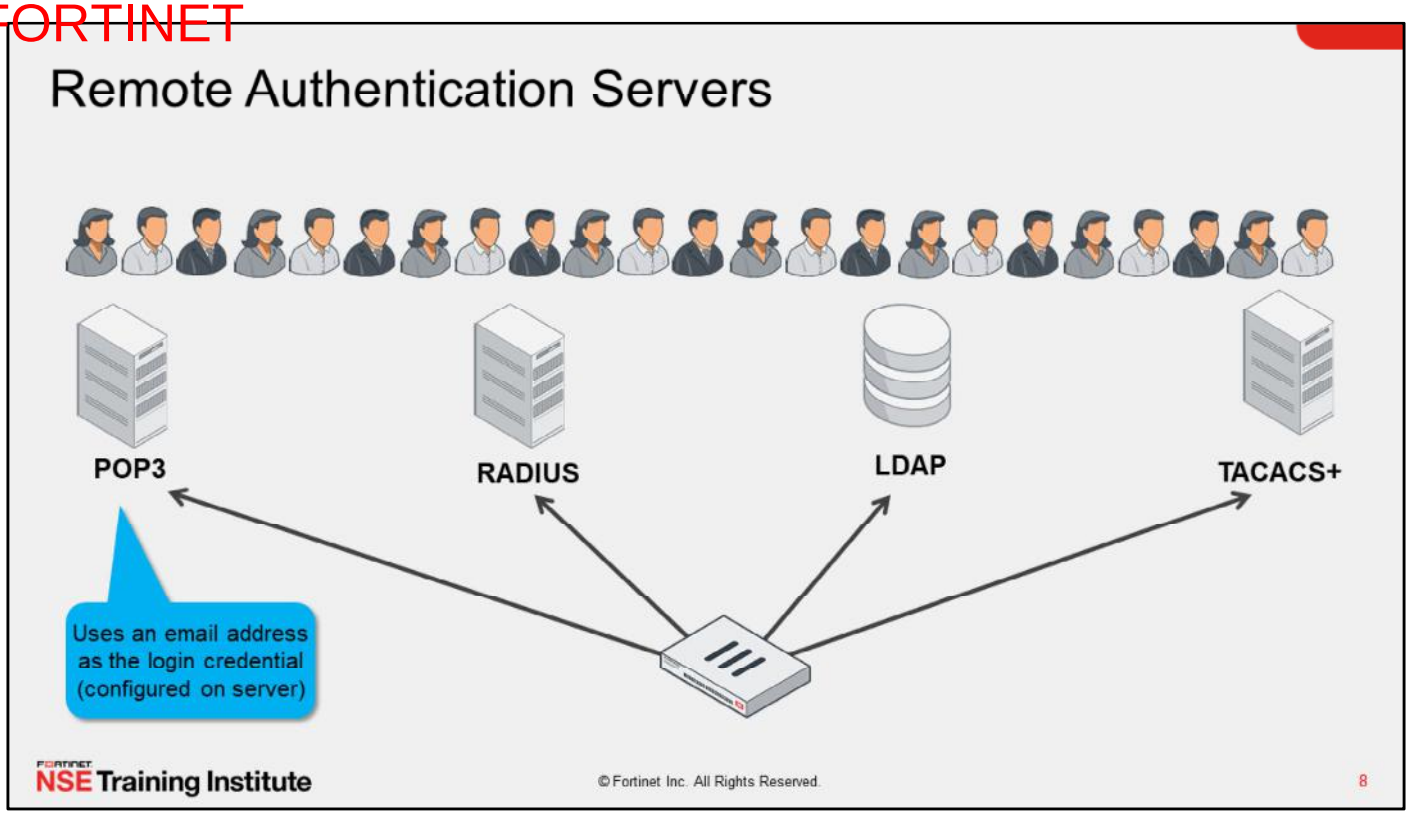

FortiGate provides support for many remote authentication servers, including POP3, RADIUS, LDAP, and TACACS+.

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.

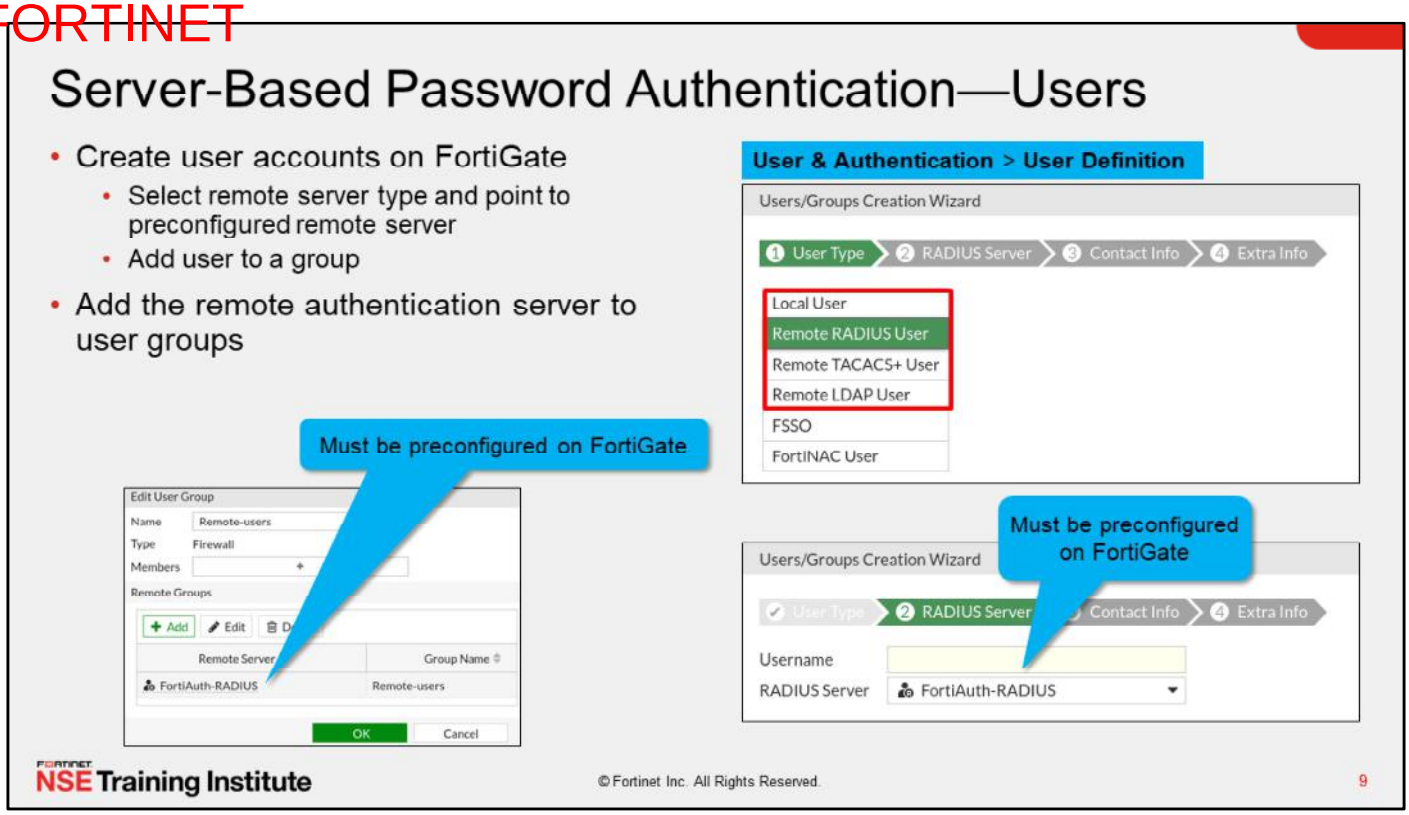

You can configure FortiGate to use external authentication servers in the following two ways:

- 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 server, 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.

Two-Factor Authentication and One-Time Passwords

- Strong authentication that improves security by preventing attacks associated with the use of static passwords alone
- Requires two independent methods of identifying a user:
  - Something you know, such as password or PIN
  - Something you have, such as a token or certificate
- One-time passwords (OTPs) can be used one time only
  - OTPs are more secure than static passwords
- Available on both user and administrator accounts
  - The user or user group is added to a firewall policy in order to authenticate
- Methods of OTP delivery include:
  - FortiToken 200 or FortiToken Mobile
    - Generates a six-digit code every 60 seconds based on a unique seed and GMT time
  - Email or SMS
    - An OTP is sent to the user's email or SMS
    - Email or SMS must be configured on the user's account
  - FortiToken mobile push
    - Supports two-factor authentication without requiring user to enter code
- NTP server recommended!

NSE Training Institute © Fortinet Inc. All Rights Reserved. 10

Traditional user authentication requires your user name plus something you know, such as a password. The weakness with this traditional method of authentication is that if someone obtains your user name, 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.

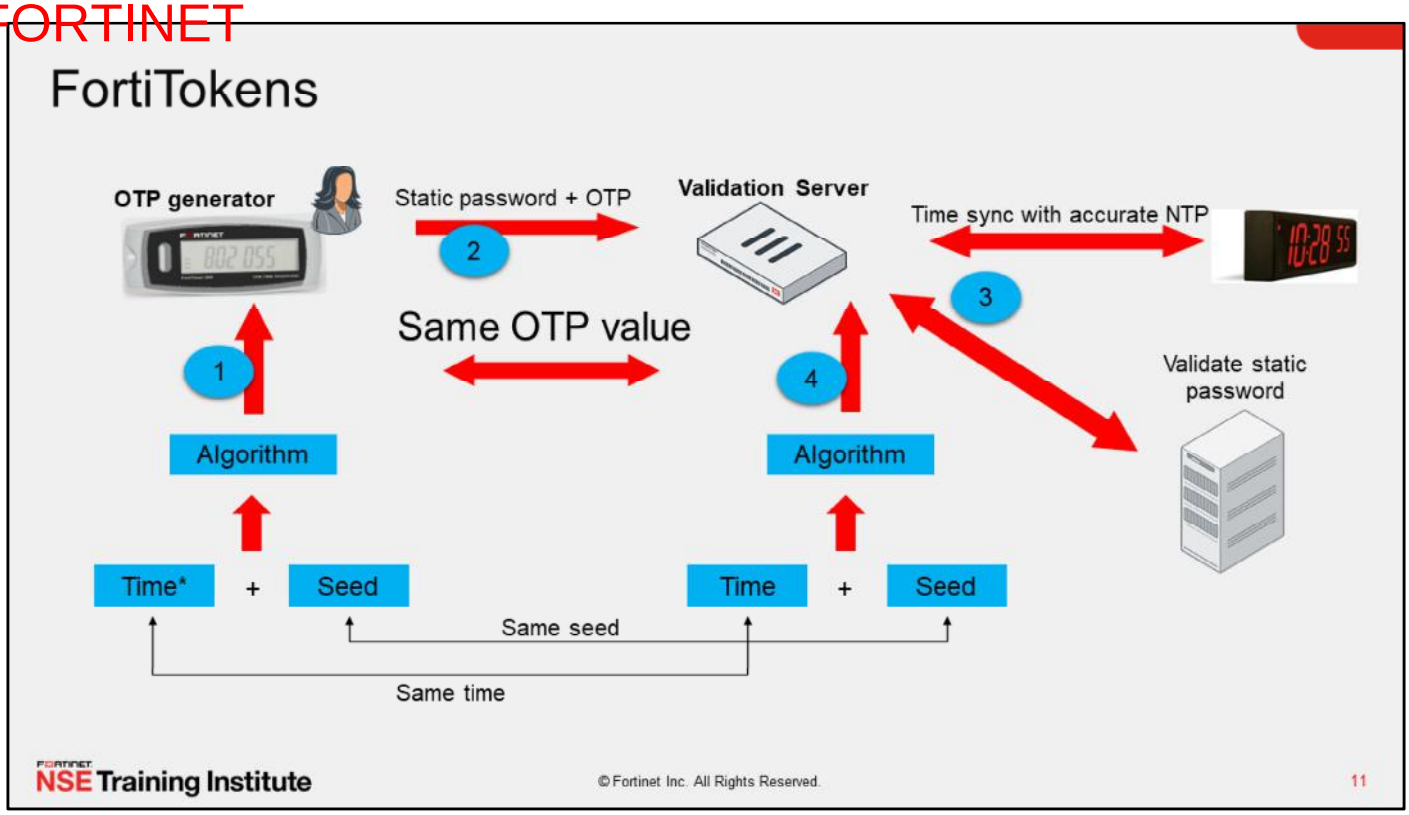

Tokens use a specific algorithm to generate an OTP. The algorithm consists of:

- 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).

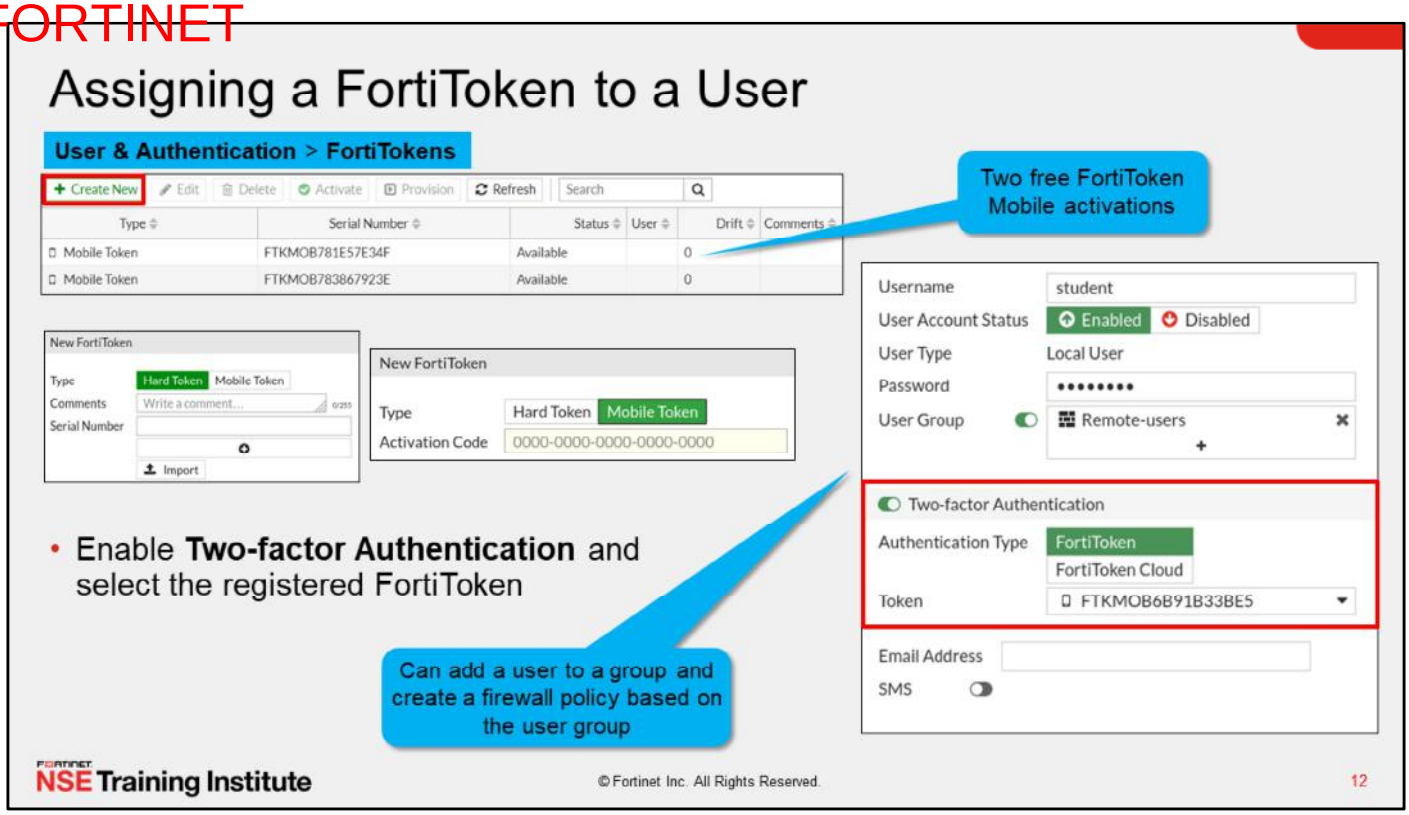

You can add a FortiToken 200 or FortiToken Mobile to FortiGate on the **FortiTokens** page.

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.

Authentication Methods and Active Authentication

- Active
  - User receives a login prompt
  - Must manually enter credentials to authenticate
  - POP3, LDAP, RADIUS, Local, and TACACS+
- Passive
  - User does not receive a login prompt from FortiGate
  - Credentials are determined automatically
    - Method varies depending on type of authentication used
  - FSSO, RSSO, and NTLM

NSE Training Institute © Fortinet Inc. All Rights Reserved. 13

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. Which type of token can generate OTPs to provide two-factor authentication to users in your network?
   - ✓ A. FortiToken Mobile
   - B. USB FortiToken

NSE Training Institute

14

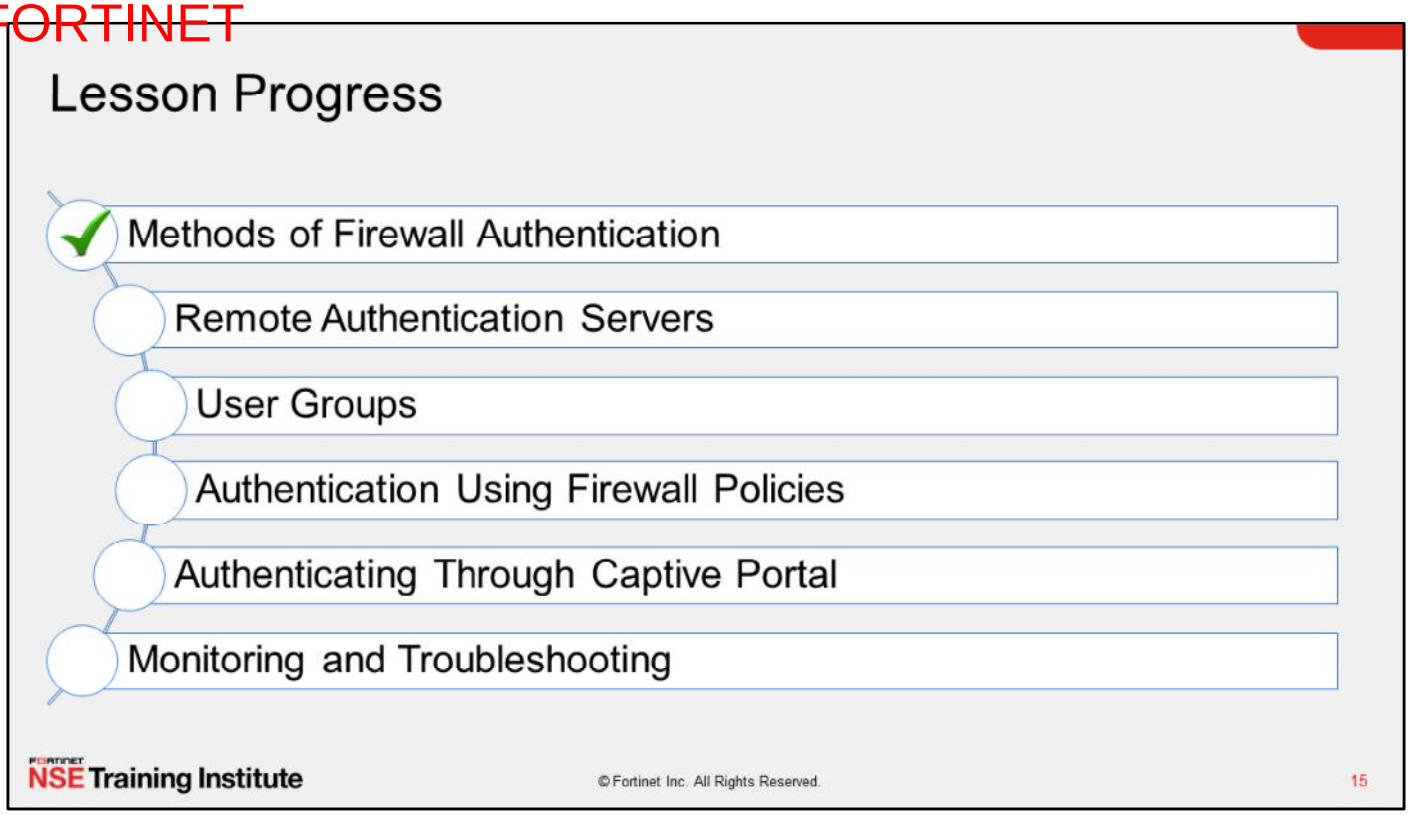

Good job! You now understand the basics of firewall authentication.

Now, you will learn about remote authentication servers.

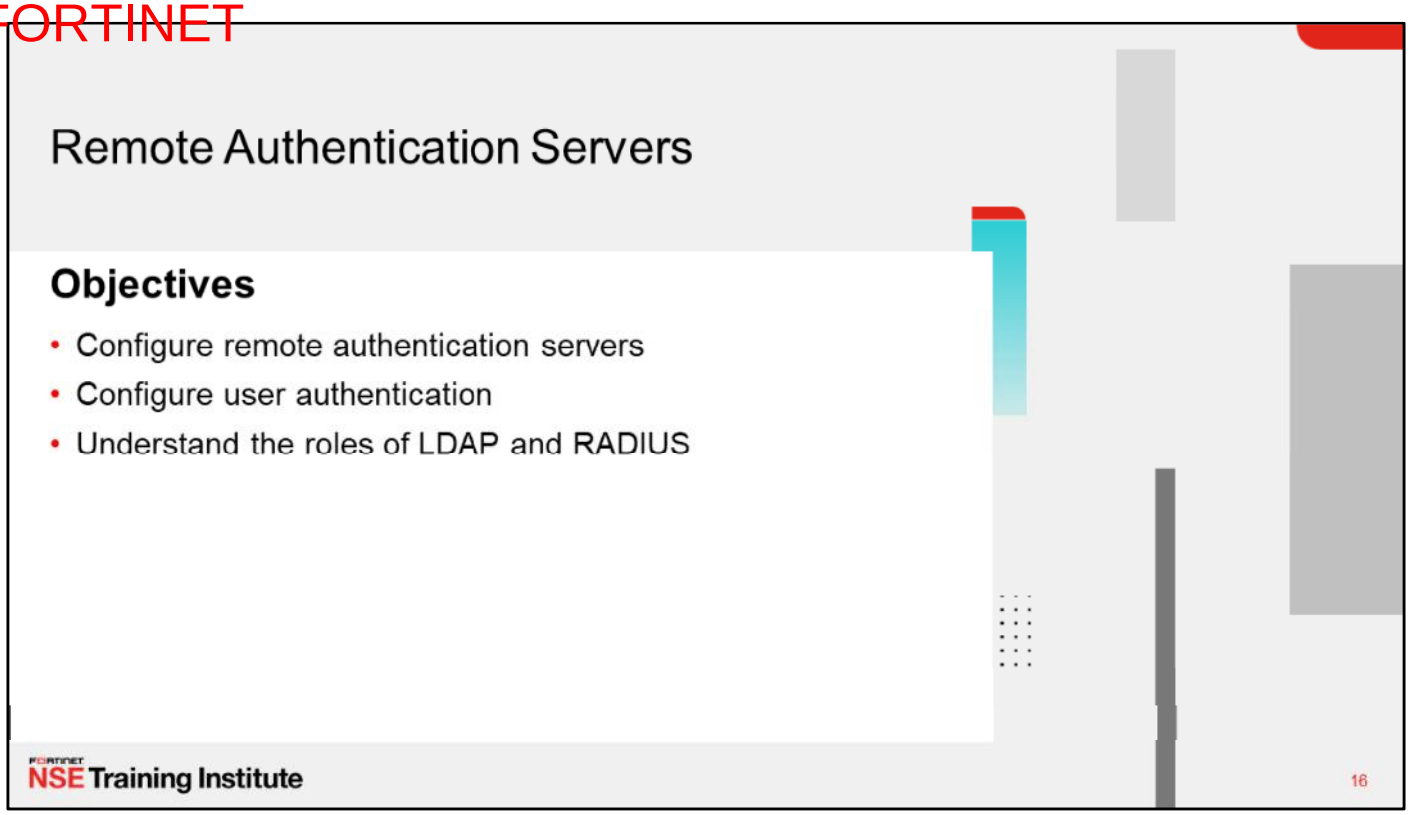

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in remote authentication servers, you will be able to configure firewall authentication using remote user accounts defined on a remote authentication server.

LDAP Overview

- LDAP is an application protocol for accessing and maintaining distributed directory information services

- LDAP maintains authentication data, including:
  - Departments, people (and 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
- Binding is the operation in which the LDAP server authenticates the user

NSE Training Institute 17

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.

LDAP Directory Tree

- The LDAP structure is similar to a tree that contains entries (objects) in each branch
- Each entry has a unique ID: the distinguished name (DN)
- Each DN has attributes
- Each attribute has a name and one or more values
- The attributes are defined in the directory schema

NSE Training Institute

18

The root of the LDAP directory tree represents the organization itself, and is defined as a domain component (DC). The DC is usually a DNS domain, such as example.com. (Because the name contains a dot, it is written as two parts separated by a comma: dc=example,dc=com.) You can add additional entries, known as objects, to the hierarchy as needed. Generally, two types of objects make up most entries: containers and leafs.

Containers are objects that can include other objects, similar to a folder in a file system. Example containers include:

- Country (represented as c)
- Organizational unit (represented as ou)
- Organization (represented as o)

Leafs are objects at the end of a branch and have no subordinate objects. Example leafs include:

- User ID (represented as uid)
- Common name (represented as cn)

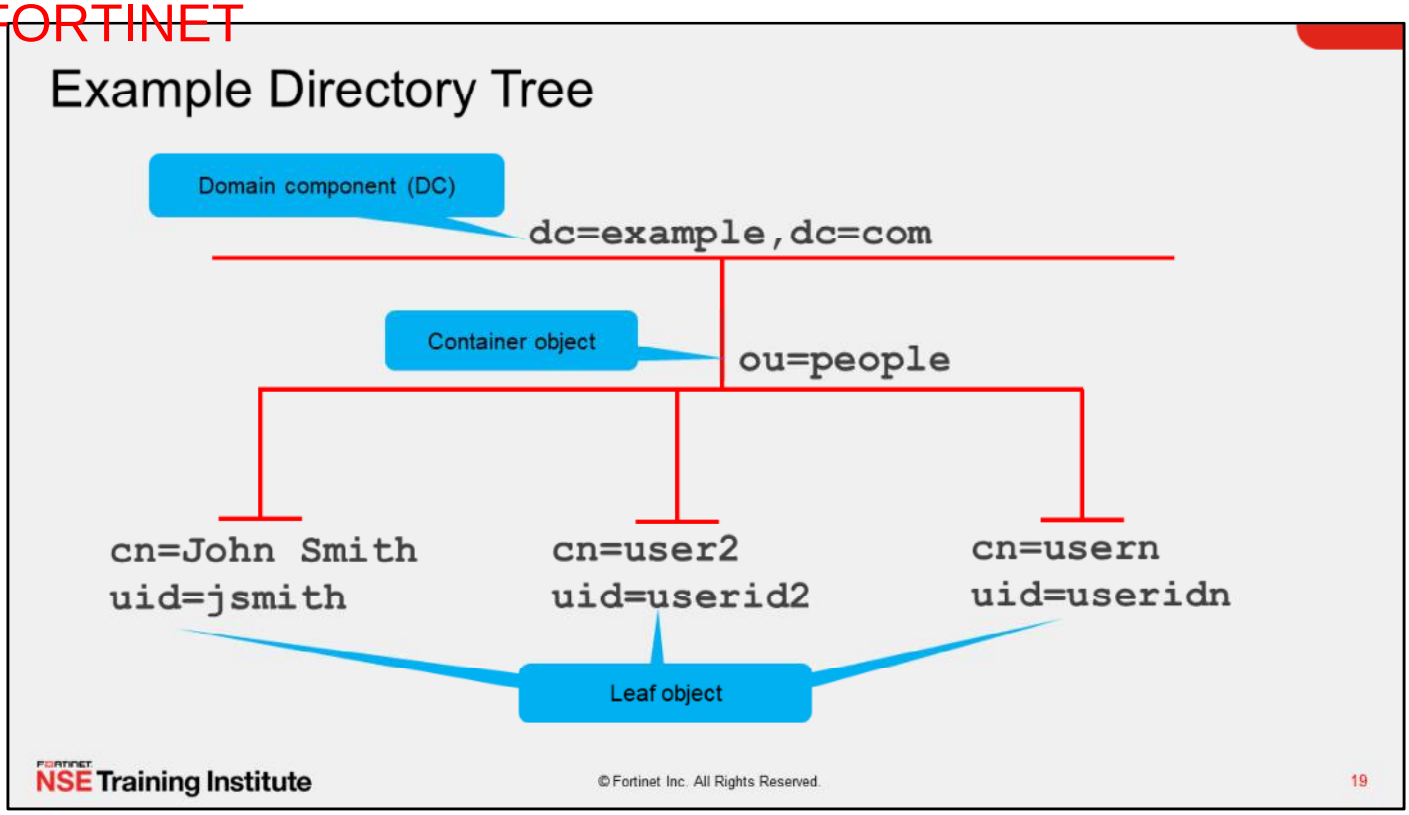

This slide shows an example of a simple LDAP hierarchy.

You must configured the FortiGate device (acting as an LDAP client) requesting authentication to address its request to the part of the hierarchy where user records exist: either the domain component, or a specific container where the record exists. Similar to users, containers have DNs, and in this example, the DN is `ou=people,dc=example,dc=com`.

The authentication request must also specify the user account entry. This can be one of many options including the common name (cn) or, on a computer network, the user ID (uid), which is the information users use to log in. Note that if the object name includes a space, such as John Smith, you must enclose the text with double quotes when testing in the CLI. For example: `cn="John Smith"`.

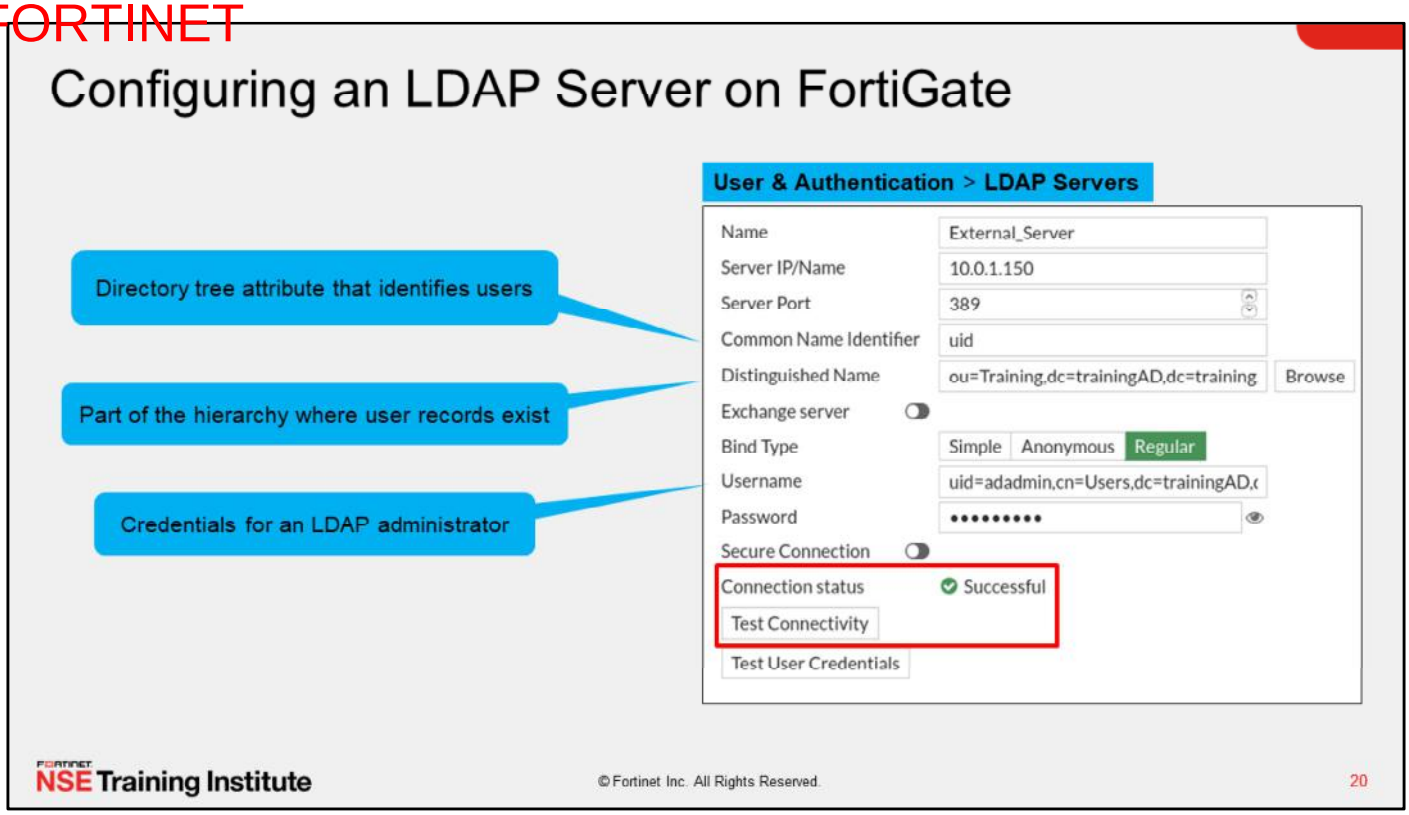

On the **LDAP Servers** page, you can configure FortiGate to point to an LDAP server for server-based password authentication. The configuration depends heavily on the server's schema and security settings. Windows Active Directory (AD) is very common.

The **Common Name Identifier** setting is the attribute name you use to find the user name. Some schemas allow you to use the attribute uid. 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.

Note that the **Test Connectivity** button only tests 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 you can use the CLI.

Testing the LDAP Query on the CLI

- `diagnose test authserver ldap <server_name> <username> <password>`
- Example:

```
# diagnose test authserver ldap External_Server aduser1 Training!

authenticate 'aduser1' against 'External_Server' succeeded!
Group membership(s) - CN=AD-users,OU=Training,DC=trainingAD,DC=training,DC=lab
```

NSE Training Institute

21

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.

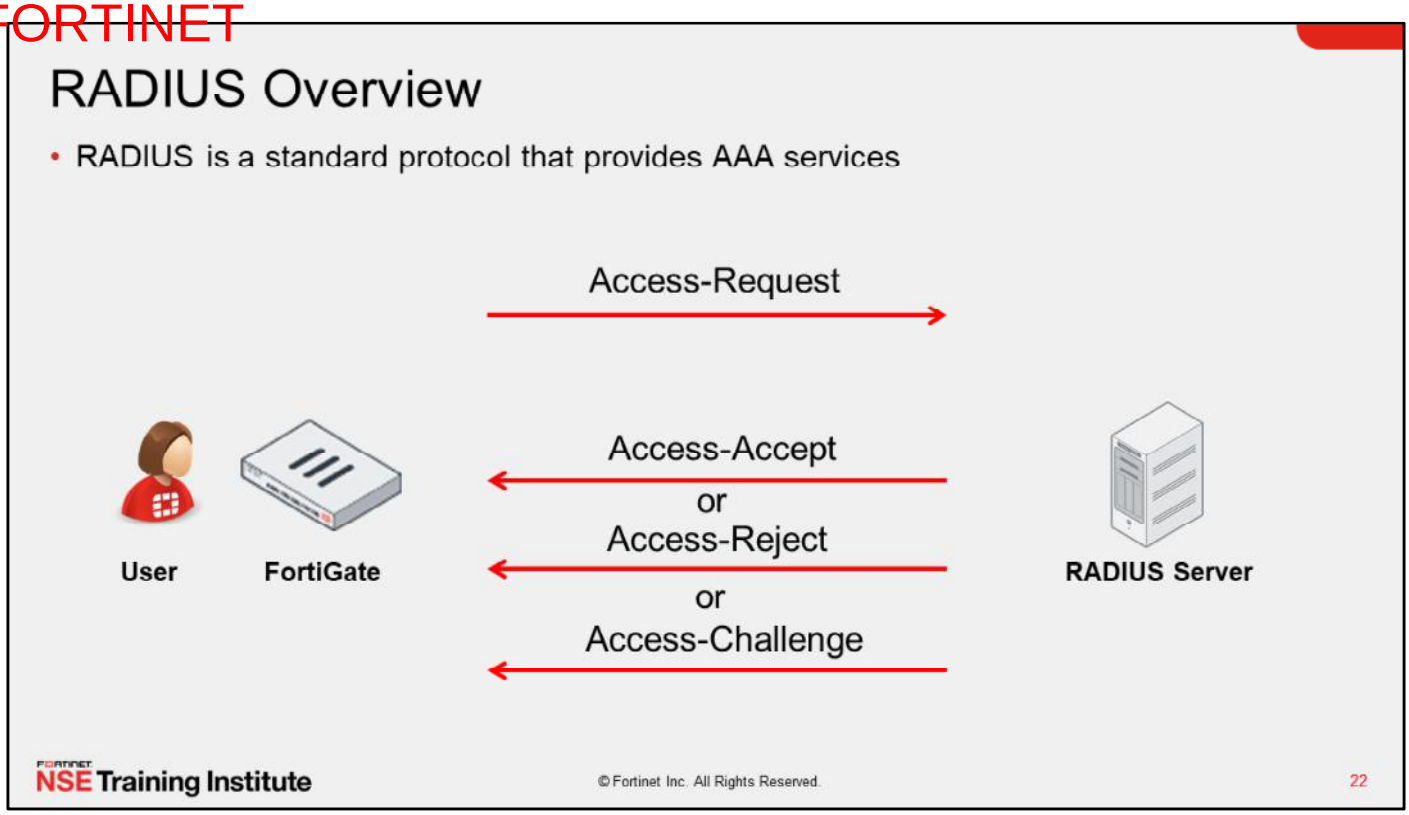

RADIUS is much different from LDAP, because there is no directory tree structure to consider. RADIUS is a standard protocol that provides authentication, authorization, and accounting (AAA) services.

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.

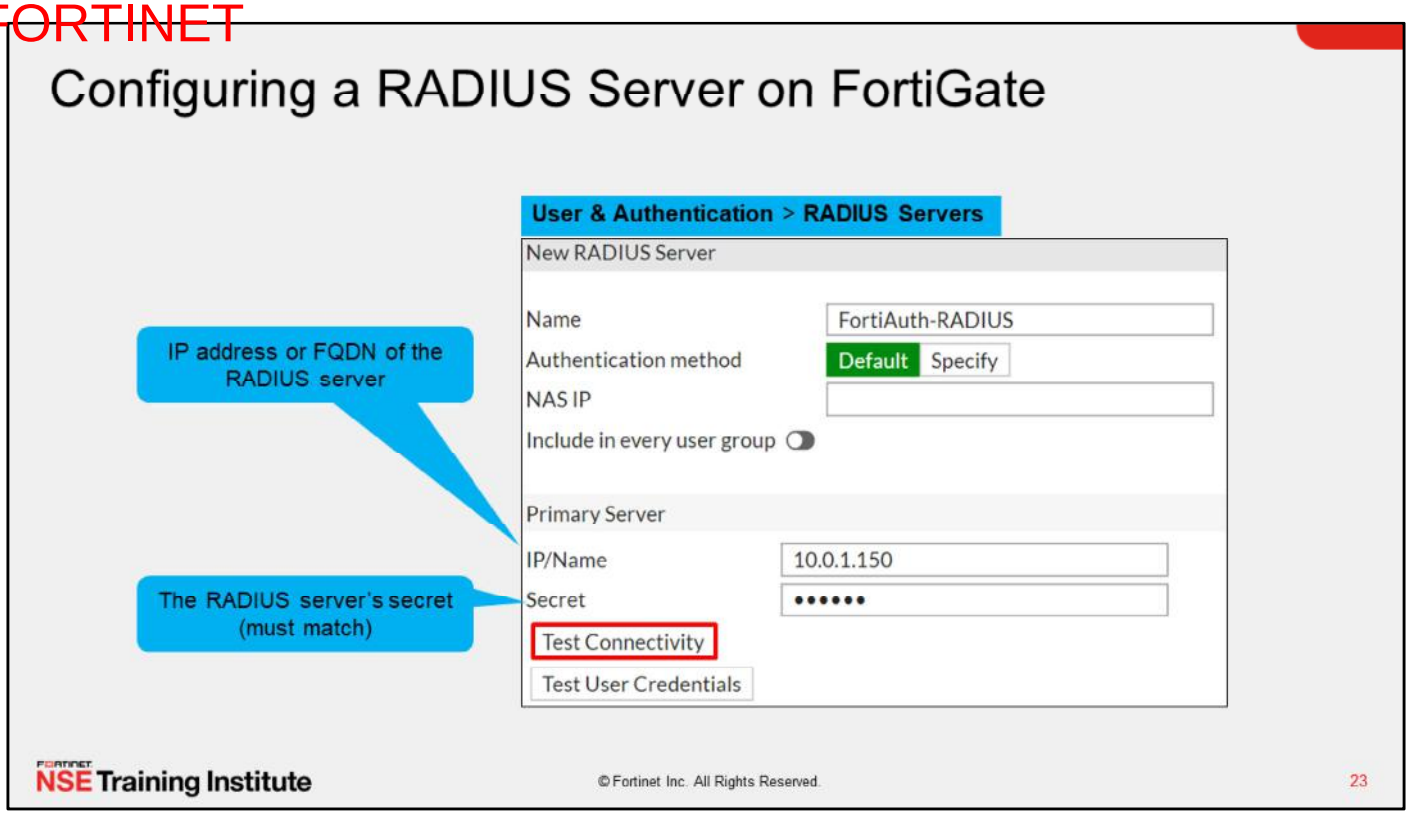

You can configure FortiGate to point to a RADIUS server for server-based password authentication through the **RADIUS Servers** page.

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).

Testing RADIUS Queries

- `diagnose test authserver radius <server_name> <scheme> <user> <password>`
- Example:

```
# diagnose test authserver radius FortiAuth-RADIUS pap student fortinet

authenticate 'student' against 'pap' succeeded, server=primary
assigned rad session id=810153440 session timeout=0 secs!
Group membership(s) - remote-RADIUS-admins
```

Group memberships are provided by vendor-specific attributes configured on the RADIUS server

NSE Training Institute

24

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`).

Knowledge Check

1. 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.

2. What is a valid reply from a RADIUS server to an ACCESS-REQUEST packet from FortiGate?
   A. ACCESS-PENDING
   ✓ B. ACCESS-REJECT

3. 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.

NSE Training Institute

25

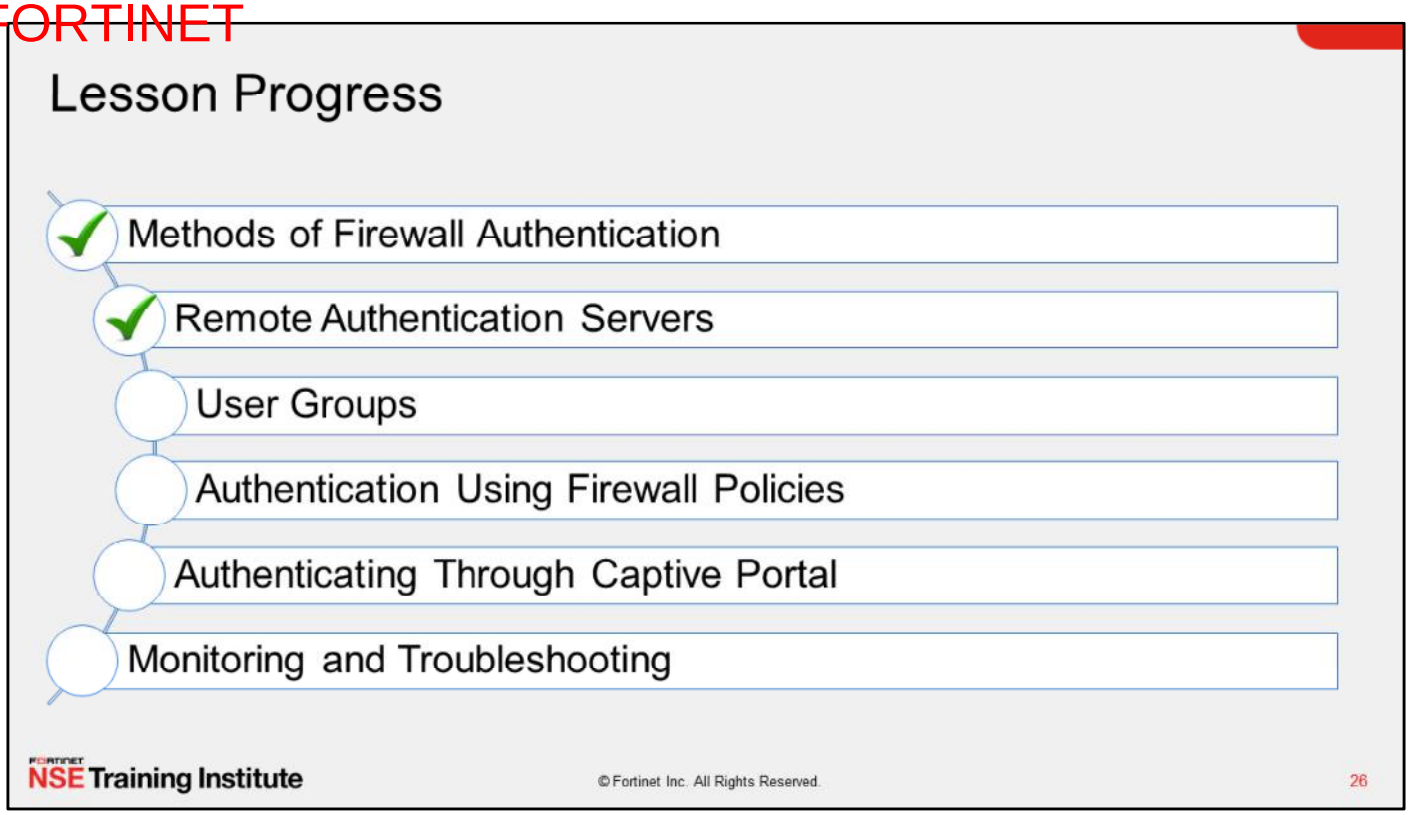

Good job! You now understand the basics of remote authentication servers.

Now, you will learn about user groups.

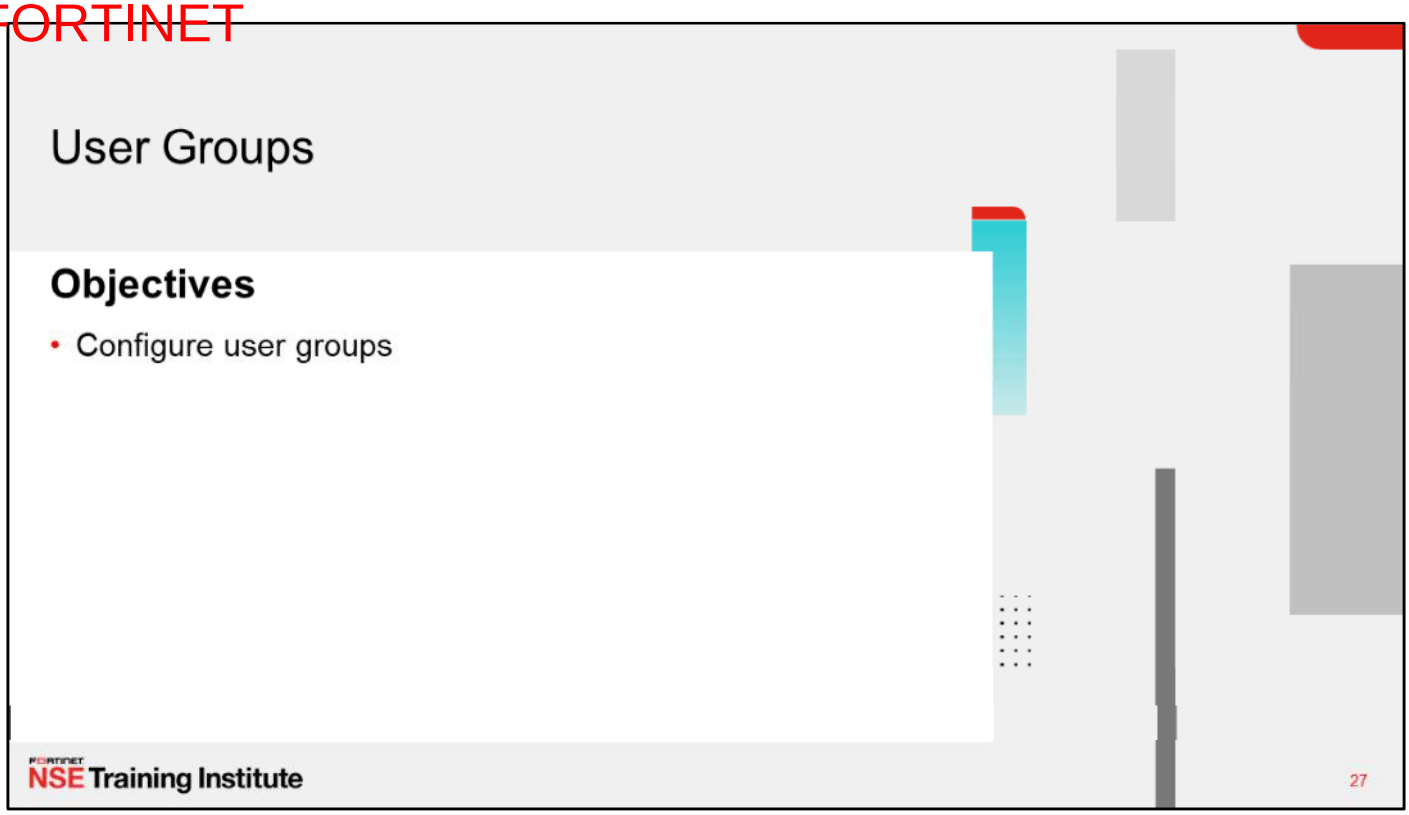

After completing this section, you should be able to achieve the objective shown on this slide.

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

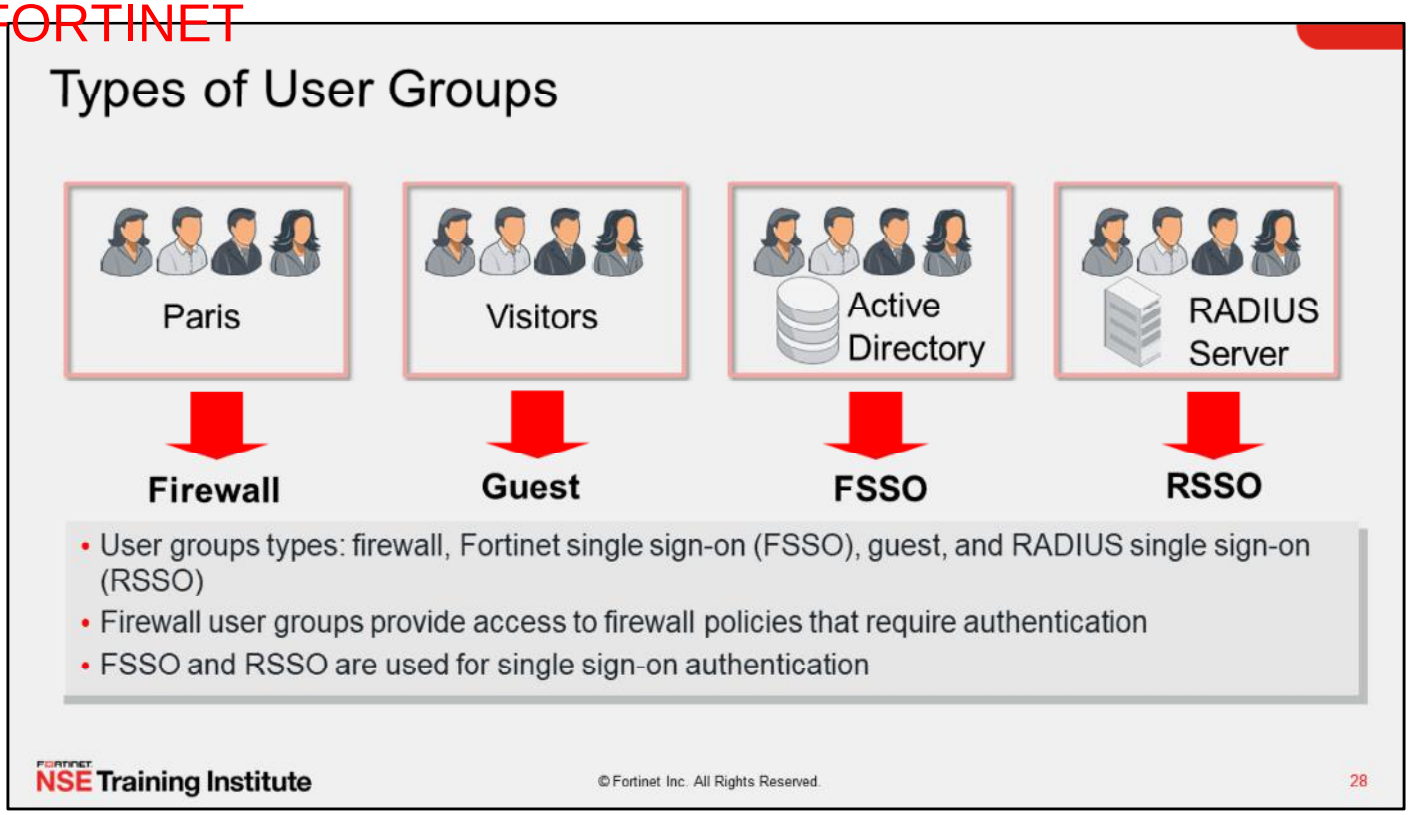

FortiGate allows administrators to assign users to groups. Usually, groups are used to more effectively manage individuals that have some kind of shared relationship. You might want to group employees by business area, such as finance or HR, or by employee type, such as contractors or guests.

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.

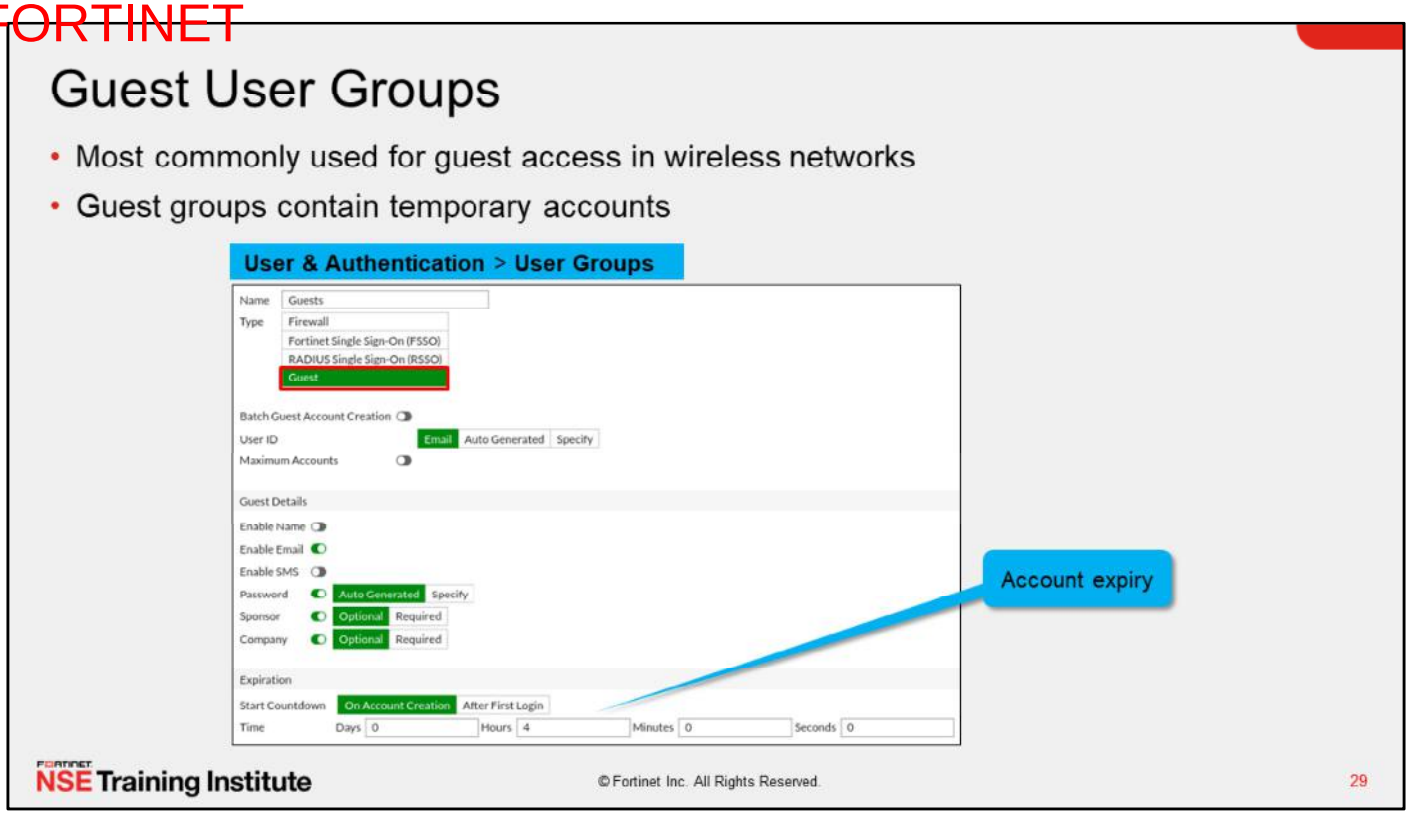

Guest user groups are different from firewall user groups because they contain exclusively temporary guest user accounts (the whole account, not just the password). Guest user groups are most commonly used in wireless networks. Guest accounts expire after a predetermined amount of time.

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 that have access only to create and manage guest user accounts.

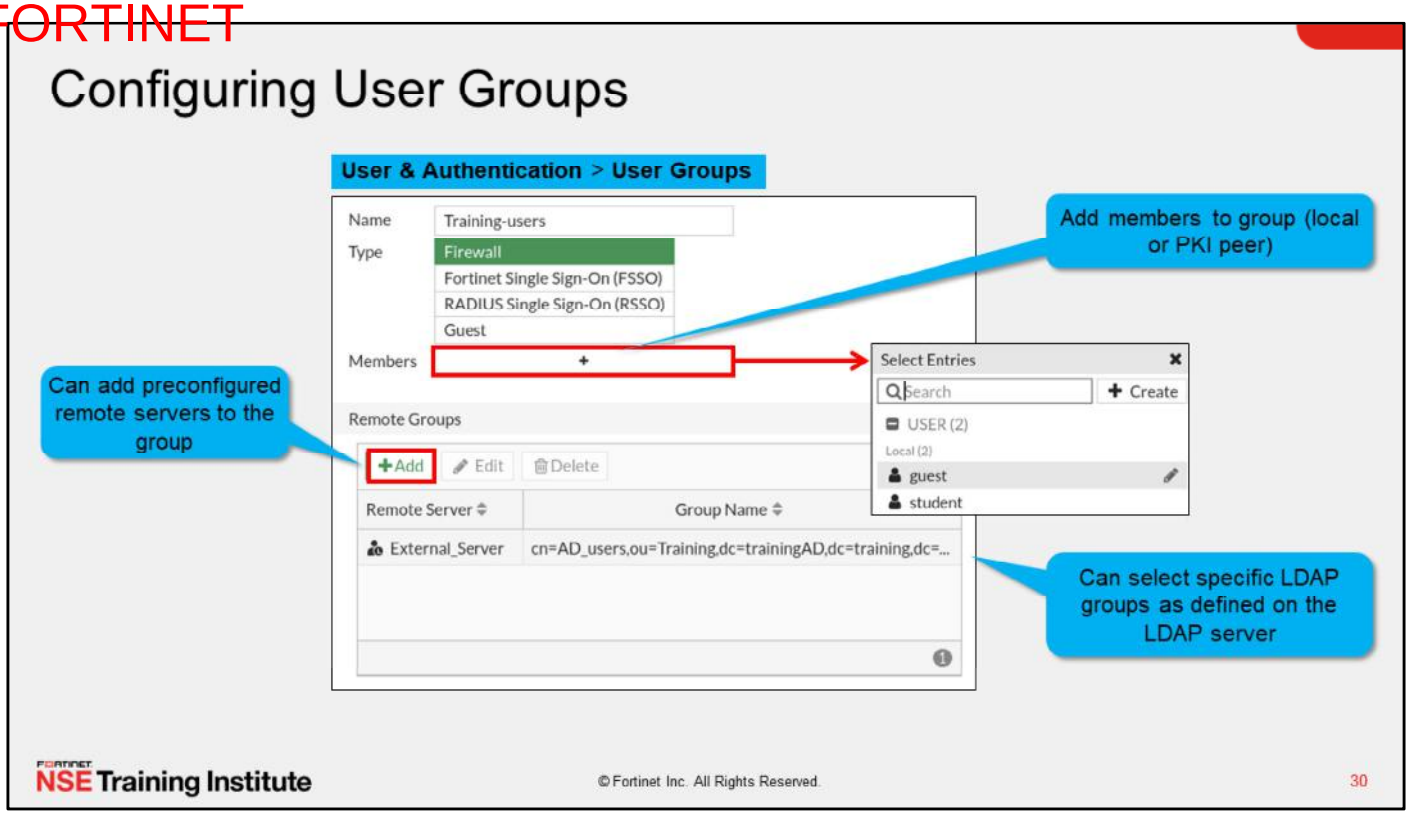

You can configure user groups on the **User Groups** page. You must specify the user group type and add users to the group. Depending on the group you create, you require different configurations. For the firewall user group, for example, members can consist of local users, PKI peer users, and users 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

NSE Training Institute

31

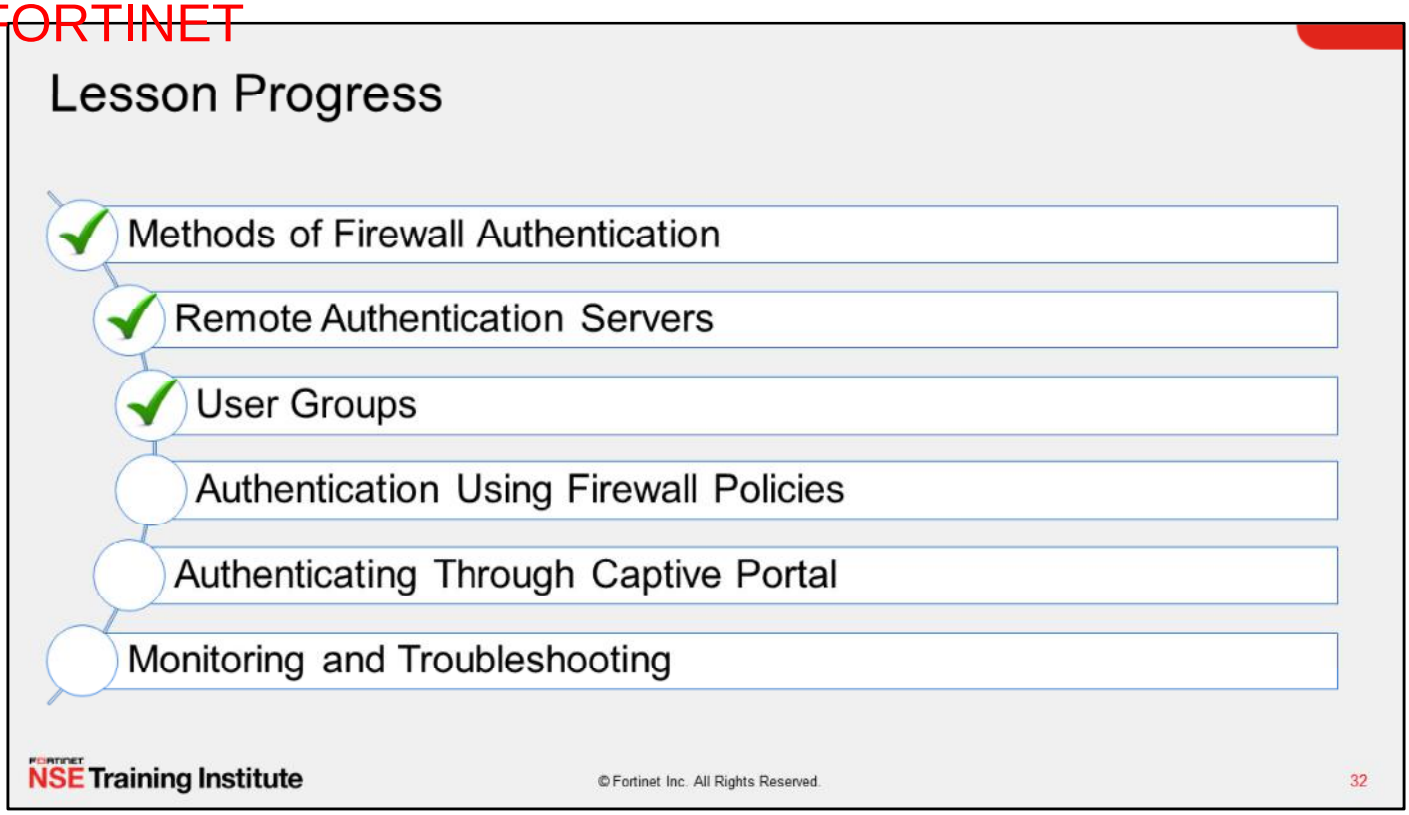

Good job! You now understand the basics of user groups.

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

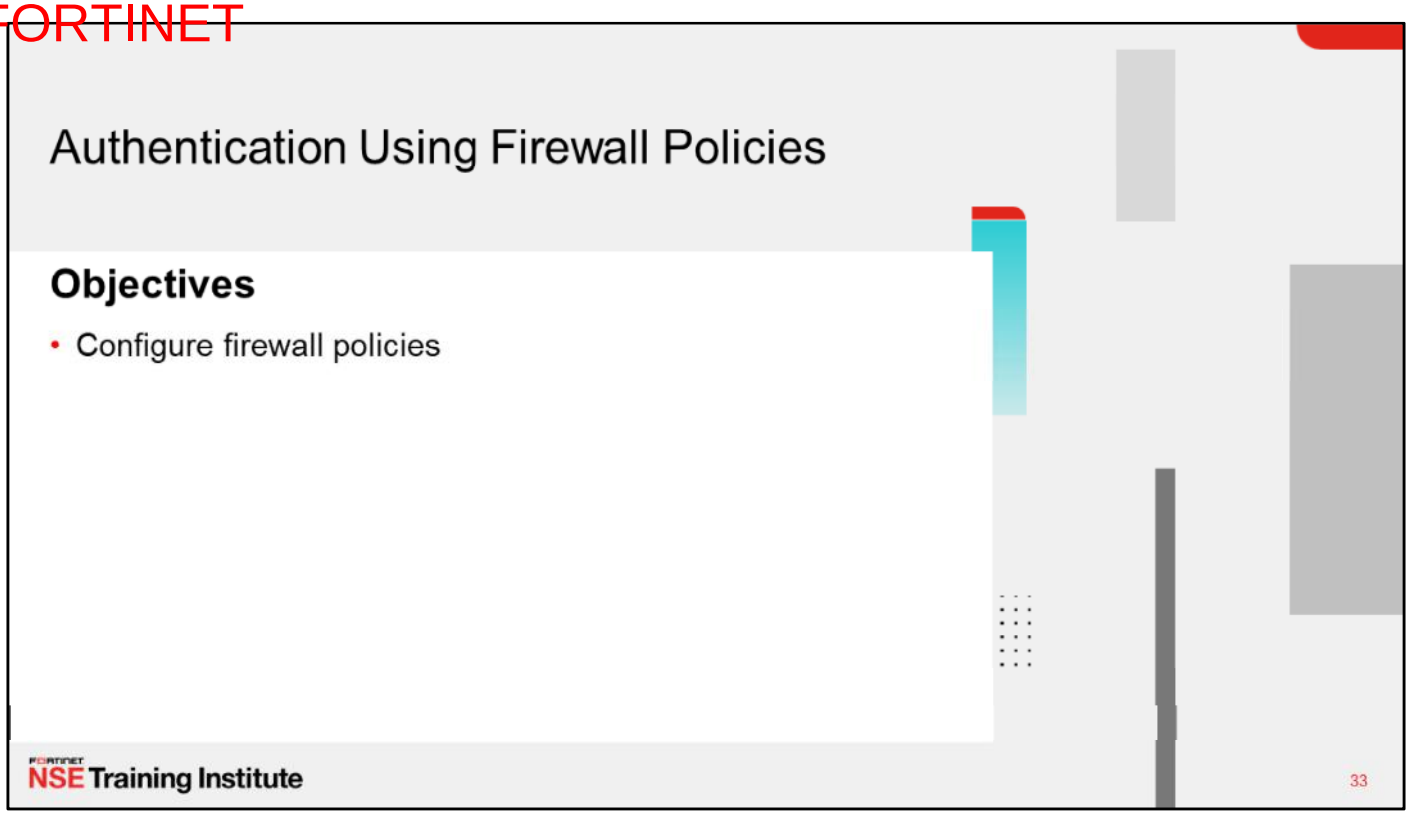

After completing this section, you should be able to achieve the objective shown on this slide.

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

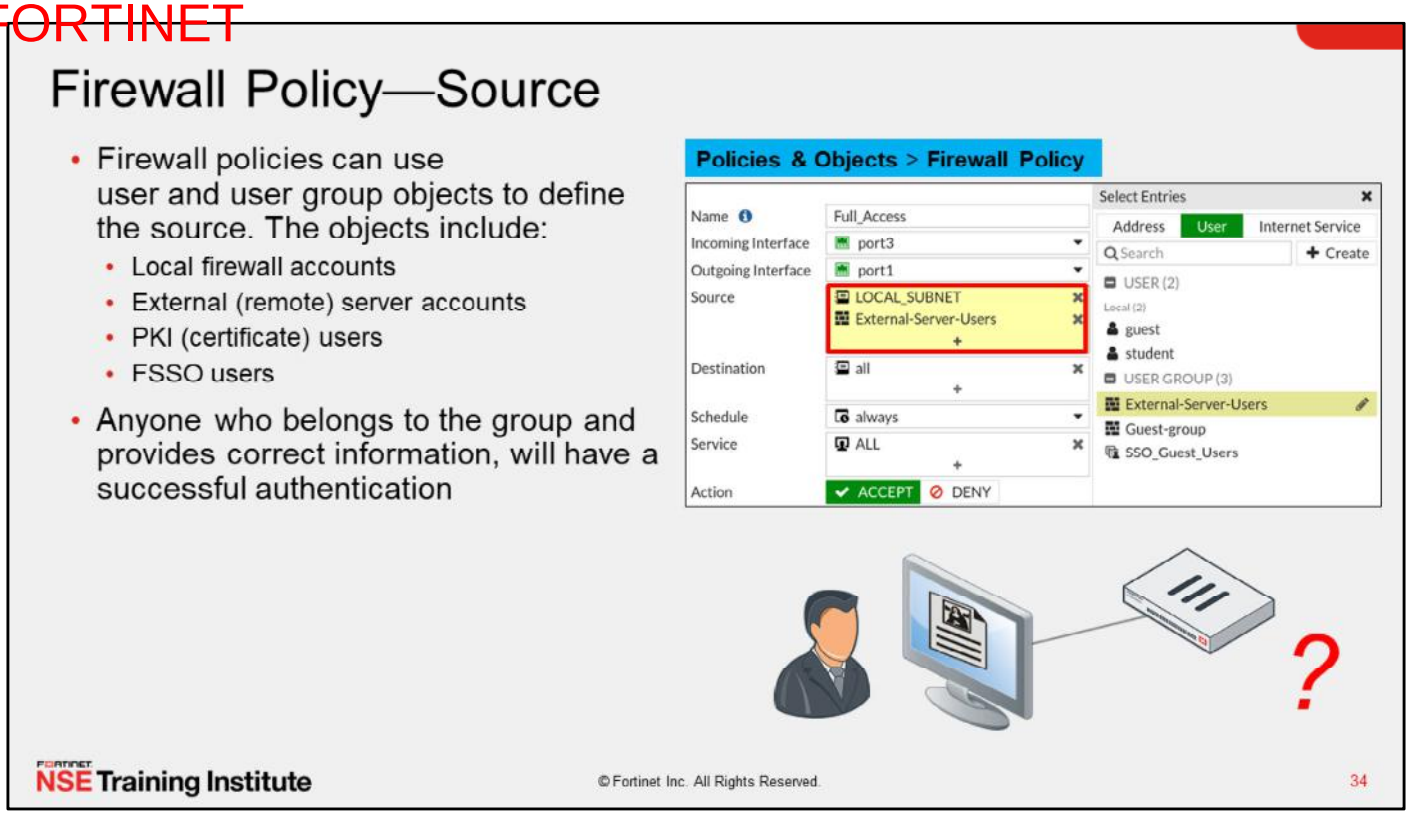

A firewall policy consists of access and inspection rules (compartmentalized sets of instructions) that tell FortiGate how to handle traffic on the interface whose traffic they filter. After the user makes an initial connection attempt, FortiGate checks the firewall policies to determine whether to accept 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.

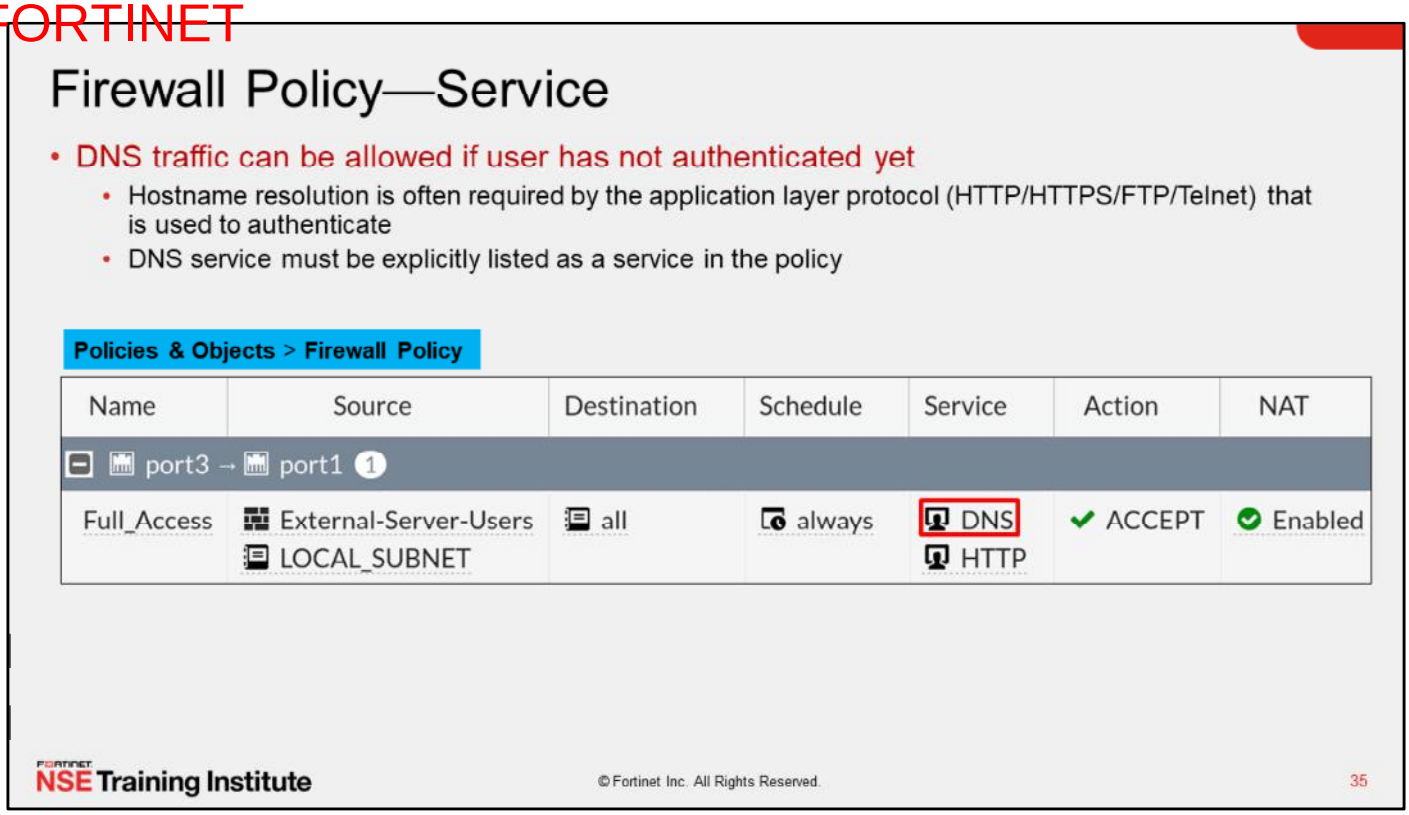

A firewall policy also checks the service in order to transport the named protocols or group of protocols. No service (with the exception of DNS) is allowed through the firewall policy before successful user authentication. DNS is usually used by HTTP so that people can use domain names for websites, 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).

Protocols

- A firewall policy must allow a protocol in order to show the authentication dialog that is used in active authentication:
  - HTTP
  - HTTPS
  - FTP
  - Telnet
- All other services are not allowed until the user has authenticated successfully through one of the protocols listed above

NSE Training Institute

36

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.

Mixing Policies

- Enabling authentication on a policy does not always force an active authentication prompt

port5 → port1

17	Guest	LOCAL_SUBNET Guest-group	all	AV Guest_AV SSL certificate-inspection	always	ALL	✔ ACCEPT	Enabled
18	Contractor	LOCAL_SUBNET Contractor	all	AV Contractor_AV SSL certificate-inspection	always	ALL	✔ ACCEPT	Enabled
19	Other	LOCAL_SUBNET	all	AV default SSL certificate-inspection	always	ALL	✔ ACCEPT	Enabled

- Three options:
  - Enable authentication on every policy that could match the traffic
  - Enforce authentication on demand option (CLI option only)
  - Enable a captive portal on the ingress interface for the traffic
- If login cannot be determined passively, then FortiGate uses active authentication
  - FortiGate does not prompt the user for login credentials when it can identify the user passively
  - By default, active authentication is intended to be used as a backup when passive authentication fails

NSE Training Institute © Fortinet Inc. All Rights Reserved. 37

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.

Active Authentication Behavior

- Enable authentication on every policy that could match the traffic:
  - All firewall policies must have authentication enabled (active or passive)
- If there is a fall-through policy in place, unauthenticated users are not prompted for authentication

- Enforce authentication on demand option:
  - CLI option only

```
# config user setting
(setting) # set auth-on-demand
<always|implicit>
Implicit - default option. It will not
trigger authentication if there is a fall
through policy.
Always - Trigger authentication prompt for
policies that have active authentication
enabled regardless of a fall through policy
```

- Provides more granular control
  - Authentication is enabled at a firewall policy level
- You must place passive authentication policies on top of active authentication policy

- Enable a captive portal on the ingress interface for the traffic:
  - Authentication happens at an interface level
  - Traffic is not allowed without valid authentication unless it matches an exemption
  - All users are prompted for authentication before they can access any resource

NSE Training Institute

38

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 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.

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.

NSE Training Institute

39

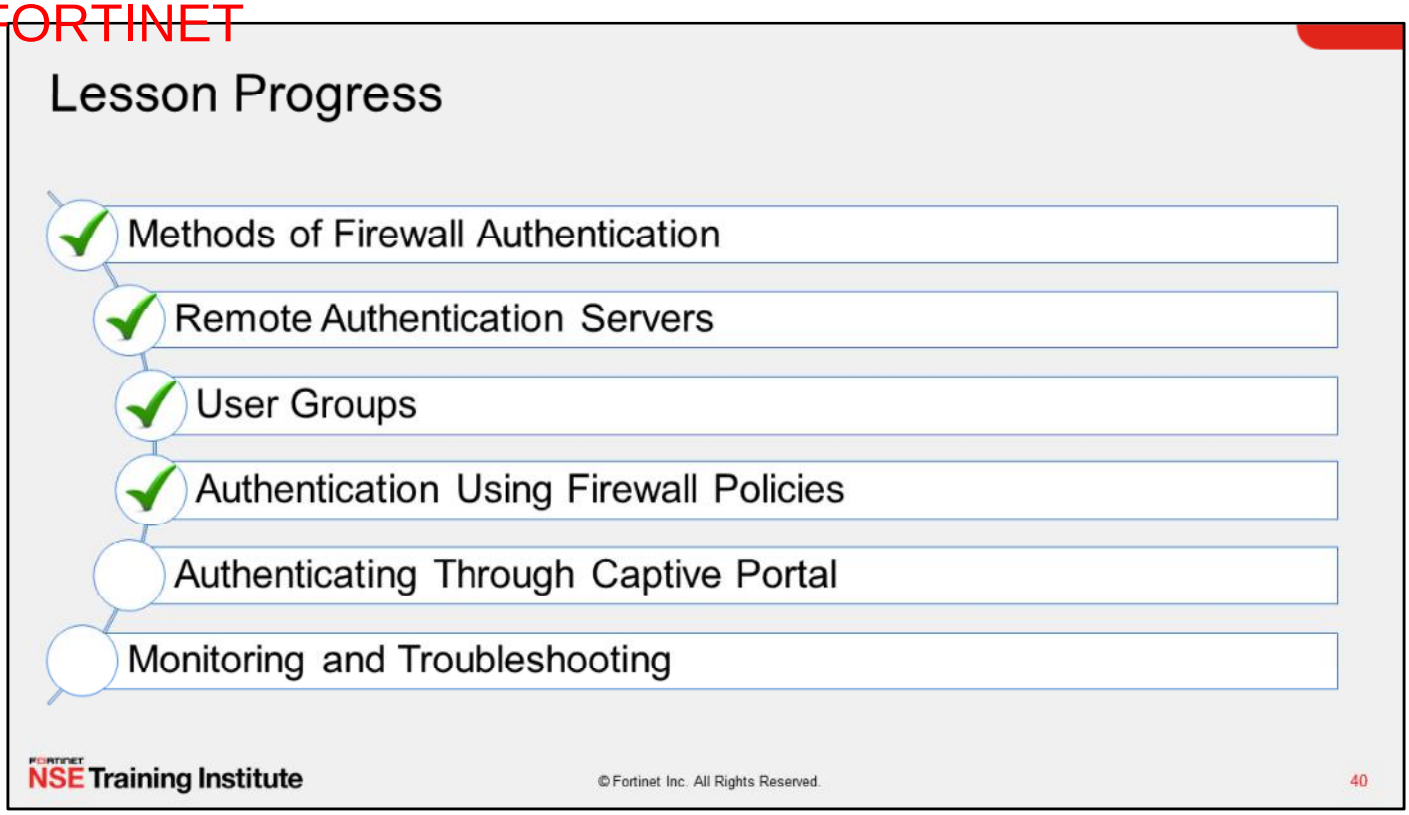

Good job! You now understand how to use firewall policies for authentication.

Now, you will learn about authenticating though captive portal.

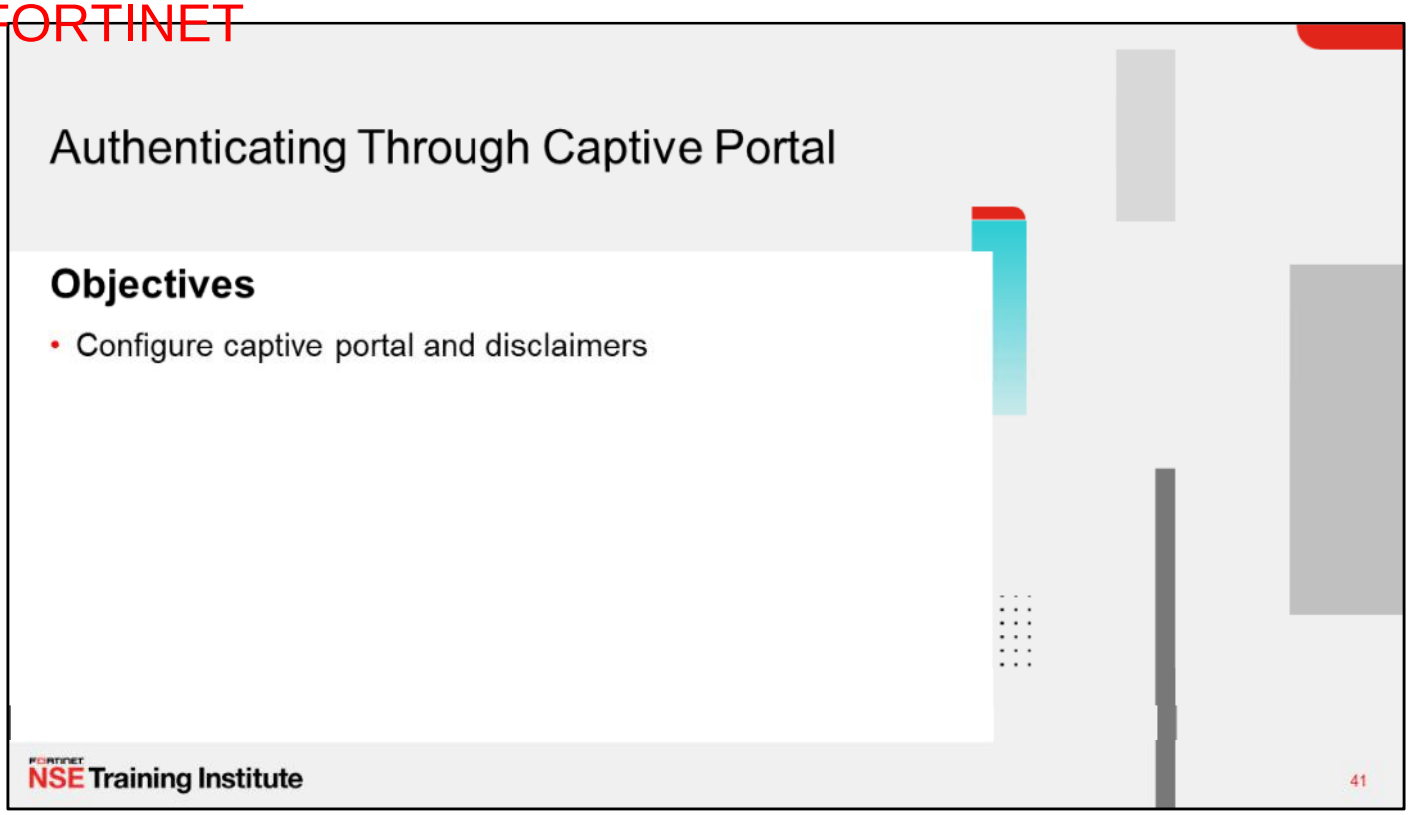

After completing this section, you should be able to achieve the objective shown on this slide.

By demonstrating competence in captive portal, you will be able to configure authentication through a captive portal.

Captive Portal

- Authenticates users on web pages that request a username and password
  - Enabled at interface level
- Only active authentication methods can use captive portal
- Can host captive portal on FortiGate or an external authentication server

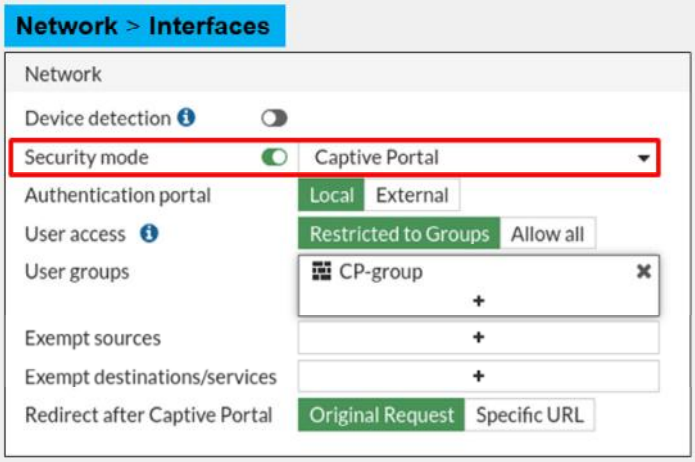

NSE Training Institute

42

If you want all users connecting to the network to be prompted for their login credentials (active authentication), you can enable captive portal. Captive portal is a convenient way to authenticate web users on wired or Wi-Fi networks through an HTML form that requests a username and password.

You can host a captive portal on FortiGate or an external authentication server, such as FortiAuthenticator.

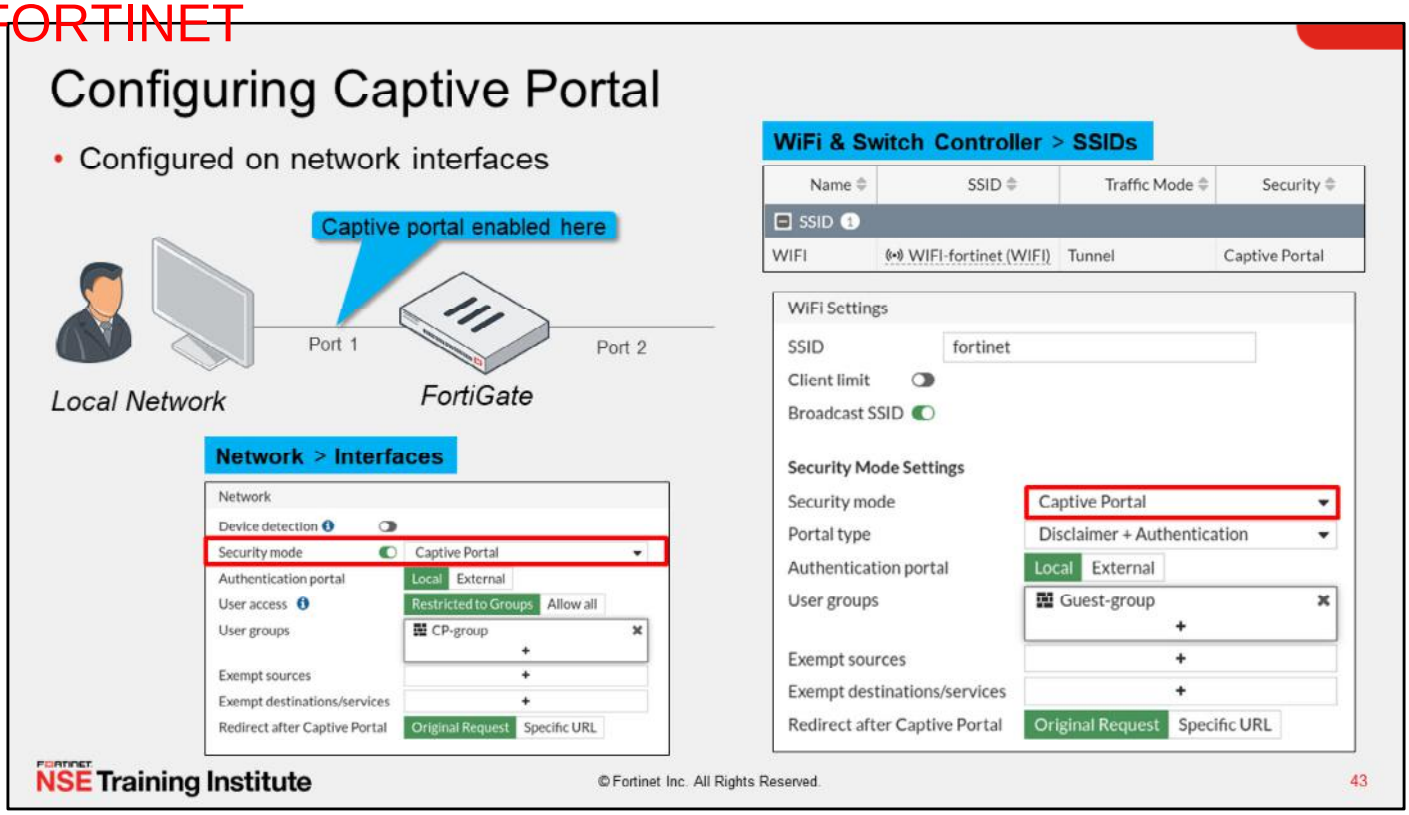

You enable captive portal, for both wired and Wi-Fi networks, at the interface level—regardless of the firewall policy that allows it or the port that it ultimately leaves by (authentication being enabled or disabled on the policy is not a factor). This is true for any network interface, including Wi-Fi and VLAN interfaces. On the local network, you must enable the captive portal setting on the incoming port.

You can configure a captive portal on the **Interfaces** page. Select the required interface. In the **Network** section, enable **Security mode**, and on the drop-down list, select **Captive Portal**. Note that if you are configuring captive portal for a Wi-Fi network, the Wi-Fi SSID must first exist.

Captive portals are not compatible with interfaces in DHCP mode.

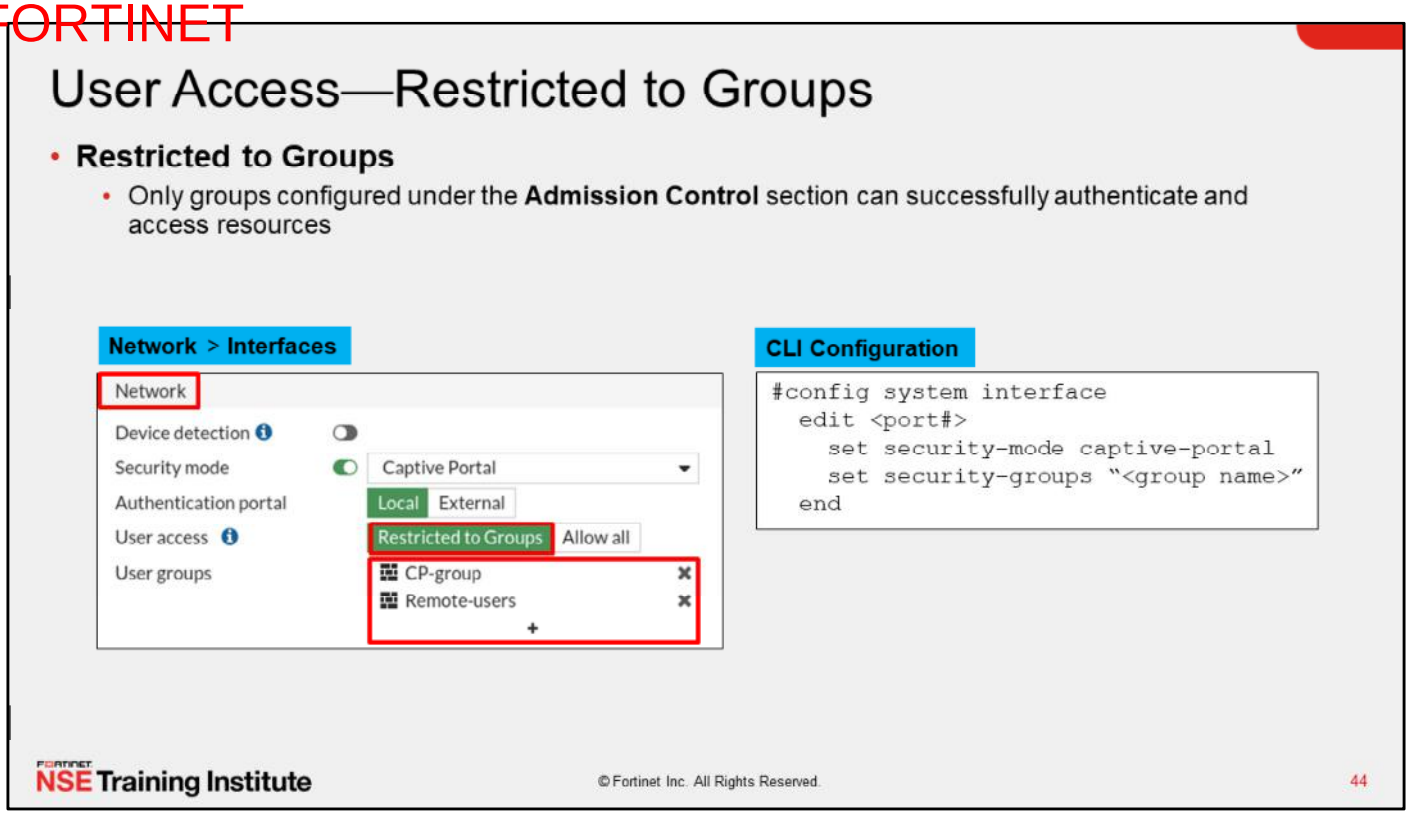

In the **Network** section, you also restrict captive portal user access.

Select **Restricted to Groups** to control the access from the captive portal configuration.

Use `the security-mode` and `security-groups` settings in port configurations to make the same changes on the CLI.

User Access—Allow All

- **Allow all:**
  - Any groups configured on the firewall policies can successfully authenticate and access resources

Network > Interfaces

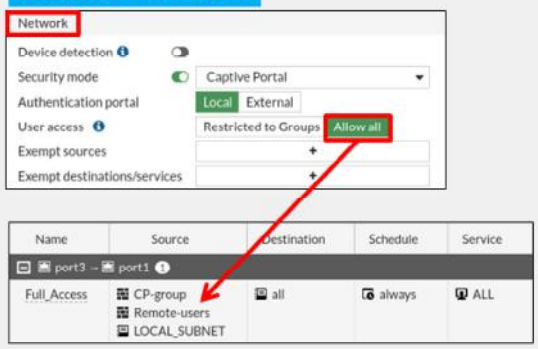

CLI Configuration

```
#config system interface
  edit <port#>
    set security-mode captive-portal
  end
```

NSE Training Institute

45

You can select **Allow all** to allow access to members of any groups configured on the firewall policies after authentication.

Use the `security-mode captive-portal` setting in port configurations to enforce captive portal access using the CLI. By omitting the security-groups setting the **User access** configuration is set to **Allow all**.

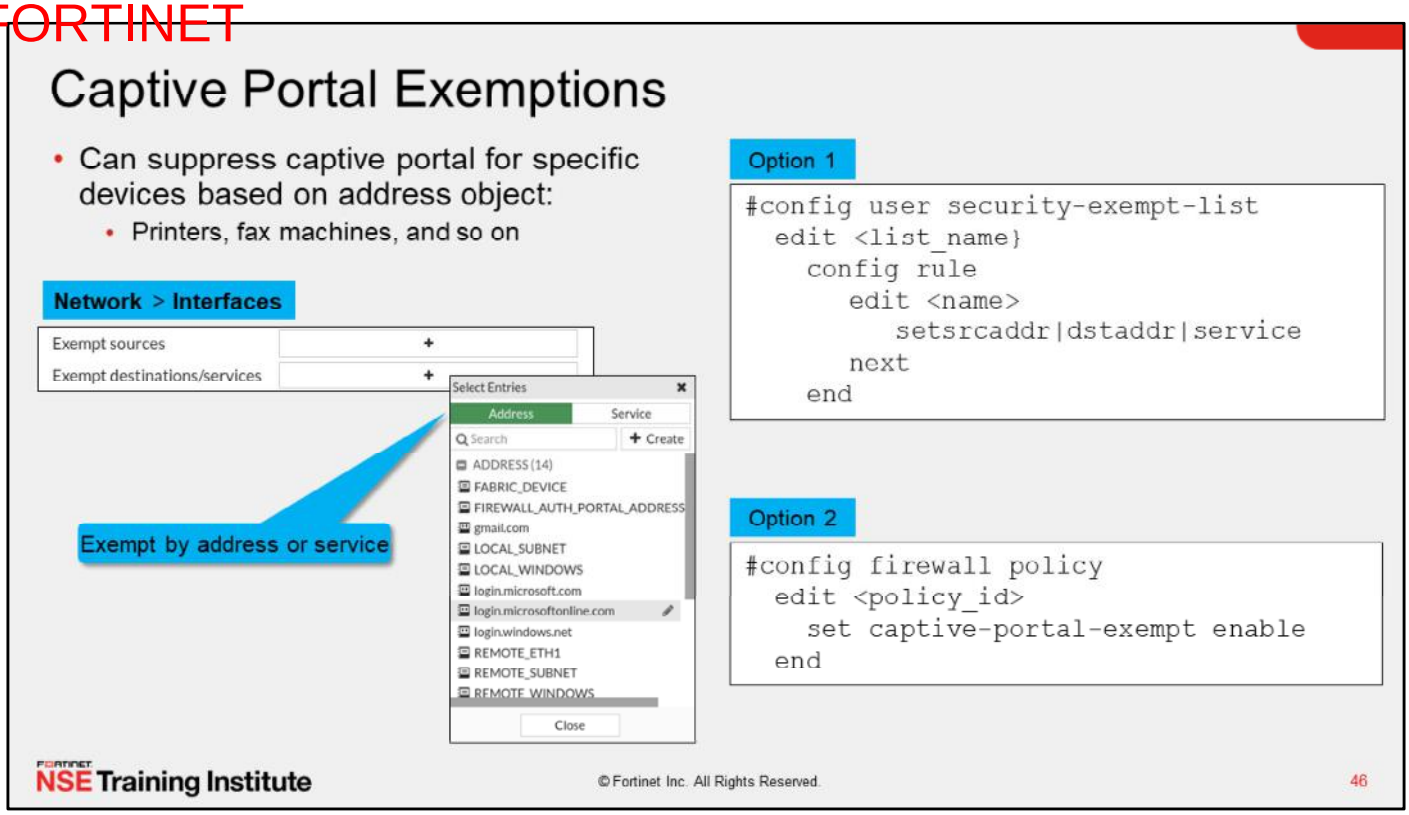

You can configure a firewall policy to suppress captive portal for specific addresses, or services. This is useful for devices that are unable to actively authenticate, such as printers and fax machines, but still need to be allowed by the firewall policy. When suppressed, traffic that matches the source or destination is not presented with the captive portal login page.

There are two ways you can bypass captive portal:

- Through a security exemption list in the GUI or the CLI under `config user security-exempt-list`
- Through the firewall policy. In the CLI, edit the policy and enter the command `set captive-portal-exempt enable`. All traffic matching this policy is now exempt from having to authenticate through the captive portal.

Terms of Service Disclaimer

- **Terms and Disclaimer Agreement** page displays before the user authenticates
  - The user must accept the disclaimer to proceed
  - After accepting, the user is directed to the intended destination

```
#config firewall policy
  edit <policy_id>
    set disclaimer enable
  end
```

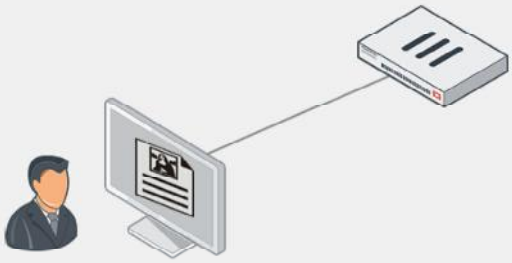

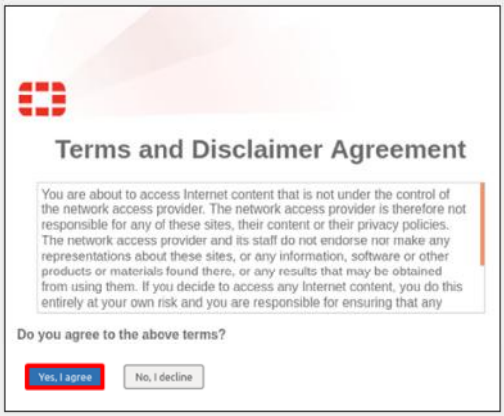

NSE Training Institute

47

If you want to enable a terms of service disclaimer to be used in combination with captive portal authentication, you can do so by using the `config firewall policy` and `set disclaimer enable` commands on the CLI. The terms of service disclaimer states the legal responsibilities of the user and the host organization. When you enable the disclaimer, the user must agree to the terms outlined in the statement in order to proceed to the requested URL. When enabled, the terms of service disclaimer opens immediately following a successful authentication.

Neither a security exemption list, nor a captive portal exemption on a firewall, can bypass a disclaimer.

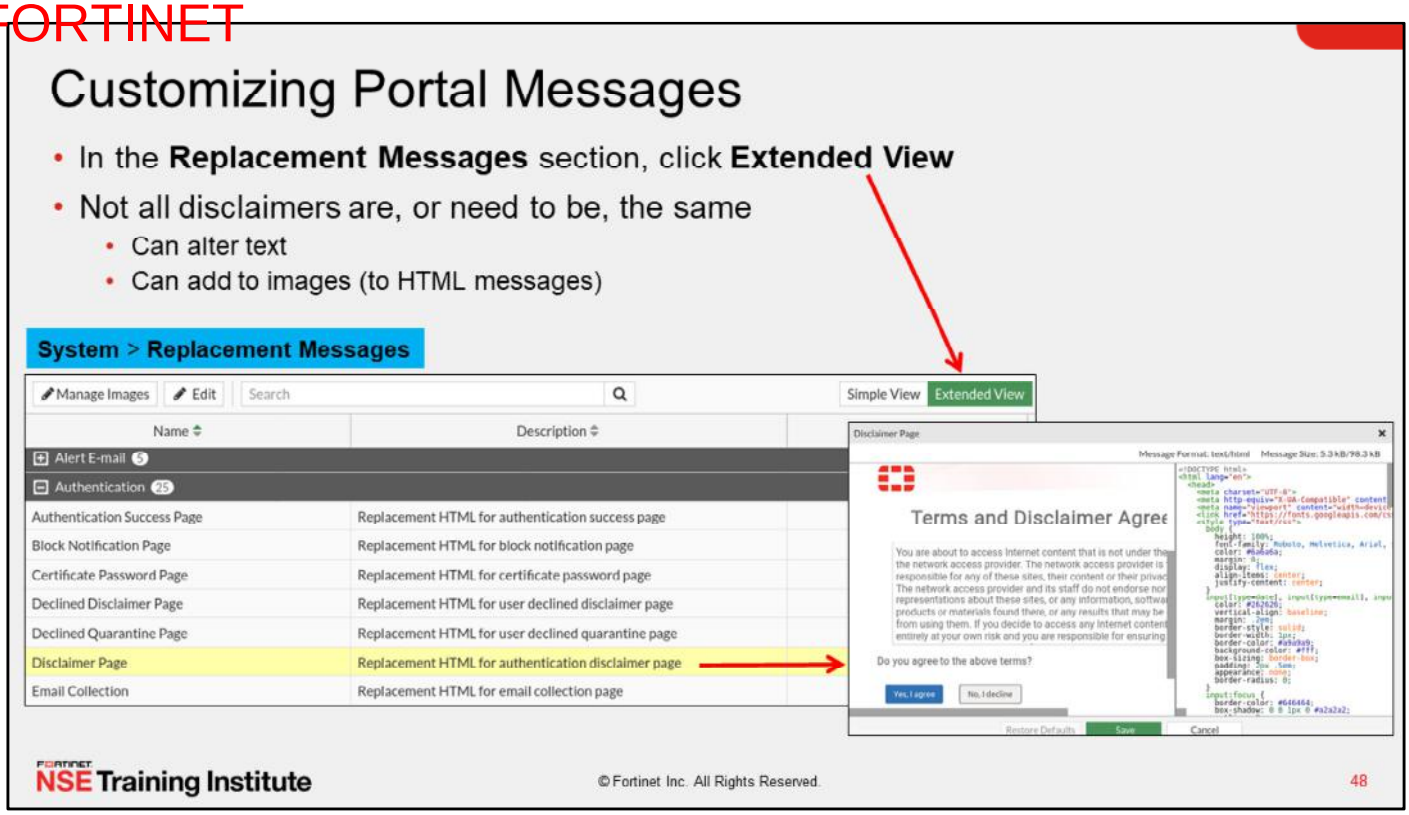

FortiGate allows you to customize portal messages, which include the login page and disclaimer page. You can customize the messages on the **Replacement Messages** page.

The disclaimer page is in HTML, so you must have knowledge of HTML in order to customize the message. The default layout is **Simple View**, which hides most of the replacement messages. Use **Extended View** to show all editable replacement messages.

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

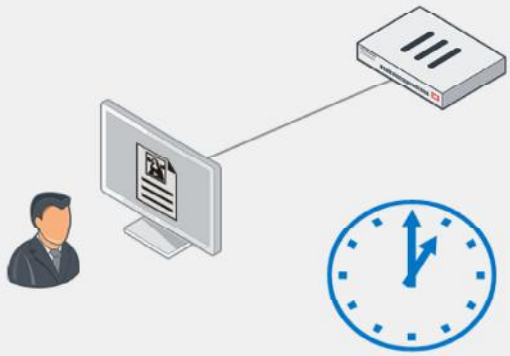

NSE Training Institute

49

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.

Knowledge Check

1. Which statement about captive portal is true?
   A. Captive portal must be hosted on a FortiGate device.
   ✓ B. Captive portal can exempt specific devices from authenticating.

2. 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 NSE Training Institute

50

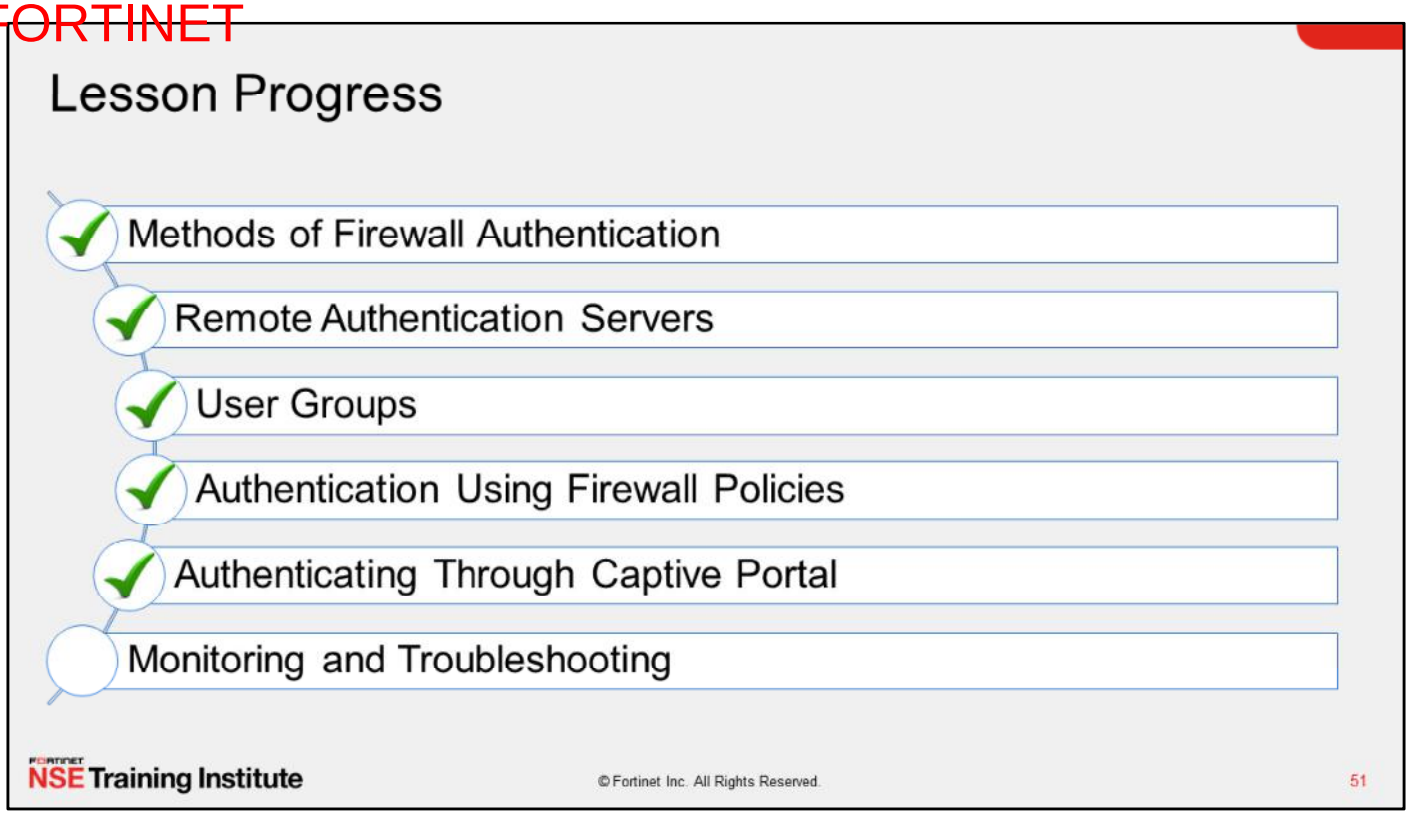

Good job! You now understand authenticating though captive portals.

Now, you will learn about monitoring and troubleshooting.

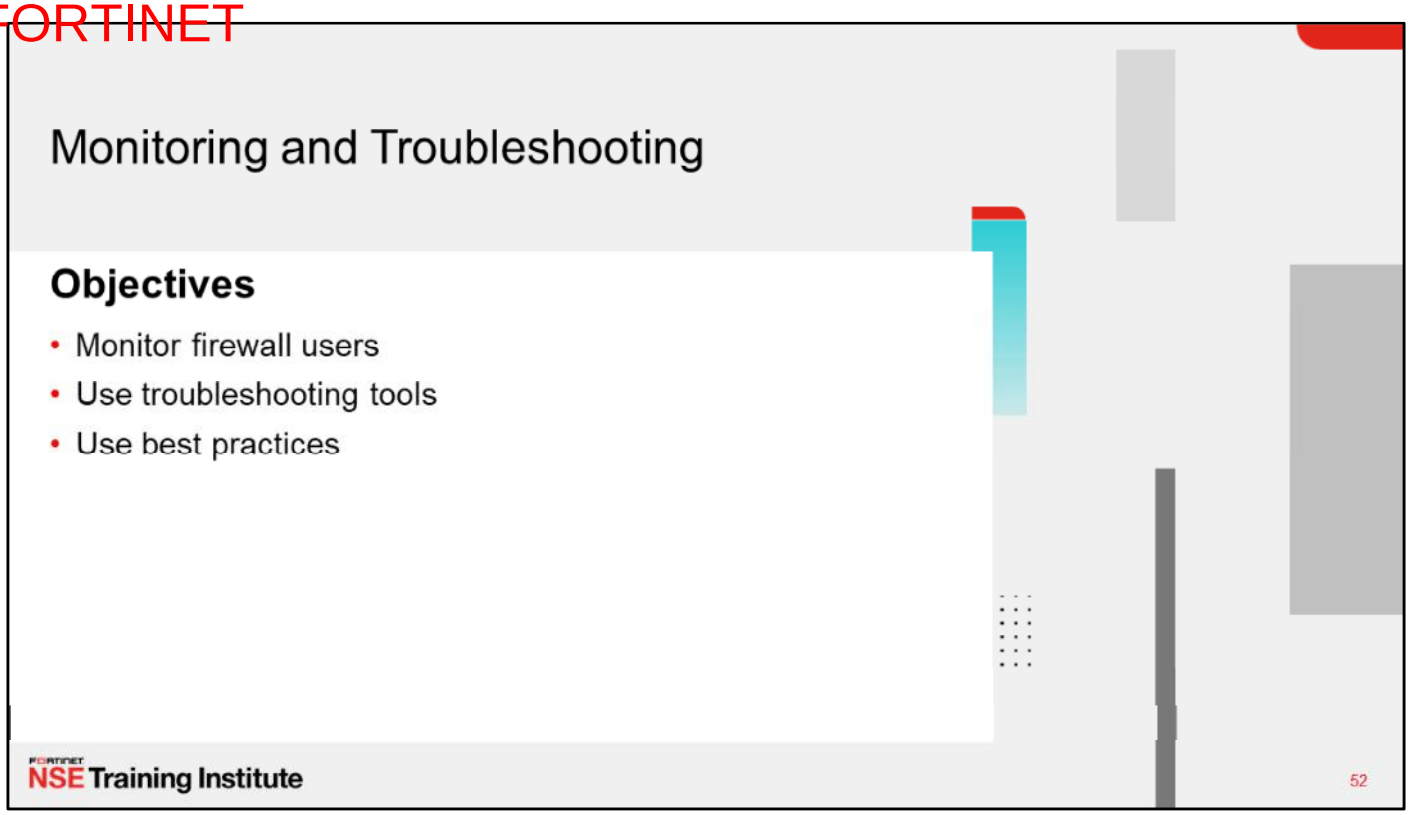

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in monitoring and troubleshooting, you will be able to monitor authenticated users and troubleshoot any issues that may occur.

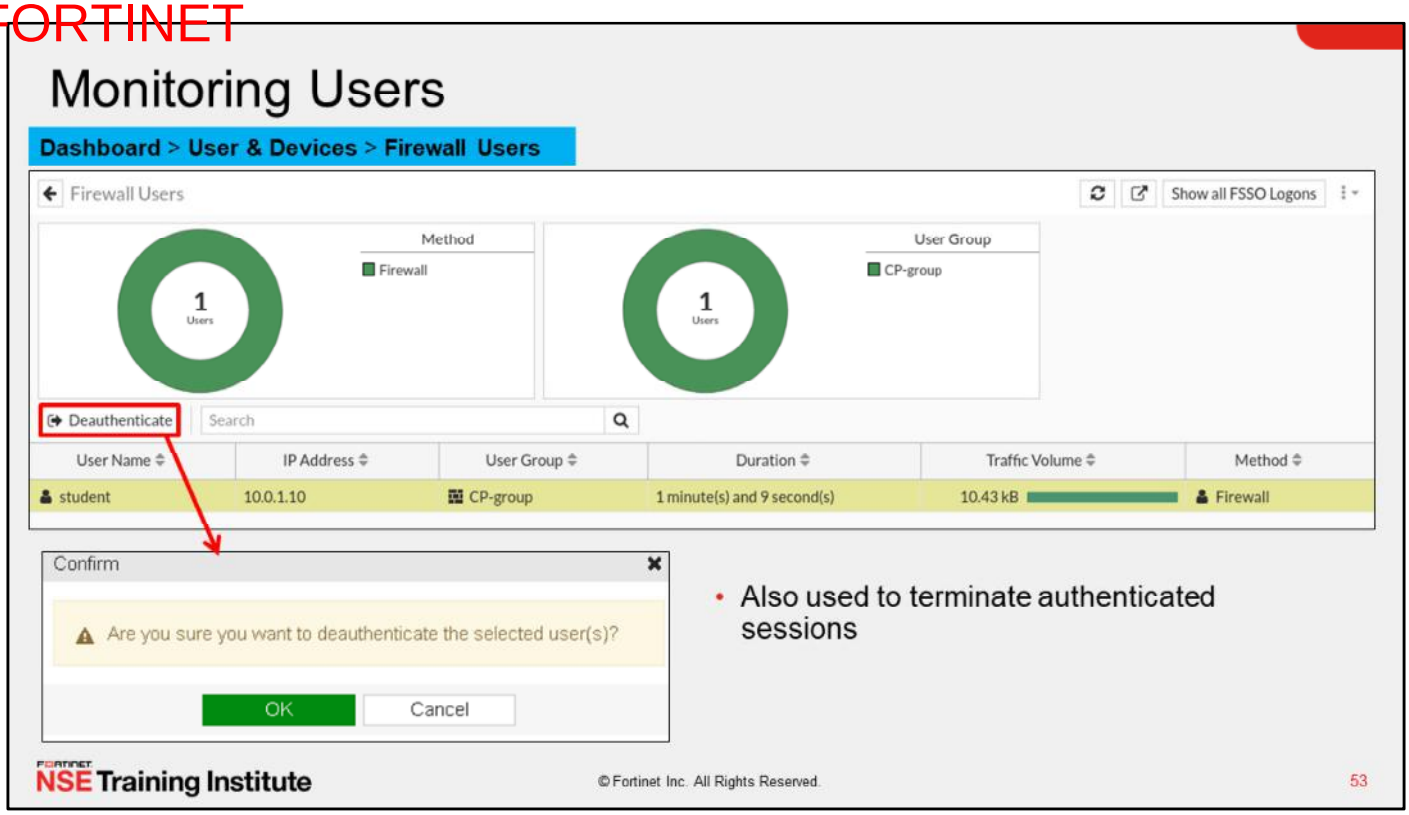

You can monitor users who authenticate through your firewall policies using the **Dashboard** > **User & Devices** > **Firewall Users** page. It displays the user, user group, duration, IP address, traffic volume, and authentication method.

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.

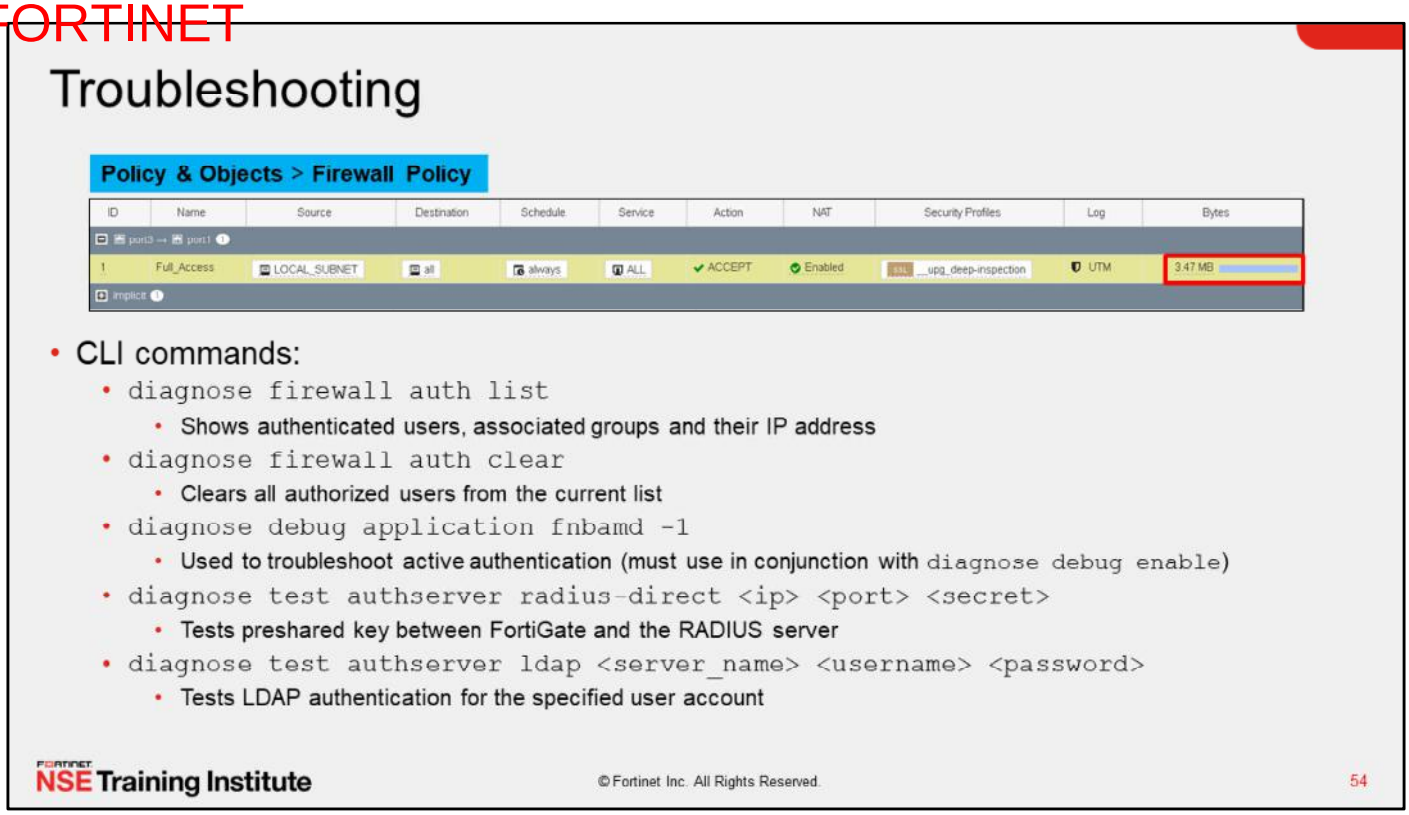

In the web-based manager, a good tool for troubleshooting is the **Bytes** column on the security policy page, which you open by clicking **Policy & Objects** > **Firewall Policy**. This column displays the number of bytes that have passed through this policy. This is valuable information to have when you are troubleshooting. When you are testing your configuration (end-to-end connectivity, user authentication, and policy use) watching the byte count for an increase can help with troubleshooting. An increase indicates if the policy in question is seeing any traffic, which is useful information if you expect a user to require authentication, but they are never prompted.

Use the following CLI commands to gather more information about users and user authentication attempts to help troubleshoot failed authentication attempts:

- `diagnose firewall auth list`: shows authenticated users and their IP address.
- `diagnose firewall auth clear`: clears all authorized users from the current list. This is useful when you need to force users to reauthenticate after system or group changes. However, this command can easily result in many users having to reauthenticate, so use it carefully.
- `diagnose debug application fnbamd -1`: use this command to troubleshoot active authentication, (You must use it in conjunction with `diagnose debug enable`.)
- `diagnose test authserver radius-direct <ip> <port> <secret>`: tests preshared key between FortiGate and the RADIUS server.
- `diagnose test authserver ldap <server_name> <username> <password>`: tests LDAP authentication for the specified user account.

Best Practices

- Set the source IP whenever the remote RADIUS server is accessed through a VPN, because most VPNs do not have an IP address associated with the VPN interface
- Servers should not go through an authentication policy. Use a dedicated, non-authentication policy for each server

Caution: Use extreme caution when selecting the **Include in every User Group** option when configuring a RADIUS server. This option places the RADIUS server, and all users who can authenticate against that server, into *every* FortiGate user group, including groups that are used for administrator access.

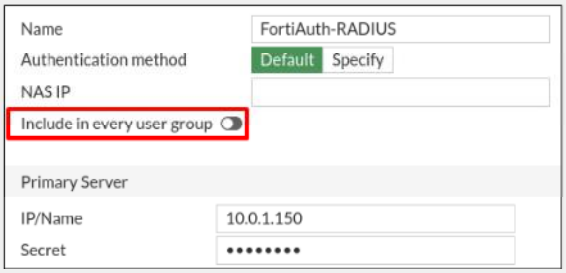

NSE Training Institute

55

Use the best practices listed on this slide to avoid unnecessary issues when configuring firewall authentication.

Knowledge Check

1. Which command would you use to identify the IP addresses of all authenticated users?
   A. `diagnose firewall auth clear`
   ✓ B. `diagnose firewall auth list`

NSE Training Institute

56

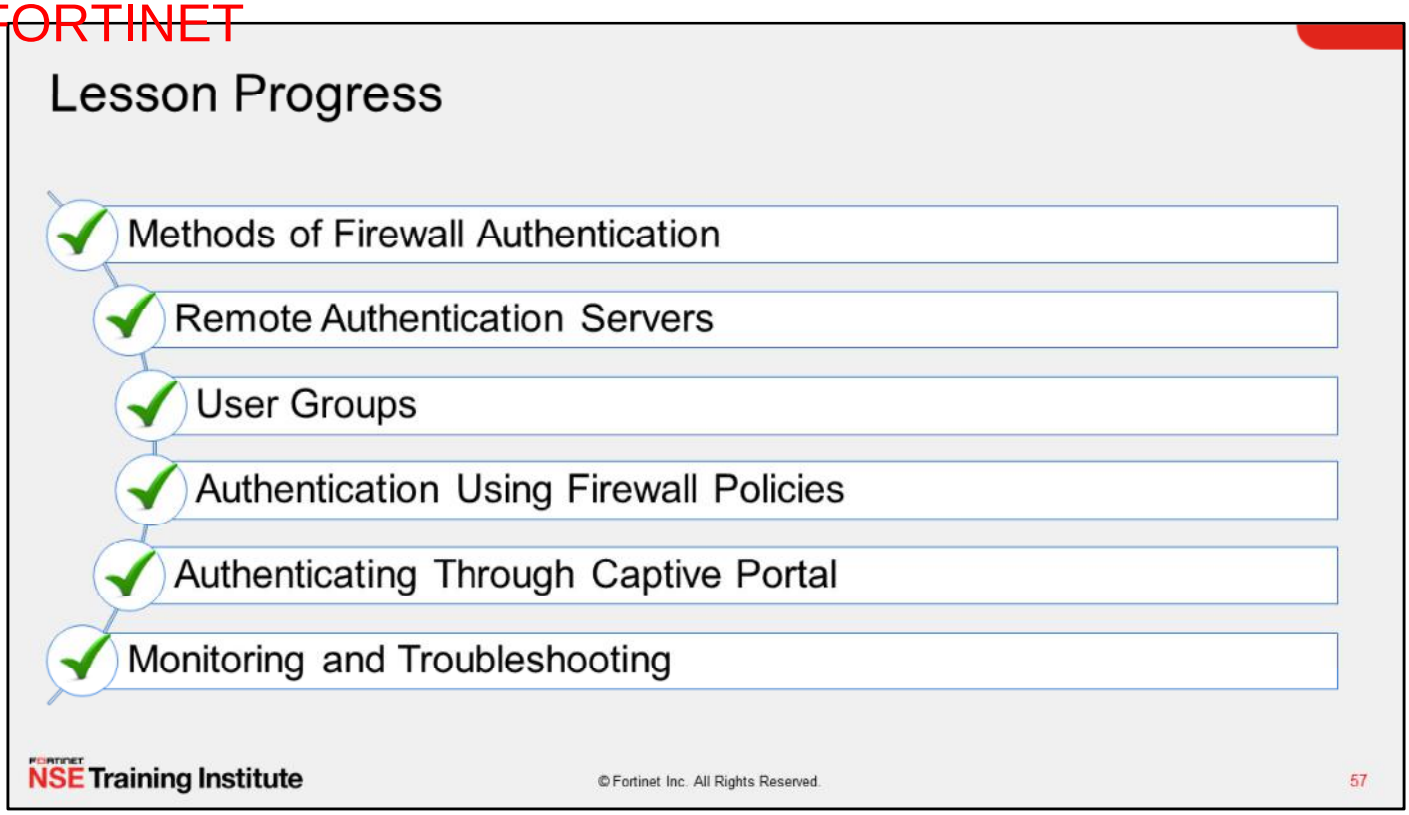

Congratulations! You have completed this lesson.

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

Review

- ✓ Describe firewall authentication
- ✓ Identify the different methods of firewall authentication available on FortiGate devices
- ✓ Identify supported remote authentication servers
- ✓ Describe active and passive authentication and the order of operations
- ✓ Configure users for local password authentication, server-based password authentication, and two-factor authentication
- ✓ Configure a remote authentication server
- ✓ Configure user authentication, firewall policies, captive portal, and disclaimers
- ✓ Monitor firewall users
- ✓ Use troubleshooting tools and best practices

NSE Training Institute

58

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.

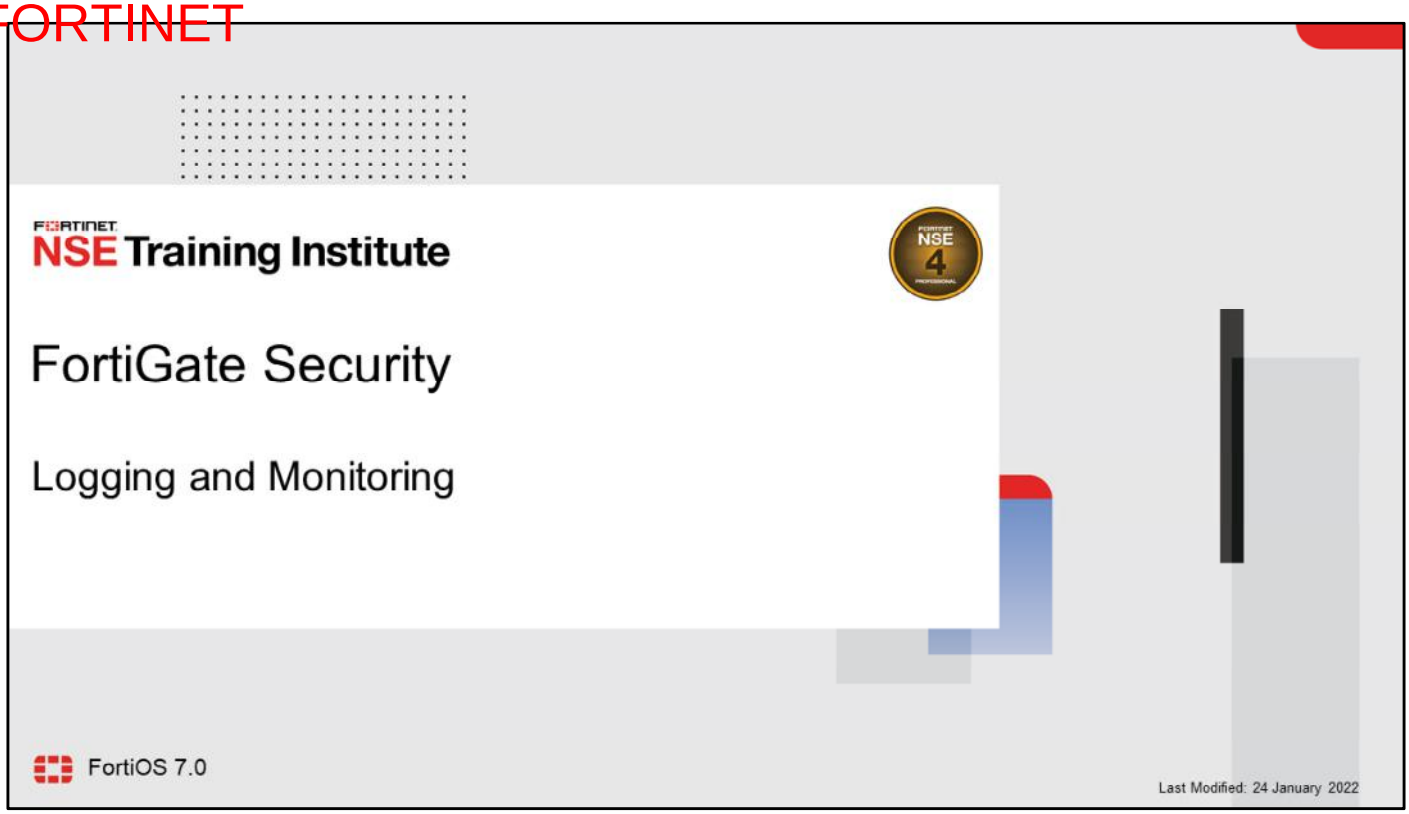

In this lesson, you will learn how to configure local and remote logging on FortiGate; view, search, and monitor logs; and protect your log data.

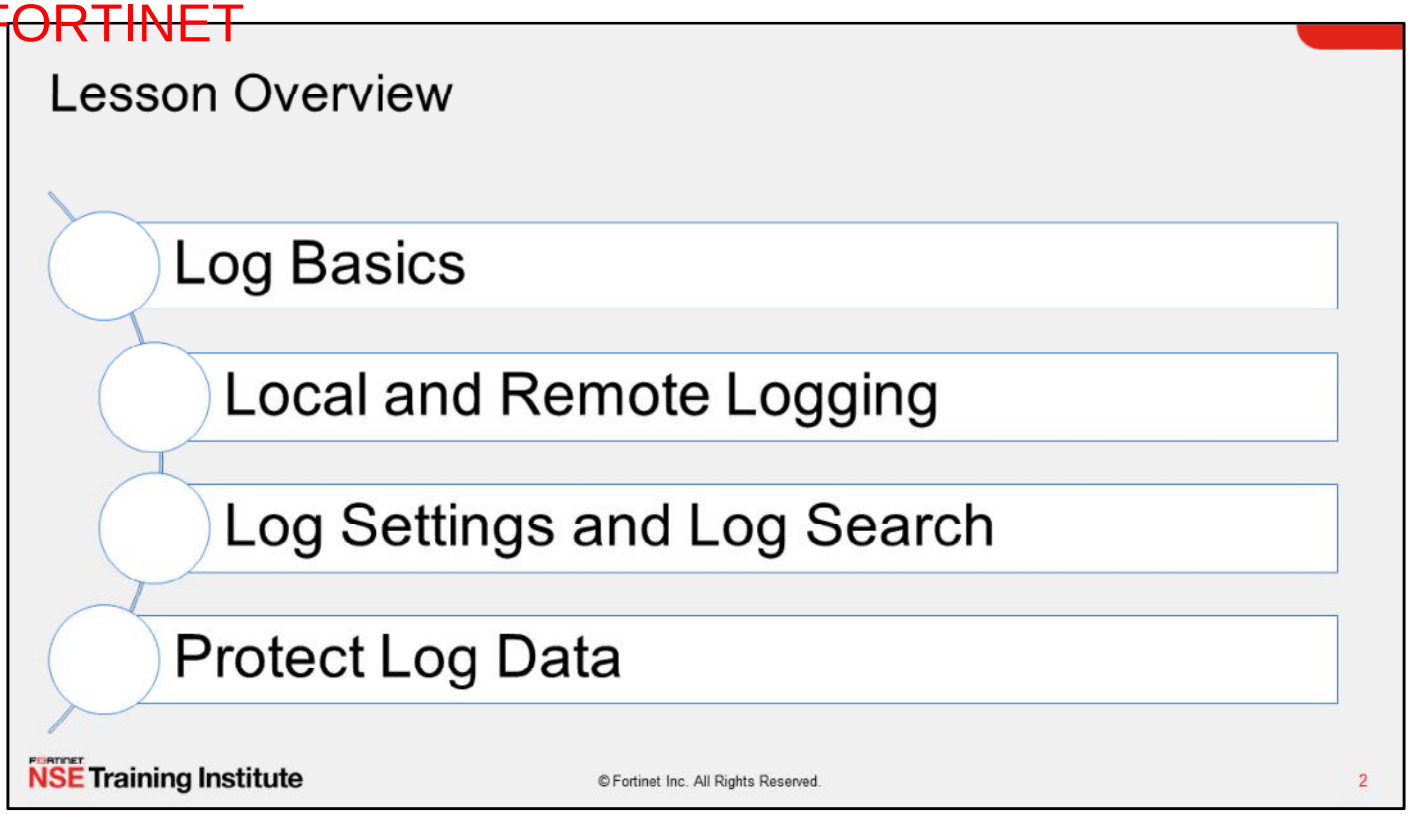

In this lesson, you will learn about the topics shown on this slide.

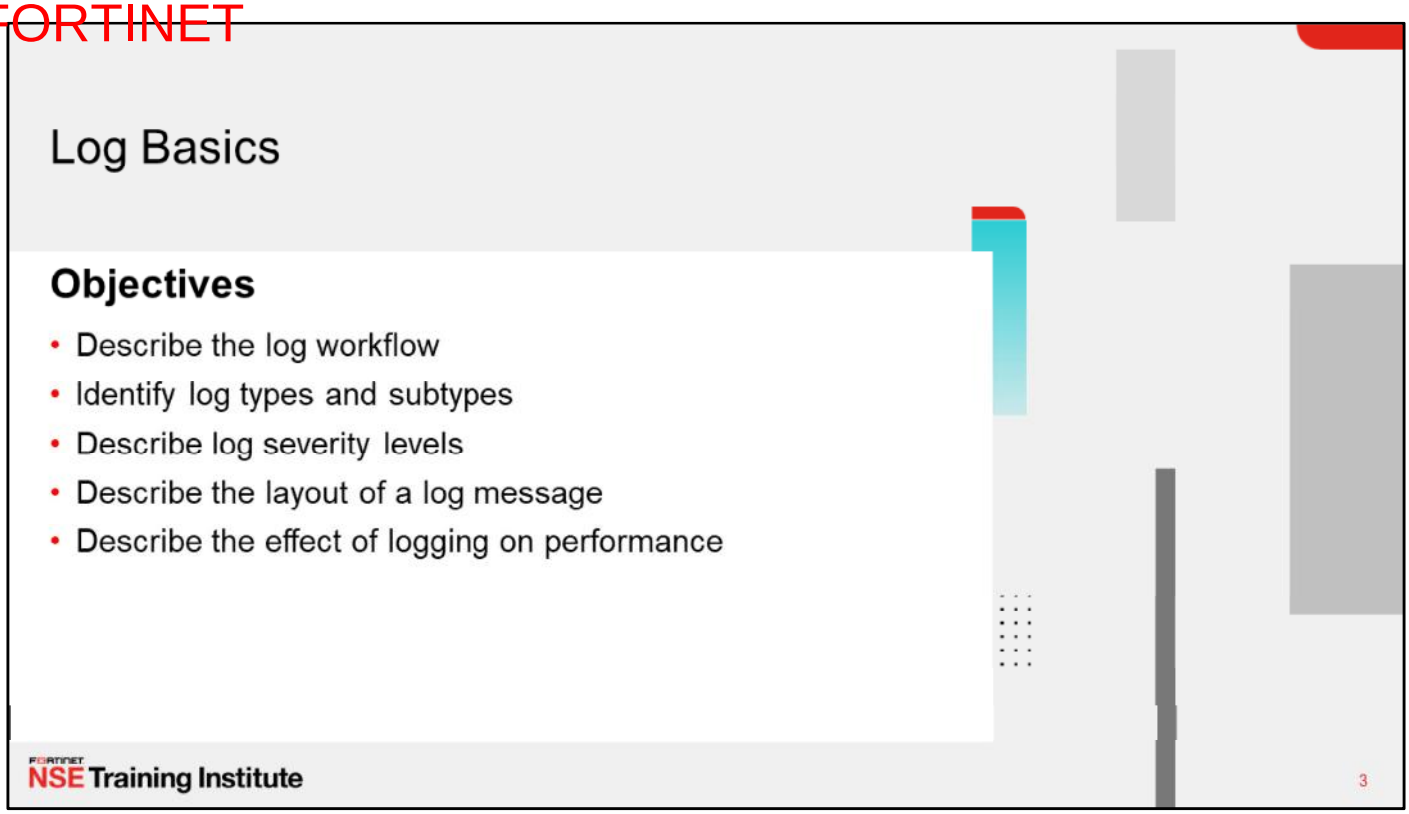

After completing this section, you should be able to achieve the objectives shown on this slide.

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

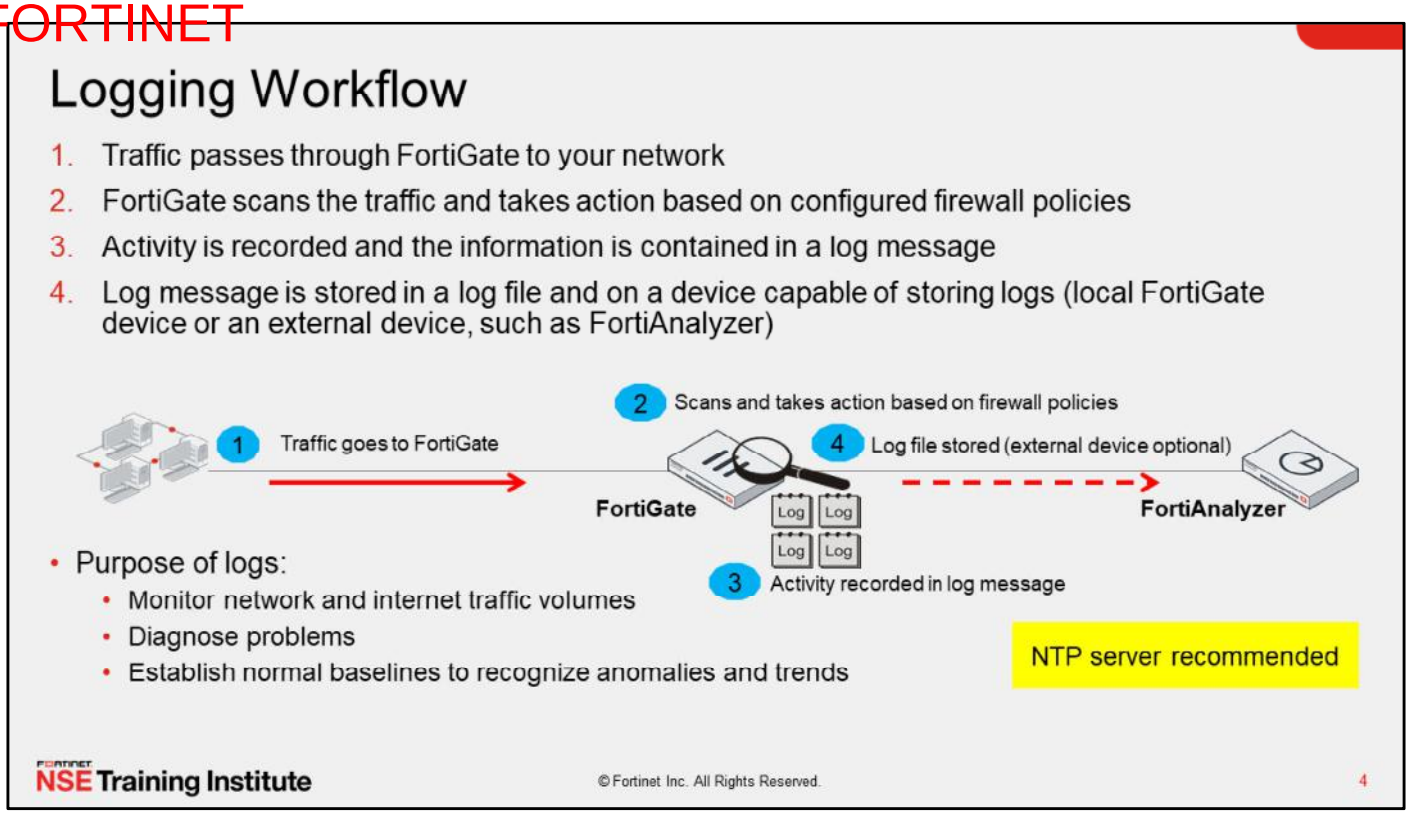

When traffic passes through FortiGate to your network, FortiGate scans the traffic, and then takes action based on the firewall policies in place. This activity is recorded, and the information is contained in a log message. The log message is stored in a log file. The log file is then stored on a device 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 make adjustments to your network security, 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)
  - If no security logs exist, the menu item does not appear in the GUI.

Traffic	Event	Security
Forward	Endpoint Control	Application Control
Local	High Availability	Antivirus
Sniffer	System	Data Leak Prevention (DLP)
	User	Anti-Spam
	Router	Web Filter
	VPN	Intrusion Prevention System (IPS)
	WAD	Anomaly (DoS-policy)
	Wireless	Web Application Firewall (WAF)

WAN optimization logs are found within traffic logs

GPRS Tunneling Protocol (GTP) logs are handled separately from default event logs

NSE 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 subtypes named endpoint control, high availability, system, user, router, VPN, WAD, and wireless.

- System event logs contain information related to operations, such as automatic FortiGuard updates and GUI logins.
- User logs contain logon and logoff events for firewall policies with user authentication.
- Router, VPN, WAD, 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 application control, antivirus, DLP, anti-spam (email filter), web filter, intrusion protection, anomaly (DoS-policy), and WAF. Security logs and subtypes are only visible in the GUI if logs are created within it—if no security logs exists, the menu item does not appear.

Log Severity Levels

- Each log entry includes a log level (also known as priority level) that ranges in order of importance
  - 0 = high importance / 6 = low importance

Levels	Description
0 – Emergency	System unstable
1 – Alert	Immediate action required
2 – Critical	Functionality effected
3 – Error	Error exists that can affect functionality
4 – Warning	Functionality could be affected
5 – Notification	Information about normal events
6 – Information	General system information
7 – Debug	Diagnostic information for investigating issues

Rarely used, unless actively investigating an issue with Fortinet Support

NSE Training Institute

6

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. It puts 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.

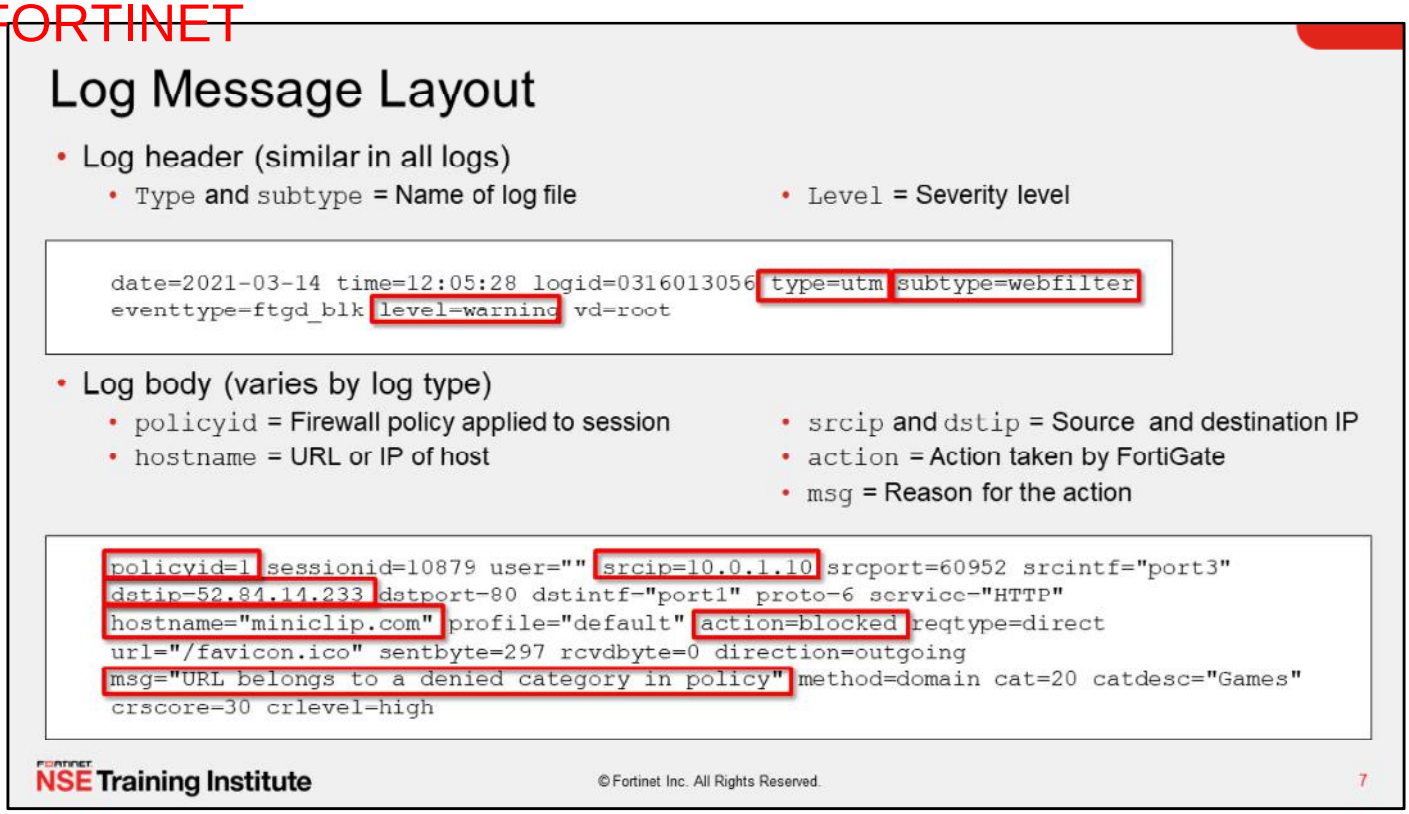

Every log message has a standard layout comprising two sections: a header and a body.

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 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.

Logging in a Security Fabric Design

- Requisite products: Two or more FortiGate devices and a FortiAnalyzer (a remote logging device)
- With FortiGate, you can enable different security features in different firewalls in the fabric
  - Ensures you do not have to scan and log the same traffic flow more than once when it passes more than one firewall
- FortiGate can share network-related information
  - Devices connected to downstream FortiGate devices will be visible on the upstream device as well (you must enable device detection on the **Interfaces** page of the FortiGate GUI)
- Administrators can view logs and devices connected to the network by logging in to the root FortiGate in the Security Fabric
  - Information is securely shared using the FortiTelemetry protocol

NSE Training Institute © Fortinet Inc. All Rights Reserved. 8

Collecting logs from the devices in your Security Fabric is important. This is why two or more FortiGate devices and a FortiAnalyzer—a remote logging device—are requisite products at the core of the Security Fabric solution. With FortiGate, you can enable different security features, like antivirus, web filtering, intrusion prevention (IPS), and application control, in different firewalls in the fabric. For example, in the Internal Segmentation firewall (ISFW), you can enable only antivirus, while in the Next Generation firewall (NGFW) facing the internet, you can enable web filtering, IPS, and application control. This means you do not have to duplicate scans and logs of the same traffic flow when it passes through multiple firewalls.

The Security Fabric can provide a network topology view (physical and logical), and FortiGate devices can share network-related information. For example, devices connected to downstream FortiGate devices will be visible on the upstream device as well (you must enable device detection on the **Interfaces** page of the FortiGate GUI). In short, administrators can view logs and devices connected to the network by logging on to the root FortiGate in the Security Fabric. This information is securely shared using the FortiTelemetry protocol.

Effect of Logging on Performance

- More logs = more CPU, memory, and disk space
- Depending on the amount of traffic you have, and the logging settings that are enabled, your traffic logs can swell and impact the performance of your firewall
- Traffic logs record every session
  - Extra information for troubleshooting
  - Some UTM events
  - More system intensive

Enable performance statistic logging for remote logging devices on FortiGate

```
# config system global
   set sys-perf-log-interval <number from 0-15>
 end
```

NSE Training Institute

9

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 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 logging settings that are enabled, your traffic logs can swell and, ultimately, impact the performance of your firewall.

From the FortiGate CLI, you can enable performance statistic logging for remote logging devices, such as FortiAnalyzer and syslog, 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, antivirus, and DLP?
   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

NSE Training Institute

10

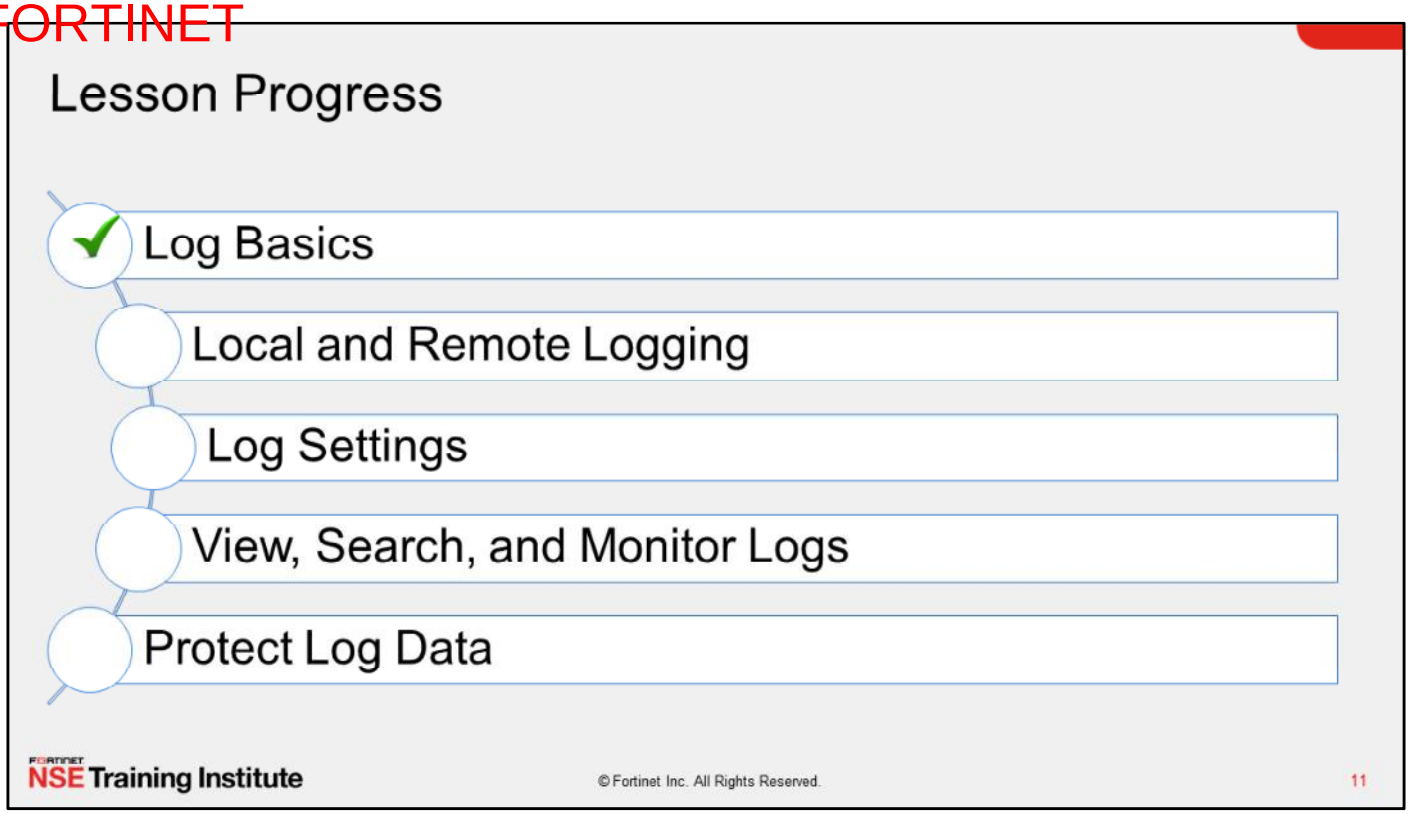

Good job! You now understand log basics.

Now, you will learn about local logging.

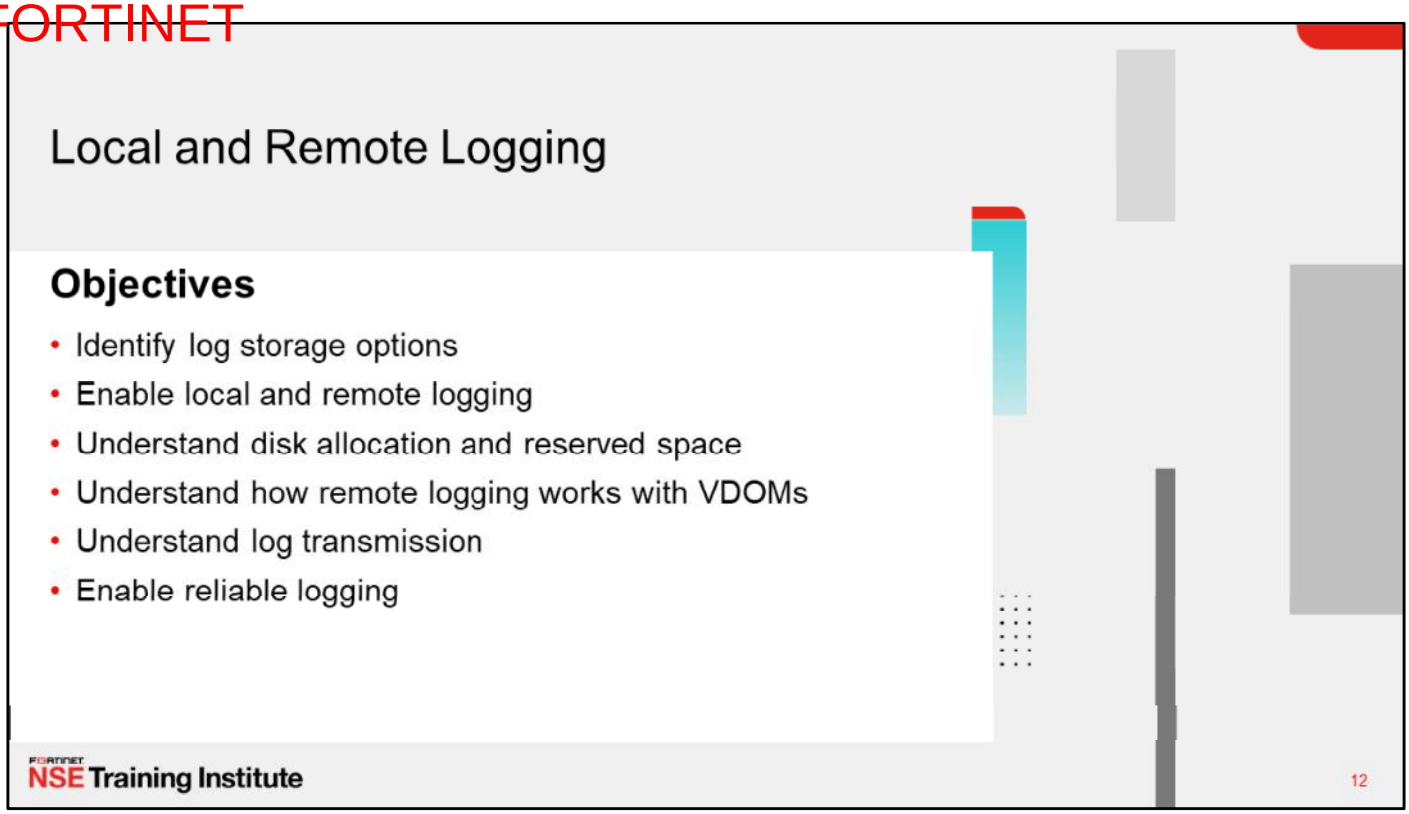

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

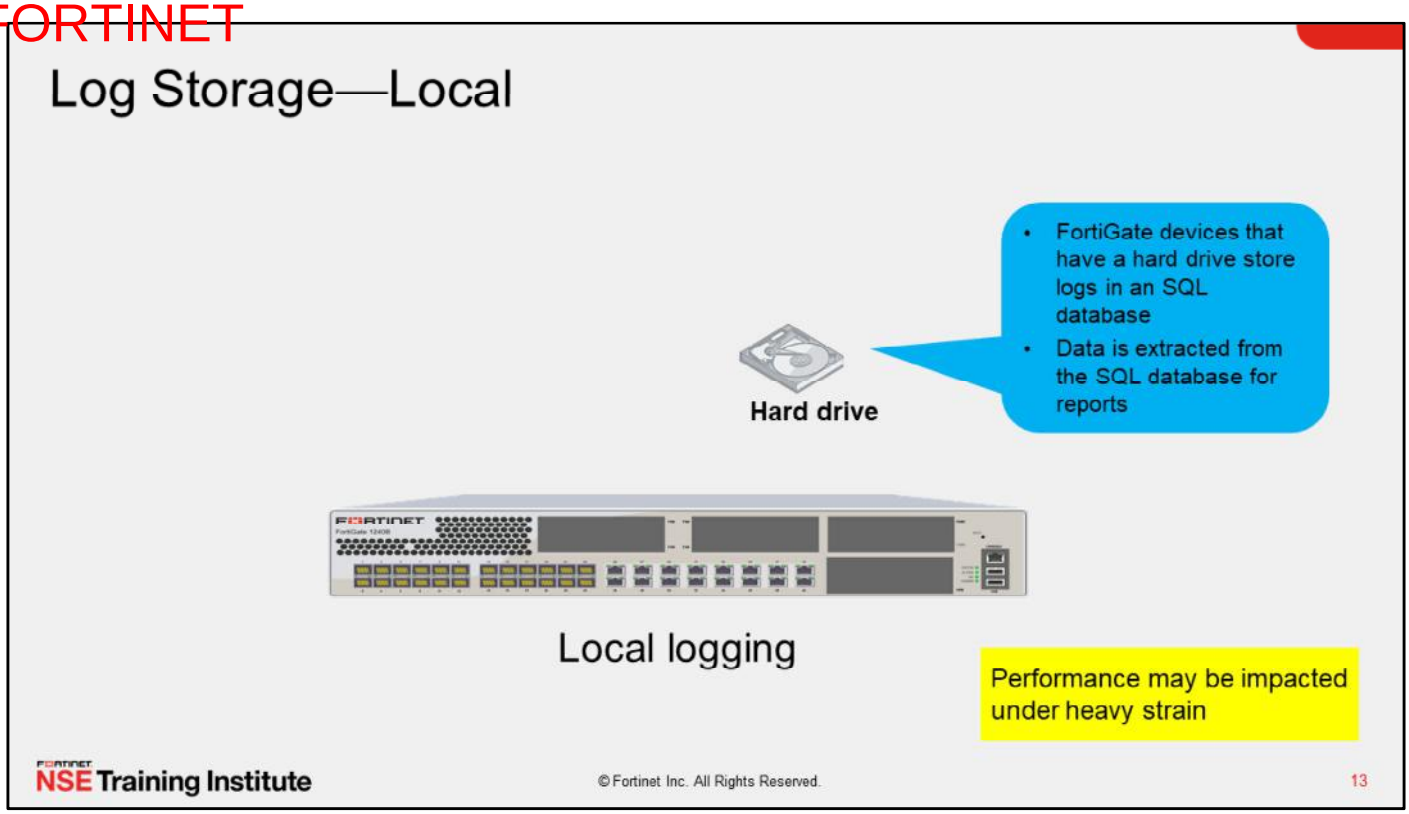

Storing logs on FortiGate is known as local logging. You can store logs to the device's hard drive.

Typically, mid to high-end FortiGates have a hard drive. Logging to a hard drive is known as 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, any logging to FortiGate—disk—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: you can use one with WAN optimization and the other for logging). If you are using the local hard disk for WAN optimization, you can log to remote FortiAnalyzer devices or syslog servers.

Enabling Local Logging

- To store logs locally on FortiGate, you must enable disk logging
- With disk logging enabled, the report daemon collects statistics used for historical FortiView from disk
  - If disk logging is disabled, FortiView logs are only available in real time
- By default, logs older than seven days are deleted from disk (configurable)

```
# config log disk setting
    set maximum-log-age <integer>
```

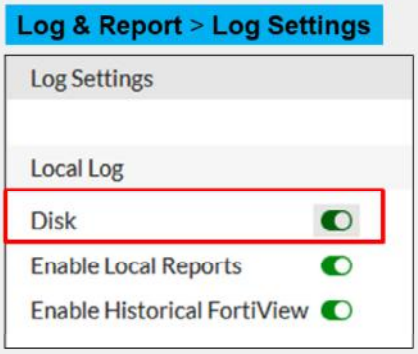


```
# config log disk setting
    set status enable
```

NSE Training Institute

14

If you want to store logs locally on FortiGate, you must enable disk logging from the **Log Settings** page. Only certain FortiGate models support disk logging. If your FortiGate does not support disk logging, you can log to an external device instead. You will learn about remote logging later in this lesson.

Disk logging must be enabled 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 (7) days are deleted from the disk (log age is configurable).

FortiGate Disk Allocation—Reserved Space

- The system reserves approximately 25% of its disk space for system usage and unexpected quota overflow
  - Only ~75% of disk space is available to store logs

```
FGT_A (global) # diagnose sys logdisk usage
Total HD usage: 208MB/118145MB
Total HD logging space: 88608MB
HD logging space usage for vdom "root": 0MB/9965MB
HD logging space usage for vdom "vdom1:" 0MB/104857MB
```

Use this command to obtain the amount of reserved space on your FortiGate

- Formulas:
  - disk - logging = reserved (i.e. 118145MB – 88608MB = 29537MB reserved)
  - reserved/disk*100 = reserved % (i.e. 29537/118145*100 = 25%)

NSE Training Institute

15

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.

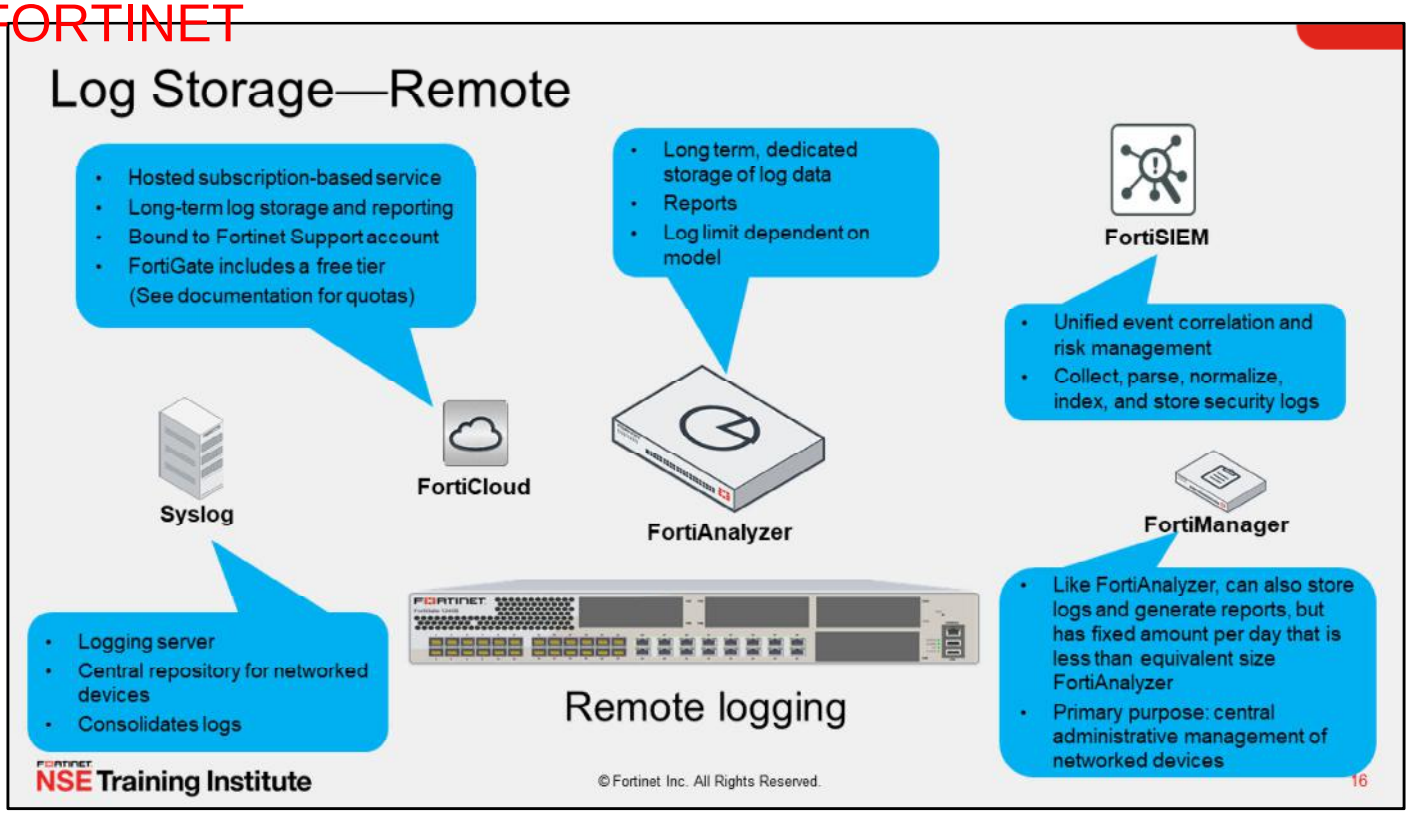

If storing logs locally does not fit your requirements, you can store logs externally. You can configure FortiGate to store logs on syslog servers, FortiCloud, FortiSIEM, FortiAnalyzer, or FortiManager. These logging devices can also be used as a backup solution.

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 disk or logging is not required for you to configure logging to FortiAnalyzer or FortiManager.

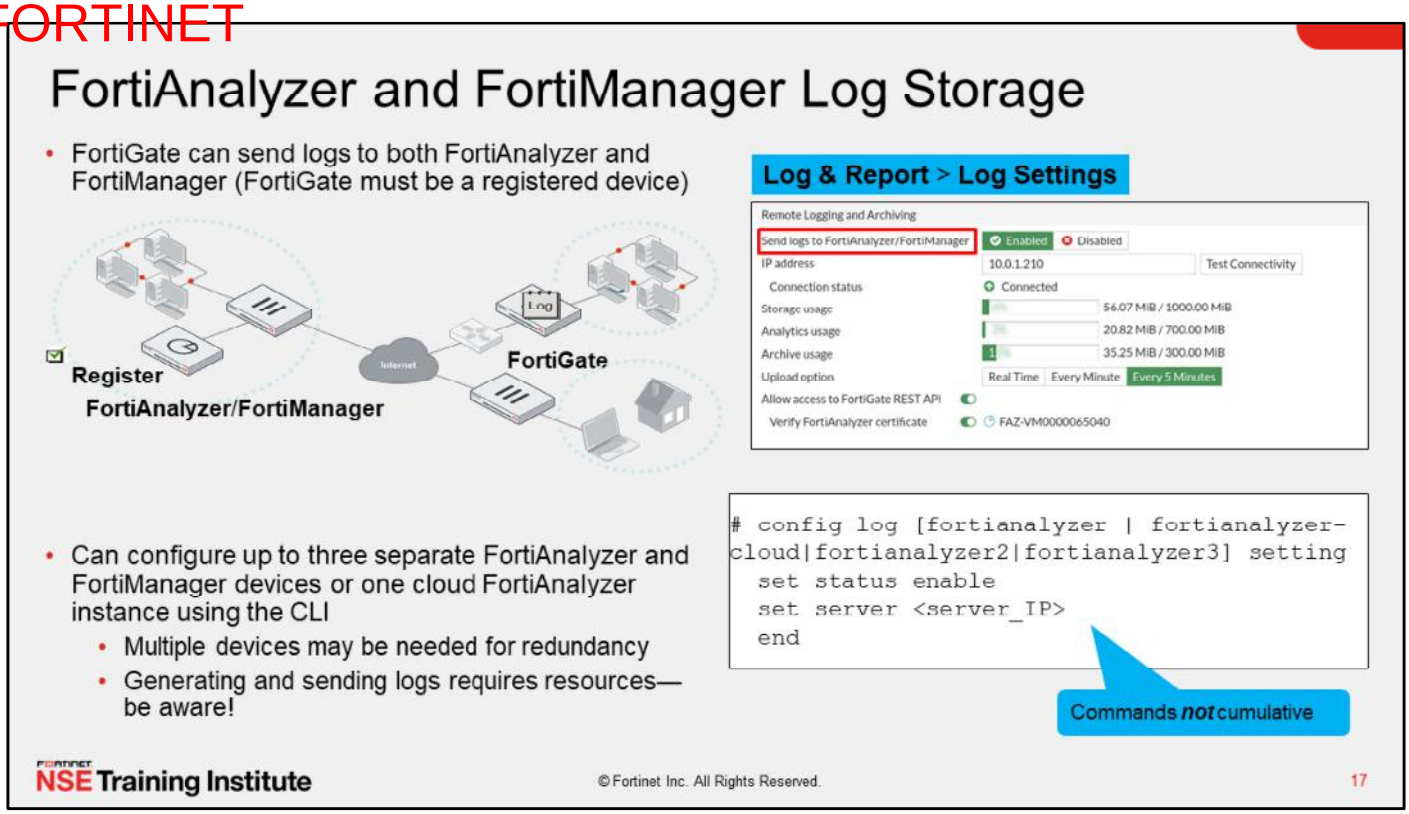

The process to configure FortiGate to send logs to FortiAnalyzer or FortiManager is identical. In order for FortiGate to send logs to either device, you must register FortiGate with FortiAnalyzer or FortiManager. After it is registered, FortiAnalyzer or FortiManager can begin to accept incoming logs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

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

- GUI: On the **Log Settings** page, enable logging to FortiAnalyzer/FortiManager, and type the IP address of the remote logging device.
- CLI: For both FortiAnalyzer and FortiManager, use the `config log fortianalyzer setting` command. Even though FortiManager isn't explicitly mentioned in the command, it is used for FortiManager as well. Using the CLI, up to three separate devices or one cloud FortiAnalyzer instance can be added to increase redundancy for the protection of log data. The commands for the three devices are not cumulative. Generating logs uses system resources, so if FortiGate frequently creates and sends logs to multiple places, CPU and RAM usage increase.

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.

Upload Option

- Near real-time uploading and consistent high-speed compression and analysis
- Configure logging options:
  - `store-and-upload` (CLI configuration only)
  - **Real Time**
  - **Every Minute**
  - **Every 5 Minutes** (default)

Log & Report > Log Settings

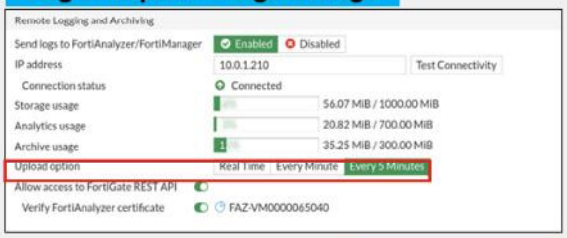


```
# configure log fortianalyzer setting
    set upload-option [store-and-upload |realtime/1-minute/5-minute]
```

`store-and-upload` only available to FortiGates with an internal hard drive

- By default, if the FortiAnalyzer disk is full, the oldest logs are overwritten; however, you can configure FortiAnalyzer to stop logging

NSE Training Institute

18

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.

FortiAnalyzer Temporarily Unavailable to FortiGate?

- The FortiGate *miglogd* process caches logs on FortiGate when FortiAnalyzer is not reachable
- When maximum cached value is reached, miglogd will drop cached logs (oldest first)
- When FortiAnalyzer connection is back, *miglogd* will send the cached logs
  - FortiGate buffer will keep logs long enough to sustain a reboot of FortiAnalyzer, but is not intended for lengthy outages
- FortiGate devices with an SSD have a configurable log buffer

```
Local-FortiGate # diagnose test application miglogd 6

mem=0, disk=0, alert=0, alarm=0, sys=0, faz=19, faz-cloud=0, webt=0, fds=0
interface-missed=0
queues in all miglogds: cur:0  total-so-far:153
global log dev statistics:
faz 0: sent=15, failed=0, cached=0, dropped=0 , relayed=0
```

Current cache size and total cache size

If there are bursts or link is overloaded, `failed` increases

```
Local-FortiGate # diagnose log kernel-stats

fgtlog: 1
fgtlog 0: total-log=32, failed-log=0 log-in
queue=0
```

If queue is full, `failed-log` value increases

NSE Training Institute

19

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 is able to 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.

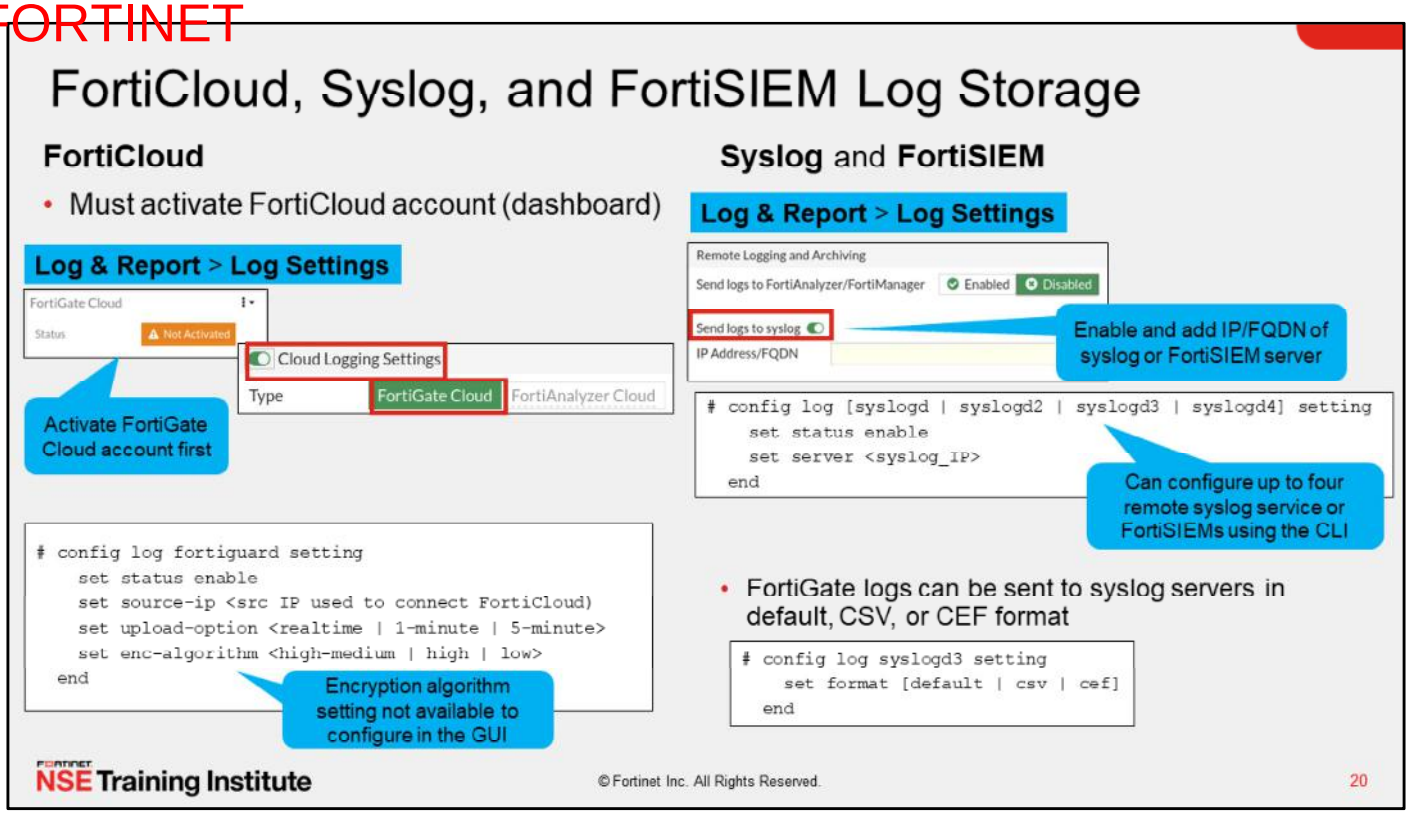

Similar to FortiAnalyzer and FortiManager, you can configure remote logging to FortiCloud on the **Log Settings** page or the CLI. However, you must first activate your FortiCloud account, so FortiGate can communicate with your FortiCloud account. Once complete, you can enable FortiCloud logging and 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, 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.

VDOMs and Remote Logging

- If you have a FortiGate with Virtual Domains (VDOMs) configured, you can globally add multiple FortiAnalyzers and syslog servers.
  - Up to three FortiAnalyzer devices
  - Up to four syslog servers

```
# config global
   config log fortianalyzer setting
       set status enable
       set server 10.0.1.1
   end
   config log fortianalyzer2 setting
      set status enable
      set server 10.0.2.1
   end
```

If override FAZ/Syslog needed, must enable it from VDOM level

```
# config vdom
  edit Training
    config log setting
      set faz-override enable
      set syslog-override enable
  end
```

NSE Training Institute

21

If you have a FortiGate with virtual domains (VDOMs) configured, you can globally add multiple FortiAnalyzers 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 (or TCP 514, if reliable logging is enabled) for log transmission

```
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
```

Controls encryption algorithm

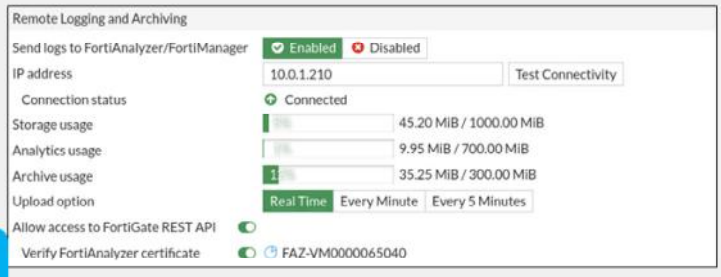

- Log messages are stored on disk and transmitted to FortiAnalyzer as plain text in LZ4 compressed format
  - Reduces disk log size and reduces log transmission time and bandwidth usage

NSE Training Institute

22

FortiGate uses UDP port 514 (or TCP port 514, if reliable logging is enabled) for log transmission.

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.

Reliable Logging and OFTPS

- Changes the log transport delivery method from UDP to TCP
- TCP provides reliable data transfer
- If you enable logging to FortiAnalyzer using the GUI, reliable logging is auto-enabled
  - If you enable logging to FortiAnalyzer using the CLI, reliable logging is not auto-enabled. You must manually enable using the CLI command:
- FortiCloud uses TCP, and you can set the encryption algorithm using the CLI (default setting is high)
- If using reliable logging, you can encrypt communications using SSL-secured OFTP (OFTPS)

```
# config log fortianalyzer setting
    set status enable
    set enc algorithm [high medium | high | low]
    set reliable enable
end
```

Reliable logging must be enabled to use OFTPS

NSE Training Institute

23

When you enable reliable logging on FortiGate, the log transport delivery method changes from UDP (User Datagram Protocol) to TCP (Transmission Control Protocol). TCP provides reliable data transfer, guaranteeing that the data transferred 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 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 encrypt communications using SSL-secured OFTP by configuring the `enc-algorithm` setting on the CLI.

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

NSE Training Institute

24

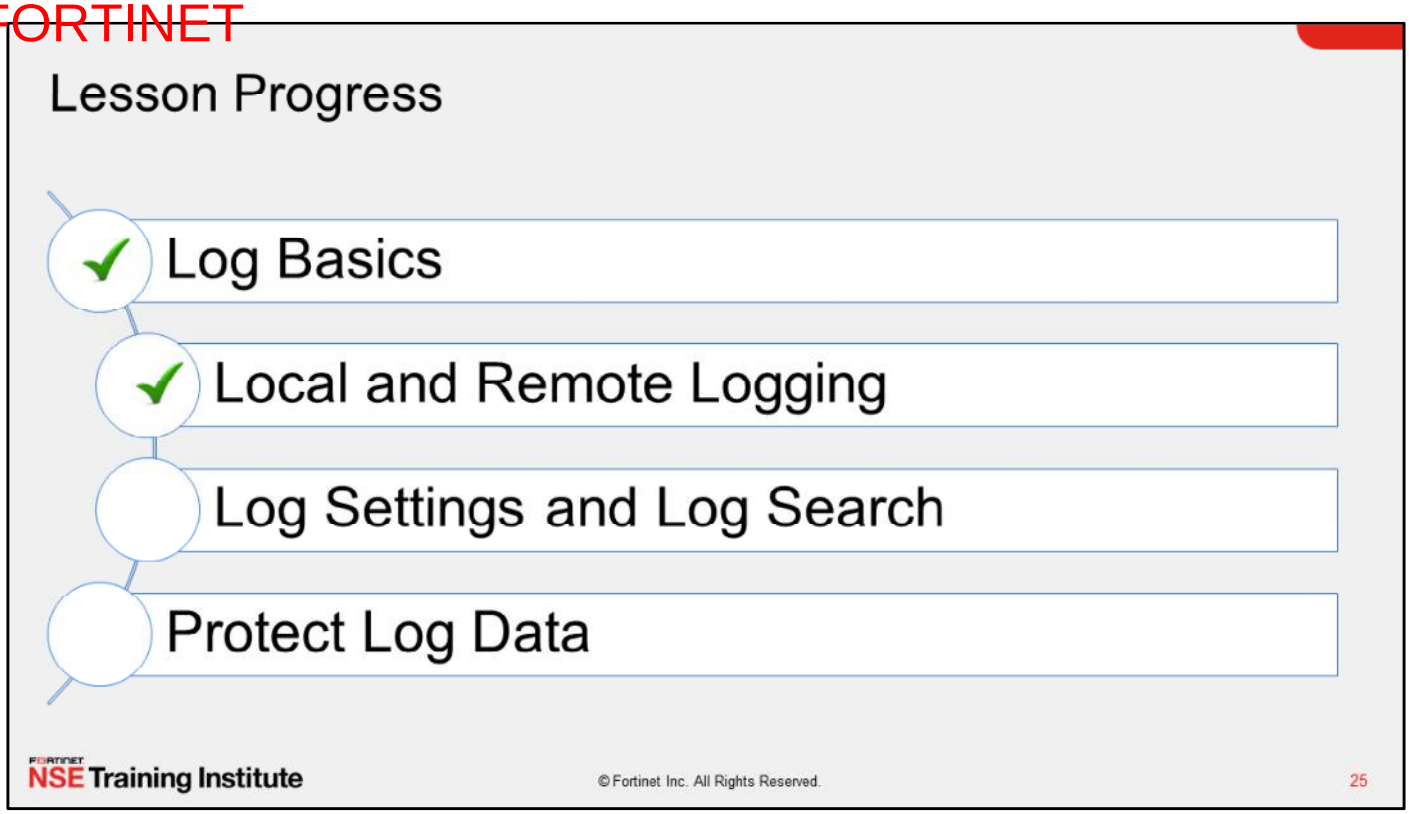

Good job! You now understand remote logging.

Now, you will learn about log settings.

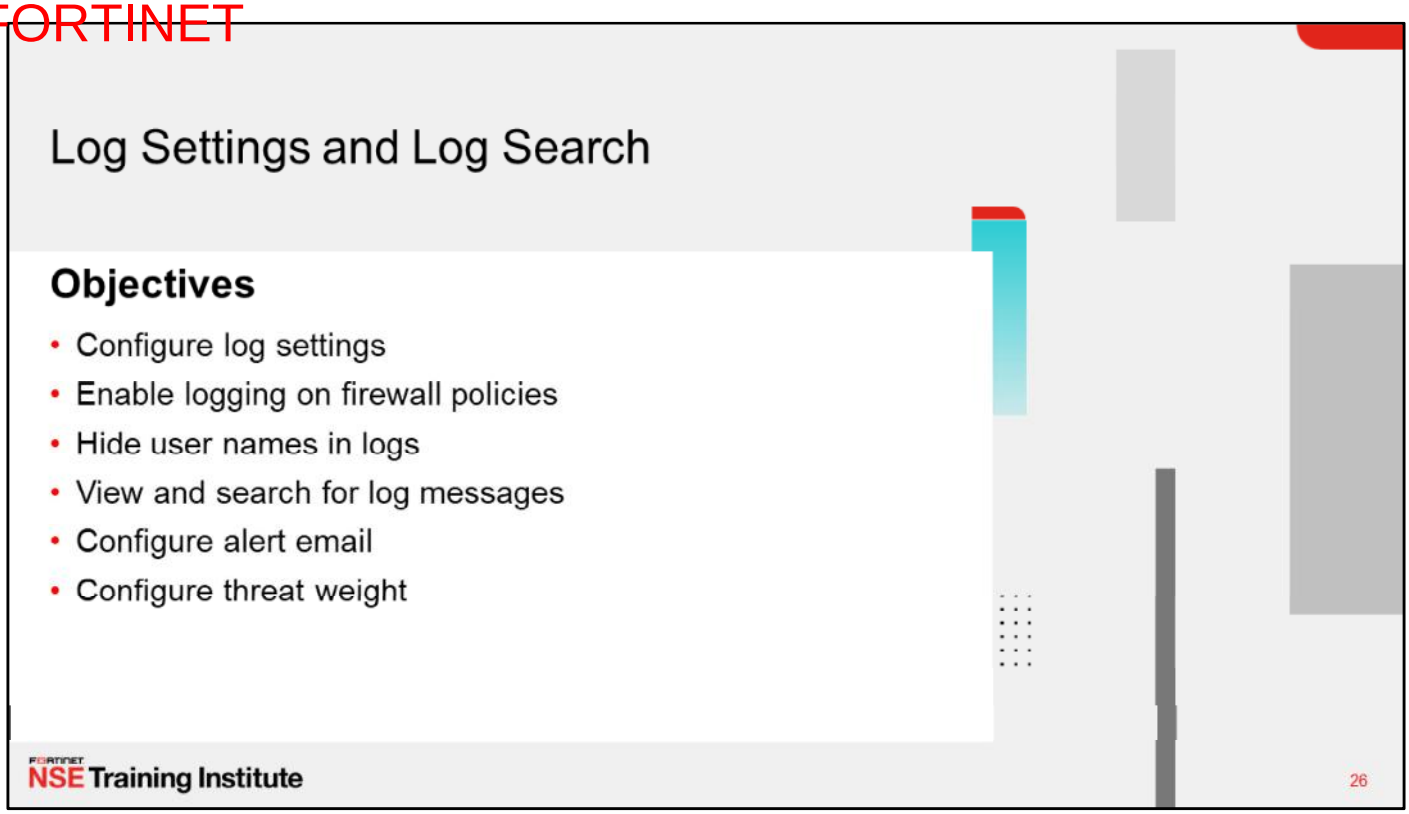

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

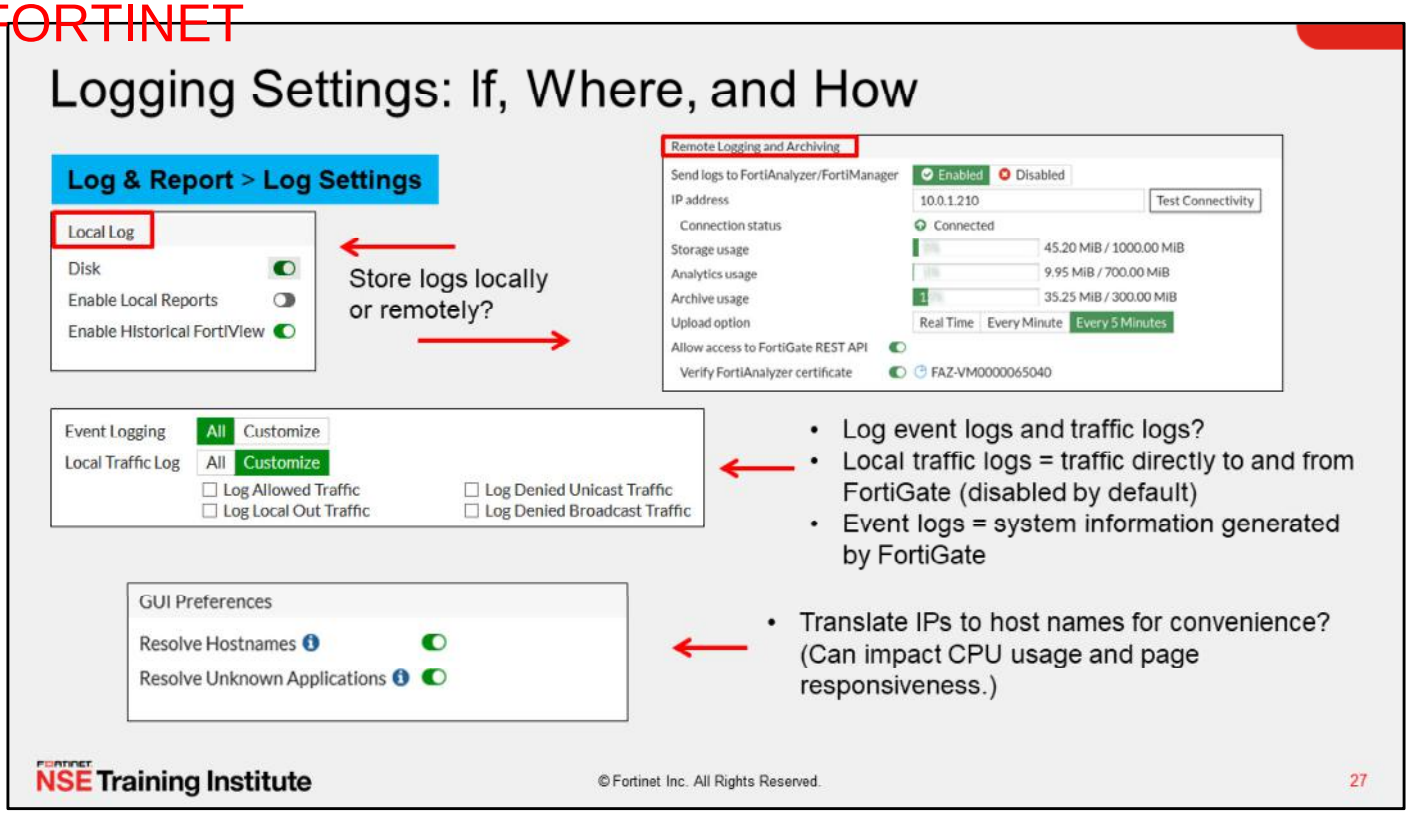

The **Log Settings** page allows you to decide if, where, and how a log is stored.

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. Local traffic logs provide information about traffic directly to and from FortiGate. 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. For example, IPsec VPNs closing, or routing protocol activity, are not caused by traffic passing through a firewall policy. One exception might be the user log, because it does record user login and logout events on traffic that passes through 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.

Log Filtering

- Configure log filter settings to determine which logs are recorded
  - Configure up to four remote syslog or FortiSIEM logging servers:

```
# config log [syslogd | syslogd2 | syslogd3 | syslogd4] filter
```

  - Configure up to three FortiAnalyzer or FortiManager devices or one cloud FortiAnalyzer instance:

```
# config log [fortianalyzer | fortianalyzer-cloud | fortianalyzer2 |
fortianalyzer3] filter
```

- Filters include:
  - Severity <level>
  - Forward traffic [enable/disable]
  - Local traffic [enable/disable]
  - Multicast traffic [enable/disable]
  - Sniffer traffic [enable/disable]
  - Anomaly [enable/disable]
  - VOIP [enable/disable]
  - DLP archive [enable/disable]
  - GTP [enable/disable]
  - Filter [string]
  - Filter type [include | exclude]

NSE Training Institute © Fortinet Inc. All Rights Reserved. 28

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

Previously, we mentioned that you can configure up to four logging services for syslog and FortiSIEM using the command `config log syslogd setting`, and up to four FortiAnalyzer or FortiManager devices using the `config log fortianalyzer setting`. You can control what 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`.

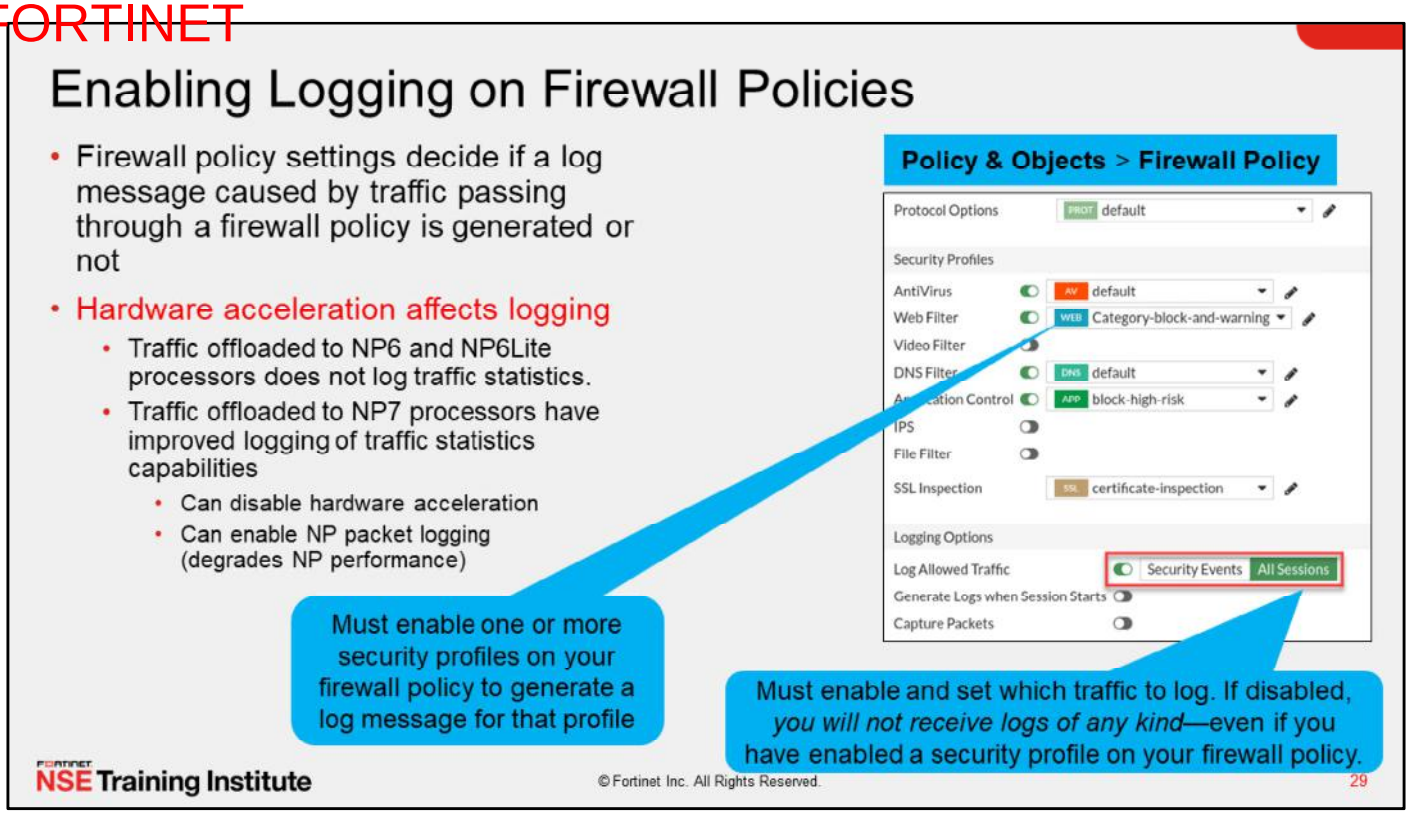

After you configure all logging settings, you can enable logging on your firewall policies. Only when enabled on a firewall policy can a log message—caused by traffic passing through that firewall policy—generate.

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 is also enabled, security log events appear in the forward traffic log and security log.

Hiding User Names in Logs

- Some laws require that usernames be anonymized
- Use the following command to hide usernames in traffic and UTM logs, so that the username appears as `anonymous`

```
# config log setting
   set user-anonymize enable
  end
```

```
date=2021-03-16 time=14:45:16 logid=0317013312 type=utm subtype=webfilter
eventtype=ftgd_allow level=notice vd="root" policyid=2 identidx=1
sessionid=31232959 user="anonymous" group="ldap_users" srcip=192.168.1.24
srcport=63355 srcintf="port2" dstip=66.171.121.44 dstport=80 dstintf="port1"
service="http" hostname="www.fortinet.com" profiletype="Webfilter_Profile"
profile="default" status="passthrough" reqtype="direct" url="/" sentbyte=304
rcvdbyte=60135 msg="URL belongs to an allowed category in policy" method=domain
class=0 cat=140 catdesc="custom1"
```

NSE Training Institute

30

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.

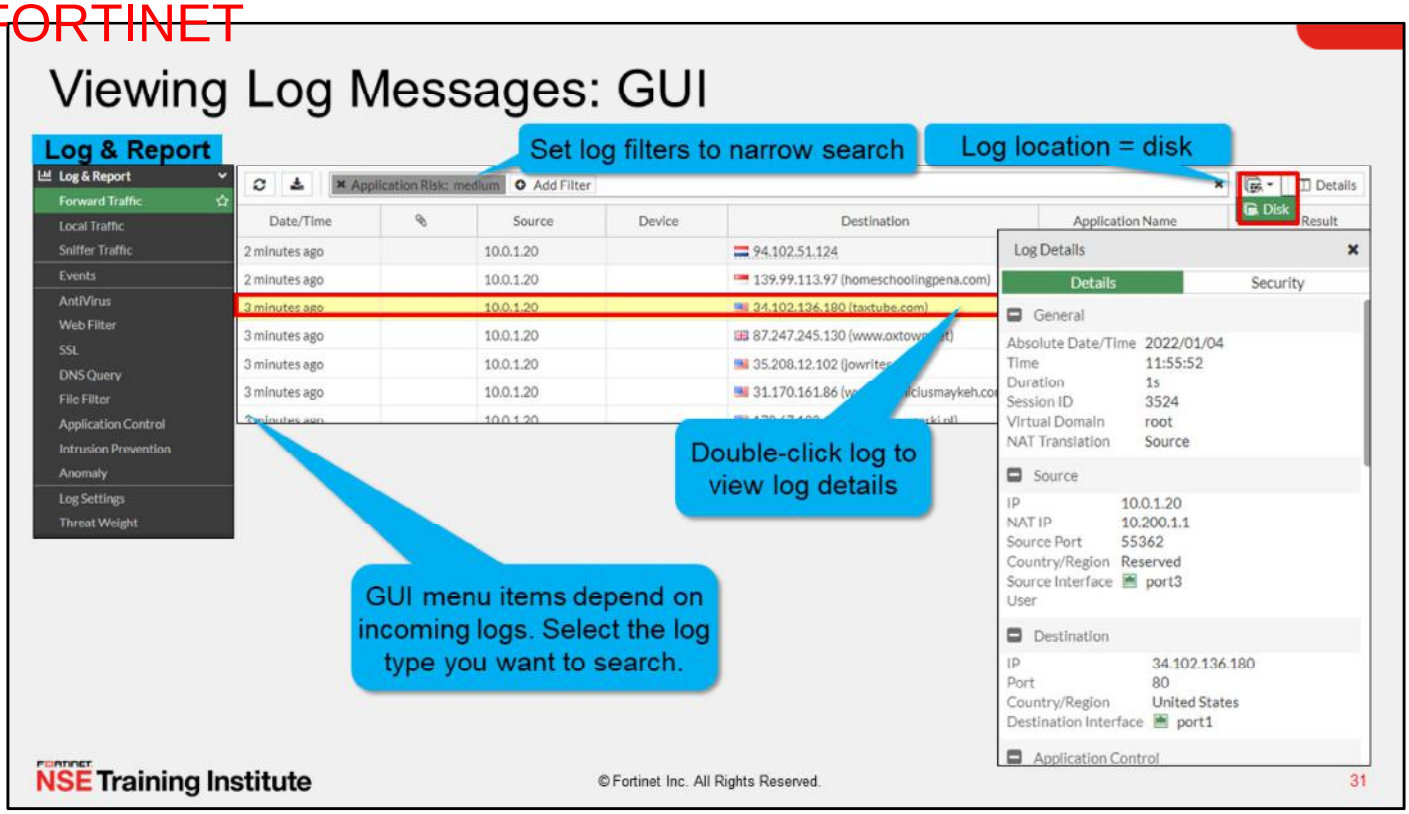

You can access your logs on the GUI in the **Log & Report** menu. The options that appear in this menu depend on your configuration. Security logs appear only if security events exist.

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 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 within 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.

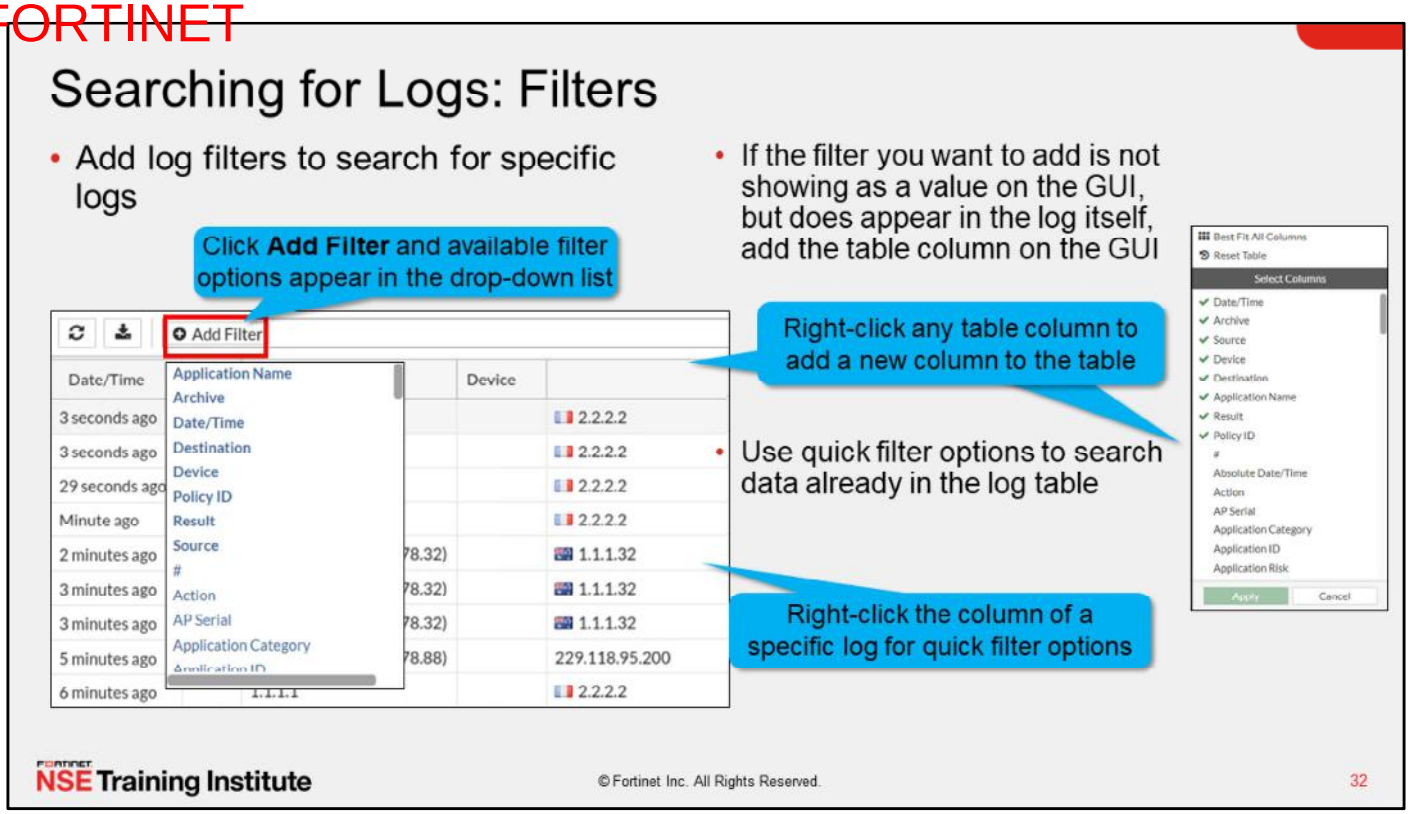

Depending on your configuration, your FortiGate might record a high volume of logs. This can make it more difficult to locate a specific log, especially during an investigation.

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 from the drop-down list that appears. If you see data that you want to filter on in a log in the table already, 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 appears that lets you filter on all logs with that botnet name.

By default, the most common columns are shown and less common columns are hidden. Accordingly, if filtering data based on a column that is hidden, be sure to add the column as a selected column. To add columns, right-click any column field, and, in the pop-up menu that appears, select the column in the **Available Columns** section.

If your search filters don't return any results when the log data does exist, the filter may be poorly formed. FortiGate looks for an exact match in the log, so you must form the search string correctly.

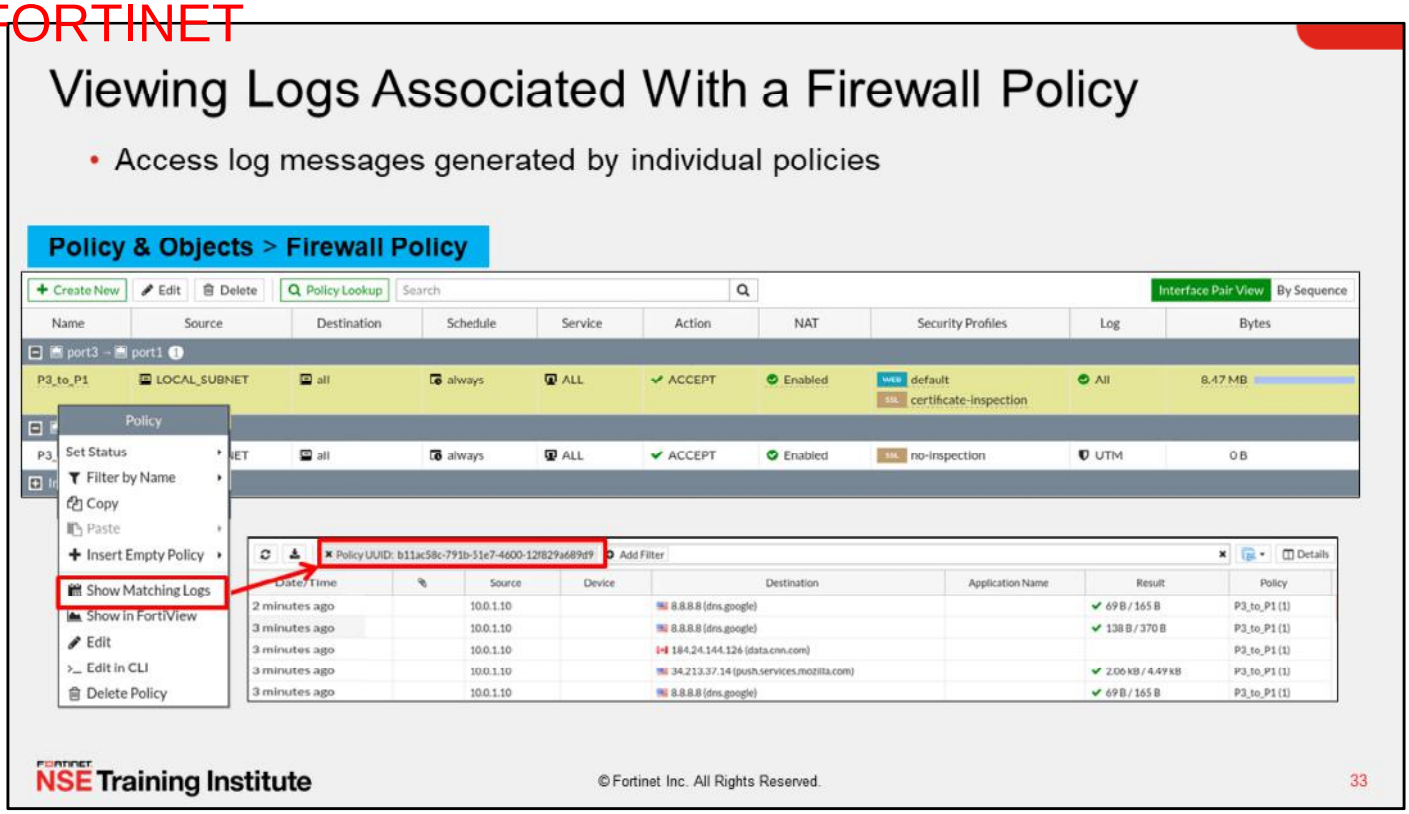

You can also access log messages generated by individual policies. Right-click the policy for which you want to view all associated logs and, in the pop-up menu, select **Show Matching Logs**. FortiGate takes you to the **Forward Traffic** page where a filter is automatically set based on the policy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

Viewing Log Message: CLI

`# execute log filter` ← 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

`# execute log display` ← Allows you to see specific log messages that you already configured within the `execute log filter` command

```
Local-FortiGate # execute log display
40 logs found.
10 logs returned.
1: date=2021-04-13 time=08:45:49 eventtime=1618328749810305885 tz="-0700" logid="0000000020"
type="traffic" subtype="forward" level="notice" vd="root" srcip=10.0.1.10 srcport=40570
srcintf="port3" srcintfrole="undefined" dstip=74.6.143.25 dstport=443 dstintf="port1"
dstintfrole="undefined" srccountry="Reserved" dstcountry="United States" sessionid=4201 proto=6
action="accept" policyid=1 policytype="policy" poluuid="b11ac58c-791b-51e7-4600-12f829a689d9"
policyname="Full Access" service="HTTPS" trandisp="snat" transip=10.200.1.10 transport=40570
duration=153 sentbyte=6623 rcvdbyte=23201 sentpkt=40 rcvdpkt=40 appcat="unscanned"
sentdelta=6623 rcvddelta=23201

2: date=2021-04-13 time=08:45:46 eventtime=1618328746107660006 tz="-0700" logid="0000000020"
type="traffic" subtype="forward" level="notice" vd="root" srcip=10.0.1.10 srcport=35908
srcintf="port3" srcintfrole="undefined" dstip=54.243.191.211 dstport=443 dstintf="port1"
dstintfrole="undefined" srccountry="Reserved" dstcountry="United States" sessionid=4255 proto=6
action="accept" policyid=1 policytype="policy" poluuid="b11ac58c-791b-51e7-4600-12f829a689d9"
policyname="Full Access" service="HTTPS" trandisp="snat" transip=10.200.1.10 transport=35908
duration=147 sentbyte=2932 rcvdbyte=8084 sentpkt=23 rcvdpkt=19 appcat="unscanned"
sentdelta=2932 rcvddelta=8084
```

NSE Training Institute © Fortinet Inc. All Rights Reserved. 34

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.

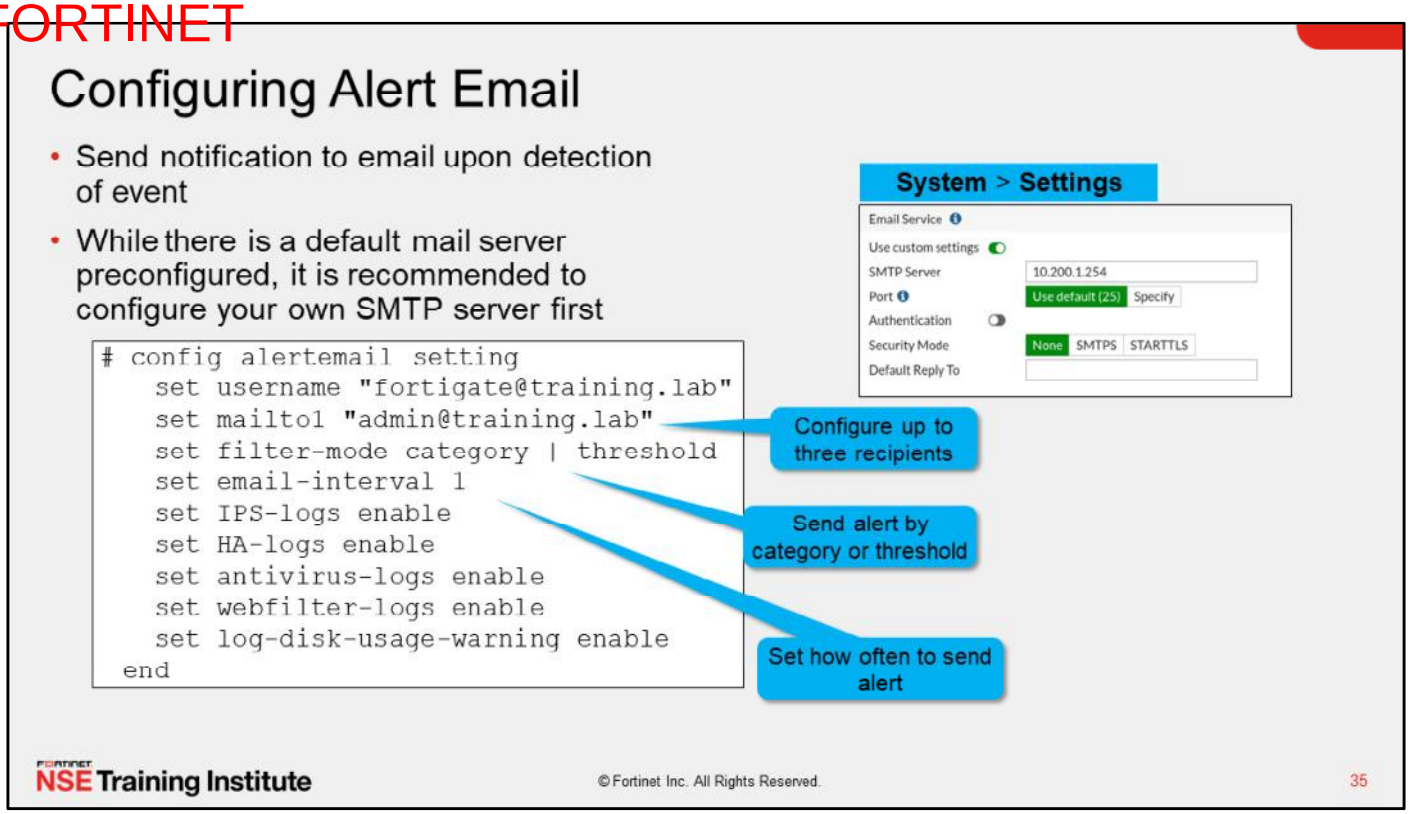

Because you can't always be physically watching the logs on the device, you can monitor events by setting up alert email. Alert emails provide an efficient and direct method of notifying an administrator of events.

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.

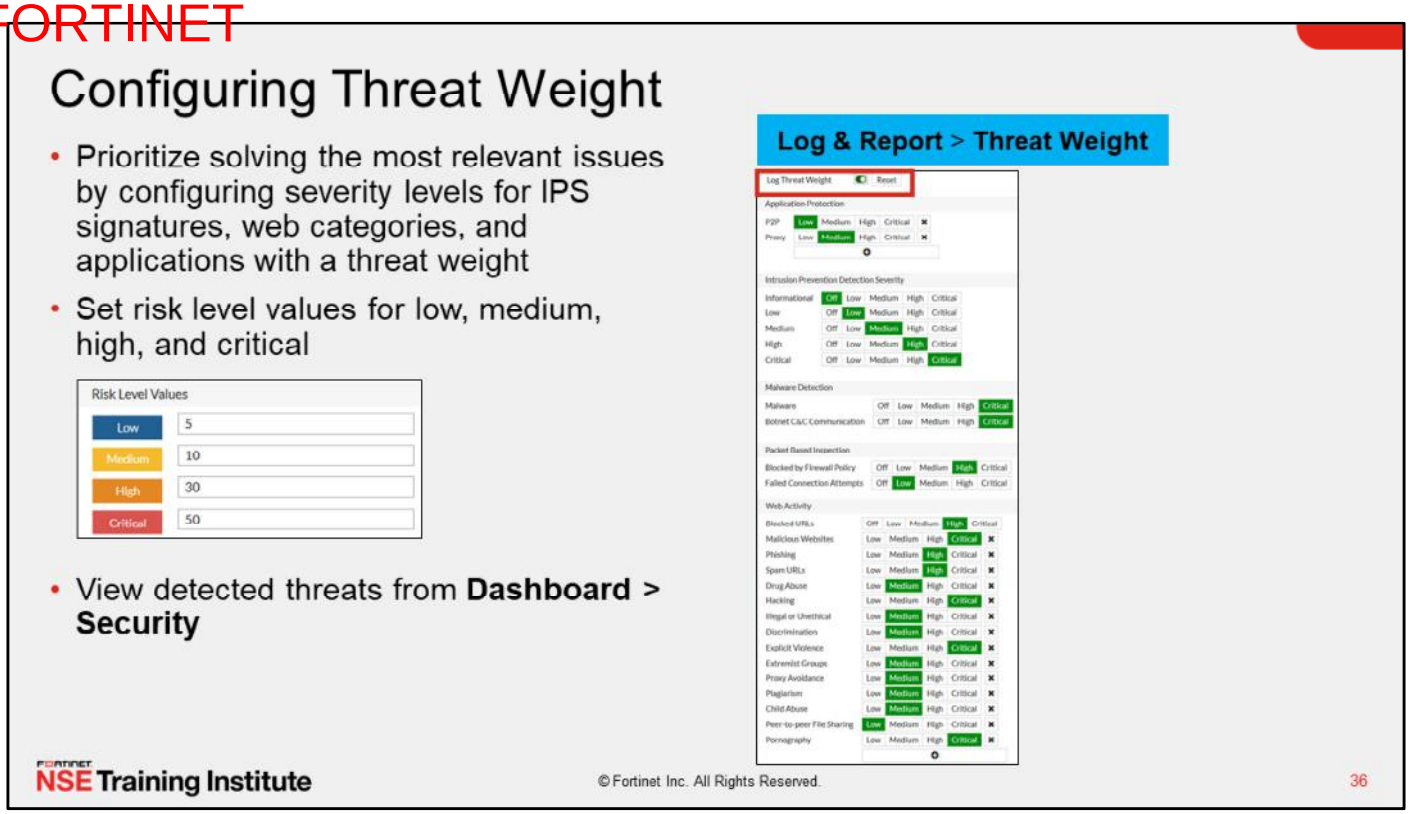

In order to prioritize solving the most relevant issues easily, you can configure severity levels for IPS signatures, web categories, and applications that are associated with a threat weight (or score).

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 on 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

NSE Training Institute

37

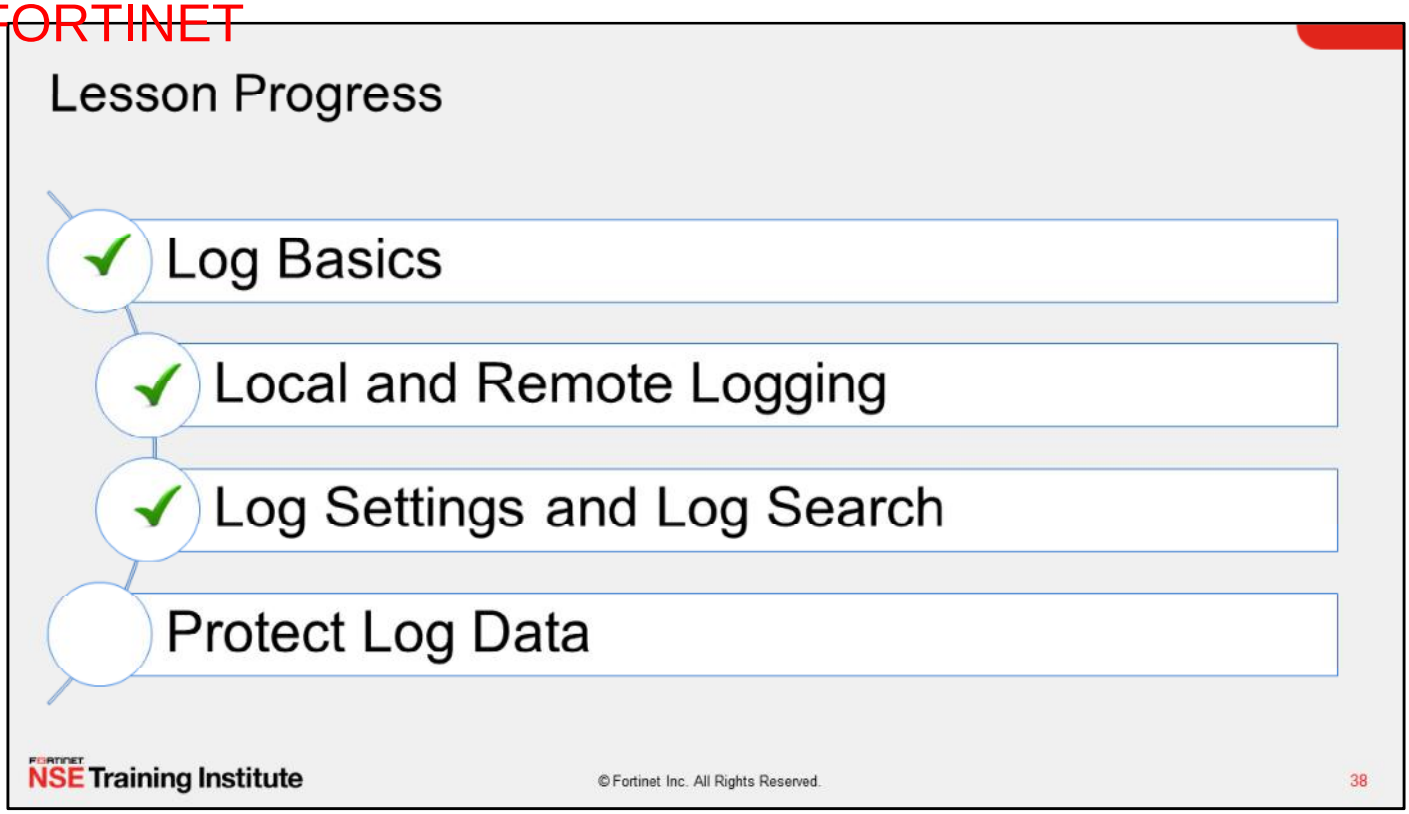

Good job! You now understand how to troubleshoot communication issues.

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

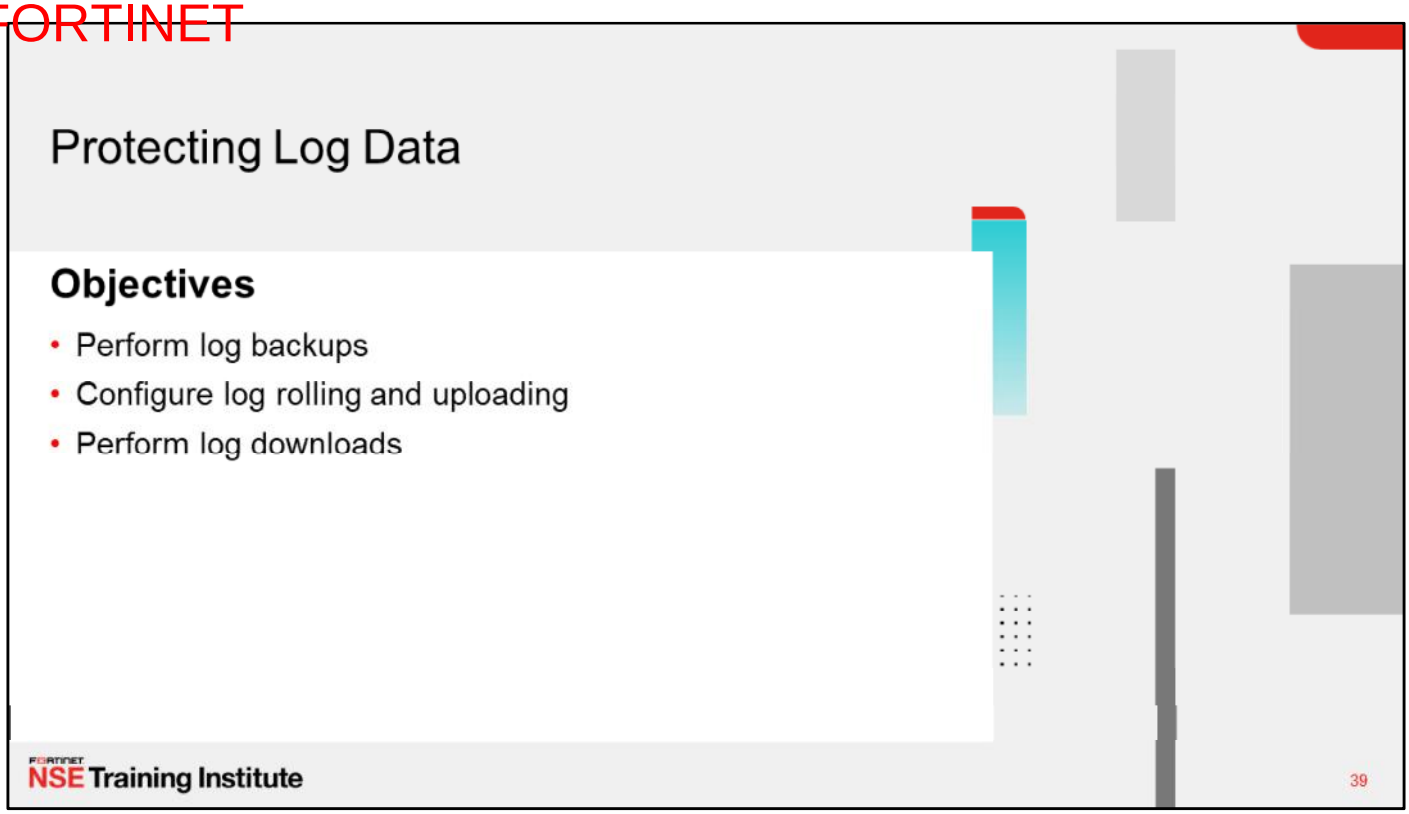

After completing this section, you should be able to achieve the objectives shown on this slide.

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

Backing Up Logs

- Export all logs to FTP, TFTP, or USB (stored as LZ4 compressed files):

```
# execute backup disk alllogs [ftp | tftp | usb]
```

- Export specific log type to FTP, TFTP, or USB (stored as LZ4 compressed files)

```
# execute backup disk log [ftp | tftp | usb] <log_type>
```

These backups cannot be restored to FortiGate devices

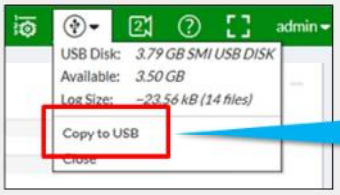

Appears as option in GUI when you insert a USB drive into the FortiGate USB port

NSE Training Institute

40

You can also protect your log data by performing log backups, which is to say copying 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 `execute backup disk log <log type>` 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 restore FortiGate backup logs to a FortiAnalyzer device and then view on both FortiGate and FortiAnalyzer devices.

You can also back up logs to USB using the GUI. The GUI menu item appears when you insert a USB drive into a FortiGate USB port.

Log Rolling and Uploading

Log rolling

- Similar to zipping a file, rolling lowers space requirements needed to contain them
- Can configure max log file size to roll (default 20 MB)
- Can configure roll schedule and time

```
# config log disk setting
    set max-log-file-size <1-100>
    set roll-schedule [daily | weekly]
    set roll-time [hh:mm]
```

Log uploading

- Can configure rolled log files to upload to an FTP server
- Can specify which types of log files to upload
- Can configure an upload schedule and time (command not shown—similar to log rolling example)
- Can delete log files after uploading (enabled by default)

```
# config log disk setting
    set upload [enable | disable]
    set upload-destination [FTP]
    set uploadip [IPv4 IP]
    set uploadport [integer]
    set source-ip [source IPv4 IP]
    set uploaduser [FTP user]
    set uploadpass [FTP user password]
    set uploaddir [remote FTP dir]
    set uploadtype [log type]
    set upload-delete-files [enable* | disable]
```

NSE Training Institute

41

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.

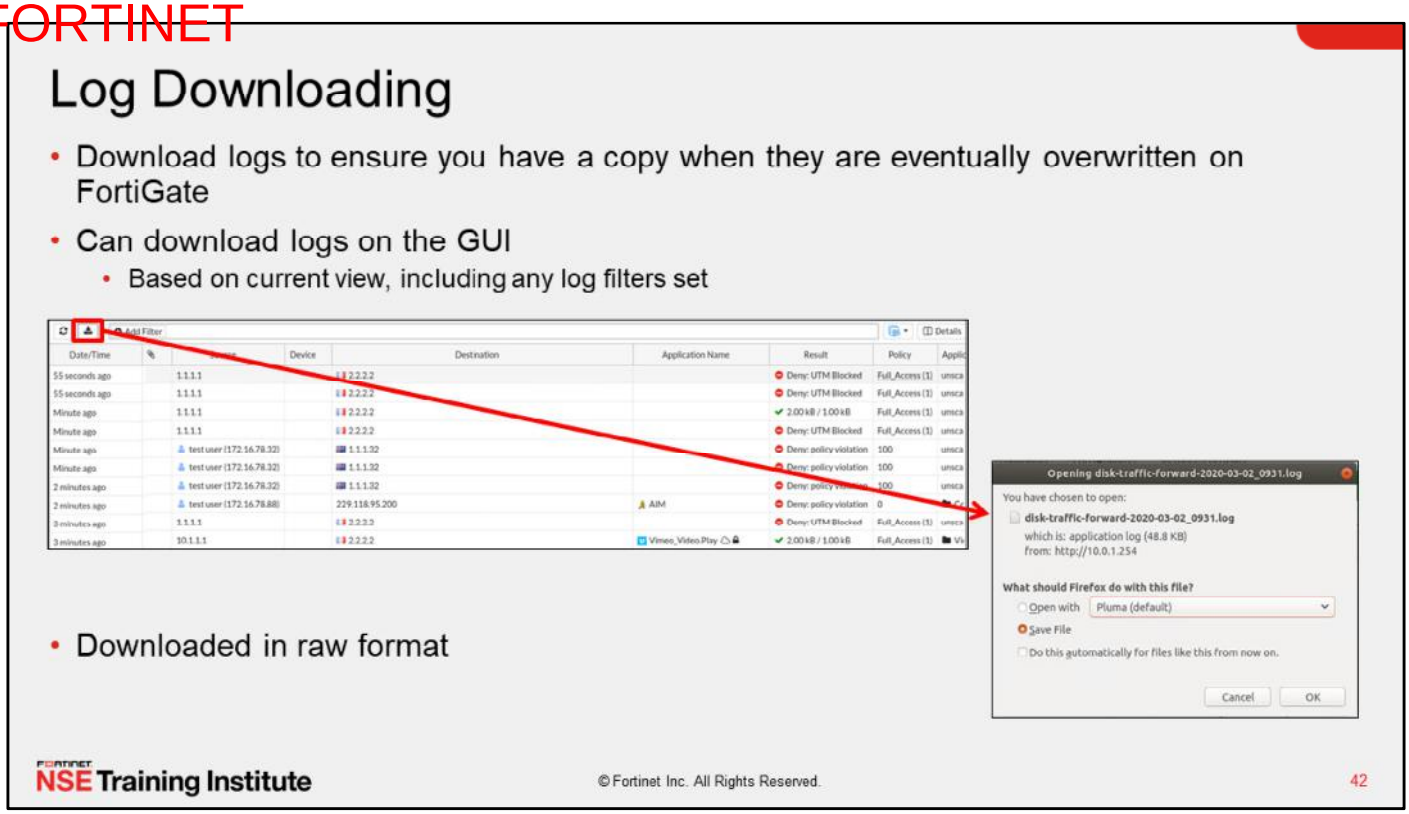

You can also download a copy of the logs from FortiGate and save them on a server or on a computer to view and access later. This ensures that you still have a copy when the originals are eventually overwritten on FortiGate.

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 the log messages on the GUI, you are downloading log messages in the raw format.

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.

NSE Training Institute

43

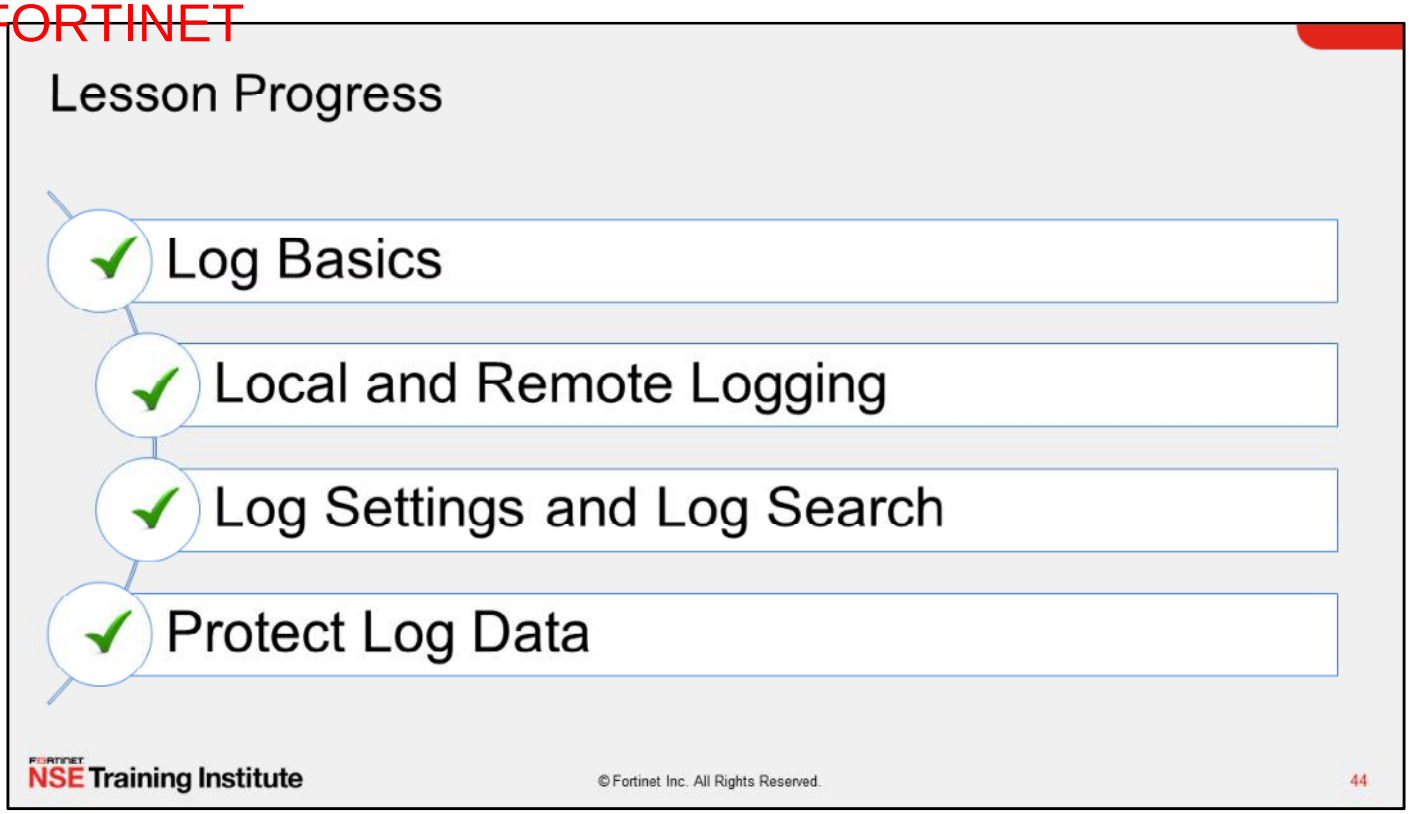

Congratulations! You have completed this lesson. Now, you will review the topics that you covered in this lesson.

Review

- ✓ Understand log basics
- ✓ Describe the effect of logging on performance
- ✓ Identify log storage options
- ✓ Configure local and remote logging
- ✓ Understand disk allocation and reserved space
- ✓ Identify external log storage options
- ✓ Configure remote logging
- ✓ Understand log transmission and how to enable reliable logging and OFTPS
- ✓ Configure logging settings
- ✓ Understand miglogd
- ✓ View and search for log messages on the GUI and CLI
- ✓ View logs on FortiView
- ✓ Configure alert email and threat weight
- ✓ Configure log backups, rolling, uploading, downloading

NSE Training Institute

45

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.

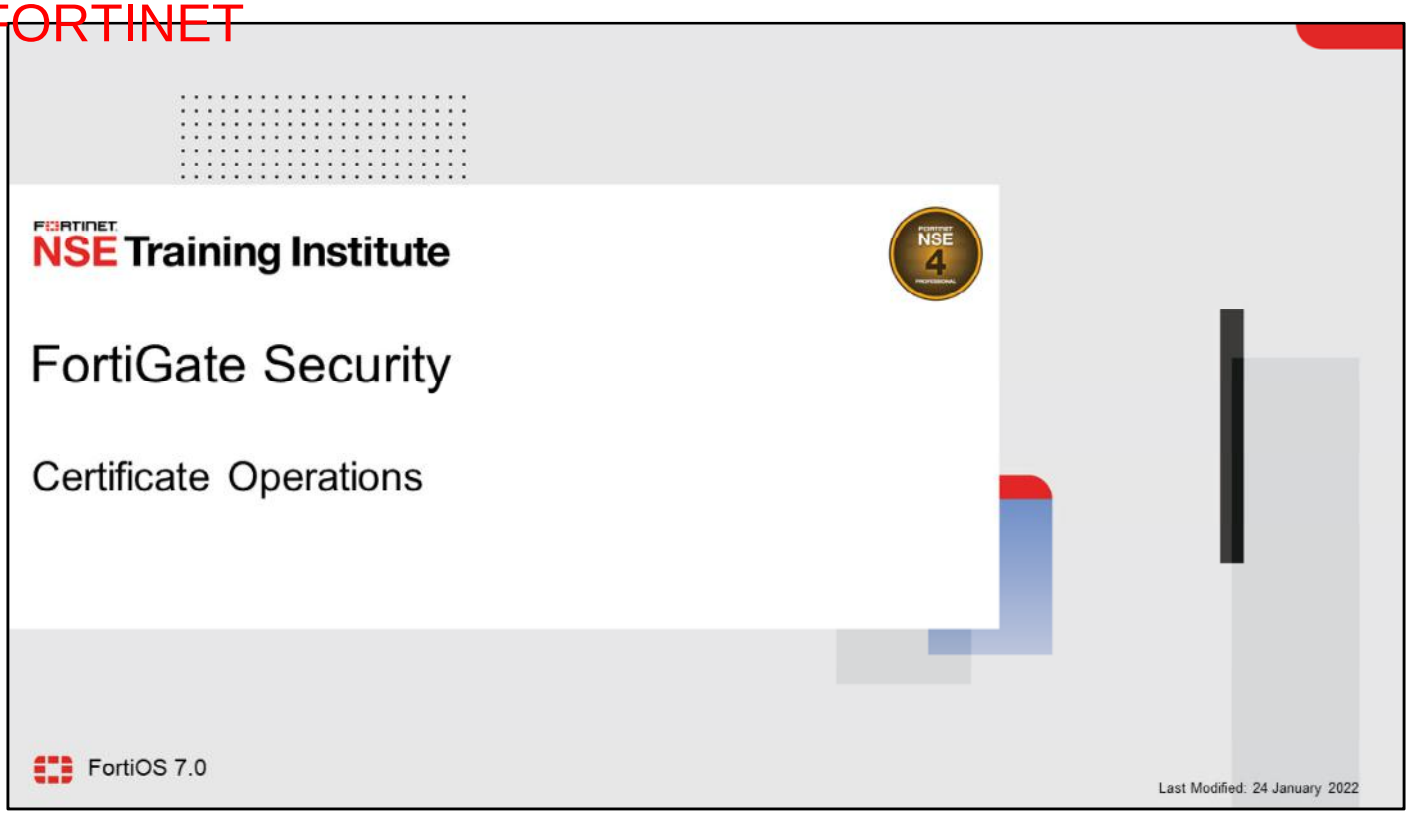

In this lesson, you will learn why FortiGate uses digital certificates, how to configure FortiGate to use certificates (including using certificates to inspect the contents of encrypted traffic), and how FortiGate manages certificates.

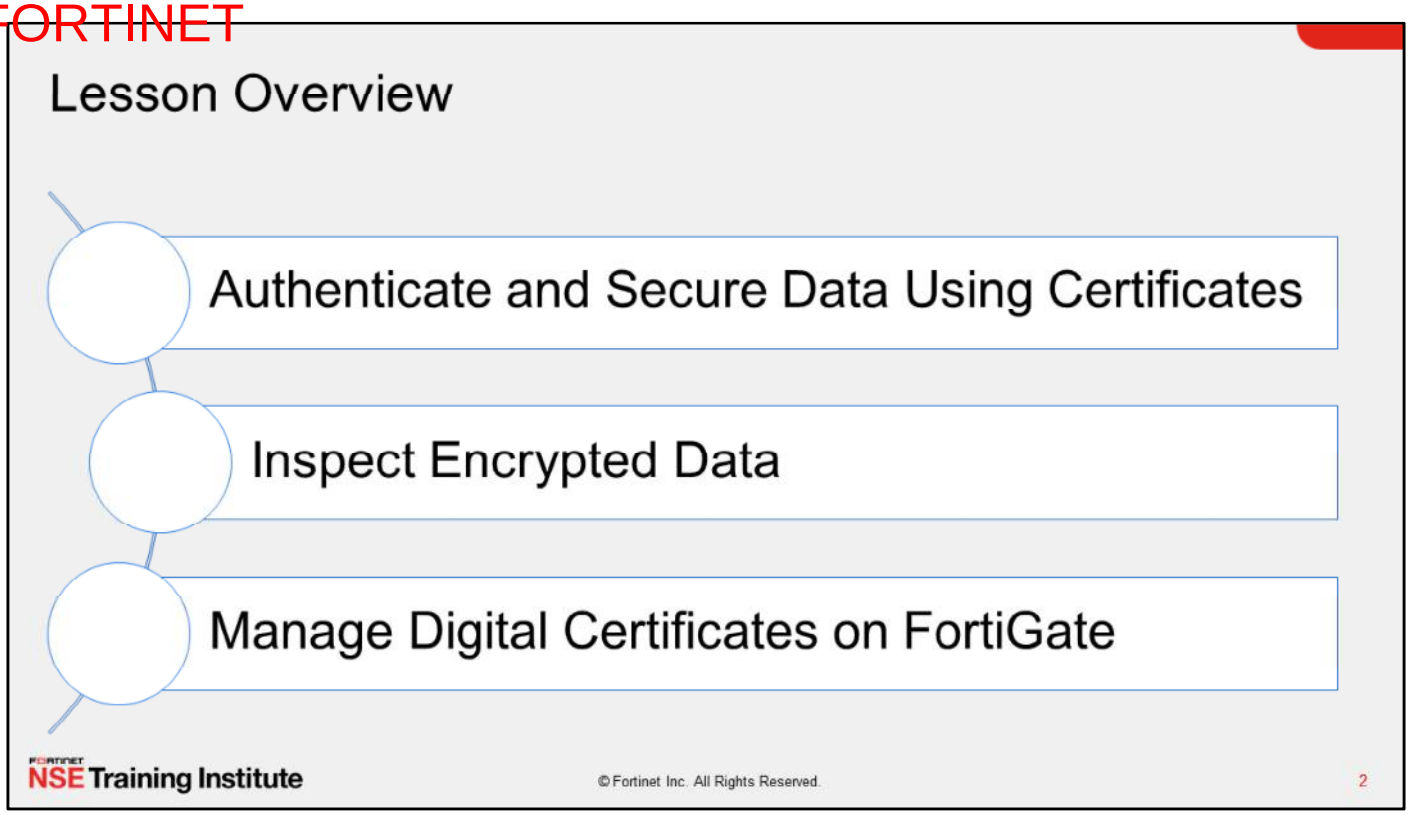

In this lesson, you will learn about the topics shown on this slide.

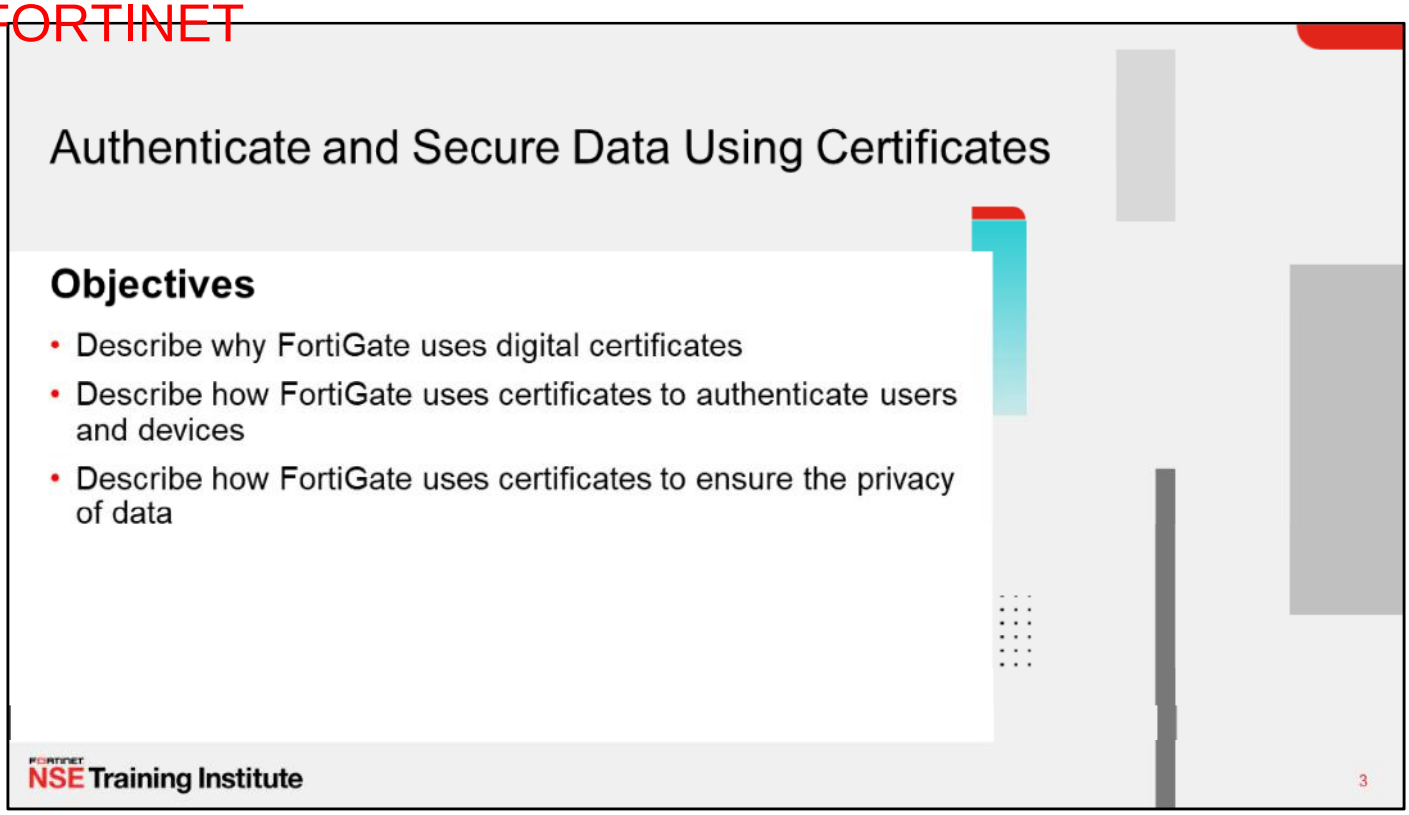

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

Why Does FortiGate Use Digital Certificates?

- Inspection
  - FortiGate dynamically generates temporary certificates to perform full SSL inspection
  - FortiGate can inspect certificates to ensure that they are trusted and valid, before permitting a client to connect to an outside device
- Privacy
  - FortiGate uses digital certificates, and their associated private keys, to establish SSL connections with other devices, such as FortiGuard
- Authentication
  - Users who have certificates issued by a trusted certificate authority (CA), can authenticate on FortiGate to access the network or to establish a VPN connection
  - Administrator users can use certificates as second-factor authentication to log in to FortiGate

NSE Training Institute © Fortinet Inc. All Rights Reserved. 4

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 device, 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.

Using Certificates to Identify a Person or Device

- What is a digital certificate?
  - A digital identity produced and signed by a CA
  - Analogy: passport or driver's license
- How does FortiGate use certificates to identify devices and people?
  - The **Subject** and **Subject Alternative Name** fields in the certificate identify the device or person associated with the certificate
- FortiGate uses the X.509v3 certificate standard

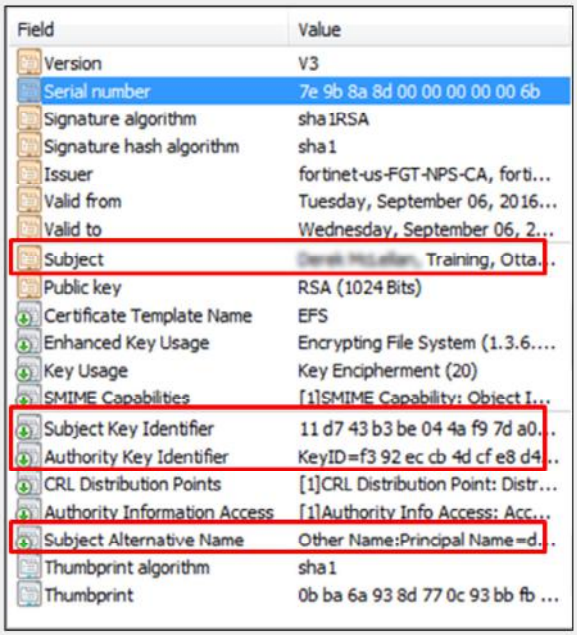

Field	Value
Version	V3
Serial number	7e 9b 8a 8d 00 00 00 00 00 6b
Signature algorithm	sha1RSA
Signature hash algorithm	sha1
Issuer	fortinet-us-FGT-NPS-CA, forti...
Valid from	Tuesday, September 06, 2016...
Valid to	Wednesday, September 06, 2...
Subject	[redacted] Training, Otta...
Public key	RSA (1024 Bits)
Certificate Template Name	EFS
Enhanced Key Usage	Encrypting File System (1.3.6....
Key Usage	Key Encipherment (20)
SMIME Capabilities	[1]SMIME Capability: Object I...
Subject Key Identifier	11 d7 43 b3 be 04 4a f9 7d a0...
Authority Key Identifier	KeyID=f3 92 ec cb 4d cf e8 d4...
CRL Distribution Points	[1]CRL Distribution Point: Distr...
Authority Information Access	[1]Authority Info Access: Acc...
Subject Alternative Name	Other Name:Principal Name=d...
Thumbprint algorithm	sha1
Thumbprint	0b ba 6a 93 8d 77 0c 93 bb fb ...

NSE Training Institute

5

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

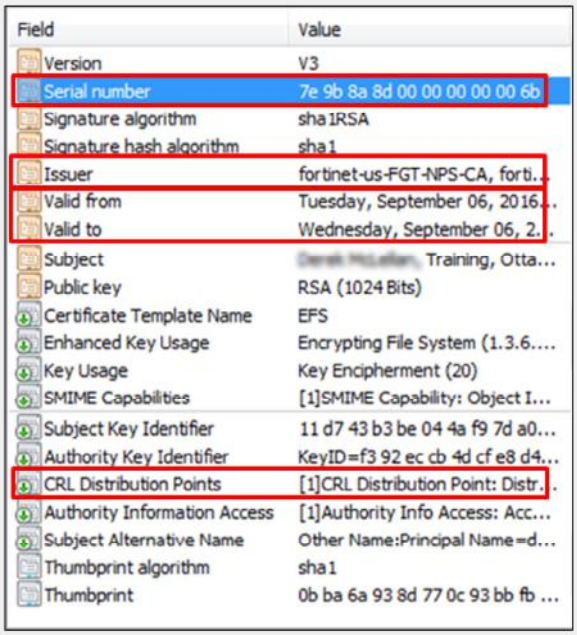

Field	Value
Version	V3
Serial number	7e 9b 8a 8d 00 00 00 00 00 6b
Signature algorithm	sha1RSA
Signature hash algorithm	sha1
Issuer	fortinet-us-FGT-NPS-CA, forti...
Valid from	Tuesday, September 06, 2016...
Valid to	Wednesday, September 06, 2...
Subject	[redacted] Training, Otta...
Public key	RSA (1024 Bits)
Certificate Template Name	EFS
Enhanced Key Usage	Encrypting File System (1.3.6....
Key Usage	Key Encipherment (20)
SMIME Capabilities	[1]SMIME Capability: Object I...
Subject Key Identifier	11 d7 43 b3 be 04 4a f9 7d a0...
Authority Key Identifier	KeyID=f3 92 ec cb 4d cf e8 d4...
CRL Distribution Points	[1]CRL Distribution Point: Distr...
Authority Information Access	[1]Authority Info Access: Acc...
Subject Alternative Name	Other Name:Principal Name=d...
Thumbprint algorithm	sha1
Thumbprint	0b ba 6a 93 8d 77 0c 93 bb fb ...

NSE 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.

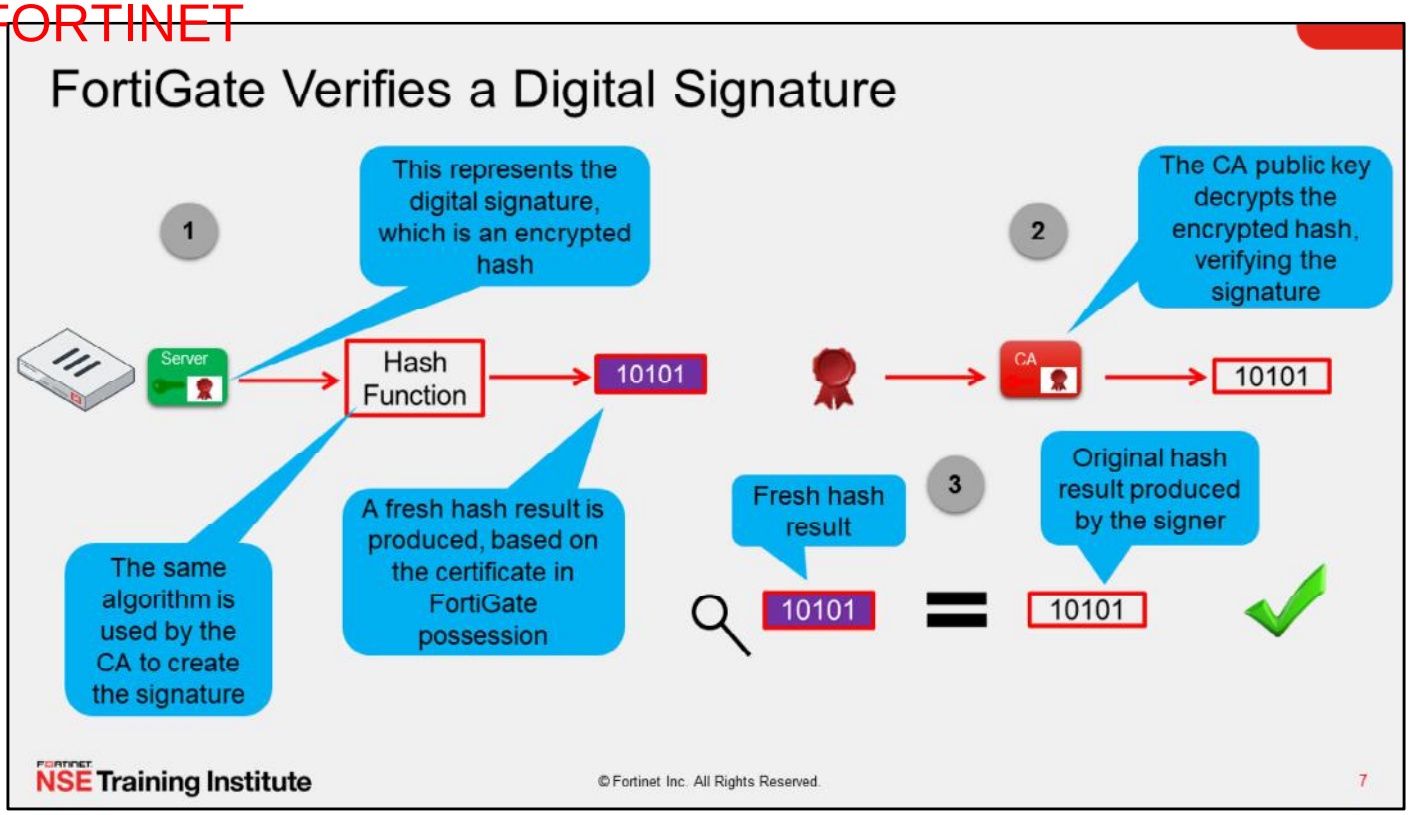

Before it generates a digital signature, the CA runs the content of the certificate through a hash function, which produces a hash result. The hash result, which is a mathematical representation of the data, is referred to as the *original hash result*. The CA encrypts the original hash result using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

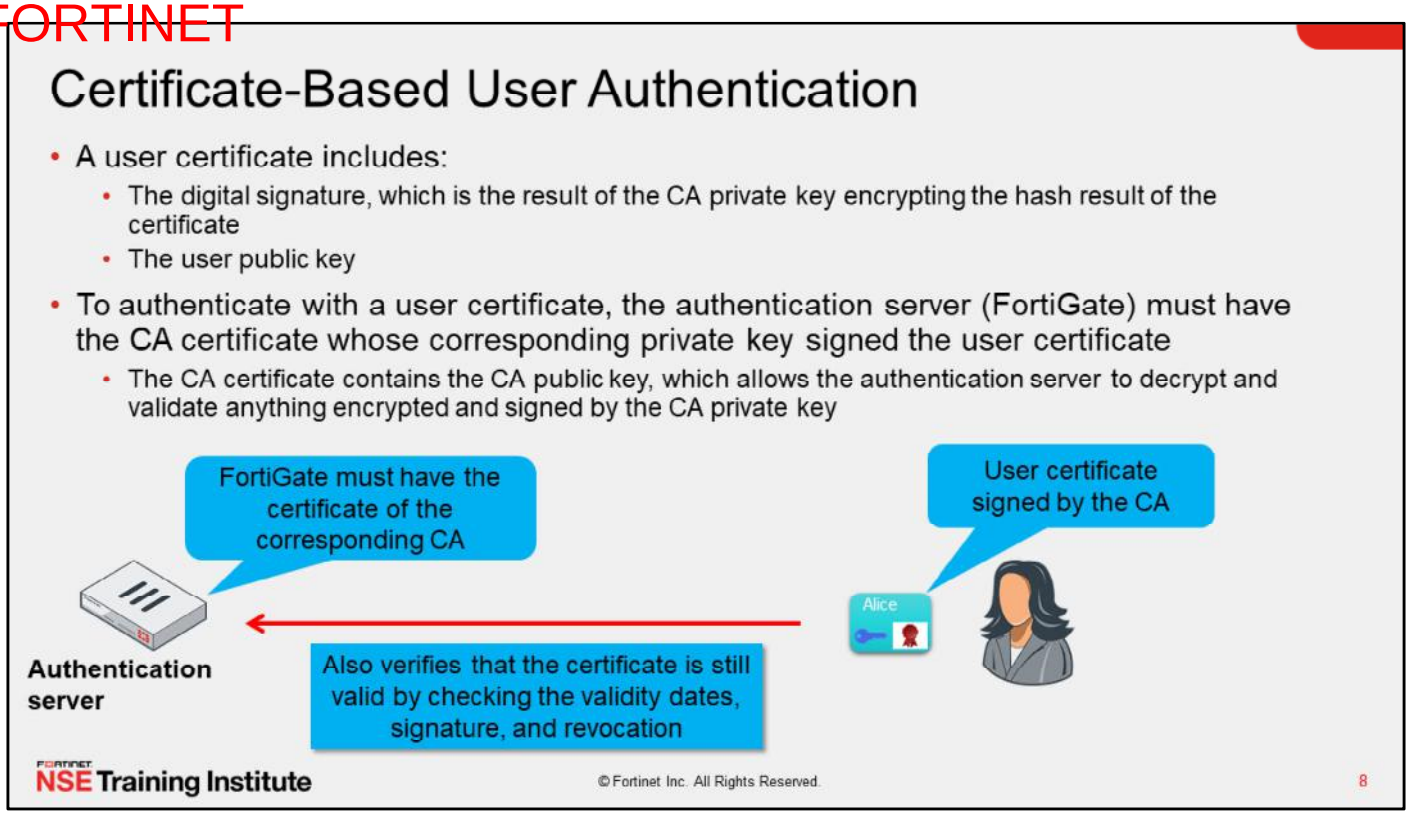

Certificate-based user authentication uses an end-entity certificate to identify the user. This certificate contains the user public key and the signature of the CA that issued the certificate. The authentication server (for example, FortiGate) must have the CA certificate whose private key signed the 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.

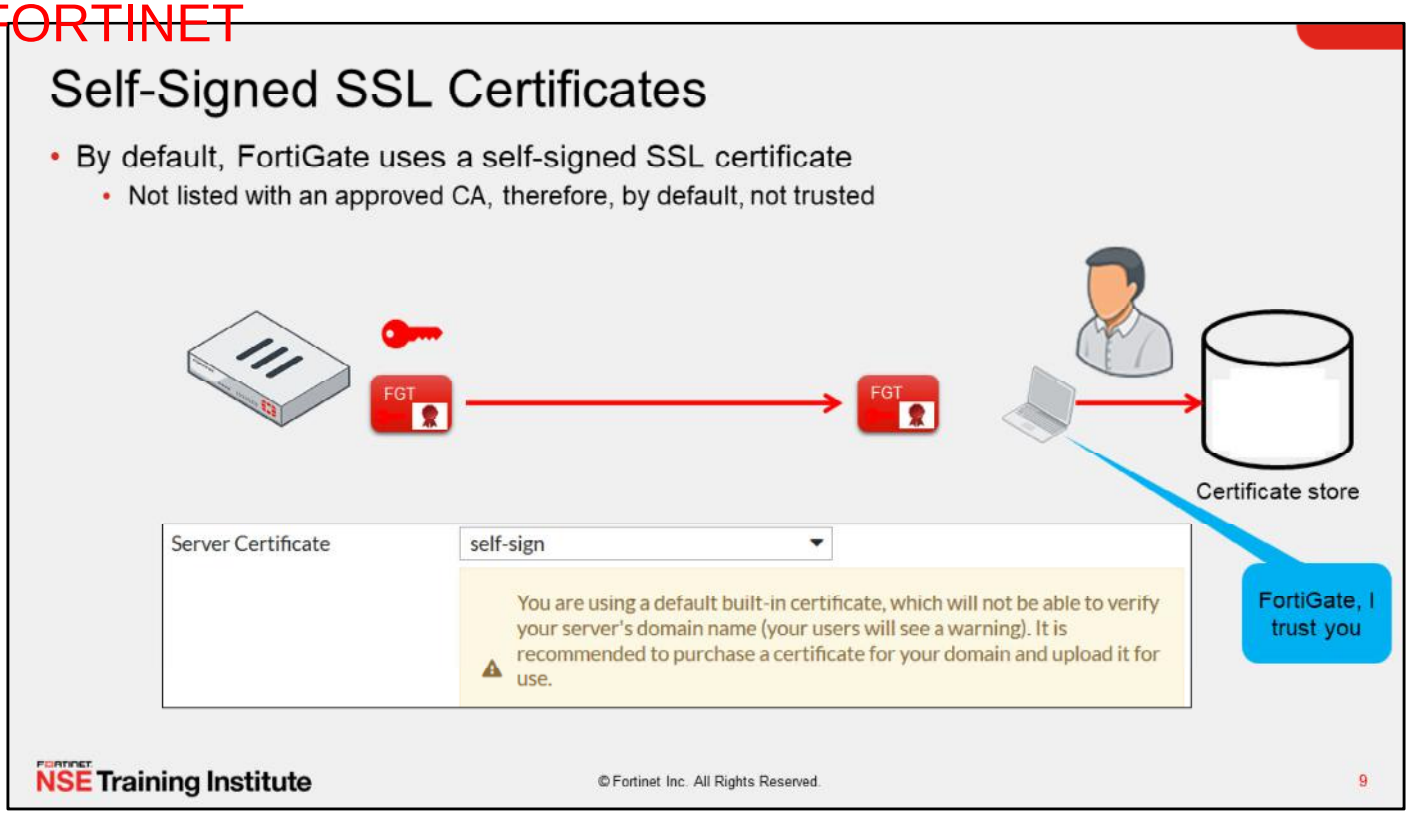

As you can see in the example shown on this slide, trust in the web model is determined by whether or not your certificate store possesses the CA certificate that is required to verify the signature on the SSL certificate. Certificate stores come prepopulated with root and subordinate CA certificates.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

- SSL features:
  - Privacy of data
  - Identifies one or both parties using certificates
  - Uses symmetric and asymmetric (public key) cryptography
- Symmetric cryptography
  - Uses the same key to encrypt and decrypt data
  - When FortiGate establishes an SSL session between itself and another device, the symmetric key (or rather the value to produce it) must be shared 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. For example, if FortiGate connects to a web server to initiate an SSL session, it would use the web server public key to encrypt a string known as the premaster secret. The web server private key would decrypt the premaster secret

NSE Training Institute © Fortinet Inc. All Rights Reserved. 10

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.

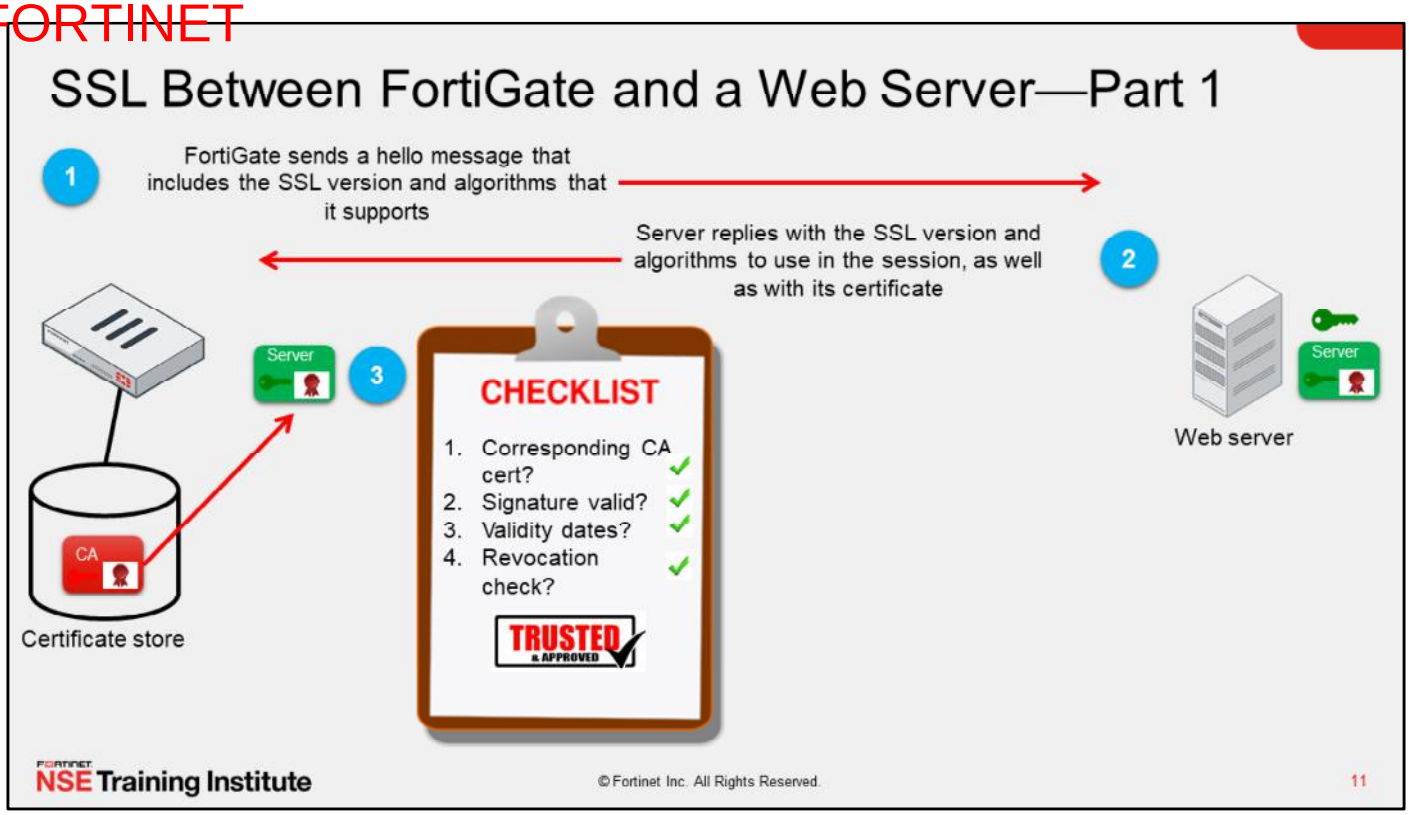

Now, you will learn more about the process of establishing an SSL session.

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.

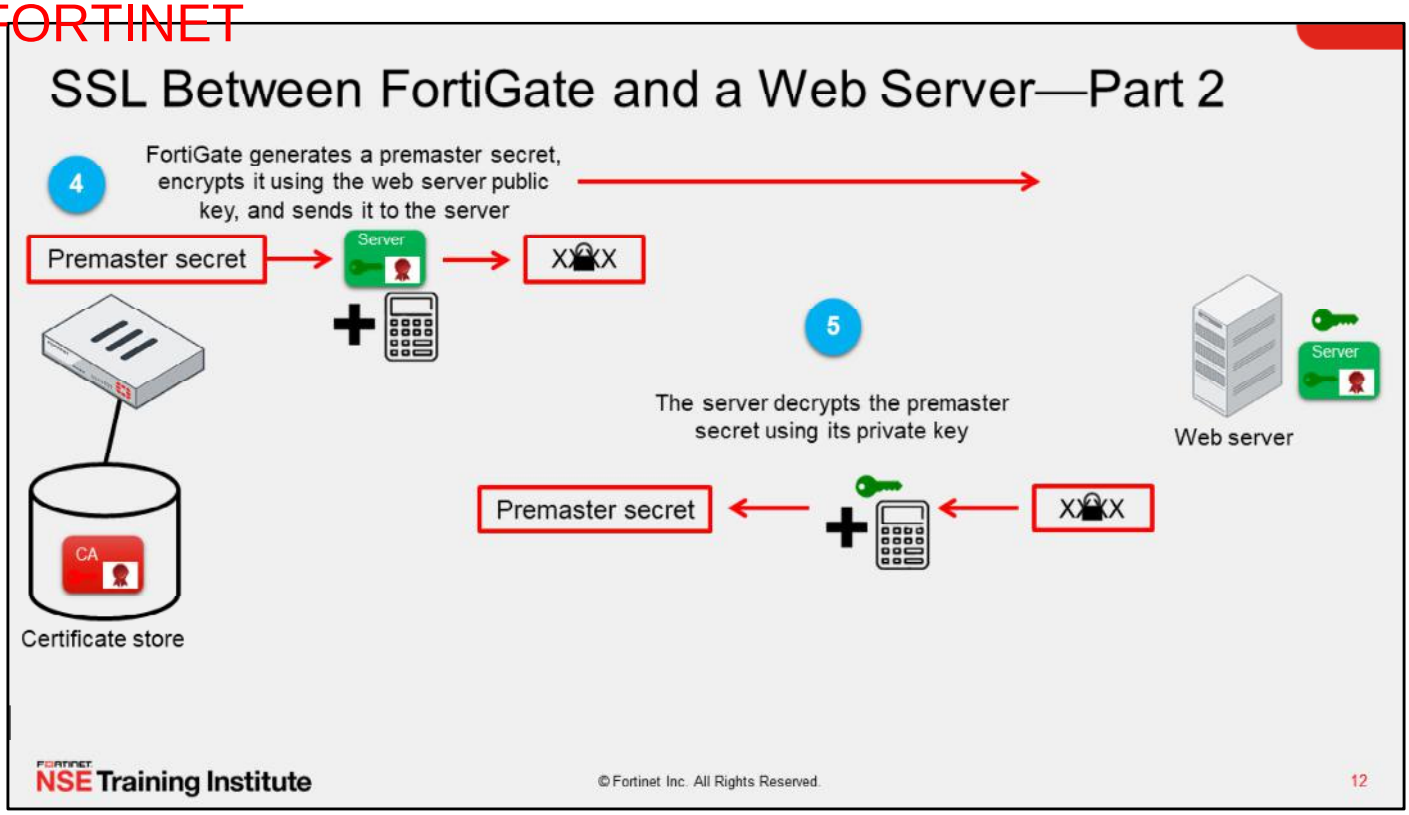

In the fourth step, FortiGate generates a value known as the premaster secret. FortiGate uses the server public key, which is in the certificate, to encrypt the premaster secret. FortiGate then sends the encrypted premaster secret to the web server. If a third-party intercepted the premaster secret, 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.

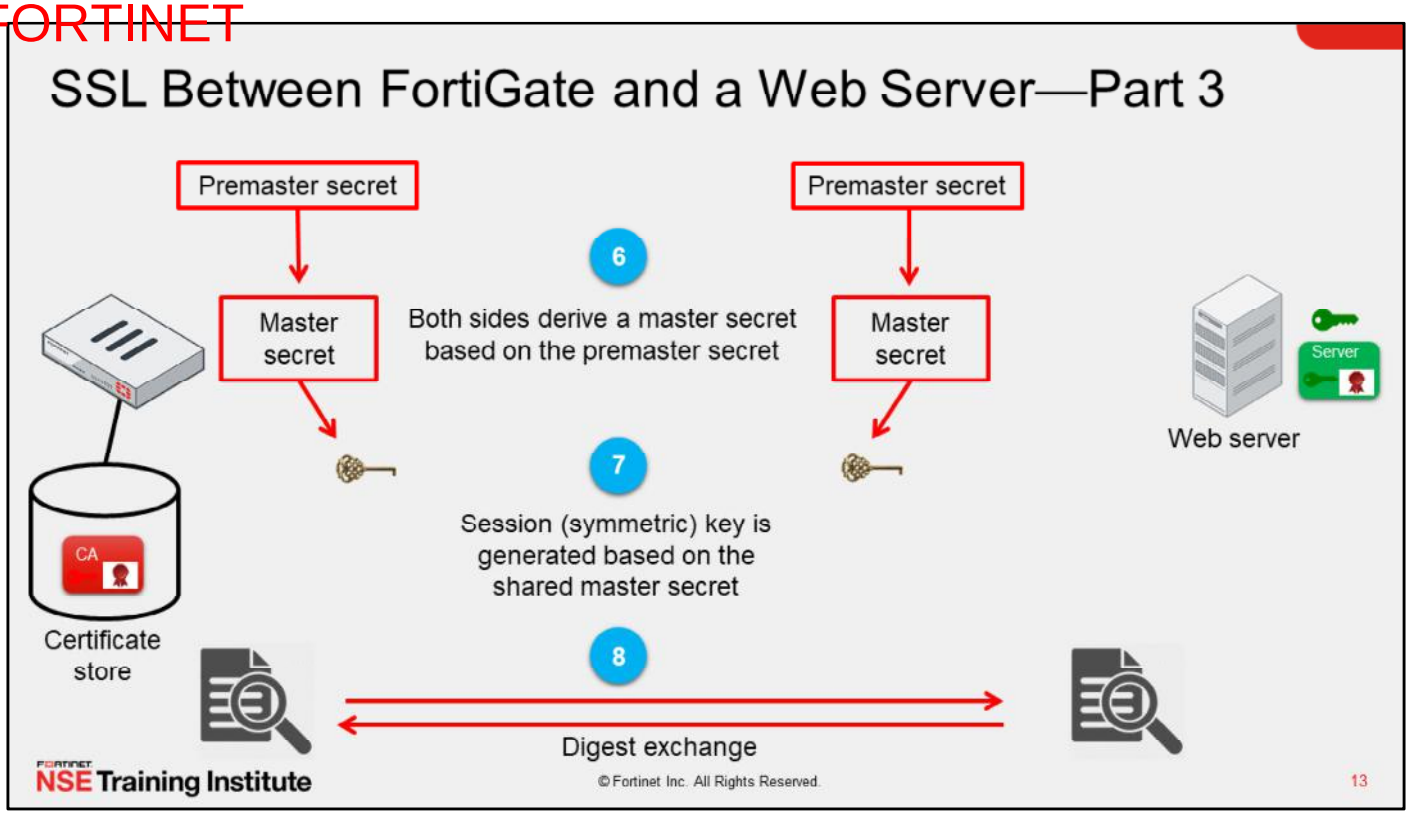

In the sixth step, both FortiGate and the web server derive the master secret based on the premaster secret.

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. 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

NSE Training Institute

14

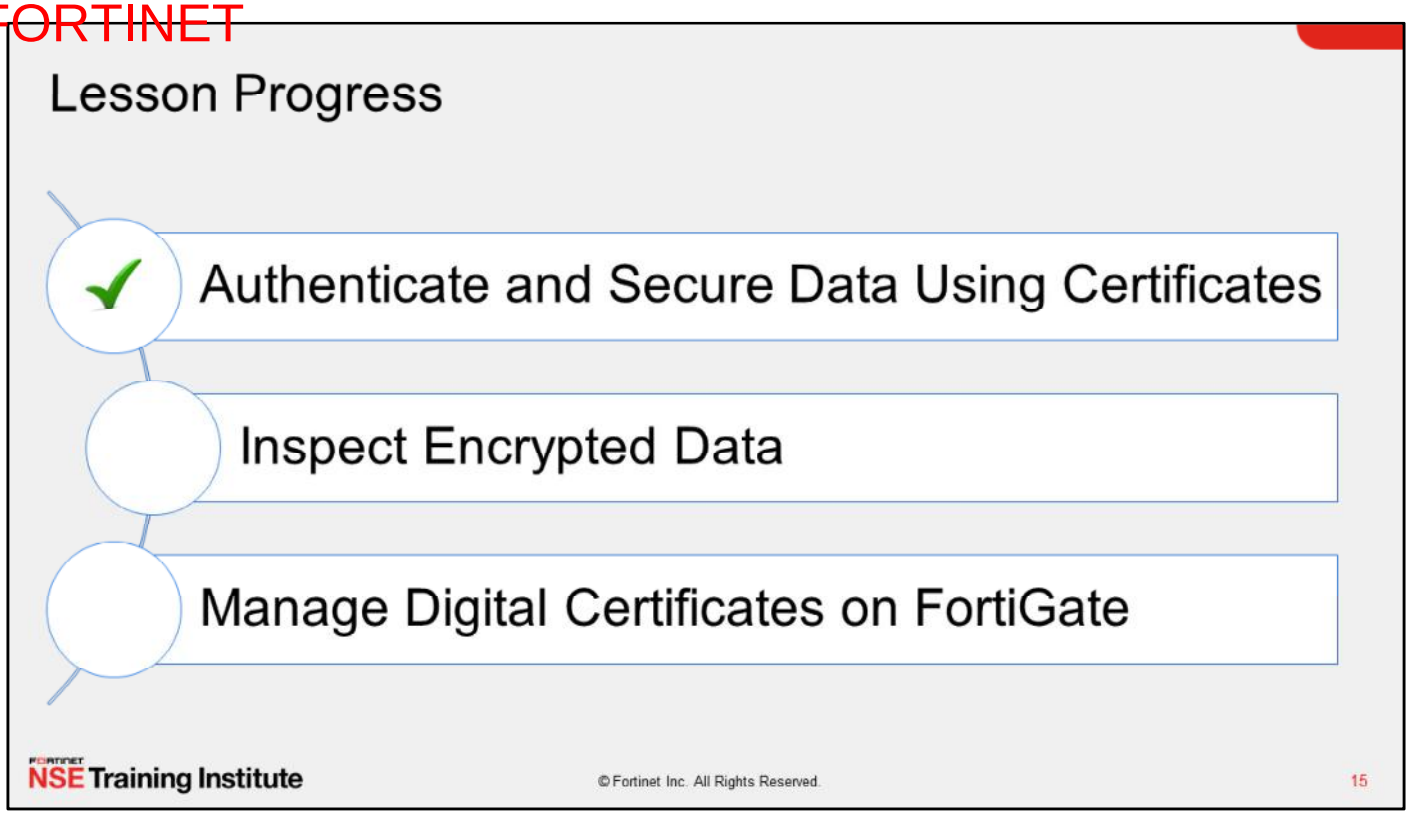

Good job! You now understand why and how FortiGate uses certificates to authenticate devices and people. You also understand how FortiGate uses certificates to ensure the privacy of data as it flows from FortiGate to another device, or from another device to FortiGate.

Now, you will learn about how to inspect encrypted data.

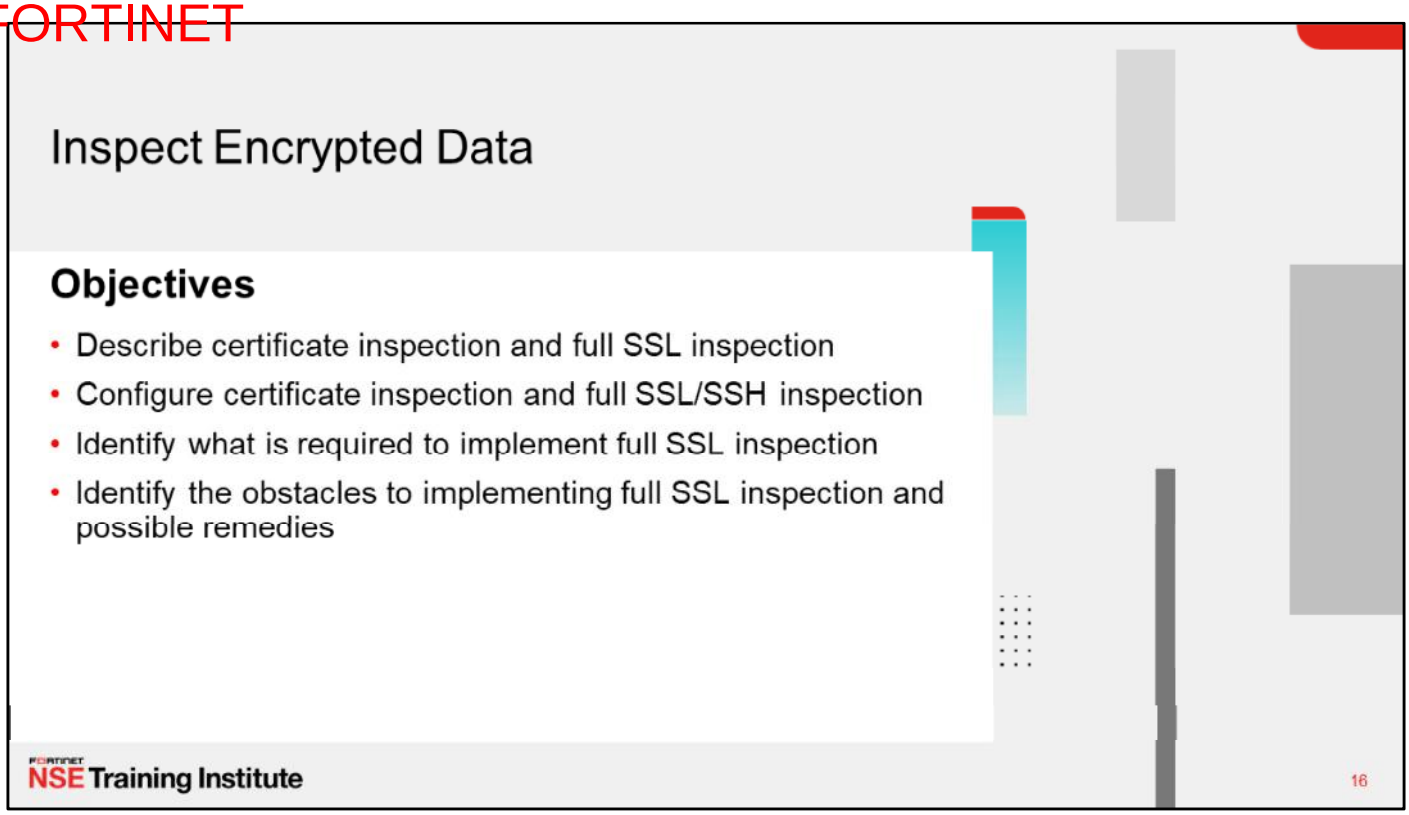

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

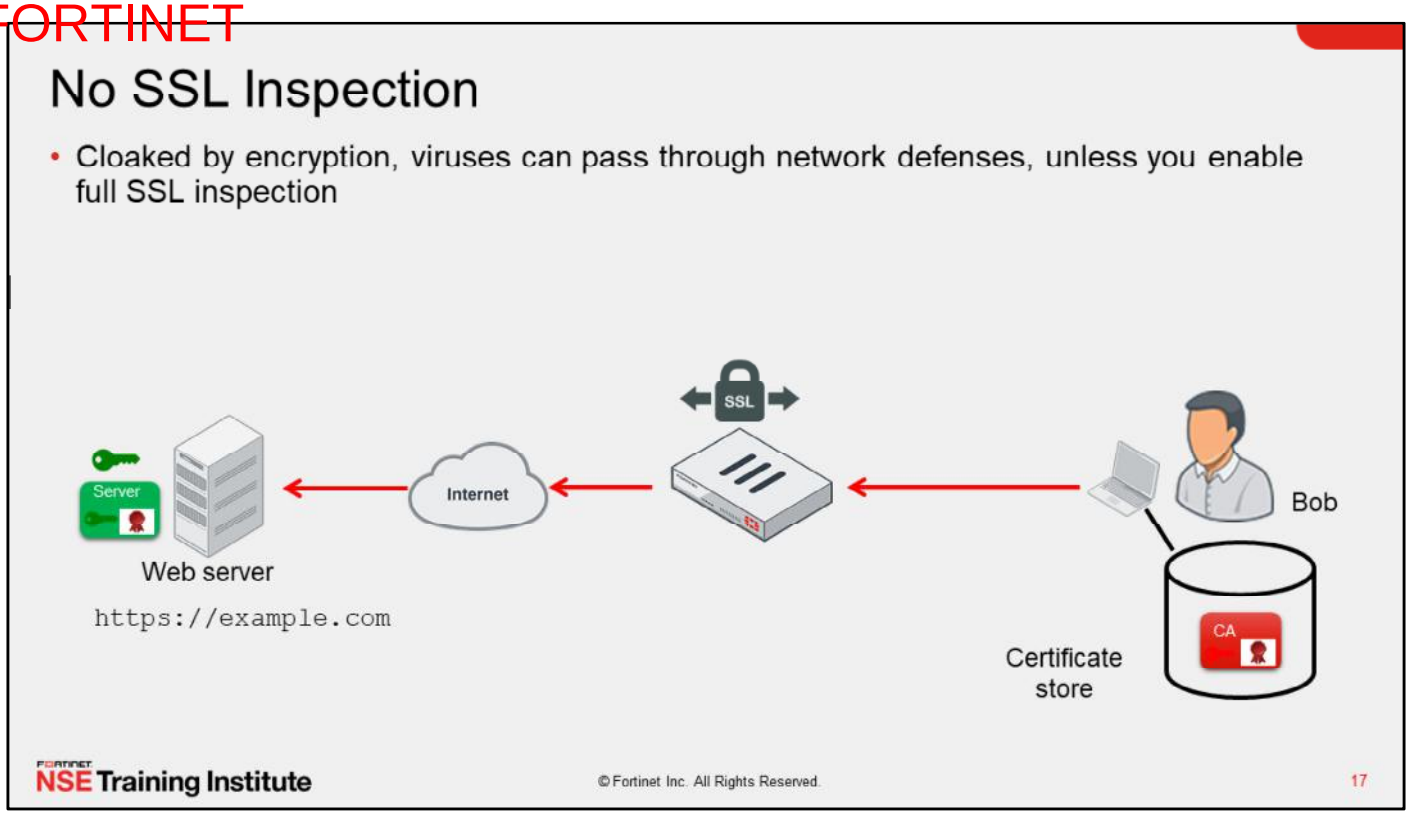

(slide contains animation)

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.

SSL Certificate Inspection

- FortiGate uses the server name indication (SNI) to discern the hostname of the SSL server at the beginning of the SSL handshake
  - If there is no SNI, FortiGate looks at the subject and subject alternative name fields
- The only security feature you can apply using SSL certificate inspection mode is web filtering
- While offering some level of security, certificate inspection does not permit the inspection of encrypted data

NSE Training Institute © Fortinet Inc. All Rights Reserved. 18

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. 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.

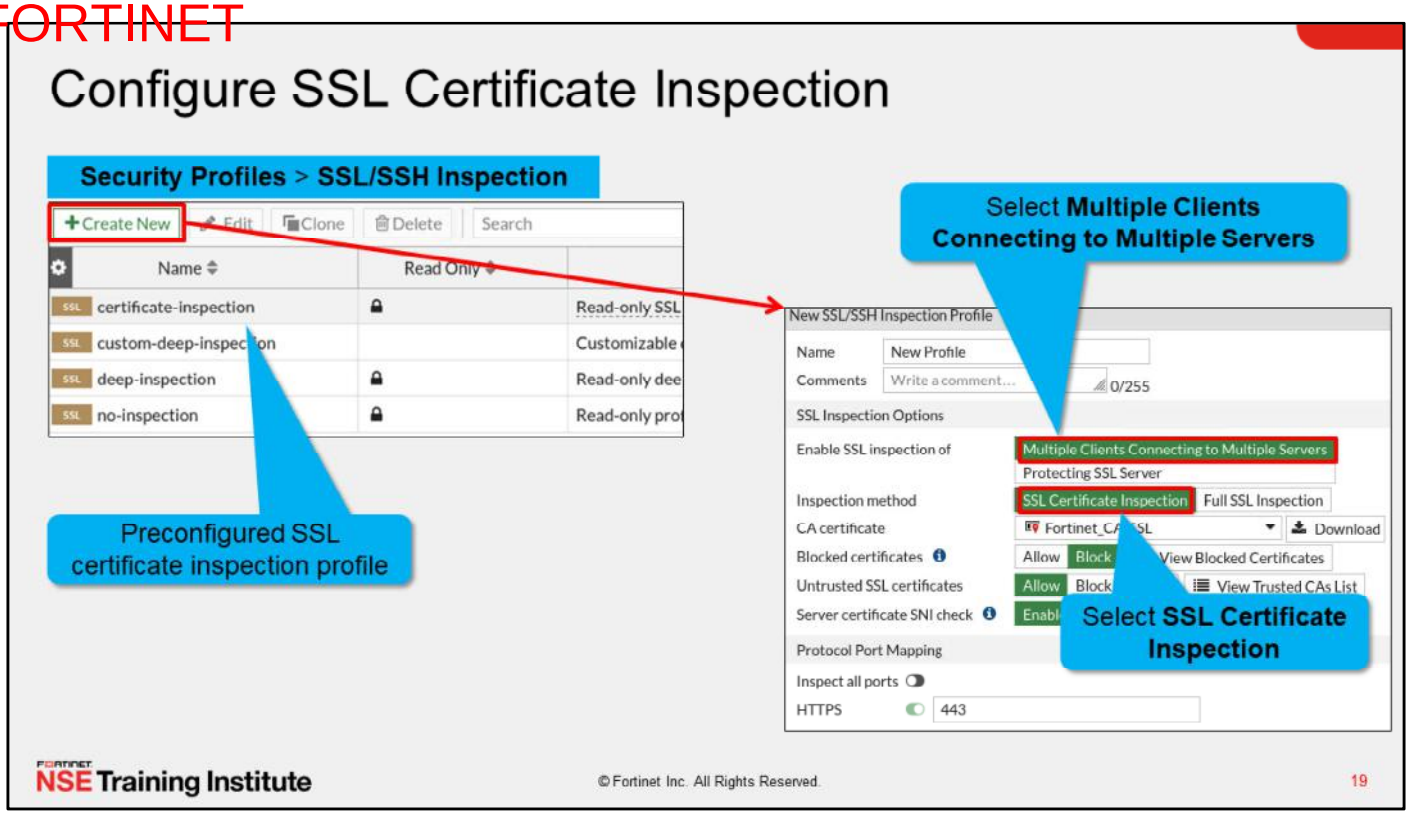

FortiGate has a read-only preconfigured profile for SSL certificate inspection named **certificate-inspection**. If you want to enable SSL certificate inspection, select this profile when configuring a firewall policy.

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**.

Full SSL Inspection—Certificate Requirements

- Full SSL inspection requires that FortiGate act as a CA to generate an SSL private key and certificate as a proxy web server
  - To be compliant with the Internet Engineering Task Force (IETF) RFC 5280, the CA certificate requires these two extensions to issue certificates:
    - cA=True
    - keyUsage=keyCertSign
- FortiGate devices that support full SSL inspection can get their CA certificate from a couple of sources:
  - A self-signed Fortinet_CA_SSL certificate from within FortiGate
  - A certificate issued by an internal CA (FortiGate then acts as a subordinate CA)
- The root CA certificate must be imported into the client machines

NSE Training Institute © Fortinet Inc. All Rights Reserved. 20

FortiGate 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.

Full SSL Inspection on Outbound Traffic—Part 1

- FortiGate requires the private key to decrypt and inspect SSL traffic
  - FortiGate intercepts traffic coming from the server and generates and signs a new certificate with the same subject name

NSE Training Institute

21

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.

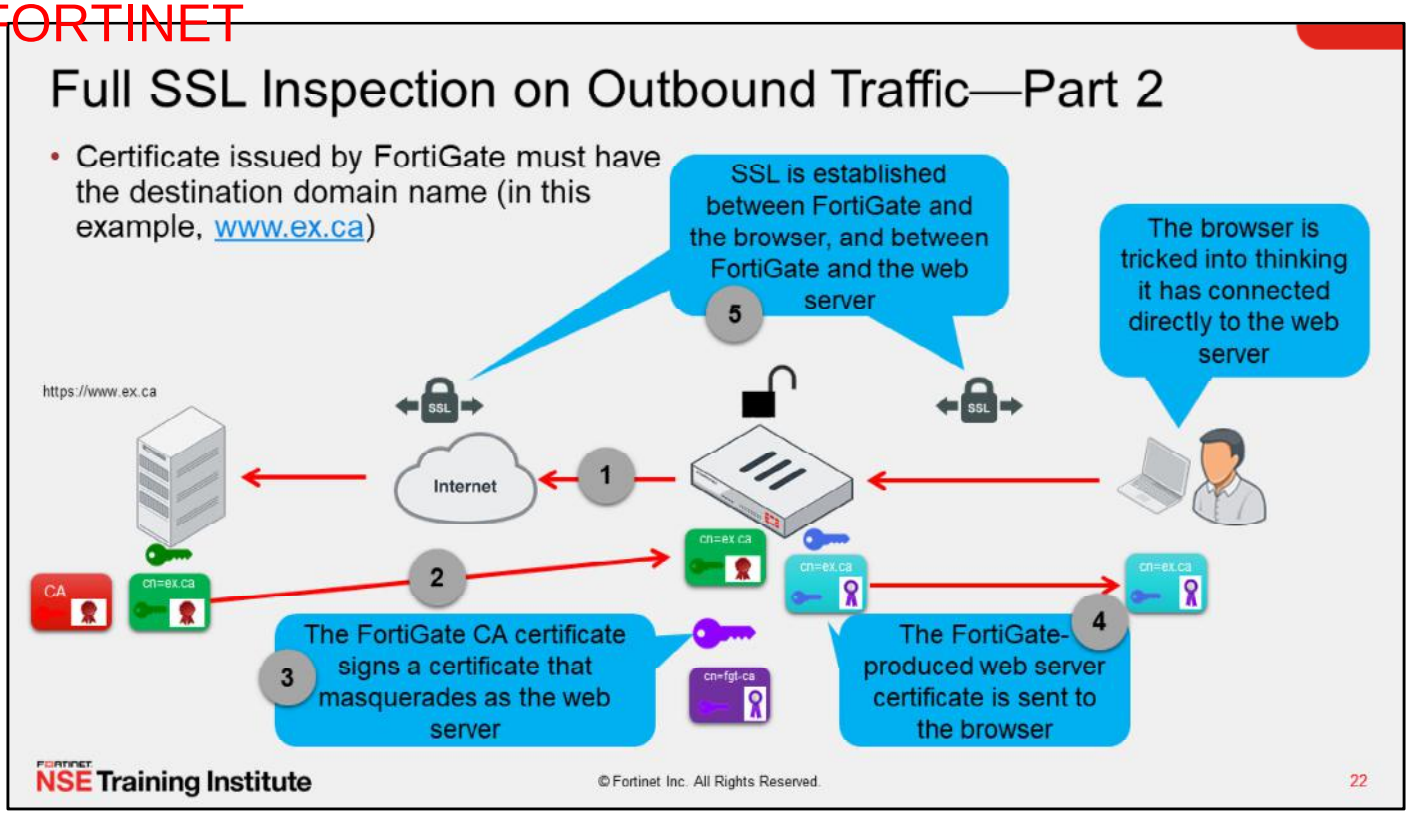

In step 1, an internal web browser connects to an SSL-enabled web server. Normally, when a browser connects to a secure site, the web server sends its certificate to the browser. However, in step 2, FortiGate intercepts the web server certificate. In step 3, the FortiGate internal CA generates a new 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.

Untrusted SSL Certificates Setting

- Allow, block, or ignore untrusted certificates (only available if **Multiple Clients Connecting to Multiple Servers** is selected)
  - **Allow**: sends the browser an untrusted temporary certificate when the server certificate is untrusted
  - **Block**: blocks the connection when an untrusted server certificate is detected
  - **Ignore**: uses a trusted FortiGate certificate to replace the server certificate always, even when the server certificate is untrusted

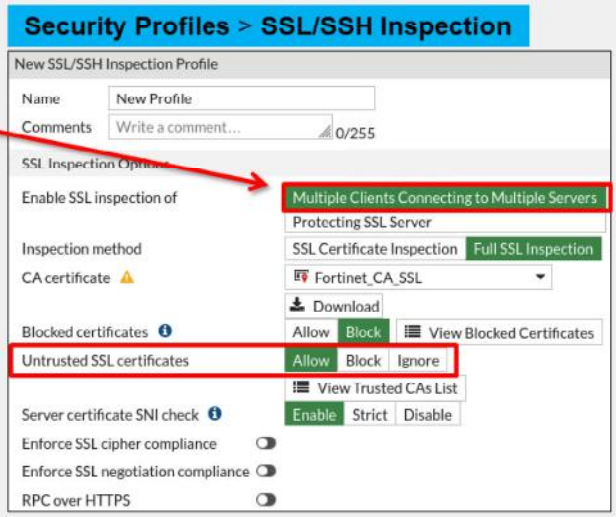

NSE Training Institute

23

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.

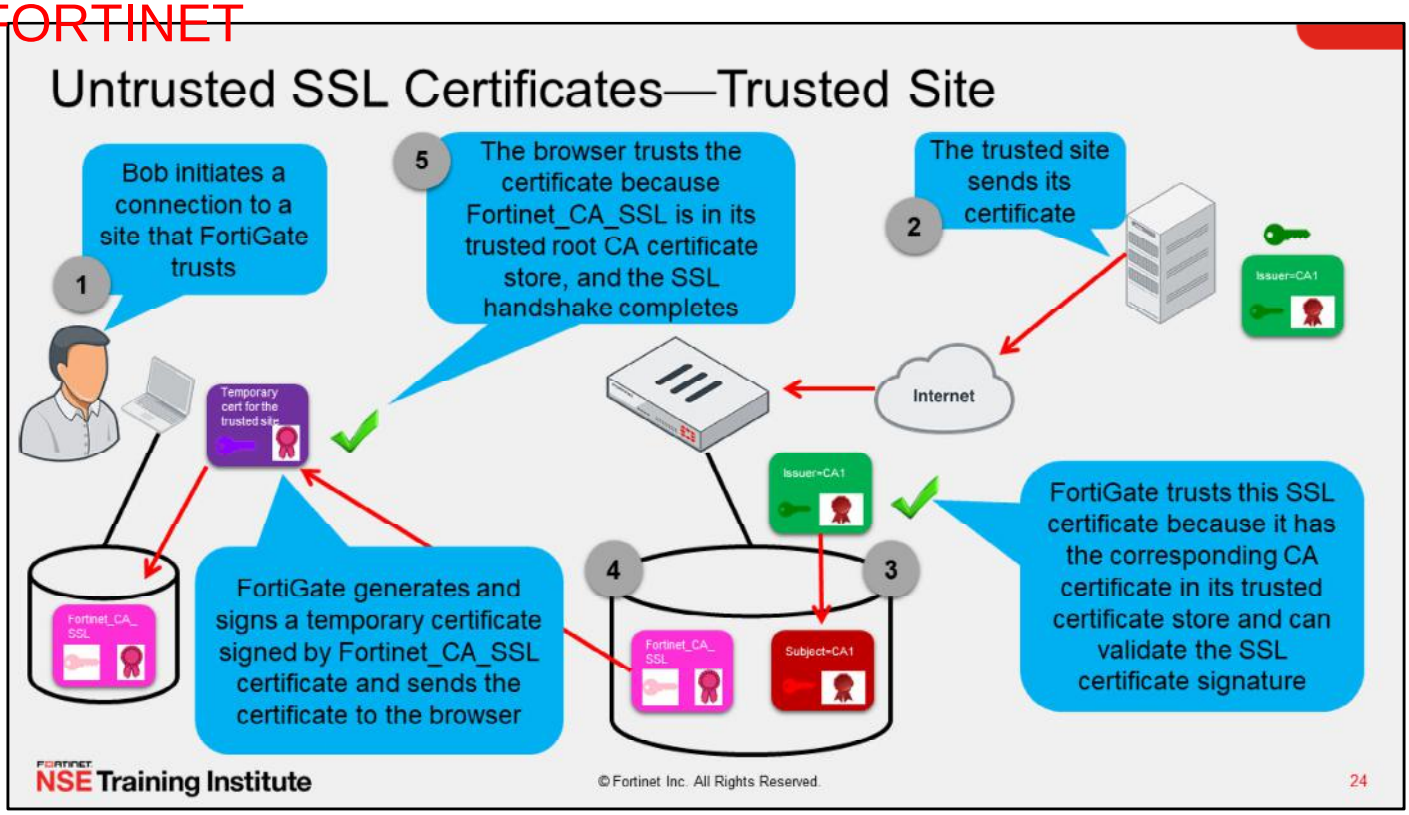

The scenario shown on this slide describes how FortiGate handles a trusted external site regardless of the **Untrusted SSL Certificate** setting.

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.

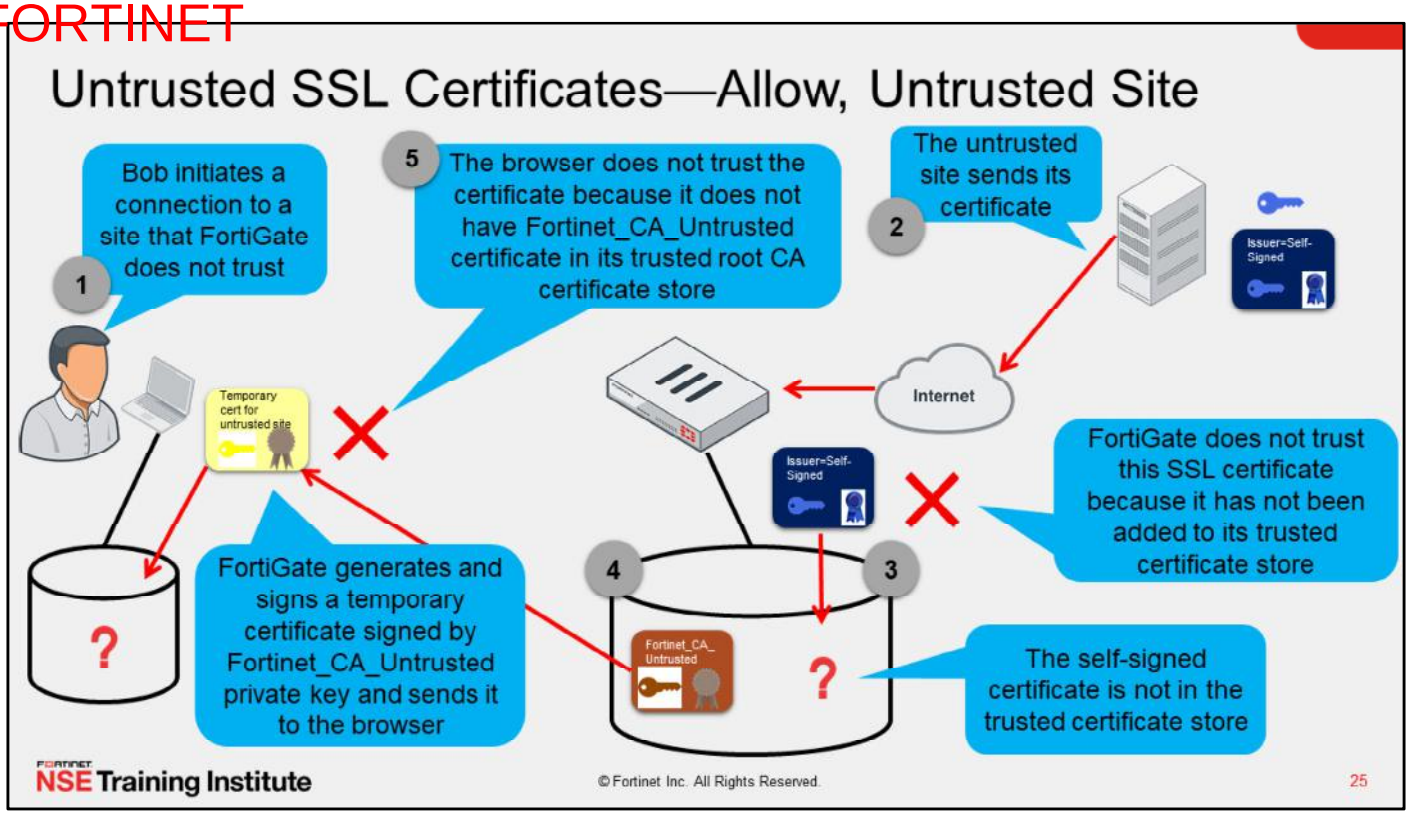

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Allow**.

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.

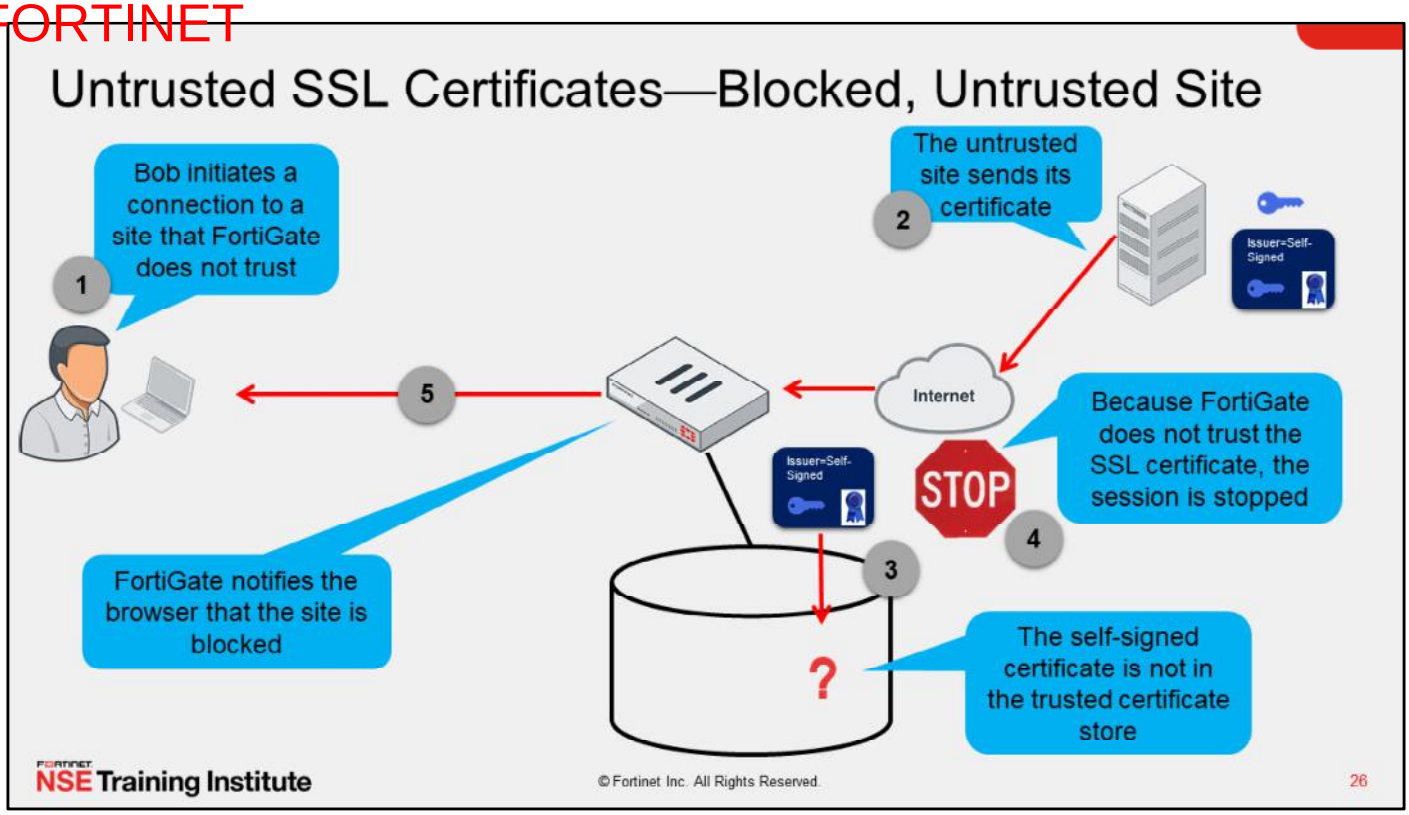

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Block**.

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.

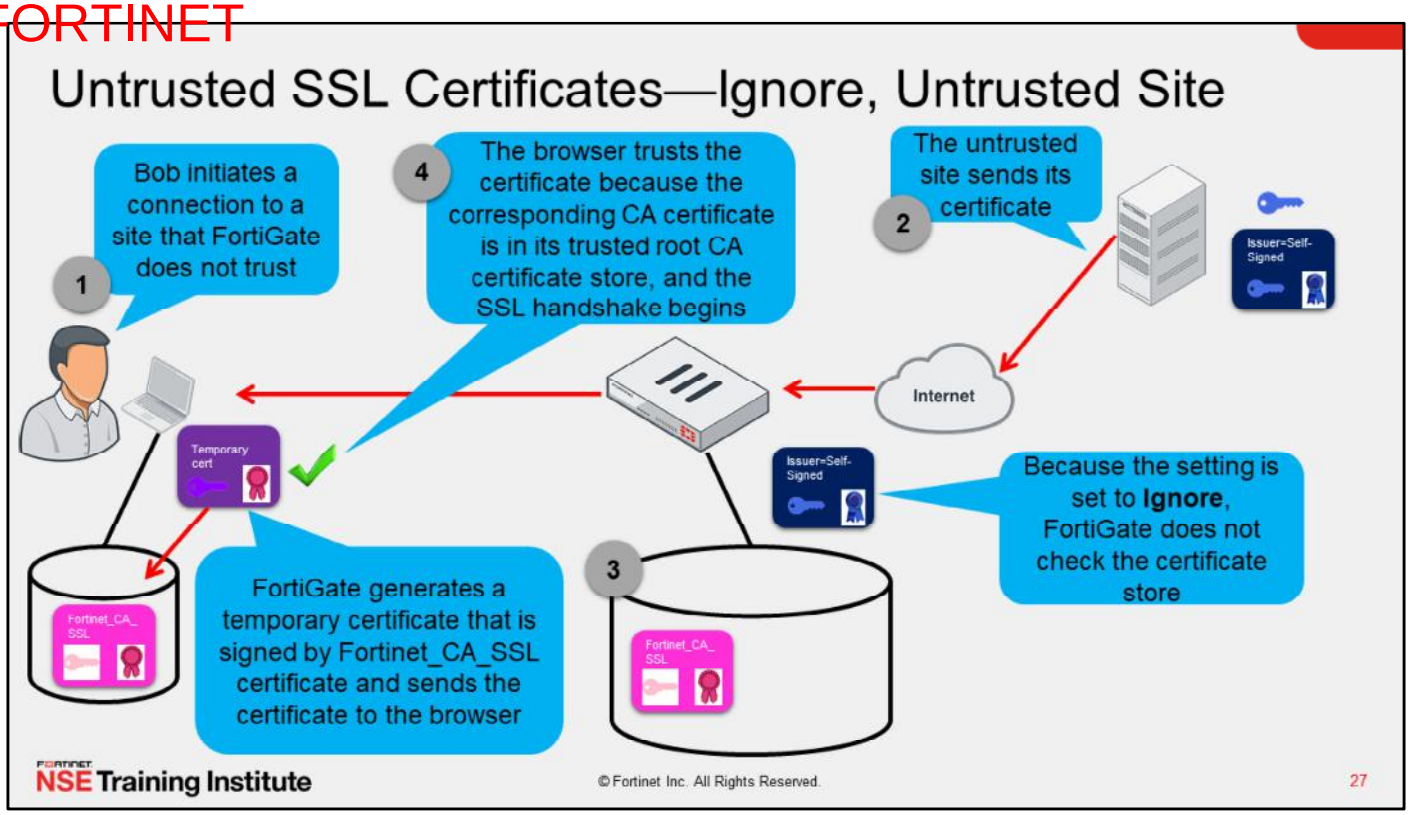

The scenario shown on this slide describes how FortiGate handles an untrusted external site when **Untrusted SSL Certificate** is set to **Ignore**.

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.

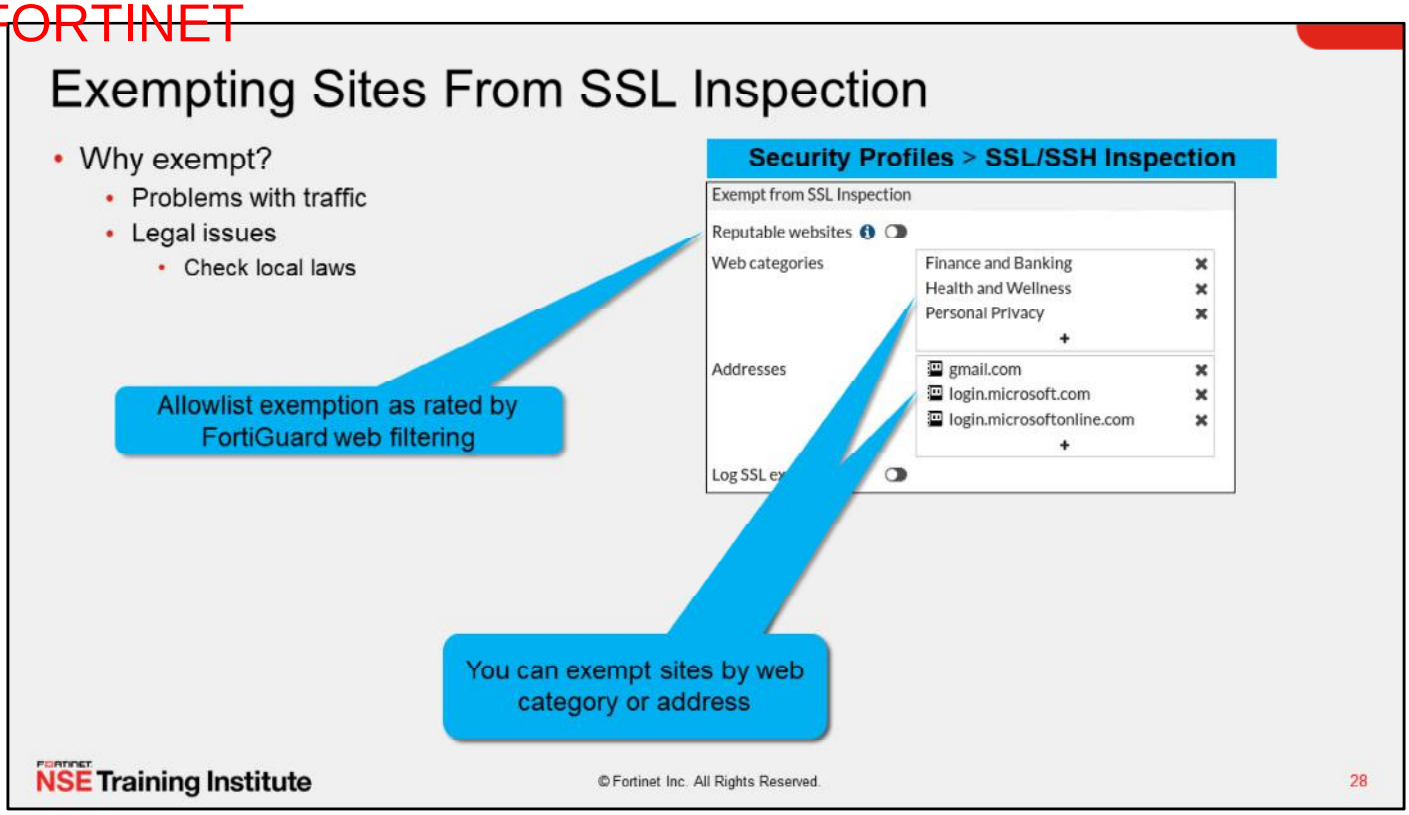

Within the full SSL inspection profile, you can also specify which SSL sites, if any, you want to exempt from SSL inspection. You may need to exempt traffic from SSL inspection if it is causing problems with traffic, or for legal reasons.

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.

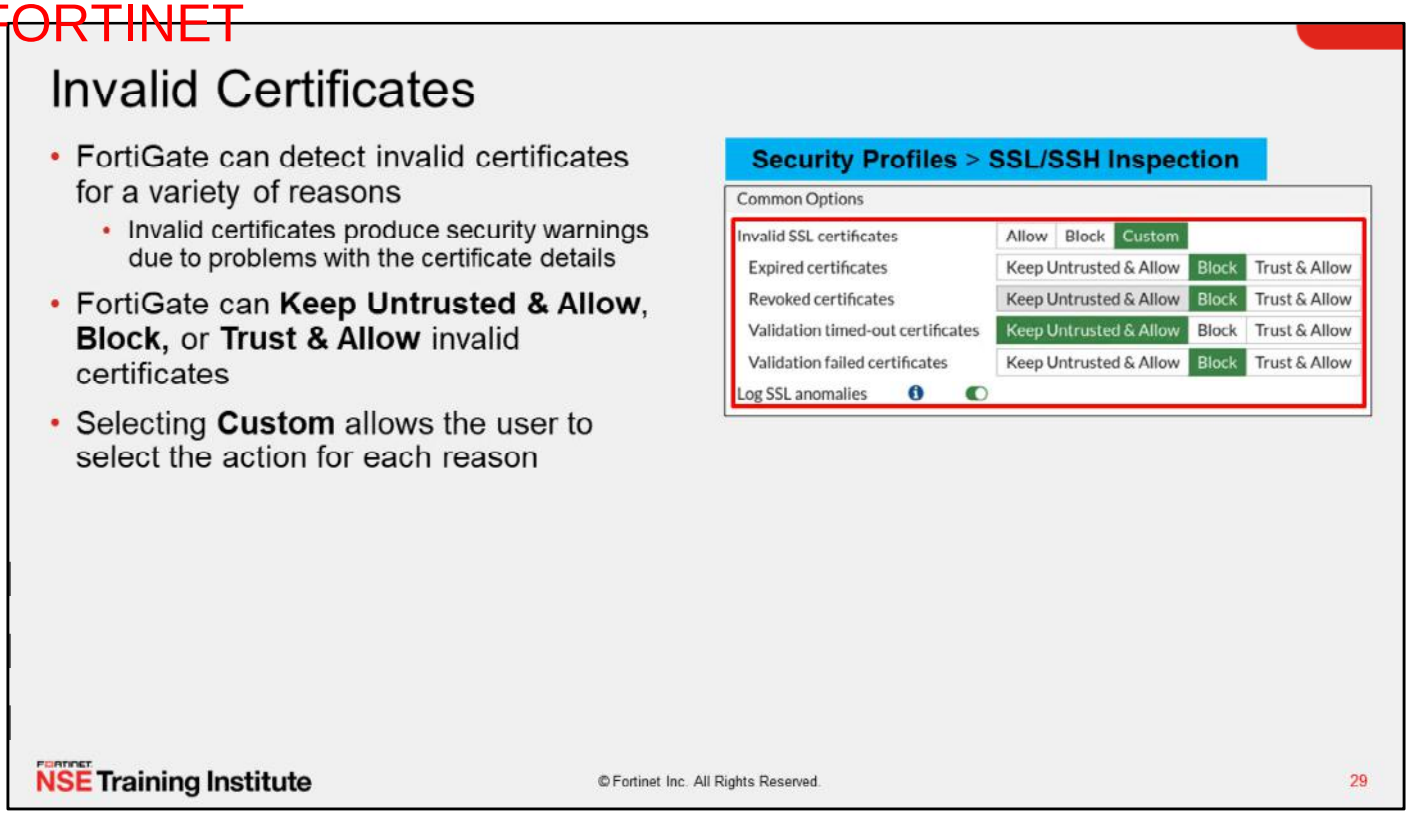

FortiGate can detect certificates that are invalid for the following reasons:

- 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.

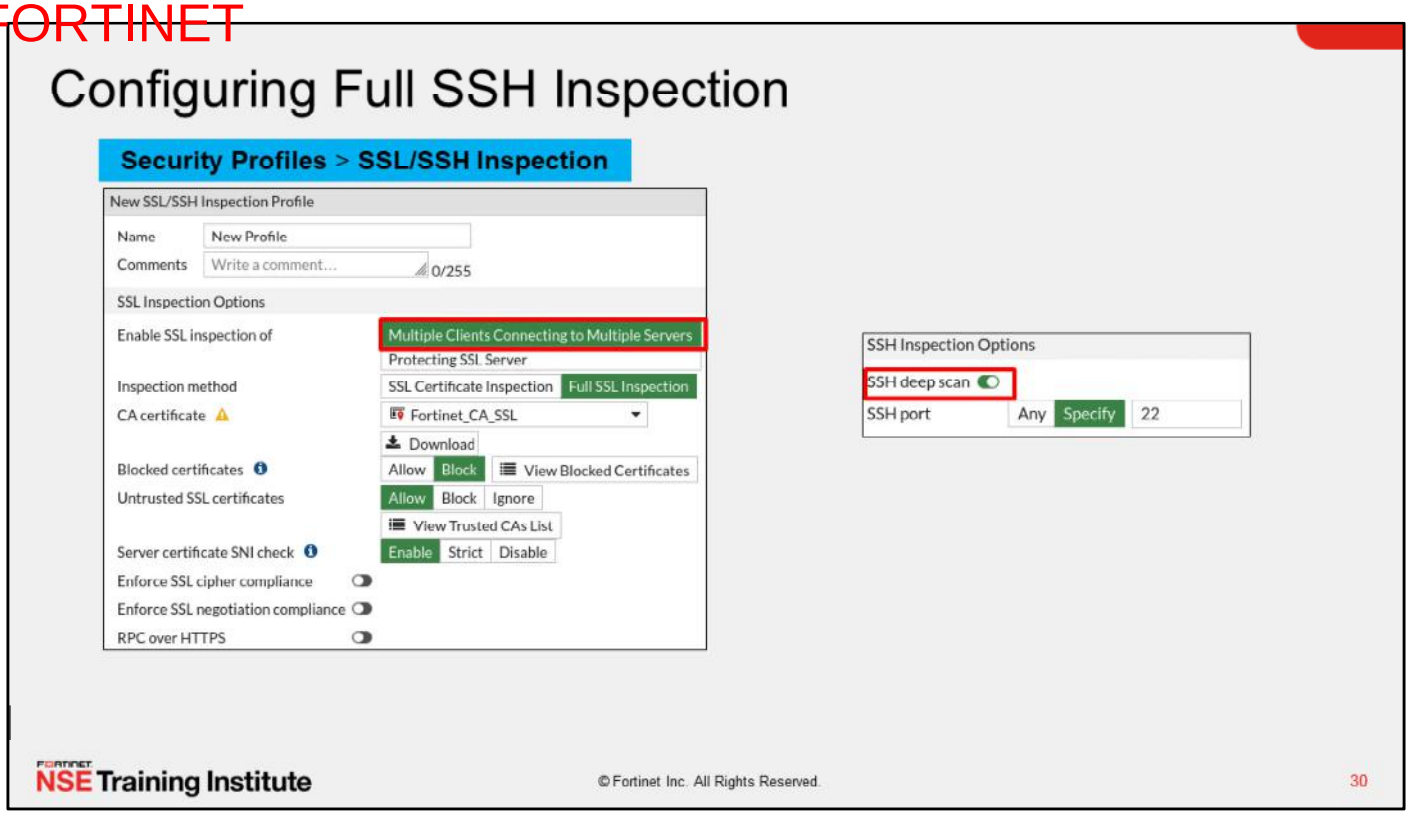

You can enable **SSH deep scan** when you select **Multiple Clients Connecting to Multiple Servers**. When you enable **SSH deep scan**, FortiGate does an MITM attack for SSH traffic. A process similar to the one done for full SSL inspection takes place. FortiGate intercepts the SSH key sent by the server, generates a new one, and sends it to the client. If the SSH client had stored the original host key, then it detects a change in the host key and warns the user. The user can then replace the original host key with the new host key generated by FortiGate.

By default, **SSH deep scan** listens on TCP port 22. You can specify a different port number, or select **Any** in the **SSH port** field. When you do this, FortiGate scans all connections to identify SSH traffic using different ports. Specifying a port for SSH traffic is not as comprehensive as searching all ports, but it is easier on the performance of the firewall.

Finally, note that **SSH deep scan** is a proxy-based inspection feature only. In addition, the only security features that use **SSH deep scan** are antivirus and data leak prevention.

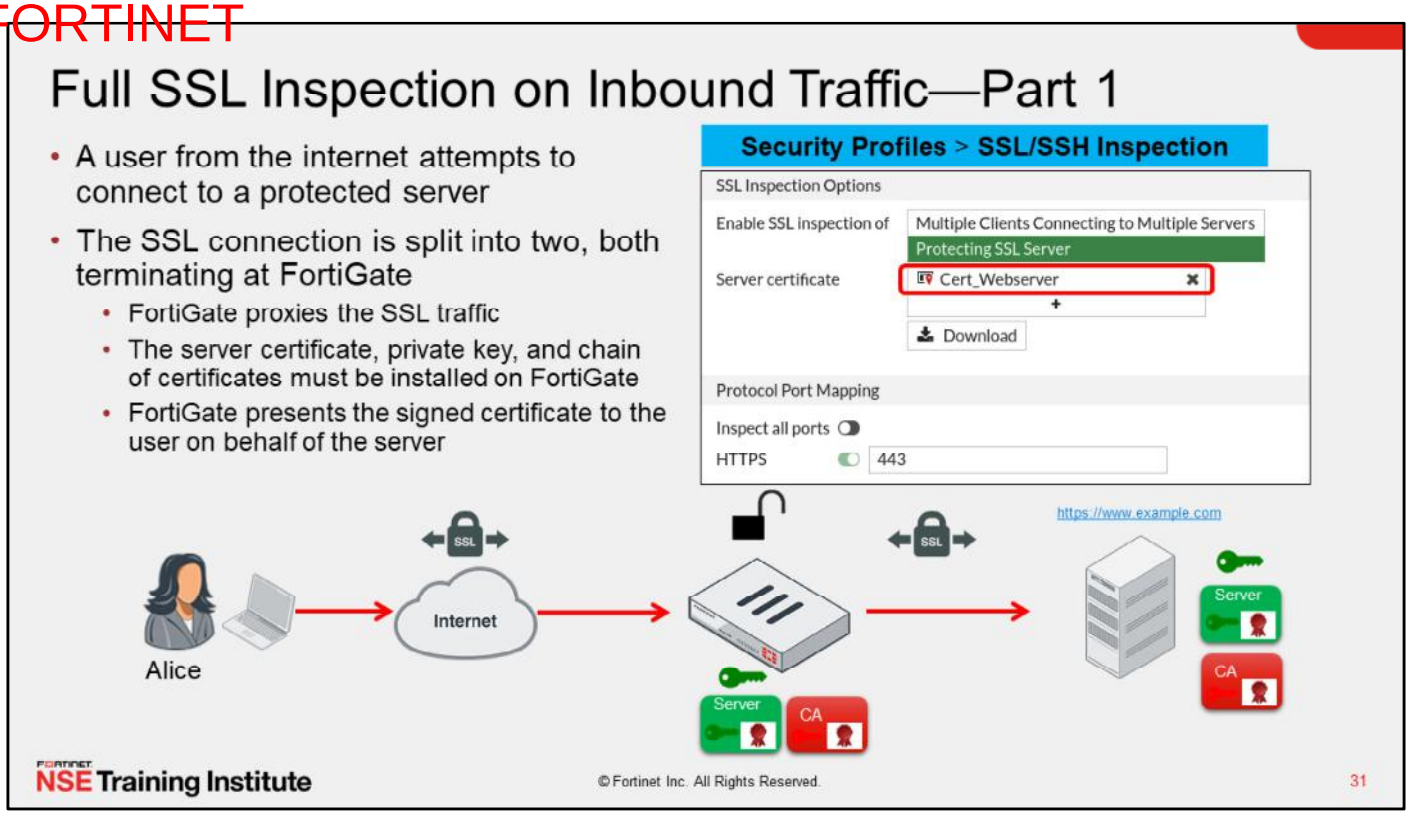

In the example shown on this slide, FortiGate is protecting a web server. This is the second configuration option for full SSL inspection. When configuring the SSL inspection profile for this server, you must select **Protecting SSL Server**, import the server key pair to FortiGate, and then select the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

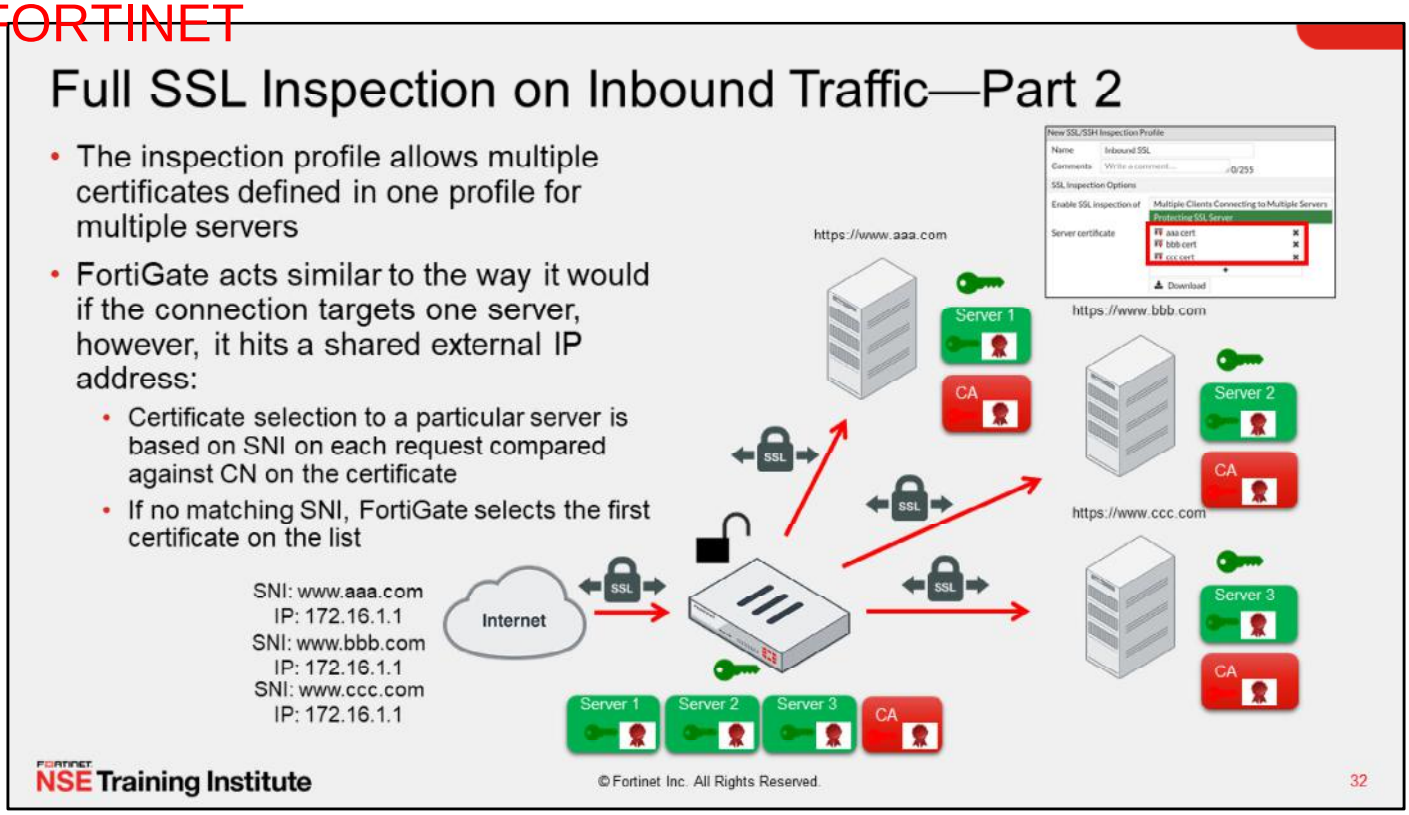

By creating a full SSL inspection profile on inbound traffic, you can configure the profile to use multiple web sites if they are approachable by the same external IP address. When FortiGate receives client and server hello messages, it selects the certificate to perform the full SSL inspection based 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.

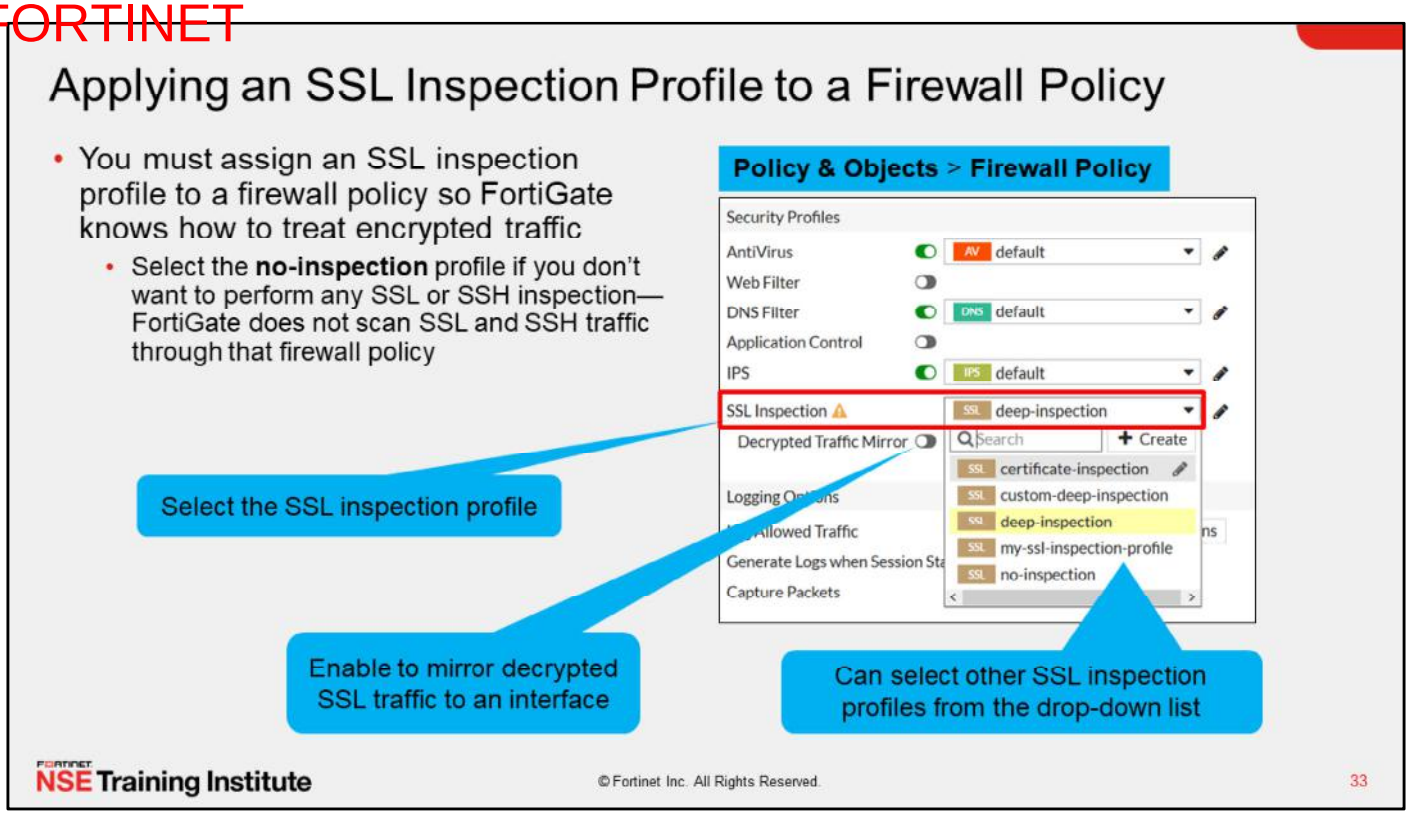

After you create and configure an SSL inspection profile, you must assign it to a firewall policy so FortiGate knows how to inspect encrypted traffic. Most of the internet traffic is being encrypted nowadays. For this reason, you usually want to enable SSL inspection to protect your network from security 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. 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.

Certificate Warnings

- The browser may display a certificate warning during SSL inspection because it does not trust the CA
- To avoid certificate warnings, do one of the following:
  - Use the Fortinet_CA_SSL certificate and install the FortiGate CA root certificate in all the browsers
  - Use an SSL certificate issued by a CA and ensure that the root CA certificate is installed on all the browsers

NSE Training Institute © Fortinet Inc. All Rights Reserved. 34

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.

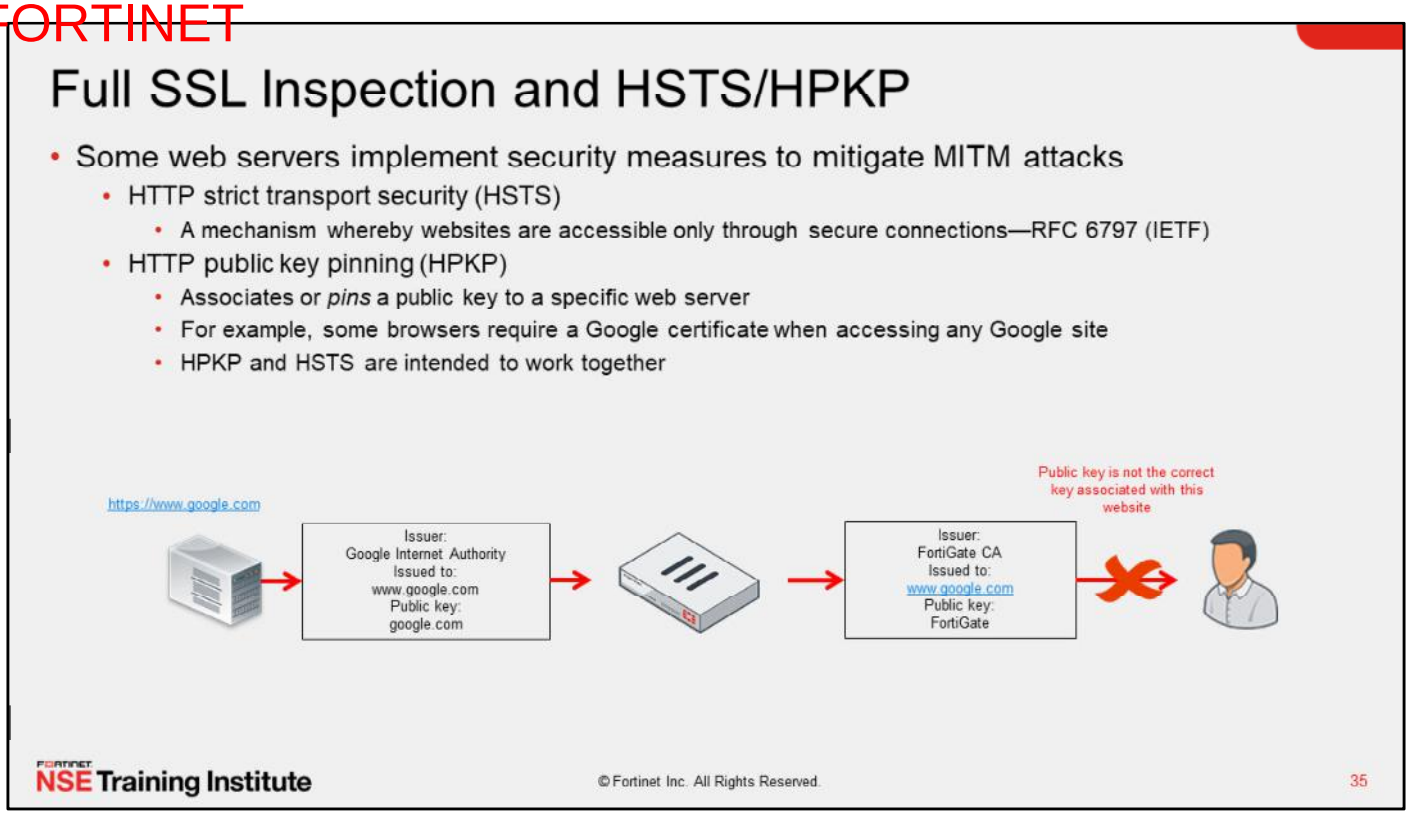

Some security measures have been introduced by the IETF to mitigate MITM attacks. Some of these measures cause problems when you implement outbound full SSL inspection.

HTTP strict transport security (HSTS) and HTTP public key pinning (HPKP) are security features designed to thwart MITM attacks. HSTS is “a mechanism enabling websites to declare themselves accessible only through secure connections …”, according to RFC 6797 of the IETF. In other words, a user from a web browser would be forced to use HTTPS when connecting to a website with this policy; there would be no option to connect using HTTP. HPKP is a security feature imposed by the web server that associates one or more public keys with the website for a specified period of time. The public key doesn’t have to be the web server public key, it could be one of the intermediate or root CA public keys, as long as it exists in the certificate chain. When the web browser visits an HPKP-enabled website, hashes of the public keys associated with a website are cached on the client machine.

Going forward, each time the web browser connects to the web server, it compares one or more of the keys presented with the cached key fingerprints. If the browser cannot match at least one of the keys, the SSL handshake terminates. This is a problem for outbound full SSL inspection. FortiGate generates a new certificate and public key to establish an SSL session with the web browser. FortiGate cannot provide an authentic certificate chain, so the connection would be rejected by the browser. Predictably, this could prevent users from connecting to many legitimate sites.

Resolving HPKP Issues

- Exempt those websites from full SSL inspection
- Use SSL certificate inspection instead
- Use a web browser that does not support HPKP, like Chrome, Internet Explorer, or Edge
- Disable the security setting in the browser (not always an option)

NSE Training Institute

36

The options available to circumvent HPKP are limited. One option is to exempt SSL inspection for those sites. Another option is to use SSL certificate inspection instead. A third option is to use a browser that does not support HPKP, like Chrome, Internet Explorer, or Edge. Last, in some browsers it is possible to disable HPKP.

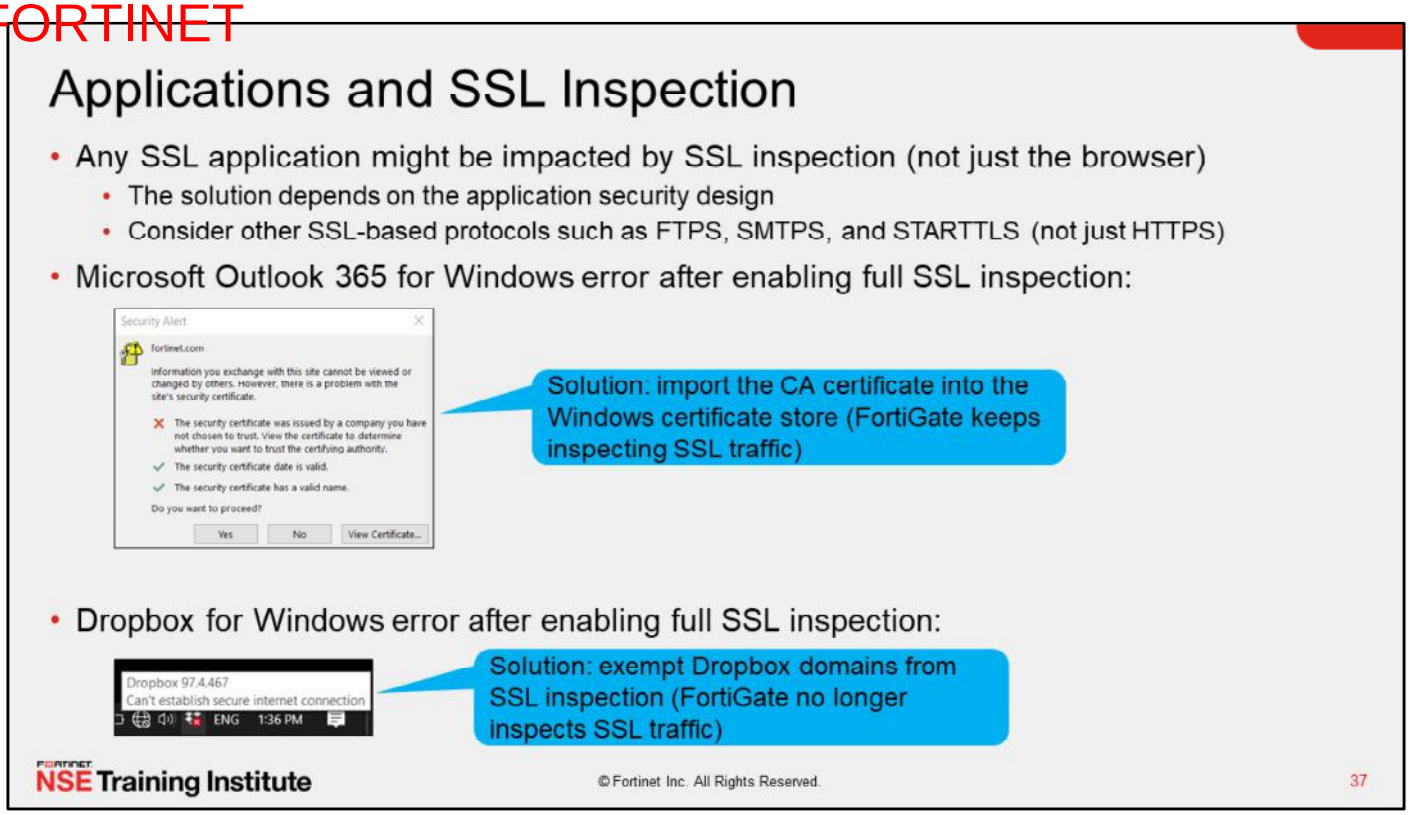

More and more applications are using SSL to securely exchange data over the internet. While most of the content in this lesson centers around the operation and impact of SSL inspection on browsers, the same applies to other applications using SSL as well. After all, the browser is just another application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

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

NSE Training Institute

38

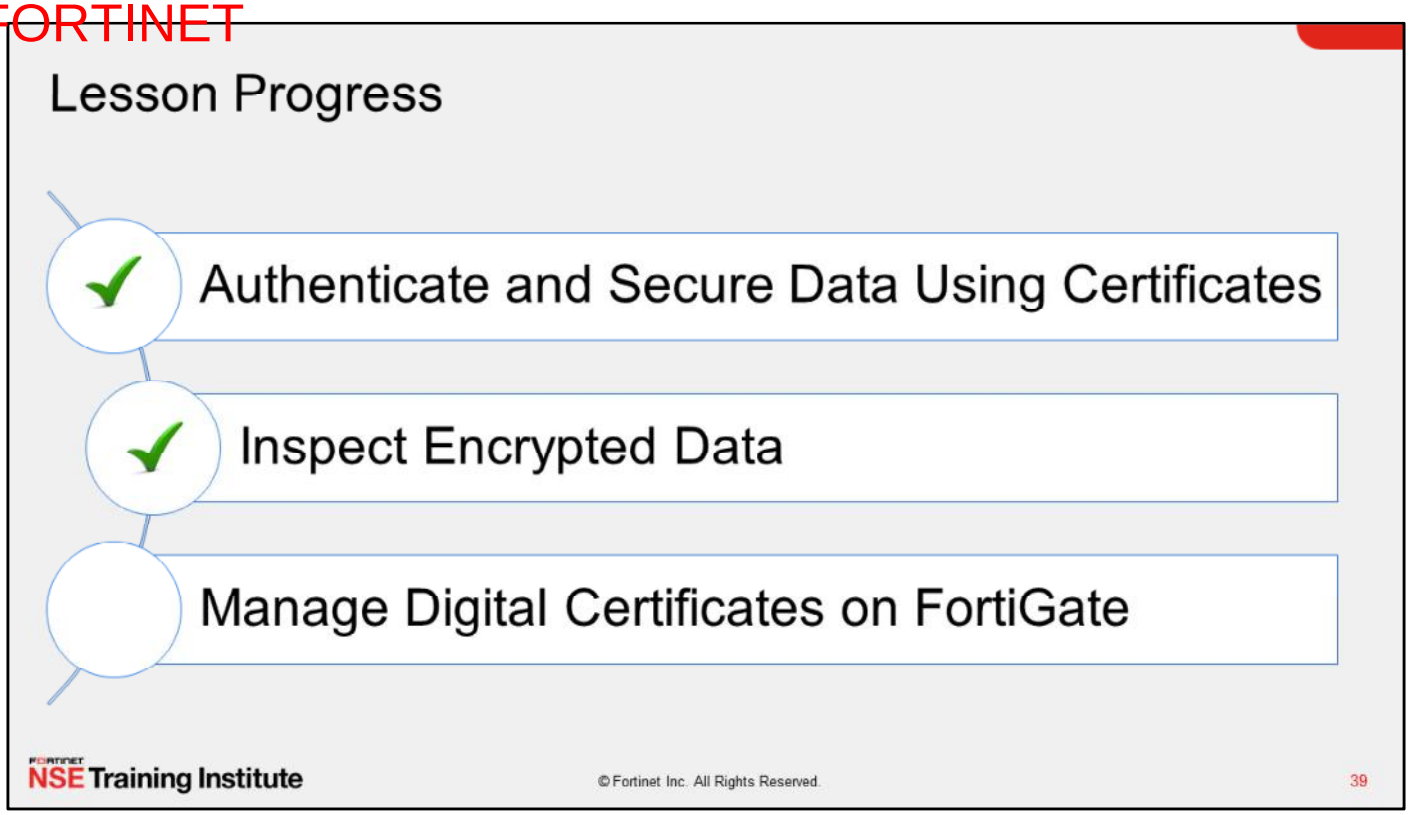

Good job! You now can describe certificate and deep inspection, and you can configure FortiGate to use either one of these options.

Now, you will learn how to manage digital certificates on FortiGate.

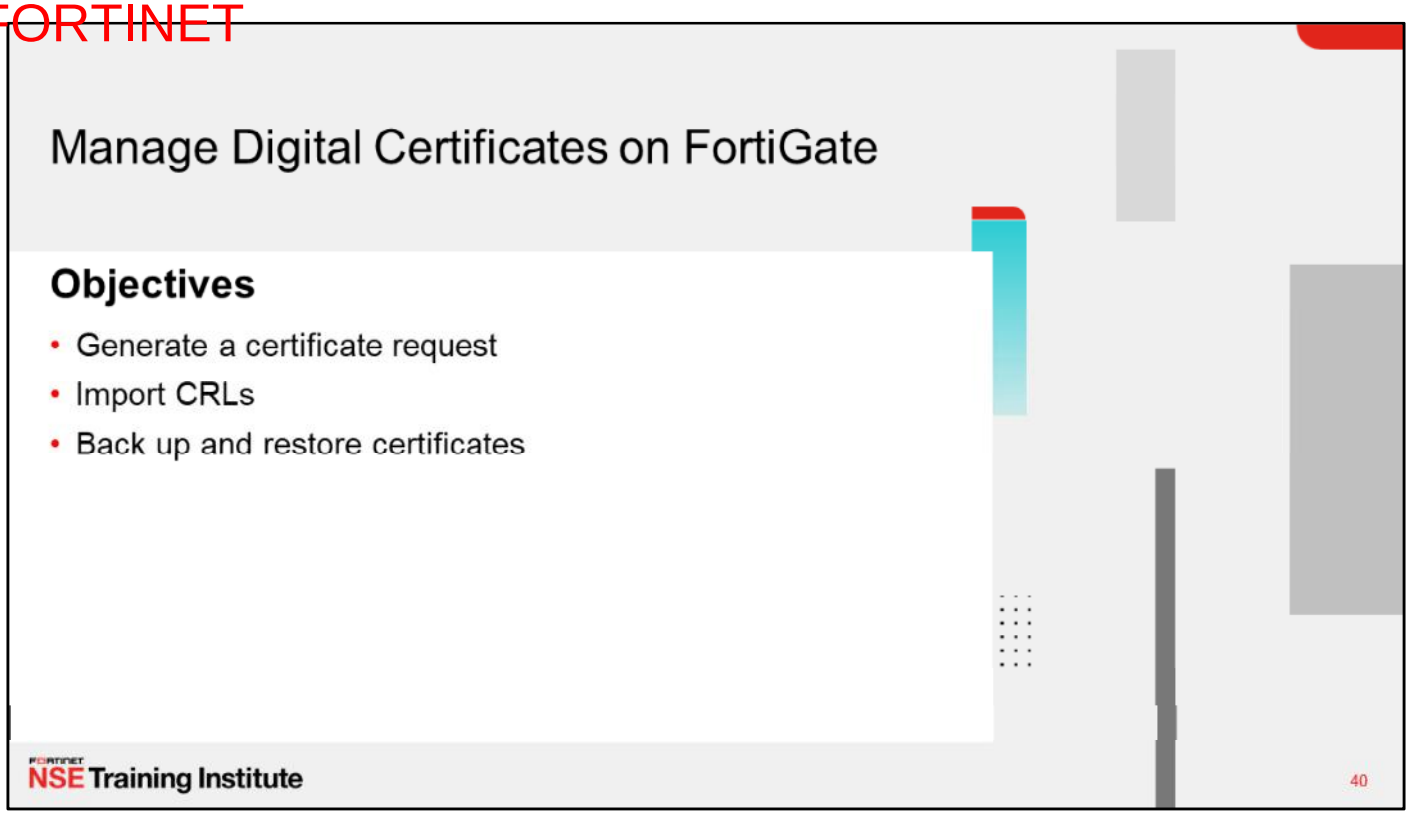

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in generating certificate requests, importing CRLs, and backing up and restoring certificates, you will be able to manage certificates on FortiGate.

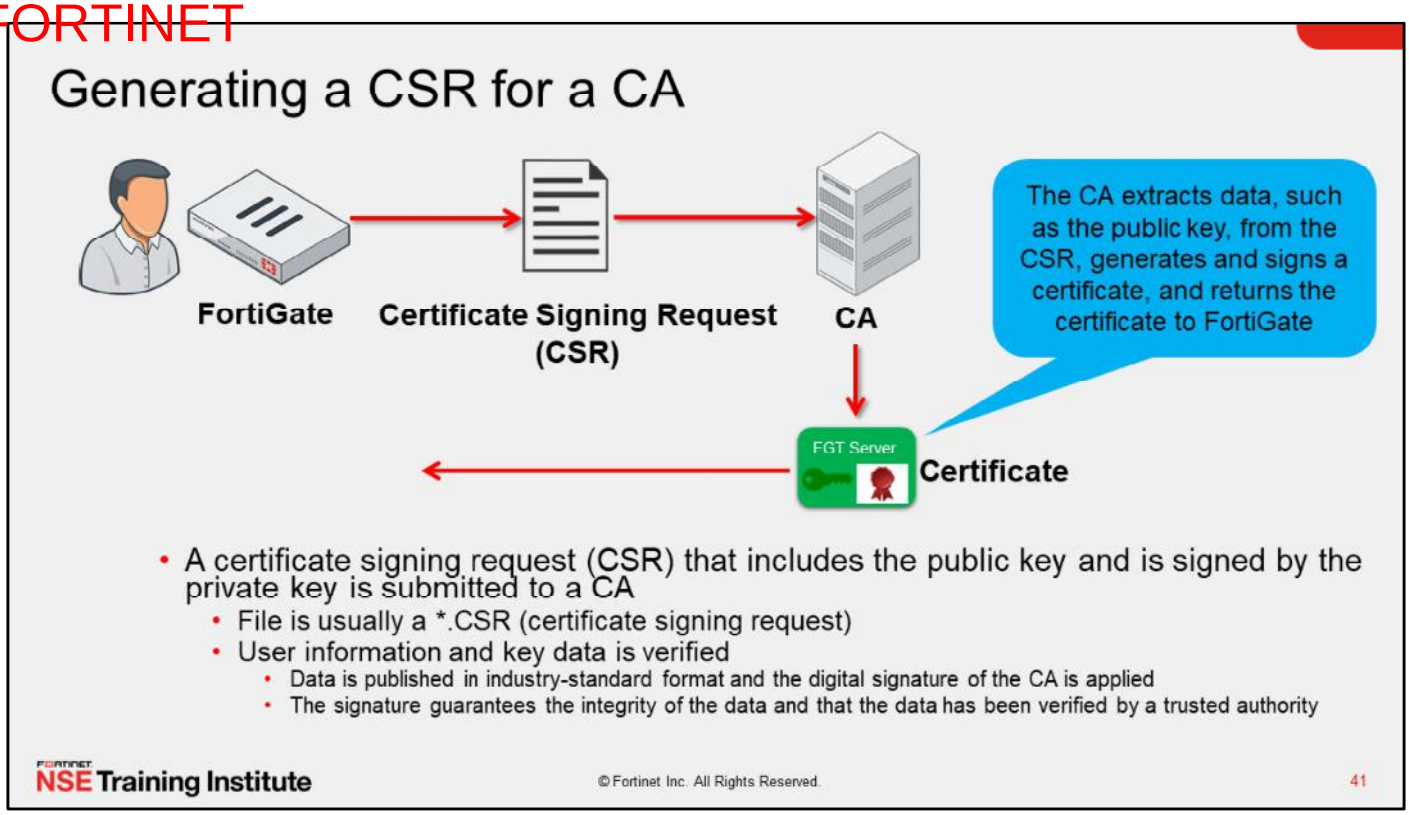

The process of obtaining a digital certificate for FortiGate begins with creating a certificate signing request (CSR). The process is as follows:

1. FortiGate generates a CSR. A private and public key pair is created for FortiGate. The CSR is signed by the FortiGate private key.
2. FortiGate submits the CSR to a CA. The CSR includes the FortiGate public key and specific information about FortiGate (IP address, distinguished name, email address, and so on). Note that the private key remains confidential on FortiGate.
3. The CA verifies that the information in the CSR is valid, and then creates a digital certificate for FortiGate. The certificate is digitally signed using the CA private key. The CA also publishes the certificate to a central repository. The certificate binds the public key to FortiGate.
4. The certificate is returned to install on FortiGate.

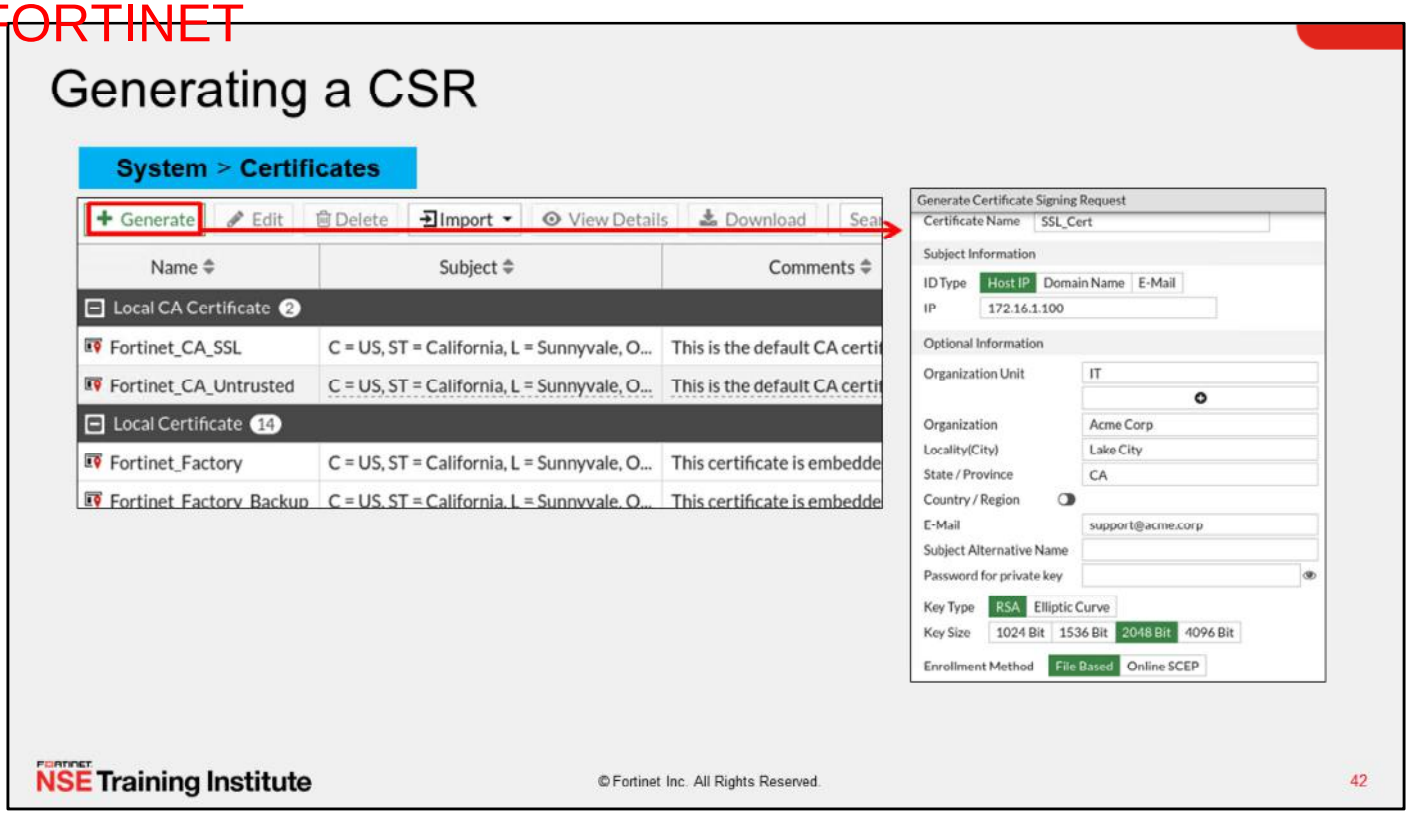

You can generate a CSR on the **Certificates** page of the GUI by clicking **Generate**. Enter all of the required information, such as the IP address (or FQDN) and company name. Ensure the key type and size fit your requirements. You can submit the CSR to a CA using either of the following methods:

- Select **File Based** to generate the CSR as a .csr file, which is then sent to the CA.
- Select **Online SCEP** to submit the CSR to the CA online using the Simple Certificate Enrollment Protocol (SCEP). For example, if you are using FortiAuthenticator as your CA, you can enable and configure SCEP on FortiAuthenticator and use this method.

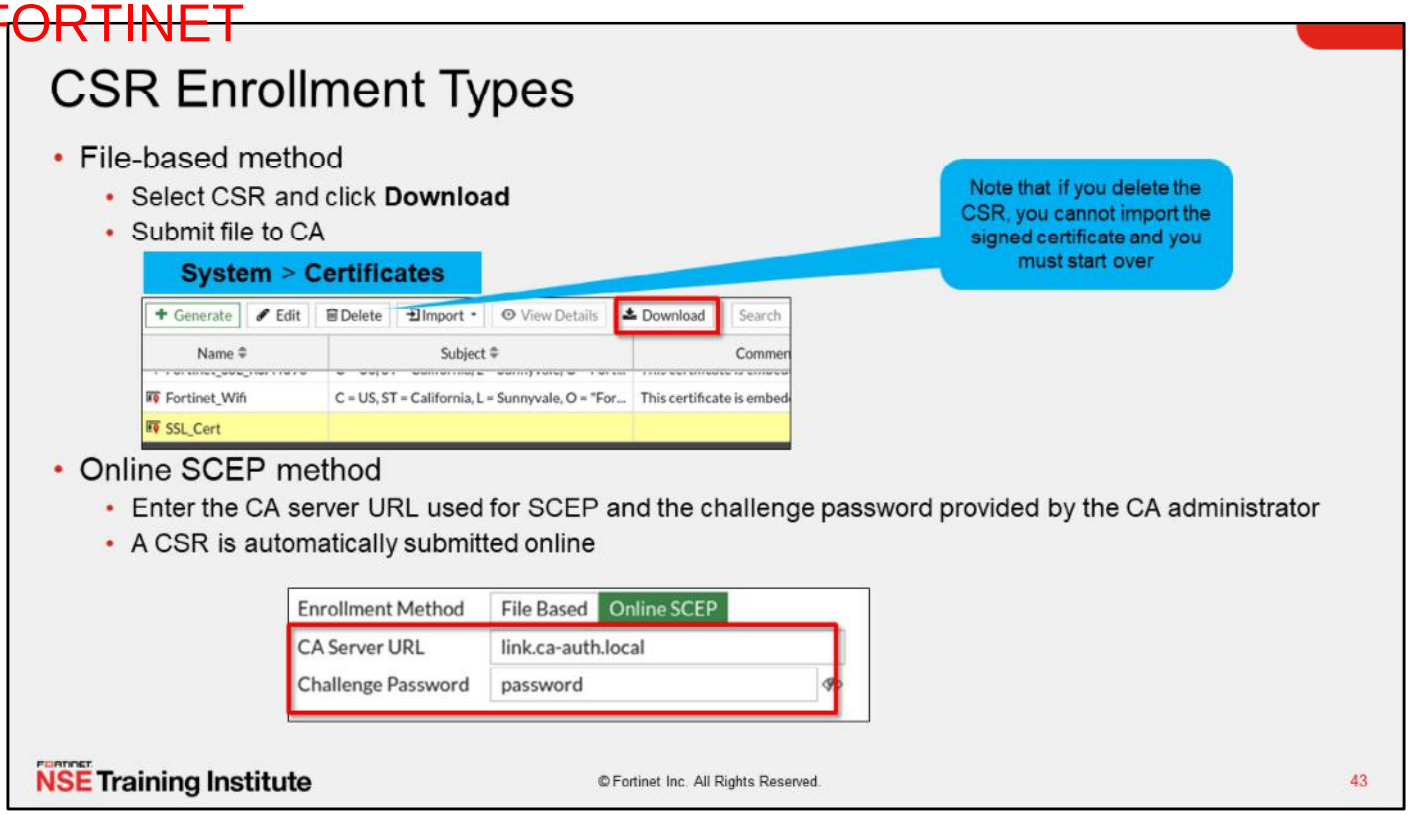

If you are using the file-based method, the CSR is added to your list of certificates on the **Certificates** page. Select the CSR and click **Download**. The administrator can now submit the file (`.csr`), which is a PKCS#10 request, to the CA. PKCS#10 is the most common format for a certificate request. The CA uses this file to generate a signed certificate.

If using the online SCEP method, enter the CA server URL used for SCEP and the challenge password provided by the CA administrator. The CSR is automatically submitted online.

After the CSR is submitted using either method, FortiGate shows the certificate status as **Pending** until the certificate is returned by the CA and imported into FortiGate. At this point, the status changes to **Valid** and the digital certificate can be used.

Note that if you delete the CSR, you cannot install the certificate and you must start over.

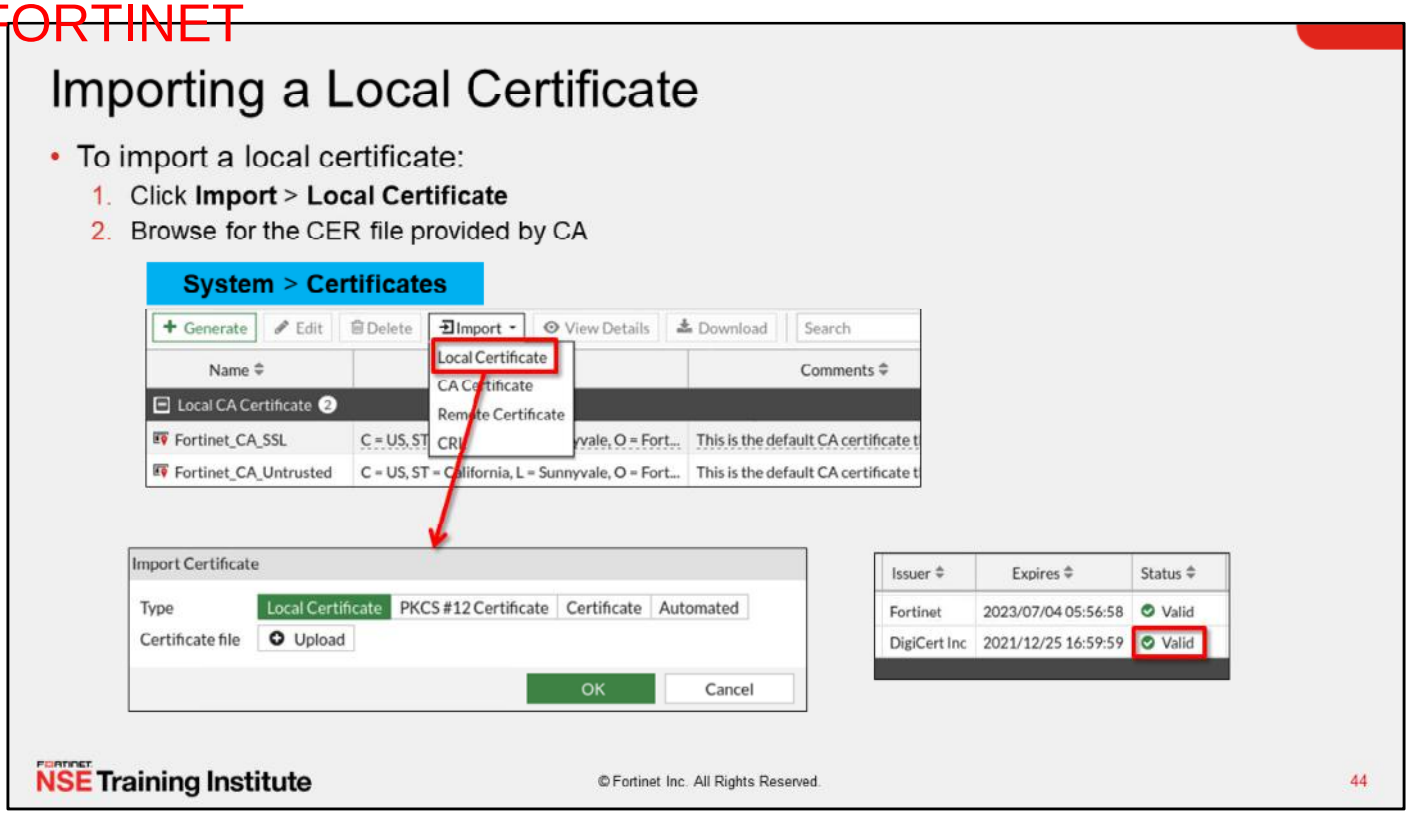

The file-based method of submitting a CSR is a manual process. The SCEP process, which occurs automatically online, requires no manual file import.

You can import the certificate from the **Certificates** page. Click **Import** and select **Local Certificate**. On the **Import Certificate** dialog, in the **Type** field, select **Local Certificate** and browse to the CER file provided by the CA.

After you import the certificate, the status changes from **Pending** to **Valid**. Note that it is possible to add a certificate that FortiGate uses in SSL communications without generating and signing a CSR. The CA can create a certificate for your FortiGate without a CSR (though the CA is responsible for providing all the certificate details for your FortiGate device). In this way, you can add a certificate using the following methods:

- Upload a PKCS#12 file, which is a single file that includes the signed certificate file and the key file
- Upload both a certificate file and the key file

An administrator user with the super_admin profile can put a password on a certificate and control access to its private key.

Importing a Local Certificate—Automated

- Automated Certificate Management Environment protocol (ACME)
  - CA management services that support ACME
- Let's Encrypt CA offers public Free SSL server certificates
- Configure FortiGate to use certificate managed by Let's Encrypt
  - Use server certificates for secure administrator log in
- To import ACME certificates:
  - FortiGate must have public IP address and FQDN
  - Public facing interface must have no VIPs on port 80 and 443
  - Subject Alternative Name (SAN) cannot be edited and filled with FortiGate FQDN

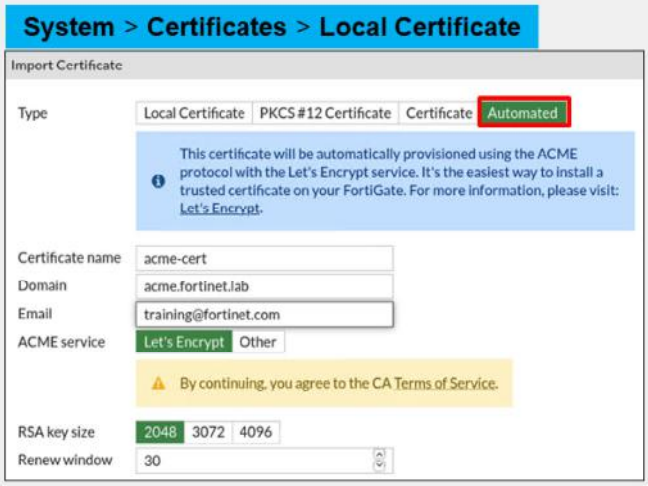

NSE Training Institute

45

You can also import local certificate which gets provisioned by an external service using the ACME protocol. Defined in RFC 8555, automated local certificate import can use the public Let's Encrypt CA to provide free SSL server certificates. You can configure FortiGate to use a certificate managed by Let's Encrypt and other certificate management services that use the ACME protocol.

Importing ACME certificates must meet the following criteria:

- FortiGate must have a public IP address and a hostname FQDN that resolves to the same public IP address.
- The public facing interface that ACME can access must have no virtual IP configured to forward port 80 (HTTP) or 443 (HTTPS). The interface has to listen for ACME update requests.
- SAN field is automatically filed with the FortiGate FQDN. It cannot be edited or allowed to have multiple SANs or wildcards.

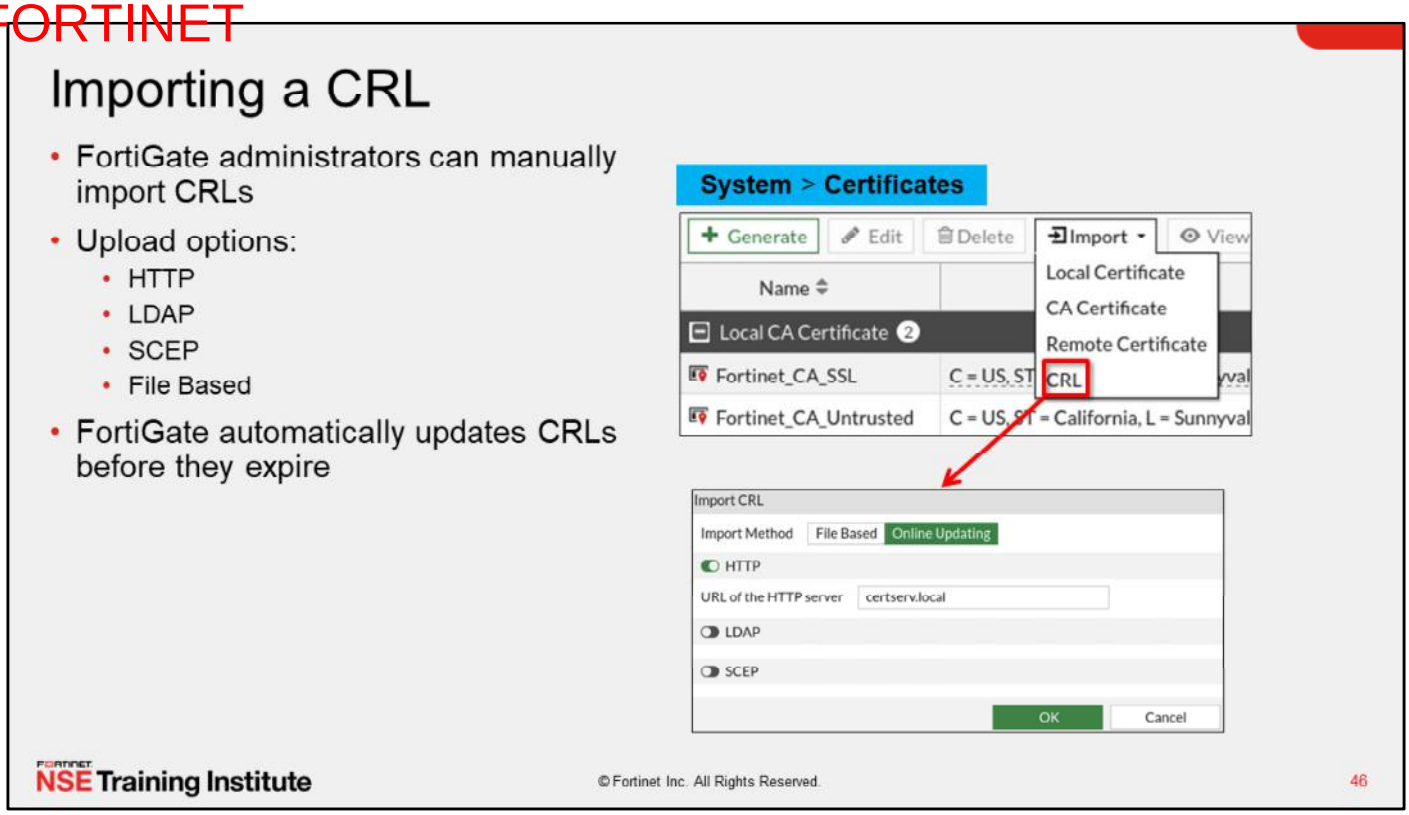

When FortiGate is validating a certificate, it checks that the certificate serial number is not listed in a CRL imported to FortiGate.

You can import a CRL from the **Certificates** page by clicking **Import** > **CRL**. In the **Import CRL** dialog, you can select one of these four import options: **HTTP**, **LDAP**, **SCEP**, and **File Based**. The first three options point to external repositories and require you to connect to the repositories to upload the CRL to FortiGate. The last option, **File Based**, requires you to have the CRL file locally stored before you can upload the CRL to FortiGate.

Before the CRL expires, FortiGate automatically retrieves the latest iteration using the protocol specified in the configuration.

Backing Up and Restoring Certificates

- Back up keys and certificates through the CLI (TFTP server required for import and export):
  - `execute vpn certificate local import tftp <file-name_str> <tftp_ip>.`
  - `execute vpn certificate local export tftp <certificate-name_str> <file-name_str> <tftp_ip>.`
- Keys and certificates are stored in the PKCS#12 file
- Configuration backup also contains the keys and certificates

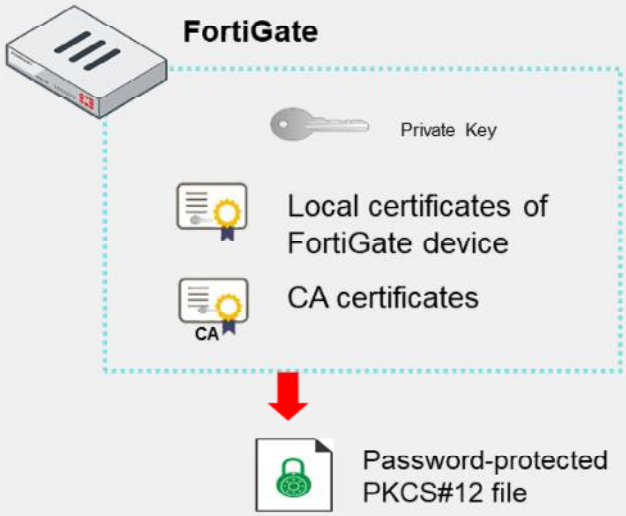

NSE Training Institute

47

When you back up the FortiGate configuration, the keys and certificates are backed up as well.

FortiGate also provides the option to store digital certificates as a PKCS#12 file, which includes the private and public keys as well as the certificate. You can restore the PKCS#12 file to a FortiGate device of any model or firmware version, or to a non-FortiGate device.

You can perform the backup and restore on the CLI only, which requires the use of a TFTP server.

Certificate Configuration—VDOM and Global

- You can configure CA and local certificates per VDOM

```
config vpn certificate local
   edit Fortinet_Factory
      set range <global/vdom>
      set source <factory/user/bundle/fortiguard>
   end
end
```

FortiGate certificates, some identified with specific signature algorithms and key lengths in their names

NSE Training Institute

48

You can configure a CA and local certificate globally or for a VDOM. If you upload a certificate to a VDOM, it is accessible only inside that VDOM. If you upload a certificate globally, it is accessible to all VDOMs and globally.

Global and VDOM-based certificate configuration includes the ability to view certificate details, as well as to download, delete, and import certificates.

Note that some of the FortiGate certificates have specific signature algorithms and key lengths in their names, such as Elliptic Curve Digital Signature Algorithm 256 (ECDSA256) and RSA2048. Policy and technical requirements may determine which certificates you use.

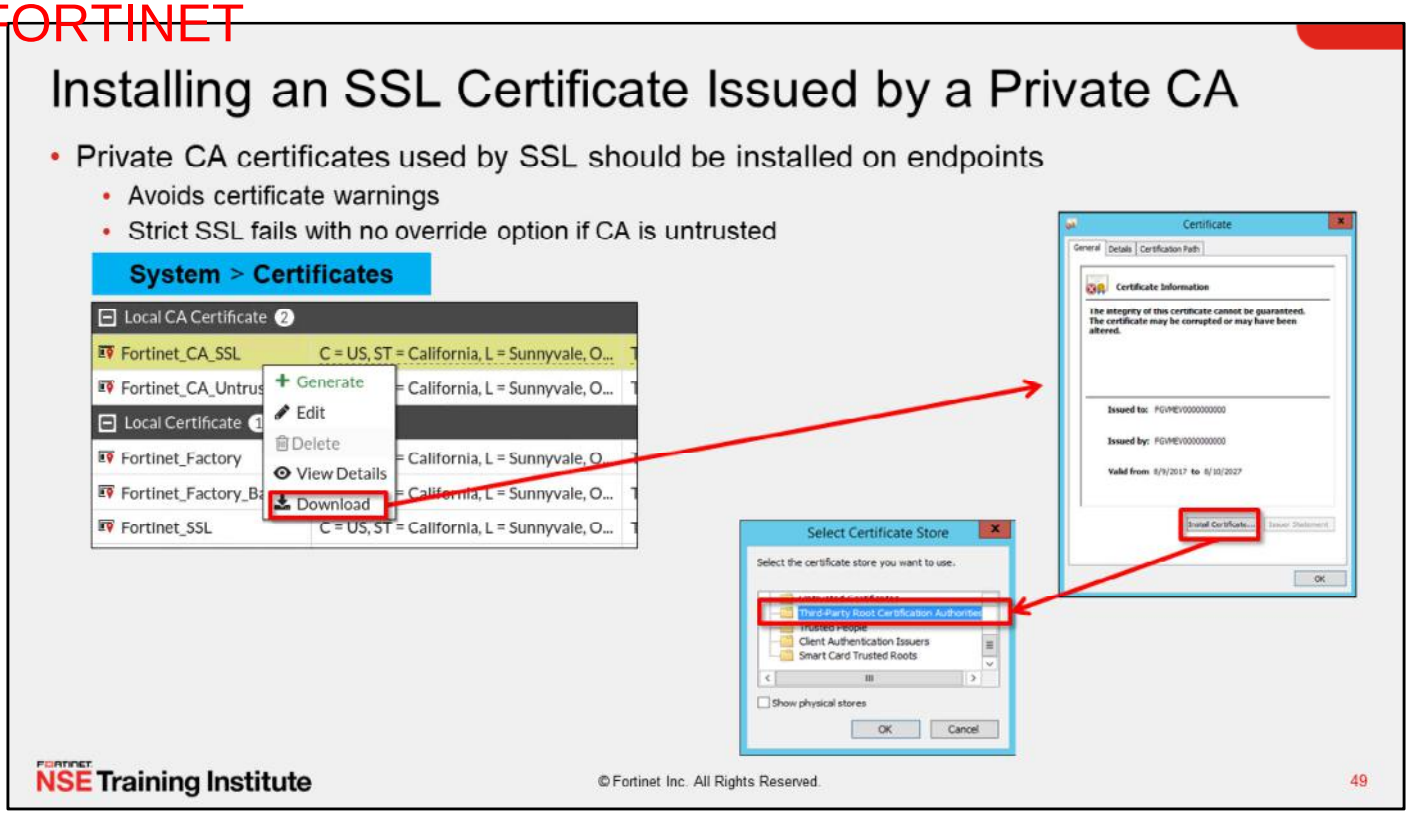

If you are using an SSL certificate issued by a private CA, you must install the CA certificate in the list of trusted CAs. If you fail to do this, a warning message appears in your web browser any time you access an HTTPS website. Encrypted communications might also fail, simply because the CA that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

Once 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 avoid certificate warnings, you need to 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 CSR enrollment method is supported by FortiGate?
   A. Enrollment over Secure Transport (EST)
   ✓ B. Simple Certificate Enrollment Protocol (SCEP)

2. After a CSR has been enrolled and imported into FortiGate, the status of the certificate should change to:
   ✓ A. Valid
   B. Pending

NSE Training Institute

50

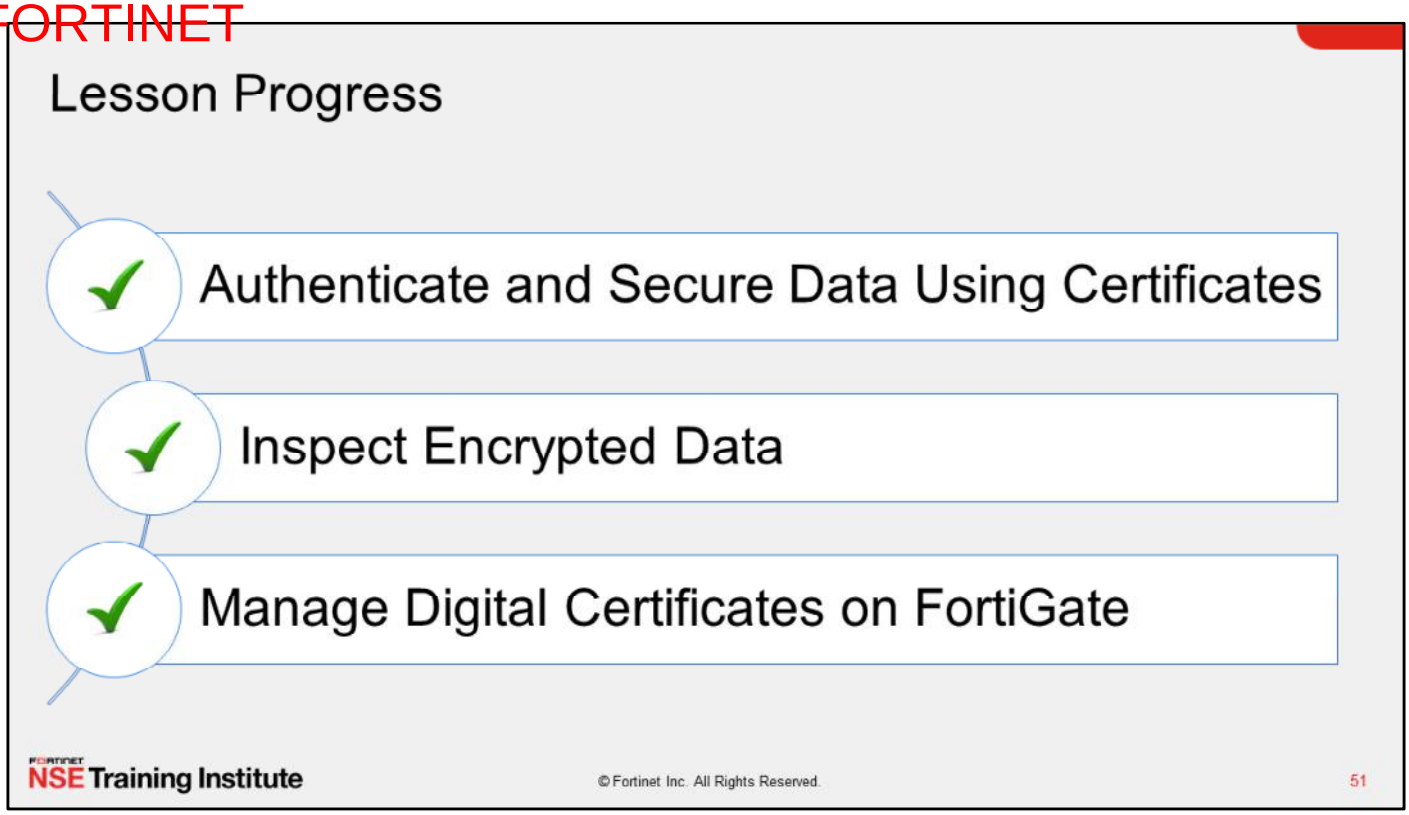

Congratulations! You have completed this lesson.

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

Review

- ✓ Describe why FortiGate uses digital certificates
- ✓ Describe how FortiGate uses certificates to authenticate users and devices
- ✓ Describe how FortiGate uses certificates to ensure the privacy of data
- ✓ Describe certificate inspection and full SSL inspection
- ✓ Configure certificate inspection and full SSL/SSH inspection
- ✓ Identify what is required to implement full SSL inspection
- ✓ Identify the obstacles to implementing full SSL inspection and possible remedies
- ✓ Generate a certificate request
- ✓ Import CRLs
- ✓ Back up and restore certificates

NSE Training Institute © Fortinet Inc. All Rights Reserved. 52

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.

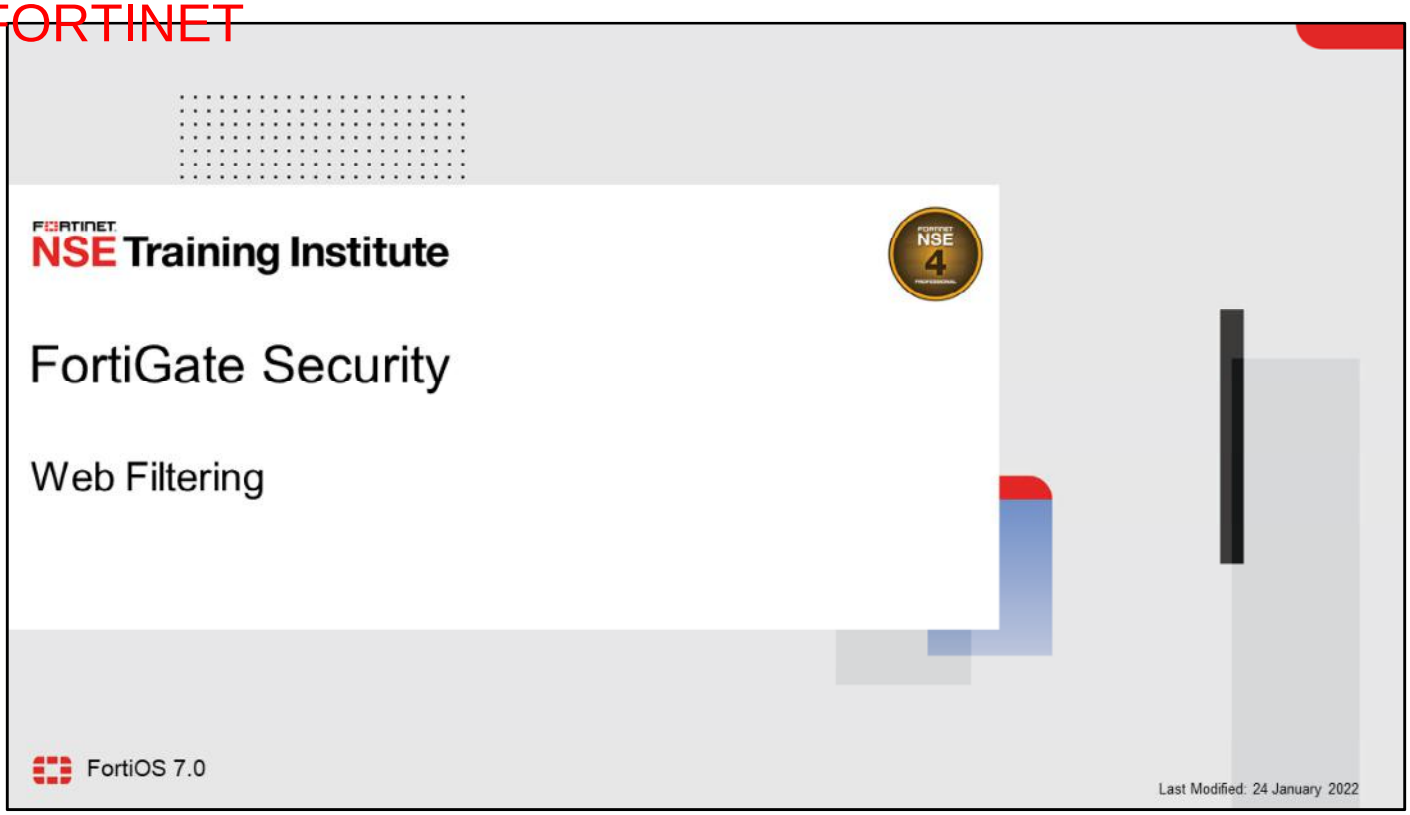

In this lesson, you will learn how to configure web filtering on FortiGate to control web traffic in your network.

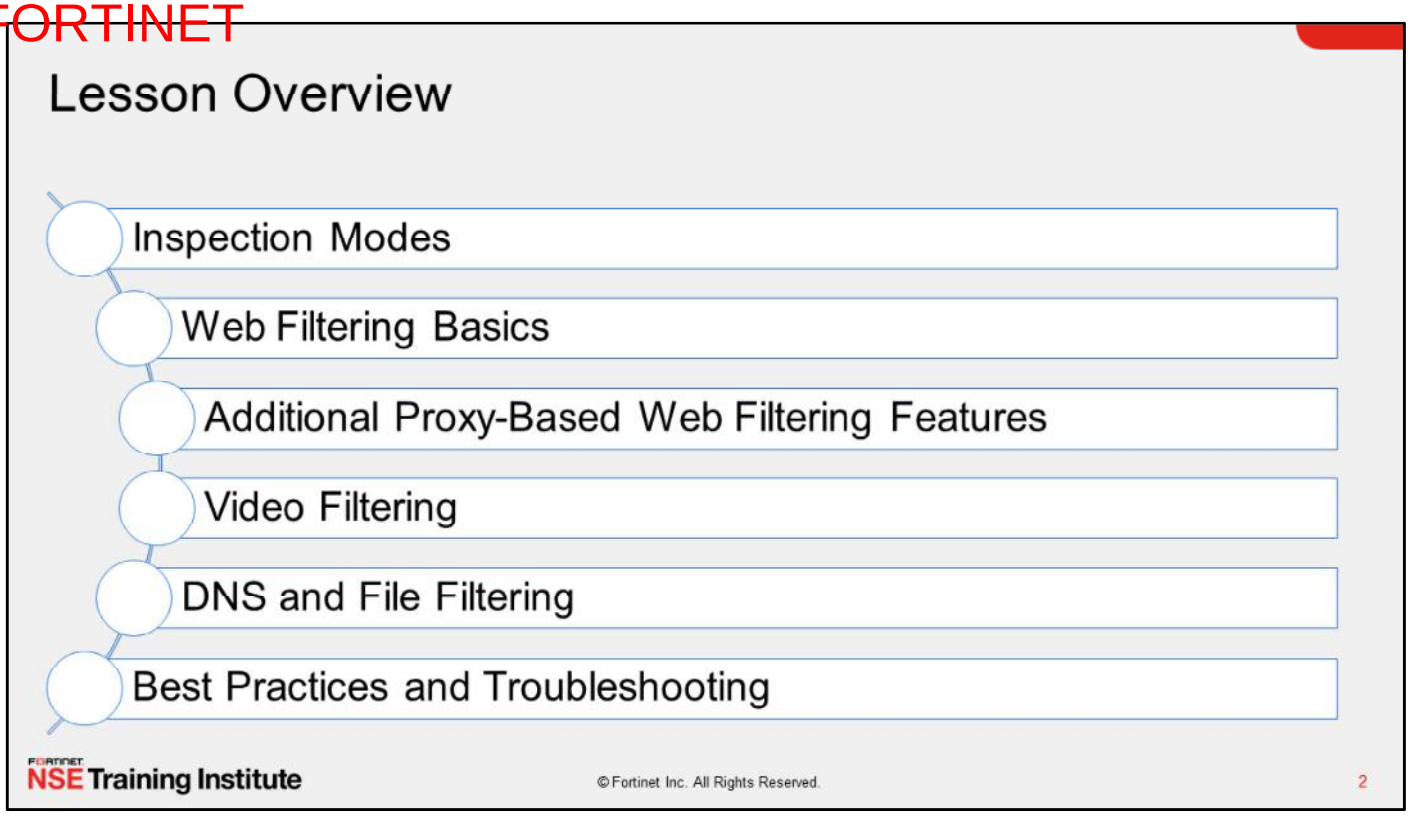

In this lesson, you will learn about the topics shown on this slide.

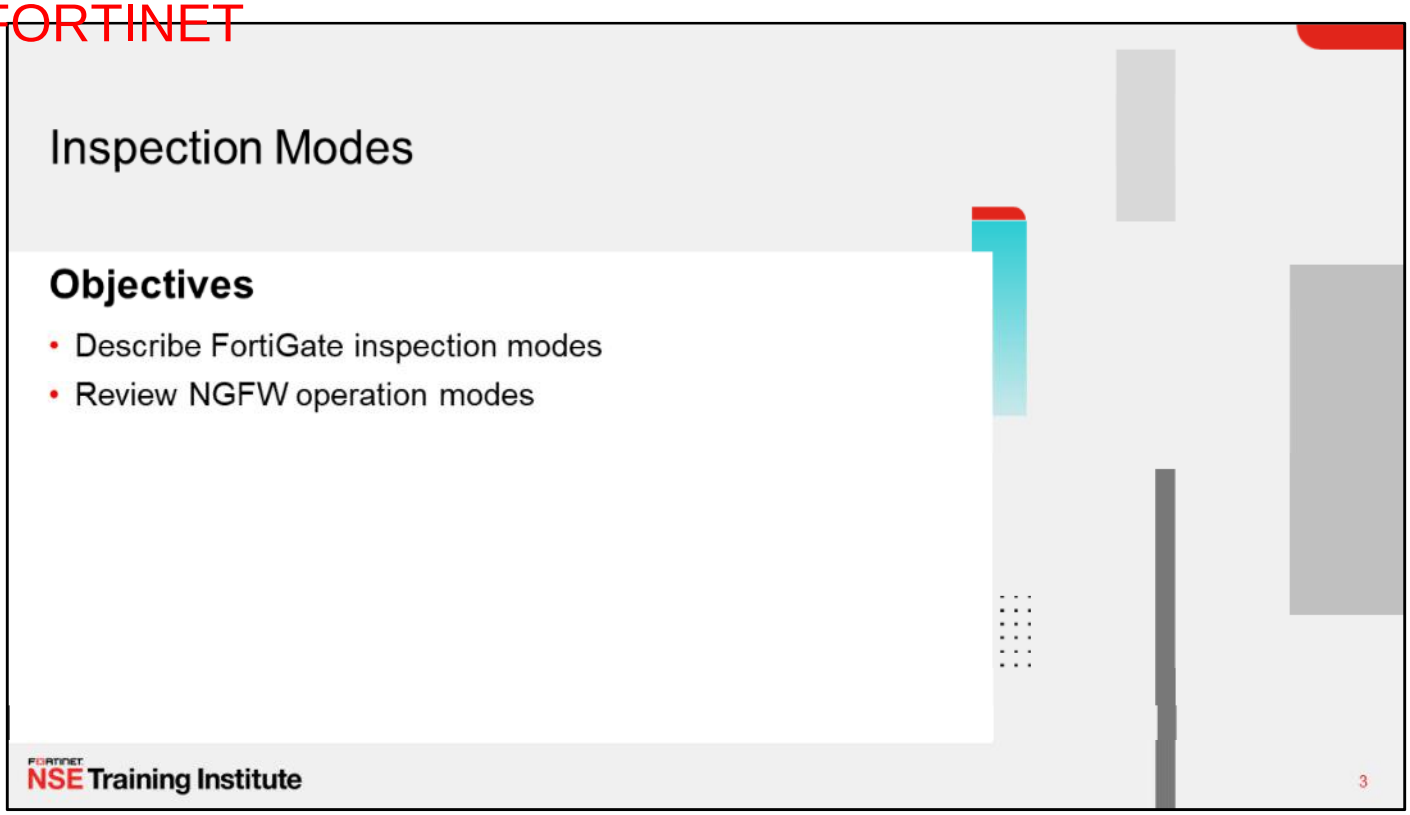

After completing this section, you should be able to achieve the objectives shown on this slide.

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

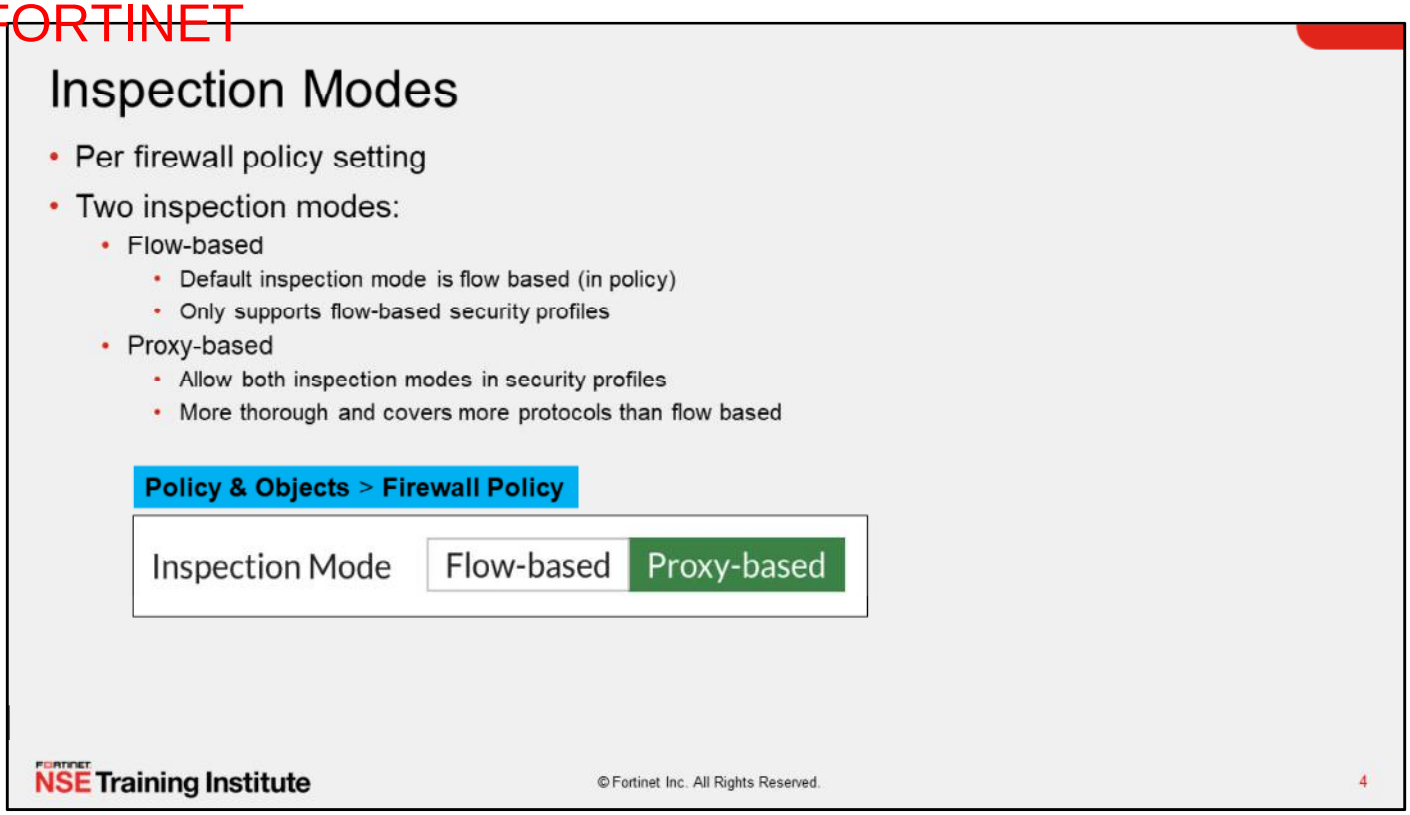

Each inspection component plays a role in processing traffic on its way to its destination. Having control over flow-based and proxy-based mode is helpful if you want to be sure that only flow-based inspection mode is used. In most cases, proxy mode is preferred because more security profile features and more configuration options are available. However, some implementations require all security profile scanning to use only flow-based inspection mode for the highest possible throughput. You can configure the firewall policy to use flow-based mode (which is the default option for a new policy), and vice versa. While both modes offer significant security, proxy-based mode is more thorough while flow-based mode is designed to optimize performance.

You can select the inspection mode in a firewall policy. Switching from flow-based to proxy-based mode will not require removing the selected security profiles on the policy. However, switching from proxy-based to flow-based mode will remove any security profiles configured to use proxy-based inspection mode.

Flow-Based Inspection

- Default inspection mode
- Uses single-pass direct filter approach (DFA) pattern matching to identify possible attacks or threats
- File is scanned on a flow basis as it passes through FortiGate
- Requires fewer processing resources
- Faster scanning

NSE Training Institute

5

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

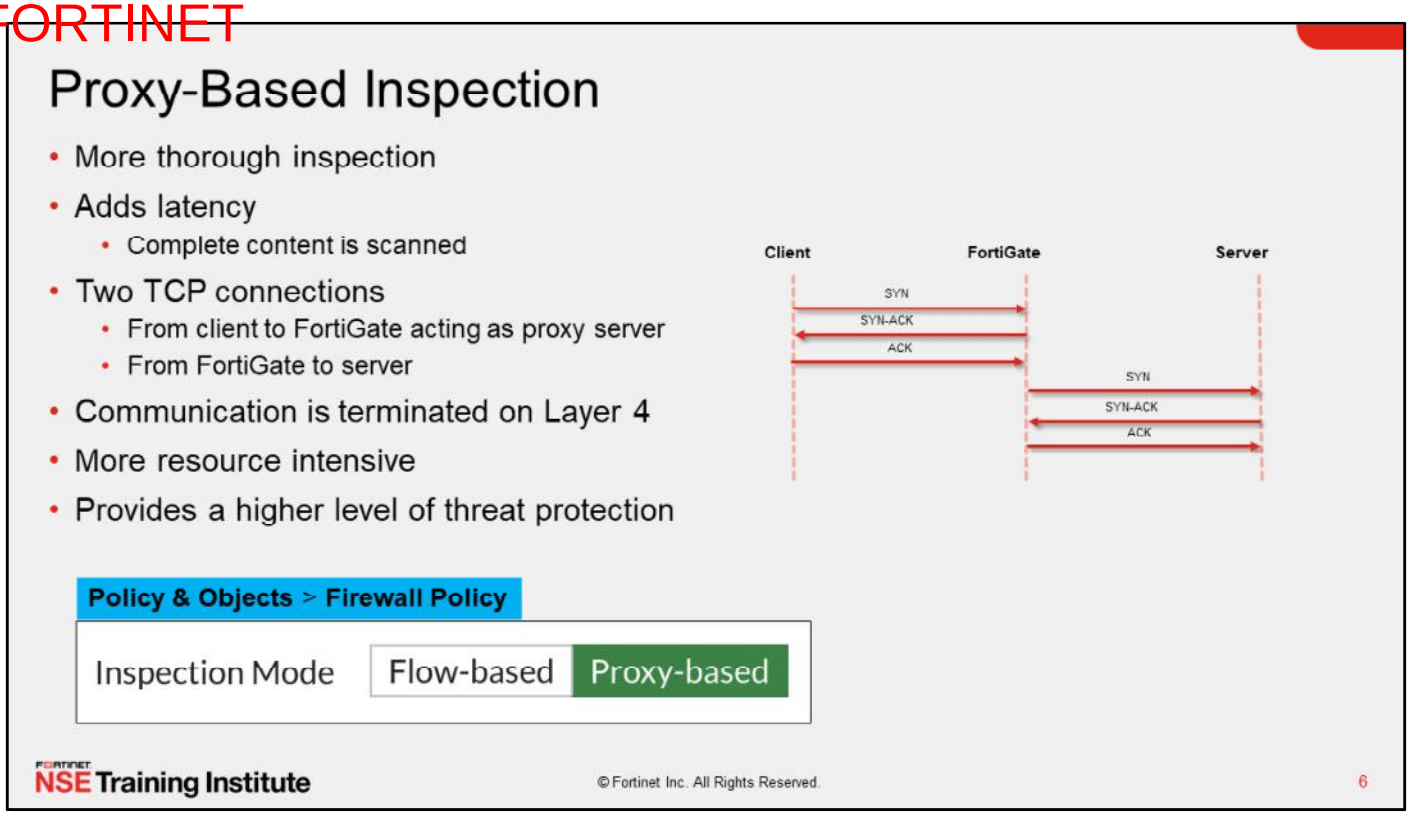

Proxy-based scanning refers to transparent proxy. It's called transparent because, at the IP layer, FortiGate is not the destination address, but FortiGate *does* intercept the traffic. When proxy-based inspection is enabled, FortiGate buffers traffic and examines it *as a whole*, before determining 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.

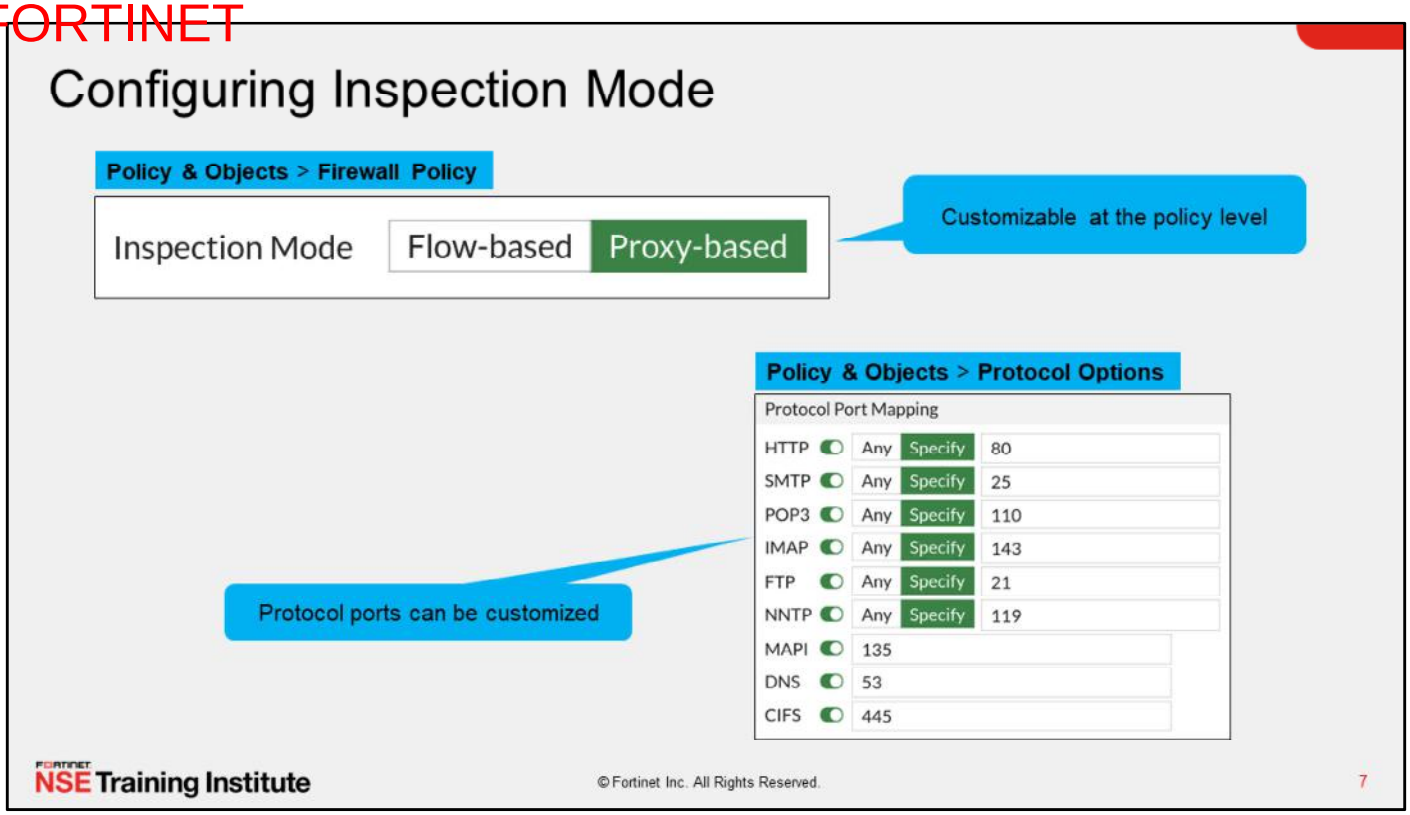

FortiGate web filters are also security profiles. The security profiles are customizable, according to the selected inspection mode. So, the first step, before setting up a web filter, is to configure the inspection mode.

The protocol options profile determines the protocols your security profiles use, for example, to inspect web or DNS traffic.

Note that HTTPS inspection port numbers, and other settings related to the handling of SSL, are defined separately in the SSL/SSH inspection profile.

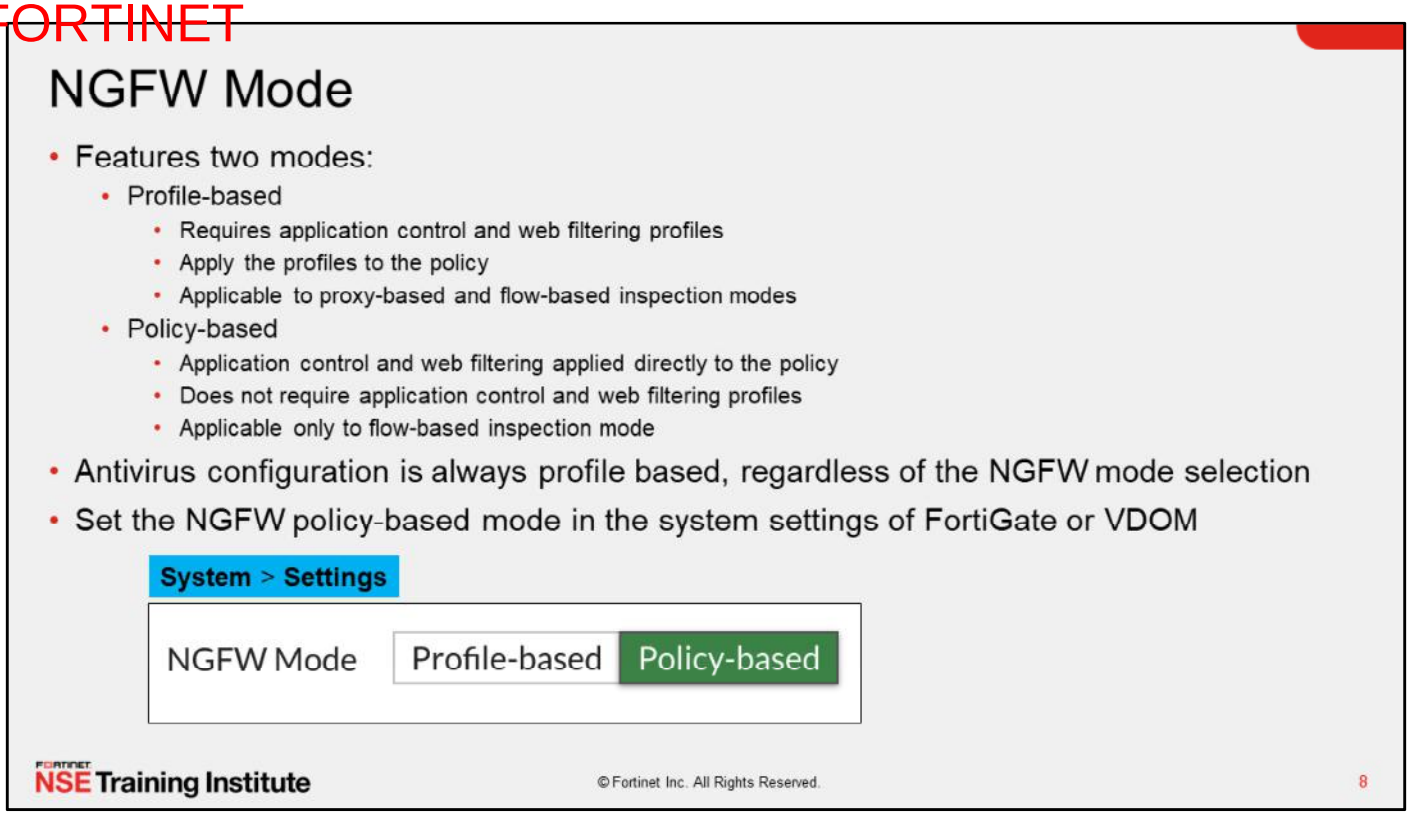

FortiGate, or the individual VDOM, has two next-generation firewall (NGFW) modes available:

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.

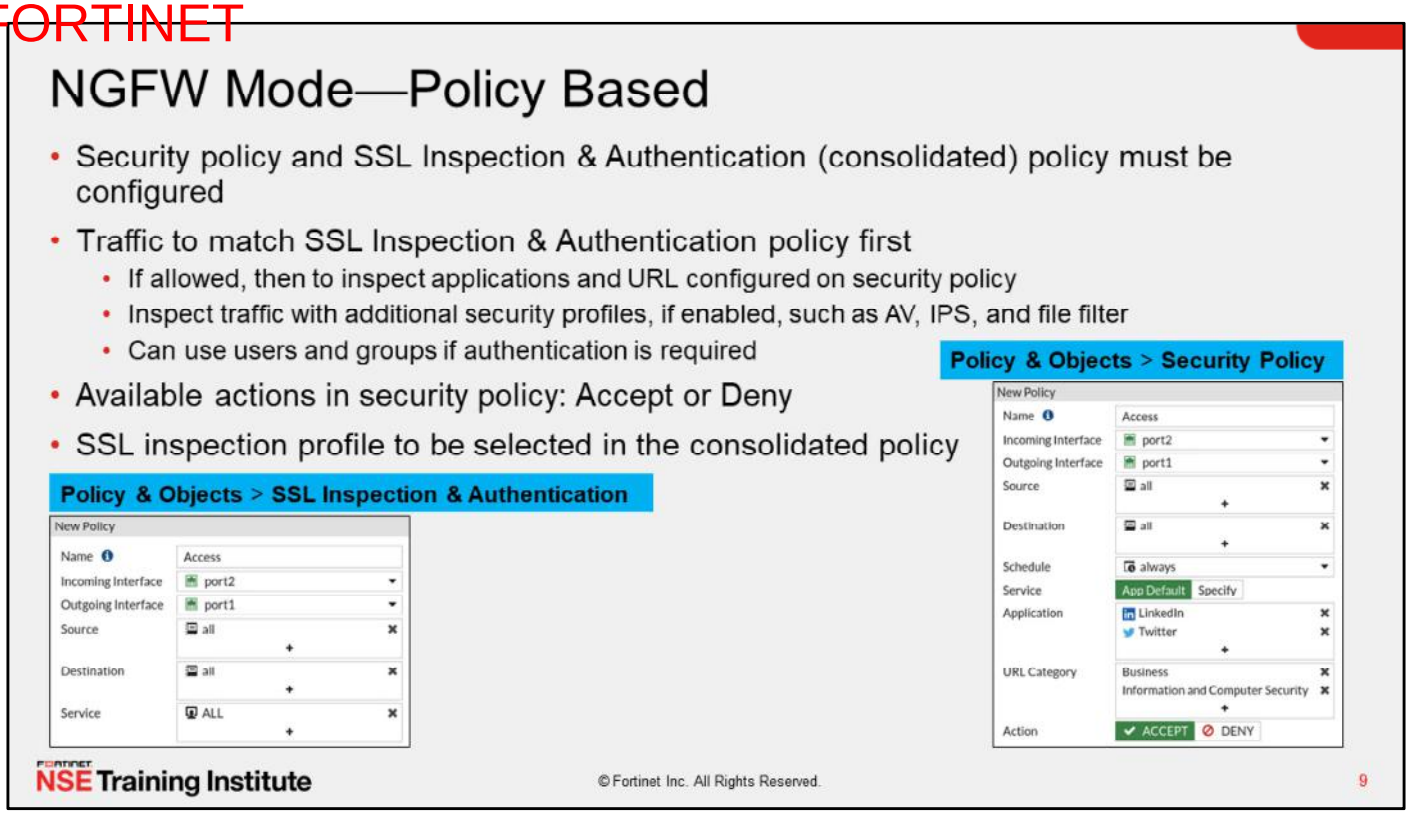

If you configured FortiGate to use NGFW policy-based mode or created a VDOM specifically to provide NGFW policy-based mode, you must configure a few policies to allow traffic.

SSL Inspection & Authentication (consolidated) policy: This 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 as per the consolidated policy, FortiGate then processes it based on the security policy to analyze additional criteria, such as URL categories 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. Which is the default inspection mode on a firewall policy?
   A. Proxy based
   ✓ B. Flow based

2. 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.

3. 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

NSE Training Institute

10

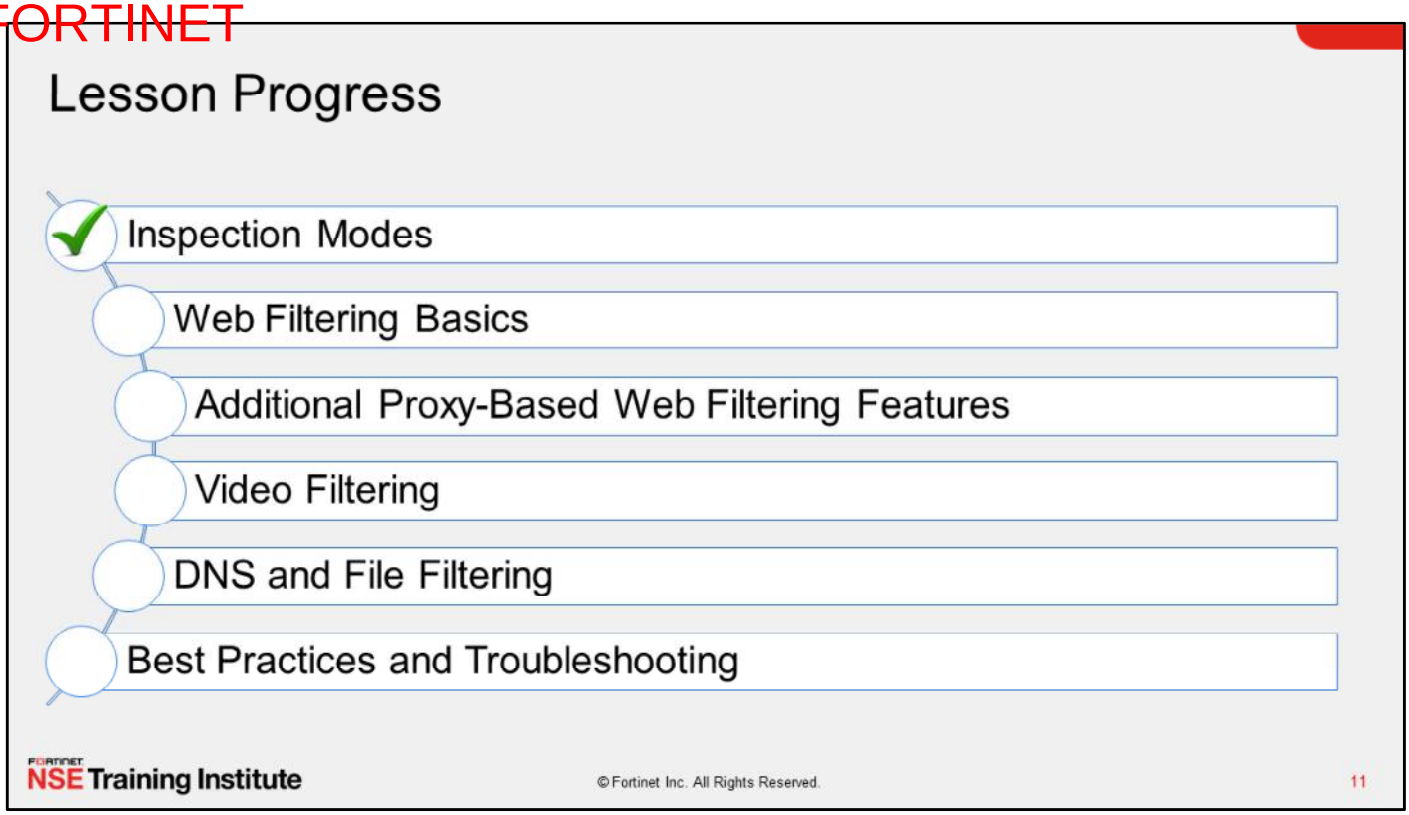

Good job! You now understand inspection modes.

Now, you will learn about web filtering basics.

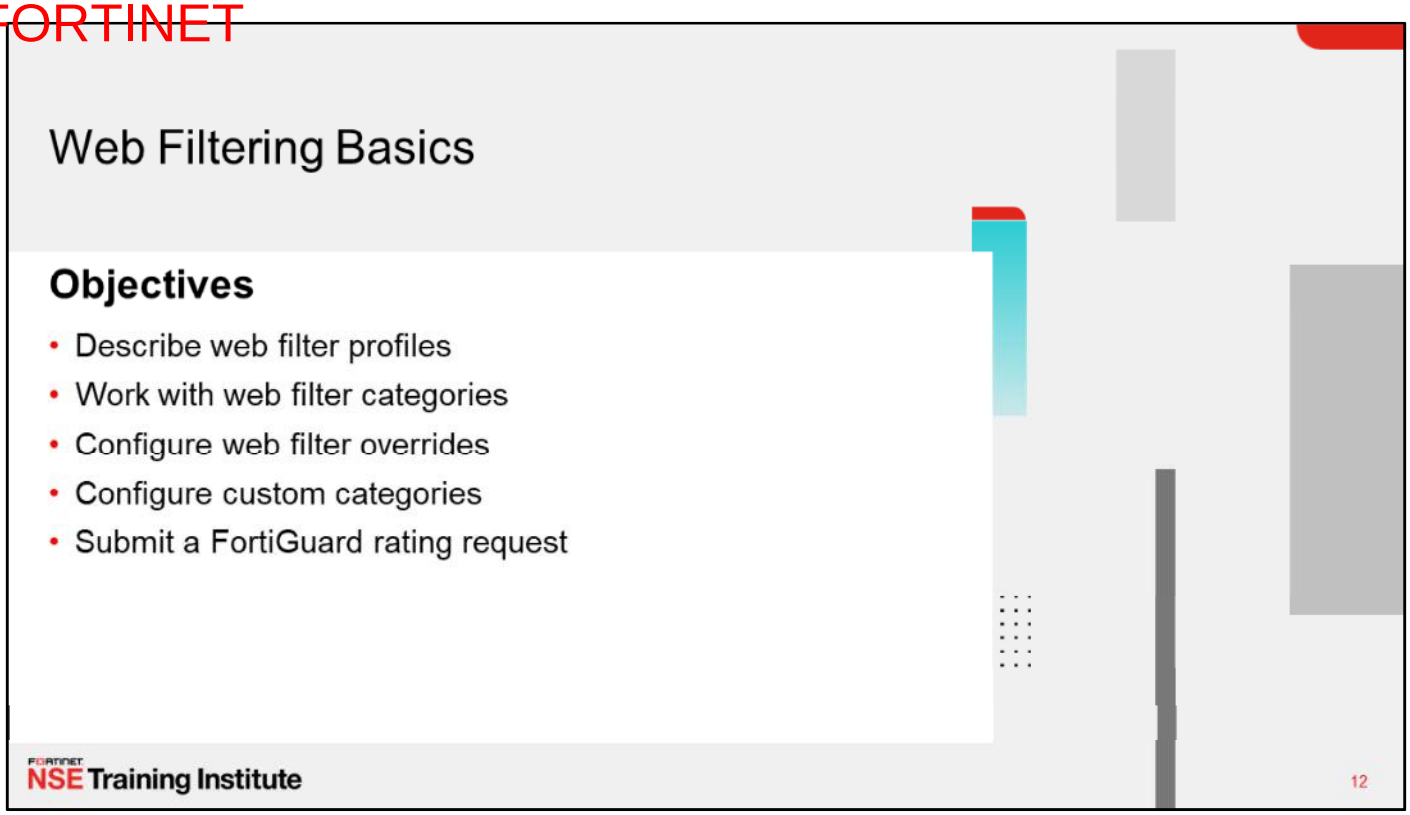

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in web filtering basics, you will be able to describe web filter profiles, use FortiGuard web filter profiles, configure web filter overrides, define custom categories, and submit FortiGuard rating requests.

Why Apply Web Filtering?

- Mitigate the negative effects of inappropriate web content
- Preserve employee productivity
- Prevent network congestion
- Prevent data loss and exposure of confidential information
- Decrease exposure to web-based threats
- Prevent copyright infringement
- Prevent viewing of inappropriate or offensive material

NSE Training Institute

13

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

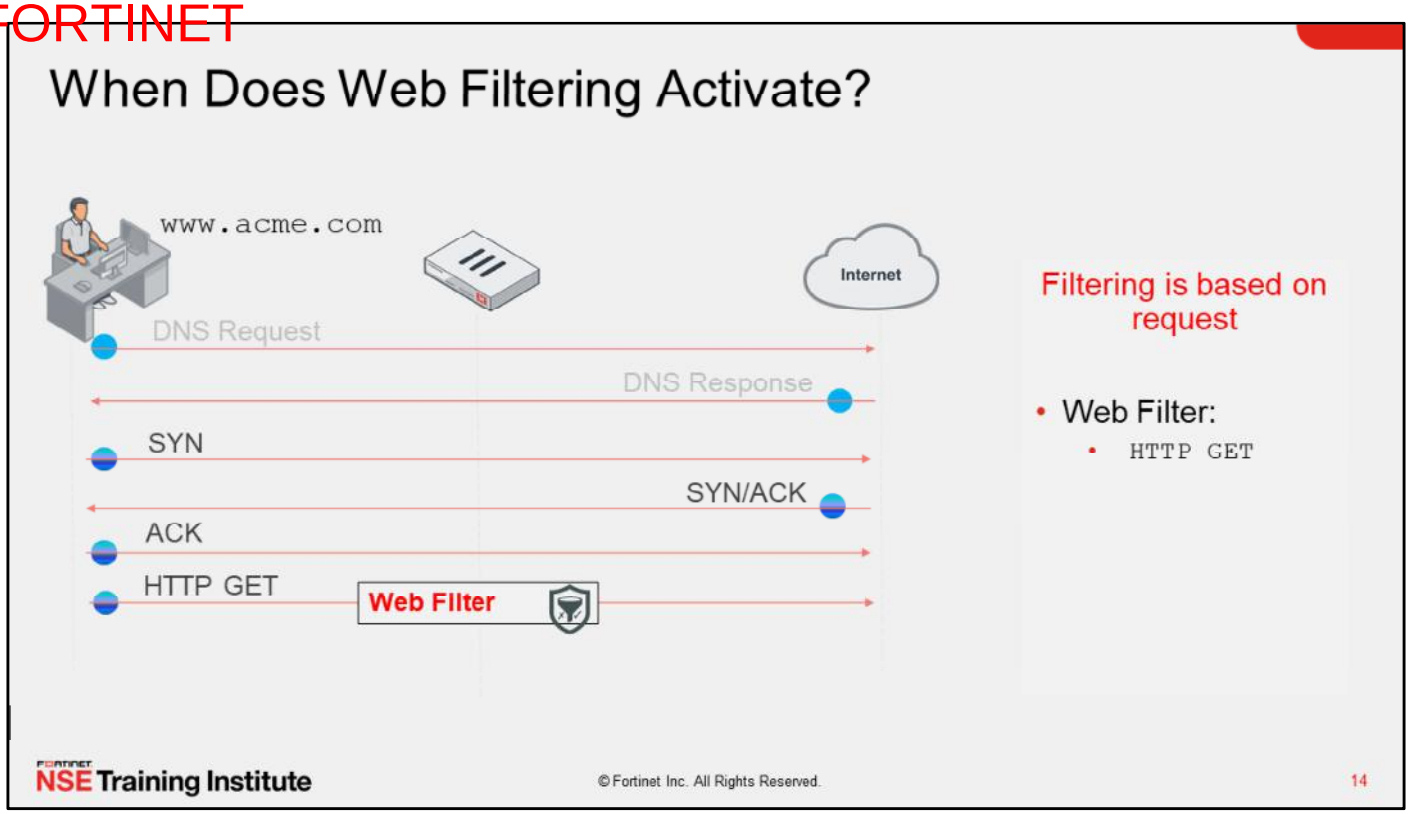

(slide contains animation)

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.

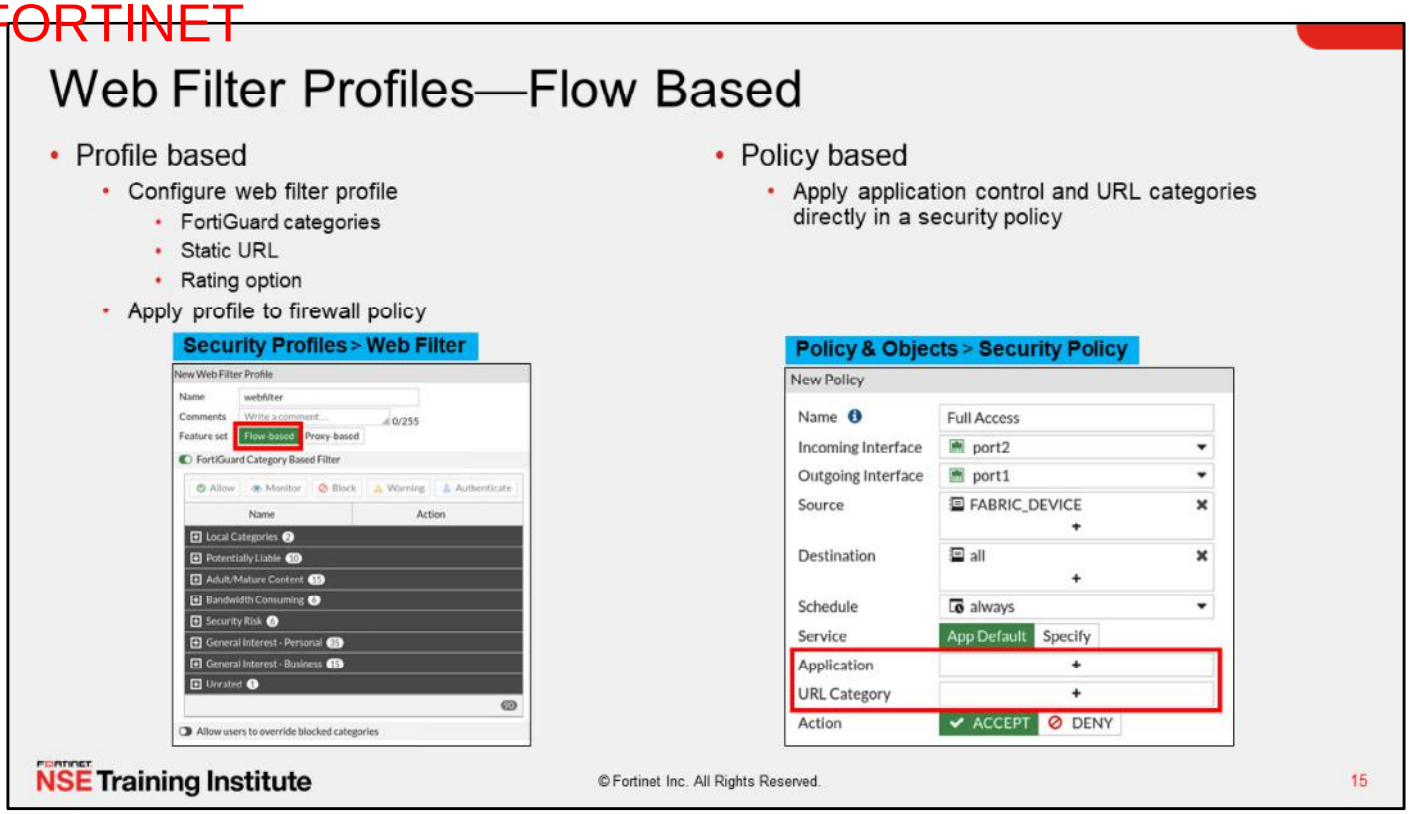

Now, you will look at the web filter profile.

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

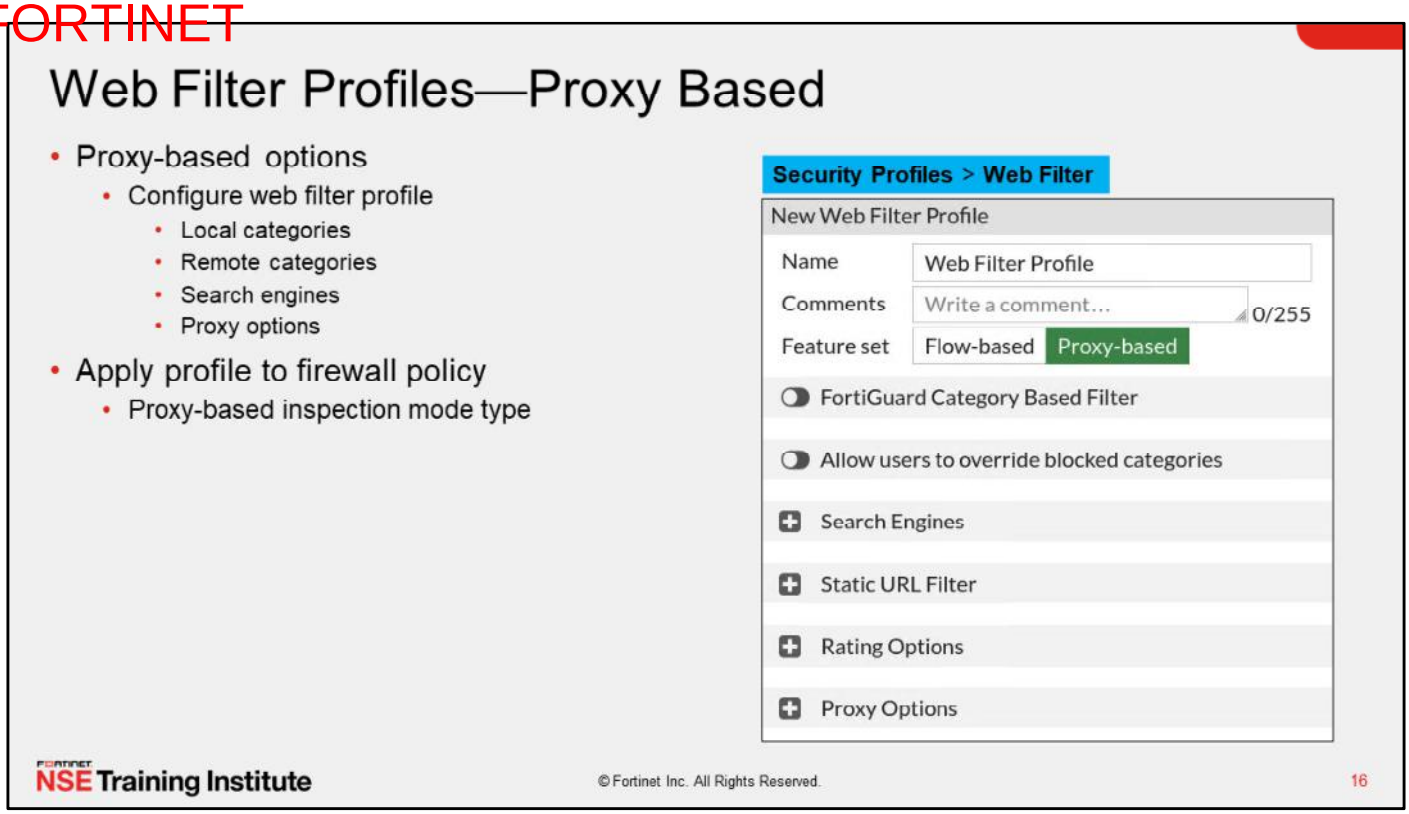

In the example shown on this slide, the security profile is configured to use a proxy-based feature set. The profile is available to a firewall policy configured to use proxy-based inspection mode. Other local options include:

- **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.

FortiGuard Category Filter

- Split into multiple categories and subcategories
  - Release new categories and subcategories compatible with updated firmware
  - Older firmware has new values mapped to existing categories
- Live connection to FortiGuard
  - Active contract required
  - Two-day grace period on expiry
- Can use FortiManager instead of FortiGuard

NSE Training Institute © Fortinet Inc. All Rights Reserved. 17

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.

How Are Categories Decided?

- FortiGate queries the FortiGuard Distribution Network (FDN) to determine a website category
- The web filter rating is determined by:
  - Human rater
  - Text analysis
  - Exploitation of web structure
- Description of categories:
  - www.fortiguard.com/webfilter/categories

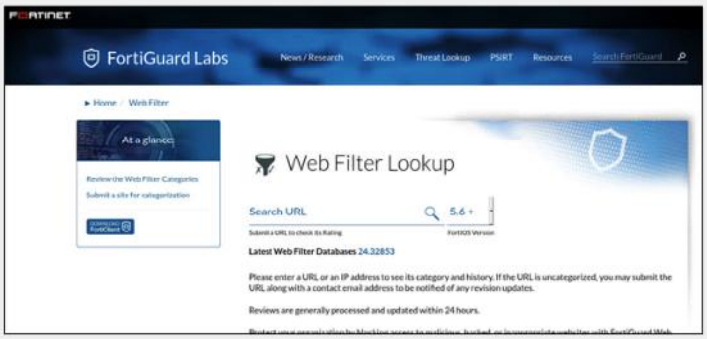

NSE Training Institute

18

Website categories are determined by both automatic and human methods. The FortiGuard team has automatic web crawlers that look at various aspects of the website in order to come up with a rating. There are also people who examine websites and look into rating requests to determine categories.

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

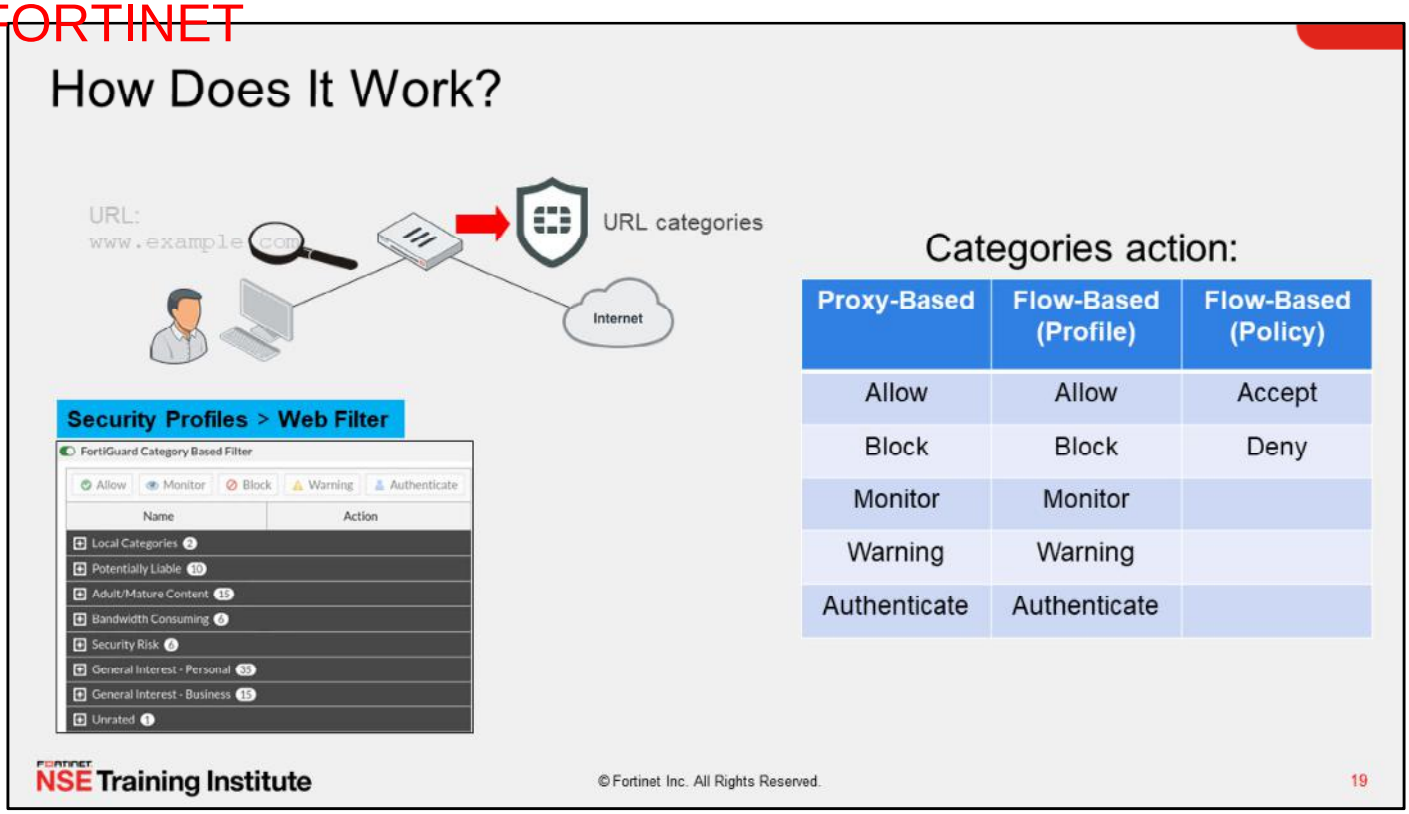

So, how does it work?

FortiGate queries the FDN—or FortiManager, if it has been configured to act as a local FortiGuard server—to determine the category of a requested web page.

When users visit websites, FortiGate uses the FortiGuard live service to determine the category that the URL belongs to and takes a configured action for that category, such as allow or block access. Using this feature, you can perform bulk URL filtering, without individually defining each website.

You can enable the FortiGuard category filtering on the web filter, or DNS filter profiles. 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)

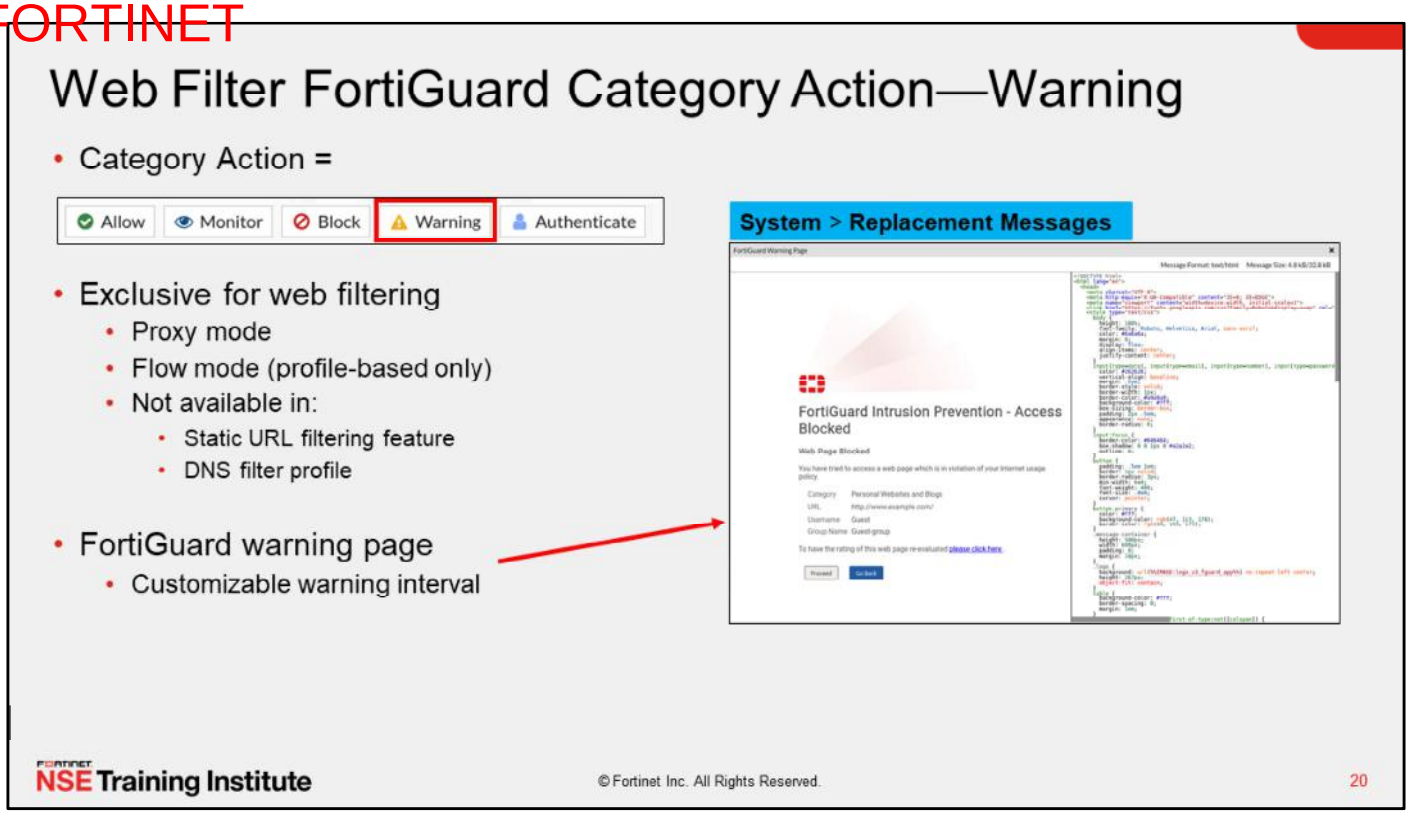

The warning action informs users that the requested website is not allowed by the internet policies. However, the action gives the user the option to proceed to the requested website, or return to the previous website.

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

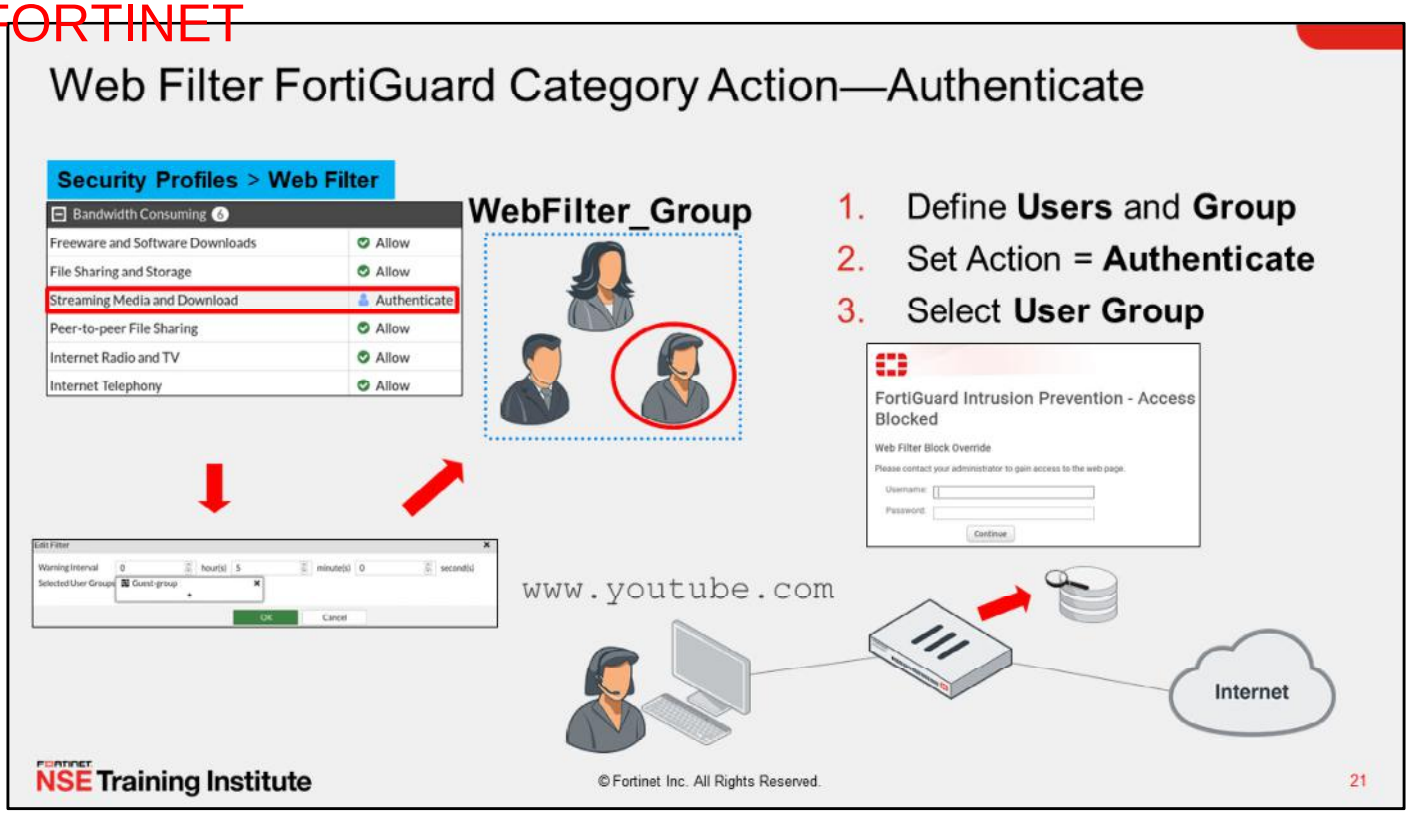

The authenticate action blocks the requested websites, unless the user enters a successful username and password.

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.

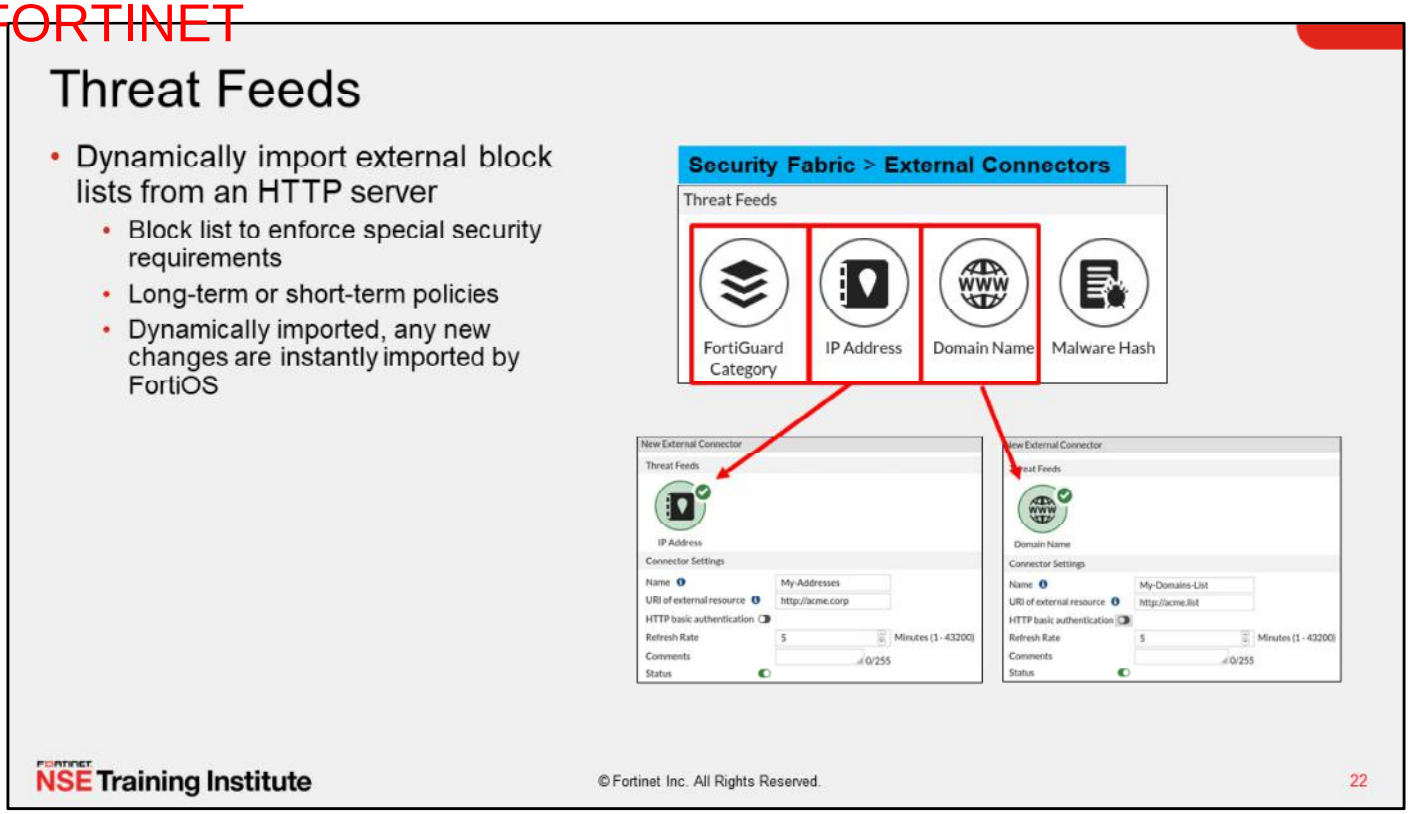

The Threat Feeds feature enables FortiGate to dynamically import external block lists from an HTTP server. You can use the block lists to enforce special security requirements specified by your organization. These requirements can include long-term policies to always block access to specific websites, or short-term requirements to block access to known compromised locations. These block lists are text files that are in plain text format, where each line contains a single URL to be blocked.

Because the lists are dynamically imported, any changes made to the list are instantly imported by FortiOS using the Security Fabric feature.

FortiGuard Category: This resource name appears as a remote category in web filter profiles and SSL inspection exemptions.
IP Address: This resource name appears as an external IP block list in DNS filter profiles and as a source/destination in IPv4 policy, IPv6, and Proxy policy.
Domain Name: This resource name will appear as a remote category in DNS filter profiles.
Refresh Rate: Using this setting, you can specify how often, in minutes, block lists can be refreshed from the external source.

The size of the block list file can be 10 MB or 128,000 lines of text, whichever is most restrictive.

Note that the DNS profile supports only IPv4 addresses and ignores IPv6 addresses.

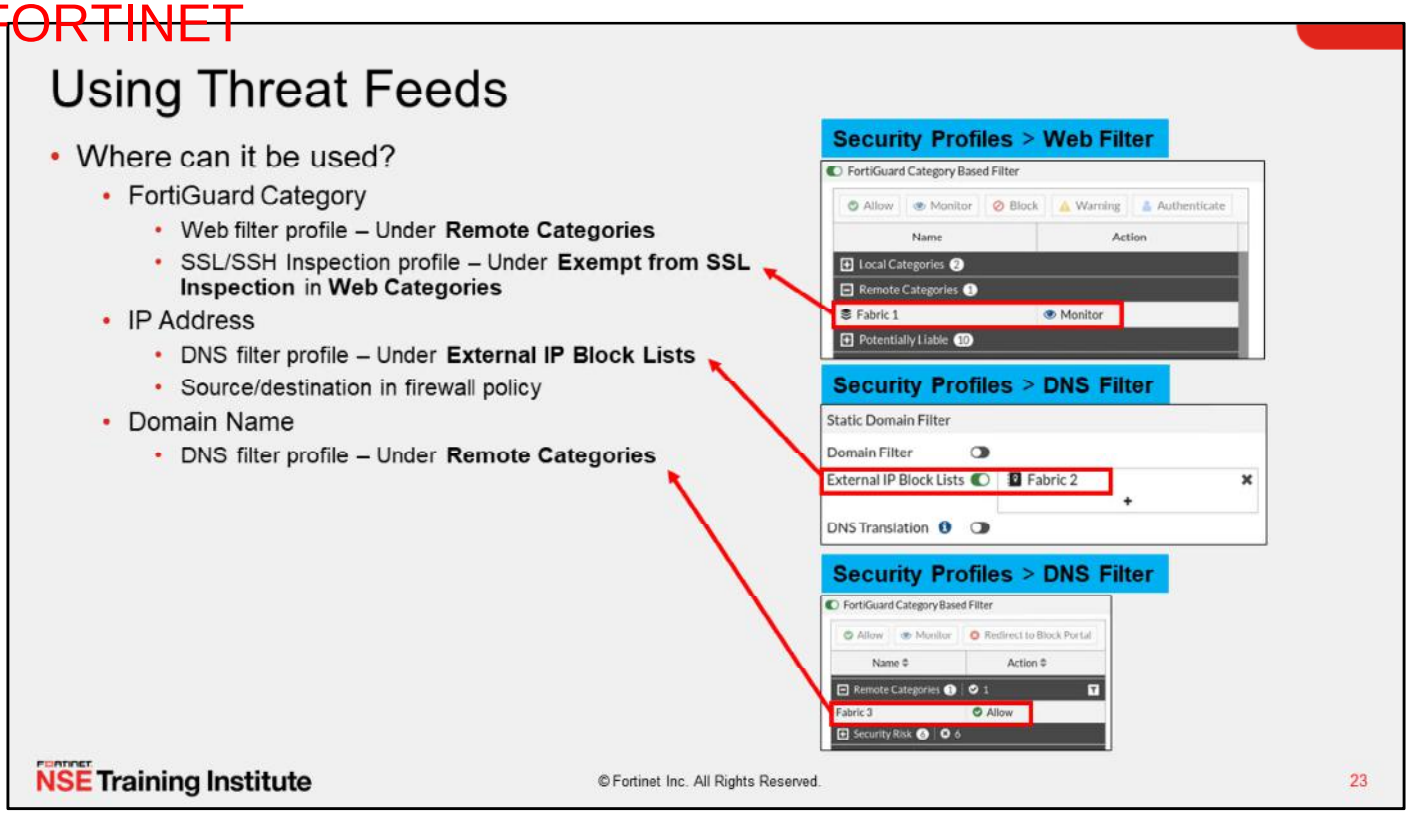

You can add dynamic block lists to:

- Web filter profiles and SSL inspection exemptions
- DNS filter profiles and source/destination addresses in firewall policies

Web Rating Override

- Override the rating applied to a host name by FortiGuard service
  - Host name reassigned to a completely different category and uses that action
  - Rating overrides are checked before contacting FortiGuard for a rating
- Override applies to FortiGate device only
  - Changes are not submitted to FortiGuard subscription services
- Host names only
  - `google.com` ✓
  - `www.google.com` ✓
  - `www.google.com/index.html` ✗
  - `google.*` ✗

NSE Training Institute © Fortinet Inc. All Rights Reserved. 24

When using FortiGuard category filtering to allow or block access to a website, one option is to make a web rating override and define the website in a different category. Web ratings are only for host names—no URLs or wildcard characters are allowed.

If the contract expires, and the two-day grace period passes, web rating overrides are not be effective. All website category rating requests are returned with a rating error.

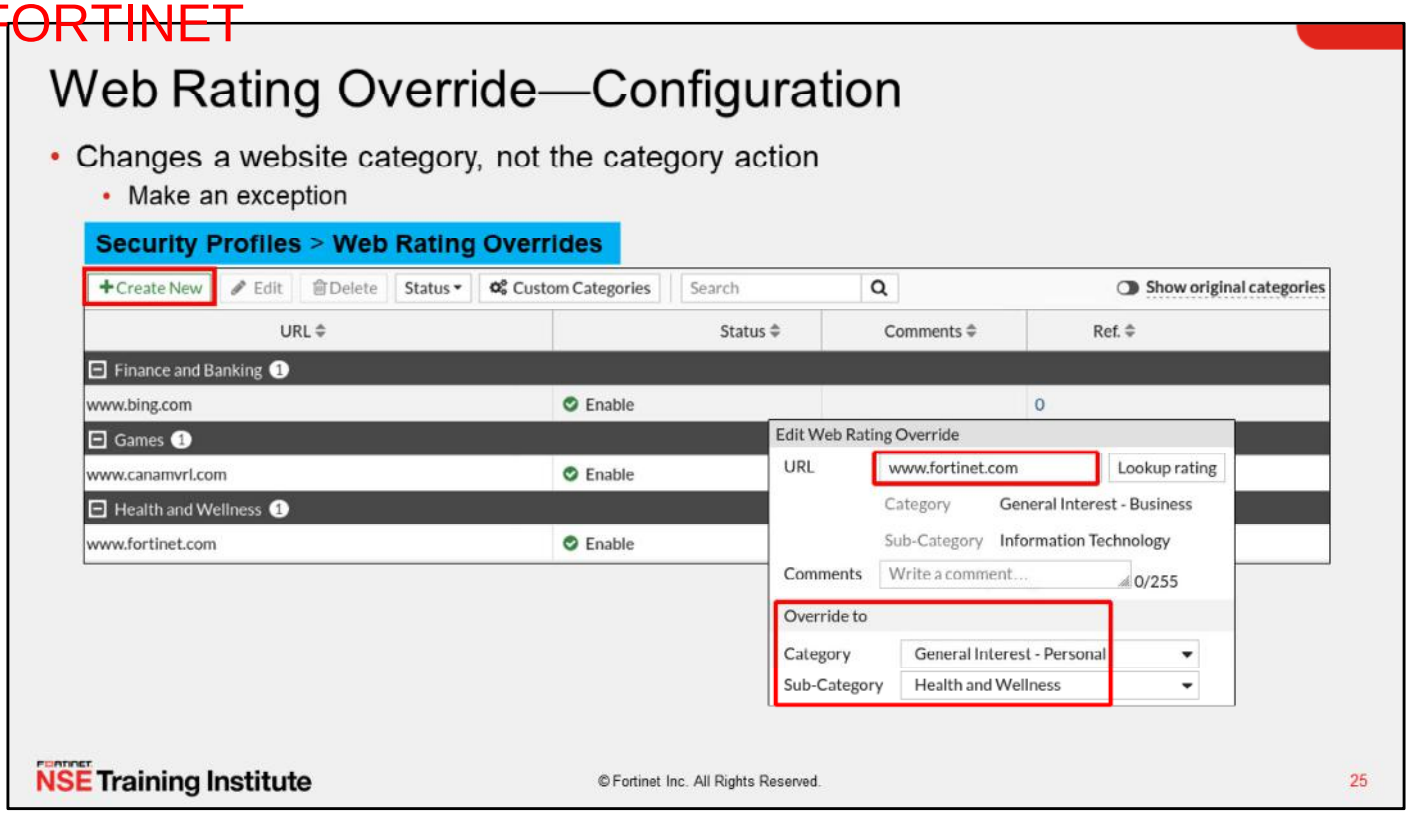

If you want to make an exception, for example, rather than unblock access to a potentially unwanted category, change the website to an allowed category. You can also do the reverse. You can block a website that belongs to an allowed category.

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.

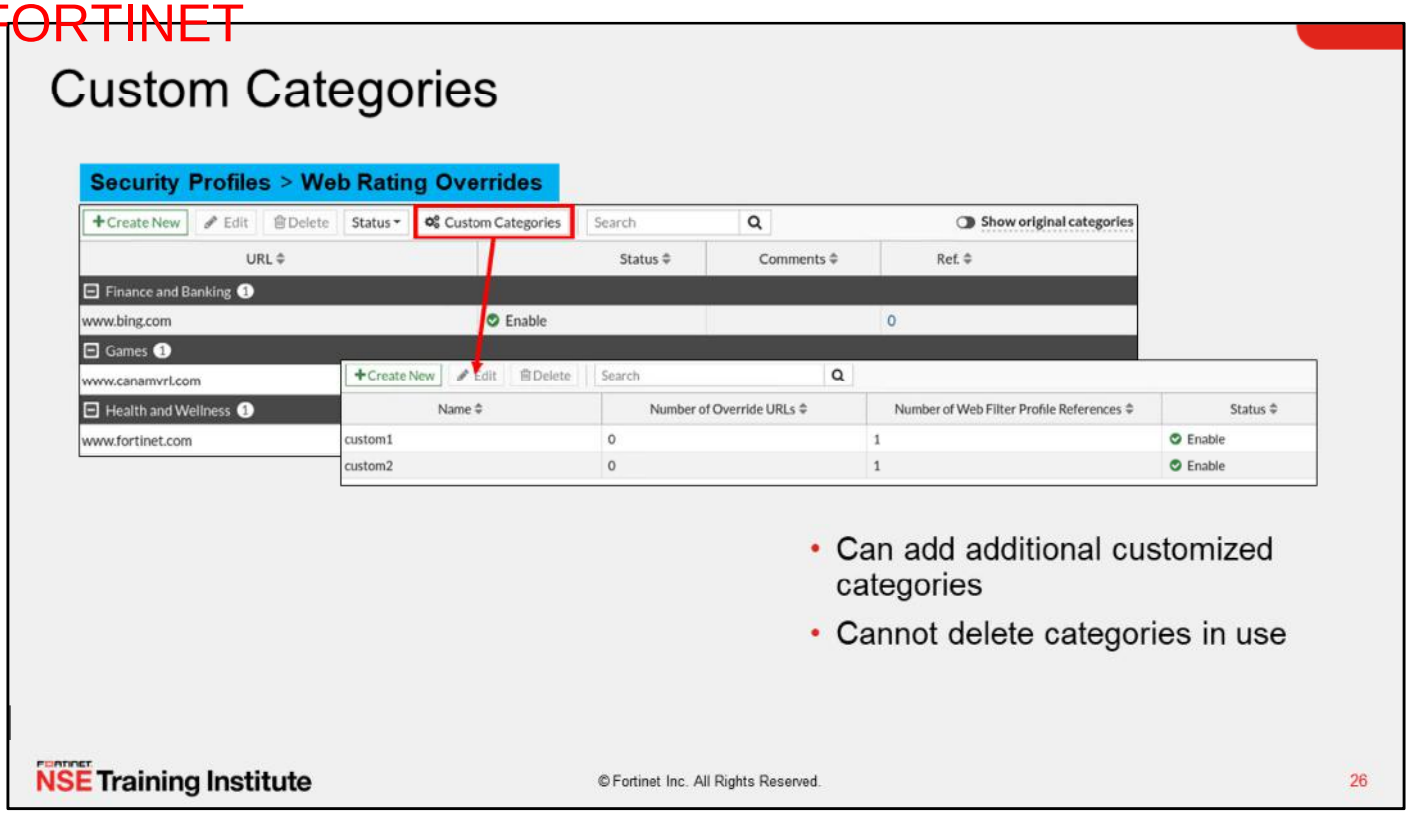

If the predefined categories in FortiGuard are not suitable for the situation, you can add additional custom categories.

You can add and delete custom categories as needed, as long as they are not in use.

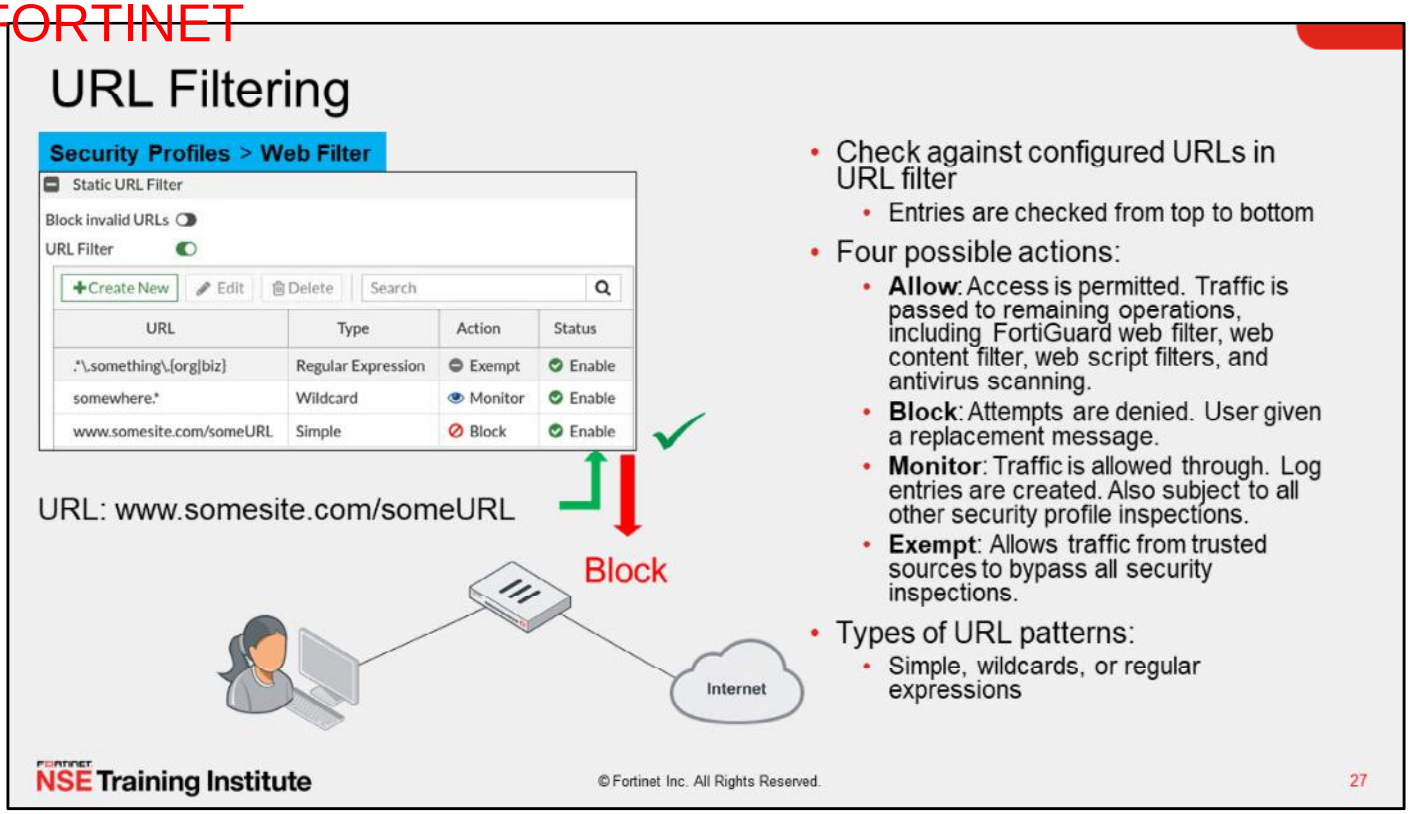

Static URL filtering is another web filter feature. Configured URLs in the URL filter are checked against the visited websites. If a match is found, the configured action is taken. URL filtering has the same patterns as static domain filtering: simple, regular expressions, and wildcard.

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.

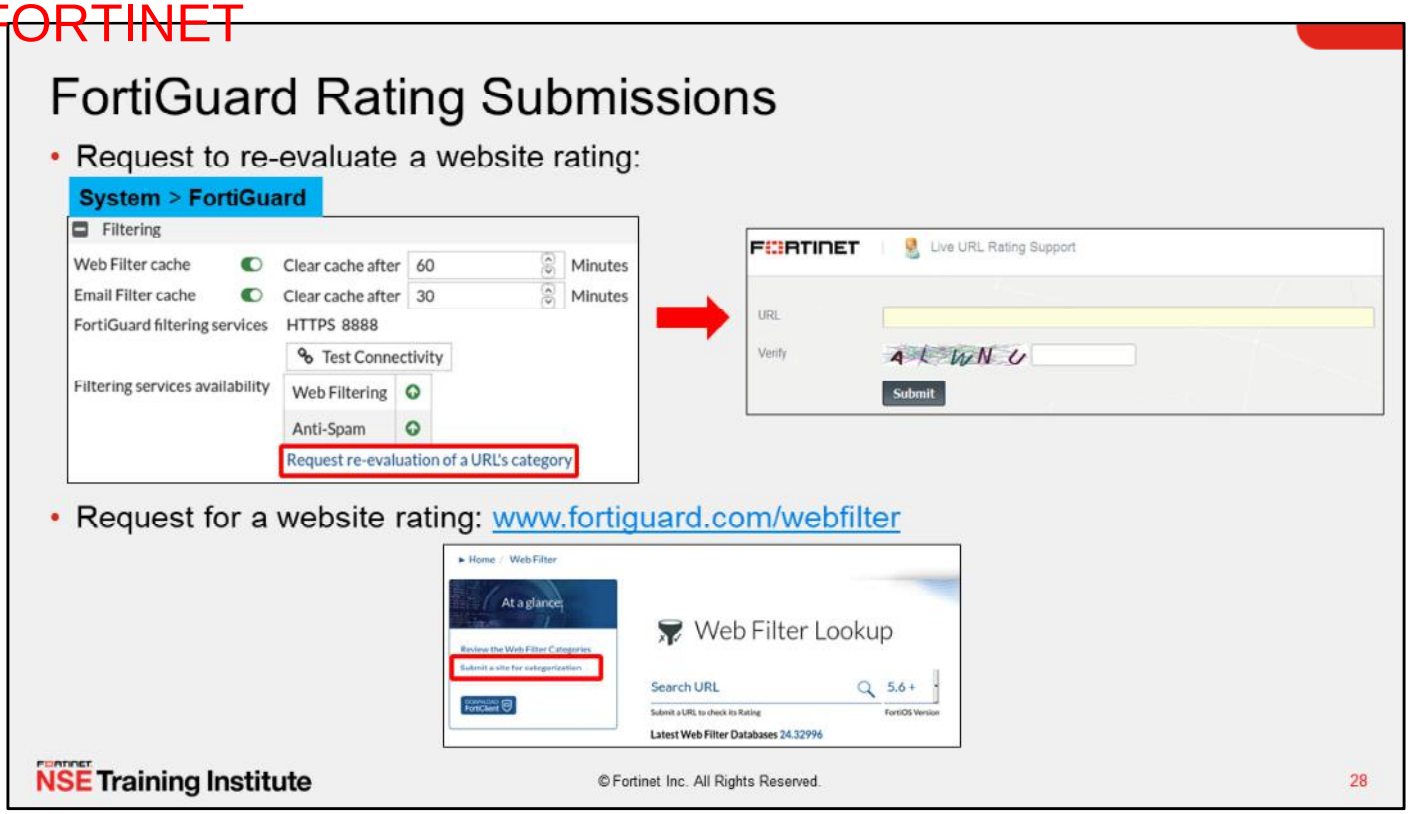

There is always the possibility for errors in ratings, or a scenario where you simply do not agree with the rating given. In that case, you can use the web portal to contact the FortiGuard team to submit a website for a new rating, or get it rated if it is not already in the database.

Knowledge Check

1. Which is a valid action for FortiGuard web category filtering?
   - ✓ A. Allow
   - B. Deny

2. Which is a valid action for static URL filtering?
   - ✓ A. Exempt
   - B. Warning

NSE Training Institute

29

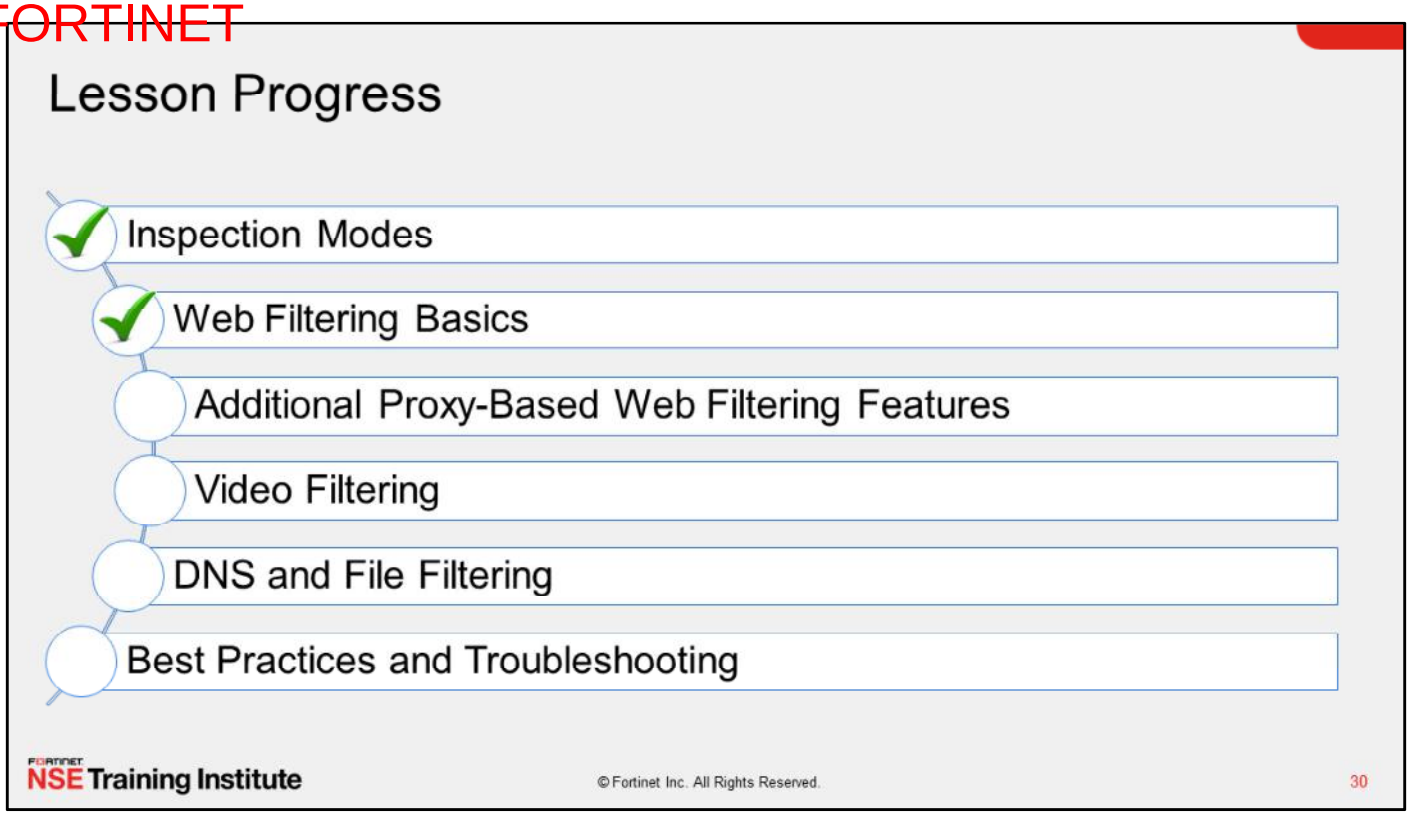

Good job! You now understand the basics of web filtering.

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

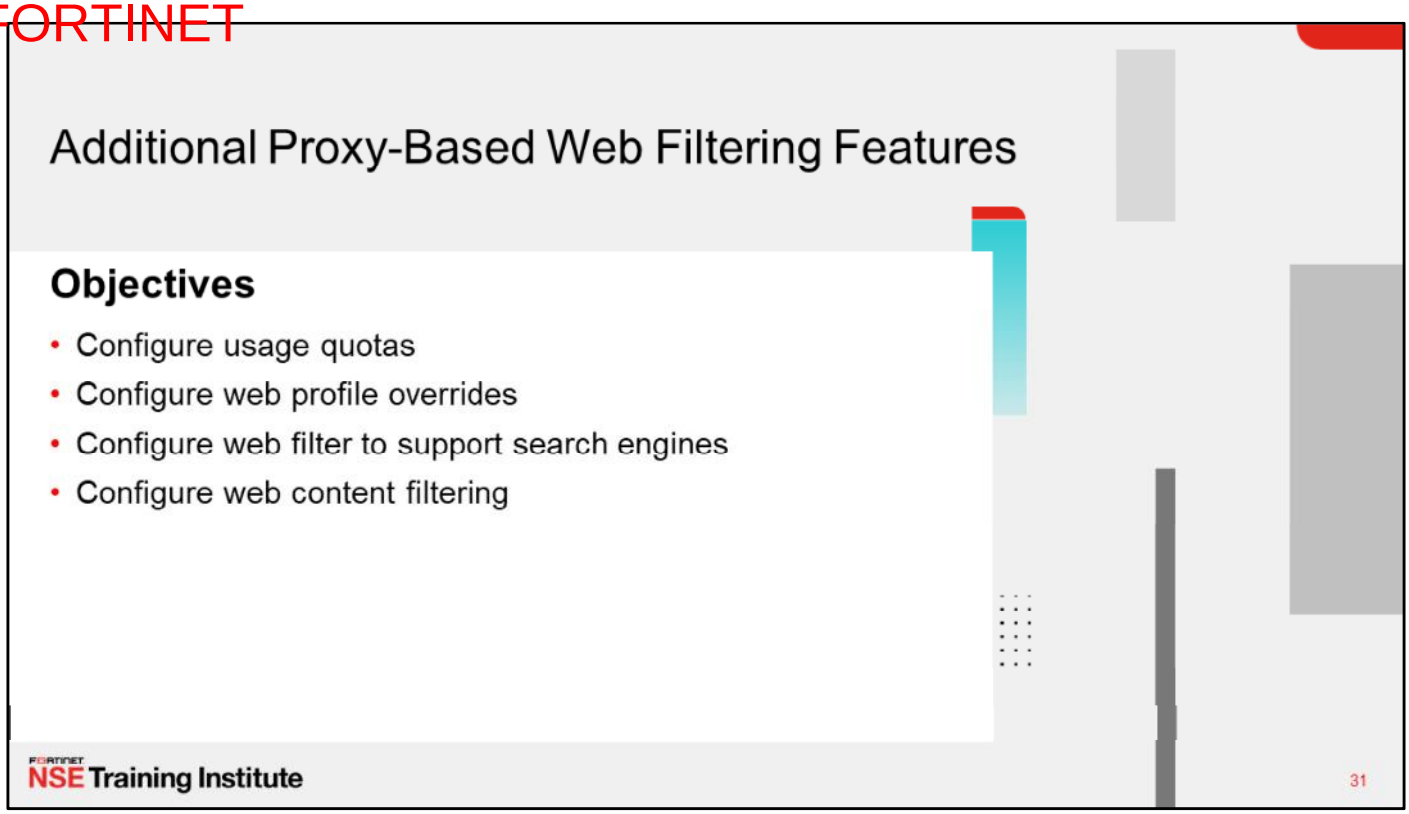

After completing this section, you should be able to achieve the objectives shown on this slide.

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

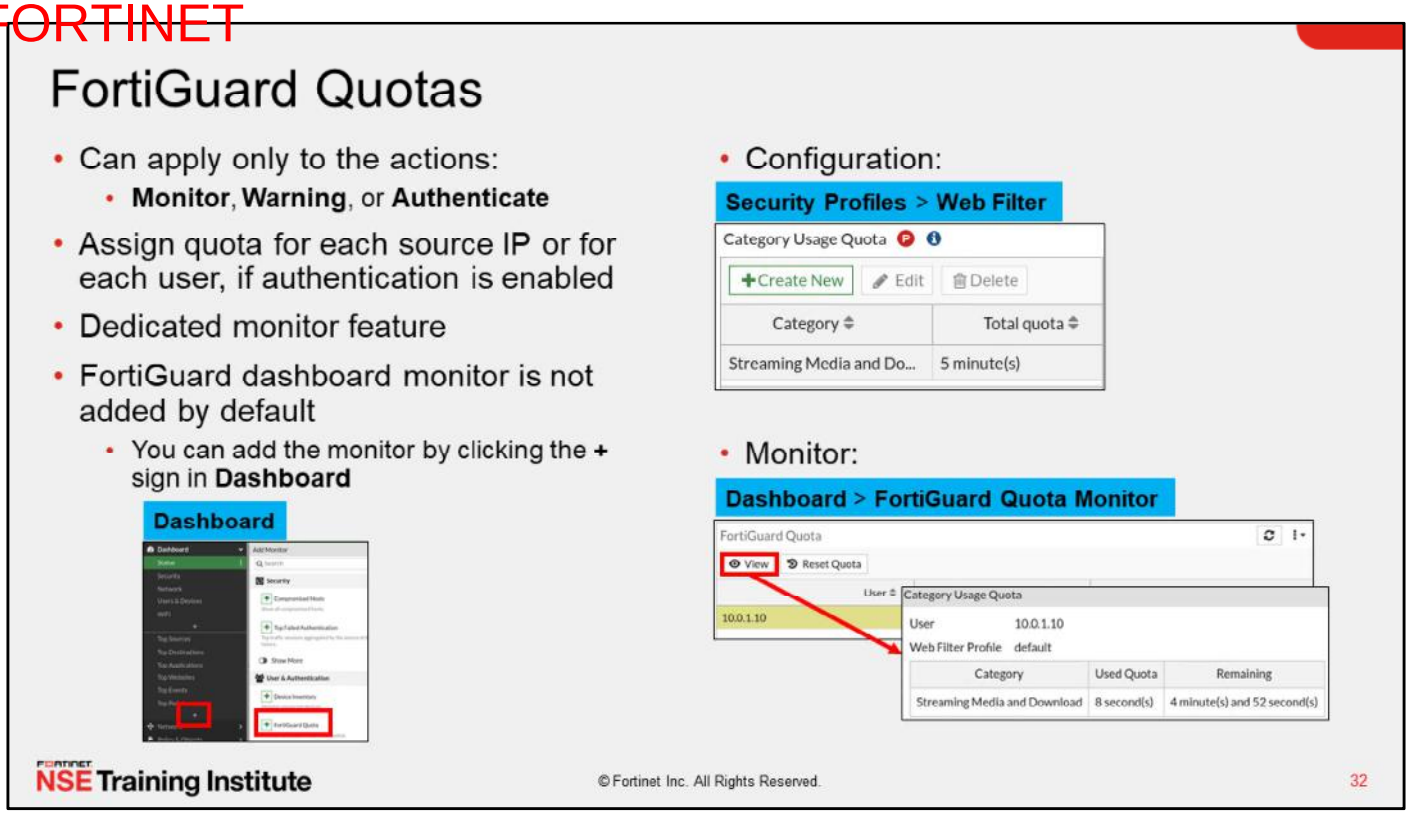

FortiGate also includes a feature to customize the quotas of time to access the categories that are set to monitor, warning, or authenticate in the web filter profile in proxy-based inspection mode.

You can customize multiple quotas (timers). Each quota can be applied to either a single category or multiple categories. If the quota applies to multiple categories, the timer is shared among all the categories instead of having a single timer for each individual category.

FortiGate automatically assigns quotas for each source IP, or each user if the authentication action is used, as the dashboard monitor shows. By default, the dashboard monitor is not added. You can click the **+** symbol to add **FortiGuard Quota** monitor to the **User & Authentication** section.

Now, take a look at how quotas work.

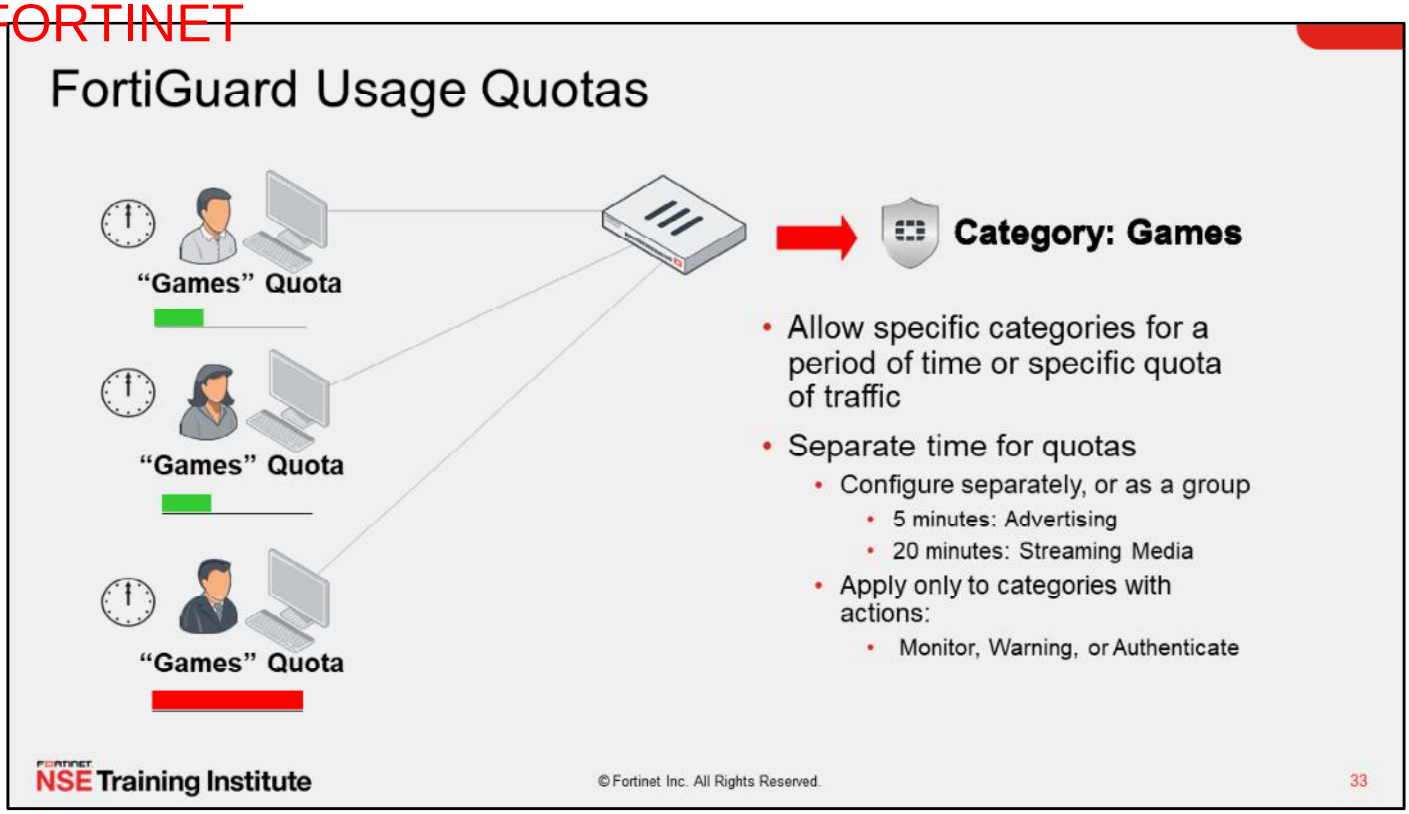

As shown on this slide, the FortiGuard quota limits the time users spend on websites, based on category. You can also set a quota, on the amount of traffic that can be allowed to a particular category.

A quota cannot redirect the user once the website is loaded in their browser. For example, if the user has 45 seconds left in their quota, and they access a website from the specified category, the selected website will likely finish loading before the remaining 45 seconds are up. Then, the user can stay on that website, and that website won't be blocked until the browser is refreshed. This scenario occurs because the connection to the website is not, usually, a live stream. After you receive the information, the connection is closed.

Note that the quota resets every 24 hours at midnight.

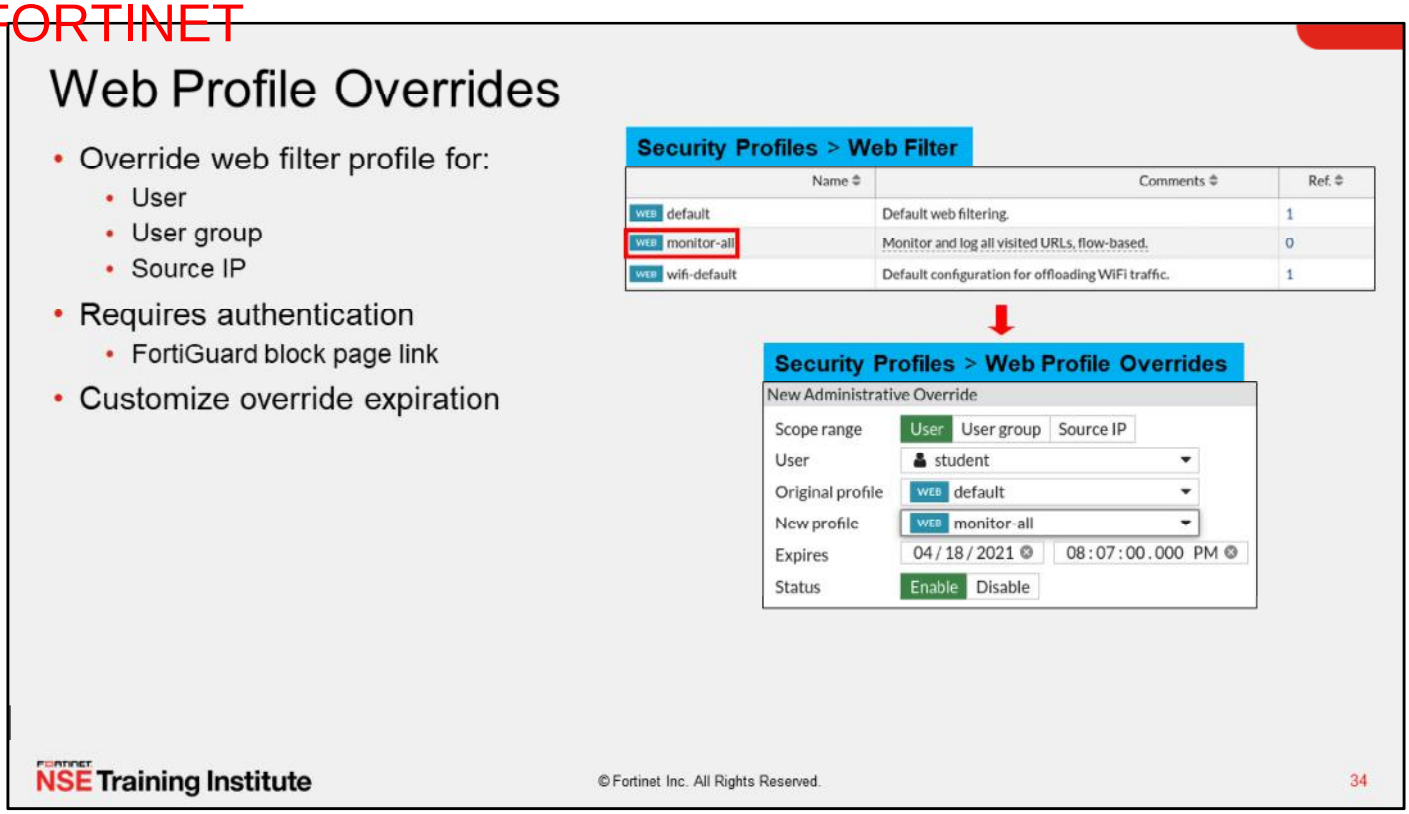

You can also override the filter profile. Web profile overrides change the rules that are used to inspect traffic. Overrides authorize specified users, user groups, or predefined source IPs, to use a different web filter profile to inspect their traffic.

In the example shown on this slide, the new profile applied to the user **student**, inspects all of that user's web traffic from the time that the new profile is applied, until the timer expires. To use this override, you must enable an override authentication. When you enable the web profile override, the FortiGuard block page shows a link you can select to activate the override.

Search Engine Filtering

- A proxy-based mode feature
- Requires FortiGate to use deep SSL inspection
  - Not supported when using certificate inspection
  - FortiGate requires full access to the application layer data
- Restricts websites or images from search results
  - Rewrites the search URL to enable safe search
    - For Google, Yahoo, Bing, and Yandex
- Logs all search keywords

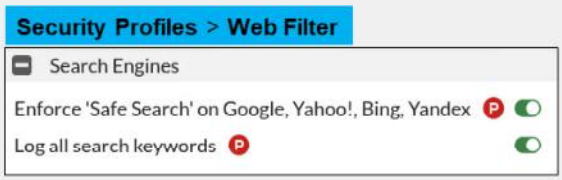


```
config webfilter profile
    edit "default"
        config web
            set safe-search url header
        end
    next
end
```

NSE Training Institute

35

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

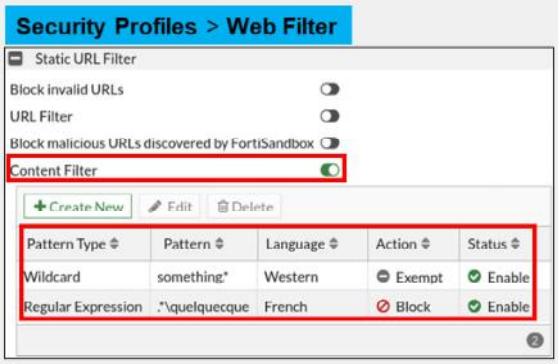

NSE Training Institute

36

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.

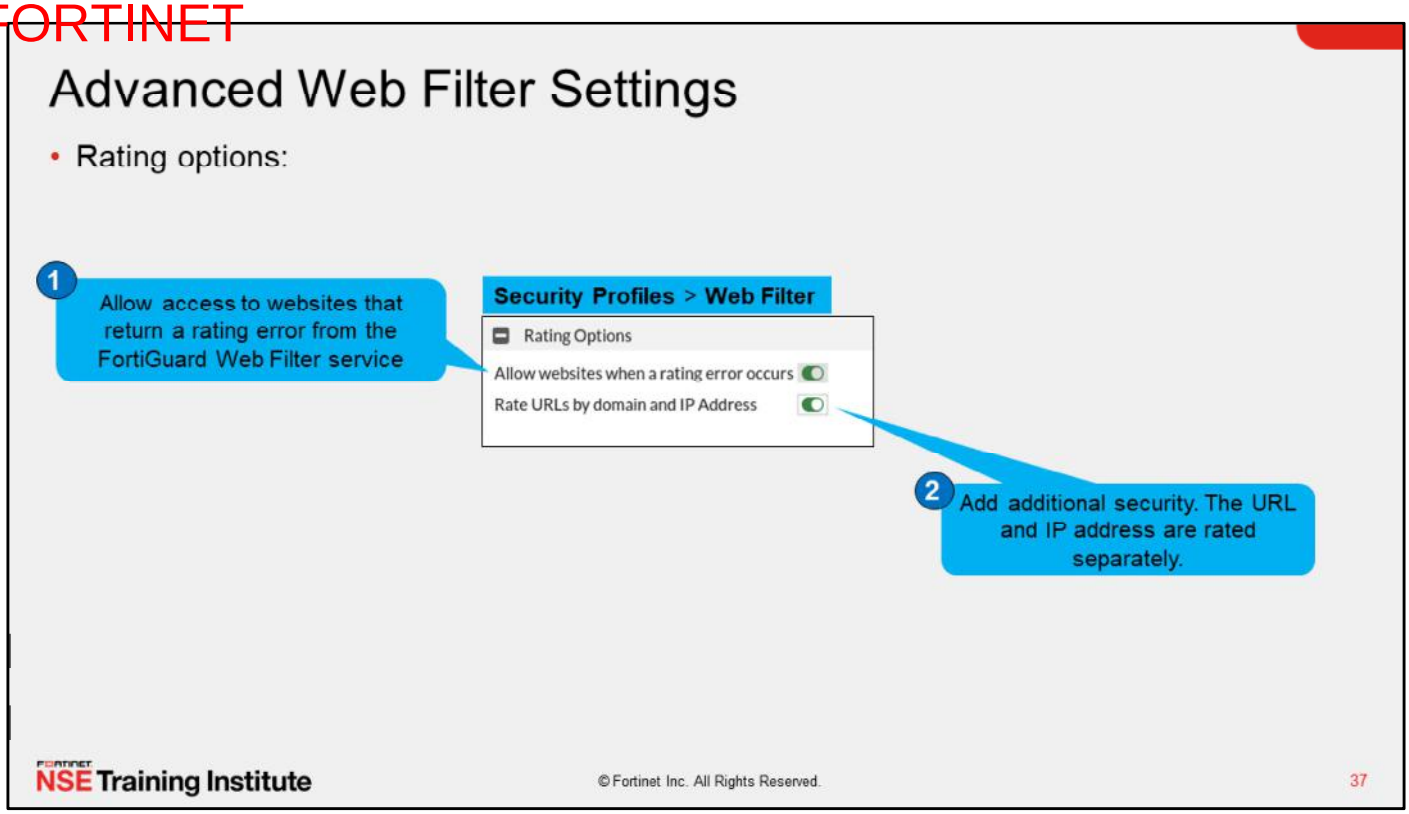

You can use advanced web filtering settings to improve the web filter.

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.

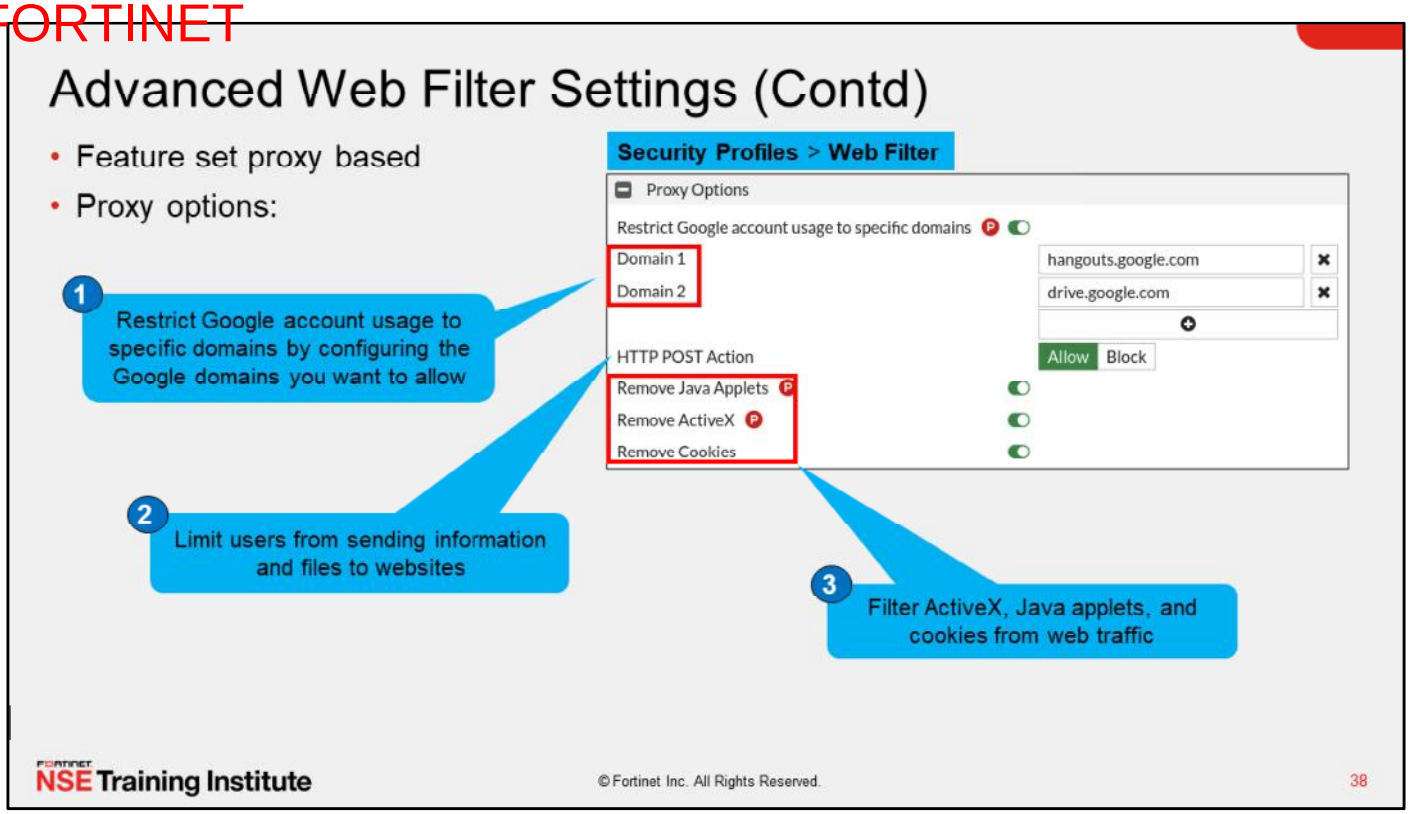

If you configure the web filter profile to use a proxy-based feature set, the advanced proxy option settings for web filtering are as follows:

1. Block access to some Google accounts and services. You can include an exception list.
2. HTTP POST is the command used by your browser when you send information, so you can limit the information and files to websites. The **Allow** option prevents a server timeout when scanning, or when other filtering processes are performed for outgoing traffic.
3. Filter cookies, Java applets, and ActiveX scripts from web traffic.

Knowledge Check

1. Which action can be used with the FortiGuard quota feature?
   - ✓ A. Monitor
   - B. Shape

2. Which statement about web profile overrides is true?
   - A. It is used to change the website category.
   - ✓ B. Configured users can activate this setting through an override link on the FortiGuard block page.

NSE Training Institute

39

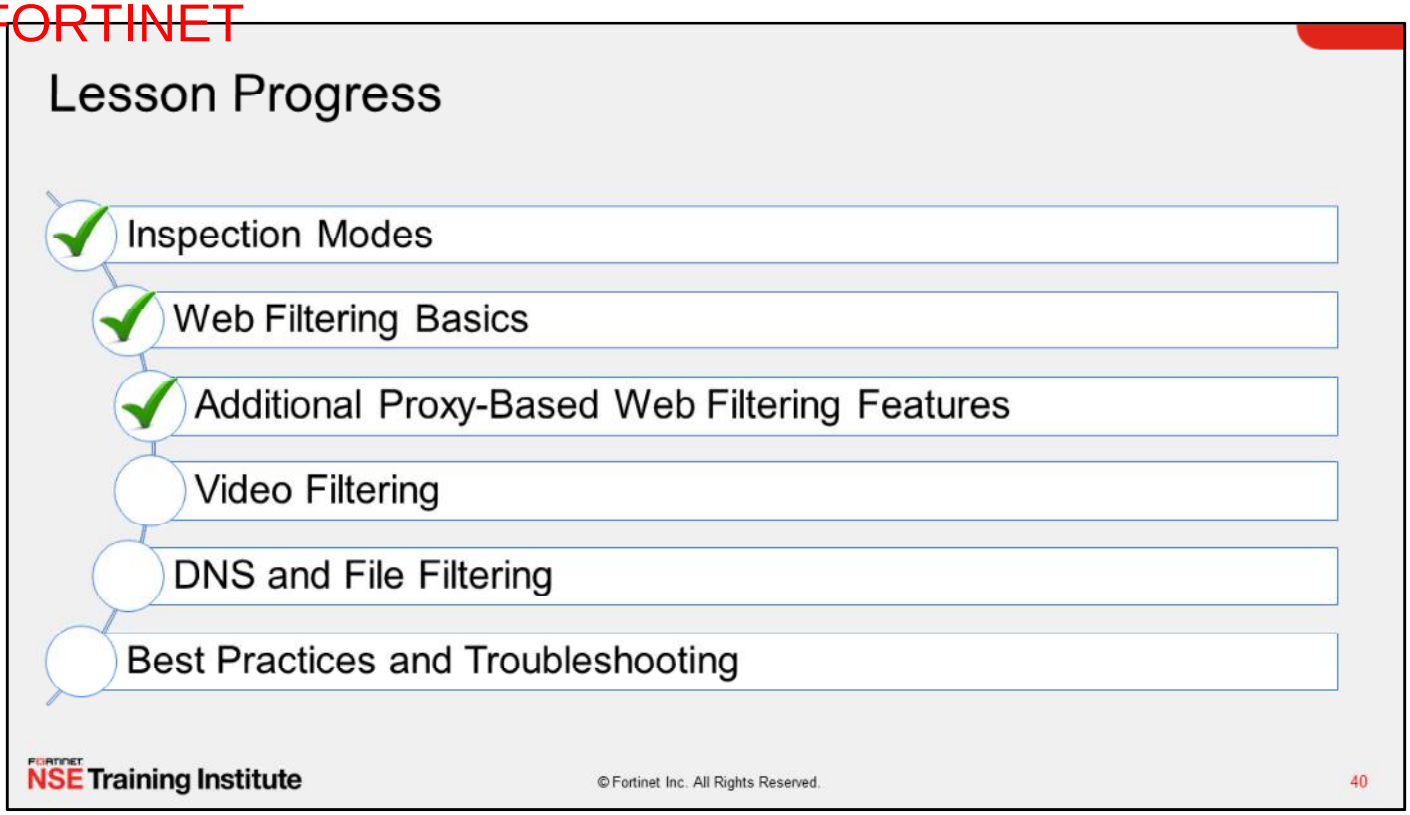

Good job! You now understand additional proxy-based web filtering features.

Now, you will learn about video filtering.

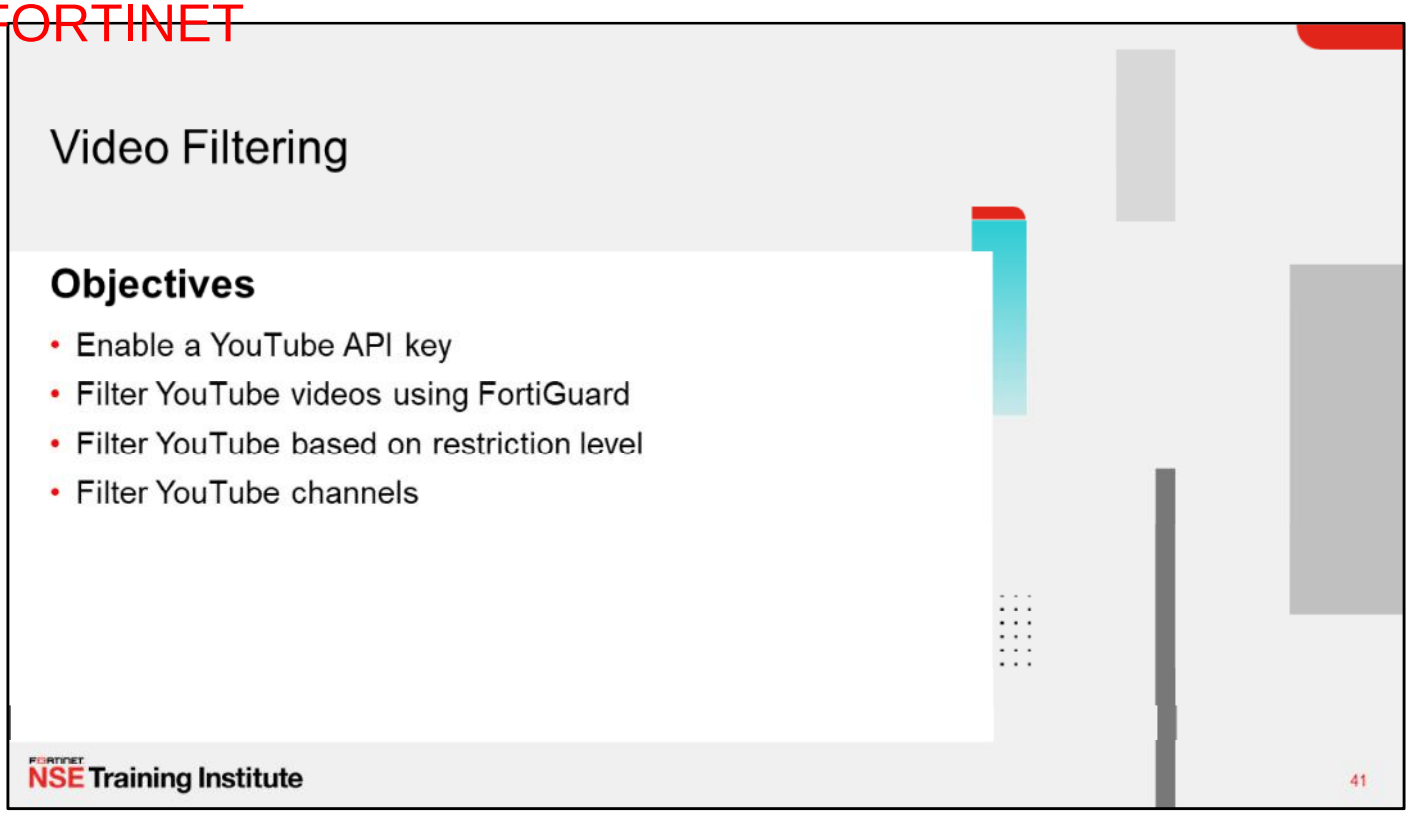

After completing this section, you should be able to achieve the objectives shown on this slide.

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

- Control 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
- Filter videos in two methods:
  - FortiGuard categories
  - Channel ID

NSE Training Institute

42

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 security FortiGuard 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.

YouTube API Key

- Access Google developer console
- Obtain YouTube API key
  - Create a new project
  - Enable YouTube data API v3
  - Create a credential
  - Copy API key
- Enable YouTube key on CLI
  - You can add multiple YouTube API keys

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

YouTube needs an active project created on a Google developer account

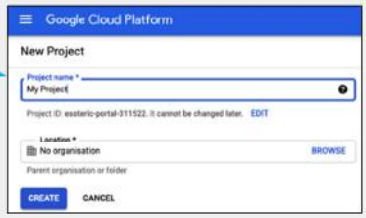

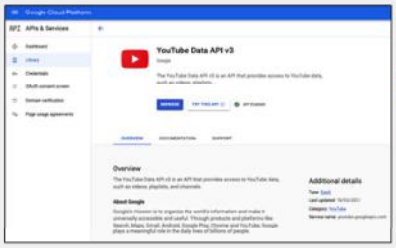

YouTube data API v3 is accessible through APIs and services library

NSE Training Institute

43

The YouTube API can help you to configure a video filter on FortiGate, and control access to content based on the content categories. You must have access to the Google developer console to obtain the API key. For more information about the Google developer console, visit *cloud.google.com*.

To obtain the YouTube API key:
1. Create a new project on the Google developer console.
2. Continue to fill in the project information by selecting your organization name and location.
3. In **APIs and services**, click **ENABLE APIS AND SERVICES** to add YouTube data API v3 from the library.
4. Enable the API and create a new credential to generate the API key.

Now the key is generated, you can add the key to FortiGate on CLI. Create a new object under `config videofilter youtube-key` and add the key using command `set key "youtube_api_key"`.

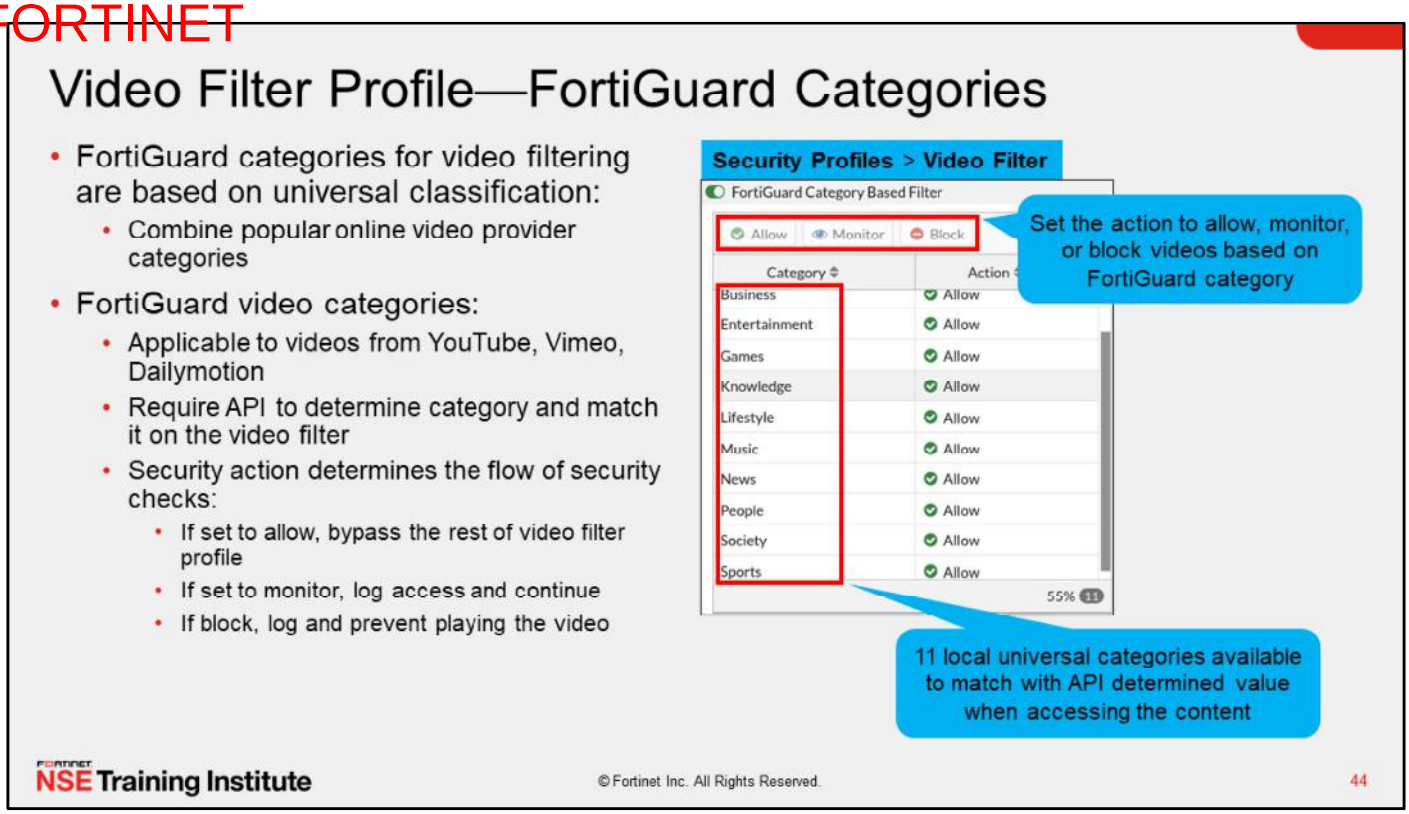

The video filter can identify videos using universal categories used by major online video content providers, such as YouTube. The generic classification combines multiple categories by these providers into one category. For example, the FortiGuate video category **Entertainment** includes YouTube categories,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

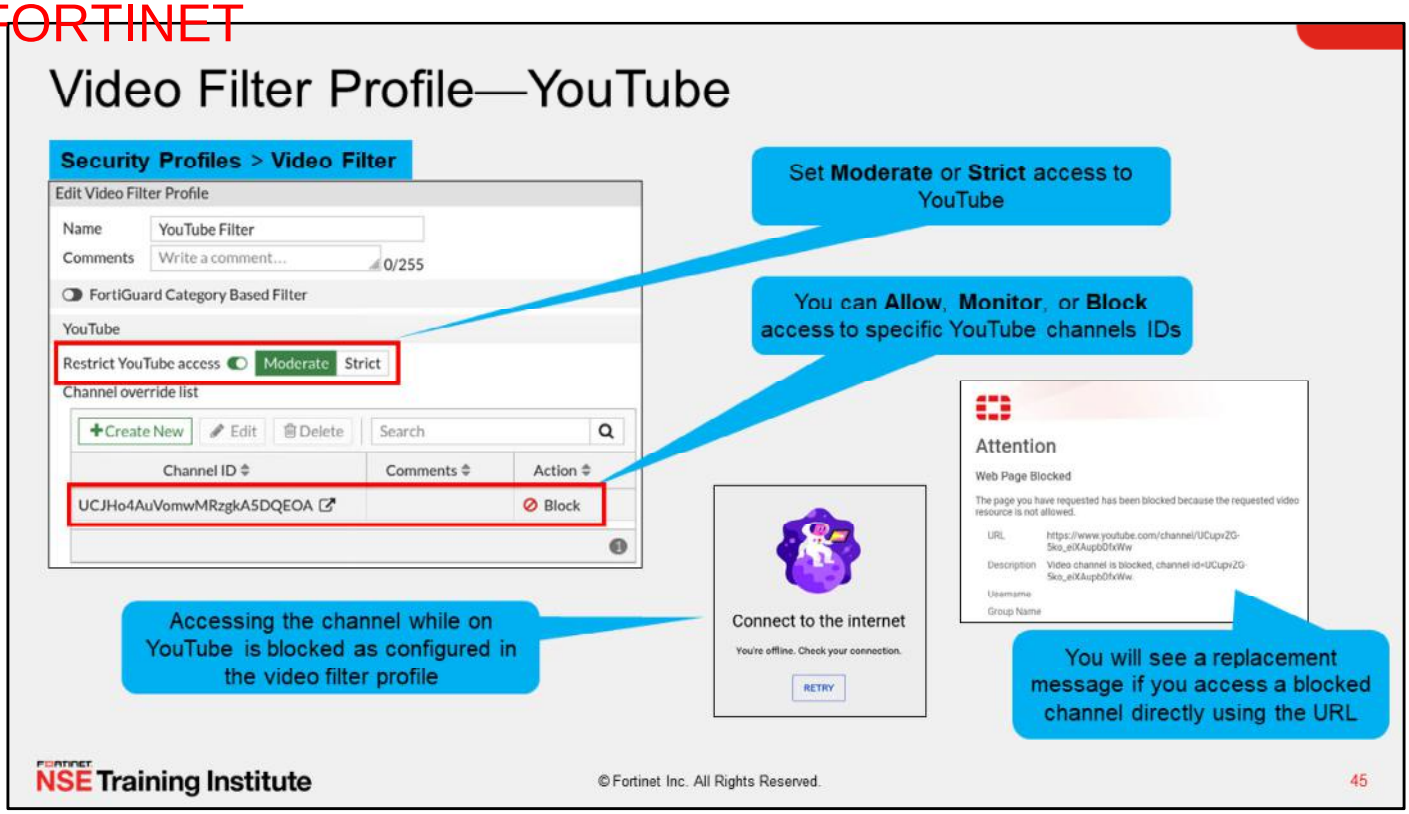

You can restrict YouTube access on a video filter by setting the restriction level to **Moderate** or **Strict**. When users access YouTube content using the firewall policy with the video filter profile applied, the users are given only content that is screened according to a filter applied by Google. 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 to configure YouTube video filtering?
   - ✓ A. YouTube API key
   - B. Username

2. Which action can be used with the video FortiGuard categories?
   - A. Authenticate
   - ✓ B. Monitor

NSE Training Institute

46

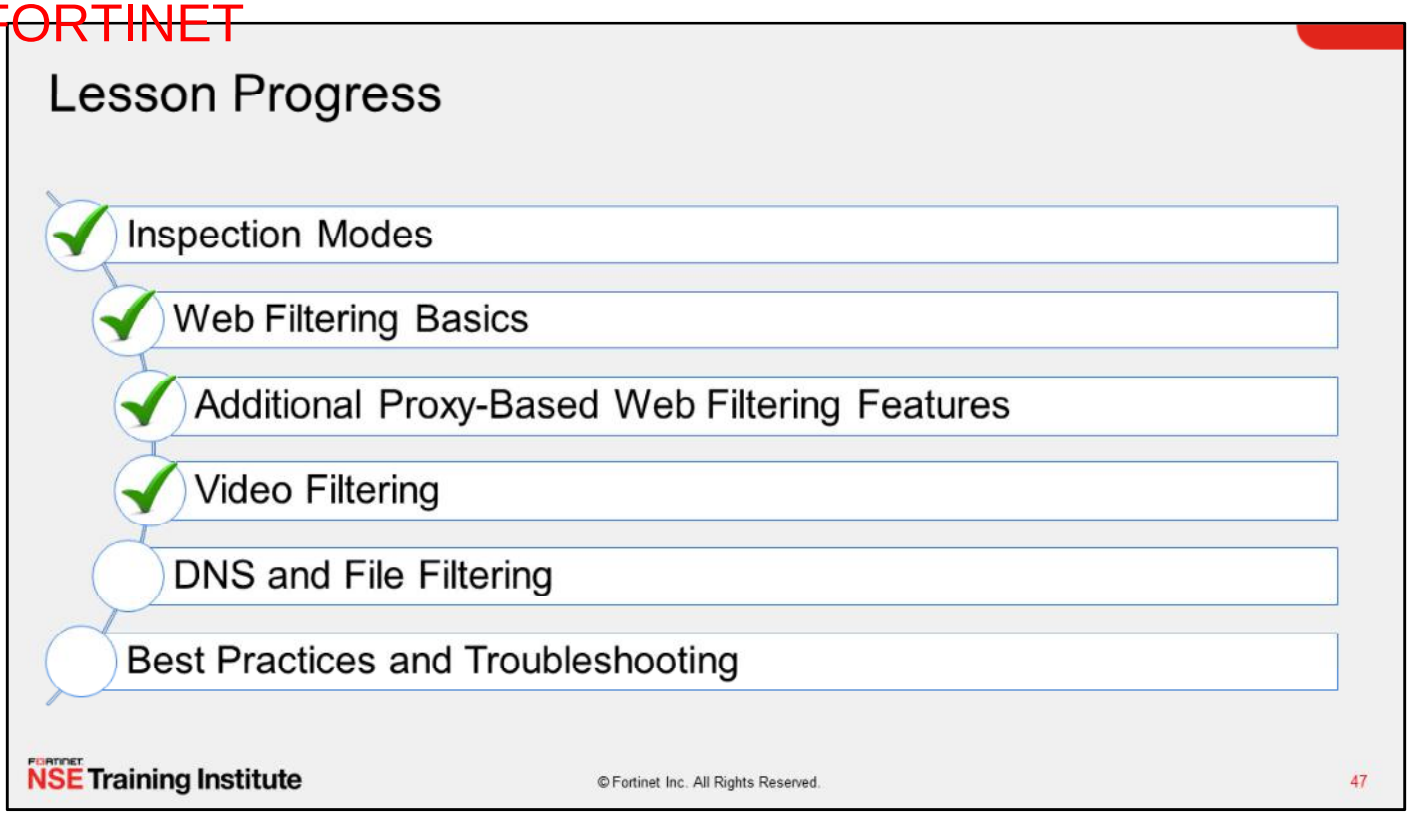

Good job! You now understand the video filtering feature.

Now, you will learn about DNS and file filtering.

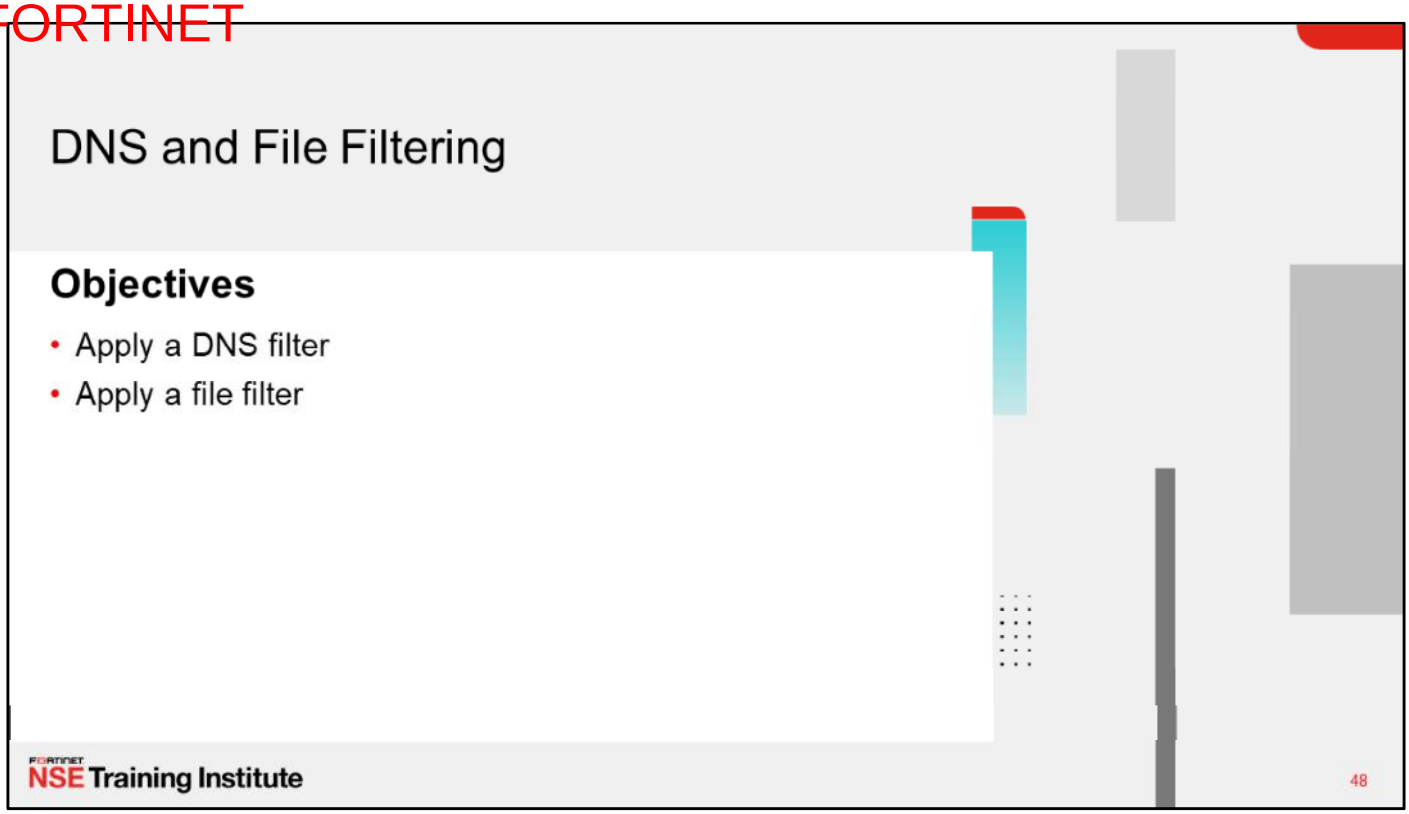

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in DNS and file filtering, you will be able to apply a DNS and file filters on FortiGate.

DNS-Based Web Filtering

- Uses FortiGuard SDNS ratings of DNS queries to decide access
- FortiGate must use FortiGuard SDNS service for DNS lookups
  - DNS queries redirected to FortiGuard SDNS server
- Lightweight
  - Lacks the precision of HTTP filtering
- SSL inspection available
  - DNS over TLS (DoT)
  - DNS over HTTPS (DoH)
- Cannot inspect a URL, only a host name
  - DNS resolves host name
- Supports URL filtering and FortiGuard category only

NSE Training Institute © Fortinet Inc. All Rights Reserved. 49

You can also inspect DNS traffic using the DNS filter security profile.

Rather than looking at the HTTP protocol, this option filters the DNS request that occurs before an `HTTP GET` request. This has the advantage of being very lightweight, but at a cost, because it lacks the precision of HTTP filtering.

Although DNS traffic is plain text, it can be encapsulated within a DNS over TLS (DoT) or DNS over HTTPS (DoH) payload. You can use the SSL inspection profile on the firewall policy to decrypt the payload and apply the DNS filter on the original DNS request.

Every protocol generates DNS requests in order to resolve a host name; therefore, this kind of filtering impacts all of the higher level protocols that depend on DNS, not just web traffic. For example, it could apply FortiGuard categories to DNS requests for FTP servers. Very few web filtering features are possible beyond host name filtering because of the amount of data available at the point of inspection.

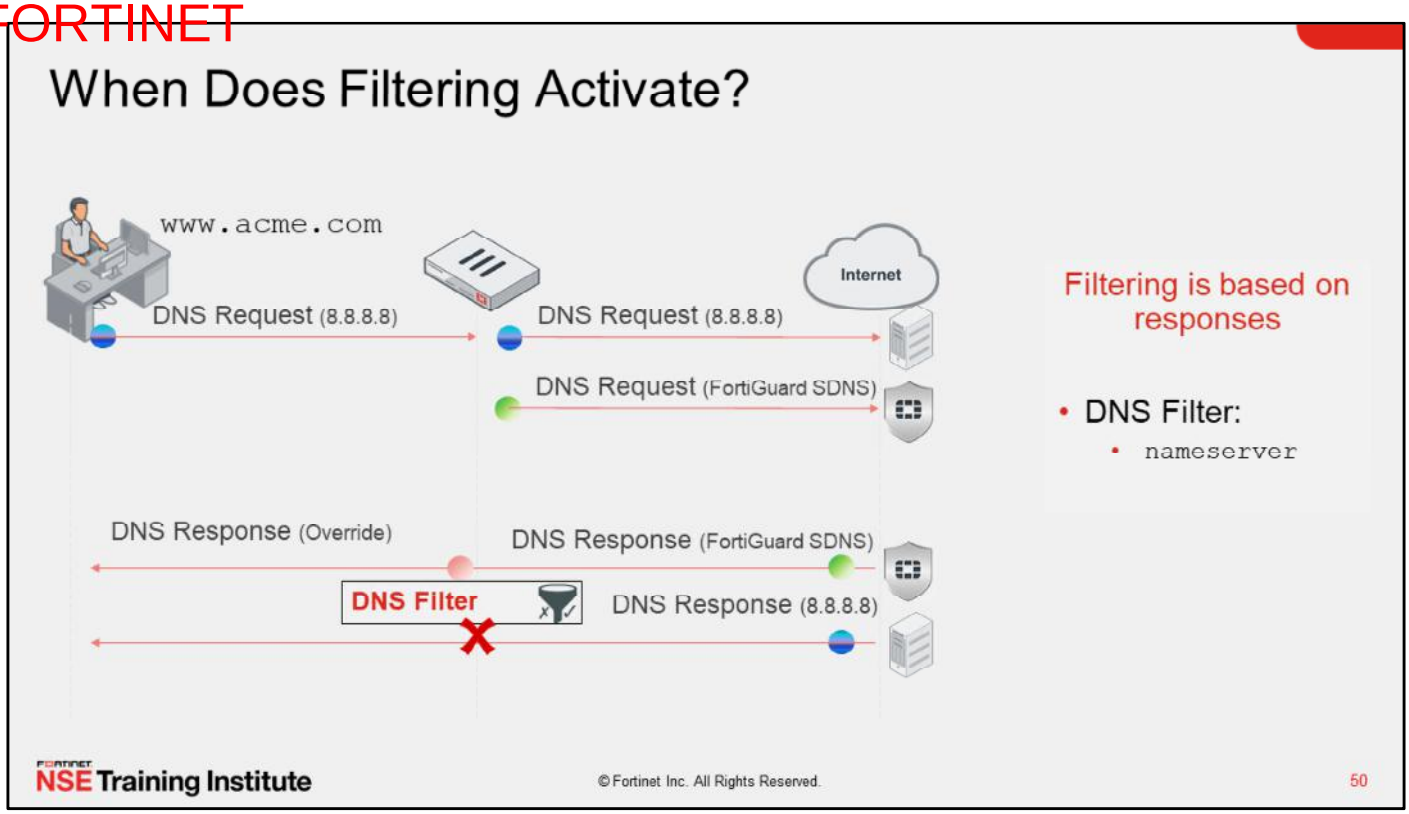

(slide contains animation)

This example illustrates filtering at the DNS lookup process.

DNS filtering looks at the `nameserver` response, which typically occurs when you connect to a website.

As discussed previously, in HTTP, the domain name and URL are separate pieces. The domain name might look like this in the header: `Host: www.acme.com`, and the URL might look like this in the header: `/index.php?login=true`.

When a device initiates a DNS lookup, it sends the FQDN information in the initial request. The DNS lookup occurs before the HTTP request can be sent.

When FortiGate receives the DNS request from the client, it sends a simultaneous request to the FortiGuard SDNS servers. With the FortiGuard SDNS service, there are two possible results:

1. Category is allowed: The original DNS response is passed and the remainder of the connection flow continues normally through to the HTTP 200 response. At this point, other web filters might be applied.
2. Category is blocked: FortiGate overrides the site IP address with the FortiGuard override address and presents a DNS error to the client.

As a result, if you are using a DNS filter, and the domain is blocked, the connection is blocked before the HTTP request is even sent.

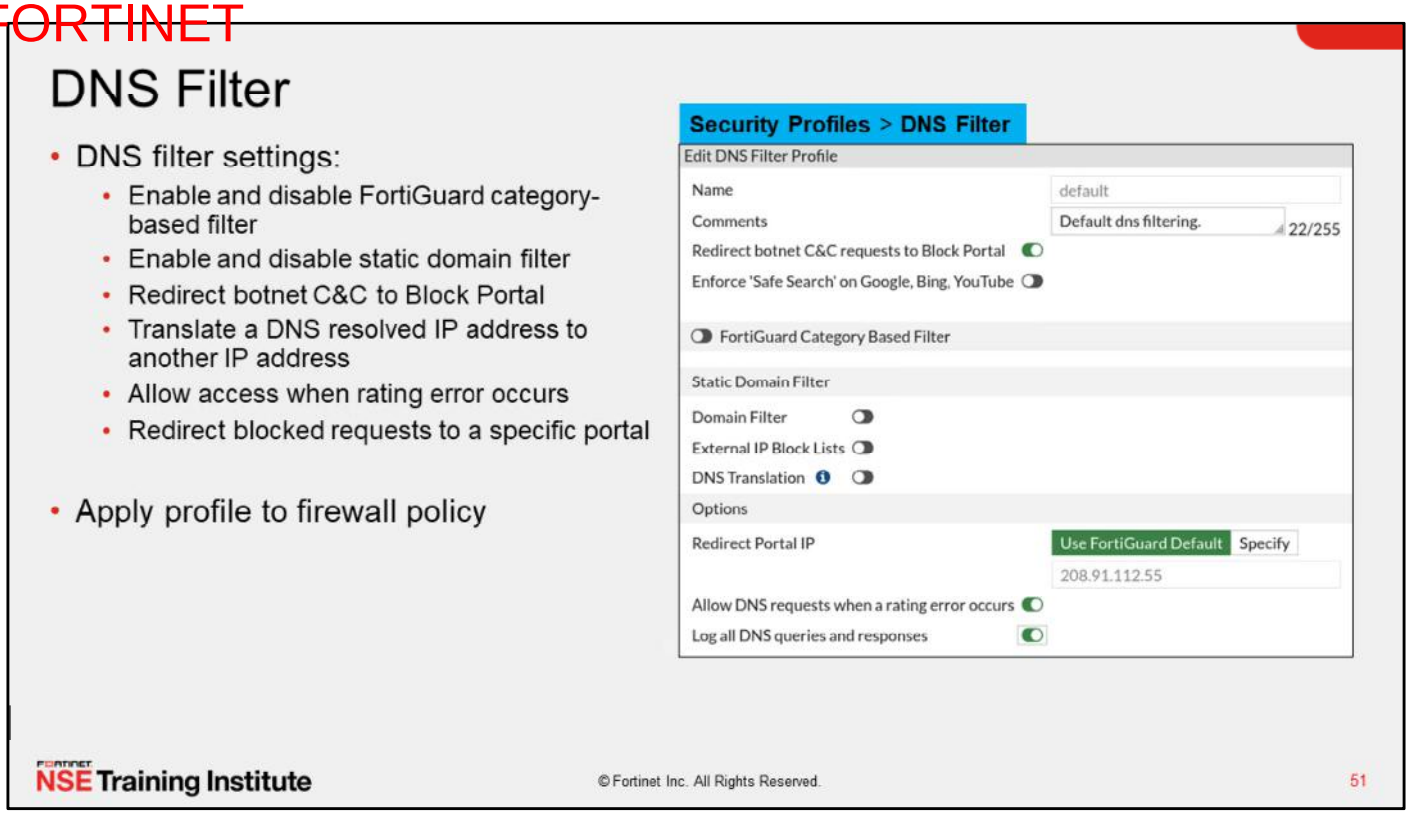

Here's a look at the DNS filter profile.

The DNS filter includes various configuration settings. You can enable or disable the FortiGuard category-based filter and the static domain filter. You also have the option to:

- Redirect botnet command and control requests to block portal
- Allow DNS requests when a rating error occurs from the FortiGuard web filter service
- Redirect blocked requests to a specific portal (the **Use FortiGuard Default** setting is recommended)

Using the **DNS Translation** feature, you can translate a DNS-resolved IP address to another IP address you specify on a per-policy basis.

After you enable and save the settings you require, remember to apply this profile to your firewall policy to activate the options. Any traffic being examined by the policy has those operations applied to it.

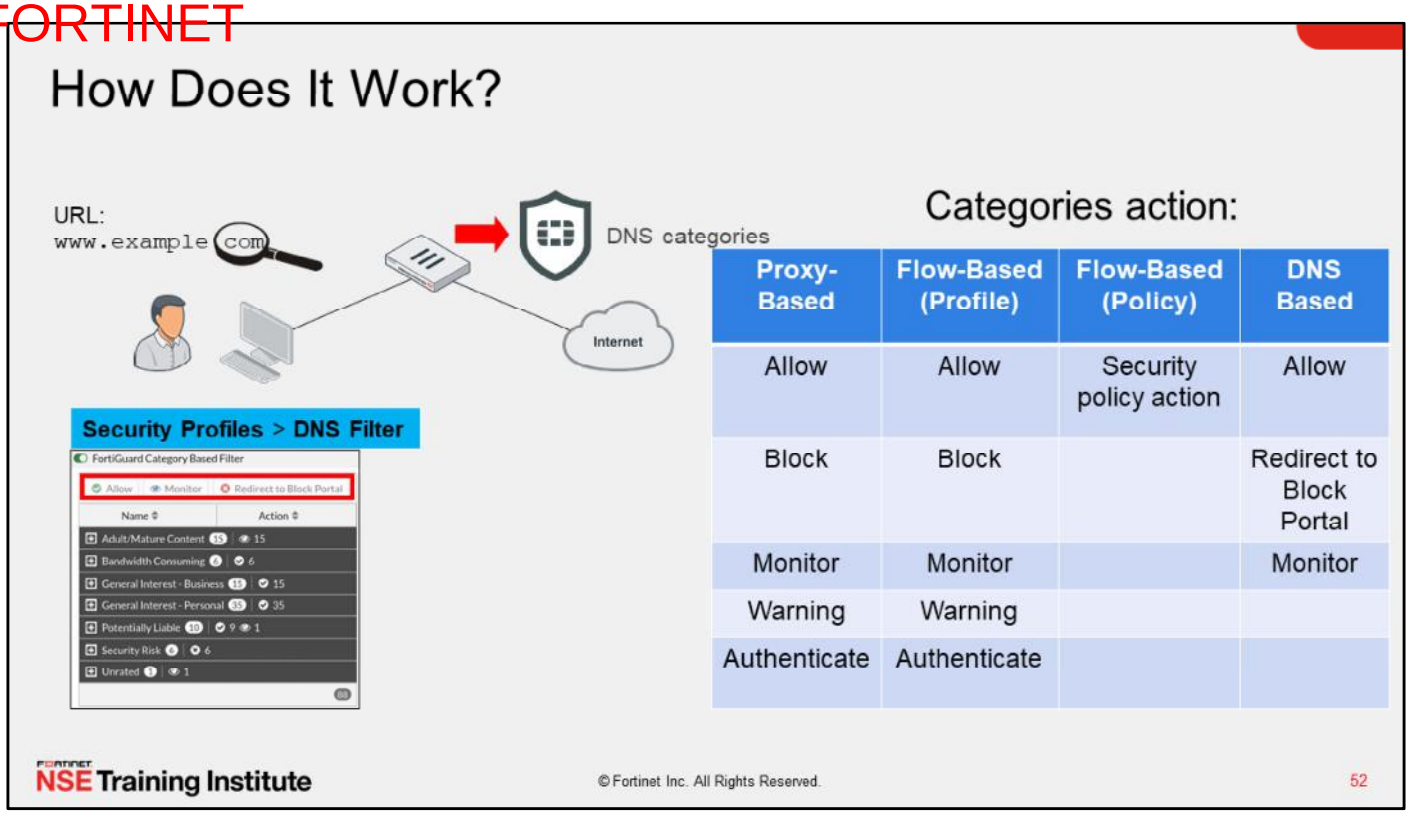

So, how does it work?

FortiGate queries the FortiGuard SDNS servers—or FortiManager, if it has been configured to act as a local FortiGuard server—to identify the category of a requested website.

When users visit websites, FortiGate uses the FortiGuard SDNS service to identify the category that the URL belongs to and takes a configured action for that category, such as allow, or redirect to a block portal. Using this feature, you can perform bulk URL filtering, without individually defining each website.

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

DNS filtering supports the following actions:
- Allow
- Redirect to block portal
- Monitor

Static Domain Filter

- Inspects DNS requests
- Actions to DNS requests
  - Redirect to Block Portal, Allow, and Monitor
- Patterns
  - Simple, wildcards, and regex

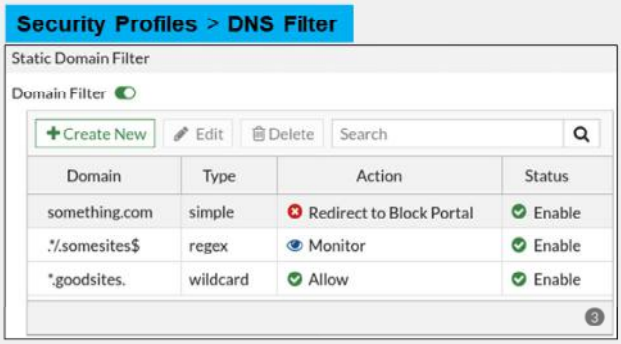

NSE Training Institute

53

You can configure DNS filtering to allow, redirect to a block portal, or monitor access to websites through the static domain filter. Entries in the domain list are checked against the DNS requests. If a match is found, the configured action is taken.

Patterns set to **Simple** are exact text matches. Patterns set to **Wildcard** allow for some flexibility in the text pattern by allowing wildcard characters and partial matching to occur. Patterns set to **Reg. Expression** allow for the use of PCRE regular expressions.

With this feature, you can prevent many HTTP requests from ever being made, because the initial lookup fails.

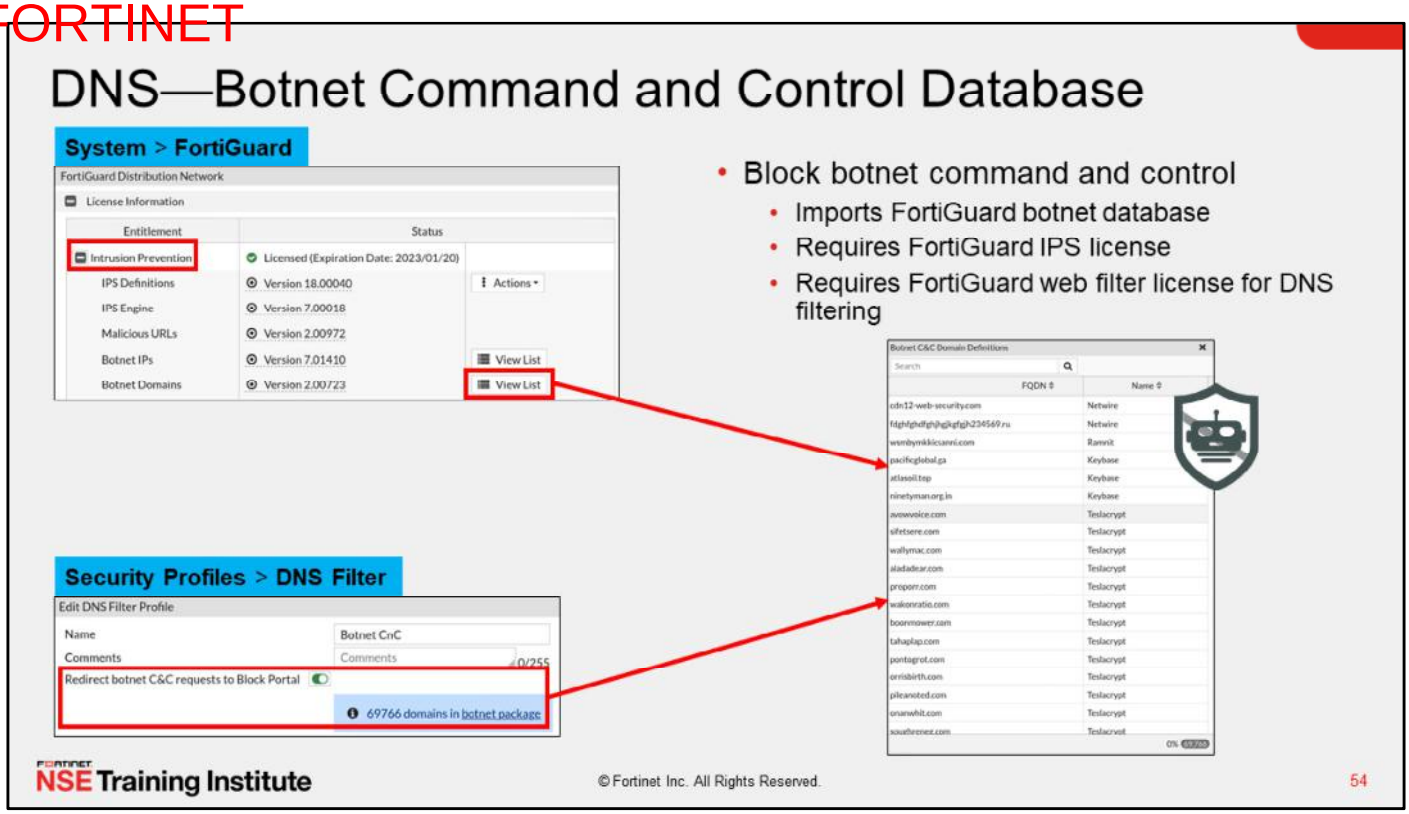

When you enable **Block DNS requests to known botnet C&C** in your DNS filter profile, DNS lookups are checked against the botnet command and control database. This database is dynamically updated from FortiGuard and stored on FortiGate.

All matching DNS lookups are blocked. Matching uses a reverse prefix match, so all subdomains are also blocked.

This service requires an active IPS and web filtering license.

File Filter

- File filter profile to inspect files over HTTP and other protocols:
  - Support most of the protocols in flow-based
  - Proxy-based covers MAPI and SSH
- Create rules to block or monitor files based on file type
- Choose which direction of traffic
  - Incoming
  - Outgoing
  - Both
- Specify files types in one rule or more—depends on traffic direction and security action needed

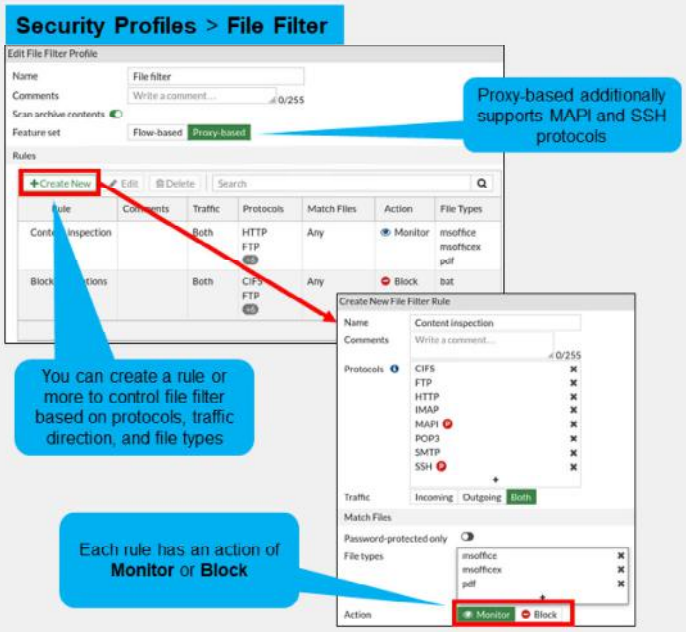

NSE Training Institute

55

The file filter profile can provide file filtering when enabled on firewall policies. It inspects files traversing over HTTP and other protocols, such as FTP and SSH.

You can select supported file types, such as compressed files and javascript content, when creating a new file filter rule. The rule can also specify whether to log or block files entering or leaving the network, or both.

Knowledge Check

1. Which statement about blocking the known botnet command and control domains is true?
   - ✓ A. DNS lookups are checked against the botnet command and control database.
   - B. The botnet command and control domains can be enabled on the web filter profile.

2. Which security profile inspects only the fully qualified domain name?
   - A. Web Filter
   - ✓ B. DNS Filter

NSE Training Institute

56

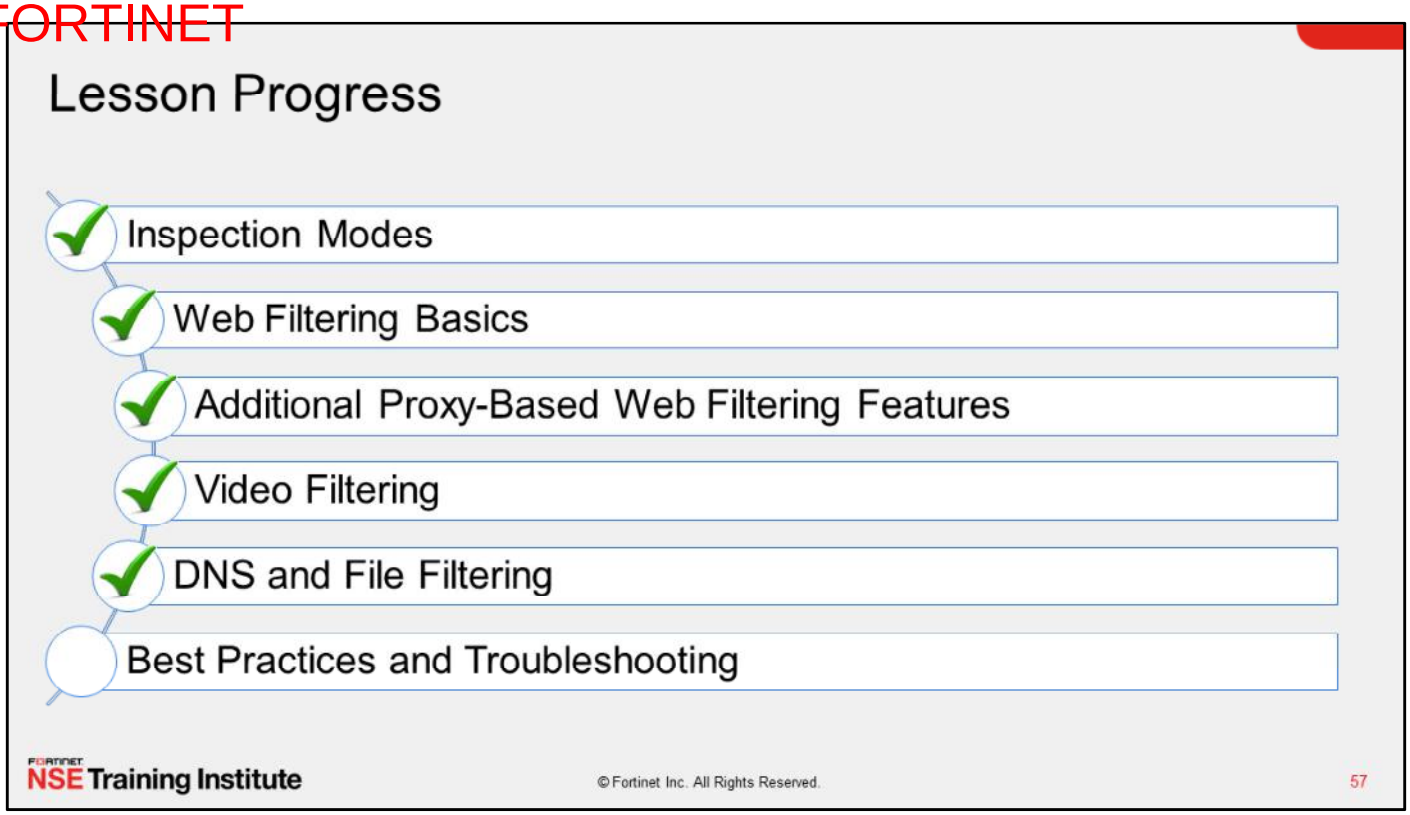

Good job! You now understand DNS and file filtering.

Now, you will learn about best practices and troubleshooting.

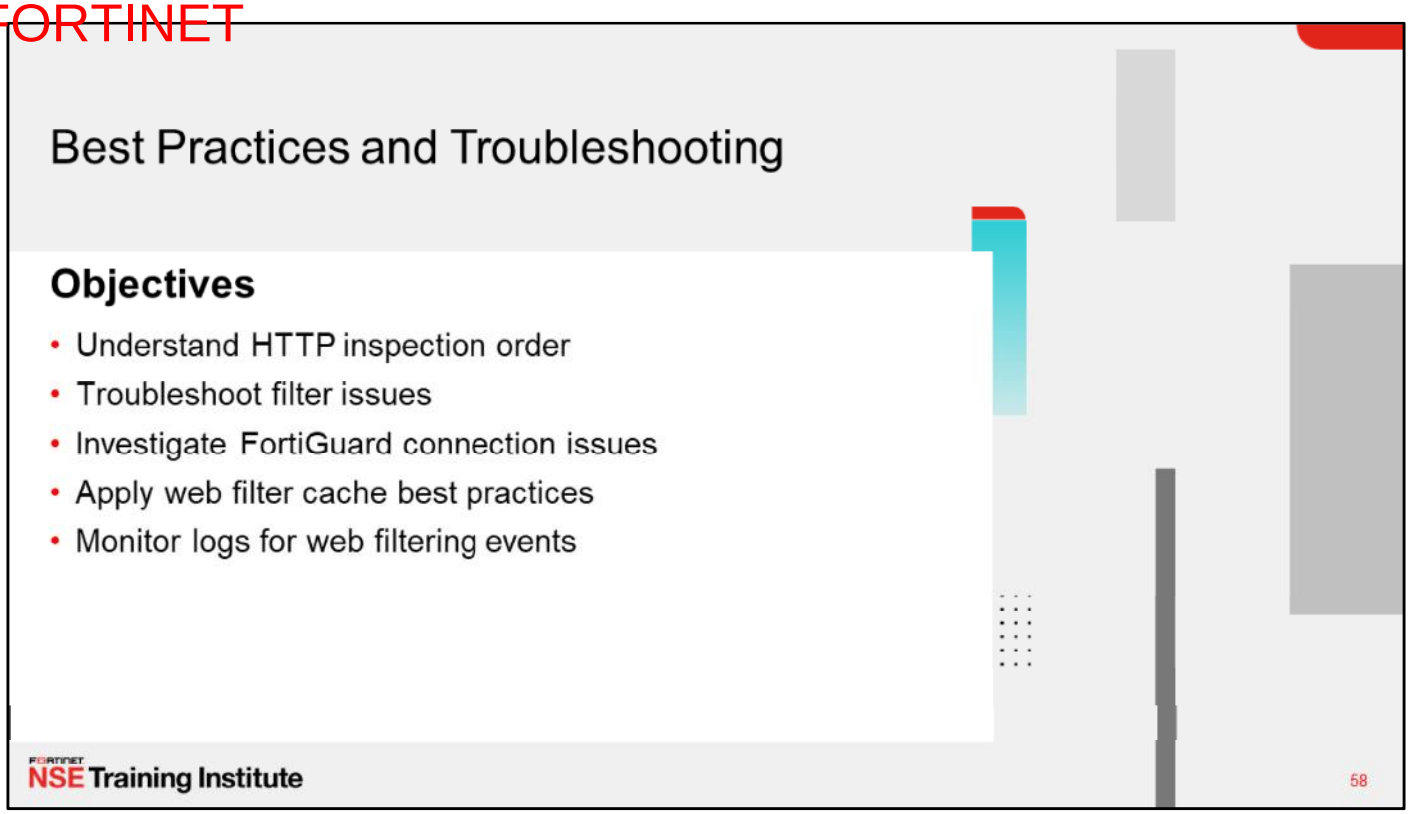

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

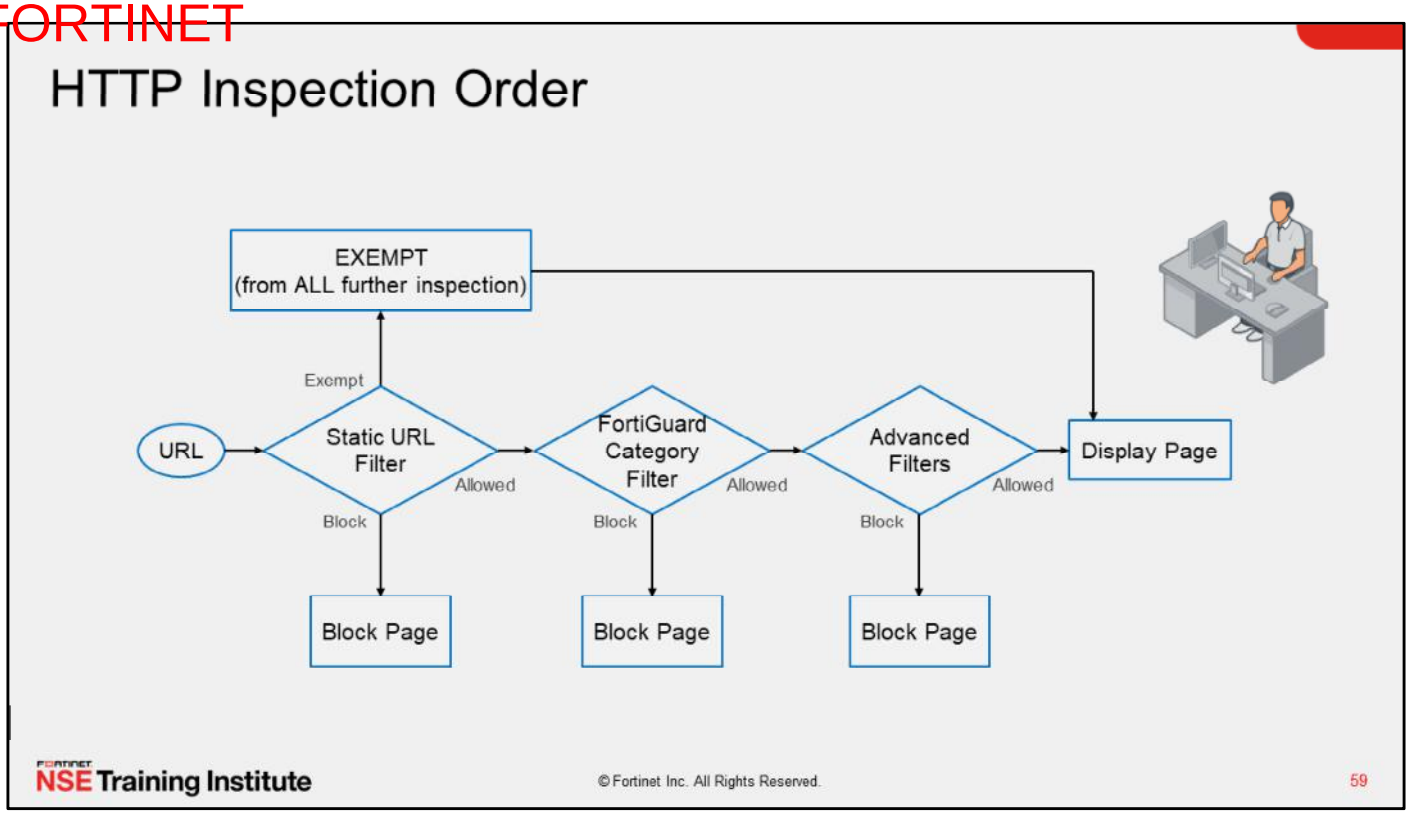

Remember that the web filtering profile has several features. So, if you have enabled many of them, the inspection order flows as follows:

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?
  - Is FortiGuard SDNS service accessible for DNS filters?

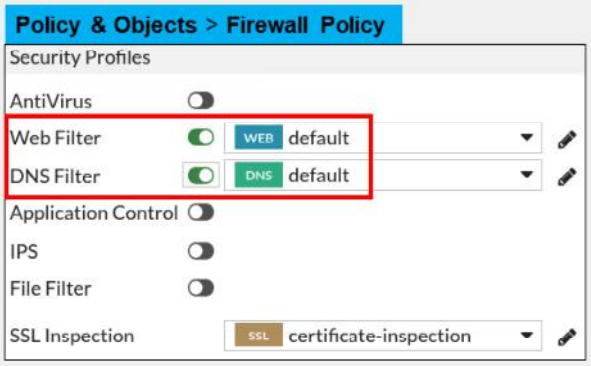


```
config firewall policy
 edit 1
  set dnsfilter-profile <profile>
  set webfilter-profile <profile>
 next
end

config firewall profile-group
 edit <group name>
  set dnsfilter-profile <profile>
  set webfilter-profile <profile>
 next
end
```

NSE Training Institute

60

You have configured your security profiles, but they are not performing web or DNS 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.

Additionally, to use the FortiGate DNS filter, you must use the FortiGuard SDNS service for DNS lookups. DNS lookup requests sent to the FortiGuard SDNS service return with an IP address and a domain rating that includes the FortiGuard category of the website. As a result, for this mechanism to work, the FortiGuard SDNS service must be reachable by FortiGate.

FortiGuard Connection

- FortiGuard category filtering requires a live connection
- Weight Calculation: default = (difference in time zone) x 10
  - Goes down over time (never below default)
  - Goes up if FortiGuard requests are lost

```
FortiGate-VM64 # diagnose debug rating
Locale       : english

Service      : Web-filter
Status       : Enable
License      : Contract
\
Num. of servers : 1
Protocol        : https
Port            : 443
Anycast         : Enable
Default servers : Included

-=- Server List (Wed Apr 21 13:59:43 2021) -=-

IP                                          Weight    RTT Flags   TZ   FortiGuard-requests  Curr Lost Total Lost            Updated Time
173.243.140.16                                 -72    101 DI       0                    36          0          0 Wed Apr 21 13:58:13 2021
```

NSE Training Institute

61

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)

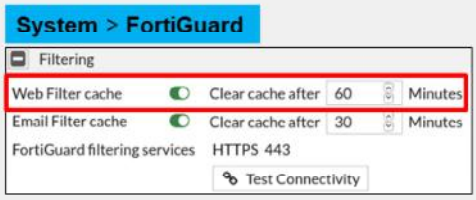


```
config system fortiguard
    set fortiguard-anycast {enable|disable}
    set protocol {udp|https}
    set port {8888|53|443}
    set webfilter-timeout {<1> - <30>}
end
```

NSE Training Institute

62

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.

Web Filter Log

- Record HTTP traffic activity, such as:
  - Action, profile used, category, URL, quota info, and so on

Log & Report > Web Filter

Add Filter | Details

Date/Time	User	Source	Action	URL	Category Description	Sent / Rec
Minute ago		10.0.1.10	passthrough	https://bat.bing.com/	Search Engines and Portals	517 B / 0
Minute ago		10.0.1.10	passthrough	https://site.fortinet.com/	Information Technology	517 B / 0
Minute ago		10.0.1.10	passthrough	https://site.fortinet.com/	Information Technology	517 B / 0

```
date=2021-04-22 time=09:32:04 eventtime=1619109124643229175 tz="-0700"
logid="0317013312" type="utm" subtype="webfilter" eventtype="ftgd_allow"
level="notice" vd="root" policyid=1 sessionid=9505 srcip=10.0.1.10 srcport=57734
srcintf="port3" srcintfrole="undefined" dstip=96.45.36.159 dstport=443
dstintf="port1" dstintfrole="undefined" proto=6 service="HTTPS"
hostname="site.fortinet.com" profile="monitor-all" action="passthrough"
reqtype="direct" url="https://site.fortinet.com/" sentbyte=517 rcvdbyte=0
direction="outgoing" msg="URL belongs to an allowed category in policy"
method="domain" cat=52 catdesc="Information Technology"
```

NSE Training Institute

63

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.

NSE Training Institute

64

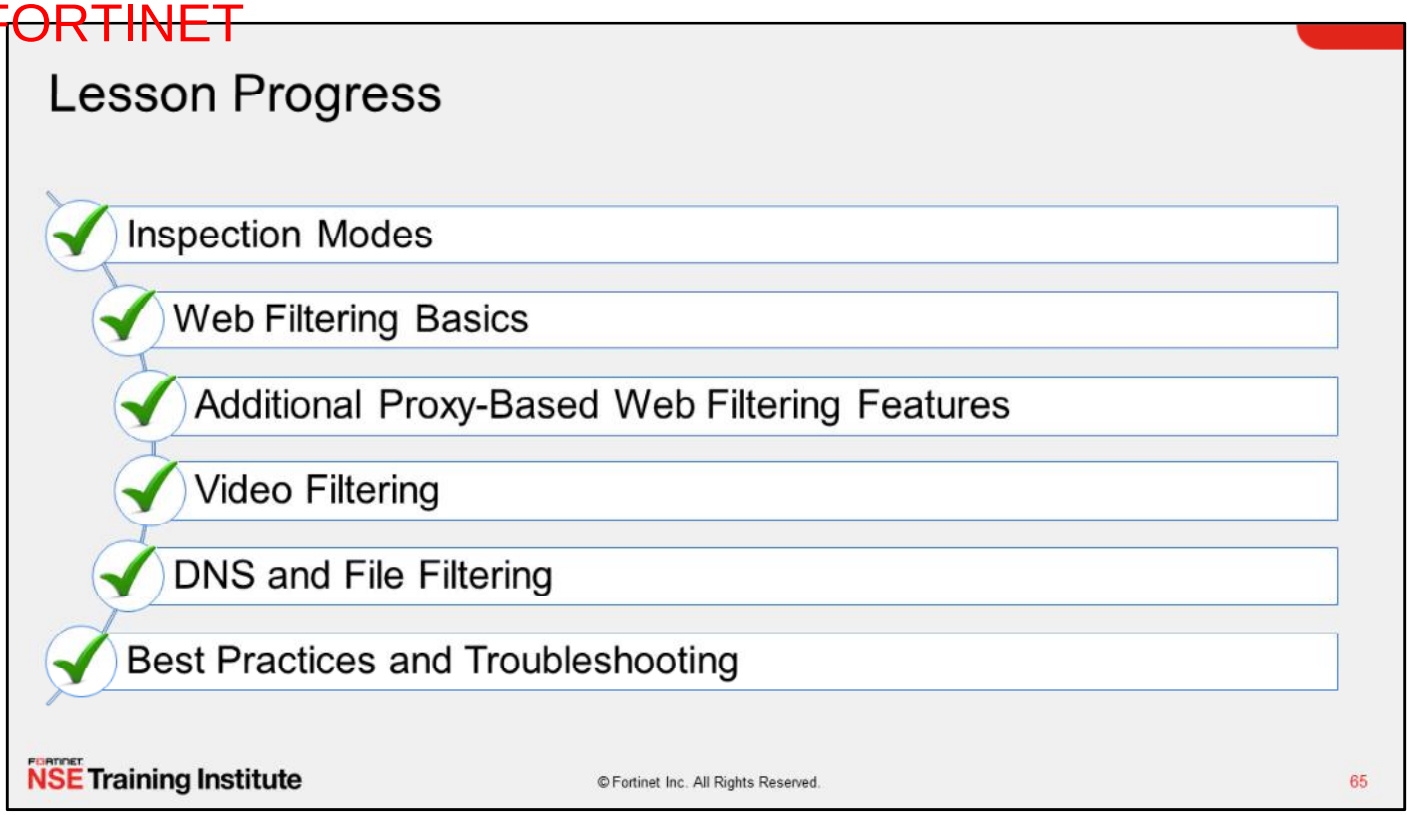

Congratulations! You have completed this lesson.

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

Review

- ✓ Describe FortiOS inspection modes
- ✓ Implement NGFW operation modes
- ✓ Work with web filter categories and custom categories
- ✓ Apply an SSL inspection profile to a firewall policy
- ✓ Exempt traffic from SSL inspection
- ✓ Configure web filter overrides and submit a FortiGuard rating request
- ✓ Configure web profile overrides
- ✓ Configure usage quotas
- ✓ Configure web filter to support search engines
- ✓ Apply video filter on proxy-based firewall policy
- ✓ Apply DNS and file filtering
- ✓ Monitor logs for web filtering events

NSE Training Institute 66

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.

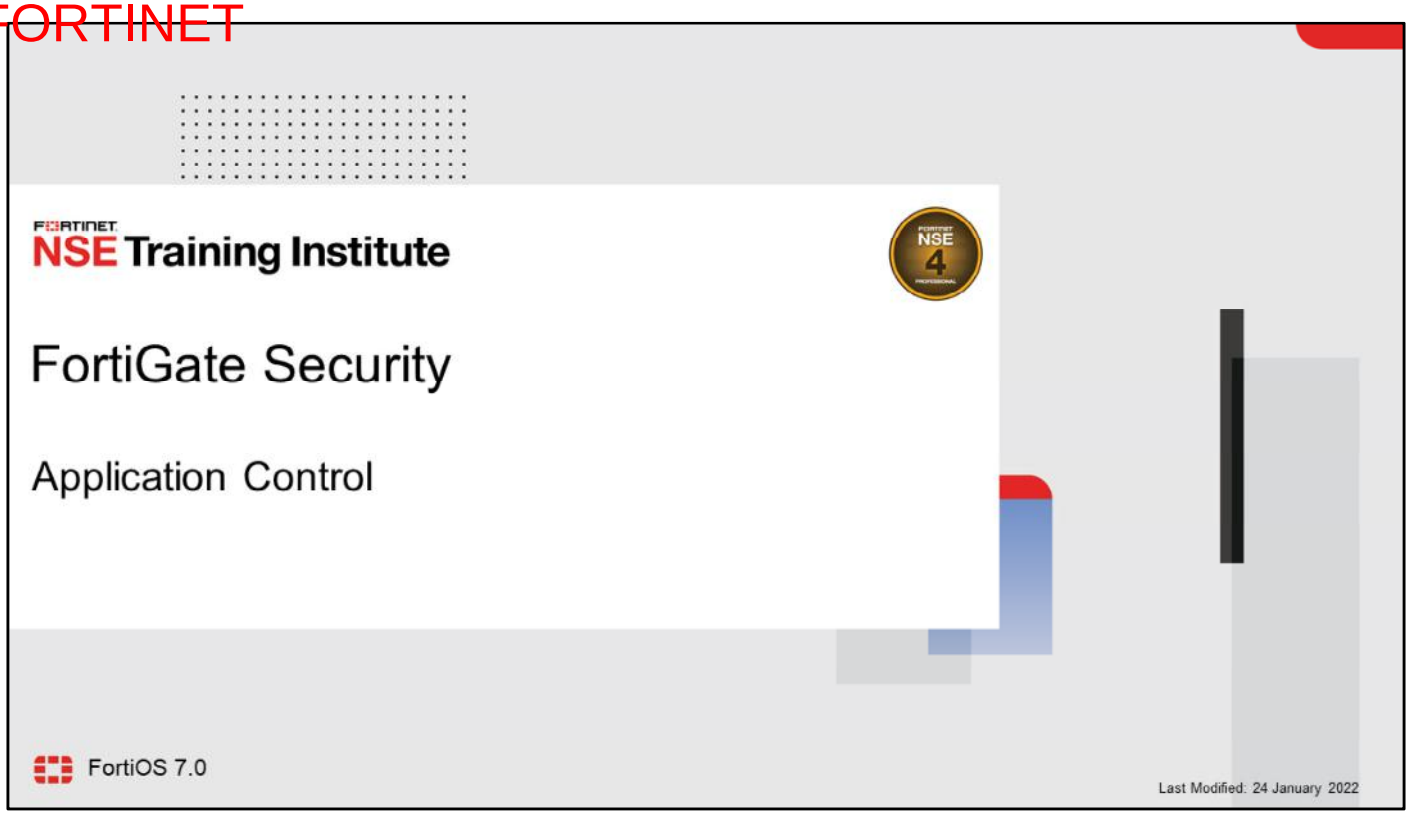

In this lesson, you will learn how to monitor and control network applications that may use standard or non-standard protocols and ports—beyond simply blocking or allowing a protocol, port number, or IP address.

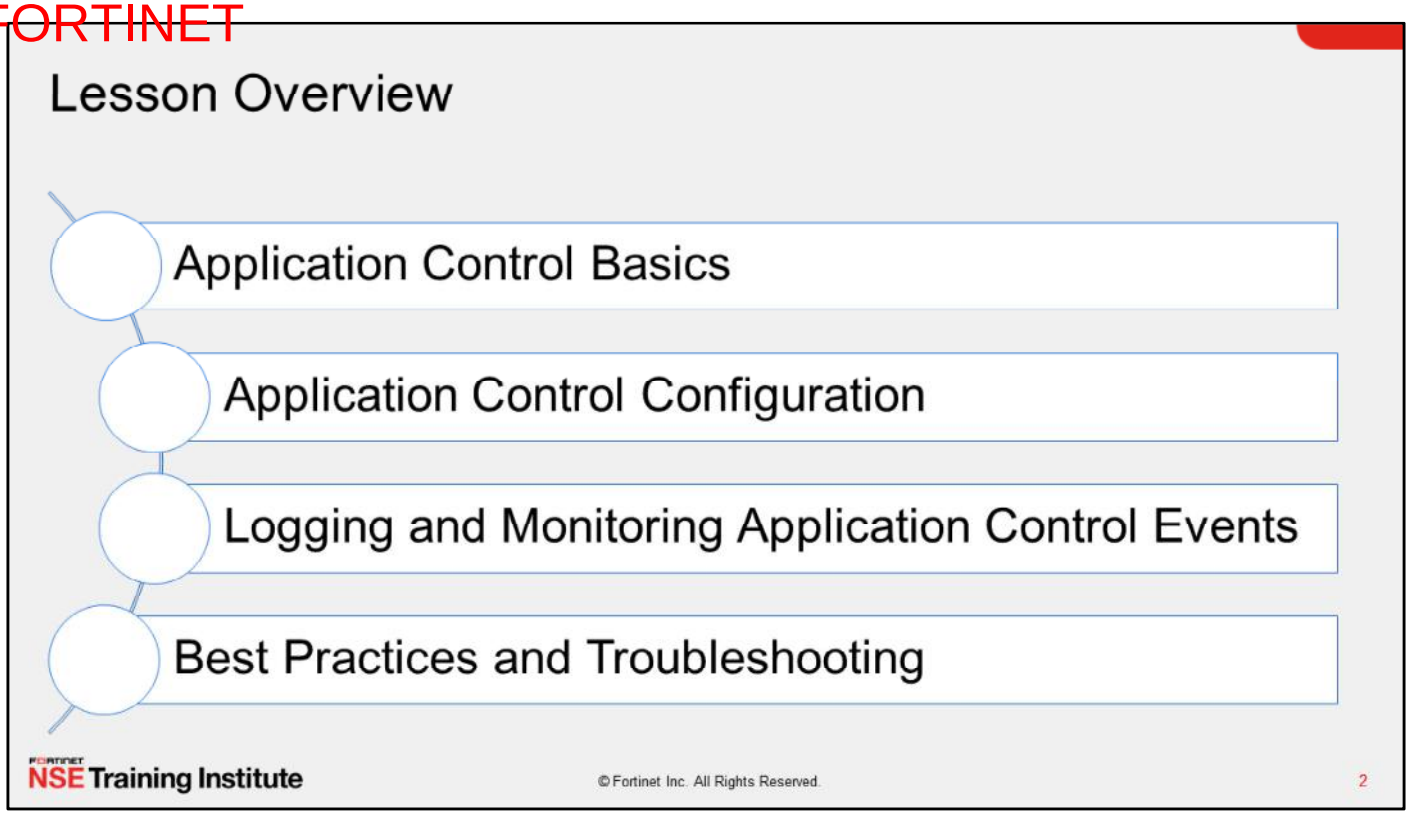

In this lesson, you will learn about the topics shown on this slide.

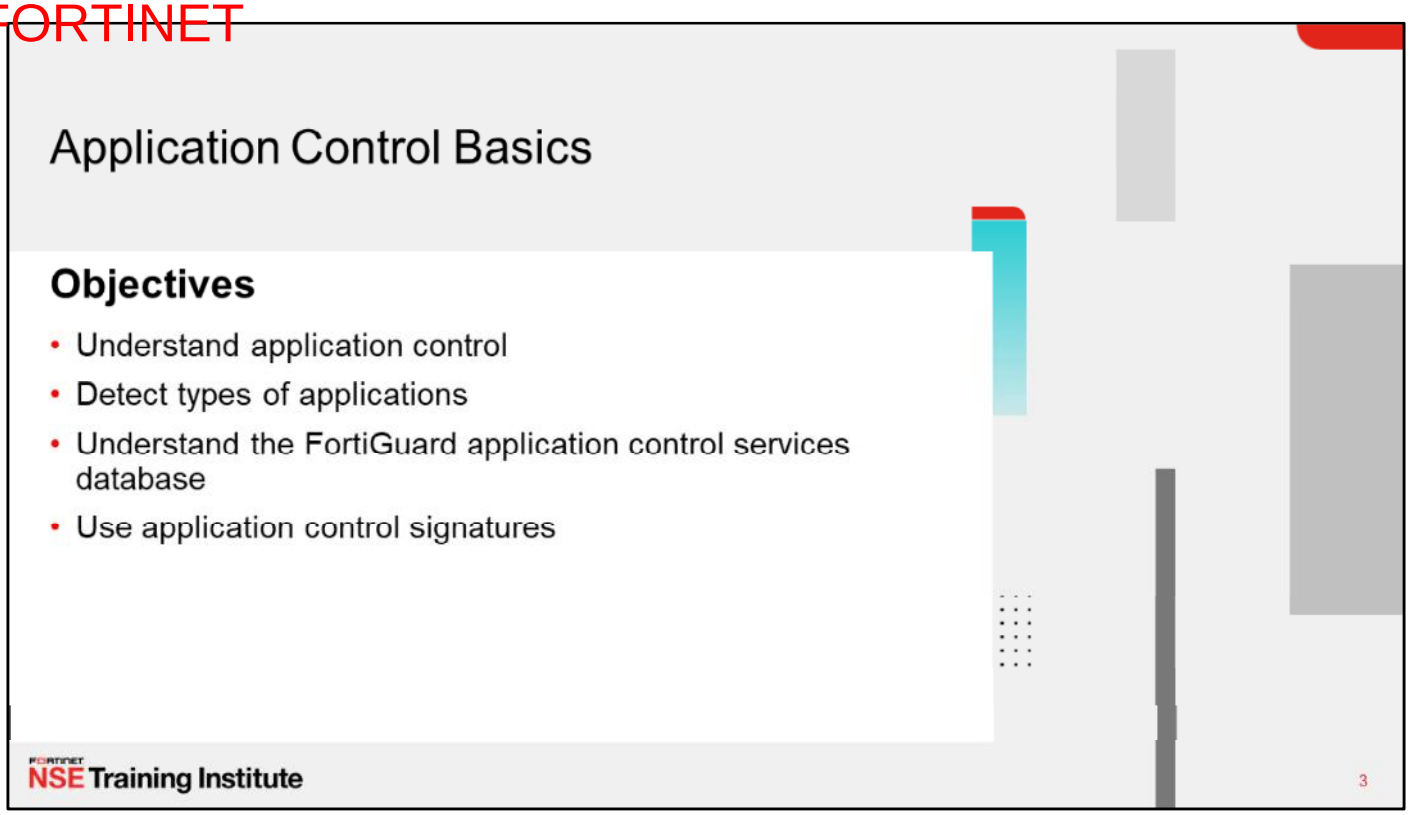

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control basics, you will be able to understand how application control works on FortiGate.

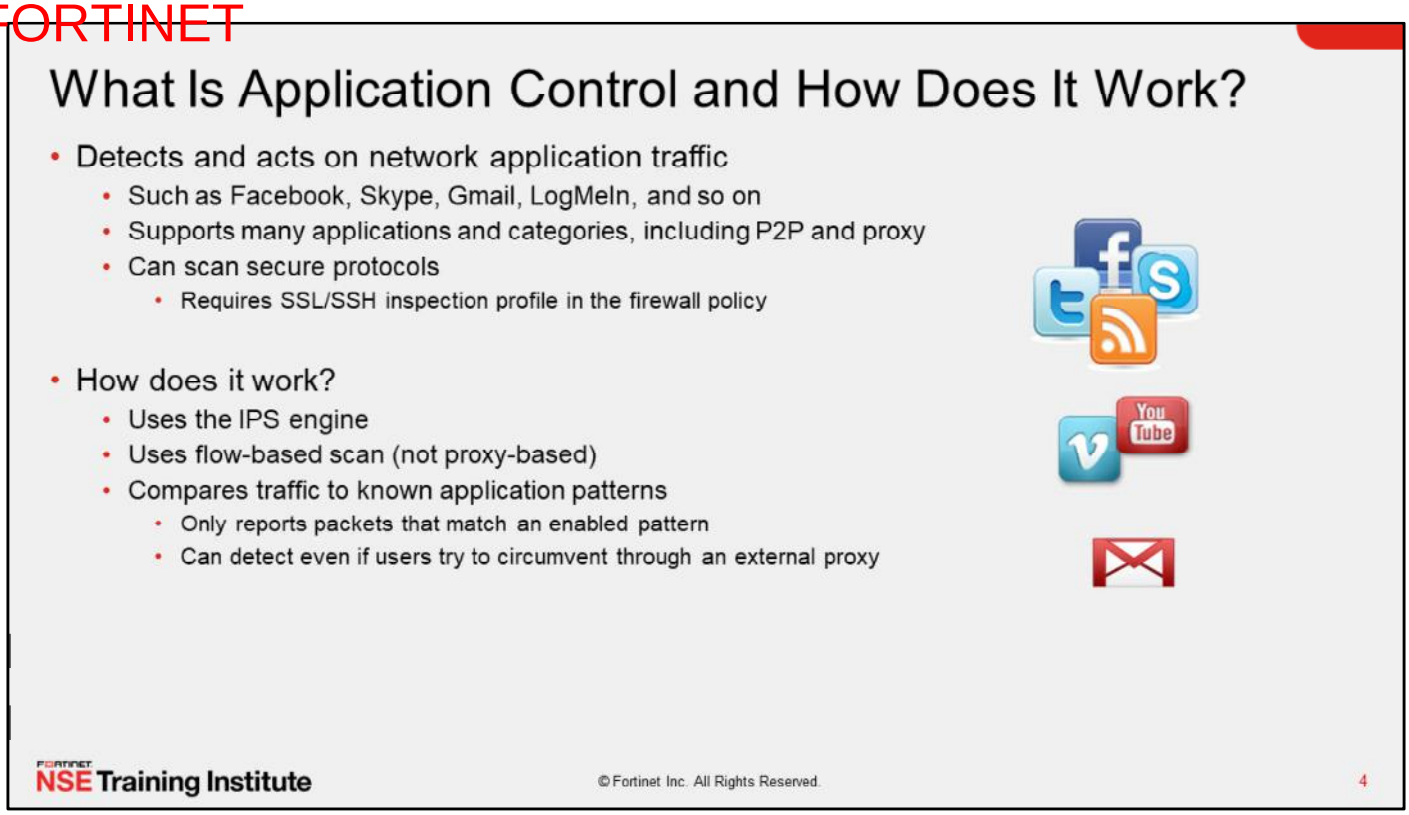

Application control detects applications—often applications that consume a lot of bandwidth—and allows you to take appropriate action related to application traffic, such as monitoring, blocking, or applying traffic shaping.

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.

Detecting Peer-to-Peer Applications

- Why is peer-to-peer (P2P) traffic so difficult to detect?
  - Traditional protocols (HTTP, FTP) have a client-server architecture
    - It uses a single server with large bandwidth for many clients
    - It requires predictable port numbers, NAT/PAT, and firewall policies
  - Peer-to-peer protocols (BitTorrent, Skype) have a distributed architecture
    - Each peer is a server with small bandwidth to share
    - They are difficult to manage multiple firewall policies to block them
    - They do not depend on port forwarding
    - They use evasive techniques to bypass these limitations

NSE Training Institute © Fortinet Inc. All Rights Reserved. 5

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.

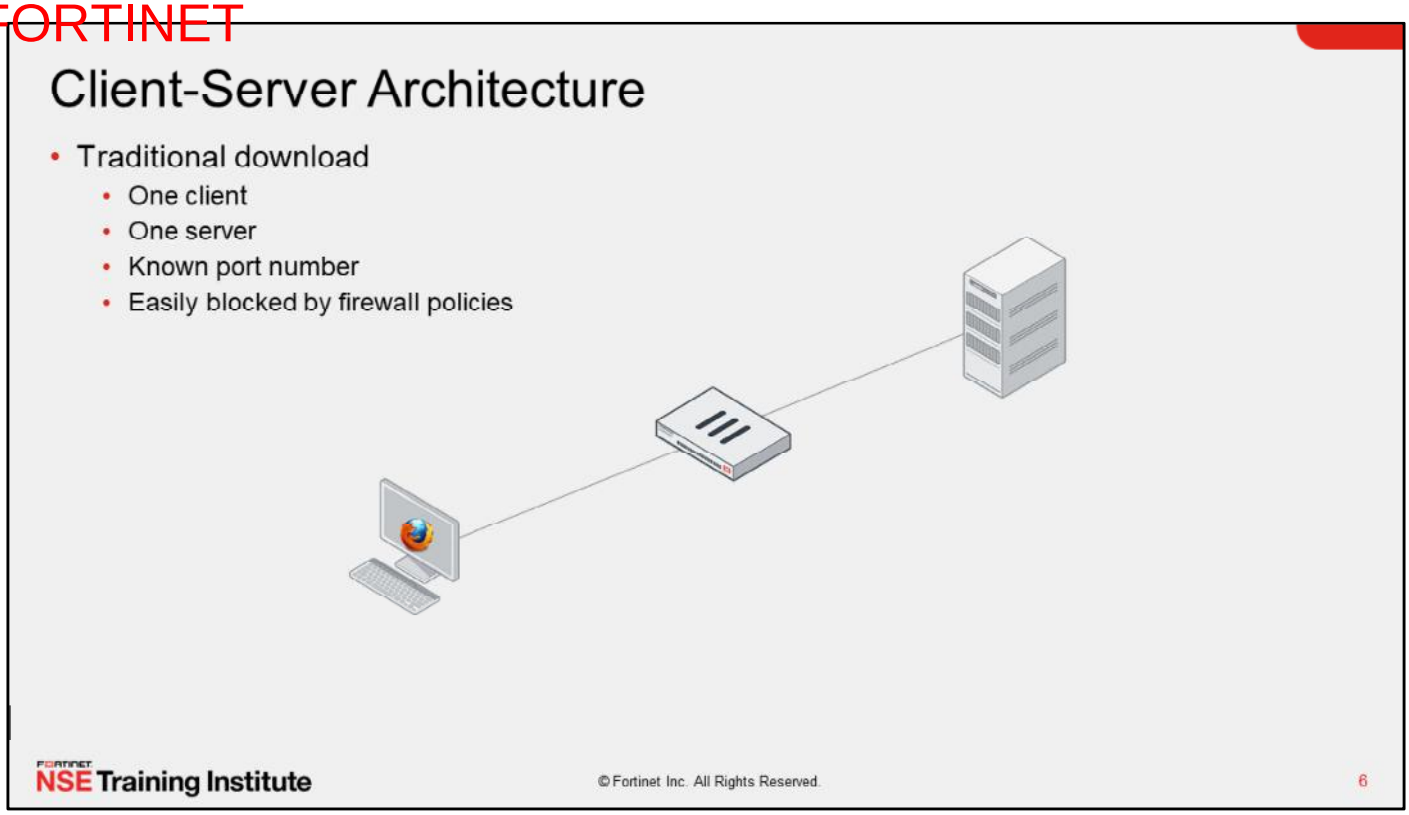

This slide shows a traditional, client-server architecture. There may be many clients of popular sites, but often, such as with an office file server, it's just one client and one server.

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?

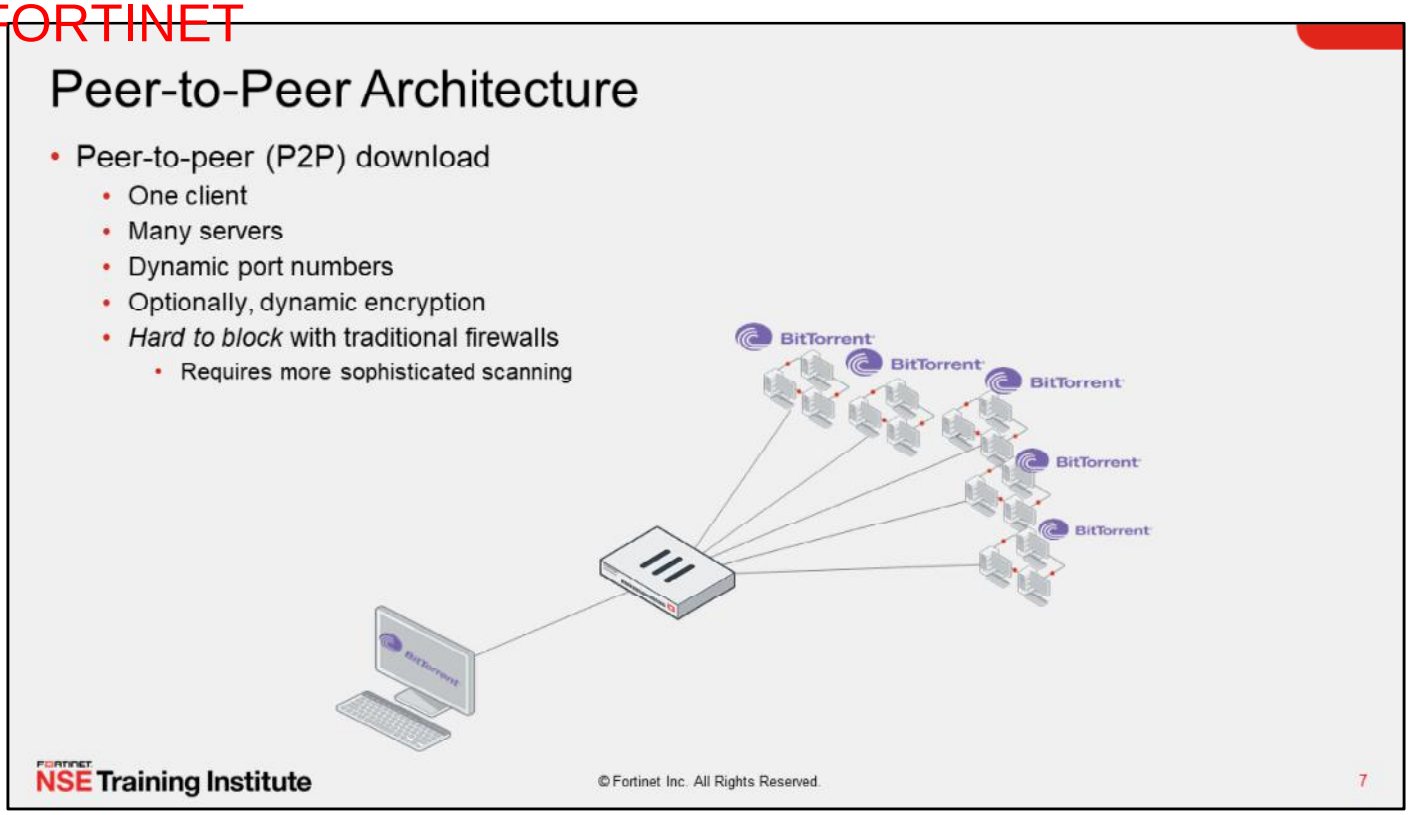

Peer-to-peer (P2P) downloads divide each file among multiple (theoretically unlimited) peers. Each peer delivers part of the file. While having many clients is a disadvantage in client-server architectures, it is an advantage for P2P architecture because, as the number of peers increases to n, the file 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.

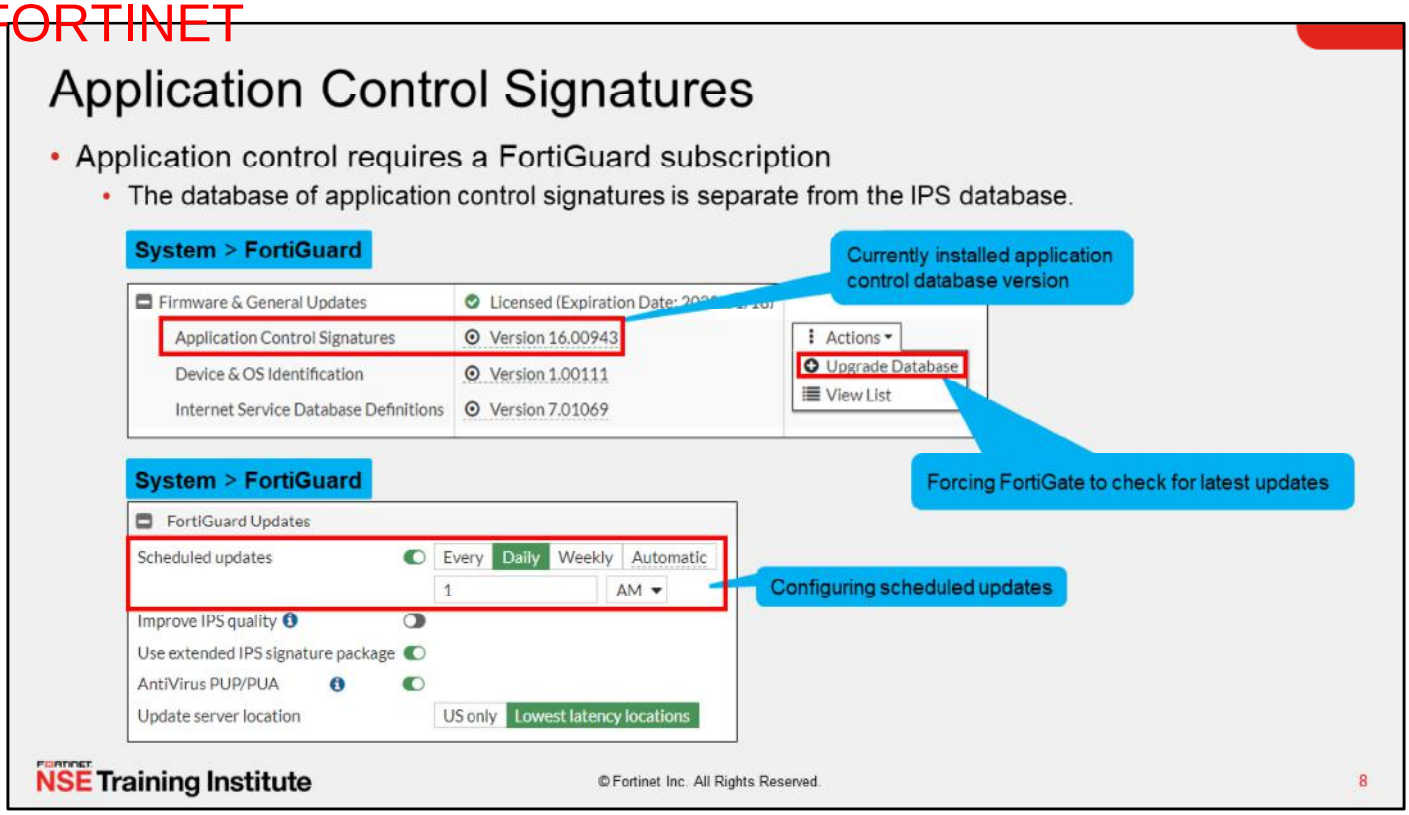

Before you try to control applications, it's important to understand the signatures used by application control.

How does application control detect the newest applications and changes to application protocols?

Application control requires a subscription to FortiGuard application control. 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.

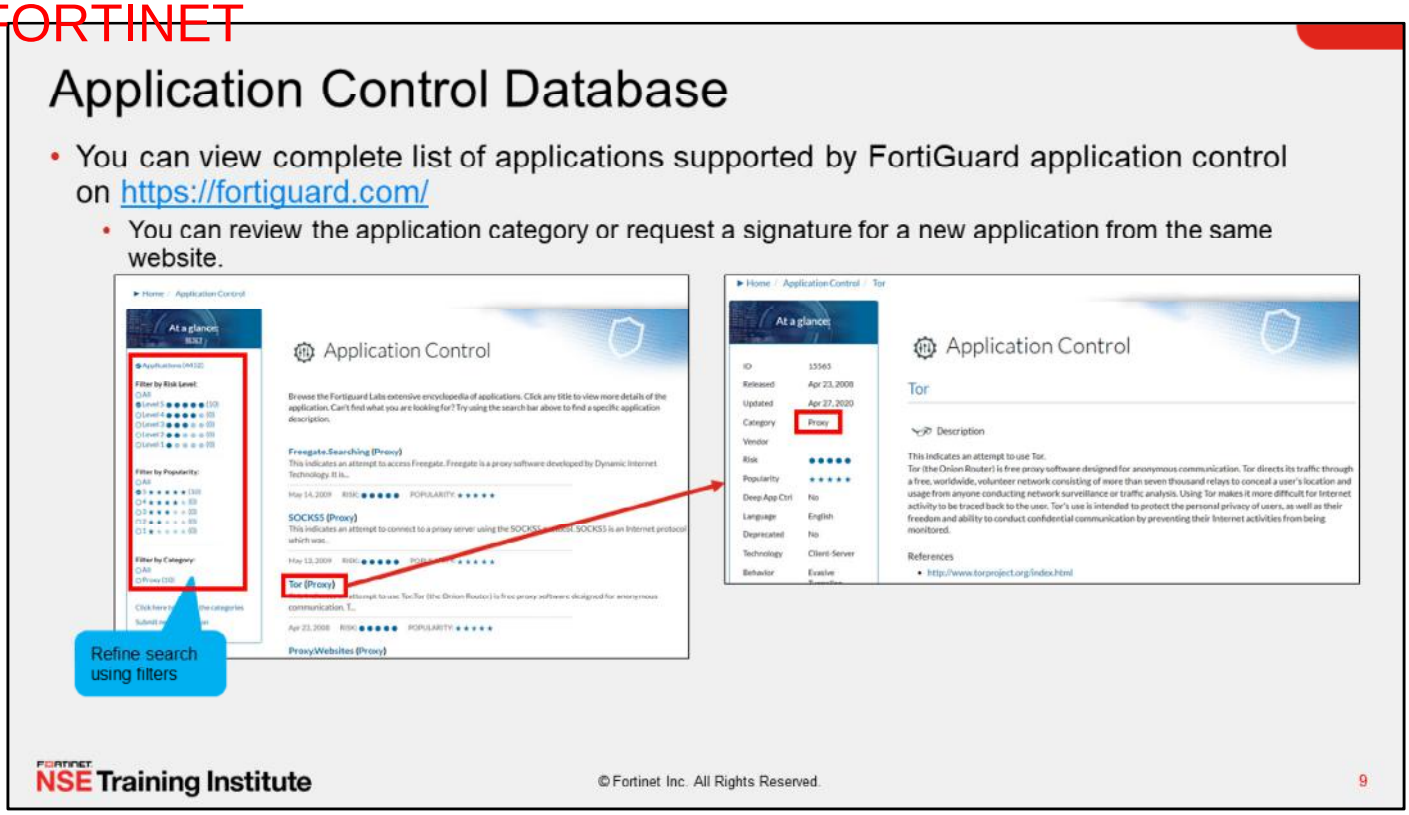

You can view the latest version of the application control database on the FortiGuard website, or by clicking an individual application signature in the application control profile.

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. The assigned risk level is 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. If you aren't aware of the specific application, this information can help you to decide if it would be wise to block an application or not.

On the FortiGuard website, you can read details about each signature's related application.

On this slide, you can see an example article for Tor is a web proxy, so it belongs in the proxy category. It is a good practice to create test policies that you can use to observe policy behavior.

If there are new applications that you need to control, and the latest update doesn't include definitions for them, you can go to the FortiGuard website and submit a request to have the new applications added. You can also submit a request to re-evaluate an application category, if you believe an application should belong to a different category.

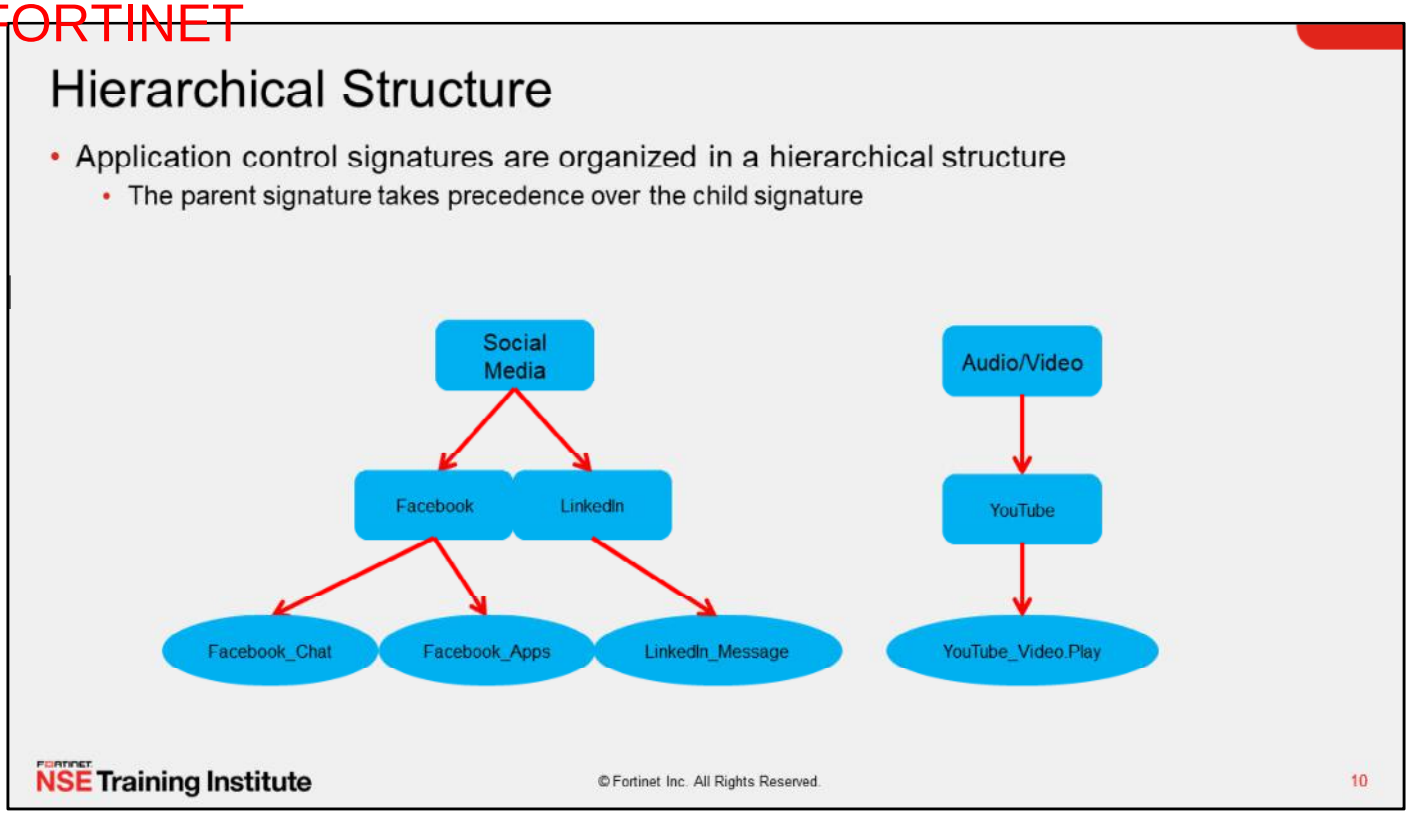

Many web applications offer functionality that can be embedded in third-party websites or applications. For example, you can embed a Facebook **Like** button at the end of an article or reference a YouTube video on an educational website. FortiOS gives administrators all the tools they need to inspect 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 apps 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.

NSE Training Institute

11

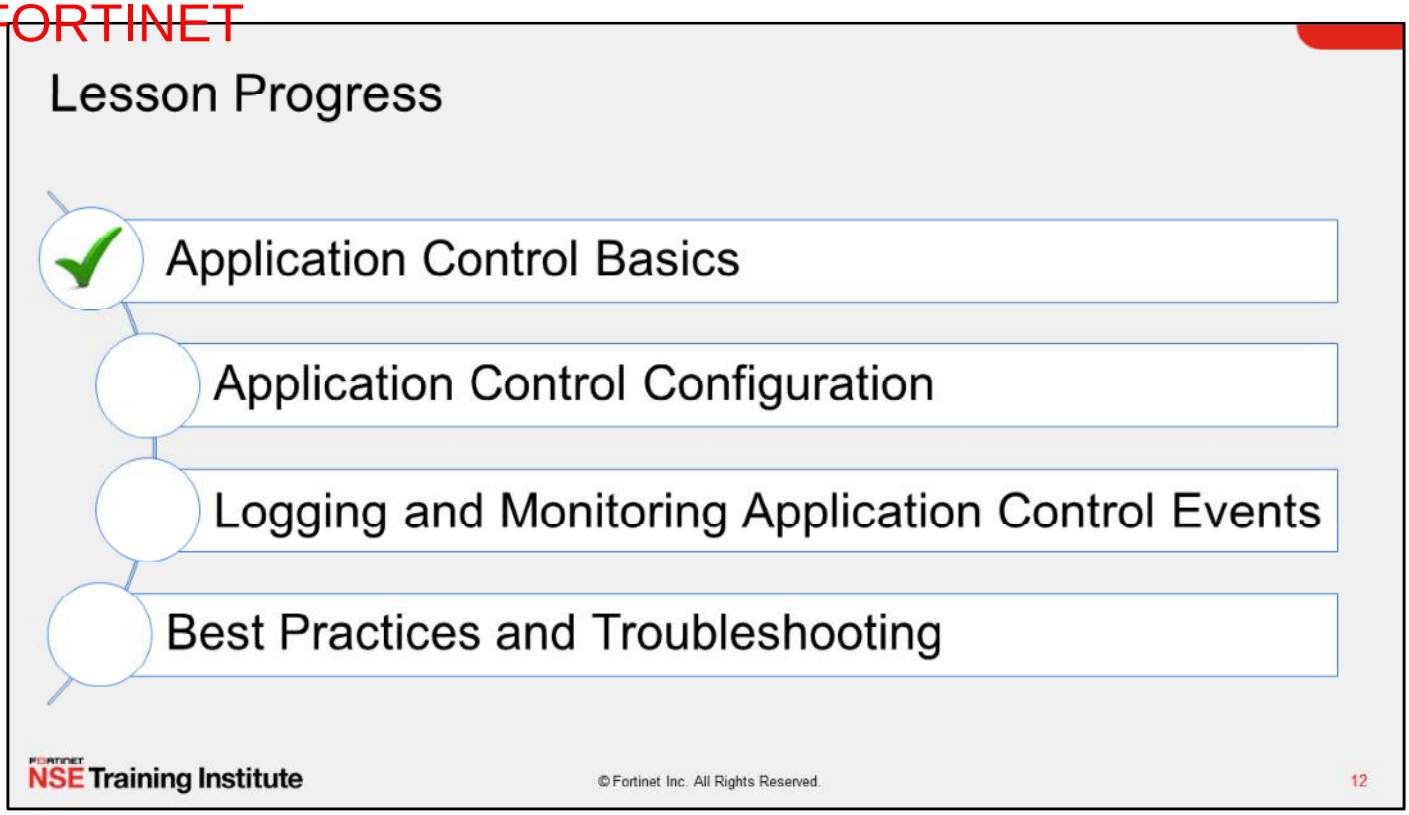

Good job! You now understand basic application control functionality.

Now, you will learn about application control configuration.

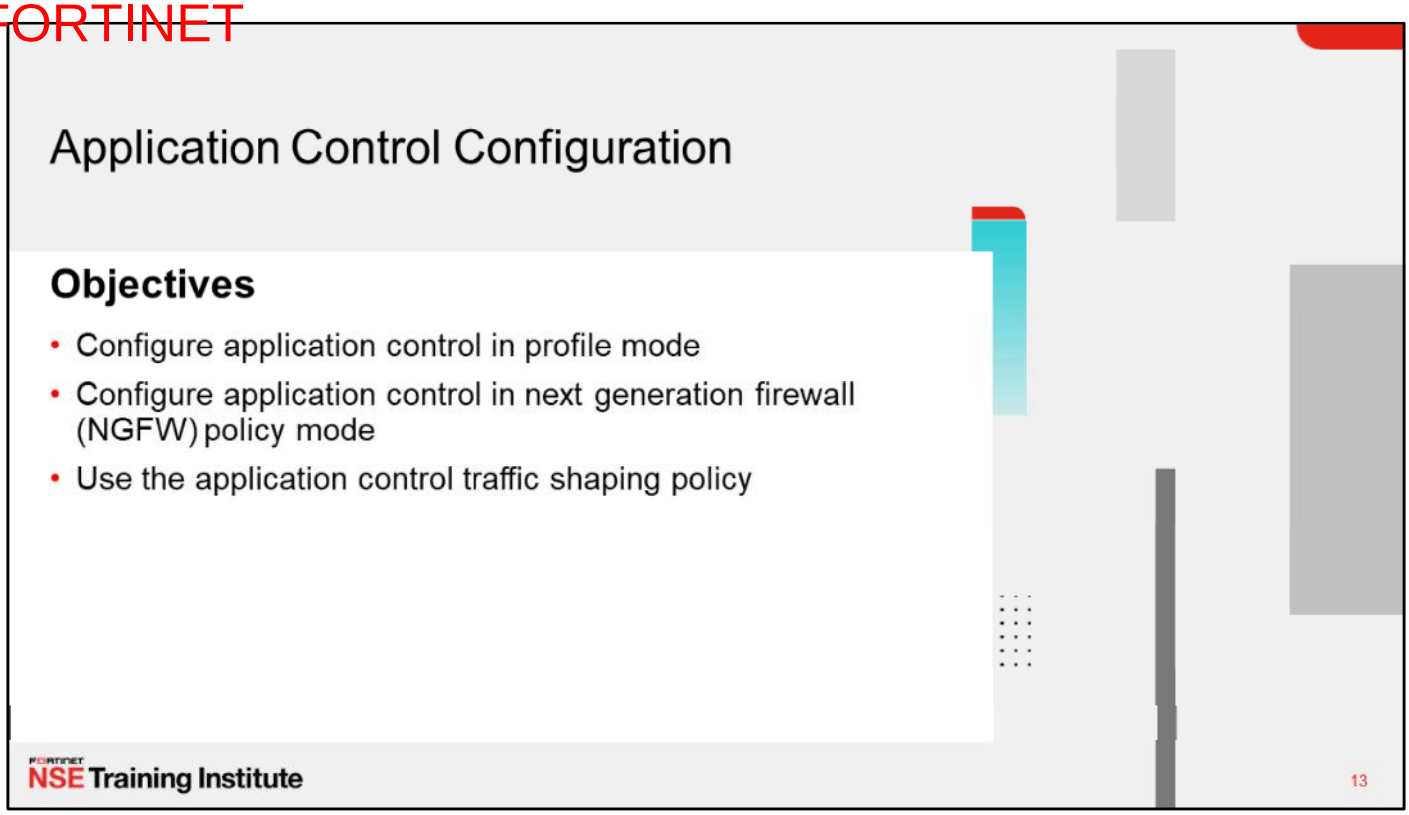

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

Application Control Profiles

- Configured when FortiGate NGFW mode is set to profile-based
- Uses flow-based scanning techniques in both inspection modes
- Allows you to filter application traffic based on:
  - Categories
    - Similar applications are grouped together
    - Can view application control signatures for that category
    - Can configure actions for predefined categories
  - Application overrides
    - Allows you to configure actions for specific signatures or applications
  - Filter overrides
    - Provides a more flexible way to create application categorization based on behavior, popularity, protocol, risk, and so on
- Must be applied to a firewall policy

NSE Training Institute © Fortinet Inc. All Rights Reserved. 14

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.

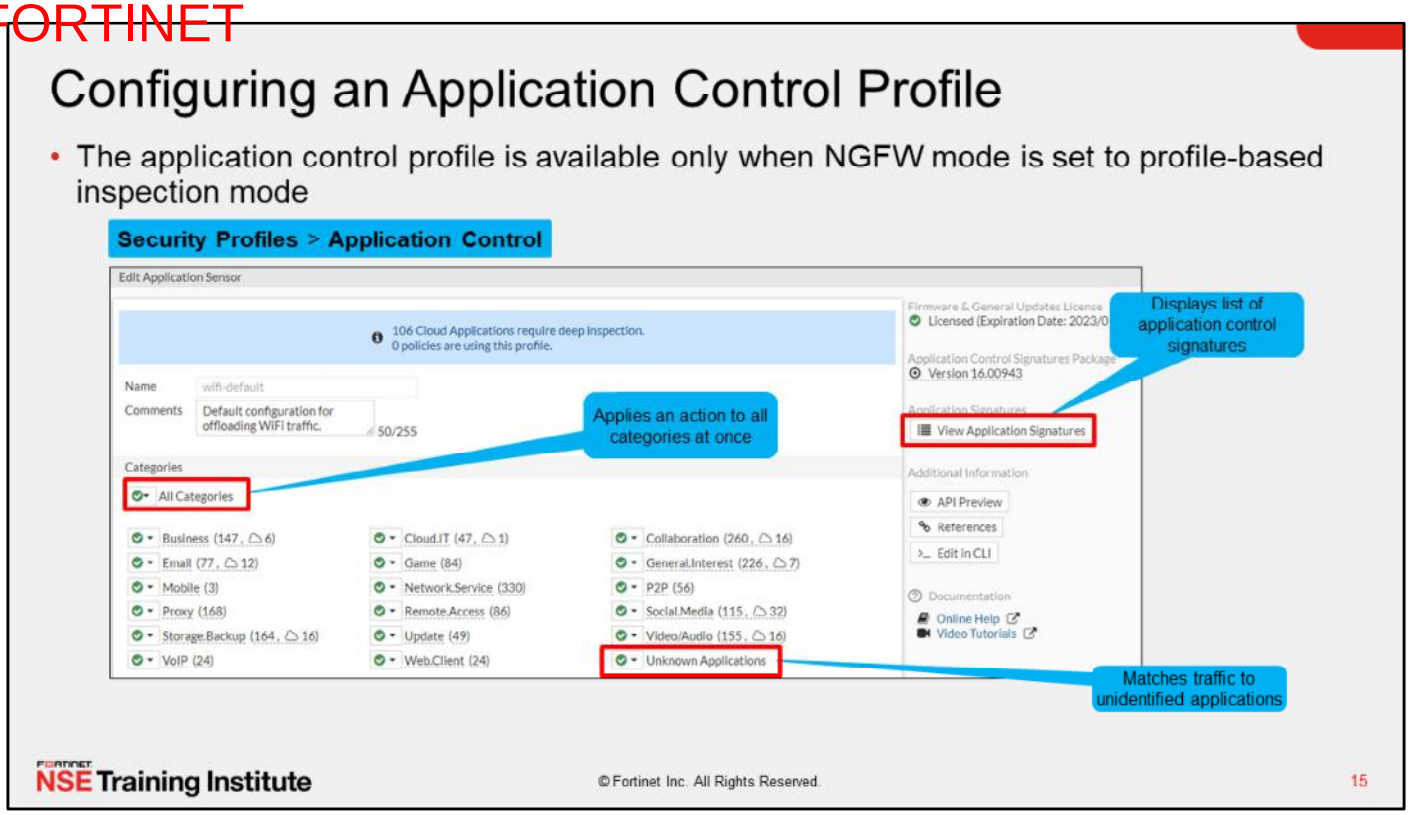

The application control profile is configured on the **Application Control** page. You can configure actions based on categories, application overrides, and filter overrides. You can also view the list of application control signatures by clicking **View Application Signatures**.

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.

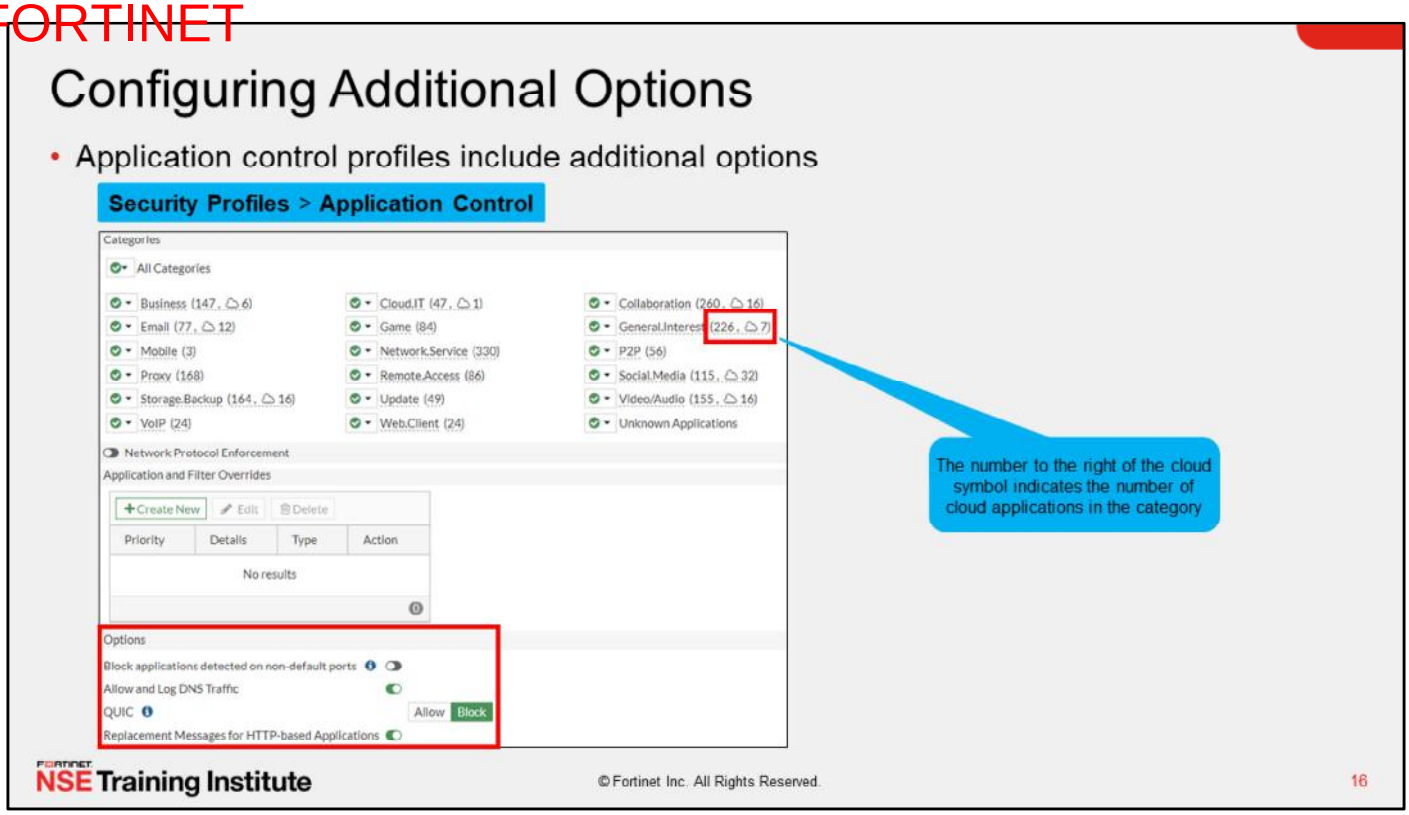

The number listed to the right of the cloud symbol indicates the number of cloud applications within a specified category.

If you need to enable **Allow and Log DNS Traffic**, you should enable it only for short periods, such as during an investigation. Depending on the application and how often it queries DNS servers, enabling this setting can use significant system resources.

QUIC is a protocol from Google. Instead of using the standard TCP connections for web access, QUIC uses UDP, which is not scanned by web filtering. Allowing QUIC instructs FortiGate to inspect Google Chrome packets for a QUIC header, and generate logs as a QUIC message. Blocking QUIC forces Google Chrome to use HTTP2/TLS1.2 and FortiGate to log QUIC as blocked. The default action for QUIC is **Block**.

The **Replacement Messages for HTTP-based Applications** setting allows you to replace blocked content with an explanation (for the user's benefit). However, for non-HTTP/HTTPS applications, you can only drop the packets or reset 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.

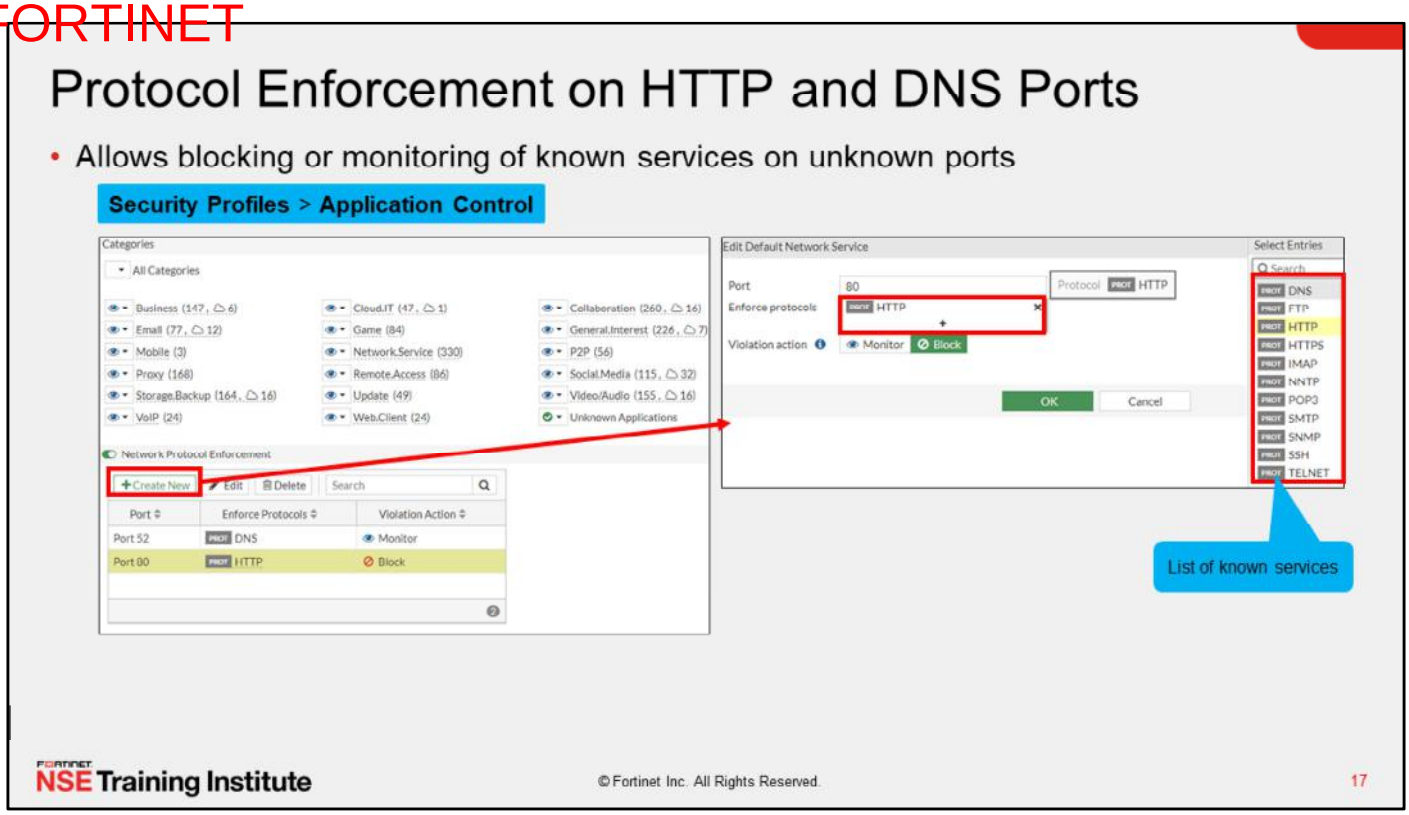

Protocol enforcement is added to the application control profile, allowing the administrator to configure network services (for example, FTP, HTTP, and HTTPS) on known ports (for example, 21, 80, and 443), while blocking those services on other ports.

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.

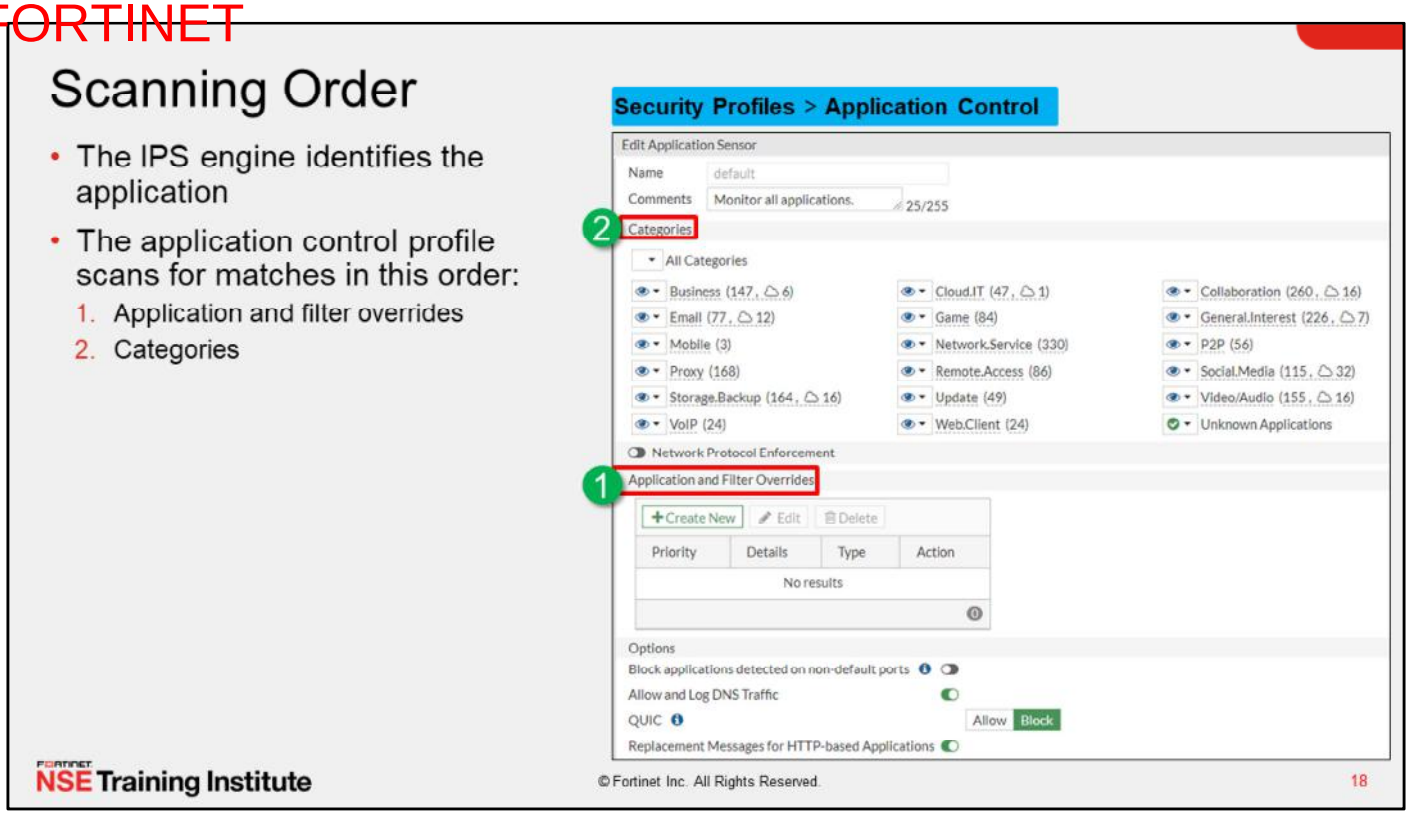

The IPS engine examines the traffic stream for a signature match.

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.

Order of Scan and Blocking Behavior (Scenario 1)

1. **Application Overrides**: Battle.Net and Dailymotion applications are set to **Monitor**
2. **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)
3. **Categories**: The **Game** and **Video/Audio** categories are set to **Block** and all other categories are set to **Monitor**

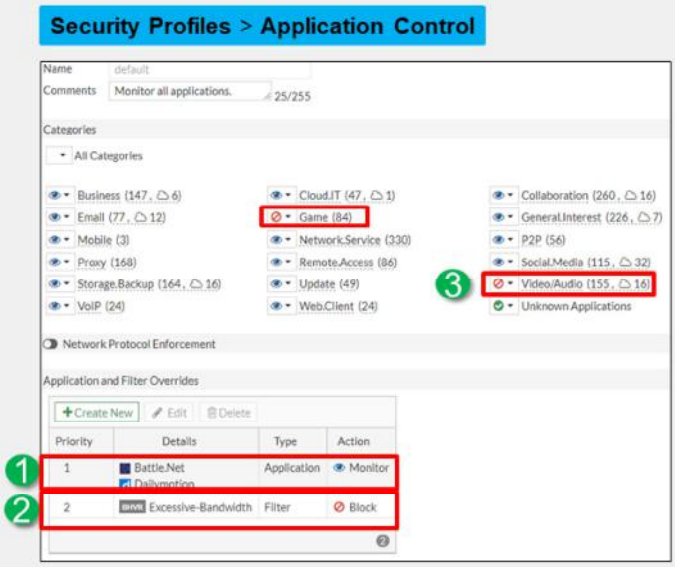

NSE Training Institute © Fortinet Inc. All Rights Reserved. 19

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**

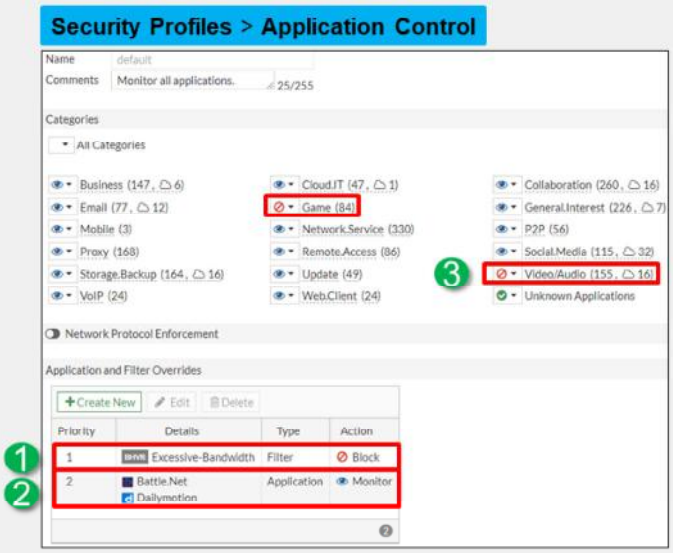

FORTINET NSE Training Institute

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.

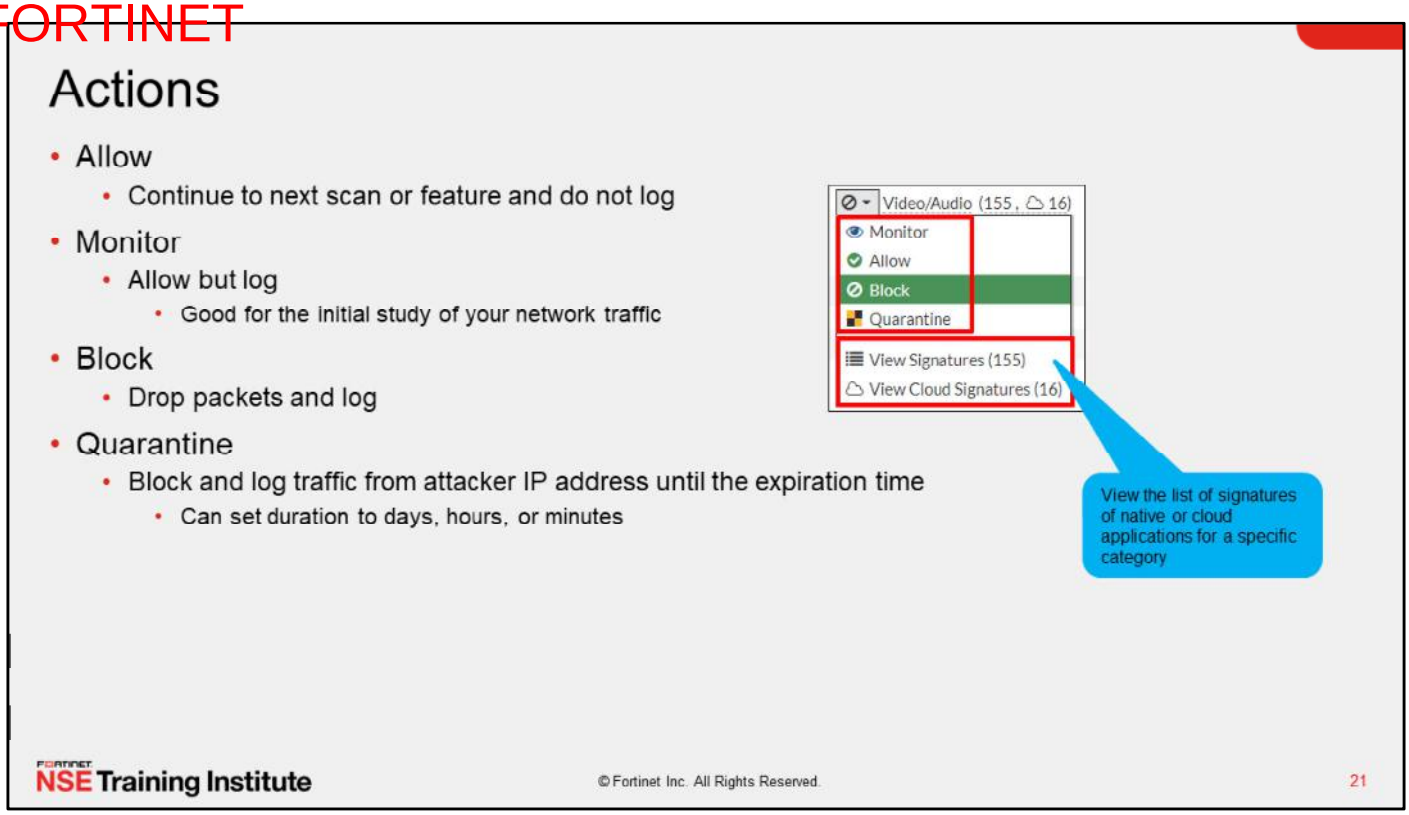

For each filter in the application control profile, you must indicate an action—what FortiGate does when traffic matches. Actions include the following:

- **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.

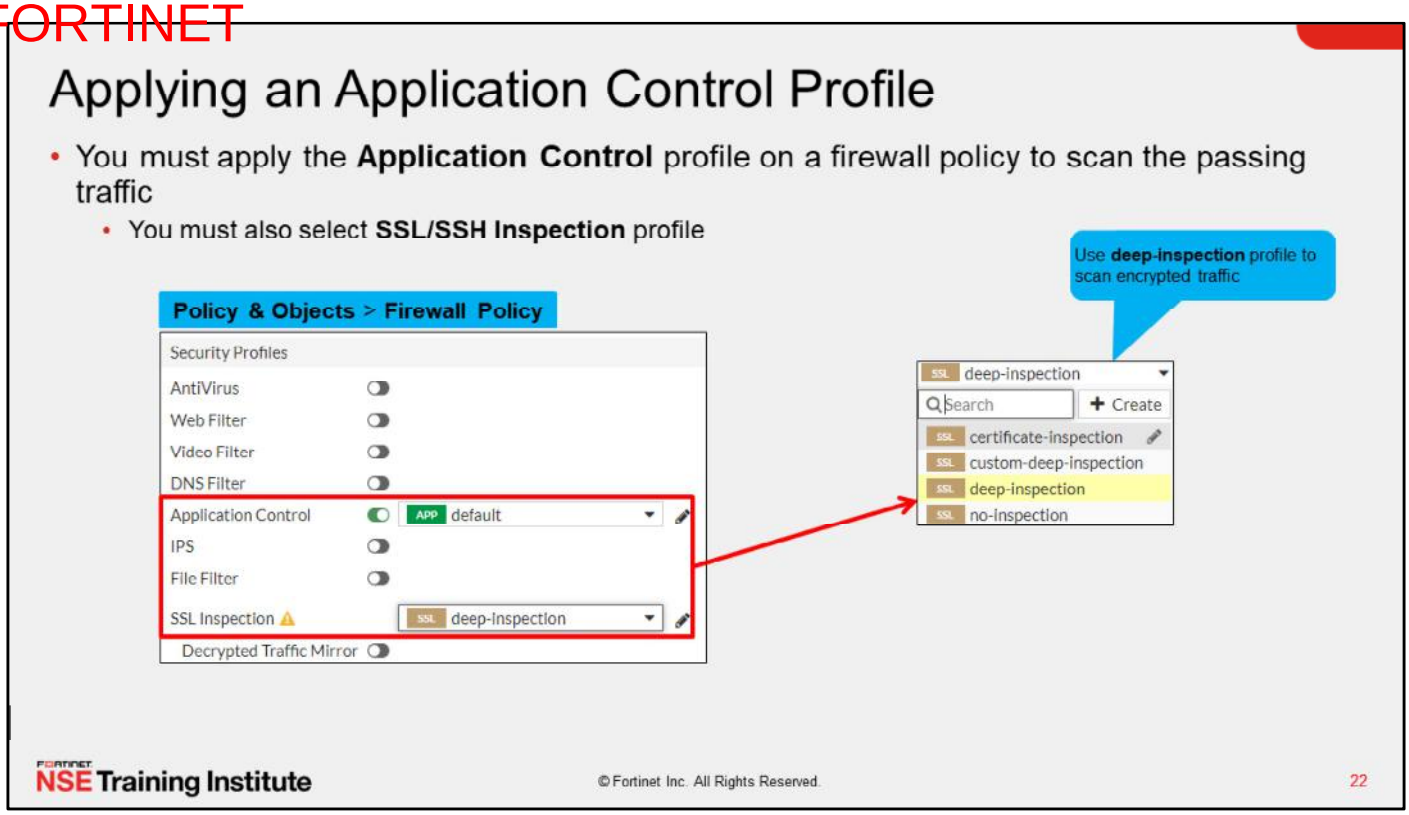

After you configure an application control profile, you must apply it to a firewall policy. This instructs FortiGate to start scanning application traffic that is subject to the firewall policy.

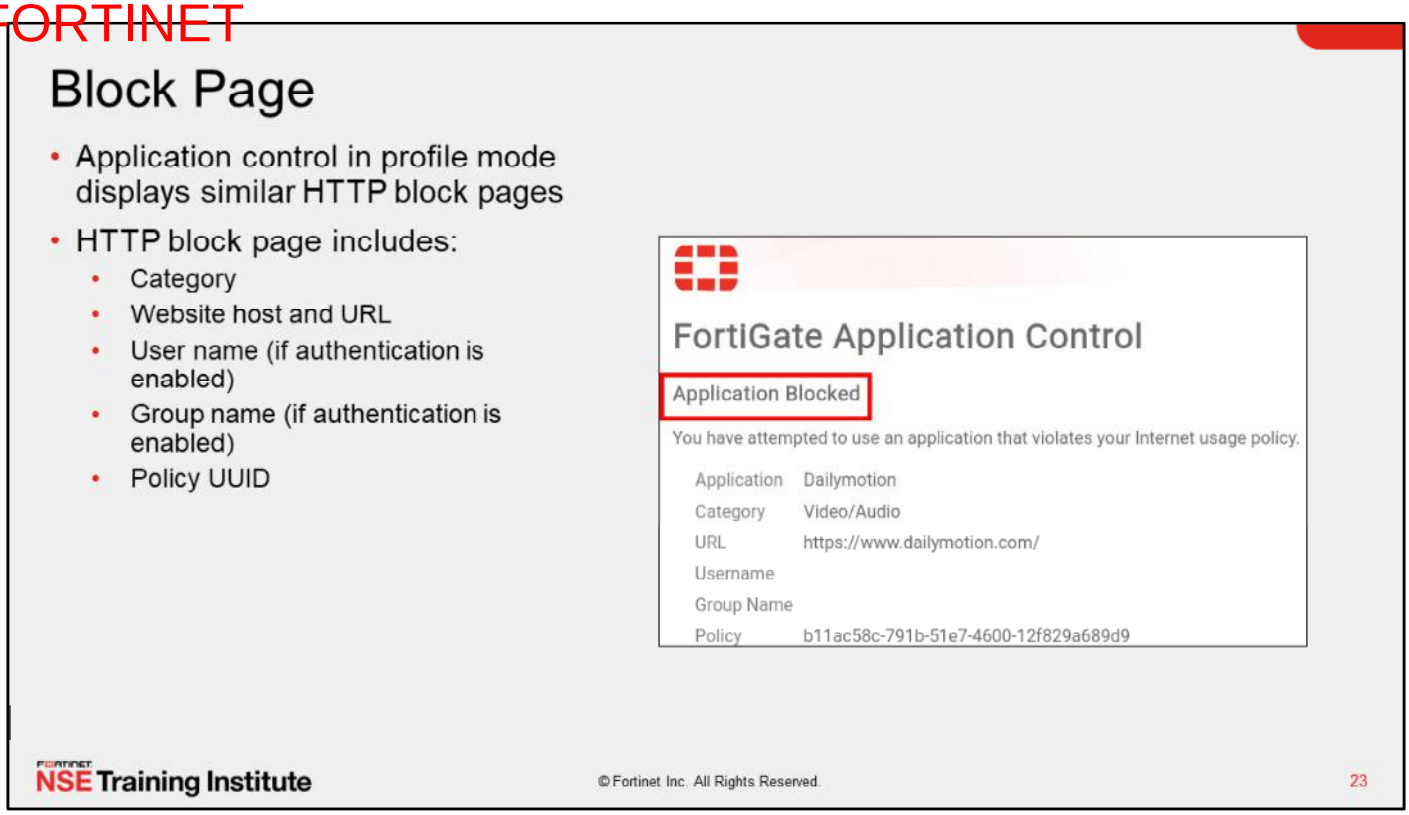

For HTTP-based applications, application control can provide feedback to the user about why their application was blocked. This is called a block page, and it is similar to the one you can configure for URLs that you block using FortiGuard web filtering.

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.

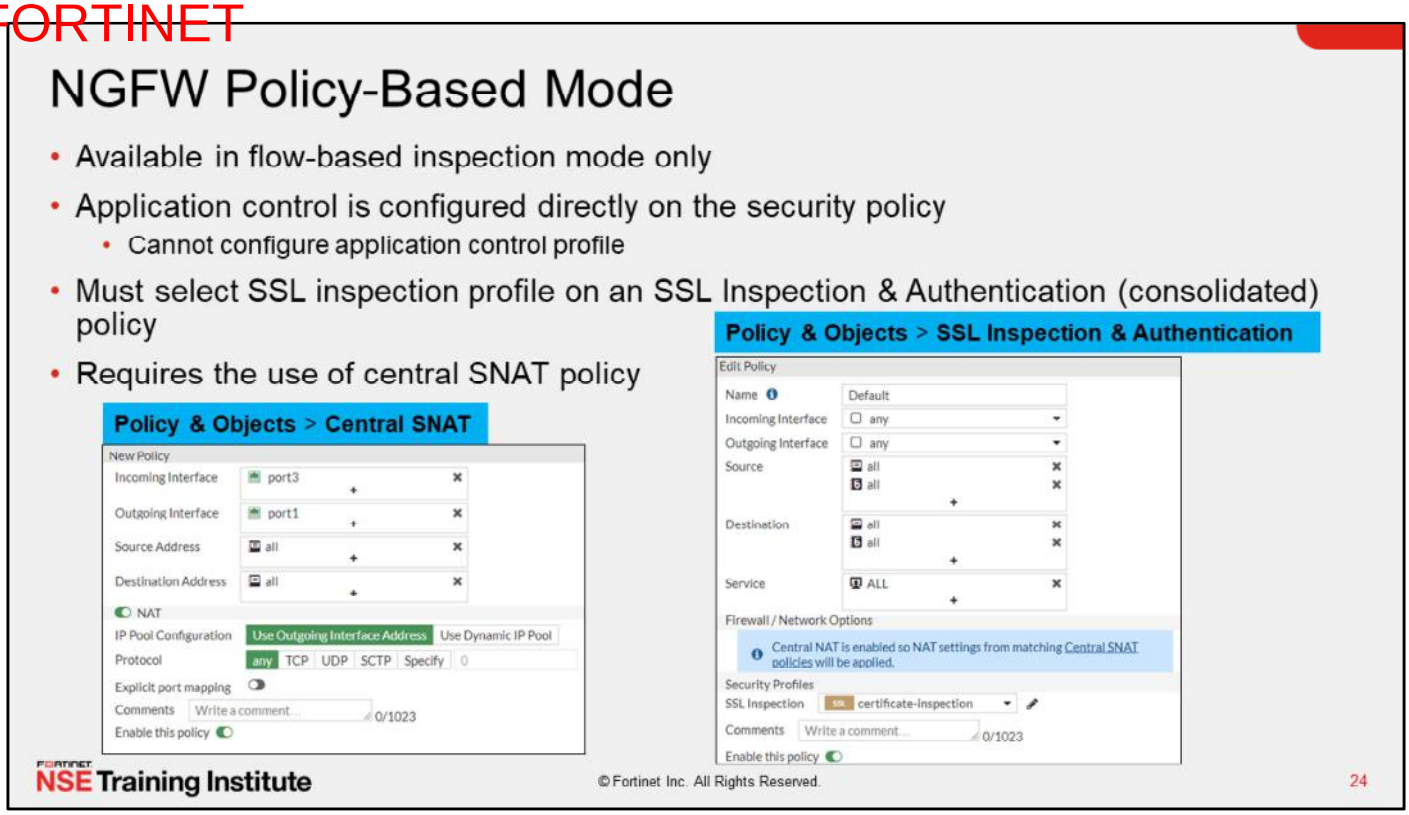

When FortiGate is operating in NGFW policy-based mode, administrators can apply application control to a security policy directly, instead of having to create an application control profile first, and then apply that to a firewall policy. Eliminating the need to use an application control profile makes 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.

NGFW Policy-Based Mode (Contd)

- You can select applications, application categories, or groups directly on a security policy
- You can apply the **ACCEPT** or **DENY** actions to allow or block selected application traffic
- If a **URL Category** is set, then applications that you add to the policy must be within the browser-based technology category
- You can apply the **AntiVirus** and **IPS** security profiles to a security policy with the action set to **ACCEPT**

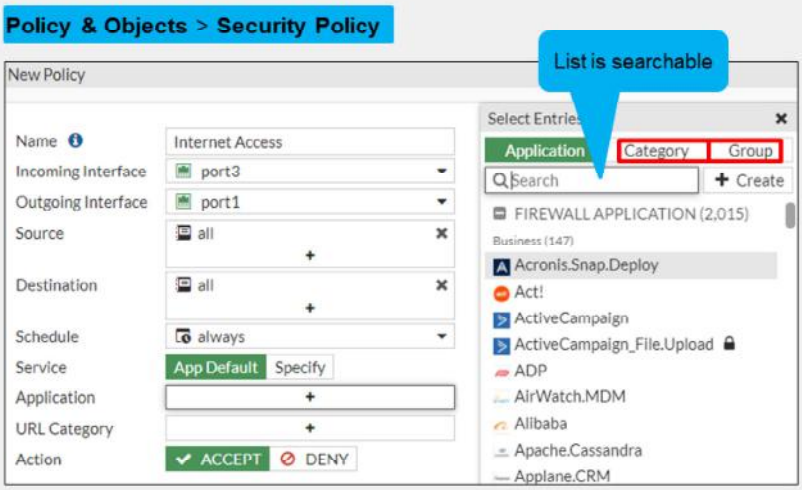

NSE Training Institute

25

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.

How Does NGFW Policy-Based Filtering Work?

- It is a three-step process:
  - Step 1–Allow all applications until they can be identified:
    - Uses only the IPv4 header information to match the policy
    - Accepts the traffic
    - Creates an entry in the session table with the `may_dirty` flag
    - Forwards all the packets to the IPS engine for inspection
  - Step 2–As soon as the IPS engine identifies the application, it adds the following to the session:
    - `dirty` flag - instructs the kernel to re-evaluate session entry
    - `app_valid` flag - indicates that IPS engine has validated the traffic
    - Application ID
  - Step 3–The dirty flag instructs the kernel to look up the security policy again:
    - This time the kernel uses the Layer 4 headers *and* the Layer 7 information to match the traffic
    - The action configured in the security policy is applied to the identified application traffic

NSE Training Institute © Fortinet Inc. All Rights Reserved. 26

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.

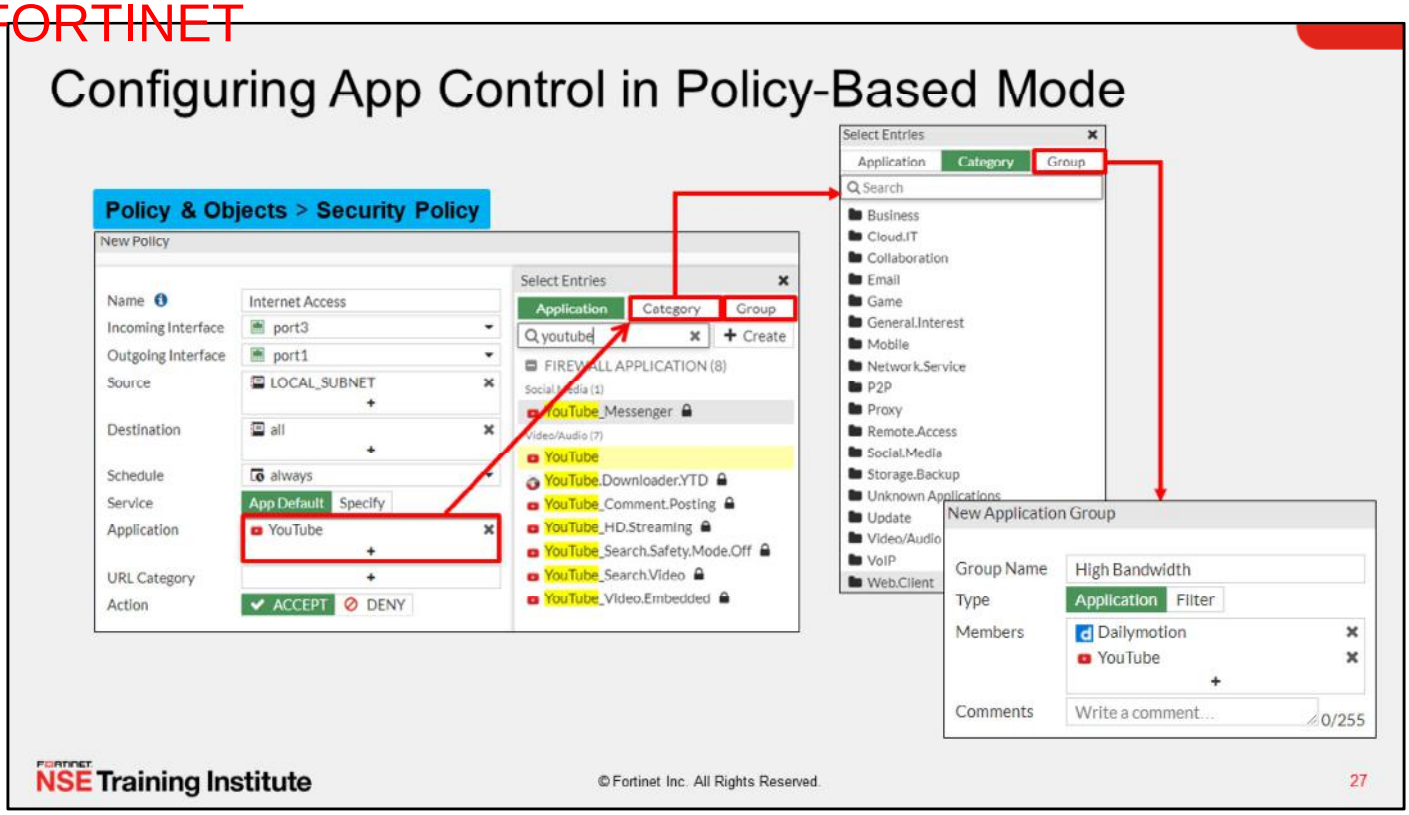

Configuring application control in NGFW policy-based mode is simple. You can create a new security policy or edit an existing security policy. In the **Application** section, select the applications, categories, or groups that you want to allow or deny, and change the security policy **Action** accordingly. On 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.

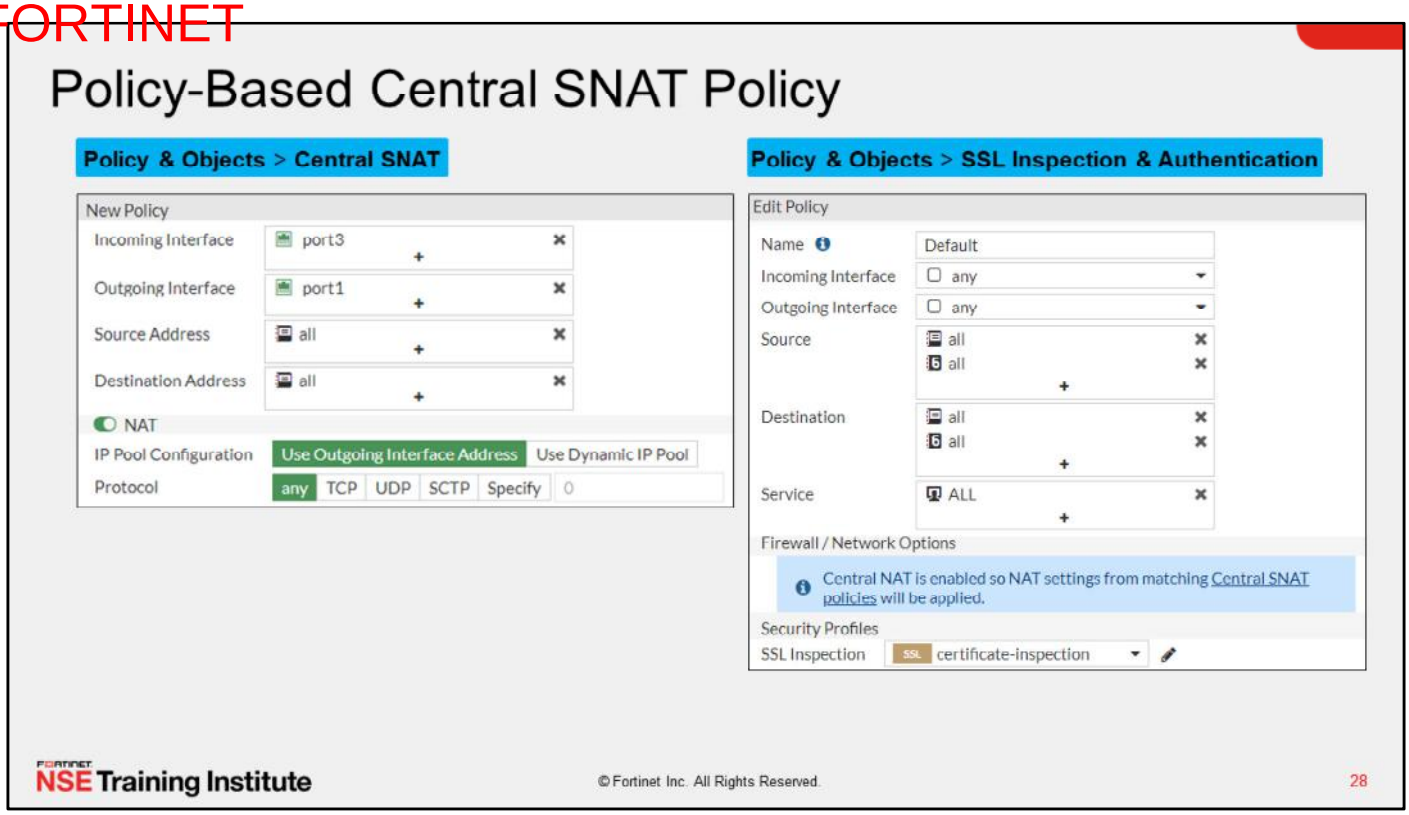

You must have a matching central SNAT policy in NGFW policy-based mode to be able to pass traffic. NAT is applied on the traffic based on criteria defined in the central SNAT policy.

It is extremely important to arrange security policies so that the more specific policies are located at the top to ensure proper use of application control.

A default **SSL Inspection & Authentication** policy is defined to inspect 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

NSE 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.

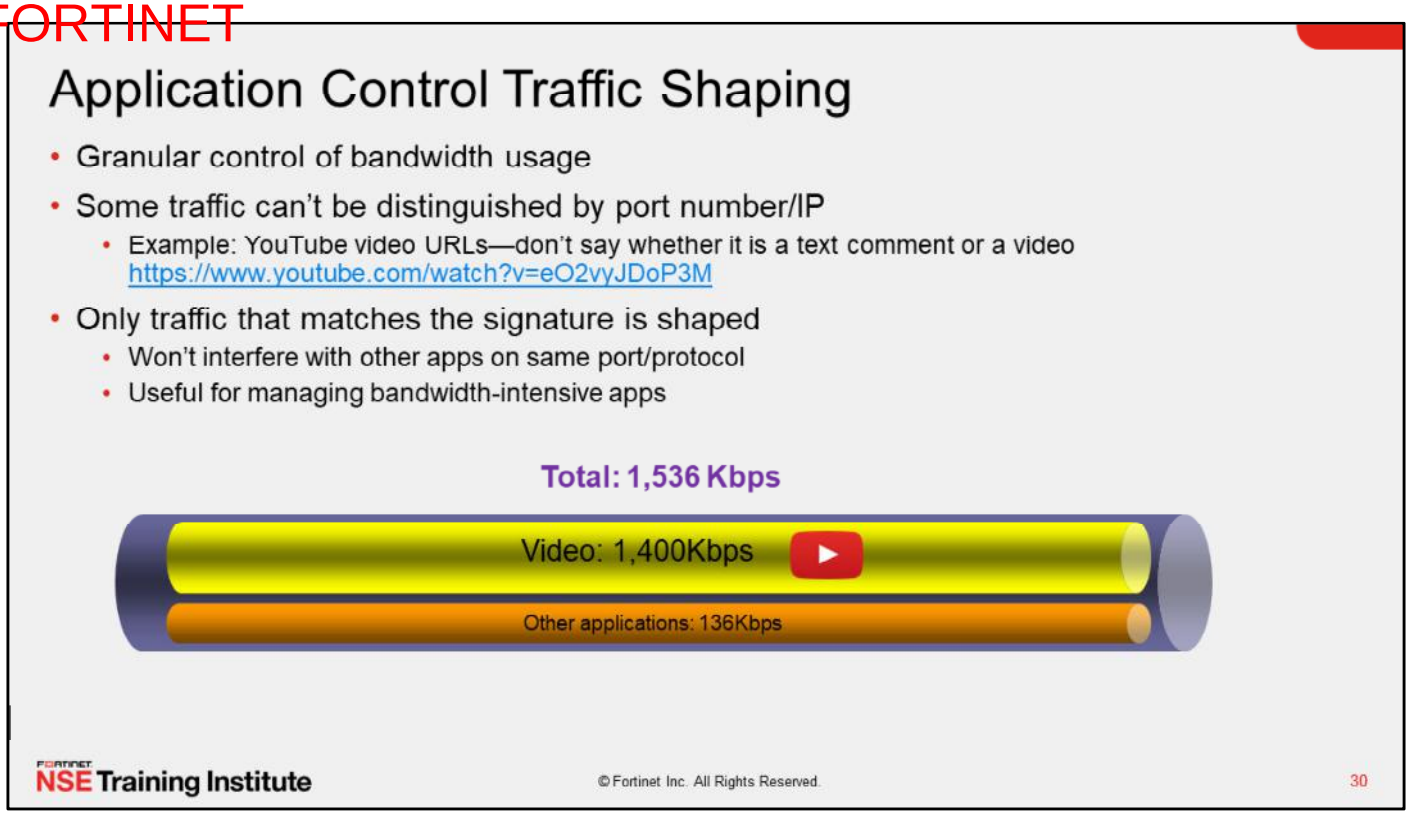

If an application is necessary, but you must prevent it from impacting bandwidth then, instead of blocking it entirely, you can apply a rate limit to the application. For example, you can rate limit applications used for storage or backup leaving enough bandwidth for more sensitive streaming applications, 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.

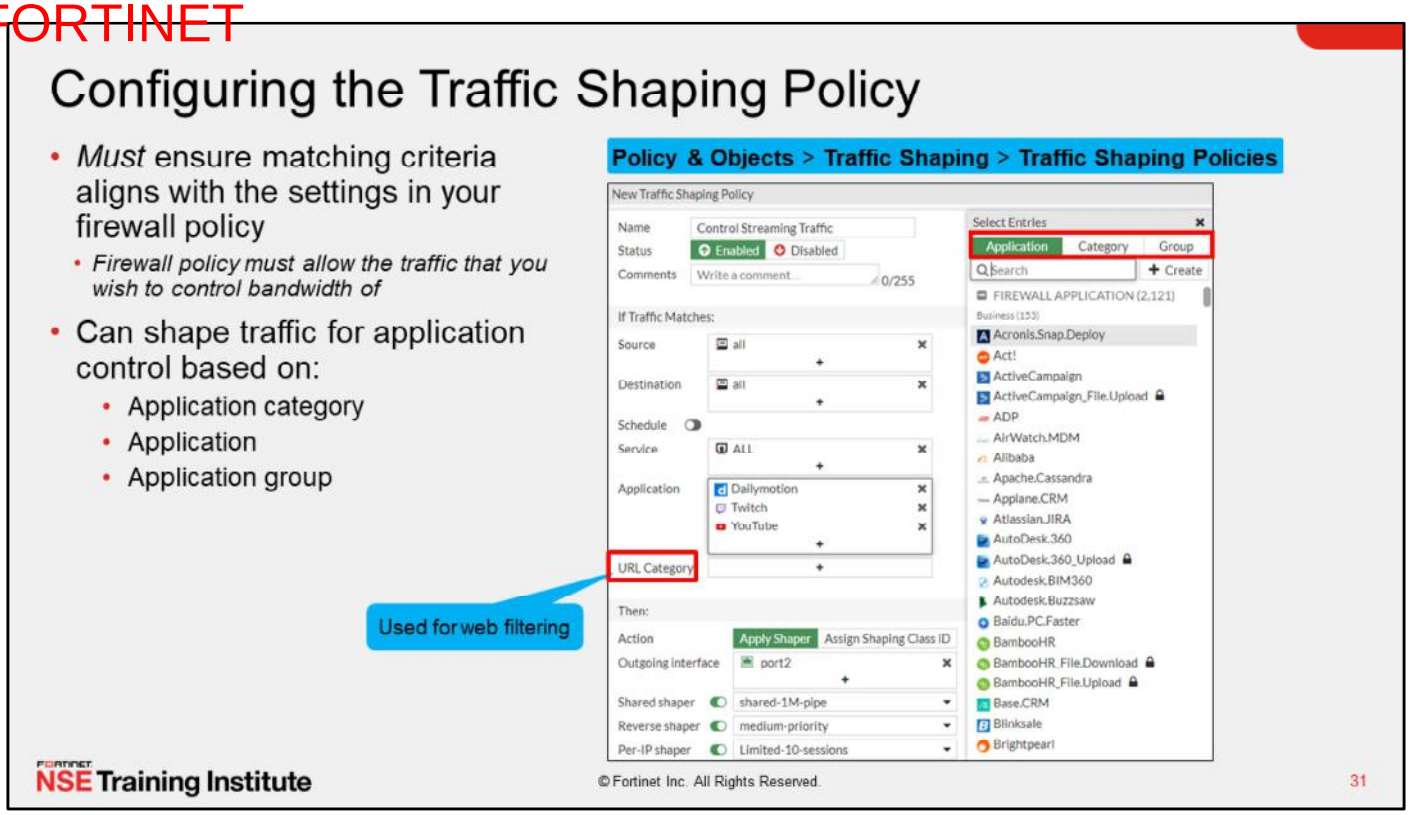

You can limit the bandwidth of an application category, application group, or specific application by configuring a traffic shaping policy. You can also apply traffic shaping to FortiGuard web filter categories and to the application group.

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.

NSE Training Institute

32

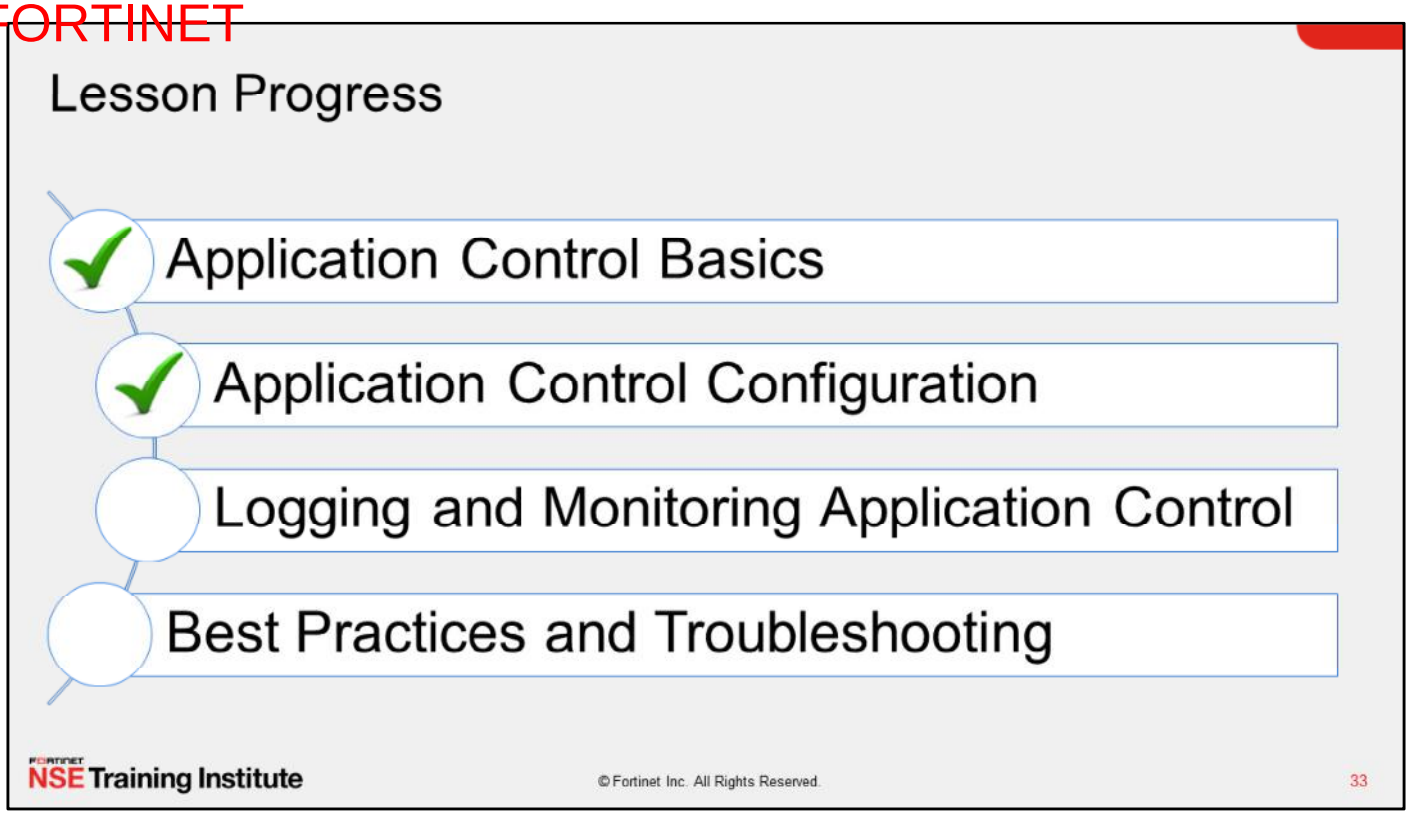

Good job! You now understand application control configuration.

Now, you will learn about logging and monitoring application control events.

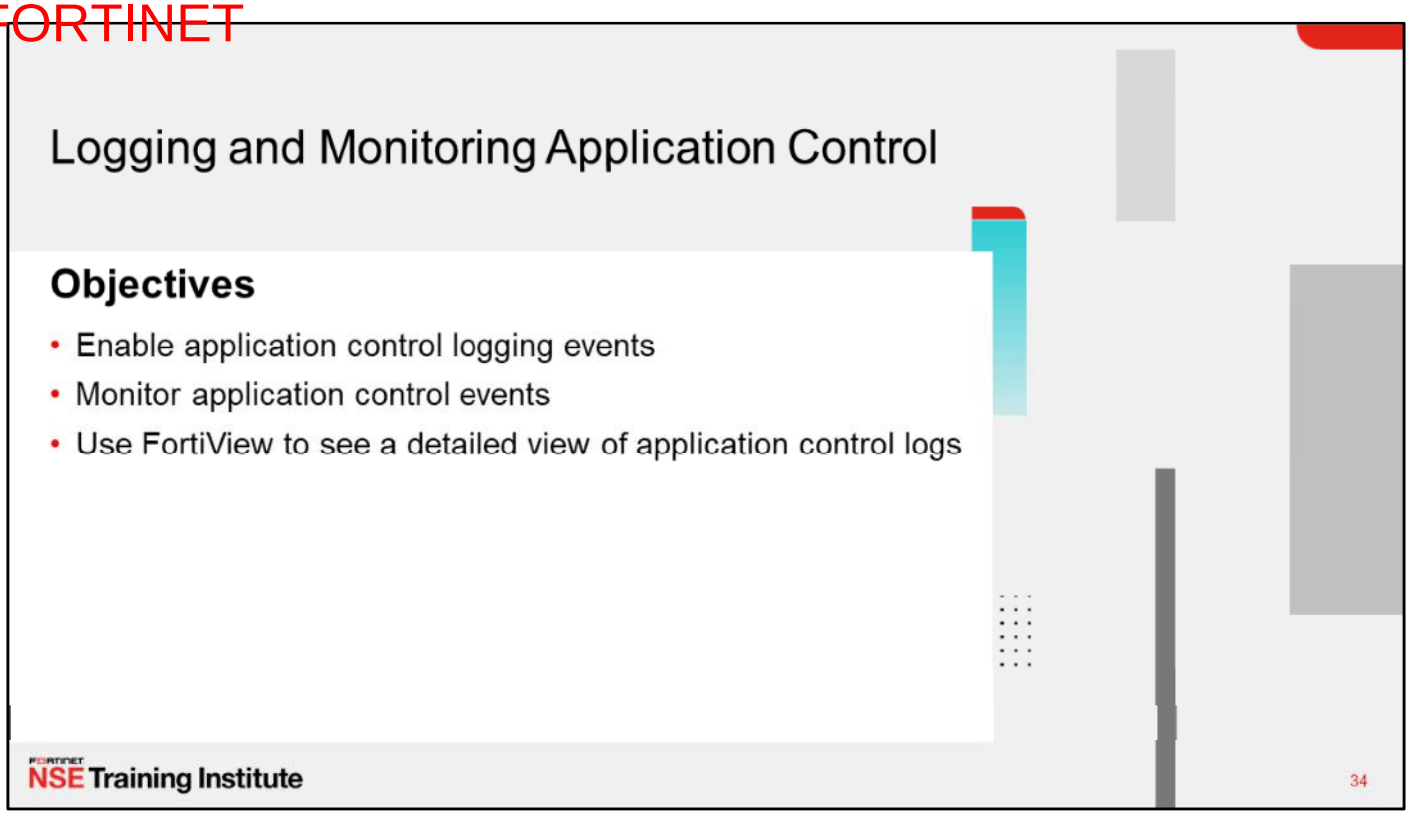

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control configuration, including reviewing application control logs, you will be able to effectively use and monitor application control events.

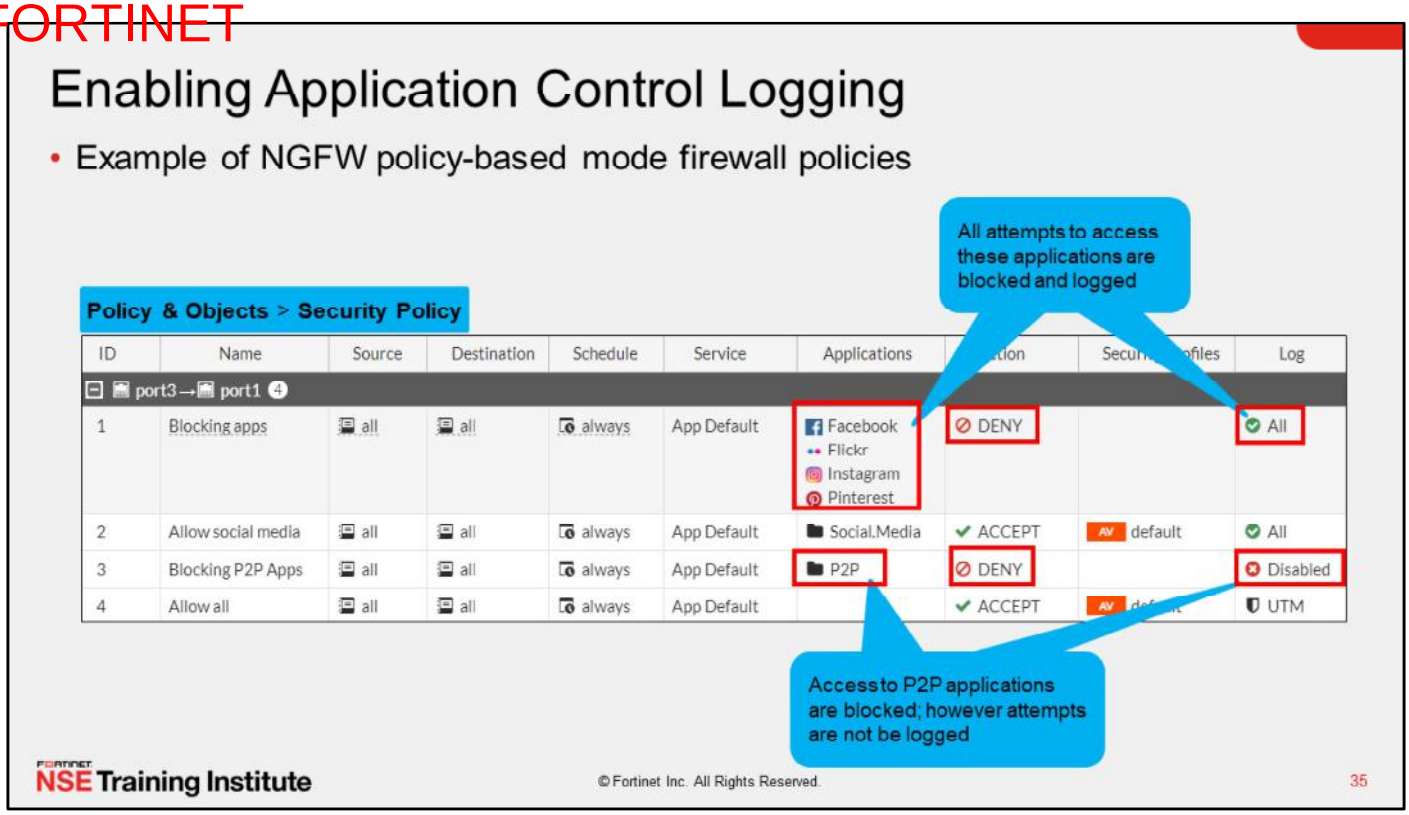

Regardless of which operation mode application control is configured in, you must enable logging on the security or firewall policy. When you enable the logging of security events or all sessions on a security or firewall policy, application control events are also logged. You must apply application 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.

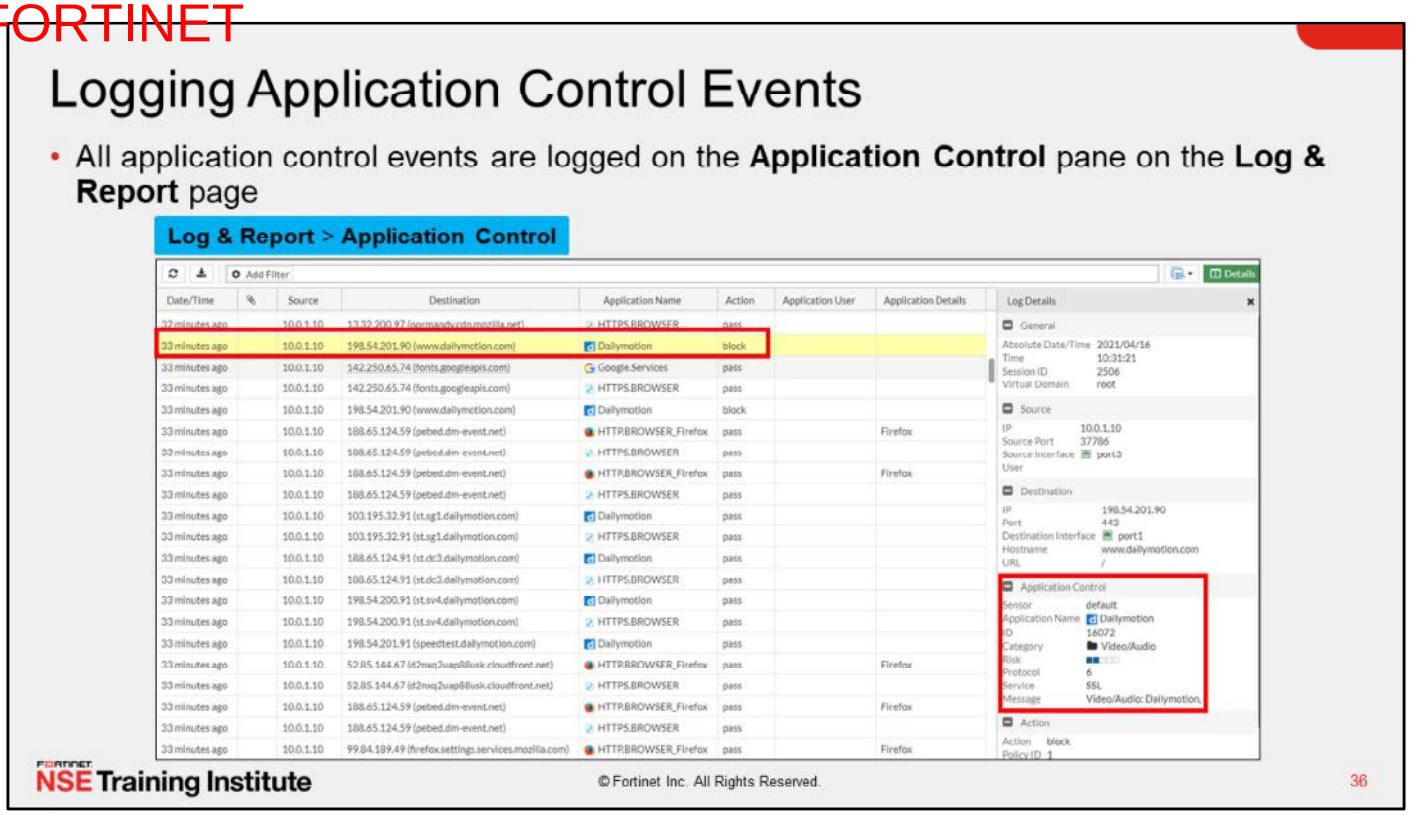

All application control events are logged on the **Application Control** pane on the **Log & Report** page. You can view details about individual logs by clicking on the log entry.

In the example shown on this slide, access to **Dailymotion** is blocked using the default application control profile. This information is available in the **Log Details** section, as well as information about the log source, destination, application, and action.

Note that this log message was generated by application control using a profile-based configuration. In an NGFW policy-based configuration, you will not find information such as application sensor name, because it does not apply. The remainder of the information and structure of the log message is the same for each log, regardless of which inspection mode FortiGate is operating in.

You can also view the details on the **Forward Traffic** logs pane. This pane is 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.

Application Control Events In Dashboard View

- Application control events are saved in a standalone dashboard on the **Top Applications** dashboard
  - Requires disk logging

Dashboard > Top Applications

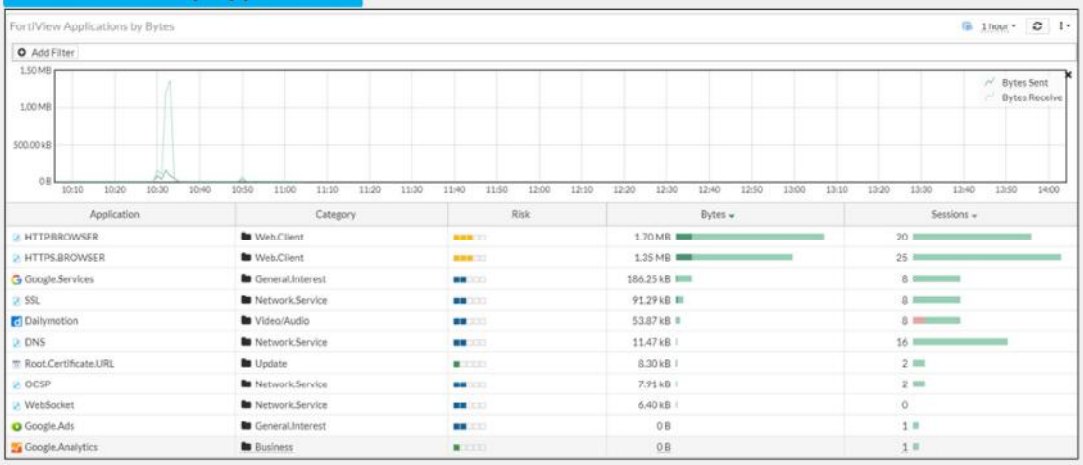

NSE Training Institute © Fortinet Inc. All Rights Reserved. 37

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

NSE Training Institute

38

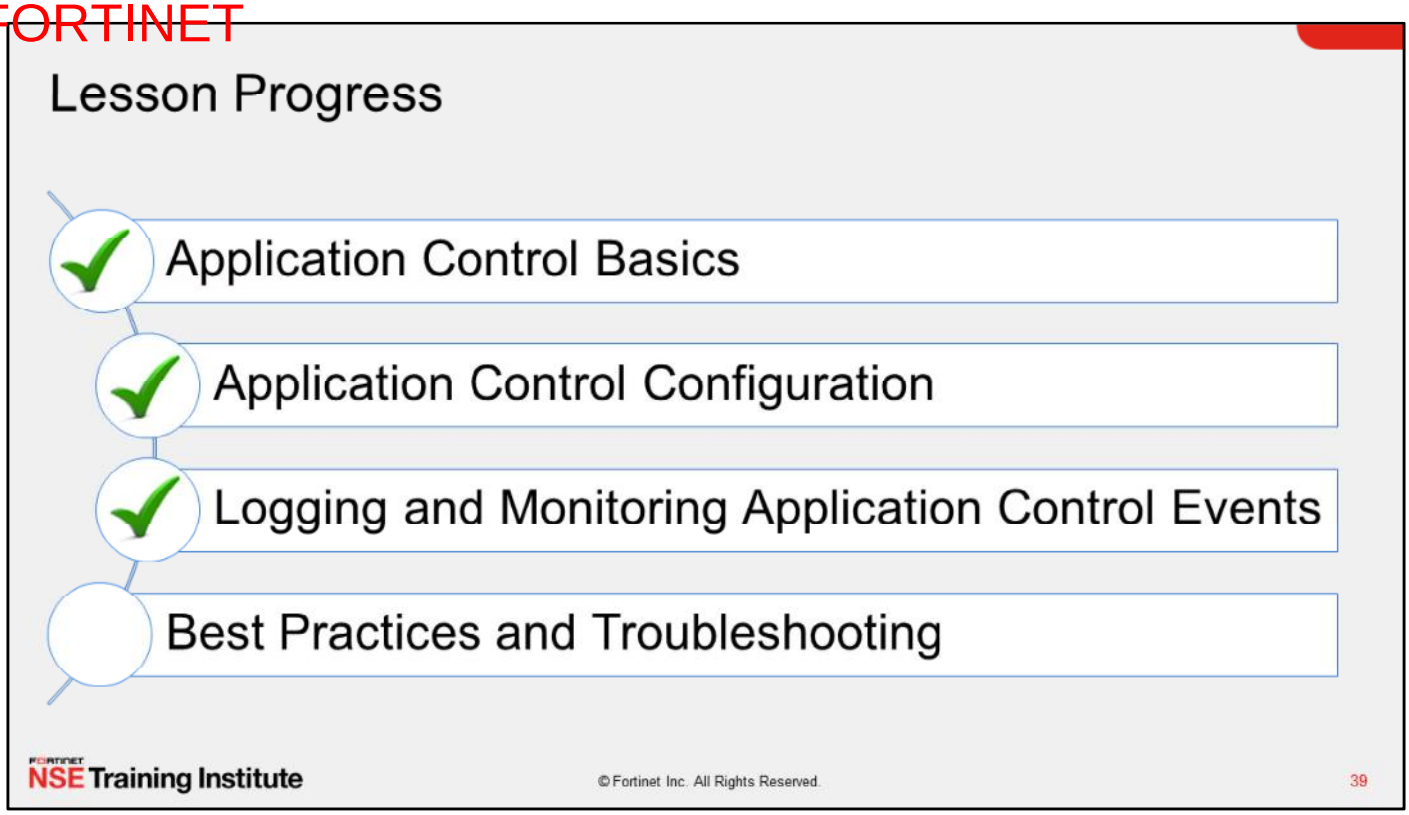

Good job! You now understand application control logging and monitoring.

Now, you will learn about application control best practices and troubleshooting.

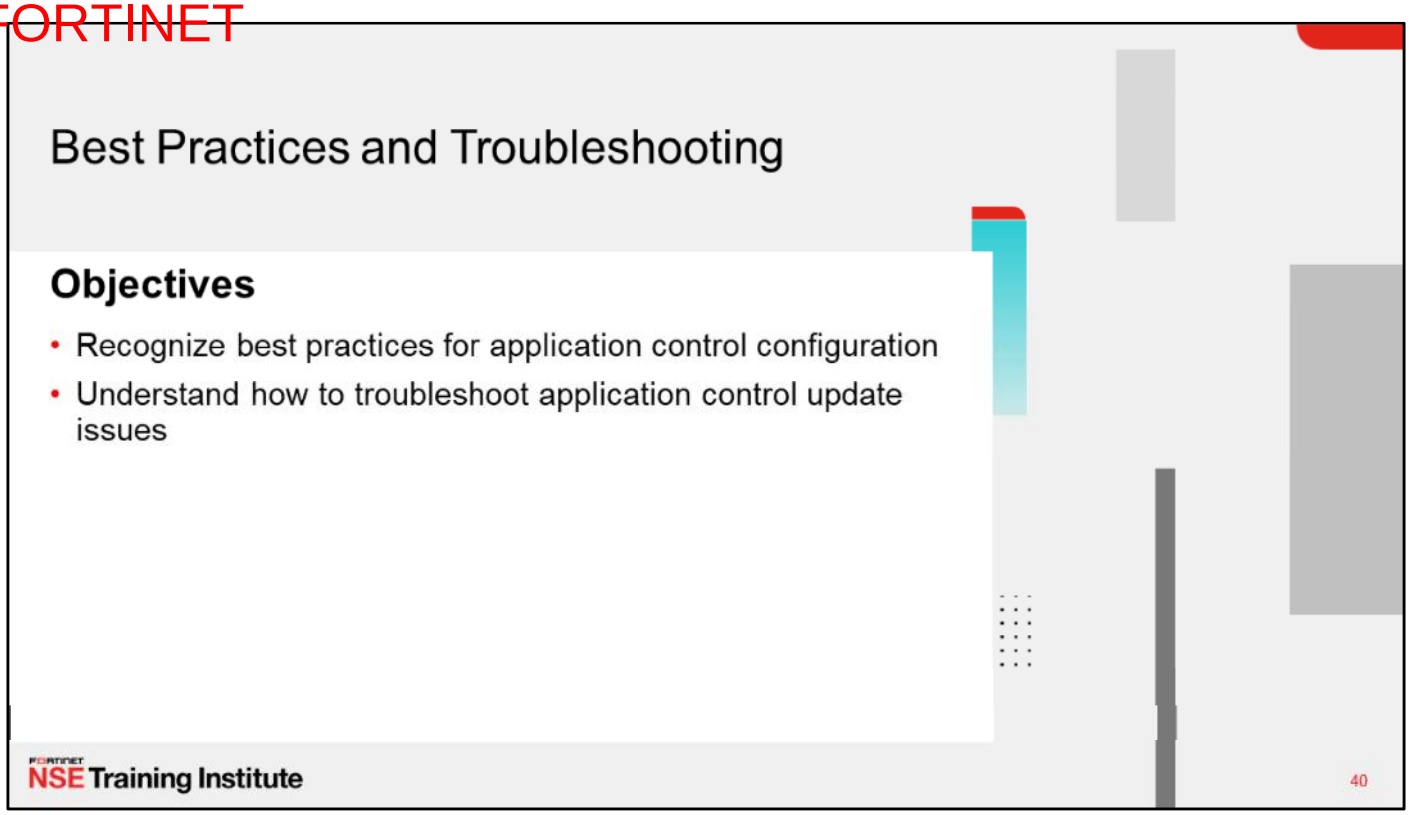

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in application control best practices and troubleshooting, you will be able to configure and maintain an effective application control solution.

Best Practices for Application Control

- Apply application control to only the traffic that requires it
  - Specify subnets (source, destination, or both) within the firewall policy, whenever possible
  - Don't apply application control to internal-to-internal traffic
- If using load balancing or failover internet connections, apply identical application control on all load balancing or redundant firewall policies
- Select **Deep-Inspection** instead of **Certificate-based** inspection as the SSL/SSH inspection method
- Use a FortiCloud account to save and view application control events in FortiView
  - FortiGate devices that don't have an internal disk for logging require FortiCloud logging to use FortiView
- Use hardware acceleration for application signature matching

NSE Training Institute 41

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.

Application Control Troubleshooting

- If FortiGuard has update issues, make sure that:
  - FortiGate has a stable connection to the internet
  - FortiGate is able to resolve DNS (`update.fortiguard.net`)
  - TCP port 443 is open
- Force FortiGate to check for new application control updates:

```
execute update-now
```

- Verify that the application control signatures database version is up-to-date with the FortiGuard website

System > FortiGuard

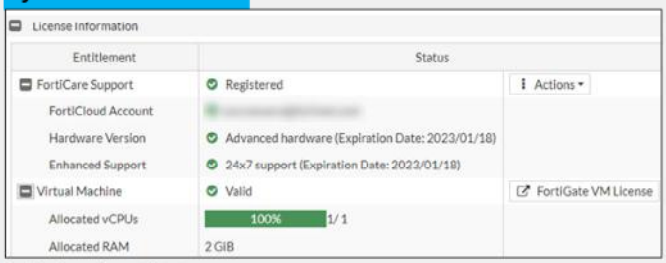

NSE Training Institute © Fortinet Inc. All Rights Reserved. 42

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

NSE Training Institute

43

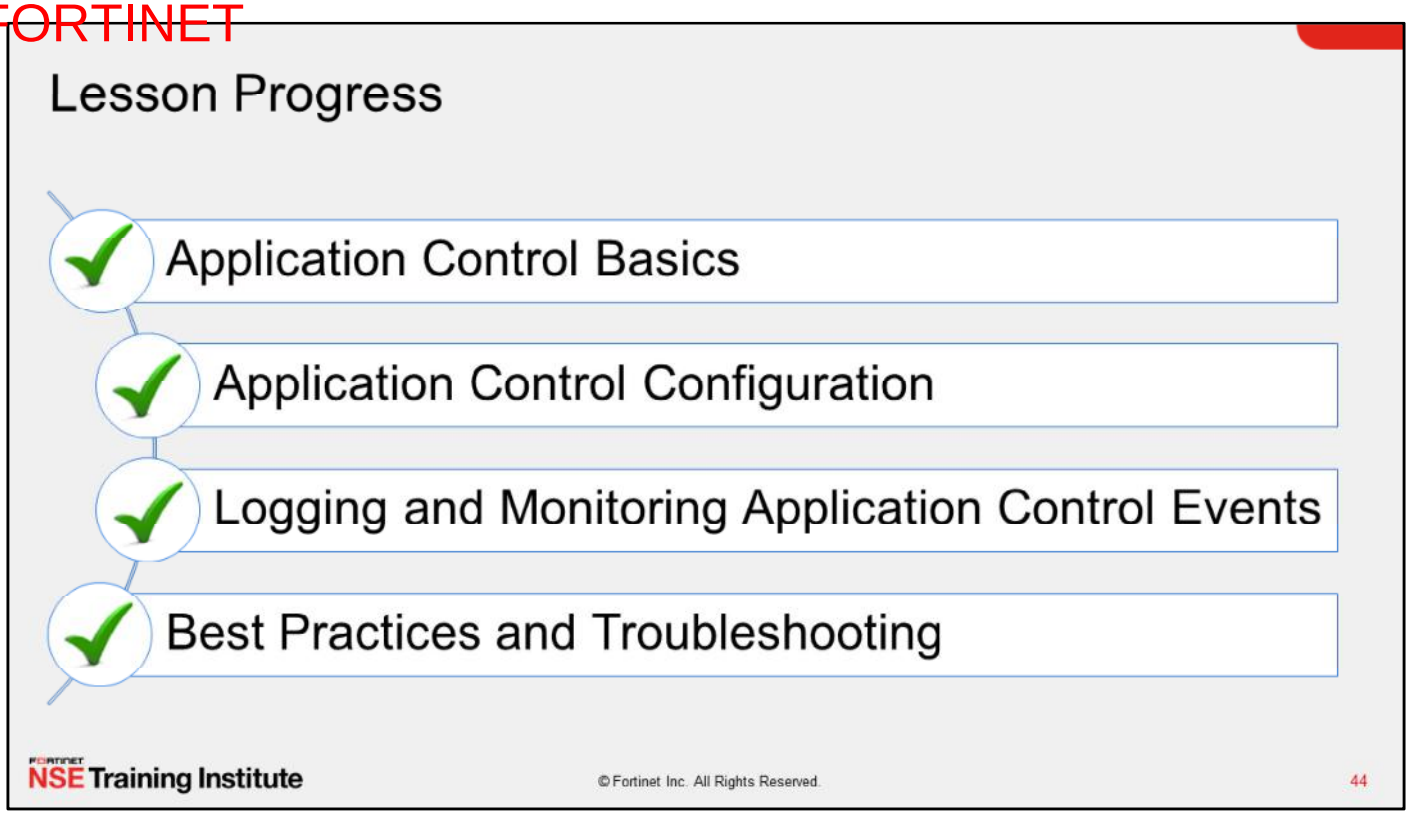

Congratulations! You have completed this lesson.

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

Review

- ✓ Understand application control
- ✓ Detect types of applications
- ✓ Understand FortiGuard application control services
- ✓ Use application control signatures
- ✓ Configure application control in profile mode
- ✓ Configure application control in NGFW policy mode
- ✓ Use the application control traffic shaping policy
- ✓ Enable application control logging events
- ✓ Monitor application control events
- ✓ Use the dashboard to see a detailed view of application control logs
- ✓ Recognize best practices for application control configuration
- ✓ Understand how to troubleshoot application control update issues

NSE Training Institute © Fortinet Inc. All Rights Reserved. 45

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.

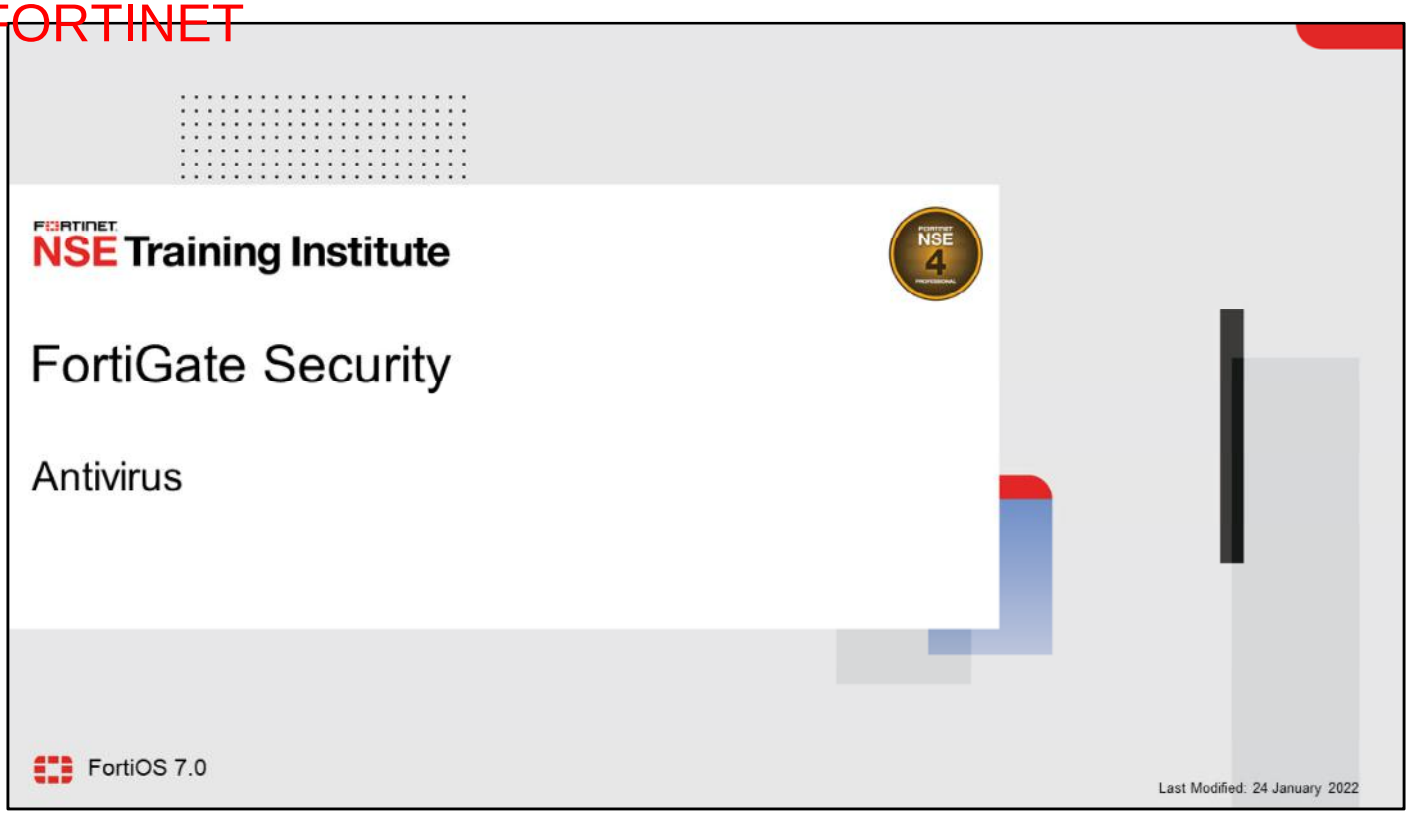

In this lesson, you will learn how to use FortiGate to protect your network against viruses.

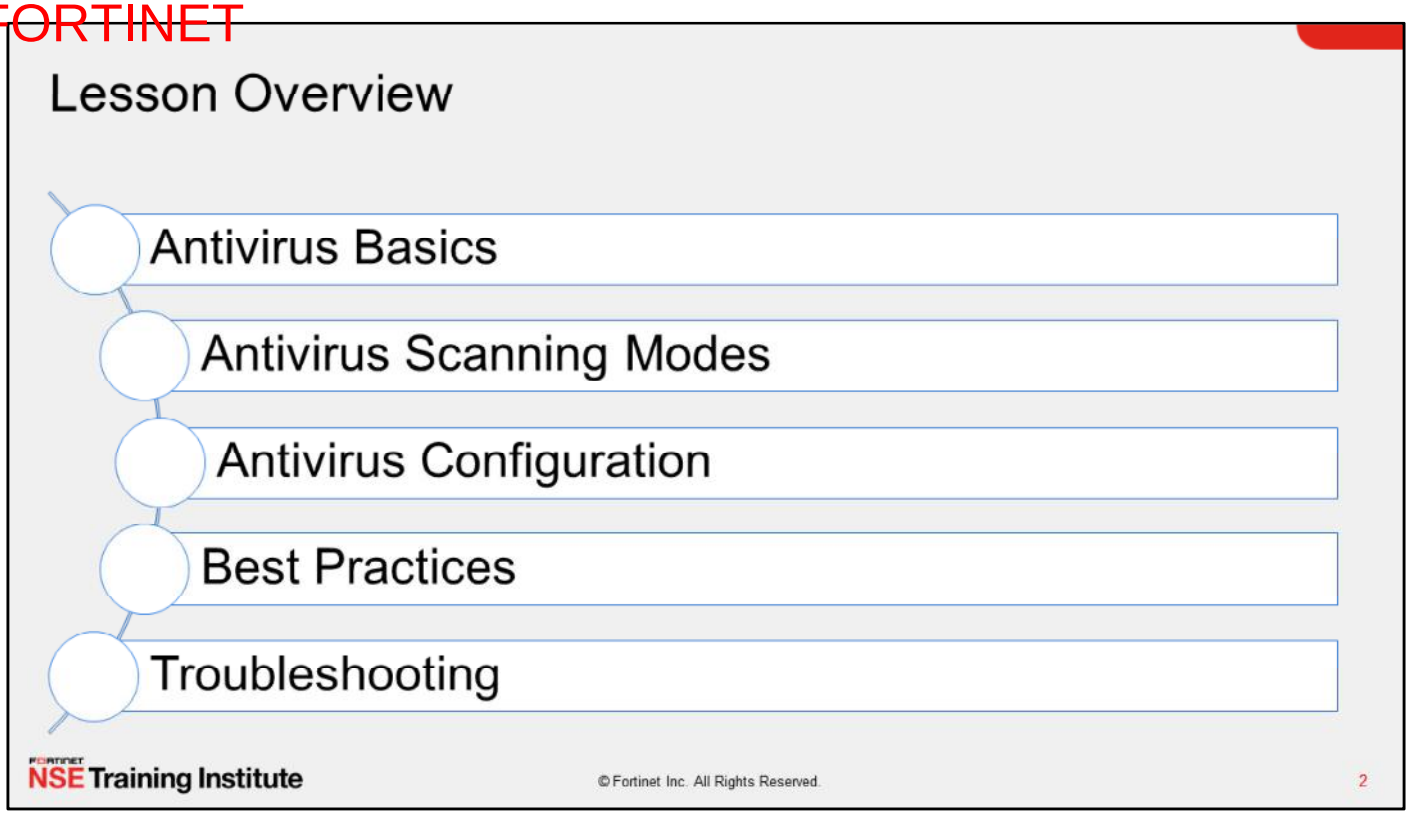

In this lesson, you will learn about the topics shown on this slide.

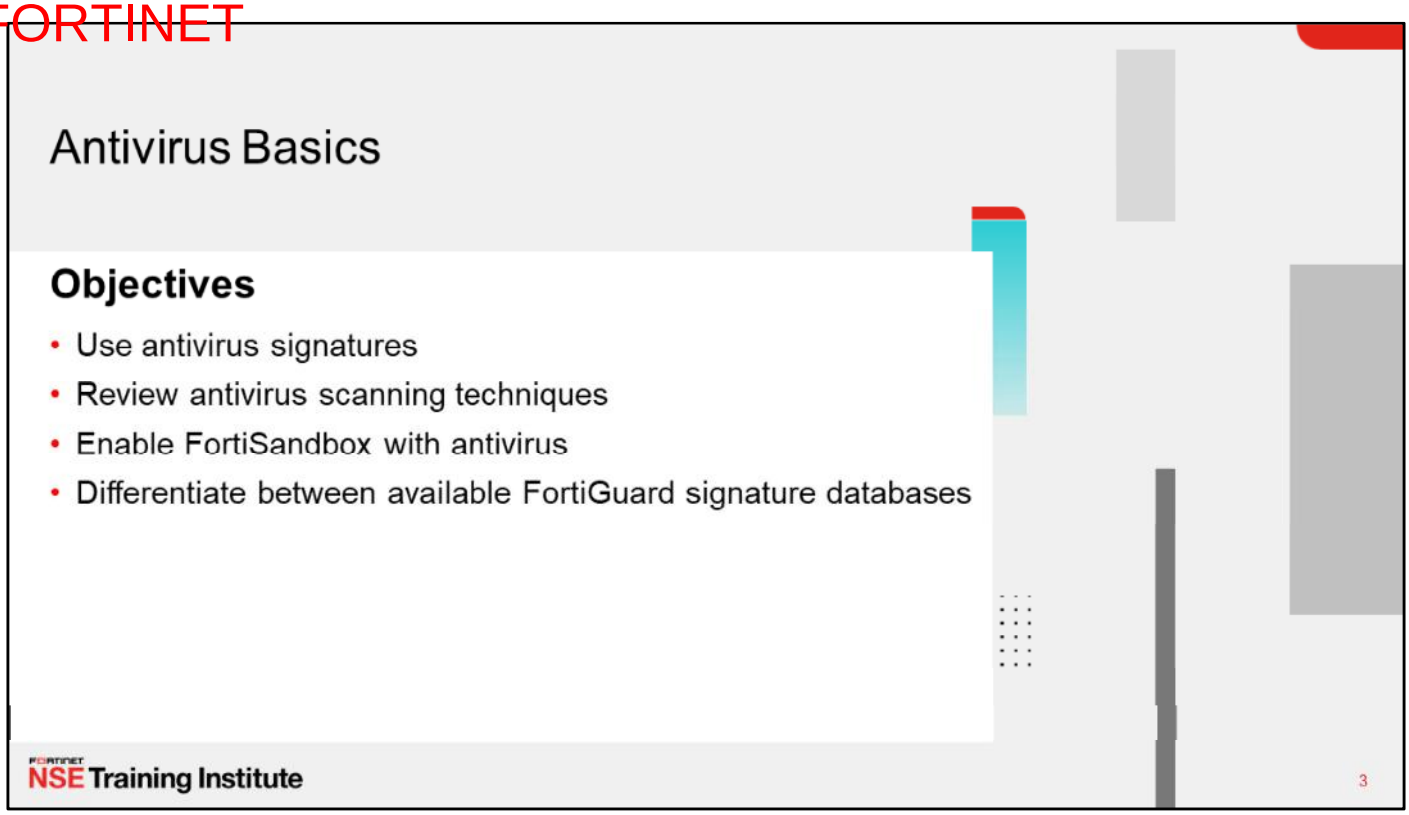

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus basics, you will be able to understand and apply antivirus on FortiGate.

What is Antivirus and How Does It Work?

- Antivirus is a database of virus signatures that is used to identify malicious code
- Virus names: `<vector>/<pattern>`
  - Example: `W32/Kryptik.EMT!tr`
    - `<vector>` for a virus will always be the same, but vendors assign different IDs for `<pattern>`
- To detect a virus, the antivirus engine must match file with pattern `<signature>`
- Each vendor uses different detection engines and signatures, such as:
  - MD5
  - CRC
  - Combinations of file attributes
  - Binary values in some areas
  - Encryption keys
  - Parts of code

NSE Training Institute © Fortinet Inc. All Rights Reserved. 4

An antivirus is a database of virus signatures that is used to identify infections. During an antivirus scan, in order to be detected as a virus, the virus must match a defined pattern called a *signature*.

Different vendors assign different names to the same virus. All vendors use the attack vector designation in the virus name. The vector comes at the beginning of the virus name. Some examples include:
- W32, which represents 32-bit Windows
- W64, which represents 64-bit Windows
- JS, which represents JavaScript (which is cross-platform)

Some vendors also use a pattern as part of the virus name. Some patterns detect only one virus per pattern. Other patterns are more flexible and can detect multiple viruses per pattern. The pattern that the vendor uses depends on the vendor's engine.

Host-based antivirus software, such as FortiClient, can help at the host level; however, host-based antivirus software cannot be installed on routers. Also, guest Wi-Fi networks and ISP customers might not have antivirus software installed.

So, how can you protect guest networks, ISP customers, and your own network from malware threats?

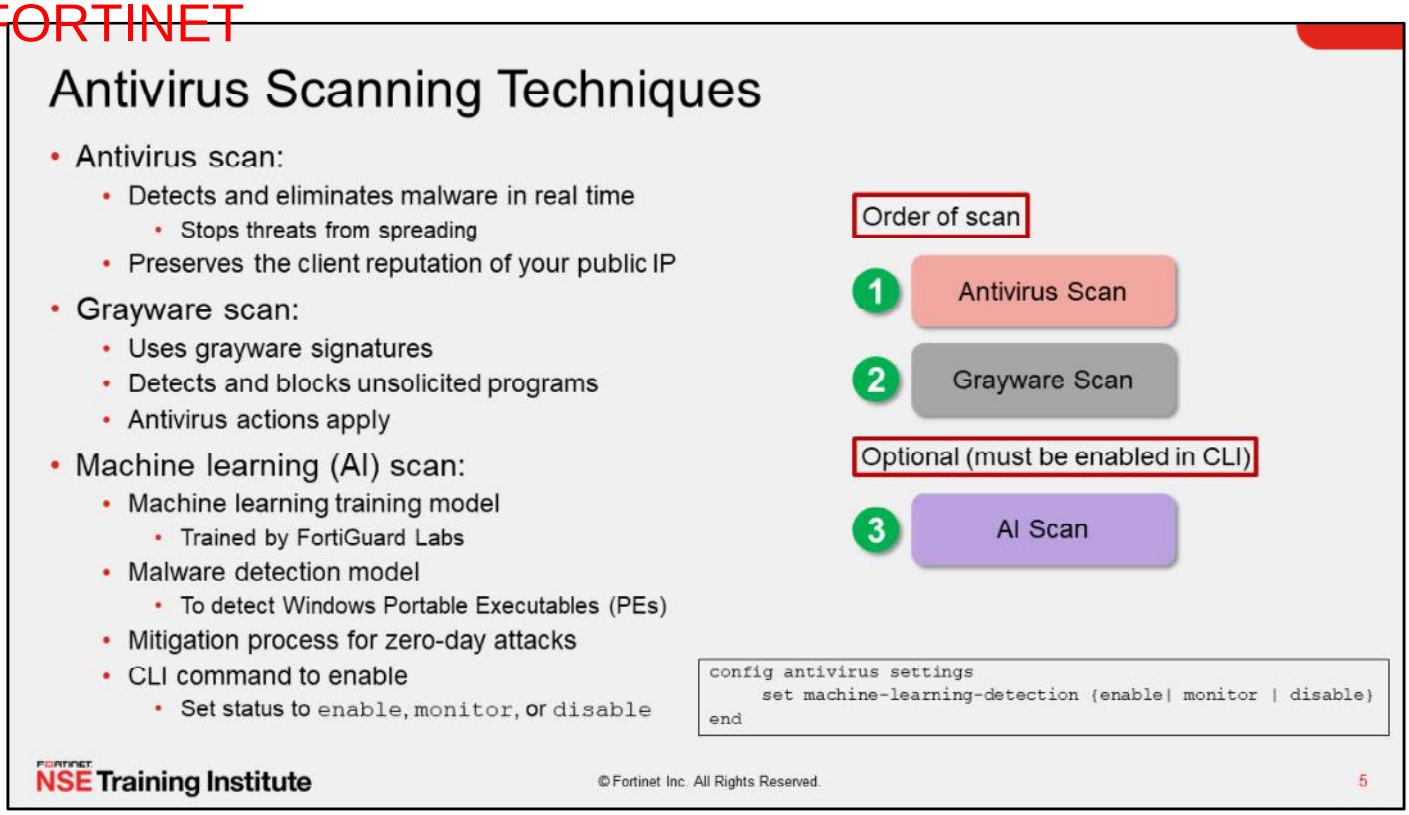

Like viruses, which use many methods to avoid detection, FortiGate uses many techniques to detect viruses. These detection techniques include:

- 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. By default, when the AI engine detects a new virus, it logs the file as **Suspicious** but does not block it. You can choose whether to block or allow suspicious files.

The AI scan is an optional feature that must be enabled in the CLI. You can configure the action for the AI scan to `enable`, `monitor`, or `disable` using the CLI command in the antivirus settings.

If all antivirus features are enabled, FortiGate applies the following scanning order: antivirus scan, followed by grayware scan, followed by AI scan.

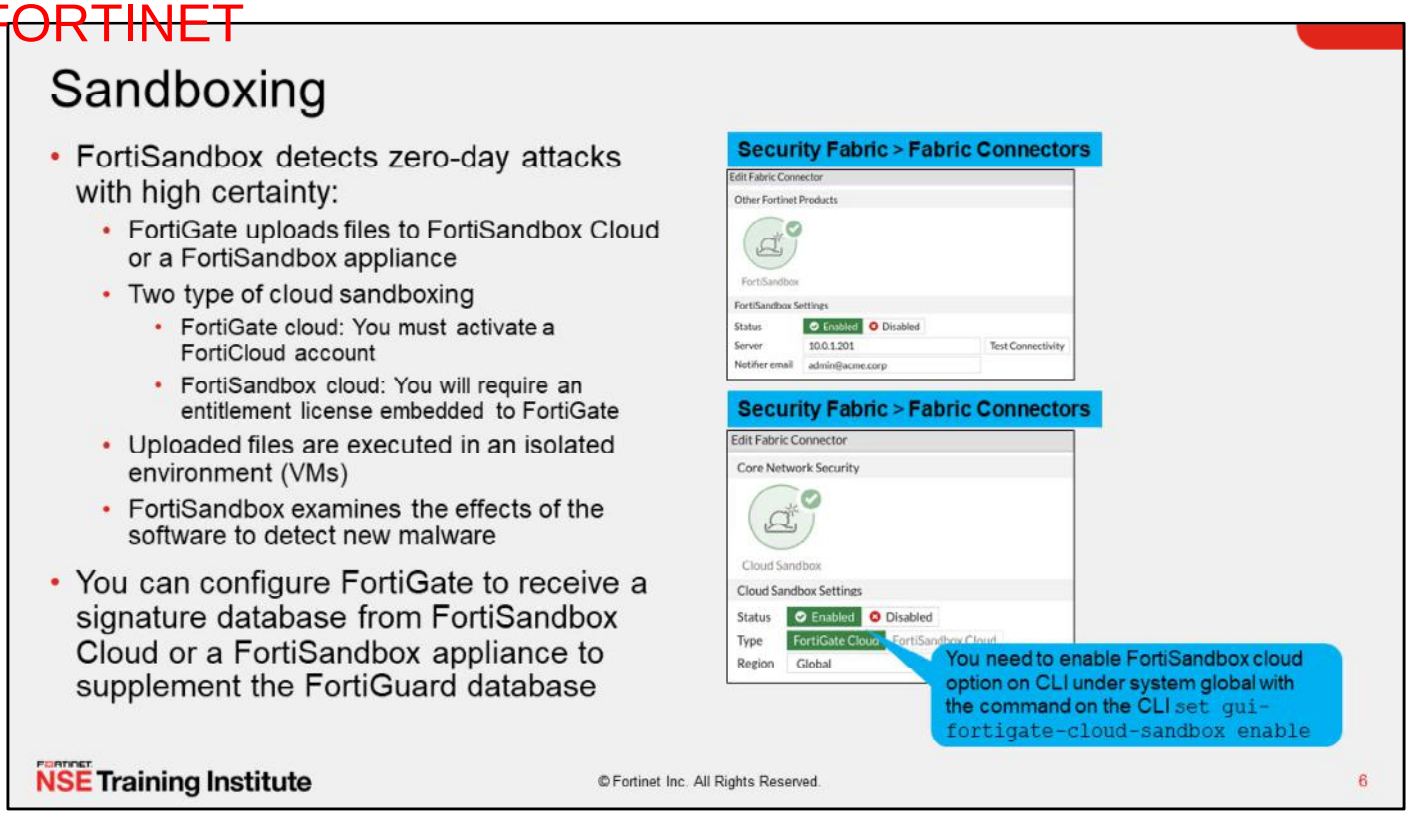

What if AI scans are too uncertain? What if you need a more sophisticated, more certain way to detect malware and find zero-day viruses?

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.

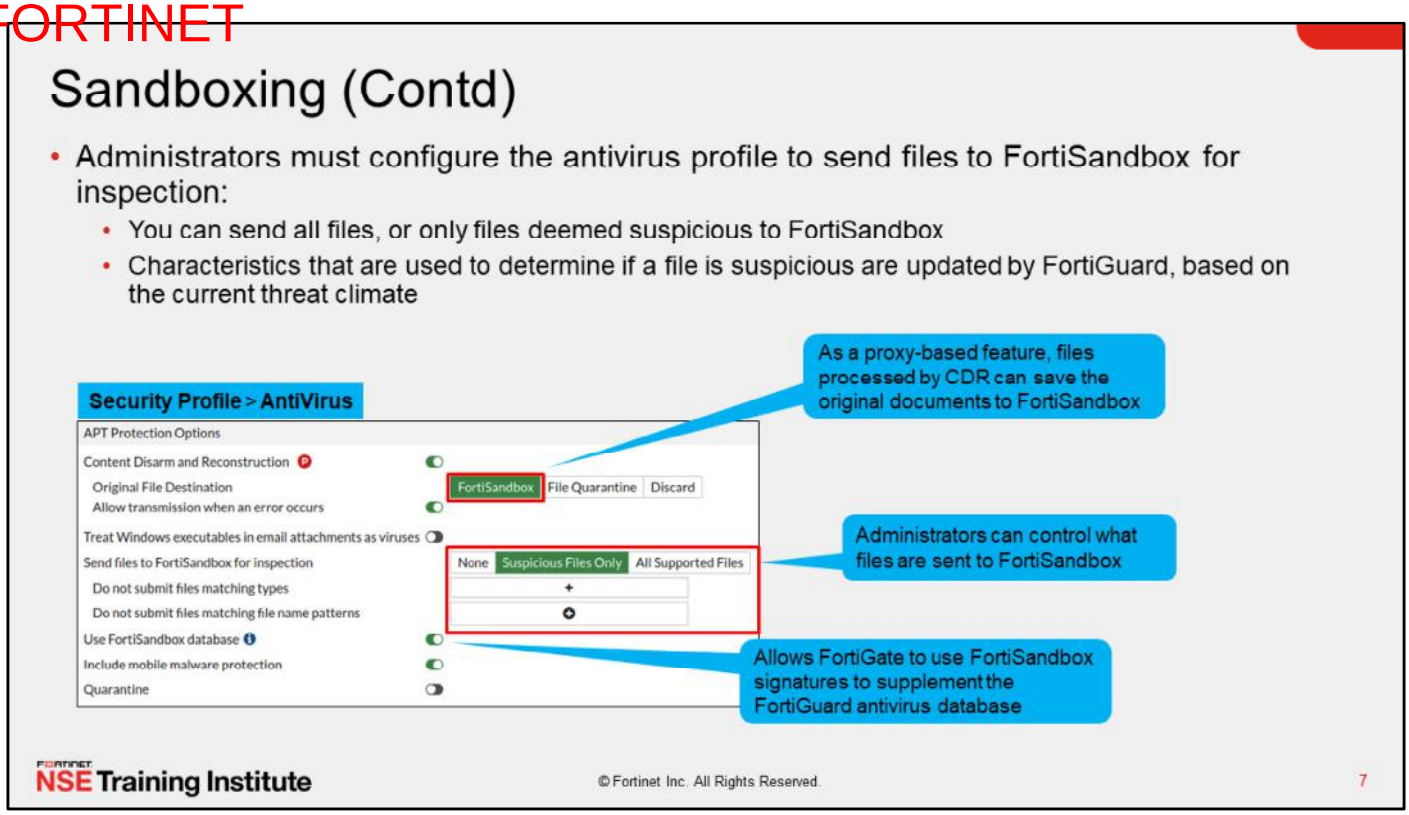

FortiOS is smart when it comes to determining what files are sent to FortiSandbox. One feature FortiOS uses for this is content disarm and reconstruction (CDR), a proxy-based feature that you will learn more about later. When CDR processes files, the original documents can be saved to FortiSandbox.

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.

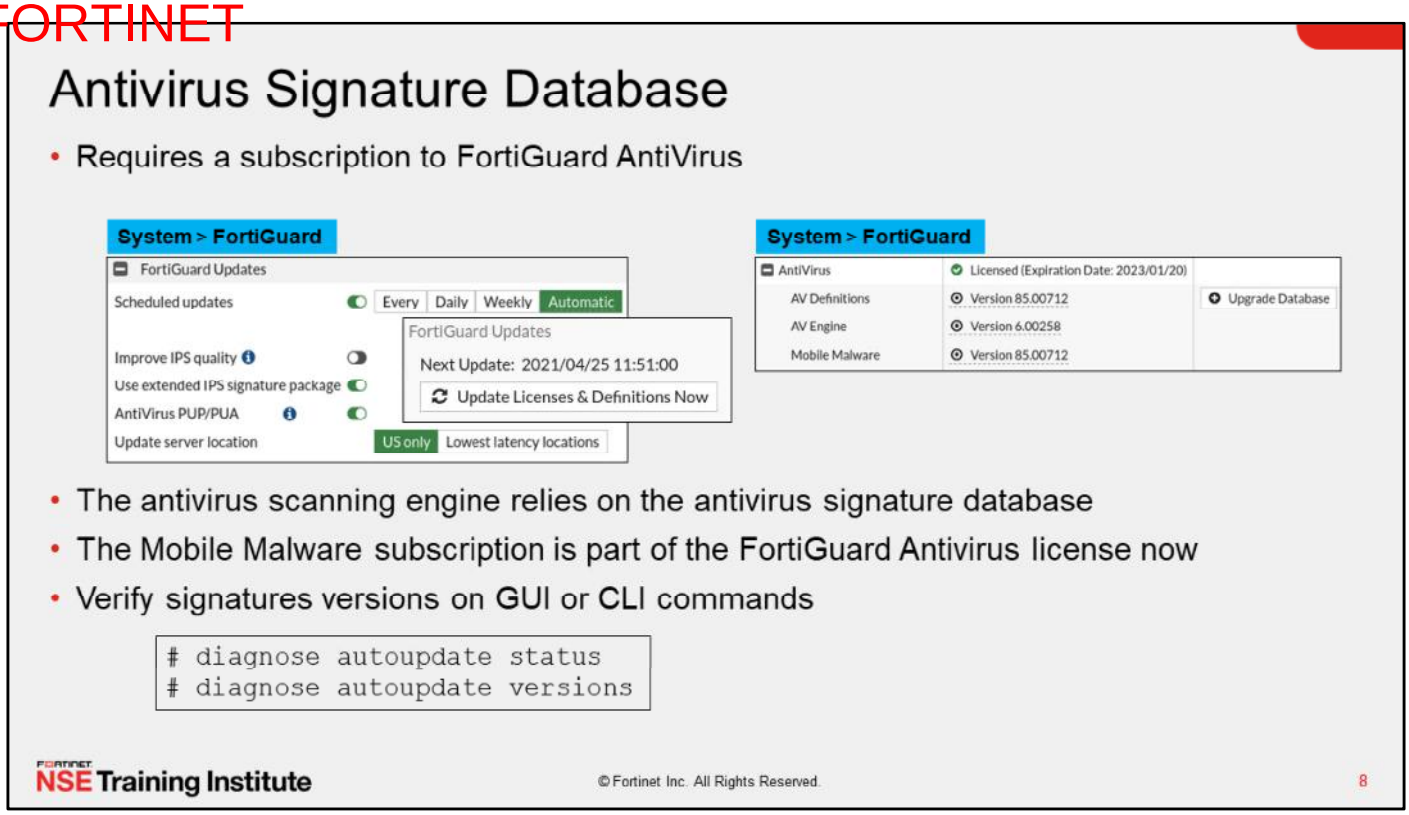

Scheduled updates allow you to configure scheduled updates at regular intervals, such as hourly, daily, weekly, or automatically within every hour. You can also enable **AntiVirus PUP/PUA**, which allows antivirus grayware checks for potentially unwanted programs and applications.

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.

Antivirus Signature Database (Contd)

- FortiGuard antivirus databases:
  - Extended: Includes common and additional recent non-active viruses
    - Available on all models
    - The default antivirus database setting
  - Extreme: Includes extended plus additional dormant viruses
    - Extreme is only available on select FortiGate models

- Choosing an antivirus signature database (CLI only)

```
config antivirus settings
  set use-extreme-db {enable | disable}
end
```

NSE Training Institute

9

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

- Content disarm and reconstruction
  - 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

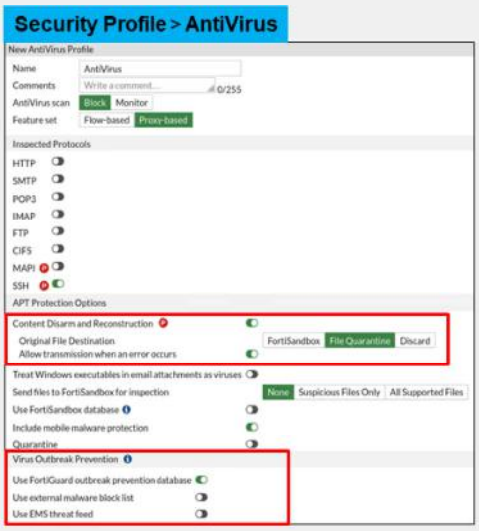

NSE Training Institute

10

Content disarm and reconstruction (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, then the original file is sent to FortiSandbox for inspection. The client can download the original file by logging in to the FortiSandbox.

Virus outbreak prevention: 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 list can be in the forms of MD5, SHA1, and SHA256 hashes, and are 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

NSE Training Institute

11

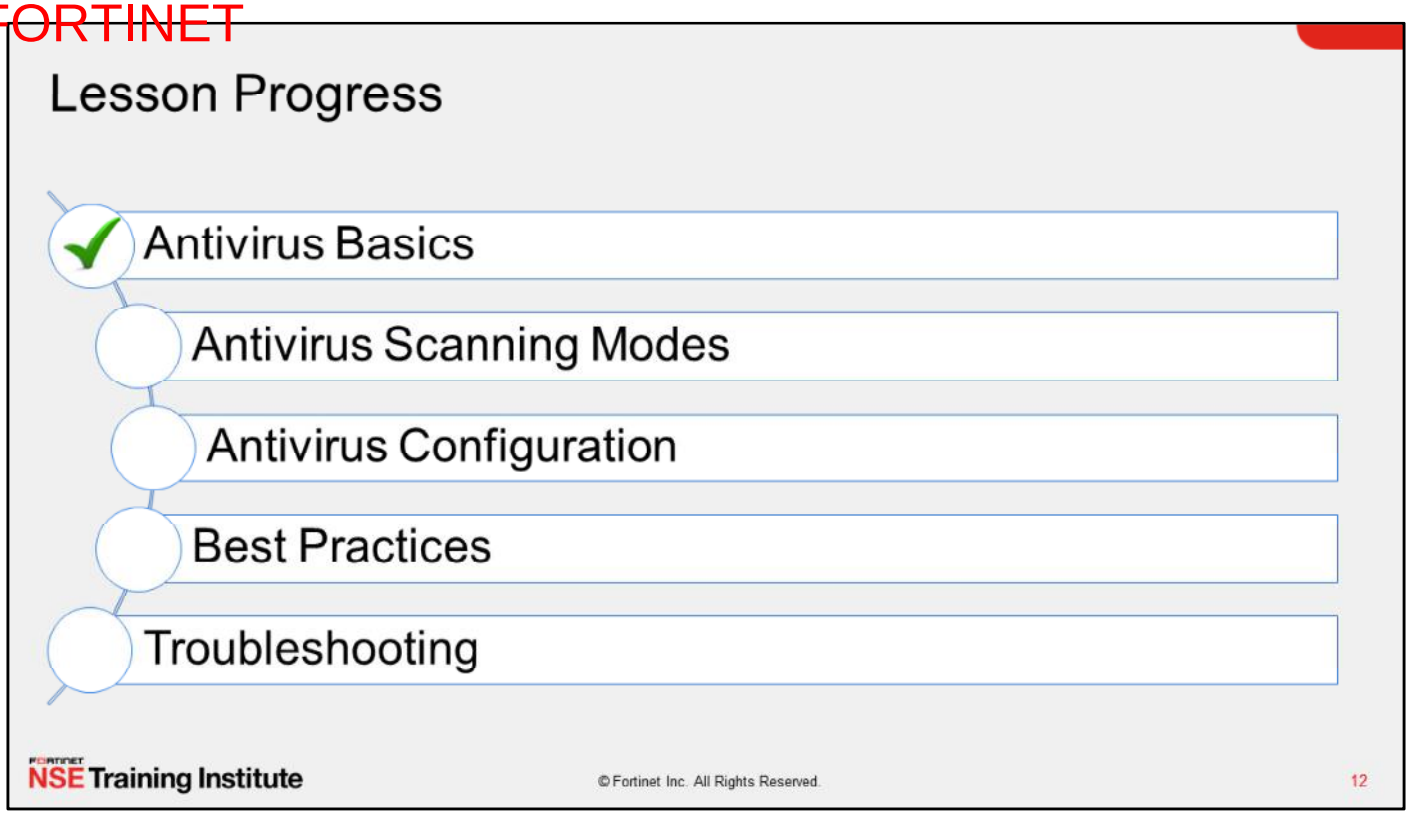

Good job! You now understand the basics of antivirus functionality.

Now, you will learn about antivirus scanning modes.

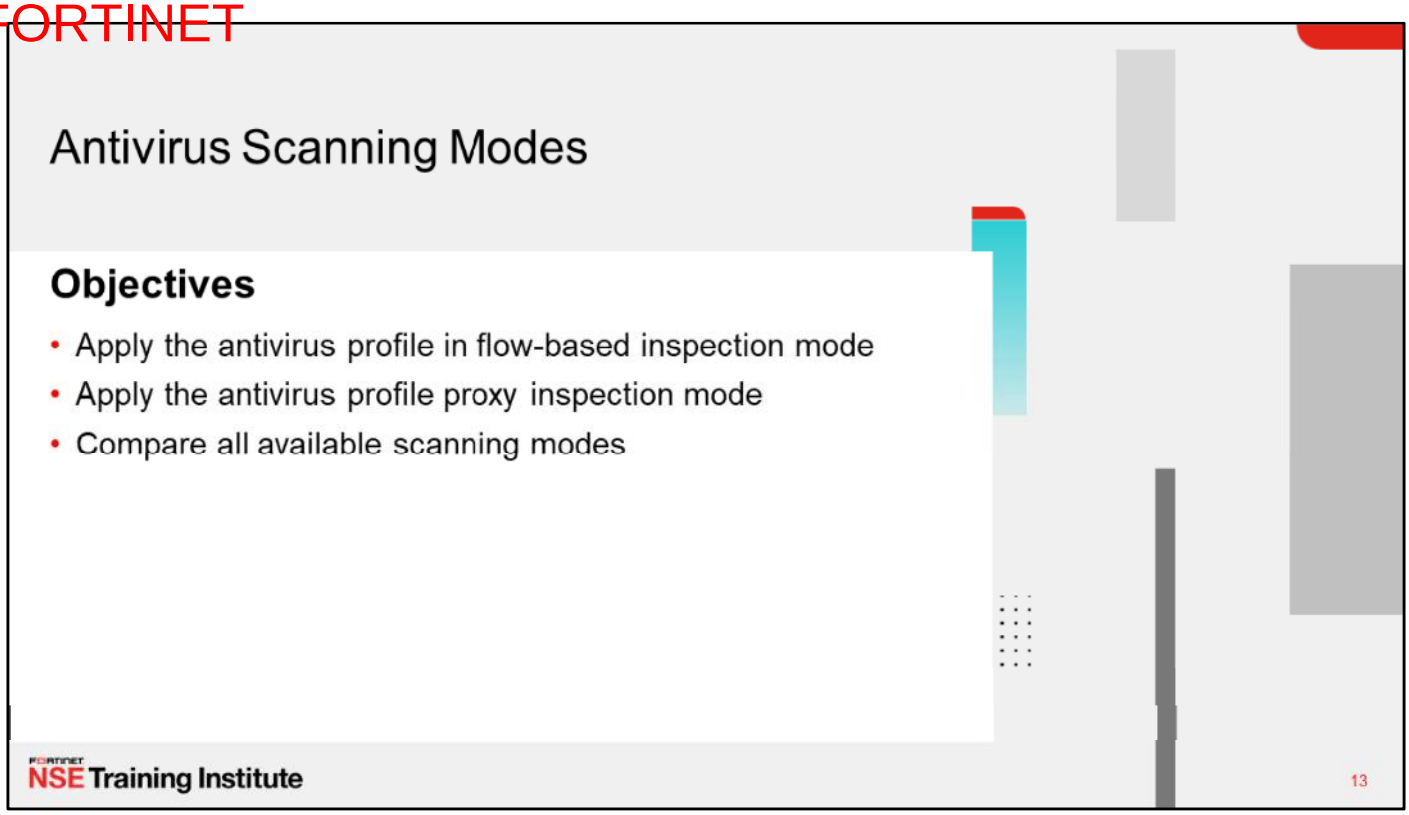

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in all antivirus scanning modes available in FortiOS, you will be able to use the antivirus profile in an effective manner.

Flow-Based Inspection Mode

- Uses the extended antivirus database by default
  - Extreme database on certain FortiGate models—depending on the CLI settings
- Optimized performance compared to proxy-based scan
  - Proxy-based offers two scanning modes: default scanning and legacy scanning
  - Flow-based is designed to use a hybrid of proxy-based scanning modes
- FortiGate buffers the whole file, but transmits to the client simultaneously
  - When the *last* packet arrives, the AV engine starts the scan
    - Files bigger than buffer size are not scanned—can enable logging of these files
    - Packets are not delayed by scan—*except last packet*
  - Lower perceived latency—data loads faster
- If a virus is detected, the last packet is dropped and the connection is reset
- If an identical request is made, the block replacement page is inserted immediately

NSE Training Institute © Fortinet Inc. All Rights Reserved. 14

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.

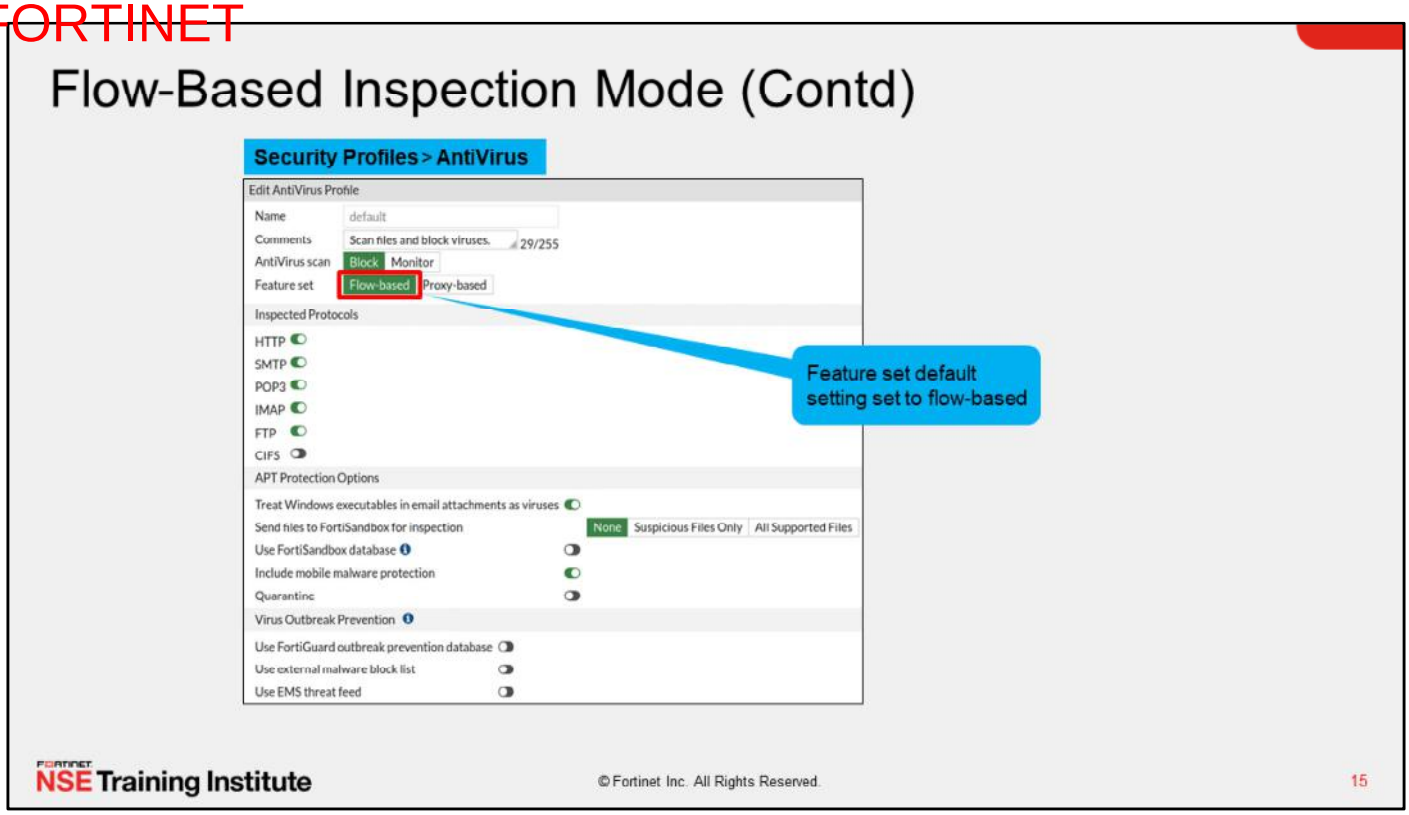

This slide shows an example of the antiVirus profile operating in flow-based inspection mode. By default, **Feature set** is set to **Flow-based**.

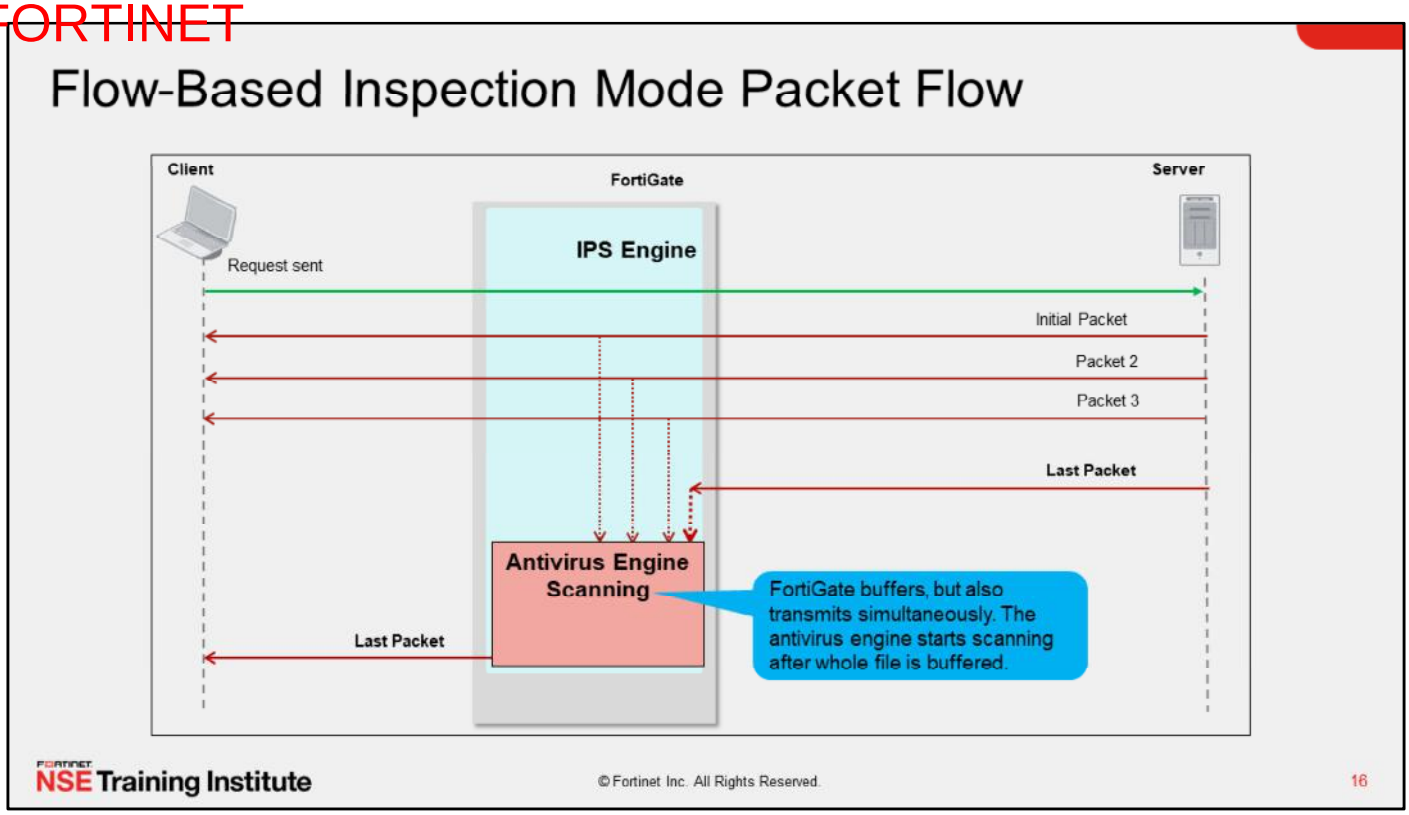

As you can see on this slide, the client sends a request and starts receiving packets immediately, but FortiGate also caches those packets at the same time. When the last packet arrives, FortiGate caches it and puts it on hold. Then, the IPS engine extracts the payload of the last packet, assembles the 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.

Proxy Inspection Mode

- Uses extended or extreme antivirus database
- Buffers the whole file
  - Antivirus engine starts scanning after the end of the file is detected
    - Files bigger than buffer size are not scanned—can configure to pass or block
  - Packets sent to the client after scan finishes—*client must wait*
  - Highest perceived latency
- Provides granularity over performance
- Weighted towards being more thorough and easily configurable
- Displays a block message immediately if a virus is detected
- Stream-based scanning supports FTP, SFTP, and SCP
  - Optimizes memory utilization for large archive files by decompressing and scanning them on the fly
  - Viruses are detected even if they are in the middle or end of the large files

NSE Training Institute © Fortinet Inc. All Rights Reserved. 17

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.

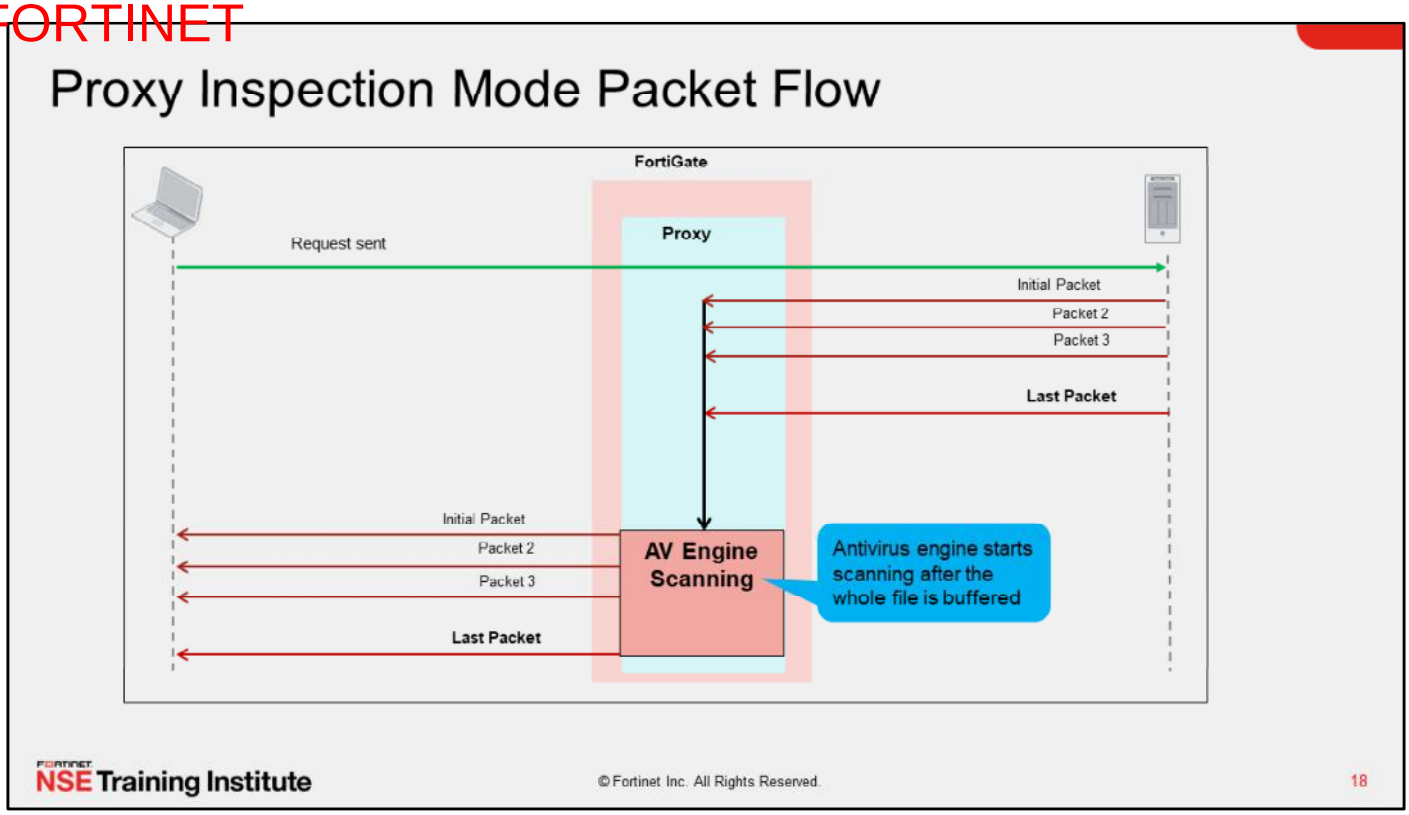

With a proxy inspection mode scan, the client sends a request and FortiGate starts buffering the whole file, then sends it to the antivirus engine for scanning. If the file is clean (without any viruses), FortiGate starts transmitting the file to the end client. If a virus is found, no packets are delivered to the end client and the proxy sends the replacement block message to the end client.

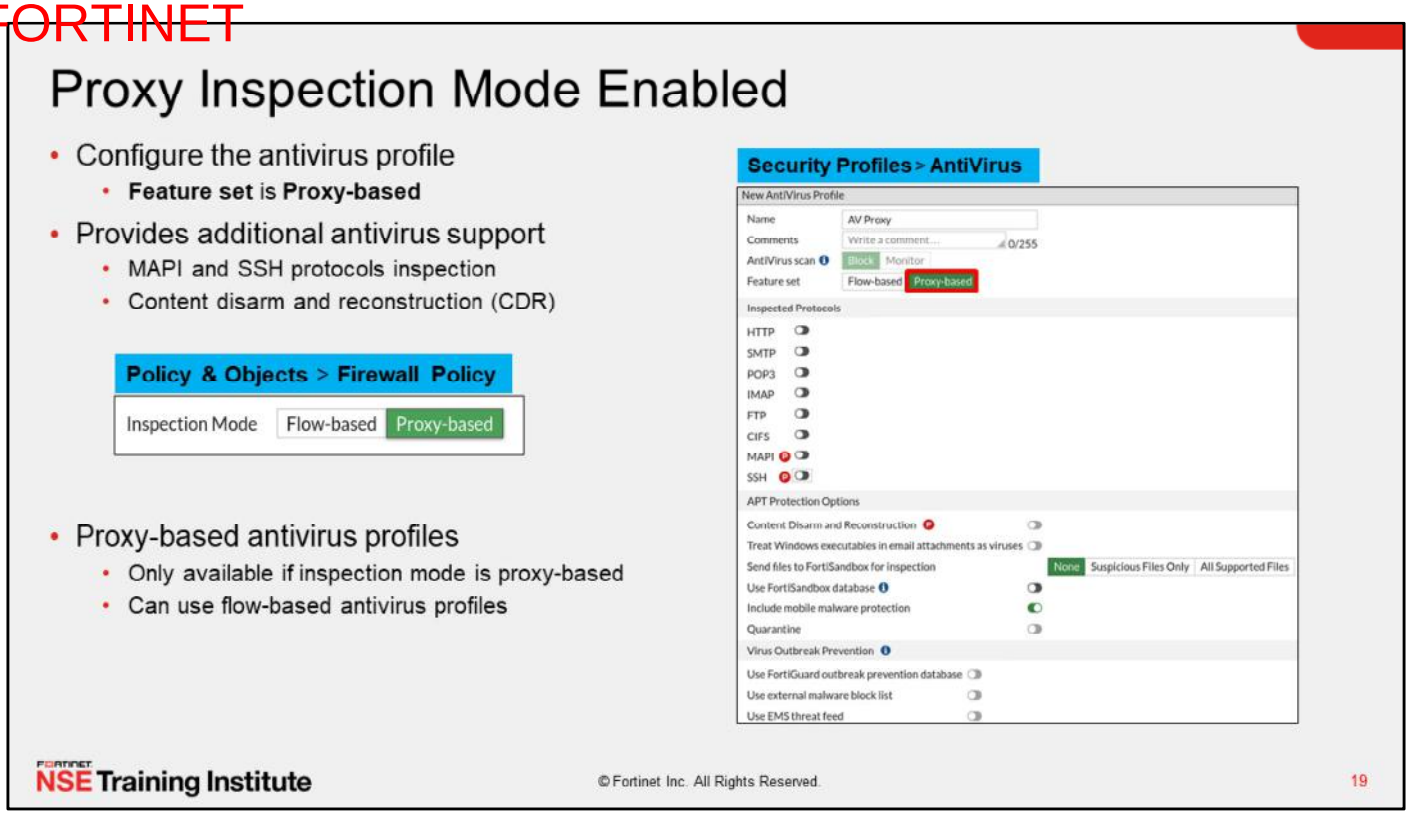

Applying a proxy-based antivirus profile requires two sections in FortiGate configuration to use non-default settings:

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**.

Antivirus Scanning Modes Comparison

	Flow-based (hybrid)	Proxy-based
Catching Rate	Highest	Highest
Sandbox Support	Yes	Yes
Advanced Heuristic	Yes	Yes
Memory	High	High
Perceived Latency	High	Highest
MAPI, NNTP Scanning	No	Yes
SMB Scanning	Yes	No
HTTP, FTP, IMAP, POP3, SMTP Scanning	Yes	Yes
Use FortiSandbox Database	Yes	Yes
Use Mobile Malware Protection Service	Yes	Yes

NSE Training Institute

20

This slide provides comparison of the different antivirus scanning modes.

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

NSE Training Institute

21

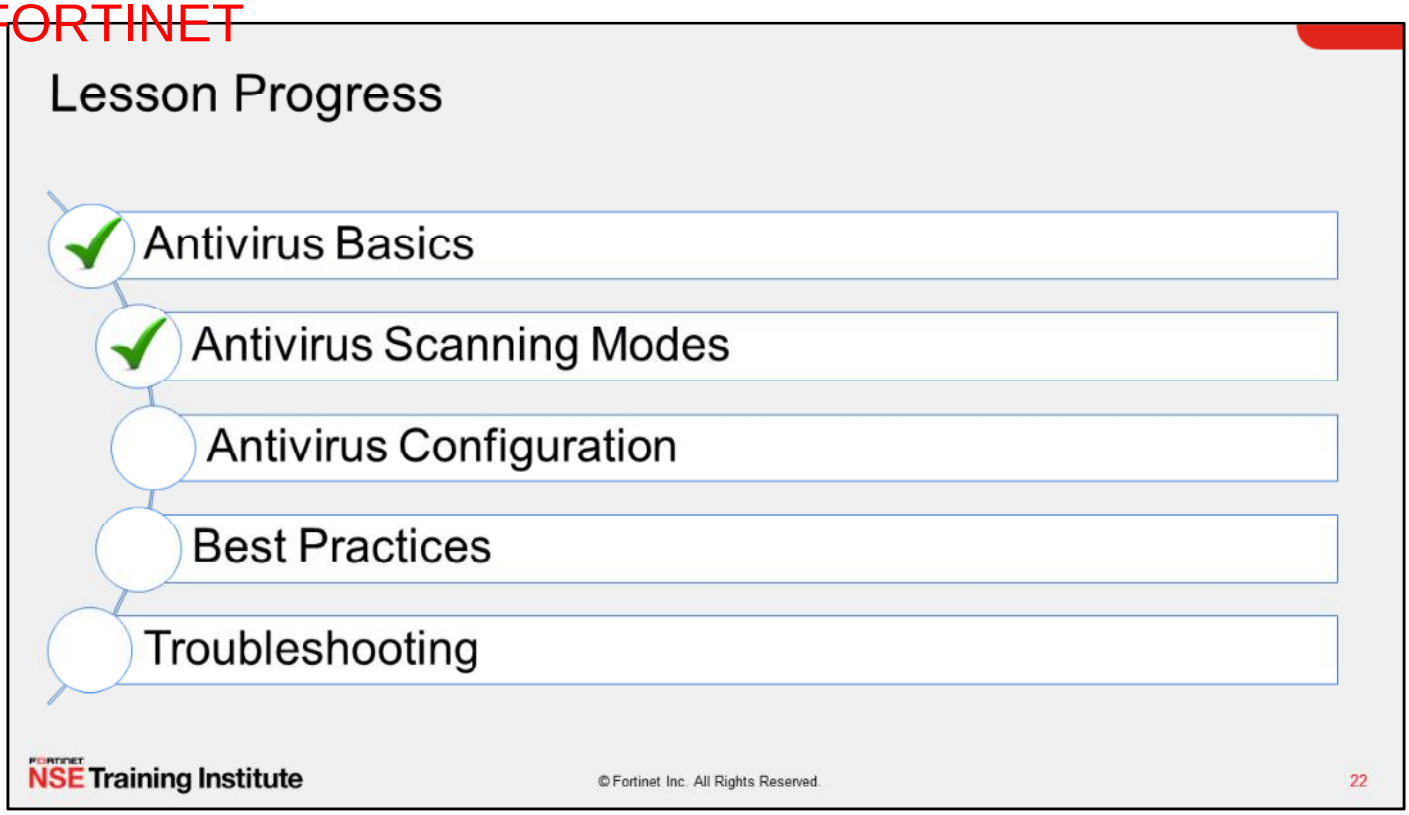

Good job! You now understand antivirus scanning modes.

Now, you will learn about antivirus configuration.

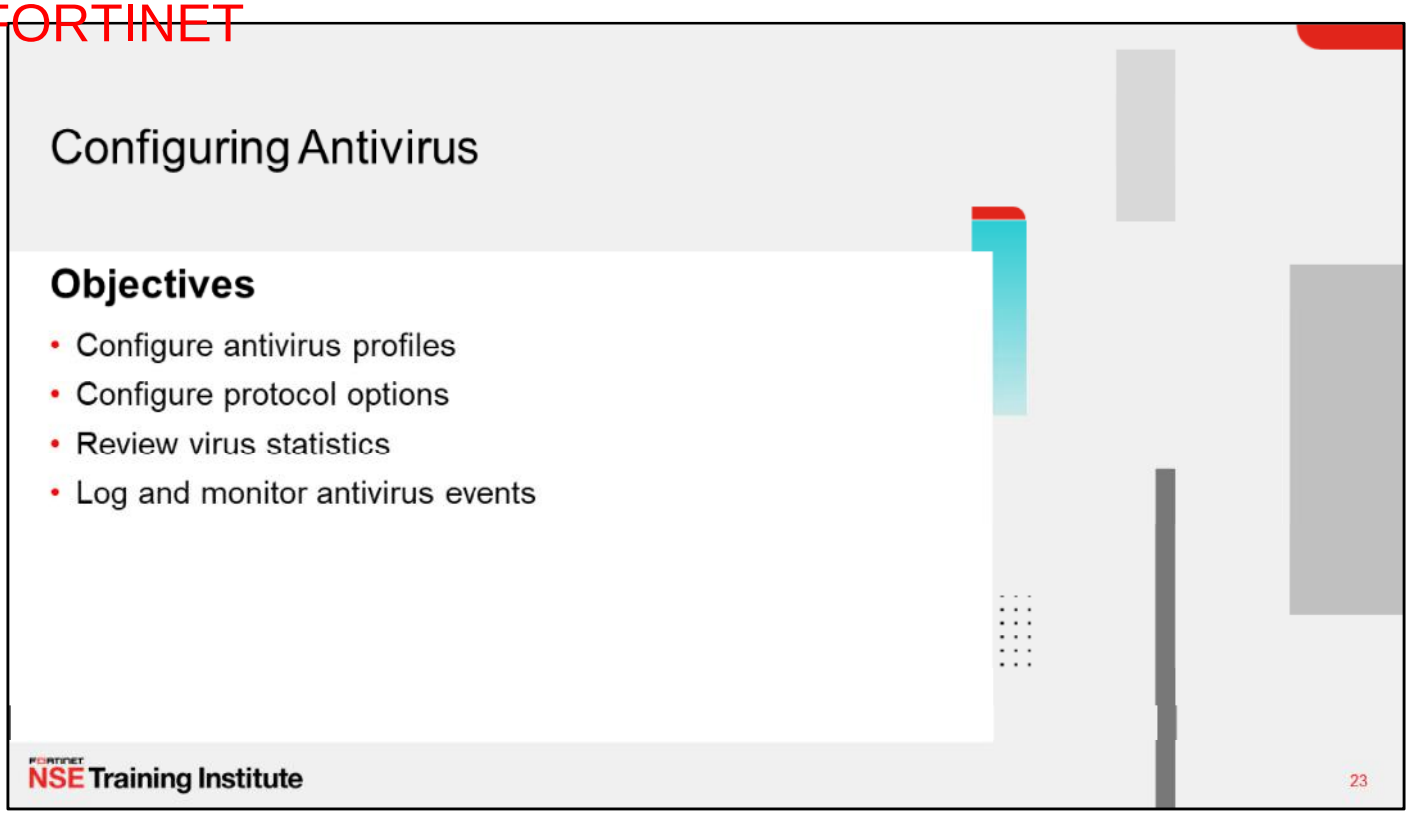

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus configuration, including reviewing antivirus logs, you will be able to use the antivirus profile in an effective manner.

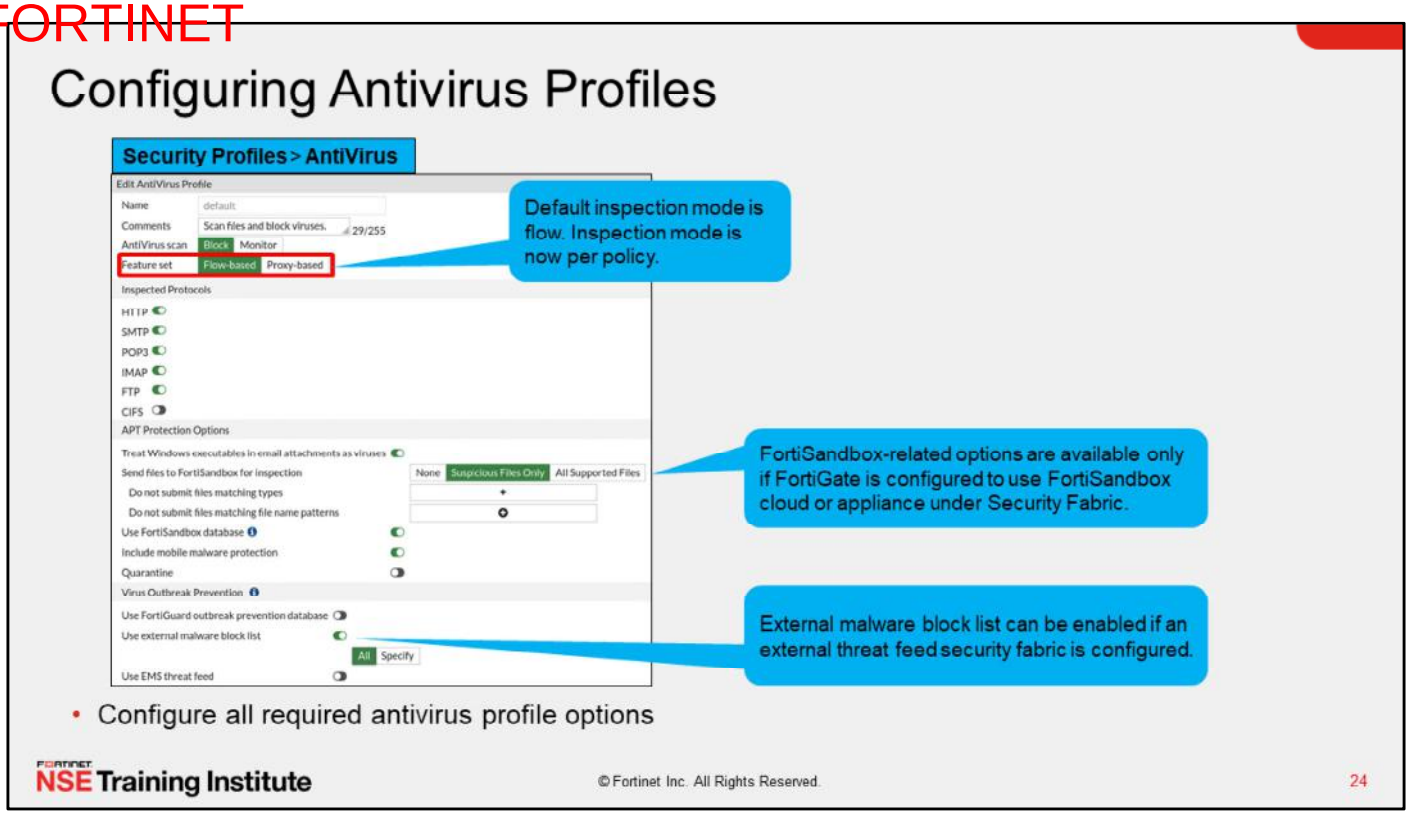

The antivirus profile can be configured on the **AntiVirus** page. Since the default inspection mode on a firewall policy is flow-based, **Feature set** is required to be set to **Flow-based**. If the inspection mode of the firewall policy is proxy-based, **Feature set** can be set to **Proxy-based,** 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>
```

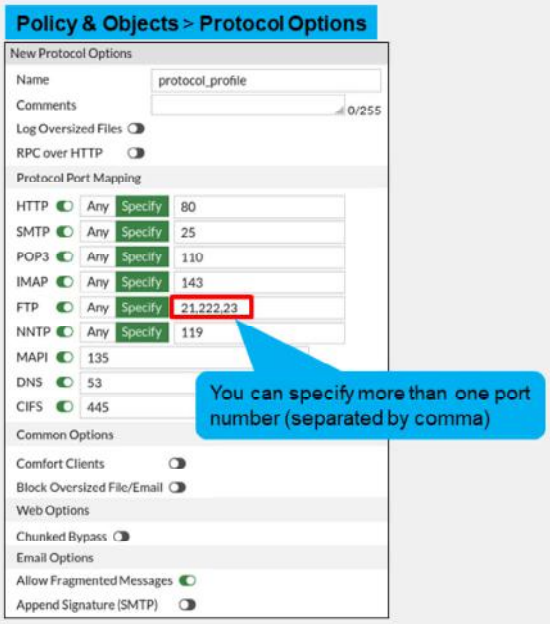

NSE Training Institute

25

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.

Protocol Options—Large Files

- By default, FortiOS allows files that are too big for the buffer size
  - Files that are bigger than oversize limit are bypassed from scanning
- You can modify this behavior for all protocols

```
config firewall profile-protocol-options
edit <profile_name>
config <protocol_name>
set options oversize
set oversize-limit <integer>
end
end
```

HTTP, FTP, and so on

Default value is 10 MB. Maximum value is hardware dependant.

- You can enable logging of oversize files using CLI

```
config firewall profile-protocol-options
edit <profile_name>
set oversize-log {enable|disable}
end
```

NSE Training Institute

26

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.

Protocol Options—Compressed Files

- Often, compression algorithms can be identified using header only
- Archives are unpacked and files and archives within are scanned separately
  - Nested archives are supported (default is 12 layers)
    - Supported formats: ZIP, TAR, GZIP, RAR, LSH, CAB, ARJ, MSC, BZIP, BZIP2, 7Z, EGG, XZ, CPIO, AR, ACE, ISO, DAA, CRX, and CHM
  - Decompressed files have a separate oversize limit
  - Limit can be configured for each protocol separately

```
config firewall profile-protocol-options
edit <profile_name>
config <protocol_name>
set uncompressed-oversize-limit [1-<model_limit>]
set uncompressed-nest-limit [1-<model_limit>]
end
end
```

HTTP, FTP, and so on

- Password-protected archives cannot be decompressed
- Increasing the size will increase memory usage!

NSE Training Institute

27

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.

Detection Rate and File Size

- Most malware is small
- Very large files require more RAM to scan completely
- Often, scanning only small files is an acceptable risk
  - Default: 10 MB threshold for `oversize`
  - Maximum size varies by model

Malware Type	1MB	2MB	3MB	4MB	5MB	6MB	7MB	8MB	9MB	10MB	∞
Exploit	99.83%	99.95%	99.97%	99.97%	99.98%	99.98%	99.99%	100%	100%	100%	100%
Mass-mailer	99.62%	99.87%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Phish	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Spyware	95.08%	97.97%	98.88%	99.47%	99.76%	99.83%	99.89%	99.91%	99.94%	99.95%	100%
Trojan	97.52%	99.24%	99.62%	99.80%	99.88%	99.93%	99.95%	99.97%	99.98%	99.98%	100%
Virus	98.27%	99.37%	99.63%	99.80%	99.88%	99.93%	99.95%	99.97%	99.98%	99.99%	100%
worm	99.08%	99.65%	99.74%	99.86%	99.89%	99.92%	99.94%	99.94%	99.95%	99.96%	100%

NSE Training Institute

28

An antivirus profile in full scan mode buffers up to your specified file size limit. The default is 10 MB. That is large enough for most files, except video files. If your FortiGate model has more RAM, you may be able to increase this threshold.

Without a limit, very large files could exhaust the scan memory. So, this threshold balances risk and performance. Is this tradeoff unique to FortiGate, or to a specific model? No. Regardless of vendor or model, you must make a choice. This is because of the difference between scans in theory, that have no limits, and scans on real-world devices, that have finite RAM. In order to detect 100% of malware regardless of file size, a firewall would need infinitely large RAM—something that no device has in the real world.

Most viruses are very small. This table shows a typical tradeoff. You can see that with the default 10 MB threshold, only 0.01% of viruses pass through.

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

NSE Training Institute

29

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.

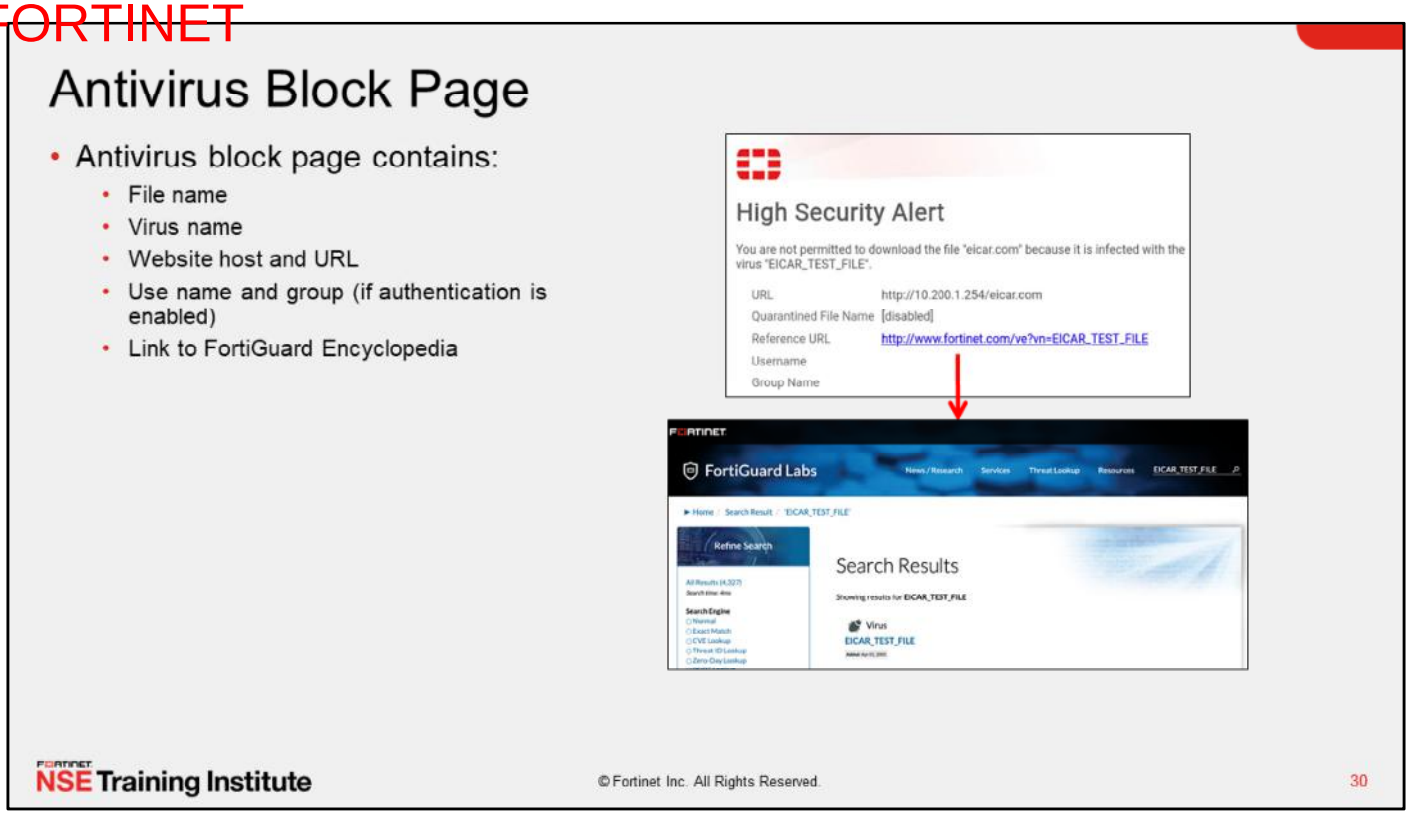

For antivirus scanning in proxy-based inspection mode (with client comforting disabled), the block replacement page is displayed *immediately* when a virus is detected.

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.

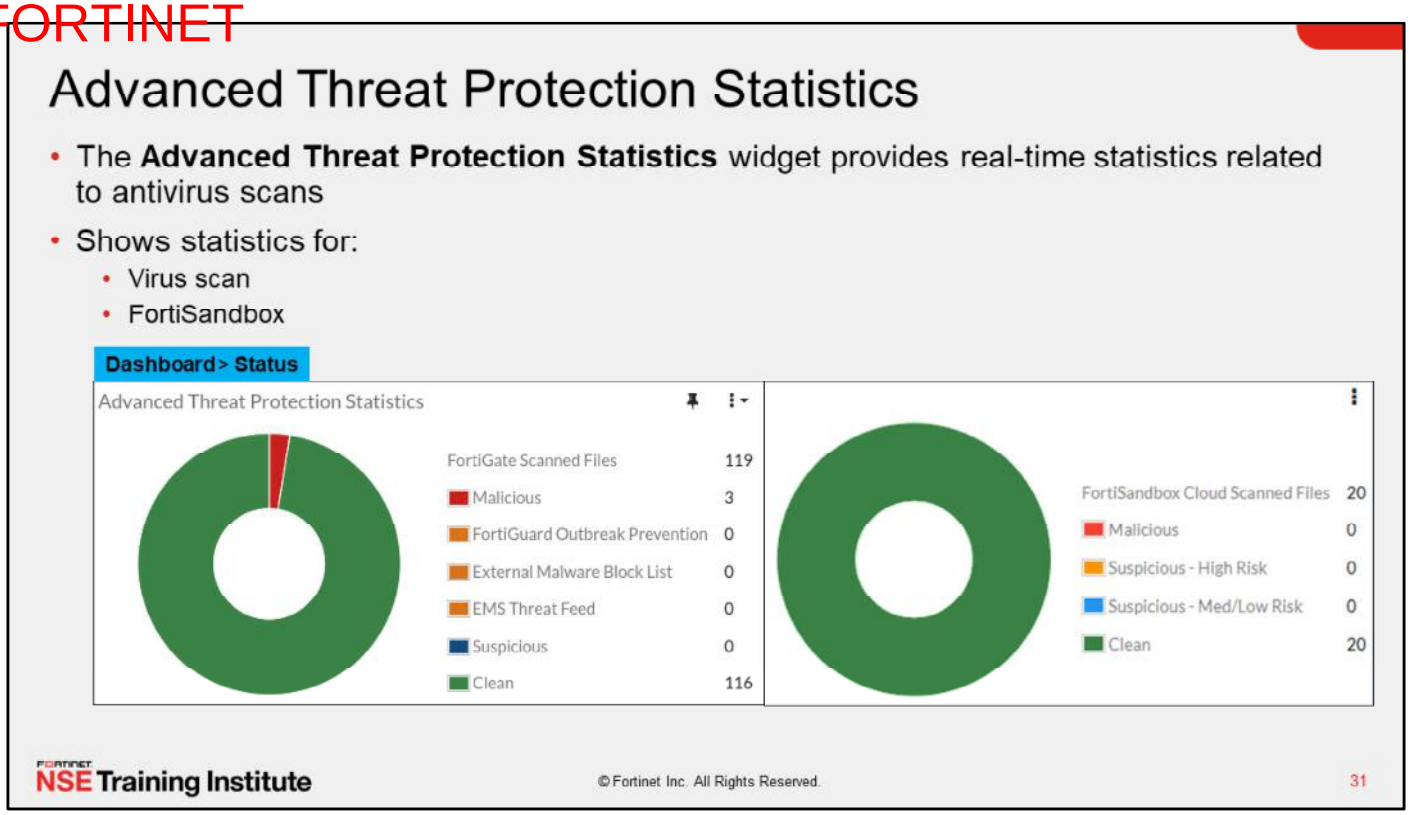

You can find virus scanning statistics on the **Advanced Threat Protection Statistics** widget on the dashboard.

If your FortiGate is submitting files for sandboxing, it keeps statistics about the number of files submitted and the results of those scans. These statistics are separate from files that are scanned locally on FortiGate.

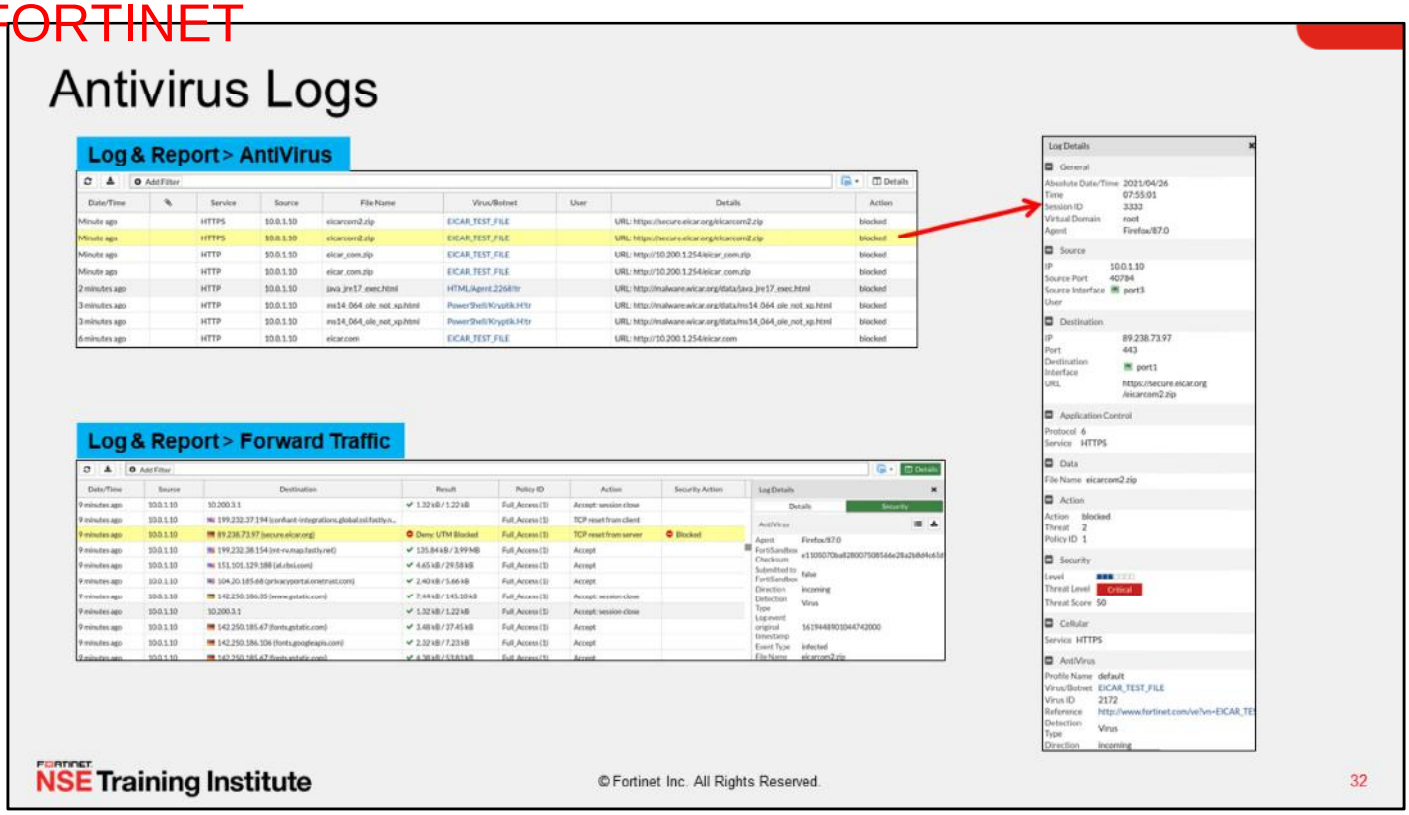

If you enable logging, you can find details on the **AntiVirus** log page.

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.

Knowledge Check

1. What is the default scanning behavior for files over 10 MB?
   - ✓ A. Allow the file without scanning
   - B. Block all large files that exceed the buffer threshold

NSE Training Institute

33

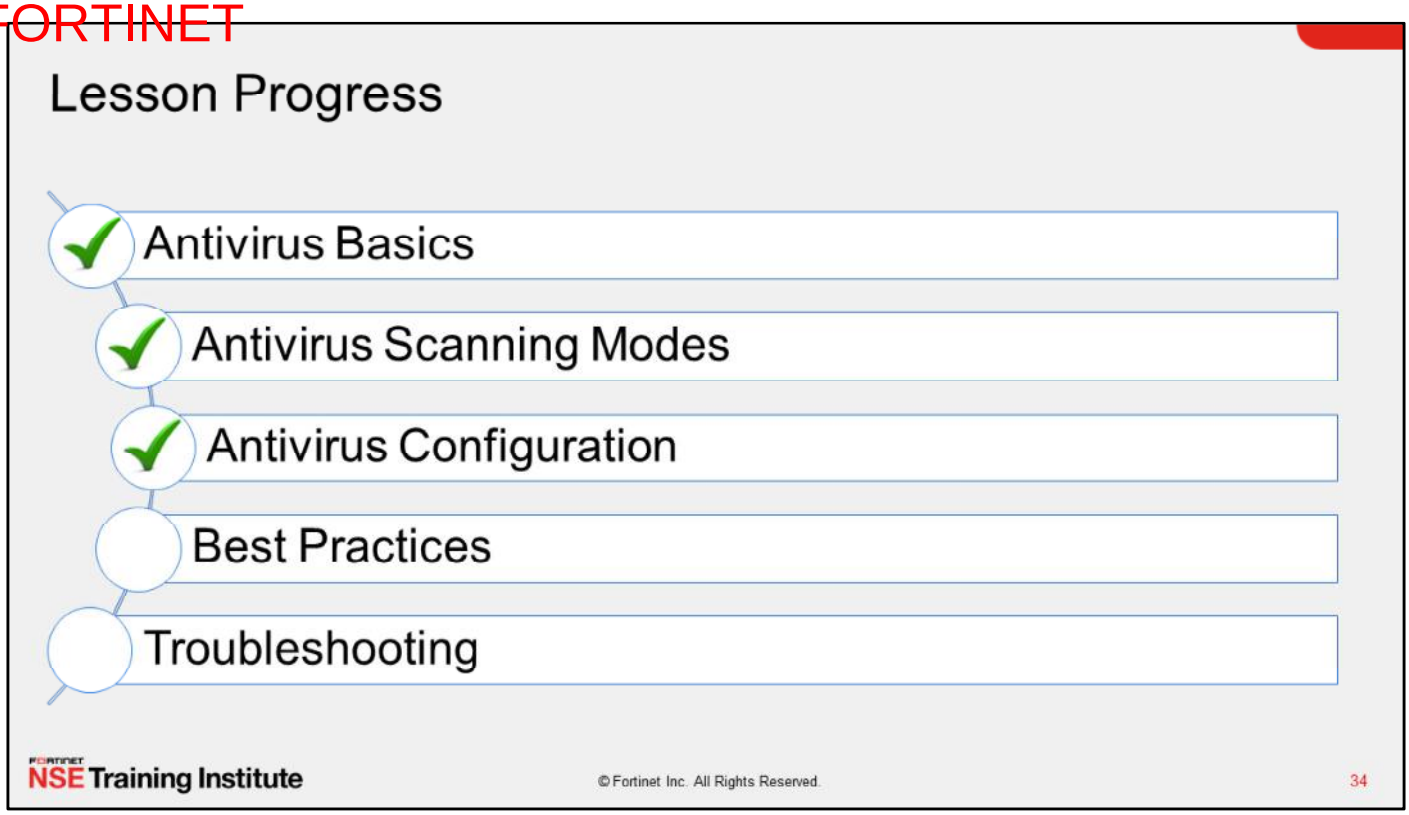

Good job! You now understand antivirus configuration.

Now, you will learn about some antivirus best practices.

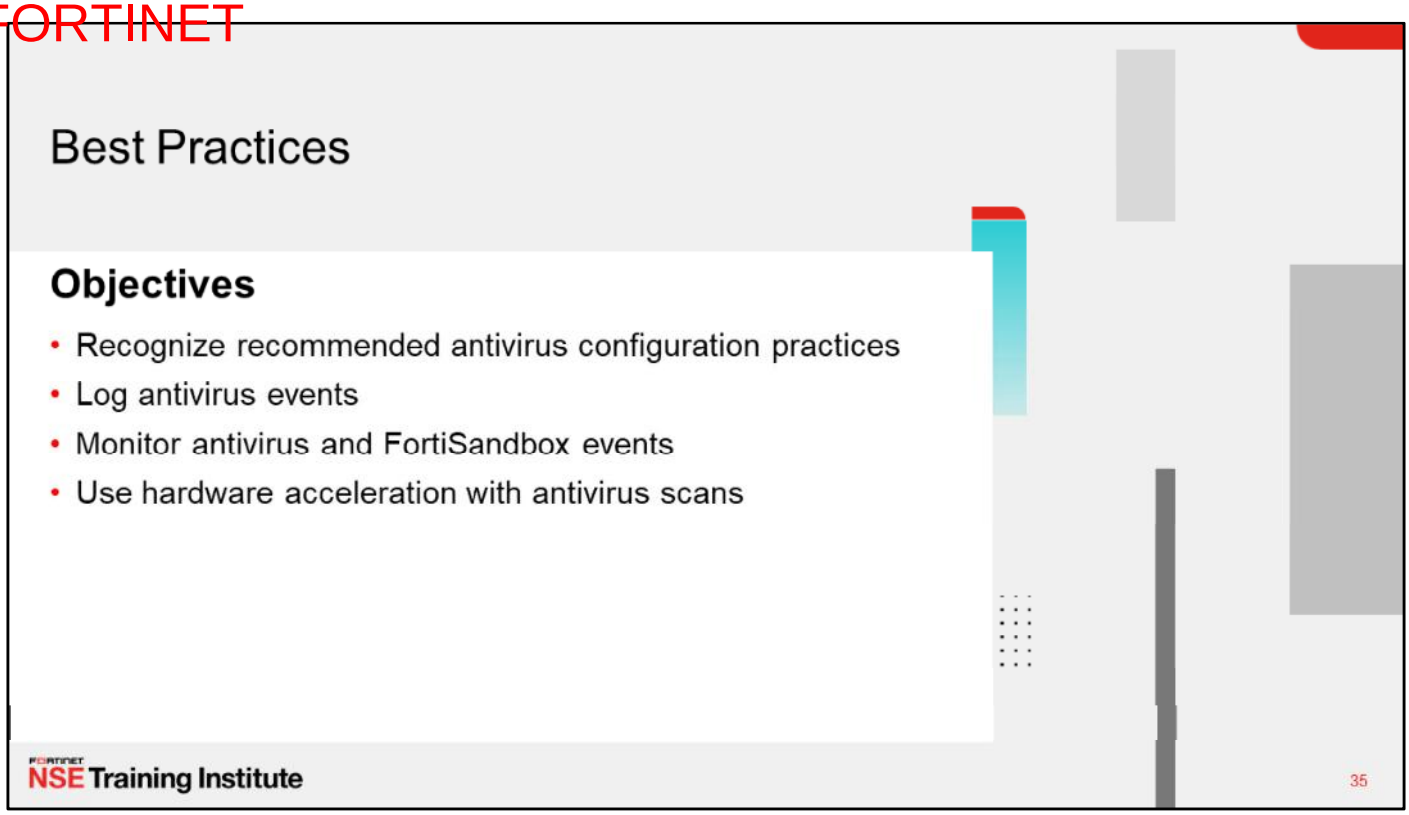

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in antivirus best practices, you will be able to configure an effective antivirus solution.

Recommended Configuration Practices

- Perform antivirus scan on all internet traffic
  - If using load balancing or redundant internet connections, ensure all internal to external firewall policies have antivirus profiles applied on them
- Use *deep-inspection* instead of *certificate-based* inspection, to ensure that full content inspection is performed
- Use FortiSandbox Cloud or a FortiSandbox device to enable sandboxing support
  - Configure the antivirus profile to use the FortiSandbox database
- Do not increase the maximum file size to be scanned, unless it is required
  - Viruses usually travel in small files
  - More scanning means more memory utilization

NSE Training Institute 36

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

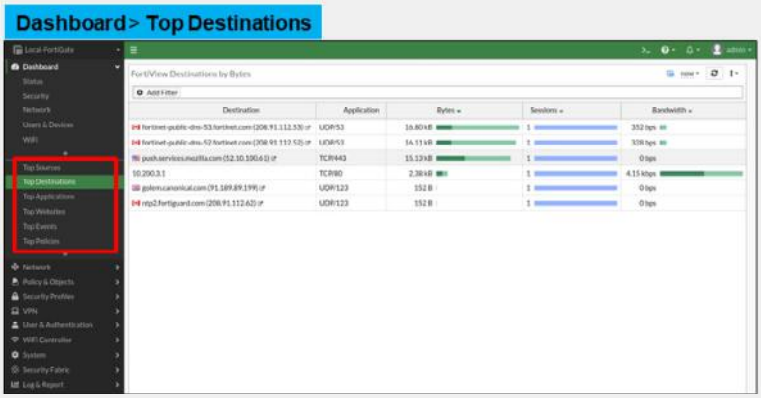

NSE Training Institute

37

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.

Hardware Acceleration for Antivirus Scanning

- Accelerates flow-based antivirus only
- FortiGate models that feature NTurbo (NP6 or NP7) can accelerate antivirus processing to enhance performance
  - SoC4 models also support 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

```
config ips global
  set np-accel-mode {none | basic}
end
```

Enable NTurbo acceleration

- Proxy inspection mode
  - Proxy-based inspection cannot be offloaded for acceleration

NSE Training Institute

38

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.

Hardware Acceleration for Antivirus Scanning (Contd)

- FortiGate models with content processors (CP8 or CP9) support offloading of flow-based pattern matching
- Flow-based pattern databases are compiled and downloaded to the content processors from the IPS engine and IPS database
  - Accelerates pattern matching while reducing the load on FortiGate CPU

```
config ips global
  set cp-accel-mode {none | basic | advanced}
end
```

Enable AV scan offloading to CP

- Proxy inspection mode
  - Proxy-based antivirus scanning cannot be offloaded for acceleration

NSE Training Institute

39

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

NSE Training Institute

40

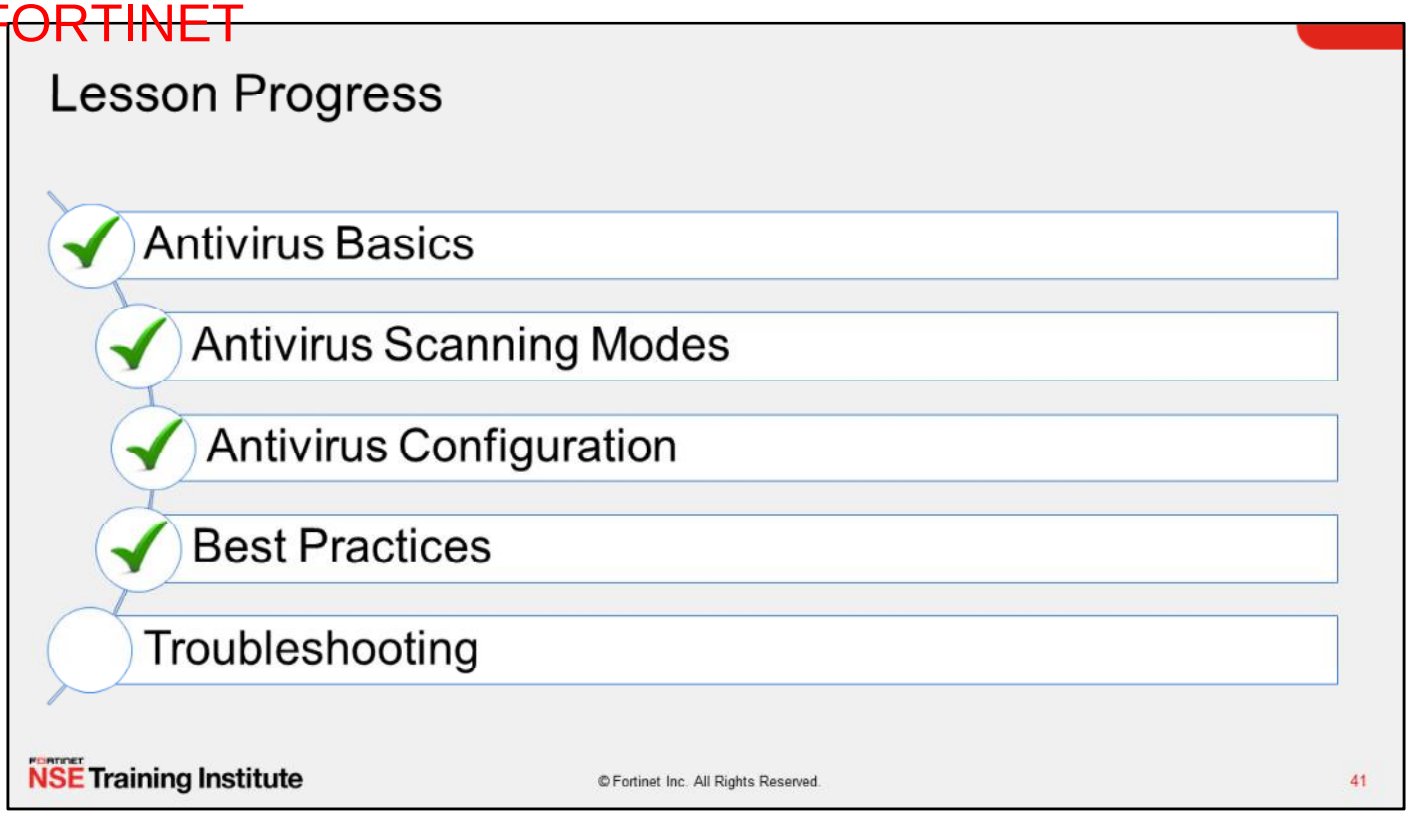

Good job! You now understand antivirus best practices.

Now, you will learn about antivirus troubleshooting.

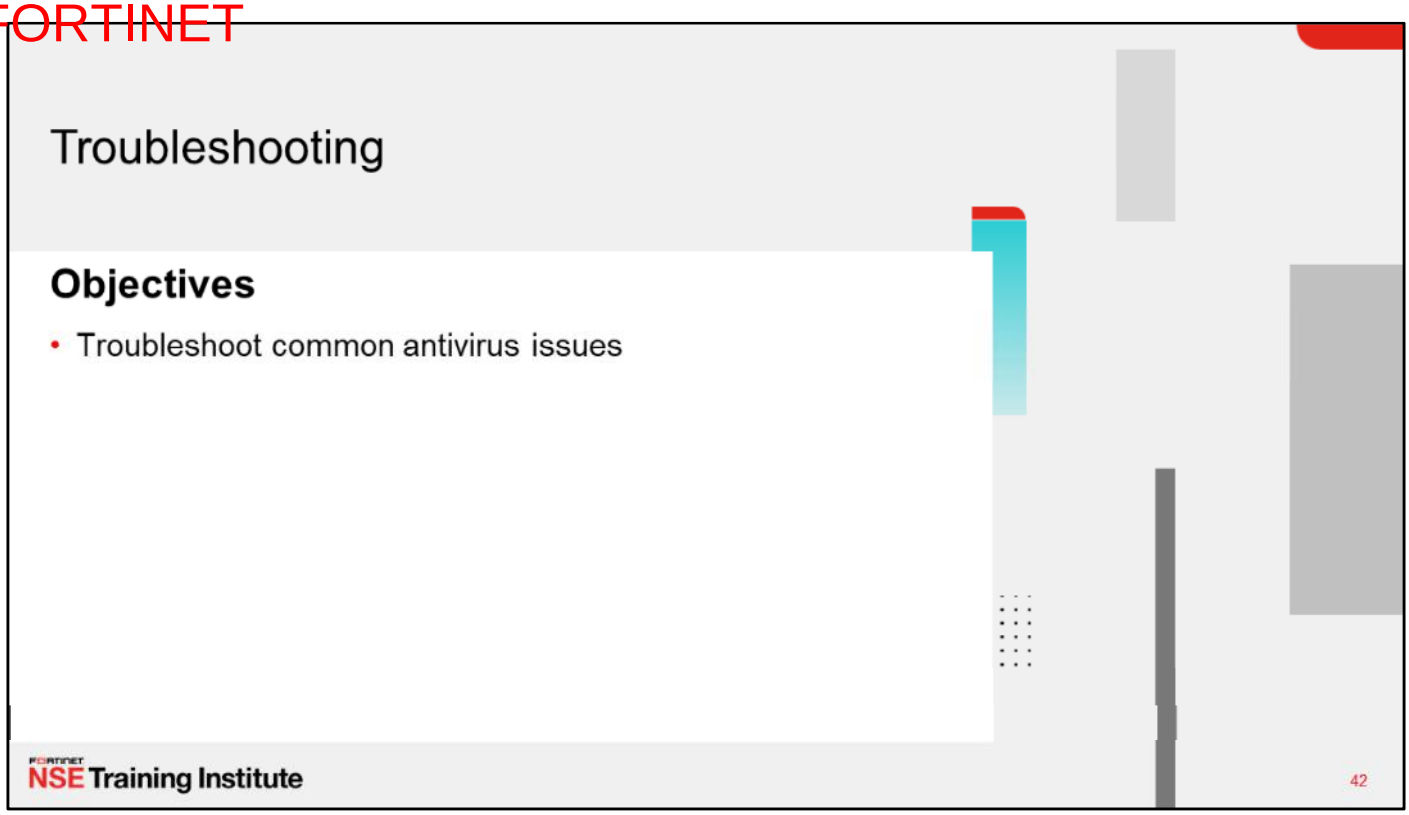

After completing this section, you should be able to troubleshoot common issues with antivirus.

By demonstrating competence in troubleshooting common antivirus issues, you will be able to configure and maintain an effective antivirus solution.

Troubleshooting Common Antivirus Issues

- Valid contract but antivirus database is out-of-date?
  - Check FortiGuard website for latest antivirus database version
    - https://fortiguard.com/updates/antivirus
  - Make sure the antivirus profile is applied on at least one firewall policy
- Run the real-time update debug to isolate update-related issues

```
# diagnose debug application update -1
# diagnose debug enable
# execute update-av
```

NSE Training Institute

43

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.

Troubleshooting Common Antivirus Issues (Contd)

- FortiGuard update issues? Make sure that:
  - FortiGate has a stable connection to the internet
  - FortiGate is able to resolve DNS (`update.fortiguard.net`)
  - TCP port 443 is open
- Force FortiGate to check for new antivirus updates

```
# execute update-av
```

- Verify that the FortiGate antivirus license is valid

NSE Training Institute

44

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.

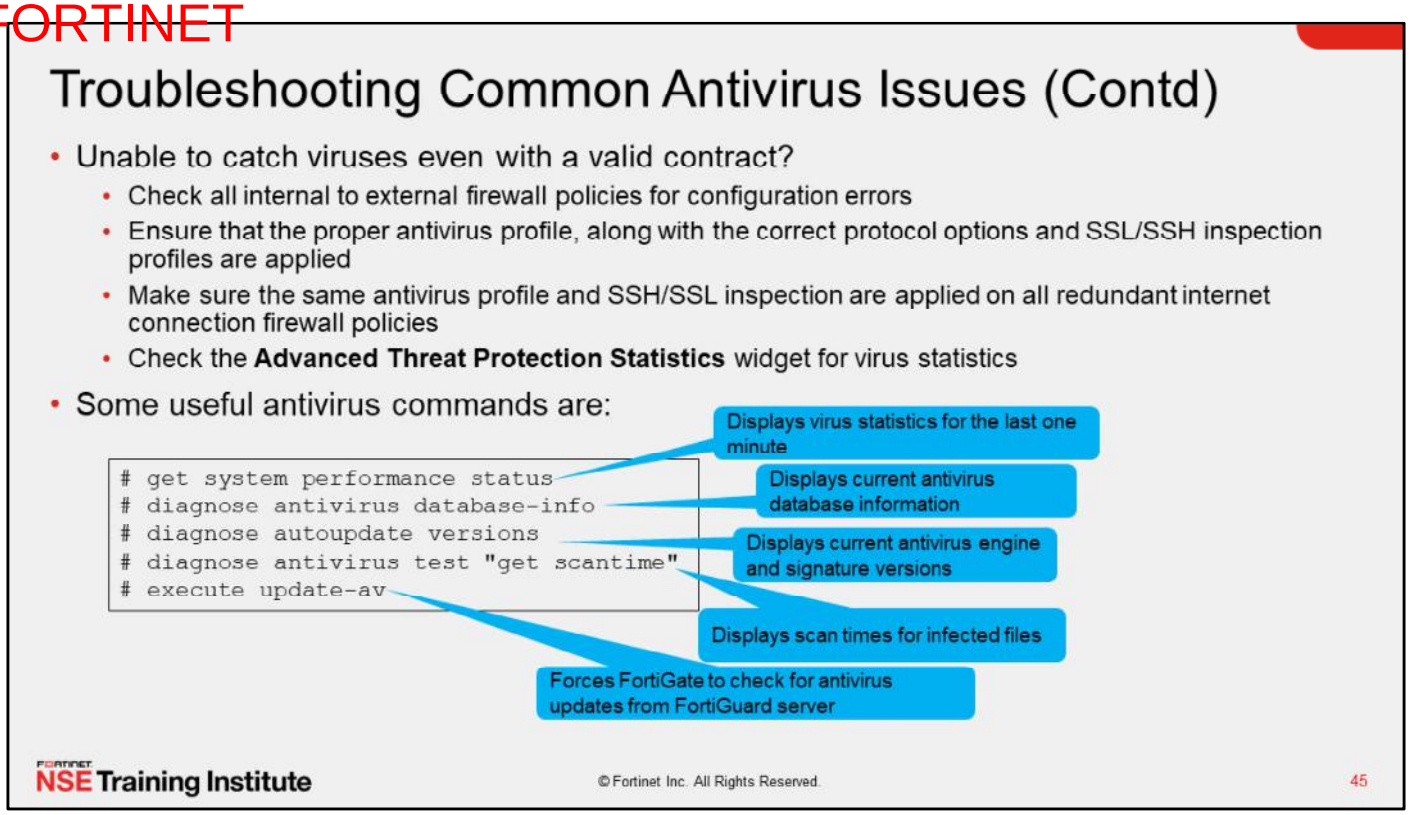

What if you have a valid contract and updated database, and you are still having issues catching viruses? Start troubleshooting for basic configuration errors. Most of the time, issues are caused by misconfiguration on the device. You can verify them as following:

- 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`

46

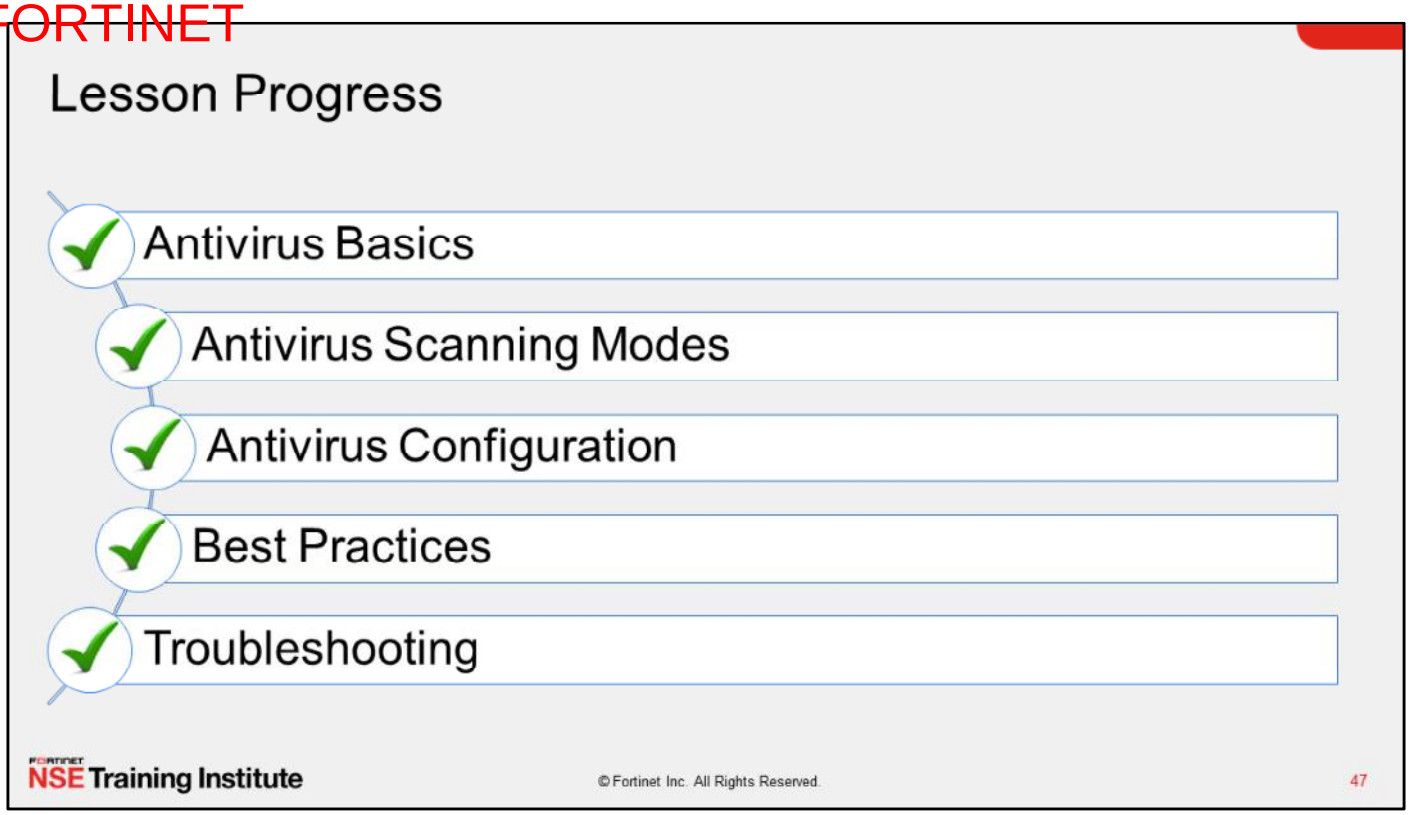

Congratulations! You have completed this lesson.

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

Review

- ✓ Use antivirus signatures
- ✓ Review antivirus scanning techniques
- ✓ Enable FortiSandbox with antivirus
- ✓ Differentiate between available FortiGuard signature databases
- ✓ Apply the antivirus profile in flow-based and proxy-based inspection modes
- ✓ Compare all available scanning modes
- ✓ Configure antivirus profiles and protocol options
- ✓ Review virus statistics
- ✓ Log and monitor antivirus events
- ✓ Recognize recommended antivirus configuration practices
- ✓ Log and monitor antivirus and FortiSandbox events
- ✓ Use hardware acceleration with antivirus scans
- ✓ Troubleshoot common antivirus issues

NSE Training Institute 48

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.

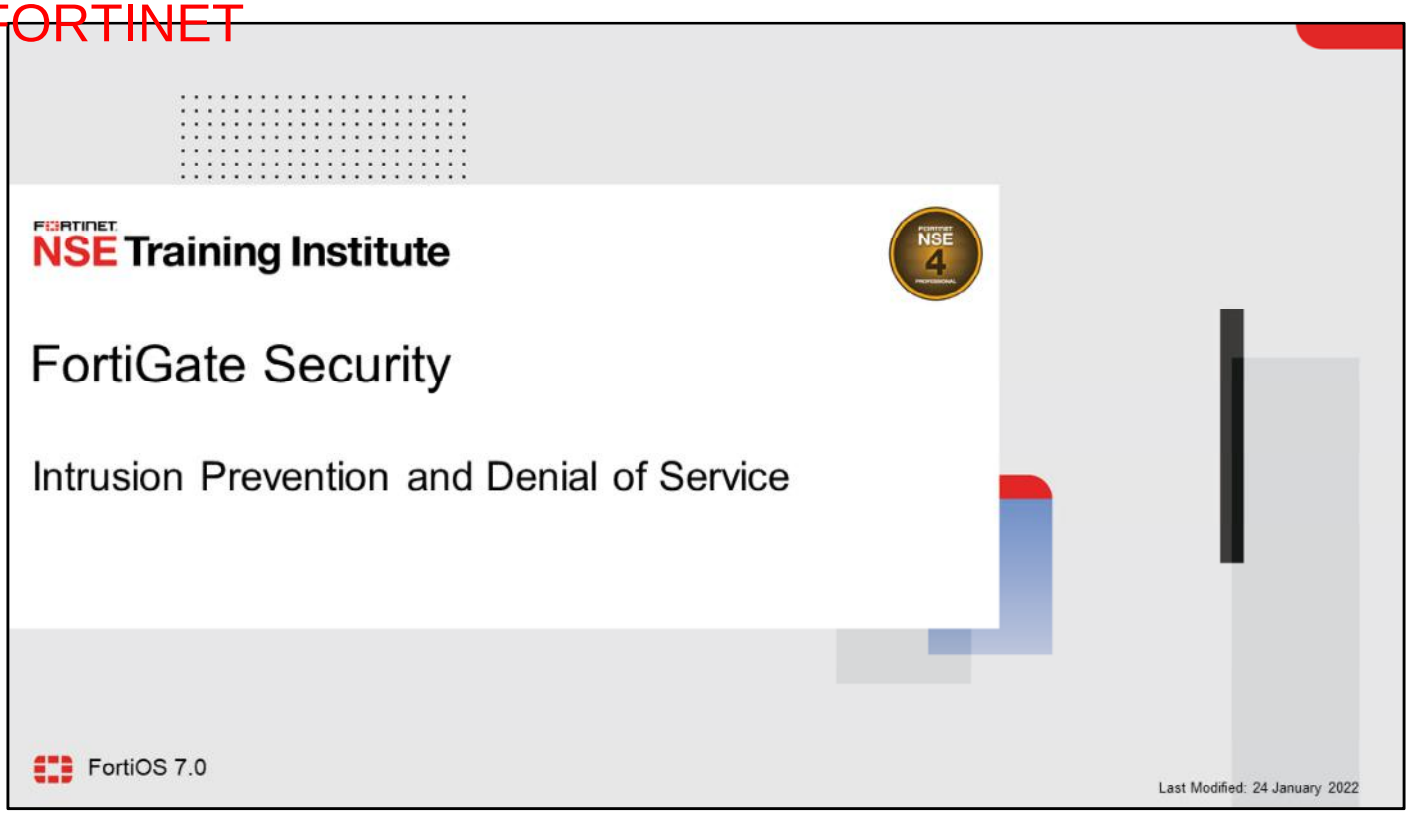

In this lesson, you will learn how to use FortiGate to protect your network against intrusions and denial of service (DoS) attacks.

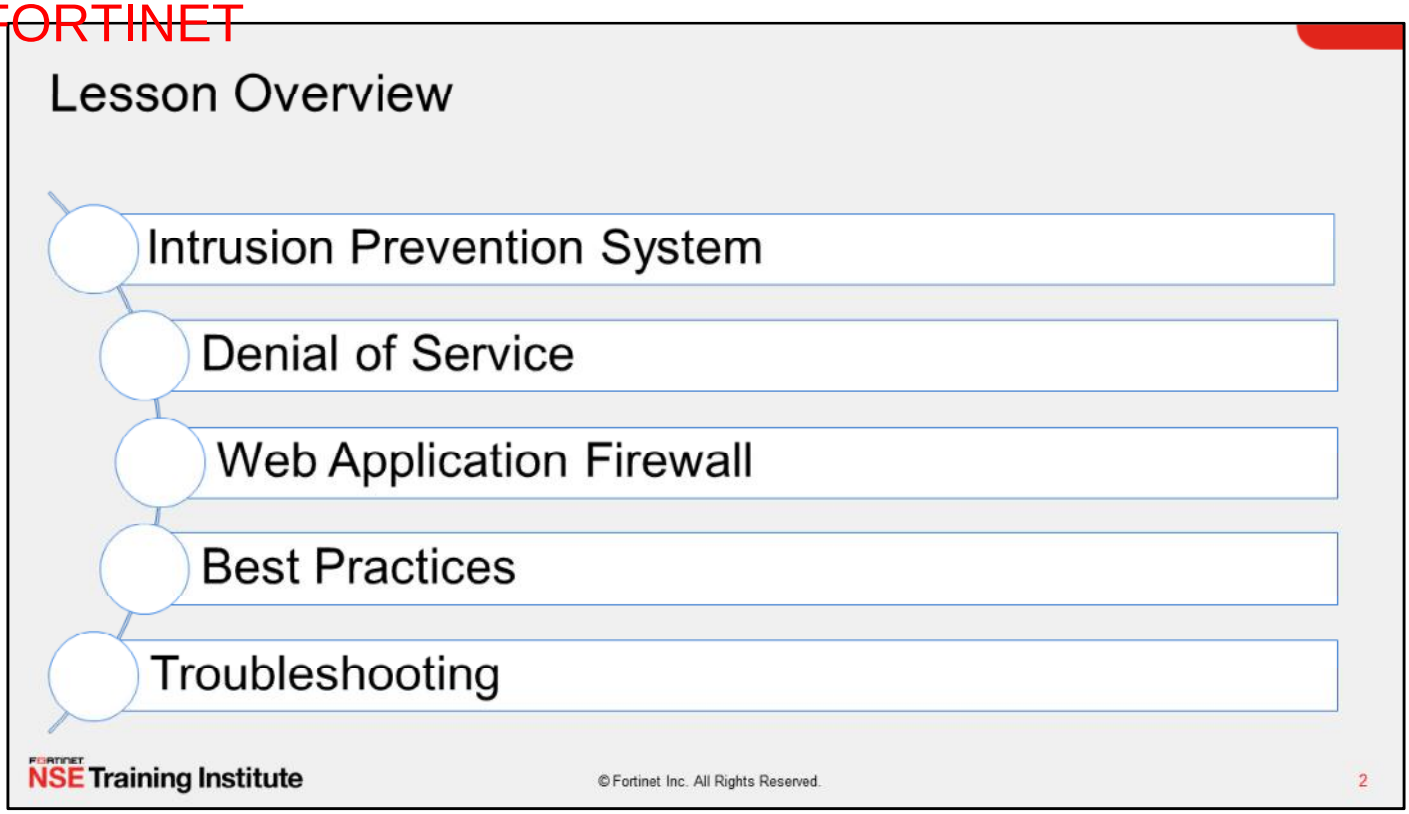

In this lesson, you will learn about the topics shown on this slide.

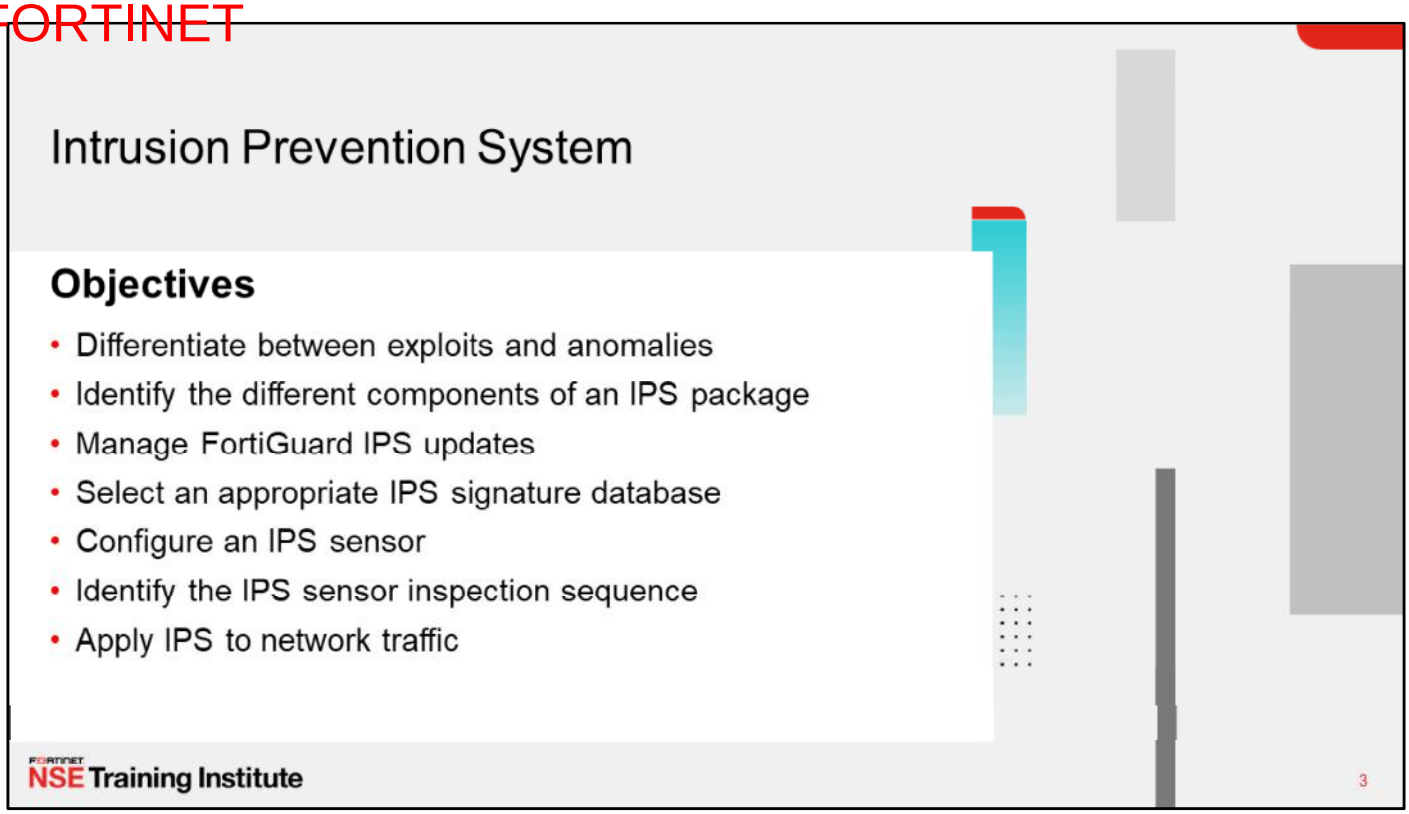

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in intrusion prevention system (IPS), you should be able to implement an effective IPS solution to protect your network from intrusion.

Why Use IPS?

- Increased volume and sophistication of attacks on organizations
  - Driven by previously successful high-profile hacks and a highly profitable black-market demand for stolen data
- More attacks against client and cloud applications
  - Attacks are no longer targeted only at servers and server-based applications only
- BYOD and remote workers increase risk of exposure

See attacks happening in real time around the world on the FortiGuard Labs live threat map.

NSE Training Institute

4

Organizations are under continuous attacks. Cybercriminals, motivated by previously successful high-profile hacks and a highly profitable black market for stolen data, continue to increase both the volume and sophistication of their attacks on organizations. Many organizations encourage BYOD and flexible working environments, which has led to the explosion of anytime, anywhere data consumption. This consumption increases the risk that sensitive data will be exposed to unauthorized access outside corporate boundaries.

Today's threat landscape requires IPS to block a wider range of threats, while minimizing false positives.

Exploits and Anomalies

Anomaly

- Can be zero-day or DoS attacks
- Detected by behavioral analysis:
  - Rate-based IPS signatures
  - DoS policies
  - Protocol constraints inspection
- Example:
  - Abnormally high rate of traffic (DoS/flood)

Exploit

- A known, confirmed attack
- Detected when a file or traffic matches a signature pattern:
  - IPS signatures
  - WAF signatures
  - Antivirus signatures
- Example:
  - Exploit of known application vulnerabilities

NSE Training Institute © Fortinet Inc. All Rights Reserved. 5

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)
    - Data leak prevention (DLP) (flow based in one-arm sniffer mode)

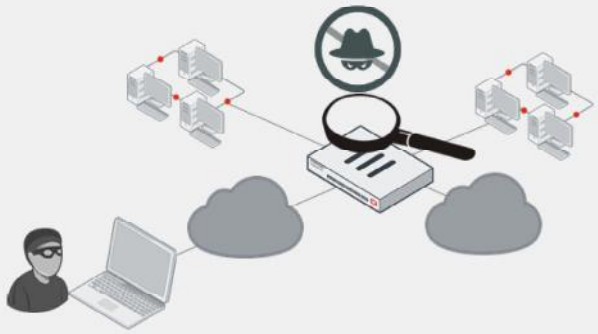

NSE Training Institute

6

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, email filtering, and flow-based DLP in one-arm sniffer mode.

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

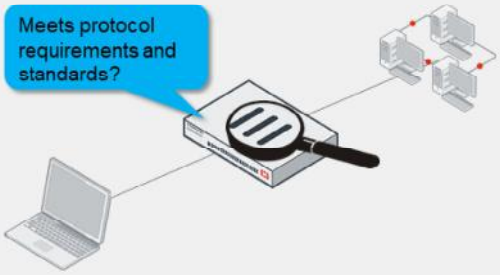

NSE Training Institute

7

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.

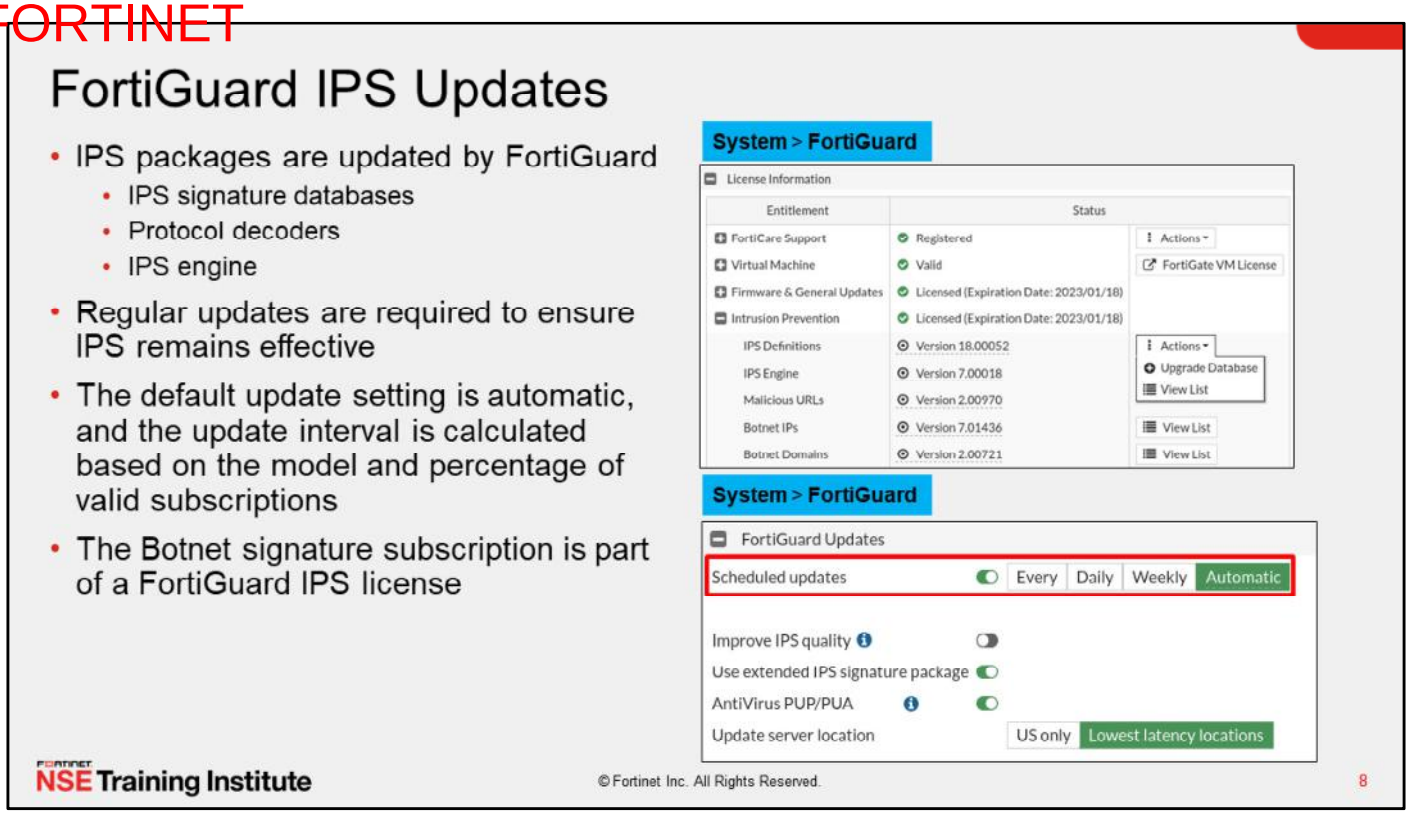

By default, an initial set of IPS signatures is included in each FortiGate firmware release. FortiGuard updates the IPS signature database with new signatures. That way, IPS remains effective against new exploits. Unless a protocol specification or RFC changes (which doesn't happen very often), protocol decoders are rarely updated. The IPS engine itself changes more frequently, but still not often.

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 go without being 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 botnet signature database. The botnet IP database is part of the ISDB updates. The botnet domains database is part of the AV updates, 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

NSE Training Institute © Fortinet Inc. All Rights Reserved. 9

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.

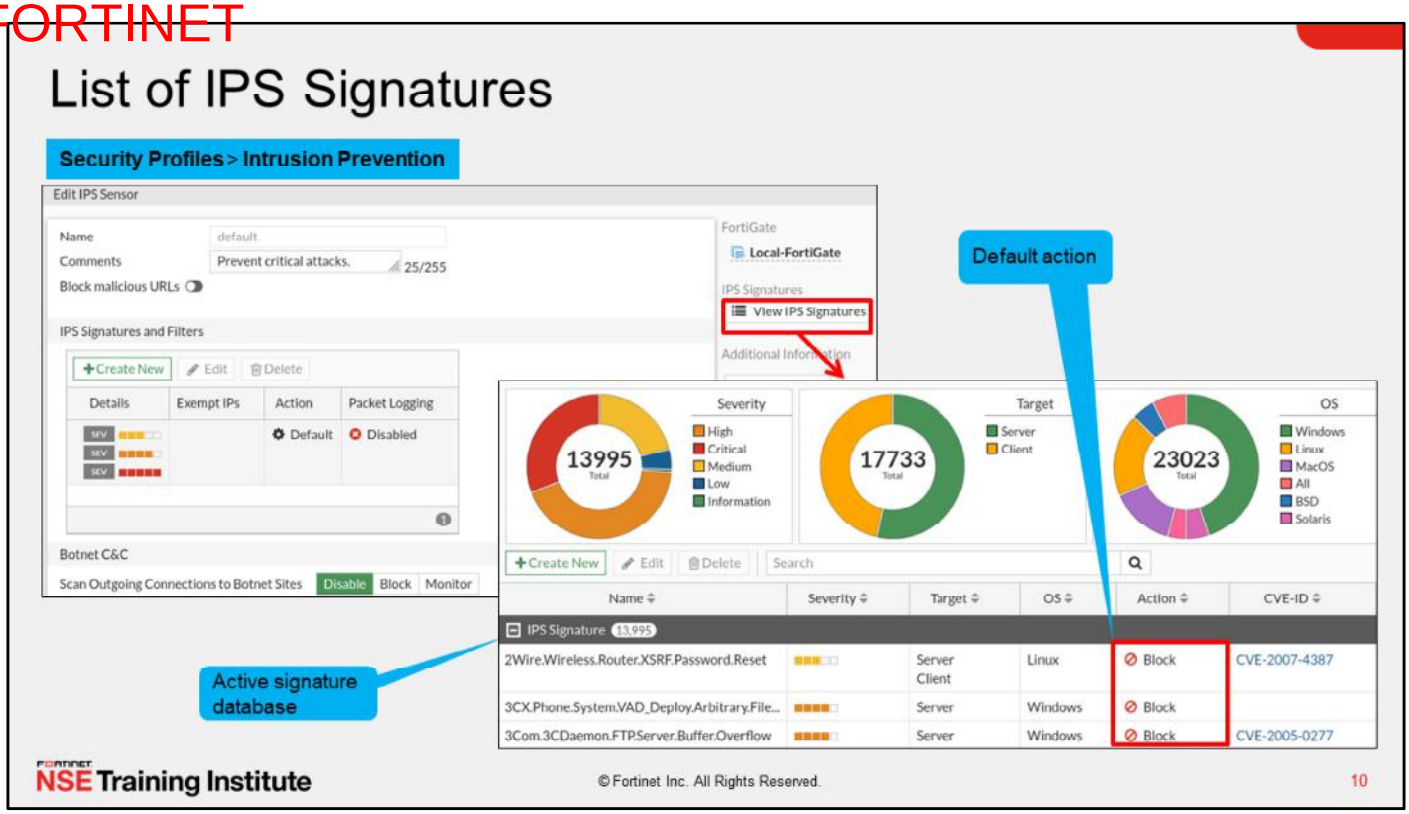

After FortiGate downloads a FortiGuard IPS package, new signatures appear in the signature list. When configuring FortiGate, you can change the **Action** setting for each sensor that uses a signature.

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.

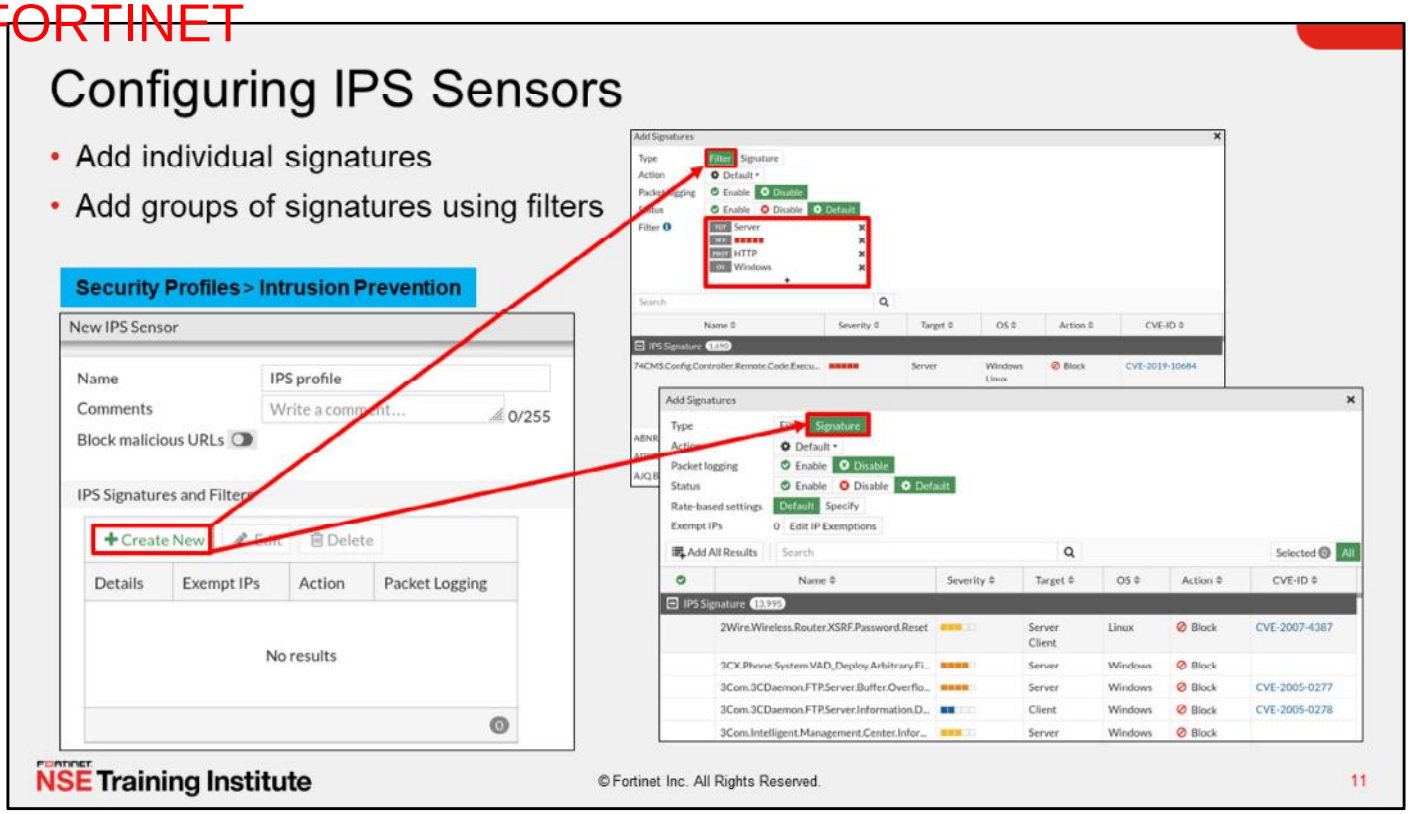

There are two ways to add predefined signatures to an IPS sensor. One way is to select the signatures individually. After you select a signature in the list, the signature is added to the sensor with its default action. Then, you can right-click the signature and change the action.

The second way to add a signature to a sensor is using filters. FortiGate adds all the signatures that match the filters.

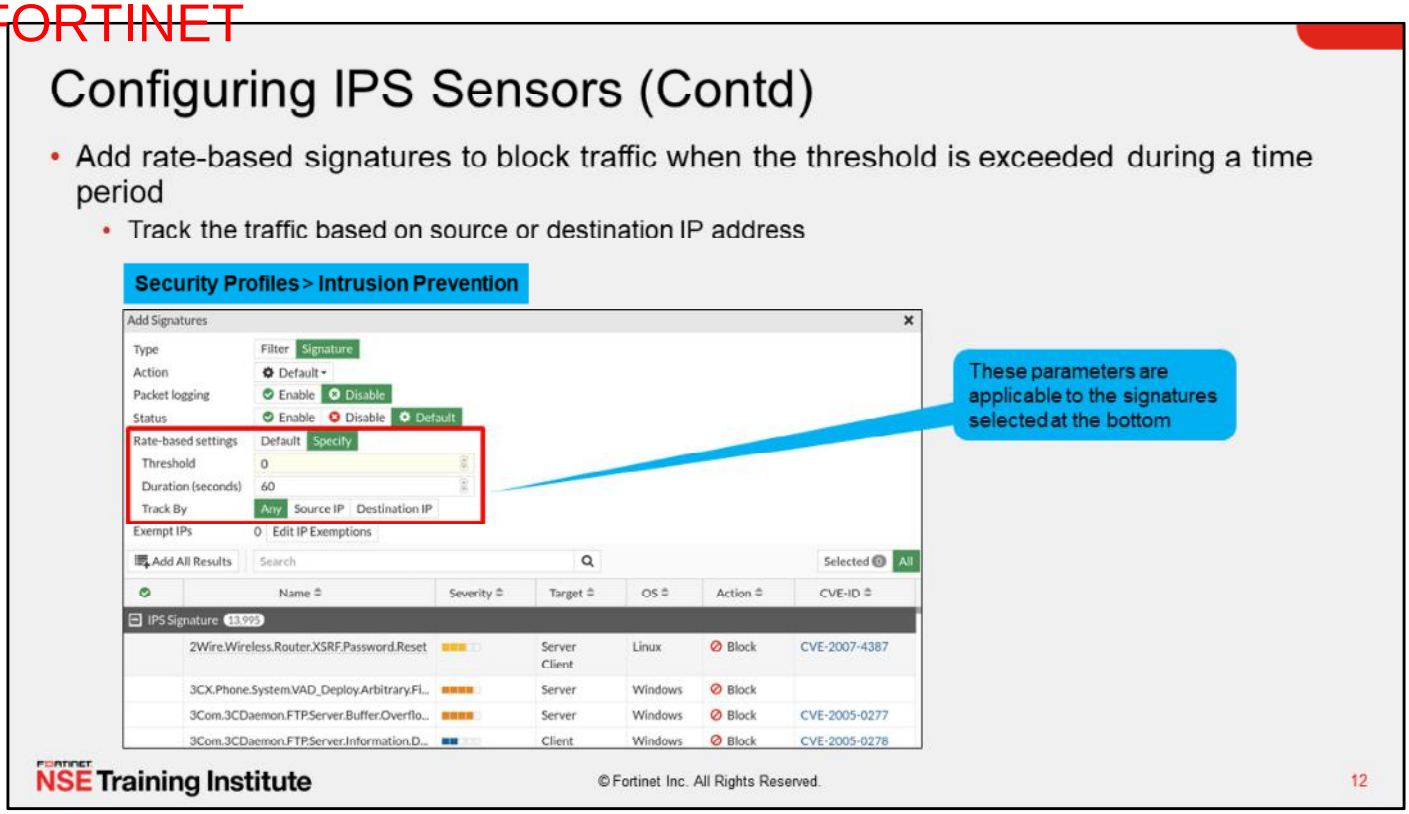

You can also add rate-based signatures to block specific traffic when the threshold is exceeded during the configured time period. You should apply rate-based signatures only to protocols you actually use. Then, configure **Duration** to block malicious clients for extended periods. This saves system resources and can discourage a repeat attack. FortiGate does not track statistics for that client while it is temporarily blocklisted.

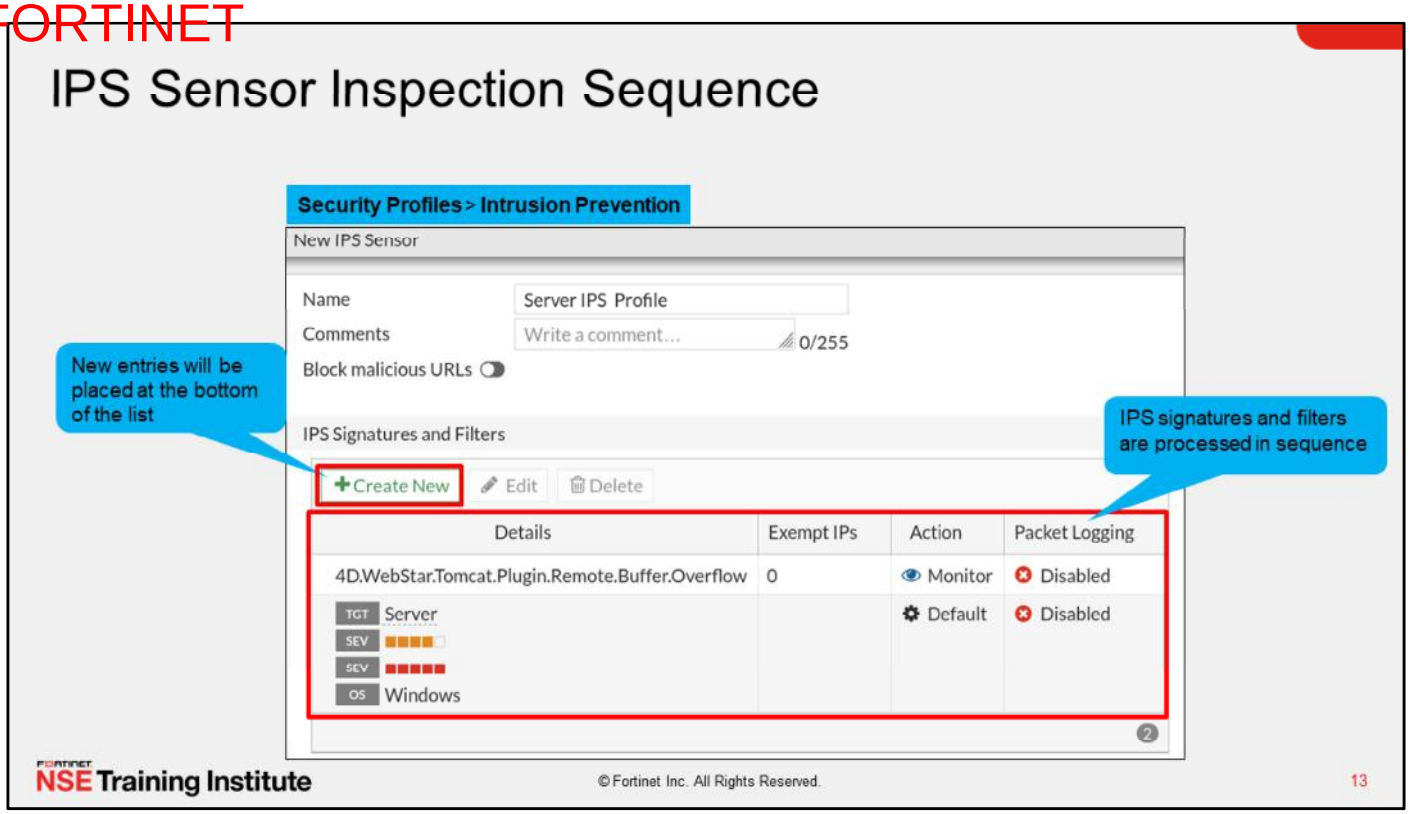

When the IPS engine compares traffic with the signatures in each filter, order matters. The rules are similar to firewall policy matching; the engine evaluates the filters and signatures at the top of the list first, and applies the first match. The engine skips subsequent filters.

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.

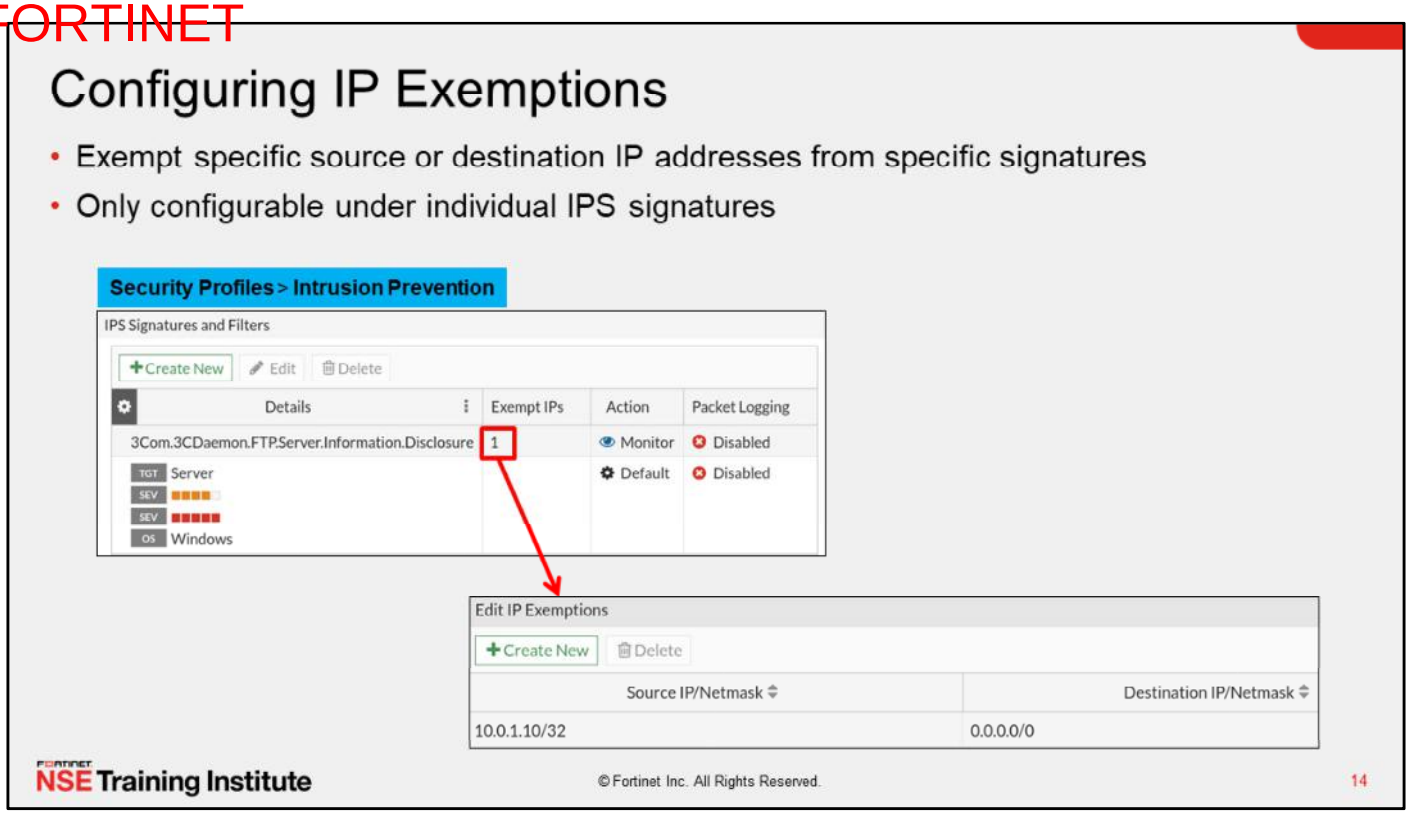

Sometimes it is necessary to exempt specific source or destination IP addresses from specific signatures. This feature is useful during false-positive outbreaks. You can temporarily bypass affected endpoints until you investigate and correct the false-positive issue.

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—Hold Time

- IPS signature filter options include hold time
  - Allows you to set the amount of time that signatures are held after a FortiGuard IPS signature update per VDOM
  - The signature mode is monitor
  - New signatures are enabled after the hold time, to avoid false positives
  - The hold time can be from 0 days and 0 hours (default) up to 7 days
  - To configure the amount of time to hold and monitor IPS signatures:

```
# config system ips
    set signature-hold-time 3d12h
    set override-signature-hold-by-id enable
  end
```

  - When a signature that is on hold is matched, the log includes the message signature is on `hold`

```
date=2021-04-06 time=00:00:57 logid="0419016384" type="utm" subtype="ips" eventtype="signature"
level="alert" vd="vd1" eventtime=1278399657778481842 tz="-0700" severity="info" srcip=10.1.100.22
srccountry="Reserved" dstip=172.16.200.55 srcintf="port13" srcintfrole="undefined" dstintf="port14"
dstintfrole="undefined" sessionid=3620 action="detected" proto=6 service="HTTP" policyid=1
attack="Eicar.Virus.Test.File" srcport=52170 dstport=80 hostname="172.16.200.55" url="/virus/eicar"
direction="incoming" attackid=29844 profile="test" ref="http://www.fortinet.com/ids/VID29844"
incidentserialno=25165825 msg="file_transfer: Eicar.Virus.Test.File, (signature is on hold)"
```

NSE Training Institute

16

You can configure the hold time option for the IPS signature filter. The hold-time option allows you to set the amount of time that signatures are held after a FortiGuard IPS signature update per VDOM. During the holding period, the signature mode is `Monitor`. The new signatures are enabled after the hold time, to avoid false positives.

The hold time can be from 0 days and 0 hours (default) up to 7 days.

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=2021-04-13 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"
```

NSE Training Institute

17

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.

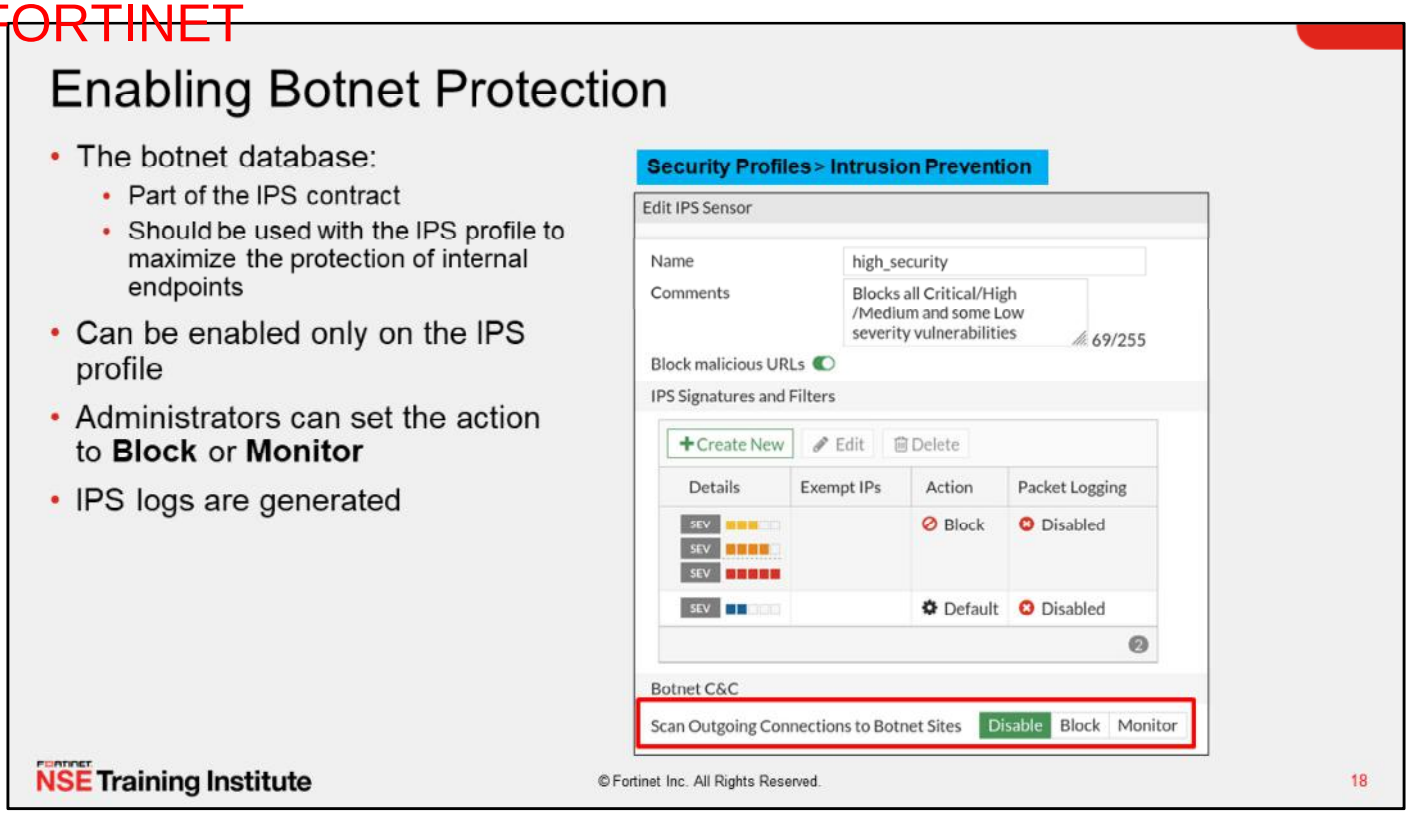

Since the botnet database is part of the FortiGuard IPS contract, administrators can enable scanning of botnet connections to maximize their internal security. You enable botnet scanning on the IPS profile that you applied the firewall policy on. You can also enable scanning of botnet connections using 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

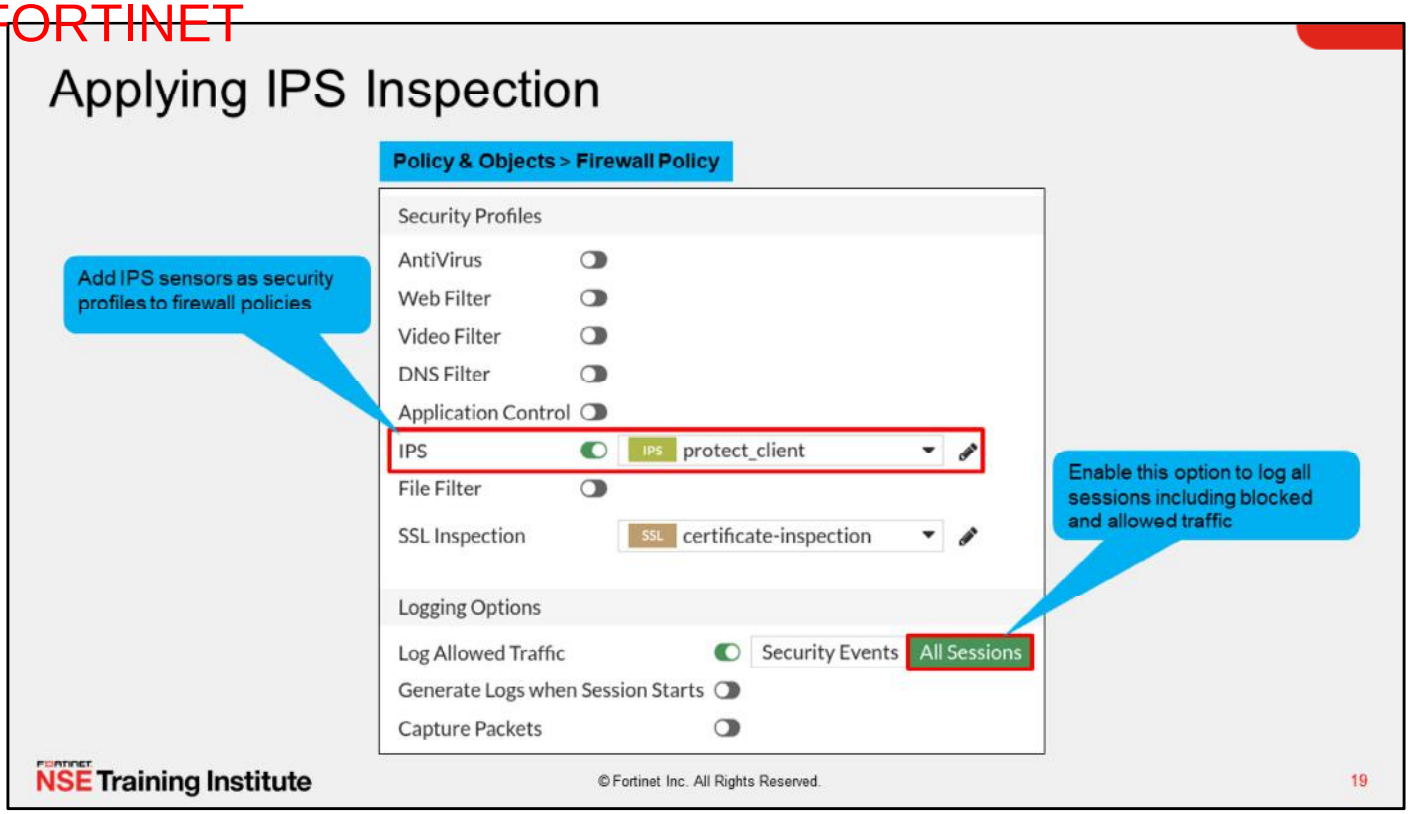

To apply an IPS sensor, you must enable **IPS** and then select the sensor in a firewall policy. By default, FortiGate logs all security events. This means you can see any traffic that is being blocked by IPS.

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.

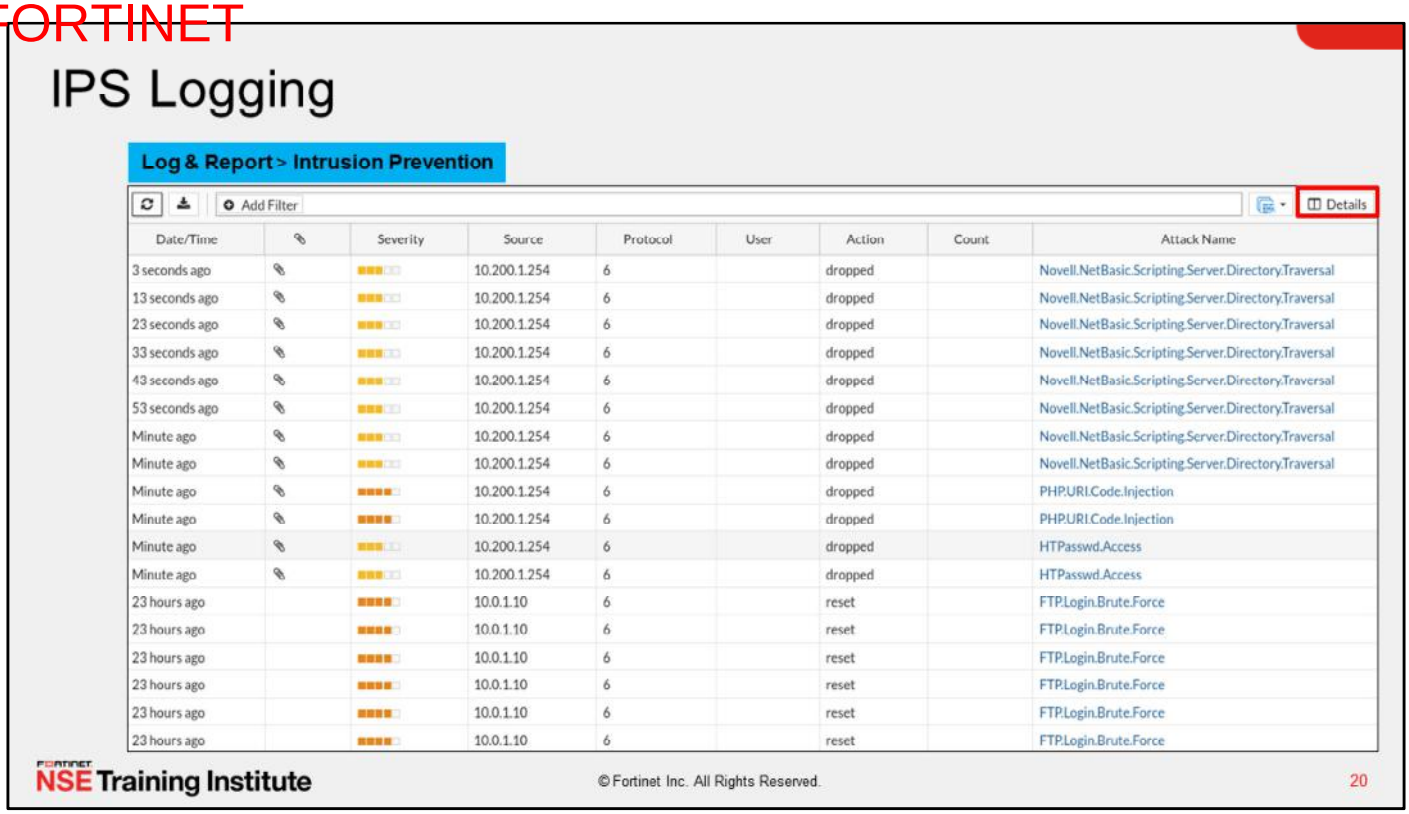

Date/Time		Severity	Source	Protocol	User	Action	Count	Attack Name
3 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
13 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
23 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
33 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
43 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
53 seconds ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
Minute ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
Minute ago			10.200.1.254	6		dropped		Novell.NetBasic.Scripting.Server.Directory.Traversal
Minute ago			10.200.1.254	6		dropped		PHP.URI.Code.Injection
Minute ago			10.200.1.254	6		dropped		PHP.URI.Code.Injection
Minute ago			10.200.1.254	6		dropped		HTPasswd.Access
Minute ago			10.200.1.254	6		dropped		HTPasswd.Access
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force
23 hours ago			10.0.1.10	6		reset		FTP.Login.Brute.Force

If you enabled security events logging in the firewall policies that apply IPS, you can view IPS events by clicking **Log & Report** > **Intrusion Prevention**. The **Intrusion Prevention** log menu appears only if FortiGate has matched attack attempts with IPS signatures.

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.

Knowledge Check

1. Which IPS action allows traffic and logs the activity?
   A. Allow
   ✓ B. Monitor

2. Which IPS component is updated most frequently?
   A. Protocol decoders
   ✓ B. IPS signature database

NSE Training Institute

21

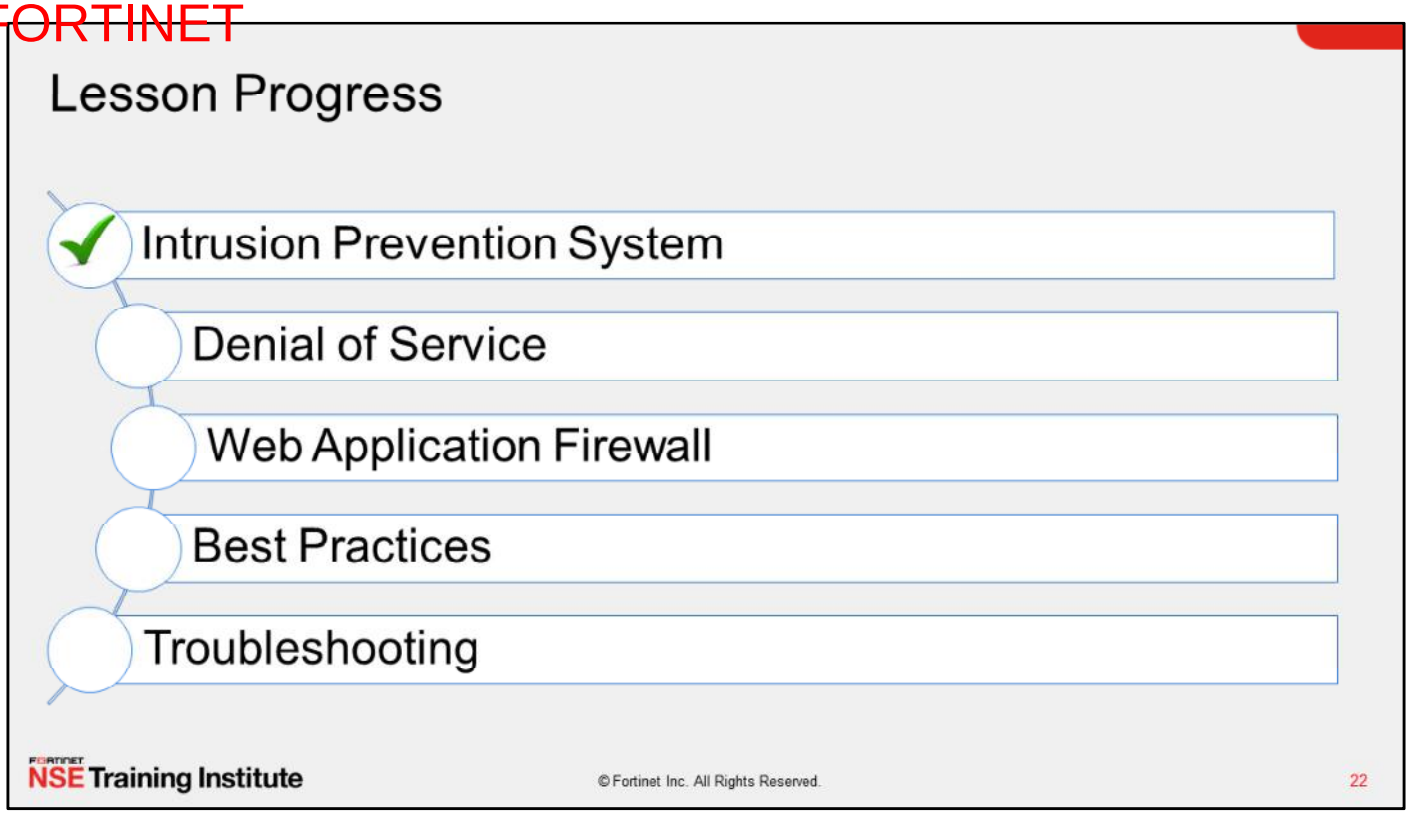

Good job! You now understand the IPS on FortiGate.

Now, you will learn about DoS.

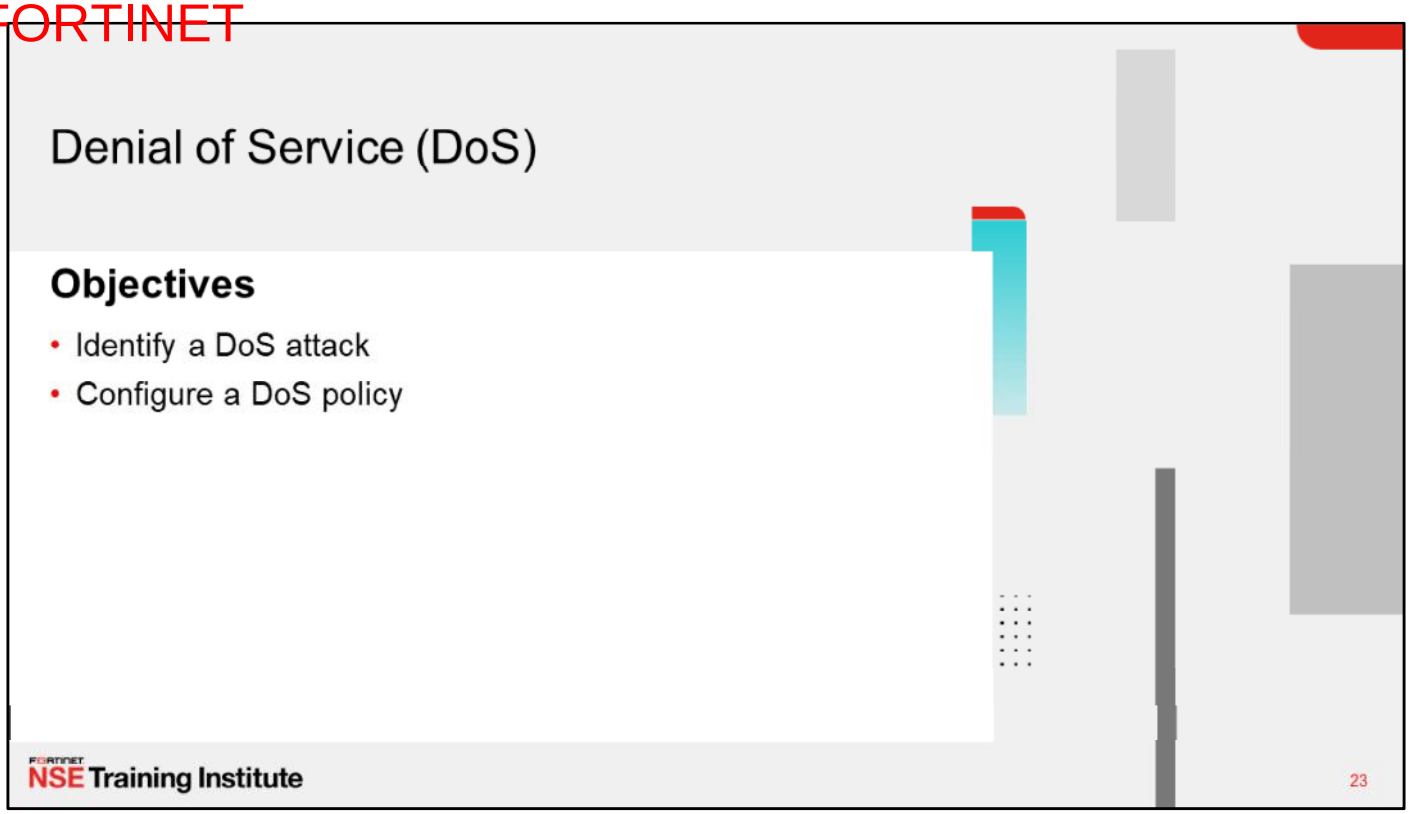

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in Denial of Service (DoS), you should be able to protect your network from common DoS attacks.

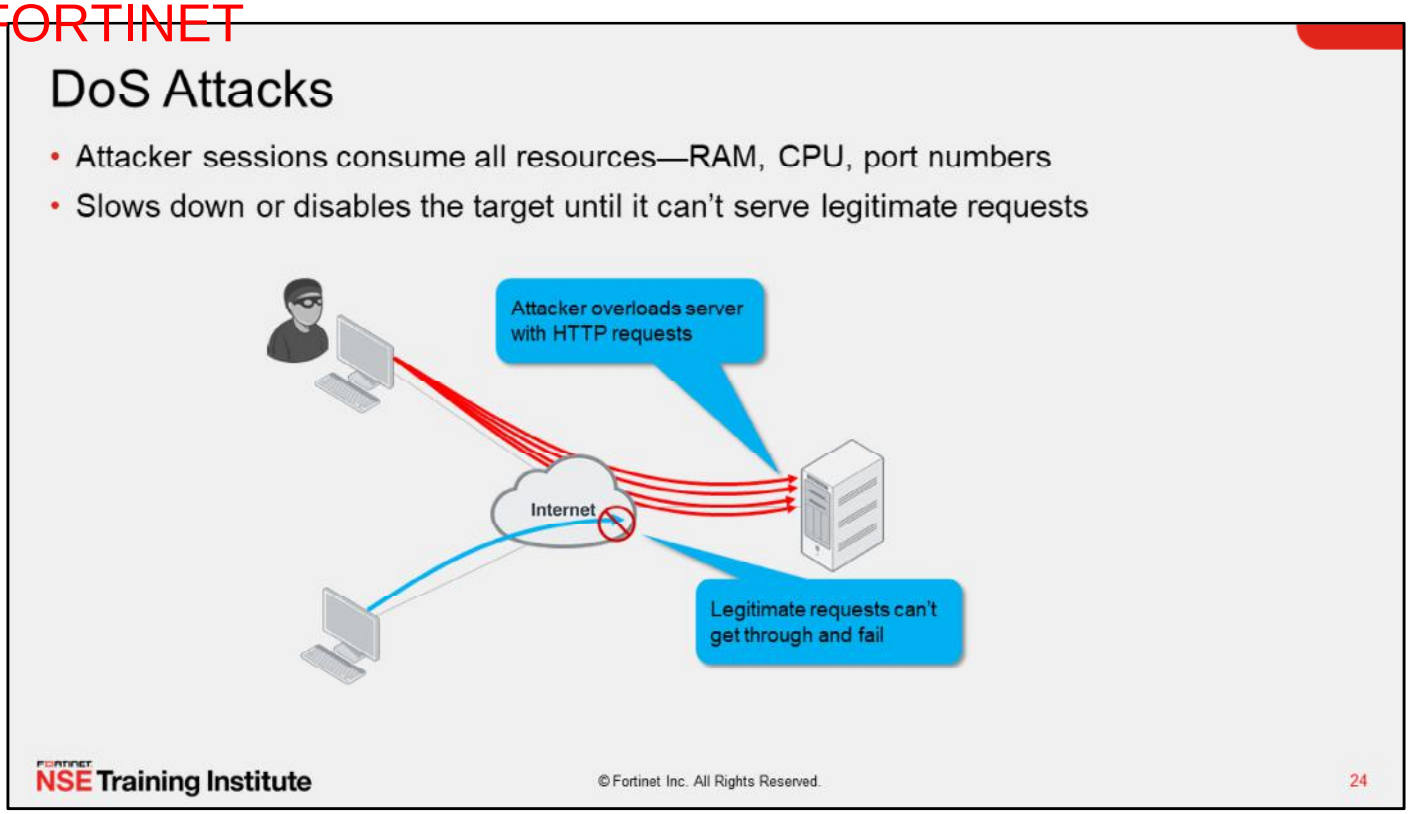

So far, you have learned about signatures that match illegal commands and invalid protocol implementations. Those are easy to confirm as attacks.

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.

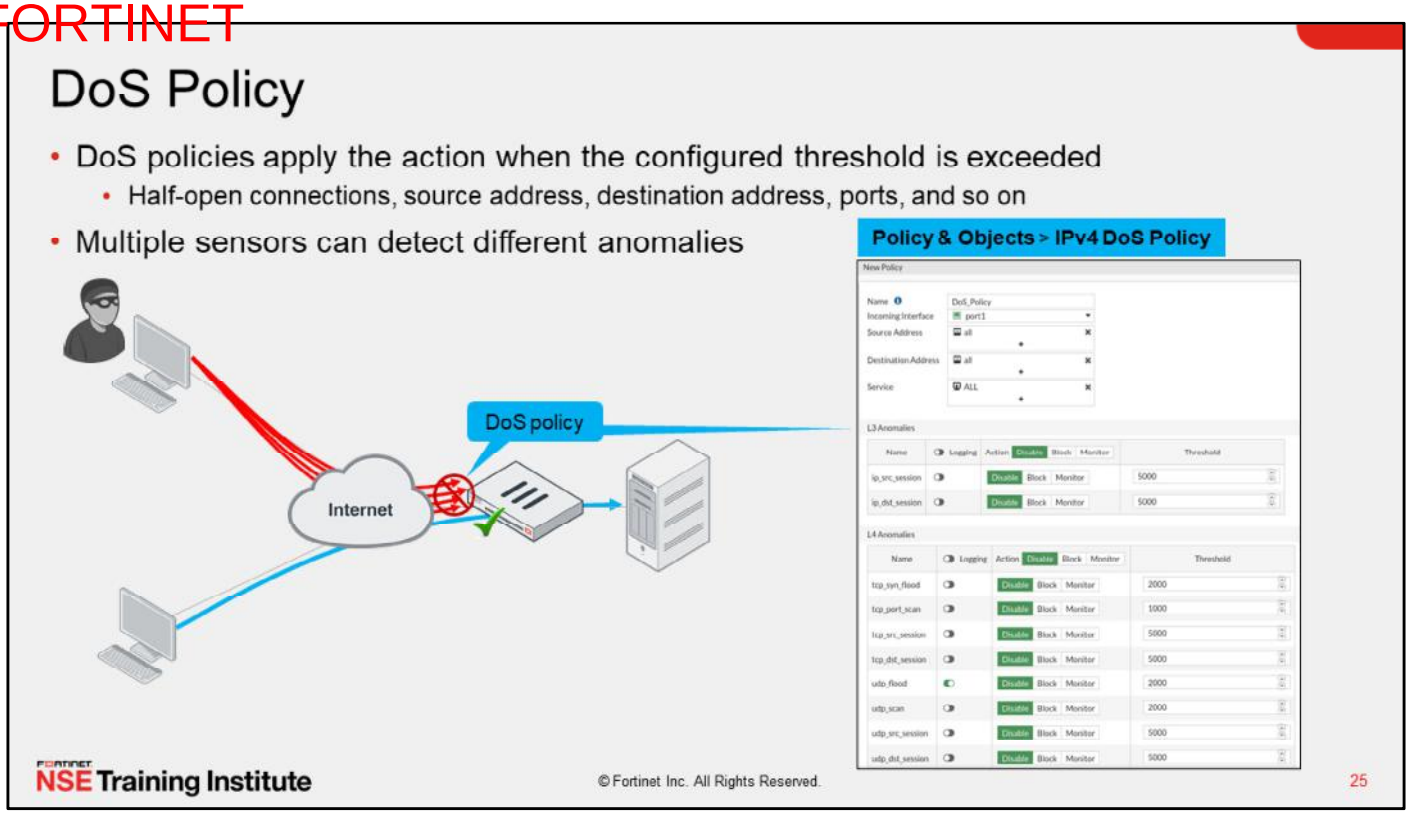

To block DoS attacks, apply a DoS policy on a FortiGate that is located between attackers and all the resources that you want to protect.

DoS filtering is done early in the packet handling process, which is handled by the kernel.

Types of DoS Attacks

- TCP SYN flood
  - Attacker floods victim with incomplete TCP/IP connection requests
  - The victim's connection table becomes full, so legitimate clients can't connect
- ICMP sweep
  - Attackers sends ICMP traffic to find targets
  - Attacker then attacks hosts that reply
- TCP port scan
  - Attacker probes a victim by sending TCP/IP connection requests to varying destination ports
  - Based on replies, attacker can map out which services are running on the victim system
  - Attacker then targets those destination ports to exploit the system

NSE Training Institute © Fortinet Inc. All Rights Reserved. 26

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.

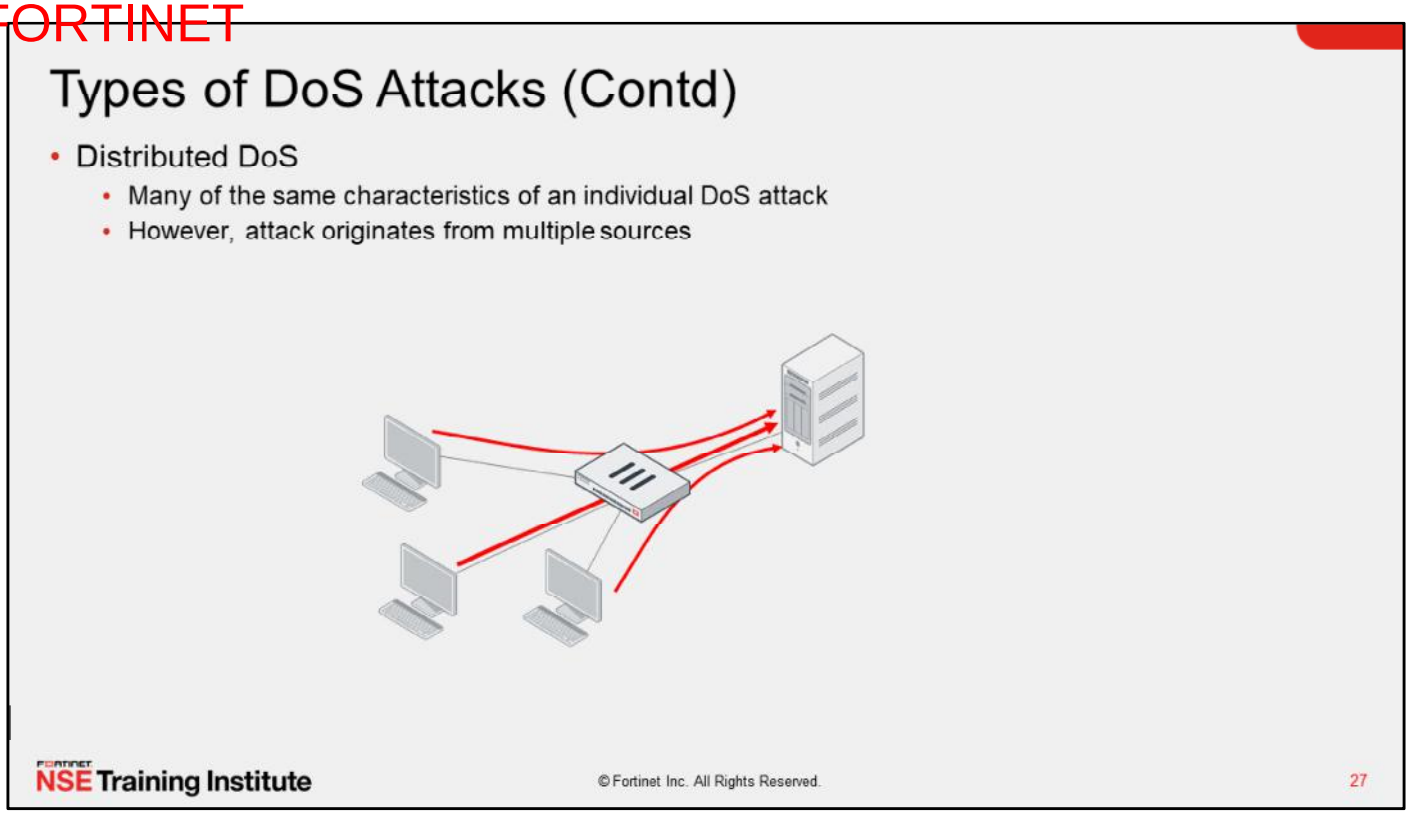

An individual DoS attack is a flood of traffic coming from a single address. It can originate from the internet, or even from your internal network. Typically, a single device makes many connections or sessions, and possibly uses much bandwidth to connect to a single location. A variation of this is 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

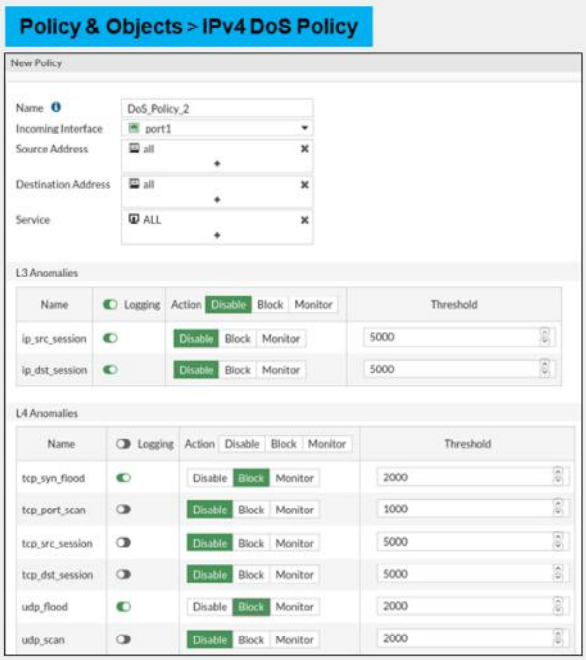

NSE Training Institute

28

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 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

2. 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

NSE Training Institute

29

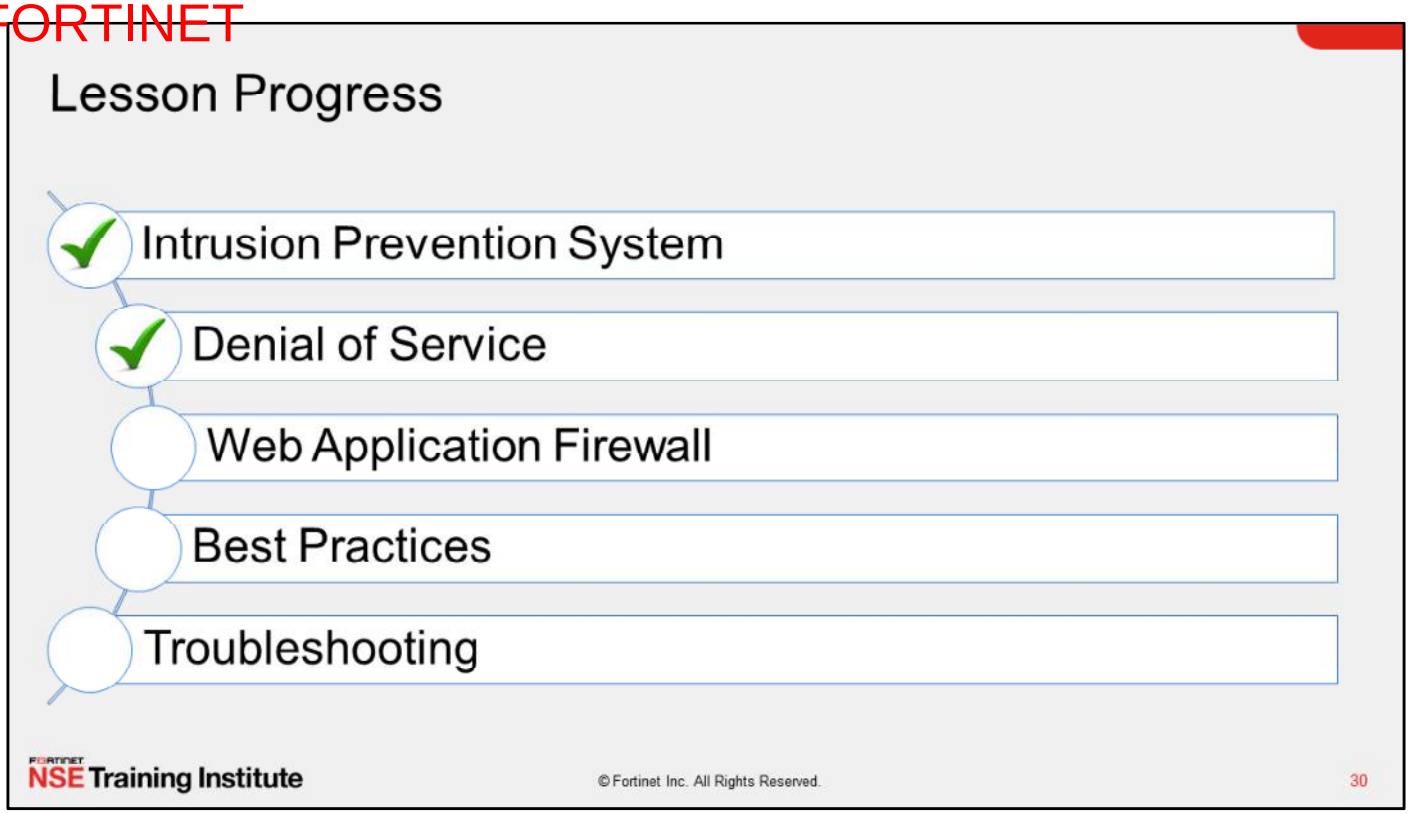

Good job! You now understand how to protect your network from DoS attacks on FortiGate.

Now, you will learn about WAF.

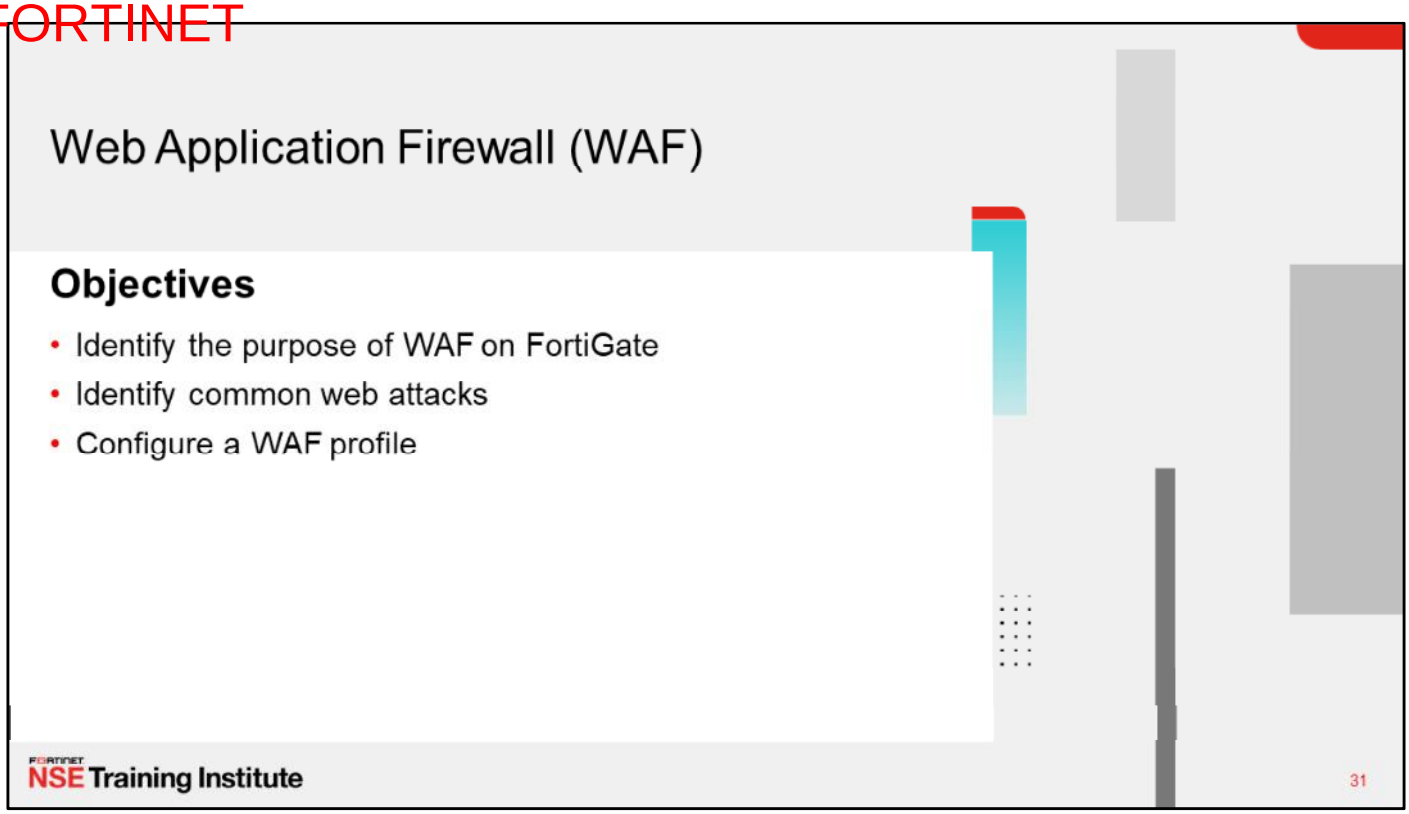

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in WAF, you should be able to apply the correct WAF inspection to protect the servers in your network.

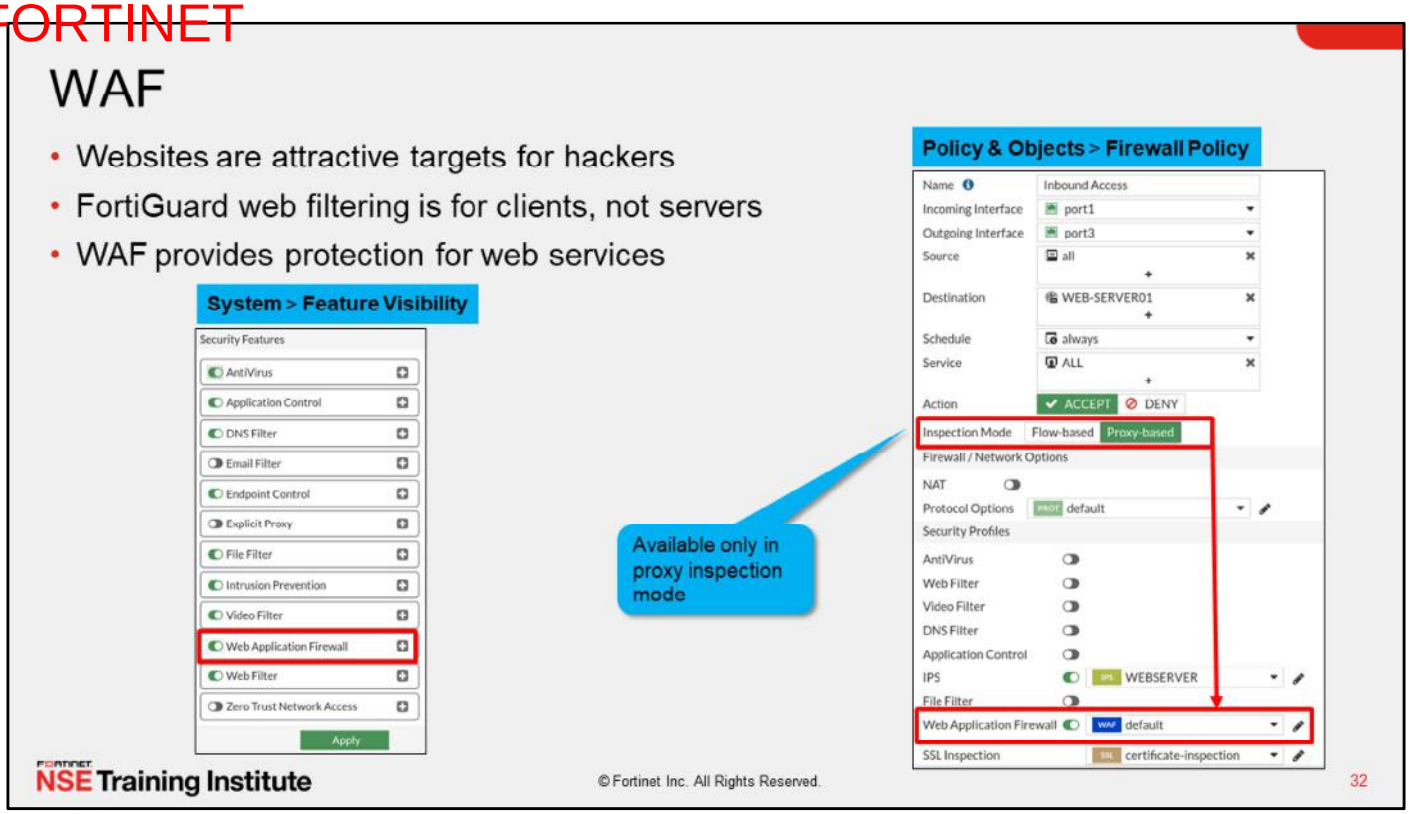

What is a WAF and why do you need it?

Some FortiGate features are meant to protect clients, not servers. For example, FortiGuard web filtering blocks requests based on the category of the server web pages. Antivirus prevents clients from accidentally downloading spyware and worms. Neither protects a server (which doesn't send requests—it receives them) from malicious scripts or SQL injections. Protecting web servers requires a different approach because they are subject to other kinds of attacks. This is where WAF applies.

The WAF feature is available only in proxy inspection mode.

Example of a Web Attack—Cross-Site Scripting

1. An attacker inputs JavaScript in an HTML form/parameter
2. The web app does not reject illegal input
3. Usually, the web app saves the input to a database
4. An innocent client requests a page that is retrieved from the database. The page:
   - Now includes malicious script
   - Can cause a client's browser to transmit to a third-party, malicious server

- The variety of attacks based on cross-site scripting (XSS) is limitless, but they commonly include transmitting private data, like authentication cookies or other session information, to the attacker

NSE Training Institute © Fortinet Inc. All Rights Reserved. 33

Take a look at some examples of attacks that target web applications specifically.

One type of attack is called cross-site scripting (XSS). If a web application doesn't sanitize its inputs and reject JavaScript, it ends up storing the XSS attack in its database. Then, when other clients request the page that reuses that data, the JavaScript is now embedded in the page.

JavaScript can do many things with a page, including rewriting the whole page and making its own requests. This is the basic mechanism of asynchronous JavaScript and XML (AJAX) apps. In this case, XSS causes innocent clients to transmit to a different server that is controlled by the attacker. This could, for example, transmit credit card information or passwords from an HTTP form to the attacker.

Example of a Web Attack—SQL Injection

- SQL statements are inserted into entry fields of a web application
- The web application doesn't reject illegal input
- When the web application connects to the database to add input, it can:
  - Download sensitive data from the database (`select * from USERS`)
  - Modify database (insert/update/delete)
  - Perform administrative operations (close management interface)

NSE Training Institute

34

Another very common web attack is a SQL injection. Just like an XSS attack, the root cause of a SQL injection is that the web application doesn't sanitize input. If the attacker enters a SQL query into an input, such as an HTML form, the web app simply accepts it, and passes it along to the database engine, which accidentally runs the query.

The SQL language can do anything to the data. It can, for example, download the table of users so that the attacker can run a password cracker. A query could add new entries for new administrator login attempts, or modify login attempts, blocking administrators from logging in.

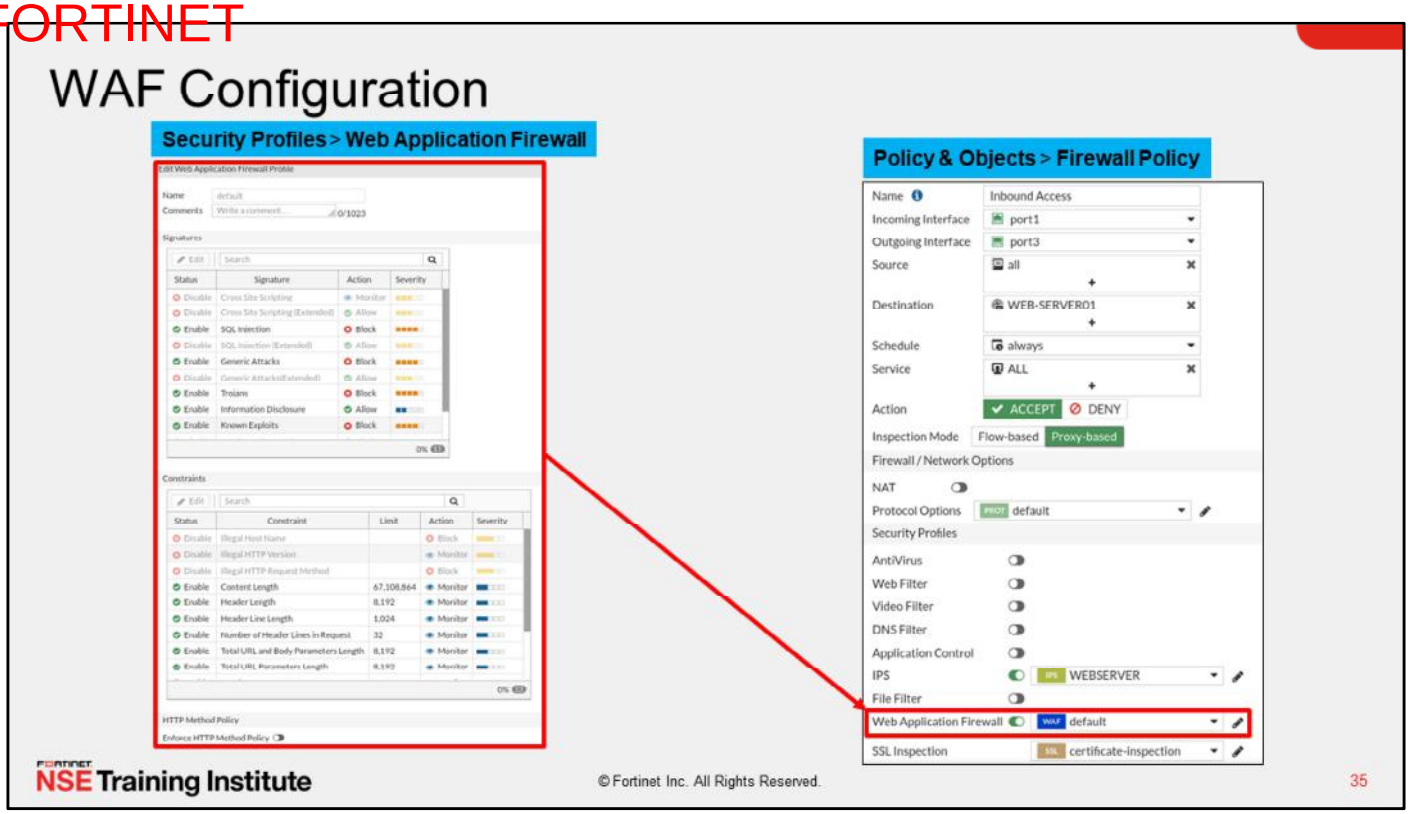

One component of a WAF profile is WAF signatures. WAF signatures work in the same way as IPS signatures. FortiGate can act on the traffic that matches any of them. Some WAF signatures are categorized as extended. They are more likely to cause false positives, but are sometimes required in high-security environments.

HTTP constraints can monitor and control the number, type, and length of many HTTP headers, which are also inputs. This prevents unexpected inputs that a malicious client could craft to compromise your server. The limits can vary by your server's software, but also by its hardware. If a server has limited RAM, for example, then it is potentially easier to overload or crash with an excessive number of headers, because parsing the headers and storing them in buffers requires RAM.

After you configure a WAF profile, it is assigned to one or more firewall policies.

FortiWeb

- Provides more specialized web server protection
- More complete protocol understanding
- HTTP state attack protection
- HTTP vulnerability scans/penetration tests
- HTTP rewriting and application delivery (basic ADC)
- Better performance for high HTTP traffic

NSE Training Institute

36

FortiWeb is a specialized WAF device. For environments where the protection of web services is critical, you can complement FortiGate with FortiWeb.

FortiWeb offers a more complete HTTP protocol understanding and state attack protection. It can perform vulnerability scans and penetration tests. It can also rewrite HTTP packets, and route traffic based on HTTP content.

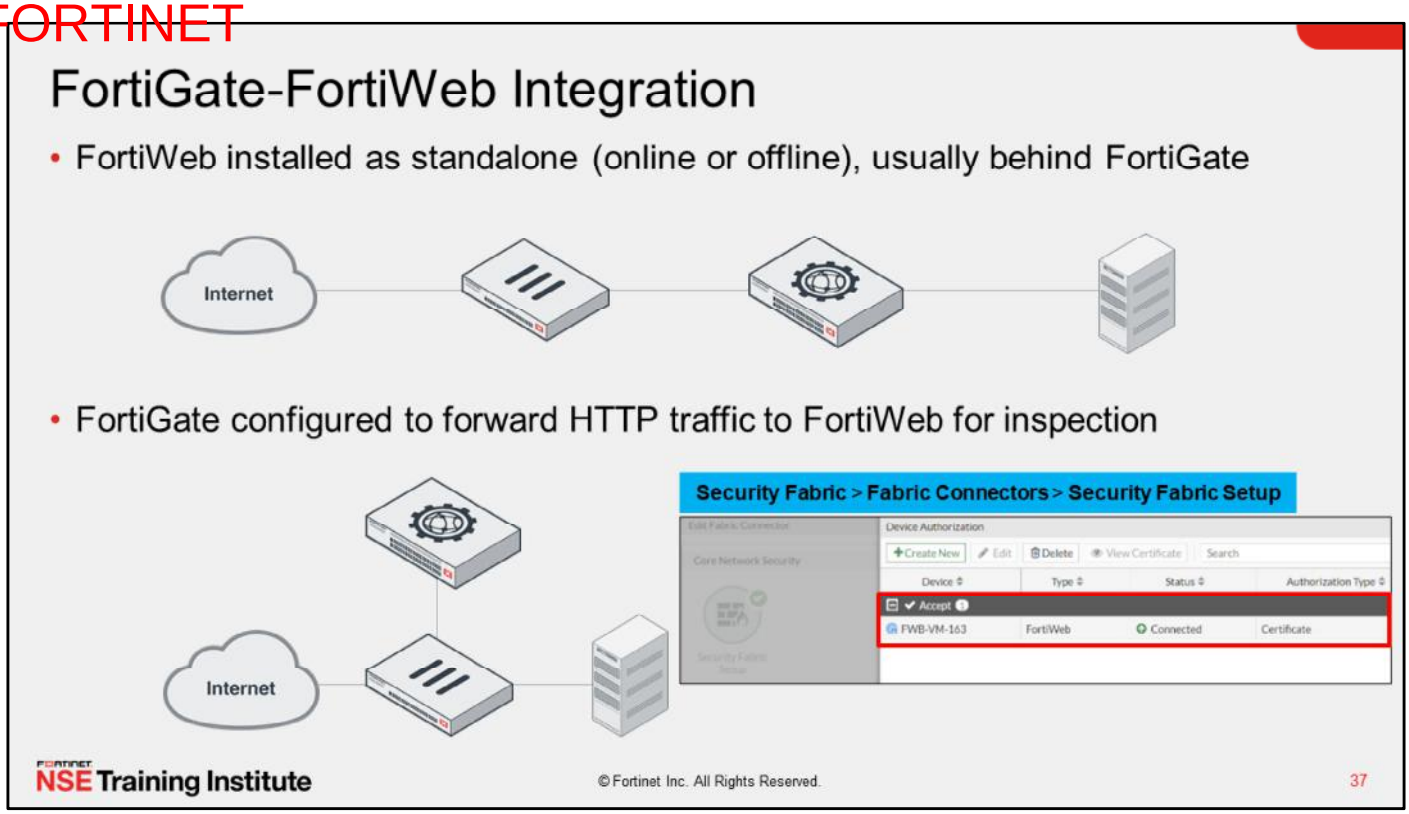

In most cases, FortiWeb is installed as a standalone device, usually located between FortiGate and the protected web servers. You can install FortiWeb online (web traffic crossing the device) or offline (device is connected as a one-arm sniffer).

Alternatively, you can configure FortiGate to forward web traffic to an external FortiWeb, where the WAF inspection happens. This is useful, for example, when you must protect servers located in multiple sites with a single FortiWeb. With this setup, FortiGate forwards all web traffic to the FortiWeb if the traffic matches a firewall policy configured with a WAF profile enabled for external inspection.

For detailed information about FortiWeb, see the NSE 6 FortiWeb training material.

Knowledge Check

1. WAF protocol constraints protect against which type of attacks?
   - ✓ A. Buffer overflow
   - B. ICMP Sweep

2. To use the WAF feature, which inspection mode should be used in the firewall policy?
   - A. Flow
   - ✓ B. Proxy

NSE Training Institute

38

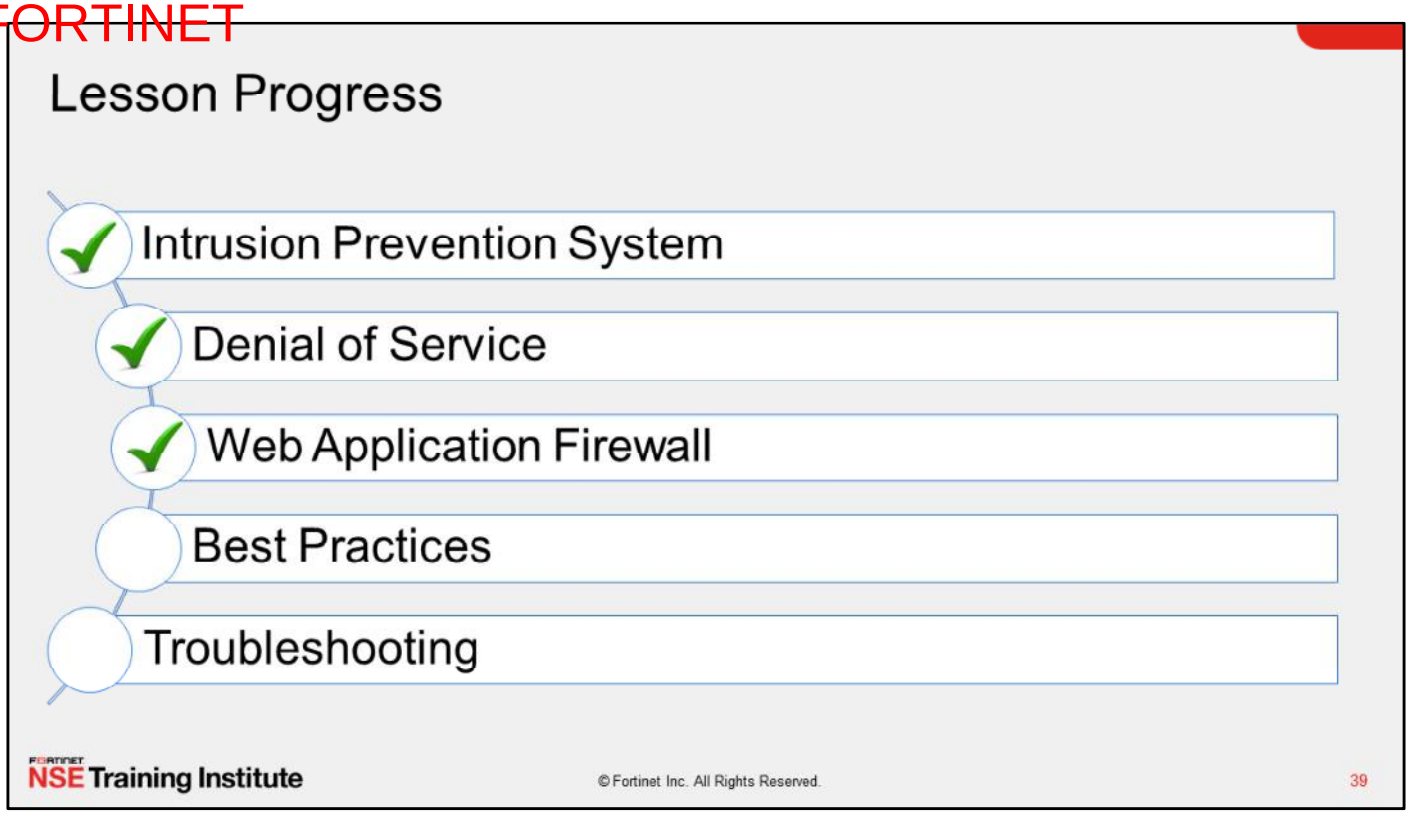

Good job! You now understand how to protect your servers using WAF on FortiGate.

Now, you will learn about IPS best practices.

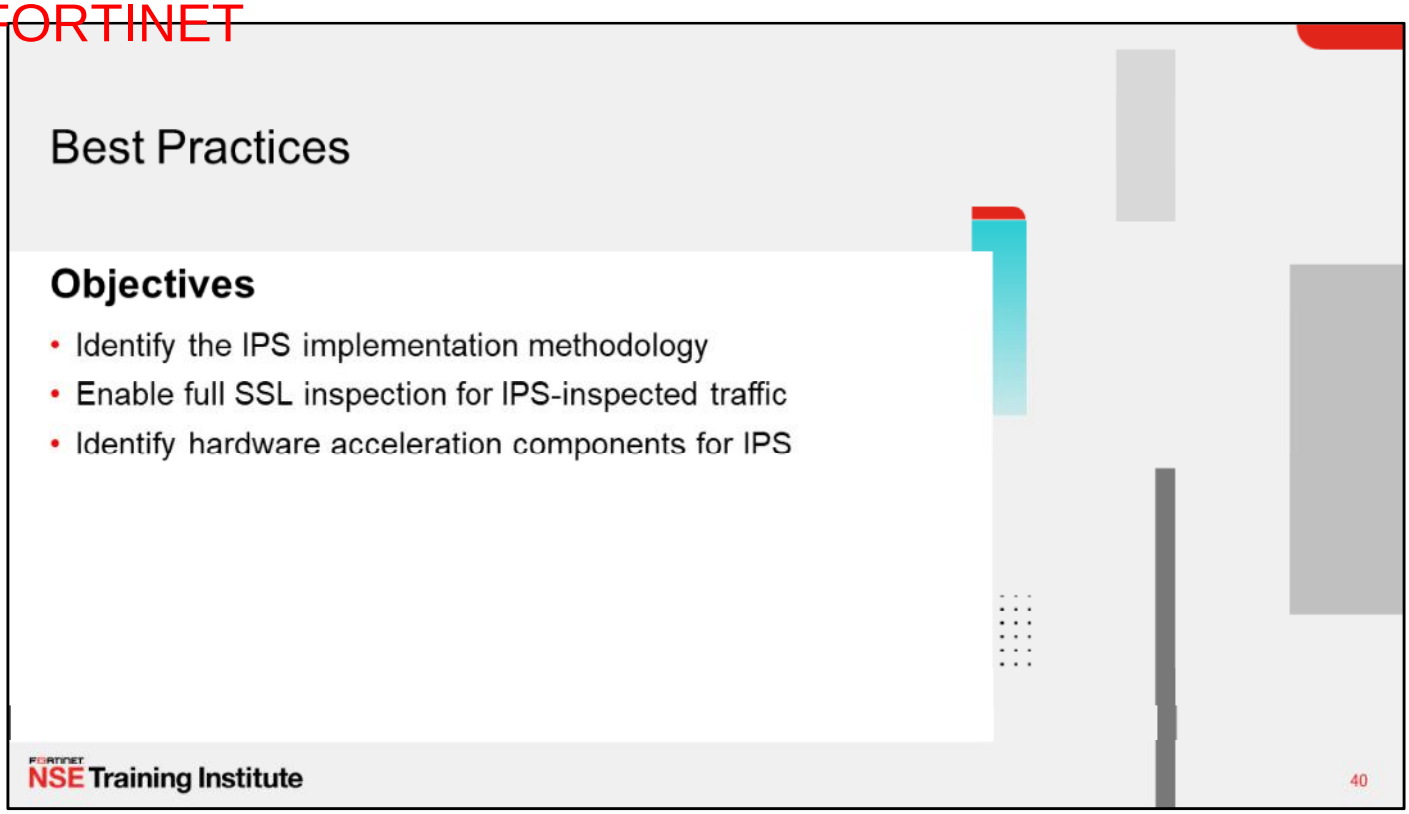

After completing this section, you should be able to achieve the objectives shown on this slide.

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.

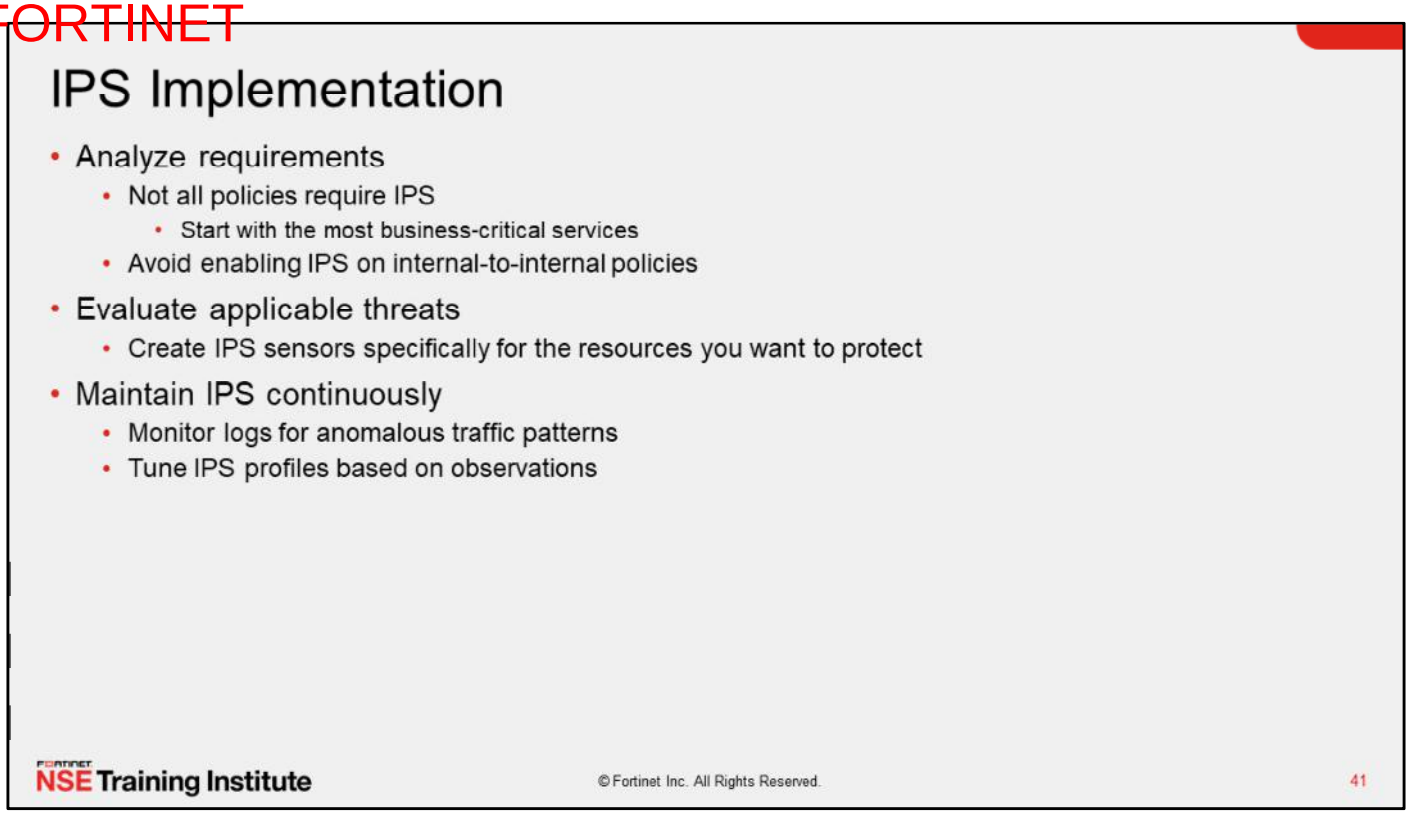

Before you implement IPS, you must analyze the needs of your network. Enabling the default profiles across all policies quickly causes issues, the least of which are false positives. Performing unnecessary inspections on all network traffic can cause high resource utilization, which can hamper the ability 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.

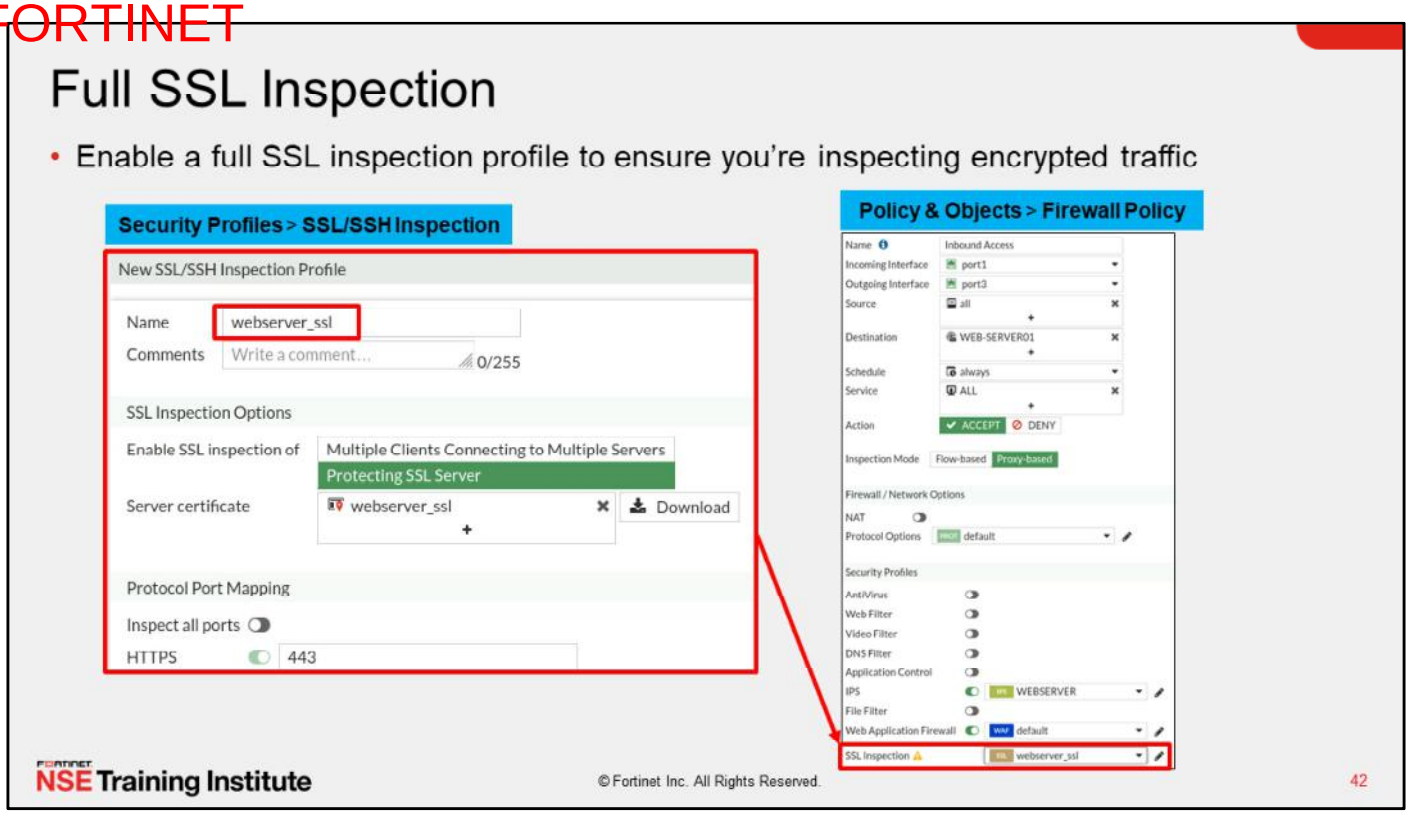

Certain vulnerabilities apply only to encrypted connections. In some of these cases, FortiGate can't identify the threat reliably if it can't parse the payload. For this reason, you must use an SSL inspection profile if you want to get the maximum benefit from your IPS and WAF features.

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 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

NSE Training Institute

43

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

NSE Training Institute

44

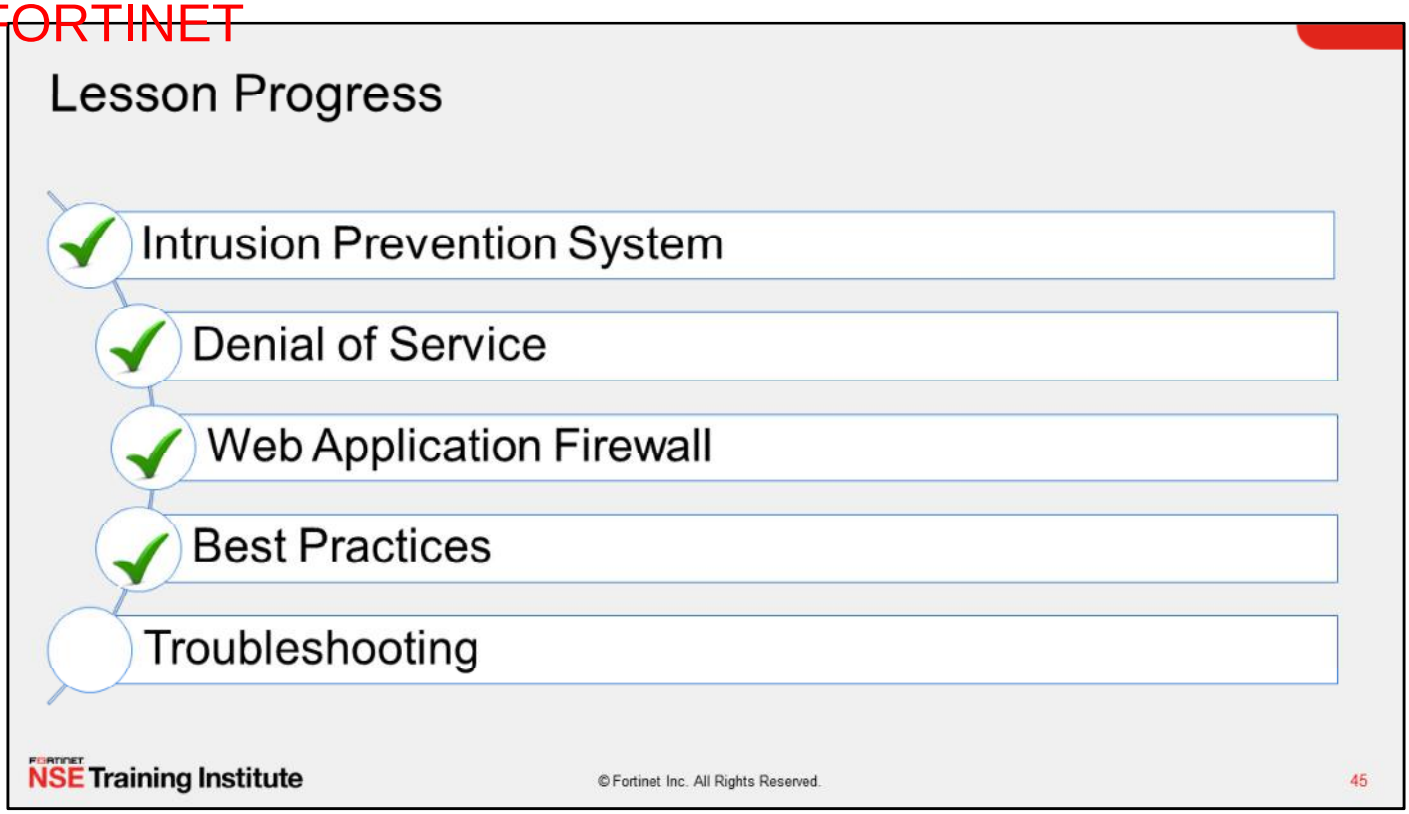

Good job! You now understand some best practices for IPS implementation on FortiGate.

Now, you will learn about IPS troubleshooting.

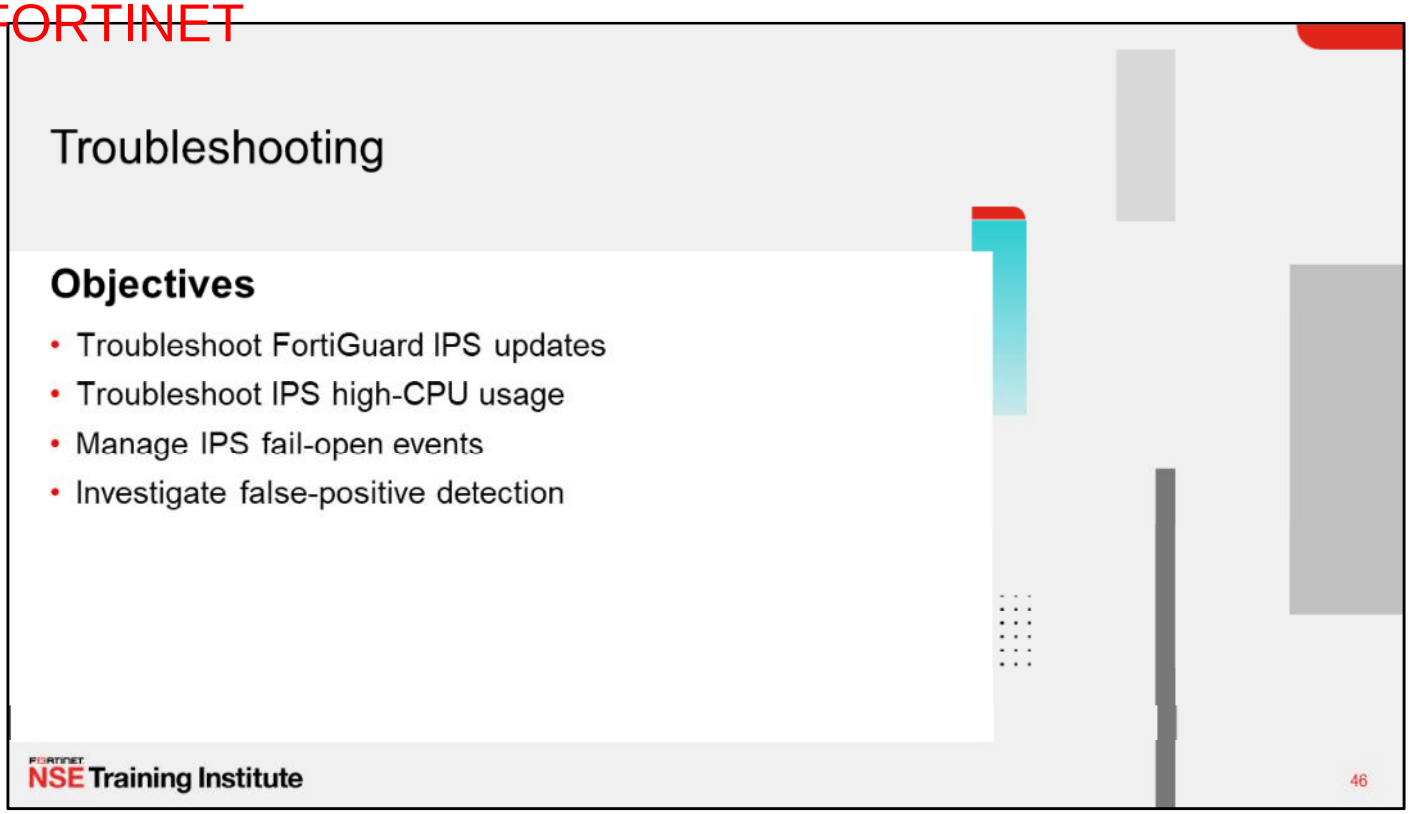

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in troubleshooting, you should be able to identify, investigate, and manage some common issues with IPS deployments on FortiGate.

FortiGuard IPS Troubleshooting

- All IPS update requests are sent to `update.fortiguard.net` on TCP port 443
  - Can be configured to connect through a web proxy (CLI only):
    - `config system autoupdate tunneling`
- Verify update status on GUI

- Enable real-time debug on CLI

```
# diagnose debug application update -1
# diagnose debug enable
# execute update-now
```

After enabling real-time debugging, force a manual update of all FortiGuard packages

NSE Training Institute © Fortinet Inc. All Rights Reserved. 47

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.

IPS and High-CPU Use

```
# diagnose test application ipsmonitor <Integer>

1: Display IPS engine information
2: Toggle IPS engine enable/disable status
3: Display restart log
4: Clear restart log
5: Toggle bypass status
6: Submit attack characteristics now
10: IPS queue length
11: Clear IPS queue length
12: IPS L7 socket statistics
13: IPS session list
14: IPS NTurbo statistics
15: IPSA statistics
97: Start all IPS engines
98: Stop all IPS engines
99: Restart all IPS engines and monitor
```

Shuts down IPS engine completely

IPS engine remains active, but does not inspect traffic

NSE Training Institute © Fortinet Inc. All Rights Reserved. 48

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 Fail Open

- Fail open is triggered when the IPS socket buffer is full and new packets can't be added for inspection

```
config ips global
 set fail-open <enable|disable>
 ...
end
```

- IPS fail open entry log:

```
date=2021-04-21 time=09:07:59 logid=0100022700 type=event subtype=system
level=critical vd="root" logdesc="IPS session scan paused" action="drop"
msg="IPS session scan, enter fail open mode"
```

Packets dropped!

- When troubleshooting IPS fail-open events, try to identify a pattern
  - Has the traffic volume increased recently?
  - Does fail open trigger at specific times during the day?
- Create IPS profiles specifically for the traffic type
  - An IPS sensor configured to protect Windows servers doesn't need Linux signatures
  - Disable IPS on internal-to-internal policies

NSE Training Institute © Fortinet Inc. All Rights Reserved. 49

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.

False-Positive Detection

- Check the logs to determine which signature is triggering the false-positive
- Use IP exemptions on the signature as a temporary bypass for the affected endpoints
- Collect samples of the traffic:
  - Use the **Packet Logging** action
- Provide the traffic samples and the IPS logs to the FortiGuard team for further investigation

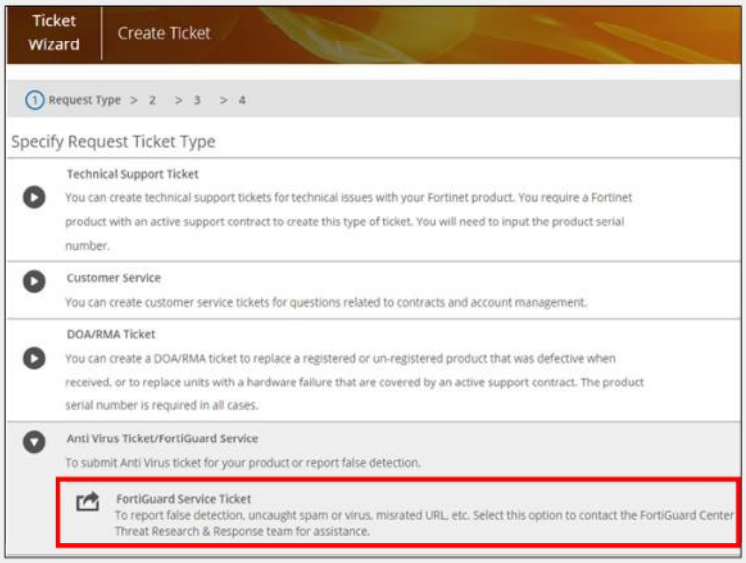

NSE Training Institute © Fortinet Inc. All Rights Reserved. 50

In the event of a false-positive detection, first identify which signature is generating them. You should also verify that the traffic is hitting the correct policy and IPS sensor. After you verify these factors, you should gather samples of the traffic. Use the **Packet Logging** action on the signature. Provide the traffic samples and the matching IPS logs to the FortiGuard team for further investigation.

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 dropped

NSE Training Institute

51

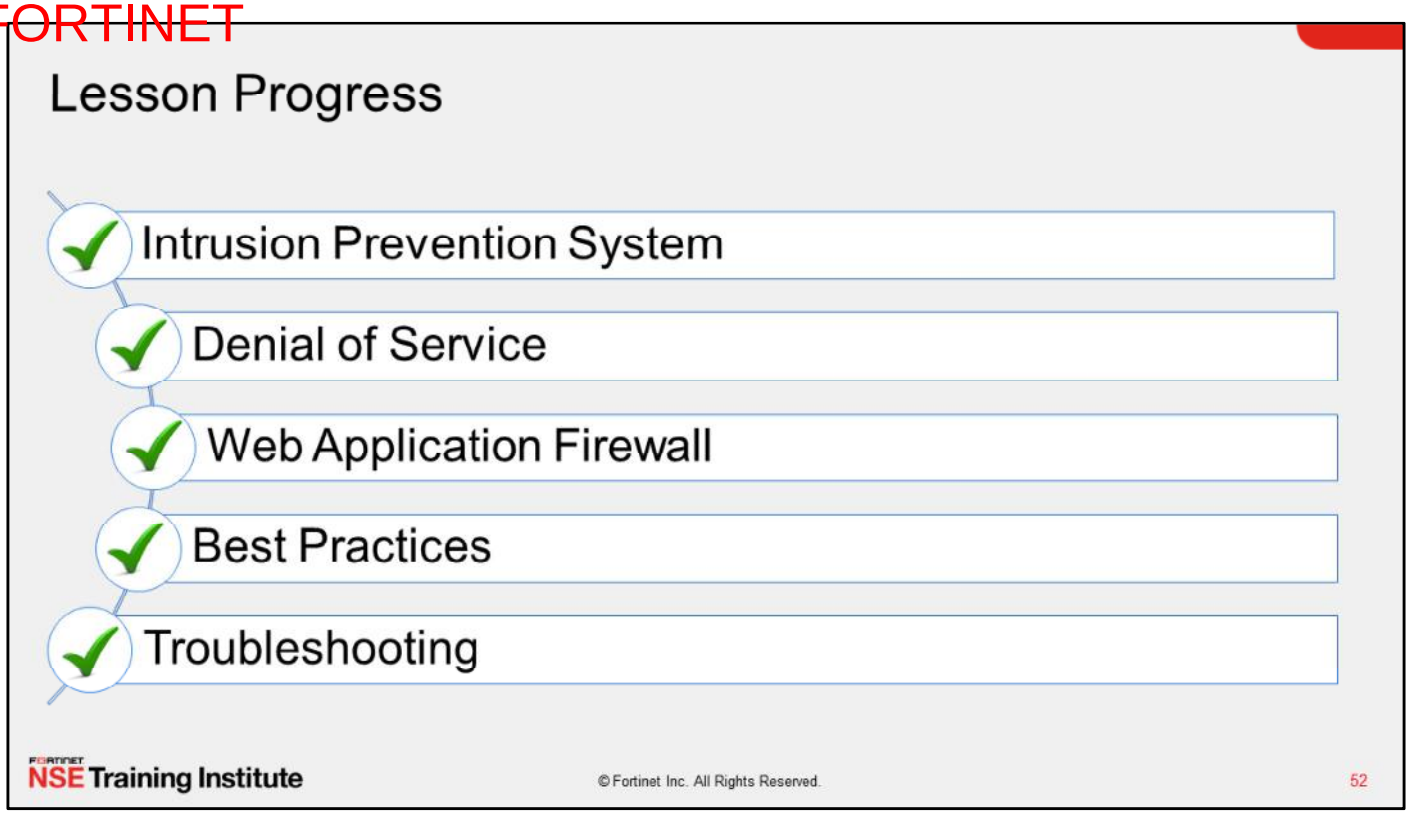

Congratulations! You have completed this lesson.

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

Review

- ✓ Manage FortiGuard IPS updates
- ✓ Configure an IPS sensor
- ✓ Apply IPS to network traffic
- ✓ Identify a DoS attack
- ✓ Configure a DoS policy
- ✓ Identify common web attacks
- ✓ Configure a WAF profile
- ✓ Identify IPS implementation methodology
- ✓ Troubleshoot common IPS issues

NSE Training Institute © Fortinet Inc. All Rights Reserved. 53

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.

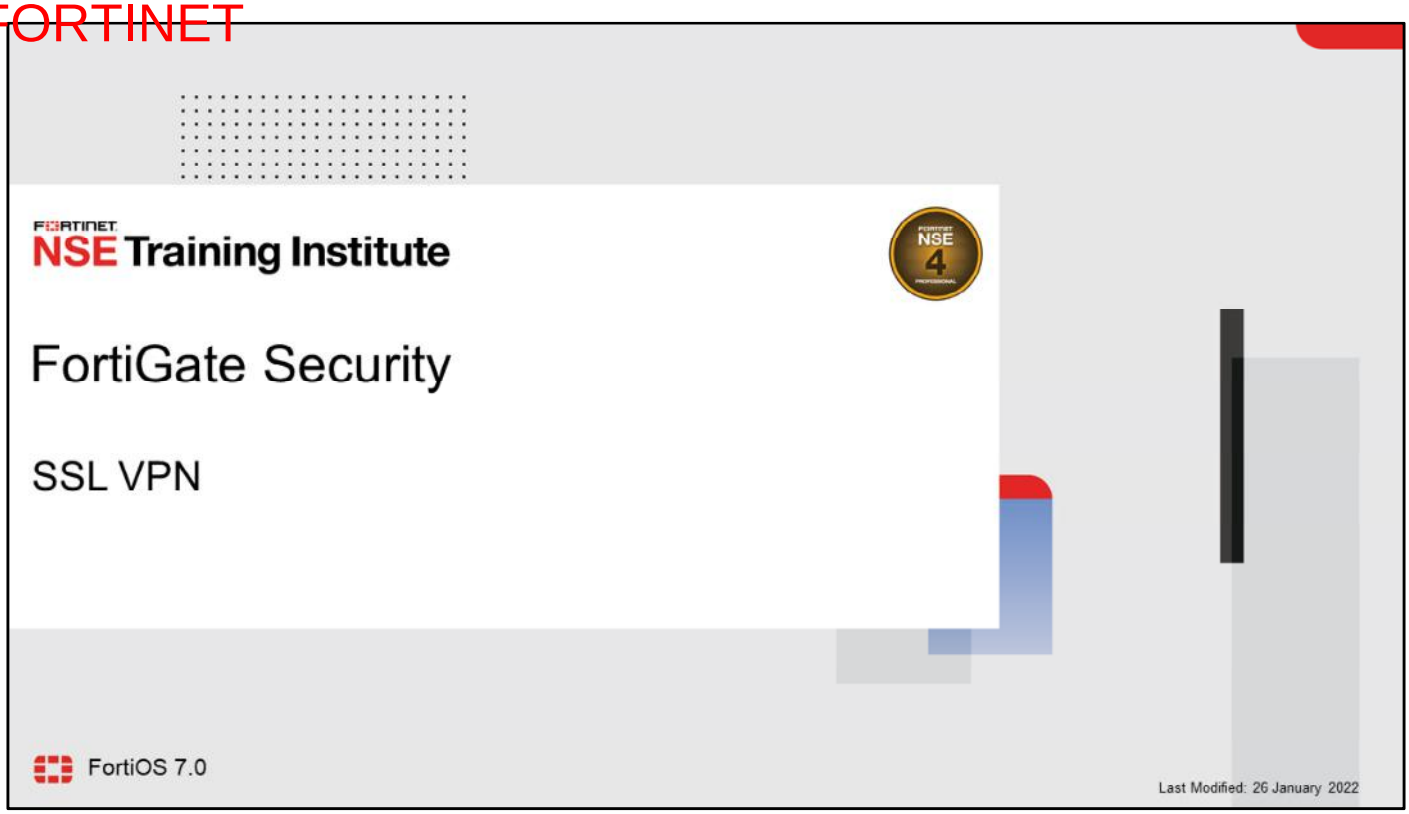

In this lesson, you will learn how to configure and use SSL VPNs. SSL VPNs are an easy way to give remote users access to your private network.

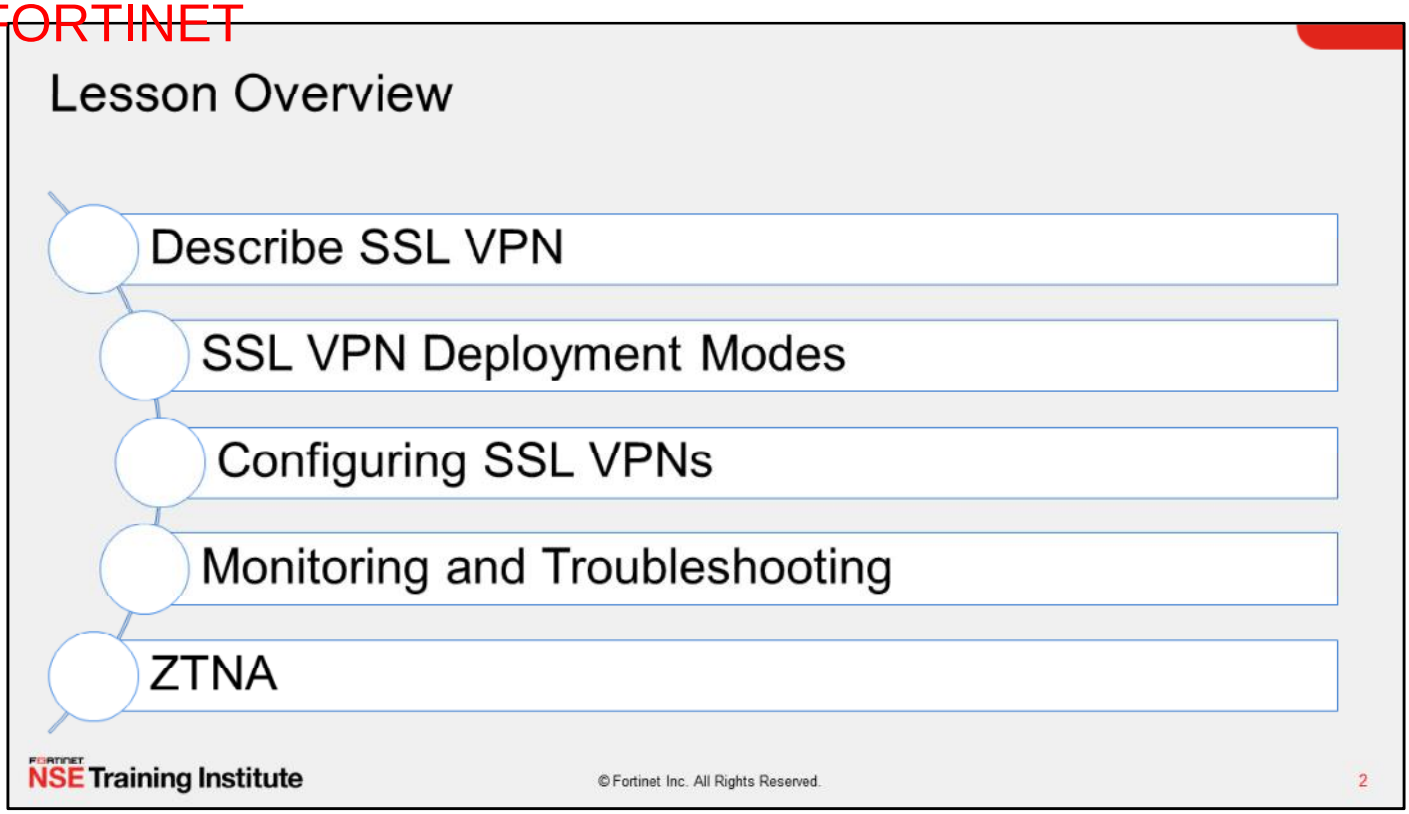

In this lesson, you will learn about the topics shown on this slide.

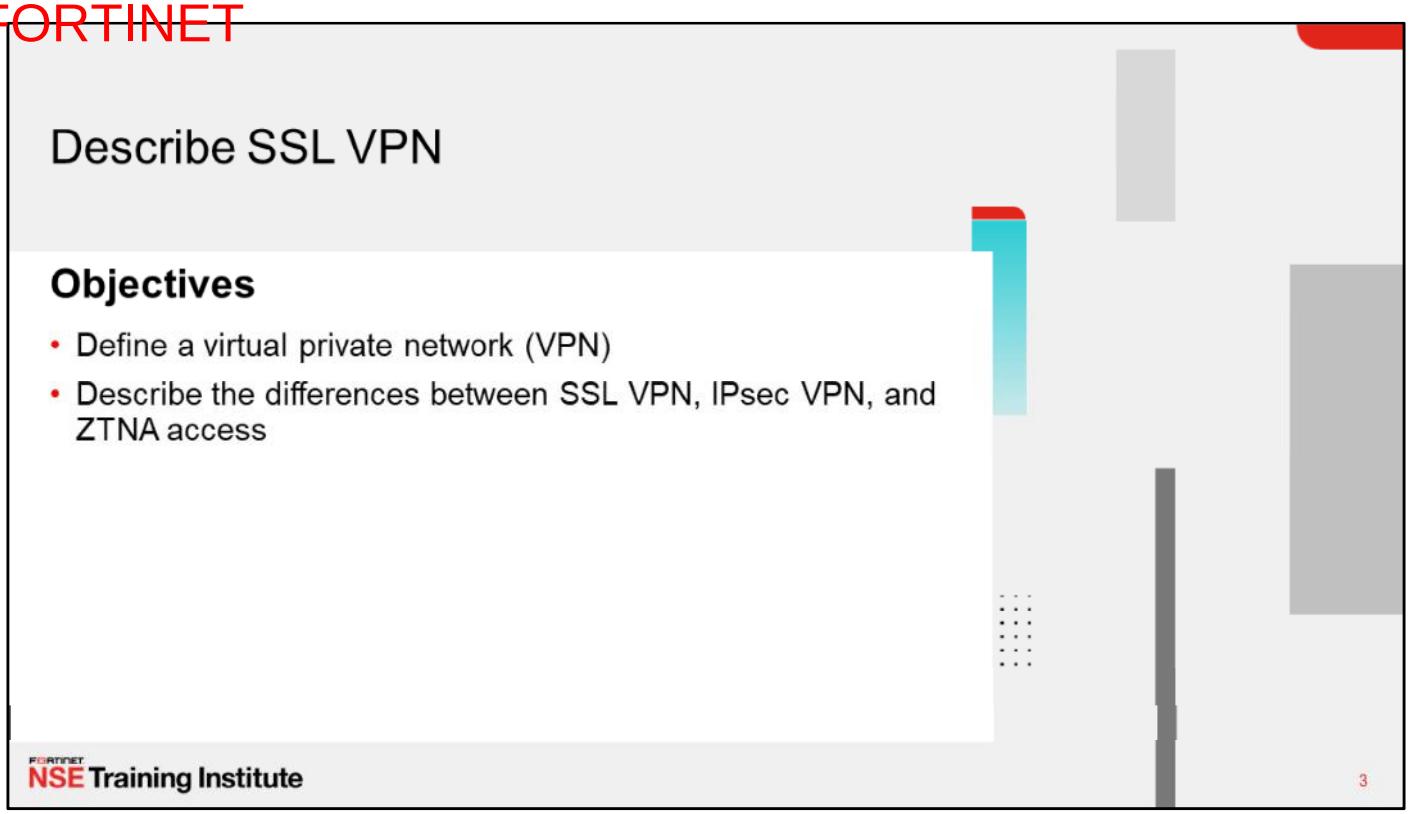

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding VPN concepts, you will be able to more effectively understand how FortiGate manages SSL VPN methods.

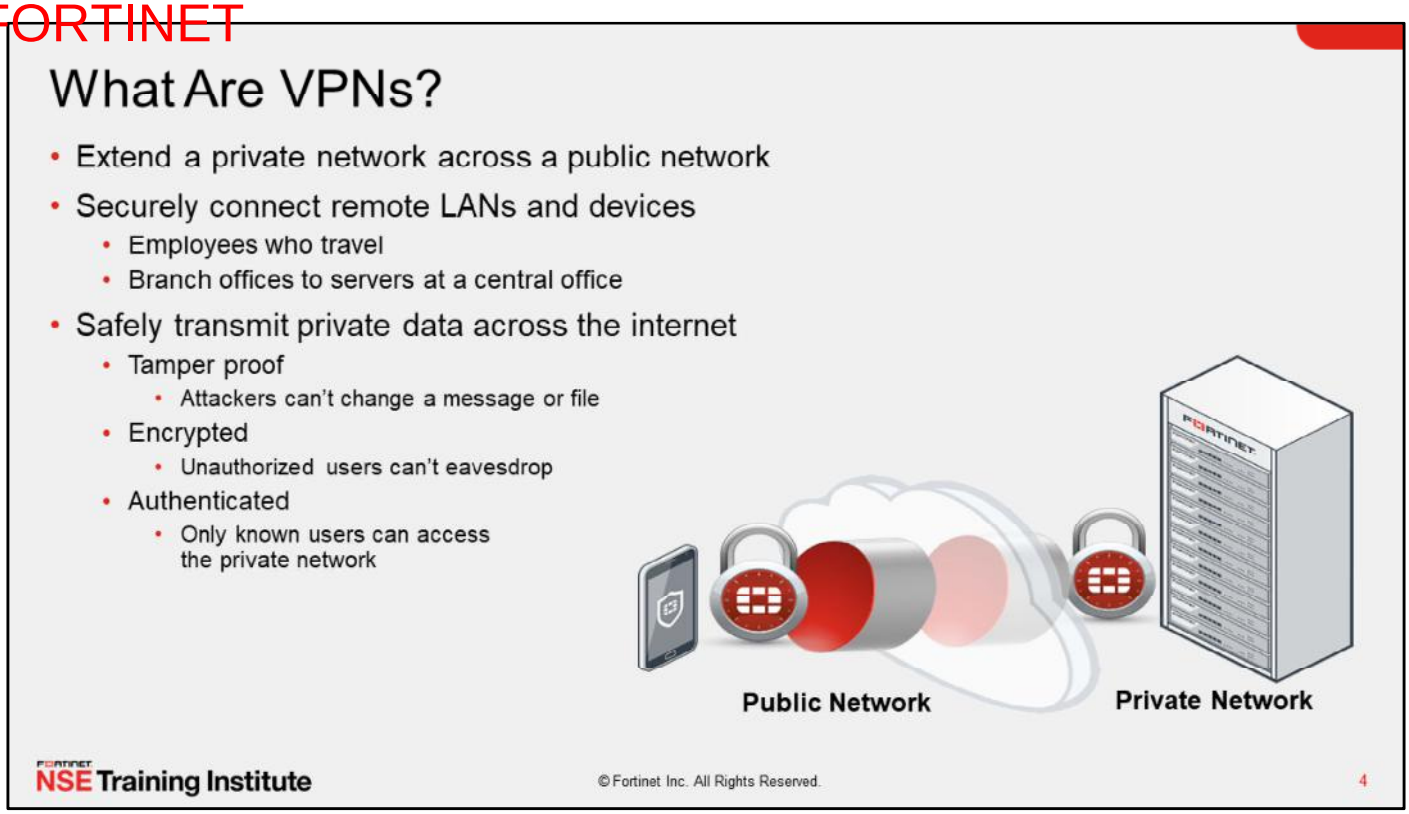

A VPN creates a tunnel that gives users or remote LANs secure access to your private network, as if they were connected to your LAN.

A VPN is often used when LANs are separated by an untrusted public network, such as the internet. As well as providing users with secure access to private networks while they are traveling, a VPN can also interconnect branch office networks located across the internet, and even on the other side of the world.

User data inside a VPN tunnel is encrypted for privacy. It cannot be read, even if it is intercepted by unauthorized users. VPNs also use security methods to ensure that only authorized users can establish a VPN and access the private network's resources. They typically also provide tamper proofing.

Most VPNs are SSL or IPsec VPNs. FortiOS supports both, as well as less common, weaker VPNs such as PPTP. In this lesson, we will focus on SSL VPNs.

Comparing SSL VPN, IPsec VPN, and ZTNA Access

	IPsec VPN	SSL VPN	ZTNA
Tunnel type:	IPsec tunnel only	Session-based OR tunnel	Session-based only
Configured between:	FortiClient and FortiGate FortiGate and FortiGate FortiGate and compatible third-party IPsec VPN gateway FortiGate and compatible third-party IPsec VPN clients	Browser and FortiGate FortiClient and FortiGate FortiGate (SSL Client) and FortiGate (SSL Server)	Browser and FortiGate FortiClient and FortiGate (TCP forwarding access)
Log in through:	IPsec client	HTTPS web page on FortiGate FortiClient FortiGate (SSL Client)	HTTPS hostname or IP and port number FortiClient (TCP forwarding access)

NSE Training Institute © Fortinet Inc. All Rights Reserved. 5

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

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

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

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

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

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

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

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

NSE Training Institute

6

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

In general, IPsec VPN is preferred when tunnels must be up continuously and interoperate with many types of devices, while SSL VPN is preferred when people travel and need to connect to the office.

Knowledge Check

1. What does a VPN do?
   - ✓ A. Extends a private network across a public network
   - B. Protects a network from external attacks

2. Which statement about SSL VPNs is true?
   - A. An SSL VPN can be established between workstation and a FortiGate device only.
   - ✓ B. An SSL VPN can be established between an end-user workstation and a FortiGate device or two FortiGate devices.

NSE Training Institute

7

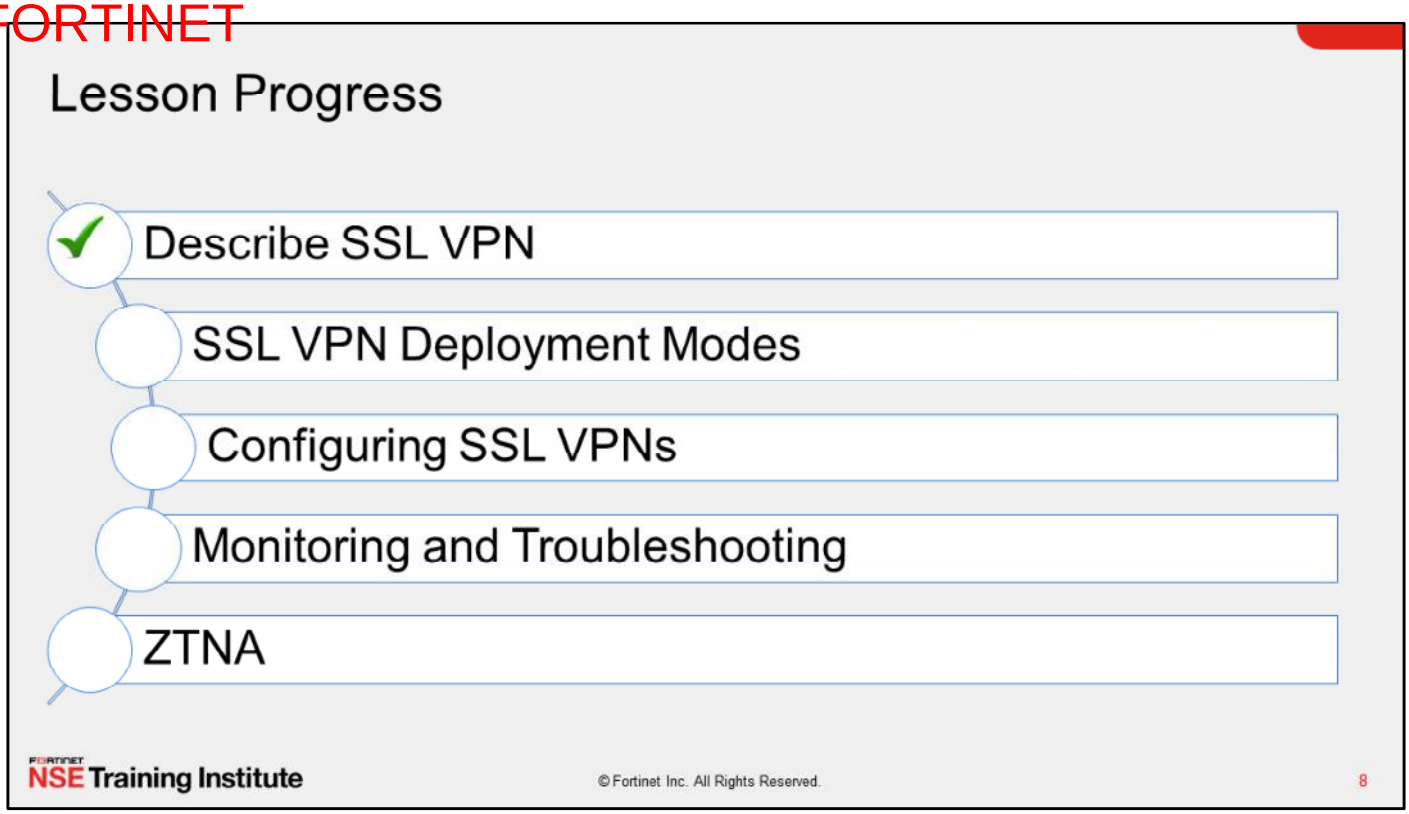

Good job! You now understand basic concepts related to the function of SSL VPNs and how SSL VPN is different from IPsec.

Now, you will learn about the SSL VPN deployment modes supported by FortiGate.

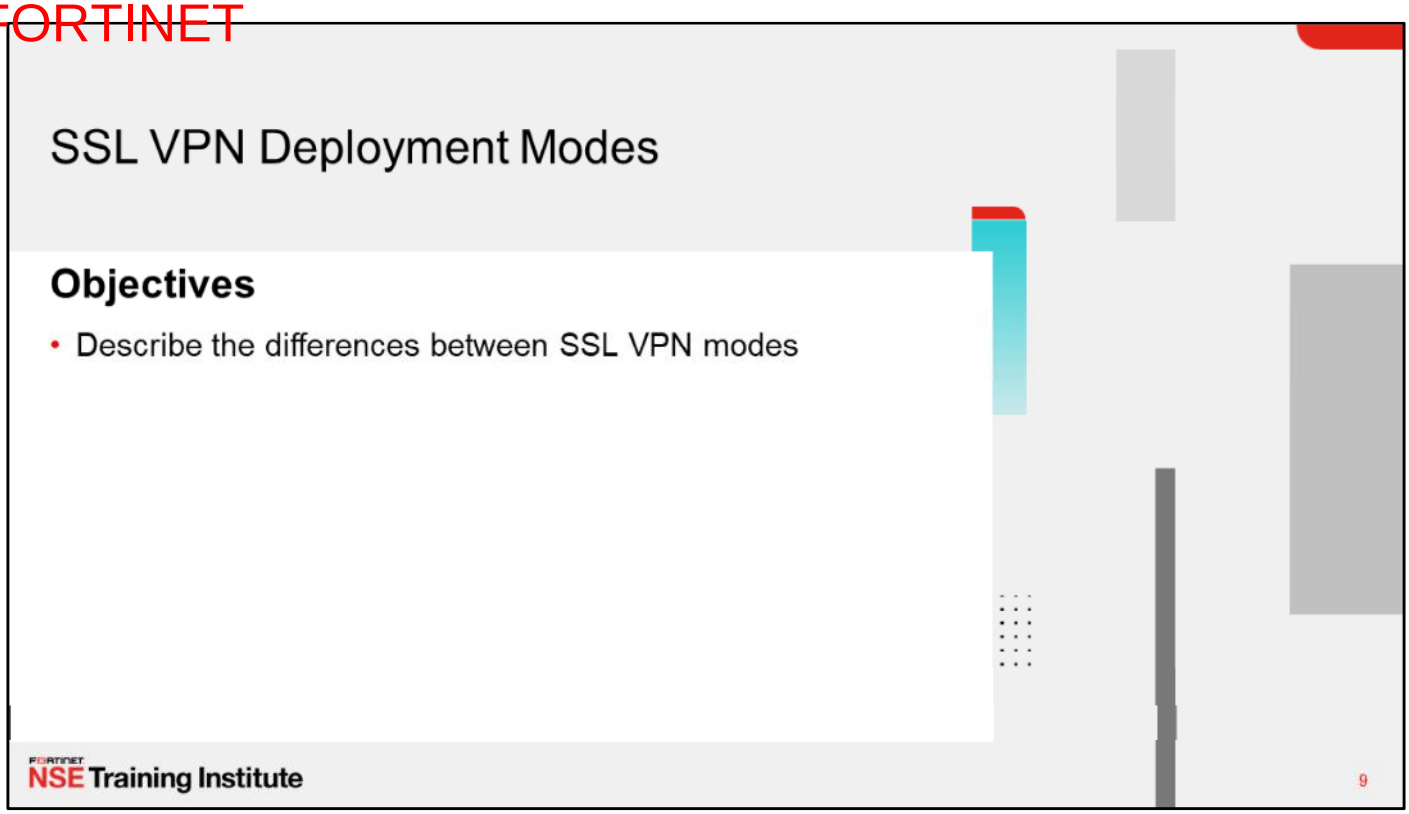

After completing this section, you should be able to achieve the objective shown on this slide. By demonstrating competence in understanding the different ways FortiGate allows SSL VPN connections, you will be able to better design the configuration of your SSL VPN.

SSL VPN Deployment Modes

- Tunnel mode
  - Accessed through a FortiClient
  - Requires a virtual adapter on the client host
- Web mode
  - Requires only a web browser
  - Supports a limited number of protocols:
    - FTP, HTTP/HTTPS, RDP, SMB/CIFS, SSH, Telnet, VNC, and Ping

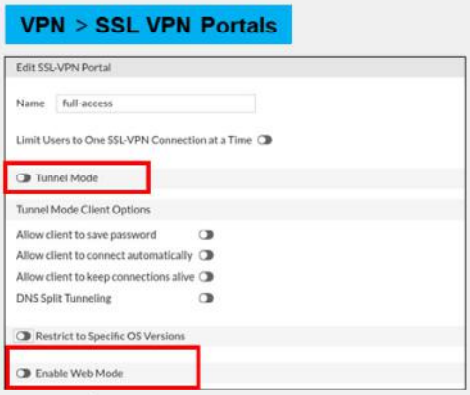


```
config vpn ssl web portal
  edit <portal-name>
   set tunnel-mode [enable|disable]
   set web-mode [enable|disable]
  end
```

NSE Training Institute

10

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

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

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

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

Web Mode

- Connect to the FortiGate SSL VPN portal from any browser
  - The web portal displays the status of SSL VPN
  - The SSL VPN stays up only while the SSL VPN portal page is open
- Access internal network resources easily using:
  - Bookmarks
  - Quick connection
- Disadvantages:
  - Interaction with the internal network exclusively by browser
    - Through the SSL VPN portal
    - External network applications cannot send data across the VPN
  - Limited number of protocols supported

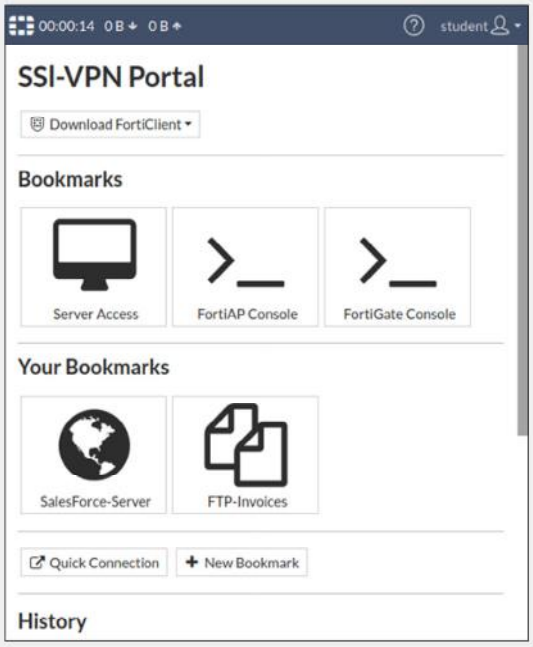

NSE Training Institute

11

Web mode is the simplest SSL VPN mode.

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

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

Web mode has two main disadvantages:

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

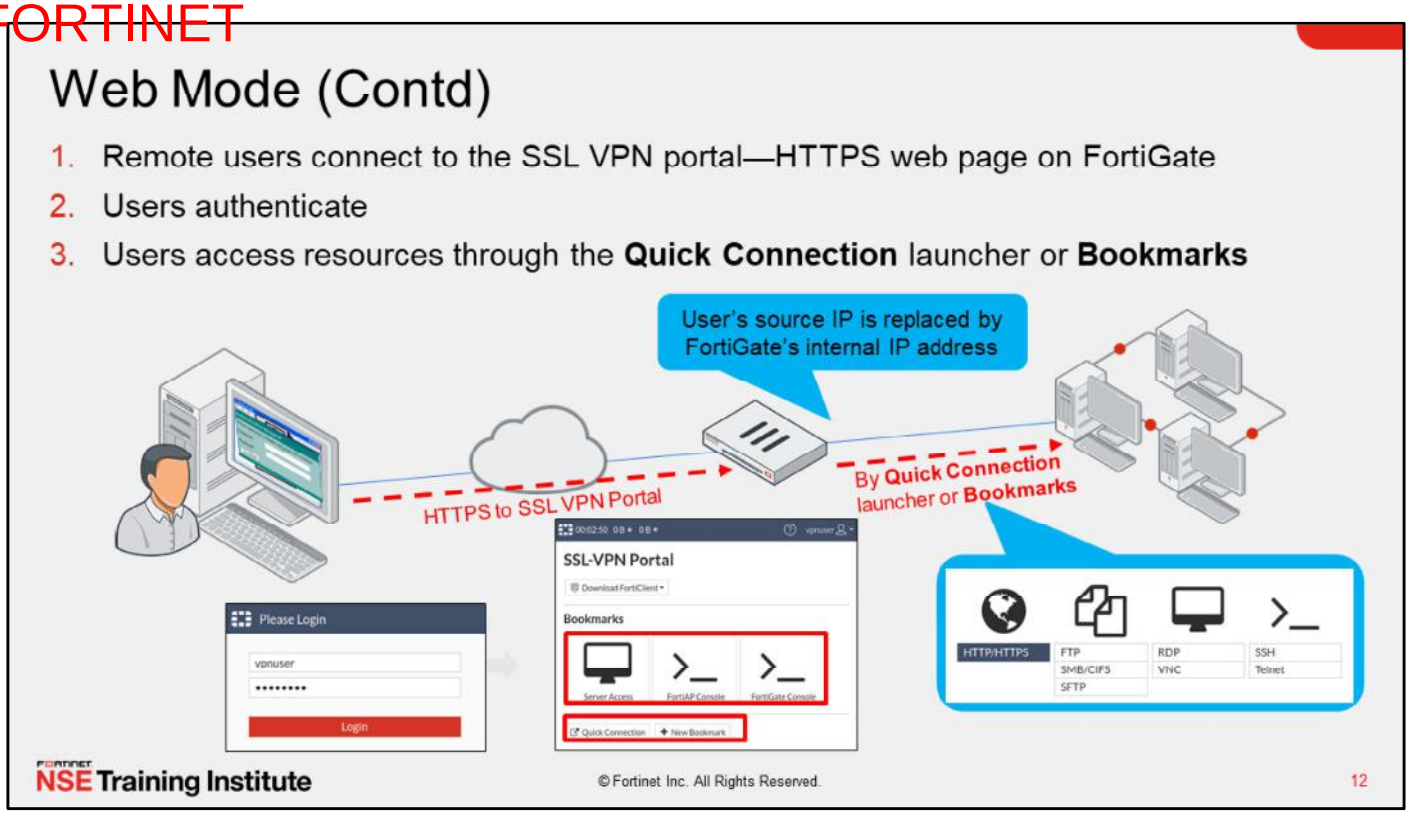

How does web mode work?

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

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

Tunnel Mode

- Connect to FortiGate through FortiClient
  - Tunnel is up only while the SSL VPN client is connected
  - FortiClient adds a virtual network adapter called `fortissl`
- FortiGate establishes the tunnel
  - Assigns a virtual IP address to the client from a pool of reserved addresses
  - All traffic is encapsulated with SSL/ TLS
- Advantage:
  - Any IP network application on the client can send traffic through the tunnel
- Disadvantage:
  - Requires the installation of a VPN client
    http://www.forticlient.com/

FortiClient
Next Generation Endpoint Protection

NSE Training Institute

13

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

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

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

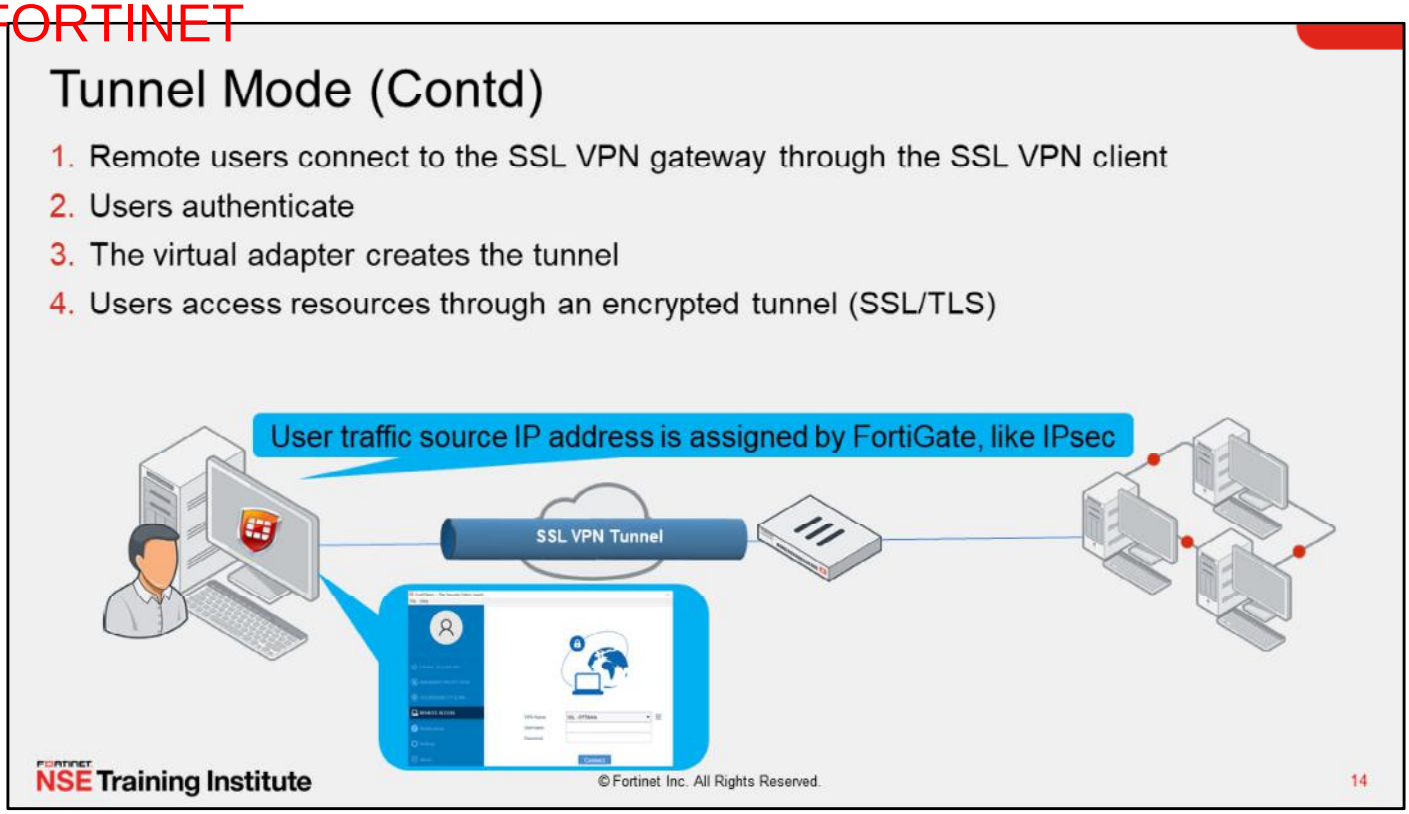

How does tunnel mode work?

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

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

Tunnel Mode—FortiGate as Client

- Connect to server FortiGate device as SSL VPN client
  - Use *SSL VPN Tunnel* interface type
  - Devices connect to client FortiGate device can access the resources behind server FortiGate
- Tunnel establishes between two FortiGate devices
  - Hub-and-spoke topology
  - Client FortiGate dynamically adds route to remote subnets
  - Assigns a virtual IP address to the client FortiGate device from a pool of reserved addresses
- Advantage:
  - Any IP network application on the user machines connect to client FortiGate device can send traffic through the tunnel
  - Useful to avoid issues caused by intermediate devices, such as:
    - ESP packets being blocked.
    - UDP ports 500 or 4500 being blocked.
    - Fragments being dropped, causing IKE negotiation that uses large certificates to fail if the peer does not support IKE fragmentation.
- Disadvantage:
  - Requires proper CA certificate on SSL VPN Server FortiGate
  - SSL VPN Client FortiGate user uses PSK and PKI client certificate to authenticate

NSE Training Institute © Fortinet Inc. All Rights Reserved. 15

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

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

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

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

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

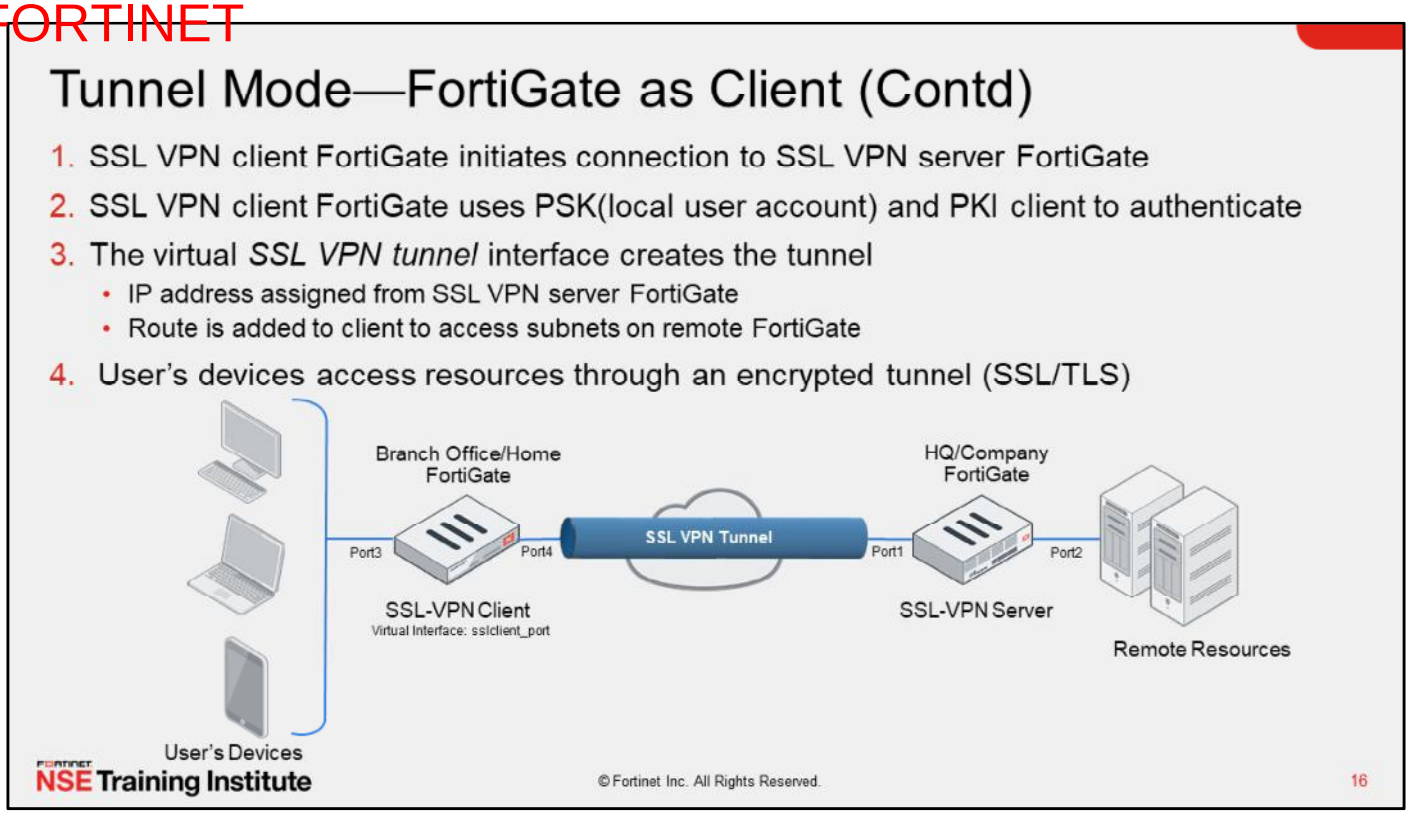

How does tunnel mode work when FortiGate is configured as client?

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

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

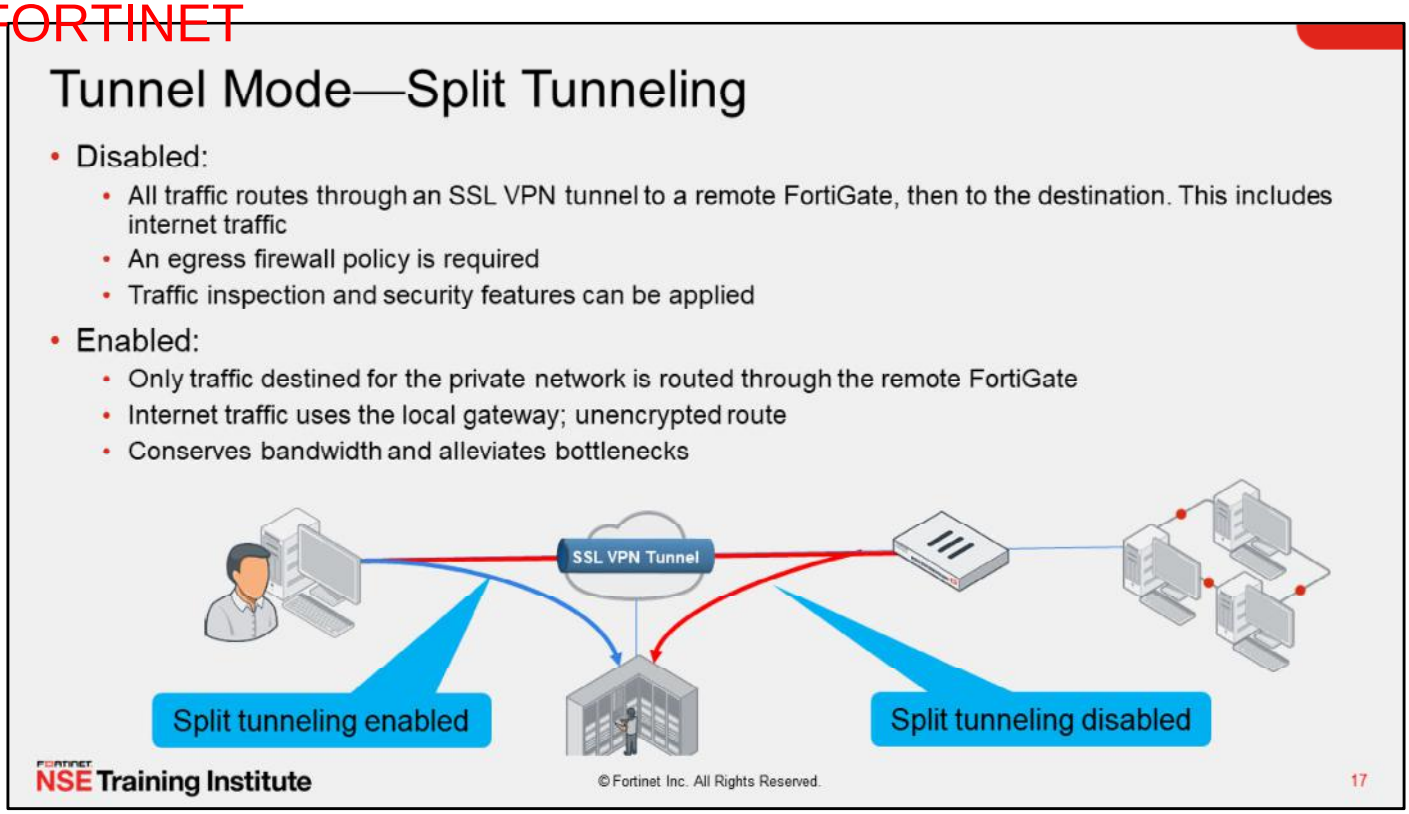

Tunnel mode also supports split tunneling.

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

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

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

When split tunneling is enabled, only traffic that is destined for the private network behind the remote FortiGate is routed through the tunnel. All other traffic is sent through the usual unencrypted route.

Split tunneling helps to conserve bandwidth and alleviates bottlenecks.

Knowledge Check

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

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

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

NSE Training Institute

18

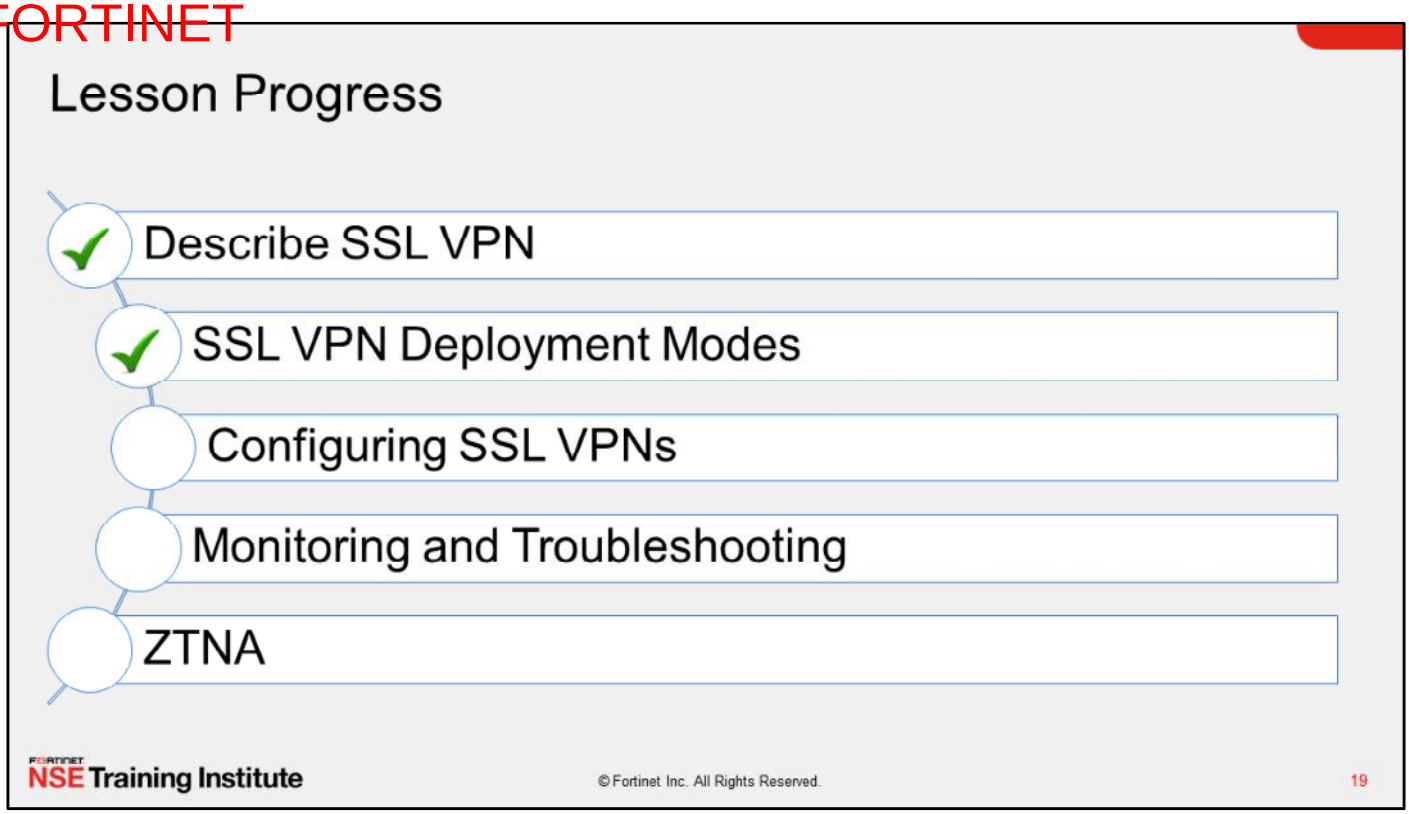

Good job! You now understand the SSL VPN operation modes supported by FortiGate.

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

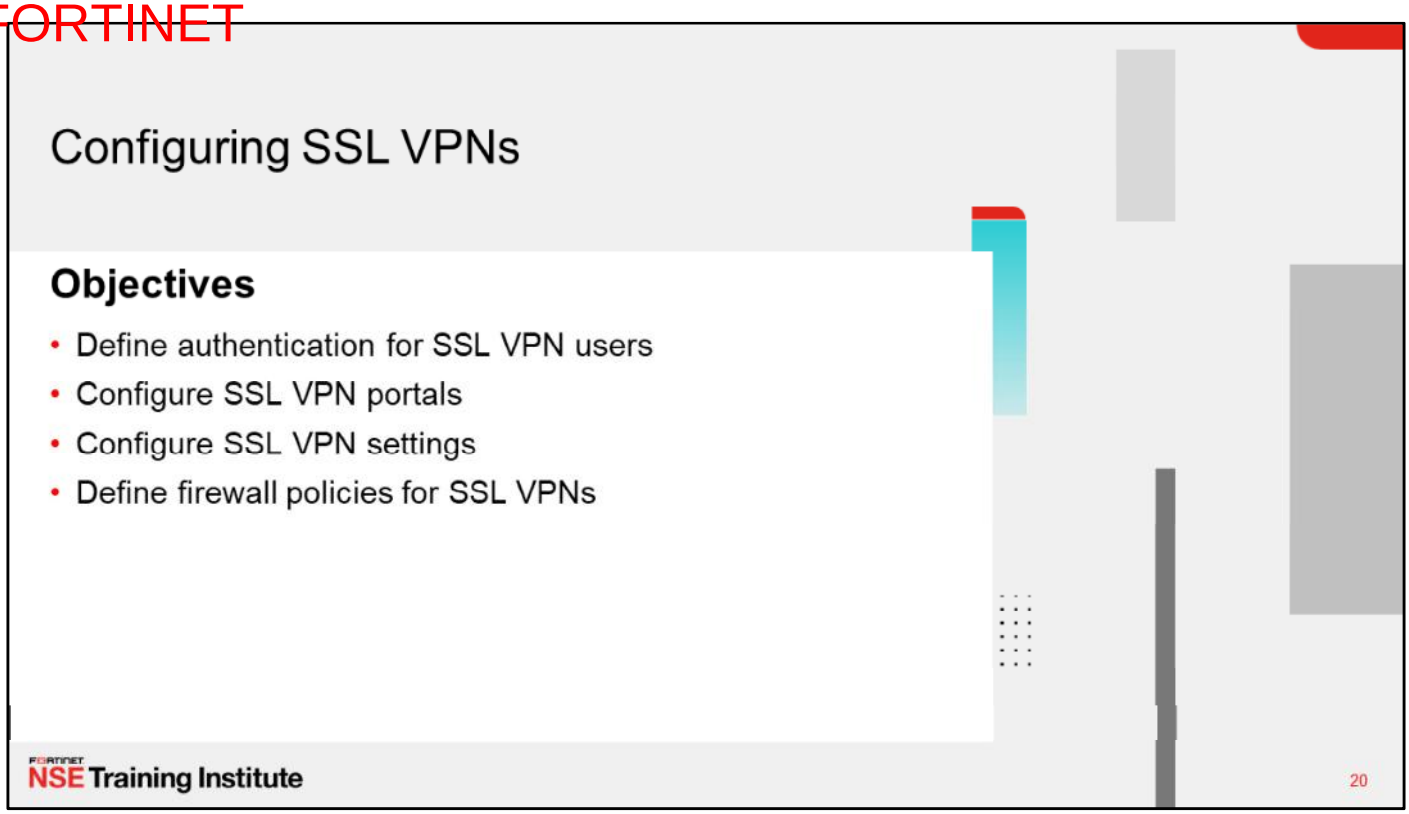

After completing this section, you should be able to achieve the objectives shown on this slide.

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

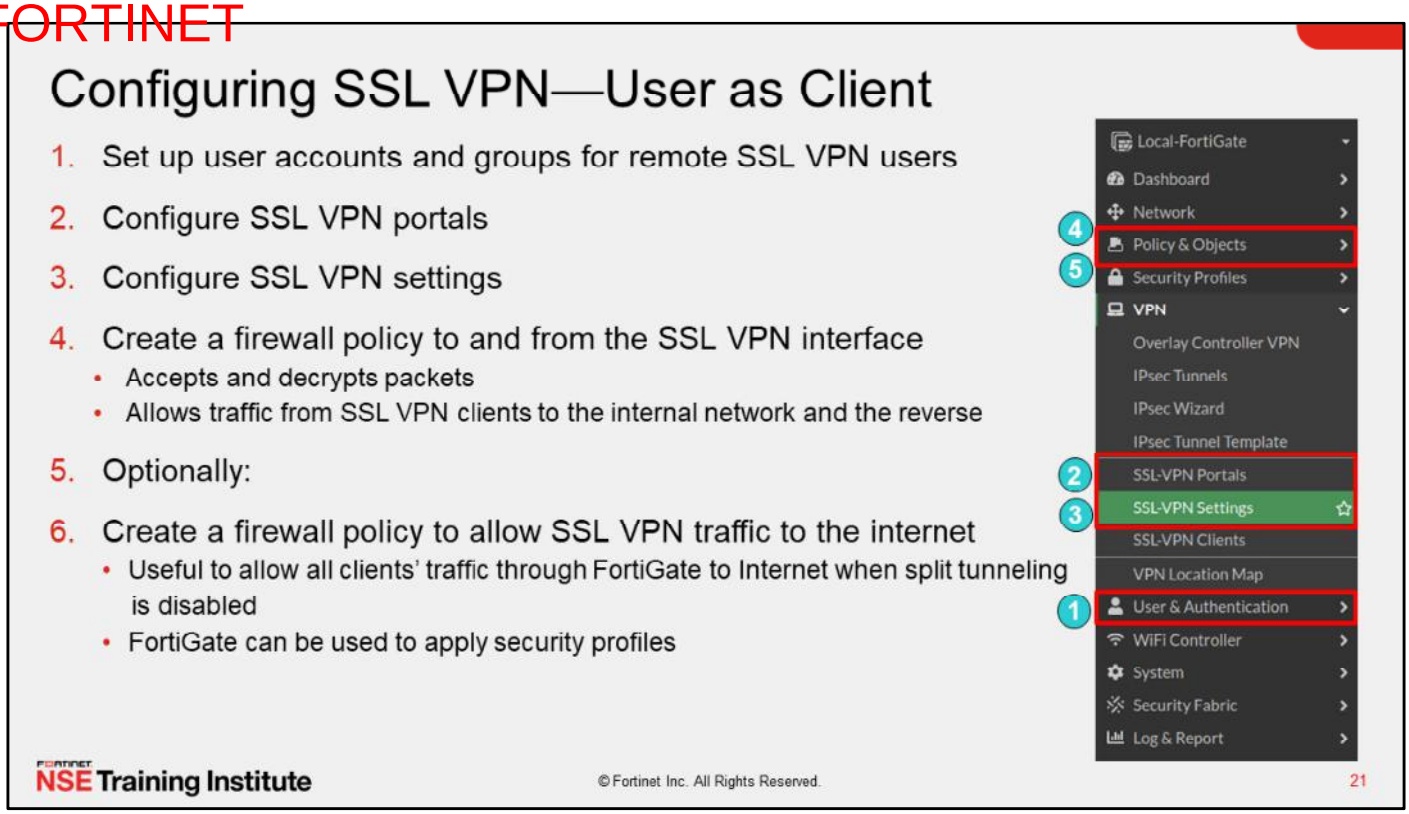

To configure SSL VPN, you must take these steps:

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

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

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

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

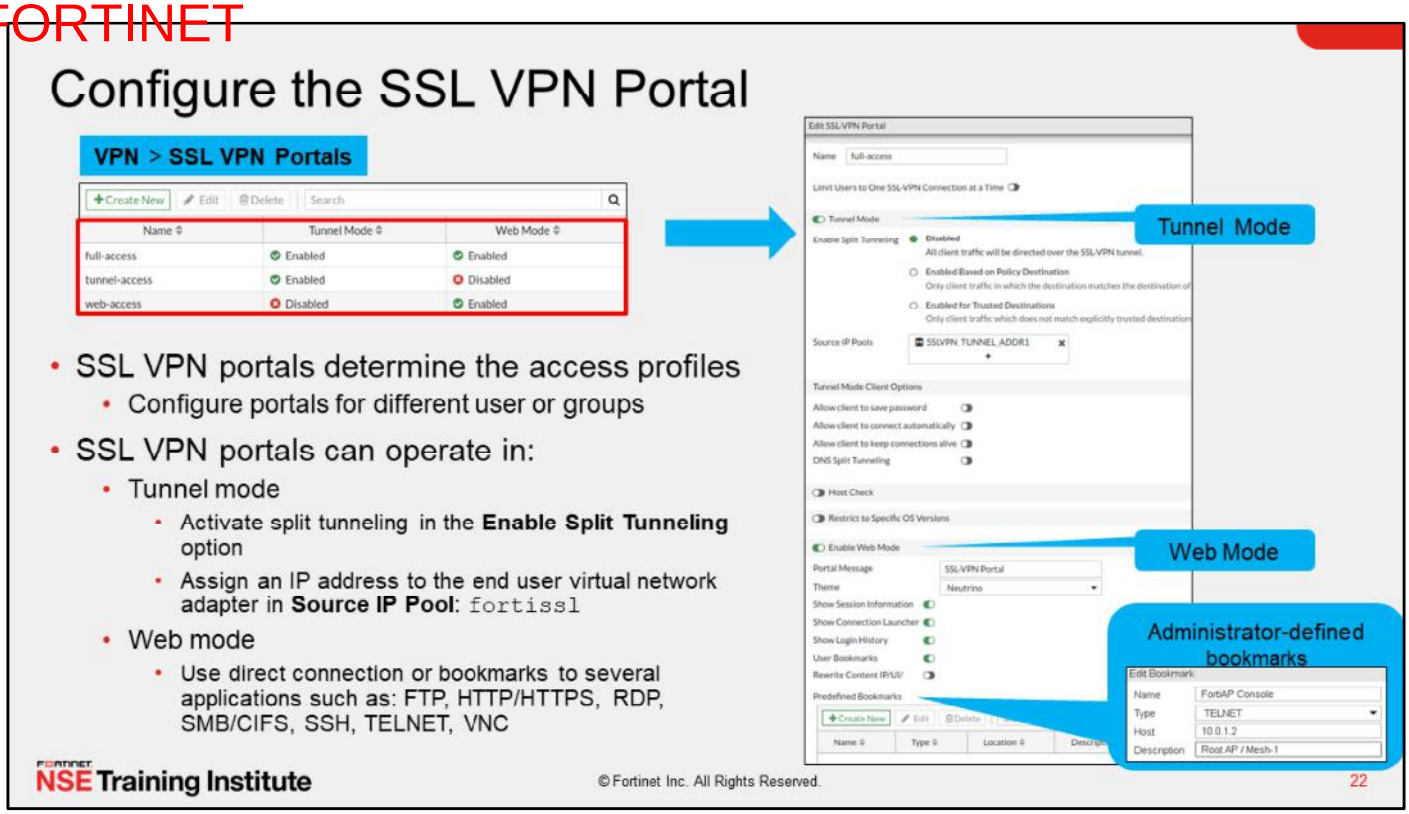

The next step is to configure the SSL VPN portal(s). An SSL VPN portal contains tools and resource links for the users to access.

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

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

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

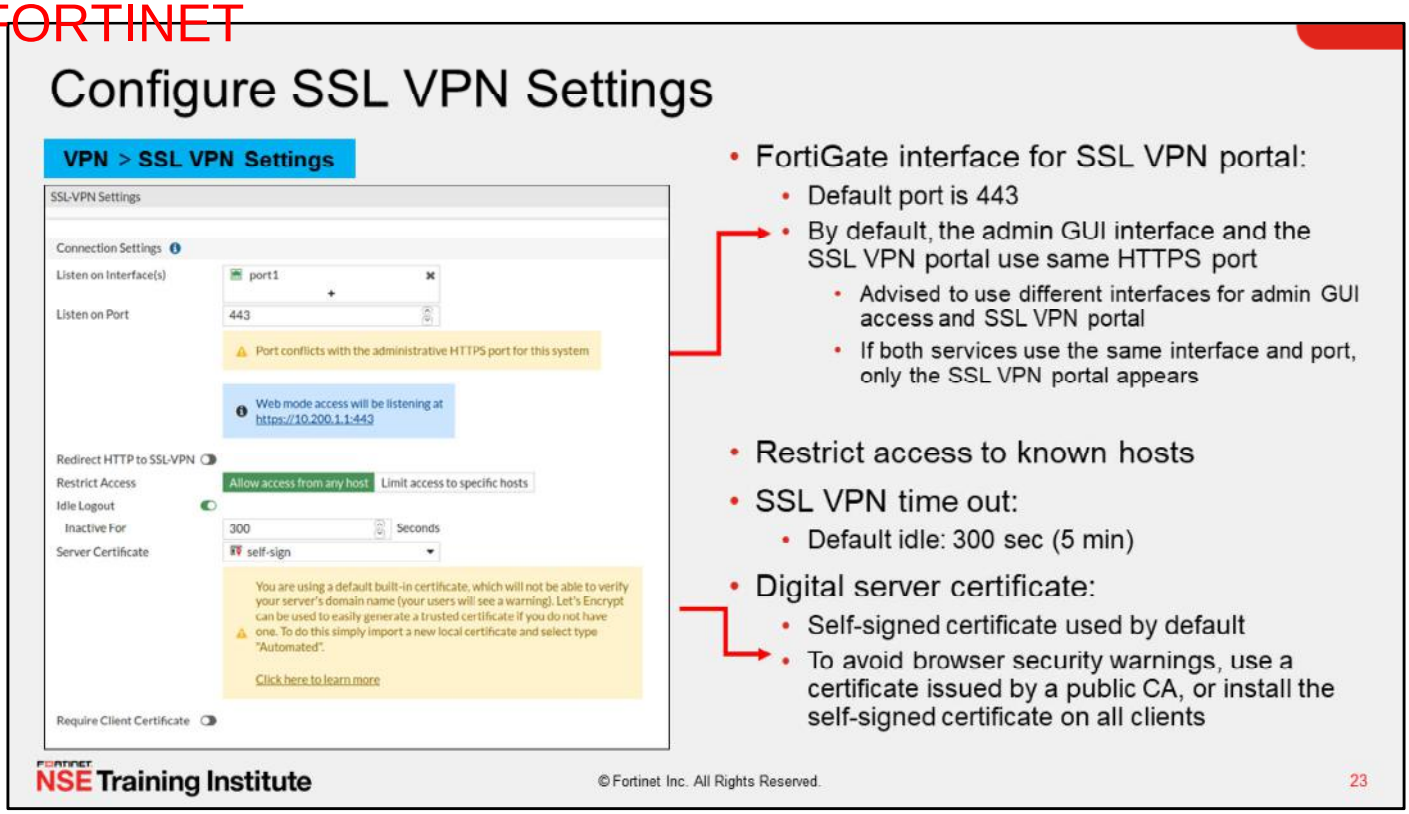

After you configure the SSL VPN portal, the next step is to configure the SSL VPN settings.

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

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

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

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

Configure SSL VPN Settings (Contd)

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

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

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

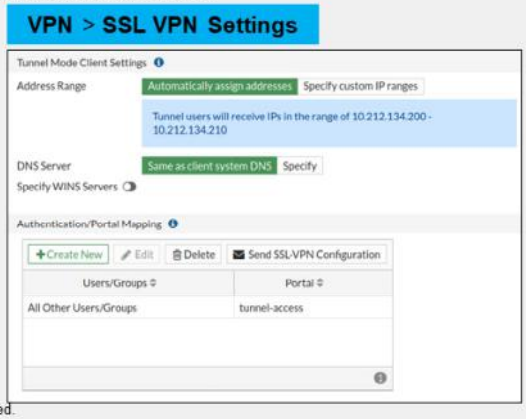

NSE Training Institute

24

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

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

- First-available (default setting)
- Round robin

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

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

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

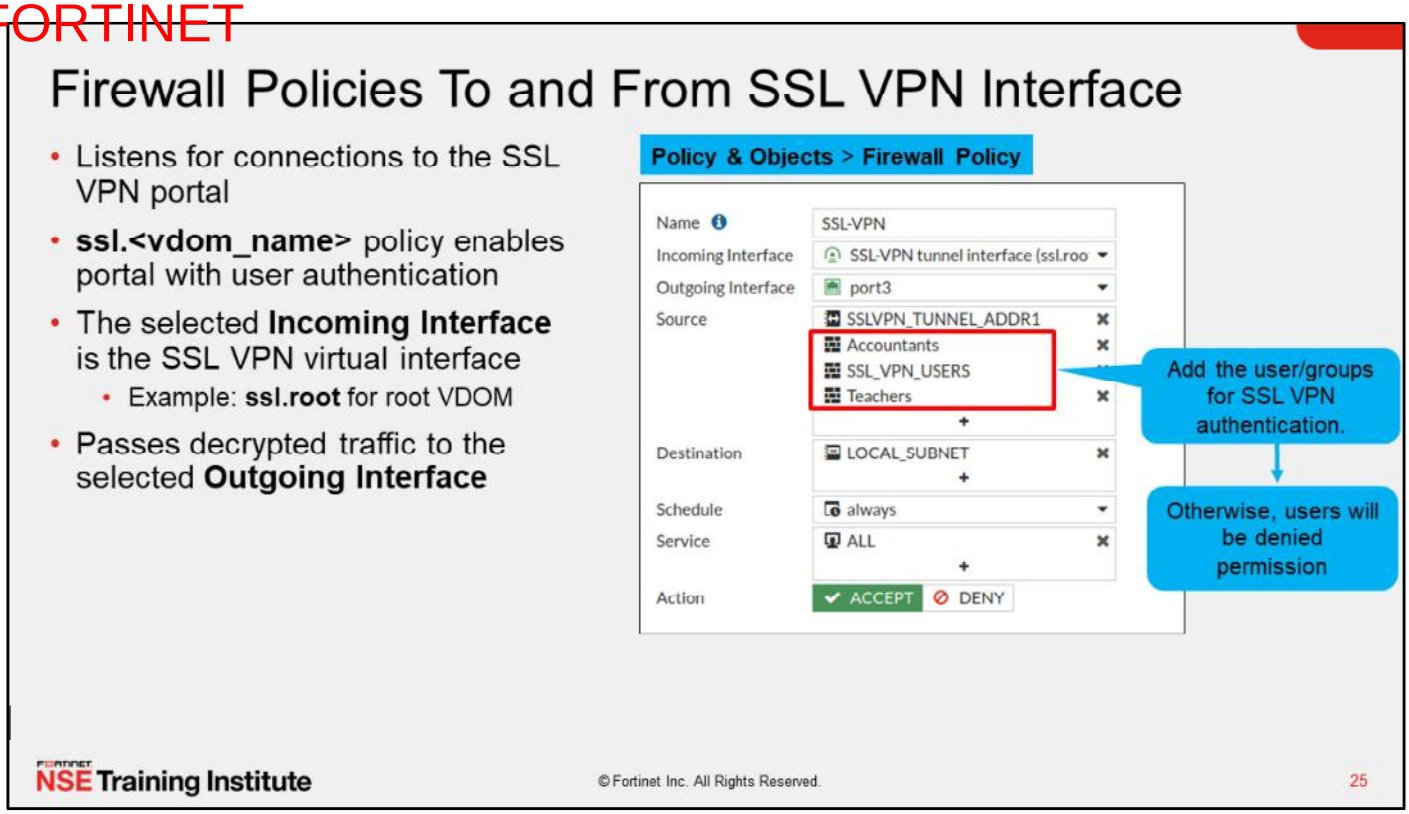

The fourth, and last, mandatory step involves creating firewall policies for logging on.

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

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

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

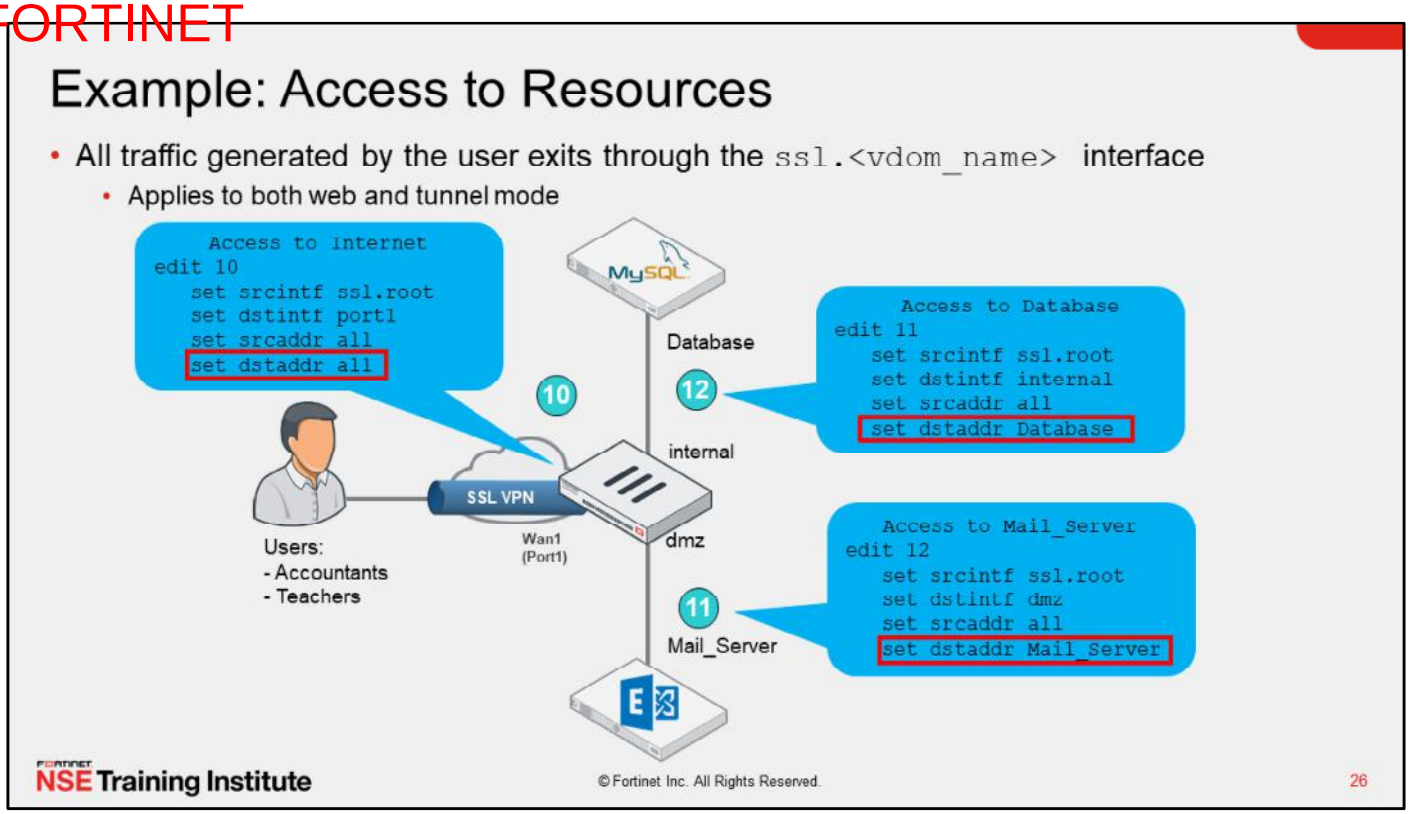

Any traffic from SSL VPN users, whether in web portal or tunnel mode, exits from the `ssl.<vdom_name>` interface.

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

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

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

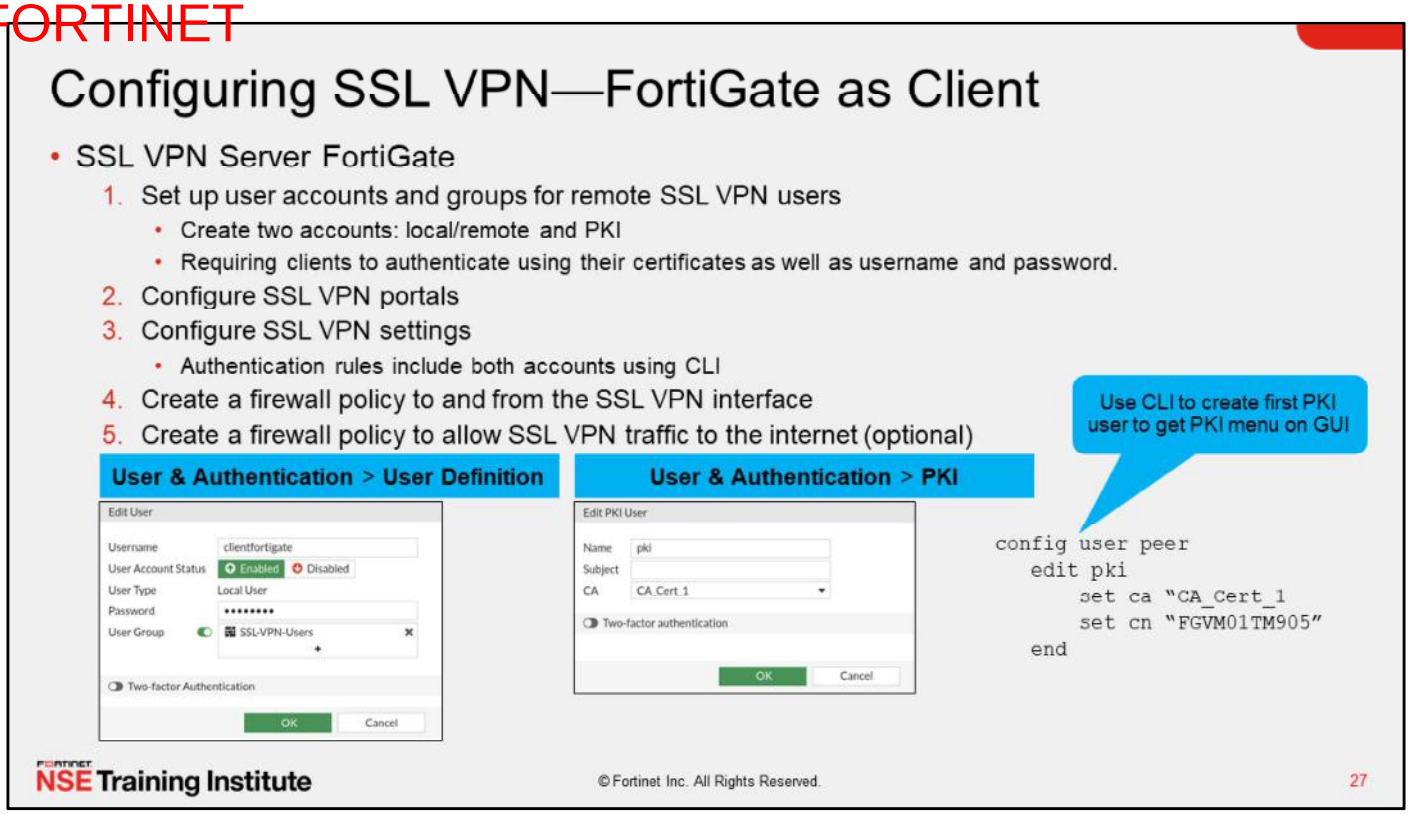

To configure SSL VPN, you must take these steps:
SSL VPN Server FortiGate:

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

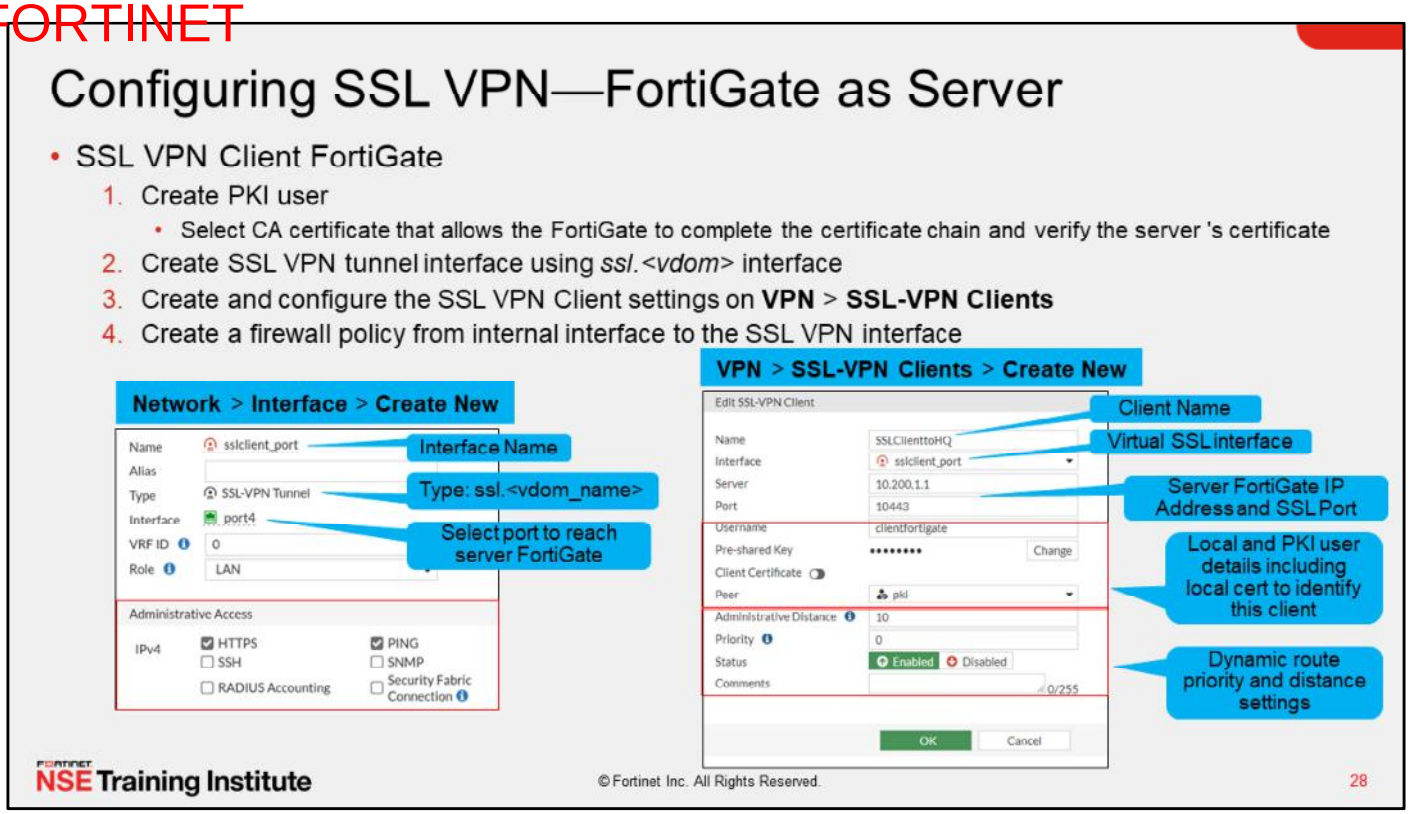

To configure SSL VPN, you must take these steps:
SSL VPN Client FortiGate:

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

Knowledge Check

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

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

NSE Training Institute

29

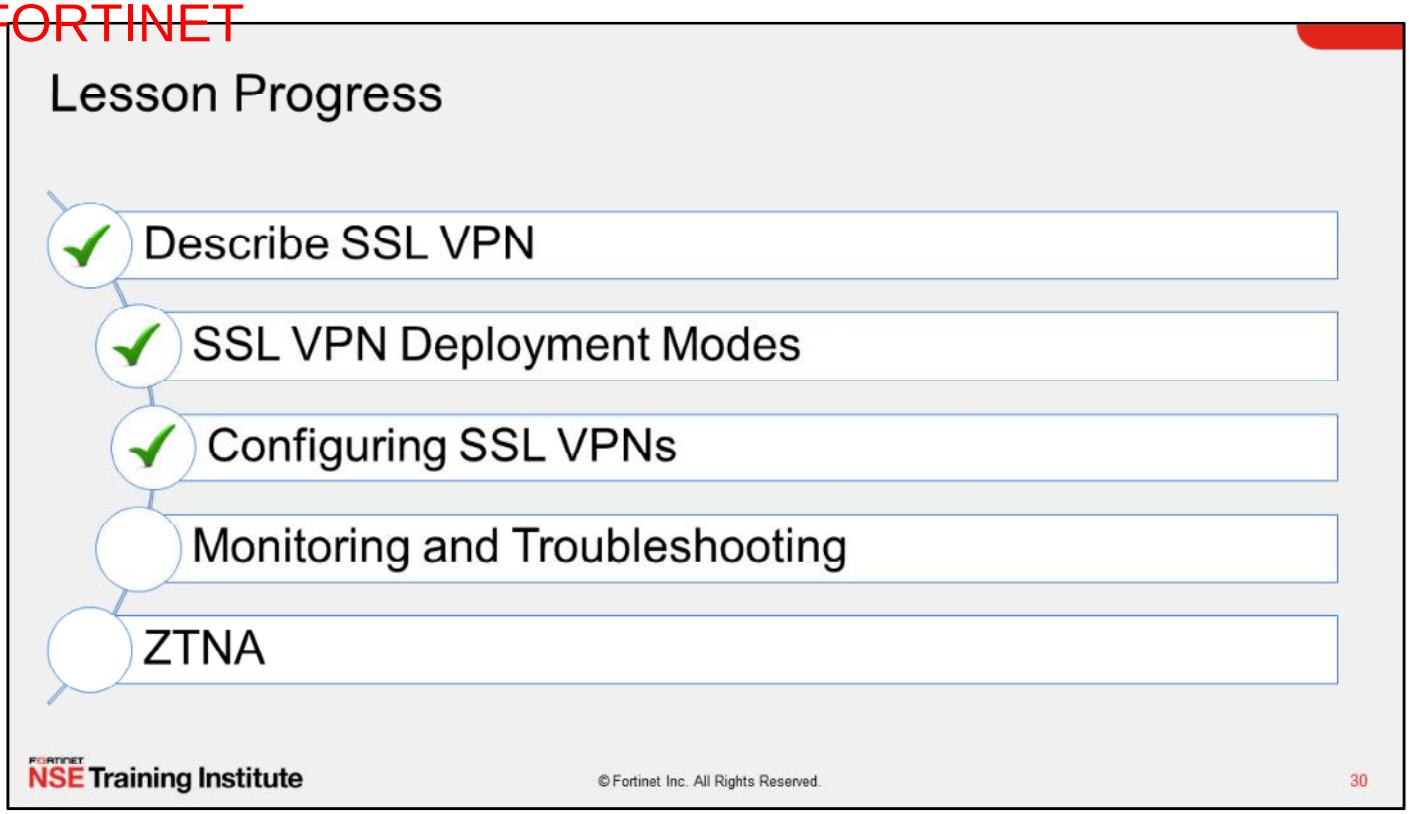

Good job! You now understand how to configure the FortiGate for SSL VPN connections.

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

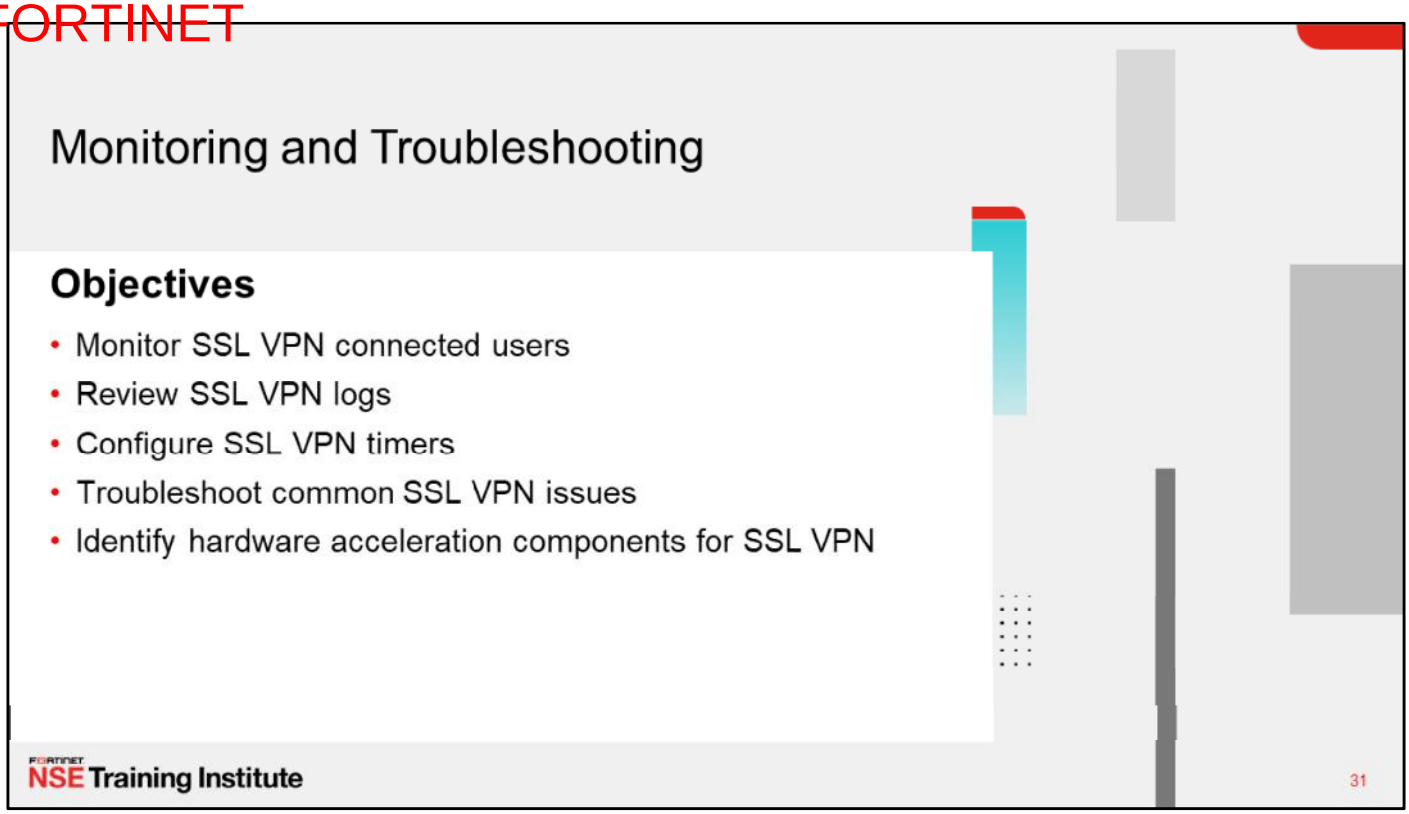

After completing this section, you should be able to achieve the objectives shown on this slide.

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

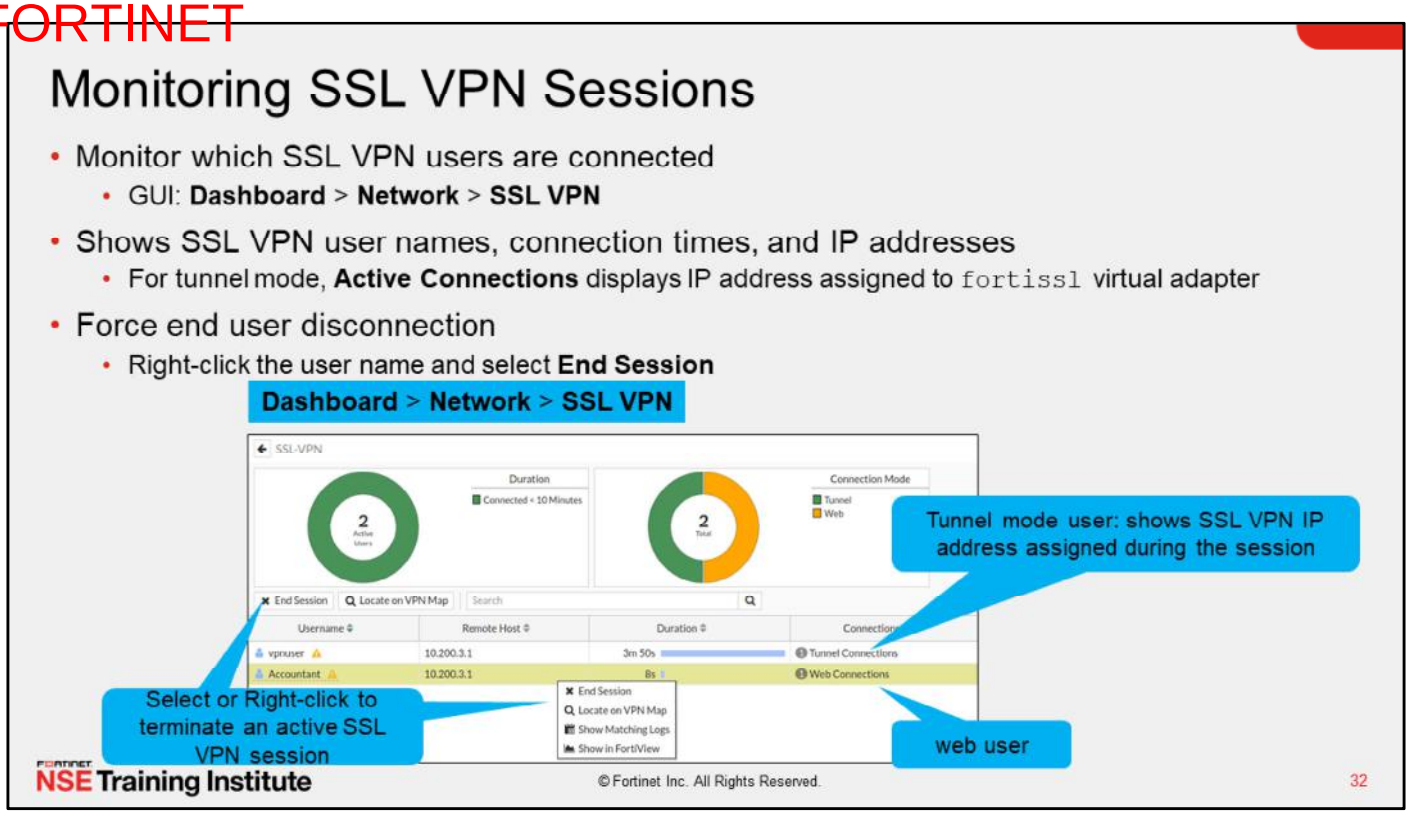

You can monitor which SSL VPN users are connected on the **SSL VPN** widget. This shows the names of all SSL VPN users that are currently connected to FortiGate, their IP addresses (both inside the tunnel and outside), and connection times.

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

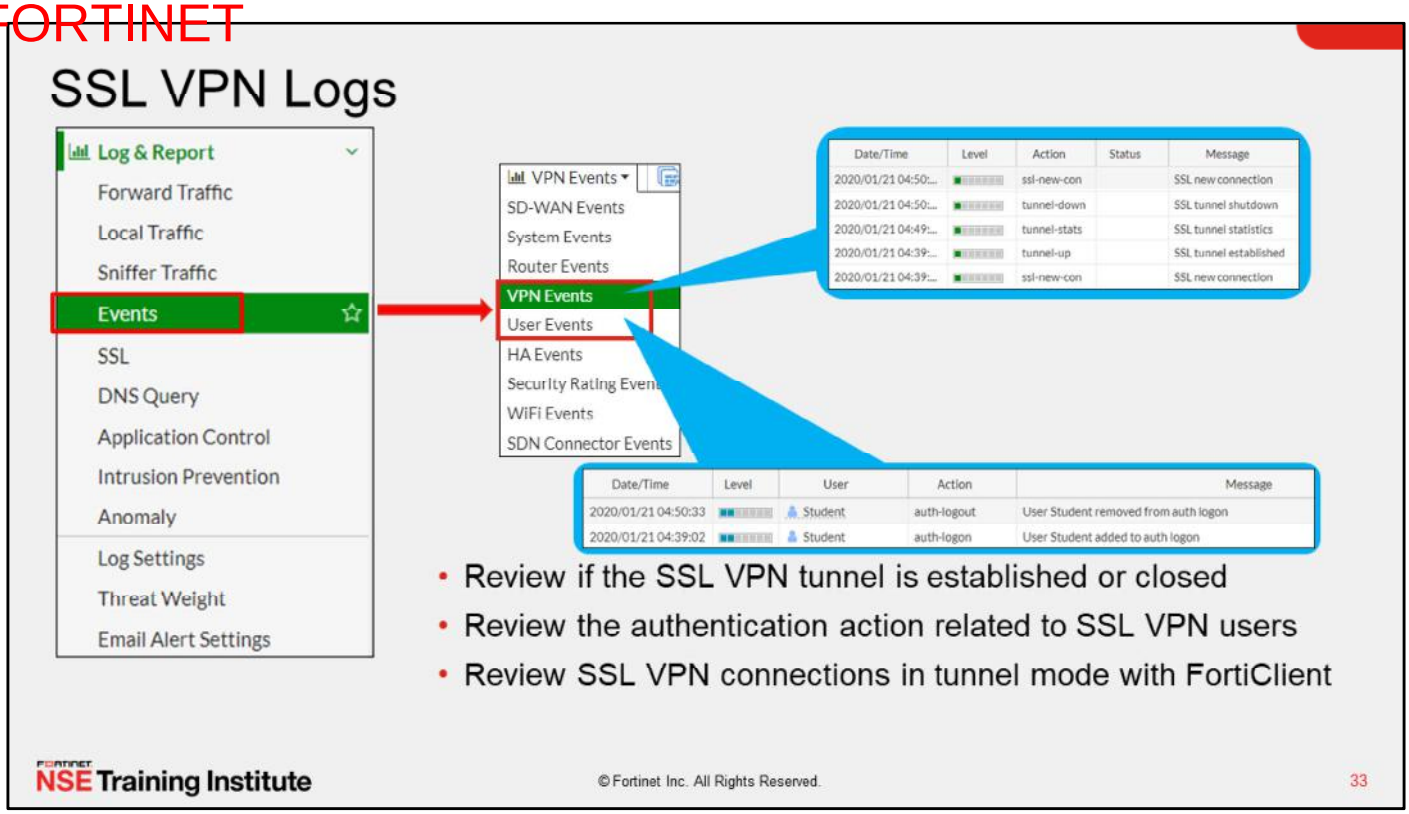

You can also review SSL VPN logs. On **Log & Report** > **Events**:

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

SSL VPN Idle Timeout vs. Authentication Session

- Firewall policy authentication session is associated with SSL VPN tunnel session
  - Firewall policy authentication session is forced to end when SSL VPN tunnel session ends
  - Prevents reuse of authenticated SSL VPN firewall sessions (not yet expired) by a different user, after the initial user terminates the SSL VPN tunnel session
- SSL VPN authentication is not subject to the firewall authentication timeout setting
  - It has a separate idle setting: default 300 seconds

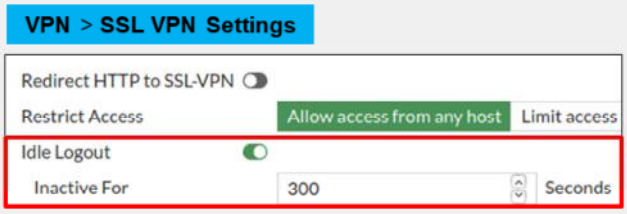


```
config vpn ssl settings
    set idle-timeout <0-259200>
end
```

NSE Training Institute

34

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

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

SSL VPN Timers

- Set up timers to avoid logouts when SSL VPN users are connected over high latency connections
  - DTLS hello timeout—default 10 seconds
  - Login timeout—default 30 seconds

```
config vpn ssl settings
   set login-timeout <10-180>
   set dtls-hello-timeout <10-60>
   set http-request-header-timeout <1-60>
   Set http-request-body-timeout <1-60>
end
```

- Timers can also help to mitigate DoS attacks within SSL VPN caused by partial HTTP requests, such as Slowloris and R-U-Dead-Yet

NSE Training Institute © Fortinet Inc. All Rights Reserved. 35

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

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

Best Practices for Common SSL VPN Issues

- For web mode connections, make sure that:
  - Cookies are enabled and the internet privacy options are set to high in your web browser
  - SSL VPN clients are following the proper URL structure: https://<FortiGateIP>:<port>
- For tunnel mode connections, make sure that:
  - The FortiClient version is compatible with the FortiOS firmware
    - Refer to release notes for product compatibility and integration
  - Split tunneling is enabled to allow internet access without backhauling all user's data to the remote network, or
  - Split tunneling is disabled and an egress firewall policy is created for SSL VPN connections
- For general SSL VPN connections, make sure that:
  - Users are connecting to the correct port number
    - To check SSL VPN port assignment, click **VPN** > **SSL VPN Settings**
  - Firewall policies include SSL VPN groups or users, and the destination address
  - The timeout timer is configured to flush inactive sessions after a short time
  - Users are encouraged to log out if they are not using the network resources only accessible by SSL VPN

NSE Training Institute © Fortinet Inc. All Rights Reserved. 36

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

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

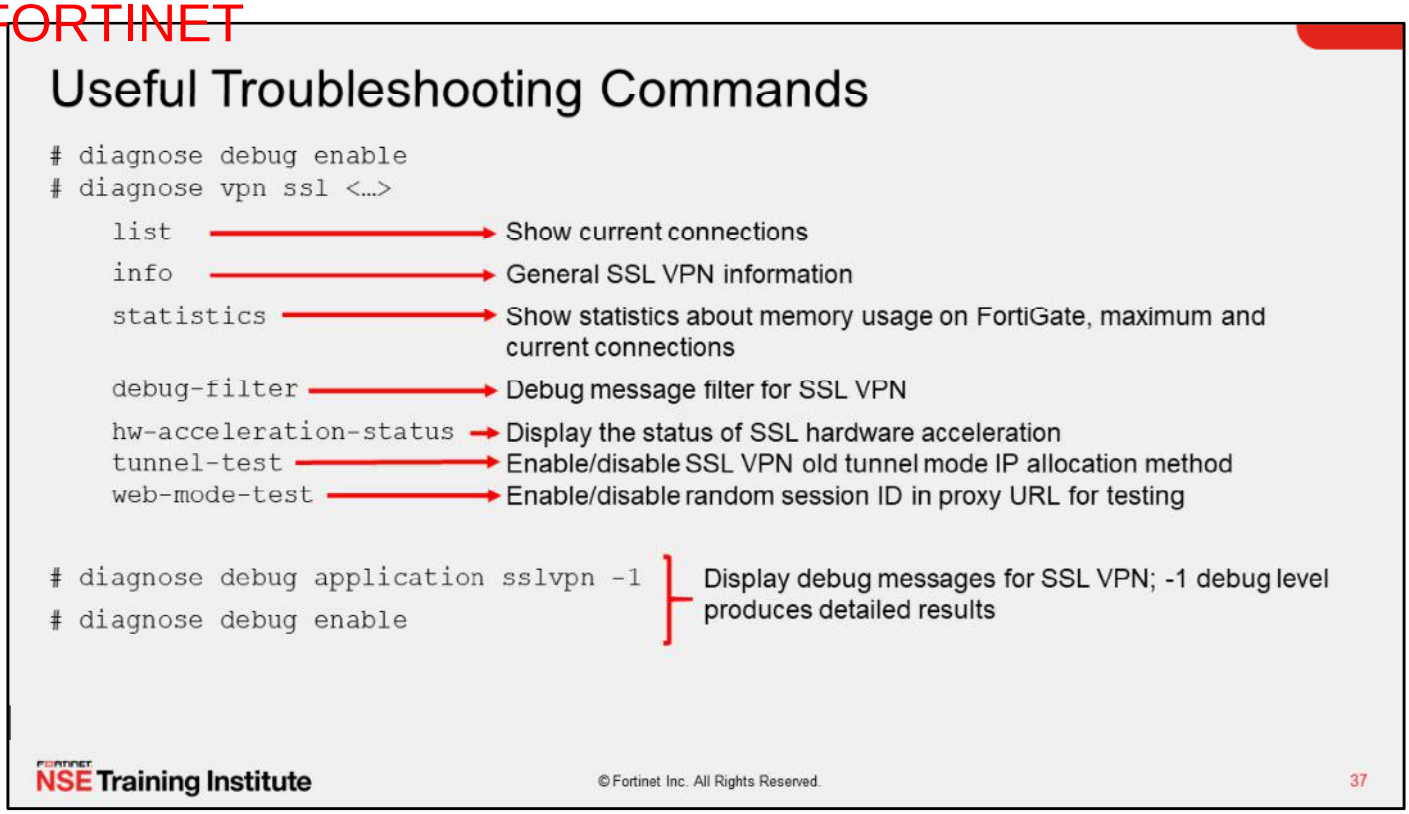

There are several useful troubleshooting commands available under `diagnose vpn ssl`. They include:

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

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

Remember, to use the commands listed above, you must first run the `diagnose debug enable` command.

Hardware Acceleration for SSL VPN

- FortiGate devices with content processors (CP8 or CP9), which offload specific CPU-intensive operations, support high-performance SSL VPN bulk data engines
  - SSL/TLS protocol processor
- Administrators can disable CP offloading through firewall policies
  - For example: test purposes

```
config firewall policy
   edit 1
   set auto-asic-offload [enable |disable]
  end
```

- To view the status of SSL VPN acceleration, use the following command:

```
get vpn status ssl hw-acceleration-status
```

```
Acceleration hardware detected: kxp=on
cipher=on
```

```
No acceleration hardware detected
```

NSE Training Institute

38

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

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

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

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

Knowledge Check

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

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

NSE Training Institute

39

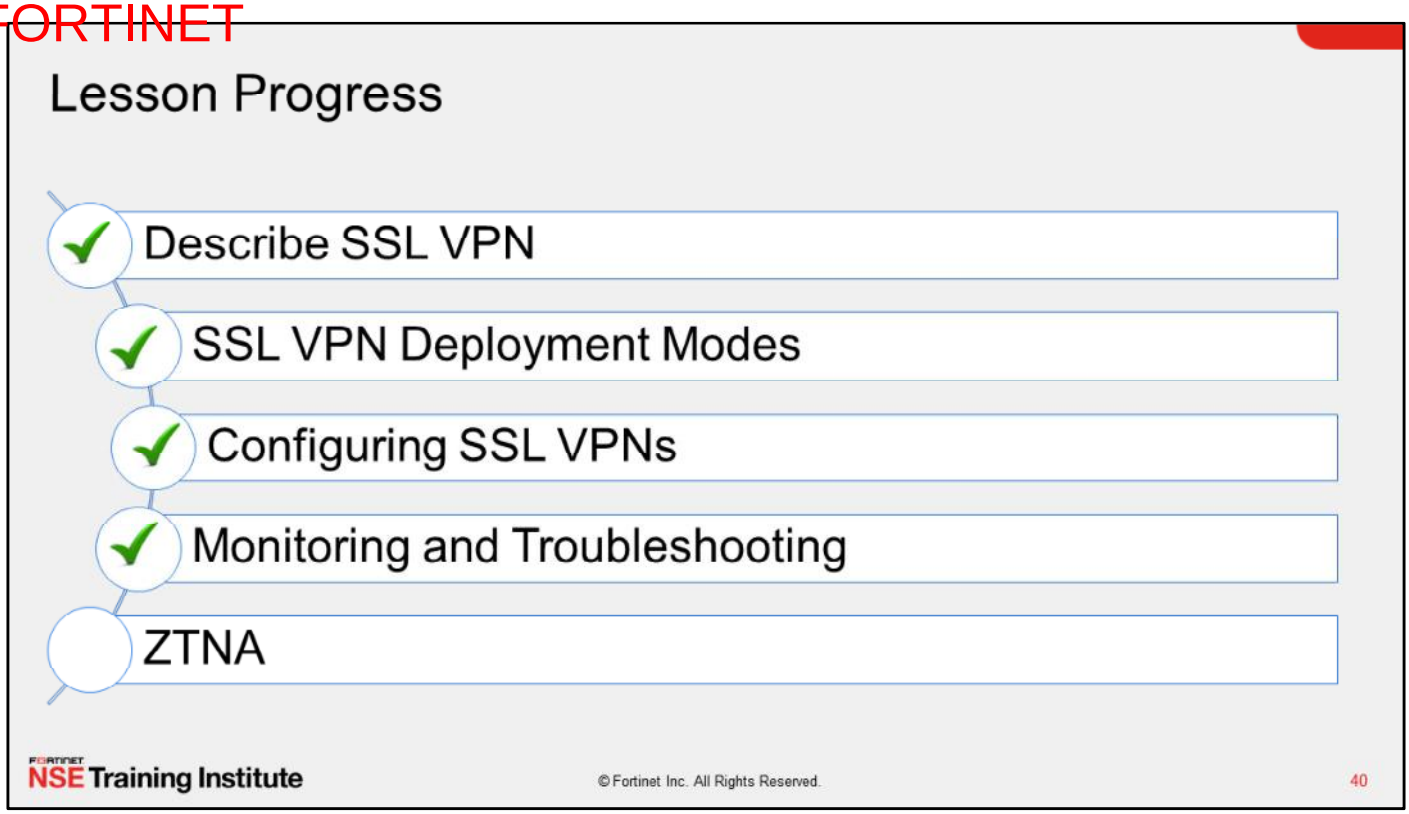

Good job! You now understand how to monitor and troubleshoot SSL VPN.

Now, you will learn the benefits and basic configuration of ZTNA.

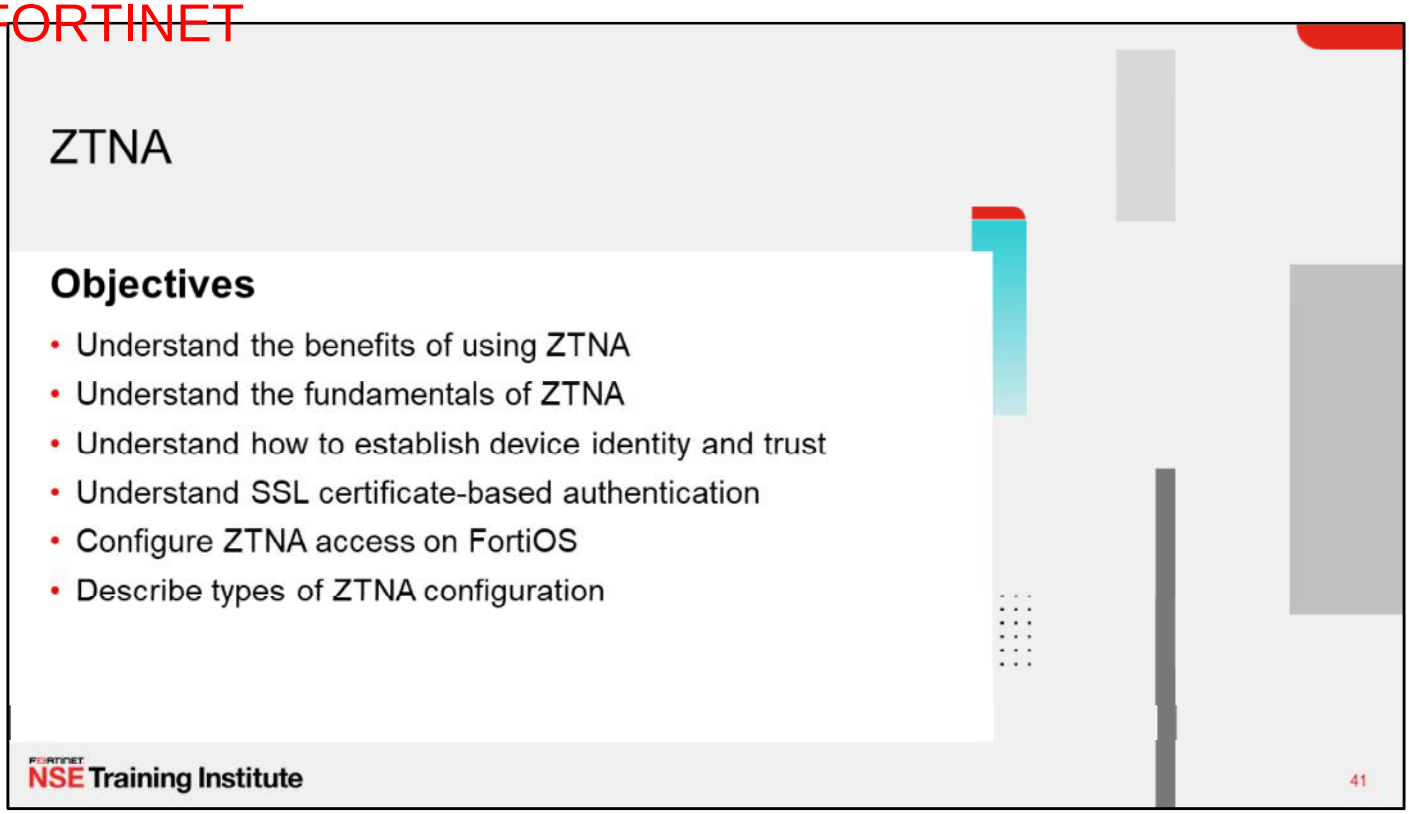

After completing this section, you should be able to achieve the objectives shown on this slide.

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

What is ZTNA?

- Access control method that provides role-based application access
- ZTNA method uses:
  - Client device identification
  - Authentication
  - Zero-trust tags
- Provides flexibility to manage both on-net and off-net users
- ZTNA has two modes:
  - **Full ZTNA (VPN alternative for remote access)**
  - IP/MAC filtering (on-fabric, devices for IT compliances and rules enforcement)

NSE Training Institute

42

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

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

ZTNA has two modes:

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

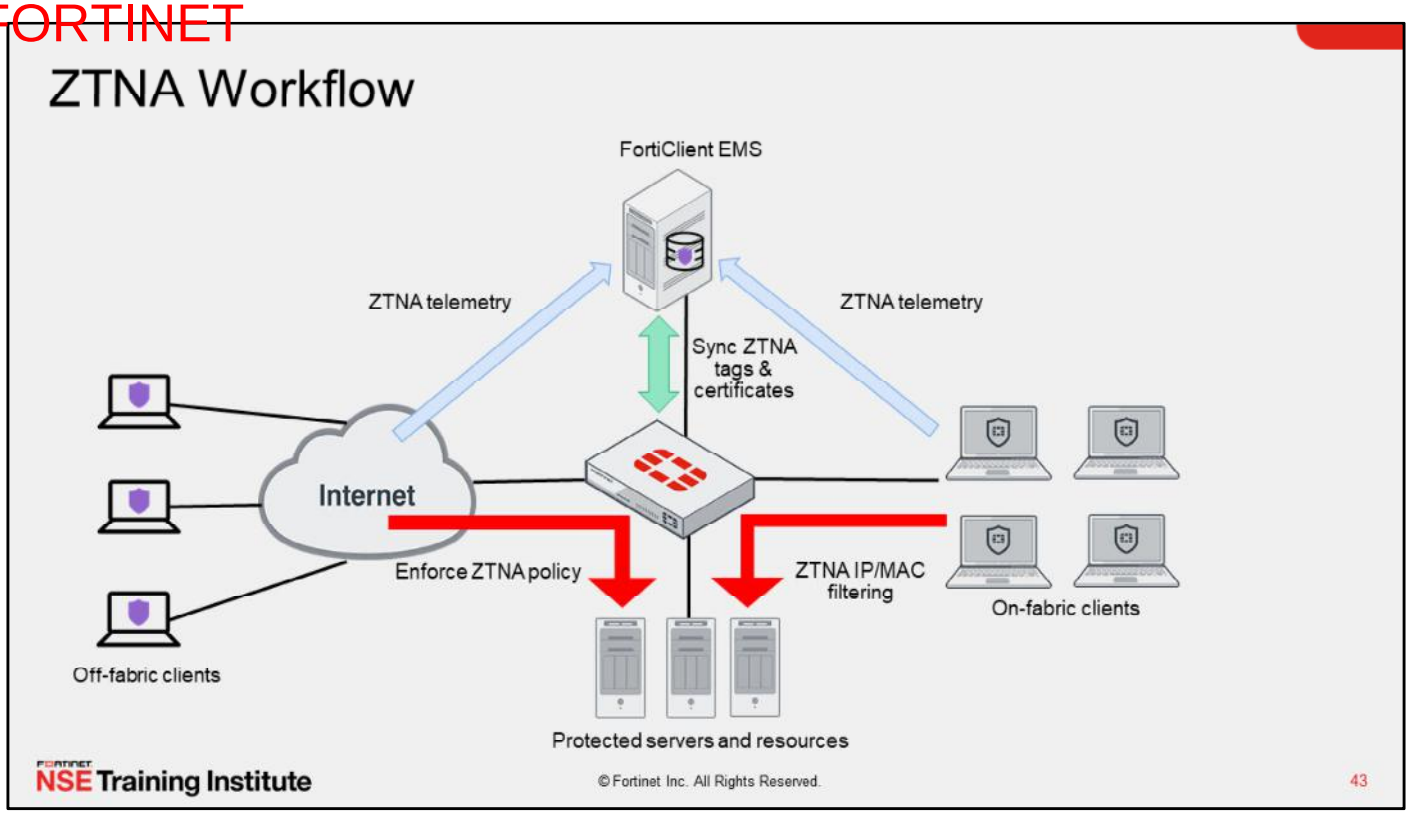

This slide demonstrates ZTNA telemetry, tags, and policy enforcement. You configure ZTNA tag conditions and policies on FortiClient EMS. FortiClient EMS shares the tag information with FortiGate through Security Fabric integration. FortiClient communicates directly with FortiClient EMS to continuously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

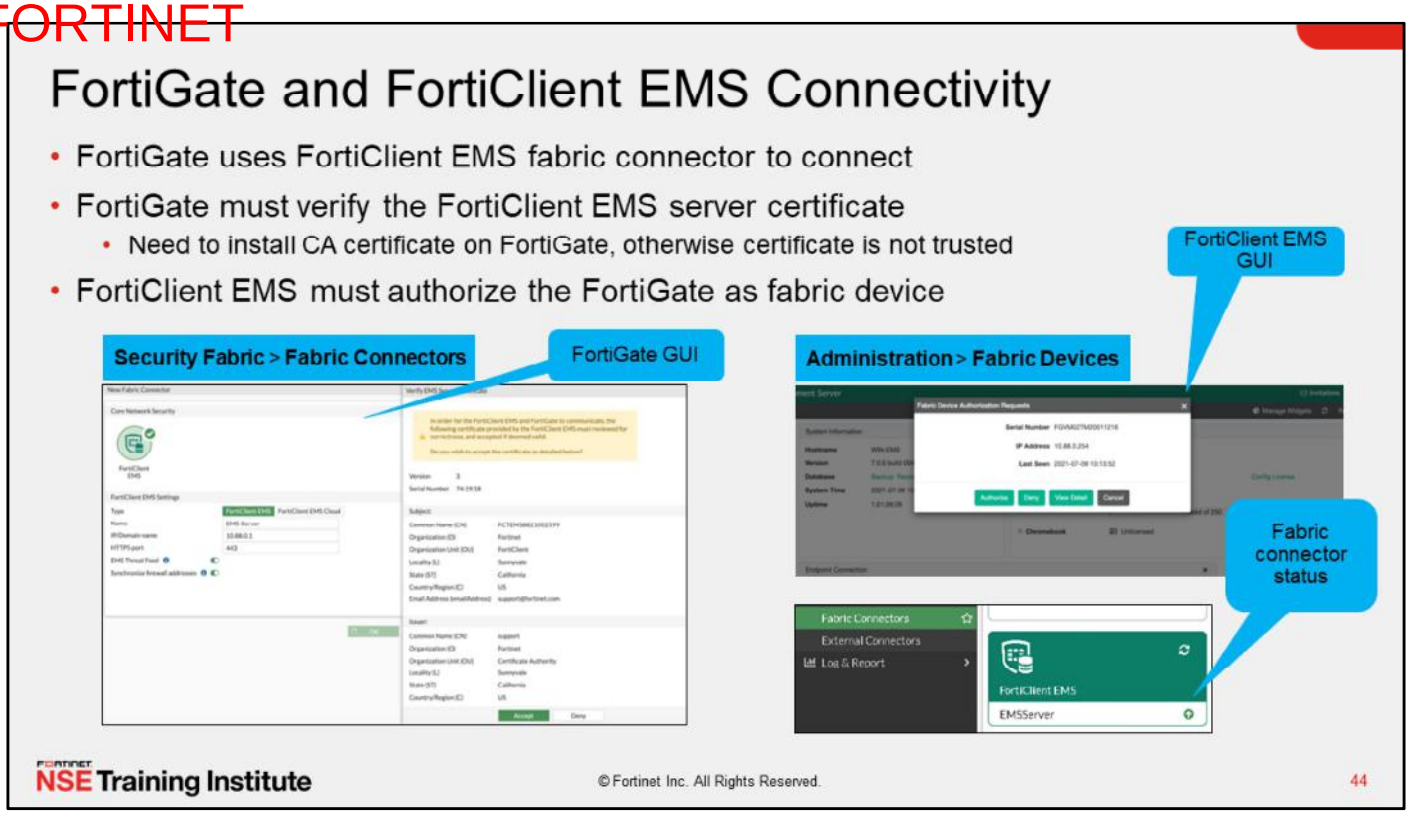

You can configure the on-premises FortiClient EMS connector on FortiGate by clicking **Security Fabric** > **Fabric Connectors**. After applying the FortiClient EMS settings, FortiGate must accept the FortiClient EMS server certificate. However, when you configure a new connection to FortiClient EMS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

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

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

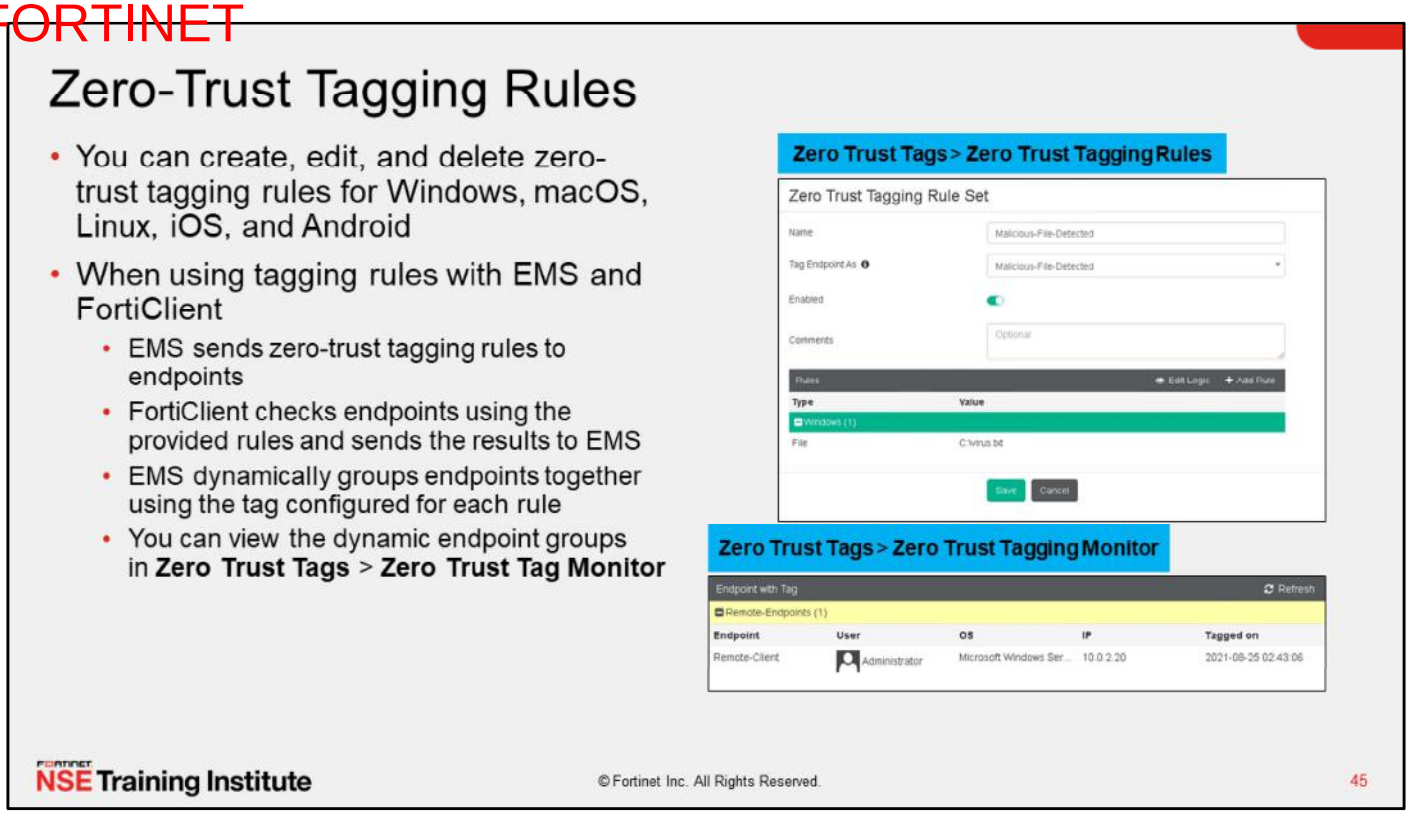

You can create, edit, and delete zero trust tagging rules for Windows, macOS, Linux, iOS, and Android endpoints. The following happens when using zero trust tagging rules with FortiClient EMS and FortiClient:

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

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

Device Roles

- Device identity and trust are integral to ZTNA
- Identity is established through client certificates
- Trust is established between:
  - FortiClient
    - Provides endpoint information (device information, logged on users, and security posture)
    - Obtains client certificate from FortiClient EMS
  - FortiClient EMS
    - Issues and signs the client certificate
    - Synchronizes certificate to FortiGate
    - Uses tagging rules to tag endpoints
  - FortiGate
    - Maintains continuous connection to FortiClient EMS to synchronize endpoint information
    - When device information changes, FortiClient EMS updates FortiGate
    - FortiGate WAD daemon uses this information when processing ZTNA traffic

NSE Training Institute © Fortinet Inc. All Rights Reserved. 46

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

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

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

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

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

FortiClient EMS Certificate Management

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

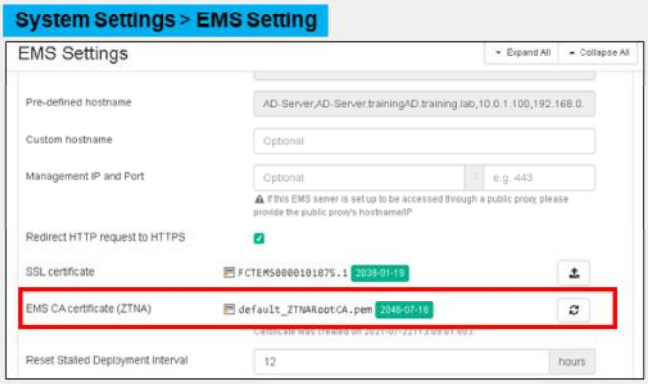

NSE Training Institute

47

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

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

FortiClient EMS Certificate Management (Contd)

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

```
HQ-FortiGate # diagnose endpoint record list
Record #1:
                IP Address = 10.0.1.100
                MAC Address = 00:50:56:a1:1b:15
                MAC list = 00:50:56:a1:19:7a;00:50:56:a1:1b:15;
                VDOM = root (0)
                EMS serial number: FCTEMS0000101875
                Client cert SN: 64415C90D08FFA3EA2A50A22DA87255FB845BA9
                Public IP address: 206.47.132.124
                Quarantined: no
                Online status: online
                Registration status: registered
                On-net status: on-net
                Gateway Interface: port3
                FortiClient version: 7.0.0
                AVDB version: 88.336
                FortiClient app signature version: 18.143
                FortiClient vulnerability scan engine version: 2.31
                FortiClient UID: 37D54E49E415407D86BAAD3F1F46BB36
                Host Name: AD-Server
                OS Type: WIN64
                OS Version: Microsoft Windows Server 2012 R2 Standard Edition, 6
4-bit  (build 9600)
                Host Description:
                Domain: trainingAD.training.lab
                Last Login User: Administrator
                Owner:
                Host Model: VMware Virtual Platform
                Host Manufacturer: VMware, Inc.
```

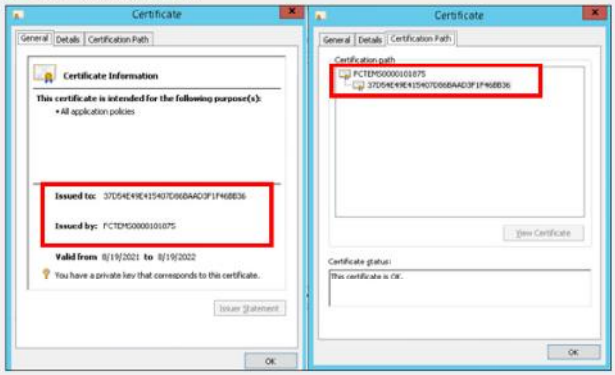

NSE Training Institute

48

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

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

This slide shows that client certificate information is synchronized to the FortiGate.

SSL Certificate-Based Authentication

- An endpoint obtains a client certificate when an it registers to FortiClient EMS
- FortiClient automatically submits CSR request
- FortiClient EMS signs and returns the client certificate
- Certificate is stored in OS certificate store
- By default:
  - Client certificate authentication is enabled on access proxy
  - Empty certificate response is set to block
  - Options can be configured on CLI only

```
config firewall access-proxy
        edit <name>
             set client-cert enable
             set empty-cert-action block
        end
```

- Currently, ZTNA supports the Microsoft Edge and Google Chrome browsers

NSE Training Institute 49

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

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

If the client responds with the correct certificate that the client UID and certificate SN can be extracted from:

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

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

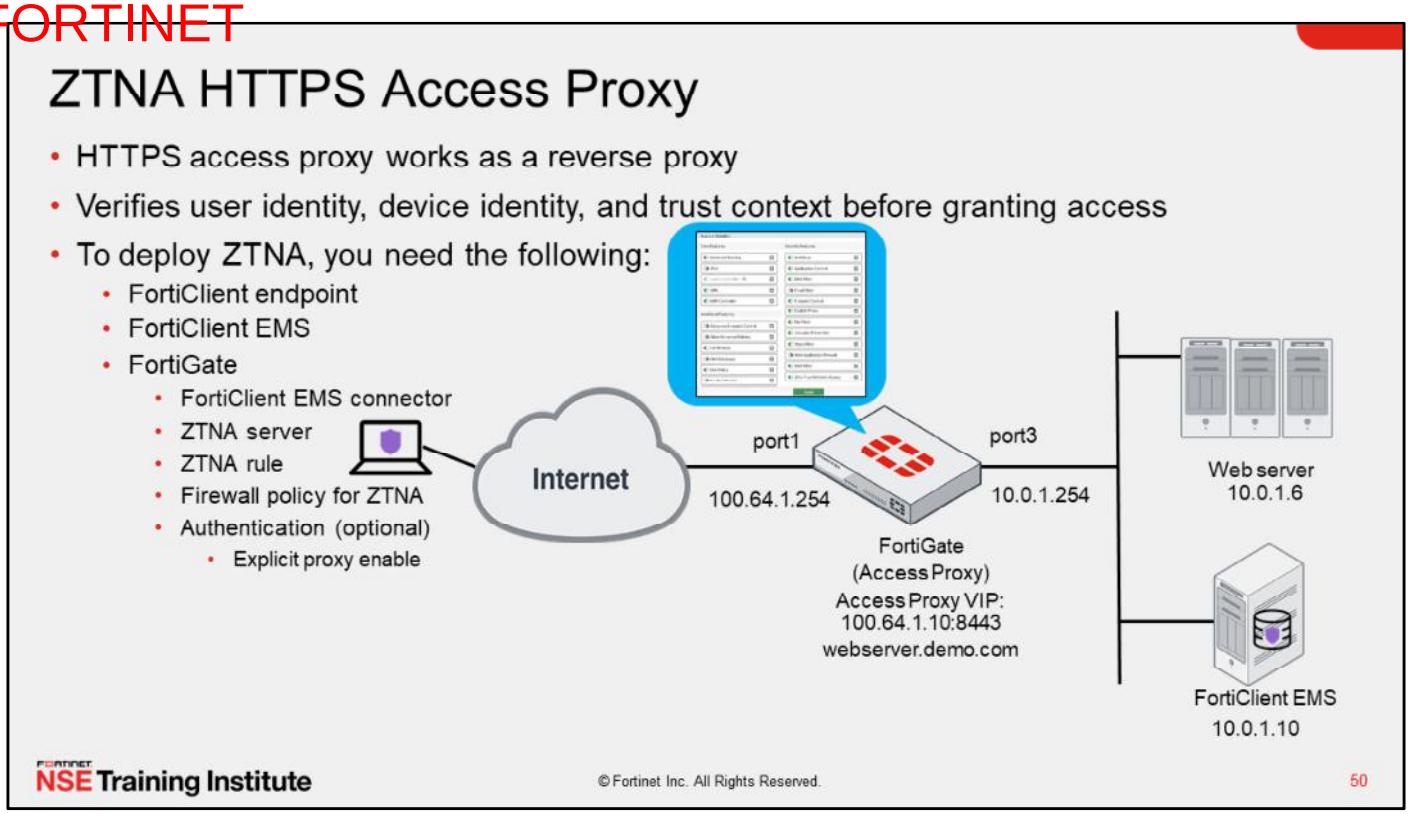

The FortiGate HTTPS access proxy works as a reverse proxy for the HTTP server. When a client connects to a web page hosted by the protected server, the address resolves to the FortiGate access proxy VIP (100.64.1.10:8443), as shown on this slide. FortiGate proxies the connection and takes steps to authenticate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

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

ZTNA configuration on the FortiGate requires the following configuration:

- FortiClient EMS adds a fabric connector in the security fabric. FortiGate maintains a continuous connection to the EMS server to synchronize endpoint device information, and also automatically synchronizes ZTNA tags. You can create groups and add tags to use in the ZTNA rules and firewall policies.
- The ZTNA server defines the access proxy VIP and the real servers that clients connect to. The firewall policy matches and redirects client requests to the access proxy VIP. You can also enable authentication.
- A ZTNA rule is a proxy policy used to enforce access control. You can define ZTNA tags or tag groups to enforce zero-trust role-based access. You can configure security profiles can be configured to protect this traffic.
- The firewall policy matches and redirects client requests to the access proxy VIP. You can define the source interface and addresses that can access the VIP can be defined. By default, the destination is any interface. UTM processing of the traffic happens at the ZTNA rule.

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

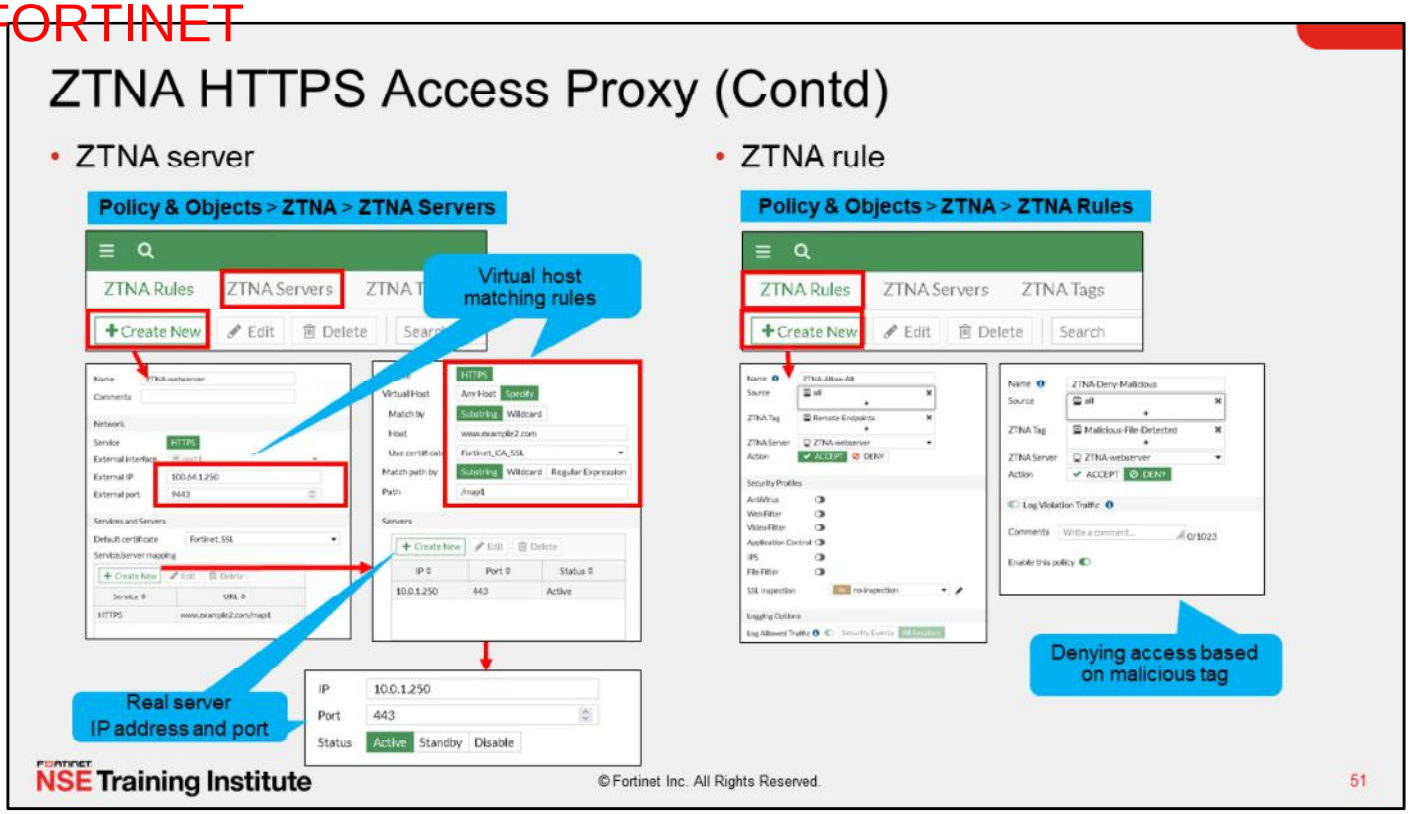

After you can FortiClient EMS as the fabric connector and you sync ZTNA tags with FortiGate, you must create a ZTNA server or access proxy. The access proxy VIP is the FortiGate ZTNA gateway that clients make HTTPS connections to. The service and server mappings define the virtual host matching rules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

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

The firewall policy matches and redirects client requests to the access proxy VIP. You can define source interface and addresses that are allowed access to the VIP. By default, the destination is any interface, so after a policy is configured for full ZTNA, the policy list is organized by sequence. The example on this slide is configured to allow **ALL** services from **all** IP addresses at **port1** as the incoming interface to **ZTNA-webserver** as the destination.

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

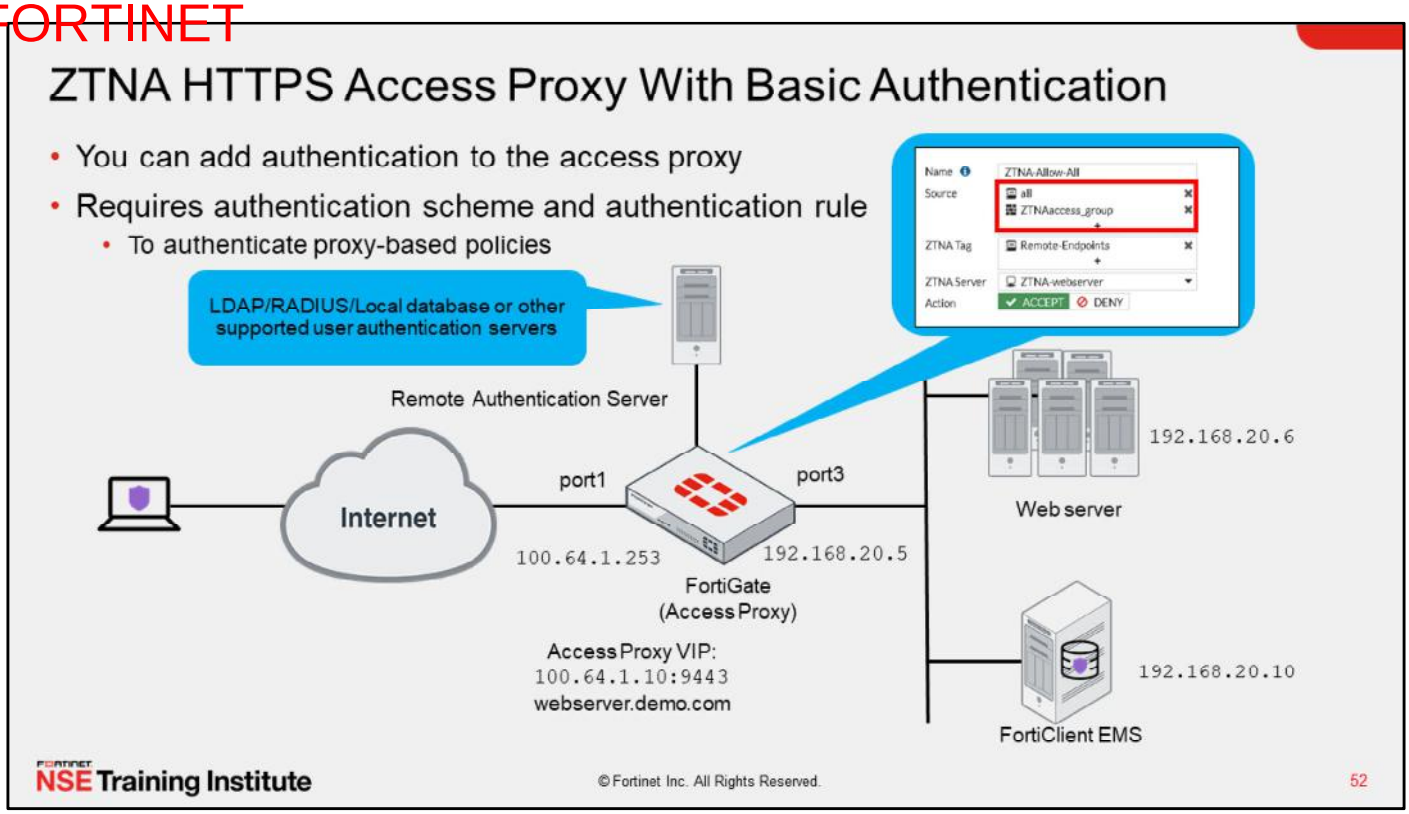

You can add authentication to the access proxy, which requires you to configure an authentication scheme and authentication rule on the FortiGate CLI. You use authentication schemes and authentication rules to authenticate proxy-based policies, similar to configuring authentication for explicit and transparent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

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

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

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

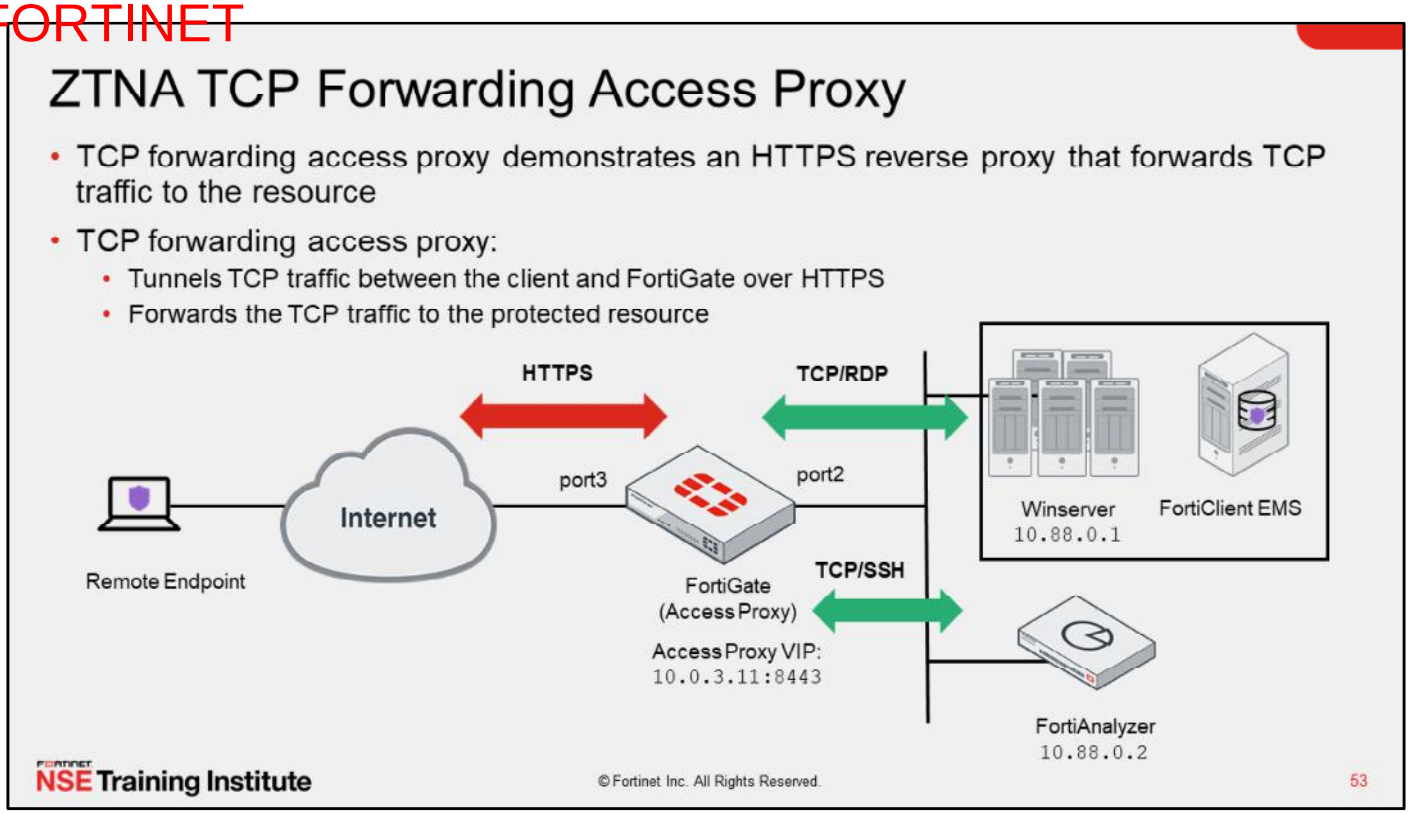

In the example shown on this slide, a TCP forwarding access proxy (TFAP) is configured to demonstrate an HTTPS reverse proxy that forwards TCP traffic to the designated resource. The access proxy tunnels TCP traffic between the client and FortiGate over HTTPS, and forwards the TCP traffic to the protected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

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

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

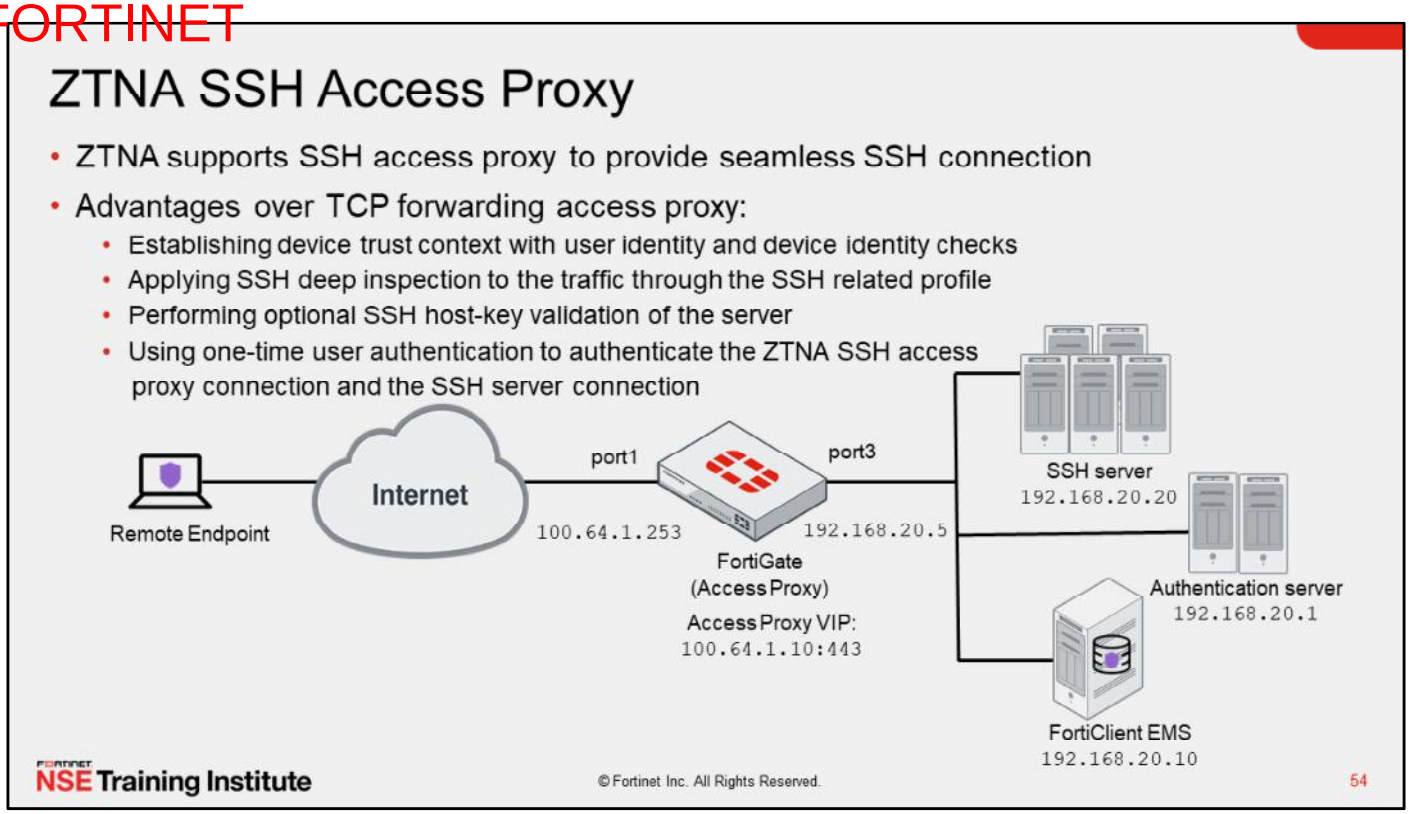

You can configure ZTNA with an SSH access proxy to provide a seamless SSH connection to the server.

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

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

Knowledge Check

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

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

3. What does FortiClient EMS integration ensure?
   - ✓ A. Device identification
   - B. User identification

NSE Training Institute

55

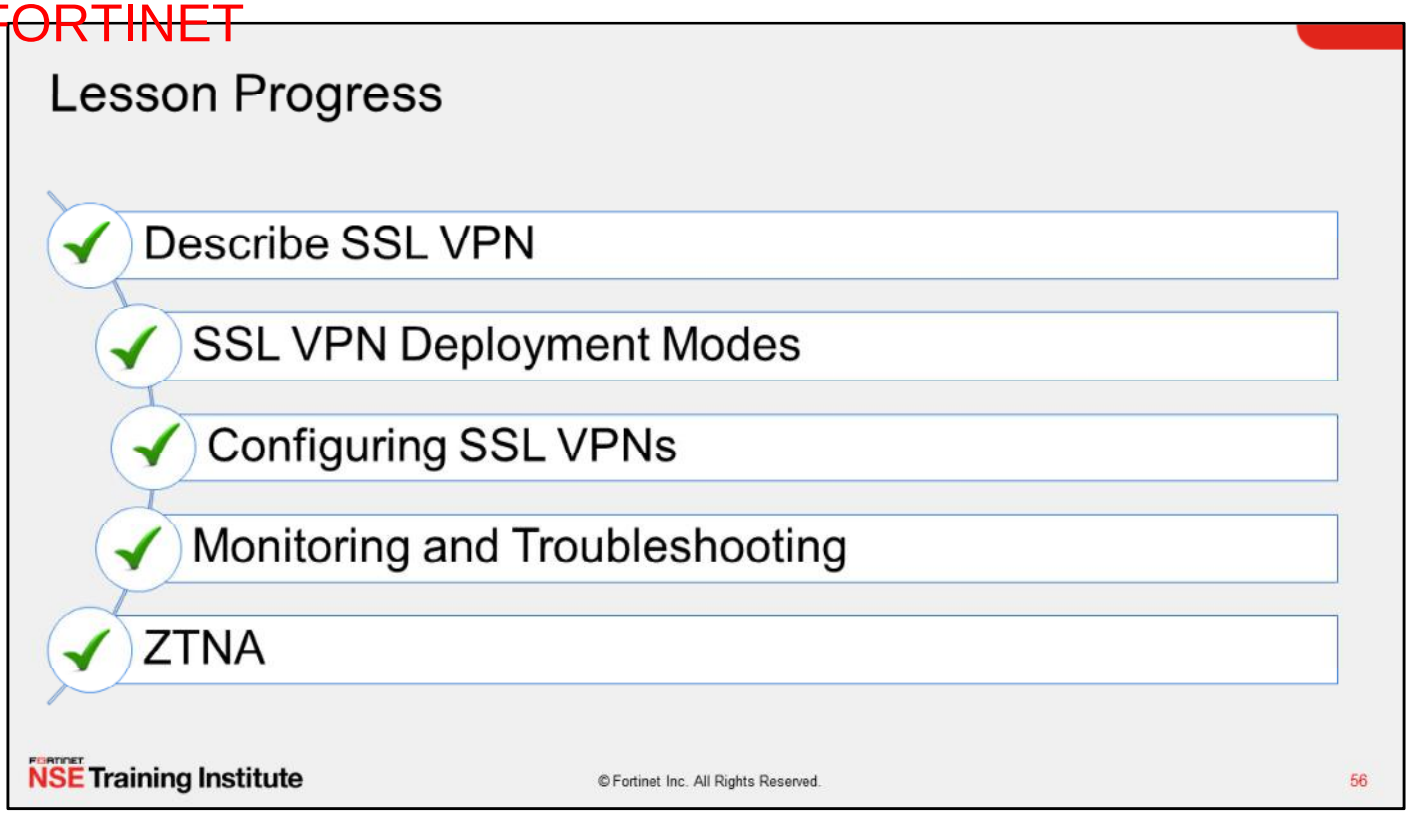

Congratulations! You have completed this lesson.

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

Review

- ✓ Define a virtual private network (VPN)
- ✓ Describe the differences between SSL VPN, IPsec VPN, and ZTNA access
- ✓ Describe the differences between SSL VPN modes
- ✓ Define authentication for SSL VPN users
- ✓ Configure SSL VPN portals
- ✓ Configure SSL VPN settings
- ✓ Define firewall policies for SSL VPN
- ✓ Monitor SSL VPN connected users
- ✓ Review SSL VPN logs
- ✓ Configure SSL VPN timers
- ✓ Troubleshoot common SSL VPN issues
- ✓ Understand the benefits and fundamentals of ZTNA
- ✓ Understand how to establish device identity and trust
- ✓ Understand SSL certificate-based authentication
- ✓ Configure ZTNA access on FortiOS
- ✓ Describe types of ZTNA configuration

NSE Training Institute 57

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

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