




In this lesson, we will show FortiGate administration basics. This includes how – and where – FortiGate fits into your existing network architecture.



Objectives

- Identify major features of FortiGate
- Differentiate between FortiGuard queries & packages
- Choose an operation mode
- Restrict administration to access via management networks
- Create administrator accounts with specific permissions
- Reset a lost “admin” password
- Back up and restore configuration files
- Install new FortiGate firmware
- Run the built-in DNS server on an interface
- Run the built-in DHCP server on an interface



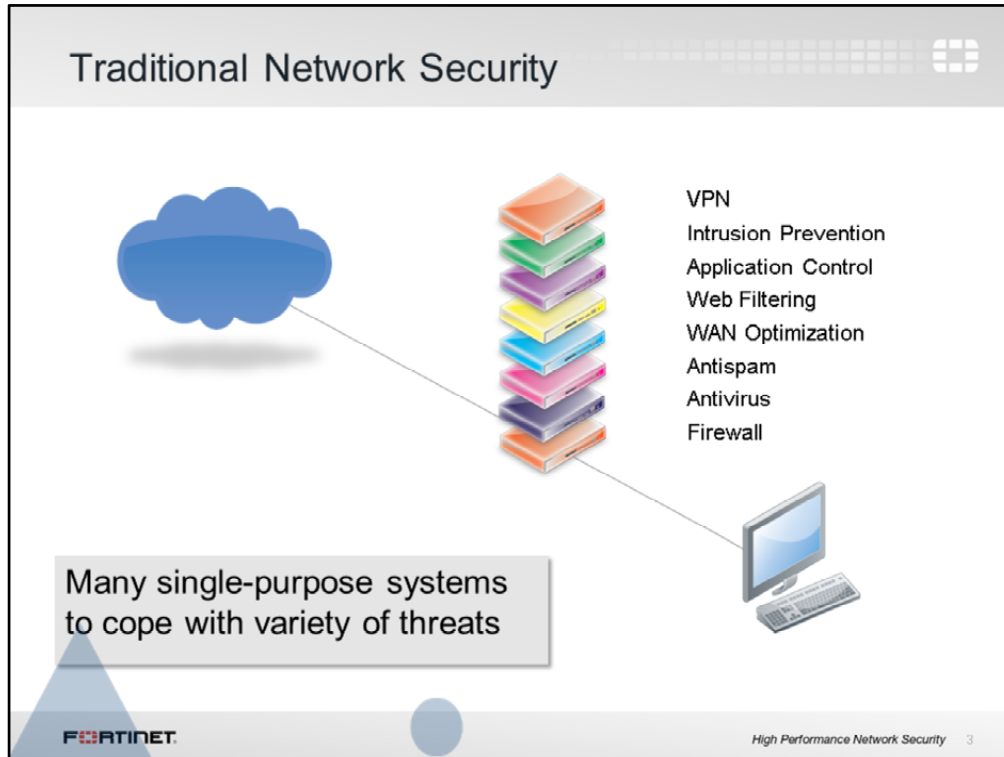
FORTINET

High Performance Network Security 2

After completing this lesson, you should have these practical skills in FortiGate administration fundamentals, such as how to log in, make administrator accounts, do basic network settings, and how to use your FortiGate's GUI or CLI.

You'll also be able to set up FortiGate to act as your local network's DNS or DHCP server.

Lab exercises can help you to test and reinforce your skills.



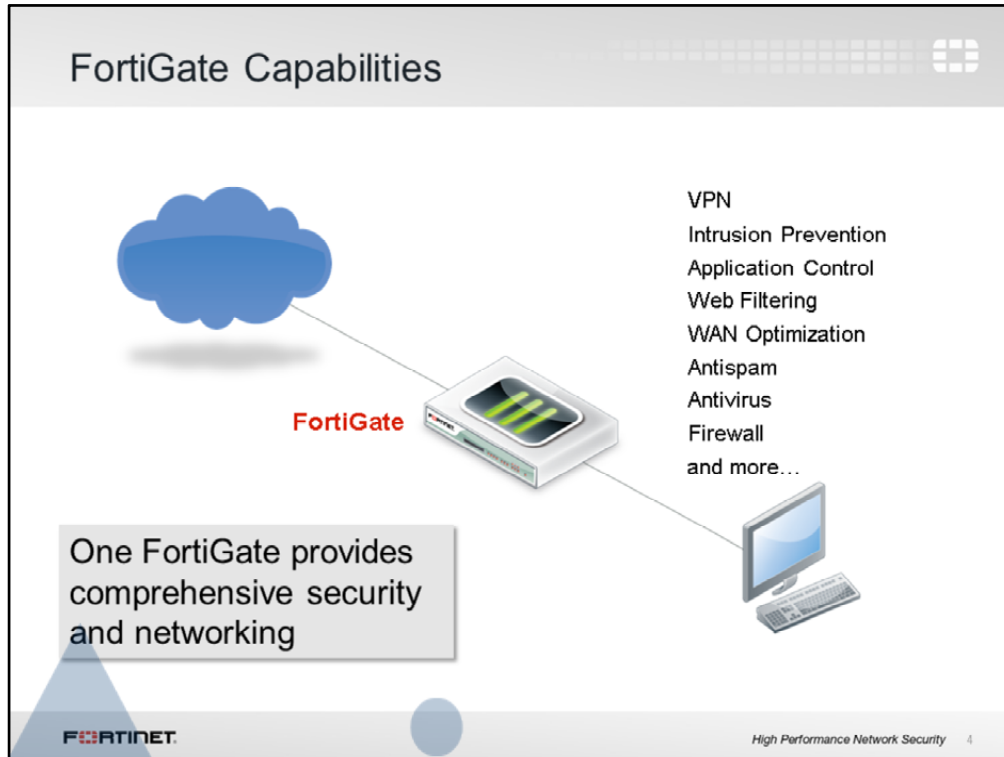
(slide contains animation)

A FortiGate is a “Unified Threat Management” device, but what exactly does this mean? Well, if we look at a typical network security solution, multiple single-purpose devices are used. Each performs a specific task. There is:

(click)

- One device acting as the firewall
- Another device that scans for viruses
- Another device filtering email
- One device to optimize WAN usage
- Another device to filter web sites
- One device for application control
- One device for intrusion prevention
- Another device to provide VPN access

That is a lot of different devices. Most likely, they all have different vendors. All of this can introduce unwanted complexity, and many potential points of failure.



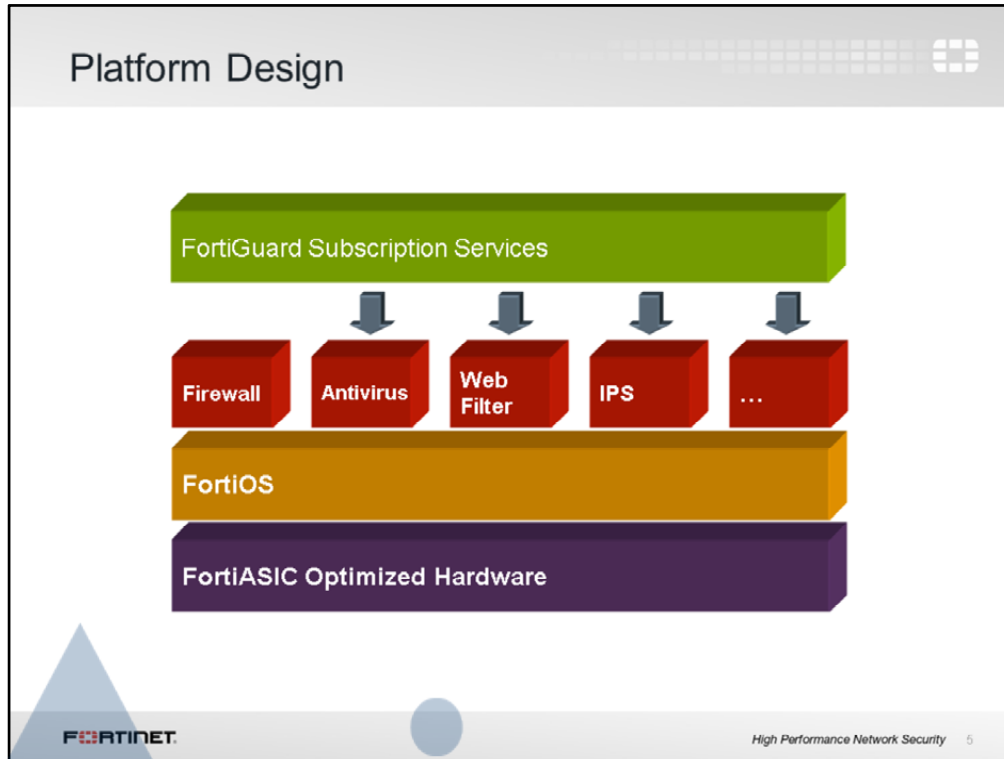
So how is FortiGate different?

FortiGate provides a *comprehensive* approach to security. It even includes some basic accessory network services such as authentication and DHCP. All this and more is combined into a single device. That way, you can reconfigure your network and security deployment by simply accessing one device. Cabling and interfaces between 10 devices? Gone. And it's all from a single vendor. Per-module licensing? Gone.

If you're familiar with Cisco ASA, you may even expect multiple management interfaces. This, too, is simpler on FortiGate. Regardless of whether you are building a VPN or applying antivirus, you can configure it all from one unified GUI or CLI.

How can FortiGate do so many things? *Shouldn't* separate functions be divided among different devices for performance reasons?

In some cases, yes. High load of *one* specific workload may be worth a dedicated device. And Fortinet offers several. But now you have the *choice* – you can specialize if *your* network requires it.






In this architecture diagram, you can see how FortiGate UTM platforms add strength without compromising on flexibility – they are still *internally* modular. Plus:

- **Devices add duplication.** Sometimes, dedication *doesn't* mean efficiency. If it's overloaded, can 1 device borrow free RAM on 9 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?
- **FortiGate hardware isn't just off-the-shelf.** It's carrier-grade. Underneath, most FortiGate models have 1 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. (The exception? Virtualization platforms – VMware, Citrix Xen, Microsoft, or Oracle Virtual Box – have general-purpose vCPUs. But virtualization might be worthwhile due to other benefits, such as distributed computing and cloud-based security.)
- **FortiGate is flexible.** If all you need is firewalling and antivirus, FortiGate won't require you to waste CPU, RAM, and electricity on others. In each firewall policy, UTM modules can be enabled or disabled. You won't pay more to add VPN seat licenses later, either. What requires a subscription? Only FortiGuard subscription services.

FortiGuard Subscription Services

- Internet connection & contract required
- Provided by FortiGuard Distribution Network (FDN)
 - Major data centers in North America, Asia, and Europe
 - Or, from FDN via your FortiManager
 - FortiGate prefers data center in nearest time zone, but will adjust by server load
- **Package updates:** FortiGuard Antivirus & IPS
 - update.fortiguard.net
 - TCP port 443 (SSL)
- **Live queries:** FortiGuard Web Filtering & Antispam
 - service.fortiguard.net
 - Proprietary protocol on UDP port 53 or 8888


High Performance Network Security 6

FortiGuard subscription services give your FortiGate access to 24 x7 security updates powered by Fortinet's researchers. Your FortiGate uses FortiGuard in 2 ways:

- By periodically requesting packages that contain a new engine and many signatures, or
- By querying the FDN on an individual URL or host name

Queries are real-time – that is, FortiGate asks the FDN every time it scans for spam or filtered web sites. Also, queries use UDP for transport – they are connectionless and the protocol is not designed for fault tolerance, but speed. So they require that your FortiGate have a reliable Internet connection.

Downloaded packages like antivirus and IPS, however, aren't that frequent. They use TCP for reliable transport. And their associated FortiGate features continue to function even if FortiGate does not have reliable Internet connectivity. Keep in mind, though, that you should still avoid interruptions. If your FortiGate must try repeatedly to download updates, it can't detect new threats during that time.




Demo Version



In this lesson, we will look at how to monitor your FortiGate, and how to log its system events and network traffic. Since you are implementing a security solution, it is important to know how to appropriately monitor the device's operation. It is vital to have logging and monitoring configured properly and to know how to read the output. Otherwise if you encounter issues, you won't have any messages from FortiGate to help you find out what is happening in your network.

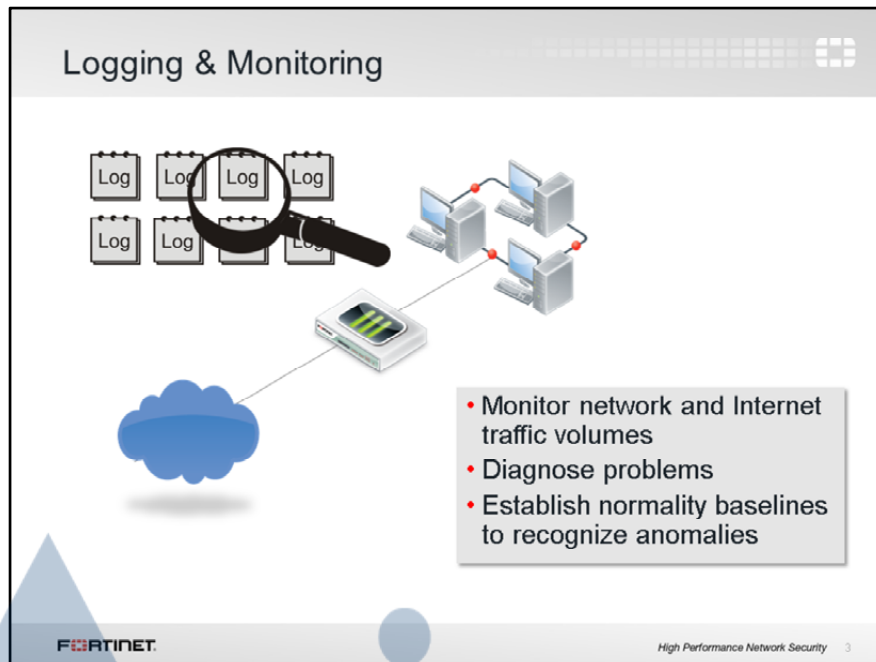
Objectives

- Understand log severity levels
- Recognize the available log storage locations
- Describe the different log types and subtypes
- Understand log structure and behavior
- Configure log settings
- Understand the impact of logs on resources
- Describe how to view log messages
- Describe how to search and interpret log messages



FORTINET High Performance Network Security 2

By the end of this lesson, you'll be able to:
Describe log severity levels
Identify where logs are stored
Describe the different types of logs
Understand log structure and behavior
Configure log settings
Understand the impact of logs on resources
Describe how to view log messages, and finally
Describe how to search and interpret log message



The basic purpose of logs is to help you monitor your network traffic levels, track down problems, establish baselines and a lot more.

Think of your own internal organization, where it is highly probable that more than one administrator has access to your FortiGate device. Since it is not practical to block other administrators from making changes to your FortiGate configuration, you can simply view the log files to find out what is happening on the device—including any changes that were made. Logs help provide you with the big picture so you can make adjustments to your network security, if necessary.

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

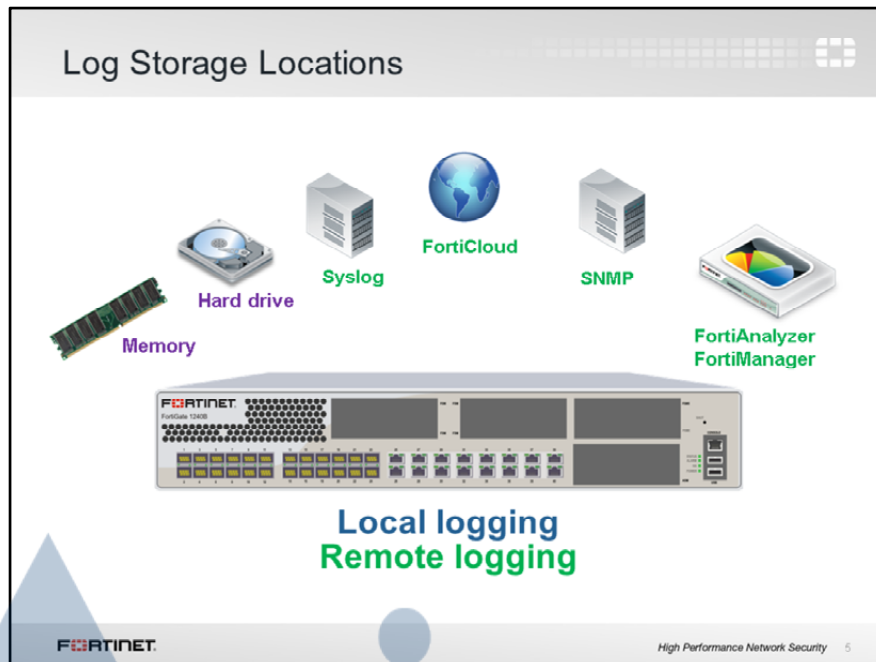
Log Severity Levels

- Administrators define what type of logs are recorded
- All log messages have a severity level to help indicate the importance of the event
 - Emergency → System unstable
 - Alert → Immediate action required
 - Critical → Functionality affected
 - Error → Error exists that can affect functionality
 - Warning → Functionality could be affected
 - Notification → Information about normal events
 - Information → General system information
 - Debug → Debug log messages

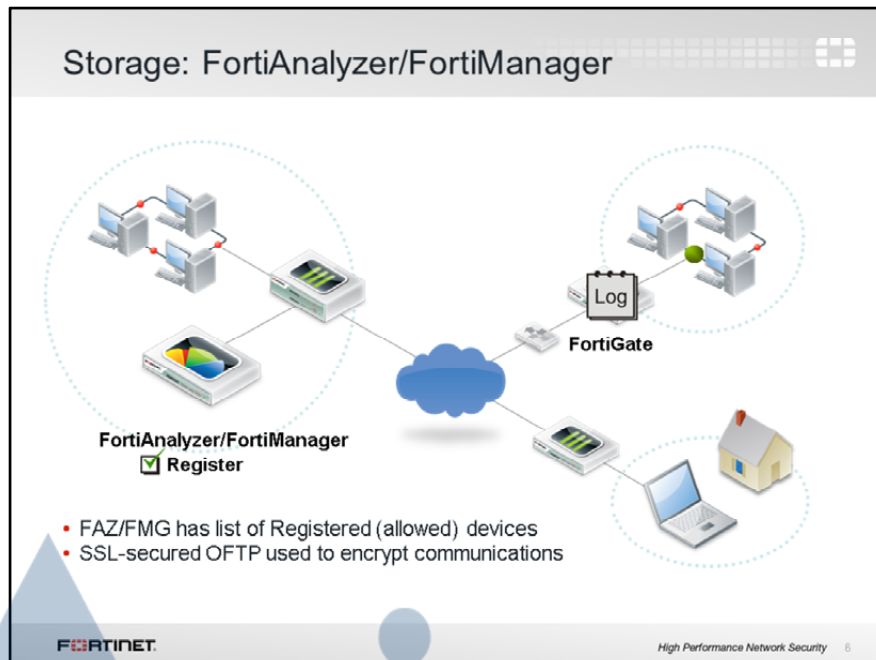
FORTINET High Performance Network Security 4

Each log entry includes a log level that ranges in order of importance from Debug to Emergency. In total there are eight levels. Debug, the lowest level, puts additional information into the event log and is worthless unless you are actively investigating something. Debug is only needed to log diagnostic data, puts more strain on the CPU resources, and requires additional resources to create. Generally the lowest level you want to use is Information.

You and your organization's policies dictate what needs to be logged.



You can choose to store logs in a variety of places both on and off the device. Locally, the FortiGate device has memory and many devices have a built-in hard drive. Externally, you can store logs on Syslog Servers, FortiCloud, SNMP, or a FortiAnalyzer device.



As an external logging device for FortiGate, a FortiAnalyzer or FortiManager is simply viewed as an IP with which the FortiGate can communicate. As a result, you can place a FortiAnalyzer or FortiManager within the same network as a FortiGate, or outside of it. However, a FortiGate can communicate with a FortiAnalyzer or FortiManager only if it is a registered device. So long as the FortiGate is properly registered with the FortiAnalyzer or FortiManager, it accepts incoming logs. Communication between the FortiGate and FortiAnalyzer or FortiManager is done via SSL-encrypted OFTP traffic, so when a log message is generated, it can be safely transmitted across an unsecured network.

Viewing Log Messages(CLI)

```

exe log display
FortiGate # exe log display
1743 logs found.
10 logs returned.
1: date=2013-12-09 time=13:24:08 logid=0001000014 type=traffic subtype=local level=notice vd=root srcip=192.168.1.112
srcport=61597 srcintf="port1" dstip=192.168.1.110 dstport=80 dstintf="root" sessionid=562 status=close policyid=0
dstcountry="Reserved" srcountry="Reserved"trandisp=noop service=HTTP proto=6 app="Web Management" duration=12
sentbyte=1045 rcvbyte=2633 sentpkt=5 rcvdpkt=6

```

- Set up log filter first

```

exe log filter
FortiGate-UM64 # exe log filter
category          Category.
device            Device to get log from.
dump              Dump current filter settings.
field             Filter by field.
ha-member
max-checklines    Maximum number of lines to check.
reset             Reset filter.
start-line        Start line to display.
view-lines        Lines per view.
FortiGate-UM64 # exe log filter _

```

FORTINET High Performance Network Security 18

Rather than look at raw logs or logs through the GUI, you can also display log messages from the CLI. This allows you to set up a number of filters on the logs that display and capture the output to a file and send it via the options you specify, such as FTP.




Demo Version




In this lesson, we will show you how to pass traffic through FortiGate, and explain how that works. At its core, FortiGate is a firewall, so almost everything that it does to your traffic is linked into your firewall rules.

Objectives

- Match traffic to firewall policies by:
 - Source IP address, device ID/type, or user
 - Interface or zone
- Reorder firewall policies for correct matching
- Identify components of firewall policies
- Explain 'implicit fall through' for authentication
- Choose between central NAT vs. source NAT in the policy
 - Apply source NAT with IP pools (overload vs. one-to-one, fixed port range and port block allocation)
- Configure destination NAT with virtual IPs or a virtual server
- Log blocked traffic





High Performance Network Security 2

After this lesson, you should be able to properly identify the different components used in a firewall policy. You'll be able to configure firewall policies and arrange them to correctly match traffic.

Objectives

- Modify the session TTL
- Describe behavior differences when processing is offloaded to network processors
- Use a SIP session helper for VoIP
- Shape traffic by shared and per-IP limits
- Compare flow-based vs. proxy-based inspection
- Enable SSL/SSH inspection
- Monitor & debug the flow of traffic through firewall policies



FORTINET High Performance Network Security 3

You'll also be able to apply UTM and other features through the firewall policy, test your policies, and monitor traffic passing through them.

What Are Firewall Policies?

- Policies define:
 - Which traffic matches them
 - How to process traffic that matches
- When packet for new IP session arrives, FortiGate looks for matching policy
 - Only first matching policy applies
 - Starts at top of list
- **Implicit deny**
 - No matching policy? FortiGate drops packet

Seq #	Source	Destination	Schedule	Service	Action	Log	NAT	Count
1	all	all	always	ALL	ACCEPT			0Packets / 0B
2	all	Windows PC	always	HTTP	ACCEPT			126,054 Packets / 75,52 MB
3	all	all	always	HTTPS	ACCEPT			180,260 Packets / 126,95 MB
4	all	Linux_ETH0	always	SSH	ACCEPT			0Packets / 0B
5	all	Linux_ETH0	always	ALL	ACCEPT			572,789 Packets / 562,34 MB
7	all	all	always	IMB	DENY			0Packets / 0B
8	all	all	always	HTTP	ACCEPT			0Packets / 0B
9	all	all	always	HTTPS	ACCEPT			0Packets / 0B
10	all	all	always	ALL	DENY			0Packets / 0B

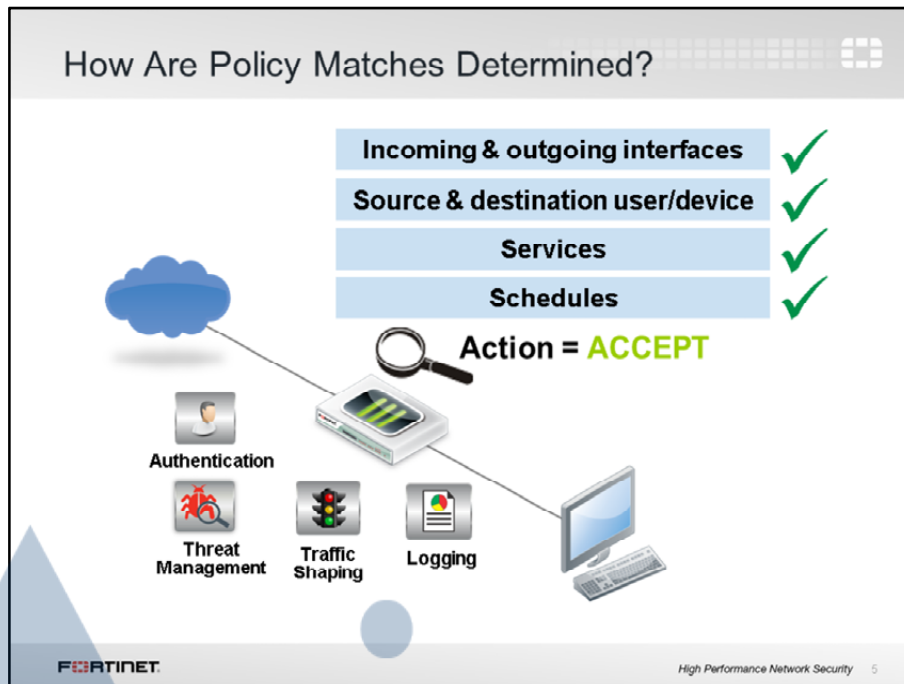
FORTINET High Performance Network Security 4

To begin, let's talk about what firewall policies are.

Firewall policies define which traffic matches, and what FortiGate will do if it does.

Should the traffic be allowed? This is decided first based on simple criteria such as the source. Then, if the policy itself does not block the traffic, FortiGate begins more computationally expensive UTM inspection, such as 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 NAT be applied? Authentication required? Firewall policies also determine that. Once processing is finished, FortiGate forwards the packet towards its destination.



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

- Ingress and egress interfaces
- Source and destination, by IP address, device ID, or user
- Network service(s) (that is, IP protocol and port number)
- Schedule

Once FortiGate finds a matching policy, it applies its settings for packet processing. Is antivirus scanning applied? Will source NAT be applied?

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, and 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) nor 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, limits the bandwidth consumption, and logs blocked connection attempts.



Firewall policies appear in an organized list. It's either organized into a section view, or global view.

Usually, it will appear in section view. Each section contains policies for that ingress-egress pair.

Alternatively, you can choose to view your policies as a single comprehensive list, by selecting *Global View* at the top of the page.

Policy sequence numbers define the order in which rules are processed. Policy IDs are identifiers. By default sequence numbers are displayed on the GUI. CLI commands, however, use policy ID: *edit <ID>*. This may confuse the administrator in to modifying the wrong policy. To avoid such errors add the policy ID to the GUI using the column settings.

```

Session Table: TCP Example

# diagnose sys session list
session info: proto=6 proto state=05 expire=89
               timeout=3600 flags=00000000 av idx=0 use=3
               bandwidth=204800/sec guaranteed_bandwidth=102400/sec
               traffic=332/sec prio=0 logtype=session ha_id=
               hakey=4450
               tunnel=/
               state=log shape may dirty
               statistic(bytes/packets/err): org=3408/38/0
               reply=3888/31/0 tuples=2
               orgin->sink: org pre->post, reply pre->post oif=3/5
               gw=192.168.11.254/10.0.5.100
               hook=post dir=org act=snat 10.0.5.100:1251-
               >192.168.11.254:22(192.168.11.105:1251)
               hook=pre dir=reply act=dnat 192.168.11.254:22-
               >192.168.11.105:1251(10.0.5.100:1251)
               pos/(before,after) 0/(0,0), 0/(0,0)
    
```

TCP state

Session TTL

Traffic shaping

Traffic counts

NAT operation

In this example, you can see the session TTL, which reflects how long FortiGate can receive no packets until it will remove the session from its table.

Proto_state for TCP is taken from its state machine, which we'll talk about next.

Traffic shaping manages your bandwidth. Traffic counters are the overall counters for the session, and determine how much data was sent and received.

NAT actions are also tracked.



Demo Version





In this lesson, we will show you how to use authentication on the firewall policies of a FortiGate.

Normal firewall policies involve separating devices based on the IP address or subnet involved. Adding authentication to firewall policies, however, provides a mechanism to make decisions on not just where the device is, but who is using the device.

Objectives

- Explain firewall authentication
- Describe the different methods of authentication available on FortiGate devices
- Identify which authentication protocols are used with each method of authentication
- Configure Two-Factor Authentication (OTP and Tokens)
- Describe authentication types (active and passive)
- Create authentication policies
- Configure Captive Portal and disclaimers
- Configure authentication timeout
- Describe and configure users/user groups:
 - LDAP, RADIUS
 - FortiGate
- Monitor firewall users





High Performance Network Security 2

After completing this lesson, you should have a solid understanding of the mechanics of authentication on a FortiGate as well as some practical skills configuring firewall authentication.

Authentication

- Confirms identity of a user or device
- Once the FortiGate identifies the user/device, FortiGate applies the right firewall policies and profiles to allow / deny access to each network resource



FORTINET

High Performance Network Security 3


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

Authentication allows action based on the user, not just the IP address. In this way, inspection rules follow individuals across multiple devices.

Methods of Authentication

You can use the following methods of authentication for firewall authentication:

- Local password authentication
- Remote password authentication
- Two-factor authentication
 - Enabled on top of an existing method
 - Requires something you know *and* something you have


High Performance Network Security 4

Not all available methods of authentication can be used for firewall authentication (for example, certificate-based authentication cannot be used). You can, however, use local password authentication, remote password authentication, and two-factor authentication. Two-factor authentication is slightly different from the others, as it is enabled on top of an existing method—it cannot be enabled without first configuring one of the other methods.

In this lesson, we will discuss all three available methods.

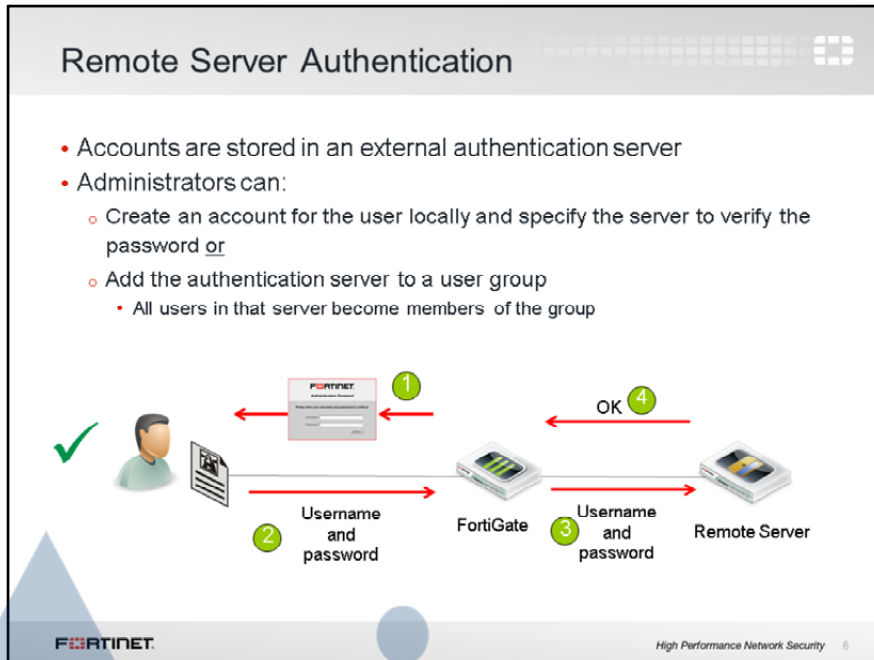
Local Password Authentication

- Local password authentication is based on user accounts stored locally on FortiGate
 - For each account, a user name and password (credentials) is stored

The diagram illustrates the local password authentication process. On the left, a user is shown with a green checkmark, indicating successful authentication. A laptop icon represents the user's device. A red arrow labeled '2' points from the laptop to the FortiGate device, labeled 'User name and password'. A second red arrow labeled '1' points from the FortiGate device back to the user, representing the response. The FortiGate device is depicted as a physical hardware unit and is labeled 'Fortigate'. Above the device is a 'Fortinet Authentication Portal' window. The Fortinet logo and 'High Performance Network Security' are visible at the bottom of the slide.

The first and simplest method of authentication is Local Password Authentication. User account information (user name and password) is stored locally on the FortiGate device, so there is no lookup to an external server for user validation.

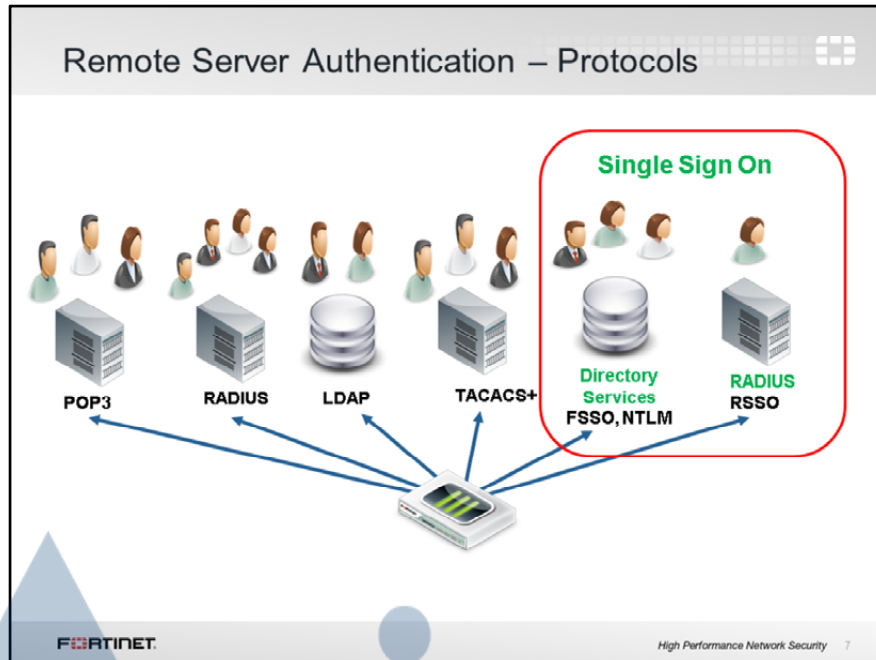
Local Password Authentication is the simplest method of authentication to configure, since you only need access to the FortiGate. Other methods of authentication are more complex, as they involve configuring the exchange of information between the FortiGate and a remote server as well as configuring the various users and user groups on the server itself. Troubleshooting in those situations becomes more complicated, as you need to examine both the FortiGate and external server. With Local Password Authentication, you need only examine the FortiGate.



The second method of authentication is remote server authentication (or server-based password authentication). This includes any form of authentication where the final decision on user credentials is made by an external server—not the FortiGate. This method is desirable when multiple FortiGate devices need to authenticate the same users or user groups.

With remote server authentication, user information is sent from the FortiGate to a remote server. The remote server then evaluates the information it receives and sends a response. The server response is examined by FortiGate and consults its configuration to deal with the traffic. However, it is the server — not the FortiGate — that has final authority over evaluating the user credentials.

With Remote Server Authentication, the FortiGate does not store all (or, in the case of some configurations, any) of the user information locally.



Multiple protocols are supported for remote user authentication, including POP3, RADIUS (includes server authentication and the single sign on method, RSO), LDAP, and TACACS+.

Single sign on (SSO) methods, such as FSSO, NTLM, and RSO, are also supported for remote user authentication.



Demo Version




In this lesson, we will show you how to use and configure SSL VPN. SSL VPNs are an easy way of providing access to your private network for remote users.



Objectives

- Understand and configure different operating modes for SSL VPN
- Configure SSL VPN options, such as bookmarks and realms
- Configure additional security for SSL VPN access
- Monitor SSL VPN connected users
- Configure firewall policies and authentication for SSL VPN



FORTINET High Performance Network Security 2

After completing this lesson, you should have these practical skills that you can use to configure an SSL VPN for your organization.

Virtual Private Networks (VPN)

- Allows users to remotely access network resources as if they were physically connected to the local network
- Used when there is the need to transmit private data across a public network
- Provides an encrypted point-to-point connection, so it cannot be intercepted by unauthorized users
- Employs different security methods to ensure that only authorized users can access the private network

FORTINET

High Performance Network Security 3


A virtual private network enables users to remotely and securely access private resources as if they were locally connected.

It is generally used to transmit private information safely between LANs separated by an untrusted public network such as the Internet, so it is not only implemented for providing access to mobile users, but also for interconnecting geographically dispersed networks across the Internet. The user data travelling inside a VPN tunnel is encrypted, so it cannot be intercepted by unauthorized users. VPNs also use security methods to ensure that only authorized users can establish the VPN and access the private network's resources.


FortiGate VPN

SSL VPN

- Typically used to secure web transactions
- HTTPS tunnel created to securely transmit application data
- Client signs on through secure web page (SSL VPN portal) on the FortiGate device




VPN



IPsec VPN

- Well suited for network-based legacy applications
- Secure tunnel created between two host devices
- IPsec VPN can be configured between FortiGate unit and most third-party IPsec VPN devices or clients

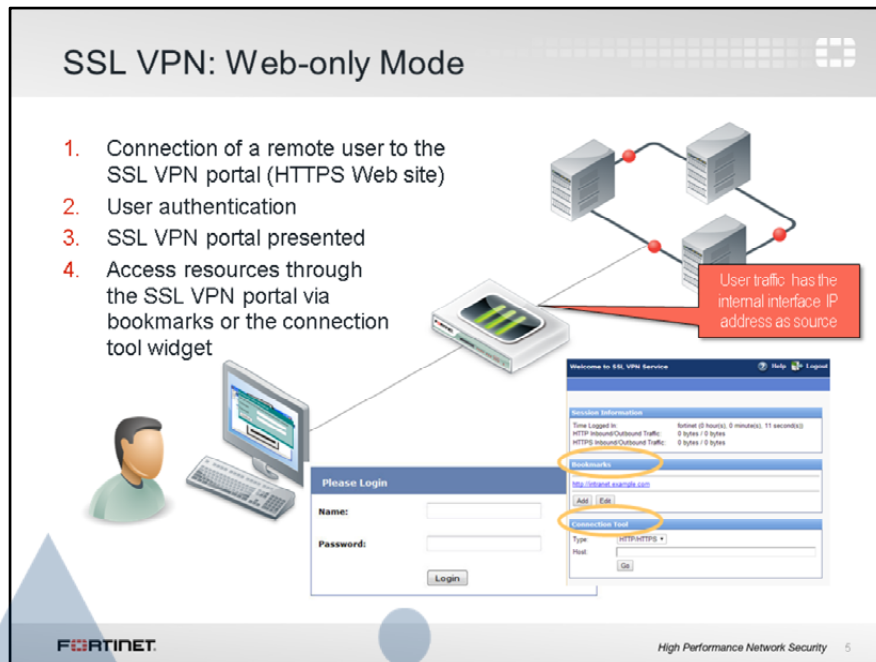

High Performance Network Security 4

The most common type of VPNs are SSL VPN and IPsec VPN.

SSL VPNs are commonly used to secure web transactions. Clients connect to a web portal and log in. It is essentially meant to connect a PC to a private network. This approach is simple in that users only need a regular web browser to connect and are not usually required to install any kind of special software or go through a complex setup. They simply need to access an HTTPS web site and log in. This makes SSL VPN an ideal solution for users who are either not technically skilled, or who need to connect from public computers.

IPsec is also used to connect a PC to a private network. However, there are some important differences. Firstly, SSL VPN access is through a web portal, whereas IPsec is not. Finally, IPsec is a standard protocol supported by most vendors, so a VPN session can be established not only between two FortiGate devices, but also between different vendor devices. By comparison, SSL VPN can only be established between a client PC and an end device.

In this lesson, we are going to focus on SSL VPN.



Web-only mode is used to connect using HTTPS to the FortiGate device from any browser. Once connected, users need credentials in order to pass an authentication check. Once authenticated, users are presented with a portal that contains possible resources for them to access. Different users can have different portals with different resources and access permissions.

One of the widgets contains links to all or some of the resources available for the user to access. Another widget allows users to type the URL or IP address of the server they want to reach. A Web-only SSL VPN user makes use of these two widgets to access the internal network. The main advantage of Web-only mode is that it is clientless. This means the user is not required to install any client VPN software to obtain access. However, Web-only mode has two main disadvantages: First, all interaction with the internal network must be done from the browser exclusively (through the web portal). External network applications running on the user's PC cannot send data across the VPN. Second, a limited number of protocols are supported, such as HTTP/HTTPS, FTP, RDP, SMB/CIFS, SSH, Telnet, VNC, Ping.



Demo Version



In this lesson, we will show you how to set up site-to-site IPsec VPN.

VPNs are heavily used in today's IT infrastructure to join private corporate networks across the Internet. IPsec is an RFC standard. Whether you have FortiGate devices only or mix in another vendor's devices, the principles are essentially the same.

Objectives

- Define the architectural components of IPsec VPN
- Identify the phases of Internet Key Exchange (IKEv1)
- Compare route-based vs. policy-based VPNs
- Deploy a site-to-site VPN between 2 FortiGates
- Monitor VPN tunnels



FORTINET


High Performance Network Security 2

After completing this lesson, you should have these practical skills that you can use to set up a simple IPsec tunnel for a site-to-site VPN.

During this, we will explain how to choose between configuring a policy-based or route-based VPN. You will also learn how to verify the status of each tunnel.

What Are Virtual Private Networks (VPN)?

- Also called “tunnels”...
- Transmits private data across untrusted networks (Internet)
- **Secure access, like local LAN connection**
 - Authorizes access
 - External device gets a private network IP
 - Encrypted
 - Cannot be read / tampered with if intercepted
- **Types:**
 - PPTP
 - L2TP
 - SSL VPN
 - IPsec...



FORTINET High Performance Network Security 3

A Virtual Private Network (VPN) allows people in remote places – separated by the Internet – to securely access resources on your local network. For example, if workers are traveling or working from home, you can use a VPN to give LAN access to them. You can also use a VPN to interconnect multiple campuses.

There are multiple types of VPN: PPTP, L2TP, SSL VPN, and IPsec are popular choices.

- PPTP is fast, but security is weak, and easily defeated.
- IPsec requires a gateway or installation of client software. So it is more complicated to set up for mobile users than SSL VPN, where they can simply utilize their web browser instead.
- SSL VPN is designed for tunnels between a single client and a LAN, not between entire offices.

Because of this, many networks now use a combination of SSL VPN – for mobile user access – and IPsec or L2TP – for tunnels between offices.

Often, “tunnel” is used as a synonym for “VPN,” although not *all* VPNs technically are tunnels, as we will see in a minute.



Demo Version




In this lesson, we will show you how to use antivirus scanning on a FortiGate.

Since antivirus scanning is one of the features that, depending on your configuration and chosen signature database, can use significant RAM, we will also show you how to resolve “conserve mode.”

Objectives

- Categorize malware types & evasion techniques
- Describe FortiGate's antivirus techniques
- Differentiate between proxy-based vs. flow-based scans
- Configure antivirus scanning, including:
 - Grayware
 - Heuristic
 - Sandboxing
- Update antivirus scans via FortiGuard services
- Submit new virus samples to Fortinet
- Identify the order of evaluations in an antivirus scan
- Diagnose causes of conserve mode




FORTINET High Performance Network Security 2

After completing this lesson, you should have these practical skills. Not only will you be able to configure antivirus, but you should have a better understanding of how virus scanning works, along with knowledge of some tools to help you optimize memory usage on your FortiGate.

What is Malware?

malware (n): a category of software that is capable of copying itself and has a detrimental effect, such as corrupting the system or destroying data

- **Viruses**
 - Behavior modeled after a biological virus
 - Often injects code into files, like biological viruses inject their own DNA into cells
 - Does not require user permission, or tricks user into giving permission
 - Infects and spreads on its own
 - Very small
- **Grayware**
 - User permission *is* required for installation
 - Spyware etc. often bundled with shareware / free software
 - Size varies



FORTINET High Performance Network Security 3

How old are viruses? In 1949, John Von Neumann gave lectures at the University of Illinois about what he called “self-replicating automata.” On ARPANET, the precursor to the Internet, the first virus, named Creeper, was detected in 1971.

Since then, malicious software has evolved into many types. Technically, although we often refer to all malware as viruses, not every piece of unwanted software behaves like a virus – malware is not always self-replicating, and sometimes users willingly install it. To include viruses, worms, Trojans, spyware and all others, we now use the term “malware.”

Malware can be divided into 2 major types: viruses, which infect the computer and spread on their own (generally via an exploit), such as Flash ad banners whose binaries contain buffer overflow code grayware which requires some kind of user interaction but convinces them that the benefit outweighs the cost, such as browser toolbars that also track the user’s activity and insert its own ads into web pages



Demo Version



In this lesson, we will show you how your web browsers can use FortiGate as an explicit proxy.

Objectives

- Enable the explicit web proxy on FortiGate
- Use a PAC file and WPAD to configure web browsers with multiple web proxies and server exemptions
- Reduce WAN bandwidth usage
- Improve responsiveness via web cache
- Authenticate multiple web proxy users that share the same source IP address
- Apply proxy policies with URL address objects
- Monitor current proxy users



FORTINET

High Performance Network Security 2

After completing this lesson, you should have these practical skills.

You will learn how to configure both FortiGate and the web browsers that will use it as an explicit proxy. Since you can alternatively use an implicit proxy, we will also explain why in some cases you might want an explicit proxy instead.

What is an Web Proxy?

- Proxy forwards requests for clients to web site
 - May cache responses
 - If cache exists, proxy is a shortcut: responds itself with cache, *doesn't* forward request to web site
- Two TCP connections:
 1. From client to proxy
 2. From proxy to server

```

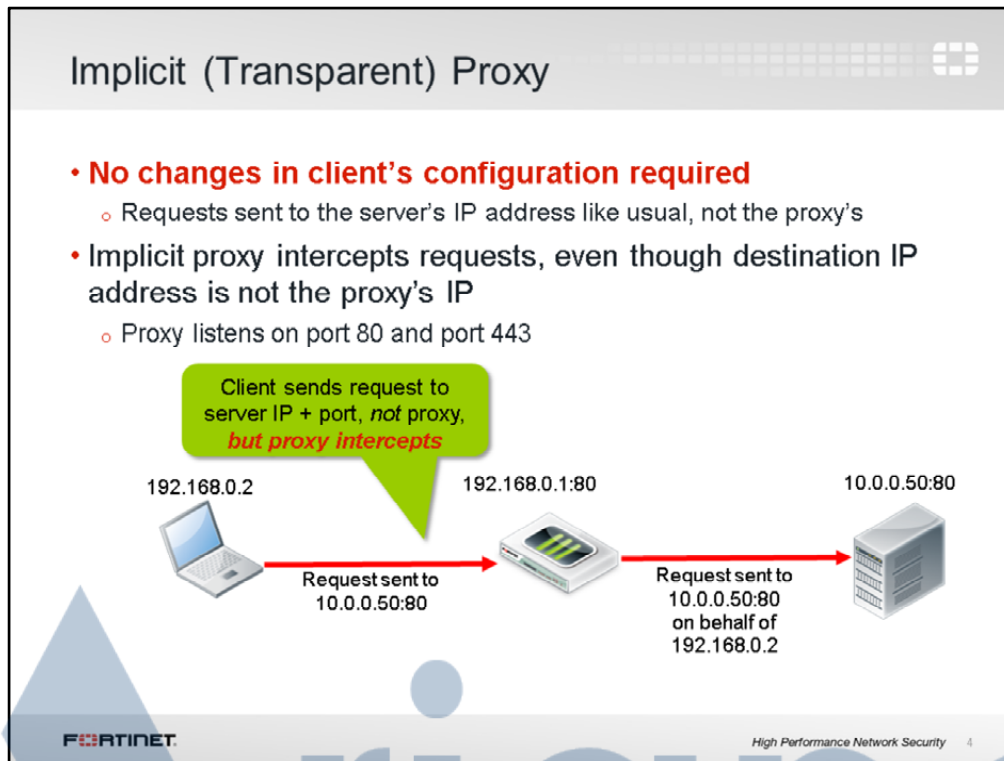
graph LR
    Client[Client] -- Connection 1 --> Proxy[HTTP Proxy]
    Proxy -- Connection 2 --> Server[Server]
            
```

FORTINET
High Performance Network Security 3

A proxy receives or intercepts requests from a client to a server. If allowed, and if no cache is available, it forwards the request to the server on behalf of the client.

Two sessions are created: one from the client to the proxy, and another one from the proxy to the server.

How is this different from an implicit proxy, sometimes called a transparent proxy?



An implicit proxy server does not require any configuration change on the clients. Clients continue to use the web just like they would without a proxy.

Clients send requests to the web server's IP address and port number. The proxy intercepts the client's requests transparently – that is, at the IP layer, the destination address doesn't change.

Does this mean that implicit proxies don't require any configuration changes, anywhere? Not necessarily.

Usually, both incoming and outgoing traffic is routed through FortiGate. As a result, web browsing is already being routed through FortiGate, where it can be intercepted by the transparent proxy. But if clients' traffic *isn't* currently routed through FortiGate, then you must reconfigure routing so that the packets will be routed through FortiGate, where the implicit proxy can intercept.



Demo Version