FORTINET CERTIFIED PROFESSIONAL
Public Cloud Security

AWS Cloud Security Administrator

Study Guide

FortiOS 7.4

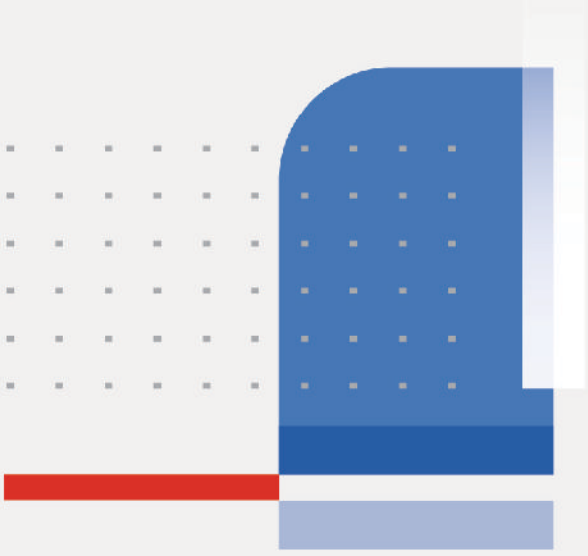

FORTINET
Training Institute

Fortinet Training Institute - Library

https://training.fortinet.com

Fortinet Product Documentation

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 Product Support

https://support.fortinet.com

FortiGuard Labs

https://www.fortiguard.com

Fortinet Training Program Information

https://www.fortinet.com/nse-training

Fortinet | Pearson VUE

https://home.pearsonvue.com/fortinet

Fortinet Training Institute Helpdesk (training questions, comments, feedback)

https://helpdesk.training.fortinet.com/support/home

FORTINET

2/28/2024

TABLE OF CONTENTS

In this lesson, you will learn about public cloud security.

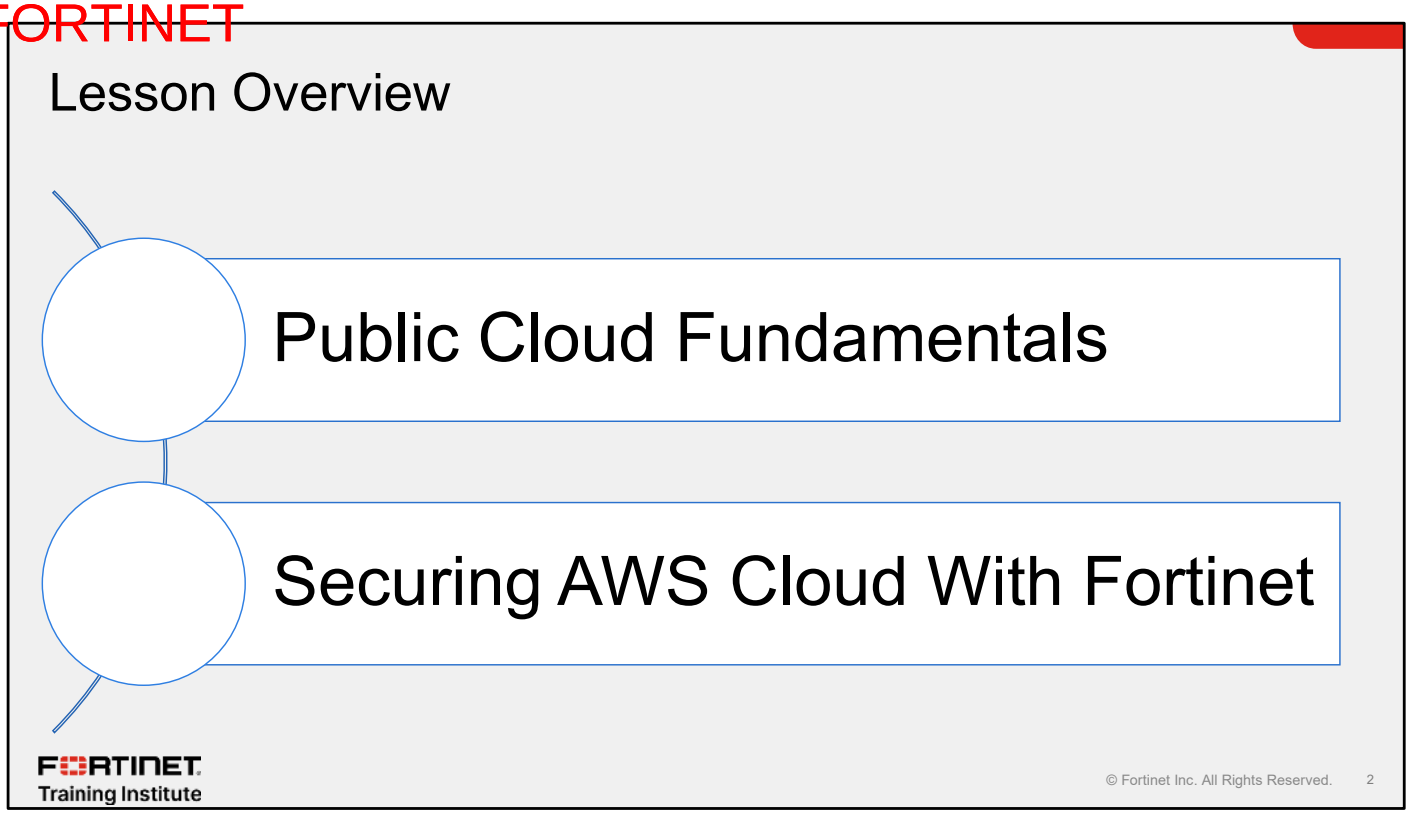

In this lesson, you will learn about the topics shown on this slide.

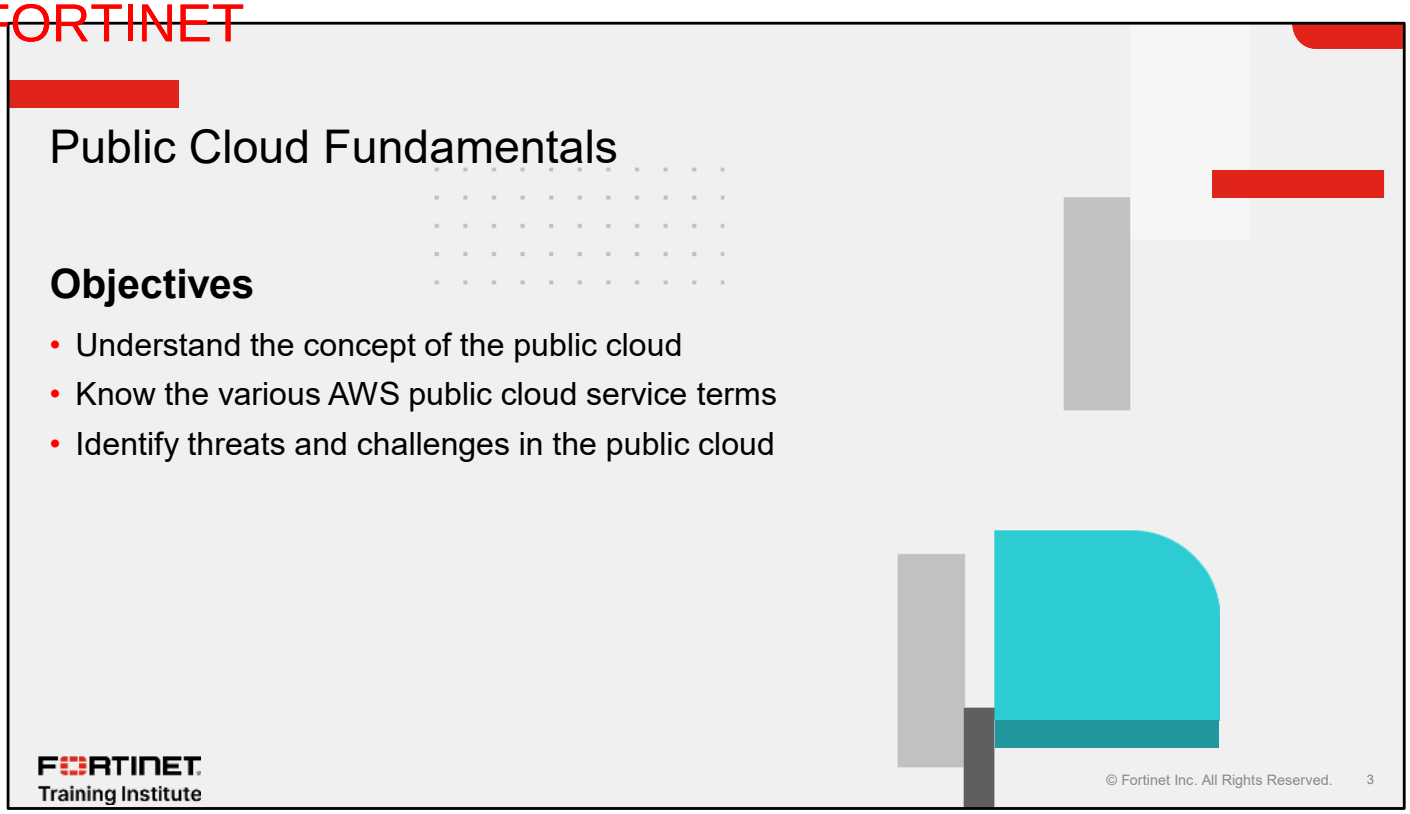

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding the fundamentals of the public cloud, you will be able to understand how the public cloud applies to your network.

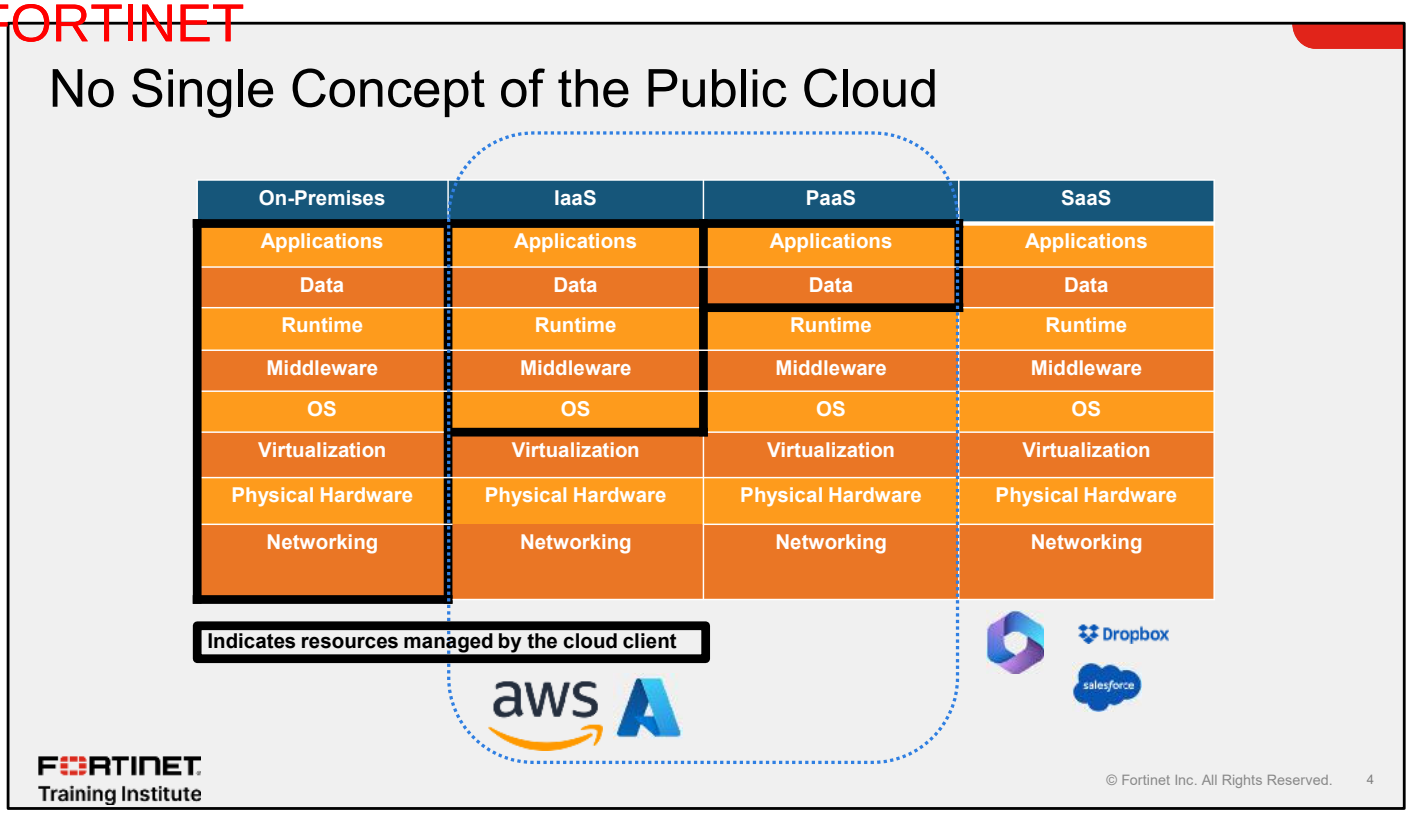

The term *public cloud* comes from the marketing world, but in the technology world, public cloud can mean one or more specific concepts. As shown on this slide, there are different service models of a public cloud. In a traditional on-premises scenario, all the servers, switches, and databases run locally, on site. The VMs that you deploy during the labs are considered to be Infrastructure-as-a-Service (IaaS). In an IaaS solution, some parts of networking and services are managed by the vendor, and other parts are managed by the customer. There is also a solution called Platform-as-a-Service (PaaS), in which the customer is responsible for programming applications and the rest of the services are managed by the vendor. Finally, in the Software-as-a- Service (SaaS) solution, the customer is using the services as a consumer, for running applications. Some examples are Dropbox, Office365, and Salesforce. This course focuses on the IaaS models.

The AWS Public Cloud

- Advantages:
  - Cost aligned with what you use
  - Automated deployments
  - Scalability
  - Security
  - Recovery
  - Reliability
- Well-suited for:
  - Data storage
  - Data archival
  - Application hosting
  - Latency-intolerant or mission-critical web tiers
  - On-demand hosting for microsite and application
  - Auto-scaling environment for large applications

FORTINET Training Institute

 5

The AWS public cloud provides services such as Amazon Elastic Compute Cloud (EC2), supplying infrastructure and services over the public internet. These infrastructure and services are available anywhere and anytime as needed. The AWS public cloud offers advantages such as low cost of ownership, automated deployments, scalability, security, recovery, and reliability.

AWS Service Names

Service category	Service	Amazon Web Services
Compute	IaaS	Amazon Elastic Compute Cloud (EC2)
	PaaS	AWS Elastic Beanstalk
	Containers	AWS Fargate, Amazon Elastic Compute Cloud Container Service
	Serverless functions	AWS Lambda
Network	Virtual networks	AWS VPC
	Load balancer	Elastic Load Balancer (ELB)
	Dedicated interconnect/peering	Direct Connect
	DNS	Amazon Route 53
	CDN	Amazon CloudFront
Storage	Object storage	Amazon Simple Storage Service (S3)
	Block storage	Amazon Elastic Block Store
	File storage	Amazon Elastic File System
	Reduced-availability storage	Amazon S3 Reduced Redundancy Storage

FORTINET Training Institute

 6

Vendor service names are vendor specific. As shown on this table, the VM is named differently for each vendor. For example, the Amazon Web Services VM is named Amazon Elastic Compute Cloud (EC2). There are also different names for DNS, for example, Amazon Route 53. The content delivery network name is also based on the vendor, such as Amazon CloudFront.

AWS Service Names (Contd)

Service category	Service	Amazon Web Services
Database	RDBMS	Amazon Relational Database Service
	NoSQL: Key-value	Amazon DynamoDB
	NoSQL: Indexed	Amazon SimpleDB
Big Data & Analytics	Batch data processing	Amazon Elastic MapReduce
	Stream data processing	Amazon Kinesis
	Stream data ingest	Amazon Kinesis
	Analytics	Amazon Redshift
Application Services	Messaging	Amazon Simple Notification Service
Management Services	Monitoring	Amazon CloudWatch
	Logging	Amazon CloudWatch
	Deployment	AWS CloudFormation

FORTINET Training Institute

 7

This slide is the continuation of the vendor service names table.

Threats and Challenges of the Public Cloud

- Threats:
  - Cloud misconfigurations
  - Malware
  - Insecure interfaces/APIs
  - Exfiltration of sensitive data
  - Unauthorized access
- Challenges:
  - Who is responsible for security?
  - Which cloud architecture should be implemented?
  - Reaching regulatory compliance

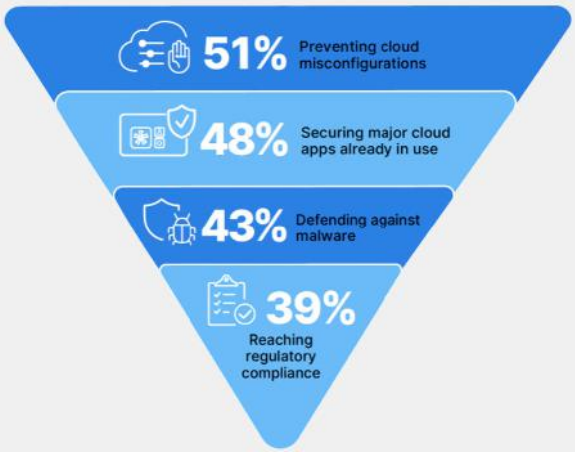

FORTINET Training Institute

 8

According to the 2023 Cloud Security Report conducted by Cybersecurity Insiders and sponsored by Fortinet, more than 51% of surveyed cybersecurity professionals deem misconfigurations of the public cloud as the number one threat. Malware, insecure interfaces and APIs, as well as exfiltration of sensitive data were also commonly reported as threats.

The top challenges of the public cloud are security responsibility, the type of cloud architecture that should be implemented, and reaching regulatory compliance.

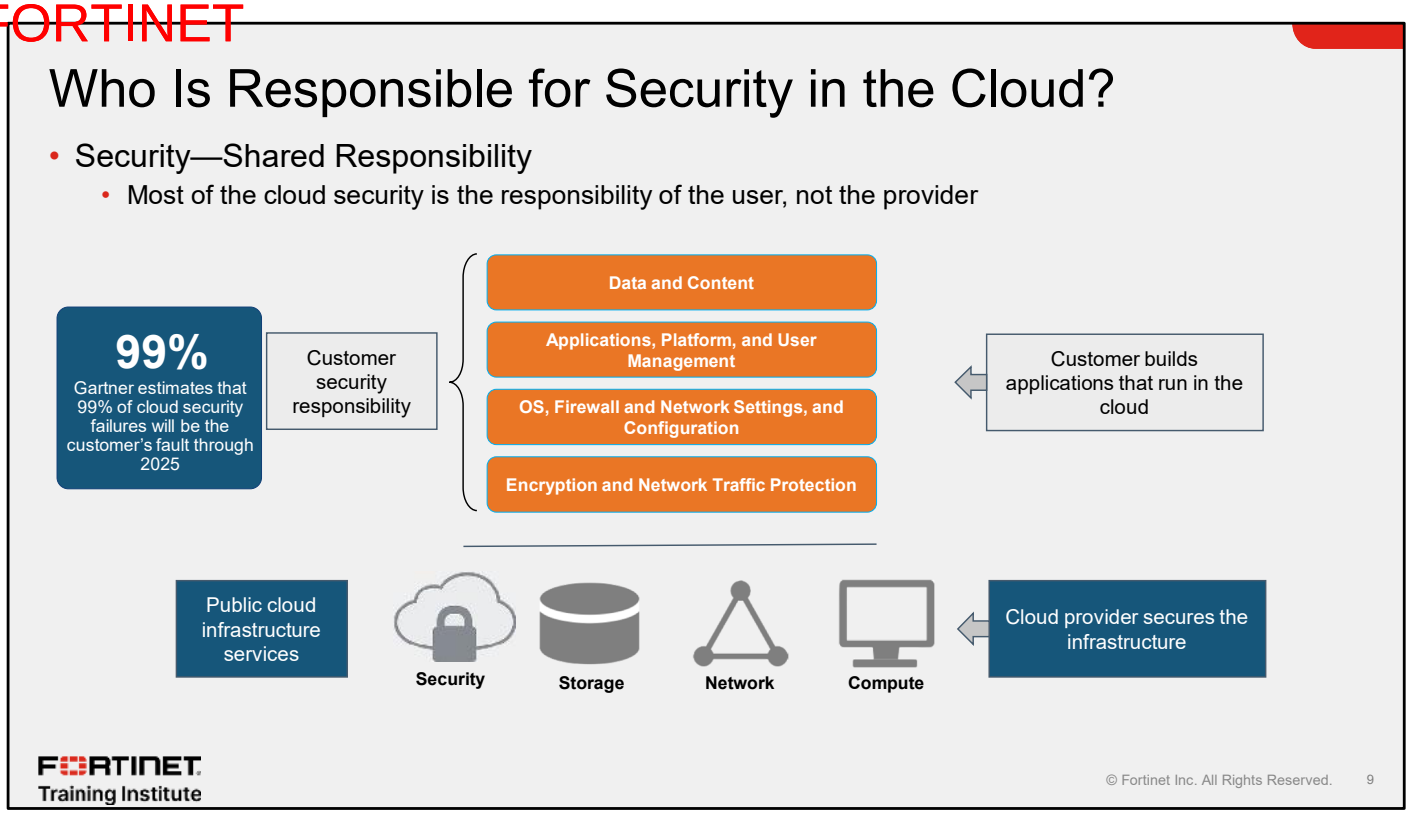

Security responsibility is arguably the number one top challenge in the public cloud. At a high level, the responsibility breaks down into two stacks: the lower stack and the upper stack. The lower stack includes the elements that are provided and, therefore, secured by the cloud service provider. Cloud *customers* are responsible for securing the remaining elements—customer network, applications, and data. The cloud security model is commonly broken down using the familiar OSI model; however, the OSI model doesn't represent the security responsibility breakdown. In some cases, cloud users build overlay networks on top of the cloud network, or layer additional services on top of existing infrastructure services. In cases like these, responsibility for the security of the modified infrastructure belongs to the customer. Essentially, if you manage it, you are responsible for it.

AWS Shared Responsibility Model

- AWS responsibility:
  - Protects infrastructure running all services offered in AWS
  - Infrastructure is composed of:
    - Hardware
    - Software
    - Cloud networking
    - Facilities
- Customer responsibility:
  - Determined by the AWS Cloud services selected
    - Amazon EC2
    - Amazon S3

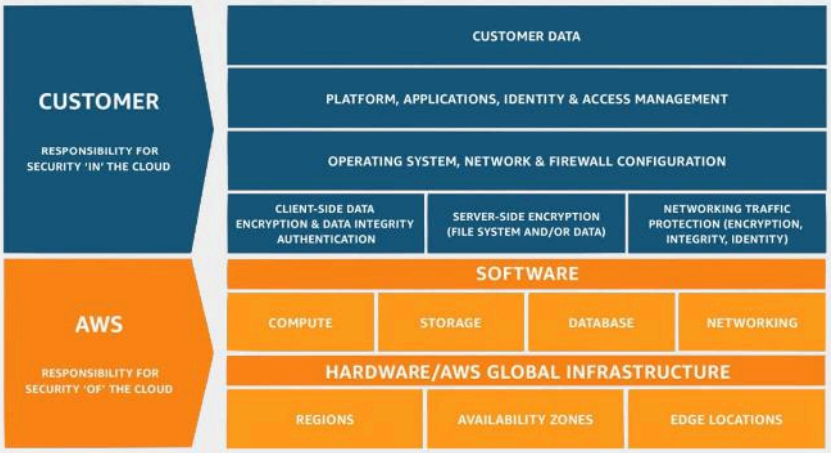

FORTINET Training Institute

 10

For AWS specifically, the shared responsibility model is composed of AWS and customer responsibilities.

The shared responsibility of AWS is composed of the infrastructure that runs all the services offered in the AWS cloud. This includes hardware, software, cloud networking, and facilities that run the AWS cloud services.

The customer's shared responsibility depends on the AWS cloud service chosen. For example, Amazon EC2, which is categorized as IaaS, requires the customer to perform all necessary security configuration and management tasks. This includes:

- Management of the guest operating system
- Application software or utilities installed on the instance
- Configuration of security groups
- Customer overlay network

On the other hand, services such as Amazon S3, require the customer to be responsible for managing their data, classifying their assets, and using identity and access management (IAM) tools to apply appropriate permissions.

In summary, a customer is responsible for security within the cloud, and AWS is responsible for the security of the cloud.

Which Cloud to Choose?

Type	Owner	End Users	Connectivity	Hardware	Cost
Public	Cloud provider	Multiple	Through the internet	Shared	Lower
Private	Organization	Within organization	Through a private network	Dedicated	Higher
Hybrid	Mixed	Mixed	Mixed	Mixed	Highest

FORTINET Training Institute

 11

The first issue to consider as you look toward the cloud is which architectural approach you want to take in adopting cloud services. The models of cloud architecture are public, private, and hybrid. Now, you will examine briefly each of the cloud models.

- **Public cloud**: Public clouds are available to any organization, and a variety of well-known vendors including Amazon, Microsoft, Google, Oracle, and Alibaba provide these public cloud environments.
- **Private cloud**: As the name suggests, private clouds are designed to be visible only to the organization that creates them. Private clouds provide many of the same benefits that a public cloud does, and still allow you to maintain ownership of the data and equipment. A private cloud is essentially a private data center that an organization creates with stacks of servers all running virtual environments, providing a consolidated, efficient platform on which to run applications and store data.
- **Hybrid cloud**: The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (for example, cloud bursting for load balancing between clouds).

Knowledge Check

1. What is the responsibility of the customer in the shared responsibility model?
   - A. Networking
   - ✓ B. Identity and access management (IAM)

2. What is an advantage of the public cloud over traditional on-premises networking?
   - A. Full traffic visibility
   - ✓ B. Scalability

FORTINET Training Institute

 12

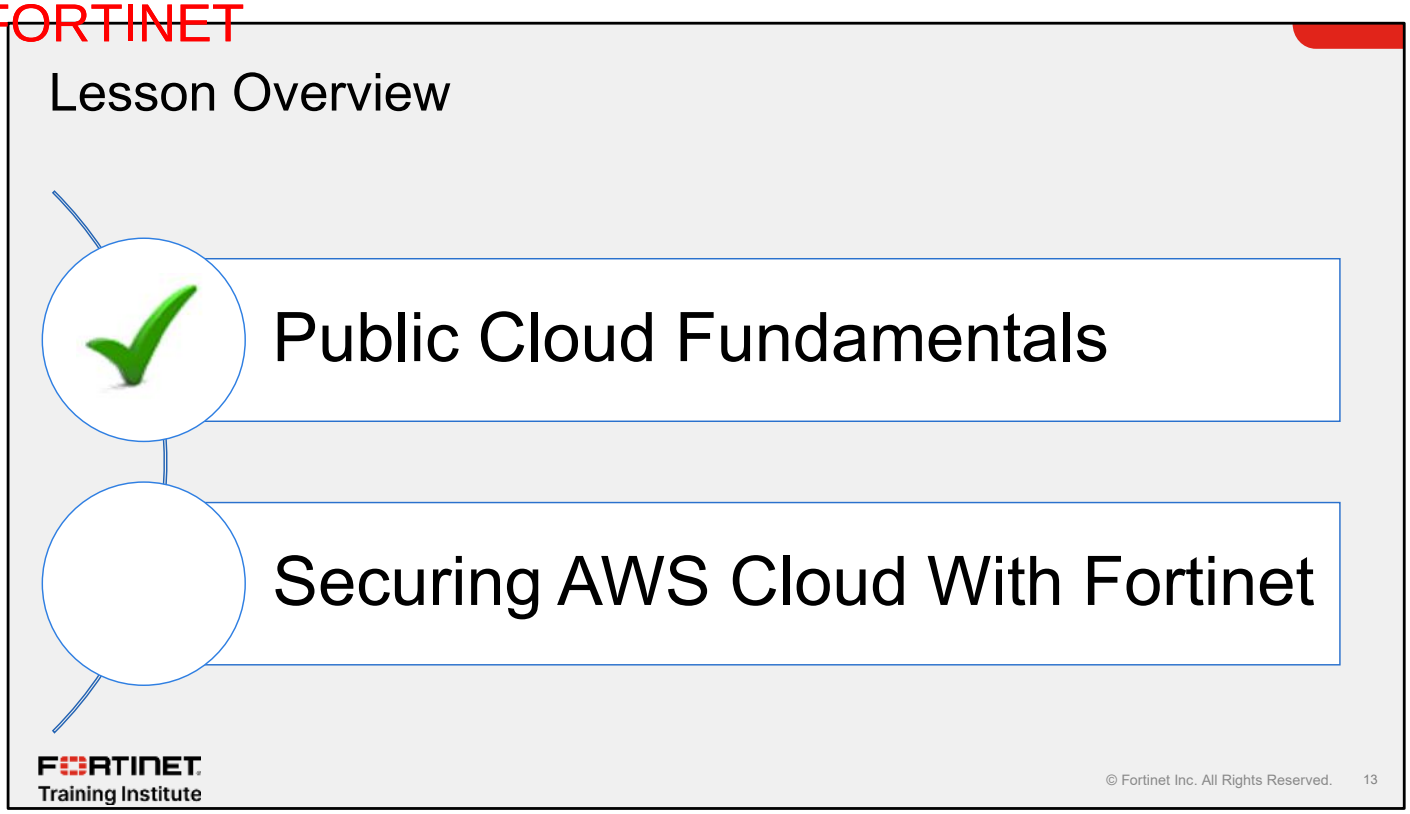

Good job! You now know about public cloud fundamentals.

Now, you will learn how to secure the AWS cloud with Fortinet.

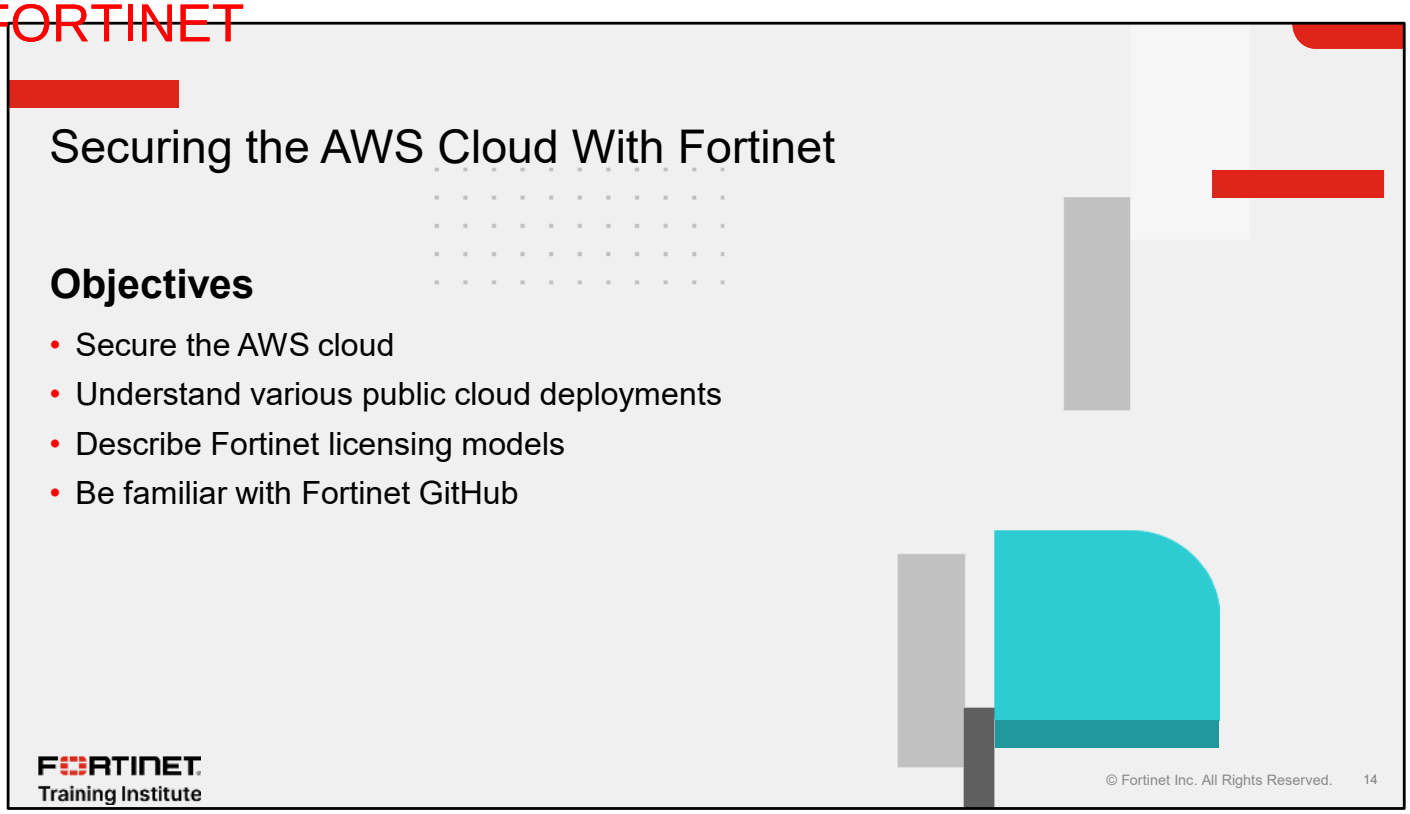

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding Fortinet solutions for the AWS cloud, you will be able to secure your cloud network using Fortinet solutions.

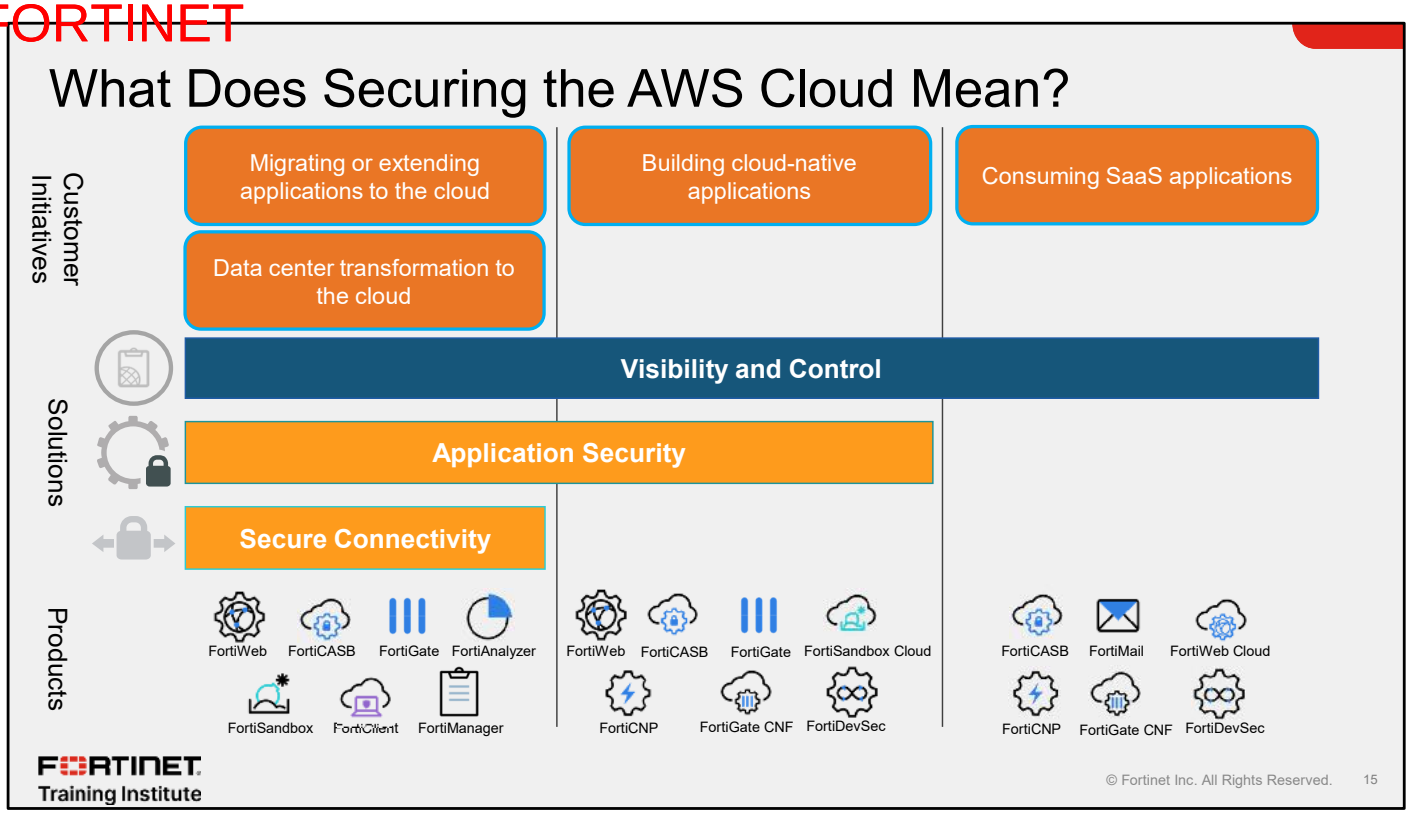

As the leader in multicloud security, Fortinet gives you the confidence to deploy any application in any cloud. Fortinet solutions provide broad protection across the entire digital attack surface, both on-premises and in public clouds, such as AWS. Native integration with each of the major cloud providers enables automated, centralized management across all clouds uniformly and seamlessly. This gives you unified visibility and control and policy management that supports risk management and compliance requirements.

There are three Fortinet solutions for securing the AWS cloud: the secure connectivity solution, which belongs to the category IaaS; application security; and visibility and control. Fortinet provides solutions for each of these categories. For example, Fortinet can provide secure connectivity for IaaS, but cannot provide the same solution for SaaS applications. So, for SaaS, Fortinet can provide only visibility and control. In other words, you cannot create an IPsec tunnel or web application firewall (WAF) to a dropbox (SaaS).

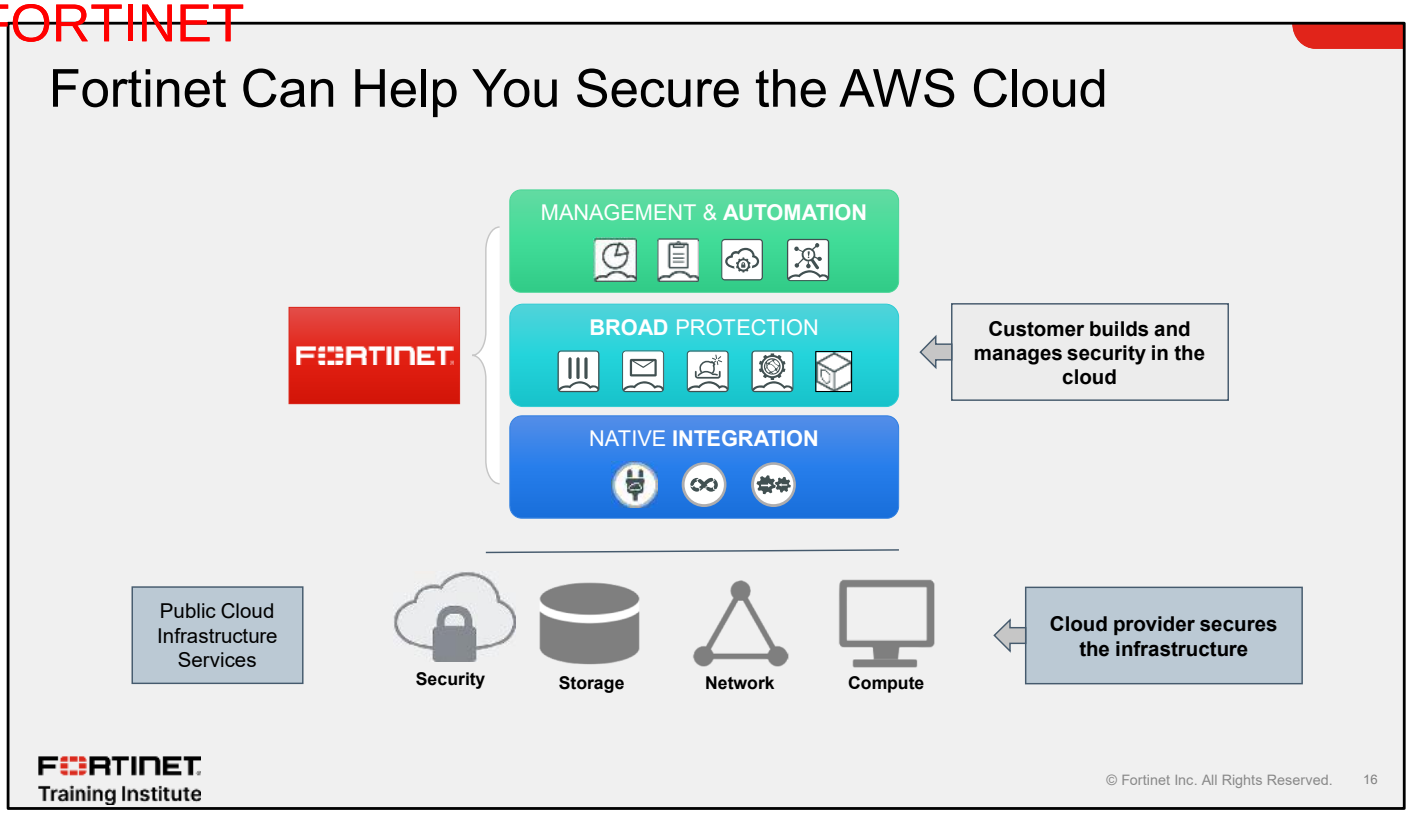

As shown on this slide, Fortinet can provide different products to secure the AWS cloud.
Management & Automation: In order to make the best use of their often limited and overstretched security personnel, Fortinet provides customers with a unique single-pane-of-glass solution that empowers them to consistently manage the broad set of protection services that is natively integrated into the cloud infrastructure. This approach also provides the ability to automate the management of these capabilities by using standard web-based APIs, as well as consume predefined automation recipes. By extending this automation framework across multiple cloud environments, customers can integrate the consumption of security services into their emerging DevOps-oriented application lifecycles, while supporting a more agile application and business operation.

Broad Protection: Offering the broadest set of security products both in and out of the cloud allows customers to consistently build the most secure infrastructures possible, regardless of deployment mode, workflow complexity, or degree of distribution and elasticity. The ability to natively integrate with the cloud infrastructure allows Fortinet to uniquely offer multiple security products in—and between—the cloud environments offered by every major cloud service provider. This helps customers build consumable and automation-ready security services to protect their cloud applications, regardless of where they choose to deploy them.

Native Integration: Integration seamlessly extends consistent security across the platforms of every major cloud provider, enabling organizations to define security similarly across their multicloud and on-premises deployments. Likewise, native integration provides the ability to natively consume cloud services by security products, providing faster and more seamless protection and response, and extends the web service-based APIs of products that are running in the cloud.

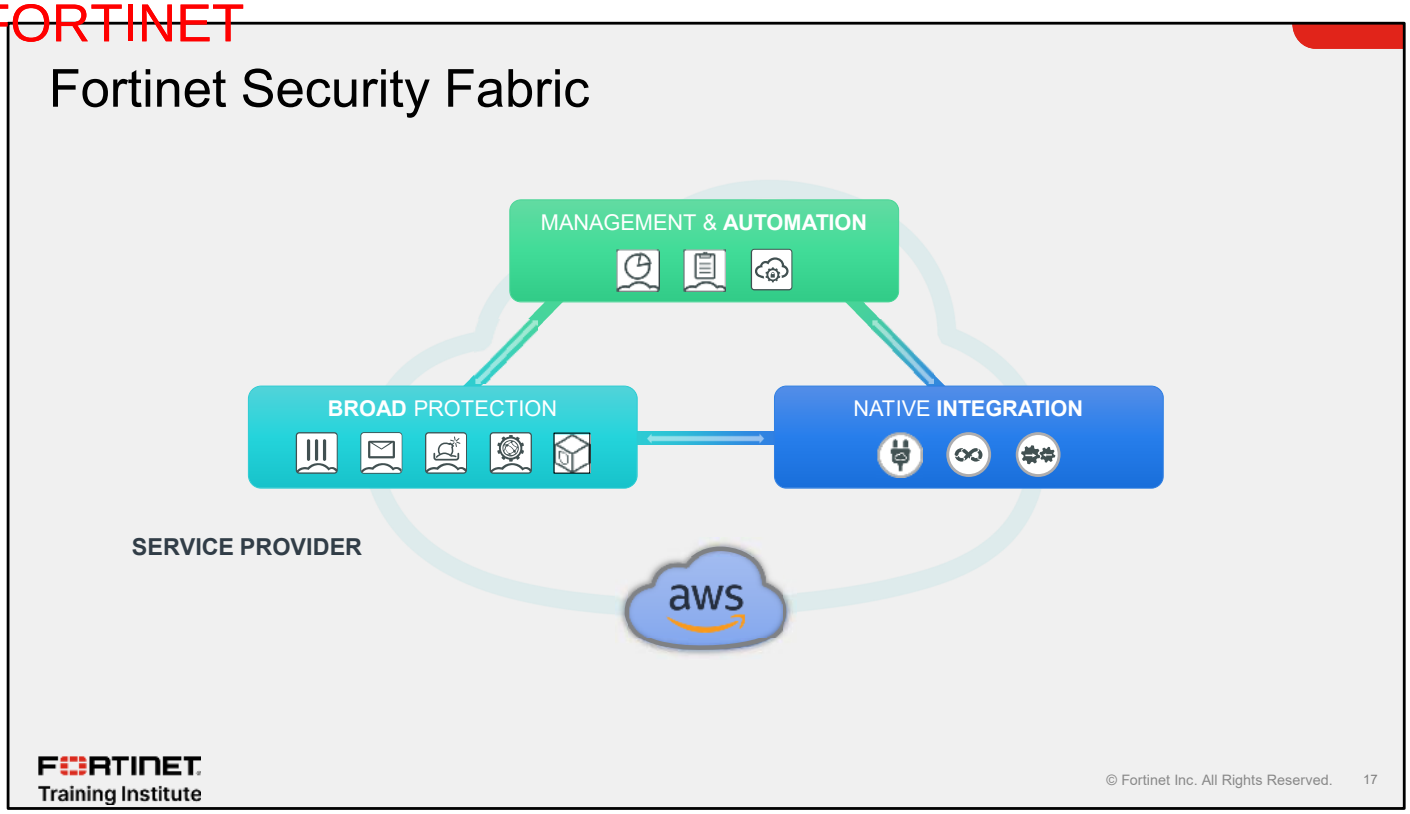

This slide shows the Fortinet Security Fabric overlaid onto the multicloud reality that was previously outlined. The key pillars are integration, protection, and management. As part of the Fortinet Security Fabric, FortiManager and FortiAnalyzer provide automation-ready, single-pane-of-glass management, transparent visibility, advanced compliance reporting, and network-aware rapid response across on-premises, cloud, and hybrid environments.

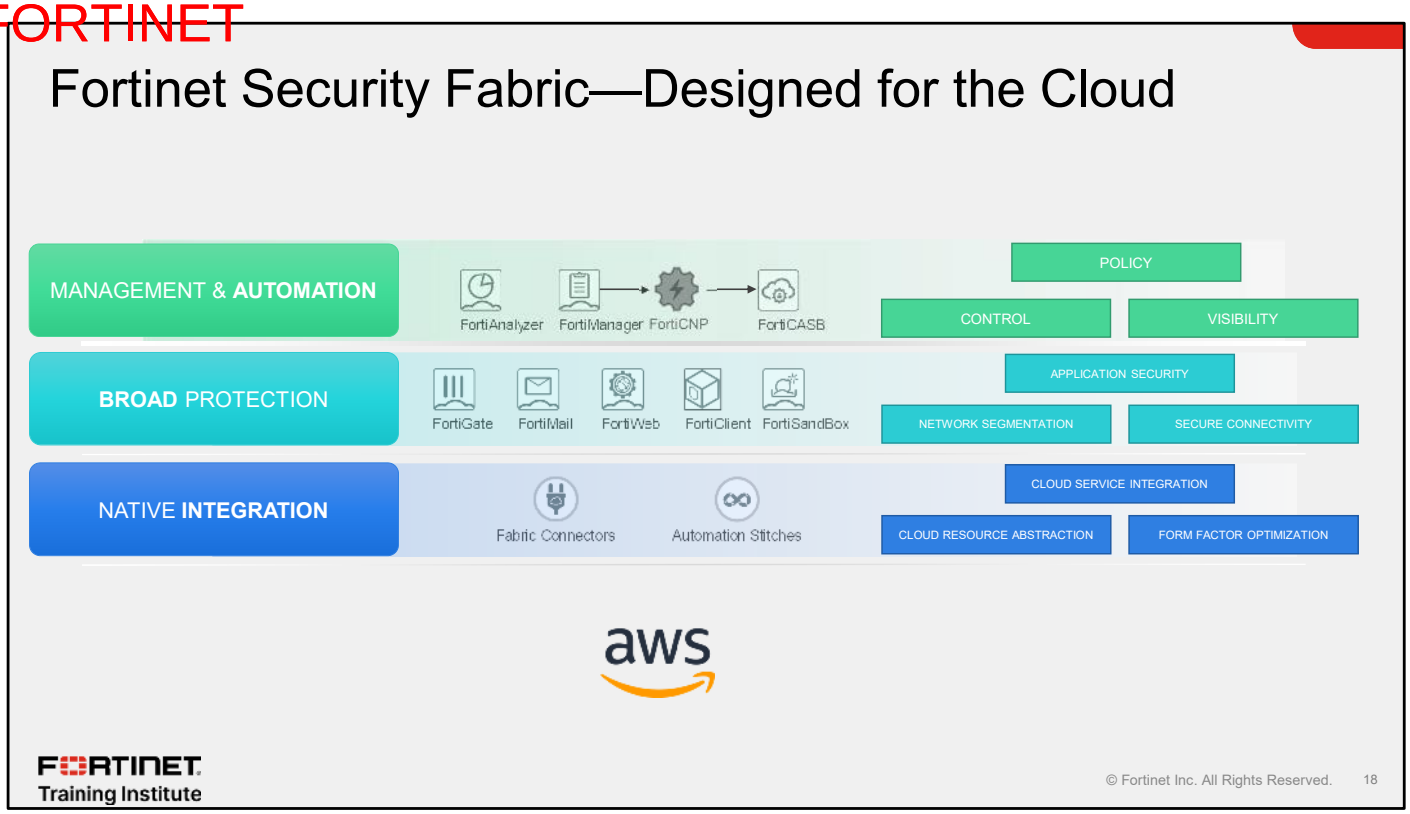

This slide shows the three pillars of the Fortinet Security Fabric for the AWS cloud, and the services and capabilities each pillar enables. Fortinet is investing in each of these pillars to provide native integration and capabilities across clouds.

The Fortinet Security Fabric enables the following services and capabilities:

- Seamless integration of separate cloud infrastructures, and use of native cloud services
- Broad protection for each product, regardless of cloud platform—effectively running virtual versions of the enforcement products on each cloud
- Management products that interact with, and manage the security of, the Fortinet products that run on each cloud

Fabric Connectors

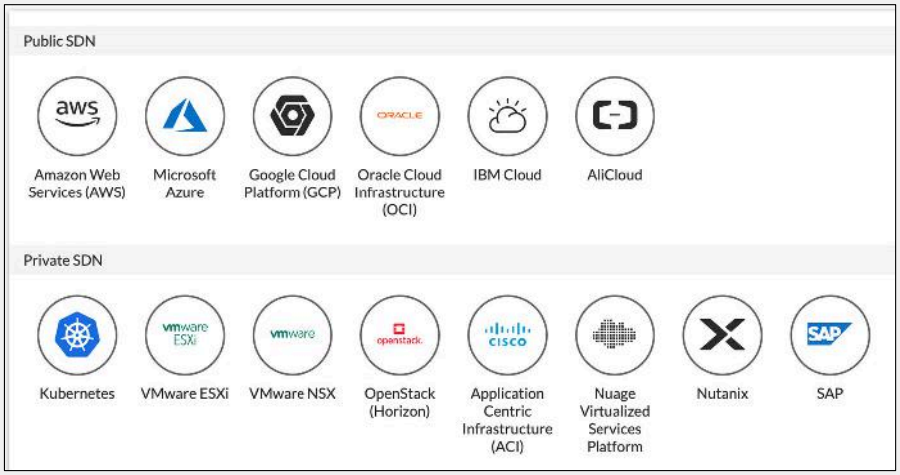

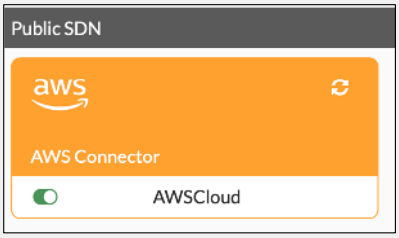

FORTINET Training Institute

 19

To address the complexities of today's digital enterprise and help reduce security gaps, Fortinet expands the openness of its Security Fabric architecture through its fabric connectors to extend security visibility and management capabilities deeper into fabric-ready partner infrastructure and applications.

Fabric connectors help customers maintain a consistent network security posture with centralized orchestration for users, applications, and data across hybrid, public, and private cloud environments. They enable automation of workflows, SOC environments, threat feeds, and security policy automation across clouds as new services and applications are deployed, removing the need for manual intervention.

Fabric connectors link into partner solutions through API integration points or through specialized engineering, and are instantly accessible to customers through easy, downloadable DevOps kits with one-click activation. The open design of the fabric connectors enables ongoing, deep integration with a growing number of ecosystem components and extends the Security Fabric capabilities into validated, third-party infrastructure.

Born in the Cloud

- No physical on-premises IT equipment
- Cost benefits with hourly billing
- Compliance regulations and policy enforcement
- Full enterprise-class security
- Broad portfolio of solutions

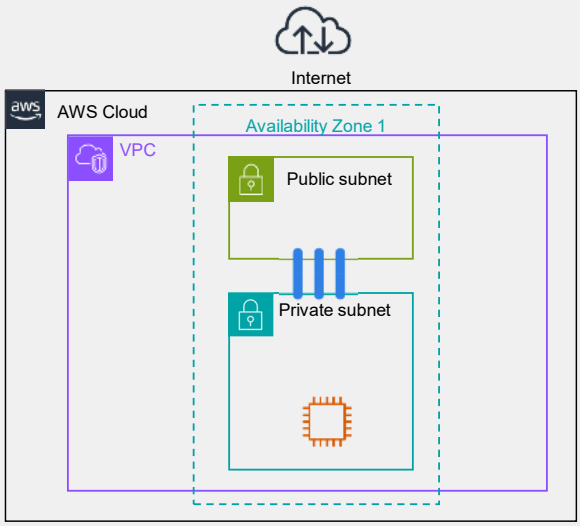

FORTINET Training Institute

 20

An interesting aspect of the Fortinet solution is that the customer can run all devices in the AWS cloud. There is no need for the customer to run physical devices on-premises. Unlike other vendors, Fortinet can offer all security products in cloud-based form, for example, FortiGate, FortiWeb, and FortiGate CNF.

Depicted on this slide is a single FortiGate-VM deployment. You will deploy a similar scenario later in the labs. The FortiGate can either be deployed as IaaS (FortiGate-VM) or SaaS (FortiGate CNF).

Extending to the AWS Cloud

- Customer on-premises infrastructure extends to the cloud through the VPN
  - From on-premises FortiGate to FortiGate-VM in the cloud

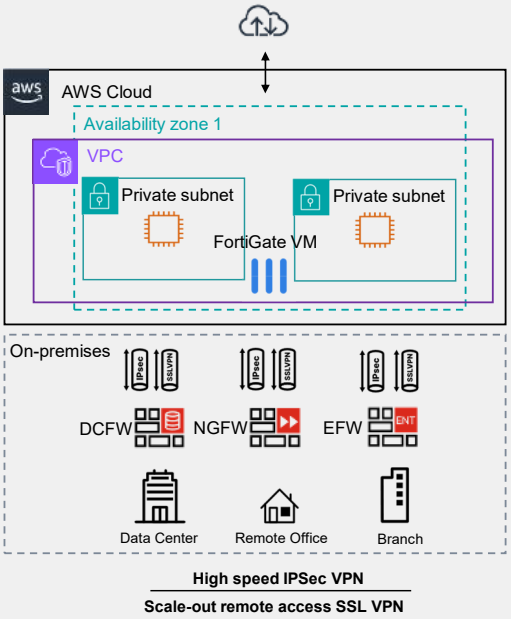

FORTINET Training Institute

 21

The customer can extend the on-premises infrastructure to the AWS cloud through the VPN. As shown on this slide, the customer can run an IPsec tunnel between the AWS cloud and the FortiGate on-premises infrastructure. On the cloud side, you can deploy a cloud vendor–native IPsec service, which is not recommended, or you can depoly a virtual FortiGate, which is highly recommended.

Fortinet Cloud Security Solution

- Extends to physical, virtual, and cloud devices with advanced security orchestration and unified threat protection
- Provides more control and visibility by identifying and setting policy by user applications, device specifications, IP addresses, and network interfaces
- Delivers a highly optimized solution that protects application workloads beyond native cloud vendor security options

FORTINET Training Institute

 22

It is important to know that the Fortinet cloud security solution is not a replacement for the existing cloud vendor security. It is just an extra layer of security in addition to the cloud vendor security solutions. The Fortinet cloud security solution provides more control and visibility and delivers a highly optimized security solution beyond native cloud vendor security options.

Challenges With Current Multicloud Deployments

- Connecting multiple public cloud networks through data center WAN edge for centralized security and routing
  - Deployment complexity
  - Degrades application performance
  - High connection costs

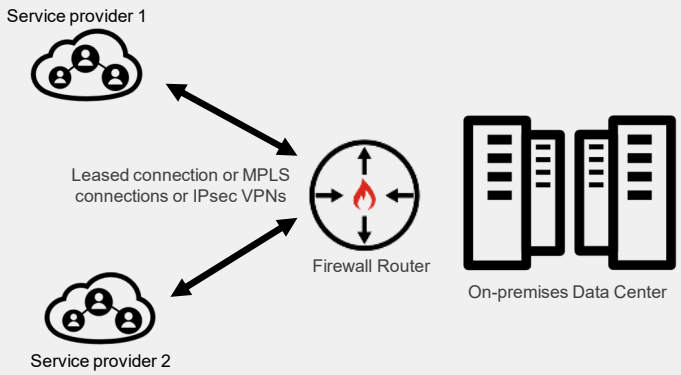

FORTINET Training Institute

 23

Managing and securing an assortment of different cloud platforms remains a challenge. Few IT teams have the expertise to manage a mixed deployment of multiple public cloud, private cloud, and on-premises environments—especially considering the ongoing lack of skilled IT and cybersecurity talent.

To address the diversity challenge, many organizations choose to connect their clouds through their on-premises data center WAN edge for centralized inspection and routing. But the use of this type of traditional WAN infrastructure approach, though secure, inhibits agility and results in deployment complexity, inconsistent network performance, and expensive connectivity.

Fortinet Secure SD-WAN for Multicloud Solution

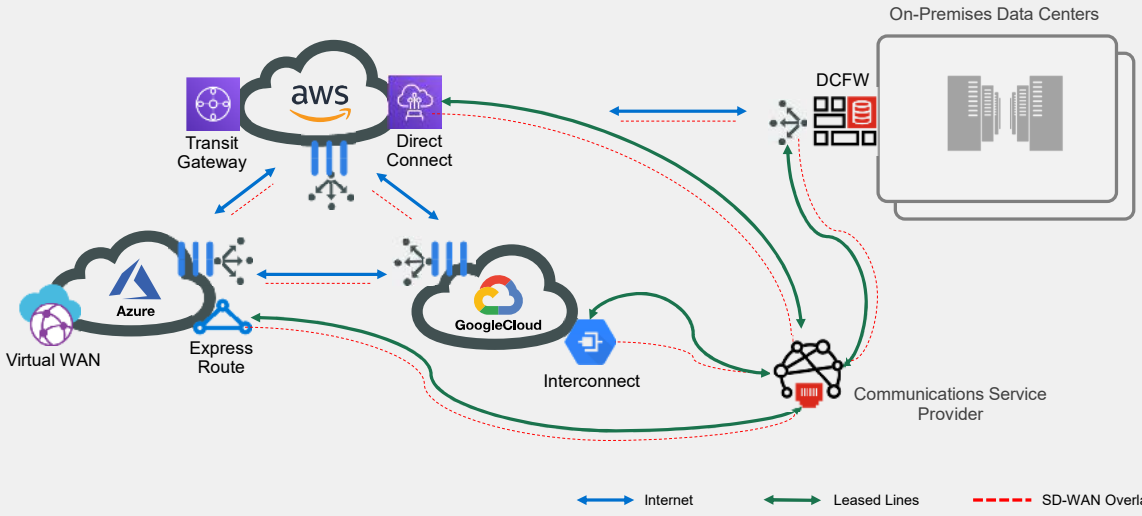

FORTINET Training Institute

 24

The Fortinet Secure SD-WAN for multicloud solution is a new approach to establishing secure and high-performance connectivity between IaaS workloads running on multiple clouds—without increasing cost and complexity. This solution enables SD-WAN between clouds and empowers enterprise IT to build a seamless cloud-to-cloud network and security architecture. The Fortinet secure SD-WAN solution delivers the following capabilities:

- Automates the deployment of a seamless overlay network across different cloud networks, reducing complexity and increasing agility to save teams time and resources.
- Offers visibility, control, and centralized management that unifies functionality across multiple cloud environments through Fortinet Security Fabric software-defined networking (SDN) connectors and cloud-native integrations.
- Securely transports cloud traffic between clouds without needing to backhaul through the data center, enabling better scaling of deployments and reducing latency.
- Intelligently selects connections based on cloud application and workload awareness, improving performance and reducing dependence on costly leased lines or MPLS connections.

FortiCloud

- Provides single point of access to all your Fortinet cloud portals
- View, manage, and access all Fortinet products and services in one place

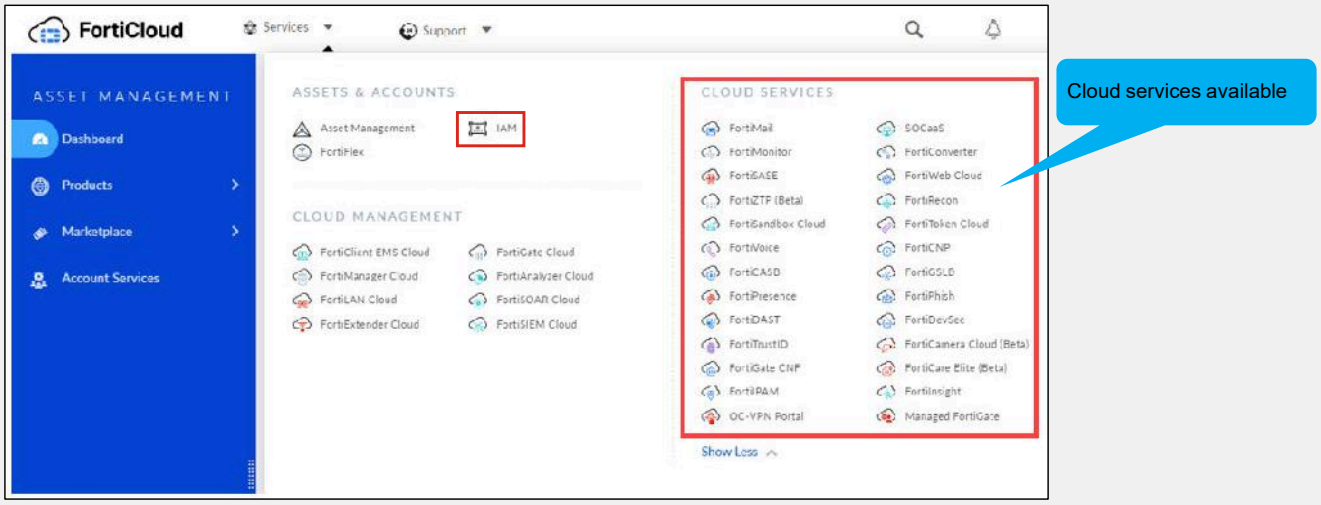

FORTINET Training Institute

 25

FortiCloud is a centralized login portal, which serves as a unified authentication mechanism for accessing all your Fortinet products and services.
FortiCloud allows you to view, manage, and access your entire Fortinet portfolio in one convenient and centralized location including but not limited to:

- FortiGate CNF
- FortiCNP
- FortiDevSec
- FortiFlex
- FortiWeb Cloud

Using the IAM portal, you can create different users with different permissions to the various portals. For example, FortiCloud administrator A has access only to the FortiWeb Cloud portal, whereas FortiCloud administrator B has access only to the FortiGate CNF portal.

FortiCloud (Contd)

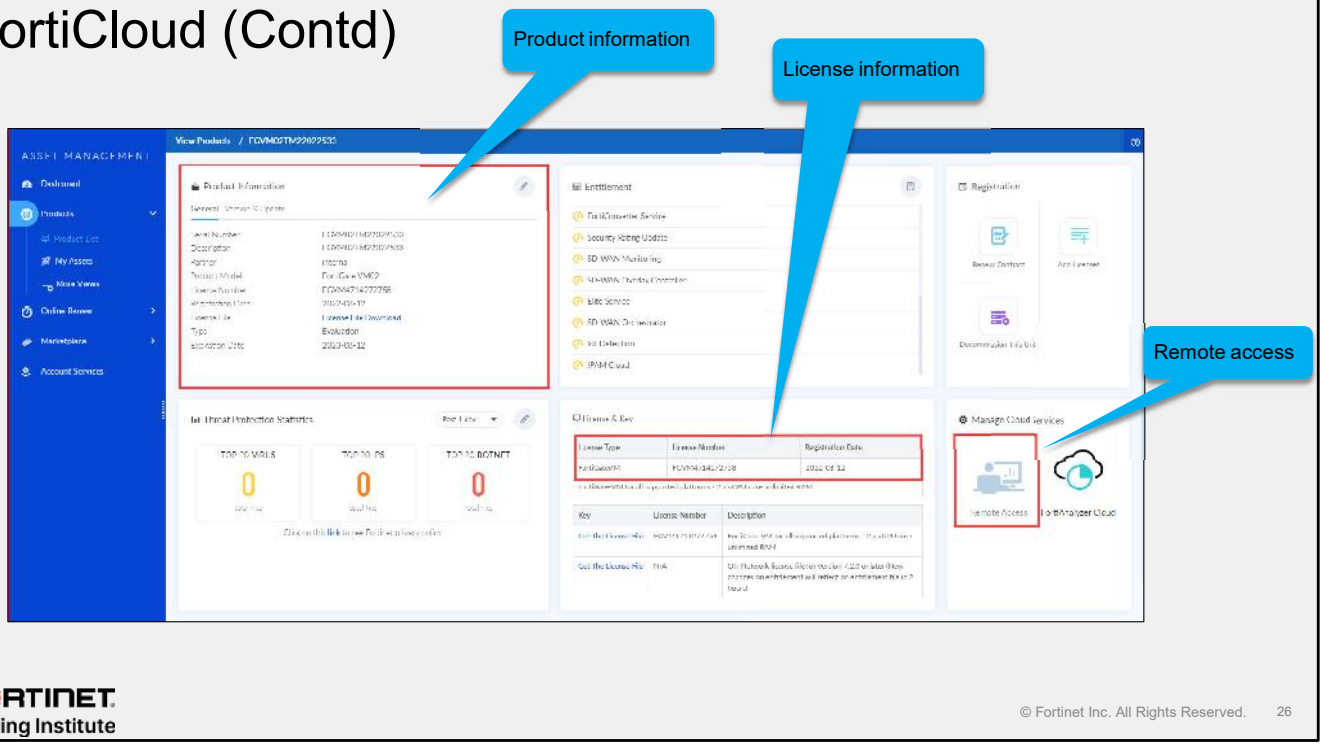

FORTINET Training Institute

 26

This slide shows an example of the information that can be collected and managed through FortiCloud. On this slide you can see a FortiGate VM showing its:

- Product information
- License information
- Remote access

Licensing Models

- Bring your own license (BYOL):
  - Just like FortiGate-VM but a different SKU
  - Acquired through partners
- Pay-as-you-go (PAYG) or on demand:
  - Hourly or yearly
  - Based on instance type
  - Paid through the cloud vendor

	aws	
BYOL	☑	☑
PAYG	☑	☑

- For both models, you must pay infrastructure running costs directly to the cloud vendor
- Switching between the two models requires redeployment

FORTINET Training Institute

 27

There are different Fortinet licensing models to select from, based on the customer requirements.

- Bring your own license: The customer pays the cloud vendor for the VMs and pays Fortinet for Fortinet products running 24/7 on the cloud. This model is recommended for VMs running all the time on the cloud. The customer gets Fortinet 24/7 support with the enterprise bundle.
- Pay-as-you-go or on demand: The customer is paying for both through the cloud vendor but pays for the service based on usage. The customer gets Fortinet 8x5 support with the UTM bundle.

In both cases, the customer must pay infrastructure running costs directly to the cloud vendor.

When switching between the two licensing models, redeployment of the VM or instances is required.

FortiFlex (FlexVM) Licensing Model

- Usage-based licensing for a wide range of Fortinet cybersecurity solutions across public cloud, hybrid cloud, and on-premises deployment
- Enterprise
  - Prepaid program
  - Points system
- MSSP
  - Postpaid program

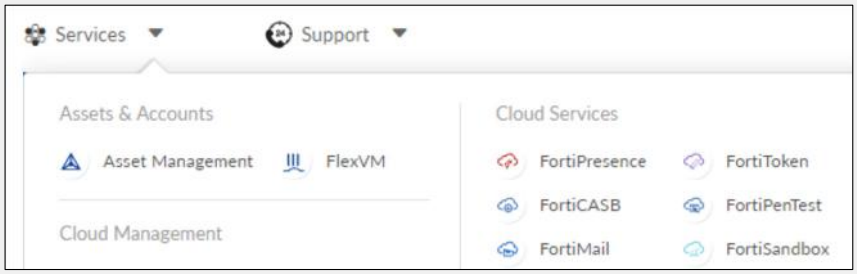

FORTINET Training Institute

 28

FortiFlex is similar to the pay-as-you-go licensing model. It differs in that customers are charged daily instead of hourly or yearly. As the name suggest it is also much more flexible. Once a Fortinet VM has been deployed, FortiFlex allows you to change the specifications on the fly. No need to deploy an entirely new VM or wait for a new license to be purchased if all that is required is an increase in RAM or CPUs.

Once a FortiFlex license has been acquired and registered on FortiCare, the FortiFlex portal can be accessed on FortiCloud.

FortiFlex offers two different programs:
- Prepaid: suitable for enterprises
- Postpaid: suitable for MSSPs

Resource consumption is based on predefined points that are calculated on a daily basis.
The points can be used to deploy any mix of cloud and VM offerings as well as services for on-premises deployments.
This allows organizations to leverage FortiFlex in various ways including:
- Reducing excessive procurement cycles for new security solutions
- Simplifying the deployment and provisioning of new services through the FortiFlex powerful APIs
- Maximizing budget and return on investment (ROI) by scaling down or pausing services as needed
- Optimizing cloud spend by using “use it or lose it” dollars committed to cloud providers to purchase FortiFlex points that organizations can redeem in the future

Note that if you change the configuration settings, the serial number of the VM changes as well. This means that you will need to create a new entitlement every time you change the settings.

FortiFlex Calculator

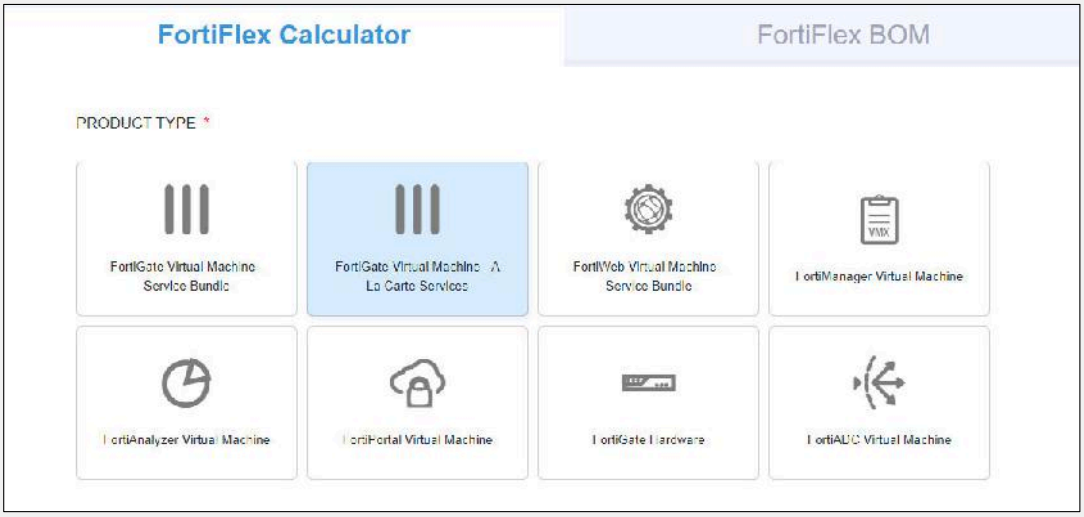

Flex-VM Calculator
https://fndn.fortinet.net/index.php?/tools/fortiflex/

FORTINET Training Institute

 29

Because this model charges users based on their use of cloud services, you can use the FortiFlex Calculator tool to calculate points consumption. Users pay for the resources consumed, such as number of VMs, CPU size, FortiGuard, and cloud services.

VMs and services charges are calculated on a daily basis. Administrators can change any resource at any time and calculation will take effect immediately. The daily charge is based on the largest configuration points of any given VM and associated services deployed on that day.

FortiFlex Calculator (Contd)

- Choose:
  - Product type
  - Number of VMs
  - CPU size
  - Service bundle
  - Number of virtual domains
- Point consumption updates dynamically
- Add to bill of materials (BOM)

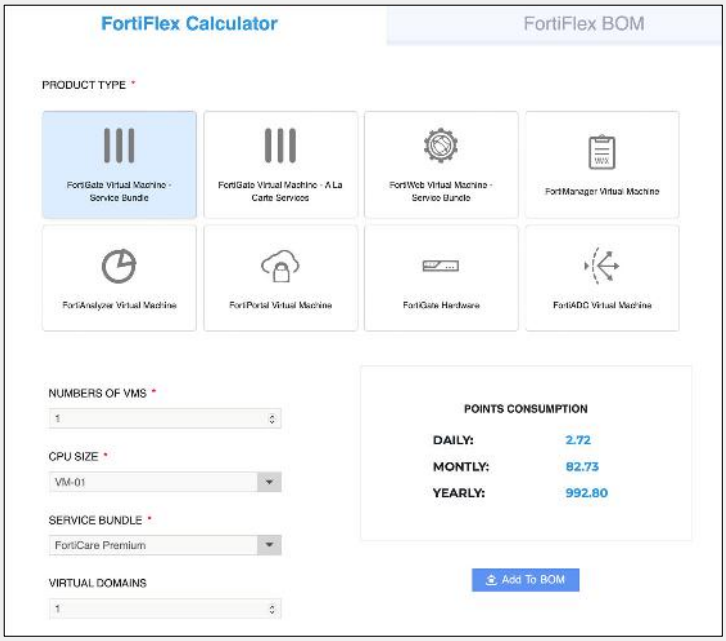

FORTINET Training Institute

 30

This slide shows a simple example of how to use the FortiFlex Calculator. In the example, a low-powered FortiGate VM with a FortiCare Premium service bundle is chosen. Any time a change is made to any parameters, the points consumption section is updated dynamically and immediately. Once done, you can click **Add To BOM**.

FortiFlex BOM

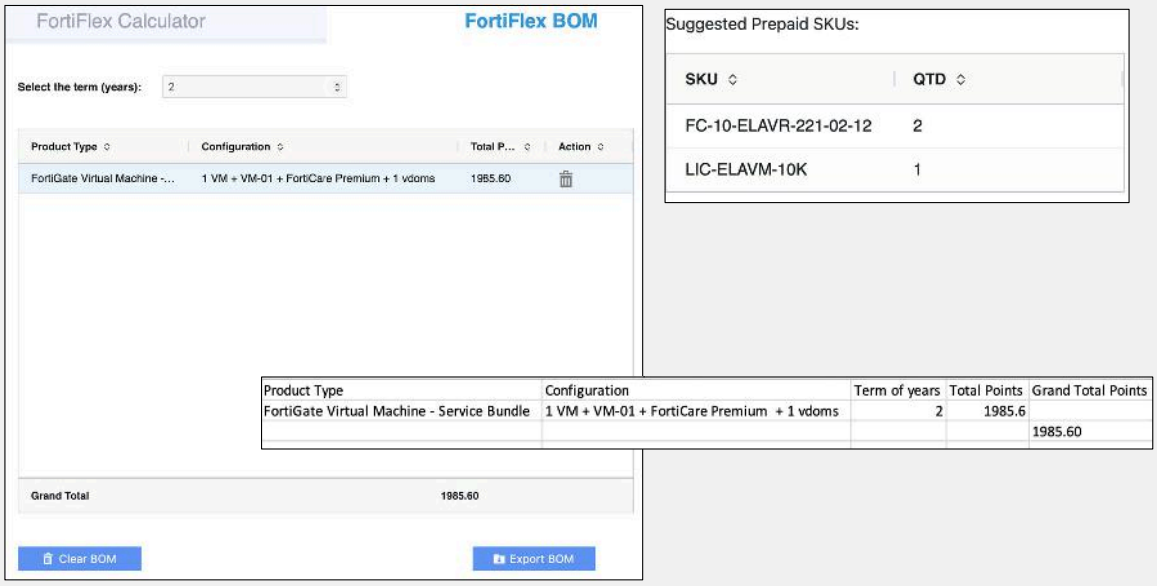

FORTINET Training Institute

 31

After the configuration is added to FortiFlex BOM, you can select the term (in years) for how long the services are required. The total points will update depending on the term chosen. FortiFlex BOM also provides the suggested SKU and the quantity. Finally, FortiFlex BOM provides an option to export the configuration as an Excel spreadsheet using the **Export BOM** button.

Fortinet GitHub

- Public cloud templates:
  - https://github.com/fortinet
- AWS cloud formation templates
- Terraform templates
- Azure templates

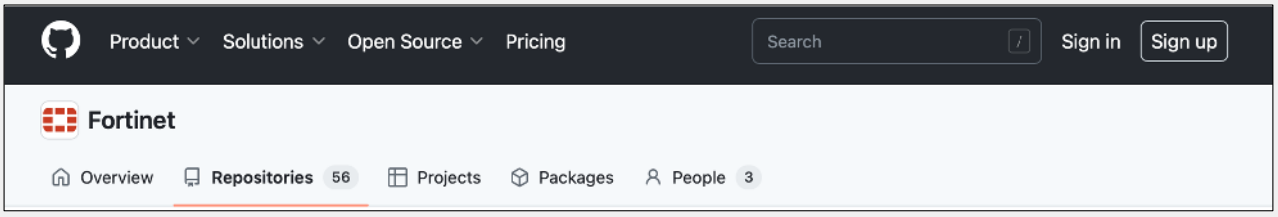

FORTINET Training Institute

 32

Fortinet GitHub is a Fortinet website where you can download various templates for your cloud security design. Some examples are AWS cloud formation templates, Terraform, and Azure templates that allow you to download preconfigured settings for the cloud security solutions. You can visit the official Fortinet GitHub at the website shown on this slide. However, during the lab you will be using a different GitHub, which is the Fortinet solution GitHub (developer GitHub).

Knowledge Check

1. What best describes FortiCloud?
   A. The Fortinet public cloud
   ✓B. Centralized portal to access Fortinet products and services

2. Which cybersecurity solutions can the FortiFlex licensing model be used for?
   A. Only public cloud
   ✓B. Public cloud, hybrid cloud, and on-premises

FORTINET Training Institute

 33

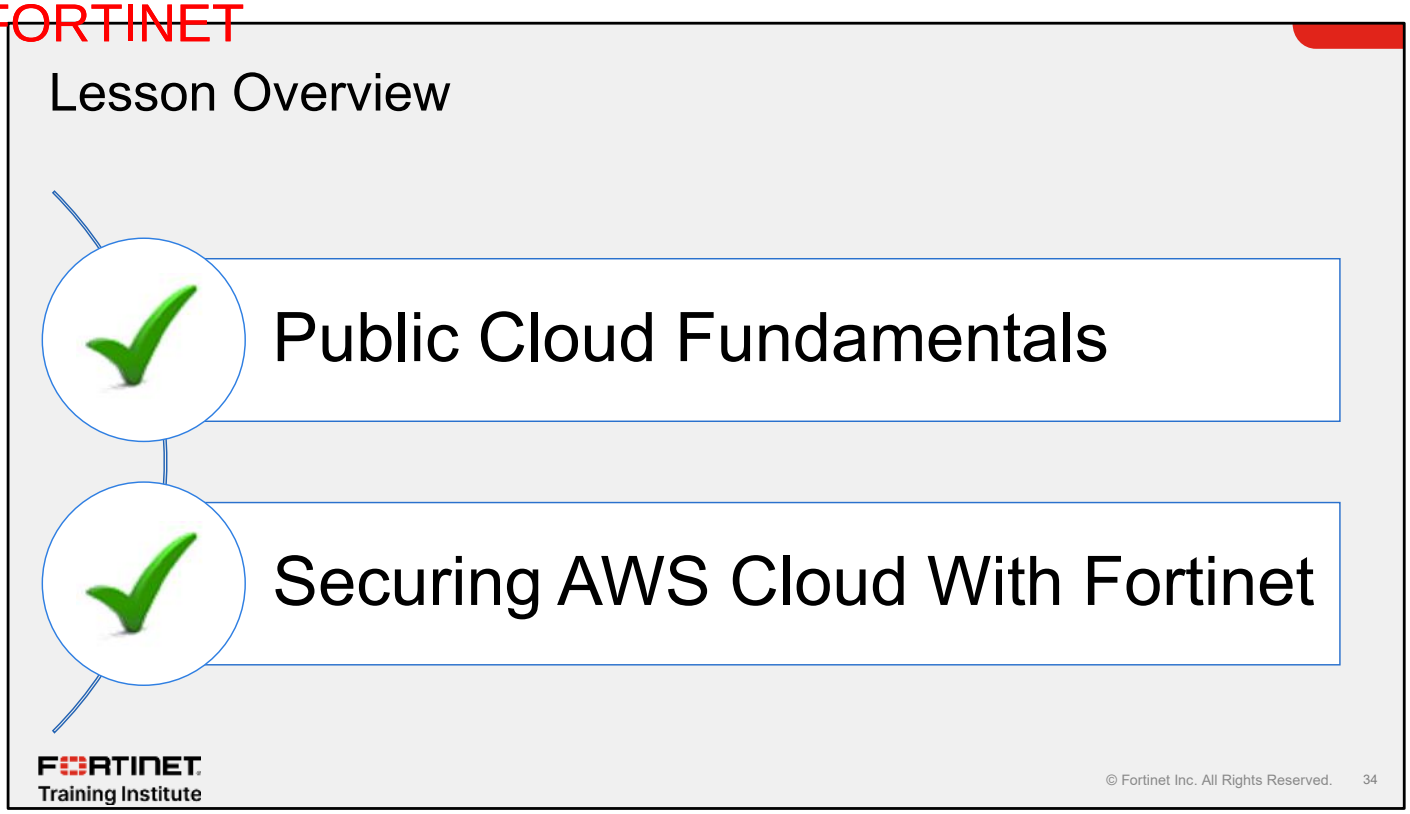

Congratulations! You have completed this lesson.

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

Review

- ✓ Understand the concept of the public cloud
- ✓ Know public cloud service terms
- ✓ Identify various threats and challenges in the public cloud
- ✓ Secure the AWS cloud
- ✓ Understand various public cloud deployments
- ✓ Describe licensing models
- ✓ Be familiar with Fortinet GitHub

FORTINET Training Institute

 35

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

By mastering the objectives covered in this lesson, you learned about the concept of the public cloud and how to use it in your network.

In this lesson, you will learn about Amazon Web Services (AWS) components.

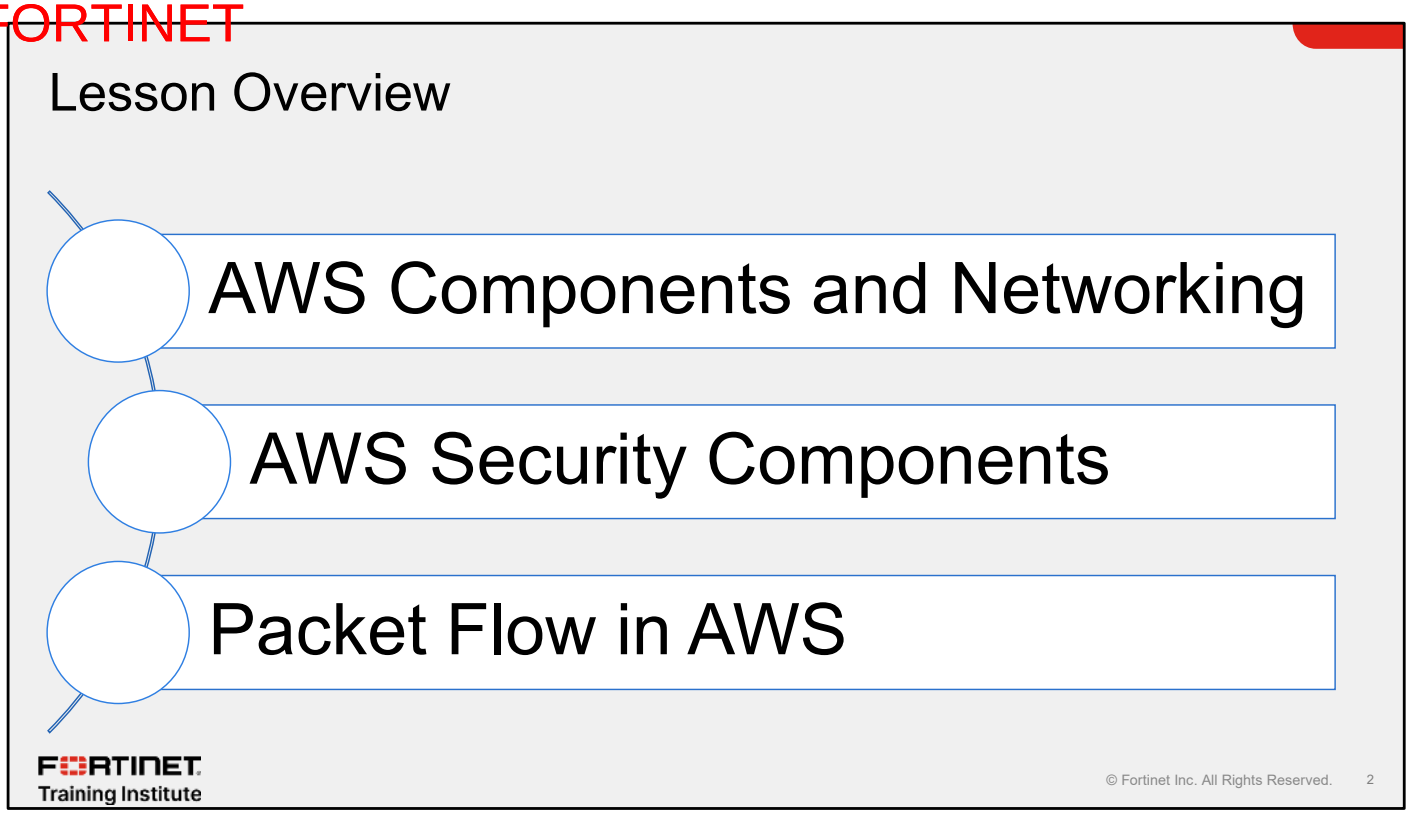

In this lesson, you will learn about the topics shown on this slide.

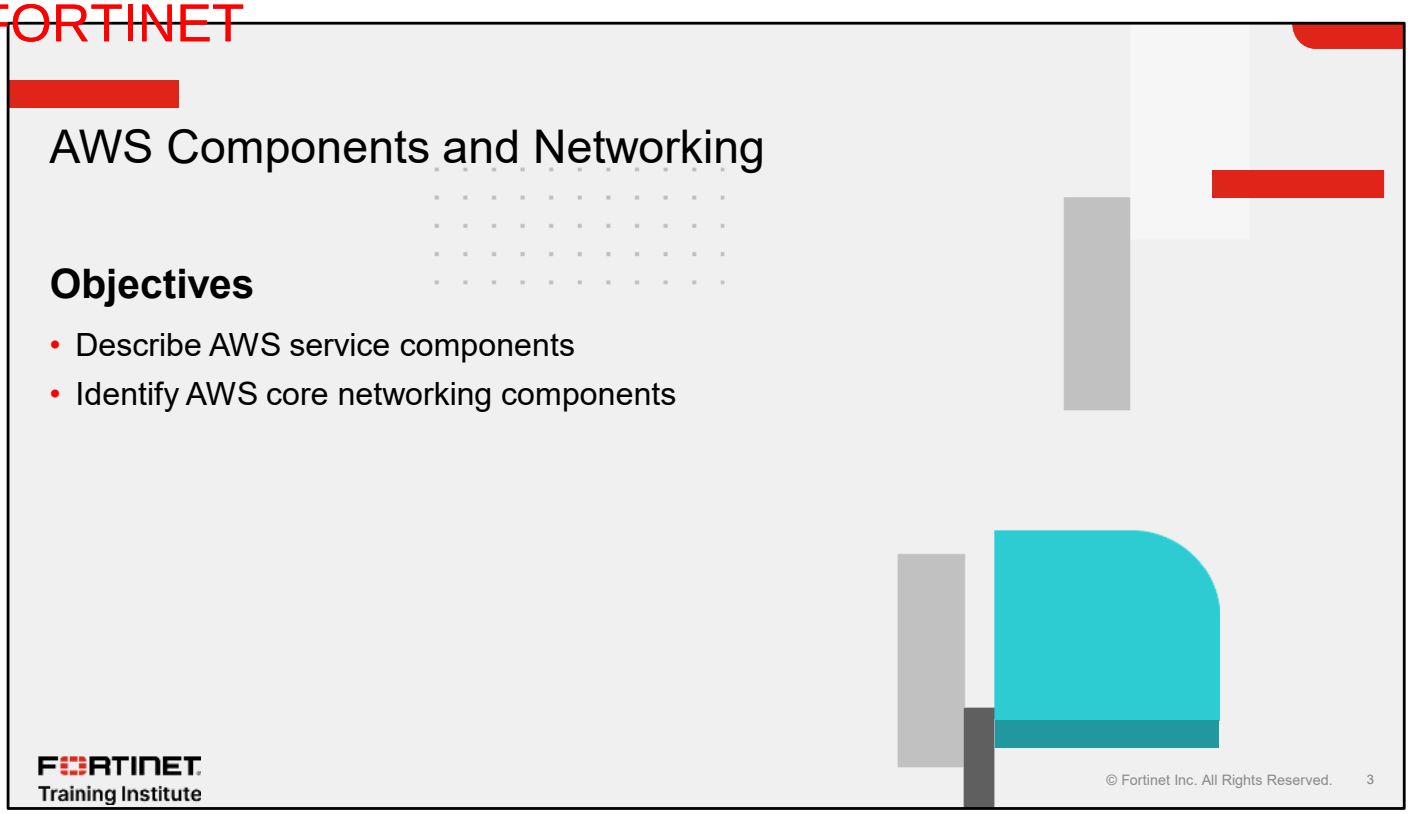

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding AWS components and networking, you will be able to use AWS successfully and efficiently to deploy your own cloud network.

AWS Regions

- AWS Cloud currently stretches across 102 availability zones (AZs) within 32 geographical regions across the globe

Source: Amazon Web Services, Inc, Global Infrastructure

FORTINET Training Institute

 4

AWS services and components are hosted in various locations worldwide. These locations are composed of regions and AZs. Each region is a separate geographic area with multiple, isolated locations known as AZs. When you view your resources, you'll see only the resources tied to the region you've specified. Regions are isolated from each other, and AWS does not automatically replicate resources across regions. There is a charge for data transfers between regions, but not all regions have the same features, functions, and offers.

Availability Zones (AZs)

- Availability Zones are isolated from each other
- AZs in a region are connected through low-latency links
- Amazon EC2 resources are either global, tied to a region, or tied to an AZ
- Number of AZs varies per region
- Minimum two AZs per region

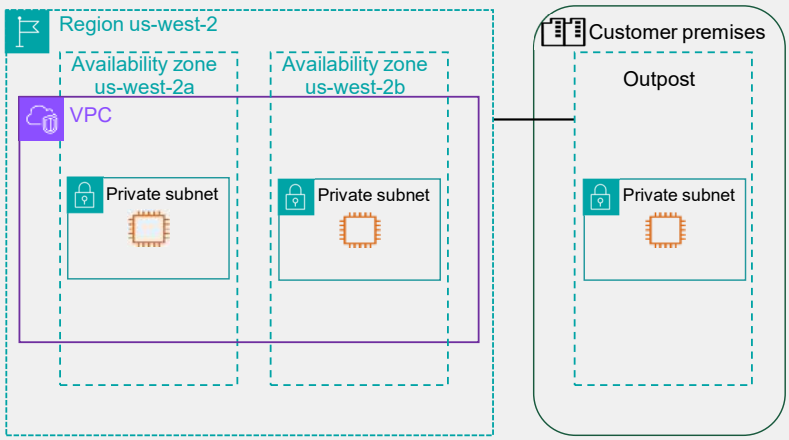

FORTINET Training Institute

 5

By launching your instances in separate AZs, you can protect your applications from a failure in a single location. Think of this redundancy as a physical hypervisor located in a different data center. If data center A fails, your workloads are redundantly deployed in data center B.

An AWS best practice is to place instances in more than one AZ. Each AZ is isolated, but the AZs in a region are connected through low-latency links.

For example, you can deploy two FortiGates in two different AZs to form an active-passive HA cluster. You will learn more about this in a later lesson.

You cannot have one FortiGate sitting between two AZs; instead, you can have a load balancer between AZs. These load balancers can be used to essentially form an active-active HA cluster.

An AZ is represented by a region code, followed by a letter identifier, for example, us-east-1a.

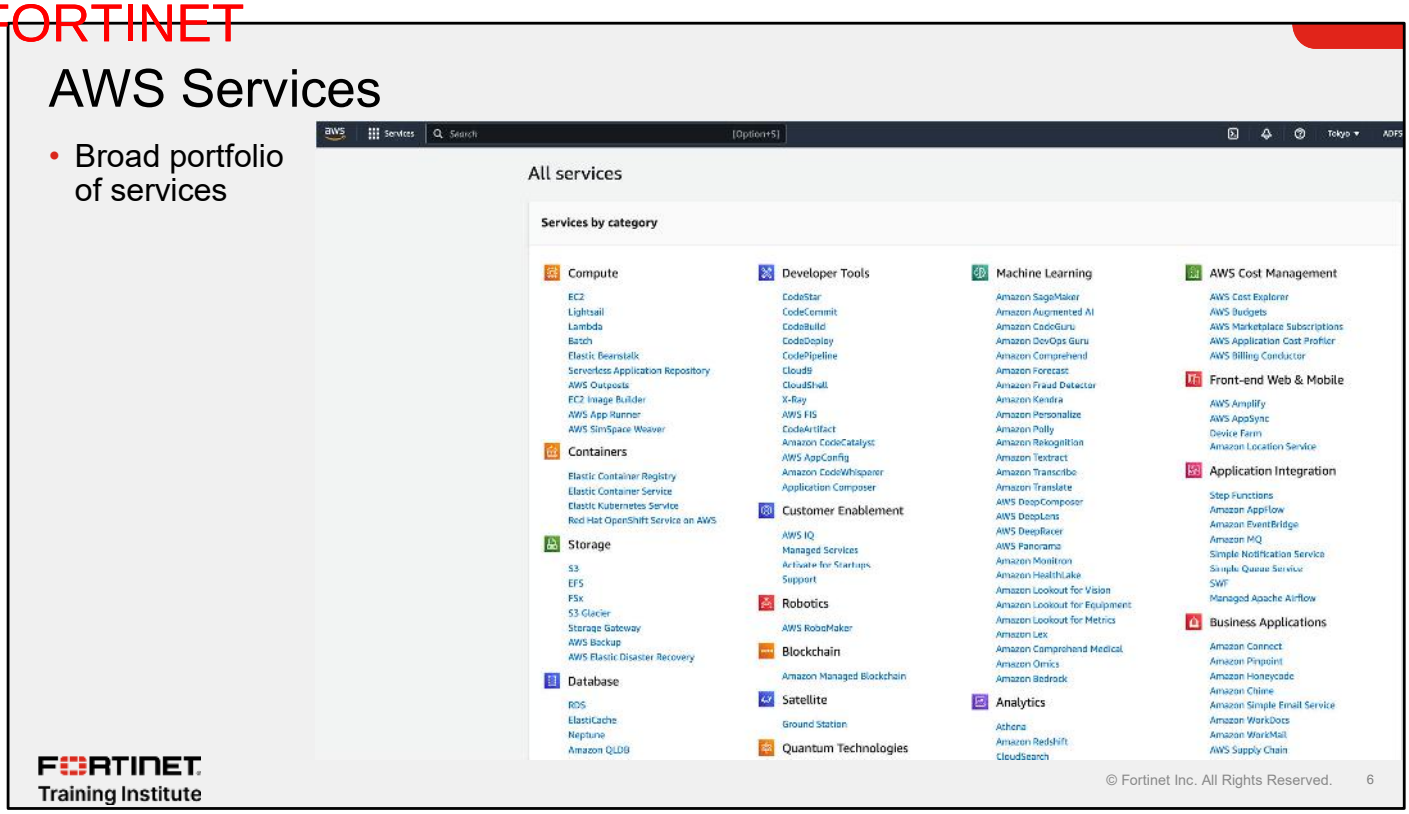

As shown on this slide, AWS has a broad portfolio of services. You will see all the available services when you click on **Services** > **All services** > **View all services** on the AWS console. However, in this lesson, you will focus on the main components including Amazon Elastic Cloud Compute (EC2) and Amazon Virtual Private Cloud (VPC).

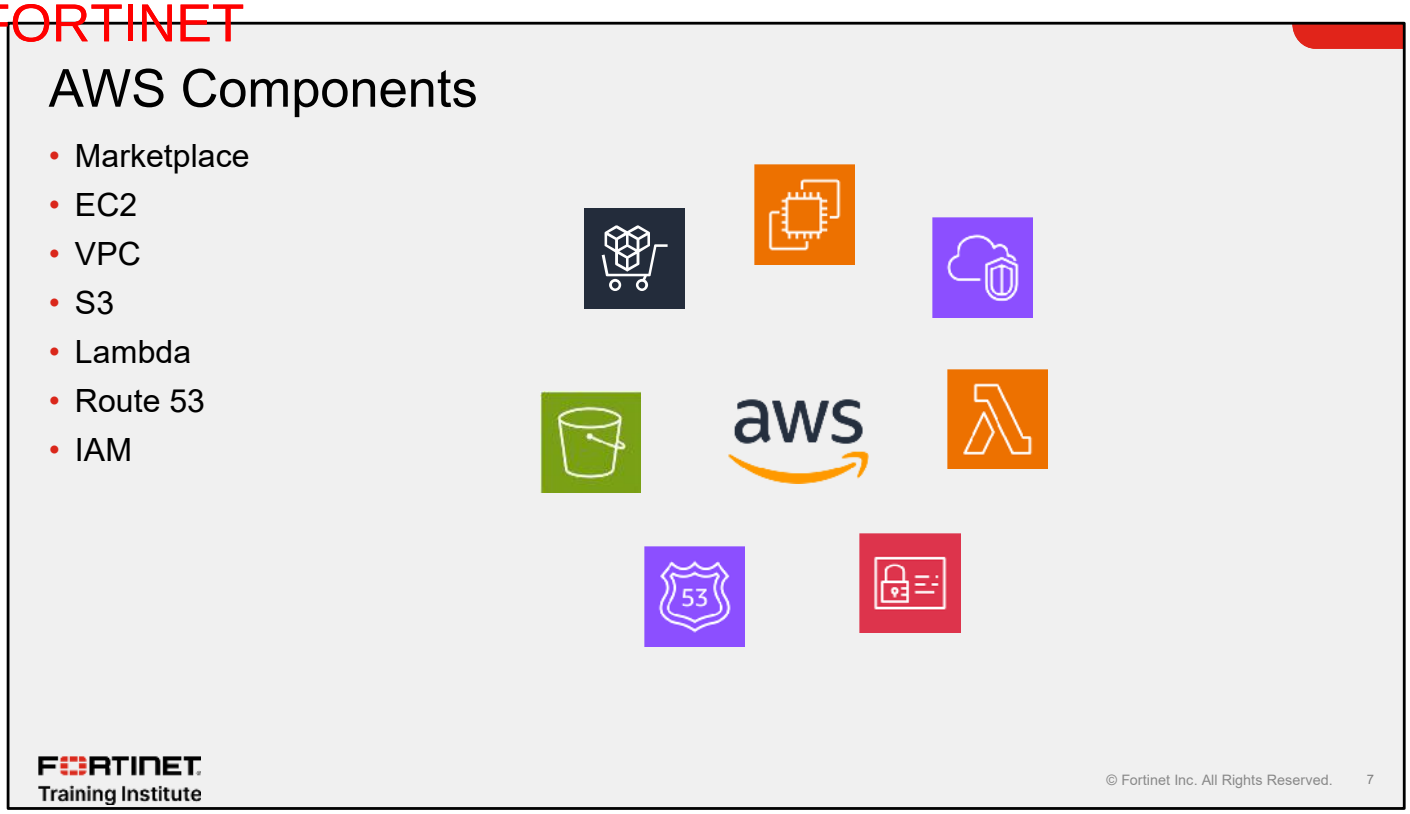

Some of the key AWS components covered in this lesson are:

- AWS Marketplace: An online store that allows customers to discover, purchase, and deploy a wide range of software and services directly on the AWS cloud platform including FortiGate, FortiWeb, FortiManager, and more.
- Amazon EC2: A scalable cloud computing service that provides resizable virtual machines (VMs) (instances) for running applications and services on the AWS cloud infrastructure.
- Amazon Virtual Private Cloud (VPC): A customizable and isolated network environment within AWS, allowing users to launch AWS resources like EC2 instances in a virtual network that they have complete control over.
- Amazon Simple Storage Service (S3): A scalable and highly durable object storage service, designed for storing and retrieving data, files, and objects in the cloud. You can use this service to bootstrap FortiGate configurations or to store VM licenses.
- AWS Lambda: A serverless compute service that allows you to run code in response to events and automatically manage the computing resources required, without the need to provision or manage servers such as a full EC2 instance. This can be useful for HA failovers or responding to intrusions.
- Amazon Route 53: A scalable and highly available domain name system (DNS) web service, designed to route incoming internet traffic to resources such as web servers, load balancers, and other AWS services.
- AWS Identity and Access Management (IAM): A service that enables you to securely control access to AWS resources by defining and managing permissions for users, groups, and roles within your AWS account.

EC2 Components

- Elastic Network Interface (ENI)
  - A virtual network interface
  - Attributes (IP/MAC/security group) follow the ENI when it is attached to or detached from an instance
  - When you move an ENI, network traffic is redirected to the new instance
  - Each instance in your VPC has a default network interface (the primary network interface) that is assigned a private IPv4 address from the IPv4 address range of your VPC
    - You cannot detach a primary network interface from an instance
  - Cannot move between AZs
- Source/destination check
  - Set by the ENI
  - Allows source and destination IP addresses that are different from the assigned IP address of the interface
  - Required on private FortiGate interface when routing traffic to other networks

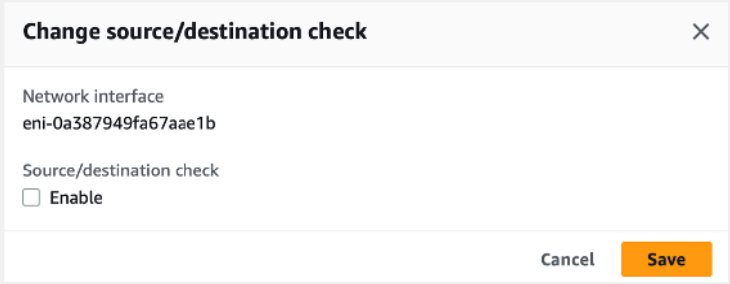

FORTINET Training Institute

 8

An ENI is required so that EC2 instances, such as FortiGate VMs, can communicate. These virtual network interfaces can be attached to, detached from, and reattached to another instance. When an ENI is reattached or moved to another instance, the network traffic destined to this interface is redirected to the new instance. You can also modify the attributes of your network interface, including changing its security groups and managing its IP addresses. Keep in mind that once an ENI is created inside the AZ, it cannot be moved outside the AZ.

The source/destination check feature is set for the network interface. If you disable source/destination checks in AWS, source and destination IP addresses that are different from the assigned IP address of the interface are allowed. In AWS, the source/destination check feature is enabled by default.

Amazon Virtual Private Cloud

- Your own network in the cloud
  - The base classless interdomain routing (CIDR) is defined here
- Allows more than one VPC
- Defined within a region
- Spans all AZs within the region in which it is defined
- Contains subnets, routing, EC2 instances, IP addressing, and so on

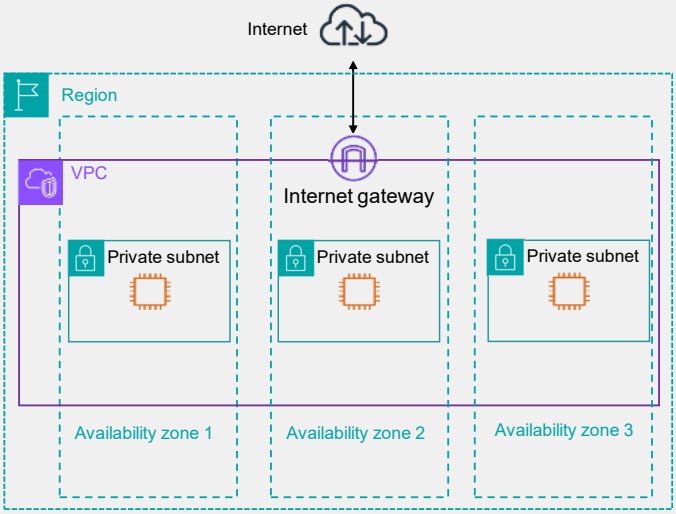

FORTINET Training Institute

 9

A VPC allows you to define a virtual network in your own logically isolated area within AWS Cloud. This is the same concept as a virtual network in Microsoft Azure. The VPC belongs to a region, and within the VPC, you can create different subnets. All subnets should be in the same CIDR block that is defined for the VPC, for example, `10.0.0.0/16`.

As shown on this slide, the VPC (the purple rectangle) belongs to a region but not to any AZs. Within the VPC, you can deploy subnets that belong to different AZs.

Elastic IP Addresses

- An Elastic IP (EIP) address is a static, public IPv4 address
- You can associate an EIP address with any EC2 instance or ENI for any VPC in your account
- You can use an EIP address to mask the failure of an instance by rapidly remapping the address to another instance in your VPC
- Associating an EIP address with an ENI instead of directly with the instance, allows all attributes of the ENI to move from one instance to another, in a single step

FORTINET Training Institute

 10

An EIP address is a static, public IPv4 address. You can associate an EIP address with any instance or network interface for any VPC in your account. You can use an EIP address to mask the failure of an instance by rapidly remapping the address to another instance in your VPC. Associating the EIP address with the network interface instead of directly with the instance, means that you can move all the attributes of the network interface from one instance to another, in a single step.

During the lab, you will see both EIP and non-EIP addresses. Keep in mind that in an active-passive HA setup, you must use an EIP address to move from one instance to another during a failover.

Subnets

- Contained to one specific AZ
- Public: internet gateway connected subnet
  - Doesn’t mean public addressing within the subnet
- Private: internal subnet without an internet gateway
  - They must follow the addressing space defined in the VPC that they belong to
- Reserved IP addresses:
  - First IP address (X.X.X.1): intrinsic router
  - Second IP address (X.X.X.2): AWS DNS
  - Third IP address (X.X.X.3): reserved for future use

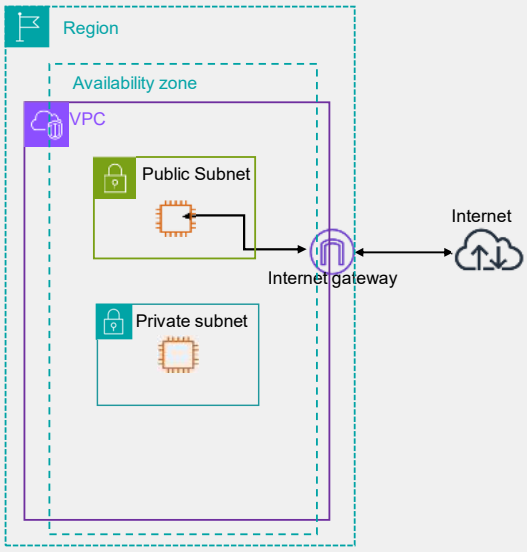

FORTINET Training Institute

 11

A subnet is contained to one specific AZ. It cannot span availability zones.

There are two different kinds of subnets: public and private. A public subnet has an internet gateway attached, and therefore has internet access. It may or may not have public IP addressing, such as an EIP. A private subnet is an internal subnet that doesn't have an internet gateway attached to it. Private subnets must follow the addressing space defined in the VPC that they belong to.

For example, if you want to deploy a FortiGate device for outgoing traffic protection, you can have one interface connected to the private subnet and the other interface connected to the public subnet. After you connect both interfaces, you define a routing table for the private subnet to route internet traffic through the FortiGate device and then to the public subnet.

The first three usable IP addresses in AWS are reserved. The first IP address is reserved for the intrinsic router, the second IP address is reserved for AWS DNS, and the third IP address is reserved for future use. If you deploy a FortiGate device, you must use the fourth usable IP address. By default, this is done automatically with DHCP addressing.

Intrinsic Router

- AWS routing service
- All subnets are connected to an intrinsic router that resides at the VPC level (in all AZs)
- It is assigned the first IP address of the subnet
- Referred to as the default gateway for the ENI of an EC2 instance

FORTINET Training Institute

 12

An intrinsic router doesn't relay traffic. It is a service that runs on hypervisors throughout the AWS infrastructure. It is for OSI layer 3 destinations of the packets as soon as the ENI of an EC2 instance sends them out.

All subnets are connected to an intrinsic router that resides at the VPC level. The default gateway or intrinsic router is always the first IP address of the subnet.

Internet Gateway

- Allows for communication between instances in your VPC and the internet
- Serves two purposes:
  - Provides a target in your VPC route tables for internet-routable traffic
  - Performs NAT for instances that have been assigned public IPv4 addresses
- To enable communication over the internet, an instance must have a public address, or an elastic IP (EIP) assigned to it

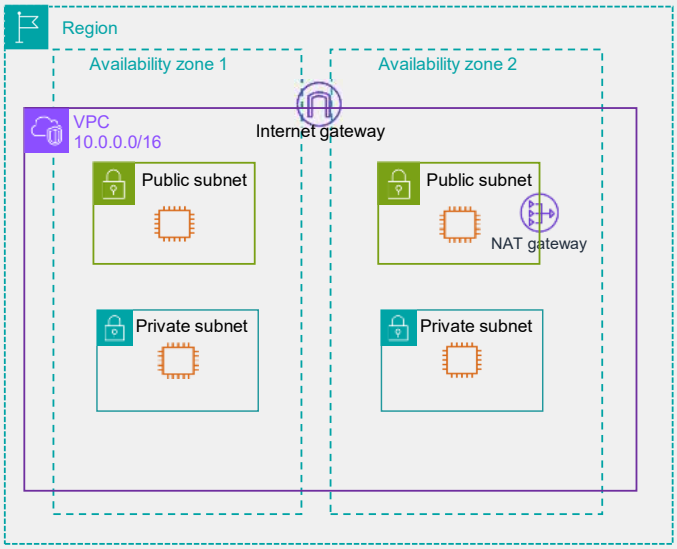

FORTINET Training Institute

 13

An internet gateway is a VPC component that facilitates communication between instances in your VPC and the internet. Internet gateway is a feature that you enable in the subnet, allowing the intrinsic router to connect to the internet. If you want to make the subnet public, you first must create an internet gateway and then attach it to the appropriate subnet routing table. It is a paid-for service.

An internet gateway serves two purposes:
- It provides a target in your VPC route tables for internet-routable traffic
- It performs network address translation (NAT) for instances that have been assigned public IP addresses

An internet gateway does not perform NAT. An instance must have a public IP address or EIP assigned to it.

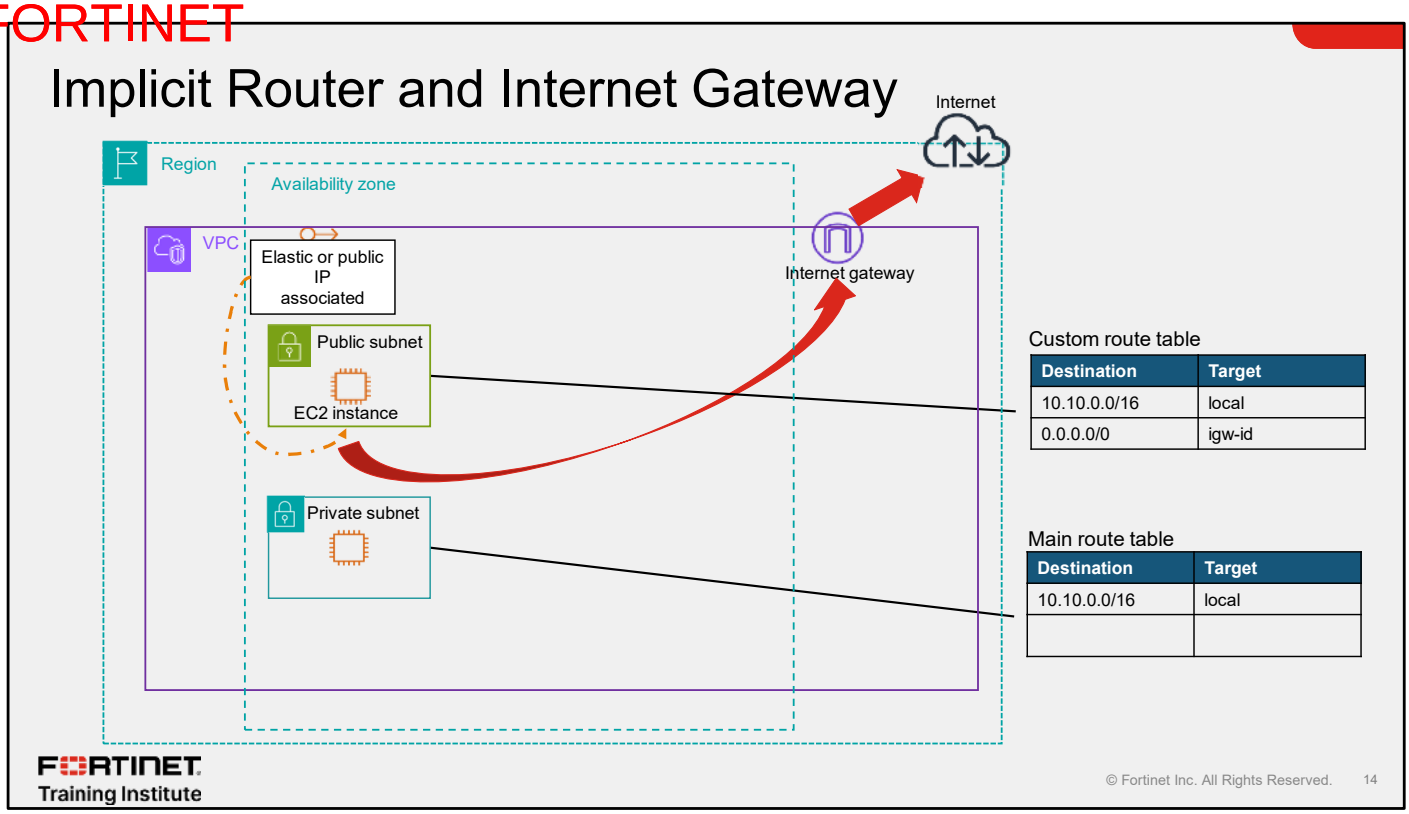

The diagram on this slide shows the routing for a VPC with an internet gateway, a public subnet, and a private subnet. The main route table comes with the VPC. The custom route table is associated with the public subnet. The custom route table has a route pointing to the internet gateway (the destination is `0.0.0.0/0`, and the target is the internet gateway).

The topology shows two subnets in the same AZ. The AZ is inside the VPC and the VPC itself is placed within a region.

The router is the intrinsic router and is responsible for routing traffic between the subnets. You can create additional routing tables, and then associate them to a subnet.

Routing Tables

- By default, subnets are associated with the main routing table
- You can create more routing tables and explicitly associate them with subnets
- A gateway is not defined by an IP address
  - Instead, it uses the ENI object
- EC2 instances always use the intrinsic router as the default gateway but they are then redirected to each gateway defined in the routing table
- You can use static or traditional routing within an instance, but automation could be affected

FORTINET Training Institute

 15

As mentioned previously in this lesson, when you create a VPC, a default main routing table is created along with it, and all subnets are associated with the main routing table. You can reassociate subnets to different route tables after you create them.

The gateway of a route uses an ENI object, internet gateway, NAT gateway, etc and is not defined by an IP address.

EC2 instances always use the intrinsic router as the default gateway, but they are then redirected to a gateway defined in the routing table.

Note that you can create and use traditional or static routes within an instance, but this might be problematic for future automation.

AWS NAT Gateway

- Allows instances in a private subnet to connect to the internet
- No need to launch a NAT instance
- A route table sends internet traffic from private subnet instances to the NAT gateway
- NAT gateway sends traffic to the IGW using the source IP address of the elastic IP address
- Instances of a private subnet hide behind the NAT gateway

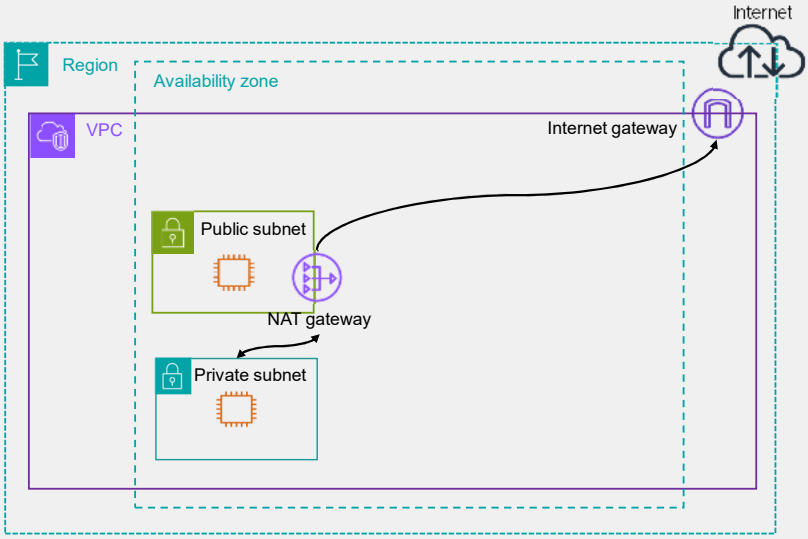

FORTINET Training Institute

 16

The NAT gateway allows instances in a private subnet to connect to the internet or other AWS services without using a NAT instance. A route table sends internet traffic from private subnet instances to the NAT gateway, then the NAT gateway sends traffic to the Internet Gateway (IGW). The advantage here is that instances of the private subnet hide behind the NAT gateway. This scenario is useful if you want to inspect only layer 4 traffic. However, if you want to inspect higher-layer traffic, you must deploy a proper firewall.

Other VPC Components

- DHCP:
  - EC2 interfaces should use DHCP
  - All addressing is defined at the AWS console level
  - You can use static IP addresses, but they must match the AWS configuration
- DNS:
  - Each EC2 instance has an internal DNS name
  - The DNS server is present on all subnets as the second reserved IP address

FORTINET Training Institute

 17

Other VPC components are the DHCP and DNS services. EC2 instance interfaces should use DHCP. By default, when you assign an IP address to an ENI, the DHCP service is automatically activated and delivers the IP address to the DHCP-enabled interface of the device. You can enable the DHCP feature on the FortiGate interface to receive the IP address. You can also create specific options inside the DHCP server.

Each EC2 instance has an internal DNS name that you can use to send traffic to. This DNS server is present on all subnets as the second reserved IP address and is the default DHCP option. Every time you deploy a network interface, it gets assigned a random DNS name without a VPC reference. You can resolve this random DNS name both inside and outside the VPC.

VPC Peering

- Connection between two VPCs that allow you to route traffic between them
- Instances in either VPC can communicate with each other as if they are in the same network
- Can create VPC peering connections between your own VPCs, or with VPCs in other AWS accounts
  - VPCs can be in different regions (also known as an *inter-region* VPC peering connection)
- No single point of failure for communication or a bandwidth bottleneck
- A one-to-one relationship between two VPCs
  - You can create multiple VPC peering connections for each VPC, but transitive peering relationships are not supported
  - You can't have any peering relationships with VPCs that your VPC is not directly peered with

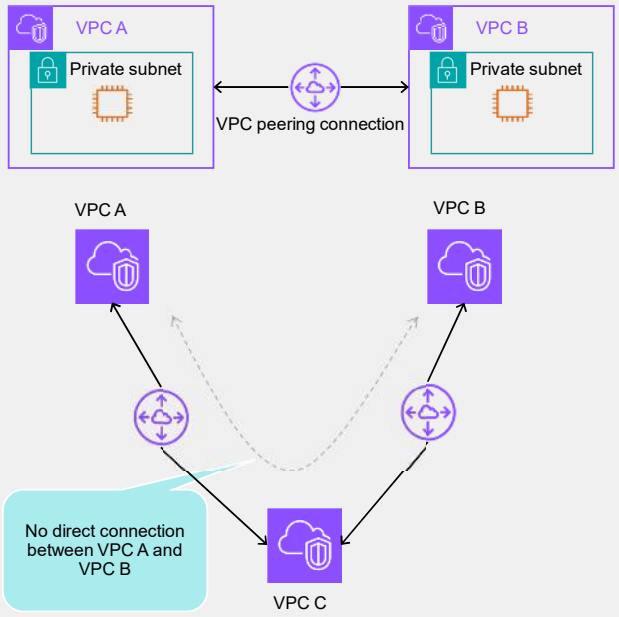

FORTINET Training Institute

 18

VPC peering is a connection between two VPCs that allows you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are in the same network, similar to a VPN connection. You can create a VPC peering connection between your own VPCs, or with VPCs in other AWS accounts. You can connect VPCs among different regions (inter-region VPC peering connection) or within the same region. It is like connecting the VPCs using a single cable. For example, you can have a route to force traffic to go to from one specific VPC to another VPC. A VPC peering connection is a one-to-one relationship between two VPCs. Note that the cost associated with VPC peering varies depending on whether they are connected within the same region or a different region.

This slide shows some VPC peering limitations. For example, you cannot route packets directly from VPC B to VPC C through VPC A. To route packets directly between VPC B and VPC C, you must create a separate VPC peering connection between them. You can think of this as a manual full-mesh or star topology.

Transit VPC With IPsec

- Good solution to reduce the complexity of VPC peering but huge administrative work in the central VPC
- Adding new VPCs requires more administrative work and introduces complexity
- Central VPC must be highly available
- Use IPsec tunnels to connect
- Not the most efficient way to route and control traffic in AWS

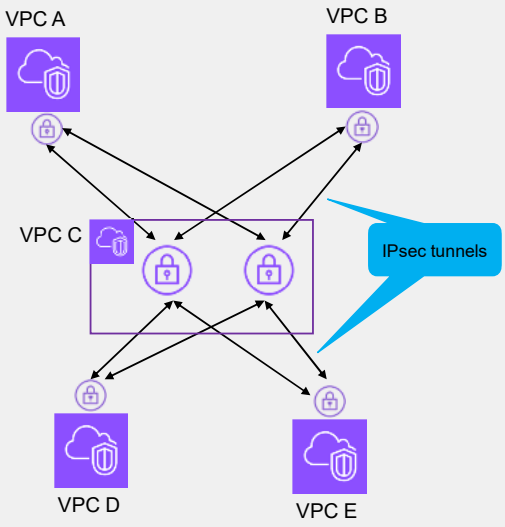

FORTINET Training Institute

 19

A star topology is not conducive if many VPCs must all be interconnected. One option to reduce the number of connections is to use a transit VPC. This most closely resembles a hub-and-spoke topology. As shown on this slide, VPCs A, B, D, and E are all interconnected using IPsec tunnels through transit VPC C.

While a transit VPC reduces the complexity of VPC peering, adding more VPCs to the existing setup introduces a huge administrative task. A growing organization that is continuously adding more VPCs should migrate to the AWS transit gateway. Another drawback of the transit VPC is the need to maintain EC2 instances, in the form of a pair of routers, which must be highly available to route traffic among VPCs.

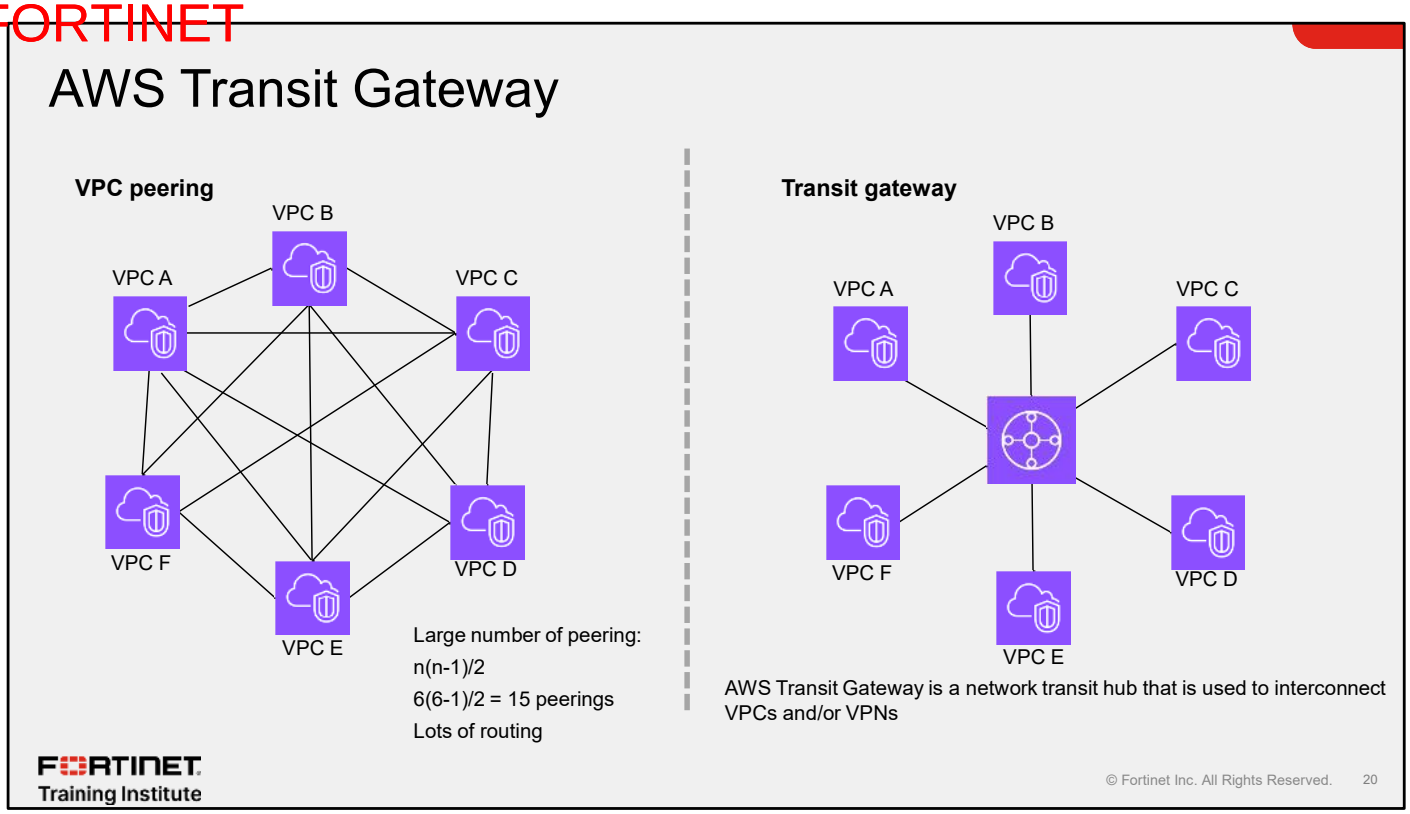

So, what are the complexities of manual VPC peering? To achieve full mesh connectivity between VPCs, you will need to use formula n(n-1)/2. As shown in the example on this slide, full mesh connectivity between six VPCs requires 15 connections. Now, imagine if you had hundreds of VPCs requiring full mesh connectivity. This would resolve in a lot of administrative overhead.

To simplify this setup and decrease this administrative overhead, we can use a transit gateway. AWS Transit Gateway (TGW) solves most of the problems introduced by VPC peering. As shown on this slide, TGW is similar to the transit VPC hub-and-spoke topology, however, AWS TGW is much simpler and more flexible.

Transit Gateway—Concept

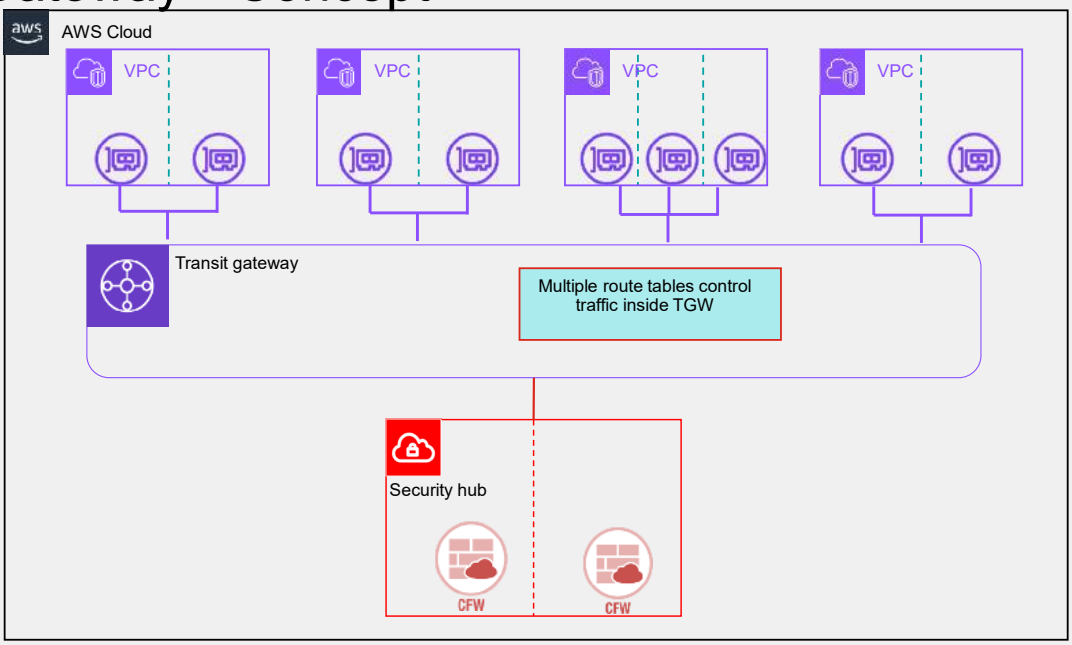

FORTINET Training Institute

 21

Now, you will learn more about the concept of AWS Transit Gateway.

As shown on this slide, many VPCs connect to the AWS TGW. From the TGW, you can define rules to route traffic between VPCs and restrict traffic. There is no need to connect VPCs with multiple IPsec tunnels. You can connect multiple VPCs to the TGW and then define rules to send traffic to the security hub VPC for traffic inspection between VPCs (east-west traffic inspection), or to send traffic directly from one VPC to another through the TGW.

Transit Gateway

- 5 TGWs per account
- 20 TGW route tables for routing control
- 10,000 static routes per TGW
- Up to 50 Gbps of bandwidth per attachment per availability zone
- Up to 5000 Amazon VPC attachments per TGW
- 1.25 Gbps per VPN connection
- Support for ECMP
- Support for multicast

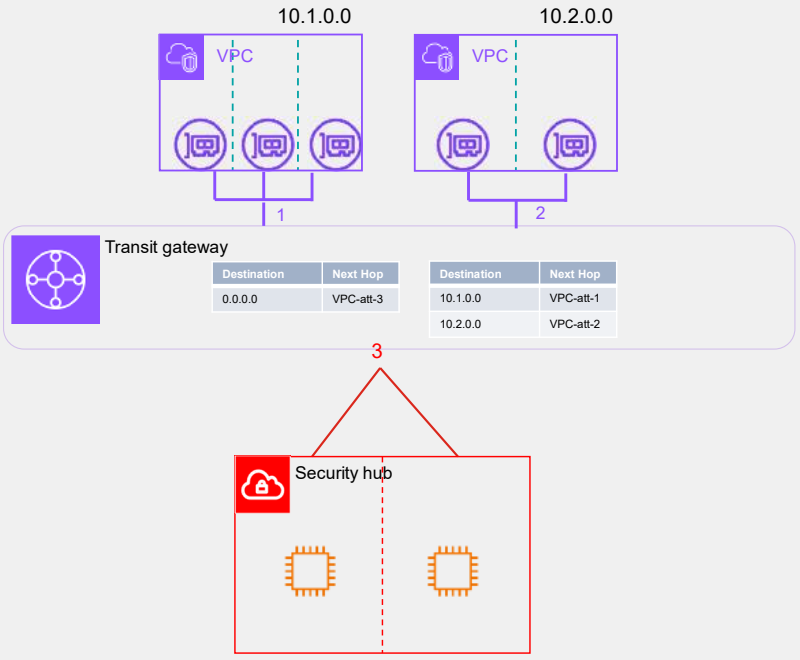

FORTINET Training Institute

 22

A TGW helps to solve multiple issues with VPC peering and transit VPC. Using TGW technology, you can create multiple transit gateway route tables inside TGW for better traffic control. As shown in the example on this slide, you can create multiple attachments based on the number of VPCs you need to connect. For example, you will need only three attachments to create all three VPCs. This eliminates the full mesh requirement that is part of the VPC peering scenario.

As shown in the example on this slide, there are two route tables inside the TGW with three attachments. Any traffic coming to the TGW, except subnets `10.1.0.0` and `10.2.0.0`, goes to the security hub VPC through attachment `VPC-att-3`. Traffic going to subnet `10.1.0.0` uses `VPC-att-1`, and subnet `10.2.0.0` uses the attachment `VPC-att-2`. This granular level of control means a lighter workload for the administrator when they are adding multiple VPCs to the existing environment.

Another main advantage is bandwidth. Customers can create multiple VPN connections from the TGW to the on-premises data center with ECMP to achieve higher bandwidth.

Aside from ECMP, TGW also has support for multicast.

Transit Gateway Connect

- Connects SD-WAN infrastructure with AWS
  - Supports Generic Routing Encapsulation (GRE) and BGP
  - Higher bandwidth performance than VPN–up to 20Gbps per attachment.

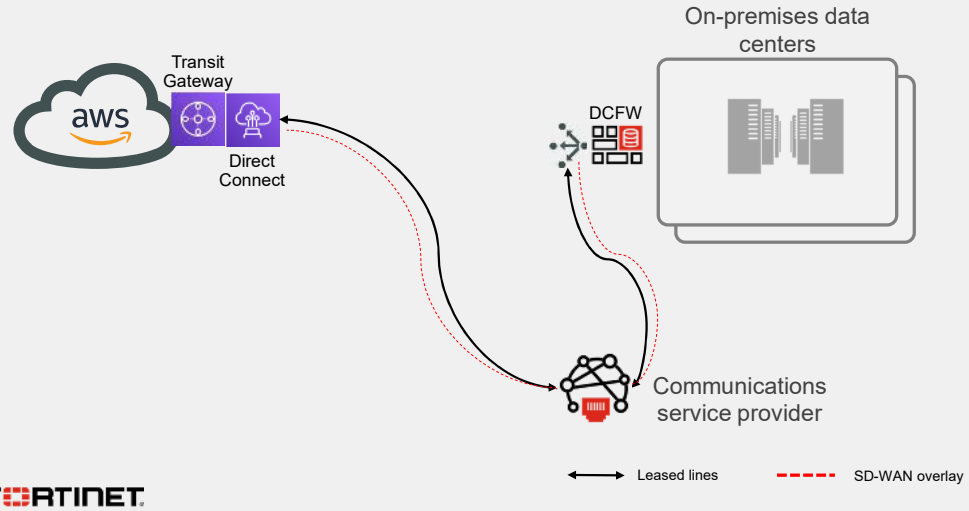

FORTINET Training Institute

 23

You can use Transit Gateway Connect to not only connect VPCs to a transit gateway, but also your on-premises data centres and branch offices.

AWS Transit Gateway Connect allows your existing SD-WAN infrastructure to be seamlessly integrated with AWS.

Transit Gateway Connect extends SD-WAN into AWS without the need to set up IPsec VPNs between the SD-WAN network device and AWS Transit Gateway. It supports GRE for higher bandwidth performance compared to a VPN connection. Transit Gateway Connect supports BGP for dynamic routing and removes the need to configure static routes.

Transit Gateway Connect—Benefits

- Centralized network hub:
  - Connect multiple VPCs and on-premises networks
  - Single point for managing and monitoring connectivity
- Simplified management:
  - Reduces complexity by managing connectivity through a single interface
- Dynamic routing:
  - Uses BGP
  - Allows for scalable and flexible network configurations
- Seamless integration with SD-WAN
- Global reach and multiregion connectivity:
  - Can be associated with multiple AWS regions
  - Allow for connectivity across different geographic locations

FORTINET Training Institute

 24

The following are the benefits of Transit Gateway Connect:

- AWS Transit Gateway Connect serves as a centralized hub for connecting multiple VPCs and on-premises networks. This centralization simplifies network architecture by providing a single point for managing and monitoring connectivity.
- The solution simplifies the management of network connectivity by offering a unified and streamlined interface. Instead of managing individual connections for each VPC, organizations can use Transit Gateway Connect to manage connectivity through a single interface, reducing complexity.
- Transit Gateway Connect uses BGP for dynamic routing. BGP facilitates automatic route updates, allowing organizations to adapt to changes in the network topology without manual intervention. This dynamic routing capability is crucial for scalable and flexible network configurations.
- AWS Transit Gateway Connect is designed to seamlessly integrate with SD-WAN solutions. This integration simplifies the deployment and management of SD-WAN connectivity to AWS resources, providing organizations with flexibility in their network architecture.
- AWS Transit Gateway Connect can be associated with multiple AWS Regions, providing global reach and connectivity. This allows organizations to connect resources across different geographic locations, supporting multi-region architectures and ensuring consistent connectivity on a global scale.

Transit Gateway Connect—Architecture

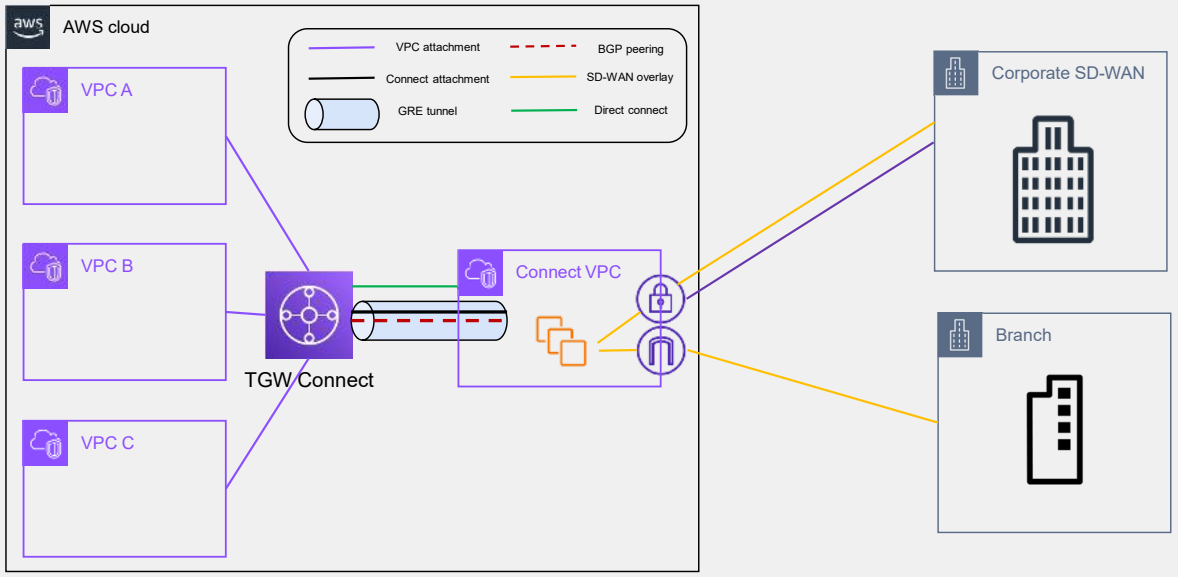

FORTINET Training Institute

 25

This slide shows a high-level architecture of TGW connect using VPC attachments with a virtual appliance. The Connect VPC contains third-party virtual device instances that are connected to branch offices using SD-WAN networks. Then, these instances link to the TGW using a Connect attachment type.

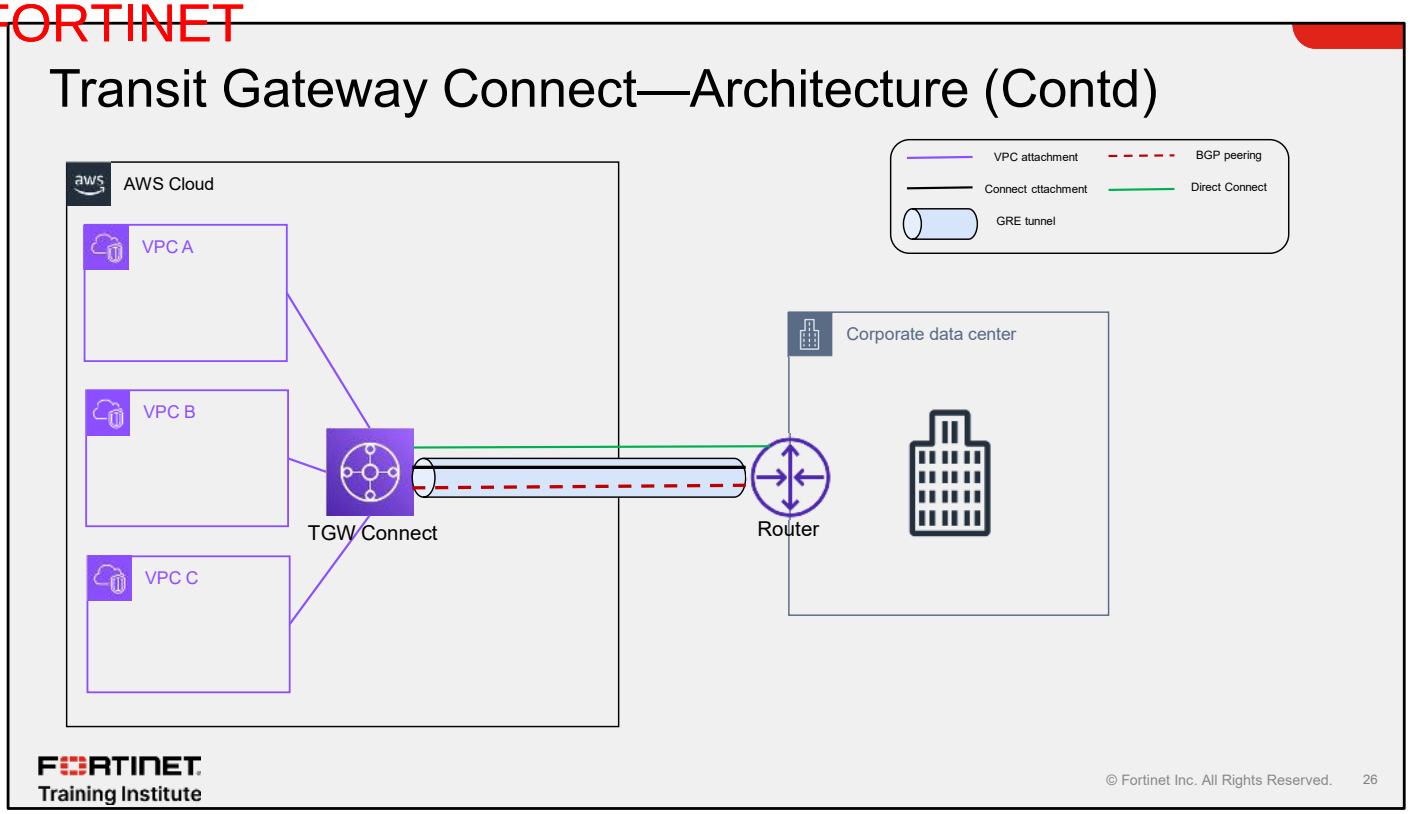

You can also use Transit Gateway Connect with a third-party branch or on-premises data center using AWS Direct Connect as transport. This slide shows a Transit Gateway Connect to AWS Direct Connect attachment. BGP peering has been established over GRE between the TGW Connect and the corporate data center using the Connect attachment. The same GRE tunnel is used to exchange traffic with the TGW.

Knowledge Check

1. In AWS, what does ENI stand for?
   A. Elastic Network IP
   ✓ B. Elastic Network Interface

2. What is a requirement for an internet gateway to route traffic to the internet?
   ✓ A. An Elastic IP address must be assigned to an instance
   B. A custom route table

FORTINET Training Institute

 27

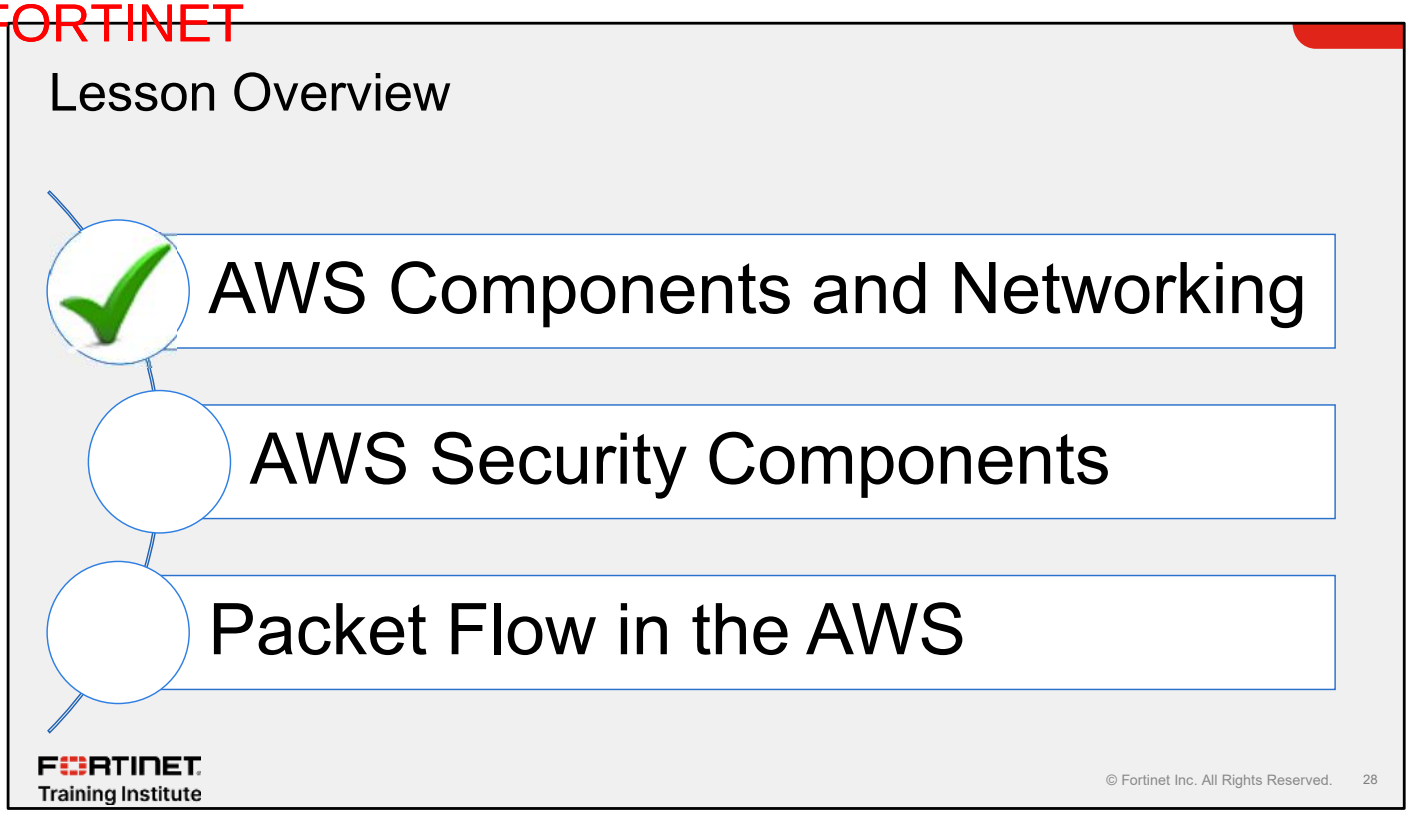

Good job! You now know about AWS components and networking.

Now, you will learn about AWS security components.

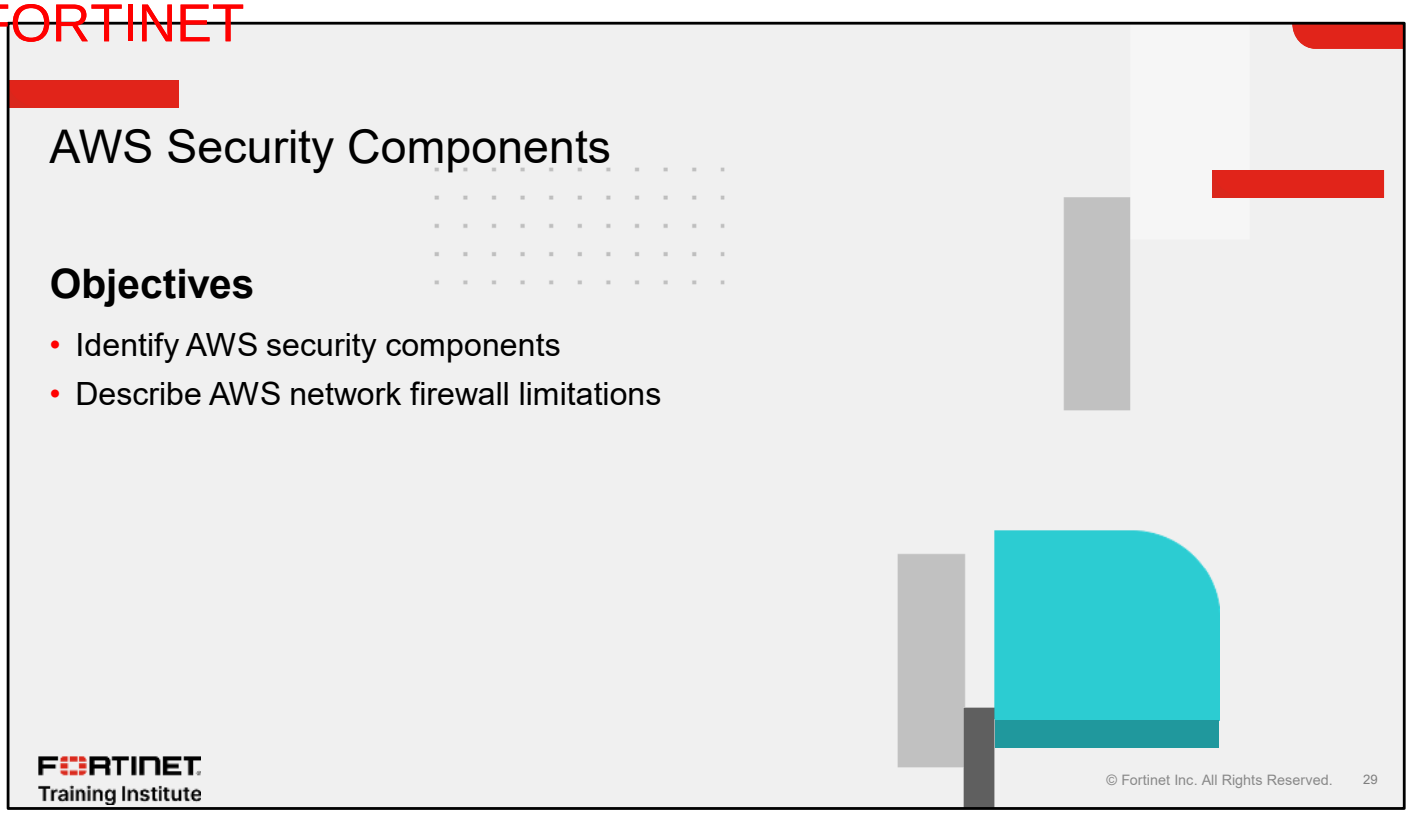

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding AWS security components, you will be able to use secure public cloud deployments.

Security Groups

- Acts as a virtual firewall that controls traffic for one or more instances
- Associated with network interfaces
  - Changing the Security Group (SG) of an instance, changes the SG associated with the primary network interface (eth0)
- Allows all outbound traffic by default
  - SG rules are always permissive; you can't create rules that deny access
  - Stateful
- Instances are automatically associated with the default SG (unless you specify a SG)
  - Allows all inbound traffic from other instances associated with the default security group
- Rules are aggregated when assigning multiple SGs to an instance
- You can create your own security groups and specify them when you launch instances
- Tip: If things aren't working, check the SGs first

FORTINET Training Institute

 30

An SG acts as a virtual firewall that controls the traffic for one or more instances. SGs are associated with network interfaces. Changing the SG of an instance changes the SG associated with the primary network interface (eth0). By default, SGs allow all outbound traffic.

Instances are automatically associated with the default SG (unless you specify an SG).

When you associate multiple SGs with an instance, the rules from each SG are effectively aggregated to create one set of rules.

You can create your own SGs and specify them when you launch your instances.

The only difference between the SGs in AWS and Azure, is that in AWS SGs are attached to the network interfaces.

Network Access Control Lists (NACLs)

- Optional layer of security for your VPC
- VPC automatically comes with a modifiable default NACL
  - By default, it allows all inbound and outbound IPv4 traffic
- You can create a custom NACL
  - Denies all inbound and outbound traffic until you add rules
- Each subnet in your VPC must be associated with an NACL
  - If you don't, it is associated with the default NACL
- Separate inbound and outbound rules
- Stateless

FORTINET Training Institute

 31

An NACL is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. Your VPC automatically comes with a modifiable default NACL.

By default, it allows all inbound and outbound IPv4 traffic. You can create a custom NACL and associate it with a subnet. By default, each custom NACL denies all inbound and outbound traffic until you add rules.

Each subnet in your VPC must be associated with an NACL. If you don't explicitly associate a subnet with an NACL, it is associated with the default NACL.

An NACL has separate inbound and outbound rules, and each rule can either allow or deny traffic.

NACLs are stateless; responses to allowed inbound traffic are subject to the rules for outbound traffic (and the reverse).

Flow Logs and CloudWatch

- Captures information about IP traffic going to and from network interfaces within a VPC
  - No additional charge for using flow logs
  - Flow log data is published to a log group in CloudWatch Logs
  - Each network interface has a unique log stream
- No cost associated with flow logs but CloudWatch is a paid service.
- Flow logs do not capture:
  - Traffic to and from `169.254.169.254`, for example, metadata
  - Traffic to and from `169.254.169.123` for the Amazon Time Sync Service
  - DHCP traffic
  - Traffic to the reserved IP address for the default VPC router
  - Flow logs do not capture real-time log streams for your network interfaces
- Use flow logs as a security tool to monitor the traffic that is reaching your instance

FORTINET Training Institute

 32

VPC flow logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. There is no additional charge for using flow logs; however, standard CloudWatch Logs charges apply. Flow log data is published to a log group in CloudWatch Logs, and each network interface has a unique log stream.

Flow logs do not capture traffic to and from `169.254.169.254`, such as metadata traffic. They also do not capture traffic to and from `169.254.169.123` for the Amazon Time Sync Service, DHCP traffic, and traffic to the reserved IP address for the default VPC router. Also, flow logs do not capture real-time log streams for your network interfaces.

You can use flow logs as a security tool to monitor the traffic that is reaching your instance. Flow logs are useful if you want to perform quick troubleshooting and to view the behavior of the security groups.

Amazon GuardDuty

- Managed threat detection service
- Scans for malicious or unauthorized behavior
- Helps to protect AWS accounts and workloads
- Monitors activity, such as:
  - Unusual API calls
  - Potentially unauthorized deployments that indicate a possible account compromise
- GuardDuty also detects potentially compromised instances or reconnaissance by attackers

FORTINET Training Institute

 33

AWS GuardDuty is a managed threat detection service that continuously monitors and analyzes AWS account activity, identifying and alerting about potentially malicious behavior or security risks.

GuardDuty employs machine—learning algorithms and threat intelligence feeds to analyze collected data. The machine learning models identify patterns, anomalies, and behaviors that may indicate malicious activities.

You can use the FortiGate threat feed feature to obtain all blocklisted IP addresses from GuardDuty and then create appropriate firewall policies to block traffic.

AWS Network Firewall

- Managed service that makes it easy to deploy essential network protections
- Define rules that provide granular control over network traffic
- Works together with AWS firewall manager
- Stateful firewall—filter and monitor traffic at protocol and port level
- Advanced threat protection using IDS/IPS
- Logging and monitoring

FORTINET Training Institute

 34

AWS Network Firewall is a managed service that provides scalable, stateful, and rule-based network traffic filtering for VPCs to protect against unwanted access and malicious activity. It enables organizations to define and enforce fine-grained security policies for both inbound and outbound traffic, enhancing the overall security posture of their AWS environments.

AWS Network Firewall is designed to work together with the AWS Firewall Manager so you can build policies based on Network Firewall rules and then centrally apply those policies across your VPCs and accounts. This is similar to FortiManager.

AWS Network firewall performs stateful inspection of traffic by inspecting the full context of network traffic, and then makes decisions on whether to allow or deny the traffic based on the configured firewall rules.

Deep packet inspection allows the network firewall service to examine packet content and headers controlling traffic across layer 4 to layer 7.

Firewall rules permit or deny traffic based on IP addresses, port numbers, protocols, and domain names.

IDS and IPS enable the network firewall service to detect and prevent common attacks, such as SQL injection, cross-site scripting and command injection.

Another key feature of AWS Network irewall is logging and monitoring. Detailed logging allows other AWS services like CloudWatch and Amazon S3 to provide monitoring and analysis of network traffic.

AWS Network Firewall Limitations

- No consistency of firewall policy management across hybrid cloud
- Limited advanced threat protection (ATP) feature set
- IPS depends on third party for signatures
- Signatures must be updated manually
- Each AWS account requires its own AWS Network Firewall instance
- Cost

Key capabilities	FortiGate CNF	AWS NF
Consistent hybrid-cloud firewall policy management	●	●
Scale and resiliency	●	●
NGFW market leadership	●	●
Depth in visibility and control	●	●
Region-wide security aggregation	●	●
Dynamic workload security policies	●	●
Cloud native console	●	●

FORTINET Training Institute

 35

AWS Network Firewall provides basic and essential firewalling capabilities, but it has several limitations that can be alleviated by either complementing it with FortiGate-VM or replacing it with FortiGate CNF.

If Fortinet security products are deployed in your on-premises or hybrid environments, use FortiGate-VM or FortiGate CNF in AWS to provide seamless integration and consistent security policies across your entire infrastructure. You cannot integrate your existing rule set for on-premises firewalls into AWS Network Firewall. With FortiGate-VM or FortiGate CNF, FortiManager can provide hybrid-cloud security management.

AWS Network Firewall's ATP feature set is limited to IDS and IPS. FortiGate CNF can be configured with all the next-generation firewall (NGFW) security features of a regular FortiGate, such as web filtering, application control, bad IP filtering, DNS filtering, and so on.

AWS Network Firewall relies on third parties for IDS and IPS signatures. Also, IDS/IPS signatures must be manually updated on AWS Network Firewall, whereas they are automatically updated on FortiGate CNF. FortiGate CNF is powered by FortiGuard Labs threat intelligence services and offers high security efficacy with real-time content updates.

One AWS Network Firewall instance is required for each AWS account. A single FortiGate CNF instance can be used across multiple AWS accounts, simplifying and centralizing management of security rules.

Finally, cost is a factor when it comes to AWS Network Firewall. AWS requires an hourly rate to be paid for each firewall endpoint it is protecting. Additionally, you must also pay for the traffic processed by each firewall endpoint. FortiGate CNF is more cost-effective, paying only for the traffic you want secure.

AWS Web Application Firewall

- Monitor HTTP and HTTPS requests that are forwarded to Amazon CloudFront or an application load balancer
- Ability to configure customizable rules
- Can use Fortinet-managed rules

FORTINET Training Institute

 36

AWS Web Application Firewall (WAF) is a managed security service that helps protect web applications from online threats by allowing you to configure customizable rules to filter and monitor incoming web traffic.

The service allows you to use managed rule sets from third-party vendors such as Fortinet. You will learn more about managed rule sets in the next lesson.

Knowledge Check

1. Which of the following is an AWS Network Firewall limitation?
   - A. Not stateful
   - ✓ B. You must manually update signatures

2. In a security group, which type of traffic is allowed by default?
   - ✓ A. Outbound traffic
   - B. Inbound traffic

FORTINET Training Institute

 37

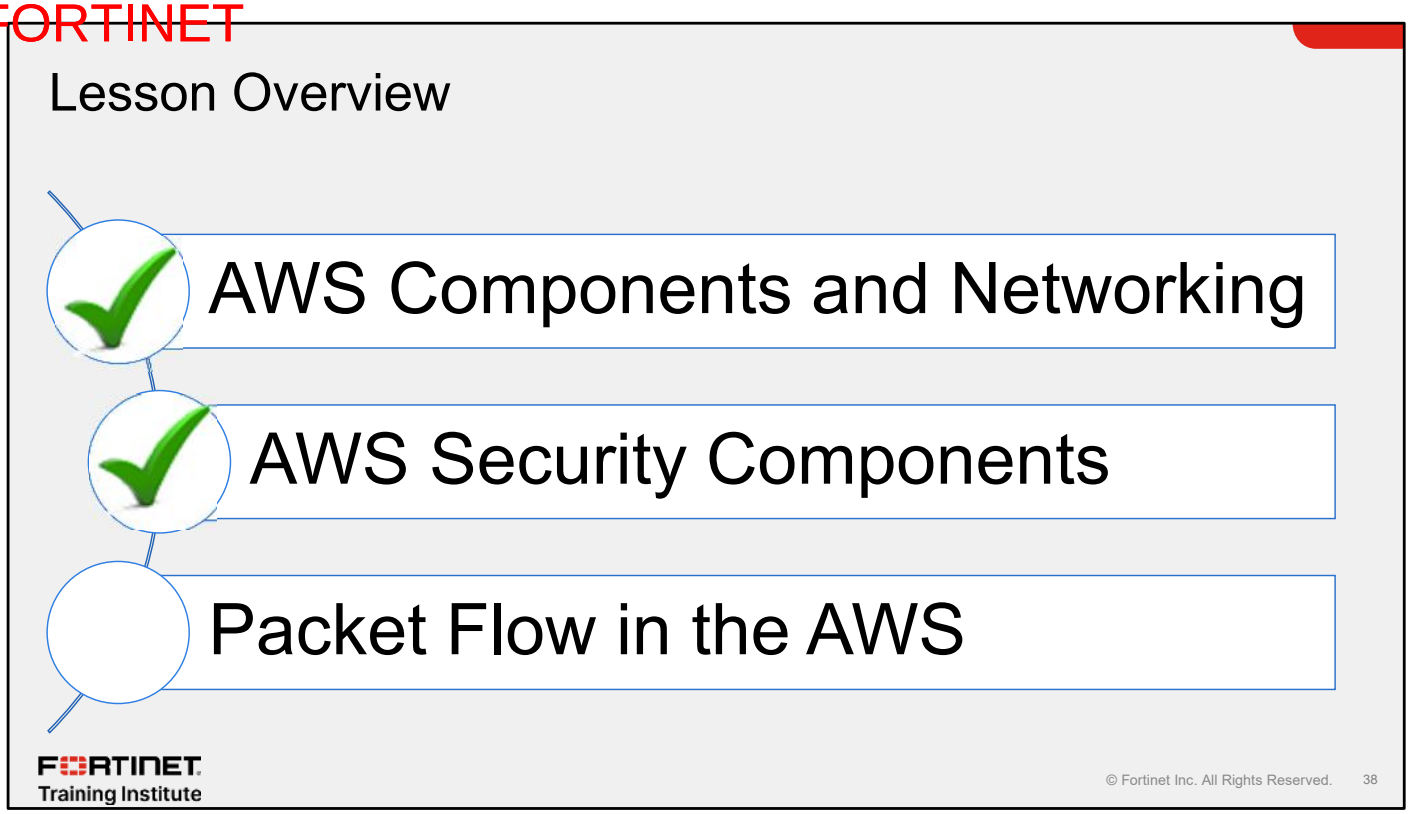

Good job! You now know about AWS security components.

Now, you will learn about network packet flow in the AWS cloud.

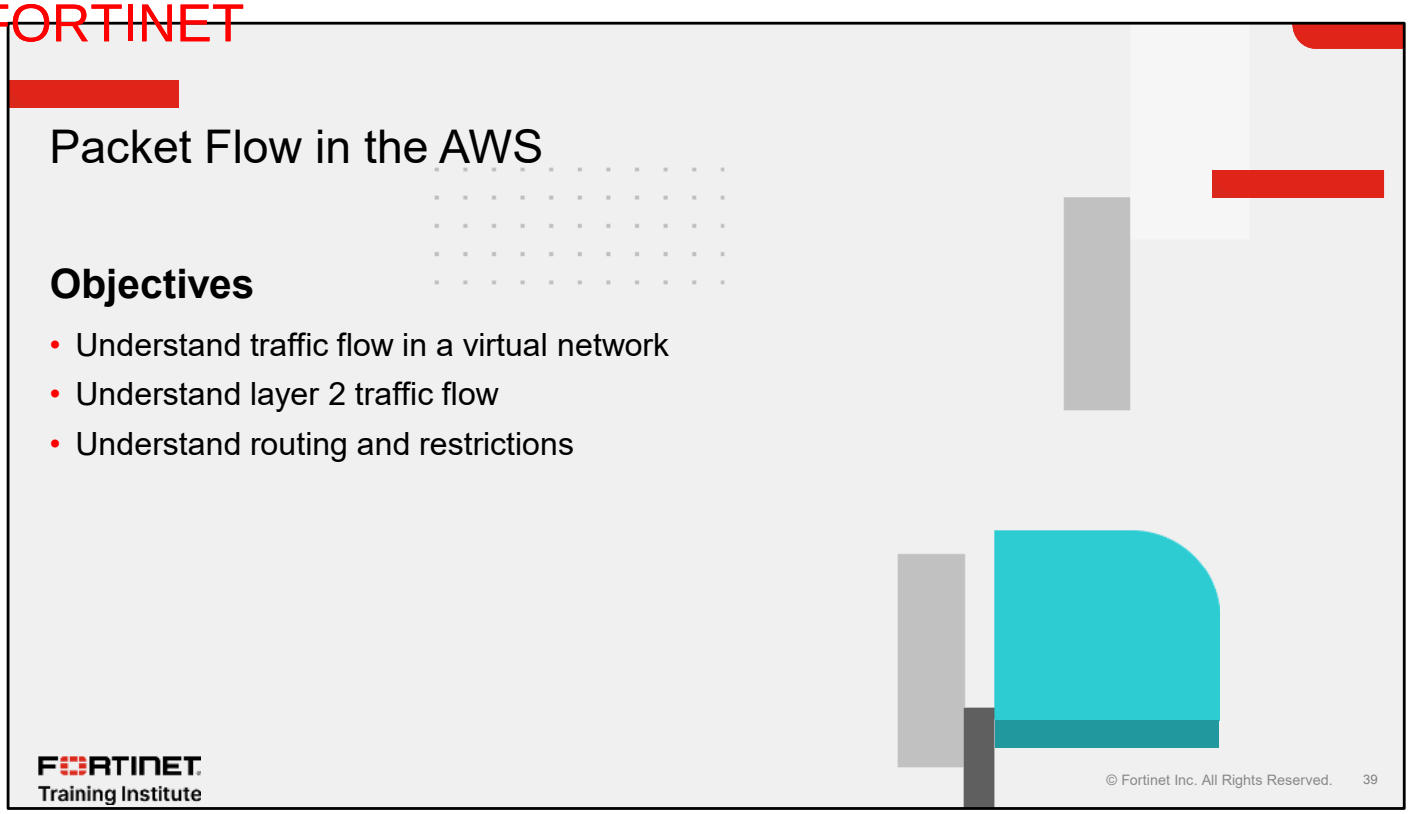

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in networking in the AWS cloud, you will be able to understand traffic flow, and how to manipulate traffic using routes in a virtual network.

VPC Traffic Flow—Scenario 1

- Within a subnet

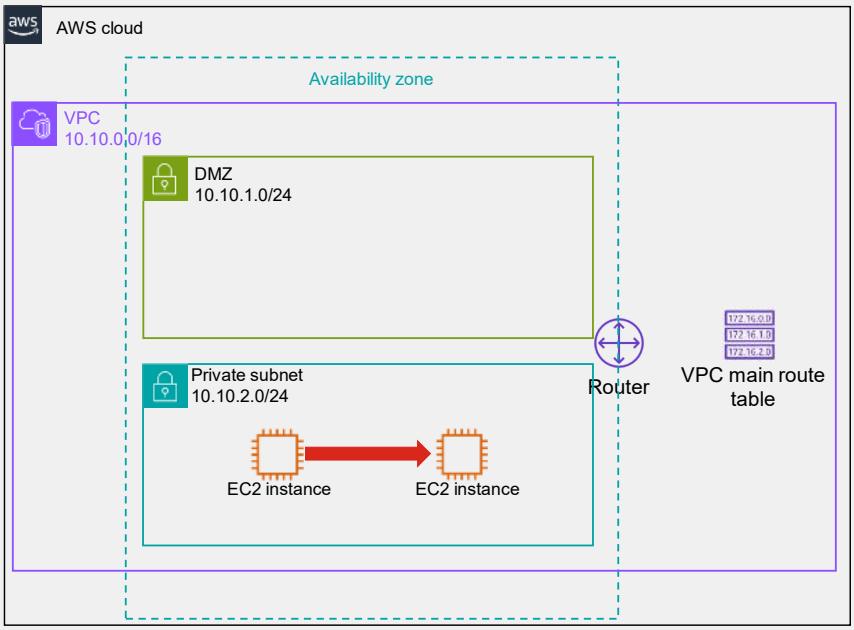

FORTINET Training Institute

 40

Now, you will learn about traffic flow in a VPC. The VPC created in AWS uses a CIDR address space. It is possible to have multiple non-continuous CIDR address spaces assigned to a VPC. However, at minimum, you must assign a single CIDR address space at VPC creation. Within the CIDR address space of the VPC you will build subnets and then connect resources such as VMs to one or more subnets. AWS refers to these instances as EC2 instances.

When you create a VPC, AWS automatically creates a routing table. This allows EC2 instances in the same VPC to send traffic directly to each other.

The connectivity appears similar to traditional Ethernet networks within a single subnet. The OS of an EC2 instance sends an ARP request to learn the MAC address of another EC2 instance and then sends packets to that destination MAC address on layer 2 and the destination IP address on layer 3. However, while traffic flow seems to work similarly to on-premises Ethernet networking, it really is different in AWS VPCs.

VPC Traffic Flow—Scenario 2

- Between subnets in the same virtual network

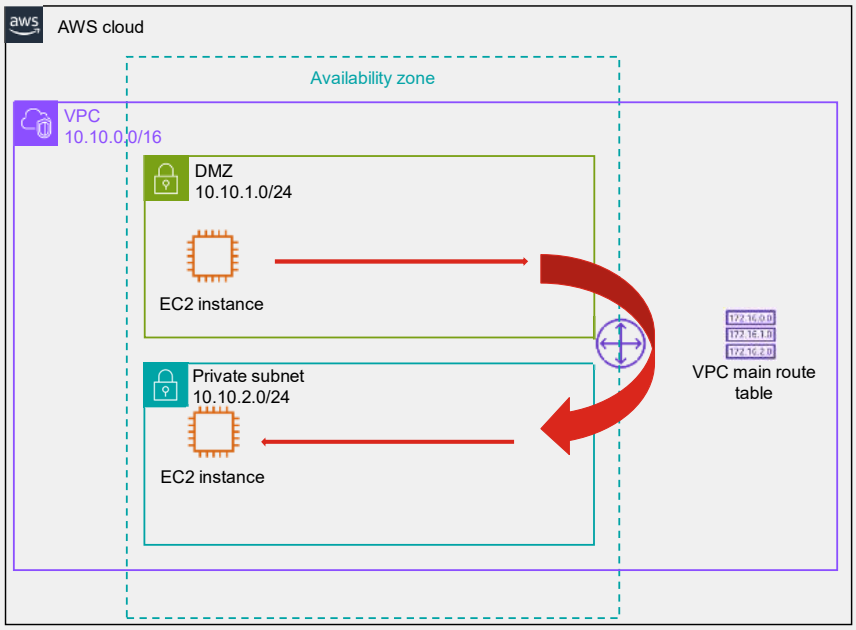

FORTINET Training Institute

 41

When using AWS VPC, it is possible to connect EC2 instances to different subnets within the same VPC. As mentioned earlier, a route table is automatically created when the VPC is created.

This route table always has a route that allows all subnets within the VPC to route to each other. No extra steps are required for configuring connectivity between subnets in the same VPC.

The OSs within the EC2 instance are not aware they are connected to a special cloud network. Each OS has a network interface card driver loaded for a virtualized NIC ethernet adapter. In AWS, the virtual NIC is referred to as an elastic network interface, or ENI.

The OS within an EC2 instance has an IP configuration for its ENI: an IP address, a subnet mask, and a default gateway. The default gateway is always the first IP address of the subnet. This is a cloud networking routing service that is always provided by AWS. This router is referred to as the intrinsic router.

VPC Traffic Flow—Scenario 2 (Contd)

- Between subnets in the same virtual network

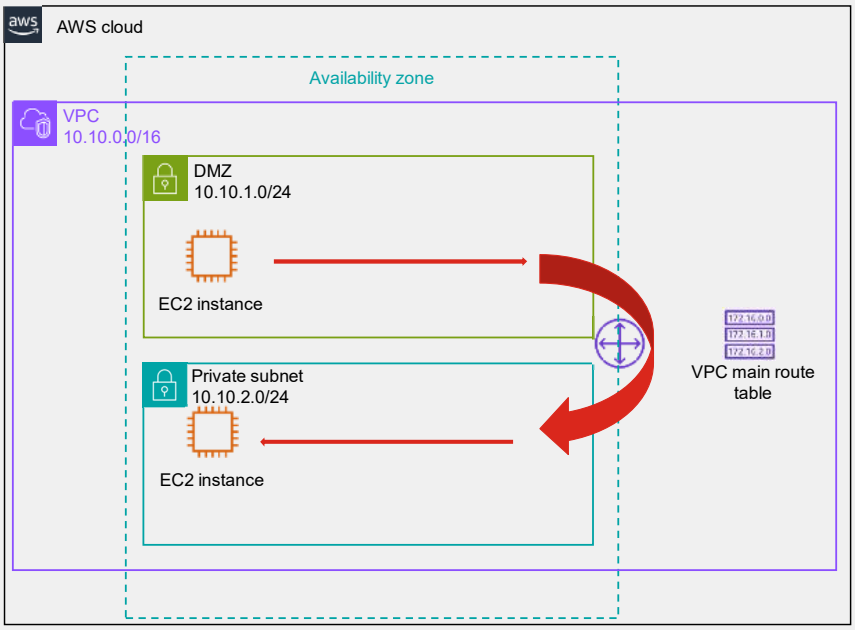

FORTINET Training Institute

 42

When an OS within an EC2 instance attempts to send packets to an IP address outside of its own subnet, it sends an ARP request to learn the MAC address of its default gateway, the intrinsic router. Then it sends packets to the layer 2 destination of the default gateway MAC address, but at layer 3, the destination of the IP address is in another subnet.

While it appears similar to on-premises IP networking in an Ethernet network with two subnets, the cloud network functions differently. As we learned, the intrinsic router is not a router that traffic relays through. It is a service running on hypervisors throughout the AWS infrastructure that routes the layer 3 destinations of the packets as soon as they are sent out of the ENI of an EC2 instance. The service finds the hypervisor of the destination VM and transfers the packet to that hypervisor. Upon arriving at the destination hypervisor, the packet is sent to the ENI of destination EC2 instance.

VPC Traffic Flow—Scenario 3

- To the internet:
  - Internet gateway associated with VPC
  - Route table associated with subnet
  - Public or elastic IP associated with the ENI of the EC2 instance

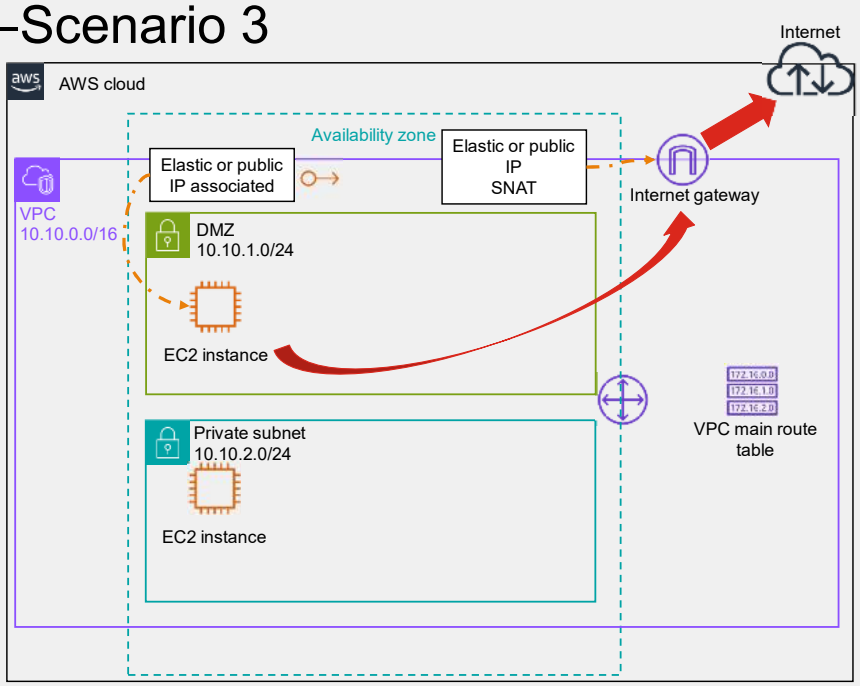

FORTINET Training Institute

 43

When an EC2 instance needs to connect to the internet, you must do the following to configure the handling of outbound traffic:

1. Configure and attach an internet gateway service to the VPC.
2. Create a route table, or edit an existing route table, with a default route that directs traffic destined to the internet toward the internet gateway.
3. Associate the route table with the subnet where the EC2 instance is connected to its ENI.
4. Associate the ENI of the EC2 instance with a public IP address.

VPC Traffic Flow—Scenario 3 (Contd)

- To the internet:
  - Internet gateway associated with VPC
  - Route table associated with subnet
  - Public or elastic IP associated with the ENI of the EC2 instance

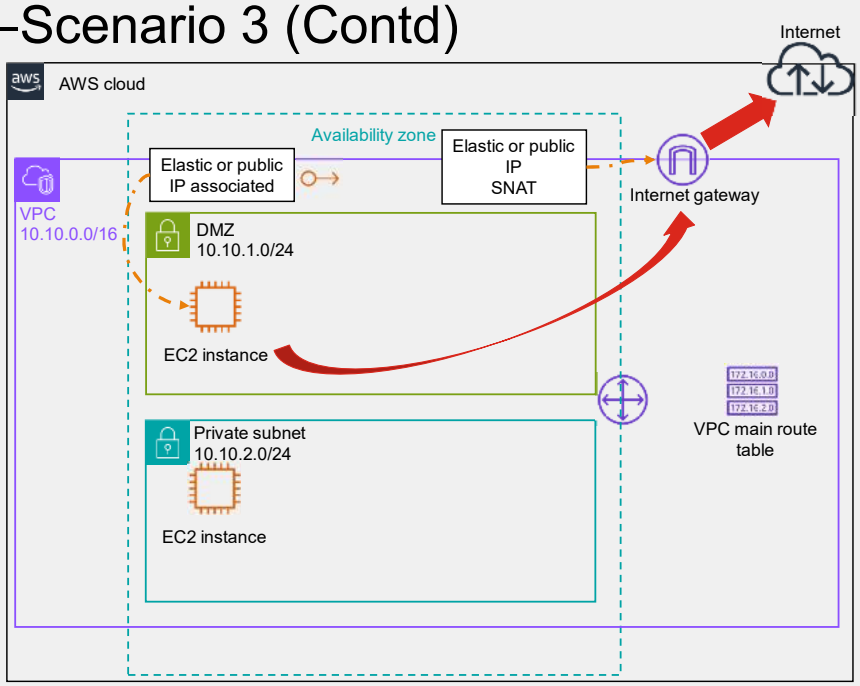

FORTINET Training Institute

 44

While the public IP address is associated with the ENI, it is important to note that the OS within the EC2 instance is not configured to use the public IP address. The EC2 instance should always use an IP address assigned to its ENI within the subnet that it is connected. In the example above, the EC2 instance in the subnet might have the IP address `10.10.1.4`.

When this EC2 instance sends a packet from its ENI with a destination IP address that is outside the VPC, a route table entry determines that this packet should be processed by the internet gateway. This routes the packet through an internet service provider (ISP) at the AWS data center. The internet gateway service performs source network address translation (SNAT) on the packet destined for the public IP address associated with the ENI as it is sent through to the ISP. When the resource on the internet replies to the public IP address, the internet gateway service performs destination network address translation (DNAT) on the reply to the `10.10.1.4` IP address of the ENI. After the gateway performs DNAT, the reply packet is transferred to this EC2 instance ENI where the OS on the EC2 instance receives it.

VPC Traffic Flow—Scenario 4

- From the internet:
  - Internet gateway associated with VPC
  - Route table associated with subnet
  - Public or elastic IP associated with EC2 instance ENI

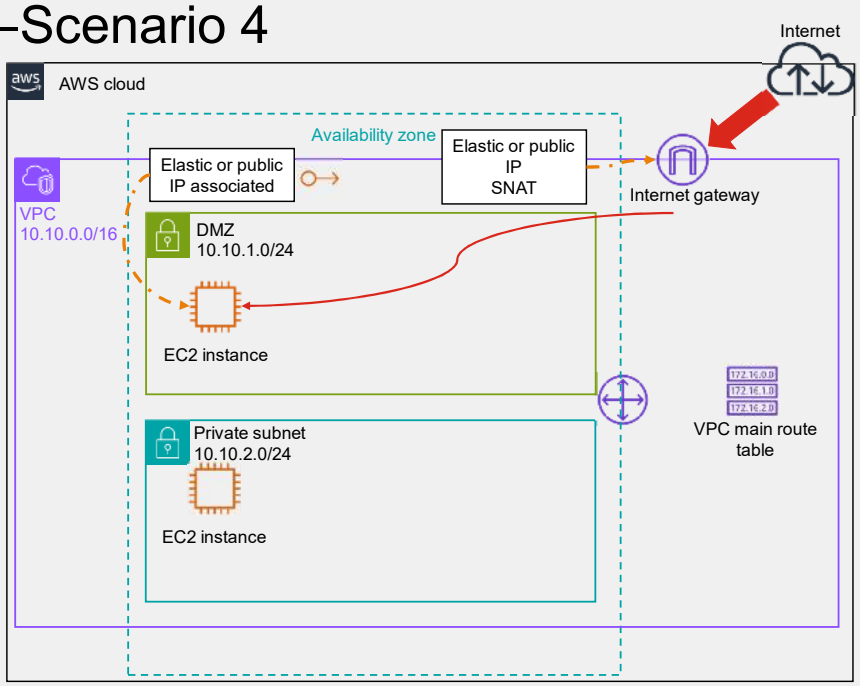

FORTINET Training Institute

 45

How does traffic initiated from the internet reach the EC2 instance? Traffic from the internet connects to the public IP address you just learned about, and then undergoes DNAT to the EC2 instance's ENI private IP address in the subnet, which it is connected to. This slide shows that the EC2 instance is connected to a VPC subnet named DMZ, and the instance might have an IP address such as `10.10.1.4`. As traffic connects to the public IP address, the internet gateway service performs DNAT on the traffic from the public IP address to the IP address `10.10.1.4`.

You must do the following to configure the handling of inbound traffic:

1. Configure and attach an internet gateway service to the VPC.
2. Create a route table or edit an existing route table with a default route that directs traffic destined to the internet towards the internet gateway.
3. Associate this route table with the subnet where the EC2 instance has its ENI connected.
4. Associate the ENI of the EC2 instance with a public IP address.

Layer 2—AWS Example for the Same Subnet Traffic

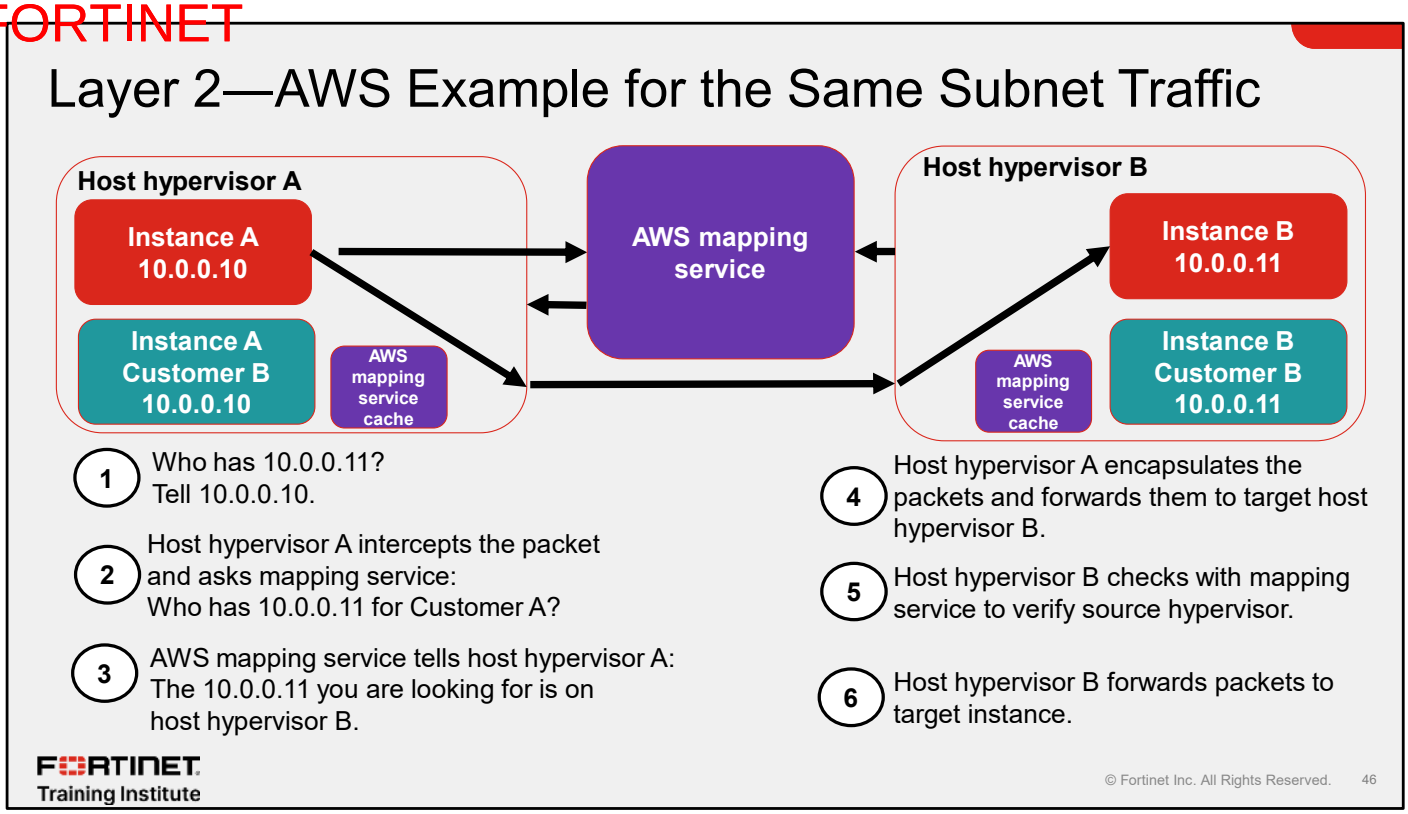

Now, you will learn about layer 2 networking in AWS cloud computing. Layer 2 networking works differently in cloud computing. How does instance A communicate with instance B? As computer nodes in a regular network, instance A sends an ARP request. Therefore, it must send a broadcast request for the destination MAC address. However, in a cloud environment there could be thousands of machines between two instances generating lots of broadcast traffic in cloud switches, which is very problematic.

The AWS mapping service minimizes large amounts of broadcast traffic in cloud computing. It contains all the MAC addresses and IP addresses of the subnet in a database accessible to all hypervisors in a region.

As shown on this slide, the AWS mapping service is responsible for capturing the request packet and replying with the correct MAC address of instance B. The AWS mapping service checks its database for the correct IP address and corresponding MAC address and then the traffic flows from the MAC address to the MAC address on instance B in a unicast fashion. No broadcast traffic travels over the network.

You must assign and declare all your VM IP addresses in the cloud portal. The cloud vendor console must sync IP address information with the VMs. If you add an IP address to a VM, you must add the IP addresses to the configuration of the cloud console. The cache service available on the physical host records all the information. If you change the IP address of the VM, the cache service may take some time to update, especially if you encounter any connectivity issues after changing the IP address of the host. Note that the cache service is time-limited.

Layer 2—Restrictions

- An instance receives traffic only on an IP address defined in the cloud console
- All static or virtual IP addresses on the VM must match the cloud console
- No traditional layer 2
  - Only packets destined for an IP address leave an instance; all other traffic is dropped. This means:
    - No FGCP
    - No gratuitous ARP or proxy ARP
    - No instant IP failover
    - No custom frames/ethertypes/layer 2 manipulation
    - No broadcast
    - No 802.1q
- All layer 2 modes are forbidden: transparent and virtual wire

FORTINET Training Institute

 47

In addition to not allowing broadcast traffic, layer 2 cloud networking sets other limitations. An instance receives the traffic only if the IP address is defined on the cloud console. If static or virtual IP addresses are configured on the VM, you must make sure that those IP addresses are also configured on the cloud console. In terms of layer 2 restrictions, there shouldn't be any traditional layer 2 traffic, such as FortiGate clustering protocol, gratuitous ARP, instant IP failover, and so on. Essentially, only unicast traffic is allowed. Also, no layer 2 modes are allowed in cloud computing, for example, transparent mode or virtual wire.

Routing Restrictions

- Traffic entering a virtual network always goes through a routing table that you can configure on the cloud console
  - This traffic passes directly to the target instance through the embedded router
- Traffic leaving a VM instance must have a route from the local subnet router or it is blackholed
- There is always an embedded router on every subnet
  - All VMs use the embedded router as the default gateway

FORTINET Training Institute

 48

There are additional limitations when it comes to layer 3 traffic in cloud networking. When traffic enters the virtual network, it must first go through the routing table, which is configured on the cloud console. At the same time, traffic leaving a VM instance must have a valid route from the local subnet router, otherwise, traffic is blackholed. Keep in mind that every subnet has an embedded router, and all VMs use the embedded router as the default gateway.

Knowledge Check

1. Which VPC component is required to provide internet access for instances with public IP addresses?
   - A. Transit gateway
   - ✓ B. Internet gateway

2. Which type of layer 2 traffic is allowed inside an AWS VPC?
   - ✓ A. Unicast traffic
   - B. Broadcast traffic

FORTINET Training Institute

 49

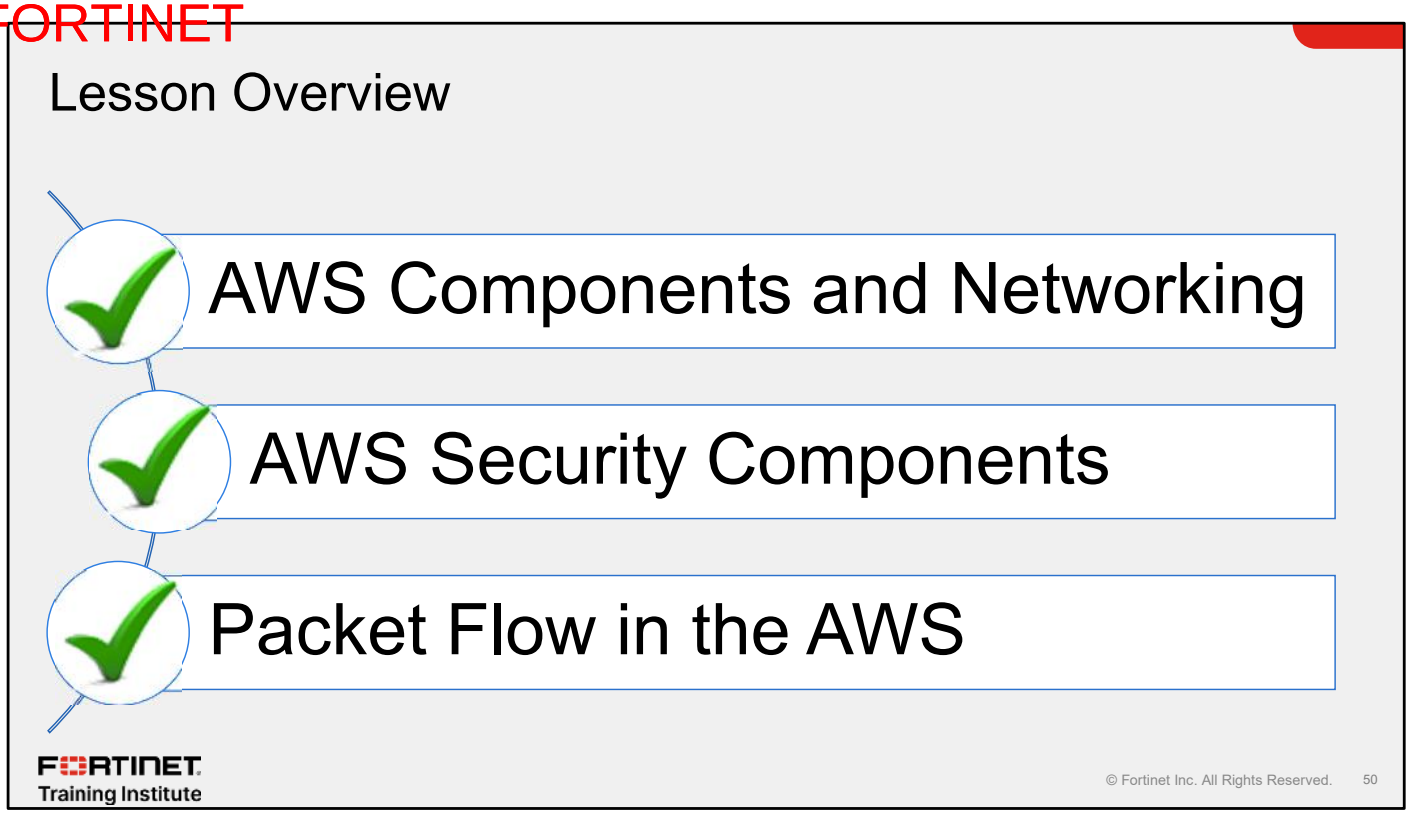

Good job! You now know about network packet flow in the AWS cloud.

Review

- ✓ Describe AWS service components
- ✓ Identify AWS core networking components
- ✓ Identify AWS security components
- ✓ Understand traffic flow in a virtual network
- ✓ Understand layer 2 traffic flow
- ✓ Understand routing and restrictions

FORTINET Training Institute

 51

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

By mastering the objectives covered in this lesson, you learned about the concept of the public cloud and how to use it in your network.

In this lesson, you will learn about the Fortinet solutions for Amazon Web Services (AWS).

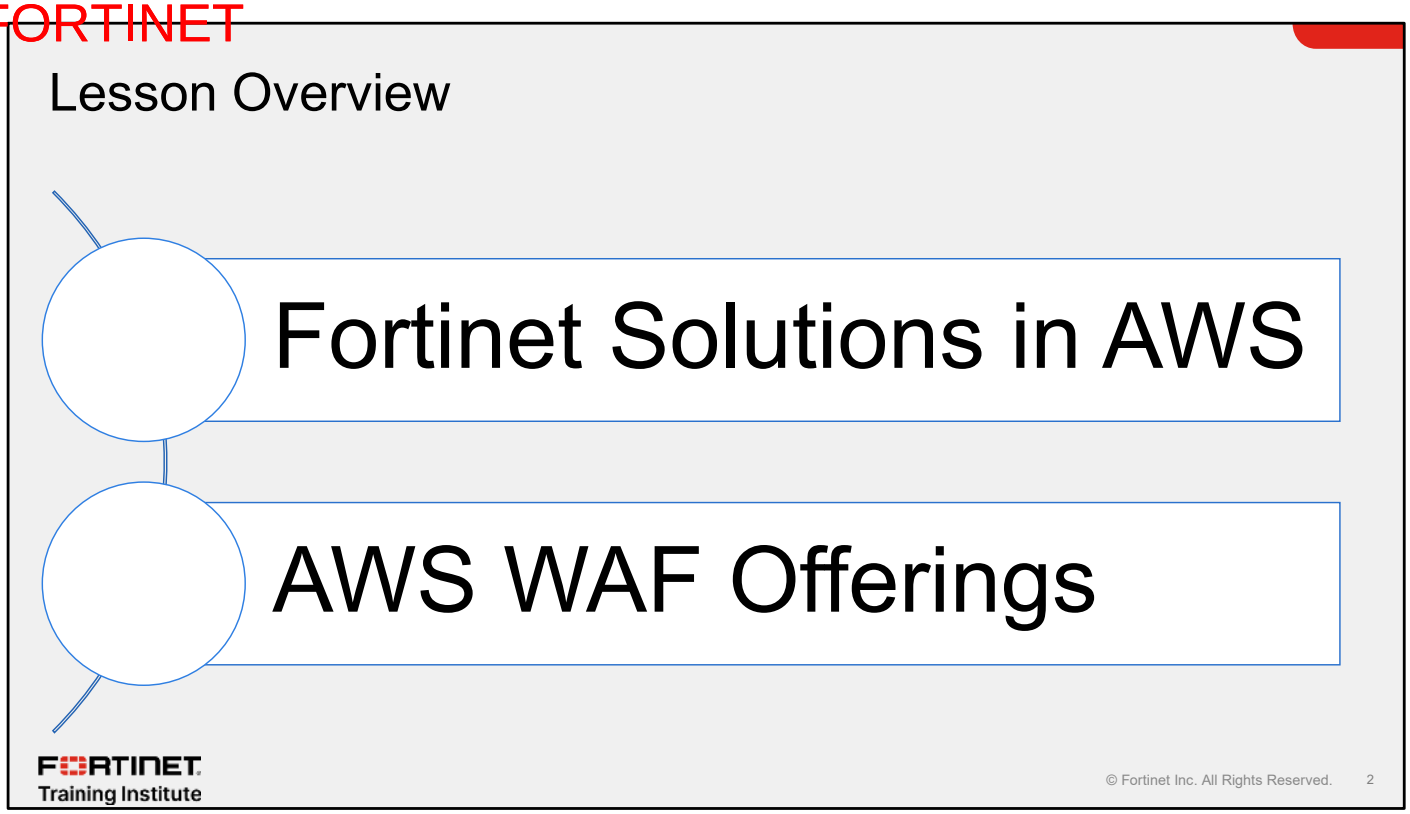

In this lesson, you will learn about the topics shown on this slide.

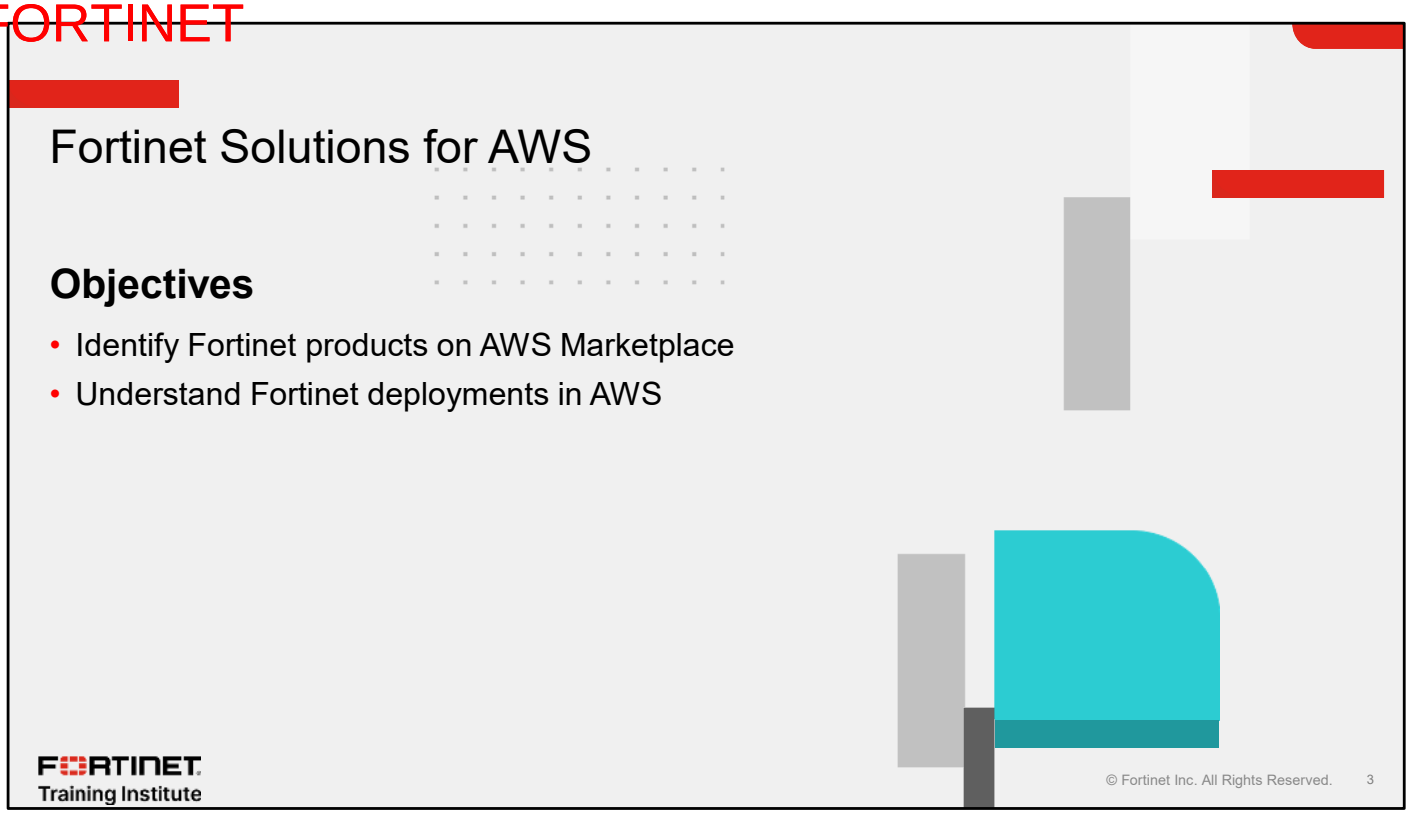

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding Fortinet solutions for AWS, you will be able to successfully use AWS with Fortinet solutions.

Fortinet Products on AWS Marketplace

FORTINET
on
aws marketplace

Product	BYOL	PAYG	SaaS
FortiGate	✓	✓	
FortiGate CNF		✓	✓
FortiWeb	✓	✓	
WAF Rules			✓
FortiWeb Cloud		✓	✓
FortiAuthenticator	✓		
FortiSandbox	✓	✓	
FortiADC	✓	✓	
FortiAnalyzer	✓	✓	
FortiManager	✓	✓	
FortiSIEM	✓		
FortiMail	✓		
FortiTester	✓		
FortiRecorder		✓	
FortiProxy	✓		
FortiVoice	✓		
FortiSOAR	✓		
FortiCNP		✓	✓

FORTINET Training Institute

 4

This slide shows the Fortinet solutions currently available in AWS. AWS is the most broadly supported cloud vendor for Fortinet products.

Keep in mind that the information shown can change at any time, based on the new support availability for Fortinet products.

By leveraging these solutions, you can enhance your organization’s security, networking, and management capabilities in their AWS cloud environments.

Fortinet Products on AWS Marketplace (Contd)

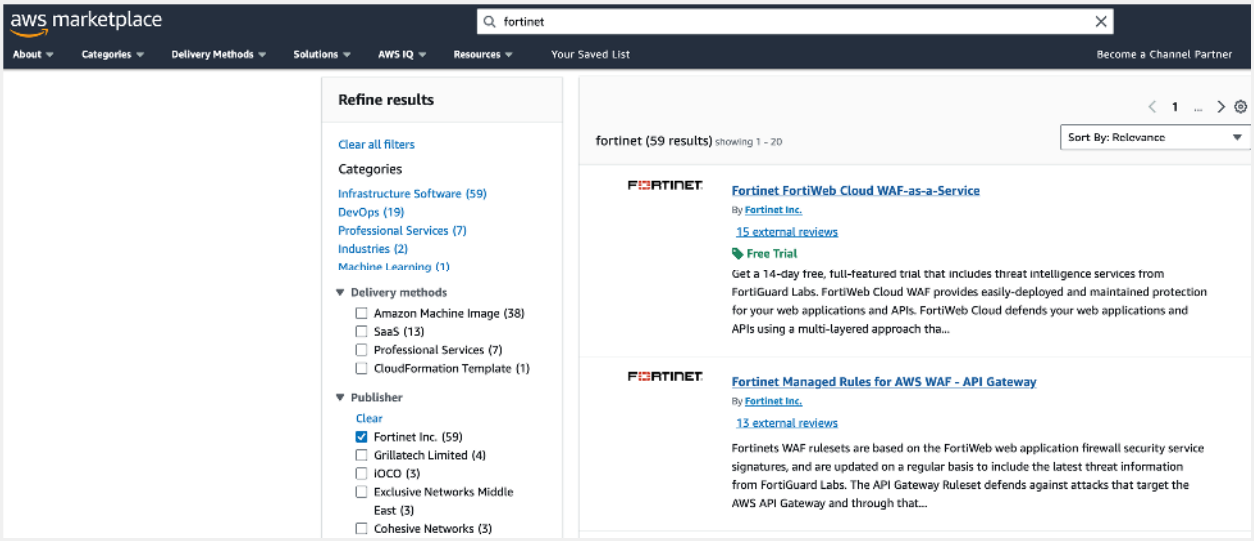

FORTINET Training Institute

 5

You can find all available Fortinet products on the AWS marketplace website.

This slide shows a search performed in the marketplace, looking for the string *fortinet*, and filtering the results by the publisher name of *Fortinet Inc.*

Fortinet solutions provide superior visibility, protection, and control for public cloud deployment options in AWS in the following key areas:

- Hybrid cloud security
- Cloud visibility and control
- Secure access VPN
- Cloud security services hub
- Container security
- Web application security
- Internet-based segmentation

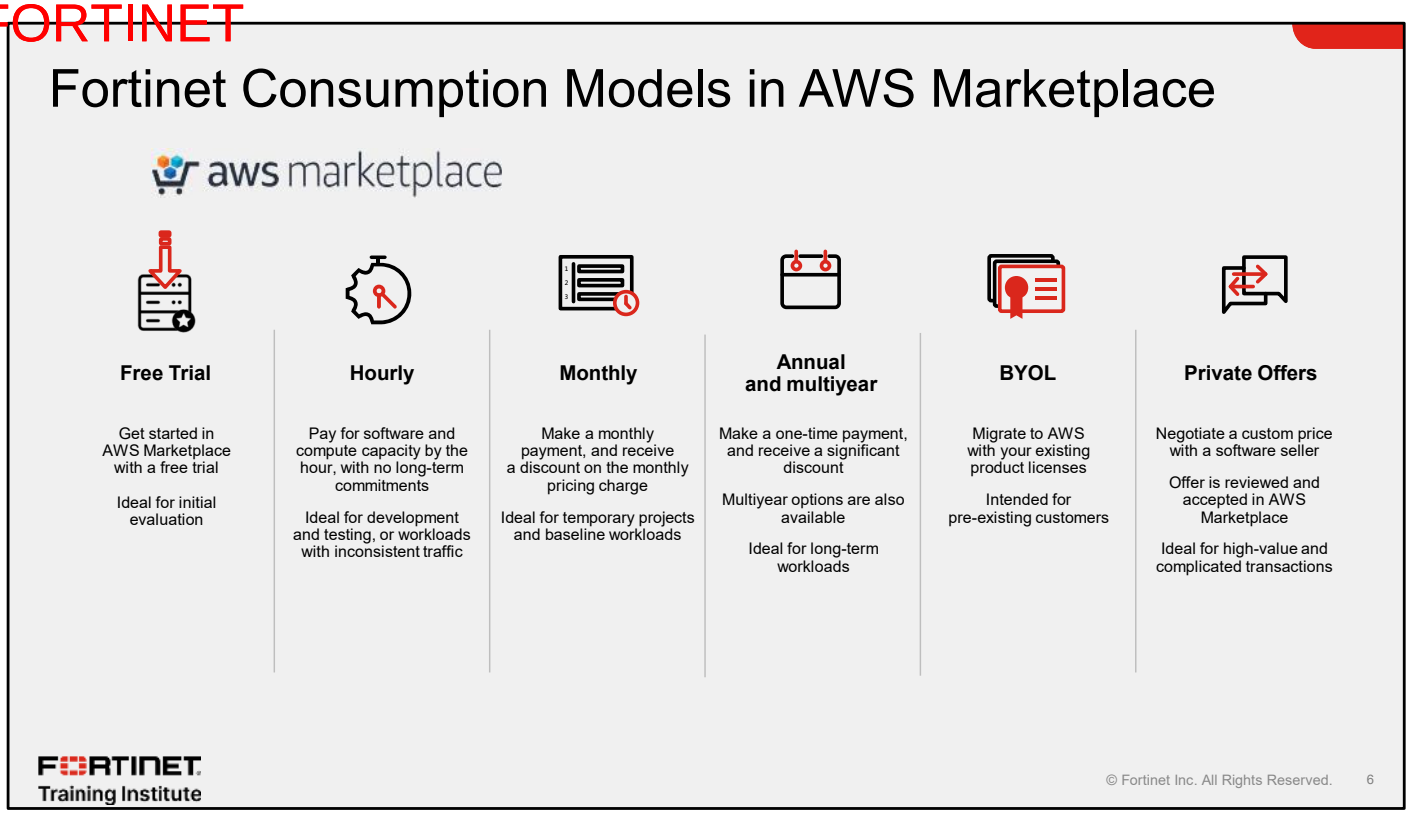

The licensing consumption models available for Fortinet products on AWS marketplace are:
- Free trial
- Hourly
- Monthly
- Annual and multiyear
- BYOL
- Private offers
- FortiFlex

Remember that Fortinet's licensing models and offerings may evolve over time, so it's crucial to check the official Fortinet website or contact Fortinet directly to get the most current information and guidance on licensing for AWS deployments.

Better Together

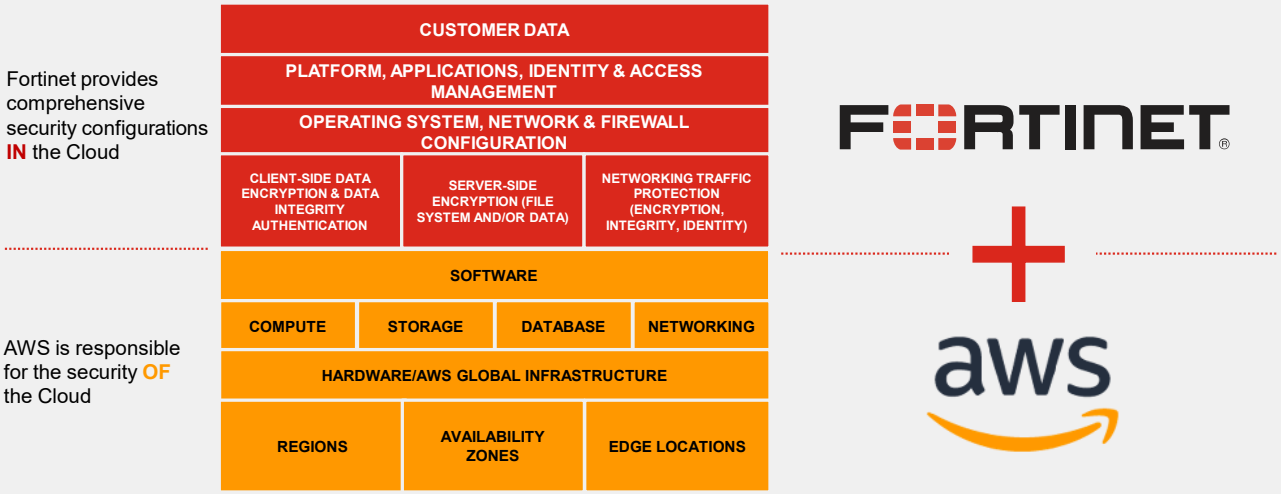

FORTINET Training Institute

 7

Fortinet products complement the security already provided by AWS native services. You can think of this collaboration as a joined shared responsibility model. Fortinet provides security in the cloud. AWS provides security of the cloud. The next few slides show how Fortinet products accomplish exactly that.

FortiGate for AWS

- Superior connectivity compared to AWS native services
- VPN
  - IPsec, SSL VPN
  - Additional supported topologies (full mesh and partial mesh)
  - OSPF over IPsec
  - SD-WAN support
  - Performance depends on instance type
  - Dynamic tunnel support using ADVPN
- SD-WAN
- Dynamic routing
  - BGP over IPsec
  - OSPF over IPsec
  - BGP over GRE

FORTINET Training Institute

 8

FortiGate VM is a powerful virtual instance that offers a wide range of advanced security, VPN, SD-WAN, and routing features. It offers robust capabilities that allow you to create complex configurations and provides granular control.
FortiGate for AWS provides critical features that are not available by default from AWS native services, such as VPN, SD-WAN, and dynamic routing functionality.

Compared to AWS cloud native services, FortiGate adds additional supported topologies for IPsec, such as full mesh and partial mesh. FortiGate also allows the use of OSPF over IPsec. SD-WAN is not supported by AWS natively, but it is supported when using FortiGate.

FortiGate also adds additional functionality when it comes to dynamic routing. You can use OSPF over IPsec, which is not supported by AWS cloud native services. All dynamic routing protocols must be encapsulated by either IPsec or GRE.

SDN Connector for AWS

- Dynamic address learning
- Scale up or down automatically

Security Fabric > External Connectors

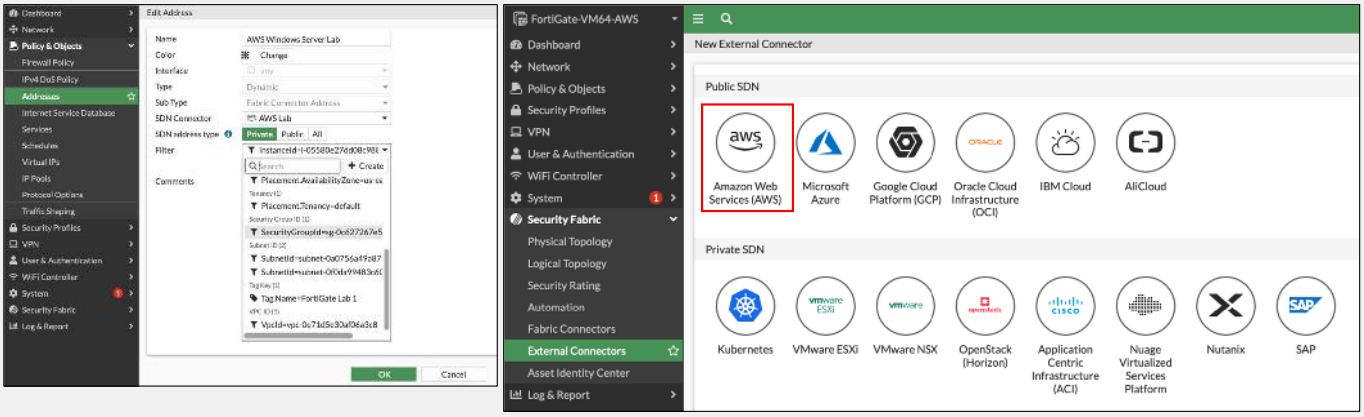

FORTINET Training Institute

 9

Fortinet expands its Security Fabric architecture through its fabric connectors to extend security visibility and management to the cloud.

Fabric connectors link into partner solutions through API integration points or through specialized engineering. The open design of the fabric connectors enables ongoing deep integration with a growing number of ecosystem components, and extends the Security Fabric capabilities into validated, third-party infrastructures.

Fortinet fabric connectors help automate security operations and policies through one-click integrations with partners, including AWS. You can pull information from AWS, addresses, VM names, and subnets, and then use this information to create firewall policies. Compared to Azure, you need less information to configure fabric connectors in AWS.

Automation Using AWS Lambda and FortiGate

Security Fabric > Automation > Action > AWS Lambda

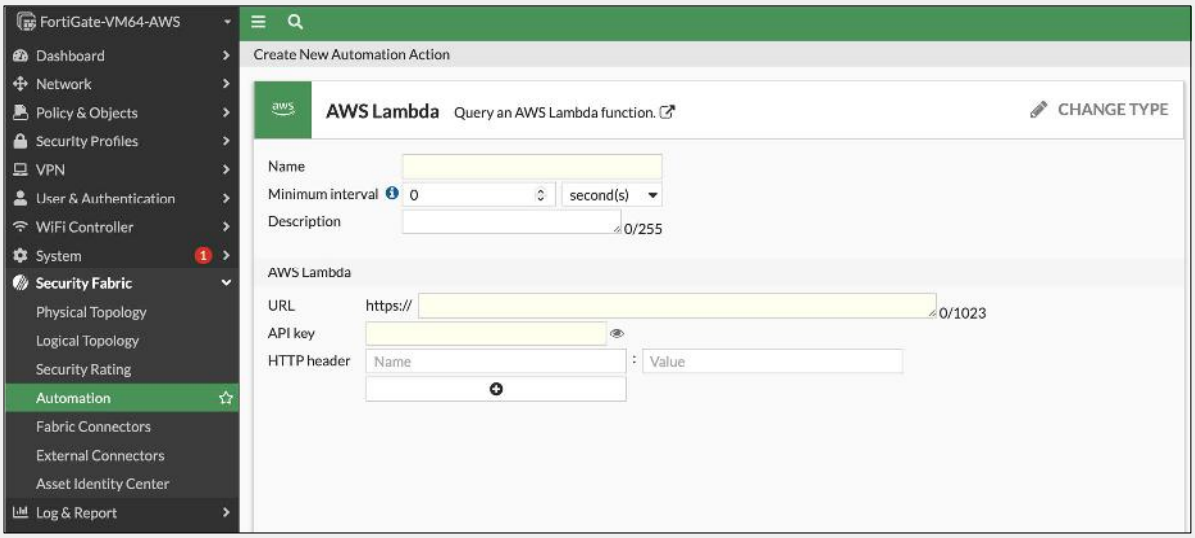

FORTINET Training Institute

 10

FortiGate has an automation stitches feature that you can combine with AWS Lambda or other vendors that invoke automation rules in the Fortinet Security Fabric. For example, you can use the Fortinet compromised host trigger feature with AWS Lambda to automatically quarantine any identified infected hosts in the network. Another example where Lambda is used are FortiGate HA deployments. Lambda is responsible to ensure a successful failover. There are many automation triggers that you can use with AWS Lamda.

GuardDuty Integration With FortiGate

- Automates security remediation for workloads running in AWS

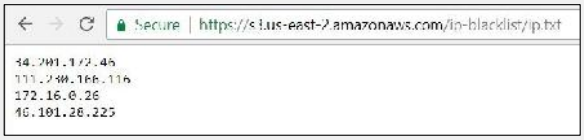

- Accelerates time-to-protection for threats detected by the AWS service and automates the creation of network firewall rules in FortiGate to mitigate threats

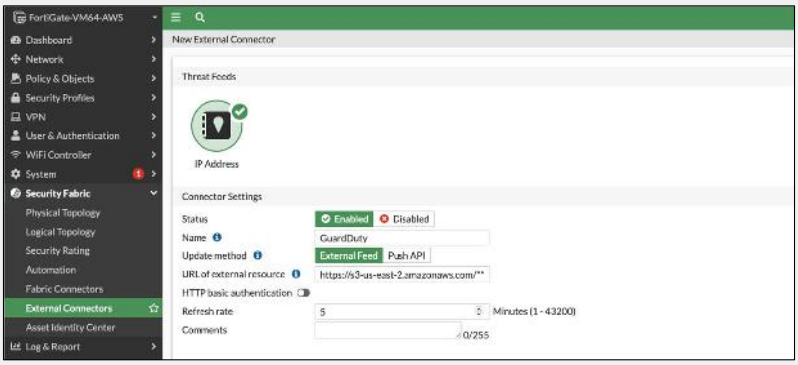

- Reduces dependency on manual incident response and human intervention

FORTINET Training Institute

 11

You already learned about Amazon GuardDuty in the previous lesson. It is a managed threat detection service offered by AWS that continuously monitors your AWS environment for malicious activity and unauthorized behavior. You can integrate Amazon GuardDuty with FortiGate to enhance your cloud security posture by leveraging GuardDuty's threat detection capabilities and FortiGate's firewall and security features.

Amazon GuardDuty integration with FortiGate automates security remediation for workloads running in AWS. It accelerates time-to-protection for threats detected by the AWS service and automates the creation of network firewall rules in FortiGate to mitigate threats. It also reduces dependency on manual incident response and human intervention. You can configure this under **Security Fabric** > **External Connectors** > **Create New** > **Threat Feeds** > **IP Address**.

FortiSandbox for AWS

- Control
  - Windows EC2s
- Extended scalability
- Addresses limitation in AWS
- VMs started and stopped outside FortiSandbox
- Most of the cost is associated with EC2 instances

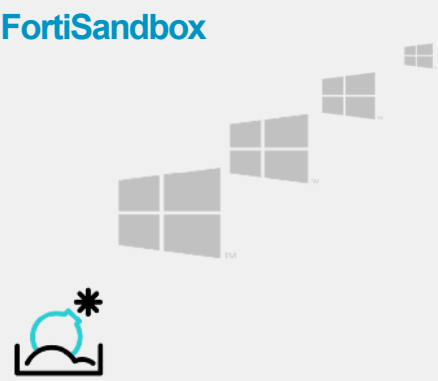

FORTINET Training Institute

 12

Now, you will learn about another one of the AWS-supported Fortinet products, FortiSandbox. FortiSandbox is not a hypervisor in AWS—it is simply a manager that analyzes the results of the sandboxing process. FortiSandbox deploys new EC2 instances with custom Windows VMs. FortiSandbox then uses these VMs to send malware to them. When the malware finishes running, the results are captured for analysis. FortiSandbox for AWS does not need more resources because it performs management and analysis tasks only. Note that the cost of EC2 instances varies based on the number of instances deployed, the size of the instances, and the duration of the running time.

FortiSandbox for AWS (Contd)

- Enables organizations to defend against advanced threats natively in the cloud
- Automated zero-day, advanced malware detection and mitigation
- Works alongside network, email, endpoint, and other security to leverage scale with complete control
- Can be installed as a standalone zero-day malware behavior analysis system
- Integrates with existing FortiGate, FortiMail, or FortiWeb AWS instances

FORTINET Training Institute

 13

FortiSandbox for AWS enables organizations to defend against advanced threats natively in the cloud.

FortiSandbox provides several important benefits, including:
- Automated zero-day, advanced malware detection and mitigation
- An addition to network, email, endpoint, and other security, or an extension to on-premises security architectures, that leverages scale with complete control

You can install FortiSandbox as a standalone zero-day malware behavior analysis system. Also, you can integrate FortiSandbox with existing FortiGate, FortiMail, and FortiWeb AWS instances.

Knowledge Check

1. Which AWS feature integrates with FortiOS to automatically quarantine any identified infected hosts in the network?
   - ✓ A. AWS Lambda
   - B. AWS WAF

2. Fortinet is responsible for providing security OF the cloud.
   - A. True
   - ✓ B. False

FORTINET Training Institute

 14

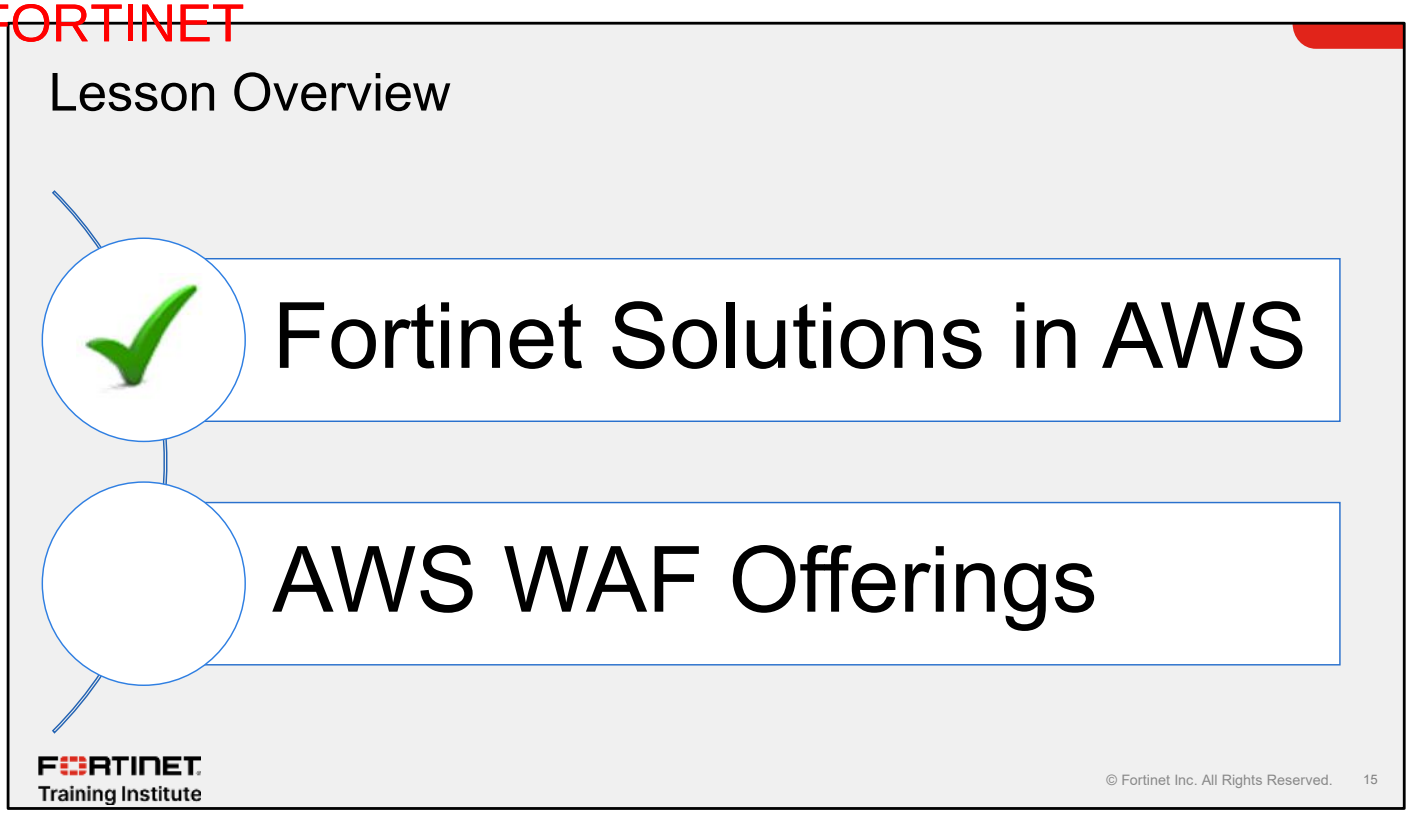

Good job! You now know about Fortinet solutions in AWS.

Now, you will learn about the different WAF offerings in AWS.

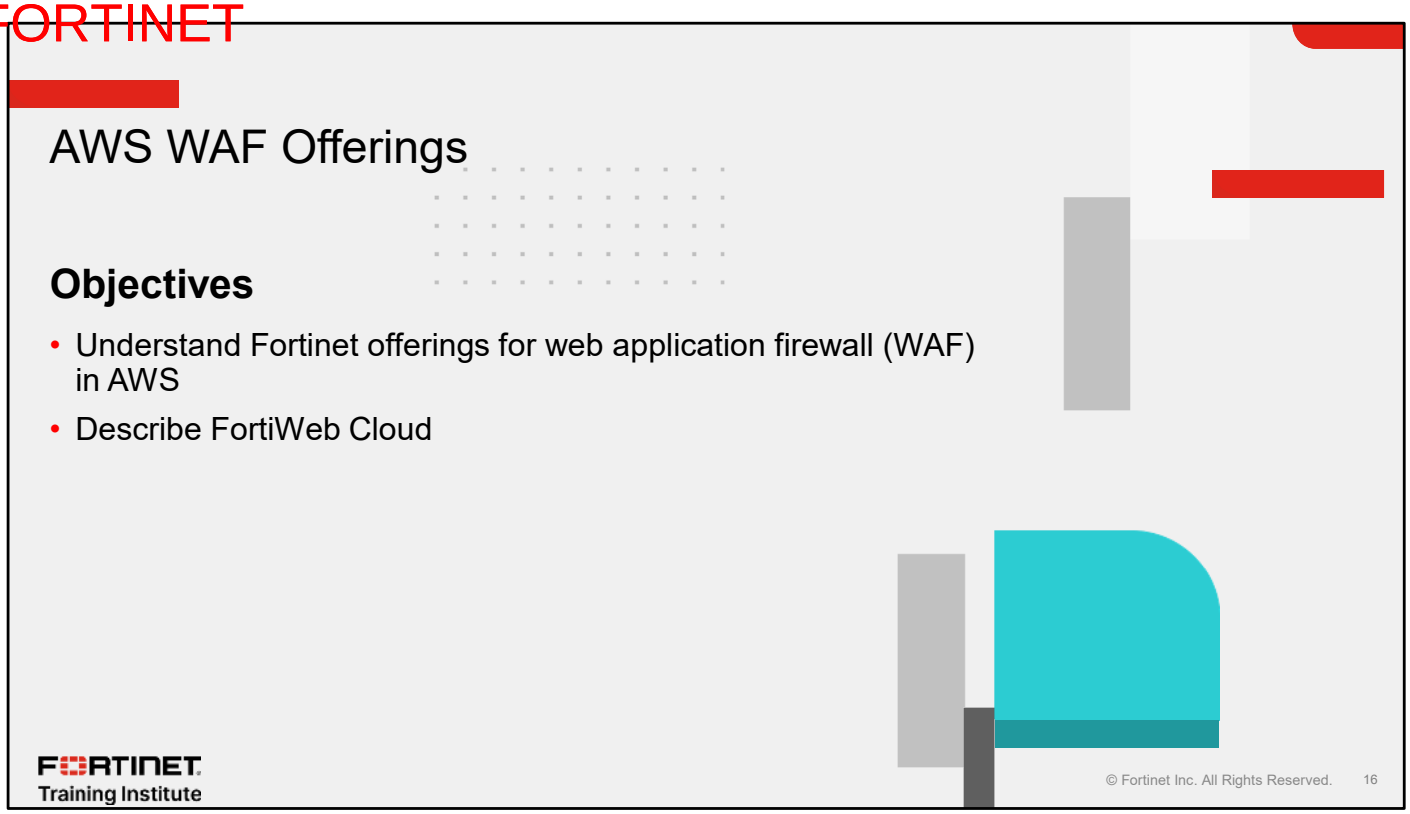

After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding AWS web application firewall (WAF) offerings, you will be able to successfully deploy various WAF solutions.

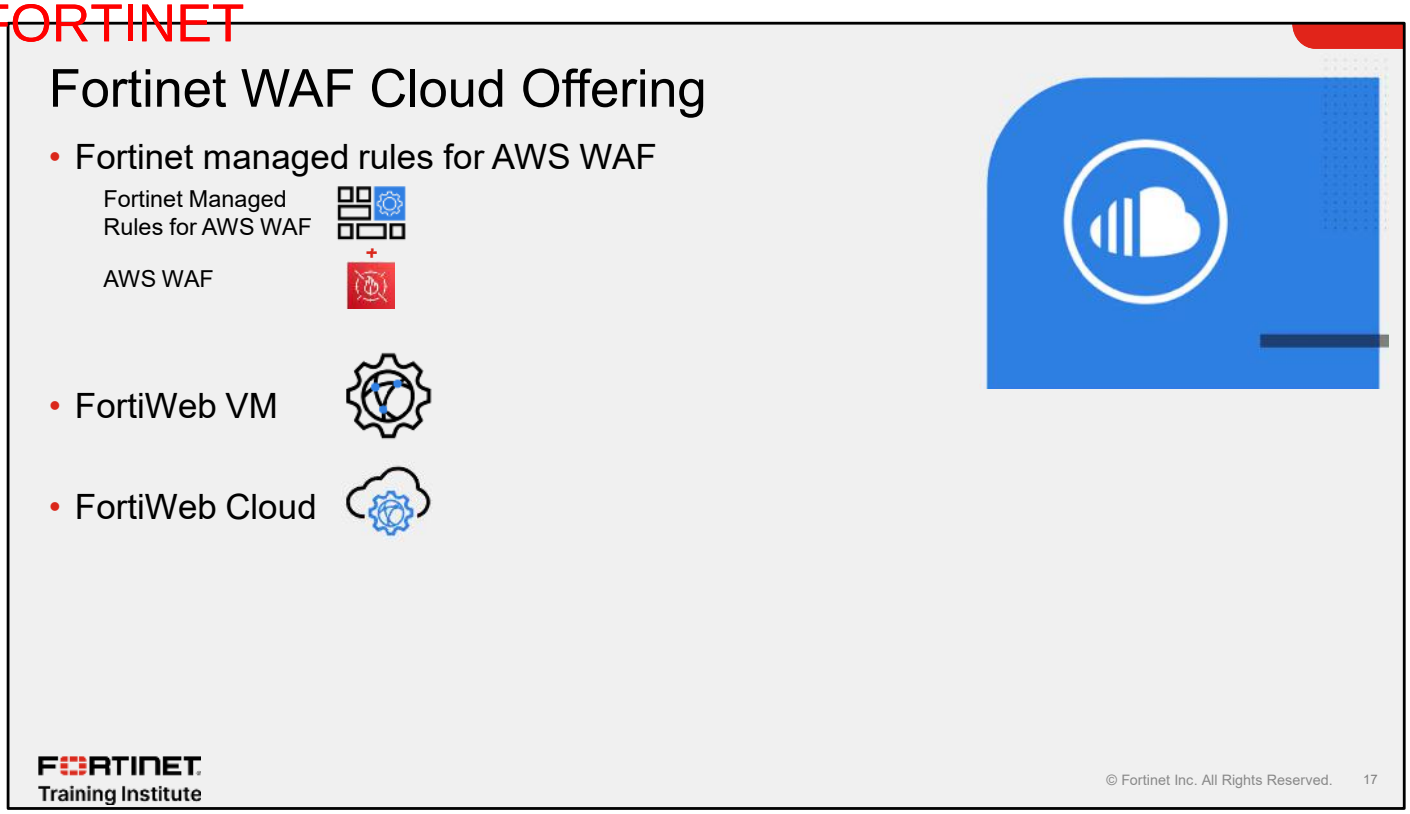

There are different Fortinet offerings that can provide WAF protection in AWS. For example, you can deploy a FortiWeb VM inside the virtual private cloud (VPC). One of the drawbacks in this scenario is that you can protect only applications going through the VPC.

You can also use FortiWeb Cloud, which is a WAF-as-a-Service hosted by Fortinet, that runs in AWS. You can use FortiWeb Cloud to protect applications that are internet facing. For example, you can have your DNS records pointing to the service, and FortiWeb Cloud will filter all web traffic before it is forwarded to the web application. The web application does not have to be hosted behind FortiWeb Cloud. It can be located anywhere.

AWS WAF Partner Rule Basics

- AWS offers basic SQLi and XSS in their current WAF offering
- Additional service to augment AWS WAF
- AWS partners with WAF vendors to offer prepackaged rule sets
- Can subscribe to up to three times the partner rule sets
- Actions at rule set level: log, alert, block
- Customer benefits:
  - Additional WAF protections from leading WAF vendors
  - Protection that is guaranteed to be up-to-date with latest signatures
  - Simplified WAF setup and management
  - Convenient availability on AWS Marketplace

FORTINET Training Institute

 18

AWS WAF partner rule groups are subscription-based, WAF signatures offered by third-party vendors to augment the basic WAF protection offered by the Amazon WAF product. These new rule groups allow AWS WAF customers to choose prepackaged WAF rules from leading IT security providers. Until now, AWS offered only SQL injection and XSS protection. With partner rule groups, vendors offer protection from a wide variety of application layer attacks, packaged in a variety of security rule sets. Some customer benefits include the following:

- Additional WAF protection from leading WAF vendors
- Protection that is guaranteed to be up-to-date with the latest signatures
- Simplified WAF setup and management
- Convenient availability on AWS Marketplace

Fortinet Managed Rules for AWS WAF

- Focus on building and delivering applications, not managing security rules
- Highlights and benefits
  - Additional layers of WAF protection
  - Updated automatically
  - No user intervention required
  - Add-on to AWS WAF

Powered by FortiGuard Labs	Deployed Directly from AWS Marketplace	Configurable in the AWS WAF Console



FORTINET Training Institute

 19

FortiWeb rule sets are additional security signatures that you can use to enhance the protection included in the base AWS WAF product without adding management or architectural complexity. They are based on FortiWeb security service signatures and are updated on a regular basis to include the latest threat information from FortiGuard Labs.

The highlights of Fortinet managed rules for AWS WAF are: additional layers of WAF protection; the rules are updated automatically without user intervention, and it serves as an add-on to the AWS WAF.

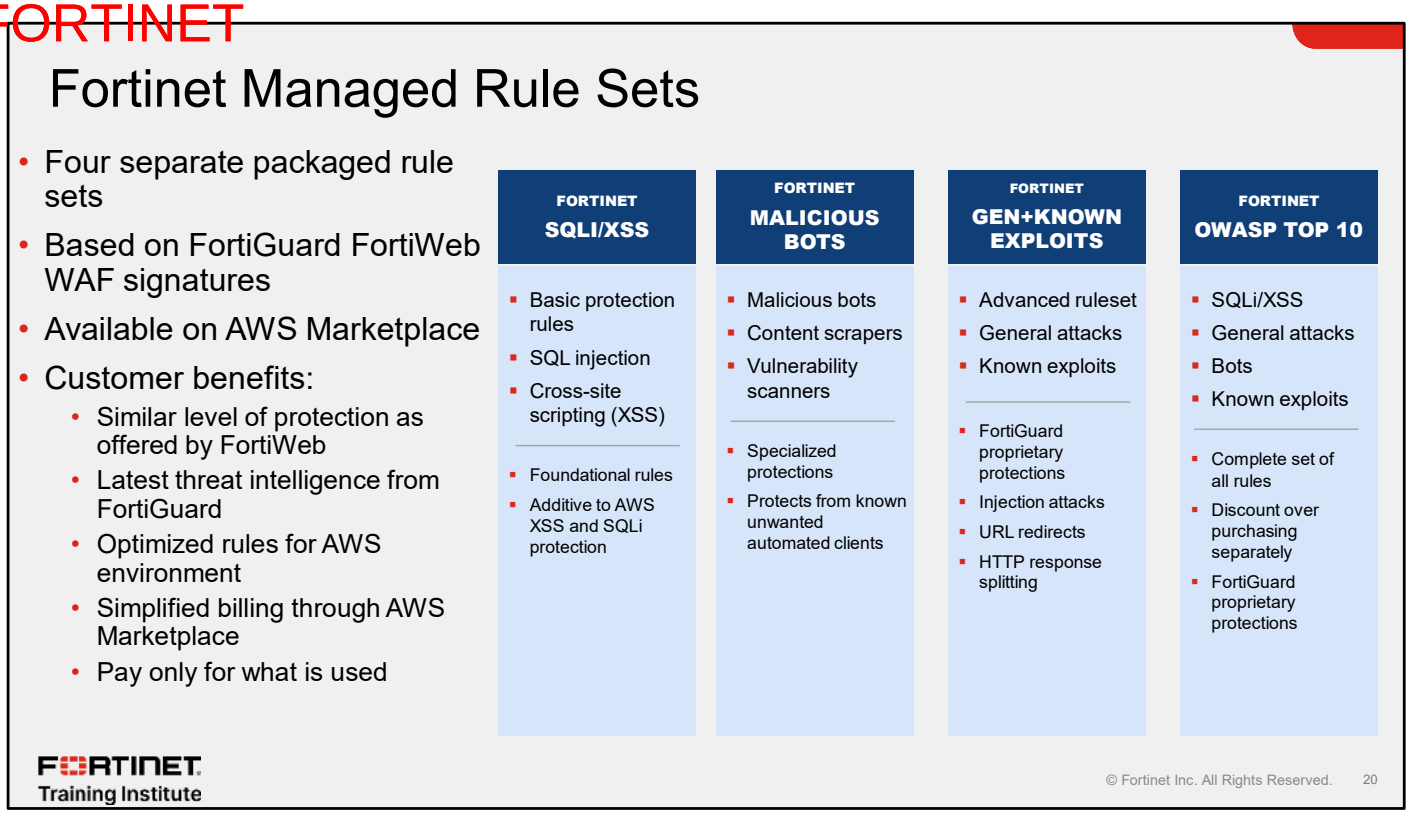

You can also purchase additional rule packages. There are four separate packaged rule sets based on FortiGuard FortiWeb WAF signatures, which are available on AWS Marketplace. These rule sets offer the same level of protection as WAF signatures on FortiWeb WAF devices (when combined, and all rules are used). Some of the benefits of Fortinet managed rule sets include:

- Access to the latest threat intelligence from FortiGuard
- Optimized rules for the AWS environment
- Simplified billing through AWS Marketplace
- Paying only for what is used

Fortinet Managed Rules for AWS WAF

- Listings accessible from AWS Marketplace

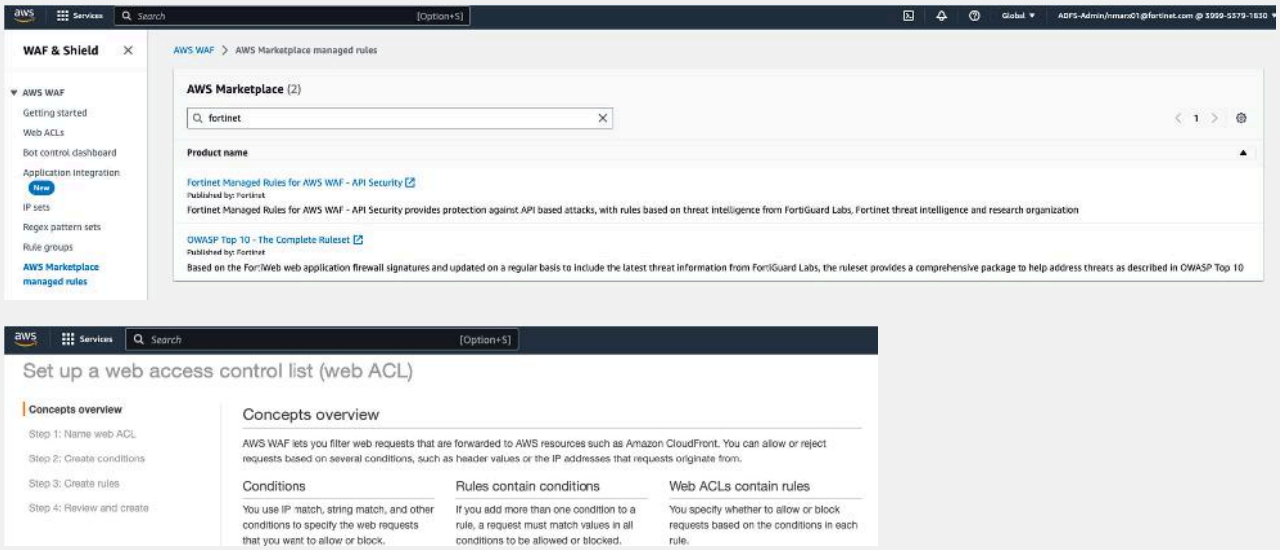

FORTINET Training Institute

 21

This slide shows how Fortinet managed rules for AWS WAF appear on AWS Marketplace. You can purchase WAF packages from AWS Marketplace and enable them in the WAF configuration.

FortiWeb vs. AWS WAF Rules

Feature	FortiWeb	AWS WAF partner rules
Web App Attack Signatures	Yes	Yes
WAF Subscription (FortiGuard)	Yes	Yes
IP Reputation (FortiGuard subscription)	Yes	No
Layer 7 DoS Protection	Yes	No
Bot and known search engine identification/protection	Yes	Yes (Partial)
Captcha	Yes	No
HTTP RFC Validation	Yes	No
Cookie Security	Yes	No
Antivirus/Antimalware	Yes	No
Behavioral Web App Attack Detection	Yes	No
Attack Correlation (protection from scanners, crawlers, scrapers)	Yes	No
Web App Vulnerability Scanner	Yes	No
Attack Alert Tuning	Yes	No
Web Defacement Protection	Yes	No
User and Device Identification	Yes	No
Brute Force Protection	Yes	No
Authentication Offload	Yes	No
Site Publishing and SSO	Yes	No
Meets PCI 6.6 Compliance	Yes	Yes
SSL Inspection	Yes	Yes

FORTINET Training Institute

 22

This slide shows a comparison between FortiWeb and AWS WAF partner rules. As you can see, there are some limitations to the AWS WAF partner rules. For example, there is no malware protection in AWS WAF partner rules because there is no engine to protect malware. Additionally, you cannot incorporate AWS WAF into the Fortinet Security Fabric, while you can incorporate FortiWeb.

WAF Product Positioning

	AWS WAF Partner Rules	FortiWeb
Primary function and focus	Skinny, simplified WAF for applications hosted on AWS using AWS WAF	Dedicated WAF (full feature set including behavior detection, customizable signatures, correlation)
Basic WAF (signatures, IP reputation, and so on)	Yes	Yes
Advanced WAF (behavioral scanning, correlation)	No	Yes
Up-to-date WAF signatures (FortiGuard)	Yes (optimized signatures)	Yes (subscription)
Customizable rules and whitelisting	No (negative security model)	Yes (positive and negative model)
Sample use cases	• 1-2 small applications • Mid-sized enterprise • Application hosted on AWS • Must use AWS WAF as base	• Mission-critical web applications • All segments including carrier/MSSP • Applications hosted any location
Pros	• Easy to deploy and manage; convenient • Pay for what is used • Optimized FortiGuard signatures	• Dedicated WAF solution • Zero-day protections • Advanced features
Cons	• Expensive for large applications • No zero-day protection • Only available for AWS WAF • No rule customizations	• Separate appliance • Increased investment • Increased setup and management
Availability	• AWS Marketplace only	• HW, VM, AWS (BYOL and On Demand

FORTINET Training Institute

 23

This slide shows WAF product positioning. It compares services between AWS WAF partner rules and FortiWeb.

Subscriptions and Services

- Managed IPS rules for AWS network firewall subscription
  - Filters malicious traffic
  - Powered by FortiGuard
- Consulting Service by Fortinet for:
  - Network and application security
  - Cloud security posture assessment
  - JumpStart for FortiGate and FortiWeb VMs, FortiCNP, FortiGate CNF

FORTINET Training Institute

 24

You can support the AWS network firewall with Fortinet-managed intrusion prevention system (IPS) rules based on the latest threat information from FortiGuard labs to filter malicious traffic at the perimeter of your VPC. The enterprise subscription is designed for customers with multiple VPCs in two or more regions. Fortinet IPS rules protect against:

- Application, network, and IoT vulnerabilities
- Malware detection
- Server and operating system vulnerabilities
- Web client vulnerabilities
- Web application vulnerabilities
- Web server vulnerabilities

Fortinet also provides the following consulting services:

- Network and application security consulting provides design, architecture, automation, and implementation services for dynamic, highly available, and scalable network (FortiGate) and application security (FortiWeb) in AWS. It also provides architecture and implementation guidance for AWS networking and FortiGate integration with AWS transit gateway, or autoscaling, or both.

- Cloud security posture assessment discovers a security posture baseline and receives actionable recommendations to reduce your risk profile and remediate existing vulnerabilities and misconfigurations.

- JumpStart consulting service is a five-day engagement to help you design the optimal architecture and successfully deploy and configure FortiGate, FortiWeb VMs, FortiCNP, or FortiGate CNF. This service lowers your total cost of ownership (TCO) with training on how to best support these products and ensures security best practices to avoid architectural missteps and reduce business risks. For more information, visit `https://aws.amazon.com/marketplace`.

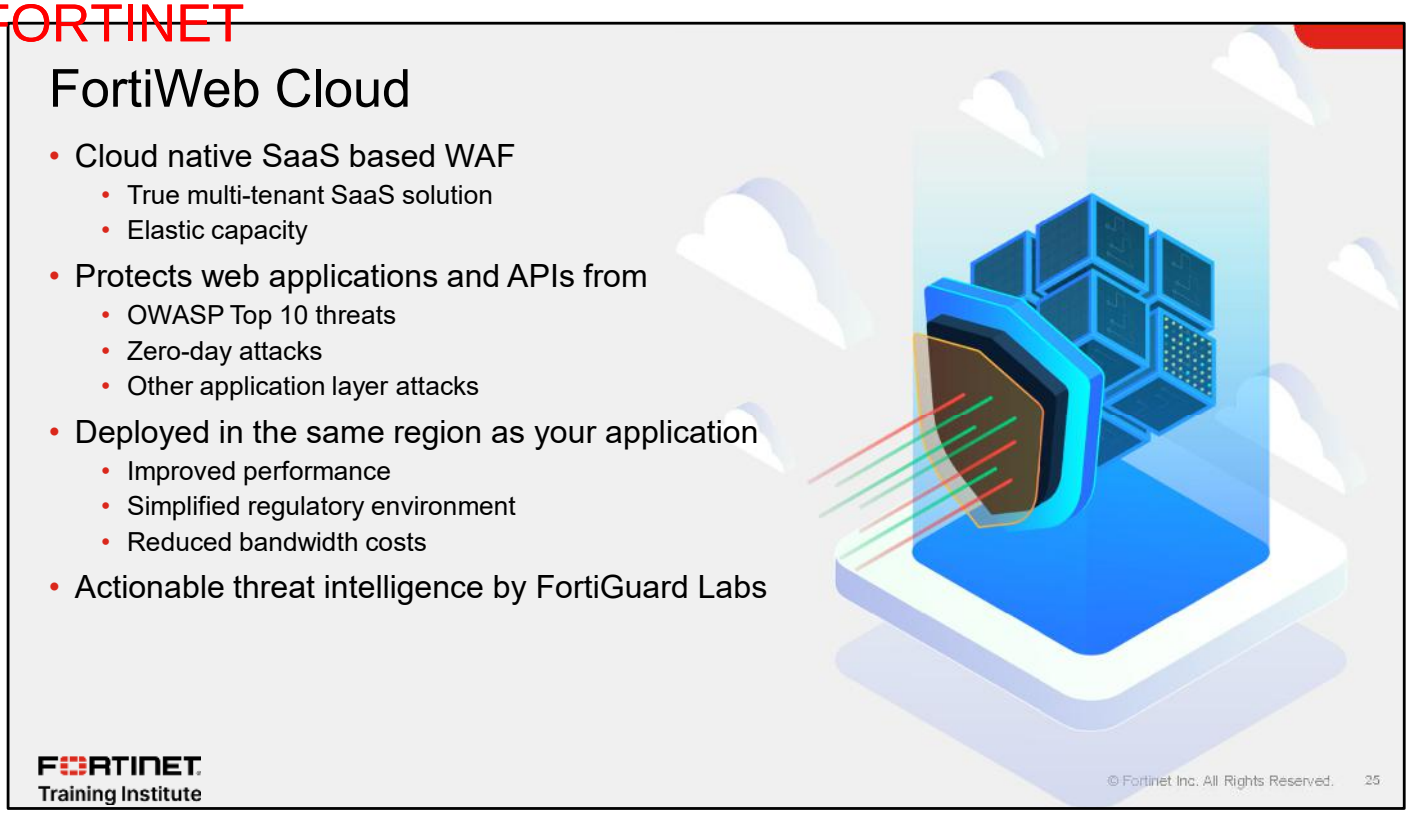

FortiWeb Cloud is a WAF-as-a-Service solution. Unlike FortiWeb, FortiWeb Cloud is a cloud native service allowing for easy and rapid deployment thereby efficiently and quickly providing protection. FortiWeb Cloud is not simply run on a few FortiWeb VMs, it is a true multi-tenant SaaS solution allowing for elastic capacity. This allows for a cost-effective way to protect small applications all the way up to enterprise-level applications.

Just like with FortiWeb, FortiWeb cloud protects your web applications and APIs from OWASP TOP10 threats, zero-day attacks, as well as other application layer attacks.

The service is deployed in the same region as your web application, which has the benefit of improved performance, a simplified regulatory environment, and reduced bandwidth costs.

FortiWeb Cloud is backed up by FortiGuard Labs providing a continuously updated service to protect web applications.

FortiWeb Cloud Deployment and Traffic Flow

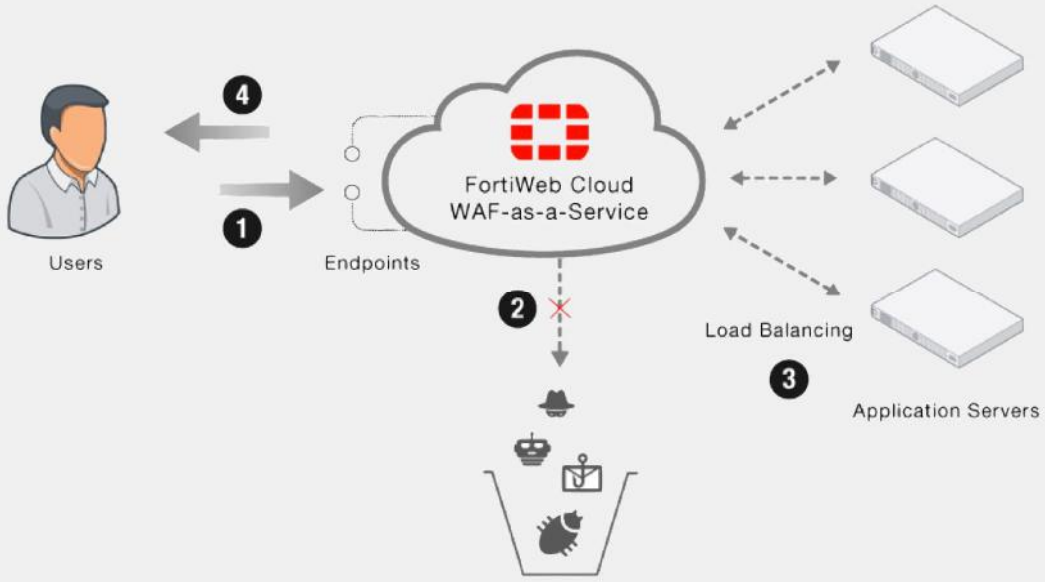

FORTINET Training Institute

 26

Before deploying FortiWeb Cloud, it is important to understand the traffic flow between the clients, FortiWeb Cloud, and the application servers.

As depicted in the illustration on this slide:

1. When users visit your application, the traffic is directed to the endpoints on FortiWeb Cloud.
2. FortiWeb Cloud filters the incoming traffic from users, blocking the OWASP Top 10 attacks, zero-day threats, and other application layer attacks.
3. Legitimate traffic arrives at origin servers. A load balancing algorithm is used to distribute traffic among servers.
4. When FortiWeb Cloud sends responses to your users, it obfuscates sensitive data such as the credit card number and other information that are likely to be used by hackers to damage your business.

You will deploy FortiWeb Cloud in the lab.

Subscribing to FortiWeb Cloud

- Steps to subscribe to FortiWeb Cloud:
  - Search for **FortiWeb Cloud WAF-as-a-Service** in AWS Marketplace
  - Review the different purchase options
  - Click **Subscribe**. After a few minutes you will be prompted to set up your account
  - Use your account, or create a new one, to log in to FortiCloud
  - Your subscription is automatically associated with your FortiCloud account

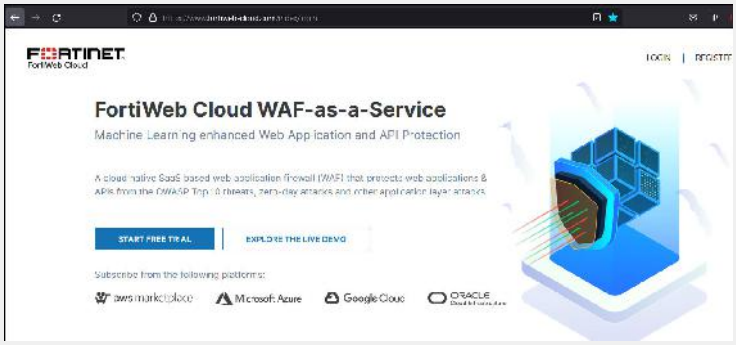

FORTINET Training Institute

 27

The process of subscribing to FortiWeb Cloud starts in AWS Marketplace and ends in FortiCloud. You can use your existing FortiCloud account or create a new one during the process.

Your AWS account is automatically associated with the FortiCloud account when you log in to FortiWeb Cloud.

Onboarding Applications

- Once you have logged in to FortiCloud, you can start onboarding applications:
  - Add the websites you want to protect, using their domain names
  - Add up to ten domains
  - You can use wildcards, for example: *.appdomain.com

FORTINET Training Institute

 28

After you finish the subscription process you can start onboarding applications from FortiCloud. Add the websites you want to protect by their domain names.

You can add up to ten domains for each application, and you can use wildcards to include domains sharing the same namespace with a single entry.

FortiCloud must be able to resolve the domain names you add. Keep in mind that the information for newly registered domains may take a few minutes to propagate through the global DNS servers.

Onboarding Applications (Contd)

- FortiWeb Cloud chooses the closest region and scrubbing center where it will be deployed, based on:
  - The IP address of the application
  - The cloud provider used by the application if hosted in AWS, Azure, OCI, or Google Cloud
  - The closest AWS center if the application is hosted on any other hosting platform
- Enable CDN to dynamically cache the application data in the scrubbing center closest to users
  - Note that using CDN may increase the operation cost

FORTINET Training Institute

 29

By default, FortiWeb Cloud chooses the closest region and scrubbing center where it is deployed based on:

- The IP address of the application.
- The same cloud provider used by the application if hosted in AWS, Azure, OCI, or Google Cloud.
- The closest AWS center (N. Virginia or Frankfurt) if the application uses any other hosting platform.

You can enable the content delivery network (CDN) to dynamically cache the application data in the scrubbing center nearest to users. When users request data from your application, they can be directed to the nearest scrubbing center and rendered with the requested data faster.

Although you can enable CDN for free, its use may increase costs due to the traffic expenses. This is not always the case, and it depends on the location of the users and the location of the data they access.

Onboarding Applications (Contd)

- Using the information displayed on FortiCloud, you must create DNS records in the domain server that hosts your application

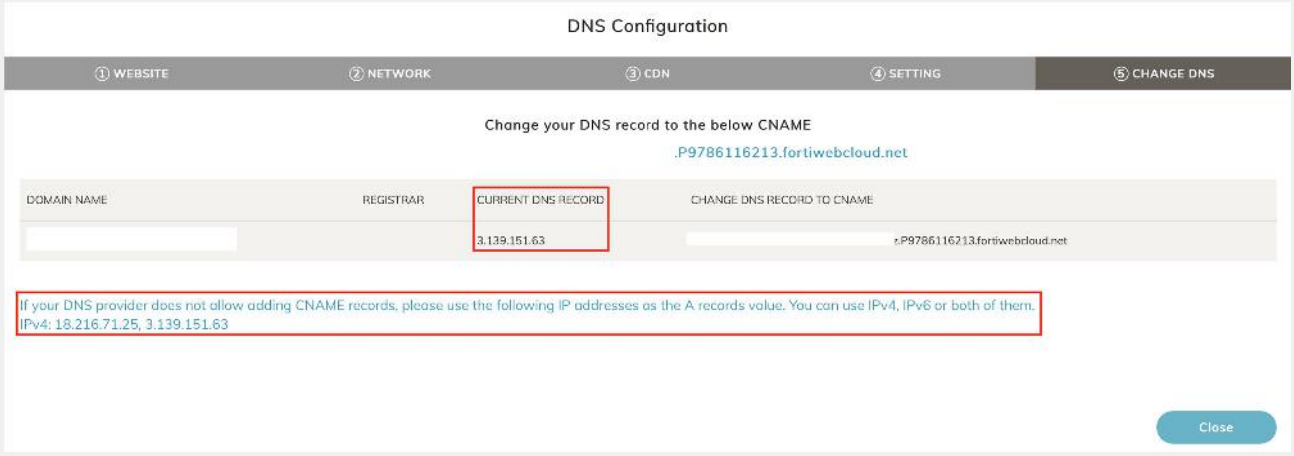

FORTINET Training Institute

 30

To finish the onboarding process, you must create DNS records in the domain server that hosts your application. FortiCloud verifies that those records exist so that the correct redirection takes place.

As shown on this slide, the GUI provides all the details about the records you must create. The GUI also displays and gives you an option for cases in which CNAME records are not allowed.

Benefits of FortiWeb Cloud

- Advanced web and API protection
- Reduced cost and time savings
  - No infrastructure or software to manage
  - Machine learning dramatically reduces manual configuration and management
  - No overprovisioning required
- Provides CDN, bot mitigation, DDoS, and API built-in protection with simple, all-in-one pricing
- Deploy applications with confidence

FORTINET Training Institute

 31

The benefits of FortiWeb Cloud include:

Advanced, enterprise-class web application security that protects against OWASP Top 10 threats, DDoS attacks, malicious bots, and zero-day attacks.
It offers:
- Reduced cost and time savings because there is no infrastructure or software to manage
- Machine learning, which dramatically reduces manual configuration and management
- Cost savings because there is no need to overprovision for idle capacity to handle spikes.

FortiWeb Cloud provides CDN, bot mitigation, DDoS, and API built-in protection with simple, all-in-one pricing. FortiWeb Cloud allows you to deploy your web applications with confidence, removing any security concerns as an impediment to rapid application deployment.

Knowledge Check

1. Which statement is a benefit of Fortinet Managed Rules?

✓A. No user intervention required
B. Natively available in AWS WAF.

1. Which statement is true about FortiWeb Cloud?

A. It needs to be deployed in front of the application server
✓B. It is not run as an EC2 instance.

FORTINET Training Institute

 32

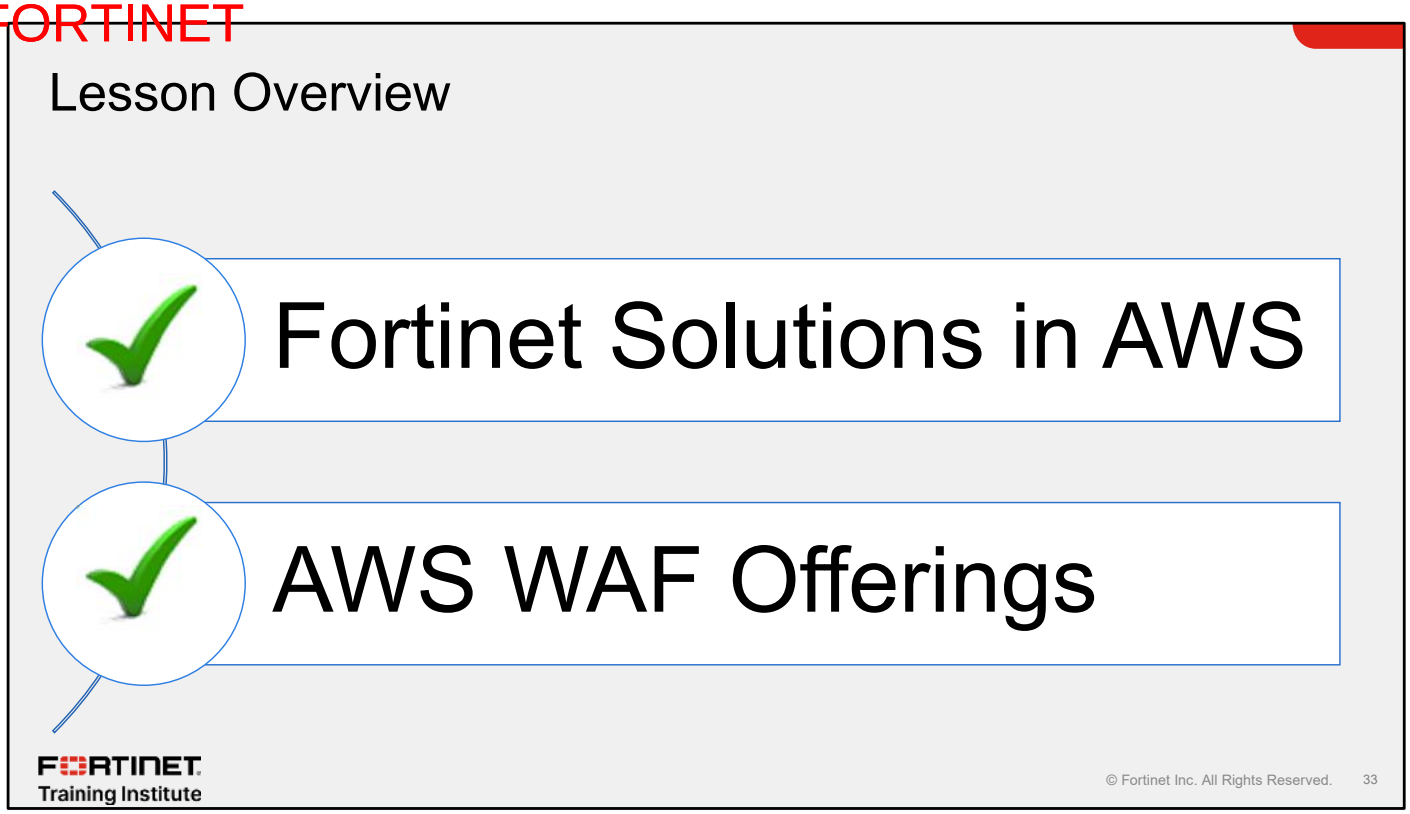

Congratulations! You have completed this lesson.

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

Review

- ✓ Be aware of the Fortinet products on AWS Marketplace
- ✓ Understand FortiGate AWS SDN integration
- ✓ Identify Fortinet WAF solutions for AWS
- ✓ Describe FortiWeb Cloud

FORTINET Training Institute

 34

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

By mastering the objectives covered in this lesson, you learned the different Fortinet products in AWS and how to deploy them.

In this lesson, you will learn about AWS cloud high availability (HA).

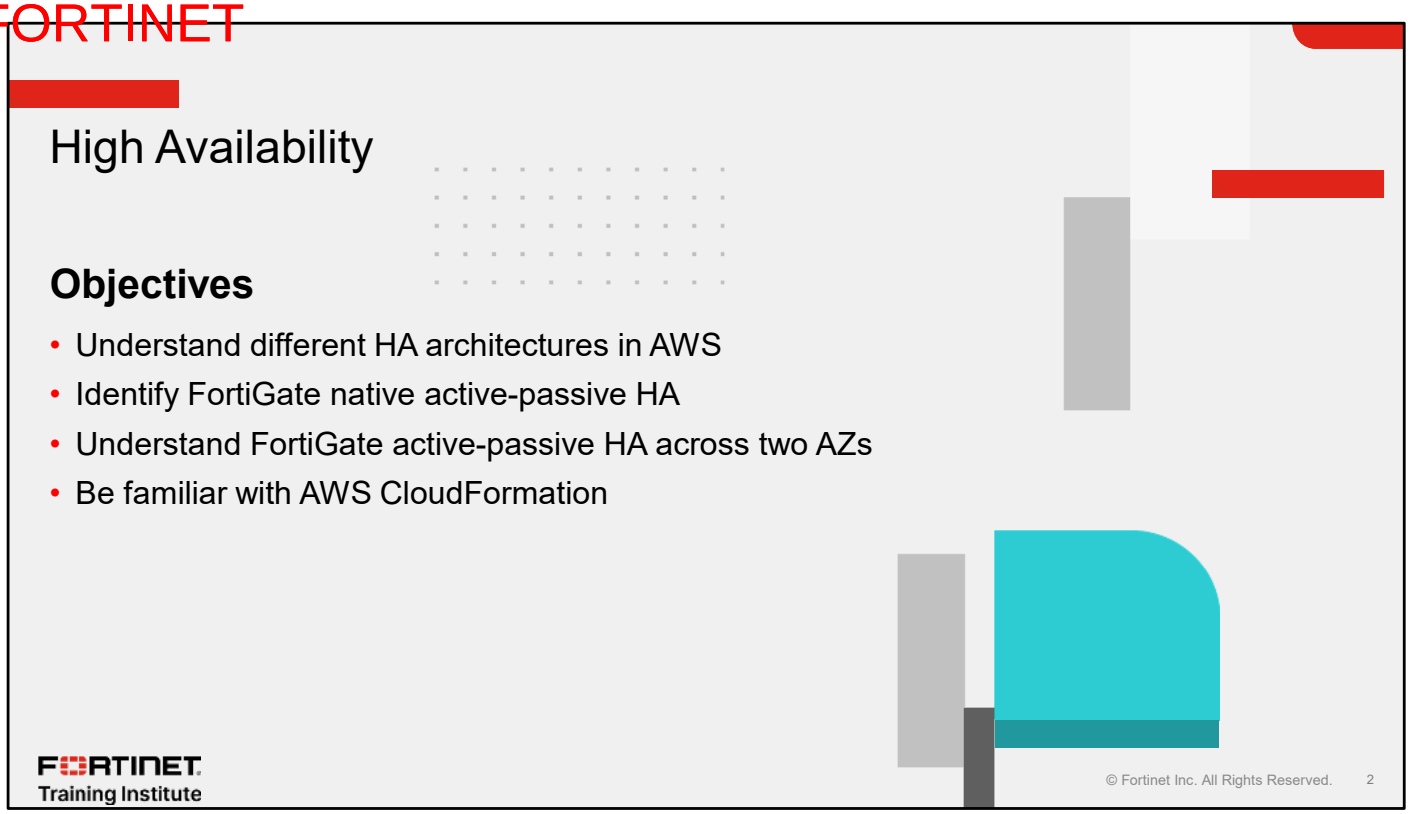

After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding HA, you will be able to implement redundancy in your cloud network using Fortinet solutions.

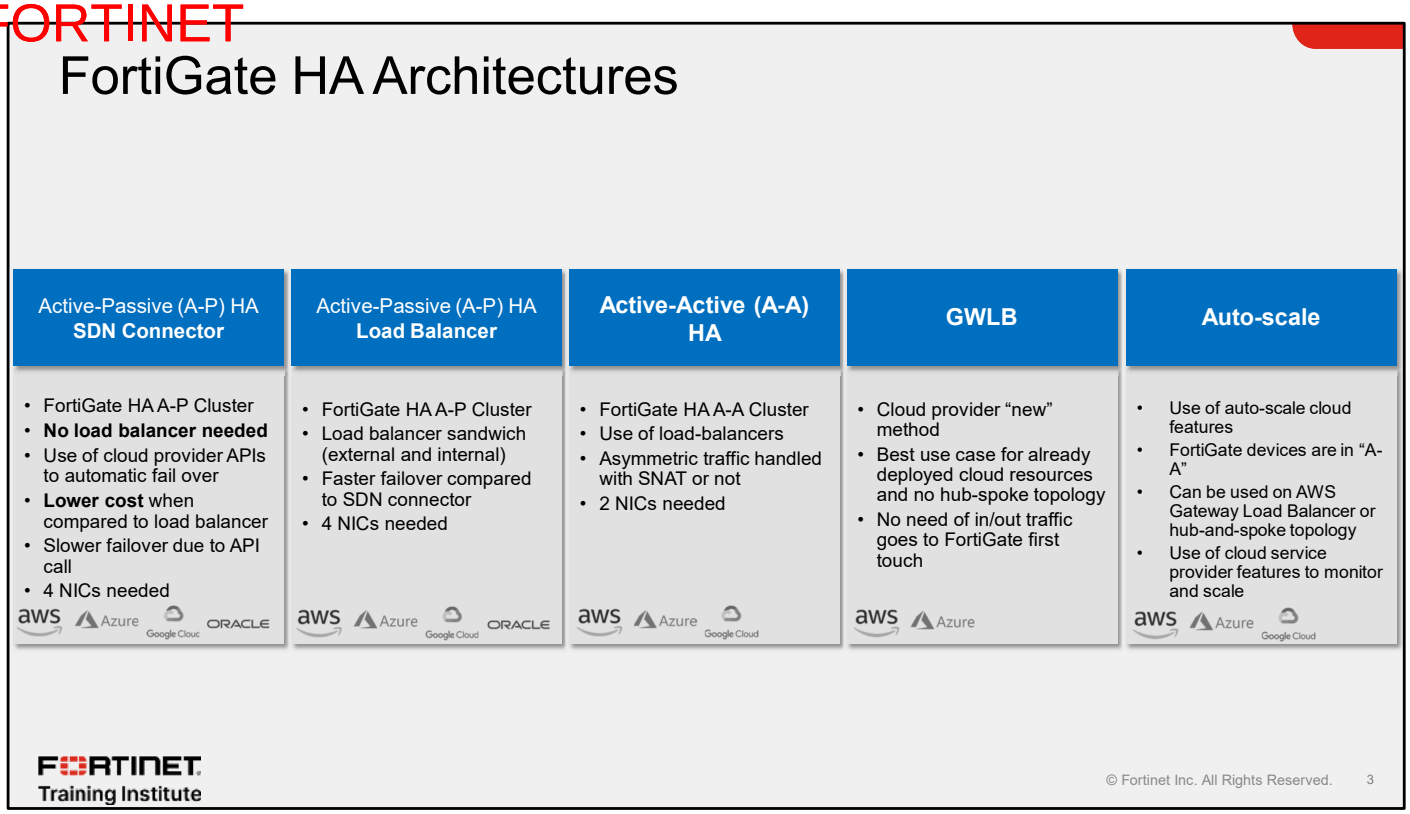

There are several different FortiGate HA architectures that can be deployed in AWS. In this lesson, you will focus on active-passive HA SDN connector. This architecture generally provides the lowest cost to deploy HA because no load balancing is required. The downside of this architecture is that the failover time can be unpredictable because of to the API calls that are involved. You also cannot use this architecture for an active-active cluster.

HA—Unicast Heartbeat (HB) CLI

- HA sync must use unicast IP to sync; cannot use layer 2
- Configuration sync works over unicast IP addresses
- Failover mechanism: commands sent directly to AWS
  - Move public IP addresses
  - Change outbound routing table
- Failover times unpredictable
  - Depends on number of items to rewrite
  - Serial change, not parallel

```
config system ha
    set group-name "CloudHA"
    set mode a-p
    set hbdev "port3" 100
    set session-pickup enable
    set session-pickup-
connectionless enable
    set ha-mgmt-status enable
    config ha-mgmt-interfaces
        edit 1
            set interface "port4"
            set gateway 10.1.3.1
        next
    end
    set override disable
    set priority 255
    set unicast-hb enable
    set unicast-hb-peerip 10.1.2.5
end
```

FORTINET Training Institute

 4

HA must use the unicast IP address to sync between cluster members. You must add the two commands that are highlighted on this slide to the traditional HA cluster configurations. These settings are unique to each cluster member because the peer IP address is the other member of the cluster. When failover happens, FortiGate uses AWS APIs to communicate to the cloud and report the failover. FortiGate sends commands to AWS to change the public IP address and the outbound routing table to the FortiGate IP address and routing table. Also, failover times are unpredictable because of the number of items to rewrite, serial changes, and so on.

HA—Active-Passive Unicast FGCP

- Four network interfaces
- Heartbeat and management interfaces are unusable for production traffic
- Management interface (port4) for administrative traffic
- Unicast heartbeat and config sync interface (port3)
- Use the commands below to stop sync of static routes, interface IP, VIP, and IP pools

```
config system vdom-exception
    edit 1
        set object router.static
    next
    edit 2
        set object firewall.vip
    next
    edit 8
        set object system.interface
    next
```

FORTINET Training Institute

 5

There is no traditional FortiGate Clustering Protocol (FGCP) to use in HA in cloud computing. The solution is to use HA active-passive unicast FGCP, which is a modified version of traditional FGCP. In this scenario, there is no multicast traffic between heartbeat interfaces; instead, there is only unicast traffic. To form HA between two FortiGate devices, you must configure the peer IP address on node. Also, there is a management interface (port4), which is unique to each cluster member and has a subnet with internet access. Each cluster member can be accessed separately through management interfaces. There are two interfaces processing traffic: external and internal. Both heartbeat and management interfaces are system virtual domains (VDOMs) that are hidden and unusable for processing production traffic.

Note that in the HA cluster, all configuration, including the management IP address, synchronizes between HA peers. This is fine if HA is deployed in traditional networking. However, for cloud HA deployments, you must configure a VDOM exception to prevent interface synchronization between the two FortiGate devices. This is because in the AWS cloud you cannot assign duplicate IP addresses. To avoid interface synchronization, you can use the commands shown on this slide.

Active-Passive FortiGate HA for AWS

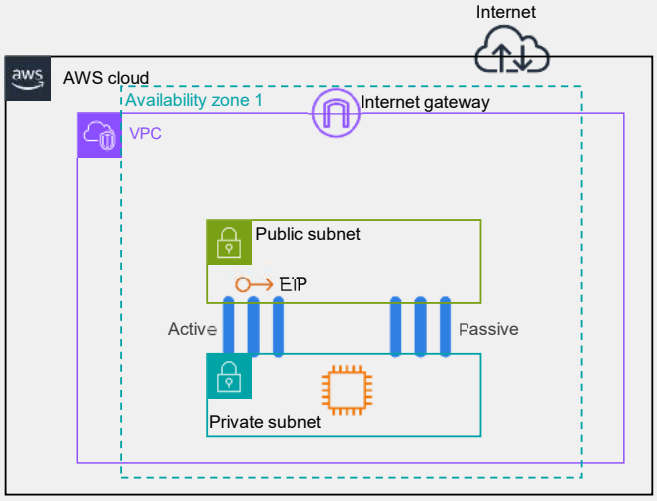

- Single availability zone (AZ)
- Native active-passive FortiGate cluster
- Uses secondary IP addresses in FortiOS
- Use of cloud provider APIs to automatically fail over
- Data plane and management HA
- Configuration synchronization
- Session synchronization
- CloudFormation template available on GitHub

FORTINET Training Institute

 6

This slide shows an example of a FortiGate active-passive HA scenario for AWS. This scenario is based on a single AZ. All HA communication is being handled by unicast traffic. The FortiGate devices act as a single logical instance and share IP addressing and subnets. The benefits of this solution include:

- Fast failover of FortiOS and AWS SDN without external automation/services
- Automatic AWS SDN updates to elastic IP addresses (EIP) and route targets
- Native FortiOS session synchronization of firewall, IPsec/SSL VPN, and VoIP sessions
- Ease of use because the cluster is treated as single, logical FortiGate

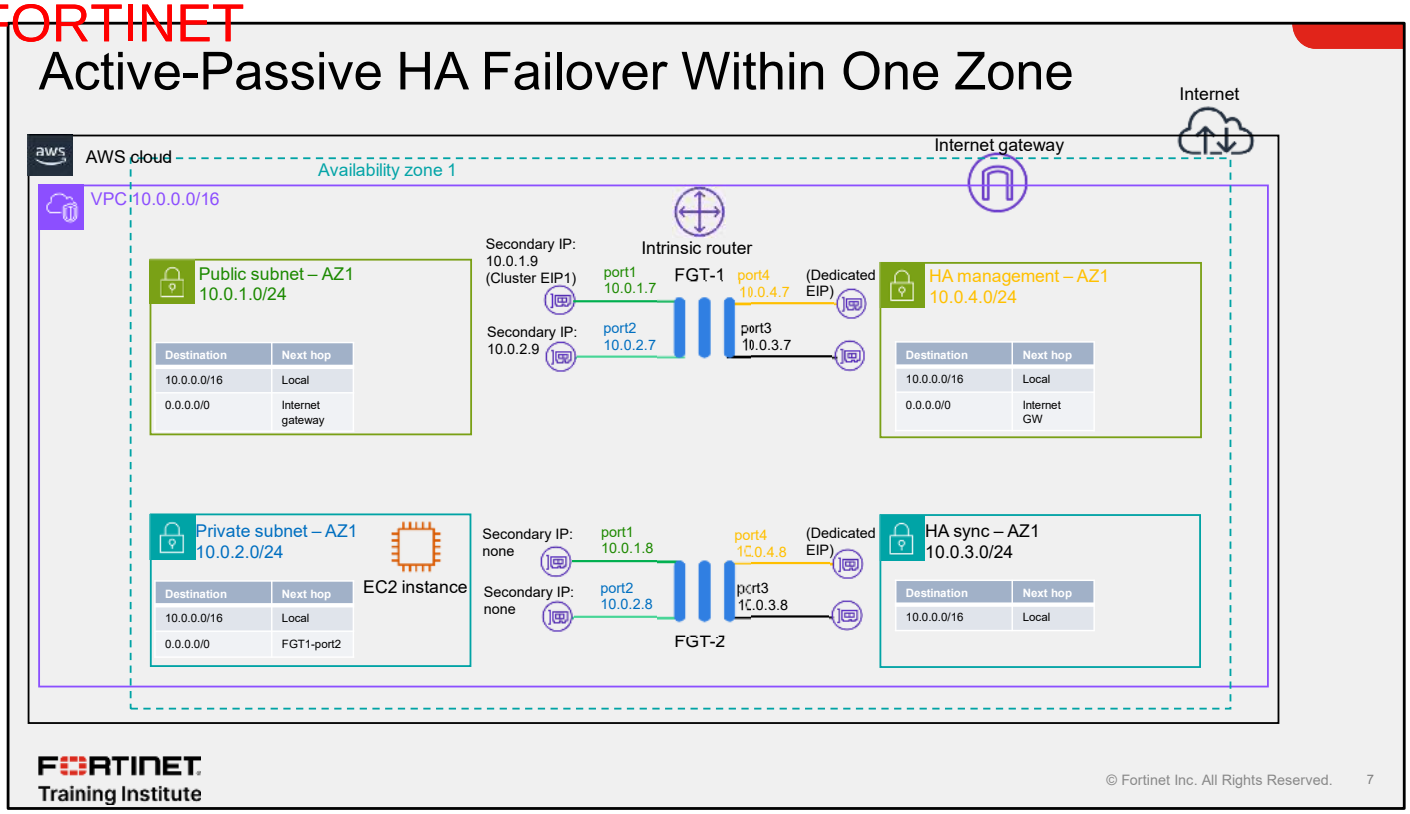

When a failover occurs on a FortiGate HA cluster within the same availability zone, the secondary node takes over. During failover, the secondary IP addresses, which are configured for port1 and port2 on the master unit, move to the slave unit. The elastic IP assigned to the port1 secondary IP address of the master unit also moves to the slave unit. The routing table is updated to forward traffic through the slave unit.

All sessions are synchronized. IPsec phase1 and phase2 are also synchronized and continue to operate during and after the failover.

The AWS SDN updates are performed by the slave unit by initiating API calls from the HA management interface through the AWS internet gateway. The HA management interfaces must be in a public subnet because the AWS EC2 API is only accessible publicly.

Active-Passive HA Failover Within One Zone (Contd)

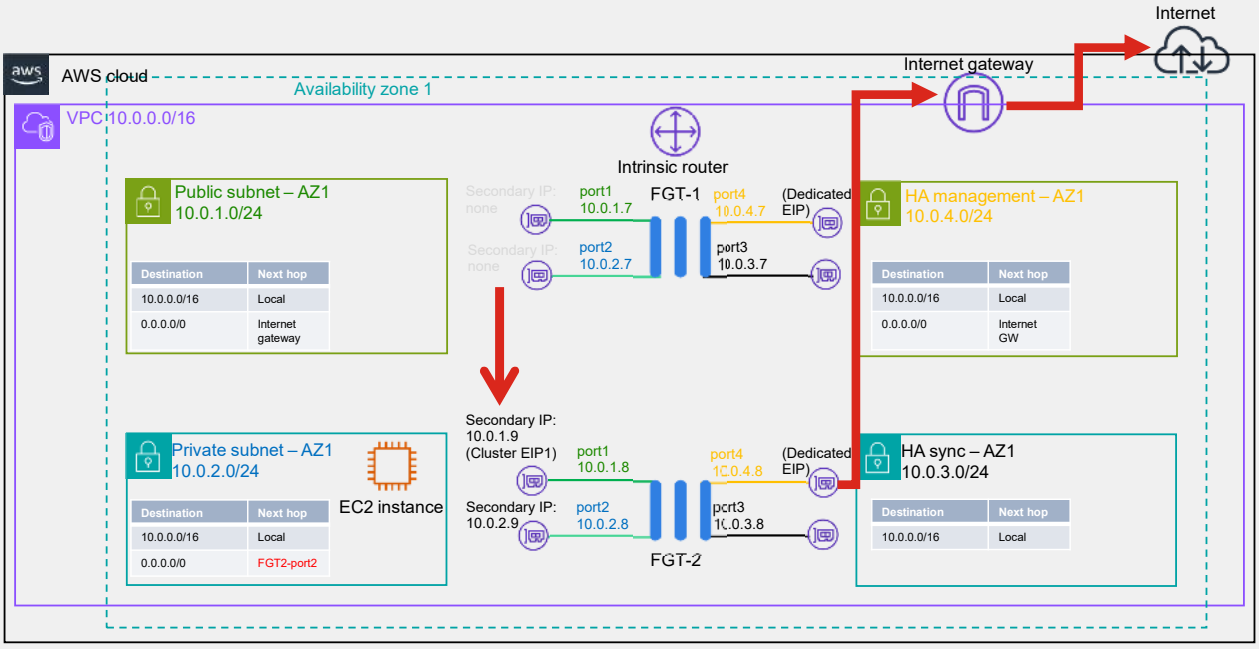

This slide shows the automatic changes a FortiGate HA cluster within the same zone undergoes to stay operational during a failover. The secondary IPs assigned to port1 and port2 are moved from FGT-1 to FGT-2, along with any elastic IPs. Additionally, the routing table in the private subnet is updated to forward all traffic to FGT-2. FGT-2 initiates all SDN updates by performing API calls through its management interface through the AWS internet gateway.

Active-Passive HA Between Two AZs

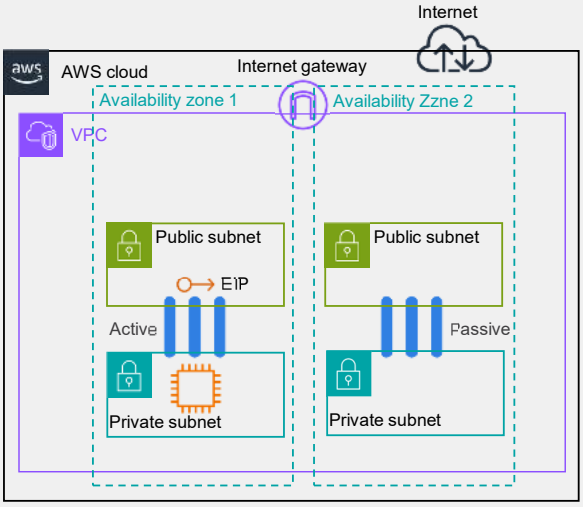

- Two availability zones (AZs)
- Does not utilize secondary IP addresses in FortiOS
- Does not share IP addressing
- Same session synchronization abilities as in single zone deployment, except for IPsec Phase1
- CloudFormation template available on GitHub

FORTINET Training Institute

 9

High availability is important for redundancy. However, what happens if an AWS availability zone goes down and that availability zone happens to be the one your FortiGate HA cluster is deployed in? HA effectively becomes useless. It is a good practice, when feasible, to deploy HA between multiple zones to prevent this issue. The scenario on this slide is based on an HA cluster between two AZs. The FortiGate devices act as a single, logical instance, just like a FortiGate cluster within a single zone, however, IP addressing and subnetting are not shared. This is because it is not possible to span a subnet across multiple availability zones. The other difference is that secondary IP addressing is not used in this scenario.

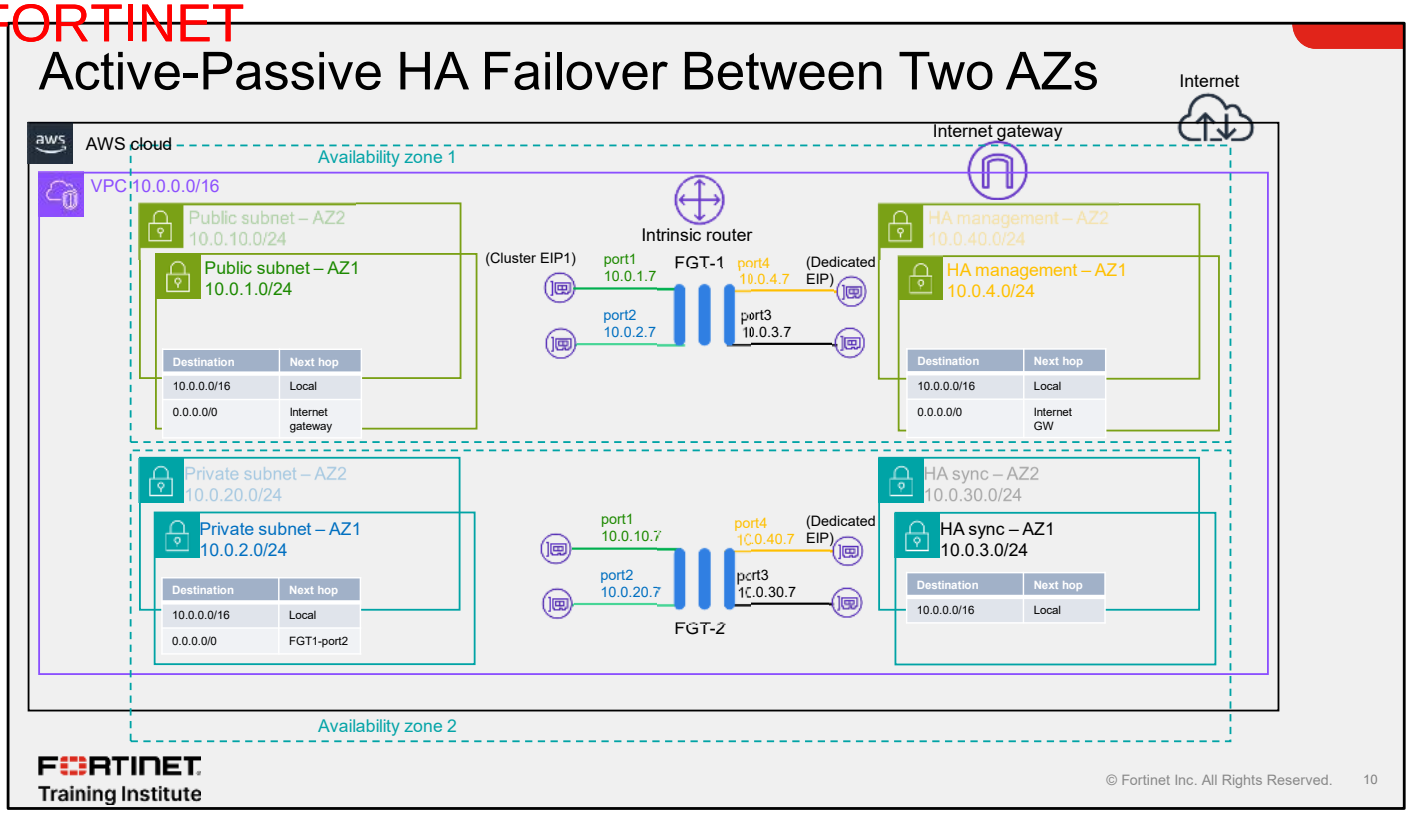

Failover for this scenario is simpler than for a failover for a cluster within the same AZ. This is because this architecture does not use secondary IP addressing. There is only one elastic IP configured for production traffic. During the failover the elastic IP, which is assigned to port1 of FGT-1, moves to FGT-2. Additionally, just like with the single AZ scenario, the private routing table for AZ1 is updated to forward all traffic through port2 of FGT-2.

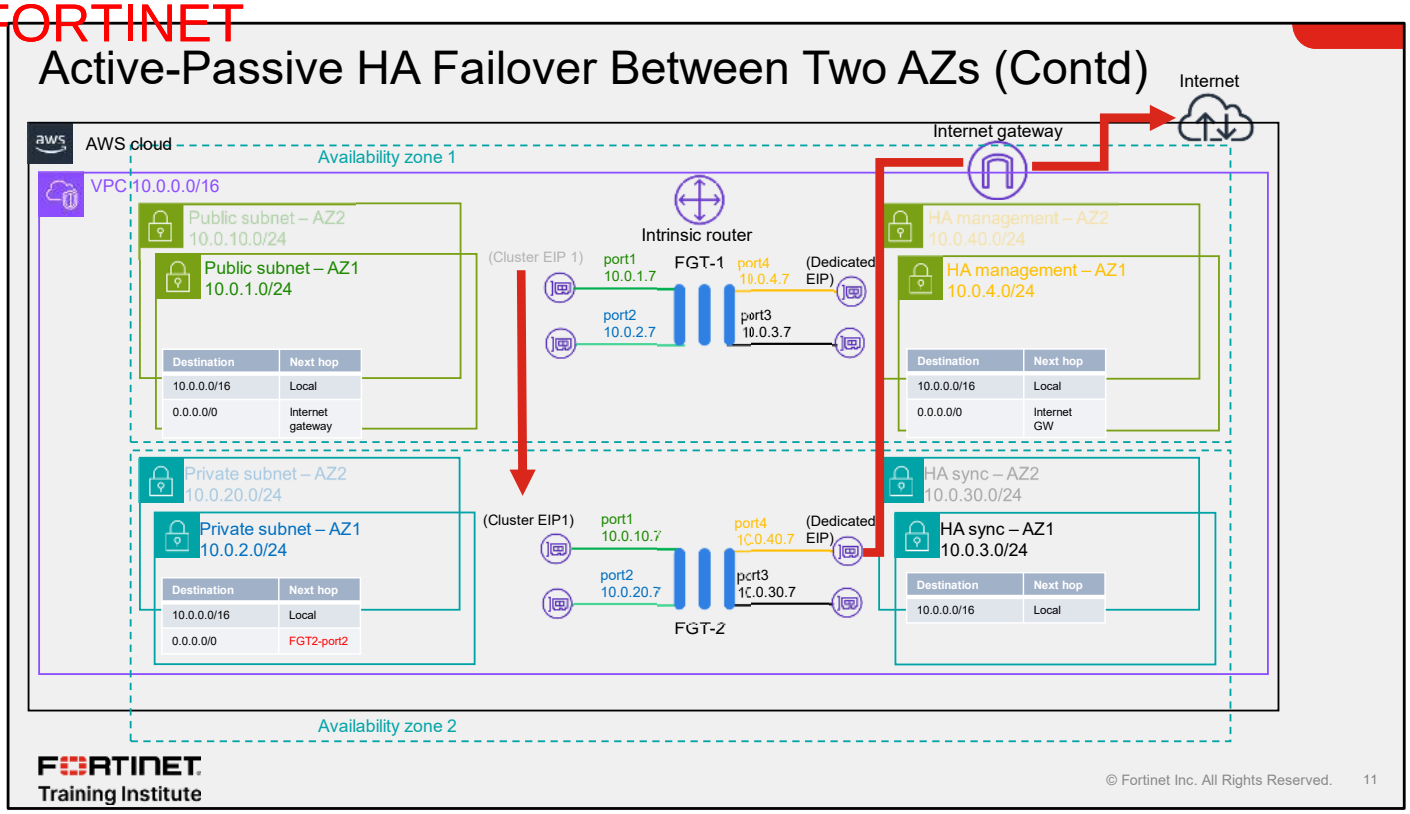

After a failover for an HA cluster between multiple availability zones, the elastic IP moves from port1 of FGT-1 to port1 of FGT-2. The private routing table for AZ1 is updated to forward all traffic through port2 of FGT-2. As with the single AZ scenario, FGT-2 initiates API calls from its dedicated HA management interface though the AWS internet gateway to perform the necessary AWS SDN updates.

You can verify on AWS that the elastic IP addresses of port1 on FGT-1 moved to the new primary FortiGate (FGT-2), and that the routing table changes to point to the internal network ENI of the secondary FortiGate.

Active-Passive HA Failover Between Two AZs (Contd)

```
slave # diagnose debug application awsd -1
slave # diagnose debug enable

...

slave # Become HA master
send_vip_arp: vd root master 1 intf port1 ip 10.0.10.7
send_vip_arp: vd root master 1 intf port2 ip 10.0.20.7
awsd get instance id i-0b29804fd38976af4
awsd get iam role LabDemoRole
awsd get region us-east-2
awsd get vpc id vpc-0ade7ea6e64befbfc
awsd doing ha failover for vdom root
awsd associate elastic ip for port1
awsd associate elastic ip allocation eipalloc-06b849dbb0f76555f to 10.0.10.7 of eni eni-0ab045a4d6dce664a
awsd associate elastic ip successfully
awsd update route table rtb-0a7b4fec57feb1a21, replace route of dst 0.0.0.0/0 to eni-0c4c085477aaff8c5
awsd update route successfully
```

FORTINET Training Institute

 12

During the failover you can observe the slave taking over the master role using the commands shown on this slide.

The output highlighted in green shows that the elastic IP has moved to `10.0.10.7`, which is the IP of port1 of FGT-2. It also shows that the routing table changes to point to the internal network ENI of the secondary FortiGate.

AWS CloudFormation

- Simplifies provisioning and management of AWS resources
- Create templates to provision “stacks” (resources/services/applications)
- Template is stored as a text file
  - JSON
  - YAML
- Fortinet HA templates available
  - Provide value instead of creating the object
  - Review all values on one page
  - Whole stack is created at once
- Templates available on GitHub

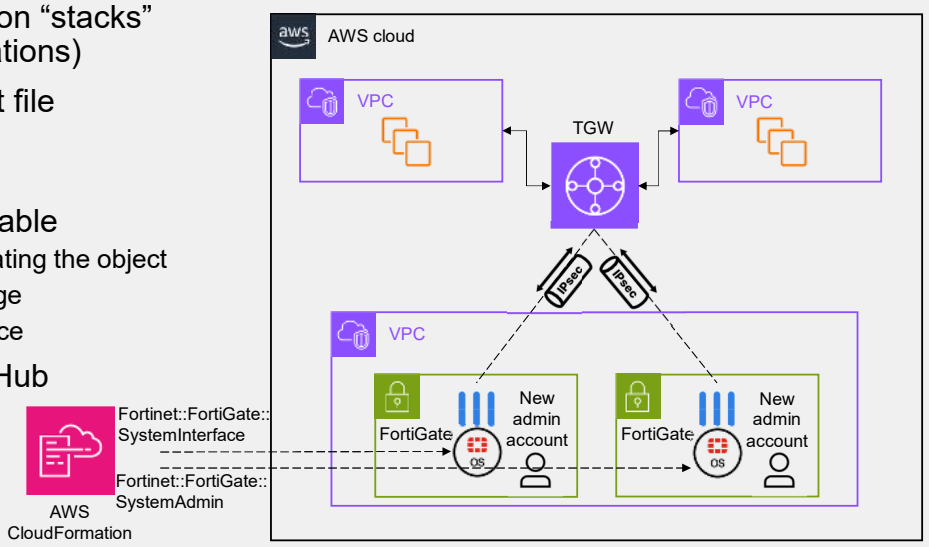

FORTINET Training Institute

 13

Manually deploying any HA setup in AWS can take a significant amount of time. Fortinet provides AWS CloudFormation templates that you can use to speed up that process. AWS CloudFormation is a service that allows you to define and provision AWS infrastructure and resources as code using templates, enabling automated and repeatable cloud environment deployments. Instead of creating every single resource or object manually, potentially introducing manual errors, you can simply run a CloudFormation template and have all the desired stacks created for you by simply providing values such as IP addressing or VPC ID. You will utilize AWS CloudFormation templates in the lab to deploy FortiGate HA.

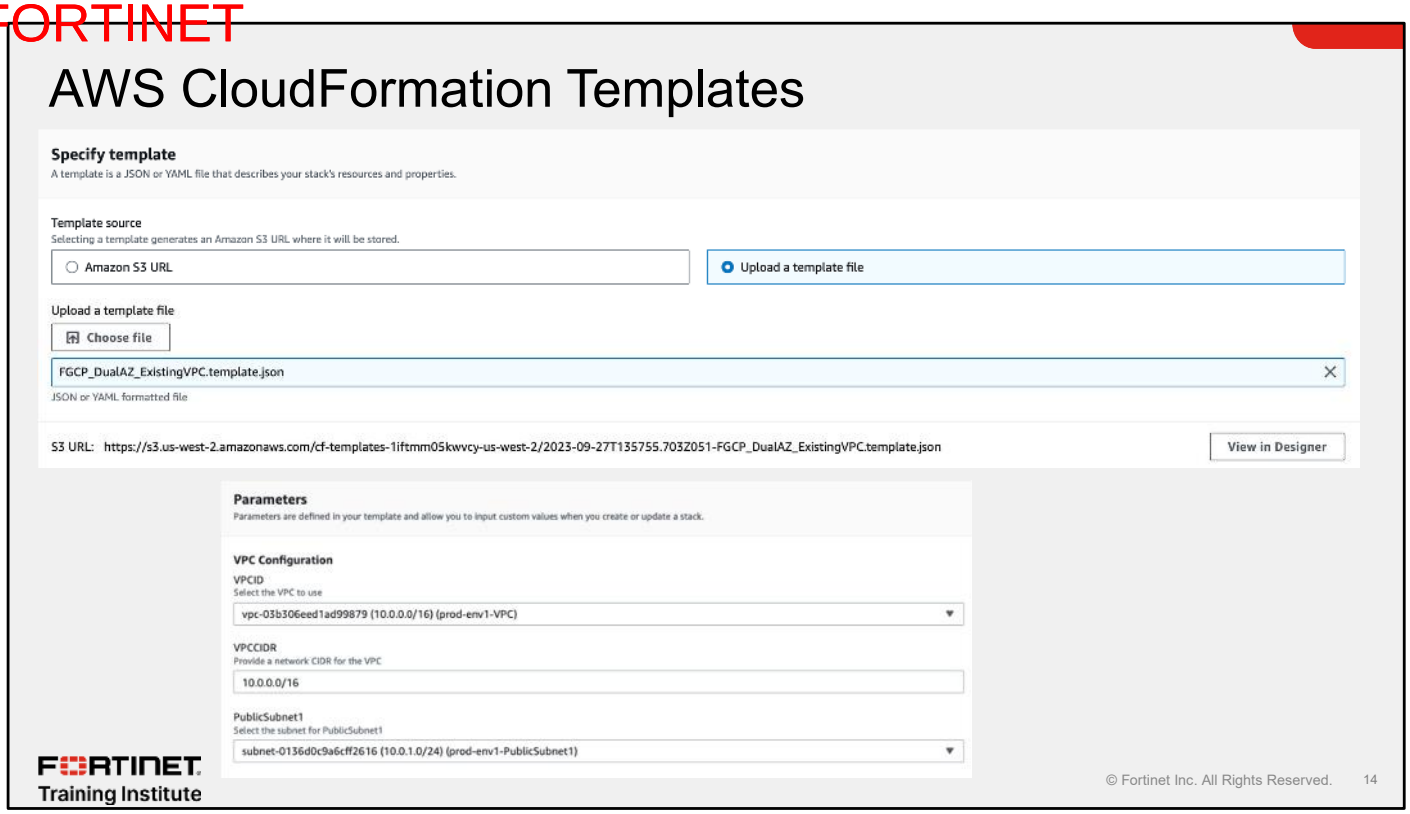

AWS CloudFormation templates provide an easy way to create and manage a collection of related AWS resources, enabling you to provision and update in an orderly and predictable fashion. You can use AWS CloudFormation sample templates or create your own templates to describe the AWS resources. A CloudFormation template is a set of code, based on JSON or YAML, that you can use to specify the kind of VMs, number of subnets, and IP addresses to deploy. After CloudFormation deploys the AWS resources, you can modify and update them in a controlled and structured way. You can apply version control to your AWS infrastructure the same way you do with your software.

Knowledge Check

1. Which high availability architecture can utilize SDN?
   A. Only single AZ HA deployments
   ✓ B. Both single and multiple AZ deployments

2. Which syntax is supported for AWS CloudFormation templates?
   A. Python
   ✓ B. JSON

FORTINET Training Institute

 15

Review

- ✓ Identify different HA architectures in AWS
- ✓ Understand FortiGate native active-passive HA
- ✓ Use Fortinet GitHub to deploy CloudFormation templates

FORTINET Training Institute

 16

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

By mastering the objectives covered in this lesson, you learned methods to deploy HA and load balancer in the AWS cloud.

In this lesson, you will learn about load balancers in AWS.

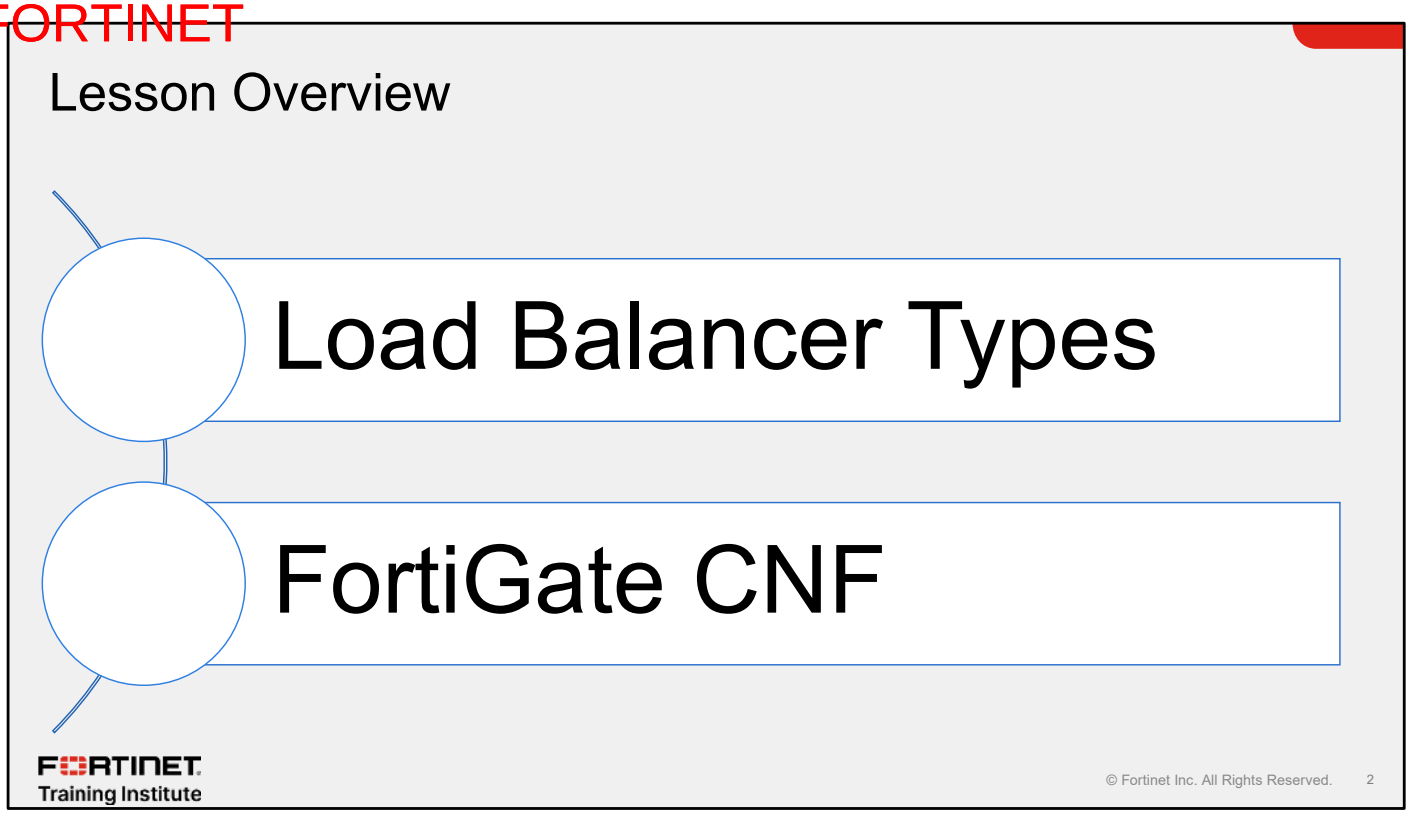

In this lesson, you will learn about the topics shown on this slide.

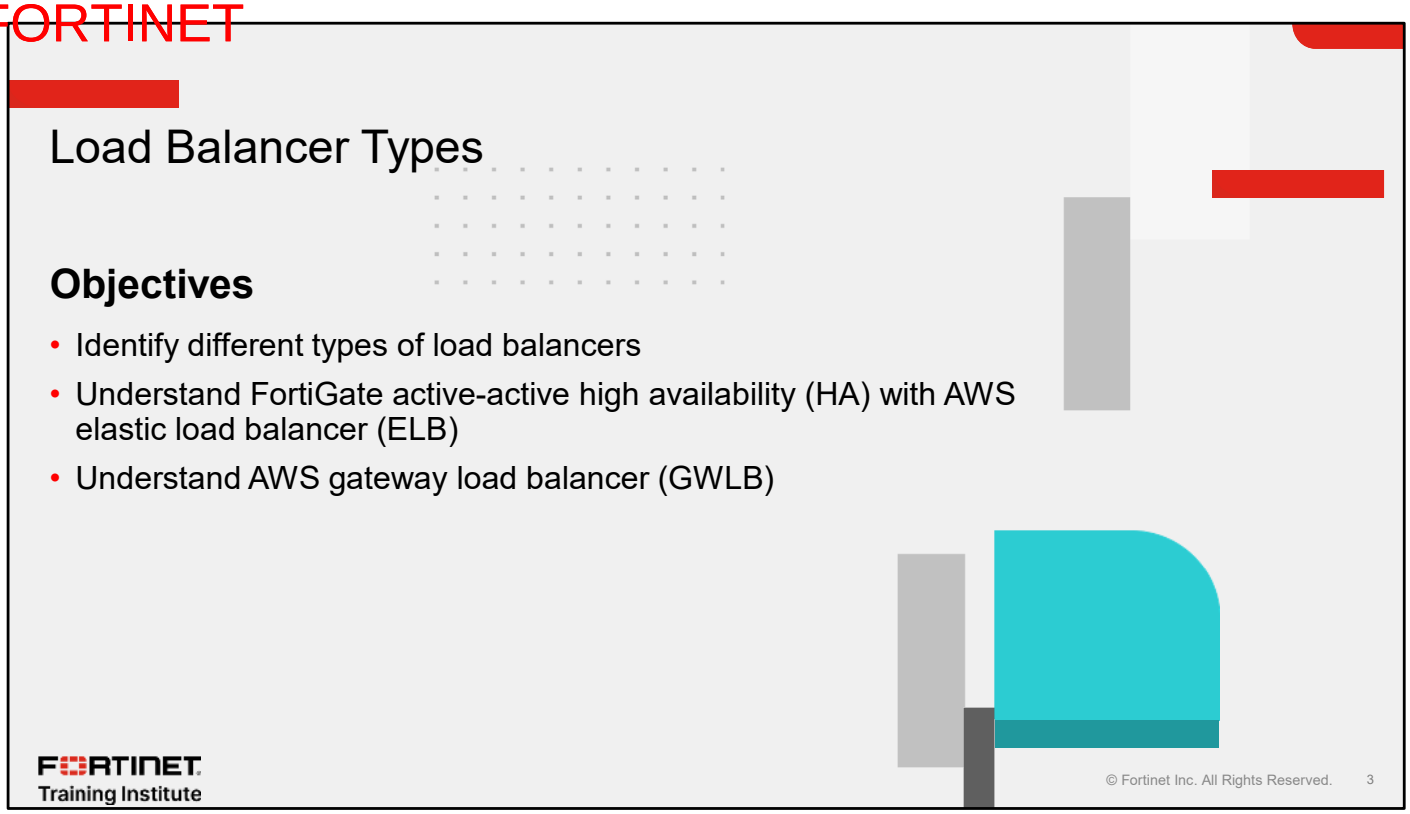

After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding load balancer types, you will be able to implement load balancing architectures in your cloud infrastructure.

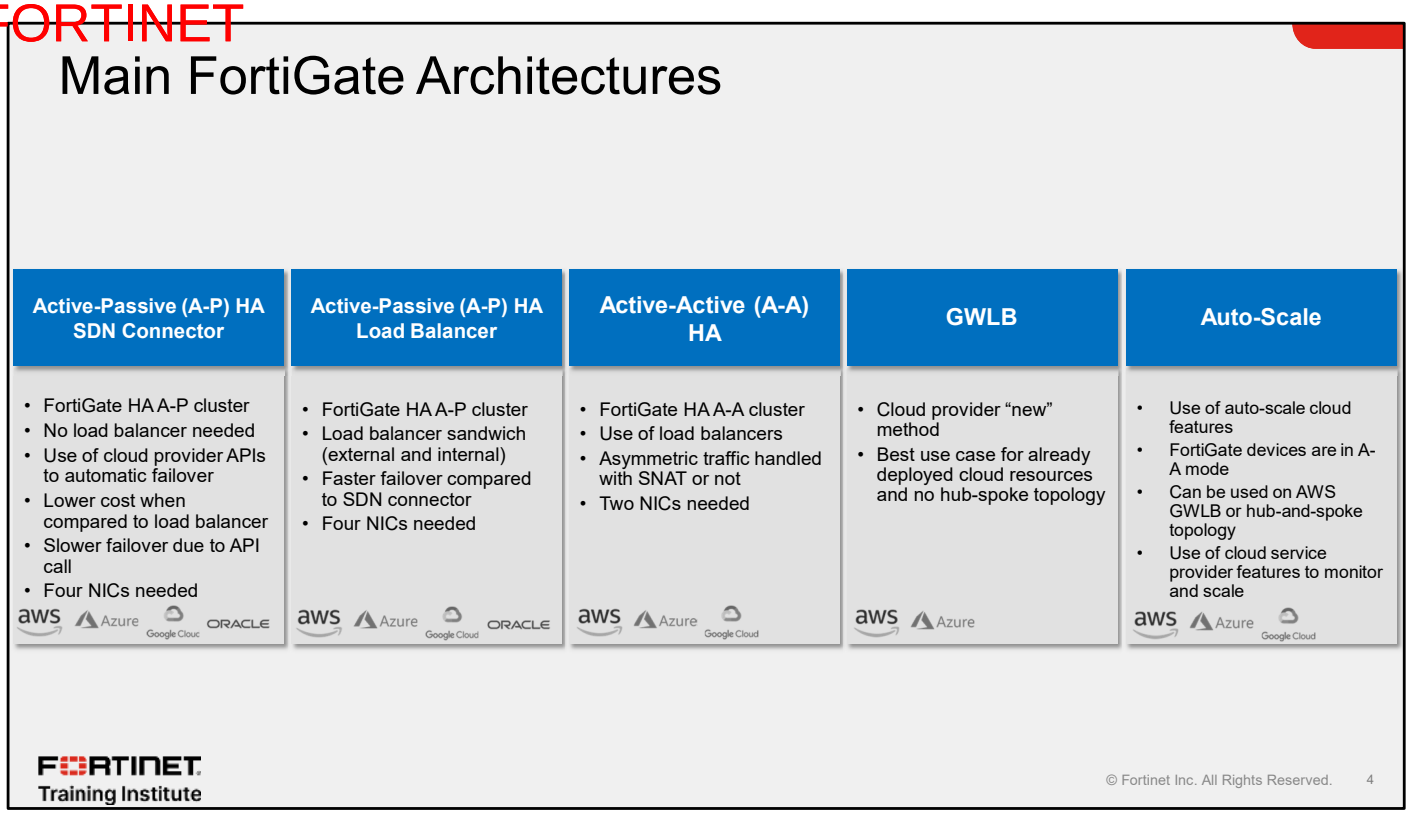

In the *High Availability* lesson, you learned how to deploy an active-passive HA cluster using the SDN connector. In this lesson, you will learn how to use load balancers to set up FortiGate HA clusters. Using AWS load balancers with FortiGate HA clusters allows for faster failovers compared to using SDN connectors.

In the next few slides, you will learn about the different types of load balancers and when to best use them.

Types of AWS Load Balancers

- ELB
  - Network load balancer (NLB)
    - Makes decision at transport layer
  - Application load balancer (ALB)
    - Makes decision at Layer 7
    - Content-based routing
  - Classic load balancer (CLB)
    - Works at Layer 3 and Layer 7
    - Supports classic EC2
    - No longer supported
- GWLB
  - Combines a transparent network gateway and a load balancer
  - Uses a GWLB endpoint, a new type of VPC endpoint
  - Load balances traffic across a fleet of VMs

Feature	NLB	ALB	CLB
Protocols	TCP	HTTP, HTTPS	TCP, SSL, HTTP, HTTPS
Platforms	VPC	VPC	EC2-Classic, VPC

FORTINET Training Institute

 5

AWS offers several types of load balancer types to distribute incoming traffic across multiple instances or resources, enhancing the availability, fault tolerance, and scalability of your applications.

The main types of load balancers covered in this lesson are ELB and the GWLB. The ELB encompasses the NLB, the ALB, and the CLB.

- The NLB makes all its decisions at the transport layer (Layer 4) and is an ideal solution for handling TCP and UDP traffic. NLB is highly scalable and is often used for applications that require high throughput and low-latency, like gaming or VoIP services. The NLB is the best sub-type of load balancer to use for FortiGate HA active-active clusters.
- The ALB operates at the application layer (Layer 7) and is designed for routing HTTP/HTTPS traffic. It provides advanced features, such as content-based routing, path-based routing, and host-based routing.
- The CLB is the oldest type of load balancer in AWS and it is no longer supported. It works at both the application and network layers (Layer 7 and Layer 3, respectively). However, it lacks some of the advanced features of the newer load balancers, which is why it was decommissioned.

The GWLB is used in conjunction with third-party virtual instances like FortiGate. It routes traffic to these instances before forwarding it to target resources or the internet for outgoing traffic.

NLB

- The NLB distributes incoming traffic across multiple targets
- Only forwards traffic to healthy endpoints
- Sits at the VPC level and can handle the varying load of your application traffic in a single availability zone (AZ) or across multiple AZs
- The optimal choice for cross AZ HA clusters
- Always uses source network address translation (SNAT)

FORTINET Training Institute

 6

The AWS network load balancer is designed to efficiently distribute incoming network traffic across multiple Amazon Elastic Compute Cloud (EC2) instances, containers, or IP addresses.

NLB can handle the varying load of your traffic in a single AZ or across multiple AZs.

It sits at the virtual private cloud (VPC) level, so they have access to different subnets in different AZs.

To have traffic and services load balanced between different AZs in a high availability setup, you must use an NLB.

HA—Active-Active With an ELB

- Ingress traffic flows through public ELB
- Public ELB hosts public IP addresses for FortiGate devices
- Load balancer health checks FortiGate instance
- FortiGate VIP translates inbound connections to the protected hosts
- FortiGate performs SNAT inbound to ensure that reply packets arrive at the same firewall
- FortiOS supports configuration sync for firewall policy and objects

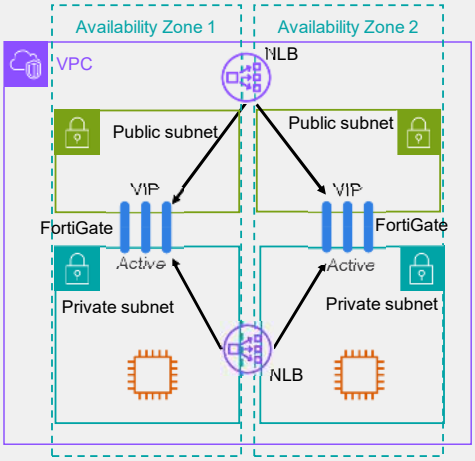

FORTINET Training Institute

 7

This slide shows an example of an active-active load balancing scenario. There are two load balancers: the public load balancer and the internal load balancer. You must pair both FortiGate WAN interfaces with the public load balancer. The internet traffic arrives at the public load balancer first, where it load balances the traffic to the two FortiGate devices. Then it goes to the internal load balancer, and finally, to the servers.

FortiGate VIP translates inbound connections to the protected hosts. FortiGate must perform SNAT on inbound connections in the firewall policy, to ensure that reply packets arrive at the same firewall. FortiOS supports configuration sync for firewall policy and objects.

You will deploy a similar scenario in the lab.

GWLB

- Scalable and highly available entry point for network traffic
- Centralized entry point
- HA
- GWLB uses gateway load balancer endpoint (GWLBe)
- Utilizes GENEVE to preserve traffic content

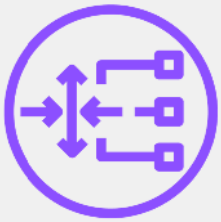

FORTINET Training Institute

 8

The AWS GWLB simplifies deployment and management of virtual instances by providing a scalable and highly available entry point for network traffic. It serves as a centralized entry point for incoming and outgoing network traffic, allowing you to route traffic to and from multiple virtual instances in your network. It is designed for HA, allowing you to distribute traffic across multiple AZs.

The GWLB deploys a gateway load balancer endpoint (GWLBe), which is used in AWS route tables to forward traffic to. Instead of routing traffic directly out to the internet gateway, you could use the GWLBe to forward traffic to a FortiGate first and and then out to the internet.

GWLB utilizes the generic network virtualization encapsulation (GENEVE) to preserve the content of the traffic in a cloud environment. This allows GWLB to create an interoperable overlay network to communicate with other virtual applications in AWS cloud, such as a FortiGate VM.

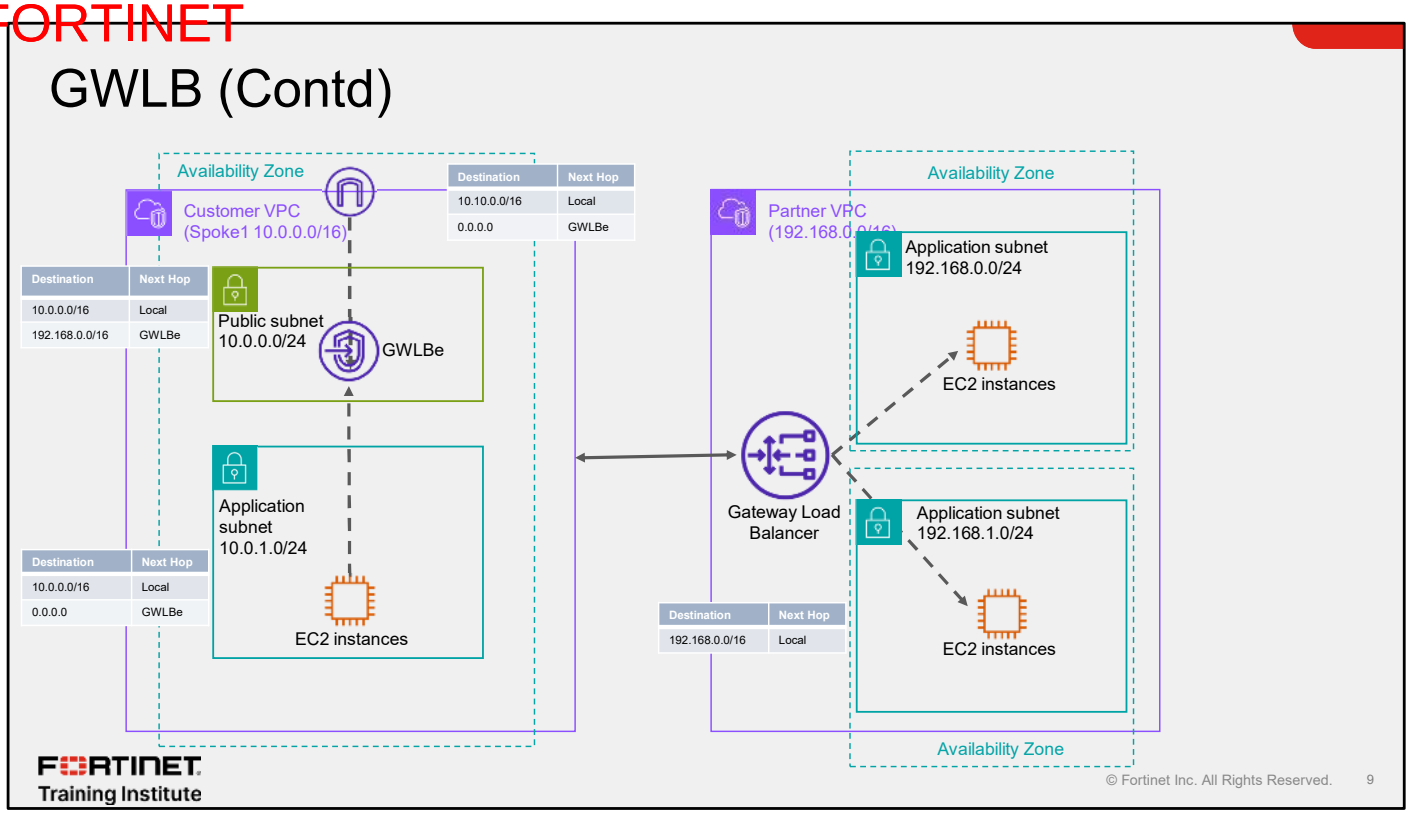

The GWLB deployment example on this slide shows two VPCs:

- Customer
- Partner

A customer VPC contains customer applications or workloads. This is a VPC where a GWLBe is also deployed. The partner VPC will have partner or third-party appliances, such as firewalls or intrusion detection system (IDS) and intrusion prevention system (IPS) devices. In this scenario, GWLBe allows for seamless inspection of traffic without having to change the traffic source or destination or requiring NAT translations.

To ensure high availability, you can use the advanced routing capabilities of GWLB to direct traffic to only healthy appliances and reroute traffic when an appliance becomes unhealthy. GWLB works across multiple VPCs and AWS user accounts, giving you the option to centralize virtual instances. The ability to use GWLB across user accounts enables partners such as Fortinet to offer their virtual instances as an AWS-hosted service that customers access from their VPCs. This reduces complexity and improves security.

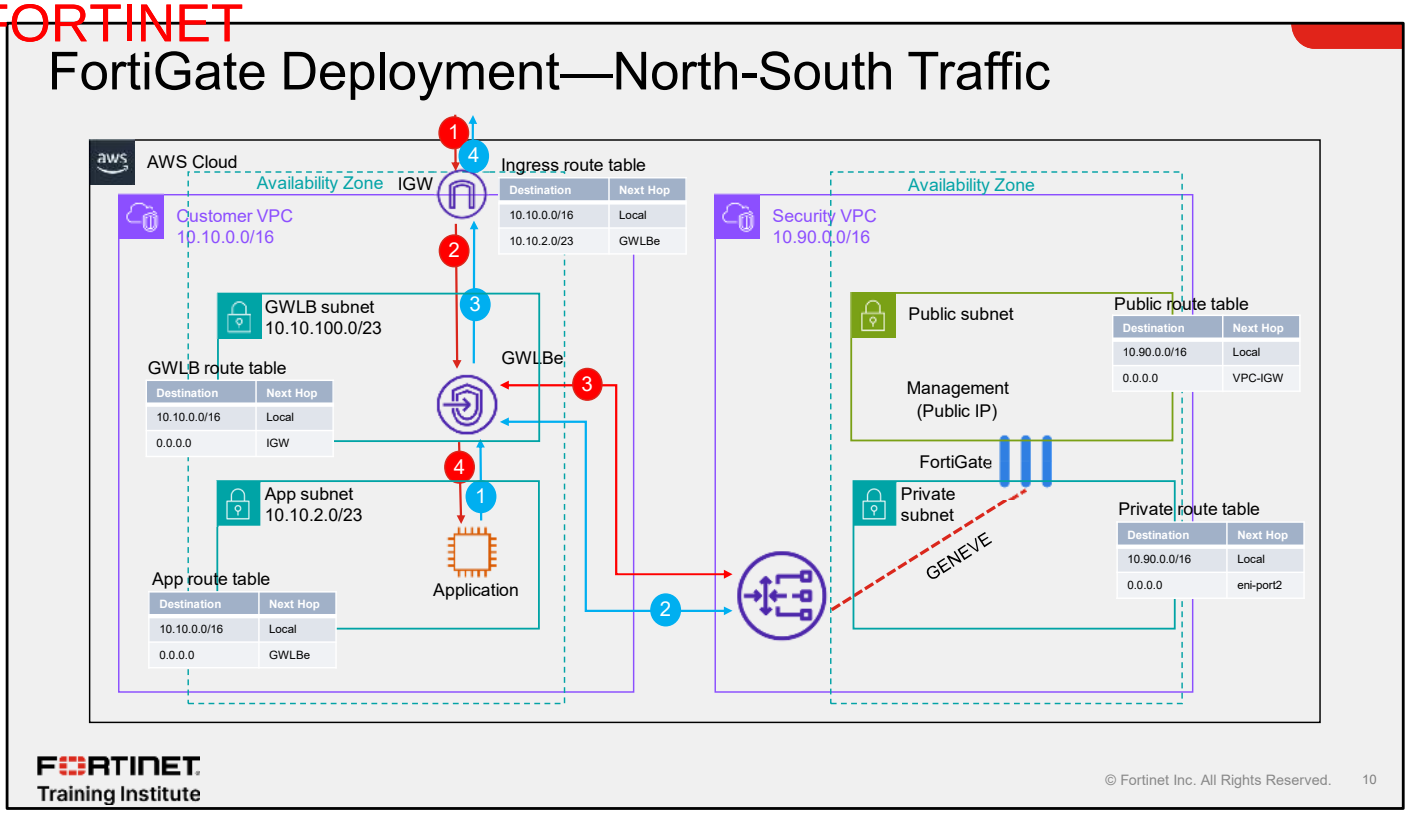

This slide shows the FortiGate deployment to inspect a customer VPC for north-south traffic. This example has two VPCs:

- Customer VPC
- Security VPC

The customer VPC hosts all the workloads and application instances. It also has two subnets. One is an application-purposed subnet, and the other is the GWLB subnet. The application-purposed subnet deploys application workloads, and FortiGate must inspect this traffic. The GWLB subnet deploys the GWLBe so that traffic is redirected to the GWLB, which then forwards the traffic to FortiGate for inspection. The security VPC deploys FortiGate.

So how does FortiGate inspect inbound and outbound traffic?

To inspect inbound traffic, the internet gateway in the customer VPC is associated with an ingress route table. This route table directs traffic for the application subnets through the GWLBe. The traffic then goes through the GWLB in the security VPC, where it is encapsulated with Geneve and sent to FortiGate. FortiGate inspects the traffic and forwards it to the application instances.

For outbound traffic, the route table that the application subnet is associated with has a default route through the GWLBe. The traffic originating from the application instance is forwarded to FortiGate through the GWLB. After inspection, FortiGate sends the traffic to the internet. You must set static routes for all of these traffic redirects after deployment.

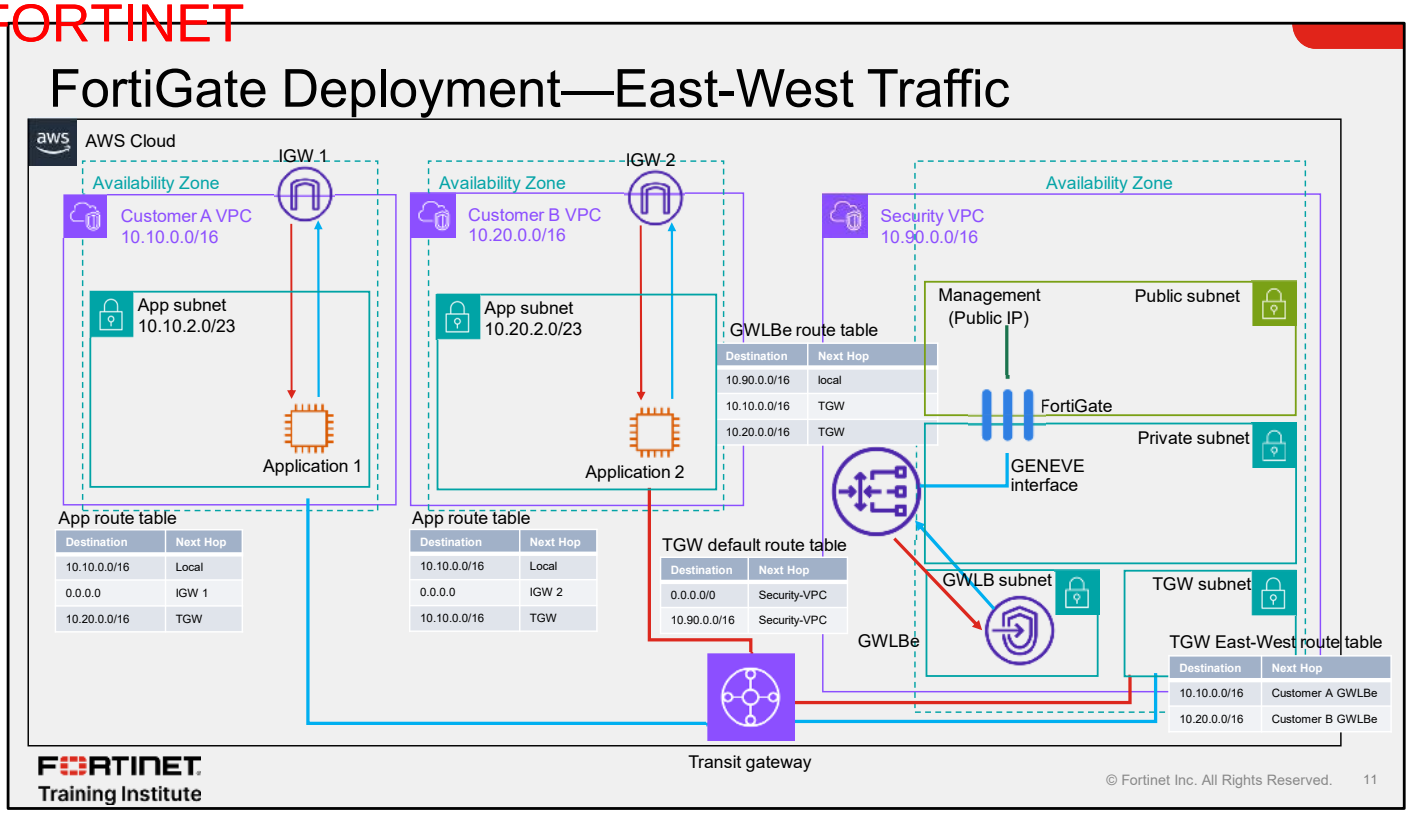

This slide shows an example of a deployment that uses GWLB for east-west security inspection between two customer VPCs. In this deployment, transit gateway is used to connect two customer VPCs to the security VPC through a transit gateway (TGW). This ensures that any access to and from the application VPC is routed through the security VPC, where FortiGate can inspect the traffic.

The security VPC has a FortiGate deployed as well as GWLBe and TGW subnets. The GWLB subnet deploys the endpoint so that traffic is redirected to the GWLB, which then forwards the traffic to FortiGate for inspection. The TGW subnet deploys the TGW and associated resources, which allows the connection of the customer VPCs to the security VPC.

When a VPC needs to communicate with another VPC, client packets are forwarded to the TGW gateway. Based on the TGW default routing table, the packets are forwarded to the security VPC. The packets are now forwarded to the security VPC either through the austomer A GWLBe or the customer B GWLBe attached subnets. TGW is configured in appliance mode (stateful mode) and will always route packets to the same zone for one established session. Both relay subnets are associated with a local routing table, forwarding all packets to the GWLBe interface located in the same local zone. The packets entering the endpoint are automatically forwarded to the local GWLB component responsible for establishing a tunnel to the local FortiGate device. Now, the local FortiGate device is receiving the packets on its unique GENEVE tunnel interface and processes them using all of its security profiles, such as antivirus, intrusion prevention system (IPS), antispam, data loss prevention (DLP), webfiltering, and so on. After inspection, FortiGate uses its local routing table, pointing to all VPC classless inter-domain routing (CIDR) to forward the packets back through the GENEVE tunnel. After the packets have left the GWLB interface through the GENEVE tunnel, they are routed back to the originating endpoint subnet. The endpoint subnet is associated with a routing table, the purpose of which is to route all traffic to TGW through the VPC attachment. After packets arrive at the TGW, it routes the packets to its destination: customer A VPC or customer B VPC.

Knowledge Check

1. Which type of AWS load balancer is best suited for FortiGate HA clusters?
   A. ALB
   ✓ B. ELB

2. Which encapsulation method is used by AWS GWLB to preserve the original packet content, when forwarding traffic?
   ✓ A. GENEVE
   B. GRE

FORTINET Training Institute

 12

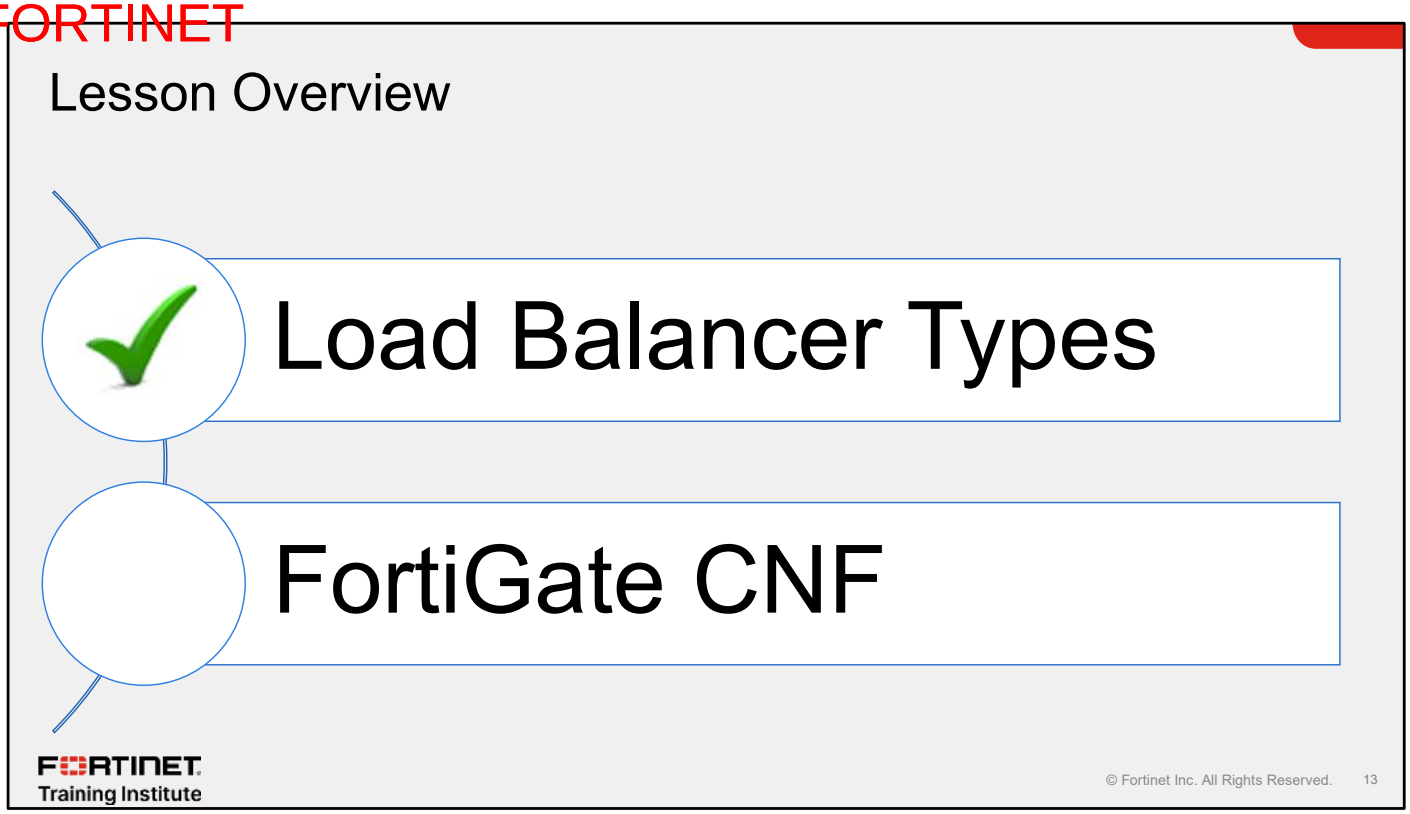

Good job. You now know about load balancer types in AWS and will be able to implement load balancing architectures in your cloud infrastructure.

Now, you will learn about FortiGate CNF.

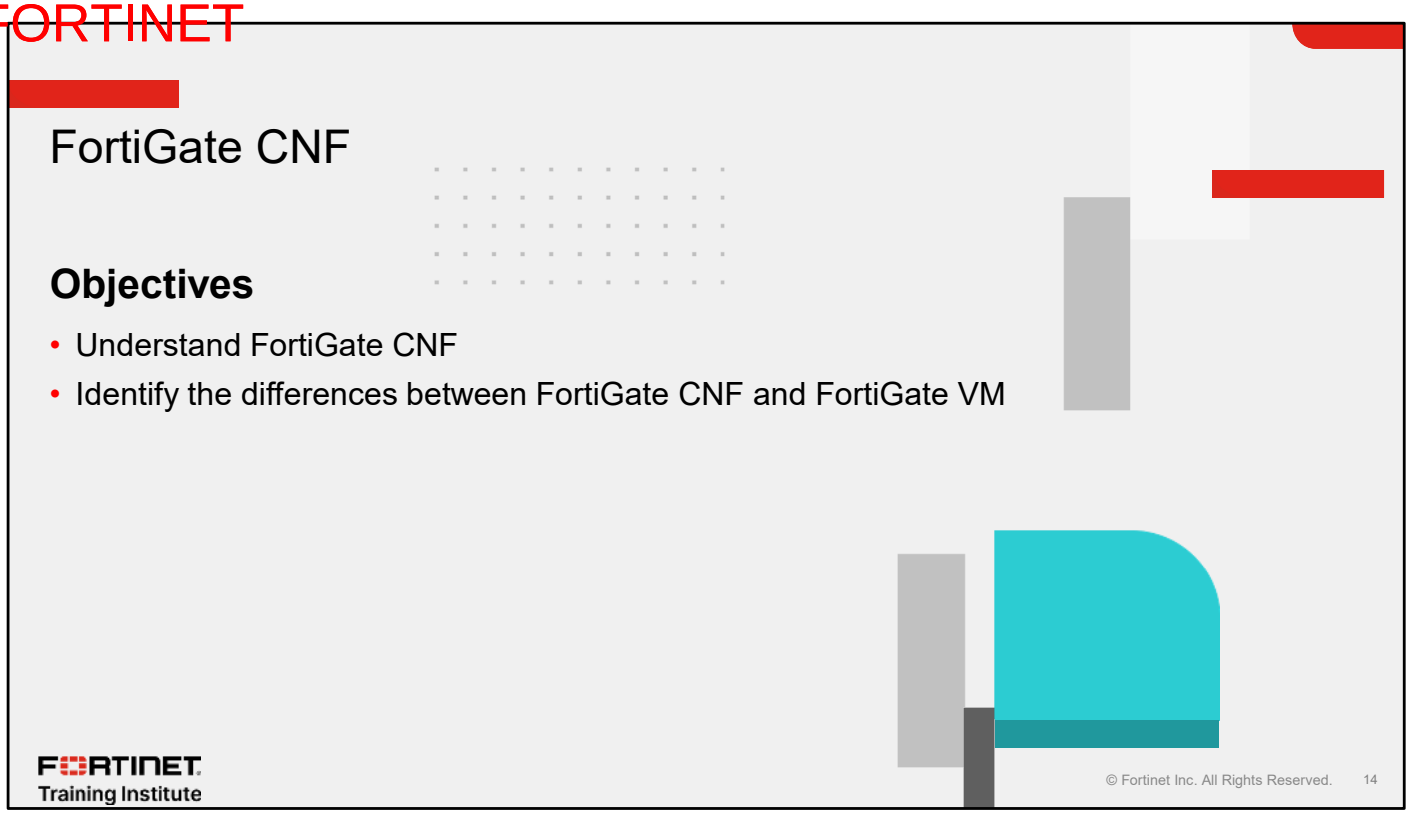

After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in understanding FortiGate CNF, you will be able to deploy a Software-as-a-Service (SaaS) version of FortiGate.

FortiGate CNF

- High performing, auto-scaling next-generation firewall (NGFW) solution to control and inspect north-south and east-west network traffic
- Firewall-as-a-Service (FWaaS)

FORTINET Training Institute

 15

Manually deploying GWLBs AND managing FortiGate devices can be a daunting task. This is why Fortinet offers FortiGate CNF. It is an FWaaS that reduces network security operations workloads. Enterprises don't have to configure, provision, or maintain any firewall software infrastructure.

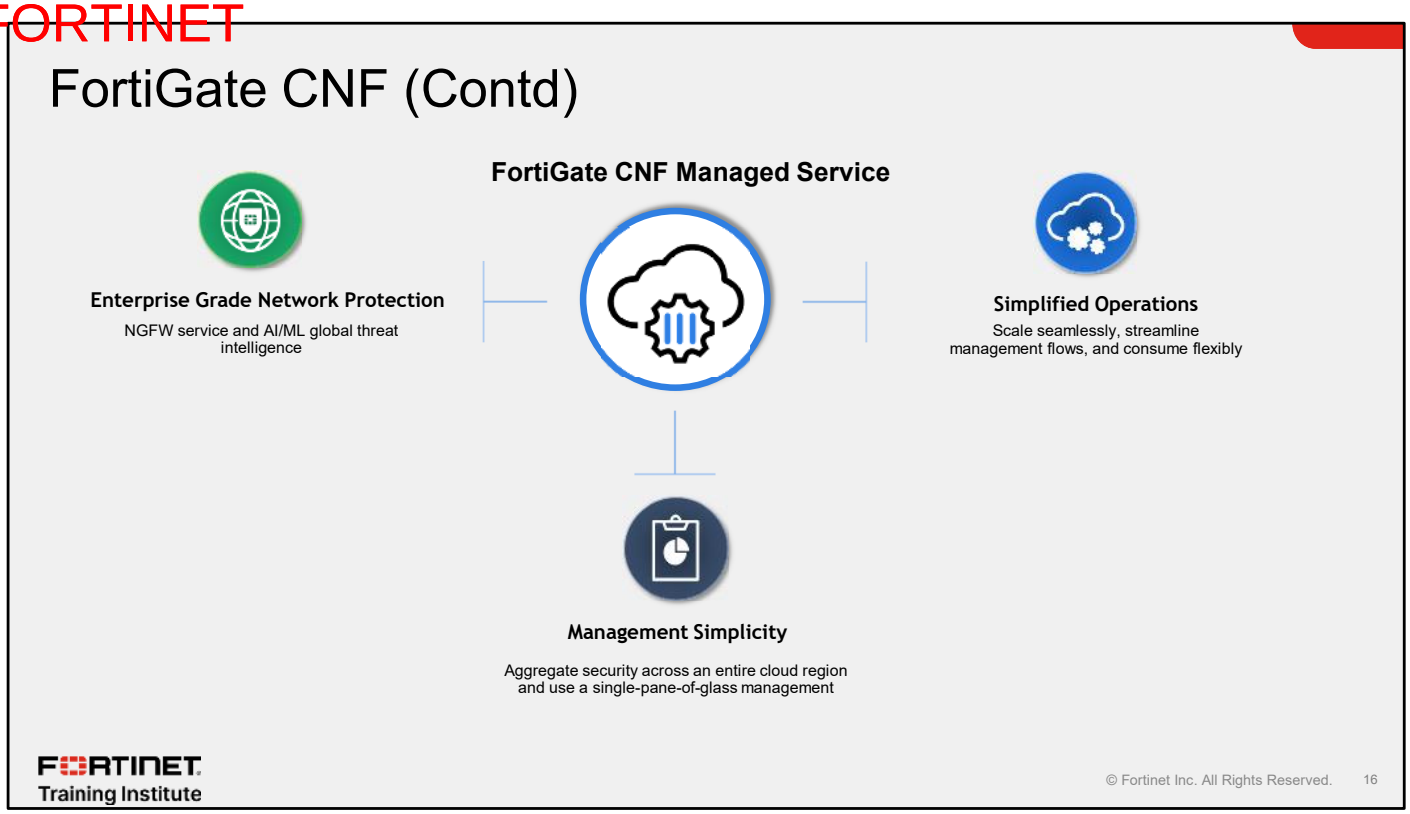

In addition, enterprises enjoy the following benefits:

- Enterprise-grade protection: FortiGate CNF supports the security inspection capabilities of an NGFW, providing deep visibility into the application layer along with advanced detection and comprehensive protection powered by artificial intelligence (AI). It includes Geo-IP blocking, advanced filtering, and threat protection. With this level of traffic inspection, customers can reduce the risks of unauthorized events on AWS workloads caused by web-based threats, vulnerability exploits, and other external and internal threat vectors.
- Zero operations overhead: FortiGate CNF simplifies security delivery by using just one FortiGate CNF instance to secure an entire AWS region. It can protect multiple accounts, subnets, VPCs, and AZs, consolidating security in a region. Cloud-native integration with AWS GWLB helps network security teams move at the speed and scale of applications teams. It eliminates do-it-yourself automation and helps easily secure Amazon VPC environments while improving HA and scaling.
- Simplified management: Cloud-native organizations can use the lightweight user interface and intuitive wizards on the FortiGate CNF console to easily create, deploy, and manage security policies for their AWS environment. For hybrid cloud deployments, you can use a centralized management tool like FortiManager to define, deploy, and manage advanced security policies, backed by the FortiGuard Global Threat Intelligence service, which operates consistently across hybrid environments. You can use integration with AWS Firewall Manager to streamline security workflows and automate security rollout, saving time and increasing efficiency.
- Lower costs: Because there is no security software infrastructure to build, deploy and operate, costs are reduced. Organizations also can save on the training and resourcing costs that would be necessary to deliver do-it-yourself security on AWS. Aggregating security across a region into a single CNF instance avoids the extra costs accrued by solutions that charge by cloud network or AZ. In addition, the FortiGate CNF service utilizes AWS Graviton instances to deliver better price performance.

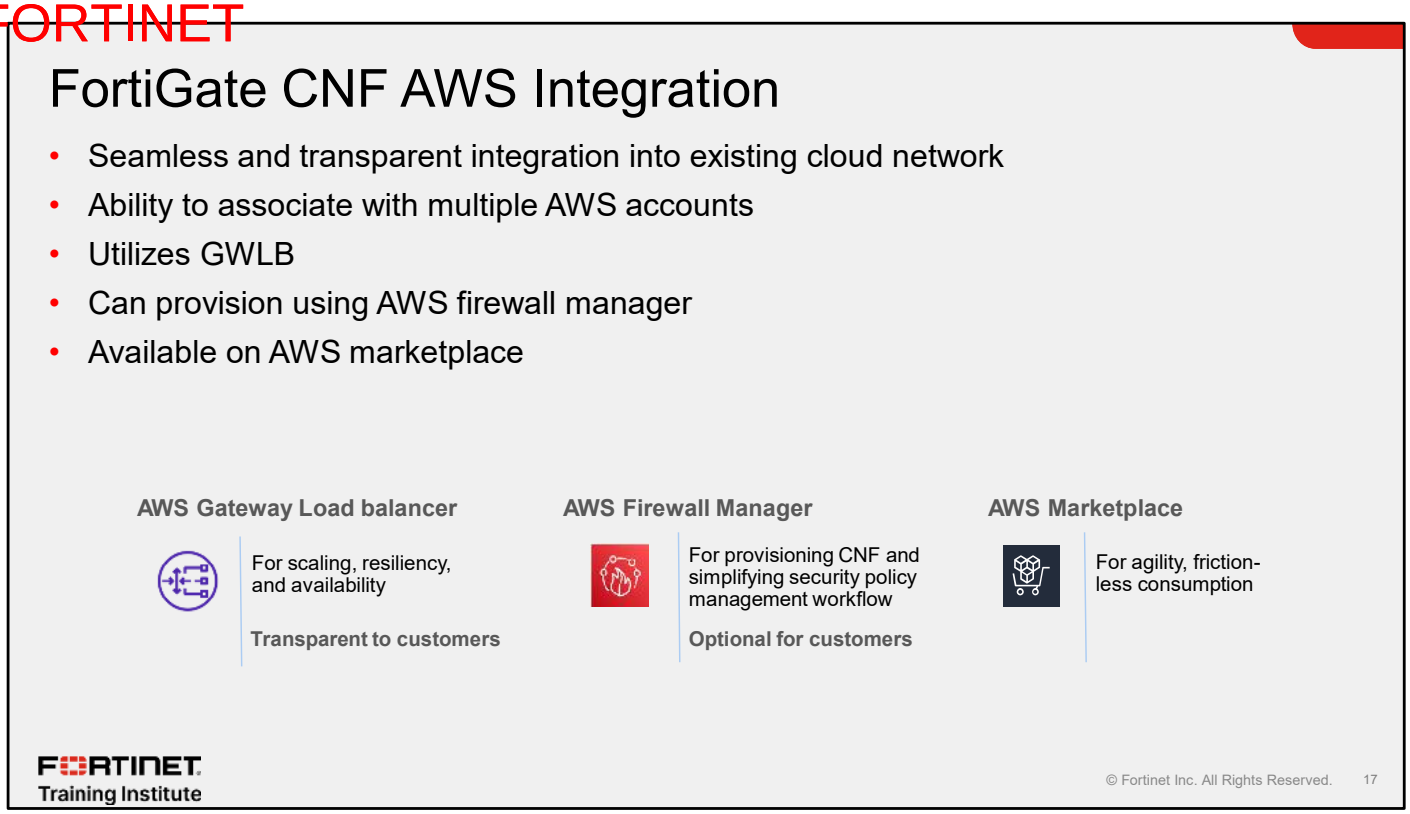

FortiGate CNF integrates seamlessly with your existing AWS cloud infrastructure. The service sits in a separate subnet from your existing cloud network. When the instance is deployed to this subnet, a GWLBe is deployed with it, which is then used to forward all traffic requiring inspection to be forwarded to FortiGate CNF.

A FortiGate CNF instance is dedicated to one customer; however, it can connect to multiple VPCs, accounts, and AZs.

Aside from using GWLB, FortiGate CNF also integrates with AWS Firewall Manager to provision instances and to simplify security policy management workflow. This is similar to how you would integrate FortiGate CNF with FortiManager.

FortiGate CNF is available on AWS Marketplace.

Finally, you can configure FortiGate CNF to log to an S3 bucket, AWS Security Lake, Syslog, and FortiAnalyzer.

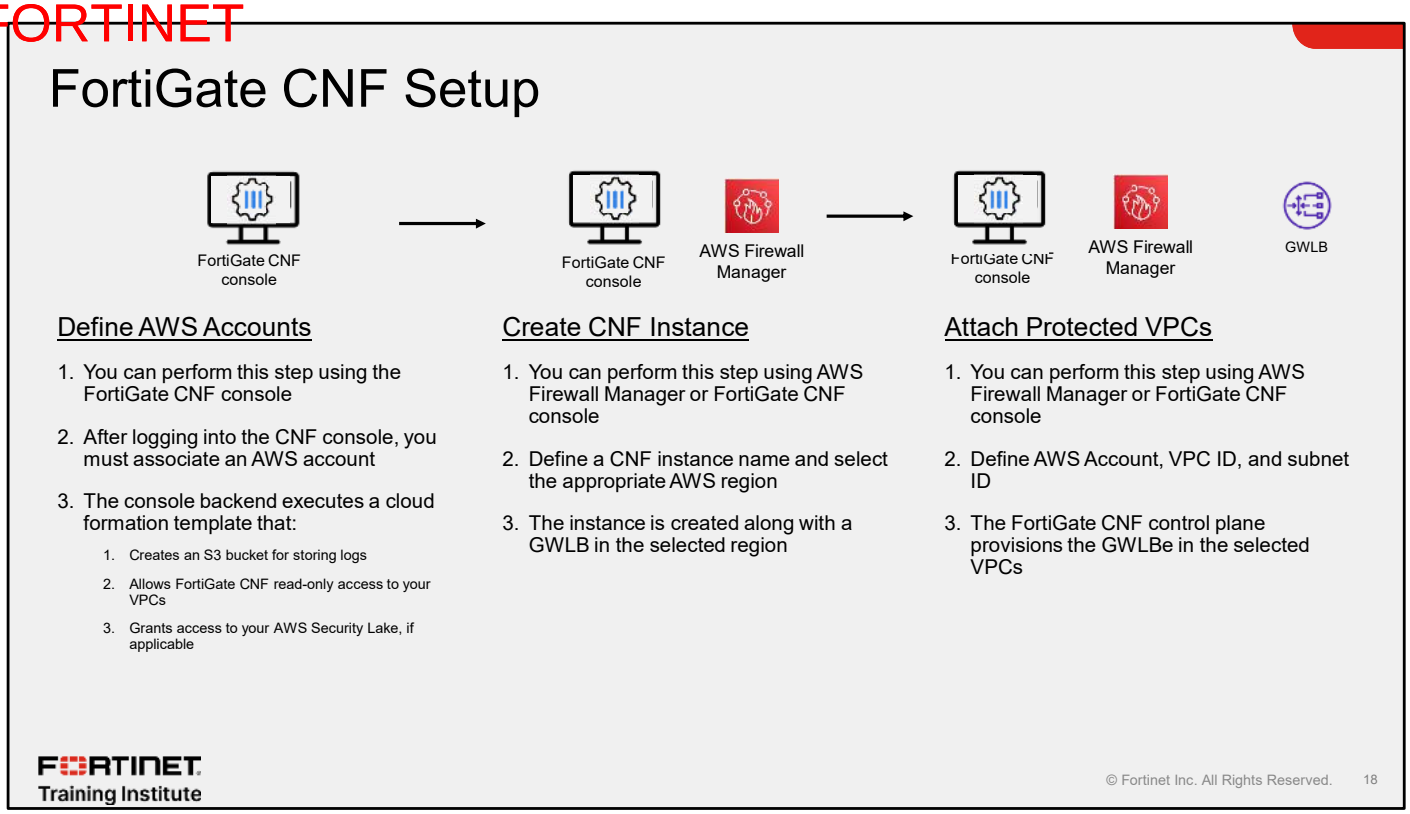

You must complete several steps to set up FortiGate CNF. You can perform all of these steps using the FortiGate CNF console, which is accessible through your FortiCloud account once you have subscribed to the service using AWS marketplace.

First, you must associate your AWS account with FortiGate CNF. Then, you will be prompted to launch a cloud formation template in the US Oregon region. This cloud formation template creates an Amazon S3 bucket for writing and storing logs, allows FortiGate CNF read-only access to your VPCs, and grants access to your AWS Security Lake, if applicable.

Second, you must create a CNF instance. You can perform this step using either the FortiGate CNF console or the AWS Firewall Manager. To create the instance, you must define an instance name and select the AWS region in which to deploy the instance. This process can take up to ten minutes. Once completed, the instance will have been created along with a GWLB in the selected AWS region.

Finally, you must attach the VPCs that you want to protect. You can perform this step using either the CNF console or AWS Firewall Manager. Here, you will define which VPCs and subnets are to be provisioned with a GWLBe.

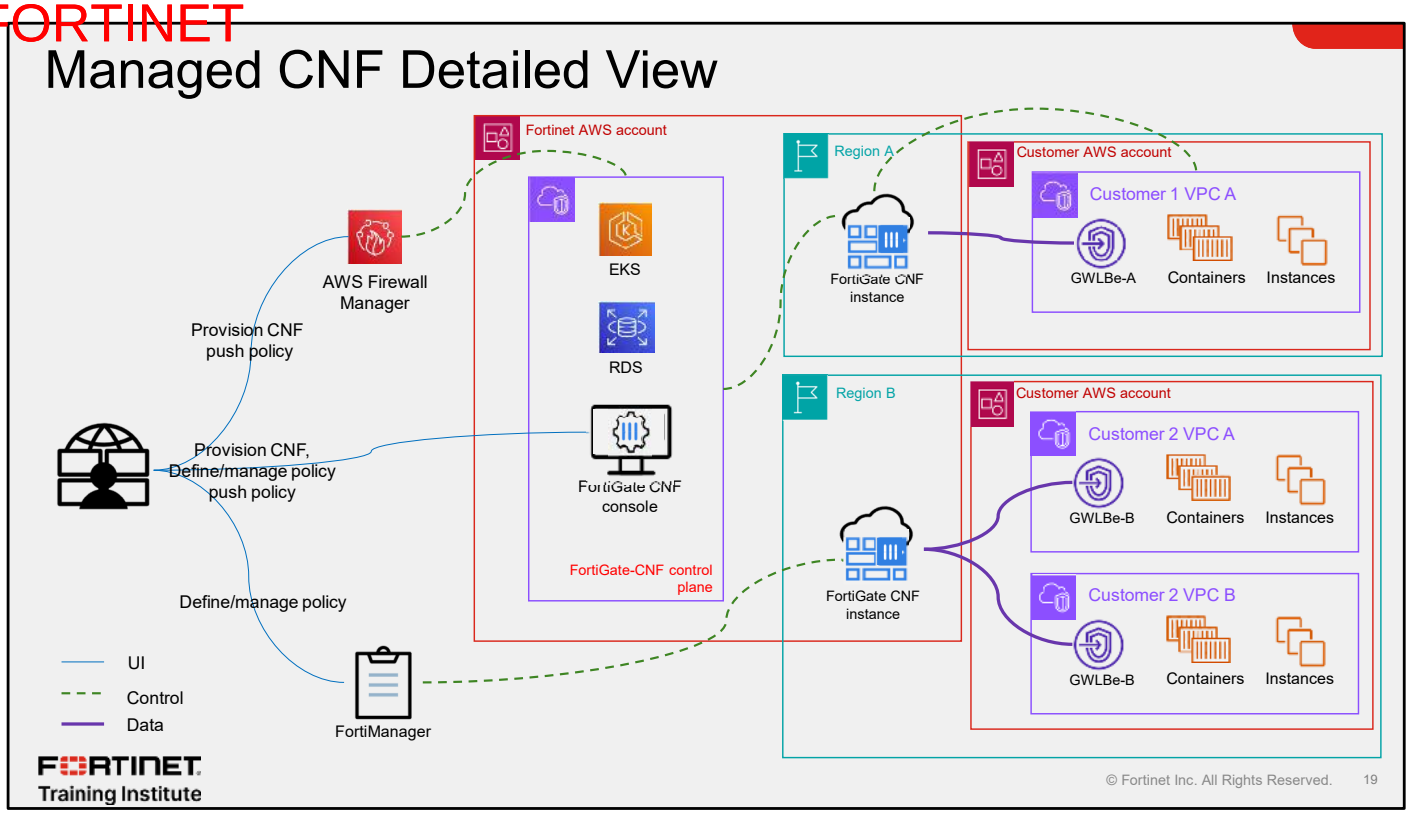

After you set up a CNF instance, you can configure it. This slide shows the management process for FortiGate CNF. You can use AWS Firewall Manager, FortiManager, or the FortiGate CNF console to configure your instances. Note that each of the three management options have varying levels of control. A CNF instance cannot span regions and you can associate a GWLBe only with the instance it its own region. When FortiGate CNF is correctly configured, the GWLBe forwards traffic to the FortiGate CNF instance for security inspection.

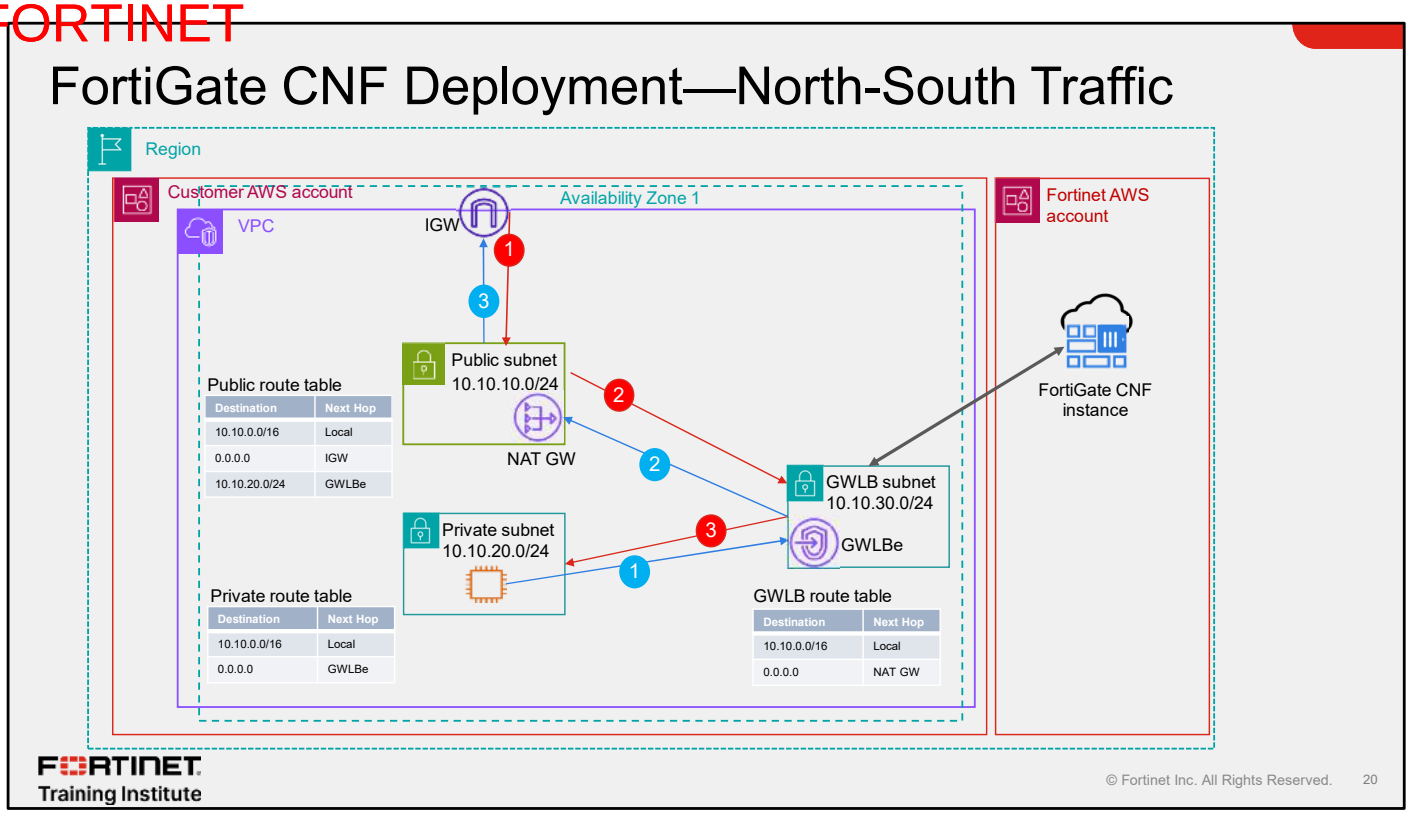

This slide shows a basic FortiGate CNF deployment for north-south traffic inspection.

For outgoing traffic to the internet, packets originating from the EC2 instance in the private subnet are sent to the GWLB subnet through the GWLBe, as defined in the private route table. The traffic is inspected by the FortiGate CNF instance, which resides in Fortinet's AWS account, and is then forwarded to the NAT gateway, as defined in the GWLB route table. When the packets arrive in the public subnet, they are forwarded to the internet gateway (IGW) and out to the internet.

For incoming traffic to the private subnet, packets arriving in the public subnet from the IGW are forwarded to the GWLB subnet, as defined in the public route table, to be inspected by FortiGate CNF. The inspected packets are then forwarded to the private subnet, as per the local route defined in the GWLB route table.

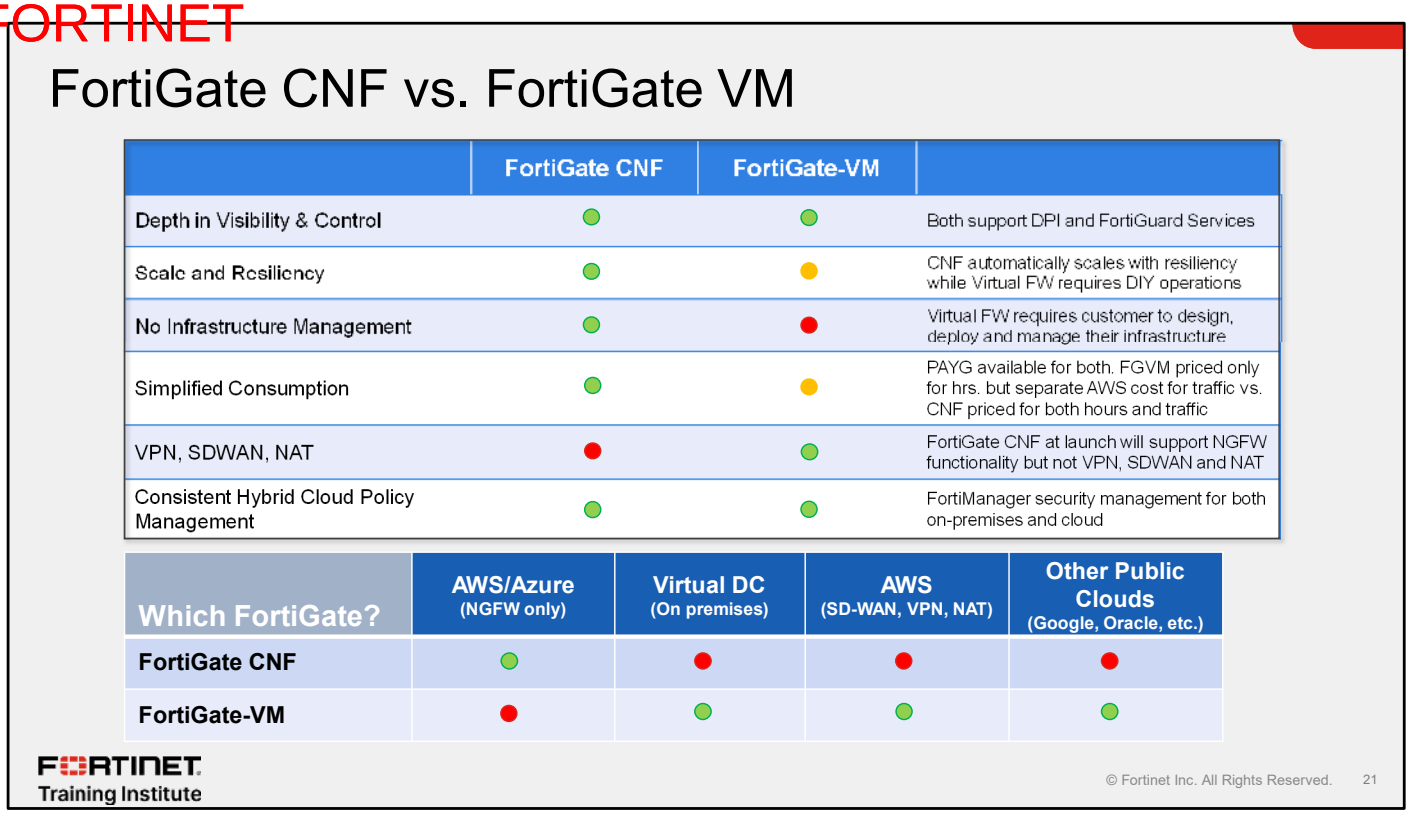

This slide compares FortiGate CNF and FortiGate VM. FortiGate CNF is the better choice if you need only NGFW features. It is also the better choice if ease of management and scalability are important factors. If you require your FortiGate to have VPN, SD-WAN, or NAT functionality, FortiGate VM is the obvious choice. You can manage both, using FortiManager.

Knowledge Check

1. Which AWS component integrates with FortiGate CNF?
   A. NLB
   ✓ B. GWLBe

2. Which statement about FortiGate CNF instances is correct?
   A. You can connect only one AWS account to it.
   ✓ B. You can provision it using AWS Firewall Manager.

FORTINET Training Institute

 22

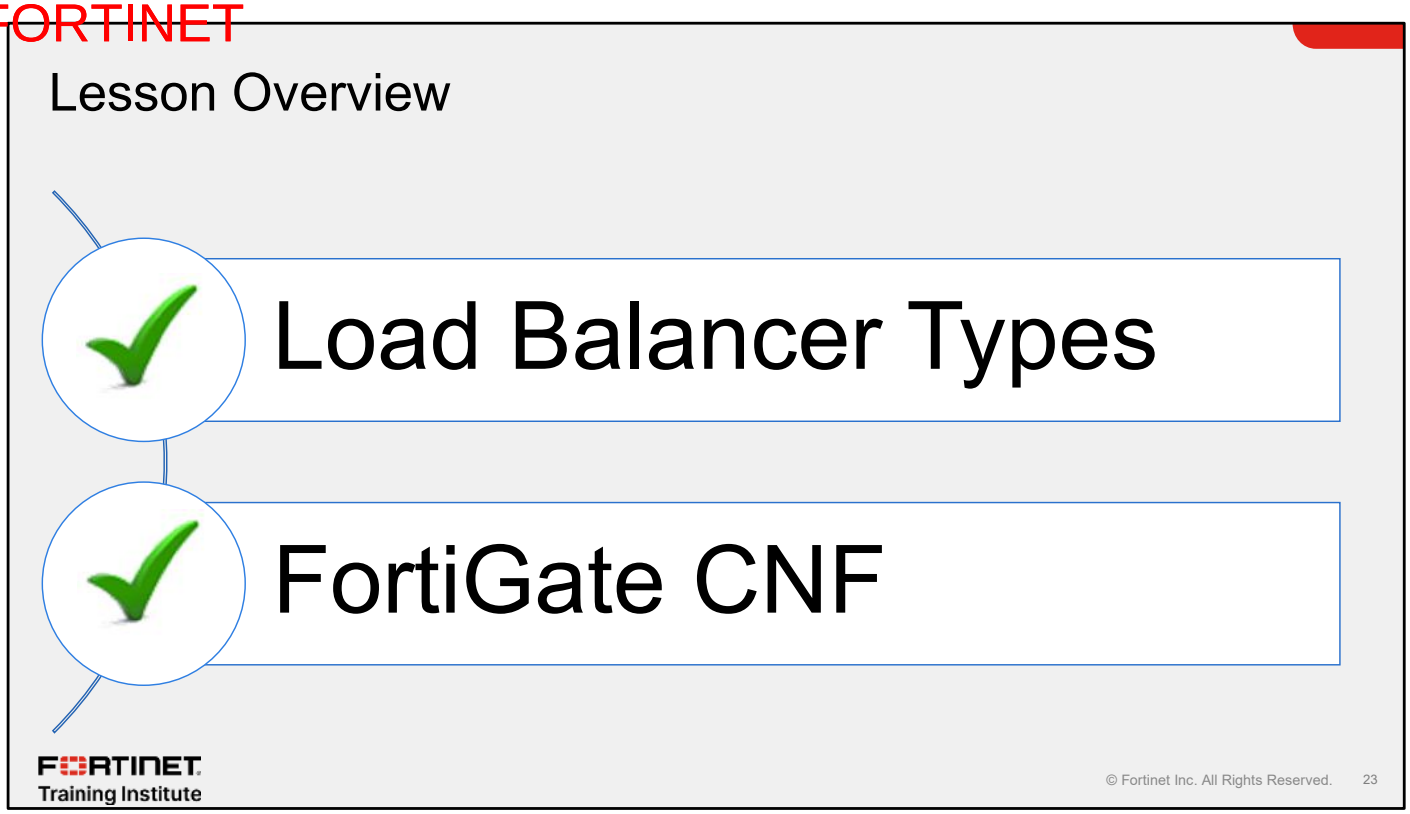

Good job. You now understand FortiGate CNF and will be able to deploy a FWaaS.

Review

- ✓ Understand AWS ELB
- ✓ Understand FortiGate active-active HA with AWS ELB
- ✓ Understand AWS GWLB
- ✓ Understand FortiGate CNF
- ✓ Identify the differences between FortiGate CNF and FortiGate VM

FORTINET Training Institute

 24

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

By mastering the objectives covered in this lesson, you learned about different types of load balancers in AWS and how to use different Fortinet solutions alongside them.

FORTINET