DO NOT REPRINT
© FORTINET



SD-WAN
Study Guide

for FortiOS 7.2

FORTINET®
Training Institute

DO NOT REPRINT
© FORTINET

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

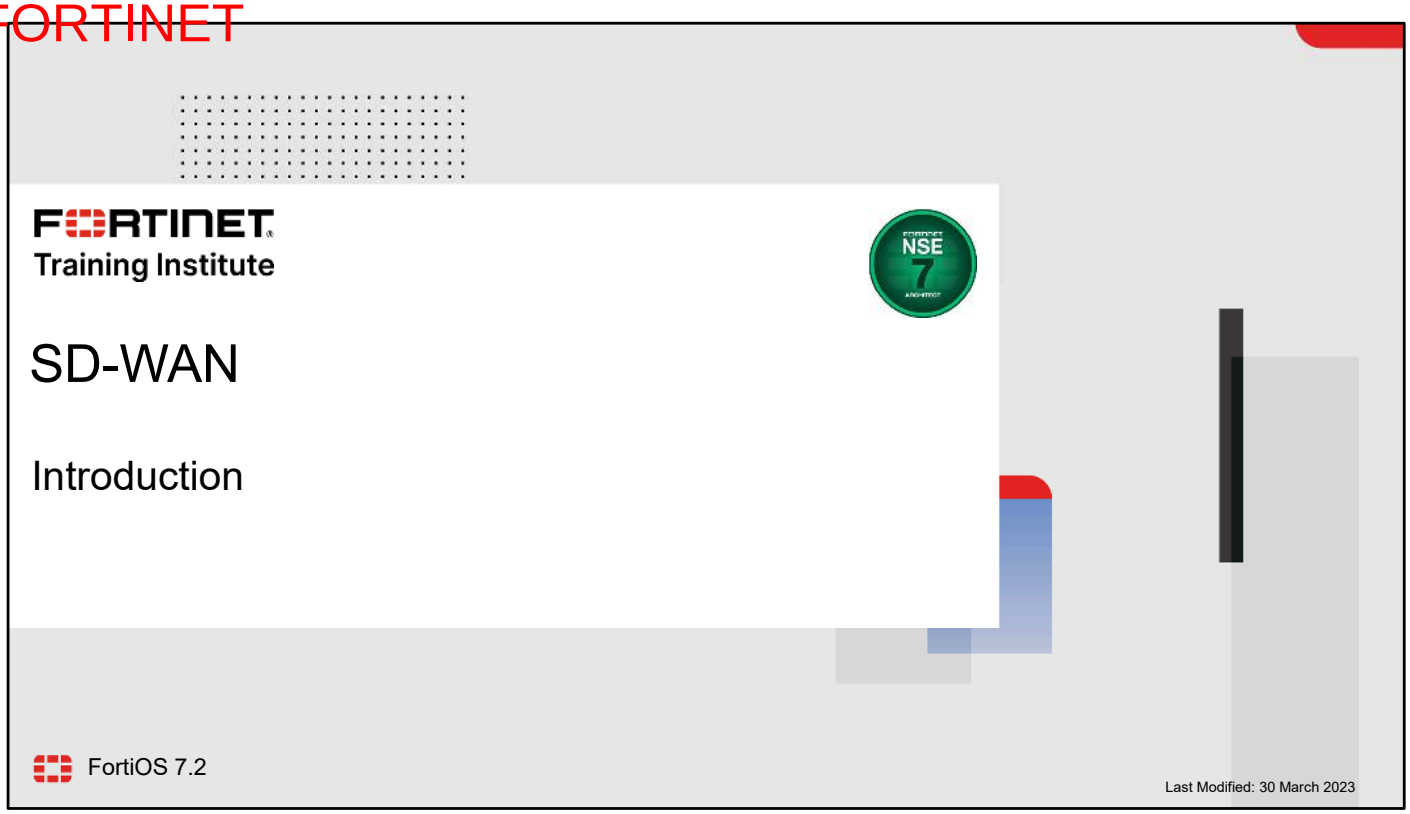
3/30/2023

DO NOT REPRINT
© FORTINET

TABLE OF CONTENTS

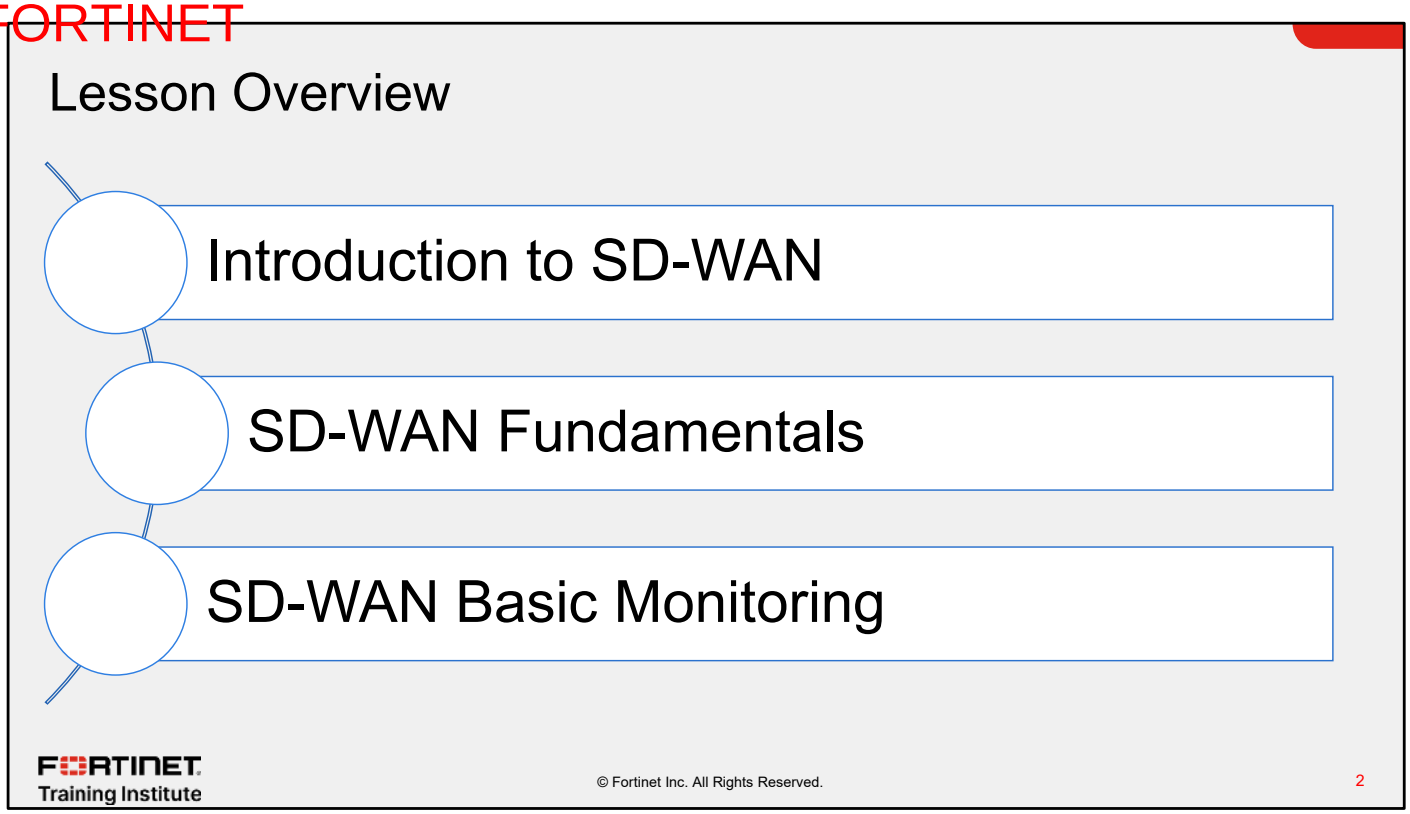
01 Introduction 4
02 Centralized Management 37
03 Members, Zones, and Performance SLAs 80
04 Routing and Sessions 122
05 Rules 167
06 SD-WAN Overlay Design and Best Practices 229
07 Monitoring and Troubleshooting 293
Solution Slides 323

DO NOT REPRINT
© FORTINET



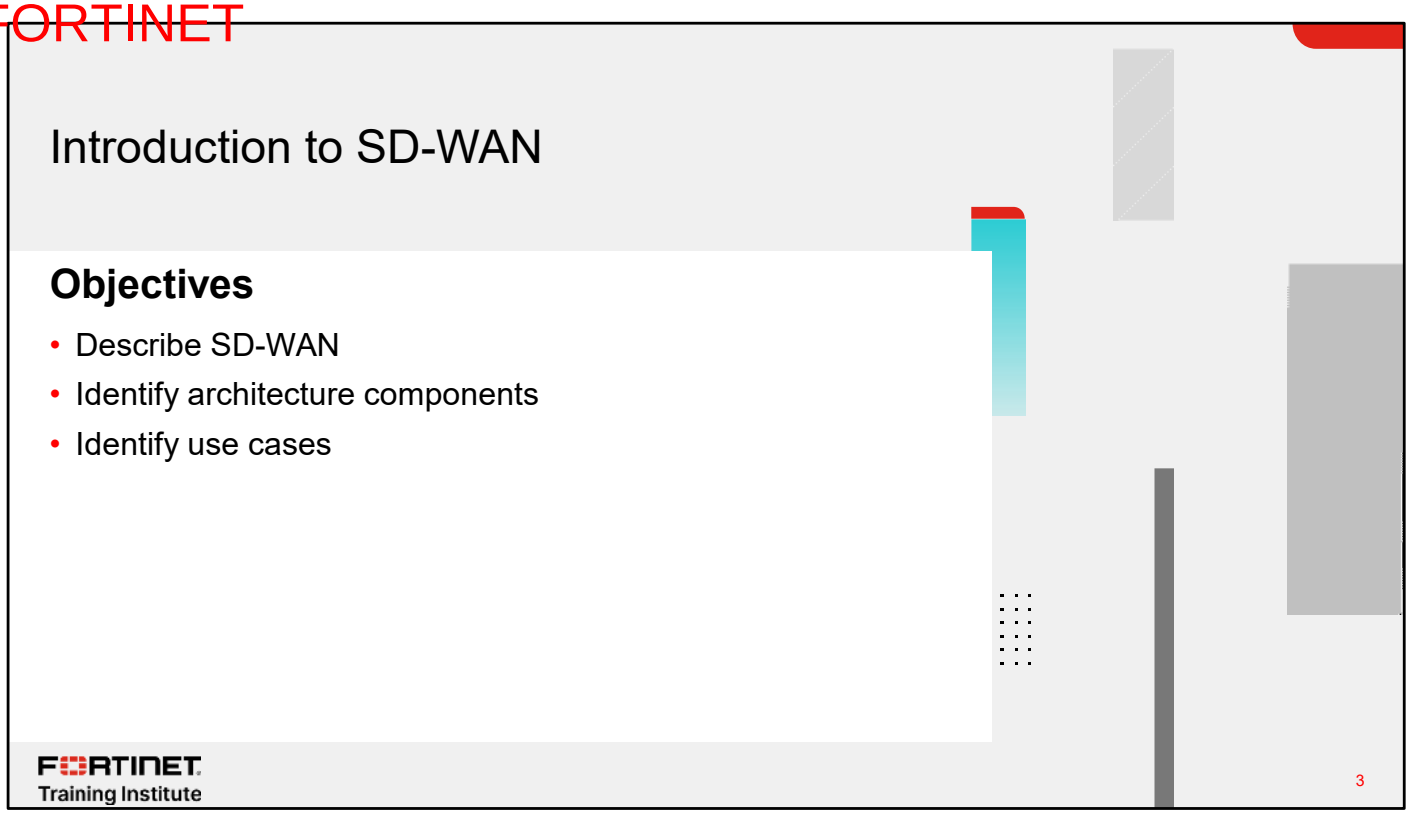
In this lesson, you will be introduced to the SD-WAN concept, use cases, and main components. You will also learn how to use the FortiGate GUI to configure and monitor a basic SD-WAN setup.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



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

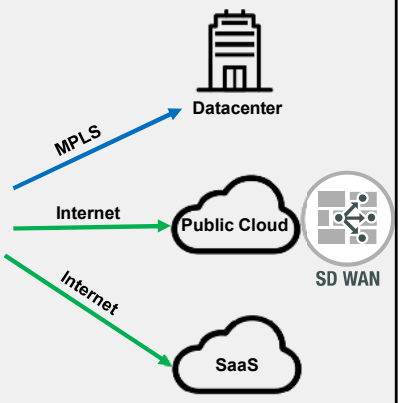
By understanding an SD-WAN solution and its use cases, you should be able to identify the most common scenarios where SD-WAN can be deployed to distribute traffic across your WAN links effectively and securely.

DO NOT REPRINT
© FORTINET

What Is SD-WAN?

- Software-defined approach to steer WAN traffic using:
  - Flexible user-defined rules
    - Protocol and service-based traffic matching
    - Application-awareness
    - Dynamic link selection
  - Controls egress traffic
- Secure SD-WAN
  - Fortinet SD-WAN implementation (built-in security)
- Benefits:
  - Effective WAN usage
  - Improved application performance
  - Cost reduction





FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

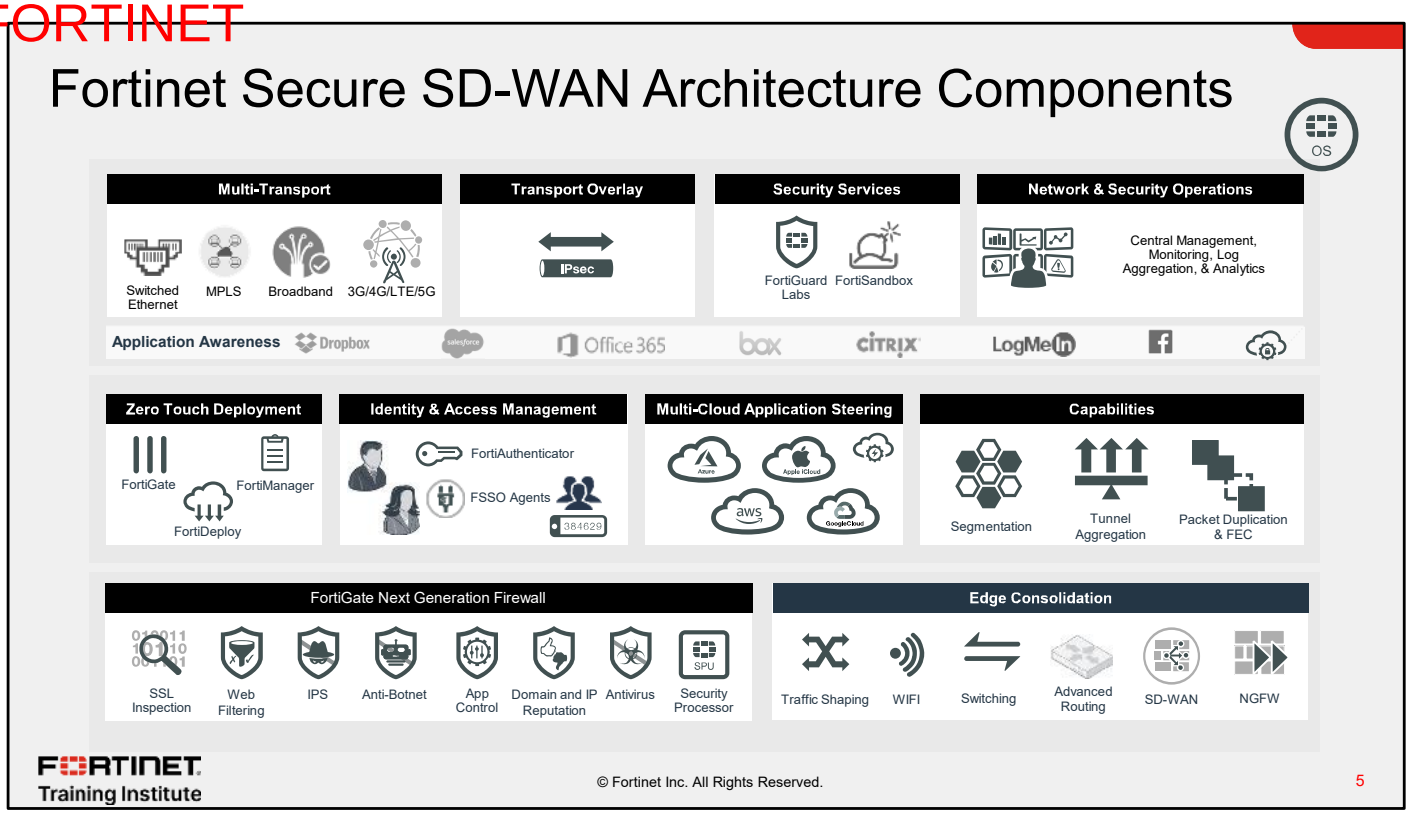
According to Gartner, software-defined WAN (SD-WAN) provides dynamic, policy-based, application path selection across multiple WAN connections and supports service chaining for additional services, such as WAN optimization and firewalls. Fortinet implementation of SD-WAN is called secure SD-WAN because it also provides security by leveraging the built-in security features available in FortiOS.

Secure SD-WAN relies on well-known FortiOS features such as IPsec, auto-discovery VPN (ADVPN), link monitoring, advanced routing, internet services database (ISDB), traffic shaping, UTM inspection, and load balancing. The administrator can then combine these features and set rules that define how FortiGate steers traffic across the WAN based on multiple factors, such as the protocol, service, or application identified for the traffic; and the quality of the links. Note that SD-WAN controls egress traffic, not ingress traffic. This means that the return traffic may use a different link from the one SD-WAN chose for egress.

One benefit of SD-WAN is effective WAN usage. That is, you can use public (for example, broadband, LTE) and private (for example, MPLS) links to securely steer traffic to different destinations: internet, public cloud, private cloud, and the corporate network. This approach of using different types of links to connect sites to private and public networks is known as hybrid WAN. Using a hybrid WAN reduces costs mainly because administrators usually steer traffic over low-cost fast internet links more than over high-cost slow private links. The result is that private links, such as MPLS links, are often used to steer critical traffic only, or as failover links for high availability.

Another benefit of SD-WAN is improved application performance because you can steer traffic through the best link that meets the application requirements. During congestion, you can leverage traffic shaping to prioritize sensitive and critical applications over less important ones. Also, the support of ADVPN shortcuts enables SD-WAN to use direct IPsec tunnels between sites to steer traffic, resulting in lower latency for traffic between the sites (spokes), and less load on the central locations (hubs).

DO NOT REPRINT
© FORTINET



This slide shows the architecture components of the Fortinet Secure SD-WAN solution.

The core component of the architecture is FortiGate. When you use FortiGate to deploy SD-WAN, you leverage the existing connectivity, management, and next-generation firewall (NGFW) features supported by FortiOS. This means that you can consolidate SD-WAN and security in a single device, thus the term Secure SD-WAN.

Because FortiGate is one of the core components in the Fortinet Security Fabric, then by extension, your SD-WAN deployment can also benefit from many of the features supported by other products in the fabric. For example, you could use FortiManager to perform zero touch provisioning for branches that require a SD-WAN setup. Similarly, you could use FortiSwitch to connect your WAN edge and LAN edge devices of your SD-WAN branch.

One key benefit of using FortiOS and FortiGate for your SD-WAN solution, is the ability to perform application steering in SD-WAN. FortiGate inspects traffic using its IPS engine and the application signatures provided by FortiGuard to identify thousands of applications. The result is that you can configure FortiGate to steer traffic based on the application detected, rather than ports, protocols, and IP addresses. The fact that FortiGate can identify the application regardless of the Layer 3 and Layer 4 information on the packet, enables you to significantly reduce administrative overhead and scale easier when deploying new applications and sites.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Direct Internet Access

- Traffic steered across multiple physical internet links
- Typical operation:
  - Critical/sensitive traffic expedited and steered over best performing links
  - Costly links used for critical traffic or failover
  - Static default routing
- Upstream and downstream speeds can be:
  - Set manually by administrator
  - Set automatically by SD-WAN WAN monitoring service

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

Direct internet access (DIA), also known as local breakout, is arguably the most common use case for SD-WAN. A site has multiple internet links (also known as underlay links), and the administrator wants FortiGate to steer internet traffic across the links. The links are connected to FortiGate using different types of physical interfaces: physical port, VLAN, link aggregation (LAG), USB modem, or through FortiExtender.

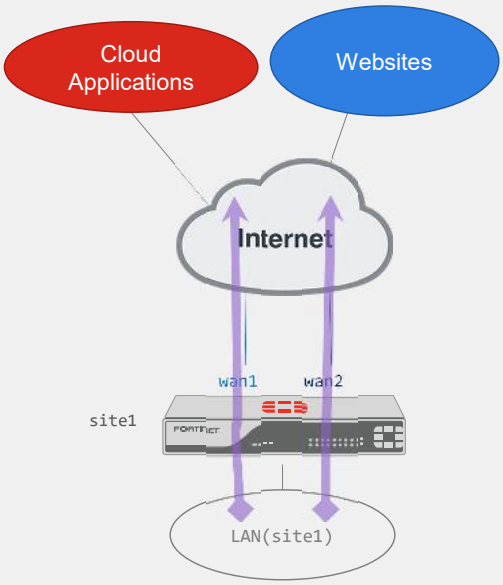
Usually, sensitive traffic is expedited and steered over the best performing links, while non-critical traffic is distributed across one or more links using a best effort approach. Costly internet links are commonly used as backup links, or to steer critical traffic only.

Because the internet traffic leaves the organization boundaries directly on the local site, administrators usually enforce strict security policies on the internet traffic. For routing, a typical configuration makes use of static default routes. However, in some cases, BGP is used between the ISP and FortiGate, especially if the site must advertise a public IP prefix.

Administrators can also manually define the upstream and downstream speeds of each link to prevent saturation during traffic distribution. Alternatively, they can configure FortiGate to use the SD-WAN bandwidth monitoring service to run speed tests against FortiGuard, and then automatically adjust the upstream and downstream speeds of the links based on the test results.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Direct Internet Access (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

This slide shows an example of DIA.

FortiGate has two internet links. One link is connected to wan1 and the other to wan2. FortiGate uses both links to steer traffic sourced from the LAN and destined to cloud applications and websites on the internet.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Site-To-Site Traffic

- Use overlay links to steer site-to-site corporate traffic and internet (RIA)
  - Overlay: IPsec, GRE, IP-in-IP (established over underlays)
  - Underlay: physical links (internet and MPLS circuits, 3G/4G/LTE modems, and so on)
- Typical operation:
  - Dial-up IPsec used for overlay
  - If using ADVPN:
    - Offload site-to-site traffic to shortcuts
    - SD-WAN checks health of shortcuts
    - Apply full security inspection on spokes
  - Dynamic routing
    - More scalable
    - IBGP recommended for ADVPN
- Hub-to-spoke upstream speed can be:
  - Measured automatically for traffic shaping

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

You can use SD-WAN to steer corporate site-to-site traffic. Usually, companies follow a hub-and-spoke topology, and use VPN tunnels—typically dial-up IPsec tunnels—to transport the traffic between the sites. The tunnels (also known as overlay links) are established over internet and MPLS links (also known as underlay links). Tunnels can also carry internet traffic from a spoke to a hub, where it then breaks out to the internet. This is also known as remote internet access (RIA), and you will learn more about it in this lesson.

SD-WAN can monitor the link quality of the tunnels and select the best performing link for sensitive and critical traffic. If using ADVPN, SD-WAN can offload the traffic from a parent tunnel to a shortcut, thus reducing latency for spoke-to-spoke traffic. SD-WAN can also monitor the health of the shortcut tunnels to ensure they meet the configured service-level agreement (SLA).

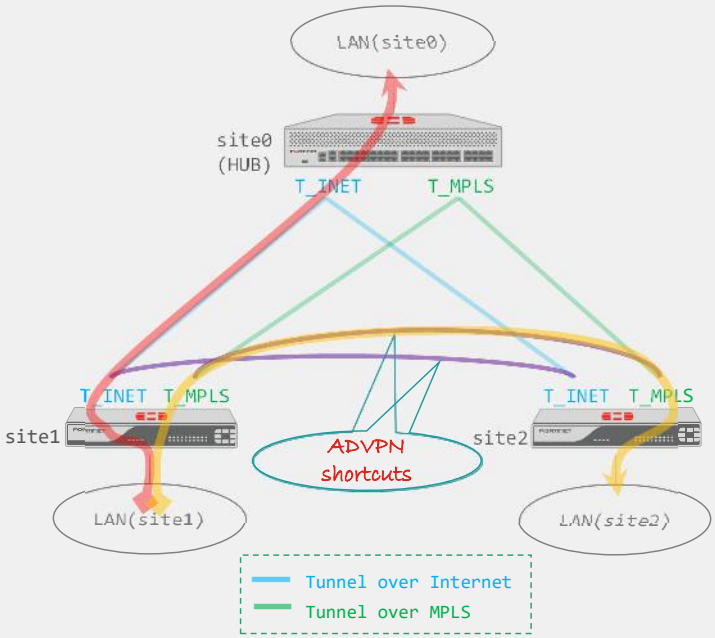
If using ADVPN, you should apply all necessary security inspection on the local site because spoke-to-spoke traffic will eventually flow directly through the shortcut and will therefore bypass any inspection enabled on the hub. If not using ADVPN, you may consider applying a less restrictive policy on the spoke provided you configure the hub to perform the additional required inspection.

For routing, a dynamic routing protocol, such as BGP, is often used to exchange routing information through the tunnels and scale easier when adding new sites. If using ADVPN, internal BGP (IBGP) is recommended to preserve next hop information.

Similar to DIA, the hub FortiGate can run speed tests against the spokes to determine the upstream speed of tunnels. The hub FortiGate can then apply the speed test result as the upstream speed on the tunnel for traffic shaping purposes.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Site-To-Site Traffic (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

This slide shows an example of a deployment that uses SD-WAN to steer site-to-site traffic.

Each site has two overlays configured, one using the internet underlay and the other the MPLS underlay. SD-WAN steers spoke-to-spoke and spoke-to-hub traffic. Because ADVPN is configured on the network, shortcuts are automatically established between spokes when spoke-to-spoke traffic is sent across the network. SD-WAN can then automatically offload the spoke-to-spoke traffic from parent tunnels to shortcuts, thus improving performance. SD-WAN also monitors the health of shortcuts and dynamically makes steering decisions based on their health and availability.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—RIA

- Internet traffic steered across overlay links to:
  - Centralize inspection on hub
  - Improve performance if DIA performance is poor
  - Provide internet access if DIA is unavailable
- Typical operation:
  - Internet traffic is steered to MPLS
  - Hub performs thorough inspection

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

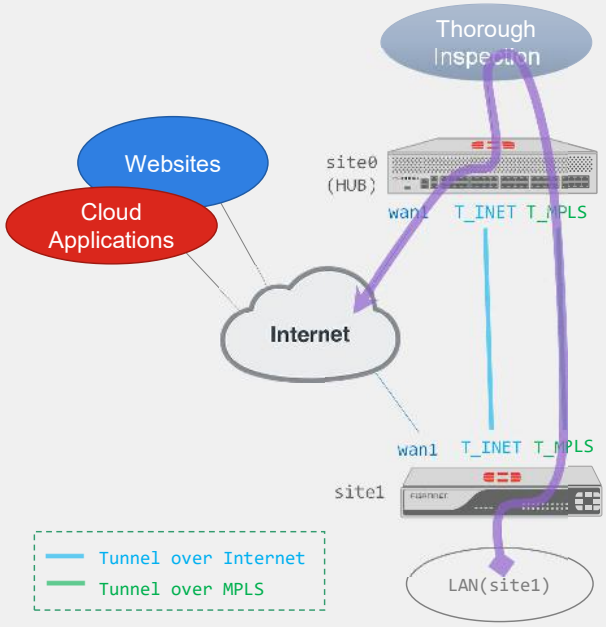
RIA, also known as remote breakout, is another use case for SD-WAN. Internet traffic from the spokes is backhauled through the WAN using overlay links. When the traffic arrives the hub, it breaks out to the internet.

The most common reason to use RIA is to centralize security inspection and internet access on the hub. For example, you can have a central high-end FortiGate device that inspects all the internet traffic that leaves the organization and that conforms with the company policy, instead of having each low-end spoke FortiGate device to inspect traffic, thus reducing costs and administrative overhead.

Another reason to use RIA is for DIA backup. For example, you could configure FortiGate to steer internet traffic through an MPLS link if the performance measured for internet applications on internet links is worse than on MPLS links, or simply if the internet links become unavailable.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—RIA (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

This slide shows an example of RIA.

Instead of using the local internet underlay to forward internet traffic, the FortiGate device on site 1 steers internet traffic to the hub through the overlay built over MPLS. Once the traffic reaches the hub, the traffic is subject to a thorough security inspection before it breaks out to the internet.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Cloud On-Ramp

- Steer cloud application traffic to overlays against cloud PoPs:
  - High performance
  - Secure access
- Establish overlays against:
  - Built-in cloud gateway
    - Unidirectional Secure SD-WAN
  - Cloud FortiGate VMs
    - Bidirectional Secure SD-WAN
- Typical operation:
  - Direct connection from spokes
    - Usually, the best performance
  - Backhaul traffic through WAN
    - Required by company policy
    - Best performance (premium WAN links)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

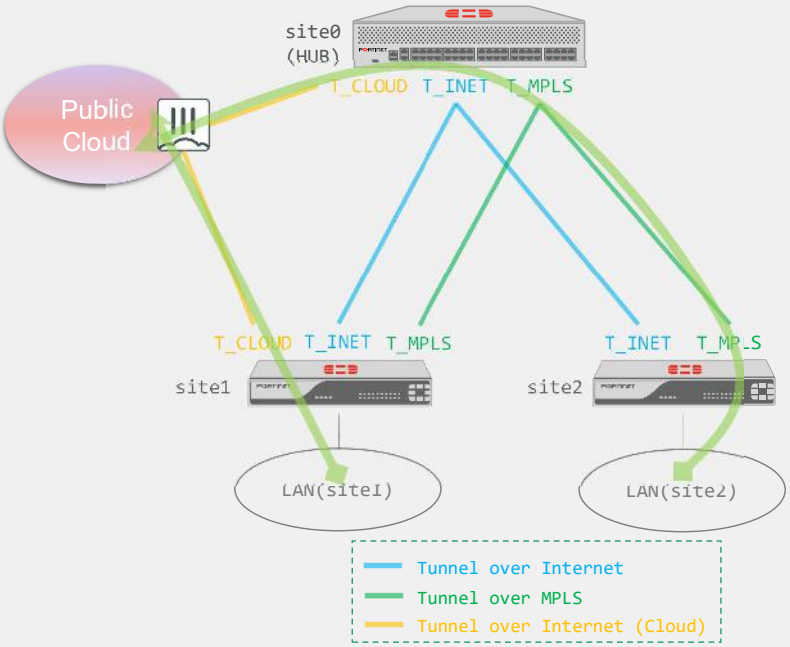
12

To improve performance of cloud applications while keeping network traffic secure, you can configure overlays against the closest point of presence (PoP) offered by the cloud provider in the area, thus reducing latency. You can configure FortiGate to connect to the cloud provider's built-in VPN gateway. Alternatively, you can deploy a FortiGate VM in the cloud and establish the overlays against it. Choosing to deploy a FortiGate VM in the cloud enables you to use FortiOS security features on traffic entering or leaving the cloud, as well as to use SD-WAN to steer sessions originated from the cloud.

For performing reasons, a cloud on-ramp connection is usually established directly from the spoke because a direct connection often results in the lowest latency. However, in some cases, traffic could be backhauled through the WAN using overlay links because of company policy, or because cloud tunnels are only available on the hub. Another reason to backhaul cloud traffic through the WAN can be performance. For example, a company may use premium WAN links between the spokes and the hub, and between the hub and the cloud. The performance of the premium links is so high that backhauling the traffic through the WAN would result in a better user experience than accessing the local PoP directly from the spoke.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Cloud On-Ramp (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

This slide shows an example of cloud on-ramp.

The FortiGate device on site 1 has a direct tunnel against a FortiGate VM running on the cloud. This tunnel is used to steer traffic of cloud applications for performance and security reasons. Also, unlike site 1, the FortiGate device on site 2 uses an overlay to the hub to send traffic of cloud applications. The hub then routes the traffic through its local cloud tunnel.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Dual Hub

- Dual SD-WAN attachment
  - Each site connects to two SD-WAN gateways—hubs
  - Geographical redundancy
- Typical operation
  - Dial-up IPsec used for overlay
  - Each SD-WAN gateway acts as dial-up IPsec server
  - Overlays grouped per device and location
  - Dynamic routing
    - More scalable
    - IBGP recommended
- ADVPN possible
  - All sites must preserve original BGP next-hop values
  - Two options
    - BGP route reflector function
    - P2 selectors for ADVPN route exchange

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

For larger deployments, you can extend the SD-WAN base design to include a second SD-WAN gateway. This second gateway can be in a different location to provide geo-redundancy.

In such topologies, branch SD-WAN devices connect to two (or more) SD-WAN gateways. FortiGate steers traffic to one gateway or the other, according to the status, link health, and SLA of the devices.

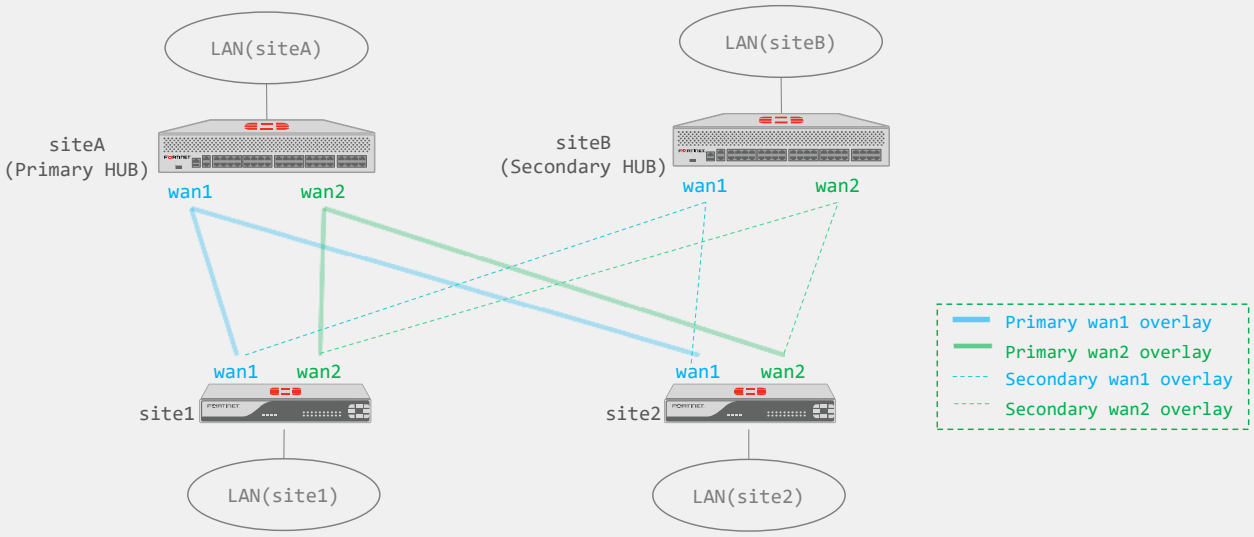
For easier management and rule definition, the administrator will group overlay links per devices and per location.

In addition to geo-redundancy, intra-site redundancy with HA cluster—FGCP or FGSP— remains possible.

You can combine geo-redundant topology with use of ADVPN. In such scenarios, all sites must preserve original BGP next-hop values. Site prefixes must remain unchanged, including their original BGP next-hop value. A common way to achieve this is with the route reflector function. Alternatively, you can use Phase2 selectors for ADVPN route exchange. You will learn more about both options in another lesson.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Dual Hub (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

This slide shows an example of a deployment that uses SD-WAN to steer traffic with dual hub geo-redundancy.

Each site has four overlays configured, two to the primary site—one on each WAN underlay link—and two to the secondary site. SD-WAN steers spoke-to-hub traffic to the primary hub and continues to monitor the health of all overlay links. According to the rules definitions, SD-WAN triggers a fall back to the secondary hub, based on link quality or site availability.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Multi-Region

- Increased scalability
- Geographically expanded solution
  - High number of spokes
  - Group devices per region
  - Full mesh connectivity between hubs
- Typical operation
  - EBGP between regional gateways
  - IBGP within regions
  - ADVPN within regions
  - Traffic across regions through regional hubs

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

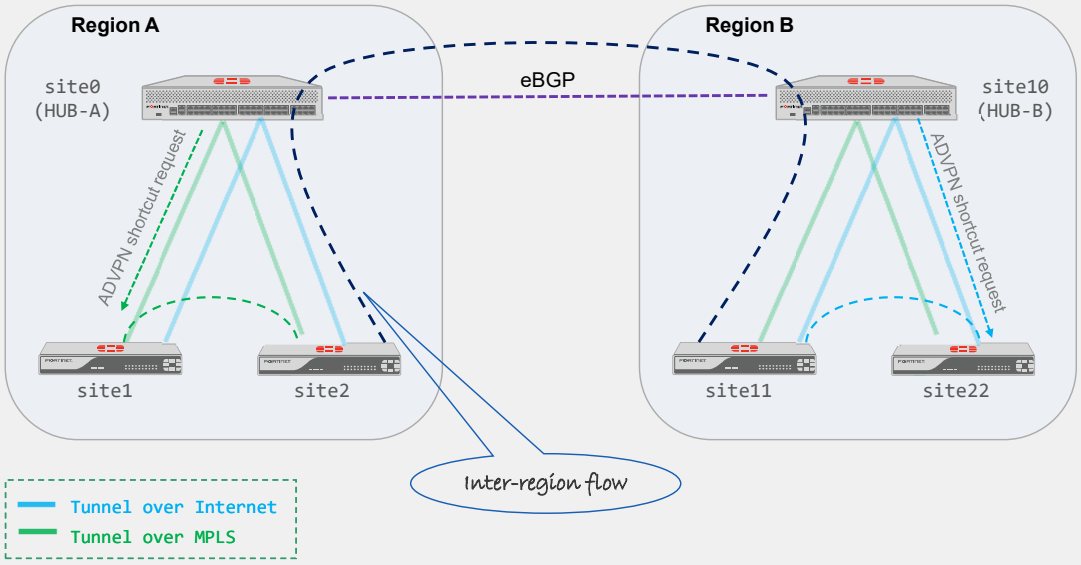
For increased scalability, as your solution expands geographically, and the number of sites grows, you might decide to segment the traffic and management of devices into multiple regions. For each region, you can define either a single-hub or dual-hub SD-WAN topology.

In each region, branch devices connect to the regional SD-WAN gateway. Then, regional gateways are interconnected between regions in a full-mesh design. When a user from one region needs to connect to a user or an application in another region, traffic will go through regional gateways.

You should consider a design like this when traffic exchanged within each region is significantly higher than traffic through regions. Usually, regions correspond to geographical area but, according to expected traffic flows, you might want to consider a different segmentation. It could be shops and factory or any other logical grouping applicable to your business.

DO NOT REPRINT
© FORTINET

SD-WAN Use Cases—Multi-Region (Contd)



FORTINET Training Institute

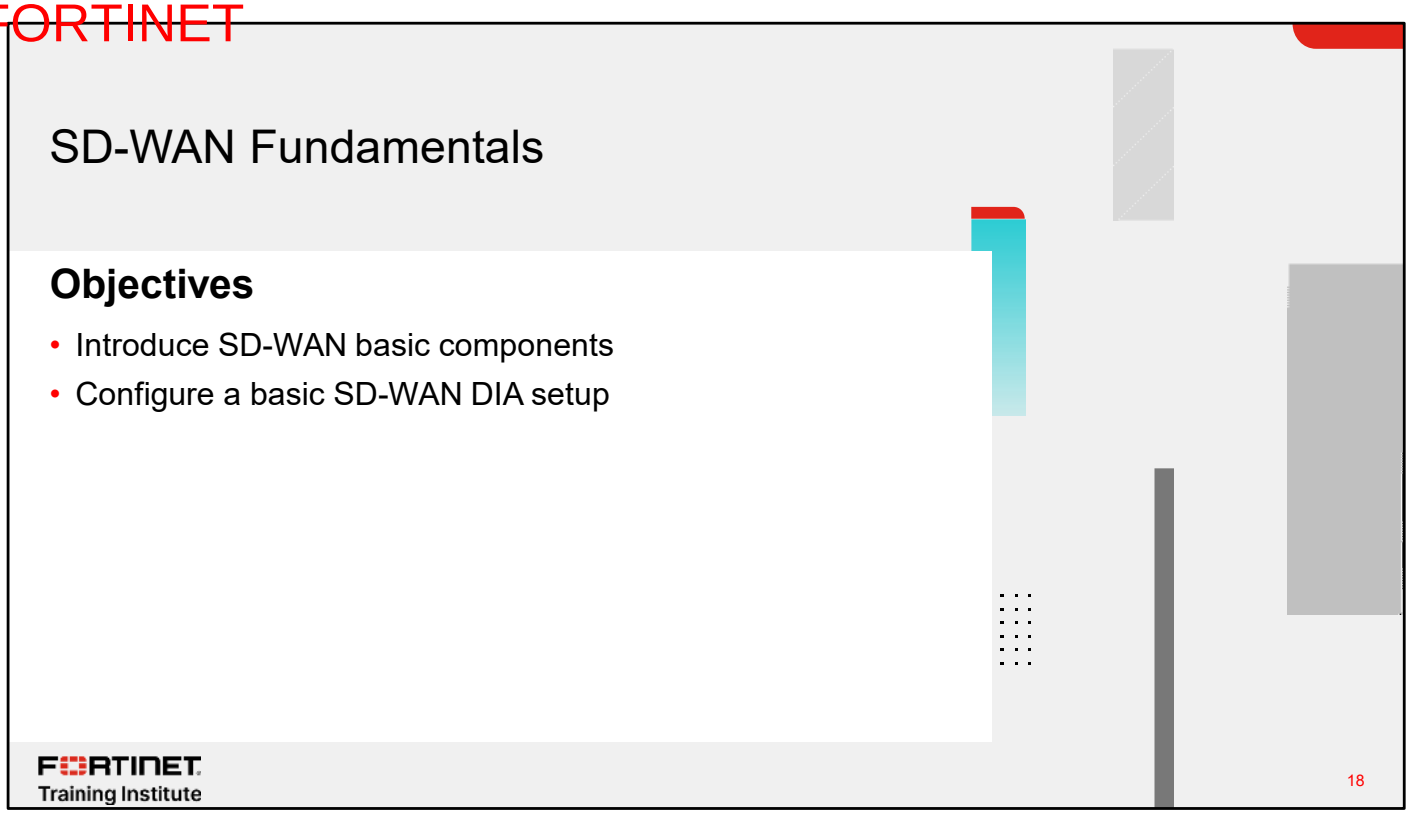
© Fortinet Inc. All Rights Reserved.

17

This slide shows an example of a multi-region topology.

Each region has its own SD-WAN topology, which can be single or dual hub. ADVPN shortcuts can be established between devices in each region, but inter-region traffic must always flow through regional hubs.

DO NOT REPRINT
© FORTINET



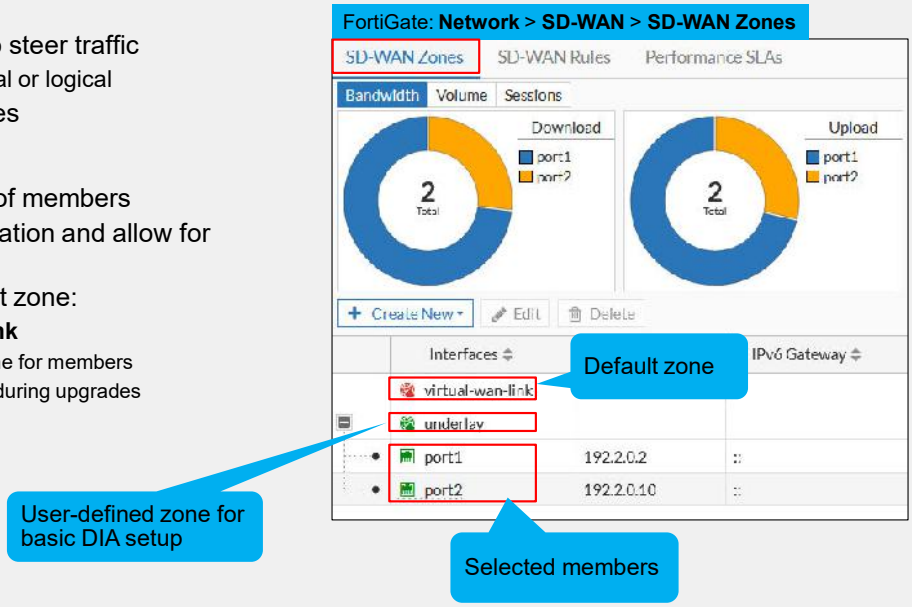
After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in SD-WAN fundamentals, you should be able to understand the basics of SD-WAN, including how to configure a basic DIA setup.

DO NOT REPRINT
© FORTINET

SD-WAN Members and Zones

- Members
  - Interfaces used to steer traffic
    - Can be physical or logical
  - Organized in zones
- Zones
  - Logical grouping of members
  - Optimize configuration and allow for segmentation
  - Predefined default zone:
    - **virtual-wan-link**
      - Default zone for members
      - Also used during upgrades



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

The first step to configure SD-WAN is to define the members and assign them to zones. This configuration is done in the **SD-WAN Zones** page.

Members (also known as links) are existing physical or logical FortiOS interfaces that you select to be part of SD-WAN. The interfaces are then used to steer traffic based on the SD-WAN rules configured.

When you configure a member in SD-WAN, you must assign it to a zone, and optionally set a gateway. Zones are logical groupings of interfaces. The interfaces in a zone have similar configuration requirements. Like FortiGate interface zones, the goal with SD-WAN zones is to reference them in the configuration instead of individual members to optimize the configuration by avoiding duplicate settings, and to achieve network segmentation. When set, the **Gateway** setting is used as the next hop to forward traffic through the member.

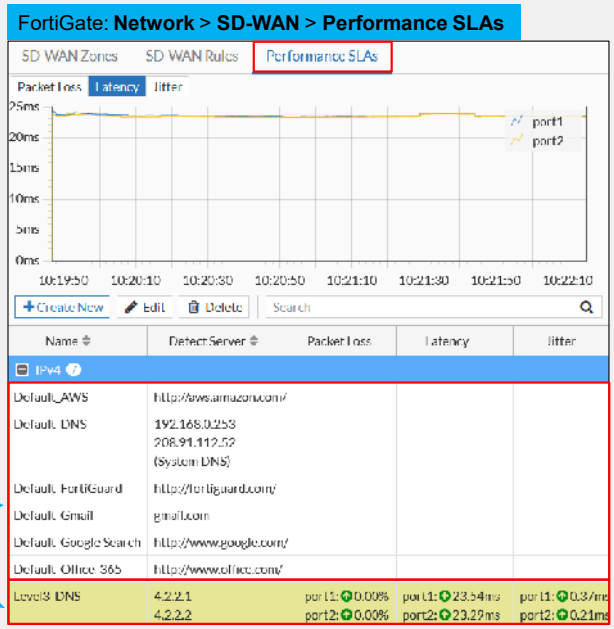
One zone is created by default: **virtual-wan-link**. The zone, **virtual-wan-link,** is the default zone where members are placed when you create them. The default zone is also used when you upgrade a FortiGate device running a version with no support for zones to a version that supports zones. During the upgrade, FortiGate places all existing members to the virtual-wan-link zone.

The example on this slide shows the default SD-WAN zone—**virtual-wan-link**—and one user-defined zone: **underlay**. The **underlay** zone contains **port1** and **port2** as members, which are used for a basic DIA setup. Note that, although the zone is named **underlay** because it contains such type of members, you can assign any name you like.

DO NOT REPRINT
© FORTINET

Performance SLAs

- Performs member health check:
  - State
    - Alive or dead
  - Performance:
    - Packet loss, latency, and jitter
- Usage:
  - Detect unavailable links
  - Monitor performance criteria
- Health can be measured:
  - Actively
    - Based on periodic probes sent to configured servers
  - Passively
    - Based on member traffic



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

After you define your SD-WAN members and assign them to zones, you will probably want to monitor the health of your SD-WAN members on the **Performance SLAs** page.

FortiGate performance SLAs monitor the state of each member—whether it is alive or dead—and measures the member packet loss, latency, and jitter. SD-WAN then uses the member health information to make traffic steering decisions based on the configured SD-WAN rules. For example, you can instruct FortiGate to steer internet traffic to a member, provided the member is alive and its latency doesn't exceed a given threshold. Performance SLAs will also detect situations where the interface is physically up, but FortiGate is unable to reach the desired destination and flags the corresponding link as dead.

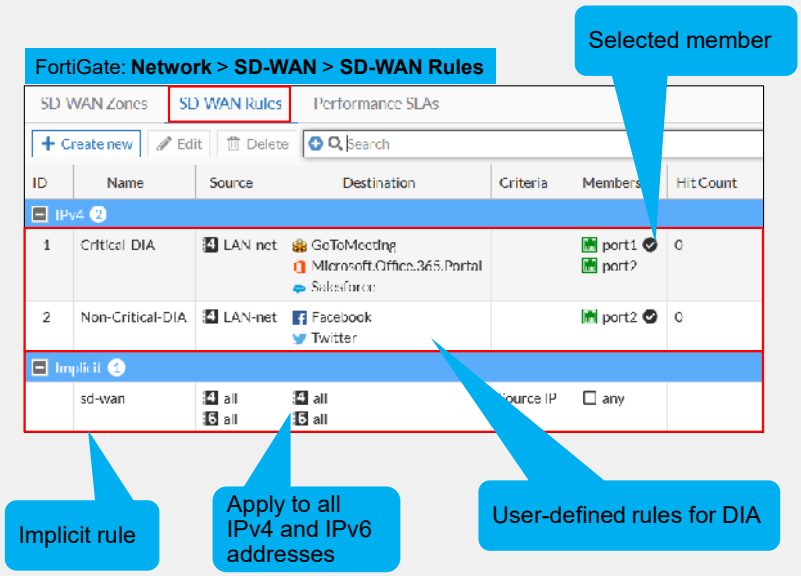
When you configure a performance SLA, there are a several entries created by default that you can choose from. The default entries measure the health of members against well-known internet services, such as FortiGuard, Google Search, and Amazon AWS. Alternatively, you can create your own entry and choose whether you want to monitor the health actively or passively. In active monitoring, the health of the member is checked by periodically—by default every 500ms— sending probes from the member to one or two servers that act as a beacon. In passive monitoring, the health of a member is determined based on the traffic passing through the member.

The example on this slide shows a new entry named **Level3_DNS**. The entry contains the well-known **4.2.2.1** and **4.2.2.2** DNS servers, both of which are used to monitor the health of **port1** and **port2**. The results show that the members are alive (green arrow), report no packet loss, and have average values for delay and jitter over the internet.

DO NOT REPRINT
© FORTINET

SD-WAN Rules

- Define steering rules based on:
  - Matching traffic criteria
    - Define pattern, ISDB, and application to match
  - Member preference
    - Define list of preferred members and zones to steer traffic to
  - Member performance
    - Define SLA to meet for members
- Evaluated from top to down:
  - Rules are used to steer traffic
    - Firewall policy required
  - Implicit rule
    - Used if user-defined rules are not matched



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

After you configure the zones, members, and performance SLAs, it's time to define the traffic steering rules for SD-WAN. This is done on the **SD-WAN Rules** page.

When you configure an SD-WAN rule, you first define the application or traffic pattern to match. After that, you indicate the preferred members, or zones, to steer the matching traffic to, and in some cases, the performance metrics that the member must meet to be eligible to receive and forward the traffic.

SD-WAN rules are evaluated in the same way as firewall policies: from top to bottom, using the first match. However, unlike firewall policies, they are used to steer traffic, *not* to allow traffic. When you use SD-WAN rules, you must configure corresponding firewall policies to allow the SD-WAN traffic.

There is an implicit SD-WAN rule created by default. If none of the user-defined SD-WAN rules are matched, then the implicit rule is used. By default, the implicit rule load balances the traffic across all available SD-WAN members.

The example on this slide shows two user-defined rules named **Critical-DIA** and **Non-Critical-DIA**, which are used to steer traffic in this basic DIA setup. The **Critical-DIA** steers **GoToMeeting**, **Microsoft.Office.365.Portal**, and **Salesforce** traffic to the member with the lowest latency, between **port1** and **port2**. The example shows that **port1** is selected because it is the member with the check mark beside it. The **Non-Critical-DIA** rule steers **Facebook** and **Twitter** traffic to **port2**.

The implicit rule, located at the bottom of the list, is used if none of the two user-defined rules are matched. You can see that it applies to all IPv4 and IPv6 source and destination addresses. The icons beside the object names—4 and 6—distinguish IPv4 objects from IPv6 objects.

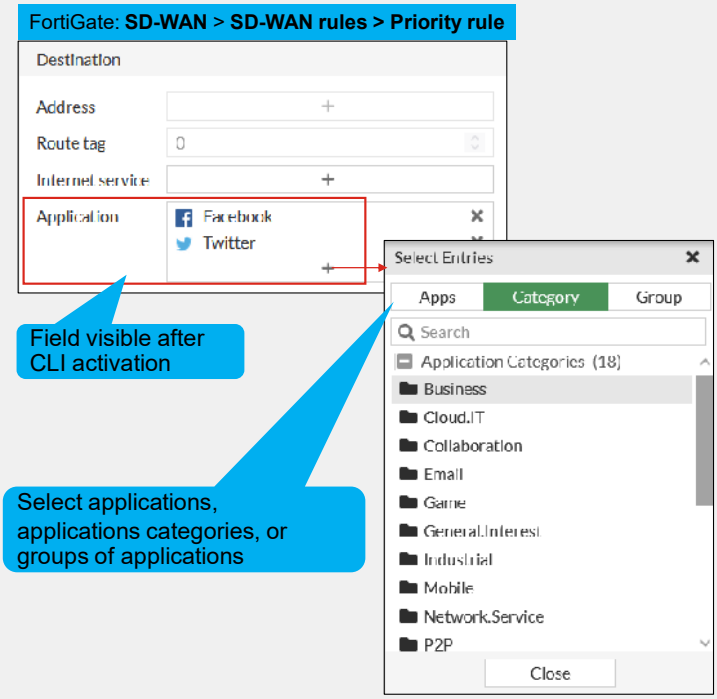
DO NOT REPRINT
© FORTINET

Application Detection

- Application detection is defined per:
  - Application
  - Application category
  - Group of applications
- Criteria visibility must be enabled with a CLI command

```
config system global
  set gui-app-detection-sdwan enable
end
```

- Enable application control policy
  - Required for application detection by SD-WAN rules



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

For well-known applications, you can rely on FortiGuard services and define SD-WAN rules to steer traffic per application or application category. For instance, you can decide to direct leisure applications like games or Facebook to low-cost links and reserve high-quality, costly links for business traffic.

Visibility of application detection criteria is, by default, hidden on the FortiGate GUI. You must enable feature visibility on the CLI using the global command `set gui-app-detection-sdwan enable`. Note that if GUI visibility is disabled, configuration for application criteria remains active and configuration is still visible on the CLI.

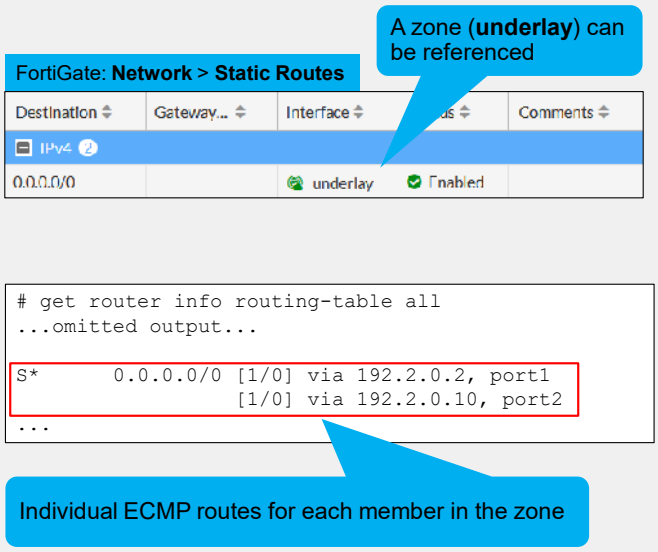
In addition to applications and application groups, you can also select application categories as SD-WAN rule destination criteria for IPv4 rules.

To determine which applications flowing through your network math the defined applications, use the CLI command `diagnose sys sdwan internet-service-app-ctrl-category-list <id>`.

DO NOT REPRINT
© FORTINET

Routing

- Valid route required for steering traffic to member
- Static and dynamic routes supported
- Static routes
  - Reference a zone
    - Common case, simplified configuration
    - Individual ECMP routes installed for each member in the zone
    - Gateway obtained from member configuration
  - Reference a member
    - More granular control



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

SD-WAN rules define the traffic steering policies in SD-WAN. However, traffic won't be forwarded to an SD-WAN member unless there is a valid route that matches the destination address of the traffic through the SD-WAN member.

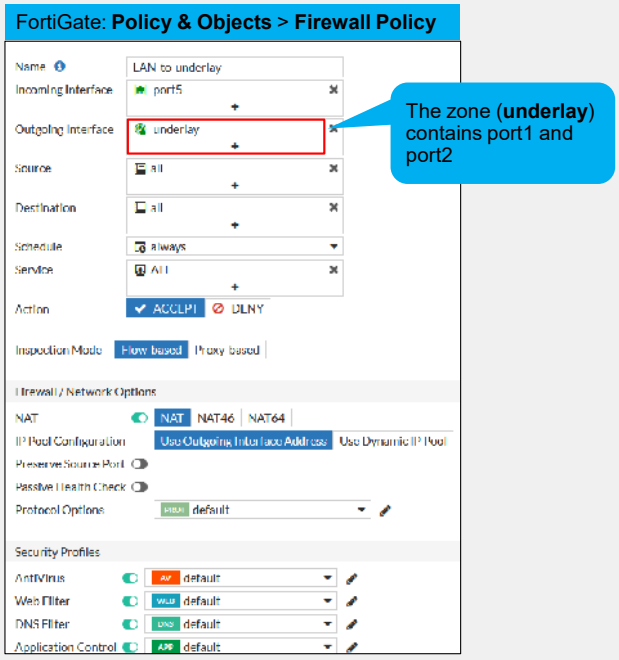
You can use static and dynamic routing in SD-WAN. This slide shows an example of a static default route configured for the **underlay** zone, which is used to route traffic in our basic DIA setup. When you configure static routes for SD-WAN, you usually reference an SD-WAN zone to simplify the configuration. However, you can also reference individual members instead, so you can have a more granular control over traffic. When you reference a zone in a static route, FortiGate installs individual routes for each member in the zone. The individual routes are then displayed in the routing table as equal cost multi-path (ECMP) routes.

Note that when you configure a static route that references a zone, you don't have to specify a gateway address because FortiGate retrieves it from the member configuration.

DO NOT REPRINT
© FORTINET

Firewall Policies

- Steered traffic must be allowed by a firewall policy
- Reference zones only
  - Simplified configuration
- Can't reference a member directly
  - Create zones with single members



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

In addition to having a valid route, you must also have a firewall policy that allows the traffic steered by SD-WAN.

You configure SD-WAN firewall policies in the same way as regular firewall policies, except that, when selecting an outgoing or incoming interface, you must reference an SD-WAN zone. When you reference a zone, you simplify the configuration by avoiding duplicate firewall policies.

You may need to reference a member in your firewall policy because you want to apply a different action on the traffic flowing through that member, such as applying different security profiles and NAT settings. Because you can't reference members in a firewall policy, a workaround is to place a single member in a separate zone, and then reference that zone in the firewall policy.

The example on this slide shows a firewall policy named **LAN-to-underlay** that references the **underlay** zone, which contains port1 and port2 as members. As a result, DIA traffic steered over port1 or port2 will be allowed by FortiGate provided it matches the firewall policy criteria, and it passes the security inspection configured on the policy.

DO NOT REPRINT
© FORTINET

SD-WAN CLI Configuration

- Zone, member, and performance SLA:

```
config system sdwan
    set status enable
    config zone
        edit "underlay"
        next
    end
    config members
        edit 1
            set interface "port1"
            set zone "underlay"
        next
        edit 2
            set interface "port2"
            set zone "underlay"
        next
    end
    config health-check
        edit "Level3_DNS"
            set server "4.2.2.1" "4.2.2.2"
            set members 1 2
        next
    end
...
```

- Rules:

```
...
    config service
        edit 1
            set name "Critical-DIA"
            set src "all"
            set internet-service enable
            set internet-service-app-ctrl 16354 41468 16920
            set priority-zone "underlay"
        next
        edit 2
            set name "Non-Critical-DIA"
            set src "all"
            set internet-service enable
            set internet-service-app-ctrl 15832 16001
            set priority-members 2
        next
    end
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

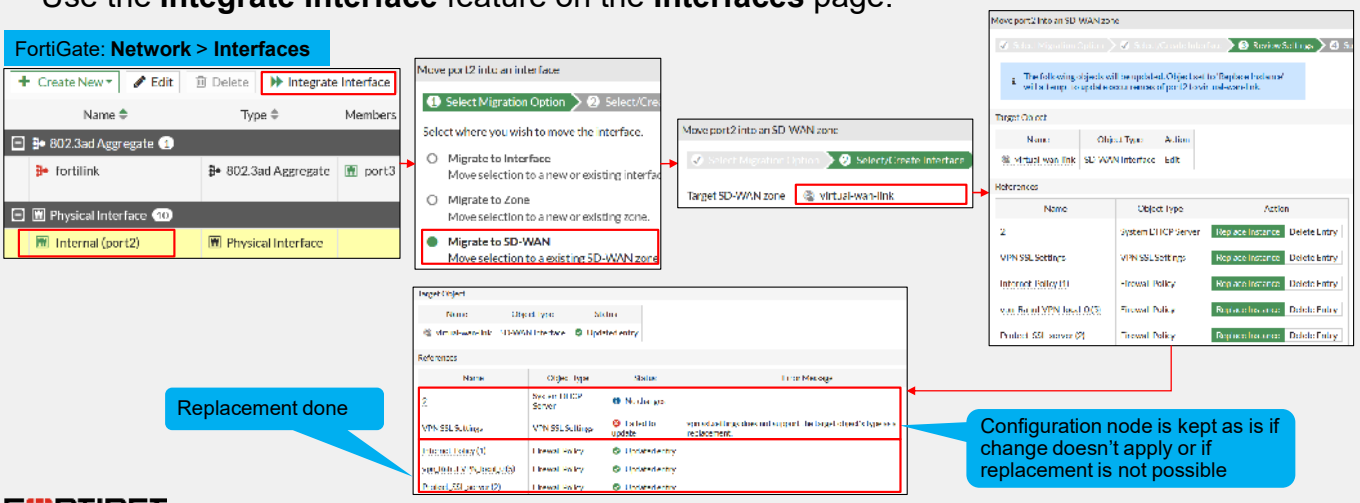
This slide shows the equivalent CLI configuration for the basic SD-WAN DIA setup described so far. The SD-WAN specific configuration is found under `config system sdwan`. Inside `config system sdwan`, there are separate configuration sections for each SD-WAN component.

The example shown on this slide breaks down the CLI configuration into two. The first portion displays the zone, member, and performance SLA configuration. The second, the SD-WAN rules configuration. Note that FortiOS uses the terms `health-check` and `service` to refer to the performance SLA and rule configuration on the CLI, respectively. You will explore in more details the CLI configuration in other lessons.

DO NOT REPRINT
© FORTINET

Migrating Interface to SD-WAN

- Replace references to individual interfaces automatically with selected SD-WAN zone
- Save time and reduce service disruption
- Use the **Integrate Interface** feature on the **Interfaces** page:



FORTINET Training Institute

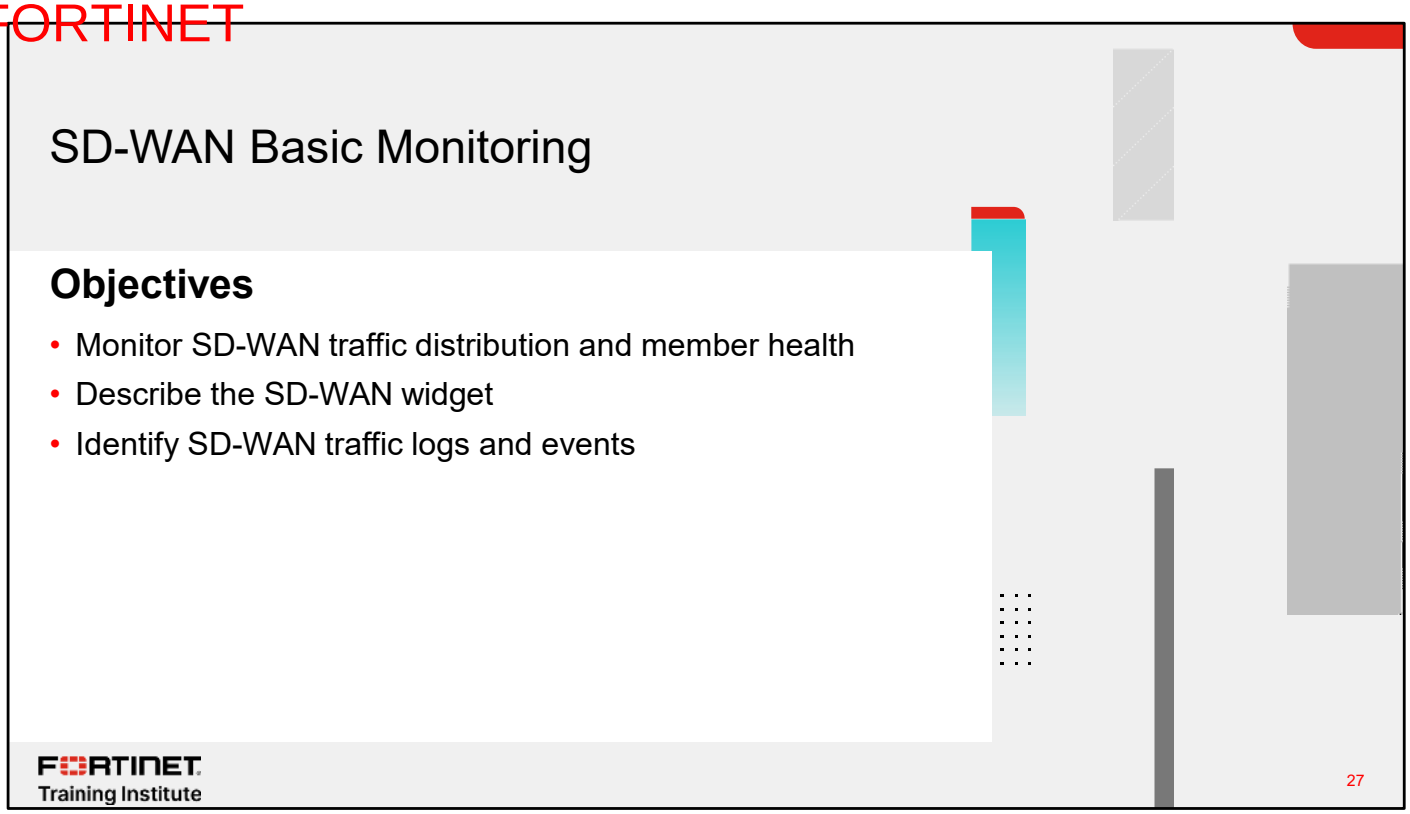
© Fortinet Inc. All Rights Reserved.

26

Before adding an interface as SD-WAN member, you must first remove any configuration references to the interface. This is fine if your configuration is simple, but if your configuration has a considerable number of references, then the reference removal and adding process can be very time consuming and disruptive to the network.

An alternative is to use the **Integrate Interface** feature available on the **Interfaces** page on the FortiGate GUI. When you use the integrate interface feature, you can instruct FortiGate to migrate an interface to SD-WAN. The result is that FortiGate automatically tries to replace the individual interface with the selected SD-WAN zone on every configuration object that references the interface. Note that If the change does not apply to a configuration node or if FortiGate can't replace the reference, then FortiGate leaves the configuration node as is.

DO NOT REPRINT
© FORTINET



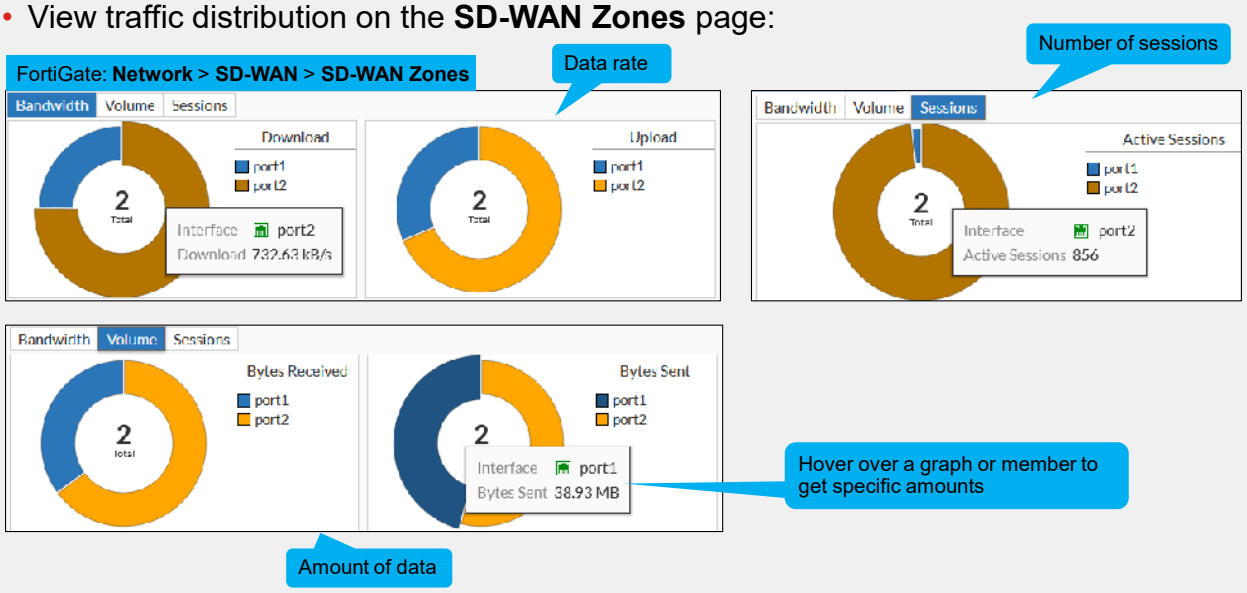
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding basic monitoring of SD-WAN, you should be able to identify the different tools available on the FortiGate GUI to check SD-WAN traffic distribution, health, and events.

DO NOT REPRINT
© FORTINET

Traffic Distribution

- View traffic distribution on the **SD-WAN Zones** page:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

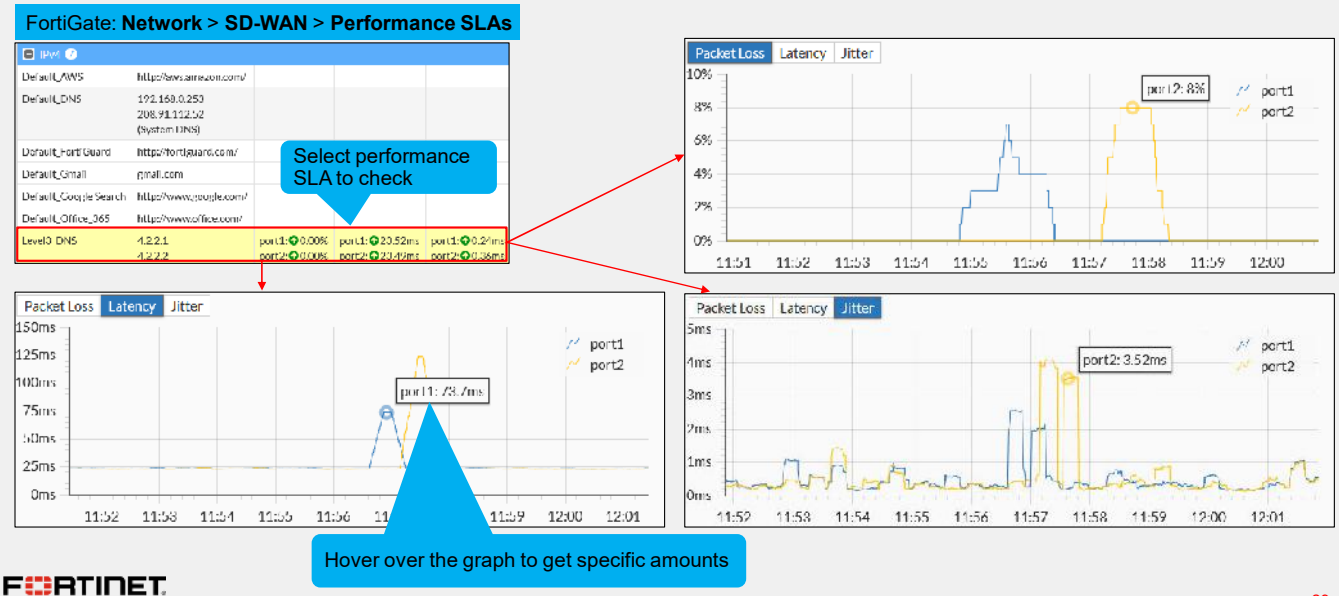
You can browse to the **SD-WAN Zones** page to monitor the traffic distribution over the SD-WAN members. The page contains graphs that display traffic distribution based on bandwidth, volume, or sessions. Note that bandwidth refers to the data rate, while volume refers to the amount of data.

You can also hover over a member or the graph to get a specific amount of bandwidth, volume, or sessions. The example on this slide shows the corresponding traffic distribution graphs of the basic DIA setup.

DO NOT REPRINT
© FORTINET

Member Health

- View member health on the **Performance SLAs** page (last 10 minutes):



You can browse to the **Performance SLAs** page to monitor the health of your members.

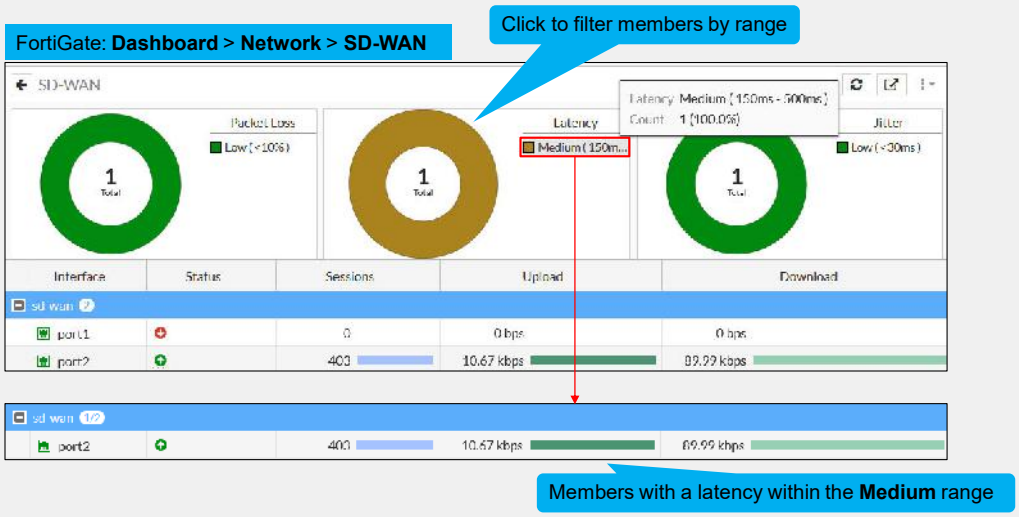
You first select the performance SLA you want to check (Level3_DNS in the example). The graphs on the page will then display the packet loss, latency, and jitter of each member using the selected performance SLA. Note that the information shown in the graphs is limited to the last 10 minutes.

You can also hover over the graph to get a specific amount of packet loss, latency, or jitter. The example on this slide shows the corresponding health graphs for the two members used in the basic DIA setup.

DO NOT REPRINT
© FORTINET

SD-WAN Widget

- Consolidated view of member health and utilization



FORTINET Training Institute

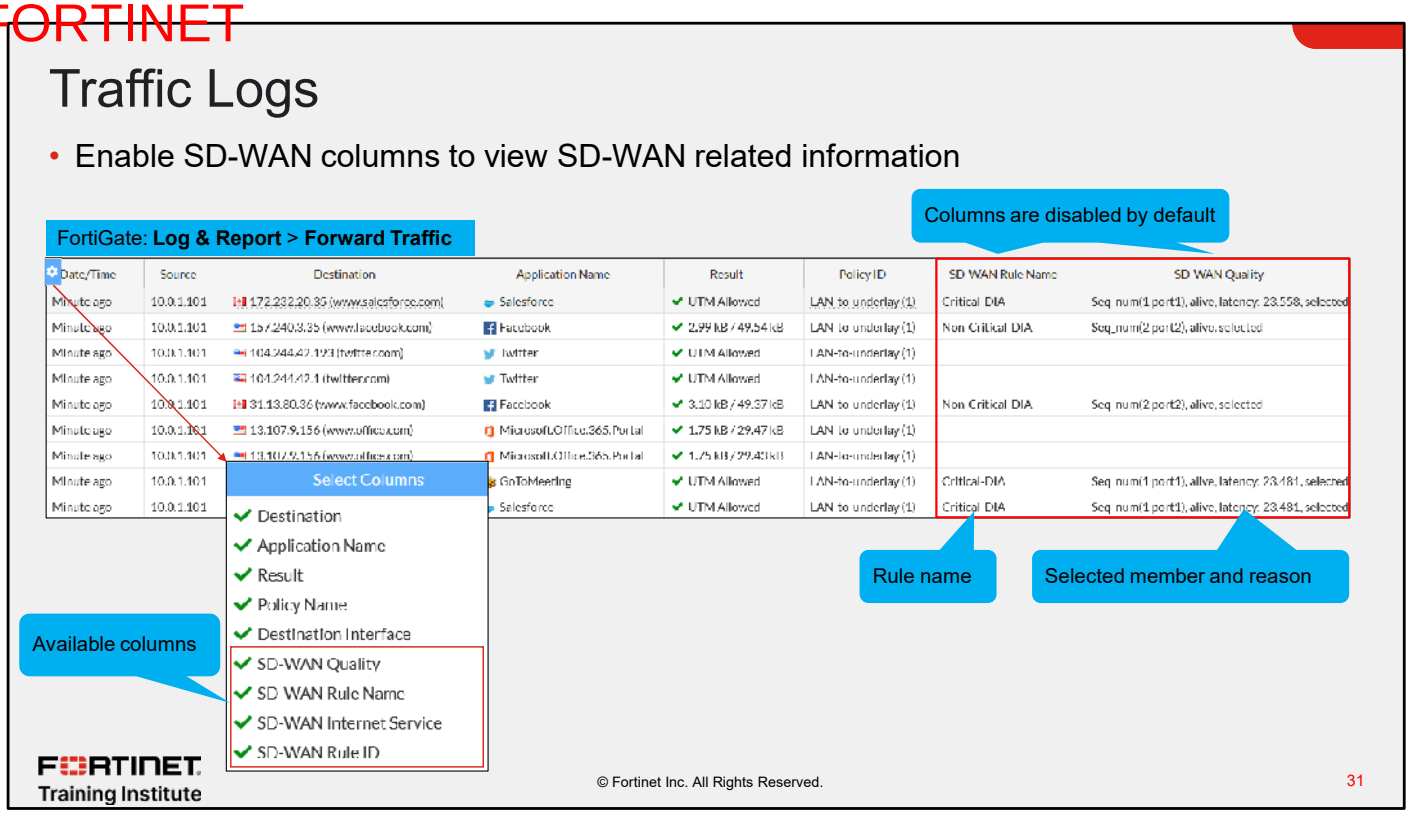
© Fortinet Inc. All Rights Reserved.

30

The SD-WAN widget offers a consolidated view of both the health of a member and its usage. The example on this slide shows two SD-WAN members configured: port1 and port2. However, port1 is currently down, which is why there is only traffic going through port2.

The graphs on the page summarize the health of the alive members—only port2 in this case—and indicate how many of them fall within some predefined ranges of packet loss, latency, and jitter. The example shows that packet loss and jitter on port2 is within the low range, and its latency is within the medium range. You can click a range to display the list of members that fall within that range.

DO NOT REPRINT
© FORTINET



The **Forward Traffic** logs page is useful to identify how sessions are distributed in SD-WAN and the reason. Make sure to enable the **SD-WAN Rule Name** and **SD-WAN Quality** columns, which are disabled by default. The former indicates the matched SD-WAN rule for a session, and the latter the member the session was steered to and the reason.

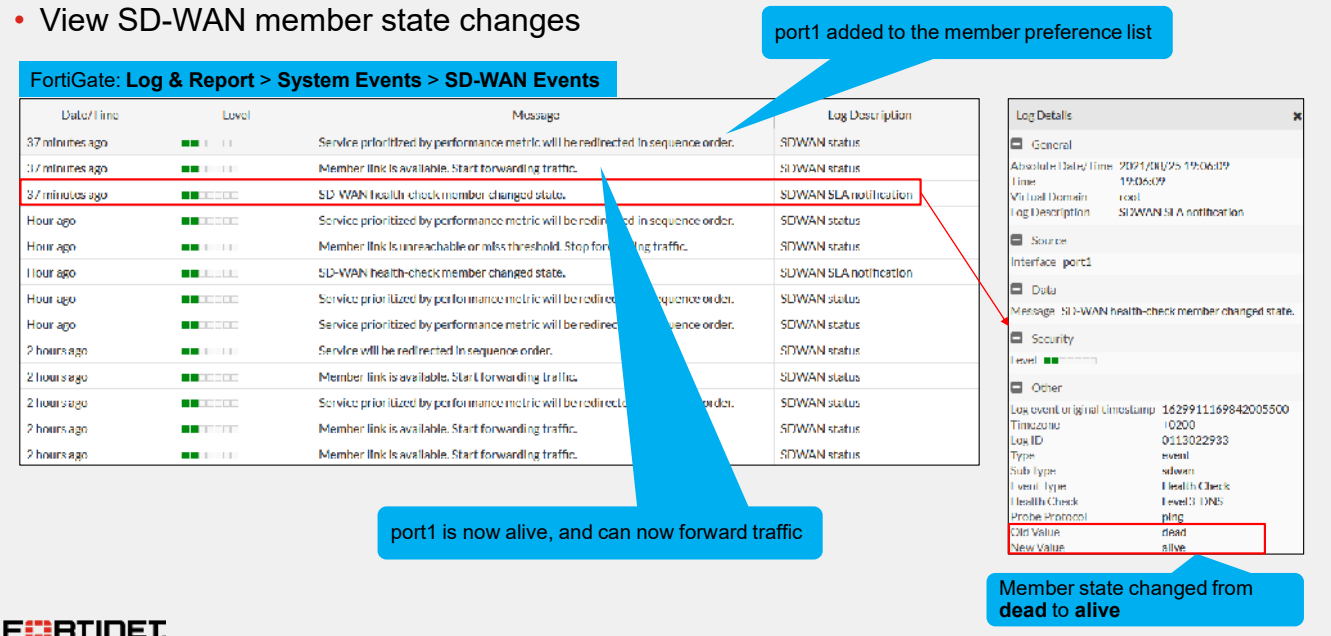
The table on this slide shows multiple sessions. The first session in the table was identified as a **Salesforce** application, matched the **Critical-DIA** rule, and was sent to port1. The reason that port1 was selected was because it had the lowest latency.

The second session in the table, which was identified as a **Facebook** application, matched the **Non-Critical-DIA** rule, and was sent to port2. The **Non-Critical-DIA** rule instructs FortiGate to steer matching traffic to port2 only, provided the port is alive. This behavior matches the reason described in the **SD-WAN Quality** column for that session.

DO NOT REPRINT
© FORTINET

SD-WAN Events

- View SD-WAN member state changes



FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

The **SD-WAN Events** sub-section on the **Events** page displays logs that report the state changes of the SD-WAN members.

In most cases, you want to click a log to fully understand the event. For example, the third log in the table indicates that the state of **port1** changed from **dead** to **alive**. Although the details of the second and first logs are not shown, the logs report that port1 is ready to forward traffic, and that the member preference in the rule that uses port1 (**Critical-DIA**) was updated to include port1.

DO NOT REPRINT
© FORTINET

Review

- ✓ Describe SD-WAN and Fortinet secure SD-WAN
- ✓ Describe DIA, site-to-site, RIA, and cloud on-ramp use cases
- ✓ Introduce SD-WAN main components: members, zones, and rules
- ✓ Describe basic SD-WAN routing and firewall policies
- ✓ Configure a basic SD-WAN DIA setup using the FortiGate GUI
- ✓ Monitor SD-WAN traffic distribution, member health, and events using the FortiGate GUI

FORTINET Training Institute

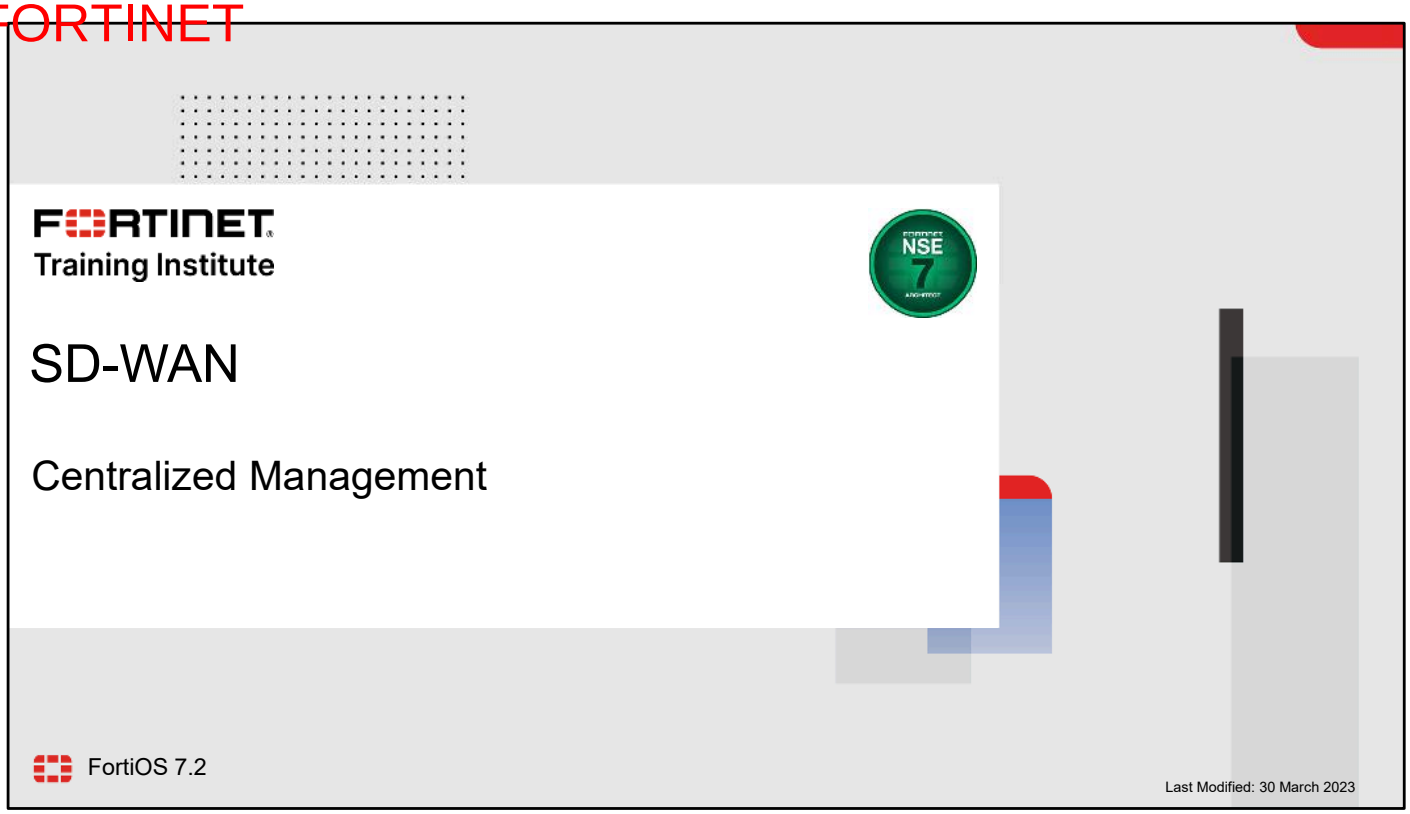
© Fortinet Inc. All Rights Reserved.

33

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

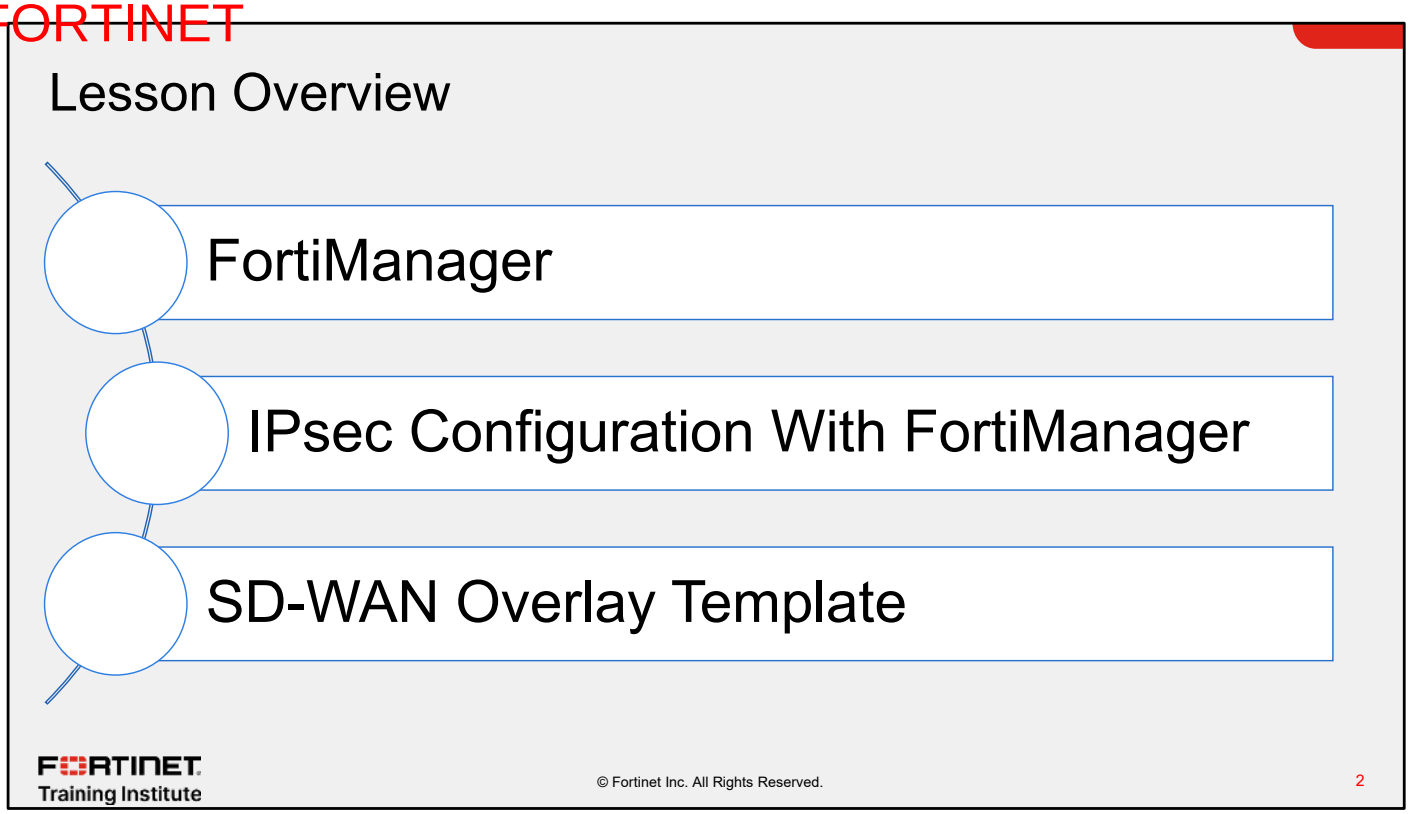
By mastering the objectives covered in this lesson, you learned the most common use cases for SD-WAN, its main components, and how to configure and monitor a basic SD-WAN DIA setup.

DO NOT REPRINT
© FORTINET



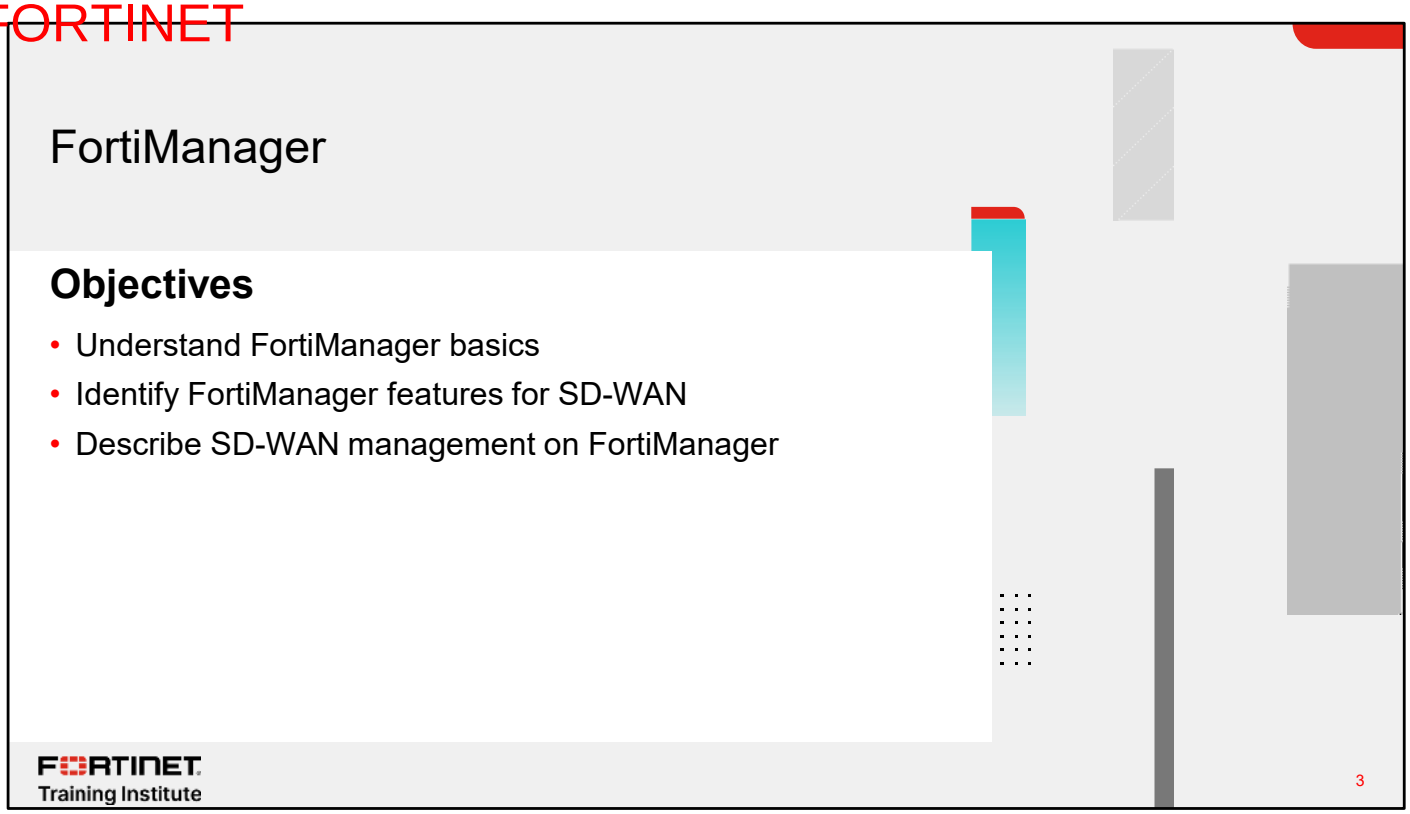
In this lesson, you will learn how FortiManager can help to deploy and maintain SD-WAN networks.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



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

By understanding how to deploy SD-WAN using FortiManager, you should be able to leverage the FortiManager centralized features to reduce operation costs when deploying and maintaining SD-WAN across a large network.

DO NOT REPRINT
© FORTINET

What Is FortiManager?

- Provides centralized management for SD-WAN
  - Single-pane-of-glass management
- Benefits for SD-WAN:
  - Provisioning of SD-WAN templates, CLI templates, and firewall policies across multiple devices
  - Per-device SD-WAN configuration support
  - SD-WAN status monitoring
  - Central repository for configuration revision control and security audits
  - Deploy and monitor complex IPsec VPN topologies (IPsec overlays)
  - Zero-touch provisioning for new SD-WAN sites
  - Support for scripts and JSON APIs to automate device provisioning and perform policy changes

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

FortiManager is a key component for deploying SD-WAN across a large network. Centralized (single-pane-of-glass) management through FortiManager can help you to more easily manage SD-WAN deployment across many devices and reduce the cost of operation.

How can FortiManager help you with your SD-WAN deployment?

- Provision SD-WAN templates, CLI templates, and firewall policies across multiple devices with the same configuration requirements.
- Enable you to configure SD-WAN on a per-device basis.
- Monitor SD-WAN status.
- Act as a central repository for configuration revision control and security audits.
- Deploy and monitor complex IPsec VPNs topologies (IPsec overlays).
- Perform zero-touch provisioning for new SD-WAN sites.
- Use scripts and JSON APIs to automate device provisioning and perform policy changes for SD-WAN.

DO NOT REPRINT
© FORTINET

Management Layers

- Global ADOM layer
  - Global objects
  - All header and footer policies
- ADOM layer
  - Common object database, devices, device groups, policy packages
  - Groups of devices sharing common objects
  - Default ADOM: root
- Device Manager layer
  - Name and type of managed devices, their IP addresses, revision history, and real-time status

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

To organize and efficiently manage a large-scale network, FortiManager has multiple management layers.

The global ADOM layer has two key pieces: the global object database and header and footer policy packages. Header and footer policy packages envelop the policies for each ADOM. Policy packages are often used in a carrier environment, where the carrier allows customer traffic to pass through their network but doesn't allow the customer to have access to their network infrastructure.

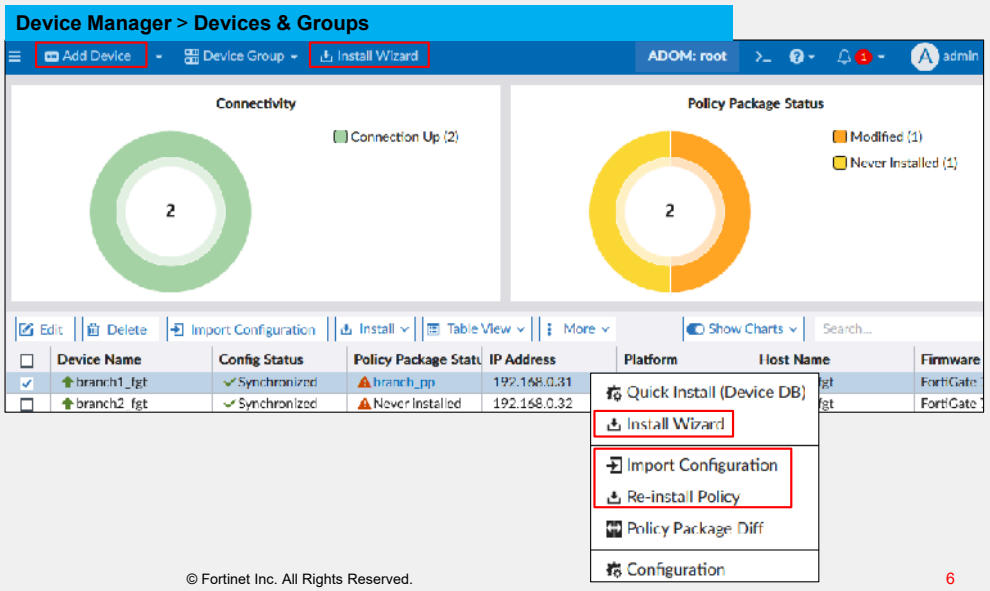
The ADOM layer is where policy packages are created, managed, and installed on managed devices or device groups. You can create multiple policy packages here. The ADOM layer includes one common object database for each ADOM. The common object database contains information such as addresses, services, and security profiles. ADOMs enable you to create groupings of devices for administrators to monitor and manage. The purpose of ADOMs is to divide administration of devices by ADOM and to control (restrict) administrator access.

The Device Manager layer records information on devices that are centrally managed by the FortiManager device, such as the name of the device, type of device, model, IP address, current firmware installed, revision history, and real-time status.

DO NOT REPRINT
© FORTINET

Wizards

- Assist with various tasks
- Main wizards:
  - Add Device
  - Install Wizard
  - Import Configuration
  - Re-install Policy



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

The **Device Manager** pane provides device and installation wizards to aid you in various administrative and maintenance tasks. Using these wizards can decrease the amount of time it takes to do many common tasks.

There are four main wizards on the **Device Manager** pane:

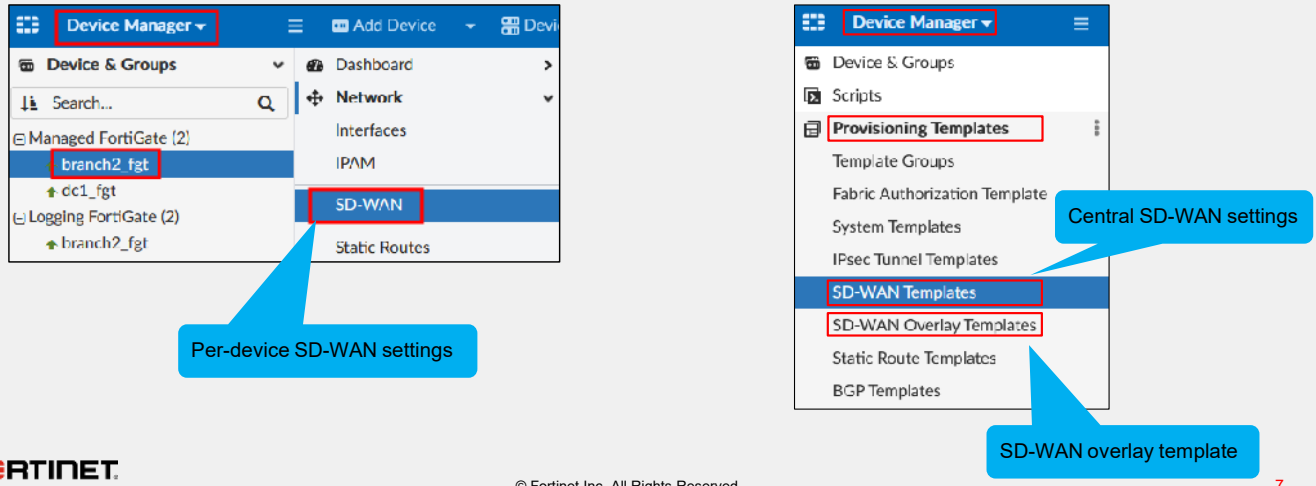
- Use **Add Device** to add devices to central management and import their configurations.
- Use **Install Wizard** to install configuration changes from the **Device Manager** pane or **Policies & Objects** pane to the managed devices. It allows you to preview the changes and, if the administrator doesn't agree with the changes, cancel and modify them.
- Use **Import Configuration** to import interface mappings, policy databases, and objects associated with the managed devices into a policy package one the **Policy & Object** page. It runs with the **Add Device** wizard, by default, and you can run it at any time from the managed device list.
- Use **Re-install Policy** to perform a quick installation of the policy package. It provides the ability to preview the changes that will be installed on the managed device.

You can open the **Import Configuration** and **Re-install Policy** wizards by right-clicking your managed device in the **Device Manager**.

DO NOT REPRINT
© FORTINET

SD-WAN Management

- Per device
  - Apply settings on individual devices
  - Devices have different SD-WAN requirements
  - SD-WAN objects stored in device settings database
- Central
  - Apply templates on multiple devices
  - Devices have the same SD-WAN requirements
  - SD-WAN objects stored in the ADOM database



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

FortiManager offers two approaches for configuring SD-WAN: per-device management and central management.

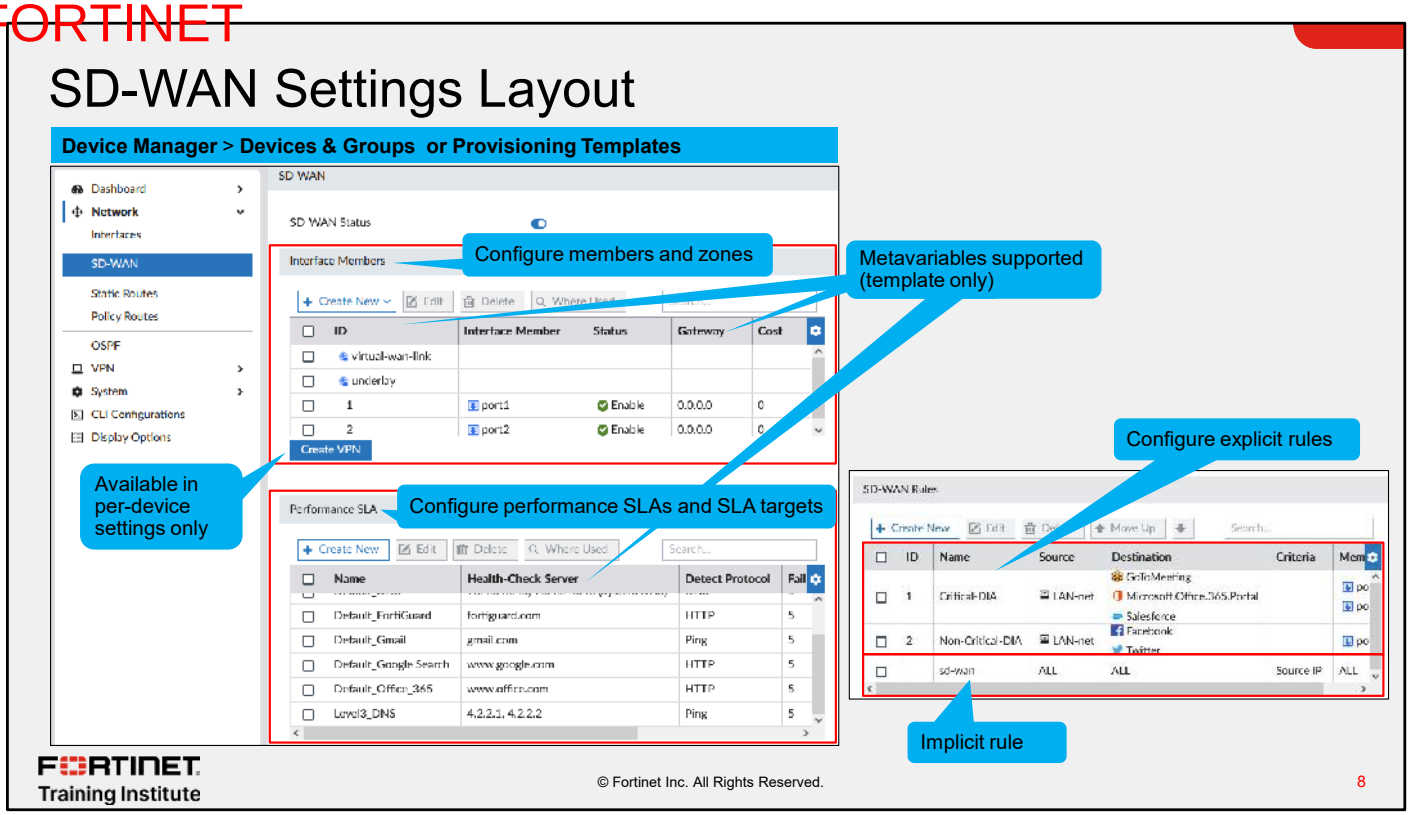
In per-device management, you configure SD-WAN settings for individual devices. You make configuration changes on the SD-WAN page of the managed device, and then install them. The per-device management approach is useful when your devices have different SD-WAN configuration requirements, and therefore, you must maintain separate settings for each device.

In central management, you configure SD-WAN templates and assign them to one or more FortiGate devices. For each SD-WAN template, you define SD-WAN members, zones, performance SLAs, rules, and so on. This approach is convenient for deploying multiple devices that use similar configurations because it reduces the administrative overhead. That is, instead of applying a change on each managed device, you apply it on the shared template. Then, when you install the template changes, FortiManager pushes the change to all target devices of the template.

The screenshots on this slide show the FortiManager pages where you apply per-device and central SD-WAN settings. While both pages are in the **Device Manager** pane, you must click the managed device to access the per-device SD-WAN settings. You can find the SD-WAN central settings under the **Provisioning Templates** section in the **Device Manager** pane.

The **SD-WAN Overlay Templates** combine, in a guided approach, configuration of SD-WAN, IPsec overlay, and ADVPN templates. You will learn more about the SD-WAN overlay templates in this lesson.

DO NOT REPRINT
© FORTINET



Whether you configure SD-WAN using per-device management or central management, the FortiManager pages for both approaches look almost the same. One difference is that the per-device settings page shows the **Create VPN** button, which enables you to quickly create an IPsec tunnel that can later be added as an SD-WAN member.

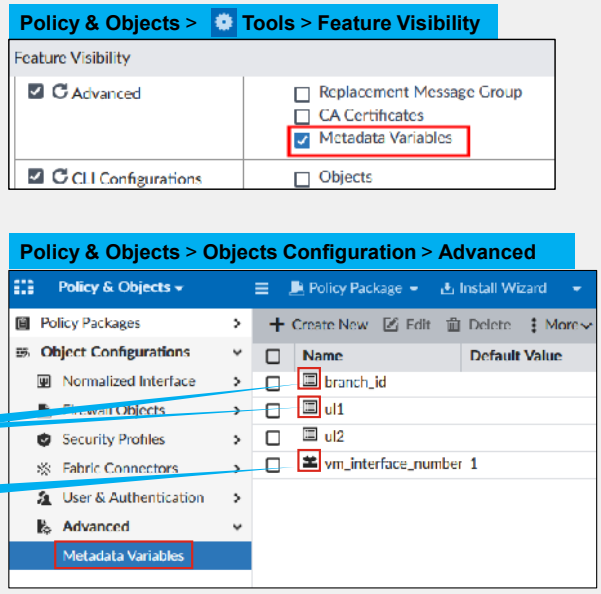
Another difference is that when you configure an SD-WAN member using an SD-WAN template (central management), you can use metavariables in the **Interface Member**, **Gateway** and **Health-Check Server** settings. Metavariables enable you to define variables that can be assigned with different values per device. Metavariable are introduced with FortiManager 7.2.0 and replace metafields. You will learn more about metavariables in this lesson.

The way you configure SD-WAN settings using FortiManager is very similar to how you configure them on the FortiGate GUI. This slide shows an example of a configuration made using per-device management. Like FortiGate, there are three sections available on the GUI: **Interface Members**, **Performance SLA**, and **SD-WAN Rules**, and they serve the same purpose as on the FortiGate GUI.

DO NOT REPRINT
© FORTINET

Metadata Variables

- ADOM level
- Available for CLI scripts, templates, and model devices
- Provide flexibility
  - Allow one variable name and per-device value
- Menu hidden by default



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

Metadata variables are ADOM-level parameters—also available for Global ADOM for per-ADOM mapping—introduced with FortiManager version 7.2.0. You can use metadata variables in CLI scripts, templates, or model devices. They provide required flexibility each time a parameter has different values on each devices.

You can access the **Metadata Variables** menu under **Policy & Objects** > **Advanced** to review and edit metadata variables.

The metadata variable name can contain only letters, numbers, and underscores.

On upgrade from version 7.0 or earlier, FortiManager creates metadata variables at the ADOM level for metafields used in the ADOM.

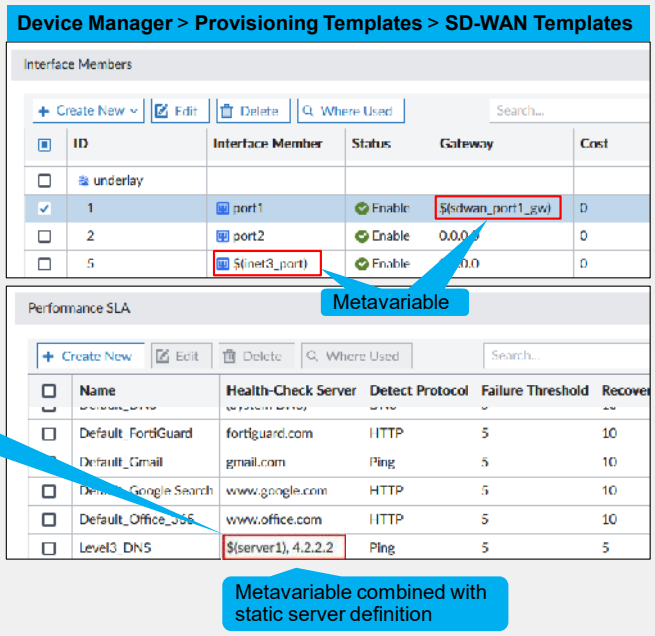
If a metafield contains characters unsupported for metadata variable names, the name is modified, and every unsupported character is replaced with an underscore "_".

System level metafields are kept on upgrade for reference. They remain visible under **System Settings** > **Advanced** > **Meta Fields**.

DO NOT REPRINT
© FORTINET

Metadata Variables for SD-WAN Members

- User-defined variables:
  - Apply different settings per device
  - Useful for SD-WAN member configuration
    - Central management mode only
      - Normalized interfaces are not supported
    - Available for
      - Interface Member
      - Gateway
      - Health-Check server



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

The FortiManager metadata variables are user-defined variables that enable you to assign different values to a setting for a given device. They are particularly useful when you configure SD-WAN members using SD-WAN templates.

You can use metadata for the **Interface Member** setting. This allows you to use different interfaces on different devices without having to create separate templates.

Another use case for metadata is the gateway setting of SD-WAN members. Even if you use the same interface name for the devices assigned to a template, their gateway is likely to differ. For this reason, you can define a metadata variable for the gateway setting that indicates the gateway IP address to push for the member on each device.

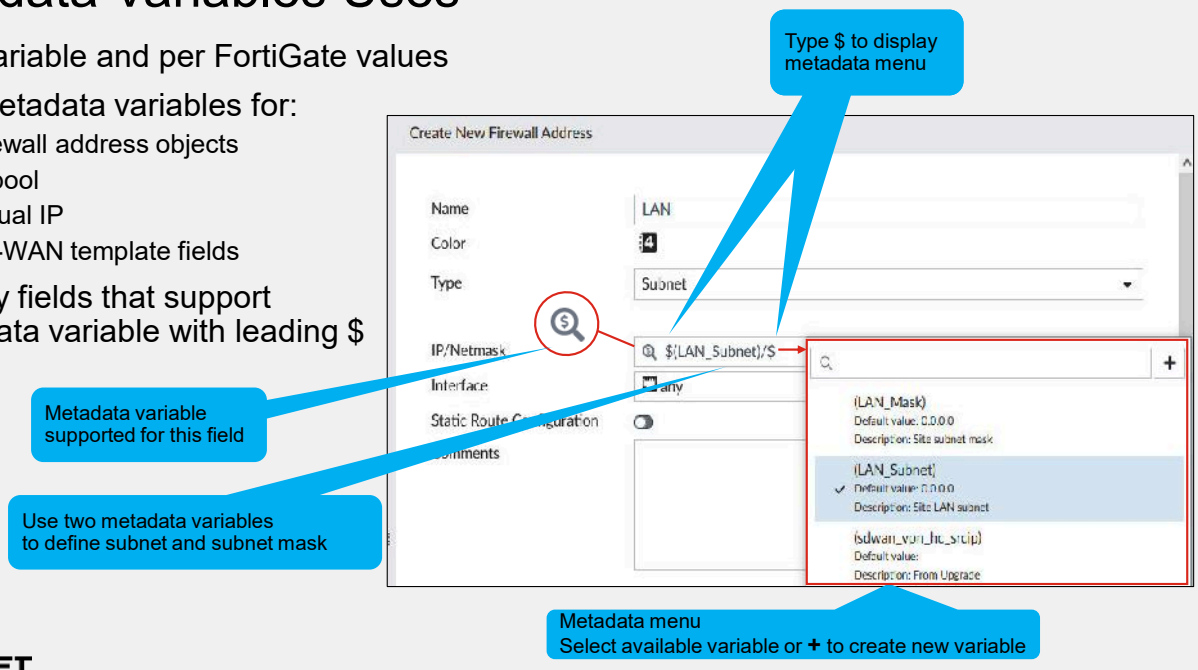
The example on this slide shows the use of two metadata variables for SD-WAN member configuration. The **sdwan_port1_gw** metadata variable is used to define a different gateway for **port1** on each managed device. That is, the gateway for **port1** on **branch1_fgt** and **branch2_fgt** will be `192.2.0.2` and `203.0.113.2`, respectively. For member ID **5** in the **underlay** zone, the **inet3_port** metadata variable is used to indicate the name of the interface to use as the member on each managed device. That is, **branch1_fgt** will use **port3** as the member, and **branch2_fgt** will use **port8**.
In **Performance SLA** section, health check servers for the Level3_DNS entry are defined with one static server IP, the same IP for every device, and one metadata variable to adjust the health check server IP according to device location.

To reference a metadata variable in the SD-WAN member configuration, type `$` at the beginning of the string so FortiManager shows a list of available metadata variables. On this pop-up menu, a "**+**" sign allows the user to create a new metadata variable, if required.

DO NOT REPRINT
© FORTINET

Metadata Variables Uses

- One variable and per FortiGate values
- Use metadata variables for:
  - Firewall address objects
  - IP pool
  - Virtual IP
  - SD-WAN template fields
- Identify fields that support metadata variable with leading $



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

The magnifying glass with a $ sign indicates fields where you can use a metadata variable.

.

In the example shown on this slide, the user can enter the IP address and netmask as usual, or enter $ to display the metadata variable menu and select one metadata variable for the subnet and one metadata variable for the subnet mask.

Note that it's mandatory to specify the subnet mask separately.

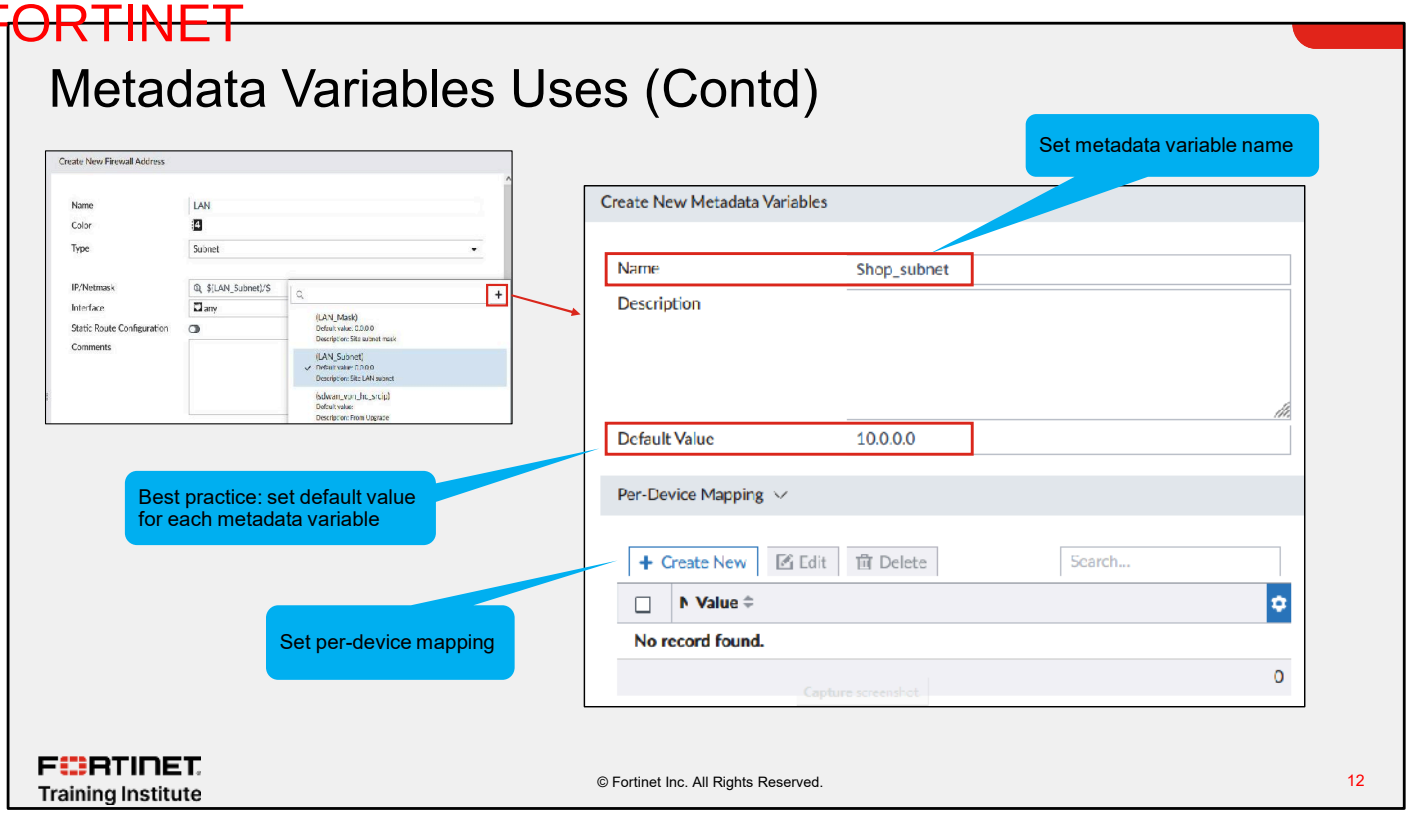
Valid options are:

- `$(LAN_Subnet)/$(LAN_Mask)`
- `$(LAN_Subnet)/255.255.255.0`

After you enter the $ sign in a field, a pop-up menu allows you to select an already defined variable. You can create a new variable with the **+** sign, or edit an existing metavariable with the pen sign displayed when you hover over a variable.

To review and edit all metadata variables at the ADOM level, go to **Policy & Objects** > **Object Configurations** > **Advanced** > **Metadata Variables**.

DO NOT REPRINT
© FORTINET



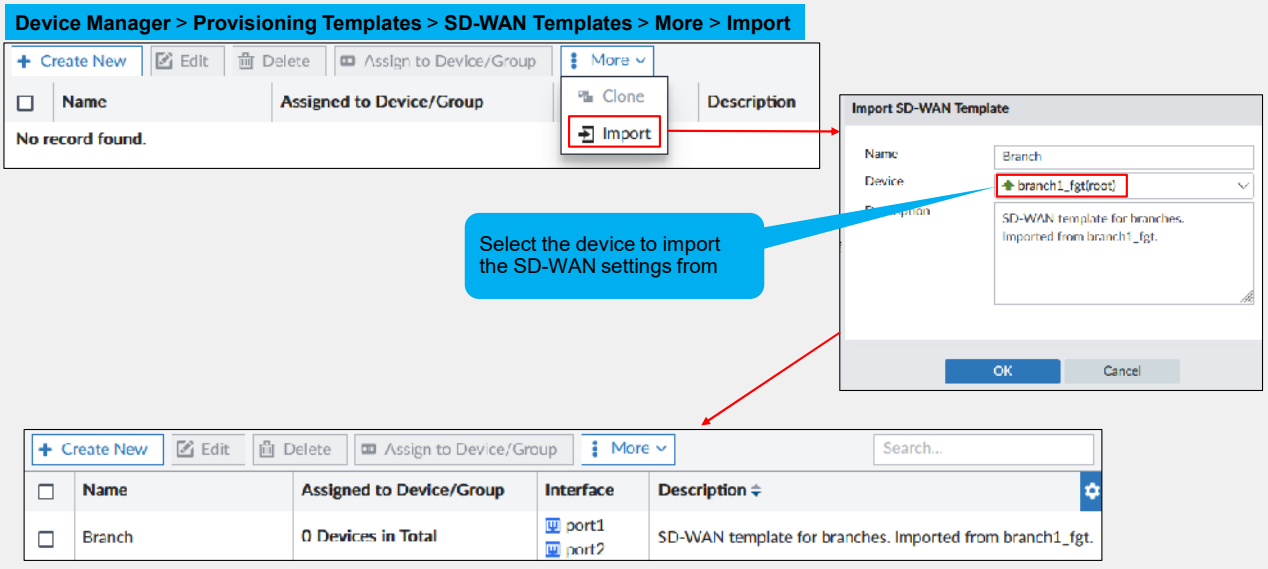
If you click the **+** sign to create a new metadata variable, FortiManager shows a window where you can create a new variable and its per-device mapping. You can then define the variable name and its values.

Note that for some fields, you cannot use a metadata variable without a default value. Therefore, it's a good practice to set a default value for each metadata variable you create. FortiManager uses the default value when installing the configuration on each FortiGate for which you didn't specify a per-device mapping. One example of a field that requires a default value is **IP/Netmask**, for a firewall address object.

DO NOT REPRINT
© FORTINET

Central Management—Importing Template Settings

- Importing SD-WAN settings from a managed device into a new template:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

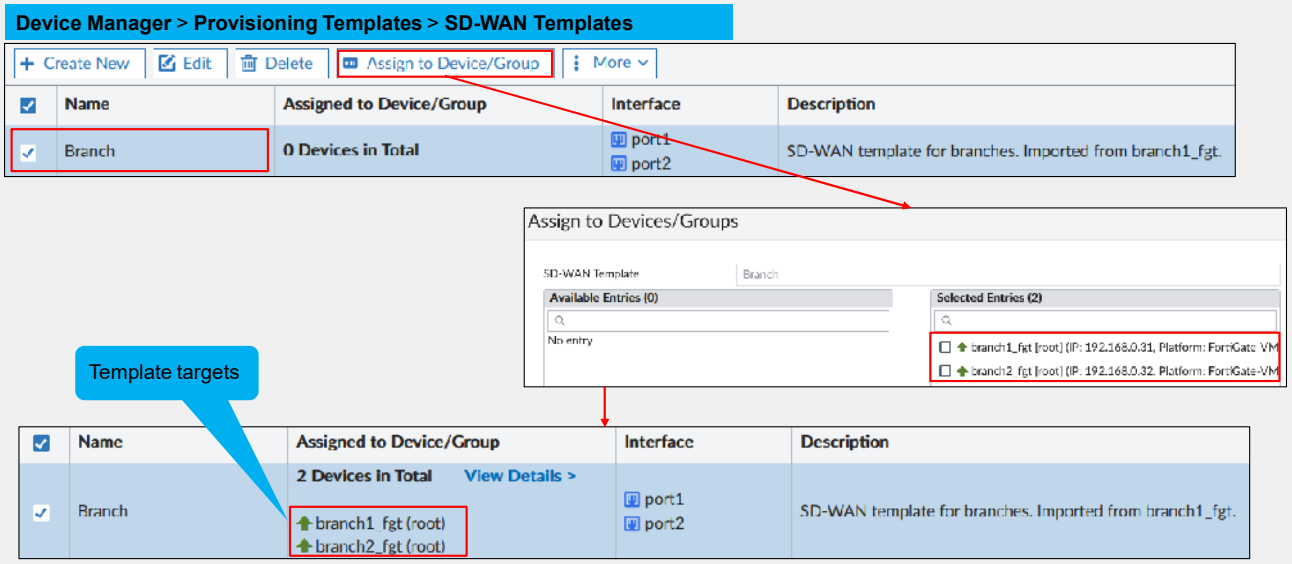
13

For SD-WAN central management, you can import the SD-WAN settings of a device into an SD-WAN template in FortiManager. You can then use the template to deploy SD-WAN on other FortiGate devices that require the same configuration.

DO NOT REPRINT
© FORTINET

Central Management—Assigning a Template

- Assigning an SD-WAN template to one or more devices:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

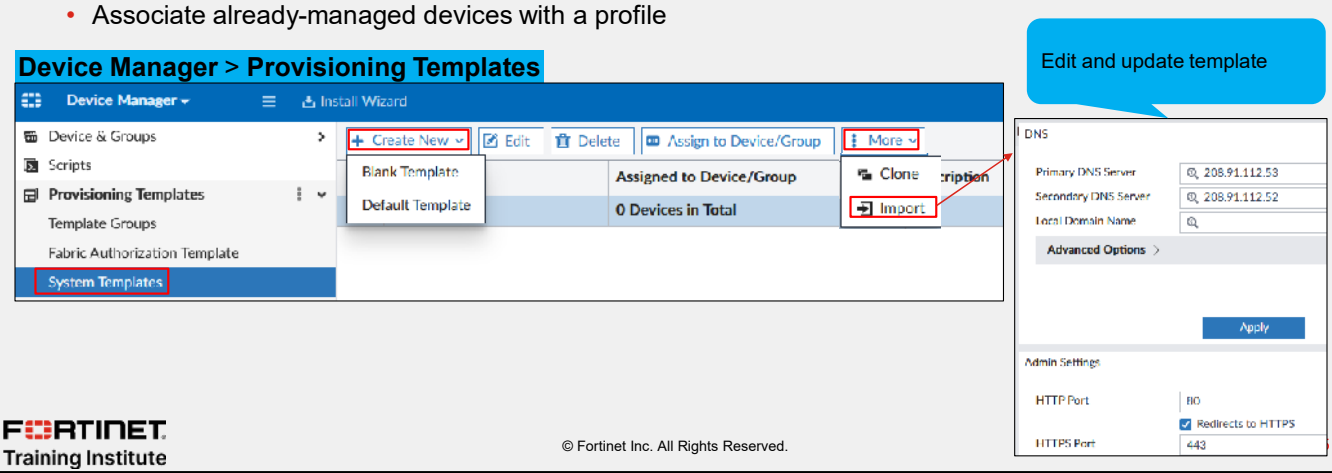
14

After you configure a new template or you import its settings from a managed device, you can then assign the template to one or more devices or device groups.

DO NOT REPRINT
© FORTINET

Provisioning Templates—System Templates

- Subset of model device configuration—system-level settings
  - Configure or modify common settings
- You can:
  - Create new templates
  - Inherit the system settings and CLI settings of a managed device using the import feature
  - Associate already-managed devices with a profile



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

In addition to configuring an SD-WAN template, you may also need to configure a system template and one or more CLI templates. A system template enables you to create and manage common system-level settings for the managed device. The **System Template** page contains one generic profile named **default**, which contains widgets for settings such as **DNS**, **Alert Email**, **Admin Settings**, **Log Settings**, and others.

You can create a new device profile and configure the settings in the widgets in that profile. You can use the **More** icon and **Import** to import the settings from a specific managed device, which inherits the system-level settings of that managed device.

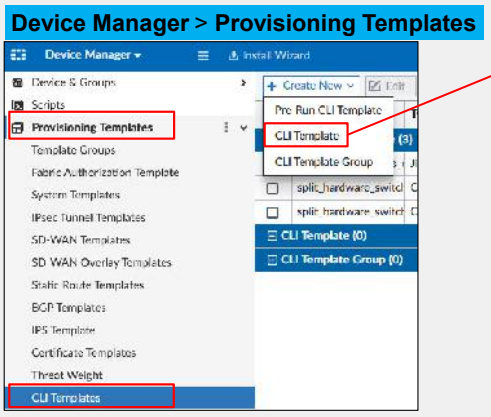
You can use the **Assign to Device/Group** tab to associate devices with a profile, or to view the list of devices already assigned to a profile.

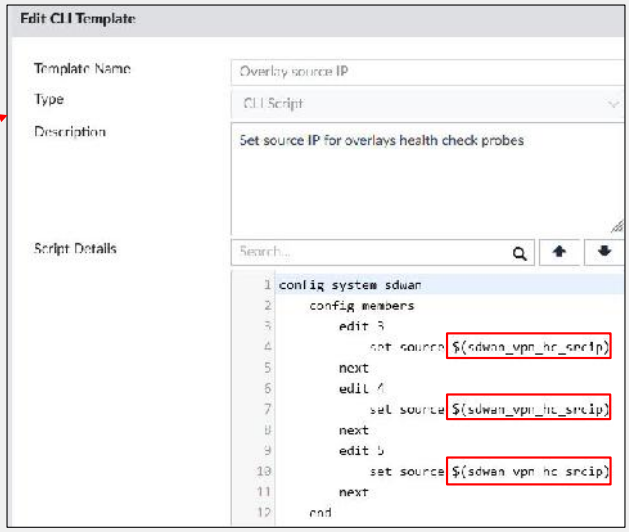
You can apply these configured profiles to multiple devices within the same ADOM, which facilitates identical device-level settings across many devices.

DO NOT REPRINT
© FORTINET

Provisioning Templates—CLI Templates

- CLI scripts or groups of CLI scripts
  - Content of script is enforced when pushing the configuration to managed device
- Useful for:
  - Configuring advanced settings
  - Referencing metadata variables





FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

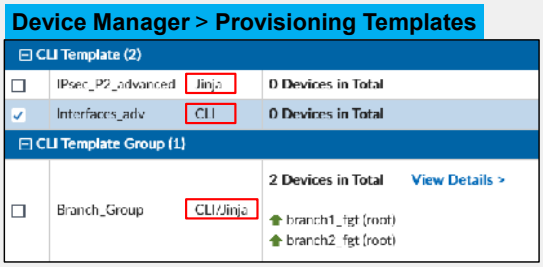
CLI templates enable you to create CLI scripts or a group of CLI scripts that you can assign to managed devices. FortiManager then enforces the content of the script when pushing the configuration to managed devices.

CLI templates are useful for pushing advanced CLI settings, or settings that reference metadata variables. For example, you can use CLI templates to push the SD-WAN settings shown on this slide, which instructs FortiManager to configure the source address for health check probes used by the SD-WAN members. Instead of indicating the source address directly, you can reference a metadata variable (**sdwan_vpn_hc_srcip** in the example). Because the metadata variable is defined with per-device mapping values, then FortiManager can push a different source address based on the target device.

DO NOT REPRINT
© FORTINET

Provisioning Templates—CLI Templates (Contd)

- CLI template type
  - CLI Script
    - FortiGate CLI commands
    - Metadata variables referenced with $
  - Jinja
    - Jinja2 is a full-featured template engine
    - Advanced scripting options
    - Python-like syntax
    - Metadata variables referenced with {{
- Can mix both types in template group
  - Scripts applied in top-down order



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

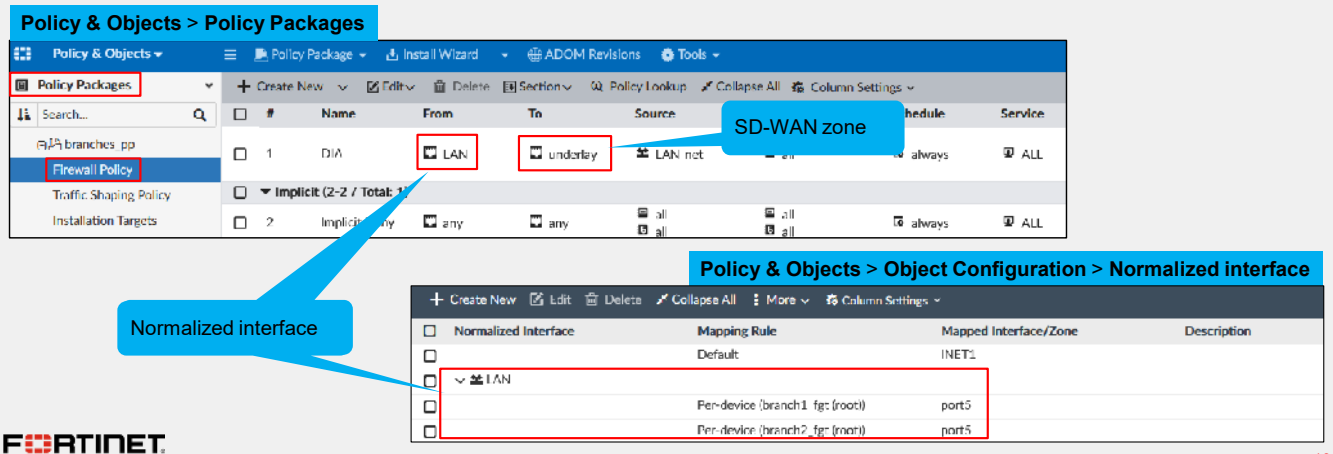
On FortiManager you can create two types of scripts: CLI scripts and Jinja scripts. For CLI scripts, you will use the FortiGate CLI commands and the $ sign to reference metadata variables. This type of script is very easy to use but does not offer advanced programming functions. With Jinja scripts, you can use Python-like syntax to configure advanced scenarios. Note that for Jinja scripts, you must reference metadata variables with a double brace symbol— {—.

You can use a template group to assign multiple CLI scripts to managed devices. The template groups can contain a mix of CLI and Jinja scripts and are applied in a top-down order.

DO NOT REPRINT
© FORTINET

Firewall Policies and Normalized Interfaces

- Configure firewall policy for SD-WAN
  - Create a policy package
    - Create firewall policy for SD-WAN zone
    - Select installation targets
  - Check mapping rules of normalized interfaces



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

18

When you configure firewall policies for SD-WAN, you must first create a policy package. The policy package then contains one or more firewall policies that are assigned to one or more managed devices. Firewall policies for SD-WAN traffic must reference SD-WAN zones and not individual members. The other interface configured in the firewall policy is usually a normalized interface, for which you must configure correct mapping rules.

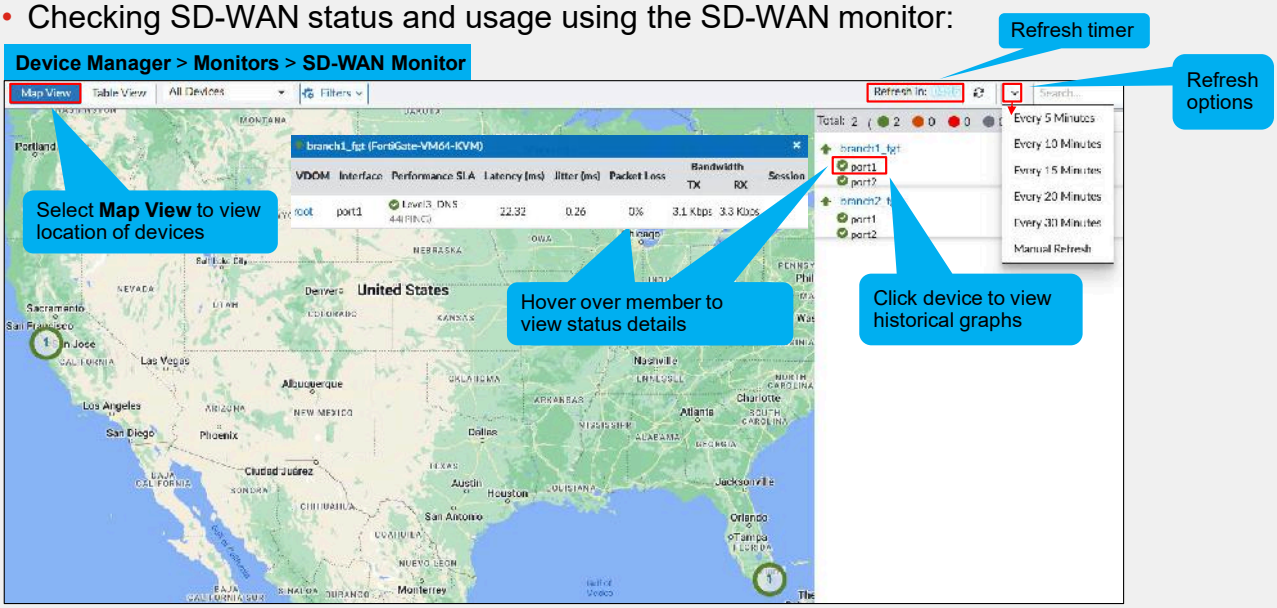
Normalized interfaces enable you to reference different interfaces on a per-device or per-platform basis. The goal is to be able to share objects, such as firewall policies, across multiple devices with different interface configurations. When FortiManager installs objects that reference a normalized interface, it reads the configured mapping rules, and then assigns the mapped interface to the pushed configuration of each target device.

In the example shown on this slide, the **branches-pp** policy package contains one firewall policy named **DIA**. The **LAN** normalized interface is configured as the incoming interface on the policy. **LAN** is mapped to **port5** on both **branch1_fgt** and **branch2_fgt** devices, but it could be mapped to different interfaces if needed.

DO NOT REPRINT
© FORTINET

Checking SD-WAN Status

- Checking SD-WAN status and usage using the SD-WAN monitor:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

After you install the SD-WAN settings and relevant configuration on FortiGate, you can use the **SD-WAN Monitor** page on FortiManager to view the status of the FortiGate devices and their SD-WAN members. Note that, by default, you must manually refresh the page to poll the latest status from the device. Alternatively, you can select an automatic refresh interval.

The **Map View** option shows the location of devices on a map. The location is based on the location settings configured for the device in FortiManager.

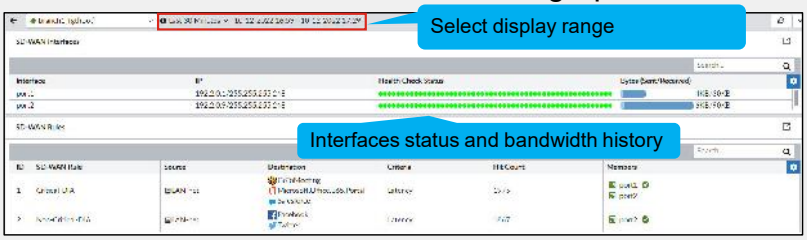
Hover over a member to view its health and utilization details. Note that the member utilization percentage is calculated based on the values configured for the `estimated-upstream-bandwidth` and `estimated-downstream-bandwidth` interface settings. You will learn more about these settings in another lesson.

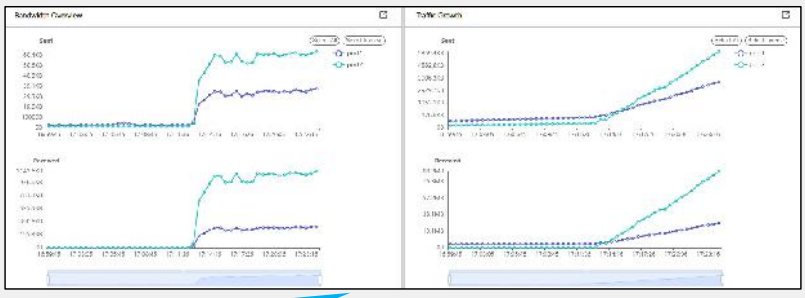
Finally, you can click a device to view historical graphs reporting on the member utilization and health. The next slide shows an example of those graphs.

DO NOT REPRINT
© FORTINET

Checking SD-WAN Status (Contd)

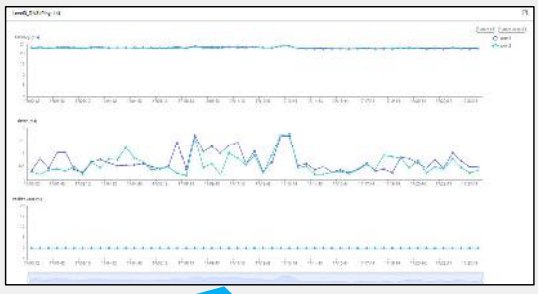
- View member utilization and health graphs:





- Default display for past 10 min
- For over extended time period enable SD-WAN history monitor with CLI commands:

```
config system admin setting
    set sdwan-monitor-history enable
end
```



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

When you click a device in the SD-WAN monitor page, FortiManager displays historical graphs reporting on the utilization and health of the SD-WAN members on that device. This slide shows an example of the utilization and health graphs displayed by FortiManager.

The user can select the time range to display.

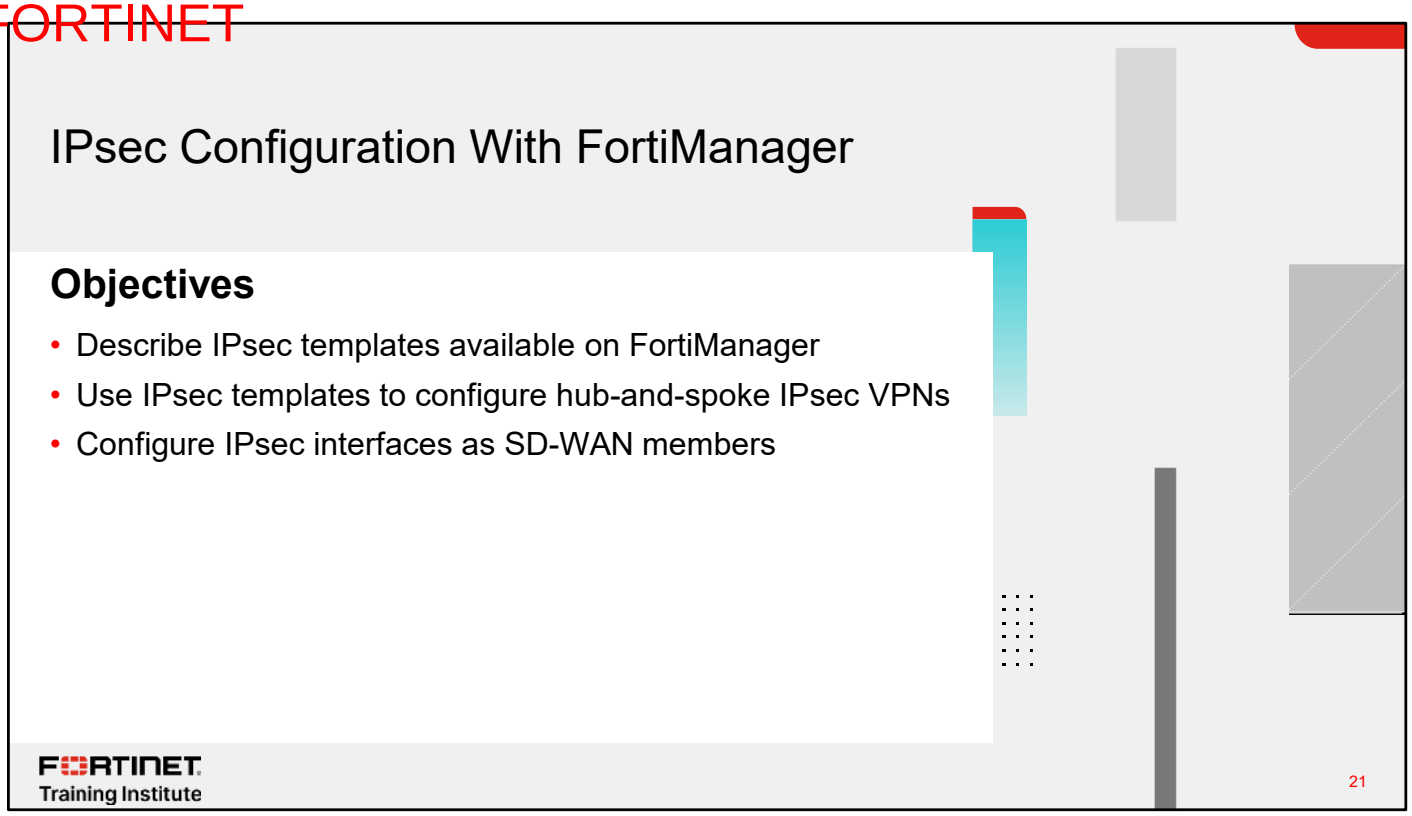
Note that, by default, FortiManager displays the data for the past 10 minutes only. To collect and display data over an extended time period you need to enable data storage on the FortiManager disk with the following CLI commands:

```
config system admin setting
      set sdwan-monitor-history enable
end
```

Once the feature is activated, FortiManager stores the data on its disk for each managed SD-WAN device for up to 180 days. To avoid excessive disk usage, you can reduce the data conservation to only a few days with the following CLI commands:

```
config system admin setting
      set rtm-max-monitor-by-days <nb of days>
end
```

DO NOT REPRINT
© FORTINET



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

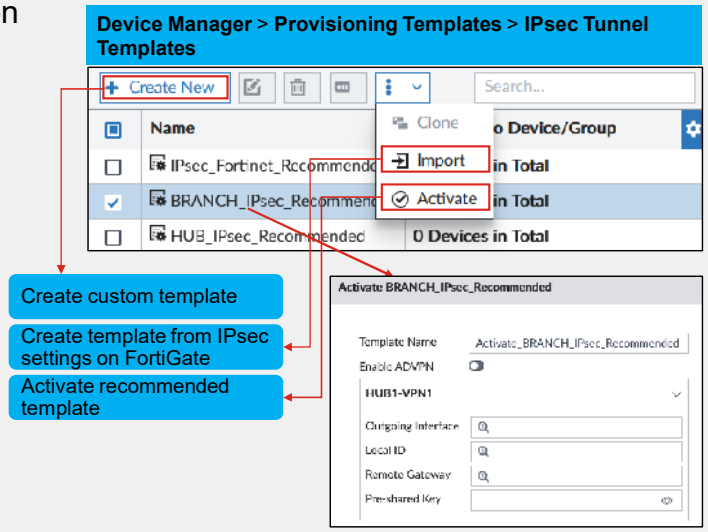
By understanding the IPsec configuration templates available on FortiManager, you will be able to easily configure IPsec VPN tunnels on FortiGate. You will then see how to use the configured IPsec tunnels as overlays in SD-WAN.

DO NOT REPRINT
© FORTINET

FortiManager IPsec Tunnel Templates

- Easy provisioning of IPsec tunnels
- Guidance for IPsec tunnel configuration
  - Use recommended parameters
  - Consistent settings between phase1 and phase2
- Applied to multiple devices
  - Simplified deployment
  - Same parameters at both end of tunnels
- Available templates
  - IPsec Fortinet Recommended
  - BRANCH IPsec Recommended
  - HUB IPsec Recommended
- One template per device
  - One or multiple tunnels per template

- Template menu



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

On FortiManager, you can use the IPsec template to easily configure consistent IPsec settings on multiple devices.

There are various ways to prepare a template:

- Create a new template and manually define all required IPsec parameters (**Create New**).
- Import settings from a managed FortiGate with an already defined IPsec tunnel (**Import**).
- Use a Fortinet Recommended template (**Activate**).

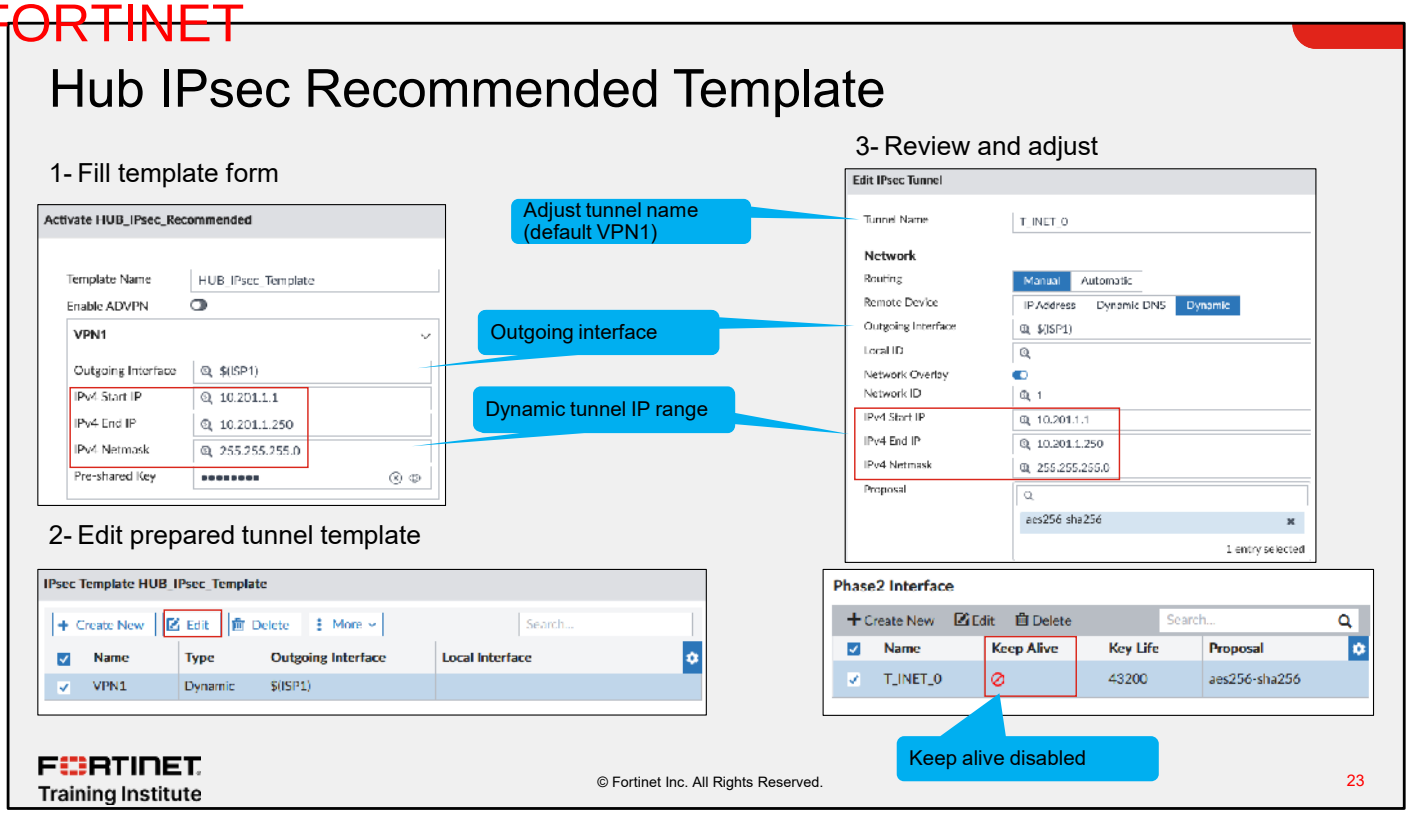
Recommended templates will allow you to prepare a template for IPsec tunnels using Fortinet recommended settings for phase1 and phase2 parameters.

- The IPsec_Fortinet_Recommended template defines a template for a static point-to-point tunnel
- The BRANCH_IPsec_Recommended template defines a template for a static tunnel (with a known remote IP address)
- The HUB_IPsec_Recommended template defines a template for a dynamic tunnel (an IPsec hub for dial-up tunnels)

This slide shows an example of a limited number of parameters that you must provide to prepare a branch tunnel configuration with the Branch_IPsec recommended template. Then, using those parameters, FortiManager prepares a template with the IPsec configuration ready to install on the device. You can still review the template and adjust it if necessary. As you already discovered for SD-WAN or CLI templates, metadata are useful to customize the template with values specific to each device.

Note that if you are already familiar with FortiManager and have used the VPN manager for your deployments, you can keep using it. However, for new deployments, you should use the IPsec recommended templates.

DO NOT REPRINT
© FORTINET



The Hub IPsec recommended template prepares a template for a dynamic IPsec VPN tunnel with an outgoing interface, pre-shared key, and IP address range for the remote tunnel specified by the administrator.

The tunnel name, network ID, and encryption proposal are automatically set with the values VPN1, ID1, and aes256-sha256.

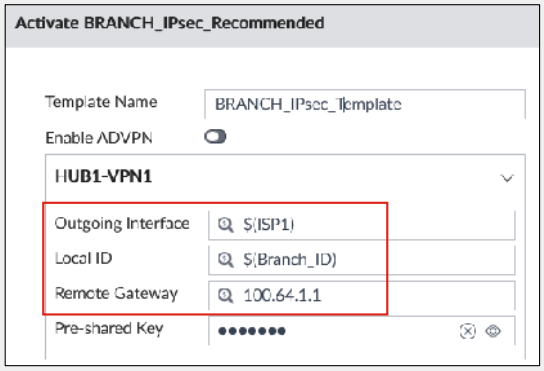
Of course, if required, you can edit and adjust the template.

Note that on template creation, **Keep Alive** is disabled. You can edit **Keep Alive** to enable it.

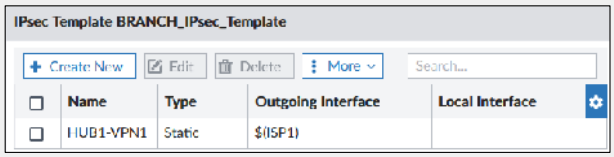
DO NOT REPRINT
© FORTINET

Branch IPsec Recommended Template

1- Fill template form

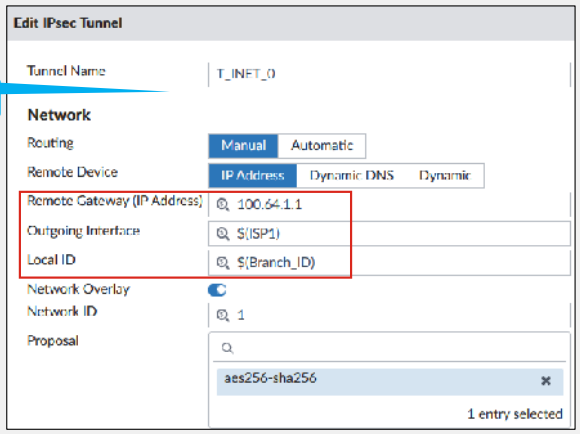


2- Edit prepared tunnel template



Adjust tunnel name (default HUB1-VPN1)

3- Review and adjust



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

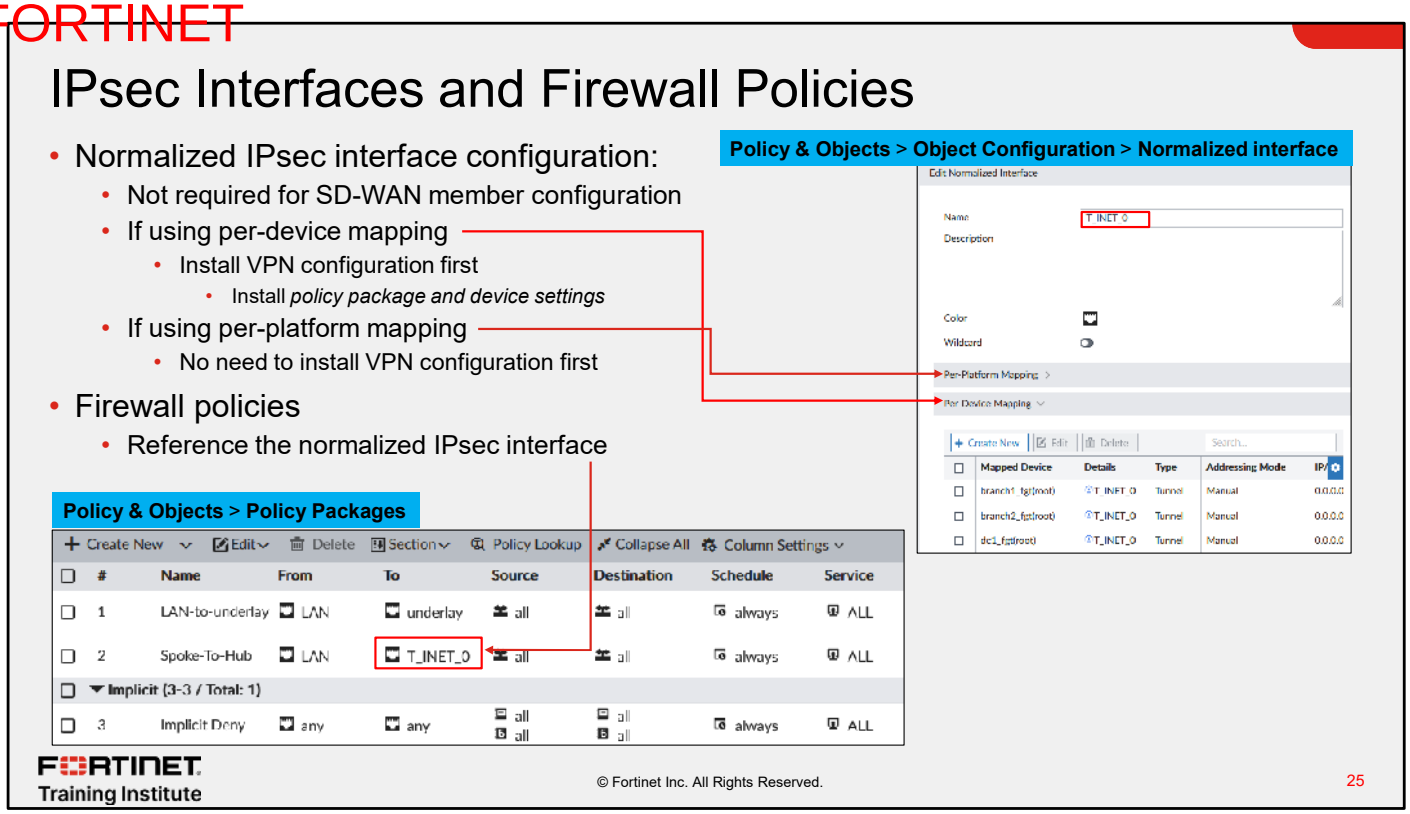
24

The Branch IPsec recommended template prepares a template for the static IPsec VPN tunnel. You must specify the IP address for the remote end of the tunnel. The template uses the outgoing interface, pre-shared key, and local branch ID specified by the administrator. The local ID must be unique for each branch site. It's therefore convenient to define it as metadata variable.

The tunnel name, network ID, and encryption proposal are automatically set to HUB1-VPN1, ID1 and, aes256-sha256. If required, you can edit and adjust those values and any other parameter from the template.

Note that on template creation, such as for hubs, **Keep Alive** is disabled. You can edit the template to enable **Keep Alive**.

DO NOT REPRINT
© FORTINET



Normalized interfaces enable you to reference different interfaces on a per-device or per-platform basis, thus simplifying the configuration and deployment process for multiple devices. Usually, FortiManager requires you to normalize interfaces. However, it's not required if you plan to configure an IPsec interface as an SD-WAN member. This is because SD-WAN members don't use normalized interfaces. Another reason is that firewall policies for SD-WAN must reference SD-WAN zones, and not individual members.

If you *don't* plan to use an IPsec interface as an SD-WAN member, then you must normalize the interface so you can reference it on firewall policies, which are required for the tunnel to come up. Keep in mind the following when normalizing the interface for IPsec:

- If you plan to use per-platform mapping, you don't need to install the VPN configuration first. This is because FortiManager asks you to type the name of the interface. You must type the correct interface name; it must correspond to a normalized interface that you have already defined.

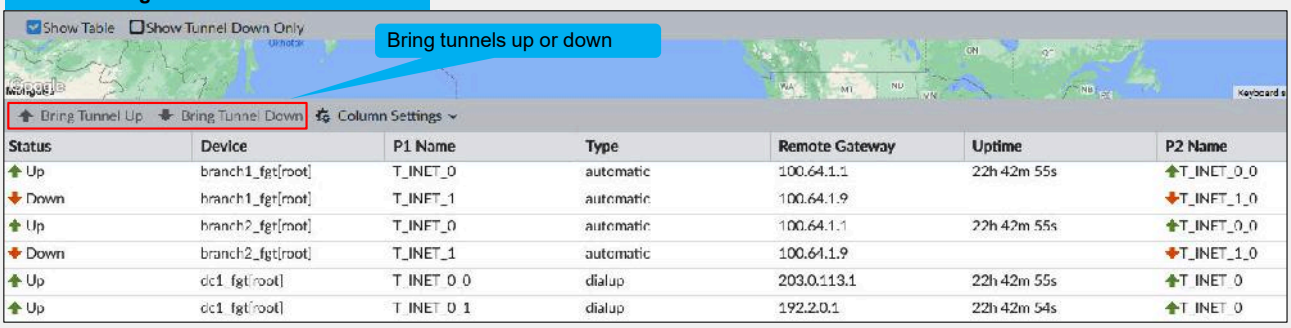
- If you plan to use per-device mapping, then you must install the VPN configuration first. This is because in per-device mapping, FortiManager looks for existing interfaces in the device settings database, which are created only after the installation is performed.

After you configure the normalized IPsec interface, you can reference the interface on firewall policies.

DO NOT REPRINT
© FORTINET

Map View and Monitor

Device Manager > Monitor > VPN Monitor



Status	Device	P1 Name	Type	Remote Gateway	Uptime	P2 Name
Up	branch1_fgt[root]	T_INET_0	automatic	100.64.1.1	22h 42m 55s	T_INET_0_0
Down	branch1_fgt[root]	T_INET_1	automatic	100.64.1.9		T_INET_1_0
Up	branch2_fgt[root]	T_INET_0	automatic	100.64.1.1	22h 42m 55s	T_INET_0_0
Down	branch2_fgt[root]	T_INET_1	automatic	100.64.1.9		T_INET_1_0
Up	dc1_fgt[root]	T_INET_0_0	dialup	203.0.113.1	22h 42m 55s	T_INET_0
Up	dc1_fgt[root]	T_INET_0_1	dialup	192.2.0.1	22h 42m 54s	T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

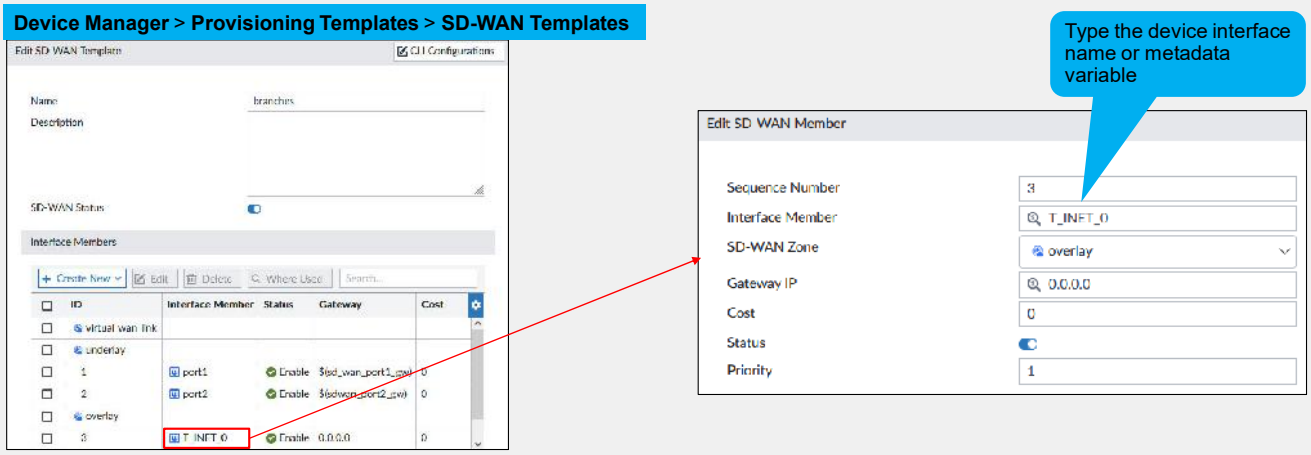
You can monitor the tunnel status and force them up or down from **VPN Monitor** page, as shown on this slide.

Note that Map View is available only for tunnels configured with VPN manager. You must select **Show Table** to monitor tunnels configured with tunnel templates.

DO NOT REPRINT
© FORTINET

IPsec Interfaces as SD-WAN Members

- Add the IPsec interface as SD-WAN member
  - Type the device interface name or meta field
    - Normalized interfaces are not used
  - Remove references to individual members on firewall policies first



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

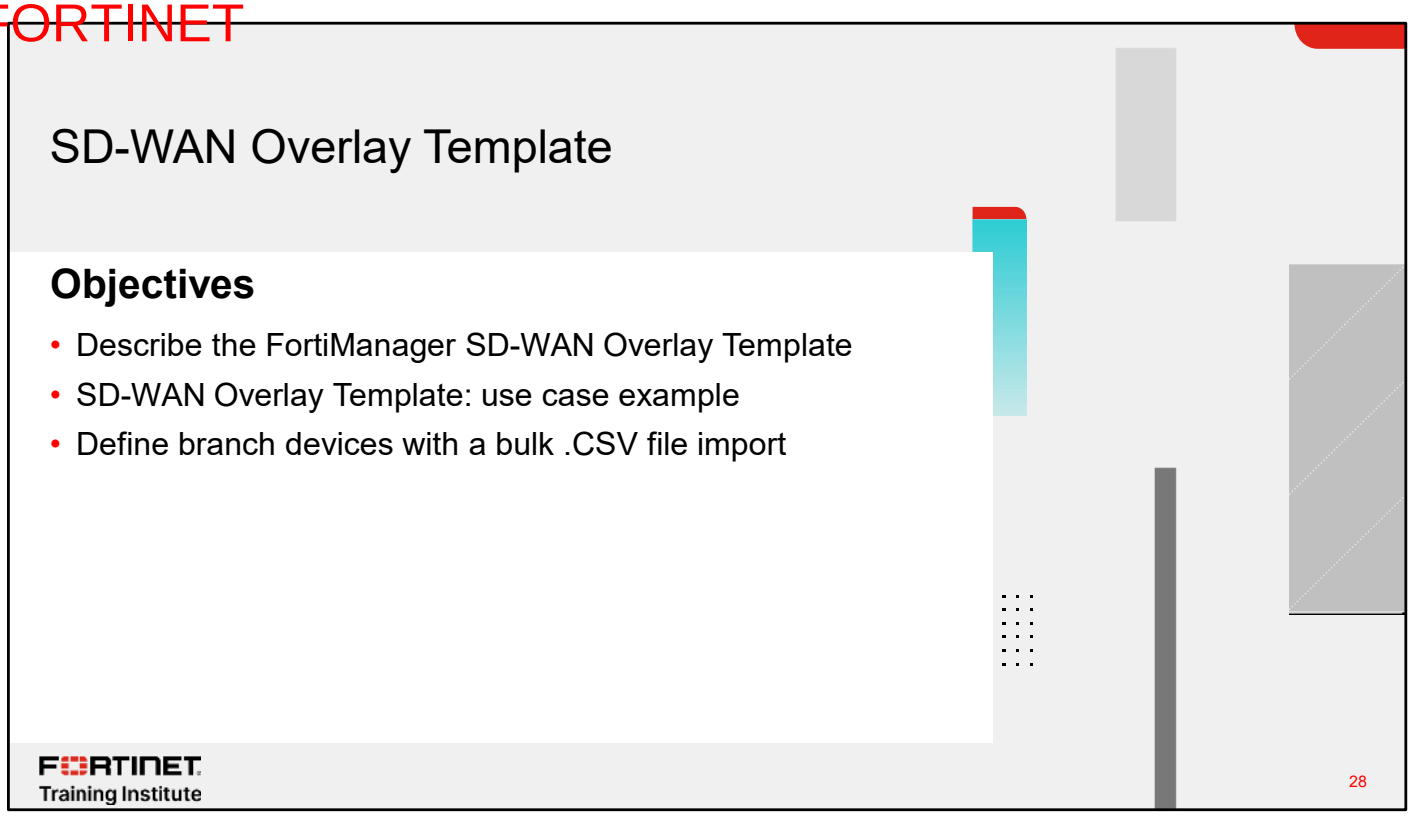
27

After you configure your IPsec tunnels, you can start using them as SD-WAN members. Just make sure to remove any existing references to the individual members on firewall policies. Otherwise, FortiManager can't install the SD-WAN configuration because you can't reference individual SD-WAN members on firewall policies. You can reference SD-WAN zones only.

Note that SD-WAN members don't use normalized interfaces. Instead, you must type the name of the interface on the device or use a metadata variable.

In the example shown on this slide, **T_INET_0** is the name of the IPsec interface on the device. The interface is configured as an SD-WAN member and placed in the **overlay** zone.

DO NOT REPRINT
© FORTINET



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

By demonstrating competence in the FortiManager SD-WAN overlay templates, you will be able to configure and deploy a large SD-WAN topology and associated overlay IPsec network with reduced effort.

DO NOT REPRINT
© FORTINET

FortiManager SD-WAN Overlay Template

- SD-WAN overlay template simplifies the configuration and deployment of secured SD-WAN topology
- Wizards guide administrator to define overlay configuration and routing for SD-WAN topology
- Available per ADOM
- Steps:
  1. Define network topology and subnets to use
  2. Initial basic device configuration (administrator settings, interface IP/subnet, and so on)
  3. Import devices to FortiManager
  4. **Follow SD-WAN overlay template wizards**
  5. Review configuration
  6. Install configuration to devices

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

Most SD-WAN deployments require complex overlay configurations for datacenter or cloud connectivity. FortiManager includes an SD-WAN overlay template with a wizard to automate and simplify the process using the IPsec and BGP settings recommended by Fortinet. SD-WAN overlay template wizards guide you through the configuration steps and generate a set of consistent configuration templates for SD-WAN, overlay, and routing configurations for hubs and branches devices.

Note that you can't use VPN manager and SD-WAN overlay templates for the same network. When applicable, for a new deployment, you should use the SD-WAN overlay template. It provides a comprehensive approach and guides you through IPsec, BGP, and SD-WAN configuration.

Do the following to use the SD-WAN Overlay Template:

1. Define the network topology and subnets to use.
2. Preconfigure your network and SD-WAN devices with administrator settings, the interface IP/subnet and so on.
3. Import devices to FortiManager.
4. Let the SD-WAN Overlay Template guide you through the configuration wizards.
5. Review the configuration templates generated.
6. Install the configurations on devices.

Note that if you need to add additional branch devices later, you can do that by adding devices to the branch device group.

DO NOT REPRINT
© FORTINET

Prerequisites and Network Planning

- Prerequisites
  - Import FortiGate devices that will make the hub (or hubs)
  - Create a device group for your branch devices
  - Configure the WAN, ISP links, and other interfaces on your imported devices
- Optional
  - Import branch devices into FortiManager
- Network planning
  - Allocate the overlay network address space (default template value 10.10.0.0/16)
  - Allocate the loopback IP address space (default template value 172.16.0.0/16)
  - Select an AS number for BGP for the new SD-WAN overlay region (default template value 65000)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

Before you can define the SD-WAN Overlay Template, the first step is to import hub devices into FortiManager—you can use model devices—and configure the necessary links and interfaces.

You can import branch devices in this preliminary phase or later, but you must create a device group that contains the branch devices before you proceed further with the SD-WAN overlay template.

Network planning is an important phase.
You must define the following elements before using the SD-WAN overlay template:

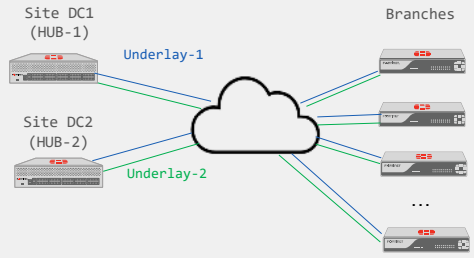
- Topology type
- Overlay network address space
- Loopback IP address space
- BGP AS number for SD-WAN overlay region

Note that to use the SD-WAN overlay template, your configuration must include a single overlay network.

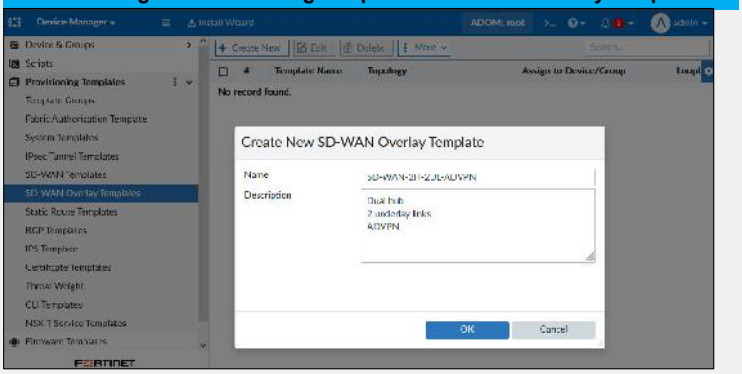
DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Example

- Through the next few slides, you will see how to use SD-WAN overlay template
- You will use the overlay template to create templates for a dual-hub topology with two underlay links on each device and ADVPN to allow shortcut tunnels between branches



Device Manager > Provisioning Templates > SD-WAN Overlay Templates



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

SD-WAN overlay templates help you to configure templates for single-hub and dual-hub topologies. Dual-hub topology offers options for primary and secondary hubs or two hubs active simultaneously (called Primary & Primary).

Over the next few slides, you will explore how to use the wizards to create templates to configure devices for topology shown on diagram.

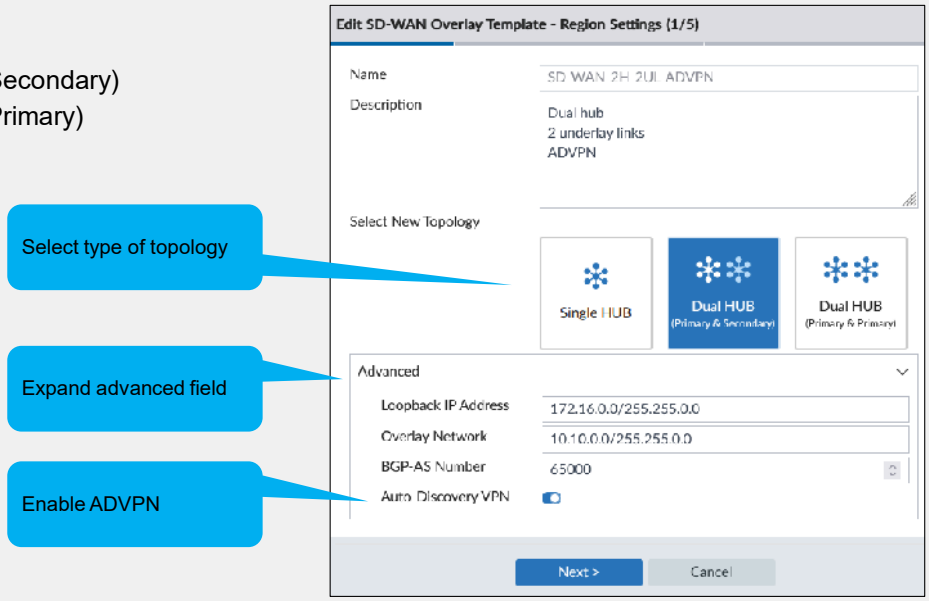
We will use following parameters:

- Dual Hub (Primary & Secondary)
- Two underlay links for each devices
- iBGP routing
- ADVPN

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Step1—Region Settings

- Define topology type
  - Single HUB
  - Dual HUB (Primary & Secondary)
  - Dual HUB (Primary & Primary)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

First you define the type of topology used, either single hub or dual hub and the type of redundancy.

In the **Advanced** section, apply settings for the network elements:

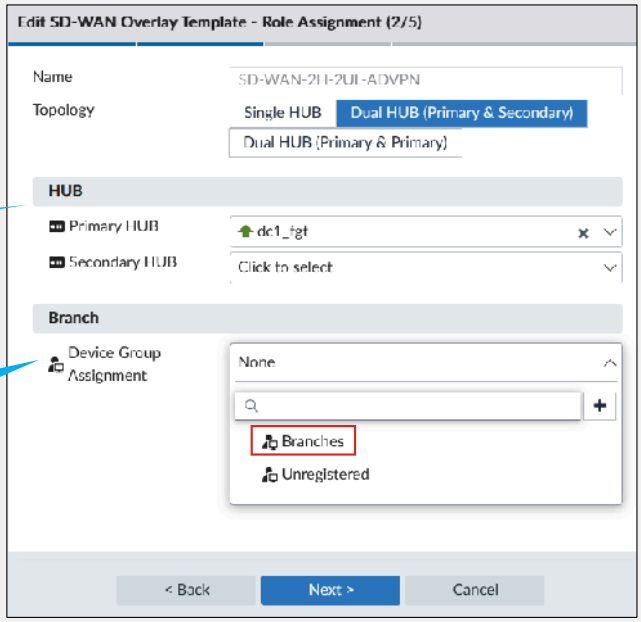
- Loopback IP Address
- Overlay Network
- BGP-AS Number
- Auto-Discovery VPN

The SD-Wan overlay template guides you through multiple wizards. FortiManager adjust the fields available on each wizard according to the option you selected during this first step.

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Step 2—Role Assignment

- Define device roles
  - Select devices with hub role
  - Select or define a device group for branch devices



Select hub devices (already on FortiManager)

Select device group for branch devices

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

At step 2, you define device roles.

For hub devices—two for a dual-hub topology—you must select from devices already discovered on FortiManager.

For branch devices, at this stage, you must define a branch devices group. This group can be empty while you are working on the template. Later in this lesson, you will learn how to perform a bulk device import from a CSV file. Later, you can add devices to the branch device group as the network evolves and new sites join the SD-WAN topology.

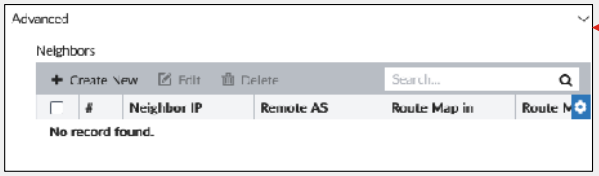
DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Step 3—Network Configuration

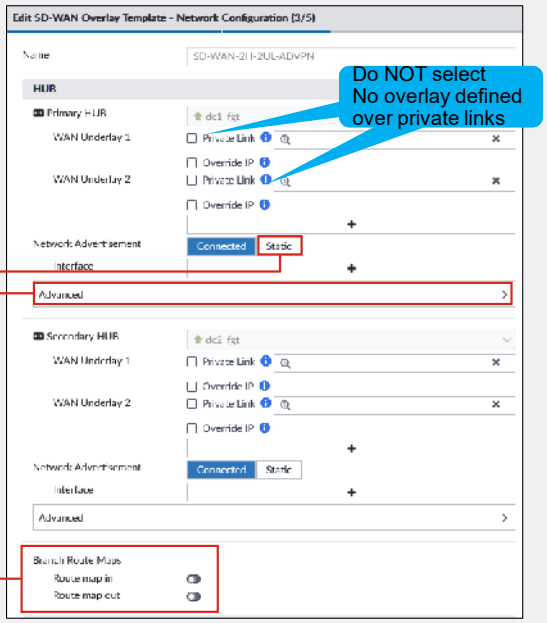
- Part 1: Network settings for hubs
  - Define underlay interfaces
  - Define type of network advertisement
    - Connected: redistribute hub-connected subnets
    - Static: distribute specified subnets



  - Optional: specify BGP neighbors



  - Route mapping in/out from branches



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

Step 3 consists of network configuration for hub and branch devices.

At this stage, you will define the network interfaces used for underlay links, usually static definition for hubs and metadata variables for branches.

The **Private Link** checkbox indicates a secure internal link. FortiManager will NOT define an overlay link through this interface.

You can also add additional underlay links at this stage (**+** sign).

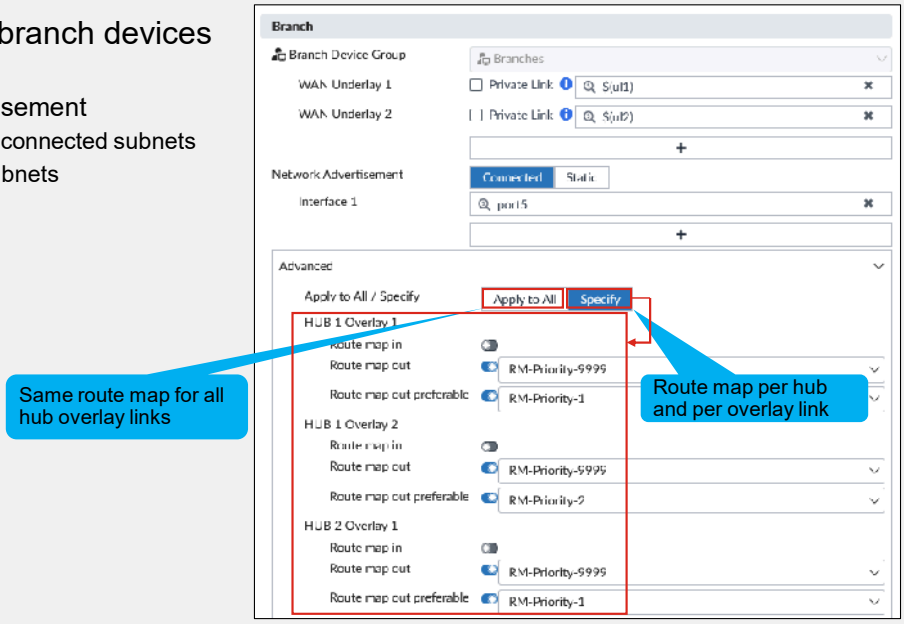
For network advertisements, you can decide between connected subnets advertisement or static network prefixes definition. For connected subnets advertisement you must select interfaces that correspond to subnets to advertise.

Branch route maps correspond to route maps that the hub applies to a branch neighbor group.

DO NOT REPRINT
© FORTINET

Step 3—Network Configuration (Contd)

- Part 2: Network settings for branch devices
  - Define underlay interfaces
  - Define type of network advertisement
    - Connected: redistribute hub-connected subnets
    - Static: distribute specified subnets
  - Define BGP route map
    - Route map in/out with hubs
    - Route map out preferable



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

Usually, to define WAN underlay interfaces for branch devices you use metadata variables. This allows you to accommodate various types of branch sites and branch devices. Similarly, you can use metadata variables for connected subnet interface definitions.

In the **Advanced** section you can define route mapping per hub overlay link or globally for all links. You can define **Route map in**, **Route map out**, and **Route map out preferable** settings.

Route map out and **Route map out preferable** allow differentiated route map advertisements. If an SD-WAN member meets the SLA threshold, FortiGate applies the route map defined in the BGP neighbor's `route-map-out-preferable` option.

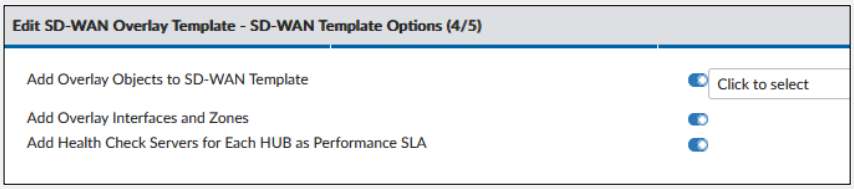
If the SD-WAN member fails to meet the SLA, FortiGate applies the route map defined in the BGP neighbor's `route-map-out` option.

This allows FortiGate to advertise the health of the SD-WAN member to its BGP neighbor by advertising different community strings based on SLA status.

You will explore route maps options in greater details in another lesson.

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Step 4—Template Options



- Add Overlay Objects to SD-WAN Template
  - Call an existing SD-WAN template or create one
- Add Overlay Interfaces and Zones
  - Adds overlay IPsec interfaces created to SD-WAN members list
  - Create zones for each hub
- Add Health Check Servers for Each HUB as Performance SLA
  - Create health checks pointed to loopback IPs on each hub

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

At step 4, you must define the SD-WAN template that the SD-WAN overlay template process uses. The process does not create it automatically.

You can call an existing SD-WAN template, and the SD-WAN overlay template process adds settings to it according to the options you selected. Usually, the process adds interface members and zones, at a minimum. If you have no predefined SD-WAN template, you must create one.

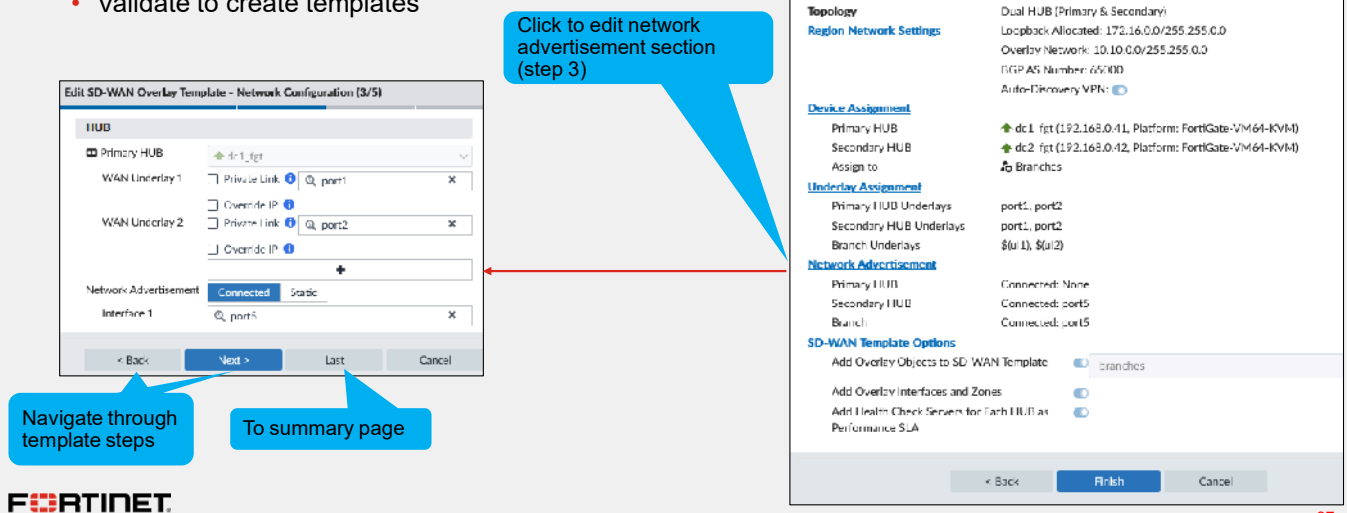
If you select **Add Health Check Servers for Each HUB as Performance SLA**, the process creates one performance SLA check for each hub loopback IP. These SLA checks are defined with the default values detection protocol Ping, failure threshold 5, and recovery threshold 5. Once created, you can edit the SD-WAN template to review and adjust the values as required.

The wizard does not create SD-WAN rules. You can add them later and refine them as requirements evolve.

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Step 5—Summary

- Last step of the process
  - Review all settings
  - Go directly to section if changes are required
  - Validate to create templates



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

37

The last step, step 5, is a review phase.

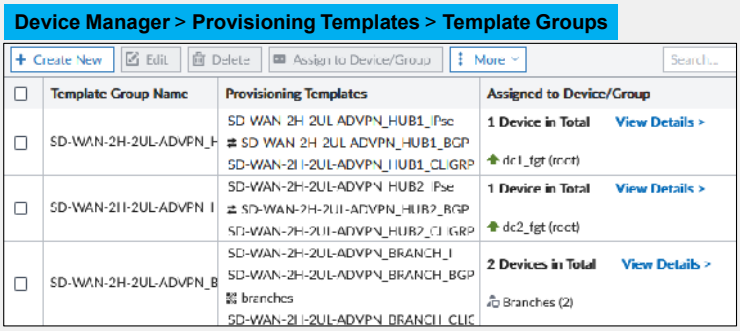
You can review all settings and, if adjustments are required, go directly to the appropriate menu. Once you are happy with the settings, click **Finish**.

You will see the results on next slide.

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Result

- SD-WAN overlay template generates multiple templates
  - SD-WAN (branches only)
  - BGP
  - IPsec tunnels
  - CLI
- It groups templates per devices
  - HUB1
  - HUB2
  - BRANCH
- Administrator can
  - Review
  - Edit
  - Install to devices



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

38

SD-WAN overlay template generates multiple templates, one per topic.

For the branch devices, it creates settings to the SD-WAN template defined.

For branches and hubs, it creates BGP, IPsec, and CLI templates to accommodate all required configuration changes.

The templates are grouped by device, and device assignment is done for you.

At this stage, you can review each template and fine tune settings as required before you apply the configurations to devices.

DO NOT REPRINT
© FORTINET

SD-WAN Overlay Template—Post-Run Tasks

- Mandatory
  - Assign metadata variables to devices
    - Created by SD-WAN overlay template
      - Branch_id
    - User-created variables
      - Device name
      - Gateway
      - Interface IP
      - and so on
  - Configure the SD-WAN rules by editing the SD-WAN template
  - Create policy packages for branch and hub devices
  - Install the changes to SD-WAN devices using the install wizard
- Optional
  - Edit SD-WAN overlay template
  - Add new branch devices

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

39

After completing the SD-WAN overlay template wizard, you must complete some additional tasks.

The SD-WAN overlay template automatically creates a metadata variable called **branch_id**. You must define a unique branch ID value for each branch device. If you created additional metadata variables through the process, you should also define corresponding values for each device. Usually, you use metadata variables for device name, interface IP, or gateway.

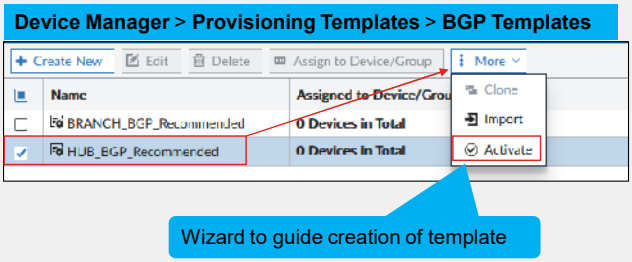
You will also edit the SD-WAN template to configure SD-WAN rules and define SLA criteria, create policy packages for your branch and hub devices, and then, install changes to the device with the install wizard.

Later, you can return to the SD-WAN overlay template to edit and modify settings, and add new branch devices.

DO NOT REPRINT
© FORTINET

Fortinet Recommended Templates

- Predefined template per topic
- Guide user for recommended settings
- Flexibility
  - Use recommended template for some topics
  - Use custom template for others
- Available recommended templates
  - IPsec tunnel templates
    - BRANCH_IPsec_Recommended
    - HUB_IPsec_Recommended
  - BGP templates
    - BRANCH_BGP_Recommended
    - HUB_BGP_Recommended



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

40

SD-WAN overlay template guides you through the creation of SD-WAN, BGP, and IPsec tunnel templates in a combined process.

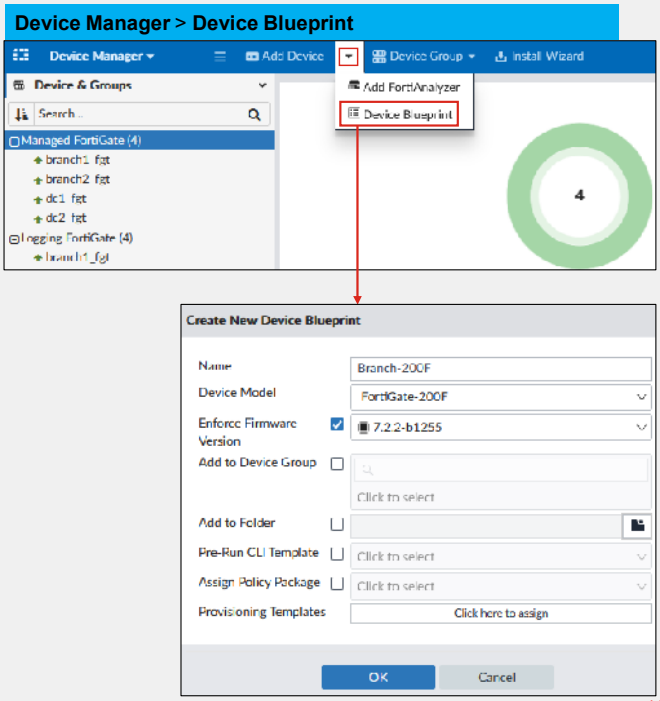
If you require guidance for some tasks only, you can use Fortinet recommended templates. They are available for IPsec and BGP.

To use a recommended template, select the desired template and activate it. It creates a personalized copy and, using a wizard, guides you through the configuration.

DO NOT REPRINT
© FORTINET

Device Blueprint

- Create model device
  - One blueprint per device model
  - Device basic device setting
- Device Blueprint can
  - Enforce firmware version
  - Add device to group
  - Pre-run CLI template
  - Assign policy package
  - Assign provisioning template



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

41

Device blueprints can simplify new deployments that use multiple devices of the same type and with similar configurations. They allow you to predefine various parameters per FortiGate device model, such as firmware version, assigned templates, or policy package.

Note that if you decide to assign the devices to a group, they will inherit all templates associated to that device group. Therefore, to avoid the risk of conflicting configuration instructions, you should either assign provisioning templates or add devices to the group to benefit from group templates. You should not do both on the same blueprint.

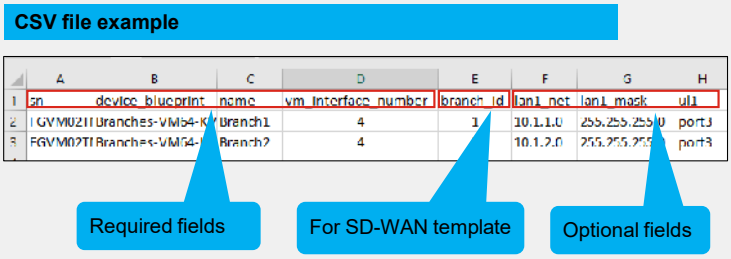
You can use a device blueprint in conjunction with a CSV file import to simplify the task of adding all branch devices to FortiManager.

DO NOT REPRINT
© FORTINET

Import Model Devices From CSV File

- Bulk import of devices to FortiManager
- Easy CSV file creation
- Useful for large deployments

- CSV file elements
  - Required
    - Device serial number
    - Blueprint
    - Device name
    - Number of ports to provision for VM models
  - Required for SD-WAN template
    - Branch_id
  - Optional
    - User-created metadata variables for each device



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

For new SD-WAN deployments, you must add to FortiManager many similar FortiGate devices used on branch sites. Adding devices one by one would be repetitive and time consuming.

To simplify this step, FortiManager provides a CSV file import feature. You must build the file as comma-separated values (CSV) with device serial numbers, device names, associated blueprints and, for VM models, the number of interfaces.

The SD-WAN overlay template automatically creates a branch_id metadata variable. If you use a CSV file to import SD-WAN branch devices, you must assign a unique branch_id value for each device. In addition to those required parameters, you can use a CSV file to define various per-device parameters with metadata variables.

The CSV file must be built with column headers and the corresponding list of values. The first columns must be: **sn**, **device_blueprint**, **name** and, for VMs, **vm_interface_number**. Additional optional column names must match exact metadata variable names.

On file import, FortiManager sets metadata variable values for every predefined value and ignores columns that match already defined metadata variables.

Note that the **Add Model Device** and **Import Model devices from CSV file** modes are intended for new FortiGate deployments, where no pre-existing configuration on the FortiGate device must be preserved. The configuration associated with the model device overwrites the configuration of the FortiGate device as part of the installation process, after FortiManager authorizes the FortiGate device.

DO NOT REPRINT
© FORTINET

Review

- ✓ Understand FortiManager basics
- ✓ Identify FortiManager features for SD-WAN
- ✓ Describe SD-WAN management on FortiManager
- ✓ Describe IPsec templates available on FortiManager
- ✓ Use IPsec templates to configure hub-and-spoke IPsec VPNs
- ✓ Configure IPsec interfaces as SD-WAN members
- ✓ Describe the FortiManager SD-WAN Overlay Template
- ✓ SD-WAN Overlay Template: use case example
- ✓ Define branch devices with a bulk .CSV file import

FORTINET Training Institute

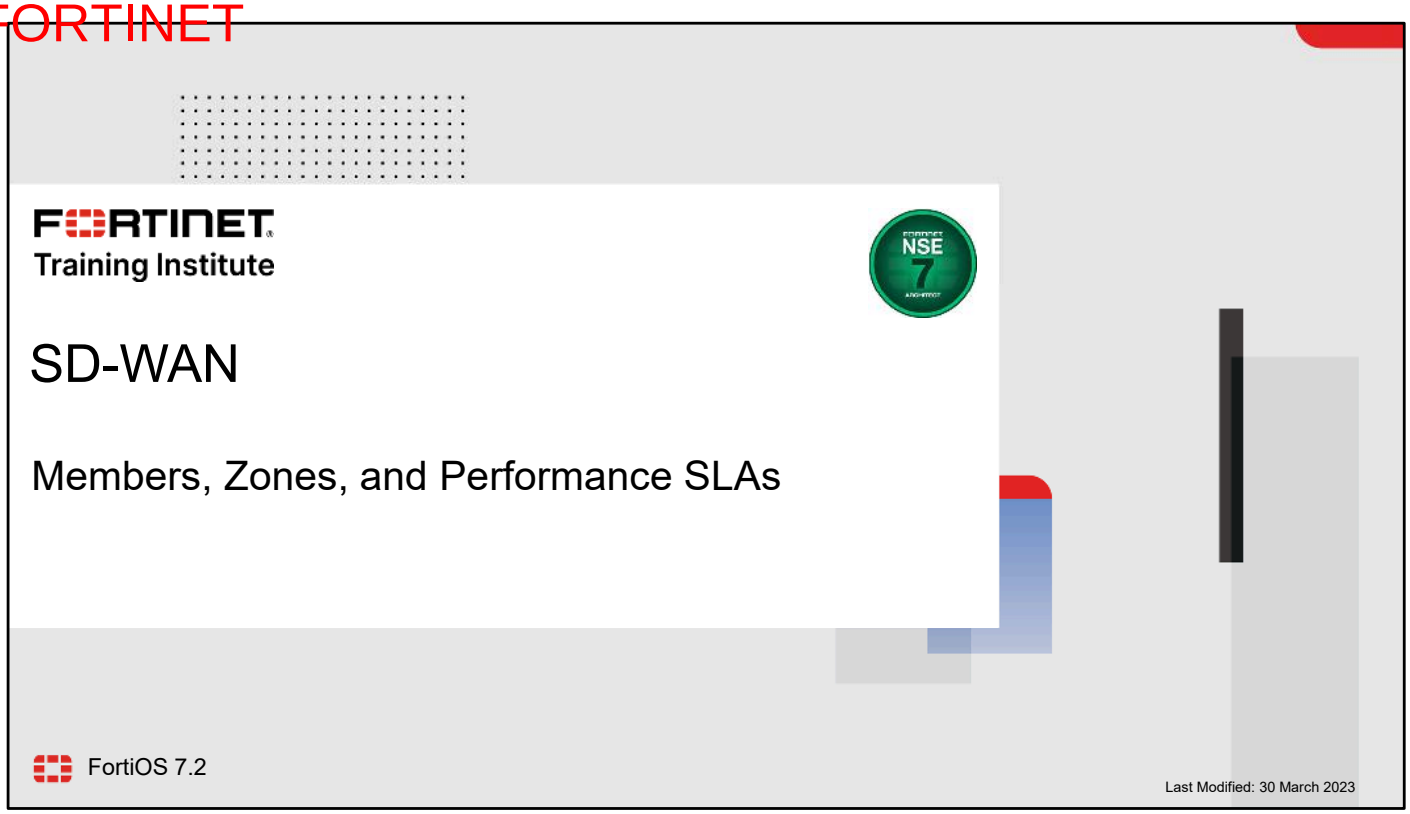
© Fortinet Inc. All Rights Reserved.

43

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

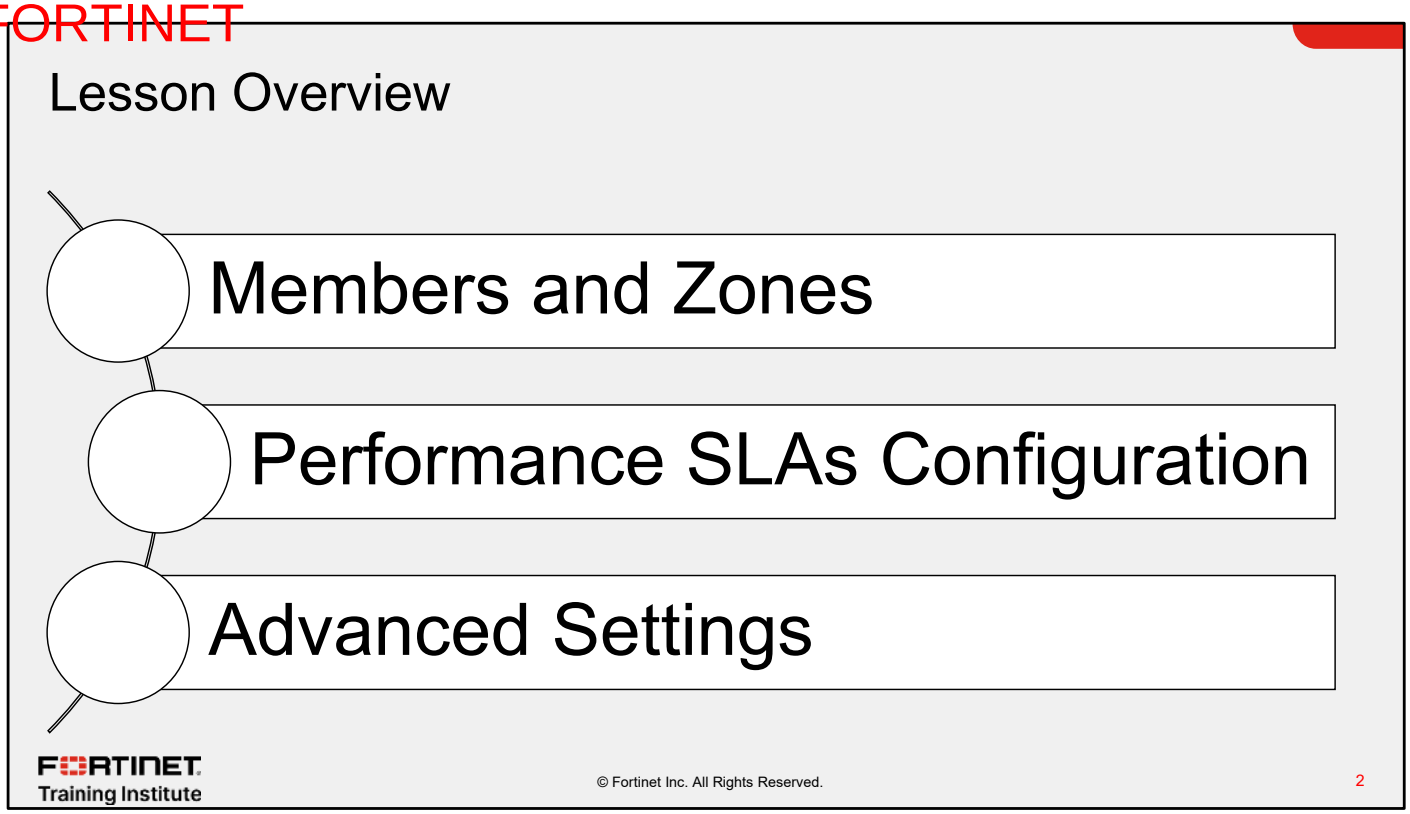
By mastering the objectives covered in this lesson, you learned the most useful features available in FortiManager for deploying SD-WAN in large networks.

DO NOT REPRINT
© FORTINET



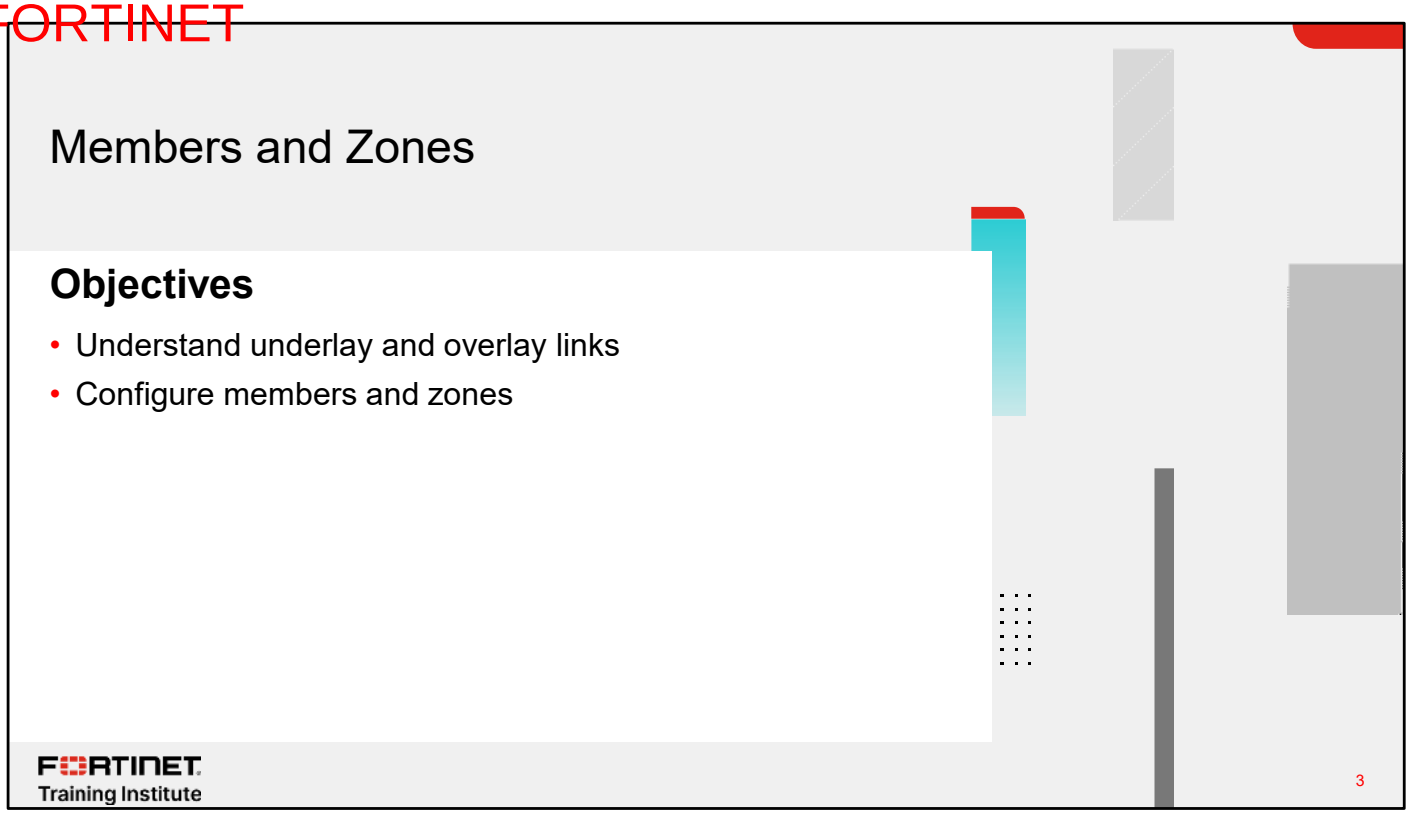
In this lesson, you will learn about members, zones, and performance SLAs, in detail.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



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

By understanding SD-WAN members and zones, you should be able to select and organize your underlay and overlay links for effective WAN usage.

DO NOT REPRINT
© FORTINET

SD-WAN Members—Underlay and Overlay Links

- Underlay:
  - Physical links provided by ISP
    - Cable, DSL, fiber, MPLS, 3G/4G/LTE, ATM
  - Restricted routing
  - No added security
- Overlay:
  - Virtual links built on top of underlay links
    - IPsec, GRE, IP-in-IP
  - Flexible routing
  - Enhanced security

Supported SD-WAN Members*	
Interface	**Type**
Physical	Underlay
VLAN	
LAG	
3G/4G/LTE USB modems	
FortiExtender	
IPsec (including ADVPN)	Overlay
GRE	
IP-in-IP	

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

The terms *underlay* and *overlay* are used to describe the link type of an SD-WAN member.

Underlays refer to the physical links that you can rent or buy from an ISP, such as cable, DSL, fiber, MPLS, 3G/4G/LTE, and ATM links. These links are part of the ISP physical infrastructure that is responsible for delivering packets across networks. The traffic that travels through underlays is restricted to the routing policies deployed by the ISP, and therefore, the packet source and destination IP addresses must be routable within the ISP network. This restriction leaves you with limited options to define your network addressing plan. In addition, traffic transmitted through underlays is usually not encrypted by the ISP network, which means that sensitive data can be accessed by unauthorized parties if the data is not encrypted by the sender.

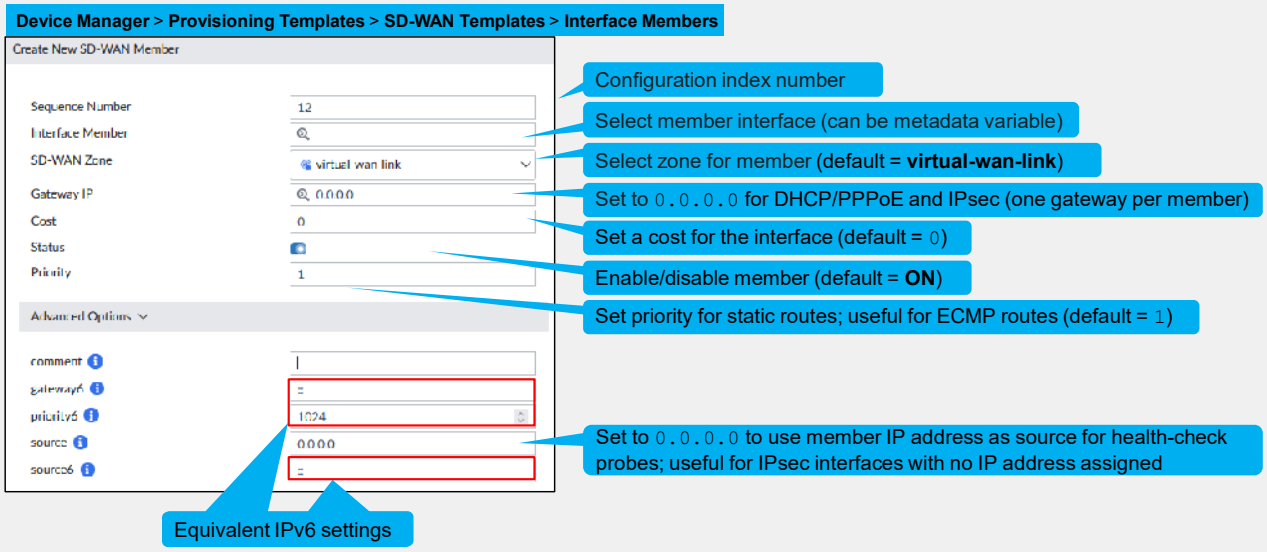
Overlays are virtual links that you build on top of underlays. A common example of an overlay is an IPsec tunnel. Because original packets are often encapsulated in ESP packets, the networks that communicate through the IPsec tunnel are no longer restricted to the routing policies of the ISP. In addition, the privacy and authentication features provided by IPsec protect your traffic from unauthorized access.

This slide shows the different underlay and overlay links supported by FortiGate as SD-WAN members.

DO NOT REPRINT
© FORTINET

SD-WAN Members—Configuration

- FortiManager GUI:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

When you configure an SD-WAN member using the **Device Manager** or the **SD-WAN provisioning template** on FortiManager, you can configure the following settings:

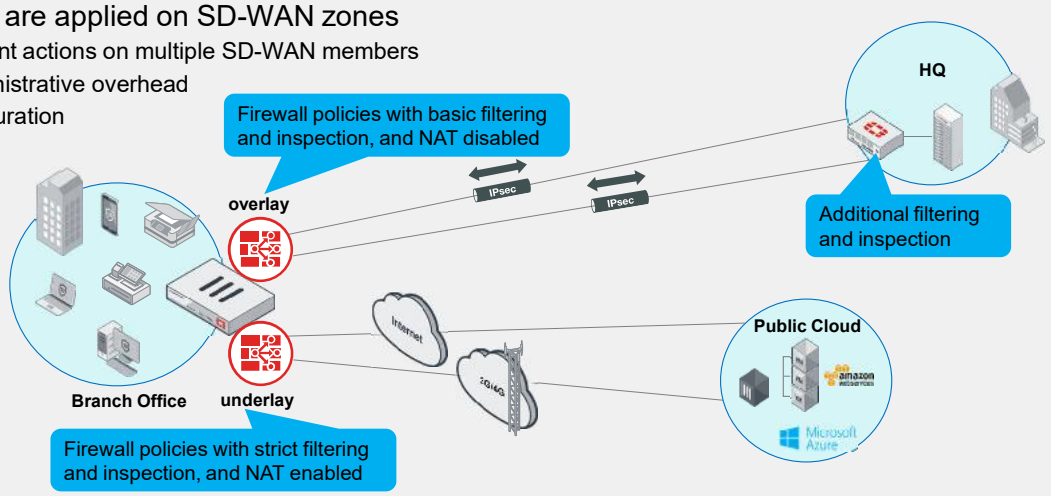
- **Sequence Number**: Sets the configuration index number of the member. It's automatically assigned by FortiManager in sequential order, but can be changed by the administrator.
- **Interface Member**: Type the name of the interface on the device to use as the SD-WAN member, or a metadata variable.
- **SD-WAN Zone**: Select the zone to place the member in. A member can be assigned to one zone only.
- **Gateway IP**: Indicates the IPv4 gateway address to use by the member. Set the address to `0.0.0.0` if the interface is DHCP, PPPoE, or IPsec. This instructs FortiGate to automatically use the gateway assigned by the ISP. The gateway is then used for static routes and health-check probes. Note that you can assign one gateway only per member.
- **Cost**: Set the cost of the member. The cost serves as a tiebreaker in SD-WAN rules using the lowest cost as strategy.
- **Status**: Enable or disable a member. When disabled, the member is not used in SD-WAN.
- **Priority**: Assigns the priority of static routes created by SD-WAN. The priority setting is useful for prioritizing one ECMP route over other ECMP routes. The default priority is 1 (it was 0 in previous versions).
- **Source**: Indicates the member source IP address for health-check probes. If set to `0.0.0.0`, FortiGate uses the primary IP address of the member interface as source. This setting is useful for IPsec interfaces with no IP address assigned.

The gateway, priority, and source settings have equivalent settings for IPv6, as shown on this slide.

DO NOT REPRINT
© FORTINET

SD-WAN Zones

- Divides SD-WAN members in groups
  - Default zone: **virtual-wan-link**
    - Can't be deleted
  - An interface can belong to one zone only
- Firewall policies are applied on SD-WAN zones
  - Perform different actions on multiple SD-WAN members
  - Reduces administrative overhead
  - Cleaner configuration
- Usually underlay and overlay links are grouped into different SD-WAN zones



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

Usually, you should apply a different set of policies based on the link type of your SD-WAN members. For example, you may want to enable NAT and apply strict security policies to internet traffic sent through underlay links, because the traffic directly leaves the site boundaries. Conversely, you may want to disable NAT and apply basic filtering and inspection to traffic sent through overlay links, because the remote site is fully routable and performs additional filtering and inspection on the traffic.

SD-WAN zones allow administrators to group together members that require a similar set of firewall policies. Usually, this means grouping underlays and overlays into different SD-WAN zones.

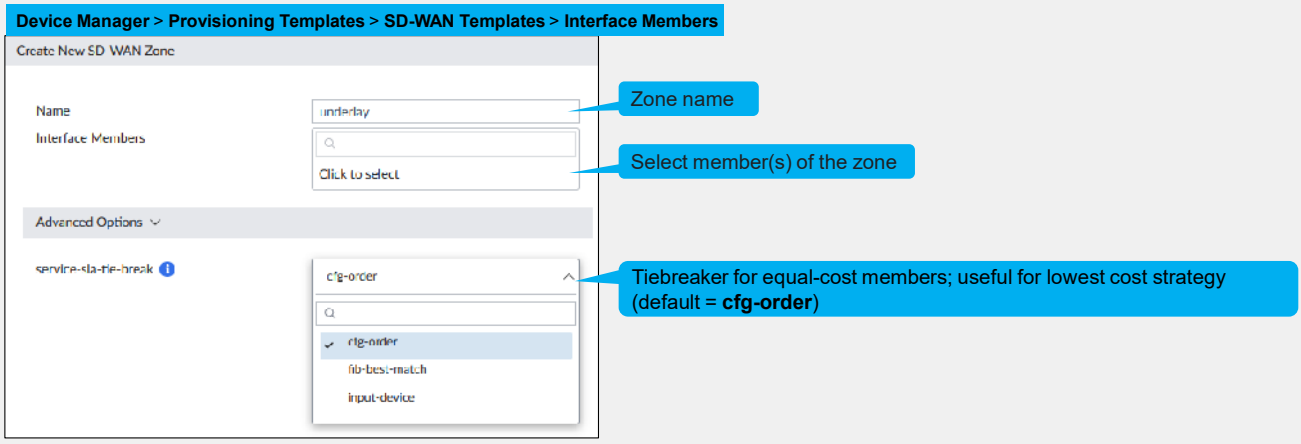
The **virtual-wan-link** SD-WAN zone is created by default and can't be deleted. It contains any SD-WAN member not explicitly assigned to a user-defined SD-WAN zone. Firewall policies defined for your SD-WAN traffic, must reference the SD-WAN zones, and cannot reference individual SD-WAN members.

The topology shown on this slide shows a branch office with two SD-WAN zones configured: overlay and underlay. The overlay SD-WAN zone is composed of IPsec tunnels and the underlay SD-WAN zone is composed of an internet link and a 3G/4G link. The branch office uses the overlays to access the headquarter networks, and the underlays to access services in the public cloud. By dividing SD-WAN members into zones, you can apply the same set of firewall policies to a zone instead of having to apply them to their individual members, thus reducing the administrative overhead and building a cleaner configuration.

DO NOT REPRINT
© FORTINET

SD-WAN Zones—Configuration

- FortiManager GUI:



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

When you configure an SD-WAN zone using **Device Manager** or the **SD-WAN provisioning template** on FortiManager, you can configure the following settings:

- **Name**: Type in a name for the zone.
- **Interface Members**: Select one or more members to be included in the zone.
- **service-sla-tie-break**: This is applicable to all SD-WAN rule strategies except **Maximize Bandwidth (SLA)**.
  - **cfg-order** instructs FortiGate to use the member configuration order as the tiebreaker for the selected member. That is, members that are configured first, have higher priority.
  - **fib-best-match** uses as tiebreaker the most specific route. That is, the member with the most specific route to the destination becomes the selected member.
  - **input-device** is used for overlay stickiness. Traffic received on one overlay should egress on same overlay, when possible. It instructs FortiGate to send the reply through the overlay link that received the incoming flow, provided it matches the SLA criteria.

Note that when you create a zone with the FortiManager SD-WAN template, FortiManager automatically creates a corresponding normalized interface.

You will learn more about SD-WAN rules and the available strategies in another lesson.

DO NOT REPRINT
© FORTINET

SD-WAN Members and Zones—CLI Configuration

- FortiGate CLI configuration:

```
config system sdwan
    config zone
        edit underlay
            set service-sla-tie-break cfg-order
        next
    end
    config members
        edit 1
            set interface port1
            set zone underlay
            set gateway 0.0.0.0
            set source 0.0.0.0
            set gateway6 ::
            set source6 ::
            set cost 0
            set priority 1
            set priority6 1024
            set status enable
        next
    end
end
```

Setting hidden for IPsec interfaces

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

This slide shows the equivalent CLI configuration for the SD-WAN member and zone created on FortiManager on the previous slides. Although this course is focused on SD-WAN deployment from FortiManager, it is useful to know the corresponding SD-WAN FortiGate CLI settings in case you want to apply the configuration using FortiManager CLI templates and scripts.

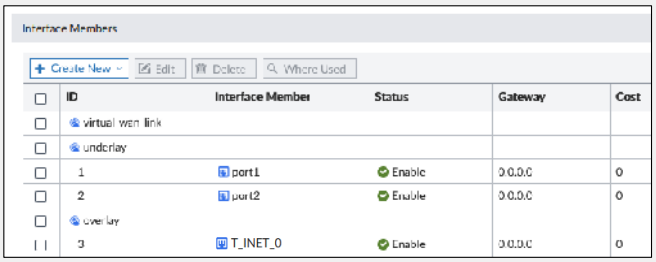
Note that the `gateway` and `gateway6` settings are not available if `interface` is set to an IPsec interface. This is because FortiOS automatically uses the tunnel ID as gateway for IPsec interfaces. Starting with FortiOS 7.0, FortiGate uses tunnel IDs to determine the next hop for IPsec traffic. The tunnel ID can be found in the output of `diagnose vpn tunnel list`. Usually, FortiOS uses as tunnel ID the IP address configured as `remote-gw` in the phase 1 settings.

Also note that the output shown on this slide has been modified to include a mix of settings with default and non-default values. On the FortiGate CLI, you can run `show full-configuration system sdwan` to display the SD-WAN settings with default values.

DO NOT REPRINT
© FORTINET

SD-WAN Members and Zones—CLI Configuration (Contd)

- Sample FortiManager configuration:



- Equivalent FortiOS CLI configuration:

```
config system sdwan
    set status enable
    config zone
        edit "virtual-wan-link"
        next
        edit "underlay"
        next
        edit "overlay"
        next
    end
    config members
        edit 1
            set interface "port1"
            set zone "underlay"
        next
        edit 2
            set interface "port2"
            set zone "underlay"
        next
        edit 3
            set interface "T_INET_0"
            set zone "overlay"
        next
    end
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

This slide shows a more elaborate member and zone configuration and its corresponding FortiOS CLI configuration. Note that, by default, the CLI configuration shows only settings with non-default values.

DO NOT REPRINT
© FORTINET

SD-WAN Members and Zones—CLI Verification

- FortiOS CLI verification:

```
# diagnose sys sdwan member
Member(1): interface: port1, flags=0x0 , gateway: 192.2.0.2, priority: 0 1024, weight: 0
Member(2): interface: port2, flags=0x0 , gateway: 192.2.0.10, priority: 0 1024, weight: 0
Member(3): interface: T_INET_0, flags=0xc , gateway: 100.64.1.1, priority: 0 1024, weight: 0

# diagnose sys sdwan zone
Zone overlay index=3
        members(1): 25(T_INET_0)
Zone underlay index=2
        members(2): 3(port1) 4(port2)
Zone virtual-wan-link index=1
        members(0):

# diagnose netlink interface list | grep index=25
if=T_INET_0 family=00 type=768 index=25 mtu=1420 link=0 master=0
```

Configuration index number

Automatic gateway detection—tunnel ID for IPsec tunnels (`diagnose vpn tunnel list`)

For volume-based load balancing

Kernel interface index number

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

In addition to checking the FortiOS CLI configuration for SD-WAN, you can also run the commands shown on this slide to verify the SD-WAN settings in place.

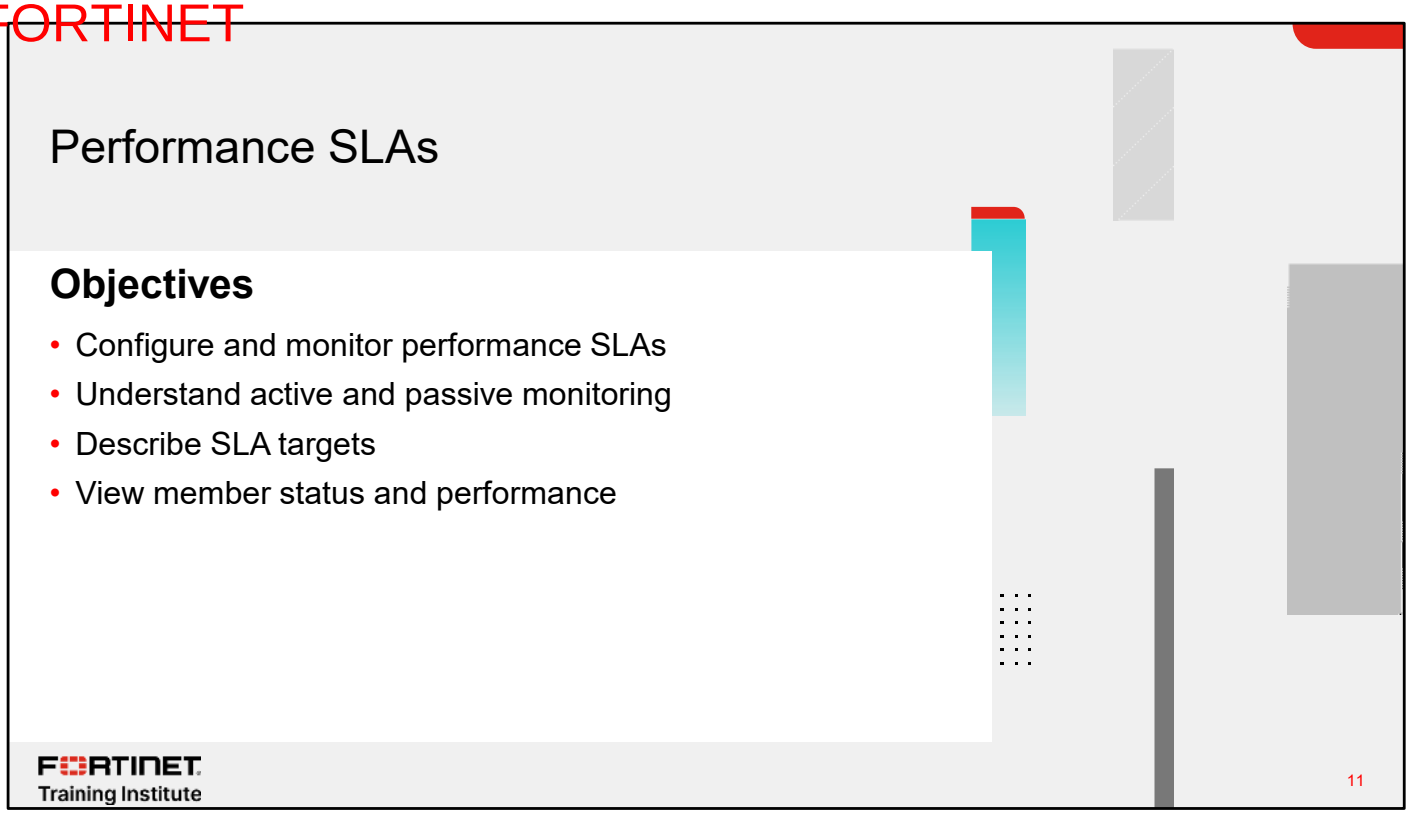
`diagnose sys sdwan member` displays the current settings of each member. Some of the settings that stand out in the output are the configuration index number, the gateway, and the weight. The configuration index number should match the member index under `config member` in `config system sdwan`.

If the member is configured with automatic gateway detection—that is, the gateway is set to `0.0.0.0`—or if the member is an IPsec interface, the output displays the gateway detected by FortiGate. If the member is an IPsec interface, FortiGate uses as gateway its tunnel ID. The tunnel ID can be found in the output of `diagnose vpn tunnel list`. Usually, FortiOS uses as tunnel ID the IP address configured as `remote-gw` in the phase 1 settings.

The weight setting applies to the volume-based load balancing algorithm only, which can be configured for the implicit SD-WAN rule. You will learn more about load balancing algorithms available for traffic that matches the implicit SD-WAN rule in another lesson.

`diagnose sys sdwan zone` displays the configured zones and their members. Note that the output indicates the kernel interface index number of a member, which should match the index displayed by `diagnose netlink interface list`.

DO NOT REPRINT
© FORTINET



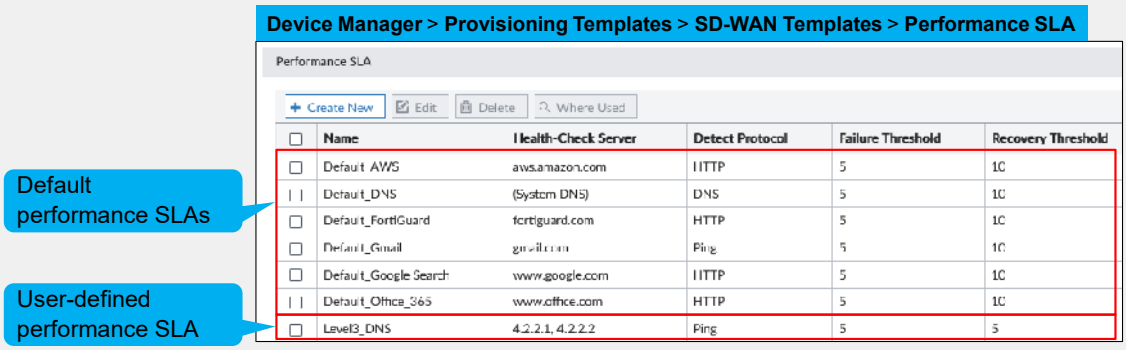
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding performance SLAs, you should be able to configure proper health check and performance monitoring for your SD-WAN members. Member health check and performance are key components in an effective SD-WAN deployment.

DO NOT REPRINT
© FORTINET

Performance SLAs

- Monitor member health:
  - State
    - Alive or dead
  - Performance
    - Packet loss, latency, and jitter
    - Mean Opinion Score (MOS)
    - SLA targets
      - Minimum performance requirements
- Health can be measured:
  - Actively
    - Based on periodic probes sent to configured servers
  - Passively
    - Based on member traffic



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

You can use performance SLAs to monitor the health and performance of members. Although configuring performance SLAs is optional, you should configure them to ensure members meet the health and performance requirements for steering traffic, which is critical for effective WAN usage with SD-WAN.

When you configure performance SLAs, you configure health checks and SLA targets. Health checks determine the state of members—alive or dead—and monitor their performance in terms of packet loss, latency, and jitter. You can also decide to combine the three criteria and determine what is called the mean opinion score or MOS. SLA targets define the minimum performance requirements for members to be eligible for steering traffic. SD-WAN then uses this information to make traffic steering decisions based on the SD-WAN rules configured. For example, you can instruct FortiGate to steer internet traffic to an alive member whose latency doesn't exceed a given threshold.

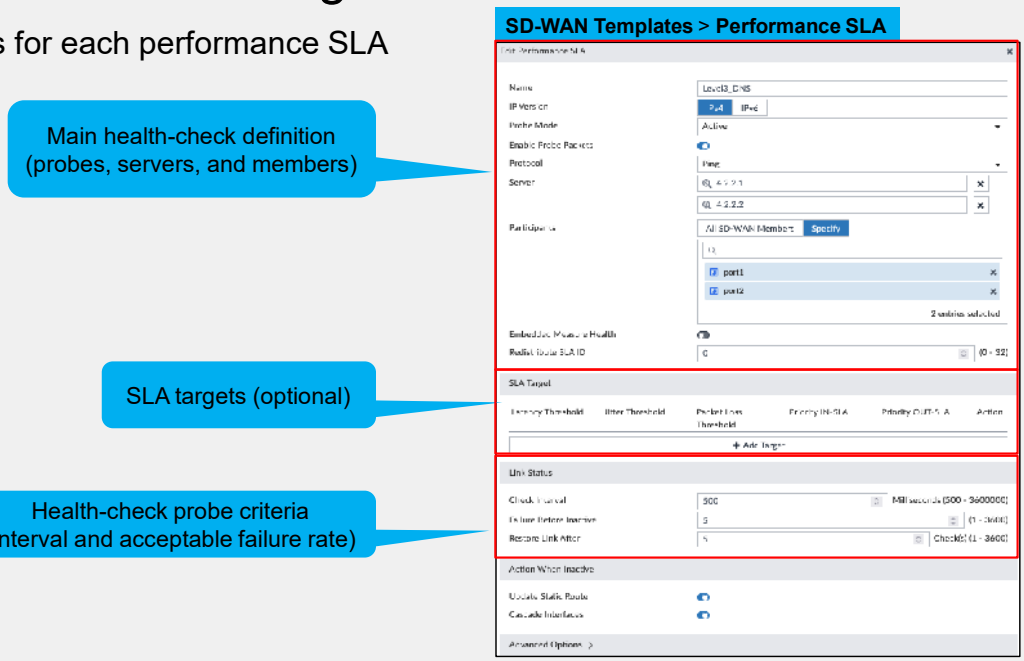
There are several performance SLAs created by default that you can use for your setup. The default performance SLAs measure the health of members against well-known internet services, such as FortiGuard, Google Search, and Amazon AWS. Alternatively, you can create your own entry and choose whether you want to monitor the member actively or passively. In active monitoring, the health of the member is checked by periodically sending probes from the member to one or two servers that act as beacons. In passive monitoring, the health of a member is determined based on the traffic passing through the member.

The example on this slide shows a user-defined performance SLA named **Level3_DNS**, which monitors the health and performance of members by sending pings to the well-known **4.2.2.1** and **4.2.2.2** DNS servers.

DO NOT REPRINT
© FORTINET

Performance SLAs Configuration Overview

- Multiple sections for each performance SLA



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

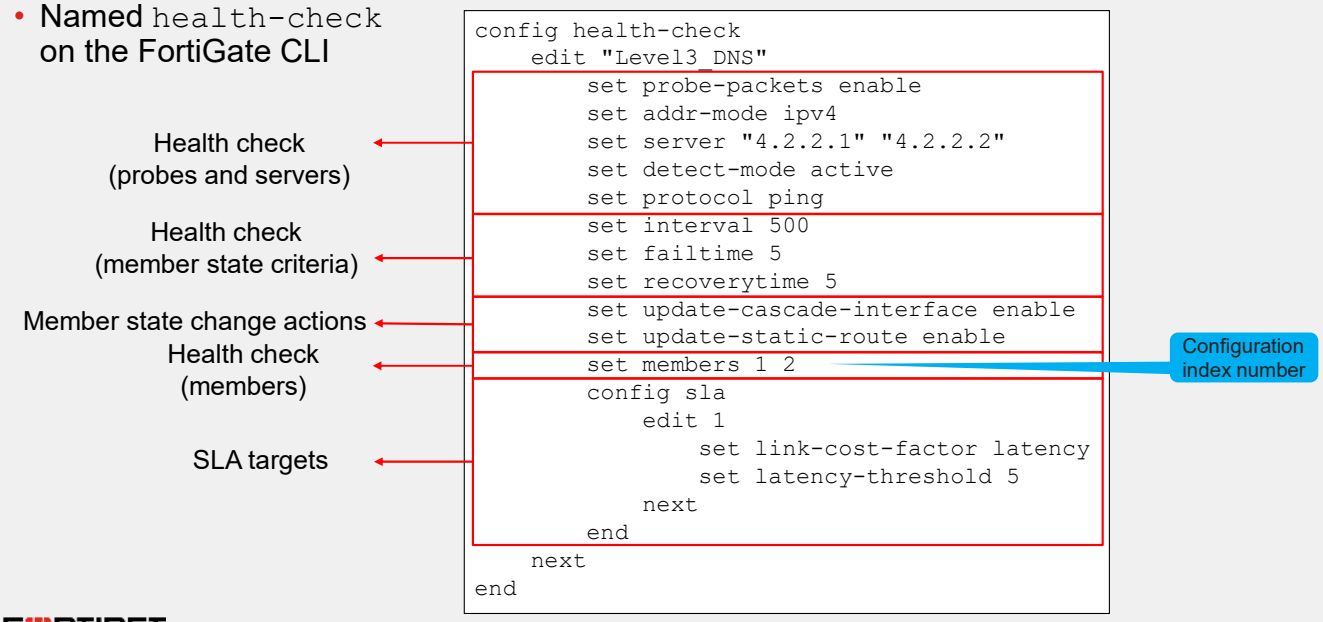
The FortiManager Performance SLA configuration page can be divided into the following sections:

- Health check: Divided into two parts, this is where you configure how the member health check is performed, which members it applies to, and the criteria for alive and dead members.
- SLA targets: Define the performance requirements of alive members so they can be eligible for traffic steering. SLA targets are used by some SD-WAN rule strategies like *lowest cost* or *Maximize Bandwidth*.
- Member state change actions: When there is a change in the state of a member (alive or dead), you can instruct FortiGate to update the static routes that match the gateway used by the member. You can also instruct FortiGate to bring down or bring up one or more alert interfaces based on the state of all members.
- Advanced options: This is where you configure additional performance SLA settings.

DO NOT REPRINT
© FORTINET

CLI Configuration for Performance SLAs

- Named `health-check` on the FortiGate CLI



FORTINET Training Institute

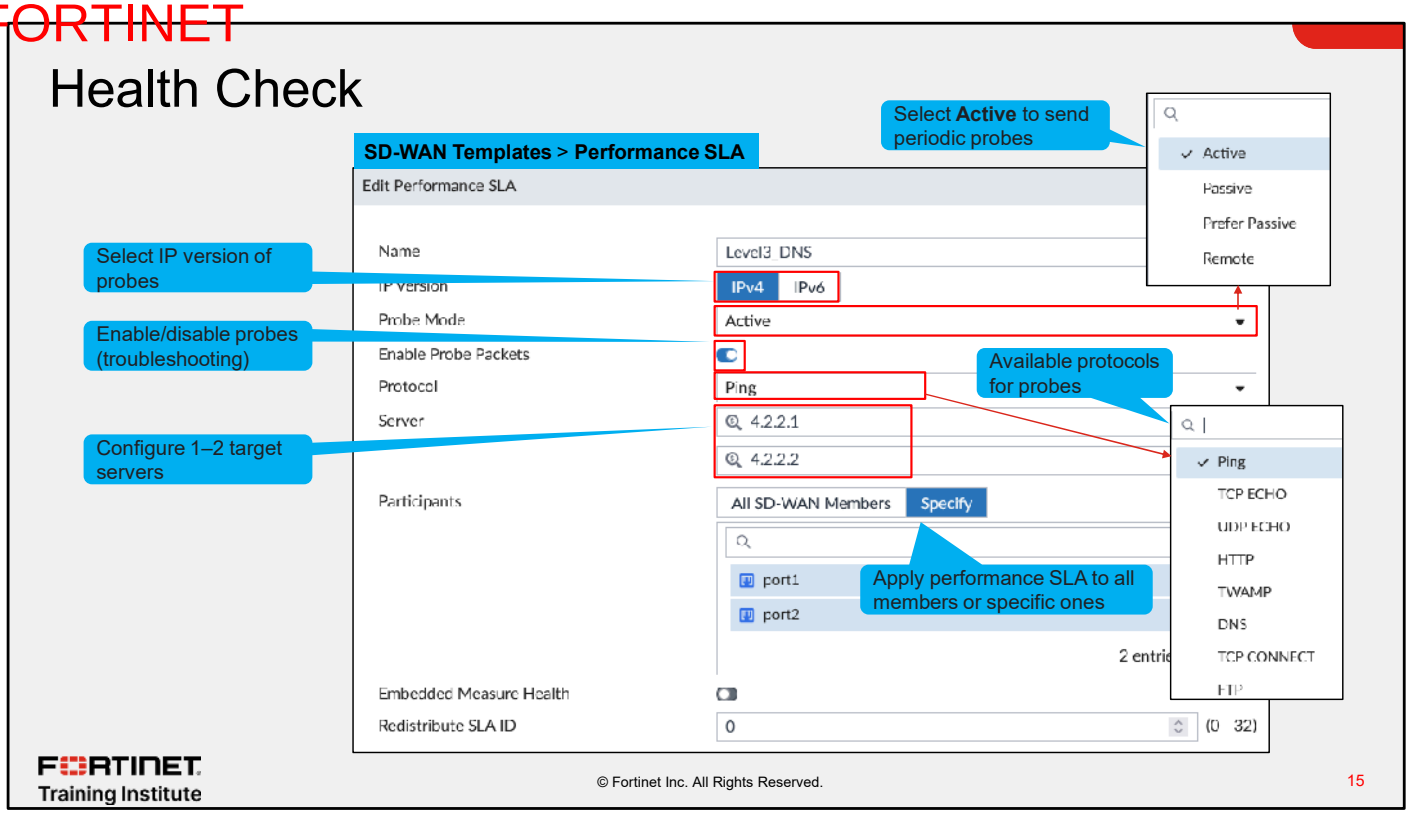
© Fortinet Inc. All Rights Reserved.

14

This slide shows the equivalent FortiOS CLI configuration to the FortiManager example configuration shown on the previous slide. Note that performance SLA is called health-check on CLI configurations.

Note also that the members are referenced using their configuration index number, and not their names. The same is done for most SD-WAN settings and diagnostics. For this reason, it's very useful to identify the configuration index numbers of your members for troubleshooting purposes.

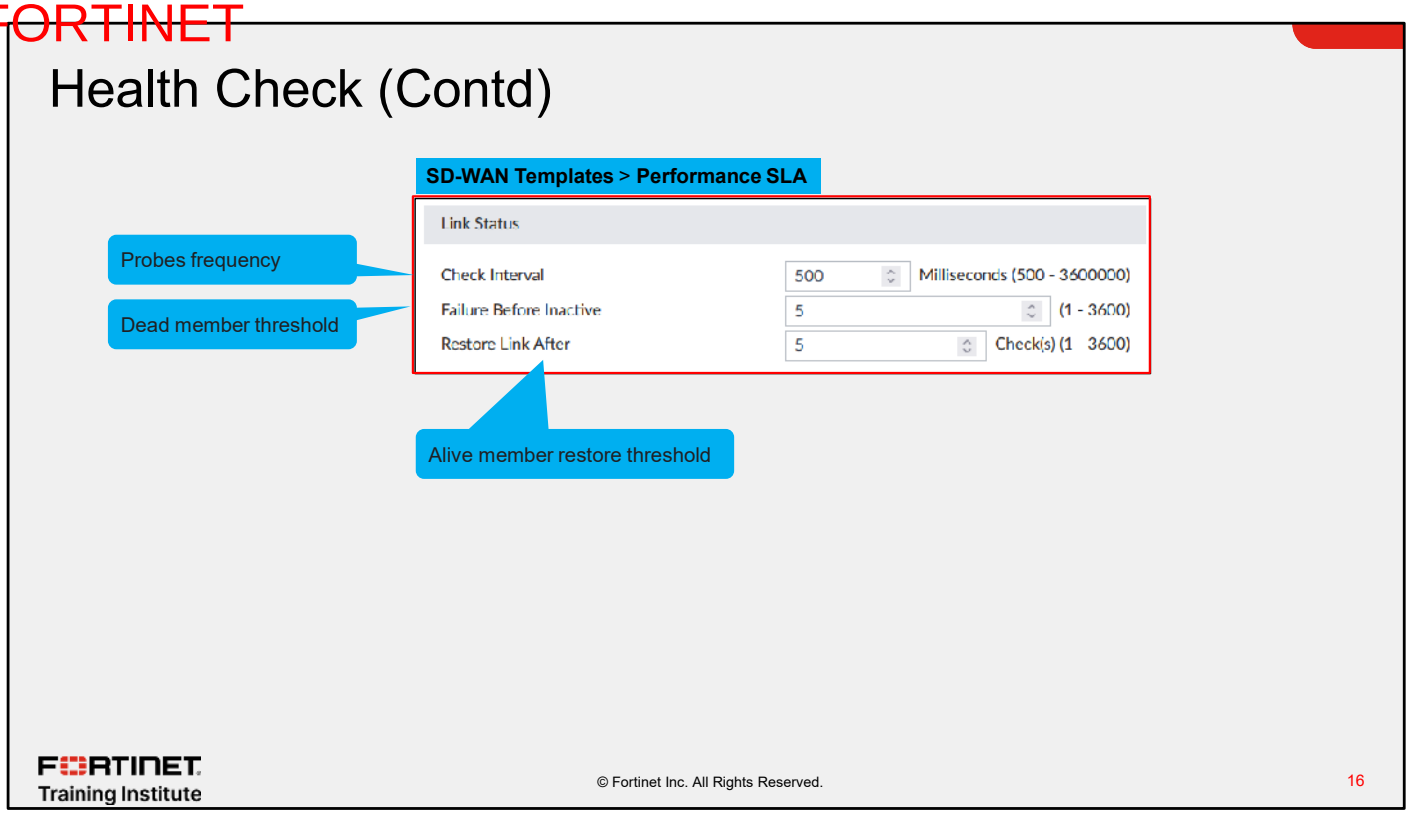
DO NOT REPRINT
© FORTINET



The first part of the health check configuration provides the following settings:

- **IP Version**: Select the IP version to use for probes. Both IPv4 and IPv6 are supported.
- **Probe Mode**: FortiGate can monitor members actively, passively, using a mix of both (*Prefer Passive*) or using information received from remote device. In **Active** mode, FortiGate sends periodic probes through the member to monitor its health and performance. In **Passive** mode, FortiGate monitors the actual network traffic flowing through the member to determine its performance. In **Prefer Passive** mode, FortiGate uses passive mode first, and then switches to active mode if the member has been idle for three minutes. With **Remote** mode, FortiGate uses SLA information received from the remote device. You will learn more about detection modes in this lesson.
- **Protocol**: Select the protocol to use for probes. Multiple protocols are supported. Although ping is often used, in some cases, it is more convenient to use a different protocol.
- **Server**: If using **Active** or **Prefer Passive** modes, configure up to two servers to send the probes to. It is best practice to configure two servers to guard against false positives caused by the server being at fault, and not the link.
- **Participants**: Select if you want to use the health check to monitor all SD-WAN members or specific ones.
- **Enable Probe Packets**: If using **Active** or **Prefer Passive** modes, disable or enable the use of probes. When you disable probes, FortiGate continues processing SD-WAN traffic using the last known measured metrics for members but stops monitoring them, which means that FortiGate won't detect new changes in the quality of links. For this reason, you should disable probes for troubleshooting purposes only.

DO NOT REPRINT
© FORTINET



The second part of the health-check configuration provides the following settings:

- **Interval**: Defines how often the probes are sent to the configured target servers. By default, probes are sent every 500 msec.
- **Failure Before Inactive**: Defines the number of consecutive failed probes to detect before a member is determined dead. By default, a member is considered dead after five consecutive failed probes.
- **Restore Link After**: Defines the number of consecutive successful probes to detect before a member changes from the dead state to the alive state. Initially, all members are considered alive. If the member is later detected dead, by default, it becomes alive again after five consecutive successful probes.

DO NOT REPRINT
© FORTINET

Active Monitoring

- Periodic probes sent to target servers to determine a member:
  - State: alive or dead
    - Initial state for members: alive
  - Performance: packet loss, latency, and jitter
    - Used by SD-WAN rules
- Probes:
  - Supported IPv4 protocols:
    - General-purpose
      - Ping, UDP echo, TCP connect, and TCP echo, TWAMP
    - Application-specific:
      - DNS, FTP, and HTTP
  - Supported IPv6 protocols:
    - General-purpose
      - Ping, UDP echo, TCP connect
    - Application-specific:
      - DNS, FTP
  - Default interval: 500 msec
  - Default failure and restore thresholds: 5

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

When you configure a performance SLA to use active monitoring, FortiGate sends probes to determine the state of the member—whether it is alive or dead—and its performance in terms of latency, packet, and jitter.

The protocols supported for probes can be divided into two groups: general-purpose and application-specific. The protocols in the former group can be used to measure the quality of a link regardless which type of application the target server is running. The application-specific protocols enable you to get more accurate performance results for applications and services. For example, you may want to prefer HTTP over ping to measure the performance of a member that is used to steer web traffic. In this case, an HTTP probe would provide more accurate results instead of ping, which could be blocked or rate-limited along the path.

For IPv4 probes, FortiGate supports the protocols listed on this slide. For IPv6 probes, FortiGate supports the same protocols as in IPv4 except TCP echo, HTTP, and Two-Way Active Measurement Protocol (TWAMP). You will learn more about these protocols in this lesson.

Initially, all configured members are assigned the alive state, and then FortiGate starts sending probes at the configured interval—500 msec by default. After that, FortiGate marks a member as dead if the number of consecutive failed probes reach the failure threshold—five probes by default. A dead member is determined alive again if the number of consecutive successful probes reach the restore threshold—also five probes by default.

The probes are also used to measure the member performance. FortiGate measures the packet loss, latency, and jitter based on requests and replies sent and received for probes. The measured performance can then be used in SD-WAN rules to identify members that meet the required performance of applications and services.

DO NOT REPRINT
© FORTINET

Active Monitoring—General-Purpose Protocols

- Ping
  - Sends ICMP echo requests and waits for ICMP echo replies

```
config system sdwan
    config health-check
        edit "Ping"
            set server "192.0.2.1" "203.0.113.1"
            set protocol ping
            set members 1 2
        next
    end
end
```

- TCP echo and UDP echo
  - Sends TCP/UDP requests on port 7
  - Any data received by the server is sent back
  - Enable TCP Echo on server side

```
config system sdwan
    config health-check
        edit "TCP_Echo"
            set server "192.0.2.1" "203.0.113.1"
            set protocol tcp-echo
            set port 0
            set members 1 2
        next
    . . .
```

Set to 0 to use default port (port 7); (default = 0)

```
config system sdwan
    config health-check
        edit "UDP_Echo"
            set server "192.0.2.1" "203.0.113.1"
            set protocol udp-echo
            set port 0
            set members 1 2
        next
    . . .
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

18

This slide describes the general-purpose protocols supported by SD-WAN when using active monitoring.

Ping is the most used network monitoring protocol because it is supported by virtually all network devices. When you use ping, FortiGate sends ICMP echo requests to the configured target servers and waits for the respective ICMP echo replies. Because some ISPs and content providers block or limit ICMP traffic on their network, you may want to switch to TCP echo, UDP echo, or TWAMP.

When you use TCP echo and UDP echo, FortiGate sends periodic packets to the configured target servers, which are listening for connections on port 7 for both TCP and UDP. Upon reception of the packets, the server sends back an identical copy of the data it received from FortiGate.

This slide shows a FortiOS CLI configuration example of performance SLAs using ping, TCP echo, and UDP echo protocols. Note that you can configure a custom port for TCP echo and UDP echo, in case you don't want to use the default port—port 7.

DO NOT REPRINT
© FORTINET

Active Monitoring—General-Purpose Protocols (Contd)

- TWAMP
  - Client-side implementation
  - Most accurate protocol
  - Two sessions:
    - Control: TCP 862 by default
      - Not seen if authentication is disabled (default=disabled)
    - Test: UDP 862 by default

```
config system sdwan
    config health-check
        edit "TWAMP"
            set server "100.64.1.1"
            set protocol twamp
            set port 0
            set security-mode authentication
            set password fortinet
        next
    end
end
```

Set to 0 to use default port (port 862); (default = 0)

- TCP connect
  - TCP connections using a custom port
    - Check connectivity of any TCP application

```
config system sdwan
    config health-check
        edit "TCP_Connect"
            set server "100.64.1.1"
            set protocol tcp-connect
            set port 22
            set quality-measured-method half-open | half-close
        next
    end
end
```

Measure latency based on connection setup

Measure latency based on connection teardown

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

TWAMP—RFC 5357— is the most accurate protocol among the five. SD-WAN uses the client-side implementation of TWAMP. There are two sessions used in TWAMP: control and test. The former is used to authenticate the endpoints, and the latter to exchange packets used to measure the performance. Note that if authentication is disabled—it is disabled by default—FortiGate generates the test session only. SD-WAN uses port 862 as default port for both control and test sessions, but you can configure a different port, as shown on this slide.

TCP connect works by initiating TCP connections to the target servers using a custom port. TCP connect enables you to test the connectivity of any TCP application running on the target servers by monitoring packets exchanged for TCP connection setup and teardown. That is, when you set `quality-measured-method` to `half-open`, FortiGate determines the latency based on the round-trip time (RTT) between the TCP SYN it sends and the TCP SYN-ACK that it receives. If you set `quality-measured-method` to `half-close`, FortiGate determines the latency based on the RTT between the TCP FIN it sends and the TCP FIN-ACK that it receives.

DO NOT REPRINT
© FORTINET

FortiGate as TWAMP Server

- Configure server-side for TWAMP:

```
config system probe-response
    set port 862
    set mode twamp
    set security-mode authentication
    set password fortinet
end
config system interface
    edit "port1"
        append allowaccess probe-response
    next
end
```

Default port: 8008

Enable probe-response access

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

You can also configure FortiGate as a TWAMP server in your network. You can then configure the performance SLAs that use TWAMP to point to the FortiGate device acting as TWAMP server.

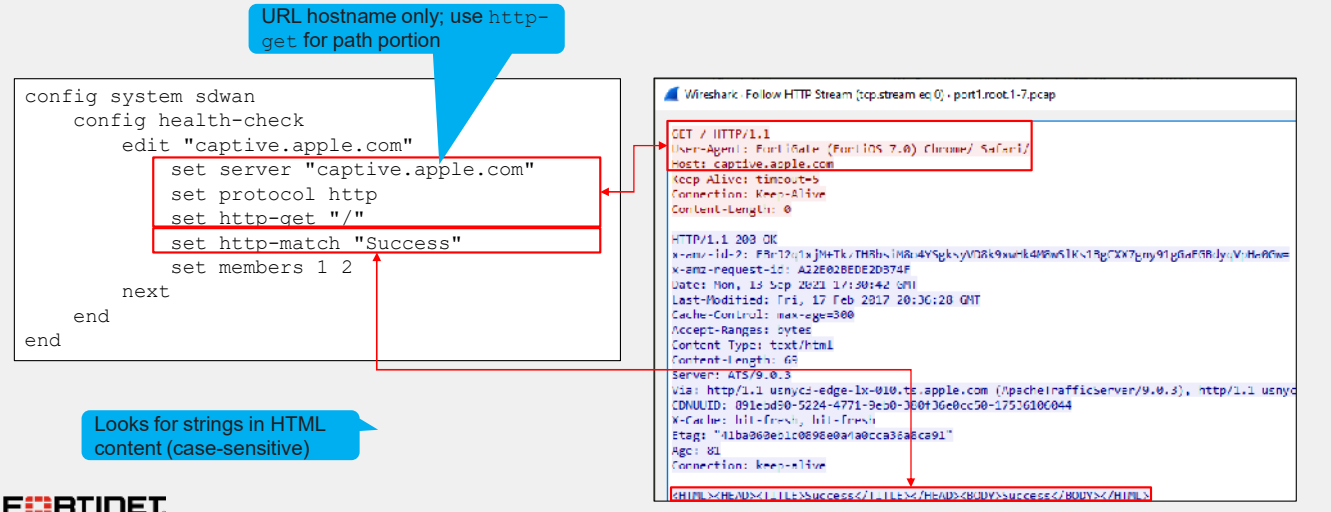
This slide shows a CLI configuration example of a FortiGate TWAMP server that has authentication enabled. To configure a TWAMP server on FortiGate, you configure a probe response service. Note that unlike the client-side implementation used in performance SLAs, the server-side implementation uses a different default port for TWAMP—port 8008.

Also, make sure to enable `probe-response` access on the interface that listens for TWAMP requests.

DO NOT REPRINT
© FORTINET

Active Monitoring—Application-Specific Protocols

- HTTP
  - Sends an HTTP GET request and waits for response
  - Optionally, checks if the response contains the configured string



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

SD-WAN also supports HTTP, DNS, and FTP as probe protocols. These protocols enable you to monitor the health and performance of members using web, DNS, and FTP servers as target servers.

When you configure HTTP as the protocol, FortiGate sends periodic HTTP GET requests to the target server, and then waits for a response. Optionally, you can configure FortiGate to check if the response contains a specific string in the HTML content.

This slide shows a CLI configuration example of an HTTP probe to `captive.portal.com`, and the HTTP stream—as decoded by Wireshark—for a packet capture containing the resulting FortiGate HTTP GET and the server response packets. Note that FortiGate performs a case-sensitive search for the string set as `http-match` in the HTML content received from the server. Also note that you indicate the path portion of the target URL in the `http-get` setting, not in the `server` setting.

For the string to match, you should configure a string that is included in the HTTP response of the web server when the service is available and properly working.

DO NOT REPRINT
© FORTINET

Active Monitoring—Application-Specific Protocols (Contd)

- DNS
  - Sends a DNS A-record query and waits for response
  - Optionally, checks if the response contains the configured IP address

```
config system sdwan
    config health-check
        edit "Google_DNS"
            set system-dns disable
            set server "8.8.8.8"
            set dns-request-domain "fortinet.com"
            set dns-match-ip 54.186.80.150
            set members 1 2
        next
    end
end
```

default requested domain is `example.com`

Source	Destination	Protocol	Length	Info
192.2.0.1	8.8.8.8	DNS	72	Standard query 0xaf00 A fortinet.com
8.8.8.8	192.2.0.1	DNS	104	Standard query response 0xaf00 A fortinet.com A 54.70.126.22 A 54.186.80.150

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

When you configure DNS as the probe protocol, FortiGate sends periodic DNS A-record queries to the configured DNS server. If no domain name is configured, FortiGate queries `example.com`. Instead of configuring a target DNS server, you can also send the queries to the system DNS servers used by FortiGate by enabling the `system-dns` option.

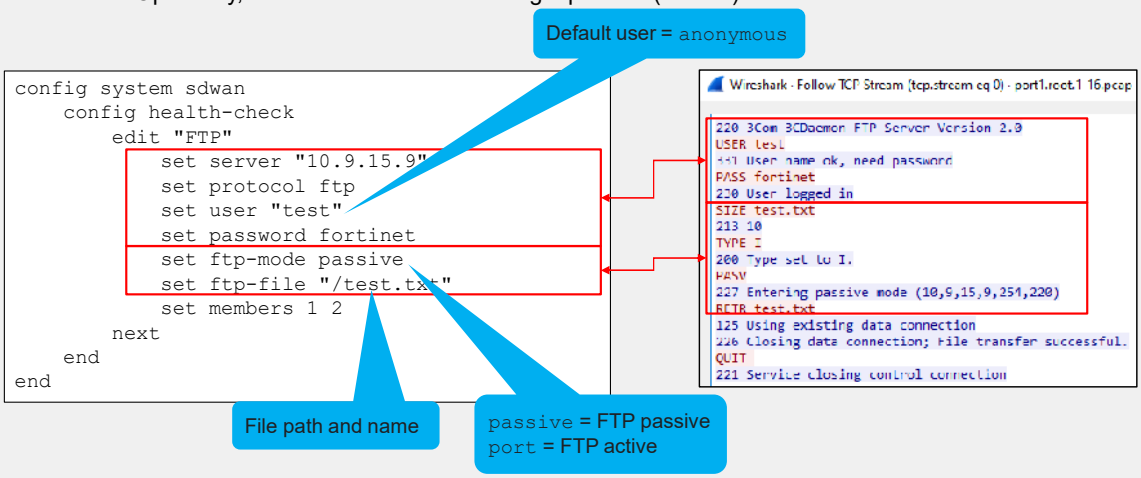
After sending the DNS query, FortiGate waits for the DNS response and optionally, checks if the IP address set as `dns-match-ip` is included in the list of resolved IP addresses.

This slide shows a CLI configuration example of a DNS probe that queries `fortinet.com` against Google's DNS server `8.8.8.8`. FortiGate also checks that the DNS response includes `54.186.80.150` in the list of resolved IP addresses.

DO NOT REPRINT
© FORTINET

Active Monitoring—Application-Specific Protocols (Contd)

- FTP
  - Connects to an FTP server and:
    - Logs in using the configured credentials
    - Optionally, downloads a file either using a passive (default) or active data connection



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

When you configure FTP as the probe protocol, FortiGate periodically connects to the configured FTP server, logs in with the configured credentials, and optionally downloads a file using either a passive or active data connection.

If you don't configure a username, FortiGate uses `anonymous`. In addition, if you specify a file to download, the file is retrieved using a passive connection by default. If you want to use an FTP active connection to download the file, then set `ftp-mode` to `port`.

This slide shows an example of an FTP passive configuration. FortiGate logs in to the FTP server with the username `test` and password `fortinet`, and then downloads the file `test.txt` using a passive data connection. This slide also shows the FTP stream—as decoded by Wireshark—for a packet capture of the resulting FortiGate FTP connection.

DO NOT REPRINT
© FORTINET

Passive Monitoring

- Probe-free performance monitoring
  - Performance based on member traffic
- Benefits
  - More accurate measuring
  - Simplified configuration
  - Reduced network traffic
- Limitations:
  - Only TCP traffic is measured
    - Latency is based on RTT during TCP setup and teardown
    - Jitter and packet loss are based on TCP headers
  - No member state detection
    - Members are always alive
  - Hardware acceleration is disabled

- CLI configuration
  - Performance SLA:

```
config system sdwan
    config health-check
        edit "Passive"
            set detect-mode passive
            set members 3 4
        next
    end
end
```

  - Firewall policy:

```
config firewall policy
    edit 5
        set passive-wan-health-measurement enable
    next
end
```

```
# show full firewall policy 5 | grep auto
        set auto-asic-offload disable
```

Hardware acceleration is automatically disabled

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

Active monitoring requires you to configure target servers and protocols for each probe, which can sometimes result in complex SD-WAN configurations. In addition, the more active health checks and members you configure, the more monitoring traffic is generated in the network.

Passive monitoring simplifies configuration and reduces network traffic by using a probe-free approach for measuring the performance of members. FortiGate measures packet loss, latency, and jitter based on TCP traffic sent and received through a member. Passive monitoring is a more accurate approach than active monitoring because you measure the member performance based on the actual traffic passing through the member, instead of measuring it based on probes that may be unrelated to the application you want to effectively steer using SD-WAN. For example, configuring a ping probe to monitor the health of a member that is used to steer web traffic won't provide more accurate performance metrics than the ones obtained by monitoring the web traffic itself.

In passive monitoring, latency is calculated based on the RTT of TCP connection setup and teardown, while jitter and packet loss are calculated based on TCP header information.

Note that passive monitoring doesn't detect dead members. That is, members are always alive. Also, note that hardware acceleration is disabled on traffic subject to passive monitoring.

This slide shows a CLI configuration example for passive monitoring. Note that, in addition to configuring the health check in passive mode (`set detect-mode passive`), you must also enable `passive-wan-health-measurement` on the firewall policies that accept traffic for the monitoring member. When you enable `passive-wan-health-measurement` on a policy, `auto-asic-offload` is automatically disabled.

DO NOT REPRINT
© FORTINET

Per-Application Passive Monitoring

- Differentiate and collect metrics for apps defined in rules
  - If multiple apps are defined, an average is calculated
- Benefits
  - Most accurate monitoring method
- CLI configuration:

```
config system sdwan
    config health-check
        edit "Passive"
            set detect-mode passive
            set members 3 4
        next
    end
end
config firewall policy
    edit 5
        set passive-wan-health-measurement enable
    next
end
```

- CLI configuration (contd):

```
config system sdwan
    config service
        edit 1
            set name "GoToMeeting-Salesforce"
            set src "all"
            set internet-service enable
            set internet-service-app-ctrl 16354 16920
            set health-check "Passive"
            set priority-member 3 4
            set passive-measurement enable
        next
    end
end
```

Application IDs:
16354: GoToMeeting
16920: Salesforce

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

When you use passive monitoring, you get more accurate member metrics because the metrics are calculated based on the actual traffic passing through the members. However, if multiple applications use the same member, then the member metrics are a result of measuring the traffic for multiple applications, including irrelevant applications. This lack of granularity may affect the steering decisions made by FortiGate for SD-WAN rules configured for specific applications.

Per-application passive monitoring instructs FortiGate to measure the member quality based on the performance of applications selected in SD-WAN rules. For SD-WAN rules, you select applications by choosing the respective internet service or application name. FortiGate then differentiates and collects the metrics for each internet service and application. If you define multiple applications in a rule, then FortiGate determines the member quality by averaging the metrics of all configured internet services and applications in the rule.

The configuration required for per-application passive monitoring is the same as regular passive monitoring, except that you must enable `passive-measurement` in one ore more SD-WAN rules that match traffic based on the detected application. These type of rules enable you to perform application steering, and you will learn more about it in another lesson.

This slide shows a CLI configuration example for per-application passive monitoring. The configuration instructs FortiGate to passively measure the quality of members 3 and 4 by averaging the metrics measured for GoToMeeting and Salesforce traffic passing through members 3 and 4.

DO NOT REPRINT
© FORTINET

Prefer Passive Monitoring

- Mix of passive and active monitoring
  - Passive when there is TCP traffic
  - Active when no TCP traffic has been detected for 3 minutes (hard-coded)
- Dead members detected during active monitoring

- CLI Configuration
  - Performance SLA:

```
config system sdwan
    config health-check
        edit "T_INET_0_Passive"
            set server "10.1.0.7"
            set protocol ping
            set detect-mode prefer-passive
            set members 3
        next
    end
end
```

  - Firewall policy:

```
config firewall policy
    edit 5
        set passive-wan-health-measurement enable
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

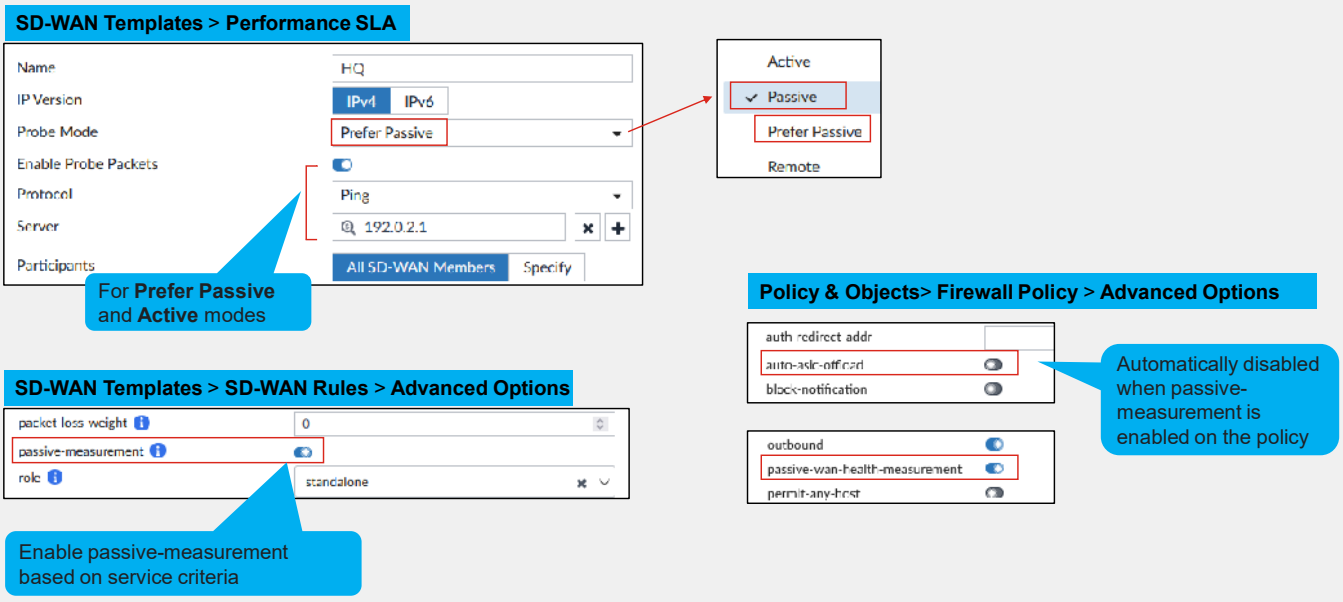
An alternative to active and passive monitoring is prefer passive monitoring.

In prefer passive monitoring, FortiGate uses passive monitoring if there is TCP traffic through the member. If no TCP traffic is detected for three minutes, then FortiGate starts using the configured probe to actively monitor the state and performance of the member. If FortiGate later detects TCP traffic through the member, it switches immediately back to passive monitoring. Note that changes in member state are detected only when FortiGate is doing active monitoring. Also note that the 3-minute TCP traffic idle timer is hard-coded and can't be changed.

This slide shows a CLI configuration example for prefer passive monitoring. FortiGate starts pinging `10.1.0.7` only after no TCP traffic has been seen through the member for three minutes. Like in passive monitoring, you must enable `passive-wan-health-measurement` on the firewall policies that accept traffic for the monitoring member. Consequently, this also results in `auto-asic-offload` being automatically disabled on the respective policies.

DO NOT REPRINT
© FORTINET

Passive Monitoring—GUI Configuration



This slide shows the GUI menus used for the configuration of passive performance SLA monitoring.

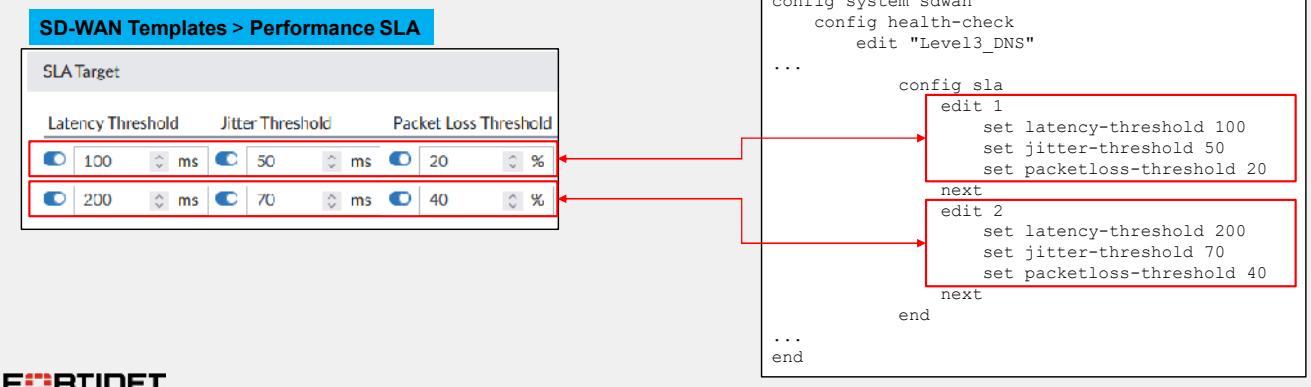
After you select **Passive** or **Prefer Passive** for a performance SLA, you must activate **passive-wan-health-measurement** for at least one policy with SD-WAN zone as the source or destination interface. This option will automatically disable **auto-asic-offload** for the policy.

If you want to specify the services used for passive measurement, you must activate **passive-measurement** on the corresponding SD-WAN rule.

DO NOT REPRINT
© FORTINET

SLA Targets

- Define member performance requirements
- Used by **Lowest Cost (SLA)** and **Maximize Bandwidth (SLA)** rules
- Member must meet SLA target for being eligible for traffic steering
- Configure one or more SLA targets per performance SLA
  - Use same health check for multiple applications



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

When you configure a performance SLA, you can optionally configure SLA targets. SLA targets define the performance requirements that alive members must meet for being eligible for traffic steering. SLA targets are required by SD-WAN rules that use **Lowest Cost (SLA)** and **Maximize Bandwidth (SLA)** as strategies. In addition to verifying that members are alive, these strategies also check the quality of the members.

Using SLA targets enable you to fine-tune SD-WAN rules for specific applications. For example, you may want to configure an SD-WAN rule that steers voice traffic through members whose latency and jitter don't exceed 300 msec and 30 msec, respectively. Otherwise, the quality of voice during calls will likely be poor.

You can configure one or more SLA targets per performance SLA. This enables you to have different performance requirements for the same monitoring mode, and then configure the SD-WAN rule to use the best suited SLA target for a given application. For example, you can monitor the performance of a member using a particular mode—active, passive, or prefer passive—and then define two SLA targets, one with stricter performance requirements than the other. You can then configure an SD-WAN rule that steers sensitive traffic to use the stricter SLA target, and another SD-WAN rule that steers non-sensitive traffic to use the more lenient SLA target.

This slide shows a configuration example using SD-WAN templates on FortiManager, and the corresponding CLI settings on FortiGate.

DO NOT REPRINT
© FORTINET

MOS—Mean Opinion Score

- Link quality for voice traffic
  - Combines latency, jitter, and packet loss
  - Formula per codec: G.711, G.729, G.722
  - Provide a score between 1 and 5
- Benefits:
  - Optimize link selection for voice traffic
  - Better user experience
- Limitations:
  - CLI configuration only
  - SD-WAN rule mode **Lowest Cost (SLA)** only

```
config system sdwan
    config health-check
        edit "Voice_health"
           set server "10.200.99.1"
           set mos-codec g711
           set members 3 4
           config sla
              edit 1
                 set link-cost-factor mos
                 set mos-threshold "3.6"
              next
     . . .
    config service
        edit 1
           set name "MOS_traffic_steering"
           set mode sla
           set dst "Corp-net"
           set src "LAN-net"
           config sla
              edit "Voice_health"
                 set id 1
              next
           end
           set priority-members 1
        next
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

The perceived voice quality is sensible to multiple factors, latency, jitter, and packet loss, with interaction between them. The MOS is a method of measuring link quality using a formula that takes into account latency, jitter, packet loss, and the codec to produce a score from zero to five (0 - 5). You can use this score to steer voice to most appropriate link for perceived voice quality.

Available codecs are G.711, G.729, and G.722 (default G.711)

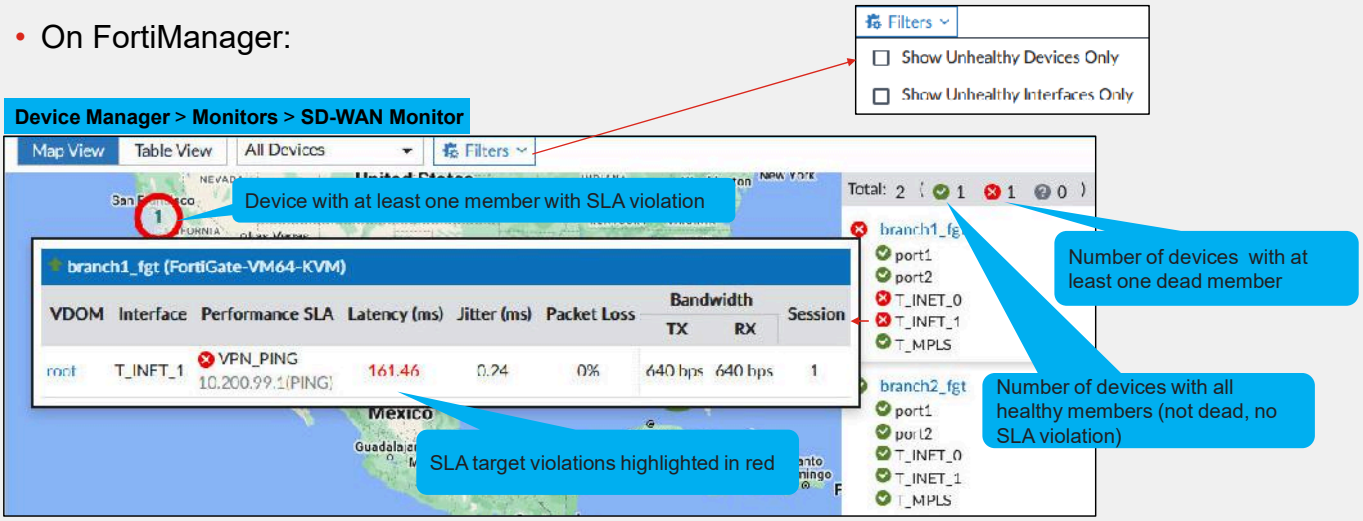
You can define performance SLA with MOS criteria using CLI commands as shown on this slide. You can use MOS when service is **Lowest Cost (SLA)**, it is not supported with service mode **Priority**.

MOS threshold indicates the minimum MOS score for the SLA to pass. It is defined as a one-digit decimal value between 1.0 and 5.0, with a default 3.6.

DO NOT REPRINT
© FORTINET

Viewing Member State and Performance

- On FortiManager:



FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

You can use FortiManager and the FortiGate GUI to view the state and performance of members.

Both FortiManager and the FortiGate GUI indicate whether the member is alive or dead. In FortiManager, alive members are shown with a green icon, and members that miss at least one SLA target are shown with a red icon.

When you hover over a member, you can view the measured performance of alive members. If you configure SLA targets, both FortiManager and FortiGate highlight the members that don't meet the configured SLA targets.

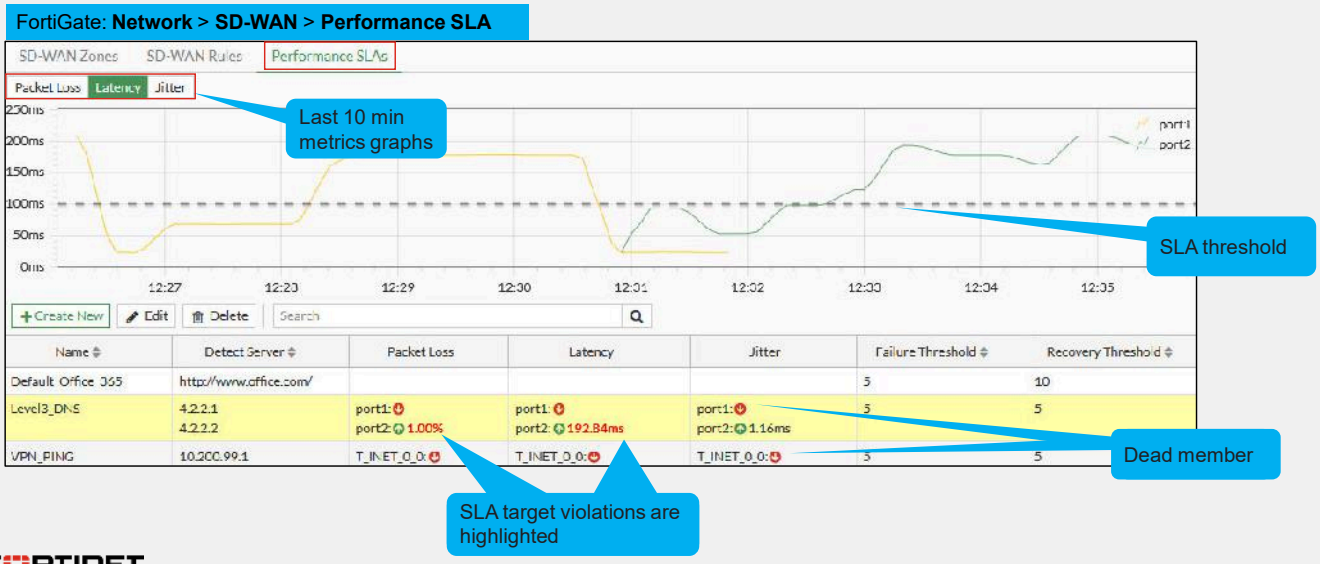
In the example shown on this slide, the **VPN_PING** performance SLA reports that **T_INET_1** doesn't meet the configured SLA target for latency criteria. Note that the SLA target configuration is not shown on this slide.

From the map view you can filter to see only healthy or only unhealthy devices.

DO NOT REPRINT
© FORTINET

Viewing Member State and Performance (Contd)

- On the FortiGate GUI:



FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

The FortiGate GUI shows state and performance of members with some additional details.

As you do on FortiManager, you will quickly see alive and dead members. The FortiGate GUI shows alive members with a green up arrow icon, and dead members with a red down arrow icon.
For a missed SLA target, the FortiGate highlight the impacted metric in red.

FortiOS stores the last 10 minutes of SLA metrics for each monitoring member. The FortiGate GUI displays the graph for the selected performance SLA and criteria.

In the example shown on this slide, the **Level3_DNS** performance SLA is selected and reports that **port2** is alive and **port1** is dead. The graph shows latency for both monitored interfaces over the past 10 minutes.

DO NOT REPRINT
© FORTINET

Viewing Member State and Performance on the FortiGate CLI

- Performance SLA:

```
# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(alive), packet-loss(9.000%) latency(42.351), jitter(0.270), mos(4.377), bandwidth-up(10217),
bandwidth-dw(10016), bandwidth-bi(20233) sla_map=0x1
Seq(2 port2): state(dead), packet-loss(100.000%) sla_map=0x0
Health Check(HQ):
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(0.000), jitter(0.000), mos(4.404), bandwidth-
up(10240), bandwidth-dw(10240), bandwidth-bi(20480) sla_map=0x0
```

Configuration index number

- Link monitor status:

```
# diagnose sys link-monitor interface port1
Interface(port1): state(up, since Tue Nov  8 08:22:49 2022), bandwidth(up:79743bps, down:1094502bps), session
count(IPv4:412, IPv6:0), tx(39720018 bytes), rx(202184681 bytes), latency(21.96), jitter(0.17), packet-
loss(0.00).
```

- Link monitor passive status:

```
# branch1_fgt # diagnose sys link-monitor-passive admin list
  T_INET_0( 19) | service=0x00000000 | latency=70.0  14:21:35 | jitter=0.0  14:21:47 | pktloss=0.0 %  14:21:47
```

Last time for criteria check

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

You can also display the state and performance of members on the FortiGate CLI, which is often more convenient for troubleshooting purposes.

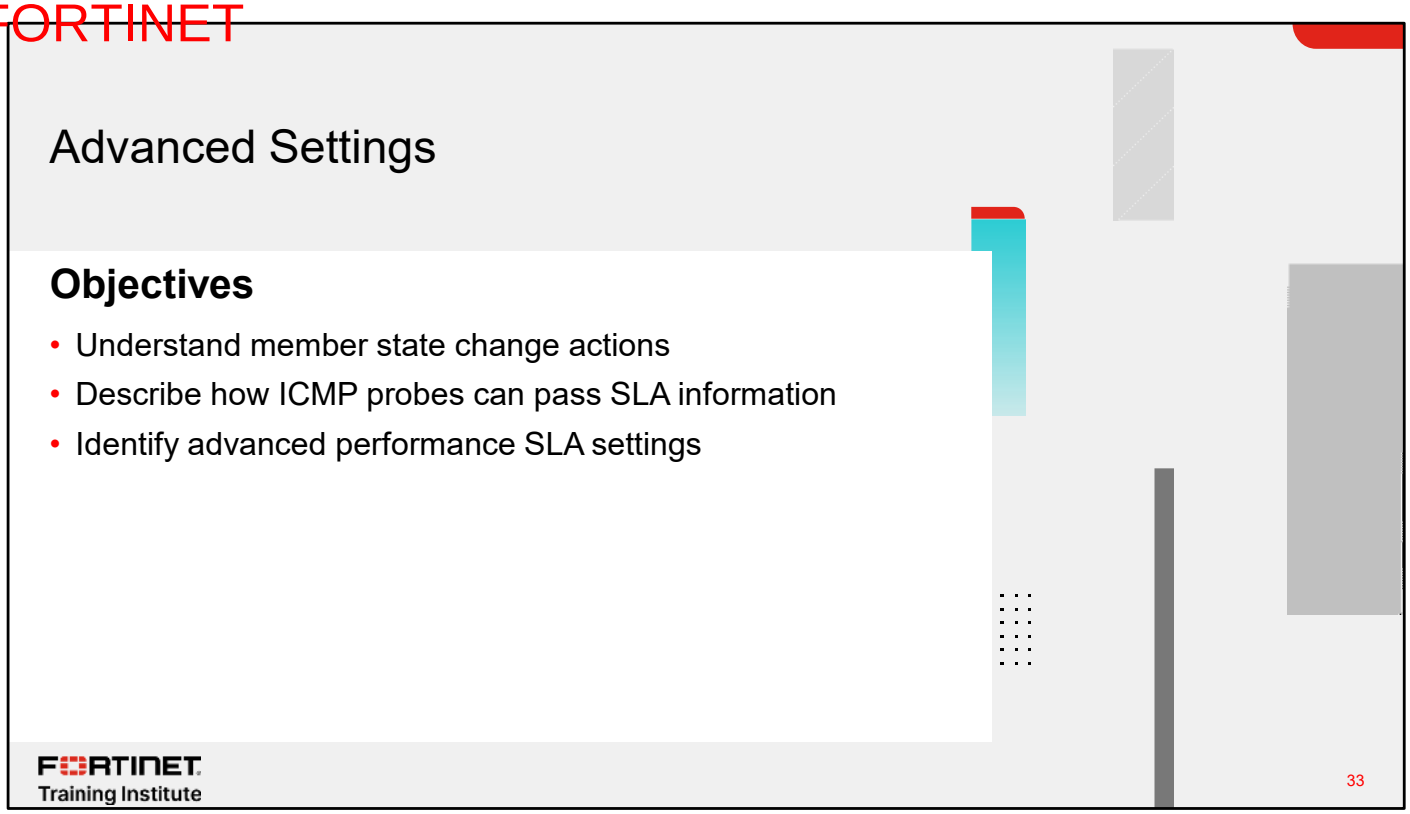
The `diagnose sys sdwan health-check status` command displays the member status information per performance SLA. The output includes both the member state and the measured performance.

The performance SLAs rely on the FortiOS link monitor process (`lnkmt`) to monitor the state and performance of members. For this reason, you can run the `diagnose sys link-monitor interface` command to display member status information that is similar to what is displayed by the `diagnose sys sdwan health-check status` command.

In passive or prefer passive mode, the link monitor passive process (`lnkmt_passive`) is responsible for gathering the data and preparing the reports used by the link monitor process to calculate packet loss, latency, and jitter. For this reason, you can run the `diagnose sys link-monitor-passive admin list` command to display the passive data collected.

This slide shows sample outputs for the three commands discussed on this slide. Note that the `Level3_DNS` performance SLA monitors `port1` and `port2` actively, while `HQ` monitors `T_INET_0` passively. Also note that `port2` is dead.

DO NOT REPRINT
© FORTINET

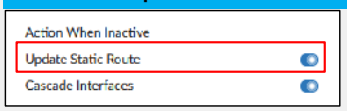


After completing this section, you should be able to achieve the objectives shown on this slide.

DO NOT REPRINT
© FORTINET

Member State Change Actions—Update Static Route

SD-WAN Templates > Performance SLA



- Update status of static routes matching the gateway of members that change their state
- Static routes of alive members are installed in the routing table

```
# diagnose sys sdwan member
Member(1): interface: port1, flags=0x0 , gateway: 192.2.0.2, priority: 1 1024, weight: 0
Member(2): interface: port2, flags=0x0 , gateway: 192.2.0.10, priority: 1 1024, weight: 0
Member(3): interface: T_INET_0, flags=0xc , gateway: 100.64.1.1, priority: 1 1024, weight: 0

# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(alive), packet-loss(2.000%) latency(22.300), jitter(0.218) ...
Seq(2 port2): state(alive), packet-loss(0.000%) latency(52.406), jitter(0.253) ...
Health Check(HQ):
Seq(3 T_INET_0): state(alive), packet-loss(3.000%) latency(1.891), jitter(0.319)...

# get router info routing-table all
...
S*      0.0.0.0/0 [1/0] via 192.2.0.2, port1, [1/0]
                  [1/0] via 192.2.0.10, port2, [1/0]
...
S       10.0.0.0/8 [10/0] via T_INET_0 tunnel 100.64.1.1, [1/0]
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

When **Update Static Route** is enabled—default setting—, FortiGate updates the status of static routes that match the gateway used by a member that changes its state—from alive to dead, and vice versa. The goal is to prevent traffic from being routed through a dead member.

When the member is alive, the static routes become active, and therefore, are installed in the routing table if they are determined the best routes to the destination. Note that static routes of members configured with automatic gateway detection are also automatically identified by FortiGate.

The example on this slide shows the status of three members: `port1`, `port2`, and `T_INET_0`. The three members are alive, and therefore, the static routes that match the member gateway is assigned an active status. The result is that the static routes are installed in the routing table.

DO NOT REPRINT
© FORTINET

Member State Change Actions—Update Static Route (Contd)

- Static routes of dead members become inactive
  - Inactive routes are removed from the routing table
  - You can view the inactive routes in the routing table database

```
# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(alive), packet-loss(4.000%) latency(70.634), jitter(1.894),..., sla_map=0x0
Seq(2 port2): state(dead), packet-loss(100.000%) sla_map=0x0
Health Check(HQ):
Seq(3 T_INET_0): state(dead), packet-loss(100.000%) sla_map=0x0

# get router info routing-table database
...
S     *> 0.0.0.0/0 [1/0] via 192.2.0.2, port1, [1/0]
       >           [1/0] via 192.2.0.10, port2 inactive, [1/0]
...
S       10.0.0.0/8 [10/0] via T_INET_0 tunnel 100.64.1.1 inactive, [1/0]
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

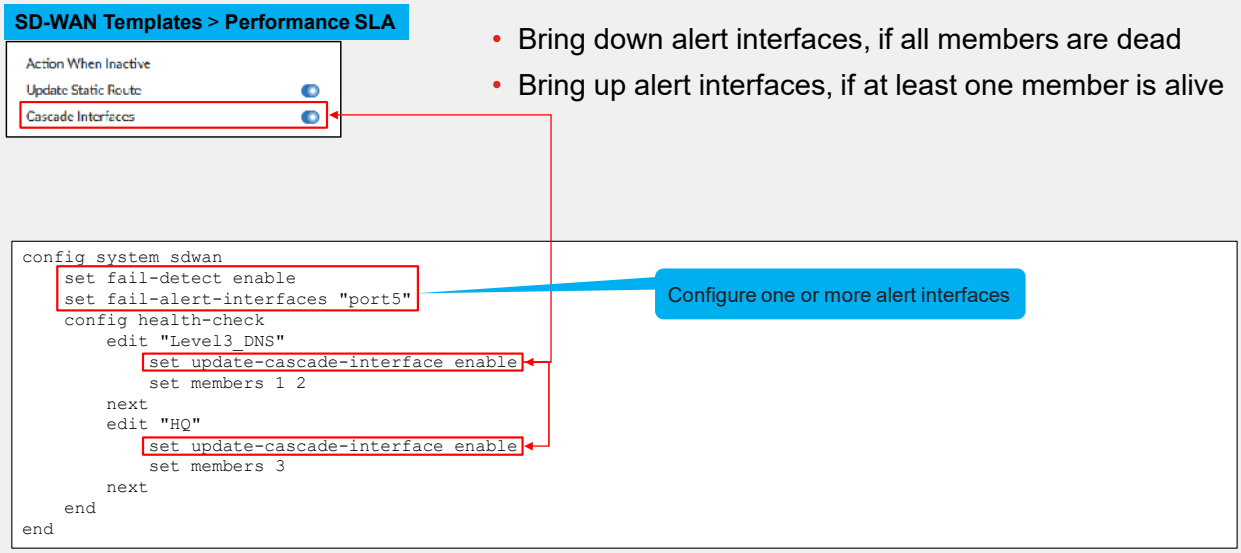
35

If a member is determined dead, the static routes that match the member gateway become inactive. The result is that the routes are not installed in the routing table.

The example on this slide shows `port2` and `T_INET_0` as dead. As a result, the static routes that match the gateways used by `port2` and `T_INET_0` become inactive. For this reason, you can view the impacted static routes in the routing table database only.

DO NOT REPRINT
© FORTINET

Member State Change Actions—Cascade Interfaces



When you enable **Cascade Interfaces** and configure one or more alert interfaces, one of the following events will occur:

- FortiGate brings down the alert interfaces, if all members are dead.
- FortiGate brings up the alert interfaces, if at least one member is alive.

Cascade Interfaces is useful to force the traffic from networks behind the alert interfaces to be routed through a different device, if all SD-WAN members are dead, which could mean that FortiGate is unable to forward traffic to the WAN.

Note that you must manually configure one or more alert interfaces. You should configure alert interfaces for critical interfaces, such as your LAN interfaces, which provide users with access to the WAN, and for which you want WAN traffic to be routed through a different device, if all SD-WAN members are dead. For example, if you are using dynamic routing or Virtual Router Redundancy Protocol (VRRP) on your LAN interface, then bringing down the interface can trigger a routing failover to a backup gateway.

This slide shows an example of the **Cascade Interfaces** option enabled on FortiManager and `port5` configured as the alert interface using the FortiGate CLI.

DO NOT REPRINT
© FORTINET

Member State Change Actions—Cascade Interfaces (Contd)

- Alert interface (`port5`) is up if there is at least one alive member:

```
# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(alive), packet-loss(73.000%) latency(75.520), jitter(2.071), ..., sla_map=0x0
Seq(2 port2): state(dead), packet-loss(100.000%) sla_map=0x0
Health Check(HQ):
Seq(3 T_INET_0): state(dead), packet-loss(99.000%) sla_map=0x0

# diagnose hardware deviceinfo nic port5 | grep "Link\|State"
State:            up
Link:             up
```

- Alert interface is down if all members are dead:

```
# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(dead), packet-loss(11.000%) sla_map=0x0
Seq(2 port2): state(dead), packet-loss(97.000%) sla_map=0x0
Health Check(HQ):
Seq(3 T_INET_0): state(dead), packet-loss(10.000%) sla_map=0x0

# diagnose hardware deviceinfo nic port5 | grep "Link\|State"
State:            down
Link:             down
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

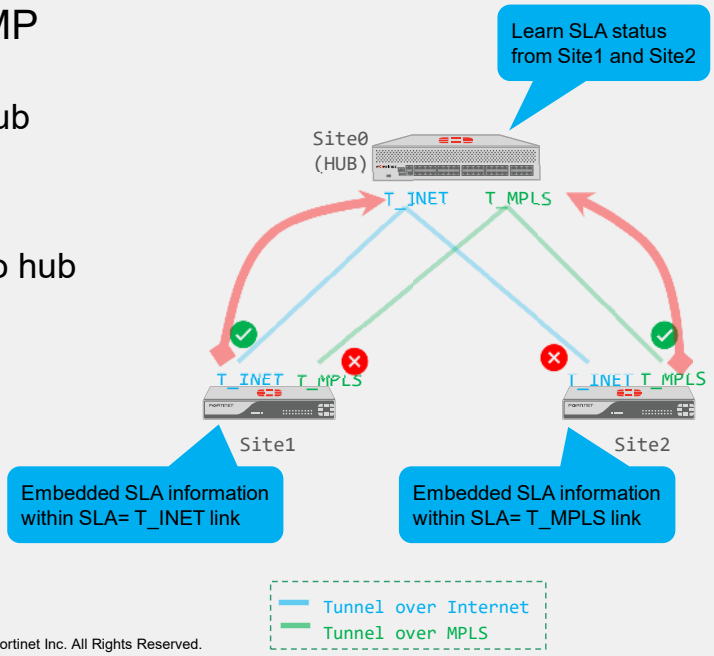
37

This slide shows the effect of **Cascade Interfaces** based on the configuration shown in the previous slide. If there is at least one alive member—`port1` in the example—the alert interface (`port5`) is up. However, if all members are dead, `port5` is brought down.

DO NOT REPRINT
© FORTINET

Inform Remote Device About SLA Status

- Embed SLA information in ICMP probes
  - Spoke provides SLA status to hub
  - Hub learns from remote devices
- Benefits
  - Symmetric routing from spoke to hub and hub to spoke
  - Compatible with static routing
- Limitations:
  - Active probe mode
  - Protocol ping



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

38

In a hub-and-spoke design, if the hub is unaware of the spoke decision-making process, it might route traffic over a link that is different from the preferred spoke link. This can cause asymmetric traffic flow and selection of a link with bad performance. To avoid this, the spoke SD-WAN device can pass SLA information to the hub with SLA information embedded into ICMP probes.

With embedded SD-WAN SLA information in ICMP probes, the spoke communicates the SLA status to the hub, for each overlay, directly through ICMP probes. The hub can use the received SLA status information to apply priorities to the IKE routes, giving routes over overlays that are inside SLAs a lower priority value than routes over overlays that are outside of SLAs.

This method is particularly useful in combination with static routing. In case of BGP routing, you can use this method or a BGP-specific method with the parameter `route-map-out-preferable`. We will cover this in another lesson.

Note that you need to use active probes and protocol ping to inform remote devices about SLA status.

DO NOT REPRINT
© FORTINET

Inform Remote Device About SLA Status (Contd)

- Spoke CLI configuration

```
config system sdwan
    config health-check
        edit "to_hub"
            set server "4.2.2.1" "4.2.2.2"
            set detect-mode active
            set protocol ping
            set embed-measured-health enable
            set members 3 4
            config sla
               edit 1
                set link-cost-factor latency
                set latency-threshold 100
            end
        next
    end
end
```

- Hub CLI configuration

```
config system sdwan
    config health-check
        edit "from_spoke"
            set detect-mode remote
            set members 1 2
            config sla
               edit 1
                set link-cost-factor latency
                set latency-threshold 100
            end

        next
    end
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

39

This slide shows example CLI configurations.

On the spoke, you activate transmission of SLA information in the ICMP probe using the command `set embed-measured-health enable`. On the hub, you use the command `set detect-mode remote` to instruct the device to get SLA information from received SLA probes.

You can use the embedded SLA mechanism with static routing or BGP per overlay topologies. However, note that, to use embedded SLA with BGP, you need to use static IP addressing for tunnels. It is not compatible with the IKE option `mode-cfg`.

DO NOT REPRINT
© FORTINET

Advanced SLA Settings

```
config system sdwan
    config health-check
        edit "Level3_DNS"
            set probe-timeout 500
            set probe-count 30
            set diffservcode 001010
            set vrf 1
            set threshold-warning-packetloss 5
            set threshold-alert-packetloss 10
            set threshold-warning-latency 100
            set threshold-alert-latency 150
            set threshold-warning-jitter 30
            set threshold-alert-jitter 50
next
    end
end
```

Time to wait for probe response (default = 500)

Number of most recent probes to use for latency and jitter calculation (default = 30)

DSCP code to be used by probes (default = 000000)

Warning and alert thresholds for metrics; used only by the FortiGate GUI for visual notification

FortiGate: **Network** > **SD-WAN** > **Performance SLAs**

Warning

Alert

Name	Detect Server	Packet Loss	Latency	Jitter
Level3_DNS	4.2.2.1	port1: 5.00%	port1: 42.26ms	port1: 0.69ms
	4.2.2.2	port2: 0.00%	port2: 177.15ms	port2: 0.65ms

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

40

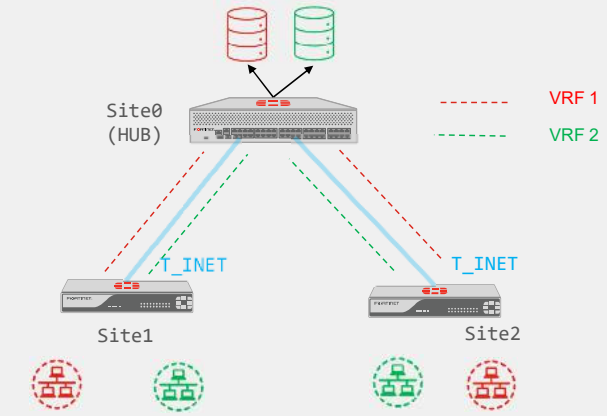
This slide shows some advanced settings that you may want to configure for performance SLAs:

- `probe-timeout`: Defines the time, in milliseconds, to wait for probe responses before the response is considered lost. The default value is `500`.
- `probe-count`: Defines the number of the most recent probes to use to calculate latency and jitter. By default, FortiGate uses the last 30 probes. Note that this setting doesn't apply to packet loss calculation. Packet loss calculation always uses the last 100 probes.
- `diffservcode`: Sets the 6-bit differentiated services code point (DSCP) marking for probes. This is useful when a member is connected to a link that uses DSCP markings to classify and prioritize traffic. For example, some MPLS links prioritize packets with DSCP markings assigned for voice services. In this scenario, setting a DSCP probe enables you get more accurate member performance results for voice traffic. The default value is `000000`, which means no marking is used.
- `threshold-warning` and `threshold-alert` settings: These settings are used only by the FortiGate GUI. They define the warning and alert thresholds for packet loss, latency, and jitter. If a metric exceeds its configured threshold, the FortiGate GUI displays a visual notification. In the example shown on this slide, the measured packet loss has reached 5% for **port1** and is reported as a warning. The measured latency for **port2** exceeds its configured alert threshold (150 msec), which is why a critical notification is shown on the FortiGate GUI. The default value is `0`, which means no threshold is used. In addition to warning and alert threshold display, when a value exceed its SLA threshold, it is highlighted in red and bold.

DO NOT REPRINT
© FORTINET

Advanced SLA Settings—VRF-Aware Health Check

- SD-WAN on multi-VRF network
  - Multi-VRF tunnels
  - Health check tagged with VRF ID
  - Specified source IP



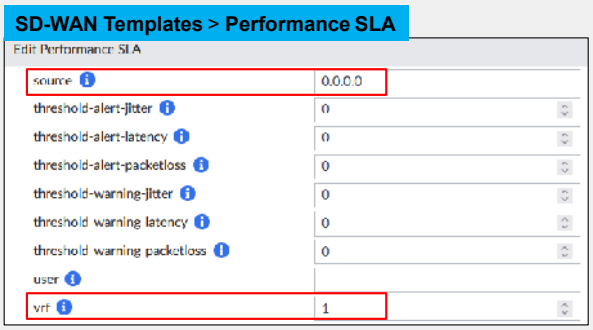
- CLI configuration

```
config system sdwan
    config health-check
        edit "Level3_DNS"
            set vrf 1
            set source <address>
        next
    end
```

VRF id

Source IP address used in health check packets to the server

- GUI



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

41

For SD-WAN on a multi-VRF network, you can decide to have health checks performed per VRF. The originating spoke can tag the health check probes with the correct VRF, when transmitting to a multi-VRF tunnel. The hub will then forward the probes to the correct health check server in the specified VRF.

You can use the advanced performance SLA options to define the VRF ID and the source IP used to forward health check packets to the server as shown on this slide

DO NOT REPRINT
© FORTINET

Review

- ✓ Describe underlay and overlay links, and identify supported interfaces
- ✓ Identify and verify member and zone configuration
- ✓ Describe the benefits of using zones
- ✓ Identify performance SLA configuration sections
- ✓ Understand member active and passive monitoring
- ✓ Describe SLA targets and use cases
- ✓ Monitor member state and performance using FortiManager and the FortiGate CLI
- ✓ Describe the update static route and cascade interfaces actions
- ✓ Identify relevant advanced performance SLA settings

FORTINET Training Institute

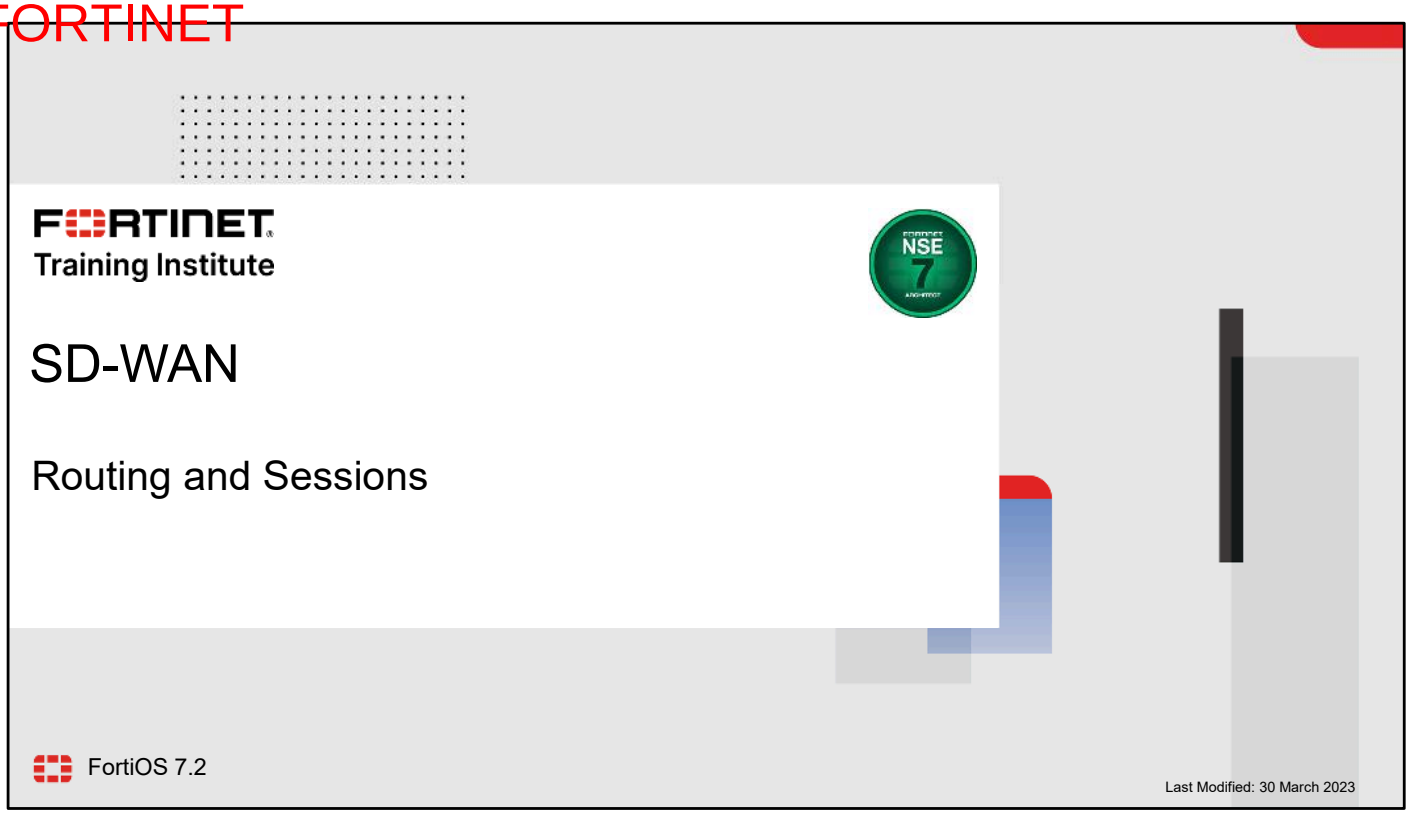
© Fortinet Inc. All Rights Reserved.

42

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

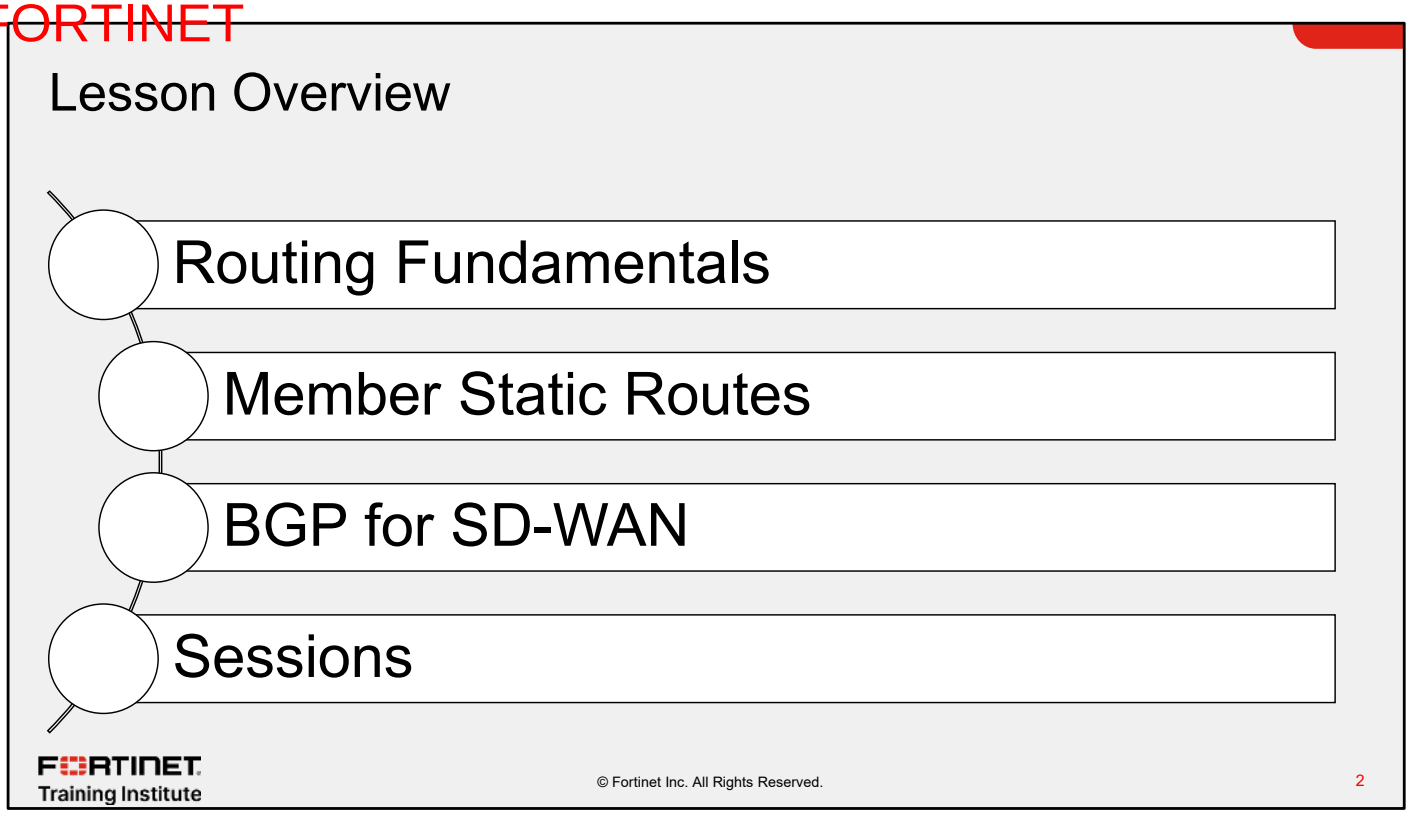
By mastering the objectives covered in this lesson, you learned about the supported interfaces for SD-WAN, the benefits of placing members in zones, and the options available to monitor the state and performance of members.

DO NOT REPRINT
© FORTINET



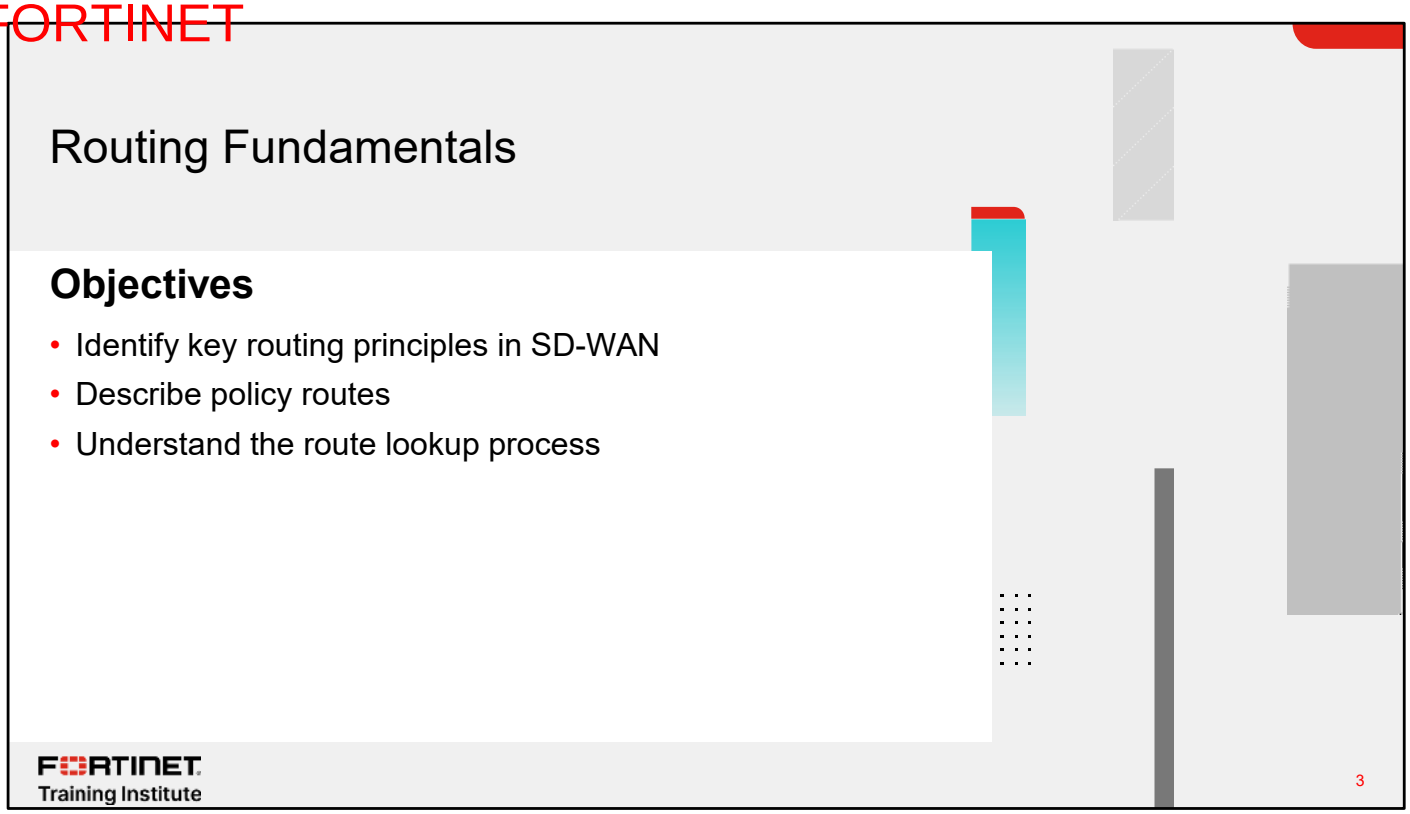
In this lesson, you will learn how FortiGate performs routing and handles sessions for SD-WAN traffic.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



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

By understanding SD-WAN routing fundamentals, you will be able to understand how FortiGate performs routing when SD-WAN is involved.

DO NOT REPRINT
© FORTINET

Key Routing Principles

1. SD-WAN rules are policy routes
2. Regular policy routes have precedence over SD-WAN rules
3. Route lookup is done for new and dirty sessions
   - For original and reply traffic
   - Includes policy route lookup
4. By default, SD-WAN rules are skipped if:
   - Best route to destination isn't an SD-WAN member
   - None of the members have a valid route to destination
     - If the preferred member doesn't have a valid route to destination, the next member in the rule is checked
5. Implicit SD-WAN rule equals standard forwarding information base (FIB) lookup
   - If lookup matches ECMP routes, traffic is load balanced using the configured algorithm

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

Routing is a core component of SD-WAN. Understanding how routing works in SD-WAN is essential for design and troubleshooting. The following are the SD-WAN key routing principles:

1. SD-WAN rules are policy routes. Like regular policy routes, SD-WAN rules route traffic based on multiple criteria. That is, when you configure an SD-WAN rule, the kernel installs a corresponding policy route that reflects the source, destination, service, and outgoing interfaces configured in the SD-WAN rule.

2. Regular policy routes have precedence over SD-WAN rules. Therefore, if you configure regular policy routes, you should ensure that their matching criteria is as narrow as possible. Otherwise, traffic that is intended to be handled by SD-WAN could end up being handled by regular policy routes instead.

3. FortiGate performs route lookup on both new and dirty sessions. A dirty session is a session that must be re-evaluated by the kernel after it is impacted by a routing, firewall policy, or interface change. FortiGate performs route lookups for both original and reply traffic. During route lookup, FortiGate also checks policy routes.

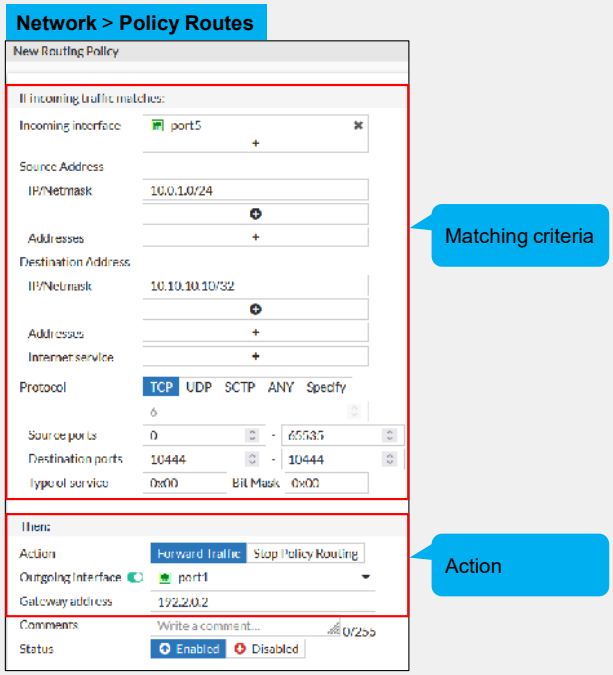
4. By default, FortiGate skips SD-WAN rules if the best route to the destination isn't an SD-WAN member. If the best route matches an SD-WAN member, then the selected member in the rule must have a valid route to the destination, otherwise FortiGate skips the member, and checks the next best member. If none of the members have a valid route to the destination, then FortiGate skips the rule.

5. The implicit SD-WAN rule equals standard FIB lookup. That is, if the traffic doesn't match any of the SD-WAN rules, then FortiGate routes the traffic using the regular process, which consists of looking for the best route in the FIB. If the best route matches equal cost multipath (ECMP) routes—usually the case—then FortiGate load balances the traffic using the configured load balancing algorithm.

DO NOT REPRINT
© FORTINET

Policy Routes

- SD-WAN rules are essentially policy routes
- Policy routes:
  - Provide more granular matching than static routes:
    - Protocol
    - Source address
    - Source ports
    - Destination ports
    - ToS marking
    - Destination internet service
  - Have precedence over SD-WAN rules and entries in the FIB
  - Best practice: narrow down matching criteria



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

When you configure an SD-WAN rule, FortiGate essentially applies a policy route in FortiOS. For this reason, before learning how routing in SD-WAN works, it is convenient that you first understand policy routes.

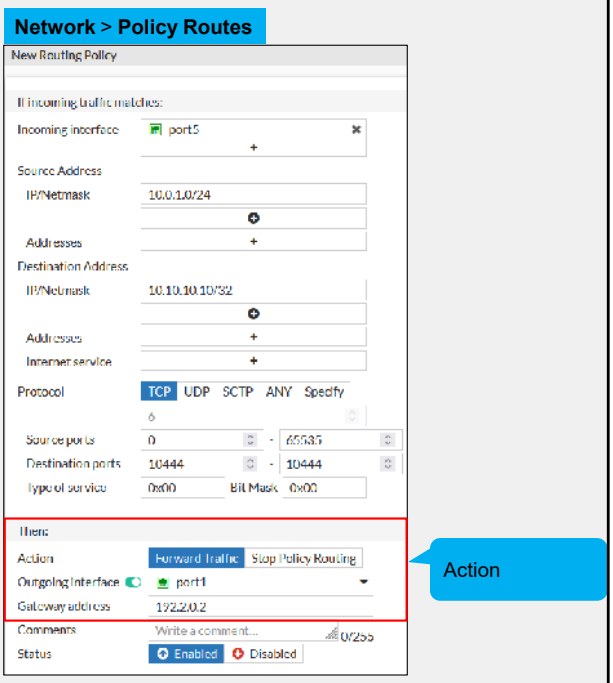
Static routes are simple and are often used in small networks. Policy routes, however, are more flexible because they can match more than just the destination IP address. For example, you can configure as matching criteria the incoming interface, the source and destination subnets, protocol, and port number. Because regular policy routes have precedence over any other routes, it is a best practice to narrow down the matching criteria as much as possible. Otherwise, traffic that is expected to be routed by SD-WAN rules or other routes in the FIB could be handled by regular policy routes instead.

This slide shows an example of a policy route configured using the FortiGate GUI. The policy route instructs FortiGate to match traffic received at **port5**, sourced from `10.0.1.0/24` and destined to the host `10.10.10.10`. The traffic must also be destined to TCP port `10444` for the policy route to match. FortiGate then forwards the traffic—**Forward Traffic** action—to **port1** through the gateway `192.2.0.2`.

DO NOT REPRINT
© FORTINET

Policy Route—Actions

- **Stop Policy Routing**
  - Skips all policy routes, uses the FIB
- **Forward Traffic**
  - Forwards traffic using the set outgoing interface and gateway
  - FIB must have a matching route; otherwise, policy route is considered invalid and skipped



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

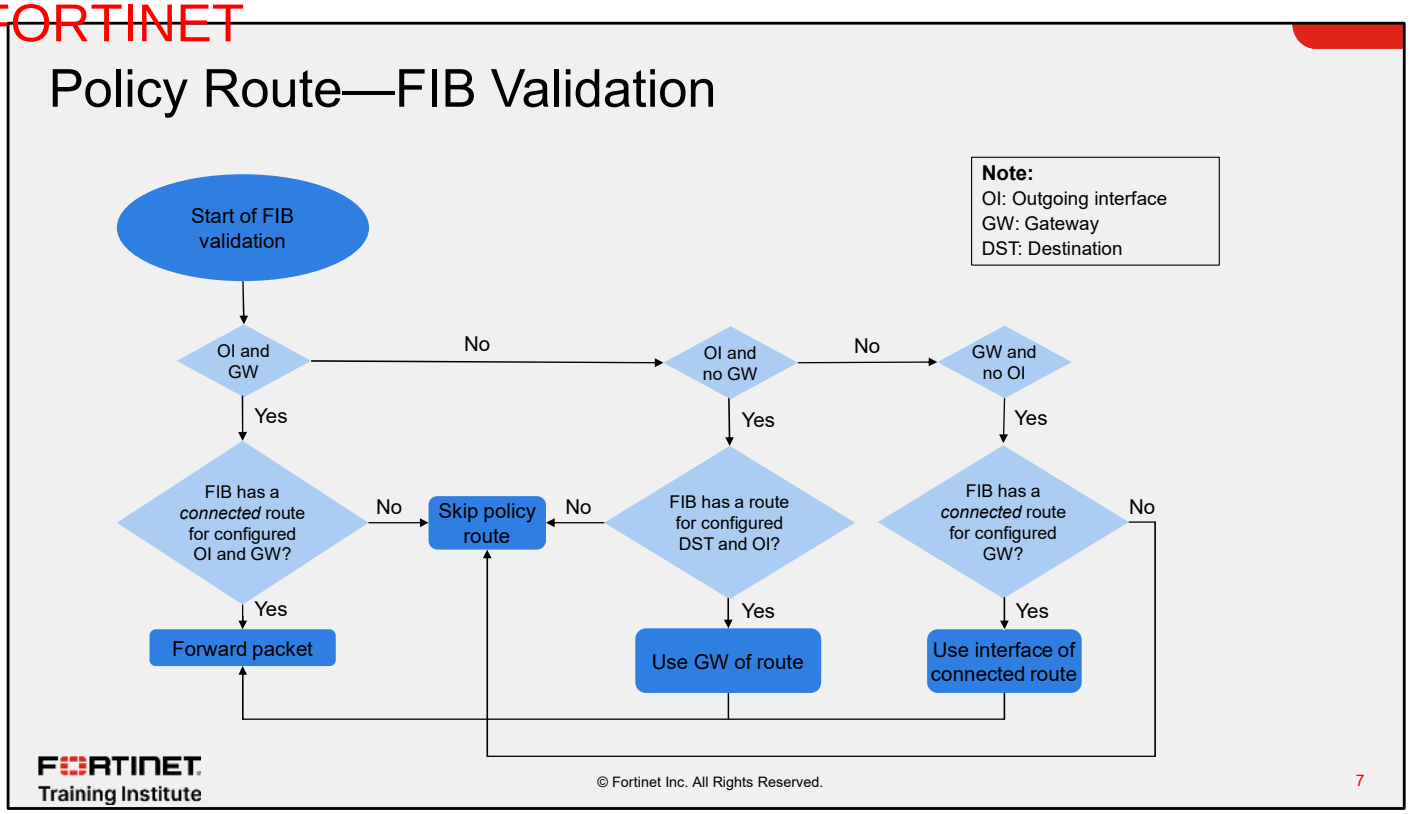
6

When a packet matches a policy route, FortiGate takes one of two actions. Either it routes the packet to the configured outgoing interface and gateway—**Forward Traffic** action—or it stops checking the policy routes—**Stop Policy Routing** action—so the packet is routed based on the FIB.

Note that when you configure **Forward Traffic** as the action, the **Destination Address**, **Outgoing interface**, and the **Gateway address** settings must match a route in the FIB. Otherwise, the policy route is considered invalid and, as a result, skipped.

Also note that policy routes have precedence over SD-WAN rules, and over any routes in the FIB. That is, if a packet matches a policy route and the policy route has a matching route in the FIB, then FortiGate doesn't check any of the configured SD-WAN rules or the routes in the FIB.

DO NOT REPRINT
© FORTINET

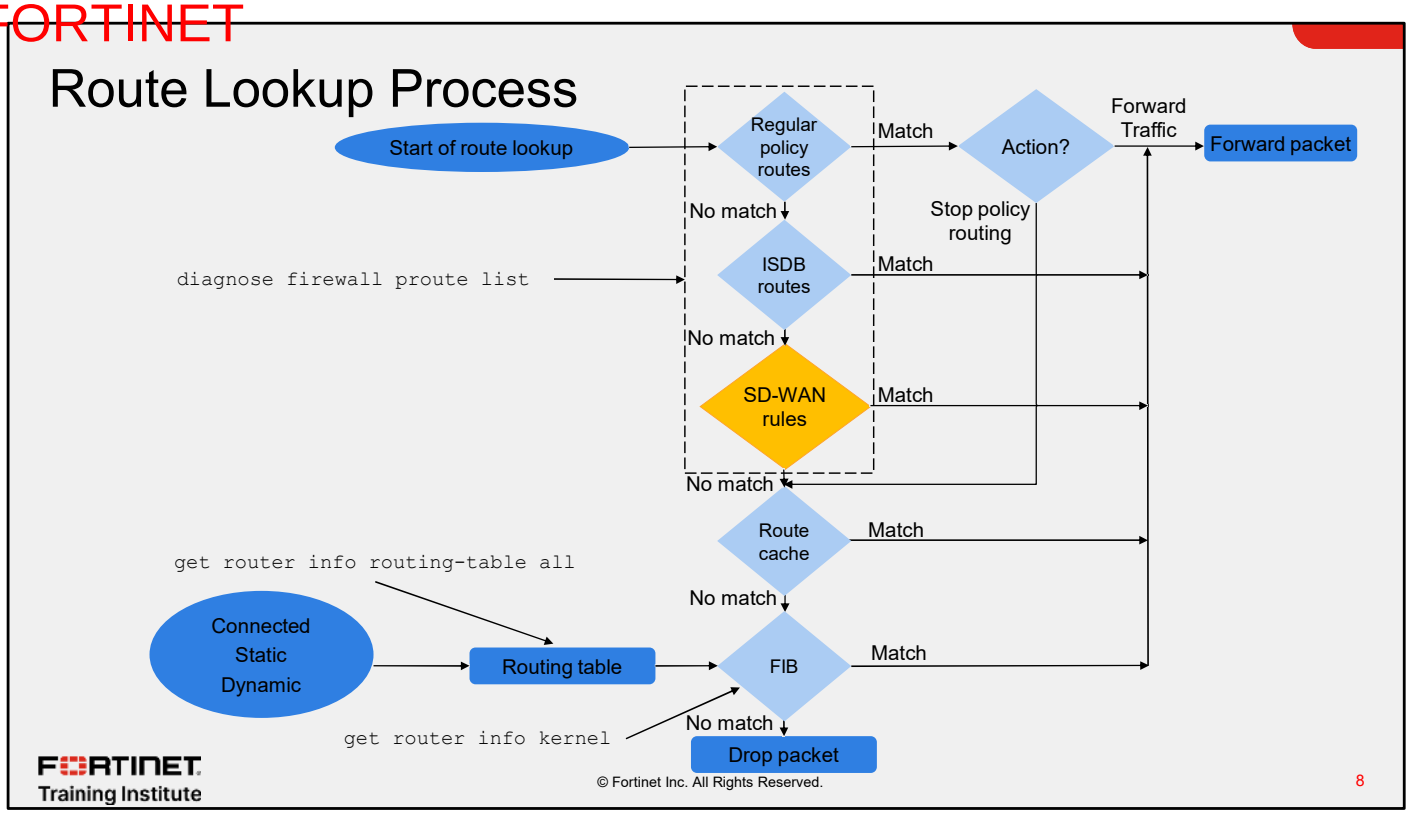


A policy route with the **Forward Traffic** action is subject to a validation against the FIB before the packet is forwarded. The goal is for FortiGate to check the configured **Destination Address**, **Outgoing interface**, and the **Gateway address** settings against the routes in the FIB, or to identify the outgoing interface and gateway to use if they are not specified in the policy route configuration.

The validation is as follows:

- If you configure the outgoing interface and the gateway, FortiGate uses the policy route if the FIB contains a connected route that matches the configured gateway and outgoing interface. If so, FortiGate routes the packet using the configured outgoing interface and gateway. Otherwise, FortiGate skips the policy route.
- If you configure the outgoing interface only—that is, the gateway is set to `0.0.0.0`—FortiGate uses the policy route if the FIB contains a route—any type of route—that matches the configured destination and outgoing interface. If so, FortiGate routes the packet using the outgoing interface of the policy route and the gateway of the matching route in the FIB. Otherwise, FortiGate skips the policy route.
- If you configure the gateway only, FortiGate uses the policy route if the FIB contains a connected route for the configured gateway. If so, FortiGate routes the packet using the interface of the connected route and the configured gateway. Otherwise, FortiGate skips the policy route.

DO NOT REPRINT
© FORTINET



The flowchart on this slide describes the route lookup process that FortiGate performs when it uses policy routes. Note that policy routes can be regular policy routes, internet-service database (ISDB) routes, or SD-WAN rules.

First, FortiGate checks the policy routes. The first type of policy routes to check is the regular policy routes. If there is a match, and the action is **Forward Traffic**, FortiGate routes the packet accordingly provided the policy route passes the FIB validation process. If the action is **Stop Policy Routing**, FortiGate moves on to check its route cache.

If the packet doesn't match any of the regular policy routes, FortiGate moves on to check the ISDB routes first, and then the SD-WAN rules. If the packet doesn't match any of the SD-WAN rules, FortiGate checks its route cache. You will learn more about the SD-WAN rule matching process in another lesson.

Next, FortiGate checks the FIB, which is the table used for performing standard routing. The FIB can be described as the routing table from the kernel point of view, and is built mostly out of routes in the routing table, but also from system-specific entries required by FortiOS. If the packet doesn't match any of the routes in the FIB, FortiGate drops the packet and sends an ICMP destination network unreachable message to the sender.

This slide also shows the FortiOS CLI commands you can use to display the policy routes, the route cache entries, the routing table entries, and the FIB entries.

DO NOT REPRINT
© FORTINET

Policy Route Table

```
# diagnose firewall proute list
list route policy info(vf=root):

id=1 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0-0 iif=7 dport=0-65535
path(1) oif=21(T_MPLS)
source(1): 10.0.1.0-10.0.1.255
destination(1): 10.0.0.0-10.255.255.255
hit_count=18 last_used=2022-11-14 05:47:21

id=2113929223 static_route=7 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0-0
iif=0 dport=1-65535 path(1) oif=3(port1) gwy=192.2.0.2
source wildcard(1): 0.0.0.0/0.0.0.0
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(1): Fortinet-FortiGuard(1245324,0,0,0)
hit_count=0 last_used=2022-11-14 06:39:07

id=2130903041(0x7f030001) vwl_service=1(Critical-DIA) vwl_mbr_seq=1 2 dscp_tag=0xff 0xff flags=0x0
tos=0x00 tos_mask=0x00 protocol=0 sport=0-65535 iif=0 dport=1-65535 path(2) oif=3(port1) oif=4(port2)
source(1): 10.0.1.0-10.0.1.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(3): GoToMeeting(4294836966,0,0,0, 16354)
Microsoft.Office.365.Portal(4294837474,0,0,0, 41468) Salesforce(4294837976,0,0,0, 16920)
hit_count=0 last_used=2022-11-14 05:46:43
```

This is a regular policy route (ID ≤ 65535)

This is an ISDB route (ID > 65535 and no `vwl_service` field)

This is an SD-WAN rule (ID > 65535 and the `vwl_service` field is present)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

FortiOS maintains a policy route table that you can view by running the `diagnose firewall proute list` command.

There are three types of policy routes displayed in the policy route table: regular policy routes, ISDB routes, and SD-WAN rules. Follow these rules to identify each type of policy route in the table:

- Regular policy routes are assigned an ID no higher than 65535. In the output shown on this slide, the first entry is assigned ID 1, which makes it a regular policy route.
- ISDB routes and SD-WAN rules are assigned an ID higher than 65535. However, SD-WAN rule entries include the `vwl_service` field, and ISDB route entries don't. The `vwl_service` field indicates the ID and the name of the rule from the SD-WAN configuration perspective. In the output shown on this slide, the second entry is an ISDB route and the third entry an SD-WAN rule.

Note that although IDs for regular policy routes are in the 1 to 65535 range, the maximum number of regular policy routes that you can configure are much lower and varies among models. For example, you can configure up to 512 regular policy routes in a FortiGate 100F device. For more information about the maximum supported values per model, refer to the `FortiOS Maximum Values Table` on `docs.fortinet.com`. Alternatively, you can run the `print tablesize` command on the FortiGate CLI to get the maximum values for your device.

DO NOT REPRINT
© FORTINET

Matching Policy Route in Debug Flow

- The debug flow output indicates the ID of the matching policy route:

```
branch1_fgt # id=65308 trace_id=145 func=print_pkt_detail line=5892 msg="vd-root:0 received a
packet(proto=1, 10.0.1.101:2124->10.0.2.101:2048) tun_id=0.0.0.0 from port5. type=8, code=0, id=2124,
seq=1."
id=65308 trace_id=145 func=init_ip_session_common line=6073 msg="allocate a new session-00050a89,
tun_id=0.0.0.0"
id=65308 trace_id=145 func=rpdb_srv_match_input line=1043 msg="Match policy routing id=2130903043: to
10.0.2.101 via ifindex-3"
...
```

Matching policy route ID (ID based on the policy route table)

FORTINET Training Institute

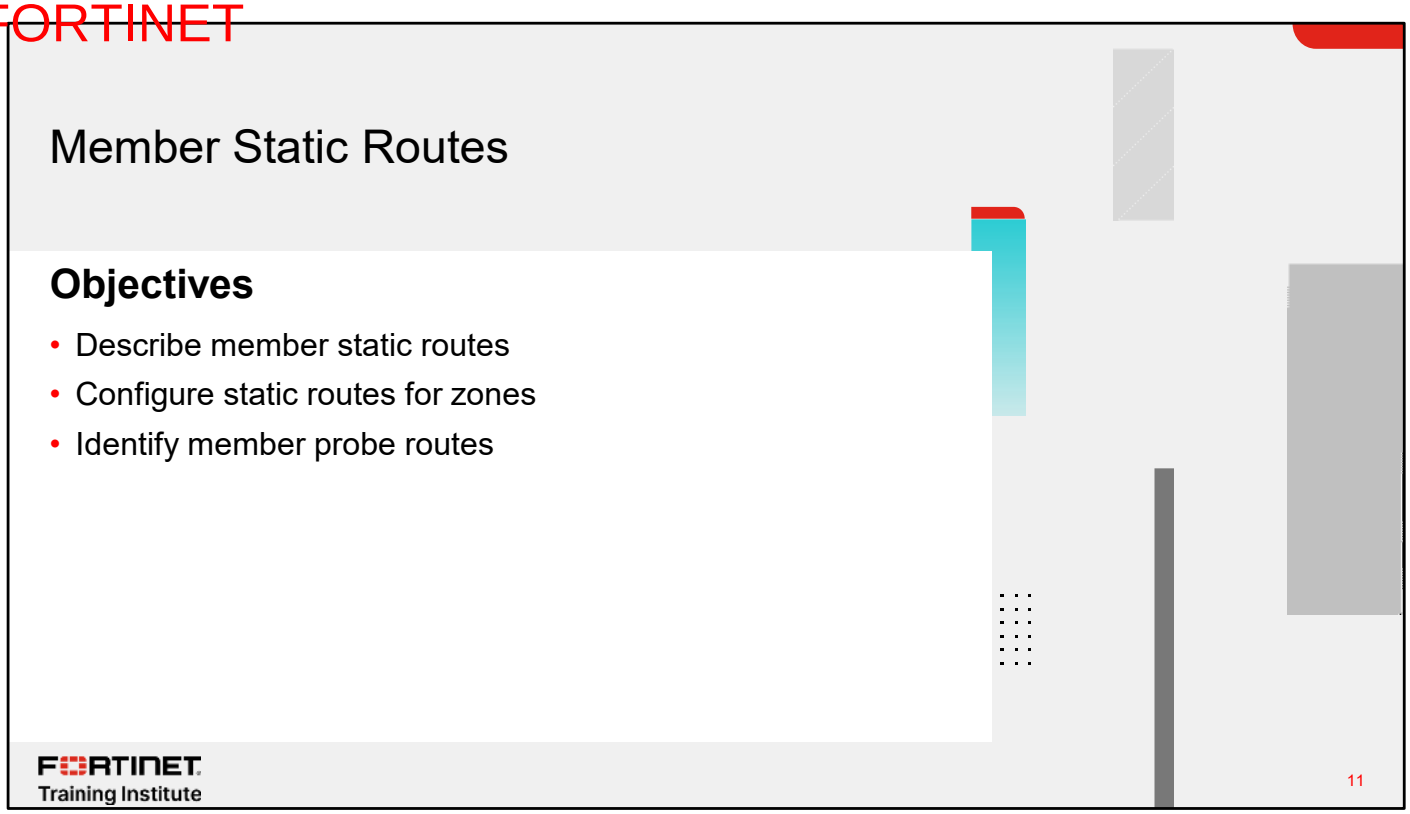
© Fortinet Inc. All Rights Reserved.

10

The debug flow tool indicates the ID of the matching policy route, which is useful for troubleshooting purposes. Note that the displayed ID corresponds to the entry ID in the policy route table.

You can collect the debug flow on the CLI, using the `diagnose debug flow` commands or from the FortiGate GUI menu: **Network** > **Diagnostics** > **Debug Flow**.

DO NOT REPRINT
© FORTINET



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

By understanding static routes for SD-WAN members, you will be able to configure static routes for members based on SD-WAN requirements and limitations.

DO NOT REPRINT
© FORTINET

Member Routes

- By default, FortiGate requires a valid route to steer traffic to a member
  - Best practice: Ensure that sites learn all possible routes to all possible destinations
- Static routes
  - Reference a SD-WAN zone
    - Common case, simplified configuration
    - Individual ECMP routes installed for each member in the zone
    - Gateway obtained from member configuration
  - Reference a member
    - More granular control
- DHCP and PPPoE interfaces
  - Automatic default static routes are removed

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

One of the key routing principles is that, by default, SD-WAN requires a valid route in the FIB so the member can be used to steer traffic. Because the goal is to have SD-WAN pick the best member to forward the traffic to, based on the SD-WAN rule criteria, it's a best practice to configure your routing setup so your SD-WAN sites know all possible routes to all possible destinations that are intended for handling by SD-WAN. Otherwise, SD-WAN may fail to choose the best member, not because it doesn't meet the application requirements, but because FortiGate doesn't have a route for the destination and member.

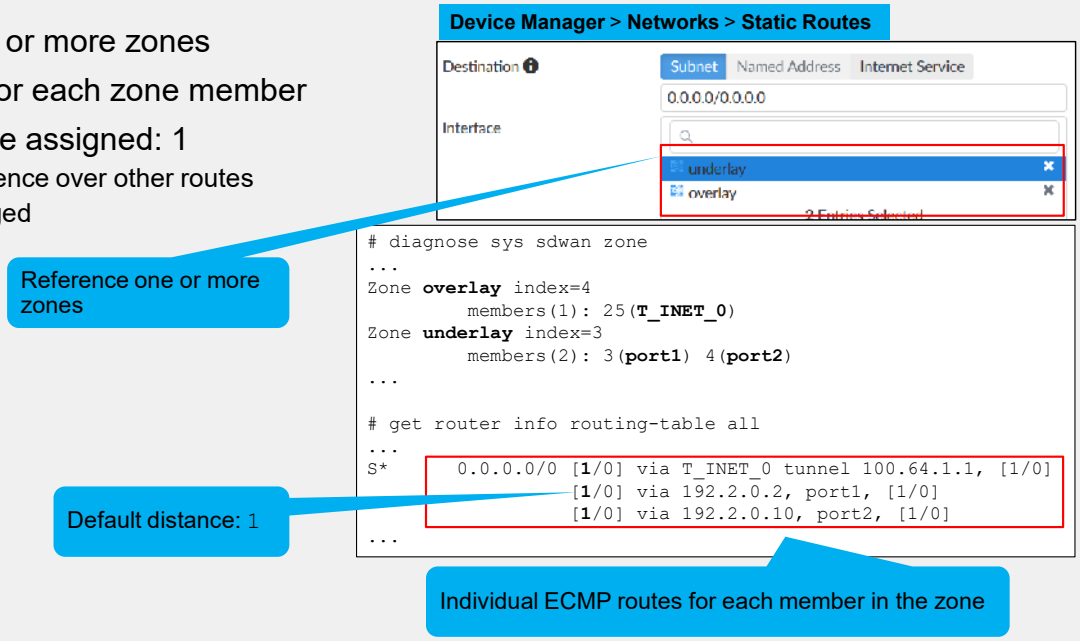
You can use static and dynamic routing in SD-WAN. This section focuses on static routing. When you configure static routes for SD-WAN, you usually reference an SD-WAN zone to simplify the configuration. However, you can also reference individual members instead, so you can have more granular control over traffic. When you reference a zone in a static route, FortiGate installs individual routes for each member in the zone. FortiGate then displays the individual routes in the routing table as equal cost multipath (ECMP) routes.

Note that when you configure a static route that references an SD-WAN zone, you don't have to specify a gateway address because FortiGate retrieves it from the member configuration. Also note that after you configure a DHCP or PPPoE interface as member, FortiGate removes the automatic default static routes for these interfaces from the routing table.

DO NOT REPRINT
© FORTINET

Static Routes for SD-WAN Zones

- Reference one or more zones
- ECMP routes for each zone member
- Default distance assigned: 1
  - Higher preference over other routes
  - Can be changed



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

You can reference one or more zones when you configure a static route. As a result, FortiOS installs a static route for every member configured in the zone in the routing table. Because the static routes share the same distance, they become ECMP routes.

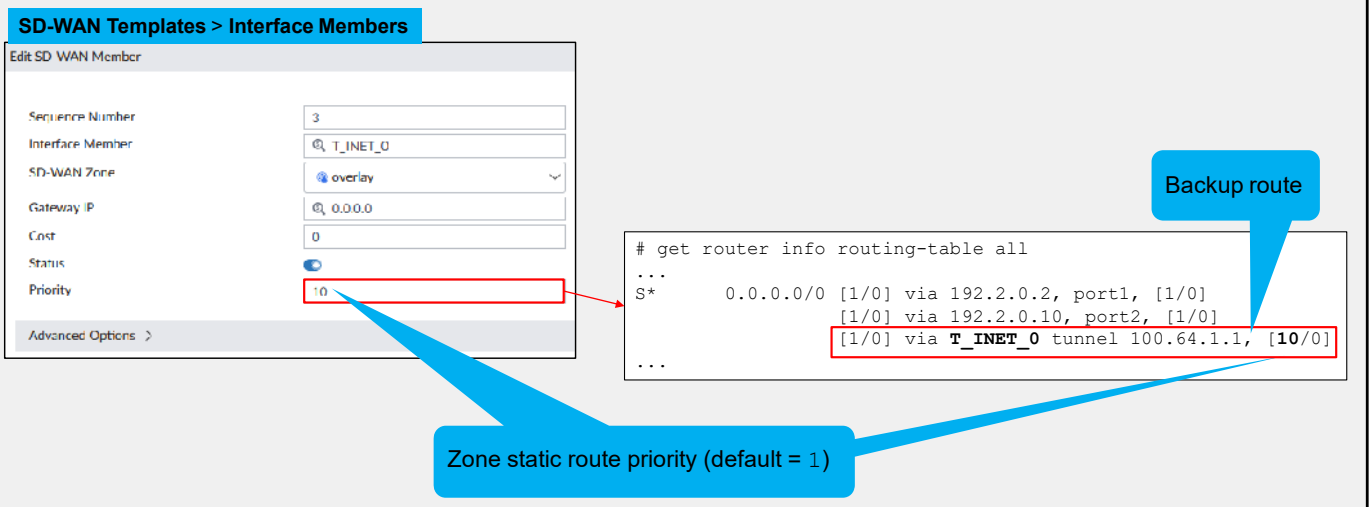
When you create a static route for a zone, FortiOS assigns the routes with a distance of `1` by default. FortiOS assigns such a low distance by default because administrators usually want their SD-WAN routes to have preference over other routes in the FIB. However, you can change the distance to a different value if required.

In the example shown on this slide, `port1` and `port2` are members of the `underlay` zone, and the IPsec tunnel `T_INET_0` is a member of the `overlay` zone. The administrator created a default static route that references both zones. The result is that the routing table displays ECMP routes for each of the three members.

DO NOT REPRINT
© FORTINET

Member Priority

- Control member preference in static routes for SD-WAN zone
- Used as tie-breaker for ECMP routes in zone (implicit SD-WAN rule)



Traffic that doesn't match an SD-WAN rule is handled by the implicit SD-WAN rule, which means that traffic is routed based on the FIB contents. If you configure a static route for a zone, which results in ECMP routes, FortiGate load balances the connections that don't match an SD-WAN rule but match the destination of the ECMP routes, across all zone members.

In the same setup as the previous slide, internet traffic that matches the implicit rule is load balanced across underlay and overlay links. But what if, in that setup, you want FortiGate to prefer the underlay links for internet traffic to leverage the lower latency of direct internet access (DIA)? And then use the overlays only if the underlays are dead? You can achieve this behavior by increasing the priority of overlay interfaces.

By default, the member priority is set to `1`. The higher the priority number, the lower the member preference. In the example shown on this slide, the overlay `T_INET_0` is assigned `10` as priority. The result is that internet traffic that doesn't match the SD-WAN rules is load balanced across `port1` and `port2`. FortiGate uses `T_INET_0` to route internet traffic only if both `port1` and `port2` are dead.

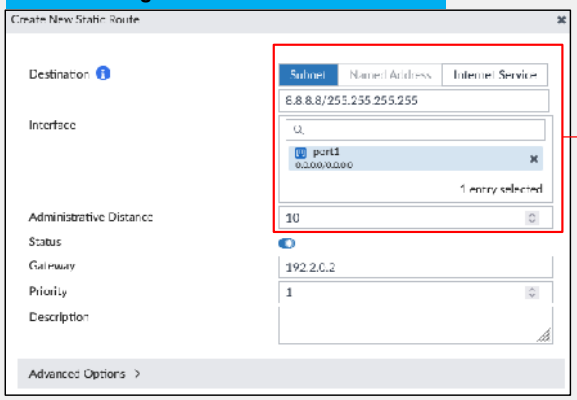
Note that the member priority setting available in the SD-WAN configuration applies only to *zone static routes*. The priority of per-member static routes is controlled by the priority setting available in the static route configuration.

DO NOT REPRINT
© FORTINET

Static Routes for Members

- Configure per-member static routes
  - More granularity
- Gateway setting may be required
  - Gateway is not retrieved from member settings

Device Manager > Network> Static Routes




```
# get router info routing-table all
...
S*      0.0.0.0/0 [1/0] via 192.2.0.2, port1, [1/0]
                  [1/0] via 192.2.0.10, port2, [1/0]
                  [1/0] via T_INET_0_tunnel_100.64.1.1, [10/0]
S       8.8.8.8/32 [10/0] via 192.2.0.2, port1, [1/0]
...
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

Alternatively, you can configure per-member static routes for more granular control over traffic. However, unlike static routes for zones, which retrieve the member gateway from the member configuration, with per-member static routes, you must specify a gateway if the interface type requires it.

In the example shown on this slide, the administrator configured a per-member static route for `8.8.8.8` through **port1**. The route can then be used by SD-WAN rules to route traffic, or by the FIB to route traffic when no rule is matched.

DO NOT REPRINT
© FORTINET

Duplicate Static Route Conflict

- Duplicate routes involving a zone and an interface are not allowed
  - Overlapped routes are allowed
  - Automatic check when configured at device level
  - Prompt for error on install if configured with template

Device Manager > Network> Static Routes

Error: router/static/2/ : You cannot have duplicated routes (2 and 1) on SD-WAN and non-SD-WAN interfaces ×

Install Wizard > Device Settings

Install Log

Starting log (Run on device)

Start installing
branch1_fgt $ config router static
branch1_fgt (static) $ edit 2
branch1_fgt (2) $ set gateway 192.2.0.2
branch1_fgt (2) $ set device "port1"
branch1_fgt (2) $ next
You cannot have duplicated routes on SD-WAN and non-SD-WAN interfaces.
object set operator error, -9100 discard the setting
Command fail. Return code 1
branch1_fgt (static) $ end

Installation fails

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

If you configure a static route for a destination that references a zone, FortiOS doesn't allow you to configure a static route for the same destination that references an interface. In SD-WAN, these routes are known as duplicate routes. You can, however, configure overlapping routes for zones and interfaces.

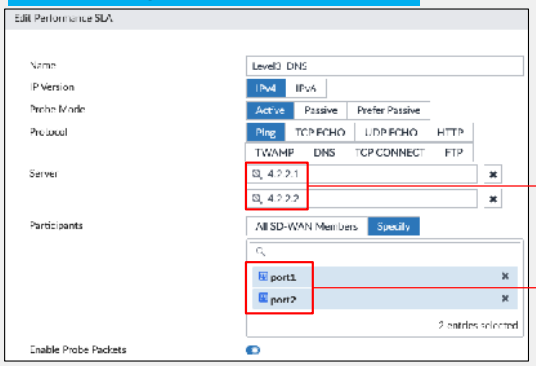
In the example shown on this slide, the administrator tries to configure a default static route for **port1** on FortiManager at the device level. However, because there is already a default static route configured for the **underlay** and **overlay** zones, FortiManager rejects the second route configuration and displays an error message. If a similar configuration is done at the template level, FortiManager detects the conflict during device settings installation, and the install log reports an error indicating that such configuration is not allowed. As a result, the installation fails.

DO NOT REPRINT
© FORTINET

Member Probes Routes

- FortiGate adds entries in FIB for probes
  - Flagged with `proto=18`

SD-WAN Templates > Performance SLA




```
# get router info kernel | grep proto=18
tab=254 vf=0 scope=0 type=1 proto=18 prio=0 192.2.0.1/255.255.255.255/0->4.2.2.1/32 pref=0.0.0.0 gwy=192.2.0.2 dev=3(port1)
tab=254 vf=0 scope=0 type=1 proto=18 prio=0 192.2.0.9/255.255.255.255/0->4.2.2.1/32 pref=0.0.0.0 gwy=192.2.0.10 dev=4(port2)
tab=254 vf=0 scope=0 type=1 proto=18 prio=0 192.2.0.1/255.255.255.255/0->4.2.2.2/32 pref=0.0.0.0 gwy=192.2.0.2 dev=3(port1)
tab=254 vf=0 scope=0 type=1 proto=18 prio=0 192.2.0.9/255.255.255.255/0->4.2.2.2/32 pref=0.0.0.0 gwy=192.2.0.10 dev=4(port2)
```

FORTINET Training Institute

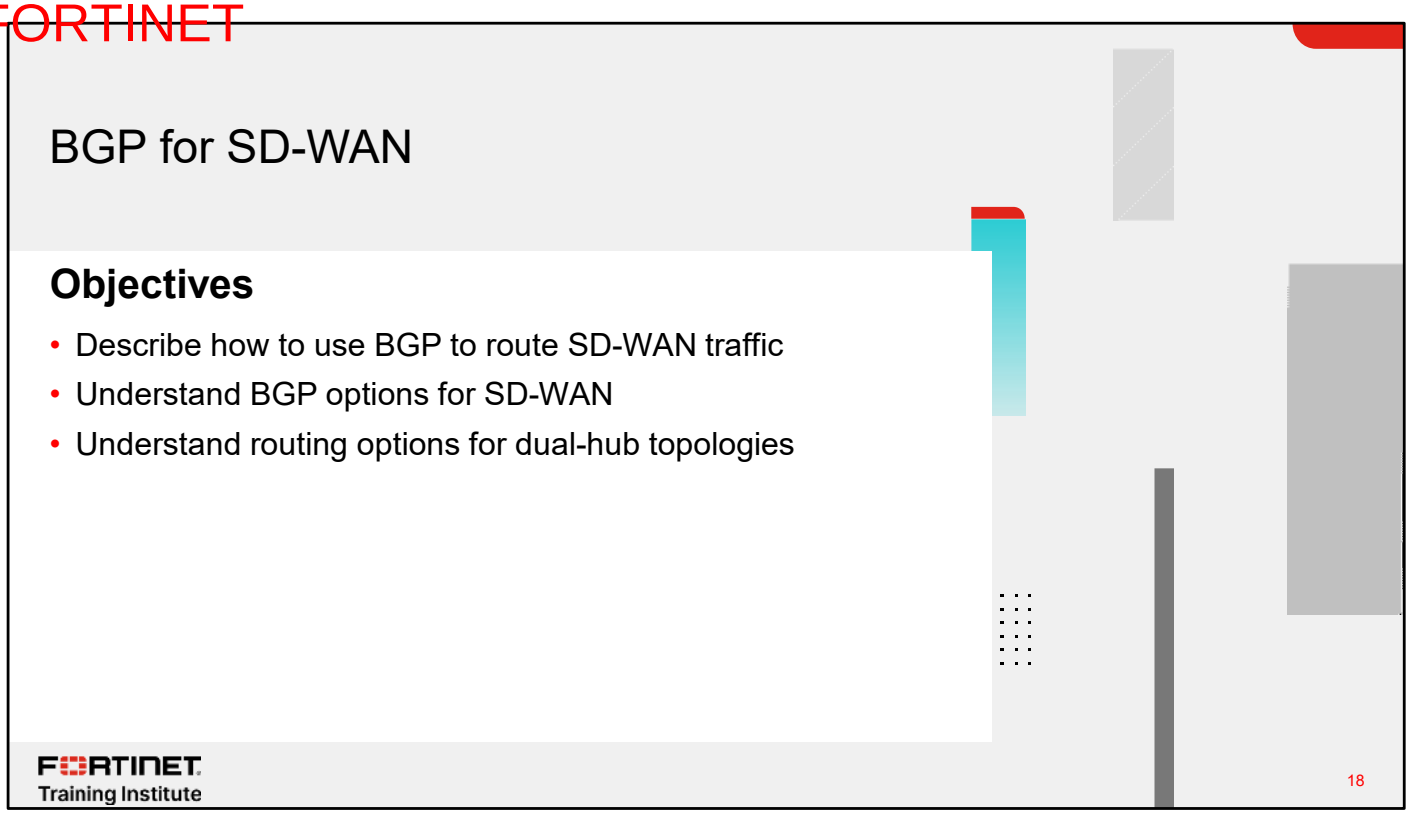
© Fortinet Inc. All Rights Reserved.

17

When you use active monitoring for measuring the performance of members, FortiGate doesn't use the entries in the routing table to route the probes. Instead, FortiGate adds a FIB route entry for the configured target servers through each monitoring member and its gateway. These kernel routes are flagged as `proto=18`, as shown on this slide.

Note that `proto=18` has nothing to do with assigned internet protocol numbers. It is an arbitrary protocol number assigned by Fortinet to identify entries in the FIB used for routing health-check traffic.

DO NOT REPRINT
© FORTINET



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

By demonstrating competence in understanding SD-WAN specific BGP options, you will be able to configure BGP routing for SD-WAN in context of single-hub and dual-hub topologies.

DO NOT REPRINT
© FORTINET

BGP for SD-WAN Traffic

- SD-WAN rules can use BGP routes to steer traffic:
  - iBGP
  - eBGP
- BGP options for SD-WAN
  - BGP tags
    - SD-WAN rules steer traffic to routes using the specified route tag
  - SD-WAN self-healing
  - Route mapping according to SLA status
    - `route-map-out-preferable`
    - `minimum-sla-meet-members`
  - BGP neighbor status
    - Primary–Secondary
    - Standalone

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

As explained previously, one of the key routing principles is that, by default, SD-WAN requires a valid route in the FIB so the member can be used to steer traffic. The required valid route can be static or dynamic. You can combine the use of SD-WAN and the main dynamic routing protocols —OSPF and BGP. Fortinet recommends the use of BGP in the overlay, because it allows routing and steering of SD-WAN traffic based on health-check results. In the next few slides, you will discover how to use health-check results to adjust BGP routing.

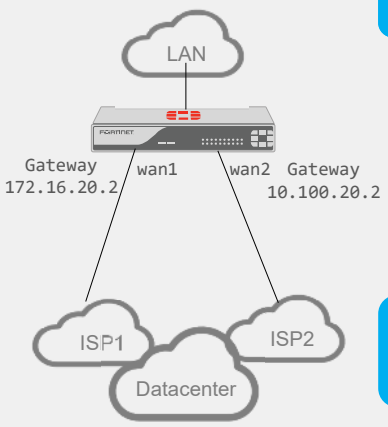
You can combine SD-WAN with iBGP or eBGP. This lesson focuses only on what is specific for SD-WAN.

When using SD-WAN with BGP routing you can:
- Configure SD-WAN rules to steer traffic to members only if routes have a specific route-tag.
- Interact with a BGP neighbor to optimize route selection—called SD-WAN self-healing.
- Use a route map that relates to the SLA status with the following parameters: `route-map-out-preferable` and `minimum-sla-meet-members`.
- Define a primary or secondary role for BGP neighbors.

DO NOT REPRINT
© FORTINET

BGP Tags

- Use BGP tags in SD-WAN rules
  - Example: Use a tag to apply a service rule for traffic to the data center



Set the tag to route, received from ISP2 to the data center.

Use the tag in the SD-WAN rule.
The rule can use only routes that have the tag set.

```
config router community-list
    edit "30:5"
      config rule
         edit 1
            set action permit
            set match "30:5"
         next
       ...
config router route-map
    edit "comm1"
      config rule
         edit 1
           set match-community "30:5"
           set set-route-tag 15
         next
       ...
config router bgp
    config neighbor
      edit 10.100.20.2
        set remote-as 65001
        set route-map-in "comm1"
        ...
config system sdwan
    config service
      edit 1
        set name "DataCenter"
        set mode manual
        set route-tag 15
        set priority-members 2
      next
      ...
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

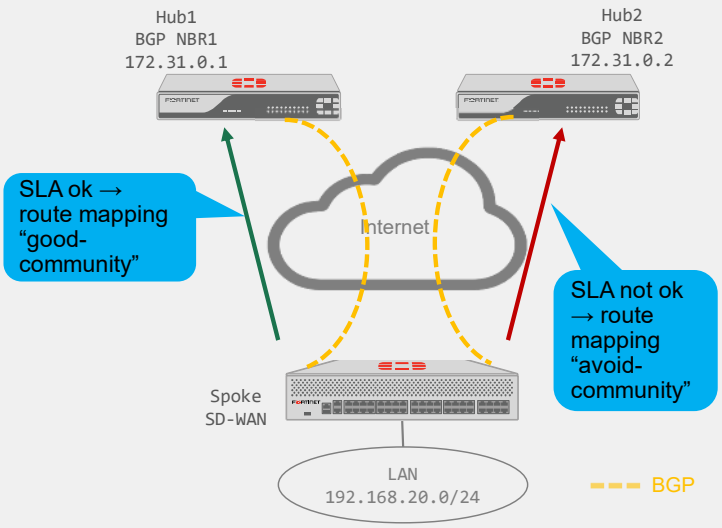
In this example shown above, a customer has two ISP connections, wan1 and wan2. wan1 is used primarily for direct access to internet applications, and wan2 is used primarily for traffic to the customer's data center. The datacenter has a dynamic IP address range, therefore the administrator cannot configure SD-WAN rules based on addresses. As an alternative option, the administrator can use BGP tags.

The CLI example shows how an administrator can use BGP tags in SD-WAN rules.
In the example, the BGP neighbor wan2 advertises the data center network range with a community number of 30:5. FortiGate use this community number to define a tag used by the SD-WAN rule. The SD-WAN rule will match traffic according to the defined criteria and can use an underlying FIB route only if it matches the specified tag.

DO NOT REPRINT
© FORTINET

SD-WAN Self-Healing With BGP

- BGP route mapping to inform BGP neighbors of link SLA status
  - BGP route-map-out adjusted from SLA status
    - `route-map-out-preferable`
      - Route mapping when SLA passed
    - `route-map-out`
      - Route mapping when SLA missed
- BGP peer
  - Non-SD-WAN device
    - Routes traffic according to standard BGP rules
    - Can be any BGP device
    - Adjusts routing from info received in BGP routing updates
  - SD-WAN device
    - Adjusts routing from info received in BGP routing updates
    - No health check needed
    - Uses rules to steer traffic from health info received



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

SD-WAN allows you to select different outbound WAN links based on performance SLAs. When a device selects a link to forward traffic, it is useful to inform BGP neighbors about this choice and allow them to adjust routing accordingly. This can limit use of bad performer links and avoid asymmetric routing. It is called SD-WAN self-healing with BGP.

On FortiGate, BGP advertisements can adapt to the SLA status of a link in the following ways:
- Apply different route-maps based on the SD-WAN health checks. For example, it can advertise different BGP community strings when SLAs are met or not.
- Forward traffic based on the active BGP neighbor.

Use the following BGP parameters to adjust route-maps according to SD-WAN health checks: `route-map-out` and `route-map-out-preferable`. When the SLAs are met, FortiGate advertises routes with settings defined with `route-map-out-preferable`. When SLAs are missed, it will setting `route-map-out`.

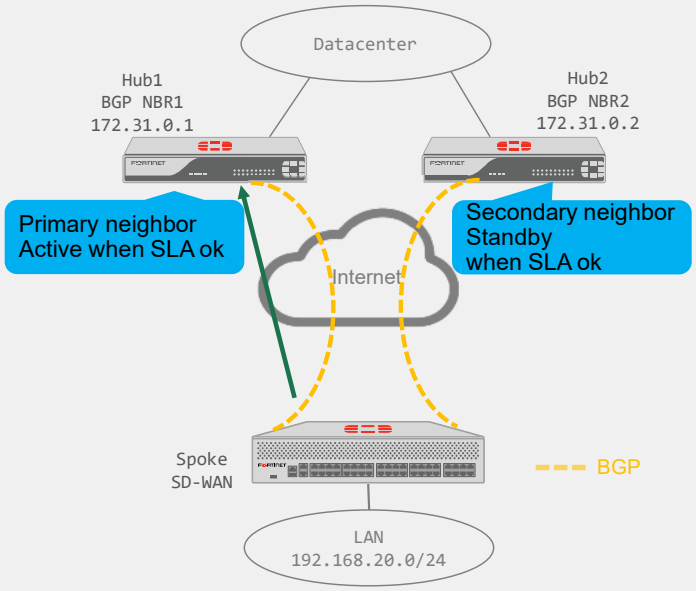
The BGP peer does not need to be compatible with SD-WAN:
- If the BGP peer is a device with standard BGP routing features only, it will adjust its routing from advertisement fields received from the SD-WAN device—usually community or route-tag.
- If the BGP peer is an SD-WAN device, it does not need to run its own health check. The administrator configure the SD-WAN service rules to match each route-tag and steer traffic to the corresponding VPN overlay when the neighbor link is healthy.

Note that, because all branches advertise the same community strings and route-tags, you need to configure only one set of service rules on the hub for all branches. No additional settings are needed when new branches join the network. This feature makes this solution easily scalable.

DO NOT REPRINT
© FORTINET

Routing for Dual-Hub Topology

- Select Primary and Secondary neighbor
  - Primary
    - Preferred neighbor if SLAs are met
  - Secondary
    - Active neighbor if primary SLAs are not met
  - Standalone
    - Default status (if primary/secondary not set)
    - Fallback mode if the SLAs for primary and secondary neighbor are not met



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

A dual-hub topology allows you to set one hub as active and the other as standby. The link SLAs determine which hub to use. To achieve this, you define a primary and a secondary BGP neighbor on the SD-WAN spoke device. The spoke direct the traffic to the primary hub when link SLAs are met. If SLAs are not met for the primary device, and SLAs for the secondary unit are still within threshold, the spoke will adjust BGP routing to direct traffic to the secondary unit.

If you don't set primary and secondary roles for SD-WAN BGP peers, the default standalone mode applies. FortiGate establishes BGP peering with both hubs without considering SD-WAN SLAs.

If you do set primary and secondary roles, if neither the primary nor secondary neighbors are active (SLAs missed), the SD-WAN neighbor status becomes standalone. Only service rules with standalone-action enabled will continue to pass traffic. This option is disabled by default.

DO NOT REPRINT
© FORTINET

Route Map and Role Configuration

- Configure a route map

```
config router route-map
    edit "comm1"
      config rule
         edit 1
            set match-ip-address "net192"
            set set-community "20:1"
         next
       ...
```

- Configure route mapping for BGP neighbors

```
config router bgp
    config neighbor
      edit "172.31.0.1"
         set route-map-out "comm5"
         set route-map-out-preferable "comm1"
      next
      edit "172.31.0.2"
         set route-map-out "comm5"
         set route-map-out-preferable "comm2"
      next
    end
```

- Configure SD-WAN neighbor:
  - Assign a role
  - Assign a health check

```
config system sdwan
    config neighbor
      edit "172.31.0.1"
         set member 1
         set role primary
         set health-check "HC1"
         set sla-id 1
      next
      edit "172.31.0.2"
         set member 2
         set role secondary
         set health-check "HC2"
         set sla-id 1
      next
    end
```

Default role = standalone

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

This slide shows the CLI configuration for route mapping based on the diagram shown on the previous slide. FortiGate advertises the route to local network 192.168.20.0/24 with a different community according SLA status.

First, you define the router access list `net192` for prefix `192.168.20.0/24` (not shown).
Next, you define the following route mapping:

- Use community 20:1 for the primary neighbor preferred route map (comm1)
- Use community 20:2 for the secondary neighbor preferred route map (comm2)
- Use community 20:5 if no SLAs are met (comm5)

For BGP neighbor configuration, use two route maps: `route-map-out-preferable` to define mapping when the SLA is met and `route-map-out` for mapping when the SLA fails.

In the SD-WAN configuration section, you can define the BGP neighbor role (primary, secondary, standalone) and the health check to determine the device role and mapping to use. If you do not set a role (default = standalone), the BGP neighbor is active regardless of SLA status, but different route maps are used: `route-map-out-preferable` (comm1 or comm2) if the SLA is met and `route-map-out` (comm5) if the SLA is missed.

DO NOT REPRINT
© FORTINET

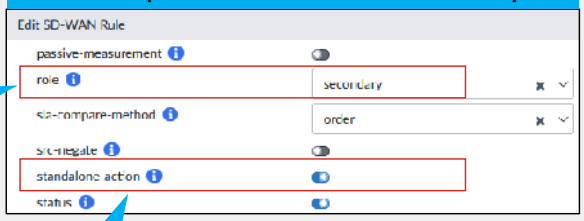
Dual-Hub—Policy Settings

- Configure the SD-WAN service
  - Active according to hub role: standalone, primary, secondary
  - Enable or disable rule after fallback to standalone mode

```
config system sdwan
    config service
       edit 1
          set name "Critical-DIA"
          set role secondary
          set standalone-action enable
          set src "LAN-net"
          set priority-members 1 2
       next
    end
end
```

Default role = standalone

SD-WAN Templates > SD-WAN Rules > Advanced Options



Allow rule after fallback to standalone role
Default = disable

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

After you set the primary and secondary roles, you can define the SD-WAN rules for each scenario. Then, you can decide if each rule remains active or not in case of fallback to standalone mode. By default, the primary and secondary rules are inactive when the SLAs for both roles are not met.

In the example shown on this slide, the rule steers traffic when the secondary hub is active. Standalone-action is enabled, so the rule will continue to accept traffic when SLAs are missed for both hubs.

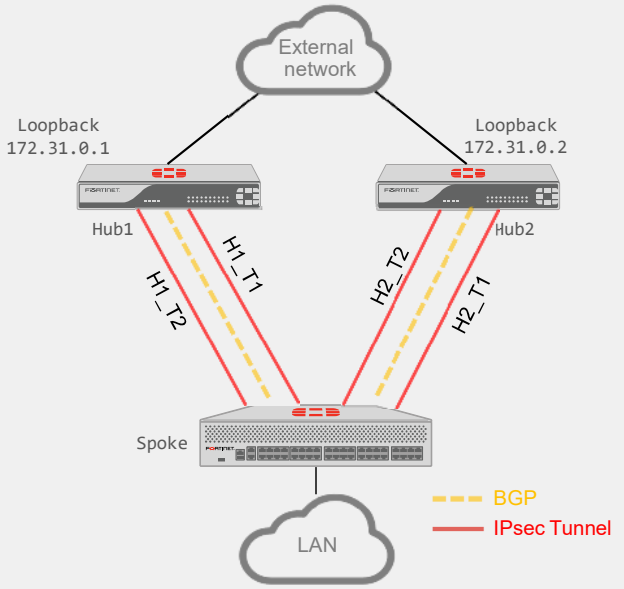
DO NOT REPRINT
© FORTINET

Multiple Members per SD-WAN Neighbor

- By default, `route-map-out-preferable` used if all members to neighbor meet SLA
- Refine with `minimum-sla-met-members`

```
# config system sdwan
    config neighbor
        edit "172.31.0.1"
            set member 1 2
            set minimum-sla-meet-members 2
            set health-check "Level3_DNS"
            set sla-id 1
        next
    end
```

Number of members out of SLA to trigger change from `route-map-out-preferable` to `route-map-out` (default=1)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

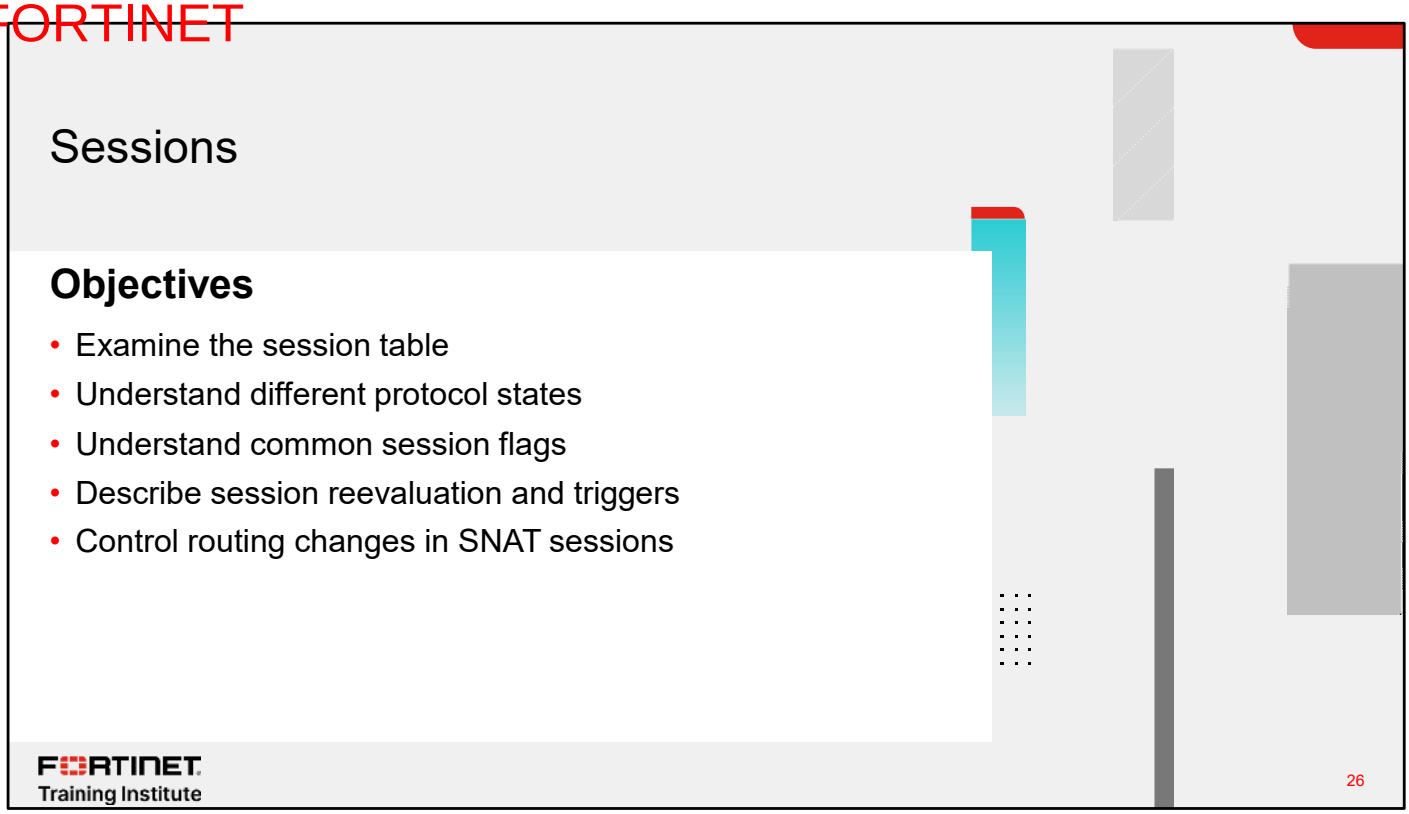
As you learned on previous slides, when the SD-WAN member meets the SLA threshold, FortiGate applies the route map defined in the BGP neighbor's `route-map-out-preferable` option. If the SD-WAN member fails to meet the SLA, FortiGate applies the route map defined in the BGP neighbor's `route-map-out` option instead.

What happens when multiple SD-WAN members connect to the same BGP neighbor?

By default, FortiGate applies `route-map-out-preferable` only when all members connected to the BGP neighbor meet SLA criteria. In some situations, where the administrator may prefer to keep sending traffic to the primary neighbor even if one link fails to meet the SLA, with a fallback to the secondary neighbor only when the situation gets worse.
Using the `minimum-sla-meet-members` parameter, you can refine the behavior and define the number of links that must fail the SLA before the routing changes from `route-map-out-preferable` to `route-map-out`.

For example, in the network shown in the diagram, you can continue to route traffic to Hub1 if one tunnel, H1_T1 or H1_T2, still meets the SLA.
When `minimum-sla-meet-members is` set to `2`, FortiGate advertises settings corresponding to `route-map-out-preferable` when at least one tunnel meets the SLA. A change to `route-map-out` occurs when two or more tunnels do not meet the SLA.

DO NOT REPRINT
© FORTINET



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

By demonstrating competence in understanding FortiOS sessions, you will be able to identify relevant SD-WAN fields in a session, the different factors that trigger a session reevaluation, and the settings that enable you to change the default behavior.

DO NOT REPRINT
© FORTINET

Sessions—Table Summary

```
# get system session status
The total number of IPv4 sessions for the current VDOM: 11

# get system session list
PROTO    EXPIRE SOURCE           SOURCE-NAT       DESTINATION      DESTINATION-NAT
tcp      3522   10.1.10.1:41418  100.64.1.1:41418 172.217.0.106:443        -
udp      151    10.1.0.1:4387    100.64.1.1:64803 208.91.112.220:53        -
udp      28     10.1.0.1:1687    100.64.1.1:62103 208.91.112.52:53         -
udp      61     100.64.1.1:3075  -                208.91.112.53:53-
udp      55     100.64.1.1:3075  -                208.91.112.52:53         -
udp      172    10.1.10.1:123    100.64.1.1:60539 209.121.129.48:123       -
udp      152    10.1.0.1:4387    100.64.1.1:64803 173.243.138.221:53       -
udp      152    10.1.0.1:4387    100.64.1.1:64803 45.75.200.89:53          -
udp      171    10.1.10.1:123    100.64.1.1:60539 208.81.1.197:123         -
udp      104    10.1.0.1:1420    -                10.1.0.254:53            -
tcp      3600   10.1.10.1:34433  -                10.1.0.254:22            -
udp      176    10.1.0.1:1900    -                10.1.0.254:53            -
tcp      3524   10.1.10.1:59972  100.64.1.1:59972 172.217.0.110:80         -
tcp      119    10.1.10.1:42824  100.64.1.1:42824 72.21.91.29:80           -
tcp      3517   10.1.10.1:42816  100.64.1.1:42816 72.21.91.29:80           -
udp      171    10.1.10.1:123    100.64.1.1:60539 144.217.65.184:123       -
udp      175    10.1.10.1:123    100.64.1.1:60539 144.217.65.182:123       -
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

The FortiGate session table contains detailed information about every IP connection that crosses, terminates, or is initiated by FortiGate. The command `get system session status` displays the total number of sessions in the current VDOM. The command `get system session list` provides a brief summary of each session. This command lists one session per line, and includes information such as protocol, source IP address, destination IP address, and port. You can use the `grep` utility to filter the output.

DO NOT REPRINT
© FORTINET

Sessions—TCP Example

```
# diagnose sys session list
session info: proto=6 proto_state=11 duration=1 expire=3599 timeout=3600 flags=00000000
socktype=0 sockport=0 av_idx=0 use=4
origin-shaper=medium prio=3 guarantee 0Bps max 134217728Bps traffic 232868Bps drops 0B
reply-shaper=medium prio=3 guarantee 0Bps max 134217728Bps traffic 232868Bps drops 0B
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr npu f00 app valid
statistic(bytes/packets/allow_err): org=1720/9/1 reply=10804/13/1 tuples=3
tx speed(Bps/kbps): 0/0 rx speed(Bps/kbps): 0/0
orgin->sink: org pre->post, reply pre->post dev=7->31/31->7 gwy=10.1.0.254/10.9.31.117
hook=post dir=org act=snat 10.9.31.117:45388->200.8.57.5:443(10.1.0.3:45388)
hook=pre dir=reply act=dnat 200.8.57.5:443->10.1.0.3:45388(10.9.31.117:45388)
hook=post dir=reply act=noop 200.8.57.5:443->10.9.31.117:45388(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 pol_uuid_idx=14720 confiauth_info=0 chk_client_info=0 vd=0
serial=0002932f tos=ff/ff app_list=2000 app=34050 url_cat=0
sdwan_mbr_seq=1 sdwan_service_id=1
rpdb_link_id=80000000 ngfwid=n/a
npu_state=0x003c94 ips_offload
npu info: flag=0x81/0x81, offload=8/8, ips_offload=1/1, epid=16/16, ipid=64/88,
vlan=0x0000/0x0000
vlifid=64/88, vtag_in=0x0000/0x0000 in_npu=1/1, out_npu=1/1, fwd_en=0/0, qid=0/0
```

SD-WAN session information

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

This slide shows an example output detailing information particular to a single session table entry. From left to right, and from top to bottom, the following information is highlighted:

- The IP protocol number and the protocol state
- The length of time until the session expires (if no more traffic matches the session)
- Traffic shaping counters
- Session flags
- Received and transmitted packet and byte counters
- The original and reply direction of the traffic. If the device is doing NAT, this portion shows the type of NAT (source or destination) for each traffic direction, and the NAT IP address.
- The ID of the matching policy
- The SD-WAN specific session information. `sdwan_mbr_seq` and `sdwan_service_id` indicate the SD-WAN member ID and the SD-WAN rule ID in use, respectively. If the session matched the SD-WAN implicit rule, and therefore was handled using standard FIB routing, these SD-WAN fields do not appear.
- Counters for hardware acceleration

Use the CLI command `diagnose sys session list` for IPv4 traffic, and the command `diagnose sys session6 list` for IPv6 traffic.
The output for both commands is similar (it displays the same fields, in the same order).

DO NOT REPRINT
© FORTINET

May Dirty Sessions

- New firewall sessions created after matching a firewall policy with `accept` action
  - A firewall policy lookup is done (top-down)
  - Flagged as `may_dirty`
- Lookup process
  - First original packet (route and firewall policy lookup)
  - First reply packet (route lookup only)
  - No additional lookups unless session is flagged as `dirty`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

When working with FortiGate, especially with SD-WAN setups, it's important to understand the concept of may_dirty and dirty sessions.

May_dirty sessions are firewall sessions that were created after matching a firewall policy with `accept` as action. For FortiGate to identify the matching policy, it performs a firewall policy lookup, selecting the first matching policy in the configuration from the top down. Most firewall sessions contain the `may_dirty` flag. However, some sessions such as expectation sessions, do not contain the `may_dirty` flag because they are not created as a result of matching a firewall policy.

For new sessions, FortiGate performs route and firewall policy lookups upon receiving the first packet (in the original direction). FortiGate also performs a route lookup—but not a firewall policy lookup—for the first reply packet. FortiGate then saves the information that results from the route lookup—the outgoing interface and gateway to use—and the firewall policy lookup—policy ID, address translation, inspection, authentication, logging, and so on—to the session.

FortiGate doesn't perform additional lookups for the session unless the session is flagged as `dirty`.

DO NOT REPRINT
© FORTINET

Dirty Sessions

- A session is flagged as `dirty` after a routing, firewall policy, or interface change
  - Each direction of a dirty session must be reevaluated
  - Routing changes are common in SD-WAN
  - Session routing information is flushed (routing change)
- Default reevaluation process after a routing change (*no SNAT*):
  - Next packet in each direction goes to the CPU (applies only to NPU-offloaded sessions)
  - Original direction: Route and firewall policy lookups for first packet
  - Reply direction: Only route lookup for first packet
  - FortiGate updates the session with new routing and firewall policy information, and removes the `dirty` flag
  - If the action is still `accept`, FortiGate forwards the packet
  - If the new action of the matching firewall is `deny`, then:
    - FortiGate flags the session as `block`, and drops the packet
    - The session remains in the session table until it expires
    - FortiGate drops further packets matching the session
  - FortiGate eventually offloads sessions again to hardware

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

By default, FortiGate flags all `may_dirty` sessions as `dirty` if there is a routing, firewall policy, or interface change. The FortiGate must reevaluate every dirty session. During reevaluation, it checks both directions of the dirty session.

Dirty sessions are common in SD-WAN because of its dynamic nature. Members that are down or performing poorly may change the entries in the FIB and policy route table, respectively. In either case, the result is a routing change that may affect multiple SD-WAN sessions.

After a routing change, the routing information of may-dirty sessions is flushed. The default reevaluation process following a routing change for sessions *without* source NAT (SNAT) is as follows:

- If the impacted session is offloaded to the NPU, FortiGate instructs the NPU to send the next packet from each direction of the session to the CPU. This ensures that the session is handled by the CPU, and thus reevaluated. If the session is not offloaded to the NPU, then the packets are always handled by the CPU, and this step is not required.
- In the original direction, FortiGate performs route and firewall policy lookups for the first packet.
- In the reply direction, FortiGate performs only a route lookup for the first packet.
- FortiGate updates the session with the new routing and firewall policy information, and removes the `dirty` flag.
- If the matching firewall policy action is still `accept`, FortiGate forwards the packet.
- If the matching firewall policy action is `deny`, FortiGate flags the session as `block`, drops the packet, and retains the session in the session table until it expires. FortiGate also drops any additional packets that match the session .
- FortiGate eventually offloads allowed sessions again to the NPU, if they still meet the NPU offloading requirements.

DO NOT REPRINT
© FORTINET

Dirty Sessions (Contd)

- Before the route change (`may_dirty` flag)

```
session info: proto=1 proto_state=00 duration=4 expire=55 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0
use=3
state=log may_dirty npu f00
orgin->sink: org pre->post, reply pre->post dev=7->19/19->7 gwy=100.64.1.1/10.0.1.101
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
sdwan_mbr_seq=3 sdwan_service_id=2
```

- After the route change (`dirty` flag; gateway is flushed)

```
session info: proto=1 proto_state=00 duration=30 expire=29 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0
use=3
state=log dirty may_dirty npu f00
orgin->sink: org pre->post, reply pre->post dev=7->19/19->7 gwy=0.0.0.0/0.0.0.0
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
sdwan_mbr_seq=3 sdwan_service_id=2
```

- After reevaluation (no `dirty` flag; outgoing interface, gateway, and SD-WAN member updated)

```
session info: proto=1 proto_state=00 duration=53 expire=56 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0
use=3
state=log may_dirty npu f00
orgin->sink: org pre->post, reply pre->post dev=7->20/20->7 gwy=100.64.1.9/10.0.1.101
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
sdwan_mbr_seq=4 sdwan_service_id=2
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

This slide shows the transition of an ICMP session from `may_dirty` to `dirty`, and then back to `may_dirty`, following a route change. Note that only relevant lines of the session are displayed.

Before the route change, the session is flagged as `may_dirty`. The session output shows the index number of the outgoing interface in use (`19`), the gateway information (`gwy`), and the configuration index number of the SD-WAN member in use (`3`).

After the route change, the session is flagged as `dirty`. Note that the `may_dirty` flag is still there, and the `dirty` flag is just added. The gateway information is also flushed.

After reevaluation, the `dirty` flag is removed, and the index number of the outgoing interface and the gateway information are updated based on the new route. Because the outgoing interface changed, the configuration index number of the SD-WAN member also changed.

Note that the firewall policy (`policy_id`) didn't change. This is because the firewall policy lookup during reevaluation resulted in the same firewall policy, but this is not always the case.

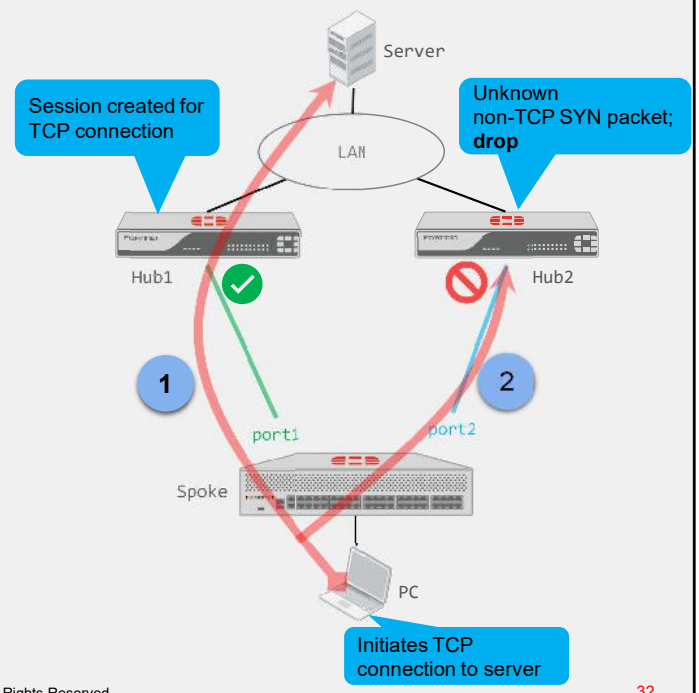
DO NOT REPRINT
© FORTINET

Interface Stickiness

- Force sessions to keep using outgoing interface and gateway after a route change:

```
config system interface
    edit <interface>
        set preserve-session-route enable
    next
end
```

- Current route must still be present in FIB
  - Otherwise FortiGate flags the session as dirty and reevaluates it



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

The reevaluation of a dirty session following a route change may result in a failover to another SD-WAN member. If the SD-WAN members are connected to different devices, especially a stateful firewall, it can cause interruption of TCP sessions.

To avoid a route change, when current route is still available but no longer the best route, you can enable the interface-level command shown on this slide. It forces the session to stay on the same SD-WAN member, provided the route in use by the session is still present in the FIB.

However, if the route is removed from the FIB, then FortiGate must flag the session as `dirty`, flush its gateway information, and reevaluate the session.

In the topology shown on this slide, if FortiGate establishes the session via port1, but due to SLA changes, the best route is now via port2, the behavior is as follow:

- With `preserve-session-route` **`disable`**, FortiGate reevaluate the session, and redirects the traffic through port2. Hub2 drops any already established TCP sessions.
- With `preserve-session-route` **`enable`**, FortiGate does not reevaluate the session, and the session remains established through port1 and hub1. Active TCP sessions do not change. FortiGate routes new sessions through port2.

DO NOT REPRINT
© FORTINET

Interface Stickiness (Contd)

- Session is flagged as `route_preserve`

```
config system interface
    edit "T_INET_0"
        set preserve-session-route enable
    next
end

# diagnose sys session list
session info: proto=1 proto_state=00 duration=4 expire=55 timeout=0 flags=00000000 socktype=0 sockport=0
av_idx=0 use=3
state=log may_dirty npu f00 route_preserve
orgin->sink: org pre->post, reply pre->post dev=7->19/19->7 gwy=100.64.1.1/10.0.1.101
sdwan_mbr_seq=3 sdwan_service_id=2

# diagnose netlink interface list | grep index=19
if=T_INET_0 family=00 type=768 index=19 mtu=1420 link=0 master=0

# diagnose sys sdwan member | grep (3)
Member(3): interface: T_INET_0, flags=0x4 , gateway: 100.64.1.1, priority: 0 1024, weight: 0
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

This slide shows the details of an ICMP session established through an interface (`T_INET_0`) that has the setting `preserve-session-route` enabled. Note that only relevant lines of the session are displayed.

Also note the presence of the `route_preserve` flag in the session. You can find out the interface name and SD-WAN member name by using the `diagnose netlink interface list` and `diagnose sys sdwan member` commands, respectively.

DO NOT REPRINT
© FORTINET

Routing Changes and SNAT Sessions

- By default, SNAT sessions are not flagged as `dirty` after a routing change
  - Exception: route in use is removed from FIB
- Force reevaluation of SNAT sessions after a routing change (default = `disable`):

```
config system global
    set snat-route-change < enable | disable >
end
```

- If SNAT IP changes during reevaluation, packet is dropped, and session is cleared

```
id=20085 trace_id=51 func=print_pkt_detail line=5746 msg="vd-root:0 received a packet(proto=1, 10.0.1.101:13106->8.8.8.8:2048) from port5. type=8, code=0, id=13106, seq=3."
id=20085 trace_id=51 func=resolve_ip_tuple_fast line=5827 msg="Find an existing session, id-00008483, original direction"
id=20085 trace_id=51 func=vf_ip_route_input_common line=2589 msg="Match policy routing id=2131230721: to 8.8.8.8 via ifindex-4"
id=20085 trace_id=51 func=vf_ip_route_input_common line=2615 msg="find a route: flag=04000000 gw-192.2.0.10 via port2"
id=20085 trace_id=51 func=get_new_addr line=1229 msg="find SNAT: IP-192.2.0.9(from IPPOOL), port-13106"
id=20085 trace_id=51 func=fw_strict_dirty_session_check line=264 msg="SNAT IP 192.2.0.1 != 192.2.0.9, drop"
```

- Enable `snat-route-change` if using the same IP pool for old and new path

Different SNAT IP; drop the packet and clear the session

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

By default, SNAT sessions are not flagged as `dirty` following a routing change that impacts the session. This is true if the route in use is still in the FIB. However, if the route is removed from the FIB, then FortiGate flags the session as dirty, flushes the outgoing interface and gateway information, and then re-evaluates the session on the next packet received.

To force the reevaluation of SNAT sessions following a related routing change, regardless of whether the route in use is still in the FIB or not, enable the `snat-route-change` setting, as shown on this slide.

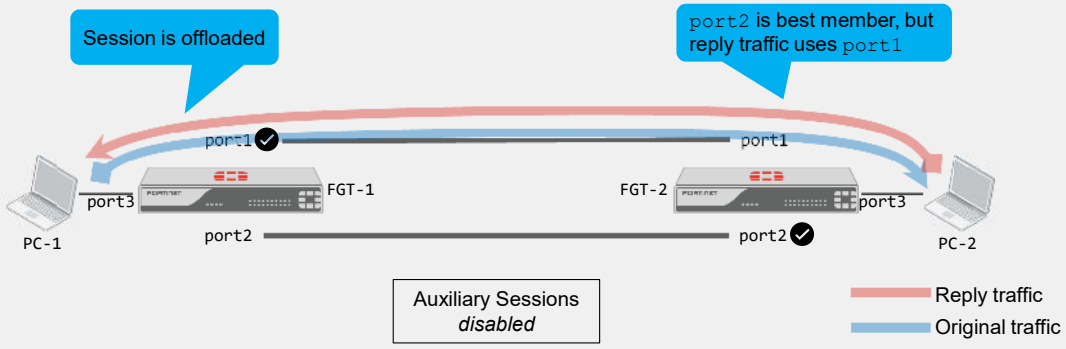
Note that during reevaluation of SNAT sessions, if the new route and firewall policy lookup results in a change of the SNAT IP, then FortiGate drops the packet and clears the session. This means that the impacted application could have to initiate a new connection to resume network connectivity, especially if the application is TCP-based.

This slide shows the debug flow output for a SNAT session during reevaluation. Because the SNAT IP of the new path (`192.2.0.9`) is different from the SNAT IP of the current path (`192.2.0.1`), FortiGate drops the packet and clears the session. This also means that, from a connectivity point of view, it makes sense to enable `snat-route-change` only when the new path for the session uses the same SNAT IP, which can be achieved using IP pools.

DO NOT REPRINT
© FORTINET

Auxiliary Sessions

- Dirty sessions are also triggered by a change in the reply traffic interface
  - Sessions handled by system CPU (no hardware offload)
- By default, route lookup for reply traffic considers routes over original ingress interface only
  - Reply traffic can't be routed over another member with better performance



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

Dirty sessions are also a result of a change in the reply traffic interface. This change is often seen in SD-WAN because a site may switch to a better performing link at any time, which could result in a change in the reply traffic interface for a session. By default, the system CPU handles dirty sessions that are triggered by reply interface changes. Therefore, in this case, hardware offloading is not used. For this reason, a huge amount of traffic handled by these dirty sessions may result in high CPU utilization and poor performance.

In addition, by default, when performing route lookups for the reply direction, FortiGate considers only routes through the same ingress interface used in the original direction. This is done to preserve session symmetry. However, this also prevents reply traffic from switching to a better performing member, which can impact some sensitive applications such as voice and video.

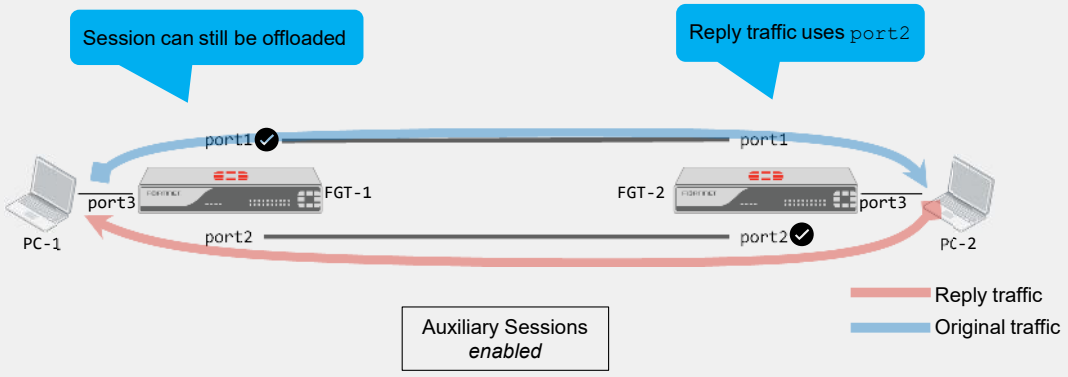
Auxiliary sessions solve the two previous issues by enabling FortiGate to:

1. Offload asymmetric sessions to hardware by creating an auxiliary session (also known as a reflect session). Symmetric traffic matches the main session, and asymmetric traffic matches the auxiliary session. Both sessions can be offloaded to hardware. For FortiGate VMs, although hardware offload does not apply, performance is improved because FortiGate does not have to reevaluate the asymmetric traffic.
2. Select the best route for reply traffic through any member, not necessarily the same interface where the original incoming traffic was received.

This slide shows two FortiGate devices with **auxiliary sessions disabled** and that are connected over two links. PC-1 initiates a connection to PC-2, and FGT-1 routes the connection over port1 because it's the best port. On FGT-2, the best port is port2. However, when FGT-2 performs a route lookup for the reply traffic, it selects the route over port1 because port1 is the incoming interface in the original direction. Also, both FortiGate devices can offload the session to hardware because the session is symmetric.

DO NOT REPRINT
© FORTINET

Auxiliary Sessions (Contd)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

If you *enable auxiliary sessions on both FortiGate* devices in the previous topology, FGT-2 can now route the reply traffic through the best performing member, which is port2. FGT-1 can continue to offload the asymmetric traffic because it matches the auxiliary session (or reflect session). If you don't enable auxiliary sessions on FGT-1, FGT-1 uses the system CPU to handle the asymmetric traffic.

DO NOT REPRINT
© FORTINET

Auxiliary Sessions (Contd)

- Enable auxiliary sessions per-VDOM on the FortiGate CLI (default = `disable`):

```
config system settings
    set auxiliary-session enable
end
```

- Debug flow sample for auxiliary session (FGT-1):

```
trace_id=51 func=print_pkt_detail line=5746 msg="vd-root:0 received a packet(proto=6, 10.0.1.7:22->10.9.31.160:7932) from port2. flag [.], seq 453029597, ack 2210383401, win 11"
trace_id=51 func=resolve_ip_tuple_fast line=5827 msg="Find an existing session, id-00045e02, reply direction"
trace_id=51 func=vf_ip_route_input_common line=2615 msg="find a route: flag=00000000 gw-10.9.31.160 via port3"
trace_id=51 func=fw_forward_dirty_handler line=384 msg="auxiliary ses proto=6 dev=7->6 10.9.31.160/7932=>10.0.1.7/22"
trace_id=51 func=npu_handle_session44 line=1184 msg="Trying to offloading session from port2 to port3, skb.npu_flag=00000400 ses.state=00010280 ses.npu_state=0x04000000"
trace_id=51 func=ip_session_install_npu_session line=353 msg="npu session installation succeeded"
trace_id=51 func=fw_forward_dirty_handler line=394 msg="state=00010280, state2=00000000, npu_state=04000800"
```

- Interface index numbers:

```
# diagnose netlink interface list
if=port1 family=00 type=1 index=5 mtu=1500 link=0 master=0
if=port2 family=00 type=1 index=6 mtu=1500 link=0 master=0
if=port3 family=00 type=1 index=7 mtu=1500 link=0 master=0
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

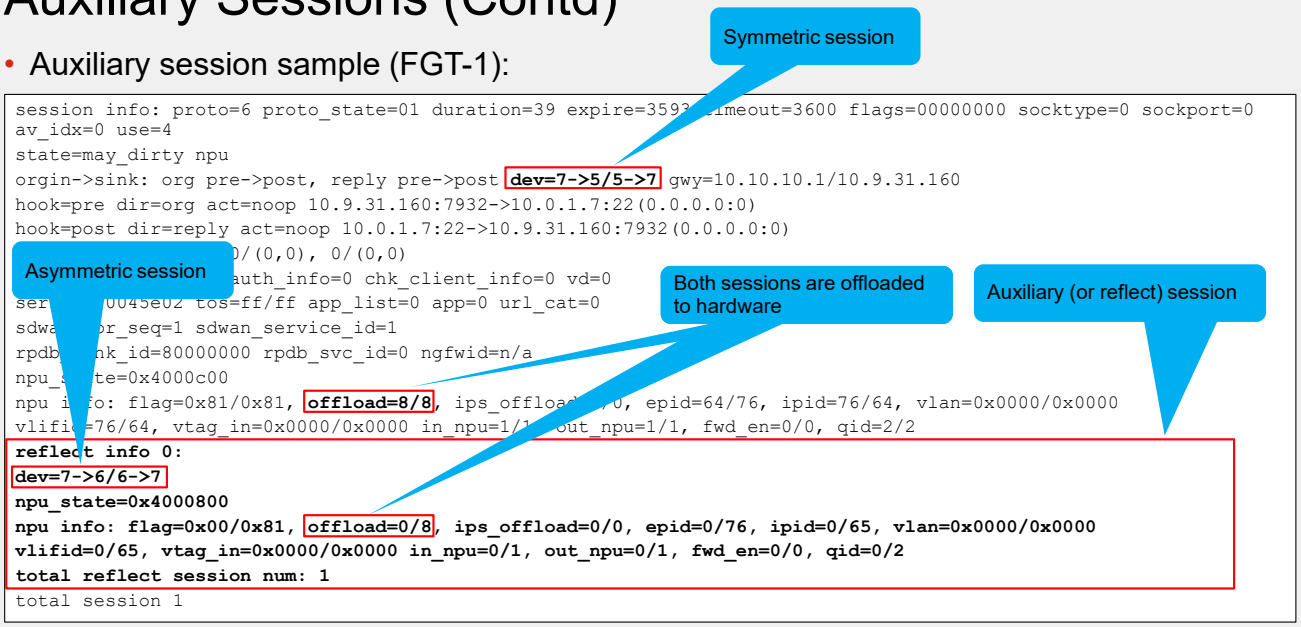
Auxiliary sessions are disabled by default. You can enable the feature per VDOM by running the CLI commands shown on this slide.

This slide also shows a debug flow sample taken on FGT-1 when an auxiliary session is created for an SSH connection. The debug flow shows a line containing the `auxiliary ses` message followed by the index number of the interfaces (or devices) that are used for creating the auxiliary session. In this example, although not displayed on this slide, the session was initially established symmetrically from `port3` to `port1`—interface index numbers `7` and `5`, respectively. After that, FortiGate received a reply packet on `port2`—the first line of the debug flow—which triggered the creation of the auxiliary session for the `port3` to `port2` direction—interface index numbers `7` and `6`, respectively.

DO NOT REPRINT
© FORTINET

Auxiliary Sessions (Contd)

- Auxiliary session sample (FGT-1):



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

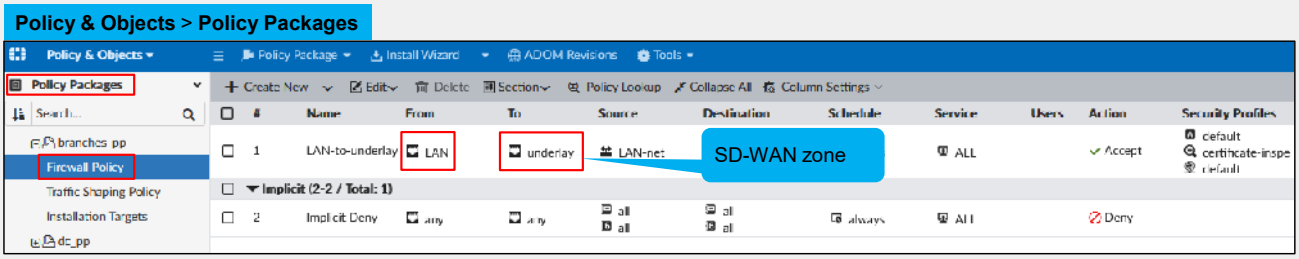
This slide shows the resulting main session and auxiliary session based on the debug flow output shown on the previous slide. Note that only relevant lines of the session are displayed.

Also note that an auxiliary session is not a separate session, but an extension of the main session used for asymmetric traffic. As a result, FortiGate can offload traffic for both sessions to hardware, which results in higher performance.

DO NOT REPRINT
© FORTINET

Firewall Policies

- Steered traffic must be allowed by a firewall policy
- Reference normalized interface for SD-WAN zones only
  - Simplified configuration
- Can't reference a member directly
  - Create zones with single members



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

You configure SD-WAN firewall policies in the same way as regular firewall policies, except that, when selecting an outgoing or incoming interface, you must reference a normalized interface that refer to an SD-WAN zone. When you reference a zone, you simplify the configuration by avoiding duplicate firewall policies.

You may need to reference a member in your firewall policy because you want to apply different action on the traffic flowing through that member, such as applying different security profiles and NAT settings. Because you can't reference members in a firewall policy, a workaround is to place a single member in a separate zone, and then reference that zone in the firewall policy.

The example on this slide shows a firewall policy named **LAN-to-underlay** that references the **underlay** zone, which contains port1 and port2 as members. As a result, DIA traffic steered over port1 or port2 is allowed by FortiGate provided it matches the firewall policy criteria, and it passes the security inspection configured on the policy.

DO NOT REPRINT
© FORTINET

Firewall Policy Changes and Sessions

- Firewall policy changes can lead to high CPU utilization
- Select which sessions in VDOM are flagged as `dirty` (default = `check-all`):

```
VDOM level
config system settings
    set firewall-session-dirty < check-all | check-new | check-policy-option >
end
```

  - `check-all`: All sessions are flagged as dirty
  - `check-new`: New sessions are flagged as dirty. Existing sessions are not affected.
  - `check-policy-option`: Follow firewall policy-level configuration
- Firewall policy-level configuration (default = `check-all`):

```
Policy Level
config firewall policy
    edit <id>
        set firewall-session-dirty < check-all | check-new >
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

40

When you make a change to a firewall policy, all existing sessions with the flag "`may-dirty`" change to status "`dirty`". FortiGate then performs a firewall policy lookup for the first packet received for each existing session; it updates session table entry and reset the flag to "`may-dirty`". Then, it processes the next packets based on the new firewall settings.

If the firewall handles a huge number of sessions, flagging all sessions as `dirty`, and performing a firewall policy lookup for the sessions may result in high CPU utilization. To prevent this, you can configure FortiGate to flag only new sessions as `dirty` by setting `firewall-session-dirty` to `check-new`. The result is that FortiGate evaluates only new sessions against the new firewall policy configuration.

The `firewall-session-dirty` setting is available on the VDOM and firewall policy levels. It is set to `check-all` by default, which instructs FortiGate to flag all sessions as `dirty`. To instruct FortiGate to follow the firewall policy-level setting, you must set the VDOM-level setting to `check-policy-option`. Note that the firewall policy-level setting is available only if the VDOM-level setting is set to `check-policy-option`.

DO NOT REPRINT
© FORTINET

Firewall Policy Changes and Sessions (Contd)

- Session is flagged as `persistent`

```
config system settings
    set firewall-session-dirty check-new
end

# diagnose sys session list
session info: proto=6 proto_state=01 duration=3 expire=3598 timeout=3600 flags=00000000 socktype=0 sockport=0
av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log npu f00 persistent
statistic(bytes/packets/allow_err): org=3369/19/1 reply=3797/17/1 tuples=2
tx speed(Bps/kbps): 1011/8 rx speed(Bps/kbps): 1140/9
orgin->sink: org pre->post, reply pre->post dev=7->19/19->7 gwy=100.64.1.1/10.0.1.101
hook=pre dir=org act=noop 10.0.1.101:47774->10.1.0.7:22(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.0.7:22->10.0.1.101:47774(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
serial=000f884f tos=ff/ff app_list=0 app=0 url_cat=0
sdwan_mbr_seq=3 sdwan_service_id=2
rpdb_link_id=ff000002 ngfwid=n/a
npu_state=00000000
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

41

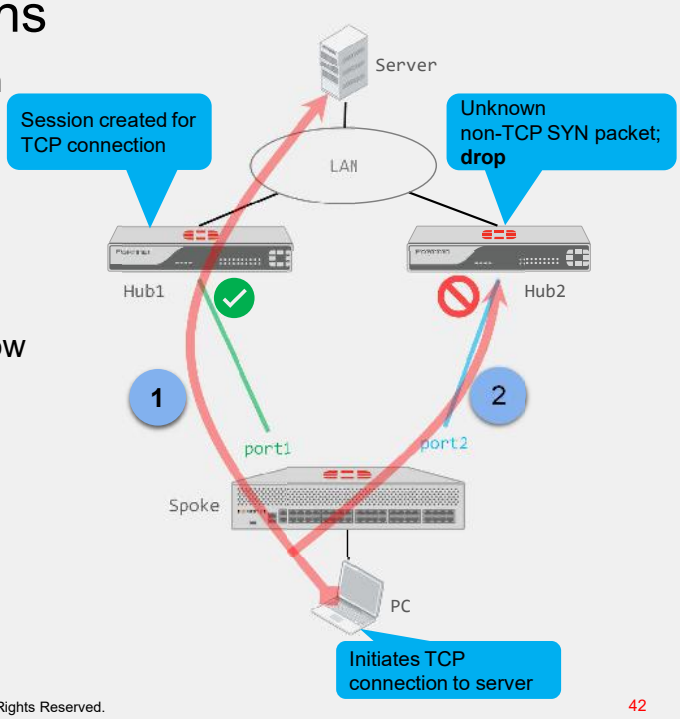
This slide shows the details of an established SSH session when `firewall-session-dirty` is set to `check-new`.

Note the presence of the `persistent` flag in the session, and the absence of the `may_dirty` flag, which indicates that FortiGate does not perform a new firewall policy lookup if there is a configuration change.

DO NOT REPRINT
© FORTINET

Cross-FortiGate TCP Sessions

- SD-WAN can steer traffic of an active session to another FortiGate
  - Example: Dual-hub and ADVPN topologies
- Cross-FortiGate TCP sessions are dropped because:
  - Packet is not TCP SYN and session is unknown to new FortiGate
- Enable `tcp-session-without-syn` to allow non-TCP SYN packets
  - Enable globally at VDOM level
  - Refine per firewall policy



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

You can configure SD-WAN to steer active sessions to a different member following a route or SLA change that affects the session. If the session in question is a TCP session, the session is dropped on the remote end if the new selected member connects to a different FortiGate device because of its stateful firewall inspection feature. Sessions that are routed across two different FortiGate devices are called cross-FortiGate sessions and are often seen in SD-WAN topologies like the one shown on this slide—dual-hub—or when using auto-discovery VPN (ADVPN). You will learn more about ADVPN in another lesson.

In the topology shown on this slide, the PC is connected behind the spoke FortiGate, and communicates using TCP with the server, which is behind the two hub FortiGate devices. If port1 on Spoke is the best member for steering traffic from the PC to the server, the packets are forwarded to Hub1. Hub1 also creates a session for the TCP connection. Then, if port2 becomes the best member while the connection is still active, Spoke starts forwarding the packets to Hub2. However, Hub2 drops the packets because they are not the initial TCP SYN packet of a three-way handshake, nor do they match any of the existing sessions.

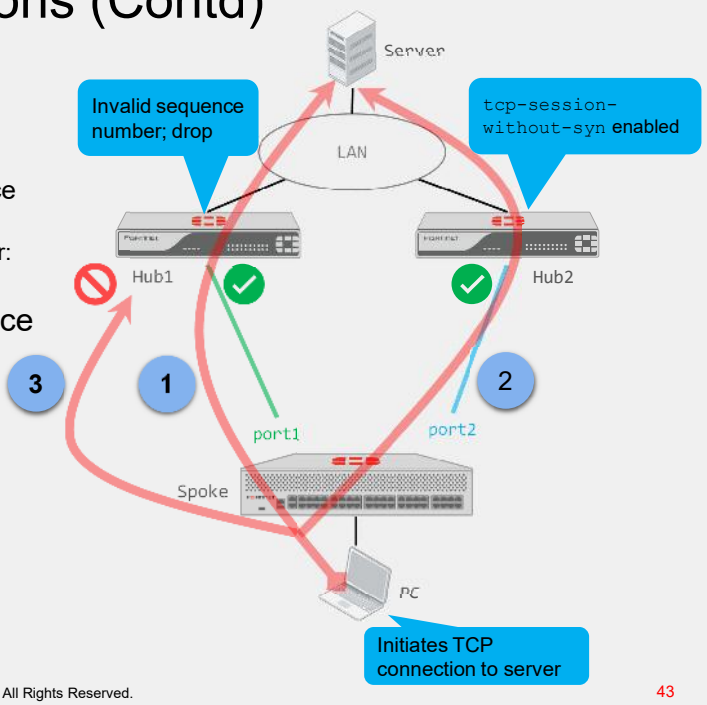
You can solve the previous issue by enabling the `tcp-session-without-syn` setting. When you enable `tcp-session-without-syn`, FortiGate accepts unknown TCP sessions whose first packet isn't a TCP SYN, provided there is a firewall policy that allows the traffic. For the use case described here, enabling `tcp-session-without-syn` on Hub2 would result in the connection staying up after the packets are forwarded to Hub2.

Note that you must first enable `tcp-session-without-syn` at the VDOM level under `config system settings` and then allow the feature for each firewall policy that requires it.

DO NOT REPRINT
© FORTINET

Cross-FortiGate TCP Sessions (Contd)

- FortiGate checks TCP sequence numbers
- If TCP session is forwarded back to original FortiGate:
  - Packets are dropped because of invalid sequence numbers
    - If old session is still active (default expiration timer: 3600 seconds)
- Disable `anti-replay` to skip TCP sequence number check
  - Enable globally at VDOM level
  - Refine per firewall policy



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

43

Continuing with the same dual-hub topology, assume that the administrator enabled `tcp-session-without-syn`, and that the TCP connection was accepted on Hub2 after the routing change on Spoke. Also assume that the connection is still active and continues to be routed to port2, and that no more than 60 minutes have passed since the routing change.

If port1 becomes again the best member, Spoke starts forwarding the packets back to Hub1. Because Hub1 still has a session for the TCP connection in its session table—the default expiration timer for established TCP sessions is 3600 seconds—the packets match the existing session but then get dropped by Hub1 because of invalid sequence numbers. The reason is that FortiGate also checks the sequence numbers of TCP packets. As a result, the sequence numbers of packets received by Hub1 do not match the expected values based on the latest session state recorded by Hub1, which happened just before the routing change to Hub2.

You can solve the previous issue by disabling the `anti-replay` setting. When you disable `anti-replay`, FortiGate doesn't check sequence numbers of TCP packets. As a result, Hub1 accepts the packets after the session is routed back to Hub1.

Note that you must first enable `anti-replay` at the VDOM level under `config system settings` and then allow the feature for each firewall policy that requires it.

DO NOT REPRINT
© FORTINET

Cross-FortiGate TCP Sessions (Contd)

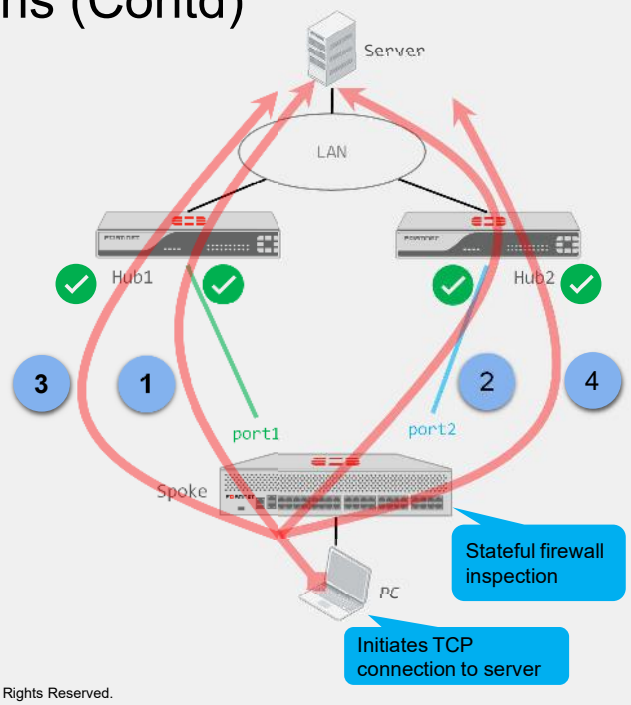
- Enable `tcp-session-without-syn` per VDOM

```
config system settings
    set tcp-session-without-syn enable
end
```

- Enable `tcp-session-without-syn` and disable `anti-replay` per policy:

```
config firewall policy
    edit <id>
        set anti-replay disable
        set tcp-session-without-syn all
    next
end
```

Setting is hidden if it is disabled on the VDOM



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

44

You can enable `tcp-session-without-syn` and disable `anti-replay` using the FortiGate CLI, as shown on this slide. The default value for each setting are `disable` and `enable`, respectively.

You must configure both settings on the firewall policy level. For the `tcp-session-without-syn` setting, you must first enable the setting per VDOM.

Continuing with the dual-hub topology example, the administrator should enable `tcp-session-without-syn` and disable `anti-replay` on both Hub1 and Hub2. The reason is that active TCP sessions could be routed back and forth between the hubs. Also, although the suggested configuration changes disable stateful firewall inspection for TCP traffic on Hub1 and Hub2, this shouldn't represent a security concern because Spoke still performs the stateful firewall inspection for the TCP traffic.

DO NOT REPRINT
© FORTINET

Review

- ✓ Describe policy route operation, actions, and FIB validation
- ✓ Identify key routing principles in SD-WAN
- ✓ Understand the route lookup process
- ✓ Configure static routes for zones
- ✓ Describe how to use BGP to route SD-WAN traffic
- ✓ Understand BGP options for SD-WAN
- ✓ Understand routing options for dual-hub topologies
- ✓ Understand common session flags
- ✓ Describe session reevaluation triggers
- ✓ Control routing changes in SNAT sessions

FORTINET Training Institute

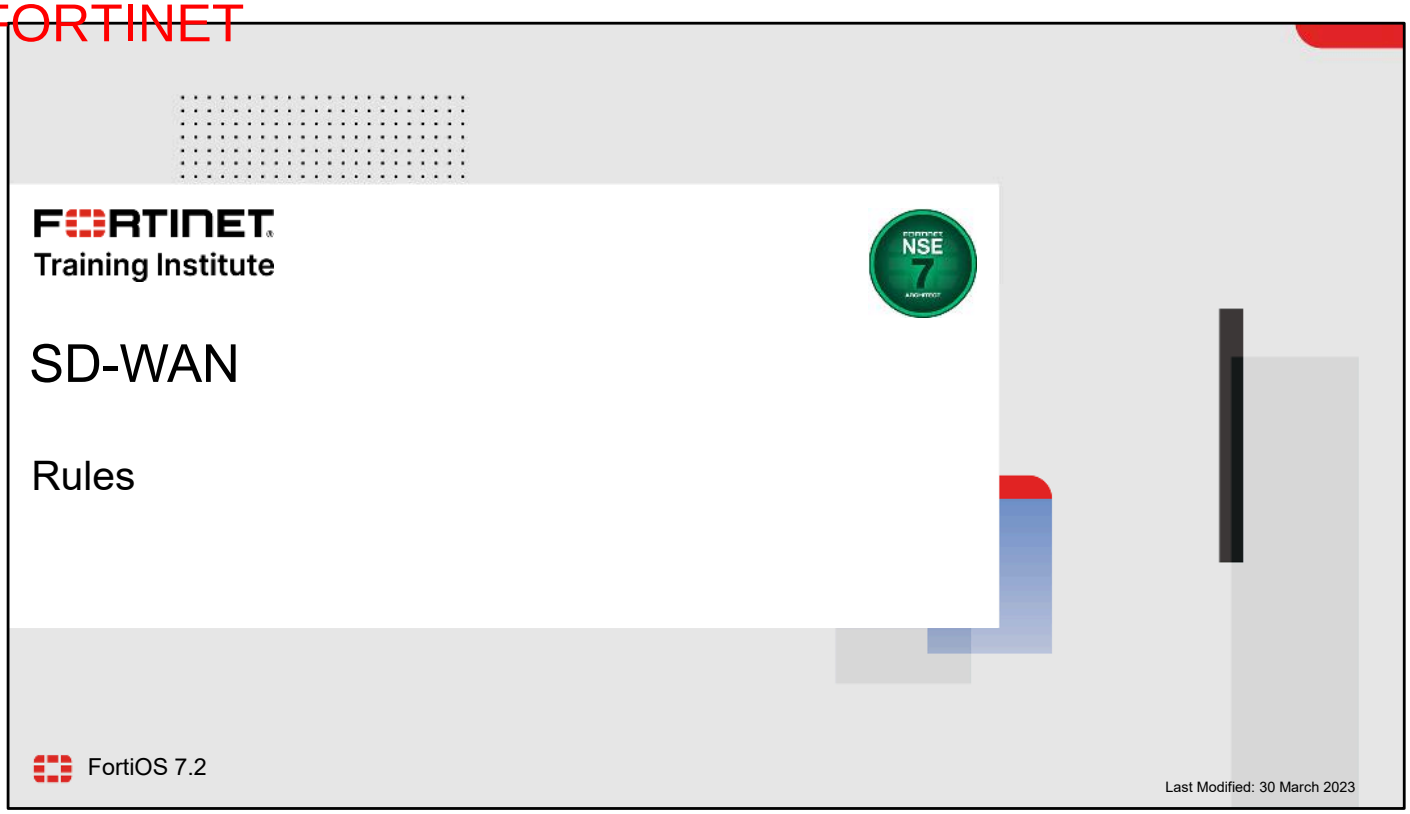
© Fortinet Inc. All Rights Reserved.

45

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

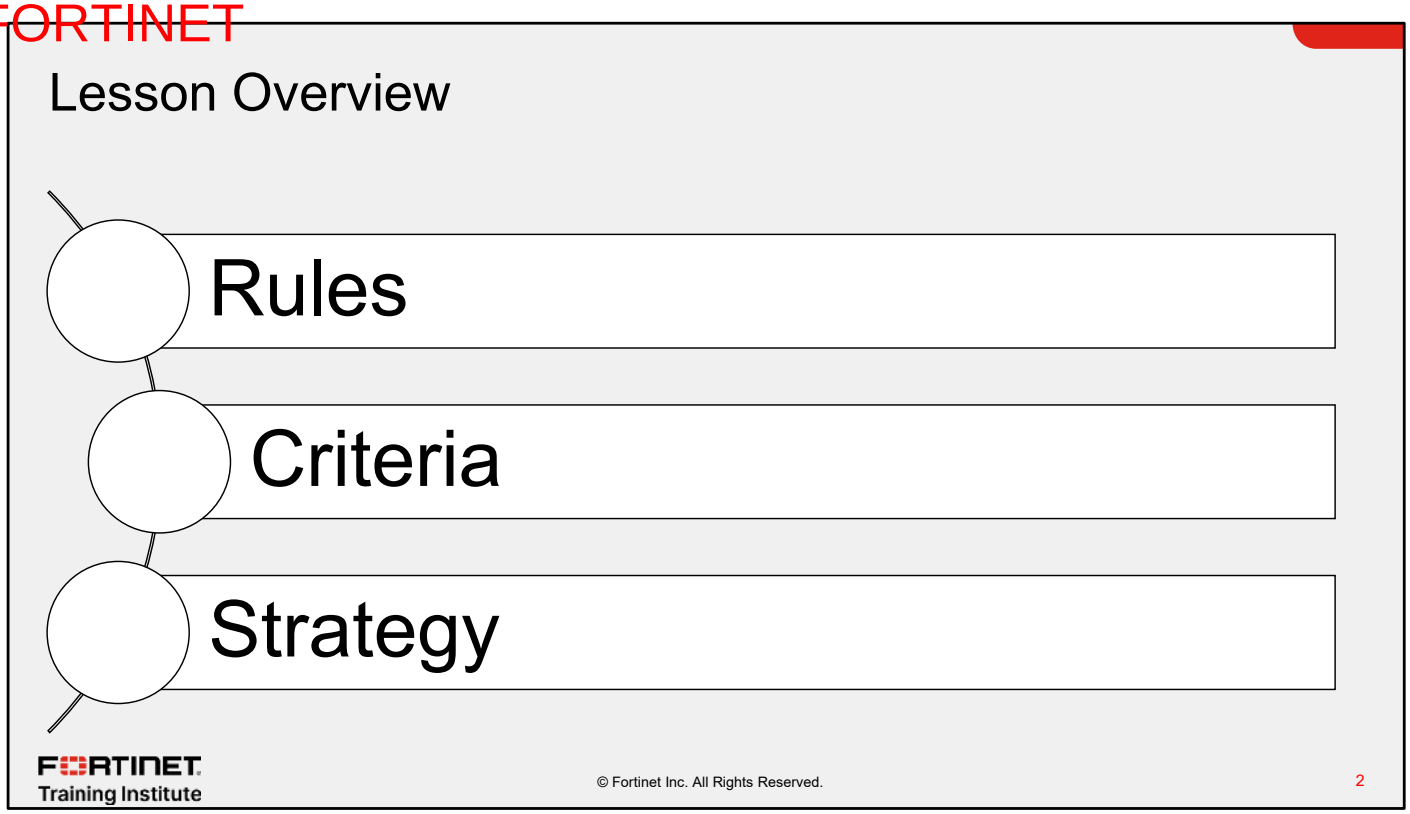
By mastering the objectives covered in this lesson, you learned about the routing process in SD-WAN, member static route configuration, BGP parameters for SD-WAN, and how to control session reevaluation caused by route, interface, or firewall policy changes.

DO NOT REPRINT
© FORTINET



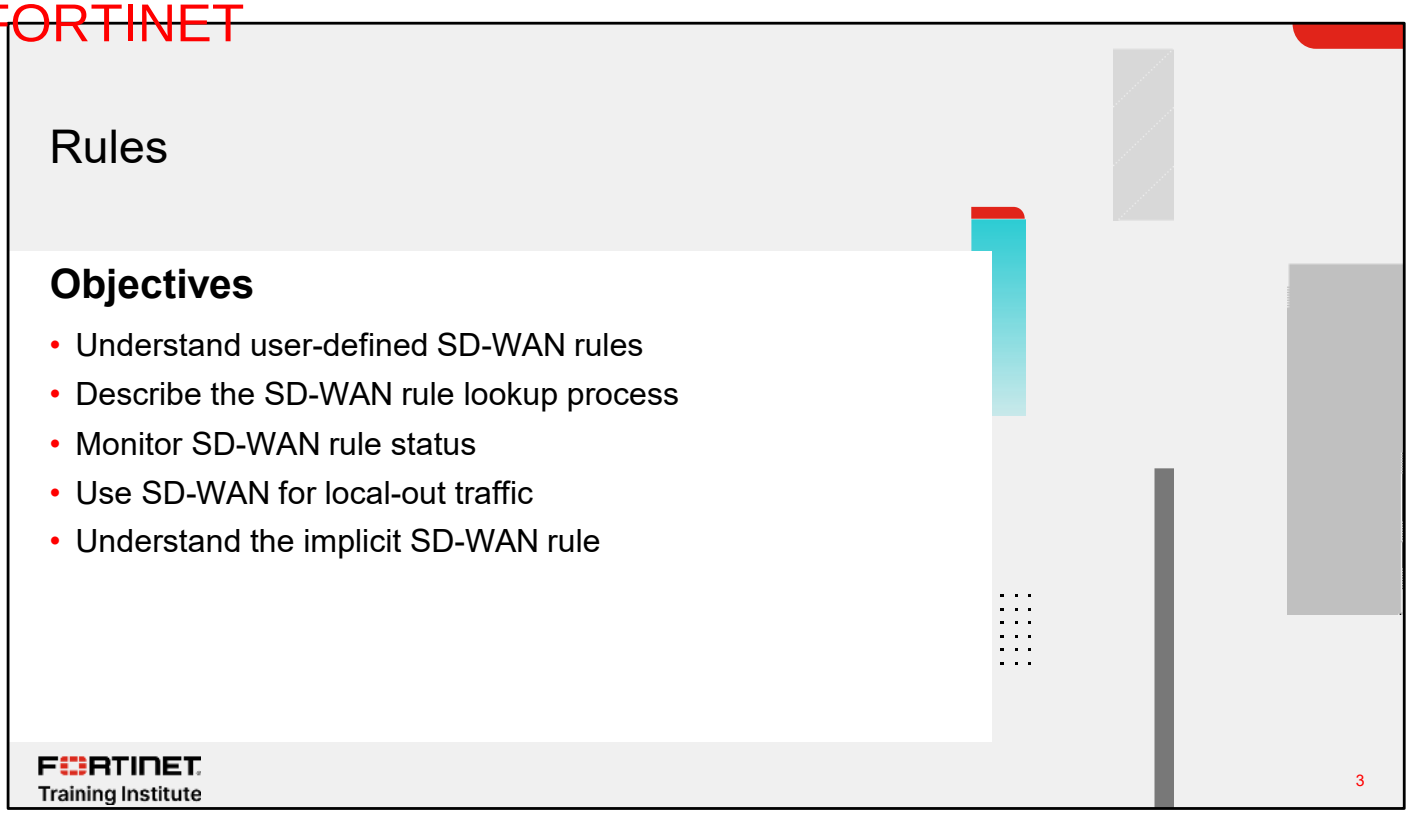
In this lesson, you will learn how SD-WAN rules work based on the criteria and strategy in use.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



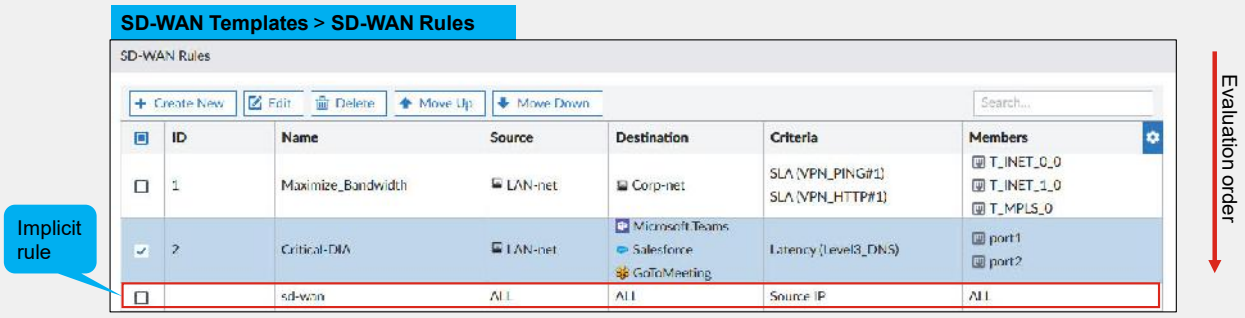
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding SD-WAN rules, you will know the fundamentals of SD-WAN rule configuration and operation.

DO NOT REPRINT
© FORTINET

Rules

- Describe administrator SD-WAN choices
- Two configuration sections:
  - Matching traffic criteria:
    - Define pattern, ISDB, and application to match
  - Forward policy:
    - Select preferred egress members and zones
    - Strategy
    - Required member performance metrics
- Evaluated in descending order:
  - Rules are used to steer traffic
    - Firewall policy required
  - Implicit rule
    - Used if user-defined rules are not matched
    - Traffic is usually load balanced



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

SD-WAN rules combine traffic matching criteria and traffic steering preferences. They describe the administrator choices related to the SD-WAN solution and the software-defined aspect of it.

You first define the criteria of the application or traffic to match. Then, you indicate the forward policy to follow for steering traffic across one or more members and zones, including the strategy to apply and the performance metrics to determine the preferred members.

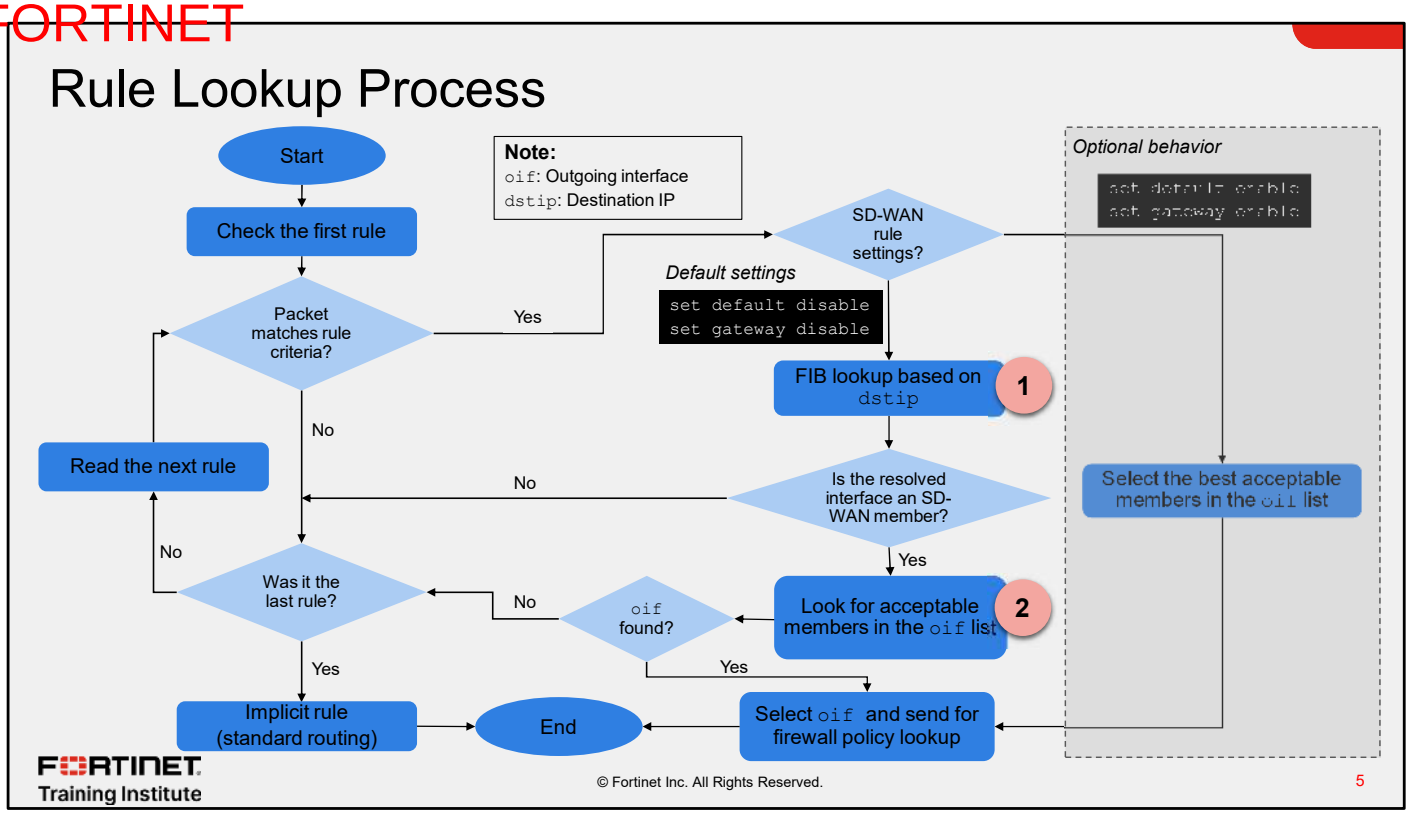
Preferred members are the best alive members based on the strategy in use. FortiGate then uses the preferred members—provided they are acceptable—to steer traffic. For all strategies except **Maximize Bandwidth (SLA)**, FortiGate always chooses a single member to steer traffic.

SD-WAN rules are evaluated in the same way as firewall policies: from top to bottom, using the first match. However, unlike firewall policies, SD-WAN rules are used to steer traffic, and *not* to allow traffic. That is, you must also configure corresponding firewall policies to allow the steered traffic.

If none of the user-defined SD-WAN rules are matched, then the implicit rule is used. The implicit rule instructs FortiGate to perform standard routing on traffic. Because SD-WAN deployments usually have multiple routes to the same destination—that is, ECMP routes—then traffic that matches the implicit rule is usually load balanced across multiple SD-WAN members.

The example on this slide shows a rule named **Critical-DIA**, which is used to steer traffic for direct internet access (DIA). The rule steers, **Microsoft Teams**, **Salesforce** and **GoToMeeting** traffic to the member with the lowest latency, between **port1** and **port2**. Note that only the most significant parts of rule configuration are shown on the output.

DO NOT REPRINT
© FORTINET



This slide shows the SD-WAN rule lookup process. SD-WAN rules are essentially policy routes. Like regular policy routes, SD-WAN rules are checked from top to bottom (first match). For each rule, FortiGate maintains an outgoing interface (`oif`) list. The `oif` list sorts the configured members by preference based on the strategy in use. The members that are placed first in the list have higher preference for steering traffic. FortiGate starts the lookup process by comparing the packet against the rule matching criteria. If the packet doesn't match the criteria, FortiGate moves on to the next rule, and so on, until finding a match. When there is a match, FortiGate proceeds as follows:

1. If both `default` and `gateway` settings are disabled—default values—FortiGate performs a FIB lookup for the packet destination IP (`dstip`). If the resolved interface for the FIB best-match isn't an SD-WAN member, then FortiGate moves on to the next rule. This behavior follows the key routing principle: *By default, SD-WAN rules are skipped if the best route to the destination isn't an SD-WAN member.*

2. If the resolved interface is an SD-WAN member, then FortiGate looks for one or more acceptable members in the `oif` list, starting from the first member in the list. An acceptable member is an alive member that has a route to the destination. This behavior follows the key routing principle: *By default, SD-WAN rules are skipped if none of the configured members in the rule have a valid route to the destination.*

If FortiGate finds an acceptable member, it forwards the packet to that member—then firewall policy check occurs— and the rule lookup process ends. Otherwise, FortiGate moves on to the next rule. If all rules are skipped, then FortiGate routes the packet using standard routing, hence the key routing principle: *The implicit SD-WAN rule equals standard FIB lookup*.

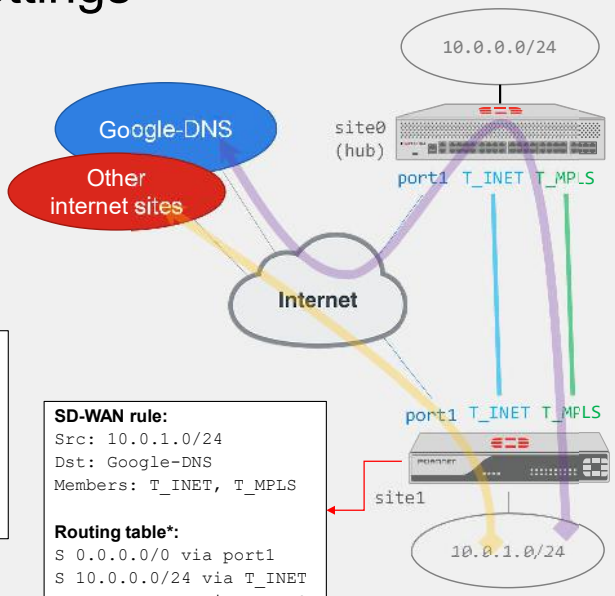
Note that when both `default` and `gateway` settings are enabled, FortiGate doesn't perform steps 1 and 2. Instead, it selects the best member or members—depending on the rule mode in use—in the `oif` list.

DO NOT REPRINT
© FORTINET

Default and Gateway Rule Settings

- Enable `default` to skip best route check for SD-WAN rules
- Enable `gateway` to skip FIB lookup for SD-WAN rule
  - Uses gateway detected for member
- Configure `default` and `gateway` settings on the FortiGate CLI (default = `disable`):

```
config system sdwan
    config service
        edit <id>
            set gateway enable
            set default enable
        next
    end
end
```



* T_INET and T_MPLS are SD-WAN members, port1 isn't.

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

By default, SD-WAN behaves according the following two key routing principles:

1. SD-WAN rules are skipped if the best route to the destination isn't an SD-WAN member.
2. SD-WAN rules are skipped if none of the configured members in the rule have a valid route to the destination.

Whenever possible, you should prefer the default behavior and configure a route to the destination (default or more specific). Nevertheless, for some specific scenario, you can change the behavior described in 1 by enabling the `default` setting in an SD-WAN rule. This instructs FortiGate to skip the best route check. You can also change the behavior described in 2 by enabling the `gateway` setting in an SD-WAN rule. This instructs FortiGate to skip the FIB lookup, and instead use the gateway address that you either manually configured in the SD-WAN member settings or that was automatically detected for DHCP, PPPoE, and IPsec interfaces.

The example on this slide shows a remote internet access (RIA) deployment. The administrator wants to use SD-WAN to steer traffic destined to the Google-DNS internet service using RIA and standard routing for all other internet traffic. The overlays are SD-WAN members, the underlay isn't. To avoid non-Google-DNS traffic to be routed over the overlays, the administrator uses default routes for the underlay only. With `default` disabled, FortiGate skips the rules because the overlays are not the best route for Google-DNS. If you then enable the `default` setting, FortiGate ignores the best route check and therefore evaluates Google-DNS traffic against the rules. However, because the overlays don't have a valid route to the destination (Google-DNS), then FortiGate skips the rule for Google-DNS. To avoid this, you can enable the `gateway` setting to instruct FortiGate to skip the FIB lookup and use as gateway the one configured or detected for the member.

DO NOT REPRINT
© FORTINET

Viewing Rule Status

- View rule status on the FortiGate CLI:

```
# diagnose sys sdwan service

Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members(2):
    1: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x1), gid(0), cfg_order(1), local cost(0), selected
  Internet Service(1): SSH(4294837953,0,0,0,0 16060)
  Src address(1):
        10.0.1.0-10.0.1.255
```

Matching criteria and rule mode

Outgoing interface list

```
# diagnose firewall proute list
list route policy info(vf=root):

id=2130968578(0x7f040002) vwl_service=2(SSH-Corp) vwl_mbr_seq=4 3 dscp_tag=0xfc despite flags=0x0
tos=0x00 tos_mask=0x00 protocol=0 sport=0-65535 iif=0(any) dport=1-65535 path(2) oif=20(T_INET_1)
oif=19(T_INET_0)
source(1): 10.0.1.0-10.0.1.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(1): SSH(4294837953,0,0,0,0, 16060)
hit_count=10 last_used=2023-02-14 01:59:22
```

ID shown in debug flow output

Hit count and last time rule was hit

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

SD-WAN rules are dynamically updated based on the member status and performance. This means that the outgoing interface list can change over time, which is why it is useful to check their current status for troubleshooting purposes.

The best way to check the status of an SD-WAN rule is by running `diagnose sys sdwan service` on the FortiGate CLI. As shown on this slide, the output indicates the matching criteria in use, the rule mode, and the outgoing interface list. You will learn more about the matching criteria, rule modes, and how FortiGate determines the preferred members in this lesson.

In addition, because SD-WAN rules are essentially policy routes, you can also run the `diagnose firewall proute list` command to display the rule settings from a policy route standpoint. SD-WAN rules are assigned a high ID number. The ID displayed in the `diagnose firewall proute list` command output is the same ID displayed in the debug flow output when a packet matches a rule. The output also includes the outgoing interface list, except that the interface preference is sorted from left to right, as opposed to the `diagnose sys sdwan service` command, which sorts the interfaces from top to bottom.

For troubleshooting purposes, the output of the `diagnose firewall proute list` command also displays the rule hit count and the last time the rule was hit.

DO NOT REPRINT
© FORTINET

Choosing the Best Route as the Preferred Member

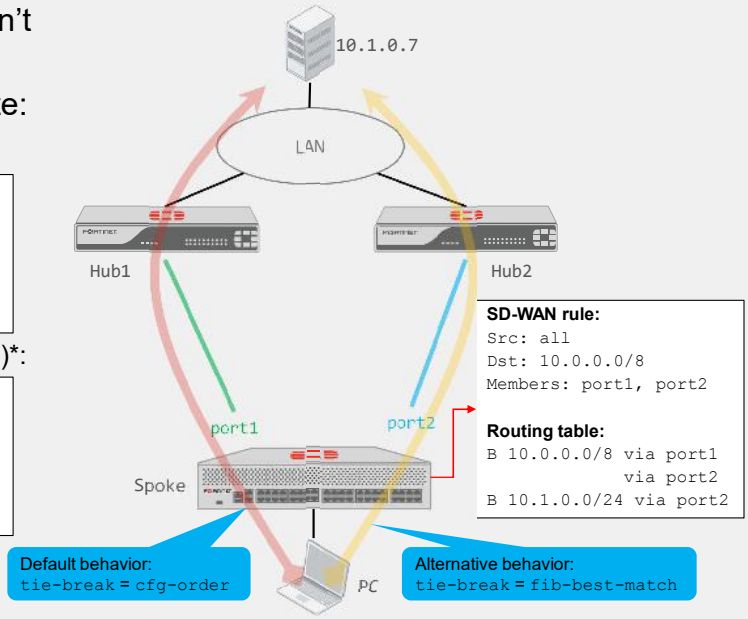
- By default, the preferred member doesn't have to be the best route
- Choose preferred member by best route:
  - Option 1: Rule level (default = `zone`)*

```
config system sdwan
    config service
        edit <id>
           set tie-break fib-best-match
        next
    end
end
```

  - Option 2: Zone level (default = `cfg-order`)*:

```
config system sdwan
    config zone
        edit <zone>
           set service-sla-tie-break fib-best-match
        next
    end
end
```

* Not supported in maximize bandwidth (SLA) rules



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

By default, during the SD-WAN rule lookup, FortiGate checks the member routes twice:

1. When looking for a matching rule: FortiGate skips SD-WAN rules if the best route to the destination isn't an SD-WAN member.
2. After the rule is matched: FortiGate skips SD-WAN rules if none of the configured members in the rule have a valid route to the destination.

From 2, you can deduce that the preferred member doesn't have to be the best route to the destination. The member just needs to have a valid route to the destination. But what if you want FortiGate to consider only members with the best route to the destination as candidates? Then you could control the preferred member election just by advertising a better route from the remote site.

The example on this slide shows a dual-hub and single-spoke topology. It includes a simplified version of the SD-WAN rule and routing table on the spoke. In the default behavior (`tie-break` set to `cfg-order`), FortiGate chooses port1 as the preferred member because it has a higher priority based on the configuration order. But if you set `tie-break` to `fib-best-match`, then FortiGate considers the best route to the destination (`10.1.0.7`) when choosing the preferred member in a rule. As a result, FortiGate chooses port2 as the only preferred member candidate because it has the most specific (best route) to the destination. Then, because port2 is the only candidate, it becomes the preferred member. This applies for all strategies. For example, in a lowest-cost strategy, FortiGate selects the preferred member only from members that have the best route to the destination.

This slide also shows how to configure FortiGate to consider the members with the best route as preferred members. You can apply the configuration at the zone level or, for additional granularity, at the rule level. By default, the rule level setting (`tie-break`) is set to `zone`, which instructs FortiGate to follow the zone-level setting (`service-sla-tie-break`), which in turn is set to configuration order (`cfg-order`) by default.

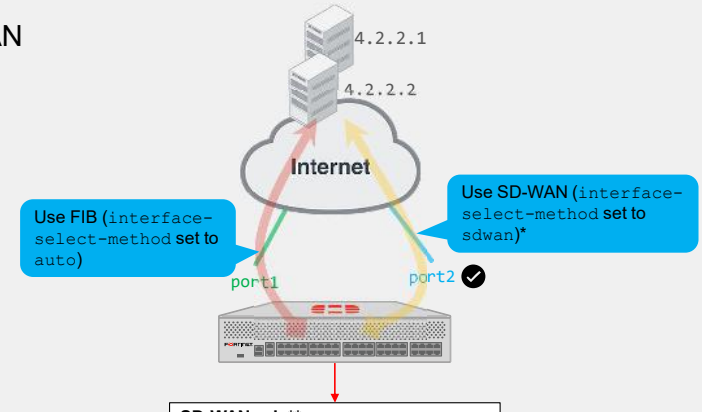
DO NOT REPRINT
© FORTINET

Use SD-WAN for Local-Out Traffic

- By default, local-out traffic doesn't use SD-WAN
- Enable SD-WAN for system DNS

```
config system dns
    set primary 4.2.2.1
    set secondary 4.2.2.2
    set interface-select-method sdwan
end
```

- `interface-select-method` is available on multiple features:
  - `config system dns | ntp | sflow | netflow`
  - `config system central-management`
  - `config system fortiguard`
  - `config user radius | ldap | fsso`
  - `config log fortianalyzer setting`
  - `config log syslogd setting`
- Enable SDWAN for ping and traceroute
  - `execute ping-options use-sdwan <yes|no>`
  - `execute traceroute-options use-sdwan <yes|no>`



SD-WAN rule:**
Src: all
Dst: all
Members: port2, port1

Routing table:
S 0.0.0.0/0 via port1
via port2 [10/0]

* SD-WAN routing principles apply
** SD-WAN rule source address must match local IP

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

Local-out traffic is defined as the traffic initiated by FortiGate, usually for management purposes. For example, when you ping a device from FortiGate, that's local-out traffic. When FortiGate connects to FortiGuard to download the latest definitions, that's also local-out traffic.

By default, local-out traffic uses the FIB to determine the best route for the destination. However, you can configure FortiGate to use SD-WAN rules for local-out traffic. For that, you must set `interface-select-method` to `sdwan` on the individual feature, so the feature uses SD-WAN to determine the best member to route the traffic to.

This slide shows an example of the system DNS configuration that uses SD-WAN to route system DNS queries. Before enabling SD-WAN for system DNS queries, FortiGate uses port1 because it has the best route for the destination in the FIB (lowest route priority). After enabling SD-WAN for system DNS queries, FortiGate starts using port2 because it's the preferred member for the matching SD-WAN rule.

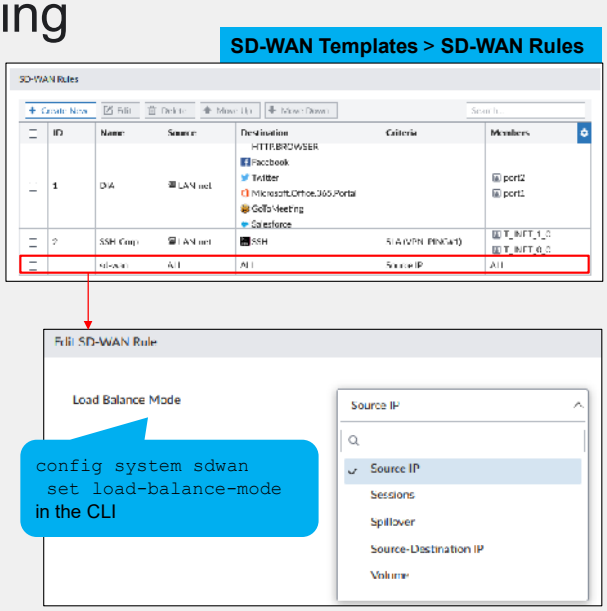
This slide also shows a list of some of the features that support the `interface-select-method` setting. For troubleshooting purposes, you can also use SD-WAN for the FortiOS ping and traceroute tools.

Note that when you enable SD-WAN for local-out traffic, the same SD-WAN routing principles apply. For example, for a member to be acceptable, the member must have a route in the FIB to the destination. In addition, note that the source address of the SD-WAN rule must match the local IP address used by FortiGate for the local-out connection. Usually this means setting the source address to all addresses to potentially match any SD-WAN member.

DO NOT REPRINT
© FORTINET

Implicit Rule and Load Balancing

- Implicit rule means standard FIB
- SD-WAN sites usually have ECMP routes
  - Result: sessions are load balanced
- Load balancing algorithms:
  - **Source IP** (default):
    - Traffic from a source IP is sent to the same member
  - **Sessions**:
    - The higher the member weight, the more sessions are sent to it
  - **Spillover**:
    - Send sessions to first member until spillover limit is reached, then send to next member
  - **Source-Destination IP**:
    - Traffic from a source IP to a destination IP address is sent to the same member
  - **Volume**:
    - The higher the member weight, the more traffic is sent to it



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

If none of the user-defined SD-WAN rules are matched, then the implicit rule is used. When a session matches the implicit rule, FortiGate performs a standard FIB lookup. Because SD-WAN sites usually have multiple routes to the same destination through multiple members, this results in the presence of ECMP routes in the FIB. By default, FortiGate load balances sessions that match ECMP routes.

To load balance sessions, FortiGate uses the algorithm configured in the implicit rule:

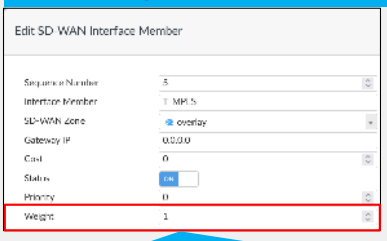
- **Source IP**: This is the default algorithm. FortiGate sends sessions sourced from an IP address to the same member. (CLI `source-ip-based`)
- **Sessions**: FortiGate load balances sessions across members based on the member weight. The higher the weight, the more sessions FortiGate sends to the member. (CLI `weight-based`)
- **Spillover**: FortiGate sends sessions to the interface of the first ECMP route until the bandwidth of the interface reaches the configured spillover limit. After the spillover limit is reached, FortiGate uses the interface of the next ECMP route. (CLI `usage-based`)
- **Source-Destination IP**: FortiGate sends sessions with the same source and destination IP address pair to the same member. (CLI `source-dest-ip-based`)
- **Volume**: FortiGate load balances sessions across members based on the measured interface volume and the member weight. Unlike **Sessions**, which doesn't consider the traffic of an interface, the **Volume** algorithm instructs FortiGate to track the cumulative number of bytes of each member and to distribute sessions based on the weight. The higher the weight, the higher the target volume of the interface and, as a result, the more traffic FortiGate sends to it. (CLI `measured-volume-based`)

DO NOT REPRINT
© FORTINET

Implicit Rule and Load Balancing (Contd)

- Weight configuration:

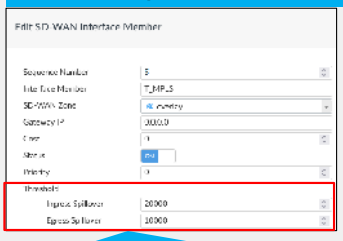
SD-WAN Templates > Interface Members



Displayed only when the load balance algorithm is **Sessions** or **Volume** (default = `1`)

- Spillover threshold:

SD-WAN Templates > Interface Members



Displayed only when the load balance algorithm is **Spillover** (default = `0`)

- Formula:

% sessions or traffic = member weight / (sum of all weights)

- If using dynamic routes, sessions are distributed equally, regardless of weight

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

When you set the implicit rule load balancing algorithm to **Sessions** or **Volume**, FortiManager assigns a default **Weight** of `1` to each member. You can change the weight by editing the member settings, as shown on this slide.

When using static routing with the implicit SD-WAN rule, FortiGate uses the weight to calculate the percentage of sessions or traffic sent to each member using the formula shown on this slide. For example, suppose that you configure two members—A and B—and you assign them a weight of 5 and 10, respectively. If you use **Sessions** as the load balancing algorithm, then out of 15 sessions, FortiGate sends 5 to member A—that is, 33.33% of sessions—and 10 to member B—that is, 66.67% of sessions. Similarly, if you use **Volume** as the load balancing algorithm, then FortiGate distributes sessions so member A handles 33.33% of the total *amount of traffic*, while B handles the rest (66.67%).

Also note that, with dynamic routing, when you use **Sessions** or **Volume** load balancing, FortiGate distributes sessions equally if ECMP routes are present. That is, FortiGate supports the weight setting only for static ECMP routes.

When you set the load balancing algorithm to **Spillover**, you should also adjust the **Ingress Spillover** and **Egress Spillover** settings of each member. Otherwise, FortiManager keeps the default value (`0`), which disables spillover check.

DO NOT REPRINT
© FORTINET

Implicit Rule and Load Balancing Accuracy

- **Sessions** distribution doesn't have to result in proportional bandwidth or volume
  - Traffic varies among sessions
- **Spillover** and **Volume** distribution may exceed the configured threshold and weight
  - Sessions are not flagged as dirty if threshold is reached
  - Sessions stay in the same member until expiration or forced route lookup (route in use is removed)
    - Including long-lived high-bandwidth sessions

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

It's important to understand the expected traffic distribution with the **Sessions**, **Spillover**, and **Volume** algorithms.

For **Sessions**, the weight assigned to members doesn't necessarily translate to the same bandwidth or volume observed on members. FortiGate load balances sessions across members based on the configured weight. Because the bandwidth and lifetime of each session are likely to be different, then the resulting bandwidth and volume are not proportional to the weight. For example, assume that FortiGate sends ten ICMP sessions to a member, and one FTP session to another member. Despite being one session, the bandwidth and volume of traffic transmitted over the member used for routing the FTP session is likely to be higher than the total bandwidth and volume observed on the member used for routing the ten ICMP sessions.

In addition, for **Spillover** and **Volume**, it's fine if the bandwidth exceeds the configured spillover thresholds, or if the amount of traffic on an interface exceeds the expected volume based on the configured weight. The reason is that unlike SLA changes in SD-WAN rules, which flag impacted sessions as dirty to force re-evaluation, sessions routed using **Spillover** and **Volume** are not flagged as dirty as a result of thresholds being reached. That is, sessions stay in the same member until they expire, or until they are subject to a forced route lookup caused by the removal of the route in use. This means that FortiGate usually doesn't change the routing of existing sessions distributed using **Spillover** and **Volume** algorithms, including that of long-lived high-bandwidth sessions that generate large amounts of traffic.

DO NOT REPRINT
© FORTINET

System Settings Algorithm vs. Implicit Rule Algorithm

- Both `v4-ecmp-mode` and `load-balance-mode` control the VDOM ECMP algorithm
  - `load-balance-mode` replaces `v4-ecmp-mode` when SD-WAN is enabled
- Differences:
  - `load-balance-mode` supports the volume algorithm, `v4-ecmp-mode` does not
  - `load-balance-mode` uses the weight defined under the SD-WAN member configuration, `v4-ecmp-mode` the weight defined in the static route
  - `load-balance-mode` uses the spillover thresholds defined under the SD-WAN member configuration, `v4-ecmp-mode` the spillover thresholds defined in the interface settings

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

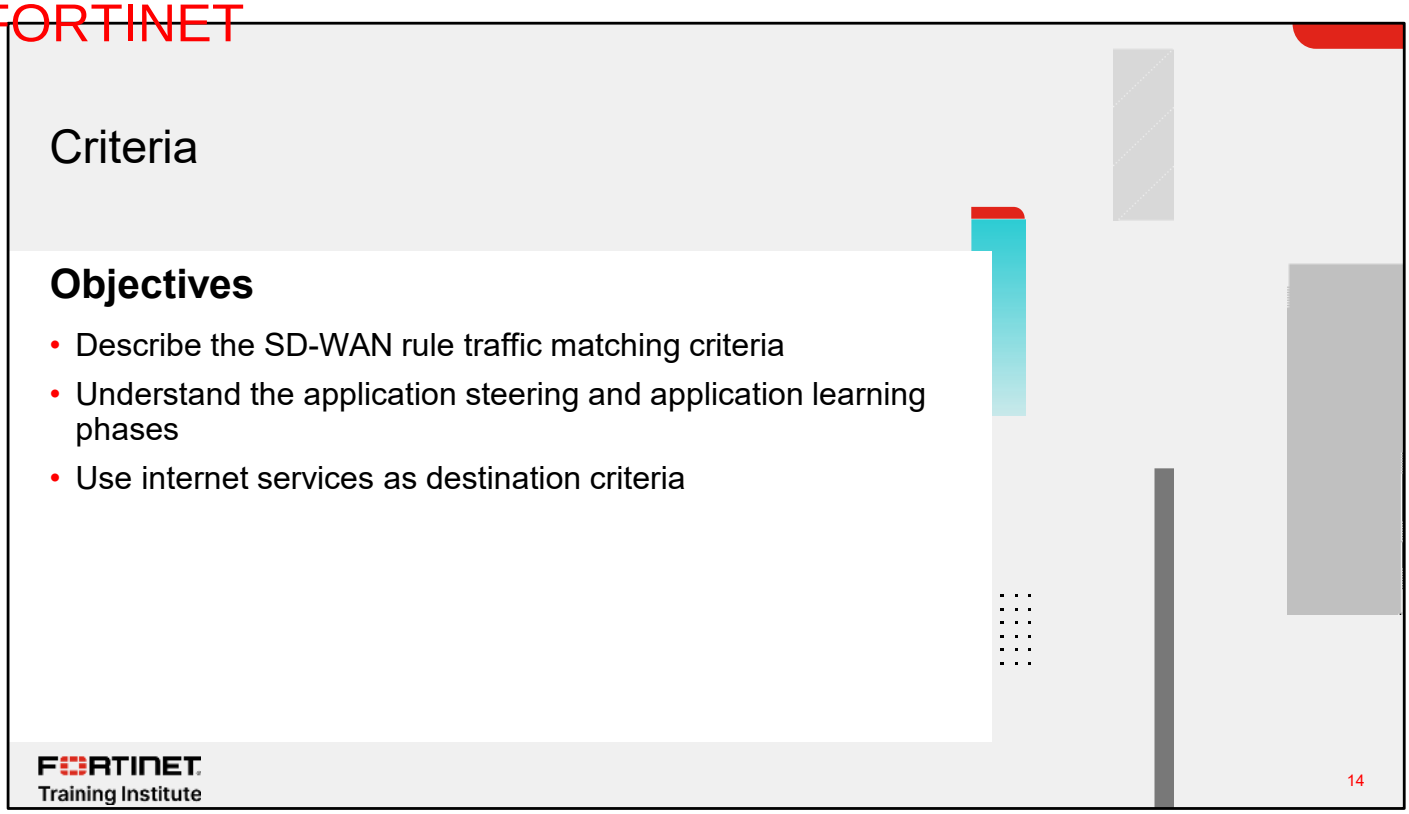
13

If you are familiar with ECMP, you probably know the `v4-ecmp-mode` setting available under `config system settings`. The `v4-ecmp-mode` setting defines the algorithm that FortiGate uses to load balance sessions that match ECMP routes in the VDOM.

However, when you enable SD-WAN on FortiGate, FortiOS hides the `v4-ecmp-mode` setting and replaces it with the `load-balance-mode` setting under `config system sdwan`. That is, after you enable SD-WAN, you now control the VDOM ECMP algorithm with the `load-balance-mode` setting.

There are some differences between the two settings. The main difference is that `load-balance-mode` supports the volume algorithm, and `v4-ecmp-mode` does not. In addition, the related settings such as weight and spillover thresholds are configured differently. That is, when you enable SD-WAN, the weight and spillover thresholds are defined on the SD-WAN member configuration. When you disable SD-WAN, the weight and spillover thresholds are defined on the static route and interface settings, respectively.

DO NOT REPRINT
© FORTINET



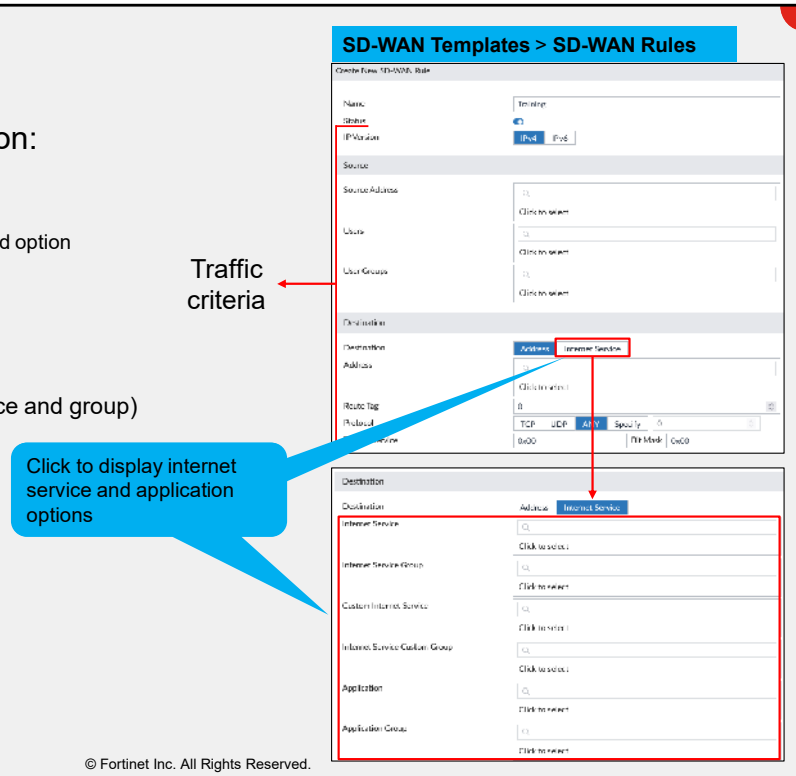
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding the supported SD-WAN rule traffic matching criteria, you will identify the best criteria to use for matching SD-WAN traffic in your deployment.

DO NOT REPRINT
© FORTINET

Criteria

- Rules can match traffic based on:
  - Source
    - IP address and interface
      - Source interface is an advanced option
    - Firewall user and user group
  - Destination:
    - IP address
    - Port number
    - Internet service (individual service and group)
    - BGP route tag
  - Application
  - IP protocol
  - ToS
- IPv4 and IPv6 support



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

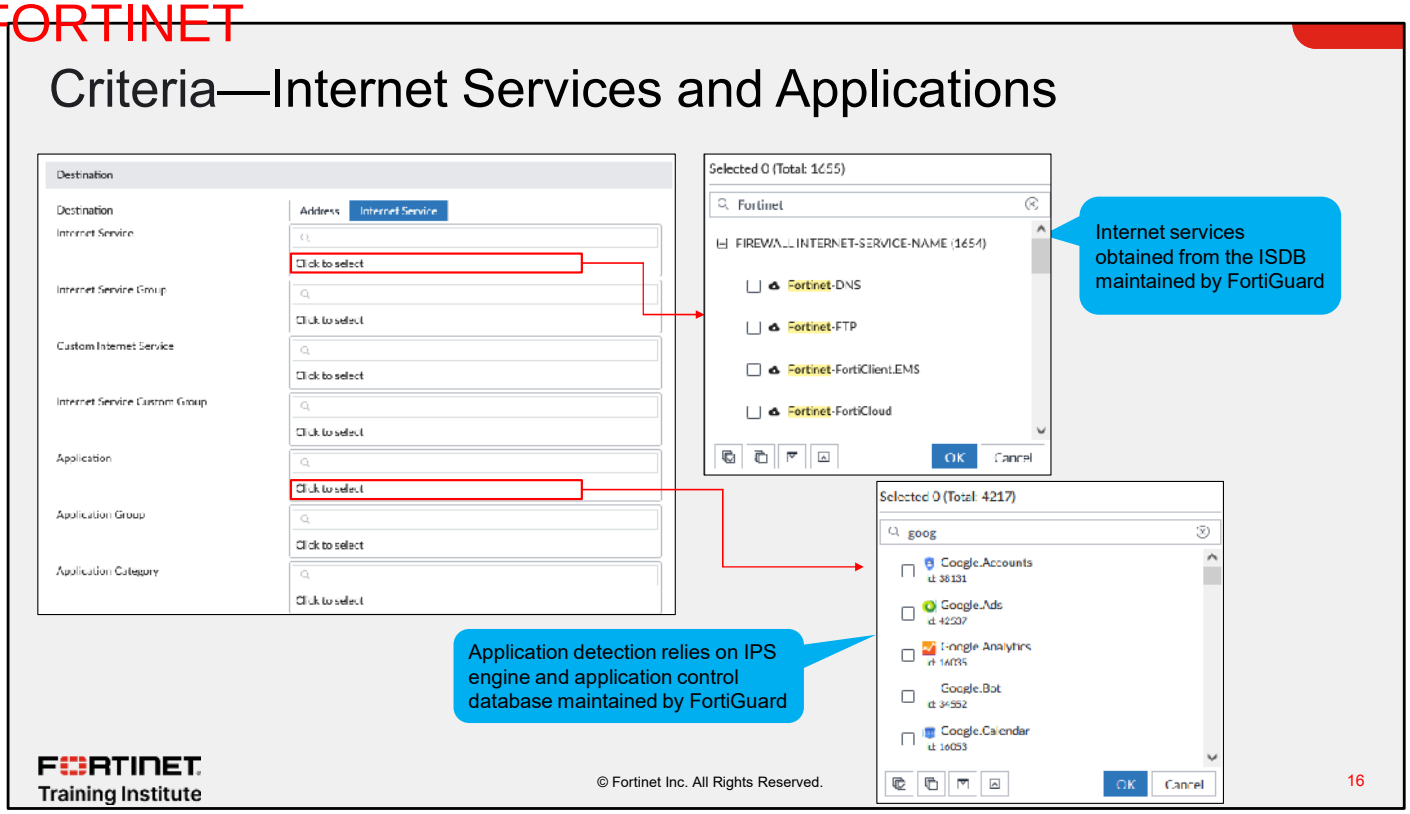
You can configure rules to match traffic based on the following criteria:

- Source IP address, source interface, firewall user, and firewall user group. The source interface option is available in **Advanced Options** section using the `input-device` and `input-device-negate` setting.
- Destination IP address, destination port number, internet service, and BGP route tag
- Application
- IP protocol
- Type of Service (ToS)

SD-WAN rules offer great flexibility for traffic matching. For example, you can match Netflix traffic sourced from specific authenticated users, or match the ICMP traffic—IP protocol 1—destined to a particular address.

The example on this slide shows the settings available when you create a new rule using FortiManager. Note that you must click **Internet Service** to see the internet service and application related options. Also note that both IPv4 and IPv6 protocols are supported, and that you can select individual internet services and applications, or groups of them.

DO NOT REPRINT
© FORTINET



You can configure a rule to match traffic based on the destination internet service and the application detected for the traffic. The destination internet service relies on the internet service database (ISDB), and application detection relies on the IPS engine and the application control database signature.

FortiGuard maintains both databases, and FortiGate periodically gets an updated copy by default. New versions of the IPS engine are also released by FortiGuard, but much less frequently than definitions and databases. When a new IPS engine version is released, FortiGate also downloads it automatically by default.

The example on this slide shows a subset of the internet services and applications available to choose from in a rule.

DO NOT REPRINT
© FORTINET

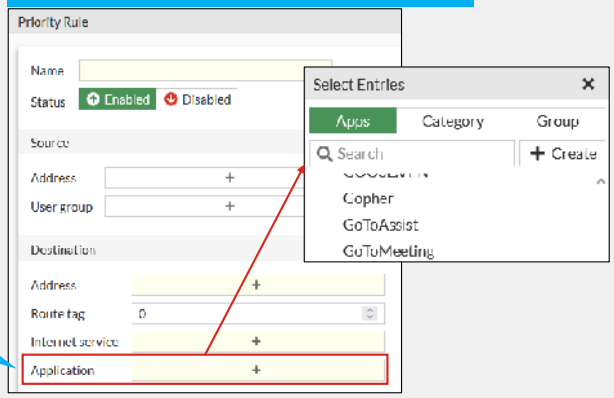
Criteria—Applications and Application Groups

- Rules can match traffic based on:
  - Applications
    - One or more application from predefined list
  - Application Group:
    - User defined group of applications
  - Application Category
    - Predefined group of applications per category
    - Business, Game, Proxy, Social.Media,…
- Application filter for SD-WAN rules
  - Visible by default on FortiManager
  - Hidden by default on FortiGate GUI

- Enable GUI visibility for application detection on FortiGate

```
config system global
    set gui-app-detection-sdwan enable
end
```

FortiGate: **Network> SD-WAN > SD-WAN Rules**



Visible after CLI activation

Training Institute

© Fortinet Inc. All Rights Reserved.

17

For application detection you can use applications from FortiGuard's predefined application list, create groups with those applications, or use application categories. Application categories group application per purpose, for example business, game, social media. You can also combine application group with specific applications.

Note that **Application**, **Application Group** and **Application Category** are always visible on the FortiManager menu —when you select **Internet Service** as the destination for the rule.

If you decide to configure SD-WAN rules directly from the FortiGate GUI, you need to enable GUI visibility for application detection with the CLI commands shown on this slide.

DO NOT REPRINT
© FORTINET

ISDB

- Maintained by FortiGuard
- Automatically retrieved by FortiGate
- Each internet service:
  - Is assigned an ID
  - Contains multiple entries indicating:
    - A range of public IP addresses
    - IP protocol
    - Destination port
  - Is loaded to the kernel and can be used in multiple features
  - Indicates the direction the entries are valid for

```
# diagnose internet-service id-summary
...
id: 1245187 name: "Fortinet-DNS"
id: 1245188 name: "Fortinet-Outbound_Email"
...

# diagnose internet-service id-summary 1245187
Version: 00007.02845
Timestamp: 202302182245
Total number of IP ranges: 1899363
Number of Groups: 23
Group(0), Singularity(90), Number of IP ranges(2985)
...
Group(22), Singularity(4), Number of IP ranges(9375)
Internet Service: 1245187(Fortinet-DNS)
Number of IP ranges: 499
Number of IP addresses: 55259
Singularity: 10
Icon Id: 19
Direction: dst
Data source: isdb
Country: 12 32 36 ...
Region: 55 132 159 169 ...
City: 615 679 818 1106 ...
```

Can be `dst` (destination), `src` (source), or `both`; rules support `dst` or `both`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

18

FortiGuard maintains the ISDB and releases updated versions frequently. FortiGate then downloads the most recent version automatically from FortiGuard servers.

Each service in the ISDB is assigned an ID that represents a well-known internet service such as Fortinet-FortiGuard, Google-ICMP, Facebook-Web, and so on. Each ID contains multiple entries, each of them indicating a range of public IP addresses, IP protocol, and destination ports associated with the service.

The internet service entries are loaded in the kernel. You can then reference an internet service in multiple parts of the configuration, such as SD-WAN rules and firewall policies.

Each internet service also indicates the direction of the traffic that the entries are valid for. The direction can be source, destination, or both. For SD-WAN rules, you can use only internet services that are valid for the destination direction, or both source and destination directions. This is because SD-WAN supports internet services for matching only the destination of a packet.

You can use the `diagnose internet-service id-summary` command to display a summary of the internet services in the ISDB. Add the internet service ID to the end of the command to view more details about that specific internet service.

DO NOT REPRINT
© FORTINET

ISDB (Contd)

- View address ranges, protocol, and ports of internet service:

```
# diagnose internet-service id 1245187
Internet Service: 1245187(Fortinet-DNS)
Version: 00007.02845
Timestamp: 202302182245
Number of IP ranges: 998
...
212.166.61.130-212.166.61.130 country(56) region(2060) city(615) blocklist(0x0) reputation(5),
popularity(5) domain(376) botnet(0) proto(6) port(53)
...
```

- Find internet service by IP, protocol, and port:

```
# diagnose internet-service info  root 6 53 212.166.61.130
Internet Service: 1245187(Fortinet-DNS) country(56 Belgium) region(2060 Walloon) city(615 Amay)
```

- Find internet service by IP:

```
# diagnose internet-service match root 212.166.61.130 255.255.255.255
Internet Service: 1245185(Fortinet-Web), matched entry num: 4, matched num: 4
Internet Service: 1245186(Fortinet-ICMP), matched entry num: 2, matched num: 2
Internet Service: 1245187(Fortinet-DNS), matched entry num: 2, matched num: 2
```

FORTINET Training Institute

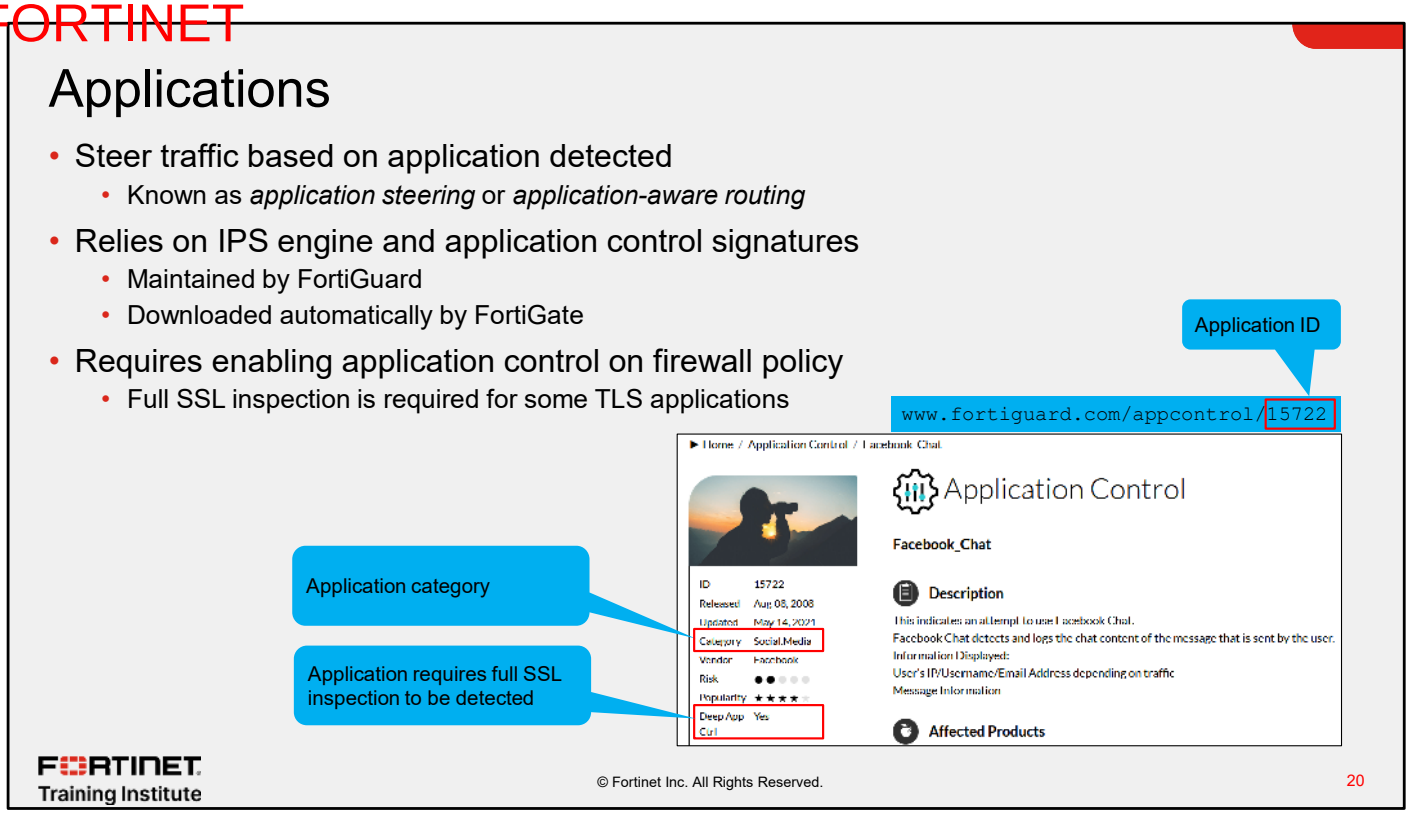
© Fortinet Inc. All Rights Reserved.

19

You can use the command `diagnose internet-service id` to display detailed information about an internet service. The output indicates the public IP ranges of the service, as well as the protocol and port used for each range.

You can also find the internet service that matches a given IP address, protocol, and port by using the `diagnose internet-service info` command. However, if you only know the IP address of the service you want to look up, you can run the `diagnose internet-service match` command instead. Note that the `diagnose internet-service info` and `diagnose internet-service match` commands are global commands and, therefore, require you to indicate the VDOM to perform the lookup on—`root` in the example.

DO NOT REPRINT
© FORTINET



You can configure rules to steer traffic based on the application detected by FortiGate. This is known as *application steering* or *application-aware routing.*

Application detection relies on the IPS engine, and the application control signatures maintained by FortiGuard. By default, FortiGate automatically downloads the new versions of the IPS engine and the application control database after they are released by FortiGuard.

For FortiGate to detect applications and, therefore, for SD-WAN to perform application steering, you must enable application control on the firewall policy that allows the SD-WAN traffic. You must also enable full SSL inspection if you want to detect applications that require it. For example, **Facebook_Chat**—a child application of Facebook—requires full SSL inspection but its parent application, Facebook, does not. Usually, FortiGate can detect a parent TLS application, like Facebook, without full SSL inspection but requires it for detecting child TLS applications like Facebook_Chat. This is because FortiGate must decrypt the TLS traffic to identity the child application.

You can identify whether an application requires full SSL inspection by checking the application information page on FortiGuard. A quick way to access the application information page is by adding the application ID to the end of `www.fortiguard.com/appcontrol/`, as shown on this slide.

This page also indicates the category in which the application is classified. This is the category that the application category filter uses.

DO NOT REPRINT
© FORTINET

Application Learning

- Application cache
  - 3-tuple: destination address, destination protocol, destination port
  - Dynamically populated from traffic
  - Multiple 3T per application ID
  - 8 hours timeout
- Learning phase
  - No entry in cache for application *App1*
  - Session not identified as application *App1*
  - Session routed by SD-WAN rule and allowed policy without application *App1* match
  - IPS engine analyze the flow to populate application cache
- Known application
  - Received flow matches 3-tuple in cache
  - Session identified as application *App1*
    - Can match SD-WAN with application steering
    - Can match policy rule with application detected

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

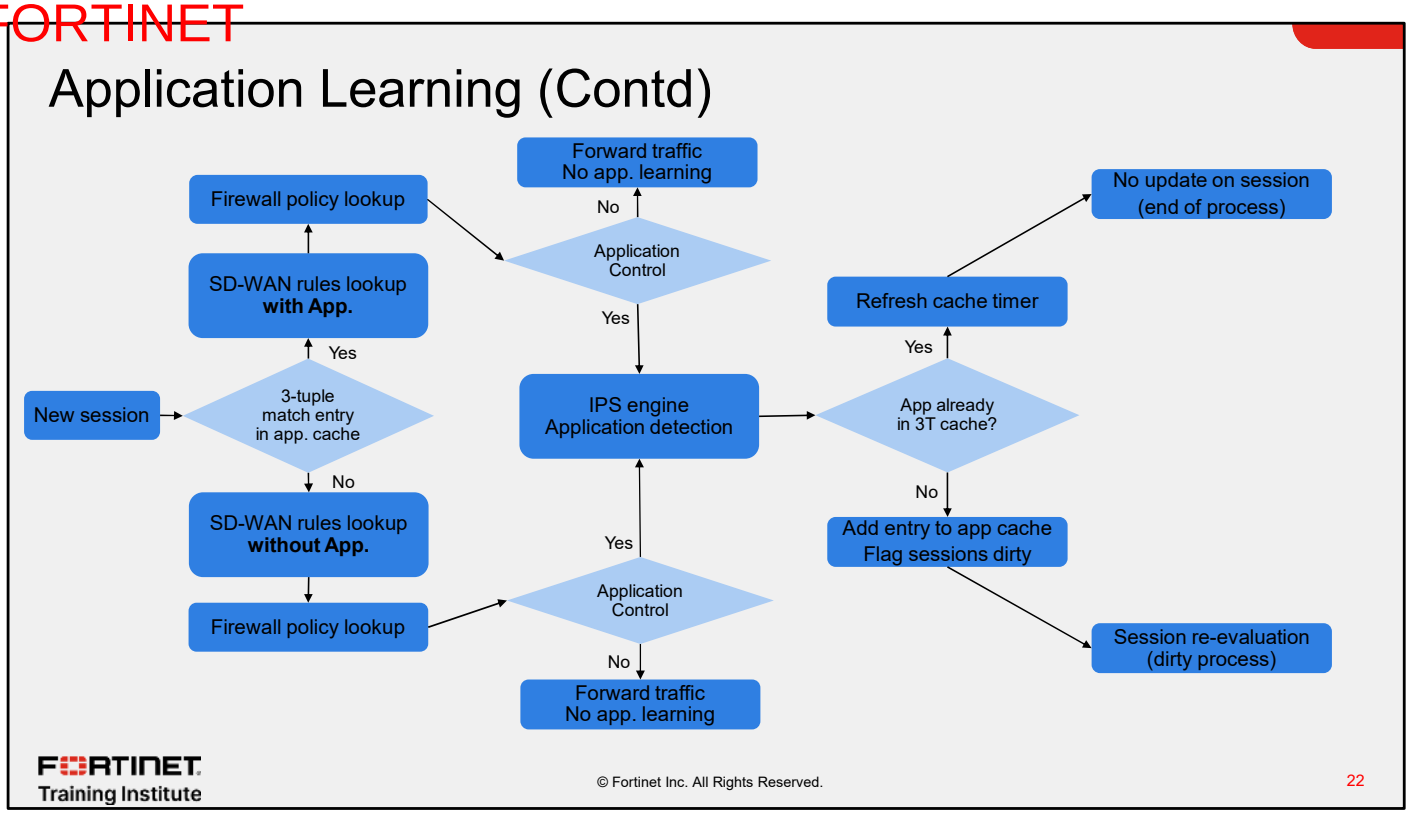
FortiGate populates its application cache dynamically by learning from the traffic that flows through it.

Each time FortiGate allows a traffic flow through with a policy that has application control enabled, it sends the traffic to the IPS engine for analysis. Once FortiGate detects an application, it populates the application cache with a 3-tuple entry. The cache consists of a list of application IDs associated with destination addresses, protocol, and port. Note that the application cache can contain multiple 3-tuples for one application to accommodate multiple servers handling the same application. Entries remain in the cache for 8 hours if they are not matched.

When FortiGate receives traffic, if the destination address, port, and protocol do not match an entry in application cache, FortiGate cannot immediately identify the application. It forwards the traffic using an SD-WAN rule and a firewall policy rule that do not require an application match. At the same time, the IPS engine analyzes the traffic to identify the application. This is called the learning phase.

When a 3-tuple from incoming traffic matches an application in the cache, FortiGate can steer the traffic according to application detected and the corresponding SD-WAN rules.

DO NOT REPRINT
© FORTINET



When you use SD-WAN for application steering, FortiGate must identify the application on the traffic before it can match the right rule. This is called the application learning phase. During this phase, FortiGate will still forward the packets with information already available. Therefore, the first packets of a session may not match the expected rule and member.

When a new session arrives, it matches SD-WAN rules with application steering only if there is a 3-tuple cache entry match. Otherwise, it matches an SD-WAN rule, based on destination address or ISDB entry (explicit or default implicit rule). At the same time, if the firewall policy that allows the traffic has application control, FortiGate begins application analysis. When FortiGate detects a new application, it updates the application cache. This trigger a reevaluation of the sessions (dirty process). If FortiGate flags the session as dirty, it reevaluates the session and can now match an SD-WAN rule using application steering.

It is important to understand that, during the learning phase, when FortiGate has not yet identified an application, the session might match a different rule and therefore might not be routed through the desired SD-WAN link.

Also, note that application learning can occurs only if firewall policies enable application detection. You must make sure application detection is enabled on policies that allow traffic during the learning phase.

DO NOT REPRINT
© FORTINET

Viewing the ISDB Application Cache

- View the ISDB application cache entries on the FortiGate CLI:

```
# diagnose sys sdwan internet-service-app-ctrl-list

SSH(16060 4294837753): 10.1.0.7 6 22 Mon Fev 06 02:44:18 2023
SSH(16060 4294837753): 10.1.0.28 6 22 Mon Fev 06 02:48:27 2023
Twitter(16001 4294838042): 104.244.42.65 6 443 Mon Fev 06 03:24:12 2023
```

Multiple 3-tuple for SSH app.

App ID | ISDB app ID | 3-tuple | Timestamp (8-hour lifetime); one application per 3T (most recent detected application is written to cache)

- SD-WAN rule correlation:

```
# diagnose sys sdwan service 2

Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members(2):
    1: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(1), cost(0), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x0), gid(0), cfg_order(0), cost(0), selected
  Internet Service(1): SSH(4294837753,0,0,0 16060)
  Src address(1):
        10.0.1.0-10.0.1.255
```

ISDB app ID | App ID (SSH)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

You can view the detected application entries in the ISDB application cache by running `diagnose sys sdwan internet-service-app-ctrl-list` on the FortiGate CLI. Each entry in the output indicates the application ID, the ISDB application ID, the 3-tuple, and the last time the entry was created. Note that FortiGate maintains one application per 3-tuple but one application can be associated with multiple 3-tuples. If multiple applications are detected for the same 3-tuple, FortiGate uses the most recent detected application for the 3-tuple.

FortiGate uses the application ID, ISDB application ID, and 3-tuple to match the SD-WAN rule. In the example shown on this slide, one cache entry is for an SSH connection destined to `10.1.0.7` on TCP port 22. A connection that matches the cache entry also matches SD-WAN rule ID 2.

Note that entries in ISDB application cache expires eight hours after last detected match.

DO NOT REPRINT
© FORTINET

Viewing the Application ID on the Session:

- Session routed using application steering:

```
session info: proto=6 proto_state=11 duration=5 expire=3596 timeout=3600 flags=00000000 socktype=0 sockport=0
av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr npu f00 f02 app_valid
statistic(bytes/packets/allow_err): org=3427/20/1 reply=4749/23/1 tuples=2
tx speed(Bps/kbps): 638/5 rx speed(Bps/kbps): 884/7
orgin->sink: org pre->post, reply pre->post dev=7->20/20->7 gwy=100.64.1.9/10.0.1.101
hook=pre dir=org act=noop 10.0.1.101:43020->10.1.0.7:22(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.0.7:22->10.0.1.101:43020(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
serial=000b2f2d tos=ff/ff app_list=2002 app=16060 url_cat=0
sdwan_mbr_seq=4 sdwan_service_id=2
rpdb_link_id=ff000002 ngfwid=n/a
npu_state=0x001008
```

App ID (SSH)

Member and rule IDs

- Look up application name by ID:

```
# get application name status | grep 16060 -B1
app-name: "SSH"
id: 16060
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

FortiGate writes the application information on the SD-WAN session. The `app` field indicates the application ID. You can look up the application name by using the `get application name status` command, as shown on this slide.

DO NOT REPRINT
© FORTINET

Unexpected Routing During Application Learning

- SSH session *without SNAT*:
  - Packet capture before application is learned (no entry in cache; unexpected member):

```
# diagnose sniffer packet any "host 10.1.0.7 and port 22" 4 | grep T_INET
...
9.612273 T_INET_0 out 10.0.1.101.45366 -> 10.1.0.7.22: syn 2174489972
9.723809 T_INET_0 in 10.1.0.7.22 -> 10.0.1.101.45366: syn 1405985100 ack 2174489973
```

  - Application is learned; entry is added in cache:

```
# diagnose sys sdwan internet-service-app-ctrl-list

SSH(16060 4294837753): 10.1.0.7 6 22 Mon Fev 06 19:31:31 2023
```

  - Packet capture after application is learned (same session):

```
...
45.728896 T_INET_1 out 10.0.1.101.45366 -> 10.1.0.7.22: psh 2174491574 ack 1405986770
45.740151 T_INET_1 in 10.1.0.7.22 -> 10.0.1.101.45366: psh 1405986770 ack 2174491722
```

  - Packet capture for new session with same 3-tuple (expected member):

```
...
59.423616 T_INET_1 out 10.0.1.101.50214 -> 10.1.0.7.22: syn 2044383473
59.425359 T_INET_1 in 10.1.0.7.22 -> 10.0.1.101.50214: syn 1902219655 ack 2044383474
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

When you use application steering, a session may match an unexpected rule and member when the session 3-tuple doesn't have an entry in the ISDB application cache, or when the session 3-tuple has an entry in the ISDB application cache, but the detected application is different. In both cases, the session may initially match the wrong rule and member.

Consider the example on this slide, which shows some packets captured for an SSH session *without SNAT*. The packets are sent to different members before and after the application learning phase is completed. The expected member is `T_INET_1`, but the packets are initially routed to `T_INET_0` because the packet didn't match an entry in the cache. After the application is learned, FortiGate starts routing the packets to `T_INET_1` and adds the 3-tuple to the cache. In addition, FortiGate routes new sessions that match the same 3-tuple in cache to `T_INET_1` right from the start of the connection.

Also, remember that SNAT conditions apply. That is, if a session is subject to SNAT, then, by default, FortiGate doesn't flush the routing information of the session after the application is detected—the session is not flagged as dirty. For example, if the SSH sessions shown on this slide were subject to SNAT, then the first session would remain routed through `T_INET_0` even after FortiGate detects the application—unless the default FortiOS SNAT routing behavior is modified. However, subsequent sessions to the same 3-tuple would then be routed to `T_INET_1`.

DO NOT REPRINT
© FORTINET

Criteria—Route Tag

- Configuration example
  - SD-WAN rule—use tag to define the service

```
config system sdwan
    config service
        edit 2
            set route-tag 10
        next
    end
end
```

- Set tag to some routes received from hub
  - Community list and route map:

```
config router community-list
    edit "hub-lan-cl"
        config rule
            edit 1
                set action permit
                set match "65000:10"
            . . .
config router route-map
    edit "hub-lan-rm"
        config rule
            edit 1
                set match-community "hub-lan-cl"
                set set-route-tag 10
...
```

  - BGP neighbor—set tag with `route-map-in`

```
config router bgp
    config neighbor
        edit "10.201.1.254"
            set interface "port1"
            set route-map-in "hub-lan-rm"
. . .
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

Route tag is one of the criteria you can use to determine which traffic matches an SD-WAN rule.

This slide shows FortiGate configuration commands for matching a community in received BGP routes, set route tags. You can then use the tag to define an SD-WAN rule.

On FortiGate, the route map assigns a route tag of `10` to prefixes that match the community indicated in the community list (`65000:10`). FortiGate then applies the route map in the inbound direction of two neighbors, each of them communicating over a different interface. In addition, the SD-WAN rule that steers corporate traffic—ID `2` in the example—is configured to use as destination the prefixes that have a route tag of `10`.

DO NOT REPRINT
© FORTINET

Checking Community, Route Tag, and Rule

- Learned prefixes by community:

```
# get router info bgp community 65000:10
...

   Network          Next Hop            Metric LocPrf Weight RouteTag Path
* i10.1.0.0/24      10.202.1.254             0    100      0       10 i <-/->
*>i                 10.201.1.254             0    100      0       10 i <-/1>

Total number of prefixes 1
```

- SD-WAN service status:

```
# diagnose sys sdwan service 2
Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members(2):
    1: Seq_num(1 port1), alive, sla(0x0), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(2 port2), alive, sla(0x0), gid(0), cfg_order(1), local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Route tag address(1):
        10.1.0.0-10.1.0.255
```

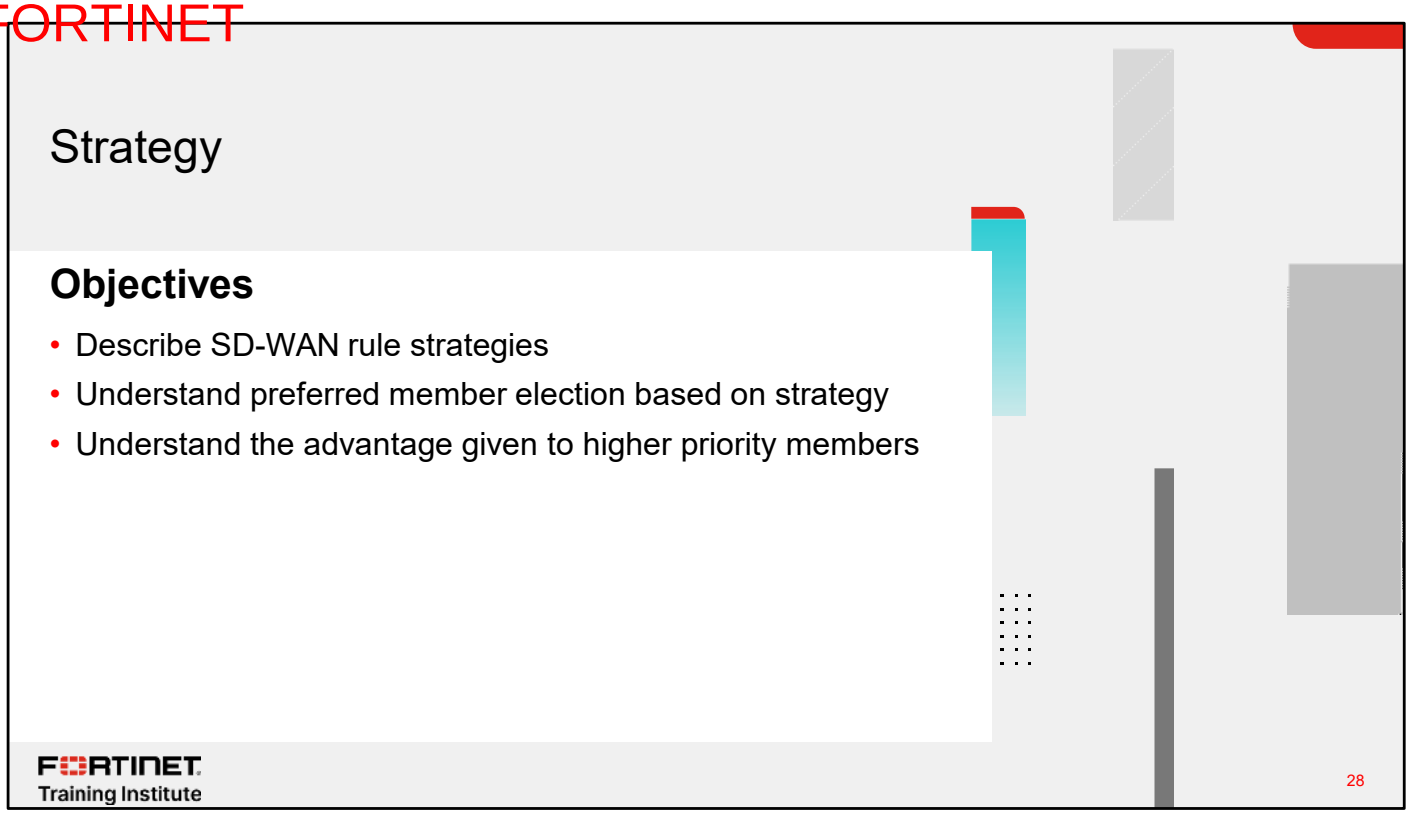
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

On a FortiGate with the settings from the previous slide, when you look up the learned prefixes by community, the output includes the `10.1.0.0/24` subnet. The output also indicates that a route tag of `10` is assigned to the prefix. The result is that when you check the SD-WAN rule status, the output includes `10.1.0.0/24` as a route tag address.

DO NOT REPRINT
© FORTINET



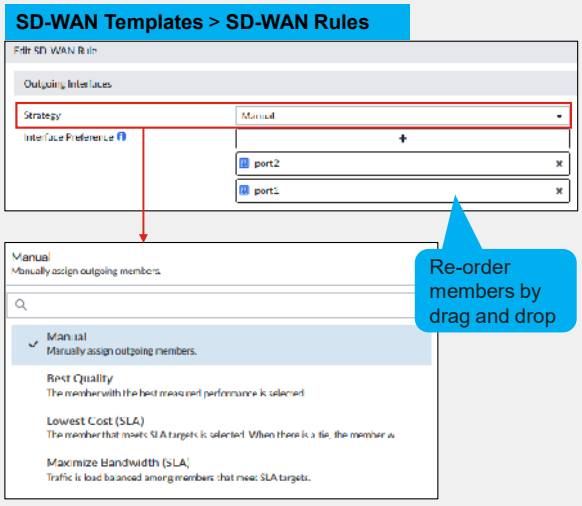
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding available SD-WAN rule strategies you should be able to select and configure the best option to address your requirements.

DO NOT REPRINT
© FORTINET

Strategies

- Define:
  - Requirements for preferred members
  - Single or multiple member traffic distribution
- Preferred members:
  - Best candidates to steer traffic
  - Are used only if they have a valid route to the destination
- Member selection:
  - **Manual**:
    - Configuration order-based preference
  - **Best Quality**:
    - Best performing member based on quality criteria
  - **Lowest Cost (SLA)**:
    - Member that meets SLA target (tiebreakers: cost and priority)
  - **Maximize Bandwidth (SLA)**:
    - Members that meet SLA target
    - Traffic is load balanced across multiple members



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

The strategy in a rule defines the requirements for preferred members, and whether the traffic is steered to one or more preferred members. The preferred members are the best members from the `oif` list—based on the strategy in use—that meet the SLA requirements (if applicable). The `oif` list sorts the configured members—**Interface Preference** list in FortiManager—by preference. That is, although the members are the same, their order in the `oif` list, and in **Interface Preference** list, can be different. There are four strategies you can chose from:

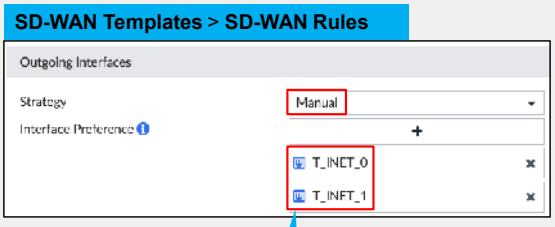
- **Manual**: FortiGate prefers members according to configuration order. Member metrics are not considered for member preference.
- **Best Quality:** FortiGate prefers the best-performing member based on the configured quality criteria.
- **Lowest Cost (SLA)**: FortiGate prefers the member that meets the configured SLA target. If multiple members meet the SLA target, member cost, followed by the configuration order, are used as tiebreakers.
- **Maximize Bandwidth (SLA)**: FortiGate prefers the members that meet the configured SLA target. If multiple members meet the SLA target, FortiGate load balances the sessions across all preferred members using the configured hashing algorithm (round-robin by default).

Note that for all strategies, by default, FortiGate must check that the preferred member has a valid route to the destination. If the member doesn't have a valid route, then FortiGate checks the next member in the `oif` list, and so on until finding an acceptable member. Moreover, all strategies, except **Manual**, consider the member metrics for member preference. In addition, **Maximize Bandwidth (SLA)** is the only strategy that supports traffic load balancing across multiple members. That is, when using any of the other three strategies, traffic is steered to one member only.

DO NOT REPRINT
© FORTINET

Manual

- First member in the configuration that is alive is preferred
  - Metrics are not considered (packet-loss, latency, …)
  - Health-check configuration is optional
  - Single preferred member (no load balancing)




```
# diagnose sys sdwan health-check status
Health Check(VPN_PING):
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(101.170), jitter(5.808), mos(4.3  , bandwidth-
up(10239), bandwidth-dw(10239), bandwidth-bi(20478) sla_map=0x3
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.420), jitter(0.289), mos(4.4  bandwidth-
up(10239), bandwidth-dw(10239), bandwidth-bi(20478) sla_map=0x3

# diagnose sys sdwan service
Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Members(2):
    1: Seq_num(3 T_INET_0), alive, selected
    2: Seq_num(4 T_INET_1), alive, selected
  Internet Service(1): SSH(4294837953,0,0,0,0 16060)
  Src address(1):
       10.0.1.0-10.0.1.255
```

T_INET_0 is preferred because is alive and the first member configured in the list

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

When you use the manual strategy, FortiGate builds the `oif` list based on the **Interface Preference** list configuration order. For building the `oif` list, FortiGate does not consider the performance of the member, such as jitter or packet loss . It considers only the member configuration order—also known as the member configuration priority—and whether the member is alive or dead. If the first configured member is dead, then the preferred member becomes the next member in the configuration list that is alive.

Note that the manual strategy *does not* require the configuration of performance SLA. However, when a performance SLA rule is configured to monitor members, it improves the detection of whether a member is alive or dead, because a member is considered alive only if the health-check can reach at least one configured server. Without a health-check, members are considered alive or dead according to the interface status (up or down).

In the example shown on this slide, `T_INET_0` and `T_INET_1` are the configured members. However, although `T_INET_0` has a much higher latency than `T_INET_1`, FortiGate still prefers `T_INET_0` for steering traffic because it is alive and has the highest configuration priority.

Do not confuse the member configuration priority with the **Priority** setting available on the SD-WAN member configuration. The latter is used for the priority of static routes for members when you configure static routes for zones. The former refers to the member priority based on the **Interface Preference** list configuration. Members that are configured first in the list have higher priority over those configured last. The **Priority** setting is used as a tiebreaker for ECMP routes when matching the implicit SD-WAN rule.

DO NOT REPRINT
© FORTINET

Manual (Contd)

- If T_INET_0 is detected dead, T_INET_1 becomes the preferred member:

```
# diagnose sys sdwan health-check status
Health Check(VPN_PING):
Seq(3 T_INET_0): state(dead), packet-loss(15.000%) sla_map=0x0
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.261), jitter(0.294), mos(4.403), bandwidth-
up(10239), bandwidth-dw(10239), bandwidth-bi(20478) sla_map=0x3

# diagnose sys sdwan service

Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(2), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Members(2):
    1: Seq_num(4 T_INET_1), alive, selected
    2: Seq_num(3 T_INET_0), dead
  Internet Service(1): SSH(4294837953,0,0,0 16060)
  Src address(1):
        10.0.1.0-10.0.1.255
```

T_INET_1 is now the first member in the oif list

Dead members are moved to the bottom of the oif list

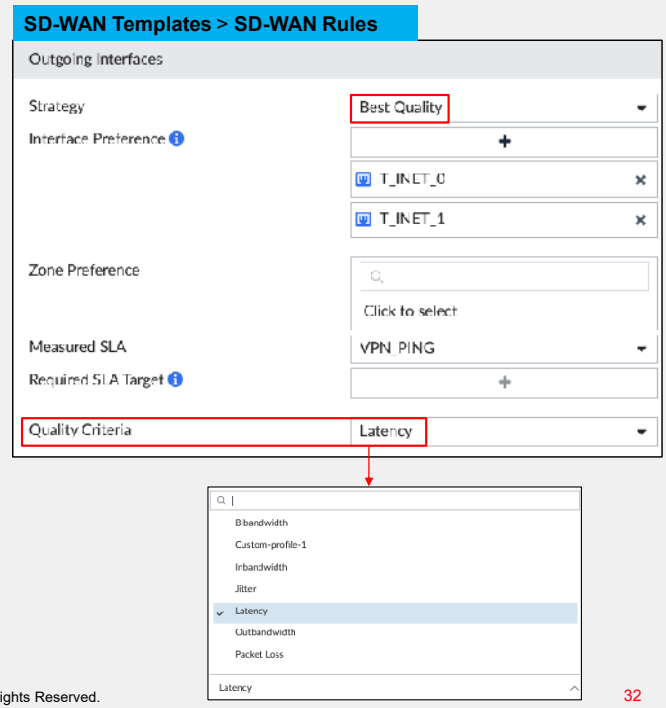
© Fortinet Inc. All Rights Reserved.

In the example shown on this slide, FortiGate detects T_INET_0 as dead, and as result, moves T_INET_1 to the top of the oif list by assigning it the sequence number 1. Also, T_INET_0 is assigned the sequence number 2 because dead members are moved to the bottom of the list.

DO NOT REPRINT
© FORTINET

Best Quality

- Member with best quality becomes preferred
- Multiple quality metrics available
  - Default: latency
- Factors considered for best quality:
  - Member priority
  - link-cost-threshold
  - Metric
- SLA targets are not considered
- One member is used (no load balancing)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

The best quality strategy instructs FortiGate to select as the preferred member the member with the best measured quality. FortiGate supports multiple metrics to measure the link quality. The default metric is latency. For increased accuracy, FortiGate calculates the quality metric on the last few health-check probes (default 30).

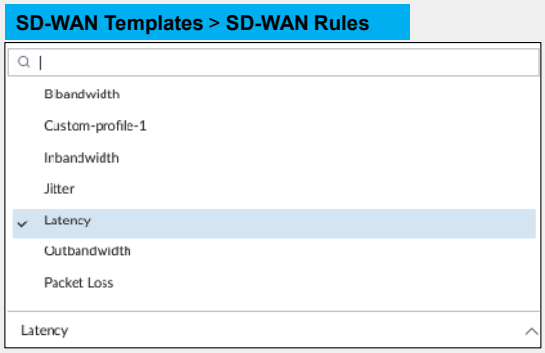
FortiGate determines the member with the best quality using three factors: member priority, the `link-cost-threshold` setting, and the value of the metric measured for the member. SLA targets are not considered.

In addition, when you use the best quality strategy, FortiGate uses one member to steer traffic. That is, it does not use balancing.

DO NOT REPRINT
© FORTINET

Best Quality—Metrics

- **Latency** (default)
  - Lower latency, higher preference
- **Jitter**
  - Lower jitter, higher preference
- Available bandwidth:
  - More available bandwidth, higher preference
  - Three measurement options:
    - **Inbandwidth** (ingress)
    - **Outbandwidth** (egress)
    - **Bibandwidth** (bidirectional)
- **Packet Loss**
  - Lower packet loss, higher preference
- **Custom-profile-1**
  - Weight-based calculation based on:
    - Latency, jitter, packet loss, and bibandwidth
    - Assign your own weights
  - Lower quality index, higher preference



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

This slide shows the metrics you can choose from to measure the quality of the members in a best quality rule.

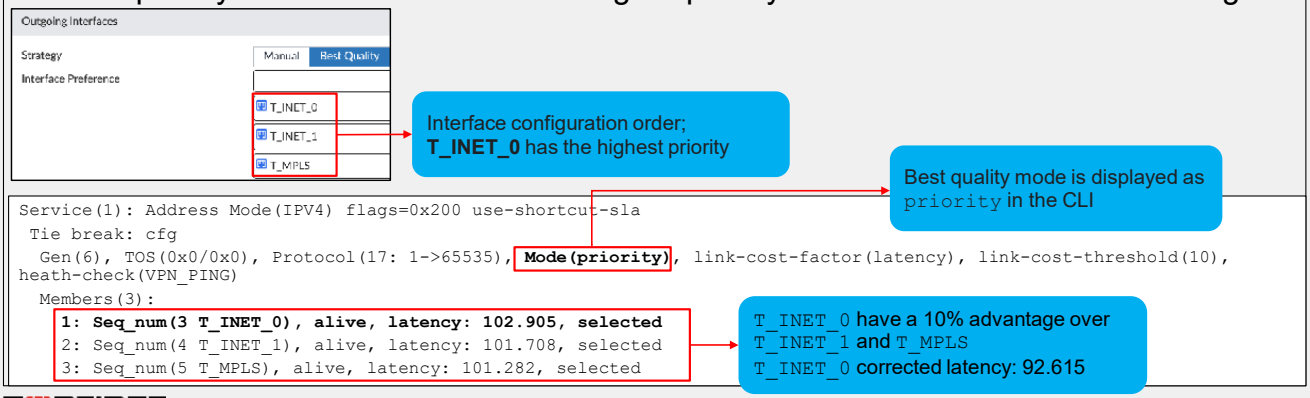
For latency, jitter, and packet loss, the lower the metric value, the higher the member preference. The opposite is true for available bandwidth: The more available bandwidth a member has, the higher its preference. Note that FortiGate can measure the available ingress bandwidth, the available egress bandwidth, or the sum of both (bidirectional bandwidth).

When you select the **Custom-profile-1** metric, FortiGate uses a weight-based formula to calculate a value—called the link quality index—that represents the quality of the member based on its latency, jitter, packet loss, and available bidirectional bandwidth. The lower the link quality index, the higher the member preference. The administrator assigns the weight for each metric.

DO NOT REPRINT
© FORTINET

Best Quality—Link Cost Threshold

- Advantage given to higher priority members
  - Logic: higher priority members have more preference
  - Prevents SLA changes due to slight metric variations
  - Default value: 10 (percentage)
  - Calculation varies among metrics
- Lower priority member is better than a higher priority member if it beats the advantage



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

In addition to the member metric value, FortiGate also considers the member priority and the `link-cost-threshold` value to determine the preferred member in a best quality rule. You can identify the preferred member by looking at the output of the `diagnose sys sdwan service` command. The member that is placed on top of the list is the preferred member.

`link-cost-threshold` instructs FortiGate to give an advantage (or handicap) to higher priority members. The member priority is determined from the interface configuration order. Members that are configured first have higher priority over those configured last. The logic behind `link-cost-threshold` is that higher priority members should have more preference than lower priority members. For this reason, for a lower priority member to be considered better than a higher priority member, the lower priority member must beat the advantage given to the higher priority member.

`link-cost-threshold` is also useful to prevent frequent SLA changes caused by slight variations in the metric, which would in turn result in more CPU activity caused by frequent session re-evaluation. The `link-cost-threshold` setting is expressed as a percentage, and its default value is **10**.

The example on this slide shows the output of the `diagnose sys sdwan service command` for a rule that uses best quality (shown in the CLI as `priority`) as the strategy. `T_MPLS` and `T_INET_1` have lower latency than `T_INET_0`, but because `link-cost-threshold` is set to `10`, then `T_INET_0`—the highest priority member—has an advantage of 10% over other members configured in the rule. The 10% advantage results in a corrected latency of 92.615 ms for `T_INET_0`, which makes it the lowest latency among the members. For this reason, `T_INET_0` becomes the preferred member.

The advantage is calculated differently depending on the metric in use.

DO NOT REPRINT
© FORTINET

Best Quality—Link Cost Threshold (Contd)

- Advanced option
  - SD-WAN rule advanced options parameter
- CLI configuration (default = `10`):

```
config system sdwan
    config service
        edit <id>
            set link-cost-threshold <threshold>
        next
    end
end
```

- If `link-cost-threshold` is set to `0`, then the metric comparison is strict
  - Use priority as tiebreaker

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

You can configure the `link-cost-threshold` parameter in the SD-WAN rule advanced options section on FortiManager or using a CLI parameter. This slide shows how to configure the setting using the FortiGate CLI.

If you set `link-cost-threshold` to **0**, then FortiGate performs a strict metric comparison. That is, there is no advantage, and the member with the best metric becomes the preferred. If two or more members have the same metric, FortiGate uses the member priority as the tiebreaker.

DO NOT REPRINT
© FORTINET

Best Quality—Latency, Jitter, and Custom-profile-1

- Objective:
  - Steer traffic to the member with the lowest metric (the advantage is considered)
- Corrected metric (CM) formula (factors in the advantage):

CM = real metric / (1 + `link-cost-threshold`/100)

- Rules for building the `oif` list (`link-cost-threshold` > 0):
  - A higher priority member (HPM) has advantage over a lower priority member (LPM)
  - An LPM is placed on top of an HPM when its metric is better than the CM of the HPM
  - An LPM is placed again below an HPM when any of the following occurs:
    - The LPM metric is equal to or worse than the CM of the highest priority member (quick recovery)
    - The LPM metric is equal to or worse than the HPM (longer recovery)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

In you configure best quality as the strategy, the way FortiGate determines the preferred member depends on the metric in use. However, for latency, jitter, and custom-profile-1, the method is the same.

When you use latency, jitter, or custom-profile-1 as the metric, FortiGate selects as the preferred member the member with the lowest metric. If `link-cost-threshold` is set to a value higher than zero, FortiGate calculates the corrected metric (CM) for higher priority members, which factors in the advantage given to them. FortiGate uses the formula shown on this slide to calculate the CM.

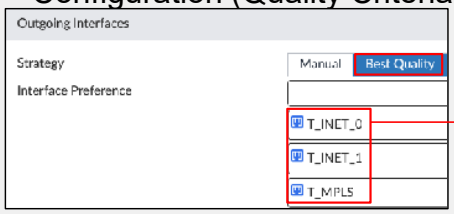
This slide shows the rules that FortiGate follows when building the `oif` list. Do not confuse the **Interface Preference** configuration list of a rule with its `oif` list. The latter is dynamic and depends on the member performance. The former refers to the order of the interfaces in the configuration. The rules are the following:

- A higher priority member (HPM) has advantage over a lower priority member (LPM). The advantage depends on the `link-cost-threshold` value.
- An LPM is placed on top of an HPM when its metric is better than the CM of the HPM.
- An LPM is placed again below an HPM when any of the following occurs:
  - The LPM metric is equal to or worse than the CM of the highest priority member. This allows for a quick recovery of the highest priority member.
  - The LPM metric is equal to or worse than the HPM. This results in a longer recovery of the HPM.

DO NOT REPRINT
© FORTINET

Best Quality Member—Latency Example 1

- Configuration (Quality Criteria = **Latency**, `link-cost-threshold` = `10`):



Interface configuration order;
T_INET_0 has the highest priority

- Moving up and down the list (highest priority member vs LPM):

```
...
1: Seq_num(3 T_INET_0), alive, latency: 101.702, selected
2: Seq_num(4 T_INET_1), alive, latency: 96.282, selected
3: Seq_num(5 T_MPLS), alive, latency: 97.236, selected
...
1: Seq_num(4 T_INET_1), alive, latency: 91.098, selected
2: Seq_num(3 T_INET_0), alive, latency: 101.652, selected
3: Seq_num(5 T_MPLS), alive, latency: 97.196, selected
...
1: Seq_num(3 T_INET_0), alive, latency: 101.699, selected
2: Seq_num(4 T_INET_1), alive, latency: 92.898, selected
3: Seq_num(5 T_MPLS), alive, latency: 97.285, selected
...
```

1. T_INET_0 CM = 101.702 / 1.10 = 92.46
   T_INET_1 latency ≥ 92.46
   T_MPLS latency ≥ 92.46
   Result: no change in order
2. T_INET_0 CM = 101.652 / 1.10 = 92.41
   T_INET_1 latency < 92.41
   T_MPLS latency ≥ 92.41
   Result: Move T_INET_1 above T_INET_0
3. T_INET_0 CM = 101.699 / 1.10 = 92.45
   T_INET_1 latency ≥ 92.45
   T_MPLS latency ≥ 92.45
   Result: Move T_INET_1 back below T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

37

This slide shows an example of the preferred member election in a best quality rule that uses **Latency** as the metric. For jitter and custom-profile-1 metrics, the same behavior applies.

The configuration shows **T_INET_0** as the highest priority member. This slide also shows the output of the `diagnose sys sdwan service` command taken at different times. Note that only relevant lines of the output are displayed. The example shows how a member competes with the highest priority member for a position in the `oif` list.

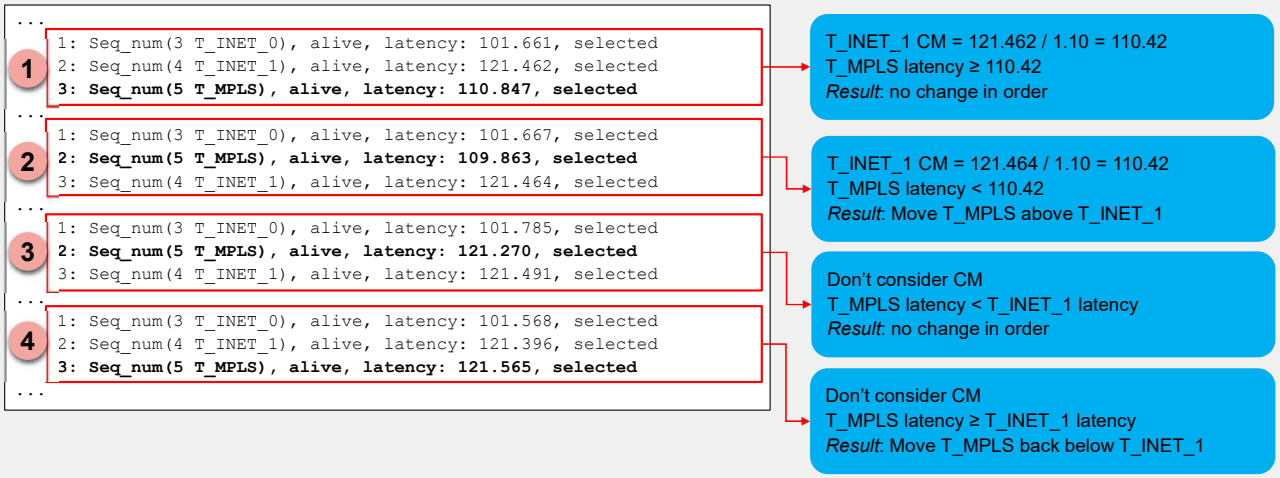
Focus on the position of T_INET_1 in the list compared to T_INET_0. In the first output, T_INET_1 is placed below T_INET_0 even though it has the lowest real latency among members. This is because FortiGate calculates a corrected metric (CM) for T_INET_0. The corrected latency for T_INET_0 is 92.46 ms, which makes it the actual lowest latency among all members, which is why FortiGate selects T_INET_0 as the preferred member.

In the second output, T_INET_1 now becomes the preferred member because its latency is now lower than the corrected latency of T_INET_0. In the third output, T_INET_1 is moved back below T_INET_0 because its latency is again equal to or higher than the corrected latency of T_INET_0.

DO NOT REPRINT
© FORTINET

Best Quality Member—Latency Example 2

- Moving up and down the list (HPM vs LPM):



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

38

This slide shows another example of the preferred member election in a best quality rule that uses latency as the metric. The configuration is the same as the one shown on the previous slide. The example shows how members other than the highest priority member, compete for a position in the `oif` list.

Focus on the position of T_MPLS in the list compared to T_INET_1. In the first output, T_MPLS is placed below T_INET_1 even though T_MPLS has a lower latency. This is because FortiGate considers the corrected latency (CM) for T_INET_1, which is 110.42 ms and equal to or lower than the real latency of T_MPLS.

In the second output, T_MPLS is now placed on top of T_INET_1 because its latency is now lower than the corrected latency of T_INET_1. In the third output, T_MPLS is still placed on top of T_INET_1 even though its latency is no longer lower than the corrected latency of T_INET_1. This is because FortiGate doesn't consider the advantage when moving a lower priority member back down the interface list. That is, FortiGate considers only the members' real latency and priority.

In the fourth output, T_MPLS is now moved back below T_INET_1 because its latency is now equal to or higher than the real latency of T_INET_1.

DO NOT REPRINT
© FORTINET

Best Quality—Bandwidth

- Objective:
  - Steer traffic to the member with the highest available bandwidth
  - Available bandwidth (ABW) is based on interface settings and usage
- Corrected bandwidth difference (CBWD) formula (factors in the advantage):

CBWD = LPM ABW * `link-cost-threshold`/100

- Available bandwidth difference (ABWD) formula:

ABWD = LPM ABW – HPM ABW

- Rules for building the `oif` list (`link-cost-threshold` > 0):
  - A higher-priority member (HPM) has advantage over a lower-priority member (LPM)
    - Higher priority, higher preference
  - An LPM is placed on top of an HPM when the ABWD is higher than the CBWD
  - An LPM is placed again below an HPM when any of the following occurs:
    - The ABWD is equal to or lower than the CBWD (quick recovery)
    - The ABW of LPM is equal to or lower than the ABW of the HPM (longer recovery)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

39

When you use inbandwidth, outbandwidth, or bibandwidth as the metric, FortiGate selects the member with the highest available bandwidth as the preferred member. The available bandwidth depends on the interface settings and usage.

If you set the `link-cost-threshold` to a value higher than zero, FortiGate calculates the corrected bandwidth difference for lower-priority members, which factors in the advantage given to the higher-priority members and is calculated using the formula shown on this slide. The formula is a basic percentage calculation of the available bandwidth of the lower priority member based on the `link-cost-threshold` setting.

FortiGate also calculates the difference in available bandwidth between a lower-priority member and a higher-priority member using the formula shown on this slide.

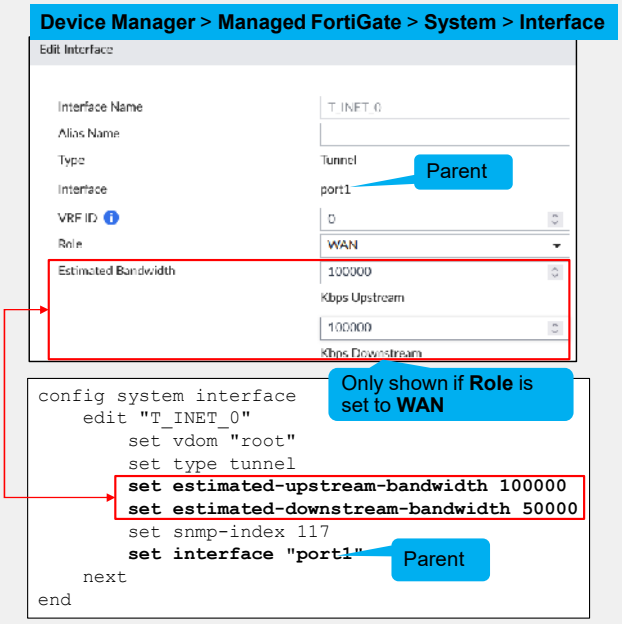
This slide also shows the rules that FortiGate follows to build the `oif` list using bandwidth as the metric. The rules are the following:

- A higher-priority member (HPM) has advantage over a lower-priority member (LPM). The advantage depends on the `link-cost-threshold` value.
- A LPM is placed on top of a HPM when the available bandwidth difference is higher than the corrected bandwidth difference.
- A LPM is placed again below a HPM when any of the following occurs:
  - The available bandwidth difference is equal to or lower than the corrected bandwidth difference. This allows for a quick recovery of the highest priority member.
  - The available bandwidth of the LPM is equal to or lower than the available bandwidth of the HPM. This results in a longer recovery for the HPM.

DO NOT REPRINT
© FORTINET

Best Quality—Member Maximum Bandwidth

- FortiGate requires maximum bandwidth for members
  - Manually defined using `estimated-upstream-bandwidth` and `estimated-downstream-bandwidth`
  - Or automatically obtained from:
    - Physical speed (physical interfaces)
    - Parent interfaces (VLAN and tunnel interfaces)
      - Use parent physical speed as last resource
  - Bibandwidth is the sum of upstream and downstream bandwidth



FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

When you use a bandwidth metric in a best quality rule, FortiGate must know the maximum bandwidth of the members to calculate their available bandwidth.

The maximum bandwidth of a member is obtained from the values defined for the `estimated-upstream-bandwidth` and `estimated-downstream-bandwidth` interface settings. If the settings are not defined, then FortiGate uses the actual physical interface speed. For example, for a 100 Mbps interface with undefined maximum bandwidth settings, FortiGate resolves to use 100 Mbps for both upstream and downstream maximum speeds. For the bidirectional bandwidth metric, FortiGate just adds the upstream and downstream speeds.

For VLAN and tunnel interfaces with undefined maximum bandwidth settings, FortiGate resolves to use the bandwidth values of the parent interfaces. For example, an IPsec tunnel that is bound to port1 and that has undefined maximum bandwidth settings obtains its maximum bandwidth from port1 settings. If port1 has no maximum bandwidth values defined, then the IPsec tunnel uses the physical speed of port1.

The example on this slide shows the maximum bandwidth settings on FortiManager, and the equivalent CLI settings on FortiGate. Note that FortiManager displays the maximum bandwidth settings only if you select **WAN** as the interface **Role**.

DO NOT REPRINT
© FORTINET

Best Quality—Member Available Bandwidth

- Outbandwidth example:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(priority), link-cost-factor(outbandwidth), link-cost-threshold(10),
heath-check(VPN_PING)
  Members(2):
    1: Seq_num(3 T_INET_0), alive, outbandwidth: 99995Kbps, selected
    2: Seq_num(4 T_INET_1), alive, outbandwidth: 99995Kbps, selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

Available bandwidth

- Bibandwidth example:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(priority), link-cost-factor(bibandwidth), link-cost-threshold(10),
heath-check(VPN_PING)
  Members(2):
    1: Seq_num(3 T_INET_0), alive, bibandwidth: 149980Kbps, selected
    2: Seq_num(4 T_INET_1), alive, bibandwidth: 149980Kbps, selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

41

This slide shows sample outputs for the `diagnose sys sdwan service` command when using outbandwidth and bibandwidth as metrics. Note that the bandwidth value displayed in the output refers to the available bandwidth.

DO NOT REPRINT
© FORTINET

Best Quality—Bandwidth Used by Member

- Available bandwidth:

```
# diagnose sys sdwan service

Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(priority), link-cost-factor(bibandwidth), link-cost-threshold(10),
heath-check(VPN_PING)
  Members(2):
    1: Seq_num(3 T_INET_0), alive, bibandwidth: 139234Kbps, selected
    2: Seq_num(4 T_INET_1), alive, bibandwidth: 149997Kbps, selected
  Src address(1):
        0.0.0.0-255.255.255.255
  Dst address(1):
        10.1.0.7-10.1.0.7
```

Available bandwidth is updated every 10 seconds

- Used bandwidth:

```
# diagnose sys link-monitor interface T_INET_0
Interface(T_INET_0): state(up, since Tue Fev 07 08:44:45 2023), bandwidth(up:519189bps, down:15097bps), session
count(IPv4:63, IPv6:0), tx(24754785 bytes), rx(23261162 bytes), latency(2.08), jitter(0.49), packet-loss(0.00).

# diagnose sys link-monitor interface T_INET_1
Interface(T_INET_1): state(up, since Tue Fev 07 08:44:45 2023), bandwidth(up:4769bps, down:15097bps), session
count(IPv4:2, IPv6:0), tx(13157864 bytes), rx(23640548 bytes), latency(1.91), jitter(0.41), packet-loss(0.00).
```

Used bandwidth

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

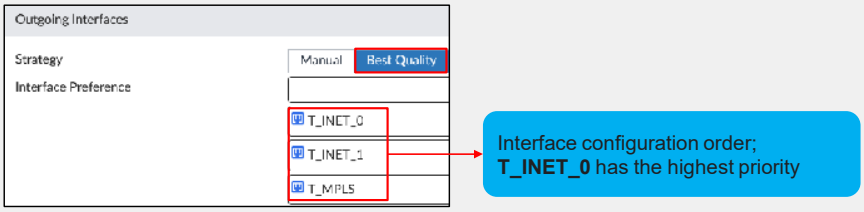
The available bandwidth displayed by the `diagnose sys sdwan service` command is updated every 10 seconds based on the used bandwidth and the maximum bandwidth of the member.

The used bandwidth is measured by the link monitor process (`lnkmtd`), and you can view the measured values by running the `diagnose sys link-monitor interface` command, as shown on this slide.

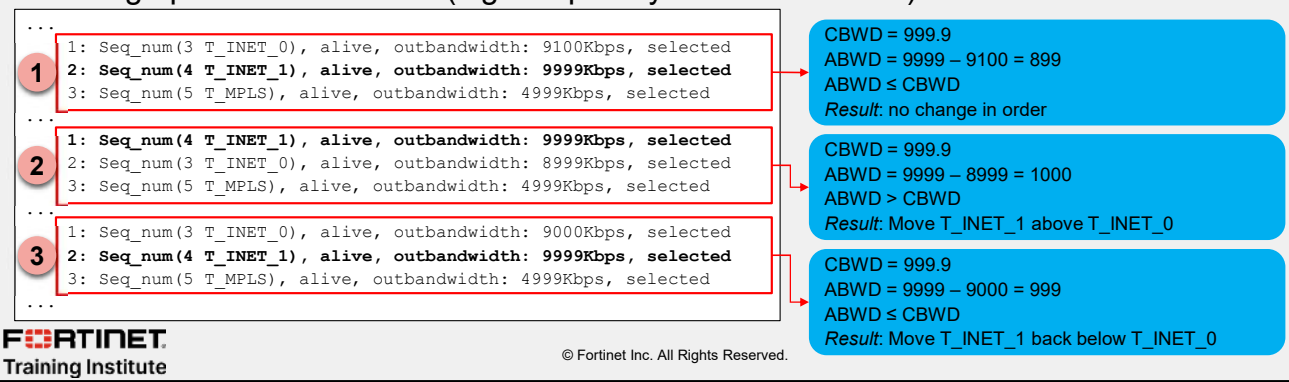
DO NOT REPRINT
© FORTINET

Best Quality Member—Bandwidth Example 1

- Configuration (Quality Criteria = **Outbandwidth**, `link-cost-threshold` = `10`)



- Moving up and down the list (highest priority member vs LPM):



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

43

This slide shows an example of the preferred member election in a best quality rule that uses **Outbandwidth** as the metric.

The configuration shows T_INET_0 as the highest priority member. This slide also shows the output of the `diagnose sys sdwan service` command taken at different times. Note that only relevant lines of the output are displayed. The example shows how a member competes with the highest priority member for a position in the `oif` list.

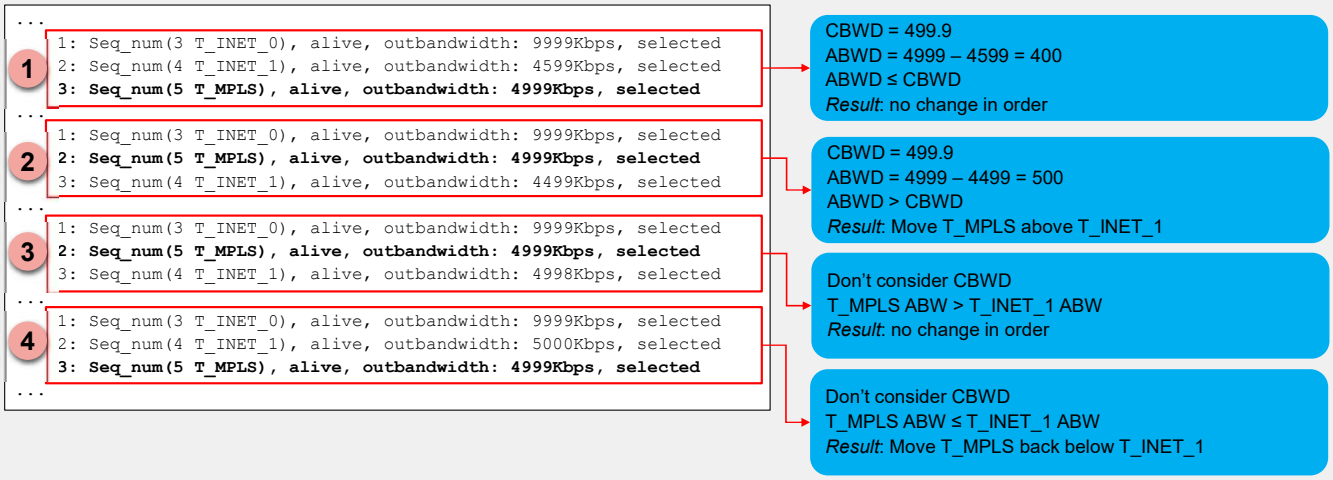
Focus on the position of T_INET_1 in the list compared to T_INET_0. In the first output, T_INET_1 is placed below T_INET_0 even though it has more available bandwidth. This is because FortiGate considers the corrected available bandwidth difference (CBWD). The CBWD is 999.9 kbps, while the ABWD is 899 kbps. Because ABWD is equal to or lower than the CBWD, then the higher priority member (T_INET_0) is considered better.

In the second output, T_INET_1 now becomes the preferred member because the ABWD is now higher than the CBWD. In the third output, T_INET_1 is moved back below T_INET_0 because the ABWD is again equal to or lower than the CBWD.

DO NOT REPRINT
© FORTINET

Best Quality Member—Bandwidth Example 2

- Moving up and down the list (HPM vs LPM):



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

44

This slide shows another example of the preferred member election in a best quality rule that uses outbandwidth as the metric. The configuration is the same as the one shown on the previous slide. The example shows how members other than the highest priority member, compete for a position in the list.

Focus on the position of T_MPLS in the list compared to T_INET_1. In the first output, T_MPLS is placed below T_INET_1 even though it has more available bandwidth (ABW). This is because FortiGate considers the corrected available bandwidth difference (CBWD). The CBWD is 499.9 kbps, while the ABWD is 400 kbps. Because ABWD is equal to or lower than the CBWD, then the higher priority member (T_INET_1) is considered better.

In the second output, T_MPLS is now placed on top of T_INET_1 because the ABWD is now higher than the CBWD. In the third output, T_MPLS is still placed on top of T_INET_1 even though the ABWD is no longer higher than the CBWD. This is because FortiGate doesn't consider the CBWD when moving a lower priority member back down the interface list. That is, FortiGate prefers the member with the highest available bandwidth.

In the fourth output, T_MPLS is now moved back below T_INET_1 because the available bandwidth is now equal to or lower than T_INET_1.

DO NOT REPRINT
© FORTINET

Best Quality—Packet Loss

- Objective:
  - Steer traffic to the member with the lowest packet loss
- Corrected packet loss (CPL) formula (adds advantage):

CPL = LPM packet loss + `link-cost-threshold`

- Rules for building the `oif` list (`link-cost-threshold` > 0):
  - A higher priority member (HPM) has advantage over a lower priority member (LPM)
    - Higher priority, higher preference
  - An LPM is placed on top of an HPM when the CPL is lower than the HPM packet loss
  - An LPM is placed again below an HPM when any of the following occurs:
    - The CPL is equal to or higher than the highest priority member packet loss (quick recovery)
    - The LPM packet loss is equal to or higher than the HPM packet loss (longer recovery)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

45

When you use packet loss as the metric, FortiGate selects as the preferred member the member with the lowest packet loss percentage.

If `link-cost-threshold` is set to a value higher than zero, FortiGate calculates the corrected packet loss (CPL) by adding the `link-cost-threshold` value to the packet loss of the LPM. The CPL factors in the advantage given to the higher priority member.

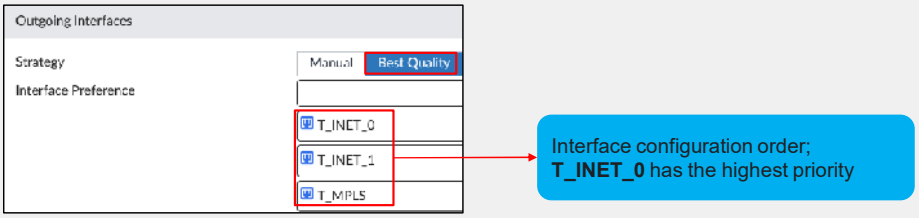
This slide also shows the rules that FortiGate follows when building the `oif` list using packet loss as the metric. The rules are the following:

- A higher priority member (HPM) has advantage over a lower priority member (LPM). The advantage depends on the `link-cost-threshold` value.
- An LPM is placed on top of an HPM when the CPL is lower than the packet loss of the HPM.
- An LPM is placed again below an HPM when any of the following occurs:
  - The CPL is equal to or higher than the packet loss of the highest priority member. This allows for a quick recovery of the highest priority member.
  - The packet loss of the LPM is equal to or higher than the packet loss of the HPM. This results in a longer recovery of the HPM.

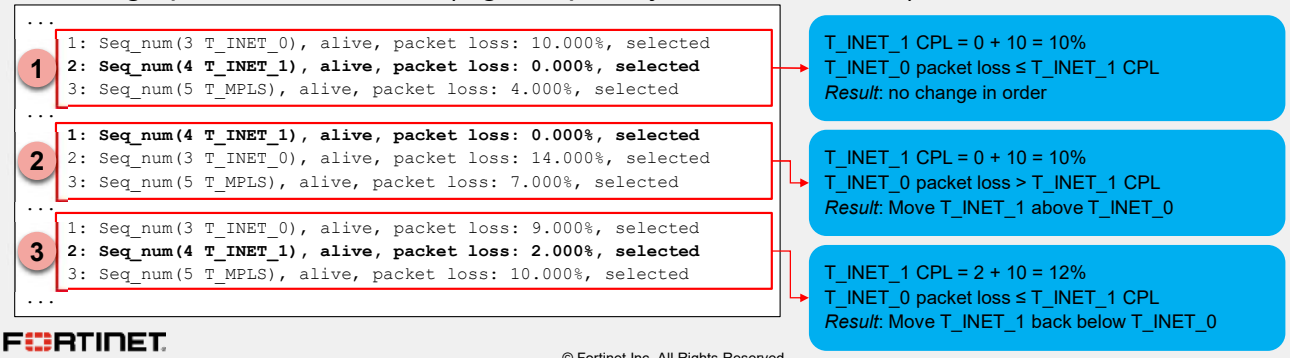
DO NOT REPRINT
© FORTINET

Best Quality Member—Packet Loss Example 1

- Configuration (Quality Criteria = **Packet Loss** `link-cost-threshold` = `10`)



- Moving up and down the list (highest priority member vs LPM):



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

46

This slide shows an example of the preferred member election in a best quality rule that uses **Packet Loss** as the metric.

The configuration shows **T_INET_0** as the highest priority member. This slide also shows the output of the `diagnose sys sdwan service` command taken at different times. Note that only relevant lines of the output are displayed. The example shows how a member competes with the highest priority member for a position in the list.

Focus on the position of T_INET_1 in the list compared to T_INET_0. In the first output, T_INET_1 is placed below T_INET_0 even though it has no packet loss. This is because FortiGate considers the corrected packet loss (CPL) for T_INET_1, which is 10%. Because the T_INET_0 packet loss is equal to or lower than the CPL, then T_INET_0 is considered better.

In the second output, T_INET_1 now becomes the preferred member because the T_INET_0 packet loss is now higher than the CPL. In the third output, T_INET_1 is moved back below T_INET_0 because T_INET_0 packet loss is again equal to or lower than the CPL.

DO NOT REPRINT
© FORTINET

Best Quality Member—Packet Loss Example 2

- Moving up and down the list (HPM vs LPM):

```
...
1: Seq_num(3 T_INET_0), alive, packet loss: 5.000%, selected
2: Seq_num(4 T_INET_1), alive, packet loss: 5.000%, selected
3: Seq_num(5 T_MPLS), alive, packet loss: 2.000%, selected
...
1: Seq_num(3 T_INET_0), alive, packet loss: 0.000%, selected
2: Seq_num(5 T_MPLS), alive, packet loss: 0.000%, selected
3: Seq_num(4 T_INET_1), alive, packet loss: 11.000%, selected
...
1: Seq_num(3 T_INET_0), alive, packet loss: 0.000%, selected
2: Seq_num(5 T_MPLS), alive, packet loss: 5.000%, selected
3: Seq_num(4 T_INET_1), alive, packet loss: 6.000%, selected
...
1: Seq_num(3 T_INET_0), alive, packet loss: 0.000%, selected
2: Seq_num(4 T_INET_1), alive, packet loss: 4.000%, selected
3: Seq_num(5 T_MPLS), alive, packet loss: 7.000%, selected
...
```

1. T_MPLS CPL = 2 + 10 = 12%
   T_INET_1 packet loss ≤ T_MPLS CPL
   Result: no change in order
2. T_MPLS CPL = 0 + 10 = 10%
   T_INET_1 packet loss > T_MPLS CPL
   Result: Move T_MPLS above T_INET_1
3. Don't consider CPL
   T_INET_1 packet loss > T_MPLS packet loss
   Result: no change in order
4. Don't consider CPL
   T_INET_1 packet loss ≤ T_MPLS packet loss
   Result: Move T_MPLS back below T_INET_1

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

47

This slide shows another example of the preferred member election in a best quality rule that uses packet loss as the metric. The configuration is the same as the one shown on the previous slide. The example shows how members other than the highest priority member, compete for a position in the list.

Focus on the position of T_MPLS in the list compared to T_INET_1. In the first output, T_MPLS is placed below T_INET_1 even though it has lower packet loss. This is because FortiGate considers the CPL, which is 12%. Because T_INET_1 packet loss is equal to or lower than the CPL, then the higher priority member (T_INET_1) is considered better.

In the second output, T_MPLS is now placed on top of T_INET_1 because its CPL is now lower than T_INET_1. In the third output, T_MPLS is still placed on top of T_INET_1 even though its CPL is no longer lower than T_INET_1. This is because FortiGate doesn't consider the CPL when moving a lower priority member back down the interface list. That is, FortiGate prefers the member with the lowest packet loss.

In the fourth output, T_MPLS is now moved back below T_INET_1 because its packet loss is now equal to or higher than T_INET_1.

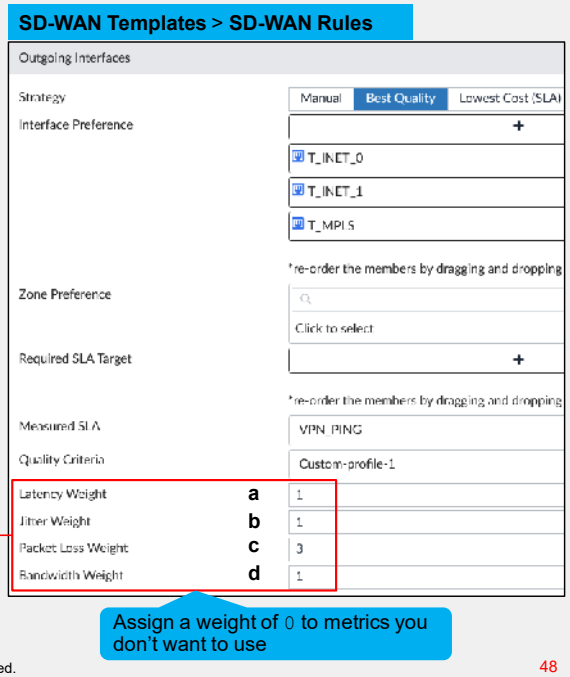
DO NOT REPRINT
© FORTINET

Best Quality—Custom-profile-1

- Link quality is a weight-based composite value of:
  - Latency
  - Jitter
  - Packet loss
  - Bibandwidth
- Formula:

Link quality index = (**a***latency)+(**b***jitter)+(**c***packet loss)+(**d**/bibandwidth)

- Lower link quality index, higher preference



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

48

The custom-profle-1 metric uses a formula that calculates a composite value out of the latency, jitter, packet loss, and bibandwidth. The composite value is called the link quality index, and the formula to calculate it is shown on this slide.

For each metric, you set the weight to use in the formula. The weight enables you to influence the link quality index by one or more metrics. The higher the metric weight, the more influence the metric has on the link quality index. For example, if you consider packet loss the most important metric, then you should assign the highest weight to it. On the contrary, if you just want to ignore a metric, assign it a weight of `0`.

In the example shown on this slide, the link quality index considers all four metrics. However, packet loss has the highest weight among the four.

DO NOT REPRINT
© FORTINET

Best Quality—Custom-profile-1 (Contd)

- Custom-profile-1 example:

```
# diagnose sys sdwan service

Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(10), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(priority), link-cost-factor(custom-profile-1), link-cost-
threshold(10), heath-check(VPN_PING)
  Members(3):
    1: Seq_num(3 T_INET_0), alive, custom1: 2.204: 0.000% 1.732 0.472 20460Kbps, selected
    2: Seq_num(4 T_INET_1), alive, custom1: 2.091: 0.000% 1.562 0.529 20460Kbps, selected
    3: Seq_num(5 T_MPLS), alive, custom1: 19.711: 6.000% 1.303 0.409 19999980Kbps, selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        10.0.0.0-10.255.255.255
```

Link quality index

Packet loss, latency, jitter, and bibandwidth

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

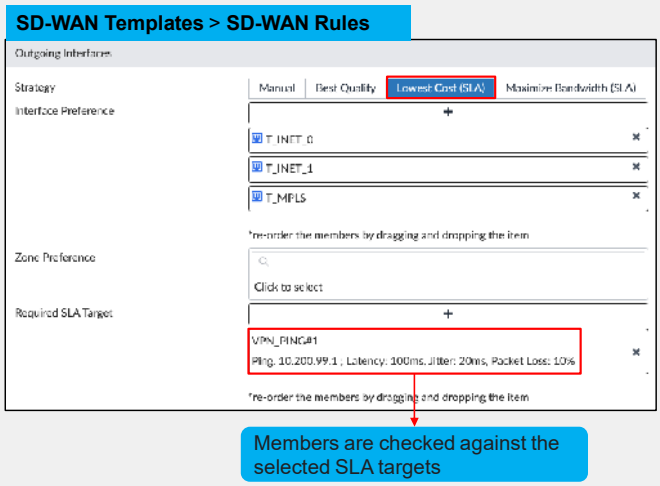
49

This slide shows a sample output for the `diagnose sys sdwan service` command when using custom-profile-1 as the metric. The output includes the link quality index value, followed by the measured values for packet loss, latency, jitter, and bibandwidth.

DO NOT REPRINT
© FORTINET

Lowest Cost (SLA)

- Factors considered for preferred member:
  - SLA targets
    - The more SLA targets a member meets, the higher its preference
  - Member cost
    - Retrieved from member configuration
  - Member priority
    - Based on the rule Interface Preference configuration
- Single preferred member (no load balancing)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

The lowest cost (SLA) strategy instructs FortiGate to select the preferred member based on the following factors:

1. SLA target: You must select one or more SLA targets. FortiGate then checks whether the members meet the selected SLA targets.

2. Member cost: The cost that you configured for the member. The default cost for members is 0.

3. Member priority: Based on the rule **Interface Preference** configuration. Members that are configured first have higher priority over those configured last.

FortiGate first checks how many SLA targets a member meets. The more SLA targets it meets, the higher its preference. Note that even though you can select one or more SLA targets in a rule, each selected SLA target must be from a different performance SLA. That is, each SLA target must point to a different server.

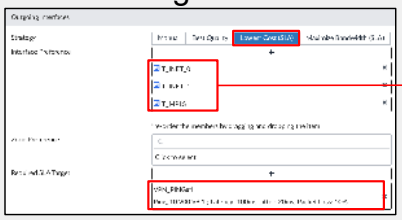
If there are two or more members that meet the same number of SLA targets, then FortiGate uses the member cost as the tiebreaker, and then the member priority as the last tiebreaker.

In addition, one member only—the preferred member (if acceptable)—is used to steer traffic. That is, no load balancing is performed. Also, if none of the members meet the SLA target, FortiGate still builds the `oif` list based on members cost and priority.

DO NOT REPRINT
© FORTINET

Lowest-Cost (SLA) Member—Example 1

- Rule configuration:



Interface configuration order; **T_INET_0** has the highest priority

- Member priority as tiebreaker:

Lowest Cost (SLA) mode is displayed as `sla` in the CLI

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(sla), sla-compare-order
  Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x1), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(1), local cost(0), selected
    3: Seq_num(5 T_MPLS), alive, sla(0x1), gid(0), cfg_order(2), local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

SLA map

- All members meet SLA target (**VPN_PING#1**) and have the same cost (`0`)
- Use priority (`cfg_order`) as tiebreaker
- Resulting order: 1: T_INET_0
  2: T_INET_1
  3: T_MPLS

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

51

This slide shows an example of the preferred member election in a lowest cost (SLA) rule.

The configuration shows **T_INET_0** as the highest priority member. The output of the `diagnose sys sdwan service` command shows `sla(0x1)` for all members. This means that all members meet the SLA target (**VPN_PING#1**)..

Because all members meet the SLA target and have the same cost, then the member priority is used as the tiebreaker. The result is that the members are ordered in the same way they are configured and, therefore, T_INET_0 becomes the preferred member.

DO NOT REPRINT
© FORTINET

Lowest-Cost (SLA) Member—Example 2

- Member cost as tiebreaker:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(5), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(sla), sla-compare-order
  Members(3):
    1: Seq_num(5 T_MPLS), alive, sla(0x1), gid(0), cfg_order(2), local cost(1), selected
    2: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(1), local cost(2), selected
    3: Seq_num(3 T_INET_0), alive, sla(0x0), gid(0), cfg_order(0), local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        10.0.0.0-10.255.255.255
```

- T_INET_0 doesn't meet the SLA target and is placed last in the list
- T_MPLS is better than T_INET_1 because it has a lower cost
- Resulting order: 1: T_MPLS
  2: T_INET_1
  3: T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

52

This slide shows another example of the preferred member election in a lowest cost (SLA) rule.

T_INET_0 doesn't meet the SLA target—reported as sla(0x0), which is why is placed at the bottom of the list. T_MPLS and T_INET_1 meet the SLA target. However, T_MPLS is placed on top because it has a lower cost (1) than T_INET_1 (2). The result is that T_MPLS becomes the preferred member.

DO NOT REPRINT
© FORTINET

Lowest-Cost (SLA) Member—Example 3

- All members fail the SLA target:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(21), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(sla), sla-compare-order
Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x0), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(5 T_MPLS), alive, sla(0x0), gid(0), cfg_order(2), local cost(1), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x0), gid(0), cfg_order(1), local cost(2), selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        10.0.0.0-10.255.255.255
```

- None of the members meet the SLA target
- Choose preferred member (T_INET_0) based on cost and priority
- Resulting order: 1: T_INET_0
  2: T_MPLS
  3: T_INET_1

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

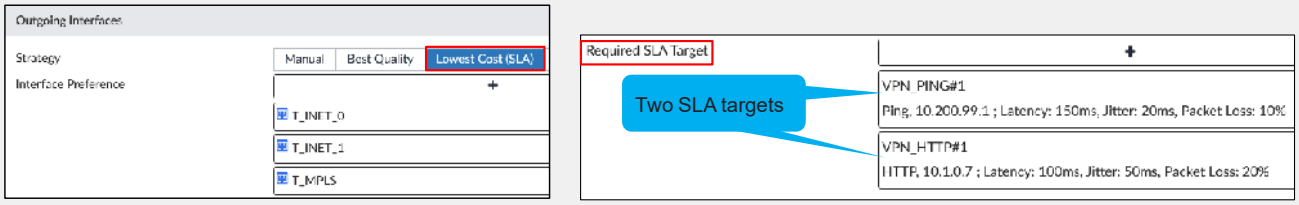
This slide shows another example of the preferred member election in a lowest cost (SLA) rule.

None of the members meet the SLA target. Therefore, the list order is decided first based on cost, and then priority. T_INET_0—the highest priority member—is placed on the top of the list, then T_MPLS (cost=1) and T_INET_1 (cost=2).

DO NOT REPRINT
© FORTINET

Lowest-Cost (SLA) Member—Example 4

- Rule configuration (two SLA targets):



- Only T_MPLS meets both SLA targets:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(7), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(sla), sla-compare-order
  Members(3):
    1: Seq_num(5 T_MPLS), alive, sla(0x3), gid(0), cfg_order(2), local cost(0), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x1), gid(0), cfg_order(0), local cost(0), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x2), gid(0), cfg_order(1), local cost(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

SLA map
(0x3) match both
(0x1) match target 1
(0x2) match target 2
(0x0) no target match

- Only T_MPLS meets both SLA targets (**VPN_PING#1** and **VPN_HTTP1#1**)
- Use priority (`cfg_order`) as tiebreaker for T_INET_0 and T_INET_1
- Resulting order: 1: T_MPLS
  2: T_INET_0
  3: T_INET_1

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

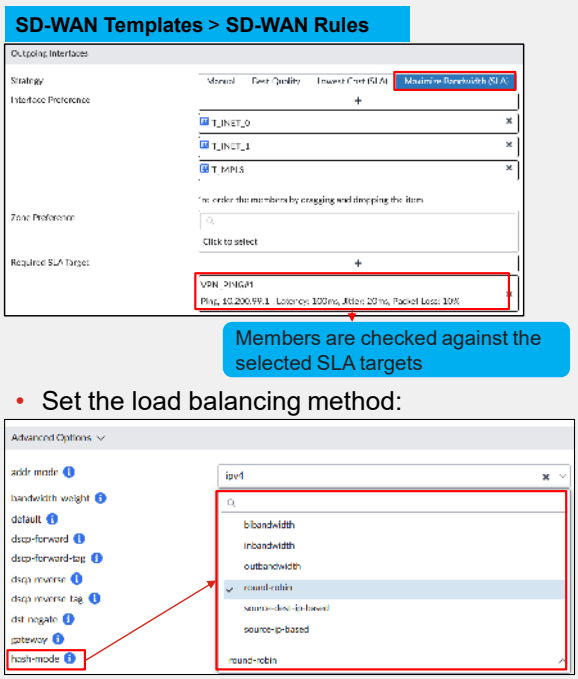
This slide shows an example of the preferred member election in a lowest-cost (SLA) rule when two SLA targets are used.

Because T_MPLS is the only member that meets both SLA targets, it becomes the preferred member. Both T_INET_0 and T_INET_1, meet only one SLA target, and both have the same cost, therefore FortiGate uses the priority (`cfg_order`) as the tiebreaker.

DO NOT REPRINT
© FORTINET

Maximize Bandwidth (SLA)

- Load balance sessions across preferred members
  - Members that meet more SLA targets are preferred
- Load balancing methods:
  - `round-robin` (default)
    - Equal and circular distribution
  - `source-ip-based`
    - Traffic from a source IP is sent to the same member
  - `source-dest-ip-based`
    - Traffic from a source IP to a destination IP is sent to the same member
  - `inbandwidth`
    - Use member with the most available incoming bandwidth
  - `outbandwidth`
    - Use member with the most available outcoming bandwidth
  - `bibandwidth`
    - Use member with the most available bidirectional bandwidth



- Set the load balancing method:

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

55

The maximize bandwidth (SLA) strategy instructs FortiGate to distribute sessions across preferred members. Member preference is based on the number of SLA targets that members meet. That is, members that meet the most SLA targets are preferred, and are therefore used for session distribution.

After FortiGate identifies the preferred members, it load balances sessions across the members using one of the following algorithms:

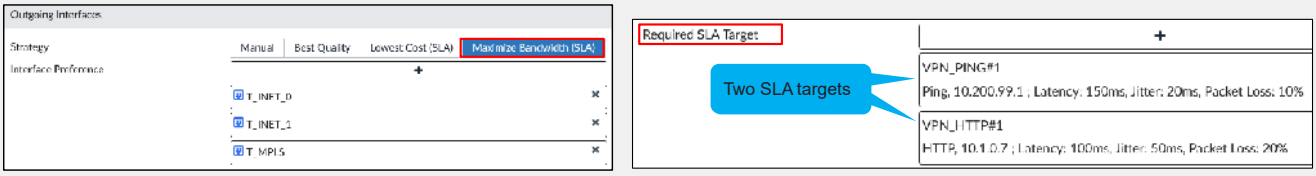
- `round-robin`: This is the default algorithm. FortiGate load balances sessions across members equally using a circular distribution.
- `source-ip-based`: FortiGate sends sessions sourced from an IP address to the same member.
- `source-dest-ip-based`: FortiGate sends sessions with the same source and destination IP pair to the same member.
- `inbandwidth`: FortiGate sends sessions to the member with the most available incoming bandwidth.
- `outbandwidth`: FortiGate sends sessions to the member with the most available outgoing bandwidth.
- `bibandwidth`: FortiGate sends sessions to the member with the most available bidirectional bandwidth.

Note that FortiGate doesn't consider the member cost and priority. FortiGate considers only the number of SLA targets the member meets. Also note that the load balancing method, called hash-mode, is an advanced option. You can configure the setting as shown on this slide.

DO NOT REPRINT
© FORTINET

Maximize Bandwidth (SLA)—Example 1

- Rule configuration:



- All members meet all SLA targets :

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(16), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Members(3):
    1: Seq_num(5 T_MPLS), alive, sla(0x3), gid(3), num of pass(2), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x3), gid(3), num of pass(2), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x3), gid(3), num of pass(2), selected
  Src address(1):
        10.0.1.0-10.0.1.255

  Dst address(1):
        10.0.0.0-10.255.255.255
```

Maximize bandwidth (SLA) mode is displayed as `load-balance` in the CLI

- All members meet the two configured SLA targets (**VPN_PING#1** and **VPN_HTTP#1**)
- Load balance sessions across all members using `round-robin`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

This slide shows an example of the election of preferred members in a maximize bandwidth (SLA) rule.

The output of the `diagnose sys sdwan service` command shows that all members meet the two configured SLA targets, and that the load balancing algorithm configured for the rule is `round-robin`. Because all members meet the same number of SLA targets, FortiGate load balances sessions across all members using the round-robin algorithm. Note that member cost and priority are irrelevant.

DO NOT REPRINT
© FORTINET

Maximize Bandwidth (SLA)—Example 1 (Contd)

- Round-robin distribution for six sessions (two sessions for each member):

```
# diagnose sniffer packet any "port 5201 and udp" 4 0 l | grep T_
Using Original Sniffing Mode
interfaces=[any]
filters=[port 5201 and udp]
1 2023-02-08 03:31:18.160185 T_INET_0 out 10.0.1.101.35239 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.161245 T_INET_0 in 10.1.0.7.5201 -> 10.0.1.101.35239: udp 4
2 2023-02-08 03:31:18.161878 T_INET_1 out 10.0.1.101.38872 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.162679 T_INET_1 in 10.1.0.7.5201 -> 10.0.1.101.38872: udp 4
3 2023-02-08 03:31:18.163186 T_MPLS out 10.0.1.101.56150 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.164248 T_MPLS in 10.1.0.7.5201 -> 10.0.1.101.56150: udp 4
4 2023-02-08 03:31:18.164760 T_INET_0 out 10.0.1.101.57251 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.165681 T_INET_0 in 10.1.0.7.5201 -> 10.0.1.101.57251: udp 4
5 2023-02-08 03:31:18.166217 T_INET_1 out 10.0.1.101.56499 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.167065 T_INET_1 in 10.1.0.7.5201 -> 10.0.1.101.56499: udp 4
6 2023-02-08 03:31:18.167559 T_MPLS out 10.0.1.101.53358 -> 10.1.0.7.5201: udp 4
  2023-02-08 03:31:18.168625 T_MPLS in 10.1.0.7.5201 -> 10.0.1.101.53358: udp 4
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

57

This slide shows the distribution of six sessions that match the rule configured in example 1. You can see that FortiGate sends one session through each member using a circular distribution (round-robin) for a total of two sessions for each member.

DO NOT REPRINT
© FORTINET

Maximize Bandwidth (SLA)—Examples 2 and 3

- One preferred member:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(18), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Members(3):
    1: Seq_num(5 T_MPLS), alive, sla(0x3), gid(3), num of pass(2), selected
    2: Seq_num(4 T_INET_1), alive, sla(0x1), gid(2), num of pass(1), selected
    3: Seq_num(3 T_INET_0), alive, sla(0x1), gid(2), num of pass(1), selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

- T_MPLS meets the most SLA targets
- Send sessions to T_MPLS only

- None of the members meet any of the SLA targets:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(3), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x0), gid(1), num of pass(0), selected
    2: Seq_num(4 T_INET_1), alive, sla(0x0), gid(1), num of pass(0), selected
    3: Seq_num(5 T_MPLS), alive, sla(0x0), gid(1), num of pass(0), selected
  Src address(1):
        10.0.1.0-10.0.1.255
  Dst address(1):
        10.0.0.0-10.255.255.255
```

- All members have the same SLA status
- Load balance sessions across all members using `round-robin`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

58

This slide shows two more examples of the preferred members selection in a maximize bandwidth (SLA) rule. The configuration is the same as the one shown on the previous slide.

The first example on this slide shows T_MPLS as the only preferred member. This is because T_MPLS is the member that meets the most SLA targets. As a result, sessions are sent to T_MPLS only.

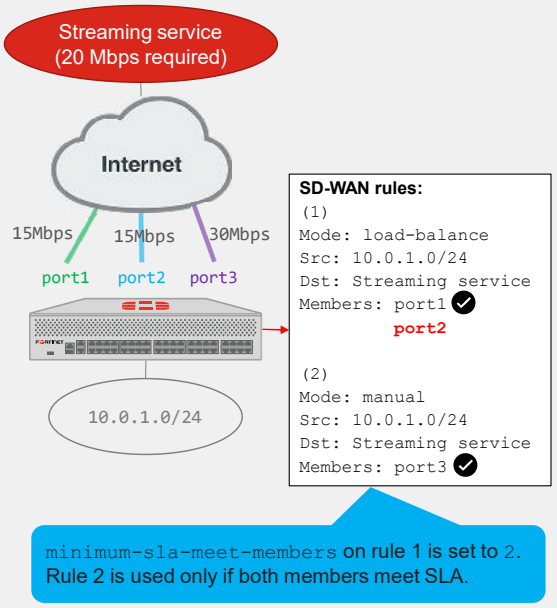
In the second example, none of the members meet any of the SLA targets. Because all members have the same SLA status, then FortiGate load balances sessions across all of them.

DO NOT REPRINT
© FORTINET

Consider Number of Members That Meet SLA

- Define minimum number of members that meet *at least one* SLA target
- Supported by lowest cost (SLA) and maximize bandwidth (SLA) modes only
- Maximum bandwidth (SLA)
  - Useful to ensure minimum bandwidth for service
- GUI configuration
  - SD-WAN rule advanced option
- CLI configuration (default = `0`; no check):

```
config system sdwan
    config service
        edit <id>
            set minimum-sla-meet-members <num>
        next
    end
end
```



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

59

For lowest cost (SLA) and maximize bandwidth (SLA) modes, you can define the minimum number of members that must meet *at least one* of the SLA targets configured in a rule, so the rule remains active. If the number of members that meet the SLA is below the minimum threshold, the rule is disabled and skipped during the rule matching stage. The setting that controls this behavior is `minimum-sla-meet-members` and it is set to `0` by default, which means that the minimum number of members is not considered.

The `minimum-sla-meet-members` setting is particularly useful for maximum bandwidth (SLA) rules, because it can be used to ensure that the bandwidth required for a given service is available. Consider the example on this slide, which shows a FortiGate with three internet underlays, each of them supporting a different speed. The example also includes a simplified version of two SD-WAN rules configured to steer traffic to an internet streaming service. Because port3 is an expensive link, the administrator wants to use port3 only as a backup for port1 and port2.

FortiGate then determines that port1 is the only preferred member for rule 1 because port1 is the only member that meets the SLA. If the administrator leaves `minimum-sla-meet-members` to `0` (default value), then FortiGate would use rule 1 for the streaming service, which would then result in a service impact because the minimum required bandwidth (20 Mbps) is not available. However, if the administrator sets `minimum-sla-meet-members` to `2`, then FortiGate disables rule 1 because the rule doesn't meet the minimum number of members that pass at least one SLA. The result is that the administrator can now have FortiGate use rule 2 for the streaming service only when rule 1 doesn't have enough bandwidth.

DO NOT REPRINT
© FORTINET

Effect of minimum-sla-meet-members

- Set to 0 (no check; only T_INET_0 is used to load balance traffic):

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla [...]
  Gen(8), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x3), gid(3), num of pass(2), selected
    2: Seq_num(5 T_MPLS), alive, sla(0x1), gid(2), num of pass(1), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x0), gid(1), num of pass(0), selected
```

T_INET_0 meets two SLA targets
T_MPLS meets one SLA target
T_INET_1 doesn't meet any SLA targets

- Set to 2 (no change in rule status):

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla [...]
  Gen(8), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x3), gid(3), num of pass(2), selected
    2: Seq_num(5 T_MPLS), alive, sla(0x1), gid(2), num of pass(1), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x0), gid(1), num of pass(0), selected
```

- Set to 3 (rule is disabled and skipped):

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla [...]
Gen(3), TOS(0x0/0x0), Protocol(17: 1->65535), Mode(load-balance  hash-mode=round-robin)
  Service disabled caused by no minimum SLA meet.
  Members(3):
    1: Seq_num(3 T_INET_0), alive, sla(0x3), gid(3), num of pass(2), selected
    2: Seq_num(5 T_MPLS), alive, sla(0x1), gid(2), num of pass(1), selected
    3: Seq_num(4 T_INET_1), alive, sla(0x0), gid(1), num of pass(0), selected
```

Rule is disabled and skipped during rule matching

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

60

This slide shows the effect of the `minimum-sla-meet-members` setting based on its value and the rule SLA status.

The rule status indicates that T_INET_0 passes two SLA targets, T_MPLS one SLA target, and T_INET_1 no SLA targets. If the number of SLA targets for each member doesn't change, then the effect of the `minimum-sla-meet-members` setting value is as follows:

- If set to 0 (default value), FortiGate doesn't enforce a minimum threshold. The rule is active and T_INET_0 becomes the preferred member because it has the best SLA status.
- If set to 2, FortiGate checks if the number of members that meet at least one SLA is equal to or higher than 2. Because T_INET_0 and T_MPLS meet at least one SLA target, then the rule passes the minimum threshold and remains active, and T_INET_0 continues to be the preferred member.
- If set to 3, FortiGate checks if the number of members that meet at least one SLA is equal to or higher than 3. Because only two members meet at least one SLA target, then the rule is disabled because it doesn't pass the minimum threshold. The rule is also skipped during the rule matching process.

DO NOT REPRINT
© FORTINET

Understand SLA Target Maps

- `sla_map` uses a bitmask to reference SLA targets and indicates their status
  - The `sla_map` value represents the hexadecimal value of a binary number
  - Number of bits = number of configured SLA targets
  - First configured SLA target is assigned bit 0, second configured SLA target bit 1, and so on
  - Bit of SLA target is set to 1 if met, otherwise to 0
- Possible `sla_map` values for three SLA targets:

sla_map (hex)	SLA Target #3 (Bit 2) Value = 2^2 = 4	SLA Target #2 (Bit 1) Value = 2^1 = 2	SLA Target #1 (Bit 0) Value = 2^0 = 1
0x7	Pass (1)	Pass (1)	Pass (1)
0x6	Pass (1)	Pass (1)	Fail (0)
0x5	Pass (1)	Fail (0)	Pass (1)
0x4	Pass (1)	Fail (0)	Fail (0)
0x3	Fail (0)	Pass (1)	Pass (1)
0x2	Fail (0)	Pass (1)	Fail (0)
0x1	Fail (0)	Fail (0)	Pass (1)
0x0*	Fail (0)	Fail (0)	Fail (0)

* Also means that there are no SLA targets configured

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

61

With the previous strategies description, Lowest Cost (SLA) and Maximize Bandwidth (SLA), you discovered the `sla_map` field in the output of the `diagnose sys sdwan health-check status` command. This field indicates whether the member meets the configured SLA targets or not. The `sla_map` field uses a bitmask representation to reference the SLA targets and their status. Follow these rules to understand the `sla_map` field:

- The `sla_map` value represents the hexadecimal value of a binary number.
- The number of bits in the binary number equals the number of configured SLA targets.
- The first configured SLA target is assigned bit 0, the second configured SLA target bit 1, and so on.
- If the member meets the SLA target, the bit of the SLA target is set to 1, otherwise to 0.

This slide shows a table with all the possible `sla_map` values when you configure three SLA targets for a member. For example, an `sla_map` of `0x6` means that SLA targets 3 and 2 are met, but not SLA target 1 (6 = 4 + 2 + 0). Expand or reduce the table according to the number of SLA targets configured in your setup.

Note that an `sla_map` of `0x0` has two possible meanings. If you configure one or more SLA targets for a member, `0x0` means that none of the SLA targets are met. However, `0x0` can also mean that you didn't configure any SLA targets for the member. Therefore, make sure to check the performance SLA configuration to identify whether there are SLA targets configured or not.

DO NOT REPRINT
© FORTINET

Review

- ✓ Describe the rule operation and lookup process
- ✓ Configure FortiGate to filter preferred members based on best route
- ✓ Describe the implicit rule and available load balancing algorithms
- ✓ Describe the supported criteria for matching SD-WAN traffic
- ✓ Understand application steering, its learning phase, and expected behavior
- ✓ Describe internet services and their role in rules
- ✓ Understand strategies and the preferred member election process
- ✓ Identify preferred members when monitoring the rule status

FORTINET Training Institute

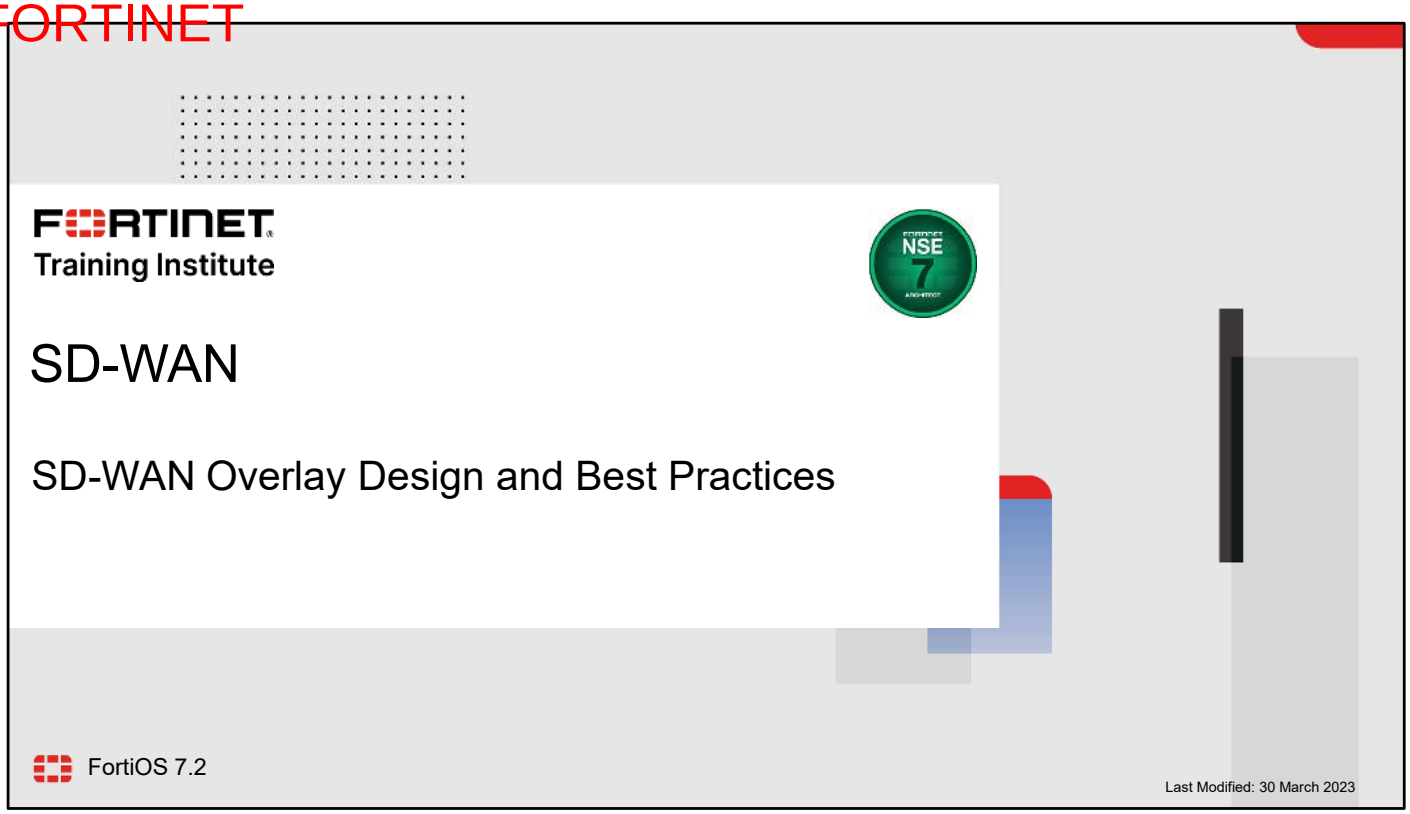
© Fortinet Inc. All Rights Reserved.

62

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

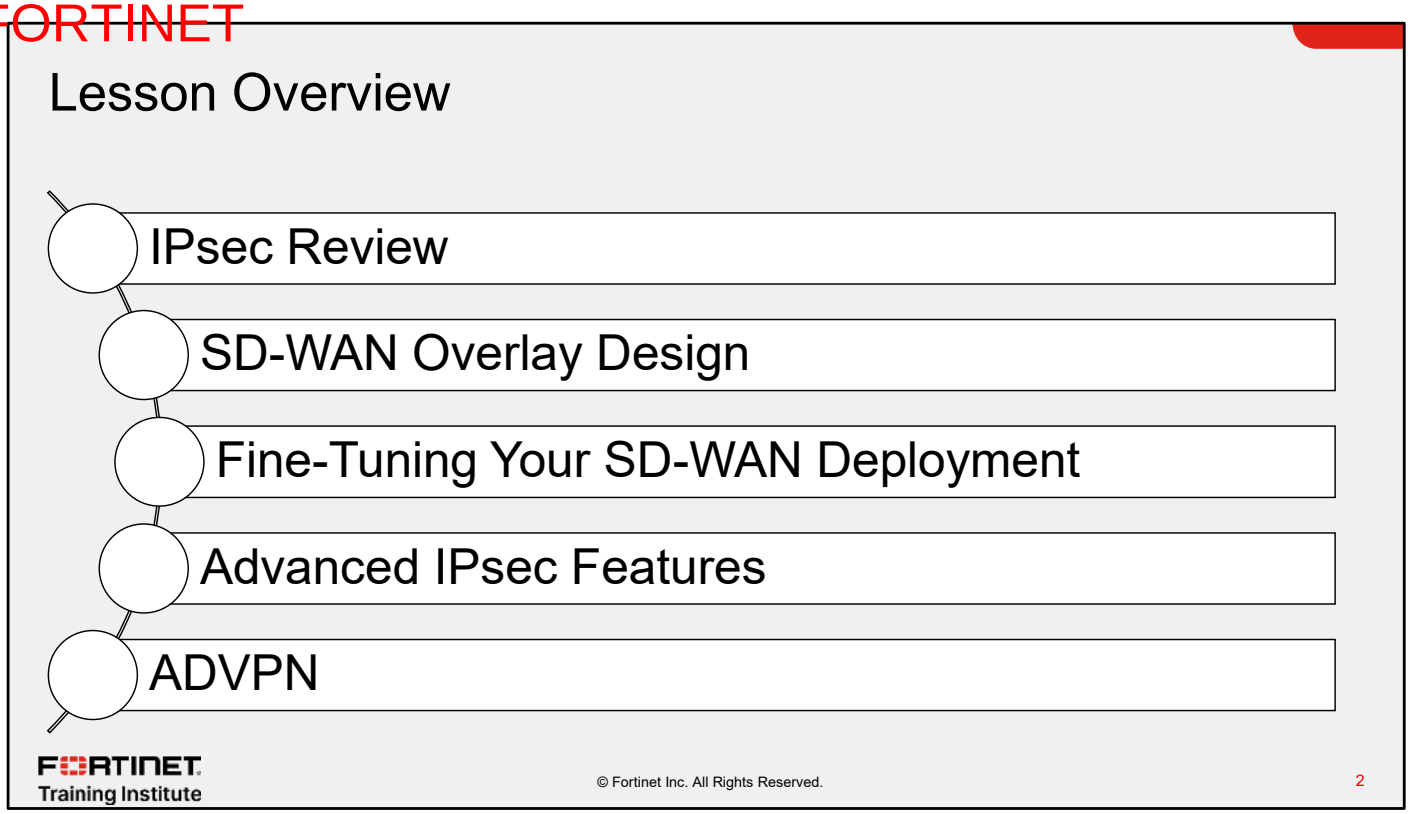
By mastering the objectives covered in this lesson, you learned how SD-WAN rules work based on the criteria and strategy configured.

DO NOT REPRINT
© FORTINET



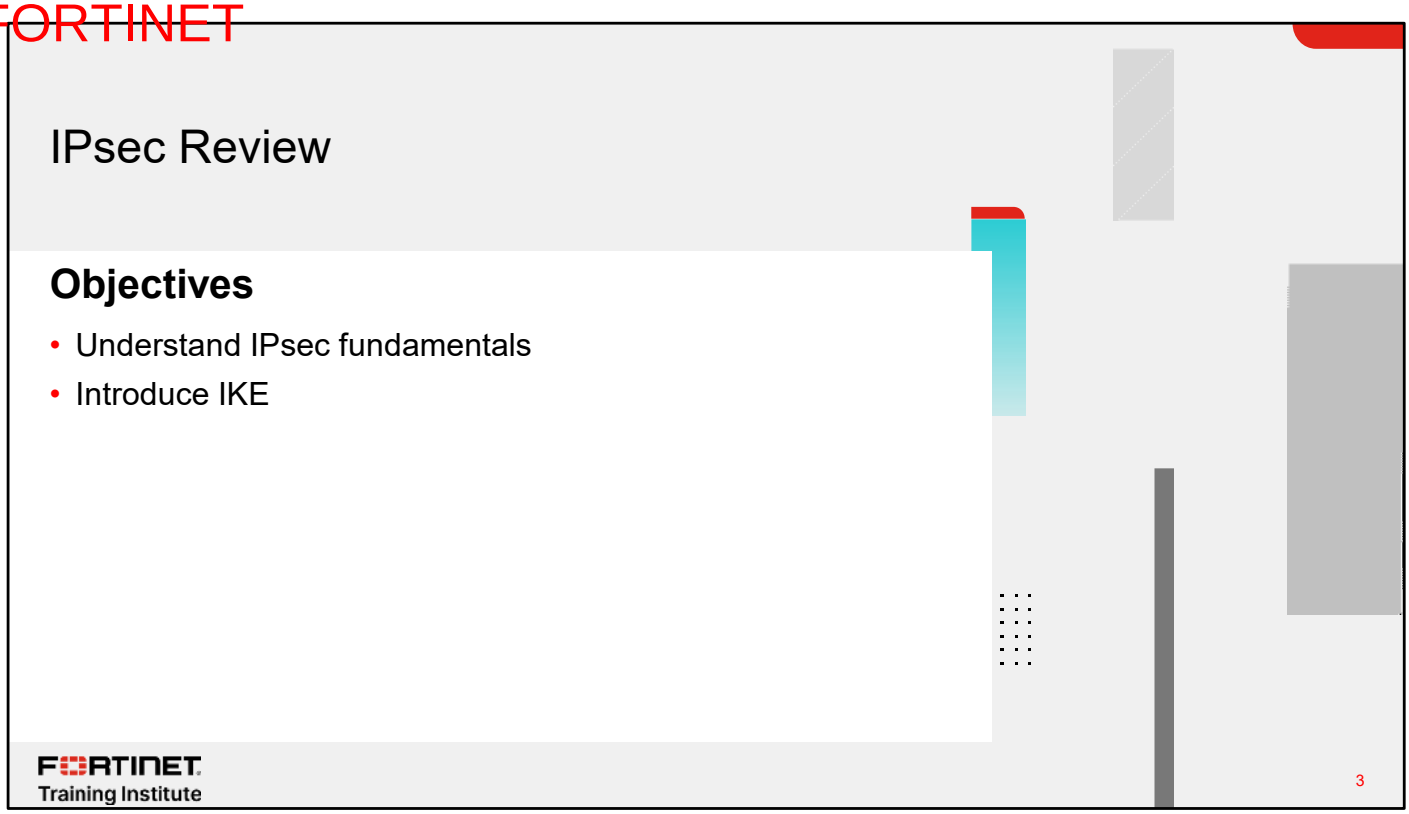
In this lesson, you will learn about the SD-WAN overlay design and advanced IPsec features that are useful for SD-WAN deployments.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



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

By reviewing IPsec, you will be able to recall the basics of the protocol.

DO NOT REPRINT
© FORTINET

IPsec Review

- Suite of protocols for securing IP communications
  - Most used tunneling protocol for SD-WAN overlays
- Authenticates and/or encrypts packets:
  - Internet key exchange (IKE)
    - UDP 500 (no NAT-T)
    - UDP 4500 (NAT-T)
  - Encapsulating security payload (ESP)
    - Provides both data integrity and encryption
    - IP proto 50 (no NAT-T)
    - UDP 4500 (NAT-T)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

IPsec is a suite of protocols for authenticating and encrypting traffic between two peers. Because of its security features, IPsec is the tunneling protocol that is most used for building SD-WAN overlays.

The two most-used protocols in the IPsec suite are:

- IKE, which does the handshake, tunnel maintenance, and disconnection. By default, IKE uses UDP port 500.
- ESP, which ensures data integrity and encryption. ESP uses IP protocol number 50.

The ESP protocol usually has problems crossing devices that are performing NAT. One of the reasons is that ESP does not use port numbers, like TCP and UDP do, to differentiate one tunnel from another. To solve this, NAT traversal (NAT-T) was added to the IPsec specifications. When NAT-T is enabled on both ends, peers can detect when NAT is performed along the path. If NAT is found, then the following occurs on both peers:

- IKE negotiation switches to using UDP port 4500.
- ESP packets are encapsulated in UDP port 4500.

DO NOT REPRINT
© FORTINET

IKE Review

- Negotiates the tunnel private keys and encryption
- Allows parties involved in a transaction to set up their security associations (SAs)
  - SAs are the basis for building security functions into IPsec
  - In normal two-way traffic, the exchange is secured by a pair of IKE SAs
- Phases:
  - Phase 1 (IKEv1) — IKE_SA_INIT and IKE_AUTH exchanges (IKEv2) → Outcome: IKE SA
  - Phase 2 (IKEv1) — CREATE_CHILD_SA exchange (IKEv2) → Outcome: IPsec SA
- Versions
  - IKEv1
    - Legacy, well known
  - IKEv2
    - Simplified negotiation process to create security association, more features
    - Recommended for SD-WAN (network ID feature for ADVPN)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

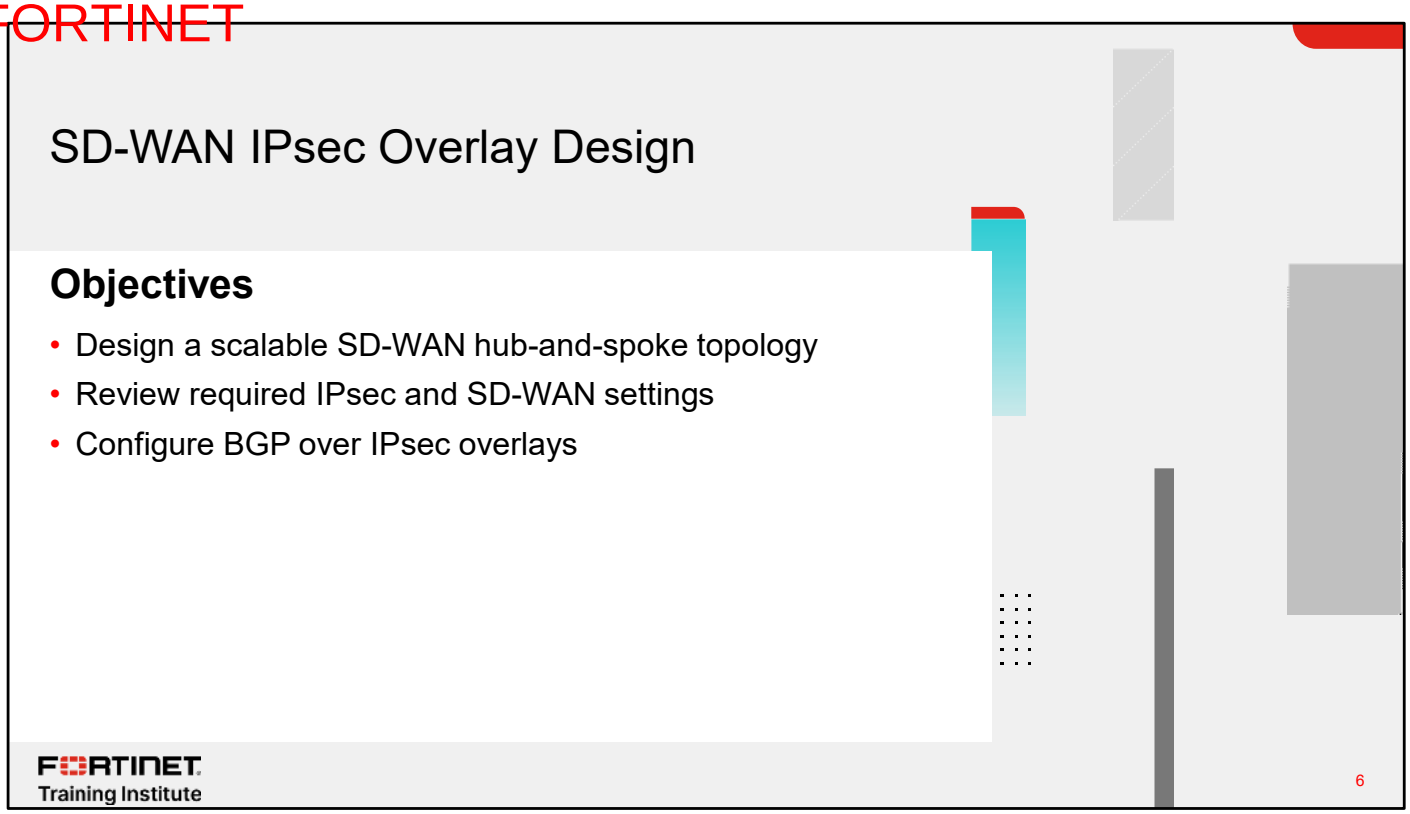
IKE negotiates the private keys and encryption that the device uses to create an IPsec tunnel.

In order to create an IPsec tunnel, both devices must establish their security associations (SAs) and secret keys, which are facilitated by the IKE protocol. The SAs are the basis for building security functions into IPsec. An SA is simply the bundle of algorithms and parameters being used to encrypt and authenticate data travelling through the tunnel. In normal two-way traffic, this exchange is secured by a pair of SAs: one for each traffic direction. Essentially, both sides of the tunnel must agree on the security rules. If both sides cannot agree on the rules for sending data and verifying each other's identity, then the tunnel is not established. SAs expire and need to be renegotiated by the peers after they have reached their lifetime.

IKE uses two distinct phases: phase 1 and phase 2. Each phase negotiates different SA types. The SA negotiated during phase 1 is called IKE SA, and the SA negotiated during phase 2 is called IPsec SA. IKE SAs are used for setting up a secure channel to negotiate IPsec SAs. IPsec SAs are used for encrypting and decrypting the data sent and received, respectively, through the tunnel.

The are two IKE versions: IKEv1 and IKEv2. IKEv1 is the legacy, well known version of the protocol. Though still present on the field, it is deprecated and should not be used for new deployments. Administrators should move to IKEv2, because it supports more features and is easier to operate. For example, IKEv2 exclusively supports the network ID feature, which enables the administrator to establish multiple tunnels between the same local and remote gateways, which can be required during failover scenarios involving SD-WAN and ADVPN. You will learn more about the network ID feature and ADVPN in this lesson.

DO NOT REPRINT
© FORTINET



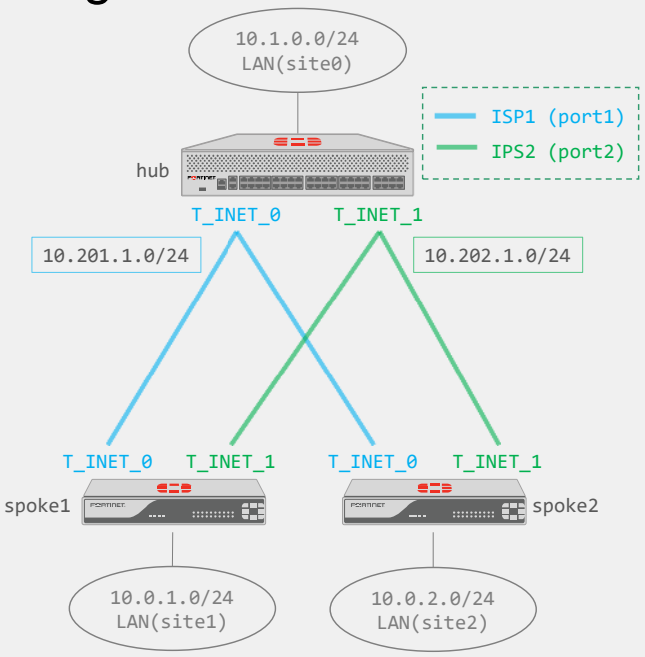
After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in SD-WAN IPsec overlay design, you will be able to understand the configuration required to deploy a scalable SD-WAN network that uses IPsec and BGP as base building blocks.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design With IPsec

- Hub-and-spoke topology
  - Dial-up tunnels
    - Hub: server
    - Spokes: client
  - Group overlays per ISP (higher performance)
- SD-WAN
  - Usually deployed on the spokes only
  - Sometimes required on the hub too
  - Support for ADVPN
- Routing
  - Dynamic routing is preferred (more scalable)
  - Multiple paths (all sites exchange their prefixes)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

The topology shown on this slide represents a standard SD-WAN deployment using IPsec overlays.

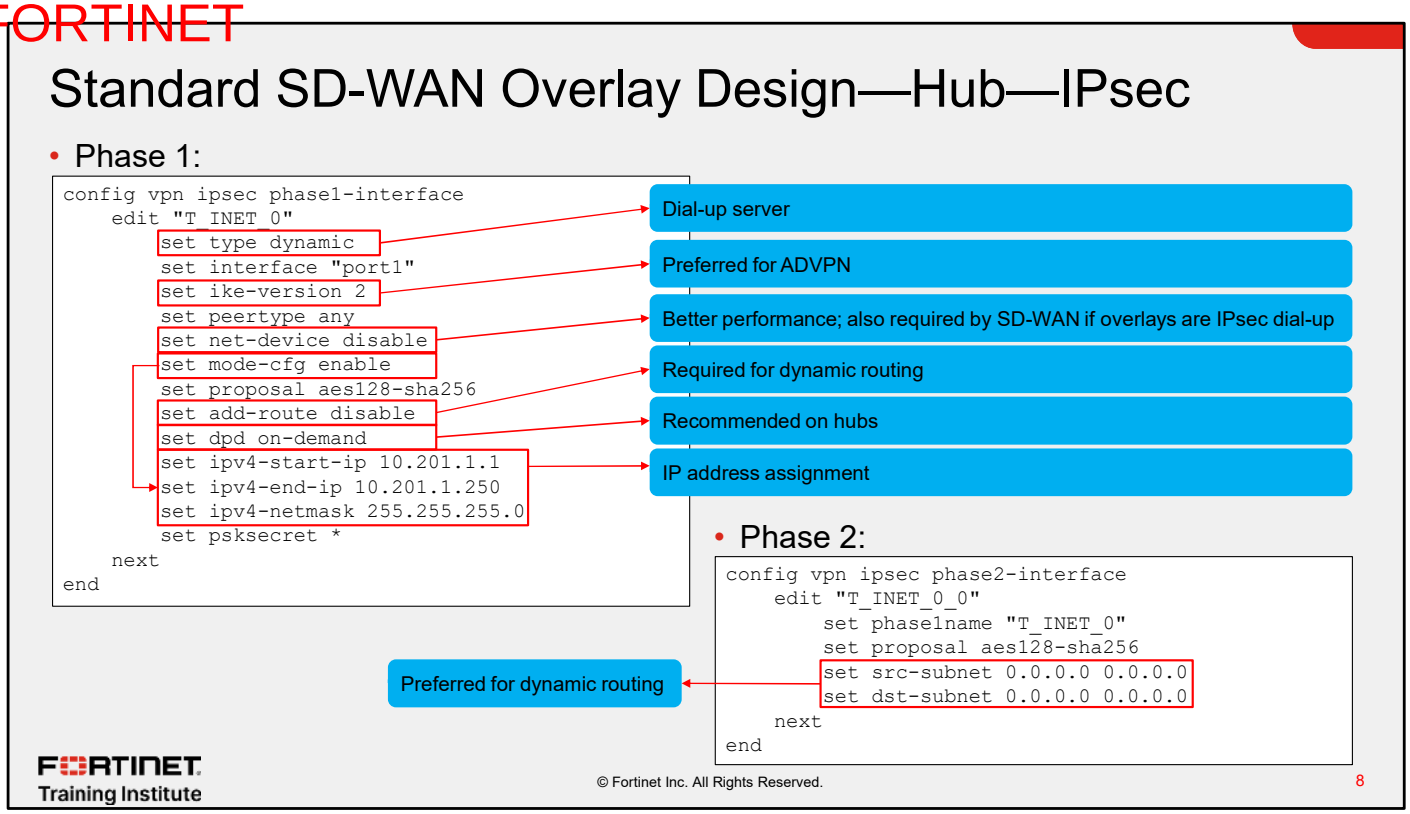
The IPsec overlays follow a hub-and-spoke design. The hub is located at the customer's central office or data center, and the spokes (also known as branches or edges) are distributed across all remote sites (branch offices, retail stores, remote workers, and so on). The hub acts as a dial-up IPsec server for the spokes and has a separated dial-up server configuration for each underlay interface. The result is that each dial-up server configuration defines a point-to-multipoint (PTMP) overlay. To have multiple alternative paths to all subnets in the network, each spoke usually connects to all overlays.

For better performance, it's recommended to group overlays per ISP. That is, it's better to deploy same-ISP overlays than cross-ISP overlays. By connecting tunnels over the same ISP, you should get better performance because the latency is usually lower than the latency of packets crossing multiple ISPs.

Most of the IPsec overlay traffic is initiated from the spokes to the hub. Spokes access workloads protected by the hub. The workloads are located on the hub, or on a network that is reachable through the hub. Because most of the traffic is initiated in the spoke-to-hub direction, SD-WAN is usually deployed on the spokes only. However, in some cases, traffic is also initiated in the hub-to-spoke direction, which is why SD-WAN may also be required on the hub side.

A dynamic routing protocol, such as BGP, is used to exchange routing information through the overlays. The topology also shows the IP subnets used for the overlays. The overlays established over the ISP1 underlay (port1) are assigned an address in the `10.201.1.0/24` subnet, and the overlays established over the ISP2 underlay (port2) are assigned an address in the `10.202.1.0/24` subnet. The goal is for all sites to exchange their prefixes over all available overlays, and to support ADVPN with SD-WAN if needed. Although static routing is possible, dynamic routing is preferred because of scalability reasons.

DO NOT REPRINT
© FORTINET



This slide shows an example of the IPsec configuration required on the hub for the dial-up ISP1 overlays. It is assumed that the reader is familiar with IPsec, so this lesson will highlight only the relevant settings for SD-WAN, without providing many details about them.

The tunnel is bound to the `port1` underlay. The configuration is the same for the dial-up ISP2 overlays, except the binding interface (port2) and the mode config address range (10.202.1.0/24). The type of the tunnel is set to `dynamic` (dial-up server). That is, tunnels can be initiated only from the spoke side. Dynamic works better for hub-and-spoke deployments because it reduces the amount of configuration required on the server side, while it enables spokes to have dynamic IP addresses or be located behind NAT devices.

The IKE version is 2. IKEv2 is recommended for increased security level, and for SD-WAN, because of its benefits for ADVPN. Both `net-device` and `add-route` are disabled. The former allows for better performance on large deployments and is also required for adding dial-up IPsec tunnels as SD-WAN members. You can disable `add-route` if you plan to use a dynamic routing protocol to exchange routing information through the tunnels.

For dead peer detection (DPD), it's recommended to use `on-demand` mode on dial-up servers, like the hub. When you use on-demand mode, FortiGate sends DPD probes if there is only outbound traffic through the tunnel, but no inbound. On-demand mode is more convenient on hubs because of the reduced overhead. The more passive a hub is, the better it scales.

To assign IP addresses to overlays on spokes, `mode-cfg` is enabled. The `ipv4-xxx` settings are also defined to indicate the IP address range and netmask to use. The phase 2 configuration is simple. Sites exchange their prefixes using dynamic routing, so it's more convenient to leave the encryption domain open. This way, SD-WAN, the FIB, and the firewall policies dictate the traffic forwarding policy.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Hub—System Interfaces

- IPsec interface:

```
config system interface
    edit "T_INET_0"
        set vdom "root"
        set ip 10.201.1.254 255.255.255.255
        set allowaccess ping
        set type tunnel
        set interface "port1"
    next
    edit "T_INET_1"
        set vdom "root"
        set ip 10.202.1.254 255.255.255.255
        set allowaccess ping
        set type tunnel
        set interface "port2"
    next
end
```

Local overlay IP; used as NH for routing advertisements

Target server configured in spokes performance SLA

- Loopback interface:

```
config system interface
    edit "lo-HC"
        set vdom "root"
        set ip 10.200.99.1 255.255.255.255
        set allowaccess ping
        set type loopback
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

This slide shows an example of the system interfaces configuration on the hub.

Each IPsec interface is assigned an IP address. This is the local IP for the overlay, and it is used by BGP as the next hop (NH) address in advertisements.

A loopback interface is created to act as the target server defined in the performance SLAs on the spokes.

DO NOT REPRINT
© FORTINET

Optional SD-WAN Settings on Hub

- For SD-WAN traffic steering from the hub:

```
config system sdwan
    config zone
        edit "overlay"
        next
    end
    config members
        edit 3
            set interface "T_INET_0"
            set zone "overlay"
        next
        edit 4
            set interface "T_INET_1"
            set zone "overlay"
        next
    end
    config health-check
        edit "spoke1-hc"
            set server "10.0.1.1"
            set members 3 4
        next
        edit "spoke2-hc"
            set server "10.0.2.1"
            set members 3 4
        next
    end
...
```

net-device must be disabled

One performance SLA and one rule per spoke*

```
...
    config service
        edit 1
            set name "spoke1-rule"
            set mode priority
            set dst "spoke1-net"
            set src "all"
            set health-check "spoke1-hc"
            set priority-zone "overlay"
        next
        edit 2
            set name "spoke2-rule"
            set mode priority
            set dst "spoke2-net"
            set src "all"
            set health-check "spoke2-hc"
            set priority-zone "overlay"
        next
    end
end
```

* Reduce the number of performance SLAs and administrative overhead by using passive monitoring and route tags, respectively.

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

This slide shows an example of an SD-WAN configuration for the IPsec dial-up tunnels on the hub. You will use a configuration like this one if you want to use SD-WAN on the hub, especially if the volume of traffic initiated by the hub to the spokes is important.

The IPsec tunnels are placed in the overlay zone. Note that because the tunnels are dynamic, you must first disable `net-device` on their phase 1 configuration. Otherwise, FortiOS does not allow you to configure the tunnels as SD-WAN members.

The rest of the configuration is fairly simple. However, note that you probably must configure at least one performance SLA and a rule for each spoke. As a result, the higher the number of spokes, the larger the configuration. You can reduce the number of performance SLAs by using passive monitoring. This way, FortiGate measures the performance of the overlay based on user traffic going through it, and thus, doesn't require you to configure a target server to send probes to. You can also reduce the administrative overhead of maintaining the destination networks on spokes by using route tags instead of fixed firewall address objects. This way, if the prefixes on the spoke change, then the destination on the rule is automatically updated without you having to make configuration changes.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Hub—BGP

- IBGP configuration:

```
config router bgp
    set as 65000
    set router-id 10.1.0.1
    set ibgp-multipath enable
    config neighbor-group
        edit "Spokes_INET_0"
            set interface "T_INET_0"
            set remote-as 65000
            set update-source "T_INET_0"
            set route-reflector-client enable
        next
        edit "Spokes_INET_1"
            set interface "T_INET_1"
            set remote-as 65000
            set update-source "T_INET_1"
            set route-reflector-client enable
        next
    end
...
```

- Allow ECMP for IBGP routes
- Standard BGP settings for each group of spokes (one group per overlay)
- Bind BGP packets to overlay
- Use, as source IP for BGP packets, the address of the selected interface
- Spokes can learn prefixes from other spokes

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

This slide shows an example of the BGP configuration on the hub. Internal BGP (IBGP) is used—both the local AS and remote AS are the same (65000). IBGP is preferred over external BGP (EBGP) because IBGP preserves the next hop for prefixes, which is convenient for ADVPN.

With SD-WAN, ECMP is expected because of the multipath nature of SD-WAN. If using IBGP, you must enable `ibgp-multipath` to support ECMP for IBGP routes.

Because the remote BGP speakers—the spokes—are dial-up clients, it's convenient to define a `neighbor-group` for overlays established over each underlay. This way, the BGP peerings are always initiated from the spokes—expected because spokes are dial-up clients. FortiGate then applies the common settings in the `neighbor-group` for each BGP peering.

Inside each `neighbor-group` entry, the `interface` and `update-source` settings indicate the source interface and source IP address to use for BGP packets, respectively. This is important because of the multipath nature of SD-WAN. FortiOS performs a route lookup to determine the outgoing interface for BGP packets. Binding the BGP packets to an interface and a source address ensures that there are no cross-overlay BGP peerings between the spoke and the hub.

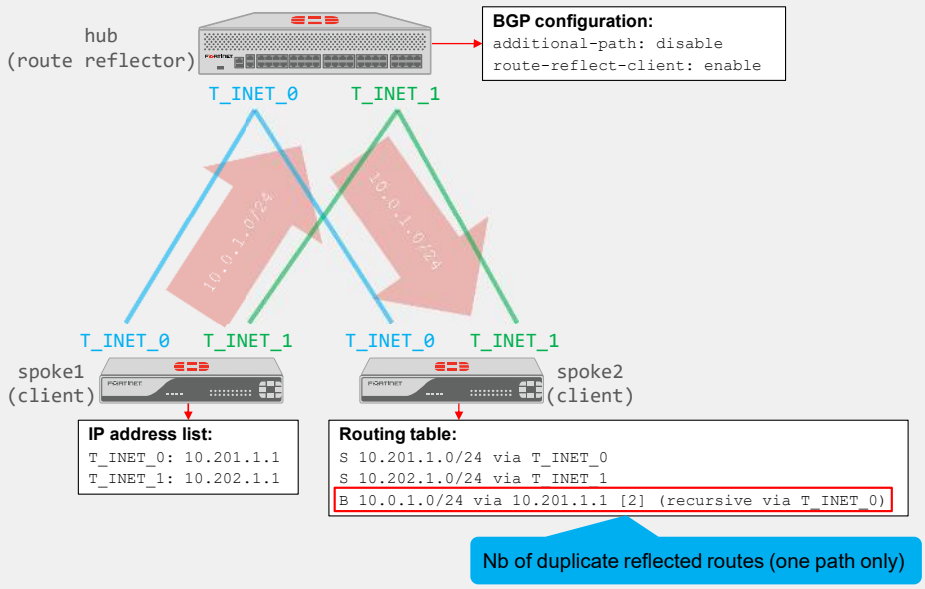
Note that in this setup, `update-source` is redundant and not really required. The `interface` setting automatically forces the source IP address to be that of the selected interface. `update-source` is more useful for cases in which the source IP address for BGP is different from the source interface. A common example is when using a loopback interface address as the source IP. However, even when redundant, it's still a good practice to configure the `update-source` setting for consistency purposes.

With `route-reflector-client` setting enabled, spokes can learn about each other's prefixes.

DO NOT REPRINT
© FORTINET

BGP Route Reflector

- Reflect learned routes from a spoke to other spokes in the group



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

By default, IBGP routers do not pass on routes learned from internal neighbors to other internal neighbors. The `route-reflector-client` setting instructs FortiGate to act as a route reflector, and therefore, to pass on (or reflect) learned routes from an IBGP client (a spoke) to other IBGP clients (other spokes in all neighbor groups). Because, in our SD-WAN topology, you want spokes to learn the prefixes of other spokes, then you must enable `route-reflector-client` on the neighbor groups.

The diagram on this slide shows how BGP route reflector works. The hub, which acts as route reflector, reflects the `10.0.1.0/24` prefix, which is advertised by spoke1, to spoke2. The hub learns the prefix through two overlays—T_INET_0 and T_INET_1—each using a different next hop—`10.201.1.1` and `10.202.1.1`, respectively. That is, the hub learns the prefix through two different paths: one with next hop `10.201.1.1` and another with next hop `10.202.1.1`.

The hub then reflects the prefix to spoke2, which in turn performs a recursive lookup to determine the outgoing interface to use for the prefixes based on the next hop. The recursive lookup matches the static route for the overlay IP subnets, which is why the BGP routes show T_INET_0 as the outgoing interface.

It is important to highlight the following:

- The routing table on spoke2 shows number two [2] to indicate duplicate routes for `10.0.1.0/24`. They are the result of IBGP preserving the next hop. Duplicate routes are visible on FIB and summarized on the routing table for better readability.
- The hub reflects only the path through `10.201.1.1`. It doesn't reflect the path through `10.202.1.1`. This is because the hub is not configured to advertise additional paths for prefixes (`set additional-path disable`).

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Hub—BGP (Contd)

- IBGP configuration:

```
...
    config neighbor-range
        edit 1
            set prefix 10.201.1.0 255.255.255.0
            set neighbor-group "Spokes_INET_0"
        next
        edit 2
            set prefix 10.202.1.0 255.255.255.0
            set neighbor-group "Spokes_INET_1"
        next
    end
    config network
        edit 1
            set prefix 10.1.0.0 255.255.255.0
        next
    end
end
```

Define IP address range for each neighbor group

BGP peers from `10.201.1.0/24` match the `Spokes_INET_0` group

Advertise IGP prefix (connected route) to peers

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

This slide shows the remaining part of the BGP configuration on the hub.

Each entry in `neighbor-range` defines the IP address range to use for each neighbor group. For example, peers with an IP address in the 10.201.1.0/24 subnet are considered part of the `Spokes_INET_0` neighbor group. Usually, the groups correspond to BGP neighbors from overlay tunnels built on the same type of underlay links (Same ISP or BGP). Neighbor-groups are useful to limit the route propagation within a single neighbor group when you activate route-reflection.

Each entry in `config network` indicates the interior gateway protocol (IGP) prefix to inject into the BGP table so it is advertised to peers. IGP refers to non-BGP routes in the routing table, such as OSPF, RIP, connected, and static. Although the hub routing table is not shown on this slide, 10.1.0.0/24 is a connected route, which is why you configure a network entry to advertise it to the spokes.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Spoke—IPsec

- Phase 1:

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        set type static
        set interface "port1"
        set ike-version 2
        set peertype any
        set net-device enable
        set mode-cfg enable
        set proposal aes128-sha256
        set add-route disable
        set dpd on-idle
        set remote-gw 100.64.1.1
        set psksecret *
    next
end
```

Dial-up client; remote gateway IP is known

Required when ADVPN is used with SD-WAN

- Phase 2:

```
config vpn ipsec phase2-interface
    edit "T_INET_0_0"
        set phase1name "T_INET_0"
        set proposal aes128-sha256
        set src-subnet 0.0.0.0 0.0.0.0
        set dst-subnet 0.0.0.0 0.0.0.0
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

This slide shows an example of the IPsec configuration required on the spoke for the ISP1 overlay. Some settings are the same used on the hub (dial-up server), so this lesson focuses on settings that are specific for the spoke.

The tunnel is bound to the `port1` underlay. The configuration is the same for the ISP2 overlay, except the binding interface (port2) and the IP address of the remote gateway. The type of the tunnel is set to `static` (dial-up client) because the IP address of the remote gateway (the hub) is known. In this case, the address corresponds to the address assigned to port1 on the hub.

`net-device` must be enabled if you want to support ADVPN shortcuts in SD-WAN. Although disabling `net-device` provides better performance, enabling the setting should not represent a concern because the number of overlays deployed on the spokes is usually low.

The phase 2 configuration is the same as in the hub. The encryption domain is open to let SD-WAN, the FIB, and the firewall policies dictate the traffic forwarding policy.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Spoke—System Interfaces

- IPsec interface:

```
config system interface
    edit "T_INET_0"
        set vdom "root"
        set allowaccess ping
        set type tunnel
        set interface "port1"
    next
    edit "T_INET_1"
        set vdom "root"
        set allowaccess ping
        set type tunnel
        set interface "port2"
    next
end
```

IP address is assigned automatically through IKE mode config

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

This slide shows an example of the system interfaces configuration on the spokes.

You don’t have to assign an IP address for the overlays. Their IP addresses are obtained using IKE mode configuration.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Spoke—SD-WAN

- SD-WAN:

```
config system sdwan
    config zone
        edit "overlay"
        next
    end
    config members
        edit 3
            set interface "T_INET_0"
            set zone "overlay"
        next
        edit 4
            set interface "T_INET_1"
            set zone "overlay"
        next
    end
    config health-check
        edit "VPN"
            set server "10.200.99.1"
            set protocol ping
            set members 3 4
        next
    end
...
```

Points to loopback address on the hub

```
...
    config service
        edit 1
            set name "Ping"
            set mode priority
            set protocol 1
            set dst "10.0.0.0/8"
            set src "10.0.1.0/24"
            set health-check "VPN"
            set priority-zone "overlay"
        next
    end
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

This slide shows an example of the SD-WAN configuration for the IPsec dial-up tunnels on the spokes.

The configuration is fairly simple. The two overlays are placed in the `overlay` zone. The `VPN` performance SLA uses ping to actively measure the health and performance of the overlays against `10.200.99.1`, which is the address of the loopback interface on the hub.

A best quality SD-WAN rule is configured to steer ICMP traffic from `10.0.1.0/24` to `10.0.0.0/8`. The rule uses the `VPN` performance SLA to determine the best quality member in the `overlay` zone.

DO NOT REPRINT
© FORTINET

Standard SD-WAN Overlay Design—Spoke—BGP

- IBGP configuration:

```
config router bgp
    set as 65000
    set router-id 10.0.1.1
    set ibgp-multipath enable
    config neighbor
        edit "10.201.1.254"
            set interface "T_INET_0"
            set remote-as 65000
            set update-source "T_INET_0"
        next
        edit "10.202.1.254"
            set interface "T_INET_1"
            set remote-as 65000
            set update-source "T_INET_1"
        next
    end
    config network
        edit 1
            set prefix 10.0.1.0 255.255.255.0
        next
    end
end
```

One neighbor per overlay

Advertise IGP prefix (connected route) to peers

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

This slide shows an example of the IBGP configuration on the spokes. Some settings are the same used on the hub, so this lesson focuses only on settings that are specific to the spoke.

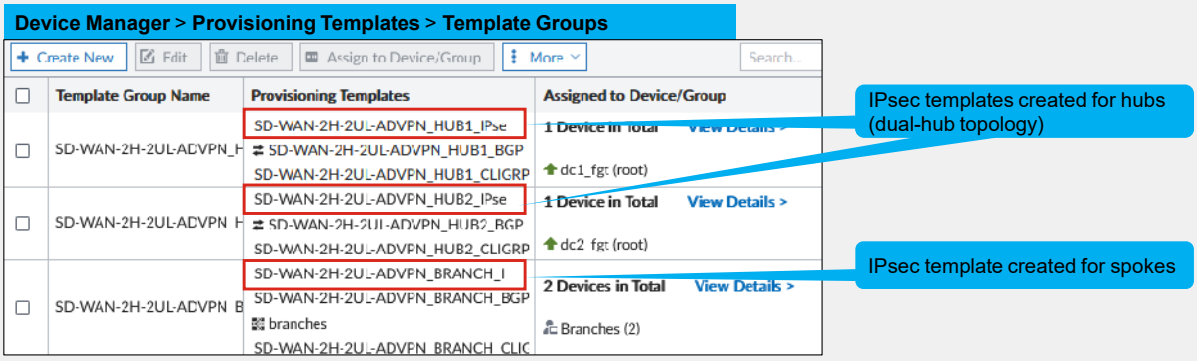
Each entry in `config neighbor` defines the IP address of the BGP neighbor for the overlay. For instance, `10.201.1.254` is the IP address of the hub overlay established over ISP1, and `10.202.1.254` is the IP address of the hub overlay established over ISP2.

The entry in `config network` instructs the spoke to inject the IGP prefix (10.0.1.0/24) into the BGP table so it is advertised to peers. Although the spoke routing table is not shown on this slide, 10.0.1.0/24 is a connected route, which is why we configure a network entry to advertise it to the spokes.

DO NOT REPRINT
© FORTINET

FortiManager—SD-WAN Overlay Template

- Configure IPsec and BGP templates for SD-WAN deployment
- IPsec template parameters and settings correspond to IPsec recommended templates for branches and hubs
- Simplified configuration with recommended settings
- Templates created by SD-WAN overlay template



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

18

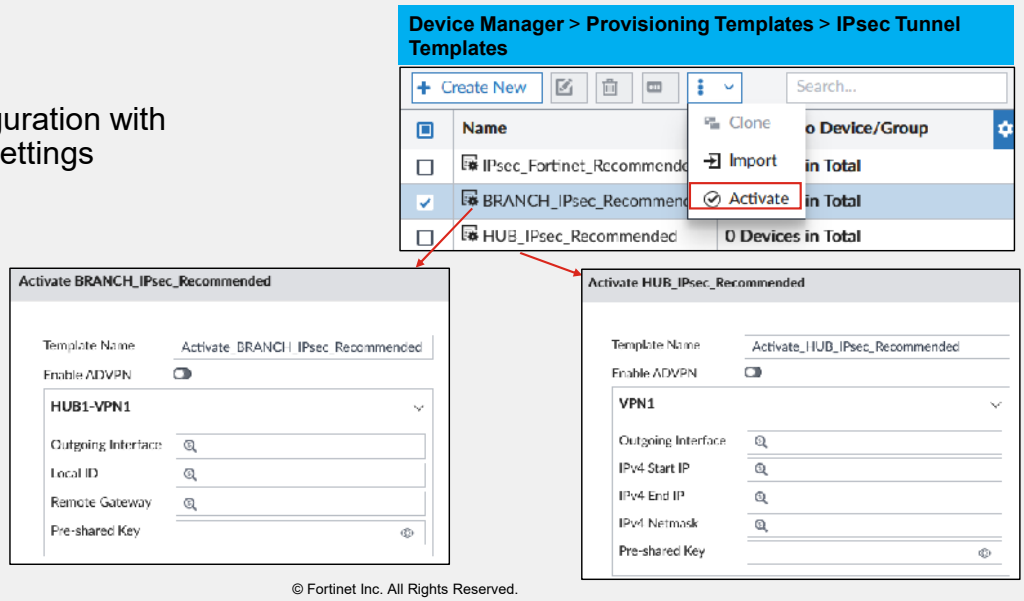
To prepare the IPsec tunnel templates and at the same time, the BGP templates, you can use the FortiManager SD-WAN overlay template.

As explained in a previous lesson, the SD-WAN overlay template will simplify the configuration. You need to enter only the parameters specific to your network, like network IP range, and FortiManager creates IPsec tunnel templates that follow Fortinet recommended best practices.

DO NOT REPRINT
© FORTINET

FortiManager—IPsec Recommended Templates

- Predefined templates
  - IPsec
  - Branch IPsec
  - Hub IPsec
- Simplified configuration with recommended settings
- Activate template



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

If you don't want to use the SD-WAN overlay template to help with IPsec VPN tunnel configuration for the hub-and-spoke topology, you can use the predefined IPsec tunnel templates available on FortiManager. In both cases FortiManager uses the same recommended settings to build the IPsec tunnels.

With the predefined templates, you need to define only the main parameters, as shown on this slide. The templates BRANCH_IPsec_Recommended and HUB_IPsec_Recommended will add the parameters discussed on previous slides, like tunnel type or net-device settings. Later, you can edit the prepared template for fine-tuning and advanced settings.

Note that you can also use those templates to enable ADVPN, and set corresponding recommended parameters, on your hub-and-spokes topology.

DO NOT REPRINT
© FORTINET

FortiManager—IPsec Recommended Templates (Contd)

- Hub template main settings
  - As described in previous slides
    - ike-version 2
    - type **dynamic**
    - net-device **disable**
    - mode-cfg enable
    - add-route disable
    - IP address assignment
  - Additional settings
    - network-overlay enable
    - network-id
- Hub parameters to add manually
  - Phase 2
    - keepalive
  - VPN interface
    - IP address
    - remote IP subnet
- Branch template main settings
  - As described in previous slides
    - ike-version 2
    - type **static**
    - net-device **enable**
    - mode-cfg enable
    - add-route disable
    - IP address assignment
  - Additional settings
    - network-overlay enable
    - network-id
    - **local ID**
    - Phase2 auto-negotiate enable
- Branch parameters to add manually
  - Phase 2
    - keepalive

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

When you use the SD-WAN overlay template or IPsec recommended templates for the hub and branches, FortiManager will prepare IPsec templates with the following parameters:

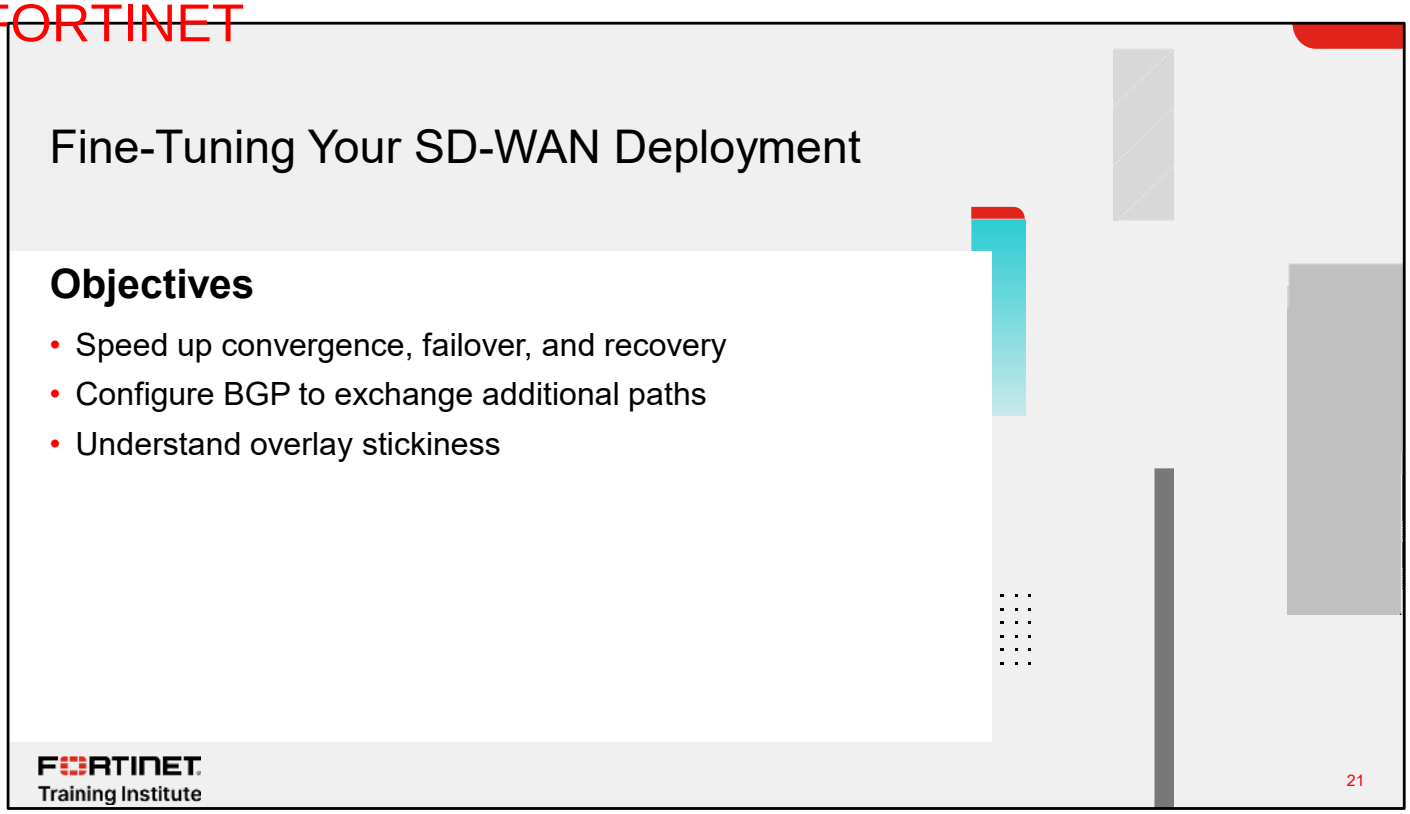
For Hubs:
- dynamic tunnel type for dial-up server
- net-device disable

For Branches:
- static tunnel type (remote end IP address is known)
- net-device enable

These parameters correspond to the recommended settings for IPsec overlay tunnel configuration, discussed earlier in this lesson.

The templates will also enable network-overlay and configure network-id for hubs, and, for branches, local-id and network-id. Optionally, you can edit the templates to add keepalives for both types of devices.
s
For the hub, you need to complete the template configuration with the VPN interface IP address and remote IP subnet settings. This can be done with the CLI template.

DO NOT REPRINT
© FORTINET



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

By understanding how to fine-tune your SD-WAN deployment, you will be able to speed up convergence, failover, and recovery. You will also learn how FortiGate devices can exchange all available paths, which is convenient for SD-WAN. Finally, you will be able to improve performances by forcing overlay stickiness.

DO NOT REPRINT
© FORTINET

Speeding Up Convergence, Failover, and Recovery

- BGP
  - Timers (seconds):
    - Keep alive (default: 60)
    - Hold (default: 180)
    - Advertisement (default: 30)
    - Connect (default: unset)
  - Link down failover (default: disable)
- IPsec
  - DPD settings:
    - Retry count (default: 3)
    - Retry interval (default: 20 seconds)

- Spoke1 configuration:

```
config router bgp
    set keepalive-timer 5
    set holdtime-timer 15
    config neighbor
        edit "10.201.1.254"
            set advertisement-interval 1
            set connect-timer 1
            set link-down-failover enable
        next
    end
end

config vpn ipsec phase1-interface
    edit "T_INET_0"
        set dpd-retrycount 2
        set dpd-retryinterval 10
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

With today's increasing service availability requirements, it's important to tune the configuration of your FortiGate devices so SD-WAN can respond faster to network changes. This slide shows the BGP and IPsec settings that you can adjust to speed up the convergence, failover, and recovery of SD-WAN overlays, what their default values are, and an example of how to configure them using the FortiGate CLI.

For BGP, you can configure FortiGate to bring down dead peerings faster by lowering the keepalive and hold timers. You can also speed up routing convergence by lowering the time that FortiGate waits between BGP updates by reducing the advertisement interval. Similarly, you can accelerate peering setup by reducing the connect timer value. Otherwise, after a failed connection attempt, FortiGate can wait a minute or more before making another connection attempt, which slows down routing convergence.

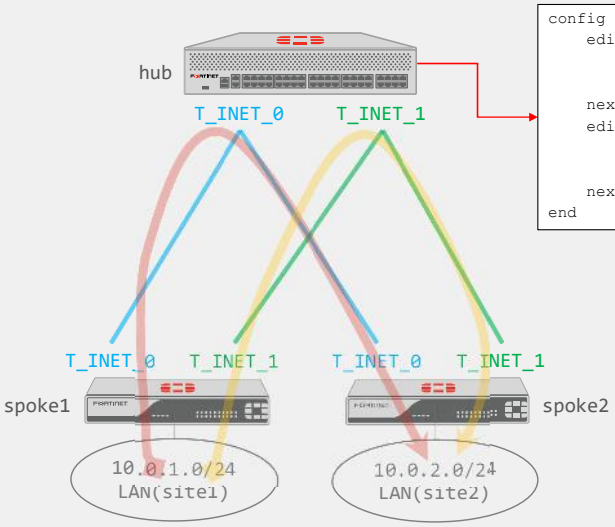
For BGP, it's also key to enable the link down failover feature. By default, FortiGate doesn't bring a peering down if the binding interface—the overlay in this case—goes down. Instead, FortiGate waits for the hold timer to expire. If you enable link down failover, then FortiGate brings down the peerings immediately after the interface they use comes down, which can then accelerate failover.

For IPsec, you can reduce the DPD retry count and retry interval settings to accelerate the detection of dead remote gateways. The DPD retry count defines the number of DPD retry attempts before considering the remote gateway dead. The DPD retry interval defines the time that FortiGate waits after sending a DPD message before considering the DPD attempt as failed (unanswered). If using DPD on-idle mode, the DPD retry interval also defines the idle time that FortiGate waits before considering the tunnel idle. With that said, with default settings, DPD takes 80 seconds to detect a dead gateway: a 20-second idle time plus 60 seconds from three unanswered DPD messages. For example, by setting the retry count and retry interval to 2 and 10, respectively, the time FortiGate takes to detect a dead gateway is reduced to 30 seconds: a 10-second idle time plus 20 seconds from two unanswered DPD messages.

DO NOT REPRINT
© FORTINET

Overlay Stickiness on the Hub

- Prefer spoke-to-spoke traffic to stay within same-ISP overlays
  - Better performance
  - Useful for ADVPN




```
config router policy
    edit 1
        set input-device "T_INET_0"
        set output-device "T_INET_0"
    next
    edit 2
        set input-device "T_INET_1"
        set output-device "T_INET_1"
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

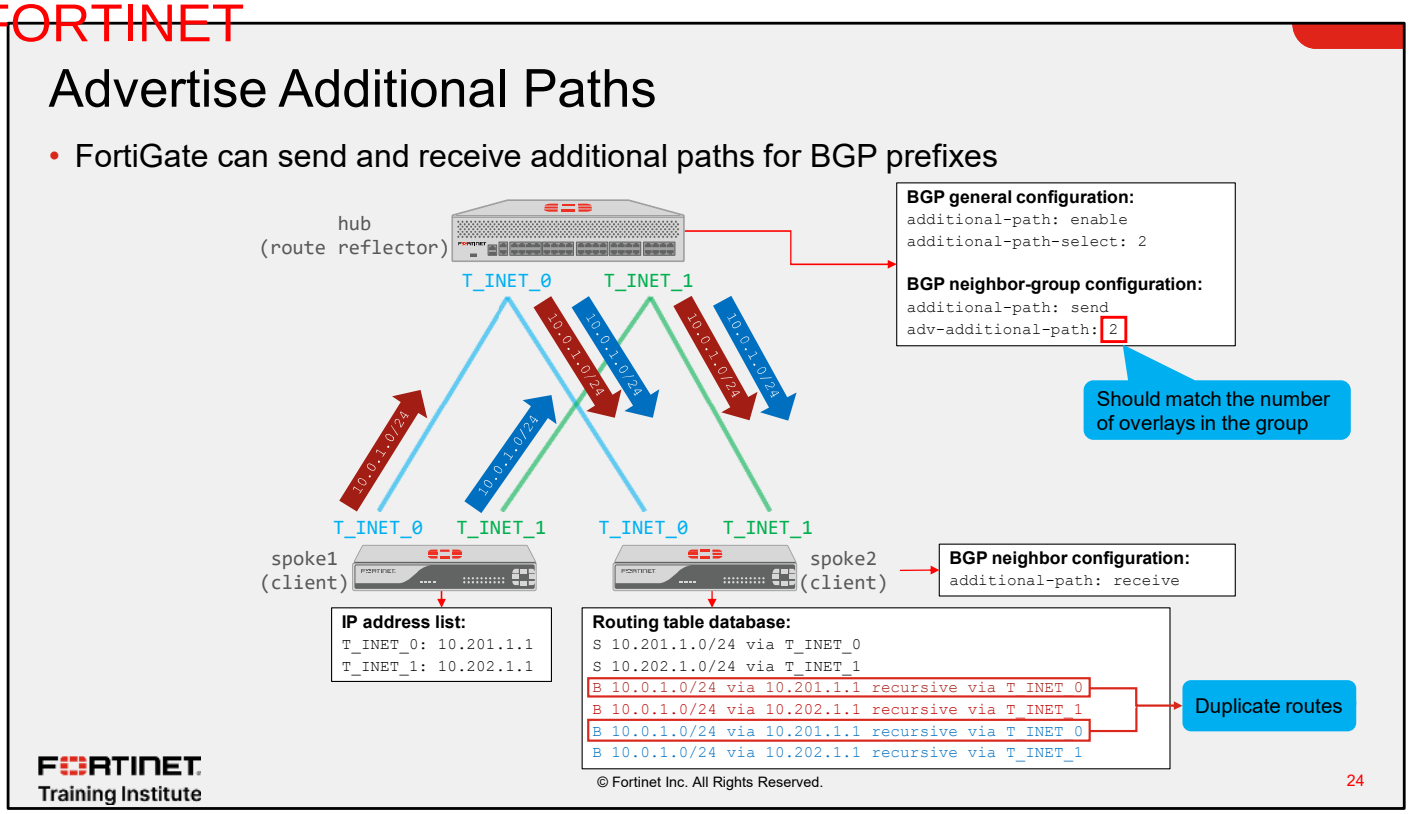
In the standard SD-WAN overlay design shown on this slide, spoke-to-spoke traffic is routed through the hub. When the hub receives the first packet of a spoke-to-spoke connection, it performs a route lookup to determine the best route and outgoing interface to forward the packet to. Based on the routing lookup results, the hub may resolve to use an overlay established over a different underlay from the one used by the incoming overlay. That is, based on our topology, if the hub receives the first packet on T_INET_0, it may decide to forward it out of T_INET_1.

Although the connection is established, the performance isn't the best because of the added latency introduced by the cross-ISP overlay path. To instruct FortiGate to prefer to keep spoke-to-spoke traffic within same-ISP overlays, you can configure the policy routes shown on this slide.

Note that the policy routes are a preference only. That is, they are used only if the FIB contains a route for the outgoing overlay. Otherwise, they are skipped and FortiGate forwards the traffic based on the best route in the FIB.

Overlay stickiness is very important for ADVPN because it helps to prevent the spokes from trying to negotiate shortcuts over unreachable underlays. You will learn more about ADVPN and overlay stickiness in this lesson.

DO NOT REPRINT
© FORTINET

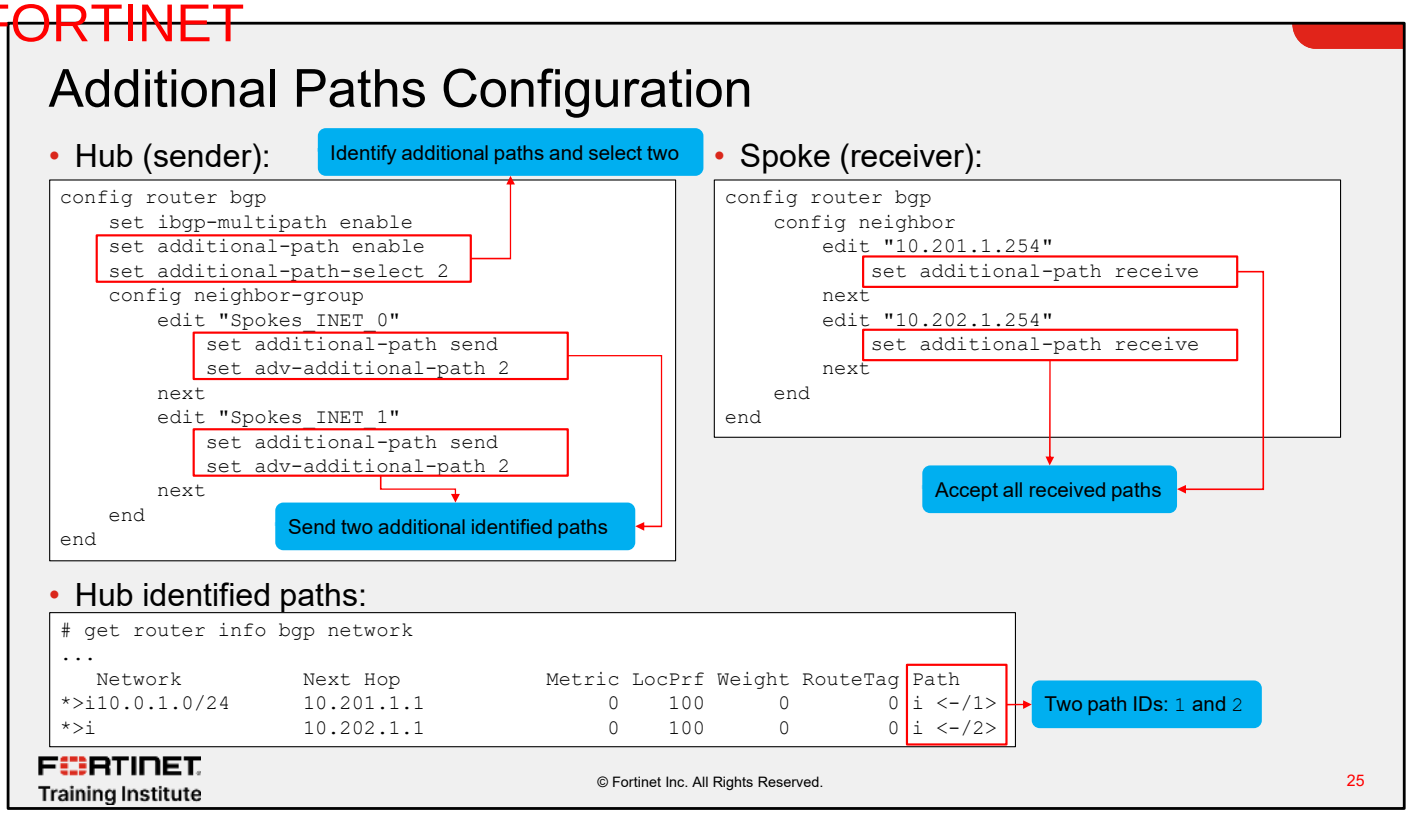


When you configure FortiGate as a BGP speaker, FortiGate advertises one path per prefix by default. In the SD-WAN overlay deployment shown on this slide, this means that the hub, acting as the route reflector, reflects only one path for each prefix despite learning two paths from the spoke client.

For SD-WAN, it is recommended to have sites learn all available paths to all possible destinations, otherwise SD-WAN rules and members may be skipped because of the absence of a route in the FIB. For this reason, it's a best practice to allow FortiGate to advertise additional paths so peers know all available paths. You should configure the hub to advertise one path per overlay.

The diagram on this slide shows the two paths that the hub learns for 10.0.1.0/24: one learned through T_INET_0 with next hop 10.201.1.1, and another through T_INET_1 with next hop 10.202.1.1. The hub is configured to send two additional paths, as well as to reflect prefixes learned from clients. In addition, the spoke is configured to receive any additional paths sent by the neighbor. The result is that spoke2 now receives both paths for 10.0.1.0/24 from the hub through the two overlays, for a total of four routes received. Because IBGP preserves the next-hop information, then the routing table database on spoke2 shows duplicate routes for the prefix.

DO NOT REPRINT
© FORTINET



This slide shows how to configure BGP additional paths using the FortiGate CLI based on the SD-WAN overlay topology used in this lesson.

For the hub to send the additional paths, it must first identify the unique paths learned from the BGP neighbor. The hub basically assigns an identifier to each unique path received for a prefix. You instruct FortiGate to identify the paths by enabling `additional-path` under `config router bgp`. You can view the identifiers and paths using the `get router info bgp network` command.

Next, the hub must select two or more identified paths that it can later advertise to neighbors. You control the number of selected identified paths with the `additional-path-select` setting.

After the hub identifies all paths, and selects a number of identified paths to advertise, you must indicate the neighbors you want the hub to advertise the additional paths to and how many. You configure this using the `additional-path` and `adv-additional-path` settings under `config neighbor-group`.

The example configuration shown on this slide for the hub instructs FortiGate to identify all paths received from neighbors, and selects two of them for advertising over BGP. In addition, the hub FortiGate advertises (or sends) two selected additional paths—that is, all selected additional paths—to peers on each neighbor group.

On the spoke, which acts as a receiver of the additional paths, you must set `additional-path` to `receive` under each neighbor entry, so the spoke FortiGate accepts all the additional paths sent by the hub.

DO NOT REPRINT
© FORTINET

Duplicate Routes Consolidation

- Routing table database shows all duplicate routes received

```
# get router info routing-table database
...
B     *> 10.0.1.0/24 [200/0] via 10.201.1.1 (recursive is directly connected, T_INET_0), 01:05:02
      *>             [200/0] via 10.202.1.1 (recursive is directly connected, T_INET_1), 01:05:02
      *>             [200/0] via 10.201.1.1 (recursive is directly connected, T_INET_0), 01:05:02
      *>             [200/0] via 10.202.1.1 (recursive is directly connected, T_INET_1), 01:05:02
```

- Duplicate routes are consolidated in the routing table

Number of duplicate routes available

```
# get router info routing-table all
...
B       10.0.1.0/24 [200/0] via 10.201.1.1 [2] (recursive is directly connected, T_INET_0), 01:03:45
                    [200/0] via 10.202.1.1 [2] (recursive is directly connected, T_INET_1), 01:03:45
```

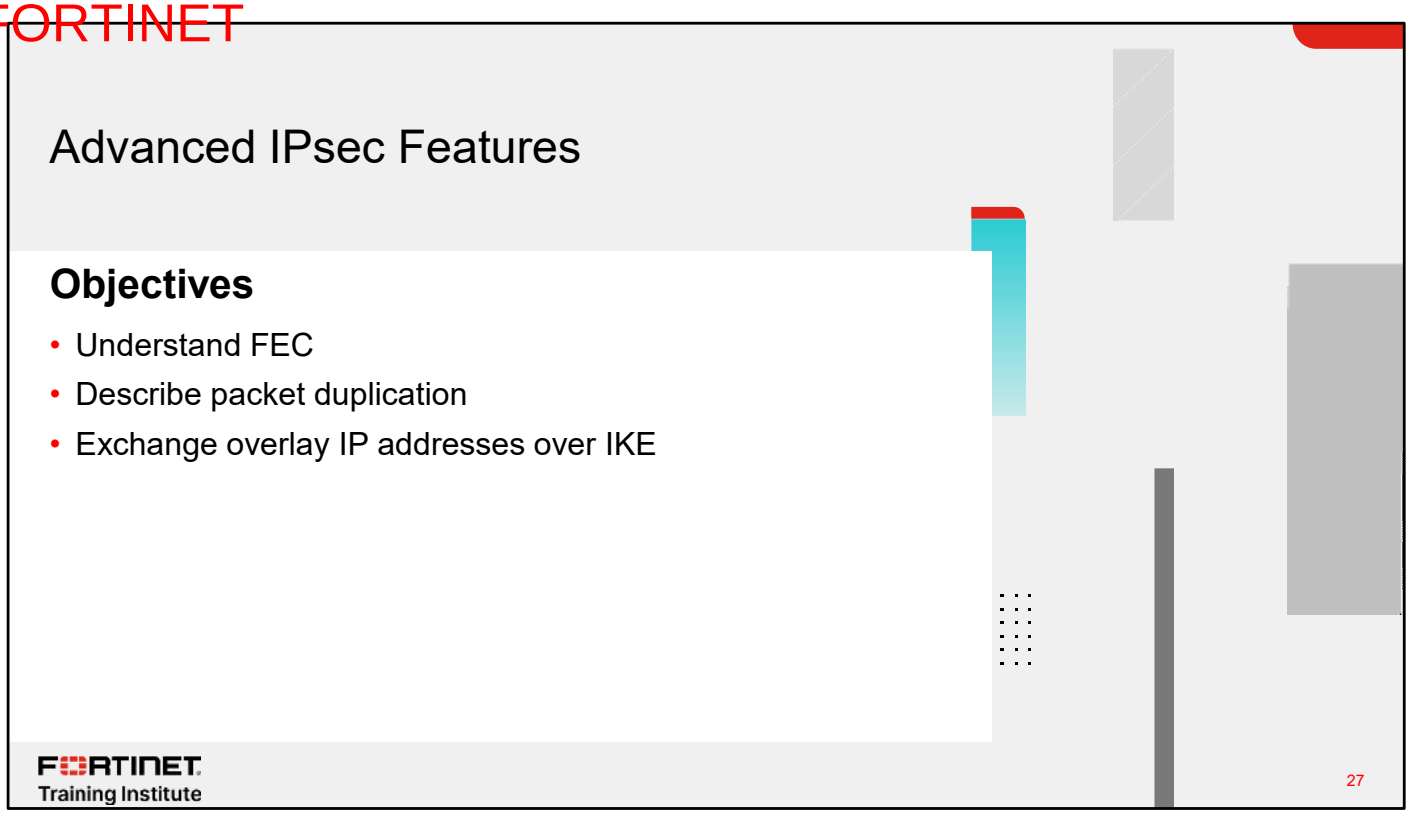
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

The routing table database show all duplicate routes received. On the routing table, for easier readability, FortiOS consolidates the duplicate routes and indicates the number of duplicate routes available. Note that the duplicate routes are not filtered, they are just consolidated for cosmetic purposes.

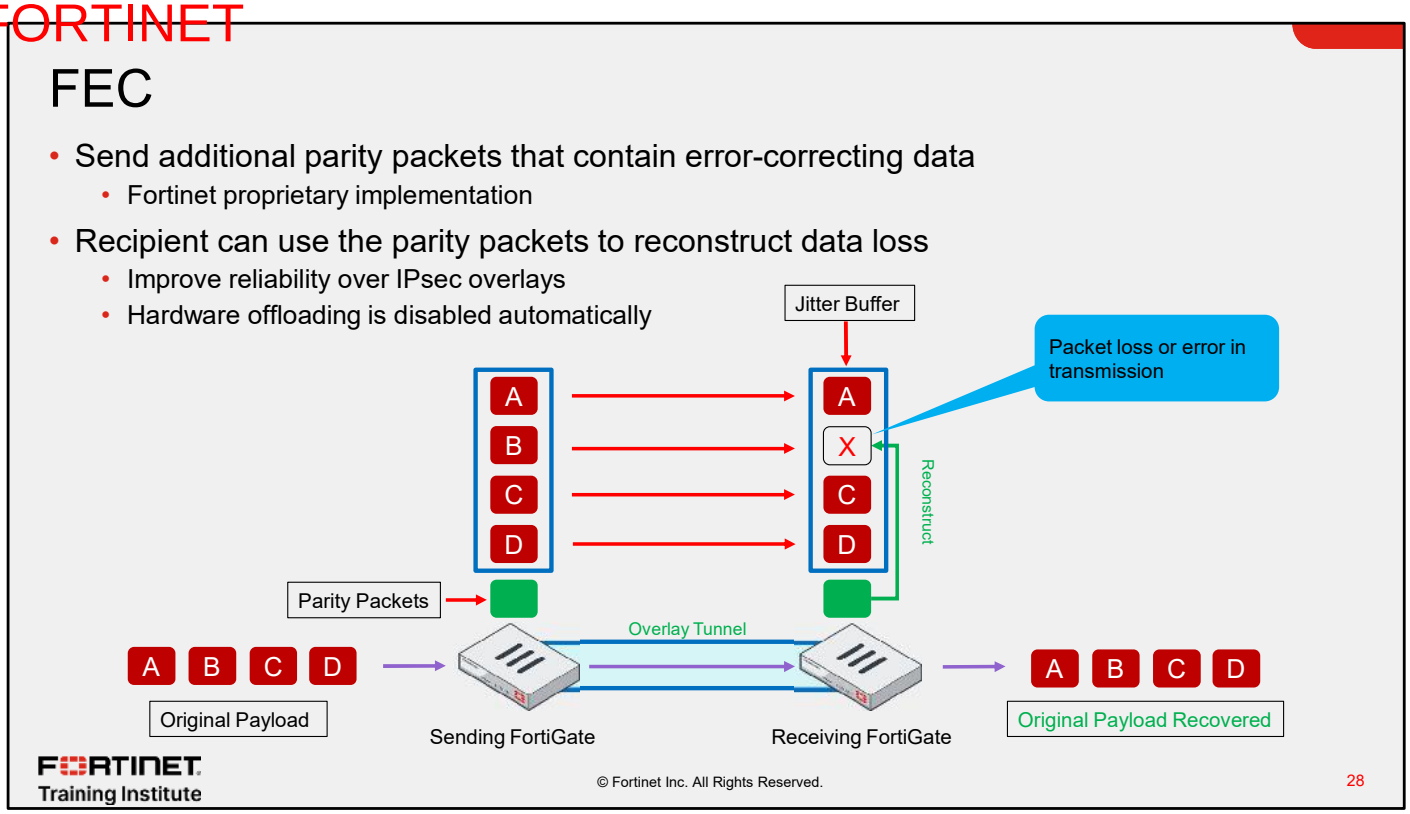
DO NOT REPRINT
© FORTINET



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

By demonstrating competence in advanced IPsec features, you will be able to incorporate useful IPsec features into SD-WAN to improve performance, reliability, and management.

DO NOT REPRINT
© FORTINET



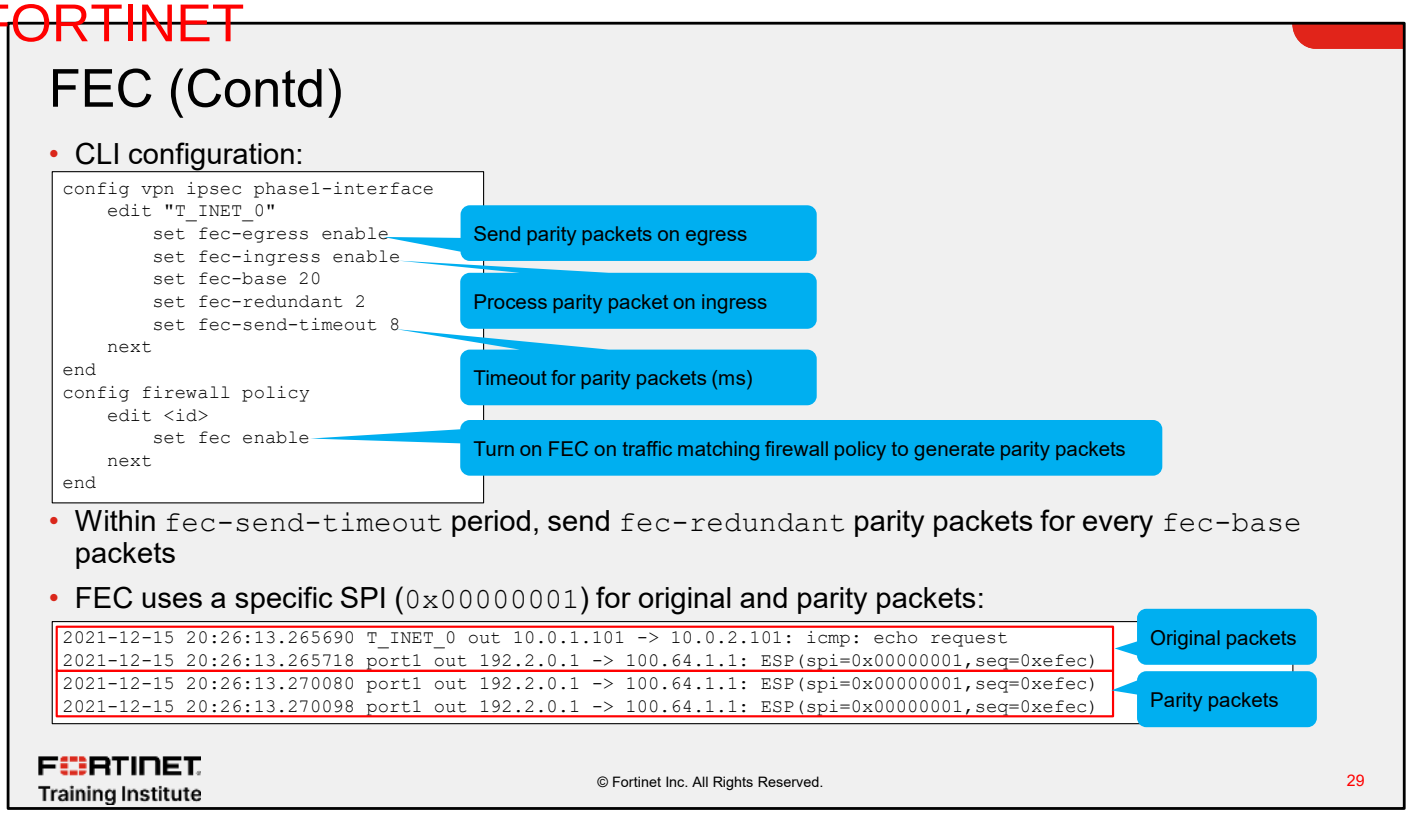
Forward Error Correction (FEC) is an error correction technique available for IPsec overlays that attempts to reduce the number of retransmissions over an IPsec tunnel. The FEC implementation on FortiGate doesn't work with third-party IPsec gateways.

When you enable FEC on an IPsec overlay, the sending FortiGate transmits additional packets—called parity packets—through the tunnel. The parity packets contain error-correcting data that is calculated based on the data contained in the original packets. The receiving FortiGate can then use the parity packets to reconstruct any lost packet, or any packet that arrived with errors.

Although FEC increases the bandwidth usage on an IPsec tunnel, it also improves the reliability of overlays by overcoming adverse WAN conditions, such as lossy or noisy underlays. FEC is useful for delivering a better user experience for business-critical applications, like voice and video services.

Note that FortiOS automatically disables hardware offload for traffic that is subject to FEC and therefore increases CPU usage.

DO NOT REPRINT
© FORTINET



This slide shows a basic configuration for FEC on an IPsec overlay using the FortiGate CLI.

FEC works by sending a `fec-redundant` number of parity packets for every `fec-base` number of packets sent within the `fec-send-timeout` period. However, if `fec-send-timeout` is reached, FortiGate sends the parity packets regardless of `fec-base`. The `fec-send-timeout` period starts after FortiGate receives the first packet destined for the tunnel. If you apply the configuration shown on this slide, then FortiGate sends two parity packets over T_INET_0 for every 20 packets sent through the tunnel over a period of eight milliseconds.

Also note that you must indicate the direction that you want FEC to run on. For example, if you want FortiGate to send parity packets on the egress direction, then you must enable `fec-egress`. Similarly, if you want FortiGate to process parity packets sent by the remote FortiGate, then you must enable `fec-ingress`. You must also enable `fec` on the firewall policy that matches the traffic that you want to generate parity packets for. Note that you don't have to manually disable hardware offloading. FortiOS will do that automatically for you on sessions that are subject to FEC.

Note that when you enable FEC, FortiGate uses the specific security payload index shown on this slide for original and parity packets.

This slide also shows a packet capture taken using the example configuration. The packet capture shows the original packet—an ICMP echo request to 10.0.2.101—egressing the T_INET_0 overlay, which is bound to port1. The packet capture also shows three more ESP packets egressing port1. Although not revealed in the output, the first ESP packet carries the original ICMP echo request packet. The following two ESP packets are parity packets of the original ICMP echo request packet.

DO NOT REPRINT
© FORTINET

Impact of FEC on Overlay Quality

- Before enabling FEC:

```
root@branch1-client-cli:~# ping 10.0.2.101 -i 0.2 -c 100
PING 10.0.2.101 (10.0.2.101) 56(84) bytes of data.
...
--- 10.0.2.101 ping statistics ---
100 packets transmitted, 73 received, 27% packet loss, time 964ms
rtt min/avg/max/mdev = 1.002/2.619/5.433/0.895 ms
```

- After enabling FEC (`fec-egress: enable`, `fec-ingress: enable`, `fec-redundant: 1`, `fec-base: 20`, `fec-send-timeout: 8`):

```
--- 10.0.2.101 ping statistics ---
100 packets transmitted, 88 received, 12% packet loss, time 914ms
rtt min/avg/max/mdev = 1.749/4.497/13.317/2.851 ms
```

- After increasing `fec-redundant` to `3`:

```
--- 10.0.2.101 ping statistics ---
100 packets transmitted, 99 received, 1% packet loss, time 891ms
rtt min/avg/max/mdev = 2.133/4.317/13.098/2.271 ms
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

This slide shows the positive impact that FEC has on the quality of IPsec overlays. The test consists of generating 100 ICMP echo request packets across a noisy overlay, and then identifying the amount of packet loss reported by the ping tool before and after enabling FEC.

The first output shows the ping statistics before enabling FEC. The reported packet loss is 27%.

The second output shows the ping statistics after the administrator enables FEC on both local and remote FortiGate devices using the settings shown on this slide. The packet loss is reduced to 12%.

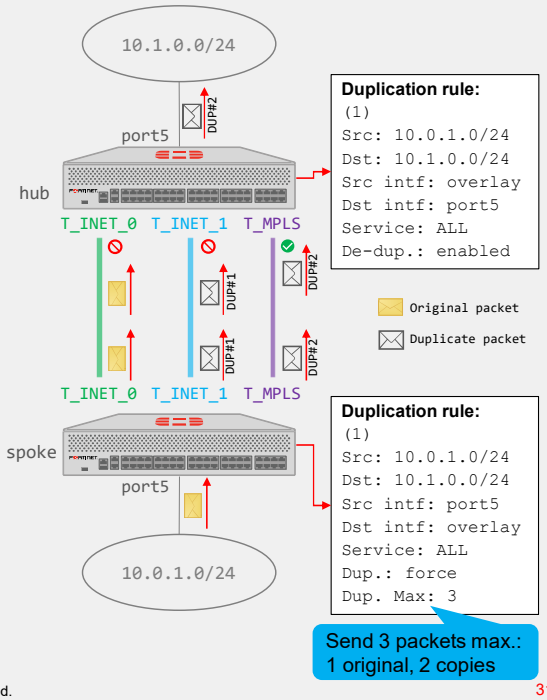
The administrator then increases the `fec-redundant` setting from `1` to `3` on both ends of the tunnels. As a result, the packet loss is reduced to 1%.

Therefore, you can increase the quality of noisy overlays by enabling FEC and increasing the `fec-redundant` setting. Just keep in mind that the higher the value of `fec-redundant`, the more bandwidth is used.

DO NOT REPRINT
© FORTINET

Packet Duplication

- Send duplicate packets through additional links
  - Can be underlay or overlay
  - Supports hardware offloading
- Requirements
  - Best route must match an SD-WAN member
  - Member must have route to destination
- Duplicate packets are exact copies of the original packet
  - Useful for data loss protection and OOB inspection
- Enable de-duplication to drop additional copies
  - First arriving packet is accepted, others dropped



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

FEC provides an efficient way to reduce data loss on IPsec tunnels. However, FEC doesn't leverage multiple IPsec tunnels. In addition, FEC doesn't support hardware offload or links that are not IPsec.

Packet duplication is an SD-WAN feature that enables you to send duplicate packets through up to three additional members of any kind, provided the best route to the destination is an SD-WAN member, and the links used for duplication have a route to the destination. Unlike the parity packets in FEC, which are not exact copies of the original packets, the duplicate packets are verbatim copies of the original packet. This way, the duplicate packets can be used for data loss protection, and for out-of-band inspection or packet capture.

You can also enable packet de-duplication on the receiving FortiGate. When you do this, the receiving FortiGate accepts only the first copy of the packet received, and drops the additional copies. The goal is to save resources at the receiving end by instructing FortiGate to forward one copy only, instead of forwarding all copies and letting the next hop discard the additional packets.

The example on this slide shows two FortiGate devices connected through three IPsec overlays, which are members of the overlay zone. On the spoke FortiGate, the administrator configured a duplication rule that instructs FortiGate to always duplicate (forced duplication) any packet sourced from 10.0.1.0/24 and destined to 10.1.0.0/24, to up to three links. On the hub FortiGate, the administrator configured a duplication rule to de-duplicate the same traffic.

The result is that when the spoke FortiGate receives a packet that matches the duplication rule, it forwards the original packet to the preferred member in the zone, plus two additional copies through two additional links, for a total of three packets. On the remote side, the hub FortiGate accepts only the packet received through T_MPLS because it was the first packet to arrive, and drops the additional copies received through T_INET_0 and T_INET_1. The result is that one packet only is forwarded to the local interface (port5).

DO NOT REPRINT
© FORTINET

Forced Packet Duplication

- Spoke CLI configuration:

```
config system sdwan
    set duplication-max-num 3
    config duplication
        edit 1
            set srcaddr "10.0.1.0/24"
            set dstaddr "10.1.0.0/24"
            set srcintf "port5"
            set dstintf "overlay"
            set service "ALL"
            set packet-duplication force
        next
    end
end
```

Always duplicate; an alternative is `on-demand`

- Hub CLI configuration:

```
config system sdwan
    config duplication
        edit 1
            set srcaddr "10.0.1.0/24"
            set dstaddr "10.1.0.0/24"
            set srcintf "overlay"
            set dstintf "port5"
            set service "ALL"
            set packet-de-duplication enable
        next
    end
end
```

- Spoke packet duplication:

```
port5 in 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_INET_1 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_MPLS out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_INET_0 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_MPLS in 10.1.0.7 -> 10.0.1.101: icmp: echo reply
port5 out 10.1.0.7 -> 10.0.1.101: icmp: echo reply
```

3 copies (original plus 2 duplicates)

- Hub packet de-duplication:

```
T_MPLS in 10.0.1.101 -> 10.1.0.7: icmp: echo request
port5 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_INET_1 in 10.0.1.101 -> 10.1.0.7: icmp: echo reques
T_INET_0 in 10.0.1.101 -> 10.1.0.7: icmp: echo reque
port5 in 10.1.0.7 -> 10.0.1.101: icmp: echo reply
T_MPLS out 10.1.0.7 -> 10.0.1.101: icmp: echo repl
```

First to arrive, only to be accepted

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

This slide shows the configuration used for the packet duplication example shown on the previous slide, which uses forced packet duplication.

The duplication rules on both spoke and hub indicate the 5-tuple of the traffic to duplicate and de-duplicate, respectively. On the spoke, `duplication-max-num` is set to `3`. This instructs FortiGate to forward up to three copies of the packet—the original packet plus two duplicates. Moreover, each copy is sent through a different member. On the spoke, `packet-duplication` is set to `force`, which instructs FortiGate to always duplicate packets. An alternative is to set `packet-duplication` to `on-demand`, which instructs FortiGate to duplicate packets based on the SLA status of the outgoing interface. You will learn more about on-demand duplication in this lesson.

On the hub side, `packet-de-duplication` is enabled to instruct FortiGate to accept one copy of the packet only—the first to arrive—and drop any additional copies.

This slide also shows the sniffer output for the corresponding duplicated and de-duplicated packets on the spoke and hub, respectively. On the spoke, the original ICMP echo request packet is received at port5, and then forwarded to T_INET_1. The spoke also forwards two more copies, one to T_MPLS, and another to T_INET_0. On the hub side, all three packets are received, but only the first packet received—the packet received at T_MPLS—is accepted and forwarded. The hub FortiGate then creates a session with T_MPLS as the incoming interface, which is why the ICMP echo reply packet is forwarded to the same interface.

DO NOT REPRINT
© FORTINET

On-Demand Packet Duplication

- Consider link quality before duplicating packets:
  - Don't duplicate if:
    - Outgoing interface meets at least one of the SLA targets
  - Duplicate packets if:
    - Outgoing interface doesn't meet any SLA targets, or no SLA targets are configured
- Example (outgoing interface is T_INET_0):
  - T_INET_0 meets at least one SLA target (no duplication):

```
# diagnose sys sdwan health-check status
Health Check(VPN):
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(1.500), jitter(0.291) sla_map=0x3
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.415), jitter(0.281) sla_map=0x3
Seq(5 T_MPLS): state(alive), packet-loss(0.000%) latency(1.116), jitter(0.225) sla_map=0x3
```

  - T_INET_0 doesn't meet any SLA targets (perform duplication):

No SLA targets met

```
# diagnose sys sdwan health-check status
Health Check(VPN):
Seq(3 T_INET_0): state(alive), packet-loss(1.000%) latency(171.729), jitter(0.340) sla_map=0x0
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.506), jitter(0.419) sla_map=0x3
Seq(5 T_MPLS): state(alive), packet-loss(0.000%) latency(1.061), jitter(0.280) sla_map=0x3
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

Forced packet duplication instructs FortiGate to always duplicate packets. But what if you want FortiGate to perform duplication only when the quality of the outgoing interface is poor? This way, you can save network resources by avoiding duplicate traffic when the outgoing interface is performing well, and the original packet is likely to be successfully delivered.

When you configure on-demand packet duplication, FortiGate considers the quality of the outgoing interface selected to route the original packet before creating any duplicates. The outgoing interface is selected based on the SD-WAN rule lookup for the packet. That is, the packet doesn't have to match a user-defined SD-WAN rule for duplication to be triggered. This is also true for forced packet duplication.

FortiGate determines the quality of the outgoing interface based on the status of the SLA targets configured. If the interface meets at least one of its configured SLA targets, then FortiGate considers the quality to be acceptable and doesn't create duplicates. If the interface doesn't meet any of its configured SLA targets, or if the interface doesn't have any SLA targets configured, then FortiGate performs duplication.

This slide shows the health of three overlays taken at different times. Assuming T_INET_0 is the selected outgoing interface for the original packet, then FortiGate doesn't create duplicates when the performance of T_INET_0 is the one shown in the first output. However, if the performance is the one shown in the second output, then FortiGate creates duplicates because T_INET_0 doesn't meet any of the SLA targets.

DO NOT REPRINT
© FORTINET

Restrict Packet Duplication to SD-WAN Rules

- Duplicate traffic that matches one or more SD-WAN rules
  - More granular control and flexibility
  - Reduced configuration
- CLI configuration:

```
config system sdwan
    set duplication-max-num 3
    config duplication
        edit 1
            set service-id 1
            set packet-duplication force
        next
    end
end
```

- SD-WAN rule ID 1 status:

```
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(1: 1->65535), Mode(manual)
  Members(3):
    1: Seq_num(3 T_INET_0), alive, selected
    2: Seq_num(4 T_INET_1), alive, selected
    3: Seq_num(5 T_MPLS), alive, selected
  Src address(1):
        0.0.0.0-255.255.255.255

  Dst address(1):
        10.1.0.7-10.1.0.7
```

- ICMP traffic to 10.1.0.7 is duplicated:

```
port5 in 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_INET_1 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_MPLS out 10.0.1.101 -> 10.1.0.7: icmp: echo request
T_INET_0 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
```

- Other traffic, like SSH, is not duplicated:

```
port5 in 10.0.1.101.55030 -> 10.1.0.7.22: syn 122428627
T_INET_0 out 10.0.1.101.55030 -> 10.1.0.7.22: syn 122428627
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

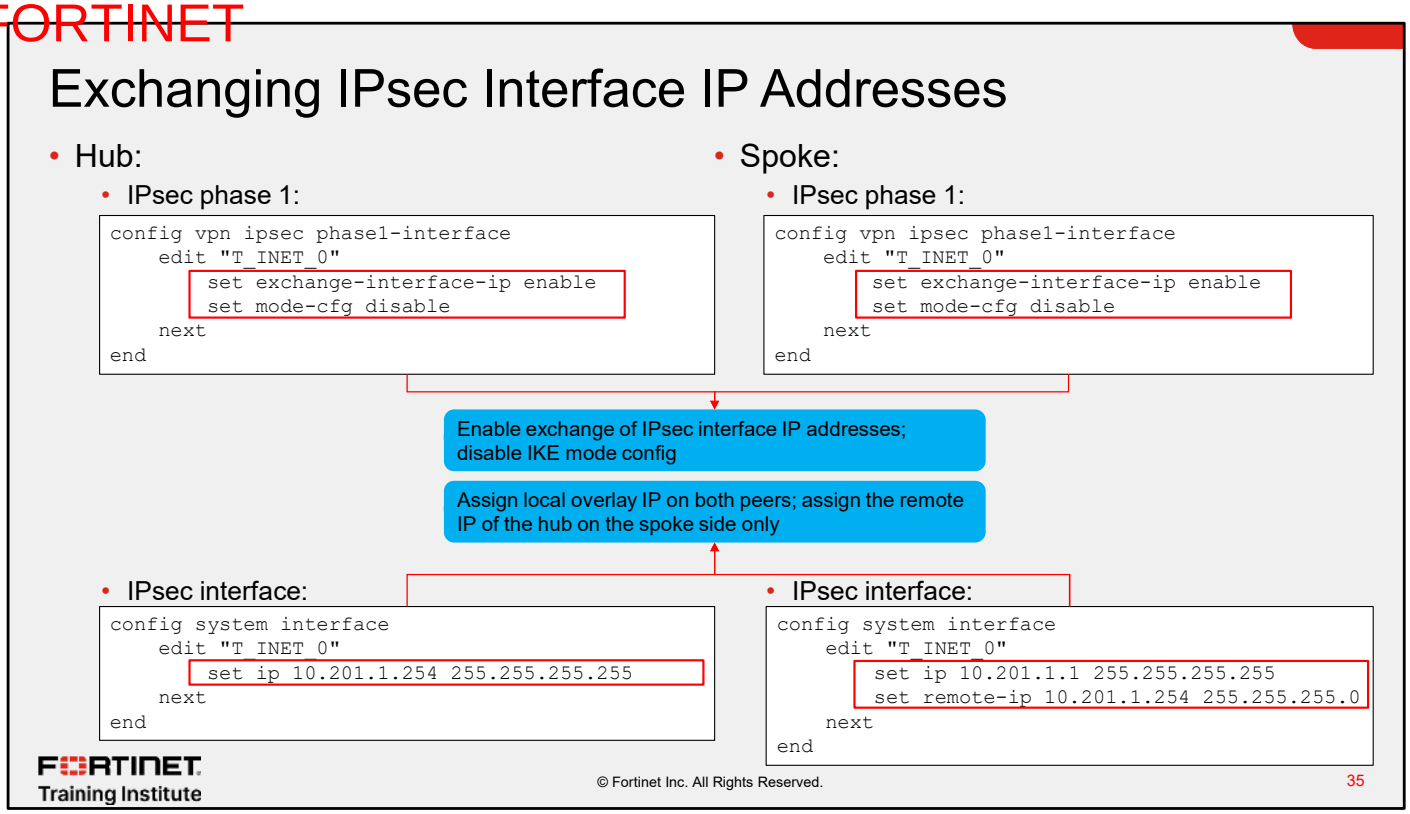
34

You can configure packet duplication rules to duplicate packets that match one or more SD-WAN rules. This enables you to have more granular control and flexibility over duplicated traffic, as well as reduce the configuration required. For example, you can leverage the application steering feature of SD-WAN rules to duplicate traffic based on the detected application, rather than using the fixed 5-tuple configuration on a duplication rule to identify the traffic. Similarly, you can leverage ISDB and route tags to identify the destination of the traffic to duplicate.

This slide shows an example of a forced packet duplication rule that matches the SD-WAN rule ID 1. Note that after you configure the SD-WAN rule ID using the service setting, the options to define the 5-tuple of the traffic on the duplicate rule are hidden. This is because the matching criteria is now defined by the SD-WAN rule.

This slide also shows the SD-WAN rule status. The output indicates that the SD-WAN rule matches ICMP traffic from any source to 10.1.0.7. The result is that FortiGate duplicates only the traffic that matches the SD-WAN rule ID 1. Other traffic, such as SSH, is not duplicated.

DO NOT REPRINT
© FORTINET



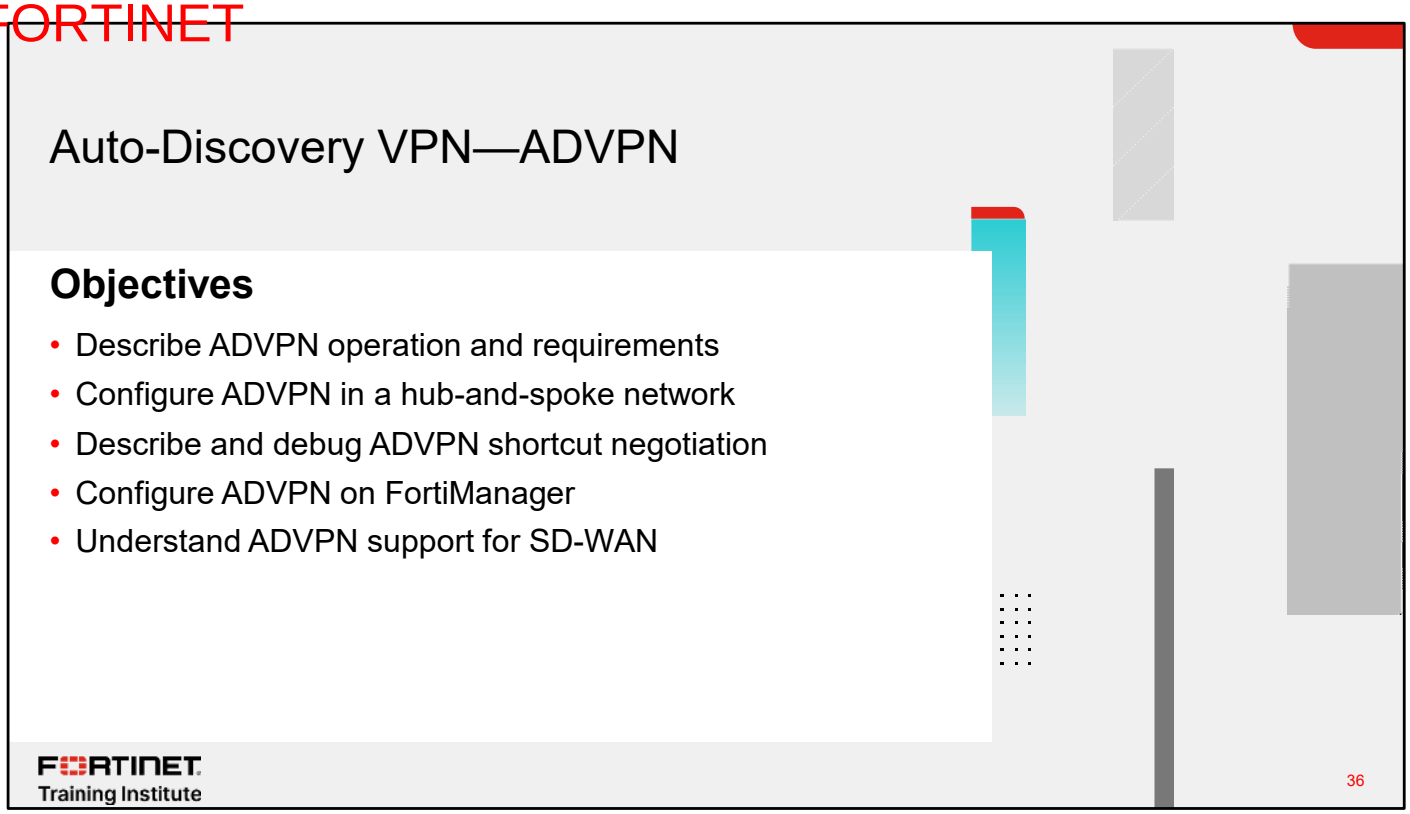
The example IPsec configuration used for the SD-WAN overlay topology shown in this lesson makes use of IKE mode config to assign the IP addresses of the spoke overlays. Although its configuration is simple, the IKE mode config doesn't allow you to define fixed IP addresses on the spoke side. That is, the address assigned to an overlay depends on the available addresses during tunnel negotiation.

An alternative to IKE mode config is the IPsec exchange interface feature. Instead of using IKE mode to assign addresses automatically, you manually assign the address on the spoke and hub sides, and then have IKE exchange them during tunnel negotiation so each gateway knows the remote overlay IP. The result is that you can assign a fixed address to a spoke overlay, which can then be used for monitoring and identification purposes.

This slide shows the changes that are required on the hub and spoke to exchange the IPsec interface IP addresses using IKE. The configuration changes are the same on each end except that you assign different local IP addresses on each side, and on the spoke FortiGate, you must also define the hub IP address. Also, when defining the hub IP address, use as the netmask the overlay subnet netmask (/24 in the example) instead of /32. Using the overlay subnet netmask is useful for next hop resolution when reflecting routes using IBGP.

Note that FortiGate uses a Fortinet proprietary attribute to exchange the IP addresses. This means that, unlike IKE mode config which is a standard IPsec feature, you can exchange IP addresses using IKE only when both gateways are FortiGate devices.

DO NOT REPRINT
© FORTINET



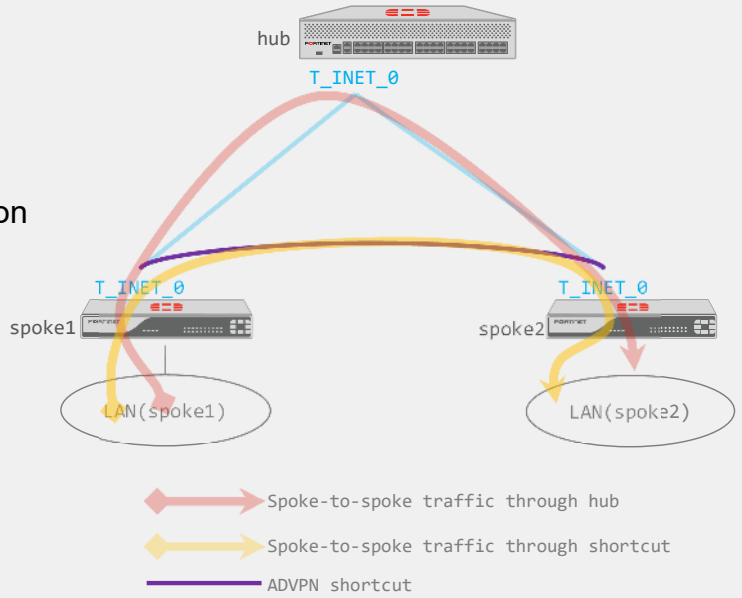
After completing this section, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in ADVPN, you will be able to configure and troubleshoot ADVPN. You will also understand SD-WAN support for ADVPN.

DO NOT REPRINT
© FORTINET

ADVPN

- Enables direct spoke-to-spoke communication
  - Fortinet proprietary
  - Based on IKE and IPsec
  - Benefits of full-mesh in a hub-and-spoke network
- Hub facilitates spokes shortcut negotiation
  - On-demand IPsec tunnels
  - Shortcuts carry spoke-to-spoke traffic
- Supported by SD-WAN
  - Automatically steers traffic through shortcuts
  - Shortcuts are monitored



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

37

Consider the simple hub-and-spoke topology shown on this slide. If a device behind spoke1 wants to communicate with a device behind spoke2, the device must do so through the hub. While a hub-and-spoke topology reduces cost and configuration overhead, the fact that a spoke must communicate with other spokes through the hub increases the delay of the communication, and therefore impacts the user experience, especially if the hub and spokes are in geographically distant locations.

An alternative is to use a full-mesh topology, which enables direct spoke-to-spoke communication at the expense of higher costs and increased administrative overhead. However, this is not always practical, or even feasible, especially in large networks.

ADVPN—auto discovery VPN— is a Fortinet-proprietary solution based on IKE and IPsec, that addresses the need for direct spoke-to-spoke communication in hub-and-spoke topologies by enabling the spokes to automatically negotiate on-demand IPsec tunnels—called shortcuts—between them without you having to make topology changes or many configuration changes. After a shortcut is established and routing has converged, spoke-to-spoke traffic no longer needs to flow through the hub.

SD-WAN supports ADVPN shortcuts. For this, SD-WAN automatically steers the traffic through shortcuts and monitors their health and performance. You add the parent tunnel as member, and after the shortcut is negotiated, SD-WAN automatically starts steering the traffic through the shortcut.

The topology on this slide also shows a shortcut established between spoke1 and spoke2. The result is that FortiGate routes the spoke-to-spoke traffic through the shortcuts instead of through the hub.

DO NOT REPRINT
© FORTINET

ADVPN (Contd)

- Supports multiple hub-and-spoke architectures
- Supports NAT for on-demand tunnels
- Requires use of dynamic routing
  - iBGP recommended
  - OSPF, and RIPv2/RIPng are supported
  - PIM/multicast is supported
- Supports both IPv4 and IPv6

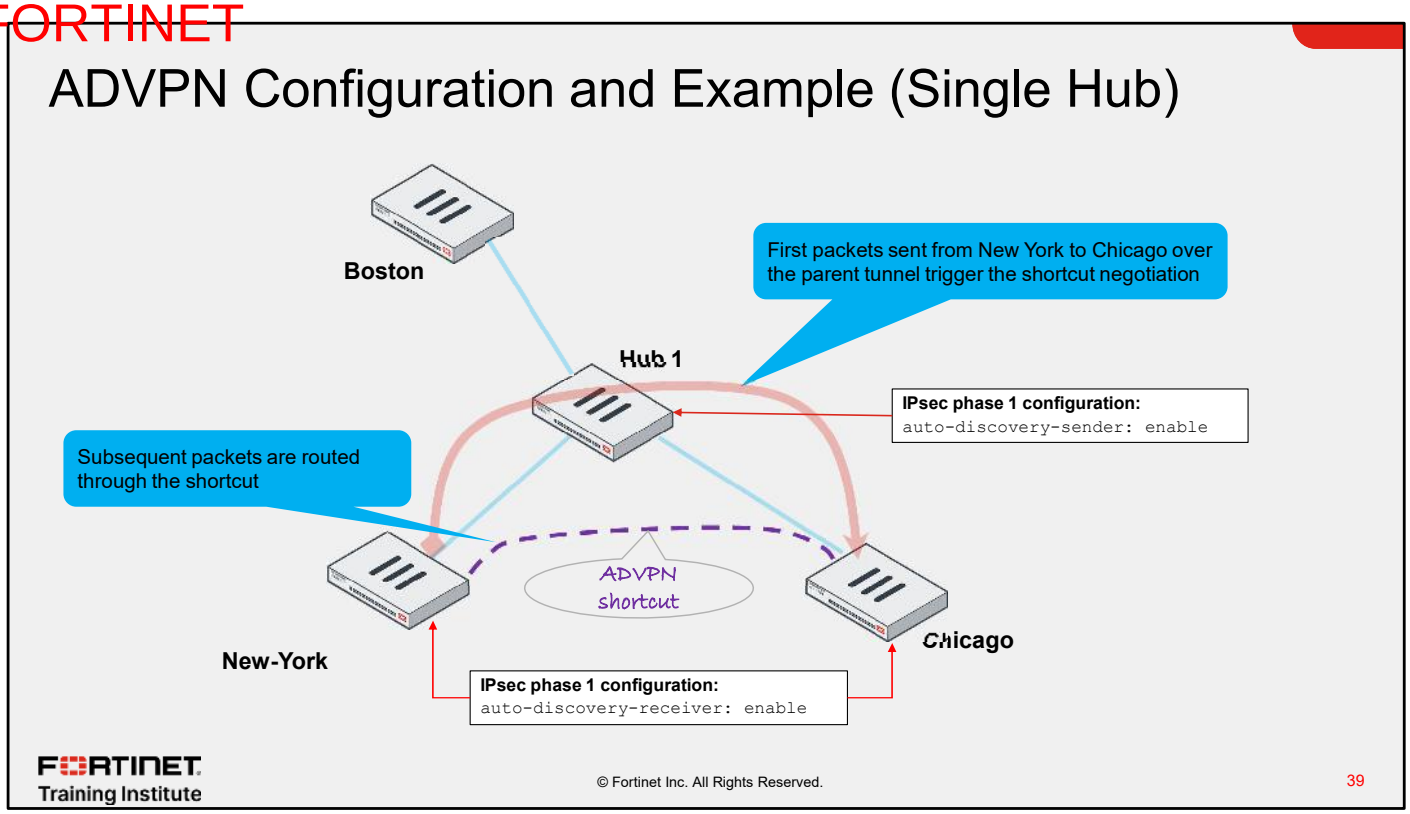
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

38

- ADVPN supports single or multiple hub architectures
- NAT is supported for the on-demand tunnels
- ADVPN requires the use of a routing protocol
  - Currently, it supports BGP, OSPF, RIPv2/RIPng, and static routing
  - It also supports PIM and multicast
- Both IPv4 and IPv6 are supported
- Fortinet recommends to use IBGP as routing protocol for ADVPN topologies

DO NOT REPRINT
© FORTINET



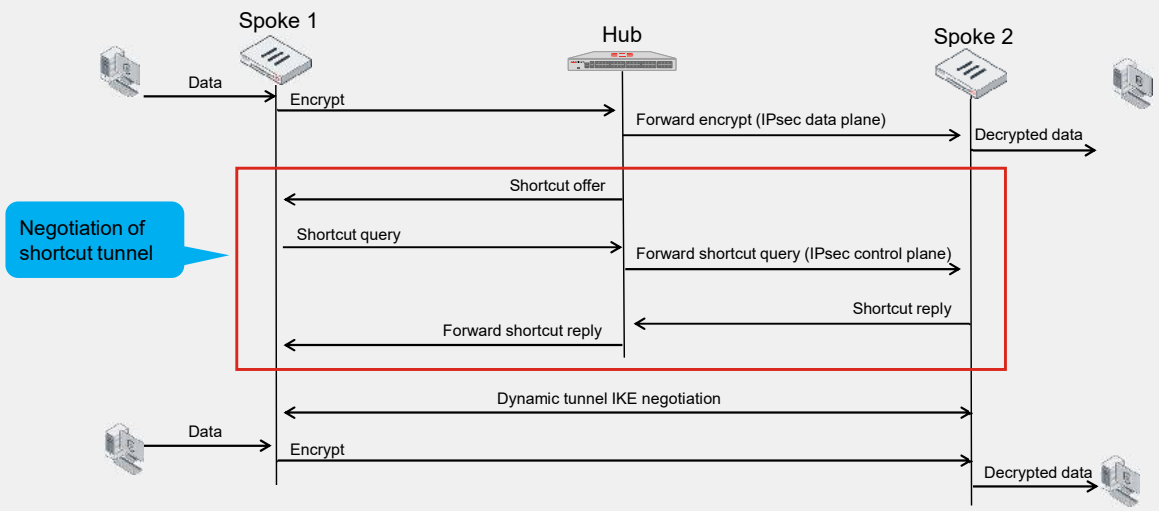
This slide shows an example of how ADVPN works in a single hub topology. Consider the basic IPsec hub-and-spoke topology shown on this slide. Overlays that connect a spoke to a hub are called parent tunnels. Direct tunnels negotiated over parent tunnels are called shortcuts. The shortcuts between spokes are created on demand based on the IPsec phase 1 auto-discovery settings configured on each peer:

- `auto-discovery-receiver` should be enabled on the spoke overlay to inform the hubs that the spoke can negotiate shortcuts.
- `auto-discovery-sender` should be enabled on the hub overlay that connects the spoke to allow for shortcut negotiation between spokes.

Now, say that a user in Boston sends traffic to Chicago. Initially, the shortcut between Boston and Chicago has not been negotiated. So, the first packets are routed through Hub 1. When Hub 1 receives those packets, it knows that ADVPN is enabled both spokes. So, Hub 1 sends an IKE message to Boston informing that it can try to negotiate a direct connection to Chicago. On receipt of this IKE message, New-York creates a FortiOS-specific IKE information message that contains its public IP address, its local subnet, the desired destination subnet (Chicago's subnet), and an auto-generated PSK (alternatively, it can also use digital certificate authentication). This IKE message is sent to Chicago through Hub 1. When Chicago receives the IKE message from New-York, it stores the PSK and replies with another IKE information message that contains Chicago's public IP address. After the reply arrives to New-York, the shortcut is negotiated between both peers. The negotiation succeeds because Chicago is expecting a connection attempt from New-York's public IP address. You will explore this in greater detail in this lesson.

DO NOT REPRINT
© FORTINET

ADVPN Message Exchange



FORTINET Training Institute

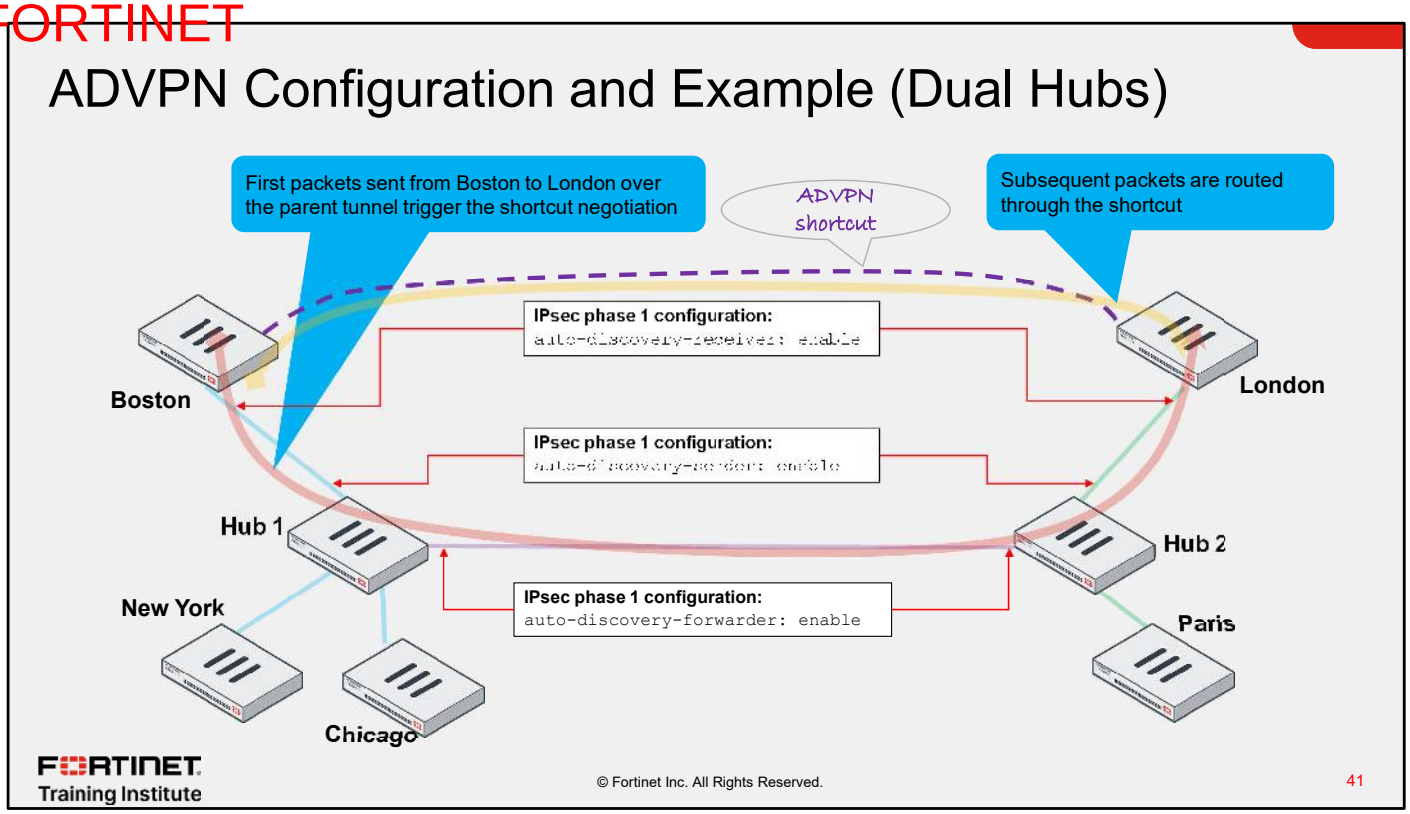
© Fortinet Inc. All Rights Reserved.

40

Now, you will examine the IKE messages that are exchanged when an on-demand tunnel is being negotiated:

1. The client behind Spoke 1 generates traffic for devices located on the Spoke 2 network.
2. Spoke 1 receives the packet, encrypts it, and sends it to the hub.
3. The hub receives the packet from Spoke 1 and forwards it to Spoke 2.
4. Spoke 2 receives the packet, decrypts it, and forwards it to the destination device.
5. The hub knows that a more direct tunnel option might be available from Spoke 1 to Spoke 2. The hub sends a shortcut offer message to Spoke 1.
6. Spoke 1 acknowledges the shortcut offer by sending a shortcut query to the hub.
7. The hub forwards the shortcut query message to Spoke 2.
8. Spoke 2 acknowledges the shortcut query and sends a shortcut reply to the hub.
9. The hub forwards the shortcut reply to Spoke 1.
10. Spoke 1 and Spoke 2 initiate the tunnel IKE negotiation.

DO NOT REPRINT
© FORTINET



Now, consider the dual-hub IPsec hub-and-spoke topology shown on this slide. In such a scenario, you must enable `auto-discovery-forwarder` on both hubs to allow ADVPN information exchange between the two hubs and permit ADVPN shortcut establishment.

- `auto-discovery-receiver` should be enabled on the spoke overlay to inform the hubs that the spoke can negotiate shortcuts.
- `auto-discovery-sender` should be enabled on the hub overlay that connects the spoke to allow for shortcut negotiation between spokes.
- `auto-discovery-forwarder` should be enabled on the hub overlay that connects to the other hub to forward ADVPN sender and receiver information between hubs.

Now, say that a user in Boston sends traffic to London. Initially, the shortcut between Boston and London has not been negotiated. So, the first packets from Boston to London are routed through Hub 1 and Hub 2. When Hub 1 receives those packets, it knows that ADVPN is enabled in all the VPNs all the way to London. So, Hub 1 sends an IKE message to Boston informing that it can try to negotiate a direct connection to London. On receipt of this IKE message, Boston creates a FortiOS-specific IKE information message that contains its public IP address, its local subnet, the desired destination subnet (London's subnet), and an auto-generated PSK (alternatively, it can also use digital certificate authentication). This IKE message is sent to London through Hub 1 and Hub 2. When London receives the IKE message from Boston, it stores the PSK and replies with another IKE information message that contains London's public IP address. After the reply arrives in Boston, the shortcut is negotiated between both peers. The negotiation succeeds because London is expecting a connection attempt from Boston's public IP address. You will explore this in greater detail in this lesson.

DO NOT REPRINT
© FORTINET

Basic SD-WAN Overlay Design—ADVPN Configuration

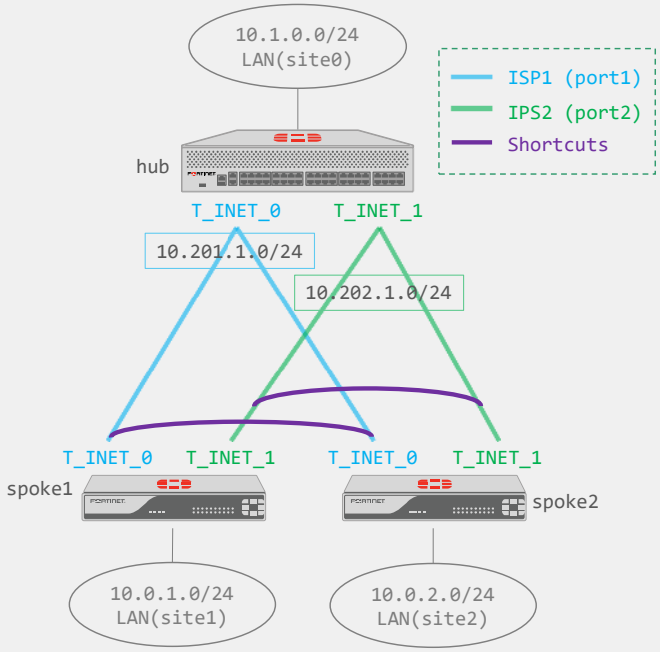
- Hub:
  - IPsec phase 1:

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        set net-device disable
        set auto-discovery-sender enable
    next
end
```

- Spoke:
  - IPsec phase 1 and interface:

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        set net-device enable
        set auto-discovery-receiver enable
    next
end
config system interface
    edit "T_INET_0"
        set allowaccess ping
    next
end
```

Required for shortcut monitoring



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

The topology on this slide shows the ADVPN configuration required on our basic SD-WAN deployment.

On the hub, you must disable `net-device` and enable `auto-discover-sender` on the phase 1 configuration of each overlay.

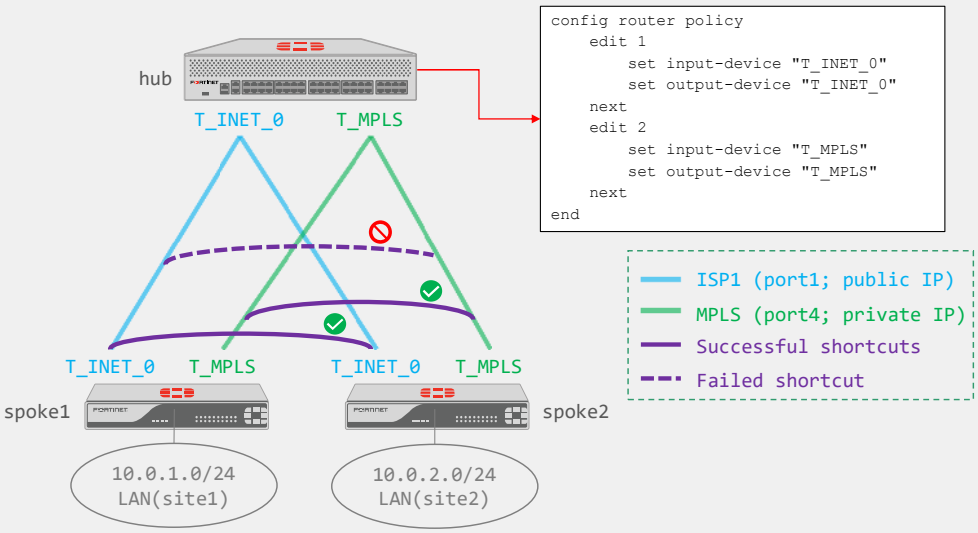
On spokes, you must enable both `net-device` and `auto-discovery-receiver` on the phase 1 configuration of each overlay. You must also allow the interface to respond to ping packets. The reason is that SD-WAN sends ping probes to the overlay IP address of the remote spoke to monitor the shortcut performance and health. For this reason, if you don't enable ping access, then ping probes fail, and shortcuts are marked as dead.

Note that `net-device` is disabled by default, so make sure to enable it on spokes.

DO NOT REPRINT
© FORTINET

Overlay Stickiness and ADVPN

- Prefer shortcut negotiation over same-ISP overlays
  - Prevents shortcut negotiation over unreachable underlays (for example, internet and MPLS)



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

43

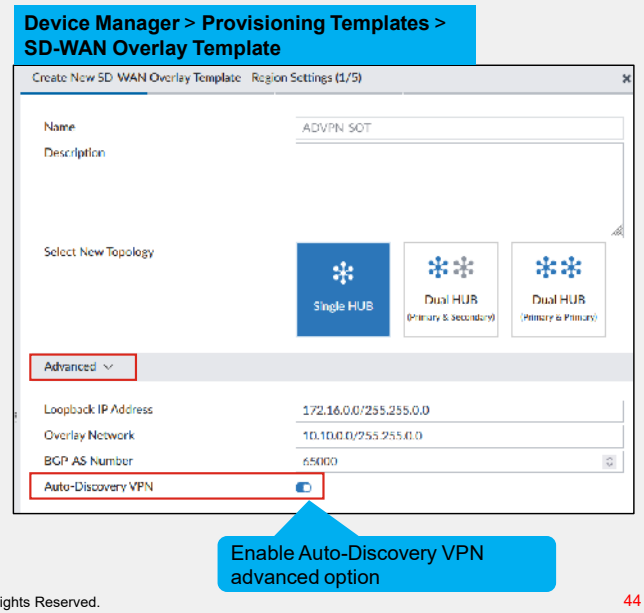
In addition to increasing performance, overlay stickiness is also important for ADVPN because it helps to prevent the spokes from trying to negotiate shortcuts over unreachable underlays.

The topology on this slide shows two overlays, one established over ISP1 and another over MPLS. ISP1 is an internet link assigned with a public IP address. MPLS is a private link assigned with a private IP address routable only within the MPLS provider network. Same-ISP overlay shortcuts are negotiated successfully because the spokes can reach the peer IP of each other through the corresponding underlays. However, cross-ISP overlays fail to establish because the ISP1 and MPLS networks are not routable between them. That is, during shortcut negotiation, a spoke using its public IP is unable to reach the private IP of the peer, and vice-versa.

DO NOT REPRINT
© FORTINET

ADVPN and FortiManager—SD-WAN Overlay Template

- ADVPN option on SD-WAN overlay template
- Hub
  - IPsec template
    - **add-route** option is disabled
    - **net-device** option is disabled
    - **auto-discovery-sender** option is enabled
- Spoke
  - IPsec template
    - **add-route** option is disabled
    - **net-device** option is enabled
    - **auto-discovery-receiver** option is enabled
- For hub and spoke
  - BGP templates
    - **recursive-next-hop** option is enabled
    - **additional-path** option is enabled



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

44

The **SD-WAN Overlay Template** on FortiManager offers the possibility to configure ADVPN with a single stitch selector. When selected, FortiManager configures the IPsec and BGP templates with the necessary parameters to configure ADVPN on the hub and spokes using Fortinet recommended settings.

For BGP, FortiManager defines the templates with `additional-path enable` and `recursive-next-hop enable`.

For IPsec, FortiManager enables `auto-discovery-sender` on the hub, `auto-discovery-receiver` on the spoke, and adjusts a few other parameters, as shown this slide.

Note that `add-route` must be disabled if you use a dynamic routing protocol. This ensures that the IKE protocol does not add routes—back to spokes to hubs—and leaves routings solely to the configured dynamic routing protocol.

DO NOT REPRINT
© FORTINET

ADVPN and FortiManager—IPsec Tunnel Templates

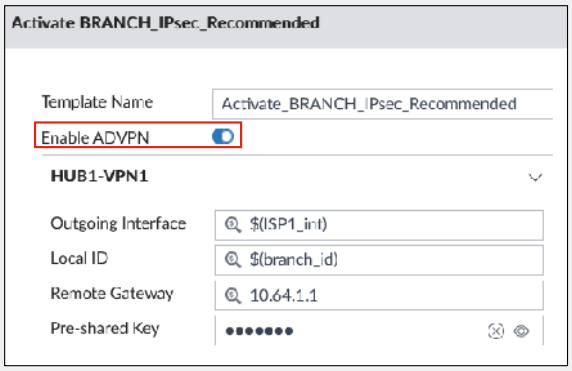
- IPsec-recommended template can set parameters required for ADVPN

- Hub
  - Enable ADVPN



  - **add-route** option is disabled
  - **net-device** option is disabled
  - **auto-discovery-sender** option is enabled

- Spoke
  - Enable ADVPN



  - **add-route** option is disabled
  - **net-device** option is enabled
  - **auto-discovery-receiver** option is enabled

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

45

If you prefer, you can use the IPsec-recommended templates, to configure ADVPN on your FortiGate. Those templates are pre-configured with recommended settings for SD-WAN, hub and spoke, and topology. When you enable the ADVPN option, the predefined recommended templates will automatically be set with the following parameters:

For hubs phase1:

- `mode-cfg` enable
- `net-device` option to disable
- `add-route` option to disable
- `auto-discovery-sender` option to enable
- `network-overlay` to enable
- `network-id` as configured

For spokes phase1:

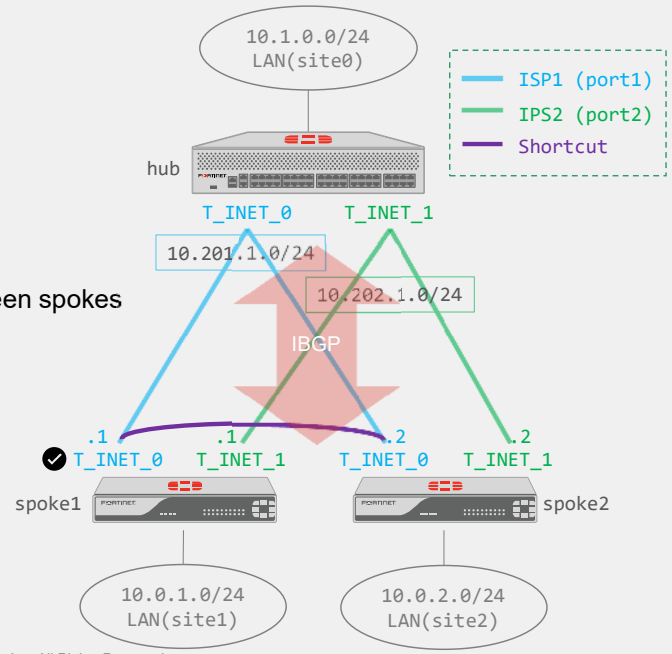
- `mode-cfg` enable
- `net-device` option to enable
- `add-route` option to disable
- `auto-discovery-receiver` option to enable
- `network-overlay` to enable
- `network-id` as configured

You can edit, review and adjust the template settings as required.
Note that `add-route` must be disabled if you use a dynamic routing protocol. This ensures that the IKE protocol does not add routes—back to spokes or back to hubs—and leaves routings solely to configured dynamic routing protocol.

DO NOT REPRINT
© FORTINET

Summary for Basic SD-WAN + ADVPN Deployment

- Hub
  - net-device: disable
  - Route reflector enabled (IBGP)
  - Overlay stickiness configured (policy route)
  - auto-discovery-sender: enable
- Spokes
  - Overlays are members of SD-WAN
  - Manual SD-WAN rule ID 1 steers traffic between spokes
  - Performance SLA
  - net-device: enable
  - auto-discovery-receiver: enable
  - allowaccess: ping
  - mode-cfg: enable



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

46

This slide shows the topology used for testing SD-WAN and ADVPN. It's essentially the same basic topology used before, except that the slide indicates the IP addresses assigned to each overlay on the spokes using IKE mode config, and a summary of the configuration applied on the hub and spokes.

The next slides show how the FortiOS routing table, the diagnostic output for SD-WAN rules, and the performance SLA change when a shortcut is established between the spokes after an endpoint behind spoke1 generates traffic to spoke2. Note that the preferred member on spoke1 is T_INET_0.

DO NOT REPRINT
© FORTINET

Before Shortcut Is Established—Spoke1

- Routing table and IPsec tunnel summary:

```
# get router info routing-table all
...
B       10.0.2.0/24 [200/0] via 10.201.1.1 [2] (recursive is directly connected, T_INET_0), 2d20h18m, [1/0]
                    [200/0] via 10.202.1.2 [2] (recursive is directly connected, T_INET_1), 2d20h18m, [1/0]
S       10.201.1.254/32 [15/0] via T_INET_0 tunnel 100.64.1.1, [1/0]
# get vpn ipsec tunnel summary
'T_INET_0' 100.64.1.1:0  selectors(total,up): 1/1  rx(pkt,err): 111933/0  tx(pkt,err): 603908/14
'T_INET_1' 100.64.1.9:0  selectors(total,up): 1/1  rx(pkt,err): 121167/0  tx(pkt,err): 627331/11
```

Recursive lookup for next hop 10.201.1.2 is T_INET_0

Installed by iked because of IKE mode config

- SD-WAN service and performance SLA:

```
# diagnose sys sdwan service

Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(20), TOS(0x0/0x0), Protocol(1: 1->65535), Mode(manual)
  Members(2):
    1: Seq_num(3 T_INET_0), alive, selected
    2: Seq_num(4 T_INET_1), alive, selected
...
# diagnose sys sdwan health-check status
Health Check(VPN_PING):
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(1.544), jitter(0.299) sla_map=0x3
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.431), jitter(0.276) sla_map=0x3
```

Preferred member

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

47

This slide shows the routing table, IPsec tunnel summary, and relevant SD-WAN diagnostic outputs before the shortcut is established.

Note that the recursive lookup for the next hop of the first IBGP route resolves to T_INET_0. This is because there is a static route to 10.201.1.0/24 that FortiOS adds on spoke1 as a result of the IKE mode config feature, which is enabled.

Note also that the SD-WAN rule is configured in manual mode, and the preferred member is T_INET_0.

DO NOT REPRINT
© FORTINET

Bringing Up Shortcut—Spoke1

- Sniffer:

```
#  diagnose sniffer packet any "host 10.0.2.101 and icmp" 4 0 l | grep T_INET_0
...
2022-12-01 19:35:37.107837 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2022-12-01 19:35:37.111114 T_INET_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2022-12-01 19:35:38.108441 T_INET_0_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2022-12-01 19:35:38.110166 T_INET_0_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2022-12-01 19:35:39.109904 T_INET_0_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2022-12-01 19:35:39.111321 T_INET_0_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

Shortcut established

- Routing table and IPsec tunnel summary:

```
# get router info routing-table all
...
B       10.0.2.0/24 [200/0] via 10.201.1.2 [2] (recursive is directly connected, T_INET_0_0), 00:06:35
                    [200/0] via 10.202.1.2 [2] (recursive is directly connected, T_INET_1), 00:06:35
C       10.201.1.0/24 is directly connected, T_INET_0
S       10.201.1.2/32 is directly connected, T_INET_0_0
# get vpn ipsec tunnel summary
'T_INET_0' 100.64.1.1:0  selectors(total,up): 1/1  rx(pkt,err): 55670/0  tx(pkt,err): 54341/88
'T_INET_1' 100.64.1.9:0  selectors(total,up): 1/1  rx(pkt,err): 41954/0  tx(pkt,err): 40524/20
'T_INET_0_0' 203.0.113.1:0  selectors(total,up): 1/1  rx(pkt,err): 73/0  tx(pkt,err): 70/2
```

Recursive lookup for next hop 10.201.1.2 is now T_INET_0_0

Installed by iked because of ADVPN

Spoke2 IP (port1)

Shortcut is up

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

48

When an endpoint with address 10.0.1.101 generates traffic to 10.0.2.101, spoke1 initially steers traffic through the parent overlay T_INET_0. At the same time, the negotiation of the shortcut between spoke1 and spoke2 begins. After the shortcut is established, spoke1 starts routing the traffic through the shortcut (T_INET_0_0).

The routing table on spoke1 also reflects the successful shortcut negotiation. FortiOS installs a static route through the shortcut to reach the overlay IP address of the remote spoke. The result is that the next hop of the IBGP routes to 10.0.2.0/24 now recursively resolves to the shortcut, instead of to the parent tunnel. Note that the IKE mode config connected route to 10.201.1.0/24 through the parent tunnel is still there. It is just that the new static route to 10.201.1.2/32 is better (more specific).

In addition, the IPsec tunnel summary output also shows the new shortcut. Note that the shortcut name is composed by the parent tunnel name followed by an underscore and a number.

DO NOT REPRINT
© FORTINET

Bringing Up Shortcut—Spoke1 (Contd)

- SD-WAN service and performance SLA:

```
# diagnose sys sdwan service
Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(6), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Member sub interface(4):
    2: seq_num(3), interface(T_INET_0):
       1: T_INET_0_0(28)
  Members(3):
    1: Seq_num(3 T_INET_0_0), alive, selected
    2: Seq_num(3 T_INET_0), alive, selected
    3: Seq_num(4 T_INET_1), alive, selected
...

# diagnose sys sdwan health-check status
Health Check(VPN):
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(1.613), jitter(0.327), ..., sla_map=0x3
Seq(3 T_INET_0_0): state(alive), packet-loss(0.000%) latency(1.358), jitter(0.292),..., sla_map=0x3
Seq(4 T_INET_1): state(alive), packet-loss(0.000%) latency(1.500), jitter(0.287), ..., sla_map=0x3

# diagnose sniffer packet T_INET_0_0 "" 4
...
0.409032 T_INET_0_0 -- 10.201.1.1 -> 10.201.1.2: icmp: echo request
0.410572 T_INET_0_0 -- 10.201.1.2 -> 10.201.1.1: icmp: echo reply
```

Shortcut listed as subinterface

Shortcut becomes preferred member

10 seconds after shortcut is established, FortiGate starts monitoring the shortcut using ping and targeting remote overlay IP (10.201.1.2)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

49

After the shortcut is established, FortiGate lists the shortcut as a subinterface in the SD-WAN rule diagnostic output. FortiGate also starts monitoring the shortcut using ping as protocol and targeting the overlay IP of the remote spoke (`10.201.1.2`), 10 seconds after the shortcut is established.

Note that FortiGate moves the shortcut to the top of the list, so it takes precedence over the parent tunnel.

DO NOT REPRINT
© FORTINET

Debugging Shortcut Negotiation

- Specify multiple IP addresses to filter the IKE real-time debug
  - Useful when debugging ADVPN shortcut messages and spoke-to-spoke negotiations

```
diagnose debug console timestamp enable
diagnose vpn ike log filter clear
diagnose vpn ike log filter mdst-addr4 <ip.of.hub> <ip.of.spoke>
diagnose debug application ike -1
diagnose debug enable
```

filter on multiple IPv4 destination address (mdst-addr6 for IPv6 address)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

50

Enable debug of IKE to debug ADVPN shortcut negotiation, as shown on this slide. In the IKE log filter, you can specify multiple IP addresses to print debug messages for. This is very useful when debugging ADVPN shortcuts and spoke-to-spoke ADVPN negotiation issues.

DO NOT REPRINT
© FORTINET

Shortcut Debug

- Hub output—hub sends shortcut offer to spoke1:

```
ike 0: shortcut T_INET_0_0:192.2.0.1:0 to T_INET_0_1:203.0.113.1:0 for 10.0.1.101->10.0.2.101
ike 0 send shortcut-offer to T_INET_0_0
ike 0:T_INET_0_0:3543: sending NOTIFY msg
ike 0:T_INET_0_0:214:3543: send informational
```

- Spoke1 output—spoke1 receives shortcut offer from hub, and then sends shortcut query to hub:

```
ike 0:T_INET_0:129: received informational request
ike 0:T_INET_0:129: processing notify type SHORTCUT_OFFER
ike 0:T_INET_0: shortcut-offer 10.0.1.101->10.0.2.101 psk 64 ppk 0 ver 2 mode 0, peer-addr 203.0.113.1:500
ike 0 looking up shortcut by addr 10.0.2.101, name T_INET_0, peer-addr 203.0.113.1:500
ike 0:T_INET_0: send shortcut-query 9065761962601467474 07409008f7fbd17e/0000000000000000 192.2.0.1
10.0.1.101->10.0.2.101 psk 64 ttl 32 nat 0 ver 2 mode 0
```

- Hub output—hub receives shortcut query from spoke1 and forwards it to spoke2:

```
ike 0:T_INET_0_0:214: received informational request
ike 0:T_INET_0_0:214: processing notify type SHORTCUT_QUERY
ike 0:T_INET_0_0: recv shortcut-query 9065761962601467474 07409008f7fbd17e/0000000000000000 192.2.0.1
10.0.1.101->10.0.2.101 psk 64 ppk 0 ttl 32 nat 0 ver 2 mode 0
ike 0:T_INET_0_0: iif 19 10.0.1.101->10.0.2.101 0 route lookup oif 20 T_INET_0_0 gwy 10.201.1.2
ike 0: shared dev tunnel lookup, tun-id=10.201.1.2
ike 0:T_INET_0_1: forward shortcut-query 9065761962601467474 07409008f7fbd17e/0000000000000000 192.2.0.1
10.0.1.101->10.0.2.101 psk 64 ppk 0 ttl 31 ver 2 mode 0, ext-mapping 192.2.0.1:500
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

51

This slide shows an example of the real-time debug output displayed by IKE during a shortcut negotiation triggered after spoke1 communicates with spoke2 over the parent tunnel.

The shortcut negotiation starts on the hub. On the hub, the tunnel to spoke1 is `T_INET_0_0`, and its remote gateway is `192.2.0.1`. The tunnel to spoke2 is `T_INET_0_1`, and its remote gateway is `203.0.113.1`. The hub identifies that a shortcut can be negotiated between the two spokes for traffic sourced from `10.0.1.101` and destined to `10.0.2.101`. For this reason, the hub sends a shortcut offer to spoke1 to initiate the shortcut negotiation.

Spoke1 then receives the shortcut offer from the hub, to which it replies with a shortcut query. The hub then receives the shortcut query from spoke1 and forwards it to spoke2.

DO NOT REPRINT
© FORTINET

Shortcut Debug (Contd)

- Spoke2—spoke2 receives shortcut query from hub, and then sends shortcut reply to hub:

```
ike 0:T_INET_0:119: received informational request
ike 0:T_INET_0:119: processing notify type SHORTCUT_QUERY
ike 0:T_INET_0: recv shortcut-query 9065761962601467474 07409008f7fbd17e/0000000000000000 192.2.0.1
10.0.1.101->10.0.2.101 psk 64 ppk 0 ttl 31 nat 0 ver 2 mode 0
ike 0:T_INET_0: iif 19 10.0.1.101->10.0.2.101 route lookup oif 7 port5 gwy 0.0.0.0
ike 0:T_INET_0: shortcut-query received from 192.2.0.1:500, local-nat=no, peer-nat=no
ike 0:T_INET_0: send shortcut-reply 9065761962601467474 07409008f7fbd17e/53f421af88136a87 203.0.113.1 to
10.0.1.101 psk 64 ppk 0 ver 2 mode 0
```

- Hub output—hub receives shortcut reply from spoke2 and forwards it to spoke1:

```
ike 0:T_INET_0_1:213: received informational request
ike 0:T_INET_0_1:213: processing notify type SHORTCUT_REPLY
ike 0:T_INET_0_1: recv shortcut-reply 9065761962601467474 07409008f7fbd17e/53f421af88136a87 203.0.113.1 to
10.0.1.101 psk 64 ppk 0 ver 2 mode 0 ext-mapping 0.0.0.0:0
ike 0:T_INET_0: iif 20 10.0.2.101->10.0.1.101 route lookup oif 20 T_INET_0 gwy 10.201.1.1
ike 0:T_INET_0_0: forward shortcut-reply 9065761962601467474 07409008f7fbd17e/53f421af88136a87 203.0.113.1
to 10.0.1.101 psk 64 ppk 0 ttl 31 ver 2 mode 0 ext-mapping 203.0.113.1:500
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

52

Spoke2 receives the shortcut query from the hub, to which it responds with a shortcut reply. The hub then receives the shortcut reply from spoke2 and forwards it to spoke1.

DO NOT REPRINT
© FORTINET

Shortcut Debug (Contd)

- Spoke1 output—spoke1 receives shortcut reply from hub, and then initiates tunnel negotiation with spoke2:

```
ike 0:T_INET_0:129: received informational request
ike 0:T_INET_0:129: processing notify type SHORTCUT_REPLY
ike 0:T_INET_0: recv shortcut-reply 9065761962601467474 07409008f7fbd17e/53f421af88136a87 203.0.113.1 to
10.0.1.101 psk 64 ppk 0 ver 2 mode 0 ext-mapping 203.0.113.1:500
ike 0:T_INET_0: iif 19 10.0.2.101->10.0.1.101 route lookup oif 7 port5 gwy 0.0.0.0
ike 0:T_INET_0: shortcut-reply received from 203.0.113.1:500, local-nat=no, peer-nat=no
ike 0:T_INET_0: created connection: 0xf6d34b0 3 192.2.0.1->203.0.113.1:500.
ike 0:T_INET_0: adding new dynamic tunnel for 203.0.113.1:500
ike 0:T_INET_0_0: tunnel created tun_id 203.0.113.1/::203.0.113.1 remote_location 0.0.0.0
ike 0:T_INET_0_0: added new dynamic tunnel for 203.0.113.1:500
ike 0:T_INET_0_0: shortcut selector added, new serial 1
```

- Spoke1 output—shortcut `T_INET_0_0` is established:

```
ike 0:T_INET_0_0:153:T_INET_0_0:3494: sending SNMP tunnel UP trap
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

53

Spoke1 receives the shortcut reply from the hub, and then starts the IKE negotiation to bring up the shortcut tunnel. Shortly after, the shortcut (`T_INET_0_0`) is established.

Note that for this example, the name of the spoke1 tunnel on the hub is the same as the shortcut tunnel on spoke1. However, the tunnels are different. That is, on the hub, T_INET_0_0 is the child tunnel that connects the hub to spoke1, and that results from configuring the hub as the IPsec dial-up server. On spoke1, T_INET_0_0 is the shortcut negotiated with spoke2. Both the child dial-up and shortcut tunnels use the same name convention.

DO NOT REPRINT
© FORTINET

Identify Shortcuts on Log Messages

- Log messages field **advpnsc**
  - Shortcut tunnel: `advpnsc=1`
  - Other tunnels: `advpnsc=0`
- Filter ADVPN shortcut log messages on FortiGate
  - `execute log filter category event`
  - `execute log filter field advpnsc 1`

```
branch1_fgt # exec log filter category event
branch1_fgt # exec log filter field advpnsc 1
branch1_fgt # exec log display
23 logs found.
10 logs returned.
27.8% of logs has been searched.

1: date=2023-01-19 time=01:10:29
eventtime=1674119429746059820 tz="-0800"
logid="0101037134" type="event" subtype="vpn"
level="notice" vd="root" logdesc="IPsec phase 1 SA
deleted" [...] vpntunnel="T_INET_0_0" advpnsc=1
```

ADVPN shortcut flag

Shortcut tunnel name

- On FortiAnalyzer
  - Column **AD-VPN Shortcut**

Log View > Event > VPN

branch1_fgt · Last 4 Hours · 07:11:50 To 11:11:49

AD-VPN Shortcut = "1" | Add Filter

#	Action	Message	VPN tunnel	AD-VPN Shortcut
1	delete_phase..	delete IPsec phase 1 SA	T_INET_0_0	1
2	phase2-down	IPsec phase 2 status ch...	T_INET_0_0	1
3	phase2-down	Psec phase 2 status ch...	T_INET_0_0	1
4	tunnel-down	ec connection status ...	T_INET_0_0	1
5	negotiate	ress IPsec phase 2	T_INET_0_0	1

ADVPN shortcut filter

ADVPN shortcut flag

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

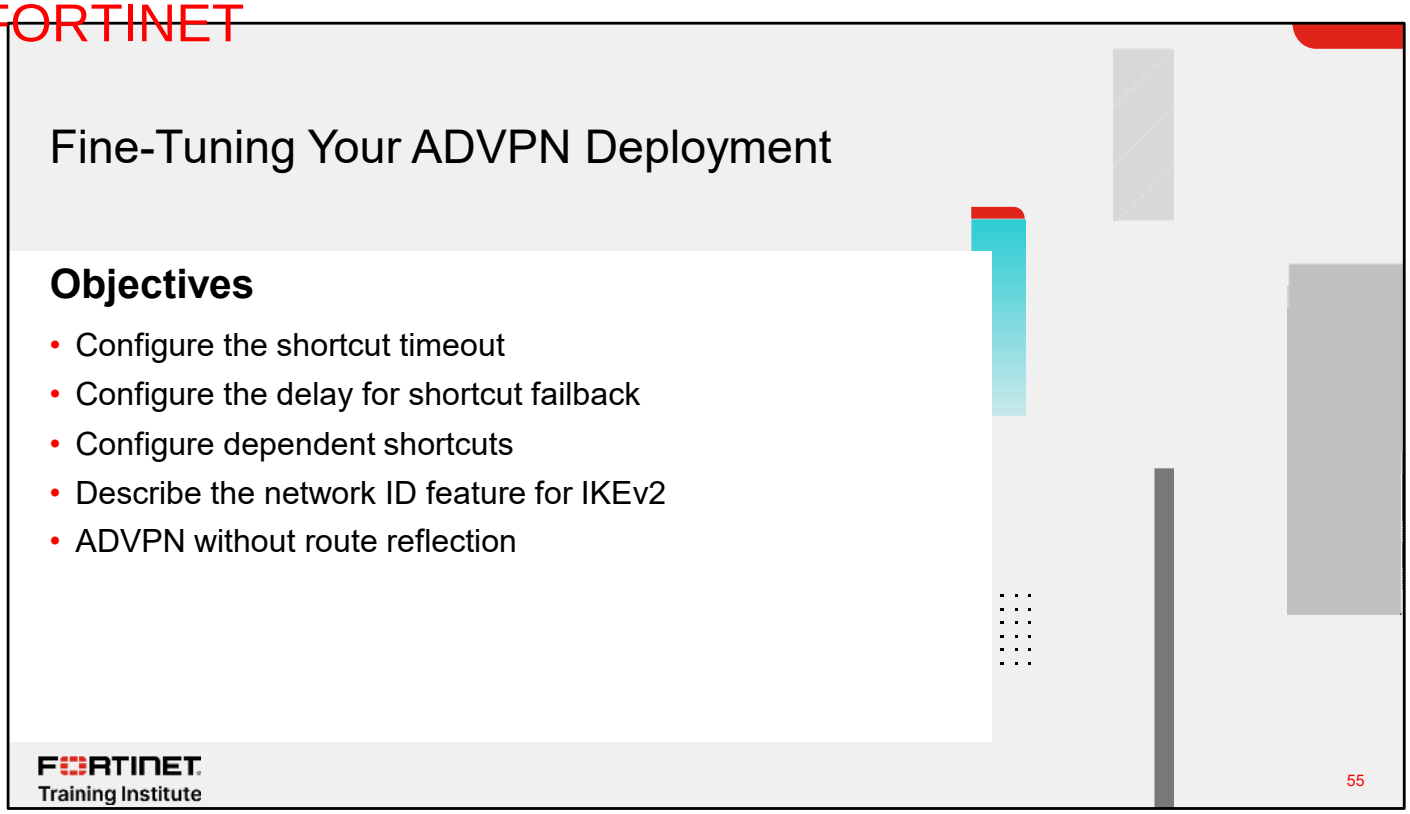
54

When reviewing VPN log messages, the field `advpnsc` will help you identify the shortcut VPN tunnels. FortiGate will set `advpnsc` value `1` for any log messages related to shortcut tunnels; for any other tunnel, the `advpnsc` value is set to `0`.

When viewing the log on FortiAnalyzer, you can display the column **AD-VPN Shortcut** to view the `advpnsc` flag.

This slide shows examples of log display, filtered to see only ADVPN-related messages on the FortiGate CLI and on the FortiAnalyzer GUI.

DO NOT REPRINT
© FORTINET



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

By demonstrating competence in fine-tuning your ADVPN deployment, you will be able to configure the network to save resources, speed up routing convergence, and improve the user experience for sensitive applications.

DO NOT REPRINT
© FORTINET

Timing Out Idle Shortcuts

- By default, shortcuts inherit lifetime settings of parents
  - Idle shortcuts may remain up for a long time
    - Waste of resources
- Set an idle timeout to shortcuts to save resources
  - FortiGate brings down shortcuts with no user traffic for a set period of time
    - Health check traffic doesn't count

- Spoke configuration:

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        set idle-timeout enable
        set idle-timeoutinterval 5
    next
end
```

- IKE debug upon timeout:

```
ike 0:T_INET_0_0: connection idle time-out
ike 0:T_INET_0_0: deleting
ike 0:T_INET_0_0: flushing
...
ike 0:T_INET_0_0: sending SNMP tunnel DOWN trap
for T_INET_0_0
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

56

By default, shortcuts inherit their lifetime from their parent tunnels. That is, if the lifetime of a parent tunnel is 1800 seconds, then its shortcuts also use the same lifetime.

Usually, the lifetime of an IPsec tunnel is set for a few hours, which means that idle shortcuts remain up for the same period of time unless they are manually flushed or are detected dead by DPD, which can lead to unnecessary resource utilization on spokes.

To reduce resource utilization, you can configure an idle timeout for shortcuts. This way, FortiGate brings down a shortcut if it hasn't been used for forwarding user traffic in a certain period of time. Note that health check traffic is not considered user traffic and, therefore, it doesn't prevent an idle tunnel from timing out.

This slide shows how to set an idle timeout for shortcuts. The example configuration sets an idle timeout of five minutes. This slide also shows some of the relevant messages seen in the IKE debug when a shortcut times out.

DO NOT REPRINT
© FORTINET

Delaying Failback to Recovered Shortcut

- Delay use of recovered members
  - Including shortcuts
- Wait until `hold-down-time` has passed
  - More accurate monitoring
  - Prevent impact to:
    - Sensitive applications during brownout conditions and SLA changes
    - CPU usage caused by session re-evaluation
- Highly recommended for SD-WAN + ADVPN + lowest cost (SLA) rules deployments

- Example configuration (default = `0`):

```
config system sdwan
    config service
        edit 1
            set hold-down-time 20
        next
    end
end
```

- SD-WAN rule status:

```
# diagnose sys sdwan service 1

Service(1): Address Mode(IPV4) flags=0x200 use-
shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(1: 1->65535),
Mode(sla), sla-compare-order
Hold down time(20) seconds, Hold start 84259
second, now 84270
...
```

Value decrease when timer active (SLA recovered for higher priority member)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

57

When you use lowest cost (SLA) rules, FortiGate uses the lowest cost member that meets the configured SLA target. If a member doesn't meet the SLA target, it is considered out of SLA and therefore, placed at the bottom of the outgoing interface list.

Links that are impacted by brownout conditions—that is, the link is up but its quality is degraded—may switch in and out of SLA states in short periods of time. This may result in poor user experience for sensitive applications such as voice or video, as well as high CPU utilization caused by frequent session re-evaluation.

To prevent immediate failback to a recovered member, including recovered shortcuts, you can configure the `hold-down-time` setting on an SD-WAN rule. The result is that FortiGate waits a set number of seconds before moving a recovered member back to its original place in the outgoing interface list. This way, FortiGate can monitor the member more accurately and ensure it meets the required SLA target for a certain amount of time before it starts to steer traffic to it.

This slide shows an example configuration of the `hold-down-time` setting, which is set to 20 seconds. The setting applies to all members configured in the rule, including ADVPN shortcuts. Configuring a value of at least 20 seconds on SD-WAN and ADVPN deployments using lowest cost (SLA) rules is highly recommended.

Note that the setting affects only the lowest cost (SLA) and best quality rules. The `hold-down-time` default value is `0`, which means there is no failback delay. When you set it to a value higher than 0, the service SD-WAN rule status reflects the change, as shown on this slide.

DO NOT REPRINT
© FORTINET

Making Shortcuts Lifetime Dependents of Parents

- By default, lifetime of a shortcut is independent of parent's
  - Shortcuts remain up if parent goes down
    - Waste of resources
    - Slows down routing convergence
- Bring down shortcuts immediately after parent goes down (default = `independent`):

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        set auto-discovery-shortcuts dependent
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

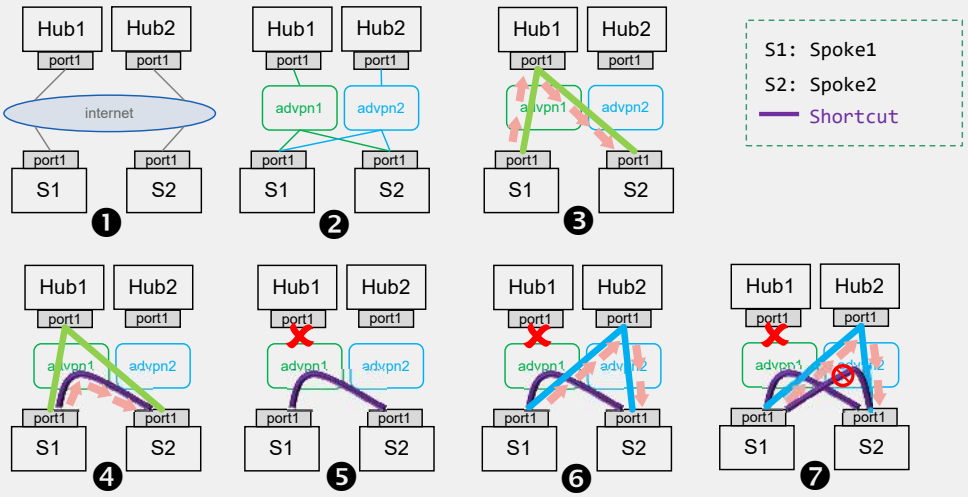
58

By default, shortcuts are not brought down if the parent tunnel goes down. That is, the lifetime of a shortcut is independent of its parent tunnel.

To save resources and obtain a faster routing convergence, you can instruct FortiGate to bring down a shortcut immediately after its parent tunnel goes down by configuring the `auto-discovery-shortcuts` to `dependent`, as shown on this slide. By default, the setting is set to `independent`.

DO NOT REPRINT
© FORTINET

Allowing Multiple Shortcuts Over Same Pair of Gateways



- Lifetime of shortcuts are independent of parent tunnels

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

59

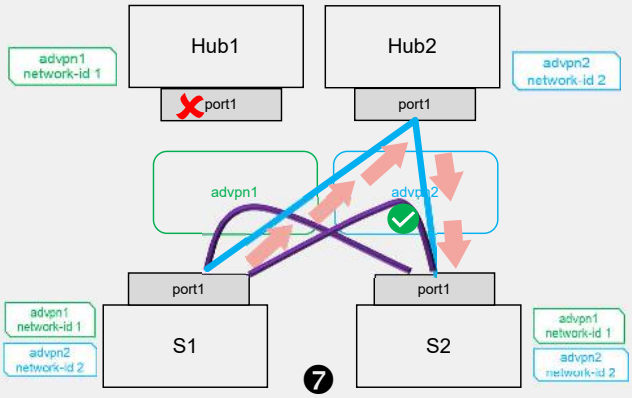
Consider the dual-hub topology shown on this slide, which shows two hubs and two spokes.

1. All devices have one internet underlay on port1.
2. The spokes have dual IPsec overlays, one to each hub. ADVPN is enabled on the overlays.
3. Traffic from spoke1 to spoke2 initially flows through the parent tunnel to hub1 and a shortcut negotiation starts.
4. A shortcut between spoke1 and spoke2 is established, and spoke-to-spoke traffic starts to flow through the shortcut.
5. On hub1, internet access is lost, which results in the parent tunnel between hub1 and spoke1 going down. The shortcut, however, remains up because its lifetime is independent of the parent tunnel (default behavior).
6. Routing information on spoke1 is updated, and now spoke2 is reachable through hub2, which triggers a session re-evaluation. Traffic from spoke1 to spoke2 starts flowing through hub2, and a shortcut negotiation over hub2 starts. At this point, the first shortcut negotiated over hub1 is still up.
7. Shortcut negotiation over hub2 fails because there is already a shortcut established between the same pair of gateways. That is, the local and remote gateway IP addresses are the same.

DO NOT REPRINT
© FORTINET

Allowing Multiple Shortcuts Over Same Pair of Gateways (Contd)

- Assign different network IDs to overlays to allow multiple shortcuts between two spokes



- Example configuration:

```
config vpn ipsec phase1-interface
    edit "advpn1"
        set ike-version 2
        set network-overlay enable
        set network-id 1
    next
    edit "advpn2"
        set ike-version 2
        set network-overlay enable
        set network-id 2
    next
end
```

- Supported by IKEv2 only
- Can be used for gateway lookup
  - Multiple IPsec dial-up VPNs on the same local gateway

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

60

To allow the spokes to establish shortcuts between the same pair of local and remote gateway IP addresses, set a network ID on each overlay, as shown on this slide. The network ID feature is supported by IKEv2 only.

When you set a different network ID for each overlay, both shortcuts, the one negotiated over hub1 and the one negotiated over hub2, can be established at the same time.

The network ID is also useful for gateway lookup when you configure multiple dial-up VPNs that terminate on the same local IP address, and you want users to connect to a particular VPN. Since the network ID is in the first packet of the initiator, it can be used to identify the correct phase 1 gateway configuration to match.

Note that you need to enable `network-overlay` to configure `network-id`.

DO NOT REPRINT
© FORTINET

ADVPN Without Route Reflection

- Allow ADVPN without BGP route reflectors—Static, OSPF, or BGP routing
- Routing over shortcuts managed by IPsec
- Each site announces its protected subnet during IPsec negotiation
- Also called ADVPN with Phase2 selector

- Advantages and limitations

	Phase 2 selector	Route-Reflection
Routing (Overlay / LAN)	Static / BGP / OSPF	BGP
Route Reflector on hub	No	Yes
BGP route aggregation	Yes	No
Prefixes redistribution	No	Yes
Definition of local networks on IPsec phase2	Yes	No
VRF aware overlay	Not compatible	Compatible

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

61

You learned how FortiGate can use BGP route-reflection to allow routing through ADVPN tunnels. What about ADVPN on network with static routing?
Before FortiOS 7.2.0, it was not possible to use auto-discovery VPN without BGP routing. Now, FortiGate can take advantage of Phase 2 selectors to allow ADVPN without route-reflection, and so, makes ADVPN compatible with static routing.

Instead of BGP route-reflection, routing over ADVPN shortcuts is managed by IPsec protocol. With the use of phase2 selectors, each site announces its protected subnets during IPsec negotiation for tunnel establishment.

Route-reflection and phase2 selector are both solution to allow routing through ADVPN shortcuts. Each has its own advantage and limitations. The main elements to take into account when selecting an option are summarized in the table shown on this slide.
With BGP routing, the two options are possible. Note that phase 2 selector allow BGP route aggregation and therefore simplify the routing table. On hubs, it reduces the BGP daemon process load and can help improving scalability on large ADVPN networks. Phase 2 selector option also allows a simplified BGP configuration (no route-reflection, no add-path), but it adds configuration requirements on IPsec phase 2. The administrator must define each local network.

DO NOT REPRINT
© FORTINET

Configure ADVPN with Phase2 Selector—Hub

- Hub configuration
  - IPsec

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
        ...
        set ike-version 2
        set net-device disable
        set auto-discovery-sender enable
    next
end
```

- with BGP routing

```
config router bgp
    ...
    set ibgp-multipath enable
    set additional-path enable
    set additional-path-select 3
    config neighbor-group
        edit "Branches_INET_0"
            ...
            set additional-path send
            set adv-additional-path 3
            set route-reflector-client disable
        next
```

keep advertisement of additional path for SD-WAN

disable route reflection (default setting)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

62

To use ADVPN with phase2 selector, IPsec configuration on the hub is similar to one you will do for ADVPN with route-reflector. On the phase1, you need to enable ADVPN `auto-discovery-sender` and should keep `net-device` to disable. There are no specific setting required for the phase2 on the hub.

You should make sure that required routing for the overlay is correctly configured. It can be static, OSPF or BGP. If you use BGP routing and SD-WAN, you should allow additional path advertisement to allow SD-WAN to learn all possible path and correctly steer the traffic. Keep the parameter `route-reflector-client` disabled (default setting).

DO NOT REPRINT
© FORTINET

Configure ADVPN with Phase2 Selector—Spoke

- Spoke configuration
  - IPsec

```
config vpn ipsec phase1-interface
    edit "T_INET_0"
         ...
         set ike-version 2
         set net-device enable
         set add-route enable
         set mode-cfg enable
         set auto-discovery-receiver enable
         set mode-cfg-allow-client-selector enable
         ...
    next
config vpn ipsec phase2-interface
    edit " T_INET_0_0 "
         ...
         set src-addr-type name
         set dst-addr-type name
         set src-name "LAN_Net"
         set dst-name "all"
    next
...
```

Allow mode-cfg client to use custom selectors

Select `name` to define src/dst addresses as address object
Select `subnet` for direct configuration as subnet

Local subnet or subnets

Both selectors must be of same type (`name` or `subnet`)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

63

To use phase2 selectors to inject IKE routes on ADVPN shortcuts tunnel you must adjust IPsec configuration on the spokes as shown on this slide.

On IPsec phase1, you must enable configuration method (`mode-cfg enable`) and allow addition of routes per the destination selector (`add-route enable`). You must also enable the parameter `mode-cfg-allow-client-selector` to allow the configuration of custom selector at phase2 level.

On the phase2, you will define local subnets that will be injected on tunnel and tunnel shortcut establishment. You can define the local subnet directly as a subnet or as firewall address object. On configuration example shown on this slide, the source subnet is defined with firewall address group LAN_net. This type of configuration with address object is required when you want to define multiple subnets as phase2 source.

ADVPN shortcuts establishment are follow the same steps similar with Phase2 selectors and route-reflection—shortcut offer, shortcut query, shortcut reply. For troubleshooting and to review shortcut tunnel establishment you can use the commands seen earlier in this lessons.

DO NOT REPRINT
© FORTINET

Review

- ✓ Describe the IPsec and BGP settings required for overlays in a standard hub-and-spoke deployment
- ✓ Configure BGP route reflector to exchange prefixes between spokes
- ✓ Adjust BGP timers to speed up convergence, failure detection, and recovery
- ✓ Adjust DPD settings to speed up detection of dead gateways
- ✓ Configure overlay stickiness to improve performance and prevent outage during shortcut negotiation
- ✓ Configure BGP to exchange all available paths
- ✓ Configure FEC to improve quality of impaired overlays
- ✓ Configure forced and on-demand packet duplication
- ✓ Describe ADVPN operation and configuration, and its support in SD-WAN
- ✓ Configure network ID to allow multiple overlays between the same pair of gateways

FORTINET Training Institute

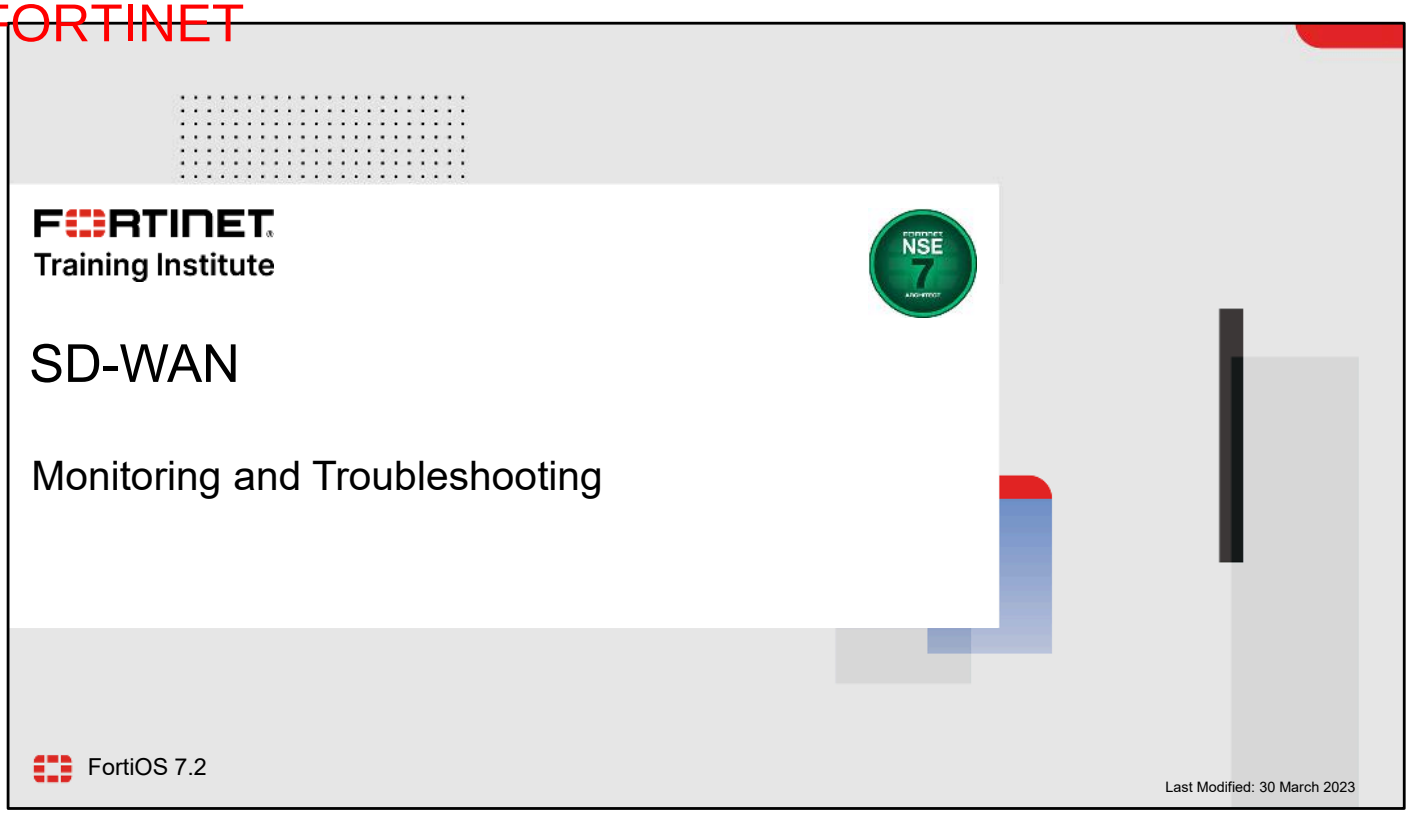
© Fortinet Inc. All Rights Reserved.

64

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

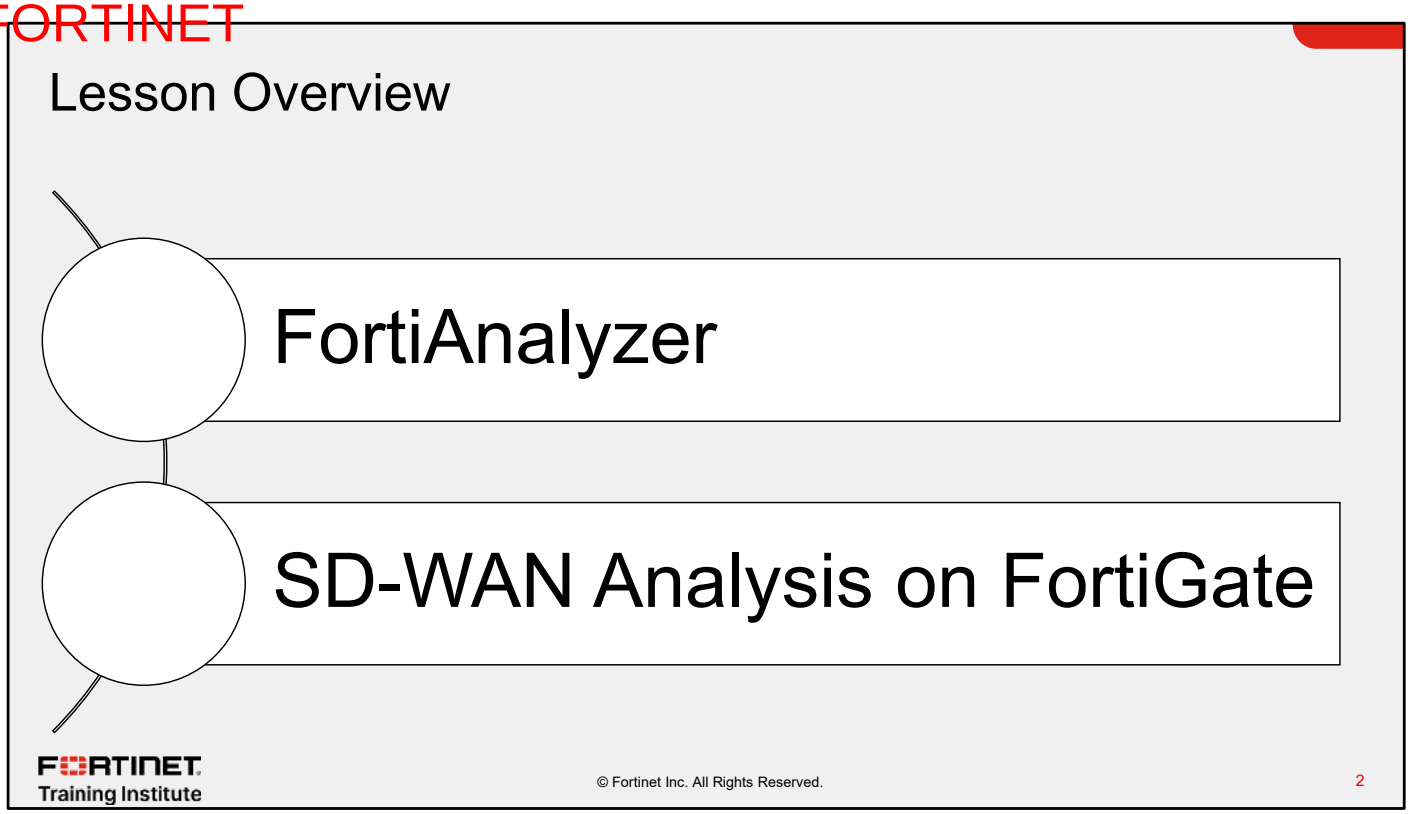
By mastering the objectives covered in this lesson, you learned how to set up SD-WAN and ADVPN on a basic hub-and-spoke topology. You also learned how to tune your IPsec, BGP, and ADVPN configuration to achieve faster routing convergence and failover.

DO NOT REPRINT
© FORTINET



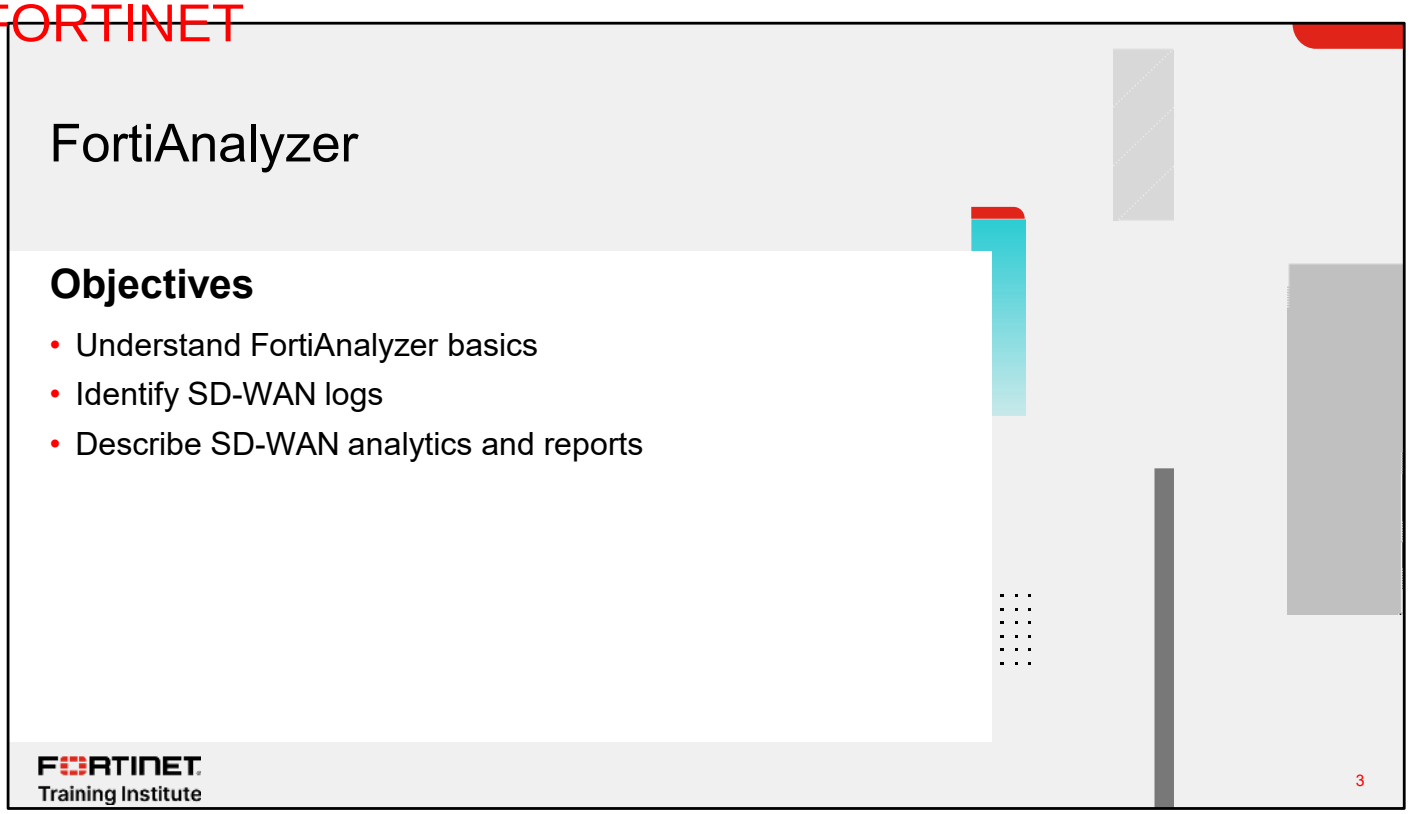
In this lesson, you will learn how FortiAnalyzer can help with SD-WAN networks monitoring and analysis. You will also learn about CLI commands available on FortiGate for monitoring and troubleshooting SD-WAN deployments.

DO NOT REPRINT
© FORTINET



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

DO NOT REPRINT
© FORTINET



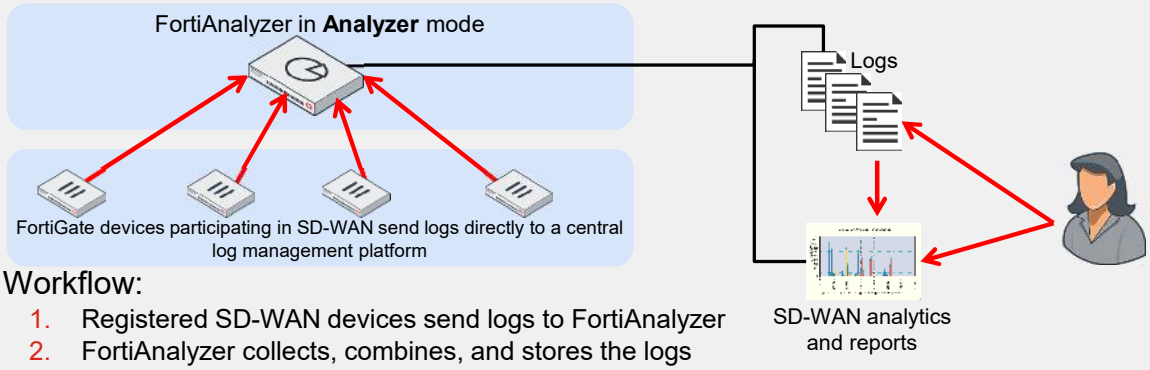
After completing this section, you should be able to achieve the objectives shown on this slide.

By understanding the logging and reporting features available on FortiAnalyzer for SD-WAN, you will be able to troubleshoot and monitor the health and performance of all FortiGate devices in your network from a single device.

DO NOT REPRINT
© FORTINET

What Is FortiAnalyzer?

- Centralized log repository
  - Aggregates SD-WAN log data from one or more Fortinet devices
  - Creates a single view of SD-WAN events taking place on a range of devices



- Workflow:
  1. Registered SD-WAN devices send logs to FortiAnalyzer
  2. FortiAnalyzer collects, combines, and stores the logs
  3. Administrators:
     - View logs for SD-WAN
     - View analytics for SD-WAN
     - Configure, request, and view reports for SD-WAN (based on log data)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

FortiAnalyzer aggregates log data from one or more Fortinet devices, including FortiGate devices that participate in SD-WAN. FortiAnalyzer acts as a centralized log repository and provides a single channel for accessing your complete SD-WAN network data, so you don't need to access multiple SD-WAN devices several times a day.

The logging and reporting workflow operates as follows:

1. Registered SD-WAN devices send logs to FortiAnalyzer.

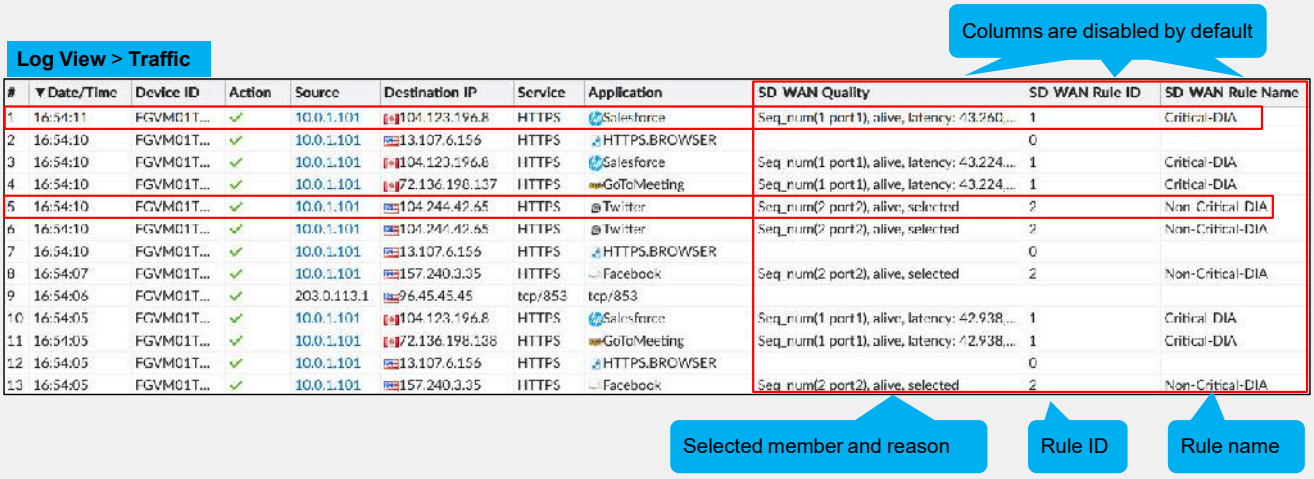
2. FortiAnalyzer collects, combines, and stores SD-WAN logs in a way that makes it easy to search and run reports.

3. Administrators can connect to the FortiAnalyzer GUI to view SD-WAN logs, analytics, and reports.

DO NOT REPRINT
© FORTINET

SD-WAN Traffic Logs

- Enable SD-WAN columns to differentiate SD-WAN sessions from non-SD-WAN sessions



#	Date/Time	Device ID	Action	Source	Destination IP	Service	Application	SD WAN Quality	SD WAN Rule ID	SD WAN Rule Name
1	16:54:11	FGVM01T...	✓	10.0.1.101	104.123.196.8	HTTPS	Salesforce	Seq_num(1 port1), alive, latency: 43.260,...	1	Critical-DIA
2	16:54:10	FGVM01T...	✓	10.0.1.101	13.107.6.156	HTTPS	HTTPS.BROWSER		0	
3	16:54:10	FGVM01T...	✓	10.0.1.101	104.123.196.8	HTTPS	Salesforce	Seq_num(1 port1), alive, latency: 43.224,...	1	Critical-DIA
4	16:54:10	FGVM01T...	✓	10.0.1.101	72.136.198.137	HTTPS	GoToMeeting	Seq_num(1 port1), alive, latency: 43.224,...	1	Critical-DIA
5	16:54:10	FGVM01T...	✓	10.0.1.101	104.244.42.65	HTTPS	Twitter	Seq_num(2 port2), alive, selected	2	Non-Critical-DIA
6	16:54:10	FGVM01T...	✓	10.0.1.101	104.244.42.65	HTTPS	Twitter	Seq_num(2 port2), alive, selected	2	Non-Critical-DIA
7	16:54:10	FGVM01T...	✓	10.0.1.101	13.107.6.156	HTTPS	HTTPS.BROWSER		0	
8	16:54:07	FGVM01T...	✓	10.0.1.101	157.240.3.35	HTTPS	Facebook	Seq_num(2 port2), alive, selected	2	Non-Critical-DIA
9	16:54:06	FGVM01T...	✓	203.0.113.1	96.45.45.45	tcp/853	tcp/853			
10	16:54:05	FGVM01T...	✓	10.0.1.101	104.123.196.8	HTTPS	Salesforce	Seq_num(1 port1), alive, latency: 42.938,...	1	Critical-DIA
11	16:54:05	FGVM01T...	✓	10.0.1.101	72.136.198.138	HTTPS	GoToMeeting	Seq_num(1 port1), alive, latency: 42.938,...	1	Critical-DIA
12	16:54:05	FGVM01T...	✓	10.0.1.101	13.107.6.156	HTTPS	HTTPS.BROWSER		0	
13	16:54:05	FGVM01T...	✓	10.0.1.101	157.240.3.35	HTTPS	Facebook	Seq_num(2 port2), alive, selected	2	Non-Critical-DIA

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

The **Forward Traffic** logs page is useful to identify how sessions are distributed in SD-WAN and the reason. Make sure to enable the **SD-WAN Rule Name** and **SD-WAN Quality** columns, which are disabled by default. The former indicates the matched SD-WAN rule for a session, and the latter indicates the member the session was steered to and the reason. You can also enable the **SD-WAN Rule ID** column, and, when applicable, the **SD-WAN Internet Service** column. Note that SD-WAN rule ID 0 corresponds to traffic steered according to the default SD-WAN rule.

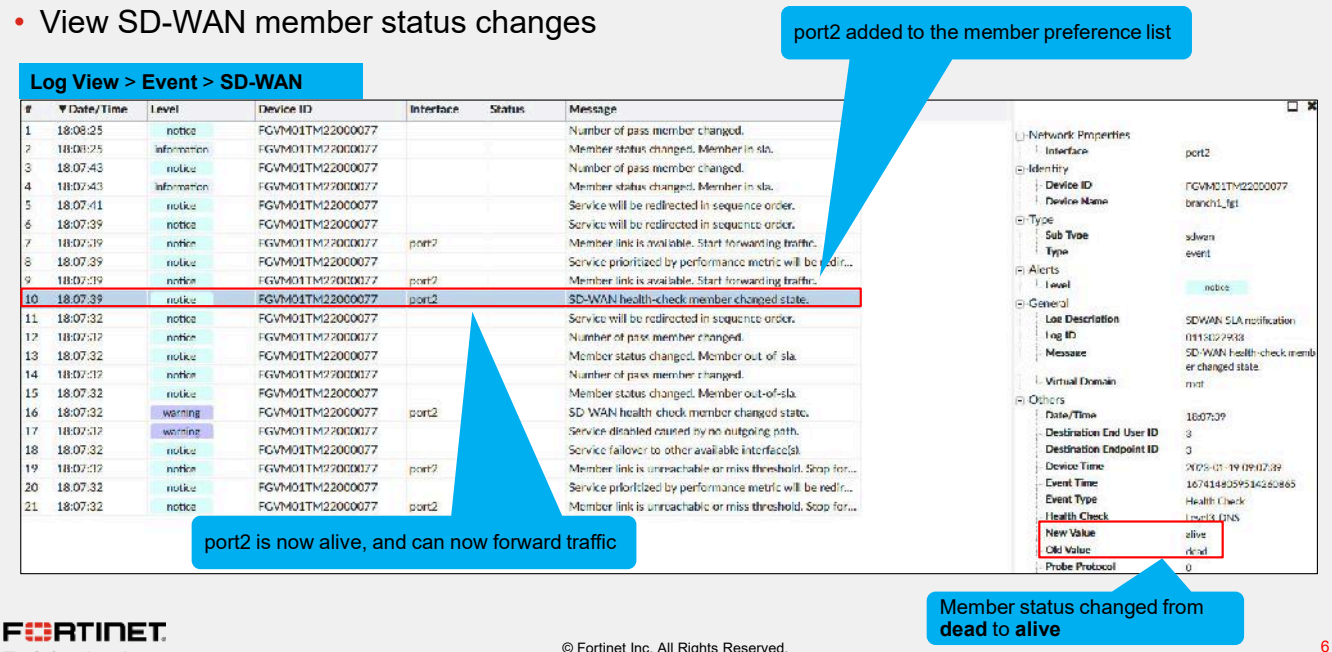
The table on this slide shows multiple sessions, including SD-WAN sessions and non-SD-WAN sessions. The SD-WAN sessions include information in the **SD-WAN Quality, SD-WAN Rule ID**, and **SD-WAN Rule Name**. The first session in this table is an SD-WAN session, and was identified as a **Salesforce** application. It matched the **Critical-DIA** rule, and was sent to port1. The reason for selecting port1 was because it had the lowest latency.

The fifth session in the table, which is also an SD-WAN session, was identified as a **Twitter** application, matched the **Non-Critical-DIA** rule, and was sent to port2. The rule **Non-Critical-DIA** instructs FortiGate to steer matching traffic to port2 only, provided the port is alive. This behavior matches the reason described in the **SD-WAN Quality** column for that session.

DO NOT REPRINT
© FORTINET

SD-WAN Events

- View SD-WAN member status changes



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

Like on the FortiGate GUI, FortiAnalyzer places the SD-WAN event logs in a separate subsection: **SD-WAN**.

The example on this slide shows several sample SD-WAN event logs. Click a log to get further details of the event. For example, the details for log ID 10 indicate that the status of **port2** changed from **dead** to **alive**, which is the reason that log ID 9 indicates that the port is ready to start forwarding traffic.

DO NOT REPRINT
© FORTINET

SD-WAN Events (Contd)

- Enable health check SLA status logs on FortiGate
  - For performance SLAs configured with SLA targets
  - Required by some FortiView graphs

Device Manager > Provisioning Templates > SD-WAN Templates > Performance SLA > Advanced Options

sla-fail-log-period	10
sla-id-redistribute	0
sla-pass-log-period	10

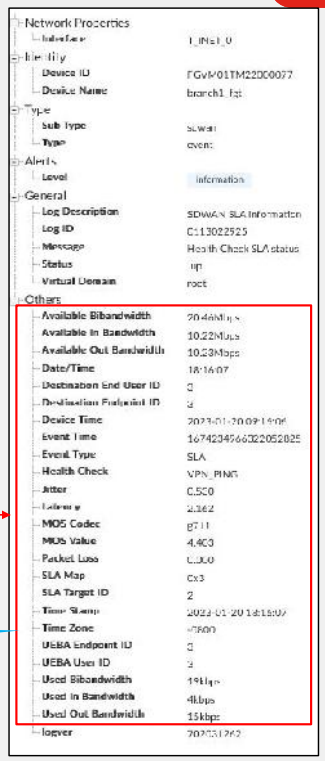
Default = 0 (no logs generated)

- Periodic health check SLA status logs in FortiAnalyzer:

Log View > Event > SD-WAN

#	▼Date/Time	Level	Device ID	Interface	Status	Message
5	18:16:12	information	FGVM01TM22000078	T_MPLS	up	Health Check SLA status.
6	18:16:12	information	FGVM01TM22000078	T_MPLS	up	Health Check SLA status.
7	18:16:07	information	FGVM01TM22000077	T_INET_0	up	Health Check SLA status.
8	18:16:07	information	FGVM01TM22000077	T_INET_0	up	Health Check SLA status.

Detailed member health and performance status



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

The analytics and reporting features of FortiAnalyzer are mainly based on the information contained in the logs received from devices. The data used by some of the SD-WAN FortiView graphs in FortiAnalyzer is taken from the health check SLA status logs generated by FortiGate. By default, these logs are not generated, and you must enable them if you want to take advantage of all the SD-WAN monitoring and reporting features provided by FortiAnalyzer.

On FortiGate, you enable the health check SLA status logs by configuring the `sla-fail-log-period` and `sla-pass-log-period` settings available on the performance SLA CLI configuration. On FortiManager, you can configure these settings under **Advanced Options** in the **Performance SLA** section.

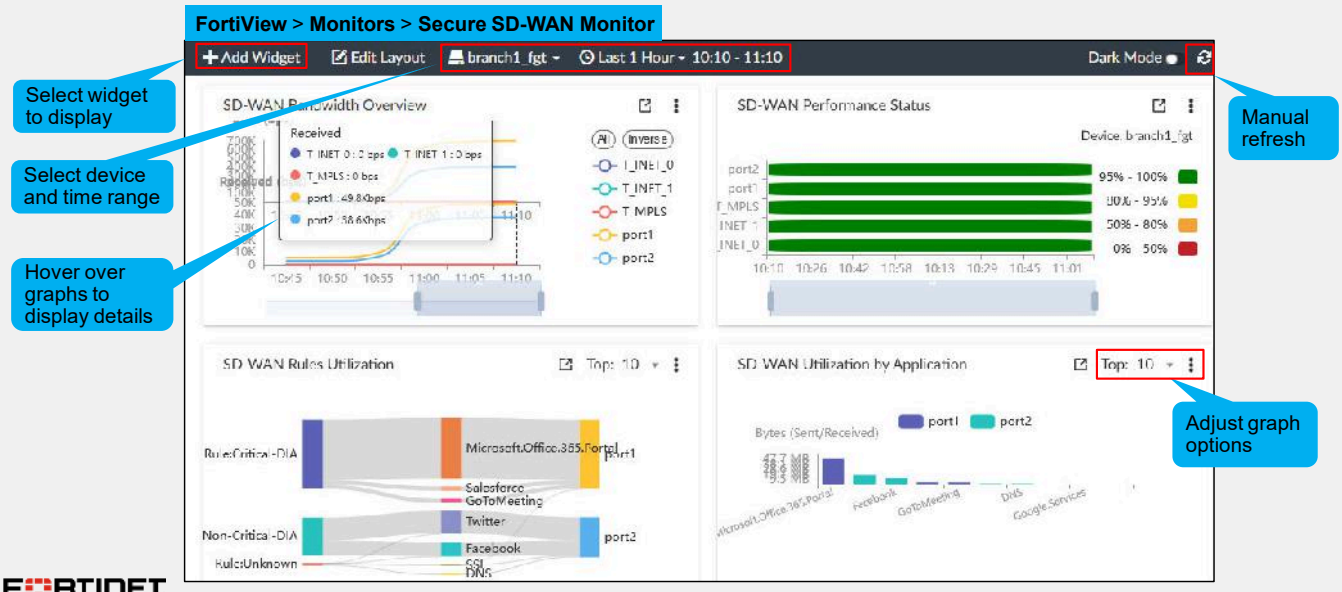
`sla-fail-log-period` and `sla-pass-log-period` indicate the interval (in seconds) that health check SLA status log messages are generated when a member doesn't meet its configured SLA target and when it does, respectively. By default, both settings are set to `0`, which means that no logs are generated.

In the example shown on this slide, both settings are set to `10` seconds. As a result, FortiAnalyzer receives health check SLA status logs every 10 seconds for each SD-WAN member configured with an SLA target. Health check SLA status logs contain detailed health and performance information about members.

DO NOT REPRINT
© FORTINET

Secure SD-WAN Monitor

- Consolidated view of real-time and historical member health and utilization per device



FortiView is a comprehensive monitoring system for your network that integrates real-time and historical data into a single view.

The **Secure SD-WAN Monitor** page on FortiView displays member health and utilization information for a particular device and time range using practical graphs. The example on this slide shows some of the graphs available on the page.

The **SD-WAN Rules Utilization** widget displays a Sankey type graph. It reports the rules utilization per applications and members. The example on this slide shows that the **Non-Critical-DIA** rule is used to forward **Twitter** and **Facebook** applications, and that the destination interface is **port2**. Hover over a graph to see details.

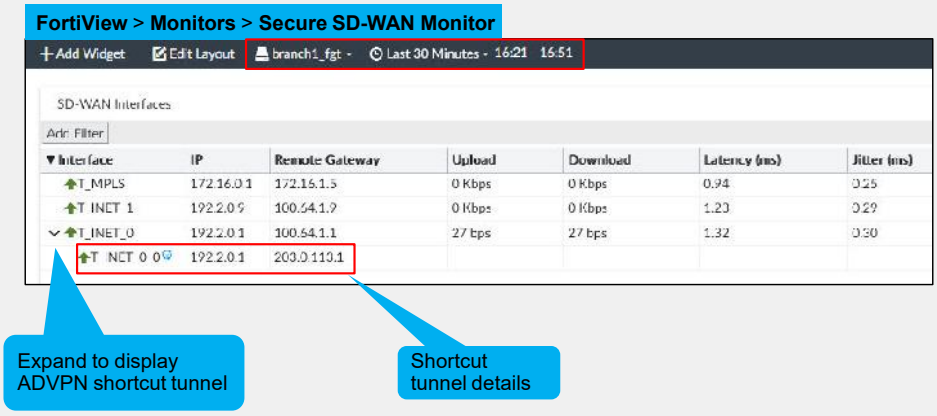
The **SD-WAN Utilization by Application** widget indicates the amount of traffic per application, and the members that the traffic was steered to. The example on this slide shows that all the **Facebook** traffic was steered to **port2**.

Note that the graphs are automatically refreshed every 5 minutes. However, you can manually refresh them by clicking the refresh button on the upper-right corner of the page.

DO NOT REPRINT
© FORTINET

Secure SD-WAN Monitor (Contd)

- ADVPN shortcut information



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

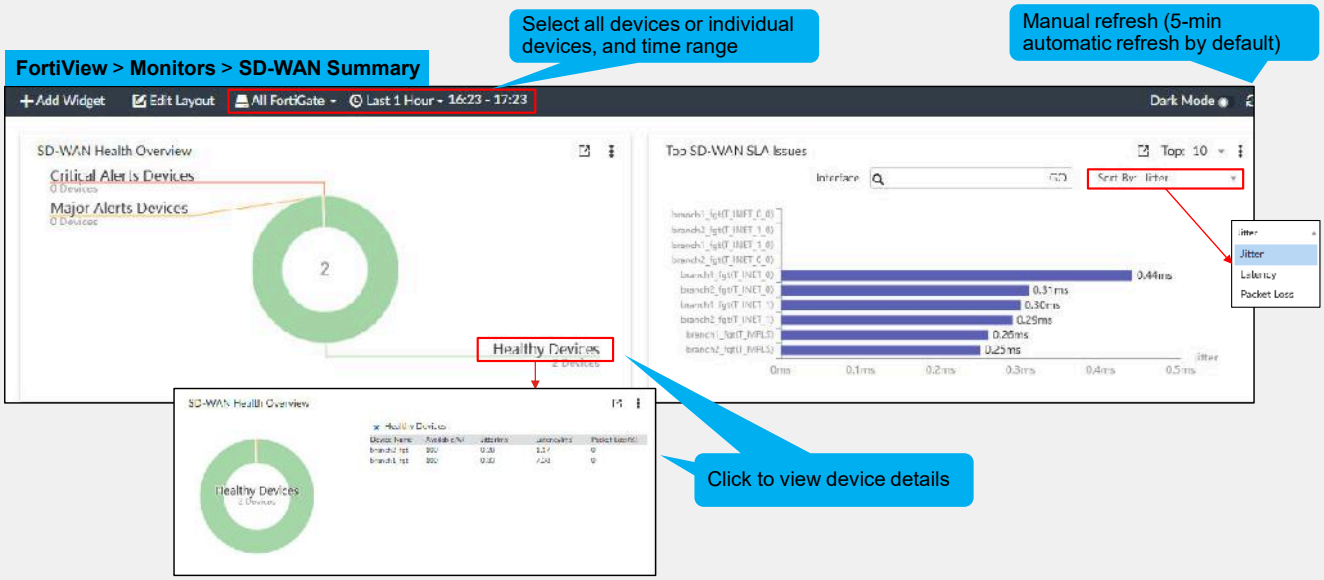
On this slide you can see an example of the SD-WAN interface widget, also available on the secure SD-WAN monitor page. It displays the bandwidth and performance information for SD-WAN interfaces. For some tunnel interfaces, you will see an expand button near the interface name, which indicates that one or more ADVPN shortcuts are established from this tunnel. Expand the display to see details for ADVPN shortcut tunnels.

Note that for VPN tunnels, the **IP** and **Remote Gateway** columns display the local IP and remote gateway IP of the VPN tunnel.

DO NOT REPRINT
© FORTINET

SD-WAN Summary

- Summarized view of member health and utilization



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

The **SD-WAN Summary** page on FortiView provides a summarized view of your SD-WAN deployment. Like the secure SD-WAN monitor page, it provides SD-WAN health and utilization information except that it can aggregate the data from multiple devices in the network.

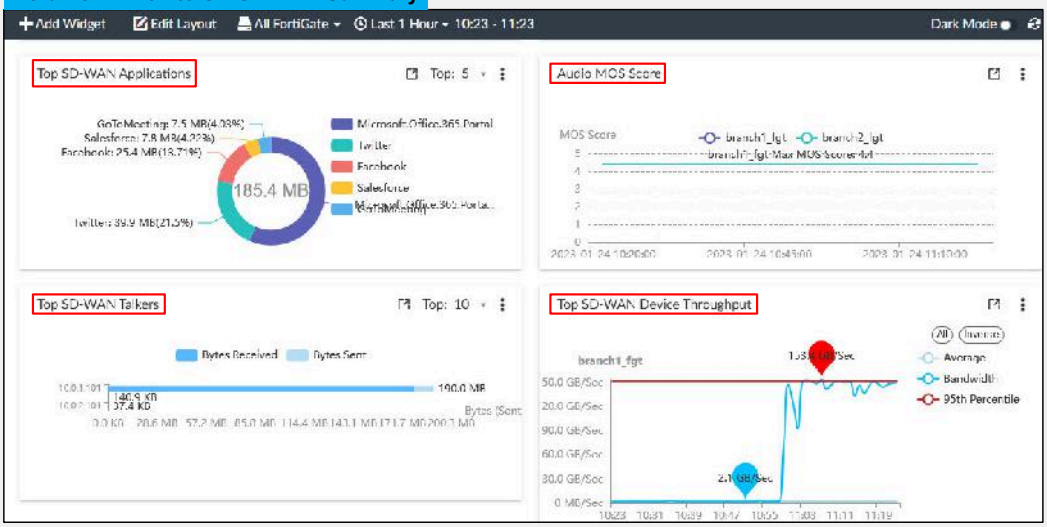
The example on this slide shows the **SD-WAN Health Overview** and **Top SD-WAN SLA Issues** widgets. The former indicates how many devices are in critical, major, and healthy states. You can also click the graph links to get the list of devices in each of those states. The **Top SD-WAN SLA Issues** widget displays the best performance interfaces based on jitter, latency, or packet loss. You can select the performance parameter from the field in the upper-right corner of the page.

DO NOT REPRINT
© FORTINET

SD-WAN Summary (Contd)

- Other widgets:

FortiView > Monitors > SD-WAN Summary



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

This slide shows additional widgets available on the **SD-WAN Summary** page.

The **Top SD-WAN Applications** and **Top SD-WAN Talkers** widgets display the applications and endpoints that generated the most traffic in SD-WAN, respectively.

The **Audio MOS Score** widget graphs MOS per codec and per FortiGate.

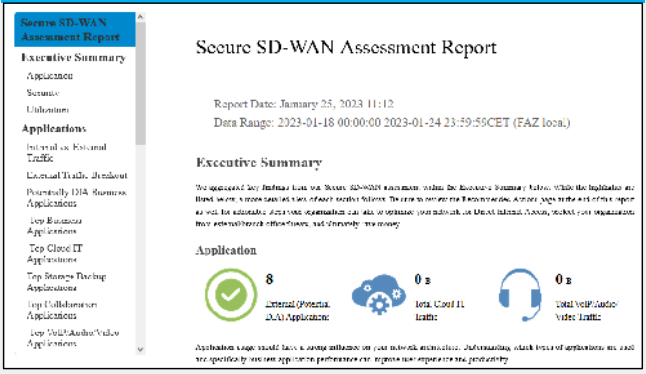
The **Top SD-WAN Device Throughput** widget graphs the bandwidth utilization for SD-WAN traffic.

DO NOT REPRINT
© FORTINET

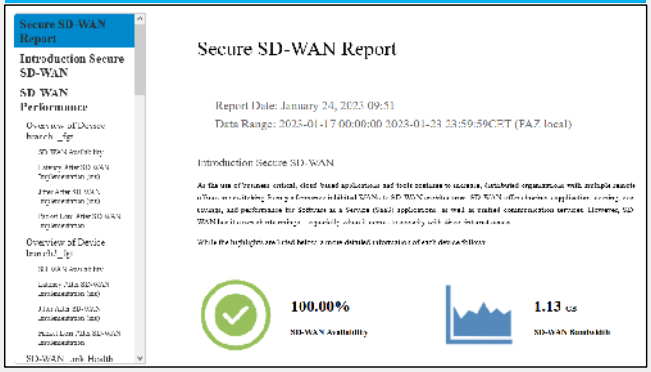
Built-In SD-WAN Reports

- Secure SD-WAN Assessment
  - For pre-SD-WAN deployments
  - Describes traffic pattern and how SD-WAN can help to improve application performance
- Secure SD-WAN
  - For post-SD-WAN deployments
  - Lists applications, and indicates status of devices and members

Reports > Reports Definitions > All Reports > Network Reports > Secure SD-WAN Assessment Report



Reports > Reports Definitions > All Reports > Network Reports > Secure SD-WAN Report



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

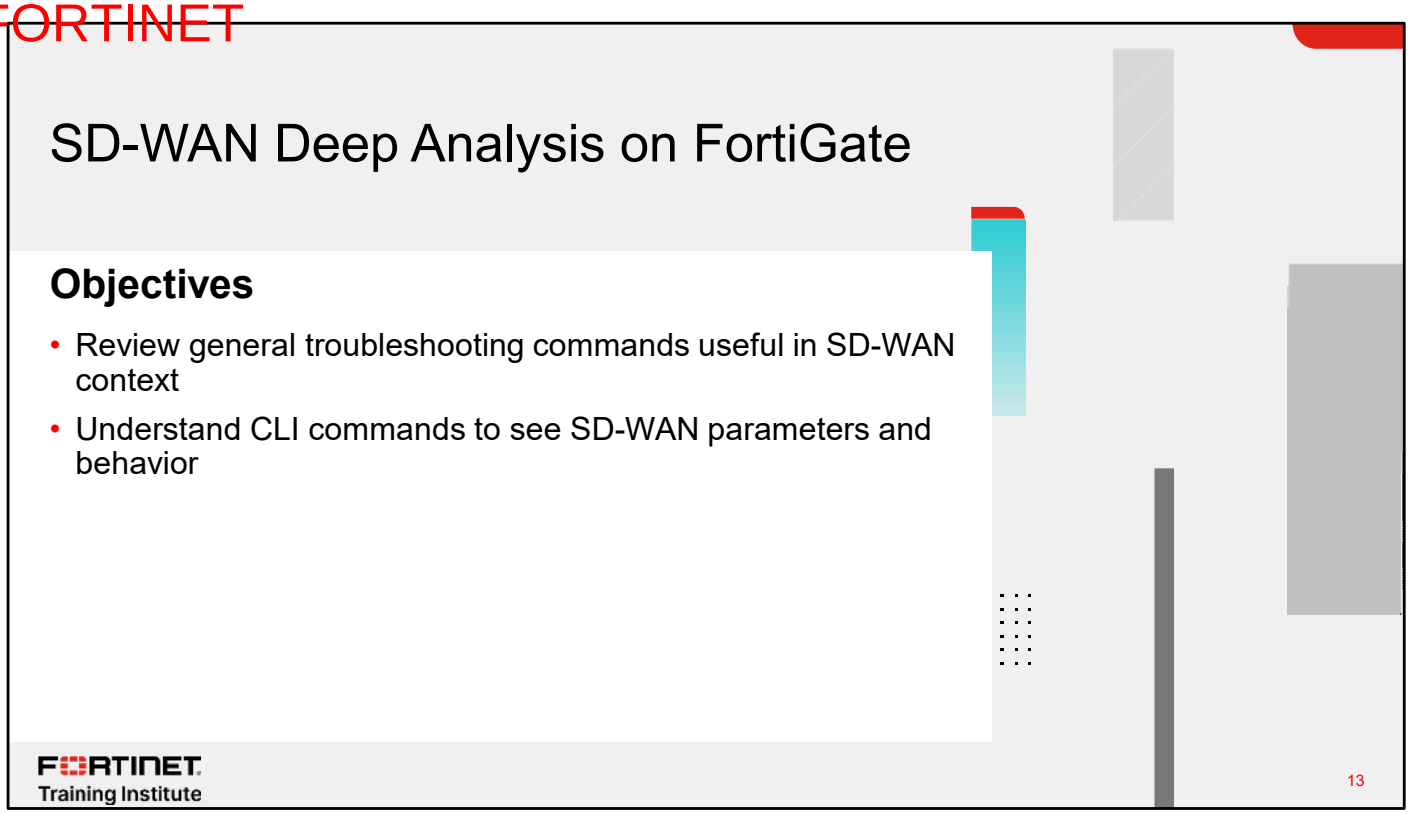
12

FortiAnalyzer comes with two preconfigured SD-WAN reports: the **Secure SD-WAN Assessment Report** and **Secure SD-WAN Report**.

The **Secure SD-WAN Assessment Report** is intended to be run before you deploy SD-WAN on the network. The report describes how much traffic is considered external or internal, the top applications in the network, and their category (business, cloud IT, social media, and so on). The report also indicates the top sources and destination based on traffic volume. The goal of the report is for administrators and management to understand the current traffic pattern, and then identify the areas where SD-WAN can help improve application performance and availability.

The **Secure SD-WAN Report** is meant to be run after you deploy SD-WAN. The report lists the applications steered by SD-WAN, and the health and performance information of devices and members. Administrators and management can use the report to check the status of SD-WAN during a given period and identify past issues.

DO NOT REPRINT
© FORTINET



After completing this section, you should be able to achieve the objectives shown on this slide. By demonstrating a good understanding of FortiGate troubleshooting commands you will be able to analyze and troubleshoot SD-WAN deployments.

DO NOT REPRINT
© FORTINET

General-Purpose Troubleshooting Commands

- General-purpose troubleshooting and monitoring commands useful in an SD-WAN context
  - Collect traffic received and sent by the FortiGate device
    - `diagnose sniffer packet`
  - View how FortiGate handles a packet and decision taken
    - `diagnose debug flow`
  - Display established sessions
    - `get system session list`
    - `diagnose system session list`
  - Route lookup commands
    - `diagnose firewall proute list`
    - `get router info routing-table all`
    - `get router info kernel`
    - `get router info bgp network`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

When troubleshooting SD-WAN on FortiGate, you can continue to use the general-purpose troubleshooting commands.

The main commands used in an SD-WAN context are:
- `diagnose sniffer packet`, to collect PCAP traces of traffic through FortiGate
- `diagnose debug flow`, for a detailed view of how FortiGate analyzes a packet and makes forwarding decisions
- `get system session list` and `diagnose sys session list`, to get a list of, and details about, established sessions
- Routing-related commands, and especially `diagnose firewall proute list`

On the next few slides, you will see SD-WAN-specific fields available for those commands.

Note that this lesson doesn't present the general principles and detailed usage for those commands. You can refer to the troubleshooting section of the *FortiGate Administration Guide* for additional details about commands usage.

DO NOT REPRINT
© FORTINET

Sniffer Capture

- Sniffer on the FortiGate CLI
  - Interface *any* or a specified interface (no zone)
  - Use trace level 4 to identify the outgoing interface selected
  - Disable traffic offloading to collect sniffer trace (when applicable)

Select any as interface to detect outgoing interface member

Filter on host, protocol,…

Trace level 4 to display header of packets with interface name

```
# diagnose sniffer packet any "host 10.0.2.101 and icmp" 4
2023-01-04 14:10:36.360797 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:36.386674 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.406708 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.406789 T_MPLS out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.408639 T_MPLS in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:37.408669 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:38.407970 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

On an SD-WAN deployment, you can collect a sniffer trace on the FortiGate CLI to examine the traffic flowing through the FortiGate device. If you collect a sniffer trace to determine the outgoing interface selected for a particular traffic flow, you will use the following settings:

- Trace level 4 to display the interface name but only the header of packet—for better readability
- Interface *any* to collect data on all possible outgoing interfaces
- Traffic filter to limit the capture to the flow you want to analyze

Remember that for FortiGate with NP2, NP4, or NP6 interfaces that are offloading traffic, you must disable offloading at the policy level before you collect a sniffer trace, otherwise the PCAP process cannot collect the traffic.

DO NOT REPRINT
© FORTINET

Debug Flow

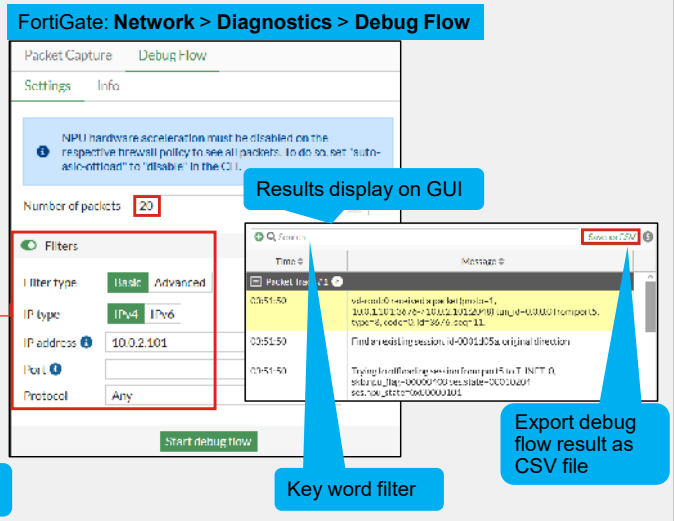
- Debug flow with CLI commands

```
# diagnose debug enable
# diagnose debug flow filter addr 10.0.2.101
# diagnose debug flow trace start 20
id=20085 trace_id=3 func=print_pkt_detail line=5783
msg="vd-root:0 received a packet(proto=1,
10.0.1.101:8605->10.0.2.101:2048) from port5. type=8,
code=0, id=8605, seq=1."
id=20085 trace_id=3 func=init_ip_session_common
line=5955 msg="allocate a new session-00000209"
id=20085 trace_id=3 func=rpdb_srv_match_input line=1030
msg="Match policy routing id=1: to 10.0.2.101 via
ifindex-21"
id=20085 trace_id=3 func=vf_ip_route_input_common
line=2605 msg="find a route: flag=04000000 gw-
172.16.1.5 via T_MPLS_0"
id=20085 trace_id=3 func=fw_forward_handler line=869
msg="Allowed by Policy-2:"
id=20085 trace_id=3 func=ipsecdev_hard_start_xmit
line=625 msg="enter IPSec interface T_MPLS_0"
...
```

Number of entries to capture

Set filters to limit the capture

- Debug flow on GUI



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

The debug flow indicates, in detail, how packets are handled by the FortiGate device. You can collect the data with CLI commands, or, starting from FortiOS 7.2.0, from the GUI as shown on this slide.

When debug flow is collected from the GUI, you can filter entries per key words and export the output as CSV file.

On the CLI, you can use the command `diagnose debug flow filter ?` to list available filter criteria and the command `diagnose debug flow filter` to review the filters in place.

DO NOT REPRINT
© FORTINET

Session List

- CLI commands
  - diag sys session filter
  - diag sys session list
  - diag sys session6 list
- SD-WAN information for the session
  - sdwan_mbr_seq
  - sdwan_service_id
  - None if traffic match default SD-WAN rule
  - None if not an SD-WAN session

```
# diagnose sys session list
session info: proto=6 proto_state=11 duration=5
expire=3596 timeout=3600 flags=00000000 socktype=0
sockport=0 av_idx=0 use=3
origin-shaper= [...]
ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr npu f00 f02 app_valid
statistic(bytes/packets/allow_err): org=3427/20/1
reply=4749/23/1 tuples=2
tx speed(Bps/kbps): 638/5 rx speed(Bps/kbps): 884/7
orgin->sink: org pre->post, reply pre->post dev=7-
>20/20->7 gwy=100.64.1.9/10.0.1.101
hook=pre dir=org act=noop 10.0.1.101:43020-
>10.1.0.7:22(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.0.7:22-
>10.0.1.101:43020(0.0.0.0:0)
misc=0 policy_id=3 auth_info=0 chk_client_info=0 vd=0
serial=000b2f2d tos=ff/ff app_list=2002 app=16060
url_cat=0
sdwan_mbr_seq=4 sdwan_service_id=2
rpdb_link_id=ff000002 ngfwid=n/a
npu_state=0x001008
```

Firewall policy ID

Member and rule IDs

App ID (SSH)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

The CLI command `diagnose sys session filter` allows you to filter the sessions to display. Then, use the command `diagnose sys session list` to display the session detail.

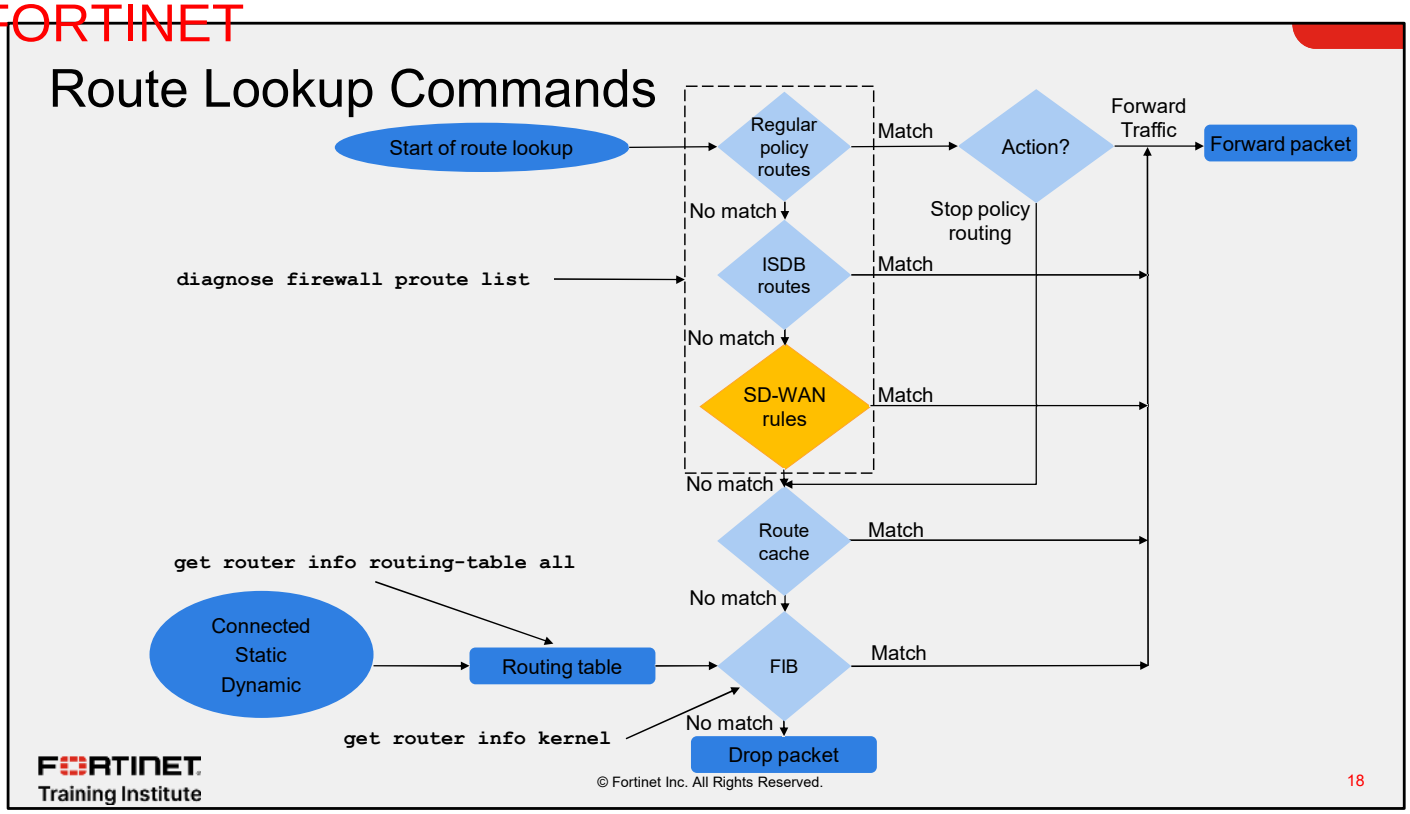
You can use `diagnose sys session filter ?` to view available filters, `diagnose sys session filter` to see active filters, and `diagnose sys session filter clear` to reset the filters settings. Use the command `diagnose sys session list` for IPv4 traffic, and `diagnose sys session6 list` for IPv6 traffic.

The right part of this slide shows an example output with detailed information about the session table entry. Note that the example does not show the traffic shaping information.

From left to right, and from top to bottom, the following information is highlighted:

- The IP protocol number and the protocol state
- The length of time until the session expires (if no more traffic matches the session)
- Session flags
- Received and transmitted packet and byte counters
- The original and reply direction of the traffic. If the device is doing NAT, this portion shows the type of NAT (source or destination) for each traffic direction, and the NAT IP address.
- The ID of the matching policy
- The SD-WAN-specific session information. `sdwan_mbr_seq` and `sdwan_service_id` indicate the SD-WAN member ID and the SD-WAN rule ID in use, respectively. If the session matched the SD-WAN implicit rule, and therefore was handled using standard FIB routing, those SD-WAN fields do not appear.

DO NOT REPRINT
© FORTINET



You already discovered the flowchart on this slide in the routing and session lesson. You can use it to review the commands available to consult routing information on the FortiGate devices.

Note that the command `diagnose firewall proute list` displays all policy routes. You can determine the type of route by looking at the route ID.

A regular policy route will get an ID between 1 and 65535. ISDB routes and routes derived from SD-WAN rules will have an ID above 65535. You will see the SD-WAN service ID field only for SD-WAN rules (field `vwl_service`).

DO NOT REPRINT
© FORTINET

Route Lookup Commands (Contd)

- SD-WAN fields in proute list

```
# diagnose firewall proute list
list route policy info(vf=root):

id=2131034113(0x7f050001) vwl_service=1(Critical-DIA) vwl_mbr_seq=1 2 dscp_tag=0xfc despite flags=0x0 tos=0x00
tos_mask=0x00 protocol=0 sport=0-65535 iif=0(any) dport=1-65535 path(2) oif=3(port1) oif=4(port2)
source(1): 10.0.1.0-10.0.1.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(3): GoToMeeting(4294836842,0,0,0,0, 16354) Microsoft.Office.365.Portal(4294837313,0,0,0,0,
41468) Salesforce(4294837785,0,0,0,0, 16920)
hit_count=34219 last_used=2023-01-24 04:04:15

id=2131034115(0x7f050003) vwl_service=3(Corp) vwl_mbr_seq=3 4 5 dscp_tag=0xfc last_used flags=0x0 tos=0x00
tos_mask=0x00 protocol=0 sport=0-65535 iif=0(any) dport=1-65535 path(3) oif=19(T_INET_0) oif=20(T_INET_1)
oif=21(T_MPLS)
source(1): 10.0.1.0-10.0.1.255
destination(1): 10.0.0.0-10.255.255.255
hit_count=13 last_used=2023-01-23 11:31:42
```

SD-WAN rule ID and rule name

SD-WAN members by order of preference

Outgoing interface by order of preference

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

This slide shows an example of a policy route list output.

Note the fields `vwl_service` and `vwl_mbr`, which indicate the SD-WAN rule that allowed the route creation and the SD-WAN member used to steer the traffic.

The ID displayed in the `diagnose firewall proute list` command output corresponds to the ID displayed in the debug flow output when a packet matches a rule. The output also includes the outgoing interface list, with the interface preference sorted from left to right.

For troubleshooting purposes, the output of the `diagnose firewall proute list` command also displays the rule hit count and the last time the rule was hit.

DO NOT REPRINT
© FORTINET

SD-WAN-Specific Commands

- Review SD-WAN overall settings
  - `get sys sdwan`
- Display SD-WAN members and zones configuration and status
  - `diagnose sys sdwan member`
  - `diagnose sys sdwan zone`
- Review health-check and link monitor status
  - `diagnose sys sdwan health-check status`
  - `diagnose sys link-monitor interface <interface_name>`
  - `diagnose sys link-monitor-passive admin list`
- Review rules status
  - `diagnose sys sdwan service`
  - `diagnose sys sdwan internet-service-app-ctrl-list`
- SLA logging
  - Activate: `sla-fail-log-period, sla-pass-log-period`
  - View performance SLA link status: `diagnose sys sdwan sla-log`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

This slide shows the main diagnostic commands that you will use when reviewing SD-WAN behavior on a FortiGate device.

Through the next few slides, you will learn how to use them and see example outputs for each of them.

DO NOT REPRINT
© FORTINET

SD-WAN Members and Zones

- Summary view of SD-WAN members and zones configuration

```
# diagnose sys sdwan member
Member(1): interface: port1, flags=0x0 , gateway: 192.2.0.2, priority: 0 1024, weight: 0
Member(2): interface: port2, flags=0x0 , gateway: 192.2.0.10, priority: 0 1024, weight: 0
Member(3): interface: T_INET_0, flags=0xc , gateway: 100.64.1.1, priority: 0 1024, weight: 0

# diagnose sys sdwan zone
Zone overlay index=3
        members(1): 25(T_INET_0)
Zone underlay index=2
        members(2): 3(port1) 4(port2)
Zone virtual-wan-link index=1
        members(0):

# diagnose netlink interface list | grep index=25
if=T_INET_0 family=00 type=768 index=25 mtu=1420 link=0 master=0
```

Configuration index number

Automatic gateway detection—tunnel ID for IPsec tunnels (`diagnose vpn tunnel list`)

For volume-based load balancing

Kernel interface index number

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

To check the SD-WAN members and zones configuration, you can use the CLI commands:

- `diagnose sys sdwan member`
- `diagnose sys sdwan zone`

Note that FortiGate adjusts the output of the command `diagnose sys sdwan member` according to the load-balance mode you configured for the rule. For instance, when you use *measured-volume-based,* the command will show the volume ratio, last reading time, and the remaining volume room.

The command `diagnose sys sdwan zone` indicates the interface kernel ID near the interface name. You can review the interface kernel IDs with the command `diagnose netlink interface list`.

DO NOT REPRINT
© FORTINET

SD-WAN Health-Check Status

- Performance SLA:

```
# diagnose sys sdwan health-check status
Health Check(Level3_DNS):
Seq(1 port1): state(alive), packet-loss(0.000%) latency(22.129), jitter(0.201), mos(4.393), bandwidth-up(10235),
bandwidth-dw(10235), bandwidth-bi(20470) sla_map=0x0
Seq(2 port2): state(alive), packet-loss(7.000%) latency(42.394), jitter(0.912), mos(4.378), bandwidth-up(10236),
bandwidth-dw(10237), bandwidth-bi(20473) sla_map=0x0
Health Check(VPN_PING):
Seq(5 T_MPLS): state(alive), packet-loss(0.000%) latency(131.336), jitter(0.199), mos(4.330), bandwidth-
up(9999999), bandwidth-dw(9999999), bandwidth-bi(19999998) sla_map=0x2
Seq(4 T_INET_1): state(alive), packet-loss(11.000%) latency(1.465), jitter(0.245), mos(4.398), bandwidth-
up(10239), bandwidth-dw(10239), bandwidth-bi(20478) sla_map=0x1
Seq(3 T_INET_0): state(alive), packet-loss(0.000%) latency(1.440), jitter(0.226), mos(4.403), bandwidth-
up(10239), bandwidth-dw(10239), bandwidth-bi(20478) sla_map=0x3
```

Configuration index number

SLA result

- Link monitor status:

```
# diagnose sys link-monitor interface port1
Interface(port1): state(up, since Tue Nov 8 08:22:49 2022), bandwidth(up:79743bps, down:1094502bps), session
count(IPv4:412, IPv6:0), tx(39720018 bytes), rx(202184681 bytes), latency(21.96), jitter(0.17), packet-
loss(0.00).
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

To check the SD-WAN health-check status and SLA performance you can use the command `diagnose sys sdwan health-check status`. This command gives, for each configured health check, the status of each member, alive or dead, and instant value for available criteria (packet-loss, latency, jitter, MOS, and bandwidth). The `sla_map` summarizes the SLA targets missed or passed. You will see on the next slide how to interpret the value presented.

Note that the health-check command displays a value for each available parameter, but the SLA process considers only the parameters for which you configured a target SLA value to determine the link SLA status and SLA map result.

The performance SLAs rely on the FortiOS link monitor process (`lnkmt`) to monitor the state and performance of members. For this reason, you can run the `diagnose sys link-monitor interface` command to display similar member status information as the one displayed by the `diagnose sys sdwan health-check status` command.

In passive or prefer passive mode, the link monitor passive process (`lnkmt_passive`) is responsible for gathering the data and preparing the reports used by the link monitor process to calculate packet loss, latency, and jitter. For this reason, you can run the `diagnose sys link-monitor-passive admin list` command to display the passive data collected.

DO NOT REPRINT
© FORTINET

SLA Target Maps

- `sla_map` uses a bitmask to reference SLA targets and indicates their status
  - The `sla_map` value represents the hexadecimal value of a binary number
  - Number of bits = number of configured SLA targets
  - First configured SLA target is assigned bit 0, second configured SLA target bit 1, and so on
  - Bit of SLA target is set to 1 if met, otherwise to 0
- Possible `sla_map` values for three SLA targets:

sla_map (hex)	SLA Target #3 (Bit 2) Value = $2^2 = 4$	SLA Target #2 (Bit 1) Value = $2^1 = 2$	SLA Target #1 (Bit 0) Value = $2^0 = 1$
0x7	Pass (1)	Pass (1)	Pass (1)
0x6	Pass (1)	Pass (1)	Fail (0)
0x5	Pass (1)	Fail (0)	Pass (1)
0x4	Pass (1)	Fail (0)	Fail (0)
0x3	Fail (0)	Pass (1)	Pass (1)
0x2	Fail (0)	Pass (1)	Fail (0)
0x1	Fail (0)	Fail (0)	Pass (1)
0x0*	Fail (0)	Fail (0)	Fail (0)

* Also means that there are no SLA targets configured

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

The `sla_map` field included in the output of the `diagnose sys sdwan health-check status` command indicates whether the member meets the configured SLA targets or not. The field uses a bitmask representation to reference the SLA targets and their status. Follow these rules to understand the `sla_map` field:

- The `sla_map` value represents the hexadecimal value of a binary number.
- The number of bits in the binary number equals the number of configured SLA targets.
- The first configured SLA target is assigned bit 0, the second configured SLA target bit 1, and so on.
- If the member meets the SLA target, the bit of the SLA target is set to 1, otherwise it is set to 0.

This slide shows a table with all the possible `sla_map` values when you configure three SLA targets for a member. For example, an `sla_map` of `0x6` means that SLA targets 3 and 2 are met, but not SLA target 1 (6 = 4 + 2 + 0). Expand or reduce the table according to the number of SLA targets configured in your setup.

Note that an `sla_map` of `0x0` has two possible meanings. If you configure one or more SLA targets for a member, `0x0` means that none of the SLA targets are met. However, `0x0` can also mean that you didn't configure any SLA targets for the member. Therefore, make sure to check the performance SLA configuration to identify whether there are SLA targets configured or not.

DO NOT REPRINT
© FORTINET

SD-WAN Link Monitor for Passive Health Check

- Link monitor passive status measured by kernel (local device measurement):

```
# diagnose sys link-monitor-passive admin list
  port1( 3) | service=0x00000000 | latency=40.0    05:04:04 | jitter=50.0    05:06:56 | pktloss=0.0  % 05:06:56
  port1( 3) | service=0x00003fe2 | latency=40.0    04:49:26 | jitter=40.0    04:49:32 | pktloss=1.8  % 04:49:32
  port2( 4) | service=0x00000000 | latency=40.0    04:49:18 | jitter=60.0    04:49:19 | pktloss=0.0  % 04:49:19
```

Kernel interface index

- Ordered by interface

```
# diagnose sys link-monitor-passive admin list by-interface

Interface port1 (3):
  Microsoft.Office.365.Portal(0x0000a1fc): latency=40.0    03:14:46, jitter=20.0    03:14:47,
pktloss=10.7 % 03:14:47
    Salesforce(0x00004218): latency=10.0    03:14:45, jitter=40.0    03:14:50, pktloss=3.5  % 03:14:50
   GoToMeeting(0x00003fe2): latency=30.0    03:14:46, jitter=30.0    03:14:51, pktloss=1.2  % 03:14:51
       Default(0x00000000): latency=50.0    03:14:04, jitter=0.0     NA      , pktloss=0.0  % NA

Interface port2 (4):
       Default(0x00000000): latency=30.0    03:14:42, jitter=60.0    03:14:43, pktloss=0.1  % 03:14:43
      Facebook(0x00003dd8): latency=40.0    03:14:49, jitter=40.0    03:14:51, pktloss=0.0  % 03:14:51
       Twitter(0x00003e81): latency=40.0    03:14:51, jitter=20.0    03:14:52, pktloss=1.3  % 03:14:52
    Salesforce(0x00004218): latency=10.0    03:14:51, jitter=60.0    03:14:52, pktloss=2.9  % 03:14:52
   GoToMeeting(0x00003fe2): latency=10.0    03:14:51, jitter=30.0    03:14:35, pktloss=0.8  % 03:14:35
```

Applications monitored on interface

Monitored parameters per applications and monitored time

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

In passive or prefer passive mode, the link monitor passive process (`lnkmt_passive`) is responsible for gathering the data and preparing the reports used by the link monitor process to calculate packet loss, latency, and jitter. When the local device performs the link monitoring—done at the kernel level—you can run the `diagnose sys link-monitor-passive admin list` command to display the passive data collected.

This command comes with various display and filtering modes:

- Overall view: `diagnose sys link-monitor-passive admin list`
- Detailed view of services monitored per interface, for all interfaces (no filter), or for a specified interface: `diagnose sys link-monitor-passive admin list by-interface [<interface name>]`
- Detailed view ordered by monitored application: `diagnose sys link-monitor-passive admin list by-application [<application ID or application name>]`

For a configuration in which FortiGate uses the SLA report measured by a remote peer device, you will use the command `diagnose sys link-monitor-passive remote list` to check SLA information.

Remember that when you want to use passive monitoring, you must enable passive measurement at the policy level with the command `set passive-wan-health-measurement enable`. Because SD-WAN passive measurement is not supported by NPU, FortiGate will automatically disable `auto-asic-offload` for a policy that has a passive health measurement.

DO NOT REPRINT
© FORTINET

SD-WAN Link Monitor for Passive Health-Check (Contd)

- Link monitor passive status:
  - Ordered by application:

```
# diagnose sys link-monitor-passive admin list by-application
Service: Unmatched (0x00000000)
         port1( 3): latency=0.0     03:31:30, jitter=50.0    03:21:56, pktloss=0.0  % 03:21:56

Service: Facebook (0x00003dd8)
         port2( 4): latency=30.0    03:31:57, jitter=40.0    03:31:55, pktloss=0.0  % 03:31:55
[...]

# diagnose sys link-monitor-passive admin list by-application Twitter
Service: Twitter (0x00003e81)
         port2( 4): latency=50.0    03:33:29, jitter=20.0    03:33:27, pktloss=0.6  % 03:33:27

# diagnose sys link-monitor-passive admin list by-application 41468
Service: Microsoft.Office.365.Portal (0x0000a1fc)
         port1( 3): latency=40.0    03:34:25, jitter=20.0    03:34:25, pktloss=10.1 % 03:34:25
         port2( 4): latency=30.0    03:19:09, jitter=20.0    03:19:10, pktloss=11.7 % 03:19:10
```

Filter by application name

Filter by application ID or hex. ID

  - Application ID mapping:

```
#    # diag sys link-monitor-passive admin app-id-map
     0: vd=0 4294836715 (0xfffe01eb) -> 15832 (0x00003dd8) Facebook
     1: vd=0 4294836842 (0xfffe026a) -> 16354 (0x00003fe2) GoToMeeting
     2: vd=0 4294837313 (0xfffe0441) -> 41468 (0x0000a1fc) Microsoft.Office.365.Portal
```

Application ID | Application Hex. ID | Application name

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

This slide shows an example of a passive link monitor output ordered by applications.

Note that you can decide to display the output for all applications or for only one. When you filter by application, you can filter on application name, application ID, or application hexadecimal value.

The command `diag sys link-monitor-passive admin app-id-map` gives you the mapping between application name, application ID, and hexadecimal ID.

DO NOT REPRINT
© FORTINET

SD-WAN Rules

- Rule status

```
# diagnose sys sdwan service

Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
 Tie break: cfg
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members(2):
    1: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(0), local cost(0), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x1), gid(0), cfg_order(1), local cost(0), selected
  Internet Service(1): SSH(4294837953,0,0,0,0 16060)
  Src address(1):
        10.0.1.0-10.0.1.255
```

Matching criteria and rule mode

Outgoing interface list

```
# diagnose sys sdwan service 1

Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service disabled caused by no destination.
  Members(2):
    1: Seq_num(4 T_INET_1_0), alive, selected
    2: Seq_num(5 T_MPLS_0), alive, selected
  Src address(1):
        10.0.1.0-10.0.1.255
```

Warning message

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

SD-WAN rules are dynamically updated based on the member status and performance. This means that the outgoing interface list can change over time, which is why it is useful to check their current status for troubleshooting purposes.

The best way to check the status of an SD-WAN rule is by running `diagnose sys sdwan service` on the FortiGate CLI. As shown on this slide, the output indicates the matching criteria in use, the rule mode, and the outgoing interface list. You can specify the service ID to display output for a single service (`diagnose sys sdwan service <service_id>`).

When you check SD-WAN service status details, pay attention to the warning messages that might appear above the member list.

Some of the warning messages you can see are:

- `service disabled caused by no destination` (no route to destination in the routing table)
- `service disabled caused by no outgoing path` (all outgoing interfaces are down)

DO NOT REPRINT
© FORTINET

Viewing the ISDB Application Cache

- View the ISDB application cache entries on the FortiGate CLI:

```
# diagnose sys sdwan internet-service-app-ctrl-list

Twitter(16001 4294838042): 104.244.42.65 6 443 Mon Nov 21 03:24:12 2022
SSH(16060 4294837753): 10.1.0.7 6 22 Mon Nov 21 02:44:18 2022
SSH(16060 4294837753): 10.1.0.28 6 22 Mon Nov 21 02:48:27 2022
```

Multiple 3-tuple for SSH app

App ID | ISDB app ID | 3-tuple | Timestamp (8-hour lifetime); one application per 3T (most recent detected application is written to cache)

- SD-WAN rule correlation:

```
# diagnose sys sdwan service 2

Service(2): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members(2):
    1: Seq_num(4 T_INET_1), alive, sla(0x1), gid(0), cfg_order(1), cost(0), selected
    2: Seq_num(3 T_INET_0), alive, sla(0x0), gid(0), cfg_order(0), cost(0), selected
  Internet Service(1): SSH(4294837753,0,0,0 16060)
  Src address(1):
        10.0.1.0-10.0.1.255
```

ISDB app ID | App ID (SSH)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

You can view the detected application entries in the ISDB application cache by running `diagnose sys sdwan internet-service-app-ctrl-list` on the FortiGate CLI. Each entry in the output indicates the application ID, the ISDB application ID, the 3-tuple, and the last time the entry was refreshed. Note that FortiGate maintains one application per 3-tuple but one application can be associated with multiple 3-tuples. If multiple applications are detected for the same 3-tuple, FortiGate uses the most recent detected application for the 3-tuple.

FortiGate uses the application ID, ISDB application ID, and 3-tuple to match the SD-WAN rule. In the example shown on this slide, one cache entry is for an SSH connection destined to `10.1.0.7` on TCP port 22. A connection that matches the cache entry also matches SD-WAN rule ID 2.

Note that entries in the ISDB application cache expire 8 hours after the last detected match.

You can run the command `diagnose sys sdwan internet-service-app-ctrl-flush` to clear all ISDB application cache entries,

DO NOT REPRINT
© FORTINET

SLA Logging

- Disabled by default
- Enable SLA fail logs and SLA pass log independently
- FortiGate stores 10 minutes of history
- Forward to FortiAnalyzer for longer history record

- Activate SLA log on CLI

```
config sys sdwan
   config health-check
      edit "VPN_PING"
          set sla-fail-log-period 10
          set sla-pass-log-period 20
      next
   end
```

Log frequency in seconds

- Activate SLA log on FortiManager SD-WAN template

Device Manager > Provisioning template > SD-WAN Templates > Performance SLA > Advanced Options



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

As already introduced in the first section of this lesson, by default, FortiGate does not log the SLA health-check result. If you want to view past SLA status, you can enable SLA fail and SLA pass log collection with the commands `set sla-fail-log-period <frequency>` and `set sla-pass-log-period <frequency>` respectively. When configured, FortiGate stores the SLA status for each measured parameter at the frequency you defined. It will keep 10 minutes of log history. For a longer retention period, you can forward the log to an external storage device like FortiAnalyzer.

DO NOT REPRINT
© FORTINET

Viewing Member SLA Logs on the FortiGate CLI

- Member metrics: `diagnose sys sdwan sla-log <SLA target> <index>`

```
# diagnose sys sdwan sla-log Level3_DNS 1
Timestamp: Wed Sep 15 00:42:07 2021, vdom root, health-check Level3_DNS, interface: port1, status: up,
latency: 34.780, jitter: 0.166, packet loss: 0.000%.
Timestamp: Wed Sep 15 00:42:07 2021, vdom root, health-check Level3_DNS, interface: port1, status: up,
latency: 34.784, jitter: 0.164, packet loss: 0.000%.
...
Timestamp: Wed Sep 15 00:52:07 2021, vdom root, health-check Level3_DNS, interface: port1, status: up,
latency: 34.766, jitter: 0.170, packet loss: 0.000%.
```

Member configuration index number

10-min period

- Member utilization: `diagnose sys sdwan intf-sla-log <interface>`

```
# diagnose sys sdwan intf-sla-log port1
Timestamp: Wed Sep 15 00:41:54 2021, used inbandwidth: 1946bps, used outbandwidth: 1788bps, used
bibandwidth: 3734bps, tx bytes: 17460310429bytes, rx bytes: 5174715499bytes.
Timestamp: Wed Sep 15 00:42:04 2021, used inbandwidth: 2827bps, used outbandwidth: 2322bps, used
bibandwidth: 5149bps, tx bytes: 17460315966bytes, rx bytes: 5174723207bytes.
...
Timestamp: Wed Sep 15 00:51:54 2021, used inbandwidth: 3128bps, used outbandwidth: 2319bps, used
bibandwidth: 5447bps, tx bytes: 17460497678bytes, rx bytes: 5174940763bytes.
```

10-min period

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

FortiOS stores the measured metrics and utilization of members for the last 10 minutes. This historical data is used by both the FortiManager and FortiGate GUIs to build their performance SLA graphs.

You can view the stored member metrics by running the `diagnose sys sdwan sla-log` command. Note that you must indicate the name of the performance SLA followed by the member configuration index number. To display the SLA logs per interface, you run the `diagnose sys sdwan intf-sla-log` command.

DO NOT REPRINT
© FORTINET

Review

- ✓ Understand FortiAnalyzer basics
- ✓ Identify SD-WAN logs
- ✓ Describe SD-WAN analytics and reports
- ✓ Review general troubleshooting commands useful in an SD-WAN context
- ✓ Understand CLI commands used to observe SD-WAN parameters and behavior

FORTINET Training Institute

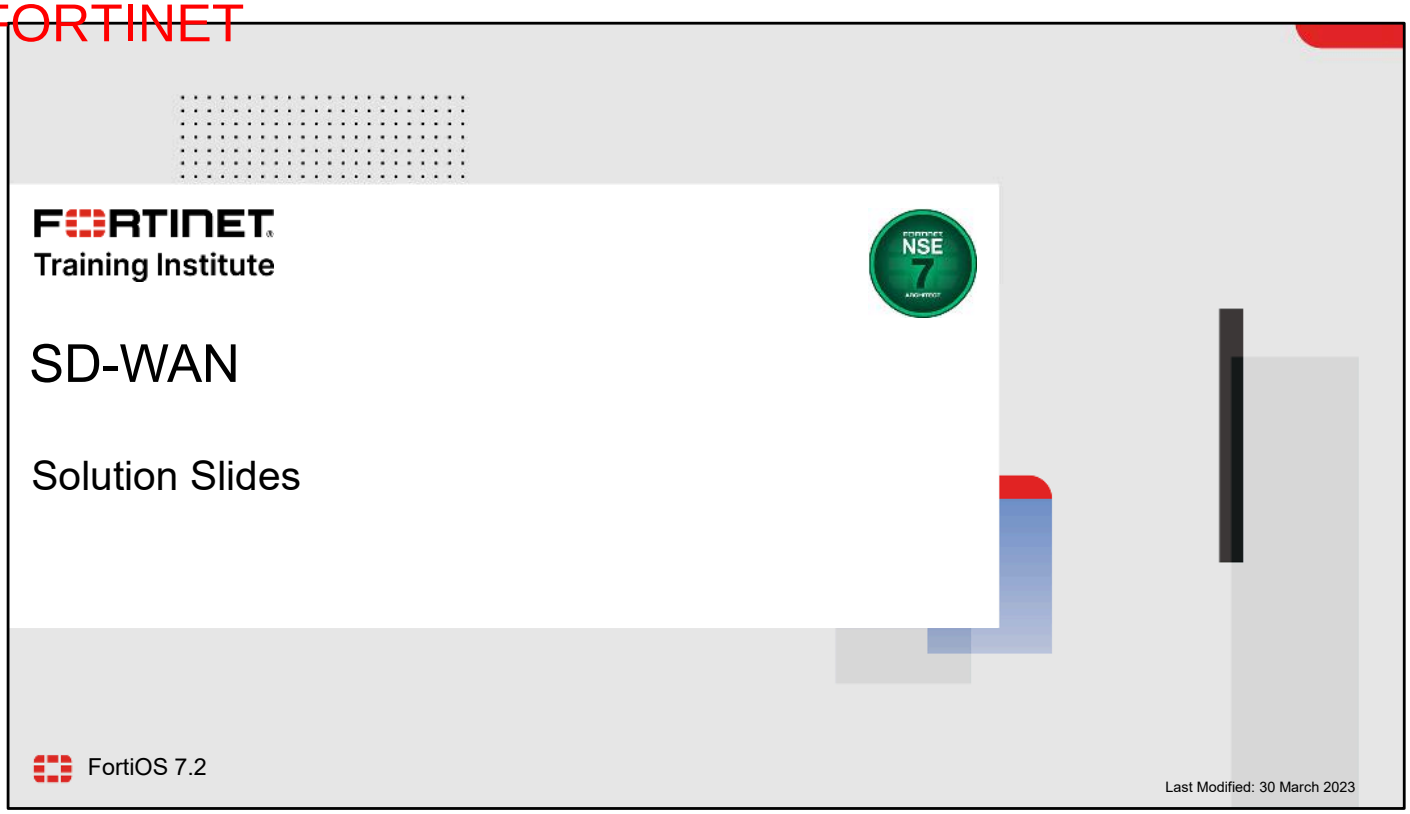
© Fortinet Inc. All Rights Reserved.

30

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

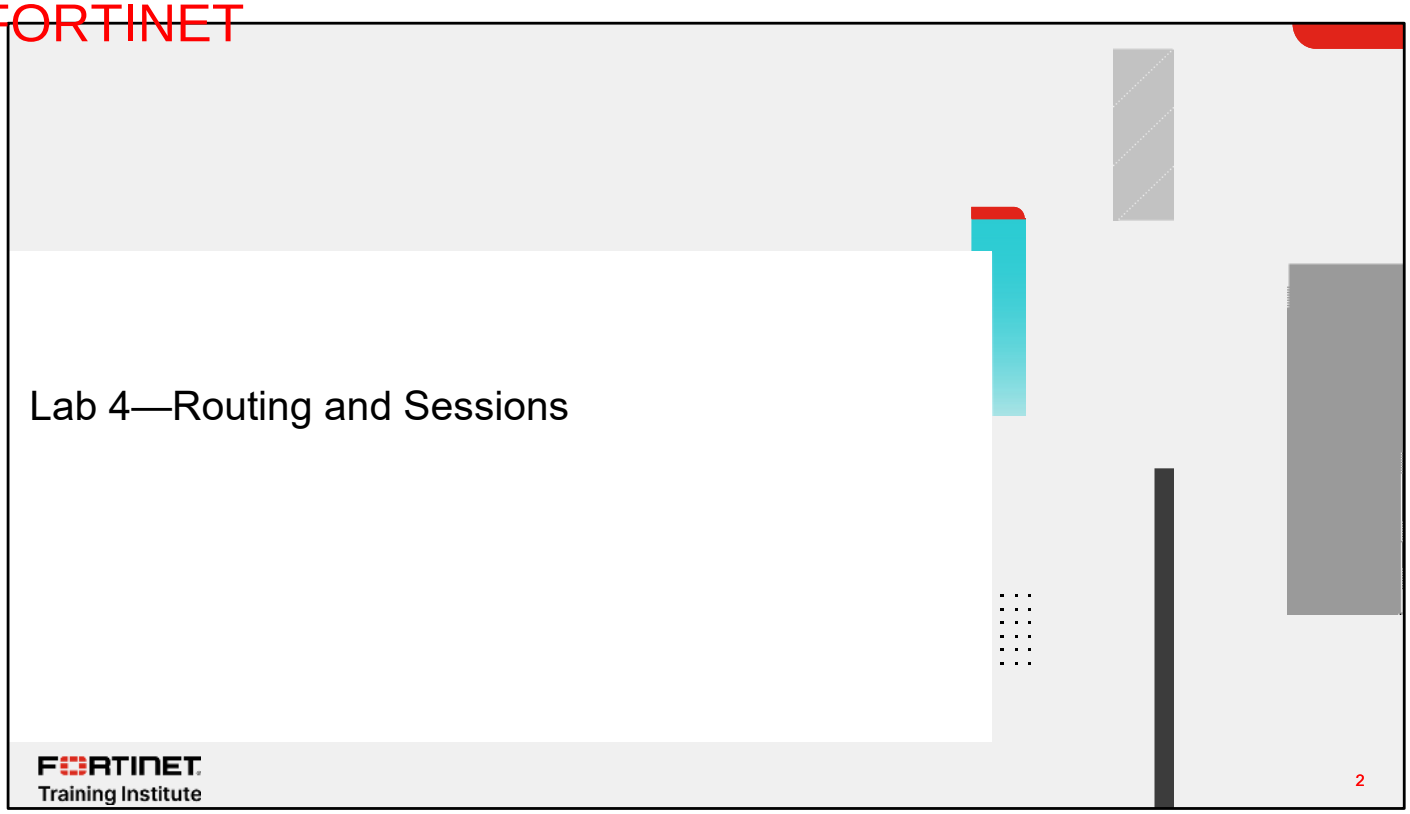
By mastering the objectives covered in this lesson, you will be able to use FortiAnalyzer to monitor your SD-WAN deployment, and use the FortiGate troubleshooting commands for a detailed view of SD-WAN behavior.

DO NOT REPRINT
© FORTINET



These slides contain the solutions to the troubleshooting exercises.

DO NOT REPRINT
© FORTINET



Now, you will review the solutions for the troubleshooting exercises in the routing and sessions lab.

DO NOT REPRINT
© FORTINET

Exercise 1—Issue Description

- Issue:
  - Traffic from branch1_client to branch2_client is always routed over T_MPLS
- Expected behavior:
  - Use T_INET_0 as primary
  - Use T_MPLS as failover (when T_INET_0 is dead)
  - After T_INET_0 recovers, fail back to T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

3

The administrator has identified the following issue:

- Traffic from branch1_client to the branch2_client is always routed over T_MPLS.

However, the administrator wants to see the following behavior:

- branch1_fgt uses T_INET_0 as the primary member to route traffic to branch2_fgt.
- branch1_fgt uses T_MPLS to route traffic to branch2_fgt, only if T_INET_0 is determined to be dead.
- After T_INET_0 recovers, spoke-to-spoke traffic is routed back through T_INET_0.

DO NOT REPRINT
© FORTINET

Exercise 1—Issue Confirmation

- Ping branch2_client (10.0.2.101) from branch1_client (10.0.1.101)
- Sniff traffic on branch1_fgt:

```
branch1_fgt # diagnose sniffer packet any "host 10.0.2.101" 4 0 l
Using Original Sniffing Mode
interfaces=[any]
filters=[host 10.0.2.101]
2023-01-04 18:47:55.886529 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 18:47:55.886600 T_MPLS out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 18:47:55.888757 T_MPLS in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 18:47:55.888781 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

- Traffic is routed over T_MPLS, but T_INET_0 is alive and the preferred member for the expected matching SD-WAN rule (ID 3):

Network > SD-WAN > Performance SLAs

VPN_PING	10.200.99.1	T_INET_0: 0.00% T_MPLS: 0.00%	T_INET_0: 1.17ms T_MPLS: 0.93ms	T_INET_0: 0.17ms T_MPLS: 0.18ms

Network > SD-WAN > SD-WAN Rules

3	Corp	LAN-net	Corp-net	T_INET_0 T_MPLS

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

You can confirm the issue by pinging 10.0.2.101 from branch1_client, while at the same time running a sniffer on the FortiGate CLI. You will see that FortiGate routes the ICMP packets over T_MPLS. However, when you check the health of T_INET_0 and the SD-WAN rules configuration on the FortiGate GUI, you will see that T_INET_0 is alive and is the preferred member for SD-WAN rule ID 3. SD-WAN rule ID 3 is the expected matching rule for spoke-to-spoke traffic.

DO NOT REPRINT
© FORTINET

Exercise 1—Troubleshooting

- Debug flow:

```
id=20085 trace_id=3 func=print_pkt_detail line=5783 msg="vd-root:0 received a packet(proto=1, 10.0.1.101:8605-
>10.0.2.101:2048) from port5. type=8, code=0, id=8605, seq=1."
id=20085 trace_id=3 func=init_ip_session_common line=5955 msg="allocate a new session-00000209"
id=20085 trace_id=3 func=rpdb_srv_match_input line=1030 msg="Match policy routing id=1: to 10.0.2.101 via ifindex-21"
id=20085 trace_id=3 func=vf_ip_route_input_common line=2605 msg="find a route: flag=04000000 gw-172.16.1.5 via T_MPLS"
id=20085 trace_id=3 func=fw_forward_handler line=869 msg="Allowed by Policy-2:"
id=20085 trace_id=3 func=ipsecdev_hard_start_xmit line=625 msg="enter IPSec interface T_MPLS"
id=20085 trace_id=3 func=_do_ipsecdev_hard_start_xmit line=229 msg="output to IPSec tunnel T_MPLS"
id=20085 trace_id=3 func=esp_output4 line=867 msg="IPsec encrypt/auth"
id=20085 trace_id=3 func=ipsec_output_finish line=546 msg="send to 172.16.0.2 via intf-port4"
...
```

- Traffic matches policy route ID 1
  - ID ≤ 65535, so policy route is a regular policy
  - Traffic should match SD-WAN rule ID 3

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

If you run a debug flow, you will notice that the traffic matches policy route ID 1. The ID is lower than 65535, which means that the policy route is a regular policy route. However, the traffic should match the SD-WAN rule ID 3 instead.

DO NOT REPRINT
© FORTINET

Exercise 1—Troubleshooting (Contd)

- Policy route table:

```
branch1_fgt # diagnose firewall proute list
list route policy info(vf=root):

id=1 dscp_tag=0xff despite flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0-0 iif=7 dport=0-65535 path(1)
oif=21(T_MPLS)
source(1): 10.0.1.0-10.0.1.255
destination(1): 10.0.0.0-10.255.255.255
hit_count=19 last_used=2023-01-04 07:40:49
...
```

- Regular policy route configuration:

Network > Policy Routes

Seq.#	Incoming Interface	Outgoing Interface	Source	Destination
1	port5	T_MPLS	LAN-net	Corp-net

- Root cause: Presence of regular policy route (precedence over other types of policy routes)
- Solution: Remove regular policy route

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

If you view the policy route table, you can see the details for regular policy route ID 1. On the FortiGate GUI, you can see the corresponding regular policy route configuration.

The presence of the regular policy route, which has precedence over other types of policy routes, including SD-WAN rules, results in traffic being sent to T_MPLS. You can solve the issue by removing the regular policy route from the configuration.

DO NOT REPRINT
© FORTINET

Exercise 1—Testing Member Selection

- Regular policy route is removed
- Ping branch2_client (`10.0.2.101`) from branch1_client (`10.0.1.101`)
- Sniff traffic on branch1_fgt:

```
2023-01-04 18:51:56.537826 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 18:51:56.537926 T_INET_1 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 18:51:56.543305 T_INET_1 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 18:51:56.543327 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

- Debug flow:

```
id=20085 trace_id=5 func=print_pkt_detail line=5783 msg="vd-root:0 received a packet(proto=1, 10.0.1.101:8608-
>10.0.2.101:2048) from port5. type=8, code=0, id=8608, seq=1."
id=20085 trace_id=5 func=init_ip_session_common line=5955 msg="allocate a new session-00000262"
id=20085 trace_id=5 func=vf_ip_route_input_common line=2605 msg="find a route: flag=04000000 gw-100.64.1.9 via T_INET_1"
id=20085 trace_id=5 func=fw_forward_handler line=869 msg="Allowed by Policy-4:"
id=20085 trace_id=5 func=ipsecdev_hard_start_xmit line=625 msg="enter IPSec interface T_INET_1"
id=20085 trace_id=5 func=_do_ipsecdev_hard_start_xmit line=229 msg="output to IPSec tunnel T_INET_1"
```

- No policy route match:
  - SD-WAN rules are skipped
  - Traffic is routed over T_INET_1 after FIB lookup (implicit rule)
  - Traffic should be routed over T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

After you remove the regular policy route, repeat the test. You will see that traffic is now routed to T_INET_1, instead of T_INET_0.

When you run a debug flow, you will see that the traffic doesn't match a policy route. Instead, the traffic is handled using standard FIB routing, and the best route resolves to T_INET_1.

DO NOT REPRINT
© FORTINET

Exercise 1—Route Lookup and SD-WAN Members

- Route lookup for `10.0.2.101`:

```
branch1_fgt # get router info routing-table details 10.0.2.101

Routing table for VRF=0
Routing entry for 10.0.2.0/24
  Known via "static", distance 10, metric 0, best
  * 100.64.1.9, via T_INET_1
```

- SD-WAN members:

T_INET_1 is not listed

```
branch1_fgt # diagnose sys sdwan member
Member(1): interface: port1, flags=0x0 , gateway: 192.2.0.2, priority: 0 1024, weight: 0
Member(2): interface: port2, flags=0x0 , gateway: 192.2.0.10, priority: 0 1024, weight: 0
Member(3): interface: T_INET_0, flags=0x4 , gateway: 100.64.1.1, priority: 0 1024, weight: 0
Member(5): interface: T_MPLS, flags=0x4 , gateway: 172.16.1.5, priority: 0 1024, weight: 0
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

When you perform a route lookup on the FortiGate CLI for `10.0.2.101`, you can confirm that T_INET_1 is the best route. If you then view the status of SD-WAN members, you will see that T_INET_1 is not included in the list.

DO NOT REPRINT
© FORTINET

Exercise 1—Checking Configured Static Routes

- Static routes configuration:

Network > Static Routes

Destination	Gateway IP	Interface	Status
0.0.0.0/0		underlay	Enabled
10.0.2.0/24	100.64.1.9	T_INET_1	Enabled
172.16.0.0/16	172.16.0.2	port4	Enabled
10.0.0.0/8		overlay	Enabled

Best route to `10.0.2.101`

- Root cause: Best route to destination isn't an SD-WAN member
  - Rule ID 3 is skipped
- Solution: Delete static route for T_INET_1
- Route lookup for `10.0.2.101`:

```
branch1_fgt # get router info routing-table details 10.0.2.101

Routing table for VRF=0
Routing entry for 10.0.0.0/8
  Known via "static", distance 1, metric 0, best
  * 172.16.1.5, via T_MPLS
```

  - Best route to destination is now an SD-WAN member

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

If you check the configured static routes, you can confirm that there is a static route configured for T_INET_1.

Because the best route to the destination isn't an SD-WAN member, by default, FortiGate skips the SD-WAN rule during the matching process. The result is that the traffic matches the SD-WAN implicit rule. That is, FortiGate performs standard FIB routing for the traffic.

You can solve the issue by removing the static route for T_INET_1. After you remove the static route and perform another router lookup for `10.0.2.101`, you will see that T_MPLS, an SD-WAN member, is now the best route.

DO NOT REPRINT
© FORTINET

Exercise 1—Testing Member Selection (Contd)

- Ping branch2_client (10.0.2.101) from branch1_client (10.0.1.101)
- Sniff traffic on branch1_fgt:

```
2023-01-04 13:24:29.262836 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 13:24:29.263035 T_MPLS out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 13:24:29.270880 T_MPLS in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 13:24:29.270924 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

- Debug flow and policy route table:

```
...
id=20085 trace_id=5 func=rpdb_srv_match_input line=1030 msg="Match policy routing id=2131427331: to 10.0.2.101 via
ifindex-21"
...
id=20085 trace_id=5 func=ipsecdev_hard_start_xmit line=625 msg="enter IPSec interface T_MPLS"
id=20085 trace_id=5 func=_do_ipsecdev_hard_start_xmit line=229 msg="output to IPSec tunnel T_MPLS"

branch1_fgt # diagnose firewall proute list
...
id=2131427331(0x7f0b0003) vwl_service=3(Corp) vwl_mbr_seq=3 5 dscp_tag=0xff valeric flags=0x0 tos=0x00 tos_mask=0x00
protocol=0 sport=0-65535 iif=0 dport=1-65535 path(2) oif=19(T_INET_0) oif=21(T_MPLS)
...
```

- SD-WAN rule ID 3 is matched, but T_MPLS is used
  - Seems that T_INET_0 is skipped

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

If you repeat the ping test, you will see that FortiGate routes the traffic over T_MPLS. If you run a debug flow, you will see that the traffic matches a policy route. The policy route corresponds to the SD-WAN rule ID 3, based on the policy route table output.

Even though the traffic now matches the expected SD-WAN rule (ID 3), FortiGate routes the traffic to T_MPLS, instead of T_INET_0. Therefore, it seems that FortiGate skips T_INET_0 during the rule lookup process.

DO NOT REPRINT
© FORTINET

Exercise 1—Routing Table and Static Routes

- Routing table:

```
branch1_fgt # get router info routing-table all
...
Routing table for VRF=0
S*      0.0.0.0/0 [1/0] via 192.2.0.2, port1
                  [1/0] via 192.2.0.10, port2
S       10.0.0.0/8 [1/0] via T_MPLS tunnel 172.16.1.5
C       10.0.1.0/24 is directly connected, port5
S       172.16.0.0/16 [10/0] via 172.16.0.2, port4
C       172.16.0.0/29 is directly connected, port4
C       192.2.0.0/29 is directly connected, port1
C       192.2.0.8/29 is directly connected, port2
C       192.168.0.0/24 is directly connected, port10
```

No routes over T_INET_0

- Static routes configuration:

Network > Static Routes

Destination	Gateway IP	Interface	Status
10.0.0.0/8		overlay	Enabled
172.16.0.0/16	172.16.0.2	port4	Enabled
0.0.0.0/0		underlay	Enabled

Static route references the **overlay** zone

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

If you view the routing table using the FortiGate CLI, you will see that there are no routes to `10.0.2.101` over T_INET_0. In fact, there are no routes over T_INET_0 at all. If you check the configured static routes, you will see that the route that matches the destination references the overlay zone.

DO NOT REPRINT
© FORTINET

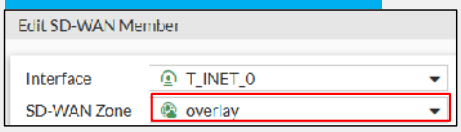
Exercise 1—SD-WAN Zones

- SD-WAN zones:

```
branch1_fgt # diagnose sys sdwan zone
Zone SASE index=2
         members(0):
Zone overlay index=4
         members(1): 21(T_MPLS)
Zone underlay index=3
         members(2): 3(port1) 4(port2)
Zone virtual-wan-link index=1
         members(1): 19(T_INET_0)
```

- Root cause: T_INET_0 doesn't have a valid route to the destination
  - T_INET_0 is not a member of the `overlay` zone
  - T_INET_0 is skipped
- Solution: Move T_INET_0 to the **overlay** zone

Network > SD-WAN > SD-WAN Zones



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

If you check the SD-WAN zone diagnostic command, you will see that T_INET_0 is a member of the virtual-wan-link zone, instead of the overlay zone.

Because the static route references the overlay zone, but T_INET_0 is not a member of that zone, then FortiOS doesn't add the corresponding static route for T_INET_0. Then, because T_INET_0 doesn't have a valid route to the destination, FortiGate skips the member during the rule lookup process.

You can solve this issue by placing T_INET_0 in the overlay zone.

DO NOT REPRINT
© FORTINET

Exercise 1—Testing Member Selection (Contd)

- Routing table:

```
branch1_fgt # get router info routing-table all
...
S        10.0.0.0/8 [1/0] via T_INET_0 tunnel 100.64.1.1
                    [1/0] via T_MPLS tunnel 172.16.1.5
...
```

Valid route to destination over T_INET_0

- Ping branch2_client (10.0.2.101) from branch1_client (10.0.1.101)
- Sniff traffic on branch1_fgt:

```
2023-01-04 14:06:59.332990 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:06:59.333146 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:06:59.341167 T_INET_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:06:59.341196 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

  - Traffic is now routed to T_INET_0

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

After placing T_INET_0 in the overlay zone, you will see that the routing table now contains a route to the destination over T_INET_0. If you repeat the test, you will see that the traffic is now routed over T_INET_0.

DO NOT REPRINT
© FORTINET

Exercise 1—Testing Failover

- Run extended ping branch2_client (10.0.2.101) from branch1_client (10.0.1.101)
- Sniff traffic on branch1_fgt:
- port1 (BR1-ISP) is up, check sniffer:

```
2023-01-04 14:10:22.334962 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:22.335090 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:22.342919 T_INET_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:22.342967 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

  - Traffic is routed to T_INET_0 (expected)
- Bring down port1 (BR1-ISP) using the WAN simulator and check the sniffer:

```
2023-01-04 14:10:36.386674 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.406708 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.406789 T_MPLS out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:37.408639 T_MPLS in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:37.408669 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:38.407970 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
```

  - Traffic fails over to T_MPLS (expected)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

You can test failover by running an extended ping to the destination while sniffing traffic on the FortiGate CLI. After you bring down port1 using the WAN simulator, you will see that traffic is now being routed over T_MPLS.

DO NOT REPRINT
© FORTINET

Exercise 1—Testing Failback

- Bring port1 back up (BR1-ISP) using the WAN simulator and check the sniffer:

```
2023-01-04 14:10:46.418326 T_MPLS out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:46.419584 T_MPLS in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:46.419602 port5 out 10.0.2.101 -> 10.0.1.101: icmp: echo reply
2023-01-04 14:10:47.419802 port5 in 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:47.419858 T_INET_0 out 10.0.1.101 -> 10.0.2.101: icmp: echo request
2023-01-04 14:10:47.421946 T_INET_0 in 10.0.2.101 -> 10.0.1.101: icmp: echo reply
```

  - Traffic fails back to T_INET_0 (expected)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

You can test fallback by bringing port1 back up, and then looking at the sniffer output. You will see that traffic is being routed over T_INET_0 again.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Description

- Issue 1:
  - If both port1 and port2 are alive, non-critical internet traffic is load balanced across port1 and port2
- Expected behavior:
  - Port2 is used only if port1 goes down

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

The administrator has identified this issue: if both port1 and port2 are alive, and the administrator generates internet traffic, critical traffic is routed to port1, which is expected. However, non-critical traffic is load balanced across port1 and port2.

The administrator wants port2 to be used for internet traffic, only if port1 goes down.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Confirmation

- Ping `8.8.8.8`, `8.8.4.4`, `4.2.2.3` and `4.2.2.4` from branch1_client (`10.0.1.101`)
- Sniff traffic on branch1_fgt, and notice that some traffic is routed over port2:

```
2023-01-04 14:45:41.128669 port5 in 10.0.1.101 -> 4.2.2.3: icmp: echo request
2023-01-04 14:45:41.129065 port2 out 192.2.0.9 -> 4.2.2.3: icmp: echo request
2023-01-04 14:45:41.131183 port2 in 4.2.2.3 -> 192.2.0.9: icmp: echo reply
2023-01-04 14:45:41.131259 port5 out 4.2.2.3 -> 10.0.1.101: icmp: echo reply
```

- But port1 and port2 are alive, and the preferred member for non-critical traffic is port1:

Network > SD-WAN > Performance SLAs

Level3 DNS	4.2.2.1 4.2.2.2	port1: 0.00% port2: 0.00%	port1: 1.81ms port2: 1.83ms	port1: 0.12ms port2: 0.12ms

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

You can confirm the issue by pinging the addresses shown on this slide from branch1_client, while at the same time running a sniffer on the FortiGate CLI.

You will see that, for some destinations, FortiGate routes ICMP packets over port2. However, when you check the health of port1 and port2, both members are alive; therefore, FortiGate should route all internet traffic over port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Troubleshooting

- Session details:

```
branch1_fgt # diagnose sys session list

session info: proto=1 proto_state=00 duration=13 expire=46 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0 use=4
...
orgin->sink: org pre->post, reply pre->post dev=7->4/4->7 gwy=192.2.0.10/10.0.1.101
hook=post dir=org act=snat 10.0.1.101:1252->4.2.2.3:8(192.2.0.9:61668)
hook=pre dir=reply act=dnat 4.2.2.3:61668->192.2.0.9:0(10.0.1.101:1252)
hook=post dir=reply act=noop 4.2.2.3:1252->10.0.1.101:0(0.0.0.0:0)
misc=0 policy_id=2 pol_uuid_idx=14721 auth_info=0 chk_client_info=0 vd=0
serial=00000355 tos=ff/ff app_list=2000 app=24466 url_cat=0
rpdb_link_id=80000000 ngfwid=n/a
...
```

No reference to SD-WAN rule.
Traffic matches implicit SD-WAN rule (standard FIB routing)

- Routing table:

```
branch1_fgt # get router info routing-table all
...
S*      0.0.0.0/0 [1/0] via 192.2.0.2, port1
                  [1/0] via 192.2.0.10, port2
...
```

Matching route is ECMP

- Root cause: Matching route is ECMP, and FortiGate load balances the traffic
- Solution: Break ECMP by increasing priority of port2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

18

If you view the session details for the internet traffic that is routed over port2, you will see that the traffic does not match explicit SD-WAN rule. This means that FortiGate performs standard FIB routing for this traffic.

If you then check the routing table, you will see that there are ECMP default routes: one for port1 and another for port2. As a result, FortiGate load balances the internet traffic that matches the implicit rule.

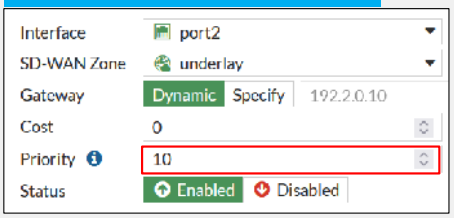
You can solve the issue by breaking the ECMP default route, by increasing the priority of port2.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Increase the Priority of port2

- Set port2 priority to 10:

Network > SD-WAN > SD-WAN Zones



- Routing table:

```
branch1_fgt # get router info routing-table all
...
S*      0.0.0.0/0 [1/0] via 192.2.0.2, port1
                  [1/0] via 192.2.0.10, port2, [10/0]
```

port1 is now preferred for standard FIB routing

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

If you increase the member priority of port2 to 10, you will see a change in the routing table. The default routes will no longer be ECMP routes, because port1 and port2 have different priorities. That is, port2 has a lower priority (higher number) than port1. The result is that FortiGate routes internet traffic over port1, because port1 has a higher priority.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Testing

- Ping the addresses again and sniff the traffic:

```
2023-01-04 15:06:50.586525 port5 in 10.0.1.101 -> 4.2.2.3: icmp: echo request
2023-01-04 15:06:50.587095 port1 out 192.2.0.1 -> 4.2.2.3: icmp: echo request
2023-01-04 15:06:50.589153 port1 in 4.2.2.3 -> 192.2.0.1: icmp: echo reply
2023-01-04 15:06:50.589238 port5 out 4.2.2.3 -> 10.0.1.101: icmp: echo reply
```

  - All internet traffic is now routed to port1 (no load balancing)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

If you repeat the test, you will see that all internet traffic, including non-critical traffic, is routed over port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Description

- Issue 2:
  - During a failover from port1 to port2 (port1 is dead), existing TCP sessions time out
  - For example, an SSH connection to the internet web server (`128.66.0.1`) from branch1_client times out during the failover
- Expected behavior:
  - Existing TCP sessions fail over to port2 successfully and to not time out

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

There is another issue to troubleshoot. During a failover from port1 to port2 (port1 is dead), existing TCP sessions time out. For example, an SSH connection to the internet web server (`128.66.0.1`) from branch1_client, times out during the failover.

The administrator wants existing TCP sessions to fail over to port2 successfully, and to not time out.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Confirmation

- Connect to `128.66.0.1` over SSH from branch1_client (`10.0.1.101`), and log in with username `root` and password `password`

```
root@branch1-client-cli:~# ssh 128.66.0.1
Warning: Permanently added '128.66.0.1' (ECDSA) to the list of known hosts.
root@128.66.0.1's password:
...
root@www-webserver-lab:~#
```

- Bring down port1 (BR1-ISP) using the WAN simulator
- Check if the SSH connection responds to user command
  - The SSH session is unresponsive

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

You can confirm the issue by connecting to `128.66.0.1` over SSH from branch1_client. You then bring down port1 using the WAN simulator. After that, you will see that the SSH connection no longer responds to user commands.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting

- Bring port1 back up
- On branch1_client, restart the SSH connection, and log in to the web server
- On branch1_fgt, check the SSH session details:

```
session info: proto=6 proto_state=11 duration=44 expire=3557 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr f00 app_valid
statistic(bytes/packets/allow_err): org=3421/20/1 reply=3777/17/1 tuples=3
tx speed(Bps/kbps): 76/0 rx speed(Bps/kbps): 84/0
orgin->sink: org pre->post, reply pre->post dev=7->3/3->7 gwy=192.2.0.2/10.0.1.101
hook=post dir=org act=snat 10.0.1.101:49262->128.66.0.1:22(192.2.0.1:49262)
hook=pre dir=reply act=dnat 128.66.0.1:22->192.2.0.1:49262(10.0.1.101:49262)
hook=post dir=reply act=noop 128.66.0.1:22->10.0.1.101:49262(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=2 pol_uuid_idx=14721 auth_info=0 chk_client_info=0 vd=0
serial=000032d9 tos=ff/ff app_list=2000 app=16060 url_cat=0
sdwan_mbr_seq=1 sdwan_service_id=2
rpdb_link_id=ff000002 ngfwid=n/a
npu_state=0x001108
```

- `may_dirty` flag is present
- Outgoing device is index `3` (port1)
- Outgoing gateway is `192.2.0.2` (port1)
- Matching firewall policy: `2`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

Bring port1 back up to repeat the test. After you connect to the web server over SSH, you can check the SSH session details on the FortiGate CLI.

You will see that the `may_dirty` flag is present, and that the outgoing device is the one with index `3`, which is port1. Although not shown on this slide, you can view the device index numbers in the output of the `diagnose netlink interface list` command.

You will also see that the outgoing gateway is `192.2.0.2`, which is the gateway used by port1, and that the matching firewall policy is `2`.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting (Contd)

- Bring down port1, and check the SSH session details:

```
session info: proto=6 proto_state=11 duration=352 expire=3249 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log dirty may_dirty ndr f00 app_valid
statistic(bytes/packets/allow_err): org=3421/20/1 reply=3777/17/1 tuples=3
tx speed(Bps/kbps): 0/0 rx speed(Bps/kbps): 0/0
orgin->sink: org pre->post, reply pre->post dev=7->3/3->7 gwy=0.0.0.0/0.0.0.0
hook=post dir=org act=snat 10.0.1.101:49262->128.66.0.1:22(192.2.0.1:49262)
hook=pre dir=reply act=dnat 128.66.0.1:22->192.2.0.1:49262(10.0.1.101:49262)
hook=post dir=reply act=noop 128.66.0.1:22->10.0.1.101:49262(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=2 pol_uuid_idx=14721 auth_info=0 chk_client_info=0 vd=0
serial=000032d9 tos=ff/ff app_list=2000 app=16060 url_cat=0
sdwan_mbr_seq=1 sdwan_service_id=2
rpdb_link_id=ff000002 rpdb_svc_id=0 ngfwid=n/a
npu_state=0x001008
```

- `dirty` flag is present
- Outgoing device is still index `3` (port1)
- Gateway information is flushed (`0.0.0.0`)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

Continue troubleshooting by bringing down port1, and then checking the SSH session details again. In the output, you will see that the `dirty` flag is present, and that the gateway information was flushed.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting (Contd)

- Set up debug flow and then check if the SSH connection responds to the user commands
  - SSH connection is unresponsive
- Observe the debug flow output:

```
id=20085 trace_id=1 func=print_pkt_detail line=5783 msg="vd-root:0 received a packet(proto=6, 10.0.1.101:49262-
>128.66.0.1:22) from port5. flag [.], seq 3783503819, ack 3805876270, win 501"
id=20085 trace_id=1 func=resolve_ip_tuple_fast line=5864 msg="Find an existing session, id-000032d9, original direction"
id=20085 trace_id=1 func=rpdb_srv_match_input line=1030 msg="Match policy routing id=2131034114: to 128.66.0.1 via
ifindex-4"
id=20085 trace_id=1 func=vf_ip_route_input_common line=2605 msg="find a route: flag=04000000 gw-192.2.0.10 via port2"
id=20085 trace_id=1 func=get_new_addr line=1230 msg="find SNAT: IP-192.2.0.9(from IPPOOL), port-49262"
id=20085 trace_id=1 func=fw_strict_dirty_session_check line=264 msg="SNAT IP 192.2.0.1 != 192.2.0.9, drop"
id=20085 trace_id=2 func=print_pkt_detail line=5783 msg="vd-root:0 received a packet(proto=6, 10.0.1.101:49262-
>128.66.0.1:22) from port5. flag [.], seq 3783503819, ack 3805876270, win 501"
id=20085 trace_id=2 func=rpdb_srv_match_input line=1030 msg="Match policy routing id=2131034114: to 128.66.0.1 via
ifindex-4"
id=20085 trace_id=2 func=vf_ip_route_input_common line=2605 msg="find a route: flag=04000000 gw-192.2.0.10 via port2"
id=20085 trace_id=2 func=fw_forward_dirty_handler line=360 msg="no session matched"
```

- FortiGate:
  - Performs route lookup for the next packet (re-evaluation, dirty flag)
  - Resolves to port2 as the new gateway, but drops the packet because its SNAT IP changed
  - Clears the session and drops subsequent packets because they don't match an existing session

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

If you set up debug flow for the test traffic, and then you enter commands on the SSH test connection, you will see the following:

- FortiGate performs a route lookup for the next packet. This is because the session is flagged as dirty, as a result of a route change. This instructs FortiOS to re-evaluate the session.
- FortiGate resolves to port2 as the new gateway. However, FortiGate drops the packet because the SNAT IP on the packet changed. It was `192.2.0.1`, and now is `192.2.0.9`.
- Because of the SNAT change, FortiGate clears the session and drops subsequent packets, because they don't match an existing session.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting (Contd)

- Check the matching firewall policy configuration:

```
branch1_fgt # show firewall policy 2
config firewall policy
    edit 2
        set name "DIA"
        set uuid 96a6df1a-729a-51ec-089b-cba0021dd4d0
        set srcintf "port5"
        set dstintf "underlay"
        set action accept
        set srcaddr "LAN-net"
        set dstaddr "all"
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set ssl-ssh-profile "certificate-inspection"
        set application-list "default"
        set logtraffic all
        set nat enable
    next
end
```

- Firewall policy references the `underlay` zone
- No IP pool, different SNAT IP per member

- Root cause: No IP pool results in a different SNAT IP per member
- Solution: Configure IP pool on firewall policy 2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

The matching firewall policy is `2`. If you check the firewall policy configuration, you will see that it references the `underlay` zone, but doesn't have IP pool enabled. This means that FortiGate uses the IP address of the selected member in the underlay zone as the SNAT IP. This explains why the SNAT IP changed during the failover.

You can solve the issue by configuring an IP pool on firewall policy 2. The goal is for FortiGate to use the same SNAT IP, regardless of the selected member.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Configure IP Pool

- View configured IP pools:

```
branch1_fgt # show firewall ippool
config firewall ippool
    edit "192.2.0.100"
        set startip 192.2.0.100
        set endip 192.2.0.100
    next
end
```

IP pool 192.2.0.100 has been preconfigured for you

- Apply IP pool on firewall policy 2:

```
config firewall policy
    edit 2
        set ippool enable
        set poolname "192.2.0.100"
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

If you view the configured IP pools, you will see that the IP pool 192.2.0.100 has been preconfigured for you. The administrator wants FortiGate to SNAT internet traffic with the 192.2.0.100 IP address, so you must apply the IP pool on the firewall policy, as shown on this slide.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Testing

- Bring port1 back up and repeat the test
- During failover, confirm that the SSH connection remains responsive, and then check the debug flow output:

```
id=20085 trace_id=59 func=print_pkt_detail line=5783 msg="vd-root:0 received a packet(proto=6,
10.0.1.101:49608->128.66.0.1:22) from port5. flag [.], seq 1331221955, ack 3504353634, win 501"
id=20085 trace_id=59 func=resolve_ip_tuple_fast line=5864 msg="Find an existing session, id-000034bd,
original direction"
id=20085 trace_id=59 func=rpdb_srv_match_input line=1030 msg="Match policy routing id=2131165186: to
128.66.0.1 via ifindex-4"
id=20085 trace_id=59 func=vf_ip_route_input_common line=2605 msg="find a route: flag=04000000 gw-
192.2.0.10 via port2"
id=20085 trace_id=59 func=get_new_addr line=1230 msg="find SNAT: IP-192.2.0.100(from IPPOOL), port-
60434"
id=20085 trace_id=59 func=ids_receive line=328 msg="send to ips"
id=20085 trace_id=59 func=__ip_session_run_tuple line=3492 msg="SNAT 10.0.1.101->192.2.0.100:49608"
```

- FortiGate now forwards the packet out (no SNAT IP change error)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

Bring port1 back up and then repeat the test. This time, you will see that the SSH connection remains responsive during the failover. Also, the debug flow no longer shows the SNAT IP address change error, which means that FortiGate is forwarding the packets successfully.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Description

- Issue 3:
  - After existing sessions fail over to port2 successfully, sessions don't fail back to port1 after port1 recovery
- Expected behavior:
  - Sessions to fail back to port1 after port1 recovery

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

There is another issue to troubleshoot. After existing sessions fail over to port2 successfully, the sessions don't fail back to port1, after port1 recovery.

The administrator wants the sessions to fail back to port1 after port1 recovery.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Confirmation

- Bring port1 back up, and repeat the test of issue 2
- Confirm the SSH connection fails over to port2 successfully
- Bring port1 back up, and then test the SSH connection while sniffing SSH traffic:

```
5.244388 port5 in 10.0.1.101.49608 -> 128.66.0.1.22: psh 1331222171 ack 3504354210
5.244739 port2 out 192.2.0.100.49608 -> 128.66.0.1.22: psh 1331222171 ack 3504354210
5.253437 port2 in 128.66.0.1.22 -> 192.2.0.100.49608: psh 3504354210 ack 1331222207
```

  - SSH connection is still responsive, but FortiGate continues to forward packets to port2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

You can confirm the issue by bringing up port1, and then repeating the test used for issue 2. You can confirm that the SSH connection fails over to port2 successfully. If you then bring port1 back up, the SSH continues to respond to user commands. However, if you sniff the SSH traffic, you will see that FortiGate continues to route the traffic to port2, instead of routing it to port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Troubleshooting

- Repeat the test
- Check the SSH session after bringing port1 back up:

```
session info: proto=6 proto_state=11 duration=21 expire=3586 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log dirty may_dirty ndr f00 app_valid
statistic(bytes/packets/allow_err): org=4077/28/1 reply=4517/22/1 tuples=3
tx speed(Bps/kbps): 0/0 rx speed(Bps/kbps): 0/0
orgin->sink: org pre->post, reply pre->post dev=7->4/4->7 gwy=192.2.0.10/10.0.1.101
hook=post dir=org act=snat 10.0.1.101:49858->128.66.0.1:22(192.2.0.100:49858)
hook=pre dir=reply act=dnat 128.66.0.1:22->192.2.0.100:49858(10.0.1.101:49858)
hook=post dir=reply act=noop 128.66.0.1:22->10.0.1.101:49858(0.0.0.0:0)
...
```

- The `dirty` flag is present
- However, the gateway info is not flushed
- Gateway address will not be updated

- Root cause: By default, FortiOS doesn't change the routing information of SNAT sessions, unless existing gateway is no longer valid
- Solution: Enable `snat-route-change` setting

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

You troubleshoot the issue by repeating the test, and then checking the SSH session details after bringing port1 back up—that is, during failback. The SSH session shows the dirty flag, but the gateway information is not flushed. That is, FortiGate does not update the gateway information for this session.

The reason is that, by default, FortiOS doesn't change the routing information of SNAT sessions, unless their existing gateway is no longer valid, which is not the case here. That is, port2, which is the device that the session uses to reach the current gateway, is still up and so are the respective routes.

You can solve this issue by enabling the `snat-route-change` setting. When you enable this setting, FortiGate always refreshes the routing information of SNAT sessions, following a routing change that impacts the sessions.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Enable `snat-route-change`

- Enable `snat-route-change`:

```
config system global
    set snat-route-change enable
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

You can enable the `snat-route-change` setting by running the commands shown on this slide.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Testing

- Repeat the test
- Check the SSH session after bringing port1 back up:

```
session info: proto=6 proto_state=11 duration=37 expire=3570 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
...
state=log dirty may_dirty ndr f00 app_valid
statistic(bytes/packets/allow_err): org=4477/35/1 reply=4925/28/1 tuples=3
tx speed(Bps/kbps): 0/0 rx speed(Bps/kbps): 0/0
orgin->sink: org pre->post, reply pre->post dev=7->4/4->7 gwy=0.0.0.0/0.0.0.0
...
```

- The `dirty` flag is present
- The gateway info is flushed
- Gateway address can be updated

- Test the SSH connection, and confirm that it remains up
- Check the SSH session again:

```
session info: proto=6 proto_state=11 duration=289 expire=3597 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
...
state=log may_dirty ndr f00 app_valid
statistic(bytes/packets/allow_err): org=5705/54/1 reply=6273/42/1 tuples=3
tx speed(Bps/kbps): 8/0 rx speed(Bps/kbps): 9/0
orgin->sink: org pre->post, reply pre->post dev=7->3/3->7 gwy=192.2.0.2/10.0.1.101
...
```

- The `dirty` flag was removed
- The outgoing device was updated (port1)
- The outgoing gateway was updated (port1)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

If you repeat the test, you will see that ForitOS now flushes the gateway information during the failback. The result is that FortiGate updates the gateway information, so it uses port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Testing (Contd)

- Test the SSH connection while sniffing SSH traffic:

```
2.715688 port5 in 10.0.1.101.50374 -> 128.66.0.1.22: psh 56438236 ack 1749555534
2.715943 port1 out 192.2.0.100.50374 -> 128.66.0.1.22: psh 56438236 ack 1749555534
2.716672 port1 in 128.66.0.1.22 -> 192.2.0.100.50374: ack 56438272
```

  - Traffic fails back to port1

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

If you also sniff the SSH traffic during the failback, you will see that FortiGate is routing packets to port1 again.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 4—Description

- Issue 4:
  - When port1 is dead, and new sessions are established through port2, the sessions don't fail over to port1 after port1 recovery, and continue using port2
- Expected behavior:
  - Sessions to fail over to port1 after port1 recovery

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

There is one last issue to troubleshoot in this exercise. When port1 is dead, and new sessions are established through port2, the sessions don't fail over to port1, after port1 recovery. That is, the sessions continue using port2.

The administrator wants these sessions to fail over to port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 4—Confirmation

- Bring down port1
- Connect to `128.66.0.1` over SSH from branch1_client (`10.0.1.101`)
- Sniff SSH traffic to confirm packets are routed over port2:

```
1.527331 port5 in 10.0.1.101.50552 -> 128.66.0.1.22: syn 1461883675
1.528100 port2 out 192.2.0.100.50552 -> 128.66.0.1.22: syn 1461883675
```

- Bring port1 back up, and then test the SSH connection while sniffing SSH traffic:

```
115.083422 port5 in 10.0.1.101.50552 -> 128.66.0.1.22: psh 1461885277 ack 2970972072
115.083928 port2 out 192.2.0.100.50552 -> 128.66.0.1.22: psh 1461885277 ack 2970972072
```

  - FortiGate continues to route packets over port2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

Confirm the issue by bringing down port1, and then connecting to the web server over SSH. Sniff the SSH traffic to confirm that packets are routed over port2. Then, bring port1 back up, and test the SSH connection. The SSH connection works, but FortiGate continues to route the packets over port2, instead of routing them again over port1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 4—Troubleshooting

- Repeat the test, and then check the SSH session before bringing port1 back up:

```
session info: proto=6 proto_state=11 duration=9 expire=3594 timeout=3600 flags=00000000 socktype=0 sockport=0 av_idx=0
use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr f00 app_valid route_preserve
...
```

- The `route_preserve` flag is present
- Interface stickiness is enabled
- Gateway address will not be updated

- Check port2 configuration:

```
branch1_fgt # show system interface port2
config system interface
    edit "port2"
...
        set preserve-session-route enable
...
```

- Root cause: Interface stickiness is enabled on port2, and gateway address is not updated unless existing gateway is no longer valid
- Solution: Disable `preserve-session-route` on port2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

37

To troubleshoot this issue, repeat the test. Check the SSH session details before bringing port1 back up (failback event). In the session, the `route_preserve` flag is present. That is, interface stickiness is enabled on the outgoing interface—port2, in this case. The result is that the gateway address will not be updated, unless the existing gateway is no longer valid, which is not the case.

If you check port2 configuration, you will see that the `preserve-session-route` setting is enabled. You can solve the issue by disabling the setting on port2.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 4—Disable Interface Stickiness

- Disable interface stickiness on port2:

```
config system interface
    edit "port2"
        set preserve-session-route disable
    next
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

38

You can disable the `preserve-session-route` setting on port2 by running the commands shown on this slide.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 4—Testing

- Repeat the test, and then check the SSH session before bringing port1 back up:

```
session info: proto=6 proto_state=11 duration=14 expire=3595 timeout=3600 flags=00000000 socktype=0
sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty ndr f00 app_valid
...
```

- `route_preserve` flag is not present

- Bring port1 back up, test the SSH connection, and sniff the traffic to confirm packets are routed over port1:

```
1.030673 port5 in 10.0.1.101.50746 -> 128.66.0.1.22: psh 931097527 ack 721063059
1.030838 port1 out 192.2.0.100.50746 -> 128.66.0.1.22: psh 931097527 ack 721063059
```

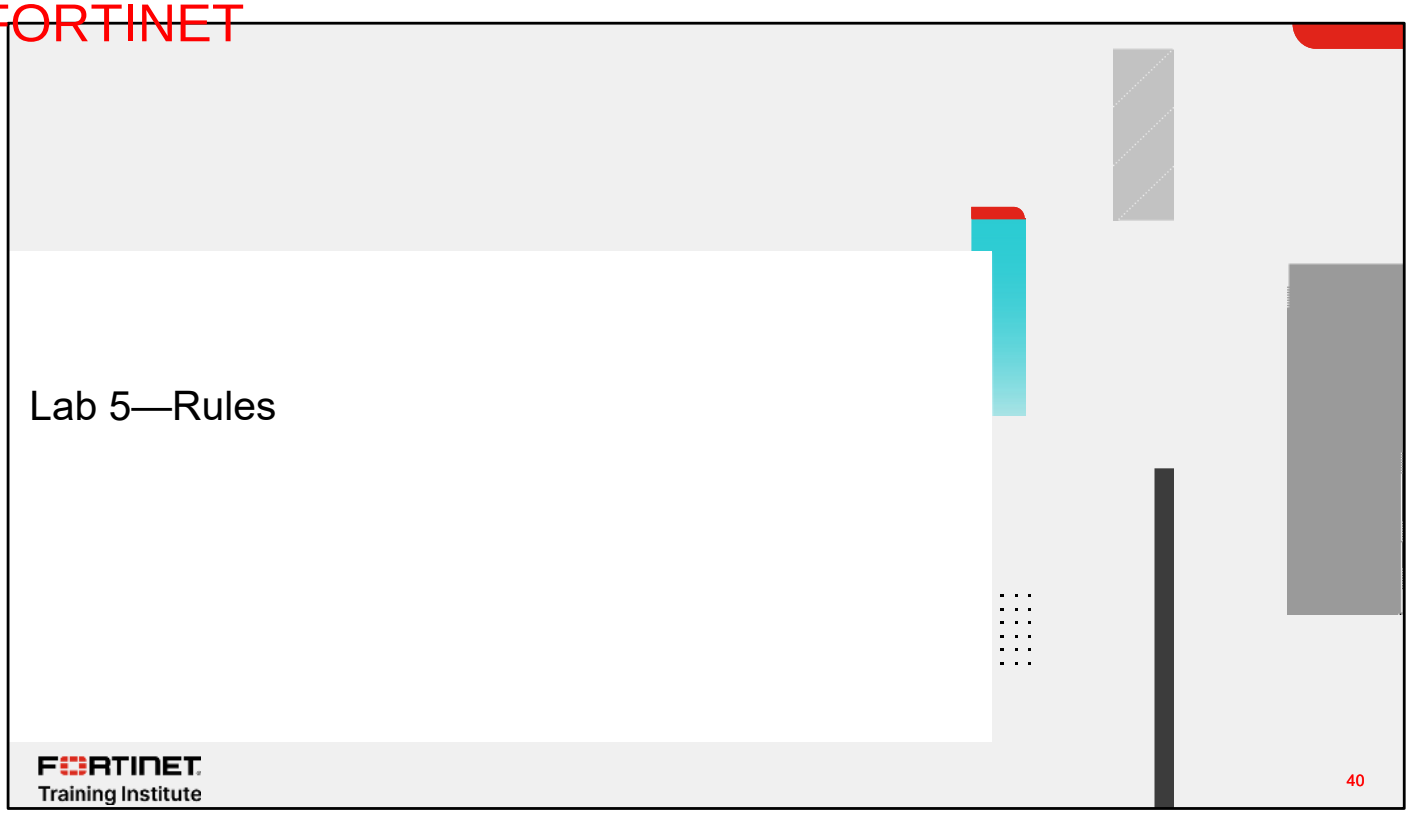
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

39

If you repeat the test, you no longer see the `route_preserve` flag in the session. The result is that FortiGate can now update the gateway information during the failover to port1; therefore, FortiGate routes the SSH traffic over port1 again.

DO NOT REPRINT
© FORTINET



Now, you will look at the solutions for the troubleshooting exercise in the rules lab.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Description

- Issue 1:
  - If port1 is up, internet traffic on branch1_fgt is routed over port1 successfully
  - However, if port1 is down, internet traffic is dropped
- Expected behavior:
  - If port1 is down, internet traffic fails over to the overlays (RIA)
    - Overlays: T_INET_1 and T_MPLS

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

41

The administrator identified this issue: if port1 is down, internet traffic sourced from branch1_client is dropped by branch1_fgt. However, the administrator wants internet traffic to fail over to the overlays to perform RIA. The overlays are T_INET_1 and T_MPLS.

Note that, if port1 is up, internet traffic on branch1_fgt is routed over port1 successfully.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Confirmation

- Ping `8.8.8.8` from branch1_client (`10.0.1.101`) and sniff the traffic:

```
2023-01-06 07:34:08.479281 port5 in 10.0.1.101 -> 8.8.8.8: icmp: echo request
2023-01-06 07:34:08.480259 port1 out 192.2.0.1 -> 8.8.8.8: icmp: echo request
2023-01-06 07:34:08.481387 port1 in 8.8.8.8 -> 192.2.0.1: icmp: echo reply
2023-01-06 07:34:08.481725 port5 out 8.8.8.8 -> 10.0.1.101: icmp: echo reply
```

  - You can ping the address, and the sniffer shows the traffic is routed over port1

- Bring down port1, and repeat the test:

```
2023-01-06 08:00:34.881429 port5 in 10.0.1.101 -> 8.8.8.8: icmp: echo request
2023-01-06 08:00:35.902681 port5 in 10.0.1.101 -> 8.8.8.8: icmp: echo request
```

  - You can't ping the address, and the branch1_client CLI shows the `Destination Net Unreachable` error
  - The sniffer shows that FortiGate drops the packets

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

You can confirm the issue by pinging `8.8.8.8` from branch1_client, while at the same time running a sniffer on the FortiGate CLI.

You will see that when port1 is up, FortiGate routes the ICMP packets over port1, which is the expected behavior. However, when port1 is down, FortiGate drops the packets, instead of routing them over the overlays. You will also see that the branch1_CLI shows the `Destination Net Unreachable` error.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Troubleshooting

- Bring down port1 and then check the routing table:

```
branch1_fgt # get router info routing-table all
...
C       10.0.1.0/24 is directly connected, port5
B       10.1.0.0/24 [200/0] via 10.202.1.254 (recursive via T_INET_1 tunnel 100.64.1.9), 00:06:15
                    [200/0] via 10.203.1.254 (recursive via T_MPLS tunnel 172.16.1.5), 00:06:15
S       10.202.1.0/24 [15/0] via T_INET_1 tunnel 100.64.1.9
C       10.202.1.1/32 is directly connected, T_INET_1
S       10.202.1.254/32 [15/0] via T_INET_1 tunnel 100.64.1.9
S       10.203.1.0/24 [15/0] via T_MPLS tunnel 172.16.1.5
C       10.203.1.1/32 is directly connected, T_MPLS
S       10.203.1.254/32 [15/0] via T_MPLS tunnel 172.16.1.5
S       100.64.1.9/32 [10/0] via 192.2.0.10, port2
S       172.16.0.0/16 [10/0] via 172.16.0.2, port4
C       172.16.0.0/29 is directly connected, port4
C       192.2.0.8/29 is directly connected, port2
C       192.168.0.0/24 is directly connected, port10
..
```

- Root cause: No routes to `8.8.8.8`, packet is dropped by FortiGate
- Solution: Enable `default` and `gateway` settings on SD-WAN rule ID 2

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

43

The destination net unreachable error message you see in the branch1_client CLI means that FortiGate doesn't have a route to the destination. As a result, FortiGate sends an ICMP destination unreachable message to branch1_client to inform the sender.

If you check the routing table on branch1_fgt when port1 is down, you can confirm that there are no routes to `8.8.8.8`.

Because you are not allowed to change the static routing configuration on the device, you must find another way to solve the issue. Another way to solve this issue is to enable the `default` and `gateway` settings on the expected matching SD-WAN rule—rule ID 2, in this case.

If you enable the `default` setting, FortiGate doesn't check whether the best route to the destination is an SD-WAN member. This enables FortiGate to evaluate the SD-WAN rules; otherwise, all the rules are skipped.

However, you must also ensure that FortiGate doesn't skip the rule members, which don't have a valid route to the destination. You can prevent members from being skipped by enabling the `gateway` setting. When this setting is enabled, FortiGate doesn't check if a member has a valid route to the destination. Instead, FortiGate uses the gateway defined or detected for the member to route the packets out.

Because the administrator wants internet traffic to match SD-WAN rule ID 2, you must enable both the `default` and `gateway` settings on the rule.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Configure SD-WAN Rule

- Enable `default` and `gateway` settings on SD-WAN rule ID 2:

```
config system sdwan
    config service
        edit 2
            set default enable
            set gateway enable
        next
    end
end
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

44

You can enable the `default` and `gateway` settings on SD-WAN rule ID 2 by running the commands shown on this slide.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 1—Testing

- Repeat the test and sniff the traffic:

```
2023-01-06 12:05:14.208419 port5 in 10.0.1.101 -> 8.8.8.8: icmp: echo request
2023-01-06 12:05:14.208540 T_INET_1 out 10.0.1.101 -> 8.8.8.8: icmp: echo request
2023-01-06 12:05:14.216421 T_INET_1 in 8.8.8.8 -> 10.0.1.101: icmp: echo reply
2023-01-06 12:05:14.216474 port5 out 8.8.8.8 -> 10.0.1.101: icmp: echo reply
```

  - Traffic is routed over the T_INET_1 overlay

- Session details:

```
session info: proto=1 proto_state=00 duration=1 expire=58 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0 use=3
...
state=log may_dirty npu f00
statistic(bytes/packets/allow_err): org=84/1/1 reply=84/1/1 tuples=2
tx speed(Bps/kbps): 61/0 rx speed(Bps/kbps): 61/0
orgin->sink: org pre->post, reply pre->post dev=7->20/20->7 gwy=100.64.1.9/10.0.1.101
hook=pre dir=org act=noop 10.0.1.101:2078->8.8.8.8:8(0.0.0.0:0)
hook=post dir=reply act=noop 8.8.8.8:2078->10.0.1.101:0(0.0.0.0:0)
misc=0 policy_id=2 pol_uuid_idx=14719 auth_info=0 chk_client_info=0 vd=0
serial=00003811 tos=ff/ff app_list=0 app=0 url_cat=0
sdwan_mbr_seq=4 sdwan_service_id=2
...
```

  - Connection matches SD-WAN rule ID 2 and member ID 4 (T_INET_1)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

45

If you repeat the test, you will see that FortiGate routes the ICMP packets over T_INET_1, which is the expected behavior when port1 is down.
Session information shows that the traffic match SD-WAN rule ID 2 and member ID 4, that correspond to T_INET_1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Description

- Issue 2:
  - SD-WAN rule ID 1 is disabled because it has no destination
  - The rule uses BGP route tags to match the rule destination
- Expected behavior:
  - The rule is enabled and matches corporate traffic

FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

There is another issue to troubleshoot. SD-WAN rule ID 1 is disabled because it has no destination. The rule uses BGP route tags to match the rule destination.

The administrator wants the rule to be enabled and match corporate traffic.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Confirmation

- SD-WAN rule ID 1 status:

```
branch1_fgt # diagnose sys sdwan service 1

Service(1): Address Mode(IPV4) flags=0x200 use-shortcut-sla
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service disabled caused by no destination.
  Members(2):
    1: Seq_num(4 T_INET_1), alive, selected
    2: Seq_num(5 T_MPLS), alive, selected
  Src address(1):
        10.0.1.0-10.0.1.255
```

  - The rule is disabled because it has no destination

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

47

You can confirm the issue by checking the status of rule ID 1. You will see that the rule is disabled because it has no destination.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting

- View the rule configuration:

```
branch1_fgt # show system sdwan
config system sdwan
...
    config service
        edit 1
            set name "Corp"
            set route-tag 10
            set src "LAN-net"
            set priority-zone "overlay"
        next
...
```

  - The destination addresses of the rule are the learned BGP prefixes assigned a route tag of `10`

- View the learned BGP prefixes that match the `65000:10` community:

```
branch1_fgt #  get router info bgp community 65000:10
...
   Network          Next Hop            Metric LocPrf Weight RouteTag Path
*>i10.1.0.0/24      10.202.1.254             0    100      0      100 i <-/1>
* i                 10.203.1.254             0    100      0      100 i <-/->
...
```

  - The route tag assigned to prefixes is `100`, but it should be `10`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

48

If you view the rule configuration, you can confirm that the rule is configured to match the destination address to the BGP prefixes assigned a route tag of `10`.

The administrator also says that the route tag is assigned to BGP prefixes matching the `65000:10` community. If you then view the learned BGP prefixes with that community, you will see that there is one prefix learned (`10.1.0.0/24`), but is assigned the route tag `100`. Therefore, there seems to be an error in the configuration.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Troubleshooting (Contd)

- View BGP neighbor configuration:

```
branch1_fgt # show router bgp
...
        edit "10.202.1.254"
            set soft-reconfiguration enable
            set interface "T_INET_1"
            set remote-as 65000
            set route-map-in "dc1-lan-rm"
            set update-source "T_INET_1"
        next
        edit "10.203.1.254"
            set soft-reconfiguration enable
            set interface "T_MPLS"
            set remote-as 65000
            set route-map-in "dc1-lan-rm"
            set update-source "T_MPLS"
        next
...
```

  - Route map `dc1-lan-rm` applied inbound

- View `dc1-lan-rm` configuration:

```
branch1_fgt # show router route-map dc1-lan-rm
config router route-map
    edit "dc1-lan-rm"
        config rule
            edit 1
                set match-community "dc1-lan-cl"
                set set-route-tag 100
            next
        end
    next
end
```

- Root cause: Incorrect route tag value
- Solution: Assign `10` as route tag

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

49

If you check the BGP configuration, you will see that the relevant BGP neighbors are assigned the `dc1-lan-rm` route map in the inbound direction. If you then check the route map configuration, you will see that it is assigned route tag `100`, and not `10`.

Therefore, you can solve the issue by updating the route map configuration to use a route tag of `10`.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Fix Route Tag in Route Map

- Update route tag to `10`, and then soft clear all BGP peerings:

```
config router route-map
    edit "dc1-lan-rm"
        config rule
            edit 1
                set set-route-tag 10
            next
        end
    next
end
# execute router clear bgp all soft
```

- View the learned BGP prefixes that match the `65000:10` community:

```
Network           Next Hop            Metric LocPrf Weight RouteTag Path
*>i10.1.0.0/24      10.202.1.254             0    100      0       10 i <-/1>
* i                 10.203.1.254             0    100      0       10 i <-/->
```

  - Prefix is now assigned a route tag of `10`

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

50

You can update the route tag by running the commands shown on this slide. You must also clear the BGP peerings, so the change takes effect. The example shown on this slide uses the soft clear command to prevent the peerings from bouncing (no reset).

If you then view the learned BGP prefixes that match the `65000:10` community again, you will see that the route tag is now `10`.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 2—Testing

- SD-WAN rule ID 1 status:

```
...
  Src address(1):
        10.0.1.0-10.0.1.255

  Route tag address(1):
        10.1.0.0-10.1.0.255
```

  - Rule is no longer disabled, and the destination address matches learned BGP prefix
- Ping `10.1.0.7` from branch1_client
  - You can ping the address
- Session details:

```
session info: proto=1 proto_state=00 duration=10 expire=49 timeout=0 flags=00000000 socktype=0 sockport=0 av_idx=0 use=3
...
orgin->sink: org pre->post, reply pre->post dev=7->20/20->7 gwy=100.64.1.9/10.0.1.101
hook=pre dir=org act=noop 10.0.1.101:2110->10.1.0.7:8(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.0.7:2110->10.0.1.101:0(0.0.0.0:0)
...
sdwan_mbr_seq=4 sdwan_service_id=1
```

  - Traffic matches SD-WAN rule ID `1` and is routed over member ID `4` (T_INET_1)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

51

If you check the status of rule ID 1 again, you will see that the rule is no longer disabled, and that the destination address matches the learned BGP prefix (`10.1.0.0/24`).

If you then ping `10.1.0.7` from branch1_client, you can ping the address. Then, if you check the session details, you will see that the traffic matches rule ID 1 and member ID 4 (T_INET_1).

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Description

- Expected behavior:
  - For traffic that matches rule ID 1, branch1_fgt prefers the member with the best route to the destination
  - Goal: Control member preference by advertising from dc1_fgt a better route over a preferred overlay
    - You can't change the rule mode

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

52

The last task is to fulfil a request by the administrator. The administrator wants branch1_fgt to prefer the member with the best route to the destination. This way, the administrator can control the member preference by advertising a better route, over its preferred overlay, from dc1_fgt, without having to change the configuration on the branch.

Remember that you can't change the rule mode, which is currently set to manual mode.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Confirmation

- SD-WAN rule ID 1 and overlay zone tiebreaker configuration:

```
branch1_fgt # config system sdwan

branch1_fgt (sdwan) # config service

branch1_fgt (service) # edit 1

branch1_fgt (1) # get | grep tie-break
tie-break           : zone

branch1_fgt (1) # end

branch1_fgt (sdwan) # config zone

branch1_fgt (zone) # edit overlay

branch1_fgt (overlay) # get | grep service-sla-tie-break
service-sla-tie-break: cfg-order
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

53

You can confirm the issue by checking the rule or zone configuration. By default, rules use the member priority as the first tiebreaker; that is, the order of the interface preference list, as the first tiebreaker in all applicable rules, which includes manual mode rules.

The setting that controls the first tiebreaker on the rule level is `tie-break`. By default, it is set to `zone`, which means that it follows the zone-level setting. The zone-level setting, `service-sla-tie-break`, is set to `cfg-order` (member priority).

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Set Tiebreaker to Best Route

- Set the tiebreaker to the best route at the rule level:

```
config system sdwan
    config service
        edit 1
            set tie-break fib-best-match
        next
    end
end
```

  - Usually, it's better to change the setting at the rule level, to minimize impact

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

54

The example on this slide shows how to change the tiebreaker to the best route at the rule level. Usually, it's better to change this setting at the rule level so that only the rule in question is impacted. If you make the change at the zone level, it will impact all the rules that reference the zone.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Testing

- Route lookup for `10.1.0.7`:

```
branch1_fgt # get router info routing-table details 10.1.0.7

Routing table for VRF=0
Routing entry for 10.1.0.0/24
  Known via "bgp", distance 200, metric 0, best
   Last update 00:20:50 ago
  * vrf 0 10.202.1.254 priority 1 (recursive via T_INET_1 tunnel 100.64.1.9)
  * vrf 0 10.203.1.254 priority 1 (recursive via T_MPLS tunnel 172.16.1.5)
```

  - Both routes are the best ones (ECMP)

- View SD-WAN rule ID 1 status:

```
...
    1: Seq_num(4 T_INET_1), alive, selected
    2: Seq_num(5 T_MPLS), alive, selected
...
```

  - T_INET_1 wins because it has a higher member priority

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

55

To test the solution, you can first perform a route lookup for the destination (`10.1.0.7`). FortiOS should report that both members have ECMP routes. This means that both members have equally good routes. Because of this, FortiOS will use the member priority as a tiebreaker to identify the preferred member, which in this case is T_INET_1.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Testing (Contd)

- Ping `10.1.0.7` from branch1_client and sniff the traffic:

```
2023-01-06 15:03:07.854051 port5 in 10.0.1.101 -> 10.1.0.7: icmp: echo request
2023-01-06 15:03:07.854168 T_INET_1 out 10.0.1.101 -> 10.1.0.7: icmp: echo request
2023-01-06 15:03:07.855574 T_INET_1 in 10.1.0.7 -> 10.0.1.101: icmp: echo reply
2023-01-06 15:03:07.855590 port5 out 10.1.0.7 -> 10.0.1.101: icmp: echo reply
```

  - Traffic is routed over T_INET_1

- Apply the following configuration on dc1_fgt to test best route as tiebreaker:

```
config router bgp
   config neighbor-group
      edit Branches_MPLS
          set route-map-out prefer-dc1_host
      next
   end
end
# execute router clear bgp all soft
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

56

If you ping `10.1.0.7` from branch1_client while sniffing the traffic, you can confirm that FortiGate routes the packets over T_INET_1.

The lab instructions indicate that you should apply the configuration shown on this slide to test the best route as a tiebreaker. After that, you must soft clear all BGP peerings so that the configuration takes effect.

DO NOT REPRINT
© FORTINET

Exercise 2—Issue 3—Testing (Contd)

- Route lookup for `10.1.0.7` and routing table:

```
branch1_fgt # get router info routing-table details 10.1.0.7
...
  * vrf 0 10.203.1.254 priority 1 (recursive via T_MPLS tunnel 172.16.1.5)

branch1_fgt # get router info routing-table all
...
B       10.1.0.0/24 [200/0] via 10.202.1.254 (recursive via T_INET_1 tunnel 100.64.1.9), 00:28:07, [1/0]
                    [200/0] via 10.203.1.254 (recursive via T_MPLS tunnel 172.16.1.5), 00:28:07, [1/0]
B       10.1.0.7/32 [200/0] via 10.203.1.254 (recursive via T_MPLS tunnel 172.16.1.5), 00:02:00, [1/0]
...
```

  - Both T_INET_1 and T_MPLS have a valid route, but best route is over T_MPLS

- Ping `10.1.0.7` from branch1_client and sniff the traffic:

```
2023-01-06 15:10:56.076849 port5 in 10.0.1.101 -> 10.1.0.7: icmp: echo request
2023-01-06 15:10:56.077000 T_MPLS out 10.0.1.101 -> 10.1.0.7: icmp: echo request
2023-01-06 15:10:56.078056 T_MPLS in 10.1.0.7 -> 10.0.1.101: icmp: echo reply
2023-01-06 15:10:56.078348 port5 out 10.1.0.7 -> 10.0.1.101: icmp: echo reply
```

  - T_MPLS is now used (best route)

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

57

After making the changes indicated on the lab instructions on dc1_fgt, you can perform another route lookup. You will now see that the best route to the destination is over T_MPLS. Note that both members have a valid route to the destination, but T_MPLS has the best one.

When you ping `10.1.0.7` again, you will see that FortiGate now routes the packets over T_MPLS, because T_MPLS has the best route.

DO NOT REPRINT
© FORTINET



No part of this publication may be reproduced in any form or by any means or used to make any derivative such as translation, transformation, or adaptation without permission from Fortinet Inc., as stipulated by the United States Copyright Act of 1976.

Copyright© 2023 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations,and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.