

The Fortinet logo, featuring the word "FORTINET" in a bold, sans-serif font. The letter "O" is stylized with a red and white grid pattern.

FORTINET®

A blurred background image of a speedometer. The needle is pointing towards the right, and the numbers 4, 5, and 6 are visible. The text "K1000rpm" is also visible.

Architecture for Enterprise

SD-WAN / SD-Branch



DEFINE / DESIGN / DEPLOY / DEMO



Table of Contents

Change Log	5
Introduction	6
Executive summary	6
Legacy WAN edge	7
Transforming the WAN edge	8
Improving inefficient routing and inferior performance	9
Fixing security gaps and bottlenecks	9
Modernizing WAN	9
Reducing risk with Secure SD-WAN	10
Supporting DX initiatives	10
Why Fortinet	10
Benefits of a controllerless-based architecture	11
Unique, unbeatable design	11
Industry leader	12
Intended Audience	12
Uses cases	13
Dynamic application steering across multiple WAN links	14
Redundant connectivity for enterprise branch	14
Reduce WAN OPEX with direct internet access	14
Secure and automated intra-site connectivity	15
Multi-cloud connectivity and cloud on-ramp	15
Application performance improvement	15
Work from anywhere	16
Solutions and technologies	17
FortiGate	17
Application identification	18
Increased performance	18
Form factor	18
Security	19
Zero touch deployment	19

API / automation	20
FortiManager	20
Single console management	21
Administrative domains	21
Centralized policy	21
Zero touch provisioning	21
Secure SD-WAN capabilities	22
Secure SD-WAN security automation	22
FortiAnalyzer	22
Security visibility	23
Administrative domains	23
Automatic security and SD-WAN reports	23
Security automation	24
FortiAP	24
Important terms for FortiAP	26
FortiSwitch	26
Important terms for FortiSwitch	29
Secure SD-WAN solution	30
Technical background	30
SD-WAN configuration	30
SD-WAN routing logic	33
Design principles	33
Underlay	35
Overlay	35
Routing	36
Security	37
SD-WAN	38
Architecture and design	40
Single datacenter (active-passive gateway)	41
Gateway components	41
SD-WAN considerations	50
Security considerations	51
Multiple datacenters (primary/secondary gateways)	52
Gateway components	53
SD-WAN considerations	63
Security considerations	64
Multi-region datacenters	65
Inter-region connectivity	66
SD-WAN Considerations	68
Supplemental designs	69
Direct internet access	70
Cloud on-ramp	72
Monitoring and reporting	75
FortiManager and SD-WAN monitoring	76
SD-WAN Monitor Map view	76
SD-WAN Table view	77
SD-WAN device monitoring	78
SD-WAN device monitoring of performance SLAs	78

Route table and device dashboards	79
FortiAnalyzer	80
ADOMs, sizing, log storage, scaling, and enforcement	80
SD-WAN logging	81
FortiAnalyzer HA recommendation	87
Evolution to secure SD-branch solution	88
Visibility	88
Attack surface reduction with network segmentation	89
Zero trust local access network	90
SD-branch simplification	91

Change Log

Date	Change Description
2021-12-01	Initial release.
2023-02-16	Updated Architecture and design on page 40 .

Introduction

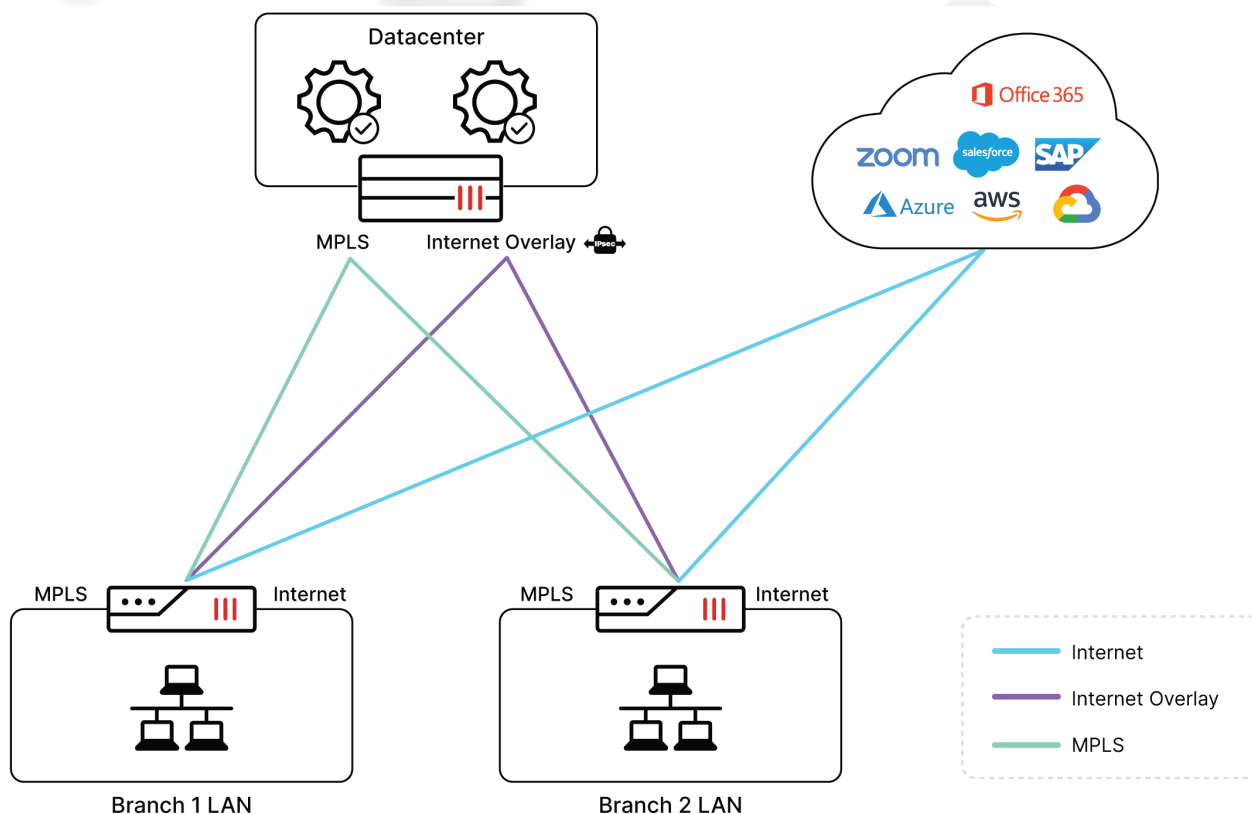
The intention of this reference architecture is to provide an overview of Fortinet SD-WAN solution, along with the components and architectures to satisfy common use cases. This document will cover the Fortinet technology involved in deploying various types of SD-WAN designs, along with considerations and best practices. Our intention is to design a highly scalable, redundant, and secure SD-WAN design that is practical for your organizational requirements.

This document is not intended to be a *step-by-step* configuration guide. Instead, it is meant to be the starting point in your network design, where you begin to draw out the architecture that will be used to meet your specific needs. Fortinet's *SD-WAN Deployment Guide* will cover the *how-to* configuration for some of the common architectures and designs covered in this document.

For more information and documentation about the topics covered in this document, please see the Fortinet Document Library at <https://docs.fortinet.com>.

Executive summary

The following image illustrates a modernized SD-WAN branch edge solution that manages a hybrid architecture inclusive of both private WAN (MPLS) and broadband internet connectivity.



First, the branch has multiple transports, or connectivity options. In this example, the corporate WAN MPLS network remains, but this organization has introduced a single broadband connection to provide direct internet access (DIA) from the branch. In addition, the organization has established an overlay network using Internet Protocol security (IPsec) tunnels between branches and the datacenter over the broadband internet transport. The result is that multiple paths are possible from the branch to both the datacenter and a multi-cloud environment.

Compare this with legacy single-path architecture with a switch connected to a simple router that has one connection to a private WAN. Essentially, there is only one option for egress traffic. But introducing DIA inherently provides for a redundant connectivity architecture. In terms of datacenter connectivity, the overlay network (IPsec tunnel) delivers an alternative path for critical applications that would normally traverse the MPLS. In the same way, the private WAN path will continue to provide its path to the internet, but is now superseded by the DIA connection.

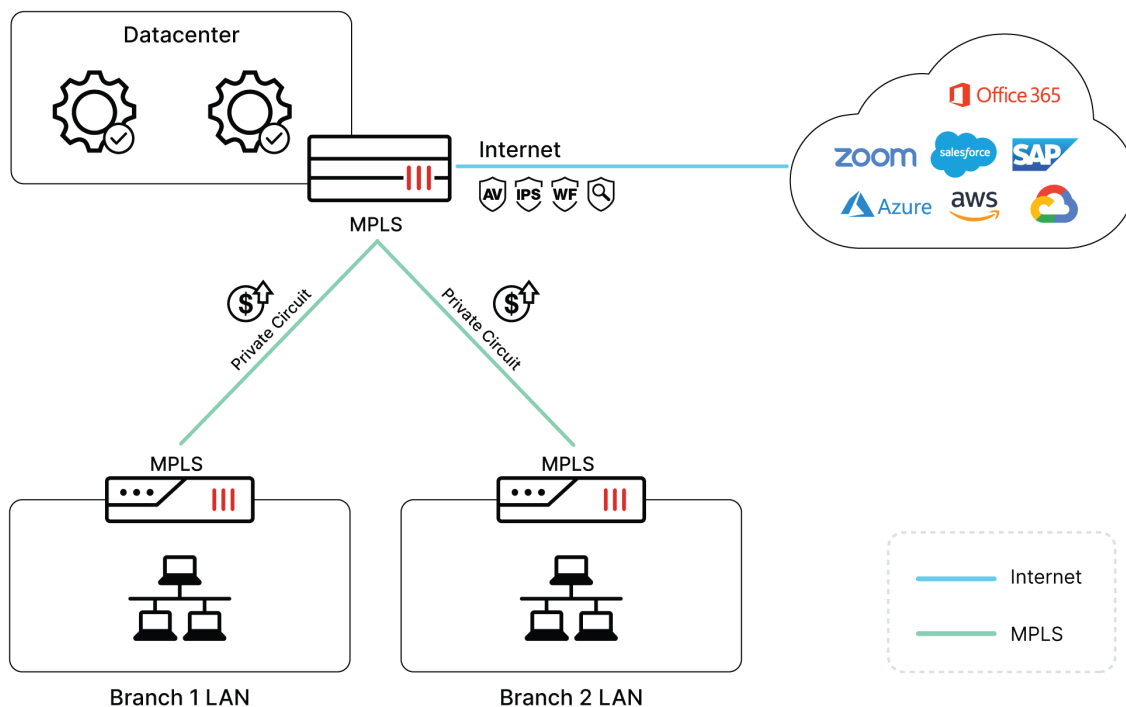
This section includes the following topics:

- [Legacy WAN edge on page 7](#)
- [Transforming the WAN edge on page 8](#)
- [Why Fortinet on page 10](#)
- [Intended Audience on page 12](#)

Legacy WAN edge

For decades, the hub-and-spoke network architecture portrayed in the following image has been commonplace. All network traffic flows through the central corporate datacenter—including traffic moving from branch locations to the internet. Branch traffic travels to the datacenter using dedicated connections, usually multiprotocol label switching (MPLS) circuits.

But a set of forces collectively known as digital transformation (DX) is quickly changing that model. These trends include the digitization of virtually everything in business, the emergence and growth of cloud-based services like Software-as-a-Service (SaaS), and the proliferation of Internet-of-Things (IoT) devices at the network edge. Together, these revolutionary changes necessitate new approaches to networking.



To address the needs of such a widely distributed network, many businesses have embraced solutions, such as a software-defined WAN (SD-WAN), alongside lower-cost connectivity options for businesses. As a result, many organizations have undertaken major WAN edge transformation projects in recent years.

To support these goals, Fortinet Secure SD-WAN leverages:

- Path failover
Moving flows from an under-performing transport to a transport that performs better
- Remote destination monitoring and steering
Detecting issues at a remote location and re-routing traffic through optimal paths
- Link aggregation
Taking advantage of multiple WAN transports
- Active path performance metrics
Viewing WAN underlay metrics and trends
- Application performance improvement
Improving user experience by using a variety of techniques, such as Forward Error Correction, Packet Duplication, QoS, and WAN Optimization

Logically speaking, Fortinet SD-WAN determines which path best meets performance expectations or service-level agreements (SLAs) for a particular application, and assigns application flows to that WAN path.

Transforming the WAN edge

The simplicity of the legacy WAN architecture is evident, specifically with routing. Its hub-and-spoke design requires each remote site to route all non-local traffic to the hub, regardless of the final destination. Legacy

WAN architectures that consist of aging hardware and software solutions continue to provide network connectivity as well as a consistent level of performance and security, and they continue to satisfy some organizations.

However, if an organization needs to add redundancy or additional bandwidth to a legacy WAN infrastructure, complexity can quickly increase. Leveraging private connectivity in a full-mesh approach, for example, would require either multiple static routes or the introduction of a dynamic routing protocol, such as Border Gateway Protocol (BGP) or equal-cost multi-path (ECMP) routing.

For a diagram of the legacy WAN architecture, see [Legacy WAN edge on page 7](#).

Improving inefficient routing and inferior performance

Even if an organization avoids the complexity of multiple static routes or a dynamic routing protocol, its network traffic is extremely inefficient. Consider a branch user's legacy path to the internet in a legacy WAN architecture. In order to arrive at Google's search engine website for a simple search, for example, the application flow would need to:

- Cross the branch WAN edge
- Navigate across the MPLS circuit
- Enter the datacenter
- Negotiate its way through a centralized security stack that includes a firewall, intrusion prevention system (IPS), antivirus/anti-malware (AV/AM), data loss prevention (DLP), web filter, and so on
- Travel to the Google website through the datacenter internet edge

The minimal infrastructure required at the branch was traditionally seen as a key benefit of legacy WAN architecture. However, it has largely fallen short of expectations concerning user experience. At a time when consumers have almost universally been using broadband connections at home for more than a decade, legacy WANs do not generally reflect typical broadband speeds. As more and more employees use cloud-based services that require more bandwidth, performance has only declined.

For a diagram of the legacy WAN architecture, see [Legacy WAN edge on page 7](#).

Fixing security gaps and bottlenecks

The ability to centralize the security stack was also previously seen as a benefit of legacy WAN architecture. Branch sites typically have a simple router for connectivity to an MPLS or other private WAN circuit. Because all flows must first traverse the WAN, it made sense to centralize advanced security capabilities at the core instead of building distributed stacks at each branch.

Unfortunately, flows failing security policy must traverse the WAN before they are inspected. As a result, infected hosts are often permitted to freely communicate throughout the enterprise network because security only exists within the datacenter, and site-to-site traffic therefore passes without inspection.

Another issue with the centralized security stack is performance. As traffic increases—especially traffic bound for the internet and cloud-based resources—security inspections can become a bottleneck, with legitimate traffic waiting in line behind traffic that may not be permitted to continue.

Modernizing WAN

Modernization of a WAN infrastructure is not just about replacing end-of-life hardware or software. WAN edge redesign is a business solution, not simply a technology requirement. Budgets are growing to accommodate digital transformation (DX) not because organizations prefer to consume cutting-edge

technology, but because their customers are demanding this technology. The hope of improved user experience and increased productivity loosens purse strings, and provides necessary budgetary resources for technology leaders to initiate WAN transformation projects.

SD-WAN is one of the primary innovations behind WAN edge modernization. Its core capabilities include multi-path control, application awareness (such as with SaaS solutions), and the resultant dynamic application steering. These capabilities enable network traffic to be routed over the public internet or over private infrastructure—whatever is most efficient for application performance and availability in a multi-cloud environment.

Reducing risk with Secure SD-WAN

Secure SD-WAN adds the advanced security capabilities of a next-generation firewall (NGFW) to the networking solution. It's no accident that the icon in the modernized SD-WAN branch edge solution that represents the SD-WAN device at the branch edge looks like a firewall. This is because introducing DIA at the branch also establishes direct connectivity to a volatile threat landscape. Such connectivity did not exist in the legacy architecture, which routed all traffic through a centralized security stack. The DIA necessitates that the centralized security stack give way to a more distributed security architecture.

In a multi-cloud environment with many SaaS solutions, it is especially important that the secure SD-WAN solution be able to distinguish between applications to leverage the full functionality of the solution. In addition to distinguishing applications and controlling a multi-path environment, a secure SD-WAN solution provides dynamic application steering (packets or sessions) to traverse available paths to the corporate WAN or the multi-cloud environment. To aid application steering, it provides active path metrics. In conjunction with customer-defined SLAs, the SD-WAN policy engine determines which paths are viable transports for each application, choosing the best path or balancing traffic between multiple viable paths.

For a diagram of the modernized SD-WAN branch edge solution, see [Introduction on page 6](#).

For a diagram of the modernized SD-WAN branch edge solution, see the *Introduction* in [SD-WAN Architecture for Enterprise](#).

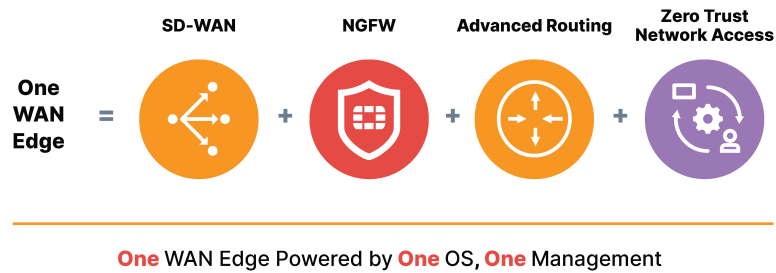
Supporting DX initiatives

In summary, for this high-level secure SD-WAN architecture example, digital transformation (DX) is the driver for branch edge modernization. Organizations are creating projects to address WAN connectivity models (such as MPLS, Long-Term Evolution (LTE), and broadband) and edge device consolidation (such as router, firewall, advanced security). Organizations are also adding SD-WAN functionality to improve branch-user experience, maintain application performance, and sustain application availability.

Why Fortinet

Fortinet offers a broad portfolio of integrated and automated security tools covering network security, cloud security, application security, access security, network operations center (NOC), and security operations center (SOC) functions.

The Fortinet Secure SD-WAN solution accelerates network and security convergence with enterprise-grade SD-WAN, advanced routing, Next-Generation Firewall, and recently added access proxy for Zero Trust Network Access (ZTNA) support. This simplifies the LAN and WAN architecture to provide a unified Fortinet WAN edge—powered by a single OS and controlled with a single management solution. Not only are we providing the best-in-class SD-WAN solution, but the technology is also integrated with network access to deliver the most secure and manageable remote branch in the industry.



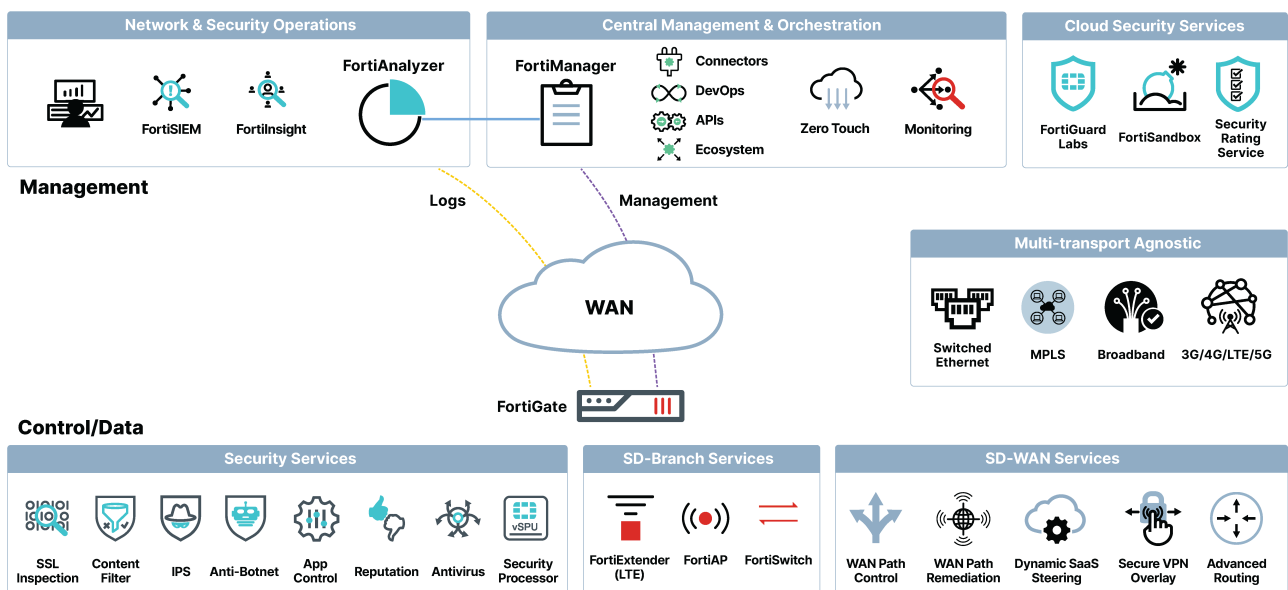
The Fortinet Secure SD-WAN goes beyond traditional SD-WAN requirements to provide a flexible and scalable fit for all enterprise sizes and requirements. The FortiGate device, with its underlying firmware FortiOS, is the basic component of the Secure SD-WAN solution. It offers self-healing capabilities without the bottleneck or single point of failure a centralized orchestrator would provide.

Benefits of a controllerless-based architecture

A major differentiator from other SD-WAN vendors, Fortinet Secure SD-WAN offers a controllerless-based architecture where each FortiGate device maintains control-plane autonomy at the branch edge. In other words, the solution does not require a centralized or cloud-based controller to provide control-plane operations for application steering. Instead, each FortiGate edge device operates independently to evaluate available path efficacy and choose the most appropriate path for applications to traverse the WAN, whether the selected link be an overlay interface (IPsec) or an underlay interface (MPLS, DIA).

Unique, unbeatable design

At the same time, this architecture maintains a centralized approach for full monitoring, management, analytics, and reporting capabilities over the entire enterprise deployment. To achieve this, FortiManager acts as a single pane of glass to simplify operations. The following image demonstrates how each FortiGate edge device communicates with centralized components, but maintains all control-plane functionality at the edge. Transport-agnostic link support, SD-WAN core capabilities and services, and NGFW services are all delivered throughout the enterprise without dependency for control-plane input from an external device.



Industry leader

Fortinet is recognized as a Gartner Leader and positioned highest for Ability to Execute in WAN Edge Infrastructure. It is also ranked number one in three out of five Critical Capabilities use cases, which is higher than any other vendor.

Fortinet is also amongst the earliest SD-WAN technology vendors to be certified by the Metro Ethernet Forum ([MEF](#)), the world's defining authority for standardized services designed to address the most demanding networking needs of today's digital transformation efforts.

Fortinet has been an active member of MEF since 2017, and is closely partnering with them to develop new SD-WAN security standards. Fortinet currently leads a key initiative in the MEF Applications Committee on application security for SD-WAN services (MEF 88), and has won two MEF 3.0 Proof of Concept awards for developing security standards for secure connections between separate SD-WAN devices, and for ensuring application security for SD-WAN services.

Intended Audience

This guide has primarily been created for a technical audience, including system architects and design engineers who want to deploy Fortinet Secure SD-WAN or Secure SD-Branch in a managed offering capacity.

It assumes the reader is familiar with the basic concepts of applications, networking, routing, security, and high availability, and has a basic understanding of network and datacenter architectures. For implementation, a working knowledge of FortiOS networking and policy configuration is ideal.

Uses cases

FortiGate SD-WAN is a flexible solution that can be used to cover a broad range of different use cases. This section covers some of the most common use cases.

Use case	Description
Dynamic application steering across multiple WAN links	Dynamic application steering across multiple WAN transports based on the business intent.
Redundant connectivity for enterprise branch	Leveraging multiple WAN transports to provide branch redundancy from failures and reduced performance. WAN links may include private circuits, public internet, LTE/5G, or satellite connectivity.
Reduce WAN OPEX with direct internet access	Secure, local internet breakout of SaaS applications and internet traffic without the need to offload to a remote location.
Secure and automated intra-site connectivity	Secure transport and connectivity of corporate traffic between branch, head quarters, datacenter, and other locations.
Multi-cloud connectivity and cloud on-ramp	Connectivity and intelligent steering of network traffic between one or more cloud locations.
Application performance improvement	Improving application and network performance by using various methodologies, including steering between best performing links, QoS, WAN remediation, WAN optimization, and more.
Work from anywhere	Remote workers connecting to corporate resources from any location.

This section includes the following topics:

- [Dynamic application steering across multiple WAN links on page 14](#)
- [Redundant connectivity for enterprise branch on page 14](#)
- [Reduce WAN OPEX with direct internet access on page 14](#)
- [Secure and automated intra-site connectivity on page 15](#)
- [Multi-cloud connectivity and cloud on-ramp on page 15](#)
- [Application performance improvement on page 15](#)
- [Work from anywhere on page 16](#)

Dynamic application steering across multiple WAN links

One of the most common SD-WAN use cases is to leverage multiple WAN transports to dynamically steer network traffic based on the business intent. The FortiGate SD-WAN's core capabilities include multi-path control, application awareness (such as with SaaS solutions), and the resultant dynamic application steering. These capabilities enable network traffic to be routed over the public internet or over private infrastructure—whatever is most efficient for application performance and availability in a multi-cloud environment.

By leveraging a variety of different methods described in this documentation, the FortiGate SD-WAN chooses the optimal path based on the business requirements and protects the application as it traverses the WAN.

Redundant connectivity for enterprise branch

Modern branch locations require maximum availability and uptime for business-critical services. A combination of private circuits (MPLS), public internet, LTE/5G wireless connectivity or satellite WAN transports may be required to achieve redundancy from WAN failures and impairments. The SD-WAN solution needs to manage a hybrid of public and private WAN connections, while also providing intelligent steering across multiple, diverse connections.

FortiGate SD-WAN provides sub-second detection of issues across all available WAN paths to provide failover based on user configuration SLA's. Using the techniques we'll review in this document, SD-WAN rules may be configured to take advantage of all available paths based on the business requirements.

Reduce WAN OPEX with direct internet access

The traditional WAN model consisted of using expensive private circuits for all connectivity to business services. This model involved sending all packets to a central location where security inspection would take place, and policies would control traffic flow. As businesses move their workloads to SaaS and cloud, the need for more bandwidth and intelligent steering is required.

In the modern WAN edge model, it is now common for branch locations to share multiple WAN links of varying transport dependencies. Secure DIA (direct internet access) provides intelligent and secure steering of network traffic based on business requirements. Applications destined to a SaaS or cloud provider can be sent directly to the internet using a public internet connection without the need of backhauling to a central location. This allows for a much more efficient use of WAN bandwidth and improved user experience. Because FortiGate SD-WAN is also a Next-Generation Firewall, internet traffic can be locally inspected and controlled without needing to offload inspection to another location.

An important consideration for this use case is that edge locations may consist of many different WAN types. The FortiGate SD-WAN solution is transport agnostic, and can be mixed and matched with several different WAN types, including MPLS through Ethernet handoff, internet, and LTE.

Secure and automated intra-site connectivity

As businesses look towards supplementing or replacing their traditional MPLS or private links, the need for secure transport connectivity over public internet becomes essential. For customers that have business applications in a datacenter or HQ location, branch locations require remote connectivity to access these resources. For redundancy, there may be multiple WAN transports and multiple gateways across geo-redundant locations that could be used to access the resource. This leads to many potential paths from the edge that must be evaluated and steered accordingly.

The Fortinet SD-WAN solution can be used to automatically set up dynamic, site-to-site connections across all corporate resources and locations. This may include branch offices, HQ sites, datacenters, or public cloud providers. FortiGate SD-WAN appliances can be configured at any location as a physical or virtual device to provide connectivity and intelligent steering across all available WAN links.

FortiManager is our central management solution that can provide *light touch* or *zero touch* provisioning to onboard new locations, and provide centralized management and monitoring. This means you can choose how much to automate during the onboarding process.

Multi-cloud connectivity and cloud on-ramp

Multi-cloud connectivity refers to the ability of the edge location to steer traffic to one or multiple cloud environments. This could span across workloads within a single cloud provider or multiple cloud providers. As business requirements evolve and cloud providers continue to offer differentiating services, a flexible solution with broad support is critical.

Because of the dynamic nature of cloud environments, the edge device needs to know in real-time where the application is hosted. SaaS or IaaS applications and services should be available via redundant paths, with appropriate routes automatically added or removed as services in the cloud change. Then all available paths should be measured and steered according to the business intent.

Application performance improvement

Application performance can be impacted for a variety of issues. Congestion, WAN link performance (such as, high latency, jitter, or packet loss), server issues, routing changes, and other network issues could all be detrimental to performance and user experience. The FortiGate SD-WAN solution incorporates a number of techniques and capabilities to improve application performance for a variety of deployment methods.

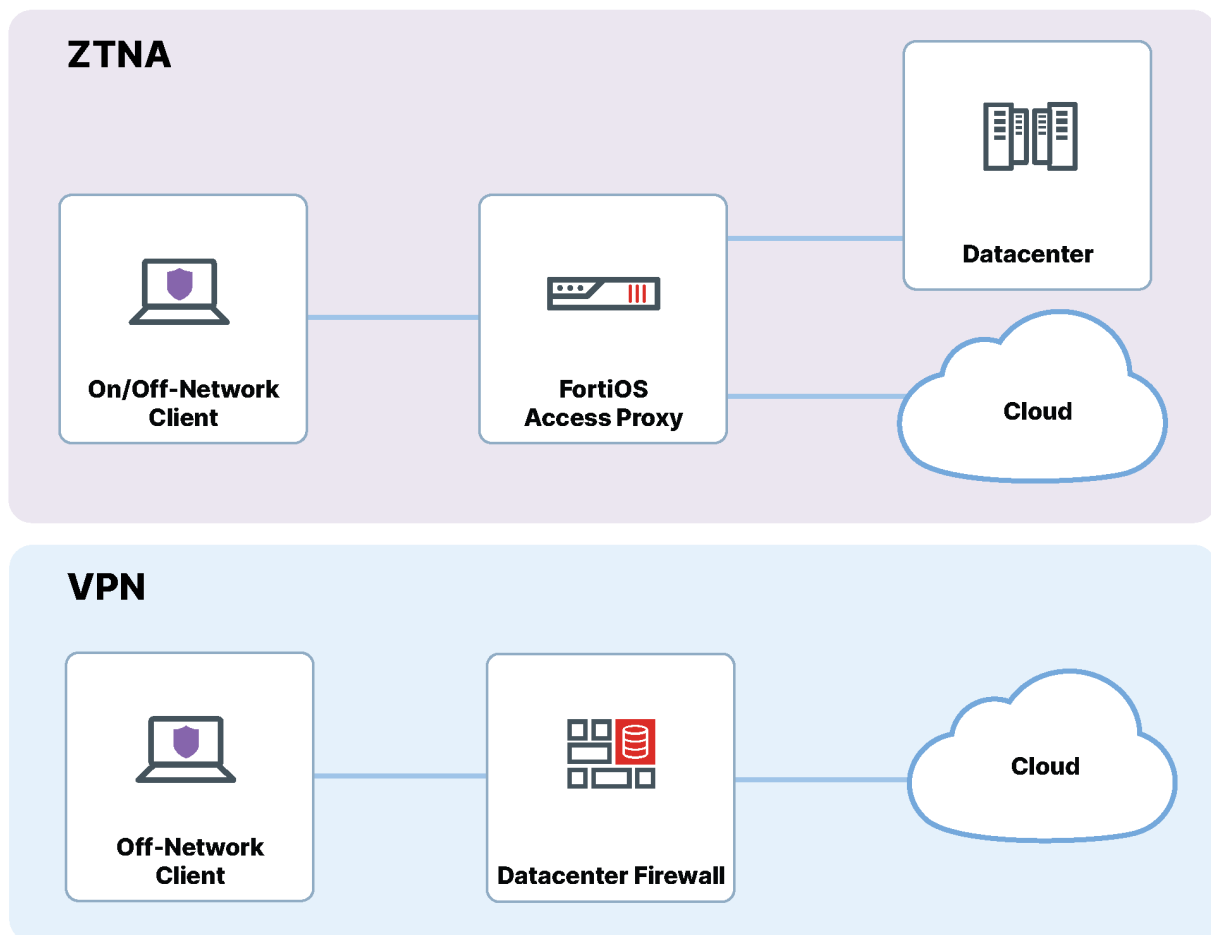
- **Application identification and steering:** Performance SLAs check all available paths to the application to determine the best performing link. Potential network issues can be detected in sub-second speed and steered to a better performing link based on the business requirement.
- **WAN Remediation (Packet loss correction):** There are situations where protecting the application from packet loss is crucial to business continuity. WAN remediation refers to a series of techniques to fix packet loss on a WAN link. Forward Error Correction (FEC) and Packet Duplication are WAN remediation techniques that can be used to protect a link from various types of impairments.
- **QoS (Quality of Service):** Because bandwidth is finite, sometimes it is necessary to prioritize how traffic is distributed based on business needs. FortiGate SD-WAN allows you to police, shape, and queue network traffic at each location. Adjustments can be made automatically, depending on available bandwidth at certain times of the day.

- **WAN Optimization:** WAN optimization is a comprehensive solution that maximizes your WAN bandwidth and improves user experience. Some techniques include protocol optimization, web caching, explicit proxy, byte caching, and more.

Work from anywhere

Remote users requiring access to corporate resources is more important than ever. Today's workers require secure, remote access to corporate resources from wherever they are located, as easily as they would inside the corporate network. Work from anywhere (WFA) solutions should be integrated into the SD-WAN device to provide consistent security policy enforcement from anywhere the user is located.

Fortinet SD-WAN devices include support for Virtual Private Network (VPNs) and Zero Trust Network Access (ZTNA) . The FortiGate SD-WAN device can act as a VPN server or access proxy, depending on your desired WFA deployment.



Solutions and technologies

Fortinet Secure SD-WAN consists of several components:

- FortiGate NGFW, which runs FortiOS, is the core of Secure SD-WAN
- Fortinet ZTNA Access Proxy, which runs natively in FortiOS, starting in FortiOS 7.0
- FortiManager for the orchestration and management plane
- FortiAnalyzer for advanced analytics and automation
- FortiPortal to provide a scalable and flexible customer self-service portal

Fortinet Secure SD-WAN solution can be extended to Secure SD-Branch. SD-Branch consists of the following components:

- FortiSwitch to provide security on the wired LAN edge
- FortiAP to provide WiFi access to users

This section includes the following topics:

- [FortiGate on page 17](#)
- [FortiManager on page 20](#)
- [FortiAnalyzer on page 22](#)
- [FortiAP on page 24](#)
- [FortiSwitch on page 26](#)

FortiGate

With its underlying FortiOS firmware, FortiGate is the product at the foundation of Fortinet's Secure SD-WAN solution. A key differentiation from other SD-WAN vendors is that the FortiGate Secure SD-WAN platform provides the following key capabilities:

- Built-in intelligence to decide the best path for a specific application
- Integrated and native Next-Generation Firewall security inspection
- Overlay network connectivity in the SD-WAN architecture

The above capabilities don't require a centralized controller as do most of the traditional SD-WAN vendors.

FortiGate is multitenant at its very core. Virtual domain (VDM) technology is a testament to this statement, enabling a single, secure gateway instance to be sliced into potentially hundreds of individual gateways.

For more details on the FortiGate SD-WAN capabilities, see [Technical background on page 30](#).

FortiGate also:

- Delivers advanced routing support (RIP, BGP, OSPF, and more)
- Participates in virtual private network (VPN) pairing as a spoke or hub (concentrator)
- Brings WAN optimization by means of protocol optimization and byte and object caching
- Supports traffic shaping and packet priority to ensure that business-critical applications take precedence

The following sections describe some of the key functionality:

- [Application identification on page 18](#)
- [Increased performance on page 18](#)
- [Form factor on page 18](#)
- [Security on page 19](#)
- [Zero touch deployment on page 19](#)
- [API / automation on page 20](#)

Application identification

Application flow definition and detection is the cornerstone of any SD-WAN solution. Policies for traffic engineering depend on precise and evolving definitions of application traffic and traffic flows.

Fortinet's [FortiGuard](#) maintains a database of more than 5,000 application definitions. Fortinet's applications detection capabilities are derived from mature data modeling created and maintained by FortiGuard Labs. FortiGate also enables the ability to define custom application flows where needed.

Increased performance

IPsec is the overlay technology recommended for the Fortinet Secure SD-WAN solution, as it provides confidentiality, integrity, and mutual site authentication. The Security Processing Unit (SPU) helps you achieve the best performance for the lowest cost, thanks in part to its IPsec offloading capabilities. The number of tunnels and encryption requirements can grow exponentially with the number of edge devices (full mesh), making the efficiency of tunnel management a critical part of the solution.

FortiGate virtual machines support all primary, generic network accelerations (SR-IOV, DPDK) to deliver fast and secure features to all possible deployments. Furthermore, we have developed vSPU capabilities to offload more features into the accelerated generic NICs (DPDK mainly) to get the most efficiency from the hypervisors and hardware.

Form factor

FortiGate is available in the following form factors:

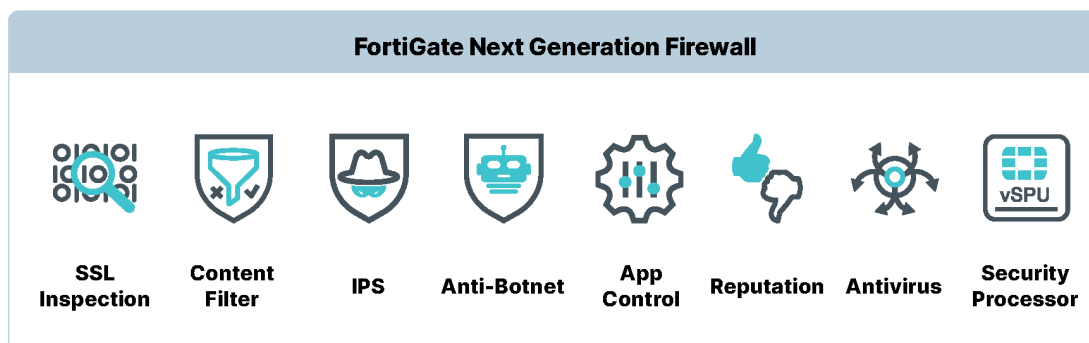
- Physical appliance
- Virtual machine for both public and private cloud environments

The physical and virtual appliances offer one-to-one feature parity, allowing your SD-WAN architecture to span from on-premises to the public cloud with the same functionality.


Security

FortiGate is a fully functioning, market-leading Next-Generation Firewall, meaning security is at the heart of the SD-WAN solution.

The following security functions are provided:



All the security features available in the FortiGate can be leveraged when SD-WAN is implemented.



The advantage of this integrated approach is in the efficiency found in processing packets in parallel for different security functions and SD-WAN in the same device, thereby reducing latency, integration, and management overhead.

Zero touch deployment

Fortinet zero touch provisioning allows a self-service type of deployment of the FortiGate. Simple cabling skills are the only technical requirement at every branch to add new devices to the SD-WAN solution. The devices also have a predefined callback to Fortinet. This enables the fully automated process of adding the device to FortiManager and maintaining the evolving SD-WAN configuration.

API / automation

Every FortiGate exposes REST API, which provides complete management and monitoring capabilities. APIs are a crucial component of the solution, allowing Fortinet Secure SD-WAN to integrate with third-party orchestration and management systems if required. More information on the FortiGate API can be found in the [Fortinet Developer Network](#).

FortiManager

FortiManager offers all the necessary tools to manage and orchestrate Fortinet Secure SD-WAN solutions. You can quickly deploy thousands of edge locations, trigger changes to entire groups of devices, and consistently define security and SD-WAN policies throughout your environment.

FortiManager reduces administration and workload costs with smart features, such as device discovery, device group creation by administration domain, audit, and management of complex SD-WAN architecture.

FortiManager is available in the following form factors:

- Physical appliance
- Virtual machine for both public and private cloud environments
- SaaS offering service directly from Fortinet

The key features are:

- **Single console management:** manage FortiGates and any subordinate FortiSwitch, FortiAP, and FortiExtender devices. Provide signature updates to FortiMail, FortiSandbox, and FortiClient.
- **Multi-tenancy and administrative domains (ADOMs):** separate customer data and manage domains with ADOMs to be compliant and operationally effective.
- **Centralized policy and device management:** centrally manage up to 100,000+ devices and policies, such as firewalls, switches, and access points.
- **Zero touch provisioning:** automate workflows and configurations for Fortinet firewalls, switches, and wireless infrastructure.
- **Secure SD-WAN provisioning and monitoring:** provision and monitor Secure SD-WAN from one console across your network, branch offices, or campuses.
- **Enterprise-grade high availability and integration:** automate backups to up to five nodes with streamlined software and security updates for all managed devices.
- **Security automation:** reduce complexity and costs by leveraging automated REST API, scripts, connectors, and automation stitches.

This section includes the following information about some of these key features:

- [Single console management on page 21](#)
- [Administrative domains on page 21](#)
- [Centralized policy on page 21](#)
- [Zero touch provisioning on page 21](#)
- [Secure SD-WAN capabilities on page 22](#)
- [Secure SD-WAN security automation on page 22](#)

Single console management

FortiManager provides insight into network traffic and threats through a single pane of glass and offers enterprise-class features and sophisticated security management for unified, end-to-end protection to contain advanced threats. FortiManager also delivers the industry's best scalability to manage up to 100,000 Fortinet devices.

Access to the equipment is secured with both administration accounts and associated profiles. Credentials for administrator accounts are determined by the associated profiles. The account may be local, which means it is specific to the equipment, or external when linked to an authentication base (LDAP, RADIUS, TACACS+, PKI...), centralizing all administrator accounts.

FortiManager, coupled with the FortiAnalyzer family of centralized logging and reporting appliances, provides a comprehensive and powerful centralized management solution for all organizations.

Administrative domains

Administrative domains (ADOMS) allow for granular device and role-based administration for deploying zero trust, multi-tenancy architecture to large enterprises by using a hierarchical objects database to facilitate the reuse of common configurations and serve multiple customers. ADOMs are often used by MSSPs, Telcos, Federal and state agencies as well as large enterprises to segment device management between customers or departments.

ADOMs are used to manage independent security environments, each with its security policies, configuration database, and SD-WAN parameters. The intuitive GUI makes it easy to view, create, clone, and manage ADOMs for each customer. It is also possible to define global objects, such as firewall objects, policies, and security profiles to share across multiple ADOMs. Granular permissions allow assigning ADOMs, devices, and policies to users based on role and responsibilities.

Centralized policy

Policies and objects are managed by means of packages that can be global or local to an ADOM (administrative domain). A policy package contains a set of security rules deployed on a unique device or a group of devices. An ADOM may contain several policy packages that can be deployed on one or more devices or VDOMs.

All objects that compose the rules (addresses, time ranges, interfaces, services, and so on) can be defined with static or dynamic values when they change for each device. The association of a policy package to a device or VDOM is performed after the creation of the policy package from the *Policy & Objects* module and the *Installation* tab.

Zero touch provisioning

Zero touch deployment uses templates to provision devices for quick, mass deployment and support firmware version enforcement. To support the zero touch configuration, FortiManager leverages the *Add Model Device* feature that allows an administrator to provision a model device and automatically apply the configuration associated with that model device, once a FortiGate with a matching identifier is registered to FortiManager.

Secure SD-WAN capabilities

FortiManager offers powerful SD-WAN management capabilities using intuitive workflow and simplified provisioning at scale. Enhanced SD-WAN analytics monitor application performance and bandwidth utilization per WAN link. Leverage application-centric SD-WAN business policies to fine-tune traffic steering decisions based on performance SLA targets for each WAN provider.

Secure SD-WAN security automation

In addition to the GUI, FortiManager can be used via REST API. RESTful API allows MSSPs/large enterprises to create customized, branded web portals for policy and object administration. Automate common tasks such as provisioning FortiGate and configuring existing devices. More information on the FortiManager API can be found in [Fortinet Developer Network](#).

FortiAnalyzer

FortiAnalyzer collects information, such as traffic and security events, and reduces the effort required to monitor the information system.

The FortiAnalyzer solution is responsible for the collection and the valuation of logs generated by FortiGate, FortiMail, FortiClient solutions, FortiWeb, FortiManager, FortiSandbox, FortiDDoS, and FortiCache. It receives logs, stores them, produces predefined and customized reports, and supports configuration of advanced alerting.

FortiAnalyzer provides two operation modes: Analyzer and Collector. Analyzer mode is the default mode that supports the full FortiAnalyzer features. The primary task of a Collector is to receive logs from connected devices and upload the logs to an Analyzer. Instead of writing logs to the database, the Collector retains them in their original (binary) format and sends them to the Analyzer.

FortiAnalyzer is available in the following form factors:

- Physical appliance
- Virtual machine for public and private cloud environments
- SaaS offering service directly from Fortinet

The key features are:

- **Security Fabric analytics:** event correlation across all logs and real-time anomaly detection, with Indicator of Compromise (IOC) service and threat detection, reducing time-to-detect.
- **Fortinet Security Fabric integration:** correlates with logs from FortiClient, FortiSandbox, FortiWeb, and FortiMail for deeper visibility and critical network insights.
- **Security automation:** Reduce complexity and leverage automation via REST API, scripts, connectors, and automation stitches to expedite security response.
- **Multi-tenancy and administrative domains (ADOMs):** separate customer data and manage domains with ADOMs to be compliant and operationally effective
- **Flexible deployment options and archival storage:** supports deployment of an appliance, VM, hosted or cloud storage. Use AWS, Azure, or Google to archive logs as a secondary storage.

This section includes the following information about some of these key features:

- [Security visibility on page 23](#)
- [Administrative domains on page 23](#)

- [Automatic security and SD-WAN reports on page 23](#)
- [Security automation on page 24](#)

Security visibility

Administrators can access the FortiAnalyzer unit from a GUI (through a web browser) without any specific software client. From the GUI, a global dashboard provides links to other main menus:

- **Dashboard** has been designed to give a detailed view of the logging activity in the managed environments. The admin can quickly appreciate the average log rate as well as the number and volume of logs collected every day over a week. This information is crucial for designing the logging policy and working around *capacity planning*.
- **FortiView** provides broad visibility on traffic, applications in use, threats, and the most visited websites in just one click. FortiView aggregates and then analyzes all data to instantaneously highlight the most relevant piece of information, and you can consult the information provided by each graph. You can also click graphs to view details of the underlying events.
- **Log View** is intuitive and easily edited to optimize access to relevant information. A powerful search engine allows the filtering of logs according to multiple criteria.
- **Event Monitor** correlates the logs to generate security alerts that the administrator can acknowledge, analyze, or delete. Double-click on a security event to show the list of all events linked to the alert. It is also possible to configure rules (send email, SNMP, or syslog) based on the log content to generate alerts, when an event or an event sequence occurs. Again, the most valuable alerts have been predefined to address the most frequent use cases.

Administrative domains

Each Fortinet device (FortiMail, FortiGate) able to send logs must be declared in FortiAnalyzer and can be associated with an administrative domain (ADOM). This distribution helps the admin segment the solution into environments different from each other. A custom storage quota per ADOM can be configured, as well as access rights for the users of the domains. This segmentation is optional.

A dedicated dashboard helps you monitor the status of the quota globally or by equipment.

FortiAnalyzer supports a disk space management feature relying on the quota allocated to the domain. When the threshold for the quota is exceeded, an alert is sent. The log files automatically begin rotating when the quota is reached.

Automatic security and SD-WAN reports

FortiAnalyzer allows administrators or business owners to generate automatic SD-WAN reports targeted to executive management. These reports provide immediate information to assess the benefits of the SD-WAN solution while at the same time aggregating critical security information. While the highlights are listed in a convenient executive summary report, each section provides a more detailed view. This includes a set of recommended actions at the end of the report, plus actionable steps an organization may take to optimize their network for DIA, protect their organization from external/branch office threats, and ultimately reduce expenditures and save money.

Security automation

In addition to the GUI, FortiAnalyzer can be used via REST API. RESTful API allows MSSPs/large enterprises to create customized, branded web portals for policy and object administration. Automate common tasks such as provisioning FortiGates and configuring existing devices. Join Fortinet Developer Network (FNDN) to access exclusive articles, how-to content for automation and customization, community-built tools, scripts, and sample code.

Complete documentation is available on FNDN at <https://fndn.fortinet.net/>.

FortiAP

The most common form of access at the LAN edge for users these days is WiFi. Wireless access points can be added to any network to provide WiFi access to employees and guests alike. The challenges of adding wireless to a deployment go far beyond the physical installation of the hardware.

The screenshot displays the FortiAP management interface. On the left, a sidebar lists various configuration sections including Dashboard, Security Fabric, Network, System, Policy & Objects, Security Profiles, VPN, User & Authentication, and WiFi & Switch Controller. The main area shows a table of managed FortiAPs with columns for Access Point, Status, SSIDs, and Channels. A detailed configuration page for 'FAP-Hallway' is open on the right, showing general information like Serial Number, Base MAC Address, and Status. Below this, a 'Radios' tab provides a comparison of Radio 1 (2.4 GHz), Radio 2 (5 GHz), and Radio 3 (Monitor) settings, including Mode, SSID, Clients, Bandwidth, Operating Channel, Channel Utilization, Channels, and Operating TX Power.

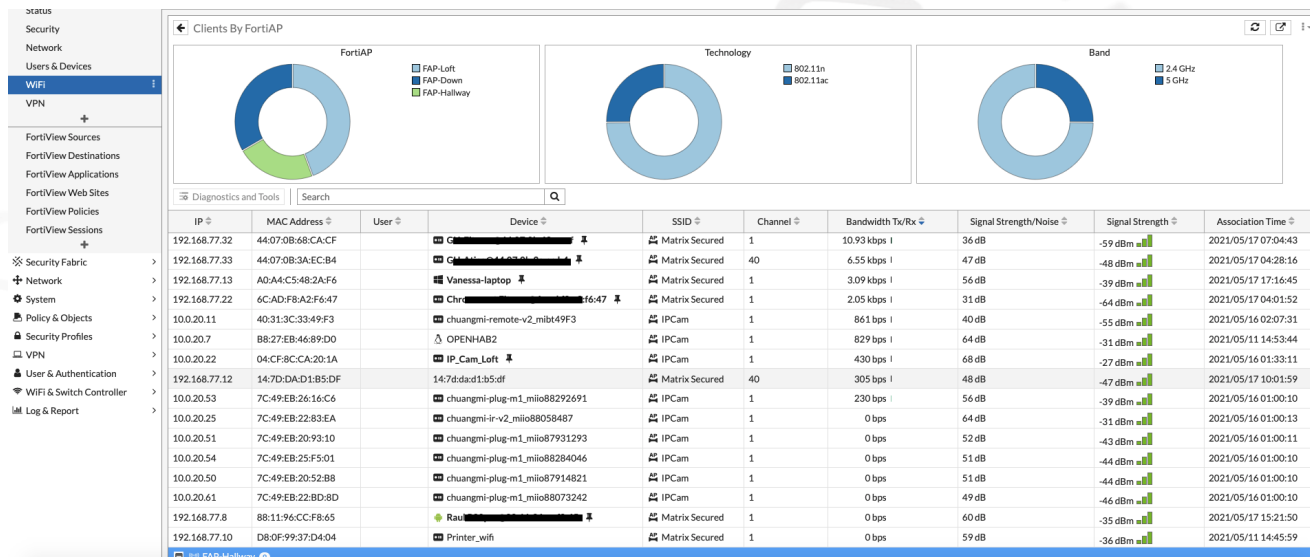
	Radio 1 - 2.4 GHz	Radio 2 - 5 GHz	Radio 3
Mode	AP	AP	Monitor
SSID	IPCam (ipcam) Red_Invitados (Guests)	MatrixSecured (wifi)	N/A
Clients	7	1	0
Bandwidth Tx	24.05 kbps	8.27 kbps	N/A
Bandwidth Rx	64.08 kbps	4.52 kbps	N/A
Operating Channel	6	36	N/A
Channel Utilization (2.4 GHz)	36%	N/A	Lowest: Channel 8 (0%) Highest: Channel 13 (5%)
Channel Utilization (5 GHz)	N/A	12%	N/A
Channels	1, 6, 11	36, 40, 44, 48, 52, 56, 60, 64, 100, 104, 108, 112, 116, 120, 124, 128	N/A
Operating TX Power	19 dBm	17 dBm	N/A
Band	802.11n/g	802.11ax/ac/n/a	N/A

Network IT demands more capability and reliable security from fewer components to save on cost and simplify the environment. Fortinet's wireless LAN equipment leverages Security-Driven Networking to provide secure wireless access for the enterprise LAN edge. Perfect for deployments from the campus to the SD-Branch, FortiAPs are Fortinet Security Fabric enabled, providing the broad visibility, automated protection, and integrated threat intelligence required to protect organizations' valuable assets and data worldwide. And that includes REST API support for most of the features used.

LAN edge equipment from Fortinet converges networking and security into a secure, simple-to-manage architecture with a single focal point for management and configuration. By leveraging Security-Driven Networking, Fortinet allows you to secure the LAN edge without the need for costly and complex licensing schemes while benefiting from all the current cutting-edge WiFi enhancements, depending on the models. From the same dashboard used to manage the Next-Generation Firewall and Policies, you also have complete visibility over the wireless client details:

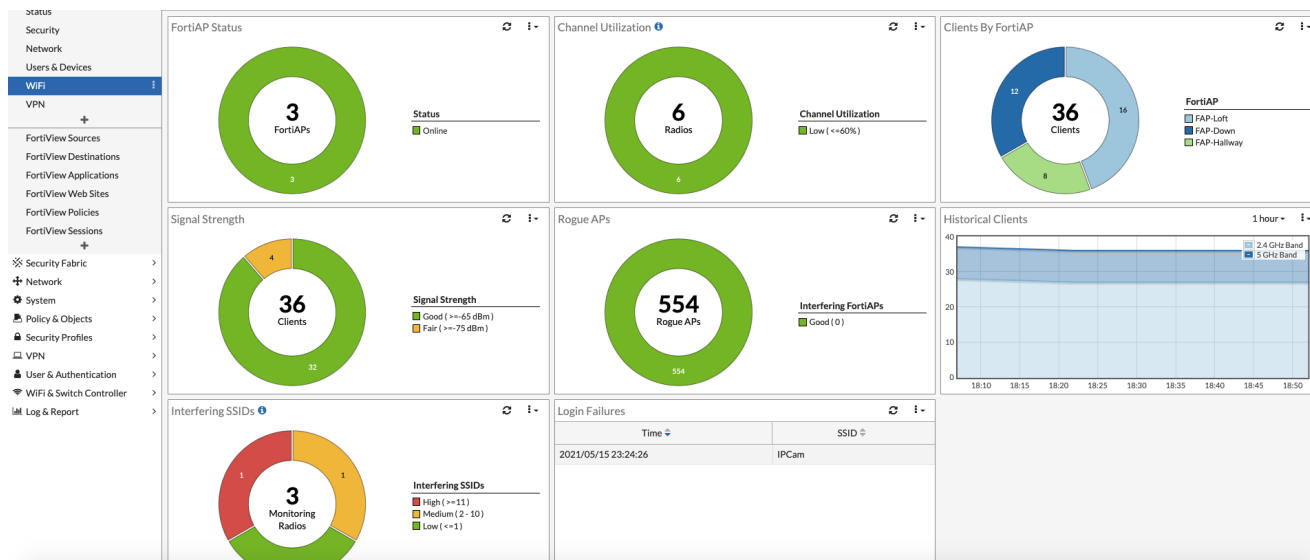
- Username
- Since when it is connected
- Type of encryption used
- Which SSID and VLAN it is connected to
- With which device (name, MAC address, IP address) using which operating system type

- On which Fortinet wireless access point (which is also displayed on the WiFi Maps)
- At what quality (signal strength, data rate, WiFi band, TX/RX bandwidth, spatial streams)



Configuring and managing access points from the same known dashboard as the security parameters also allows immediate visibility and troubleshooting advantages. One can very quickly understand:

- Which access points are online or down
- The last join time and failure reason
- How many wireless clients are connected to each AP
- The WiFi channels used, the TX power, and at what utilization percentage of the channel they operate
- The SSIDs being advertised and in which mode (tunneled, bridged, mesh)
- If the regulatory requirements are being met
- Which wireless IDS profiles are being used



All of that allows for the easy operation of a live and evolving secure wired and wireless network for administrators and a trusted infrastructure for users to perform their daily job without worrying about the underlying connectivity.

The key features of FortiAPs are:

- **Advanced security protection:** Wireless LAN security done right, from the leader in network security. Integrated Firewall, IPS, Application Control, and Web Filtering protect the wireless LAN from the latest security threats, with SSIDs that can natively be scheduled for availability.
- **Integrated WIDS and rogue AP suppression:** Protects the network from advanced wireless threats and satisfies PCI DSS compliance with the integrated Wireless Intrusion Detection System to report and suppress phishing SSIDs.
- **Deep application control:** Fortinet goes above Wireless Multimedia Extensions (WME) by offering deep Layer 7 inspection to precisely control applications and bandwidth usage.
- **Dynamic Automatic Radio Resource Provisioning (DARRP) and RX-SOP:** Advanced wireless techniques for optimized throughput to eliminate sticky clients and maximize channel efficiency in all wireless environments are applied at the AP and managed from the FortiGate and FortiManager interfaces.
- **Multiple PSK, voice enterprise certifications, and agile multi-band operations:** To ensure that Internet-of-Things devices, regular clients, and all smart devices reach maximum capacity in the most secure way available to them with user and traffic segmentation.

See also [Important terms for FortiAP on page 26](#).

Important terms for FortiAP

The following terms are important to understand FortiAP:

- **FortiAP** is the hardware used to aggregate the wireless connections on the LAN edge, providing different access modes, radio configuration capabilities, and all the current cutting-edge WiFi enhancements (depending on the model.)
- **FortiAP firmware** is the operating system, CLI, and control system of FortiAP.
- **Tunnel mode** is the default mode for a FortiAP. A FortiAP in tunnel mode uses a wireless-only subnet for wireless traffic and transports the traffic from the AP to the FortiGate in an encapsulated way.
- **Bridge mode** When a FortiAP is in Bridge mode, the Ethernet and WiFi interfaces are connected (or bridged), allowing wired and wireless networks to be on the same subnet. In essence, the WiFi traffic will be mapped with one or multiple VLANs on the FortiSwitches.
- **Segmentation or SSID** can easily be applied as the capability to create multiple VLANs and SSIDs. An SSID is a WiFi LAN identifier to separate different network segments, achieving a better network design and minimizing the spread of potential breaches at Layer 2. Each SSID can be used in Tunnelled or Bridge mode. FortiSwitch VLANs can be automatically populated in this case by using the embedded NAC to activate the port with the correct settings.

FortiSwitch

FortiSwitch can be adopted as a natural extension of SD-WAN to provide security on the wired LAN edge.

FortiSwitch is an essential cornerstone to the software-defined branch (SD-branch) that completes the SD-WAN architecture by enabling security into the access through FortiLink, consolidating all the connectivity in the branches, and enabling the management and power of the FortiAPs.


In addition to the above, the simplification of networking tasks, from the potentially complex topology designs to the lack of staff in the remote locations, by adding a layer of auto-discovery and automation allows the security teams to carry out the deployment in the branches seamlessly.

FortiSwitch facilitates and enhances network visibility as a first step in grabbing control of the network—under the umbrella of FortiGate, with FortiManager functioning as a single pane of glass.

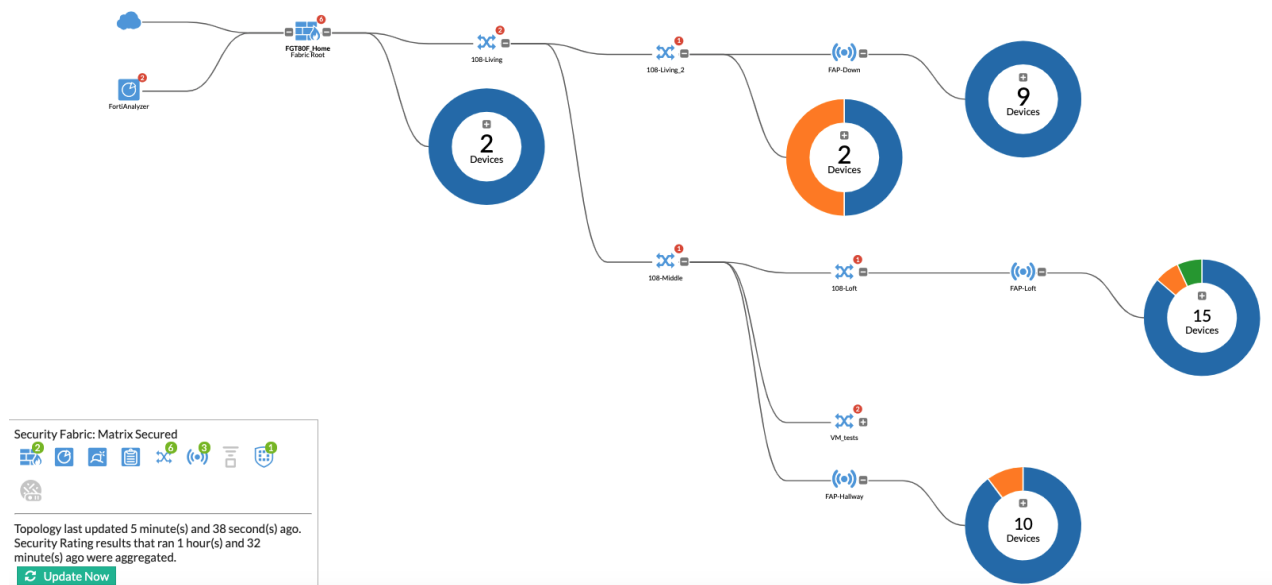
We will describe the FortiSwitch and FortiLink management setup, the simple provisioning of the Secure Access Layer as part of the architecture of the SD-Branch, and the implementation of the security use cases that complement the SD-WAN.

The key features of FortiSwitch are:

- **FortiLink management simplification:** With FortiLink, FortiSwitches are discovered and managed automatically from the FortiGate, automatically forming the topology and the Layer 2 aggregation, relieving the administrators from configuring low-level parameters, such as the Spanning Tree Protocol.



- **Single pane of glass:** From the FortiGate or FortiManager GUIs, you can access and configure the FortiSwitches managed by FortiLink, and obtain thorough information about the devices connected to them.



- **Security to the LAN:** security features associated with LAN environments—from basic VLAN segmentation to 802.1X authentication policies. VLANs become part of FortiGate's interfaces and can be easily integrated into Security Fabric policies.

FortiLink Interface

Managed FortiSwitch

FortiSwitch VLANs ☆

FortiSwitch Ports

Name	VLAN ID
vsw.flink-br1	1
Corporate_VLAN	10
Guest_VLAN	20
IoT_devices	30
OT_pre-prod	40
OT-production	60
vl_lan	1000

Policy & Objects ▾

IPv4 Policy ☆

Name

Incoming Interface

Outgoing Interface

Source

Destination

Schedule

Service

Action

Inspection Mode

Corporate-OT-pre-prod

fortinet (Corporate) ✕
Corporate_VLAN ✕
+

OT_pre-prod ✕
+

Corporate address ✕
+

FABRIC_DEVICE ✕
+

always ▾

HTTPS ✕
SAMBA ✕
Web Access ✕
+

☒ ACCEPT ☐ DENY

Flow-based Proxy-based

- **Power up you POE devices:** FortiSwitch models with PoE support lets you provide power to other devices, such as FortiAPs. PoE control and budget are also available from the FortiGate GUI. Current FortiSwitches support PoE, PoE+, and PoE++ (depending on the models).

Diagnostics and Tools - 108-Middle

108-Middle	
Name	108-Middle
Serial Number	S108EP4N17000237
Version	S108EP-v6.4.6-build470,210211 (GA)
Model	S108EP
FortiLink Interface	fortilink
IP Address	169.254.1.4
Join Time	2021/02/19 20:30:34
Status	Connected
Registration	Not Registered

Actions ▾ Edit

General Good

- 18% CPU Usage
- 60% Memory Usage
- 67 day(s) Connection Uptime
- Temperature
- 76% PoE Power Budget Remaining

Faceplate

See also [Important terms for FortiSwitch](#) on page 29.

Important terms for FortiSwitch

We are going to use the following terms in the rest of the document:

- **FortiSwitch** is the hardware used to aggregate the wired and wireless connections on the LAN edge, providing different layouts of physical ethernet or modular (SFP) ports and Power-over-Ethernet (PoE) capabilities, depending on the models.
- **FortiSwitchOS** is the operating system, CLI, and control system of the FortiSwitches.
- **FortiLink** is Fortinet's proprietary protocol that secures communications and implements the controls for configuring FortiSwitches from the FortiGate.
- **VLAN** or virtual local-area network is a smart virtual wire that interconnects devices of the same network. When managed from a FortiGate, the VLANs created on the FortiSwitches become network interfaces used in the Firewall Policies.
- **Segmentation** can create multiple VLANs to separate different network segments, thereby achieving a better network design, and minimizing the spread of potential breaches at Layer 2.

Secure SD-WAN solution

It is essential to distinguish between *Secure SD-WAN functionality* and the *Secure SD-WAN solution*. Secure SD-WAN functionality can be configured on any FortiGate device without requiring a separate license or additional products and components. In other words, any FortiGate device can provide this functionality in a completely autonomous manner, including traffic steering intelligence, monitoring, and of course, security.

This chapter will explain how to transform a group of autonomous devices providing local Secure SD-WAN functionality into the most critical element of your Secure SD-WAN solution. FortiGate devices can act as intelligent edge devices, providing secure connectivity across all your sites, cloud services, and the internet over the most optimal available path. The following chapters will teach you how to complete your Secure SD-WAN solution by centralizing its management (provisioning, monitoring, and reporting).

However, before we discuss the design of the Secure SD-WAN solution, we must spend some time describing the *SD-WAN functionality* itself. This section includes the following topics:

- [Technical background on page 30](#)
- [Design principles on page 33](#)

Technical background

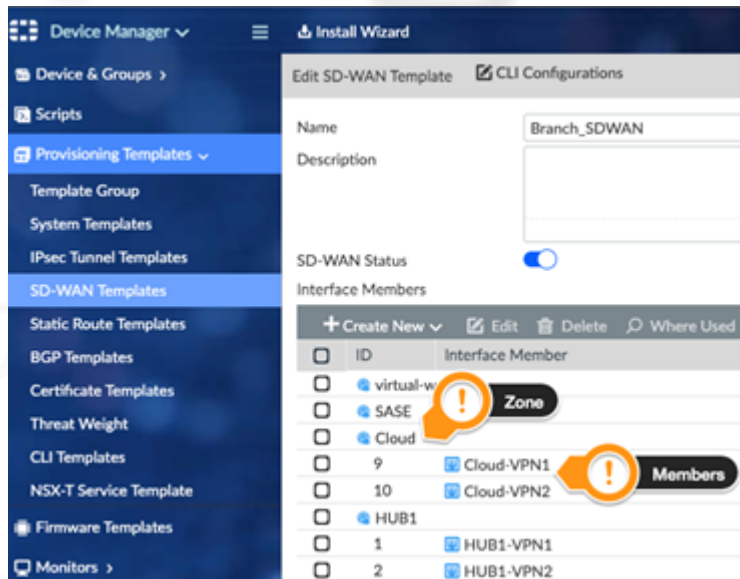
The technical background covers the following topics:

- [SD-WAN configuration on page 30](#)
- [SD-WAN routing logic on page 33](#)

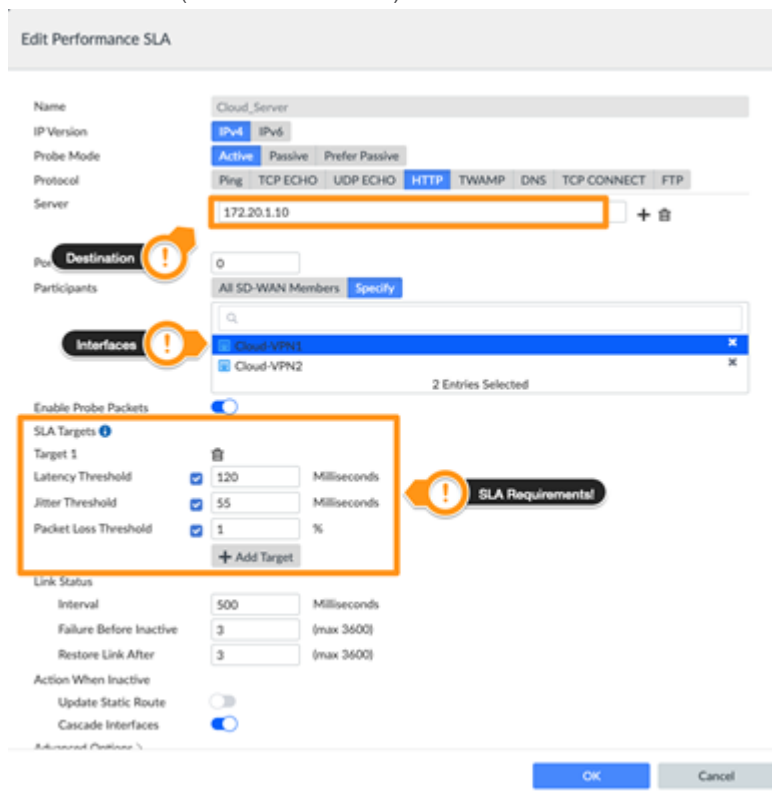
SD-WAN configuration

Fortinet SD-WAN configuration includes the following main steps:

1. **SD-WAN interface members** define your *SD-WAN bundle*. They are the interfaces that will be controlled by SD-WAN and where traffic can potentially flow. Almost any interface supported by FortiGate devices can become an SD-WAN member (including physical ports, VLAN interfaces, LAGs, IPsec/GRE/IPIP tunnels, and even FortiExtender interfaces). Often it will include both your underlays and overlays, but this is not a requirement. For example, you can configure the overlays to be your SD-WAN members while keeping the underlay outside. We will look into these options in the design examples. For convenience, the SD-WAN members are grouped into SD-WAN zones.



2. **Performance SLA** are the health-check probes used by the edge devices to actively measure the health of each available path. You can define what server to probe and what protocol to use (including Ping, HTTP, TCP/UDP Echo, TWAMP, or DNS). Each probe will measure latency, jitter, and packet loss percentage over the configured subset of the SD-WAN members. In addition, you can configure multiple SLA targets for each probe. Together, these metrics will allow SD-WAN to compare the health of different available paths, and even determine which paths are acceptable for a particular application and which are not (called *out of SLA*).



3. **SD-WAN rules** combine all the elements. These are the actual set of business rules used to steer a particular application to a specific SD-WAN member while considering its current health and SLA status. Each rule has the following logical parts:
 - **Matching Criteria** defines what applications or what kind of traffic will match this rule. We can match based on a large variety of inputs, including:


- IP Address
- Applications
- Internet Service Database (ISDB)
- User Identity
- DSCP/ToS fields
- Route Tags

Edit SD-WAN Rule

Name: Branch_to_Cloud-Server

IP Version: IPv4

Source

Source Address: [Search] **Source** 

Branch-NET
IP/Netmask: 10.1.0.0/255.255.0.0
1 Entry Selected

Users: Click here to select

User Groups: Click here to select


Destination

Address: [Search] **Destination** 

Cloud-Server
IP/Netmask: 172.20.1.10/255.255.255.255
1 Entry Selected


Route Tag: 0

Protocol: TCP UDP **ANY** Specify 0


Type of Service: 0x00 Bit Mask: 0x00 **Steering Strategy** 

Outgoing Interfaces

Strategy: Manual Best Quality **Lowest Cost (SLA)** Maximize Bandwidth (SLA)

Interface Preference: [Search] **Interface(s)** 

Cloud-VPN1
Cloud-VPN2
2 Entries Selected

Required SLA Target: [Search] **Performance SLA** 

Cloud_Server#1
HTTP: 172.20.1.10; Latency: 120ms, Jitter: 5ms, Packet Loss: 1%
1 Entry Selected

Advanced Options >

OK Cancel

- **SD-WAN Strategy** defines the logic applied to select one of the SD-WAN members to steer this traffic. The following strategies can be configured:
 - **Best Quality**—select an SD-WAN member with the best measured quality.
 - **Lowest Cost (SLA)**—select the cheapest SD-WAN member that meets a given SLA target.
 - **Maximize Bandwidth (SLA)**—load-balance across all SD-WAN members that meet a given SLA target.
 - **Manual**—manually specify an SD-WAN member to select.

SD-WAN Rules

ID	Name	Source	Destination	Criteria	Members
7	Cloud_Applications	Branch-NET	Cloud_Applications	Cloud_HCF1	Cloud-VPN1 Cloud-VPN2
5	Critical_DIA	Branch-NET	Critical_Applications	Latency (Internet)	port1 port2
4	Critical_Video	Branch-NET	Critical_Video	Passive_SLAF1	port1 port2
3	Branch_to_Cloud-Server	Branch-NET	Cloud-Server	Cloud_Server#1	Cloud-VPN1 Cloud-VPN2
1	Branch_to_DC	Branch-NET	Datacenter Branch-NET	HUB1_HCF1 HUB2_HCF1	HUB1-VPN1 HUB1-VPN2 HUB2-VPN1 HUB2-VPN2
2	Non_Critical_DIA	Branch-NET	rfc-1918	Latency	port2
	sd-wan	ALL	ALL	Source IP	ALL

The SD-WAN rules probably remind you of the Firewall rules to some extent, and, indeed, many of the same matching criteria are used. The SD-WAN rules are also evaluated in the order of their configuration—just like

Firewall rules. But they serve two complementary goals (which will be discussed in more detail in the next chapter):

- Firewall rules define how to secure a particular application, should a particular path be selected.
- SD-WAN rules define how to select a particular path for a particular application.

Having both rulesets rely on the same inputs (such as Application Control Database, Internet Service Database [ISDB], same User Identity providers, and so on) significantly improves integration between different pillars and the consistency of the overall solution.

SD-WAN routing logic

Once configured, SD-WAN takes the responsibility of intelligent traffic steering. But how does it interact with the traditional routing subsystem?

The following main rules apply by default:

1. SD-WAN rules are matched only if the best route to the destination points to SD-WAN.

The best route to the destination must point to any SD-WAN Member—not necessarily the one selected to forward the traffic. This check allows you to easily fit SD-WAN functionality into your existing network topology without disrupting services that are not supposed to be handled by SD-WAN. For example, you may have an out-of-band management network or a group of sites that have not (yet) migrated to SD-WAN. If the best route to the destination does not point to your SD-WAN *bundle*, the traffic will be handled by *conventional* routing.

2. SD-WAN member is selected only if it has a route to the destination.

This check happens at a later stage when an SD-WAN rule is already matched and evaluated. Based on the configured strategy, one of the listed SD-WAN members will be preferred. But the traffic will only be forwarded via that member if there is a route to the destination through that path. Otherwise, the member will be skipped, and the next optimal member will be checked.



This does not have be the best route this time!

As you can see, routing information serves as one of the inputs for SD-WAN intelligence.

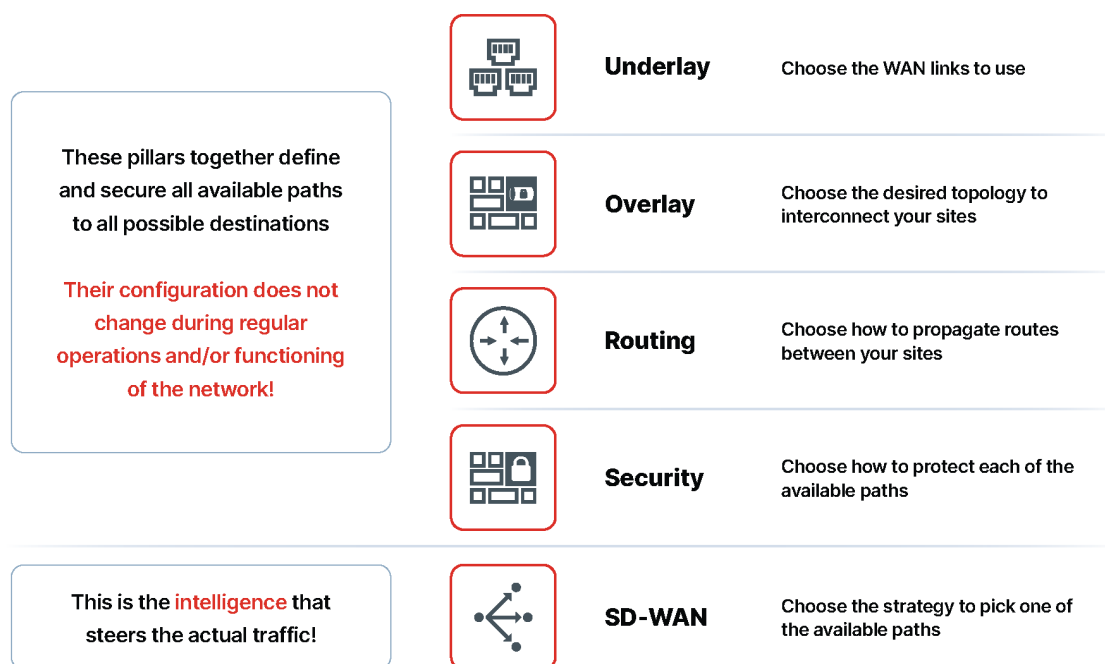
The above behavior can be overridden: It is possible to configure an SD-WAN rule that will completely bypass route lookup. This option can help in specific scenarios, but it must be used with care!

Finally, what happens if none of the SD-WAN rules can forward the traffic? This can happen either because none of the rules could match the traffic or because none of the Members of the matching rules had a route to the destination. In this case, the traffic is forwarded using *conventional* routing (often called an implicit rule).

This concludes our overview of the SD-WAN functionality on FortiGate devices. Let us now turn to our main topic and see how we can build a complete Secure SD-WAN solution!

Design principles

When designing your Secure SD-WAN Solution, we recommend that you utilize the following **Five-Pillar Approach**:



As you can see in the above diagram, the goal of the first four pillars (**Underlay, Overlay, Routing, and Security**) is to define and secure all available paths to all possible destinations. In other words, at this stage, there is still no decision about where specific traffic will flow, but all the edge (CPE) devices are aware of all the options. These four pillars should not require human intervention during regular operations and network functions. And this is despite the fact that the set of available paths and destinations in the network can change dynamically due to network failures, planned migrations, or even changes in traffic patterns.

The Zero-Touch nature of the first four pillars is achieved using two dynamic technologies that, once configured, do not require further operator intervention:

1. Our dynamic tunneling technology—**Auto-Discovery VPN (ADVPN)**—automatically builds direct IPsec tunnels between the sites willing to communicate. These tunnels (also called shortcuts) immediately become part of the overlay topology of your SD-WAN solution. And once the communication between the sites is over, these shortcuts can be automatically torn down to free up the resources.
2. We also use industry-standard dynamic routing protocols (**BGP** being a typical choice), to exchange currently available paths between sites, automatically adapting to all topology changes.

Once all available paths to all possible destinations are defined and secured, it is time for the fifth pillar (**SD-WAN**). This intelligence decides which available path will be selected *at a given moment and for a given application*. This pillar is a combination of administratively configured business rules and dynamically measured metrics.

Note that all the control plane technologies mentioned above (ADVPN, BGP, and SD-WAN) are distributed across all the edge (CPE) devices, making the overall design highly scalable.

Before we move on to design examples, let us discuss each of the five pillars in more detail:

- [Underlay on page 35](#)
- [Overlay on page 35](#)
- [Routing on page 36](#)
- [Security on page 37](#)
- [SD-WAN on page 38](#)

Underlay

First, you must decide what **underlay** links you will use to connect all participating sites and the public internet. Do you have multiple internet connections? Or an internet connection and an MPLS link? Or will it be a broadband internet connection and an LTE modem?

How will edge devices get their IP addresses—via DHCP or static configuration? Will there be a need for VLAN tagging? For Link Aggregation?

The same questions also apply to the LAN side. How will the local LAN network connect to the edge device? Are there any additional services needed. For example, will the edge device act as a DHCP server for the local network?

Since all edge devices are full-featured FortiGate devices, the range of possibilities is extensive. While each site can, in principle, be designed and configured differently from the others, we highly recommend defining a limited number of *groups of sites* with identical configurations within each group. This will simplify provisioning and the operation of your SD-WAN solution.

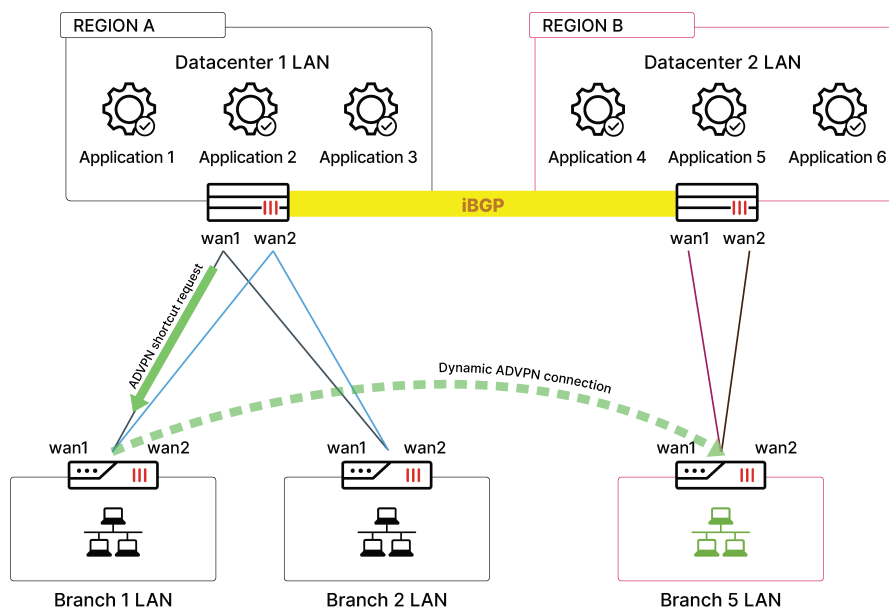
Overlay

Second, you must decide on the topology to interconnect your sites. In most cases, you will build **IPsec overlays** over all the underlay transports to most likely form a set of hub-and-spoke topologies. This way, you can secure your corporate (site-to-site) traffic, and provide confidentiality, integrity, and mutual site authentication, as expected from an industry-standard IPsec suite.

Hub-and-spoke topologies are highly scalable, and they have a crucial zero-touch property: When adding or removing a spoke, the configuration of all other devices remains untouched. Hub-and-spoke topologies can also be enhanced with redundancy options (such as dual-hub). They can be extended to multiple regions (multi-regional hub-and-spoke topologies interconnected together) for large-scale deployments.

ADVPN—our dynamic tunneling technology—can be enabled in your hub-and-spoke topologies. As mentioned earlier, ADVPN can dynamically build direct spoke-to-spoke tunnels (called *shortcuts*) when they are needed. It preserves the zero-touch property of hub-and-spoke while providing advantages of direct site-to-site communication without bottlenecks.

For multi-regional deployments, you can optionally allow cross-regional ADVPN shortcuts, making your topology even more dynamic. We walk you through these options in the design examples section.



To conclude, although other overlay topologies can be used (such as a static hub-and-spoke or even a full-mesh), we recommend ADVPN as the most generic, dynamically adjustable topology for your overlays.

It is worth highlighting at this point that overlays are *optional* in our SD-WAN solution. The traffic can be steered both to the underlays and the overlays, with broadly similar SD-WAN functionality. We return to this topic when we discuss the SD-WAN pillar. See [SD-WAN on page 38](#).

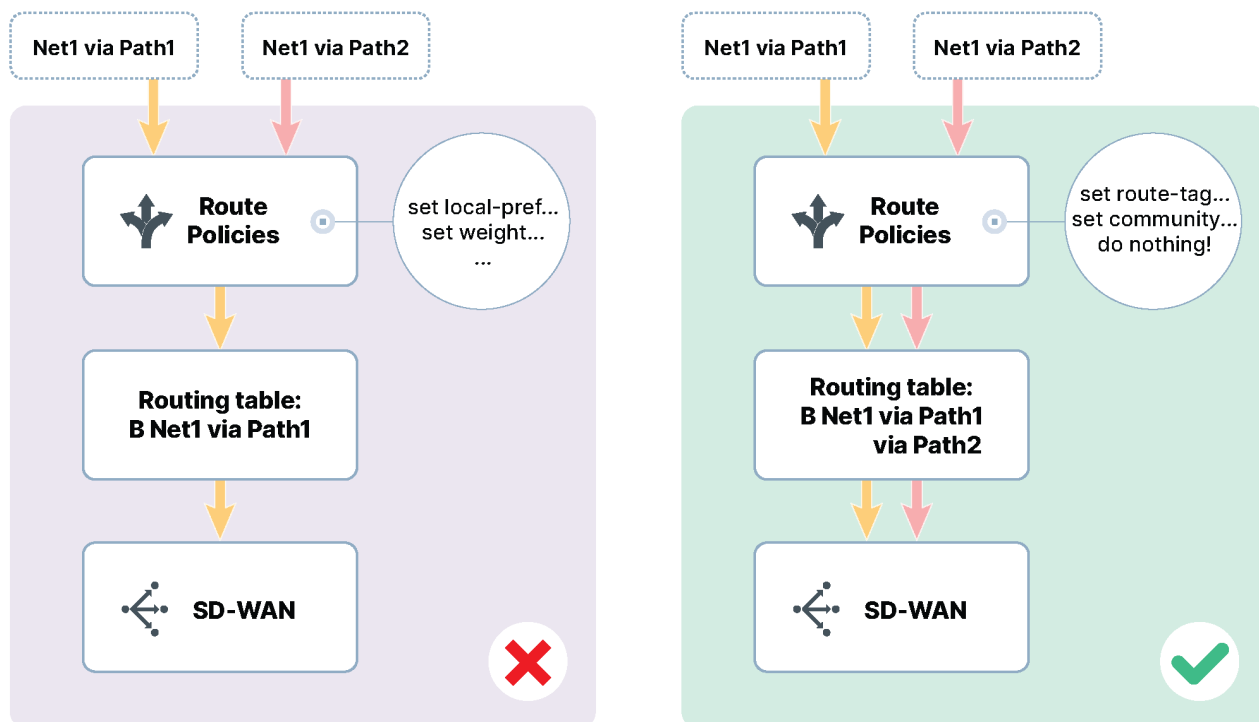
Routing

The overlays provide us with multiple paths between the sites (over different underlay transports). Still, we must also ensure that all edge devices have the correct routing information needed to use these paths. We recommend using BGP to exchange routes between all sites over the overlays.

BGP fits well into hub-and-spoke overlay topologies, and it is also the recommended routing protocol to use with ADVPN. As we will show in design examples, the hubs will act as BGP route reflectors (RR) so that the spokes will not have to peer directly with each other—not even over ADVPN shortcuts! This design is in-line with the zero touch strategy: once again, when adding or removing a spoke, the BGP configuration of all other devices remains untouched.

A crucial difference between a traditional design and our SD-WAN solution is in the *role* of the routing pillar. In a conventional design, routing oversees the steering of traffic. It is, therefore, the responsibility of routing to select the best path out of all available options. Multiple route policy techniques can be used to achieve this—some are protocol-agnostic (for example, weight), and others are protocol-specific (for example, BGP local-preference, MED, AS_PATH prepending, and so on). While all these techniques remain available on a full-featured FortiGate edge device, we must recall that our goal is *only to learn about all available paths to all possible destinations*!

Remember that the duty to steer the traffic in our solution is delegated to the fifth pillar—the SD-WAN. Therefore, it is (generally) not recommended to apply any route policy techniques to the routes learned via BGP. Rather than selecting a single best route, we would like to end up with equal-cost multi-path (ECMP) routes to all remote sites via all available overlays.



Security

By now, we have learned about all available paths to all possible destinations. It is time to define how to secure each of these paths. Here again, we are not deciding (yet) which of the paths will be selected. We are only deciding how to secure the traffic *should a particular path be chosen*.

Quite often, different security features must be applied to different paths. The most common example is the difference between direct and remote internet access. In the former case, the traffic breaks out directly from the edge device (through one or more underlay links), making it crucial to apply the necessary level of security before it leaves the site boundaries. In the latter case, the traffic might undergo additional security inspection in the central location or use a cloud-based security solution before breaking out to the public internet. As a result, the edge device has to apply a different set of security features, depending on which of the two internet access methods was selected for a particular session.

We achieve this granular security in our solution by grouping different interfaces into *SD-WAN zones* and defining firewall rules on a per-zone basis. In the above example, we would define two SD-WAN zones named *overlay* and *underlay*, and we would define separate firewall rules for the internet traffic exiting through each one of them.

The general principle that you should follow when preparing the firewall ruleset for your SD-WAN solution with hub-and-spoke topology is that security should be applied at the originating site. To better understand the rationale behind this principle, consider the following:

- When using ADVPN, spoke-to-spoke traffic will eventually flow via a direct shortcut, completely bypassing the hubs. Therefore, it is crucial to apply all the necessary security inspections at the spokes.
- With direct internet access, client traffic leaves the boundaries of your SD-WAN solution right at the edge of the originating site. Therefore, it is the only opportunity to properly secure that traffic.

Based on the above, the hubs in your topology will generally have a *permissive* policy for spoke-to-spoke traffic, as they act only as transit devices for that traffic. However, we should highlight that the hubs may

still be responsible for securing other types of traffic. For example, they could apply additional inspection to incoming traffic to better secure the workloads hosted behind them or for remote internet access.

Remember that the actual decision on where the traffic flows will be taken by the fifth pillar (the SD-WAN), *at a given moment and for a given application*. But whenever the decision is made, the firewall rules will be in place to secure the traffic appropriately.

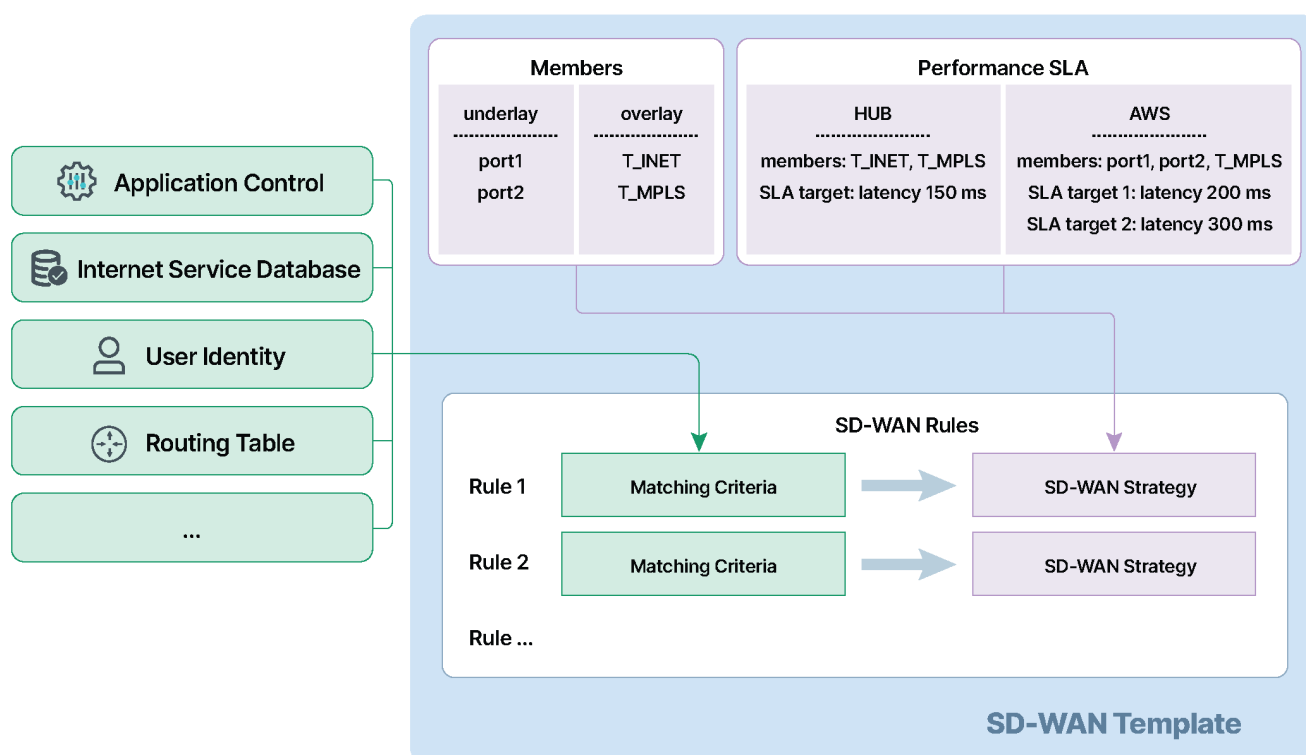
SD-WAN

And now we are reaching the fifth pillar—**SD-WAN**. In a nutshell, this is the intelligence that will be applied to each outgoing session to determine the optimal path at a given moment. It will consider all the available paths to the requested destination, compare their measured health, and then apply a business strategy configured for a particular application to make the optimal choice. Health measurement continues in real time. If the conditions change, both new and existing sessions can quickly switch over to another path.

As we have covered earlier, SD-WAN configuration typically consists of the following elements:

- SD-WAN interface members
- Performance SLAs
- SD-WAN rules

When using FortiManager to configure your SD-WAN solution, all the above elements are conveniently packed into an SD-WAN *template* that can be applied to (a group of) your sites. As usual, although you could apply an individual SD-WAN template to each edge device, we highly recommend grouping similar sites, and applying a single SD-WAN template to the entire group. This will significantly simplify your operations, and make your SD-WAN solution consistent.



For example, you could have a single SD-WAN template for all your branch offices and another SD-WAN template for your central datacenters. This would allow you to apply changes quickly and consistently

without the need to reconfigure each site individually. And this is one of the most important goals of an SD-WAN solution!

This is the main point of focus for your network operations. You can adjust the relevant SD-WAN templates to instruct your edge devices to accommodate the changes whenever business requirements change. The configuration of the other four pillars will typically remain unchanged.

Remember that the edge devices already know about all available paths to all possible destinations, and they dynamically adapt to the topology changes. The only input that cannot be obtained without operator intervention is the actual set of business rules to be applied.

For the optimal configuration of your SD-WAN solution, you must understand and use the following recommended principles:

- The *originating site* should take the steering decision—that is, by the SD-WAN rules of the edge device located at the site originating the session. If the decision is to break out locally, the traffic will leave the boundaries of the SD-WAN solution. Otherwise, the traffic will flow via one of the active overlays. Hence it will pass through one or more additional FortiGate devices that are part of your solution. All those devices are expected to “respect” the SD-WAN choice made by the originating site. For example, in a hub-and-spoke topology, if the originating site has selected an overlay over MPLS transport as its next hop to the hub, the hub should prefer using the overlay over MPLS transport to forward the traffic further toward the destination site. We also call this property the *overlay stickiness*.



It has particular importance for ADVPN since shortcut offers follow the routing decisions. If the traffic does not preserve the overlay end-to-end, this can cause an attempt to establish a shortcut between two physically disconnected transports, such as the internet and MPLS. This attempt will, of course, fail!

- The same applies to the reply traffic as well. We recommend preserving symmetrical traffic flows so that reply traffic returns via the same overlays from which the traffic in the original direction arrives. While it is possible to configure FortiGate devices to support asymmetrical replies, we advise keeping the default configuration that respects the choice of the session originator.
- As can be derived from the above two principles, transit devices (such as hubs) generally do not require SD-WAN configuration since they do not act as originating sites for traffic. They must only *respect* the steering decisions made by other sites in both directions.
- >We discuss more principles in the context of complete design examples in the following sections.

To conclude, the SD-WAN pillar allows you to define a fine-grained set of business rules to control your application traffic. It operates on top of the four other pillars—Underlay, Overlay, Routing, and Security—each of those by itself offering a wide range of possibilities to fit your needs. This degree of flexibility is no wonder since all the edge devices are full-featured FortiGate devices. But it is precisely for this reason that planning your design carefully and following our proven best practices is crucial to building a highly scalable and easy-to-operate Secure SD-WAN solution!

Architecture and design

This section will cover some of the most common SD-WAN architectures and considerations when planning your network. It's important to note that there are many different ways to design your network that are not covered in this document. This section will introduce you to some of the most commonly used architectures to cover many different use cases.

Design	Business Application Location (s)	Common Use Cases:
Single datacenter	Single datacenter location	Private workloads and applications where stability is preferred
Multiple datacenters	Geo-redundant datacenter locations	Redundant, private workloads where datacenter location is preferred, and other locations are backups
Multi-region datacenters	Geo-redundant datacenters across different regions	Connectivity to other regions for some applications and services
Cloud on-ramp* (for static cloud environments)	Static cloud infrastructure	Connectivity to cloud services with a static gateway
Cloud on-ramp* (for dynamic cloud environments)	Dynamic cloud infrastructure	Connectivity to cloud services with a dynamic gateway

* Designs like direct internet access (DIA) and cloud on-ramp are typically used in conjunction with other designs. For example, The multiple datacenter designs can include DIA and cloud on-ramp.


This section includes the following topics:

- [Single datacenter \(active-passive gateway\) on page 41](#)
- [Multiple datacenters \(primary/secondary gateways\) on page 52](#)
- [Multi-region datacenters on page 65](#)
- [Supplemental designs on page 69:](#)
 - [Direct internet access on page 70](#)
 - [Cloud on-ramp on page 72](#)

See also [Cloud on-ramp for static environments on page 73](#) and [Cloud on-ramp for dynamic environments on page 74](#).

Single datacenter (active-passive gateway)

Traditionally referred to as *hub and spoke*, this design is the fundamental building block of our solution. The more advanced multi-datacenter and multi-region examples will essentially be extensions of the single datacenter design. In this design, the SD-WAN Gateway (or sometimes referred to as the *hub*) acts as a headend into the business application or private workload. SD-WAN gateways can be located in a single datacenter or central office, and typically provide connectivity for remote locations.



This section includes the following topics:


- [Gateway components on page 41](#)
- [SD-WAN considerations on page 50](#)
- [Security considerations on page 51](#)

Gateway components

The SD-WAN gateway can provide one or multiple services:

- Act as the IPsec dialup server for branch locations
- Provide centralized routing information and orchestrate dynamic branch-to-branch communication (ADVPN)
- Protect the datacenter resources and private workloads by utilizing FortiGate Next-Generation Firewall services
- Provide remote internet breakout for branch locations

Like other SD-WAN architectures, the core functionality of the gateway is to provide connectivity and routing information between branch locations and private resources in the datacenter or the cloud



This section includes the following topics:

- [Redundancy on page 42](#)
- [IPsec overlays on page 44](#)
- [Route exchange on page 45](#)
- [ADVPN on page 46](#)
- [Traffic flow on page 47](#)

Redundancy

In this design, the SD-WAN gateway may offer intra-site redundancy by joining two or more devices into a high availability (HA) cluster. FortiGate HA offers several solutions for adding redundancy in the case where a failure occurs on the FortiGate, or is detected by the FortiGate through monitored links, routes, and other health checks. These solutions support fast failover to avoid lengthy network outages and disruptions to your traffic.

FortiGate HA options:

- [FortiGate Cluster Protocol \(FGCP\)*](#)
 - Active/passive
 - Active/active
- [FortiGate Session Life Support Protocol \(FGSP\)](#)
 - Session and configuration synchronization across standalone FortiGate or HA clusters

* In this document, we will focus on utilizing the FortiGate Cluster Protocol (FGCP) on our SD-WAN gateways to accomplish high availability.



There are more advanced use cases and scenarios where FGSP may be used to horizontally scale across local or geo-redundant location. This is beyond the scope of this document.

Intra-datacenter gateway redundancy

For most use cases, it is generally recommended to utilize active-passive HA for SD-WAN gateways at a datacenter or HQ location. If active-active is desired, it will not change our overall SD-WAN design outlined below. Both HA modes will be designed in the same matter as described in this section.

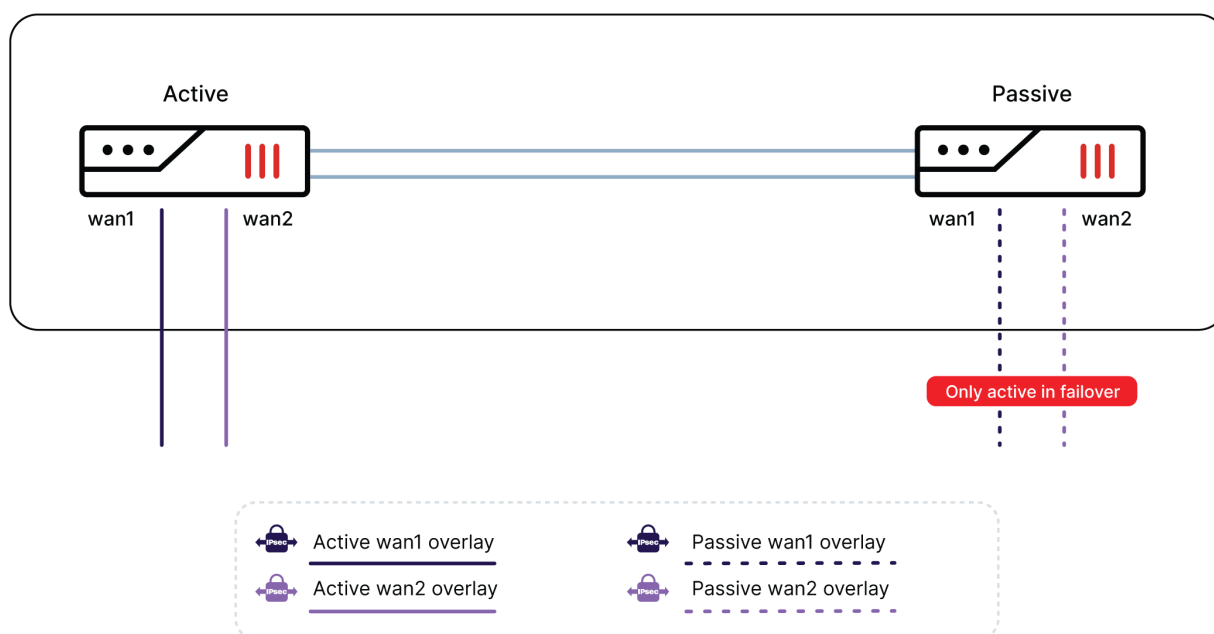
In active-passive HA mode, there are at least two devices in the cluster, with only one device acting as the primary device. To the rest of the network, including remote branch locations, the active-passive cluster appears to be a single device that shares a *floating* IP address between the active members. Remote branch locations terminate their overlays to the active device in the cluster.

SINGLE DATACENTER (ACTIVE-PASSIVE GATEWAY)

The active-passive gateway model provides redundancy inside the datacenter, while operating as a single device to outside resources. Branch locations terminate their overlay connections to the active member, while being unaware the gateway is a cluster with multiple members.

Gateway Redundancy	Benefits	Considerations
Active-passive HA cluster	<ul style="list-style-type: none">• Intra-datacenter redundancy• Logically seen as a single device on the network• Simple setup	<ul style="list-style-type: none">• Does not provide performance improvement on security inspection
Active-active HA cluster	<ul style="list-style-type: none">• Intra-datacenter redundancy• Logically seen as a single device on the network• Performance improvement on security inspection	<ul style="list-style-type: none">• More complex troubleshooting requirements due to the nature of active-active load balancing

For more information on HA design and consideration, see the latest [FortiOS Administration Guide](#).



In the event of a failover at the datacenter between active-passive members, traffic should switch over to the next healthy member.

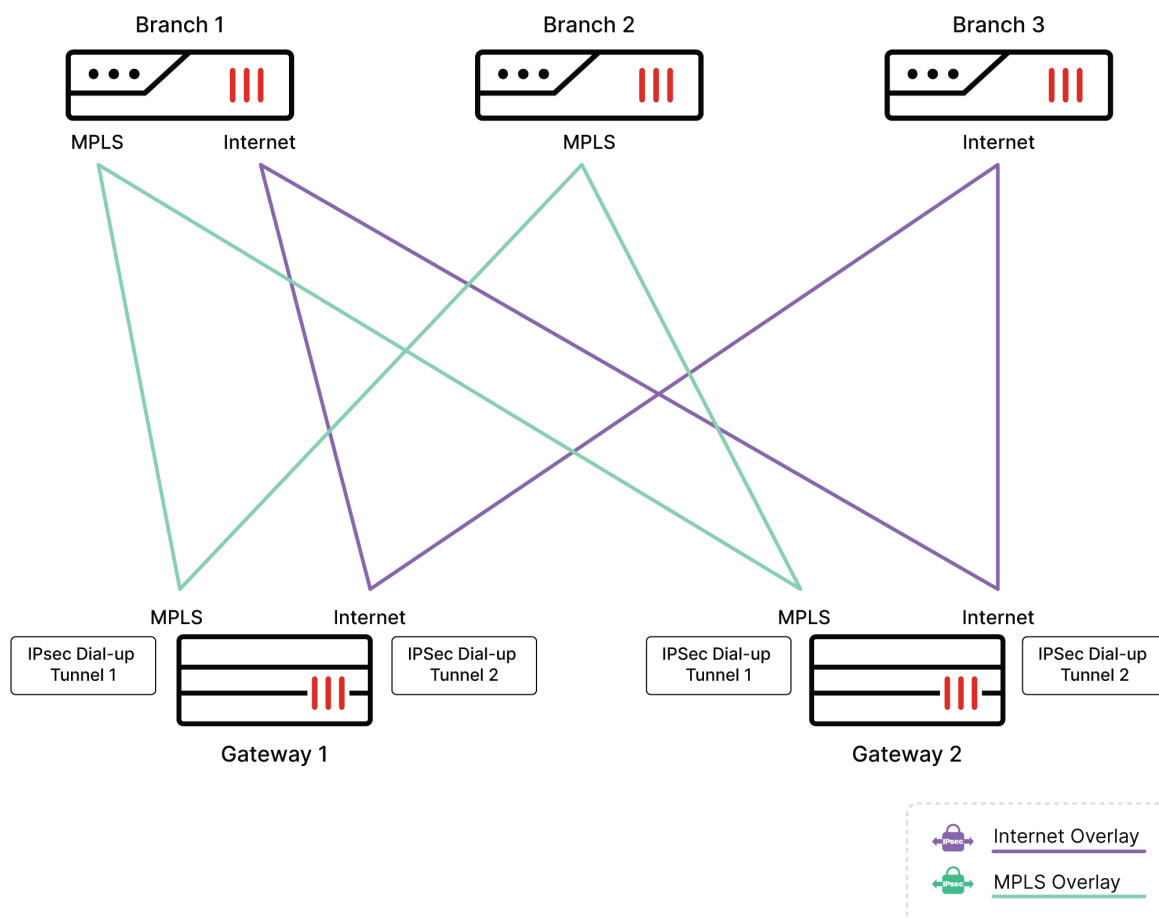
To minimize or eliminate traffic interruption during failover, it is recommended to consider the following:

- Enabling BGP graceful restart on the gateway and branch
- Enabling `route-ttl` on the HA settings to ensure the FortiGate cluster maintains the cached routes during failover
- Fine-tune BGP timers as necessary

For more information on these three components, see this [KB article](#).

IPsec overlays

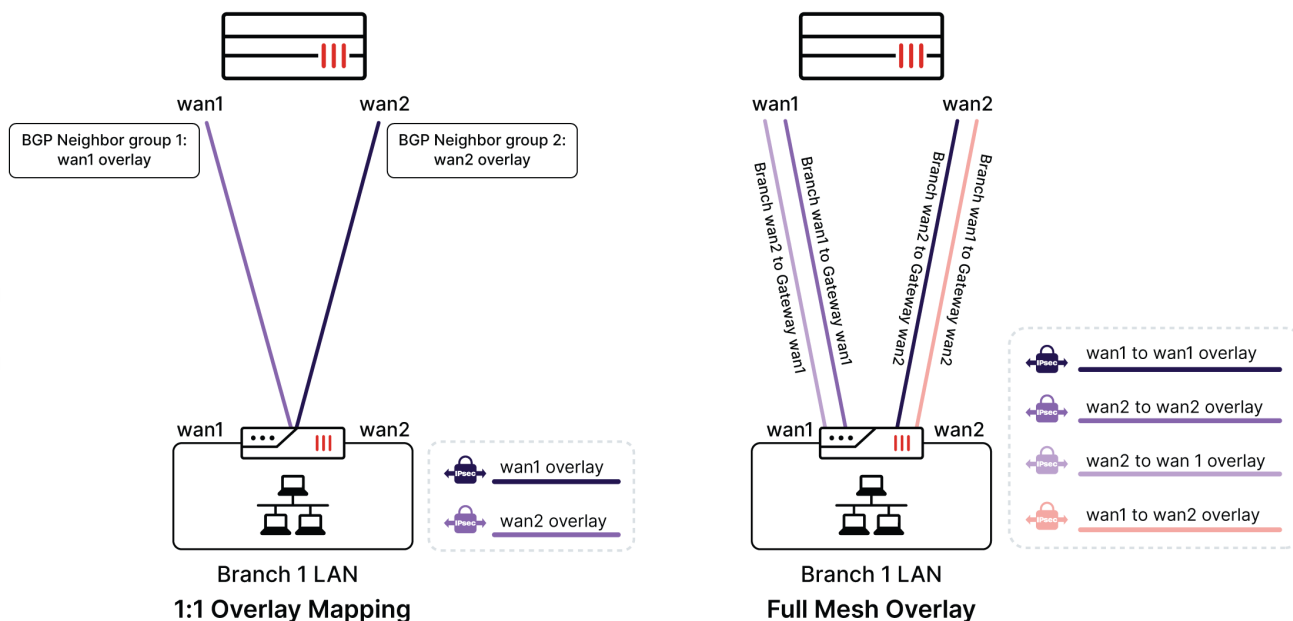
The SD-WAN gateway acts as a dial-up IPsec server for the spokes, having a separate dial-up IPsec endpoint terminated on each underlay interface. Branches will typically build overlays over all available wan ports to have multiple paths available to the gateway. However, it can also happen that some of the branches do not have a similar wan transport. Hence, they will be able to connect only to a subset of the overlays.



When considering the IPsec overlay design between the branch locations and the gateway, it is important to determine how redundancy should occur between all available links.

Consider the following options:

- **One-to-one overlay mapping per underlay:** in this design, each branch underlay terminates a new IPsec tunnel to one—and only one—gateway underlay. This is the most common overlay design, and simplifies our configuration, but also provides less redundancy than the subsequent full mesh.
- **Full mesh overlay mapping:** in this design, each branch underlay terminates a new IPsec tunnel to each WAN underlay on the gateway. This design provides more available paths for traffic to flow through during an outage, but can add complexities to our design. This design is only recommended if full-mesh redundancy is required.



Route exchange


The spokes establish separate IBGP sessions to the gateway over each overlay. The BGP Neighbor Group feature is used on the gateway for this peering. Each spoke then advertises its local site prefix(es) over each of the IBGP sessions. The gateway acts as a BGP Route Reflector (RR), readvertising the prefixes to all other spokes when ADVPN is used. Additionally, the gateway advertises its prefixes (such as the datacenter LANs) to each branch location. At the end of this process, all the sites exchange their routes over all available overlays.

Routing table for VRF=0

```

B      10.1.1.0/24 [200/0] via 10.10.10.1
                        [200/0] via 10.10.11.1
B      10.1.2.0/24 [200/0] via 10.10.10.2
                        [200/0] via 10.10.11.2

```



Additional Routing Notes

- Routing inside the datacenter is typically handled by BGP or OSPF. Datacenter networks are readvertised to the branches through BGP if needed.
- IBGP sessions are terminated on the IPsec overlays. Hence they are using the tunnel IP addresses as BGP next-hops (NH). This requires IP addresses to be configured on the tunnel interfaces. The hub can automatically allocate tunnel IP addresses to the spokes using the *IKE Mode Config* feature to simplify provisioning and administrative overhead.
- Since the spokes establish separate IBGP sessions with the hub over each overlay, there are multiple BGP routes for each prefix. To keep all the routes available, the following two BGP features must be enabled on all participating devices (hub and spokes):
 - *BGP Multipath* ensures that all the available routes are installed into the routing tables.
 - *BGP ADD-PATH* ensures that the hub between the spokes reflects all available routes.

ADVPN

For the correct operation of ADVPN, it is required to preserve all sites' prefixes unchanged, including their original BGP next-hop values. Hence, it is impossible to replace the specific routes with summaries (unlike in a static hub-and-spoke topology). Hence, the BGP RR function is mandatory: the gateway must reflect the original routes between the spokes without altering them.

FortiOS 6.4 and earlier:

- We have already mentioned the critical property of overlay stickiness that we must guarantee for proper ADVPN shortcut creation. For example, if spoke-1 sends traffic to spoke-2 using an internet overlay through the hub, the hub must select the same internet overlay for the second half of the path. Failing to

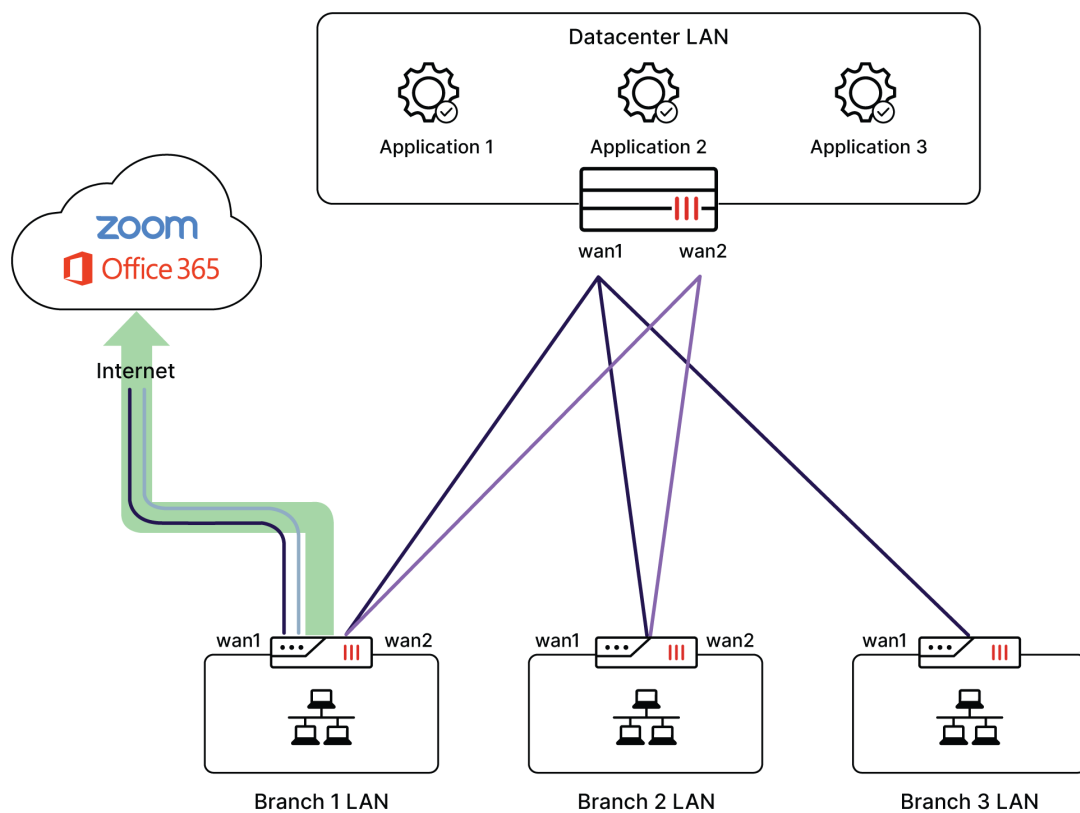
preserve the overlay might result in an attempt to create an ADVPN shortcut between two physically disconnected transports (such as the internet and MPLS), and this attempt would, of course, fail. The overlay stickiness is achieved using policy routes (PBR) on the hub.

Traffic flow

Once all the routes have been distributed across all the sites, the application traffic flow can be controlled by SD-WAN rules according to the design principles described in the previous chapter. SD-WAN rules may dictate how traffic is steered based on the business requirement and desired redundancy.

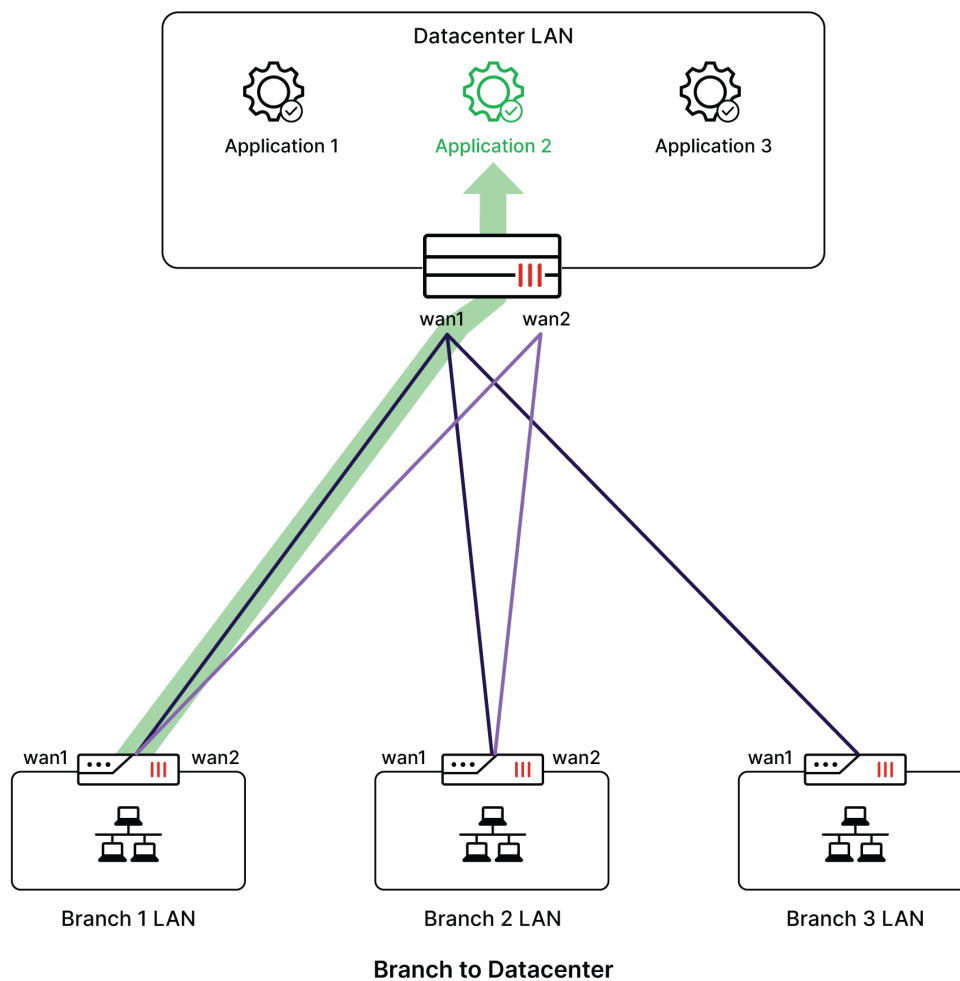
- **Direct internet access (DIA):** used when local internet breakout at a location is required. In this scenario, the business application(s), such as a SaaS application or website, is located on the internet, and the SD-WAN appliance is needed to decide the best path between multiple WAN links.

Traffic is routed directly to the internet by using the preferred method in the SD-WAN rule.

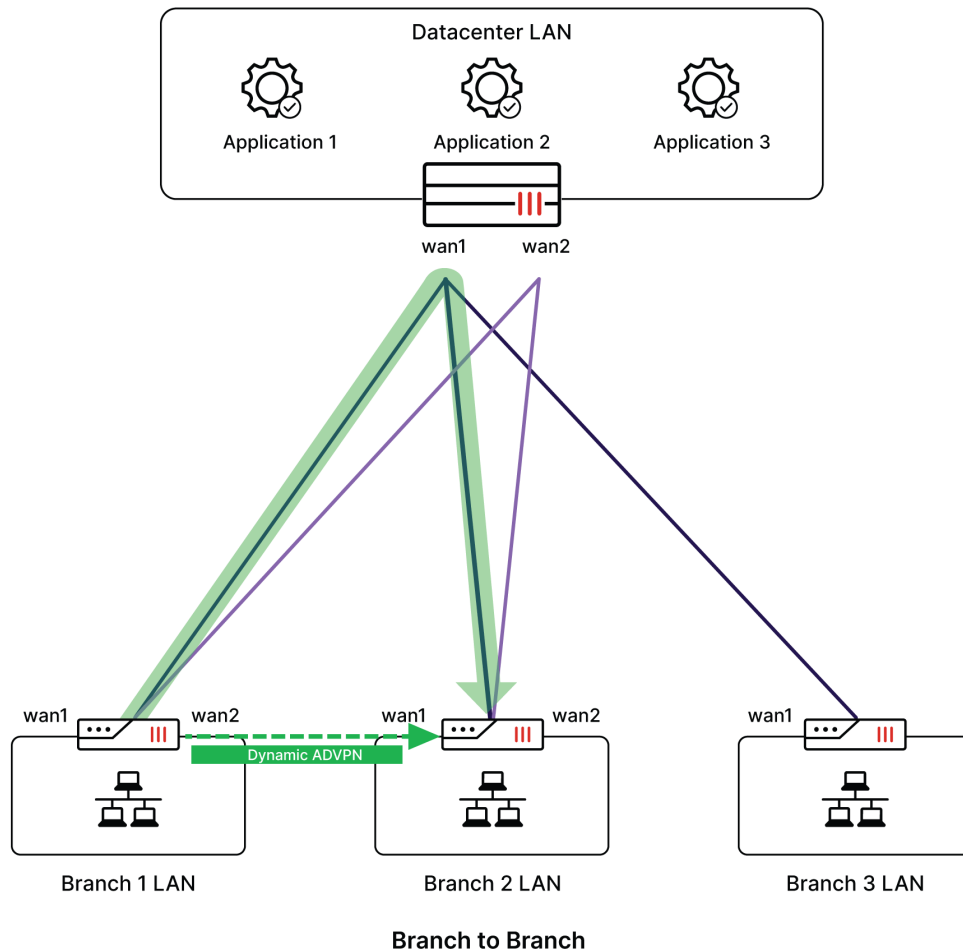


Direct Internet Access (DIA)

- **Branch to datacenter:** used when branch users require connectivity to an application or workload that is located behind the gateway in the datacenter. The branch SD-WAN device should monitor all available overlay links, and choose the best path according the business requirement.




- **Branch to Branch:** used when branch-to-branch communication is required. Traffic can flow through the SD-WAN gateway or dynamically between branch locations by means of ADVPN. Auto Discovery VPN, or ADVPN, is made possible by the SD-WAN gateway providing the routing and IPsec information to the branch where the first request originates. In this case, only the first few packets will flow through the hub, until an ADVPN shortcut is built.



- Remote internet access (RIA):** used when internet traffic from the branch is backhauled to the gateway for inspection and external breakout. RIA is typically used when a private link (such as MPLS) to the gateway is available as an alternative path to the internet for branch locations. In situations where the local internet degrades, traffic can be backhauled through the private link and broken out at the Gateway.

Another use for RIA exists where security inspection is required for internet traffic. For locations where local security inspection is not possible, traffic can be offloaded to the SD-WAN gateway and inspected before being sent to its destination.



SD-WAN considerations

SD-WAN Member	SD-WAN Zone	Performance SLA	SD-WAN Rule	Firewall Policy
All overlay interfaces to the single Datacenter location	Overlays should be grouped by device and location	Health check server: Business critical applications and resources	Dependent on business intention and availability requirements.	References the SD-WAN zone(s) and appropriate security inspection
Example: overlay1_wan1, overlay2_wan2	Example: Datacenter (overlay1_wan1, overlay2_wan2)	Example: Health-check: Datacenter_Server Members: overlay1_wan1, overlay2_wan2	Example: Destination: Datacenter_LAN Steering Strategy: Lowest Quality SLA	Example: Source: Branch_LAN Destination: Datacenter_LAN Destination Interface: Datacenter Security Inspection: Branch_Group

Additional details

- SD-WAN members:
 - Each overlay will be an SD-WAN member
 - If one member or WAN is preferred over another, assign a lower cost to the preferred interface, which can be used later by means of a *Lowest Cost SLA* SD-WAN rule.
- SD-WAN zone:
 - Zones should be grouped by device and location. For example, in the design for single datacenter with an active-passive gateway (see [Single datacenter \(active-passive gateway\) on page 41](#)), all overlays to the gateway can share a single Zone.
 - In active-passive designs, all overlays will terminate to a single, active device. As a result, no additional considerations are necessary.
- Performance SLA:
 - The health-check destination IP address should be strategically pointed to locations where the business application or workload is located.
 - Loopbacks on the gateway can also be created to serve as a health check server.
- SD-WAN rules
 - The appropriate route must be in the routing table in order for an SD-WAN rule to be active. If the route is not installed correctly or misconfigured, the SD-WAN rule will be considered *inactive* and skipped.
 - Branch to corporate traffic (datacenter or other branch locations) should contain all available overlays and their appropriate SLA.
 - Since corporate traffic is typically known, it is often preferred to match destination by the object or `route-tag` as opposed to application or internet service.
- Firewall policy
 - The firewall policy should reference the datacenter zone(s) with the appropriate rule and security profiles enabled.
 - Policies can only reference zones and not individual members. If you need different policies or inspection per WAN, consider creating a SD-WAN zone per overlay member.

Security considerations

As part of the *Zero Trust Security* model, we don't want to assume corporate traffic in other segments is trusted without applying appropriate security controls and inspection. Following is a list of security considerations for your design:

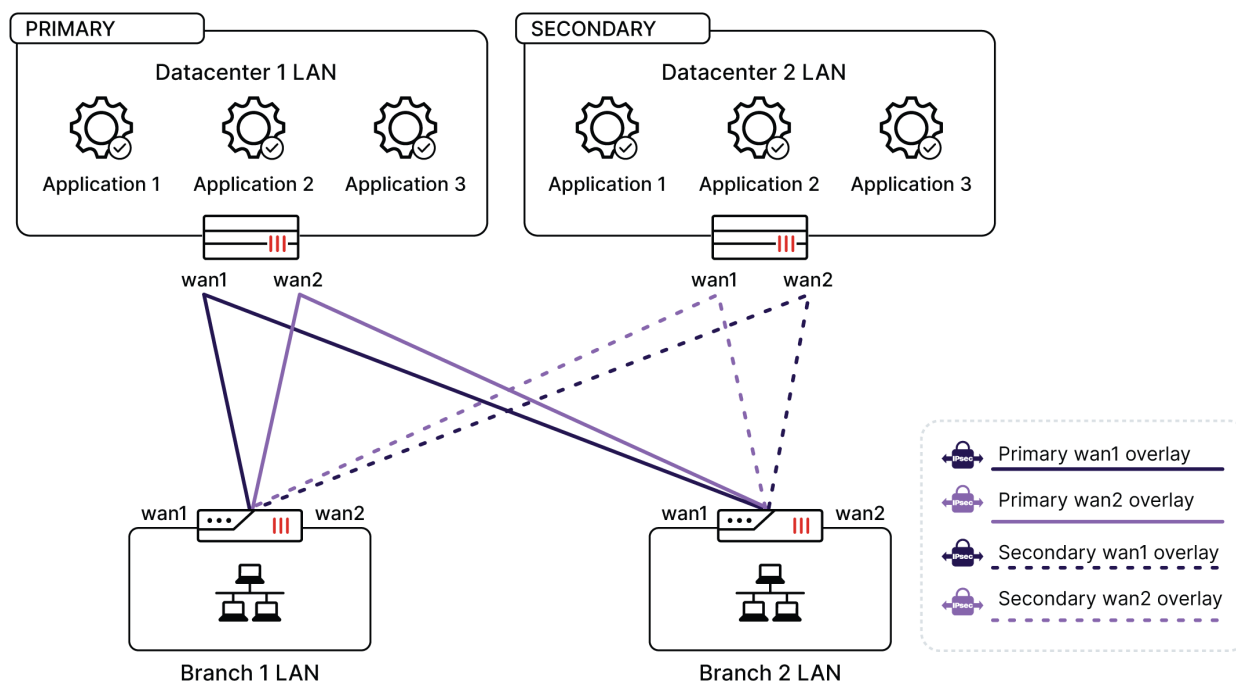
Risk	Mitigation	Considerations
Malware	Antimalware	Enabled on all external traffic from the datacenter, remote internet breakout designs, and file transfers between corporate resources
Malicious websites	Web filtering	Enabled on all HTTP/HTTPS traffic for remote internet breakout scenarios
Application visibility and reporting	Application control	Enabled on all network traffic
Server-side attacks	Intrusion prevention	Enable IPS signatures for server targets on appropriate targets

Risk	Mitigation	Considerations
Data loss	Data loss and prevention (DLP)	Enabled where applicable for sensitive data that should not transfer across unauthorized boundaries
Unauthorized access	Role-based access control (RBAC) and Zero Trust Network Access (ZTNA)	Lock down inbound policies as much as possible Utilize ZTNA

Multiple datacenters (primary/secondary gateways)

Deployments requiring geo-redundant datacenters (or private workload) connectivity will extend our base design to include a second SD-WAN gateway. For information about the base design, see [Single datacenter \(active-passive gateway\) on page 41](#).

In this design, branch SD-WAN devices now have two (or more) gateways at separate geo-redundant locations from which to steer traffic. A primary gateway is usually located in the preferred datacenter location while the secondary gateway is at the redundant location. Traffic will flow through the primary gateway under normal conditions and utilize the secondary gateway as a backup. While this section will refer to a traditional datacenter, the gateway can be located in a public or private cloud, customer premise location, or any other location where FortiGate physical or virtual appliance can be configured.



This section includes the following topics:

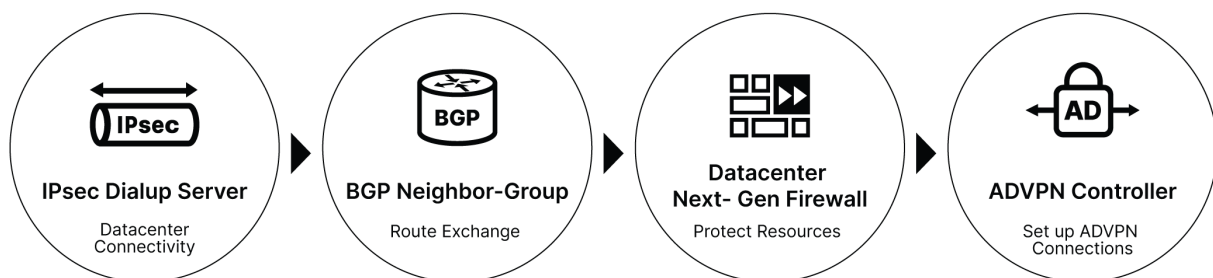
- [Gateway components on page 53](#)
- [SD-WAN considerations on page 63](#)
- [Security considerations on page 64](#)

Gateway components

In this design, each hub acts precisely as in the base design, and the hubs are independent of each other. The spokes connect to the dial-up IPsec endpoints of both hubs, over all available underlay transports. Effectively, each of the hubs defines its own set of point-to-multipoint overlays.

Each SD-WAN Gateway may provide one or multiple services:

- Act as the IPsec dialup server for branch locations
- Provide centralized routing information and orchestrate dynamic branch-to-branch communication (ADVPN)
- Protect the datacenter resources and private workloads by utilizing FortiGate Next-Generation Firewall services
- Provide remote internet breakout for branch locations



This section includes the following topics:

- [Redundancy on page 53](#)
- [Intra-datacenter failover on page 54](#)
- [Inter-datacenter failover on page 55](#)
- [IPsec overlays on page 56](#)
- [Route Exchange on page 58](#)
- [Traffic flow on page 60](#)

Redundancy

This design includes multiple SD-WAN Gateways located at geo-redundant datacenter locations that provides inter-datacenter and intra-datacenter redundancy. Intra-datacenter redundancy refers to the gateway redundancy that occurs inside a datacenter location. Inter-datacenter redundancy refers to the SD-WAN branch device steering between geo-redundant locations based on the predefined SLAs and steering rules.

In the following sections, we will discuss how to accomplish intra-site and inter-site redundancy to maximize business continuity for various use cases. For more information on the various high availability mechanism and modes, please refer to the [FortiOS Administration Guide](#).

Gateway Redundancy	Benefits	Considerations
Intra-datacenter redundancy with active-passive HA cluster	<p>Local datacenter protection from equipment failures and other datacenter impairments</p> <p>Logically seen as a single device per datacenter</p>	Offers redundancy at the local datacenter or HQ location from various impairments

Gateway Redundancy	Benefits	Considerations
Inter-datacenter redundancy with primary/secondary gateways at geo-redundant locations	Geo-redundant datacenter protection and steering Logically seen as different devices at different locations SD-WAN rules will determine optimal path	Offers redundancy from issues at the datacenter location Allows you to choose the best performing SD-WAN gateway across locations for improved application performance
Inter-datacenter redundancy with primary/primary gateways at geo-redundant locations	Like primary/secondary	

Intra-datacenter failover

SD-WAN gateways at each datacenter operate as independent HA clusters to offer intra-site redundancy from failures and issues at their location.

FortiGate HA offers several solutions for adding redundancy in the case where a failure occurs on the FortiGate, or is detected by the FortiGate through monitored links, routes, and other health checks. These solutions support fast failover to avoid lengthy network outages and disruptions to your traffic.

FortiGate HA options:

- [FortiGate Cluster Protocol \(FGCP\)*](#)
 - Active/passive
 - Active/active
- [FortiGate Session Life Support Protocol \(FGSP\)](#)
 - Session and configuration synchronization across standalone FortiGate or HA clusters

* In this document, we will focus on utilizing the FortiGate Cluster Protocol (FGCP) on our SD-WAN gateways to accomplish high availability.



There are more advanced use cases and scenarios where FGSP may be used to sync sessions between FortiGate clusters at different datacenter locations. This is beyond the scope of this document.

Utilizing FGCP for intra-datacenter HA

For most use cases, it is generally recommended to utilize active-passive HA for SD-WAN gateways at a datacenter or HQ location. If active-active is desired, it will not change our overall SD-WAN design outlined below. Both HA modes will be designed in the same matter as described in this section.

In active-passive HA mode, there are at least two devices in the cluster, with only one device acting as the primary device. To the rest of the network, including remote branch locations, the active-passive cluster appears to be a single device that shares a *floating* IP address between the active members. Remote branch locations terminate their overlays to the active device in the cluster.

The active-passive gateway model provides redundancy inside the datacenter, while operating as a single device to outside resources. Branch locations terminate their overlay connections to the active member, while being unaware the gateway is a cluster with multiple members.

Gateway Redundancy	Benefits	Considerations
Active-passive HA cluster	<ul style="list-style-type: none"> Intra-datacenter redundancy Logically seen as a single device on the network Simple setup 	<ul style="list-style-type: none"> Does not provide performance improvement on security inspection
Active-active HA cluster	<ul style="list-style-type: none"> Intra-datacenter redundancy Logically seen as a single device on the network Performance improvement on security inspection 	<ul style="list-style-type: none"> More complex troubleshooting requirements due to the nature of active-active load balancing

For more information on HA design and consideration, refer to the latest [FortiOS Admin Guide](#).


To minimize or eliminate traffic interruption during failover, it is recommended to consider the following:

- Enabling BGP graceful restart on the gateway and branch
- Enabling `route-ttl` on the HA settings to ensure the FortiGate cluster maintains the cached routes during failover
- Fine-tune BGP timers as necessary

For more information on these three components, see this [KB article](#).

Inter-datacenter failover

Gateways at each datacenter location operate as independent clusters. Traffic from the branch is steered between the primary and secondary datacenter locations using SD-WAN rules. The SD-WAN device at the branch location may detect issues at a datacenter location using its performance SLAs and steer traffic based on its preconfigured SD-WAN rules.



Example of a *Lowest Cost SLA* rule on the branch to steer between *Datacenter1 (HUB1)* and *Datacenter2 (HUB2)*:

MULTIPLE DATACENTERS (PRIMARY/SECONDARY GATEWAYS)

1. Overlays to *Datacenter 1* (HUB1-VPN1 and HUB1-VPN2) are given a lower cost (0).

+ Create New ▾ Edit Delete Where Used Column Settings				
<input type="checkbox"/>	ID	Interface Member	Status	Cost
<input type="checkbox"/>	virtual-w			
<input type="checkbox"/>	SASE			
<input type="checkbox"/>	Cloud			
<input type="checkbox"/>	9	Cloud-VPN1	✓ Enable	0
<input type="checkbox"/>	10	Cloud-VPN2	✓ Enable	0
<input type="checkbox"/>	HUB1			
<input type="checkbox"/>	1	HUB1-VPN1	✓ Enable	0
<input type="checkbox"/>	2	HUB1-VPN2	✓ Enable	0
<input type="checkbox"/>	HUB2			
<input type="checkbox"/>	4	HUB2-VPN1	✓ Enable	10
<input type="checkbox"/>	5	HUB2-VPN2	✓ Enable	10
<input type="checkbox"/>	WAN1			
<input type="checkbox"/>	7	port1	✓ Enable	0
<input type="checkbox"/>	WAN2			
<input type="checkbox"/>	8	port2	✓ Enable	0

! DC1 cost=0

! DC2 cost=10

2. The rule will prefer *Datacenter*, as long as its SLAs are met:

Edit SD-WAN Rule

Name: Branch_to_DC

IP Version: IPv4

Source

Source Address: Branch-NET (10.1.0.0/255.255.0.0) 1 Entry Selected

Users: Click here to select

User Groups: Click here to select

Destination

Address: Datacenter (192.168.1.0/255.255.0.0) Branch-NET (10.1.0.0/255.255.0.0) 2 Entries Selected

Route Tag: 0

Protocol: TCP UDP ANY Specify 0

Type of Service: 0x00 Bit Mask: 0x00

Outgoing Interfaces

Strategy: Manual Best Quality Lowest Cost (SLA) Maximize Bandwidth (SLA)

Interface Preference: HUB1-VPN1 HUB1-VPN2 HUB2-VPN1 HUB2-VPN2

! Datacenter LAN

! Datacenter Overlays

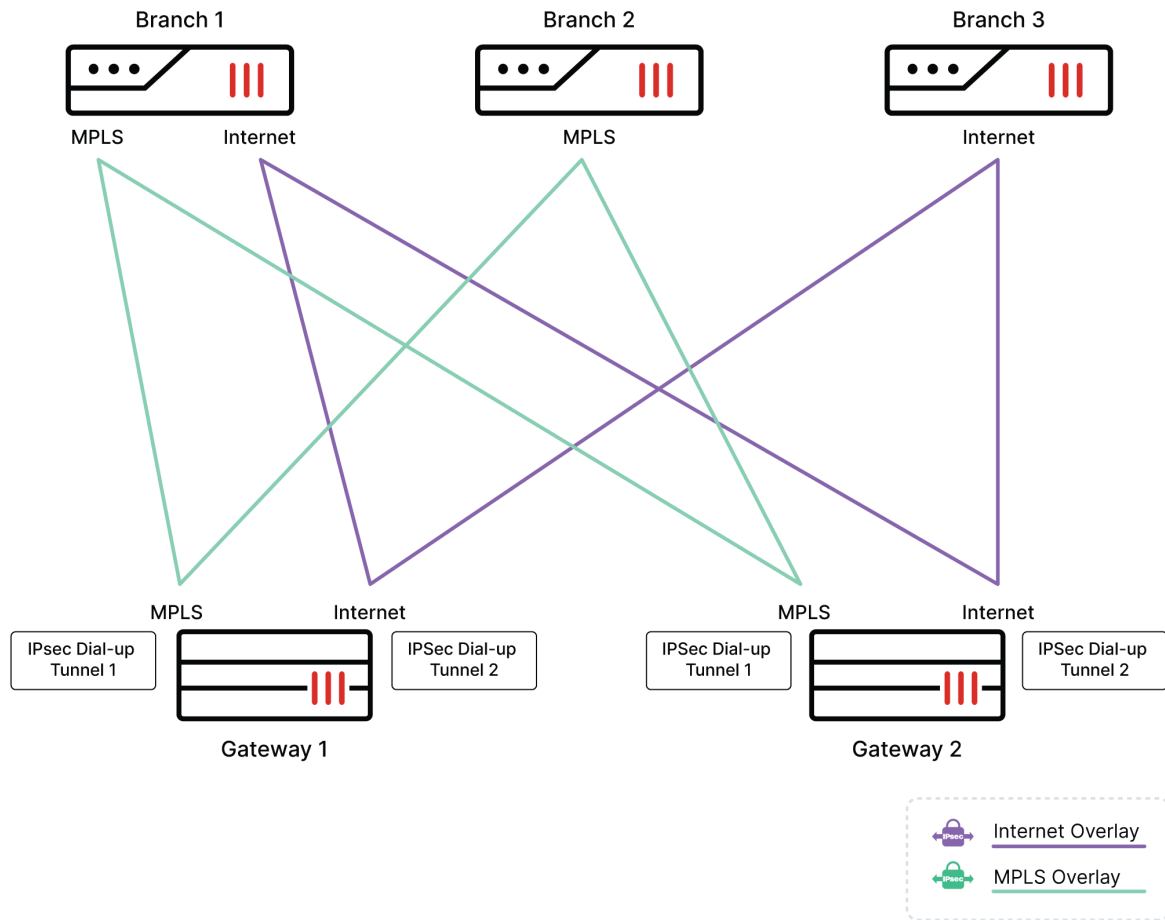
! Strategy

OK Cancel

When SLAs are not met or failure is detected, traffic will failover to *Datacenter 2*.

IPsec overlays

Each SD-WAN gateway acts as a dial-up IPsec server for the spokes, having a separate dial-up IPsec endpoint terminate on each underlay interface. Branches will typically build overlays over all available WAN ports to have multiple paths available to the gateway. However, it can also happen that some of the branches do not have a similar WAN transport. Hence, they will be able to connect only to a subset of the overlays.

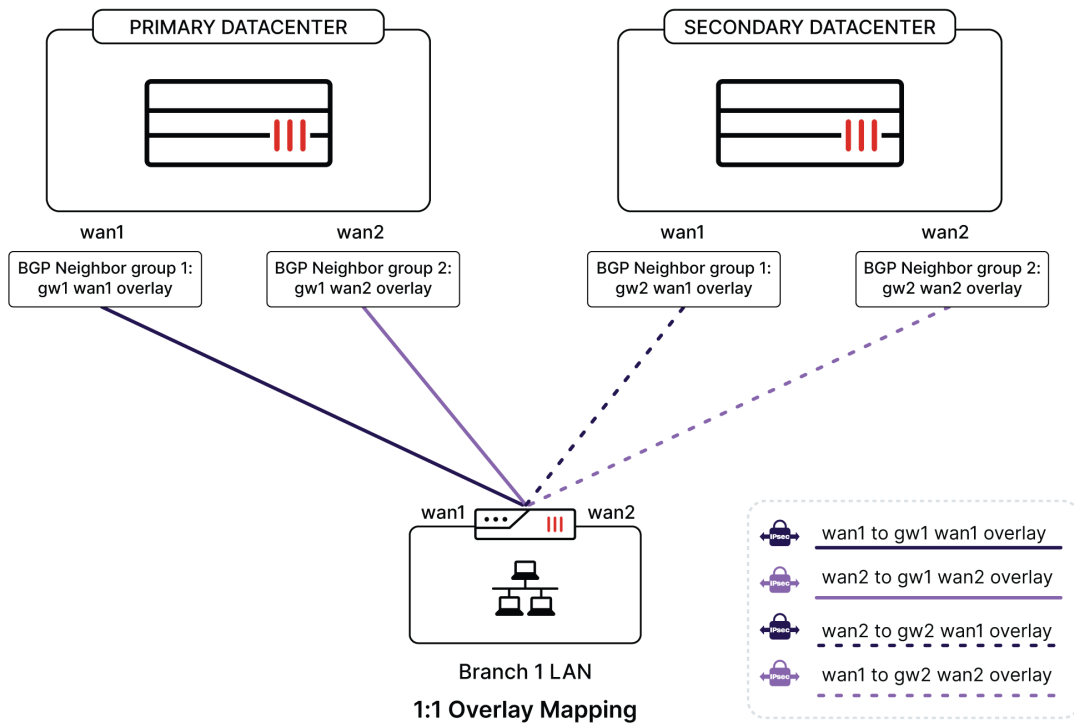


When considering the IPsec overlay design between the branch locations and the gateway, it is important to determine how redundancy should occur between all available links.

Consider the following options:

- **One-to-one overlay mapping per underlay:** in this design, each branch underlay terminates a new IPsec tunnel to one—and only one—gateway underlay. This is the most common overlay design, and simplifies our configuration, but also provides less redundancy than the subsequent full mesh.

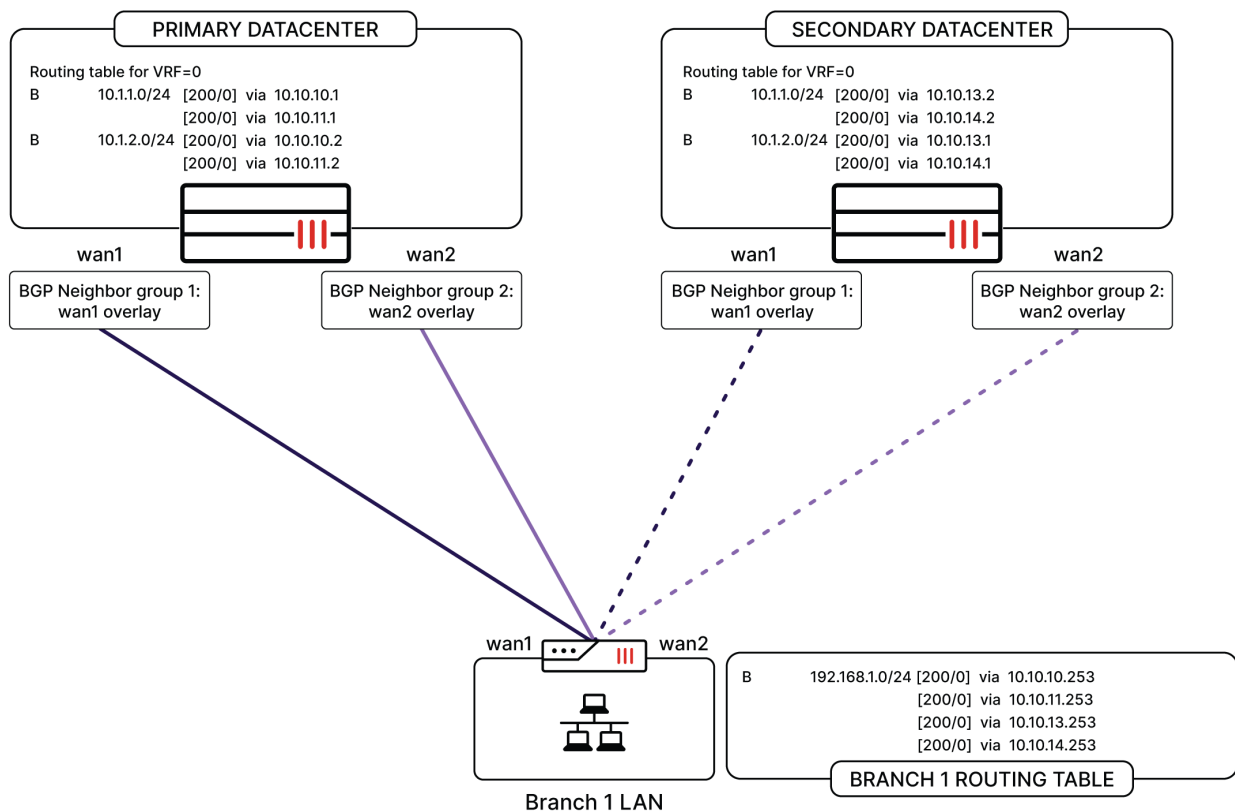
MULTIPLE DATACENTERS (PRIMARY/SECONDARY GATEWAYS)



- **Full mesh overlay mapping** is generally not recommended for multi-datacenter deployments, unless there is a specific use case by which this may be required.

Route Exchange

The spokes establish separate IBGP sessions to each gateway over each overlay. Using the previous sections' example for one-to-one overlay mapping, that means our branch device should have four separate BGP session via four available overlays to Datacenter 1 and Datacenter 2:



The BGP Neighbor Group feature is used on the gateway for this peering. Each spoke then advertises its local site prefix(es) over each of the IBGP sessions. The gateway acts as a BGP Route Reflector (RR), readvertising the prefixes to all other spokes when ADVPN is used. Additionally, each gateway advertises its prefixes (such as the datacenter LANs) to every branch location. At the end of this process, all the sites exchange their routes over all available overlays.

Additional Routing Notes

- Routing inside in the datacenter is typically handled by each gateway via BGP or OSPF. Datacenter networks are readvertised to the Branches via BGP.
- IBGP sessions are terminated on the IPsec overlays, and hence, they are using the tunnel IPs as BGP next-hops (NH). This requires IP addresses to be configured on the tunnel interfaces. Each gateway will have its own overlay network that can automatically allocate tunnel IPs to the spokes using the IKE Mode Config feature to simplify provisioning and administrative overhead.
- Since the spokes establish separate IBGP sessions with the hub over each overlay, there are multiple BGP routes for each prefix. To keep all the routes available, the following two BGP features must be enabled on all participating devices (hub and spokes):
- BGP Multipath ensures that all the available routes are installed into the routing tables
- BGP ADD-PATH ensures that the hub between the spokes reflects all available routes

ADVPN

For the correct operation of ADVPN, it is required to preserve all sites' prefixes unchanged, including their original BGP next-hop values. Hence, it is impossible to replace the specific routes with summaries (unlike in a static hub-and-spoke topology). Hence, the BGP RR function is mandatory: the gateway must reflect the original routes between the spokes without altering them.

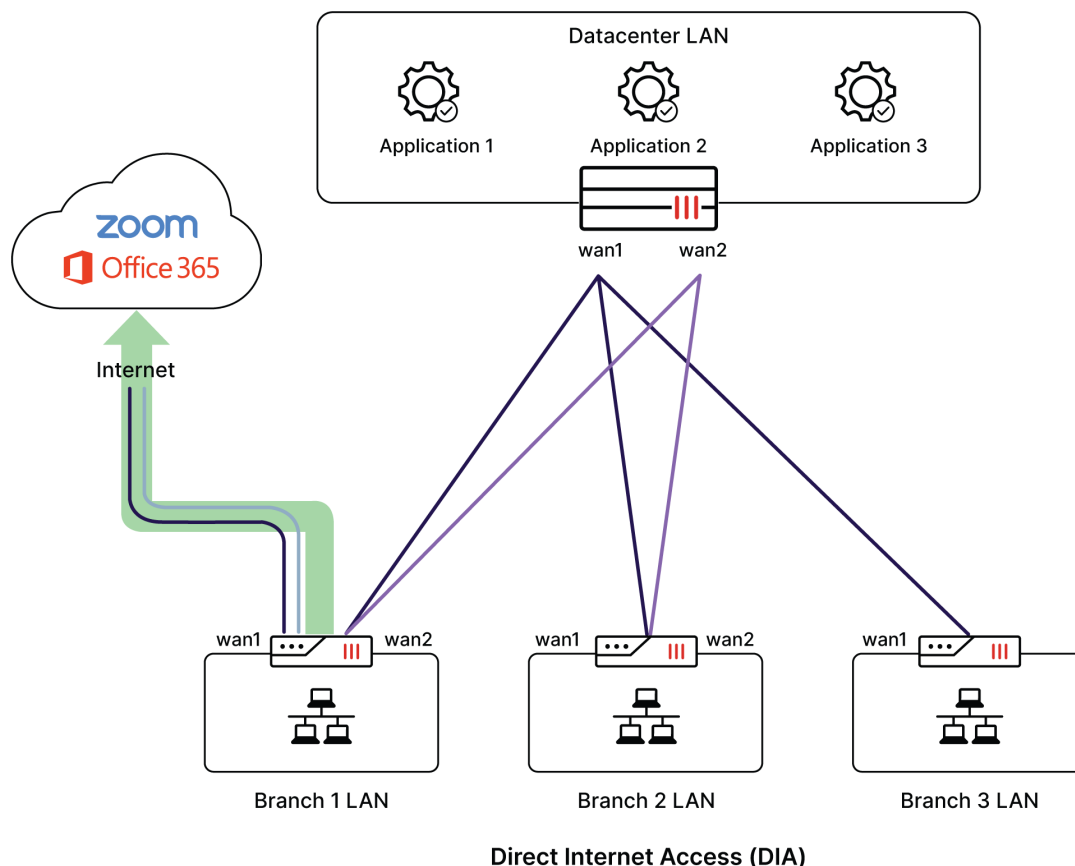
FortiOS 6.4 and earlier:

We have already mentioned the critical property of overlay stickiness that we must guarantee for proper ADVPN shortcut creation. For example, if Branch1 sends traffic to Branch2 using an internet overlay via the hub, the hub must select the same internet overlay for the second half of the path. Failing to preserve the overlay might result in an attempt to create an ADVPN shortcut between two physically disconnected transports (such as the internet and MPLS), and this attempt would, of course, fail. The overlay stickiness is achieved using Policy Routes (PBR) on each of the gateways.

Traffic flow

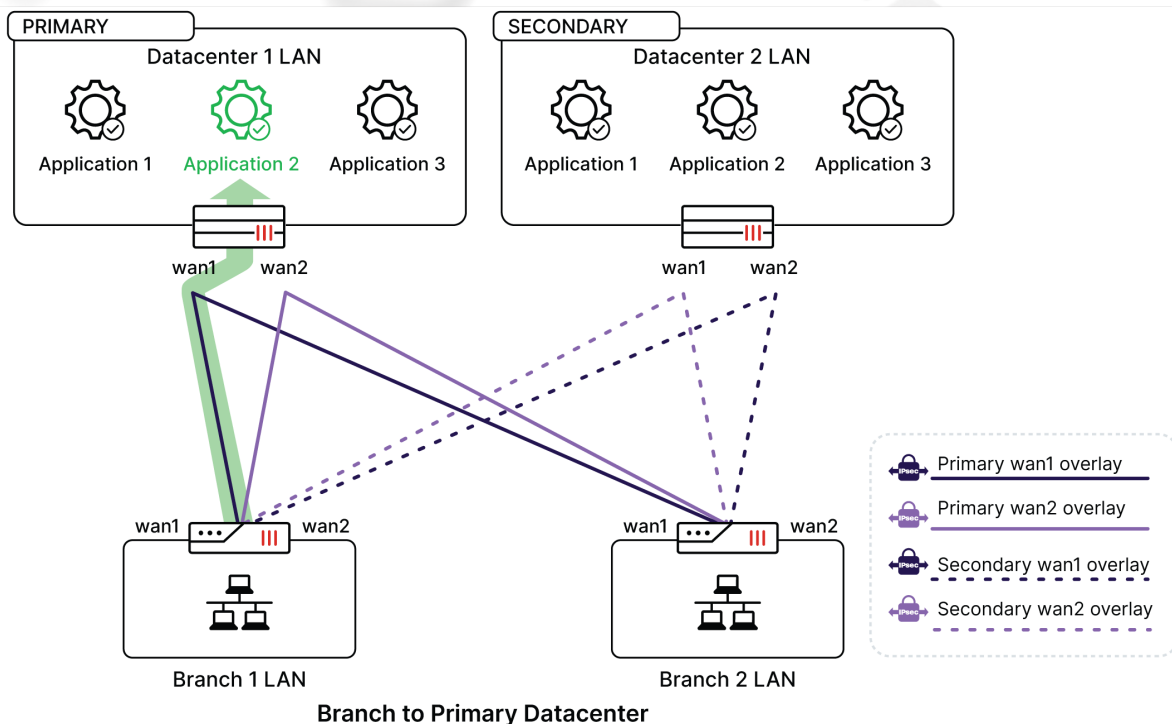
Once all the routes have been distributed across all the sites, the application traffic flow can be controlled by SD-WAN rules according to the design principles described in the previous chapter. SD-WAN rules may dictate how traffic is steered based on the business requirement and desired redundancy.

- **Direct internet access (DIA):** used when local internet breakout at a location is required. In this scenario, the business application(s), such as a SaaS application or website, is located on the internet, and the SD-WAN appliance is needed to decide the best path between multiple WAN links. Traffic is routed directly to the internet by using the preferred method in the SD-WAN rule.

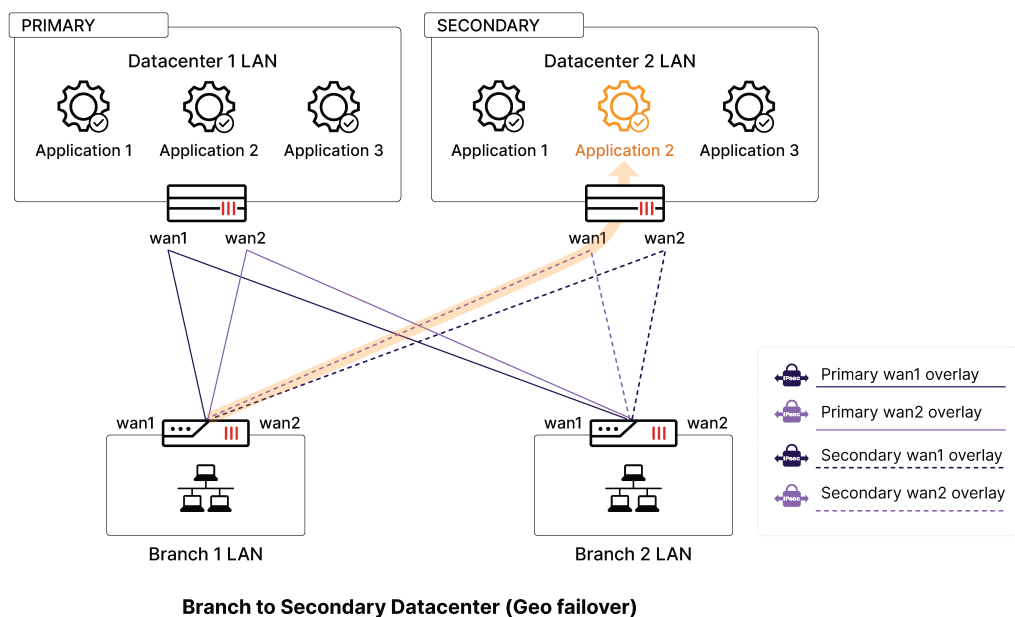


- **Branch to primary datacenter:** used when branch users require connectivity to an application or workload located behind the gateway at the primary datacenter. The secondary gateway located in the secondary datacenter will only be used as a backup. The branch SD-WAN device should monitor all available overlay links, and choose the best path according the business requirements.

MULTIPLE DATACENTERS (PRIMARY/SECONDARY GATEWAYS)

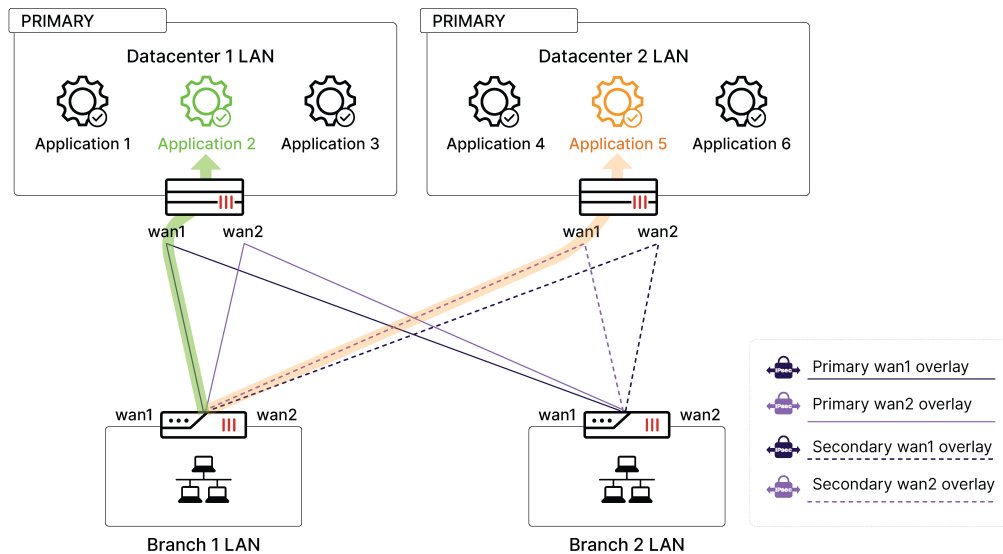


- **Branch to secondary datacenter (geo-redundant datacenter failover):** A catastrophic failover at the primary datacenter location causes traffic to route through the gateway located in the secondary datacenter. In this scenario, the desired application either lives in both datacenter locations, or the gateway has an alternative path to the primary datacenter.



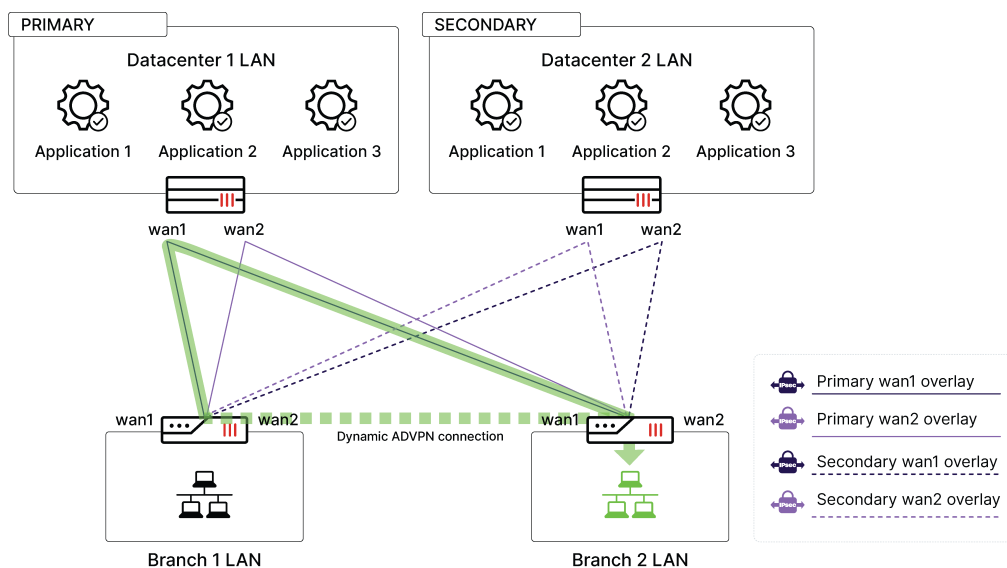
- **Branch to Datacenter1 or 2 LAN access:** It is common for redundant datacenter locations to host applications or services locally. In this scenario, Datacenter 1 has its own network space that is advertised to all branch locations with applications and services being offered. Datacenter2 also has its own network space with application and services that are independent from Datacenter1. Branch users require access to both at any given time and must use the optimal path to access their resources.

MULTIPLE DATACENTERS (PRIMARY/SECONDARY GATEWAYS)



Branch to Datacenter 1 and Datacenter 2

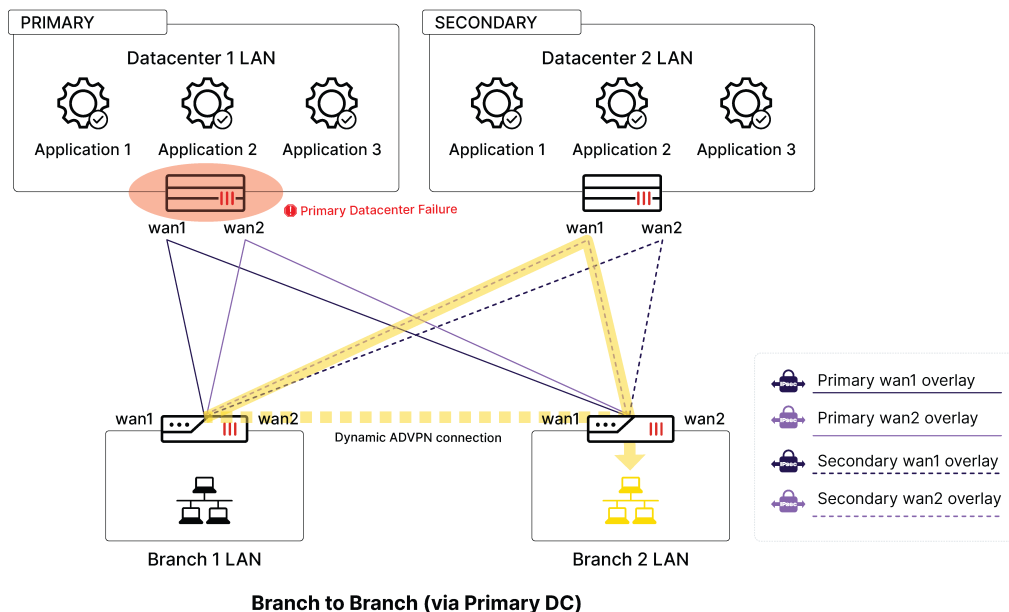
- **Branch to branch:** when ADVPN is used for dynamic branch-to-branch communication, both gateways may provide the routing and IPsec tunnel information necessary for direct communication. The gateway selected by the SD-WAN rule will dictate who becomes the *ADVPN Sender*. In this example, the gateway at our primary datacenter is the ADVPN sender under normal conditions.



Branch to Branch (via Primary DC)

In the event of a major failure at the primary datacenter, the branch SD-WAN will send traffic to the secondary datacenter, which will then become the ADVPN sender.

MULTIPLE DATACENTERS (PRIMARY/SECONDARY GATEWAYS)



SD-WAN considerations

SD-WAN Member	SD-WAN Zone	Performance SLA	SD-WAN Rule	Firewall Policy
All overlay interfaces to both Datacenter locations	Overlays should be grouped by device and location	Health check server: Business critical applications and resources	Dependent on business intention and availability requirements.	References the SD-WAN zone(s) and appropriate security inspection
Example: Cost: 10 (preferred) dc1_overlay1_wan1, dc1_overlay2_wan2 Cost: 20 dc2_overlay1_wan1, dc2_overlay2_wan2	Example: Datacenter1 (dc1_overlay1_wan1, dc1_overlay2_wan2) Datacenter2 (dc2_overlay1_wan1, dc2_overlay2_wan2)	Example: Health-check: App1_DC1 Members: dc1_overlay1_wan1, dc1_overlay2_wan2, App2_DC2 Members: dc2_overlay1_wan1, dc2_overlay2_wan2	Example: Steering Strategy: Lowest Quality SLA Members: dc1_overlay1_wan1, dc1_overlay2_wan2 dc2_overlay1_wan1 dc2_overlay2_wan2	Example: Source: Branch_LAN Destination: Datacenter_LAN Destination Interface: Datacenter1, Datacenter2

Additional details

- SD-WAN member:
 - All overlays to each gateway at Datacenter1 and Datacenter2
 - Assign the preferred datacenter overlays at a lower cost than the secondary datacenters. For example:
 Cost: 10 (preferred)
 dc1_overlay1_wan1, dc1_overlay2_wan2
 Cost: 20

dc2_overlay1_wan1, dc2_overlay2_wan2

- SD-WAN zone:
 - Zones could be grouped by datacenter location.
For example, Datacenter1 overlays should be grouped together, while Datacenter2 overlays should be grouped separately.
If you need to use packet duplication, this ensures that duplicated packets will be kept to the same gateway.
- Performance SLA:
 - Loopbacks should be created on each gateway and used as health-check servers for the performance SLA on the branch.
 - Additional performance SLAs could be created for business critical applications or services in the datacenter. This will allow you to have individual SLA requirements for each application or service.
- SD-WAN rules:
 - The appropriate route must be in the routing table in order for an SD-WAN rule to be active. If the route is not installed correctly or misconfigured, the SD-WAN rule will be considered *inactive* and skipped.
 - Branch to corporate traffic (datacenter or other branch locations) will now have multiple overlay paths to reach its destination.
 - When deciding which rule to use for steering between geo-redundant locations, stability of traffic flow is generally preferred over performance.
The *Best Quality Steering Strategy* will select the best performing path, which may fluctuate throughout the day.
The *Lowest Cost SLA Steering Strategy* will prefer the interfaces with the lowest cost (that is, the primary datacenter), as long as it meets the minimum SLA thresholds.
This is generally recommended for steering between datacenters.
- Firewall policy
 - The firewall policy should reference the datacenter zone(s) with the appropriate rule and security profiles enabled.
 - Policies can only reference zones and not individual members. If you need different policies or inspection per WAN, consider creating a SD-WAN zone per overlay member.

Security considerations

As part of the *Zero Trust Security* model, we don't want to assume corporate traffic in other segments is trusted without applying appropriate security controls and inspection. Following is a list of security considerations for your design:

Risk	Mitigation	Considerations
Malware	Antimalware	Enabled on all external traffic from the datacenter, remote internet breakout designs, and file transfers between corporate resources
Malicious websites	Web filtering	Enabled on all HTTP/HTTPS traffic for remote internet breakout scenarios

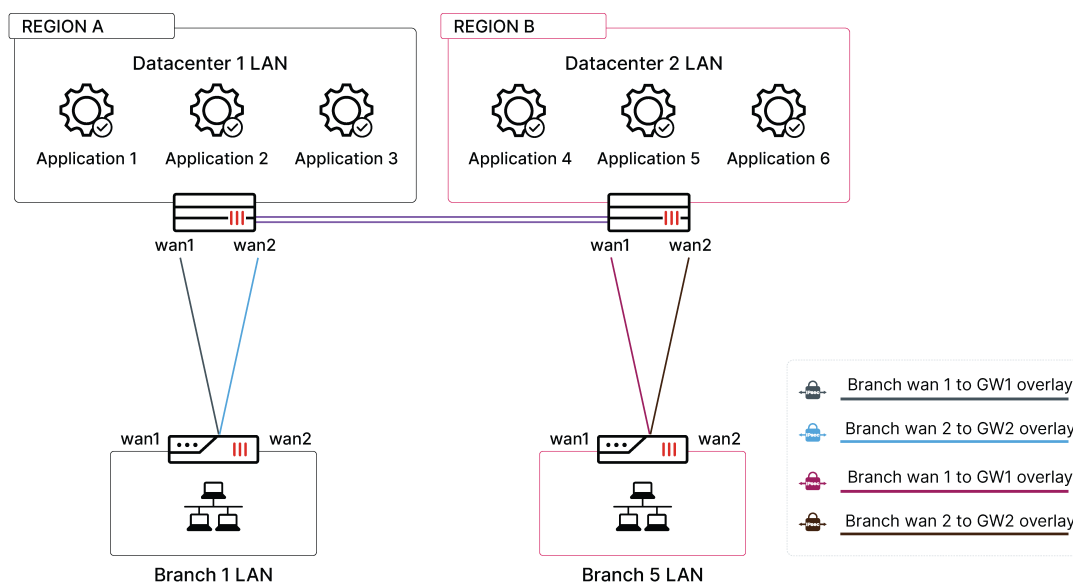
Risk	Mitigation	Considerations
Application visibility and reporting	Application control	Enabled on all network traffic
Server-side attacks	Intrusion prevention	Enable IPS signatures for server targets on appropriate targets
Data loss	Data loss and prevention (DLP)	Enabled where applicable for sensitive data that should not transfer across unauthorized boundaries
Unauthorized access	Role-based access control (RBAC) and Zero Trust Network Access (ZTNA)	Lock down inbound policies as much as possible Utilize ZTNA

Multi-region datacenters

As your solution expands geographically, and the number of sites grows, it becomes reasonable to define multiple regions. Each region would be comprised of a hub-and-spoke topology as described in one of the previous examples:

- [Single datacenter \(active-passive gateway\) on page 41](#)
- [Multiple datacenters \(primary/secondary gateways\) on page 52](#)

Branch devices connect to the gateway devices in their region, and gateways are interconnected between regions in a full mesh design.



This section contains the following topics:

- [Inter-region connectivity on page 66](#)
- [SD-WAN Considerations on page 68](#)

Inter-region connectivity

Gateways are defined for each geographical area, and all other sites in the area will only connect to these regional gateways. This includes both IPsec overlays and BGP sessions. As already discussed, this would be enough to provide connectivity within each region. In addition, all the regional hubs are interconnected between them, forming a full-mesh topology with BGP sessions exchanging the routes between all the regions.

Two recommended methods exist that define the routing configuration between the regions:

- eBGP
- iBGP

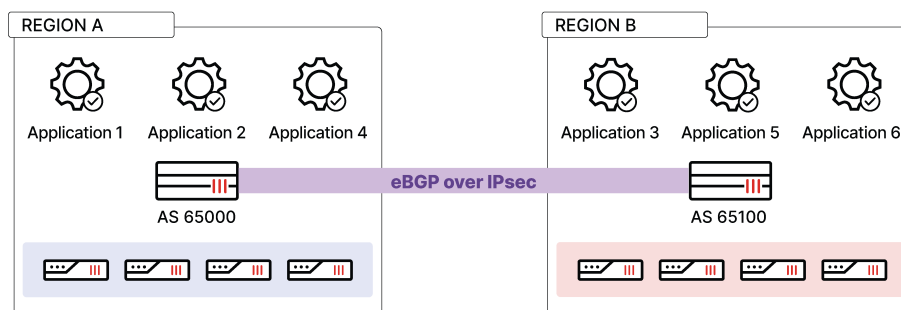
If cross-regional ADVPN is required, iBGP between regions will be required. If there is no requirement for ADVPN between regions, eBGP is preferred.

This section contains the following topics:

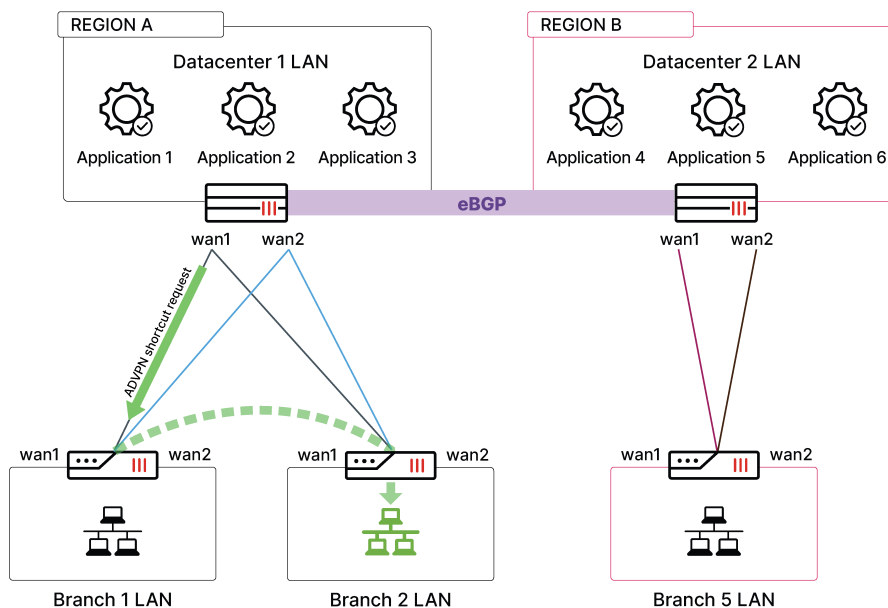
- [Using EBGP between regions with intra-region ADVPN on page 66](#)
- [Using IBGP between regions with inter-region ADVPN on page 67](#)

Using EBGP between regions with intra-region ADVPN

A straightforward approach is to use EBGP between the regional gateways. With EBGP, each gateway advertises a summary route of all regional prefixes to all remote regions. Those will, in turn, advertise default routes to their branches. A branch willing to communicate to a remote region will always send traffic to its local, regional gateway, which will use the correct summary route to forward the traffic to the remote regional gateway.




ADVPN will be used only for branch-to-branch traffic within each region, while the traffic across the regions will always flow via the regional hubs.

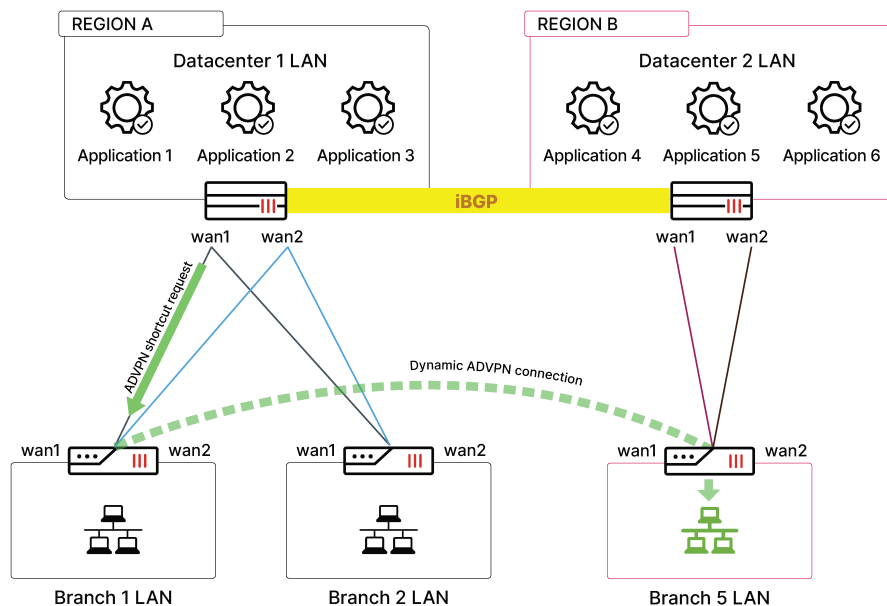


Using IBGP between regions with inter-region ADVPN

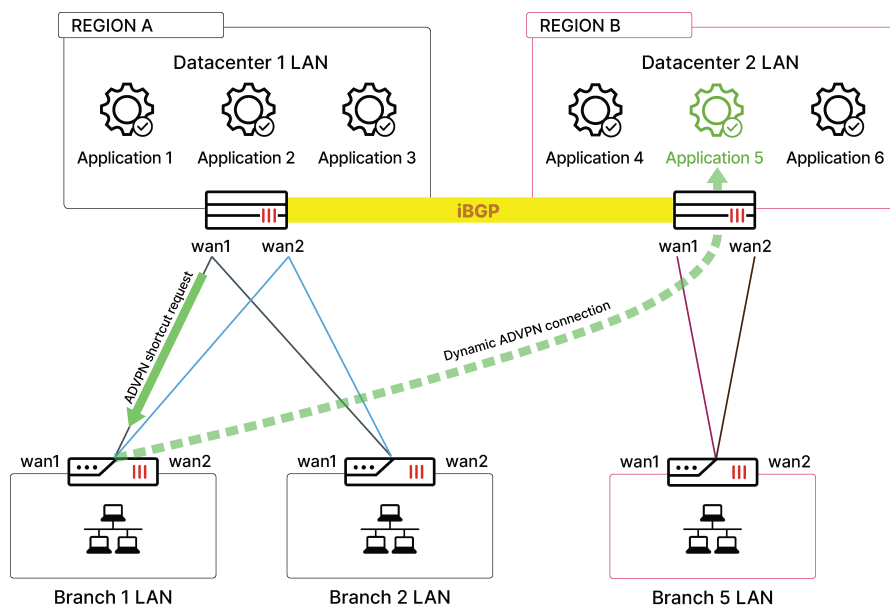
A more complex option is to implement cross-regional ADVPN, which requires preserving specific prefixes (including their original BGP next-hop values) between spokes belonging to different regions. Hence, IBGP must also be used between the regional gateways, just like it is used between gateway and branches. Each regional gateway will reflect prefixes to its branches and remote regional gateway. As a result, all the sites throughout the entire solution will learn each other's prefixes. This allows the use of ADVPN across the regions, dynamically building direct IPsec tunnels between any two sites willing to communicate.



Cross-regional branch-to-branch shortcuts will be built when two spokes belonging to different regions start communicating. The traffic will flow directly between them, thus bypassing both regional gateways on the way.



Cross-regional spoke-to-hub shortcuts will be built when a spoke is willing to reach a network behind a remote regional hub. The traffic will then bypass its local, regional hub.



SD-WAN Considerations

The SD-WAN configuration of the spokes in a multi-regional solution remains identical to the one described in the single-region examples. Note that the spokes are only connected to their local, regional hub overlays, and only those overlays are configured as SD-WAN members. Therefore, only those overlays will be used in SD-WAN rules for all the corporate traffic (including cross-regional ones).

This is true for both described methods: whether cross-regional ADVPN is used or not, the SD-WAN configuration on a spoke remains the same.

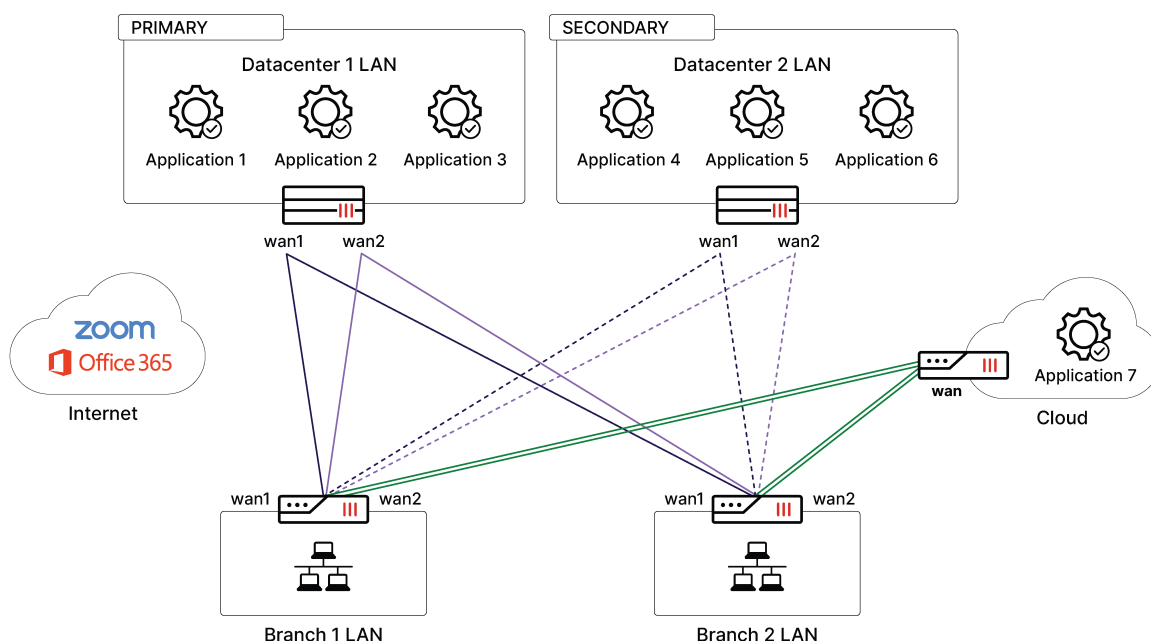
SD-WAN Member	SD-WAN Zone	Performance SLA	SD-WAN Rule	Firewall Policy
Based on per-region topology	Based on per-region topology	Health check server: Business critical applications or resource in the other region Health-check: RegB_DC1_App5 Members: dc1_overlay1_wan1, dc1_overlay2_wan2, dc2_overlay1_wan1, dc2_overlay2_wan2	Based on per-region topology. If inter-region traffic should be steered differently than intra-region traffic, more specific rules can be made for greater granularity	Granular firewall policies can be made for intra-region and inter-region control and inspection

Supplemental designs

This section will cover common SD-WAN scenarios and designs that build off of our previous designs and architectures covered in the previous sections. These designs will cover common SD-WAN use cases that are typically used in conjunction with any of the aforementioned architectures:

- [Single datacenter \(active-passive gateway\) on page 41](#)
- [Multiple datacenters \(primary/secondary gateways\) on page 52](#)
- [Multi-region datacenters on page 65](#)

However, they could also be utilized independently of any traditional SD-WAN topology.



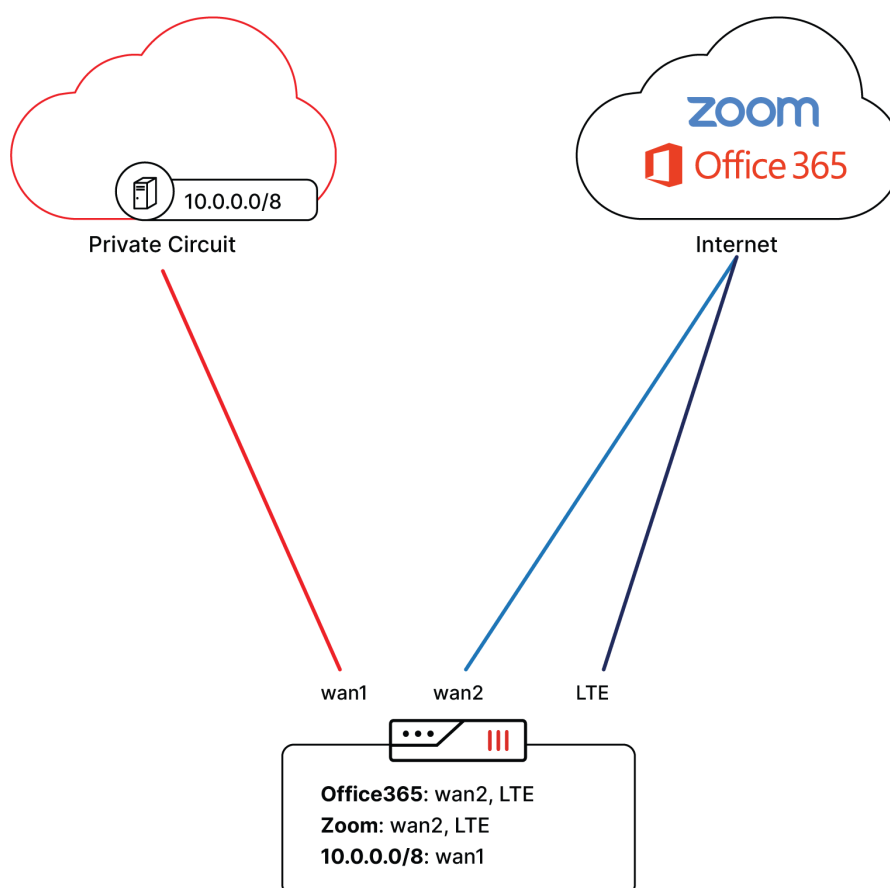
Use Case	Description
Direct internet access (DIA)	Secure, local internet breakout of SaaS applications and internet traffic without the need to offload to a remote location.
Cloud on-ramp	Connectivity and intelligent steering of network traffic between one of more cloud locations. This section will cover cloud on-ramp for static and dynamic environments.

This section contains the following topics:

- [Direct internet access on page 70](#)
- [Cloud on-ramp on page 72](#)

Direct internet access

The direct internet access (DIA) model is typically used when local internet breakout at a location is required. In this scenario, the business application(s), such as a SaaS application or website, is located on the internet, and the SD-WAN appliance is needed to decide the best path between multiple WAN links.



This section contains the following topics:

- SD-WAN considerations on page 71
- Security considerations on page 72

SD-WAN considerations

SD-WAN Member	SD-WAN Zone	Performance SLA	SD-WAN Rule	Firewall Policy
Configured for each WAN port	Zone exclusively for WAN ports	Health check server is a public server, such as Google's 8.8.8.8	Destination options are: all, application or internet service	References the SD-WAN zone(s) and appropriate security inspection

Additional details

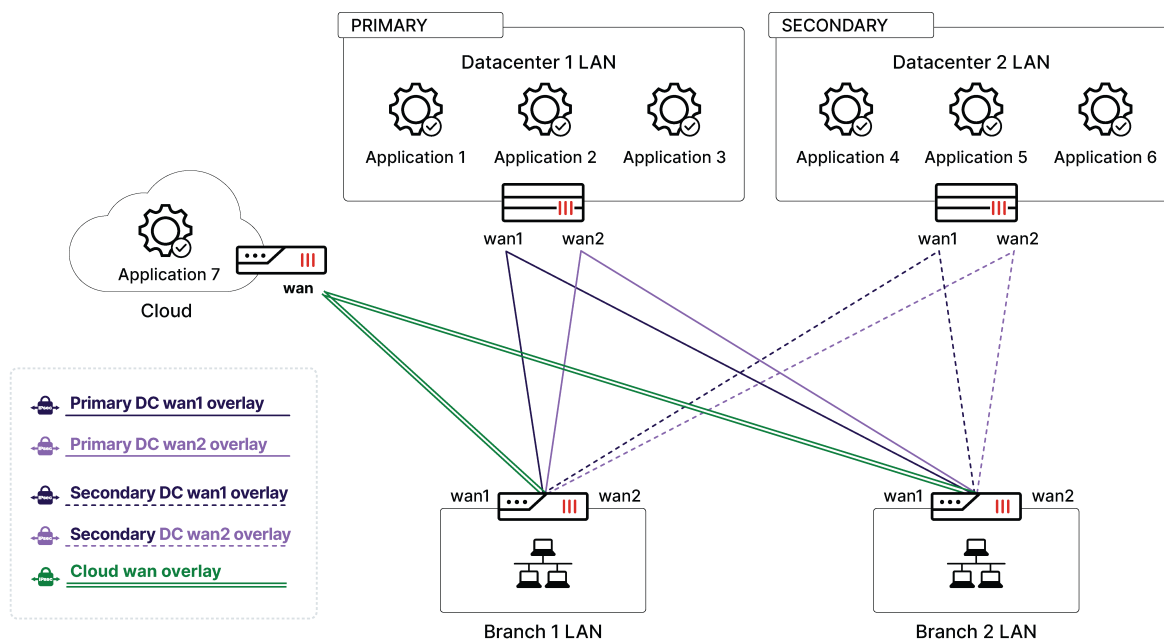
- SD-WAN member:
 - Each WAN interface should be added as an SD-WAN member.
 - If one member is preferred over another, you may assign a lower cost to the preferred interface, and reference it later in the SD-WAN steering strategy.
- SD-WAN zone:
 - Members can be added to their own individual zone or to a zone that contains multiple members.
 - The advantage of assigning each member to its own zone is that you have more granularity in your firewall policies later. For example, private links, such as MPLS, may require different security inspection than public internet links.
- Performance SLA:
 - For internet rules, a publicly available health-check server should be used, such as www.fortinet.com.
 - This rule should only utilize SD-WAN members that have public internet access.
 - SLA targets may need to be adjusted over time as you learn what is normal in your environment.
- SD-WAN rules:
 - Internet rules should contain the WAN members and their appropriate SLA.
 - More specific rules for internet breakout, such as application(s) or internet service, should be placed above more generic ones.
 - Destination options:
 - **Application:** Utilizes application control to identify the application on the network and steer appropriately. Good for granularity.
 - **Internet Service:** Utilizes the FortiGuard list of destination and port number mapping. Good for performance.
 - **Address:** Typically used for more generic rules, such as all or non-RFC statements.
- Firewall policy
 - The firewall policy should reference the datacenter zone(s) with the appropriate rule and security profiles enabled.
 - Policies can only reference zones and not individual members. If you need different policies or inspection per WAN, consider creating a SD-WAN zone per overlay member.

Security considerations

Risk	Mitigation	Considerations
Malware	Antimalware	Enabled on all user traffic
Malicious websites	Web filtering	Enabled on all HTTP/HTTPS traffic
Proxy avoidance, botnet and security circumventions	Application control	Enabled on all network traffic
Client-side web attacks	Intrusion prevention	Enable IPS signatures for client target on all web traffic
Unauthorized access	Role-based access control (RBAC) and Zero Trust Network Access (ZTNA)	Lock down inbound policies as much as possible Utilize ZTNA or VPN services for remote access

Cloud on-ramp

Cloud on-ramp provides optimized access to SaaS or IaaS workloads running in one or multiple cloud providers. Rather than accessing the cloud services through the public internet, a secure overlay can be established to the closest cloud point of presence (POP). FortiGate SD-WAN overlays use IPsec-compliant overlays that integrate with most cloud built-in gateways. However, deploying a FortiGate VM in the cloud as the gateway has many advantages, and is generally recommended for branch-to-cloud connectivity.



Branch communication to the cloud resources typically happens directly through secure overlays, similar to a datacenter gateway in previous designs. Branches will connect to the gateway through secure overlays and steer by using the best path according to the SD-WAN overlays. However, the design required for the cloud will vary based on the type of cloud environment.

This section contains the following topics:

- [Cloud design considerations on page 73](#)

Cloud design considerations

FortiGate VMs can be deployed as SD-WAN gateways for all major public and private cloud providers, including AWS, Azure, GCP, Oracle and AliCloud. While every cloud environment may differ, there are two fundamental ways a FortiGate VM is typically used:

- **Static environments:** the FortiGate VM will be assigned a traditional, static IP that does not change, and is expected to provide the head-end connectivity to the cloud workload.
- **Dynamic environments:** services or deployments where the FortiGate VM is expected to change, and/or the FortiGate VM is never the same.

The type of cloud design used will depend on the environment in which the FortiGate VM is deployed.

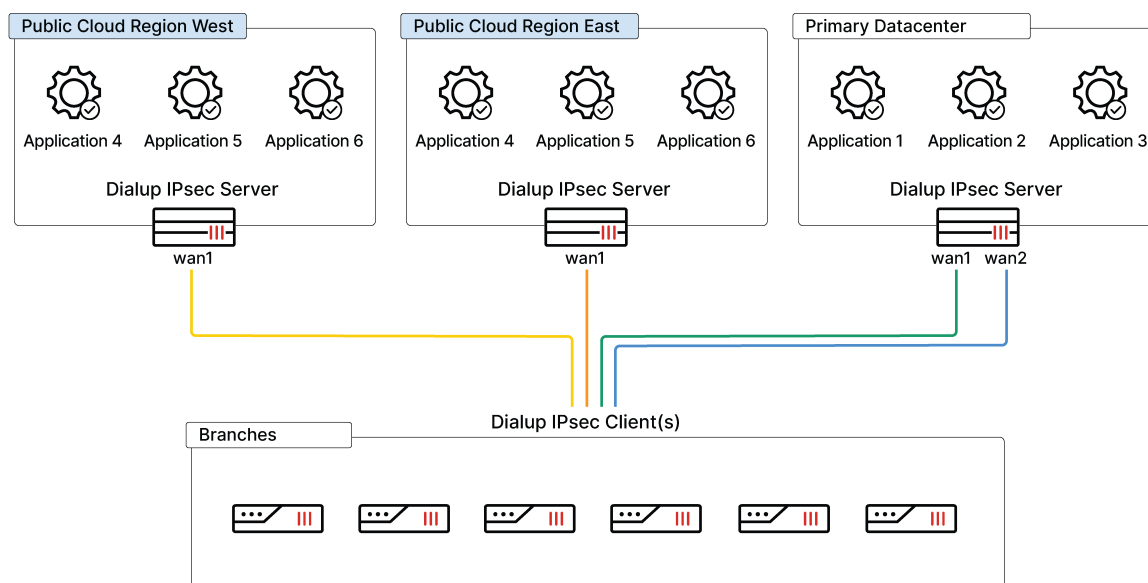
This section contains the following topics:

- [Cloud on-ramp for static environments on page 73](#)
- [Cloud on-ramp for dynamic environments on page 74](#)

Cloud on-ramp for static environments

A FortiGate VM may be deployed as a traditional virtual machine instance, where the IP address is static and not expected to change. In these environments, the FortiGate VM acts more like a traditional SD-WAN gateway at a datacenter that provides head-end connectivity to the branches and protection to the resources behind it (in other words, the cloud environment).

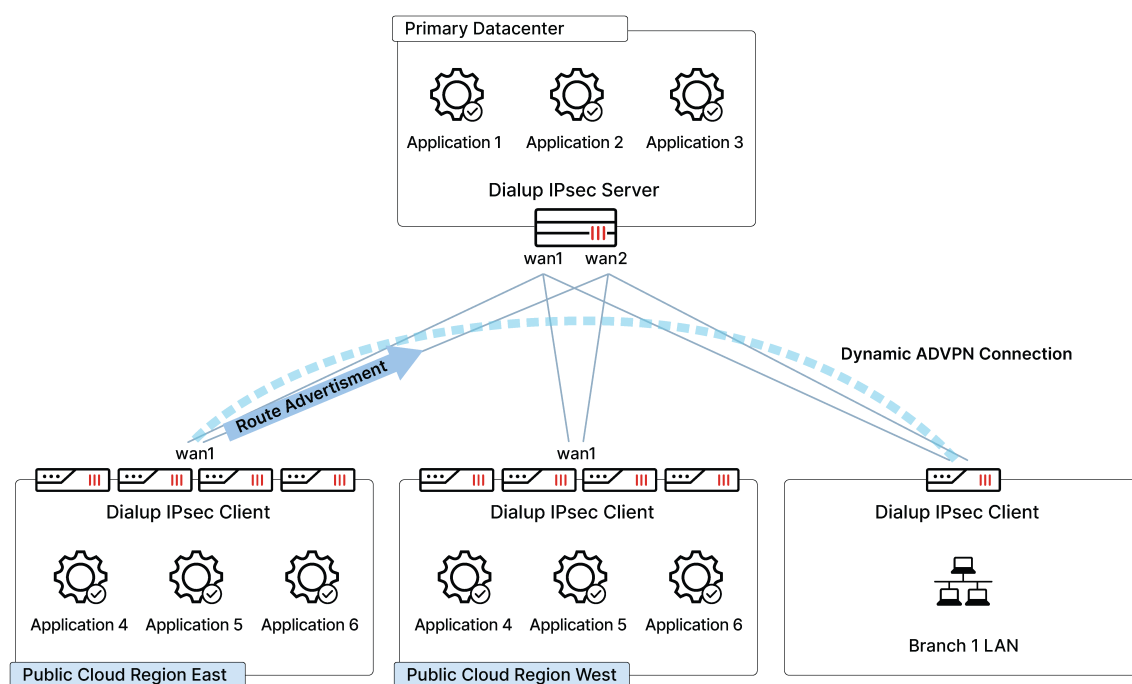
Since the IP address is not expected to change, the gateway may be configured as a traditional IPsec dialup server to which branch locations may connect for cloud services. Branch locations will have a new P1 IPsec definition for each new cloud gateway that is required. From the SD-WAN perspective, the traffic flow is similar to a corporate branch-to-gateway, since it will leave the site edge through one of the available overlays.



Cloud on-ramp for dynamic environments

Dynamic environments refer to deployments and services where the IP address and location of the FortiGate VM changes. Auto-scaling deployments where a new FortiGate VM is spun up and deleted on-demand is an example of such an environment. Since there is no static IP address to use as the destination for branches, a new design is required.

Instead of configuring the FortiGate VM to act as a dialup server, these environments will instead act as a dial up client to another gateway (such as the datacenter SD-WAN gateway). ADVPN will then be utilized for branch-to-cloud dynamic connectivity. This allows the cloud gateway(s) to spin up from any location with various different IP addresses and still provide access to branch sites through the dynamic nature of the ADVPN protocol.

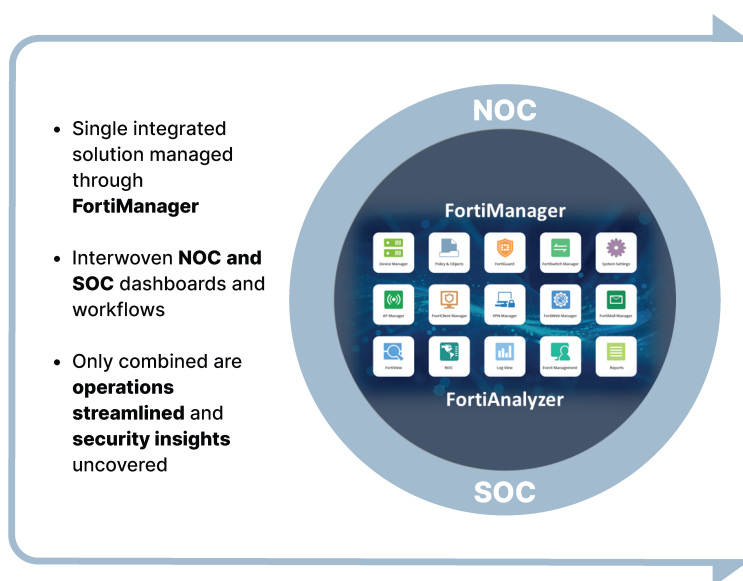


Considerations:

- At least one static gateway is required to act as ADVPN sender. This gateway, or gateways, will act as the dialup server for branches and cloud gateways.
- In this design, cloud gateways are treated similarly to a branch dialup client. This lets the gateway initiate the IPsec connection to the gateway when it comes online.
- Cloud gateways advertise their networks to the ADVPN sender (the datacenter gateway in this diagram), which in turn makes those networks accessible to other branches by using BGP.
- When a branch needs a service or resource behind the cloud gateway, it will dynamically connect through ADVPN directly.
- For the correct operation of ADVPN, it is required to preserve all sites' prefixes unchanged, including their original BGP next-hop values. Hence, it is impossible to replace the specific routes with summaries (unlike in a static hub-and-spoke topology). Hence, the BGP RR function is mandatory: the gateway must reflect the original routes between the spokes without altering them.

Monitoring and reporting

As detailed earlier in this document, FortiAnalyzer and FortiManager are tools we can use to monitor and manage our SD-WAN and SD-Branch devices from a single pane of glass. FortiManager is used for operational features, while the FortiAnalyzer provides deeper analytics and reporting. When FortiAnalyzer is integrated with FortiManager, you can use a single pane of glass for your network and security operations.



FortiManager monitoring and FortiAnalyzer analytics provide complimentary features that satisfy both NOC and SOC requirements. Following is a very high-level overview of monitoring versus analytic capabilities.

	FortiManager Monitoring	FortiAnalyzer Analytics
Communication Method	API	Logging
Device uptime and availability	x	
System health monitoring	x	
Bandwidth overview	x	
FortiGate route table monitoring	x	

	FortiManager Monitoring	FortiAnalyzer Analytics
Communication Method	API	Logging
IPsec tunnel monitoring	x	
ADVPN monitoring	x	
SD-WAN rule selection and steering monitoring	x	
Performance SLA monitoring	x	
Application performance	x	x
Per-application bandwidth consumption		x
Single page network summary		x
SD-WAN rule and interface utilization		x
Advanced SD-WAN analytics (per device and network)		x
Pre-built or custom reporting		x
Custom log dashboard		x
Long term log storage		x

While FortiManager has basic FortiAnalyzer capabilities, it is generally recommended to leverage a dedicated FortiAnalyzer appliance for deeper analytics, log storage, and reporting.

This section includes the following topics:

- [FortiManager and SD-WAN monitoring on page 76](#)
- [FortiAnalyzer on page 80](#)

FortiManager and SD-WAN monitoring

As mentioned previously, FortiManager provides central management and operations functionality from a single pane of glass. In this section, we will focus on the SD-WAN monitoring capabilities for some operational scenarios.

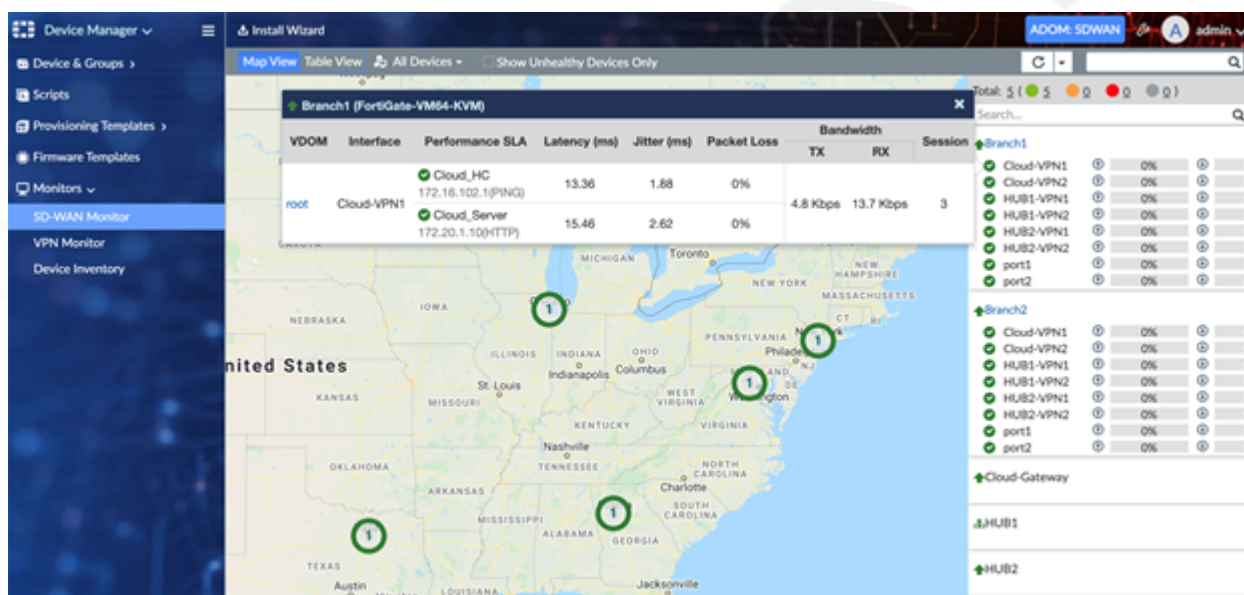
This section contains the following topics:

- [SD-WAN Monitor Map view on page 76](#)
- [SD-WAN Table view on page 77](#)
- [SD-WAN device monitoring on page 78](#)
- [SD-WAN device monitoring of performance SLAs on page 78](#)
- [Route table and device dashboards on page 79](#)

SD-WAN Monitor Map view

From the SD-WAN Monitor Map, you can view your network availability and performance from a single glance. SD-WAN branch and gateways are represented based on the geo-coordinates configured for their

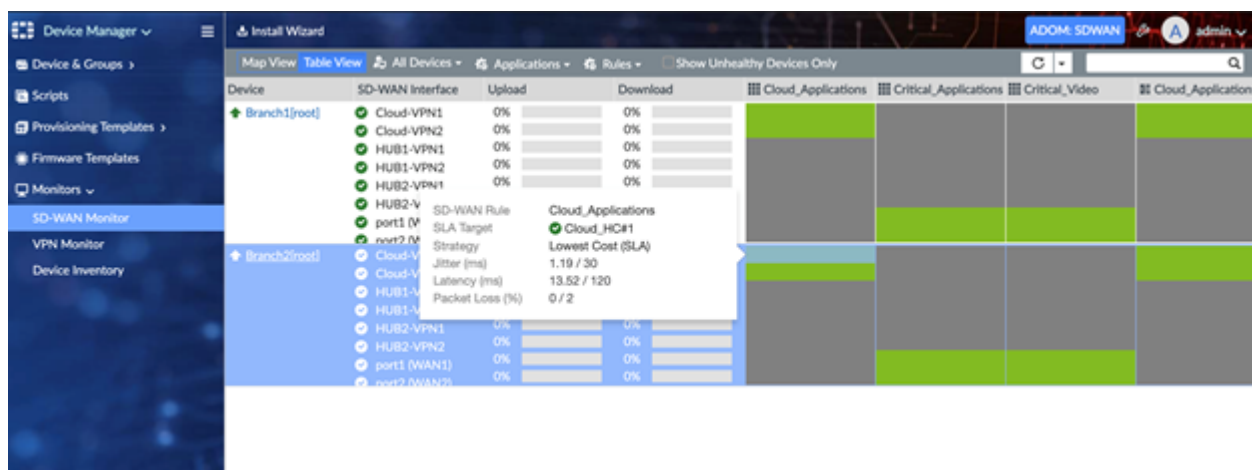
locations, and colored indicators notify you of any potential issues with the SD-WAN interfaces or SLAs. Only devices with performance SLAs are displayed.



- **Green Indicators:** all performance SLAs for the given interface are being met.
- **Yellow:** one of more performance SLAs is not currently meeting minimum requirements.
- **Red:** one of more performance SLAs is down or unreachable.

SD-WAN Table view

When Table View is selected, all devices in the network are displayed in a table format. This gives you a quick snapshot of link performance across all devices. Like Monitor Map view, interface colors represent the status of a given interface, according to its performance SLA. Hovering over a given performance SLA, you can see more granular detail about how that SLA is performing according to its configuration.



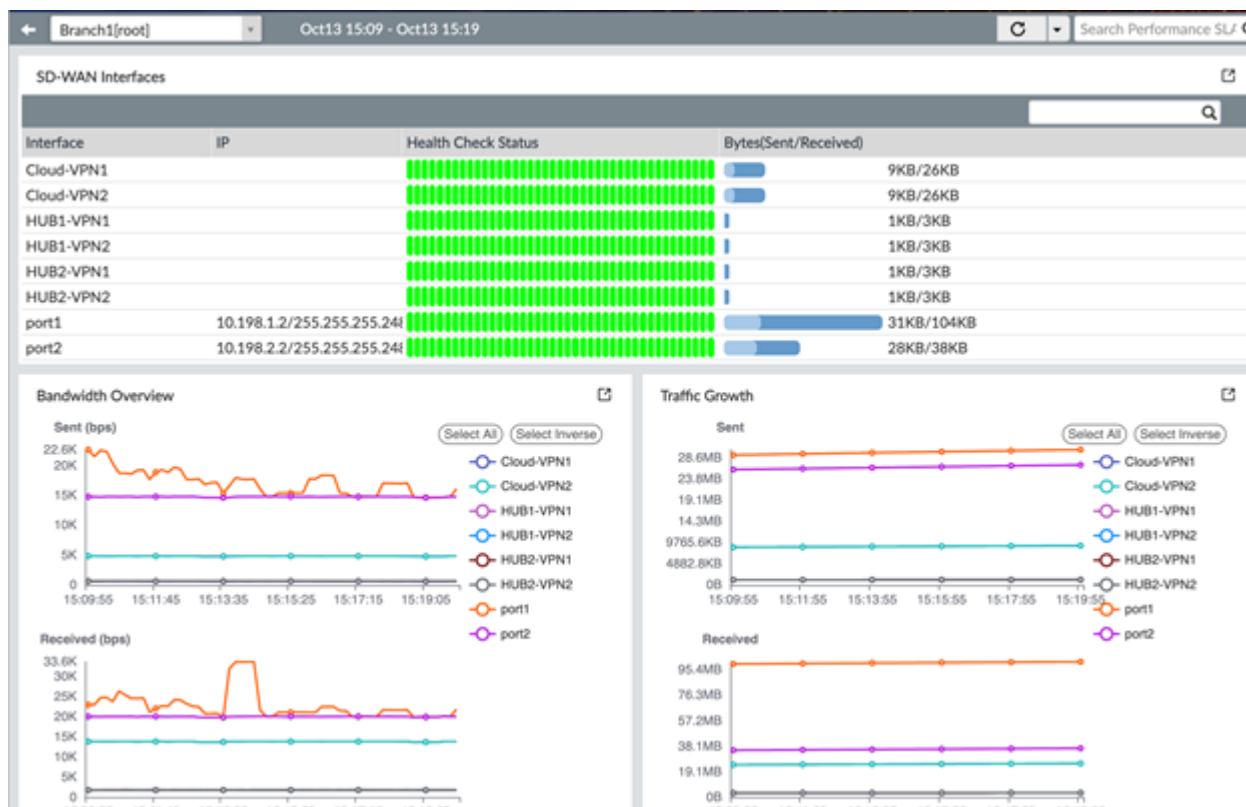
- **Green Indicators:** all performance SLAs for the given interface are being met.
- **Yellow:** one of more performance SLAs is not currently meeting minimum requirements.
- **Red:** one of more performance SLAs is down or unreachable.

SD-WAN device monitoring

Selecting a device from the Monitor Map or Table view will give you details about that specific device. This is often one of the first places to check when you are troubleshooting a potential network issue. Any links not meeting their minimum defined SLAs are displayed, which helps you narrow down the scope of your investigation.

Within the selected device, you can see per-device metrics, such as:

- Interfaces participating in SD-WAN
- Health-check status over the specified time frame
- Bytes sent/received
- Bandwidth overview per interface
- Traffic growth per interface



SD-WAN device monitoring of performance SLAs

Further down the window, you will also see all performance SLAs configured for a given device. The data is displayed according to the timeframe that you have selected. By default, the timeframe is 10 minutes, but can be extended by enabling `sdwan-monitor-history` by using the CLI. All data for the selected device will refresh according to the refresh interval selected at the top.



Route table and device dashboards

There will be occasions where you may need to see more granular information on the device, such as looking up a route or checking IPsec details. Managed FortiGate SD-WAN devices will provide visibility into live monitors that can be quickly accessed from FortiManager. These are equivalent to the dashboards that can be configured and viewed from the device Admin page on the local box.

Device Manager

Device & Groups

Managed FortiGate (5)

- Branch1
- Branch2
- Cloud-Gateway
- HUB1
- HUB2

Managed FortiAnalyzer (1)

- Branches (2)
- Cloud (1)
- Datacenter (2)

Scripts

Provisioning Templates

Firmware Templates

Monitors

Dashboard

Network Monitors

SDWAN

System

Router

WAN Opt. & Cache

Security Profiles

VPN

CLI Configurations

Display Options

IPsec VPN

Column Settings

Name	Remote Gateway	Peer ID	Incoming Data	Outgoing Data	Phase 1
root (0)					
Cloud-VPN1	10.198.9.2		25.4 MB	8.9 MB	Cloud-VPN1
Cloud-VPN2	10.198.9.2		25.4 MB	8.9 MB	Cloud-VPN2
HUB1-VPN1	10.198.5.2		3.3 MB	1.2 MB	HUB1-VPN1
HUB1-VPN2	10.198.6.2		3.3 MB	1.2 MB	HUB1-VPN2
HUB2-VPN1	10.198.7.2		3.2 MB	1.2 MB	HUB2-VPN1
HUB2-VPN2	10.198.8.2		3.2 MB	1.2 MB	HUB2-VPN2

Routing - Static & Dynamic

Route Lookup

Column Settings

Network	Gateway IP	Interfaces	Distance	Type
root (36)				
0.0.0.0/0	10.198.1.1	port1 (WAN1)	10	static
0.0.0.0/0	10.198.2.1	port2 (WAN2)	10	static
10.1.1.0/24	0.0.0.0	port3 (LAN)	0	Connected
10.1.2.0/24	10.198.5.2	HUB1-VPN1	200	BGP
10.1.2.0/24	10.198.6.2	HUB1-VPN2	200	BGP
10.1.2.0/24	10.198.7.2	HUB2-VPN1	200	BGP
10.1.2.0/24	10.198.8.2	HUB2-VPN2	200	BGP
10.10.10.0/24	10.198.5.2	HUB1-VPN1	15	static

FortiAnalyzer

This section about FortiAnalyzer covers the following topics:

- ADOMs, sizing, log storage, scaling, and enforcement on page 80
- SD-WAN logging on page 81
- FortiAnalyzer HA recommendation on page 87

ADOMs, sizing, log storage, scaling, and enforcement

FortiAnalyzer is the central log correlation engine for many Fortinet technologies, including SD-Branch (FortiGate, FortiSwitch, FortiAP), FortiClient, FortiSandbox, FortiMail, and others providing a centralized intelligence center with each of these components sending logs to FortiAnalyzer. FortiAnalyzer is responsible for log indexing (online logs) and archival (compressed logs), which can all be specified on a per-customer (ADOM) basis.

When deploying a multitenant FortiAnalyzer, MSPs should standardize on maximum log analytics (60 days in the below example) and archival periods (365 days in the below example) for each ADOM. With FortiAnalyzer being licensed based on GB of logs per day (a system-wide limit) and ADOMs (when using the FortiAnalyzer subscription license), this standardization ensures MSPs know the maximum number of customer tenants accommodated by the shared platform.

Furthermore, the MSP should also factor in the maximum number of recommended ADOMs, based on the deployed license and minimum server specification. FortiAnalyzer minimum system requirements are available at docs.fortinet.com.

System Settings ▾

- Dashboard
- Logging Topology
- All ADOMs
- Storage Info
- Network
- HA
- Admin ▾
 - Administrators
 - Profile
 - Remote Authentication Server
 - Admin Settings
 - SAML SSO
- Certificates ▾
 - Local Certificates
 - CA Certificates
 - CRL
 - Remote Certificates
- Log Forwarding
- Fetcher Management

Create New ADOM

Name:

Type:

Comments:

Devices

Name	IP Address	Platform
No Device.		

Data Policy

Keep Logs for Analytics: Days

Keep Logs for Archive: Days

Disk Utilization

Allocated: MB

Analytics : Archive:

Alert and Delete When Usage Reaches:

Maximum Available: 742.2 GB

☐ Modify

*If analytic or archive log usages exceed the configured disk quota before the retention period expires, the oldest logs will be deleted.

The above image shows the creation of an ADOM called *MSP_SD-Branch_CustomerA*, where parameters such as analytics, log archival, and disk space are defined on a per-customer basis.

When standardizing on a multitenant platform, the MSP should ensure the parameters detailed above are then written into the overall service level agreement between MSP and end-customer.

This standardization ensures platform sizing and scalability are tested and documented, and avoids situations where non-standard target customers could impact others on the shared platform. For example, imagine a shared FortiAnalyzer whereby one tenant (ADOM) manages a 20-site SD-WAN deployment. Each branch site caters to 20 concurrent users, which is representative of a typical customer on the multitenant platform. Suppose a non-standard customer requires a 2,000-branch SD-WAN solution, with each branch having 100 concurrent users. In that case, they will consume a disproportionate amount of the shared platform resource, causing performance and bottleneck issues for the remaining tenants. Subsequently, this large tenant should be deemed as non-standard and therefore not placed on the shared platform.

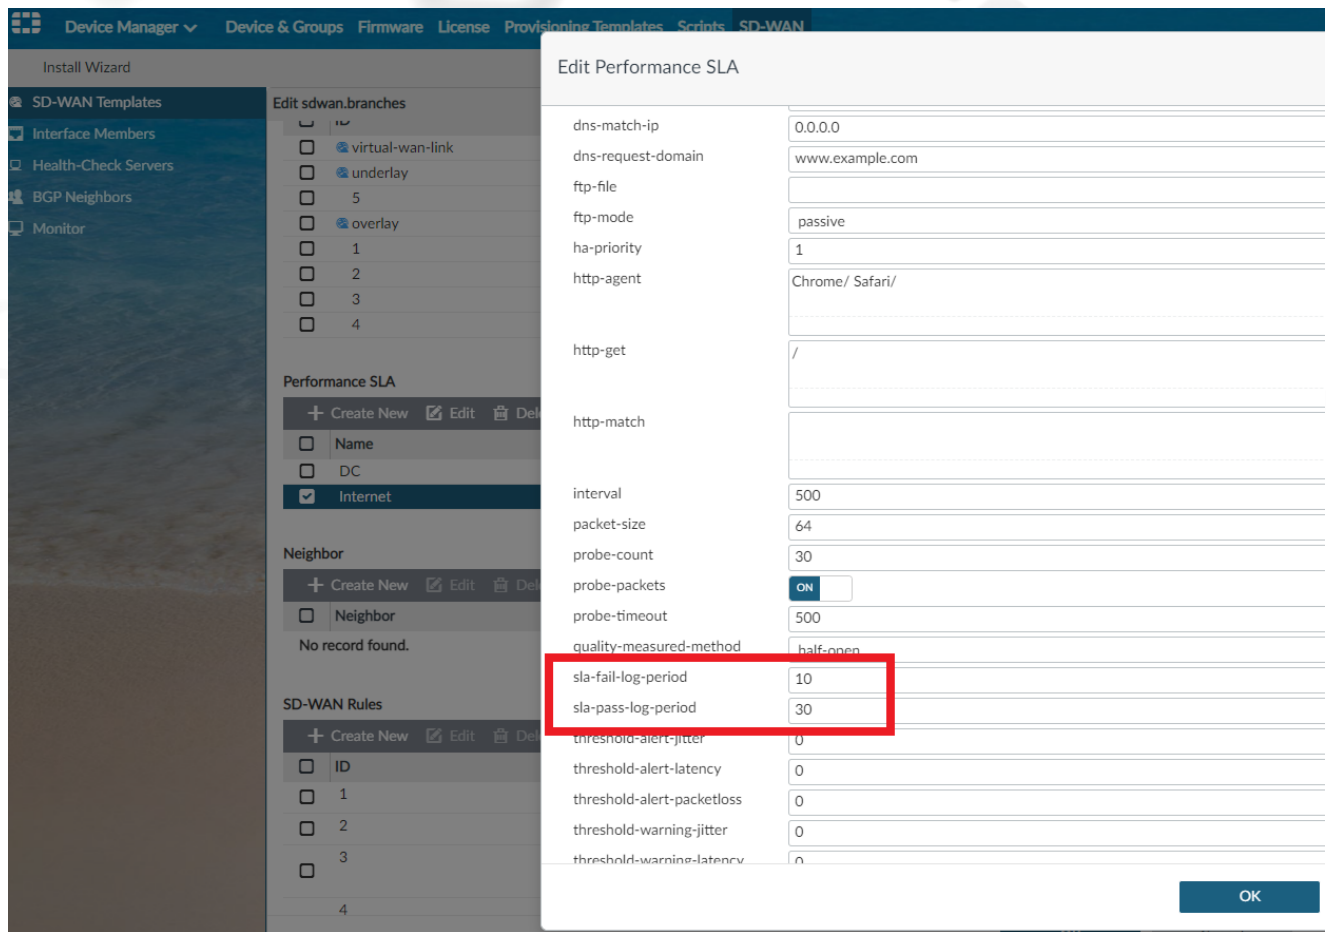
FortiAnalyzer logs are sized based on analytic and archival logs. Analytic logs are classified as indexed/non-compressed, active, and available for log querying through FortiView and reporting. These analytics logs are sized at 400 bytes per log. Archived/compressed logs are offline and sized at 40 bytes. Therefore, these log size variables should be added into a common equation across all ADOMs when sizing the multitenant FortiAnalyzer.

Fortinet Partners have access to the FortiAnalyzer sizing tool hosted on [Fortinet Developer Network \(FNDN\)](#). It can aid in estimating logging rates inclusive of storage on a per-customer basis. The partner can use known logging rates or estimates based on known customer parameters, such as the number of users, sessions per second, and office hours. Furthermore, the sizing tool can also add layered security service logging, such as application control and web filtering, into the overall calculation.

SD-WAN logging

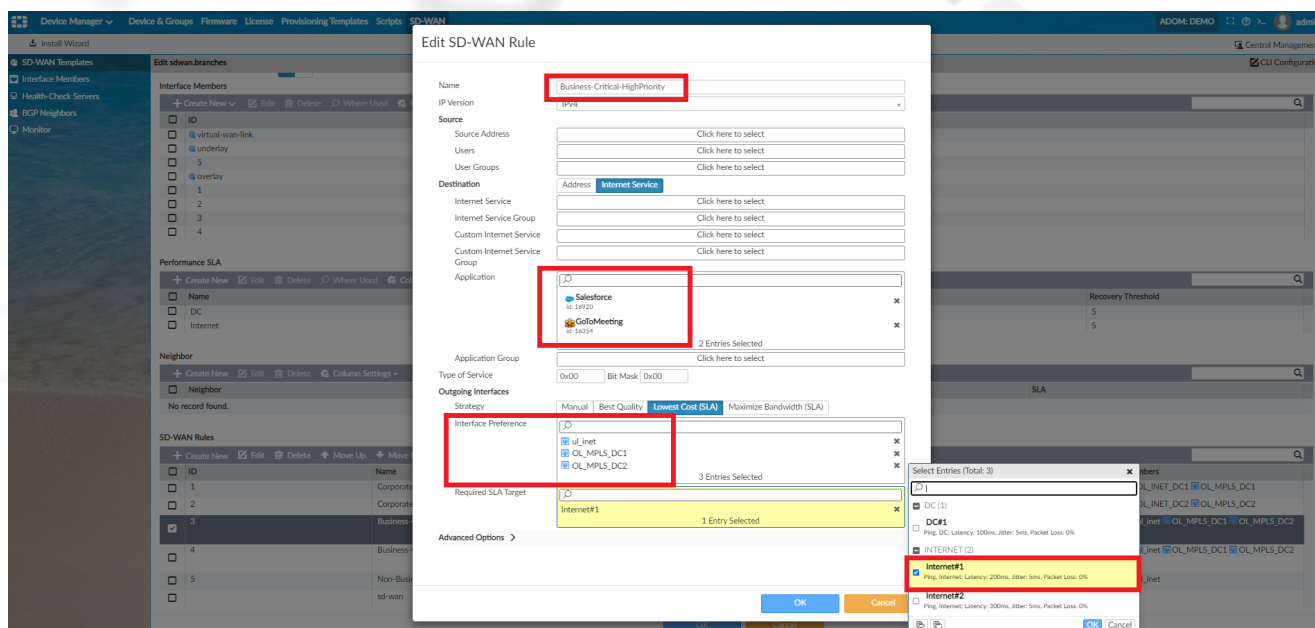
Now that we understand FortiAnalyzer acts as the central monitoring platform, let's look at the log types. As outlined in previous sections, FortiGate acts as the branch CPE in the SD-WAN solution. It utilizes SLA probes across the overlays to record latency, jitter, and packet loss.

FortiAnalyzer requires logs from the branch FortiGate with latency, jitter, and packet loss information to create and display SD-WAN graphs. It is mandatory to specify the sending interval, which is configured in the FortiManager SD-WAN template. The sending interval is configured using `set-fail-log-period` (seconds) and `set-pass-log-period` (seconds). The below example shows that the value is set to 30 seconds for passing probes and 10 seconds for failing probes. This means that when the SLA is above target (pass), FortiGate will send a log every 30 seconds with information on pass SLA. When the SLA is below target (fail), FortiGate will send a log every 10 seconds, with information on fail SLA.



In the next example below, SD-WAN rule *Business_Critical-HighPriority* uses the SLA *Internet#1*, which has 200ms latency and 5ms jitter set as thresholds. This means that a probe (a ping, DNS, HTTP, or others) is being sent at a specified time period, every 500ms being the default, across SD-WAN member interfaces listed in the SD-WAN rule. Traffic matching *GoToMeeting* and *Salesforce* are being sent via the native direct internet access (DIA) interface, called *ul_inet*, as a priority before trying the two overlay links over MPLS, which will break out centrally should the DIA either fail or hit a brownout.

The default SD-WAN interface selection method for the SD-WAN criteria *Lowest Cost SLA*, where cost is not defined on the member interfaces, is always top-down. Therefore, this rule will try *OL_MPLS_DC1* first (if currently within SLA) should the native *ul_inet* interface be in a brownout state, and then *OL_MPLS_DC2*, but only if both *ul_inet* and *OL_MPLS_DC1* are still out of SLA.



Let's look at how the various logs sent from FortiGate to FortiAnalyzer look from the CLI.

When a performance SLA detects a link failure, it will record a log:

- ```
date=2021-02-18 time=09:38:41 id=6930520380335456274 itime=2021-02-18 09:38:41 euid=3
epid=3 dsteuid=3 dstepid=3 logid=0100022921 type=event subtype=system level=critical
msg=Static route on interface BBI may be removed by health-check nonBC_streaming.
Route: (82.197.160.199->52.213.155.117 ping-down) (82.197.160.199->172.217.168.14 ping-down)
```

When health-check detects a recovery, it will record a log:

- ```
date=2021-02-18 time=09:38:50 id=6930520427580096515 itime=2021-02-18 09:38:52 euid=3
epid=3 dsteuid=3 dstepid=3 logid=0100022921 type=event subtype=system level=critical
msg=Static route on interface BBI may be added by health-check nonBC_streaming. Route:
(82.197.160.199->52.213.155.117 ping-down) (82.197.160.199->172.217.168.14 ping-up)
```

When health-check has an SLA target, and detects SLA changes, and changes to fail:

- ```
date=2020-04-11 time=11:48:39 logid=" 0113022923 " type="event" subtype="sdwan"
level="notice" vd="root" eventtime=1555008519816639290 logdesc="Virtual WAN Link
status" msg="SD-WAN Health Check(ping) SLA(1): number of pass members changes from 2 to
1."
```

When health-check has an SLA target, and detects SLA changes, and changes to pass:

- ```
date=2020-04-11 time=11:49:46 logid=" 0113022923 " type="event" subtype="sdwan"
level="notice" vd="root" eventtime=1555008586149038471 logdesc="Virtual WAN Link
status" msg="SD-WAN Health Check(ping) SLA(1): number of pass members changes from 1 to
2."
```

Now let's look at where logs are displayed in FortiAnalyzer, and how they are used in the various monitors.

Navigating to the *FortiAnalyzer* > *Log View* > *Event-SD-WAN*, we can see the logs being received across all overlays for all managed devices within the FortiAnalyzer ADOM named *DEMO*. This provides a wealth of detail on performance.

OL_MPLS_21 overlay is highlighted in the below image. It shows jitter/latency/packet loss, together with additional log details on the right.

Log View

All FortiGate Last 1 Hour 13:17:59 To 14:17:58

Add Filter

#	Date/Time	Level	Device ID	Interface	Status	Message	Jitter	Latency	Packet Loss
1	14:17:34	Information	FGVM02TM20011214	port1	up	Health Chec...	2.269	27.717	0.000%
2	14:17:34	Information	FGVM02TM20011214	port1	up	Health Chec...	2.269	27.717	0.000%
3	14:17:34	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.633	13.723	0.000%
4	14:17:34	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.633	13.723	0.000%
5	14:17:34	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.552	13.646	0.000%
6	14:17:34	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.552	13.646	0.000%
7	14:17:34	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.374	2.786	0.000%
8	14:17:34	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.201	2.673	0.000%
9	14:17:29	Information	FGVM02TM20011162	port1	up	Health Chec...	1.200	11.792	0.000%
10	14:17:29	Information	FGVM02TM20011162	port1	up	Health Chec...	1.200	11.792	0.000%
11	14:17:29	Information	FGVM02TM20011162	OL_MPLS_22	up	Health Chec...	1.449	13.907	0.000%
12	14:17:29	Information	FGVM02TM20011162	OL_MPLS_22	up	Health Chec...	1.449	13.907	0.000%
13	14:17:29	Information	FGVM02TM20011162	OL_MPLS_21	up	Health Chec...	1.421	13.879	0.000%
14	14:17:29	Information	FGVM02TM20011162	OL_MPLS_21	up	Health Chec...	1.421	13.879	0.000%
15	14:17:29	Information	FGVM02TM20011162	OL_MPLS_22	up	Health Chec...	1.565	3.030	0.000%
16	14:17:29	Information	FGVM02TM20011162	OL_MPLS_21	up	Health Chec...	1.498	3.099	0.000%
17	14:17:29	Information	FGVM02TM20011162	OL_INET_12	up	Health Chec...	1.226	3.205	0.000%
18	14:17:29	Information	FGVM02TM20011162	OL_INET_11	up	Health Chec...	1.586	3.572	0.000%
19	14:17:28	Information	FGVM02TM20011162	OL_INET_12	up	Health Chec...	0.721	18.321	0.000%
20	14:17:28	Information	FGVM02TM20011162	OL_INET_11	up	Health Chec...	2.829	19.474	0.000%
21	14:17:22	Information	FGVM02TM20011214	port1	up	Health Chec...	3.065	28.212	0.000%
22	14:17:22	Information	FGVM02TM20011214	port1	up	Health Chec...	3.065	28.212	0.000%
23	14:17:22	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.857	14.070	0.000%
24	14:17:22	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.857	14.070	0.000%
25	14:17:22	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.741	13.982	0.000%
26	14:17:22	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.741	13.982	0.000%
27	14:17:22	Information	FGVM02TM20011214	OL_MPLS_22	up	Health Chec...	1.317	2.729	0.000%
28	14:17:22	Information	FGVM02TM20011214	OL_MPLS_21	up	Health Chec...	1.209	2.597	0.000%
29	14:17:19	Information	FGVM02TM20011162	port1	up	Health Chec...	0.590	11.573	0.000%
30	14:17:19	Information	FGVM02TM20011162	port1	up	Health Chec...	0.590	11.573	0.000%

Total logs for analytics: 1 hour.

50 Items per page 1 2 3 4 5 0.079 Second

Network Properties

Interface

Identity

Device ID

Device Name

Type

Sub Type

Level

General

Log Description

Log ID

Message

Status

Virtual Domain

Others

Bandwidth

Date/Time

Destination End User ID

Destination Endpoint ID

Device Time

Event Time

Event Type

Health Check

In Bandwidth

Jitter

Latency

Out Bandwidth

Packet Loss

SLA Map

SLA Target ID

Time Stamp

Time Zone

UEBA Endpoint ID

UEBA User ID


Used Bandwidth

These logs are then used to populate the following displays within the *FortiView* > *SD-WAN Monitor* section:

- **SD-WAN Bandwidth Overview:** Bandwidth usage overview per interface
- **SD-WAN Rule Utilization:** SD-WAN rule traffic utilization by interface and application
- **SD-WAN Performance Status:** Performance of the SD-WAN and each WAN link in the network over time
- **Jitter:** Number of seconds for disruption in the data flow across the network for each WAN link over time
- **Latency:** Number of seconds for a packet of data to travel across the network for each WAN link over time
- **Packet loss:** Percentage of network data that failed to reach its intended destination for each WAN link over time
- **Bandwidth Utilization by SD-WAN Rules:** Share of bandwidth utilization for each configured SD-WAN rule
- **SD-WAN Utilization by Application:** Share of bandwidth utilization by application for each WAN link
- **SD-WAN High and Critical Events:** Existing alarms on path, connection, or individual WAN links for their states (*Information*, *Notice*, and *Warning*)

But also to populate the *FortiView* *SD-WAN Summary* page, which provides a global view of all devices:

- **SD-WAN Health Overview:** Overview of the device health status (*Healthy*, *Major Alerts*, *Critical Alerts*)
- **Top SD-WAN SLA Issues:** Worst SLA amongst all the branches
- **Top SD-WAN Applications:** Most bandwidth-consuming applications
- **Top SD-WAN Device Throughout:** Most bandwidth-consuming branches
- **Top SD-WAN Talkers:** Most bandwidth-consuming clients



Following the introduction of the **Passive WAN Health Measurement** feature, FortiAnalyzer can also display a chart of passively monitored applications and the associated telemetry.

FortiAnalyzer provides a comprehensive SD-WAN reporting section, all the reports are fully customizable to meet both MSP and end-customer branding. The following image shows some of the reports included in FortiAnalyzer as well as one page of the SD-WAN report as an example.


SLA Rules Link Percentage Within Jitter Threshold

#	SLA Rules	Links	Jitter Within Threshold
1	HUB	H1_ISP1	100.00%
		H1_MPLS	100.00%
2	Internet	H1_MPLS	100.00%
		port1	100.00%


SLA Rules Link Percentage Within Packet Loss Threshold

#	SLA Rules	Links	Packet Loss Within Threshold
1	HUB	H1_MPLS	100.00%
		H1_ISP1	99.88%
2	Internet	port1	99.87%
		H1_MPLS	99.86%


Latency by SLA Rule Over Time (ms)



Jitter by SLA Rule Over Time (ms)



Packet Loss by SLA Rule Over Time



FortiAnalyzer HA recommendation

When deploying FortiAnalyzer in a multitenant environment, high availability (HA) should be considered. This HA consists of a minimum of two FortiAnalyzer units to a maximum of four FortiAnalyzer units. These FortiAnalyzer units are configured in VRRP HA. The FortiGate(s) send the logs to the VIP or FQDN set on the FortiAnalyzer VRRP HA deployment.

A FortiAnalyzer HA cluster provides the following features:

- Provides real-time redundancy in case a FortiAnalyzer primary unit fails. If the primary unit fails, another unit in the cluster is selected as the primary unit.
- Synchronizes logs and data securely among multiple FortiAnalyzer units. Some system and configuration settings are also synchronized.
- Alleviates the load on the primary unit by using secondary (backup) units for processes, such as running reports and FortiView dashboards.

A FortiAnalyzer HA cluster can have a maximum of four units, one primary unit with up to three secondary units. All units in the cluster must be the same FortiAnalyzer model. They need to be in the same network and running in the same operation mode: Analyzer or Collector.

For more details on the Analyzer or Collector mode, see the [FortiAnalyzer Admin Guide](#).

Cluster Status

Role	Serial Number	IP	Host Name	Uptime/Downtime	Initial Logs Sync	Configuration Sync
Primary	FAZ-VM20013011	10.1.1.1	dut_faz	03h 23m 39s	-	Config will be synced to
Secondary	FAZ-VM20014390	10.1.1.2	FAZVM64-KVM	03h 23m 38s	Done	In-Sync

Cluster Settings

Operation Mode: ☐ Standalone ☒ High Availability

Preferred Role: ☒ Primary ☐ Secondary

Cluster Virtual IP

Interface: port1

IP Address: 192.168.244.222

Cluster Settings

Peer IP and Peer SN	Peer IP	Peer SN
	10.1.1.2	FAZ-VM20014390

Group Name: fortinet

Group ID: 1 (1-255)

Password:

Heart Beat Interval: 1 Seconds

Failover Threshold:

Priority: 100 (80-120)

Even though it is an active/passive HA setup, the secondary FortiAnalyzer(s) still participates in a round-robin load share for report creation and SQL query—used to populate the various FortiView dashboards. The main benefit of this mode is the overall performance improvement.

Evolution to secure SD-branch solution

The branch office itself, usually without on-site IT staff, needs to be monitored and protected. Today's next-generation branch offices not only require the same functionality, but they also suffer from the same risks as the rest of the distributed network. Direct access to the internet and SaaS applications, for example, significantly expand the potential attack surface of the branch, as does the growing proliferation of IoT and BYOD devices, creating multiple network edges beyond the WAN edge. This explosion of edges, which all must be secured, is causing many organizations to struggle to implement adequate security throughout their distributed enterprises, including at the new branch. The complexity of managing these edges—including often complicated and overlapping point products and appliances—adds an additional challenge. As a result, organizations adopting SD-WAN are finding that they need to find a vendor that can more tightly integrate their SD-WAN security and management functionality into their branch networks.

Fortinet is delivering the industry's first complete Secure SD-Branch solution to combat this challenge, enabling customers to converge security and network access, and extend the Fortinet Security Fabric to the branch. This new SD-branch solution is comprised of the following elements:

- **FortiGate Next-Generation Firewall:** provides robust security, connectivity, and management across the branch environment. The FortiGate NGFW also includes the industry's first purpose-built SD-WAN processor, combined with advanced network traffic management functionality, such as application steering to ensure high application performance on any WAN link. The FortiGate solution also includes advanced sensor functionality for increased device visibility and traffic anomaly detection without additional hardware.
- **FortiSwitch and FortiAP:** provide consolidation of branch services through the convergence of security and network access with FortiLink. FortiSwitch and FortiAP integrate with FortiGate to extend SD-WAN's benefits into the network access layer. This enables network administrators to create and enforce the same network security policies across the enterprise, including out to the network branch.

With the combination of the above technologies, a more comprehensive number of use cases are enabled:

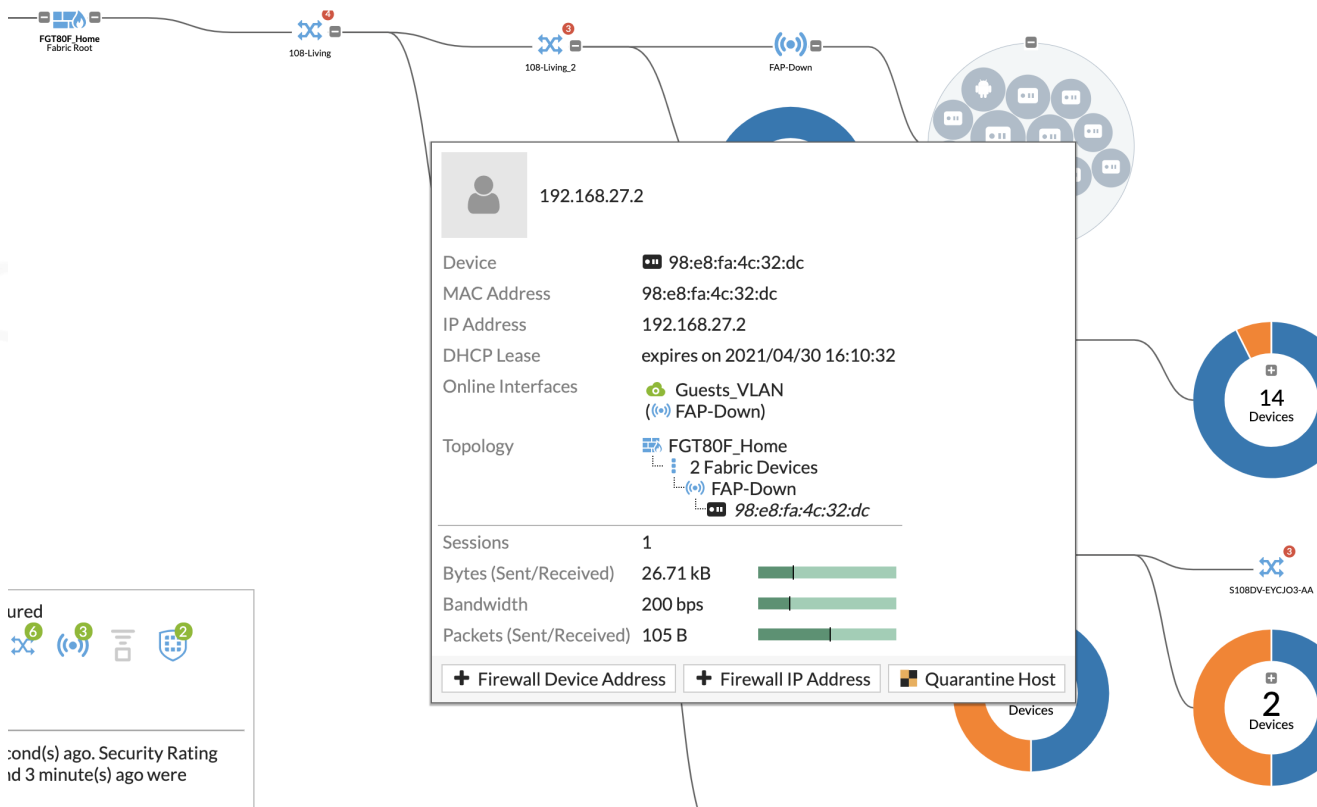
- [Visibility on page 88](#)
- [Attack surface reduction with network segmentation on page 89](#)
- [Zero trust local access network on page 90](#)
- [SD-branch simplification on page 91](#)

Visibility

Security starts with visibility. Adding an additional source of information under the single-pane-of-glass management system provides insightful information about the devices connected on the LAN edge. This rich information includes:

- SSIDs, ports, MAC addresses, OS information, Hostname, IP addresses, Device Type, Hardware vendor, User
- Extended search capabilities
- Indications of compromise at first sight

Controlling what's connected to the network is the first step to secure it.



Attack surface reduction with network segmentation

This is an essential part when it comes to securing the LAN edge. Being able to divide the network into different segments helps reduce the attack surface instantly, and minimizes the potential spread of a security breach and lateral movements.

With network segmentation, each VLAN becomes its own realm. And by being considered another FortiGate interface, it can be used in the firewall policies to enable communications control. Moreover, leveraging the interface consideration, the FortiGate can also extend different levels of prioritization for different segments into the SD-WAN.

ZERO TRUST LOCAL ACCESS NETWORK

The screenshot displays the FortiSwitch configuration interface. At the top, there are tabs for various network profiles: default, quarantine, voice, video, rspan, onboarding, vlan100, office-vlan200, voice-vlan250, fac-radius, guest-vlan, authenticated, auth-fail, employees, it-vlan, lab-vlan, fap-vlan172, dot1x-vlan249, Fortilink Over Layer3 (Fol3), no-ip, rg-test-1800, rg-test-1801, sector-1759, sector-1615, FPOC_HQ_WAN, test-rg, FNAC_ETH1_ISO, and FNAC_ISOLATION. Below these tabs, there are sections for different switch models and their port status:

- FS1D243Z14000285:** SFP+ ports 1-24, MGMT port. Status: Connected.
- FS1D243Z14000301:** SFP+ ports 1-24, MGMT port. Status: Connected.
- S108EF4N17000370:** POE+ ports 1-10, SFP+ ports 9-10. Status: Offline.
- S108EN4N17000487:** SFP ports 1-10. Status: Connected.
- S108EP5918000293:** POE+ ports 1-10, SFP ports 9-10. Status: Connected.

At the bottom, there is a PoE Total Power budget section showing 65.00W and 61.60W Unallocated. Below this is a table of SSIDs:

Name	SSID	Traffic Mode	Security	Schedule	Status	Ref
Guests	Red_Invitados (Guests)	Local Bridge	WPA2 Personal	always	Up	3
ipcam	IPCam (ipcam)	Local Bridge	WPA2 Personal	always	Up	3
wifi	Matrix Secured (wifi)	Local Bridge	WPA2 Personal	always	Up	4

Taking this one step forward, the FortiSwitch enables **microsegmentation** to isolate every device, even within the same VLAN. No direct visibility among the devices is allowed, and all flows are forced through the FortiGate, where communications decisions can be made based on policy.

Zero trust local access network

Implementing security access control is straightforward with FortiSwitch, dynamically preventing unknown devices from gaining access to the network.

There are several features that could help to achieve this goal:

- **FortiGate NAC:** this built-in capability works alongside FortiSwitch and does not require any additional license. It enables the mapping of devices into VLANs depending on the device type. Unrecognized devices can be assigned to a guest VLAN with limited access. Moreover, it allows the dynamic configuration of ports based on the matching criteria of different parameters (MAC address, OS, device type, user). Multiple policies can be applied to map different devices to their corresponding settings: LLDP profile, 802.1x, QoS, VLAN.
- **User authentication with 802.1X:** implementing a user or MAC address bypass at the port or MAC level allows different devices to connect by authenticating them against a RADIUS server or FortiAuthenticator.

- **LLDP profiles:** configuring devices detected by LLDP automatically, assigning them to specific VLANs and QoS marking.

Dashboard >

Security Fabric >

Network >

System >

Policy & Objects >

Security Profiles >

VPN >

User & Authentication >

WiFi & Switch Controller >

Managed FortiAPs

WiFi Clients

WiFi Maps

SSIDs

FortiAP Profiles

WIDS Profiles

FortiLink Interface

Managed FortiSwitch

FortiSwitch VLANs

FortiSwitch Ports

FortiSwitch NAC Policies ☆

FortiSwitch Security Policies

Log & Report >

Edit NAC Policy

NameAndroid_devices

Status

Enabled

Disabled

FortiSwitches

All

Specify

Description

0/63

If device matches all of the following patterns:

Category

Device

User

EMS Tag

MAC address

Hardware vendor

Device family

Type

Phone

Operating system

Android

User

Then:

Select an action that will be performed to the matched device.

Assign VLAN

Assign a specific VLAN to a device matching above patterns.

Apply Port Specific Settings

Apply a LLDP Profile, QoS Policy, 802.1X Policy, or VLAN Policy.

LLDP profile

QoS policy

802.1X policy

802.1X

802-1X-policy-default

VLAN policy

VLAN Policy

Guest_LAN

SD-branch simplification

When addressing the SD-branch deployment, one of the primary considerations is to make it easy and fast by taking advantage of zero touch provisioning approaches.

Thanks to the integration of FortiSwitch and FortiAP in FortiManager, the normalization of a configuration can be defined once and then replicated throughout all the branches of a given corporation. This implies that all branches should be similar to maximize their benefits.

The following scenario describes the ideal situation:

1. Creation of templates per SD-Branch on FortiManager using variables and model devices.
2. Shipping corresponding gear to remote sites, and having someone with no networking or security background connect the devices.
3. Remote devices power up, and automatically trigger a call-home procedure to reach the FortiManager.
4. Once discovered by FortiManager, the devices get provisioned according to their preconfigured setup. That is the end of the deployment.

If standardization for SD-Branches is not possible, FortiManager also supports per-device configuration for FortiSwitch, which provides the capability to manage each FortiSwitch independently, as if directly configured from a FortiGate. It can also define specific SSID Groups to be distributed on some sites and not on others.

All the benefits described above are also present on FortiManager. All elements can be deployed through its single-pane-of-glass console, and connected devices can be displayed in its Security Fabric views.



www.fortinet.com



Copyright© 2024 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., 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 Chief Legal Officer, 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. 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.