FÜRTINET®

Architecture for MSSPs

SD-WAN / SD-Branch

4D

DEFINE / DESIGN / DEPLOY / DEMO

Table of Contents

Change Log

Date	Change Description
2021-09-27	Initial release.

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 an *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.

This section contains the following topics:

Legacy SD-WAN

For decades now, the traditional architecture of *hub-and-spoke* has served distributed enterprises of all sizes well. Typically, in such architectures, network traffic flows through a central corporate data center — including traffic moving from branch locations to the public internet.

The challenges we are witnessing today with hub-and-spoke designs are confined to the transport mechanisms, not the fundamental structure of hub-and-spoke. The WAN has drastically changed in its consumption, availability, and scalability.

More commonly referred to as SD-WAN, the software-defined WAN sets out to challenge the network designs that have largely remained unchanged for decades. Technological advancements to WAN technologies, alongside a shift in network consumption dynamics, have forced network designers to revisit the tried and tested hub-and-spoke architecture.

The unprecedented growth of SD-WAN is not surprising. Technological transformation of any kind should save money for a business or increase a performance metric. SD-WAN sets out to do both. When implemented correctly, these benefits can be realized within 12 months. This is precisely why the year-on-year growth of SD-WAN shows little sign of stagnation.

Cloud adoption, be it software, infrastructure, or Platform-as-a-Service, and the proliferation of Internet-of-Things (IoT) devices at the network edge are finding a permanent footing within organizations. This contrasts with the traditional HQ data center, where applications, services, and platforms traditionally reside. While the legacy data centers and their traffic flows will remain for the foreseeable future, broader cloud adoption and the proliferation of IoT devices mean the enterprise boundaries are expanding beyond the traditional on-premises data center. This requires a new approach to both network and security architecture.

SD-WAN solutions are not created equally. The first generation of SD-WAN only addressed the transition to cost-effective public WAN options.

Private WAN connections are often expensive. For example, the first generation of SD-WAN technology witnessed many enterprises switch MPLS and P2P lease line connections to broadband or DSL. These connection types are referred to as public WAN connections, which are cost-effective. The first generation of SD-WAN would simply create overlay tunnels ensuring the public WAN connection was secure. This very quickly saved organizations money when compared to the cost of private WAN offering a similar service.
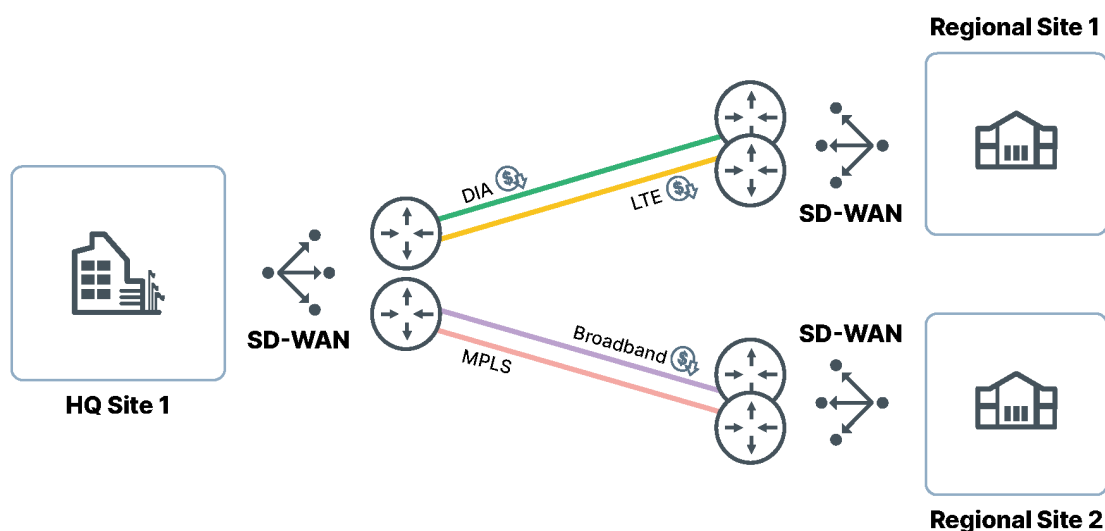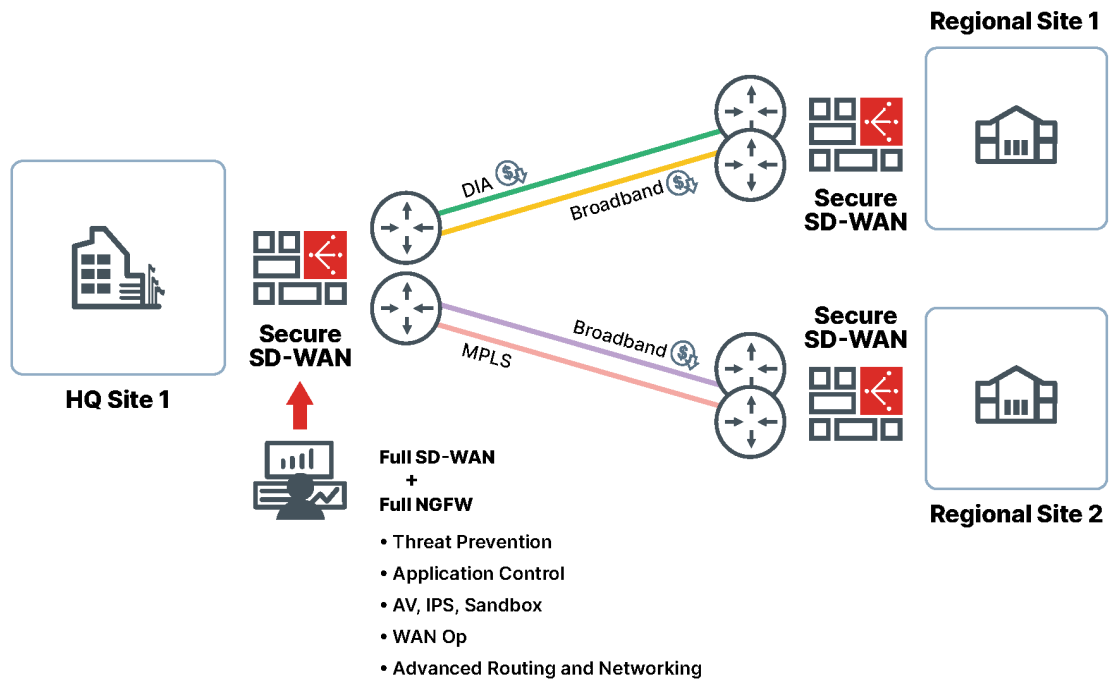

However, aside from cost savings, there was not much more to be gained from the first SD-WAN solutions. There was no technology convergence. In contrast, additional technology was being introduced. The SD-WAN solution of choice also required being managed and maintained. For the modern enterprise, such additional management overheads can quickly offset any savings derived from updating the WAN.

The next generation of SD-WAN witnessed vendors address this technology sprawl issue by combining SD-WAN solutions with other technologies, such as Next-Generation Firewalls. This approach, in theory, is great for enterprises. You save on WAN costs and combine some existing technologies.

The challenges surface when you unpack and validate these combined or embedded technologies. Historically, incorporating technologies that were not designed and built to coexist together does not work well. This is especially true when solutions are under the duress of complex workloads such as SSL and multi-threat inspection workloads.

Ideally, a single solution should be sought to manage SD-WAN and edge security. Such solutions should not be comprised of *stitched*, *bolted*, or *embedded* offerings from technology acquisitions and mergers. Solutions that combine multivendor technology should be avoided or at least thoroughly tested. A single-vendor secure SD-WAN solution should be sought.

HQ Site 1

Secure SD-WAN

DIA 💲

Broadband 💲

Broadband 💲

MPLS

Secure SD-WAN — Regional Site 1

Secure SD-WAN — Regional Site 2

Full SD-WAN + Full NGFW

- Threat Prevention
- Application Control
- AV, IPS, Sandbox
- WAN Op
- Advanced Routing and Networking

Introducing secure SD-branch

After simplifying WAN connectivity and security, there's often an opportunity to streamline enterprise technology stacks—specifically, wireless and LAN. Distributed enterprises usually manage these solutions on an ad hoc basis, introducing operational and commercial complexity.

Businesses are increasingly looking to replace their isolated WAN and LAN infrastructures in favor of a consolidated networking solution that delivers deeper integration and simplified operations at branch office locations.

An effective SD-Branch managed service should consolidate WAN and LAN capabilities to simplify remote office infrastructure and optimize operations without introducing new risks.

A fundamental starting point for SD-Branch is the delivery of SD-WAN as-a-Service. When selecting the right SD-Branch solution, service providers have multiple options. They need to weigh them carefully—factors such as orchestration, management, TCO, and security impact ARPU potential over time.
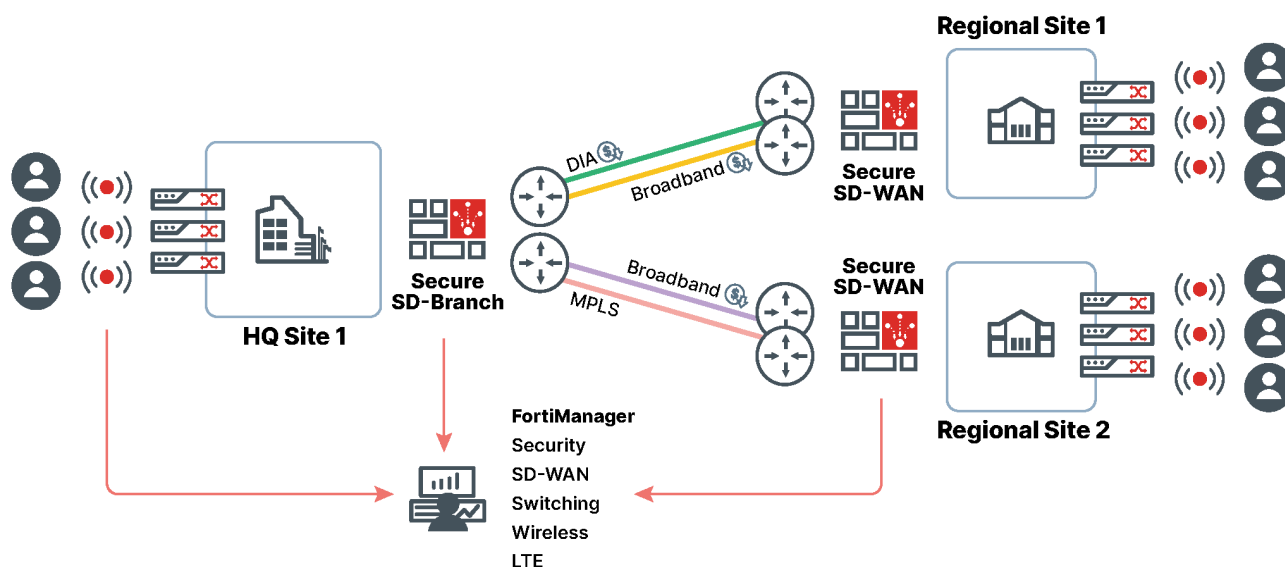
Why Fortinet

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.

One WAN Edge Powered by **One** OS, **One** Management
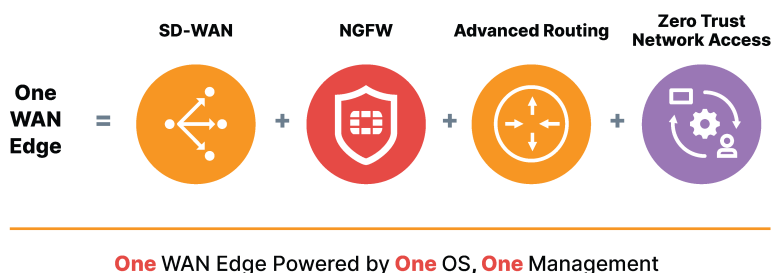
Secure SD-Branch follows the Fortinet security-driven networking philosophy, powered by FortiLink, which integrates wired and wireless services into the security infrastructure through FortiOS. Fortinet SD-Branch enables customers to converge their security, WAN, and LAN, extending the benefits of the Fortinet Security Fabric to their distributed branches.

Fortinet is among 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 MEF 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.

This certification demonstrates the ability of Fortinet Secure SD-WAN to comply with the highest industry standards required by service providers to deliver SD-WAN services. The Fortinet MEF SD-WAN product certification Test Report is available to learn more about test environments and testing methodology.

How can MSSP and SPs help

Customers are looking to industry experts to see where and how SD-WAN can simplify operations, improve uptime, and lower costs. For a distributed enterprise, this often starts with the current WAN supplier, the service provider.

As the existing providers of WAN connectivity, MSSPs and service providers are favorably positioned in helping clients understand and leverage SD-WAN technology.

While SD-WAN promises WAN cost-effective advantages over MPLS and private circuits, from our observations and findings, distributed enterprises are unlikely to drop private circuits entirely in favor of cost-effective public WAN. This further supports the notion that WAN suppliers should act quickly in building and offering Secure SD-WAN options to their customers, creating a single umbrella for WAN-related buying decisions.

The ability to bridge security and networking on the same platform is a significant advantage for service providers, enabling a broad, single-provider solution that increases average revenue per user. Choosing the underlying technology for managed SD-WAN and SD-Branch services is crucial. It is a critical factor determining the service's scope, addressable markets, potential revenue, and size of margins.

Beyond pure managed SD-WAN and SD-Branch services, service providers should consider the solution's comprehensive and consolidated functionality. The best options will provide the broadest range of extended value-added services (VAS) options for customers—including built-in security.

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.

Solutions and technologies

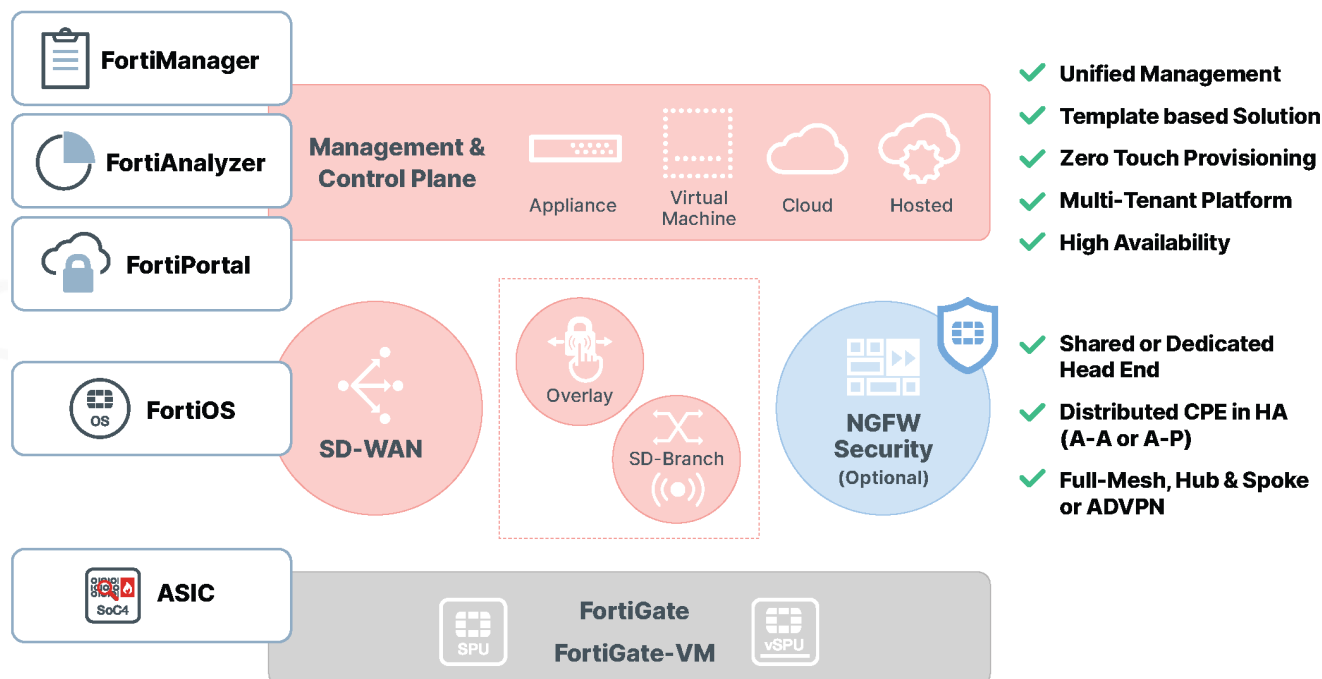Fortinet Secure SD-WAN consists of several components:

- 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

Another essential component is FortiDeploy, which allows zero touch deployment and provisioning of the entire SD-WAN and SD-Branch solution.

The following image shows the components of the Fortinet Secure SD-WAN and SD-Branch solution. Fortinet is the only vendor that comprehensively covers the entire solution, with a consistent security posture for WAN security and access layer security.

- Unified Management
- Template based Solution
- Zero Touch Provisioning
- Multi-Tenant Platform
- High Availability
- Shared or Dedicated Head End
- Distributed CPE in HA (A-A or A-P)
- Full-Mesh, Hub & Spoke or ADVPN

The design principles and use cases described in this document are based on the following software version for each component:

- FortiGate, FortiManager, FortiAnalyzer, FortiAP, and FortiSwitch version 6.4
- FortiPortal version 6.0

This section includes the following topics:

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:

- The intelligence to decide the best path for a specific application
- The ability to build the most efficient overlay network 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 (VDOM) technology is a testament to this statement, enabling a single, secure gateway instance to be sliced into potentially hundreds of individual gateways.

VDOM technology enables us to create a separation of duty between infrastructure and general management and the separation of access policies for all tenants. This is of paramount importance for managed service providers (MSPs), pivoting away from a single instance per customer model (which is still

possible with FortiGate). VDOM technology enables MSPs to benefit from the unique hardware acceleration exclusive to Fortinet FortiGates.

For more details on the FortiGate SD-WAN capabilities, see Technical background on page 25.

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

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.
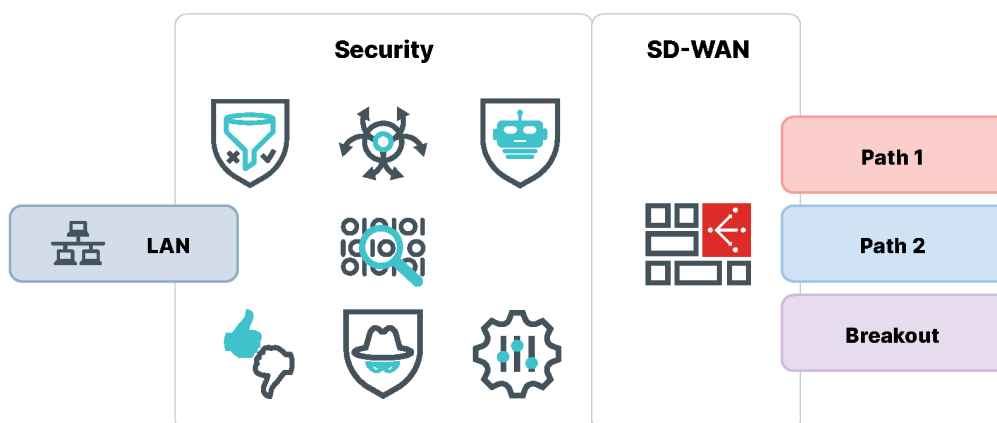
Security

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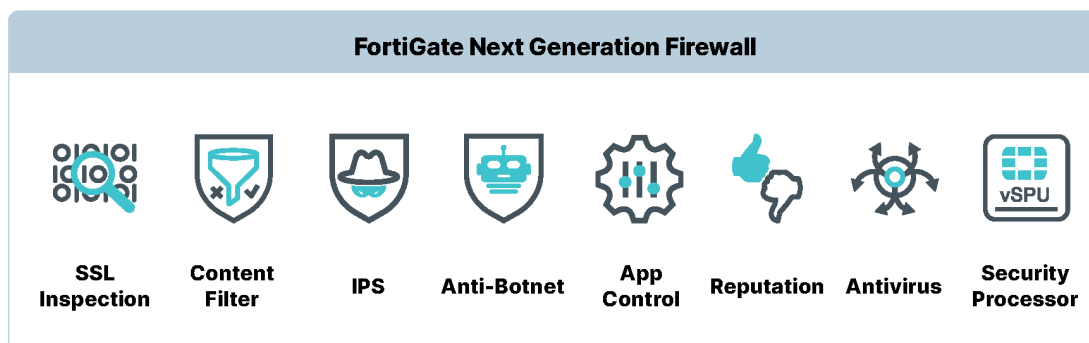
All the security features available in the FortiGate can be leveraged when SD-WAN is implemented.

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 can be leveraged as a physical appliance in the capacity of an on-premises high-performance manager. FortiManager can also be leveraged as a virtual machine in public or private cloud implementations. This enables the ability to scale when needed. Lastly, FortiManager can be leveraged as a SaaS offering service directly from Fortinet—ideal for those looking for an OpEx-enabled implementation.

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

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

FortiManager provides MSSP admins with 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 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 FortiGates 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 for on-premises high-performance deployments
- Virtual machine for public and private cloud environments for flexible and scalable deployments
- Fortinet SaaS for a complete capital expenditure turnkey solution

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 16
- Administrative domains on page 17
- Automatic security and SD-WAN reports on page 17
- Security automation on page 17

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/.

FortiPortal

FortiPortal is a comprehensive end-user self-service portal designed for enterprises, education institutions, and governments—specifically optimized for service providers. FortiPortal enables MSSPs to assign common firewall configurations and monitoring tasks to users in different geographies.

FortiPortal enables service providers to delegate the configuration and analytics to end-customers, business units, and departments in a multitenant environment, allowing them to monitor client bandwidth usage and monetize through automation.

This easy to deploy, turnkey portal delivers customer and device views that streamline adding clients and devices. Customers can quickly see usage, traffic, and other valuable summaries with an audit view and a customizable dashboard containing intuitive widgets. At the same time, self-branding enables a customer to retain their brand on the UI.

FortiPortal assists service providers in delivering the quickest time-to-market services by avoiding the need to develop costly portals.

FortiPortal provides end-users with an easy-to-use self-service customer portal, giving them access to secure SD-WAN capabilities, such as SD-WAN monitoring and templates, policies and objects, analytical dashboards, views and reports, WiFi, audit, and additional resources, such as documentation and links.

FortiPortal is delivered as an on-premises virtual machine hosted in the service provider datacenter. Optionally, it can be installed as a management extension within FortiManager.

In this guide, we will focus on the monitoring functionality provided by FortiPortal and do not give any details or recommendations on the configuration part.

The key features of FortiPortal are:

- **Dashboard**: FortiPortal's customer dashboard provides helpful visualizations to give users a first glimpse and overall picture of their network traffic and security posture, including filterable and intuitive widgets for Top Countries, Top Threats, Top Sources, Top Destinations, Top Applications, and Policy Hits as well as graphics to show Admin Logins, System Events, and Resource Usage. Users can drill down on the widgets to explore and investigate further.
- **Device Manager**: use the Device Manager tab to configure VPN IPsec, routers, and secure software-defined wide-area network (SD-WAN). Authentication servers support LDAP with password reset support, RADIUS, TACACS, and local authentication. Users can also view and manage their DHCP servers.
- **SD-WAN**: FortiPortal makes it easy for enterprises and service providers to configure secure SD-WAN and monitor SD-WAN interfaces, with intuitive tables, widgets, and map views that provide diagnostic visibility of interfaces, jitter, latency, volume, bandwidth, sessions, and more, enabling organizations to modernize their traditional WAN networks to meet the growing needs of the digital evolution. The Device Manager module has been updated with an enhanced SD-WAN monitoring map. A new drilldown history view improves users' capability to analyze overall bandwidth utilization, traffic trends, and application performance.
- **Views**: FortiPortal's enhanced View page provides comprehensive visibility into device and network configurations. It displays security and event information by application, source, or destination, with filters and controls that allow easy navigation. Customers gain meaningful insights into network activity and threats in the Monitor view with new secure SD-WAN widgets, VPN with History timelines for SSL and Dialup, and site-to-site IPsec. Top Threats and Top Sources views provide intuitive interfaces that can be filtered and searched with *text search*. Users can further research Log View with FortiGate Event and UTM data, Traffic, IPS, Antivirus, DNS, Application Control, Web Filter, and Event and Sandbox logs. Users can also filter by log type, site, date, and time, and choose from many columns and fields for a more thorough analysis.

- **Policy and Objects**: Policy and Objects views allow users to centrally manage and configure the devices managed by the FortiManager unit. This function includes the basic network settings to connect the device to the corporate network, antivirus definitions, intrusion protection signatures, access rules, and managing and updating firmware for the devices. FortiPortal's Policy and Objects views provide users with a transparent view and access to policies and objects for any configuration tasks that the service provider delegates and that have been defined on FortiManager.

- **WiFi**: use the WiFi tab for updating or deleting managed access points (APs), monitoring rogue access points, Fortinet access points (FAPs), and SSIDs, and updating or deleting access point profiles and SSIDs.

- **Audit**: the Audit tab displays a log of user activity on the administrative GUI by date, which can be searched by level, user, type, IP, or message description, which can be exported to a CSV file for further analyses.

- **Reports**: FortiPortal Reports provides administrators with the ability to easily create and assign reports to users. The Reports page displays a list of available FortiPortal or FortiAnalyzer reports. It allows users to select reports, set filters, run reports, and search reports based on the roles assigned by the service provider.

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.

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.

FORTIAP

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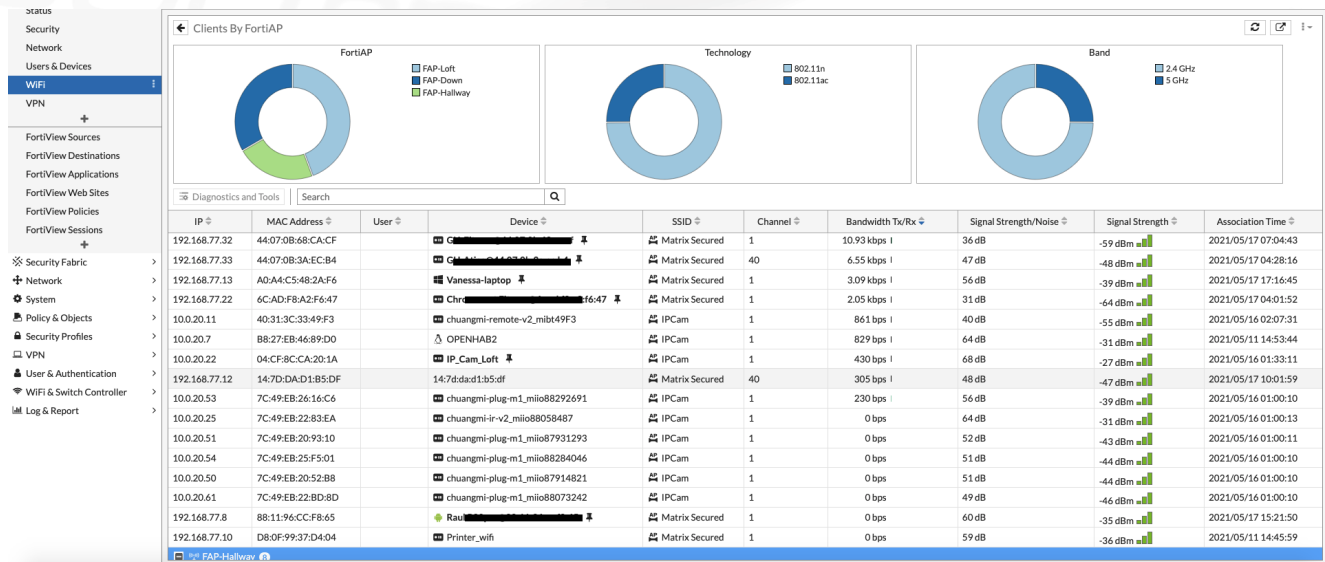
Configuring and managing access points from the same known dashboard as the security parameters also allows immediate visibility and troubleshooting advantages. One can very quickly understand:

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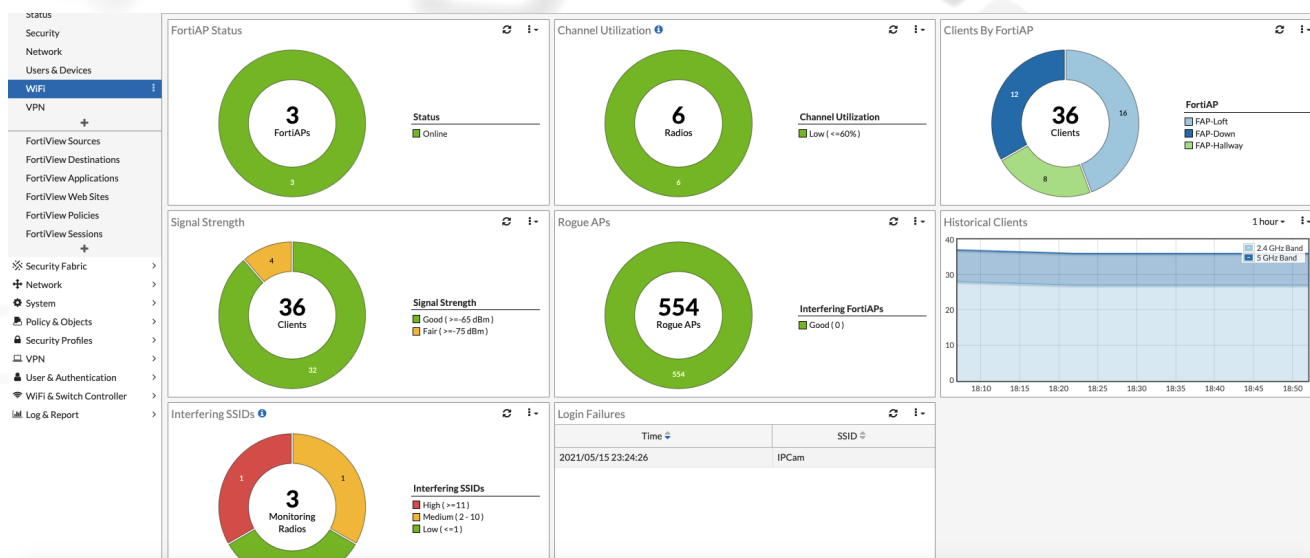
All of that allows for the easy operation of a live and evolving secure wired and wireless network for administrators and a trusted infrastructure for users to perform their daily job without worrying about the underlying connectivity.

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

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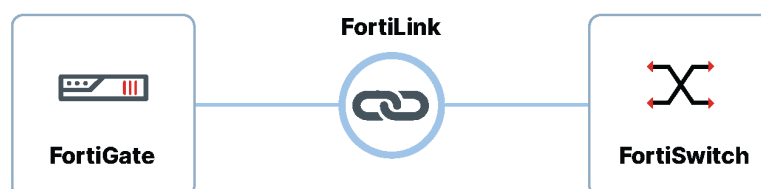
- **Single pane of glass**: From the FortiGate or FortiManager GUIs, you can access and configure the FortiSwitches managed by FortiLink, and obtain thorough information about the devices connected to them.

Security Fabric: Matrix Secured

Topology last updated 5 minute(s) and 38 second(s) ago.
Security Rating results that ran 1 hour(s) and 32 minute(s) ago were aggregated.

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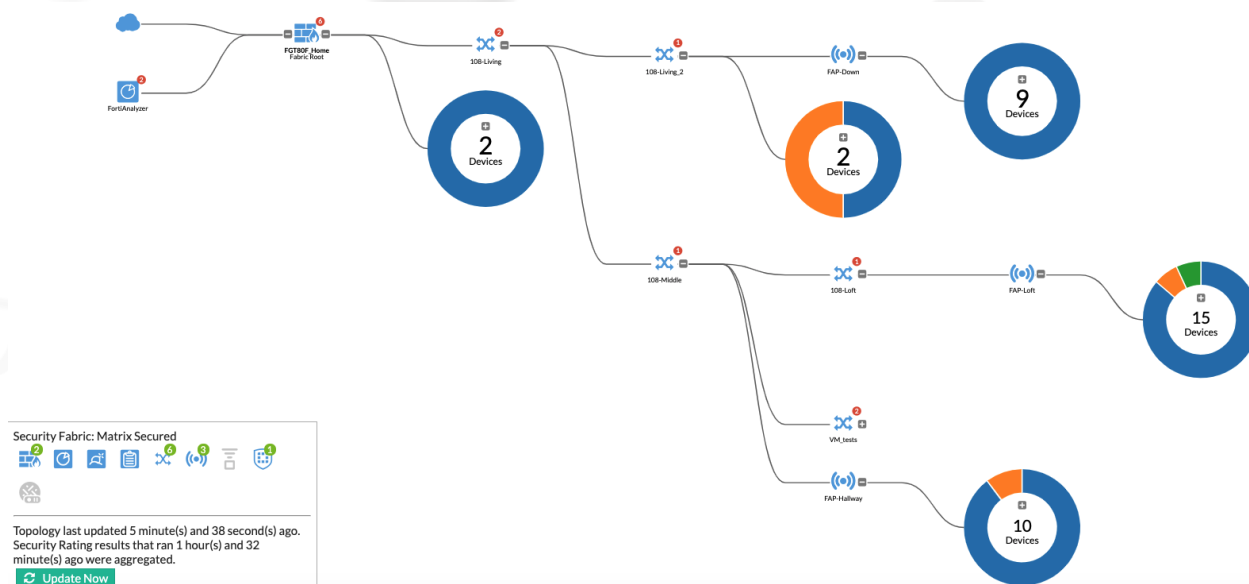
- **Security to the LAN**: security features associated with LAN environments—from basic VLAN segmentation to 802.1X authentication policies. VLANs become part of FortiGate's interfaces and can be easily integrated into Security Fabric policies.

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

- **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).

See also Important terms for FortiSwitch on page 24.

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

The technical background covers the following topics:

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*).

Edit Performance SLA

Name	Internet
IP Version	IPv4 ▾
Detect Protocol	Ping ▾
Health-Check Server	🖥 Internet ▾
Participants	All SD-WAN Members **Specify**

🔍

🖳 W_MPLS_H1 ✖
🖳 W_MPLS_H2 ✖
🖳 ul_inet ✖

3 Entries Selected

Enable Probe Packets **ON**

SLA

+ Create New   ✏ Edit   🗑 Delete   ⚙ Column Settings ▾   🔍

	ID	Latency Threshold (Milliseconds	Jitter Threshold (Milliseconds)	Packet Loss Threshold (%)
☐	1	200	-	-
☐	2	300	-	-

Link Status

Interval	500	Milliseconds
Failure Before Inactive	5	(max 3600)
Restore Link After	5	(max 3600)

Action When Inactive

Update Static Route	**ON**
Cascade Interfaces	**ON**

Advanced Options ❯

[OK]   [Cancel]

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 signature-based L7 application detection (Application Control Database), dynamic feeds (internet Service Database—ISDB), multiple User Identity providers, DSCP/ToS fields, Route Tags, and of course, all based on simple L3/L4 criteria!

- **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
☐	1	Corporate-Primary	🖥 CORP_LAN	🖥 CORP_LAN	DC#1	🌐 W_INET_H1 🌐 W_MPLS_H1
☐	2	Corporate-Secondary	🖥 CORP_LAN	🖥 CORP_LAN	DC#1	🌐 W_INET_H2 🌐 W_MPLS_H2
☐	3	Business-Critical-HighPriori ALL		☁ Salesforce 🧩 GoToMeeting	Internet#1	🌐 ul_inet 🌐 W_MPLS_H1 🌐 W_MPLS_H2
☐	4	Business-Critical-MedPriori ALL		Microsoft.Office.365 Microsoft.Office.365	Internet#2	🌐 ul_inet 🌐 W_MPLS_H1 🌐 W_MPLS_H2
☐	5	Non-Business-Critical	ALL	🖥 all	Latency	🌐 ul_inet
☐		sd-wan	ALL	ALL	Source IP Based	ALL

Edit SD-WAN Rule

Name	Corporate-Primary
IP Version	IPv4 ▾
Source	

Source

Source Address	🔍
	🖥 CORP_LAN IP/Netmask:10.0.0.0/255.0.0.0 ✖
	1 Entry Selected
Users	Click here to select
User Groups	Click here to select

Destination [Address | Internet Service]

Address	🔍
	🖥 CORP_LAN IP/Netmask:10.0.0.0/255.0.0.0 ✖
	1 Entry 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	🔍
	🌐 W_INET_H1 ✖
	🌐 W_MPLS_H1 ✖
	2 Entries Selected
Required SLA Target	🔍
	DC#1 ✖
	1 Entry Selected

Advanced Options ❯

[OK] [Cancel]

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**:

These pillars together define
and secure all available paths
to all possible destinations

Their configuration does not
change during regular
operations and/or functioning
of the network!

This is the intelligence that
steers the actual traffic!

Underlay — Choose the WAN links to use

Overlay — Choose the desired topology to interconnect your sites

Routing — Choose how to propagate routes between your sites

Security — Choose how to protect each of the available paths

SD-WAN — Choose the strategy to pick one of the available paths

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

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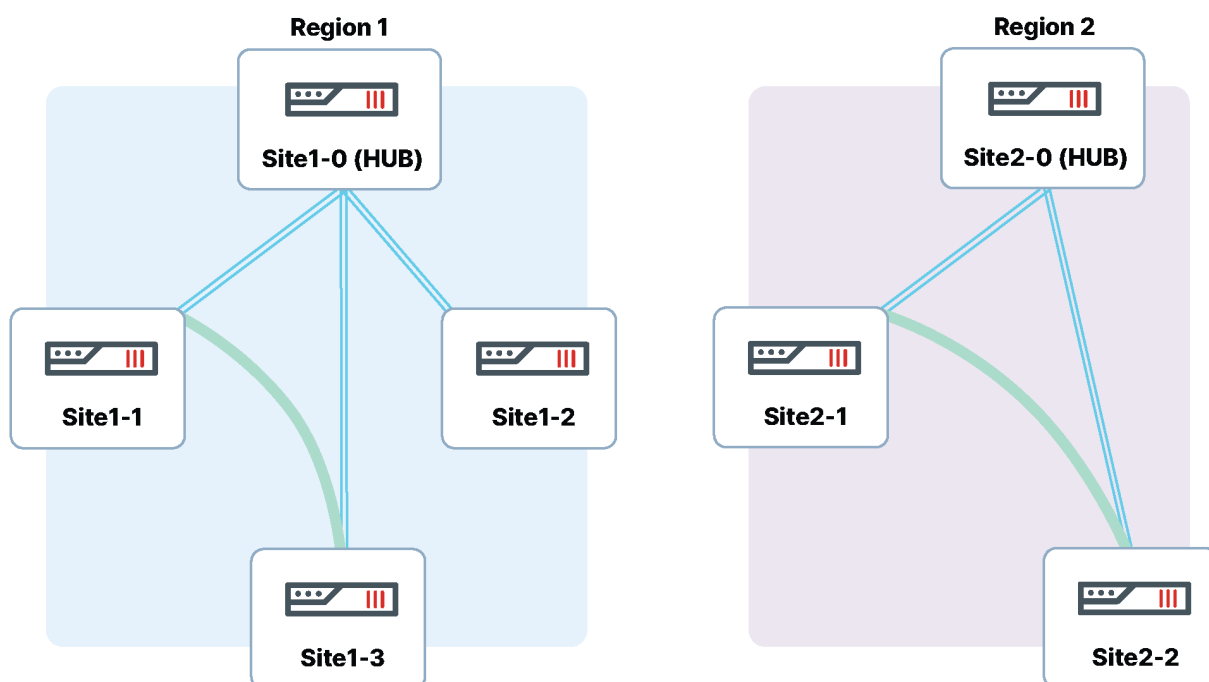
To conclude, although other overlay topologies can be used (such as a static hub-and-spoke or even a full-mesh), we recommend ADVPN as the most generic, dynamically adjustable topology for your overlays.

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 34.

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.
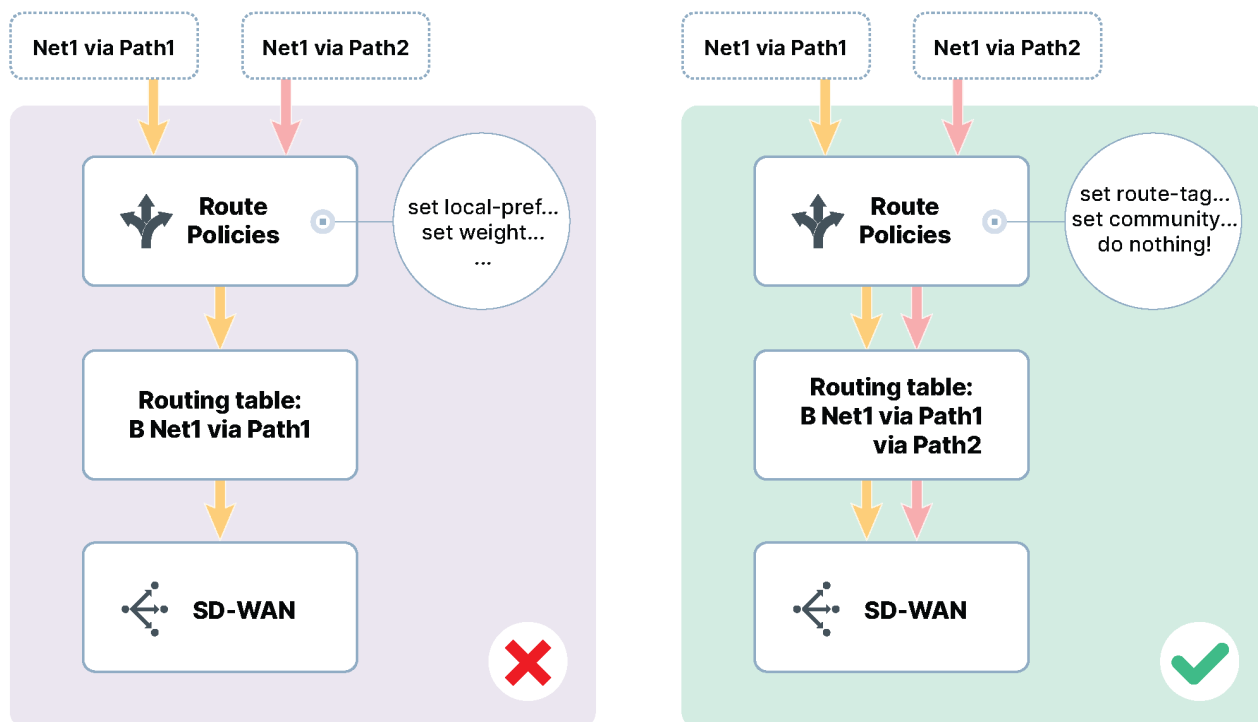
Security

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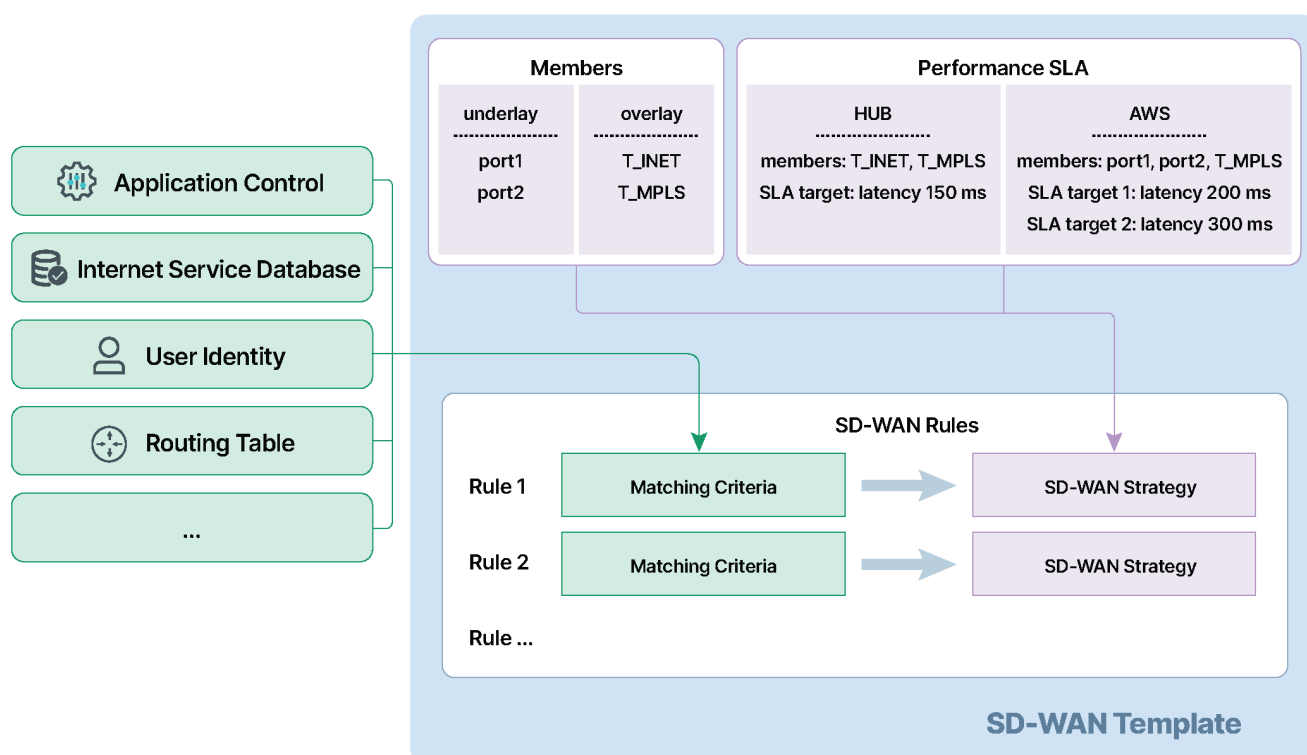
For example, you could have a single SD-WAN template for all your branch offices and another SD-WAN template for your central datacenters. This would allow you to apply changes quickly and consistently

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!

Design example - basic SD-WAN/ADVPN

This design is the most fundamental building block of our solution. The more advanced multi-hub and multi-regional examples that we cover later will essentially be extensions of basic SD-WAN/ADVPN.

Optional - LAN (Site 0)

Site 0 (HUB)

INET
MPLS

T_INET   T_MPLS

T_INET   T_MPLS   T_INET   T_MPLS   T_INET

Site 1   **Site 2**   **Site 3**

LAN (Site 1)   **LAN (Site 2)**   **LAN (Site 3)**
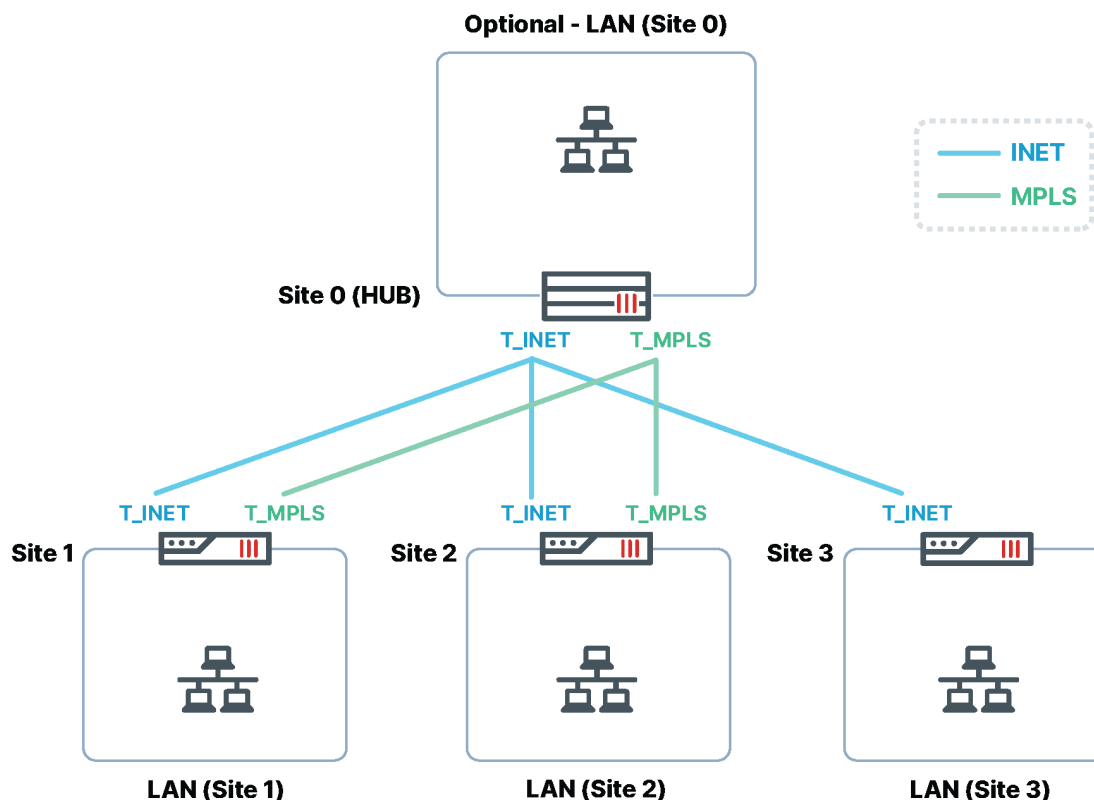
The sites are interconnected by IPsec overlays, forming hub-and-spoke topology. Two primary flavors can be distinguished:

- In the *Enterprise* flavor, the hub is located at the customer's central office or a datacenter. The spokes (edges) are distributed across all remote sites (branch offices, retail stores, homeworkers, and so on). Most traffic is either spoke-to-hub (sites accessing workloads hosted in the datacenter, remote internet access through centralized breakout, cloud on-ramp services, and so on) or direct internet access from spokes (workloads hosted on public clouds—*aaS, non-business-critical internet browsing, and so on). Occasional spoke-to-spoke communication is flowing through direct ADVPN shortcuts.

- The *Managed Service* flavor is suitable for managed security service providers (MSSPs). The hub is located at MSSP premises, while the spokes (edges) are distributed across all end-customer sites, acting as CPEs. All types of end-customer sites are spokes in this flavor, including branch offices, datacenters, and so on. As a result, nearly all the traffic is either spoke-to-spoke (flowing through ADVPN shortcuts) or direct internet access from the spokes. There are often no workloads behind the hubs in MSSP premises, which reduces traffic flow through the hubs to a minimum, leaving them mainly with a control plane function. For this reason, the hubs are often called *SD-WAN Gateways* in this flavor. Optionally, they can be shared between multiple endcustomers (also called *tenants*), using either virtual routing and forwarding (VRF) instances or virtual domains (VDOMs).

Regardless of the chosen flavor, the rest of the design remains the same.

The hub acts as a dial-up IPsec server for the spokes, having a separate dial-up IPsec endpoint terminate on each underlay interface. Each endpoint defines a point-to-multipoint overlay. Every spoke will typically connect to all the overlays to have multiple alternative paths. However, it can also happen that some of the spokes do not have all the underlay transports available. Hence, they will be able to connect only to a subset of the overlays.

The spokes establish separate IBGP sessions to the hub over each overlay. BGP Neighbor Group feature is used on the hub for this peering. Each spoke then advertises its local site prefix(es) over each of the IBGP

sessions. The hub acts as a BGP Route Reflector (RR), readvertising the prefixes to all other spokes. Additionally, the hub advertises its prefixes (such as DC LAN in the Enterprise flavor). At the end of this process, all the sites exchange their routes over all available overlays.
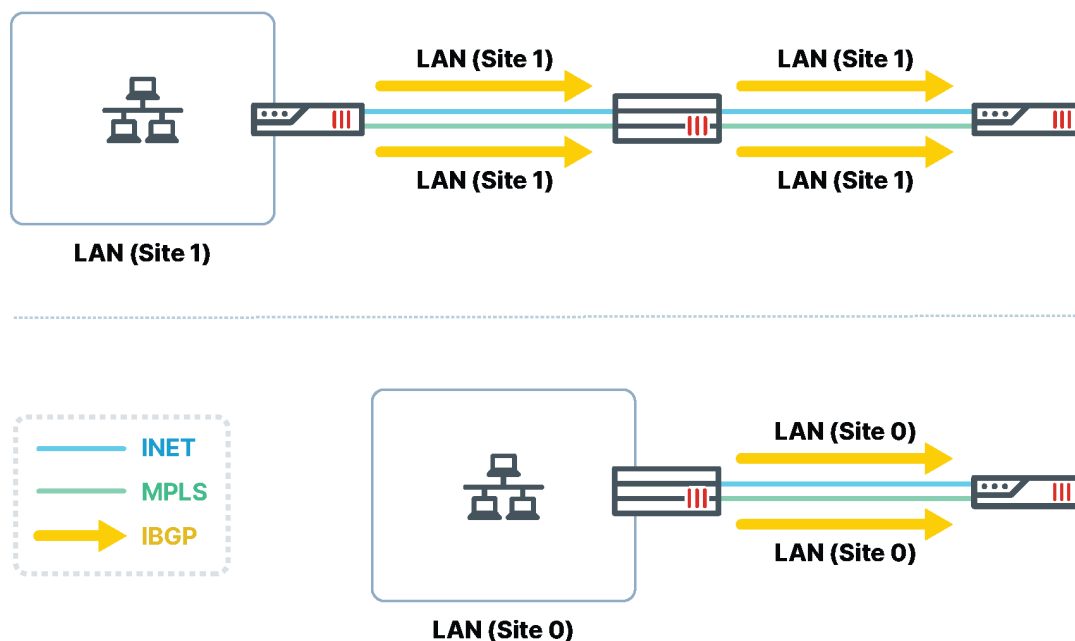
Once all the routes have been distributed across all the sites, the application traffic flow can be controlled by SD-WAN rules according to the design principles described in the previous chapter.

Let us list different types of traffic flows present in this topology. See Traffic flows on page 37.

Traffic flows

Corporate traffic (site-to-site) never leaves organization boundaries. It can be spoke-to-spoke, spoke-to-hub (when there are workloads behind the hub), or—rarely—hub-to-spoke. This traffic will usually travel via one of the available overlays protected by the IPsec suite. In spoke-to-spoke traffic, the IPsec tunnel will be dynamically built by ADVPN to provide direct communication. In this case, only the first few packets will flow through the hub until an ADVPN shortcut is built.
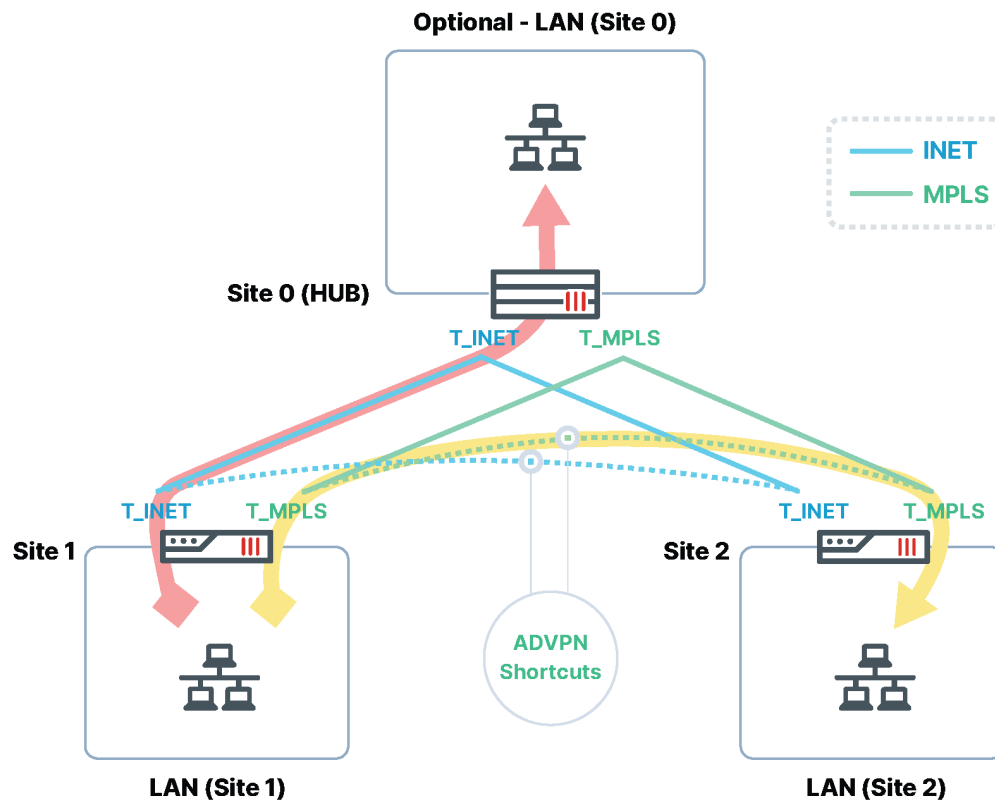
Direct internet access (DIA) is also known as local breakout, and this traffic leaves organization boundaries directly from the site edge. Usually, it will follow a default route, and there will be no dynamic routing peering involved. Depending on what is available on a particular site, it can use one or more underlay links connecting to the public internet.
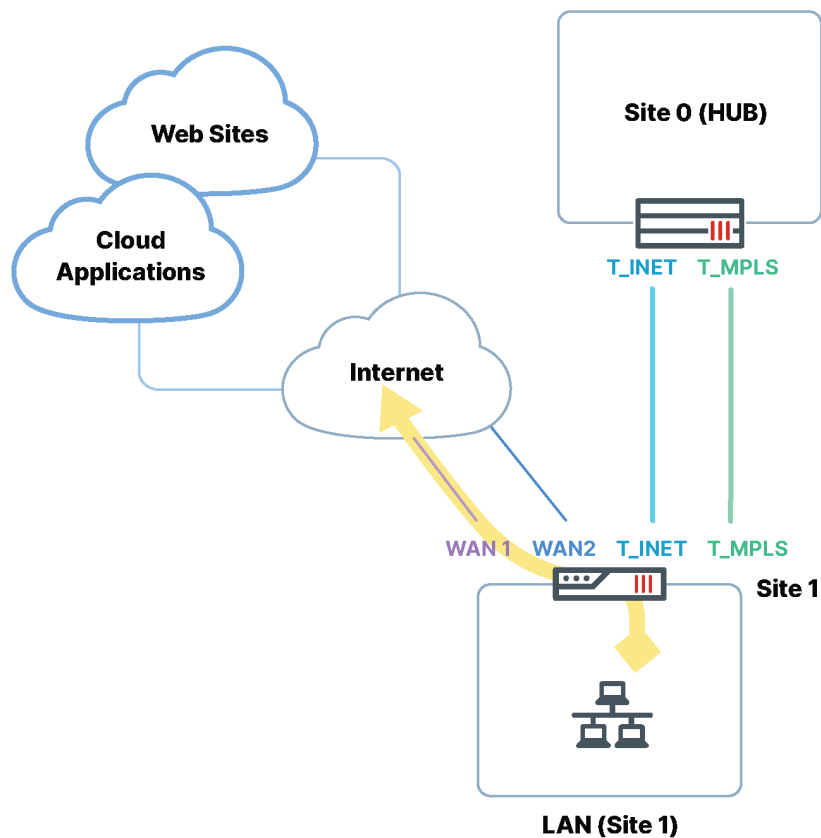
Remote internet access (RIA) is also known as remote breakout, and it means that the traffic must be backhauled through the hub (located in a datacenter or on MSSP premises). In that case, the traffic will flow through one of the available overlays toward the hub, similar to the corporate spoke-to-hub traffic. Customer policies can mandate RIA, but they can also be used as a better alternative to DIA for business-critical traffic, for example, when the quality of the local internet access degrades. At the same time, premium MPLS services can guarantee better application performance.
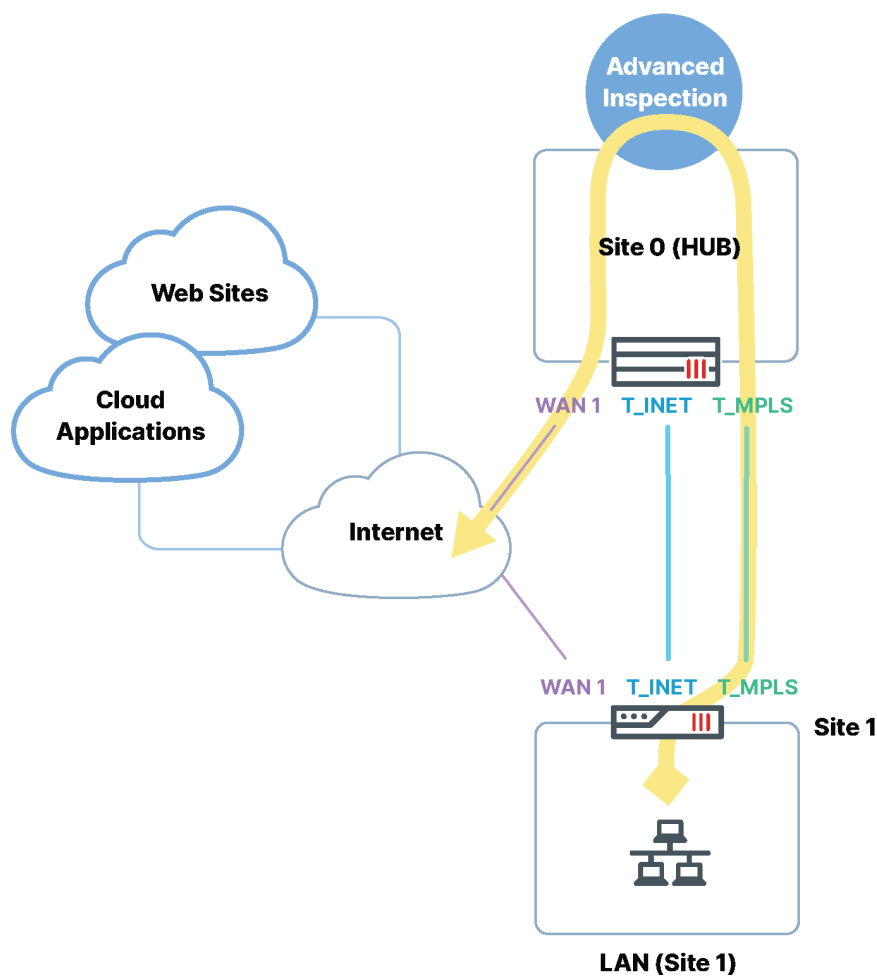
Cloud on-ramp provides optimized access to the workloads running in the cloud. Rather than accessing the cloud services through the public internet, an overlay can be established to the closest cloud POP in the area. Cloud providers offer optimized access using their built-in gateways, but deploying FortiGate VMs in the cloud is also possible. Cloud on-ramp connections can be from a central location (behind the hub) or directly from the spoke sites. From the SD-WAN's perspective, this type of traffic flow is similar to a corporate spoke-to-hub and RIA since it will leave the site edge through one of the available overlays.
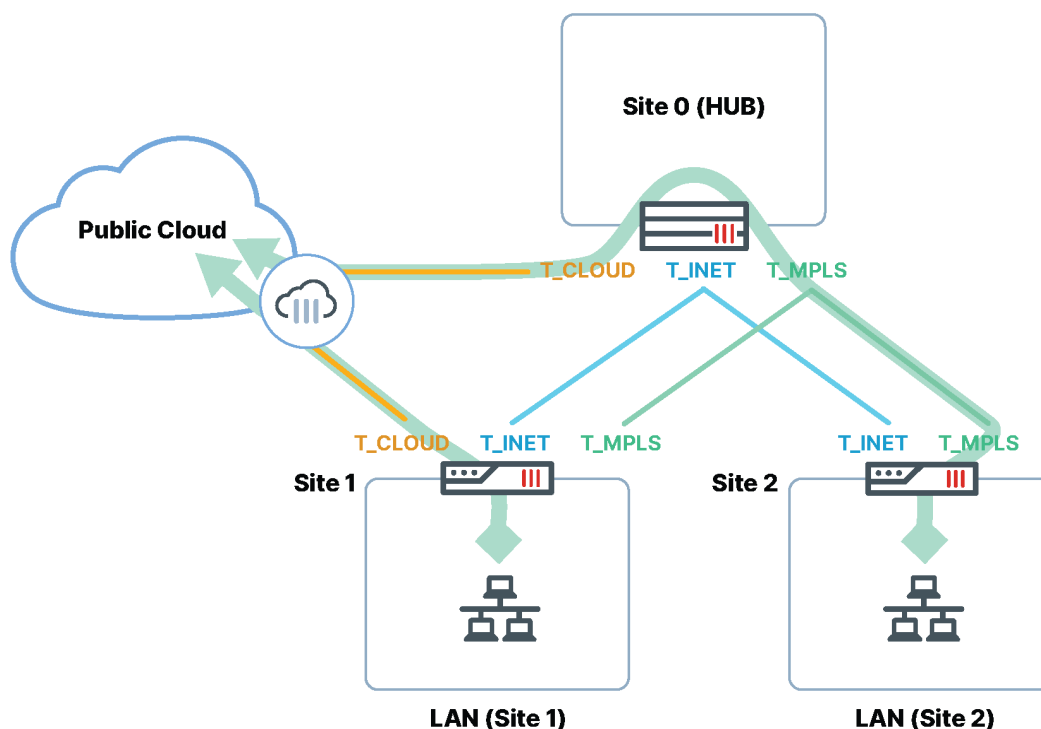
Let us now take a closer look at a typical SD-WAN configuration for this topology. See SD-WAN configuration on page 41.

SD-WAN configuration

As we have already mentioned in the previous section, SD-WAN decisions shall be made by the originating sites, typically the spokes. The hub usually acts either as a destination site or as a transit site, and, as such, it does not require SD-WAN configuration.

- **What interfaces on the spokes should become SD-WAN members?** All the overlays do. But we recommend adding the underlays too. Suppose only direct internet access is needed, and all the sites have just a single internet connection. In that case, you *could* handle internet access with the conventional routing, and that way, keep it outside of the SD-WAN solution. However, adding the underlays to the SD-WAN *bundle* allows you to measure their health (using one or more performance SLAs). Hence, even though you do not make steering decisions based on these measurements, you can still benefit from the improved visibility! Furthermore, if your business requirements change—for example, you add a new cloud on-ramp service, or you decide to backhaul at least some of the internet traffic through a central location—it will be just a matter of updating your SD-WAN rules to apply this change.

- **What probe destinations should you choose for performance SLAs?** Because it largely depends on your traffic patterns for internet traffic, it is difficult to provide a general recommendation. For example, for public cloud traffic, it is generally recommended to probe the respective cloud provider. For general internet browsing, probing a public DNS could be a good option. As for the corporate (site-to-site) traffic, we recommend configuring a loopback interface on the hub, which the spokes will probe over all available overlays. This will allow them to compare the quality of different available transports without maintaining a dedicated health-check server behind the hub.

ADVPN shortcuts are automatically monitored, using settings and SLA targets defined in the performance SLA of the respective *parent* overlay. No manual configuration and no external health-check servers are required for this functionality. ADVPN Shortcut Monitoring provides more accurate measurement for the spoke-to-spoke traffic.

Technical highlights

- Adding or removing spokes does not require altering either the hub's configuration (thanks to dial-up IPsec and BGP Neighbor Group features) or the configuration of the other spokes (thanks to the nature of hub-and-spoke topologies).
- 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. The hub can automatically allocate tunnel IPs to the spokes using the IKE Mode Config feature to simplify provisioning.
- 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
- For the correct operation of ADVPN, it is required to preserve all sites' prefixes unchanged, including their original BGP NH 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 hub must reflect the original routes between the spokes without altering them.
- 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.

Design example - dual-hub

Customers willing to provide geographic redundancy to their SD-WAN solution will typically extend the previous design to include a secondary hub.

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.
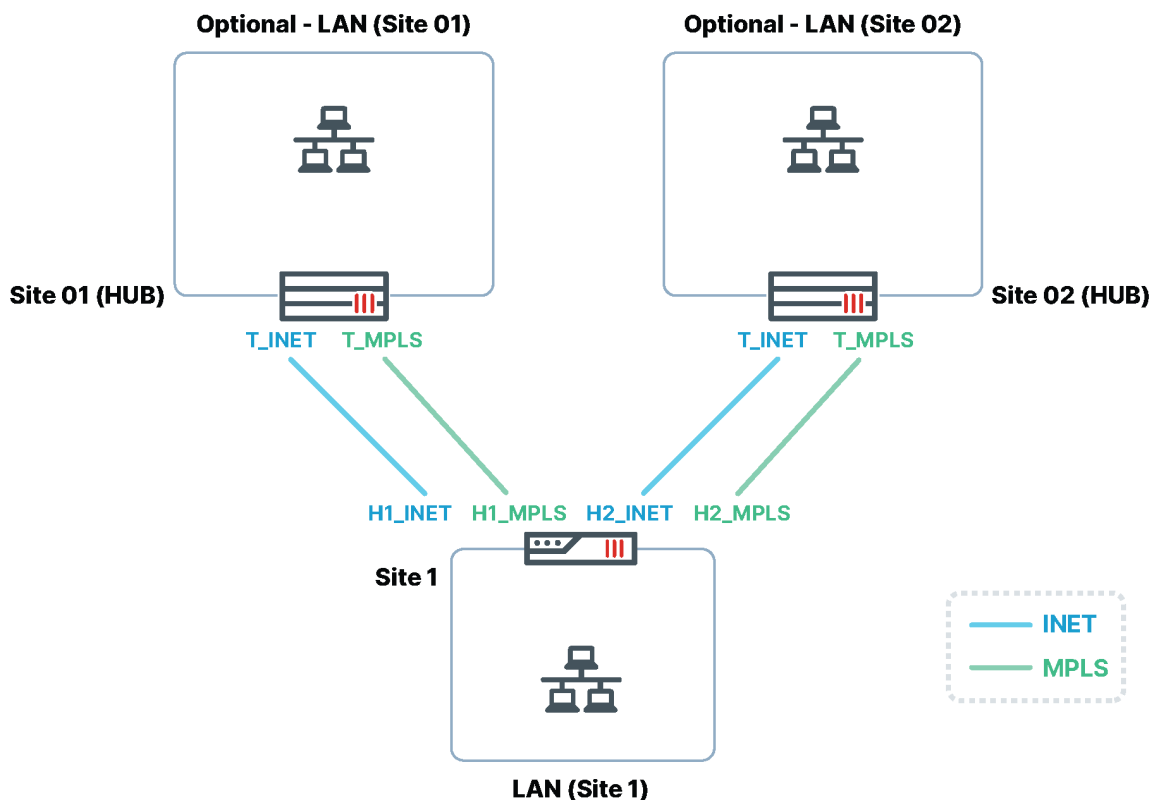
After connecting to all the overlays, the spokes also establish separate IBGP sessions to both hubs through each of the overlays. The spokes then advertise their local site prefix(es) to both hubs, and each of the hubs acts as an independent BGP route reflector. As a result of this route exchange, all the sites learn each other's prefixes by all available overlays through both hubs. Following the described principles, all these routes should be installed into the routing tables (ECMP).
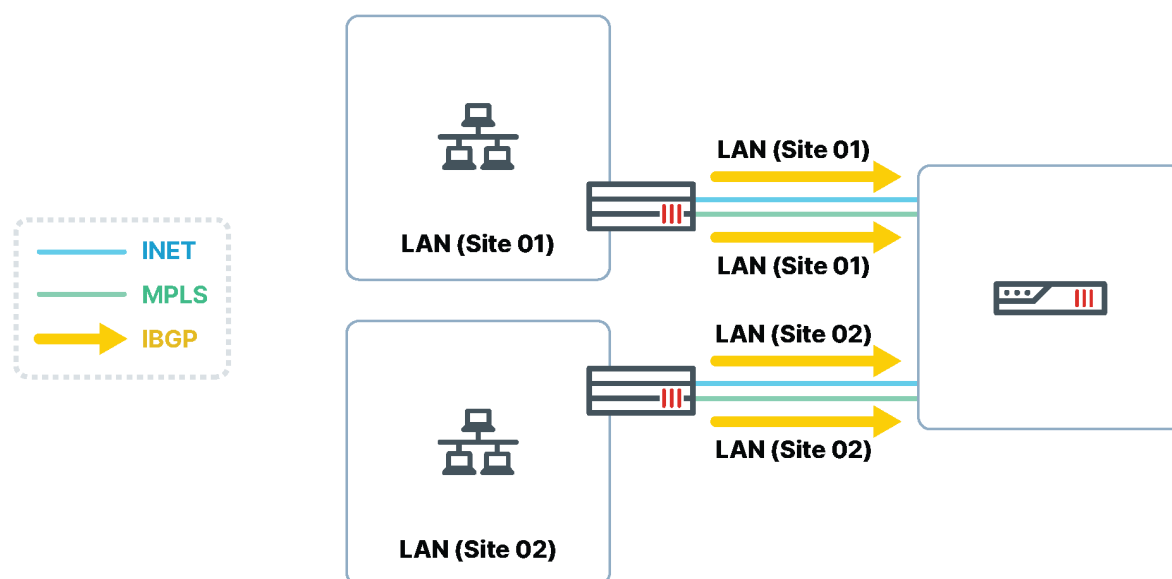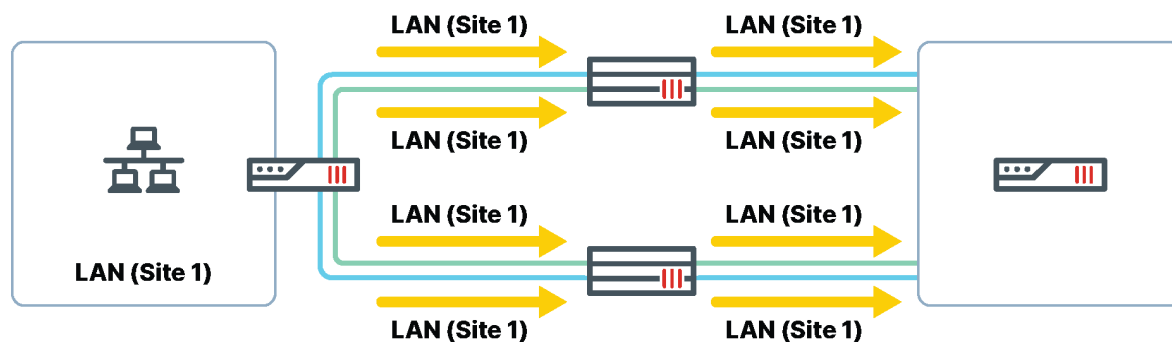
Once all the routes have been learned, it's time to see how the SD-WAN rules define the exact redundancy model.

SD-WAN configuration

The SD-WAN configuration principles remain very similar to the basic design. Let us recap with the necessary adjustments:

- **What interfaces on the spokes should become SD-WAN members?** As before, we recommend adding all underlays and overlays to the SD-WAN *bundle*. This time it will include the overlays belonging to both hubs.
- **What probe destinations should you choose for performance SLAs?** To simplify the configuration, you can configure an identical loopback address on both hubs. This will allow you to use a single performance SLA definition, while effectively, the probes will be sent to different hubs, depending on which overlay they are sent over. As shown below, you can easily refer to both hubs in the same SD-WAN rule if needed.

But the following question requires more attention: **How to implement the desired redundancy model with SD-WAN rules?**

We recommend one of the following approaches:

Architecture for MSSPs

- **Active-passive hub**: In this approach, the secondary hub will be used only when the primary hub is out of service (down or unreachable). This means that even if all the overlays toward the primary hub are out of SLA, the secondary hub will not be used. To implement this approach, you need to define two SD-WAN rules: The first rule includes the primary hub overlays, and the second rule includes the secondary hub overlays. Note that, due to the operation of SD-WAN rules, the second rule will be matched only when the first rule cannot be used to forward the traffic. This is only true when the primary hub is entirely out of service. This is shown in the following figure:

- **Active-passive underlay**: All the overlays are listed in a single SD-WAN rule in this approach. The actual path selection depends on the configured rule strategy. For example, the following figure demonstrates an active-passive underlay:
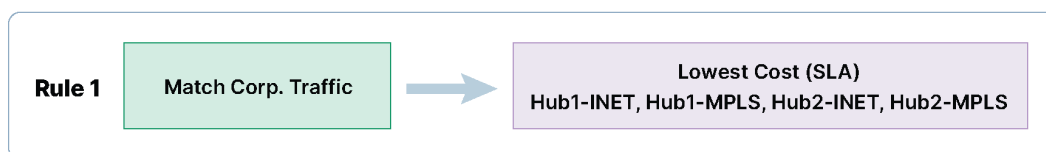
As can be seen, the rule is configured with the *Lowest Cost (SLA)* strategy, preferring the members in the order they are listed. Hence, the overlays of the Primary hub will be preferred. However, the difference with the active-passive hub approach is that if both primary hub overlays are out of SLA (although still in service), the secondary hub will be selected.

Technical highlights

What happens to spoke-to-hub traffic when the destination is only reachable through one of the hubs? This can easily happen if the two hubs are located in two different datacenters, with different workloads hosted behind them. To answer this question, recall that the SD-WAN uses routing information as one of its inputs. By default, SD-WAN members can only be selected when they have a route to the destination. Hence, even if the SD-WAN rule prefers the primary hub, it will be skipped for destinations only reachable through the secondary hub.

Design example - multi-regional design

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 (either dual-hub or single-hub).
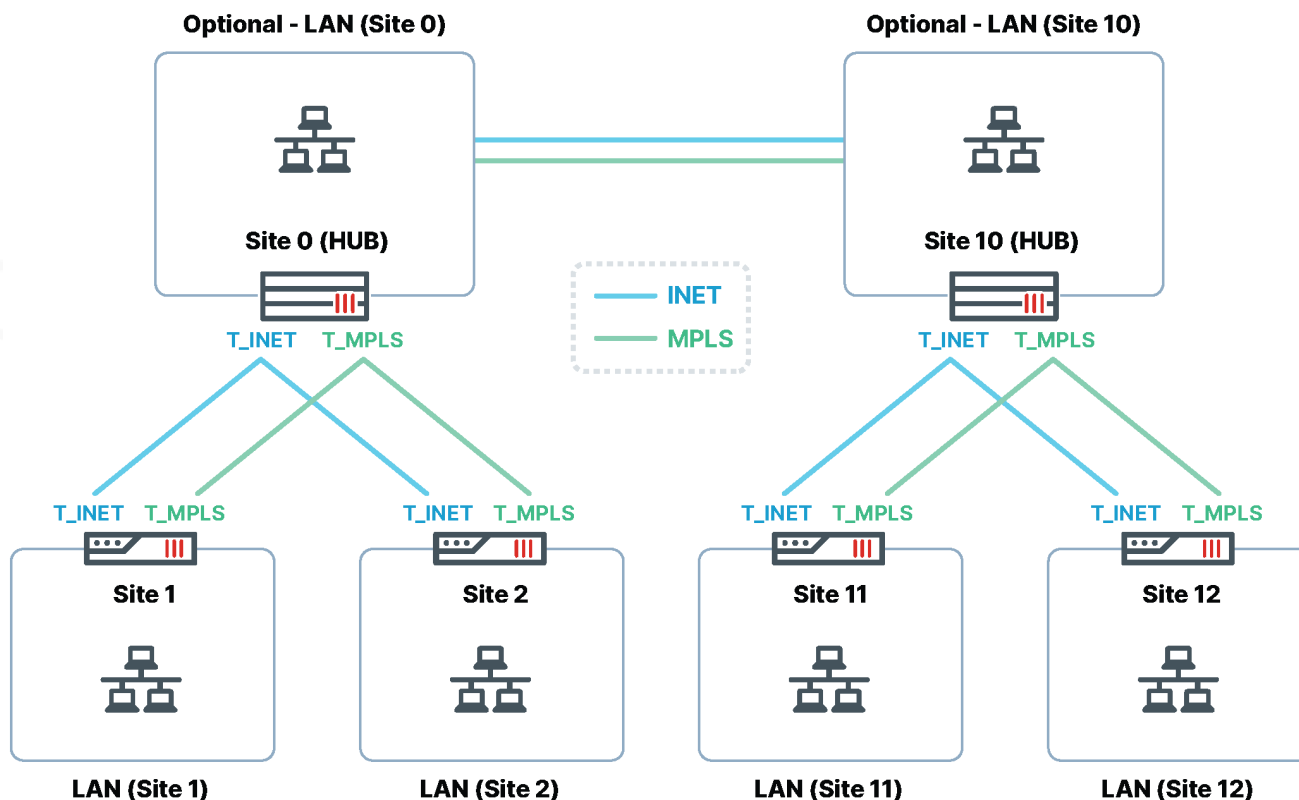
To achieve that, we define regional hubs in each geographical area, and let all other sites in the area connect only to these regional hubs. This includes both IPsec overlays and BGP sessions. As already discussed, this would be enough to provide connectivity within each region. In addition, all the regional hubs are interconnected between them, forming a full-mesh topology with BGP sessions exchanging the routes between all the regions.

The recommended approach is to use **EBGP** between the regional hubs so that each hub advertises a *summary route* of all regional prefixes to all remote regional hubs. Those will, in turn, advertise default routes to their spokes. A spoke willing to communicate to a remote region will always send traffic to its local, regional hub, which will use the correct summary route to forward the traffic to the remote regional hub. Note that ADVPN will be used only for spoke-to-spoke traffic within each region, while the traffic across the regions will always flow through the regional hubs.
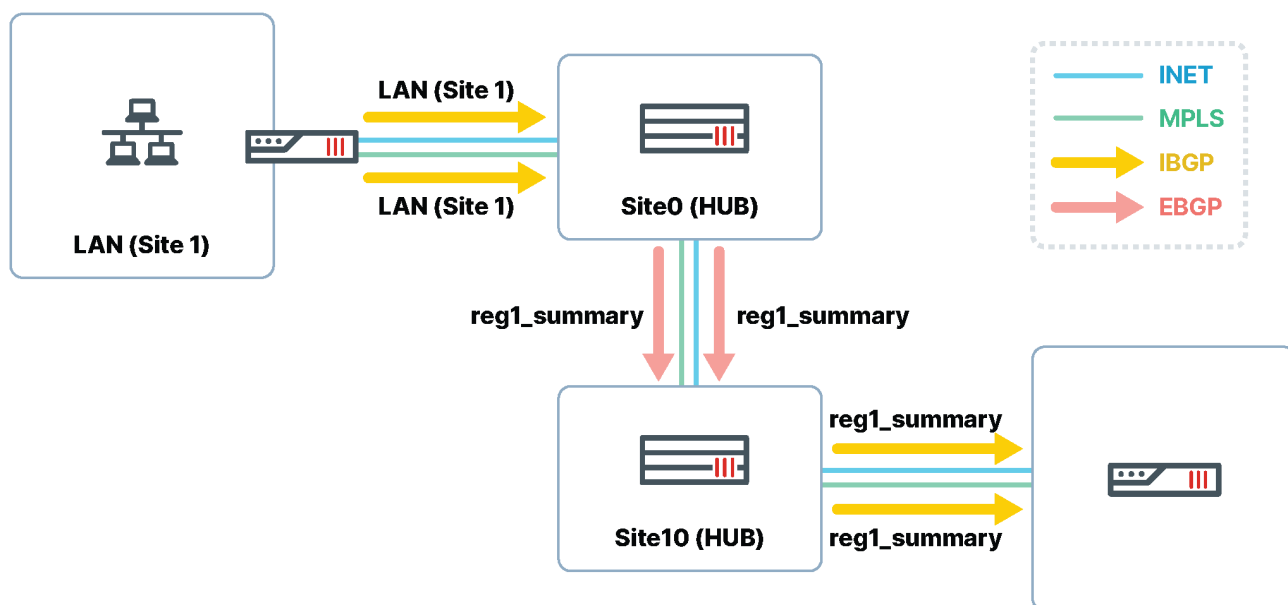
SD-WAN configuration

The SD-WAN configuration of the spokes in a multi-regional solution remains identical to the one described in the single-regional 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 flavors: whether cross-regional ADVPN is used or not, the SD-WAN configuration on a spoke remains the same.

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

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.
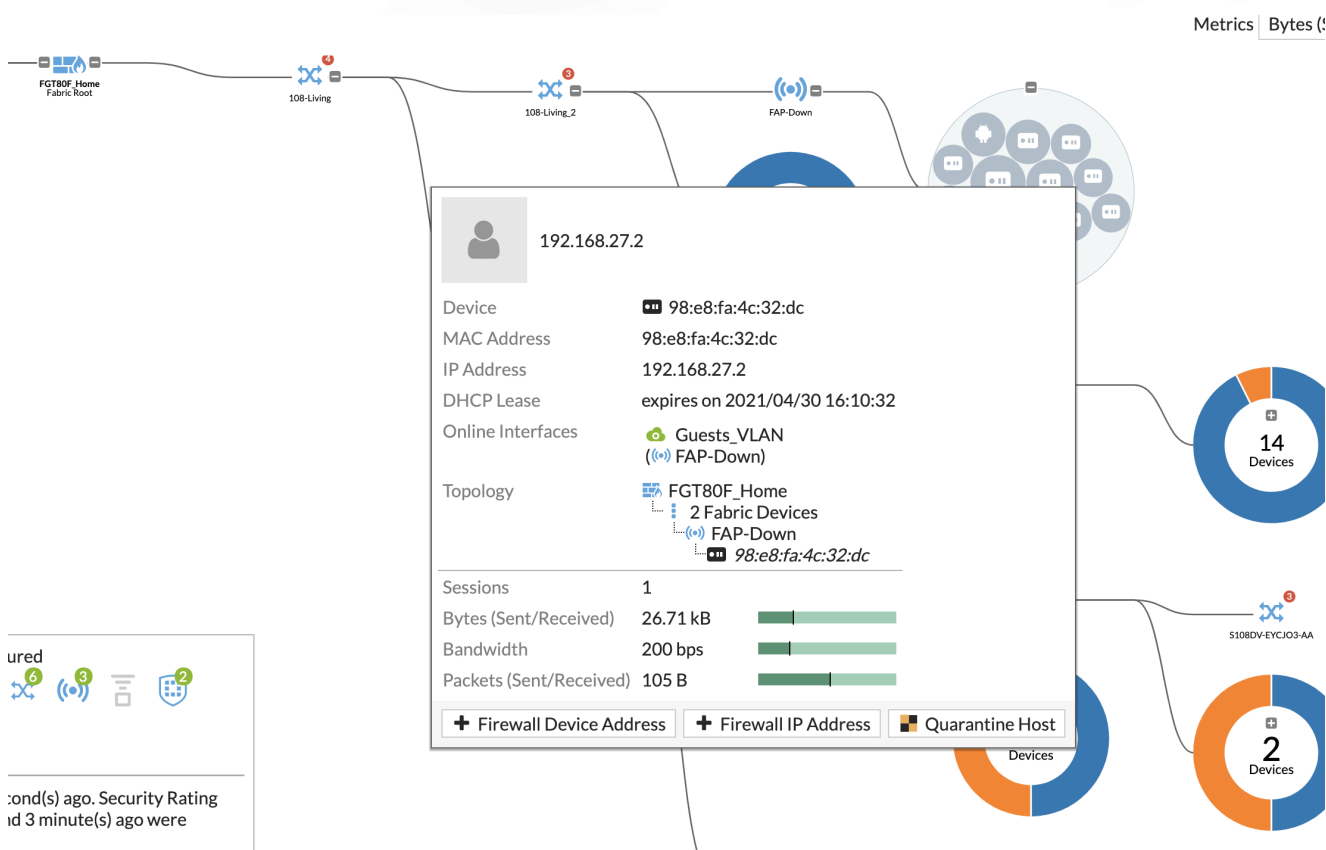
Attack surface reduction with network segmentation

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.

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.

The 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.

Menu	Edit NAC Policy
Dashboard >	
Security Fabric >	Name Android_devices
Network >	Status [⬆ Enabled] [⬇ Disabled]
System >	FortiSwitches [All] Specify
Policy & Objects >	Description [_____] 0/63
Security Profiles >	
VPN >	**If device matches all of the following patterns:**
User & Authentication >	Category [Device] User EMS Tag
WiFi & Switch Controller ⌄	MAC address ◯
Managed FortiAPs	Hardware vendor ◯
WiFi Clients	Device family ◯
WiFi Maps	Type ● [Phone]
SSIDs	Operating system ● [Android]
FortiAP Profiles	User ◯
WIDS Profiles	
FortiLink Interface	**Then:**
Managed FortiSwitch	Select an action that will be performed to the matched device.
FortiSwitch VLANs	◯ **Assign VLAN**
FortiSwitch Ports	Assign a specific VLAN to a device matching above patterns.
FortiSwitch NAC Policies ☆	● **Apply Port Specific Settings**
FortiSwitch Security Policies	Apply a LLDP Profile, QoS Policy, 802.1X Policy, or VLAN Policy.
Log & Report >	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.

Monitoring and reporting

Managed service providers (MSPs) require the ability to offer both fully managed and co-managed services, and therefore needs to be conscious of deploying a solution that can be both multi-tenanted, and allow per-tenant deliverables to be customized.

As detailed earlier in this document, FortiAnalyzer and FortiManager are the tools used by MSPs to provide complete configuration, reporting, and monitoring of the Secure SD-WAN and SD-Branch environment. However, these platforms are typically deployed within the MSP infrastructure, and best practice dictates that an end-customer should not have direct access to multi-tenant management platforms.

However, before looking at the customer presentation layer, we need to bring it back to the FortiAnalyzer and FortiManager architecture, initially focusing on sizing.

This section includes the following topics:

- FortiAnalyzer on page 53
- FortiPortal for managed service providers on page 59

FortiAnalyzer

This section about FortiAnalyzer covers the following topics:

- ADOMs, sizing, log storage, scaling, and enforcement on page 53
- SD-WAN logging on page 55
- FortiAnalyzer HA recommendation on page 58

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 systemwide 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.

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 nonstandard 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 via 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 Develope 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.

4D Accelerator

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.
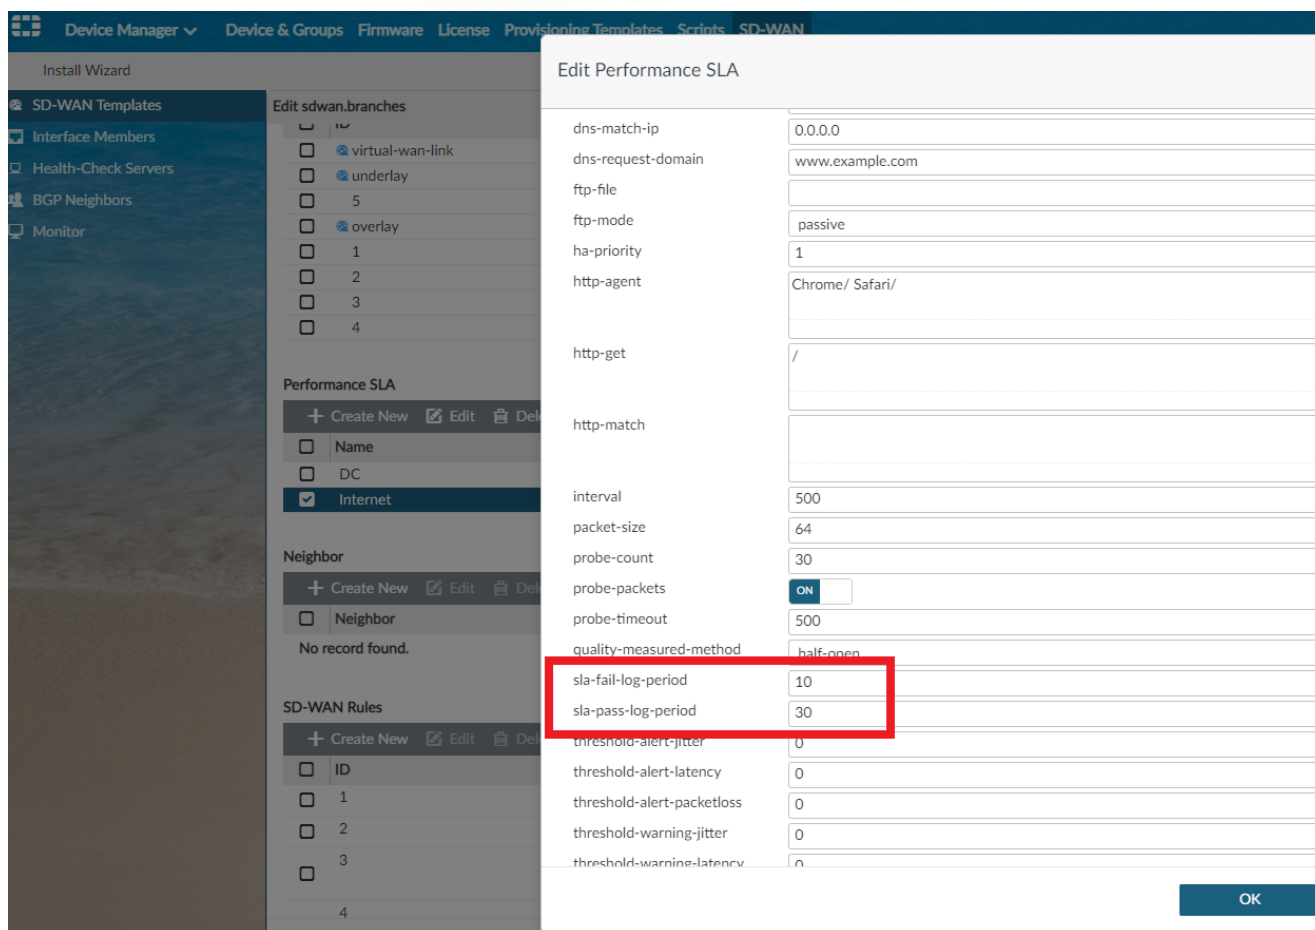
In the next example below, SD-WAN rule *Business_Critical-HighPriority* uses the SLA *Internet#1*, which has 200ms latency and 5ms jitter set as thresholds. This means that a probe (a ping, DNS, HTTP, or others) is being sent at a specified time period, every 500ms being the default, across SD-WAN member 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.
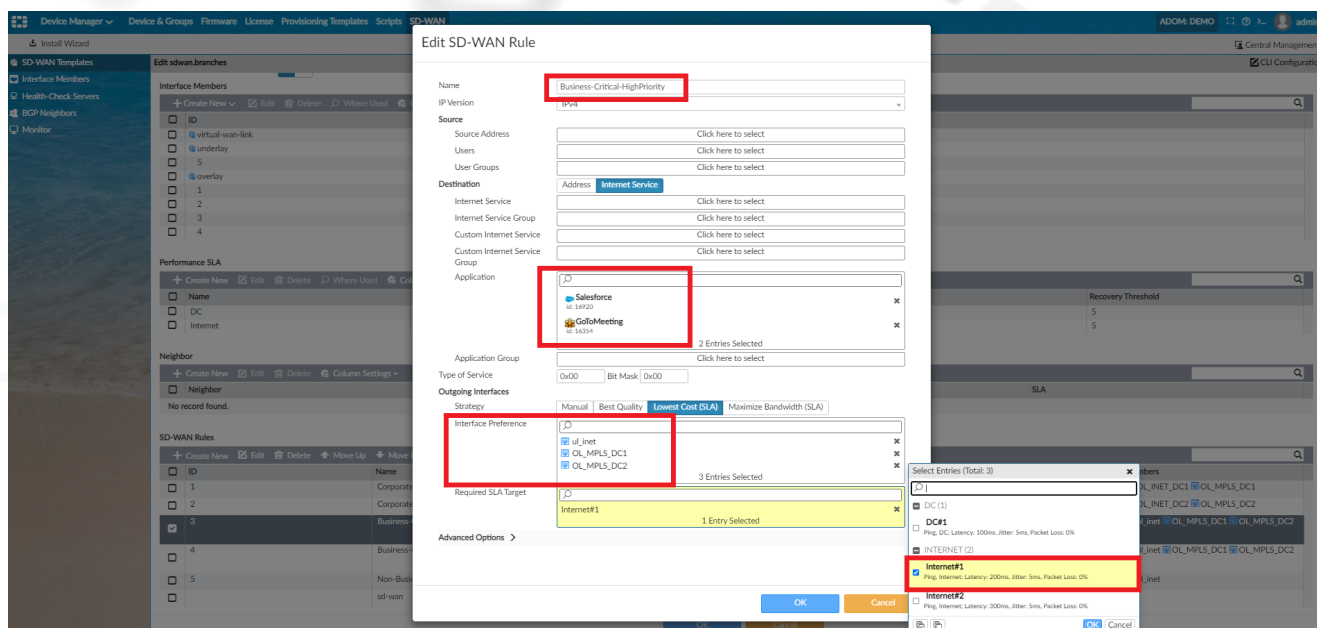
Let's look at how the various logs sent from FortiGate to FortiAnalyzer look from the CLI.

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.

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

- **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**: the 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 High and Critical Events**: existing alarms on a path, connection, or individual WAN links for their states (Information, Notice, and Warning)
- **SD-WAN Utilization by Application**: the share of bandwidth utilization by application for each WAN link
- **SD-WAN Rules Utilization**: SD-WAN rules traffic utilization by interface and application
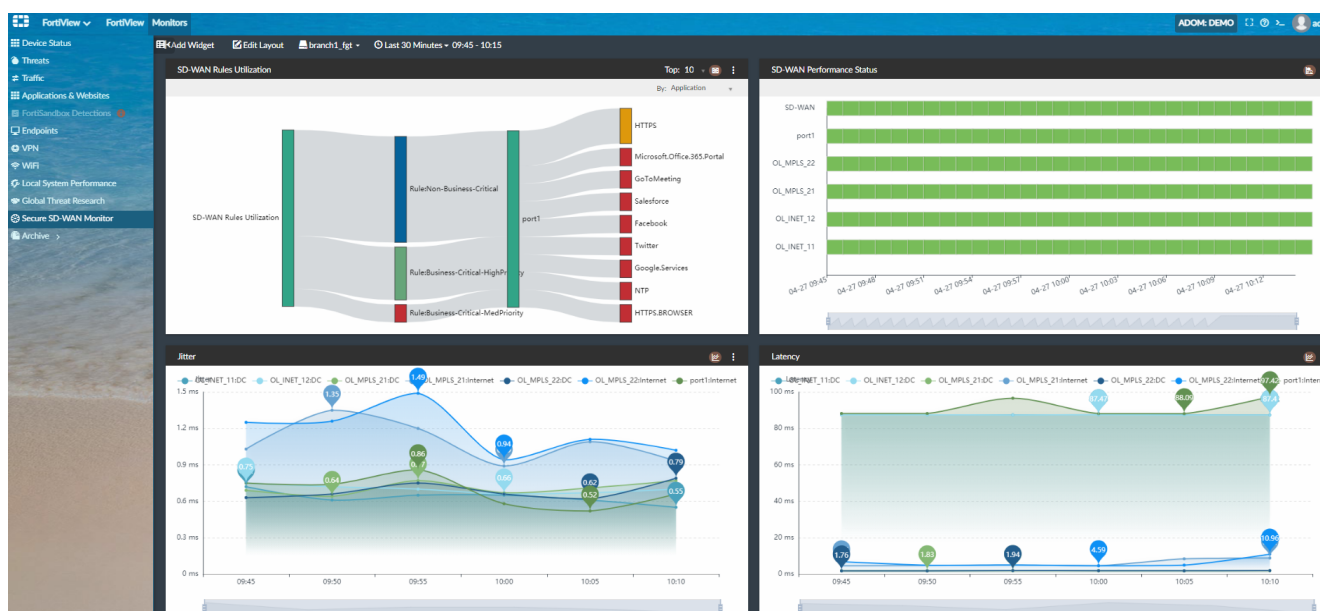
FortiAnalyzer provides a comprehensive SD-WAN reporting section. All the reports are fully customizable to meet both MSP and end-customer branding. The below image shows some of the reports included in FortiAnalyzer and one page of the SD-WAN report as an example.
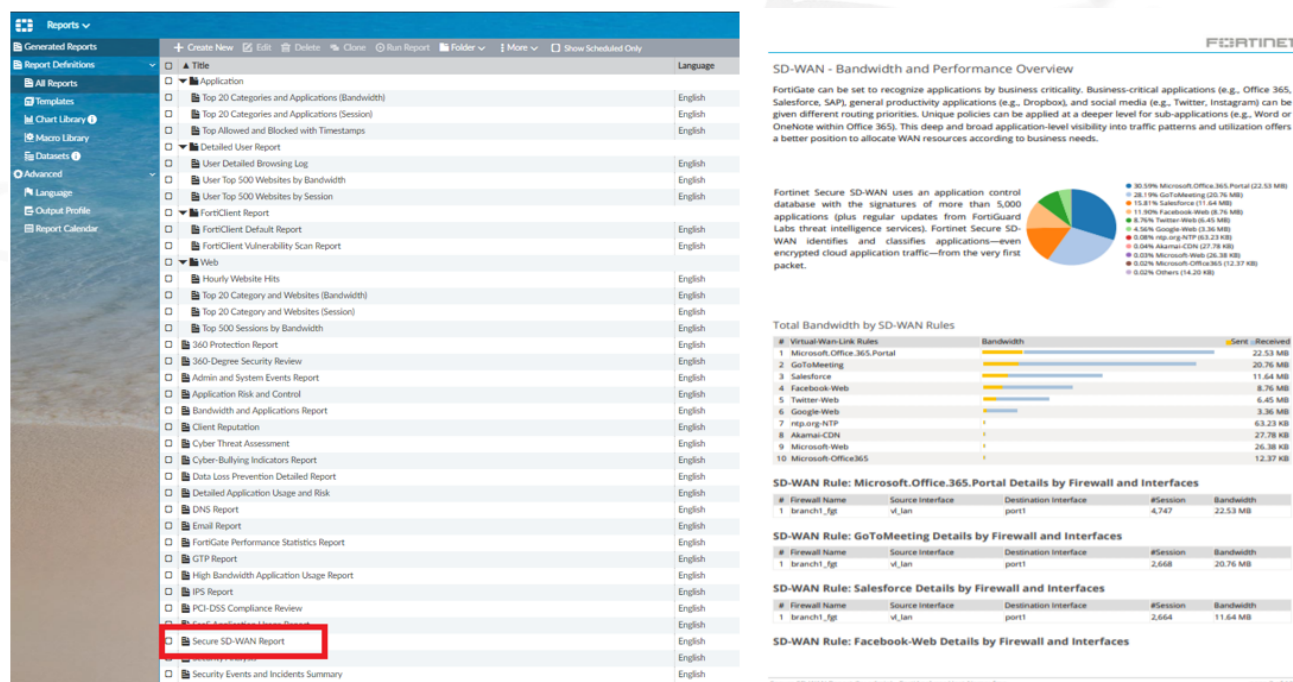
FortiAnalyzer HA recommendation

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.

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.

When introducing FortiPortal as the presentation layer of the solution, FortiPortal will use the VIP of the HA FortiAnalyzer to retrieve the logging information. The following example illustrates that cluster VIP 192.168.244.222 is used to manage FortiAnalyzer in FortiPortal.

FortiPortal for managed service providers

When an MSP offers an SD-WAN managed service, FortiManager and FortiAnalyzer are typically under complete control of the MSP. When FortiManager and FortiAnalyzer are part of a multitenant and shared platform, it is not recommended to provide administrative access to end-customers.

Therefore, the use of a customer portal in front of FortiManager and FortiAnalyzer is recommended. FortiPortal is the customer portal that leverages FortiManager and FortiAnalyzer APIs to expose information and configurations to the end-customer.

FortiPortal is the presentation layer for the end-customer and is multitenant by design. This allows the MSP to customize each customer portal, and include only access to configuration views, images, and reports relevant to the customer.

As mentioned earlier, FortiPortal requires connectivity to both FortiAnalyzer and FortiManager. The following image shows the FortiManager *admin* user account, which FortiPortal will use to connect. Notice that *JSON API Access* is set to *Read-Write*.
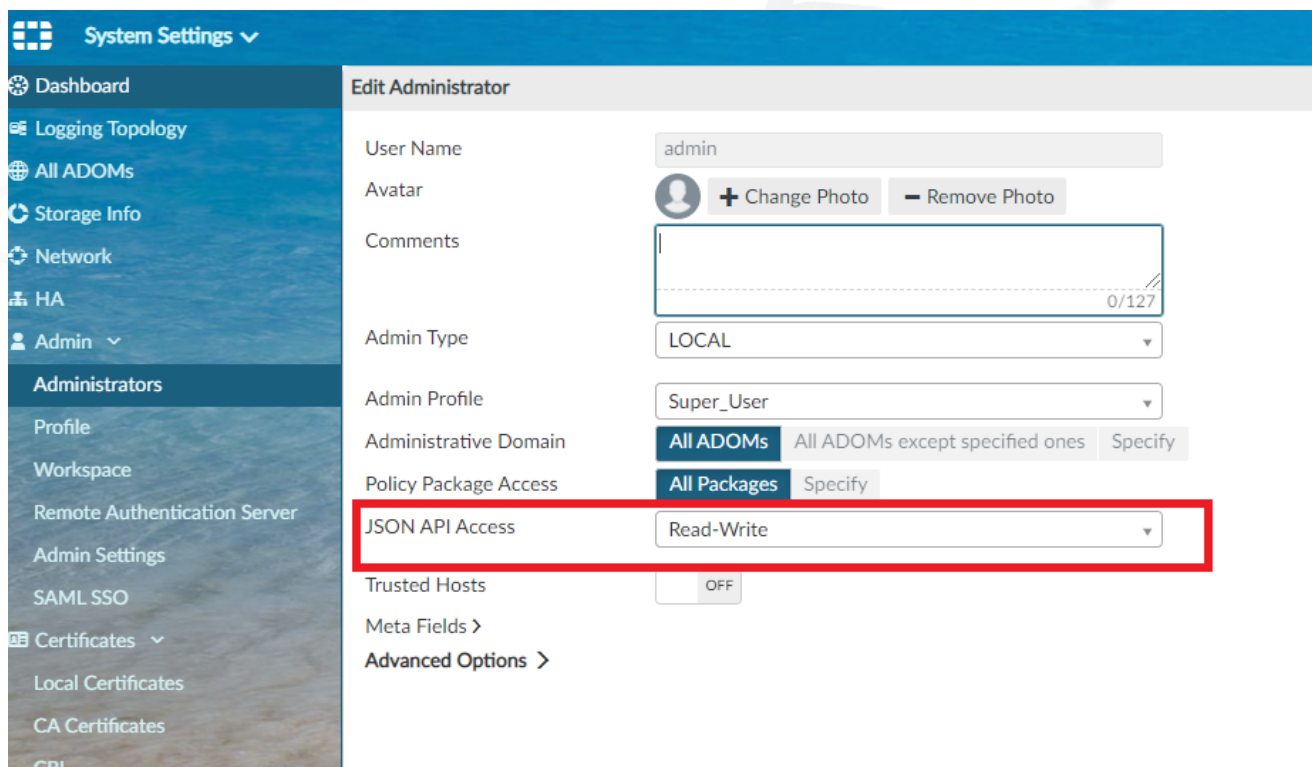
This section includes the following topics:

Connecting to FortiManager and FortiAnalyzer

When the MSP admin logs in for the first time, they need to connect FortiPortal to FortiManager and FortiAnalyzer.

The following image shows that the API-enabled administrator account in FortiManager is used within FortiPortal. This user allows FortiPortal to administer changes through the underlying APIs. The exact same process is followed when adding FortiAnalyzer to FortiPortal.

Once FortiManager and FortiAnalyzer are connected, the MSP admin creates customer tenants by navigating to *Customers > Create a new customer*. In the example below, *AcmeCorp* is the customer created.

Customer creation and role-based access

FortiPortal provides role-based access with very granular controls. The MSP admin can create different profiles and downstream access to end-customers.

After a customer tenant is created, the MSP admin needs to provide information about what FortiManager ADOM(s) and FortiGate devices the customer can access.

The MSP admin can also select what configuration options to expose to the end-customer.

Looking at *Policy and Object Edit Permission* and *Widget Permission* in the following image shows how easy it is to customize the different configuration options exposed to the customer tenant. In this example, all monitors have been selected, with the ADOM named *DEMO* being added.
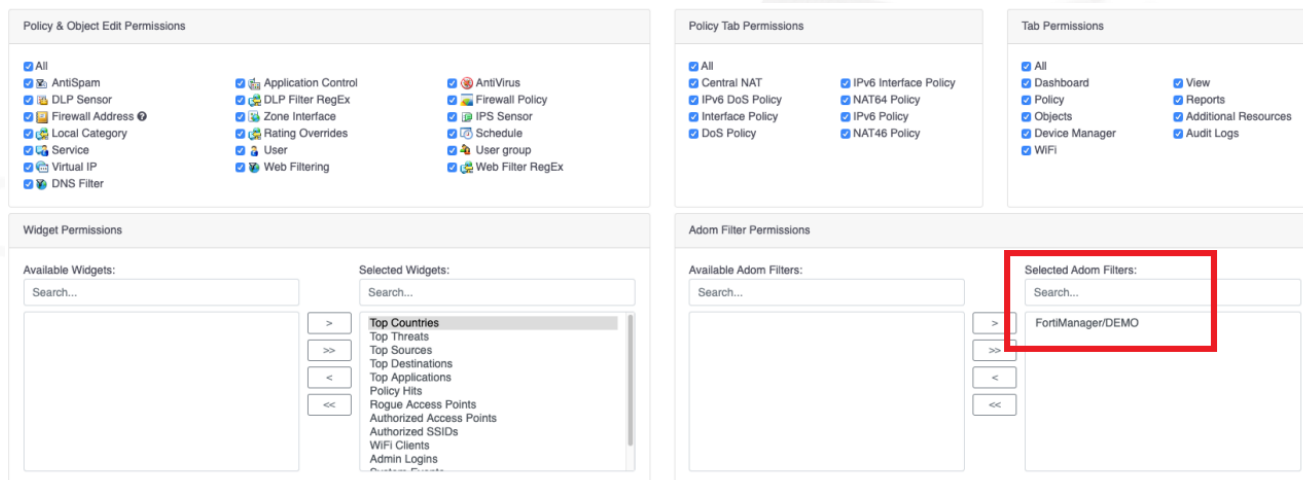
Now that the *AcmeCorp* customer has been created, we must assign devices within the ADOM named *DEMO* on FortiManager to *AcmeCorp*. It is essential to highlight that a customer can have several assigned ADOMs. If the customer admin needs to have access to multiple ADOMs, these ADOMs must be selected to provide complete visibility and access control. The following image shows how to choose available devices for a specific customer.
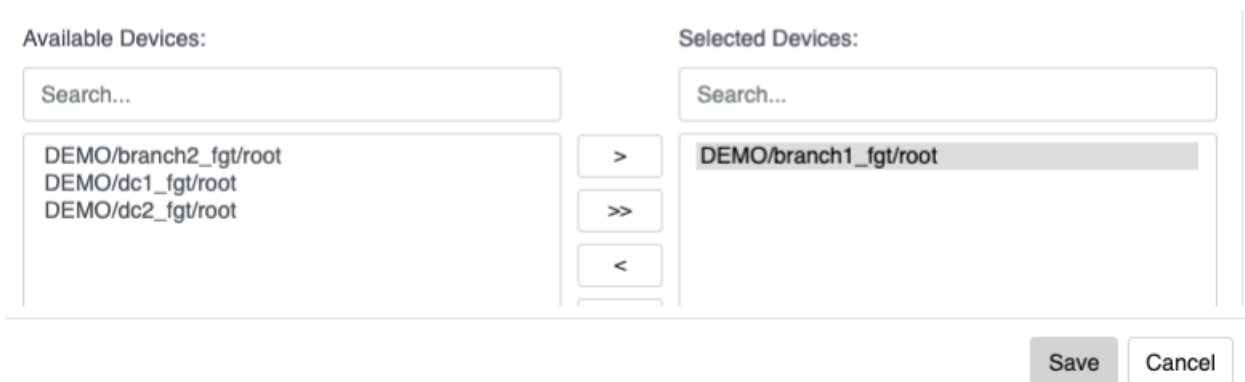
Customer view

Now that *AcmeCorp* has been created, together with *AcmeCorp* admin accounts, they will be able to log into FortiPortal using their unique credentials. It is also possible to set up administrator access using multi-factor authentication.

Logging into FortiPortal with unique *AcmeCorp* credentials will take the *AcmeCorp* admin into their portal, displaying all the views and configuration options previously defined by the MSP admin.

FortiPortal tree view on the left side shows the configuration and monitoring views available to AcmeCorp. As illustrated in the following image, the admin selects *View-Monitors* from the tree view, displaying several available monitors.

Like FortiAnalyzer, this view allows a user to drop down to look at other monitors such as Applications, Websites, and Top Threats, to name a few.
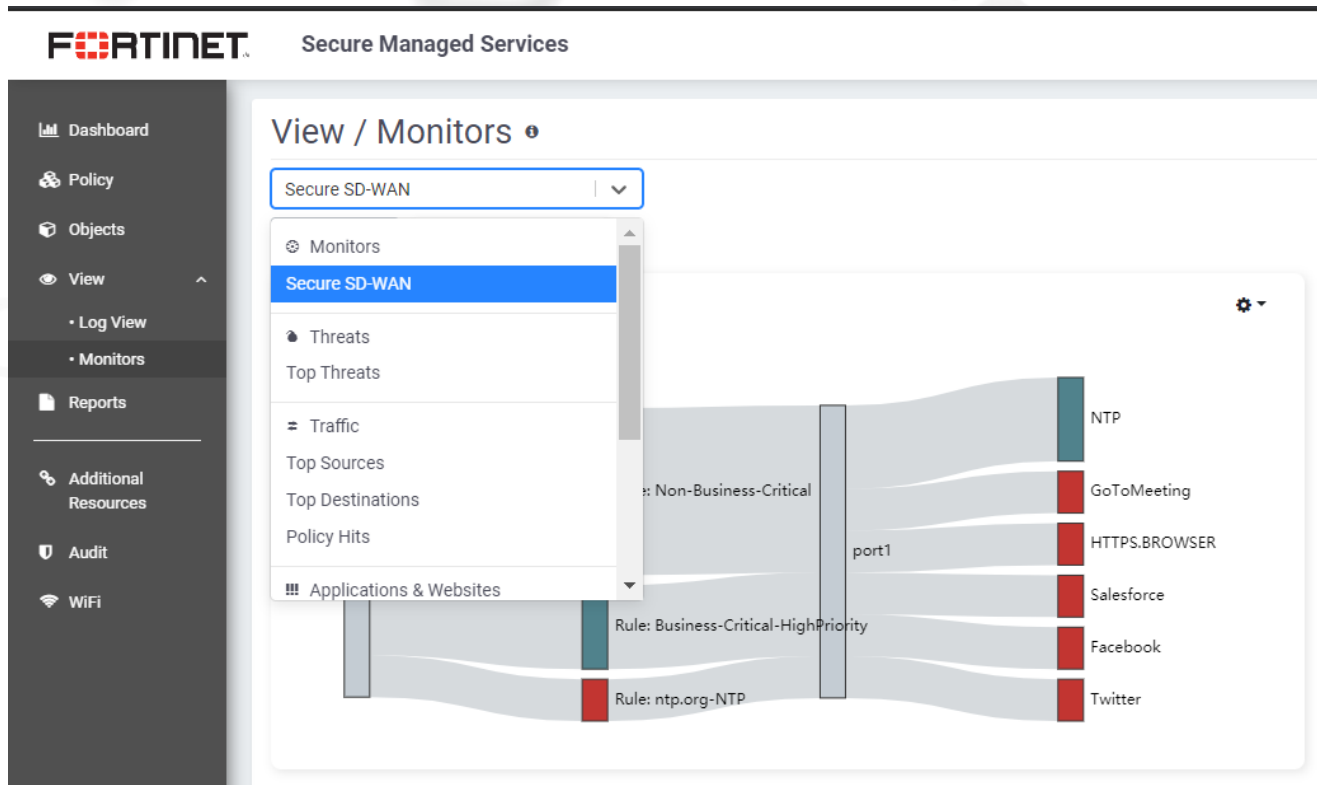
Selecting *Secure SD-WAN* within the monitoring view brings back the same SD-WAN monitoring widgets shown previously in the FortiAnalyzer monitoring section.
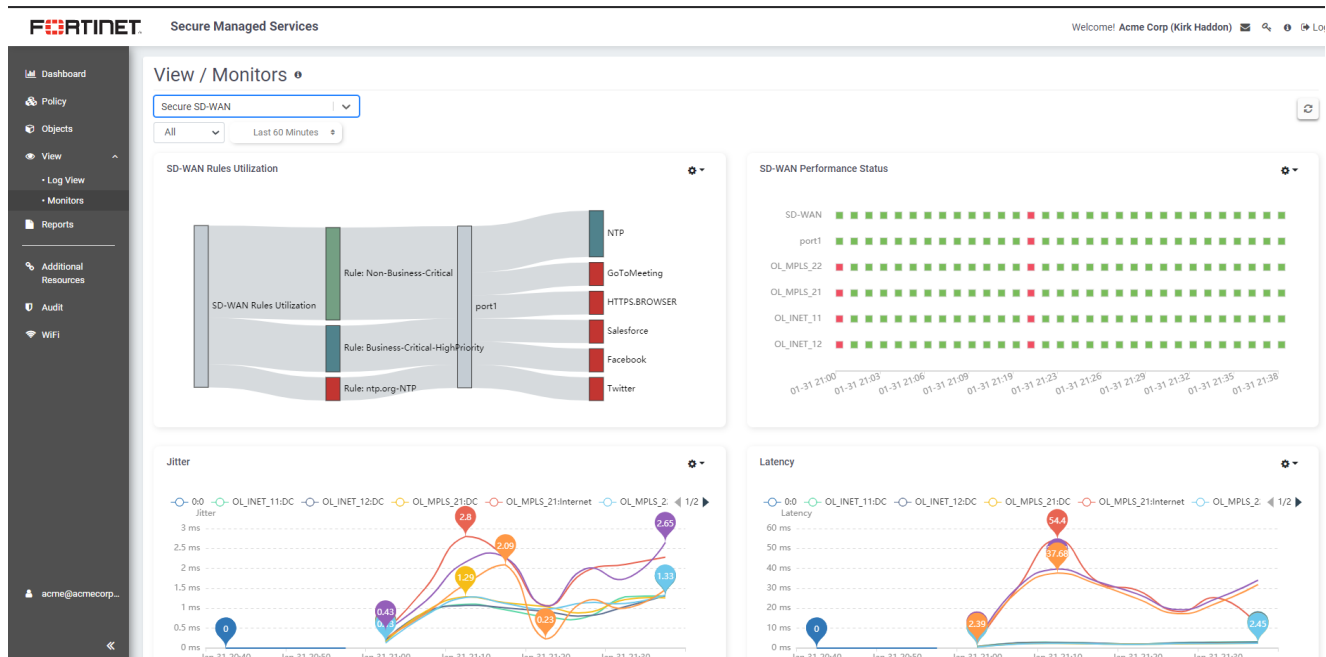
Moving to *Top Website Domains* will show browsing time, threat score, and session count.

Accelerator
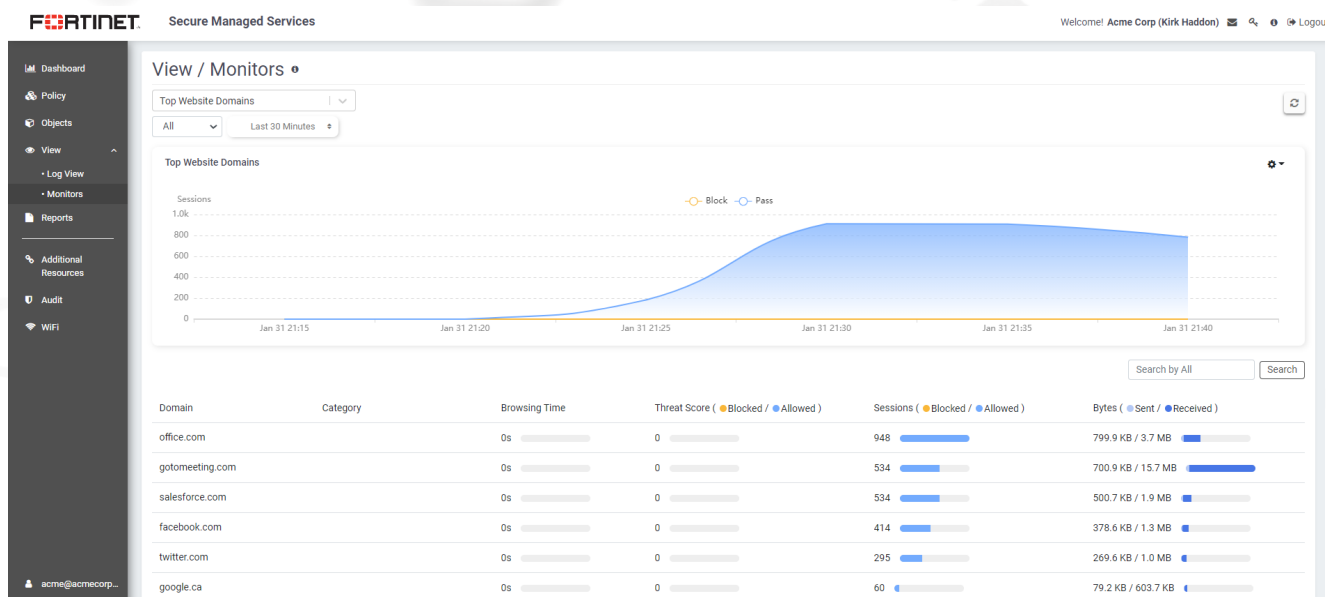
Navigating to *Log View* allows the user to look into traffic logs and all security service logs enabled on a per-device basis.

Clicking on *Map view* allows the customer to select each device and see whether the devices are in and out of SLA. This information is color-coded: green when within SLA and orange when out of SLA.

Switching to *Table view*, more detailed, per-device information is deployed.

The above views allow a customer to look initially at higher-level SD-WAN performance before honing in on specific devices. The following example provides details of a single SD-WAN FortiGate called *branch1_ FortiGate*. It shows bandwidth and traffic overview graphs, together with jitter, latency, and packet loss charts. These charts also provide SLA thresholds value, defined within the SD-WAN template. This provides the user with a simple way to see if any traffic has breached an SLA.
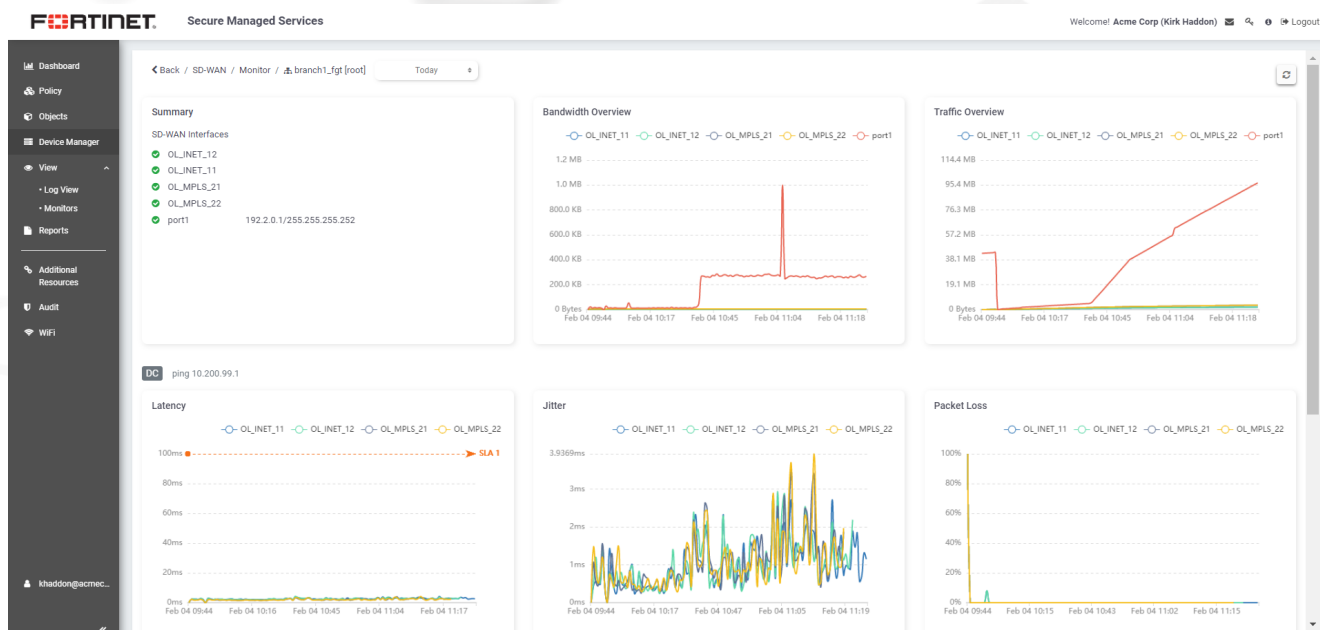
For example, within the DC SLA probe, the line is set to 100 ms of latency, and we can see that in this particular time frame, traffic has not violated the SLA.