



DEFINE • **DESIGN** • DEPLOY

FortiSwitchOS

Switching Reference Architecture Guide

Version 7.2.3

FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO GUIDE

<https://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTINET TRAINING & CERTIFICATION PROGRAM

<https://www.fortinet.com/training-certification>

NSE INSTITUTE

<https://training.fortinet.com>

FORTIGUARD CENTER

<https://www.fortiguard.com>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdoc@fortinet.com



March 20, 2023

FortiSwitchOS 7.2.3 Switching Reference Architecture Guide

11-723-840780-20230320

TABLE OF CONTENTS

Change log	4
Introduction	5
Intended audience	7
About this guide	7
Campus architectures	8
Wired local area network basics	10
Secured LAN	12
Network access control	14
Reference architectures	16
Security Fabric integration through FortiLink	16
MCLAG	16
Tiered architecture	17
Leaf-and-spine data center architecture	18
Network design principles	20
Dimensioning	20
Quality of service	21
Resiliency	22
Tier-1/core layer resiliency	22
Tier-2/aggregation layer resiliency	22
Tier-3/access layer resiliency	23
Future proofing	24
Core layer	26
Core layer platforms	27
Aggregation layer	29
Aggregation layer platforms	30
Access layer	31
Access-layer deployment recommendations	31
Access layer platforms	35
Management	37
Final design	39
SD-branch architectures	40
Small SD-branch	41
Medium SD-branch	42
Large SD-branch	43
Management	43
Appendix: Documentation references	44

Change log

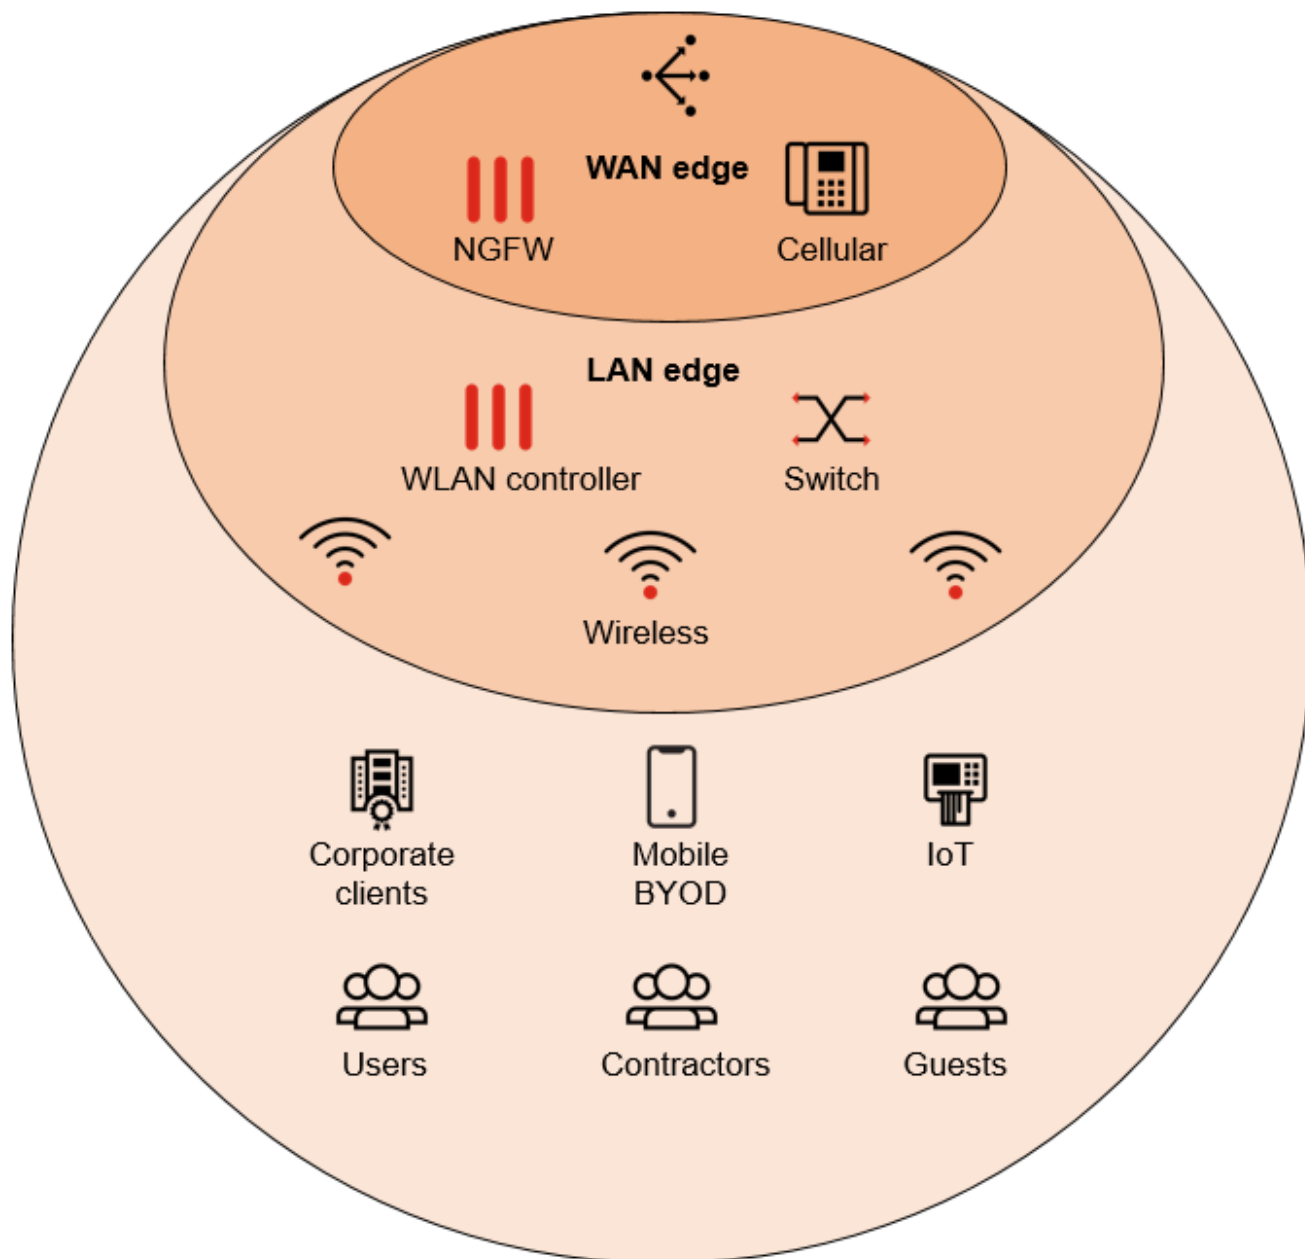
Date	Change Description
March 20, 2023	Initial release

Introduction

This document provides reference architectures for configuring networks for small campuses, large campuses, small software-defined (SD) branches, medium SD-branches, and large SD-branches.

- “Campus” covers a wide range of networks and locations, from multiple floors in an office tower to a university campus of a couple hundred acres. However, the Fortinet recommended architecture remains the same for this wide range. Generally speaking, a single location with 100 to 100,000 devices will fit this campus model.
- An effective SD-branch design should consolidate WAN and LAN capabilities to simplify remote office infrastructure and optimize operations without introducing new risks.

The nature of a campus deployment is not just the large physical size of the network, but the greater complexity of who and what is using that network. As the network grows in size, security needs become more complex with more categories of end devices. Fortinet calls this network configuration the *local area network (LAN) edge*.



Physically, the LAN edge is just the access layer, but securing the access layer now has to account for a bewildering mix of devices: enterprise owned, end-user owned, guest users, known users, Internet of things (IoT) devices with no associated users, and so on. Unlike other vendors, Fortinet takes a security-first approach to wired LANs with the FortiSwitch line and wireless LANs (WLANs) for *security-driven networking*.

Fortinet's security-driven networking strategy tightly integrates an organization's network infrastructure and security architecture, enabling the network to scale and change without compromising security. This next-generation approach is essential for effectively defending today's highly dynamic environments—not only by providing consistent enforcement across today's highly flexible perimeters, but by also weaving security deep into the network itself in a Security Fabric.

Intended audience

This guide is intended for network and security architects and engineers who are interested in deploying Fortinet's FortiSwitch units in a new environment or in replacing their equipment in an existing environment. Readers are expected to have a firm understanding of networking and security concepts.

About this guide

This guide provides reference architectures for secure campuses and SD-branch network topologies. These reference architectures offer examples of the best way to configure your network for each use case. The Fortinet solutions in this guide make it easier to configure these challenging topologies.

This document is not intended to be a step-by-step configuration guide. Instead, it is meant to be the starting point in your network design, where you begin to draw out the architecture that will be used to meet your specific needs. For more information and documentation about the topics covered in this document, refer to the Fortinet Document Library at <https://docs.fortinet.com>.

Campus architectures

A *campus* is a sizable network composed of a large building or multiple buildings with different purposes. The density of ports and users depends on the use case, and, even in the same organization, there can be multiple different implementations of the campus. For example, a campus can be any of the following:

- A university
- A conference hall
- An airport
- A Navy ship
- A site with manufacturing facilities, shipping docks, warehouses, and administrative offices
- A financial services headquarters

Because of the diverse use cases, a campus network cannot be considered a one-size-fits-all implementation. Nevertheless, the essence of the design proposed by Fortinet in campus environments is based on the following factors:

- Security
- Usability
- Scalability
- Flexibility
- Resiliency

Whether you require a very high availability of the network or just continued availability with lower levels of resiliency, this guide discusses the options for the recommended designs proposed here, with a focus on the most critical environments because they are the most difficult to grasp, handle, and deploy. You can choose less demanding network designs if waiting for a firewall, switch, or access point (AP) replacement is acceptable or if you rely on spare units stocked at some facility. The proposed designs in this guide aim for seamless operation for high availability, taking into consideration multiple hardware issues and denial-of-service (DoS) attacks on specific elements in the network.

You can also provide wired and wireless secure guest access for your visitors, contractors, and temporary workers, such as interns, with the possibility to segment the network and isolate their traffic.

The trend towards Everything-as-a-Service has dramatically disrupted the traffic path in the campus. The campus needs to be flexible and scalable now. This includes voice-conferencing and video-conferencing between people in the same building, each person using an individual device instead of congregating in meeting rooms, and communicating with remote workers who are not in the office everyday. It is now commonplace for people to launch Zoom or GoTo Meeting, for example, in a one-to-one meeting, just to be able to share their screens, hence forwarding all the traffic not to one another, but to a central location hosted in the cloud through the firewall. Even in the case of Zoom's on-premise solution, the traffic flows through a central meeting connector VM in the company's private cloud, not between users directly. Telephony connections still need to use the Zoom Cloud, as explained in the [Zoom Help Center](#).

Because of this move away from having most applications accessible through the traditional data center, the edge network needs to be even more secure and intelligent to prevent the applications (also) hosted in the cloud from escaping and bypassing controls.

This guide recommends disassociating the logical path from the underlying physical layout of the network to cope with these emerging trends that are reducing the east-west traffic in a campus to very limited use cases of direct VoIP phone calls and direct printing (without a network-attached printer spooling through a print server).

Because of the pervasiveness of 802.1X supplicants in both wired and wireless clients (and in more devices like wireless access points), as well as the adoption of network access control (NAC) in network and security elements, the design of a campus must be able to dramatically evolve and simplify the tasks of deploying a new switch or AP.

This section covers the following topics:

- [Wired local area network basics on page 10](#)
- [Secured LAN on page 12](#)
- [Reference architectures on page 16](#)
- [Network design principles on page 20](#)
- [Core layer on page 26](#)
- [Aggregation layer on page 29](#)
- [Access layer on page 31](#)
- [Management on page 37](#)
- [Final design on page 39](#)

Wired local area network basics

Switched LANs provide the basic access for network devices to communicate with each other and with resources locally adjacent (in the same room, same floor, same building, and same campus) without having to cross a wide area network (WAN) between sites. Interconnecting a group of LANs requires a network with full connectivity to the internal resources (such as the data center, phones, and printers) through a set of inter-switch links of different types. For scalability purposes, LANs are often segmented using virtual LANs, while, for security purposes, the traffic often has to be policed and filtered to only allow interactions between users and resources previously authorized. Therefore, the vast majority of the traffic arriving at a switch port is sent to an uplink (trunk) to another switch, which forwards the traffic to another device through one of its uplinks (with another switch or a firewall or a router, depending on the size of the network). The network could be a few floors in a building, a single building, or a group of buildings located near each other. A subnet becomes a set of users with similar roles, for example, users who work for the same department in a company.

This hierarchical physical design of a secure campus wired LAN is very common and involves two or three levels between the access switch and the core equipment, such as a firewall or a router. It allows the network to grow, minimizes the number of uplinks, provides the potential for reliability, and overcomes the 100-meter Ethernet link limits over copper by cascading the high-bandwidth fiber optic connections between switches.

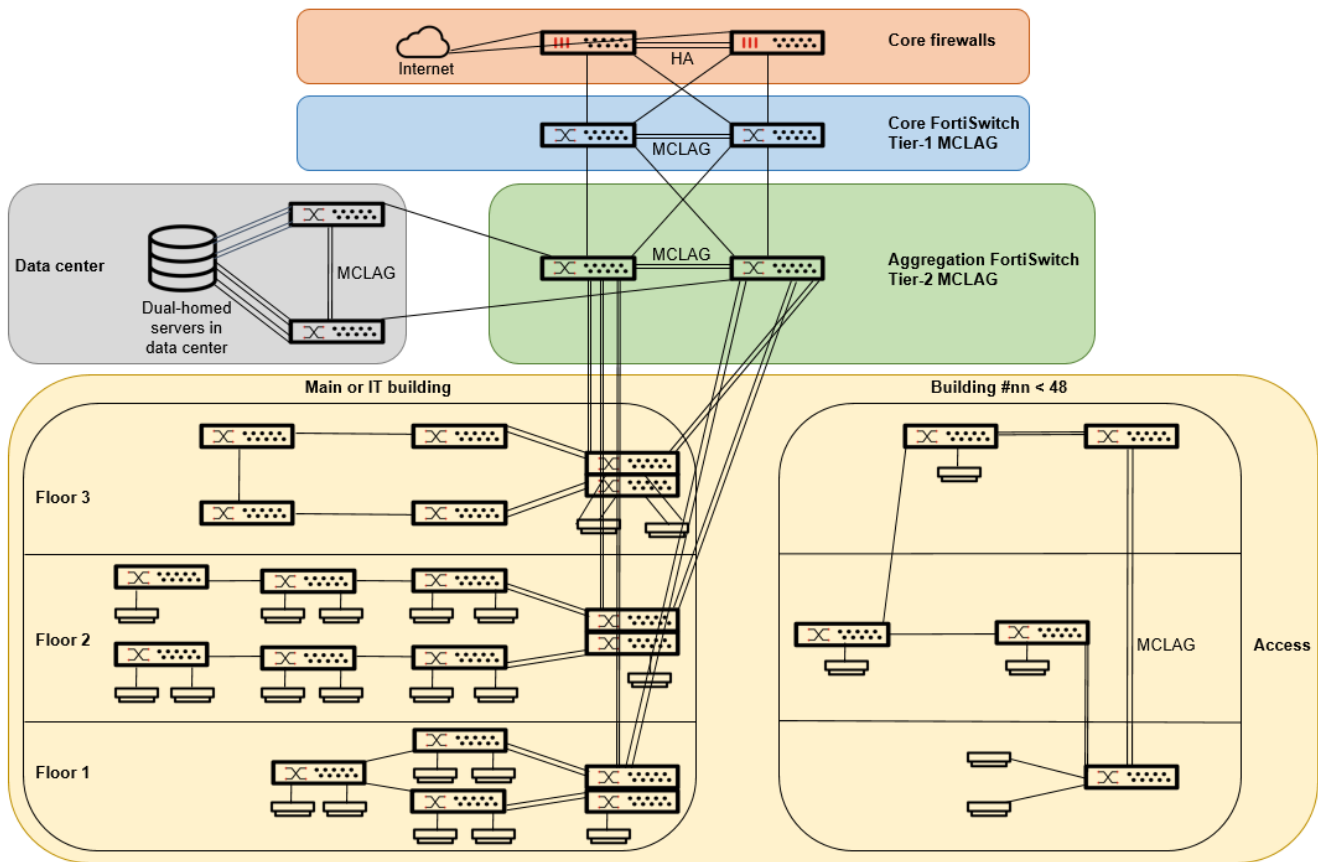
This hierarchical design model breaks down the design into different layers, which simplifies the deployment, extensibility, and management of the network and allows each layer to implement specific functions. At the same time, this design helps constrain operational changes to a subset of the devices and helps in troubleshooting by layering the detection of issues. The induced modularity allows you to create design elements that can be replicated and, therefore, is a straightforward way to scale.

The following are the typical three layers:

- **Access**—This layer provides direct wired connectivity to the network for endpoints and users.
- **Aggregation**—This layer aggregates access layers and provides connectivity among them, to data centers, and to other services.
- **Core**—This layer supplies connectivity to security resources such as hardware-accelerated Secure Sockets Layer (SSL) inspection, filtering of communications between segments, access to the Internet, as well as the connection between aggregation layers for large LAN environments.

Each layer provides different functions and capabilities to the network. Depending on the size of the campus, you might need to collapse the core and aggregation layers (especially when WiFi is the primary access for all end-user devices in the campus) or to use all three layers.

The following figure shows an example of a three-tier hierarchical LAN design with multiple buildings and one data center.



Secured LAN

The access layer is where the first security measures get enforced on the end devices when access must be revoked, granted, or restricted. This layer is where it is most important to apply network access control, with or without 802.1X authentication being applied on the access ports. Detection, visibility, and authentication are very important for the security of the organization so that users and devices get categorized and protected, as well as different access rules applied (which can also be tracked in case of an attack), while protecting the LAN edge of the network. A potential attacker might enter a building to plug a device into an enabled network port and try to gain access to the network. Potential attackers might even attack remotely, for example, by using wireless extenders to provide access from the parking lot or from a neighboring building. To get access to resources, you need to require more than just physical connections to the network.

Fortinet recommends using 802.1X authentication for wired and wireless access. 802.1X is an IEEE standard used for restricting unauthorized access to the network by making users (and devices if needed) authenticate before they are allowed onto the network. It relies on an authentication server (usually using RADIUS, such as the FortiAuthenticator server) to validate the credentials based on a local or remote database, which is often linked to Microsoft Active Directory through the Lightweight Directory Access Protocol (LDAP) in typical office environments.

Managed FortiSwitch units support EAP-PEAP, EAP-TTLS, EAP-TLS, and EAP-MD5. The FortiSwitch units communicate the user information (credentials or a certificate) from the supplicant on the client device as authenticators through the FortiGate device using FortiLink. The FortiSwitch units open only the limited ports required for such communication and nothing else until the RADIUS/Diameter server confirms the user identity and validates it, usually with an attached VLAN linked to that user in the Access-Accept response. This VLAN can then be shared between wired and wireless access, making it easier to implement segmentation and policies based on the user itself.

When employees change jobs, you change the VLAN associated with the employee as a single point of configuration to inherit the rules defined for the other colleagues in the same department or job role. You can even configure a guest VLAN for unauthorized users and a VLAN for users whose authentication was unsuccessful, leading to very limited network access or a walled garden space. Using VLANs removes the burden and cost of having to deploy a layer-3 infrastructure at the access or aggregation layer in the network and helps simplify applying rules and restrictions in parallel with reducing the broadcast domain per type of user (or per team) instead of geographically.

This design provides location-independent network access with the following advantages:

- Allows for a simpler design than complex layer-3 infrastructures because a user will be part of the same VLAN wherever they access the network (wired or wireless) across the entire campus without the need to change the user's subnet.
- Provides the necessary secure option of forwarding all user traffic to the firewalls for inspection before the traffic is allowed to communicate with devices and users in the same VLAN (intra-VLAN traffic).
- Offers real-time updated security with smart application identification (optional) being performed through deep-packet inspection, even for apps that are encrypted (and for several thousands of them).
- Improves the employee experience and productivity by applying the same set of rules through the network at the intranet level.

Obviously, there could be some endpoints that do not support 802.1X authentication (such as headless devices, IoT, and old endpoints), while still needing network connectivity and being authenticated/placed in their corresponding network segments to operate and develop their functions. To accommodate those, the access layer can be flexible and dynamic enough to accommodate the access for these devices without compromising the security. This balance is achieved through the use of dynamic port policies (DPPs), a native function of FortiSwitch units when managed by a FortiGate device. Features like embedded NAC when managed through FortiLink are also possible, even if the

implementation of a full NAC solution is not chosen. MAC-based authentication mechanisms are not recommended because they are less secure, but they are still available.

It is also in the access layer that devices can be prevented from taking over roles of other devices on the network, such as Dynamic Host Configuration Protocol (DHCP) servers, which could greatly disrupt the operations of other endpoints in the same VLAN or, even worse, provide them with the wrong settings and divert the traffic through unsanctioned devices where potential threat actors might try to intercept it.

The FortiSwitch units, managed by a FortiGate device using FortiLink, offer a solution to tackle all access requirements in a seamless, headache-free, and smooth implementation. The following are some of the built-in features on the FortiSwitch units:

- DHCP snooping to prevent devices from becoming rogue DHCP servers in the network.
- Dynamic Address Resolution Protocol (ARP) inspection (DAI) and lockdown prevents man-in-the-middle attacks and IP address spoofing by checking that packets from untrusted ports have valid IP-MAC-address binding. DAI allows only valid ARP requests and responses to be forwarded.
- Port flap guard detects how many times a port changes status during a specified number of seconds, and the system shuts down the port if necessary. This helps to avoid unwanted network topology changes and convergence actions due to threat actors trying different things on the port or connecting rogue routing/switching devices that could pose a risk to the network convergence and topology.
- Storm control uses the data rate (packets/second) of the link to measure traffic activity, preventing traffic on a LAN from being disrupted by broadcast, multicast, or unicast storms on a port that could be dropped. When the data rate exceeds the configured threshold, storm control drops excess traffic.
- A configurable learning limit for dynamic MAC addresses on ports, trunks, and VLANs (port security).
- A new FortiGuard service identifies IoT devices through an IoT Detection Service license.
- Static, dynamic, and scheduled access list control (ACL) available to control the traffic flows in very low-level detail. You can use ACL to redirect traffic to different ports, read/mark the quality of service (QoS), rate-limit specific traffic, or even quarantine a device.

Loops in a layer-2 network result in broadcast storms that have far-reaching and unwanted effects. Fortinet loop guard and Bridge Data Unit Protocol (BPDU) guard help to prevent loops. When loop guard is enabled on a switch port, the port monitors its subtending network for any downstream loops, for example, if an unmanaged device that does not support the Spanning Tree Protocol (STP) gets connected to the infrastructure.

Ultimately, you can block intra-VLAN traffic (traffic within the same VLAN or broadcast domain) by aggregating traffic using the FortiGate device that exclusively manages the whole fabric of switches. This prevents direct client-to-client traffic visibility at the layer-2 VLAN layer. Clients can only communicate with the FortiGate firewall. After the client traffic reaches the FortiGate controller, the firewall then determines whether to allow various levels of access to the client by shifting the client's network VLAN as appropriate, if allowed by a firewall policy and proxy ARP is enabled. This is useful in critical or very sensitive environments, such as operational technology (OT) networks, where any potential security breach that might spread easily with lateral movements must be rapidly contained or preferably prevented.

Automation can also play a key role in constraining and remediating security breaches in real time. In that regard, the FortiSwitch units can enforce dynamic and flexible quarantine policies directly from the access ports or WiFi access. From simply isolating a compromised device to rate-limiting its traffic or forwarding its user to a captive portal, an endpoint can be confined for further inspection when it triggers a suspicious activity. Moreover, this quarantine can be automated when the FortiSwitch unit and/or FortiAP unit managed by a FortiGate device works in collaboration with FortiAnalyzer unit. If the FortiAnalyzer unit has a valid license for Indicators of Compromise (IoC), it will inspect the device's traffic logs and issue a notification to the FortiGate device whenever those IoCs are found on the endpoints. This event can trigger an automatic action on the FortiGate device and place that device into quarantine on the access layer (wired and wireless).

Optionally, FortiSwitch ports can also be shared between virtual domains (VDOMs). VDOMs allow you to divide a FortiGate device with two or more virtual firewalls that create multiple independent units. VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations. VDOMs require careful planning, but they allow for designs (out of scope for this guide) with multitenancy in the same building, even when using a common network and security infrastructure. A single company inside a campus can use VDOMs to isolate traffic from independent business units or departments without having to multiply the network security and infrastructure elements. VDOMs allow for greater flexibility in the life cycle of the institution.

Network access control

Using the FortiGate device as a switch controller for the network adds the built-in NAC without the need of any license. The FortiGate device is responsible for assigning a device to the right VLAN based on the NAC policy when a device first connects to a switch port or when a device goes from offline to online. NAC provides a secure way to close down ports and allow ports to pass traffic only after devices get connected. Together, NAC and 802.1X authentication provide the most important part of zero-trust network access (ZTNA) at the access level without requiring a complex and dedicated solution from a third-party vendor and without forcing the use of FortiNAC for extended scenarios (which is primarily used in a multiple-vendor environment). Without FortiGate NAC (such as in competitors' switch deployment recommendations), the on-site access and device layer remains unprotected and an important potential threat.

You can configure a FortiSwitch NAC policy or DPP within FortiOS to match devices automatically based on the specified criteria, devices belonging to a specified user group, or devices with a specified FortiClient EMS tag. Devices that match the policy are assigned to a specific VLAN or have port-specific settings applied to them (such as QoS, LLDP profiles, VLAN policies, and 802.1X profiles).

NAC includes Fortinet access points (FortiAP units), which can be automatically recognized and provisioned through the use of the embedded and free NAC without requiring any external program or license. This greatly simplifies the deployment of Fortinet APs on a FortiSwitch fabric infrastructure.

Moreover, an integrated FortiGate NAC function is also provided to the FortiAP networks; this function uses a shared set of NAC policies with FortiSwitch units to simplify the deployment as a fabric solution. The NAC policy can be applied based on data from the user device list. There is also a wizard to help with configuring FortiGate NAC settings and defining a FortiSwitch NAC VLAN.

The following figure shows examples of patterns in FortiSwitch NAC policies that devices must match.

Name	Patterns	Assign	Matched Devices	Ref.
FortiAP	Operating System: FortiAP OS	VLAN Policy: FortiAP	FP431FTF20020843	1
Light_bulbs	Hardware Vendor: Xiaomi Device Family: Light	VLAN: iot_devices		0
Corporate-Laptops	Device Family: Lenovo	802.1X: 802-1X-policy-default		0
Corporate-Mobile And	Hardware Vendor: Huawei Device Family: Android Type: Mobile	VLAN: Home_vlan		0
Outlets	Hardware Vendor: Xiaomi Device Family: Plug	VLAN: iot_devices	chuangmi-plug-m1_mio87931293 chuangmi-plug-m1_mio88284046 chuangmi-lr-v2_mio88058487 chuangmi-plug-m1_mio88292691 chuangmi-plug-m1_mio88073242	5
iot_controller	Hardware Vendor: Xiaomi Device Family: Controller	VLAN: iot_devices		0
Automation_controller	Hardware Vendor: Raspberry Pi Device Family: Raspberry Operating System: Debian	VLAN: iot_devices	openHAB2	1
Brother_printers	Hardware Vendor: Brother Type: Printer	VLAN: Home_vlan	Printer_wifi	1
Windows	User: vanes	VLAN: Home_vlan	Vanessa-laptop	1
Corporate-Guests	Device Family: *	VLAN Policy: Guest_LAN	MacBook-Pro_Raul	1
Onboarding VLAN		VLAN: onboarding		

The following figure shows a list of devices that matched the patterns in FortiSwitch NAC policies.

Matched Devices ✕					
Refresh	Search				
MAC Address	Matched NAC Policy	Assigned VLAN	Status	Last Known Switch	Last Known IP
openHAB2	Automation_controller	iot_devices	Enable	108-Attic (S108EP4N17000266)	
Printer_wifi	Brother_printers		Disable	108-Attic (S108EP4N17000266)	
MacBook-Pro_Raul	Corporate-Guests		Disable	108-Attic (S108EP4N17000266)	
FP431FTF20020843	FortiAP		Enable	108-Attic (S108EP4N17000266)	
chuangmi-plug-m1_miio87931293	Outlets	iot_devices	Enable	108-Attic (S108EP4N17000266)	
chuangmi-plug-m1_miio88284046	Outlets	iot_devices	Enable	108-Attic (S108EP4N17000266)	
chuangmi-ir-v2_miio88058487	Outlets	iot_devices	Enable	108-Attic (S108EP4N17000266)	
chuangmi-plug-m1_miio88292691	Outlets	iot_devices	Enable	108-Attic (S108EP4N17000266)	
chuangmi-plug-m1_miio88073242	Outlets	iot_devices	Enable	108-Attic (S108EP4N17000266)	
Vanessa-laptop	Windows		Disable	108-Attic (S108EP4N17000266)	

The following figure shows a NAC policy applied to a FortiSwitch port to which a FortiAP unit is connected.

Port Trunk Faceplates										
Port	Description	Access Mode	Enabled Features	Native VLAN	Allowed VLANs	PoE	Device Information	DHCP Snooping	Transceiver	Bytes (Sent/R)
108-Attic - S108EP4N17000266										
port1	Forti AP Attic	NAC	Edge Port Spanning Tree Protocol	onboarding	Home_vlan iot_devices qtn.fortilink Guests_VLAN	Powered 10.90W	MacBook-Pro_Raul chuangmi-plug-m1_miio87931293 chuangmi-ir-v2_miio88058487	Untrusted		52.96 GB
port2		Normal	Edge Port Spanning Tree Protocol	FortiAP	Home_vlan iot_devices qtn.fortilink	Unpowered		Untrusted		0 B
port3		Normal	Edge Port Spanning Tree Protocol	Guests_VLAN	qtn.fortilink	Unpowered		Untrusted		0 B

Reference architectures

This section describes how the reference architecture for campus deployments can be implemented with the Fortinet solution, leveraging the tiered recommendation described in [Wired local area network basics on page 10](#) (access, aggregation, and core layers).

The solution is based on the FortiSwitch units being managed by a FortiGate device using FortiLink.

Security Fabric integration through FortiLink

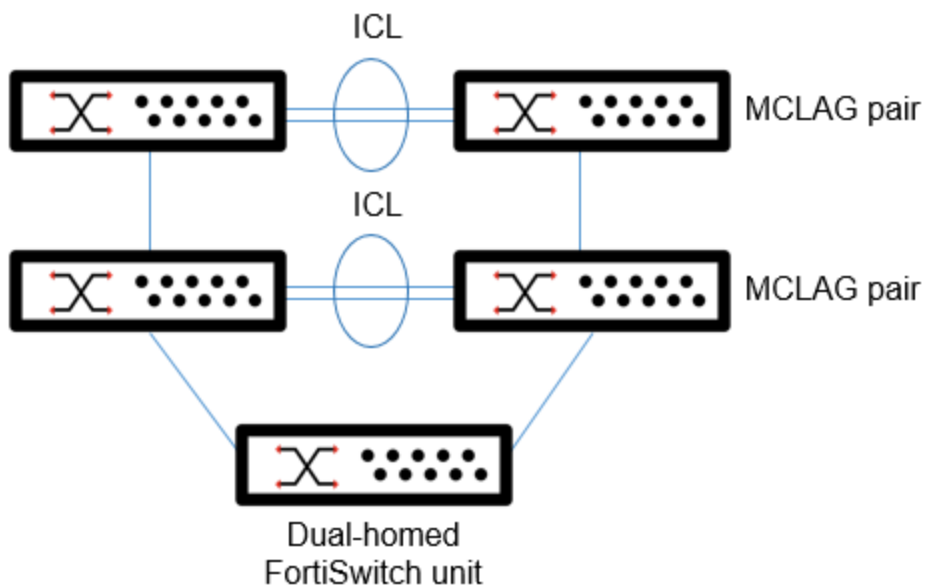
All the proposed designs rely on using FortiSwitch management from the central FortiGate platform through FortiLink or from a FortiManager unit managing FortiGate devices. FortiLink is Fortinet's solution to extend the Fortinet Security Fabric to the Ethernet switch port level. This link allows the same policies configured and applied to FortiGate interfaces to be applied to the FortiSwitch Ethernet ports, reducing the complexity and decreasing the management cost. With network visibility, it is easy for system administrators to implement and manage security and access-layer functions enabled and managed through a single console and centralized policy management (including role-based access and control). Users and devices can be authenticated against the same database and have the same security policy applied regardless of how or where they connect to the network. FortiManager provides the single-pane-of-glass network operations center (NOC) view of the overall deployment, encompassing in the same place the security and network elements of both the campus and multiple remote branch offices, especially when using SD-branch. This allows network and security teams to work hand-in-hand with the same view of the network and all elements of the security through the Fabric. This is a major differentiator for Fortinet in the market.

MCLAG

The key to scaling and optimizing the reference architecture is the multichassis link aggregation group (MCLAG) . It allows a switch pair to appear as a single device to other network elements. MCLAG has the advantage of allowing all uplinks between switches to be active and passing traffic, resulting in higher capacity and availability. It helps make the network agnostic of STP, which would block redundant links instead of using their capacity. Because it usually takes up to 45 seconds for a traditional (not rapid per-VLAN spanning tree) STP port to transition from blocking to forwarding, MCLAG also boosts the convergence time for much quicker recovery from any incident on the forwarding path of active traffic. All uplinks can therefore be used at capacity to meet network demands and provide better oversubscription ratios.

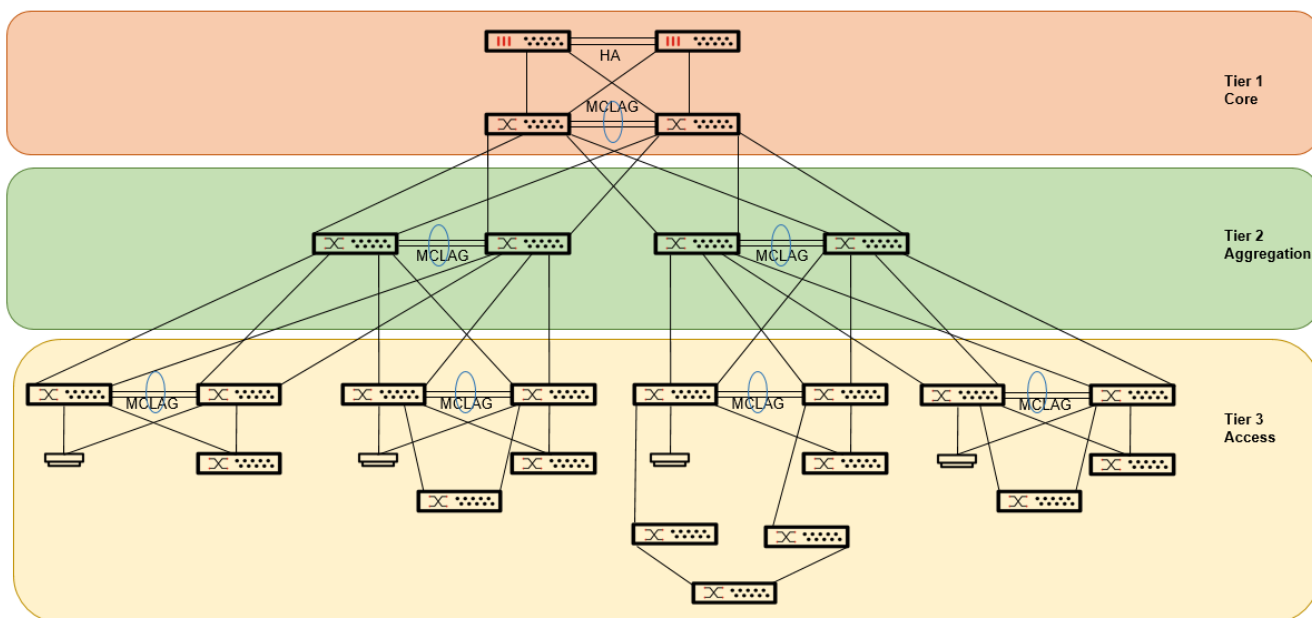
MCLAG switches are interconnected through dedicated links called interchassis links (ICLs). Fortinet recommends using at least two links for ICL redundancy, which can be, depending on the need, 10-GbE, 40-GbE, or 100-GbE ports. If a MCLAG member switch fails, the other member continues to operate without any change.

The following figure shows two FortiSwitch MCLAG pairs and two ICLs.



Tiered architecture

The following figure shows the tiered reference architecture for the campus and indicates which components are included in the core, distribution, and access layers, as well as in the three tiered layers of MCLAG.



Tier-1/core layer:

- The core layer comprises one, two, or up to four FortiGate devices acting as the FortiSwitch and FortiAP controllers. The layer-3 routing features of FortiGate devices are used for all inter-VLAN traffic. The FortiGate devices provide the NGFW for enforcing security measures for internal traffic.
- A maximum of two FortiSwitch units at this level forward all layer-2 traffic. Only east-west links are used for layer-2 traffic flowing from one aggregation stack to another. To maximize convergence, resiliency, and bandwidth, Fortinet recommends a pair of FortiSwitch units configured as an MCLAG pair and connected with the FortiGate devices using full mesh.
- This layer could be used to connect servers and headquarters resources shared across the network.
- Typical links used at this tier would be 10 GbE or an aggregation of various 10-GbE, 25-GbE, 40-GbE, or 100-GbE interfaces.

Tier-2/aggregation layer:

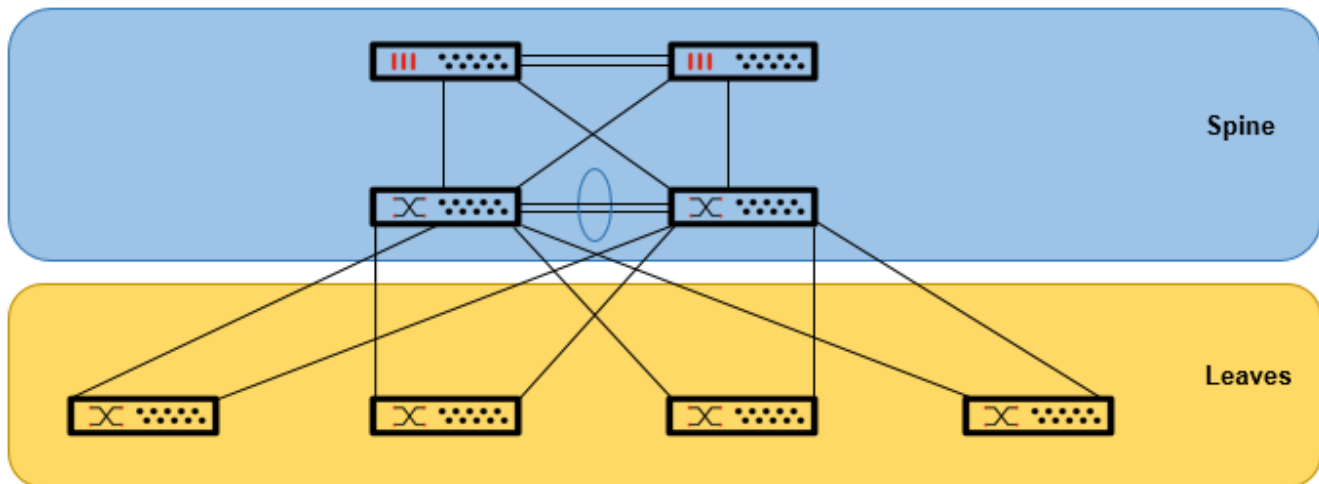
- The aggregation layer comprises one or multiple FortiSwitch MCLAG pairs, which consolidate all traffic coming from their corresponding access stacks. Only east-west links are used for layer-2 traffic flowing from one access stack to another.
- Depending on the final design, you could use this layer as a top-of-rack (ToR) level or delegate this function to the next layer below. The tier-2/aggregation layer can be the entry to different buildings or geographical locations within the campus.
- Typical links used at this tier would be 10 GbE or an aggregation of various 10-GbE links to connect towards the access layer; 10-GbE, 25-GbE, 40-GbE, or 100-GbE interfaces would be used as uplinks towards the core layer.

Tier-3/access layer:

- The access layer comprises one or multiple FortiSwitch MCLAG pairs, which provide full traffic capacity towards the aggregation layer.
- These MCLAG pairs can also be used as TOR levels, to which more FortiSwitch units could be connected in a dual-homed topology or to form additional rings, which would extend the number of ports required at this point. The tier-3/access layer can be the entry to different floors within the building, where final endpoints are connected to the subsequent stack of FortiSwitch units, or FortiAP units can be connected to leverage the PoE capabilities of the FortiSwitch units.
- Typical links used at this tier are 10 GbE or an aggregation of various 10-GbE links towards the aggregation layer; 10-GbE, 5-GbE, 2.5-GbE, or 1-GbE links would be used towards the rest of the access FortiSwitch units, FortiAP units, and endpoints.

Leaf-and-spine data center architecture

Although this architecture is not normally used in campus deployments, it could be useful for data center designs. You can implement the Fortinet leaf-and-spine topology following a similar methodology as in [Tiered architecture on page 17](#) and as shown in the following figure of a leaf-and-spine reference architecture for a campus. This design provides deterministic latencies and performance to critical services in the network. Fortinet adds a bonus by also providing 360 degrees of end-to-end security, all the way down to the LAN ports.

**Spine:**

- This layer comprises the routing elements for layer 3, the FortiGate devices, and a pair of high-capacity FortiSwitch units configured as an MCLAG. This layer is the same as the tier-1/core layer.
- All traffic is processed and secured by this combination at this level; therefore, high-performance FortiGate devices are required with native 40-G or 100-G interfaces facing the FortiSwitch units.

Leaves:

- This layer comprises a myriad of different FortiSwitch units, depending on the requirements of the connecting devices. If specific requirements are needed to provide a highly resilient service, such as priority-based flow control or ingress pause metering, then the FortiSwitch models must belong to the high end of the family.
- You can have separate FortiSwitch units connected in a dual-homed topology into the spine, or you can reduce complexity by having pairs of MCLAG FortiSwitch units to provide high availability and resiliency for servers that might require redundant paths towards the spine.
- The “leaves” layer can be implemented only for the data center or alongside the traditional tier-2/aggregation layer being deployed for campus switching.
- Typically, FortiSwitch units from the 500 series and higher are used in this layer to provide 40-G connectivity with the spine.

Network design principles

There has always been a trade-off when designing networks. A network cannot allocate the maximum traffic that could be generated by each connected device at peak all at the same time. Oversubscription must be considered when designing a network.

Additionally, the network design needs to consider the current and future traffic demands to avoid becoming obsolete in a short period of time. This is a critical factor to consider with the introduction of more and more wired and wireless devices connected to the networks, the newest WiFi 6E (802.11ax) spectrum that could potentially offer multigigabit access to a single network access device, and even the adoption of access ports for end devices at 2.5 GbE or 5 GbE. In the data center, server connectivity at the 25-GbE speed is very common.

Redundancy and resiliency form both the third and fourth pillars of network design. The requirements for fast convergence, assurance of critical services, and avoiding single points of failure while maintaining the service level agreements (SLAs) must be clearly considered in detail when designing any network.

Dimensioning

When you build a multi-tiered network, you need to consider the bandwidth oversubscription ratios for every layer of the switching hierarchy. The upstream bandwidth at each layer must provide enough bandwidth for the added traffic of each active device under it. Nevertheless, not all devices transmit or receive data at their nominal bandwidth, nor at the same time, so the ratios that make the total size of the uplink do not need to be the sum of the total amount of downstream links. This "oversubscription" ratio of uplinks must be closely followed to avoid network bottlenecks, which could cause poor network connectivity for downstream devices.

A ratio of 20:1 used to be a common downlink-to-uplink access ratio. Downlink-to-uplink aggregation ratios were 4:1 when using 1 GbE at the access layer with 2x10-GbE uplinks to the aggregation layer and 2x10-GbE uplinks to the core layer. This was when networks were designed with multiple Ethernet ports per desk and multiple ports in conference rooms and shared areas. In networks where more than 75 percent of the devices are now wireless and IoT devices continue to enter the enterprise, the number of wired ports in the network is getting close to one per user.

Moreover, with the advent of new hardware, cheaper fiber runs, and higher performing chipsets, common deployments of access switches now use 4x10-GbE links (sometimes even 2x40-GbE links) to a set of aggregation switches. The aggregation switches then send traffic from the aggregation layer to a core layer through up to 4x100-GbE links (towards two core switches) and then connect the core switches to the FortiGate devices for the core security services; the routing uses 100-Gbps links.

Hence, the common access-to-aggregation ratios are now more around 10:1 when you use gigabit Ethernet ports at the access layer, and the common aggregation-to-core ratios are around 2:1. Fortinet recommends these ratios when you design modern campuses for a refresh or a new building.

The calculation becomes more difficult when mixing multigigabit ports for access points and gigabit Ethernet for access devices like video surveillance cameras, IP phones, and video conferencing setups.

The most important consideration is what the ratio becomes if one of the devices fails. As you add more switches to a floor, you must keep in mind the distribution of the uplinks across switches, and the impact on oversubscription during failure. For access points that are dual attached to two different switches (see the figure in [Tiered architecture on page 17](#)), the network connectivity is kept. In a scenario where you lose one aggregation switch in the path, you are fully

covered by the design described here because you have redundant links from the access layer to the aggregation layer (using 2x40-Gbps links or 4x10-Gbps links). Losing one of the core switches is also covered because the links from the aggregation layer to the core layer are 2x100 Gbps or 4x100 Gbps, still providing an acceptable 4:1 statistical ratio for the duration of the outage instead of a ratio of 2:1.

Quality of service

QoS allows you to set particular priorities for different applications, users, or data flows. FortiSwitch units support the following QoS capabilities:

- Mapping the IEEE 802.1p and layer-3 QoS values (differentiated services and IP precedence) to an outbound QoS queue number.
- Providing eight egress queues on each port.
- Policing the maximum data rate of egress traffic on the interface.
- Security is provided as part of the zero-trust network access because QoS coming from the devices is not trusted. Each traffic packet is remarked according to the switch classification. DoS potential threats are minimized by implementing this QoS control mechanism.

QoS involves the following elements:

- *Classification*—This process determines the priority of a packet. This can be as simple as trusting the QoS markings in the packet header when it is received so that the packet is accepted. Alternatively, it can use such criteria as the incoming port, VLAN, or service that are defined by the network administrator.
- *Marking*—This process involves setting bits in the packet header to indicate the priority of this packet.
- *Queuing*—This process involves defining priority queues to ensure that packets marked as high priority take precedence over those marked as lower priority. If network congestion becomes so severe that dropping packet is necessary, the queuing process selects which packets to drop.

Fortinet recommends implementing queuing on switches because oversubscription might result in congestion and cause packet loss, especially during an outage. Queuing ensures that your higher priority traffic like voice conferencing and video conferencing is statistically less impacted than lower priority traffic, which is more immune to packet loss. If you select weighted-random-early-detection for the drop policy, you can enable explicit congestion notification (ECN) marking to indicate that congestion is occurring without just dropping packets.

FortiSwitch units can enable many different scenarios where QoS can be adjusted for rich multimedia and multicast handling in your campus network.

FortiSwitch units can parse LLDP (LLDP-MED) messages from voice devices such as FortiFone units and pass this information to a FortiGate device for device detection. FortiSwitchOS proposes an automatic FortiLink voice default VLAN for voice devices, and you can use a dynamic port policy to assign a device to an LLDP profile, QoS policy, and VLAN policy. When a detected device is matched to the dynamic port policy, the corresponding policy actions are applied on the switch port.

Fortinet's ability to inspect all traffic at the core, even encrypted packets, allows you to excel at application awareness and high quality of experience (QoE) delivery for your workforce.

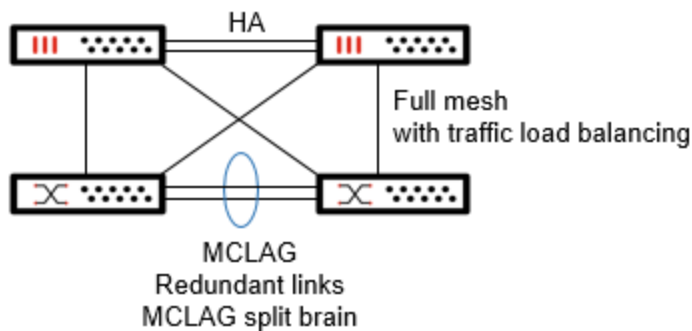
Resiliency

When considering design recommendations, a critical point is resiliency: avoiding single points of failure and ensuring a fast convergence to minimize service downtime. The following sections show how resiliency is accomplished in each layer of the reference architecture described in [Tiered architecture on page 17](#).

Tier-1/core layer resiliency

At this layer, resiliency is achieved by many different features. At the FortiGate level, the FortiGate Clustering Protocol (FGCP) provides failover protection because a cluster can provide FortiGate services even when one of the devices in the cluster encounters a problem (such as link failure, power loss, or memory or SSD failures). FortiGate devices can be configured in active-passive or active-active modes; refer to [High Availability](#) for further information. FortiGate clusters can be integrated into the load-balancing configuration using the FortiGate Session Life Support Protocol (FGSP) in a network where traffic is load balanced.

The following figure shows the redundancy mechanisms for the tier-1/core layer.



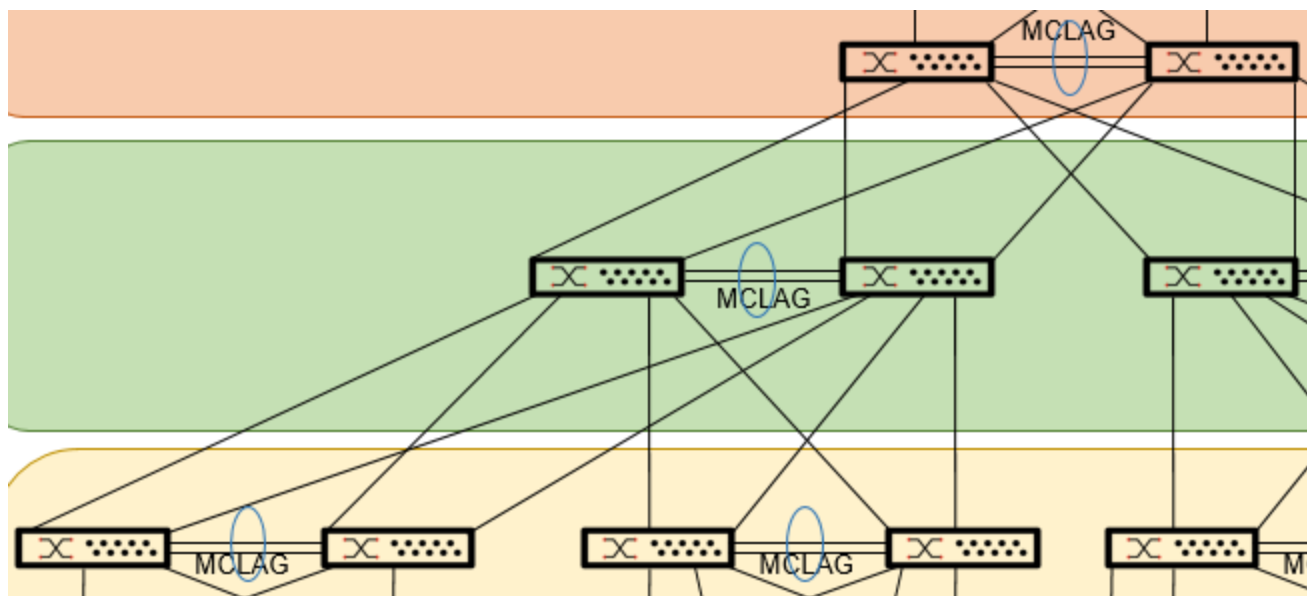
From the FortiSwitch perspective, resiliency is achieved by the following:

- Building MCLAG pairs. All links are enabled and forward traffic without loops, reducing the convergence times in case a link fails. To protect the MCLAG configuration, Fortinet suggests using two links between the FortiSwitch units and enabling *MCLAG split interface* on the MCLAG pair.
- Traffic load balancing. By default, FortiSwitch units load balance the traffic among all available uplinks. The default algorithm is based on the source-destination IP addresses, but it can be modified for other supported methods. Load balancing will work alongside the FGSP.
- Configuring a full mesh between this first layer of FortiSwitch units and the FortiGate devices, regardless of the HA configuration, results in the maximum bandwidth and minimum convergence time.
- Uplinks could be also formed by the aggregation of several links (link aggregation group [LAG]), which will increase the capacity and resiliency.

Tier-2/aggregation layer resiliency

At this layer, formed entirely by FortiSwitch units in MCLAG pairs, similar recommendations to the ones stated for the core layer are applied: enabling MCLAG pairs with link redundancy and full mesh wiring towards the core and access layers with LAGs.

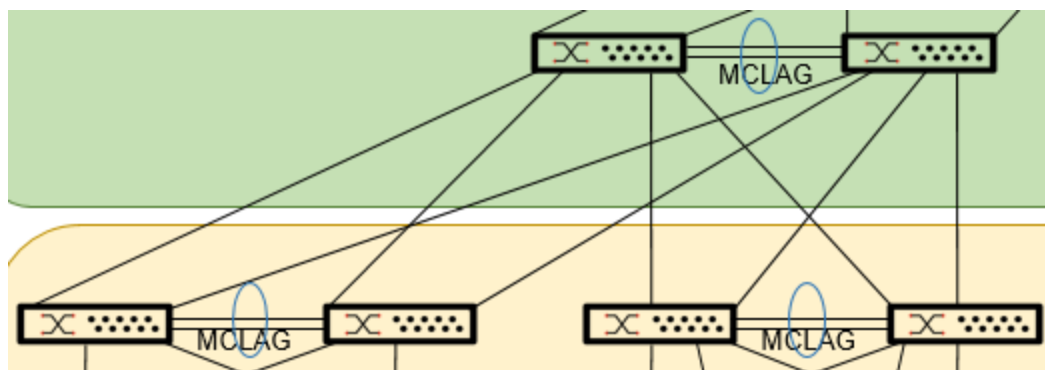
The following figure shows the redundancy mechanisms of the tier-2/aggregation layer.



Tier-3/access layer resiliency

For the first pair of FortiSwitch units forming an MCLAG pair (the tier-3 MCLAG pair), the same recommendations as in the previous two layers are applied.

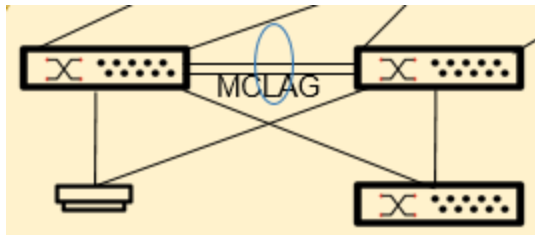
The following figure shows the redundancy mechanisms of the tier-3/access layer.



For the last meter's design:

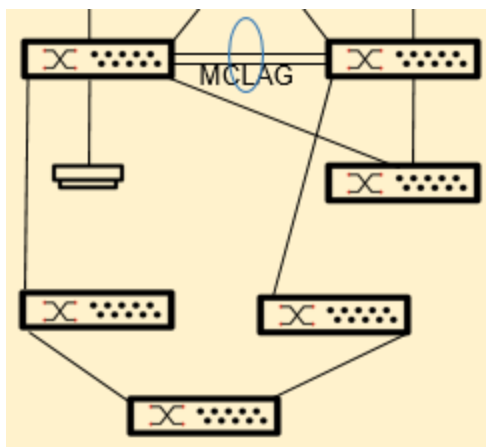
- If the last meter's design comprises a single FortiSwitch unit and single FortiAP unit, Fortinet recommends to connect them dual homed to the access MCLAG pair with single links or LAGs.

The following figure shows the last meter's redundancy for a single FortiSwitch unit and single FortiAP unit.



- If the last meter's design comprises a ring of FortiSwitch units, FortiLink automatically forms for ring topologies and takes care of the spanning tree configuration by default. The ring should have each end connected to a different access MLAG FortiSwitch unit, and the inter-switch links could comprise one or multiple interfaces (LAGs). Other mechanisms can be put in place to harden the stability and resiliency of the access layer, such as loop protection and STP BPDU guard.

The following figure shows the last meter's redundancy for a ring of FortiSwitch units.



Future proofing

The number of connected devices and the bandwidth consumption increases at a very fast pace. According to Nielsen's Law for Internet Bandwidth, Internet bandwidth is supposed to grow up to 50% year over year.

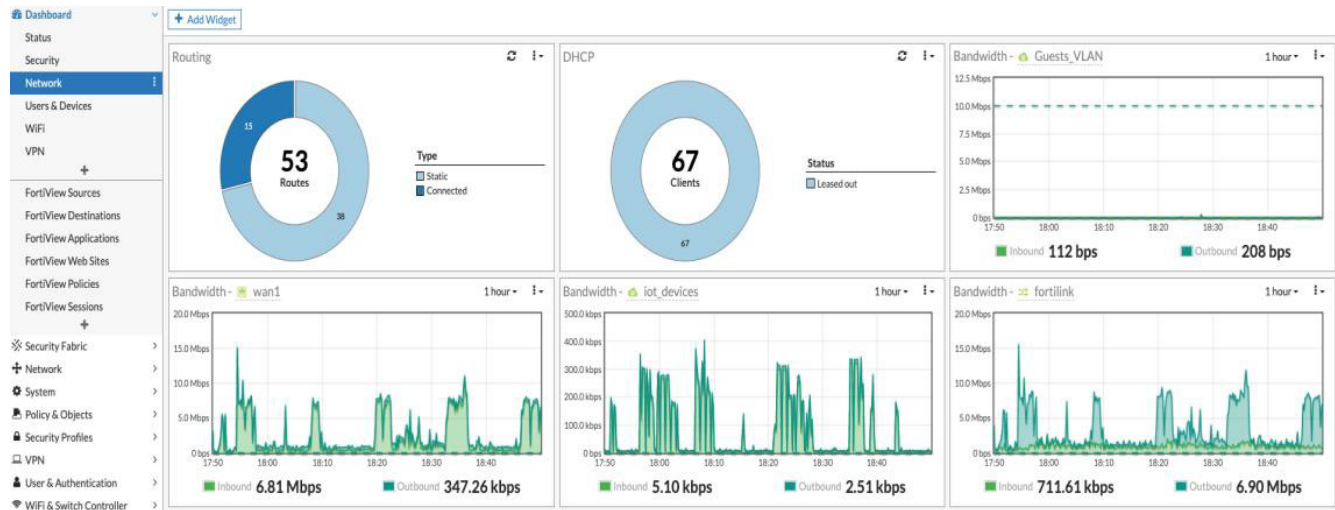
That number could be overwhelming in some situations, especially when planning the design for a network, but there are many details to be considered. Each network and its needs will vary from deployment to deployment.

When designing campus networks, there are many factors that impact the overall performance. The design must be planned to support the demands of the near future, benefit from the return of investment, and try to delay the changes needed in the near future.

- Consider the device connectivity needs for the next 5 years, such as wired devices using native 2.5 GbE or 5 GbE and the rising popularity of wireless devices.
- In the data center, server connectivity at the 25-GbE speed.
- Make it easily scalable for wired and wireless access. Adding more capacity, at least at the access layer should not be difficult with this design, nor technically complicated. With FortiAP and FortiSwitch units, adding new elements to the topology is done seamlessly and easily, thanks to the single-pane-of-glass control from the FortiGate device and FortiManager unit.

- Replacing existing infrastructure to increase performance or network capacity should not require starting from the beginning. Replacing existing FortiGate devices, FortiSwitch units, or FortiAP units is as simple as executing a few commands, and you do not have to start the configuration all over.
- The enhanced visibility obtained by the tight integration of the Fortinet Security Fabric also helps to monitor the network trends and plan in advance for the near-future needs.

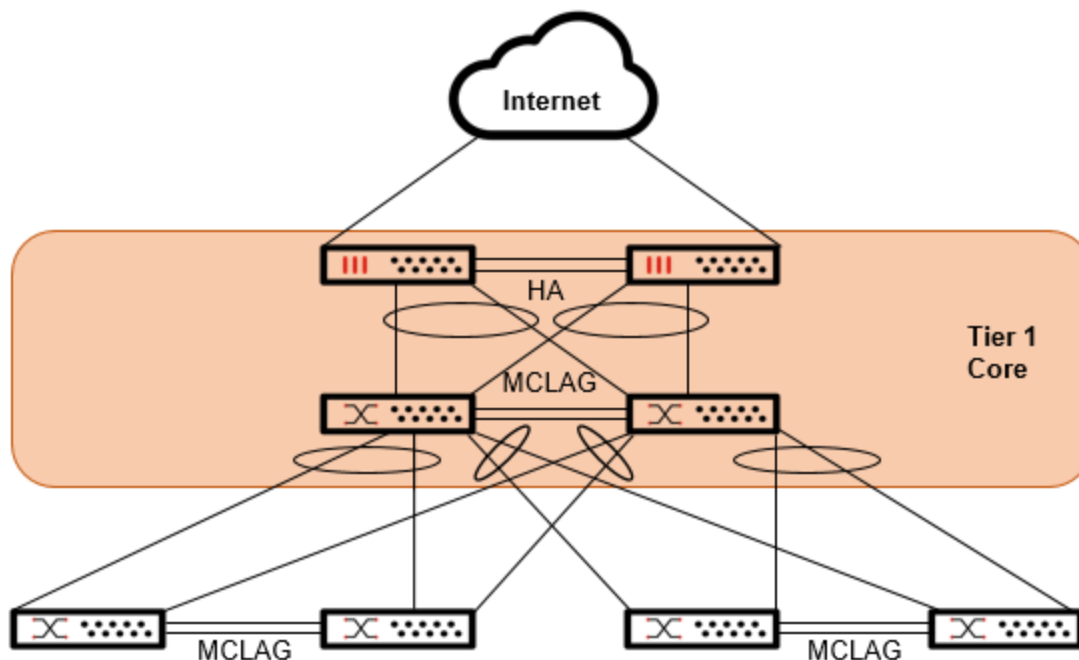
The following figure shows how you can use the FortiGate single-pane-of-glass interface to monitor network usage.



Core layer

With the Fortinet solution for integrated networking using FortiLink, the core layer always comprises a set of two to four FortiGate devices and two very high-speed FortiSwitch units, which support a large number of 100-GbE and/or 40-GbE ports with enough capacity to grow the links between them and with the aggregation layer(s). This design provides the maximum level of redundancy and resiliency towards the nonstop forwarding goal that this layer proposes to provide for the entire wired and wireless networks and with the external connectivity out of the campus. This layer of redundancy actually reduces the network complexity, especially when networks have three or more aggregation switch pairs (including the data center switches because these are usually considered as part of the aggregation layer, especially with dual-attached servers to the data center Top-of-the-Rack switches). Without a set of core switches for n aggregation switches, the redundant links to fully provide a mesh between all aggregation switches would equal $[n \times (n-1)]/2$. On the other end, with a set of two core switches in between, it is only $(n \times 2) + 2$. So, if $n=6$, you must provision for 14 links instead of 15. Each aggregation switch requires 5x100-GbE links to the other switches. With the use of a core layer, each aggregation switch only needs 2x100-GbE links, and the core layer is the only place where you need large numbers of 100-GbE ports. For example, if you have $n=10$, then you have 22 links instead of 45. In a large campus deployment, it is not practical to run that many optical fibers between buildings.

The core layer is critical, yet very simple to design, and allows for network evolution quite easily. Point-to-point links are used between each element, and Fortinet recommends using the MCLAG and dual ICLs between the core switches. The following figure shows the fully distributed set of links meshed between the core FortiGate devices, core FortiSwitch units, and the next aggregation layer of FortiSwitch units.



The FortiGate devices in the core layer can use FGCP in active-passive mode with two to four firewalls or in active-active mode for increased performance through HA load-balancing. FGCP provides failover protection, meaning that a cluster can provide FortiGate services even when one of the devices in the cluster encounters a problem that would result in the complete loss of connectivity for a stand-alone FortiGate unit (no HA). Failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in mission-critical environments. FGCP supports failover protection in four ways:

- If a link fails
- If a device loses power
- If a solid-state drive (SSD) fails
- If memory use exceeds the threshold for a specified amount of time

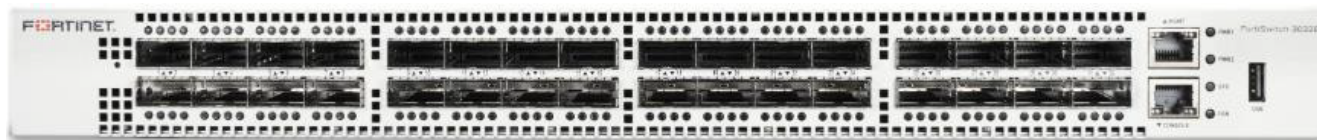
When session-pickup is enabled in the HA settings, existing TCP sessions are kept, and users on the network are not impacted by downtime because the traffic can be passed without reestablishing the sessions.

From a spanning tree standpoint (when MCLAG is used in both layers), the core-to-aggregation layer looks like a single link, removing all loops in the topology. This prevents failures at the switch or link level to cause a reconvergence and some unavailability time.

Core layer platforms

The core layer comprises a combination of two FortiSwitch units and up to four FortiGate devices. The most appropriate FortiSwitch unit to form the core layer must have many 100 gigabit Ethernet ports to address the aggregation layer and distribute a few 100-GbE ports towards the core FortiGate devices.

The following figure shows an FS-3032E core-layer switch.



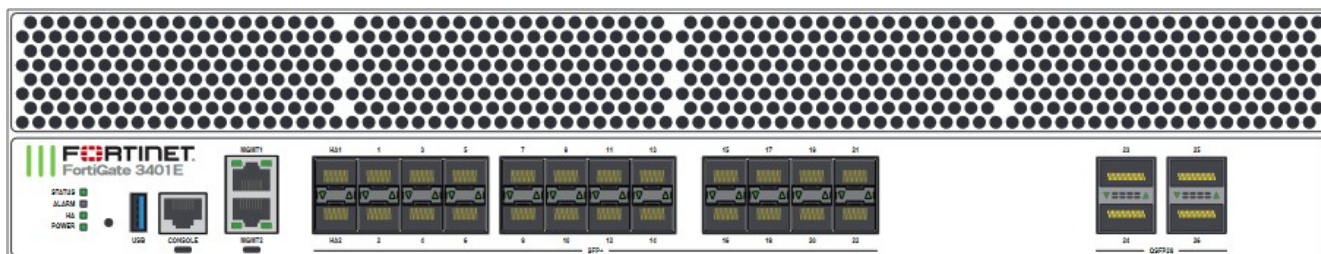
The FortiSwitch model FS-3032E with 32x100-GbE QSFP28 ports ensures that many aggregation switches can connect simultaneously with little to no possibility of oversubscription and has enough room for multiple 100-GbE ports towards the firewall. The capacity of 6.4 Tbps and its integrated dual AC power supplies with field-replaceable fans in a one rack unit form factor provide the necessary characteristics to handle the fully redundant switching requirements of this layer.

When it comes to the core FortiGate devices, the choice depends on the type of traffic and if there is a requirement to handle all requests on the campus to be secured and filtered at the core level or if some of the intranet traffic can be allowed between devices of the same VLAN directly. The capacity to handle the number of switches and access points as a fabric controller should be considered. The price-to-performance ratio has also been taken into account in recommending a limited number of choices for you with two levels: price optimized and performance optimized.

The *price-optimized* option consists of a pair of FG-3400E FortiGate devices. With 4x100-GbE QSFP28 slots, it provides enough capacity to directly connect to the two core FortiSwitch units and still allow for expansion because it also has 24x25-GbE ports (with two HA ports for communication between the FortiGate devices). Using the security processing units (SPUs) NP6 and CP9, it accelerates in hardware many of the security and networking protocols (like CAPWAP) to achieve 240 Gbps in firewall capacity, 44 Gbps in the intrusion prevention system (IPS), 34 Gbps in NGFW, and 25 Gbps in threat protection. It supports control using FortiLink of 300 FortiSwitch units, terminates CAPWAP at 57 Gbps, and also allows control of 4,096 WiFi access points in bridge mode or 2,048 APs in tunnel mode. 2,048 access points can

cover around 400,000 square meters of open-space offices, while 4,096 APs can cover up to 800,000 square meters. Based on the International Labour Organization (ILO) recommendations of 6 square meter minimum per employee, a campus of between 6,500 and 13,000 employees can be covered with such AP and switching capacity at the core.

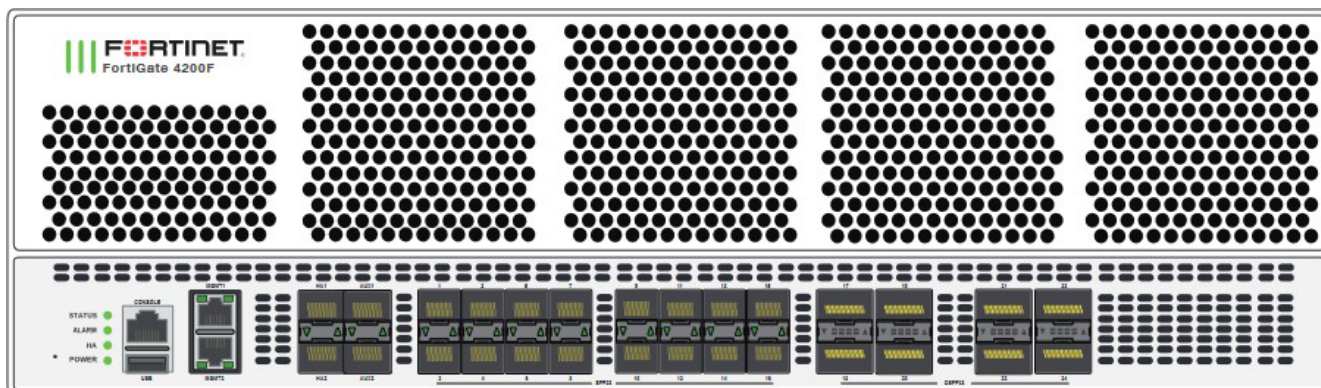
The following figure shows the FG-3401E, which provides a core firewall, secure wireless, and switch controller.



You can also use virtual machine FortiGate models like FG-VM16, FG-VM32, or FG-VMUL and their “V” versions to achieve the same number of switches and APs supported in the core layer. The requirement of several 100-GbE links towards the core FortiSwitch units on top of the performance requirements to handle that amount of traffic would probably undermine the VM options in this secure campus scenario as a core platform.

The *performance-optimized* choice consists of a pair of FG-4200F units. This design can even grow up to four FortiGate devices at the core level to process even more traffic of next-generation firewalls in parallel. With 8x100-GbE QSFP28 slots per FortiGate unit, it provides enough capacity to directly connect with 2x100-GbE ports to each of the two core FortiSwitch units at a nonstop forwarding capacity of up to 800-Gbps firewall performance using NP7 and CP9 hardware accelerated security processing. Each unit has an additional 18x25-GbE/10-GbE ports and is capable of 52 Gbps of IPS, 47 Gbps in NGFW, and 45 Gbps in threat protection. 300 FortiSwitch units can be controlled as a network controller, as well as 8,192 access points in bridge mode or 4,096 APs in tunnel mode for traffic of up to 47 Gbps in CAPWAP. As calculated previously, this covers the need of 13,000 to 26,000 employees in the campus with such capacity at the core layer.

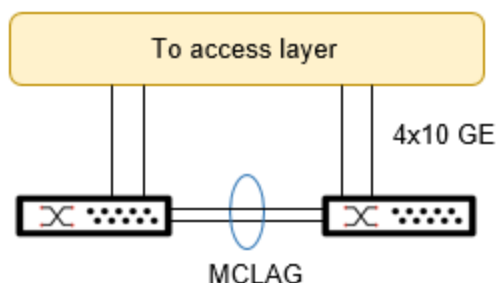
The following figure shows an FG-4200F, which provides a core firewall, secure wireless, and switch controller.



Aggregation layer

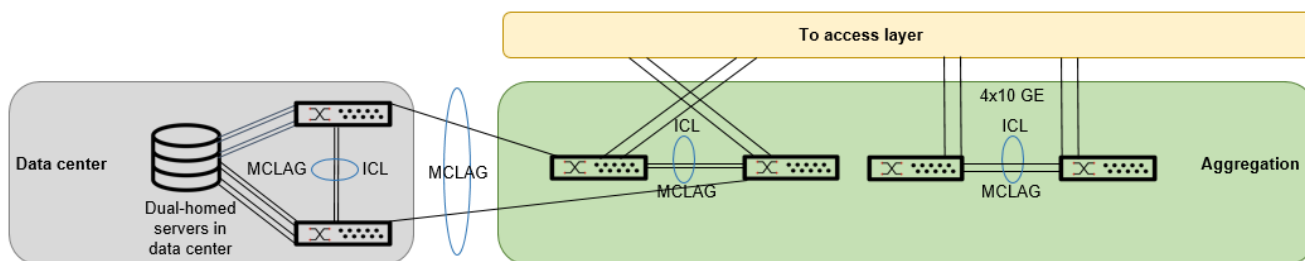
The aggregation (sometimes also called distribution) layer is a real crossroad. Its primary goal is to increase network scalability by providing a single place to interconnect multiple access switches and the core layer. It facilitates the connectivity because it would rapidly become impractical to interconnect all access switches in a full mesh of links without relying on an aggregation point for these multiple access-layer switches. The potential geographic distribution of access switches across many buildings in a larger campus would also require more fiber optics to interconnect if the aggregation layer was not there. An aggregation layer usually comprises a few blocks of two switches in MCLAG. By design, it therefore provides resiliency because it will always be deployed in pairs of switches and comes with a recommendation to deploy only dual hot swappable power supplies and redundant fans in each switch to augment reliability.

The following figure shows an aggregation-layer building block.



Multiple blocks of pairs of aggregation switches extend the design of this key layer if there are more than 24 floors or buildings in the campus. This layer is also where data center services are provided. Even though it is not going to be used for the same purpose, it is common that the connection of dual-attached servers, storage, and other network-based services connect at this level or have their own block of switches that connect to the high-speed core switches in the same manner in the main intermediate distribution frame (IDF). This main IDF usually also houses the core block. Even if the layers are collocated there, it is important to recognize the role of each layer in this three-tier hierarchical model to ensure scale and reliability and to limit human errors and malicious attacks. This model allows the aggregation switches to easily accommodate thousands of devices passing through this layer while simplifying the design, maintenance, and operations.

The following figure shows the aggregation-layer design, including the data center building block.



Fortinet recommends that no access devices (including wireless access points, surveillance cameras, IP phones, or laptops) are connected to this layer.

Aggregation layer platforms

The most appropriate FortiSwitch unit to form the aggregation layer comprises many 10/40 gigabit Ethernet ports to address the access layer and a few 100-GbE ports towards the core layer.

The following figure shows an FS-1048E aggregation-layer switch.



This is exactly what the FS-1048E provides: 48x GbE/10-GbE SFP+ ports and 6x40-GbE QSFP+ ports or 4x100-GE QSFP28 ports, 1760-Gbps switching capacity in a 1 RU rack-mounted form factor. A 24x10-GbE ports version also exists if fewer ports are required towards the access layer. Having 4x100-GbE ports allows for two ports to go to the core switches and two ports to connect the aggregation layer in MCLAG together (ICL) at very high speed. Those links can still run at 40 Gbps or 25 Gbps, depending on the fiber/SFP combination used if oversubscription in the case of the transient failure of one uplink to the core is not an issue.

Access layer

The access layer is where endpoints (such as phones, laptops, video-conferencing sets, printers, IoT sensors, IP cameras, and servers) are primarily connecting to the network. Wireless access points are also connected here and provide further access. FortiSwitch units distribute the ports to plugs distributed through the ceilings, the floor, the walls, or the desks from IDFs (or technical cabinets) spread throughout the floors. The IDF can also be at a central location, depending on the size of the floor due to the 100-meter copper cable limit. Through the use of wireless access points, a single Ethernet port can distribute the access to tens or hundreds of access devices in this layer. Wired connectivity, though, is usually provided to a single endpoint per gigabit (or more) Ethernet port, except if a PC is reaching the network through an IP phone with multiple GigE ports for example. Dual-attached endpoints like servers are typically not connected here but to the aggregation layer. Redundant uplinks from the access layer to the aggregation layer ensure security and resiliency for the entire network. The access layer acts as a collection point for high-performance wired and wireless devices and must have enough capacity to support the power and bandwidth needs of today as well as to scale for the future while the number of devices grows.

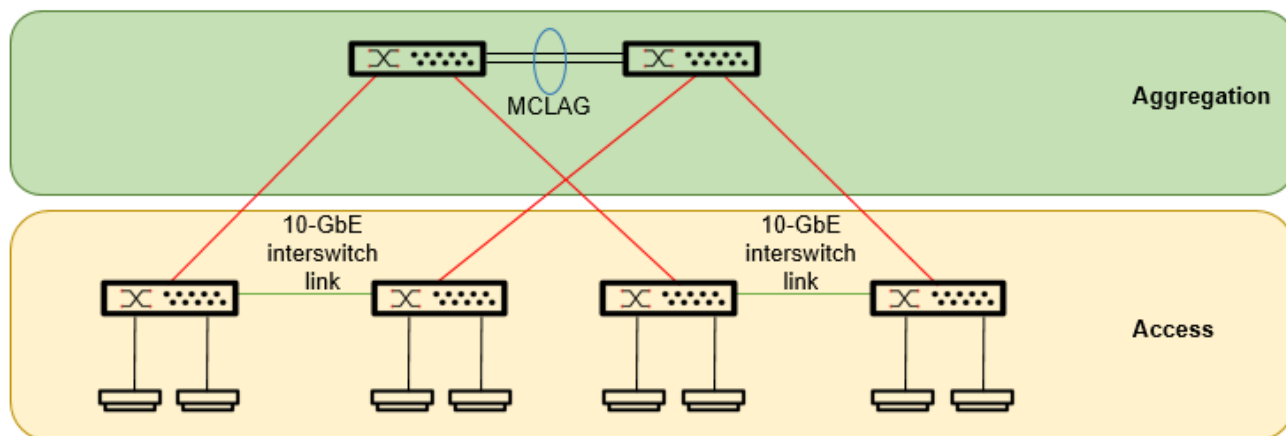
Access-layer deployment recommendations

Depicting what the access layer should look like is difficult because it depends on several physical and logical factors that make up the campus. These examples are offered as guidance for building a multi-tiered network supporting all the aforementioned design principles. You will need to adapt them to your specific environment and layout.

These network design proposals are for a network implemented with FortiSwitch units and FortiAP units. A FortiGate device (or a FortiManager unit) is the single-pane-of-glass network manager and controller for the FortiSwitch and FortiAP units.

This section considers use cases for a potentially large campus. Taking into consideration the assumptions made in the designing principles, the initial setup will use 4x10-GbE links between the access layer and the aggregation layer.

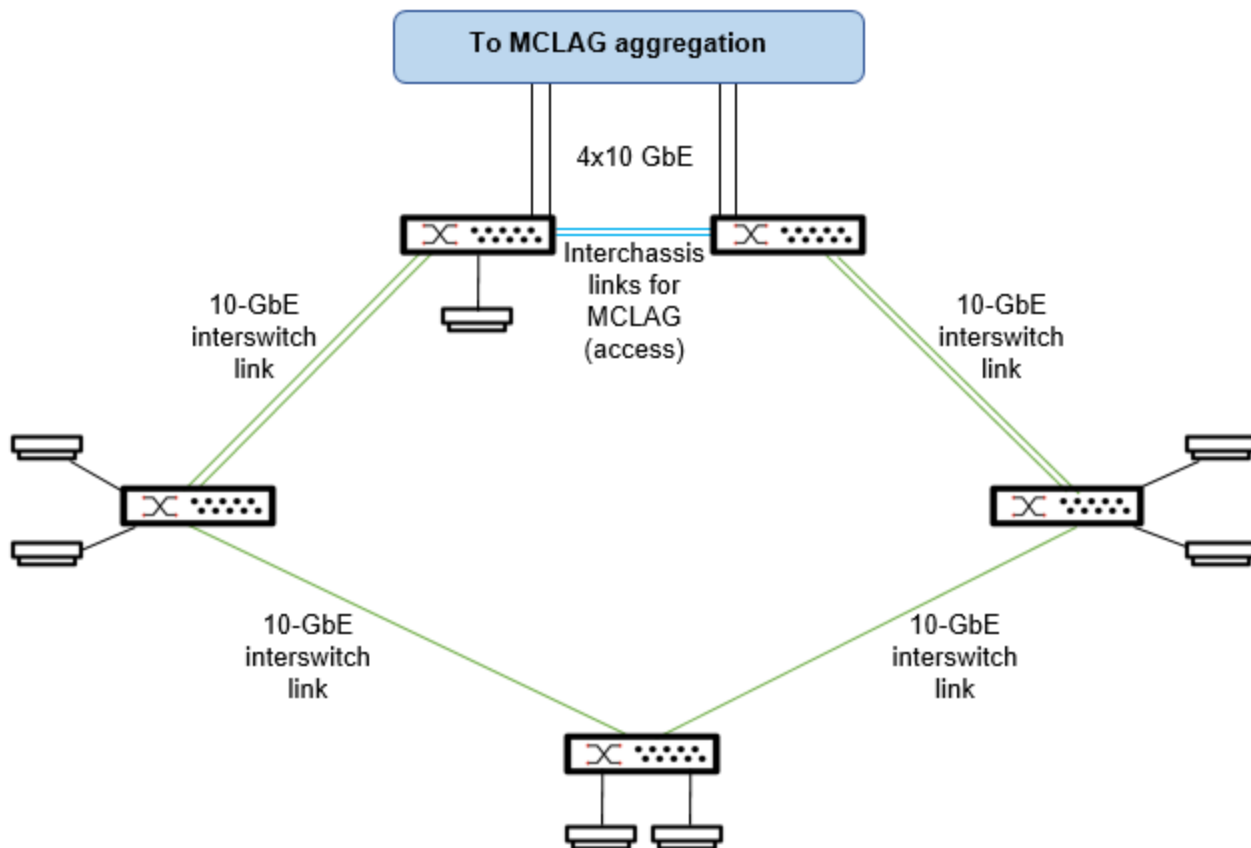
The following figure shows a typical deployment of a single floor with up to four switches and the connection of the access layer with the aggregation layer.



Each of the FortiSwitch units is managed using FortiLink by a FortiGate device in the core layer as part of the Fabric and connects with a 10-GbE aggregation link. To provide link redundancy, two sets of two switches have an inter-switch link to their closest FortiSwitch neighbor with a second 10-GbE link, forming two distinct loops. If any of the red or green links in the figure fail, the connectivity is still active, and the oversubscription ratio of 10:1 is respected as 96-GbE ports are uplinked with 10 Gbps.

In a scenario where the floor requires five switches or more, redundancy becomes even more critical. To address this requirement, the MCLAG is extended down to the access layer by using a pair of high-end switches to create redundant uplinks to the aggregation layer. Enough bandwidth is left to the additional access switches on the floor through multiple 10-GbE links to fuel the ring behind.

The following figure shows the typical deployment of a single floor with five switches and single-attached wireless access points.



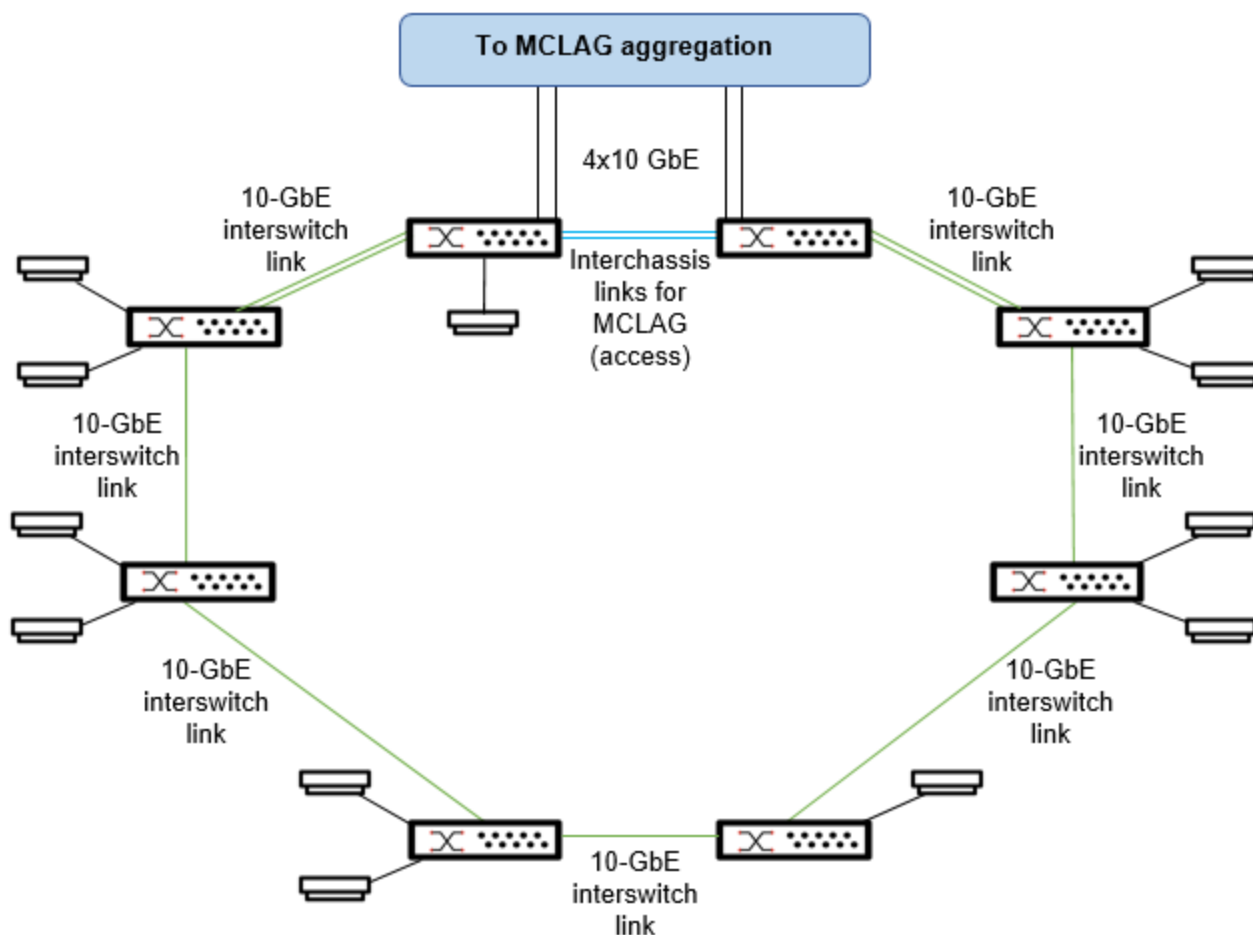
This design creates a full loop between the access and aggregation layers. Both layers need to implement MCLAG to be able to use all links simultaneously, to not rely on the STP to block ports, and to add delay when needing to converge. MCLAG must be supported for this design even at the access layer, which calls for FortiSwitch 400E and 500D Series models acting as ToR switches.

This design can also be applied with two, three, or four switches per floor, especially because it would help to make the network much more agnostic to the spanning tree already at the access layer.

Both FortiSwitch units managed with FortiLink acting as ToR in the access layer have 2x10-GbE ports directly toward the aggregation layer (each 10-GbE link to a different MCLAG unit on each side). The remaining switches (not connected directly) have at least one 10-GbE link going each to the closest two switches, creating a ring as seen in the preceding figure. You can double (LAG) the green links between the first switch in the ring and the second one to provide more

throughput overall (especially in case of a transient outage) and more resiliency. Using this design, you can go up to eight switches and never need more than 4x10-GbE ports per switch to interconnect other access-layer switches or the aggregation layer. Having eight switches on a floor is compatible with wired and wireless access for a floor of 3,200 square meters (approximately 56 m x 56 m) in an open-space environment. If more than this is needed, then you can replicate the design in another IDF to still stay within the 10:1 oversubscription ratio (and with the 100-m copper cable limitations).

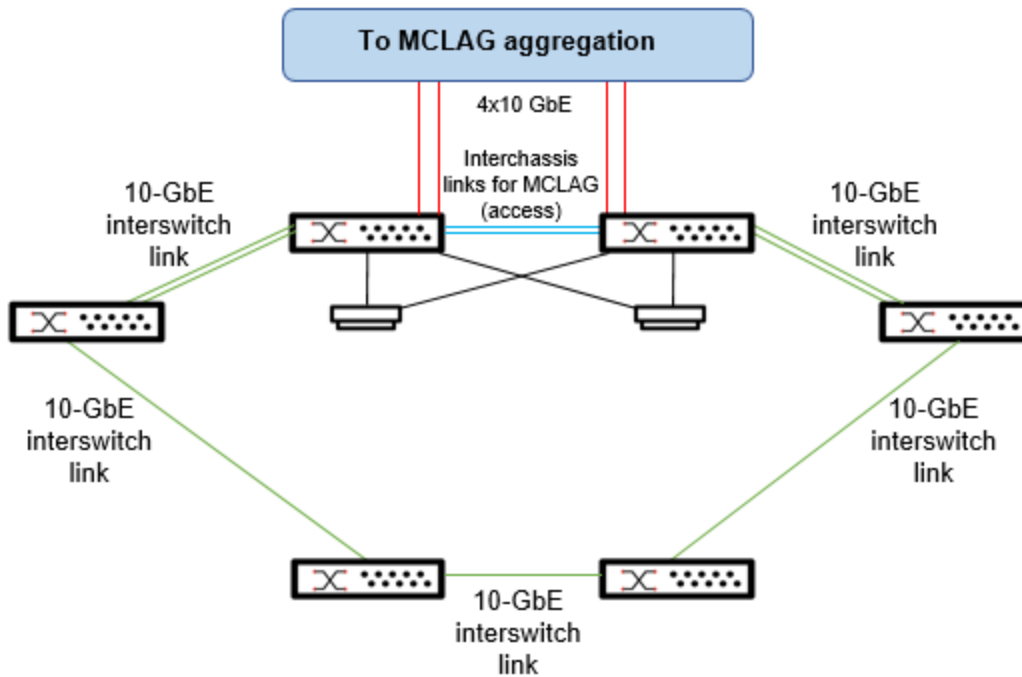
The following figure shows a typical deployment of a single floor with eight switches and single-attached wireless access points.



To extend the design to a floor with nine switches or more (modulo 8 each time), Fortinet recommends splitting the floor into sets of the previous designs with up to eight access switches per group. This design protects from oversubscription because with 8x48-GbE ports and the 10:1 oversubscription ratio, you need 38.4 Gbps on the uplink; with 9x48-GbE ports, you need 43.2 Gbps, exceeding the 40-Gbps maximum available with the 4x10-Gbps uplinks.

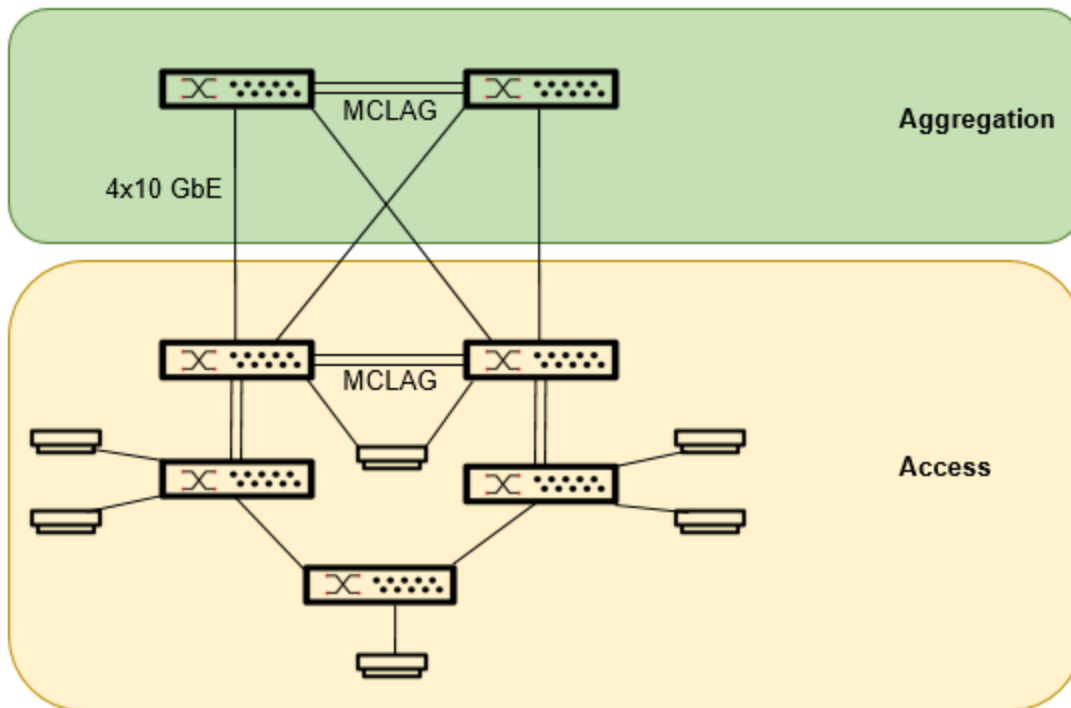
In the previous designs, access points are connected to a GbE or mGig port on a single switch. The mGig ports would also be within the expected limits with 2.5 Gbps per access point.

In an even higher availability scenario where the wireless is critical and losing wireless connectivity of an access point is not possible, the design scenario for five to eight switches is still fully adapted because it allows for dual attachment of the APs to the ToR MCLAG switches, as shown in the following figure.



In the preceding figure, each AP is dual-attached to the MLAG switches, between which the blue interchassis link can even be doubled for extra redundancy. If oversubscription occurs in a normal situation, the traffic of 384 GbE ports is forwarded towards the aggregation layer using 40 Gbps, which is in the 10:1 ratio. Losing one of the switches connected to the aggregation layer directly removes 48 ports to be uplinked, which roughly provides a 15:1 ratio. In the worst-case scenario of one of the red links failing but not the switch, the network still adheres to the 20:1 ratio in a temporary outage. If you want more than four redundant access switches, you need to deploy new sets of 10-GbE uplinks to the aggregation layer to respect the oversubscription recommendations.

The following figure shows how to interconnect 10-GbE uplinks between the access layer and aggregation layer with MCLAG at both layers for resiliency.



Access layer platforms

Tightly integrated into the Fortinet Security Fabric with FortiLink, FortiSwitch units can be managed directly from the FortiGate interface and through FortiManager, offering a single pane of glass that you can use to control and deploy the campus. This network topology offers next-generation firewalls with advanced security, core routing, switching, and wireless. This single-pane-of-glass management provides complete visibility and control of users and devices on the network, regardless of how they connect.

This network topology also includes the detection and control of IoT devices to ensure they are not compromised. IoT devices represent a serious risk to networks because they are built more for function than security. IoT devices are beginning to multiply tremendously in all connected buildings and organizations.

From an equipment basis as well as long-term contracts and services, this secure, simple, and scalable campus solution offers simplified operations, increased value for your money, and optimized total cost of ownership (TCO).

At the access layer, several options are provided, primarily depending on the level of feature richness and uplink choices. This design minimizes the number of different products to make maintenance, troubleshooting, and operations simpler for IT, security, and networking staff.

The following figure shows an FS-148F-FPOE access-layer switch.



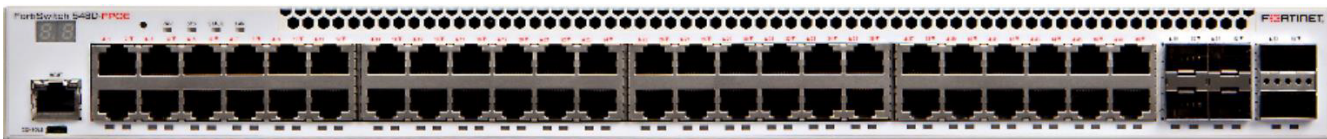
The 100F Series switches are featured entry-level access switches that provide full PoE options for 24 and 48 gigabits Ethernet ports, with 4x10-Gb Ethernet SFP+ uplinks. An eight-port version also exists with 2-GbE SFP uplinks for more complex architectures. Because this switch series does not support the MCLAG, these switches are pure access switches, which are recommended for single-attached end devices. The switches themselves can be dual-homed to Top of the Rack or aggregation switches, and they can allow for multiple links to be aggregated towards the same destination and to create rings. The 100F Series switches should not be used when access devices need to be attached to two different access switches; instead, you can use 400E Series switches or 500D Series switches.

The following figure shows an FS-M426E-FPOE access-layer MCLAG-capable switch.



The 400E Series switches are enterprise switches with full functionality, including MCLAG, and come in 24-GbE and 48-GbE ports versions with full PoE options. The FortiSwitch M426E-FPOE unit has 16-GbE FPoE plus 8x2.5-GbE FPoE ports plus 2x5-GbE ports with 4x10-GbE SFP+ uplinks, allowing for powering access points and providing multi-gigabit connectivity.

The following figure shows an FS-548D-FPOE access-layer MCLAG-capable switch.



The 500D Series switches are high-end access switches with 4x10-G SFP+ ports uplinks (including the hardware support of MACsec), and 2x40-G links. If needed, these ports allow future expansion of the uplinks toward the aggregation layer, potentially dividing the uplink ratio by 2 (80 Gbps instead of 40 Gbps). Expanding the uplinks would mean that the aggregation layer has enough 40-GbE ports on the other end and would mean an uplift from 1000E Series to 3000E Series FortiSwitch units.

Management

The FortiManager unit is the centralized management platform for many Fortinet products. Not only does it provide a single pane to manage switches, access points, security appliances, and so on, but it also helps with compliance reporting and network automation for secured wireless and campus switching deployments.

The average enterprise today has too many networking and security point products. As a result, you probably have many management consoles for security operations and others for network operations. Having different tools for managing security products increases risk, inefficiency, and cost—both in terms of operational and capital expenditures. Having a central management for many of the Fortinet security and network platforms, as well as Security Fabric-partner products, allows you to unify and efficiently manage configurations and policies across different devices. The open and extensible architecture of fabric connectors also enables an “integrated approach” to securing the enterprise and the sharing of threat intelligence with all fabric elements that you might already have. Device life-cycle management of network and security products enable scripted and automated provisioning, firewall AP and switch updates, migration, and deprovisioning easily.

The following figure shows the FortiSwitch Manager GUI.

FortiSwitch Name	Serial Number	Platform	Status	FortiLink	FortiGate	Connected Via	OS Version	Join Time	Comments	Template
tier-1-1	FS1E4BT419000051	FortiSwitch-104BE	Online	fortilink	.FGT500E-1[root]	172.16.16.1	FS1E4B-v7.0.2-build049.210823 (GA)	Wed Dec 1 16:06:41 2C		
FS1E4BT419000052	FS1E4BT419000052	FortiSwitch-104BE	Online	fortilink	.FGT500E-1[root]	172.16.16.7	FS1E4B-v7.0.2-build049.210823 (GA)	Wed Dec 1 16:08:16 2C		
S426EFTF19000243	S426EFTF19000243	FortiSwitch-M426E-F	Online	WAN-fortilink	.FGT500E-1[root]	172.16.17.3	S426EF-v7.0.1-build038.210616 (GA)	Wed Dec 1 16:07:14 2C		
S426EFTF19000296	S426EFTF19000296	FortiSwitch-M426E-F	Online	WAN-fortilink	.FGT500E-1[root]	172.16.17.4	S426EF-v7.0.1-build038.210616 (GA)	Wed Dec 1 16:06:40 2C		
S448DNTF18001121	S448DNTF18001121	FortiSwitch-448D	Online	fortilink	.FGT500E-1[root]	172.16.16.4	S448DN-v7.0.2-build049.210823 (GA)	Wed Dec 1 16:07:18 2C		
S448DNTF18002336	S448DNTF18002336	FortiSwitch-448D	Online	WAN-fortilink	.FGT500E-1[root]	172.16.17.2	S448DN-v7.0.1-build038.210616 (GA)	Wed Dec 1 16:07:20 2C		
S448DNTF18002391	S448DNTF18002391	FortiSwitch-448D	Online	fortilink	.FGT500E-1[root]	172.16.16.3	S448DN-v7.0.2-build049.210823 (GA)	Wed Dec 1 16:07:21 2C		
FS1D243Z14000285	FS1D243Z14000285	FortiSwitch-1024D	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.4	FS1D24-v7.0.3-build058.211130 (GA)	Thu Jan 13 16:03:12 20		Import
FS1D243Z14000301	FS1D243Z14000301	FortiSwitch-1024D	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.3	FS1D24-v7.0.3-build058.211130 (GA)	Thu Jan 13 16:04:18 20		Import
S108EP5918000293	S108EP5918000293	FortiSwitch-108E-PO	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.6	S108EP-v7.0.3-build058.211130 (GA)	Thu Jan 13 16:06:35 20		
S124DP3X16000413	S124DP3X16000413	FortiSwitch-124D-PC	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.12	S124DP-v6.2.0-build168.190415 (GA)	Thu Jan 13 16:01:58 20		
S124DP3X16000522	S124DP3X16000522	FortiSwitch-124D-PC	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.10	S124DP-v3.6.11-build432.191108 (GA)	Thu Jan 13 16:01:22 20		
S124EF4N17000100	S124EF4N17000100	FortiSwitch-124E-FP	Offline	fortilink	.FGT100F-FSW-1[root]					
S124EF4N17000106	S124EF4N17000106	FortiSwitch-124E-FP	Offline	port1	.FGT100F-FSW-1[root]	20.254.0.100	S124EF-v7.0.0-build022.210415 (GA)	Thu Jan 13 14:10:24 20		
S224DF3X16000207	S224DF3X16000207	FortiSwitch-224D-FP	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.9	S224DF-v6.4.7-build478.210503 (GA)	Thu Jan 13 16:01:52 20		
S224DF3X16000218	S224DF3X16000218	FortiSwitch-224D-FP	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.7	S224DF-v6.4.7-build478.210503 (GA)	Thu Jan 13 16:03:05 20		
S224DF3X16000224	S224DF3X16000224	FortiSwitch-224D-FP	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.8	S224DF-v6.4.7-build478.210503 (GA)	Thu Jan 13 16:01:59 20		
S224ENTF18000134	S224ENTF18000134	FortiSwitch-224E	Offline	fortilink	.FGT100F-FSW-1[root]	169.254.1.13	S224EN-v6.4.7-build478.210503 (GA)	Wed Nov 17 15:42:35 2		
S224ENTF18000118	S224ENTF18000118	FortiSwitch-224E-PO	Offline	fortilink	.FGT100F-FSW-1[root]	169.254.1.11	S224EN-v6.4.6-build470.210211 (GA)	Wed Nov 17 15:40:32 2		
S424DP3X17000351	S424DP3X17000351	FortiSwitch-424D-PC	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.5	S424DP-v7.0.2-build049.210823 (GA)	Thu Jan 13 16:02:01 20		Import
S424DP3X17000388	S424DP3X17000388	FortiSwitch-424D-PC	Online	fortilink	.FGT100F-FSW-1[root]	169.254.1.2	S424DP-v6.4.7-build478.210503 (GA)	Thu Jan 13 16:03:17 20		
SR12DPDT18000128	SR12DPDT18000128	FortiSwitchRugged-1	Offline	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.2	SR12DP-v7.0.1-build038.210616 (GA)	Fri Dec 31 09:20:29 20		Import
SR12DPDT20000166	SR12DPDT20000166	FortiSwitchRugged-1	Offline	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.3	SR12DP-v7.0.1-build038.210616 (GA)	Fri Dec 31 09:22:16 20		Import
SR12DPDT20000169	SR12DPDT20000169	FortiSwitchRugged-1	Offline	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.4	SR12DP-v7.0.1-build038.210616 (GA)	Fri Dec 31 09:18:53 20		Import
SR24DN4417000011	SR24DN4417000011	FortiSwitchRugged-1	Online	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.6	S24DN-v7.0.1-build038.210616 (GA)	Thu Jan 13 16:00:38 20		Import
SR24DN4417000028	SR24DN4417000028	FortiSwitchRugged-1	Offline	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.5	S24DN-v7.0.1-build038.210616 (GA)	Fri Dec 31 09:22:51 20		Import
SR24DN4417000043	SR24DN4417000043	FortiSwitchRugged-1	Online	OT_FSR	.FGT100F-FSW-1[root]	169.254.2.7	S24DN-v7.0.3-build058.211130 (GA)	Thu Jan 13 16:07:53 20		Import

Based on Fortinet’s extensive experience in the security industry, Fortinet has integrated the most important security “best practices” into a series of tests that can be run on Fortinet technology. Consequently, you now have a quantifiable measure of your security posture. By charting your security scores over time, the FortiManager unit helps you easily see if your security posture is improving or declining, discern trends, and compare your security score with your peers. The FortiManager unit, working with the FortiAnalyzer unit, tracks and stores security ratings over time.

The FortiManager unit reduces administration and workload costs with smart features, such as device discovery, device group creation by administration domain, audit, and management of complex architectures. You can quickly deploy hundreds of switches and APs in your campus, thousands of edge locations, trigger changes to entire groups of devices, and consistently define security and SD-branch policies throughout your campus and branch environments.

The FortiManager unit manages FortiGate devices and any subordinate FortiSwitch, FortiAP, and FortiExtender units and provides signature updates to FortiMail, FortiSandbox, and FortiClient units. It provides enterprise-grade high

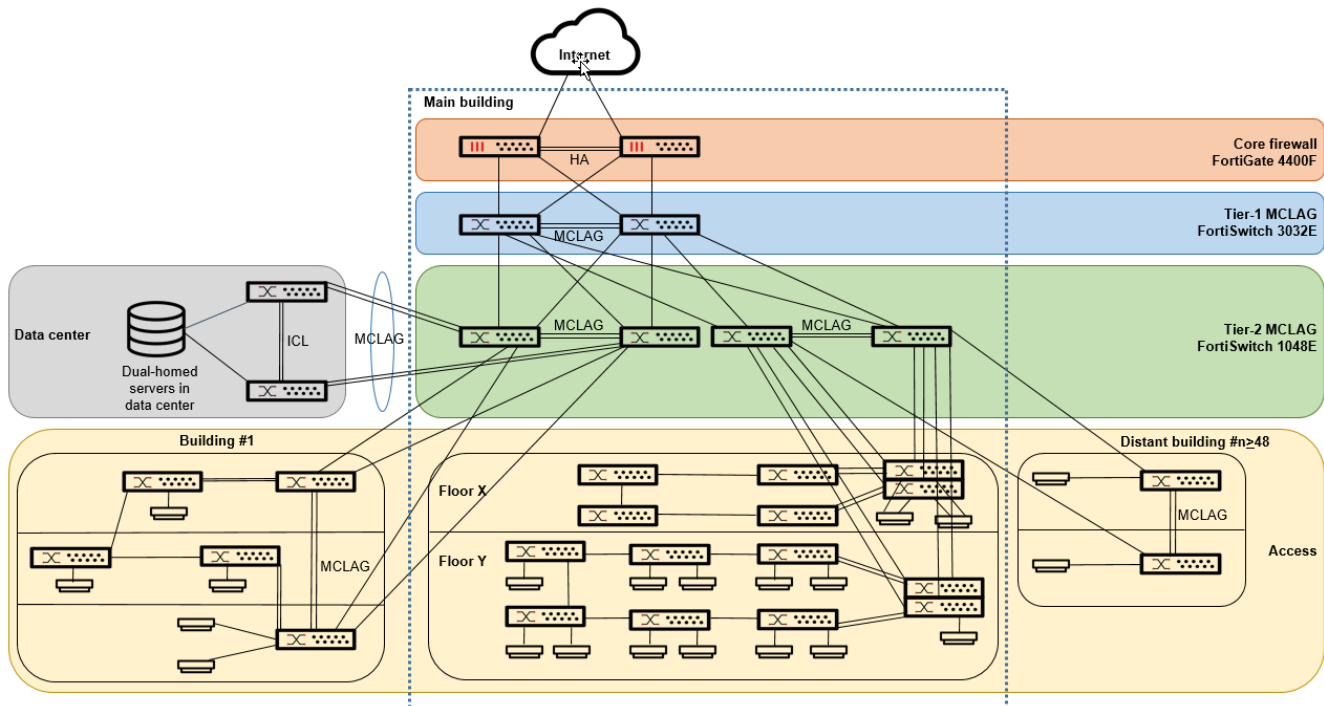
availability and integration, automating backups to up to five nodes with streamlined software and security updates for all managed devices.

Finally, the FortiManager unit facilitates rapid or automated response using FortiOS automation stitches. An automation stitch is a simple way to define actions on triggers. Think of it as “if this happens, then do that” in a single console. The FortiManager unit is making next-generation infrastructure management less painful for the MIS and IT teams with these automation-ready tools, not only removing the chance of error or omission due to manual processing, but also reducing demand on sparse resources.

The FortiManager unit can manage up to 10,000 devices and VDOMs and up to 120,000 FortiClient agents from a single FortiManager interface (depending on the model). The FortiManager unit can segregate the management of large deployments easily and securely by grouping devices and agents into geographic or functional administrative domains (ADOMs). The FortiManager product selection is based on the number of Fortinet devices it can manage and the FortiManager unit's storage capacity.

Final design

Now that all the layers and main features have been explained and depicted in detail, the following figure summarizes the overall fabric solution recommended by Fortinet for wired and wireless access on multiple floors and in multiple buildings.



If you have fewer than 48 sets of access layers, which can be distributed per floor or distant building, depending on the type of campus, you deploy a set of tier-2 MCLAG switches in the aggregation layer. If you have more than 48 sets of access layers, then you just deploy a new set of two tier-2 MCLAG switches in the aggregation layer and keep the rest of the design exactly the same in the core and access layers, up to six sets, which represents at maximum of 2,300 access FortiSwitch units or around 110,000 GbE switch ports in the campus. You would probably have to segment the network differently with many more core building blocks, even if you have fewer than the maximum number of switch ports, because you would have to account for other physical elements and WAN access pipe sizing to accommodate the amount of traffic that would be generated.

For more details on deploying a large-campus switching network, refer to the [Large Campus Switching Deployment Guide](#). You can also contact Fortinet accredited partners and distributors or Fortinet Professional Services to help you further in designing your network following these guidelines.

SD-branch architectures

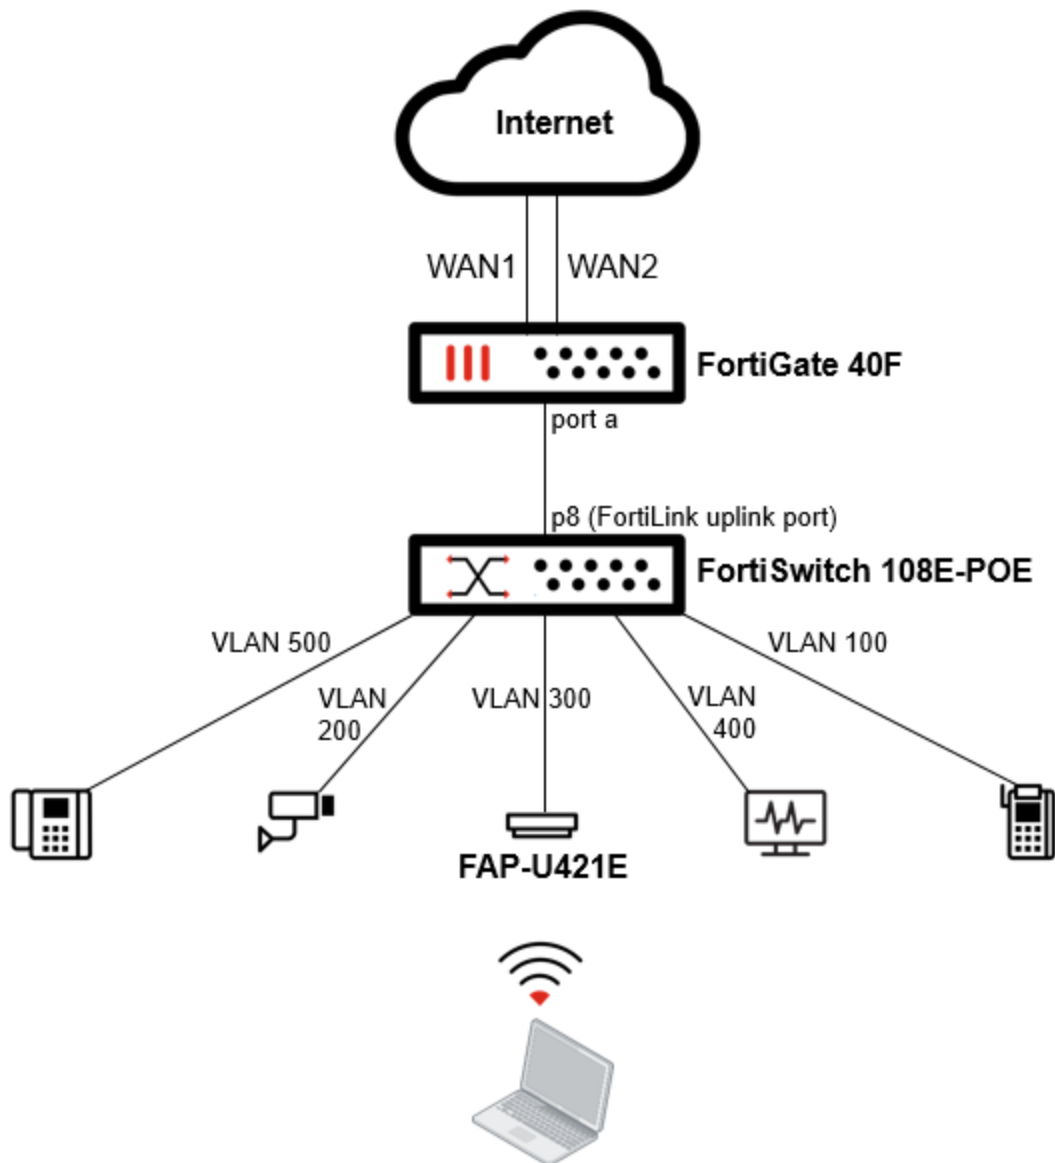
FortiSwitch units can be adopted as a natural extension of SD-WAN to provide security on the wired LAN edge.

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

In addition, 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 units facilitate and enhance network visibility as a first step in grabbing control of the network—under the umbrella of the FortiGate device with the FortiManager unit functioning as a single pane of glass.

The following figure shows a typical SD-branch architecture, where a FortiGate device manages various devices through a FortiSwitch unit. Each device is assigned a different VLAN and port.

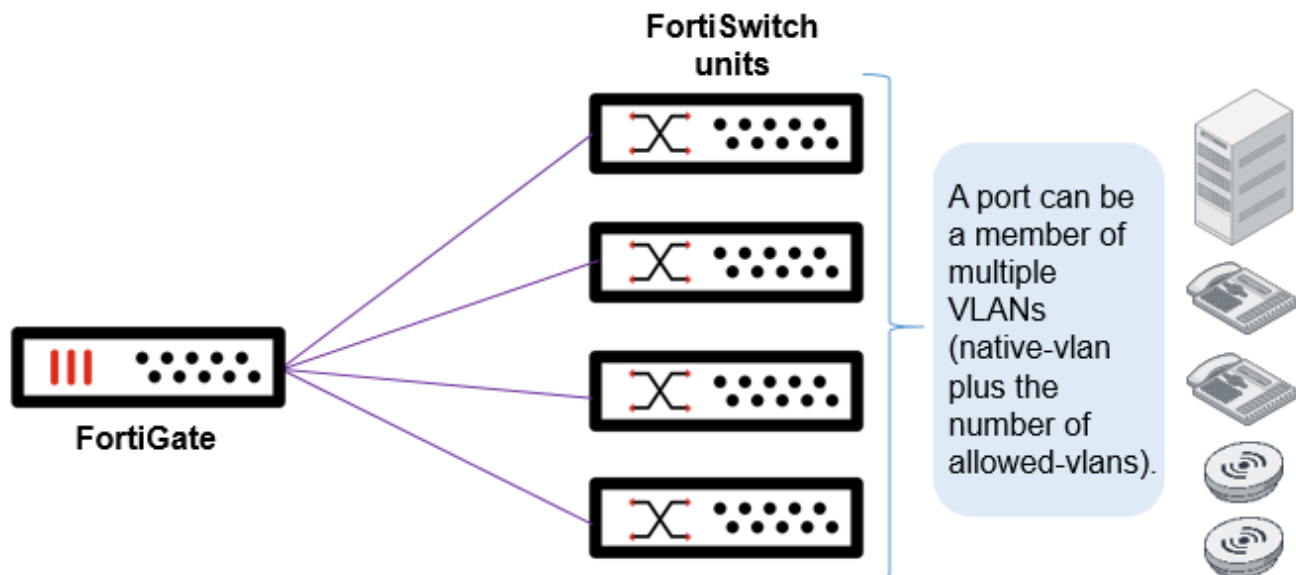


You can replace one of the WAN links from the FortiGate device to the Internet with a Long Term Evolution (LTE) connection and use it for failover if an outage occurs on the primary link. See [FortiExtender documentation](#) for further details.

Small SD-branch

In the following small SD-branch architecture, a single FortiGate device manages multiple FortiSwitch units. The FortiGate device connects directly to each FortiSwitch unit. Each of these FortiLink ports is added to the logical interface on the FortiGate unit.

Optionally, you can connect other devices to the FortiGate logical interface. These devices, which must support IEEE 802.1q VLAN tagging, will have layer-2 connectivity with the FortiSwitch ports.

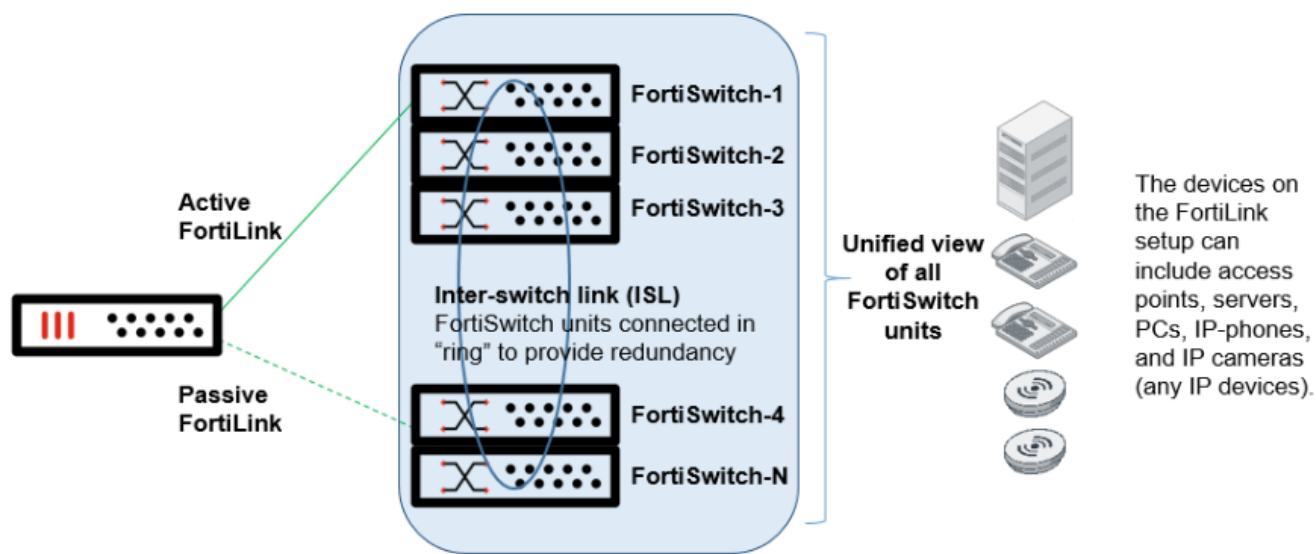


Medium SD-branch

In the following medium SD-branch architecture, a single FortiGate device manages a stack of several FortiSwitch units.

The FortiGate device connects directly to one FortiSwitch unit using a physical or aggregate interface. The remaining FortiSwitch units connect in a ring using inter-switch links (ISLs).

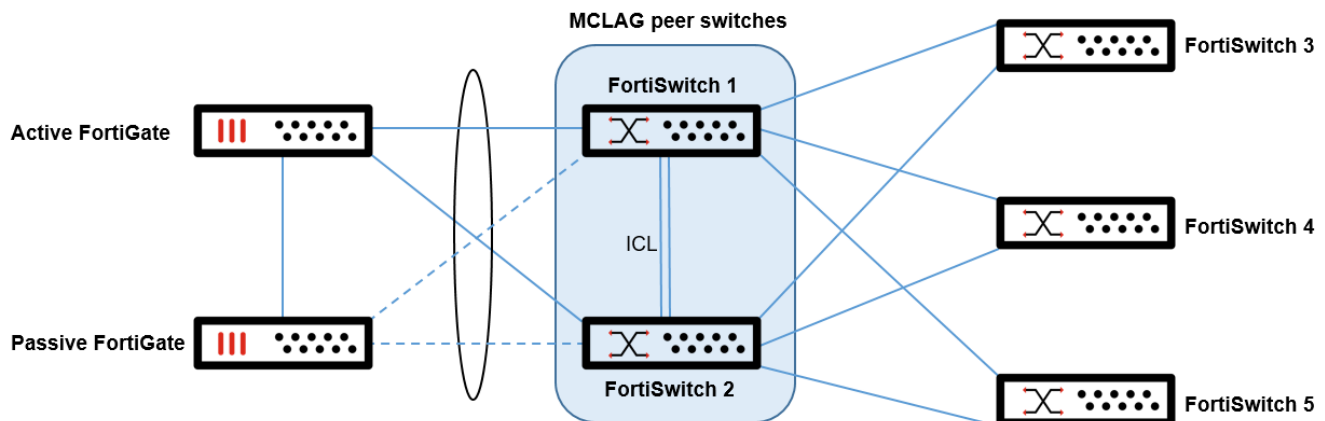
Optionally, you can connect a standby FortiLink connection to the last FortiSwitch unit. For this configuration, you create a FortiLink split interface (an aggregate interface that contains one active link and one standby link).



Large SD-branch

The following architecture contains HA-mode FortiGate devices with dual-homed FortiSwitch access.

After the MCLAG peer group is created between FortiSwitch 1 and FortiSwitch 2, the MCLAG trunks are automatically established with the access switches (FortiSwitch 3, FortiSwitch 4, and FortiSwitch 5).



Management

The FortiManager unit is also used to centrally manage SD-branch architectures. It allows the efficient deployment at scale with zero-touch provisioning. FortiGate devices can initiate the connection to the FortiManager unit, which pushes the relevant configuration, including the switching configuration for all managed switches and other extension devices (FortiAP wireless access points and FortiExtender WAN) in the branch, making deployment simple and accomplished within minutes. This architecture can be replicated to multiple sites to have an efficient and consistent deployment at scale.

Appendix: Documentation references

For more information, use the following resources:

- Deployment guides
 - [LAN Edge Deployment Guide](#)
 - [Large Campus Switching Deployment Guide](#)
 - [FortiLink over Multi-vendor Networks Using VXLAN Deployment Guide](#)
- Solutions hub
 - [Secure Access](#)



www.fortinet.com

Copyright© 2023 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.