

# Arista Zero Trust Campus Network with Multi-Domain Segmentation Services

Version 1.0

April 2024

<b>Introduction</b>	3
Definitions and acronyms	3
Zero Trust requirements	4
<i>Dynamic multidimensional nature of lateral-movement prevention policies</i>	4
<i>Open design goals</i>	5
MSS Technology	5
<i>MSS Group example for campus network</i>	6
<i>Arista campus products for MSS</i>	8
<i>MSS Enforcement Actions</i>	8
<i>MSS Policy example for campus network</i>	9
<i>MSS Open Architecture</i>	12
Prevalence of wireless endpoints in campus networks	12
Bridge- versus tunnel-based architectures for wireless AP forwarding	12
<i>Bridge-based</i>	13
<i>Tunnel-based</i>	13
<i>Comparison of bridge- and tunnel-based approach</i>	14
<b>Incremental insertion of Arista MSS in a brownfield bridged-based wireless architecture</b>	15
Bridge-based Wireless Approach in Traditional Campus Network	15
<i>Inflexibility of traditional campus macro-segmentation technology</i>	17
Arista Campus Leaf insertion	17
MSS Leaf Enforcement	18
MSS Services for wired and wireless traffic	18
<i>Client isolation in wireless access points with bridge-based forwarding</i>	20
<i>Secure bridge-based forwarding for local E-W wireless communication</i>	20

# Table of contents

Overlay network design requirements	20
MSS Leaf Enforcement with overlay network design	20
Integration with Arista bridge-based Wireless Solution	22
<i>Client isolation in Arista access points with bridge-based forwarding</i>	23
Extending the solution with Arista Campus Spine	24
<b>Incremental insertion of Arista MSS in a brownfield tunnel-based wireless architecture</b>	26
Tunnel-based Wireless Approach in Traditional Campus Network	26
Arista Campus Spine insertion	27
MSS Spine Enforcement	28
<i>MSS Spine Services for Wireless Traffic</i>	28
<i>MSS Spine Services for Wired Traffic</i>	30
<i>MSS Spine Services for both Wireless and Wired Traffic</i>	31
Integration with Arista tunnel-based wireless solution	32
Extending the solution with Arista Campus Leaf	32
<b>Insertion of stateful MSS nodes in campus network</b>	34
<b>References</b>	35

## Introduction

This document describes a set of network design options to build a Zero Trust campus network based on Arista Multi-domain Segmentation Services (MSS) technology, combined with different wireless/wired network and security function implementations, using both third party and Arista products.

## Definitions and acronyms

The following table defines in alphabetical order the technical terms and acronyms used throughout this document.

Technical term	Description
<b>ACL</b>	An Access-Control List is a list of rules that specify what network resources and communication types are permitted or denied on a given interface
<b>AGNI</b>	Arista Guardian for Network Identity is Arista next generation cloud-native solution that delivers identity-based NAC
<b>ARP</b>	Address Resolution Protocol is a network protocol used for discovering the layer-2 address (MAC) associated with the layer-3 address (IP) of a device
<b>CAPWAP</b>	Control and Provisioning of Wireless Access Points is a protocol used by a wireless controller to manage access points, that can use an IP UDP tunnel for transporting wireless data traffic
<b>CMDB</b>	Configuration Management Database
<b>CSV</b>	Comma Separated Values is a text file format that uses commas to separate values, and newlines to separate records, also defined in RFC-4180
<b>DMZ</b>	Demilitarized Zone is a physical or logical portion of the corporate network that exposes external-facing services
<b>EVPN</b>	Ethernet Virtual Private Network is a control-plane technology for carrying layer-2 or layer-3 traffic as a virtual private network using transport network protocols such as VXLAN
<b>GRE</b>	Generic Tunnel Encapsulation is an IP protocol for tunneling data packets that can be used for transporting wireless data traffic
<b>IEEE</b>	Institute of Electrical and Electronics Engineers
<b>IETF</b>	Internet Engineering Task Force
<b>IP</b>	Internet Protocol is the most prevalent layer-3 protocol
<b>IPAM</b>	An IP Address Management system is used for administering network services that assign and resolve IP addresses in a network
<b>L2PTV3</b>	Layer-2 Tunneling Protocol Version 3 is an IP UDP protocol for transparently tunneling layer-2 packets that can be used for transporting wireless data traffic
<b>Lateral vs Vertical or E-W vs N-S</b>	In the context of campus networks, Vertical or North-South indicates the communication between campus endpoints and non-campus resources reachable over the public internet network or located in corporate data centers. It contrasts with Lateral or East-West communication that indicates communication patterns between campus devices.
<b>layer-2</b>	Data-link layer in OSI model
<b>layer-3</b>	Network layer in OSI model
<b>NAC</b>	Network Access Control
<b>NDR</b>	NDR is an Arista platform that delivers continuous diagnostics for automatic network anomaly detection
<b>OSI</b>	Open Systems Interconnection model for network communication
<b>PBR</b>	Policy-Based Routing performs a layer-3 forwarding operation based on a multi-field filter instead than uniquely on the destination address used by regular routing
<b>Proxy-ARP</b>	Local Proxy-ARP is the mechanism by which a gateway on a given VLAN answers the ARP queries for a local device IP address and offers its own MAC address as the resolved destination
<b>PVLAN</b>	Private VLAN is a layer-2 isolation mechanism that uses the combination of a primary and a secondary VLAN to limit peer-to-peer communication otherwise possible in a regular VLAN

Technical term	Description
<b>RFC</b>	Request For Comments are used by IETF for documenting standards
<b>SSH</b>	Secure Shell Protocol is a cryptographic network protocol for remote secure access over an unsecured network based on IP TCP
<b>SSID</b>	A Service Set Identifier is the logical network instance used for a group of wireless devices as defined by IEEE 802.11 standard
<b>TCAM</b>	The Ternary Content-Addressable Memory (TCAM) is an expensive high-speed memory component in modern switch ASICs to execute multi-field searches prior to packet forwarding decisions
<b>TCP</b>	Transmission Control Protocol is a common session-based protocol used for IP communication
<b>UDP</b>	User Datagram Protocol is a common session-less protocol used for IP communication
<b>VLAN</b>	Virtual Local Area Network is a layer-2 ethernet segmentation construct defined by IEEE 802.1Q standard
<b>VNI</b>	Virtual Network Identifier is 24-bit value that identifies a virtual network in the VXLAN data plane, typically corresponding to a VLAN or VRF
<b>VRF</b>	Virtual Routing and Forwarding is a layer-3 segmentation technology that allows multiple instances of a routing table to co-exist within the same router
<b>VXLAN</b>	Virtual eXtensible LAN is a network virtualization or overlay technology that encapsulates data packets in an IP UDP tunnel
<b>Whitelist</b>	An explicit list or register of entities that are trusted to receive access to a particular service or resource
<b>YAML</b>	Ain't Markup Language (originally Yet Another Markup Language) is a human-readable data serialization language

## Zero Trust requirements

Zero Trust is a security model founded on the principle that trusted communication is never granted implicitly and must be continually evaluated. A Zero Trust Architecture [insert link with NIST reference] requires the ability of dynamically classifying the enterprise devices in smaller trust zones (micro-perimeters) than traditional network segmentation constructs, where lateral communication can be granularly controlled.

Campus networks primarily provide IP connectivity between campus endpoints (corporate resources, employee, guest or customer devices) and non-campus resources reachable over the public internet network or located in corporate data centers. This prevalent N-S traffic is typically secured by internal and perimeter firewalls that allow only desired conversations based on the source IP address, destination IP address and type of communication (protocol, transport ports, application). This system works in conjunction with a NAC solution that authorizes campus devices to communicate using a specific IP address as source, based on authentication protocols like IEEE 802.1X [insert link with ref.] and classification policies.

Campus networks are subject to lateral movements because their devices are usually segmented in the network using layer-2 constructs called VLANs, that by default allow peer-to-peer communication, either directly within the VLAN domain (bridged communication) or thru a layer-3 gateway (routed communication) implemented inside the internal firewall perimeter. This makes campus networks prone to lateral-movement attacks that exploit the fact that by taking ownership of an authorized device, the attacker can silently expand its control to more privileged devices in the campus network and eventually gain access to critical destination resources outside the campus network.

### Dynamic multidimensional nature of lateral-movement prevention policies

A micro-perimeter is a network-agnostic construct that defines a logical subset of homogeneous devices. Naturally, devices can be characterized with different and orthogonal properties, which can be inherent to the device itself like the model and software version, or inherited from the person operating the device, like the user role or group, or from the characteristics of the physical and logical environment where the device is present, like the physical location or the deployment environment (production vs. non-production).



This means that the micro-perimeter is conceptually multi-dimensional and shall be implemented by associating multiple property values to devices, so that such multi-dimensionality can be expressed by individual policies that operate individually on a particular dimension.

It is also expected that some of these attributes that define multi-dimensional micro-perimeters may vary over time: for example, if a particular operating system version becomes vulnerable to a security threat, and some affected devices do not support a remote software upgrade, it is necessary to tag them accordingly so that they can be matched by a security policy that restricts their lateral communication. As a second example, the device location varies because of the roaming capabilities of the wireless network that do not require IP address renumbering. As a final obvious example, in the context of an IP network, a device or an endpoint is uniquely identified by its IP address, which can be dynamically assigned during authentication or onboarding.

For the above considerations, zero trust network solutions require a system that can abstract multidimensional segmentation policies based on diverse attributes or tags that are associated with IP addresses or prefixes, where these tag-prefix mappings are either dynamically maintained or can be promptly redefined.

#### Open design goals

The initial micro-perimeter segmentation solutions implemented on switches by third-party network vendors have received interest because of their wire-speed performance, but have eventually encountered resistance due to their inconsistent architecture between campus and data center, as well as their dependency on specific data-plane tagging protocols, which often superimpose a mandatory network design (for example a network overlay) and break interoperability with traditional firewalls and infrastructure of other network vendors.

As such, a crucial requirement of a zero-trust network architecture is its openness: that is to say being based on open standards, as much as possible agnostic to network design, extensible from data center to campus networks and deployable in a multi-vendor environment.

#### MSS Technology

This paragraph provides only a summary of Arista MSS features and operational workflows. For a more complete analysis of MSS architecture, please refer to the Arista MSS Technical Whitepaper document in the [References](#) section.

Arista MSS technology provides the capability of dynamically discovering micro-perimeters, baselining their session traffic mapping, and finally recommending whitelist policy rules that can be enforced on Arista products inserted in an existing network environment.

A micro-perimeter in MSS is abstracted by a security group value or tag that identifies a group of devices as a list of IP prefixes or their individual IP addresses. In a Zero Trust campus architecture context, tags can be manually defined or automatically discovered thru the integration with the following systems:

- Arista AGNI or third-party NAC systems, which authenticate endpoints using standard protocols like IEEE 802.1X
- External databases like CMDBs, IPAM systems, CSV or YAML files
- Arista NDR, which derives tags by classifying the actual device traffic.

In MSS architecture, any endpoint can be labeled with more than one tag, thus security policies can be expressed flexibly in terms of diverse properties, like device type, access group type, security risk level, deployment environment (production vs non-production), and location.

None of these tags are associated with a network isolation construct by default: the fact that a device (in the form of its IP address) is member of a particular group value does not implicitly imply that all devices with the same tag value can mutually communicate.

In fact, following Zero Trust principles, the only default behavior is that no traffic is allowed, and thus, specific lateral and vertical traffic movements need to be explicitly defined for certain tags and specific network protocols.

MSS Group example for campus network

As represented in the following diagram, a fictitious company A has four terminal card-accessible devices located at the entrance of one of their buildings, which are restricted to use one server (Server-1) as their network gateway, located at an upper floor of the same building, in order to communicate with public internet. Adjacent on the same floor there is the equipment used by the development team responsible for the terminal application.




Figure 1: An example of terminal devices connecting to the internet via a server

The diagram below describes the logical connectivity viewpoint.

The four production terminals and the server Server-1 connect to the campus network using the same VLAN (VLAN-2). The development environment, represented in yellow, which includes terminal devices, servers and workstations, uses primarily a separate VLAN (VLAN-1), although for some experiments, the IP address of one of the terminal devices (T-3) has been relocated to the subnet used by the production terminals. Server-1 is the sole endpoint to have connectivity to both the production and the development VLANs and in fact it uses two different IP addresses. The diagram also shows how the two VLANs are routed thru a network gateway that connects to the public internet via an adjacent perimeter firewall.




Figure 2: Logical connectivity diagrams of the terminal devices and the development environment

The IP address assignment for each campus endpoint is as following:

Devices (optional interface)	IP address or prefix/mask
<b>All campus network devices</b>	192.168.0.0/16
<b>Production network</b>	192.168.20.0/24
T-4	192.168.20.104
T-5	192.168.20.105
T-6	192.168.20.106
T-7	192.168.20.107
T-3	192.168.20.23
<b>Server-1 (interface:production)</b>	192.168.20.254
<b>Development network</b>	192.168.10.0/24
<b>Server-1 (interface:development)</b>	192.168.10.254
O-1	192.168.10.11
T-1	192.168.10.21
T-2	192.168.10.22
<b>Server-2</b>	192.168.10.253

From a security perspective, these devices need to be described by a multi-dimensional tag system that, regardless of the network assignments, takes into account the role of the device and its deployment environment. With Arista MSS technology, this can be efficiently expressed with the following tag definitions and mappings, also visually represented in the next figure:

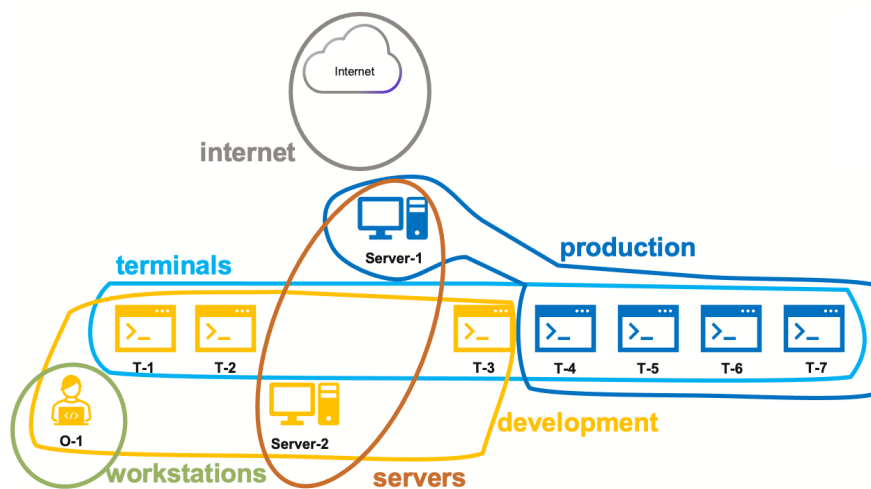


Figure 3: Security-driven classification of the devices in the example

Tag or Group	IP address or prefix/mask	Notes
<b>internet</b>	0.0.0.0/0	
	except 192.168.0.0/16	Exclude corporate campus intranet from 0/0
<b>production</b>	192.168.20.0/24	
	except 192.168.20.23/32	Exclude T-3
	192.168.10.254/32	Include development interface on Server-1
<b>development</b>	192.168.10.0/24	
	except 192.168.10.254/32	Exclude Server-1 development interface
	192.168.20.23	Include T-3
<b>terminals</b>	192.168.10.21	
	192.168.10.22	
	192.168.20.23	
	192.168.20.104	
	192.168.20.105	
	192.168.20.106	
	192.168.20.107	
<b>servers</b>	192.168.10.253	Server-2
	192.168.10.254	Server-1 development interface
	192.168.20.254	Server-1 production interface
<b>workstations</b>	192.168.10.11	

#### Arista campus products for MSS

At the time of writing this document, MSS is supported on the CCS-720, the 7050X3 and the 7280R3 switch families, and requires the EOS V2 add-on license. MSS stateful services are supported on the ZTX-7250S appliance. Specific model support and differences in scalability and capability can be reviewed in the MSS Datasheet, listed in the [References](#) paragraph.

MSS Orchestration and Telemetry are included in Arista Cloud Vision, which offers a dedicated dashboard for MSS policy manager, the integration with external tools for dynamic microperimeter management, the rule recommendation engine and the session log database.

#### MSS Enforcement Actions

MSS policy rules can enforce on supported Arista switches, at line rate and without performance impact, the actions described in the following table and drawing:

Action Name	Enforcement Result on matched traffic	Purpose
<b>FORWARD<sup>1</sup></b>	Traffic is permitted	Explicitly allow lateral and vertical traffic
<b>DROP<sup>1</sup></b>	Traffic is denied	In a whitelist approach, it is used as a default zero-trust rule to negate traffic not explicitly permitted, or to create specific drop exceptions for whitelisted traffic
<b>REDIRECT<sup>2</sup></b>	Traffic is redirected to a third-party firewall	Force lateral traffic to hairpin to a firewall gateway for L4-L7 inspection. Not relevant for vertical traffic normally routed to a firewall gateway, which can be whitelisted using FORWARD action
<b>MONITOR<sup>3</sup></b>	A truncated copy of each traffic packet is forwarded to a MSS Node configured in Monitor Mode	Used as standalone action: monitor a network zone to derive session traffic mapping used for automatic policy recommendations. Used in combination with a DROP action: provide visibility into dropped packets.
<b>INSPECT<sup>4</sup></b>	Traffic is redirected to a MSS Node configured in Inspect Mode	Hairpin traffic to an MSS Node for L4-L7 inspection. When used in combination with DROP action, traffic is denied after inspection.

With very few exceptions, the enforcement rules expressed with a whitelist approach are generally bidirectional.




Figure 4: Graphical representation of MSS enforcement actions

MSS Policy example for campus network

Using the aforementioned example of the terminal devices, the security team of company A has assessed that the perimeter firewall rules already present are sufficient to control N-S traffic as described in the following diagram. For instance, there are rules based on the device IP address, for which internet access is granted (green arrow) to the development environment and, for the production devices, only from/to Server-1 (green arrow) and denied (red arrow) from/to terminals.

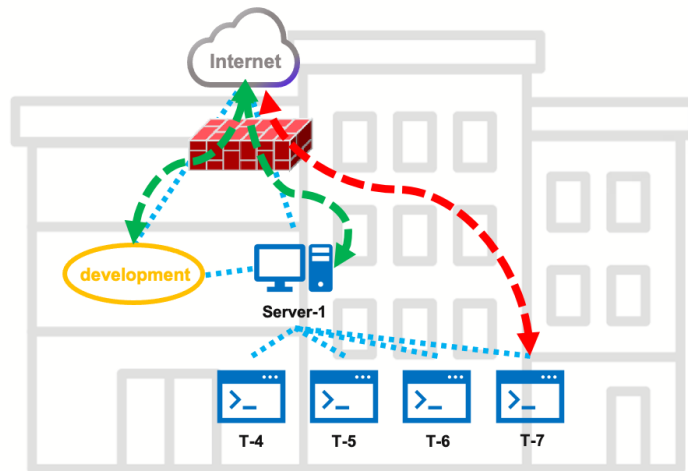


Figure 5: Graphical representation of N-S security rules for the example

What the current centralized firewall architecture cannot protect, are the device conversations that happen inside the building network.

The first security concern is that an attack, aimed at causing a denial of service or at collecting card/user information, can be perpetrated on one terminal device and propagate through the layer-2 segment to all the other production terminals.

Secondly, there is a high-security risk of lateral movements that are possible from the publicly accessible terminals on the production network to other privileged devices that have access to the internet and potentially other important company assets and data. In particular the figure below illustrates three examples of lateral movements:

1. between a production and a development terminal on the production network
2. between a production terminal and Server-1, which has access to the development network via the dedicated interface
3. between a production terminal and any device of the development network via the network gateway

In fact, the above communication cases are made possible by the existing campus network by bridging (1 and 2) and routing (3) operations.

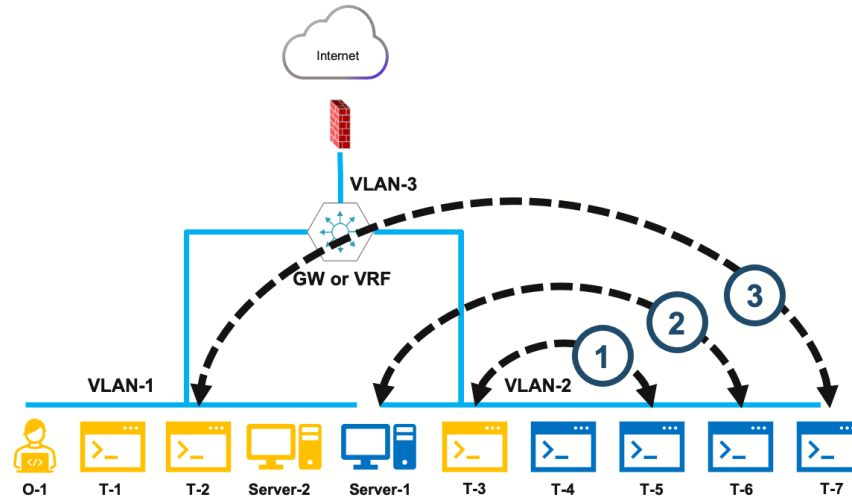


Figure 6: Potential lateral communication for terminal devices

Assuming the campus of company A is powered by Arista switches, using the MSS tag taxonomy explained [earlier](#), this campus network can be secured by implementing on the Arista switches the MSS group definitions along with an MSS whitelist policy, inclusive of the following bidirectional enforcement rules:

Source	Destination	Network Service	Enforcement Action	Notes
development	internet	*	FORWARD	N-S development traffic normally routed to the firewall
production servers	internet	WEB	FORWARD	N-S production traffic normally routed to the firewall. Source combines two tag values (AND operation), to exclude direct traffic from terminals.
production	development	*	REDIRECT to firewall	
production	production	*	REDIRECT to firewall	
development	development	*	FORWARD	
*	*	*	DROP MONITOR	default zero-trust rule

Note that the presence of the last drop-anything-else rule in the above policy, necessary for Zero Trust compliance, requires that legit N-S traffic is explicitly whitelisted, as it happens with the first two rules. In reality, as N-S traffic is granularly inspected by the firewall, N-S rules implemented with MSS do not need to be detailed. In this example they could be replaced by a single coarser rule as in:

Source	Destination	Network Service	Enforcement Action	Notes
*	internet	*	FORWARD	N-S development traffic normally routed to the firewall

Note that the presence of the last drop-anything-else rule in the above policy, necessary for Zero Trust compliance, requires that legit N-S traffic is explicitly whitelisted, as it happens with the first two rules. In reality, as N-S traffic is granularly inspected by the firewall, N-S rules implemented with MSS do not need to be detailed. In this example they could be replaced by a single coarser rule as in:

Note that the presence of the last drop-anything-else rule in the above policy, necessary for Zero Trust compliance, requires that legit N-S traffic is explicitly whitelisted, as it happens with the first two rules. In reality, as N-S traffic is granularly inspected by the firewall, N-S rules implemented with MSS do not need to be detailed. In this example they could be replaced by a single coarser rule as in:

Source	Destination	Network Service	Enforcement Action	Notes
*	internet	*	FORWARD	
production	development	*	REDIRECT to firewall	
production	production	*	REDIRECT to firewall	
development	development	SSH	REDIRECT to firewall	address new vulnerability
development	development	*	FORWARD	
*	*	*	DROP MONITOR	default zero-trust rule

It is noticeable that the above security rules are completely network independent and will not vary if the resources change their network properties or if new resources are added to the existing groups. For example, if company A expands to new adjacent offices, whose networks communicate with the existing gateway, it is sufficient to add the IP address or prefix of the new devices or networks to the defined MSS groups configured on the Arista switches.

Few final comments are related to the efficiency of the MSS solution in terms of hardware resource utilization on Arista switches.

The hardware entries used to implement MSS rules in the switch TCAM use source and destination labels that correspond in hardware to the configured MSS groups or tags, hence they are not affected by the number of prefixes that are associated with these groups. For instance, the following definitions of the production environment are equivalent in terms of TCAM utilization, despite the number of prefixes differs:

Source	IP address or prefix/mask	Notes
production	192.168.20.0/24	Production network
	except 192.168.20.23/32	Exclude T-3
	192.168.10.254/32	Include development interface on Server-1

Source	IP address or prefix/mask	Notes
production	192.168.20.104/32	T-4
	192.168.20.105/32	T-5
	192.168.20.106/32	T-6
	192.168.20.107/32	T-7
	192.168.20.254/32	production interface on Server-1
	192.168.10.254/32	development interface on Server-1

A final consideration regards the fact that MSS tags that are defined on a switch, but not referenced as source or destination by any MSS rule, do not consume hardware resources. In the existing example, there are three groups that are not present in the MSS policy rules:

Source	IP address or prefix/mask	Notes
<b>terminals</b>	192.168.10.21	
	192.168.10.22	
	192.168.20.23	
	192.168.20.104	
	192.168.20.105	
	192.168.20.106	
	192.168.20.107	
<b>servers</b>	192.168.10.253	Server-2
	192.168.10.254	Server-1 development interface
	192.168.20.254	Server-1 production interface
<b>workstations</b>	192.168.10.11	

In such a case, the MSS groups named terminals, servers and workstations and their associated prefixes do not use any hardware resource on the switch.

#### MSS Open Architecture

MSS Architecture adheres to open standard protocols and industry best practices. It is predicated on the principle that network segmentation and policy enforcement can be accomplished on networks as small as a single switch and do not require either multiple network layers or physical separation between ingress classification elements and egress enforcement elements, nor specific network designs or special data-plane implementations. As such, zero-trust solutions that use a consistent and fully interoperable architecture like Arista MSS, can span different network locations, for example distributed data centers and campus networks and be deployed in a multi-vendor network environment.

#### Prevalence of wireless endpoints in campus networks

Campus networks provide network connectivity to many devices located in buildings, like offices, universities, hospitals, industrial plants, etc. These endpoint devices can be operated by a person, like workstations, personal devices, laptops, terminals, or can be systems that transmit and receive data like sensors, security cameras, printers, and Operational Technology (OT) machines.

Depending on its capability, the endpoint connectivity to a campus network can be either wired, thru an ethernet cable or fiber that directly connects to an access switch, or wireless, thru an ethernet network element wired to the access switch, called wireless access point (AP), which provides the wireless local area network (WLAN) connectivity to the wireless endpoints using a technology based on IEEE 802.11 standards [add link in References].

Wireless device connectivity has become predominant in campus, because of endpoint mobility requirements and reduced cost of connectivity compared to wired connectivity: most campus devices are nowadays wireless, while wired connectivity is reserved for a minority of devices which are either legacy appliances or systems that have specific latency, bandwidth or connection media requirements that cannot be satisfied by wireless technology.

This guide evaluates multiple campus network design options, and given the dominance of wireless traffic, it is logically organized around wireless traffic architectures and according to each individual wireless architecture option it proposes a specific zero-trust solution design.

#### Bridge- versus tunnel-based architectures for wireless AP forwarding

Although wireless AP implementations differ across networking vendors in terms of control- and data-plane architecture, there are only two fundamental forwarding models: bridge-based and tunnel-based.



Bridge-based

Conceptually, the bridged-based forwarding architecture implements for the wireless traffic a similar segmentation approach to the one used for the wired traffic.

Wireless access points configured following this design option, are connected to access switches using an IEEE 802.1Q [insert link with ref.] trunk and hand off traffic from/to wireless devices with a VLAN tag that is specific to each endpoint category. The following figure abstracts this concept and represents two wireless endpoint categories, personal devices and printers, whose traffic is separated on the ethernet link between the access point and the access switch.

To secure the communication between personal devices and printers, the traffic segregation implemented in the 802.1Q trunk needs to be extended throughout the entire corporate network to a firewall. This is typically achieved in two ways:

1. The 802.1Q VLANs are extended to the firewall, which functions as a gateway for each layer-2 segment.
2. The gateways of the 802.1Q VLANs are implemented in the corporate network in different routing domains, for example using distinct Virtual Routing and Forwarding (VRF) instances. These routing domains are then communicating through the firewall, which for example acts as default gateway for each VRF.

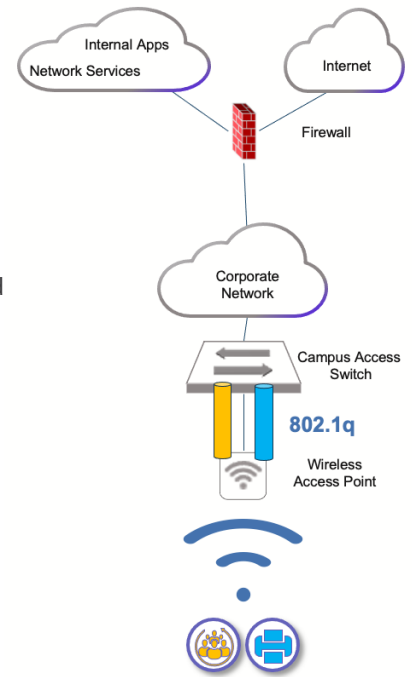


Figure 7: Bridged-based wireless architecture

Tunnel-based

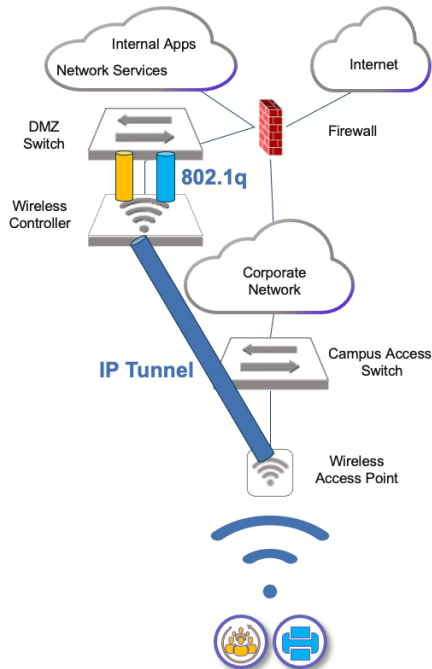


Figure 8: Tunnel-based wireless architecture

With this approach, an IP Tunnel is implemented between each access point and a centralized wireless controller, based on an IP protocol that varies with the wireless network vendor: most common protocols used are UDP (in the case of Cisco CAPWAP), GRE, L2PTV3.

This tunneling mechanism, represented in the following diagram, allows the traffic of different wireless endpoint categories to traverse the corporate network in almost its entirety without implementing a segmentation technology, like 802.1Q or VRF, usually with the sole exception of the DMZ switch (pair) that physically interconnects the wireless controller to the firewall. On that switch, where the wireless traffic is handed off with VLAN tags, the traffic is segregated using one of the two methods described in the previous chapter, most likely by implementing the VLAN gateways on the firewall.

The design described may become more complex in case the wireless controller and the firewall are not in proximity, requiring a segregation mechanism to be put in place on a larger network. Another factor to consider is that scale requirements may include in the design multiple wireless controllers or a hierarchy of controllers, making the solution more elaborated.

## Comparison of bridge- and tunnel-based approach

The following comparative table summarizes the differences between the two wireless forwarding models:

bridge-based wireless forwarding	tunnel-based wireless forwarding
Access switches receive wireless traffic with 802.1Q encapsulation and manage forwarding information of wireless clients like VLAN identifier, layer-2 and 3 address	Access switches receive wireless traffic with IP tunnel encapsulation and do require to store wireless client information
Consistent architecture: same forwarding model as 802.1Q wired traffic	Less consistent architecture: different forwarding model than 802.1Q wired traffic
Operationally inefficient: provisioning of wireless VLANs is required on campus switches	Operationally efficient: provisioning of wireless VLANs is not required on campus switches
Inherently open to lateral movement because of default any-to-any behavior of bridging and routing	Inherently limiting lateral movement because of client-isolation mechanism and tunnel hair-pinning

The choice between bridge- and tunnel-based models can be ascribable to personal preference of the network transport and security designers or their predecessors; or they can be motivated by specific characteristics of a customer environment/vertical. For the latter case, it is worth mentioning two instances:

1. edge-compute systems for industrial and healthcare applications with low-latency requirements for localized same-floor communication, which make centralized hair-pinning a suboptimal choice, and thus the decentralized forwarding model of bridge-based APs highly preferable. This is also true when such type of high-performance communication is required between wireless and wired devices.
2. Universities, event venues, hospitality customers, which periodically introduce new wireless segments to increase endpoint scale or to put in place new types of endpoints. In such a case, the tunnel-based approach is more suitable, as it avoids the operational effort of configuring VLANs on access switches.

*MSS Technology can be deployed for both wireless architecture types and concomitantly implement a Zero Trust design inclusive of both wireless and wired traffic.*

The following two sections discuss how MSS technology can be inserted in a bridge-based and tunnel-based wireless architecture, to implement a campus micro-perimeter segmentation strategy for both wireless and wired traffic.

## Incremental insertion of Arista MSS in a brownfield bridged-based wireless architecture

This section describes how MSS Technology can be inserted in an existing campus network designed using a bridged-based wireless architecture.

### Bridge-based Wireless Approach in Traditional Campus Network

The following diagram represents a traditional multi-building campus network, based on a two-layer topology (access-and-distribution or leaf-and-spine) using generic network vendors: both the wired and wireless traffic of each building enters the bottom layer with a specific VLAN tag. The gateways of these VLANs are typically centralized on the upper layers of the network topology and in some scale-limited cases can be also located in the firewall: the example in the diagram uses the distribution switches in each building to act as the layer-2/3 demarcation point.

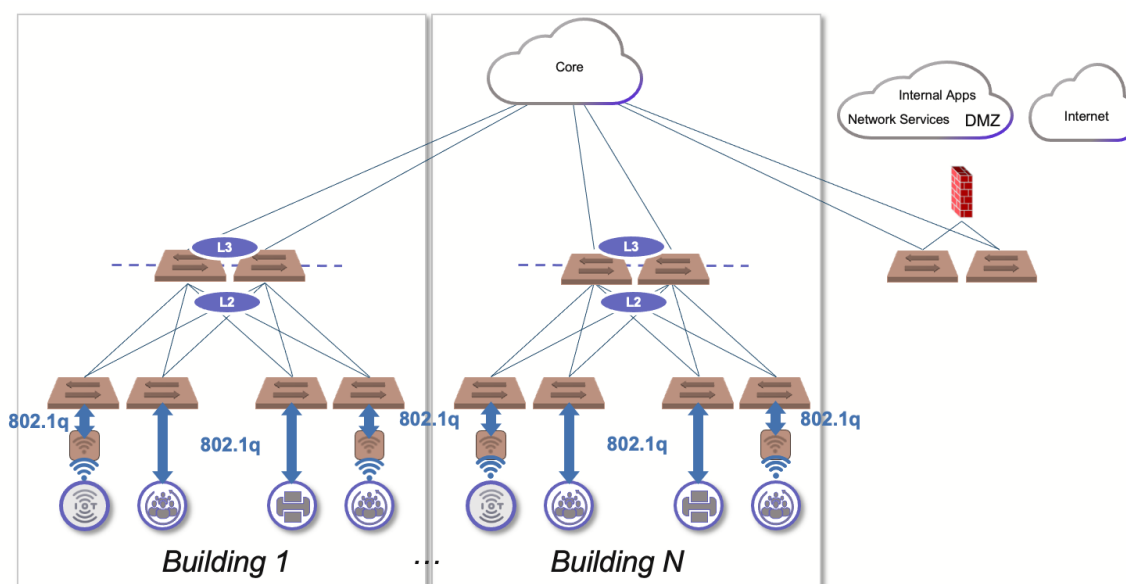


Figure 9: Traditional campus network with bridged-based wireless architecture

The buildings are interconnected through a generic core network, whose topology is not defined, that routes the inter-building traffic to the corporate data center network where connectivity to the internal services and applications and to the public internet and cloud services is provided and secured by a firewall function.

The next diagrams expand the example with a detailed view of layer-3 segmentation implemented in the upper layers, where VRFs are used to isolate the traffic of a group of VLANs preventing it from being routed in the campus network to other VLANs that are member of a different VRF.

In the first diagram, the following basic examples of VRF segmentation are used:

VRF name	Purpose
VRF-A	In-band control-plane and management-plane traffic of campus network appliances, like for example wireless access points
VRF-B	Workstations
VRF-C	Printers
VRF-X	Wireless corporate devices
VRF-Y	Wireless IoT devices
VRF-Z	Wireless guest devices

VLANs are represented by continuous segments following the same color code of each VRF, while the dotted lines represent the default route for each VRF.




Figure 10: Layer-3 segmentation in a traditional campus network

Each VRF has usually a default route with a next-hop address owned by the firewall, either announced by a dynamic routing protocol or using a static route. This segmentation approach prevents any campus VRF from communicating to other campus VRFs without proper firewall inspection. However, devices residing in different campus buildings that are part of the same VRF are mutually reachable via the core network and thus, they can directly communicate using any protocol unless a specific ACL prevents it along their switched or routed path.

To secure building-to-building communication thru the firewall for one or more VRFs, a possible strategy, represented in the next figure, consists in a two-level design obtained by creating in the core network a unique VRF instance for each building.

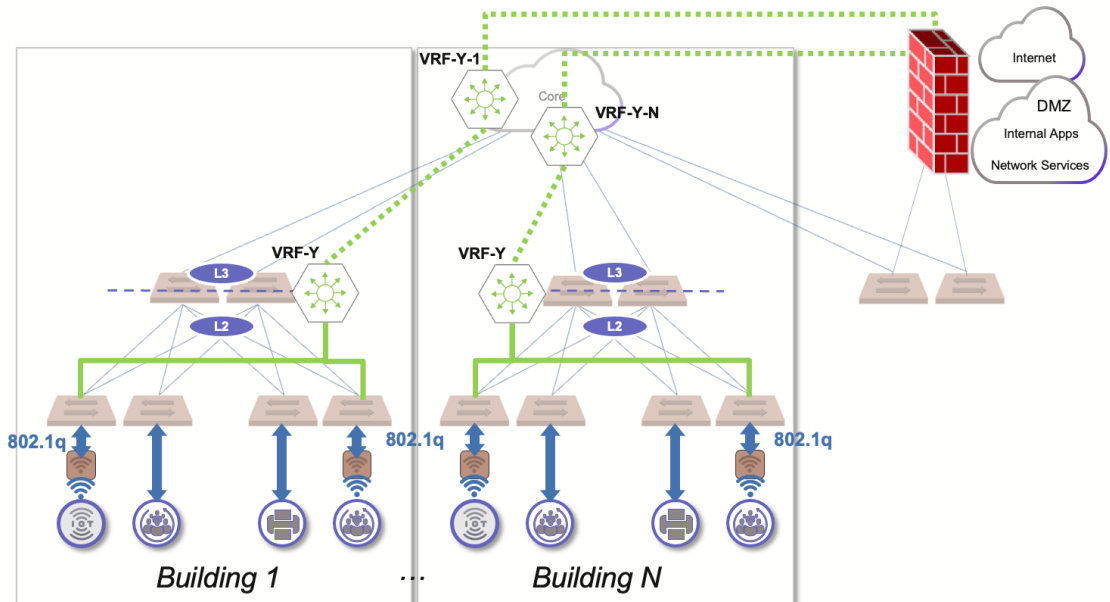


Figure 11: Inter-building isolation in a traditional campus network

In the specific example, the communication between wireless IoT devices (part of VRF-Y) located in multiple buildings is hair-pinned to the firewall without exception.

The approach just described is suitable to the nature of campus traffic patterns, which are mostly directed to a location outside the campus network (or N-S).

It is, however, worth noting that the proliferation of VRF instances in the core network can have a negative impact on control plane scale and operational scale.

#### Inflexibility of traditional campus macro-segmentation technology

Considering the goal of lateral-movement avoidance, the VRF-based solution described provides inter-building isolation for a particular VRF but does not prevent the devices of such VRF to communicate when they are inside the same building.

A possible improvement used by the industry consists in implementing a layer-2 isolation mechanism called Private VLAN (PVLAN) coupled with local proxy-ARP configured in the aggregation switches where the layer-2/3 demarcation is implemented, and finally IP ACLs or PBR to either block intra-building traffic or hairpin it to the firewall. As on switches proxy-ARP can be control-plane intensive and ACLs are often limited in scale, it is not uncommon that certain solutions opt for a layer-2 design with PVLAN for the entire core network, by placing the VLAN gateways with local proxy-ARP on the firewall for both wired and wireless endpoints. This of course may not be a universally viable solution as layer-2 stretching may not be implementable in certain network transport conditions or for scale reasons.

Given the lack of agility and flexibility of VLAN- and VRF-based segmentation, we conclude that it is inadequate to respond to zero-trust network requirements, which change dynamically as new security risks are discovered and in general require more granular and network-independent segmentation.

#### Arista Campus Leaf insertion

Adding dynamic and network-independent micro-perimeter segmentation capabilities to an existing campus network is possible with Arista MSS technology. MSS services can be enabled as in the following network diagram, by just inserting in the wired access layer the Arista Campus Leaf switches, like the Cognitive Campus Series CCS-720XP based on Trident3 chipset by Broadcom, without the need of replacing the rest of the wired network layers and the bridge-based wireless solution already in place.




Figure 12: Arista Campus Leaf insertion in a traditional campus network

**MSS Leaf Enforcement**

The MSS enforcement point in the proposed insertion strategy is obviously on the Campus Leaf layer and applies to both wired and wireless traffic, as logically represented in the figure below.




Figure 13: MSS Leaf enforcement in a traditional campus network

**MSS Services for wired and wireless traffic**

As explained at the beginning of this section, in the existing brownfield design traffic is already segmented using VLANs and VRFs constructs. While both VLANs and VRFs are implemented in the upper portion of the topology, from the perspective of the Arista Campus Leaf layer, the segmentation objects configured are limited to VLANs, as described in the following figure.




Figure 14: MSS Leaf enforcement in default VRF

This means that for a given VRF, a Campus Leaf switch pair can bridge traffic locally between endpoints connected downstream that are part of the same VLAN (case 1). Instead, a conversation between endpoints part of different VLANs of the same VRF (case 2), or part of the same VLAN but attached to different Campus Leaf pairs (case 3), will be necessarily forwarded to upper network layers where it will be either routed or bridged to destination.

The following graphical example illustrates the three cases just described: five wired workstations belonging to VRF-B are connected to two Campus Leaf pairs and are distributed over two VLANs: VLAN-B-1 is present only on the left pair and VLAN-B-2 is present on both pairs. These two VLANs have gateways on the third-party spine pair, part of VRF-B.




Figure 15: Example of lateral conversations: intra-VLAN intra-leaf (1), inter-VLAN inter-leaf (2), intra-VLAN inter-leaf (3)

Regardless of the cases, MSS policies can be enforced on Campus Leaf layer, to control traffic of VRF-B based on properties that are orthogonal to network constructs like VLAN membership. For example, as in the following graphical view, a security risk assessment has determined that the operating system running on workstations WS-1 and WS-5, which are obsolete models, represents a potential threat. For this reason, the security administrator is required to isolate these endpoints from communicating with any of the workstations that are part of VRF-B.




Figure 16: Example of MSS enforcement to limit lateral communication of obsolete devices



This requirement translates into three logical steps:

Defining a security tag “**WORKSTATIONS**” that includes the list of IP addresses of all five workstations.

Defining a security tag “**OBSOLETE**” that includes the IP addresses of the two obsolete workstations.

Blocking some of the conversation cases described earlier, and others not represented in the figure, based on the following security rule, part of a policy that is applied globally to the Campus Leaf switches:

Source	Destination	Network Service	Enforcement Action	Notes
<b>OBSOLETE</b>	<b>WORKSTATIONS</b>	*	DROP MONITOR	

#### Client isolation in wireless access points with bridge-based forwarding

An important consideration concerns how to properly configure third-party access points to fully enable the MSS Leaf enforcement architecture for wireless traffic. Most wireless vendors provide a security feature that commonly goes under the name of Client Isolation, which is complementary to the MSS Leaf enforcement and needs to be enabled. With Client Isolation activated, the traffic sourced by wireless endpoints is always hair-pinned to the upstream switch, preventing that direct peer-to-peer communication in the access point which would otherwise bypass the MSS policy enforcement.

#### Secure bridge-based forwarding for local E-W wireless communication

By default, wireless traffic received by a Campus Leaf from an access point in a particular VLAN cannot be directly bridged back to the same access point: this design complements the intent of avoiding lateral traffic, as it prevents peer-to-peer communication between wireless devices connected to the same access point and SSID.

For certain campus applications, it may be desirable to allow local peer-to-peer communication for performance requirements. The recommendation in such a case is to explicitly whitelist such traffic with an MSS rule that specifies the tags and network services involved and, second, to disable port source filtering from the physical interfaces that connect to the access points.

### Overlay network design requirements

A variant of the network design discussed so far, where both wired and wireless traffic enters the Campus Leaf switches with an 802.1Q encapsulation, consists in implementing on the Campus Leaf switches an overlay mechanism, based on VXLAN and EVPN protocols.

An overlay translates or tunnels VLANs and VRFs into identifiers (VNI) that are encoded into an UDP envelope, which can be transported using a single routed network instance, called underlay, avoiding the extension of VLAN and VRF segmentation to upstream network layers.

This results in a simplification of the aggregation and core network design, which as a unified underlay network can just be implemented with a single VRF, removing the need of running multiple routing instances and improving the overall layer-3 scale of the design.

#### MSS Leaf Enforcement with overlay network design

A feature of a VXLAN-EVPN design consists in building the layer-2/layer-3 demarcation on the Campus Leaf switches, by configuring an anycast gateway for each VLAN within a VRF context, allowing the traffic between VLANs of the same VRF to be routed on the first hop without being hair-pinned to upstream switches. This improves the forwarding performance of the network as far as E-W traffic, which is extremely valuable for campus deployments that require high-speed and/or low-latency same-floor communication.



The components of Network Provisioning are detailed as follow failed device that may be out of date.




Figure 17: MSS Leaf enforcement in overlay campus network

As illustrated in the previous diagram, the MSS enforcement does not vary when implementing an overlay network design and can be distributed on Campus Leaf layer for both wired and wireless traffic. The sole consideration is that whitelist entries for vertical traffic, would essentially be applicable only in the direction from the campus device to the external network. In general, this does not pose any security concern, because in the opposite direction the traffic is already secured by a firewall. Instead, this knowledge could be used to express whitelists for vertical communication only in one direction rather than bidirectionally, improving the scale of the solution.

Unlike the non-overlay case, the presence of the anycast gateway on the switches that act as MSS enforcement point, allows the policies to be applied to specific VRFs rather than to a global context, as in the following diagram.




Figure 18: MSS Leaf enforcement in multi-VRF overlay

One advantage of the overlay design is that the multi-VRF MSS enforcement organically increases the security group scale, because each VRF has its own MSS tagging namespace.

**Integration with Arista bridge-based Wireless Solution**

The network diagrams used in the document so far have been focused, for the wireless architecture discussion, entirely on the forwarding aspects of the design and not on how access points allow wireless devices to connect and roam to the network. Most of the access point solutions in the industry require in fact a wireless controller network, which can be cloud-based, but in most cases relies on the presence of wireless controller appliances distributed in the campus network.

The following diagram represents a typical situation of a two-tier wireless controller network, where secondary wireless controllers are deployed in each building and a primary wireless controller is deployed centrally in the data center.

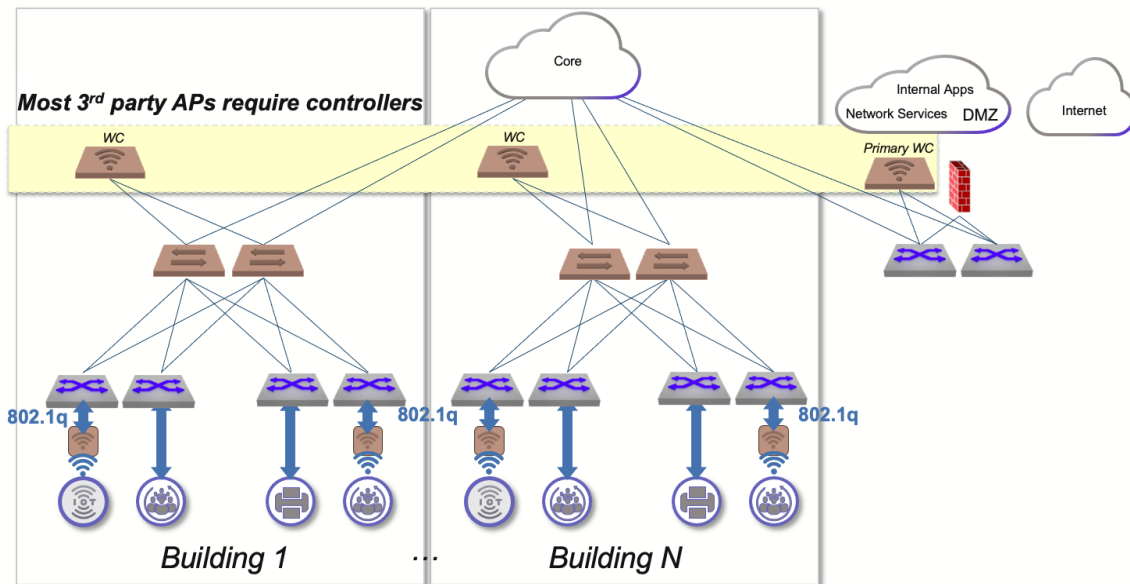


Figure 19: Wireless controller network

Arista wireless solution uses a distributed wireless controller function that is implemented in the access point software and does not require external physical appliances or cloud-based software to control how wireless devices connect to the network.




Figure 20: Arista Wireless solution combined with MSS Leaf enforcement

The diagram above shows how the discussed solution can be further enhanced by adding the Arista wireless network to the existing design or using it to replace the existing third-party access points.

There are in fact two main areas of improvement, where Arista access points are proven to be advantageous.

First, this gives the opportunity to reduce the complexity of the solution by avoiding wireless controllers and the control-plane limitations of their architectures, and improve the overall wireless connectivity and mobility performance, as well as the high availability.

Second, it provides to the administrator a single pane of glass for managing and monitoring both wired and wireless networks.

Client isolation in Arista access points with bridge-based forwarding

In case of Arista access points, client isolation, which is necessary for implementing MSS Leaf Enforcement, is implemented by configuring a feature called Layer-2 Traffic Inspection and Filtering (L2TIF) [add reference link]. As in the following picture, this approach emulates for wireless endpoints the same behavior of wired endpoints, allowing a single MSS enforcement stack for both wired and wireless traffic.

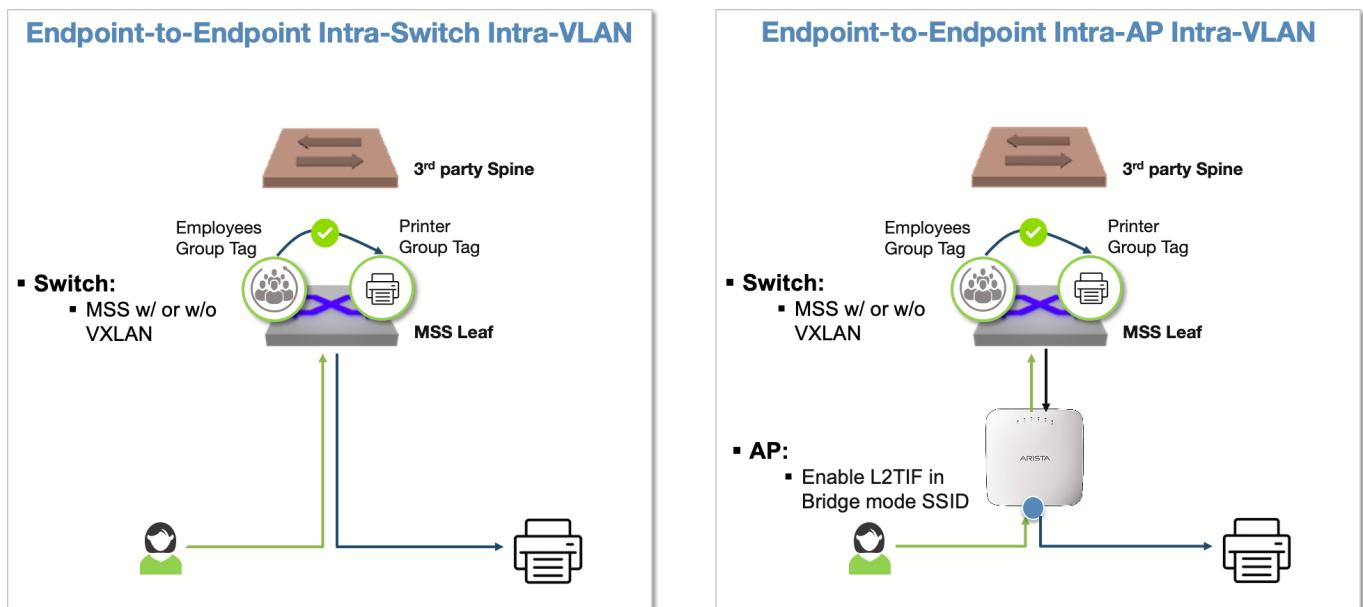


Figure 21: Symmetry of bridged traffic patterns for wired and wireless endpoints with Arista L2TIF

Arista wireless solution is compatible with both traditional and overlay campus network designs, where MSS Leaf Enforcement is applied globally or per VRF respectively, as illustrated in the following two diagrams.




Figure 22: MSS Leaf enforcement in default VRF integrated with Arista Wireless

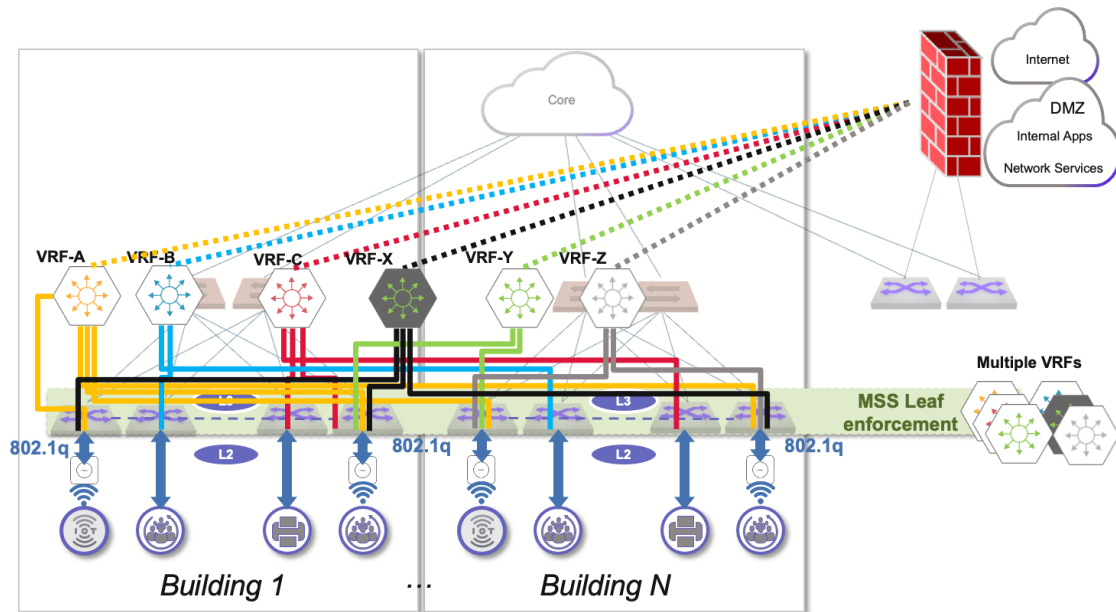


Figure 23: MSS Leaf enforcement in multi-VRF overlay integrated with Arista Wireless

### Extending the solution with Arista Campus Spine

This chapter describes the option of a greenfield design for an entire Arista campus network, or how the design discussed previously can evolve to include pairs of Campus Spine switches that distribute traffic for each building.

A campus solution based entirely on Arista switches, referenced in the picture below, can support both overlay and traditional network design options, and with a bridge-based forwarding approach for wireless traffic, can implement a zero-trust architecture with MSS Leaf enforcement.




Figure 24: MSS Leaf enforcement in all-Arista campus network

The obvious advantage of an all-Arista solution is the prospect of using a single unified management system and providing end-to-end visibility for wired and wireless traffic.

Additionally, Arista Campus Spine pairs can be used to insert MSS Node appliances and act as intelligent load balancers for stateful traffic.

### Incremental insertion of Arista MSS in a brownfield tunnel-based wireless architecture

This section describes how MSS Technology can be inserted in an existing campus network designed using a tunnel-based wireless architecture.

In contrast with the cases discussed in the previous section, there is a difference between the forwarding design for wired and wireless traffic.

#### Tunnel-based Wireless Approach in Traditional Campus Network

The following diagram represents a traditional multi-building campus network, based on a two-layer topology (access-and-distribution or leaf-and-spine) using generic network vendors: the wired traffic of each building enters the bottom layer with a specific VLAN tag. These wired-traffic VLANs are typically terminated on the upper layers of the network topology and in some scale-limited cases can be also located in the firewall: the example in the diagram uses the distribution switches in each building to act as the layer-2/3 demarcation point.




Figure 25: Traditional campus network with tunnel-based wireless architecture

The wireless traffic is instead encapsulated using an IP tunnel technology between multiple access points and a wireless controller, which is either deployed in the aggregation layer of each building, or centralized as a primary wireless controller in proximity to the firewall, as in the following example diagram:

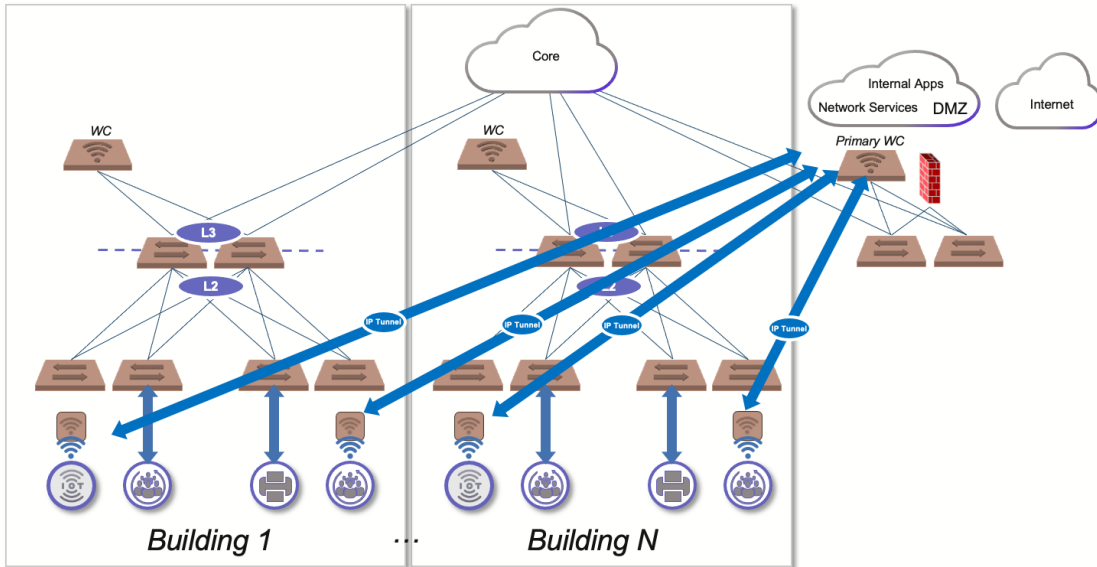


Figure 26: IP tunnels between wireless APs and centralized controller

One of the effects of a tunnel-based forwarding model is that the switches that are downstream of the wireless controllers that act as tunnel termination points, are completely unaware of the IP address space and VLANs used by the wireless devices, as they are configured only to bridge or route packets from/to the IP addresses assigned to the access points. This first of all eliminates the operational effort associated with VLAN and VRF provisioning in the majority of the switches, necessary for the segmentation of wireless traffic. Secondly, it achieves the purpose of reducing or eliminating lateral communication between wireless endpoints, because the traffic is always forced to a central location. This final observation is also corroborated by the fact that tunnel-based forwarding is always activated together with client-isolation. The result is that peer-to-peer communication between wireless endpoints assigned to the same VLAN cannot happen neither on the access point nor on the wireless controller.

**Arista Campus Spine insertion**

In the campus design options discussed in this section, MSS services can be enabled as in the following network diagram, by just inserting in the wired distribution or aggregation layer the Arista Campus Spine switches, like the 7050X3 fixed series based on Trident3 chipset by Broadcom, or like the 7280R3 modular and fixed switches based on Jericho2 chipset by Broadcom, without the need of replacing the rest of the wired network layers and the tunnel-based wireless solution already in place.




Figure 27: Insertion of Arista Campus Spine in a traditional campus network



### MSS Spine Enforcement

The MSS enforcement point in the proposed insertion strategy is obviously on the Campus Spine layer and applies to both wired and wireless traffic, as logically represented in the figure below.




Figure 28: MSS Spine enforcement in a traditional campus network

### MSS Spine Services for Wireless Traffic

With tunnel-based wireless forwarding, wireless traffic is handed off by a centralized wireless controller, like the Primary WC in the diagram below, to a pair of Campus Spine switches, using 802.1Q VLANs to segment the different types of traffic. Usually such VLANs have gateways located on the adjacent firewall, however, to provide some symmetry between the tunnel-based and the bridge-based wireless models, it is assumed a more elaborated design where each VLAN gateway is on the switches, isolated in a VRF that has a default route to the adjacent firewall.

For example the following 3 VRFs are used:

VRF name	Purpose
VRF-X	Wireless corporate devices
VRF-Y	Wireless IoT devices
VRF-Z	Wireless guest devices



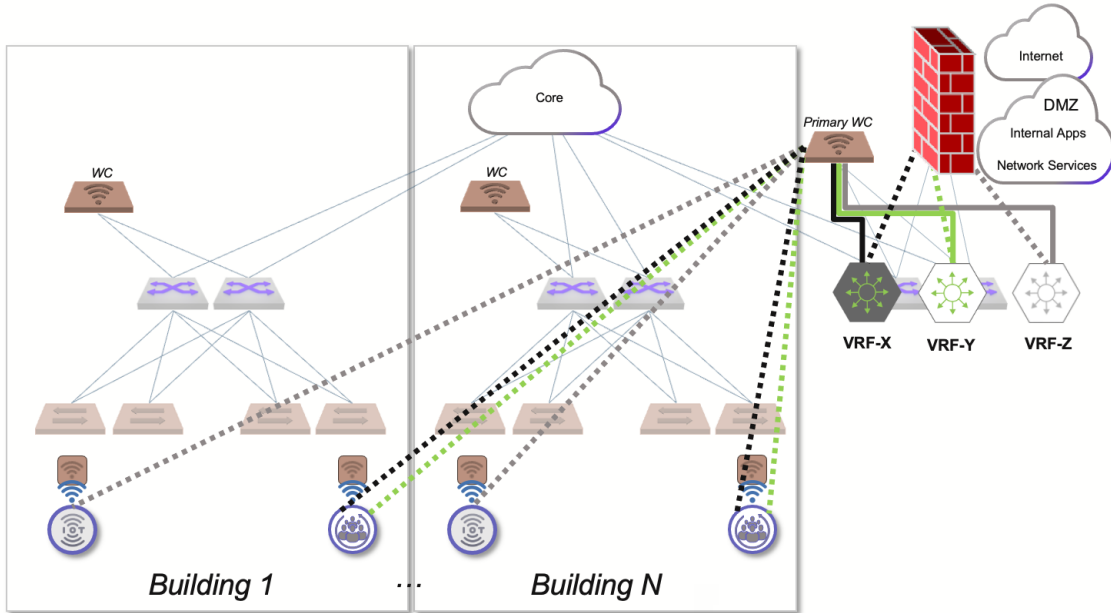


Figure 29: Segmentation of wireless endpoints in tunnel-based architecture

The most common goal of using MSS services on the Campus Spine switches that front-end the wireless controller, is enforcing the traffic to flow exclusively in the N-S direction, as in the following whitelist example simulated for VRF-Z, which segregates the wireless guest traffic:

Source	Destination	Network Service	Enforcement Action	Notes
*	internet	*	FORWARD	N-S wireless guest traffic normally routed to the firewall
*	*	*	DROP MONITOR	default zero-trust rule (preventing lateral communication between wireless guest devices)

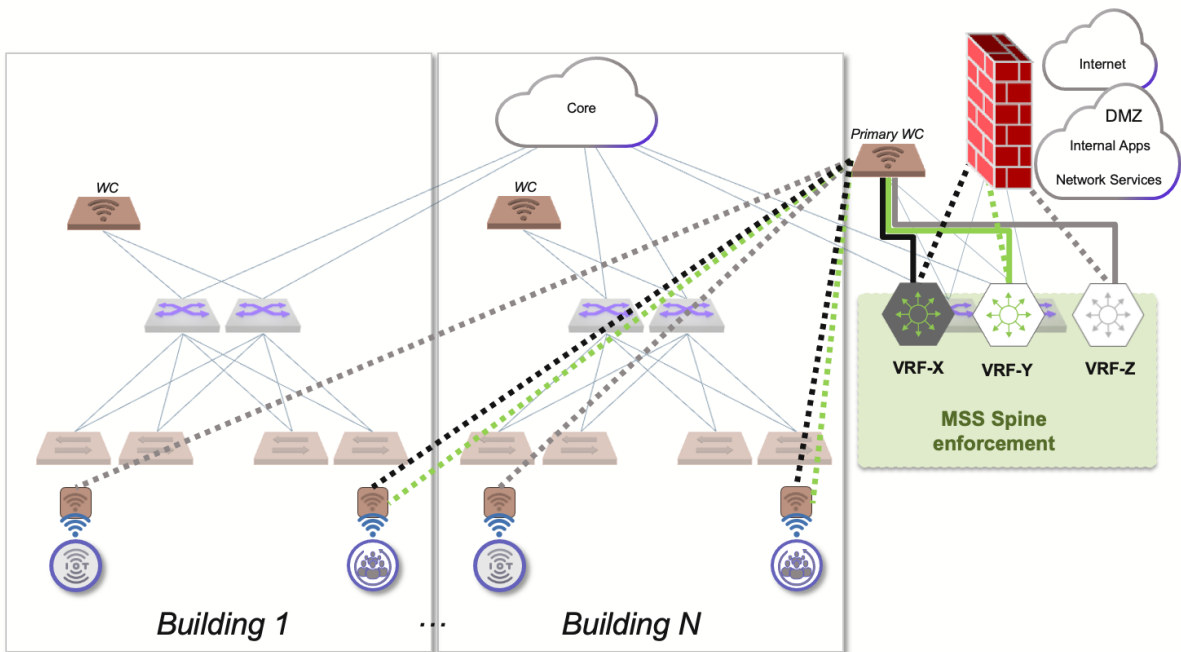


Figure 30: MSS Spine enforcement for wireless endpoints

## MSS Spine Services for Wired Traffic

As for wired traffic, a typical campus network design relies on a two-level multi-VRF strategy to isolate traffic of different wired environments and prevent building-to-building communication. This is identical to what has been discussed in the [previous](#) section and illustrated by the following table and diagram:

VRF name	Purpose
VRF-A	In-band control-plane and management-plane traffic of campus network appliances, like for example wireless access points
VRF-B	Workstations
VRF-C	Printers

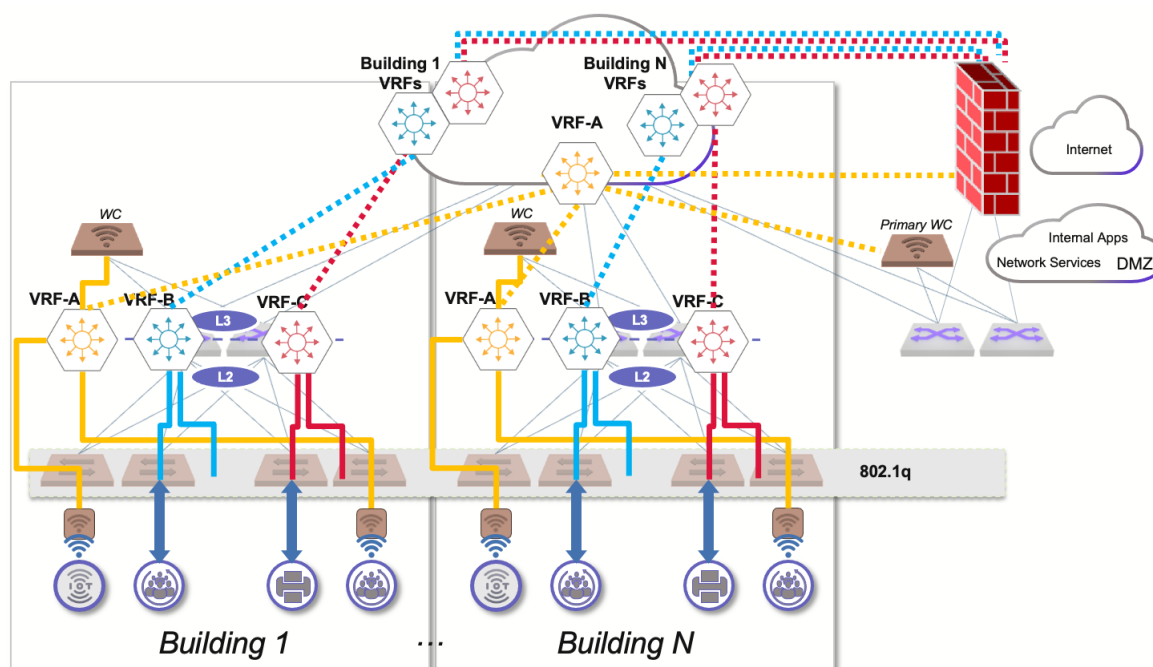


Figure 31: Layer-2 and layer-3 segmentation of wired endpoints

As discussed, using VRF segmentation by itself, other than being inflexible, does not prevent lateral communication within the VRF. The common industry adjustment in this case consists in additionally implementing two isolation mechanisms to either block lateral traffic or hairpin it to the firewall:

1. PVLAN for layer-2 isolation
2. local proxy-ARP with IP ACLs or PBR on layer-2/3 demarcation switches

As it is known that on switches proxy-ARP can be control-plane intensive and ACLs are often limited in scale, Arista has developed the MSS technology in a way that eliminates the need of local proxy-ARP while also compressing ACL and PBR entries with a more efficient and unified engine. This means that the MSS Spine enforcement removes the need of implementing mechanism 2 (local proxy-ARP with ACL and PBR), as illustrated in the following diagram.




Figure 32: MSS Spine enforcement for wired endpoints

MSS Spine Services for both Wireless and Wired Traffic

The combined design of MSS Spine enforcement for wireless and wired traffic is represented by the diagram below.




Figure 33: MSS Spine enforcement for both wired and wireless endpoints

Details about Arista PVLAN implementation can be found in the [References](#) paragraph.

### Integration with Arista tunnel-based wireless solution

The diagram below shows how the discussed solution can be further enhanced by adding the Arista wireless network to the existing design or using it to replace the existing third-party access points.

As explained in the previous section, Arista wireless solution uses a distributed wireless controller function that is implemented in the access point software and does not require external physical appliances or cloud-based software to control how wireless devices connect to the network.

Arista access points support tunnel-based forwarding mode, and use standard VXLAN as tunnel encapsulation. The centralized tunnel termination point is typically located on the Campus Spine pair that provides the MSS Spine enforcement for wireless traffic.

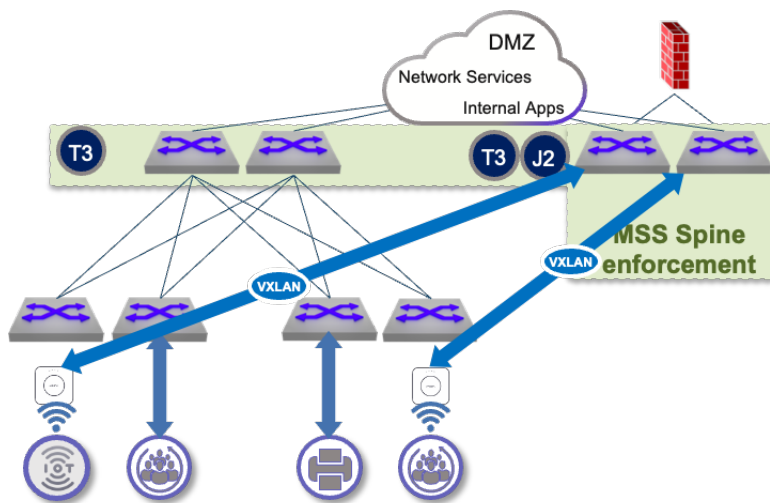


Figure 34: Integration with Arista tunnel-based wireless solution

The Arista wireless architecture reduces the complexity of the solution by avoiding wireless controllers and the control-plane limitations of their architectures, and improves the overall wireless connectivity and mobility performance, as well as the high availability.

Second, it provides to the administrator a single pane of glass for managing and monitoring both wired and wireless networks.

### Extending the solution with Arista Campus Leaf

This chapter describes the option of a greenfield design for an entire Arista campus network, or how the design discussed previously can evolve to include pairs of Campus Leaf switches as wired and wireless access switches.

A campus solution based entirely on Arista switches, referenced in the picture below, can support both overlay and traditional network design options, and with a tunnel-based forwarding approach for wireless traffic, can implement a zero-trust architecture with MSS Spine enforcement.




Figure 35: Extending the solution with Arista Campus Leaf

The obvious advantage of an all-Arista solution is the prospect of using a single unified management system and providing end-to-end visibility for wired and wireless traffic.

Additionally, an all-Arista solution provides an easier and more efficient implementation in case an overlay approach is desirable for wired traffic.

Like for the [MSS Leaf case](#), the PVLAN overlay design variant of the MSS Spine design is based as well on [VXLAN and EVPN protocols](#). And it results as well in a [simplification](#) of the aggregation and core network design, which as a unified underlay network can just be implemented with a single VRF, removing the need of running multiple routing instances and improving the overall layer-3 scale of the design.

Arista overlays offer also an enhancement called E-TREE [add reference link] that works in tandem with PVLAN to force Campus Leaf wired traffic in the same VLAN (VNI) to hairpin to a centralized location, which finally provides a single centralized MSS enforcement option for micro-perimeter segmentation across the entire campus network.

**Insertion of stateful MSS nodes in campus network**

Stateful MSS nodes like ZTX-7250S are used as target devices for monitoring and inspection rule actions. They can support different high-availability modes like active-standby or active-active pairs or scale-out clusters, so they are suitable for providing stateful monitoring or inspection service in a centralized location, for example attached to the Campus Spines located in a DMZ or in each campus building.

The following diagram provides a generic deployment example.

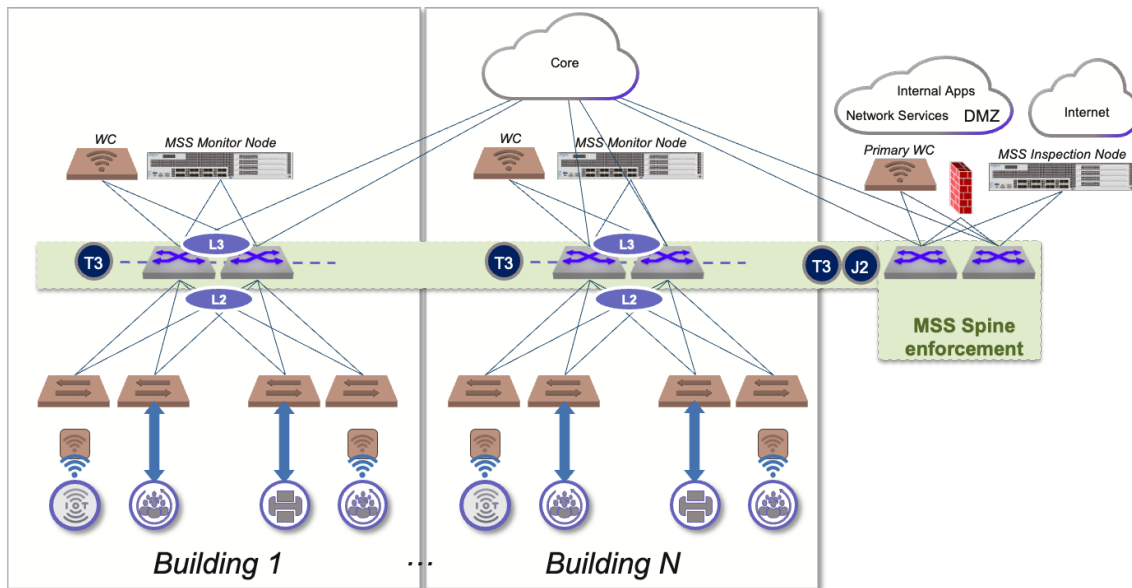


Figure 36: Insertion of MSS stateful nodes

## References

NIST Zero Trust Architecture

<https://www.nist.gov/publications/zero-trust-architecture>

Arista MSS Technical Whitepaper

<https://www.arista.com/assets/data/pdf/Whitepapers/MSS-Segmentation-Technical-WP.pdf>

Arista MSS Datasheet

<https://www.arista.com/assets/data/pdf/Datasheets/MSS-Datasheet.pdf>

Arista ZTX-7250S Datasheet

<https://www.arista.com/assets/data/pdf/Datasheets/ZTX-7250S-MSS-Traffic-Mapper-Datasheet.pdf>

Arista PVLAN Design

<https://www.arista.com/en/support/toi/eos-4-25-0f/14609-support-for-private-vlan>

Arista PVLAN with VXLAN Design

<https://www.arista.com/en/support/toi/eos-4-26-1f/14786-pvlan-vxlan-evpn>

Arista E-TREE with MPLS and VXLAN Design

<https://www.arista.com/en/support/toi/eos-4-24-0f/14481-evpn-e-tree-for-mpls>

Arista Layer-2 Traffic Inspection and Filtering (L2TIF)

<https://wifihelp.arista.com/post/managing-inter-client-communication-on-wifi-access-points>

IEEE 802.11 - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications

<https://ieeexplore.ieee.org/document/9363693>

IEEE 802.1Q - Standard for Local and Metropolitan Area Networks - Bridges and Bridged Networks

<https://ieeexplore.ieee.org/document/10004498>

IEEE 802.1X - Standard for Local and Metropolitan Area Networks - Port-Based Network Access Control

<https://ieeexplore.ieee.org/document/9018454>

### Santa Clara—Corporate Headquarters

5453 Great America Parkway,  
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: [info@arista.com](mailto:info@arista.com)

### Ireland—International Headquarters

3130 Atlantic Avenue  
Westpark Business Campus  
Shannon, Co. Clare  
Ireland

### Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300  
Burnaby, British Columbia  
Canada V5J 5J8

### San Francisco—R&D and Sales Office 1390

Market Street, Suite 800  
San Francisco, CA 94102

### India—R&D Office

Global Tech Park, Tower A, 11th Floor  
Marathahalli Outer Ring Road  
Devarabeesanahalli Village, Varthur Hobli  
Bangalore, India 560103

### Singapore—APAC Administrative Office

9 Temasek Boulevard  
#29-01, Suntec Tower Two  
Singapore 038989

### Nashua—R&D Office

10 Tara Boulevard  
Nashua, NH 03062



Copyright © 2024 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. April 24, 2024 07-0016-01