# Zero Trust Design Guide



Alcatel·Lucent
Enterprise

# Table of contents

# 1. ALE Zero Trust design overview

## 1.1 What Zero Trust means in an ALE network

The concept of Zero Trust is defined in NIST Special Publication 800-207 as an architecture where no implicit trust is granted to assets or user accounts based solely on their physical or network location. Trust must be continuously evaluated, and access must be explicitly authorized.

This Design Guide focuses specifically on the network segmentation and micro-segmentation approach to Zero Trust. It does not attempt to redefine identity governance, endpoint security or application-layer controls. Instead, it addresses how Zero Trust principles are implemented within the network infrastructure using authentication, role-based access control and segmentation.

In this context, Zero Trust means that network access is never granted implicitly. Every device connecting to the network is authenticated, classified and assigned a role before it is allowed to communicate. Simply being physically connected or located inside the enterprise network does not provide any level of trust.

Zero Trust also means that devices sharing the same segment are not automatically allowed to communicate with each other. Lateral movement is restricted by default, and communication is permitted only when it is explicitly required for the assigned role and business use case.

Enforcement is applied at the appropriate layer. The access layer controls whether a device can communicate at all and prevents uncontrolled east-west traffic. The firewall controls communication between segments and toward shared services or the internet. This separation ensures consistent enforcement without embedding complex security logic in the network fabric.

## 1.2 ALE building blocks used in this design

This design relies on ALE components, each with a clearly defined role. The Alcatel-Lucent OmniVista® User Profile and Authentication Manager (UPAM) is used to authenticate users and devices and to provide identity information to the network. OmniVista is used to define roles, access policies and their association with network segments and to consistently push these definitions to the infrastructure. Alcatel-Lucent OmniSwitch® access switches and Alcatel-Lucent OmniAccess® Stellar access points enforce role-based policies at the network edge where devices connect and lateral movement can be controlled. A firewall is used to enforce communication between segments and to provide controlled access to shared services and the internet.

## 1.3 High-level architecture

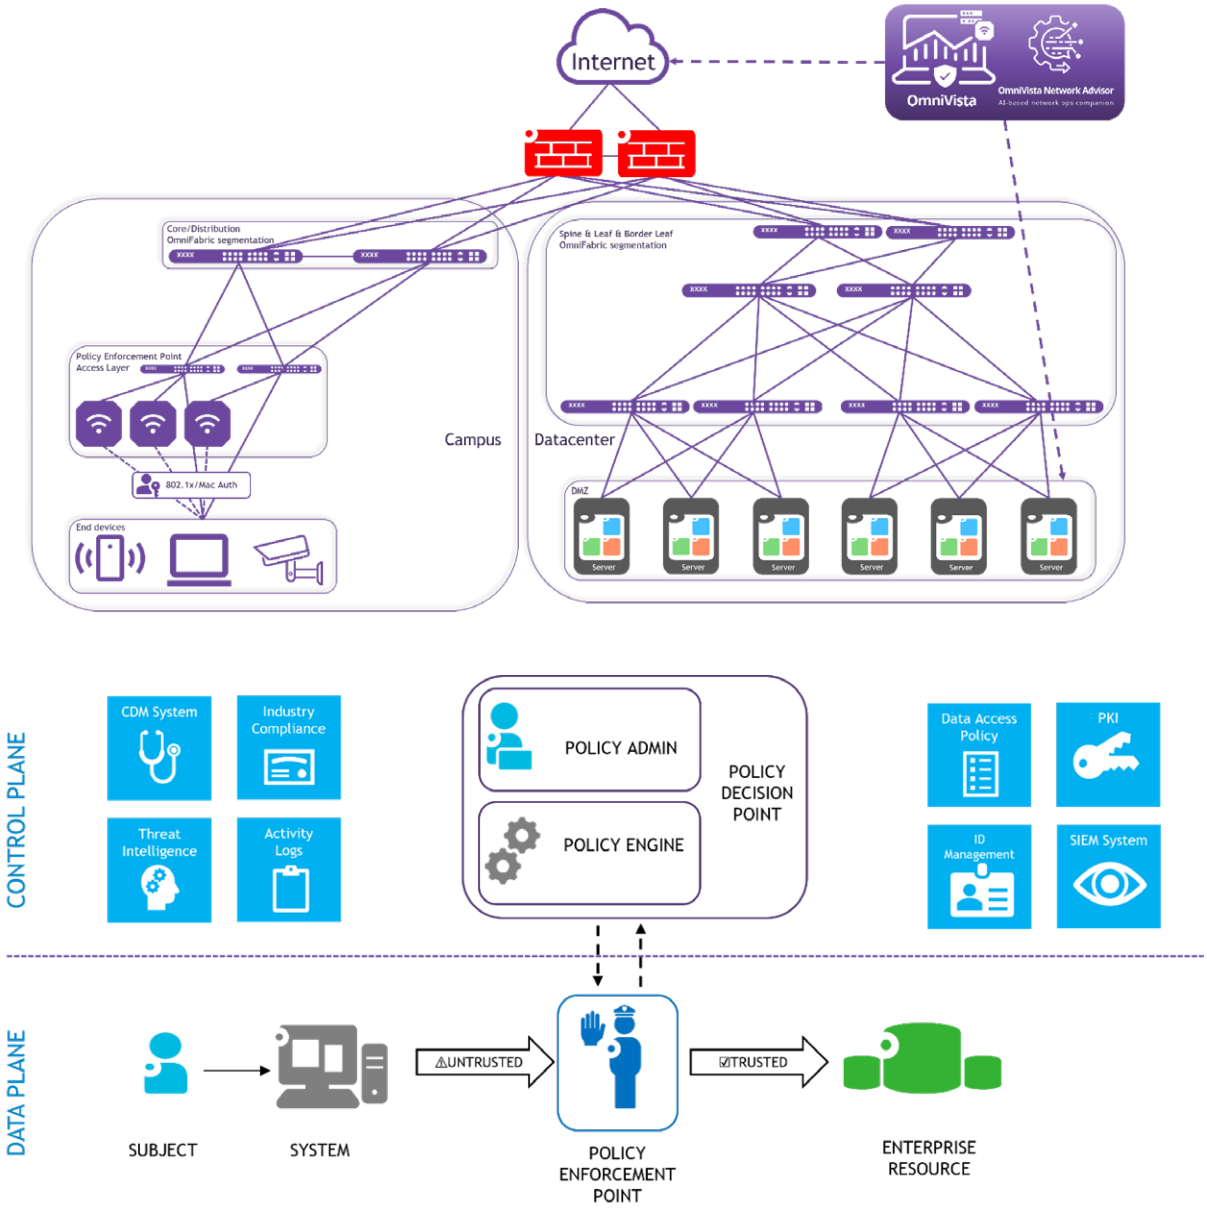


# 2. Brownfield Zero Trust adoption approach

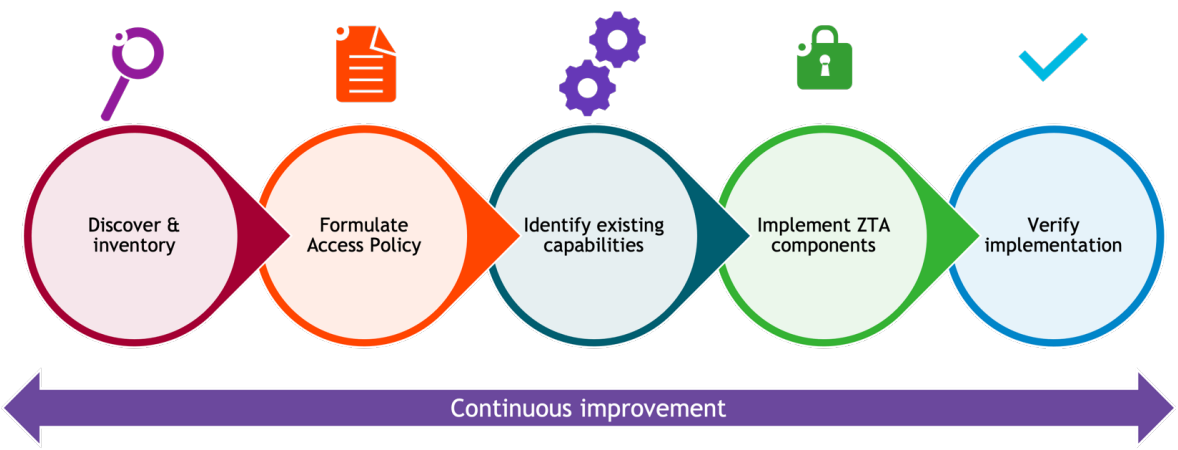

Zero Trust is rarely implemented in a greenfield environment. Most organizations already operate a production network with existing users, devices, applications and operational constraints. The challenge is therefore not to redesign everything from scratch, but to incrementally improve security while preserving business continuity and existing investments.

In such environments, Zero Trust adoption must be approached as an iterative process rather than a single transformation step. The first practical requirement is visibility. Before defining policies, the organization must understand what is connected to the network and why. OmniVista provides inventory and discovery capabilities, including IoT inventory, that allow devices and endpoints to be identified and classified. This inventory serves as the foundation for all subsequent design decisions.

Once assets are identified, access policies can be formulated based on actual business use cases. Devices or services that do not support a clear business function should be questioned, as they increase the attack surface without providing value. For required assets, policies should be defined explicitly, starting with the most critical users, devices and applications.

Existing security capabilities should then be leveraged where possible. Zero Trust does not require replacing all tools, but rather integrating them coherently. This typically means combining strong authentication, role-based access at the edge, segmentation and firewall enforcement.

Policy enforcement should be introduced progressively. During initial phases, policies can be deployed in a non-blocking or observation mode, using counters and logs to validate behavior before enforcement is activated. This reduces the risk of disrupting business processes and allows policies to be refined based on real traffic patterns.

Zero Trust is not a one-time project. As applications, devices and threats evolve, policies must be reviewed and adjusted. The objective is not to reach a theoretical "optimal" state immediately, but to continuously reduce implicit trust and limit exposure over time in a controlled and operationally sustainable way.

# 3. Identity and role model

## 3.1 Authentication model

The most secure authentication mechanism is 802.1X. It provides a consistent way to authenticate both wired and wireless devices before granting any network access and allows access decisions to be tied to identity rather than physical connectivity.

For corporate-managed endpoints, authentication is typically split between machine authentication and user authentication. Machine authentication is used to establish device trust and allow the operating system to interact with Active Directory, while user authentication determines the final access rights once a user session is established. This separation allows access to evolve based on authentication state without granting unnecessary privileges too early.

IoT and other headless devices often do not support 802.1X. For these devices, MAC-based authentication is used as a fallback, or optionally device fingerprinting to improve classification and reduce reliance on static identifiers. These mechanisms are used only when stronger authentication methods are not available.

BYOD and Guest access may use captive portal–based authentication, particularly during onboarding and for low-assurance access scenarios. Captive portals provide a practical mechanism for user identity verification, policy acceptance and role assignment for unmanaged devices.

For BYOD users requiring persistent or higher-assurance access, onboarding can be used to establish a managed association between the user and the device, enabling revocation when the device is lost, the employee leaves the company or policy changes. In such cases, authentication may transition to mechanisms that provide stronger enforcement characteristics, such as identity-bound or device-bound credentials.

Guest access remains intentionally limited and time-bound, with authentication mechanisms and network segmentation reflecting the lower trust level associated with unmanaged and non-corporate devices. The network-level enforcement of Guest isolation can be implemented in two ways, and the choice should be made explicitly at design time. In the standard model, Guest devices are placed on a dedicated VLAN and all traffic is routed through the firewall, where role-based policies restrict access to the captive portal during onboarding and to the proxy after authentication. This model is operationally simple but relies entirely on correct firewall policy to prevent Guest traffic from reaching internal segments. A stronger alternative is to place the Guest VLAN in a separate VRF on the core or distribution layer so that no routing path to internal segments exists at the network level regardless of firewall state. This eliminates the risk of accidental reachability due to policy misconfiguration and is recommended where regulatory requirements or internal policy mandate strict network-level separation. In either case, the firewall policy tables defined in section 6.2 apply. For environments where Guest traffic must be centralized at a specific aggregation point, GRE-based tunneling as described in section 4 is an additional option.

## 3.2 Authentication flow

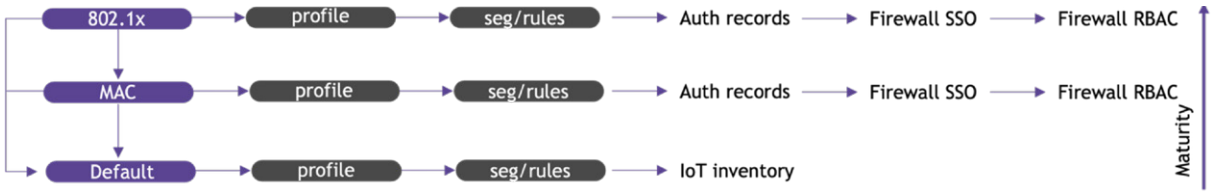


When supported, authentication should be performed using 802.1X. Upon successful authentication, the RADIUS server returns a profile that determines the segmentation and policy rules applied to the device. Authentication events generate accounting information that can be shared with other enforcement points, such as a firewall, to enable identity-aware policies beyond the access layer.

If 802.1X authentication is not attempted or does not succeed, MAC-based authentication can be used as a fallback. While MAC authentication provides a lower level of assurance, it still allows the device to be associated with a role and a controlled policy set rather than being treated as completely unknown.

If no authentication method results in a role assignment, the device is mapped to a default profile. In the early stages of deployment, this default profile typically preserves existing network behavior while providing visibility and logging. As the design matures, the default role can be progressively restricted.

This flow is designed to support a phased and non-disruptive deployment. Authentication policies can initially operate in a permissive, monitoring-focused mode and later transition to stricter enforcement as confidence in role definitions and policies increases.

## 3.3 Authentication path

In this design, network access control is enforced at the access edge using the strongest method each endpoint can support. What varies across deployments is where the RADIUS decision is made and which system provides the final authorization result. The choice of authentication path has a direct impact on operational complexity, scalability, and integration with existing identity systems.

### UPAM as RADIUS proxy to an external RADIUS server

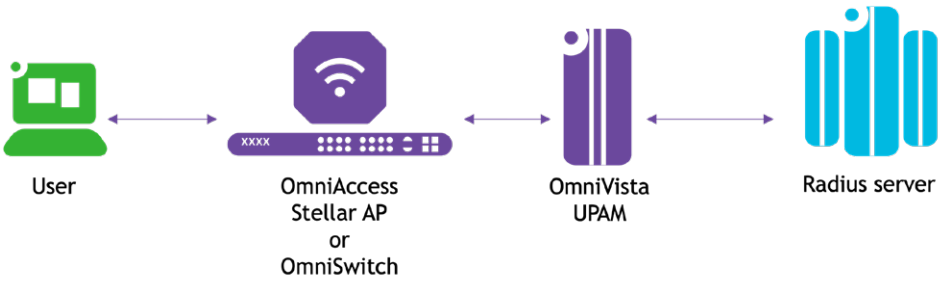


This is the most common model in enterprise environments. The access switch or access point sends the 802.1X authentication request to UPAM using RADIUS. UPAM does not make the authentication decision itself but acts as a proxy and forwards the request to an external RADIUS server, typically Microsoft NPS integrated with Active Directory. The external RADIUS validates the credentials and returns the result to UPAM, which then applies role mapping and policy enforcement consistently across wired and wireless access. Using UPAM as a proxy significantly simplifies operations: only UPAM needs to be declared as a Network Access Device on the external RADIUS server, while all OmniSwitches and OmniAccess Stellar Access Points managed by OmniVista are automatically handled.

### UPAM as RADIUS server with an external identity provider (Azure Entra ID)

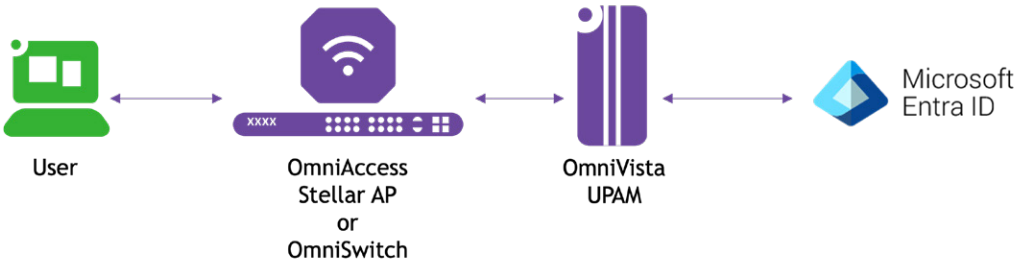


In this model, UPAM acts as the primary RADIUS server for the network. User or device authentication is performed using certificates or credentials, and UPAM delegates identity validation to an external identity provider: Azure Entra ID. The access switch or access point sends the RADIUS request to UPAM, UPAM validates the identity through the external IdP, and returns the authorization decision along with role information. This model centralizes network access control in OmniVista while allowing organizations to leverage cloud-based identity without exposing the network directly to external services.

### Direct authentication to an external RADIUS server (without UPAM)

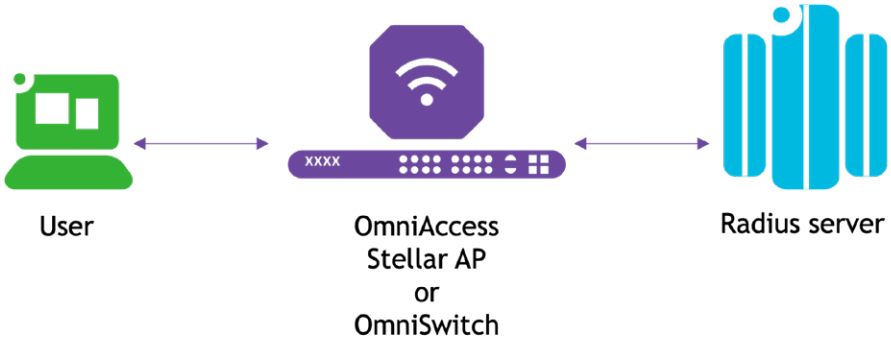


In this model, access switches and access points communicate directly with an external RADIUS server. Authentication and authorization decisions are made entirely outside OmniVista. This approach introduces operational overhead, as every access device must be individually defined as a Network Access Device on the RADIUS server. It also limits centralized visibility and policy orchestration from OmniVista, making consistent role enforcement across wired and wireless more difficult to maintain at scale.

UPAM can also operate with its local user database, but this is generally reserved for small or isolated environments. In larger enterprise deployments, an external identity provider is typically preferred to align network access with corporate identity lifecycle and governance.

Application Notes are available on Spacewalkers.com describing validated integrations with external RADIUS servers and identity providers.

## 3.4 Role definitions used in this guide

The Employee role represents an authenticated user session on a corporate-managed device. The Employee machine role represents the same device when it is authenticated using its machine identity, prior to user authentication. BYOD refers to personally owned devices accessing limited internal resources and the internet. Contractor refers to third-party users accessing the network under restricted conditions defined by company policy. Guest refers to unauthenticated or lightly authenticated users with internet-only access.

IoT roles are defined by device function rather than ownership. IoT Cameras represent video surveillance devices communicating with a video management system. IoT Sensors represent telemetry devices publishing data to an IoT broker. IoT HVAC represents building control devices communicating with a building management system.

These role definitions are used consistently across authentication, segmentation, and policy examples to avoid ambiguity and ensure a common understanding of access intent.

## 3.5 Securing access point connectivity

Access points are often installed in publicly accessible areas, where an attacker could physically connect to the Ethernet link or attempt to intercept traffic, making it essential to protect against unauthorized access and man-in-the-middle attacks at the access edge.

OmniAccess Stellar Access Points (AP) support an authentication-based uplink model where the access switch validates the AP identity using 802.1X, and the link can then be secured through MACsec encryption. This ensures that an AP cannot simply be connected to an access port and gain network connectivity without authorization.

On the wired edge, this is implemented through Secure AP Mode, which ties AP authorization directly into the same access control framework used for user authentication.

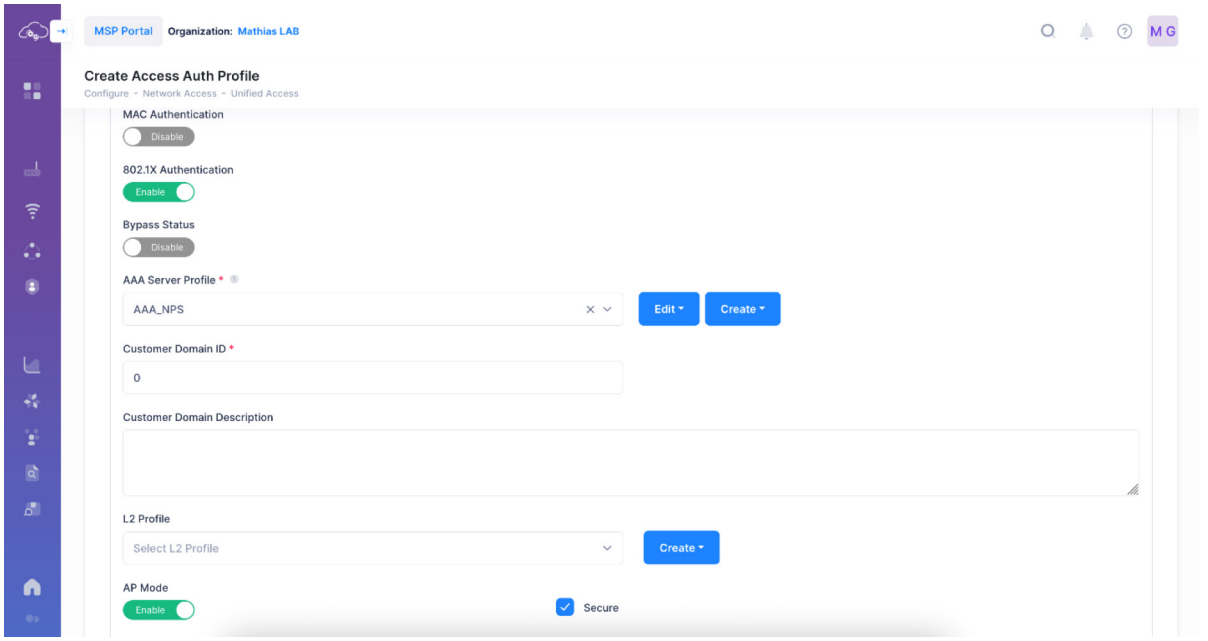


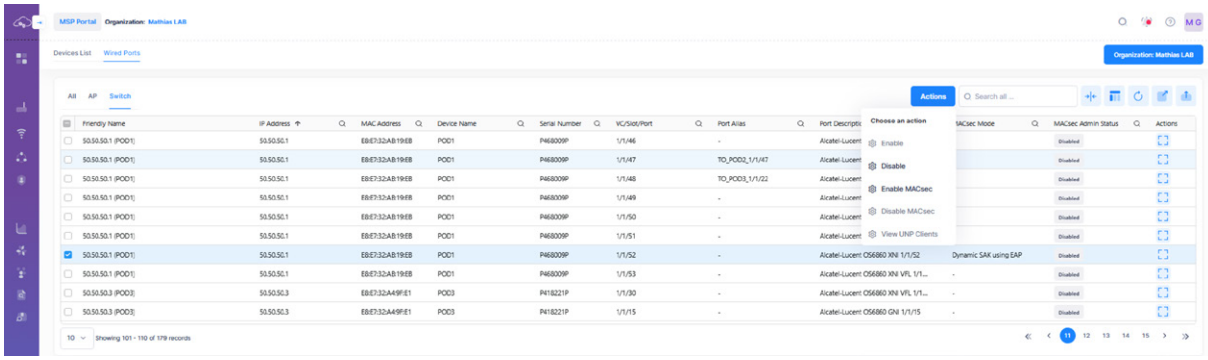

Once enabled, OmniVista automatically pushes the corresponding secure uplink configuration to the switch port using a dedicated UNP port template.

Example of resulting switch configuration:

```
unp port-template “MACsec”
unp port-template “MACsec” aaa-profile “MACsec”
unp port-template “MACsec” classification trust-tag ap-mode secure
unp port-template “MACsec” 802.1x-authentication
unp port-template “MACsec” 802.1x-authentication pass-alternate “Stellar-AP”
unpport 1/1/1 port-type bridge
unpport 1/1/1 port-template “MACsec”
interfaces port 1/1/1 macsec mode dynamic radius encryption
interfaces port 1/1/1 macsec admin-state enable
```

Note: A valid MACsec license is required on the switch to enable MACsec functionality.

This ensures that the AP uplink is authenticated before being accepted and that the AP is placed into its expected role during the bootstrapping phase.

On the AP side, the uplink authentication and MACsec encryption are enabled in the provisioning profile.

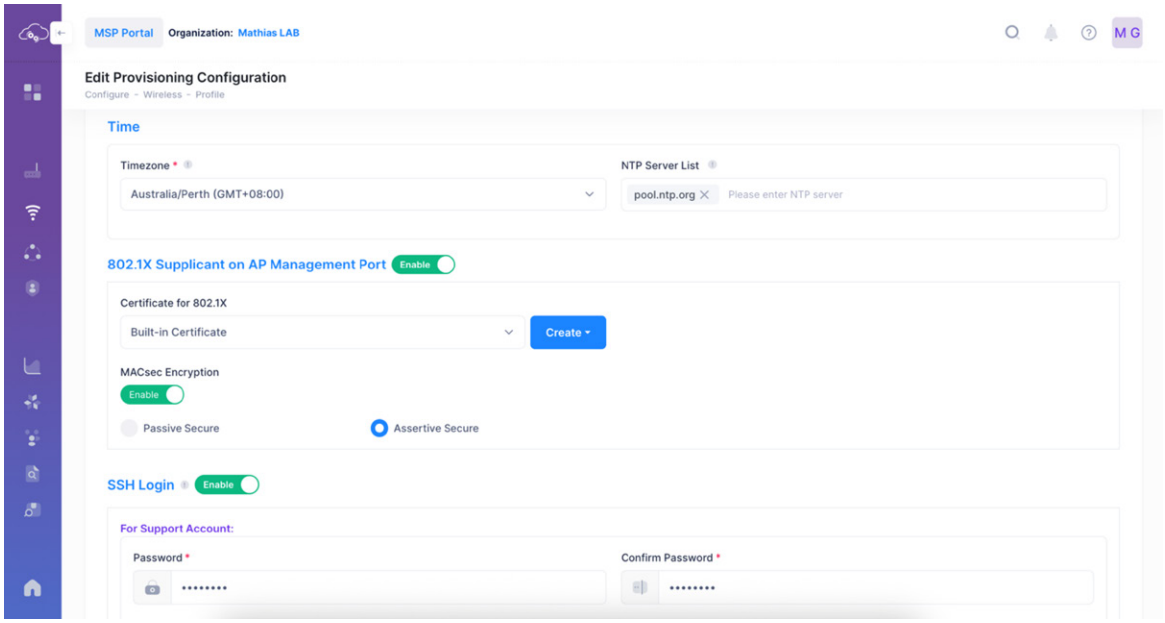


## 4. Segmentation model

Zero Trust can be implemented through different complementary approaches, including identity governance, endpoint security, application controls and network enforcement. This Design Guide focuses specifically on the network segmentation approach: using authentication-driven roles, macro-segmentation between major domains, and micro-segmentation at the access edge to prevent implicit trust and limit lateral movement. Other pillars of Zero Trust remain important but are outside the scope of this document.

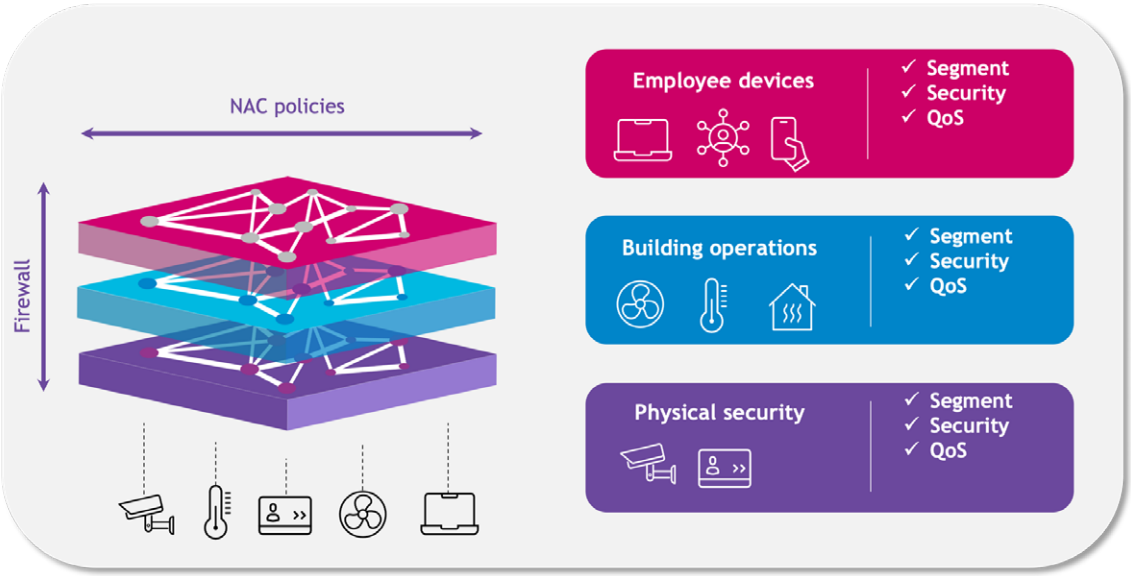

### 4.1 Macro-segmentation model

Macro-segmentation in this design is implemented using VLAN-based or SPB-based segmentation. Each major role category is assigned to its own segment, ensuring that users, guests, contractors and different classes of IoT devices are separated at the network level.

Communication between segments is not permitted directly within the network fabric. All inter-segment traffic is routed through the firewall, where access is explicitly controlled based on destination, service and, where available, identity context. This approach keeps segmentation simple, predictable and consistent across the network, while centralizing inter-segment policy enforcement at a single control point.

### 4.2 Micro-segmentation at the access layer

Within a given segment, devices are not implicitly trusted. Even when endpoints share the same VLAN or fabric service, unrestricted lateral communication is not allowed. Blocking lateral movement inside a segment reduces the impact of compromised devices and prevents attackers from moving freely between endpoints that happen to share the same network location.

This control is enforced at the access layer, where devices first connect to the network and where identity and role information is available. Enforcing policies at this point allows traffic to be restricted before it enters the network fabric, keeping security decisions close to the source and limiting unnecessary propagation of unwanted traffic.

Role-based policies are not enforced at the distribution or core layers. These layers are used strictly for forwarding and routing and must remain free of access control logic. Centralizing enforcement at the access layer avoids inconsistent behavior, simplifies operations and prevents security rules from being duplicated or applied in places where identity context is no longer available.

## 5. Authentication and role enforcement constructs in OmniVista

Authentication and role enforcement in OmniVista rely on a clear separation of responsibilities between three configuration objects: the AAA Server Profile, the Access Authentication Profile and the Access Role Profile. Each object has a distinct purpose, and understanding how they relate is essential to designing predictable and scalable access control.

The **AAA Server Profile** defines where authentication decisions are made. It represents the external identity source used by the network, typically UPAM, Active Directory via NPS or another external RADIUS. This profile specifies the authentication and accounting servers used for 802.1X, MAC authentication or captive portal flows. It is intentionally reused across multiple access policies to ensure consistent authentication behavior.

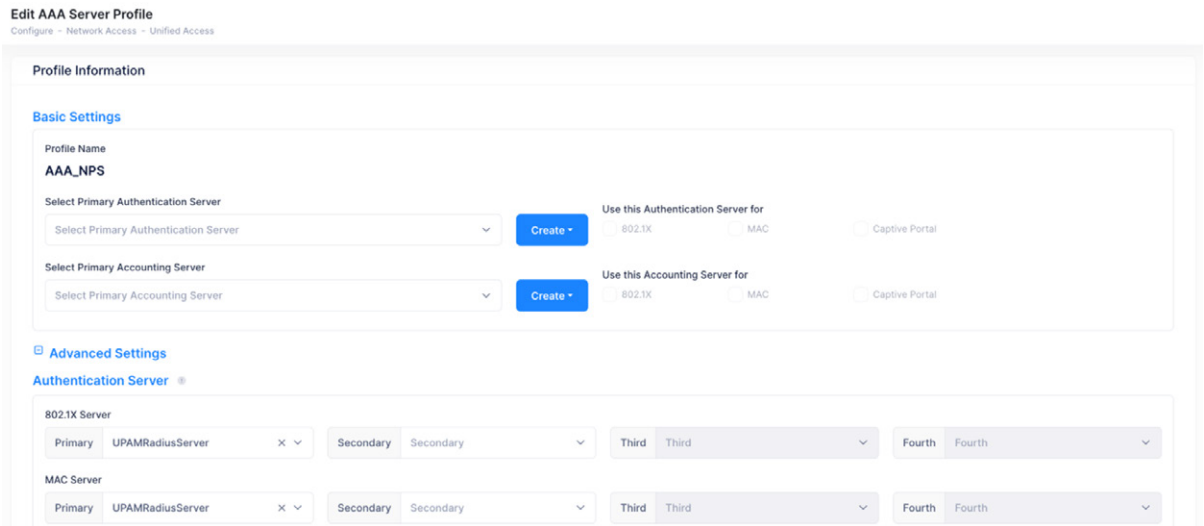

The **Access Authentication Profile** defines how devices authenticate on the network. It controls which authentication methods are enabled on a wired port, such as 802.1X, MAC authentication, or bypass mechanisms. This profile also determines how the switch or access point behaves during authentication events, including fallback behavior and interaction with the AAA Server Profile.

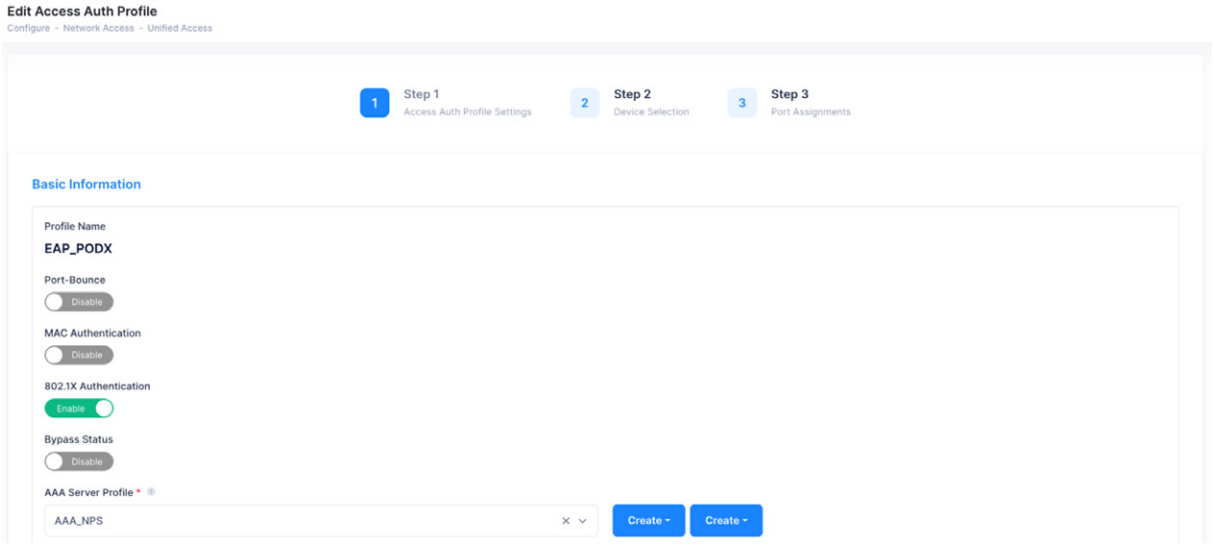


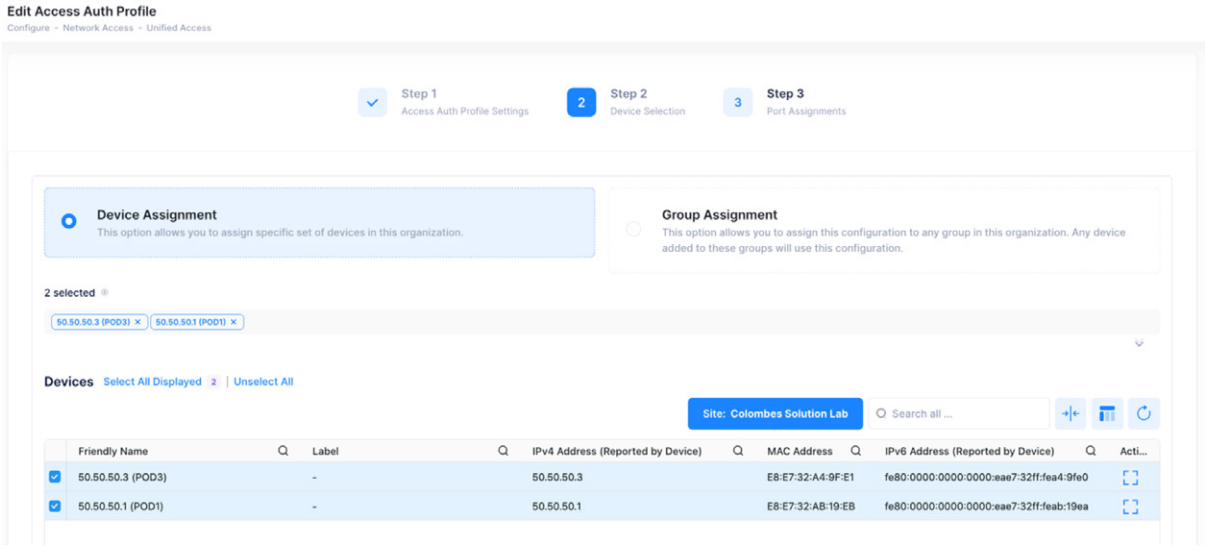


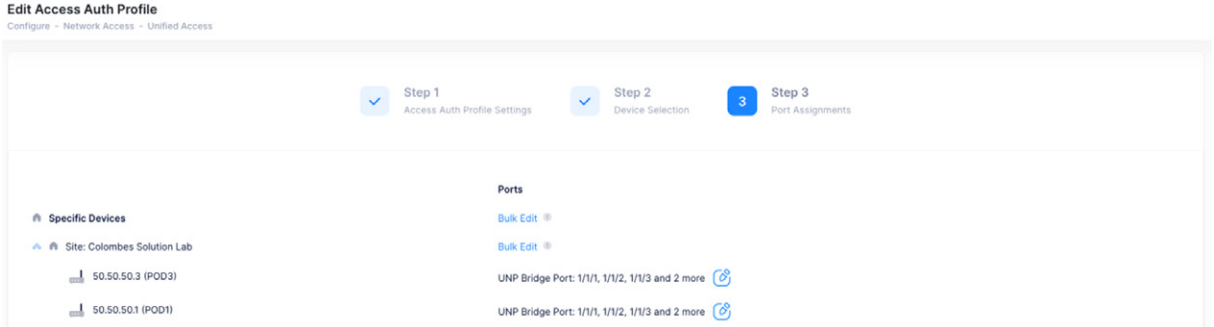

For wireless access, authentication is defined at the **SSID** level rather than through an Access Authentication Profile. The SSID specifies the AAA servers and authentication behavior.

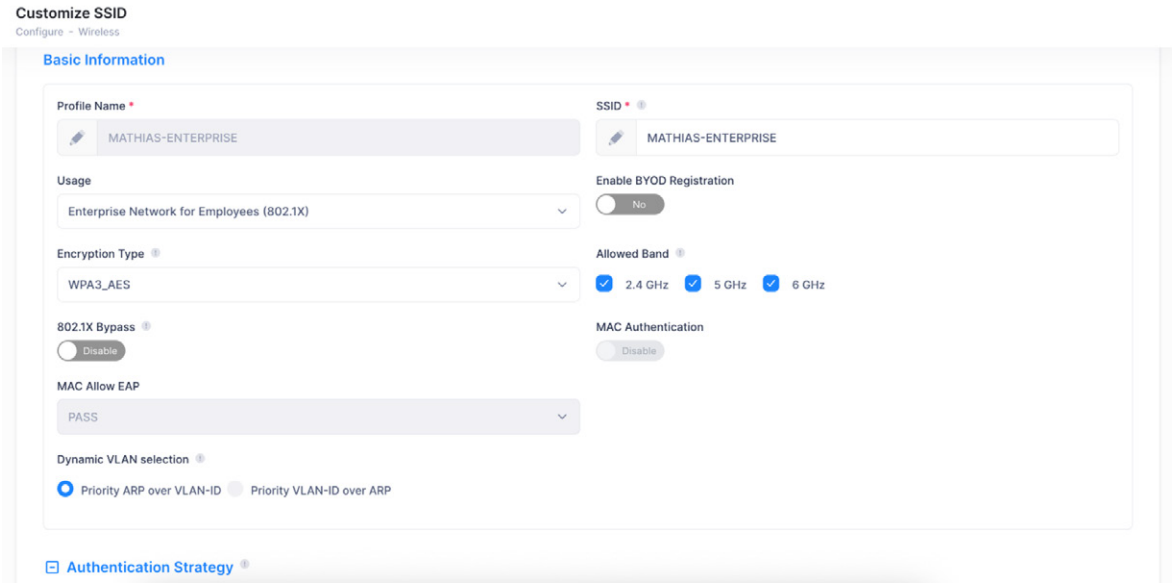


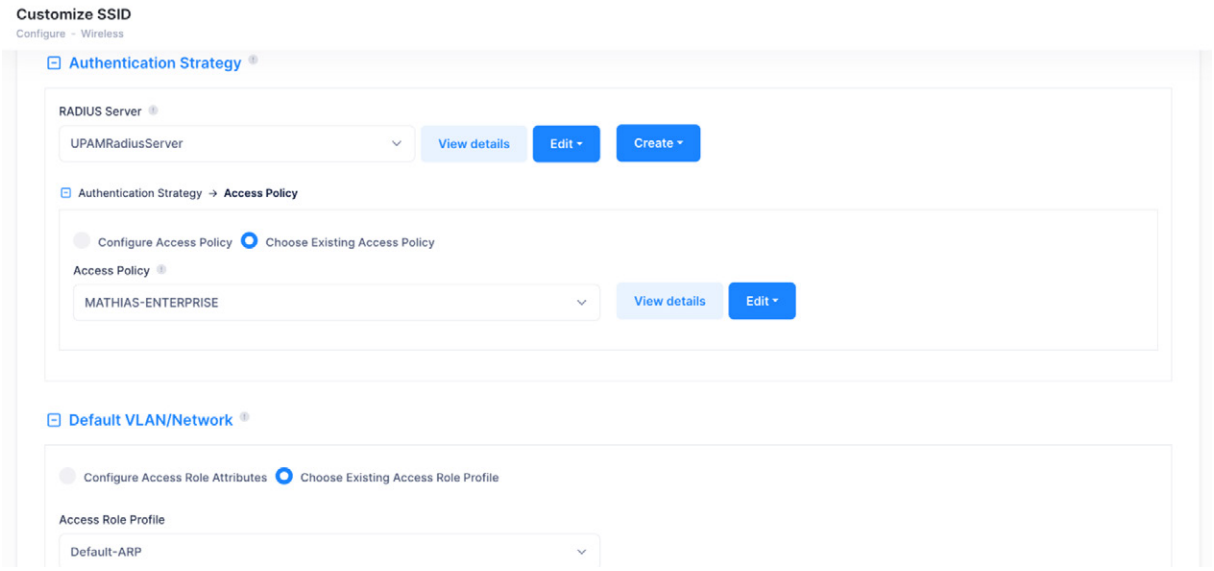


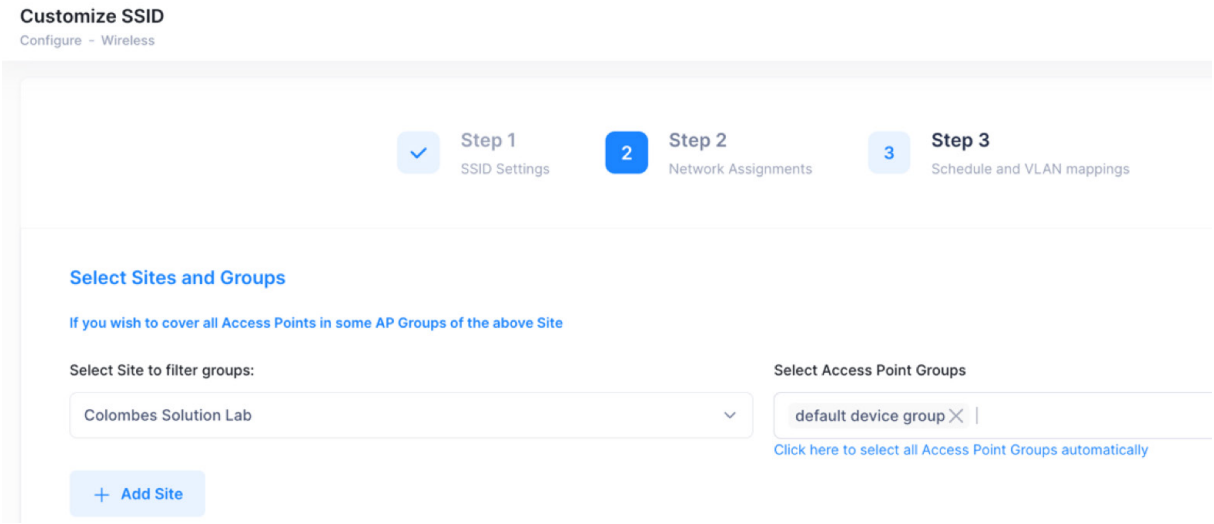

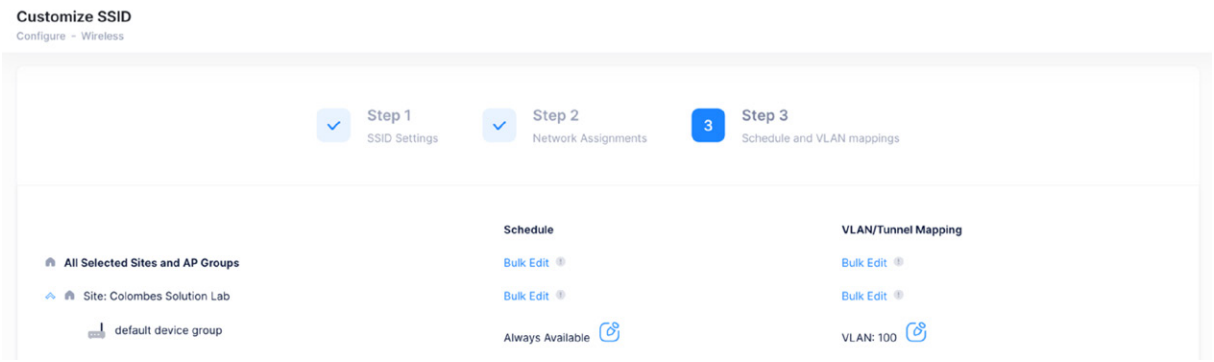


The **Access Role Profile** defines what happens after authentication succeeds. It represents the role assigned to the user or device and ties identity to enforcement. An Access Role Profile associates a Unified Policy List (QoS and ACL rules) with a network placement such as a VLAN or service. This is where segmentation and access restrictions are applied in a deterministic way.

The same role can be reused consistently across multiple devices and sites while being mapped to different VLANs or services depending on the local design, without changing the policy itself. This allows identity-based access rules to remain stable even when segmentation differs between locations.

There is also no strict one-to-one relationship between roles and segments. Multiple roles can be placed into the same VLAN or service while still receiving different policy enforcement through their associated Unified Policy Lists. This makes it possible to separate users or device categories logically, even when they share the same underlying network segment.

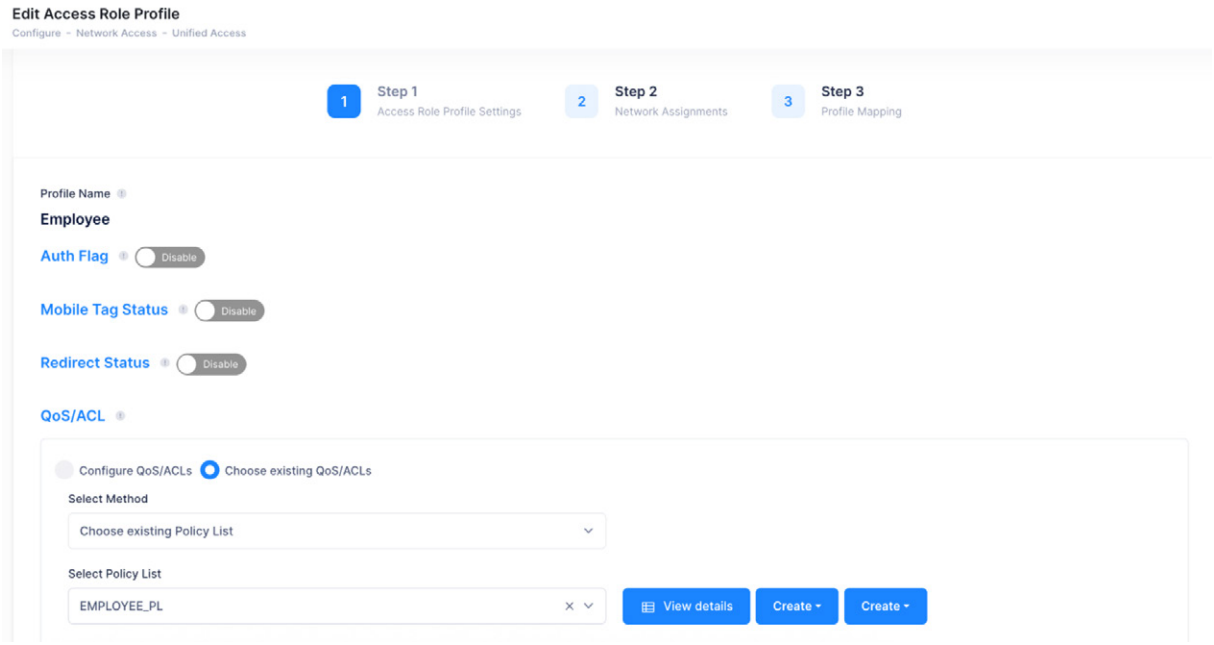

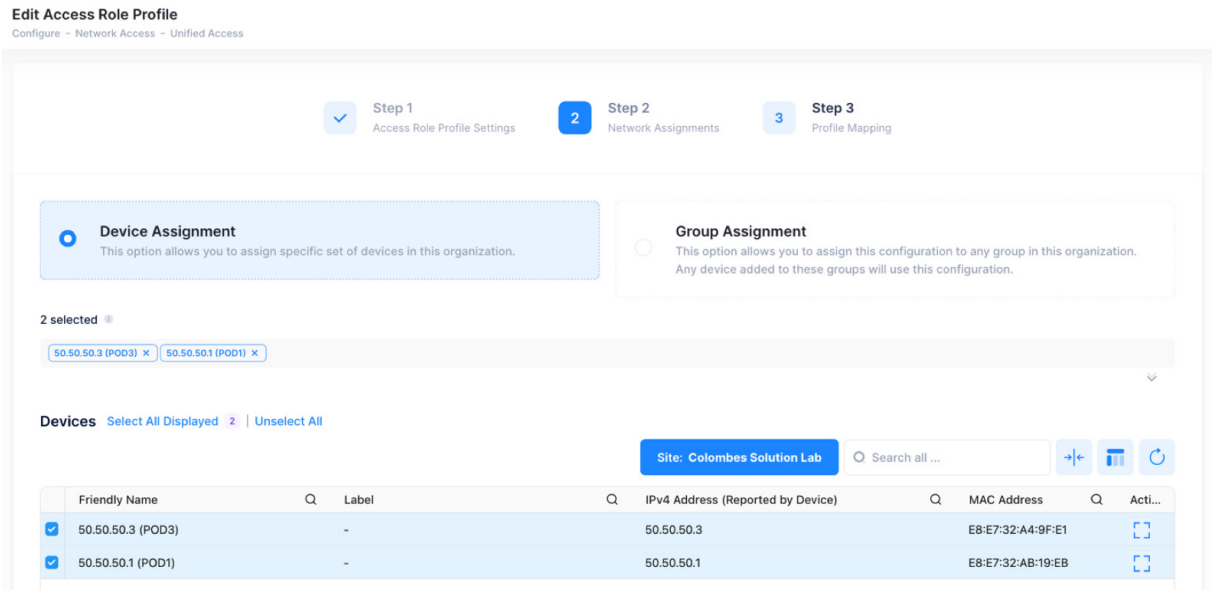


Associated with a VLAN:

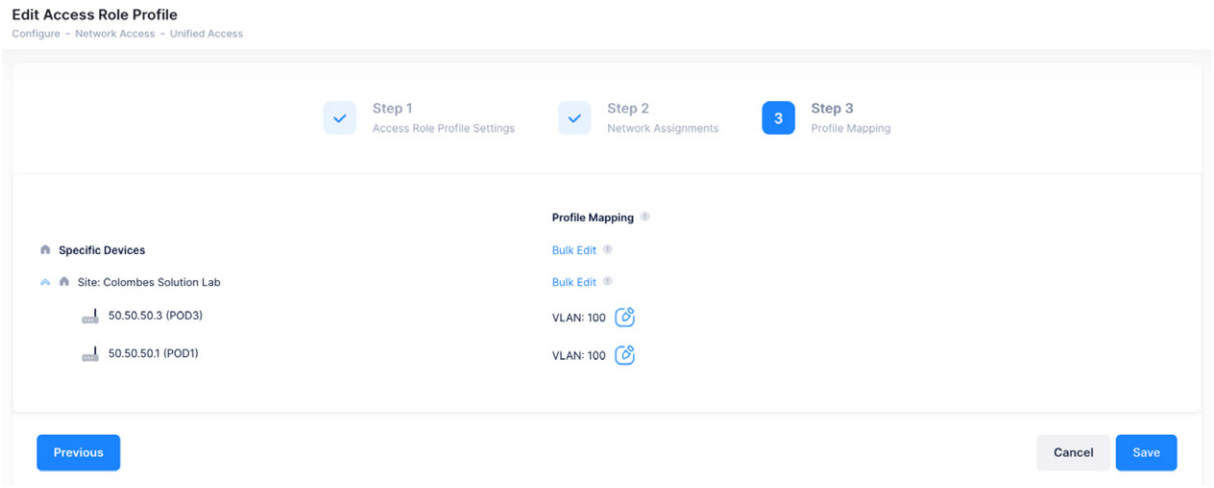


Or with an SPB service:

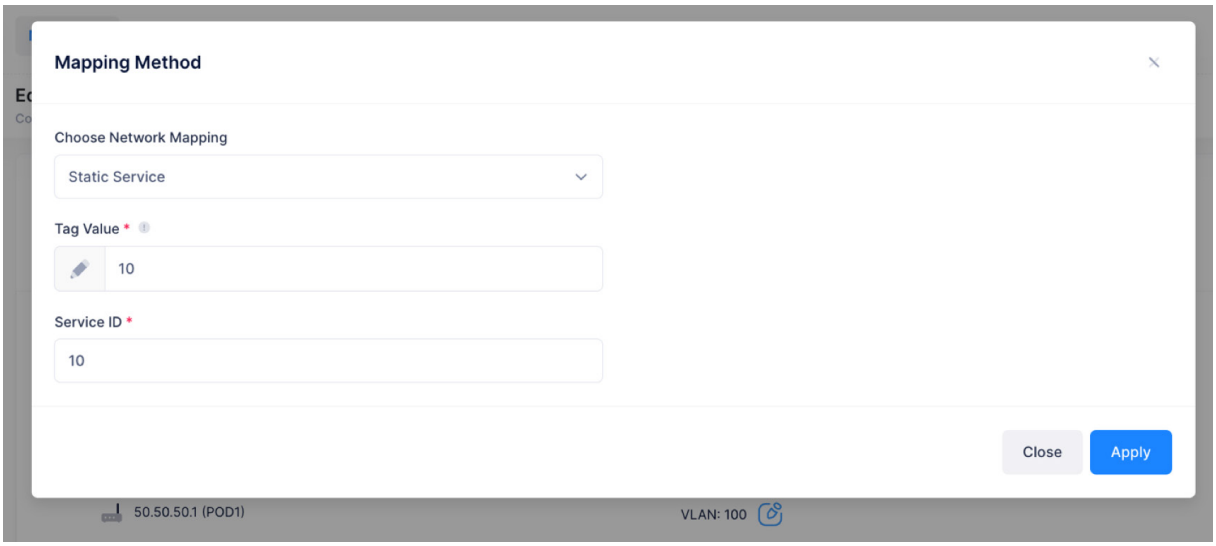

Or with a GRE Tunnel:

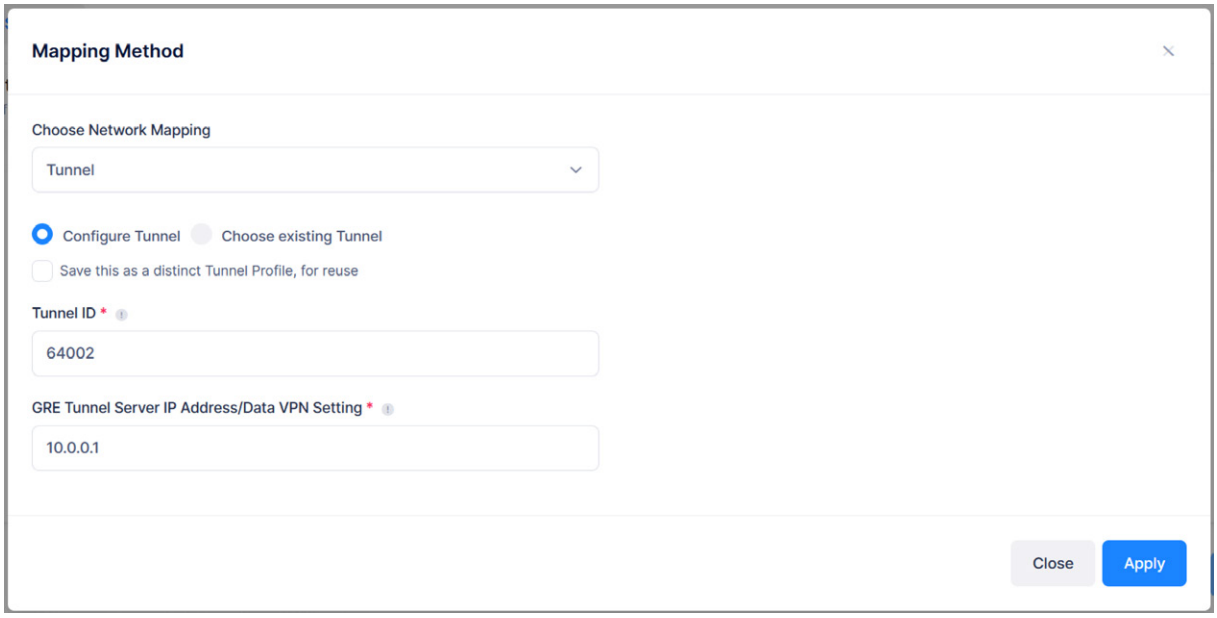


This can be useful when migrating from a controller-based architecture to tunnel all SSIDs to a central switch, replicating the behavior previously provided by the controller.

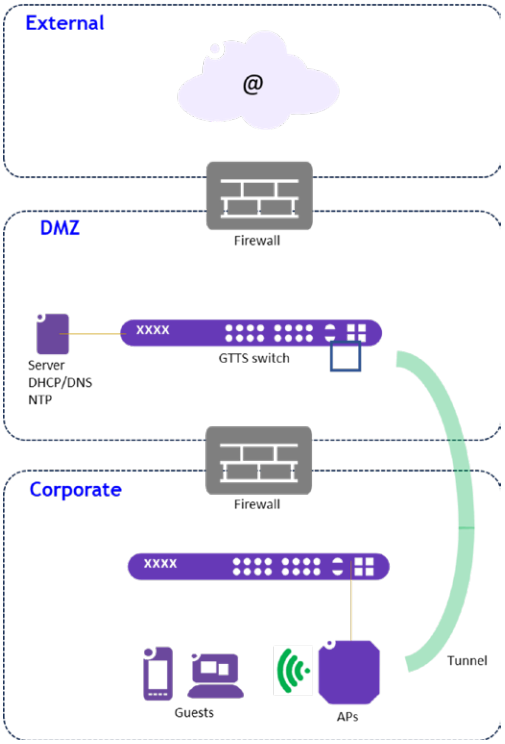


For specific use cases such as tunneling guest traffic to the core, GRE can be used to avoid mixing guest traffic with other user traffic and to maintain strictly limited, internet-only access. For more details, an application note on Guest Traffic Tunneling Service (GTTS) is available in the References and additional resources section of this document.

Example of the resulting configuration on an OmniSwitch:

```
unp profile “Employee” qos-policy-list “EMPLOYEE_PL”
unp profile “Employee” map vlan 100
unp port-template “EAP_PODX” direction both aaa-profile “AAA_NPS” default-profile
“Restricted” admin-state enable
unp port-template “EAP_PODX” 802.1x-authentication
unp port-template “EAP_PODX” 802.1x-authentication supp-timeout 10
unp port 1/1/1-4 port-type bridge
unp port 1/1/1-4 port-template “EAP_PODX”
unp port 1/1/6 port-type bridge
unp port 1/1/6 port-template “EAP_PODX”
```

Access Role Profiles in OmniVista are implemented on the switch using User Network Profiles (UNP). Despite the historical name, UNP should be understood as Universal Network Profiles: they are not limited to users and can be applied to endpoints of any type, including servers, IoT devices and virtual machines.

A UNP represents a role from the switch perspective. It combines network placement (VLAN or service mapping) with a policy list that enforces access control and QoS.

# 6. Access micro-segmentation design

## 6.1 L2 policy logic

| Rule name | Match | Action |
|---|---|---|
| L2 destination MAC | dst = GW MAC | Allow |
| (Optional) peer exception | Rainbow, RTP, etc. | Allow |
| Default | any | Deny |

At the access, the same Layer 2 policy logic is applied. This logic is intentionally simple and deterministic and is designed to prevent uncontrolled communication between endpoints connected to the same segment.

All IP traffic destined to the default gateway MAC address is allowed so that devices can reach routing and higher-layer enforcement points. This ensures that north-south communication remains functional for all roles.

All other traffic is denied by default. This blocks lateral communication between devices within the same segment and prevents endpoints from communicating directly with each other without passing through a controlled enforcement point.

An optional peer-to-peer exception may be introduced for specific use cases where limited east-west communication is required. Such exceptions must be narrowly scoped and justified by functional requirements.

## 6.2 Policy configuration example

The policy list can be created in OmniVista and automatically pushed to all OmniAccess Stellar APs and OmniSwitches.

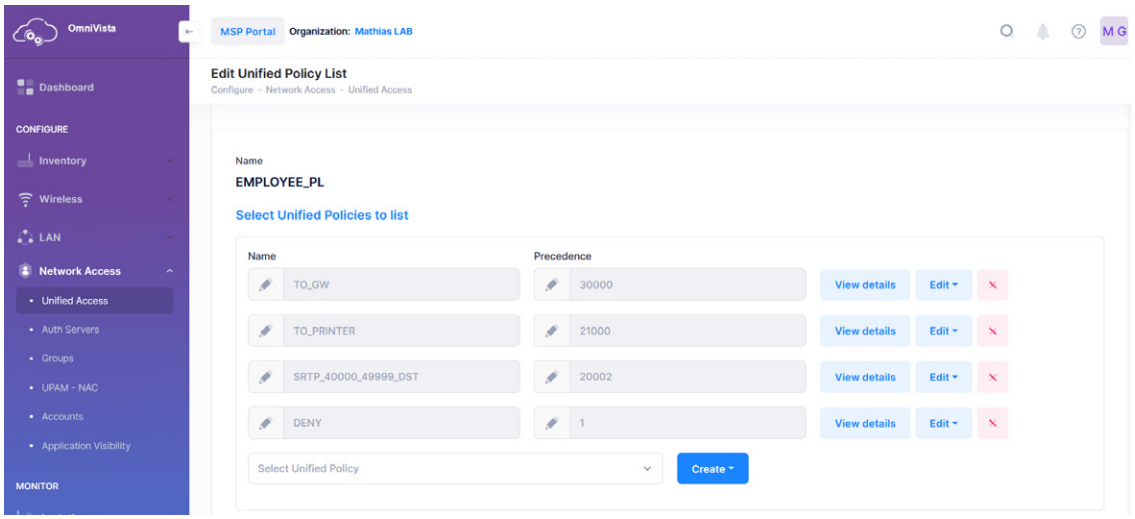


Example of switch policy configuration:

```
policy condition __TO_GW destination mac 00:11:22:33:44:55
policy condition __TO_PRINTER destination mac AA:BB:CC:DD:EE:FF ip-protocol 6
destination ip-port 443
policy condition __SRTP_UDP_40000_49999_DST ip-protocol 17 destination ip-port
40000-49999
policy condition __DENY source ip Any destination ip Any

policy action ALLOW disposition accept
policy action DROP disposition drop

policy rule TO_GW precedence 30000 condition __TO_GW action ALLOW log  no
default-list
policy rule TO_PRINTER precedence 21000 condition __TO_PRINTER action ALLOW log
no default-list
policy rule SRTP_40000_49999_DST precedence 20002 condition __SRTP_
UDP_40000_49999_DST action ALLOW log  no default-list

policy rule DENY precedence 1 condition __DENY action DROP log  no default-list

policy list EMPLOYEE_PL type unp enable
policy list EMPLOYEE_PL rules FROM_GW SRTP_40000_49999_DST DENY
```

The policy list attached to a role is enforced directly in hardware on the access switch. These policies are stateless and are evaluated at packet ingress. Each packet is processed independently as it enters the switch port, without any notion of session or flow tracking.

This has important design implications. Policy rules must always be written from the perspective of the traffic originating from the device to which the policy is applied. For example, when a client device communicates with a web server on TCP port 80, the policy applied to the client role must explicitly allow TCP port 80 as a destination, in addition to basic infrastructure traffic such as ARP and DHCP. If the server is also protected by a policy list, its own policy must explicitly allow TCP port 80 as a source, because the policy is evaluated on traffic leaving the server.

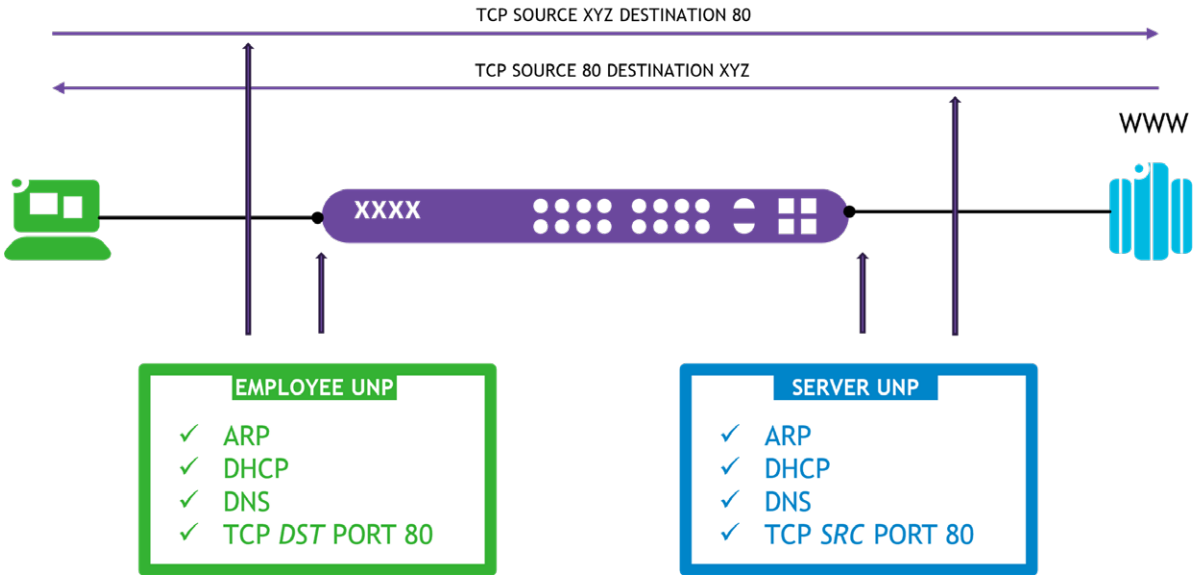


It is therefore possible, and sometimes intentional, to apply policy lists only on one side of a communication. If a device does not have a policy list applied at the access layer, it implicitly allows all traffic at that layer. This behavior applies regardless of whether devices are connected to the same switch or different switches and whether segmentation is based on VLANs or service-based designs such as SPB.

This stateless model is fundamentally different from firewall enforcement. Firewalls are stateful: once an initial packet is permitted, return traffic belonging to the same session is automatically allowed. With access layer policy lists, both directions of a communication must be considered explicitly when policies are applied to multiple endpoints.

This deterministic, ingress-based enforcement model is what makes access layer policies predictable and scalable, while reserving session-aware inspection and complex decision-making for the firewall.

When authentication is enforced on access ports, an authentication server down UNP should also be configured. This ensures that devices attempting to authenticate while the RADIUS server is temporarily unreachable, such as during a switch reload or network convergence, are placed in a temporary profile and automatically re-authenticated once the authentication infrastructure becomes available.

```
unp auth-server-down profile1 down_unp
unp auth-server-down-timeout 120
```

## 6.3 L3 access layer policy for routed campus designs

In larger campus designs, the default gateway for user VLANs may reside on a core or distribution switch rather than on the firewall. In this case, allowing traffic only to the gateway MAC address is not sufficient. Traffic must be permitted at Layer 3 toward the upstream routed domain while still preventing uncontrolled lateral movement inside the local segment.

In this design, the access switch enforces a strict allow-list. Devices are allowed to communicate north–south only toward a defined DMZ supernet where shared services and application servers reside. Fine-grained filtering to specific services such as DNS, NTP, Active Directory, intranet applications or proxy is still performed on the firewall. The access layer remains simple and deterministic: it allows only what is required to reach the next enforcement point and blocks everything else.

Local exceptions remain explicitly defined. For example, access to a printer within the same subnet can be permitted by allowing HTTPS to the specific printer IP address. Similarly, peer-to-peer media flows such as SRTP can be allowed by explicitly permitting the required UDP port range as a destination. Because UNP policies are stateless and applied at ingress, only destination-based rules are required from the perspective of the protected endpoint.

All other traffic is denied by default.

Example of Employee role in a routed campus:

Assume:

```
DMZ supernet: 10.30.0.0/16
Local printer: 10.10.10.50 (HTTPS management)
SRTP media range: UDP 40000–49999
```

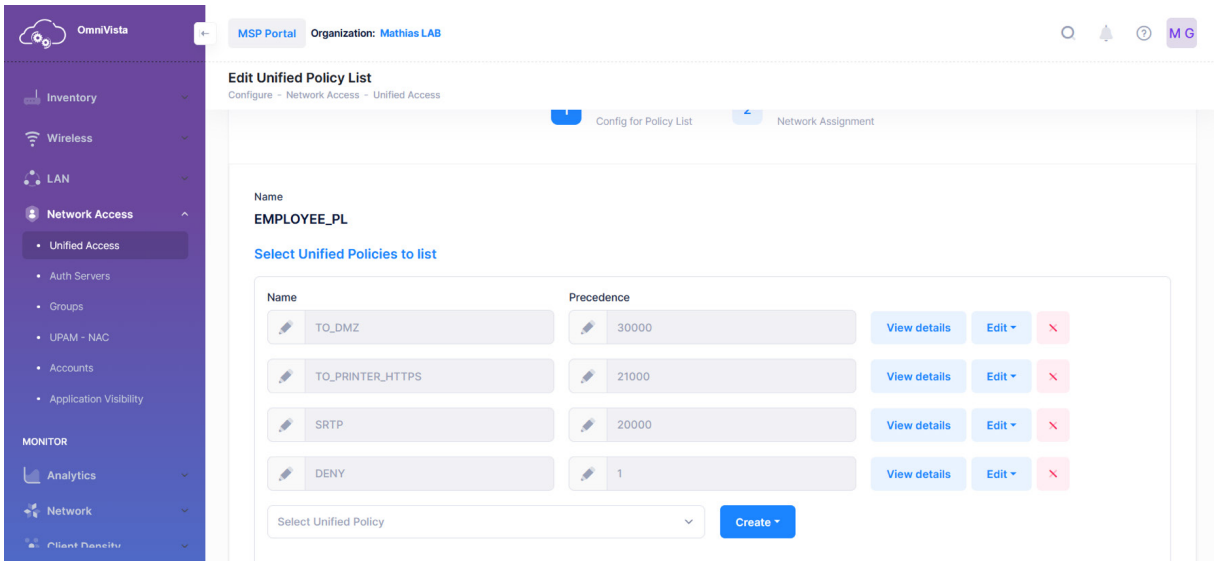

Switch policy example:

```
policy condition __TO_DMZ destination ip 10.30.0.0/16
policy condition __TO_PRINTER_HTTPS destination ip 10.10.10.50 ip-protocol 6
destination ip-port 443
policy condition __SRTP_UDP_40000_49999 ip-protocol 17 destination ip-port
40000-49999
policy condition __DENY source ip Any destination ip Any
policy action ALLOW disposition accept
policy action DROP disposition drop

policy rule TO_DMZ precedence 30000 condition __TO_DMZ action ALLOW log no
default-list
policy rule TO_PRINTER precedence 21000 condition __TO_PRINTER_HTTPS action
ALLOW log no default-list
policy rule SRTP precedence 20000 condition __SRTP_UDP_40000_49999 action ALLOW
log no default-list
policy rule DENY precedence 1 condition __DENY action DROP log no default-list

policy list EMPLOYEE_PL type unp enable
policy list EMPLOYEE_PL rules TO_DMZ TO_PRINTER SRTP DENY
```

# 7. Firewall macro-segmentation by role

## 7.1 Firewall enforcement principles

While the access switch controls communication inside a segment, the firewall controls communication between segments and towards shared services or the internet.

Each role is granted access only to the resources it explicitly requires. All other communication is denied by default.

The tables below define the expected firewall behavior for each role. These tables are written so that they can be directly translated into firewall rules.

Firewall single sign-on integration can be leverage to bind these policies to the IP addresses of currently authenticated users or devices. This allows firewall rules to remain role-based and dynamic, instead of relying only on static subnets.

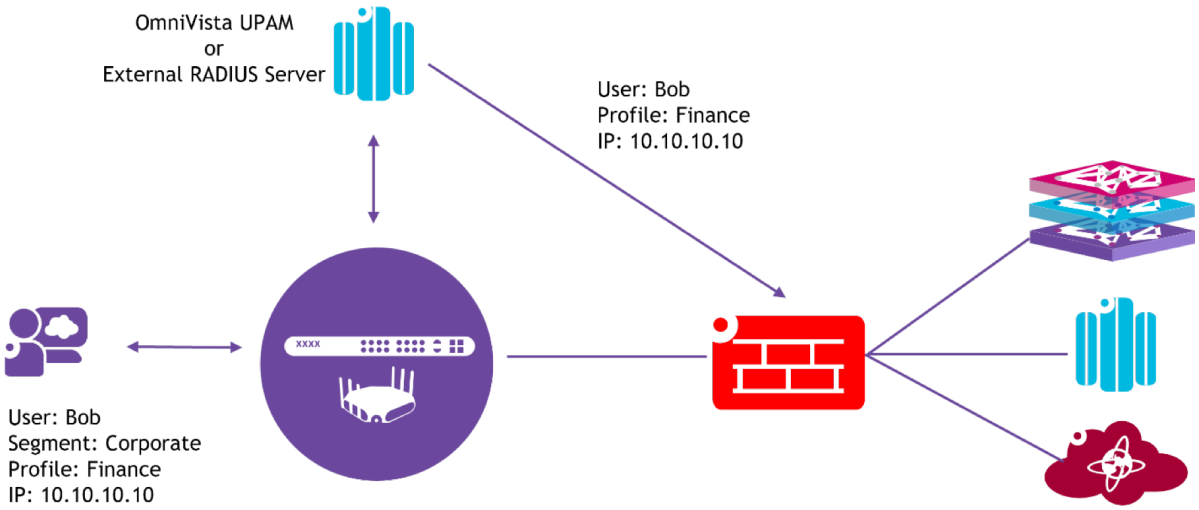

For example, the firewall can know that a specific IP address is currently used by an authenticated employee and apply the corresponding policy, while traffic from the same segment but associated with a different role is treated differently. For users authenticated through Active Directory, this information is commonly derived from AD integration. For devices and users that do not have an AD identity, such as IoT devices, identity information can be provided through integration with UPAM. Application notes and validated examples are available in the References and additional resources section of this document for major firewall vendors to support these integration models.

## 7.2 Role-based firewall policy examples

### Employee (AD, intranet and proxy access)

Employee devices are corporate-managed endpoints. They require access to Active Directory for authentication and policy enforcement, access to internal business applications and controlled internet access through an explicit proxy. Communication to any other internal segment is denied.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Employee | DNS | DNS | 53 TCP/UDP | Allow |
| Employee | NTP | NTP | 123 UDP | Allow |
| Employee | AD/DC | Kerberos/LDAP/SMB | 88, 389, 445 | Allow |
| Employee | Intranet apps | HTTPS | 443 TCP | Allow |
| Employee | Proxy | HTTP(S) proxy | 3128, 443 TCP | Allow |
| Employee | Any other segment | Any | Any | Deny |

### Employee machine profile (certificate lifecycle only)

This machine profile is applied to corporate-managed endpoints when they authenticate using a machine certificate. Its purpose is strictly limited to allowing the device to interact with Active Directory for certificate enrollment, renewal and basic Active Directory. No access to internal applications and no internet access is granted at this stage.

By separating the machine profile from the user profile, the network ensures that a device never receives broader access than required for its current authentication state. Full access is granted only after successful user authentication.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Employee | DNS | DNS | 53 TCP/UDP | Allow |
| Employee | NTP | NTP | 123 UDP | Allow |
| Employee | AD/DC | Kerberos/LDAP/SMB | 88, 389, 445 | Allow |
| Employee | Any other segment | Any | Any | Deny |

### BYOD - pre-onboarding (portal access)

Personally owned device not yet onboarded. Only what is required to reach OmniVista/UPAM and complete onboarding, plus basic infrastructure.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Employee | DNS | DNS | 53 TCP/UDP | Allow |
| Employee | NTP | NTP | 123 UDP | Allow |
| Employee | AD/DC | Kerberos/LDAP/SMB | 88, 389, 445 | Allow |
| Employee | Any other segment | Any | Any | Deny |

### BYOD — post-onboarding (steady state)

Device is onboarded and mapped to the BYOD role. Limited internal access + internet via proxy. No Active Directory. We maintain access to OmniVista/UPAM for posture, re-onboarding or portal access.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| BYOD | DNS | DNS | 53 TCP/UDP | Allow |
| BYOD | NTP | NTP | 123 UDP | Allow |
| BYOD | Limited Intranet apps | HTTPS | 443 TCP | Allow |
| BYOD | Proxy | HTTP(S) proxy | 3128, 443 TCP | Allow |
| BYOD | OmniVista/UPAM (Guest Portal) | HTTPS | 443 TCP | Allow |
| BYOD | Any other segment | Any | Any | Deny |

### Contractor (no intranet access)

Contractor access policies depend on the organization's security model. In some environments, contractors are issued corporate-managed devices and are integrated into Active Directory, in which case their access may closely resemble that of employees.

In this example, contractors are assumed to use their own devices. As a result, access is intentionally limited to a small set of required internal applications and internet access through the proxy. No access to Active Directory services is permitted, ensuring a clear separation between third-party access and core corporate resources.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Contractor | DNS | DNS | 53 TCP/UDP | Allow |
| Contractor | NTP | NTP | 123 UDP | Allow |
| Contractor | Limited Intranet apps | HTTPS | 443 TCP | Deny |
| Contractor | Proxy | HTTP(S) proxy | 3128, 443 | Allow |
| Contractor | Any other segment | Any | Any | Deny |

### Guest — pre-onboarding (guest portal access)

Untrusted user before guest acceptance/credential entry. Only portal access + basic infra.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Guest | DNS | DNS | 53 TCP/UDP | Allow |
| Guest | NTP | NTP | 123 UDP | Allow |
| Guest | OmniVista/UPAM | HTTPS | 443 TCP | Allow |
| Guest | Any internal segment | Any | Any | Deny |

### Guest — post-onboarding (internet only)

After portal success, guest role is applied. Internet goes only through the proxy. No internal access. We keep access to OmniVista/UPAM for re-auth or portal access.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| Guest | DNS | DNS | 53 TCP/UDP | Allow |
| Guest | NTP | NTP | 123 UDP | Allow |
| Guest | Proxy | HTTP(S) proxy | 3128, 443 | Allow |
| Guest | OmniVista/UPAM | HTTPS | 443 TCP | Allow |
| Guest | Any internal segment | Any | Any | Deny |

### IoT — Video cameras

Video cameras are single-purpose devices. They are allowed to communicate only with the video management system and required infrastructure services. All other communication, including lateral communication between cameras, is denied.

Access is intentionally kept very limited because these devices often do not support strong authentication mechanisms and are physically deployed in exposed locations. This increases the risk that a device or port could be misused, so network access is restricted to the minimum required for operation.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| IoT Cameras | DNS | DNS | 53 TCP/UDP | Allow |
| IoT Cameras | NTP | NTP | 123 UDP | Allow |
| IoT Cameras | VMS | RTSP | 554 | Allow |
| IoT Cameras | VMS | HTTPS | 443 | Allow |
| IoT Cameras | Any other segment | Any | Any | Deny |

### IoT — Temperature sensors

Temperature sensors publish telemetry to a central IoT broker. They do not require access to any other internal or external services. The policy reflects this limited communication model.

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| IoT Sensors | DNS | DNS | 53 TCP/UDP | Allow |
| IoT Sensors | NTP | NTP | 123 UDP | Allow |
| IoT Sensors | IoT Broker | MQTT over TLS | 8883 | Allow |
| IoT Sensors | Any other segment | Any | Any | Deny |

### IoT — HVAC

HVAC controllers communicate exclusively with a building management system. Two common patterns are shown: management over HTTPS and BACnet/IP. Only one of these patterns would typically be used in a given deployment.

#### Option A — BMS over IP (HTTPS)

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| IoT HVAC | DNS | DNS | 53 TCP/UDP | Allow |
| IoT HVAC | NTP | NTP | 123 UDP | Allow |
| IoT HVAC | BMS | HTTPS | 443 | Allow |
| IoT HVAC | Any other segment | Any | Any | Deny |

**Option B — BACnet/IP**

| From | To | Service | Ports | Allow/Deny |
|---|---|---|---|---|
| IoT HVAC | DNS | DNS | 53 TCP/UDP | Allow |
| IoT HVAC | NTP | NTP | 123 UDP | Allow |
| IoT HVAC | BMS | BACnet/IP | UDP 47808 | Allow |
| IoT HVAC | Any other segment | Any | Any | Deny |

# 8. Monitoring and operational considerations

## 8.1 What is monitored and why

Monitoring starts with knowing what is actually connected to the network. Before policies can be designed or enforced, the organization must have a clear and accurate view of devices, users and usage patterns. This inventory phase is essential to assess business needs and to reduce unnecessary exposure.

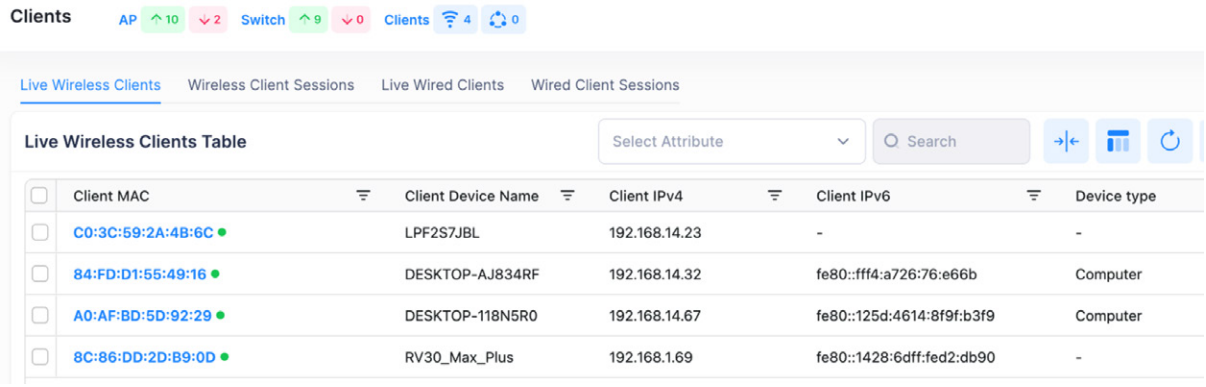


This is achieved through device fingerprinting. Fingerprinting does not authenticate a device, it classifies it by analyzing multiple observable characteristics. These include the MAC OUI (the first bytes of the MAC address identifying the manufacturer), DHCP option patterns and HTTP user-agent signatures. The combination of these attributes forms a “fingerprint” that allows the system to identify the device type and operating system.

The client inventory provided by OmniVista allows operators to see all connected wired and wireless devices, along with attributes such as MAC address, IP address, device name and device type when available. This visibility is used to identify legitimate users and devices but also to detect shadow IoT or unmanaged equipment that has no clear business justification. Devices that cannot be justified should be removed as they unnecessarily increase the attack surface. Devices that are required must then be explicitly secured through authentication and role-based policies. This visibility also enables operators to distinguish between different categories of IoT devices and to detect unexpected or unauthorized equipment connected to the network.

Once the different categories of users and devices are identified, monitoring is used to support policy design. On switches and access points that support application visibility, traffic can be observed at the application level. When authentication and roles are in place, this visibility can be correlated with user or device roles, allowing application usage to be analyzed per role rather than only per IP address or port.

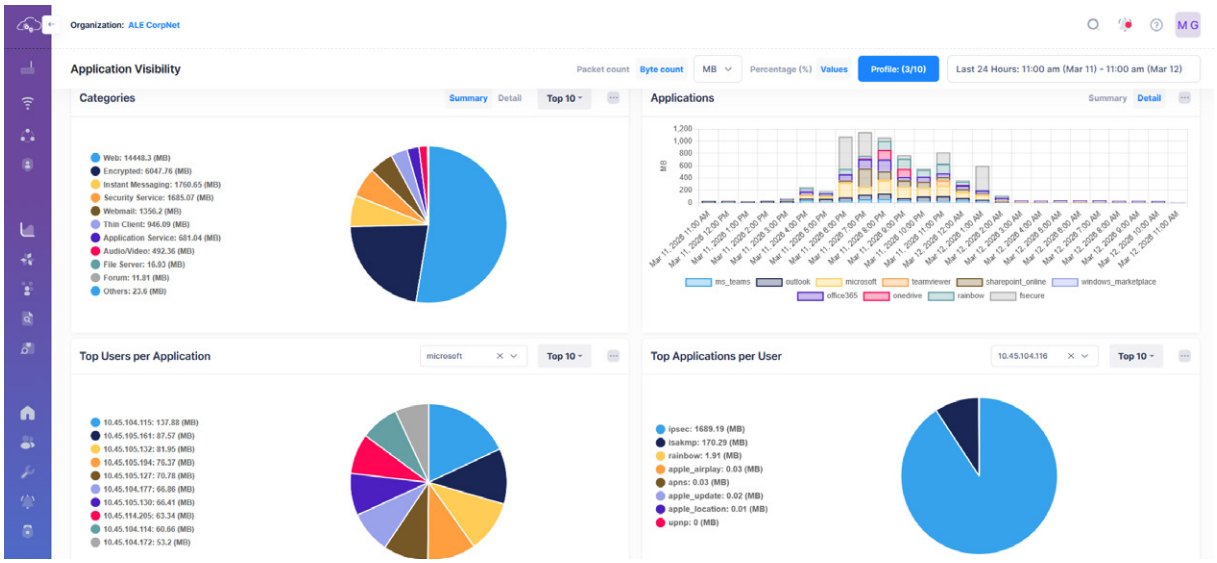


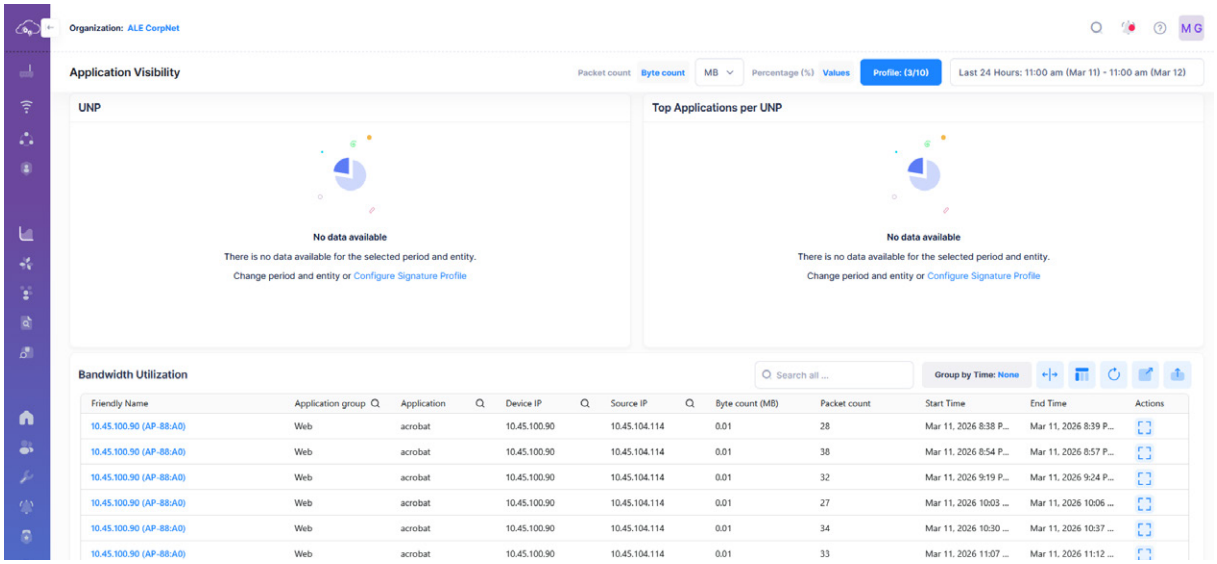


This role-aware visibility helps validate design assumptions and refine policies. It makes it possible to confirm which applications are actually used by employees, BYOD users or specific IoT device classes and to distinguish required communication from unnecessary or unexpected traffic.

Finally, before policies are fully enforced, they should be validated in operation. During testing phases, policy lists can be configured so that the last rule permits traffic instead of denying it. This allows operators to observe which flows would otherwise be blocked and to identify missing allow rules.

Policy hit counters and logs provide direct feedback on how traffic is matching the policy. By reviewing these counters, it becomes clear whether traffic is correctly matching the intended rules or falling through to the final statement. If traffic consistently matches the final permit rule, this indicates that required flows have not yet been explicitly allowed and must be investigated.

The following example shows active policy statistics on an access switch during testing:

```
POD1 -> show active policy list

 Group Name                          From  Type    Enabled  Entries
Matches

 ---------------------------------+-----+-------+--------+---------------------
---------+----------

 EMPLOYEE_PL                         cli   unp     Yes      10_TO_GW
17458

                                                           21_SRTP_40000_49999
70



POD1 ->
```

## 8.2 OVNA for anomaly detection

OmniVista Network Advisor (OVNA) is used to detect abnormal network behavior based on syslog messages generated by network devices. Rather than inspecting traffic inline, OVNA analyzes event patterns to identify conditions that deviate from normal operation, such as authentication failures, access anomalies or unexpected device behavior. When such patterns are detected, OVNA generates alerts that can be sent to network operators or IT teams through Rainbow or Teams. This allows issues to be identified and investigated quickly, without introducing additional enforcement complexity into the data path.

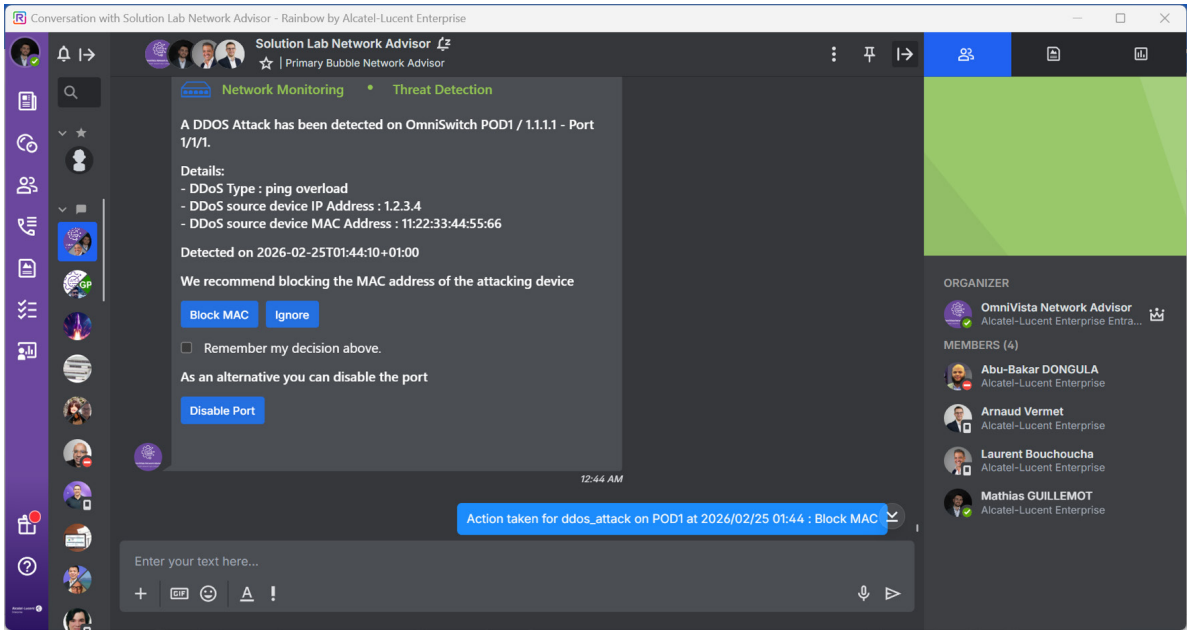

# 9. References and additional resources

The following application notes and resources complement this Design Guide with validated configuration examples, vendor-specific integration details and deployment procedures. These documents are available on Spacewalkers.com.

## 9.1 RADIUS and identity provider integration

These application notes cover the authentication path models described in section 3.3, including UPAM operating as a RADIUS proxy and as a primary RADIUS server with an external identity provider.

HPE Aruba ClearPass

Validated integration between UPAM and HPE Aruba ClearPass as an external RADIUS server

Cisco ISE

Validated integration between OmniAccess Stellar, OmniSwitch and Cisco Identity Services Engine

Microsoft Azure Entra ID

Integration of UPAM with Azure Entra ID as a cloud identity provider for user authentication

## 9.2 Firewall SSO integration

These application notes cover UPAM integration with firewall platforms for identity-aware policy enforcement, as referenced in section 7.1. They describe how authentication events from UPAM are shared with the firewall to bind role-based policies to authenticated IP addresses dynamically.

Fortinet Single Sign-On

Integration of OmniVista UPAM with FortiGate for identity-aware firewall enforcement

Palo Alto Networks Single Sign-On

Integration of UPAM with Palo Alto Networks firewalls for dynamic user-to-IP mapping and role-based policy enforcement

## 9.3 Guest Traffic Tunneling

Guest Traffic Tunneling Service (GTTS)

Describes the GRE-based model referenced in sections 3.1 and 4 for centralizing guest traffic at a core aggregation point, isolating it from user traffic without mixing it in the campus fabric

## 9.4 Standards

NIST Special Publication 800-207 — Zero Trust Architecture

Defines the Zero Trust principles referenced in section 1.1



Alcatel·Lucent Enterprise