Deployment Guide

Shortest Path Bridging (SPB)

Campus

Version 1.1 · June 2026

Alcatel·Lucent

Table of Contents

Glossary

Acronym	Description
ACC	Access Switch
BCB	Backbone Core Bridge
BEB	Backbone Edge Bridge
BSR	Bootstrap Router
BVLAN	Backbone VLAN
DHL	Dual Homed Link
ECT	Equal-Cost Tree
GRT	Global Routing Table
I-SID	Instance Service Identifier
IGMP	Internet Group Management Protocol
IPMS	IP Multicast Switching
IS-IS	Intermediate System to Intermediate System
LAG	Link Aggregation Group
LACP	Link Aggregation Control Protocol
LBD	Loopback Detection
MAC	Media Access Control
OSPF	Open Shortest Path First
OV	OmniVista
PIM	Protocol Independent Multicast
RP	Rendezvous Point
RPFC	Reverse Path Forwarding Check
RTR	Router
SAP	Service Access Point
SPB	Shortest Path Bridging (IEEE 802.1aq)
SPF	Shortest Path First
STP	Spanning Tree Protocol
UNP	Universal Network Profile
VC	Virtual Chassis
VRF	Virtual Routing and Forwarding

About this document

ATTENTION This document presents a specific configuration that was tested and validated in a controlled lab environment. It describes what was exercised in that particular setup and is provided as a validated reference for guidance only. Nothing in this document should be interpreted as a product specification, a scalability limit, a guaranteed performance figure, or a committed maximum. Actual behaviour and supported limits vary with topology, hardware variant, software release, traffic profile, and configuration. For supported maximums and committed specifications, always refer to the applicable Alcatel-Lucent Enterprise OmniSwitch platform data sheet and release notes.

Purpose

This guide describes the recommended approach for deploying Shortest Path Bridging (SPB, IEEE 802.1aq) as the network fabric, including coexistence with existing Layer 2/Layer 3 design. It provides a structured, step-by-step deployment workflow and captures baseline configurations and best practices required for a successful implementation. The guide also includes verification activities to confirm that traffic separation, resiliency, and operational manageability are achieved and maintained throughout deployment.

Use the most recent version of this document to ensure alignment with current validated practices and supported capabilities.

Scope

This document focuses on a validated and recommended deployment model. It is not intended to cover every SPB feature or explore all possible architectural variants. Instead, it provides the conceptual overview and configuration guidance necessary to implement the target architecture, which includes:

- BVLANs
- Layer 2 Services
- Layer 3 Services
- Static and Dynamic SAPs
- Static and Dynamic Services
- Interoperability with Layer 2 & Layer 3 networks

Also, this document demonstrates configuration related to single-pass inline routing only. This is not supported on all platforms. Refer to "AOS Network Configuration User Guide" for configuring two-pass routing with external or internal loopbacks as needed.

Audience

This document is intended for customers and business partners, including network architects, implementation engineers, and operations teams involved in designing and deploying enterprise networks with SPB. Readers are expected to have a working knowledge of general networking concepts and a foundational understanding of SPB.

For additional background on the SPB solution and design considerations, refer to the "SPB Design Guide" by Alcatel-Lucent Enterprise. Additional resources are listed in the Reference Documents section.

Introduction

Shortest Path Bridging (SPB)

Shortest Path Bridging (SPB, IEEE 802.1aq) is a standards-based Ethernet fabric technology designed to simplify and modernize campus and data-center Layer 2/Layer 3 networking. SPB builds a loop-free, resilient fabric by using the IS-IS link-state control plane to automatically discover the topology and compute optimal forwarding paths. Services are provisioned as logical segments across the fabric (commonly via I-SIDs), allowing the network to provide scalable Layer 2 and Layer 3 connectivity with consistent policy and predictable behavior.

In an SPB fabric, the core becomes a single operational domain that behaves like a 'network cloud': access connectivity and services are extended where needed, while the fabric automatically handles path selection, convergence, and redundancy.

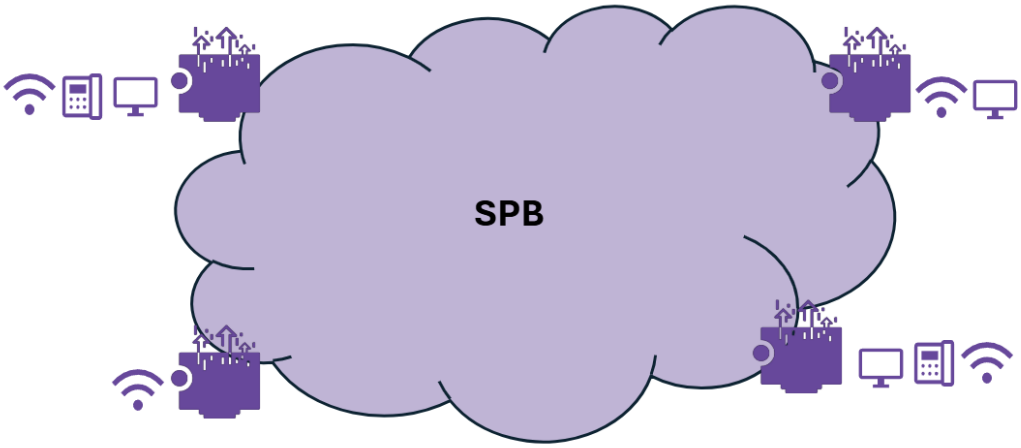

Figure 1. SPB as 'network cloud' with access connectivity extended where needed

The problems SPB solves

Traditional enterprise networks often rely on designs that become increasingly complex as they scale:

- **Spanning Tree limitations:** To prevent loops, Spanning Tree blocks links, which wastes bandwidth and creates unpredictable traffic patterns during failures or topology changes.
- **Operational complexity at scale:** Large Layer 2 domains, VLAN sprawl, and manual trunk configuration increase deployment time and troubleshooting effort.
- **Slow or disruptive convergence:** Failures can trigger reconvergence events that impact user experience, especially when multiple protocols and overlays interact.
- **Fragmented segmentation:** Enforcing consistent segmentation and service separation across many switches can lead to configuration drift and inconsistent outcomes.
- **Challenging migrations:** Introducing new services or expanding the network often requires careful manual planning to avoid loops, outages, and misconfigurations.

SPB addresses these issues by replacing blocked-link loop prevention with a link-state approach that enables all links to be active while keeping forwarding deterministic and loop-free.

Why consider SPB: key advantages

SPB is designed to deliver the benefits of a fabric architecture without introducing proprietary lock-in. Key advantages include:

- **All links active (better bandwidth utilization):** SPB enables equal-cost multipath behavior and makes efficient use of available links, avoiding the 'blocked uplink' behavior common with Spanning Tree.
- **Fast convergence and resilient behavior:** Because SPB uses a link-state control plane, the fabric can adapt quickly to failures and topology changes, improving service continuity for users and applications.
- **Simplified operations and reduced configuration overhead:** SPB reduces reliance on complex trunking and manual VLAN propagation. Services can be defined once and extended through the fabric in a consistent, controlled manner.
- **Cleaner segmentation with service-based provisioning:** Service identifiers (I-SIDs) help separate and extend services without requiring large-scale VLAN sprawl, supporting clearer intent and easier lifecycle management.
- **Predictable traffic paths and deterministic forwarding:** The fabric computes shortest paths based on a global view of the topology, producing stable forwarding behavior that is easier to reason about and troubleshoot.
- **Scalable foundation for growth and change:** SPB is well suited for expanding campus networks, adding new buildings, integrating new services, and supporting evolving connectivity needs with fewer design compromises.
- **Standards-based fabric approach:** SPB is based on an IEEE standard, enabling a widely understood model and reducing dependence on proprietary mechanisms.

Design

To demonstrate SPB deployment concepts and key capabilities, this guide uses a reference network (defined in the next section) that is designed and implemented throughout the document. The configuration is presented incrementally, introducing SPB features in logical stages so each component can be deployed, verified, and understood before proceeding.

The reference network is intentionally constructed to illustrate the recommended design options covered in this guide and to provide a practical, repeatable deployment model. By following the procedures and validation checkpoints, readers can bring up and validate an SPB fabric in their own environments and adapt the approach to match their specific topology, services, and operational requirements.

High level architecture

The reference network uses an SPB fabric as a shared transport layer to interconnect multiple buildings within a site, as shown in Figure 2. Services are introduced at the building edges and extended across the SPB domain, which operates as a single logical fabric.

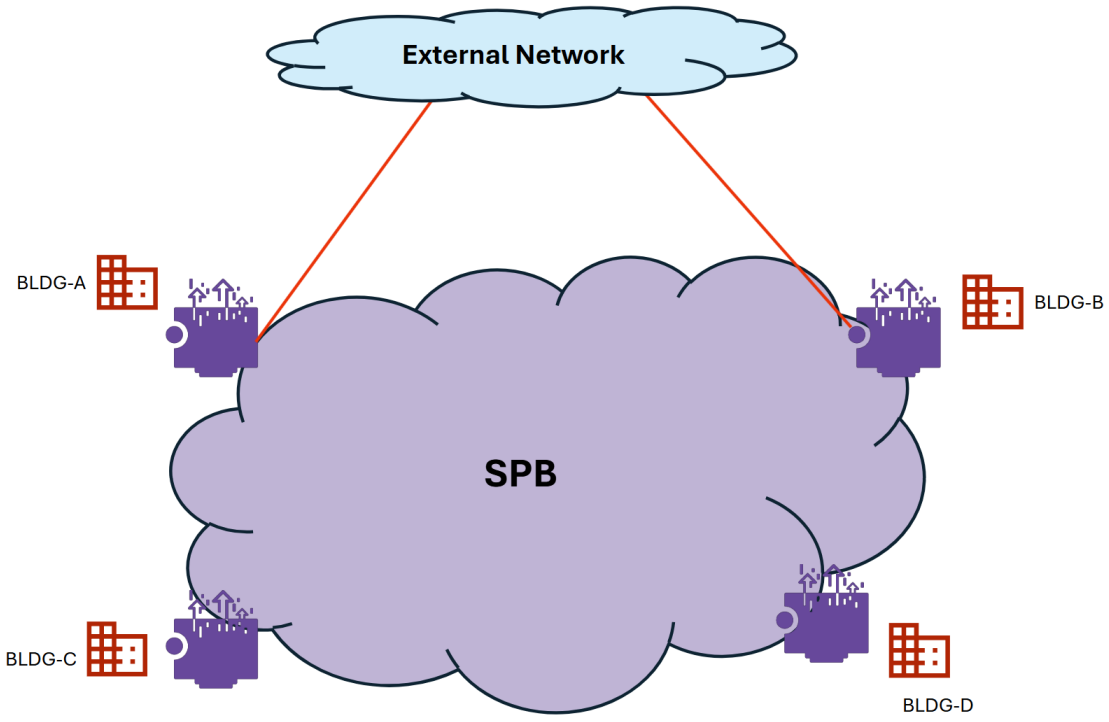

Figure 2. High level architecture

User groups (tenants) are isolated end-to-end using VRFs, providing separate routing tables and policy boundaries while sharing the same physical infrastructure. In addition, the design includes a dedicated Shared VRF to provide connectivity to common infrastructure services—such as Internet, DNS, NTP, AAA, and logging—that must be reachable across the environment without duplicating these services per tenant.

Where inter-VRF communication is required (typically tenant access to shared services), the design uses controlled route leaking with explicit import/export policies. Only approved prefixes are exchanged, keeping tenant-to-tenant connectivity blocked by default while enabling necessary exceptions in a predictable and auditable manner.

The External Network represents upstream connectivity (for example WAN/Internet/MPLS) and may include non-SPB devices such as routers, firewalls, load balancers, and provider edge equipment.

Detailed architecture

As shown in Figure 3, the reference deployment is built around a resilient SPB backbone that interconnects four building zones: BLDG-A, BLDG-B, BLDG-C, and BLDG-D. The design intentionally separates northbound (external) connectivity, core transport, and local access so each layer can be operated, scaled, and troubleshot independently.

Connectivity to the External Network is provided through two routers, RTR-1 and RTR-2. These routers exchange routes with the SPB domain using a VPN-lite approach. OSPF is used as the routing protocol between the routers and the SPB edge, with adjacencies formed toward BEB-3 and BEB-4.

The SPB backbone provides the primary east–west transport across the site. It consists of:

- Two BCBs (BCB-1 and BCB-2) acting as the fabric core/transport nodes
- Five BEBs (BEB-3 through BEB-7) acting as service edges where L2/L3 services are instantiated

Backbone connectivity is implemented as a full-mesh between BCBs and BEBs to maximize path diversity and resiliency. This topology is intended to support fast convergence and service continuity during individual link or node failures.

Each building connects to the SPB fabric through one or more BEBs:

- BLDG-A is served by BEB-3
- BLDG-B is served by BEB-4
- BLDG-C is served by BEB-5
- BLDG-D is dual-homed to BEB-6 and BEB-7 for resiliency

At the access layer, each building includes dedicated access/aggregation switches:

- BLDG-A: ACC-31, ACC-32
- BLDG-B: ACC-41, ACC-42
- BLDG-C: ACC-51
- BLDG-D: ACC-61, ACC-62

Each building is associated with a distinct IP address range (defined in the Addressing Plan section). Inter-building communication is enabled through L3VPN, while user-group separation is enforced using VRFs.

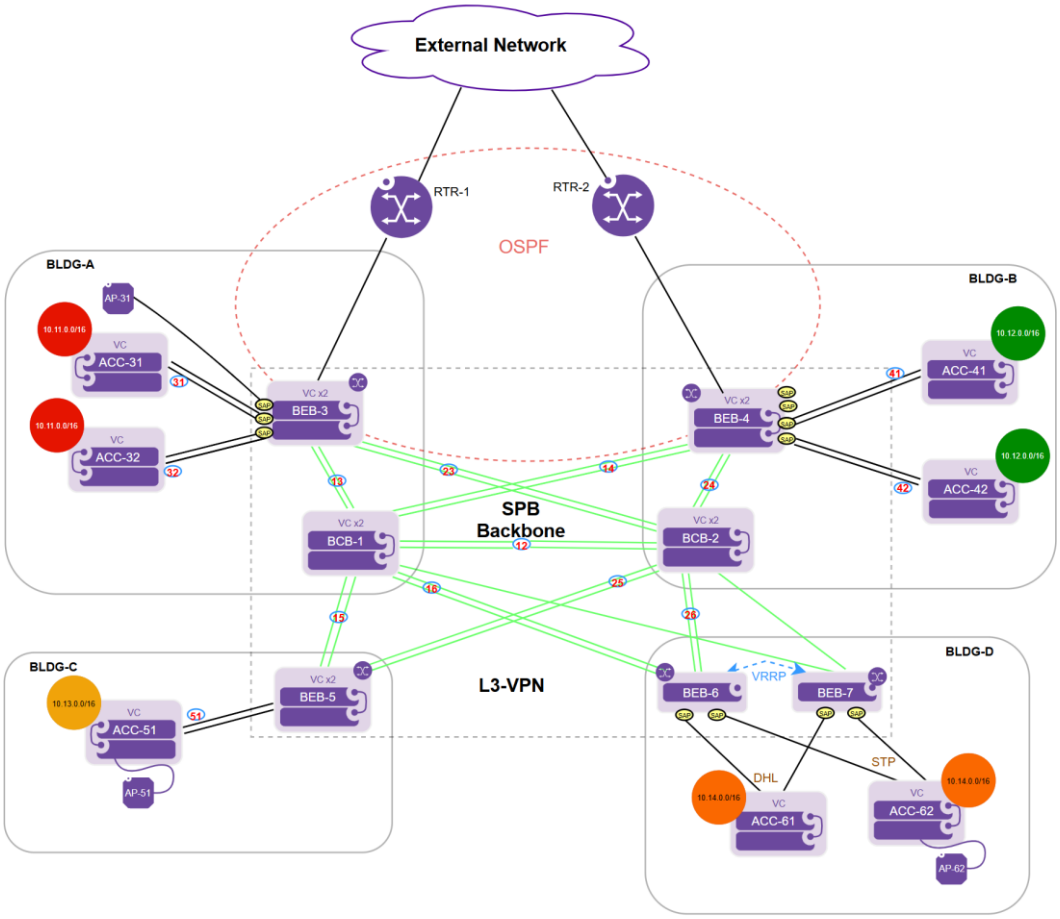

Figure 3. Detailed architecture

The reference topology includes the following VRFs:

- Corp, Guest, Voice, Util — user/service segmentation
- Mgt VRF — device management traffic (with dedicated VLAN/service used for switch and AP management)
- Shared Services VRF — common services and interconnect functions. This VRF is used to exchange routes with RTR-1/RTR-2, and selected prefixes are then leaked into other VRFs as required and redistributed within the SPB domain.
- Landlord VRF — Camera, Building IoT and Security Layer 2 services. The Security and Building IoT services use E-Tree service model, with BEB-3 and BEB-4 acting as root nodes and the remaining BEBs acting as leaf nodes. Camera service is configured to support multicast streaming.

The design uses a mix of static and dynamic services, static and dynamic SAPs, and UNP profiles, applied selectively per building and use case.

Design conventions

The following conventions are used throughout this guide to keep the examples consistent and easy to follow. These conventions are not mandatory, but they provide a simple and intuitive mental model that reduces ambiguity when reading configurations and troubleshooting. In production deployments, you may adopt a different convention to meet scalability, operational, or organizational requirements—ensure the approach you choose is consistent and can be applied at scale.

VLAN-Service-I-SID mapping

VLAN IDs, service identifiers, and I-SIDs use a one-to-one mapping. The same numeric value is used for all three constructs:

- VLAN N maps to Service N on all switches
- Service N maps to I-SID N in the SPB fabric

Example: VLAN 20 is provisioned as Service 20 on the relevant switches and is carried across the SPB fabric using I-SID 20.

> **NOTE** This convention deviates when using dynamic services where service numbers and/or I-SIDs are calculated dynamically.

Switch naming

Switch names reflect the node role and a numeric identifier.

- BCB-# : Backbone Core Bridge nodes in the SPB backbone. Example: BCB-1 and BCB-2.
- BEB-# : Backbone Edge Bridge nodes where services are provisioned. BEB numbering starts after the last BCB number. Example: BEB-3 through BEB-7.
- ACC-XY : Access switches connected to a BEB. X identifies the BEB the access switch uplinks to, Y identifies the access switch sequence number for that BEB. Example: ACC-62 is the 2nd access switch connected to BEB-6.

Link aggregation (LAG) numbering

Link aggregation group (LAG) IDs are assigned to make peer relationships obvious from the number.

- Backbone LAGs: Format Linkagg XY, where X = local BCB identifier, Y = remote BCB or BEB identifier. Example: Linkagg 23 connects BCB-2 to BEB-3.
- Access LAGs: Format Linkagg XX, where XX = the ACC-XX value of the connected access switch. Example: Linkagg 51 connects BEB-5 to ACC-51.

Addressing plan

This addressing plan supports simple route summarization, predictable operational boundaries, and easier troubleshooting.

User addressing

Users	VRF	VLAN:I-SID	BVLAN	Gateway	IP Range
Access Points	Mgt	1101	4001	10.11.0.1	10.11.0.0/20
Employee	Corp	1102	4002	10.11.16.1	10.11.16.0/20
Employee - Wifi	Corp	1103	4001	10.11.32.1	10.11.32.0/20
Guest - Wifi	Guest	1104	4002	10.11.48.1	10.11.48.0/20
Contractor	Guest	1105	4001	10.11.64.1	10.11.64.0/20
Voice	Voice	1106	4002	10.11.80.1	10.11.80.0/20
Storage	Util	1107	4001	10.11.96.1	10.11.96.0/20
Sensors	Util	1108	4002	10.11.112.1	10.11.112.0/20
Printers	Util	1109	4001	10.11.128.1	10.11.128.0/20

Table 1: BLDG-A - 10.11.0.0/16

Users	VRF	VLAN:I-SID	BVLAN	Gateway	IP Range
Access Points	Mgt	1201	4001	10.12.0.1	10.12.0.0/20
Employee	Corp	1202	4002	10.12.16.1	10.12.16.0/20
Employee – Wifi	Corp	1203	4001	10.12.32.1	10.12.32.0/20
Guest - Wifi	Guest	1204	4002	10.12.48.1	10.12.48.0/20
Contractor	Guest	1205	4001	10.12.64.1	10.12.64.0/20
Voice	Voice	1206	4002	10.12.80.1	10.12.80.0/20
Storage	Util	1207	4001	10.12.96.1	10.12.96.0/20
Sensors	Util	1208	4002	10.12.112.1	10.12.112.0/20
Printers	Util	1209	4001	10.12.128.1	10.12.128.0/20

Table 2: BLDG-B – 10.12.0.0/16

Users	VRF	VLAN:I-SID	BVLAN	Gateway	IP Range
Access Points	Mgt	1301	4001	10.13.0.1	10.13.0.0/20
Employee	Corp	1302	4002	10.13.16.1	10.13.16.0/20
Employee – Wifi	Corp	1303	4001	10.13.32.1	10.13.32.0/20
Guest - Wifi	Guest	1304	4002	10.13.48.1	10.13.48.0/20
Contractor	Guest	1305	4001	10.13.64.1	10.13.64.0/20
Voice	Voice	1306	4002	10.13.80.1	10.13.80.0/20
Storage	Util	1307	4001	10.13.96.1	10.13.96.0/20
Sensors	Util	1308	4002	10.13.112.1	10.13.112.0/20
Printers	Util	1309	4001	10.13.128.1	10.13.128.0/20

Table 3: BLDG-C – 10.13.0.0/16

Users	VRF	VLAN:I-SID	BVLAN	Gateway	IP Range
Access Points	Mgt	1401	4001	10.14.0.1	10.14.0.0/20
Employee	Corp	1402	4002	10.14.16.1	10.14.16.0/20
Employee – Wifi	Corp	1403	4001	10.14.32.1	10.14.32.0/20
Guest - Wifi	Guest	1404	4002	10.14.48.1	10.14.48.0/20
Contractor	Guest	1405	4001	10.14.64.1	10.14.64.0/20
Voice	Voice	1406	4002	10.14.80.1	10.14.80.0/20
Storage	Util	1407	4001	10.14.96.1	10.14.96.0/20
Sensors	Util	1408	4002	10.14.112.1	10.14.112.0/20
Printers	Util	1409	4001	10.14.128.1	10.14.128.0/20

Table 4: BLDG-D – 10.14.0.0/16

Users	VLAN:I-SID	BVLAN	Gateway	IP Range
Camera	2001	4002	10.20.32.1	10.20.32.0/19
Building - IoT	2002	4001	10.20.64.1	10.20.64.0/19
Security	2003	4002	10.20.96.1	10.20.96.0/19

Table 5: Landlord – 10.20.0.0/16

WAN interfaces

VRF	VLAN:I-SID	BVLAN	Network	IP address for BEB-#
Mgt	1000	4001	10.10.0.0/24	10.10.0.#
Corp	1001	4002	10.10.1.0/24	10.10.1.#
Guest	1002	4001	10.10.2.0/24	10.10.2.#
Voice	1003	4002	10.10.3.0/24	10.10.3.#
Util	1004	4001	10.10.4.0/24	10.10.4.#
Shared	1005	4002	10.10.5.0/24	10.10.5.#

Table 6: L3 VPN interfaces – 10.10.0.0/16

SPB backbone

This section describes the baseline configuration required to bring up the SPB backbone (fabric underlay) for the reference topology. The backbone configuration establishes the SPB control plane and the backbone VLANs that the fabric uses to form adjacencies, exchange topology information, and compute shortest paths. These steps must be completed before provisioning any Layer 2/Layer 3 services over SPB.

Configuring the SPB backbone entails the following steps:

- Creating BVLANs
- Defining the control BVLAN
- Defining SPB IS-IS interfaces
- Enabling the SPB IS-IS protocol

The SPB backbone used in the reference design is illustrated in Figure 4. Deploying fabric nodes as Virtual Chassis (VC) systems is recommended to improve resiliency and simplify operations through chassis-level redundancy. In the reference topology, most nodes are implemented as VCs, with BEB-6 and BEB-7 deployed as standalone switches.

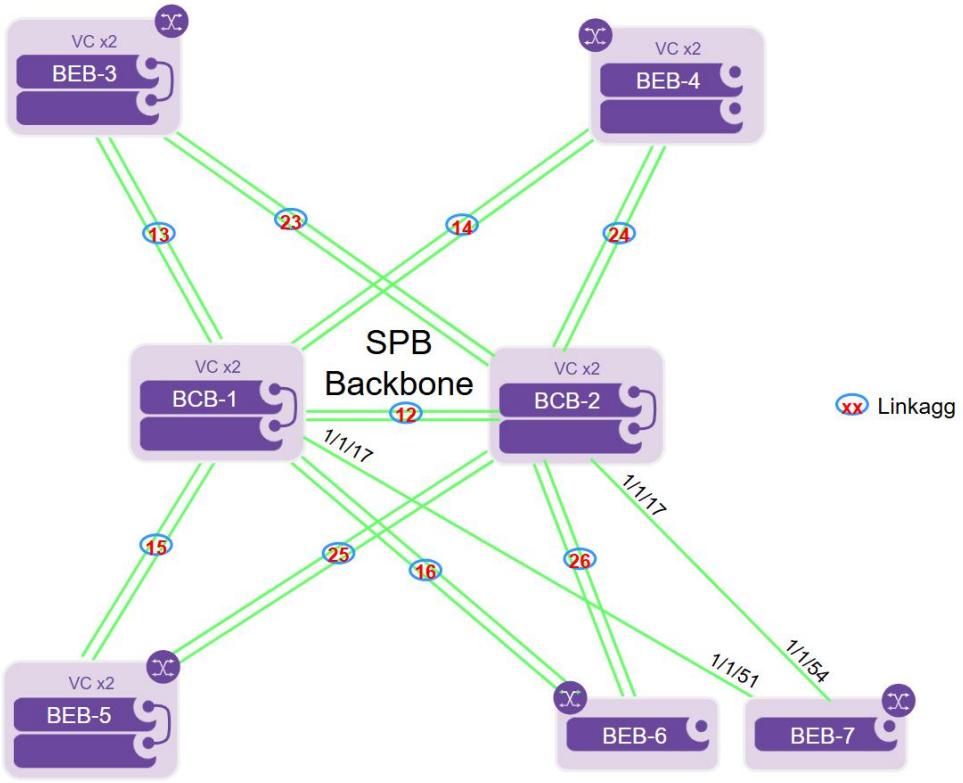

Figure 4. SPB Backbone Topology

BVLANs

Backbone VLANs (BVLANs) are used within the SPB backbone to transport encapsulated service traffic across the fabric. SPB supports up to 16 BVLANs. As a general design guideline, the number of BVLANs should align with the number of equal-cost shortest paths available in the topology to enable efficient load sharing across the fabric. Although multiple BVLANs can be created, it is highly recommended to limit the number of BVLANs to a maximum of four. Doing so reduces the scale of address updates across the control plane and improves network scalability, stability, and convergence.

A Control BVLAN is the BVLAN selected to carry SPB control-plane traffic (IS-IS messaging used to form adjacencies, advertise topology, and maintain fabric state). The control BVLAN is required for SPB operation. It does not need to be dedicated exclusively to control traffic; it may also carry data-plane traffic if desired, depending on the design.

In the reference topology, two BVLANs (4001–4002) are used to match the two equal-cost shortest paths in the fabric. BVLAN 4000 is designated as the control BVLAN.

Use the following commands to configure BVLANs and control BVLAN on all backbone nodes.

1. All BCBs and BEBs · CLI

```
spb bvlan 4000-4002
spb isis control-bvlan 4000
```

Each node computes a separate shortest path tree per BVLAN. An ECT-ID (Equal Cost Tree Identifier) is associated with each BVLAN and is used as a tie-break mechanism so that different BVLANs can produce different equal-cost trees, improving load distribution across the fabric. When BVLANs are created using the commands above, AOS automatically assigns ECT-IDs sequentially to each BVLAN.

SPB interfaces

The next step is to configure the physical uplinks that form the backbone—either individual ports or link aggregation groups (LAGs)—as SPB (IS-IS) interfaces. These interfaces participate in the SPB control plane and are used to establish adjacencies across the fabric.

In the reference topology, BEB-7 connects to each BCB using single physical links, while the remaining BEBs connect to the BCBs using link aggregation for additional bandwidth and resiliency.

SPB uses a default link metric of 10 to calculate the Spanning Tree irrespective of the link speed. Modifying this can help utilize links more efficiently. The table below shows a sample of how link metrics can be set for different link speeds.

Speed	Metric
800G	125
400G	250
200G	500
100G	1000
50G	2000
25G	4000
10G	10000
1G	100000
100M	1000000

Table 7: Link metric sample

Use the commands below to configure the required SPB interfaces on the applicable ports and LAGs for each device:

2. Example on BCB-1 · CLI

```
spb isis interface port 1/1/17 metric 10000
spb isis interface linkagg 12 metric 5000
spb isis interface linkagg 13 metric 5000
spb isis interface linkagg 14 metric 1250
spb isis interface linkagg 15 metric 5000
spb isis interface linkagg 16 metric 5000
```

After all SPB interfaces are configured, enable SPB IS-IS on all backbone nodes using the following command:

3. All BCBs and BEBs · CLI

```
spb isis admin-state enable
```

Validation

Now that the SPB backbone is configured, validate fabric health before provisioning any services. The commands in this section confirm that SPB interfaces are up, adjacencies are formed, the complete set of nodes is discovered, and BVLAN/ECT behavior is as expected.

Verify that all backbone links (ports/LAGs) are enabled for SPB and are operationally up:

4. Example on BCB-1 · CLI

```
BCB-1-> show spb isis interface
SPB ISIS Interfaces:                Oper    Admin    Link       Hello    Hello   Circ
Interface        Level    CircID     state   state    Metric     Intvl    Mult    Type
---------------+-------+----------+------+-------+----------+-------+------+------------
 1/1/17           L1       9          UP      UP       10000      9        3       Pt-to-Pt
 0/12             L1       1          UP      UP       5000       9        3       Pt-to-Pt
 0/13             L1       2          UP      UP       5000       9        3       Pt-to-Pt
 0/14             L1       3          UP      UP       1250       9        3       Pt-to-Pt
 0/15             L1       4          UP      UP       5000       9        3       Pt-to-Pt
 0/16             L1       5          UP      UP       5000       9        3       Pt-to-Pt

Interfaces : 6
```

The output lists interfaces participating in SPB IS-IS (Level-1) and their status. Confirm all expected interfaces are administratively up and operationally up.

Next, verify that the node has discovered all other SPB nodes in the network:

5. Example on BCB-1 · CLI

```
BCB-1-> show spb isis nodes
SPB ISIS Nodes:
System Name          System Id      SourceID BridgePriority
--------------------+--------------+--------+--------------
 BEB-5                2cfa.a271.c2cd  0x1c2cd  32768 (0x8000)
 BCB-1                2cfa.a2a9.c91f  0x9c91f  32768 (0x8000)
 BEB-4                7824.5974.93ae  0x493ae  32768 (0x8000)
 BEB-3                7824.59a2.2447  0x22447  32768 (0x8000)
 BEB-7                9424.e13a.309d  0xa309d  32768 (0x8000)
 BEB-6                9424.e159.706d  0x9706d  32768 (0x8000)
 BCB-2                dc08.5610.0dc1  0x00dc1  32768 (0x8000)

Total SPB Nodes : 7
```

Ensure all BCBs and BEBs (including the local node) are listed. Repeat on multiple nodes if needed to confirm consistent fabric visibility.

Verify that IS-IS adjacencies are established with all directly connected neighbors:

6. Example on BCB-1 · CLI

```
BCB-1-> show spb isis adjacency
SPB ISIS Adjacency:
(Name : SystemId)                      Type   State   Hold  Interface
-------------------------------------+------+-------+------+----------
 BEB-7               : 9424.e13a.309d  L1     UP        22      1/1/17
 BCB-2               : dc08.5610.0dc1  L1     UP        22        0/12
 BEB-3               : 7824.59a2.2447  L1     UP        23        0/13
 BEB-4               : 7824.5974.93ae  L1     UP        26        0/14
 BEB-5               : 2cfa.a271.c2cd  L1     UP        20        0/15
 BEB-6               : 9424.e159.706d  L1     UP        19        0/16

Adjacencies : 6
```

Confirm an adjacency exists for each physically connected SPB neighbor and that the adjacency is formed over the expected interface (port/LAG).

Verify the configured BVLANs, their ECT settings, and usage:

7. Example on BCB-1 · CLI

```
BCB-1-> show spb isis bvlans
SPB ISIS BVLANS:
                                    Services  Num    Tandem     Root Bridge
 BVLAN   ECT-algorithm     In Use   mapped    ISIDS  Multicast  (Name : MAC Address)
-------+-----------------+-------+---------+------+----------+---------------------------
  4000  00-80-c2-01       YES     NO            0   SGMODE                          L1L2
  4001  00-80-c2-02       NO      NO            0   SGMODE                          L1L2
  4002  00-80-c2-03       NO      NO            0   SGMODE                          L1L2

BVLANs:     3
```

The output should show:

- All configured BVLANs (including the control BVLAN)
- The ECT algorithm/ID associated with each BVLAN
- Current usage/mappings (services / I-SIDs)

At this stage, service provisioning has not been completed, so it is expected to see no I-SIDs/services mapped to BVLANs. The control BVLAN (4000) should show as 'in use' because it carries control-plane exchanges. Confirm BVLANs and ECT assignments are consistent across all backbone nodes.

Use the following command to verify the outbound interface used to reach a specific destination node:

8. Example on BEB-6 · CLI

```
BEB-6-> show spb isis unicast-table
SPB ISIS Unicast MAC Table:
            Destination                        Outbound    Level
  BVLAN     Name : MAC Address)                Interface
------+-----------------------------------+-----------+-------
  4000  BEB-5               : 2c:fa:a2:71:c2:cd       0/16     1
  4000  BCB-1               : 2c:fa:a2:a9:c9:1f       0/16     1
  4000  BEB-4               : 78:24:59:74:93:ae       0/16     1
  4000  BEB-3               : 78:24:59:a2:24:47       0/16     1
  4000  BEB-7               : 94:24:e1:3a:30:9d       0/16     1
  4000  BCB-2               : dc:08:56:10:0d:c1       0/26     1
  4001  BEB-5               : 2c:fa:a2:71:c2:cd       0/26     1
  4001  BCB-1               : 2c:fa:a2:a9:c9:1f       0/16     1
  4001  BEB-4               : 78:24:59:74:93:ae       0/26     1
  4001  BEB-3               : 78:24:59:a2:24:47       0/26     1
  4001  BEB-7               : 94:24:e1:3a:30:9d       0/16     1
  4001  BCB-2               : dc:08:56:10:0d:c1       0/26     1
  4002  BEB-5               : 2c:fa:a2:71:c2:cd       0/26     1
  4002  BCB-1               : 2c:fa:a2:a9:c9:1f       0/16     1
  4002  BEB-4               : 78:24:59:74:93:ae       0/26     1
  4002  BEB-3               : 78:24:59:a2:24:47       0/26     1
  4002  BEB-7               : 94:24:e1:3a:30:9d       0/16     1
  4002  BCB-2               : dc:08:56:10:0d:c1       0/26     1

MAC Addresses: 18
```

Note that the outbound interface may differ by BVLAN because each BVLAN can compute a different equal-cost tree. For example, traffic to BEB-5 uses Linkagg 16 on BVLAN 4000, but uses Linkagg 26 on BVLAN 4001 and 4002.

Verify the SPF tree for a given BVLAN:

9. Example on BEB-6 · CLI

```
BEB-6-> show spb isis spf bvlan 4000
SPB ISIS Path Table:
 Destination                          Outbound   Next Hop                          SPB     Num
 (Name : BMAC)                        Interface  (Name : BMAC)                     Metric  Hops
--------------------------------------+----------+---------------------------------+-------+------
 BEB-5             : 2c:fa:a2:71:c2:cd       0/16  BCB-1             : 2c:fa:a2:a9:c9:1f  10000    2
 BCB-1             : 2c:fa:a2:a9:c9:1f       0/16  BCB-1             : 2c:fa:a2:a9:c9:1f   5000    1
```

```
 BEB-4               : 78:24:59:74:93:ae       0/16  BCB-1               : 2c:fa:a2:a9:c9:1f   6250    2
 BEB-3               : 78:24:59:a2:24:47       0/16  BCB-1               : 2c:fa:a2:a9:c9:1f  10000    2
 BEB-7               : 94:24:e1:3a:30:9d       0/16  BCB-1               : 2c:fa:a2:a9:c9:1f  15000    2
 BCB-2               : dc:08:56:10:0d:c1       0/26  BCB-2               : dc:08:56:10:0d:c1   5000    1

SPF Path count: 6

BEB-6-> show spb isis spf bvlan 4001
SPB ISIS Path Table:
 Destination                              Outbound   Next Hop                              SPB     Num
 (Name : BMAC)                            Interface  (Name : BMAC)                         Metric  Hops
-----------------------------------------+----------+-------------------------------------+-------+------
 BEB-5               : 2c:fa:a2:71:c2:cd       0/26  BCB-2               : dc:08:56:10:0d:c1  10000    2
 BCB-1               : 2c:fa:a2:a9:c9:1f       0/16  BCB-1               : 2c:fa:a2:a9:c9:1f   5000    1
 BEB-4               : 78:24:59:74:93:ae       0/26  BCB-2               : dc:08:56:10:0d:c1   6250    2
 BEB-3               : 78:24:59:a2:24:47       0/26  BCB-2               : dc:08:56:10:0d:c1  10000    2
 BEB-7               : 94:24:e1:3a:30:9d       0/26  BCB-2               : dc:08:56:10:0d:c1  15000    2
 BCB-2               : dc:08:56:10:0d:c1       0/26  BCB-2               : dc:08:56:10:0d:c1   5000    1

SPF Path count: 6

BEB-6-> show spb isis spf bvlan 4002
SPB ISIS Path Table:
 Destination                              Outbound   Next Hop                              SPB     Num
 (Name : BMAC)                            Interface  (Name : BMAC)                         Metric  Hops
-----------------------------------------+----------+-------------------------------------+-------+------
 BEB-5               : 2c:fa:a2:71:c2:cd       0/26  BCB-2               : dc:08:56:10:0d:c1  10000    2
 BCB-1               : 2c:fa:a2:a9:c9:1f       0/16  BCB-1               : 2c:fa:a2:a9:c9:1f   5000    1
 BEB-4               : 78:24:59:74:93:ae       0/26  BCB-2               : dc:08:56:10:0d:c1   6250    2
 BEB-3               : 78:24:59:a2:24:47       0/26  BCB-2               : dc:08:56:10:0d:c1  10000    2
 BEB-7               : 94:24:e1:3a:30:9d       0/26  BCB-2               : dc:08:56:10:0d:c1  15000    2
 BCB-2               : dc:08:56:10:0d:c1       0/26  BCB-2               : dc:08:56:10:0d:c1   5000    1

SPF Path count: 6
```

The output shows, for the specified BVLAN, the next hop and SPB metric to reach each node. The metric may remain the same across BVLANs, while the next hop can differ due to different ECT trees. For example, the next hop toward BEB-5 is BCB-1 on BVLAN 4000, but BCB-2 on BVLAN 4001/4002.

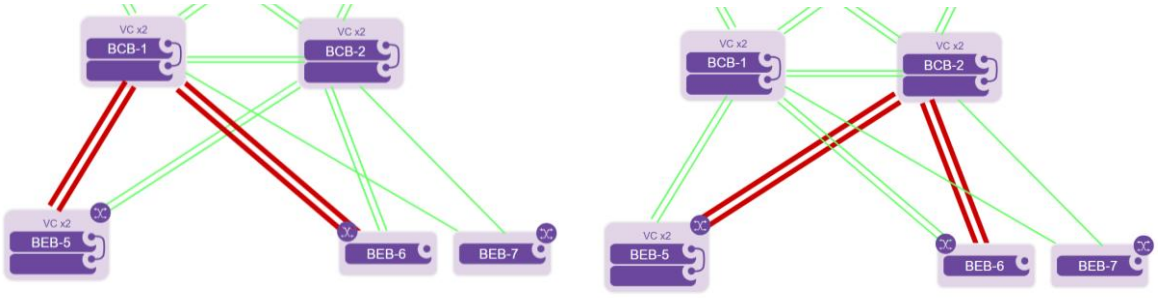

Figure 5. Equal cost path from BEB-6 to BEB-5

Validation using OV

OmniVista (OV) does not currently support provisioning the SPB backbone configuration, but it can be used to validate the fabric state and confirm that the control plane is operating as expected.

To access the SPB fabric view in OV, navigate to: Configure → LAN → Advanced Network

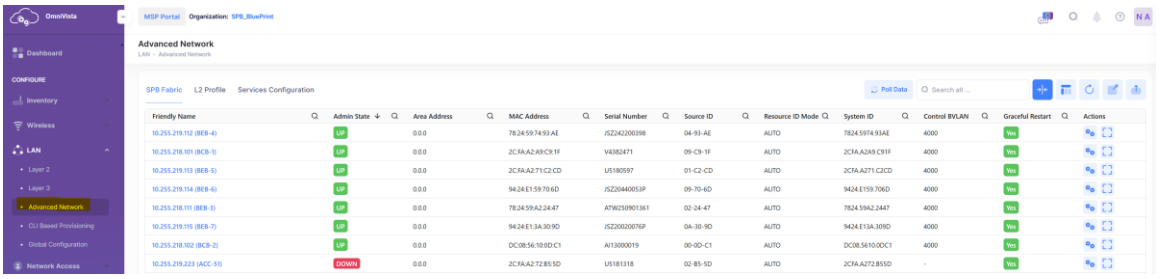

Figure 6. OV Advanced Network menu

This page provides a consolidated view of SPB status across all onboarded switches, including key fabric information such as the Control BVLAN.

To view device-specific SPB details, click the Friendly Name of a switch. The device-level view includes information such as configured BVLANs, IS-IS interfaces, and IS-IS adjacency status for that node.

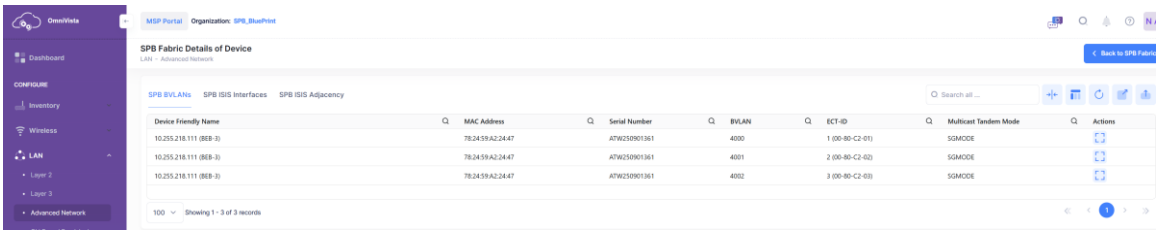

Figure 7. Example on BEB-3

Layer 2 services

A Layer 2 (L2) service in an SPB network provides end-to-end Ethernet connectivity across the fabric by transporting customer frames using MAC-in-MAC encapsulation. With the SPB backbone already in place, this chapter focuses on how to provision and validate L2 services in the reference topology.

In addition to basic L2 service creation, this chapter introduces operational service constructs commonly used in enterprise deployments, including dynamic SAPs, dynamic services, SAP handling for Stellar AP deployments, and multicast services. The chapter concludes with an example of the E-Tree service model.

Creation of SPB services generally involves the following steps:

- Creating VLANs on access switches (where applicable)
- Creating Services and binding them to appropriate I-SID and BVLAN
- Creating appropriate L2 Profiles
- Enabling service access interfaces
- Defining SAPs and binding them to respective services

Client connectivity to an SPB-based L2 network typically falls into one of the following models:

- Clients connected through dedicated access switches: Clients attach to an access switch where the required VLANs are configured and client-facing ports are assigned to the appropriate VLAN. Uplinks from the access switch carry VLAN-tagged traffic to the BEB. At the BEB, the traffic is mapped to the corresponding Service Access Point (SAP) and then bridged across the SPB fabric to the destination.
- Clients connected through access switches that also function as SPB BEBs: In this model, the access switch also serves as the SPB BEB. VLAN configuration on the BEB is not required for client attachment. Instead, client traffic is mapped directly to the appropriate SAP and bridged across the SPB fabric. Removing VLAN configuration from the BEB can simplify operations and reduce the need for Spanning Tree considerations at the fabric edge.

Figure 8. below illustrates these two scenarios.

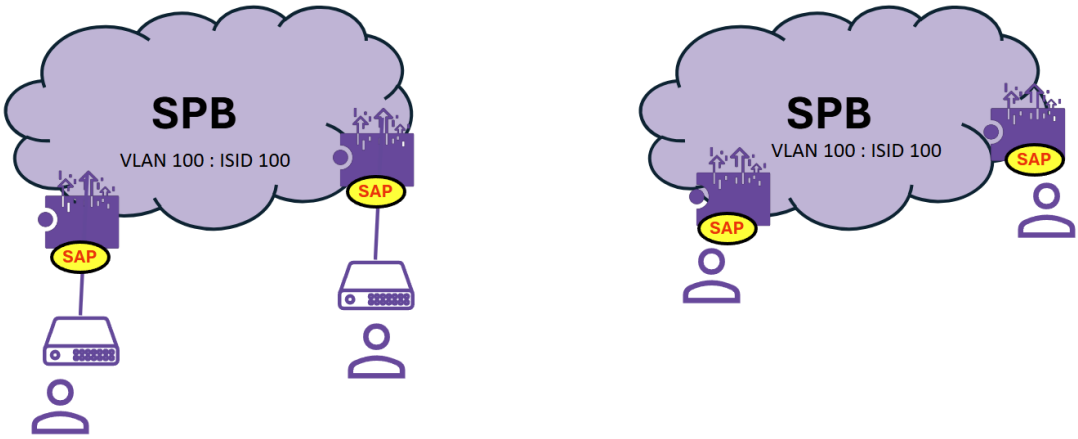

Figure 8. L2 service scenarios

VLANs

In the reference topology, selected VLANs are extended to the access layer as defined in the Addressing Plan. Configure the required VLANs on all ACC-## switches and ensure the uplink interfaces toward the BEBs are correctly tagged for those VLANs.

Use the following commands to create the VLAN(s) and tag the uplink interface(s) on each access switch:

10. Example on ACC-31 · CLI

```
vlan 1101 name "Access Points"
vlan 1102 name "Employee"
vlan 1103 name "Employee-Wifi"
vlan 1104 name "Guest-Wifi"
vlan 1105 name "Contractor"
vlan 1106 name "Voice"
vlan 1107 name "Storage"
vlan 1108 name "Sensors"
vlan 1109 name "Printers"
vlan 2001 name "Cameras"
vlan 2002 name "IoT"
vlan 2003 name "Security"
vlan 1101-1109 members linkagg 31 tagged
vlan 2001-2003 members linkagg 31 tagged
```

After configuration, verify VLAN creation and tagging using:

11. Example on ACC-31 · CLI

```
ACC-31-> show vlan
 vlan    type   admin   oper    ip    mtu          name
------+-------+-------+------+------+------+-----------------
 1      std       Dis     Dis   Dis    1500    VLAN 1
 1101   std       Ena     Ena   Dis    1500    Access Points
 1102   std       Ena     Ena   Dis    1500    Employee
 1103   std       Ena     Ena   Dis    1500    Employee-Wifi
 1104   std       Ena     Ena   Dis    1500    Guest-Wifi
 1105   std       Ena     Ena   Dis    1500    Contractor
 1106   std       Ena     Ena   Dis    1500    Voice
 1107   std       Ena     Ena   Dis    1500    Storage
 1108   std       Ena     Ena   Dis    1500    Sensors
 1109   std       Ena     Ena   Dis    1500    Printers
 2001   std       Ena     Ena   Dis    1500    Cameras
 2002   std       Ena     Ena   Dis    1500    IoT
 2003   std       Ena     Ena   Dis    1500    Security
 4094   vcm       Ena     Dis   Dis    1500    VCM IPC
```

Confirm that:

- All expected VLAN IDs are present on the switch
- The correct uplink ports/LAGs are tagged for the VLANs

SPB services

The next step is to configure SPB services on the BEB switches. Services are provisioned only on BEBs where they are required. BCBs do not host services and operate purely as fabric transit nodes.

An SPB service is identified by a service ID and represents the association between a Backbone Service Instance Identifier (I-SID) and an existing BVLAN. In practical terms, creating a service binds an I-SID to the SPB backbone, enabling traffic mapped to that I-SID to be encapsulated and transported across the fabric on the associated BVLAN(s) toward the destination BEB.

As mentioned earlier, this guide uses a 1:1:1 numbering convention for VLAN:Service:I-SID numbering for simplicity. In addition, vlan translation is enabled on all services so that the egress frames are handled according to the local SAP configuration. This enables communication between clients using different traffic encapsulations. For example, a server may transmit tagged traffic while connected clients expect to receive untagged traffic. Depending on the platform, VLAN translation may also need to be enabled at the access interface level—refer to the “AOS Network Configuration User guide” for more information.

This section focuses on static service provisioning. SPB services can also be created dynamically (demonstrated later on BEB-5 for services 2001-2003). Dynamic service creation is covered in the “Service Access Points (SAP)” section later in this chapter.

Configure the required services on the appropriate BEB-## switches according to the Addressing Plan:

12. Example on BEB-3 · CLI

```
service 1000 spb isid 1000 bvlan 4001 description "Mgt" vlan-xlation enable
service 1001 spb isid 1001 bvlan 4002 description "Corp" vlan-xlation enable
service 1002 spb isid 1002 bvlan 4001 description "Guest" vlan-xlation enable
service 1003 spb isid 1003 bvlan 4002 description "Voice" vlan-xlation enable
service 1004 spb isid 1004 bvlan 4001 description "Util" vlan-xlation enable
service 1101 spb isid 1101 bvlan 4001 vlan-xlation enable
service 1102 spb isid 1102 bvlan 4002 vlan-xlation enable
service 1103 spb isid 1103 bvlan 4001 vlan-xlation enable
service 1104 spb isid 1104 bvlan 4002 vlan-xlation enable
service 1105 spb isid 1105 bvlan 4001 vlan-xlation enable
service 1106 spb isid 1106 bvlan 4002 vlan-xlation enable
service 1107 spb isid 1107 bvlan 4001 vlan-xlation enable
service 1108 spb isid 1108 bvlan 4002 vlan-xlation enable
service 1109 spb isid 1109 bvlan 4001 vlan-xlation enable
service 2001 spb isid 2001 bvlan 4002 vlan-xlation enable
service 2002 spb isid 2002 bvlan 4001 vlan-xlation enable
service 2003 spb isid 2003 bvlan 4002 vlan-xlation enable
```

After services are configured, verify them using:

13. Example on BEB-3 · CLI

```
BEB-3-> show service spb
Legend: * denotes a dynamic object
SPB Service Info
  SystemId : 7824.59a2.2447,   SrcId : 0x22447,    SystemName : BEB-3

                            SAP     Bind                    MCast
ServiceId   Adm  Oper Stats Count   Count   Isid      BVlan Mode     (T/R)
-----------+----+----+-----+-------+-------+---------+-----+--------------
1000        Up   Up    N    0       4       1000      4001  Headend  (0/0)
1001        Up   Up    N    0       4       1001      4002  Headend  (0/0)
1002        Up   Up    N    0       4       1002      4001  Headend  (0/0)
1003        Up   Up    N    0       4       1003      4002  Headend  (0/0)
1004        Up   Up    N    0       4       1004      4001  Headend  (0/0)
1101        Up   Down  N    0       0       1101      4001  Headend  (0/0)
1102        Up   Down  N    0       0       1102      4002  Headend  (0/0)
1103        Up   Down  N    0       0       1103      4001  Headend  (0/0)
1104        Up   Down  N    0       0       1104      4002  Headend  (0/0)
1105        Up   Down  N    0       0       1105      4001  Headend  (0/0)
1106        Up   Down  N    0       0       1106      4002  Headend  (0/0)
1107        Up   Down  N    0       0       1107      4001  Headend  (0/0)
1108        Up   Down  N    0       0       1108      4002  Headend  (0/0)
1109        Up   Down  N    0       0       1109      4001  Headend  (0/0)
2001        Up   Up    N    0       4       2001      4002  Headend  (0/0)
2002        Up   Up    N    0       4       2002      4001  Headend  (0/0)
2003        Up   Up    N    0       4       2003      4002  Headend  (0/0)

Total Services: 17
```

In the verification output, Bind Count indicates the number of SPB bindings associated with the service across the fabric. For example, services 1000-1004 show a bind count of 4 because they are also provisioned on BEB-4, BEB-5, BEB-6, and BEB-7.

To remove a service, first delete any SAPs and interfaces associated with it. Then set the service to an administratively disabled state and delete the service. Use the following commands to remove a service after the associated SAPs and interfaces have been deleted:

14. Deleting a service · CLI

```
service 300 admin-state disable
no service 300
```

L2 profile and service access interfaces

Access interfaces are customer-facing ports or linkaggs on BEB switches. Traffic received on these interfaces is classified into one or more SAPs and then forwarded across the SPB fabric using the associated SPB service.

An L2 access profile is applied to service access interfaces to control how specific ingressing Layer 2 control frames are handled.

Although a default profile is available, it drops 802.1AB frames used for link layer device discovery and 802.1X frames which are used for authentication. Link layer discovery protocol (LLDP) is often desirable during initial rollout and can be restricted later based on security policy. To support both discovery and authentication in this reference deployment, an L2 access profile named “l2prof” is created on all BEBs to permit 802.1AB and 802.1X, and is applied consistently to all service access interfaces.

Create the L2 profile on the BEBs:

15. All BEBs · CLI

```
service l2profile "l2prof" 802.1ab peer 802.1x peer
```

Next, define the required service access interfaces on the relevant BEB-## switches and apply the L2 profile:

16. Example on BEB-3 · CLI

```
service access linkagg 31 vlan-xlation enable l2profile "l2prof"
service access linkagg 32 vlan-xlation enable l2profile "l2prof"
```

Verify that the correct ports/LAGs are configured as service access interfaces and that l2prof is applied:

17. Example on BEB-3 · CLI

```
BEB-3-> show service access
Legend: (-)ERP Ring (*) Hybrid
Port       Link  SAP      SAP      Vlan Id
  Status Type    Count    Xlation L2Profile                    Description
---------+------+-------+-------+-------+--------------------------------+-----------
  0/31       Up      Manual  0        Y       l2prof
  0/32       Up      Manual  0        Y       l2prof
Total Access Ports: 2
```

L2 profile using OV

OmniVista provides the ability to view and define Layer 2 profiles from a central management interface. This enables administrators to validate existing L2 profiles and create new profiles in OmniVista for later use during service provisioning.

To access the L2 Profile page, navigate to: Configure → LAN → Advanced Network → L2 Profile

The L2 Profile page displays the list of profiles currently configured in OmniVista.

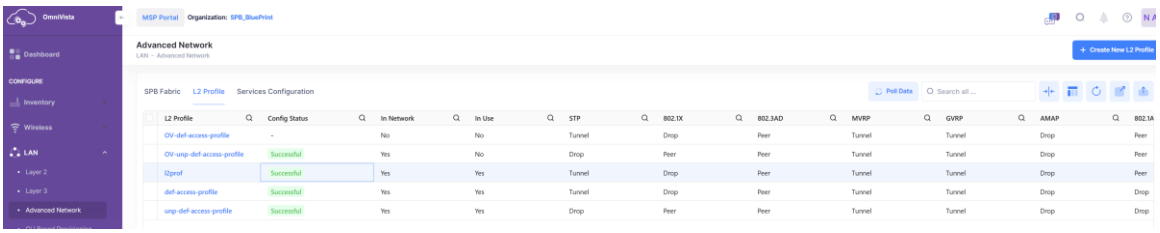

Figure 9. L2 Profile screen on OV

To create a new L2 profile, click Create New L2 Profile in the upper-right corner of the page. Complete the required fields, then click Create.

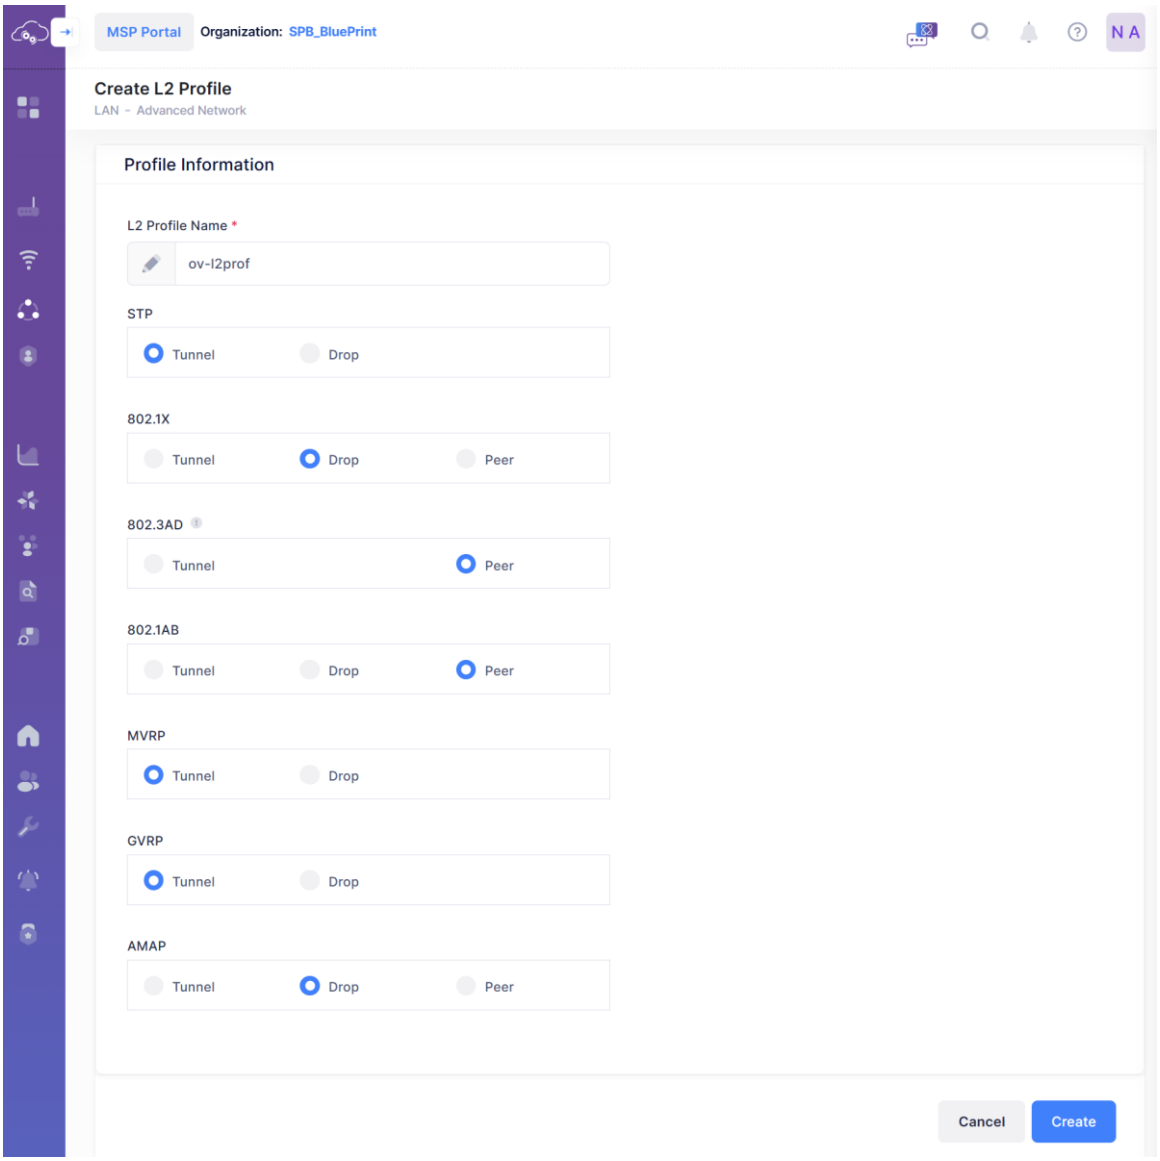

Figure 10. L2 Profile creation using OV

The new L2 profile is created in OmniVista and becomes available for use during service configuration. The profile is not immediately pushed to the switch. Instead, the corresponding L2 profile configuration is applied to the device when it is referenced as part of a service configuration.

Service Access Points (SAP)

A Service Access Point (SAP) is the UNI-side logical attachment that binds a physical interface (port or LAG) and a specific encapsulation type (untagged, single-tagged, double-tagged, or 'all') to an SPB service. It represents the point where customer traffic is classified and associated with a service, after which it can be carried across the SPB fabric.

A single physical interface can host multiple SAPs, allowing different traffic types (for example, multiple VLANs on a trunk) to be multiplexed and mapped to different SPB services.

SAPs can be created in two ways:

- Static SAPs (manually configured)
- Dynamic SAPs (created via UNP profiles)

As a general recommendation:

- Use static SAPs on BEB-facing links toward access switches, where VLAN tagging must be preserved and the service attachment is predictable.
- Use dynamic SAPs on ports where end devices connect directly to BEBs, to reduce repetitive configuration and improve consistency.

Static SAP

In the reference topology, access switches connect to BEBs using tagged VLANs. To preserve VLAN tagging, create static SAPs with an encapsulation value matching the customer VLAN.

Configure SAPs on the relevant BEB interfaces according to the Addressing Plan:

18. Example on BEB-3 · CLI

```
service 1101 sap linkagg 31:1101
service 1101 sap linkagg 32:1101
service 1102 sap linkagg 31:1102
service 1102 sap linkagg 32:1102
service 1103 sap linkagg 31:1103
service 1103 sap linkagg 32:1103
service 1104 sap linkagg 31:1104
service 1104 sap linkagg 32:1104
service 1105 sap linkagg 31:1105
service 1105 sap linkagg 32:1105
service 1106 sap linkagg 31:1106
service 1106 sap linkagg 32:1106
service 1107 sap linkagg 31:1107
service 1107 sap linkagg 32:1107
service 1108 sap linkagg 31:1108
service 1108 sap linkagg 32:1108
service 1109 sap linkagg 31:1109
service 1109 sap linkagg 32:1109
service 2001 sap linkagg 31:2001
service 2001 sap linkagg 32:2001
service 2002 sap linkagg 31:2002
service 2002 sap linkagg 32:2002
service 2003 sap linkagg 31:2003
service 2003 sap linkagg 32:2003
```

Conceptually, this maps traffic arriving on a given port/LAG with VLAN X to service X.

Verify SAP creation using:

19. Example on BEB-3 · CLI

```
BEB-3-> show service access
Legend: (-)ERP Ring (*) Hybrid
Port      Link  SAP     SAP     Vlan Id
          Status Type    Count   Xlation L2Profile                    Description
---------+------+-------+-------+-------+-------------------------------+-----------
 0/31      Up     Manual  12       Y       l2prof
 0/32      Up     Manual  12       Y       l2prof

Total Access Ports: 2
```

The output summarizes SAPs configured per interface. To view SAPs associated with a specific service (including remote endpoints), use:

20. Example on BEB-3 · CLI

```
BEB-3-> show service 1101 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 1101 Info
  Admin : Up,      Oper : Up,   Stats     : N,        Mtu    : 9194,   VlanXlation : Y,
  ISID  : 1101,    BVlan : 4001,MCast-Mode : Headend, Tx/Rx  : 0/0,    RemoveIngTag: N

                                    Sap Trusted:Priority/         Sap Description /
Identifier           Adm  Oper Stats Sdp SystemId:BVlan   Intf     Sdp SystemName

---------------------+----+----+-----+-------------------+--------+----------
sap:0/31:1101          Up   Up    N           Y:x           0/31        -
```

```
sap:0/32:1101          Up    Up     N           Y:x           0/32        -

Total Ports: 2
```

This output should confirm that required services (for example, service 1101) have SAPs on the expected BEBs across the fabric.

Dynamic SAP

Static SAP creation can become repetitive when many access ports require similar configuration. For example, provisioning dozens of employee access ports as individual SAPs increases the number of commands and the risk of configuration drift.

Dynamic SAPs address this by using UNP profiles, which map traffic to services based on tag handling and classification rules. UNP profiles can be applied:

- Statically to ports, or
- Dynamically based on authentication (802.1X, MAC) and/or classification (for example, VLAN tag)

Dynamic SAP configuration typically includes:

- Create a UNP profile and map it to correct encapsulation/tag value and servie-id
- Define the interfaces as UNP access ports
- Applying UNP profiles statically or dynamically

In the reference topology, ports 1/1/20-1/1/40 of BEB-3 are used for employee endpoints. These devices send untagged traffic that must be mapped to service 1102 (which is already provisioned on BEB-3).

Create the UNP profile(s) and define the UNP access ports, applying l2prof as the L2 profile:

21. BEB-3 · CLI

```
unp profile Employee map service-type static tag-value 0 service-id 1102
unp port 1/1/20-40 port-type access
unp port 1/1/20-40 l2-profile l2prof
```

Static profile assignment

In the reference topology, the UNP profile is statically applied to ports 1/1/20-1/1/30:

22. BEB-3 · CLI

```
unp port 1/1/20-30 profile Employee
```

Dynamic profile assignment (example)

For ports 1/1/31-1/1/40, the profile is assigned dynamically based on authentication/classification rules. The following is a simplified example to illustrate the approach. (Items such as quarantine profiles and full AAA policy are out of scope for this guide; refer to the “AOS Network Configuration User Guide” for more information.)

23. BEB-3 · CLI

```
unp port-template SAMPLE FLOW direction both aaa-profile "UPAM" default-profile Quarantine
unp port-template "SAMPLE FLOW" 802.1x-authentication
unp port-template "SAMPLE FLOW" 802.1x-authentication pass-alternate "Quarantine"
unp port-template "SAMPLE_FLOW" mac-authentication
unp port-template "SAMPLE FLOW" mac-authentication pass-alternate "Quarantine"
unp port-template "SAMPLE FLOW" l2-profile "l2prof"
unp port 1/1/31-40 port-template SAMPLE_FLOW
```

Verify dynamic SAP creation using the same SAP verification commands used earlier.

24. BEB-3 · CLI

```
BEB-3-> show service access
Legend: (-)ERP Ring (*) Hybrid
```

```
Port       Link  SAP     SAP     Vlan
Id        Status Type    Count   Xlation L2Profile                          Description
---------+------+-------+-------+-------+--------------------------------+----------------------
----------
1/1/20    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/21    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/22    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/23    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/24    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/25    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/26    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/27    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/28    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/29    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/30    Down   Dynamic 1       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/31    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/32    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/33    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/34    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/35    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/36    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/37    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/38    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/39    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
1/1/40    Down   Dynamic 0       Y       l2prof                UNP Dynamic Access Port(1X-Auth)
0/31      Up     Manual  9       Y       l2prof
0/32      Up     Manual  9       Y       l2prof

Total Access Ports: 23
```

The output shows “Dynamic” SAP type in addition to the manual SAPs created earlier. Note the “SAP Count” column, it shows the SAPs are only created for port 1/1/20-30 but not for port 1/1/31-40. This is because the UNP profile “Employee” is applied to the former set of ports but not yet to the latter.

The command below highlights this difference.

25. BEB-3 · CLI

```
BEB-3-> show service 1102 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 1102 Info
  Admin : Up,    Oper  : Up,     Stats    : N,          Mtu    : 9194,   VlanXlation : Y,
  ISID  : 1102,  BVlan : 4002,   MCast-Mode : Headend, Tx/Rx   : 0/0,    RemoveIngTag: N

                                      Sap Trusted:Priority/          Sap Description /
Identifier             Adm  Oper Stats Sdp SystemId:BVlan   Intf     Sdp SystemName
----------------------+----+----+-----+------------------+--------+--------------------
------sap:1/1/20:0*          Up   Down  N           Y:x           1/1/20   Dynamic SAP for
UNP
sap:1/1/21:0*          Up   Down  N           Y:x           1/1/21   Dynamic SAP for UNP
sap:1/1/22:0*          Up   Down  N           Y:x           1/1/22   Dynamic SAP for UNP
sap:1/1/23:0*          Up   Down  N           Y:x           1/1/23   Dynamic SAP for UNP
sap:1/1/24:0*          Up   Down  N           Y:x           1/1/24   Dynamic SAP for UNP
sap:1/1/25:0*          Up   Down  N           Y:x           1/1/25   Dynamic SAP for UNP
sap:1/1/26:0*          Up   Down  N           Y:x           1/1/26   Dynamic SAP for UNP
sap:1/1/27:0*          Up   Down  N           Y:x           1/1/27   Dynamic SAP for UNP
sap:1/1/28:0*          Up   Down  N           Y:x           1/1/28   Dynamic SAP for UNP
sap:1/1/29:0*          Up   Down  N           Y:x           1/1/29   Dynamic SAP for UNP
sap:1/1/30:0*          Up   Down  N           Y:x           1/1/30   Dynamic SAP for UNP
sap:0/31:1102          Up   Up    N           Y:x           0/31       -
sap:0/32:1102          Up   Up    N           Y:x           0/32       -

Total Ports: 13
```

Dynamic SAP and dynamic services

Static and Dynamic SAP creation discussed till now require the pre-configuration of respective services on the BEB switches. This however may not be optimal. Consider a scenario where in future an additional BEB

switch is added to the reference topology in BLDG-C. All services would have to be replicated on this new BEB. This can lead to unnecessary configuration complexity and security risks.

For such scenarios, AOS offers dynamic services feature. On a high level, this feature allows users to define UNP profiles with appropriate I-SID and tag-value. Whenever this UNP profile gets assigned to a port, UNP first checks if a SAP already exists for the given ISID and tag-value. If not, a dynamic service and SAP are created for the given access port.

In the reference topology, services (security, camera, IoT) on BEB-5 are introduced using dynamic services (and later used as E-Tree services, covered in the E-Tree section).

Configure the UNP profiles required for dynamic service creation:

26. BEB-5 · CLI

```
unp profile "IoT"
unp profile "Camera"
unp profile "Security"
unp profile "Camera" map service-type spb tag-value 0 isid 2001 bvlan 4002 multicast-mode
tandem vlan-xlation
unp profile "IoT" map service-type spb tag-value 0 isid 2002 bvlan 4001 e-tree multicast-
mode headend vlan-xlation
unp profile "Security" map service-type spb tag-value 0 isid 2003 bvlan 4002 e-tree
multicast-mode headend vlan-xlation
unp port 1/1/10 port-type access
unp port 1/1/10 l2-profile l2prof
unp port 1/1/10 profile Security
```

In this example, the UNP profile is statically assigned to port 1/1/10 (it can also be dynamically assigned using the methods shown earlier).

Verify that the dynamic service has been created:

27. BEB-5 · CLI

```
BEB-5-> show service
Legend: * denotes a dynamic object
All Service Info
            Svc                    SAP     Bind
ServiceId   Type  Adm  Oper Stats Count   Count   Description
-----------+-----+----+----+-----+-------+-------+---------------------------
1000        SPB   Up   Up    N    0       4       Mgt
1001        SPB   Up   Up    N    0       4       Corp
1002        SPB   Up   Up    N    0       4       Guest
1003        SPB   Up   Up    N    0       4       Voice
1004        SPB   Up   Up    N    0       4       Util
1301        SPB   Up   Up    N    1       0
1302        SPB   Up   Up    N    1       0
1303        SPB   Up   Up    N    1       0
1304        SPB   Up   Up    N    1       0
1305        SPB   Up   Up    N    1       0
1306        SPB   Up   Up    N    1       0
1307        SPB   Up   Up    N    1       0
1308        SPB   Up   Up    N    1       0
1309        SPB   Up   Up    N    1       0
32781*      SPB   Up   Up    N    1       2       Dynamic Service isid=2003 for UNP

Total Services: 15
```

In the output, note that service 32781 is dynamically created and associated with I-SID 2003. To view more details:

28. BEB-5 · CLI

```
BEB-5-> show service 32781 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 32781 (Dynamic Service isid=2003 for UNP)
  Admin : Up,     Oper  : Up,     Stats   : N,          Mtu     : 9194,   VlanXlation : Y,
  ISID  : 2003,   BVlan : 4002,   MCast-Mode : Headend, Tx/Rx   : 0/0,    RemoveIngTag: N,
  ETree-Enb: Y
```

```
                                        Sap Trusted:Priority/          Sap Description /
Identifier              Adm   Oper Stats Sdp SystemId:BVlan    Intf      Sdp SystemName
----------------------+----+----+-----+--------------------+--------+--------------------
-----------
sap:1/1/10:0*           Up    Down  N           Y:x           1/1/10    Dynamic SAP for UNP
sdp:32781:32781*        Up    Up    N     7824.5974.93ae:4002  0/25     BEB-4
sdp:32782:32781*        Up    Up    N     7824.59a2.2447:4002  0/25     BEB-3

Total Ports: 3
```

The output shows more details about this dynamically created service. A dynamic SAP is also created for the given UNP port. Remote SDPs on BEB-3 & 4 are also learnt for this service. Note that the service-id used on BEB-3 & 4 is 2003 which is different from the one dynamically assigned here. Yet, this works because the service-id identifier is local to the BEB. Both services however use the same I-SID 2003 and BVLAN 4002 which remain consistent.

Dynamic UNP services

In some environments, defining UNP profiles ahead of time may not be practical—such as highly dynamic VM environments where VLANs appear frequently and unpredictably. For these scenarios, AOS supports a mode where UNP profiles and services are created automatically based on observed traffic. The service ID, I-SID, and BVLAN are derived using a deterministic calculation so that the same VLAN is mapped consistently across BEBs.

By default, the service modulo (used in I-SID calculation) is 512, which can cause multiple VLANs to map to the same I-SID. To prevent collisions, this guide recommends setting the service modulo to 4096 on all BEBs to maintain consistent behavior across the fabric.

Modify the service modulo on all BEBs:

29. All BEBs · CLI

```
unp system-default service-mod 4096
```

In the reference topology, dynamic services are enabled on UNP ports 1/1/20-1/1/22 of BEB-5, BEB-6, and BEB-7. Enable the dynamic services feature on the relevant BEBs:

30. BEB-5, 6 & 7 · CLI

```
unp port 1/1/20-22 port-type access
unp port 1/1/20-22 l2-profile l2prof
unp port 1/1/20-22 dynamic-service spb
```

Service and SAP configuration using OV

OmniVista (OV) can be used to create and manage SPB services and SAPs from a centralized interface. To access this workflow, navigate to:

Configure → LAN → Advanced Network → Service Configuration

This page lists existing SPB services and key attributes such as Service ID, BVLAN, assigned switches, and multicast mode.

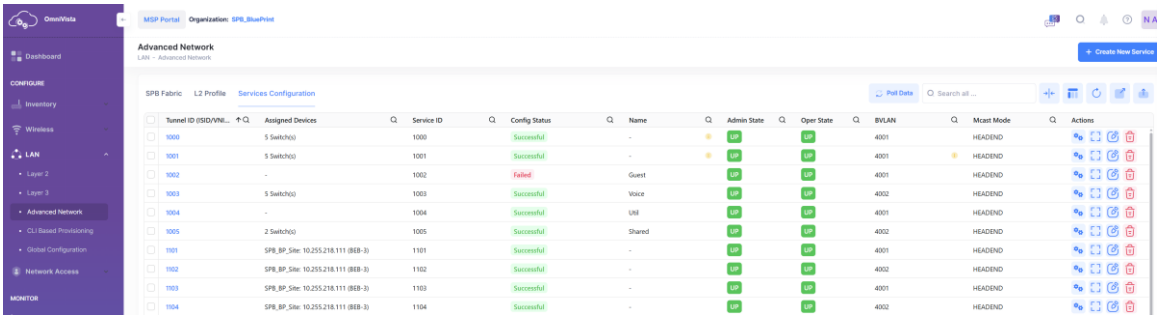

Figure 11. Service Configuration screen on OV

This page lists existing SPB services and key attributes such as Service ID, BVLAN, assigned switches, and multicast mode.

At present, OV supports static service and SAP provisioning. SAPs can be created on both UNP and non-UNP access ports. The OV workflow is organized into the following steps:

- SPB Service Settings
- Device Selection
- SAP Configuration
- UNP Access Port Configuration
- Configure UNP Mapping
- Preview and confirm

Create a new service by clicking on the Create New Service button on the top right of this screen.

Enter the desired values for service ID, I-SID and BVLAN. Choose the desired values for other options.

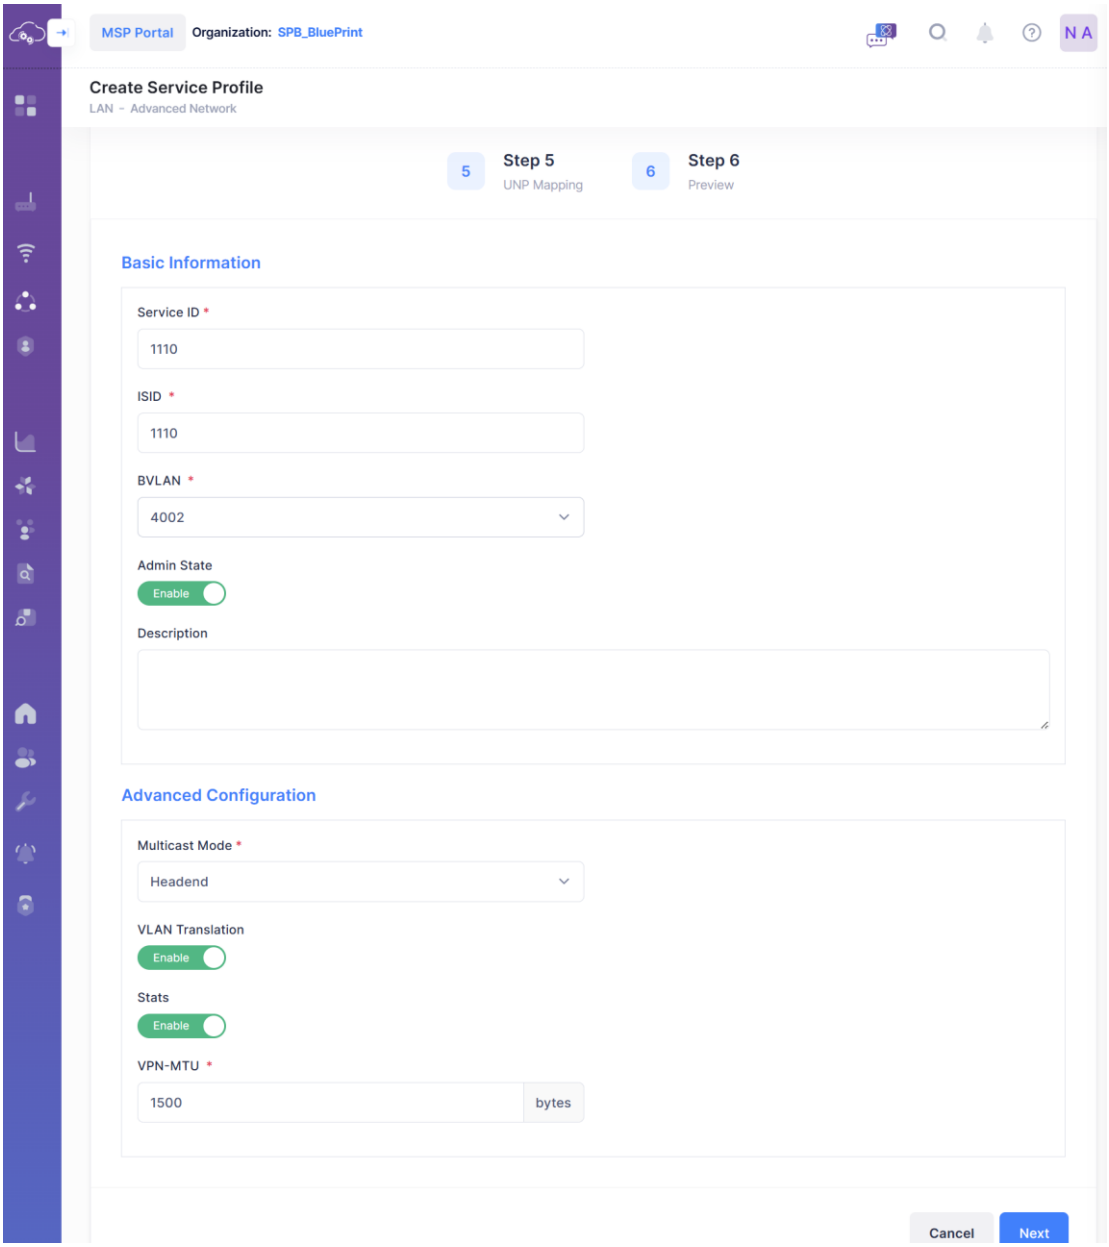

Figure 12. Service definition

On the next page, select the switch(es) where the service should be created.

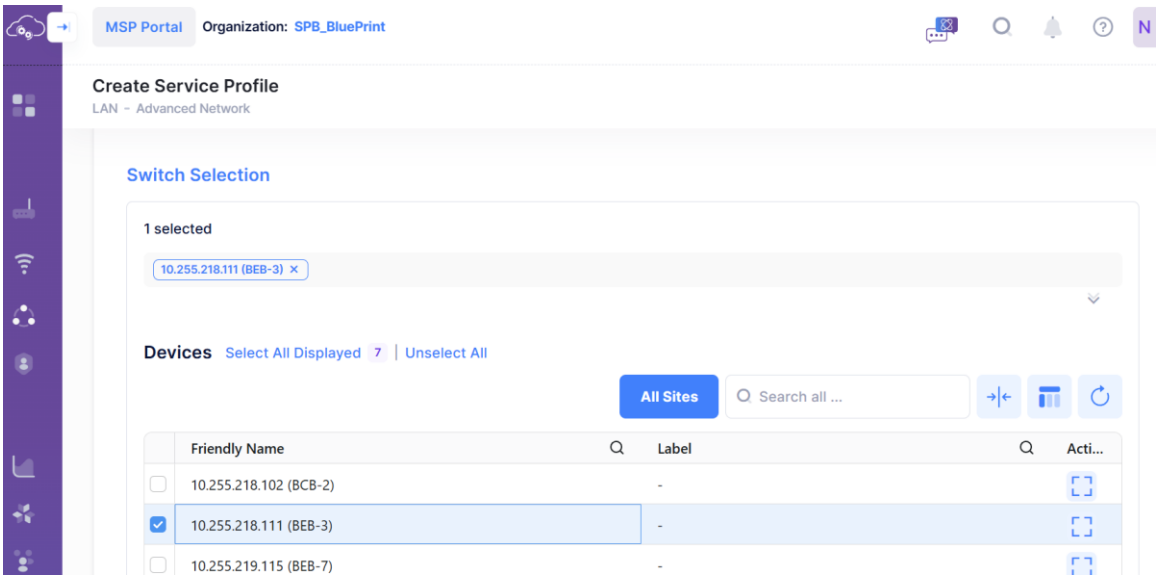

Figure 13. Device selection

This next step is used to create SAPs on non-UNP access ports. This is typically preferred for BEB interfaces that connect to access switches. If SAPs are not required on non-UNP ports, this step can be skipped.

On this page:

- Under Global SAP Settings, define the tag value used for SAP creation.
- If the interface is not already configured as a service access interface, enable Automatic Access Interface Configuration to have OV apply the required access interface configuration automatically.
- Select the appropriate L2 profile and other options. If the required L2 profile does not exist, it can be created by selecting Create next to the L2 profile field.
- Under Select Ports, choose the interface(s) where SAPs should be created.

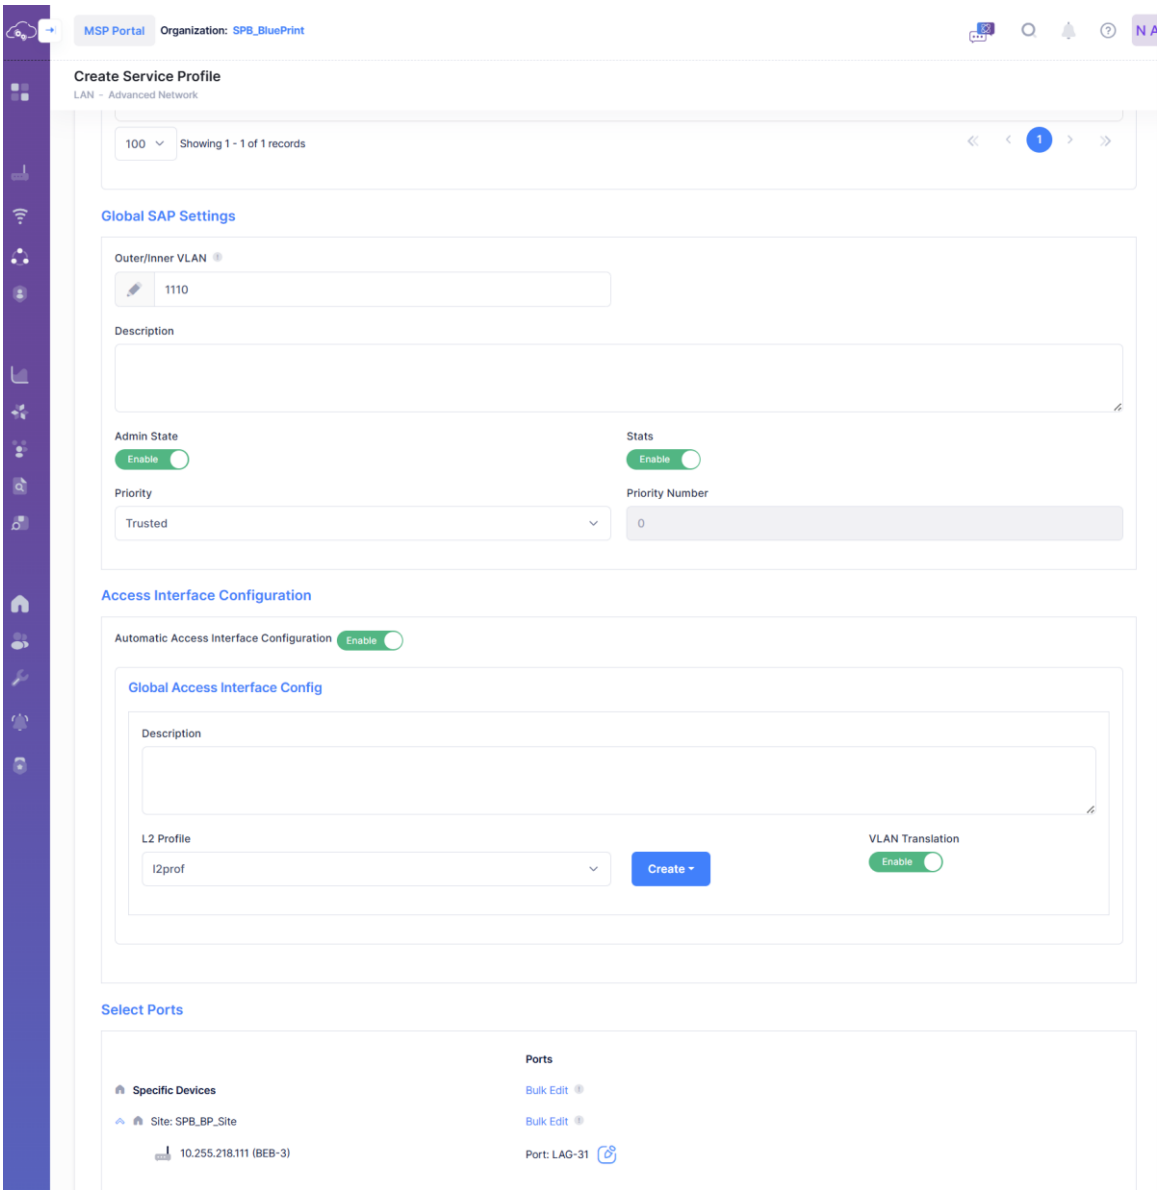

Figure 14. Non-UNP Port SAP creation

Next, configure the UNP access ports on the selected devices:

- Select the appropriate UNP port template to apply. If one does not exist, create it by selecting Create next to the field.
- Select the ports that should be defined as UNP access ports.

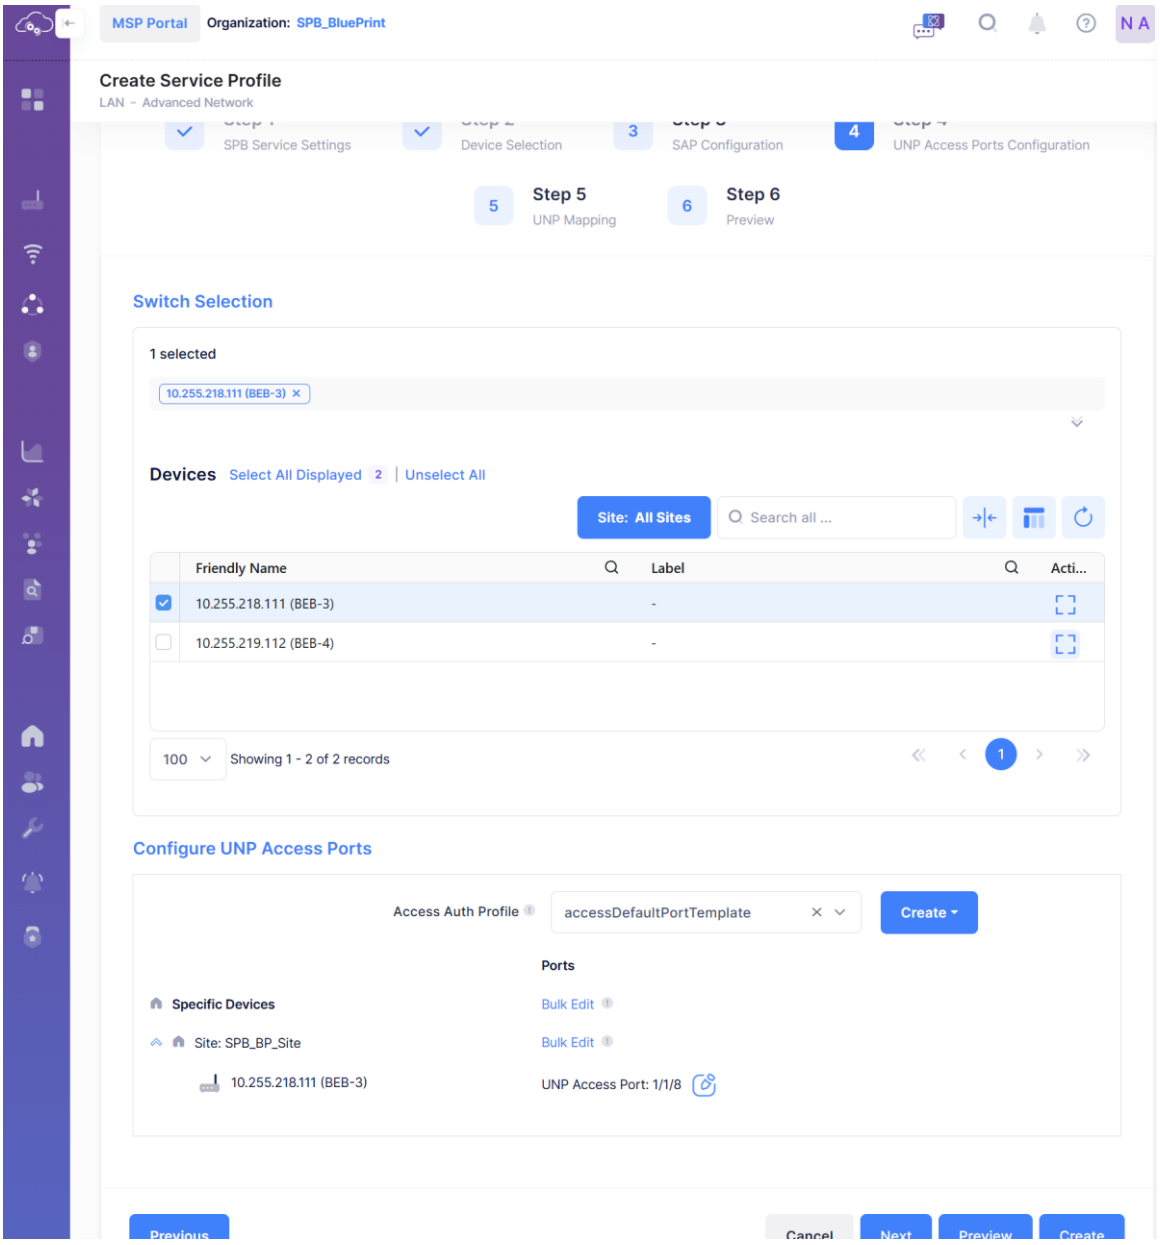

Figure 15. UNP Access Port creation

Next, configure how tagged/untagged client traffic is mapped to services on the UNP ports:

- Select the appropriate tag value and UNP profile.
- If the required UNP profile does not exist, create it by selecting Create next to the field.

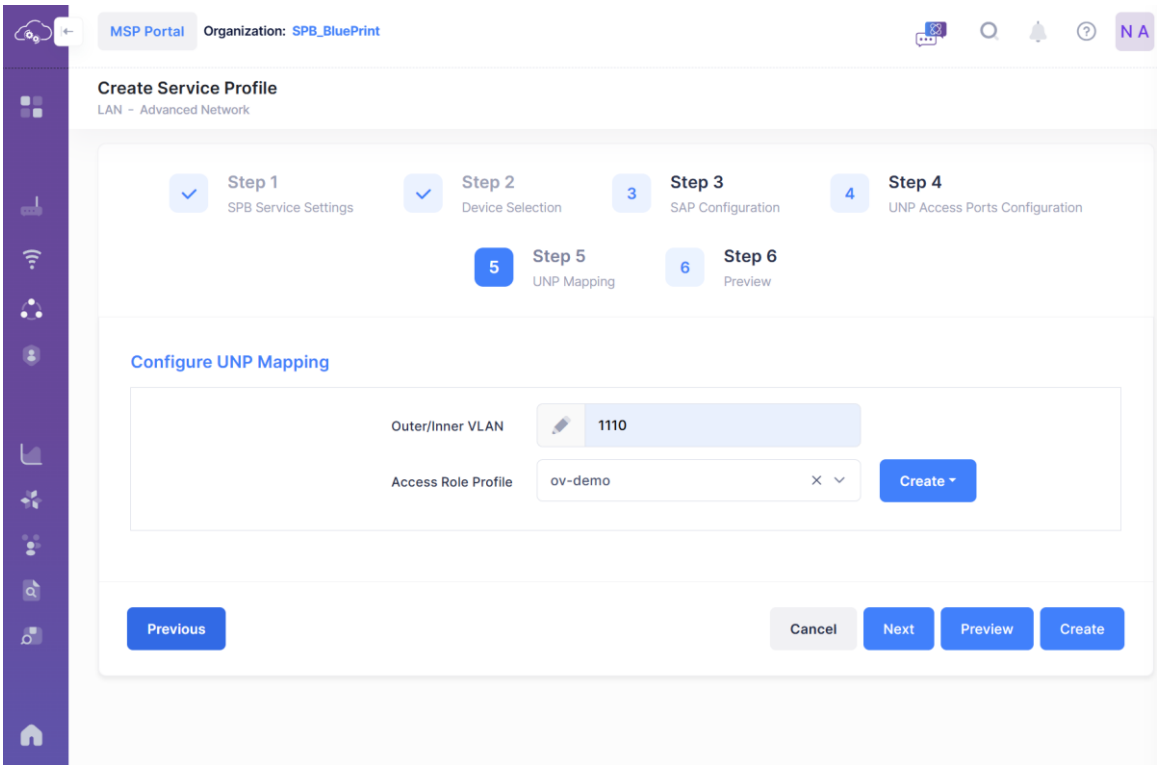

Figure 16. UNP profile mapping

In the final step, review the configuration summary and click Create to apply the changes.

Once completed, OV provisions the selected configuration elements on the target device(s), including the service, SAPs, UNP access port settings, and UNP profiles/mappings (as applicable). The deployment can be verified by returning to the “Service Configuration” view and checking the service and port status.

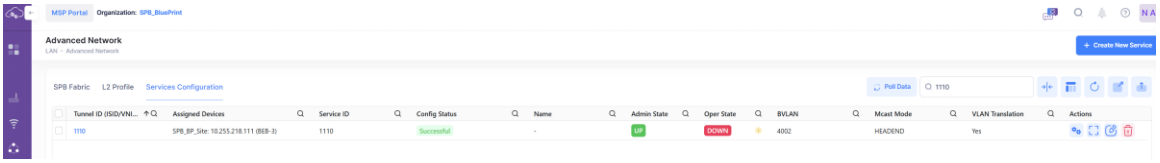

Figure 17. Service creation confirmation

Stellar Access Points

AOS provides a streamlined method to integrate Stellar Access Points (APs) into an SPB fabric. When a Stellar AP is connected to an OmniSwitch BEB, the switch can automatically discover and learn the AP and place it under management. After discovery, the AP is associated with a preconfigured SPB management service, which is used for AP control and management connectivity.

Once the AP is operational, it forwards client traffic as VLAN-tagged frames toward the BEB. The BEB then attempts to classify that tagged client traffic into a UNP service profile whose configured service tag value matches the client VLAN tag. If a matching UNP profile is not found, the traffic is handled by the System Default profile, which can dynamically create the required SAP (and, if enabled, the service) to forward AP client traffic.

> **NOTE** APs may also be connected to access switches with VLANs configured at the access layer. In that model, the access switch forwards tagged traffic to the BEB, and the BEB maps the traffic to SPB services in the same way as any other VLAN-based service described earlier.

The configuration for this feature involves the following steps:

- Configure a Layer 2 profile with a 'peer' action defined for 802.1AB and 802.1X control frames.
- Configure the port as UNP access port and assign the Layer 2 profile created in the previous step.

- If necessary, enable the UNP AP mode for the UNP access port.
- By default, 802.1X and MAC authentication are enabled on UNP ports. If authentication of an AP device is not required, disable one or both of these options.
- Map appropriate SPB service parameters to the built-in 'defaultWLANAccessProfile'.
- Create UNP profiles mapped to SPB service parameters for learning AP client MAC addresses.
- Create UNP classification rules to capture and assign tagged AP client traffic into the UNP service profile configured with a matching VLAN tag value.

In the reference topology, service 1101, 1201, 1301 and 1401 are used as the management service for APs in respective BEBs. AP client traffic is then mapped to the appropriate services based on building and role, using UNP profiles and classification rules. The following image shows the logical mapping for an AP connected to BEB-3.

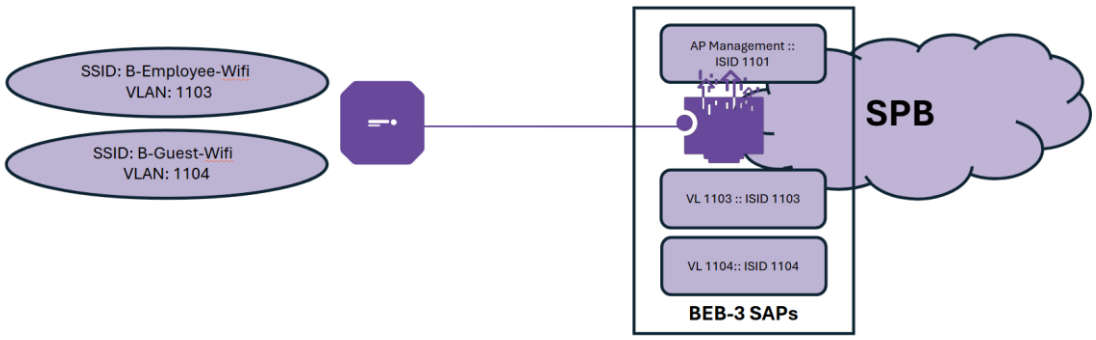

Figure 18. AP integration with SPB network

The l2prof profile created earlier is reused for AP-facing UNP ports in this guide. (If stronger separation is required, a dedicated AP-specific L2 profile may be used.) Define the UNP access ports and apply the L2 profile:

31. Example on BEB-3 · CLI

```
unp port 1/1/19 port-type access
unp port 1/1/19 l2-profile l2prof
```

Map the AP management service to the built-in defaultWLANAccessProfile:

32. Example on BEB-3 · CLI

```
unp profile "defaultWLANAccessProfile" map service-type spb tag-value 0 isid 1101 bvlan
4001 vlan-xlation
```

The command above will dynamically create the service/SAP for the specified tag value and I-SID if it does not already exist. If service 1101 is statically provisioned, the following command may be used instead:

33. Example on BEB-3 · CLI

```
service 1101 spb isid 1101 bvlan 4001 description "Access Points" vlan-xlation enable
unp profile "defaultWLANAccessProfile" map service-type static tag-value 0 service-id 1101
```

Create UNP profiles for AP client traffic and map them to the corresponding services:

34. Example on BEB-3 · CLI

```
unp profile "Employee-Wifi"
unp profile "Employee-Wifi" map service-type spb tag-value 1103 isid 1103 bvlan 4001 vlan-
xlation
unp profile "Guest-Wifi"
unp profile "Guest-Wifi" map service-type spb tag-value 1104 isid 1104 bvlan 4002 vlan-
xlation
```

The UNP profiles instead can be mapped to preconfigured services if they already exisit on the switch.

Finally, create UNP classification rules to assign tagged AP client traffic to the UNP profile that matches the VLAN tag:

35. Example on BEB-3 · CLI

```
unp classification vlan-tag 1103 profile1 "Employee-Wifi"
unp classification vlan-tag 1104 profile1 "Guest-Wifi"
```

To validate the configuration, connect a Stellar AP to the configured UNP access port and enable PoE. Then verify that a SAP has been created for the AP:

36. Example on BEB-3 · CLI

```
BEB-3-> show service 1101 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 1101 Info
  Admin : Up,     Oper  : Up,   Stats     : N,       Mtu   : 9194,   VlanXlation : Y,
  ISID  : 1101,   BVlan : 4001, MCast-Mode: Headend, Tx/Rx : 0/0,    RemoveIngTag: N

                                      Sap Trusted:Priority/          Sap Description /
Identifier             Adm  Oper Stats Sdp SystemId:BVlan   Intf     Sdp SystemName
-----------------------+----+----+-----+-------------------+--------+---------------------
-sap:1/1/19:0*          Up   Up    N           Y:x           1/1/19   Dynamic SAP for UNP
sap:0/31:1101          Up   Up    N           Y:x           0/31       -
sap:0/32:1101          Up   Up    N           Y:x           0/32       -

Total Ports: 3
```

The output should show a dynamic SAP created for the AP using service 1101. If dynamic services are used instead, the service ID would be dynamically assigned on the BEB; service interoperability is maintained as long as the I-SID and BVLAN selection remain consistent across the fabric.

Multicast switching

SPB supports two methods for replicating and forwarding BUM traffic: head-end replication and tandem replication.

In head-end mode, traffic replication is performed by the ingress BEB. The ingress BEB creates a separate copy of the traffic for each destination and forwards it accordingly. In tandem mode, traffic is replicated only at the required fork points along the path toward the destination BEBs. For additional details on the behavior and design considerations for each mode, refer to the “SPB Design Guide”.

Head-end mode typically consumes more bandwidth because replicated traffic is sent from the ingress BEB toward each destination. For this reason, it is generally preferred when the number of receivers is low. Tandem mode is more bandwidth-efficient because replication occurs only where required in the SPB fabric. However, it consumes additional switch resources and is therefore better suited for deployments with a larger number of receivers.

In the reference topology, security cameras transmit multicast video streams on service 2001. This service is configured to use tandem multicast mode.

Use the following command to create service 2001 with tandem multicast mode on all BEBs.

37. All BEBs · CLI

```
service 2001 spb isid 2001 bvlan 4002 multicast-mode tandem vlan-xlation enable
```

Note that service 2001 on BEB-5 is created as a dynamic service, as described earlier in the “Dynamic SAP and Dynamic Service” section of this chapter.

Within the Layer 2 network, multicast traffic distribution is optimized by enabling IGMP snooping. IGMP snooping ensures that multicast traffic is forwarded only to devices that have explicitly requested the multicast stream. In AOS, the IGMP snooping implementation is referred to as IP Multicast Switching (IPMS).

IPMS must be enabled on all devices where multicast traffic is expected to be received or forwarded. IPMS can be enabled either per service/VLAN or globally at the switch level.

Use the following command to enable IPMS for service/VLAN 2001 on all BEB and ACC switches:

38. All BEBs · CLI

```
ip multicast service 2001 admin-state enable
```

39. All ACC switches · CLI

```
ip multicast vlan 2001 admin-state enable
```

Additional IGMP parameters, such as IGMP querying and IGMP querier forwarding, must be configured on the appropriate devices based on the multicast design. For more information, refer to the “AOS Network Configuration User Guide”.

To verify the multicast service configuration, use the following command:

40. Example on BEB-3 (output truncated) · CLI

```
BEB-3-> show service 2001
SPB Service Detailed Info
  Service Id       : 2001,                 Description      :                    ,
  ISID             : 2001,                 BVlan            : 4002,
  Multicast-Mode   : Tandem,               Tx/Rx Bits       : 1/1,
  Admin Status     : Up,                   Oper Status      : Up,
  Stats Status     : No,                   Vlan Translation : Y,
  Service Type     : SPB,                  Allocation Type  : Static,
```

In the command output, verify that the multicast mode is shown as Tandem on all BEBs.

To display the multicast forwarding information, use the following command:

41. Example on BEB-3 · CLI

```
BEB-3-> show spb isis multicast-table
Legend: MCAST Source * indicates any source in GMODE bvlans
SPB ISIS Multicast MAC Table:
                                    MCAST Source                Inbound     Outbound
 ISID   BVLAN  MCAST Group Address  (Name : BMAC)               Interface   Interfaces
------+-------+-------------------+-----------------------------------------+-----------+---
  2001   4002   23:24:47:00:07:d1  BEB-3     : 78:24:59:a2:24:47                  0/23
  2001   4002   43:93:ae:00:07:d1  BEB-4     : 78:24:59:74:93:ae        0/23
  2001   4002   93:70:6d:00:07:d1  BEB-6     : 94:24:e1:59:70:6d        0/23
  2001   4002   a3:30:9d:00:07:d1  BEB-7     : 94:24:e1:3a:30:9d        0/23

Mac Addresses: 4
```

In the output, note that the forwarding table does not contain an entry for BEB-5. This is expected because the corresponding dynamic service has not yet been instantiated on BEB-5.

E-Tree service model

SPB supports three service models: E-Line, E-LAN, and E-Tree. Unless explicitly configured otherwise, SPB services operate in the E-LAN model (any-to-any connectivity within the service).

In the reference topology, most services are configured as E-LAN. Services associated with the landlord domain—Security and IoT—are configured as E-Tree to enforce a hub-and-spoke communication pattern. In an E-Tree service:

- Leaf SAPs can communicate only with Root SAPs
- Root SAPs can communicate with both Leaf SAPs and other Root SAPs

This design is particularly well suited for environments in which endpoints such as IoT devices need access to centralized servers, controllers, or management platforms, but do not require direct device-to-device communication. Structuring these services as E-Tree provides a form of network segmentation that limits

east-west traffic between endpoint devices and reduces the attack surface. In the event of a compromise, such as malware infection or unauthorized access to a single endpoint, this model helps contain potential lateral movement and prevents direct communication with peer devices on the same service. For these reasons, using E-Tree for landlord-operated Security and IoT networks is considered a best-practice approach for improving security and operational control.

In this deployment, the E-Tree services are implemented with:

- Leaf nodes: BEB-5, BEB-6, BEB-7
- Root nodes: BEB-3, BEB-4

NOTE: BEB-5 uses dynamic service creation, as described earlier in the “Dynamic Services” section of this chapter. E-tree configuration is hence part of the UNP profile in this case.

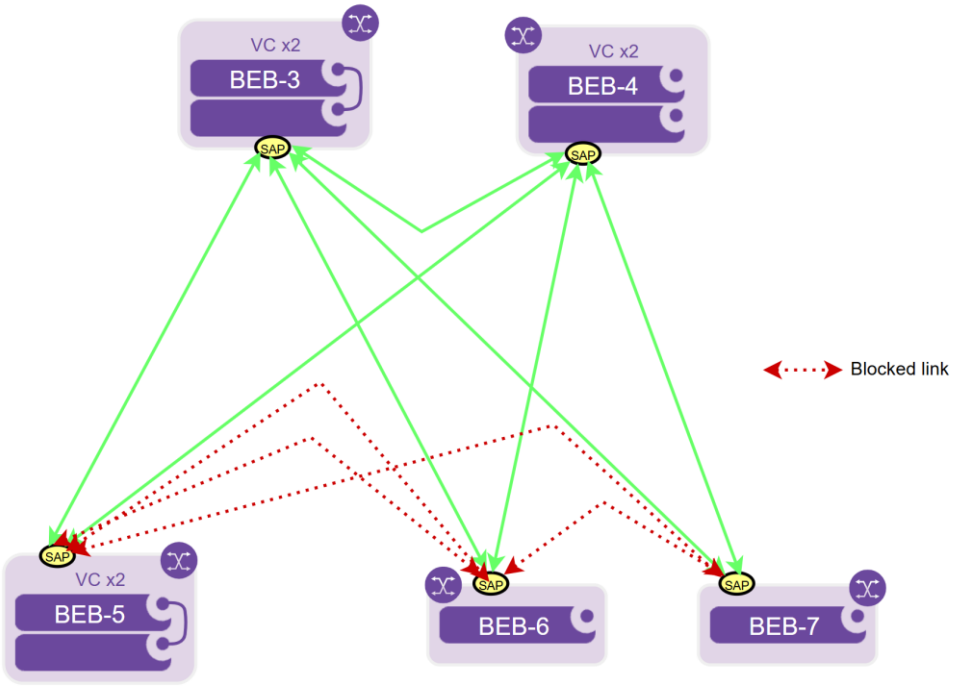

Figure 19. E-Tree services

Configure the E-Tree services on leaf nodes (BEB-6 and BEB-7):

42. BEB-6 & 7 (leaf nodes) · CLI

```
service 2002 spb isid 2002 bvlan 4001 vlan-xlation enable e-tree enable
service 2003 spb isid 2003 bvlan 4002 vlan-xlation enable e-tree enable
```

Configure the E-Tree service on root nodes (BEB-3 and BEB-4):

43. BEB-3 & 4 (root nodes) · CLI

```
service 2002 spb isid 2002 bvlan 4001 vlan-xlation enable
service 2003 spb isid 2003 bvlan 4002 vlan-xlation enable
```

Verify the E-Tree configuration using:

44. BEB-3 · CLI

```
BEB-3-> show service 2003 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 2003 Info
  Admin : Up,     Oper  : Up,   Stats      : N,       Mtu     : 9194,   VlanXlation : Y,
  ISID  : 2003,   BVlan : 4002, MCast-Mode : Headend, Tx/Rx   : 0/0,    RemoveIngTag: N
```

```
                                      Sap Trusted:Priority/          Sap Description /
Identifier              Adm  Oper Stats Sdp SystemId:BVlan   Intf      Sdp SystemName
----------------------+----+----+-----+-------------------+--------+--------------------
-----------
sdp:32782:2003*         Up   Up    N    7824.5974.93ae:4002  0/23      BEB-4
sdp:32783:2003*         Up   Up    N    9424.e13a.309d:4002  0/13      BEB-7
sdp:32784:2003*         Up   Up    N    9424.e159.706d:4002  0/23      BEB-6
sdp:32788:2003*         Up   Up    N    2cfa.a271.c2cd:4002  0/23      BEB-5

Total Ports: 4
```

On a root node (for example, BEB-3), the output would show remote SDPs learned from both root and leaf nodes (for example, SDPs on BEB-4 through BEB-7) because a root participates in communication with all nodes in the service.

Now run the same verification on a leaf node (for example, BEB-6):

45. BEB-6 · CLI

```
BEB-6-> show service 2003 ports
Legend: (*)Dyn unicast (+)Remote Mcast (#)Local Mcast (-)ERP Ring
SPB Service 2003 Info
  Admin : Up,     Oper  : Up,   Stats      : N,       Mtu     : 9194,   VlanXlation : Y,
  ISID  : 2003,   BVlan : 4002, MCast-Mode : Headend, Tx/Rx   : 0/0,    RemoveIngTag: N,
  ETree-Enb: Y

                                      Sap Trusted:Priority/          Sap Description /
Identifier              Adm  Oper Stats Sdp SystemId:BVlan   Intf      Sdp SystemName
----------------------+----+----+-----+-------------------+--------+--------------------
-----------
sdp:32780:2003*         Up   Up    Y    7824.5974.93ae:4002  0/26      BEB-4
sdp:32788:2003*         Up   Up    Y    7824.59a2.2447:4002  0/26      BEB-3

Total Ports: 2
```

Because BEB-6 is a leaf, it would learn SDPs only from root nodes—in this case, SDPs associated with BEB-3 and BEB-4—and not from other leaf nodes.

Validation

A straightforward method to validate the Layer 2 service configuration is to create temporary IP interfaces on the switches and verify IP connectivity between them using ping.

This section demonstrates a simple procedure to validate Layer 2 connectivity for the services configured in this chapter.

In this example, IP interfaces are created in the Employee VLAN 1102 on ACC-31 and ACC-32. Use the following commands to create the corresponding IP interfaces:

46. ACC-31 · CLI

```
ip interface "test" address 10.11.16.10 mask 255.255.240.0 vlan 1102
```

47. ACC-32 · CLI

```
ip interface "test" address 10.11.16.11 mask 255.255.240.0 vlan 1102
```

From each access switch, ping the IP interface configured on the peer access switch:

48. ACC-32 (Output truncated) · CLI

```
ACC-32-> ping 10.11.16.10
PING 10.11.16.10 (10.11.16.10) 56(84) bytes of data.
64 bytes from 10.11.16.10: icmp seq=1 ttl=64 time=0.625 ms
...
6 packets transmitted, 6 received, 0% packet loss, time 5045ms
rtt min/avg/max/mdev = 0.609/0.674/0.720/0.050 ms
```

Successful ping responses confirm that Layer 2 connectivity is operational for the service. Note that VLAN 1102 is not directly stretched between the two access switches. Instead, traffic is bridged across the SPB network through the configured service.

After completing the validation, remove the temporary test interfaces from both switches using the following command:

49. ACC-31 & ACC-32 · CLI

```
no ip interface "test"
```

Ping connectivity between same user group but located in different buildings would fail as they are mapped to different VLANS and services. To enable communication between these user groups, traffic must be routed between the corresponding networks. Inter-service routing is covered in the next chapter.

Layer 3 services

A Layer 3 service refers to a type of VPN service connecting multiple sites in a single any-to-any routing domain. Different sites utilize different subnets and require routing to communicate. For multi-tenancy, and to keep different users isolated at Layer 3, each user group is associated to its own VRF instance.

Layer 3 services in SPB can be offered either using L3VPN or VPN-lite solutions. VPN-lite requires the configuration of a routing protocol whereas L3VPN uses the underlying IS-IS. L3VPN is generally the better choice due to simplicity and scalability. It is recommended to use L3VPN within the SPB domain and VPN-lite at border BEBs which interact with external non IS-IS networks.

In the reference topology, each building uses different subnets and hence requires routing for communication. This is achieved using L3VPN. VPN-lite is used at border BEBs BEB-3 & BEB-4 for external connectivity.

Configuring a L3VPN solution involves the following steps:

- Creating Layer 2 service for WAN interfaces
- Creating Virtual Routing and Forwarding instances (VRF)
- Defining LAN side and WAN side IP interfaces
- Route import/export between VRF and SPB I-SID instance
- Binding WAN side interface to Layer 2 SPB service's I-SID

Layer 2 service for WAN

The L3VPN (ISIS-SPB) approach requires configuring a VRF-I-SID binding to identify the L3VPN interface that will exchange routes between the VLAN domain (VRF instance) and the SPB service domain (I-SID). This SPB service will be used to exchange routes among BEBs.

In the reference topology, there are seven such services each for Mgt, Corp, Guest, Voice, Util, Shared and Landlord VRFs. Landlord VRF only offers Layer 2 services and hence does not require L3VPN. Landlord VRF is hence created only on BEB-3 & BEB-4 which act as VRRP routers for the respective services. Shared VRF primarily is used to exchange routes with external networks, and hence is only configured on border BEBs – BEB-3 & BEB-4.

Use the following commands to configure the services on all BEBs.

50. Example on BEB-3 · CLI

```
service 1000 spb isid 1000 bvlan 4001 description "Mgt" vlan-xlation enable
service 1001 spb isid 1001 bvlan 4002 description "Corp" vlan-xlation enable
service 1002 spb isid 1002 bvlan 4001 description "Guest" vlan-xlation enable
service 1003 spb isid 1003 bvlan 4002 description "Voice" vlan-xlation enable
service 1004 spb isid 1004 bvlan 4001 description "Util" vlan-xlation enable
service 1005 spb isid 1005 bvlan 4002 description "Shared" vlan-xlation enable
```

Virtual Routing and Forwarding (VRF)

To ensure traffic isolation, VRFs are configured on BEBs. In the reference topology, VRFs are created for each group of users.

Use the following commands to configure the VRFs on all BEBs (Landlord and Shared VRF only one BEB-3 and BEB4):

51. All BEBs · CLI

```
vrf create Mgt
vrf create Corp
vrf create Guest
vrf create Voice
vrf create Util
vrf create Shared
```

```
vrf create Landlord
```

Verify the VRF creation using the following command:

52. Example on BEB-3 · CLI

```
BEB-3-> show vrf
Virtual Routers     Profile Protocols
--------------------+-------+------------------
default              default OSPF VRRP
Mgt                  max     VRRP
Corp                 max     VRRP
Guest                max     VRRP
Voice                max     VRRP
Util                 max     VRRP
Shared               max     VRRP
Landlord             max     VRRP
Total Number of Virtual Routers: 8
```

Layer 3 IP interfaces

An L3 IP interface serves as a gateway to access remote networks and is required when using either the VPN-Lite or L3 VPN method. Each BEB that requires routing needs to have an interface configured in that specific VRF.

> **NOTE** This document specifies the configuration related to service based inline routing only. Refer to the 'AOS Network Configuration User Guide' for external loopback configuration.

In the reference topology, each site uses a different subnet. Hence, each BEB would have its own set of LAN side and WAN side IP interfaces. LAN side interfaces serve as the gateway for users in a given VLAN/service. These would hence be created for each VLAN/service as mentioned in the addressing plan in their respective VRFs. WAN side IP interfaces on the other hand, function as gateways to access remote networks and are bind to the VRFs and associated I-SIDs. These are hence created per VRF on all BEBs.

Use the following command to create IP interfaces and associate them with their respective services.

53. Example on BEB-3 · CLI

```
vrf Mgt ip interface "Mgt-WAN" address 10.10.0.3 mask 255.255.255.0 service 1000
vrf Mgt ip interface "Access-Points" address 10.11.0.1 mask 255.255.240.0 service 1101

vrf Corp ip interface "Corp-WAN" address 10.10.1.3 mask 255.255.255.0 service 1001
vrf Corp ip interface "Employee" address 10.11.16.1 mask 255.255.240.0 service 1102
vrf Corp ip interface "Employee-Wifi" address 10.11.32.1 mask 255.255.240.0 service 1103

vrf Guest ip interface "Guest-WAN" address 10.10.2.3 mask 255.255.255.0 service 1002
vrf Guest ip interface "Guest-Wifi" address 10.11.48.1 mask 255.255.240.0 service 1104
vrf Guest ip interface "Contractor" address 10.11.64.1 mask 255.255.240.0 service 1105

vrf Voice ip interface "Voice-WAN" address 10.10.3.3 mask 255.255.255.0 service 1003
vrf Voice ip interface "Voice" address 10.11.80.1 mask 255.255.240.0 service 1106

vrf Util ip interface "Util-WAN" address 10.10.4.3 mask 255.255.255.0 service 1004
vrf Util ip interface "Storage" address 10.11.96.1 mask 255.255.240.0 service 1107
vrf Util ip interface "Sensors" address 10.11.112.1 mask 255.255.240.0 service 1108
vrf Util ip interface "Printers" address 10.11.128.1 mask 255.255.240.0 service 1109

vrf Shared ip interface "Shared-WAN" address 10.10.5.3 mask 255.255.255.0 service 1005

vrf Landlord ip interface "Camera" address 10.20.32.2 mask 255.255.224.0 service 2001
vrf Landlord ip interface "IoT" address 10.20.64.2 mask 255.255.224.0 service 2002
vrf Landlord ip interface "Security" address 10.20.96.2 mask 255.255.224.0 service 2003
```

Verify the IP interfaces are created with correct details:

54. Example on BEB-3 · CLI

```
BEB-3-> vrf Mgt show ip interface
Total 2 interfaces
 Flags (D=Directly-bound)
```

```
       (A=Anycast IP)

       Name         IP Address     Subnet Mask      Status Forward  Device           Flags
------------------------------+--------------+---------------+------+-------+----------
-Access-Points       10.11.0.1      255.255.240.0    UP      YES      service 1101
Mgt-WAN             10.10.0.3      255.255.255.0    UP      YES      service 1000
```

Virtual Router Redundancy Protocol (VRRP)

VRRP is a standard router redundancy protocol that provides high availability by allowing multiple routers to act as a single virtual router. This is a recommended practice for high availability scenarios and is not a requirement for Layer 3 services configuration.

In the reference topology, BEB-6 & BEB-7 serve the same networks. IP VRRP is configured on these BEBs for redundancy. Similarly, BEB-3 & BEB-4 serve as routers for the services in Landlord VRF. Hence these are also configured as VRRP.

Use the following commands to setup VRRP on BEB-6 and BEB-7:

55. Example on BEB-6 · CLI

```
vrf Mgt ip vrrp 1 interface "Access-Points" priority 120
vrf Mgt ip vrrp 1 interface "Access-Points" address 10.14.0.1
vrf Mgt ip vrrp 1 interface "Access-Points" admin-state enable

vrf Corp ip vrrp 1 interface "Employee"
vrf Corp ip vrrp 1 interface "Employee" address 10.14.16.1
vrf Corp ip vrrp 1 interface "Employee" admin-state enable
vrf Corp ip vrrp 2 interface "Employee-Wifi"
vrf Corp ip vrrp 2 interface "Employee-Wifi" address 10.14.32.1
vrf Corp ip vrrp 2 interface "Employee-Wifi" admin-state enable

vrf Guest ip vrrp 1 interface "Guest-Wifi" priority 120
vrf Guest ip vrrp 1 interface "Guest-Wifi" address 10.14.48.1
vrf Guest ip vrrp 1 interface "Guest-Wifi" admin-state enable
vrf Guest ip vrrp 2 interface "Contractor" priority 120
vrf Guest ip vrrp 2 interface "Contractor" address 10.14.64.1
vrf Guest ip vrrp 2 interface "Contractor" admin-state enable

vrf Voice ip vrrp 1 interface "Voice"
vrf Voice ip vrrp 1 interface "Voice" address 10.14.80.1
vrf Voice ip vrrp 1 interface "Voice" admin-state enable

vrf Util ip vrrp 1 interface "Storage" priority 120
vrf Util ip vrrp 1 interface "Storage" address 10.14.96.1
vrf Util ip vrrp 1 interface "Storage" admin-state enable
vrf Util ip vrrp 2 interface "Sensors" priority 120
vrf Util ip vrrp 2 interface "Sensors" address 10.14.112.1
vrf Util ip vrrp 2 interface "Sensors" admin-state enable
vrf Util ip vrrp 3 interface "Printers" priority 120
vrf Util ip vrrp 3 interface "Printers" address 10.14.128.1
vrf Util ip vrrp 3 interface "Printers" admin-state enable
```

Note that in the above configuration, interfaces on VRFs Mgt, Guest and Util are configured with a priority of 120 instead of the default value 100. Parallely, interfaces on VRFs Corp and Voice are configured with priority 120 on BEB-7. This is done for load balancing between the two BEBs. The BEB interface with highest priority is elected as the master.

To verify the VRRP status use the following command:

56. Example on BEB-6 · CLI

```
BEB-6-> vrf Mgt show ip vrrp statistics
Checksum Errors :          0,
Version Errors  :          0,
VRID Errors     :          0

               Interface
VRID             Name                  State      UpTime   Become Master Adv. Rcvd
----+--------------------------------+----------+----------+-------------+----------
   1 Access-Points                   Master         176842            1          4
```

The output shows that the IP interface in Mgt VRF was chosen as Master as intended.

VRRP Tracking

VRRP can cause traffic blackholing if the uplinks connecting the Master router are affected. Use VRRP tracking feature to monitor specific objects (interfaces or routes) on the master router and automatically reduce its priority if a failure occurs, forcing a failover to a backup router.

In the reference topology, VRRP tracking is configured on BEB-6 & BEB-7. The WAN interfaces in each VRF tracked to detect any uplink failures. Use the following commands to create VRRP tracking policies:

57. Example on BEB-6 · CLI

```
vrf Mgt ip vrrp track 1 admin-state enable priority 25 ipv4-interface "Mgt-WAN"
vrf Guest ip vrrp track 1 admin-state enable priority 25 ipv4-interface "Guest-WAN"
vrf Util ip vrrp track 1 admin-state enable priority 25 ipv4-interface "Util-WAN"
```

These VRRP tracking policies need to be associated with the respective interfaces. VRRP needs to be disabled before making any modifications. Use the following commands:

58. Example on BEB-6 · CLI

```
vrf Mgt ip vrrp admin-state disable
vrf Guest ip vrrp admin-state disable
vrf Util ip vrrp admin-state disble

vrf Mgt ip vrrp 1 interface "Access-Points" track-association 1
vrf Guest ip vrrp 1 interface "Guest-Wifi" track-association 1
vrf Guest ip vrrp 2 interface "Contractor" track-association 1
vrf Util ip vrrp 1 interface "Storage" track-association 1
vrf Util ip vrrp 2 interface "Sensors" track-association 1
vrf Util ip vrrp 3 interface "Printers" track-association 1

vrf Mgt ip vrrp admin-state enable
vrf Guest ip vrrp admin-state enable
vrf Util ip vrrp admin-state enable
```

Route import and export

The routes that will be exchanged between the VRF and the SPB service need to be specified. VRF import and export commands are used to achieve this.

In the reference topology, all routes are imported and exported within a VRF. Use the following commands on all BEBs to import and export routes within each VRF:

59. All BEBs · CLI

```
vrf Mgt ip import isid 1000 all-routes
vrf Corp ip import isid 1001 all-routes
vrf Guest ip import isid 1002 all-routes
vrf Voice ip import isid 1003 all-routes
vrf Util ip import isid 1004 all-routes
vrf Shared ip import isid 1005 all-routes

vrf Mgt ip export all-routes
vrf Corp ip export all-routes
vrf Guest ip export all-routes
vrf Voice ip export all-routes
vrf Util ip export all-routes
vrf Shared ip export route-map Shared-into-SPB
```

In some situations, sharing all routes may not be needed. For such cases, route-maps can be defined and specified with the import and export commands. Route-map “Shared-into-SPB” is used to filer routes exported from the Shared VRF. This is further explained in the “SPB Interworking With External Networks” section of this document.

VRF-I-SID binding

The L3VPN (ISIS-SPB) approach requires configuring a VRF-I-SID binding that will exchange routes between the VRF and SPB service. The IP interface that will be used as the gateway also needs to be defined here.

Use the following commands to bind the VRFs to their respective I-SIDs on all BEBs:

60. Example on BEB-3 · CLI

```
spb ipvpn bind vrf Mgt isid 1000 gateway 10.10.0.3 all-routes
spb ipvpn bind vrf Corp isid 1001 gateway 10.10.1.3 all-routes
spb ipvpn bind vrf Guest isid 1002 gateway 10.10.2.3 all-routes
spb ipvpn bind vrf Voice isid 1003 gateway 10.10.3.3 all-routes
spb ipvpn bind vrf Util isid 1004 gateway 10.10.4.3 all-routes
spb ipvpn bind vrf Shared isid 1005 gateway 10.10.5.3 all-routes
```

To check the binding status, use the following command:

61. Example on BEB-3 · CLI

```
BEB-3-> show spb ipvpn bind
Legend: * indicates bind entry is active
SPB IPVPN Bind Table:
       VRF               ISID      Gateway         Route-Map
---------------------+--------+--------------+----------------------
 * Corp                1001     10.10.1.3       all-routes
 * Guest               1002     10.10.2.3       all-routes
 * Mgt                 1000     10.10.0.3       all-routes
 * Shared              1005     10.10.5.3       all-routes
 * Util                1004     10.10.4.3       all-routes
 * Voice               1003     10.10.3.3       all-routes

Total Bind Entries: 6
```

Route leaking

It is common for some clients in a VRF to need access to those in other VRFs. Route leaking can be implemented to allow this. It is achieved by importing/exporting appropriate routes and redistributing them where needed. Route-maps can be used for granular control of routes to be imported/exported.

- Configuring route leaking involves the following steps:
- Configure appropriate route-maps to selectively import/export routes
- Export appropriate routes to the Global Routing Table (GRT)
- Import external routes from the respective I-SID
- Import local routes from the respective VRF

In the reference topology, Corp users need access to certain resources in the Util VRF like printers and storage devices. This means that all routes in the Corp VRF need to be imported into the Util VRF, but only printer and storage network routes need to be imported into Corp VRF.

A route-map is defined to allow Corp users to only access required networks in the Util VRF. IP access lists are further used for defining multiple IP addresses into a single entity. Use the following commands to configure IP access list and route-maps on all BEBs:

62. All BEBs · CLI

```
vrf Corp ip access-list "Util"
vrf Corp ip access-list "Util" address 10.11.96.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.11.128.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.12.96.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.12.128.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.13.128.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.13.96.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.14.96.0/20 action permit redist-control all-subnets
vrf Corp ip access-list "Util" address 10.14.128.0/20 action permit redist-control all-subnets

vrf Corp ip route-map "Util-into-Corp" sequence-number 10 action permit
```

```
vrf Corp ip route-map "Util-into-Corp" sequence-number 10 match ip-address "Util"
```

Next, the involved routes need to be exported from the VRF to the GRT. This was already done in the previous section while configuring L3 VPN. All routes from all VRFs were exported to the GRT. Verify the routes are present in the global routing table using the following command:

63. Example on BEB-3 (Output truncated to show only routes related to Corp and Util VRFs)· CLI

```
BEB-3-> show ip global-route-table
Type  Source               Destination        Gateway         Metric     Tag
------+--------------------+------------------+---------------+----------+----------
isid   1001                 10.10.1.0/24       10.10.1.4                1          0
isid   1001                 10.10.1.0/24       10.10.1.5                1          0
isid   1001                 10.10.1.0/24       10.10.1.6                1          0
isid   1001                 10.10.1.0/24       10.10.1.7                1          0
isid   1001                 10.12.16.0/20      10.10.1.4                1          0
isid   1001                 10.12.32.0/20      10.10.1.4                1          0
isid   1001                 10.13.16.0/20      10.10.1.5                1          0
isid   1001                 10.13.32.0/20      10.10.1.5                1          0
isid   1001                 10.14.16.0/20      10.10.1.6                1          0
isid   1001                 10.14.16.0/20      10.10.1.7                1          0
isid   1001                 10.14.32.0/20      10.10.1.6                1          0
isid   1001                 10.14.32.0/20      10.10.1.7                1          0
isid   1004                 10.10.4.0/24       10.10.4.4                1          0
isid   1004                 10.10.4.0/24       10.10.4.5                1          0
isid   1004                 10.10.4.0/24       10.10.4.6                1          0
isid   1004                 10.10.4.0/24       10.10.4.7                1          0
isid   1004                 10.12.96.0/20      10.10.4.4                1          0
isid   1004                 10.12.112.0/20     10.10.4.4                1          0
isid   1004                 10.12.128.0/20     10.10.4.4                1          0
isid   1004                 10.13.96.0/20      10.10.4.5                1          0
isid   1004                 10.13.112.0/20     10.10.4.5                1          0
isid   1004                 10.13.128.0/20     10.10.4.5                1          0
isid   1004                 10.14.96.0/20      10.10.4.6                1          0
isid   1004                 10.14.96.0/20      10.10.4.7                1          0
isid   1004                 10.14.112.0/20     10.10.4.6                1          0
isid   1004                 10.14.112.0/20     10.10.4.7                1          0
isid   1004                 10.14.128.0/20     10.10.4.6                1          0
isid   1004                 10.14.128.0/20     10.10.4.7                1          0
vrf    Corp                 10.10.1.0/24       10.10.1.3                1          0
vrf    Corp                 10.11.16.0/20      10.11.16.1               1          0
vrf    Corp                 10.11.32.0/20      10.11.32.1               1          0
vrf    Util                 10.10.4.0/24       10.10.4.3                1          0
vrf    Util                 10.11.96.0/20      10.11.96.1               1          0
vrf    Util                 10.11.112.0/20     10.11.112.1              1          0
vrf    Util                 10.11.128.0/20     10.11.128.1              1          0
```

Next, import external routes from respective I-SIDs into the respective VRFs using the following commands:

64. All BEBs · CLI

```
vrf Corp ip import isid 1004 route-map Util-into-Corp
vrf Util ip import isid 1001 all-routes
```

Finally, import local routes from respective VRFs using the following commands:

65. All BEBs · CLI

```
vrf Corp ip import vrf Util route-map Util-into-Corp
vrf Util ip import vrf Corp all-routes
```

Verify the routes are correctly imported by checking the forwarding database and router database of respective VRFs.

Use the following command to check the forwarding database:

66. Example on BEB-3 · CLI

```
BEB-3-> vrf Corp show ip routes
```

```
 + = Equal cost multipath routes
 Total 22 routes

  Dest Address        Gateway Addr         Age         Protocol
------------------+------------------+----------+-----------
  10.10.1.0/24         10.10.1.3           13d 1h    LOCAL
  10.11.16.0/20        10.11.16.1          13d 1h    LOCAL
  10.11.32.0/20        10.11.32.1          13d 1h    LOCAL
  10.11.96.0/20        10.11.96.1         01:44:48   IMPORT
  10.11.128.0/20       10.11.128.1        01:44:48   IMPORT
  10.12.16.0/20        10.10.1.4          22:37:40   IMPORT
  10.12.32.0/20        10.10.1.4          22:37:40   IMPORT
  10.12.96.0/20        10.10.4.4          01:54:06   IMPORT
  10.12.128.0/20       10.10.4.4          01:53:27   IMPORT
  10.13.16.0/20        10.10.1.5          02:35:16   IMPORT
  10.13.32.0/20        10.10.1.5          02:35:16   IMPORT
  10.13.96.0/20        10.10.4.5          01:53:01   IMPORT
  10.13.128.0/20       10.10.4.5          01:53:09   IMPORT
  10.14.16.0/20       +10.10.1.6            13d 0h   IMPORT
                      +10.10.1.7          03:13:26   IMPORT
  10.14.32.0/20       +10.10.1.6            13d 0h   IMPORT
                      +10.10.1.7          03:13:26   IMPORT
  10.14.96.0/20       +10.10.4.6          01:52:56   IMPORT
                      +10.10.4.7          01:52:56   IMPORT
  10.14.128.0/20      +10.10.4.6          01:52:50   IMPORT
                      +10.10.4.7          01:52:50   IMPORT
  127.0.0.1/32         127.0.0.1            13d 1h   LOCAL
```

Notice that only routes related to printers and storage devices (10.1x.96.0/20 & 10.1x.128.0/20) are imported from Util VRF and I-SID.

Use the following command to verify the routes on the routing database:

67. Example on BEB-3 · CLI

```
BEB-3-> vrf Util show ip router database
Legend: + indicates routes in-use
        b indicates BFD-enabled static route
        i indicates interface static route
        p indicates profinet static route
        r indicates recursive static route, with following address in brackets

Total IPRM IPv4 routes: 32

  Destination         Gateway                  Interface              Protocol  Metric     Tag
Misc-Info
--------------------+--------------+--------------------------------+--------+-------+------+
+  10.10.1.0/24       10.10.1.3       Corp-WAN          IMPORT         1          0
+  10.10.4.0/24       10.10.4.3       Util-WAN          LOCAL          1          0
   10.10.4.0/24       10.10.4.4       Util-WAN          IMPORT         1          0 (backup)
   10.10.4.0/24       10.10.4.5       Util-WAN          IMPORT         1          0 (backup)
   10.10.4.0/24       10.10.4.6       Util-WAN          IMPORT         1          0 (backup)
   10.10.4.0/24       10.10.4.7       Util-WAN          IMPORT         1          0 (backup)
+  10.11.16.0/20      10.11.16.1      Employee          IMPORT         1          0
+  10.11.32.0/20      10.11.32.1      Employee-Wifi     IMPORT         1          0
+  10.11.96.0/20      10.11.96.1      Storage           LOCAL          1          0
+  10.11.112.0/20     10.11.112.1     Sensors           LOCAL          1          0
+  10.11.128.0/20     10.11.128.1     Printers          LOCAL          1          0
+  10.12.16.0/20      10.10.1.4       Corp-WAN          IMPORT         1          0
+  10.12.32.0/20      10.10.1.4       Corp-WAN          IMPORT         1          0
+  10.12.96.0/20      10.10.4.4       Util-WAN          IMPORT         1          0
+  10.12.112.0/20     10.10.4.4       Util-WAN          IMPORT         1          0
+  10.12.128.0/20     10.10.4.4       Util-WAN          IMPORT         1          0
+  10.13.16.0/20      10.10.1.5       Corp-WAN          IMPORT         1          0
+  10.13.32.0/20      10.10.1.5       Corp-WAN          IMPORT         1          0
+  10.13.96.0/20      10.10.4.5       Util-WAN          IMPORT         1          0
+  10.13.112.0/20     10.10.4.5       Util-WAN          IMPORT         1          0
+  10.13.128.0/20     10.10.4.5       Util-WAN          IMPORT         1          0
+  10.14.16.0/20      10.10.1.6       Corp-WAN          IMPORT         1          0
+  10.14.16.0/20      10.10.1.7       Corp-WAN          IMPORT         1          0
+  10.14.32.0/20      10.10.1.6       Corp-WAN          IMPORT         1          0
```

```
+  10.14.32.0/20      10.10.1.7       Corp-WAN          IMPORT         1          0
+  10.14.96.0/20      10.10.4.6       Util-WAN          IMPORT         1          0
+  10.14.96.0/20      10.10.4.7       Util-WAN          IMPORT         1          0
+  10.14.112.0/20     10.10.4.6       Util-WAN          IMPORT         1          0
+  10.14.112.0/20     10.10.4.7       Util-WAN          IMPORT         1          0
+  10.14.128.0/20     10.10.4.6       Util-WAN          IMPORT         1          0
+  10.14.128.0/20     10.10.4.7       Util-WAN          IMPORT         1          0
+  127.0.0.1/32       127.0.0.1       Loopback          LOCAL          1          0


Inactive Static Routes
   Destination       Gateway           Metric        Tag    Misc-Info
--------------------+-----------------+------+----------+-----------------
```

In this output observe that all routes from the Corp VRF and I-SID are imported into the Util routing database.

Note that although routes are present, pinging the interfaces from the switch is blocked for security reasons. However, clients in the respective subnets should be able to ping the interfaces of imported networks. This is demonstrated in the next section.

Validation

To validate, the following test IP interfaces are created on ACC-31 and ACC-41:

Switch	IP address	Gateway	VLAN	User
ACC-31	10.11.16.10/20	10.11.16.1	1102	BLDG-A Employee
ACC-41	10.12.128.10/20	10.12.128.1	1209	BLDG-B Printers

Table 8: Test IP interfaces for Layer 3 validation

In addition to creating IP interfaces, static routes with respective gateways need to be configured on the access switches. Use the following commands to create the IP interface and static route:

68. Example on ACC-31 · CLI

```
ip interface "test" address 10.11.16.10 mask 255.255.240.0 vlan 1102
ip static-route 10.0.0.0/8 gateway 10.11.16.1
```

Ping from ACC-31 to the ACC-41 should now work.

69. Example on ACC-31 (output truncated) · CLI

```
ACC-31-> ping 10.12.128.10
PING 10.12.128.10 (10.12.128.10) 56(84) bytes of data.
64 bytes from 10.12.128.10: icmp seq=1 ttl=62 time=15.2 ms
.
--- 10.12.128.10 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5007ms
rtt min/avg/max/mdev = 0.649/3.582/15.270/5.308 ms
```

Ping from ACC-31 to gateway interface of BLDG-B printer group would work but not to the gateway interface of BLDG-A printer group.

70. Example on BEB-3 (Output Truncated) · CLI

```
ACC-31-> ping 10.12.128.1
PING 10.12.128.1 (10.12.128.1) 56(84) bytes of data.
64 bytes from 10.12.128.1: icmp seq=1 ttl=63 time=3.46 ms
.
--- 10.12.128.1 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5008ms
rtt min/avg/max/mdev = 1.575/2.520/3.511/0.772 ms
ACC-31->
ACC-31-> ping 10.11.128.1
PING 10.11.128.1 (10.11.128.1) 56(84) bytes of data.
--- 10.11.128.1 ping statistics ---
```

```
6 packets transmitted, 0 received, 100% packet loss, time 5101ms
```

Ping from BEB-3 to both gateway interfaces and created client interfaces is blocked. This can be seen in the following output:

71. Example on BEB-3 · CLI

```
BEB-3-> vrf Corp ping 10.12.128.1
Please wait...
connect: Network is unreachable
BEB-3-> vrf Corp ping 10.12.128.10
Please wait...
connect: Network is unreachable
BEB-3-> vrf Corp ping 10.11.128.1
Please wait...
connect: Network is unreachable
```

After validation, remove the “test” IP interface and the default route using the following commands:

72. Example on ACC-31 · CLI

```
no ip interface test
no ip static-route 10.0.0.0/8 gateway 10.11.16.1
```

Multicast routing

Multicast switching was introduced in the previous chapter. In many deployments, multicast sources and receivers may reside in different Layer 2 network domains. In these scenarios, multicast routing is required to forward multicast traffic between the appropriate routed interfaces and deliver it to the intended receivers.

On AOS, multicast routing can be implemented using multicast routing protocols such as DVMRP or PIM. For additional information about multicast routing protocols and configuration options, refer to the “AOS Advanced Routing User Guide”.

In the reference topology, security cameras generate multicast traffic in the network. Service 2001 is dedicated to the camera multicast streams, and its Layer 2 configuration was described in the previous chapter. Service 2003 is used by the security systems that require access to those multicast streams. To support multicast forwarding between these services, PIM Sparse Mode (PIM-SM) is used.

PIM must be configured on the routers associated with the two services. Both services belong to the Landlord VRF and are routed on BEB-3 and BEB-4, which operate as a VRRP pair.

To configure PIM, the corresponding module needs to be manually loaded to memory. Then appropriate IP interfaces need to be enabled with PIM. Use the following commands to load and configure PIM interfaces:

73. BEB-3 & BEB-4 · CLI

```
vrf Landlord ip load pim
vrf Landlord ip pim interface "Camera"
vrf Landlord ip pim interface "Security"
```

Next, configure the Candidate Rendezvous Point (C-RP) and Candidate Bootstrap Router (C-BSR) functions on the appropriate devices in the network. In this reference topology, BEB-3 and BEB-4 are configured as both C-RP and C-BSR candidates.

Use the following commands to configure C-RP and C-BSR on BEB-3 and BEB-4:

74. Example on BEB-3 · CLI

```
vrf Landlord ip pim cbsr 10.20.32.2
vrf Landlord ip pim candidate-rp 10.20.32.2 225.0.0.0/24
```

Finally, enable PIM Sparse Mode on the corresponding routed interfaces:

75. BEB-3 & BEB-4 · CLI

```
vrf Landlord ip pim sparse admin-state enable
```

To verify the PIM configuration, use the following command:

76. Example on BEB-3 · CLI

```
Landlord::BEB-3-> show ip pim interface
Total 2 Interfaces
Interface Name                   IP Address      Designated      Hello    J/P      Oper
                                                  Router          Interval Interval Status
---------------------------------+---------------+---------------+--------+--------+-------
-
Camera                           10.20.32.2      10.20.32.3      30       60       enabled
Security                         10.20.96.2      10.20.96.3      30       60       enabled
```

In the command output, verify that PIM is enabled on all required devices. Also confirm that the Designated Router (DR) and Bootstrap Router (BSR) elections are consistent across the multicast routing domain.

SPB interworking with external networks

SPB fabrics rarely operate in isolation—most deployments must interconnect with existing Layer 2 and Layer 3 domains, upstream WAN/Internet connectivity, and security or service appliances. This chapter describes how the reference SPB fabric interworks with external networks focusing on STP and OSPF.

External STP network

SPB uses IS-IS and loop-prevention mechanisms (including RPFC) within the fabric to maintain loop-free forwarding. However, Layer 2 loops can still occur in external Ethernet domains connected at the edge (for example, in access/aggregation networks). It is therefore important to ensure that any Layer 2 domain attached to a BEB is protected with a well-defined and predictable STP design, so loops do not propagate toward the SPB edge.

Spanning Tree Protocol (STP) is commonly used to prevent Layer 2 loops. While STP is largely self-forming, production deployments should explicitly control key parameters—especially root bridge selection and port preference—to achieve stable and deterministic behavior. (For additional platform-specific detail, refer to the “AOS Network Configuration User Guide”.)

A critical design decision is the selection of the STP root bridge for VLANs that extend between access switches and a BEB. When VLANs are stretched up to a BEB, placing the root at the BEB typically provides a clean boundary and predictable forwarding toward the fabric edge.

In the reference topology, VLANs in BLDG-C are extended to BEB-5. Therefore, BEB-5 is configured as the preferred STP root by lowering its bridge priority (default is 32768):

77. BEB-5 · CLI

```
spantree vlan 1301-1309 priority 20480
```

Similarly, bridge priorities on ACC-31, ACC-41, and ACC-61 are adjusted so the intended device becomes root for their respective Layer 2 domains/VLANs.

In some designs, STP path selection may not align with the intended forwarding preference—particularly when uplinks terminate on an SPB edge and traffic is transported through the fabric. For example, an access switch with two uplinks may block a higher-speed link and forward on a lower-speed link if STP costs or priorities result in that outcome.

To enforce the preferred forwarding path, adjust STP port priority (or path cost, depending on your operational standard). In the reference topology, port 1/1/25 is preferred due to higher bandwidth:

78. ACC-61 · CLI

```
spantree vlan 1401-1409 port 1/1/24 priority 8
```

To protect the network from miswiring and rogue or unexpected BPDUs, consider enabling STP protection features such as BPDU Guard, Root Guard, and Loop Guard where appropriate. Some of these controls are discussed further in the “Security and Features” section of this guide.

External OSPF router

SPB fabrics commonly need to interconnect with non-SPB Layer 3 networks such as WAN/Internet edge, data center cores, or upstream campus routing domains. In this design, interworking is achieved using a VPN-lite approach: routing separation is maintained with VRFs, and routes are exchanged with external routers using a standard routing protocol such as OSPF.

In the reference topology, the border BEBs (BEB-3 and BEB-4) peer with external routers (RTR-1 and RTR-2) using OSPF. Configuring OSPF interworking using VPN-lite typically involves the following steps:

- Create Shared VRF and interfaces for uplinks

- Enable and configure OSPF
- Exchange and redistribute routes

Shared VRF and IP interfaces

In the reference topology, a dedicated Shared VRF is created on BEB-3 and BEB-4 and is used exclusively for northbound peering toward RTR-1 and RTR-2. Limiting this VRF to border nodes keeps the external routing boundary well-defined and reduces unnecessary control-plane exposure within the fabric. The OSPF routes exchanged by the router are tagged with a tag value 10. This further helps in controlling route redistribution as will be explained in the next section.

Create the Shared VRF using the following command (skip if already configured):

79. BEB-3 & BEB-4 · CLI

```
vrf create Shared
```

The BEB-router connections use point-to-point links with the addressing shown below:

	VLAN	BEB IP address (Shared VRF)	RTR IP address
BEB-3 to RTR-1	201	10.20.0.2/31	10.20.0.3/31
BEB-4 to RTR-2	203	10.20.0.8/31	10.20.0.9/31

Table 9: IP address for point-to-point links

Create the corresponding IP interfaces using:

80. Example on BEB-3 · CLI

```
vrf Shared ip interface "To-RTR1" address 10.20.0.2/31 vlan 201 rtr-port port 1/1/9 tagged
```

The rtr-port option is used when creating these interfaces. This disables Spanning Tree for the VLAN and prevents the VLAN from being used on other ports, which is appropriate for routed point-to-point uplinks and helps avoid unintended Layer 2 extension.

OSPF

In this design, both the border BEBs and the external routers participate in OSPF Area 0. OSPF configuration follows a standard sequence:

- Loading OSPF
- Creating OSPF area
- Defining OSPF interfaces
- Associating interfaces with the area
- Enabling OSPF

Use the following commands to configure OSPF on the border BEBs (and the corresponding OSPF configuration on the routers, as applicable):

81. Example on BEB-3 · CLI

```
vrf Shared ip load ospf
vrf Shared ip ospf area 0.0.0.0
vrf Shared ip ospf interface "To-RTR1"
vrf Shared ip ospf interface "To-RTR1" area 0.0.0.0
vrf Shared ip ospf interface "To-RTR1" type point-to-point
vrf Shared ip ospf interface "To-RTR1" admin-state enable
vrf Shared ip ospf admin-state enable
```

The OSPF interfaces are configured as point-to-point. This removes the need for DR/BDR elections, reduces protocol overhead, and typically improves convergence on routed uplinks.

Route exchange and redistribution

With OSPF adjacencies established in the Shared VRF, the next step is to control how routes are exchanged between:

- the external OSPF domain (RTR-1 / RTR-2),
- the Shared VRF on the border BEBs, and
- the internal User VRFs carried across the SPB fabric.

The reference design uses a policy-driven approach to keep routing predictable and secure: only the required prefixes are imported/exported, and redistribution is performed explicitly.

Route exchange in this guide is implemented with the following sequence:

- Export selected OSPF-learned routes from the Shared VRF
- Import Shared VRF routes into the required User VRFs (route leaking)
- Redistribute Shared VRF routes so remote BEBs learn them in their User VRF instances
- Import User VRF routes (local and remote) into the Shared VRF
- Redistribute imported User VRF routes into OSPF toward RTR-1 / RTR-2

Once the border BEBs participate in OSPF, they learn multiple prefixes from the external domain. Typically, not all external routes are required inside the SPB fabric. Use a route-map to export only approved prefixes from the Shared VRF:

82. BEB-3 and BEB-4 · CLI

```
vrf Shared ip access-list "Shared"
vrf Shared ip access-list "Shared" address 0.0.0.0/0 action permit redist-control no-
subnets
vrf Shared ip route-map "Shared-into-SPB" sequence-number 10 action permit
vrf Shared ip route-map "Shared-into-SPB" sequence-number 10 match ip-address "Shared"
vrf Shared ip export route-map Shared-into-SPB
```

After export, leak the required Shared VRF routes into the appropriate User VRFs (for example, Corp/Guest/Voice/Util) using import policy:

83. BEB-3 and BEB-4 · CLI

```
vrf Corp ip import vrf Shared all-routes
vrf Guest ip import vrf Shared all-routes
vrf Voice ip import vrf Shared all-routes
vrf Util ip import vrf Shared all-routes
```

The Shared VRF routes should now be visible in the relevant User VRFs. Use the following command to verify the imported routes:

84. Example on BEB-3 · CLI

```
BEB-3-> vrf Guest show ip routes

 + = Equal cost multipath routes
 Total 14 routes

  Dest Address       Gateway Addr        Age        Protocol
-----------------+------------------+----------+-----------
  0.0.0.0/0          +10.10.2.4        00:00:05   IMPORT
                     +10.20.0.3        00:07:09   IMPORT
  10.10.2.0/24        10.10.2.3        00:12:58   LOCAL
  10.11.48.0/20       10.11.48.1       00:13:07   LOCAL
  10.11.64.0/20       10.11.64.1       00:13:07   LOCAL
  10.12.48.0/20       10.10.2.4        00:00:05   IMPORT
  10.12.64.0/20       10.10.2.4        00:00:05   IMPORT
  10.13.48.0/20       10.10.2.5        00:00:05   IMPORT
  10.13.64.0/20       10.10.2.5        00:00:05   IMPORT
  10.14.48.0/20      +10.10.2.6        00:00:05   IMPORT
```

```
                      +10.10.2.7          00:00:05   IMPORT
  10.14.64.0/20       +10.10.2.6          00:00:05   IMPORT
                      +10.10.2.7          00:00:05   IMPORT
  127.0.0.1/32         127.0.0.1          00:15:45   LOCAL
```

In the forwarding table, note that the imported Shared route may appear with two next-hop paths: one through the OSPF interface on BEB-3 and another through BEB-4. This behavior is undesirable because it can create a routing loop, where traffic is forwarded back and forth between BEB-3 and BEB-4.

To prevent this condition, a route-map similar to the following can be used on all applicable VRFs:

85. Example on BEB-3 for Guest VRF · CLI

```
vrf Guest ip route-map "Default" sequence-number 50 action permit
vrf Guest ip route-map "Default" sequence-number 50 match tag 10
vrf Guest ip route-map "Default" sequence-number 50 set metric 10 effect replace
vrf Guest ip route-map "Default" sequence-number 60 action permit
vrf Guest ip route-map "Default" sequence-number 60 match ip-address 0.0.0.0/0 redist-
control all-subnets permit
vrf Guest ip import isid 1002 route-map Default
```

This route-map sets the metric value to 10 for all routes that match tag value 10. As described earlier, all OSPF routes exchanged with the external router are assigned tag value 10.

By increasing the metric of the imported route, the switch prefers the lower-metric route learned directly through its local OSPF interface. If the local OSPF link fails, the alternate path through the peer BEB remains available and is used as a backup route.

Verify the forwarding table again to confirm that the preferred route is selected as expected:

86. Example on BEB-3 · CLI

```
BEB-3-> vrf Guest show ip routes

 + = Equal cost multipath routes
 Total 13 routes

  Dest Address       Gateway Addr        Age        Protocol
-----------------+------------------+----------+-----------
  0.0.0.0/0           10.20.0.3         00:25:28   IMPORT
  10.10.2.0/24        10.10.2.3         00:31:17   LOCAL
  10.11.48.0/20       10.11.48.1        00:31:26   LOCAL
  10.11.64.0/20       10.11.64.1        00:31:26   LOCAL
  10.12.48.0/20       10.10.2.4         00:09:20   IMPORT
  10.12.64.0/20       10.10.2.4         00:09:20   IMPORT
  10.13.48.0/20       10.10.2.5         00:09:20   IMPORT
  10.13.64.0/20       10.10.2.5         00:09:20   IMPORT
  10.14.48.0/20      +10.10.2.6         00:09:20   IMPORT
                     +10.10.2.7         00:09:20   IMPORT
  10.14.64.0/20      +10.10.2.6         00:09:20   IMPORT
                     +10.10.2.7         00:09:20   IMPORT
  127.0.0.1/32        127.0.0.1         00:34:04   LOCAL
```

To make these imported Shared VRF routes available on non-border BEBs, redistribute them so they propagate across the fabric to remote instances of the User VRFs:

87. Example on BEB-3 · CLI

```
spb ipvpn redist source-vrf Shared destination-isid 1001 all-routes
spb ipvpn redist source-vrf Shared destination-isid 1002 all-routes
spb ipvpn redist source-vrf Shared destination-isid 1003 all-routes
spb ipvpn redist source-vrf Shared destination-isid 1004 all-routes
```

At this point, the Shared/External routes should be visible on all BEBs within the relevant VRFs. Verify using:

88. Example on BEB-6 (Output truncated) · CLI

```
BEB-6-> vrf Corp show ip routes
```

```
 + = Equal cost multipath routes
 Total 20 routes

  Dest Address       Gateway Addr        Age         Protocol
------------------+------------------+----------+-----------
  0.0.0.0/0          +10.10.1.3         01:31:46   IMPORT
                     +10.10.1.4         01:24:02   IMPORT
  10.10.1.0/24        10.10.1.6            2d 0h   LOCAL
```

Next, import User VRF routes into the Shared VRF so they can be advertised northbound into OSPF. In this guide, this is accomplished by importing routes learned via the relevant I-SID context (as applicable to the design):

89. Example on BEB-3 · CLI

```
vrf Shared ip import isid 1001 all-routes
vrf Shared ip import isid 1002 all-routes
vrf Shared ip import isid 1003 all-routes
vrf Shared ip import isid 1004 all-routes
```

After the User VRF prefixes are present in the Shared VRF, redistribute them into OSPF toward the external routers:

90. Example on BEB-3 · CLI

```
vrf Shared ip redist import into ospf all-routes admin-state enable
```

Verify that the expected routes are present on RTR-1/RTR-2:

91. Example on RTR-1 (Output truncated) · CLI

```
RTR-1-> show ip routes

 + = Equal cost multipath routes
 Total 44 routes

  Dest Address       Gateway Addr        Age         Protocol
------------------+------------------+----------+-----------
  0.0.0.0/0           10.20.0.11           1d 0h   OSPF
  10.10.1.0/24        10.20.0.2         01:35:47   OSPF
  10.10.2.0/24        10.20.0.2         01:43:39   OSPF
  10.10.3.0/24        10.20.0.2         01:43:39   OSPF
  10.10.4.0/24        10.20.0.2         01:43:39   OSPF
  10.11.16.0/20       10.20.0.2         01:43:39   OSPF
  10.11.32.0/20       10.20.0.2         01:43:39   OSPF
  10.11.48.0/20       10.20.0.2         01:43:39   OSPF
  10.11.64.0/20       10.20.0.2         01:43:39   OSPF
  10.11.80.0/20       10.20.0.2         01:43:39   OSPF
  10.11.96.0/20       10.20.0.2         01:43:39   OSPF
  10.11.112.0/20      10.20.0.2         01:43:39   OSPF
  10.11.128.0/20      10.20.0.2         01:43:39   OSPF
```

TIP Route-maps can be applied at multiple points (export, import, redistribution) to tightly control which prefixes are exchanged and to prevent accidental route leaks.

NOTE Even when routes are present, reachability tests (e.g., ping/traceroute) may fail if inter-VRF security policy, ACLs, or firewall policy blocks ICMP or the traffic type. Route visibility confirms control-plane learning; end-to-end testing should also validate the data-plane policy.

DHCP relay

DHCP relay enables DHCP clients and servers to communicate across different VLANs/subnets by converting client broadcast messages into unicast messages toward a configured DHCP server. This avoids the need to deploy a DHCP server in every VLAN while still providing centralized address management.

NOTE : In an inter-VRF DHCP relay deployment, DHCP requests originating from a local client VRF on the same switch are not relayed into another local VRF.

In the reference topology, BEB-3 and BEB-4 provide client connectivity for BLDG-A and BLDG-B respectively. If the Shared VRF with DHCP server reachability is configured directly on these same BEBs, DHCP requests from the locally connected client VRFs are not forwarded into the Shared VRF. As a result, clients in BLDG-A and BLDG-B may not receive DHCP addresses with this relay placement.

The recommended workaround is to deploy dedicated BEB switches, such as BEB-8 and BEB-9 shown in Figure 20, where no DHCP clients are connected. Configure the shared VRF on these BEBs and leak the required DHCP server routes into the user VRFs. The leaked routes can then be redistributed within the network, allowing DHCP relay traffic from the client-facing BEBs to reach the DHCP server through the external forwarding path.

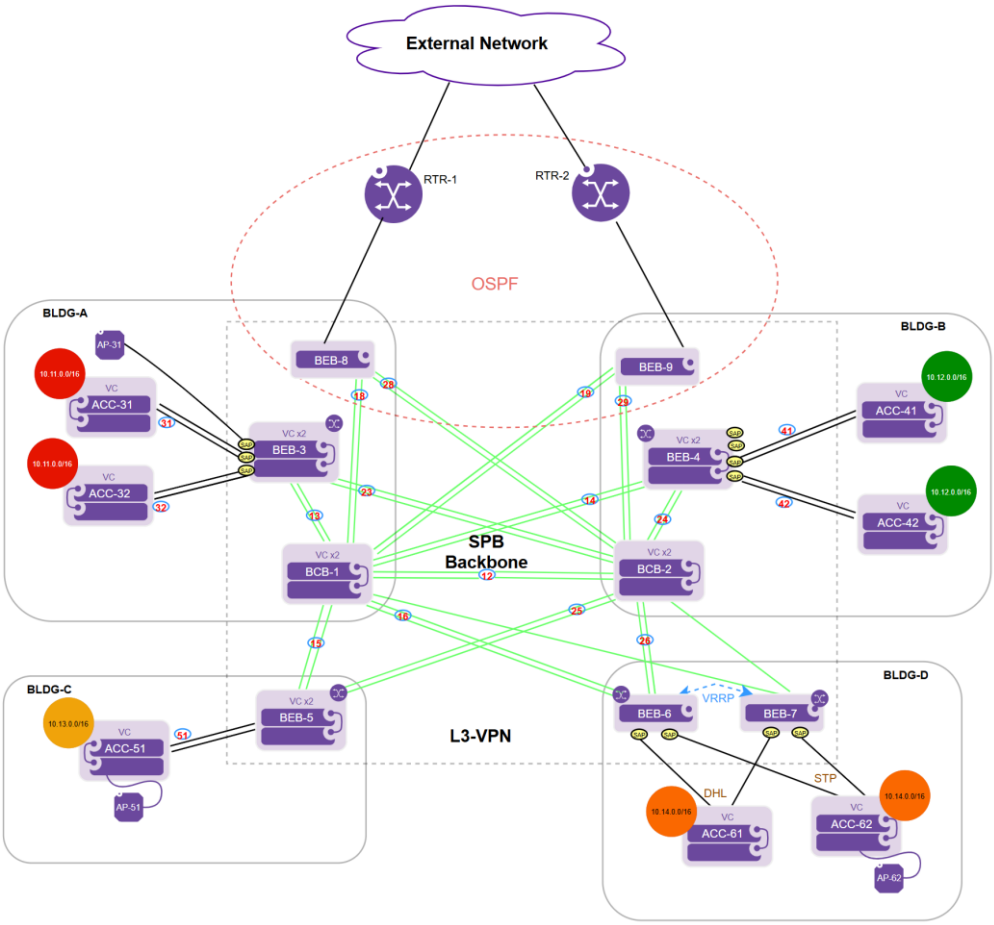

Figure 20. BEB-8 & BEB-9 with no DHCP clients

In the reference topology, a single DHCP server at 172.20.0.50 provides addressing for multiple user VLANs. Reachability to this server is available via the Shared VRF, and DHCP relay is enabled on the default gateway interface for each client VLAN so requests can be forwarded to the server.

Configure DHCP relay on the relevant gateway interfaces on all BEBs using the following commands:

92. All BEBs · CLI

```
vrf Corp ip dhcp relay admin-state enable
vrf Corp ip dhcp relay per-interface-mode
vrf Corp ip dhcp relay interface "Employee" destination 172.20.0.50
vrf Corp ip dhcp relay interface "Employee" admin-state enable
vrf Corp ip dhcp relay interface "Employee-Wifi" destination 172.20.0.50
vrf Corp ip dhcp relay interface "Employee-Wifi" admin-state enable
vrf Guest ip dhcp relay admin-state enable
vrf Guest ip dhcp relay per-interface-mode
vrf Guest ip dhcp relay interface "Guest-Wifi" destination 172.20.0.50
vrf Guest ip dhcp relay interface "Guest-Wifi" admin-state enable
vrf Guest ip dhcp relay interface "Contractor" destination 172.20.0.50
vrf Guest ip dhcp relay interface "Contractor" admin-state enable
```

```
vrf Voice ip dhcp relay admin-state enable
vrf Voice ip dhcp relay per-interface-mode
vrf Voice ip dhcp relay interface "Voice" destination 172.20.0.50
vrf Voice ip dhcp relay interface "Voice" admin-state enable
vrf Util ip dhcp relay admin-state enable
vrf Util ip dhcp relay per-interface-mode
vrf Util ip dhcp relay interface "Storage" destination 172.20.0.50
vrf Util ip dhcp relay interface "Storage" admin-state enable
vrf Util ip dhcp relay interface "Sensors" destination 172.20.0.50
vrf Util ip dhcp relay interface "Sensors" admin-state enable
vrf Util ip dhcp relay interface "Printers" destination 172.20.0.50
vrf Util ip dhcp relay interface "Printers" admin-state enable
```

To validate DHCP relay, create a DHCP client interface on an access switch (ACC) in the target VLAN and confirm that it receives a lease.

Create a DHCP client interface:

93. Example on ACC-51 · CLI

```
ip interface dhcp-client vlan 1306
```

Verify that the interface receives an IP address from the DHCP server:

94. Example on ACC-51 · CLI

```
ACC-51-> show ip interface
Total 4 interfaces
 Flags (D=Directly-bound)
       (A=Anycast IP)

       Name      IP Address      Subnet Mask     Status Forward  Device             Flags
--------------------------------+---------------+---------------+------+-------+----------
-EMP-CMMA-CHAS1   10.255.219.223  255.255.252.0   UP     NO      EMP
Loopback         127.0.0.1       255.255.255.255 UP     NO      Loopback
dhcp-client      10.13.80.151    255.255.255.0   UP     YES     vlan 1306
```

Redundancy options

Virtual Chassis

Virtual Chassis is an AOS feature which allows a group of switches to be connected together to function as a single switch. It is a simple and effective way to provide redundancy to the access layer. Access layer switches connect to the VC using a LAG where each link is connected to a different chassis of the VC, providing both link and switch level redundancy.

In the reference topology, all access switches connected to BEB-3, BEB-4 and BEB-5 are configured as a VC with LAG.

TIP A VC can be automatically set up using Auto-VC feature of AOS. Refer to 'AOS Switch Management User Guide' for further details.

Link aggregates using LACP

Link aggregation provides an active-active redundancy by grouping multiple links to behave as a single logical link. However, this by itself only provides link level redundancy. This would fail if multiple links of the LAG are connected to a switch, and a failure arises at the switch level.

It is recommended to use link aggregates on all links connecting to critical devices. This also provides the added benefit of extra bandwidth between the two connected devices.

In the reference topology LAGs are used for all links in the SPB backbone expect links connecting BEB-7. All access switches except ACC-61 and ACC-62 also connect to respective BEBs using LAGs.

By default, non-unicast traffic is forwarded only through the primary port of the LAG. It is recommend to enable hashing for non-unicast traffic on all switches. It can be done using the following command:

95. Example on BCB-2 · CLI

```
hash-control load-balance non-ucast enable
hash-control load-balance non-ucast tunnel-protocol enable
```

In addition, the default hashing logic might be inefficient when using tunneling protocols like SPB. To improve the hashing logic, AOS allows users to configure the hashing algorithm to use SPB payload data. This needs to be enabled globally and on the link aggregate level. Use the following command to modify the hashing algorithm:

96. Example on BCB-2 · CLI

```
hash-control extended spb-payload
linkagg lacp agg 12 size 2 hash tunnel-protocol admin-state enable
```

NOTE The availability of these features varies across different switch models. Refer to the 'AOS CLI Reference User Guide' for more information.

Dual Home Link (DHL)

DHL is an AOS feature which allows active-active redundancy for multihomed switches. VLANs are grouped and mapped to either of the two uplinks. DHL allows for load balancing and quick failover when links fail.

In the reference topology, DHL is used on ACC-61 to load balance the VLANs across the two uplinks to BEB-6 & BEB-7. Port 1/1/25 offers 10G link and hence is preferred for essential VLANs.

Configuring DHL involves creating a DHL session, defining the uplinks and mapping VLANs to them. Use the following commands to configure DHL:

97. ACC-61 · CLI

```
dhl 1
dhl 1 linka port 1/1/24 linkb port 1/1/25
dhl 1 admin-state enable
dhl 1 vlan-map linkb 1401-1403
```

```
dhl 1 vlan-map linkb 1406-1409
```

Configuring a MAC address flush method (MVRP or Raw flooding) is recommended if the DHL ports span across switch modules. This improves convergence times.

DHL status can be verfied by checking the VLAN members for the involved ports.

98. ACC-61 · CLI

```
ACC-61-> show vlan members port 1/1/25
  vlan      type        status
--------+----------+--------------
     1    untagged    dhl-blocking
  1401    tagged      forwarding
  1402    tagged      forwarding
  1403    tagged      forwarding
  1404    tagged      dhl-blocking
  1405    tagged      dhl-blocking
  1406    tagged      forwarding
  1407    tagged      forwarding
  1408    tagged      forwarding
  1409    tagged      forwarding
```

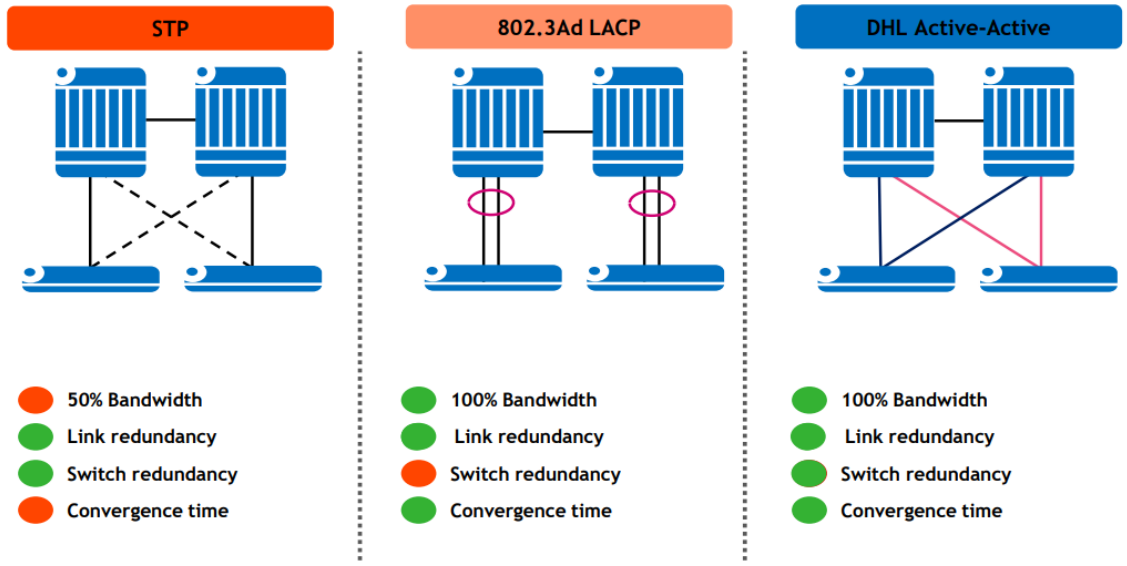

Figure 21. Comparing STP, LACP and DHL

Features & security

Loopback detection

Loops in SPB network are prevented by IS-IS protocol and Reverse-Path Forwarding Check(RPFC). However, loops may be created due to misconfiguration at Access layer. By default, SAPs forward STP Bridge Protocol Data Units (BPDUs) allowing redundantly-attached VLAN-domain access layer to use STP for loop prevention. To protect the SPB backbone from broadcast storms, loops involving SAPs must be detected and broken.

Loopback Detection (LBD) is an AOS feature designed to identify and mitigate network loops. It operates in addition to other mechanisms like DHL and STP. The switch periodically sends out LBD frames from LBD-enabled ports and concludes that the port is looped back if it receives the frame on any of the LBD-enabled ports. In which case the port is disabled and goes into a shutdown state. A trap is sent and the event is logged.

Loopback-Detection must first be enabled globally and then applied to all UNI ports (VLAN and SAP).

This configuration is implemented between BEBs and ACC switches SAPs.

Use the following commands to enable Loopback Detection globally and on the SAPs connected to access switches:

99. Example on BEB-3 · CLI

```
loopback-detection enable
loopback-detection service-access linkagg 31-32 enable
```

Learned Port Security

Learned Port Security (LPS) provides a mechanism to control which hosts are allowed to connect to the network through a physical port or SAP. LPS allows the administrator to define the maximum number of MAC addresses that can be learned on an interface. When the configured threshold is exceeded, the port or SAP can be placed into a restricted or shutdown state, depending on the selected violation action.

LPS can also control how unauthorized traffic is handled when a violation occurs. The following actions are supported:

- Block traffic that violates the LPS restrictions while continuing to forward authorized traffic.
- Disable MAC learning on the LPS-enabled port when unauthorized traffic is detected.
- Administratively shut down the LPS-enabled port, which stops all traffic on the interface.

In addition, LPS supports a learning window during which MAC addresses can be learned dynamically and then converted into static entries. These entries can be saved to the configuration file for future use. This allows the administrator to explicitly define which MAC addresses are permitted on a port or SAP. LPS also provides logging and notification capabilities when a violation occurs.

LPS configuration differs depending on whether it is applied to a static SAP or a dynamic SAP.

For a static SAP, the SAP and its mapping to the SPB service must already be configured before LPS is applied. The following example shows an LPS configuration for a static SAP:

100. Sample code · CLI

```
port-security sap 1/1/1:10 maximum 1000
```

For a dynamic SAP, the dynamic SAP and its service mapping must already exist on the UNP port before LPS can be applied. This can be achieved by configuring a persistent profile on the UNP access port or link aggregation.

Assigning a UNP service profile statically to a UNP port creates a persistent SAP. Unlike a dynamically created SAP, a persistent SAP does not age out when there is no activity on the port. This behavior is useful for supporting silent devices in the UNP service domain, where traffic may not be generated immediately after the device is connected.

DHCP snooping

DHCP Snooping enhances network security by inspecting DHCP messages received on access-facing interfaces and filtering unauthorized or unexpected DHCP traffic. It also builds and maintains a DHCP binding table, which records client access information such as MAC address, IP address, VLAN or service association, and interface details.

DHCP Snooping can be enabled either globally at the switch level or selectively at the VLAN or service level. The feature must be enabled on all user-facing UNI interfaces, including VLAN ports and SAPs, where DHCP client traffic is expected.

Use the following commands to enable DHCP Snooping globally:

101. Example on BEB-3 · CLI

```
dhcp-snooping admin-state enable
```

Alternatively, DHCP Snooping can be enabled for a specific VLAN or service. Use the following commands to enable DHCP Snooping at the VLAN or service level:

102. Sample code · CLI

```
dhcp-snooping vlan 1403 admin-state enable
dhcp-snooping service 1403 admin-state enable
```

When DHCP Snooping is enabled for a service, all SAP ports allow only DHCP client packets by default, while SDP ports allow all DHCP packets. Port-level DHCP Snooping properties cannot be configured for service ports.

When DHCP Snooping is enabled for a VLAN, all ports allow only DHCP client packets by default. Therefore, uplink ports toward the DHCP server must be configured as trusted so that DHCP server packets are allowed.

Use the following command to configure an uplink port as trusted and permit all DHCP traffic:

103. Example on ACC-61 · CLI

```
dhcp-snooping port 1/1/24-25 trust
```

Testing & validation results

ATTENTION The figures presented in this section reflect the specific scale, topology, platforms, and software release exercised during this validation exercise. They describe what was tested in this particular environment and are provided for illustration only. They must not be interpreted as scalability maximums, guaranteed convergence times, performance limits, or product specifications. Behaviour in other deployments will vary with topology, hardware variant, software release, traffic profile, and configuration. For supported maximums and committed specifications, always refer to the applicable ALCATEL-LUCENT OmniSwitch platform data sheet and release notes.

The reference topology presented in this document was thoroughly tested and validated in the lab. The validation covered multiple areas, including convergence behavior, scalability, traffic forwarding, and feature reliability.

This chapter summarizes the test environment, traffic model, scale parameters, and validation results used to qualify the proposed design.

Traffic configuration

The tests were performed by simulating user traffic in the network using an Ixia traffic generator. To validate the scalability of the proposed architecture, the lab topology was extended beyond the baseline design.

The additional scale configuration used for testing is summarized below:

- In addition to the baseline Mgt, Corp, Guest, Voice, and Util VRFs, 15 additional VRFs were configured on all BEBs.
- Each of the additional VRFs was configured with six services, resulting in a total of 99 user services per building.
- The service-to-I-SID-to-VLAN mapping was maintained consistently, resulting in 99 user VLANs per building.
- Layer 3 services were extended to each VRF using L3VPN, enabling routed connectivity between users belonging to the same VRF across different buildings.
- In addition to the existing ACC-## access switches, one additional access switch was simulated using Ixia. This required the configuration of additional SAPs on the BEBs.
- Ten end clients were configured per service on each access switch, resulting in a total of 10,800 simulated clients.
- Continuous bidirectional traffic was generated between users. The traffic profile included both Layer 2 traffic between users within the same building and Layer 3 traffic between users in the same VRF across different buildings.
- The total number of configured traffic streams was approximately 750,000.

The following table summarizes the network configuration:

	BLDG-A (BEB-3)	BLDG-B (BEB-4)	BLDG-C (BEB-5)	BLDG-D (BEB-6 & BEB-7)
BVLANs	3 BVLANs – 1 Control BVLAN (4000) + 2 User BVLANs (4001-4002)			
VRFs	23 VRFs		21 VRFs (No Shared and Landlord VRFs)	
Services:I-SID:VLANs	124 1000-1020 1101-1199 2000-2003	124 1000-1020 1201-1299 2000-2003	120 1000-1004 1006-1020 1301-1399	123 1000-1004 1006-1020 1401-1499

Access Switches	ACC-31, ACC-32 ACC-33 (simulated)	ACC-41, ACC-42 ACC-43 (simulated)	ACC-51 ACC-52 (simulated)	ACC-61, ACC-62 ACC-63, ACC-64 (simulated)
SAPs	99*3 = 297	99*3 = 297	120 (External Loopback)	99*3 = 297 each
Number of clients	2700	2700	1800	3600

Table 10: Testing configuration summary

SPB resiliency tests

One of the key advantages of SPB is its ability to provide resilient forwarding with low latency and sub-second convergence. To validate the resiliency of the SPB backbone, convergence behavior was tested under multiple failure and recovery scenarios.

The following SPB backbone resiliency scenarios were tested:

- BCB reload
- BCB link failure
- BCB and BEB VC-takeover

For each test case, convergence was measured separately for the failure event and the recovery event. This approach provides visibility into the traffic impact when the device or link goes down, as well as when it comes back into service.

BCB reload

BCB reload testing was performed to validate traffic convergence when a backbone core bridge becomes unavailable and when it subsequently rejoins the SPB fabric.

The convergence results are summarized in the following table:

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Reload - Failure	BCB-1	114	254.63	213.53	27.50	682.50
Reload - Recovery	BCB-1	112	131.98	57.25	1.10	805.00
Reload - Failure	BCB-2	126	328.78	275.50	16.60	1,075.80
Reload - Recovery	BCB-2	72	99.15	63.47	1.90	331.35

Table 11: Convergence during BCB reload

In almost all test cases, traffic converged within one second. During the failure event, affected traffic was rerouted through the alternate BCB, confirming that the SPB backbone provided the expected path redundancy.

BCB link failure

Link failure testing was performed by toggling the SPB backbone links between the BCBs and BEBs. Convergence times were measured independently for both link-down and link-up events.

The results are summarized in the following table:

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Linkagg-12- Failure & Recovery	BCB-1	0				
Linkagg-13 - Failure	BCB-1	54	136.303	141.25	31.3	261.45
Linkagg-13 - Recovery	BCB-1	39	72.2799	60.05	9.567	187.9
Linkagg-14 - Failure	BCB-1	54	230.9169	221.983	92.467	380.7
Linkagg-14 - Recovery	BCB-1	39	151.7026	124.367	53.367	314.9
Linkagg-15 - Failure	BCB-1	27	96.7302	54.4	4	284.5
Linkagg-15 - Recovery	BCB-1	14	58.7690	19.283	1.2	198.8
Linkagg-16 - Failure	BCB-1	72	119.8951	108.2	5.75	279
Linkagg-16 - Recovery	BCB-1	25	31.0713	35.5	0.367	77.25
Linkagg 23 – failure	BCB-2	54	134.4259	114.9165	73.55	234.9
Linkagg 23 - Recovery	BCB-2	42	72.3393	66.5665	6.733	198.95
Linkagg 24 – failure	BCB-2	60	220.0808	173.7665	99.633	488.7
Linkagg 24 - Recovery	BCB-2	45	117.2207	94.9	33.2	271.1
Port 1/1/17 – Failure	BCB-1	72	85.2127	60	26.35	193.9
Port 1/1/17 - Recovery	BCB-1	21	45.5634	59.433	3.033	93
Port 2/1/17 – Failure	BCB-2	64	55.7010	58.15	28.2	90.05
Port 2/1/17 - Recovery	BCB-2	14	26.7869	25.875	6.1	60

Table 12: Convergence during BCB link toggle

Traffic was observed to converge within one second in all test cases. The highest convergence time observed for an individual flow was 488.7 ms.

VC takeover

VC takeover testing was performed to validate convergence behavior during chassis failover events. In the tested topology, BCB-1, BCB-2, BEB-3, and BEB-5 in the SPB backbone were configured as Virtual Chassis systems.

BEB-5 was excluded from this test case because it uses an external loopback for VLAN-to-service translation. This external loopback uses one port on each chassis in the Virtual Chassis, which makes the BEB-5 failure scenario different from the other VC-based backbone devices.

For the remaining applicable devices, convergence times were measured separately for chassis-down and chassis-recovery events. The results are summarized in the following table:

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
VC Takeover - Failure	BCB-1	24	91.2	38.55	0.9	524.3
VC Takeover - Recovery	BCB-1	24	32.5333	30.25	3.1	81.9
VC Takeover – Failure	BCB-2	12	55.35	49.95	6.5	144.8
VC Takeover - Recovery	BCB-2	12	36.7666	18.95	6.5	107.7
VC-Takeover – Failure	BEB-3	24	201.3166	107.1	3.3	541

VC Takeover - Recovery	BEB-3	24	29.375	26.7	7.3	72.9

Table13: Convergence during VC-takeover

During VC takeover, traffic converged within one second. The highest convergence time observed for an individual flow was 541 ms.

Access resiliency tests

Access-level redundancy is demonstrated in BLDG-D, where two BEBs provide redundant connectivity to the access switches. To prevent Layer 2 loops in this dual-homed access design, either Spanning Tree Protocol (STP) or Dual Homed Link (DHL) can be used.

The reference topology was validated with both mechanisms. The following sections summarize the convergence behavior observed during the access resiliency tests.

Spanning Tree

In this scenario, STP is used to prevent Layer 2 loops by blocking one of the redundant uplinks on each access switch. AOS uses Rapid Spanning Tree Protocol (RSTP) by default.

Convergence times were measured under the following scenarios:

- BEB reload
- Access link toggle

The following table shows the convergence times observed when one of the BEBs was reloaded. During the failure event, STP unblocked the redundant link and forwarded traffic through the alternate BEB.

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Reload – Failure	BEB-6	72	10,268.23	6,230.55	7.45	30,393.10
Reload – Recovery	BEB-6	72	9,525.40	10,018.65	628.70	48,739.20
Reload – Failure	BEB-7	72	1,395.33	1,545.94	93.37	2,305.25
Reload - Recovery	BEB-7	72	1,933.55	616.98	109.03	5,405.50

Table 14: STP convergence during BEB reload

Traffic converged in approximately 10 seconds on average when BEB-6 was reloaded, and within approximately 2 seconds on average when BEB-7 was reloaded. The difference in convergence behavior may be attributed to the different device models used for the two BEBs.

Next, convergence was measured during link-toggle events on the links between the BEBs and access switches. Link-down and link-up convergence times were measured separately.

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Port 1/1/7 – Failure	BEB-6	0				
Port 1/1/7 – Recovery	BEB-6	0				
Port 1/1/17 – Failure	BEB-6	72	6,659.50	5,296.40	1,380.37	14,202.40
Port 1/1/17 – Recovery	BEB-6	63	354.65	290.30	1.20	1,102.77
Port 1/1/52 – Failure	BEB-7	36	785.59	533.28	59.85	1,826.45
Port 1/1/52 – Recovery	BEB-7	36	29,250.79	383.00	153.50	82,833.75

Table 15: STP convergence during access link toggle

In the tested STP design, all VLANs on an access switch used a single active uplink. This approach was used to prefer the higher-bandwidth link, as described earlier in the STP configuration section of this guide. If required, VLANs can also be load-balanced across the redundant uplinks by configuring different STP priorities for different VLANs.

Dual Homed Link

Dual Homed Link (DHL) provides an alternative mechanism for preventing Layer 2 loops in a dual-homed access topology. DHL blocks the redundant link under normal operating conditions and unblocks it when the active link fails.

DHL is preemptive. When the original active link becomes available again, traffic is moved back to that link after the preemption timer expires. In this test, the preemption timer was configured to 10 seconds.

The following table shows the convergence times observed during BEB reload events.

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Reload - Failure	BEB-6	63	681.64	326.50	1.05	2,127.00
Reload - Recovery	BEB-6	63	2,839.67	1,100.65	3.55	8,720.27
Reload - Failure	BEB-7	72	1,867.98	1,655.68	95.07	4,245.90
Reload - Recovery	BEB-7	72	2,193.69	307.03	0.40	5,932.85

Table 16: DHL convergence during reload

Traffic converged within approximately 3 seconds on average, with the highest observed convergence time being approximately 9 seconds.

The following table shows the convergence times observed during link-toggle events on the links between the BEBs and access switches.

Scenario	DUT	No. of affected flows	Average (ms)	Median (ms)	Best (ms)	Worst (ms)
Port 1/1/7 - Failure	BEB-6	24	438.16	220.25	2.30	1,265.20
Port 1/1/7 - Recovery	BEB-6	24	91.16	10.52	2.80	382.33
Port 1/1/17 - Failure	BEB-6	36	717.31	777.38	1.05	1,796.30
Port 1/1/17 - Recovery	BEB-6	36	99.60	11.45	0.75	402.45
Port 1/1/52 - Failure	BEB-7	24	393.48	349.13	19.45	896.80
Port 1/1/52 - Recovery	BEB-7	24	419.52	340.08	6.85	965.00
Port 1/1/53 - Failure	BEB-7	36	227.08	47.20	0.50	796.30
Port 1/1/53 - Recovery	BEB-7	36	284.08	15.40	1.30	964.60

Table 17: DHL convergence during access link toggle

With DHL, traffic was load-balanced across the two redundant links, with each link carrying approximately half of the VLANs. Traffic converged within 1 second on average, with the highest observed convergence time being approximately 2 seconds.

Conclusion

Shortest Path Bridging (SPB) provides a standards-based fabric architecture that simplifies campus network design while delivering scalable Layer 2 and Layer 3 services across a resilient backbone. By using service-based provisioning, VRF-based segmentation, and automated path computation through the SPB control plane, the network can extend connectivity where needed without relying on complex VLAN propagation or traditional blocked-link designs.

The reference architecture presented in this guide demonstrates how SPB can be used as a flexible foundation for mid-sized enterprise campus deployments. The design includes SPB backbone provisioning, Layer 2 and Layer 3 services, route leaking, multicast switching and routing, access integration, external network interworking, redundancy options, and security features such as DHCP Snooping and Learned Port Security. OmniVista integration further supports centralized visibility and service provisioning workflows, helping reduce operational complexity.

The lab validation confirms that the proposed topology can support scaled traffic conditions while maintaining predictable forwarding and resilient behavior during backbone and access failure scenarios. These results provide confidence that the design can be used as a validated deployment baseline and adapted to meet site-specific requirements such as building count, access redundancy model, segmentation policy, multicast requirements, and upstream routing design.

Overall, the architecture described in this guide offers a practical and extensible SPB deployment model for enterprise campus networks. As requirements evolve, the same design principles can be expanded to larger environments, multi-campus deployments, and deeper integration with external routing, security, and management systems.

Reference documents

- SPB Design Guide
- AOS Network Configuration User Guide
- AOS Advanced Routing User Guide
- AOS CLI Reference User Guide
- AOS Specifications Guide
- AOS Switch Management User Guide
- OmniVista® Documentation: (https://docs.ovcirrus.com/)

Appendix

External loopback configuration

External loopback is a two-pass processing mechanism that requires a physical cable to connect a regular port to a service access port. The regular port is tagged with an IP interface VLAN; the service access port is associated with an SPB Service Access Point (SAP). The VLAN-based IP interface serves as the L3 VPN interface.

An L3 VPN loopback interface configuration consists of the following components:

- An IP interface created in a specific VRF instance and bound to a VLAN ID.
- A regular switch port or link aggregate tagged with the IP interface VLAN.
- A service access port that is assigned to an SPB SAP. The SAP encapsulation is configured with the VLAN ID that is tagged on the regular switch port or link aggregate.
- A physical cable that connects the VLAN port with the service access port to form the loopback configuration. The service access port handles Layer 2 bridging into the service domain and the VLAN port handles the Layer 3 routing into the VLAN domain.

104. Example on BEB-5 · CLI

```
vlan 1306 name "Voice"
vlan 1306 members port 1/1/28 tagged
service 1306 spb isid 1306 bvlan 4002 vlan-xlation enable
service 1306 sap port 2/1/28:1306
ip interface "Voice" address 10.13.80.1 mask 255.255.240.0 vlan 1306
```