Shortest Path Bridging Design Guide

Guide

Alcatel·Lucent

Table of contents

About this document

Purpose

The purpose of this design guide is to present Shortest Path Bridging (SPB) 802.1aq networking concepts along with design guidelines such as architecture, topology requirements and high-level design choices.

Audience

The intended audience for this document includes customer and Business Partner networking professionals involved in the design and deployment of enterprise and campus networks.

Scope

This document does not attempt to cover every aspect or architecture option but only the most common, validated and recommended architectures.

Since the Alcatel-Lucent Enterprise products and solutions are under periodic evolution and enhancements, this document is based on the **Alcatel-Lucent Operating System (AOS) 8.10R4** software release and will be updated once new features are added.

This document assumes the Alcatel-Lucent OmniSwitch® product used for SPB configuration supports single-pass inline routing, which is supported in the latest generation ASICs. This is the case for most SPB-supported OmniSwitch products currently sold. This means that IP interfaces can be associated to an SPB service directly, and traffic can be routed between two SPB services or between a VLAN and an SPB service in a single-pass operation without physical or front-panel loopbacks. A slightly different variation of the configuration steps is required for products supporting only external physical or front-panel loopback routing.

Please refer to the related documentation for additional details, options and deployment guidelines, as referenced in the Related documents section.

Why the network needs to evolve

Local Area Networks (LANs) have traditionally relied on Spanning Tree Protocol (STP) and its variants (RSTP, MSTP)—collectively referred to as "STP" for simplicity—for loop prevention. STP achieves a loop-free topology by electing a "root bridge" and building a least-cost tree linking the root bridge with other non-root nodes. This least-cost tree is created by pruning (disabling) all branches (links) which are not in the least-cost path towards the root. STP's design principle presents several drawbacks for modern Enterprise networks:

- **Unused links**: Creating a loop-free topology by disabling network links results in inefficient bandwidth use and low Return on Investment (ROI).
- **Sub-optimal paths**: While communication to and from the root bridge follows the least- cost path, communication between non-root bridges may need to traverse a sub-optimal route that transits the root-bridge instead of a better route over links that have been disabled.
- **Slow convergence**: STP is a decades-old protocol designed when network devices were far less powerful than they are today. Even with the "rapid" version of STP, typical convergence times are in the order of seconds. While STP re-converges to a new topology, transient loops may form, resulting in packet drops, link saturation and session timeouts.

Figure 1 - Problems with STP

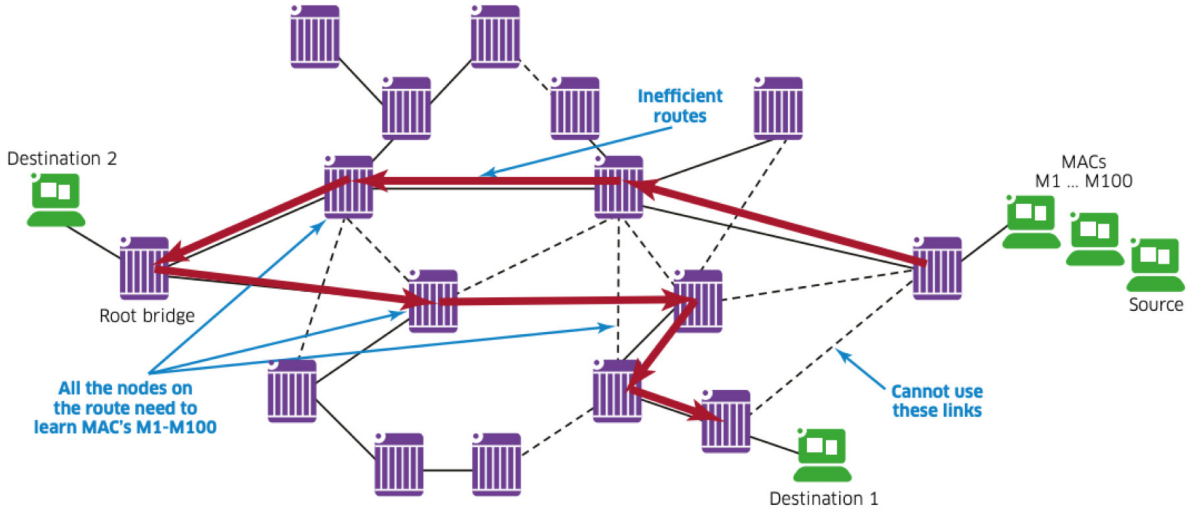

In addition to STP's weaknesses, Ethernet's scalability beyond the LAN is limited by its lack of a coordinated control plane and use of a flat (as opposed to hierarchical) address space. Legacy Ethernet networks present the following challenges:

- **Flooding**: Ethernet's "flood-and-learn" address learning floods unknown-unicast traffic until the destination address is learned from return traffic.
- **MAC Learning**: All nodes in the LAN learn all end-device MAC addresses, which createsa scalability challenge.
- Lastly, IEEE 802.1ad (Provider Bridging or Q-in-Q) is limited to a maximum of 4096 service instances.

Introducing Shortest Path Bridging (SPB)

802.1aq Shortest Path Bridging (SPB) is an IEEE networking standard whose primary focus was addressing the challenges in STP. However, SPB is much more than STP's evolution: SPB provides MPLS-like VPN services but is significantly simpler to deploy and maintain. Additionally, unlike Multiprotocol Label Switching (MPLS), which requires a "stack" of protocols (LDP, OSPF, MP-BGP and others), SPB relies on a single protocol to provide this functionality: IS-IS (Intermediate System to Intermediate System). IS-IS is the only control plane protocol required to build a multi-path topology, perform address learning and carry VPN routes across the backbone. Alcatel-Lucent Enterprise's SPB implementation brings further simplification by automating client device attachment and dynamic service instantiation. Because of this simplicity and automation, an ALE-powered SPB solution offers high-end services for a lower Total Cost of Ownership (TCO). The following sections explain the benefits of SPB in more detail.

Scalable, fast-converging, multi-path fabric

SPB's loop-free topology is built by a link-state routing protocol running Dijkstra's Shortest Path First (SPF) algorithm: IS-IS. With IS-IS, no network link is disabled, all paths are available, and traffic between any pair of nodes follows the shortest path. In addition, with MAC-in-MAC encapsulation, backbone nodes do not learn any end-device MAC addresses, increasing network scalability and stability. With IS-IS and MAC-in-MAC encapsulation, SPB creates an any-to-any, scalable, fast-converging "fabric" that supports multiple active optimal paths for both bridged and routed traffic.

Figure 2 - Addressing challenges of STP

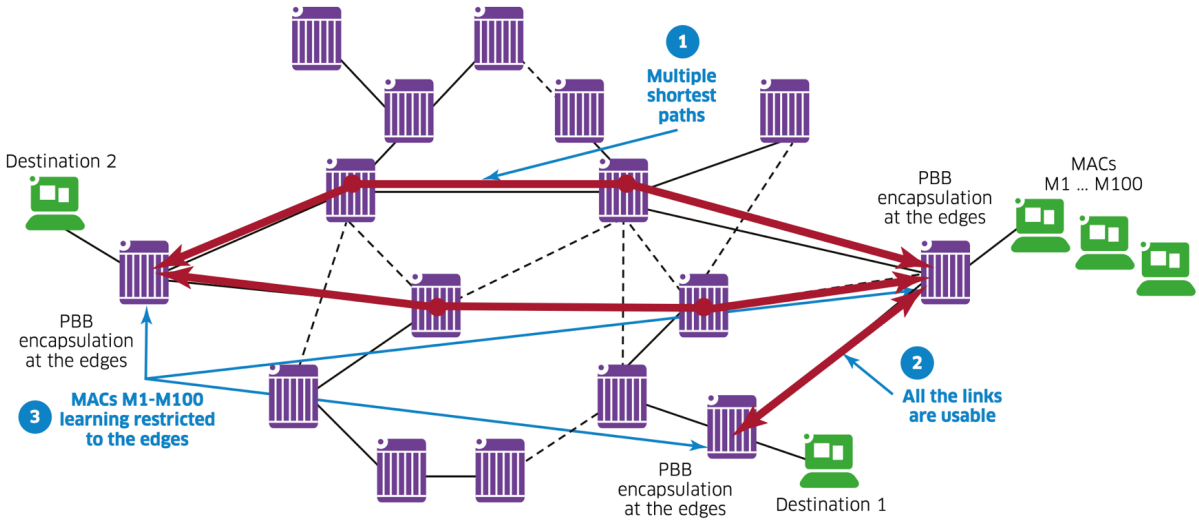

Multi-tenancy

SPB natively supports multi-tenancy: the physical network is partitioned into multiple virtual "slices" referred to as VPNs, "containers" or "communities". Customers or IoT device groups segregated into different VPNs are isolated and do not interfere with one another. In fact, they can use overlapping address space without conflict. Inter-VPN communication, if needed, is tightly controlled by firewall policies. This multi-tenancy capability makes SPB suitable for use cases such as smart cities, transportation, higher education, video surveillance or data centers. SPB's scalability is not limited to 4096 tenants because its service identifier (the ISID) is a 24-bit field which can differentiate up to approximately 16 million services.

Figure 3 - Multi-tenancy

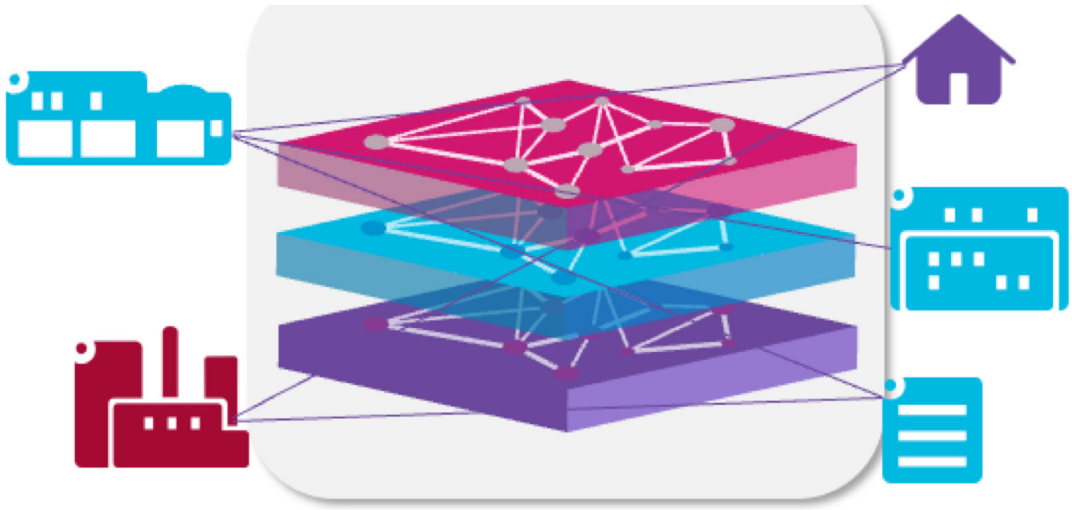

Dynamic service instantiation

SPB services do not need to be statically bound to a switch port. SPB is tightly integrated with ALE's classification and Network Admission Control (NAC) framework called Access Guardian (AG). Upon connection, end devices can be classified (for example, based on the MAC OUI or IoT "fingerprint" rules) or authenticated (for example through 802.1X or MAC) against a RADIUS server. The appropriate service is dynamically instantiated according to the device or user classification, or role attribute returned by the RADIUS server. In the same manner, this user-to-service binding is removed when the user or device disconnects. This dynamic service instantiation has the following advantages:

- **User/device mobility**: The network configuration dynamically adapts to mobile users and devices or Virtual Machine (VM) migrations without requiringn move, add or change requests.
- **Increased security**: Services are instantiated on an as-needed basis and for authenticated devices/users only, if applicable. This association is maintained for as long as the user/device remains connected and/or authenticated and is brought down on disconnection/log-off. These ephemeral services are inherently more secure: they cannot be scanned, subjected to a Denial of Service (DoS) or otherwise hacked, while they are not active.
- **Device templates**: The dynamic instantiation of network services easily lends itself into template-based configuration of network nodes. Edge nodes can share the same base configuration template and dynamically adjust the service configurations on the fly.

Edge-only service provisioning

Whether statically or dynamically instantiated, SPB services need to be provisioned only on edge nodes, not on core nodes. Core nodes are effectively isolated from service moves, adds and changes and require no touch while these activities are performed. In fact, service MACs can be conducted during business hours and do not require a maintenance window to be scheduled, reducing time-to-service.

Micro-segmentation

Firewalls filter and control communication between different VPN "tenants" or "containers", but how do you secure communication within the same VPN? For instance, if one device were compromised, how do you prevent lateral movement to other resources within the same VPN?

When users/devices are dynamically bound to a service, they are also mapped to a Universal Network Profile (UNP). The UNP is a set of Access Control Lists (ACLs) and Quality of Service (QoS) policies which are applied to the device/user according to the device category or user role. In the case of CCTV cameras, for example, ACLs contained in the UNP can allow communication between the camera and surveillance servers but block camera-to-camera communication, preventing the spread of malware, "pivoting" and other hacking techniques that rely on lateral movement.

Figure 4 - Micro-segmentation

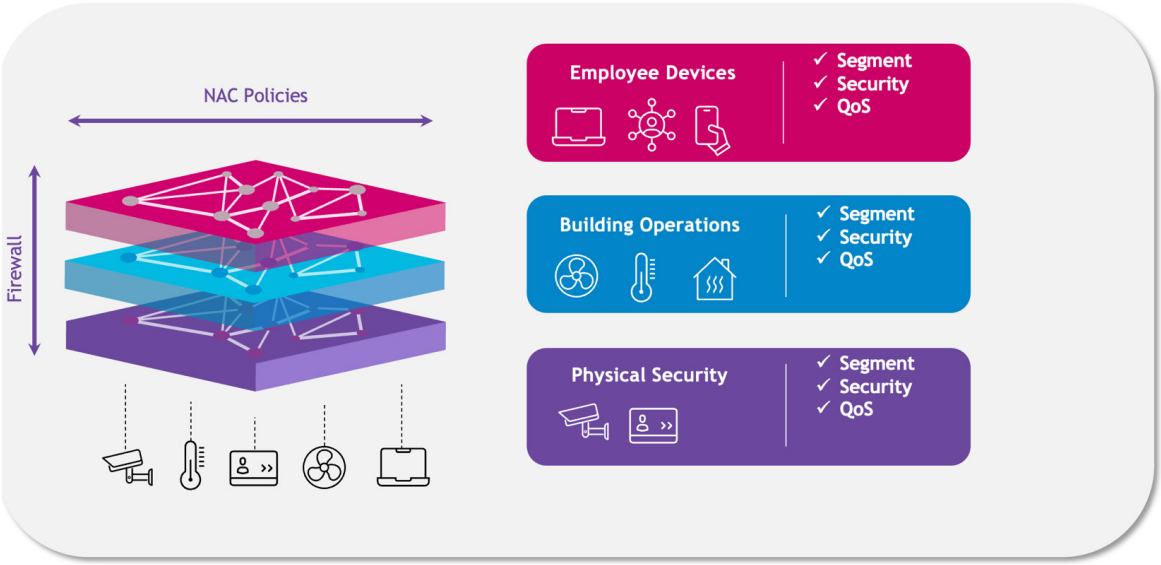

Non-IP core

Even when providing Layer 3 services to IP packets, SPB core nodes do not route traffic, they bridge it. In fact, SPB core nodes do not have IP addresses, and the IS-IS control protocol, unlike OSPF and BGP, does not run on top of IP. This makes the network core inherently more secure and protects it from IP-based attacks such as scanning, spoofing, DoS and others. Of course, SPB nodes still need an IP address for management purposes, but the management IP interface is isolated in its own service and Virtual Routing and Forwarding (VRF) instances, not in-line with user traffic.

Multi-topology technology

SPB offers a comprehensive range of Metro Ethernet service types, including E-Line, E-LAN and E-Tree. These diverse service options provide flexibility when selecting the appropriate topology for a specific use case. Detailed information on these service types is available in the SPB service types section.

Data Plane: IEEE 802.1ah PBB

The Data Plane (DP) forwards user traffic between different ports. The DP makes no decisions as to what port a frame should be forwarded to; it simply performs lookups on the Forwarding Data Base (FDB). FDB entries indicate what port or group of ports each frame should be forwarded to and what encapsulation to use. Building or populating entries in the FDB is a function of the Control Plane (CP), which is discussed in the next section.

The SPB data plane employs IEEE 802.1ah Provider Backbone Bridging (PBB) (also known as MAC-in-MAC encapsulation). The PBB header includes the following fields:

- **B-VID: Backbone VLAN (BVLAN) ID**. This VLAN transports the SPB service instances and connects SPB bridges together through Shortest Path Tree (SPT) sets. Unlike the standard VLAN domain, which uses "flood-and-learn" or source learning in the DP to populate the FDB, the BVLAN domain's FDB is pre-populated by the CP.
- **ISID: Service Instance Identifier**. The ISID is a 24-bit number that designates the service instance, tenant, container or VPN. Different SPB services are assigned different ISIDs and isolated from one another. Each SPB service or ISID is bound to a BVLAN.
- **B-SA and B-DA: Backbone Source and Destination MAC Addresses**. These are the MAC addresses associated with SPB nodes (BMACs). Within the SPB backbone, traffic is forwarded based on the destination BMAC (B-DA). Inner customer MACs are not learned or used for forwarding within the backbone.

- **Ethertype 0x88E7**: Upon entering the SPB domain, the PBB header is wrapped around the incoming frame which can be un-tagged, single-tagged (IEEE 802.1Q) or double-tagged (IEEE 802.1ad). Figure 5 illustrates a double-tagged (Q-in-Q) frame. In this diagram, MAC and BMAC addresses are shortened to 2 bytes for simplicity.

Figure 5 - PBB Data Plane

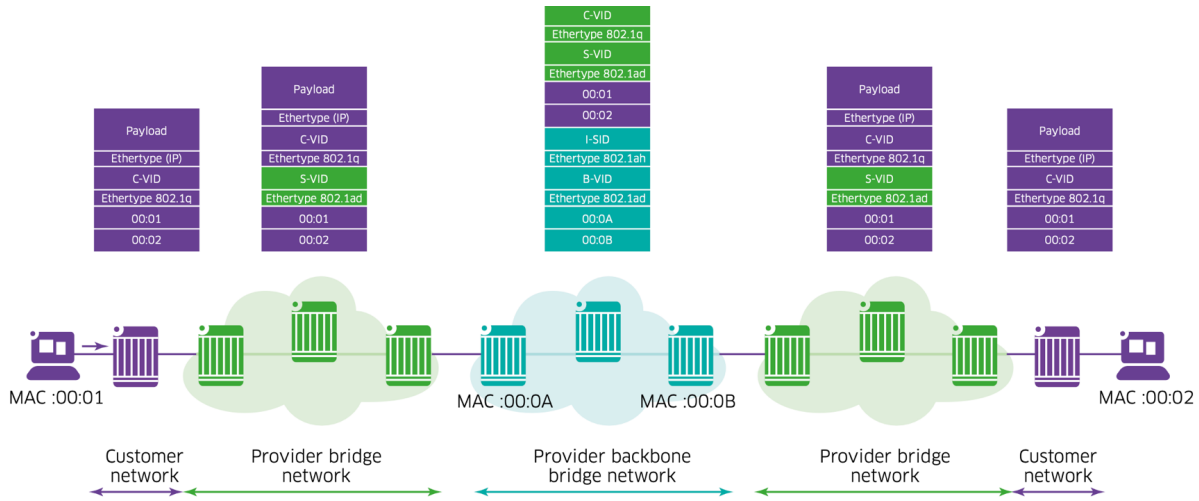

Key terms:

- **BEB**: An SPB switch positioned at the edge of the PBB network that learns and encapsulates (adds an 802.1ah backbone header to) "customer" frames for transport across the backbone network. The BEB interconnects the customer network space with PBB network space.
- **BCB**: An SPB node that resides inside the PBB network core. The BCB employs the same BVLAN on two or more network ports. This BVLAN does not terminate on the switch itself; traffic received on an SPB network port is switched to other SPB network ports. As a result, the BCB does not have to learn any of the customer MAC addresses. It mainly serves as a transit bridge for the PBB network.

Within the SPB domain (that is, between BEB and BCB nodes) frame forwarding depends entirely on the outer PBB 802.1ah header (BMAC and BVLAN) and not on the inner header or "customer" MAC addresses (CMAC). In fact, the SPB backbone nodes do not learn CMACs. This makes SPB networks more scalable and stable because CMACs are not learned and therefore do not need to be flushed and re-learned when they change or move.

The DP implements an additional loop mitigation mechanism by which a node will not accept unexpected frames from their neighbours. This additional loop mitigation mechanism is faster during topology changes. In summary, SPB implements two loop avoidance mechanisms: loop prevention and loop mitigation.

Control Plane: RFC 6329 IS-IS ECTs

The CP populates the FDB tables used by the DP. SPB uses IS-IS (ISO/IEC 10589:2002), a well-known, proven and widely-deployed protocol, particularly in service provider backbones. IS-IS is responsible for topology and service discovery. IS-IS is an extensible link-state protocol which implements Dijkstra's SPF algorithm for path computation. IS-IS extensions for SPB are described in RFC 6329 and include a new Network Layer Protocol Identifier (NLPID), as well as a set of Type-Length-Values (TLVs). These extensions add support for multiple topologies, allowing load sharing over multiple equal-cost paths, and service-membership discovery. In other words, they communicate what services are enabled on each SPB node.

Figure 6 - RFC 6329 IS-IS extensions

SPB-ISIS	New!	SPB extensions	NLPI, TLVs, PDUs
	Existing!	Discovery and computation	Discovery – Hello and LSP packets, Computation – SPF and SPT

Unlike STP, which creates a single tree rooted at the root bridge, in SPB networks, every node builds a topology tree rooted on itself. This is the key reason why, in an SPB network, traffic between any pair of nodes always travels along the shortest path. When using STP, traffic between two nodes does not necessarily travel over the shortest path unless one of the two nodes involved is the root bridge. This is illustrated in Figure 7 in which B1 is the root bridge. Traffic between nodes B5 and B2 (neither of which is the root bridge) could not use the direct single-hop path because that link is disabled by STP. Traffic between these two nodes must take a three-hop detour traversing the root bridge.

Figure 7 - Multiple trees

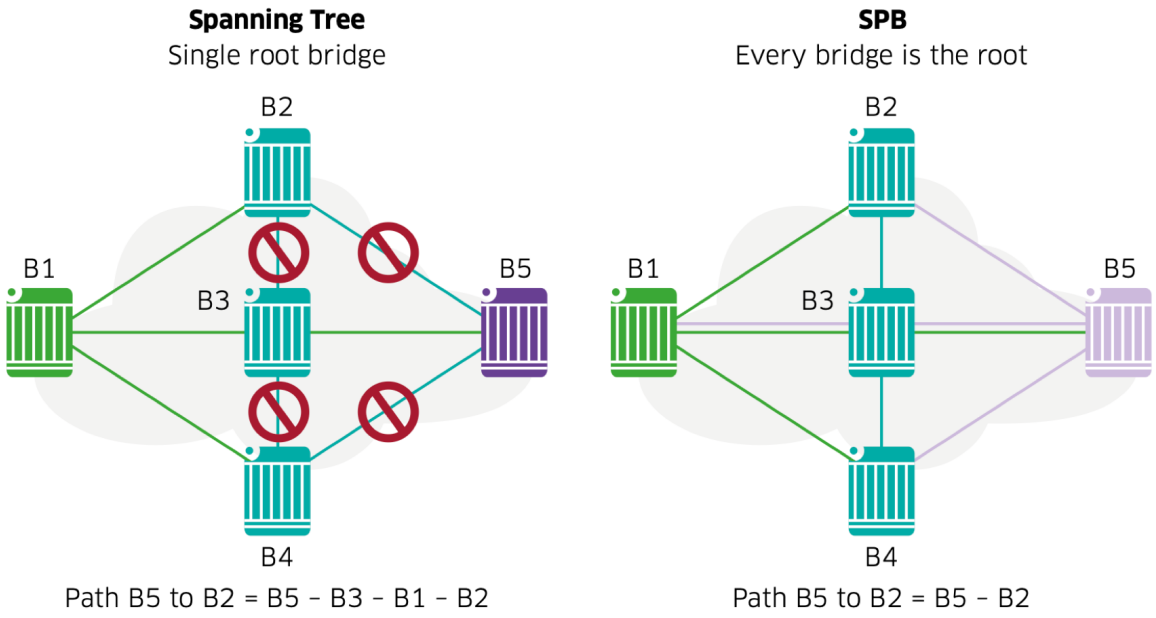

In contrast, when using SPB, no link is disabled: each node is the root of its own tree. Nodes B2 and B5 can simply communicate over the direct single-hop path while at the same time they can communicate with other nodes over different paths (for example, between B4 and B5). SPB's support for multiple trees and multiple active paths unlocks utilization of bandwidth in optimal paths that would otherwise be wasted, increasing throughput and reducing latency.

An SPB network supports up to 16 BVLANs, and each node builds an SPF tree for each BVLAN. Load balancing is accomplished by mapping different tenant services (ISIDs) to different BVLANs. Service traffic between any node pair uses a single path, which only changes if the topology changes (for instance, on node or link failure and subsequent path re-computation). In other words, SPB networks do not balance loads on a packet-by-packet basis like IP networks do. Provided the physical topology supports multiple shortest paths (same cost and same hop count) between two nodes, different BVLANs can build different trees, and services mapped to those BVLANs can use different paths, which will remain the same for as long as the topology remains the same. An important property of SPB networks is that network paths are deterministic, and frames are delivered in the order they were sent. This property is important for certain applications such as storage and real-time application traffic.

Figure 8 - One tree per node and per BVLAN

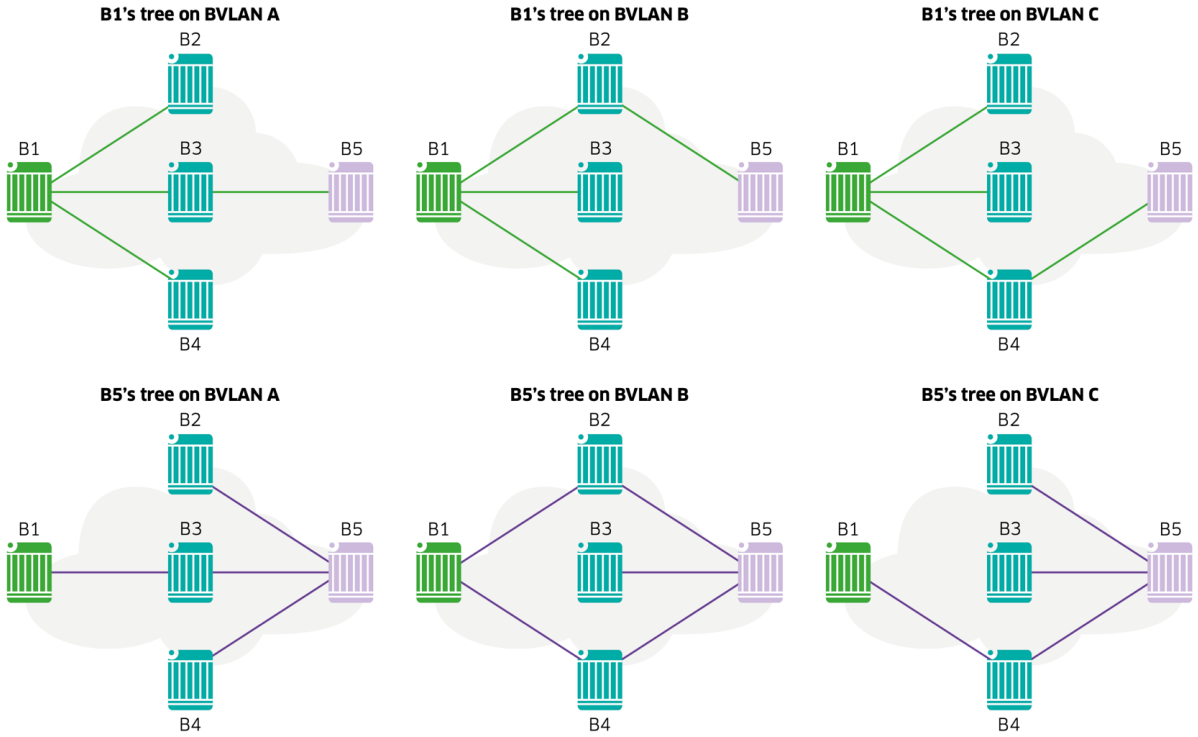

The trees shown in Figure 8 are SPB equal-cost trees (ECTs). Each node builds a tree per BVLAN and the cost to reach other nodes is the same across all BVLANs. The ECT-ID is a number assigned to each BVLAN at the time of BVLAN creation and is used for tie-breaking during path computation. Assigning different ECT-IDs to different BVLANs helps those BVLANs build different trees, provided the underlying topology supports multiple equal-cost or shortest paths.

All bridges use predefined ECT algorithms to calculate Layer 2 congruency and symmetry for switching. The standard provides 17 pre-defined ECT algorithms with the first one (ECT-MASK[0]) reserved for a Common Spanning Tree (CST) algorithm. The algorithms defined by the standard are (only index 1-16 is used):

ECT-MASK[17] = { 0x00, 0x00, 0xFF, 0x88, 0x77, 0x44, 0x33, 0xCC, 0xBB, 0x22, 0x11, 0x66, 0x55, 0xAA, 0x99, 0xDD, 0xEE }

The calculation for the SPT is done by calculating a byte-by-byte XOR operation on the ECT-MASK (16 masks to provide 16 ECTs) for all nodes excluding source and destination. This provides a sorted list of Bridge IDs computed as ECTs. The Bridge ID is populated as explained below:

Note: The System ID in this calculation is the system MAC address, and the default priority is 32768 (configurable).

Bridge ID = System ID (6 bytes) + Priority (2 bytes)

Another method to influence the SPT calculation is to modify the bridge priority for the switch. or change the link cost metric for the SPB interface connection between two switches. SPB interfaces are associated with a link metric cost that is configurable, which provides the ability to change the logical topology created by the ISIS-SPB instance. However, if different metric values are configured on each side of a link, ISIS-SPB will choose the higher-valued one as the metric to use for both sides. This is necessary to enforce the symmetry of SPT calculations in both directions across the link.

The SPT calculation precedence is highlighted in the following table:

Priority	Name	Default	Description
1	Shortest metric	10	Lowest link metric has highest priority. SPB path calculations use the maximum value of the two nodes when the metric is different.
2	Lowest hop count	N/A	When link metric is the same, hop count is used.
3	ECT-ID	1	When multiple links have an equal cost (metric and hop count), ECT-ID is used.

Another important property of SPB networks is path symmetry. In Figure 9, the path from node X to node Y is identical to the path from node Y to node X. Path symmetry is key to Operations and Maintenance (OAM). For instance, one-way delay calculations can easily be derived from roundtrip delay measurements. Note that this is not the case for other IP-based technologies such as MPLS in which the reverse path may differ.

Figure 9 - Symmetric paths with per-BVLAN load balancing

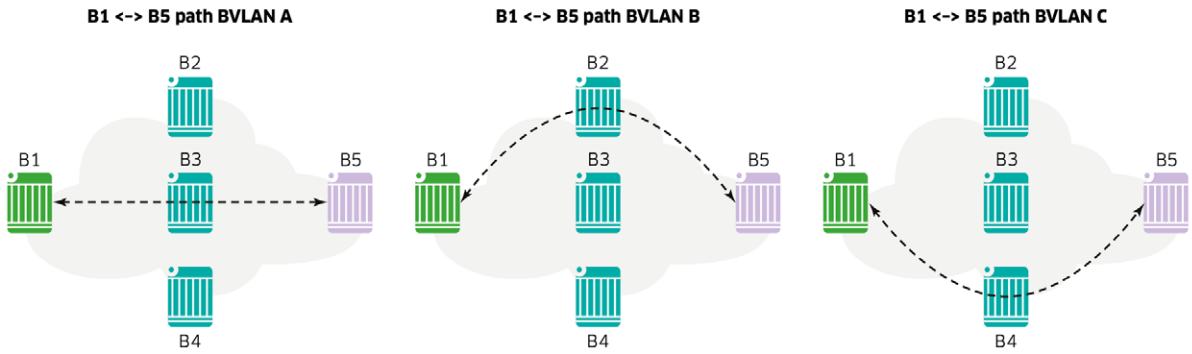

The result of IS-IS path computation for each BVLAN and node is the FDB, which is used by the data plane for frame forwarding. Figure 10 shows B5's unicast FDB. The multicast FDB will be discussed in the BUM traffic section.

Figure 10 - B5's unicast FDB

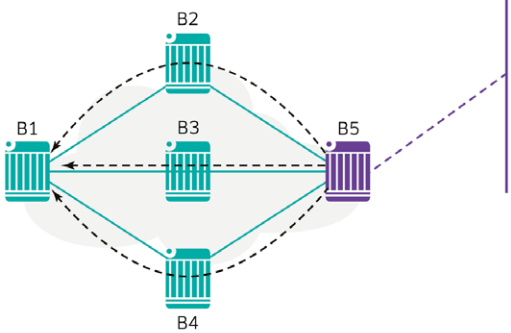

BVID	Node	Outbound port
BVID A	B1	Port 1
BVID B	B1	Port 2
BVID C	B1	Port 3
BVID A	B2	Port 1
BVID B	B2	Port 1
BVID C	B2	Port 1
BVID A	B3	Port 2
BVID B	B3	Port 2
BVID C	B3	Port 2
BVID A	B4	Port 3
BVID B	B4	Port 3
BVID C	B4	Port 3

Service framework

An SPB service represents a VPN or tenant and is uniquely identified by its service identifier (ISID). An SPB service needs only to be created or instantiated on BEB nodes (not on BCB nodes) and only on those BEB nodes servicing locations associated to the service. SPB service membership information is shared across the SPB backbone by way of IS-IS TLVs so all SPB nodes have a consistent view of the active services on each BEB. Each node then builds a service database.

Figure 11 - Service database

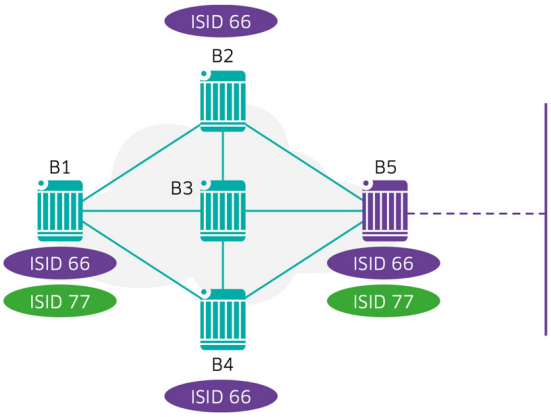

ISID	BVID	Node
66	BVID A	B1
66	BVID A	B2
66	BVID A	B4
66	BVID A	B5
77	BVID B	B1
77	BVID B	B5

In each BEB node there are two kinds of virtual ports:

- **Service Access Point (SAP)**: The SAP is a UNI-side logical port that binds a physical port and specific customer traffic types (untagged, single-tagged, double-tagged or all) to an SPB service. Multiple SAPs can be associated to the same physical port thus multiplexing and mapping different customer traffic encapsulations to different SPB services.
- **Service Distribution Point (SDP)**: The SDP is an NNI-side logical port that binds an SPB service to a far-end BEB on which the service is instantiated. SDPs are dynamically created in the CP and only for those far-end BEBs with SAPs for the specific service.

SAPs are mapped to service instances (ISIDs). A SAP always maps to a single service, but a given service can accommodate multiple SAPs. Services are also bound to SDPs. SDPs distribute the service connectivity to other BEBs through SPTs. Multiple services could be bound to a single SDP for multiplexing service traffic. It is always possible to have a single service bound to multiple SDPs providing E-LAN connectivity for customer frames. There are other types of supported service types such as E-Tree and E-Line, which are covered in the SPB service types section.

Customer traffic enters a BEB through Service Access Ports, which can be regular, physical Ethernet ports or link aggregate ports (static and LACP ports). These are not to be confused with logical Service Access Points (SAPs).

In Figure 12, B5 terminates 2 SPB services: one is associated to ISID 66 and the other to ISID 77. There are two SAP ports, one for each service. SAP 1:1 is defined on port 1, matches traffic tagged with VLAN 1 and binds it to service 66. SAP 2:2 is defined on port 2, matches traffic tagged with VLAN 2 and binds it to ISID 77.

ISID 66 is also enabled on nodes B1, B2 and B4 while ISID 77 is also enabled on node B1.

Figure 12 - Service framework

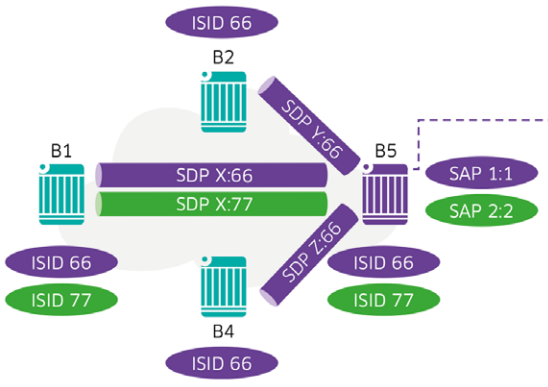

Identifier	BVID	Node
SAP 1:1	–	–
SAP 2:2	–	–
SDP 32769: 66	BVID A	B1
SDP 32768: 66	BVID A	B2
SDP 32767: 66	BVID A	B4
SDP 32766: 77	BVID B	B1

While BMAC address learning is performed in the CP (for example, not through "flood-and-learn"), CMAC address learning is performed in the BEB's DP through flood-and-learn. Near-end CMACs are bound to SAP ports, and far-end CMACs are bound to SDP ports. BCB nodes have neither SAP nor SDP ports and therefore do not learn any CMACs.

Figure 13 expands this example by adding some end-customer sites and CMACs associated to those customers. Two-byte MAC addresses are still used, for simplicity. In Figure 13, near-end CMAC addresses are bound to SAP ports while far-end CMAC addresses are bound to SDP ports. Within the service domain, a BEB performs CMAC source address learning like a standard Ethernet switch, except there is no "flooding" of BUM traffic. BUM traffic is discussed in the BUM traffic section.

Figure 13 - Customer MAC address learning

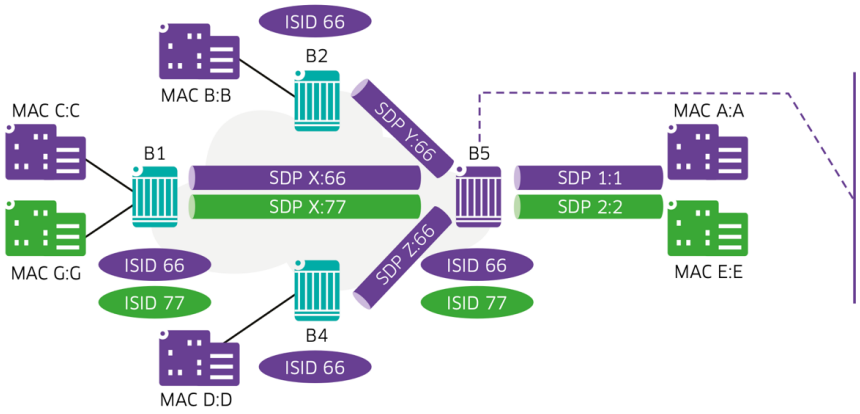

Identifier	ISID	CMAC
SAP 1:1	66	A:A
SAP 2:2	77	E:E
SDP 32769: 66	66	C:C
SDP 32768: 66	66	B:B
SDP 32767: 66	66	D:D
SDP 32766: 77	77	G:G

SAP identifiers define the rule for classification of customer traffic into services. The encapsulation of incoming traffic can be un-tagged, single-tagged (802.1Q) or double-tagged (Q-in-Q). SAP identifiers provide a very flexible mechanism to classify the traffic based on the encapsulation of incoming traffic.

For example, SAPs can be created on a given access port to specify that all double-tagged frames with the combination of outer tag of 100 and inner tag of 200 should go into service 1000, every single tagged frame with a tag of 100 should go to service 200, and all other traffic should be classified into service 2000. The SAP identifiers in the above case would be represented by:

BEB-1

```
BEB-1> service 2000 sap port 1/1/48:0
BEB-1> service 200 sap port 1/1/48:100
BEB-1> service 1000 sap port 1/1/48:100.200
```

A strict precedence rule will make sure there is no ambiguity in the classification scheme or its interpretation. The SAP classification will follow the precedence order as stated below:

1. Q-in-Q (Outer VLAN + Inner VLAN) – **highest**
2. Q-in-Q (Outer VLAN + *)
3. 802.1Q (VLAN)
4. 802.1Q (*)
5. Untagged – **lowest**

A given service may require different encapsulations on different SAPs. For instance, a server may tag traffic with a specific VLAN while client devices may require untagged SAPs. In this situation, VLAN translation can be enabled to allow both devices to communicate. Translation will need to be enabled per access port. This will implicitly enable the translation for all the SAPs of this port. The SAPs on an access port can each belong to different services, and it may not be desired to have translation for all these services. In that case, VLAN translation can be provided on a per service basis.

In summary, support for VLAN translation needs to be enabled at two levels: first, on the access port and then on the service. In this way all the SAPs of an access port that belong to the same service can have the translation property.

SAPs and services can also be dynamically created, as explained in the Automation section.

Hybrid port

Hybrid port can be configured for scenarios where a customer needs some VLANs (CVLANs) for services and others for bridging/routing on BEBs. This will allow a single port to function both as an access port and a bridging port. Without hybrid port configuration, separate physical ports would be required for SAP ports and standard VLAN (802.1Q) ports. A hybrid configured port can be understood as a bridge port with a default VLAN and tagged VLAN for bridging, and the user can configure SAPs for services with mapped tagged VLANs. Another use case for this feature is easier migration of VLAN access ports to service access ports.

Figure 14 shows an example of traffic treatment on a hybrid configured port. When different domain (SAP/VLAN/Default-VLAN) traffic from customer network is received on a hybrid configured port at BEB from an aggregator switch, the traffic gets classified to their respective domain. SAP VLAN tagged traffic is processed in the SPB service domain, and the regular VLAN tagged (default) packet gets processed in VLAN domain.

Figure 14 - Hybrid port

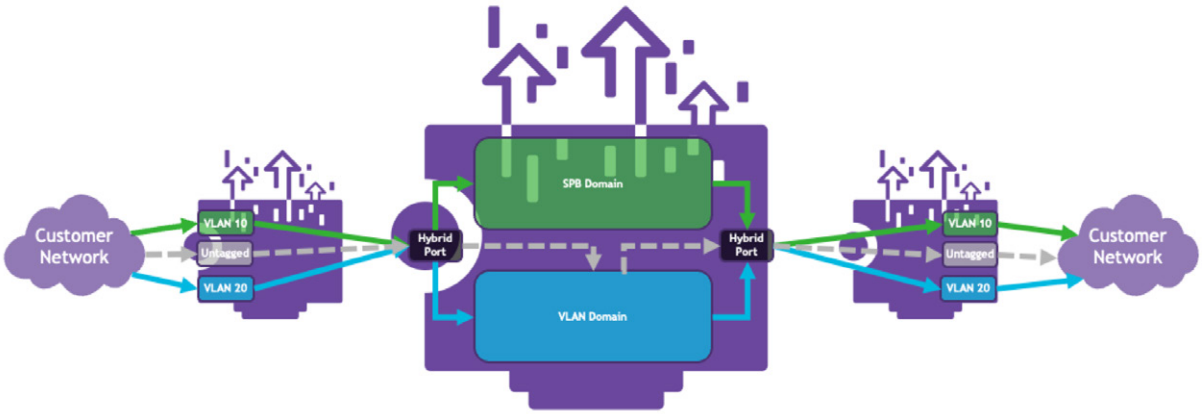

On a hybrid configured access port, single-tagged and double-tagged SAP will be supported but untagged and catch-all SAP will not be supported

Following is a sample configuration where a service access port can be configured as a hybrid port. The same can also be configured for link aggregation service access ports:

BEB-1

```
BEB-1> service access port 1/1/1 hybrid enable
```

BUM traffic

SPB supports 3 BUM (Broadcast, Unknown Unicast and Multicast) traffic replication and forwarding methods: Head-end, Tandem (S,G), and Tandem (*,G).

Head-end: In this mode, BUM traffic received on a SAP port is replicated at the ingress BEB and converted to multiple unicast frames. A replica is created for every other BEB in the same ISID, and these replicas have the BEB BMACs as the B-DA and are forwarded using the unicast FDB. For this reason, head-end replication can be inefficient in terms of bandwidth consumption but is efficient in terms of resource usage because it does not require a separate tree. However, head-end replication can be optimal in some circumstances, particularly when combined with IGMP Snooping. The OmniSwitch implementation of IGMP Snooping is called IP Multicast Switching (IPMS). Head-end replicated BUM traffic simply uses the unicast FDB and therefore travels along the same path. This property is known as congruency.

Figure 15 - Head-end BUM replication

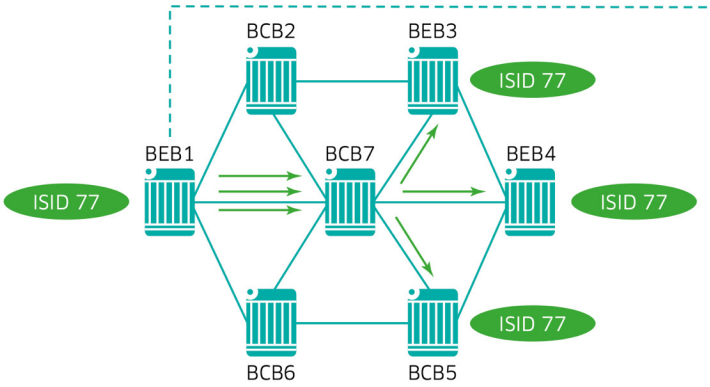

Type	B-MAC	Port	BVID
U	00:02	Port 2	BVID B
U	00:03	Port 7	BVID B
U	00:04	Port 7	BVID B
U	00:05	Port 7	BVID B
U	00:06	Port 6	BVID B
U	00:07	Port 7	BVID B

Tandem (S,G): In this mode, a separate multicast SPT and FDB are created. The multicast SPT is also congruent with the unicast SPT, however the B-DAs in the multicast FDB are multicast addresses constructed as a combination of ISID and source BEB BMAC. This special multicast group address will be explained next. When a BUM frame is received on a BEB, it is MAC-in-MAC encapsulated with this special BMAC as the B-DA and forwarded according to the multicast FDB. A backbone node can use the unicast FDB to check if it is in the SPT between a source BEB and other BEBs in the same ISID. If the backbone node happens to be in the SPT, it will populate the multicast FDB such that the frame is replicated and forwarded as needed to other BEBs connecting the same service (ISID). Tandem Replication is very efficient in terms of bandwidth use because it will only send a single replica on any given link; however, it is less efficient in terms of resource use because it requires an additional SPT and multicast FDB per ISID.

Figure 16 - Tandem (S,G) BUM replication

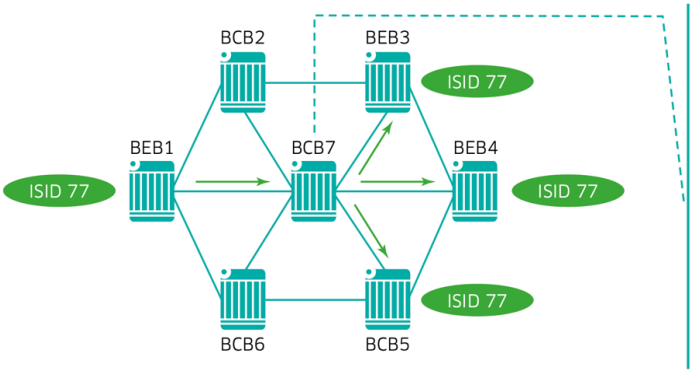

Type	B-MAC	Port	BVID	Out Intf
U	00:01	Port 1	BVID B	
U	00:02	Port 2	BVID B	
U	00:03	Port 3	BVID B	
U	00:04	Port 4	BVID B	
U	00:05	Port 5	BVID B	
U	00:06	Port 6	BVID B	
M	00:0W	Port 1	BVID B	Port 3/4/5
M	00:0X	Port 3	BVID B	Port 1/5
M	00:0Y	Port 4	BVID B	Port 1
M	00:0Z	Port 5	BVID B	Port 1/3

Tandem (*,G): In this mode, a separate multicast tree is created. This tree is not an SPT and is not congruent with the unicast SPT. A multicast (*,G) is created for every BVLAN using Tandem (*,G) multicast replication. This (*,G) tree is similar to a Spanning Tree and is rooted at one backbone node according to the bridge priority. In this mode, there is a single tree for the BVLAN and not one tree for every node. Therefore, traffic will not generally follow the shortest path. This mode is a compromise between bandwidth and resource usage but it can be a good option when all traffic is sourced or destined towards the root bridge.

The following table compares these three modes.

Table 1 - Multicast replication modes and suggested uses

	Head-end	Tandem (S,G)	Tandem (*,G)
Operation	BUM traffic replicated at ingress BEB and forwarded using the unicast FDB	BUM traffic forwarded per the multicast FDB and replicated as needed at the SPT's fork-out points	BUM traffic forwarded using a shared, non-SP tree and replicated at fork-out points
Bandwidth efficiency	Low	High	High
Resource efficiency	High	Low	Medium
Congruency	Yes	Yes	No
Suggested use	• Low multicast bandwidth • Many sources and few receivers*	• High multicast bandwidth • Few sources and many receivers*	• When root bridge is source or receiver of most multicast traffic and congruency is not required • When inter-operation with third-party equipment is required

* when combined with IGMP Snooping

As mentioned earlier, Tandem replication modes use a special B-DA multicast address. This special multicast address is derived from the B-DA unicast address and the ISID information as shown in Figure 17.

Figure 17 - Tandem replication multicast group BMAC

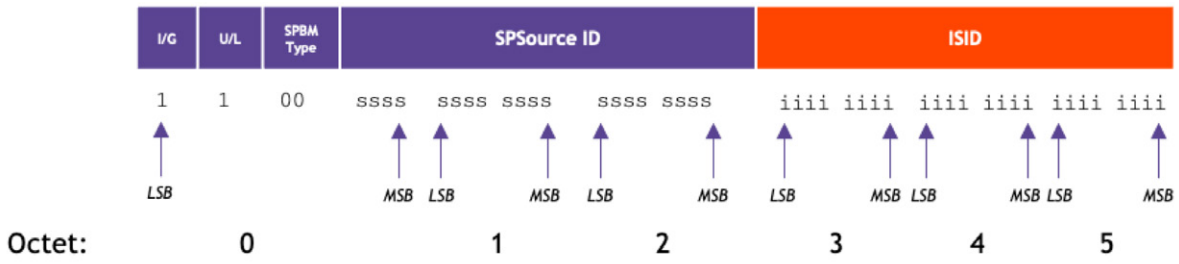

The least significant bit (LSB) and the next to LSBs of the first octet of the address, the Individual/Group (I/G) and Universally/Locally (U/L) administered bits, are not used and both set denoting a locally administered group address as follows:

- I/G (multicast bit) = 1
- U/L (local bit) = 1
- SPBM type = 00
- SPSourceID = 20-bit 'short-form' Bridge ID of the source of the multicast frames
- ISID = 24-bit ISID encoded from the ISID value

By default, the last three least significant bytes of the system ID are used for the source ID.

Creating an SPB backbone

This section describes a sample SPB backbone configuration and refers to Figure 18 as a sample topology. Nodes BEB-1 through BEB-4 are called “BEB” nodes because services will be added to these nodes later. Node “BCB” will remain as a pure transit node and not terminate any services.

This topology provides up to three shortest paths, for example, between nodes BEB-1 and BEB-3 or between nodes BEB-2 and BEB-4. To take advantage of those three diverse paths for traffic load balancing, we need to create a minimum of three BVLANs. In this example, we will dedicate one BVLAN for control traffic and therefore we will create a total of four BVLANs. This is not strictly necessary, as the control BVLAN can also be used for services.

Backbone configuration includes these tasks:

- Creating one or more BVLANs with their associated ECT-IDs. ECT-IDs need not be explicitly defined, default ECT-IDs are applied
- Defining the control BVLAN
- Defining one or more SPB IS-IS interfaces
- Enabling the SPB IS-IS protocol

Figure 18 - Sample backbone topology

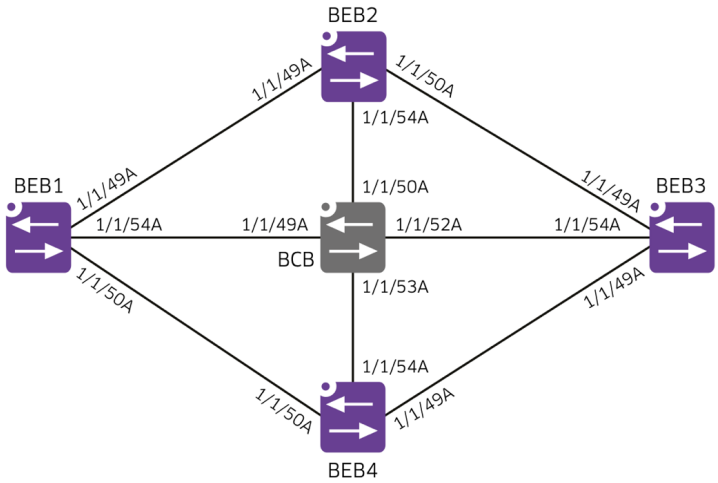

Following are the sample configuration snippets BEBs and BCB:

BEB-1 to BEB-4

```
BEB-X> spb bvlan 4000-4003
BEB-X> spb isis control-bvlan 4000
BEB-X> spb isis interface port 1/1/49A
BEB-X> spb isis interface port 1/1/50A
BEB-X> spb isis interface port 1/1/54A
BEB-X> spb isis admin-state enable
```

BCB

```
BCB> spb bvlan 4000-4003
BCB> spb isis control-bvlan 4000
BCB> spb isis interface port 1/1/49A
BCB> spb isis interface port 1/1/50A
BCB> spb isis interface port 1/1/52A
BCB> spb isis interface port 1/1/53A
BCB> spb isis admin-state enable
```

Through this configuration, VLANs 4000 through 4003 are defined as SPB backbone VLANs and will therefore not use any form of spanning tree protocol. AOS automatically assigns a different ECT-ID to each BVLAN. This maximizes the chance that different BVLANs will create different SPTs, up to the maximum number of shortest paths supported by the physical topology. Nodes will exchange IS-IS "Hello" messages over the control BVLAN (4000 in this example) and form point-to-point adjacencies. IS-IS Link State Packets (LSPs) are exchanged, a topology database is created and one SPT is built for each BVLAN.

Here is the configuration with some show commands.

BEB-1

```
BEB-1> show spb isis interface
SPB ISIS Interfaces:
                                  Oper    Admin    Link       Hello    Hello   Circ
 Interface       Level   CircID   state   state    Metric     Intvl    Mult    Type
---------------+-------+---------+------+-------+---------+-------+------+------------
 1/1/49A         L1      1        UP      UP       10         9        3       Pt-to-Pt
 1/1/50A         L1      2        UP      UP       10         9        3       Pt-to-Pt
 1/1/54A         L1      3        UP      UP       10         9        3       Pt-to-Pt

Interfaces : 3
```

In the "show spb isis interface" command output, three interfaces are SPB IS-IS enabled for L1 adjacencies. All three interfaces are both administratively and operationally up. By default, the link metric is 10 regardless of link speed. "Hello" messages are sent at nine second intervals, and adjacencies are declared lost if no "Hello" message is received for three consecutive intervals (for example; 27 seconds).

BEB-1

```
BEB-1> show spb isis nodes
SPB ISIS Nodes:
 System Name          System Id      SourceID BridgePriority
--------------------+--------------+--------+---------------
 BEB-2                dc08.5610.80f9  0x080f9  32768 (0x8000)
 BEB-4                dc08.5610.78d9  0x078d9  32768 (0x8000)
 BCB                  dc08.5610.7f19  0x07f19  32768 (0x8000)
 BEB-1                dc08.5610.8559  0x08559  32768 (0x8000)
 BEB-3                dc08.5610.7249  0x07249  32768 (0x8000)

Total SPB Nodes : 5
```

The "show spb isis nodes" command output indicates all discovered SPB IS-IS nodes including the local node. For each node, we can see the system or host name, the system ID (the BMAC) andthe source ID and the bridge priority. The source ID is a 20-bit identifier that designates the node as the origin of BUM traffic It is derived from the system ID's least significant bytes. The source ID is relevant when using tandem BUM replication. The bridge priority is 16-bit identifier and is used as a tie-breaker during path computation.

BEB-1

```
BEB-1> show spb isis adjacency
SPB ISIS Adjacency:
 System
 (Name : SystemId)                      Type   State   Hold  Interface
-------------------------------------+------+-------+------+----------
 BEB-2               : dc08.5610.80f9  L1    UP        20     1/1/49A
 BEB-4               : dc08.5610.78d9  L1    UP        20     1/1/50A
 BCB                 : dc08.5610.7f19  L1    UP        20     1/1/54A

Adjacencies : 3
```

The "show spb isis adjacency" command output indicates all SPB IS-IS adjacencies established by the local node. For each adjacency, we can see the system or host name, the system ID (the BMAC), the type (always L1 for SPB IS-IS), the state, the hold timer (number of seconds until the adjacency is declared lost if no "Hello" messages are received) and the interface over which the adjacency is formed.

BEB-1

```
BEB-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
  4001  00-80-c2-02       NO      NO            0  SGMODE
  4002  00-80-c2-03       NO      NO            0  SGMODE
  4003  00-80-c2-04       NO      NO            0  SGMODE

BVLANs:     4
```

The "show spb isis bvlans" command output indicates, for each configured BVLAN, the ECT algorithm in use and whether the BVLAN is in use and has services mapped to it. So far, we have not configured any service, therefore the only BVLAN in use is the control BVLAN, which is used for IS-IS CP messaging. We can also observe the number of ISIDs mapped to the BVLAN. For services using tandem BUM replication, we can observe whether this is (S,G), which is the default, or (*,G). Note that while the choice of head-end versus tandem replication is done on a per-service basis, the choice between (S,G) and (*,G) tandem replication is done on a per-BVLAN basis. Lastly, the root bridge BMAC is shown only for those BVLANs using (*,G) tandem replication.

BEB-1

```
BEB-1> show spb isis unicast-table
SPB ISIS Unicast MAC Table:
        Destination                                    Outbound
 BVLAN  (Name : MAC Address)                           Interface
------+--------------------------------------+-----------
  4000  BEB-2                : dc:08:56:10:80:f9    1/1/49A
  4000  BEB-4                : dc:08:56:10:78:d9    1/1/50A
  4000  BCB                  : dc:08:56:10:7f:19    1/1/54A
  4000  BEB-3                : dc:08:56:10:72:49    1/1/49A
  4001  BEB-2                : dc:08:56:10:80:f9    1/1/49A
  4001  BEB-4                : dc:08:56:10:78:d9    1/1/50A
  4001  BCB                  : dc:08:56:10:7f:19    1/1/54A
  4001  BEB-3                : dc:08:56:10:72:49    1/1/54A
  4002  BEB-2                : dc:08:56:10:80:f9    1/1/49A
  4002  BEB-4                : dc:08:56:10:78:d9    1/1/50A
  4002  BCB                  : dc:08:56:10:7f:19    1/1/54A
  4002  BEB-3                : dc:08:56:10:72:49    1/1/54A
  4003  BEB-2                : dc:08:56:10:80:f9    1/1/49A
  4003  BEB-4                : dc:08:56:10:78:d9    1/1/50A
  4003  BCB                  : dc:08:56:10:7f:19    1/1/54A
  4003  BEB-3                : dc:08:56:10:72:49    1/1/50A

MAC Addresses: 16
```

The “show spb isis unicast-table” command output indicates, for each node, the outbound interface used when sending unicast traffic to that node. Note that the outbound interface can be different for different BVLANs because different BVLANs can build different SPTs. For example, the path to BEB-3 goes through interface 1/1/49A in the case of BVLAN 4000, interface 1/1/54A in the case of BVLANs 40001 and 4002 and interface 1/1/50A in the case of BVLAN 4003.

BEB-1

```
BEB-1> show spb isis spf bvlan 4000
SPB ISIS Path Table:
 Destination                          Outbound   Next Hop                             SPB     Num
 (Name : BMAC)                        Interface  (Name : BMAC)                        Metric  Hops
--------------------------------------+----------+------------------------------------+------+------
 BEB-2              : dc:08:56:10:80:f9   1/1/49A  BEB-2              : dc:08:56:10:80:f9    10    1
 BEB-4              : dc:08:56:10:78:d9   1/1/50A  BEB-4              : dc:08:56:10:78:d9    10    1
 BCB                : dc:08:56:10:7f:19   1/1/54A  BCB                : dc:08:56:10:7f:19    10    1
 BEB-3              : dc:08:56:10:72:49   1/1/49A  BEB-2              : dc:08:56:10:80:f9    20    2

SPF Path count: 4

BEB-1> show spb isis spf bvlan 4001
SPB ISIS Path Table:
 Destination                          Outbound   Next Hop                             SPB     Num
 (Name : BMAC)                        Interface  (Name : BMAC)                        Metric  Hops
--------------------------------------+----------+------------------------------------+------+------
 BEB-2              : dc:08:56:10:80:f9   1/1/49A  BEB-2              : dc:08:56:10:80:f9    10    1
 BEB-4              : dc:08:56:10:78:d9   1/1/50A  BEB-4              : dc:08:56:10:78:d9    10    1
 BCB                : dc:08:56:10:7f:19   1/1/54A  BCB                : dc:08:56:10:7f:19    10    1
 BEB-3              : dc:08:56:10:72:49   1/1/54A  BCB                : dc:08:56:10:7f:19    20    2

SPF Path count: 4

BEB-1> show spb isis spf bvlan 4002
SPB ISIS Path Table:
 Destination                          Outbound   Next Hop                             SPB     Num
 (Name : BMAC)                        Interface  (Name : BMAC)                        Metric  Hops
--------------------------------------+----------+------------------------------------+------+------
 BEB-2              : dc:08:56:10:80:f9   1/1/49A  BEB-2              : dc:08:56:10:80:f9    10    1
 BEB-4              : dc:08:56:10:78:d9   1/1/50A  BEB-4              : dc:08:56:10:78:d9    10    1
 BCB                : dc:08:56:10:7f:19   1/1/54A  BCB                : dc:08:56:10:7f:19    10    1
 BEB-3              : dc:08:56:10:72:49   1/1/54A  BCB                : dc:08:56:10:7f:19    20    2

SPF Path count: 4

BEB-1> show spb isis spf bvlan 4003
SPB ISIS Path Table:
 Destination                          Outbound   Next Hop                             SPB     Num
 (Name : BMAC)                        Interface  (Name : BMAC)                        Metric  Hops
--------------------------------------+----------+------------------------------------+------+------
 BEB-2              : dc:08:56:10:80:f9   1/1/49A  BEB-2              : dc:08:56:10:80:f9    10    1
 BEB-4              : dc:08:56:10:78:d9   1/1/50A  BEB-4              : dc:08:56:10:78:d9    10    1
 BCB                : dc:08:56:10:7f:19   1/1/54A  BCB                : dc:08:56:10:7f:19    10    1
 BEB-3              : dc:08:56:10:72:49   1/1/50A  BEB-4              : dc:08:56:10:78:d9    20    2

SPF Path count: 4
```

The "show spb isis spb bvlan" command output indicates, for a given BVLAN, the outbound interface, the next hop node, the SPB metric and the total number of hops required to reach a destination node. Traffic destined towards BEB-3 will transit BCB-1 in the case of BVLAN 4000, BCB-4 in the case of BVLANs 4001 and 4002 and BCB-1 in the case of BVLAN 4003.

L2 services

A L2 service refers to a type of VPN service connecting multiple sites in a single any-to-any bridging domain. This section continues building on the previous example to create a L2 service on top of the previously created backbone configuration.

Services need only be created on BEBs, not on BCBs, and only on those BEBs where the service needs to be delivered. Creating an SPB service includes the following tasks:

- Creating a service and associating the service to an IS-IS and BVLAN – the specified BVLAN's SPF will be used for the service traffic
- Defining a Service Access Port (SAP)
- Defining SAPs matching specific customer traffic

Figure 19 - L2 service

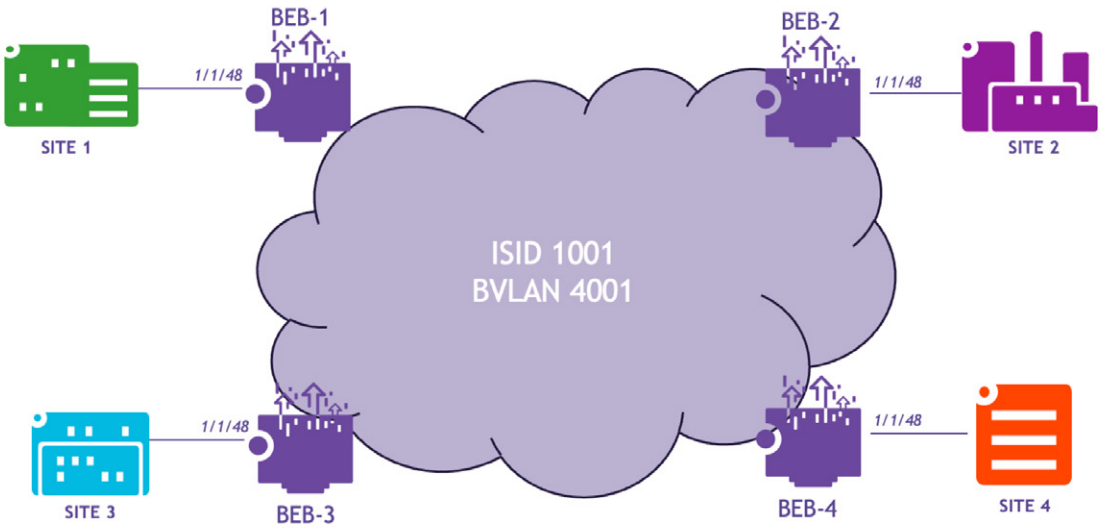

With regard to Figure 19, we provide BEB configurations in the snippet that follows. As well, note the following:

- The service number is only locally significant and can differ across different BEBs
- The ISID number is globally significant and must match across all BEBs connecting a given service
- The BVLAN that the service is mapped must also match across all BEBs connecting a given service
- Different services can be mapped to different BVLANs to achieve traffic load balancing

BEB-1 to BEB-4

```
BEB-1> service access port 1/1/48
BEB-1> service 1 spb isid 1001 bvlan 4001
BEB-1> service 1 sap port 1/1/48:0
```

The configuration snippet above indicates the following:

- Service 1 is associated to ISID 1001 and mapped to BVLAN 4001's SPF tree
- Port 1/1/48 is defined as a SAP
- A SAP is defined on port 1/1/48 mapping untagged traffic (:0) to service 1

BEB-1

```
BEB-1> show service spb
Legend: * denotes a dynamic object
SPB Service Info
  SystemId : dc08.5610.8559,   SrcId : 0x08559,    SystemName : BEB-1

                            SAP     Bind                    MCast
ServiceId   Adm  Oper Stats Count   Count   Isid      BVlan Mode      (T/R)
-----------+----+----+-----+-------+-------+---------+-----+--------------
1           Up   Up    N    1       3       1001      4001  Headend  (0/0)

Total Services: 1
```

The "show service spb" command output indicates, for a given BEB, the locally defined SPB services, their administrative and operational status, the number of (local) SAPs and (remote) SDPs along with the ISID and BVLAN number that the service is mapped to. It also shows the multicast replication mode, which is head-end by default. The multicast replication mode can be changed to tandem on a per-service basis.

BCB

```
BCB> show service spb
Legend: * denotes a dynamic object
SPB Service Info
  SystemId : dc08.5610.7f19,   SrcId : 0x07f19,    SystemName : BCB-1

                            SAP     Bind                    MCast
ServiceId   Adm  Oper Stats Count   Count   Isid      BVlan Mode      (T/R)
-----------+----+----+-----+-------+-------+---------+-----+--------------

Total Services: 0
```

The "show service spb" command output indicates that, by definition, a BCB does not have locally defined services.

BEB-1

```
BEB-1> show spb isis services
Legend: * indicates locally configured ISID
SPB ISIS Services Info:
                      System
    ISID      BVLAN   (Name : BMAC)                              MCAST(T/R)
------------+-------+-----------------------------------------+-----------
*     1001     4001   BEB-3               : dc:08:56:10:72:49
*     1001     4001   BEB-4               : dc:08:56:10:78:d9
*     1001     4001   BEB-2               : dc:08:56:10:80:f9
*     1001     4001   BEB-1               : dc:08:56:10:85:59

ISIDs:     4
```

The “show spb isis services” command output shows SPB services known to the node along with their ISID and BVLAN number and the node name, and BMACs that the service is enabled on. We should note that these services are learned thanks to the IS-IS CP. A “*” denotes that the service also matches a service locally created on the BEB.

BCB

```
BCB> show spb isis services
Legend: * indicates locally configured ISID
SPB ISIS Services Info:
                      System
    ISID      BVLAN    (Name : BMAC)                              MCAST(T/R)
------------+-------+-----------------------------------------+-----------
      1001      4001   BEB-3                 : dc:08:56:10:72:49
      1001      4001   BEB-4                 : dc:08:56:10:78:d9
      1001      4001   BEB-2                 : dc:08:56:10:80:f9
      1001      4001   BEB-1                 : dc:08:56:10:85:59

ISIDs:     4
```

The “show spb isis services” command output shows the same output from the perspective of a BCB. A BCB is still aware of all existing services with the IS-IS CP.

BEB-1

```
BEB-1> show service spb 1
SPB Service Detailed Info
  Service Id       : 1,                    Description      :                    ,
  ISID             : 1001,                 BVlan            : 4001,
  Multicast-Mode   : Headend,              Tx/Rx Bits       : 0/0,
  Admin Status     : Up,                   Oper Status      : Up,
  Stats Status     : No,                   Vlan Translation : No,
  Service Type     : SPB,                  Allocation Type  : Static,
  MTU              : 9194,                 VPN IP-MTU       : 1500,
  SAP Count        : 1,                    SDP Bind Count   : 3,
  RemoveIngressTag : No,                   Option           : None,
  Mgmt Change      : 02/01/2025 01:59:32,  Status Change    : 02/01/2025 01:59:32
```

The “show service spb” command output provides some additional details about a given SPB service, including the following:

- **RemoveIngressTag**: As explained in the Data Plane section, by default, a PBB frame includes all the frame’s original tags. However, we can remove those tags with the “service service_id remove-ingress-tag enable” command.
- **VLAN Translation**: This was explained in the Service framework section. VLAN translation must be enabled both at the service level with the command “service service_id vlan-translation enable” and on the SAP with the command “service access port vlan-xlation enable”.
- **Allocation Type**: Services can be either statically or dynamically created. Dynamic service creation will be explained in the Dynamic services section.

BEB-1

```
BEB-1> show service access
Legend: (~)Internal User Port Loopback (-)ERP Ring

Port       Link  SAP     SAP     Vlan
Id         Status Type    Count   Xlation L2Profile                       Description
---------+------+-------+-------+-------+-------------------------------+--------------------------------
1/1/48     Up     Manual  1       N       def-access-profile

Total Access Ports: 1
```

The "show service access" command output indicates, for a given BEB, the list of SAPs along with their type (manual or dynamic), the number of defined SAPs and whether VLAN translation is enabled or not. It shows the L2Profile assigned to the SAP. The L2Profile defines how L2 control protocol frames received on a SAP will be handled. Traffic can be peered, dropped or tunnelled. Default L2 profile settings are shown in the following table. Additional L2 profiles can be created with the command "service l2profile name stp action 802.1x action 802.3ad action mvrp action gvrp action amap action 802.1ab action" and assigned to the SAP with the command "service access l2profile name".

Protocol	def-access-profile	unp-def-access-profile
STP	tunnel	drop
802.1X	drop	peer
802.3ad	peer	peer
MVRP	tunnel	tunnel
GVRP	tunnel	tunnel
AMAP	drop	drop
802.1ab	drop	drop

BEB-1

```
BEB-1> show service spb 1 ports
Legend: (*)Dyn Unicast (+)Remote Mcast (#)Local Mcast (~)Internal User Port Loopback (-)ERP Ring
SPB Service 1 Info
  Admin : Up,        Oper  : Up,     Stats      : N,          Mtu     : 9194,   VlanXlation : N,
  ISID  : 1001,      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:1/1/48:0           Up   Up    N           Y:x           1/1/48       -
sdp:32769:1*           Up   Up    Y    dc08.5610.80f9:4001  1/1/49A  BEB-2
sdp:32773:1*           Up   Up    Y    dc08.5610.7249:4001  1/1/54A  BEB-3
sdp:32781:1*           Up   Up    Y    dc08.5610.78d9:4001  1/1/50A  BEB-4

Total Ports: 4
```

The "show service spb ports" command output indicates local (SAP) and remote (SDP) ports for a given service. For each port, we can see administrative and operational status, the system ID (BMAC) and BVLAN, as well as the system name and associated local interface. SDP ports will always display a "*" next to them because SDP ports are always dynamically created by the IS-IS CP. The name of an SDP is a combination of a dynamically generated number, followed by a colon and the service number.

BEB-1

```
BEB-1> show service mesh-sdp spb
Legend: * denotes a dynamic object
SPB Mesh-SDP Info
SvcId    SdpId             Isid      FarEnd SysId:BVlan   Oper Intf     FarEnd SystemName
--------+-----------------+---------+--------------------+----+--------+------------------
1        32769:1*          1001      dc08.5610.80f9:4001  Up   1/1/49A  BEB-2
1        32773:1*          1001      dc08.5610.7249:4001  Up   1/1/54A  BEB-3
1        32781:1*          1001      dc08.5610.78d9:4001  Up   1/1/50A  BEB-4

Total Mesh-SDPs: 3
```

The “show service mesh-sdp spb” command output indicates far-end SDPs for each service along with the ISID number and the far-end system ID (BMAC), BVLAN, system name and associated interface.

BEB-1

```
BEB-1> show mac-learning domain spb
Legend: Mac Address: * = address not valid,

        Mac Address: & = duplicate static address,

   Domain     Vlan/SrvcId[ISId/vnId]     Mac Address           Type          Operation          Interface
------------+----------------------+------------------+------------------+-------------+-------------------------
       SPB                   1:1001   00:50:56:85:98:df            dynamic    servicing               sap:1/1/48
       SPB                   1:1001   00:50:56:85:d9:de            dynamic    servicing              sdp:32769:1
       SPB                   1:1001   00:50:56:85:27:09            dynamic    servicing              sdp:32773:1
       SPB                   1:1001   00:50:56:85:4c:a4            dynamic    servicing              sdp:32781:1

 Total number of Valid MAC addresses above = 4
```

The “show mac-learning domain spb” command output indicates the list of CMAC addresses learned in the SPB domain along with the service number and ISID, as well as the interface (SAP or SDP) port that the CMAC address is bound to.

BCB

```
BCB> show mac-learning domain spb
Legend: Mac Address: * = address not valid,

        Mac Address: & = duplicate static address,

   Domain     Vlan/SrvcId[ISId/vnId]     Mac Address           Type          Operation          Interface
------------+----------------------+------------------+------------------+-------------+-------------------------

 Total number of Valid MAC addresses above = 0
```

The “show mac-learning domain spb” command output shows the same output from the point of view of a BCB node. As expected, BCB nodes do not learn any CMACs.

L3 services

A L3 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 customers isolated at L3, each customer service is associated to its own VRF instance.

Figure 20 - Customer A’s L3 service

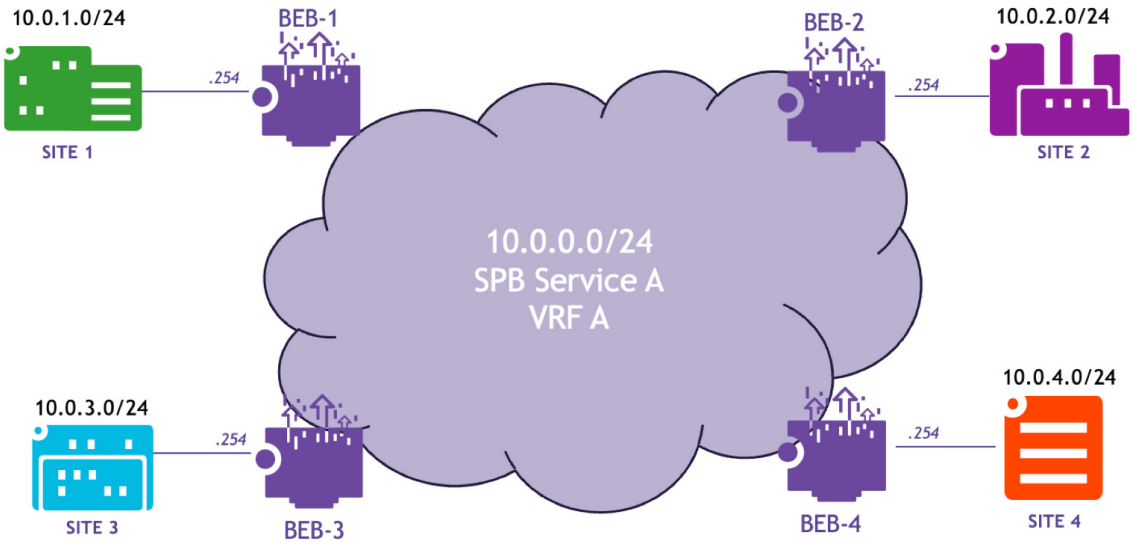

Figure 20 illustrates an example of a L3 service connecting four of customer A's sites (Sites 1 through 4). Each site uses a different subnet and therefore, inter-site routing is required. BEB nodes connecting customer sites are represented with router icons for simplicity. These BEBs have a "LAN"-facing interface which acts as the local site default gateway, as well as a "WAN"-facing interface to reach remote sites. All "WAN" interfaces are bound to a single SPB service and are on the same "WAN" subnet. Last, all the LAN and WAN IP interfaces associated to customer A are bound to the same customer A VRF to provide L3 isolation between different customers.

SPB-based L3 VPN services rely on edge routing: routing is only performed at ingress and egress BEBs and bridged between these. At L3, the WAN represents a single L3 hop regardless of the number of intermediate L2 hops (BCBs) in between. SPB simply bridges traffic from ingress BEB to egress BEB along the shortest path.

Up to this point, we have only described the DP. What about the CP? At the CP level, L3 VPN services come in two variants: VPN Lite and L3 VPN.

VPN Lite

A VPN Lite L3 service is created by overlaying a L3 routing protocol on top of the L2 WAN SPB service. This routing protocol can be OSPF, BGP or even static routing. The routing protocol runs inside the customer's VRF, and a separate instance and associated configuration is created for each customer. Figure 21 shows an example of how customer A's L3 service can be created as a VPN Lite service by running OSPF on BEB nodes.

Figure 21 - Customer A's VPN Lite service

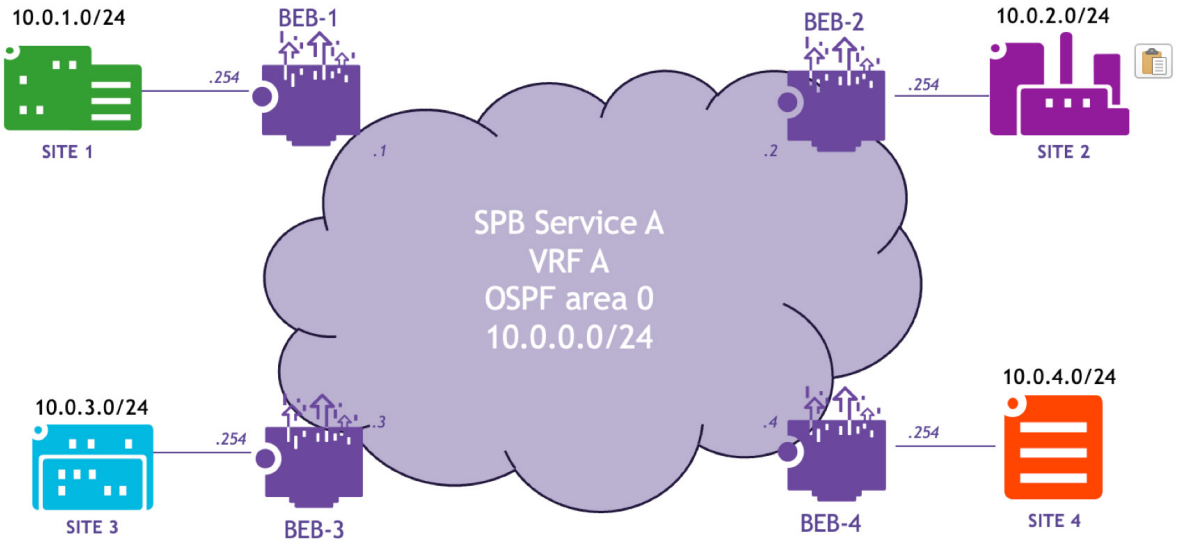

In a VPN Lite L3 service, the L2 SPB service simply provides L2 connectivity to the "WAN" IP interfaces. Continuing with OSPF as an example, this means that OSPF is configured as usual. Also, since all WAN IP interfaces are connected to a single L2 SPB service, in the case of OSPF, a DR/BDR election will take place as usual.

L3 VPN

SPB L3 VPN leverages the existing SPB IS-IS instance to carry customer VPN routes without requiring an additional routing protocol such as OSPF. This is accomplished with additional IS-IS TLVs extensions. Each customer or tenant is still associated to its own VRF and IS-IS TLVs reference the customer's ISID to preserve L3 isolation between different customers or tenants. This mechanism is described in an IETF draft [1] and shown in Figure 22.

MPLS and EVPN rely on an IGP (for example, OSPF or IS-IS) for backbone node reachability and MP-BGP (RFC 4760) for customer VPN route transport. In SPB L3 VPN, IS-IS can provide both backbone node reachability and customer VPN route transport. Using a single protocol instead of two results in a network that is simpler to deploy and operate.

In addition, when comparing SPB and MPLS, SPB BEB nodes play a role similar to MPLS PE nodes while SPB BCB nodes are similar to MPLS P nodes. In particular, SPB BCB nodes do not learn any customer VPN routes and require no VRFs to be created on them. VRFs need only be created on BEB nodes and customer VPN routes are only learned on the BEBs that those customers connect to.

Figure 22 - Customer A's L3 VPN service

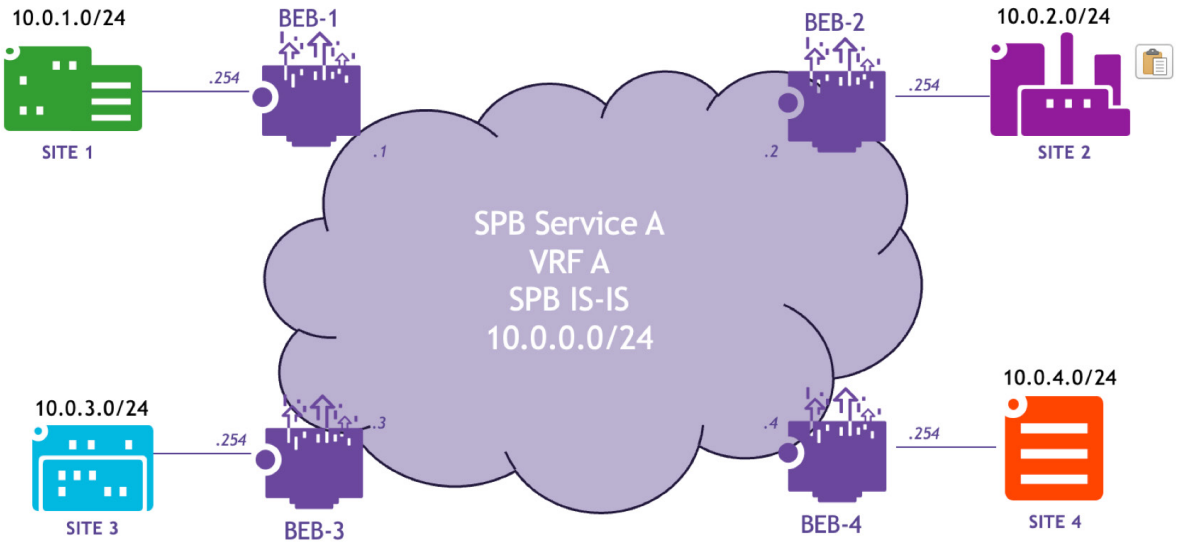

Unlike the case of a VPN Lite service, an SPB L3 VPN does not require the addition of any routing protocol. The customer's VRF routes are exported to the SPB IS-IS instance, associated to the customer's ISID and bound to the WAN IP as a gateway address. Far-end BEBs will import those routes into their local VRF routing table. Therefore, those routes will point to the WAN IP address as next-hop. This mechanism is applicable and identical for both IPv4 and IPv6. Figure 23 shows this from the perspective of BEB-1. Route maps can be used for fine-grained route filtering.

Figure 23 - Route import/export

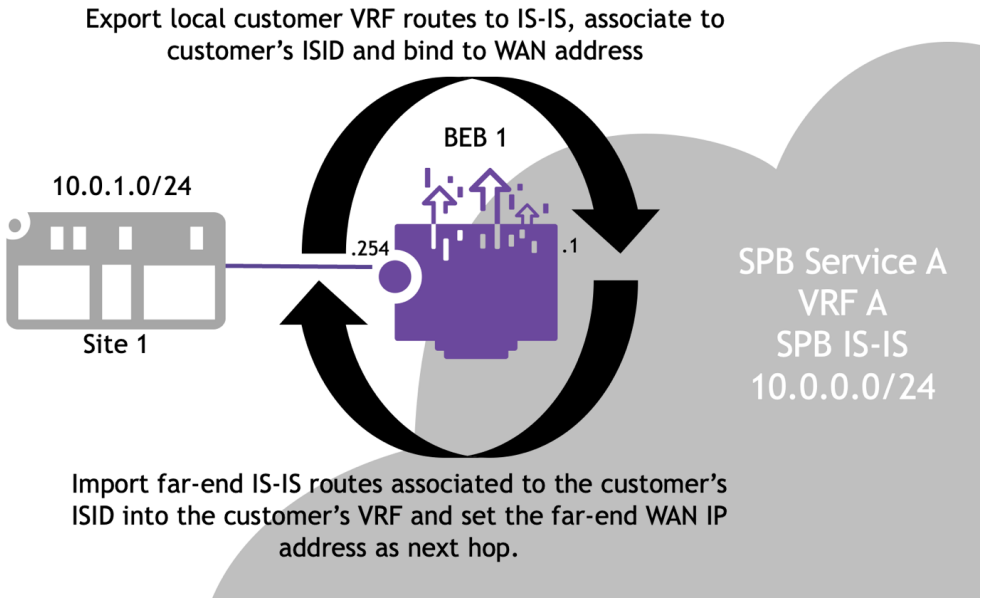

An L3 VPN service builds upon a L2 service and involves the following steps:

- Creating an L2 SPB service
- Creating a tenant VRF
- Creating LAN-side and WAN-side IP interfaces on the tenant VRF. LAN-side IP interfaces normally reside on a VLAN. WAN-side IP interfaces can reside directly on the SPB services itself.
- Route import/export between local VRF routing table and SPB IS-IS ISID instance
- Binding the WAN IP interface to the L2 SPB service's ISID

Figure 24 shows a sample topology on which we will configure an L3 VPN service.

Figure 24 - L3 VPN service example

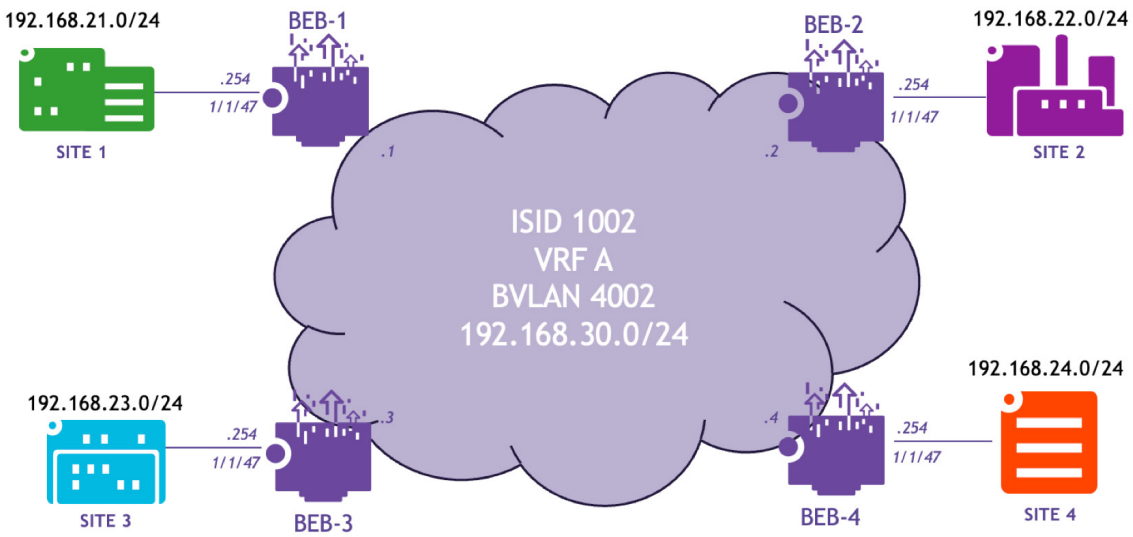

Like their L2 counterpart, L3 VPN services require no configuration on BCBs.

- Customer sites connect to their local BEB though interface 1/1/47
- LAN-side, or site default-gateway IP interfaces are bound to VLAN 3001, which is the default VLAN on port 1/1/47
- WAN-side IP interfaces are bound directly to the L2 SPB service's ISID

BEB-1

```
BEB-1> vlan 3001 name "Site_1"
BEB-1> vlan 3001 members port 1/1/47 untagged
BEB-1> service 2 spb isid 1002 bvlan 4002
BEB-1> vrf create Customer_A
Customer_A: :BEB-1> ip interface "LAN" address 192.168.21.254 mask 255.255.255.0 vlan 3001
Customer_A: :BEB-1> ip interface "WAN" address 192.168.30.1 mask 255.255.255.0 service 2
Customer_A: :BEB-1> ip export all-routes
Customer_A: :BEB-1> ip import isid 1002 all-routes
BEB-1> spb ipvpn bind vrf Customer_A isid 1002 gateway 192.168.30.1 all-routes
```

BEB-2

```
BEB-2> vlan 3001 name "Site_2"
BEB-2> vlan 3001 members port 1/1/47 untagged
BEB-2> service 2 spb isid 1002 bvlan 4002
BEB-2> vrf create Customer_A
Customer_A: :BEB-2> ip interface "LAN" address 192.168.22.254 mask 255.255.255.0 vlan 3001
Customer_A: :BEB-2> ip interface "WAN" address 192.168.30.2 mask 255.255.255.0 service 2
Customer_A: :BEB-2> ip export all-routes
Customer_A: :BEB-2> ip import isid 1002 all-routes
BEB-2> spb ipvpn bind vrf Customer_A isid 1002 gateway 192.168.30.2 all-routes
```

BEB-3

```
BEB-3> vlan 3001 name "Site_3"
BEB-3> vlan 3001 members port 1/1/47 untagged
BEB-3> service 2 spb isid 1002 bvlan 4002
BEB-3> vrf create Customer_A
Customer_A: :BEB-3> ip interface "LAN" address 192.168.23.254 mask 255.255.255.0 vlan 3001
Customer_A: :BEB-3> ip interface "WAN" address 192.168.30.3 mask 255.255.255.0 service 2
Customer_A: :BEB-3> ip export all-routes
Customer_A: :BEB-3> ip import isid 1002 all-routes
BEB-3> spb ipvpn bind vrf Customer_A isid 1002 gateway 192.168.30.3 all-routes
```

BEB-4

```
BEB-4> vlan 3001 name "Site_4"
BEB-4> vlan 3001 members port 1/1/47 untagged
BEB-4> service 2 spb isid 1002 bvlan 4002
BEB-4> vrf create Customer_A
Customer_A: :BEB-4> ip interface "LAN" address 192.168.24.254 mask 255.255.255.0 vlan 3001
Customer_A: :BEB-4> ip interface "WAN" address 192.168.30.4 mask 255.255.255.0 service 2
Customer_A: :BEB-4> ip export all-routes
Customer_A: :BEB-4> ip import isid 1002 all-routes
BEB-4> spb ipvpn bind vrf Customer_A isid 1002 gateway 192.168.30.4 all-routes
```

Having created the L3 VPN service on all nodes, we can now proceed to verify it with show commands. Let's start by verifying correct route import and export. The below snippet shows routes in BEB-1's VRF "Customer_A". Both local LAN and WAN subnets are LOCAL routes while far-end LAN subnets are IMPORT routes whose next-hop gateway address is the WAN address of the remote BEB.

BEB-1

```
Customer_A::BEB-1> show ip routes

 + = Equal cost multipath routes
 Total 6 routes

  Dest Address        Gateway Addr         Age         Protocol
------------------+-------------------+----------+-----------
  127.0.0.1/32          127.0.0.1          01:34:43    LOCAL
  192.168.21.0/24       192.168.21.254     01:29:43    LOCAL
  192.168.22.0/24       192.168.30.2       01:26:07    IMPORT
  192.168.23.0/24       192.168.30.3       01:25:37    IMPORT
  192.168.24.0/24       192.168.30.4       01:25:11    IMPORT
  192.168.30.0/24       192.168.30.1       01:27:13    LOCAL
```

The below snippet shows arp entries in BEB-1's VRF "Customer_A". Far-end WAN gateway addresses are dynamically learned.

BEB-1

```
Customer_A::BEB-1> show arp

Total 3 arp entries
 Flags (P=Proxy, A=Authentication, V=VRRP, B=BFD, H=HAVLAN, I=INTF, M=Managed)

 IP Addr           Hardware Addr       Type       Flags   Port              Interface   Name
-----------------+-------------------+----------+-------+-----------------+-----------+-------
--------------------------
 192.168.30.2      dc:08:56:10:80:f9   DYNAMIC                    1/1/47   WAN
 192.168.30.3      dc:08:56:10:72:49   DYNAMIC                    1/1/47   WAN
 192.168.30.4      dc:08:56:10:78:d9   DYNAMIC                    1/1/47   WAN
```

In addition to these L3-related verification steps, all steps covered in the L2 services section can be used to verify the underlying L2 service.

Choosing VPN Lite or L3 VPN

L3 VPN offers several advantages:

- **Simplicity**: L3 VPN does not require routing protocol configuration as it simply leverages the existing SPB IS-IS instance. VPN Lite on the other hand requires one routing protocol instance per tenant/VRF and BEB. For example, if using OSPF, 4 customer services spanning 8 BEB nodes require 4 x OSPF instances per node for a total of 32 x OSPF configurations across all nodes. If dual-stack IPv4 and IPv6 support is required, this translates to an OSPFv2 and an OSPFv3 instance per BEB and VRF for a total of 64 x OSPF configurations, all nodes included. More routing protocol configurations result in longer service provisioning times and increased chances of making mistakes.
- **Scalability**: L3 VPN is significantly more efficient than VPN Lite from a CP point of view as it uses a single routing instance. This results in lighter CP load and allows for greater scalability than VPN Lite.
- **Convergence**: L3 VPN convergence can be faster than VPN Lite because it relies on a single protocol. VPN Lite convergence can be slower because the stacking of routing protocols has a compounding effect over convergence time: IS-IS must converge before OSPF can converge.

While L3 VPN is the recommended option within the SPB domain, L3 VPN relies on SPB IS-IS and cannot directly interoperate with external networks. VPN Lite, on the other hand, can be configured on border BEB nodes linking the SPB domain to external, non-SPB capable networks. These border BEB nodes use L3 VPN to communicate with other BEB nodes and VPN Lite to interoperate with external non-SPB nodes through common routing protocols such as OSPF or BGPv4.

For these reasons, L3 VPN is recommended within the SPB domain and VPN Lite is needed only on border nodes connecting to the outside world.

Shared services VPN and route leaking

In L3 VPN designs in which each VPN maps to its own VRF, it is common for certain services such as DHCP, DNS and Internet access to be shared across two or more of those VPNs. This can be implemented through VRF leaking.

Figure 25 shows the same diagram as before, but now with two customers, A and B. Each customer is associated to its own ISID (1002 for Customer A and 1003 for Customer B) and VRF (Customer_A and Customer_B) on BEBs 1 through 4. Routes are propagated across the backbone as explained in the L3 VPN section.

Figure 25 - Shared services

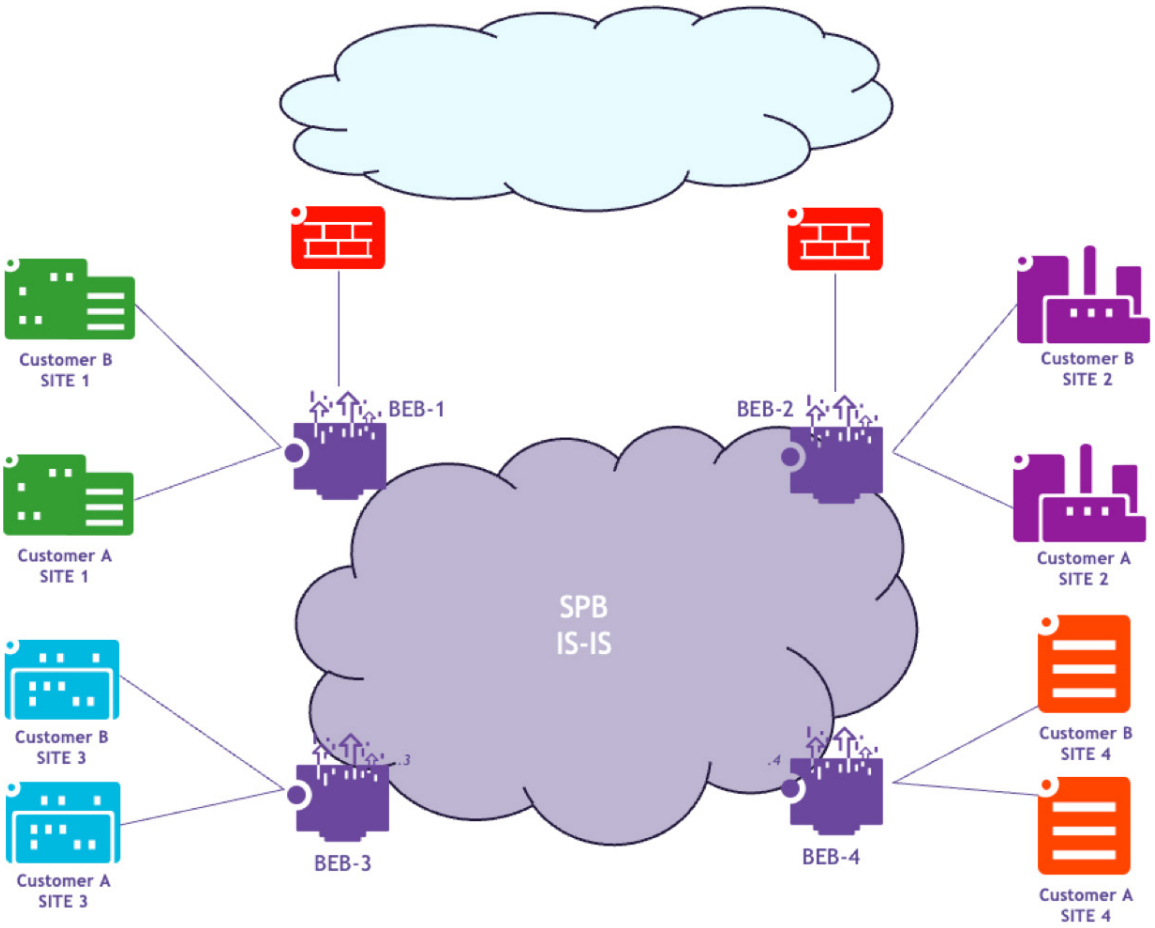

Let's now imagine that these customers need to also access some shared services and Internet access. An additional L3_VPN is created on BEB1 and BEB2, the "border" BEBs. These are the nodes that those shared services are accessed through. The "shared_services" L3VPN is associated to its own ISID (1004) and VRF (shared_services). Note that this L3VPN need not be stretched to BEBs 3 and 4.

BEB1 and BEB2 can exchange routes with external entities, such as the firewalls, using a standard protocol, such as BGP4. Those routes can be leaked to customer A's and B's VRFs. In turn, customer A's and B's VRF routes can be leaked to the "shared_services" VRF. As a pre- requisite, customer A's and B's address space must not overlap with each other or with the shared services.

BEB-1

```
! Export shared services routes to global IP routing table
vrf shared_services ip export route-map ebgp_routes_only

! Import shared services routes from global IP routing table into customer VRFs
vrf Customer_A ip import vrf shared_services all-routes
vrf Customer_B ip import vrf shared_services all-routes

! Import local customer VRF routes from global IP routing table into the shared services VRF
vrf shared_services ip import vrf Customer_A all-routes
vrf shared_services ip import vrf Customer_B all-routes

! Import remote customer routes from SPB ISIS instance into the shared services VRF
vrf shared_services ip import isid 1002 all-routes
vrf shared_services ip import isid 1003 all-routes

! Redistribute shared services routes to remote Customer sites through SPB ISIS Instance
spb ipvpn redist source-vrf shared_services destination-isid 1002 all-routes
spb ipvpn redist source-vrf shared_services destination-isid 1003 all-routes
```

The above snippet provides the commands required to accomplish this on the border BEBs, BEB1 and BEB2.

In summary:

- Shared routes are exported from the shared_services VRF and into the Global Routing Table (GRT). When doing so, a route-map filters routes such that only external routes are exported. This is to prevent re-export of routes imported from the other border BEB.
- Shared routes are imported from the GRT and into the customer VRFs. This step is necessary only if customer sites are connected to the border BEBs.
- Customer routes are imported from the GRT and into the shared_services VRF. This step is necessary only if customer sites are connected to the border BEBs.
- Remote customer routes are imported from the SPB IS-IS instance and ISID associated to those customers and into the shared_services VRF.
- Shared routes are redistributed from the shared_services VRF to the SPB IS-IS instance and ISIDs associated to customers. These routes will then be propagated across the backbone and imported into customer VRFs at remote BEBs.

SPB service types

So far, the service type discussed includes providing customers at various sites with a multipoint-to-multipoint (E-LAN) connectivity solution using SPB; however, various other service types are supported. These service types are defined by MEF Forum, based on their architecture, and include:

- E-Line
- E-LAN
- E-Tree

These service types are shown in Figure 26 below.

Figure 26 - MEF Service Types

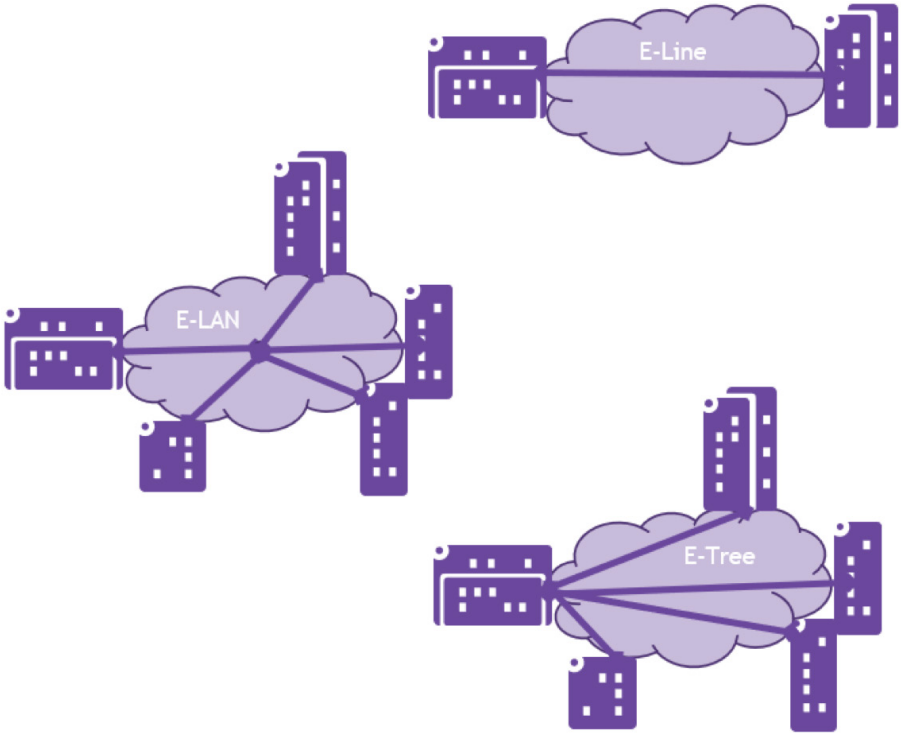

E-LAN

An E-LAN is a multipoint-to-multipoint service that connects two or more UNIs and provides full mesh connectivity for those sites. It emulates a bridged Ethernet network topology connecting the different sites. They are used to create Layer 2 VPN services.

This is the default service type.

E-Line

With E-LAN service type, as the number of users increases, so does the number of MAC addresses that are dynamically learned to minimize the amount of traffic flooding in the network.

To help reduce the use of system resources and prevent MAC address explosion, an SPB service can be configured as a pseudo-wire type of service to provide a single point-to-point (E-Line) connection between two SAP attachment points. The attachment points to customer edge (CE) devices can be between two local SAPs associated with the same service or between two SAPs across the SPB network, as shown in Figure 27.

Figure 27 - SPB pseudo-wire service

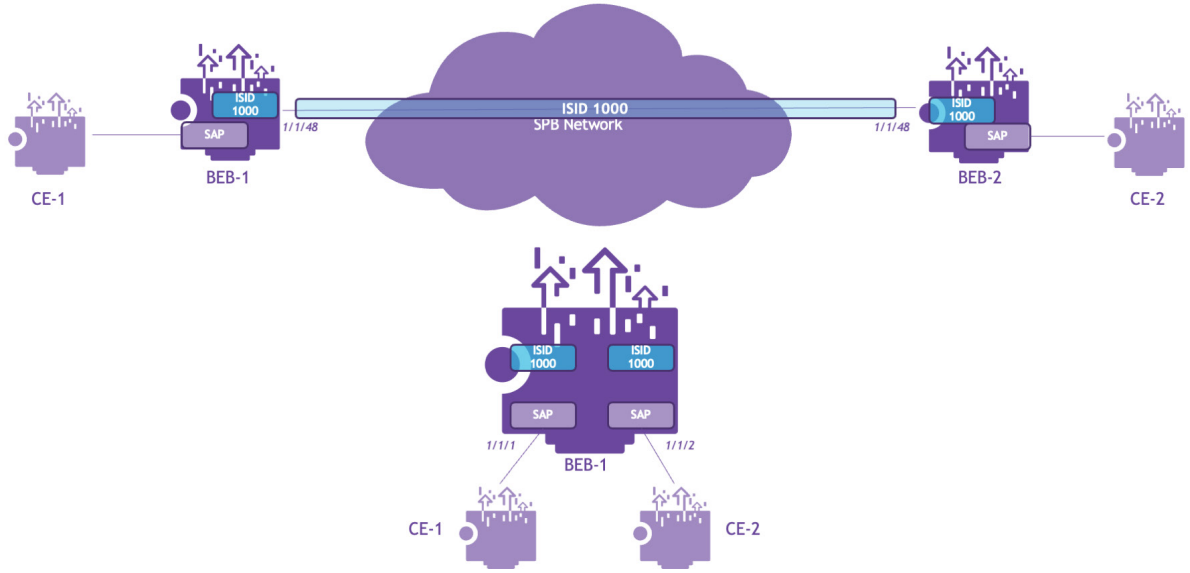

An SPB pseudo-wire service reduces the number of customer MAC addresses learned and simplifies the flow of user traffic, as follows:

- Packets are transparently forwarded (MAC addresses are not learned) between two local SAPs.
- Packets from one SAP are encapsulated and transparently forwarded out of the SPB service network port.
- The encapsulation is removed from packets received from the SPB network port and are transparently forwarded out of the SAP.
- MAC address learning is automatically disabled, since there is no forwarding decision to be made. Packets entering one SAP attachment point will simply egress the other SAP attachment point of the pseudo-wire unchanged.
- Flooding and replication is not necessary since only two virtual ports are involved (SAP to SAP for a local service or SAP to a network port for the same service instance across the SPB backbone).

To configure an SPB pseudo-wire connection, identify which SPB nodes will serve as a pseudo-wire endpoint and enable the pseudo-wire keyword.

BEB-1/BEB-2

```
BEB-1> service 100 spb isid 1000 bvlan 4001 pseudo-wire enable
```

Additional guidelines must be considered to when configuring E-Line services. For further details, refer to the OmniSwitch AOS Network Configuration User Guide referenced in the Related documents section.

E-Tree

E-Tree is a rooted multipoint service similar to a hub-and-spoke service. The hub is called the “root” and the spokes are called “leaves”. A Leaf UNI can only communicate with a Root UNI, while a Root UNI can communicate with any other Leaf UNI.

Primary use cases for E-Tree service type are security and scalability, for example, in a service provider environment where shared business services can be delivered to customers while providing customer isolation. In addition, scalability can be increased as tunnels are not created between BEBs that only support Leaf SAPs for a service.

In an E-Tree service, SAPs are designated as either “Leaf” SAP or “Root” SAP. A leaf SAP cannot communicate with another Leaf SAP in the service spanning multiple BEBs whereas Leaf SAP to Root SAP traffic is allowed. Root SAPs can communicate to all Leaf SAPs and Root SAPs.

Figure 28 - SPB E-Tree service

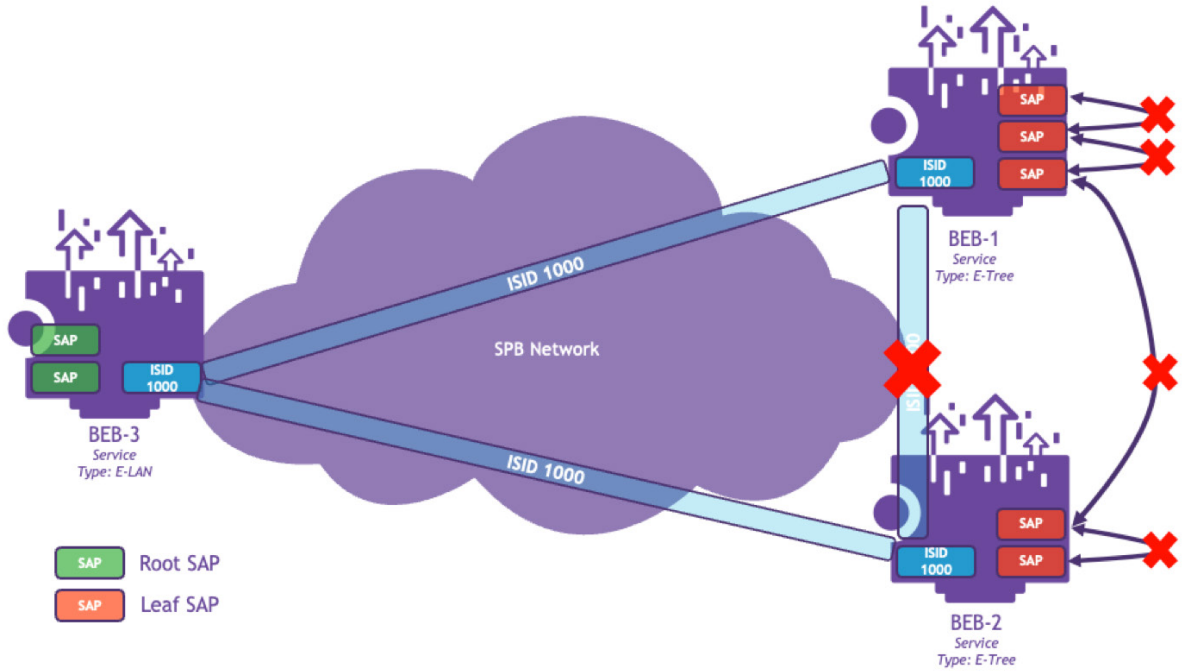

Figure 28 shows an example of E-Tree deployment. On all BEBs, a service with ISID 1000 is created. On BEB-3 the service is E-LAN, and on BEBs 1 and 2 the service is E-Tree.

All SAPs associated with the service with ISID 1000 on BEBs 1 and 2 are Leaf SAPs and all SAPs associated with the service with ISID 1000 on BEB-3 are Root SAPs.

The traffic that ingress on the Leaf SAPs is allowed to egress only on the Root SAPs. For this reason, a tunnel for the ISID 1000 is not created between the BEBs 1 and 2. Tunnels between BEB-3 and BEB-1; BEB-3 and BEB-2 are created. Also, traffic flow between the Leaf SAPs of a BEB is not allowed.

BEB-1/2

```
BEB-1> service 10 spb isid 1000 bvlan 4000 e-tree enable
```

BEB-3

```
BEB-3> service 10 spb isid 1000 bvlan 4000
```

The E-Tree feature can be enabled on UNP created services when creating the profile:

BEB-1/2

```
BEB-1> unp profile "BEB1-SAP1" map service-type spb tag-value 10 isid 1000 bvlan 4000 e-tree
```

BEB-3

```
BEB-3> unp profile "BEB3-SAP1" map service-type spb tag-value 10 isid 1000 bvlan 4000
```

SD-LAN

SD-LAN (Software-Defined LAN) is a modern LAN architecture that applies the principles of Software-Defined Networking (SDN) to wired and wireless networks, which includes:

- Policy-driven architecture
- Identity-driven architecture
- Dynamic segmentation
- Centralized management

SD-LAN simplifies operations, improves security, reduces configuration errors and provides a consistent user experience compared to traditional campus networks.

SPB along with the Access Guardian (AG) framework can provide SD-LAN services with dynamic services, micro-segmentation, IoT fingerprinting, role-based access control (RBAC) and authentication features. Centralized management can be done using the Alcatel-Lucent OmniVista® Network Management System (NMS).

Devices and users can be mapped to their UNP profile, which includes the required security policies, in different ways:

- Static: Port, VLAN
- IoT device profiling: MAC OUI, DHCP options, HTTPS agent
- Authentication: MAC authentication, 802.1X authentication

This Design Guide has already covered static configuration. The Automation section will describe dynamic SAPs and dynamic services in conjunction with authentication which allows for dynamic mapping of users and devices to the service.

IoT device profiling can be creating locally on the switch, but the best practice is to centralize the configuration on OmniVista NMS.

Automation

Dynamic SAPs

Previously in this Design Guide, we have shown how to configure SAPs statically and manually; however, SAPs can be automatically and dynamically configured using the Universal Network Profile (UNP) feature in conjunction with authentication (802.1X, MAC) or classification rules (for example, VLAN tag).

Dynamically-created SAPs can map traffic to a manually created service. Dynamically-created SAPs can also map traffic to a dynamically-created service for a fully dynamic configuration, which is covered in the next section.

The snippet shown in this section uses a sample configuration. This example refers to the case of L2 services in which any required routing, such as default gateway, DHCP relay, is performed on a central node, which can be a switch or a firewall. Either way, service and SAP configuration on the central L3 device is static. Dynamic configuration is useful at the edge nodes where client devices are added, moved and changed on a regular basis.

Six UNP profiles named “EMPLOYEE”, “IoT”, “GUEST”, “WLAN”, “CCTV” and “RESTRICTED” are created, each mapping to a different ISID. There are a total of four BVLANs, 4000 through 4003. BVLAN 4000 is reserved as a control BVLAN, so services can be mapped to BVLANs 4001 through 4003. As a result, each BVLAN carries traffic for two different services. These UNP profiles use head-end replication and have VLAN translation enabled; these are default behaviours which are explained elsewhere in this Design Guide.

So far, this describes the services but does not describe how ports or client devices will be mapped to those services. This mapping can be either static or dynamic. In the dynamic case, ports 1/1/10 through 1/1/16

are defined as UNP "access" ports. This means that they map traffic to an SPB service, as opposed to a UNP "bridge" port which maps traffic to a VLAN. These ports utilize the "SAMPLE_FLOW" port template. This template is defined as follows:

- 802.1X supplicants are authenticated against the "UPAM" radius server. If successful, the RADIUS server returns a "Filter-ID" attribute which matches one of the locally defined UNPs (for example; EMPLOYEE, IoT, among others).
- As a fall-back mechanism for non-802.1X capable devices, such devices can use MAC authentication. If successful, the RADIUS server also returns a "Filter-ID" attribute which matches one of the locally defined UNPs (for example; EMPLOYEE, IoT, among others).
- In both 802.1X or MAC authentication cases, it may happen that the RADIUS server does not return a "Filter-ID" or that the returned "Filter-ID" value does not match any of the locally defined UNPs. In such case, those devices are bound to a "RESTRICTED" UNP.
- The RESTRICTED UNP is also defined as the default UNP which is used in case of authentication failure. When bound to this RESTRICTED UNP, devices will receive an IP address through DHCP but will be very limited in their access to network resources. This is controlled at the central L3 node or firewall. This allows for these devices to have minimal network connectivity such that they can be onboarded (for example a digital certificate can be applied), and they can successfully authenticate next time they connect.

With this configuration in place, devices connected to ports 1/1/10 through 1/1/16 will be authenticated and dynamically bound to an SPB service according to their type or user identity. This means that the SPB service will automatically adapt and change as devices connect, disconnect, move or otherwise change without manual intervention.

In some cases, it may be necessary to statically bind these UNP services to a port. This is particularly useful if authentication is not used or when the device is a "silent" device (for example, a device that does not transmit traffic for extended periods of time because it goes into power-save mode. These periods of inactivity can result in a loss of service binding, making the device effectively unreachable (for example, for a WAKE-ON-LAN packet). This problem can be avoided by statically binding the UNP profile to the port. We have applied static UNP binding to ports 1/1/5 through 1/1/9 such that the service is statically bound to those ports even if the device disconnects or stops communicating for extended periods of time.

While statically binding a SAP, as opposed to a UNP, also offers a solution to the silent device problem, statically binding a UNP instead of a SAP allows the exact same UNP constructs to be used for both silent and non-silent devices. This results in a more standardized configuration which is easier to create and maintain with fewer mistakes when configurations need to change and is considered a best practice.

BEB-1

```
BEB-1> unp profile "EMPLOYEE"
BEB-1> unp profile "IoT"
BEB-1> unp profile "GUEST"
BEB-1> unp profile "WLAN"
BEB-1> unp profile "CCTV"
BEB-1> unp profile "RESTRICTED"
BEB-1> unp profile "EMPLOYEE" map service-type spb tag-value 0 isid 1001 bvlan 4001
BEB-1> unp profile "IoT" map service-type spb tag-value 0 isid 1002 bvlan 4002
BEB-1> unp profile "GUEST" map service-type spb tag-value 0 isid 1003 bvlan 4003
BEB-1> unp profile "WLAN" map service-type spb tag-value 0 isid 1004 bvlan 4001
BEB-1> unp profile "CCTV" map service-type spb tag-value 0 isid 1005 bvlan 4002
BEB-1> unp profile "RESTRICTED" map service-type spb tag-value 0 isid 1006 bvlan 4003
BEB-1> unp port-template SAMPLE_FLOW direction both aaa-profile "UPAM" default profile
"RESTRICTED" admin-state enable
BEB-1> unp port-template SAMPLE_FLOW 802.1x-authentication
BEB-1> unp port-template SAMPLE_FLOW 802.1x-authentication pass-alternate "RESTRICTED"
BEB-1> unp port-template SAMPLE_FLOW mac-authentication
BEB-1> unp port-template SAMPLE_FLOW mac-authentication pass-alternate "RESTRICTED"
BEB-1> unp port 1/1/5-16 port-type access
BEB-1> unp port 1/1/5-16 admin-state enable
BEB-1> unp port 1/1/5 profile "CCTV"
BEB-1> unp port 1/1/6 profile "WLAN"
BEB-1> unp port 1/1/7 profile "EMPLOYEE"
BEB-1> unp port 1/1/8 profile "GUEST"
BEB-1> unp port 1/1/9 profile "IoT"
BEB-1> unp port 1/1/10-16 port-template SAMPLE_FLOW
```

Dynamic services

In the previous section, we explained how SAPs can be dynamically configured to accommodate mobile users and devices and highly dynamic environments. This same mechanism is applicable to VMs in a data center. As VMs are created, turned on or off or migrated from one hypervisor to another, SAPs can be automatically and dynamically created to adapt to those events on the fly without network manager intervention.

For instance, classification rules can match VM traffic based on the VLAN tag (configured in the hypervisor) and create the required SAPs dynamically and automatically. This is a best practice compared to statically enabling all possible SAPs on all access ports because it reduces the broadcast domain footprint to only the required ports, thus eliminating unnecessary broadcast traffic and MAC learning.

However, with the features that we have described so far, even if the SAPs can dynamically adapt, the service UNP would need to be manually created. In certain scenarios, the network administrator does not know the required parameters beforehand. For instance, the server manager may create, change and delete VLANs on the hypervisor's vSwitch on a regular basis. It may be tempting to pre-provision services for all 4096 VLANs, but this practice creates an unnecessary load on the control plane.

The best practice for that type of environment is to use AOS' Dynamic Services feature. With Dynamic Services, UNPs can be dynamically created, on the fly, based on the VLAN tag seen on UNP ports.

BEB-1

```
BEB-1> unp port 10-16 dynamic-service spb
```

The maximum number of UNP users that can be created using this feature is hardware dependant. Please refer to the OmniSwitch AOS Specifications Guide referenced in the Related documents section.

Upon receiving a frame on a UNP access port, the OmniSwitch automatically creates a dynamic SAP and a dynamic UNP profile defining the SPB service that traffic will be mapped to. The following snippet provides an example of such a dynamically created UNP profile. The profile in the snippet is created upon reception of traffic tagged with VLAN 101. How does the AOS select the ISID and BVLAN to be used in the newly created service? It uses the formulas below where '%' denotes the "modulo" division: the reminder of the integer division.

- ISID Number = Base Service Number (BSN) + Domain ID + (VLAN Number % Service Modulo)
- BVLAN Index = ISID Number % (Total number of BVLANs)

By default:

- BSN = 10,000,000
- Domain ID = 0
- Service Modulo = 512

Let's also assume that BVLANs 4000-4003 are created and calculate the ISID and BVLAN number manually.

- ISID Number = 10,000,000 + 0 + (101 % 512) = 10,000,000 + 101 = 10,000,101
- BVLAN Index = 10,000,101 % 4 = 1

The formula does not provide the BVLAN number directly but the BVLAN index: the position in a BVLAN array sorted in ascending order where the lowest numbered BVLAN is in position 0 and the highest numbered BVLAN is in position N-1. Therefore, in our example, with BVLANs 4000-4003, BVLAN index 1 maps to BVLAN 4001.

BEB-1

```
BEB-1> unp profile "systemDefault10000101" map service-type spb tag-value 101 isid
10000101 bvlan 4001 multicast-mode headend vlan-xlation
```

It is important to understand that with 4096 possible VLAN tags, using the default Service Modulo of 512 can result in up to eight different VLAN tags being mapped to the same service. This is not the desired outcome most of the time because it will result in different VLAN traffic being bridged in the same L2 domain. To ensure L2 isolation, we can change the Service Modulo to 4096 as shown in the snippet below.

BEB-1

```
BEB-1> unp system-default service-mod 4096
```

Let's now focus on another parameter used in the ISID calculation formula: Domain ID. The Domain ID is useful in a multi-tenanted environment. For example, consider a network providing services to three different customers: A, B, and C. These customers can use multiple VLANs and some of those VLANs may overlap. How do you ensure customer traffic isolation in the SPB domain? Isolation is achieved by creating a Domain ID for each customer and by the mapping customer's UNI ports to the domain. The example in below snippet illustrates this configuration. Domains 1 through 3 are created for customers A through C. Ports 1/1/1-10 connecting customer A's devices are mapped to domain 1, ports 1/1/11-21 connecting customer B's devices are mapped to domain 2, and so on. This configuration preserves customer isolation even when services and SAPs are dynamically and automatically configured on the fly in response to VLAN tags in incoming traffic.

BEB-1

```
BEB-1> unp system-default service-mod 4096
BEB-1> unp domain 1 description "Customer_A"
BEB-1> unp domain 2 description "Customer_B"
BEB-1> unp domain 3 description "Customer_C"
BEB-1> unp port 1/1/1-10 domain 1
BEB-1> unp port 1/1/11-20 domain 2
BEB-1> unp port 1/1/21-30 domain 3
```

Last, the BSN enables manual and dynamic service coexistence without conflict. Dynamically created services map to ISIDs greater than or equal to the BSN. Manually created services should use ISID numbers lower than the BSN.

Alcatel-Lucent OmniAccess Stellar AP integration

Access Guardian provides the integration framework that enables automatic discovery, classification, and management of Alcatel-Lucent OmniAccess® Stellar Access Points (APs) connected to an OmniSwitch operating in an SPB environment. Once detected, wireless client traffic is forwarded from the AP to the OmniSwitch and transported across the SPB service domain using dynamically created SPB SAPs. This integration delivers a unified wired and wireless access architecture within the SPB fabric, minimizing manual configuration while ensuring consistent service-based forwarding.

This framework applies to both UNP bridge ports (for VLAN domain) and UNP access ports (for SPB service domain). This section explains only the SPB service domain relevant details as UNP access ports were explained in the Dynamic SAPs section.

The OmniSwitch boots up with specific default configuration and operational settings that trigger the following process to detect, learn and classify connected Stellar AP devices:

1. The OmniSwitch and any Stellar AP device that is connected to the UNP access port initially exchange Link Layer Detection Protocol (LLDP) TLV packets. Through this exchange of LLDP packets, the switch identifies and learns the device MAC address as an AP.
2. If the AP mode is enabled for the UNP port and an AP device is detected on that port, the transmission of LLDP Port VLAN ID and AP Location TLVs is operationally enabled on the UNP access port.
3. Once the AP MAC address is detected and learned, a built-in LLDP UNP classification rule for APs classifies the AP device into the built-in default profile "**defaultWLANAccessProfile**" for AP devices connected to UNP access ports. The profile is mapped to SPB service parameters. When the AP device is classified into this profile, an SPB SAP is dynamically created. A SAP-port association is established between the UNP access port and the SAP on which the AP MAC address is learned and forwarded.
4. After the AP device connection is established, classified, and the SPB service is assigned, any of the following actions can occur:

- The AP device sends DHCP packets.
- The switch transmits LLDP packets to the AP device to advertise the SPB service and AP location information.
- On a UNP access port, the AP device starts to send client-tagged traffic (tagged with the SSID VLAN). The switch will attempt to assign the AP client traffic to a UNP service profile configured with an SPB service tag value that matches the VLAN tag of the client traffic. If a matching UNP service profile does not exist, then the AP client traffic is assigned to the System Default profile that will dynamically create the SPB service SAP on which to forward the AP client traffic.

The OmniSwitch detection and integration of OmniAccess Stellar APs results in a switch configuration that includes an SPB service for the AP device and additional SPB services for wireless client-tagged traffic that is forwarded by the AP onto the wired network.

The configuration for this process includes the following tasks:

- Configure a Layer 2 profile with a "peer" action defined for 802.1AB control frames.
- Configure any switch port that will connect to a Stellar AP device as a UNP access port and assign the Layer 2 profile created in the previous step to that port.
- If necessary, enable the UNP AP mode for the UNP access port. By default, the global AP mode status is enabled, and this step is only required if the global AP mode status is disabled.
- 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 SPB service parameters to the built-in "defaultWLANAccessProfile".
- (Optional) Configure the QoS policy list, authentication flag status or mobile tag status for the "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.

The following example illustrates the framework and how it works. In Figure 29, the goal is to enable automatic discovery, classification and management of the OmniAccess Stellar AP "AP-1" connected to the OmniSwitch "BEB-1" connected to the SPB backbone. Once detected, the AP will be assigned to the "defaultWLANAccessProfile". The SPB service that is mapped to this profile will serve as the management service for the classified AP devices. AP client traffic tagged with SSID VLAN 10 and 20 will be assigned to the "spb10" and "spb20" profiles, respectively. A SAP is then dynamically created based on each profiles service parameters to carry the client traffic through the SPB service domain.

Figure 29 - OmniSwitch and OmniAccess Stellar integration

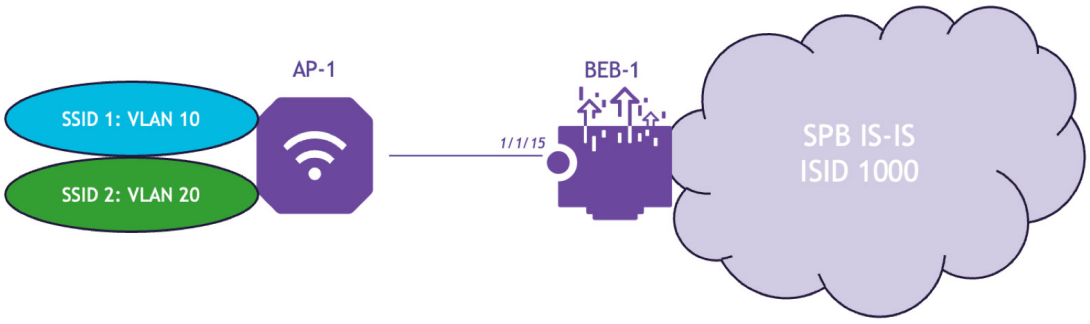

The following is a sample configuration snippet for BEB-1:

BEB-1

```
BEB-1> service l2profile "ap-SvcUnp" 802.1ab peer
BEB-1> unp port 1/1/15 port-type access
BEB-1> unp port 1/1/15 l2-profile ap-SvcUnp
BEB-1> unp port 1/1/15 ap-mode
BEB-1> unp profile defaultWLANAccessProfile map service-type spb tag-value 0 isid 1000 bvlan 4000
BEB-1> unp profile spb10
BEB-1> unp profile spb10 map service-type spb tag-value 10 isid 1010 bvlan 4001
BEB-1> unp profile spb20
BEB-1> unp profile spb20 map service-type spb tag-value 20 isid 1020 bvlan 4002
BEB-1> unp classification vlan-tag 10 profile1 spb10
BEB-1> unp classification vlan-tag 20 profile1 spb20
```

Management

As explained in the Non-IP core section, SPB IS-IS is not an IP protocol. BCB nodes do not require IP interfaces. BEB nodes supporting L2 services only do not require IP interfaces either. BEB nodes require IP interfaces only when supporting an L3 service (for example, L3 VPN or VPN Lite). However, all SPB nodes whether BCB or BEB, require IP interfaces for management purposes.

There are different ways of managing SPB nodes:

- **Out of Band Management (OOBM)**: OOBM is applicable to any network architecture and will not be discussed further.
- **Dedicated Management Service**: An SPB service and VRF are dedicated to management.
- **In-band Management**: Management IP interfaces can be created directly on the control BVLAN. The management network or stations attach to one or more gateway nodes through VLAN-domain interfaces. IP interfaces created on the control BVLAN do not support configuration of any routing protocol or function (for example, OSPF or VRRP) and do not rely on ARP for IP-to-MAC resolution because there are no broadcasts on the SPB domain. IP-to-MAC mapping is resolved through IS-IS TLVs. IS-IS TLVs also carry management routes through the SPB backbone. VLAN-domain and SPB-domain management routes can be cross-redistributed at gateways nodes. The "spb-mgmt" protocol is associated to SPB-domain management routes.

Figure 30 - In-band management

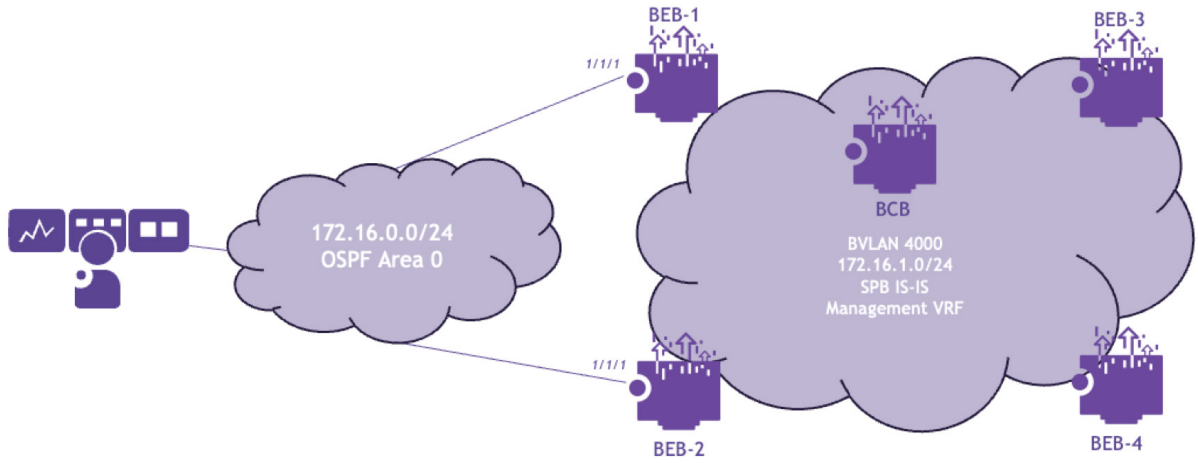

In the in-band management example shown in Figure 30, nodes BEB-1 and BEB-2 are gateways nodes linking the SPB-management domain and the VLAN-management domain. The SPB control BVLAN is 4000. The VLAN-management subnet is 172.16.0.0/24 and the SPB-management subnet is 172.16.1.0/24. OSPF is used in the management network. Nodes BEB-1 and BEB-2 redistribute routes between OSPF and SPB-MGMT protocols. Route maps prevent circular route redistribution between these two protocols.

BEB-1

```
BEB-1> vlan 1000 name Management-VLAN
BEB-1> vlan 1000 members port 1/1/1 untagged
BEB-1> vrf create Management
Management: :BEB-1> ip interface "Management-SPB" address 172.16.1.1 mask 255.255.255.0 vlan 4000
Management: :BEB-1> ip interface "Management-VLAN" address 172.16.0.1 mask 255.255.255.0 vlan 1000
Management: :BEB-1> ip service all admin-state enable
Management: :BEB-1> ip redist ospf into spb-mgmt route-map vlan-mgmt-routes
Management: :BEB-1> ip redist spb-mgmt into ospf route-map spb-mgmt-routes
```

BEB-2

```
BEB-2> vlan 1000 name Management-VLAN
BEB-2> vlan 1000 members port 1/1/1 untagged
BEB-2> vrf create Management
Management: :BEB-2> ip interface "Management-SPB" address 172.16.1.2 mask 255.255.255.0 vlan 4000
Management: :BEB-2> ip interface "Management-VLAN" address 172.16.0.2 mask 255.255.255.0 vlan 1000
Management: :BEB-2> ip service all admin-state enable
Management: :BEB-2> ip redist ospf into spb-mgmt route-map vlan-mgmt-routes
Management: :BEB-2> ip redist spb-mgmt into ospf route-map spb-mgmt-routes
```

BEB-3

```
BEB-3> vrf create Management
Management: :BEB-3> ip interface "Management-SPB" address 172.16.1.3 mask 255.255.255.0 vlan 4000
Management: :BEB-3> ip service all admin-state enable
```

BEB-4

```
BEB-4> vrf create Management
Management: :BEB-4> ip interface "Management-SPB" address 172.16.1.4 mask 255.255.255.0 vlan 4000
Management: :BEB-4> ip service all admin-state enable
```

BCB

```
BCB> vrf create Management
Management: :BCB> ip interface "Management-SPB" address 172.16.1.5 mask 255.255.255.0 vlan 4000
Management: :BCB> ip service all admin-state enable
```

In-band management configuration examples are provided in the previous snippets. OSPF and route-map configuration in BEBs 1 and 2 is excluded from these snippets.

Operation and maintenance

Connectivity Fault Management: 802.1ag

Connectivity Fault Management (CFM) in an SPB network is useful for performing L2 trace and L2 ping for analysis and troubleshooting. Other aspects of CFM such as fault detection, which are important in PBB, are less important in SPB because SPB has an IS-IS control plane. These functions (Continuity Check Messages (CCM)) are not currently supported in conjunction with SPB.

Ethernet Operations, Administration and Maintenance (OAM) is supported at the BVLAN level (see Figure 31). Virtual Maintenance End Points (MEPs) must be configured for all BVLANs and BEBs and, optionally, for BCBs (such that a L2 PING or L2 trace test can be initiated from any node to any other node). MEPs are configured manually on BEBs, while Maintenance Intermediate Points (MIPs) are created automatically only if the node participates in the OAM domain and association. If the intermediate BCB is not part of the OAM domain/association, it does not create MIPs.

Since there is no CCM function to map system names, link trace commands and output will reference the BMACs.

Figure 31 - OAM in BVLAN and VLAN domains

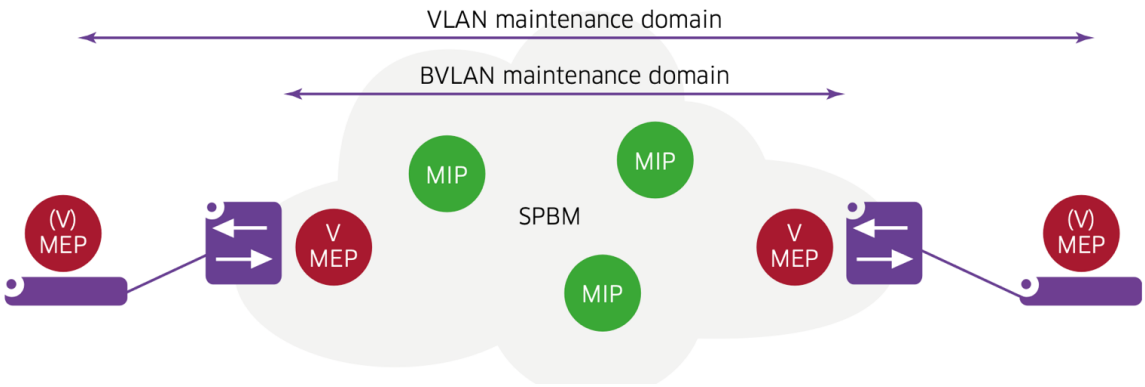

OAM is also supported at the VLAN level or between L2 access switches connected to BEBs over SAP UNIs. This is useful in a L2 deployment for testing end-to-end service connectivity between sites. OAM at the VLAN level must be set at a higher maintenance domain level than BVLAN OAM.

Figure 32 shows a practical example of how OAM can be used to verify connectivity between BEBs by means of Loopback Message (LBM) and Loopback Reply (LBR) and checking the route with Link Trace Message (LTM) and Link Trace Reply (LTR).

Figure 32 - L2 ping and L2 trace

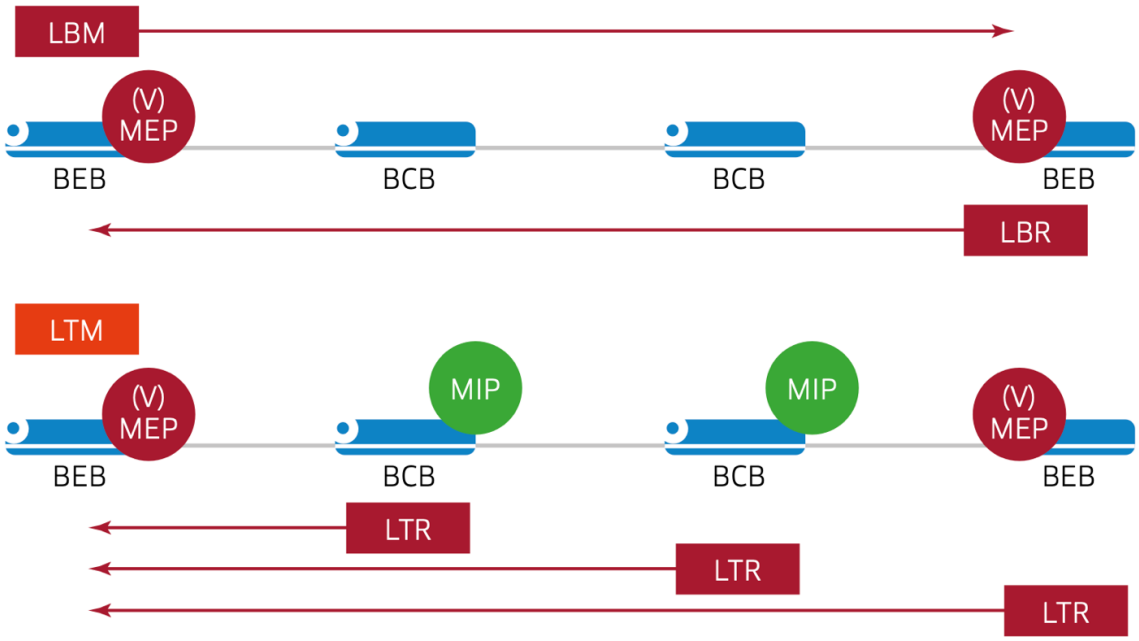

The below configuration snippet provides a sample OAM configuration for service BVLANs 4001-4003 for the MEP on the BEBs and for the MIP on the BCBs.

BEB-1

```
BEB-1> ethoam domain ALE format string level 3
BEB-1> ethoam domain ALE mhf default
BEB-1> ethoam domain ALE id-permission chassisid
BEB-1> ethoam association BVLAN4001 format string domain ALE
BEB-1> ethoam association BVLAN4001 domain ALE primary-vlan 4001
BEB-1> ethoam association BVLAN4001 domain ALE mhf default
BEB-1> ethoam association BVLAN4001 domain ALE id-permission chassisid
BEB-1> ethoam association BVLAN4001 domain ALE ccm-interval interval1s
BEB-1> ethoam association BVLAN4001 domain ALE endpoint-list 11-14
BEB-1> ethoam association BVLAN4002 format string domain ALE
BEB-1> ethoam association BVLAN4002 domain ALE primary-vlan 4002
BEB-1> ethoam association BVLAN4002 domain ALE mhf default
BEB-1> ethoam association BVLAN4002 domain ALE id-permission chassisid
BEB-1> ethoam association BVLAN4002 domain ALE ccm-interval interval1s
BEB-1> ethoam association BVLAN4002 domain ALE endpoint-list 21-24
BEB-1> ethoam association BVLAN4003 format string domain ALE
BEB-1> ethoam association BVLAN4003 domain ALE primary-vlan 4003
BEB-1> ethoam association BVLAN4003 domain ALE mhf default
BEB-1> ethoam association BVLAN4003 domain ALE id-permission chassisid
BEB-1> ethoam association BVLAN4003 domain ALE ccm-interval interval1s
BEB-1> ethoam association BVLAN4003 domain ALE endpoint-list 31-34
BEB-1> ethoam endpoint 11 domain ALE association BVLAN4001 direction up port virtual primary-vlan 4001
BEB-1> ethoam endpoint 11 domain ALE association BVLAN4001 admin-state enable
BEB-1> ethoam endpoint 11 domain ALE association BVLAN4001 priority 5
BEB-1> ethoam endpoint 11 domain ALE association BVLAN4001 lowest-defect-priority all-defect
BEB-1> ethoam endpoint 21 domain ALE association BVLAN4002 direction up port virtual primary-vlan 4002
BEB-1> ethoam endpoint 21 domain ALE association BVLAN4002 admin-state enable
BEB-1> ethoam endpoint 21 domain ALE association BVLAN4002 priority 5
BEB-1> ethoam endpoint 21 domain ALE association BVLAN4002 lowest-defect-priority all-defect
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4003 direction up port virtual primary-vlan 4003
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4003 admin-state enable
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4003 priority 5
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4003 lowest-defect-priority all-defect
```

BCB-1

```
BCB-1> ethoam domain ALE format string level 3
BCB-1> ethoam domain ALE mhf default
BCB-1> ethoam domain ALE id-permission chassisid
BCB-1> ethoam association BVLAN4001 format string domain ALE
BCB-1> ethoam association BVLAN4001 domain ALE primary-vlan 4001
BCB-1> ethoam association BVLAN4001 domain ALE mhf default
BCB-1> ethoam association BVLAN4001 domain ALE id-permission chassisid
BCB-1> ethoam association BVLAN4001 domain ALE ccm-interval interval1s
BCB-1> ethoam association BVLAN4001 domain ALE endpoint-list 11-14
BCB-1> ethoam association BVLAN4002 format string domain ALE
BCB-1> ethoam association BVLAN4002 domain ALE primary-vlan 4002
BCB-1> ethoam association BVLAN4002 domain ALE mhf default
BCB-1> ethoam association BVLAN4002 domain ALE id-permission chassisid
BCB-1> ethoam association BVLAN4002 domain ALE ccm-interval interval1s
BCB-1> ethoam association BVLAN4002 domain ALE endpoint-list 21-24
BCB-1> ethoam association BVLAN4003 format string domain ALE
BCB-1> ethoam association BVLAN4003 domain ALE primary-vlan 4003
BCB-1> ethoam association BVLAN4003 domain ALE mhf default
BCB-1> ethoam association BVLAN4003 domain ALE id-permission chassisid
BCB-1> ethoam association BVLAN4003 domain ALE ccm-interval interval1s
BCB-1> ethoam association BVLAN4003 domain ALE endpoint-list 31-34
```

The below snippet provides sample configuration and output for an L2 trace test. As shown in the snippet, the trace provides, among other elements, BMACs for all transit nodes as well as ingress and egress interfaces used.

BEB-1

```
BEB-1> ethoam linktrace target-macaddress dc:08:56:10:80:f9 source-endpoint 11 domain ALE
association BVLAN4001

Transaction Id:635723059

BEB-1> show ethoam linktrace-reply domain ALE association BVLAN4001 endpoint 11 tran-id 635723059
LTM operation unsuccessful. Target is reachable.
Ttl : 62,
  LTM Forwarded : yes,
  Terminal MEP : no,
  Last Egress Identifier : 00:00:DC:08:56:10:85:59,
  Next Egress Identifier : 00:00:DC:08:56:10:7F:19,
  Relay Action : RLY_FDB,
  Chassis ID Subtype : LOCALLY_ASSIGNED,
  Chassis ID : BCB-1,
  Ingress Action : ING_OK,
  Ingress Mac : DC:08:56:10:7F:58,
  Ingress Port ID Subtype : LOCALLY_ASSIGNED,
  Ingress Port ID : 1/1/51A,
  Egress Action : EGR_OK,
  Egress Mac : DC:08:56:10:7F:5C,
  Egress Port ID Subtype : LOCALLY_ASSIGNED,
  Egress Port ID : 1/1/52A


Ttl : 61,
  LTM Forwarded : no,
  Terminal MEP : yes,
  Last Egress Identifier : 00:00:DC:08:56:10:7F:19,
  Next Egress Identifier : 00:00:DC:08:56:10:80:F9,
  Relay Action : RLY_HIT,
  Chassis ID Subtype : NONE,
  Chassis ID : none,
  Ingress Action : ING_OK,
  Ingress Mac : DC:08:56:10:80:40,
  Ingress Port ID Subtype : LOCALLY_ASSIGNED,
  Ingress Port ID : 1/1/53A,
  Egress Action : EGR_NONE,
  Egress Mac : 00:00:00:00:00:00,
  Egress Port ID Subtype : NONE,
  Egress Port ID : none
```

Network performance: Service Assurance Agent

With Service Assurance Agent (SAA), users can verify service guarantees, increase network reliability by validating network performance and proactively identify network issues. SAA uses active monitoring to generate traffic in a continuous, reliable and predictable manner, enabling the measurement of network performance and health. This is done using standard IP PING packets, proprietary MAC pings and Ethernet OAM tests. A large number of test sessions can be configured on the switch, with each test having the ability to send notification traps and provide a method for determining network performance.

Each SAA test can specify threshold values for jitter and round-trip-time (RTT). When SAA processes an iteration of a test session, it will compare the results against the following criteria to see if an SNMP trap should be sent. A trap with the session name is sent if:

- At least one packet is lost
- Warning: Average RTT/Jitter crosses 90% of threshold
- Critical: Average RTT/Jitter at or above threshold

Latency, jitter and packet loss SAA tests can be automatically set up between all BEBs and BCBs and across all BVLANs with the “saa auto-create” command. The following snippets show the configuration and sample statistics.

BEB-1

```
BEB-1> saa spb auto-create auto-start
```

BEB-1

```
BEB-1> show saa statistics aggregate
Legend: eth-lb  = ethoam-loopback
        eth-dmm = ethoam-two-way-delay
        - = Delay or jitter value not available

Aggregate Record:
             SAA                    Owner    Type         Time of Last-Run       RTT     RTT     RTT     RTT    Jitter  Jitter  Jitter  Jitter  Packets    Description
                                                                                 Min     Avg     Max     Thr    Min     Avg     Max     Thr   Sent Rcvd
------------------------------+--------+---------+--------------------+-------+-------+-------+-------+-------+-------+-------+-------+-----+----+-----------------
SPB-4000-dc-08-56-10-72-49      SPB      mac-ping      2025-11-20,08:58:38     173     187     205       0       6      15      26       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-74-29      SPB      mac-ping      2025-11-20,08:58:33     137     180     210       0      23      28      50       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-77-e9      SPB      mac-ping      2025-11-20,08:58:28     133     147     163       0       7      14      30       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-78-d9      SPB      mac-ping      2025-11-20,08:58:23     135     161     184       0       0      24      49       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-7f-19      SPB      mac-ping      2025-11-20,08:58:18     123     154     181       0       7      15      50       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-80-f9      SPB      mac-ping      2025-11-20,08:58:13     149     158     167       0       2       6      13       0     5    5 DEFAULT
SPB-4000-dc-08-56-10-86-49      SPB      mac-ping      2025-11-20,08:58:08     129     152     165       0       4       9      26       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-72-49      SPB      mac-ping      2025-11-20,08:59:13     159     181     202       0       1      19      43       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-74-29      SPB      mac-ping      2025-11-20,08:59:08     144     194     220       0       3      31      76       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-77-e9      SPB      mac-ping      2025-11-20,08:59:03     143     162     183       0       7      11      32       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-78-d9      SPB      mac-ping      2025-11-20,08:58:58     126     162     202       0      22      29      54       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-7f-19      SPB      mac-ping      2025-11-20,08:58:53     116     156     187       0      17      32      57       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-80-f9      SPB      mac-ping      2025-11-20,08:58:48     119     139     161       0       1      14      42       0     5    5 DEFAULT
SPB-4001-dc-08-56-10-86-49      SPB      mac-ping      2025-11-20,08:58:43     116     154     170       0       0      25      54       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-72-49      SPB      mac-ping      2025-11-20,08:59:48     164     177     200       0      22      21      35       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-74-29      SPB      mac-ping      2025-11-20,08:59:43     163     189     223       0       6      15      48       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-77-e9      SPB      mac-ping      2025-11-20,08:59:38     168     181     188       0       2       6      14       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-78-d9      SPB      mac-ping      2025-11-20,08:59:33     174     198     235       0      14      25      57       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-7f-19      SPB      mac-ping      2025-11-20,08:59:28     163     176     209       0       4      12      43       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-80-f9      SPB      mac-ping      2025-11-20,08:59:23     158     182     201       0       7      13      23       0     5    5 DEFAULT
SPB-4002-dc-08-56-10-86-49      SPB      mac-ping      2025-11-20,08:59:18     193     211     239       0       2      18      44       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-72-49      SPB      mac-ping      2025-11-20,09:00:23     119     153     190       0      35      31      47       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-74-29      SPB      mac-ping      2025-11-20,09:00:18     144     158     182       0       9      15      36       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-77-e9      SPB      mac-ping      2025-11-20,09:00:13     156     196     226       0      20      32      70       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-78-d9      SPB      mac-ping      2025-11-20,09:00:08     175     183     196       0       1       7      21       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-7f-19      SPB      mac-ping      2025-11-20,09:00:03     174     202     221       0       3      17      47       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-80-f9      SPB      mac-ping      2025-11-20,08:59:58     168     180     193       0       1       8      25       0     5    5 DEFAULT
SPB-4003-dc-08-56-10-86-49      SPB      mac-ping      2025-11-20,08:59:53     139     155     181       0       4      19      38       0     5    5 DEFAULT
```

The SAA feature also provides the ability to periodically record the last five iterations of all SAA sessions to an XML file on the local switch. The name of the XML file and the logging time interval are configurable SAA XML parameters.

Network maintenance

Overload state

SPB provides a graceful way to remove a node from service for maintenance and transition traffic to an alternate path (if there is one) with minimal disruption. This is the "overload state."

Setting the overload state on the node will signal other nodes not to use it as a transit node and use alternate paths instead. This is similar to increasing the metric on all the links but is a much quicker way of achieving this outcome. Note, however, once the overload state is enabled on a node no traffic will transit through the node even if there are no alternative paths.

The overload state can be set indefinitely (until removed) or it can revert after a timer expires.

Graceful restart

SPB IS-IS supports graceful restart in a VC or physical chassis with redundant control modules.

Without graceful restart, a VC master or Chassis Management Module (CMM) takeover event would require neighbor nodes to tear down and re-establish adjacencies with the restarting node and re-build the topology database, resulting in some disruption to traffic flows.

When graceful restart is enabled, and with the help of a neighbor node, the node undergoing a takeover will announce this condition to its neighbors by setting the RR (restart request) in a TLV message and continue using its existing FDB while restarting. The neighbor nodes will maintain their adjacencies with the restarting node during this process and send their complete LSP database information to the restarting node once the process is complete.

This makes the transition a much smoother process because disruption to traffic forwarding is minimized and the topology database is rebuilt in a much shorter time.

Service attachment redundancy

When redundant links and nodes exist in the SPB domain, path computation in the event of a failure or restoration event is handled by the IS-IS protocol. Access or CE devices connected to BEB nodes do not run SPB IS-IS, so other solutions are required when redundancy is needed. This section describes options for the different service types.

The simplest way of achieving redundant CE to BEB attachment is to use VC at the BEB and to attach the CE device to the BEB through a LAG. This redundancy option is applicable to any service type (L2 or L3).

There are alternate redundancy options other than VC+LAG.

For the L2 services depicted in Figure 33, there are the following options:

- **Non-redundant**: The CE is attached to a single BEB through a single link. Link, BEB or CE failure will result in loss of service to the site.
- **Redundant links**: The CE is attached to a single BEB through a link aggregate (LAG). This adds protection from single-link failure. Note that fiber runs should use diverse physical paths to protect against fiber cuts which would typically interrupt both links otherwise.
- **Redundant links and nodes**: The CE is attached to two different BEBs through two different links. This adds protection from BEB failure. When possible, both links should use physically diverse paths such that link failure events are not correlated. Dual-Home Link (DHL) is a high availability feature that provides fast failover without implementing Spanning Tree or Link Aggregation. Refer to the OmniSwitch AOS Network Configuration User Guide referenced in the Related documents section for further details.

- **Fully redundant**: This option adds CE device redundancy. MSTP can be used to avoid loops in this redundant connection. By default, SPB floods STP BPDUs messaging over SPB services. When using MSTP, different sites must use different MSTP regions to avoid creating a large MSTP region spanning all sites.

VC can be combined with all the options above to increase resiliency.

Figure 33 - L2 service attachment

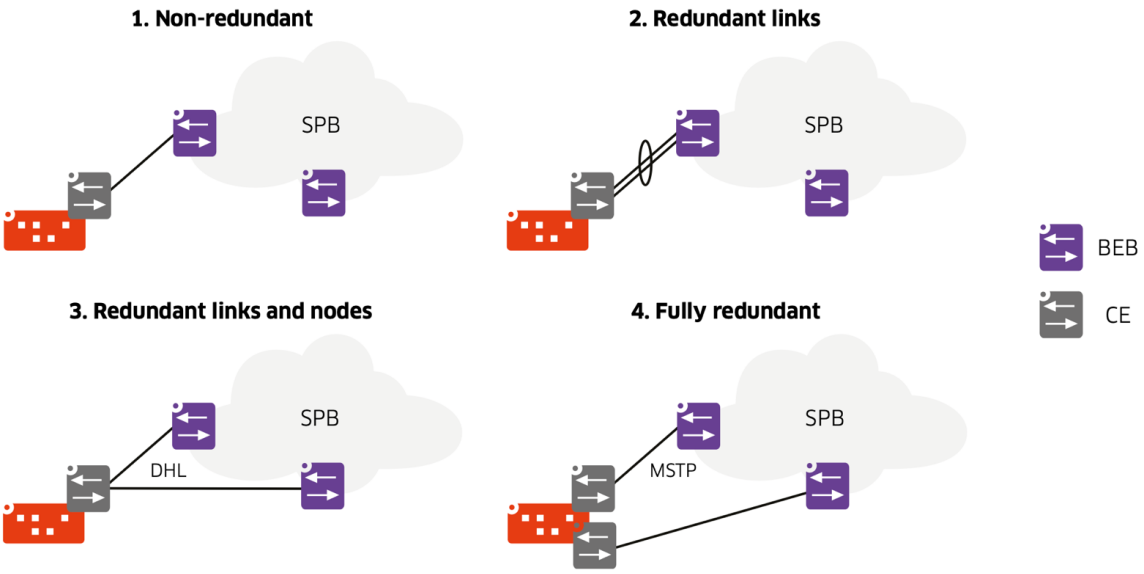

Let's now continue with L3 services. We can distinguish two sub-variants: L3 CE and L2 CE. A L3 CE can exchange routes with the BEBs by using any supported routing protocol as well as static or default routes. A L2 CE on the other hand will completely delegate routing to the BEB, which will act as a default gateway for local devices. These two sub-variants are illustrated in Figure 34 and 35.

L3 service attachment with L3 CE options:

- **Non-redundant**: The site is attached to a single BEB through a single link. Link, BEB or CE failure will result in loss of service to the site.
- **Redundant links**: The site is attached to a single BEB through a LAG. This adds protection from single-link failure. Note that fibre runs should use diverse physical paths to protect against fibre cuts which would typically interrupt both links otherwise.
- **Redundant links and nodes**: The site is attached to two different BEBs through two different links. This adds protection from BEB failure. When possible, both links should use physically diverse paths such that link failure events are not correlated. A dynamic routing protocol such as OSPF is used between BEBs and CEs to exchange routing information. Import/Export and re-distribution of routes must be carefully planned to avoid circular re-distribution of routes. This is accomplished with route tags and route-maps. This will be further discussed in the SPB L3VPN route tags section.
- **Fully redundant**: This option adds CE device redundancy.

Figure 34 - L3 service attachment - L3 CE

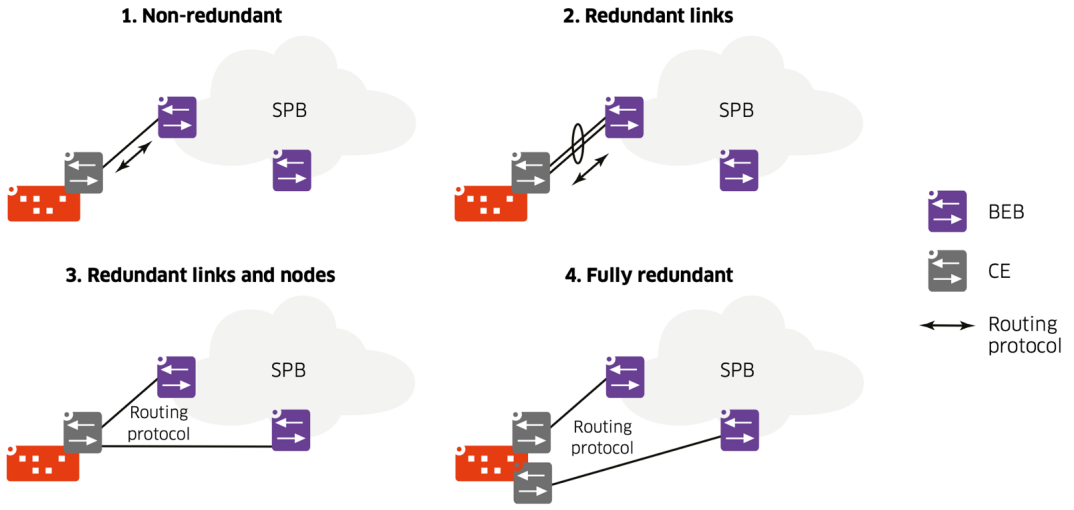

The case of L3 service attachment with a L2 CE is almost identical to the case of L2 service attachment but, since the routing function is delegated to the BEB, VRRP is required when CEs attach to redundant BEBs. This requires access VLANs to be extended across both BEBs. If BEBs are directly connected, the access VLANs can be simply tagged on the link interconnecting both BEBs. However, if there is no direct connection between the BEB pair, a dedicated SPB service can be created to this effect.

In addition, when using a L2 CE in a L3 service, there is no routing protocol between CE and BEB. In this case, the associated VRF can be configured as a "low profile" VRF. Low profile VRFs have routing capabilities restricted to static and/or imported routes, which is sufficient for such a situation. Low profile VRFs take up less BEB resources than "max profile" VRFs, allowing for the creation of more VRFs on the BEB.

As in the case of L2 service attachment, all options can be combined with VC and LAG.

Figure 35 - L3 service attachment - L2 CE

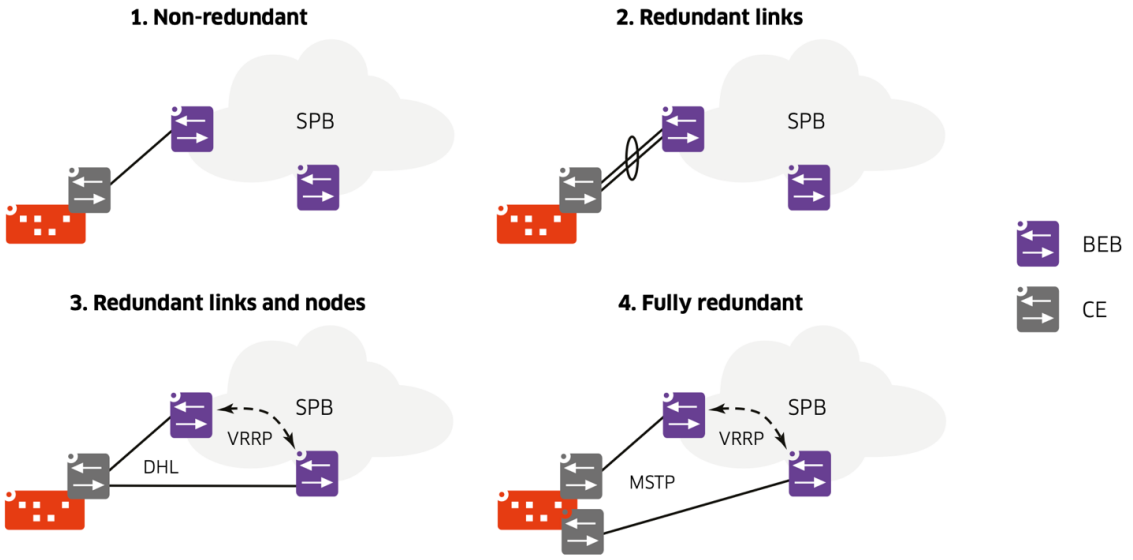

Loop avoidance and suppression

In the CP, loops are avoided with IS-IS, a link-state routing protocol. In the DP, a node will not accept unexpected frames from its neighbours.

However, short-lived transient loops may form in the event of a topology change and until network convergence is attained. Loops pose a serious threat to the network stability.

In the DP, SPB incorporates an additional loop mitigation technique to detect and break these transient loops:

- **Reverse-path Forwarding Check (RPFC)**: RPFC exploits SPB's symmetry and congruence properties. RPFC verifies that incoming traffic's source BMAC is indeed reachable over the ingress interface according to the local FDB and discards non-conforming frames.

In addition, the SPB backbone must be protected from loops that may be created due to failures and misconfiguration at the VLAN-domain 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.

There is always a chance however that STP may be misconfigured, fail, or not be enabled at all. Configuration faults in customer networks can result in loops spanning both the SPB backbone and customer access network. This can result in broadcast storms. To protect the SPB backbone from broadcast storms, loops involving SAPs must be detected and broken.

AOS supports an additional loop mitigation mechanism to detect and break access layer loops: Loopback Detection (LBD). LBD can detect and protect the backbone network from forwarding loops created at the VLAN-domain customer-access layer. LBD operates in addition to other mechanisms such as DHL or STP. When a loop is detected, the port is disabled and goes into a shutdown state. A trap is sent and the event is logged.

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.

LBD can be used on both VLAN UNI and SAP UNI ports. In the case of SAP UNI ports, LBD frames will be sent on all SAPs because different access VLANs may have different logical topologies. However, if a loop is detected on a SAP, the entire physical port will be shut down. LBD should be enabled on all UNI ports.

Figure 36 illustrates situations in which LBD can detect and break loops.

Figure 36 - Loopback detection

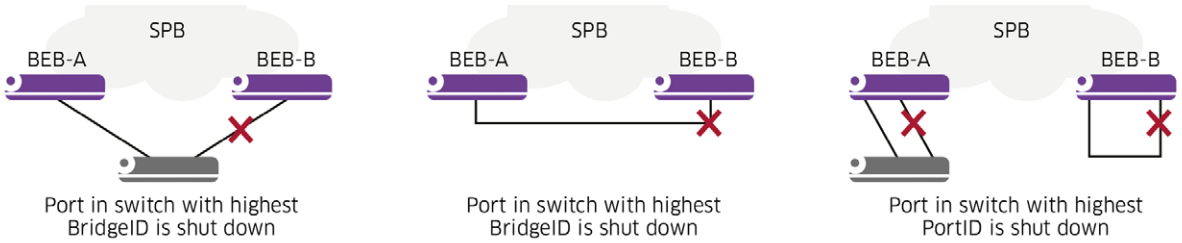

By default, LBD is disabled for the switch and on all service-access ports. Enable LBD globally on the switch and in specific service-access ports or linkages as shown in the below snippet.

BEB-1

```
BEB-1> loopback-detection enable
BEB-1> loopback-detection service-access port 1/1/1 enable
BEB-1> loopback-detection service-access linkagg 1 enable
```

AOS incorporates storm control through flood rate limiting of BUM traffic. To mitigate the consequences of a loop in case all protection mechanism were broken, the flood rate limitation can be used to restrict flooded BUM traffic from harming the network. A high threshold rate is configured in megabits-per-second (mbps), packets-per-second (pps) or as a percentage of the port speed. When the threshold value is reached, packets are dropped or, the port is shutdown. Storm control is enabled by default with pre-defined rates. Refer to the OmniSwitch AOS Network Configuration User Guide referenced in the Related documents section for further details.

Advanced SPB designs

SPB over multi-access networks

By default, ISIS-SPB operates over point-to-point (PtP) links which allows only one adjacency on an SPB network interface; however, an SPB network interface can be configured to allow multiple adjacencies to form on the interface. This is particularly useful for extending an SPB backbone over a shared Ethernet domain, such as a service provider network or even to connect to another ISIS-SPB domain.

An SPB multiple access (multi-access) network interface is configured on SPB BEBs that connect directly to a shared network instead of to SPB BCBs. Each BEB forms ISIS-SPB adjacencies over the shared network with all the other BEBs on the multi-access network interfaces.

Participating BEBs elect one of the multi-access network interfaces to serve as the Designated Intermediate System (DIS). The DIS represents all of the multi-access links as a virtual SPB node (pseudo-node). The purpose of the DIS is to create and manage the pseudonode, synchronizing the Link State Database (LSDB) among the BEBs, and for flooding LSPs over the network.

Figure 37 shows an example of an SPB backbone extended over a service provider network.

Figure 37 - SPB over multi-access networks

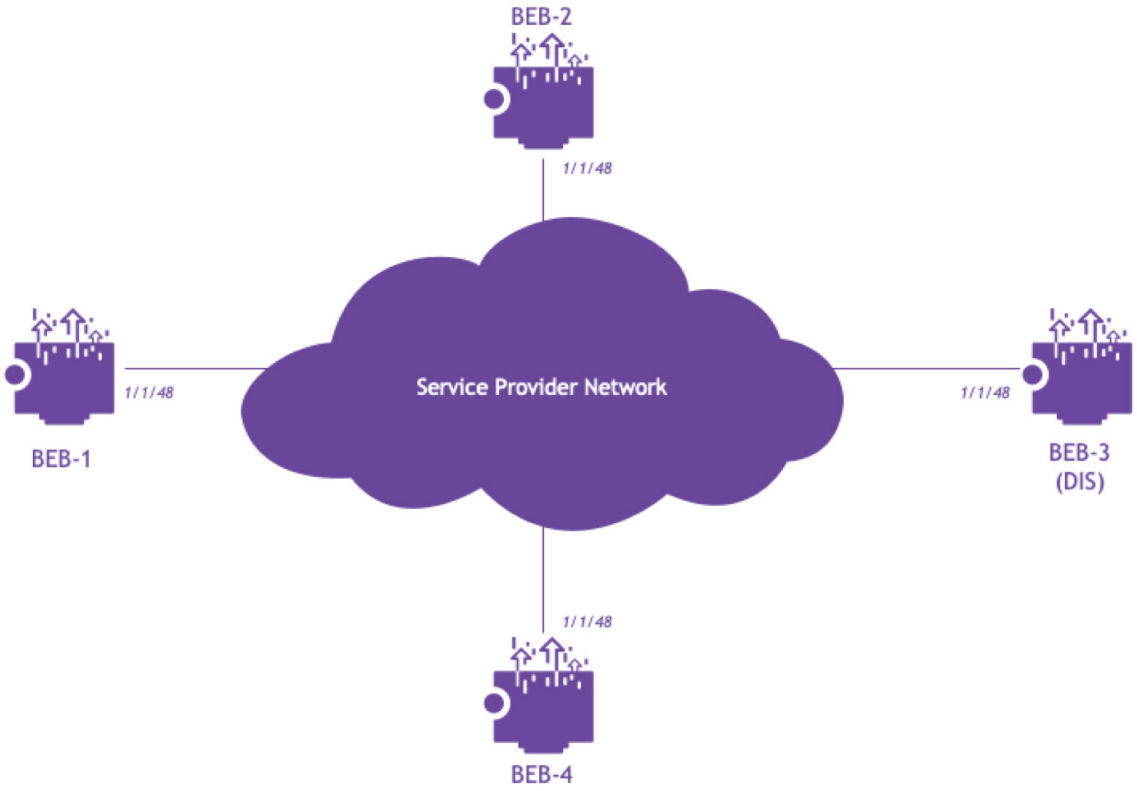

In order for the BEBs in Figure 37 to communicate, they need to form adjacencies with each of the other BEBs over the service provider network. This is not possible with a PtP configuration, so each SPB network interface port is configured as a multi-access LAN interface to allow multiple adjacencies to form across the broadcast network domain.

To the rest of the network, the SPB multi-access links are seen as a virtual node that is defined and represented by the DIS through pseudo-node LSPs. Figure 38 provides a logical depiction of how the SPB pseudo-node is interpreted by IS-IS with a DIS:

Figure 38 - SPB over multi-access networks - logical view

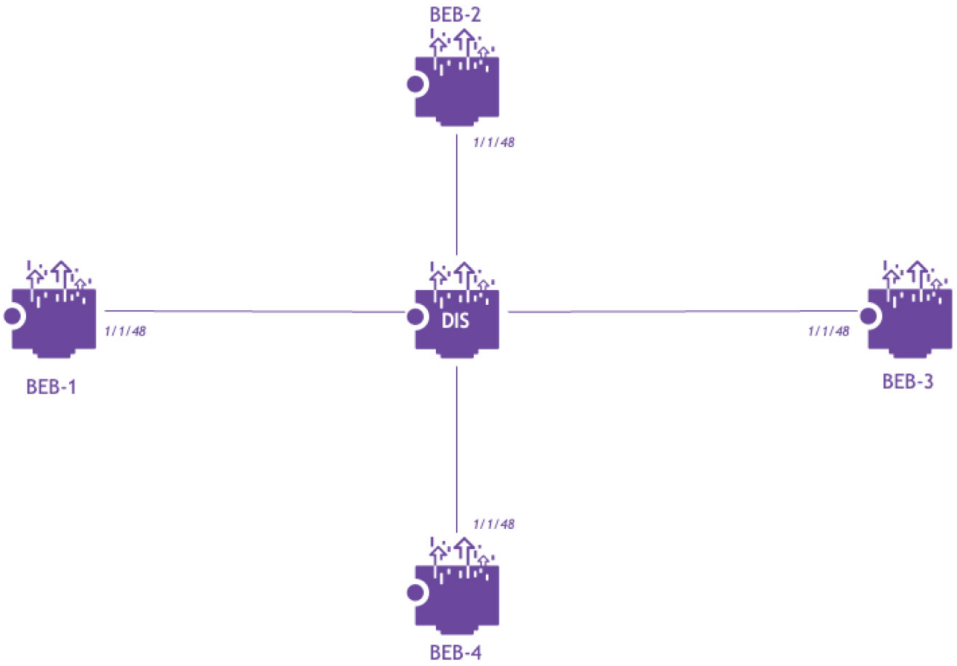

By default, the SPB network interface type is set to PtP. To configure a multi-access network interface, the below command can be used on the SPB ISIS interface level. For example:

BEB-1

```
BEB-1> spb isis interface port 1/1/48 type multi-access
```

By default, the priority value for a multi-access interface is set to 64. This value is used to determine which multi-access interface is elected as the DIS. This value can be configured to choose which BEB will be elected as the DIS. For example:

BEB-3

```
BEB-3> spb isis interface port 1/1/48 type multi-access priority 100
```

If there is a tie in priority, the DIS is elected based on the numerically highest MAC address.

If the current DIS fails, another router is immediately elected to play the role, and no backup role is assumed. Periodic database synchronization on broadcast links allows preemption of the existing DIS without significant disruption of IS-IS operation on such media. This implies that an elected router is not guaranteed to remain the DIS if a new switch with a higher priority shows up on the LAN.

The multi-access network can also be an OmniAccess Stellar Mesh network. In the example shown in Figure 39, the Main HQ is running an SPB network, and would like to extend the SPB backbone over a remote location. This can be achieved using an OmniAccess Stellar mesh network.

Figure 39 - SPB over multi-access network – OmniAccess Stellar mesh

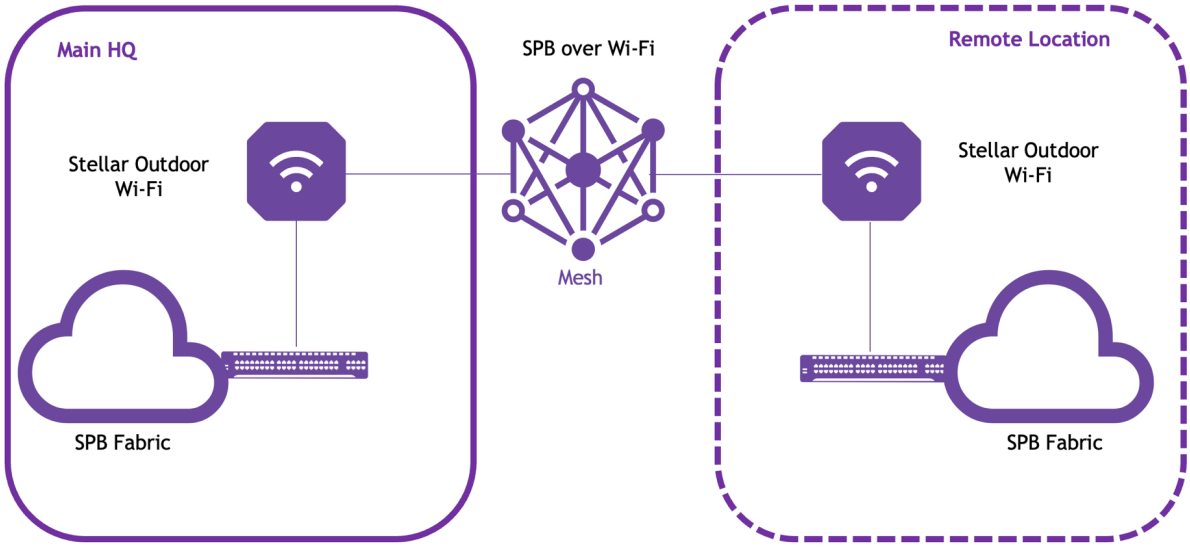

MTU handling of SPB over VXLAN tunneled traffic

When overlay network traffic is tunneled using VxLAN the encapsulated traffic could be dropped by the service provide network if the tunneled traffic exceeds the MTU size supported by the service provider tunnel. The OmniSwitch allows the TCP Maximum Segment Size (MSS) carried in the TCP SYN/SYN-ACK frames to be configured to a value supported by the tunnel. A value in range defined below or a default size profile of sbp (1380) or ethernet (1402) can be configured:

BEB

```
BEB> service 1 sap port 1/1/1:0 tcp-mss overlay-profile spb
```

BEB

```
BEB> service 1 sap port 1/1/1:0 tcp-mss overlay-profile spb
```

ERP over SPB interworking

Ethernet Ring Protection (ERP) over SPB interworking feature can be used to allow seamless connectivity between an access ERP ring and an SPBM aggregation network by enabling ERP in the BEB of the SPB domain. The feature allows ERP protected VLANs to be mapped dynamically and manually to a service on SPB network on the same SAP. For more details on ERP, refer to the Ethernet Ring Protection Switching Application Note referenced in the Related documents section.

There are two supported topologies as shown in Figure 40:

- An ERP ring connecting to a single SPBM backbone (when both sides of the ERP ring end at the same BEB of the SPB domain)
- An ERP ring using an intermediate SPBM network as transport, access ERP ring connected to two different BEBs on the SPB backbone

Figure 40 - ERP over SPB supported topologies

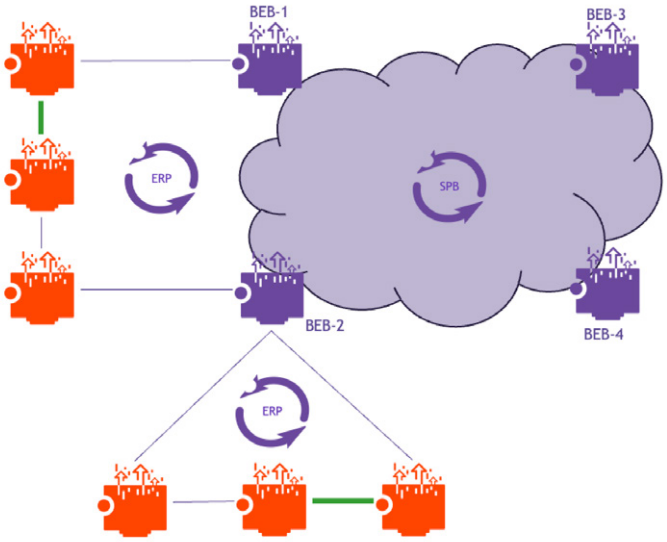

A few points to note:

- Only two ERP type NNI associations are allowed per Service VLAN (SVLAN).
- Configuring an ERP ring on 802.1q tagged port associations with SVLANs is not allowed.
- Configuring an ERP ring on an STP type NNI association with an SVLAN is not allowed.
- BEB cannot be a Ring Protection Link (RPL) node. Hence, RPL port shall not be configured on SPB network. The RPL port cannot be configured as a SAP neighbor.
- For the SPB Service associated with the ERP Service VLAN, it has to be configured in the Control BVLAN of the SPB network. This will ensure reachability to all nodes of the SPB network.
- If the underlay network needs to support more than one ERP ring, then ensure there is no overlap in the VLAN range supported within the ERP rings. That is, each ERP ring must have an exclusive range of VLANs including the service VLAN relative to the other ERP rings.
- In the underlay network, the services association with the ERP VLANs is exclusive to each ring. That is, the service IDs (for example, ISID, VNID) cannot extend across/into the other ERP rings.

Figure 41 - ERP over SPB - single BEB attachment

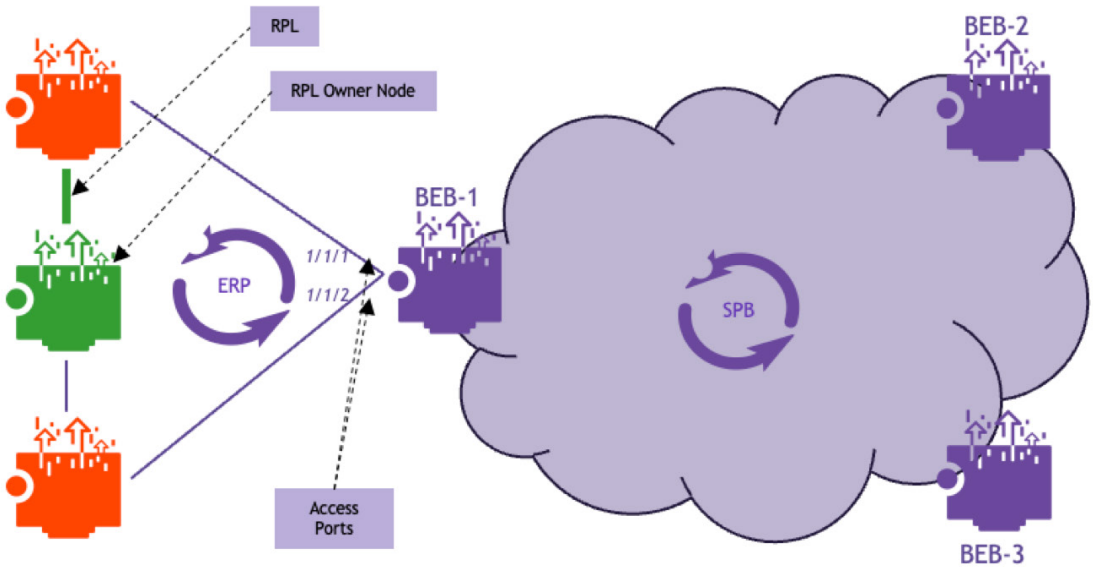

In the initial use case depicted in Figure 41, where an ERP ring connects to a single BEB, ERP is enabled on both access ports. Upon an access port failure, rather than blocking the port when the physical port goes down within a VLAN domain, all SAP associated with the access port shall remain blocked until the Wait To Restore (WTR) timer expiration. Additionally, the MAC address will be flushed on the SAP. Furthermore, when there is a failure in the ERP access ring, the MAC address on the SAP will be flushed in accordance with the ERP protocol.

Below is a sample configuration on BEB-1. If the ERP service VLAN (control VLAN) is mapped to tagged SAP, configure the ERP ring as **access-tagged** using the erp-ring command as following on the SPB domain, with the other required SAP configurations. Otherwise, **access-untagged** can be used.

BEB-1

```
BEB-1> erp-ring 50 port1 access-tagged 1/1/1 port2 access-tagged 1/1/2 service-vlan 1000 level 1
BEB-1> erp-ring 50 enable
```

Figure 42 - ERP over SPB - dual BEB attachment

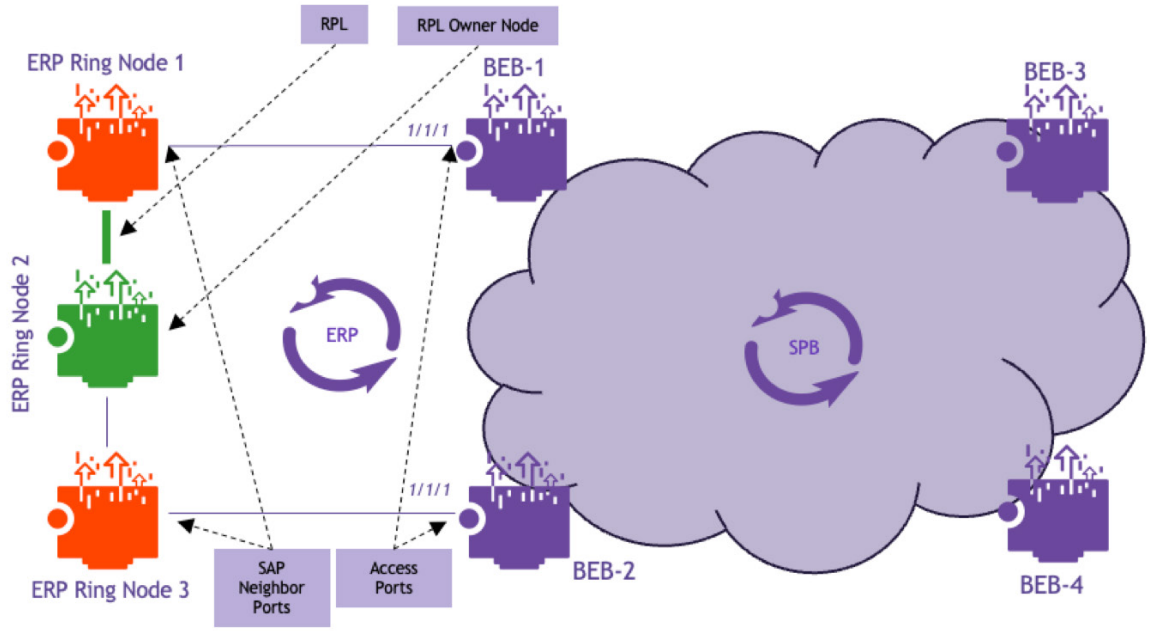

In the second use case shown in Figure 42 where an ERP ring is connected to two different BEBs on the SPB backbone, ERP shall be configured on both BEBs (BEB-1 and BEB-2) that terminate the ERP ring. On both BEBs, one of the ports (port 1/1/1 in this example) shall be the access port, that is, connected to the ERP Ring and the second port shall represent the remote system of the tunnel where the other side of the ERP ring is connected.

On ERP Ring Node 1, the SAP neighbor port is configured for the port that is adjacent to the SAP port of BEB-1. Similarly, on ERP Ring Node 3, the SAP neighbor port is configured for the port that is adjacent to the SAP port of BEB-2. The SAP neighbor cannot be configured as the RPL port.

When there is a failure in the SPB network/tunnel, the ERP RPL node in the ERP ring shall be informed to unblock the RPL port by both BEBs, but the access port on both the end of the BEBs shall be kept forwarding so that client on the ERP ring shall have connectivity to the SPB network.

When the SPB network/Tunnel is up again, instead of blocking the port (in this case SDP/Tunnel) that comes up in regular ERP, until the WTR timer expiry, the access port in one of the BEBs (either on BEB-1 and BEB-2), based on the local system ID and user configured remote system ID, shall be blocked and the SBPM network/Tunnel shall be provisioned. The BEB with higher system ID shall block the access port. The access port also generates a flush since blocking the access port would trigger a topology change.

After the WTR timer expiry on the RPL node, the BEB shall receive an ERP message to unblock the port after which the access port shall be unblocked.

Below is a sample configuration on the ERP Ring Node 1, ERP Ring Node 2 and BEB-1. A similar configuration can be done on ERP Ring Node 3 and BEB-2.

ERP Ring Node 1

```
ERP_NODE1> erp-ring 50 port1 1/1/1 port2 1/1/2 service-vlan 1000 level 1
ERP_NODE1> erp-ring 50 sap-neighbor port 1/1/2
ERP_NODE1> erp-ring 50 enable
```

ERP Ring Node 2

```
ERP_NODE2> erp-ring 50 port1 1/1/1 port2 1/1/2 service-vlan 1000 level 1
ERP_NODE2> erp-ring 50 rpl-node port 1/1/1
ERP_NODE2> erp-ring 50 enable
```

BEB-1

```
BEB-1> erp-ring 50 port1 access-tagged 1/1/1 spb-remote-system dc:08:56:10:72:49 service-vlan 1000 level 1
BEB-1> erp-ring 50 enable
```

A feature which should be considered in ERP over SPB interworking scenarios is the SPB Remote MAC Flush feature.

Refering to Figure 43, consider a network scenario involving an ERP network with a four-node ring topology. Two nodes function as BEBs and are interconnected using SDP tunnels. An ERP mechanism is deployed to prevent loops.

Figure 43 - SPB Remote MAC Flush For ERP

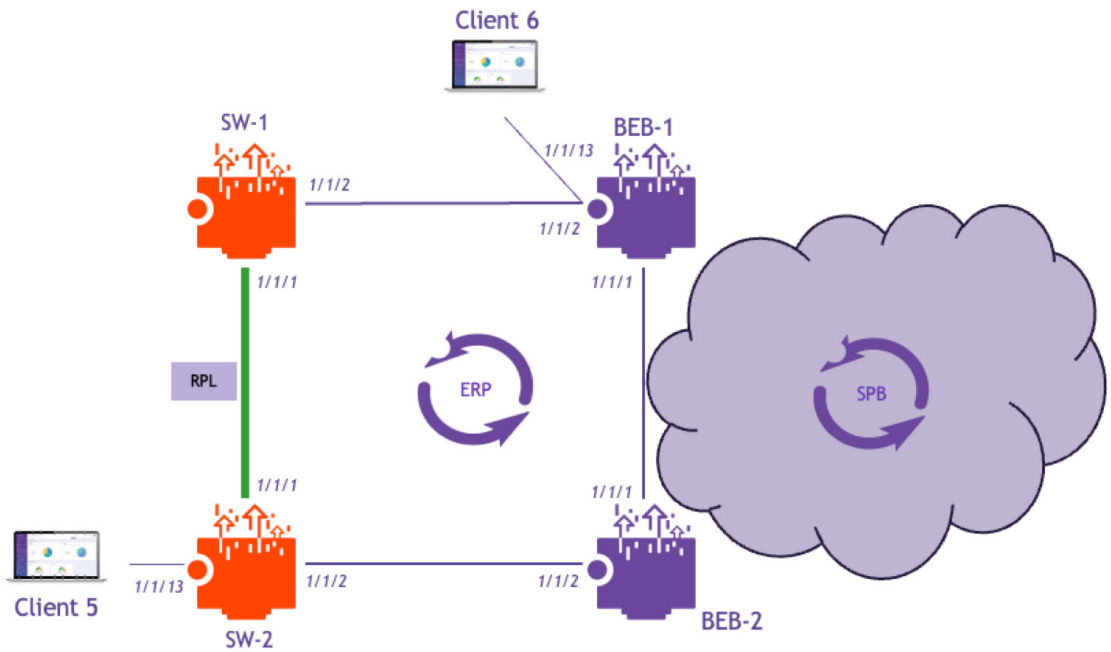

Ping traffic is initiated from client 6 (connected to BEB-1) towards client 5 (connected to SW-2). When the link between BEB-2 and SW-2 (1/1/2) goes operational down, the RPL port (1/1/1 on SW-2) transitions to the forwarding state. At this point, traffic drops entirely, and the same ping fails. BEB-1 receives a signal failure message, which triggers MAC flush events on its access ports (1/1/1 and 1/1/2); however, the MAC addresses on the SDP/tunnel ports are not flushed. As a result, the traffic fails to converge on the alternate path. Despite BEB-1 receiving the message from the RPL port towards BEB-2, only the SAP learned MACs are flushed, while the SDP learned MACs remain in the table. This prevents proper MAC relearning on the new path and leads to continued traffic loss.

When a failure occurs on a port in a BEB node, ERP reacts by forwarding the previously blocked RPL port and triggering a MAC flush across all nodes to remove stale forwarding entries, including SAP and SDP ports, ensuring traffic is re-learned through the new path. However, in this scenario, the MAC flush is not propagating to SDP ports on the BEB nodes, leading to stale MAC entries and subsequent traffic drop. This occurs with unidirectional unicast traffic, where MAC learning does not happen dynamically, leading to persistent stale entries.

Stale MAC entries flush can be achieved by enabling the SPB remote flush feature for MAC flush. This approach ensures that MAC flush events propagate correctly to SDP ports, preventing traffic drops in unidirectional unicast traffic scenarios.

To enable or disable SPB remote flush feature for MAC flush, use the "erp-ring spb-remote-flush" command. By default, SPB remote flush is disabled.

BEB-1

```
BEB-1> erp-ring 50 spb-remote-flush enable
```

SPB L3 VPN route tags

In certain scenarios, route tags are essential for preventing routing loops and implementing administrative routing policies, such as filtering or summarization. This is particularly crucial in dual-homing scenarios where mutual route redistribution occurs between different routing protocols. The Service Attachment Redundancy section covers these scenarios.

The SPB L3 VPN route advertisement feature allows the route-tag to be exchanged and injected across the SPB network.

Below is an example scenario which demonstrates the feature (see Figure 44).

We have an OSPF router which is dual-homed to BEB-1 and BEB-4 and is redistributing a local prefix 10.10.10.0/24. BEB-1 and BEB-4 are doing mutual route redistribution between OSPF and SPB. In order to avoid the prefix 10.10.10.0/24 from being redistributed back from SPB to OSPF, we will use SPB L3 VPN route-tag feature.

In the OSPF router, we are setting the tag of the prefix route 10.10.10.0/24 to 10. The below steps are performed on BEB-1 and BEB-4:

- Enable bidirectional exchange of routes between Customer VRF and ISID 1001 via GRT and set the far-end WAN IP address as the next hop
- Export OSPF routes with tag 10 from Customer VRF into ISID 1001
- Import remote SPB-ISIS routes associated with ISID 1001 from GRT into Customer VRF
- Redistribute imported routes from Customer VRF to OSPF while denying routes with tag 10

These are shown in the snippets below:

Figure 44 - SPB L3 VPN route-tag feature

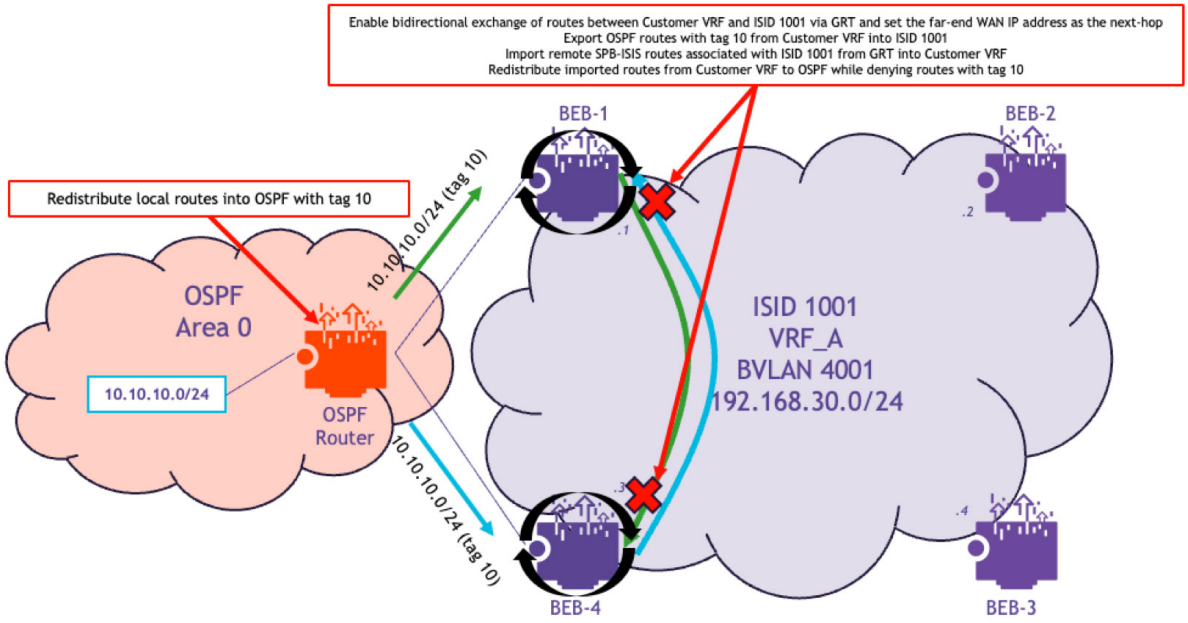

BEB-1

```
BEB-1> vlan 1000 name "OSPF_Router"
BEB-1> vlan 1000 member port 1/1/47 untagged
BEB-1> service 1 spb isid 1001 bvlan 4001
BEB-1> vrf create Customer_A
Customer_A: :BEB-1> ip interface "WAN" address 192.168.30.1 mask 255.255.255.0 service 1
Customer_A: :BEB-1> ip interface "LAN" address 192.168.21.2 mask 255.255.255.0 vlan 1000
Customer_A: :BEB-1> ip load ospf
Customer_A: :BEB-1> ip ospf interface "LAN"
Customer_A: :BEB-1> ip ospf interface "LAN" area 0.0.0.0
Customer_A: :BEB-1> ip ospf interface "LAN" type point-to-point
Customer_A: :BEB-1> ip ospf interface "LAN" admin-state enable
Customer_A: :BEB-1> ip ospf admin-state enable
Customer_A: :BEB-1> ip route-map "route-tag-ospf" sequence-number 10 action permit
Customer_A: :BEB-1> ip route-map "route-tag-ospf" sequence-number 10 match tag 10
Customer_A: :BEB-1> ip route-map "route-tag-ospf" sequence-number 10 match protocol ospf
Customer_A: :BEB-1> ip route-map "deny-route-tag-ospf" sequence-number 10 action deny
Customer_A: :BEB-1> ip route-map "deny-route-tag-ospf" sequence-number 10 match tag 10
Customer_A: :BEB-1> ip route-map "deny-route-tag-ospf" sequence-number 20 action permit
Customer_A: :BEB-1> ip redist import into ospf route-map "deny-route-tag-ospf" admin-state enable
Customer_A: :BEB-1> ip export route-map route-tag-ospf
Customer_A: :BEB-1> ip import isid 1001 all-routes
BEB-1> spb ipvpn bind vrf Customer_A isid 1001 gateway 192.168.30.1 all-routes
```

BEB-4

```
BEB-4> vlan 2000 name "OSPF_Router"
BEB-4> vlan 2000 member port 1/1/47 untagged
BEB-4> service 1 spb isid 1001 bvlan 4001
BEB-4> vrf create Customer_A
Customer_A: :BEB-4> ip interface "WAN" address 192.168.30.4 mask 255.255.255.0 service 1
Customer_A: :BEB-4> ip interface "LAN" address 192.168.22.2 mask 255.255.255.0 vlan 2000
Customer_A: :BEB-4> ip load ospf
Customer_A: :BEB-4> ip ospf interface "LAN"
Customer_A: :BEB-4> ip ospf interface "LAN" area 0.0.0.0
Customer_A: :BEB-4> ip ospf interface "LAN" type point-to-point
Customer_A: :BEB-4> ip ospf interface "LAN" admin-state enable
Customer_A: :BEB-4> ip ospf admin-state enable
Customer_A: :BEB-4> ip route-map "route-tag-ospf" sequence-number 10 action permit
Customer_A: :BEB-4> ip route-map "route-tag-ospf" sequence-number 10 match tag 10
Customer_A: :BEB-4> ip route-map "route-tag-ospf" sequence-number 10 match protocol ospf
Customer_A: :BEB-4> ip route-map "deny-route-tag-ospf" sequence-number 10 action deny
Customer_A: :BEB-4> ip route-map "deny-route-tag-ospf" sequence-number 10 match tag 10
Customer_A: :BEB-4> ip route-map "deny-route-tag-ospf" sequence-number 20 action permit
Customer_A: :BEB-4> ip redist import into ospf route-map "deny-route-tag-ospf" admin-state enable
Customer_A: :BEB-4> ip export route-map route-tag-ospf
Customer_A: :BEB-4> ip import isid 1001 all-routes
BEB-4> spb ipvpn bind vrf Customer_A isid 1001 gateway 192.168.30.4 all-routes
```

The "show ip global-route-table" command output as shown from BEB-1, indicates the prefix with the route-tag "10" learned from the customer VRF and from ISID 1001, which is exported by BEB-4. (The gateway or next-hop is the WAN IP interface of BEB-4.)

BEB-1

```
BEB-1> show ip global-route-table
Type  Source               Destination        Gateway         Metric     Tag
------+--------------------+------------------+---------------+----------+----------
isid   1001                 10.10.10.0/24      192.168.30.4             1          10
isid   1001                 192.168.22.0/24    192.168.30.2             1           0
isid   1001                 192.168.30.0/24    192.168.30.2             1           0
vrf    Customer_A           10.10.10.0/24      192.168.21.1             1          10
```

The "show spb ipvpn route-table" command outputindicates the prefix with the route-tag "10" and the source BEB displayed as well.

BEB-1

```
BEB-1> show spb ipvpn route-table
Legend: * indicates IPVPN route has matching locally configured ISID
SPB IPVPN Route Table:
                                                Source Bridge
     ISID   Destination          Gateway         (Name : BMAC)                          Metric      Tag    Site-Id
----------+------------------+---------------+--------------------------------------+---------+-----------|----------
*    1001   10.10.10.0/24       192.168.30.1    BEB-1              : dc:08:56:10:85:59       1          10  128: 0: 0(0x0)
*    1001   10.10.10.0/24       192.168.30.4    BEB-4              : dc:08:56:10:78:d9       1          10  128: 0: 0(0x0)
*    1001   192.168.22.0/24     192.168.30.2    BEB-2              : dc:08:56:10:80:f9       1           0  128: 0: 0(0x0)
*    1001   192.168.30.0/24     192.168.30.2    BEB-2              : dc:08:56:10:80:f9       1           0  128: 0: 0(0x0)
     1002   192.168.23.0/24     192.168.30.3    BEB-3              : dc:08:56:10:72:49       1           0  128: 0: 0(0x0)
     1002   192.168.30.0/24     192.168.30.3    BEB-3              : dc:08:56:10:72:49       1           0  128: 0: 0(0x0)

Routes: 6
```

This also provides a failover path over the SPB network on BEB-1 and BEB-4 in case the OSPF link fails. Since imported routes have a lower preference than OSPF routes, the BEBs will use the OSPF routes as the primary routes and failover to the SPB imported routes in case of link failure towards the OSPF network.

In the "show ip router database" command output, which displays the routing table, inside the customer VRF context, the prefix 10.10.10.0/24 is showing as an OSPF route learnt from the OSPF router and as an imported route from the GRT.

BEB-1

```
Customer_A::BEB-1> 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: 8

  Destination         Gateway                 Interface              Protocol  Metric     Tag      Misc-Info
--------------------+---------------+-------------------------------+--------+-------+----------+-----------------
+  10.10.10.0/24      192.168.21.1    LAN                            OSPF          1         10
   10.10.10.0/24      192.168.30.4    WAN                            IMPORT        1         10  (backup)
+  127.0.0.1/32       127.0.0.1       Loopback                       LOCAL         1          0
+  192.168.21.0/30    192.168.21.2    LAN                            LOCAL         1          0
+  192.168.22.0/24    192.168.30.2    WAN                            IMPORT        1          0
+  192.168.22.0/30    192.168.21.1    LAN                            OSPF          2          0
+  192.168.30.0/24    192.168.30.1    WAN                            LOCAL         1          0
   192.168.30.0/24    192.168.30.2    WAN                            IMPORT        1          0  (backup)


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

Last, in the OSPF Router, we can verify that the route-maps are filtering any export of routes with a route-tag of "10", and we can see the prefix learned only on the local interface.

OSPF Router

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

Total IPRM IPv4 routes: 6

  Destination         Gateway                 Interface              Protocol  Metric     Tag      Misc-Info
--------------------+---------------+-------------------------------+--------+-------+----------+-----------------
+  10.10.10.0/24      10.10.10.1      OSPF                           LOCAL         1          0
+  127.0.0.1/32       127.0.0.1       Loopback                       LOCAL         1          0
+  192.168.21.0/30    192.168.21.1    BEB-1                          LOCAL         1          0
+  192.168.22.0/24    192.168.21.2    BEB-1                          OSPF          1          0
+  192.168.22.0/30    192.168.22.1    BEB-4                          LOCAL         1          0

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

General design guidelines

Design guidelines have been provided throughout this document. In this section, we provide additional design guidelines to assist the network architect in designing SPB networks.

BVLANs

As described in the Control Plane section, SPB networks load balance traffic on a per-service basis. This load balancing is achieved by mapping different services to different BVLANs. An SPB network supports up to 16 BVLANs, however, most real-world physical topologies do not support 16 equal-cost paths. There is no advantage in creating more BVLANs than the number of equal-cost paths in the physical topology. Moreover, since a SPT must be computed for each BVLAN, having more BVLANs than equal-cost paths in the physical topology creates an additional unnecessary load in the CP, which results in increased resource utilization and convergence times.

In short, only create as many BVLANs as there are equal-cost paths in the physical topology. As of AOS 8.7R1 and later releases, only four BVLANs are created by default when using auto-SPB.

VLAN-to-Service mapping

When creating a SAP, AOS allows mapping of multiple or all VLAN tags to the same SPB service.

As a general guideline, to preserve L2 isolation between VLANs, different VLANs should be mapped to different services (for example, through different SAPs).

Mapping different VLANs to the same SPB service makes inter-VLAN bridging possible, which defeats the purpose of having different VLANs in the first place.

In addition, there is a risk of having duplicate MAC addresses. In theory, there should be no duplicate MAC addresses but, in reality, it can happen, particularly in virtualized environments. Duplicate MAC addresses in different VLANs do not collide, however, if these VLANs are mapped to the same SPB service and the client devices are connected to different SAPs, those MACs will be constantly learned, re-learned and flushed. This is known as a "mac-move" and should be avoided to maintain stability. To avoid mac-move, we strongly recommend mapping different VLANs to different SPB services (ISIDs). This will require one SAP and ISID per access VLAN.

There are some situations in which mapping different VLANs to the same SPB service (ISID) is acceptable, but we will not elaborate on those situations.

In short, as a general guideline, map different VLANs to different SPB services by using specific SAPs for each VLAN.

Virtual chassis

Virtual chassis (VC) is a feature that combines multiple "stackable" switches into a single logical "virtual chassis" such that each physical switch becomes a virtual "slot" in the virtually modular chassis. A virtual chassis is a single logical entity managed as one device and with single control and management planes.

VC provides many benefits such as network architecture and management simplification. It greatly simplifies redundant service attachment. Customer CE access devices can be dual-homed to diverse slots in a BEB through a link aggregate. This eliminates the need to configure other L2 or L3 redundancy mechanisms such as DHL or VRRP.

When using VC in the SPB backbone, logical link aggregates (LAGs) are recommended to interconnect the VC to all its SPB neighbours such that one member (physical) port connects to every slot in the VC as seen in Figure 45. This is not mandatory but is recommended and will improve the network convergence time in the event of slot failure because the need to update tables during the control plane takeover is greatly reduced. In addition, dual homing nodes to a VC reduces the need to forward traffic across the Virtual Fabric Link (VFL) because traffic forwarding in a LAG prioritizes the use of local linkagg member ports over remote (across the VFL) member ports.

Figure 45 - VC and SPB

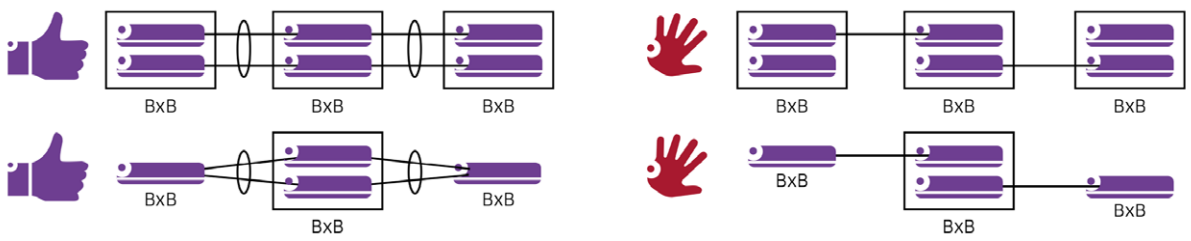

Link aggregation

Combining multiple physical links into a LAG improves resiliency and increases total available bandwidth on the logical link.

In a LAG, traffic is load balanced across member ports in one of two ways:

- MAC hash (brief mode)
- IP + TCP/UDP port hash (extended mode)

However, SPB backbone ports use MAC-in-MAC encapsulation which means MAC addresses are the BMACs of BEB and BCB nodes while IP addresses and port numbers are not visible to the hashing logic. In most cases this does not create enough entropy and the load will not be spread evenly across all different physical links.

Since AOS 8.3.1R01, a "tunnel-protocol" option can be selected such that the hashing can use CMACs or IP addresses + TCP/UDP ports.

It is recommended that this option be enabled on all SPB nodes using LAG. The choice of MAC (brief) or IP+TCP/UDP ports (extended) is a global setting which will apply to all LAGs. Refer to the OmniSwitch AOS CLI Reference User Guide referenced in the Related documents section for further details.

Link metric

SPB uses the link metric as a measure of a link's cost to reach another node. By default, all link metrics are set to 10 regardless of link speed. The link metric is an integer in the 1-16M range.

The link metric can be adjusted to influence the SPT calculations. For instance, the metric can be changed to reflect the link speed. It should be noted that the metric must be adjusted on both sides of a link. Nodes will become adjacent even when the metrics are different, but the highest metric will be used in the SPT calculations.

Changing the link metric to reflect the link speed will help steer traffic towards links with higher capacity and away from lower capacity ones, making the best use of the total available bandwidth and improving performance. The following table shows a way in which the metric can be set to be inversely proportional to the link speed.

Speed	Suggested Metric
800G	125
400G	250
200G	500
100G	1000
50G	2000
40G	2500
25G	4000
10G	10000
1G	100000
100M	1000000

QoS

In an SPB network, traffic is classified at the SAP and the classification does not change as traffic traverses the backbone and until it exits through another SAP at the destination BEB.

SAPs on an SPB Service Access Port can be configured to be trusted or un-trusted. Trusted SAPs copy CoS markings from the incoming VLAN tag onto the BVLAN tag. In case the SAP is trusted and the access port is double-tagged, then it is derived from the outer VLAN tag. If incoming traffic is not tagged, then the port's default priority is used. Un-trusted SAPs set the CoS markings to a user-defined value.

By default, the SAP is trusted with the priority set to best effort (zero). These default values are set when a port is configured as an access port and then associated with the SAP.

No further classification based on inner L2-L4 conditions is possible within the SPB backbone due to the MAC-in-MAC encapsulation.

Untagged L2 Control Packets (BPDU, GVRP, AMAP, etc...) are always tunneled (if enabled) through the SPB domain with the default EXP bits set at 7, so that they can arrive at the destination CPU at highest priority of 7. Trusted/untrusted SAPs configured on the Access Ports will not affect the L2 Control Packets priority assignment on the Access ports.

Security guidelines

This section provides additional design guidelines specific to the security domain. This is not an exhaustive list of recommendations, but focuses on certain guidelines specific to SPB deployments. We will go through different AOS features and how they can be used to improve security in an SPB network. Other more general security guidelines can be found in the Security Guidelines document referenced in the Related documents section.

Management VRF

As explained in the Non-IP core section, SPB relies on a non-IP protocol for path computation. For this reason, BCB nodes and BEB nodes supporting L2 services only do not require an IP address. The only case in which an SPB node requires an IP address is the case of a BEB node supporting a L3 service or feature such as L3 VPN, VPN Lite, or VRRP, among others.

The Management section covered different SBP management options. Management IP addresses can be bound to:

- The EMP port, in case of OOBM
- A standard VLAN port, in the case of OOBM
- The control BVLAN, in the case of in-band management
- A management SPB service

No matter what management option is chosen, management IP addresses should use a different VRF from the VRF used for service or customer traffic. This is already the case when using the EMP port for OOBM. One possibility is creating a dedicated management VRF and enabling the required management protocols on this VRF as shown in the Management section configuration snippets.

Another possibility is using the default VRF for management, under the condition of not using it for anything other than management.

MACsec

Data integrity and confidentiality must be protected while in transit through the network. MACsec is an IEEE standard (802.1AE) which provides point-to-point authentication and optional encryption between MACSec-capable devices such as switches. MACsec can prevent various threats such as man-in-the-middle (MiTM), sniffing, spoofing, and playback attacks.

Because MACsec operates at the MAC layer, it transparently secures all upper layer traffic transiting through MACsec-enabled links. This includes both application-layer data, as well as control-plane and management-plane communication. In addition, unlike IPsec, MACsec is implemented in hardware at wire-speed and does not introduce additional latency or bandwidth limitations.

NAC

The Dynamic SAPs section explained how users and devices can be dynamically mapped to their services based on their identity. Enabling authentication on every front-panel port ensures only authorized users and devices can access network services. One additional benefit of creating dynamic SAPs through NAC is that no service is instantiated on a BEB until an authorized user successfully authenticates and is mapped to the service: The service is instantiated on demand.

This is an additional layer of security compared to static SAPs because no service is connected if no authorized user is connected. It is clearly more difficult to hack, attack or otherwise disrupt a service when it is not even connected.

Router authentication

As explained in the VPN Lite section, an SPB network can exchange routes with external non-SPB entities by using the VPN Lite feature. This means that one or more SPB BEB nodes will run a routing protocol such as OSPF or BGP with external entities. Any learnt route may be imported into the SPB backbone and propagated to other BEB nodes by way of IS-IS TLVs.

This creates an opportunity for a bad actor to inject malicious routes and poison the routing table to carry out DoS, MiTM or other attacks.

This risk can be mitigated by enabling routing protocol authentication (for example, MD5 for OSPF or BGP).

Learned Port Security

The Learned Port Security (LPS) feature allows the switch administrator to ensure that only certain hosts are capable of connecting to the switch. LPS provides control on the maximum number of MACs that can be learned on a physical port or SAP (based on configuration) and then configure the port/SAP to go into a restricted or shutdown mode when that threshold is exceeded. Additional LPS functionality allows the user to specify how the LPS port handles unauthorized traffic. The following options are available for this purpose:

- Block traffic that violates LPS port restrictions; authorized traffic is forwarded on the port.
- Disable learning on the LPS port when unauthorized traffic is received.
- Administratively shutdown the LPS port when unauthorized traffic is received; all traffic is stopped.

LPS also provides a learning time window to learn MACs and convert them to static entries and to be able to save them to the configuration file for use later on. This allows an administrator to control what specific MACs are allowed on a port/SAP and to save that configuration for later. LPS also provides logging and notification if a rule violation occurs.

LPS is configured differently for static and dynamic SAPs. A static SAP and its mapping to the SPB service must be pre-configured before LPS can be configured on the respective SAP. The below snippet shows an example of static SAP LPS configuration:

BEB-1

```
BEB-1> service 100 sap port 1/1/1:10
BEB-1> port-security sap 1/1/1:10 maximum 1000
```

If LPS is to be configured on a dynamic SAP, the dynamic SAP (and its service mapping) on the UNP port, must pre-exist by creating it upfront through persistent-profile configuration on the UNP access port or linkagg. Statically assigning a UNP service profile to a UNP port creates a persistent SAP that will not age out when there is no activity on the port. This solution is particularly useful for access to silent devices in the UNP service domain.

DHCP Snooping

Using DHCP Snooping improves network security by filtering DHCP messages received from devices outside the network and building and maintaining a binding table (database) to track access information for such devices.

In order to identify DHCP traffic that originates from outside the network, DHCP Snooping categorizes ports as either trusted or untrusted. A port is trusted if it is connected to a device inside the network, such as a DHCP server. A port is untrusted if it is connected to a device outside the network, such as a customer switch or workstation.

DHCP Snooping can be enabled either on global level or service level. SAP ports will be client only ports and all the SDP ports will be trust ports in a service by default.

BEB-1

```
BEB-1> dhcp-snooping service 23 admin-state enable
```

OmniFabric

Alcatel-Lucent Enterprise OmniFabric is a multi-technology network fabric ensuring end-to-end security in a Zero Trust Network architecture with automated segmentation for both IT and OT environments.

Supporting SPB, MPLS and EVPN within a single AOS platform, OmniFabric network fabric provides unmatched flexibility, integrating seamlessly into diverse vendor ecosystems to prevent lock-in.

OmniFabric's advanced automation reduces manual tasks, minimizes errors and simplifies complex network management. It provides a cost-effective solution that includes essential cybersecurity measures, such as micro-segmentation and AI-powered analytics, at no extra charge. From smart buildings to critical infrastructure, OmniFabric adapts to various use cases, delivering reliable, end-to-end security that lowers TCO while enabling IT/OT convergence.

OmniFabric provides many benefits, including:

- **Multi-technologi Integration**: OmniFabric is the only solution in the market that supports SPB, MPLS and EVPN within the same AOS operating system, offering businesses unparalleled flexibility, performance and reliability.
- **Enhanced cybersecurity**: OmniFabric delivers robust cybersecurity measures to protect data integrity and prevent unauthorized access, with support for Zero Trust Networks and micro-segmentation.
- **Built-in automation**: Advanced automation features streamline network operations, reducing manual intervention, minimizing human error and simplifying the use of multiple technologies.
- **IoT connectivity**: IoT devices are automatically detected, classified and contained in virtual segments. This enables operational technology teams to connect devices to the network without increasing exposure to cyberattacks. This capability is included at no extra charge.
- **Flexibility and interoperability**: OmniFabric supports interoperability in brownfield environments

where equipment from multiple vendors coexist. It adapts to any underlying architecture, from edge to data center, offering increased choices and freedom to customers, eliminating vendor lock-in.

- **Simplified operations**: OmniFabric ensures an easy learning curve and management with all protocols integrated into the same AOS and monitored with a single pane of glass, augmented with AI-powered analytics. This is particularly beneficial for customers with limited resources.
- **Optimized TCO**: With no hidden fees, simple procurement, ease of learning and unified management through Alcatel-Lucent OmniVista, OmniFabric offers reduced TCO.
- **Customizable solutions**: OmniFabric offers the choice of multiple technologies to be used depending on the area or architecture—SPB in campus networks, EVPN in data centers and MPLS in metropolitan area networks (MAN).

Conclusion

Shortest Path Bridging is a powerful technology yet simple when compared to others such as MPLS or EVPN. SPB is broadly supported across the Alcatel-Lucent OmniSwitch portfolio with products in multiple formats, from stackable to modular chassis and even industrial-grade ruggedized variants. This product breadth, coupled with SPB's service-oriented framework, results in a network architecture that can deliver the required service to the right location with minimal network configuration changes or even in a fully automated manner.

List of abbreviations

ACL	Access Control List
AG	Access Guardian
AOS	Alcatel-Lucent Operating System
ASIC	Application-Specific Integrated Circuit
B-DA	Backbone Destination Address
B-SA	Backbone Source Address
B-VID	Backbone VLAN ID
BCB	Backbone Core Bridge
BDR	Backup Designation Router
BEB	Backbone Edge Bridge
BGP	Border Gateway Protocol
BMAC	Backbone MAC
BPDU	Bridge Protocol Data Unit
BSN	Base Service Number
BUM	Broadcast, Unknown Unicast and Multicast
BVLAN	Backbone VLAN
CCM	Continuity Check Messages
CCTV	Closed-Circuit Television
CE	Customer Edge
CFM	Connectivity Fault Management
CMAC	Customer MAC
CMM	Chassis Management Module
CP	Control Plane
CST	Common Spanning Tree
CVLAN	Customer VLAN
DHCP	Dynamic Host Configuration Protocol
DHL	Dual-Home Link
DIS	Designated Intermediate System
DNS	Domain Name System
DoS	Denial of Service
DP	Data Plane

DR	Designated Router
ECT	Equal-Cost Tree
EVPN	Ethernet VPN
FDB	Forwarding Database
GRT	Global Routing Table
HTTPS	HyperText Transfer Protocol Secure
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IEFT	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IoT	Internet of Things
IP	Internet Protocol
IPMS	IP Multicast Switching
IS-IS	Intermediate System to Intermediate System
ISID	Instance Service Identifier
ISO	International Organization for Standardization
LAN	Local Area Network
LBD	Loopback Detection
LBM	Loopback Message
LBR	Loopback Reply
LDP	Label Distribution Protocol
LPS	Learned Port Security
LSB	Least Significant Bit
LSDB	Link State Database
LSP	Link State Packets
LTM	Link Trace Message
LTR	Link Trace Reply
MAC	Media Access Control
MD5	Message-Digest Algorithm 5
MEF	Metro Ethernet Forum
MEP	Maintenance End Point
MIP	Maintenance Intermediate Point

MiTM	Man-in-The-Middle
MP-BGP	Multi-Protocol BGP
MPLS	Multiprotocol Label Switching
MSB	Most Significant Bit
MSTP	IEEE 802.1s Multiple Spanning Tree Protocol
NAC	Network Admission Control
NLPID	Network Layer Protocol Identifier
NMS	Network Management System
NNI	Network-to-Network Interface
OAM	Operations, Administration and Maintenance
OOBM	Out Of Band Management
OSPF	Open Shortest Path First
OUI	Organizationally Unique Identifier
PBB	IEEE 802.1ah Provider Backbone Bridging
PE	Provider Edge
PtP	Point-to-Point
Q-in-Q	IEEE 802.1ad Provider Bridging
QoS	Quality of Service
RADIUS	Remote Access Dial-In User Service
RBAC	Role-Based Access Control
RFC	Request For Comments
ROI	Return on Investment
RPFC	Reverse Path Forwarding Check
RPL	Ring Protection Link
RSTP	IEEE 802.1w Rapid Spanning Tree Protocol
RTT	Round Trip Time
SAA	Service Assurance Agent
SAP	Service Access Point
SD-LAN	Software-Defined LAN
SDN	Software-Defined Networking
SDP	Service Distribution Point
SNMP	Simple Network Management Protocol

SPB	IEEE 802.1aq Shortest Path Bridging
SPB-M	SPB MAC-in-MAC
SPB-V	SPB Q-in-Q
SPF	Shortest Path First
SPT	Shortest Path Tree
STP	IEEE 802.1D Spanning Tree Protocol
SVLAN	Service VLAN
TCO	Total Cost of Ownership
TLV	Type, Length, Value
UNI	User-Network Interface
UNP	Universal Network Profile
VFL	Virtual Fabric Link
VM	Virtual Machine
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VRRP	Virtual Router Redundancy Protocol
WAN	Wide Area Network
WTR	Wait To Restore

Related documents

[1] IP/IPVPN services with IEEE 802.1aq SPB networks - draft-unbehagen-spb-ip-ipvpn-00.txt - https://datatracker.ietf.org/doc/html/draft-unbehagen-spb-ip-ipvpn-00

[2] RFC 6329 - IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging - https://datatracker.ietf.org/doc/html/rfc6329

[3] OmniSwitch AOS Release 810R04 CLI Reference User Guide

[4] OmniSwitch AOS Release 810R04 Network Configuration User Guide

[5] OmniSwitch AOS Release 810R04 Switch Management User Guide

[6] OmniSwitch AOS Release 810R04 Specifications Guide

[7] Network Infrastructure Solutions Security Best Practices – https://al-enterprise.com/-/media/assets/internet/documents/network-infrastructure-solution-security-tech-brief-en.pdf

[8] Ethernet Ring Protection Switching Application Note - https://www.al-enterprise.com/-/media/assets/internet/documents/ethernet-ring-protection-switching-application-note-en.pdf