Shortest Path Bridging Architecture guide

Tech Brief
Shortest Path Bridging Architecture guide

Alcatel·Lucent
Enterprise

Table of Contents

1. About this architecture guide

1.1 Purpose

The purpose of this architecture guide is to present SPB (802.1aq) networking concepts along with design and deployment guidelines. It does not attempt to cover every aspect, nor every possible architecture option, only the most common, validated and recommended architectures. You are encouraged to refer to the Alcatel-Lucent Operating Software (AOS) documentation for additional details, options and guidelines.

1.2 Audience

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

1.3 Glossary

AG	Access Guardian
BCB	Backbone Core Bridge
B-DA	Backbone Destination Address
BEB	Backbone Edge Bridge
BGP	Border Gateway Protocol
BMAC	Backbone MAC
B-SA	Backbone Source Address
BSN	Base Service Number
B-VID	Backbone VLAN ID
BVLAN	Backbone VLAN
CMAC	Customer MAC
CP	Control Plane
DoS	Denial of Service
DP	Data Plane
ECT	Equal-Cost Tree
FDB	Forwarding Data Base
IETF	Internet Engineering Task Force
iFab	Intelligent Fabric
IGP	Interior Gateway Protocol
ISID	Instance Service Identifier
IS-IS	Intermediate System to Intermediate System
LDP	Label Distribution Protocol
MAC	Media Access Control
MACs	Moves Adds and Changes
MP-BGP	Multi-Protocol BGP

MSTP	IEEE 802.1s Multiple Spanning Tree Protocol
NAC	Network Admission Control
OSPF	Open Shortest Path First
PBB	IEEE 802.1ah Provider Backbone Bridging
Q-in-Q	IEEE 802.1ad Provider Bridging
RADIUS	Remote Access Dial-In User Service
ROI	Return on Investment
RSTP	IEEE 802.1w Rapid Spanning Tree Protocol
SAP	Service Access Point
SDP	Service Distribution Point
SPB	IEEE 802.1aq Shortest Path Bridging
SPB-M	SPB MAC-in-MAC
SPB-V	SPB Q-in-Q
SPF	Shortest Path First
STP	IEEE 802.1D Spanning Tree Protocol
TLV	Type, Length, Value
UNP	User Network Profile

1.4 References

[1] IP/IPVPN services with IEEE 802.1aq SPB networks - draft-unbehagen-spb-ip-ipvpn-00.txt

[2] Alcatel-Lucent OmniSwitch® Template Based Provisioning with Alcatel-Lucent OmniVista® 2500 Network Management System (NMS)

[3] Network infrastructure security best practices

2. The network needs to evolve

Local Area Networks (LAN) 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 transiting the root-bridge instead of alternative better routes 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. The problems with STP

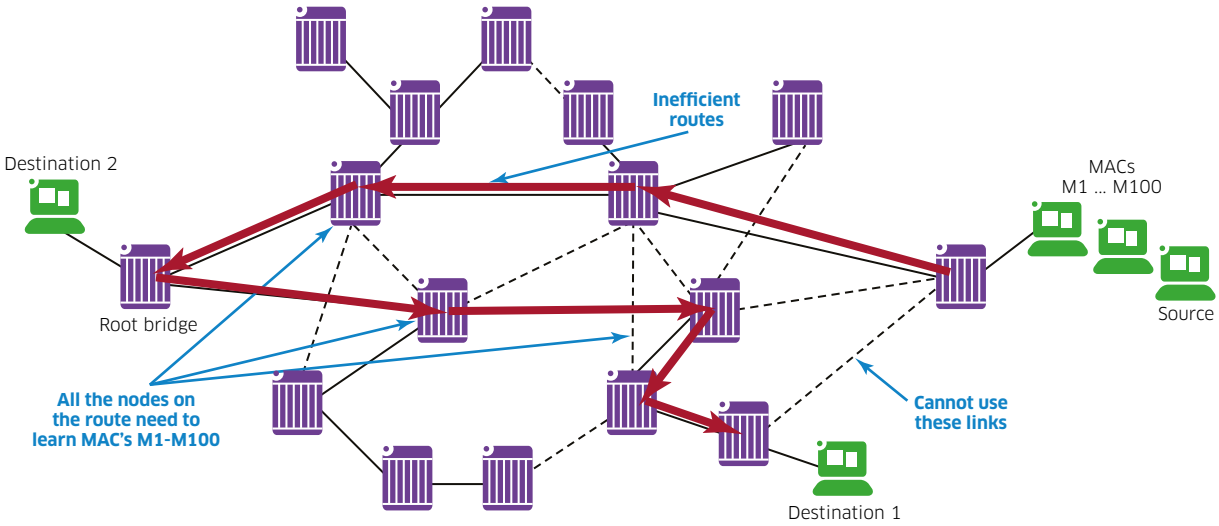

In addition to STP's weaknesses, Ethernet's scalability beyond the LAN is limited by its lack of a coordinated control plane and use of a flat (as opposed to hierarchical) address space. Legacy Ethernet networks present the following challenges:

- **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 thus posing a scalability challenge

Lastly, IEEE 802.1ad (Provider Bridging, or Q-in-Q) is limited to a maximum of 4096 service instances.

3. Introducing SPB

802.1aq Shortest Path Bridging (SPB) is an IEEE networking standard whose primary focus was addressing the challenges in STP. But SPB is much more than STP's evolution: SPB provides MPLS-like VPN services but is significantly simpler to deploy and maintain. And unlike MPLS, which requires a "stack" of protocols (for example: LDP, OSPF, MP-BGP, among 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 Intelligent Fabric (iFab) brings further simplification by automating network node provisioning, 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). Let's analyse SPB's benefits in further detail.

3.1 Scalable, fast-converging, multi-path fabric

Figure 2. Addressing STP's challenges

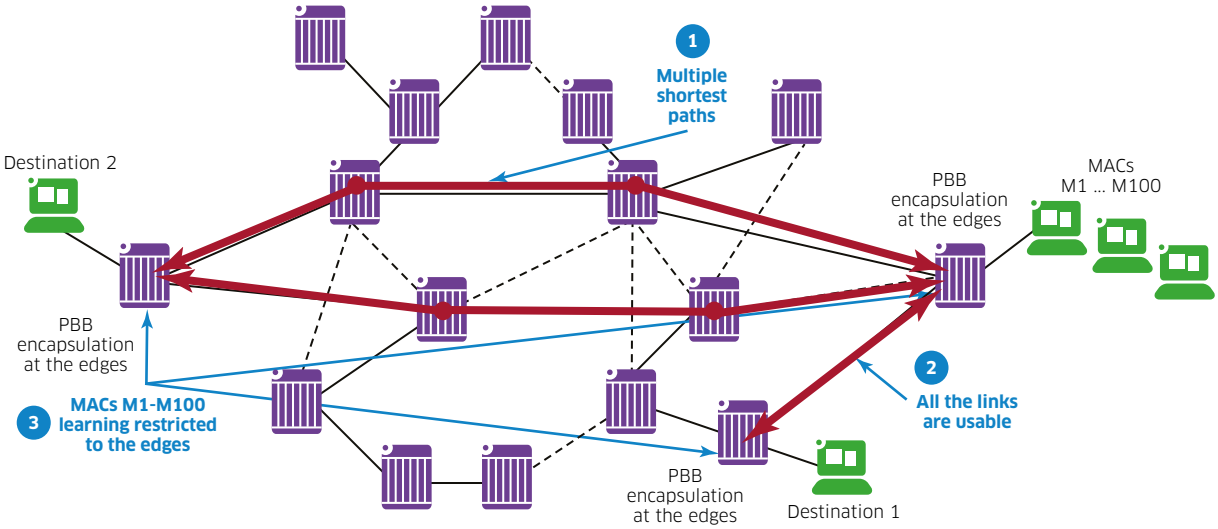

SPB's loop-free topology is built by a link-state routing protocol running Dijkstra's Shortest Path First (SPF) algorithm: IS-IS. With IS-IS, no network link is disabled, all paths are available and traffic between any pair of nodes follows the shortest path. In addition, with MAC-in-MAC encapsulation,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

3.2 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 centres, to name a few. 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 16M services.

Figure 3. Multi-tenancy

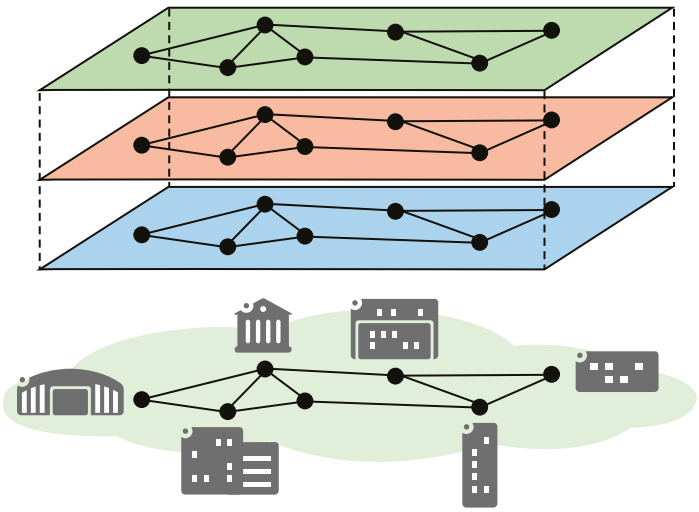

3.3 Dynamic service instantiation

SPB services do not need to be statically bound to a switch port. SPB is tightly integrated with Alcatel-Lucent Enterprise's classification and Network Admission Control (NAC) framework known as 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/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 Machines (VMs) migrations without need for Move, Add or Change requests
- **Increased security**: Services are instantiated on an as-needed basis only, 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, DoSd, or otherwise hacked, while they're not active.
- **Device templates**: This dynamic instantiation of network services easily lends itself into template-based configuration of network nodes. Edge nodes can all share the same base configuration template and dynamically adjust the service configurations on the fly.

3.4 Edge-only service provisioning

Whether statically or dynamically instantiated, SPB services need only be provisioned 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.

3.5 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 User 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. Let's take CCTV cameras as an example: ACLs contained in the UNP can allow communication between the camera and surveillance servers but at the same time block camera-to-camera communication, preventing the spread of malware, "pivoting" and other hacking techniques which rely on lateral movement.

Figure 4. Micro-segmentation

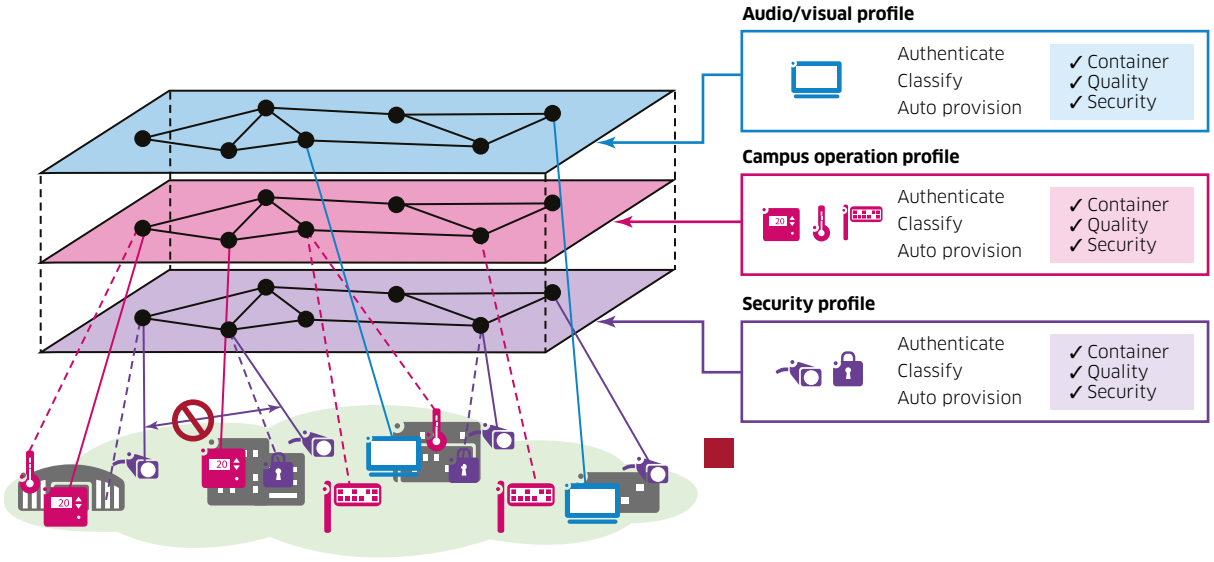

3.6 Non-IP core

Even when providing L3 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 VRF, not in-line with user traffic.

4. The Data Plane: IEEE 802.1ah Provider backbone bridging

The Data Plane's (DP) mission is to forward 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 utilizes IEEE 802.1ah Provider Backbone Bridging (PBB), aka MAC-in-MAC, encapsulation. The PBB header includes de following fields:

B-VID: Or Backbone VLAN (BVLAN) ID. A VLAN that serves as a transport VLAN for the SPB service instances and to connect SPB bridges together through 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: Or Backbone source and destination MAC addresses. 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 learnt 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 the case of a double-tagged (Q-in-Q) frame. Note that MAC and BMAC addresses are shortened to 2 bytes for simplicity in this diagram.

Figure 5. PBB Data Plane

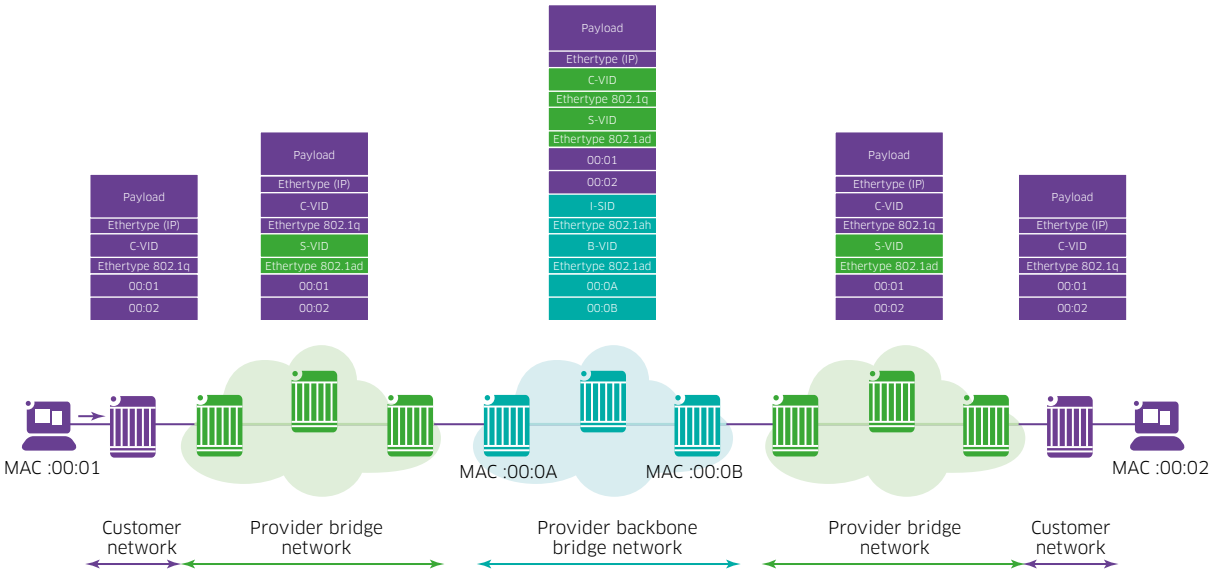

Let's define a few key terms.

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 and this makes SPB networks more scalable and stable (CMACs are not learnt and therefore do not need to be flushed and re-learnt 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.

5. The Control Plane: RFC 6329 IS-IS Equal-cost trees

As stated earlier, the role of the CP is to populate the FDB tables used by the DP. SPB uses IS-IS, or Intermediate System to Intermediate System (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 Shortest Path 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). In a nutshell, these extensions add support for multiple topologies, allowing load sharing over multiple equal-cost paths, and service-membership discovery, or in other words: Communicating 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 for instance, none of which is the root bridge, cannot use the direct single-hop path because that link is disabled by STP. Traffic between these two nodes must take a 3-hop detour traversing the root bridge.

Figure 7. Multiple trees

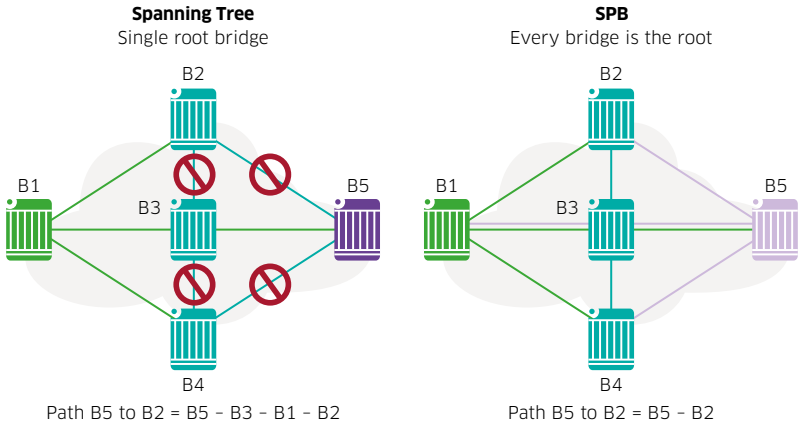

In contrast, when using SPB, no link is disabled: each node is the root of its own tree. Nodes B2 and B5 can simply communicate over the direct single-hop path while at the same time they can communicate with other nodes over different paths (for example; between B4 and B5). SPB's support for multiple 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 a 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 and this path 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. And, those paths 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

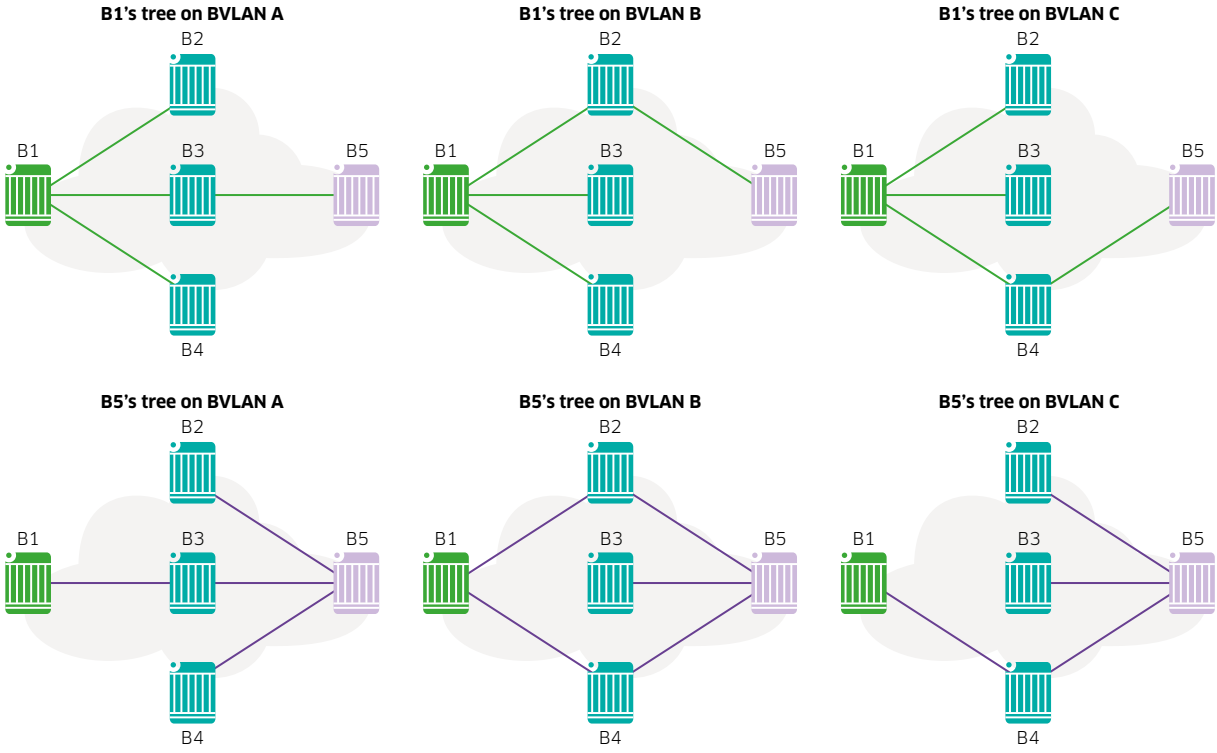

The trees shown in figure 8 are SPB's equal-cost trees (ECTs). Each node builds a tree per BVLAN and the cost to reach other nodes is the same across all BVLANs. The ECT-ID is a number assigned to each BVLAN at the time of BVLAN creation and is used for tie breaking during path computation. Assigning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

Another important property of SPB networks is path symmetry. If you closely examine the picture above, you will notice that 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 be easily 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, per-BVLAN load balancing

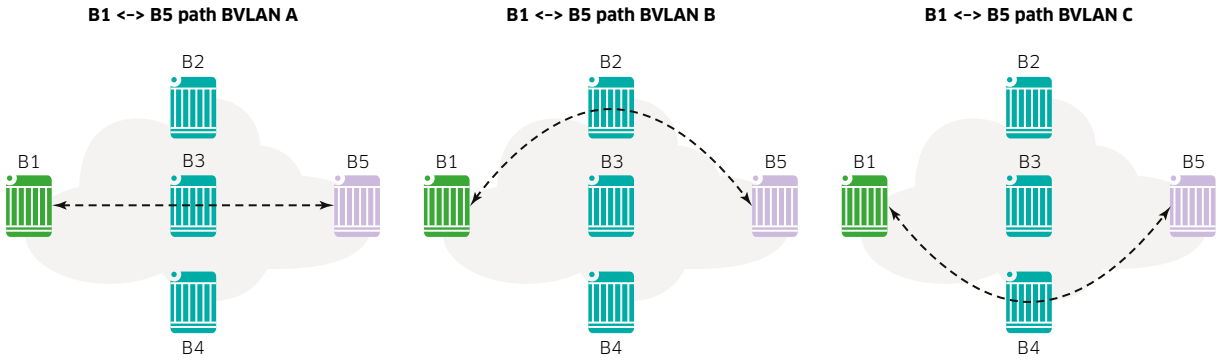

The result of IS-IS path computation for each BVLAN and node is the FDB which is used by the data plane for frame forwarding. Figure 10 shows BEB5's unicast FDB. The multicast FDB will be discussed in Section 7.

Figure 10. B5's Unicast FDB

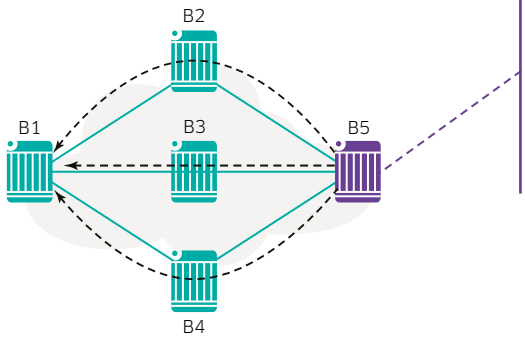

BVID	Node	Outbound port
BVID A	B1	Port 1
BVID B	B1	Port 2
BVID C	B1	Port 3
BVID A	B2	Port 1
BVID B	B2	Port 1
BVID C	B2	Port 1
BVID A	B3	Port 2
BVID B	B3	Port 2
BVID C	B3	Port 2
BVID A	B4	Port 3
BVID B	B4	Port 3
BVID C	B4	Port 3

6. The service framework

An SPB service represents a VPN, or tenant, and is uniquely identified by its service identifier, the ISID. An SPB service needs only 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 such that all SPB nodes have a consistent view of the services which are active on each BEB. Each node then builds a service database.

Figure 11. The service database

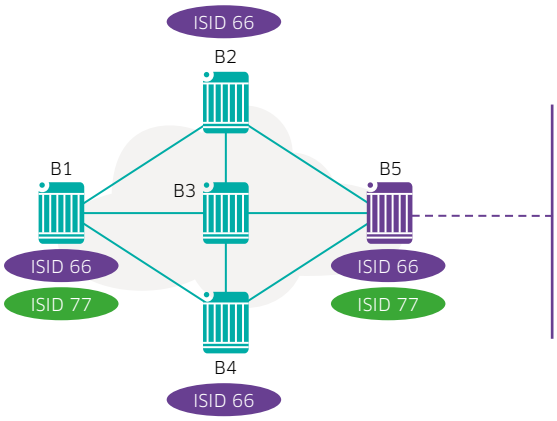

ISID	BVID	Node
66	BVID A	B1
66	BVID A	B2
66	BVID A	B4
66	BVID A	B5
77	BVID B	B1
77	BVID B	B5

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

Service Access Point: The SAP is a UNI-side logical port which 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: The SDP is an NNI-side logical port which 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.

Let's look at figure 12. In this diagram, 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. The service framework

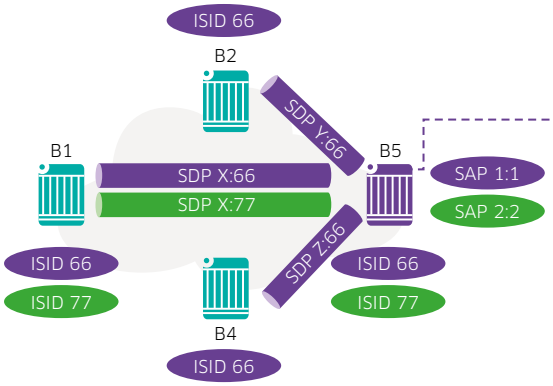

Identifier	BVID	Node
SAP 1:1	–	–
SAP 2:2	–	–
SDP 32769: 66	BVID A	B1
SDP 32768: 66	BVID A	B2
SDP 32767: 66	BVID A	B4
SDP 32766: 77	BVID B	B1

It should be noted that 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.

Let's expand this example by adding some end customer sites and CMACs associated to those customers. We will keep using 2-byte MAC addresses 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 next section.

Figure 13. Customer MAC address learning

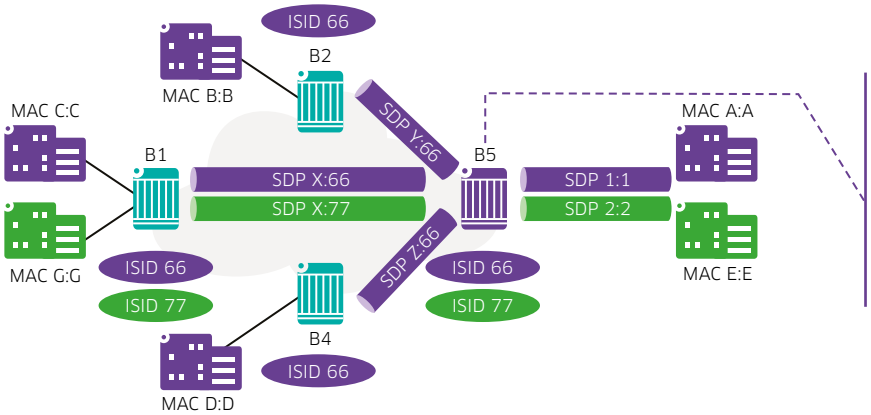

Identifier	ISID	CMAC
SAP 1:1	66	A:A
SAP 2:2	77	E:E
SDP 32769: 66	66	C:C
SDP 32768: 66	66	B:B
SDP 32767: 66	66	D:D
SDP 32766: 77	77	G:G

7. BUM traffic

SPB supports 3 BUM (broadcast, unknown unicast, and multicast) traffic replication and forwarding methods:

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. Head-end replicated BUM traffic simply uses the unicast FDB and therefore travels along the same path. This property is known as congruency.

Figure 14. Head-end BUM replication

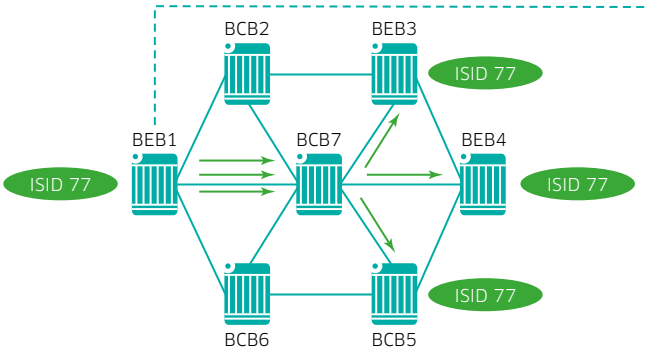

Type	B-MAC	Port	BVID
U	00:02	Port 2	BVID B
U	00:03	Port 7	BVID B
U	00:04	Port 7	BVID B
U	00:05	Port 7	BVID B
U	00:06	Port 6	BVID B
U	00:07	Port 7	BVID B

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. 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 B 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 B 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 15. Tandem (S,G) BUM replication

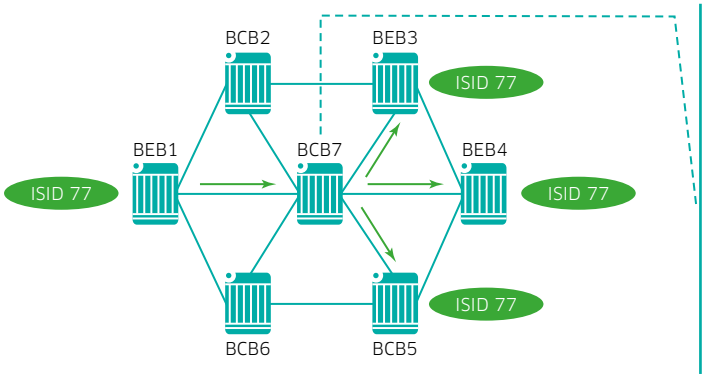

Type	B-MAC	Port	BVID	Out Intf
U	00:01	Port 1	BVID B	
U	00:02	Port 2	BVID B	
U	00:03	Port 3	BVID B	
U	00:04	Port 4	BVID B	
U	00:05	Port 5	BVID B	
U	00:06	Port 6	BVID B	
M	00:0W	Port 1	BVID B	Port 3/4/5
M	00:0X	Port 3	BVID B	Port 1/5
M	00:0Y	Port 4	BVID B	Port 1
M	00:0Z	Port 5	BVID B	Port 1/3

Tandem (*,G): In this mode, a separate multicast tree is created. This tree is not a Shortest Path tree 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 B 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, however, it can be a good option when all traffic is sourced or destined towards the root bridge.

Refer to Table 1 to compare 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 required to inter-operate with third-party equipment

* When combined with IGMP Snooping.

8. Creating an SPB backbone

In this section, we provide a sample SPB Backbone configuration and refer to figure 16 as a sample topology. We will continue using this sample topology throughout the rest of this document. Nodes BEB-1 through BEB-4 are called "BEB" nodes because we will add services to these nodes later. Node BCB will remain a pure transit node and not terminate any service.

If you observe this topology, you will notice that it provides up to 3 shortest paths, for example, between nodes BEB-1 and BEB-3, or between nodes BEB-2 and BEB-4. To take advantage of those 3 diverse paths for traffic load balancing, we need to create a minimum of 3 BVLANs. In this example, we will however, dedicate one BVLAN purely for control traffic and therefore we will create a total of 4 BVLANs. However, it should be noted that this is not strictly necessary, the control BVLAN can also be used for services.

Backbone configuration entails the following 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 16. Sample backbone topology

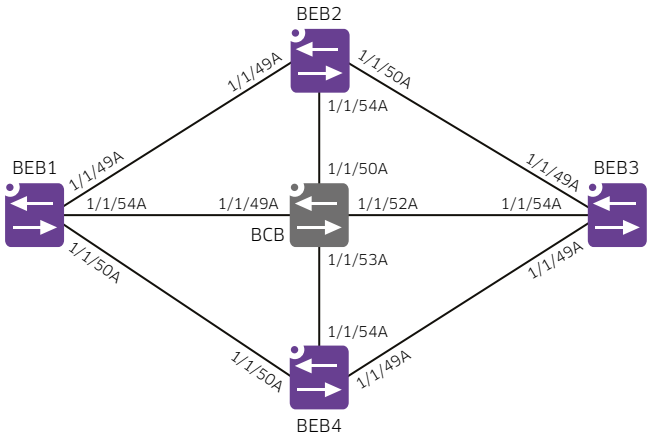

Following are the configuration snippets for all nodes.

Snippet 1. BEB-1's backbone configuration

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

Snippet 2. BEB-2's backbone configuration

```
BEB-2> spb bvlan 4000-4003
BEB-2> spb isis control-bvlan 4000
BEB-2> spb isis interface port 1/1/49A
BEB-2> spb isis interface port 1/1/50A
BEB-2> spb isis interface port 1/1/54A
BEB-2> spb isis admin-state enable
```

Snippet 3. BEB-3's backbone configuration

```
BEB-3> spb bvlan 4000-4003
BEB-3> spb isis control-bvlan 4000
BEB-3> spb isis interface port 1/1/49A
BEB-3> spb isis interface port 1/1/50A
BEB-3> spb isis interface port 1/1/54A
BEB-3> spb isis admin-state enable
```

Snippet 4. BEB-4's backbone configuration

```
BEB-4> spb bvlan 4000-4003
BEB-4> spb isis control-bvlan 4000
BEB-4> spb isis interface port 1/1/49A
BEB-4> spb isis interface port 1/1/50A
BEB-4> spb isis interface port 1/1/54A
BEB-4> spb isis admin-state enable
```

Snippet 5. BCB backbone configuration

```
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 and this maximises 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 (such as, 4000 in this example) and form point-to-point adjacencies. LSPs are exchanged, a topology database is created and one SPT is built for each BVLAN.

Let's review this configuration with some show commands.

Snippet 6. "show SPB isis interface"

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

In the "show spb isis interface" command output we can observe 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).

Snippet 7. "show SPB isis nodes"

```
BEB-1> show spb isis nodes
SPB ISIS Nodes:
 System Name          System Id       SourceID BridgePriority
--------------------+---------------+--------+---------------
 BEB-2                dc08.5610.7429  0x07429  32768 (0x8000)
 BEB-4                dc08.5610.77e9  0x077e9  32768 (0x8000)
 BCB                  dc08.5610.78d9  0x078d9  32768 (0x8000)
 BEB-1                dc08.5610.7f19  0x07f19  32768 (0x8000)
 BEB-3                dc08.5610.8649  0x08649  32768 (0x8000)
Total SPB Nodes : 5
```

In the "show spb isis nodes" command output we can observe 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), as well as the source ID and the bridge priority. The source ID is a 20-bit identifier which designates the node as the origin of BUM traffic and 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.

Snippet 8. "show SPB isis adjacency"

```
BEB-1> show spb isis adjacency
SPB ISIS Adjacency:
 System
 (Name : SystemId)                    Type   State   Hold  Interface
-------------------------------------+------+-------+------+----------
 BEB-2               : dc08.5610.7429  L1    UP        25     1/1/49A
 BEB-4               : dc08.5610.77e9  L1    UP        24     1/1/50A
 BCB                 : dc08.5610.78d9  L1    UP        18     1/1/54A
Adjacencies : 3
```

In the "show spb isis adjacency" command output we can observe 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), as well as 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.

Snippet 9. "show SPB isis bvlans"

```
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
```

In the "show spb isis bvlans" command output we can observe, 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.

Snippet 10. "show SPB isis unicast-table"

```
BEB-1> show spb isis unicast-table
SPB ISIS Unicast MAC Table:
        Destination                              Outbound
 BVLAN  (Name : MAC Address)                     Interface
------+---------------------------------------+-----------
  4000  BEB-2              : dc:08:56:10:74:29    1/1/49A
  4000  BEB-4              : dc:08:56:10:77:e9    1/1/50A
  4000  BCB                : dc:08:56:10:78:d9    1/1/54A
  4000  BEB-3              : dc:08:56:10:86:49    1/1/49A
  4001  BEB-2              : dc:08:56:10:74:29    1/1/49A
  4001  BEB-4              : dc:08:56:10:77:e9    1/1/50A
  4001  BCB                : dc:08:56:10:78:d9    1/1/54A
  4001  BEB-3              : dc:08:56:10:86:49    1/1/54A
  4002  BEB-2              : dc:08:56:10:74:29    1/1/49A
  4002  BEB-4              : dc:08:56:10:77:e9    1/1/50A
  4002  BCB                : dc:08:56:10:78:d9    1/1/54A
  4002  BEB-3              : dc:08:56:10:86:49    1/1/54A
  4003  BEB-2              : dc:08:56:10:74:29    1/1/49A
  4003  BEB-4              : dc:08:56:10:77:e9    1/1/50A
  4003  BCB                : dc:08:56:10:78:d9    1/1/54A
  4003  BEB-3              : dc:08:56:10:86:49    1/1/50A
MAC Addresses: 16
```

In the “show spb isis unicast-table” command output we can observe, 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.

Snippet 11. “show SPB isis spf bvlan”

```
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:08:56:10:74:29    1/1/49A  BEB-2              : dc:08:56:10:74:29     10    1
 BEB-4              : dc:08:56:10:77:e9    1/1/50A  BEB-4              : dc:08:56:10:77:e9     10    1
 BCB                : dc:08:56:10:78:d9    1/1/54A  BCB                : dc:08:56:10:78:d9     10    1
 BEB-3              : dc:08:56:10:86:49    1/1/49A  BEB-2              : dc:08:56:10:74:29     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:08:56:10:74:29    1/1/49A  BEB-2              : dc:08:56:10:74:29     10    1
 BEB-4              : dc:08:56:10:77:e9    1/1/50A  BEB-4              : dc:08:56:10:77:e9     10    1
 BCB                : dc:08:56:10:78:d9    1/1/54A  BCB                : dc:08:56:10:78:d9     10    1
 BEB-3              : dc:08:56:10:86:49    1/1/54A  BCB                : dc:08:56:10:78:d9     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:08:56:10:74:29    1/1/49A  BEB-2              : dc:08:56:10:74:29     10    1
 BEB-4              : dc:08:56:10:77:e9    1/1/50A  BEB-4              : dc:08:56:10:77:e9     10    1
 BCB                : dc:08:56:10:78:d9    1/1/54A  BCB                : dc:08:56:10:78:d9     10    1
 BEB-3              : dc:08:56:10:86:49    1/1/54A  BCB                : dc:08:56:10:78:d9     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:08:56:10:74:29    1/1/49A  BEB-2              : dc:08:56:10:74:29     10    1
 BEB-4              : dc:08:56:10:77:e9    1/1/50A  BEB-4              : dc:08:56:10:77:e9     10    1
 BCB                : dc:08:56:10:78:d9    1/1/54A  BCB                : dc:08:56:10:78:d9     10    1
 BEB-3              : dc:08:56:10:86:49    1/1/50A  BEB-4              : dc:08:56:10:77:e9     20    2
```

In the “show spb isis spb bvlan” command output we can observe, for a given BVLAN, the outbound interface, the next hop node, as well as the SPB metric and total number of hops required to reach a destination node. We can observe in this output that traffic destined towards BEB-3 will transit BEB-2 in the case of BVLAN 4000, BCB in the case of BVLANs 4001 and 4002, and BEB-4 in the case of BVLAN 4003.

9. L2 services

A L2 service refers to a type of VPN service connecting multiple sites in a single any-to-any bridging domain. In this section, we continue building upon the previous example and 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 entails 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 17. L2 service

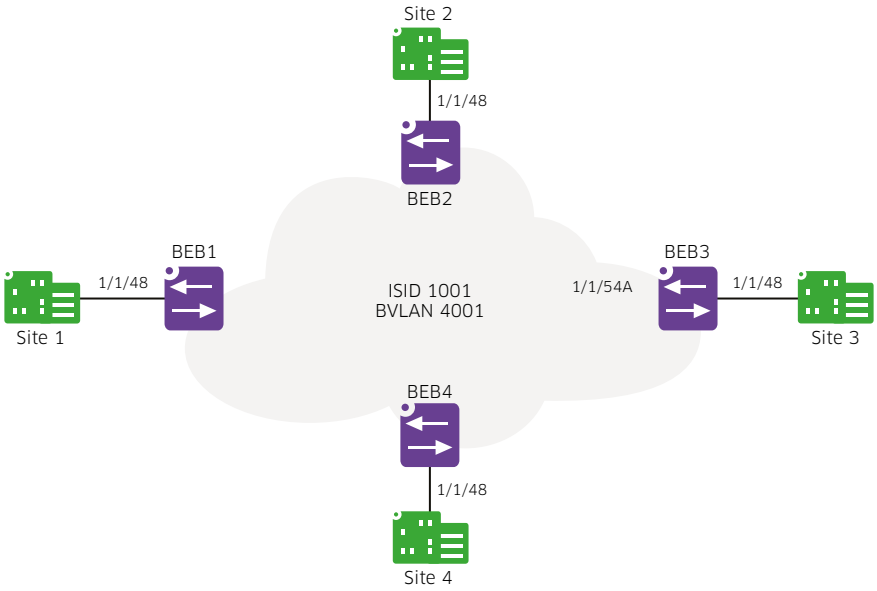

With regard to figure 17, we provide BEB configurations in the snippets that follow. As well, please note:

- 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

Snippet 12. BEB-1's service configuration

```
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
```

Snippet 13. BEB-2's service configuration

```
BEB-2> service access port 1/1/48
BEB-2> service 1 spb isid 1001 bvlan 4001
BEB-2> service 1 sap port 1/1/48:0
```

Snippet 14. BEB-3's service configuration

```
BEB-3> service access port 1/1/48
BEB-3> service 1 spb isid 1001 bvlan 4001
BEB-3> service 1 sap port 1/1/48:0
```

Snippet 15. BEB-4's service configuration

```
BEB-4> service access port 1/1/48
BEB-4> service 1 spb isid 1001 bvlan 4001
BEB-4> service 1 sap port 1/1/48:0
```

In the four configuration snippets above we can observe 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

Let's now proceed to verify the service status.

Snippet 16. "show service spb" – BEB view

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

In the "show service spb" command output we can observe, 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. We can also observe the multicast replication mode, which is head-end by default. The multicast replication mode can be changed to tandem on a per-service basis.

Snippet 17. "show service spb" – BCB view

```
BCB> show service spb
Legend: * denotes a dynamic object
SPB Service Info
  SystemId : dc08.5610.78d9,   SrcId : 0x078d9,    SystemName : BCB
                              SAP     Bind                    MCast
ServiceId   Adm  Oper Stats Count   Count   Isid      BVlan Mode      (T/R)
-----------+----+----+-----+-------+-------+---------+-----+--------------
Total Services: 0
```

In the "show service spb" command output we can observe that, by definition, a BCB does not have locally defined services.

Snippet 18. "show spb isis services" – BEB view

```
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-2               : dc:08:56:10:74:29
*     1001     4001   BEB-4               : dc:08:56:10:77:e9
*     1001     4001   BEB-1               : dc:08:56:10:7f:19
*     1001     4001   BEB-3               : dc:08:56:10:86:49
ISIDs:     4
```

In the "show spb isis services" command output we can observe 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 learnt thanks to the IS-IS CP. A "*" denotes that the service also matches a service locally created on the BEB.

Snippet 19. "show spb isis services" – BCB view

```
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-2                : dc:08:56:10:74:29
      1001     4001   BEB-4                : dc:08:56:10:77:e9
      1001     4001   BEB-1                : dc:08:56:10:7f:19
      1001     4001   BEB-3                : dc:08:56:10:86:49
ISIDs:     4
```

In the "show spb isis services" command output we can observe the same output now from the perspective of a BCB. We should note that a BCB is still aware of all existing services with the IS-IS CP.

Snippet 20. "show service spb"

```
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,
  Ingress Pkts     : 0,                    Ingress Bytes    : 0,
  Egress Pkts      : 0,                    Egress Bytes     : 0,
  Mgmt Change      : 06/23/2020 16:11:30,  Status Change    : 06/23/2020 16:11:30
```

The "show service spb" command output provides some additional details about a given SPB service. We can highlight the following:

- **RemoveIngressTag**: As explained in section 3, by default, a PBB frame includes all the frame's original tags. However, we can choose to remove those tags with the "service *service_id* remove-ingress-tag enable" command.
- **VLAN Translation**: 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 such situation, VLAN translation can be enabled to allow both devices to communicate. We should note that VLAN translation must be enabled both at 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. We will cover dynamic service creation in section 13.3.

Snippet 21. "show service access"

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

In the “show service access” command output we can observe, 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. We will cover dynamic SAP creation in section 12.2. We can also observe 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 Table 2. 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”. We will cover unp SAPs and profiles in section 12.2.

Table 2. Default L2 profiles

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

Snippet 22. “show service spb ports”

```
BEB-1> show service spb 1 ports
Legend: (*)Dyn Unicast (+)Remote Mcast (#)Local Mcast (~)Internal User Port Loopback
($)Internal Linkagg Loopback
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    N    dc08.5610.7429:4001  1/1/49A  BEB-2
sdp:32773:1*           Up   Up    N    dc08.5610.8649:4001  1/1/54A  BEB-3
sdp:32781:1*           Up   Up    N    dc08.5610.77e9:4001  1/1/50A  BEB-4
Total Ports: 4
```

In the “show service spb ports” command output, we can observe local (SAP) as well 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.

Snippet 23. "show service mesh-sdp spb"

```
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      dc08.5610.7429:4001  Up   1/1/49A  BEB-2
1         32773:1*         1001      dc08.5610.8649:4001  Up   1/1/54A  BEB-3
1         32781:1*         1001      dc08.5610.77e9:4001  Up   1/1/50A  BEB-4
Total Mesh-SDPs: 3
```

In the "show service mesh-sdp spb" command output we can observe far-end SDPs for each service along with the ISID number and the far-end system ID (BMAC), BVLAN, system name and associated interface.

Snippet 24. "show mac-learning domain spb" – BEB view

```
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
```

In the "show mac-learning domain spb" command output we can observe the list of CMAC addresses learnt 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.

Snippet 25. "show mac-learning domain spb" – BCB view

```
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
```

In the "show mac-learning domain spb" command output we can observe the same output now from the point of view of a BCB node. As expected, BCB nodes do not learn any CMACs.

10. Routing concepts

Before delving into L3 services, which are covered in the next section, we need to discuss certain routing concepts in relation to SPB. The Alcatel-Lucent OmniSwitch® product line has supported SPB since AOS 7.3.1, released in 2012. Since then, multiple SPB-enabled platforms have been launched and each new platform incorporated new advancements in ASICs.

First generation ASICs were not capable of routing and performing MAC-in-MAC encapsulation in a single-pass operation. Consequently, routing between IP interfaces associated to two different SPB services, or to a VLAN and an SPB service, had to traverse the switch fabric twice. This required an external physical loopback connecting two different switch ports: one port in the VLAN domain and another SAP in the SPB domain. IP interfaces could only be associated to a VLAN, not directly to an SPB service. It should be noted that these physical loopbacks can be either physical ports or linkaggs. When using VC, linkagg member ports can span different units in the VC for redundancy. We refer to this as two-pass routing with external physical loopback.

Newer generation ASICs support a concept similar to an external physical loopback without requiring a cable connection. One or more physical ports' bandwidth is dedicated to the loopback function without requiring a cable to be attached. Multiple ports can be dedicated to this function for additional bandwidth and redundancy. When using multiple ports, ports are configured as a linkagg and, when using VC, linkagg member ports can span different units in the VC. We refer to this as two-pass routing with internal front-panel loopback. One additional difference between the internal front-panel loopback and the external physical loopback described in the previous paragraph is that the internal front-panel loopback is a single logical port, not two ports (a VLAN port and a SAP) as in the case of the external physical loopback. However, even in the single logical port, there is a "VLAN" function and a "SAP" function. This will become clearer when looking at the configuration snippets later in this section.

Latest generation ASICs support integrated routing and bridging in the SPB domain in the exact same manner as in the VLAN domain. 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 loopbacks. We refer to this as single-pass inline routing.

Figure 18. Routing options – Physical view

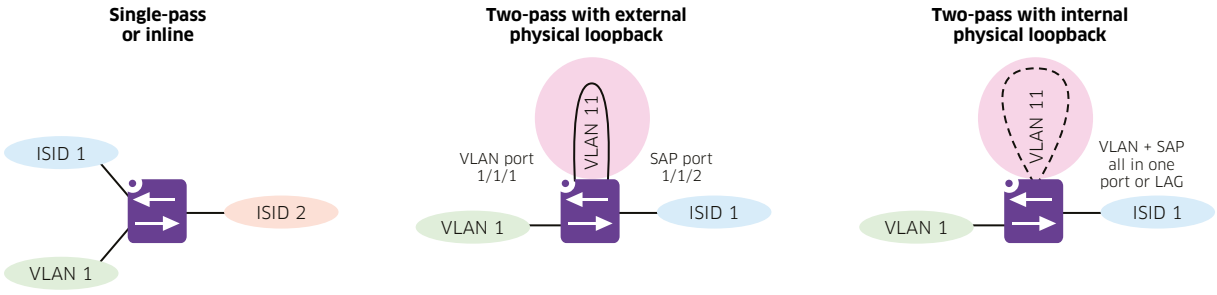

Figure 18 provides a physical view of these routing options. The leftmost diagram represents a switch supporting single-pass inline routing. This example shows a bridge with 2 SPB services, designated by their ISIDs, and one VLAN. IP interfaces are represented by dots. As we can see, the IP interfaces are bound to either VLANs or services and the switch performs inter-VLAN, inter-Service or inter-VLAN-Service routing directly in a single operation.

The diagram in the middle illustrates the case of two-pass routing with a physical hairpin. In this diagram, you can observe that IP interfaces are bound to VLAN 1 and VLAN 11, but not directly to the service. The external physical loopback cable creates the link between the service and the "dummy" VLAN, VLAN 11 in this example, where the IP interface resides. This external physical loopback is configured with a SAP-side, where SAPs are defined for each service requiring routing, and a VLAN-side, where dummy VLANs mapping to those services are tagged.

The right diagram illustrates the case of two-pass routing with internal front-panel loopback. In this diagram, the dotted line represents an imaginary physical external loopback, which is not required. In addition to not requiring an external physical loopback cable, the front-panel loopback requires a minimum of one port only. CLI configuration is different between physical external loopback and front-panel internal loopback. However, the concepts are very similar. You should still think about the front-panel internal loopback port or ports as having a SAP function and a VLAN function all in one port or linkagg.

Figure 19. Routing Options - Logical view

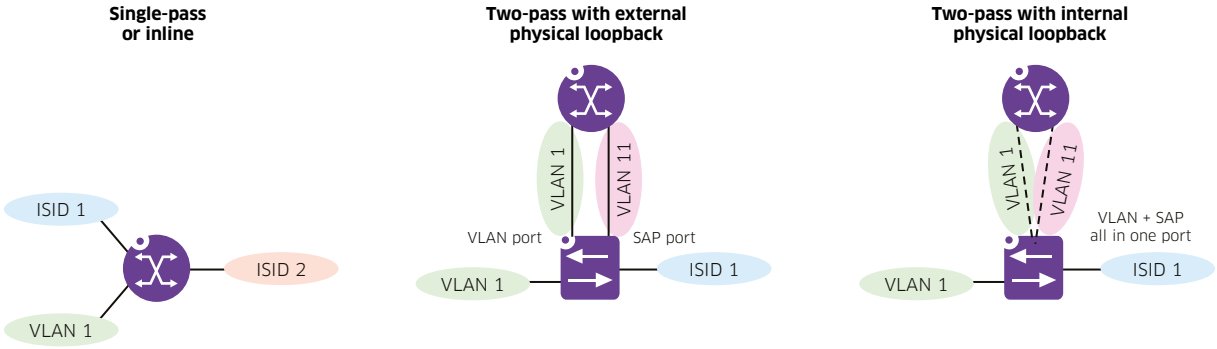

Figure 19 provides a logical representation of these options. The left diagram represents the case of single-pass or inline routing. In these products, routing and bridging functions are fully integrated in the service domain in the exact same manner as they are integrated in the VLAN domain. For this reason, these products are represented with a router icon.

The diagram in the middle represents the case of two-pass routing with an external physical loopback. In these products, routing and bridging functions are separate and represented by router and bridge icons. You can observe that the router function, where dots representing IP interfaces exist, connects to the bridge function using a VLAN port and a SAP.

The right diagram illustrates the case of two-pass routing with internal front-panel loopback. As you can see, this case is almost the same to the case of two-pass routing with an external physical loopback from a logical standpoint. However, the routing function attaches to the bridging function using a single port or group of ports. This front-panel loopback port or group of ports still performs a SAP function and a VLAN function. In addition, this connection between routing and bridging functions is created internally in the switch ASIC and does not require an external cable.

Let's review some configuration examples to commit these concepts.

Snippet 26. Single-pass or inline routing example

```
BEB> ip interface vlan_1 address 10.1.2.1/24 vlan 1
BEB> ip interface service_1 address 10.1.1.1/24 service 1
BEB> ip interface service_2 address 10.1.3.1/24 service 2
```

The configuration snippet 26 shows that, in products supporting single-pass or inline routing, IP interfaces can be bound to services just like they can be bound to VLANs. The switch simply performs routing in the same domain (VLAN or Service) or between different domains (VLAN and Service). Note that the backbone and service configuration is not shown in this example.

Snippet 27. Two-pass routing with external physical loopback example

```
BEB> ip interface vlan_1 address 10.1.2.1/24 vlan 1
BEB> vlan 11 name "Dummy VLAN Service 1"
BEB> vlan 12 name "Dummy VLAN Service 2"
BEB> ip interface vlan_11 address 10.1.1.1/24 vlan 11 rtr-port port 1/1/1 tagged
BEB> ip interface vlan_12 address 10.1.3.1/24 vlan 12 rtr-port port 1/1/1 tagged
BEB> service access port 1/1/2
BEB> service spb 1 sap port 1/1/2:11
BEB> service spb 2 sap port 1/1/2:12
```

The configuration snippet 27 shows the equivalent configuration for products supporting two-pass routing with external physical loopback. Since IP interfaces cannot be bound to a service directly, we create 2 additional "dummy" VLANs to bind these interfaces to. VLAN 11 will be associated to service 1 and VLAN 12 will be associated to service 2. The external physical loopback uses port 1/1/1 as VLAN port and port 1/1/2 as SAP. When creating the IP interfaces bound to those dummy VLANs, we use the rtr-port option. This prevents those VLANs from being bound to other ports and disables STP on those VLANs. Note that as explained previously, linkaggs can be used instead of single ports and linkagg member ports can span diverse units in a VC for redundancy.

Snippet 28. Two-pass routing with internal front-panel loopback example

```
BEB> interfaces port 1/1/51A loopback
BEB> ip interface vlan_1 address 10.1.2.1/24 vlan 1
BEB> vlan 11 name "Dummy VLAN Service 1"
BEB> vlan 12 name "Dummy VLAN Service 2"
BEB> service access port 1/1/51A
BEB> service 1 sap port 1/1/51A:11
BEB> service 2 sap port 1/1/51A:12
BEB> ip interface service_1 address 10.1.1.1/24 vlan 11 rtr-port port 1/1/51A tagged
BEB> ip interface service_2 address 10.1.3.1/24 vlan 12 rtr-port port 1/1/51A tagged
```

The configuration snippet 28 shows the equivalent configuration for products supporting two-pass routing with internal front-panel loopback. Firstly, port 1/1/51A is designated as the front-panel loopback port. Dummy VLANs are created and SAPs linking those dummy VLANs to their associated services are defined on the loopback port. When creating the IP interfaces bound to the dummy VLANs, we use the rtr-port option and reference the loopback port. Once again, the example shows the case of single front-panel loopback port but linkaggs can be used for additional bandwidth and resiliency in the case of VC.

11. 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

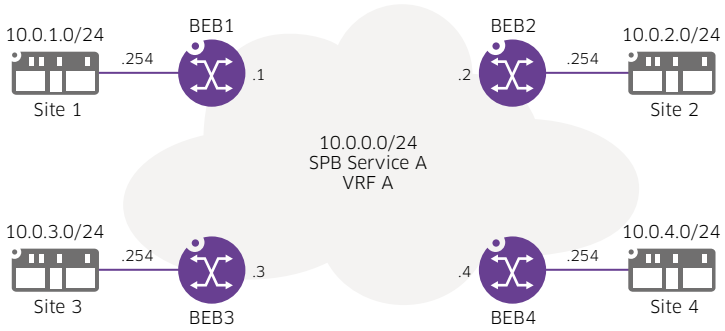

Figure 20 illustrates an example of a L3 Service connecting four of customer A's sites: Sites 1 through 4. You will notice that each site uses a different subnet and therefore, inter-site routing is required. BEB nodes connecting customer sites are represented with router icons for simplicity. These 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. Lastly, 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. Let's elaborate on these two variants.

11.1 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

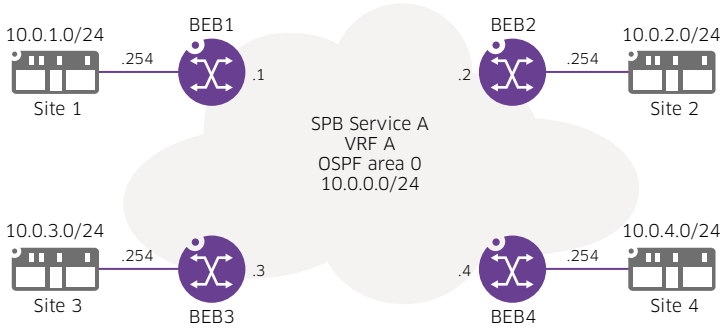

We should highlight that, in a VPN Lite type of L3 service, the L2 SPB service simply provides L2 connectivity to the "WAN" IP interfaces. Continuing with OSPF as an example, this means that OSPF is configured as usual. Also, since all WAN IP interfaces are connected to a single L2 SPB service, in the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

11.2 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. We should note that 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]. Refer to figure 22.

For those familiar with MPLS or EVPN, those technologies 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 play both of those roles; 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 learnt on the BEBs that those customers connect to.

Figure 22. Customer A's L3 VPN service

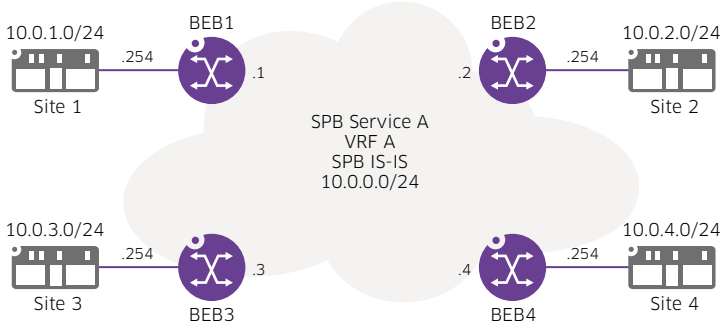

Unlike the case of a VPN Lite, an SPB L3 VPN does not require the addition of any routing protocol. Customer's VRF routes are exported to the SPB IS-IS instance, associated to the customer's ISID, and bound to the WAN IP as a gateway address. Far-end BEBs will import those routes into their local VRF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

Figure 23. Route Import/Export

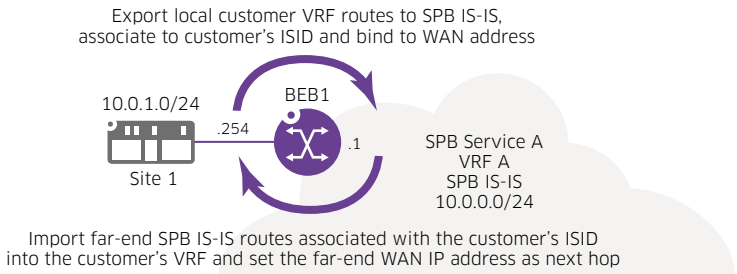

A L3 VPN service builds upon a L2 service and involves the following steps:

- 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 on products supporting single-pass inline routing, or on a “dummy” VLAN on products requiring external physical or internal front-panel loopback.
- Binding the WAN IP interface to the L2 SPB service’s ISID
- Route import/export between local VRF routing table and SPB IS-IS ISID instance

Let’s go back to the sample topology used for L2 services in section 9 and configure a L3 VPN service so we can have a look at the configuration. We will look at devices supporting internal front-panel loopback.

Figure 24. L3 VPN service example

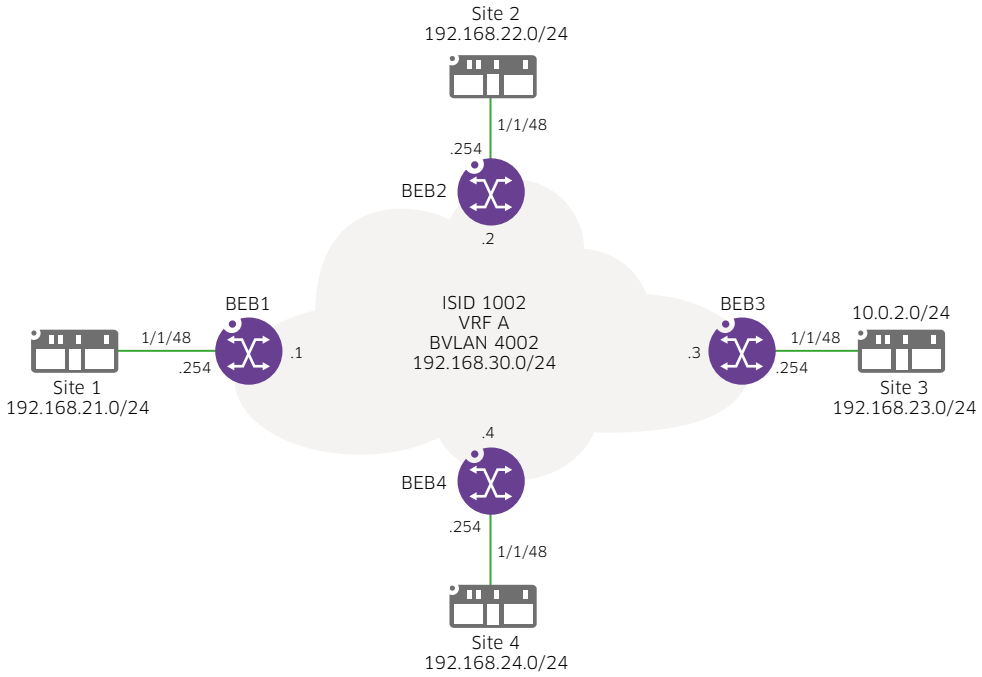

We will now provide configuration snippets for all BEBs. Like their L2 counterpart, L3 VPN services require no configuration on BCBs. Let’s provide some details about this example:

- Customer sites connect to their local BEB though interface 1/1/48
- LAN-side, or site default-gateway IP interfaces are bound to VLAN 3001, which is the default VLAN on port 1/1/48
- Port 1/1/54A is designated as a loopback port
- WAN-side IP interfaces are bound to dummy VLAN 3100

Snippet 29. L3 VPN example – BEB-1

```
BEB-1> vlan 3001 name "Site_1"
BEB-1> vlan 3100 name "Dummy VLAN Customer A"
BEB-1> vlan 3001 members port 1/1/48 untagged
BEB-1> interfaces port 1/1/51A loopback
BEB-1> service 2 spb isid 1002 bvlan 4002
BEB-1> service access port 1/1/51A
BEB-1> service 2 sap port 1/1/51A:3100
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 vlan
3100 rtr-port port 1/1/51A tagged
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
```

Snippet 30. L3 VPN example – BEB-2

```
BEB-2> vlan 3001 name "Site_2"
BEB-2> vlan 3100 name "Dummy VLAN Customer A"
BEB-2> vlan 3001 members port 1/1/48 untagged
BEB-2> interfaces port 1/1/51A loopback
BEB-2> service 2 spb isid 1002 bvlan 4002
BEB-2> service access port 1/1/51A
BEB-2> service 2 sap port 1/1/51A:3100
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 vlan
3100 rtr-port port 1/1/51A tagged
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
```

Snippet 31. L3 VPN example – BEB-3

```
BEB-3> vlan 3001 name "Site_3"
BEB-3> vlan 3100 name "Dummy VLAN Customer A"
BEB-3> vlan 3001 members port 1/1/48 untagged
BEB-3> interfaces port 1/1/51A loopback
BEB-3> service 2 spb isid 1002 bvlan 4002
BEB-3> service access port 1/1/51A
BEB-3> service 2 sap port 1/1/51A:3100
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 vlan
3100 rtr-port port 1/1/51A tagged
Customer_A: :BEB-3> ip export all-routes
Customer_A: :BEB-3> ip import isid 1002 all-routes
BEB-1> spb ipvpn bind vrf Customer_A isid 1002 gateway 192.168.30.3 all-routes
```

Snippet 32. L3 VPN example – BEB-4

```
BEB-4> vlan 3001 name "Site_4"
BEB-4> vlan 3100 name "Dummy VLAN Customer A"
BEB-4> vlan 3001 members port 1/1/48 untagged
BEB-4> interfaces port 1/1/51A loopback
BEB-4> service 2 spb isid 1002 bvlan 4002
BEB-4> service access port 1/1/51A
BEB-4> service 2 sap port 1/1/51A:3100
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 vlan
3100 rtr-port port 1/1/51A tagged
Customer_A: :BEB-4> ip export all-routes
Customer_A: :BEB-4> ip import isid 1002 all-routes
BEB-1> 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. Snippet 33 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.

Snippet 33. L3 VPN example – Verifying route import/export

```
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         00:23:18   LOCAL
  192.168.21.0/24      192.168.21.254    00:21:47   LOCAL
  192.168.22.0/24      192.168.30.2      00:20:30   IMPORT
  192.168.23.0/24      192.168.30.3      00:19:17   IMPORT
  192.168.24.0/24      192.168.30.4      00:18:09   IMPORT
  192.168.30.0/24      192.168.30.1      00:21:50   LOCAL
```

Snippet 34 shows arp entries in BEB-1's VRF "Customer_A". Far-end WAN gateway addresses are dynamically learnt.

Snippet 34. L3 VPN example – Verifying gateway L2 reachability

```
Customer_A::BEB-1> show arp
Total 4 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
-----------------+-------------------+----------+-------+-----------------+---------
--+-----------------------------------
 192.168.21.1       00:50:56:85:4d:1b   DYNAMIC                      1/1/48  LAN
 192.168.30.2       dc:08:56:10:74:29   DYNAMIC                     1/1/51A  WAN
 192.168.30.3       dc:08:56:10:86:49   DYNAMIC                     1/1/51A  WAN
 192.168.30.4       dc:08:56:10:77:e9   DYNAMIC                     1/1/51A  WAN
```

In addition to these L3-related verification steps, all steps covered in section 9 can be used to verify the underlying L2 service.

11.3 VPN Lite versus L3 VPN

Having presented VPN Lite and L3 VPN, we can now discuss the pros and cons and provide guidelines to help you choose one versus the other.

Let's start with the advantages of L3 VPN:

- **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: A total of 32 x OSPF configurations across all nodes. In case dual stack IPv4 and IPv6 support is required, this translates to an OSPFv2 and an OSPFv3 instance per BEB and VRF: 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.

With such compelling arguments in favour of L3 VPN, you may wonder why anyone would choose to use VPN Lite instead. The reason is that, 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. This is where VPN Lite comes in. VPN Lite 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.

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

12. 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 familiar diagram that we have been using so far, 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 section 11.2.

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 2 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 nor with the shared services.

Snippet 35. Route leaking

```
! 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 L3_VPN_A all-routes
vrf shared_services ip import vrf L3_VPN_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
```

Snippet 35 provides the commands required to accomplish this on the border BEBs, BEB1 and BEB2.

We can summarize the process as below:

- Shared routes are exported from the shared_services VRF and into the global IP routing table. 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 global IP routing table and into the customer VRFs. Note that this step is only necessary if customer sites are connected to the border BEBs.
- Customer routes are imported from the global IP routing table and into the shared_services VRF. Note that this step is only necessary 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.

Figure 25. Shared services

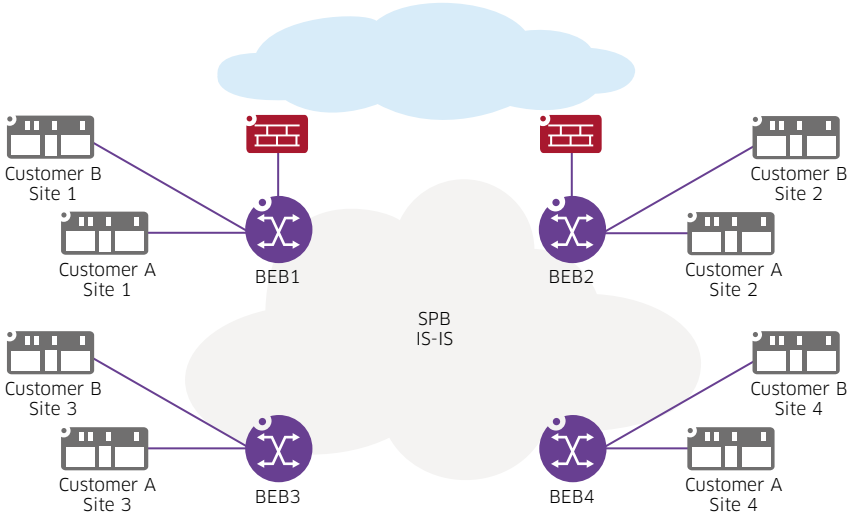

13. Automation

Up to this point, we have explained SPB concepts and configured the SPB backbone and services manually. However, AOS incorporates features that can build both the SPB backbone and services automatically. In this section, we will explain the various mechanisms that make a near zero-touch SPB network possible. A factory-default Alcatel-Lucent OmniSwitch has these mechanisms enabled by default and will automatically attempt to create an SPB backbone and services as explained in the subsequent subsections, unless these automation features are explicitly disabled. This set of features is sometimes referred to as "Intelligent Fabric" or "iFab" for short. In this section, we provide a simplified, high-level overview of these features. For a detailed description, please refer to the Alcatel-Lucent OmniSwitch Switch Management Guide.

13.1 Auto-Fabric

Figure 26 is a simplified view of a factory-default OmniSwitch bootup process. For a more detailed flow chart, please refer to the Alcatel-Lucent OmniSwitch Switch Management Guide.

This process involves the following stages:

- Auto Virtual Chassis (VC)
- Auto Remote Configuration Download (RCD)
- Auto LACP
- Auto SPB
- Auto MVRP
- Auto IP

Auto-Fabric features are enabled by default on a factory-default OmniSwitch. These features can however be disabled in their entirety, or, on a per-protocol or per-port basis. By default, automatically learnt and created configuration is not saved to the vcboot.cfg file but this option can be enabled.

Figure 26. Bootup state diagram

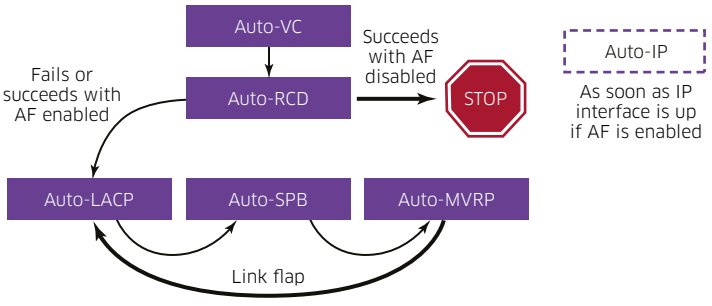

Let's describe these stages one-by-one.

13.1.2 Auto-VC

On bootup, and in absence of the vcsetup.cfg file, an OmniSwitch uses LLDP to detect other VC-compatible nodes connected to the default auto-VFL ports. Default auto-VFL ports depend on the product family. Some families such as the Alcatel-Lucent OmniSwitch® 6860 Stackable LAN Switch have 2 designated VFL ports which default to this role. In other families such as the Alcatel-Lucent OmniSwitch® 6900 Stackable LAN Switch, which support VC of up to 6 units, the last 5 VFL-eligible ports default to auto-VFL ports. If other products in the same family are detected at the other end, they will attempt to automatically create a VC. A Master node will be chosen through an election mechanism and non-Master nodes will reboot. Since this process creates a vcsetup.cfg file on all involved nodes, auto-VC will not kick-in in subsequent node reboot events.

13.1.3 Auto-RCD

Next, and in absence of a vcboot.cfg file, an OmniSwitch attempts to obtain an IP address through DHCP on any of its operational non-VFL ports. It will try this using the untagged default VLAN and tagged VLAN 127 and it will retry three times. If the switch succeeds in obtaining an IP address, and depending on the DHCP options in the lease, the switch will subsequently attempt to fetch an instruction file from a TFTP server or it will contact the Alcatel-Lucent OmniVista® 2500 Network Management System. Next, the switch will attempt to download firmware and vcboot.cfg from either an FTP/SFTP server or OmniVista. If the switch succeeds at obtaining its firmware and configuration, it will reboot and load its configuration. Depending on the configured options, the switch may or may not continue with the subsequent stages. Please refer to the AOS Switch Management Guide and to [2] for further details.

13.1.4 Auto-LACP

All non-VFL ports are auto-LACP enabled by default. Auto-LACP kicks in on a factory-default switch or a non-factory-default switch, unless explicitly disabled. Auto-LACP can be disabled globally or only on specific ports.

During the auto-LACP stage, a switch uses LLDP to identify switches connected to auto-LACP-enabled ports. Any LACP-compatible ports linking the same pair of switches will be automatically added to a linkagg. Even if there is only a single link connecting two nodes, it will still be configured as a linkagg because this allows additional links to be added later on without requiring configuration changes. For instance, by creating a linkagg of 1 member port and by referencing the (logical) linkagg as opposed to the (physical) port in other configuration commands, those configuration commands do not need to change when additional member ports are added to the linkagg. This is a best practice.

Note that, even if the remote switch is not an OmniSwitch, but is (manually) configured for LACP, the OmniSwitch detects LACP PDUs and automatically configures its side of the linkagg. This simplifies deployment even when 3rd party switches are used.

13.1.5 Auto-SPB

All non-VFL ports and linkaggs are auto-SPB enabled by default. Auto-SPB kicks in on a factory-default switch or a non-factory-default switch, unless explicitly disabled. Auto-SPB can be disabled globally or only on specific ports or linkaggs.

Auto-SPB also uses LLDP to detect presence of SPB-capable switches. When an SPB-capable switch is detected, the switch will attempt to configure the port or linkagg as an SPB backbone interface. When doing so it will use certain defaults.

On switches running AOS release 8.7R1 and later these defaults are:

- BVLANs 4000 through 4003 are created and mapped to ECT IDs 1 through 4 respectively
- BVLAN 4000 is designated as the control BVLAN

If the switch succeeds in establishing at least one SPB adjacency, all remaining non-VFL and non-SPB backbone ports are automatically configured as auto UNP access ports, unless explicitly disabled. Please refer to section 13.3 for details on auto UNP access ports.

13.1.6 Auto-MVRP

Auto-MVRP is enabled on factory-default switches. On switches booting from a vcboot.cfg file however, this feature needs to be explicitly enabled. When auto-MVRP is enabled, and if the switch fails to establish any SPB adjacency, MVRP will be enabled on all remaining and operational non-VFL ports. This enables the dynamic instantiation of VLANs learnt from neighbouring switches.

13.1.7 Auto-IP

The Auto-IP features runs in parallel with other features described in this section and, when enabled, it kicks-in as soon as an IP interface is created. Auto-IP listens for routing protocol (OSPFv2, OSPFv3 or IS-IS) “Hello” packets from neighbour devices and automatically creates local routing configuration matching parameters in the received “Hello” packets such that an adjacency can be formed. For example, reception of an OSPF “Hello” packet with area 1, Hello timer of 5 and Dead timer of 20 will result in matching configuration on the local device such that the two devices become neighbours and an adjacency is established.

13.2 Dynamic SAPs

Up to this point, we have shown how to configure SAPs statically and manually. However, SAPs can be automatically and dynamically configured using the User 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.

Let’s analyse the sample configuration in snippet 36 . 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 control BVLAN and therefore 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 document.

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. Let’s start by analysing 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 utilise the “SAMPLE_FLOW” port template. This template is defined such that:

- 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. A “silent” device is a device that does not transmit traffic for extended periods of time because it goes into power-save mode for instance. These periods of inactivity can result in a loss of service binding, thus 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.

It should be noted that statically binding a SAP, as opposed to a UNP, also offers a solution to the silent device problem. However, by statically binding a UNP instead of a SAP, the exact same UNP constructs can 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. This is considered a best practice.

Snippet 36. Dynamic SAPs – L2 services

```
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
multicast-mode headend vlan-xlation
BEB-1> unp profile "IoT" map service-type spb tag-value 0 isid 1002 bvlan 4002
multicast-mode headend vlan-xlation
BEB-1> unp profile "GUEST" map service-type spb tag-value 0 isid 1003 bvlan 4003
multicast-mode headend vlan-xlation
BEB-1> unp profile "WLAN" map service-type spb tag-value 0 isid 1004 bvlan 4001
multicast-mode headend vlan-xlation
BEB-1> unp profile "CCTV" map service-type spb tag-value 0 isid 1005 bvlan 4002
multicast-mode headend vlan-xlation
BEB-1> unp profile "RESTRICTED" map service-type spb tag-value 0 isid 1006 bvlan
4003 multicast-mode headend vlan-xlation
BEB-1> unp port-template SAMPLE_FLOW direction both aaa-profile "UPAM" default-
profile "RESTRICTED" classification ap-mode 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 port-type access
BEB-1> unp port 1/1/5 classification ap-mode dynamic-service spb
BEB-1> unp port 1/1/5 admin-state enable
BEB-1> unp port 1/1/5 profile "CCTV"
BEB-1> unp port 1/1/6 port-type access
BEB-1> unp port 1/1/6 classification ap-mode dynamic-service spb
BEB-1> unp port 1/1/6 admin-state enable
BEB-1> unp port 1/1/6 profile "WLAN"
BEB-1> unp port 1/1/7 port-type access
BEB-1> unp port 1/1/7 classification ap-mode dynamic-service spb
BEB-1> unp port 1/1/7 admin-state enable
BEB-1> unp port 1/1/7 profile "EMPLOYEE"
BEB-1> unp port 1/1/8 port-type access
BEB-1> unp port 1/1/8 classification ap-mode dynamic-service spb
BEB-1> unp port 1/1/8 admin-state enable
BEB-1> unp port 1/1/8 profile "GUEST"
BEB-1> unp port 1/1/9 port-type access
BEB-1> unp port 1/1/9 classification ap-mode dynamic-service spb
BEB-1> unp port 1/1/9 admin-state enable
BEB-1> unp port 1/1/9 profile "IoT"
BEB-1> unp port 1/1/10-16 port-type access
BEB-1> unp port 1/1/10-16 port-template SAMPLE_FLOW
```

Let's analyse the L3 Service case for this example. What this means is that, rather than routing at a centralized switch or firewall, edge routing is performed. Furthermore, let's consider the case of devices which attach to a standard VLAN port (for example not a SAP) and BEBs supporting front-end-panel loopback routing. Since VLAN-to-Service mapping happens at the loopback port, in this case we need to create bridge-type (VLAN) UNPs instead of access-type UNPs. The SPB configuration will be statically defined. Configuration snippets are split in three parts for convenience. Snippet 37 contains the VLAN-domain part of the configuration, snippet 38 contains the IP-domain part of the configuration, and snippet 39 contains the Service-domain part of the configuration.

We should note that devices placed in the "RESTRICTED" role do not normally need to communicate with other such devices. However, the configuration snippet allows for all routes in the RESTRICTED VRF to be imported. This can be modified with the addition of a route-map permitting routes to a central BEB or firewall only. Furthermore, a policy list can be attached t o the RESTRICTED UNP definition such that those devices can only communicate with certain head-end resources and can only use certain ports or applications. We will leave this exercise for you to complete.

Snippet 37. Dynamic SAPs – L3 services – VLAN Domain

```
BEB-1> interfaces port 1/1/51A loopback
BEB-1> vlan 101 name "EMPLOYEE_LAN_SIDE"
BEB-1> vlan 102 name "IoT_LAN_SIDE"
BEB-1> vlan 103 name "GUEST_LAN_SIDE"
BEB-1> vlan 104 name "WLAN_LAN_SIDE"
BEB-1> vlan 105 name "CCTV_LAN_SIDE"
BEB-1> vlan 106 name "RESTRICTED_LAN_SIDE"
BEB-1> vlan 1001 name "EMPLOYEE_WAN_SIDE"
BEB-1> vlan 1002 name "IoT_WAN_SIDE"
BEB-1> vlan 1003 name "GUEST_WAN_SIDE"
BEB-1> vlan 1004 name "WLAN_WAN_SIDE"
BEB-1> vlan 1005 name "CCTV_WAN_SIDE"
BEB-1> vlan 1006 name "RESTRICTED_WAN_SIDE"
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 vlan 101
BEB-1> unp profile "IoT" map vlan 102
BEB-1> unp profile "GUEST" map vlan 103
BEB-1> unp profile "WLAN" map vlan 104
BEB-1> unp profile "CCTV" map vlan 105
BEB-1> unp profile "RESTRICTED" map vlan 106
BEB-1> unp port-template SAMPLE_FLOW direction both aaa-profile "UPAM" default-
profile "RESTRICTED" classification ap-mode 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 bridge
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
```

Snippet 38. Dynamic SAPs – L3 services – IP Domain

```
BEB-1> vrf create EMPLOYEE
EMPLOYEE: :BEB-1> ip interface "LAN" address 192.168.101.254 mask 255.255.255.0 vlan
101
EMPLOYEE: :BEB-1> ip interface "WAN" address 192.168.1.1 mask 255.255.255.0 vlan
1001 rtr-port port 1/1/51A tagged
EMPLOYEE: :BEB-1> ip export all-routes
EMPLOYEE: :BEB-1> ip import isid 1001 all-routes
BEB-1> spb ipvpn bind vrf EMPLOYEE isid 1001 gateway 192.168.1.1 all-routes
BEB-1> vrf create IOT
IOT: :BEB-1> ip interface "LAN" address 192.168.102.254 mask 255.255.255.0 vlan 102
IOT: :BEB-1> ip interface "WAN" address 192.168.2.1 mask 255.255.255.0 vlan 1002
rtr-port port 1/1/51A tagged
IOT: :BEB-1> ip export all-routes
IOT: :BEB-1> ip import isid 1002 all-routes
BEB-1> spb ipvpn bind vrf IOT isid 1002 gateway 192.168.2.1 all-routes
BEB-1> vrf create GUEST
GUEST: :BEB-1> ip interface "LAN" address 192.168.103.254 mask 255.255.255.0 vlan
103
GUEST: :BEB-1> ip interface "WAN" address 192.168.3.1 mask 255.255.255.0 vlan 1003
rtr-port port 1/1/51A tagged
GUEST: :BEB-1> ip export all-routes
GUEST: :BEB-1> ip import isid 1003 all-routes
BEB-1> spb ipvpn bind vrf GUEST isid 1003 gateway 192.168.3.1 all-routes
BEB-1> vrf create WLAN
WLAN: :BEB-1> ip interface "LAN" address 192.168.104.254 mask 255.255.255.0 vlan 104
WLAN: :BEB-1> ip interface "WAN" address 192.168.4.1 mask 255.255.255.0 vlan 1004
rtr-port port 1/1/51A tagged
WLAN: :BEB-1> ip export all-routes
WLAN: :BEB-1> ip import isid 1004 all-routes
BEB-1> spb ipvpn bind vrf WLAN isid 1004 gateway 192.168.4.1 all-routes
BEB-1> vrf create CCTV
CCTV: :BEB-1> ip interface "LAN" address 192.168.105.254 mask 255.255.255.0 vlan 105
CCTV: :BEB-1> ip interface "WAN" address 192.168.5.1 mask 255.255.255.0 vlan 1005
rtr-port port 1/1/51A tagged
CCTV: :BEB-1> ip export all-routes
CCTV: :BEB-1> ip import isid 1005 all-routes
BEB-1> spb ipvpn bind vrf CCTV isid 1005 gateway 192.168.5.1 all-routes
BEB-1> vrf create RESTRICTED
RESTRICTED: :BEB-1> ip interface "LAN" address 192.168.106.254 mask 255.255.255.0
vlan 106
RESTRICTED: :BEB-1> ip interface "WAN" address 192.168.6.1 mask 255.255.255.0 vlan
1006 rtr-port port 1/1/51A tagged
RESTRICTED: :BEB-1> ip export all-routes
RESTRICTED: :BEB-1> ip import isid 1006 all-routes
BEB-1> spb ipvpn bind vrf RESTRICTED isid 1006 gateway 192.168.6.1 all-routes
```

Snippet 39. Dynamic SAPs – L3 services – Service Domain

```
BEB-1> service access port 1/1/51A
BEB-1> service 1 spb isid 1001 bvlan 4001 description EMPLOYEE
BEB-1> service 1 sap port 1/1/51A:1001
BEB-1> service 2 spb isid 1002 bvlan 4002 description IOT
BEB-1> service 2 sap port 1/1/51A:1002
BEB-1> service 3 spb isid 1003 bvlan 4003 description GUEST
BEB-1> service 3 sap port 1/1/51A:1003
BEB-1> service 4 spb isid 1004 bvlan 4001 description WLAN
BEB-1> service 4 sap port 1/1/51A:1004
BEB-1> service 5 spb isid 1005 bvlan 4002 description CCTV
BEB-1> service 5 sap port 1/1/51A:1005
BEB-1> service 6 spb isid 1006 bvlan 4003 description RESTRICTED
BEB-1> service 6 sap port 1/1/51A:1006
```

13.3 Dynamic Services

In the preceding 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 centre. 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, this would require that the service UNP 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 is a poor practice as it 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. This feature is enabled by default on factory-default switches.

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. Snippet 40 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 + Domain ID + (VLAN Number % Service Modulo)
- BVLAN Index = ISID Number % (Total number of BVLANs)

By default:

- Base Service Number = 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.

Snippet 40. Dynamic services – Dynamic UNP

```
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 8 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 Snippet 41.

Snippet 41. Dynamic services – Dynamic UNP – Service Modulo

```
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, let's 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 Snippet 42 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.

Snippet 42. Dynamic services – Dynamic UNP – Multi-tenancy

```
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
```

Lastly, the Base Service Number (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.

14. Management

As explained in section 3.6, 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. This is a good option if all nodes support single-pass inline routing. However, nodes that do not support single-pass inline routing will require an external physical or internal front-panel loopback for this purpose even if they would not require it otherwise (for example, because they are BCBs).
- **In-band Management**: In-band management is applicable to all SPB nodes regardless of their routing capabilities (such as, single-pass inline, external physical, or internal front-panel loopback). Management IP interfaces can be created directly on the control BVLAN, therefore,

no loopback of any kind is required. The management network or stations attach to one or more gateway nodes through VLAN-domain interfaces. We should note that 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 27. In-band management

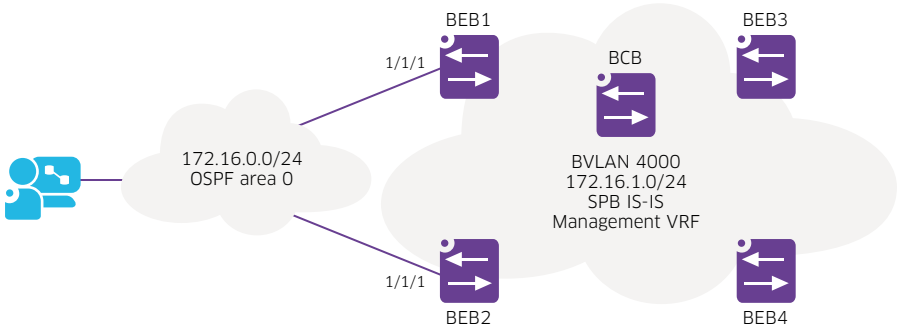

Let’s examine the in-band management example in figure 27. In this example, nodes BEB-1 and BEB-2 are gateways nodes linking the SPB-management domain and the VLAN-management domain. The VLAN-management subnet is 172.16.0.0/24 and the SPB-management subnet is 172.16.1.0/24. OSPF is used in the Management 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.

Snippet 43. In-band management – 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
```

Snippet 44. In-band management – 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
```

Snippet 45. In-band management – 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
```

Snippet 46. In-band management – 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
```

Snippet 47. In-band management – 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 snippets Snippet 43 through Snippet 47. OSPF and route-map configuration in BEBs 1 and 2 is excluded from these snippets.

15. Operation and Maintenance

15.1 Connectivity Fault Management: 802.1ag

CFM in an SPB network is most useful to perform 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 (CCM) are not currently supported in conjunction with SPB.

OAM is supported at the BVLAN level, refer to figure 28. Virtual MEPs must be configured for all BVLANs and BEBs and, optionally, also for BCBs (such that a L2 PING or L2 trace test can be initiated from any node to any other node). MIPs are automatically created and do not need to be explicitly configured.

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

Figure 28. OAM in BVLAN and VLAN Domains

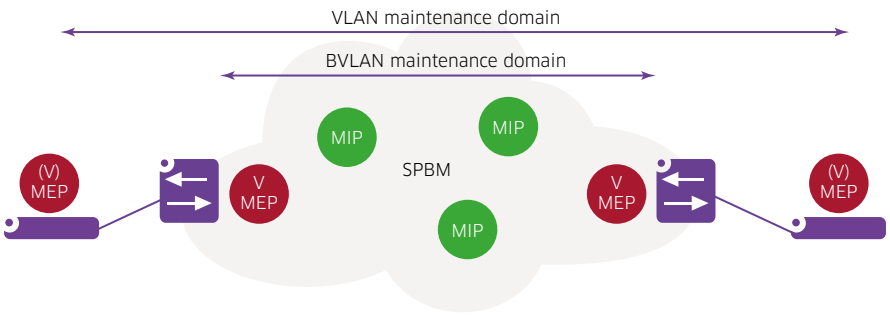

OAM is also supported at the VLAN level or between L2 access switches connected to BEBs over SAP UNIs. This is useful in a L2 deployment for testing end-to-end service connectivity between sites. OAM at the VLAN level must be set at a higher maintenance domain level than BVLAN OAM.

Figure 29 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 29. L2 ping and L2 trace

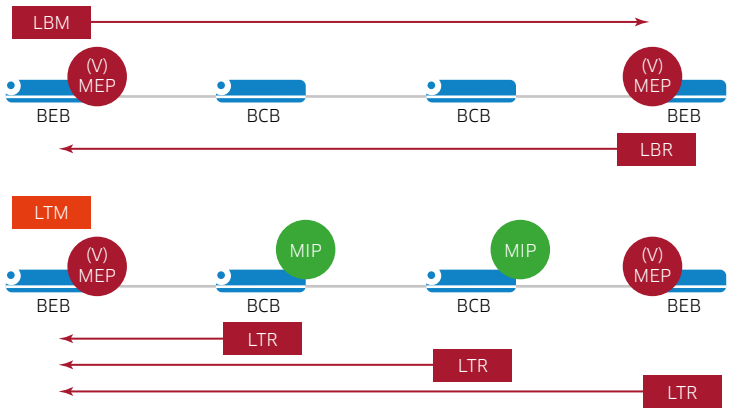

Configuration Snippet 48 provides a sample OAM configuration for service BVLANs 4001-4003.

Snippet 48. OAM

```
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 intervalls
BEB-1> ethoam association BVLAN4001 domain ALE endpoint-list 11-15
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 endpoint-list 21-25
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 endpoint-list 31-35
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 BVLAN4002 direction up port virtual primary-vlan 4003
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4002 admin-state enable
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4002 priority 5
BEB-1> ethoam endpoint 31 domain ALE association BVLAN4002 lowest-defect-priority all-defect
```

Snippet 49 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.

Snippet 49. L2 trace

```
BEB-1> ethoam linktrace target-macaddress 2c:fa:a2:1c:05:c1 source-endpoint 11 domain ALE association
BVLAN4001

Transaction Id:1681692778

BEB-1> show ethoam linktrace-reply domain ALE association BVLAN4001 endpoint 11 tran-id 1681692778


LTM operation successful. Target is reachable.
Ttl : 63,
  LTM Forwarded : yes,
  Terminal MEP : no,
  Last Egress Identifier : 00:00:2C:FA:A2:1C:09:81,
  Next Egress Identifier : 00:00:2C:FA:A2:1C:09:81,
  Relay Action : RLY_FDB,
  Chassis ID Subtype : NONE,
  Chassis ID : none,
  Ingress Action : ING_NONE,
  Ingress Mac : 00:00:00:00:00:00,
  Ingress Port ID Subtype : NONE,
  Ingress Port ID : none,
  Egress Action : EGR_OK,
  Egress Mac : 2C:FA:A2:1C:09:B6,
  Egress Port ID Subtype : LOCALLY_ASSIGNED,
  Egress Port ID : 1/1/47
Ttl : 62,
  LTM Forwarded : yes,
  Terminal MEP : no,
  Last Egress Identifier : 00:00:2C:FA:A2:1C:09:81,
  Next Egress Identifier : 00:00:2C:FA:A2:02:DE:B1,
  Relay Action : RLY_FDB,
  Chassis ID Subtype : NONE,
  Chassis ID : none,
  Ingress Action : ING_OK,
  Ingress Mac : 2C:FA:A2:02:DE:B8,
  Ingress Port ID Subtype : LOCALLY_ASSIGNED,
  Ingress Port ID : 1/1/1A,
  Egress Action : EGR_OK,
  Egress Mac : 2C:FA:A2:02:DE:B9,
  Egress Port ID Subtype : LOCALLY_ASSIGNED,
  Egress Port ID : 1/1/1B
Ttl : 61,
  LTM Forwarded : no,
  Terminal MEP : yes,
  Last Egress Identifier : 00:00:2C:FA:A2:02:DE:B1,
  Next Egress Identifier : 00:00:2C:FA:A2:1C:05:C1,
  Relay Action : RLY_HIT,
  Chassis ID Subtype : NONE,
  Chassis ID : none,
  Ingress Action : ING_OK,
  Ingress Mac : 2C:FA:A2:1C:05:F7,
  Ingress Port ID Subtype : LOCALLY_ASSIGNED,
  Ingress Port ID : 1/1/48,
  Egress Action : EGR_NONE,
  Egress Mac : 00:00:00:00:00:00,
  Egress Port ID Subtype : NONE,
  Egress Port ID : none
```

15.2 Network performance: Service Assurance Agent

Latency, jitter and packet loss SAA tests are automatically set-up between all BEBs and BCBs and across all BVLANs with the “saa auto-create” command. Refer to Snippet 50 showing the configuration and Snippet 51 showing sample statistics.

Snippet 50. Service Assurance Agent configuration

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

Snippet 51. Service Assurance Agent stats

```
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      hurl     Avg      Max          Thr
Sent Rcvd
-------------------------------------------------------------------------------------------------------------------------------------------------------------
SPB-4001-2c-fa-a2-02-de-b1            SPB        mac-ping      2017-01-24,05:37:21.0          107      109      117      0        2        6        10           0
5     5 DEFAULT
SPB-4001-2c-fa-a2-03-70-a9            SPB        mac-ping      2017-01-24,05:37:15.0          103      108      111      0        1        4        8 0          0
5     5 DEFAULT
SPB-4001-2c-fa-a2-1c-05-c1            SPB        mac-ping      2017-01-24,05:37:09.0          108      121      134      0        5        15       26           0
5     5 DEFAULT
SPB-4002-2c-fa-a2-02-de-b1            SPB        mac-ping      2017-01-24,05:37:39.0          105      109      113      0        3 ,      4        5            0
5     5 DEFAULT
SPB-4002-2c-fa-a2-03-70-a9            SPB        mac-ping      2017-01-24,05:37:33.0          102      111      118      0        0        7        16           0
5     5 DEFAULT
SPB-4002-2c-fa-a2-1c-05-c1            SPB        mac-ping      2017-01-24,05:37:27.0          108      111      114      0        0        2        4            0
5     5 DEFAULT
```

15.3 Network maintenance

Two features in SPB can assist in network maintenance tasks: Overload state and graceful restart.

15.3.1 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.

15.3.2 Graceful restart

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

Without graceful restart, a VC master or CMM takeover event would require neighbour 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 neighbour node, the node undergoing a takeover will announce this condition to its neighbours by setting the RR (restart request) in a TLV message and continue using its existing FDB while restarting. The neighbour 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 re-built in a much shorter time.

16. 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. But, access or Customer Edge (CE) devices connected to BEB nodes do not run SPB IS-IS and therefore other solutions are required when redundancy is needed. In this section we will present the different options for the different service types.

We start by highlighting that 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).

We will now present alternate redundancy options other than VC+LAG.

Let's start with L2 Services in figure 30 below. We can consider 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 fibre runs should use diverse physical paths to protect against fibre 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. Please refer to the "AOS 8 Network Configuration Guide" for further details.
- **Fully redundant**: This option adds CE device redundancy. MSTP (Multiple Spanning Tree Protocol) 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.

Note that Virtual Chassis (VC) can be combined with all the options above to increase resiliency.

Figure 30. L2 Service attachment

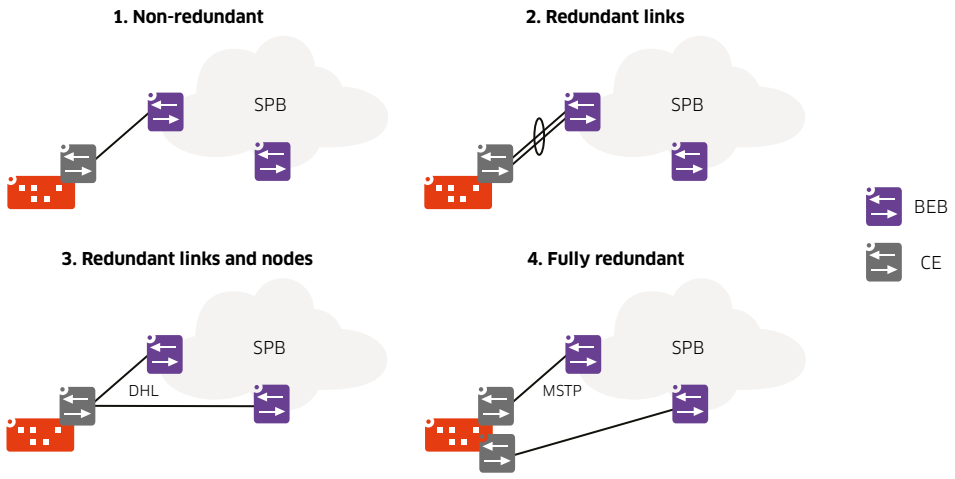

Let's now continue with L3 services. We can distinguish two sub-variants: L3 CE and L2 CE. A L3 CE can exchange routes with the BEBs by using any supported routing protocol as well as static or default routes. A L2 CE on the other hand will completely delegate routing to the BEB, which will act as a default gateway for local devices. These two sub-variants are illustrated in figure 31 and figure 32. Note that hairpins, when required, are not shown for simplicity.

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 link aggregate (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 maps.
- **Fully redundant**: This option adds CE device redundancy

Figure 31. L3 Service attachment - L3 CE

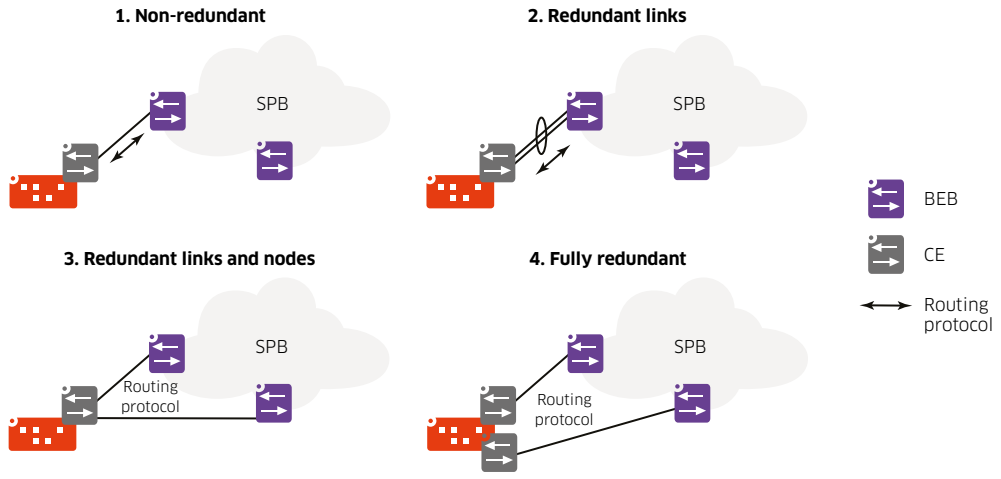

You may notice that the case of L3 Service attachment with a L2 CE is almost identical to the case of L2 Service attachment. However, since the routing function is delegated to the BEB, VRRP is required when CEs attach to redundant BEBs. This requires access VLANs to be extended across both BEBs. If 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, note that when using a L2 CE in a L3 Service, there is no routing protocol between CE and BEB. In such a 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 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 32. L3 service attachment - L2 CE

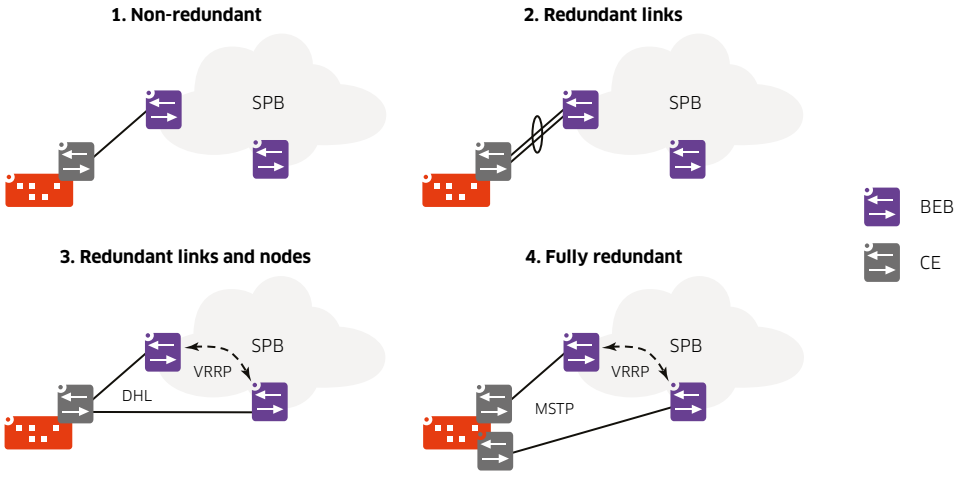

17. Loop avoidance and suppression

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 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 33 illustrates situations in which LBD can detect and break loops.

Figure 33. Loopback detection

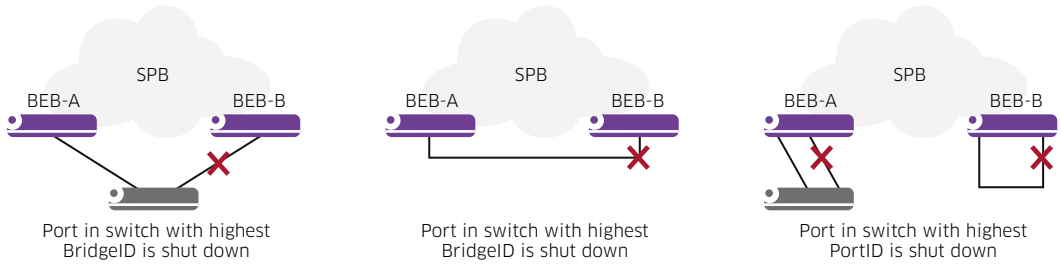

By default, LBD is disabled for the switch and on all service-access ports. Enable LBD globally on the switch and in specific service-access ports or linkaggs as shown in Snippet 52.

Snippet 52. Loopback detection

```
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 broadcast, multicast and unknown unicast traffic. 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. Please refer to the AOS Network Configuration Guide for further details.

18. 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.

18.1 BVLANs

As described in section 5, 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.

18.2 VLAN-to-Service mapping

When creating a SAP, AOS allows mapping multiple or all VLAN tags to the same SPB service. We want to stress that, 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, thus defeating 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; 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.

18.3 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.

Virtual chassis provides many benefits such as network architecture and management simplification. VC 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 virtual chassis 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 34. 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 VFL because traffic forwarding in a LAG prioritizes the use of local linkagg member ports over remote (across the VFL) member ports.

Figure 34. VC and SPB

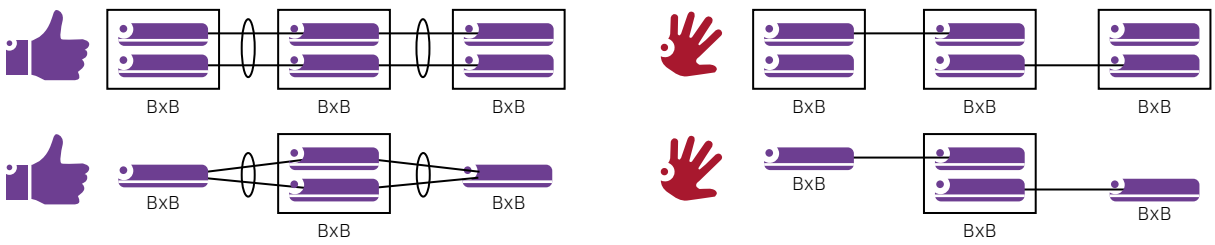

18.4 Link Aggregation

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. Please refer to the AOS Command Line Interface Guide for further details.

18.5 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. Table 3 shows a way in which the metric can be set to be inversely proportional to the link speed.

Table 3. Recommended Link Metric

Speed	Suggested Metric
100G	1000
50G	2000
40G	2500
25G	4000
10G	10000
1G	100000
100M	1000000

18.6 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.

Trusted SAPs copy CoS markings from the incoming VLAN tag onto the BVLAN 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.

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

When using an external or two-pass routing (external physical or internal front-panel loopback), the standard VLAN port must best set to trust and use CoS and not DSCP to preserve CoS markings end-to-end.

19. Security guidelines

In this section, we will provide some additional design guidelines specific to the security domain. This is not an exhaustive list of recommendations, rather, we will focus 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 [3].

19.1 Management VRF

As explained in section 3.6, 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.

We have covered different SBP management options in section 14. Management IP addresses can be bound to:

- The EMP port, in case of OOBM
- To a standard VLAN port, in the case of OOBM
- The control BVLAN, in the case of in-band management
- A Management SPB service, directly in the case of products supporting single-pass inline routing
- A Management SPB service, indirectly in the case of products supporting external physical or internal front-panel loopback

We want to point out that 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 Snippet 43 through Snippet 47.

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

19.2 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, 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.

19.3 NAC

In section 13.2, we 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.

19.4 Router authentication

As explained in section 11.1, 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 (e.g. MD5 for OSPF or BGP).

20. 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.

Alcatel·Lucent Enterprise