Alcatel·Lucent
Enterprise

Transportation Network Architecture Guide

Guide
Transportation Network Architecture Guide

Table of Contents

1 Introduction

1.1 Purpose

The purpose of this design guide is to present the requirements and considerations relevant to the transportation vertical along with design options, best practices and configuration guidelines. While this design guide is focused on Shortest Path Bridging (SPB-M) based networks, it also addresses data networking solutions applicable for this transportation vertical.

1.2 Audience

This design guide is intended for network architects and network engineers involved in the design, implementation and maintenance of networks in the transportation vertical.

To take advantage of this document, it is expected that the reader will be familiar with Shortest Path Bridging and will have a solid understanding of various networking technologies at the ACPS or similar level.

1.3 Scope

Although several modes of transport exist, this document focuses on rail (metro, heavy, light) and road transport. More specifically, this document focuses on the fixed network infrastructure underpinning the multiple systems that enable the safe and reliable operation of rail and intelligent transport systems.

Rail transport involves critical signaling and control systems which usually require sub-50ms convergence time. Although Shortest Path Bridging convergence time is roughly in the order of 200-300ms, a first significant improvement is expected in Software Release 8.5R2/3 to bring it below the 100ms. Note that ALE's ERPv2 (Ethernet ring protection switching) implementation meets the sub 50ms convergence time.

On-board, in-vehicle, ground-to-train, ground-to-vehicle communication is described in the "on-board communication solution" chapter. Indoor Location Based Services (LBS) technology is explained using an airport use case.

2 Transportation systems overview

Transportation services rely on multiple systems to keep traffic flowing or services running to schedule while ensuring drivers and passengers are safe and informed of any disruptions. These systems are enabled by a network and can be classified into four main categories: Control, safety, communications, and information. Please refer to Table 1 and Table 2 for examples of systems commonly found in ITS and rail/metro, respectively.

Table 1. Intelligent Transportation Systems

Intelligent transportation systems			
Control	**Safety**	**Communications**	**Information**
Signaling	Video surveillance	Telephony	Traveler information system
Traffic management system	Emergency call	Wireless LAN	
Variable speed limit	Access control	Toll collection	
	Fire/alarm detection		

Table 2. Rail and metro systems

Rail and metro systems			
Control	**Safety**	**Communications**	**Information**
Signaling	Video surveillance	Telephony	Passenger information system
Automatic train control	Emergency Call	Wireless LAN	Passenger announcement
	Access Control	Ticketing	Infotainment
	Fire/alarm detection	Fare collection	Internet

2.1 System description

Many of these systems are found in both rail/metro and ITS while others are more specific to either rail/metro or ITS. We provide a brief description of these systems below.

Signaling: The signaling system is responsible for directing road or railway traffic to avoid collisions. In its most basic form, the intersection or rail section is reserved for use in one direction at a time. More advanced forms of signaling improves the use of available capacity by detecting vehicle or train presence and adapting accordingly. Signaling is particularly critical in the case of trains: They require a long distance to stop due to the momentum associated with their large mass.

Automatic train control (ATC): Automatic Train Control comprises three sub-systems: Automatic train protection (ATP), automatic train operation (ATO) and automatic train supervision (ATS). ATP is responsible for keeping trains a safe distance apart. ATO is responsible for stopping the train at the right place such that all coaches are accessible from the platform. ATS monitors the system status detecting deviation from normal operation and schedules and dynamically adjusting to them.

Video surveillance (VS): CCTV is paramount to ensure personnel and passenger safety and to monitor critical assets as well as the state of the transportation network. High quality CCTV cameras can be installed at stations, intersections, tunnels, on-board the train, inside vehicles and along tunnels. Multiple parties can access video feeds in real time to improve response time in the event of an incident.

Access control (AC): These systems control access to restricted premises through badge or fingerprint scanners such that only authorized personnel are admitted.

Emergency call: Push-to-talk buttons at stations, on-board the trains and along the road connect to aid in the event of incidents, accidents or crime.

Fire/alarm detection: Fire, smoke and other alarms are reported to security staff at the operations control center and/or station.

Telephony: The telephony system is used for staff communications and as the underlying infrastructure for emergency call and public address. The telephony system also links to emergency responders (police, fire, and ambulance).

Passenger information system (PIS): Provides real-time information about service status, departure/arrival times and any delays or disruptions.

Traffic management system (TMS): In a manner that is very similar to air traffic control, TMS systems regulate the flow of vehicles with the goal of lessening or eliminating congestion and, in this way, improving road safety and efficiency. Sensors are embedded in the surface of the road or mounted on equipment such as poles or signs. Cameras are mounted on overpasses and other vantage points. They feed data and video back to the traffic operations center where it is processed and monitored and the resulting decisions are used to manage traffic.

Variable speed limit (VSL): Speed limits are adjusted to respond to various traffic (for example, congestion or crashes), and weather conditions (for example, fog or ice) and displayed on electronic signs.

Toll collection: This system includes manned stations as well as automatic smart tolling based on Dedicated Short Range Communications (DSRC), RFID tags or license plate recognition.

Traveler information system (TIS): Provides visual information about traffic and weather conditions, special events, incidents and disruptions.

Ticketing and automatic fare collection (AFC): This system includes ticket vending machines, manned ticket booths, fare gates and tap on/off smart card scanners.

Public address system (PAS): Provides audible information about service status, service departure times, schedule changes, etc.

Infotainment: Provides information such as the weather forecast, news and advertising at stations or on board the train or vehicle and can be a source of non-fare revenue.

Internet: Internet access at stations and on board the train or vehicle can improve passenger satisfaction. It can also generate ancillary revenues through access fees, advertising or other commercial arrangements.

2.2 Network requirements

This section will present and discuss the main requirements driving the network architecture and design choices.

2.2.1 Virtualization

One network will carry traffic for multiple systems over a common infrastructure. These systems will communicate various disparate, often proprietary, devices and applications that may be operated and maintained by different groups or vendors and may require communication with third parties. The network must be able to support multi-tenancy and virtual segregation such that systems and tenants do not interfere with one another. Virtual private networks (VPN) enable secure separation and bandwidth allocation for system and tenant traffic. As seen in Figure 1, all devices, systems and tenants connect to the same physical network. However, the network is logically partitioned into VPNs or so called “containers.”

Figure 1. Virtual Private Networks or Containers

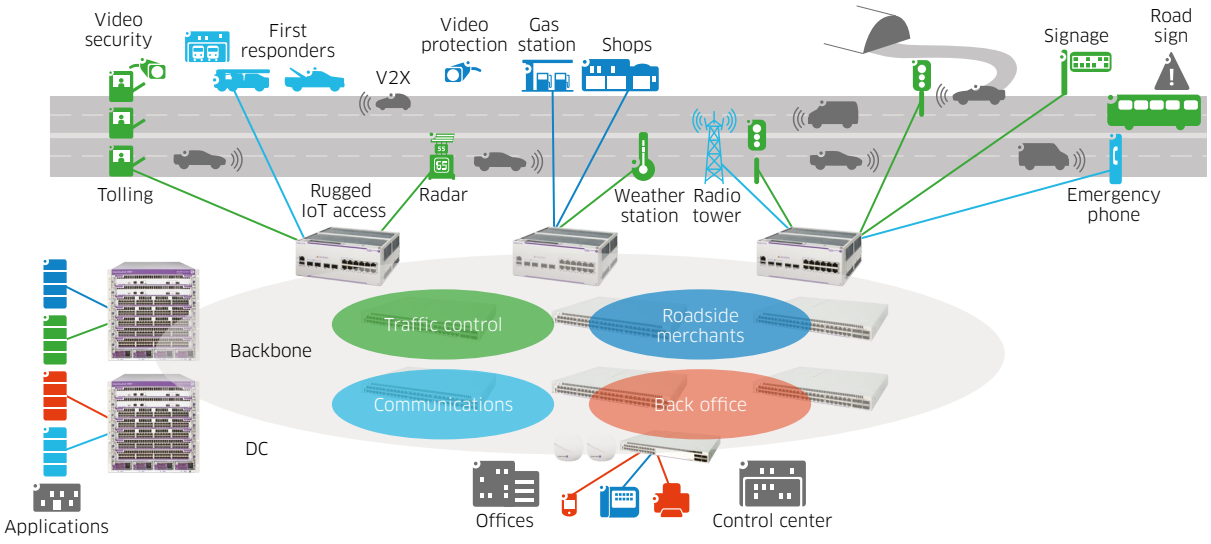

2.2.2 Availability

High availability is a fundamental requirement for a network carrying mission-critical system traffic. Redundancy without single point of failure (SPOF) is required at the network and system level such that recovery upon a failure event is automatic and maintenance tasks can be performed in-service. When a network is redundant without SPOF, the duration of an outage is equal to the convergence time needed to setup an alternative communication path. This design guide considers transportation networks with sub-second convergence time requirements. As previously mentioned, current convergence time of SPB network is roughly in the order of 200 ms. Software release 8.5R2 brings code improvements that should lower this convergence time to 100 ms.

2.2.3 Scalability

The network must be capable of scaling to support the required:

- Systems and tenants (VPNs)
- Network nodes
- End devices
- Multicast flows
- Bandwidth

Some of the largest transportation systems might include dozens of systems, hundreds of nodes, thousands of end devices and multicast flows.

2.2.4 Performance and quality of service

Not all systems are critical and each system has different performance requirements. At one end of the spectrum, an emergency call is safety-critical with low bandwidth requirements and at the other end of the spectrum, internet and infotainment are not critical but have moderate bandwidth requirements. The ability to prioritize certain systems over others is important when the network is congested or when traffic is re-routed due to any reason. When these conditions occur, the network will need to manage the congestion by allocating bandwidth and prioritizing traffic on a per-traffic-class basis. Systems will be mapped to traffic classes based on their level of priority and service level requirements. Every VPN service will carry traffic for a single system and will be mapped to a single traffic class. Therefore, single-level QoS is adequate and multi-level or hierarchical QoS will not be mandatory.

2.2.5 Security

Security is paramount to networks supporting critical transport infrastructure. In addition to segregation of system traffic into VPNs or containers, the following security requirements must be addressed:

- Network node security: Network nodes must be hardened and protected from attacks such as DDoS attacks.
- Network admission control and role-based access: Access to network resources will only be granted upon successful user or device authentication/ authorization and privileges will be set according to user or device role.
- Quarantine: The network must be able to isolate a compromised / tampered device.
- Integrity: Accuracy, consistency and trustworthiness must be protected while in transit through the network. Data must be protected from modification by unauthorized parties (for example, a man in the middle attack.)
- Confidentiality: Data traffic content must be protected from exposure to unauthorized parties while in transit through the network.

2.2.6 Environmental

Trackside and roadside equipment will be subject to harsh conditions such as extreme temperatures, vibrations, shock, dust and electrical/ magnetic emissions.

Trackside equipment must be immune to electromagnetic emissions and keep its own emissions within certain limits to not adversely affect signaling and other systems. These limits are defined by European standard EN 50121-4 (standard for rail applications).

Roadside equipment must be compliant with NEMA TS-2 standard to be mounted in NEMA roadside cabinets.

In brief, trackside and roadside equipment must be fan-less, rugged and compliant with industrial-grade temperature, vibration, shock and EMI/EMC specifications.

3 SPB Intelligent Fabric in transportation

An Alcatel-Lucent Enterprise SPB-enabled Intelligent Fabric addresses the requirements outlined in the preceding sections as follows:

- L2 and L3 VPNs or IoT containers
- High availability through self-healing redundancy and sub-second convergence time
- Scalability to thousands of nodes, services, devices and multicast flows
- High performance through shortest path and QoS
- Security through network admission control (NAC), role-based access, quarantine, DDoS protection, operating system hardening as well as data authentication and encryption through MACsec.
- Hardened Alcatel-Lucent OmniSwitch® 6465 and 6865 (SPB enabled) are suitable for roadside and trackside deployment

In addition, an ALE SPB-enabled Intelligent Fabric greatly simplifies network operations with:

- Single control protocol (IS-IS)
- Plug and play end-point provisioning through profiles
- Service provisioning only at the edge, not at the core
- Single operating system
- Single network management system
- OAM (operations and maintenance) support for monitoring and troubleshooting
- Analytics

Figure 2. SPB iFab benefits

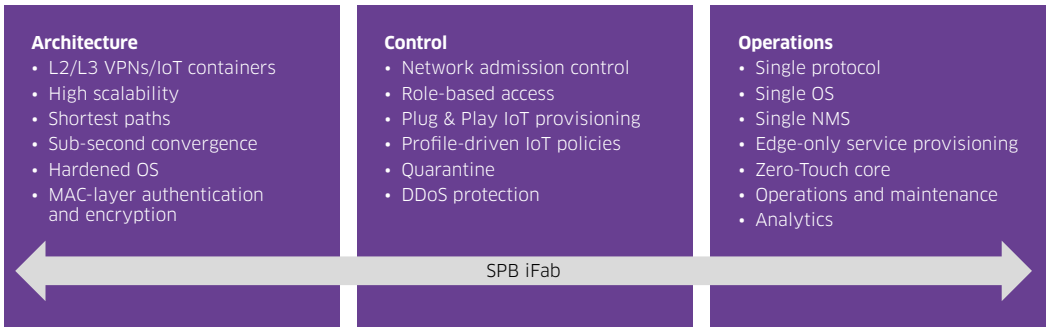

4 Architecture

This section presents control, backbone, site and site attachment architecture.

Note: Throughout this document the term “site” refers to a distinct physical location such as a station in railway or metro. The term site is preferred as it is more generic and applicable to ITS.

4.1 Overview

Rings are the natural topology to redundantly interconnect network nodes along a road, highway, metro or railway line. A sample ring topology is show in Figure 3 for a simple light rail line with 17 stations, a control center and a backup control center.

Figure 3. Ring topology

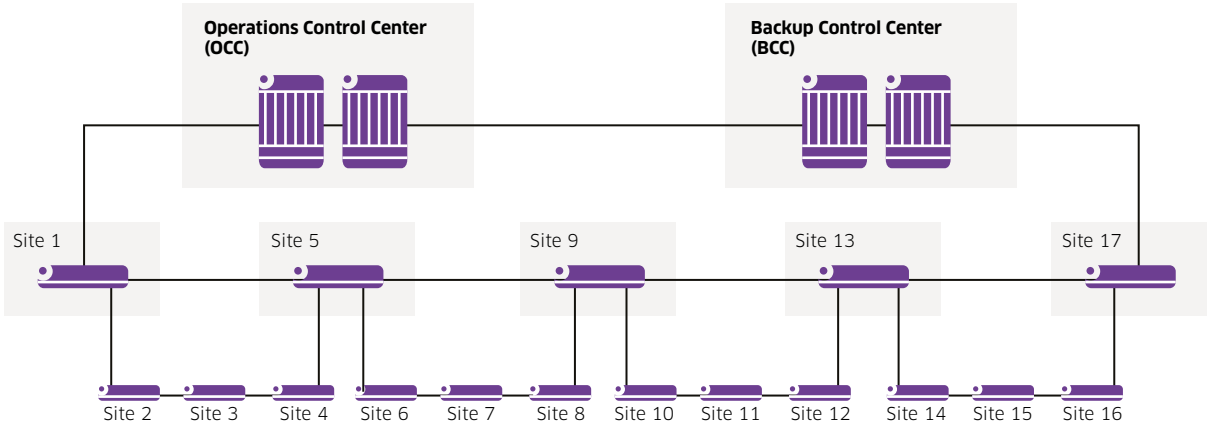

The operations control center (OCC) is the primary location where all aspects of the transportation system are supervised and controlled and responses to incidents are coordinated. The OCC hosts multi-disciplinary teams as well as various systems, applications, databases and interfaces with third parties such as emergency responders. The backup control center (BCC) hosts redundant infrastructure and resources such that it can replace the OCC in the event of a disaster or during maintenance. OCC and BCC can operate in active/active or active/standby mode.

4.2 Control architecture

When multiple lines are operated by a single entity, control of the individual lines can be fully centralized as seen in Figure 4 or hierarchical as seen in Figure 5. For the rest of this document, we will focus on individual lines within a centralized or hierarchical control architecture.

Figure 4. Centralized control

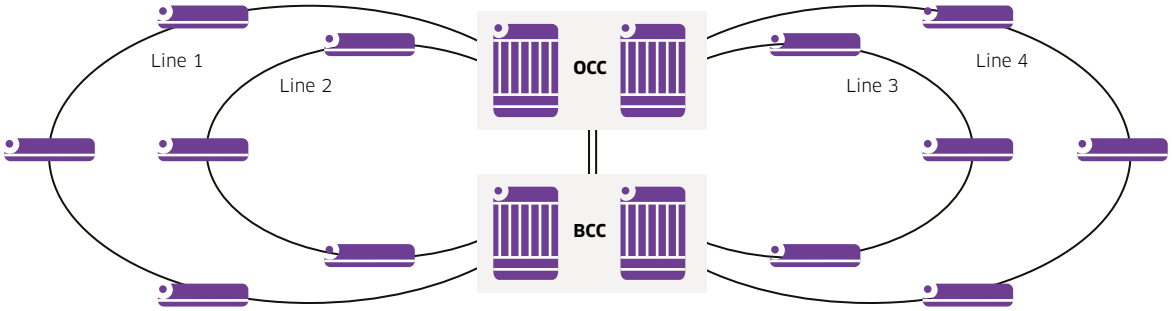

Figure 5. Hierarchical control

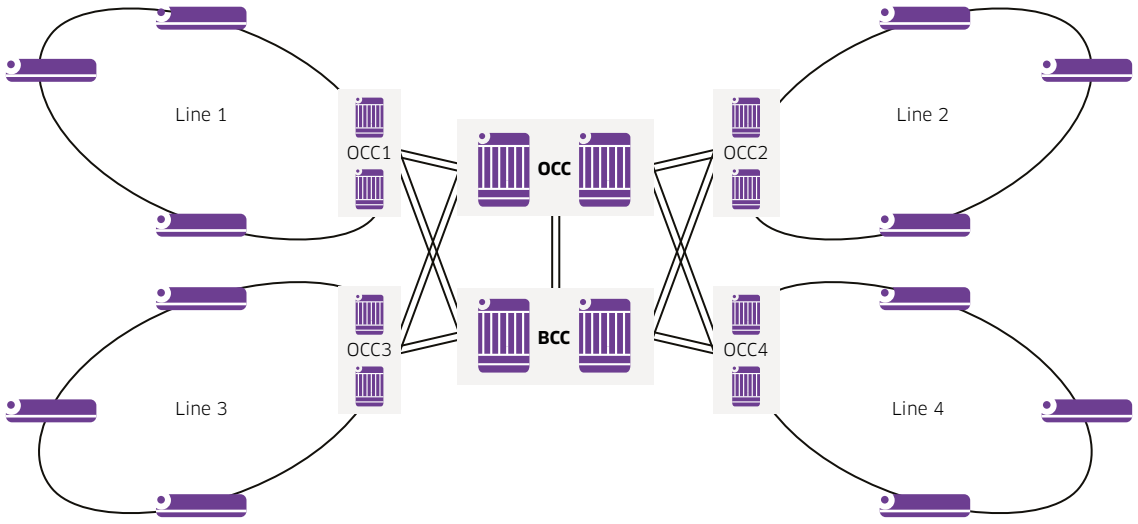

4.3 Site architecture

Site architecture can be classified as L2 or L3. At the site access level, the topology may also be a ring, or a spine-and-leaf architecture. While SPB can also be used within the site, in this guide we will consider the more general case in which the site access network is based on standard Ethernet such as STP or ERP. Therefore, we will not discuss the architecture at the site access level but rather its point of attachment to the backbone.

4.3.1 L2 VPN

In a L2 VPN architecture, no routing is performed at the station level or at the station's interconnection point to the backbone, referred to as the BEB (Backbone Edge Bridge.) Site VLANs will be mapped to SPB Services (ISIDs) at the backbone BEBs through SAPs (Service Access Point) as seen in Figure 6.

Site VLANs and ISIDs will normally be local to the site. In other words, each site will have its own set of VLANs and ISIDs which will not be shared across sites. Site VLANs will be mapped to ISIDs on a one-to-one basis through SAPs. All VLANs and ISIDs will be enabled at the OCC and BCC BEBs. As we will see in Section 5.1.2, this will limit the total number of VLANs that can be supported backbone-wide to the total number of ISIDs that can be supported at the OCC and BCC BEBs.

Sharing of VLANs and ISIDs across multiple sites is possible, however, this is not recommended because of implications on broadcast and multicast traffic. Please refer to Section 5.7.1 for a discussion of multicast replication modes and their impact on backbone bandwidth consumption.

All routing will be performed at the OCC and BCC sites which are set up as a VRRP redundant pair. This includes intra-site as well as inter-site routed traffic. This must be taken into consideration when sizing backbone links which will be discussed in Section 5.8.

Figure 6. L2 design

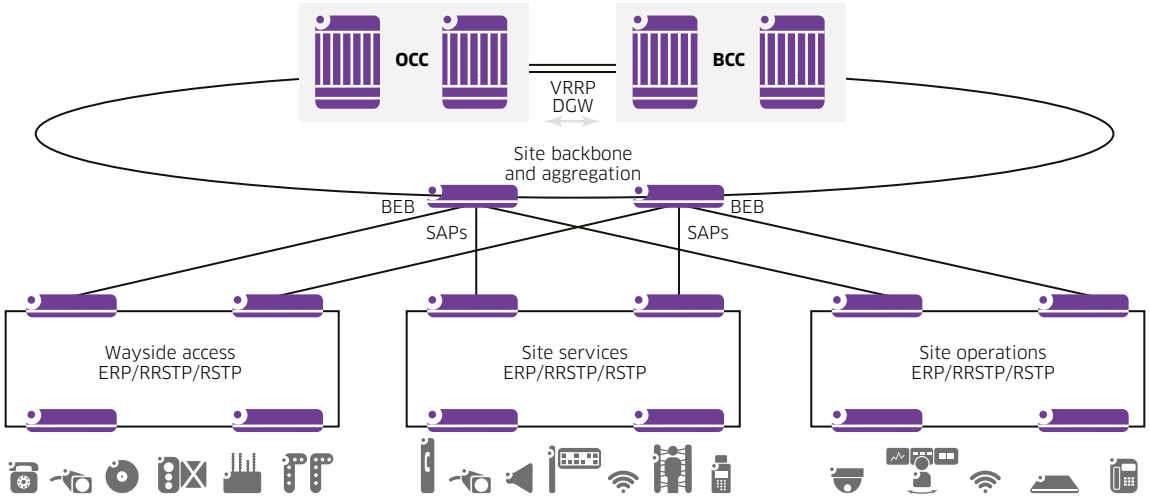

4.3.2 L3 VPN

Generally, a L3 design is more scalable than a L2 design because the IP address space is hierarchical and allows for summarization while the MAC address space is flat and cannot be summarized. Transportation networks are no exception on this rule.

In a L3 VPN architecture, routing is performed at the site BEBs. The site BEBs will be set up as a VRRP redundant pair and act as default gateway for site devices. Please refer to Figure 7.

Figure 7. L3 design

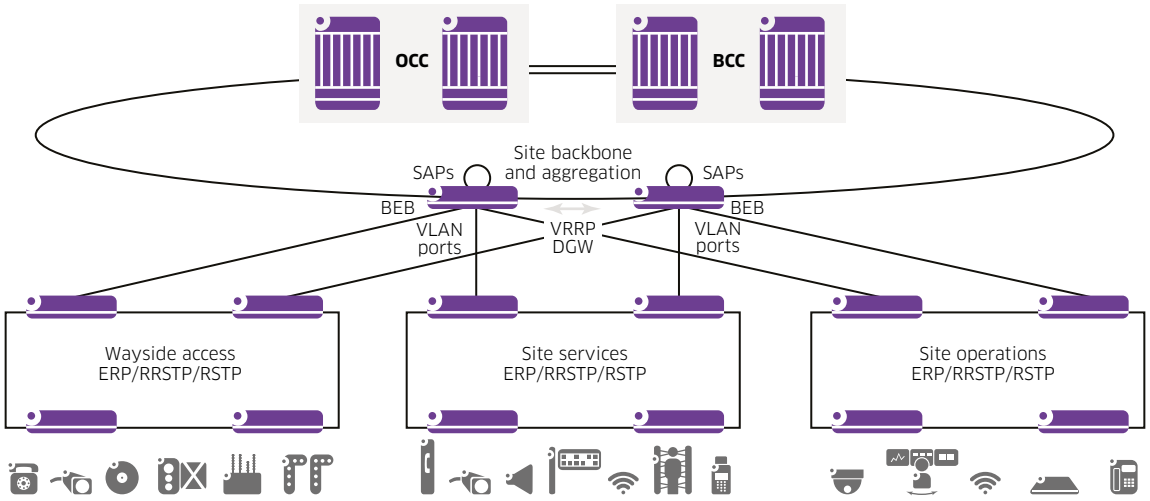

Every site will have its own set of site access VLANs which will not be shared nor exposed among other sites and are only local to the site. These site access VLANs can be the same or different across sites because they are only locally significant. The site access network will attach to the site BEBs through standard VLAN ports where these site access VLANs will be enabled. IP interfaces will reside in these site access VLANs and will be configured with VRRP to provide default gateway redundancy to site devices. For this reason, site access VLANs will need to be enabled on the link between both BEBs alongside BVLANs.

A different set of VLANs will be mapped to SPB Services (ISIDs) through SAPs on a hairpin loop. We will refer to this second set of VLANs as site uplink VLANs. Site uplink VLANs will be shared among all sites. From a site access VLAN, routes outside of the site will point to an IP interface residing on a site uplink VLAN. No site devices will be mapped to site uplink VLANs. Site uplink VLANs will be used only for routing on the VLAN-side of the hairpin loop. Please refer to Figure 8. We will refer to this kind of hairpin as a VLAN UNI attachment hairpin.

Please note that the speed of the hairpin ports must be the same or higher than the speed of the backbone NNI ports.

Figure 8. L3 VPN hairpin

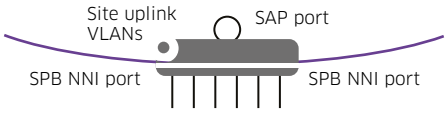

At the BEBs, VRFs and ISIDs will be created for all systems requiring L3 isolation. These VRFs and ISIDs will be the same across all sites. By means of the L3 VPN feature, site access subnet and site uplink subnet routes within a designated VRF will be exported to and imported from the ISID. This means that these routes will be propagated throughout the backbone by IS-IS using special TLVs as defined in the IETF draft [3] and no additional routing protocol will be required.

Compared to the L2 VPN design, a L3 VPN design is much more scalable in terms of total number of sites because site access VLANs are only locally significant and are not mapped to SPB services. Only site uplink VLANs are mapped to SPB services and these are shared among all sites. A L3 VPN design is also much more scalable in terms of total number of end devices because their MAC addresses will not be known outside of the site that they are in.

4.4 Site attachment

To avoid a single point of failure at the point of attachment to the backbone, the site access network will be redundantly attached to diverse BEBs. These BEBs can both reside within the same site as seen in Figure 9 or, alternatively, the site access network can be attached to a local BEB as well as a remote BEB as seen in Figure 10. The first alternative requires double the number of BEBs but no additional fiber cores and is simpler from an operational point of view while the second alternative requires no additional BEBs but requires additional fiber cores and is more complex from an operational perspective.

Figure 9. Intra-site redundant attachment

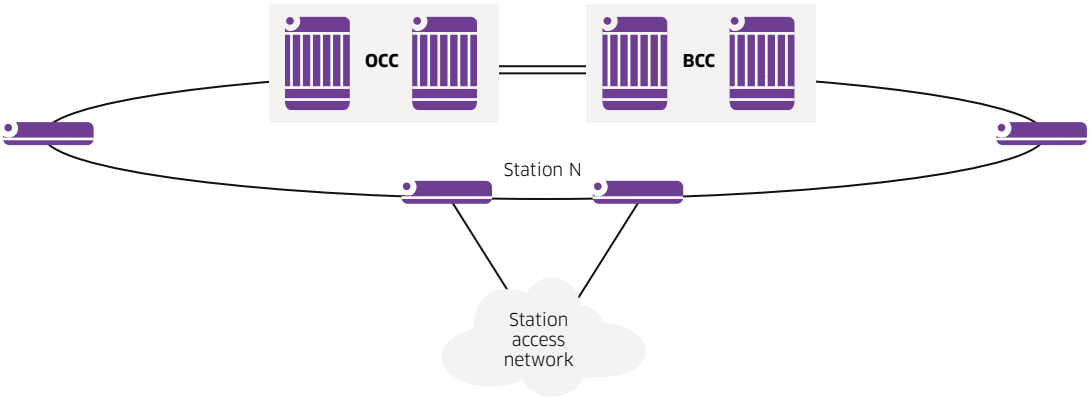

Figure 10. Inter-site redundant attachment

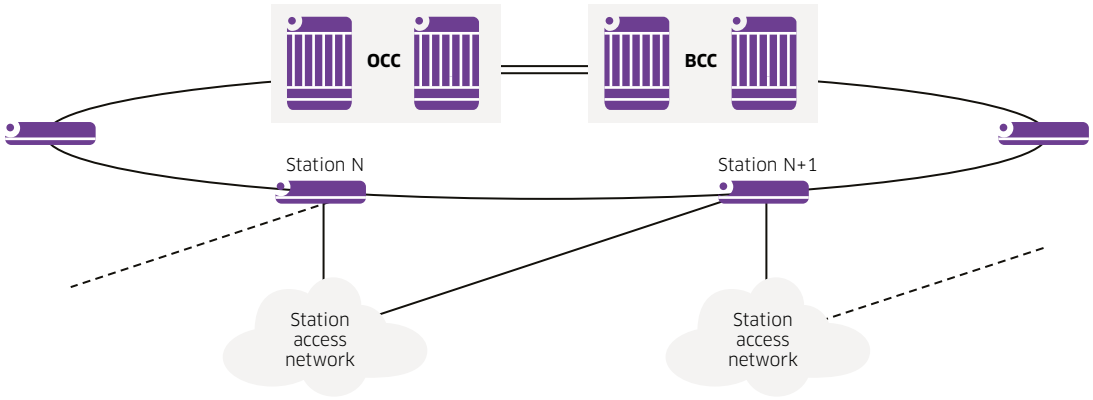

When a site access network is attached to multiple BEBs, there is potential for loops to be created. Section 5.9 will discuss how loops can be avoided or mitigated.

5 Design considerations and guidelines

In this section, we look at various aspects that need to be considered when designing SPB-based networks for the transportation vertical.

5.1 Scalability

In this section, we provide guidelines on the size that the network can scale to and how this relates to switch specifications and design choices.

5.1.1 SPB nodes

As the number of SPB nodes on the network grows, so does the amount of state information in every SPB node (the Link State Database), SPT recalculations happen more frequently and the convergence time increases because SPB uses point-to-point adjacencies which means updates are relayed hop-by-hop. In a ring topology, the convergence time will grow with the number of nodes in the ring.

It should be noted that not all nodes need to be SPB-enabled. In larger networks, SPB will be used only in the backbone while the site access network will be based on traditional Ethernet technologies such as Spanning Tree, Ethernet Ring Protection and Link Aggregation.

5.1.2 SPB services

The PBB header allows for 16M ISIDs, however, there are limits in the number of ISIDs that a single SPB node can support. This specification is most relevant in L2 designs.

Since an ISID represents a broadcast domain, site access VLANs need to be mapped to ISIDs on a one-to-one basis to preserve L2 isolation in a L2 design. As a result, in a L2 design, site access VLANs are globally significant and the total number of site access VLANs network-wide is limited by the total number of ISIDs supported on the OCC/BCC nodes as those nodes will have all VLANs and ISIDs enabled on them. Therefore, L2 designs are limited in total number of site access VLANs that can be supported across all sites.

However, in a L3 design, site access VLANs are only locally significant and it is only the site uplink VLANs which will be mapped to ISIDs on a one-to-one basis and, since these VLANs and ISIDs are shared among all sites, the number of supported ISIDs will not be a limiting factor.

Table 3 specifies the number of ISIDs supported across the range of SPB-enabled OmniSwitch products.

Table 3. Service specification

	OS10K	OS9900	OmniSwitch 6900	OmniSwitch 6860	OmniSwitch 6865
ISIDs	1K	Future	X20/X40/T20/T40: 1K X72/Q32: 8K	2K	2K

5.1.3 FDB

The FDB, or forwarding database, is the L2 (MAC) table. This specification is most relevant in L2 designs.

The size of the FDB becomes relevant in L2 design because the BEBs at OCC/BCC will have all VLANs and ISIDs enabled on them. The size of the FDB is not so relevant in a L3 design because site access VLANs are only locally significant and end-device MAC addresses are only known within the site that they reside in.

Table 4 specifies the FDB size across the range of SPB-enabled OmniSwitch products.

Table 4. FDB specification

	OS10K	OS9900	OmniSwitch 6900	OmniSwitch 6860	OmniSwitch 6865
FDB	32K/slot	128K/slot	X20/X40/T20/T40: 128K X72/Q32: 228K	48K	48K

5.1.4 L3 specifications

L3 specifications are relevant both in L2 and L3 designs because routing is always performed, if not at the site BEB, at the OCC/BCC BEB.

Both in L2 and L3 designs, a route is associated to each site access VLAN. The difference is that, while in L2 designs these routes will only exist at the OCC and BCC BEBs, in a L3 design these routes will normally exist on every BEB unless filtered out with route-maps.

In a L3 design BEBs propagate ARP entries for every end device within the site while in a L2 design the OCC and BCC BEBs will propagate entries for all end devices across all sites which again will limit the scalability of L2 designs.

Systems requiring L3 isolation will be placed in separate VRFs. In a L2 design those VRFs exist in OCC and BCC BEBs only while in L3 designs those VRFs will exist in every BEB.

Table 5 provides relevant L3 specifications across the range of SPB-enabled OmniSwitch products.

Table 5. L3 specifications

	OS10K	OS9900	OmniSwitch 6900	OmniSwitch 6860	OmniSwitch 6865
L3 table	C48, U48, U32E: 16K U32S: 12K	512K	X20/X40/T20/T40: 16K X72/Q32: 12K	64K	64K
ARP table	C48, U48, U32E: 16K U32S: 8K (lowest across all modules)	24K	X20/X40: 8K T10/T40: 16K X72/Q32: 48K In a VC, lowest across all modules (central) or sum (distributed)	16K	16K
VRFs	64	64	64	64	64

5.2 Paths and BVLANs

In a SPB network, a Shortest Path Tree is built for every BVLAN. The ECT-ID influences the SPT tie-breaking logic in a way such that BVLANs assigned with different ECT-IDs will build different SPTs, provided multiple equal-cost paths exist. Load balancing is achieved by mapping different services (ISIDs) to different BVLANs.

As shown in Figure 11, there is a single shortest path between most node pairs in a ring topology. Only nodes located at the antipodes of the ring can communicate over two equal-cost paths.

For this reason, there is no gain in having more than two BVLANs in a ring topology. Furthermore, since every BVLAN builds its own SPT, resulting in additional consumption of resources and CPU cycles, this should be avoided.

Moreover, in transportation networks, the ring topology is used for redundancy and not for the extra bandwidth associated to the alternative path: In the event of a failure, a single path must be able to carry the entire traffic load. It should also be noted that in the event of a failure, both SPTs will be recalculated (no fast re-route) and therefore having two paths will not improve re-convergence time. In fact, it will negatively impact the re-convergence time because double the number of CPU cycles are required to re-compute two SPTs.

In addition, a second BVLAN will only improve link use marginally because there is a single path between most node pairs.

In summary, a single BVLAN is recommended in a ring topology. Two BVLANs will provide a marginal improvement in link use during normal (no failure) conditions but this will come at the expense of increased resource and CPU use.

Figure 11. Shortest paths on a ring topology

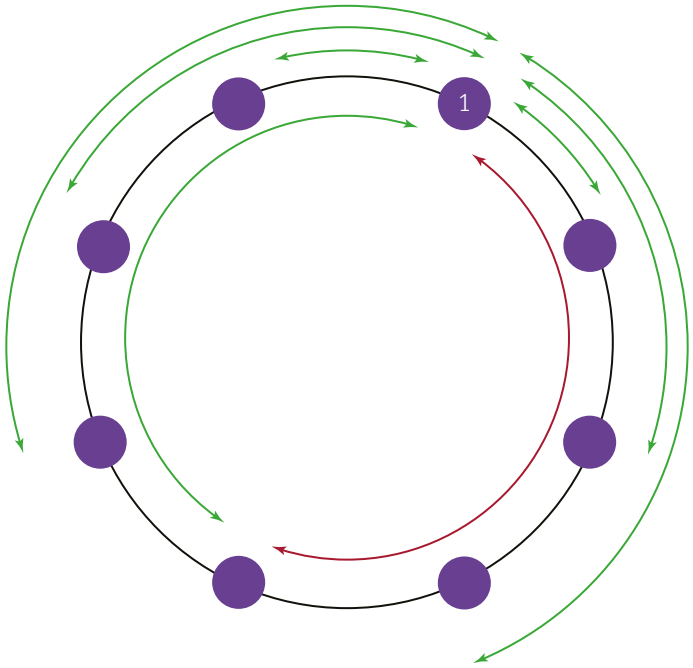

5.3 Link aggregation

Combining multiple physical links into a logical link aggregate (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 some situations, this will not provide enough randomness 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 backbone LAGs. 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.

5.4 Virtual chassis

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

Virtual chassis provides many benefits such as network architecture and management simplification. The designs presented in this document do not make use of the VC feature because control plane independence is preferred such that nodes are failure-independent at the control plane and can run different software releases if needed for maintenance reasons.

VC can be used within the SPB backbone at the data centers and site aggregation or access networks if required.

When using virtual chassis in the SPB backbone, LAGs are recommended to interconnect the VC to all its SPB neighbors, such that one-member (physical) port connects to every slot in the VC as depicted in Figure 12. This is not mandatory, but recommended and will improve the network convergence time in the event of VC unit failure because the need to update tables during the control plane takeover is greatly reduced.

Figure 12. VC and SPB

5.5 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. The metric of the link between OCC and BCC sites can be increased in such a way that site-to-site traffic does not transit through this link as shown in Figure 13.

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 to those links with higher capacity and away from lower capacity ones, thus making the best use of the total available bandwidth and improving performance. Table 6 below shows a way in which the metric can be set to be inversely proportional to the link speed.

Table 6. Suggested link metric

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

Figure 13. Influencing SPT with link metric

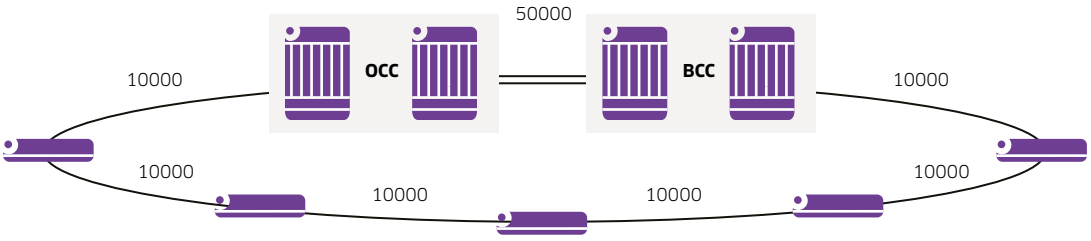

5.5.1 Link metric in a LAG

Another aspect to consider is that the link metric will not change when member ports in a LAG fails. If a LAG is used purely for resiliency but not for bandwidth, the LAG metric should reflect the speed of the member ports without adjustment for the extra bandwidth.

If the LAG is used for the extra bandwidth and one of the member ports fail, the metric will not adjust automatically and traffic will continue using the LAG (if it is in the shortest path) which may lead to saturation even though an alternative path with higher bandwidth may exist. In this case, a simple Python script can be created to dynamically adjust the link metric when the member ports fail or recover.

Alternatively, instead of using LAG, a single higher-capacity link can be used or multiple non-aggregated links can be used and load can be spread across multiple paths on a per-BVLAN basis.

5.6 Quality of service

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

Trusted SAPs copy CoS markings from the incoming CVLAN 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 L2-L4 conditions is possible within the SPB backbone due to the MAC-in-MAC encapsulation.

When using an external router or hairpin loop for routing, the standard VLAN port side of the hairpin must best set to trust and use CoS and not DSCP to preserve CoS markings end-to-end.

Please refer to Table 7 for an example of how various transportation systems and applications can be mapped to traffic classes and per-hop behaviors.

Table 7. Traffic classes and per-hop behaviors

Traffic class	PHB	CoS	Queuing	WRED	Example systems/applications
Network management	AF	7	WFQ	No	SSH, SNMP, HTTPS
Network control	AF	--	WFQ	No	IS-IS, OAM
Real time	EF	5	SP	No	Telephony
Business critical	AF	4	WFQ	No	Ticketing, tolling, admission control, traffic management system, fire and alarm detection
Broadcast	AF	3	WFQ	No	Passenger announcement, passenger information system
Streaming	AF	2	WFQ	No	Video surveillance
Bulk	AF	1	WFQ	Yes	Infotainment
Best effort	BE	0	WFQ	Yes	Internet

Please refer to the AOS Network Configuration Guide [9] for platform-specific QoS details.

5.7 Multicast

In transportation, systems such as video surveillance, passenger announcement and passenger information system can use multicast. Therefore, it is important to discuss how multicast traffic is handled in the SPB network. In this section, we will cover the basics of multicast in a SPB network at L2 and L3. Please refer to Section 5.7.3 for a practical example.

5.7.1 L2 multicast

First, let's review L2 multicast. There are three multicast replication modes in a SPB network. This is applicable not only to multicast traffic, but also broadcast and unknown unicast traffic. We will refer to this traffic as BUM traffic.

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.

Tandem (S,G): In this mode, a separate multicast SPT and FDB are created. The multicast SPT is 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.

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, as can be the case in a transportation network (OCC can be the root bridge).

We can now compare these three modes, please refer to Table 8.

Table 8. Multicast replication modes and suggested uses

	Head-end	Tandem (S,G)	Tandem (*,G)
Operation	Frames replicated at the ingress BEB and forwarded as unicast using the SPT	Frames forwarded as multicast and replicated at the SPT's fork-out points	Frames forwarded as multicast using a shared tree and replicated at fork-out points
Bandwidth efficiency	Low	High	High
Resource use	Low	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 interoperate with third party equipment.

*: When combined with IGMP Snooping, available in AOS 8.4.1R01 and later. Tandem mode can be used as an alternative otherwise.

5.7.2 L3 multicast

Having discussed multicast at L2, we can now discuss multicast at L3.

L3 multicast is based on Protocol Independent multicast (PIM). OmniSwitch products support PIM Sparse, Dense, BIDIR as well as Source Specific multicast. You can refer to the AOS Network Configuration Guide for a description of these multicast modes.

In a L2 deployment, L3 multicast is used when sources and receivers do not reside in the same Subnet. Even if traffic is not routed at the site BEB, it will be routed at the OCC BEB.

Frame forwarding and replication at L2 will depend on the type of multicast forwarding used at L3. Please refer to Section 5.7.3 for practical examples.

5.7.3 Multicast replication examples

In this section, we will look at multicast replication in various scenarios. This is important to plan for ring link capacity, which we will do in Section 5.8.

In transportation, we normally find multicast used in two ways:

- Multicast source is located at the OCC and the multicast receivers are located at the sites. Passenger information, passenger announcement and infotainment systems may fit into this category.
- Multicast sources are located at the sites and multicast receivers are located at the OCC. Video surveillance fits into this category.

There is also the possibility of sources and receivers both residing in the same site. An example of this may be a video surveillance console displaying local CCTV imagery. In a L3 deployment, this traffic will be forwarded locally and not traverse the ring therefore we will not count it for

link sizing purposes. In a L2 deployment, we will only consider the case in which source and receiver both reside in the same VLAN and therefore this traffic will also be forwarded locally and not impact ring link utilization.

In L3 deployments, we will consider PIM Sparse Mode with the rendezvous point located at the OCC because either the sources or the receivers will be located at the OCC and therefore all multicast traffic will flow through the OCC. In addition, Source Specific multicast, when supported by the application, can be enabled such that subscribers receive traffic directly from the source without a RP. This helps when both sources and subscribers are co-located at the site.

In L2 deployments, we can consider two distinct cases:

- All VLAN/ISIDs are shared across all sites
- VLANs/ISIDs are specific to each site

Now we can summarize all possibilities in Table 9 and Table 10.

Table 9. Multicast replication for source at OCC

Multicast source	Design	IP replication	SPB replication	Total traffic
One source at OCC	L2 with Different VLANs/ ISIDs for every site	Once per VLAN/site (N)	Tandem: No	N
One source at OCC	L2 with Different VLANs/ ISIDs for every site	Once per VLAN/site (N)	Head-End: Twice (once for the site and once for BCC)	2N
One source at OCC	L2 with same VLANs/ ISIDs for all sites	No	Tandem: No	1
One source at OCC	L2 with same VLANs/ ISIDs for all sites	No	Head-End: N+1 (once for every site plus the BCC)	N+1
One source at OCC	L3 VPN with RP at OCC	No	Tandem: No	1
One source at OCC	L3 VPN with RP at OCC	No	Head-End: Once per site plus once for BCC	N+1

Table 10. Multicast Replication for Site Sources

Multicast source	Design	IP replication	SPB replication	Total traffic
One source at every site (N)	L2 with Different VLANs/ ISIDs for every site	No	Tandem: No	N
One source at every site (N)	L2 with Different VLANs/ ISIDs for every site	No	Head-End: Twice (once for OCC and once for BCC)	2N
One source at every site (N)	L2 with same VLANs/ ISIDs for all sites	No	Tandem: No	N
One source at every site (N)	L2 with same VLANs/ ISIDs for all sites	No	Head-End: N+1 (once for every other site plus OCC and BCC)	N x (N+1)
One source at every site (N)	L3 VPN with RP at OCC	No	Tandem: No	N
One source at every site (N)	L3 VPN with RP at OCC	No	Head-End: No	N

Referring to Table 9 and Table 10 above, it should be clear why head-end replication is inefficient in terms of bandwidth consumption and tandem replication is recommended instead for multicast-intensive systems.

In a L2 design with multicast sources at the sites and multicast receivers at the OCC, if VLANs and ISIDs are shared among all sites, traffic will explode with the square of the number of sites if using head-end replication. Therefore, tandem replication mode is recommended in this case.

This design is highly discouraged anyway because it has poor scalability in terms of end devices and even in the absence of multicast traffic, broadcast and Unknown unicast traffic still needs to be flooded throughout the network.

5.8 Link sizing and capacity planning

Link sizing must consider the worst case scenario, in which one of the nodes or ring links has failed and a single path remains.

Depending on the specific system, traffic can flow:
- From OCC/BCC to sites (for example, passenger announcement, passenger information system)
- From sites to OCC/BCC (for example, video surveillance archiving)
- Between sites (for example, telephony)
- Local to the site (for example, live video surveillance display)

Link sizing must consider the type multicast replication:
- L2/L3 unicast is not replicated in the ring
- L2 multicast using head-end replication is replicated once per BEB in the same ISID
- L2 multicast using tandem replication sends a single copy on any one link
- L3 multicast is sent to RP as unicast and will not be replicated in the ring
- L3 multicast from RP will be sent as L2 multicast and will be replicated in the ring according to the multicast replication mode.

5.8.1 Scenario: Light rail project overview

In this section, we will show how to do capacity planning by looking at a specific light rail scenario. We will show how bandwidth requirements can be estimated for sizing the required link capacity. Note that this is for reference only. Requirements will depend on the specific solution being deployed (for example, the specific CCTV solution).

This is a single-line light rail consisting of:
- OCC and BCC sites
- 20 stations

This is a L3 VPN design and as such, station access VLANs are local to the station while station uplink VLANs are shared across all stations.

We will consider the systems below:
- CCTV
- Public address (PA)
- Passenger information system (PIS)
- Telephony (Tel)
- Automatic fare collection and ticketing system (AFC)
- Access control system (ACS)
- Time distribution System (TDS)

In this L3 design, PIM Sparse Mode is used and the rendezvous point is located at the OCC. In addition, source specific multicast will be enabled such that CCTV subscribers receive traffic directly from the Source. This helps when both sources and subscribers are co-located at the station. Tandem (S, G) replication is used in ISIDs carrying multicast-intensive system traffic, and in this case, CCTV.

Over the next few sections, we will evaluate some of these systems from a bandwidth point of view.

5.8.1.1 CCTV system

In many transportation networks, CCTV, or video surveillance, is the system generating the most traffic and placing the heaviest load on the network and hence we will discuss it in more detail.

The CCTV system is comprised of:

- 24 IP cameras per station (480 Total)
- 4 x 6-way displays and 2 operator consoles per station
- 12 x quad displays and 8 operator consoles at the OCC

Each CCTV camera generates two streams:

- 4 Mb/s stream for live viewing
- 2 Mb/s compressed stream for archiving

We consider a 20% network overhead on top of traffic requirements above.

CCTV archiving is centralized at the OCC and BCC. Archiving streams will be sent as unicast traffic to the RP at the OCC and from there onwards it will be distributed within the OCC and towards the BCC also. In this manner, a single stream will circulate on the ring. We can calculate this traffic as 20 x 24 x 2 Mb/s + 20% = 1152 Mb/s

CCTV viewing is distributed. We can further distinguish between live viewing and archive viewing.

In this example, live viewing is based on multicast. Local live viewing at the station is forwarded locally with SSM and does not impact ring bandwidth. Central live viewing at the OCC using tandem (S, G) replication will require one copy per stream to be sent over the ring. We can calculate this traffic as 20 x 24 x 4Mb/s + 20% = 2304 Mb/s

Archive viewing is based on unicast.

Two consoles and one of the six-way displays at each station can subscribe to archived streams. We can calculate this traffic as 20 x (2 + 6) x 2 Mb/s + 20% = 384 Mb/s.

Eight consoles and 12 quad-way displays at the OCC can subscribe to archived streams. Normally, those archives would be local at the OCC, however, in the event of failure or maintenance, it may become necessary to access archives stored at the BCC. For this reason, we will calculate this traffic as (8 + 12 x 4) x 2 Mb/s + 20% = 134.4Mb/s.

This is summarized in Table 11.

Table 11. CCTV bandwidth requirements

CCTV system	Streams	Rate	BW	Source	Destination	BW from OCC	BW to OCC
Archiving	480	2 Mb/s	1152 Mb/s	Stations	OCC & BCC		1152 Mb/s
Live viewing @ OCC	480	4 Mb/s	2304 Mb/s	Stations	OCC & BCC		2304 Mb/s
Archive vViewing @ Station	160	2 Mb/s	384 Mb/s	OCC/BCC	Station	384 Mb/s	
Archive viewing @ OCC	56	2 Mb/s	134.4 Mb/s	BCC	OCC		134.4 Mb/s
Sub-total						384 Mb/s	3590 Mb/s

5.8.1.2 Passenger announcement

Passenger announcement uses both live and pre-recorded messages.

Pre-recorded messages are sent from OCC to stations as a file and stored locally at the station.

Live messages originate from the station or from the OCC.

Live messages originating from the station do not consume ring bandwidth.

Live messages originating from OCC consume 128 Kb/s each.

100 Mb/s is reserved for this system such that a 50 MB pre-recorded message can be transferred in five seconds (accounting for 25% overhead) when the network is congested.

There are eight operators at OCC and therefore a maximum of 8 x 128 Kb/s = 1 Mb/s for live messages, this is negligible.

5.8.1.3 Passenger information system

Video files are transferred from the OCC and stored locally at the stations 100 Mb/s are reserved for PIS.

A 1 GB file can be transferred in 100 seconds, accounting for 25% overhead.

5.8.1.4 Other systems

Systems below require minimal bandwidth, 100 Mb/s bidirectional bandwidth is budgeted per system.

- Automatic fare collection and ticketing
- Access control system
- Time distribution system

Telephony

- Two operator telephones and six emergency telephones per station
- 128 Kb/s
- (2 + 6) x 20 x 128 Kb/s = 20 Mb/s

Nevertheless, 100 Mb/s of bidirectional bandwidth is still budgeted for telephony.

5.8.1.5 Traffic matrix

Table 12 below summarizes the traffic requirements for all systems. A 30% buffer is added on top for future requirements.

Table 12. Traffic matrix

System	BW from OCC	BW to OCC
CCTV	384 Mb/s	3590 Mb/s
Passenger announcement	100 Mb/s	~NIL
Passenger information	100 Mb/s	~NIL
Telephony	100 Mb/s	100 Mb/s
Automatic fare collection	100 Mb/s	100 Mb/s
Access control	100 Mb/s	100 Mb/s
Time distribution	100 Mb/s	100 Mb/s
Total system traffic	984 Mb/s	3990 Mb/s
Buffer 30%	295 Mb/s	1197 Mb/s
Grand total	1279 Mb/s	5188 Mb/s

As can be seen in Table 12 above, 10 Gb/s links will be sufficient in this case. However, a simple change in video codec can have tremendous impact on bandwidth consumption.

5.9 Site network attachment

The site access network is attached to diverse BEBs for redundancy. The site access network can be based on ERP or Spanning Tree Protocol. In this section, we will consider both alternatives from a loop prevention point of view.

5.9.1 Ethernet ring protection attachment

There are two different ways in which this can be achieved. Please refer to the AOS Network Configuration Guide for an introduction to ERP and ERPv2.

Let's start by introducing the more general way. In this case, the site access network topology can be an ERPv2 sub-ring as shown in Figure 14. This sub-ring is attached to two BEBs through SAP ports. The sub-ring should not be closed with an additional ring or SAP ports. The sub-ring should only be closed through the SPB backbone. The sub-ring can use R-APS or non R-APS virtual channel. R-APS PDUs will be tunneled through the SPB backbone to other BEBs connecting the service (ISID). Whether it is tagged or un-tagged, the sub-ring's service VLAN must be matched by SAPs at the BEBs. 50 ms should not be expected in every failure mode.

Figure 14. ERPv2 sub-ring attachment through SAP UNI

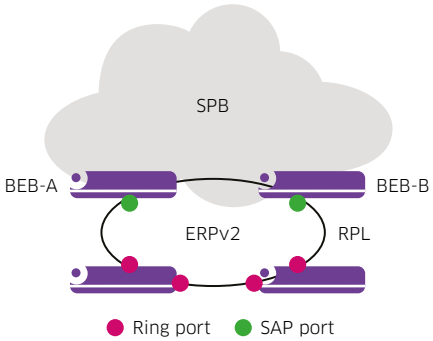

Let's now introduce the second alternative. In this case, we will take advantage of the fact that the SPB backbone topology is also a ring and there is a direct link between both BEBs as shown in Figure 15. ERP-protected VLANs as well as the ERP service VLAN will run alongside BVLANs on the link between both BEBs. In this manner, the ERP ring is closed with ring ports. But since SAP ports cannot be ring ports, a hairpin is used to map the ERP VLANs to SAPs. This method can be used in both L3 and L2 designs. The advantage of this method is that it can deliver 50ms convergence time in the event of sub-ring failure.

Figure 15. ERP ring attachment through ERP UNI

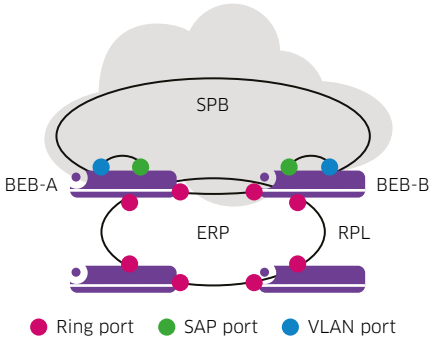

5.9.2 Spanning Tree attachment

When the topology within the site is not a ring, or when the site access network is based on third-party equipment that does not support ERP, Spanning Tree Protocol can be used.

On NNI ports, STP is automatically disabled for all BVLANs. If standard VLANs run alongside BVLANs on NNI interfaces, then STP can be used on those standard VLANs.

The default L2 profile on SAP ports is to tunnel STP BPDUs.

In a L2 design with SAP port attachment, this is appropriate when every site has its own set of VLANs and ISIDs. However, if all sites share VLANs and ISIDs, all site access networks will be in the same STP domain. This is not a good idea. This situation can be avoided by placing each site in a different MSTP region and making use of the max-hop parameter.

In a L3 design with VLAN UNI port attachment, site access VLANs are only local to the site and STP can be used to prevent loops within the site. Site uplink VLANs are shared across all sites and loops will be prevented with STP.

5.9.3 Loopback detection

An SPB backbone with a set of multiple interconnected switches can be logically viewed as a big switch. The big switch connects to the site access network through SAP or VLAN UNI ports.

Miss-configurations and faults at the site access network can create loops spanning both the site access network and the SPB backbone. This can result in broadcast storms. To protect the SPB backbone from broadcast storms, these loops must be detected and broken.

Loopback detection (LBD) can detect and protect the backbone network from forwarding loops created at the site. LBD operates in addition to STP or ERP. 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 loop-back-detection 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 site 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 16 illustrates situations in which LBD can detect and break loops.

Figure 16. Loopback detection

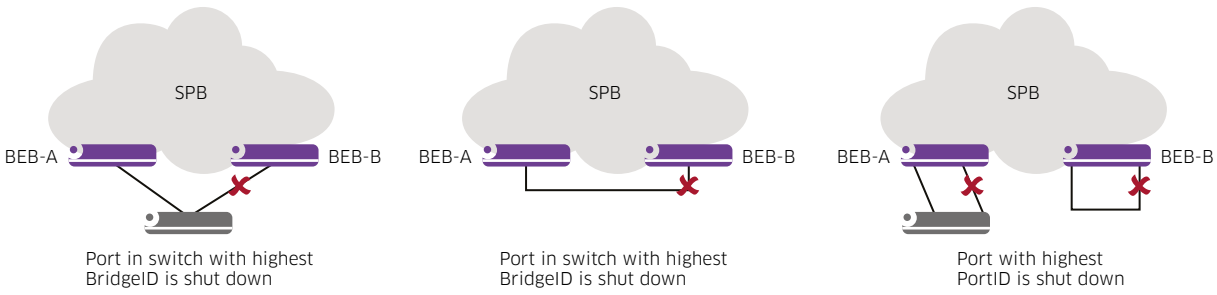

5.10 Provisioning end devices and services – Network profiles

A transportation network can comprise dozens of systems and thousands of end devices. Each device needs to be mapped to the right service, or container. Also, the service needs to be enabled on the access switch that the device is connecting. In addition, differentiated security (ACL) and QoS policies are often required for the various systems and devices.

Figure 17. Authentication, provisioning and policies

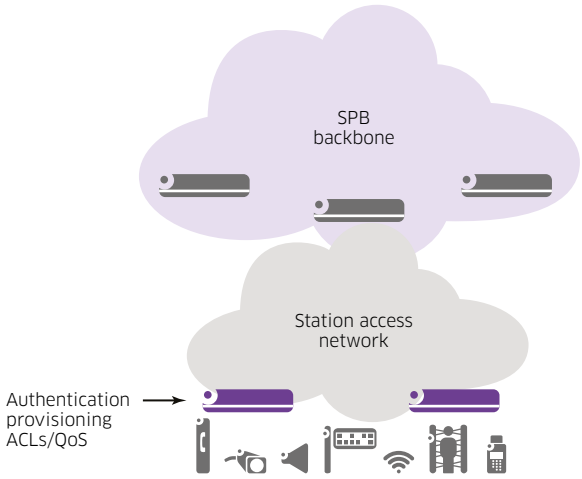

A network profile (NP) is a set of rules that specifies how a device will be bound to a VLAN or SPB service. This binding can be based on a MAC address or range, IP address, VLAN tag or authentication (802.1x or MAC) with a RADIUS server. If the VLAN or SPB service that the NP refers to does not exist on the access switch, it can be dynamically created as well. The NP can also contain ACL and QoS policies such that different security and SLAs can be applied to different systems and device types.

Figure 18. Network profiles

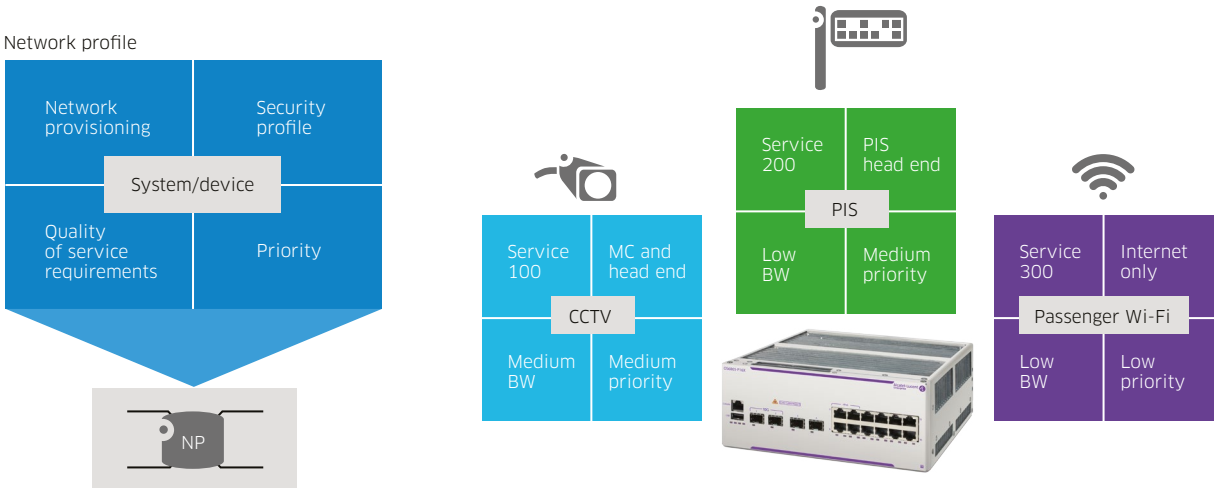

At the site access switch, end devices are mapped to VLANs based on their MAC address or range, IP address or authentication (802.1x or MAC) against a RADIUS server. If the VLAN does not exist on the access switch, it can be dynamically created and added to the uplink through MVRP, provided the VLAN exists on another switch.

At the BEB switch, there are two possibilities depending on the type of attachment (VLAN or SAP UNI). When attaching on a SAP UNI port, a SPB NP profile will bind the traffic to a SPB service based on the incoming traffic VLAN tag. If the SPB service does not already exist on the BEB, it can be dynamically created. When attaching on a VLAN UNI port, the port will be simply configured as a trunk and all required VLANs will be enabled on the port. When using a hairpin loop, both sides of the hairpin will be statically configured.

Network profiles reduce deployment and operations cost because manual configuration tasks are significant reduced, and moves, adds and changes are automated. In addition, network profiles improve security and service levels because only authenticated devices are allowed and differentiated policies are dynamically applied without the burden of manual configuration.

5.11 Network management, monitoring and operations

This section will discuss several aspects related to the management, monitoring, operations and maintenance of an SPB network.

5.11.1 OmniVista 2500 Network Management System

Starting with version 4.3 of Alcatel-Lucent OmniVista® 2500, SPB Topology will be included in the current topology view and will have the ability to:

- Show SPB view of the current map.
- Select a node on the SPB view and show the unicast links of a VBLAN
- Select a node on the SPB view and show the multicast links of a VBLAN
- Show the SPB service ports of all nodes in the SPB network

5.11.2 Element management

OmniSwitch products can be managed through a Console port as well as Telnet, SSH and Web (HTTP/HTTPS) interface (Webview).

Remote management through Telnet, SSH or Web requires IP connectivity. IP connectivity is also required to communicate with a RADIUS server, both for AAA and end-device authentication (NAC) through 802.1x or MAC.

There are various ways in which this can be accomplished:

- Out-of-band management: An out-of-band management network (OOBMN) is a dedicated, physically separate network that is used for management purposes only and not for user traffic. Depending on the specific device, this OOBMN can connect to:
  - ¬ Ethernet management port: The EMP is a dedicated physical port which is present on certain OmniSwitch products such as 10K, 9900 and 6860E. The EMP can only be used for management and not for user traffic.
  - ¬ Standard port: This is a standard VLAN port. A VLAN is dedicated for management purposes and a Loopback IP address will reside in that management VLAN. The standard VLAN port can be locked-down with ACLs to ensure that it can only be used for management. Routing protocols must be configured not to exchange routing updates or route user traffic through the management port.

- In-band management: A special management VLAN is also required in this case and this is where the loopback IP address will reside. The difference is that management traffic will not run on a physically separate network. Again, there are three different ways in which this can be accomplished:
  - ¬ A special management ISID is created and the management VLAN is mapped to this ISID with a hairpin loop and SAP.
  - ¬ The management VLAN runs alongside BVLANs on backbone interfaces. The management VLAN uses STP or ERP.
  - ¬ A special management ISID is created and the loopback management IP resides in this ISID but there is no need to use a hairpin loop. This is called “in-line routing” and is a roadmap feature on OmniSwitch 9900 and 6900. Please contact ALE for availability of this feature.

5.11.3 Operations and maintenance: 802.1ag

OAM in a SPB network is most useful to perform L2 trace and L2 ping for analysis and troubleshooting. Other aspects of OAM such as fault detection, which are important in PBB, are not so 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, please refer to Figure 19. 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 19. OAM in BVLAN and VLAN domains

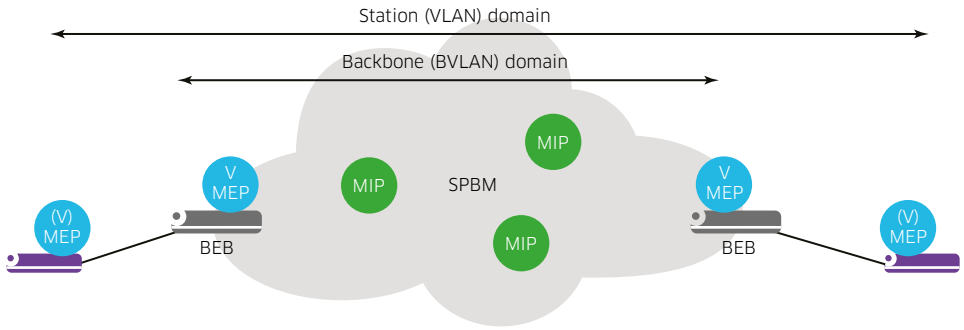

OAM is also supported at the CVLAN level or between site access switches. This is useful in a L2 deployment for testing end-to-end service connectivity between sites or between sites and OCC/BCC. OAM at the CVLAN level must be set at a higher maintenance domain level than BVLAN OAM.

Figure 20 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 20. L2 ping and L2 trace

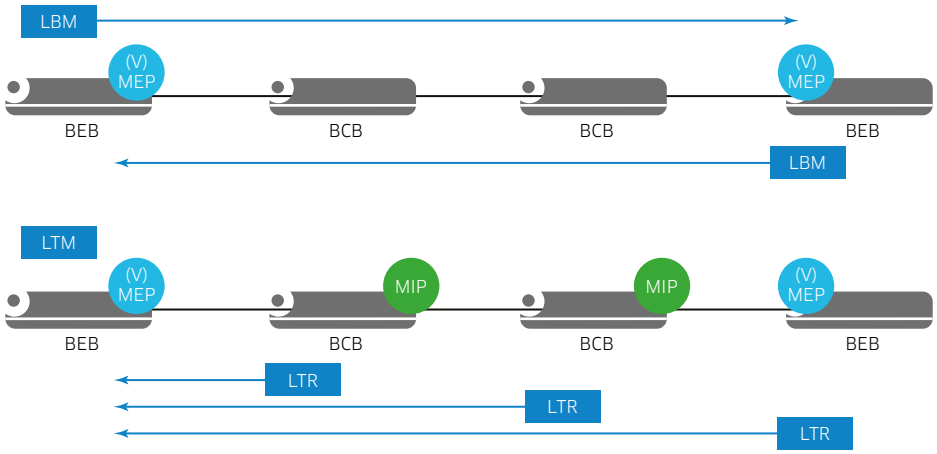

5.11.4 Network performance: Service assurance agent

Latency, jitter and packet loss SAA tests are automatically set-up between all BEBs and BCBs and overall BVLANs with the saa auto-create command.

5.11.5 Network maintenance

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

5.11.5.1 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 equivalent to increasing the metric on all the links but is a much quicker way of achieving this outcome.

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

5.11.5.2 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 neighbor nodes to tear down and re-establish adjacencies with the restarting node and re-build the topology database, resulting in some disruption to traffic flows.

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

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

6 Security considerations

Securing mission-critical transport infrastructure requires a layered “defense-in-depth” approach that combines both proactive and reactive defense mechanisms. Please refer to Figure 21 for snapshot of the various applicable mechanisms.

Figure 21. Defense in depth

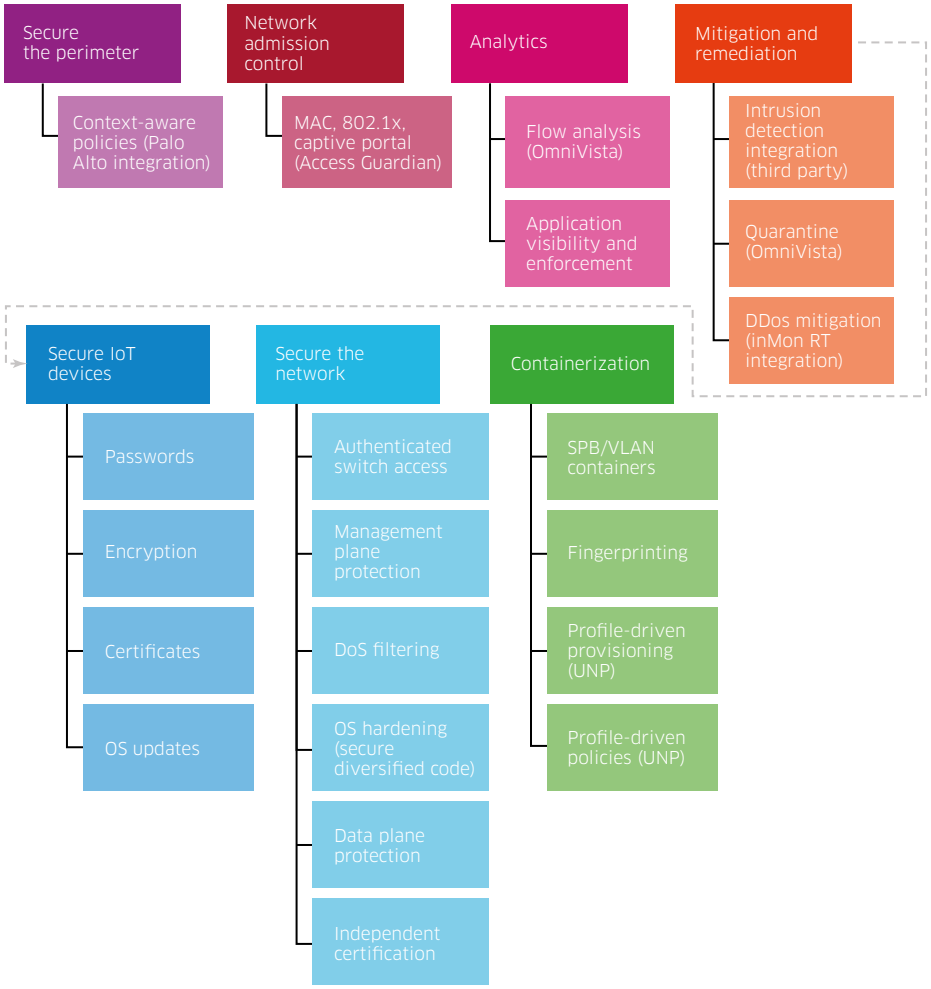

Let’s examine these mechanisms in more detail.

6.1 Securing IoT devices

Transportation systems use multiple Internet of Things (IoT) devices such as sensors, cameras, information displays and ticketing machines. These IoT devices are vulnerable to attack just like other IT assets. A compromised IoT device can also become an attack vector into other devices and systems. Because of the high number of IoT devices, the impact of such attacks can be very high.

These devices must be securely configured and managed. The exact measures will depend on the device capabilities and manufacturer recommendations, but some of these are listed below.

- **Passwords**: Password complexity, renewal and authentication against a central database.
- **Certificates**: X.509 certificates can be installed on IoT devices allowing for mutual authentication between the IoT device and the Server. Certificates can also be used for Network Admission Control (NAC).
- **Encryption**: Secure protocols such as Transport Layer Security (TLS) must be used when managing these devices and any unsecure protocol must be disabled.
- **OS Updates**: IoT devices must be updated and patched according to manufacturer specifications to prevent exploitation of known vulnerabilities.

6.2 Securing the perimeter

When communication between different systems is required, it must only be allowed through firewalls and controlled by fine and specific policies. Different systems may have physical or virtual firewalls of their own if operated by different organizations. Please refer to [4] for an explanation of how integration between Alcatel-Lucent Enterprise's Unified Access solution and Palo Alto Networks firewall can be used to create dynamic fine-grained policies that take the user's identity, device, application, location and other situational factors into consideration.

6.3 Securing the network

Several steps must be taken to secure the network infrastructure itself. An overview of network security mechanisms is given below.

6.3.1 Authenticated switch access and logging

Switch security features increase the security of the basic switch login process by allowing management only through specific interfaces for users with specific privileges. Login information and privileges should be stored on an external server such as Radius or LDAP. External servers should also be used for accounting, which includes logging statistics about user sessions. Admin authentication against the local database may be allowed on the console port only as a failover mechanism in case the external servers become unavailable or in case of misconfiguration.

6.3.2 Management plane protection

Management protocols must be secured, as detailed below:

- Insecure protocols such as Telnet, FTP/TFTP and SNMP must be disabled
- SSHv2 should be used with keys larger than 2048 bits
- HTTPS can be used if required for RESTful API/WebServices access and if so, it should be set up with X.509 certificates for mutual authentication between network node and Server.
- SNMPv3 with authentication and privacy options should be used for monitoring
- TLS 1.1. or 1.2 (preferred) should be used when connecting to remote RADIUS, LDAP or Syslog servers.

6.3.3 Denial of Service (DoS) filtering

By default, an OmniSwitch filters denial of service (DoS) attacks. Some DoS attacks aim at system bugs or vulnerability, while other types of attacks involve generating large volumes of traffic so that network service is denied to legitimate network users. These attacks include the following:

- ICMP Ping of Death—Ping packets that exceed the largest IP datagram size (65535 bytes) are sent to a host and crash the system.

- Land attack—Spoofed packets are sent with the SYN flag set to a host on any open port that is listening. The machine can crash or reboot while attempting to respond.
- ARP flood attack—Floods a switch with many ARP requests, resulting in the switch using a large amount of the CPU time to respond to these requests. If the number of ARP requests exceeds the preset value of 500 per second, an attack is detected.
- Invalid IP attack—Packets with invalid source or destination IP addresses are received by the switch. When such an invalid-IP attack is detected, the packets are dropped, and SNMP traps are generated. Following are few examples of invalid source and destination IP addresses:
  - ¬ Invalid source IP address
    - 0.x.x.x.
    - 255.255.255.255.
    - Subnet broadcast, that is, 172.28.255.255, for an existing IP interface 172.28.0.0/16.
    - In the range 224.x.x.x - 255.255.255.254.
    - Source IP address equals one of switch IP interface addresses.
  - ¬ Invalid destination IP address
    - 127.x.x.x.
    - In the range 240.x.x.x - 255.255.255.254.
    - 0.0.0.0 (valid exceptions- certain DHCP packets).
    - 172.28.0.0 for a router network 172.28.4.11/16.
    - 0.x.x.x.
- Multicast IP and MAC address mismatch—This attack is detected when:
  - ¬ the source MAC address of a packet received by a switch is a multicast MAC address.
  - ¬ the destination IP and MAC addresses of a packet received by a switch is same as the multicast IP and MAC addresses, but the multicast IP and the multicast MAC addresses do not match.
  - ¬ the destination IP is a unicast IP and the destination MAC address is either a broadcast or multicast address. In such a condition, an event is recorded in the DoS statistics. No SNMP traps are generated as valid packets can also fall under this category.
- Ping overload—Floods a switch with many ICMP packets, resulting in the switch using a large amount of CPU time to respond to these packets. If the number of ICMP packets exceed 100 per second, a DoS attack is detected. By default, the detection of attack is disabled.
- Packets with loopback source IP address—Packets with an invalid source address of 127.0.0.0/8 (loopback network) are received by the switch. When such packets are detected, they are dropped, and SNMP traps are generated.
- The switch can be set to detect various types of port scans by monitoring for TCP or UDP packets sent to open or closed ports.

6.3.4 OS hardening – Secure diversified code

Secure diversified code is a technology which mitigates risks at the source, enabling an enhanced security profile through:

- Independent verification and validation of OmniSwitch source code
- Software diversification of OmniSwitch object code to prevent exploitation
- Secure delivery of OmniSwitch software

OmniSwitch AOS secure diversified code protects networks from intrinsic vulnerabilities, code exploits, embedded malware, and potential back doors that could compromise mission-critical operations. The secure diversified code technology is continuously applied on every new AOS release, therefore it will address both current and future threats.

6.3.5 Data plane protection - MACsec

Data integrity and confidentiality must be protected whilst 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.

Please refer to Table 13 for details on hardware and software support for MACsec on OmniSwitch products.

Table 13. MACsec hardware and software support

OmniSwitch model	Hardware support	Software support
OmniSwitch 6465-P6/P12	All ports	AOS 8.5R01
OmniSwitch 6465-P28	All 1G ports and 2 out of 4 10G ports	AOS 8.5R02
OmniSwitch 6860 (24/48/P24/P48)	On 10G ports only	AOS 8.4.1R03
OmniSwitch 6860-P24Z8	On first 16 1G ports and on 10G ports	AOS 8.4.1R03
OmniSwitch 6860E-P24	On 1G and 10G ports	AOS 8.4.1R03
OmniSwitch 6860E (24/48/P48/U28)	On 10G ports only	AOS 8.4.1R03
OmniSwitch 9900	All ports (OS9900-GNI-48/P48 OS9900-XNI-48/P48 OS9900-XNI-U48 OS9900-XNI-P48Z16) except on CNI-U8 module and CMM ports when operating at 40G	AOS 8.4.1R03

In transportation, the main application for MACsec is protecting data integrity and confidentiality while transiting over public spaces outside of the physical security perimeter, where it can be subject to tapping and other malicious activity.

MACsec hardware support in current ALE OmniSwitch portfolio allows for protection of data transiting between site backbone nodes at 10 G speeds when using OmniSwitch 6860 or OmniSwitch 9900 as a site backbone node and between site access nodes when using OmniSwitch 6860 or OmniSwitch 6465 wether site nodes required hardened devices such as OmniSwitch 6465 nodes or not.

Please refer to Figure 22 and Figure 23.

Figure 22. MACsec data plane protection in ERP networks

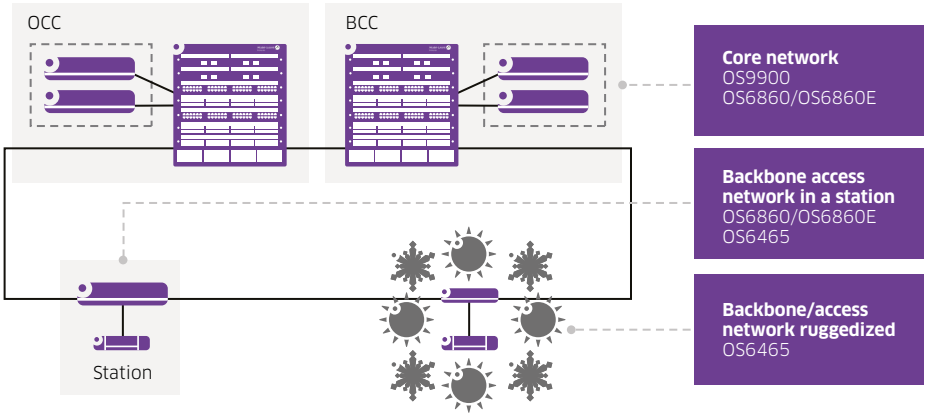

Figure 23. MACsec data plane protection with SPB network

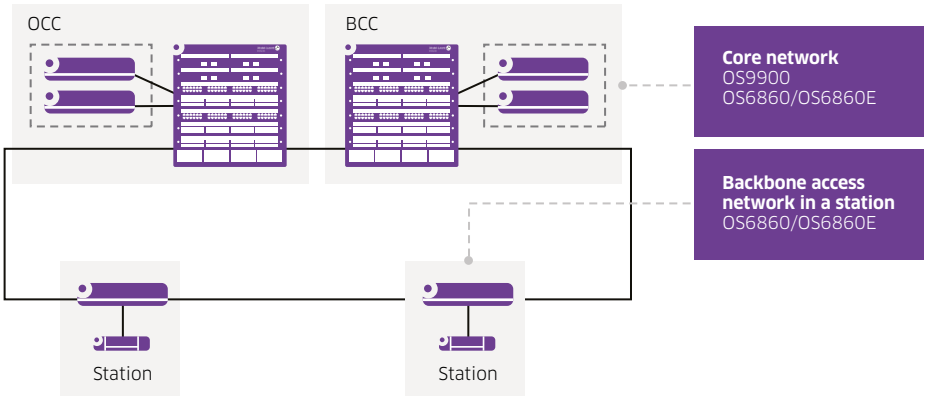

6.3.6 Independent certification

OmniSwitch products have been independently certified to comply with rigorous security standards. Compliance with these standards is mandated in the public sector, but is valuable beyond as well. Independent certification is an objective benchmark to compare security features in different products.

OmniSwitch products are certified to comply with the following standards.

Figure 24. Security certifications

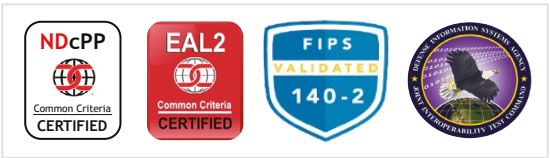

6.3.6.1 Common Criteria

The Common Criteria for Information Technology Security Evaluation, normally referred to as Common Criteria for short, is a set of specifications which serves as a framework in the evaluation of security products.

There are two important aspects in the Common Criteria certification:

- The Protection Profile: This defines the specific security requirements which are relevant to a specific device class such as a Network Device or a Firewall.
- The Evaluation Assurance Level (EAL): This is a number from 1 to 7 representing how thoroughly the product has been tested. Higher EAL levels do not necessarily mean more secure products but rather more thorough testing.

OmniSwitch products 6250, 6350 and 6450 with AOS 6.7.1.79R04 and 6860, 6865, 6900, 9900 and 10K with AOS 8.3.1.348.R01 are EAL-2 (Structurally Tested) against the Network Device Collaborative Protection Profile (NDcPP) set of requirements. Please refer to [10] and [11] for the EAL-2 and NDcPP certificates, respectively.

6.3.6.2 FIPS 140-2

The Federal Information Processing Standard 140-2 (FIPS 140-2) is a U.S. and Canadian security standard focusing on hardware and software solutions using cryptography.

FIPS 140-2 certificates for AOS 6.7.1R04 (used in OmniSwitch 6350 and 6450) and AOS 8.3.1R01 (used in OmniSwitch 6860, 6865, 6900 and 9900) can be found in [11] and [12], respectively.

6.3.6.3 JITC

The Joint Interoperability Test Command (JITC) is the U.S. Department of Defense's (DoD) Joint Interoperability Certifier and only non-Service Operational Test Agency for Information Technology (IT)/National Security Systems. JITC provides risk based test, evaluation and certification services, tools, and environments to ensure joint warfighting IT capabilities are interoperable and support mission needs. OmniSwitch 6860, 6865, 6900 and 9900 are certified, please refer to [14].

6.4 Network Admission Control – Access Guardian

Physical devices attached to a LAN port on the switch can be authenticated using port-based network access control. The following options for authentication are available:

- 802.1X authentication for supplicants.

Uses Extensible Authentication Protocol (EAP) between an end device and a network device (NAS) to authenticate the supplicant through a RADIUS server.

- MAC-based authentication for non-supplicants.

MAC-based authentication does not require any agent or special protocol on the non-supplicant device; the source MAC address of the device is verified through a RADIUS server. The switch sends RADIUS frames to the server with the source MAC address embedded in the username and password attributes.

Internal or external captive portal

Captive portal authentication enables web-based authentication for guest and BYOD users.

The authentication server may return a user network profile (UNP). UNPs map devices to VLANs or services and may contain security and QoS policies. If no UNP is returned, a default UNP is applied.

6.5 Containerization

Containers are virtual network segments where IoT devices and applications that control them are isolated from other devices and applications. This segmentation facilitates enforcement of security policies and limits the damage in the event of a security breach.

SPB intrinsically segments IoT devices into containers by way of MAC-in-MAC tunneling and the ISID field designates the container. Any communication outside of the container will be controlled by a firewall.

At the access layer, IoT devices are authenticated and assigned a network profile based on the device type. This network profile defines the VLAN that the device will be provisioned on to as well as device-type specific QoS and security policies (ACLs).

Links between the site access switch and site backbone switch are 802.1 Q tagged. At the site backbone switch BEB UNI port, VLAN tags are mapped to SPB containers through SAPs.

6.6 Threat mitigation and remediation

So far, we have focused on best practices and preventive measures that can proactively stop a security incident from happening in the first place. But due to the fast-evolving nature of security threats, no security strategy is complete without reactive mechanisms to thwart or dampen the effect of those threats that cannot be proactively avoided.

OmniVista 2500 NMS can be integrated with external intrusion detection systems (IDS) through Syslog. When intrusion or malware attacks are detected by the IDS, a Syslog message is sent to OmniVista. The Syslog message includes the intruder's or attacker's address. OmniVista can use this information to locate the switch and port or AP that the device is connected to and quarantine it by shutting down the port or applying a quarantine profile (restrictive VLAN and ACLs) such that the malicious activity is stopped and remediation activities (for example, OS patching and cleanup) can be performed.

An additional DDoS mitigation method is described in [7]. InMON sFlow-RT collects sFlow data from switches and detects DDoS attacks in real time. InMON notifies the DDoS Mitigation SDN application which in turn instructs the SDN controller to push the necessary rules to drop traffic associated to the attack.

6.7 Analytics

Smart Analytics in OmniVista 2500 brings unprecedented visibility into the network status and usage patterns up to the application level. Understanding patterns assists in fine tuning and enforcing security policies to drive compliance.

OmniVista 2500 collects data from all OmniSwitch products through SNMP and SFlow. Various reports transform this data into valuable information.

Application visibility extends this view to the application layer through deep packet inspection. This means applications can be detected even if running on standard web ports such as TCP ports 80 and 443.

When using the OmniSwitch 6860E as a site access node, application visibility and policy enforcement can be enabled on standard (VLAN) access ports.

7 Product selection

In this section, we will look at the different platforms that can be used to fulfill the various roles in a transportation network. Please refer to datasheets for full details on these products.

7.1 Portfolio overview

Type	Sub-type	Suggested product	Key environmental features	Key network features	Port types	Market-specific certifications
Core node	Modular	OmniSwitch 9900	N/A	• SPB • Virtual Chassis (2) • ISSU • MACsec	• 10/100/1000 Base-T • 10/100/1000 Base-T HPoE • 10 GBase-T • 1/2.5/5/10 GBase-T • 10G SFP+ • 25G QSFP28 • 40G QSFP • 10G/25G/40G/100G QSFP28	N/A
Core node	Stackable	OmniSwitch 6900	N/A	• SPB • Virtual Chassis (6) • ISSU • IEEE 1588v2	• 1/10 GBaseT • 10G SFP+ • 10/25G SFP28 • 10G/40G QSFP • 10/25/40/50/100 G QSFP28	N/A
Backbone node	Modular	OmniSwitch 9900	N/A	• SPB • Virtual Chassis (2) • ISSU • MACsec	• 10/100/1000 Base-T • 10/100/1000 Base-T HPoE • 10 GBase-T • 1/2.5/5/10 GBase-T • 10 G SFP+ • 10 G/40 G QSFP • 10 G/25 G/40 G/100 G QSFP28	N/A
Backbone node	Stackable	OmniSwitch 6900	N/A	• SPB • Virtual Chassis (6) • ISSU • IEEE 1588v2	• 1/10 GBaseT • 10 G SFP+ • 10/25 G SFP28 • 10 G/40 G QSFP • 10/25/40/50/100 G QSFP28	NA
Backbone node	Stackable	OmniSwitch 6860E	N/A	• SPB • Virtual Chassis (8) • ISSU • MACsec • IEEE 1588v2	• 1/10 GBaseT • 10G SFP+ • 10/100/1000/2.5 GBase-T	N/A
Advanced access or collapsed access + backbone	Hardened	OmniSwitch 6865	• Fan-less • Shock, vibration, temperature • DIN mountable • IP-30 rating	• SPB • Virtual Chassis (2) • ISSU • NAC • SDN • IEEE 1588v2	• 10/100/1000 Base-T • 10/100/1000 Base-T HPOE • 100/1000 SFP • 1G/10 G SFP+	• EN 50121 • NEMA-TS2 • DNV
Advanced access or collapsed access + backbone	Non-hardened	OmniSwitch 6860	N/A	• SPB • Virtual Chassis (8) • ISSU • MACsec • NAC • SDN • IEEE 1588v2	• 1/10 GBase-T • 10 G SFP+ • 10/100/1000/2.5 GBase-T POE+/HPOE	N/A

Type	Sub-type	Suggested product	Key environmental features	Key network features	Port types	Market-specific certifications
Value access	Hardened	OmniSwitch 6465	• Fan-less • Shock, vibration, temperature • DIN/wall mountable • IP-30 rating	• ERPv2 • Metro Ethernet • NAC • MACsec • IEEE 1588v2	• 10/100/1000 Base-T POE+/HPOE • 100FX/1 G SFP • 1 G/10 G SFP+ • Alarm relay inputs	• EN 50121 • NEMA-TS2 • DNV
Value access	Non-hardened Gig	OmniSwitch 6450	• Fan-Less (selected models)	• ERPv2 • Metro Ethernet • NAC • Stacking (8) • SDN	• 10/100/1000 Base-T • 10/100/1000 Base-T POE+ • 100/1000Base-X SFP • 1G/10G SFP+	N/A
Value access	Non-hardened mGig	OmniSwitch 6560	N/A	• Stacking	• 10G SFP+ • 10/100/1000/2.5 GBase-T POE/HPOE	N/A
Basic manageable access	Non-hardened	OmniSwitch 6350	• Fan-Less (selected models)	• QoS • NAC • Stacking (selected models)	• 10/100/1000 Base-T • 10/100/1000 Base-T POE+ • 1 G SFP	N/A
Data center core	Modular	OmniSwitch 9900	N/A	• SPB • Virtual Chassis (2) • ISSU • MACsec • SDN	• 10/100/1000 Base-T • 10 GBase-T • 1/2.5/5/10 GBase-T • 10 G SFP+ • 25 G QSFP28 • 0 G QSFP • 10 G/25 G/40 G/100 G QSFP28	N/A
Data center core	Stackable	OmniSwitch 6900	N/A	• SPB • Virtual Chassis (6) • ISSU • SDN • FCOE • VXLAN • IEEE 1588v2	• 1/10 GBase-T • 10 G SFP+ • 10/25 G SFP28 • 10 G/40 G QSFP • 10/25/40/50/100 G QSFP28 • FCOE	N/A
Data center top of rack (TOR)	Stackable	OmniSwitch 6900	N/A	• SPB • Virtual Chassis (6) • ISSU • SDN • FCOE • VXLAN • IEEE 1588v2	• 1/10 GBase-T • 10 G SFP+ • 10/25 G SFP28 • 10 G/40 G QSFP • 10/25/40/50/100 G QSFP28 • FCOE	N/A

7.2 New hardened switch - OmniSwitch 6465

OmniSwitch 6465 is a new switch that completes the product line of hardened devices for IoT and industrial deployments within the OmniSwitch portfolio. It comes in three flavors, P6 and P12 models are available with release 8.5R1 while the P28 model will be available in the 8.5R2 release. OmniSwitch 6465 is generally considered a value switch with advanced L2 features and perfectly complements the already available advanced L3 hardened switch OmniSwitch 6865.

OmniSwitch 6465 is a family of compact, fixed configuration, DIN mount models that are designed for industrial applications capable of operating in a temperature range between -40C to 75C degrees, therefore suitable for extreme temperature conditions. Both models are fan-less supporting MACsec for additional data plane security and IEEE1588v2 on all ports. HPoE (60W) is supported, in addition to IEEE 802.3at (30W) and IEEE 802.3af (15W), to power IoT devices such as PTZ camera for video surveillance.

Figure 25. OmniSwitch 6465 P6/P12 models

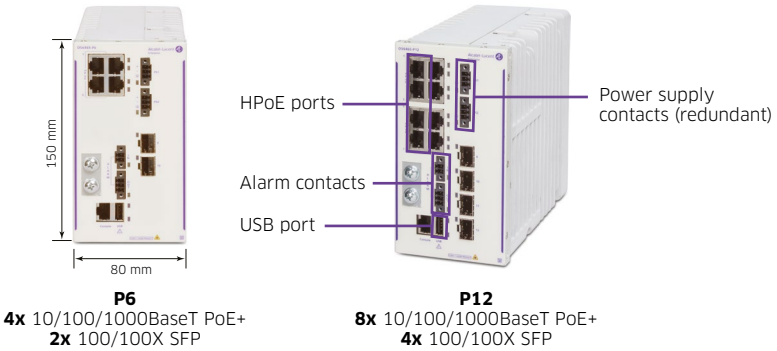

There are two new hardware features implemented on OmniSwitch 6465 family: Switch backup and restore through USB port and alarm contacts.

7.2.1 USB drive as backup and restore

If a USB drive is plugged in, switch will, by default, store the image files as well as the configuration files to the USB storage drive. Whenever “write memory” or copy running-config working are applied, the config will be stored to USB drive as well.

The switch config and image should be able to be restored from the USB drive. It is possible to install the image and config present in the USB to a new switch. A facility should be provided to prompt the user to boot with the images and configuration present on the USB, if a USB is present on boot up.

This feature enables non-expert to efficiently replace a faulty switch in the field with a new one using USB storage drive.

7.2.2 Alarm relay configuration

Alarm relay out is user configurable for associating a failure event to the alarm relay output. The events like power supply failure, port failure, temperature exceeded condition and an alarm if session console is disabled and someone tries to access the switch can be associated with the alarm relay output. Ideally any failure condition either physical or software related (which results in a SNMP tray,) could be associated to alarm relay output.

External input is connected to Alarm Relay Input. AOS provides configuration to the user to map this Alarm Relay input event to Trap, Syslog or Alarm Relay Output.

8 Long-term support (10Y)

Transportation networks are long-term investments and therefore require long-term support – usually longer than 10 years – which extends after the End-Of-Life (EOL) period. ALE commits to providing long-term support contracts providing the same components and associated services that extend after the End-of-Life (EOL). Long-term support contracts of ten years is available for all the products running 8.5R1 and higher versions (OmniSwitch 6465, OmniSwitch 6865, OmniSwitch 6560, OmniSwitch 6860, OmniSwitch 6900, OmniSwitch 9900, OmniVista) and include:

- Hardware support
- Remote technical support (phone, e-mail, webmail) 24x7
- Access to upgrades and updates
- One maintenance release per year, after end-of-maintenance
- Access to knowledge base
- SR status access

Long-term support must be ordered at least six months before the last order day (LOD) for products of interest.

Figure 26. Long term support overview

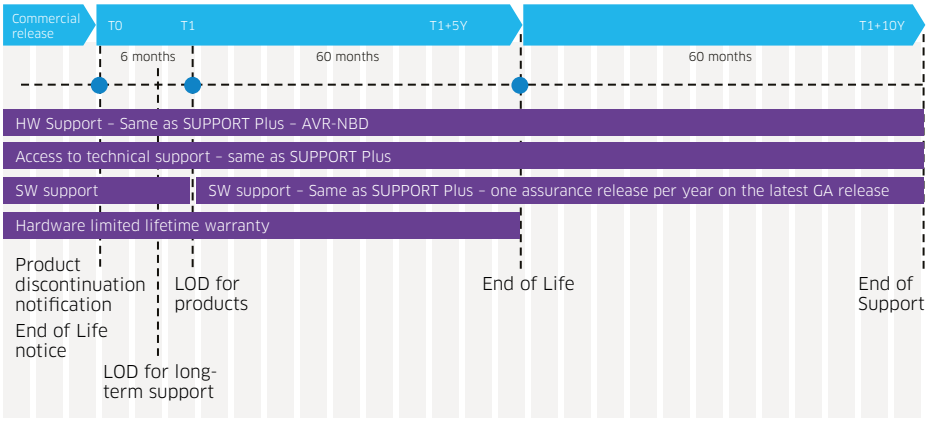

9 Location based services (LBS)

There are multiple reasons why airports are facing increased demand for deploying location-based services today:

- Improving customers experience by providing information about available parking spots, the fastest way to the gate, security lines with the shortest waiting time, accessible restaurants and shops within waiting time or on customers' path, locating people of interest, etc.
- By increasing the number of travelers' airports increase their revenue since about 50% of the revenue comes from parking and retail.
- Retail would be able to send push notifications with special offers and discounts targeting and influencing consumers for unplanned shopping or recognizing customer loyalty and other programs of interest
- Airport security, paramedics, maintenance services and others can use the tracker to locate and mobilize staff quickly, should they be required

9.1 LBS use case for airports

Indoor location based services ALE provides include geo-positioning and wayfinding, geofencing notifications and people tracking. Asset tracking service is also planned to be supported in Q4 2018. OmniAccess Stellar LBS Cloud Manager offers multiple dashboards and powerful location analytics providing essential input for business analytics.

Here is a typical use case for airports:

1. When a traveller arrives at an airport parking lot, the Smart Park feature of the mobile application that starts automatically and stores the car's location.
2. Notifications are sent to the traveler when they enter the airport about their departure gate and boarding time.
3. When a traveler opens the app, they can see where they are in the terminal and can find the nearest check-in and security line with the least amount of people.
4. Once through security, the traveler can enter the flight number to get directions to their gate. They are also provided with an estimated time of how long it will take to get there.
5. While walking near the duty-free shop, restaurants, or other retailers, promotion coupons or notifications can be triggered.
6. Travelers can browse the map and click on a point of interest (POI) to get directions or look for POI in the search bar then see the location on the map.
7. Travelers can share locations with their Facebook or LinkedIn community
8. When back from a trip, travelers will retrieve their car with the Smart Park feature.

Figure 27. Airport use case for LBS

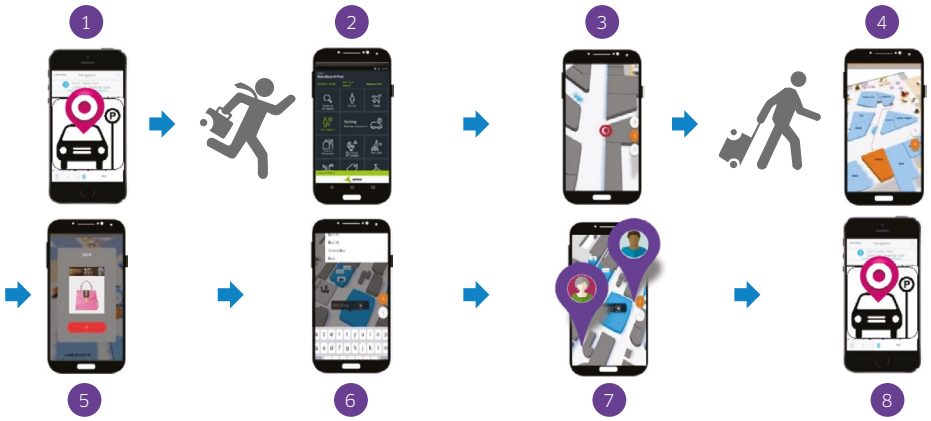

9.2 ALE LBS components and location technology

There are three major components needed for an LBS solution:

1. Beacons, as a signal source, are needed for location calculation and can be delivered in the form of:
   a. BLE (Bluetooth Low Energy) beacon, which is a standalone beacon that can be delivered with an indoor or outdoor installation kit. The external kit includes waterproof silicon case and strapping kits. A BlueSpot beacon from ALE is equipped with a military grade battery that should last from 5 to 10 years.
   b. OmniAccess Stellar BLE dongle connected to a Stellar AP's USB or any other fixed device with a USB interface. The USB interface is used solely for powering up the BLE dongle. It is not recommended to connect the dongle to devices that are easily moved such as a desk phone.
   c. Stellar AP 1230AP which is equipped with an integrated beacon for LBS

2. User application which is used for calculating and positioning on the map from the calculated signal strength received from surrounding beacons and geofencing notifications. Calculations are done by the user application on the phone itself. Therefore, loss of communication with management will not interrupt the location service.

3. Management tools such as LBS installer and LBS Cloud Manager. These are used to help with the installation of the beacons, to do calibration and to set up the notifications. The monitoring and analytics are done from the cloud.

Figure 28. LBS components

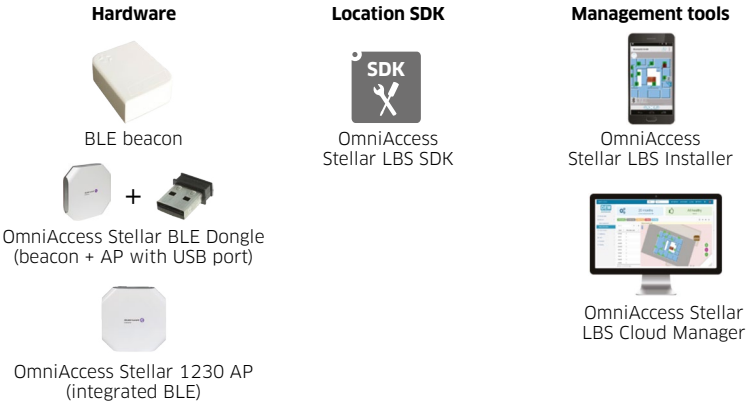

The above-mentioned components work together in the following way to provide LBS services:

- A customer mobile application installed on the mobile phone (or another terminal) must be integrated with the Stellar Location SDK and Mapping SDK. This mobile application is used to receive and process information to calculate the location locally on the mobile device. Information received is signal strength received from deployed beacons.
- The user device captures geofencing events within a building, or a zone within a building, and stores the data in the user device until the data is uploaded periodically to the LBS Cloud Manager.
- The data is stored and can be displayed on a variety of analytic dashboards – overview dashboard, user's dashboard, visit dashboard, and a zone's visit dashboard.
- The upload intervals can be configured and if there is an Internet connection service interruption, the device will continue to capture and store the geofencing events data.
- Once the Internet connection service is up and running again, the device will automatically upload the data to the LBS Cloud Manager to update the analytics with the information received.
- The Analytics module in the LBS Cloud Manager is responsible for helping organizations understand patterns, analyze how certain businesses are doing, optimize LBS in physical areas, maximize the use of staff, and so on.

9.3 LBS project for an airport sample configuration

The rough estimate of project costs (list price) and activities are given in the following example for an airport with location and geofencing services which area size is 500,000 square meters (sqm) and (average) beacon density of four beacons per 1000 sqm is required.

Calculating that three of four beacons will be deployed as standalone beacons while one of four will be USB dongles installed in the APs (AP price is not calculated here), the price for beacons results in $132k.

Professional services cost for beacon installation, calibration and test, project management, project configuration on the cloud and integration support are roughly estimated at $66k. These costs summarize upfront CAPEX cost of $198k.

Annual recurring fees are divided between LBS licensing (Usage of the location and geofencing capabilities of the SDK, Cloud Manager and Installer Application) estimated to $350k and initial 2D mapping with the access to cloud editor to do map updates estimated to $27k. Total annual cost would be $377k in this case.

10 On-board communication solutions

On-board communication, whether it is required on trains and metros or vehicles such as a bus, police car or fire track, could be provided with the latest communication on-board platforms. H2-Automotive+ router is an in-vehicle router providing Wi-Fi and 3G/4G connectivity for services such as increased security and safety of passengers and driver, entertainment services and Internet access as an additional benefit for passengers as well as monitoring and management of the vehicle itself. H2-Rail platform is a specific communications platform for highspeed, commuter, metro, light rail, subway, and freight trains. This platform improves passengers' experience offering a range of entertainment applications such as music, games and movies, internet access and increased security on trains. Perfect additions to the H2-Rail platform are access points (APR-222ac) which are used to extended the prior-mentioned services throughout rolling-stock.

Both communication platforms, H2-Automotive+ and H2-Rail, share the same platform and features but differ in connectors and certifications specific for transportation mean.

10.1 Positioning

Both communication platforms are well positioned. They target more complex solutions and provide multiple services over a single on-board unit compared to low-cost solutions available that provide support for a single service such as on-board PC or 3G communication only. Although initial investment in these platforms are higher, it allows for easy expansion by adding additional SIM cards or Wi-Fi modules or introduction of additional services over the same platform while still managed with the same tool in an intelligent way.

Figure 29. On-board communication platform positioning

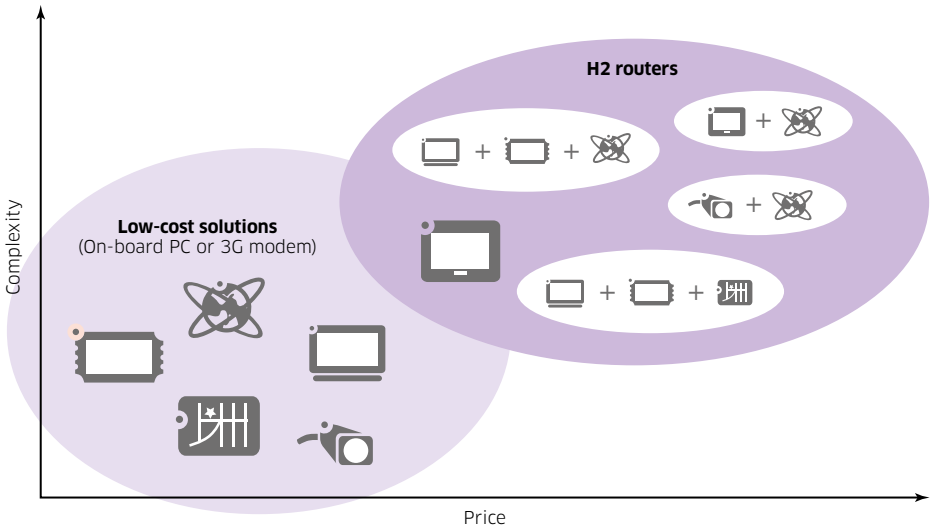

Some of advanced features that could be used as a key selling points and differentiators of H2Auto+ and H2Rail from low-cost products are:

- High-availability by dual SIM or dual Wi-Fi modules that could be used in load balancing or backup mode. Advanced quality link monitoring could be used as a trigger for switching between WAN options.
- Redundancy options by supporting VRRP (Virtual Router Redundancy protocol)
- Secure communications (such as CCTV, alarms) between vehicle and OCC/BCC over IPSec tunnels
- High device throughput with and without encryption
- LTE aggregation
- AP to Wi-Fi client transition according to geofencing capability – AP which is used in vehicle for providing Wi-Fi services to passengers' wireless clients (smart phones, tablets) turns into Wi-Fi client mode once entering the area of vehicle depot to download new entertainment content from depot Wi-Fi network
- Delayed power-off capability – While vehicle is temporary powered-off at station all passengers and driver still benefit of active entertainment and security services. On-board router can still perform maintenance upgrade or upload entertainment content although vehicle is turned – off and without supervising
- Isolation of specific services by supporting VLANs, VRFs, QoS
- Easy deployment and monitoring via Colibri platform

The features just listed are used to build some of the services that will improve a customer's experience. These services can be jointly offered over a single on-board router as follows:

- Dynamic marketing – pushed marketing messages based on location of vehicle route
- On-board Wi-Fi
- On-board entertainment (TV, gaming, shopping, etc.)
- CCTV for security of drivers and passengers
- Ticketing
- Fleet tracking (for vehicles)
- Predicted maintenance

10.2 Use cases

One deployed use case is the deployment of H2 Automotive routers in police cars in a European country. The laptop used by police officers within the police car is connected to the Ethernet port of H2 router. Each H2 router is equipped with two SIM cards. One SIM card is meant for the general access to the internet while the other card is using communication over IPSec protocol with authentication and encryption of data to the private police data center where a VPN concentrator is installed. Although it is possible to configure secure and unsecure communication over a single SIM card, having two SIMs increases the security level.

Another typical use case is on-board Wi-Fi for passengers and internet access through an LTE interface. While a bus is on the road, a Wi-Fi interface in access point mode allows passengers access to the internet. The same interface automatically changes to client mode once the bus reaches the depot.

The router will use this mode to upload and download information between the bus and an internal network of the bus company using Wi-Fi rather than the LTE interface for higher speeds, security and more affordable communication. The Wi-Fi interface changes the mode based on probes for Wi-Fi network reachability or using geofencing capabilities of the management platform Colibri.

A CCTV camera installed for the security of passengers, the bus driver and ticketing device are connected to Ethernet ports of the router.

Figure 30. H2 Automotive+ use case

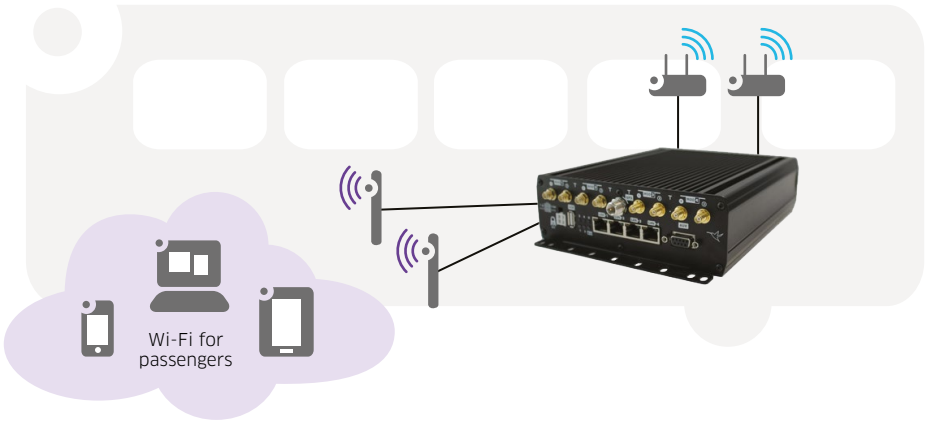

10.3 H2 routers and Colibri NetManager management tool

As previously mentioned both routers are built on the same platform sharing the same hardware capabilities:

- WAN capabilities:
  - ¬ Up to two LTE radios (4 for H2 Rail)
  - ¬ One Wi-Fi radio as WAN (with simultaneous AP)
- LAN/management:
  - ¬ Four GigE ports
  - ¬ One Serial RS-232
- WLAN: Up to two Wi-Fi radios (a/b/g/n/ac)
- Extended temperature: -25˚C to 70˚C
- Performance: 470 Mb/s clear and 70 Mb/s encrypted
- Fan-less devices designed for high vibrations and shock tests
- Power supply protection for direct feeding from battery

As well as some advanced features such as:

- Multi-WWAN fall back, load-balancing and aggregation
- Advanced Link Monitoring features
- GPS with state of the art features and dynamic configurations according GPS positioning
- Advanced troubleshooting and management features

Figure 31. H2 automotive + and H2 rail routers

H2 Automotive + H2 Rail

Colibri NetManager tool is a cloud based network management tool used for easy connecting and configuring routers and services. It provides visibility on deployed devices and status of the performance and devices to ease operations. In addition, it gives a more granular view of which service generates data traffic as well as total consumption per SIM, threshold and alerts for data consumption.

Figure 32. Colibri NetManager functionalities

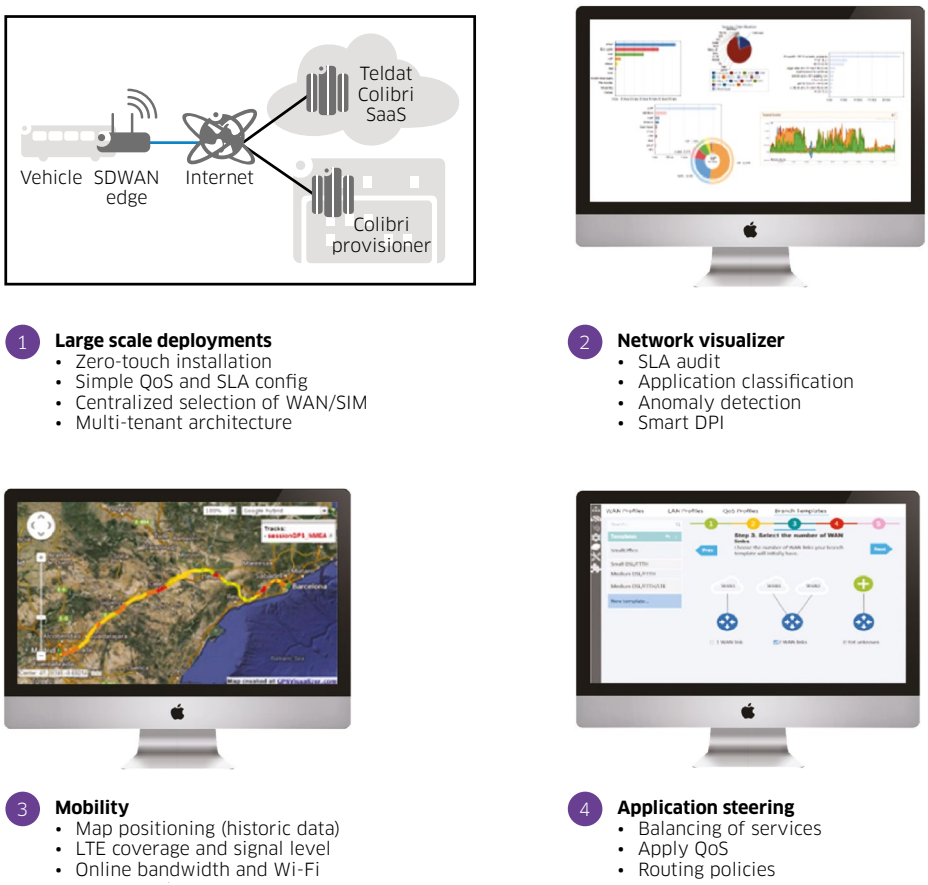

11 Conclusion

Transportation networks are under pressure to improve safety, service operation and to keep drivers and passengers informed and connected. Trends in mobility and internet of (IoT)things not only increase demands for bandwidth and power but also introduce new security and operational challenges.

Alcatel-Lucent Enterprise SPB-based Intelligent Fabric technology creates a single converged network that meets the present and future needs of Transportation operators with simplified operations and reduced TCO.

This design guide has provided practical guidelines that will assist the network architect and network engineer in designing and managing ALE iFab-based transportation networks.

12 Acronyms

AC	access control
ACPS	Alcatel-Lucent Enterprise Certified Pre-Sales
AFC	automatic fare collection
ATC	automatic train control
ATO	automatic train operation
ATP	automatic train protection

ATS	automatic train supervision
BCB	backbone core bridge
BEB	backbone edge bridge
B-DA	backbone destination address in 802.1ah PBB header
B-MAC	backbone MAC address
B-SA	backbone source address in 802.1ah PBB header
B-VID	backbone VLAN ID in 802.1ah PBB header
B-VLAN	backbone virtual LAN
BLE	Bluetooth Low Energy
Bridge ID	64-bit quantity = (Bridge Priority:16) <<48 | SYSID:48
Bridge Priority	16-bit relative priority of a node for tie breaking
BUM	broadcast, unknown unicast and multicast
C-MAC	customer MAC. Inner MAC in 802.1ah PBB header
C-VID	customer VLAN ID
C-VLAN	customer virtual LAN
DDoS	Distributed DoS
DoS	Denial of service
EAP	Extensible Authentication protocol
EC	Emergency call
ECT-ALGORITHM	32-bit unique id of an SPF tie breaking set of rules
EOL	End-of-life
ERPv2	Ethernet Ring Protection version 2
FDB	filtering database: {DA/VID}->{next hops}
HPOE	High PoE
IDS	Intrusion detection system
IOT	Internet of Things
I-SID	logical grouping identifier for E-LAN/LINE/TREE UNIs
ISSU	In-service software upgrades
ITS	intelligent transportation system
LAN	local area network
LBS	Location based services
LOD	Last order date

LSDB	link state database
LSP	link state packet
MAC-IN-MAC	Ethernet-in-Ethernet framing as per 802.1ah[PBB]
MDT	multicast distribution tree
MT-ISIS	multi topology IS-IS as used in [MT]
MT	multi topology. As used in [MT]
NLPID	Network Layer Protocol Identifier: IEEE 802.1aq= 0xC1
OAM	operations and maintenance (802.1ag)
OOBMN	out-of-band management network
Q-in-Q, QinQ	additional S-VLAN after a C-VLAN (802.1ad) [PB]
PA	public address
PBB	provider backbone bridge - forwards using PBB
PIS	passenger information system
POE	Power over Ethernet
SDN	Software-defined networking
(S,G)	source and group - identity of a source specific tree
(*,G)	any source and group - identity of a shared tree
SPB	Shortest Path Bridging – 802.1aq
SPBM, SPB-M	Shortest Path Bridging – Mac-in-Mac mode
SPOF	single point of failure
SPT	shortest path tree computed by one ECT-ALGORITHM
S-VLAN	Service VLAN
TC	toll collection
TIS	traveler information system
TMS	traffic management system
VS	video surveillance
VSL	variable speed limit

13 Related documents

[1] 802.1aq-2012 - IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks--Amendment 20: Shortest Path Bridging: https://standards.ieee.org/standard/802_1aq-2012.html

[2] IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging draft-ietf-isis-ieee-aq-05.txt: https://tools.ietf.org/html/draft-ietf-isis-ieee-aq-05

[3] IEEE P802.1ah.D4.2, Supplement to Virtual Bridged Local Area Networks: Provider Backbone Bridges, March 26, 2008: https://ieeexplore.ieee.org/document/4489712

[4] OmniSwitch Common Criteria EAL-2 Certificate: https://www.commoncriteriaportal.org/files/epfiles/Certification%20Report%20Omniswitch.pdf

[5] OmniSwitch Common Criteria NDcPP Certificate: https://www.commoncriteriaportal.org/files/epfiles/Certification%20Report%20ALE%20NDcPP.pdf

[6] OmniSwitch AOS 6.7.1R04 FIPS 140-2 Certificate: https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/3071

[7] OmniSwitch AOS 8.3.1R01 FIPS 140-2 Certificate: https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/2996