



EVOLVING ENTERPRISE NETWORKS WITH SPB-M

APPLICATION NOTE

EXECUTIVE SUMMARY

Enterprise network managers are being forced to do more with less. Their networks are growing in size and complexity. They need streamlined solutions to improve the performance and simplicity of their networks. Alcatel-Lucent Enterprise has added many new technologies to address these needs. Key among them is Shortest Path Bridging MAC in MAC mode (SPB-M). SPB-M helps improve network performance and stability while offering great flexibility.

SPB-M offers many improvements over older bridging and routing technologies including greater scalability, better bandwidth through improved link utilization, automation, while being compatible with existing technology. SPB-M on the OmniSwitch works in conjunction with Alcatel-Lucent Enterprise smart network edge and access features while providing MPLS like scalability and service separation.

TABLE OF CONTENTS

1. New network requirements call for new network technologies / 1
2. Introducing SPB-M / 1
 - 2.1 What is SPB-M? / 1
 - 2.2 Why SPB-M? / 2
3. Deploying SPB-M in the Enterprise: An example / 3
 - 3.1 Phase One: The Datacenter / 3
 - 3.2 Phase Two: The datacenter and distribution / 4
 - 3.3 Phase Three: Bringing in the access networks / 5
4. Conclusion / 6

1. NEW NETWORK REQUIREMENTS CALL FOR NEW NETWORK TECHNOLOGIES

Intelligent techniques are required for networks to meet these demands and maximize the available bandwidth, minimize latency and ensure stable connectivity. In response, Alcatel-Lucent Enterprise has included Shortest Path Bridging MAC in MAC mode (SPB-M) as part of the OmniSwitch™ family of products. SPB-M (IEEE 802.1aq) is a new bridging standard that combines forwarding decisions based on the Intermediate System to Intermediate System (IS-IS) routing protocol with a media access control (MAC)-in-MAC data plane (IEEE 802.1ah) to allow network operators to build larger and better-performing bridged networks compared to classical networks based on Spanning Tree Protocol (STP).

2. INTRODUCING SPB-M

2.1 What is SPB-M?

Shortest Path Bridging-MAC (SPB-M) essentially combines two separate technologies: Provider Backbone Bridging (commonly referred to as MAC-in-MAC) and the IS-IS routing protocol, which together help build a better 'bridged' network.

In a traditional bridged network, STP calculates and ensures a loop-free forwarding topology in the network. Frames are initially flooded and then learned by bridges to provide connectivity between other bridges, endpoints and routers.

SPB changes this paradigm by encapsulating frames from the edge of the SPB-M network in a MAC-in-MAC header, and then transmitting them within a specially designated backbone virtual local area network (BVLAN), which in turn is controlled by a special SPB-M IS-IS instance. IS-IS controls traffic forwarding through the SPB-M network. Each ingress frame to the SPB-M network includes a service identifier as part of its MAC-in-MAC header, called an I-SID. The I-SID denotes the network source of the frame, and represents a single bridging domain. The term for configuring matching I-SID and VLAN on a port is called the Service Access Point (SAP). When frames egress the SPB-M network, the MAC-in-MAC header is removed and the frames are forwarded using their native VLAN header.

Compared to traditional layer-2 networking technologies, there is no flooding or learning in the SPB-M backbone network because IS-IS determines the optimal shortest path connection between all of the bridges, and enforces correct traffic forwarding by controlling the MAC entries in the BVLANs of all the bridges in the backbone network. Frames that need to be transmitted to all points of a specific service are actually multicast to all member nodes of that service. The bridge at the edge of the SPB-M network will construct a multicast MAC address based on the system MAC address of that bridge and the I-SID. IS-IS maintains multicast forwarding entries for all of the services and source edge bridges in the backbone network.

2.2 Why SPB-M?

SPB-M in the Alcatel-Lucent Enterprise OmniSwitch offers many advantages over classic and other modern bridged networking technologies.

- Through new extensions to IS-IS, SPB-M builds a loop-free and stable topology that uses all optimal links in the network to achieve the best utilization.
- SPB-M supports up to 16 distinct virtual forwarding topologies in the same network. This enables SPB-M to further increase overall network performance.
- SPB-M BVLANS only forward traffic injected by established neighboring links. Other traffic type on the BVLAN is dropped. The same links could simultaneously transport other 802.1Q tagged or double-tag vlans.
- 802.1Q VLANs in the backbone network can be used for other purposes, such as management or IP routing.
- SPB-M is deployed in a service-based approach. Each service is created and bound to local access ports as needed by creating SAPs. IS-IS distributes the service information and automatically builds the topologies to connect all the endpoints to the service. This is done in a non-destructive manner, allowing networks to grow in complexity without rebuilding the network itself.
- Each SPB-M service represents a single layer-2 virtual network, and the protocol can scale up to 16.7 million separate services using its 24 bit I-SID field. This easily enables highly virtualized networks that far exceed the 4K limit of the traditional VLAN tag format.
- The MAC-in-MAC encapsulation protects bridges from having to learn the entire MAC table within each service. Bridges at the edge of the network see only the MAC addresses of the services they are terminating.
- SPB-M has built-in multicast capabilities and calculates a forwarding tree per service to distribute multicast, flooded, and unknown unicast traffic. SPB-M only transmits copies of frames to the end-points of the network that have ports attached to the same service.
- SPB-M is an official standard. It fully interoperates with older standards-compliant equipment, both within the SPB-M network and when connecting other equipment to the SPB-M network. This allows administrators to deploy SPB-M in a phased approach, avoiding ripping and replacing installed networks.
- The shape and size of SPB-M backbone networks topologies are very scalable and flexible. SPB-M can be deployed in small pockets of the network or can be used to connect networks across a wide area. In fact, SPB-M can ‘flatten’ the entire network, as it can scale from the datacenter room itself, across the distribution network, to the user edge using a single protocol.
- SPB-M in the OmniSwitch takes advantage of the Alcatel-Lucent Enterprise intelligent and automated network edge configuration and security features. The features include Edge Virtual Bridging (IEEE 802.1Qbg ‘EVB’), Virtual Network Profiles (vNP) and User Network Profiles (uNP), and Alcatel-Lucent Enterprise OmniVista™ 2500 Virtual Machine Manager. With this suite of tools, configuring and integrating SPB-M edge networks and attached devices is highly automated and flexible.

3. DEPLOYING SPB-M IN THE ENTERPRISE: AN EXAMPLE

Enterprise networks are commonly comprised of several different purpose-built networks linked together by various means. An enterprise may have one or more datacenters; core and/or distribution networks; and many wired and wireless access networks of varying types.

Figure 1. An enterprise network

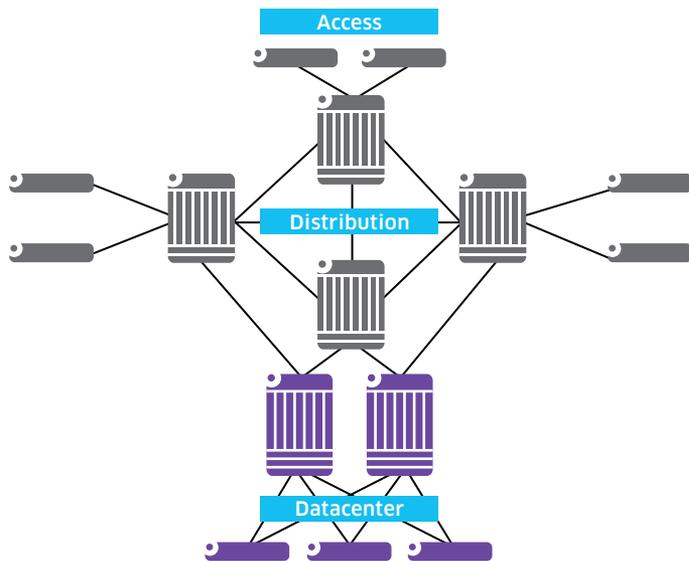


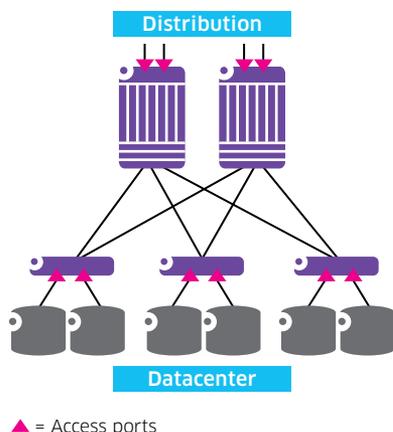
Figure 1 illustrates a typical enterprise network. There are three distinct networks (datacenter, distribution and access) that are configured differently and may be managed separately. The datacenter may be a flat layer-2 network in which virtual networks overlay the physical infrastructure. The distribution network may be configured in many ways, including Virtual Routing and Forwarding (VRF), or bridging based on Ethernet Ring Protection Switching (G8032). Access networks can be routed or bridged networks that incorporate secure edge features such as the Alcatel-Lucent Enterprise Access Guardian™ to authenticate end-points and user networks. Each of these networks are suitable targets for deploying SPB-M. As described earlier, SPB-M can be deployed in a phased approach, which will now be outlined in greater detail.

3.1 Phase One: The Datacenter

One logical application for SPB-M is in the datacenter, where it can create an optimized fabric with powerful virtualization capabilities.

Administrators have several options on the server-facing ports. If the datacenter configuration is relatively static, all the necessary networks for each server can be bound by configuring SAPs on the switches via the management interfaces. If the datacenter uses a dynamic virtual machine management or a cloud orchestration system, SPB-M on the OmniSwitch can be integrated with the Alcatel-Lucent Enterprise OmniVista™ 2500 Virtual Machine Manager (VMM), which communicates with those external systems to coordinate the placement of virtual compute resources into their corresponding network containers. On datacenter interfaces that interconnect to other networks, such as the distribution network, or on networking devices that do not support SPB-M, administrators can bind SAPs as needed to ensure those networks and devices are connected to the correct resources inside the datacenter.

Figure 2. Datacenter network with SPB-M



3.2 Phase Two: The datacenter and distribution

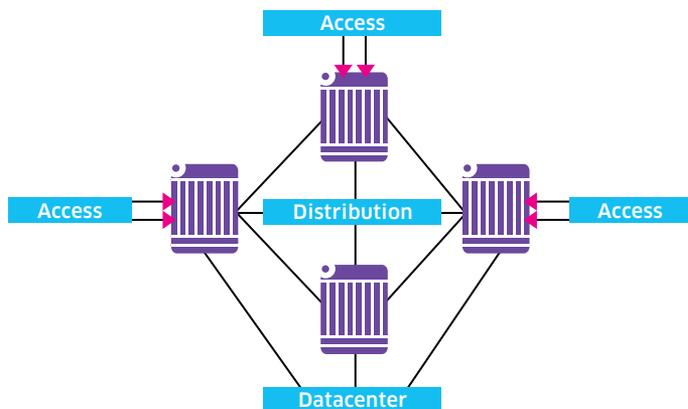
After migrating the datacenter network to SPB-M, it is a relatively straightforward process to also incorporate the distribution network. The switches in the distribution network will be configured as SPB-M bridges using the same BVLANS and IS-IS settings as switches in the datacenter.

The interfaces between the datacenter and distribution networks are now configured as SPB-M backbone interfaces.

The interfaces connecting the access networks now become SPB-M edge ports and the networks on those interfaces will be mapped to services based on their VLAN and SAP configurations.

It is likely there will be VLANs and networks in the distribution network that do not – or should not – have direct access to the datacenter. This is supported by allocating separate services to these networks that are not bound to any interfaces within the datacenter network. SPB-M does not leak traffic between services, so the two domains are securely separated by using different I-SID values.

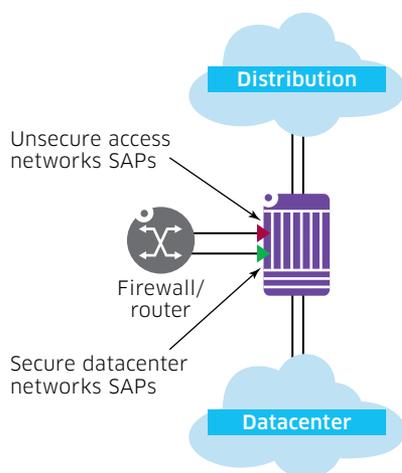
Figure 3. Distribution network with SPB-M connected to datacenter



Another key supported scenario occurs when network devices such as routers or firewalls that do not support SPB-M are required for correct network functionality. Distribution and access networks often need to traverse routers and firewalls to access other networks in the same network segment or inside the datacenter (e.g., to protect the datacenter from the rest of the enterprise via firewalls). The design for this typically involves physically separating the datacenter and non-datacenter networks by placing firewalls between them. While this would still be possible in an SPB-M context, it is unnecessary and creates two separate SPB-M networks.

An administrator can build and maintain a single SPB-M network for both secured and unsecured networks. The administrator configures such a scenario by creating services that are 'unsecure' for the networks that need to pass through the firewall, and 'secure' services for the internal datacenter network. Firewalls can have one or more interfaces for both the unsecured and secured networks. On the bridging interfaces where the firewall connects, the administrator binds the SAPs to ports for secured and unsecured services. For example, I-SID 1100 may be a network for employee PCs that needs to access company resources and file servers connected to I-SID 2100 in the datacenter. Binding I-SID 1100 to the unsecured interface of the firewall and I-SID 2100 on the secured firewall interface ensures all traffic between the two services still passes through the firewall where the appropriate enterprise security policy can be applied to either allow or deny access to the datacenter network (Figure 4).

Figure 4. Example of inline firewall connection



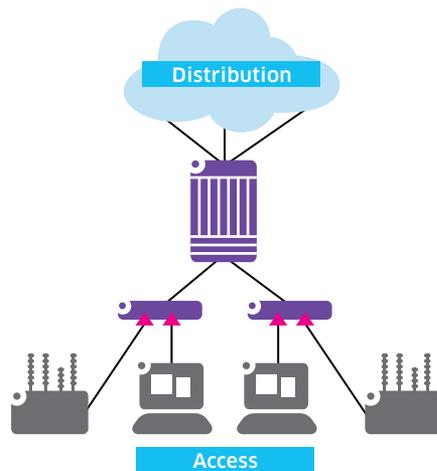
3.3 Phase Three: Bringing in the access networks

With distribution and datacenter networks successfully migrated to SPB-M, the access networks are still operating as normal, enforcing edge security and user classification through the desired mechanisms (e.g., 802.1X or MAC Authentication). With OmniSwitch SPB-M, however, it is still possible to use the same security mechanisms at the edge of the network. As a result, the entire network can be operated under a single end-to-end protocol with no feature gaps to compromise either performance or control.

The user-facing ports of the access network, whether they are LAN or WLAN endpoints, can be programmed to dynamically map devices into their target SPB-M services following OmniSwitch uNP or vNP authentication and processing.

The interfaces between the access and distribution networks will now become SPB-M backbone interfaces.

Figure 5. Access network with SPB-M



4. CONCLUSION

SPB-M implementation in the Alcatel-Lucent Enterprise OmniSwitch offers compelling advantages compared to other technologies and can be deployed in many ways to match an enterprise's current and future needs. SPB-M builds highly optimized, virtualization-ready networks that enable greater flexibility, scalability and performance, while maintaining secured access to distributed resources.