



TECHNOLOGY STRATEGY BRIEF

Enterprise L3 MPLS support by the Enterasys S-Series

Enterprise L3 MPLS support by the Enterasys S-Series

Introduction

Multiprotocol Label Switching (MPLS) is a mechanism typically used in high-performance service provider networks to direct data from one network node to the next, based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. Labels are used to identify virtual links (paths) between distant nodes rather than endpoints. The term “MPLS” is often used to broadly refer to an array of applications and solutions, beyond the simple label encapsulation process, which utilize signaling and control protocols, to offer specific services to end users.

MPLS can be used in many ways, but some of the most common applications are:

- Layer 2 or Layer 3 Virtual Private Networks (VPNs)

VPN services allow a service provider or public network carrier to provide the equivalent of dedicated private network services to multiple customers over a common network infrastructure, while ensuring the segregation and separation of the traffic from the multiple customers. While there are multiple approaches to creating VPN services using MPLS, these multiple options can be classified as either Layer 2 switched services or Layer 3 routed services. A major reason that MPLS is emerging as a consideration for campus enterprise network customers is that customers are looking to take advantage of the benefits and services they see being implemented in service provider networks.

- Traffic Engineering (MPLS-TE)

Traffic Engineering is the most common application for MPLS in service provider networks. Traffic Engineering is the use of MPLS in large, complex, wide area networks to optimize the utilization of expensive long haul links and enhance the control of traffic flows in these networks. MPLS Traffic Engineering is an application whose primary value is to service providers with very large national or global wide area networks.

MPLS operates at an OSI layer that is generally considered to lie between traditional definitions of Layer 2 (data link layer) and Layer 3 (network layer), and thus is often referred to as a “Layer 2.5” protocol.



Figure 1: MPLS header Layer 2.5 with Bottom and Top Label

While MPLS technology is commonplace in service provider networks, it is not as prevalent in enterprise networks. Enterprises often use service providers’ MPLS networks to connect remote locations, but typically the MPLS infrastructure is completely transparent to the enterprise network. The enterprise traffic will be sent in its native format (no MPLS labels) to the service provider’s router. The service provider is responsible for adding and removing labels, and the enterprise network has no knowledge of the service provider’s MPLS infrastructure.

Benefits

Multi tenancy support

- Enables larger campus deployments with separation of customers, business units in VPNs

Reduced complexity

- Allows a VRF deployment to scale across a larger backbone, without the need to provision the core for each and every new VPN

Reduced cost

- Embedded in standard core and data center switch/routers without the need for dedicated devices

A VPN service can be enhanced by an architecture which separates the L3 VPN service layer from an underlying transport layer. In this model, the Provider Edge routers are VPN aware and the core routers provide transport services for the VPN traffic without needing to be aware of VPNs. The transport layer provides connectivity and high availability services through multi-path and redundancy capabilities for the VPN service layer. The Enterasys S-Series will deliver Layer 3 VPN functionality in multiple stages. This brief will describe both current and future technologies on the S-Series to provide secure isolated networks of all sizes.

VRF Overview

While MPLS can be used to implement a campus wide virtual private network, a simpler approach is to use Virtual Routing and Forwarding (VRF). VRF is the capability to have separate routing tables and routing processes in a single physical router to create separate routing domains and it is typically used within the PE (provider edge) routers of a service provider MPLS network. Within these routing domains, one can run any of the familiar IP routing applications the customer has already implemented such as OSPF, BGP, or RIP but without the complexity of having to introduce an entirely new set of protocols and technologies into the network. A data center can segregate server resources by placing them into separate routing domains, ensuring there is a complete separation of connectivity between various groups, but still leveraging a single set of physical routers and using a common management interface. Medical or industrial applications can be segregated from the rest of the network and provided with dedicated network resources, ensuring security and control of access to sensitive applications.

VRF provides a simple solution for campus LAN applications that are limited in network diameter. No new complex protocols or architectures need to be introduced into the LAN (such as MPLS, RSVP, LDP, iBGP, etc.) so a customer's current engineering and operations staff needs minimal new training. No new management or diagnostic tools are needed – all the familiar protocols, such as OSPF and RIP, work the same within their respective routing domains and existing network protocol analyzers and appliances don't need to be upgraded to support the MPLS protocols.

VRFs can be deployed end-to-end through the network or in conjunction with GRE tunnels or MPLS labels. VRF that does not use MPLS is often referred to as VRF-lite, or Multi-VRF Customer Edge, and are thought of as "lightweight" versions of MPLS as there are no labels or label distribution protocol.

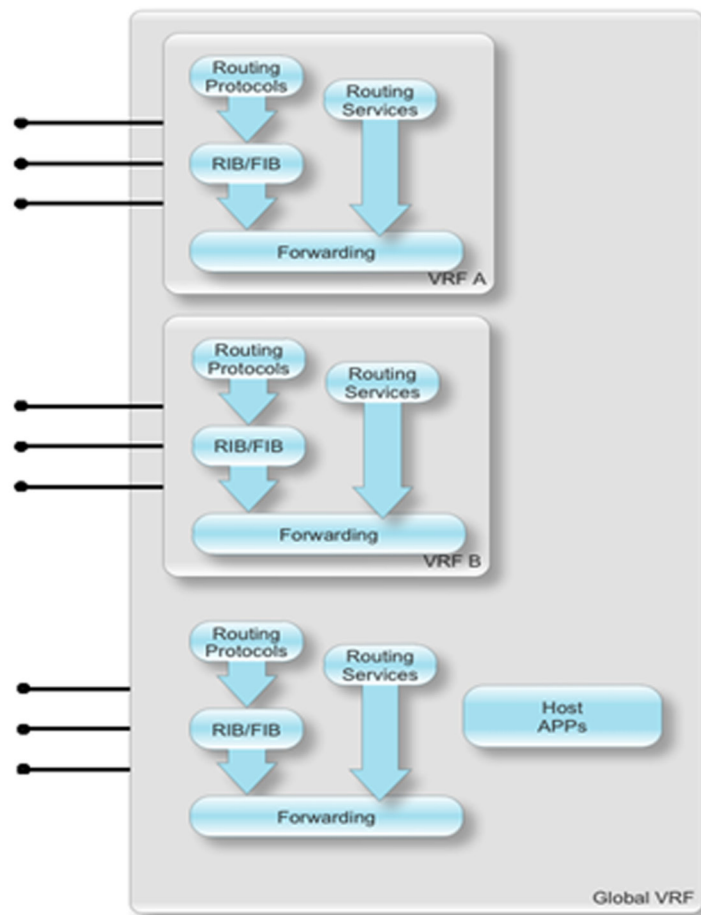


Figure 2: VRF separation inside the router

S-Series VRF Support

The Enterasys S-Series supports end-to-end VRF, VRF over IPv4 and IPv6 GRE tunnels.

In the end-to-end model, each routed interface belongs to one VRF and must be manually configured on all routers that will participate in the VRF which can become tedious and cumbersome to manage. A typical rule of thumb is that end-to-end VRF is appropriate for networks with less than four router hops, from edge-to-edge.


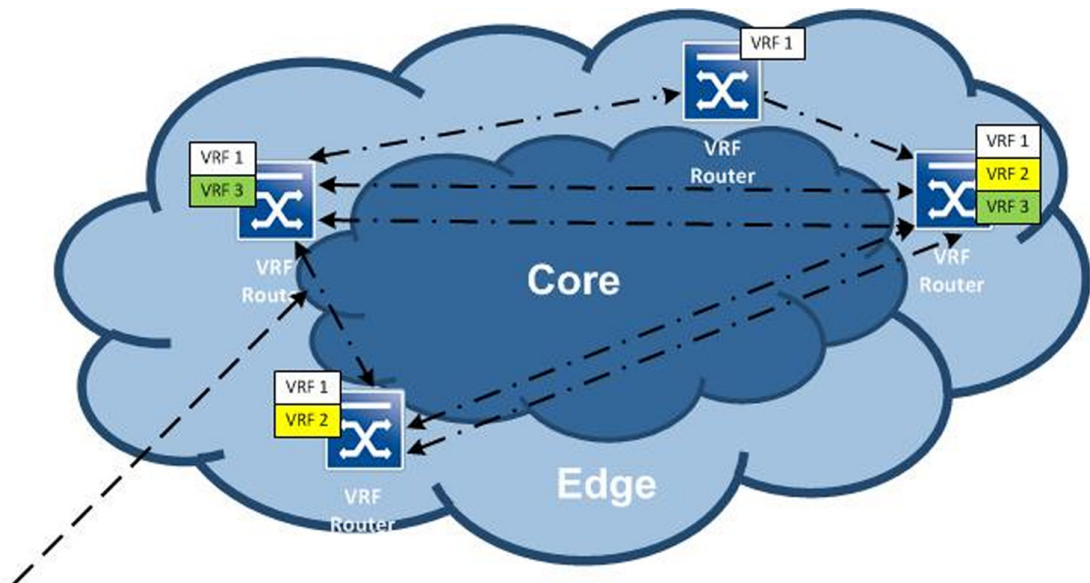


Figure 3: VRF core router deployment

To simplify and enhance the scaling of a L3 VPN service deployment, it is desirable to minimize the configuration of a VRF instance on routers which do not need to support that VRF for any edge services. To reduce configuration complexity, VRFs can be supported over GRE tunnels between PE routers, simplifying the configuration of the core P (provider) routers which do not need to be VRF aware. A GRE tunnel is provisioned for each VRF on a PE router to all the other VRF instances on PE routers in the network. This eliminates the need to provision VRF instances on each core router and so leads to a simplified IP core routing infrastructure.



GRE mesh per VRF

Figure 4: VRF over GRE

Both end-to-end VRF and VRF over GRE tunnels can provide secure, dedicated routing resources for critical applications and provide a simple solution for campus LAN applications that are limited in network diameter, without introducing a new network technology such as MPLS.

MPLS

An MPLS architecture based on a Provider Provisioned VPN (PPVPN) model uses the notion of specific nodes in each layer of the network:

- **P** – *Provider Core Node* – As the name implies, this device resides in the core of the network, and in the context of MPLS is making forwarding decisions based on MPLS labels and has no knowledge of the customer's routing infrastructure.
- **PE** – *Provider Edge Node* – This device interfaces between the Provider Core Node and the customer edge. The PE is responsible for inserting MPLS labels into traffic entering the provider network, and removing labels from traffic before entering the CE.
- **CE** – *Customer Edge* – This device resides at the edge of the network and has no knowledge of the MPLS infrastructure. It does not send or receive traffic with MPLS labels.

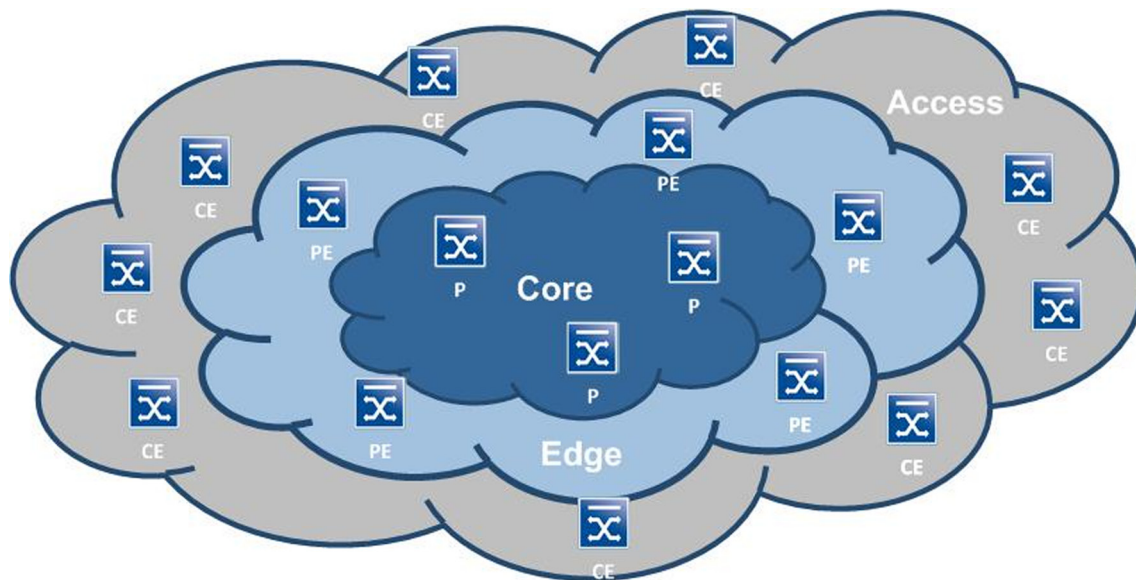


Figure 5: MPLS network overview

Different VPN Models can be deployed within such architecture:

- Virtual Private Router Networks (VPRN, i.e. RFC4364 or RFC4023) which allows for a virtual Layer 3 routed network.
- Virtual Leased Line Services (VLL, RFC 2764) which do provide an emulation of a point-to-point link.
- Virtual Private LAN Segment (VPLS)/Transparent LAN Service which do emulate a Layer 2 bridged LAN.

MPLS VPN Overview

RFC4364 and RFC4023 describe a method of supporting IP VPNs using VRF, MP-BGP and MPLS. In this context, a VRF equates to a VPN. Routes are distributed using BGP, and an MPLS label is used to identify VPNs. BGP is used to isolate traffic and then apply a corresponding MPLS label to identify the VRF. This MPLS packet is further encapsulated with either another MPLS label or with an IP or Generic Routing Encapsulation (GRE) tunnel header [MPLS-in-IP-GRE] so that it gets tunneled across the backbone to the proper edge router. Thus, the backbone core routers do not need to know the VPN routes of each VRF.

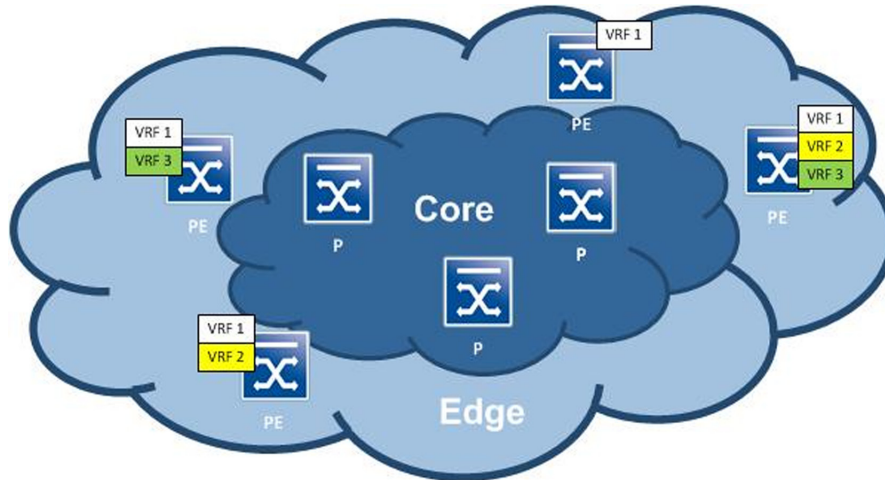


Figure 6: P routers with no VRF knowledge

S-Series MPLS VPN Support

The S-Series platform will support MPLS in a manner that is focused on enterprise customers that are deploying MPLS in a multi-phased approach.

Phase 1

MPLS BGP L3 VPNs / GRE (RFC4023) – As the number of L3 VPN domains and PE nodes increase, the need to provision a separate GRE tunnel for each VRF can add complexity to the configuration. To further enhance scaling over the IP core infrastructure, MPLS can be introduced at the L3 VPN services layer. As defined in the Proposed Standard RFC 4023, MPLS can be used to provide label multiplexing for multiple VRF instances over a SINGLE provisioned GRE tunnel, independent on the number of VRFs between TWO PE router peers. iBGP can be utilized to simplify the exchange of VPN routing information and reduce the number of IGP routing instances, such as OSPF or RIP required to traverse the core over the GRE tunnels. A single instance of iBGP peering simplifies configuration and management of the PE routers compared to a VRF over GRE implementation which would require a GRE tunnel for EVERY common VRF between TWO PE routers.

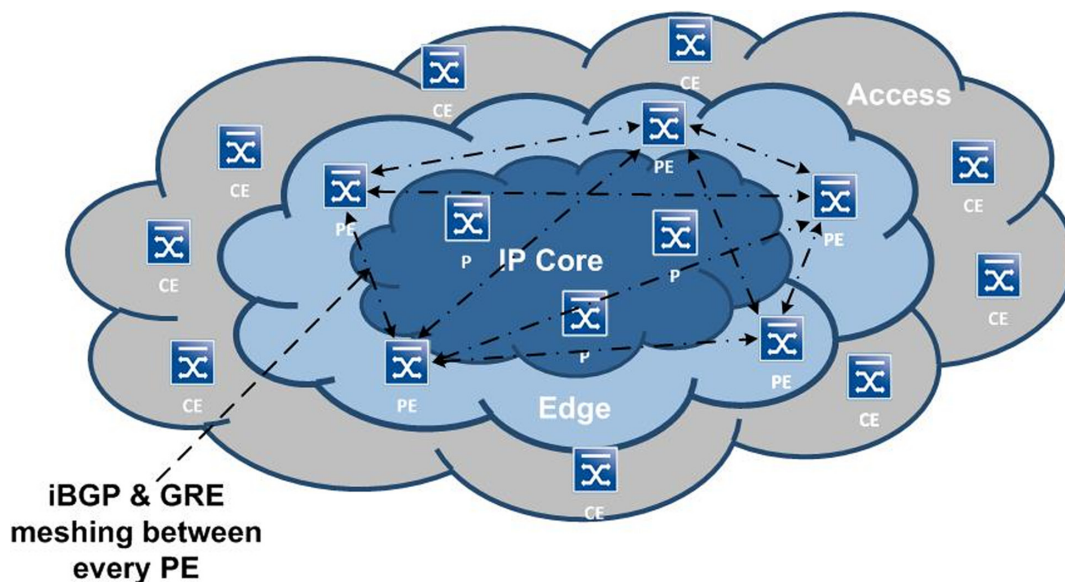


Figure 7: MPLS over GRE

Phase 2

MPLS BGP L3 VPNs / Dynamically Provisioned Transport – For very large scale VPN deployments, provisioning GRE tunnels between large numbers of PE routers may become complex and difficult to manage. The next step for enhanced scaling and service levels will be to take advantage of a dynamically provisioned infrastructure which does not require manual provisioning of transport for the L3 VPN services. There are some alternative transport technologies which can be used for this purpose.

- *MPLS Core Transport* – RFC 4364 defines the architecture for BGP/MPLS IP VPNs which utilizes an MPLS enabled core to dynamically provision LSP tunnels between PE routers. As new PE routers are provisioned, or new VPN instances added, the dynamic signaling of the MPLS infrastructure provisions the mesh of LSP tunnels over which the MPLS IP VPNs services are transported. RFC 4364 IP VPN services are commonly deployed in large-scale service provider networks, which already have an MPLS enabled core infrastructure deployed.

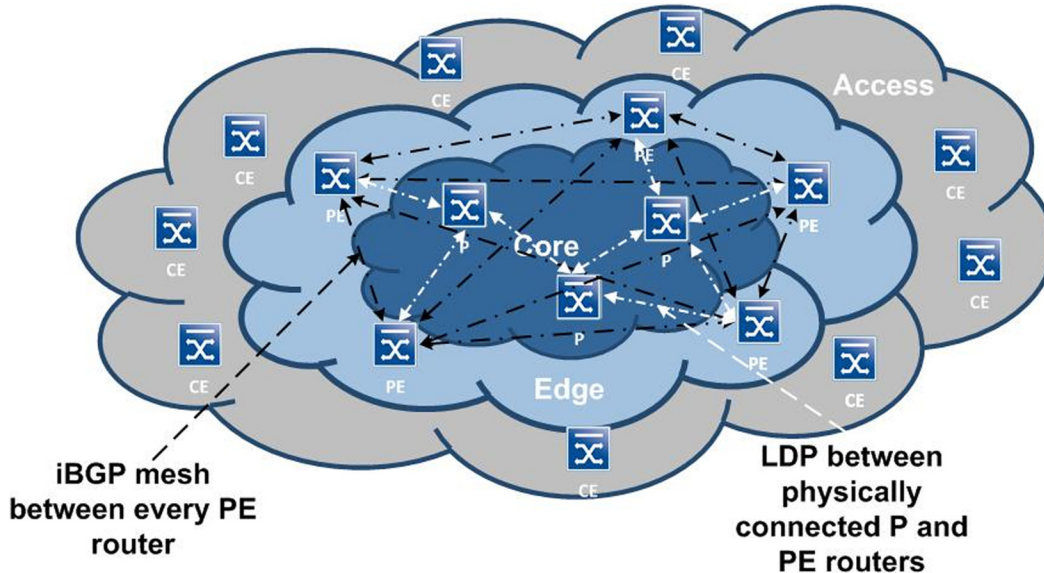


Figure 8: MPLS L3 BGP VPN using RFC 4364

- *Shortest Path Bridging Transport* – The emerging IEEE 802.1aq Shortest Path Bridging (SPB) standard defines a dynamically provisioned transport service, scalable for very large topologies. SPB, as an extension of Ethernet-based networking, is naturally targeted for enterprise deployments in the rapidly growing data center and Core network applications. Since SPB provides a highly scalable, dynamically provisioned Ethernet-based infrastructure, with multi-path and high availability capabilities, SPB is an ideal candidate as an MPLS BGP L3 VPN transport service for the Enterprise. In SPB, the notion of BEB Backbone Edge Bridges and BCB Backbone Core Bridges is used. The BEB and the PE become a single switch/router in this model. Since SPB natively provides multi-point transparent LAN services, without the added complexity of VPLS, the leveraging of SPB for enabling L3 VPN services provides an optimal architecture for the Ethernet-centric network environment of the enterprise services on top.

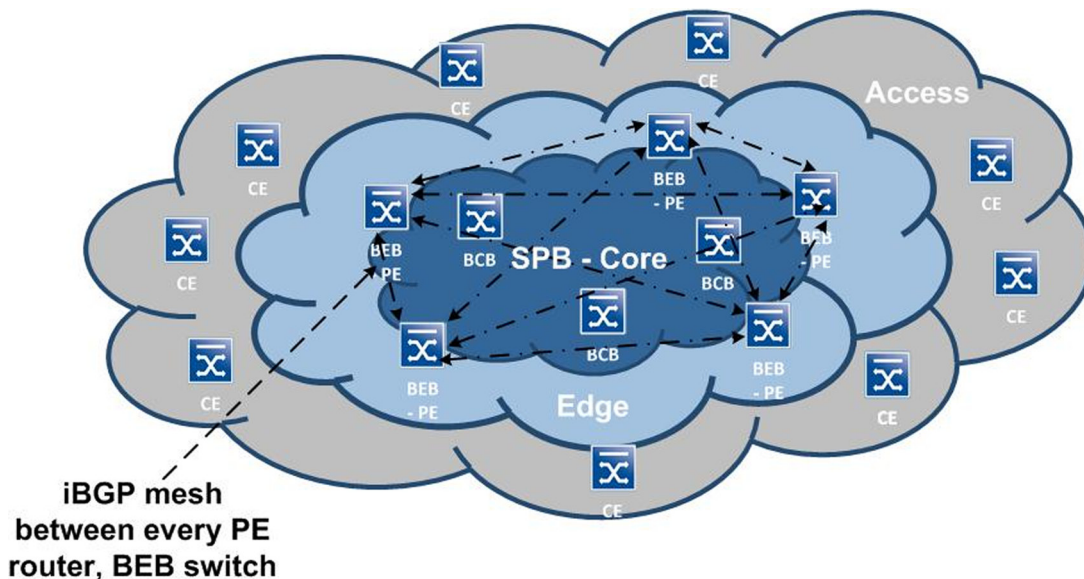


Figure 9: MPLS over SPB

Summary

MPLS provides an option to simplify VRF (Lite) deployments in larger enterprise infrastructures. The Enterasys S-Series with its CoreFlow2 ASIC architecture will support this technology, along with SPB Shortest Path Bridging in a multi-phased approach via software upgrades (license fees may apply restricted to specific S-Series modules). This solidifies the positioning of the S Series as a premier switch/router solution for the data center, core, aggregation and high-value edge in the OneFabric architecture.

Contact Us

For more information, call Enterasys Networks toll free at **1-877-801-7082**, or +1-978-684-1000 and visit us on the Web at enterasys.com



© 2011 Enterasys Networks, Inc. All rights reserved. Enterasys Networks reserves the right to change specifications without notice. Please contact your representative to confirm current specifications. Please visit <http://www.enterasys.com/company/trademarks.aspx> for trademark information.

