Improving End-to-End Network Scalability and Resilience with MPLS at the Edge

A White Paper from Telco Systems

Nir Daube

March, 2012

Network operators and service providers are facing enormous challenges. Networks must be engineered, built, and maintained for bandwidth optimization, resiliency, high quality of experience for end-users, and ease of management, end-to-end.

This paper will address the advantages of extending MPLS to the edge network to achieve lower operating costs and create a more scalable and resilient service offering.

Improving End-to-End Network Scalability and Resilience with MPLS

While Ethernet can deliver a cost effective network, it doesn't always scale well and sometimes fails to deliver on the quality of experience that users have come to expect. How, then, can providers ensure that they can deliver the services that their customers are demanding without costly investment in their network infrastructure?

MPLS has transformed the service providers' core network over the past decade and is quickly moving into the Ethernet access network to extend MPLS benefits like scalability, traffic engineering, quality of service, and resiliency closer to the customer. This is creating a more intelligent access network to support end-to-end services.

Moving MPLS closer to the edge basically integrates access and aggregation networks together with the core onto a single MPLS-managed network to create significant operational advantages.

Advantages of extending MPLS from the core to the edge:

- Offloads the core since many routing decisions can be made in the access
- Enables fast service creation/delivery that supports legacy and future services
- Optimizes bandwidth utilization throughout the network
- Scale the network beyond VLAN limitation to practically unlimited
- Ensures service delivery during Moves/Adds/Changes in the network
- Eases management and maintenance by using a single technology end-to-end
- Increases number of Classes of Service using Hierarchical QoS
- Provides five 9's availability
- Supports OAM at various layers to prevent unnecessary truck rolls

When technologies are different in the core and access networks, there are distinct hurdles that must be overcome in order to deliver on defined SLAs and quality of service.

Service Creation

Using multiple network technologies to deliver services requires delicate planning and administration to ensure seamless delivery. Typical networks use Ethernet at the access and aggregation, MPLS in the core. While each technology has its advantages, having to provision and manage both causes delays and complexity in service turn-up as well as delivery. Many times, engineering and operations departments are separate for the various parts of the network which adds further delay.

Scalability

Demand for network services continues to rise. As the number of users increase, the number of applications also increase, and the amount of bandwidth consumed by those users/applications increases. The network must scale to hundreds of thousands of nodes to support these growing requirements, and scalability is very different in Ethernet (using VLAN and q-in-q) vs. MPLS (using labels and label stacking).

Ethernet based networks are limited to 4096 VLANs. In many cases multiple VLANs are assigned per customer by service, so the use of VLANs becomes a limiting factor in network design. When multiple customers are running the same VLAN, tagging and prioritizing the service at any point in the network becomes impossible and may create conflicts. It is at this point that we need the ability to tag *the service and the customer* from the demarcation to the core. MPLS with million labels (which can be reused) makes the networking practically scalable with no growth limitation.

Optimize core resources.

When intelligence resides only at the core, all packets need to go first to the core and then back to the customer, even if the connection is within the same access network. Requiring all traffic to pass through the core is analogous to traveling from Boston to New York by way of San Francisco!

If every packet needs to reach the core in order to reach a service node, network resources are being wasted and used inefficiently. The core router will need more ports, more switching capacity, and more processing power to support traffic that could have been handled locally.

By implementing intelligence in the access network, core network resources, where per-port costs are very high, can be optimized to support significantly more services and create a more deterministic network. If traffic management and routing decisions can be made within the access network, some traffic may not need to traverse the core network at all.

Recovery and restoration mechanisms

If a link breaks, how quickly can traffic flow be restored? How does a failure affect the various services? Do all services require the same restoration mechanisms? Using Ethernet as the network transport requires the use of Ethernet restoration, which means some variation of Spanning Tree. Standards based STP restoration can take several seconds as opposed to MPLS Fast ReRoute restoration which is on the order of 50ms. When running both Ethernet and MPLS, certain failures might trigger contention among restoration schemes.

QoS End-to-end

HQoS (Hierarchical Quality of Service) allows the carrier to make sure each service gets the network resources it needs while coexisting on the same network infrastructure with other types of services. HQoS also enables unused bandwidth resources to be shared between different services. This optimizes network resources while keeping the same level of service without tying up bandwidth to unused services.

HQoS allows a wide range of traffic-manipulation possibilities at line rate *in hardware*. The HQoS implementation enhances the 'flat' QoS model by introducing multi-level hierarchy for both flow classification and traffic management by providing per customer/service hierarchical queuing, scheduling and shaping for both service ingress and service egress.

This model allows an SLA to be defined on both customer and service levels, where multiple customers can be connected to each port, and each customer is subscribed to multiple services. Per customer per service QoS enforcement using HQoS means that each service can be identified and assured throughout the network.

Different technologies have different ways of implementing QoS. Ethernet supports prioritization and priority queuing. MPLS supports these as well as bandwidth reservation and traffic engineering.

Why Extend MPLS to the Edge?

A seamless network improves service creation time – Seamless control plane

The best way to improve service creation/delivery time is to have a network that operates independently of the services, yet can support any service deployment scenario. The network should not have transport boundaries which limit access to services.

A "seamless" MPLS architecture is one that inherently has no boundaries and hence decouples the service layer from the transport layer. This decoupling increases the flexibility to define and introduce new services by allowing service nodes to be placed at optimal locations in the network rather than at the "boundary nodes". Although both the service and transport layers use the same MPLS packet formats, the difference is in the use of MPLS, especially on the control plane. Using the MPLS control plane end-to-end enables a management system to select the endpoints of the service, then triggers signaling to set up the service across the network between the endpoints.

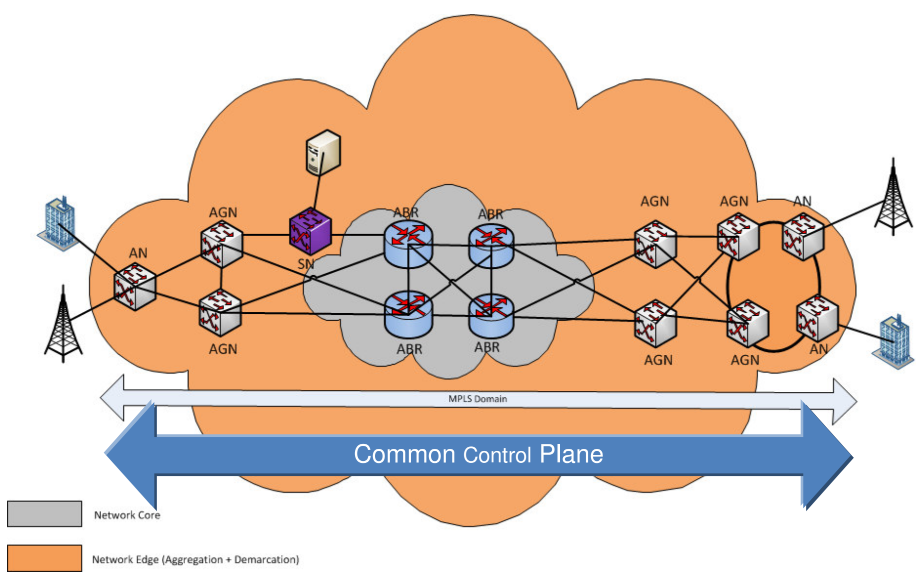

Traffic-engineering vs. Over-engineering

Delivering a specific amount of bandwidth for an aggregate of traffic while also delivering a specific amount of bandwidth for a single traffic stream can be difficult. Most providers over-engineer their networks so that the customer does not notice any performance issues. Metering and policing may be performed, but this does not guarantee resources in the network. As traffic requirements continue to grow, and specialized services are required for the various types of traffic, while operators are struggling to keep costs low, over-engineering becomes an expensive and cumbersome option to manage. This option requires extra links, ports, switching capacity, power, space, etc. The approach moving forward should be to engineer traffic as to maximize the use of the facilities already in place.

The only technology successful at doing this is MPLS. Carrier Ethernet networks was not built with the proper mechanisms and has very limited traffic engineering capabilities. In L2 carrier Ethernet all packets take the best path, which may be the most congested while other paths will remain empty and will only be used in case of failure.

The traffic engineering capabilities inherent in MPLS, combined with HQoS (Hierarchal QoS), allow for guaranteed bandwidth for an aggregate of traffic as well as individual flows. Paths can also be engineered such that delay-sensitive traffic will take the path with the least delay. As the customers' traffic mix changes, these traffic-engineered paths can easily be modified to support those changes.

Should a failure occur in the network, MPLS traffic engineering enables the provider to decide how to share the backup resources. With the over-engineering approach in L2 carrier Ethernet (depending upon the failure and its location), there may be enough bandwidth for all traffic to be supported, OR the failure could create a congestion point forcing all traffic to compete for limited backup resources. This competition will cause the performance for *all* traffic to suffer. It's an all or nothing approach. In contrast, MPLS traffic-engineering along with HQoS allows for high-priority, high-revenue services to maintain service quality while only the low-priority traffic suffers performance degradation.

Designing the network for resiliency

A resilient network is one that offers 100% availability. Redundancy must be designed in the network from the customer site, across the access and aggregation networks, and across the core. Multi-homing the customer site to different access nodes in the provider network is key for end-to-end resiliency.

There are several options for doing this depending on the access facility. Primary/backup links, multi-chassis LAG, and other mechanisms restrict the use of the backup bandwidth for failure scenarios only. Having multi-homed links, then using MPLS-TE to allocate bandwidth to customers uses the bandwidth of both links simultaneously. If a failure occurs, MPLS FRR (Fast ReRoute) handles the switchover in less than 50ms. BFD (Bidirectional Forwarding Detection) in combination with various protocols will improve time to detect failures thus invoking restoration measures even faster.

Support of TDM and other legacy services

Because MPLS is essentially a tunneling protocol, it supports the transport of ANY service available today - TDM, Ethernet, Frame Relay, ATM, PPP, HDLC, IP. Pseudo-wires are used to transport any protocol over MPLS. These services and protocols are encapsulated with MPLS labels, and switched to the destination which may be another customer site or a service within the provider network.

Because of the very strict timing requirements of TDM, traffic engineering is a must. Trying to support this without proper traffic engineering will subject the TDM to the bursts and delays of all other traffic carried on the links. MPLS with traffic engineering can guarantee dedicated bandwidth for the TDM service.

Maximizing Scalability

MPLS is highly scalable. The 20-bit label allows for over one million LSPs per node. Stacking n labels allows for $1M^n$ tunnels. By using VPLS and/or H-VPLS, this network can support thousands of customers, and each customer can have a different logical topology.

In contrast, Ethernet's 12-bit VLAN tags allow for 4K VLANs per switch. VLAN stacking (q-in-q) allows for 4K customer VLANs to be carried in 4K provider VLANs. Since each customer is likely to use multiple VLAN IDs, the number of customers that can be supported is quite limited.

L2VPN vs. L3VPN?

MPLS does not distinguish between a L2 and L3 service and can support the transport of multiple layer 2 protocols, including Ethernet, ATM and Frame Relay. It offers the scalability of Layer 3 routing with the performance of Layer 2 switching, enabling quality of service guarantees and overcoming the packet loss and latency issues in Frame Relay and Internet VPNs.

However, deploying L3 VPNs is an expensive and complex solution. Each supporting device needs to maintain routing tables, creating a full mesh routing cloud.

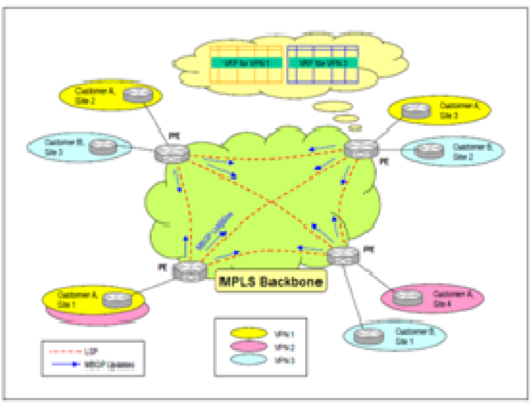

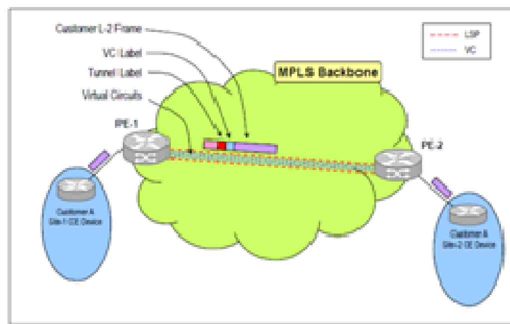

MPLS L3 VPN - Relies on taking customer IP datagrams from a given site, looking up the destination IP address of the datagram in a forwarding table, then sending that datagram to its destination across the provider's network using an LSP

VPLS / H-VPLS – A layer-2 approach that provides complete separation between the provider's network and the customer's network, i.e., there is no route exchange between the PE devices and the CE devices. Hence, the approach follows the overlay model of VPNs.

Using L2VPNs supported by VPLS and HVPLS from the demarcation to the core reduces the amount of resources required in the core routing equipment and places them in the access where they are simpler and more cost effective to purchase and operate.

Virtual private LAN service (VPLS)

VPLS is a way to provide an Ethernet-based multipoint to multipoint connection over an MPLS network. It allows geographically dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudo-wires technology like Ethernet over MPLS.

VPLS is basically just a very large virtual switch. Each entry point is a “port” in this switch. By connecting Site A to Site B, it will seem as though they are connected to the same switch. A VPLS implementation *requires a full mesh between all the devices in the network based on MAC addresses of the user at the entry point*. This means every device needs a high level of resources and a strong CPU to support MAC tables, and routing protocols.

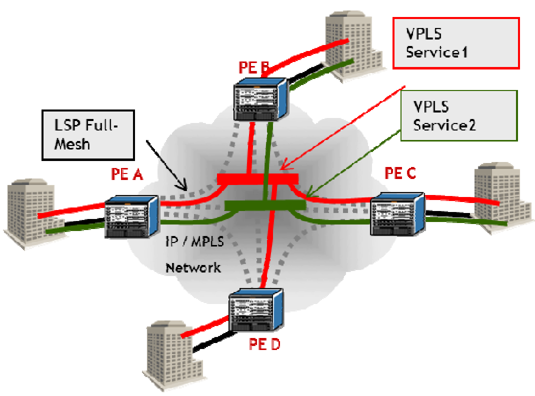

In a VPLS mesh network, each Provider Edge router must be identified and given the addresses of other PEs participating in the same VPLS. LDP sessions are then established between these PEs creating a mesh of pseudo-wires between devices. This mesh network between connected PEs enables traffic to be automatically re-routed along available backup paths, achieving much faster failover than can be supported with Spanning Tree. This creates a more reliable solution for connecting Ethernet networks over a WAN.

So, why do we need H-VPLS

H-VPLS, or hierarchical VPLS, simplifies VPLS by defining subdividing a VPLS into smaller domains utilizing spoke devices closer to the customer. Each spoke device sees only the upper PE device. A spoke device, known as the multi-tenant unit (MTU) switch adds an MPLS and VPLS label (similar to an LER) and aggregates multiple customers to a single provider edge (PE) device. This reduces the amount of work required by the PE device making it simpler and more cost effective and reduced the resources required in the core network.

Ease of Management

Since MPLS is already used in the core, it makes sense to use MPLS in the access and aggregation networks. A single protocol end-to-end allows for service nodes to be placed in the most cost-effective locations, and allows modifications to the service to be decoupled from the network.

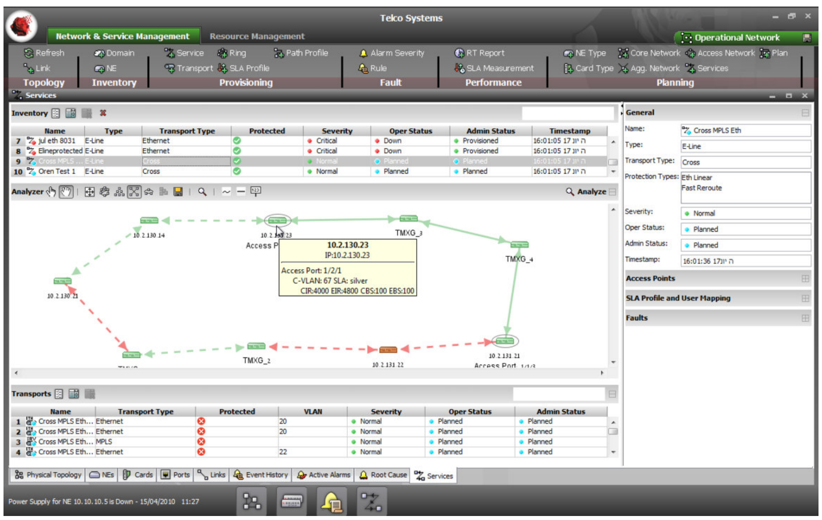

Provisioning can be done quicker using the common control plane end-to-end, and common protection mechanisms make fault tolerance and resilience easier to plan and implement.

Designing the traffic-engineered network takes some forethought, but many of the mechanisms are automated and there are many tools to assist with the planning. Management systems are available to assist with the design and they can perform "what-if" scenarios to predict behavior in the event of certain failures.

Once it is up and running, MPLS allows for testing to be done for each tunnel, and for each customer traffic flow carried inside the tunnel so that operation and connectivity can be verified. LSP ping/trace and VCCV (Virtual Circuit Connectivity Verification) will test connectivity for the MPLS label switched paths as well as the pseudo-wires.

Various OAM tools can also be implemented to give <u>*both*</u> service topology and network topology views. This will aid in management as operations will be able to know where in the network the customer traffic is flowing. In a dynamic network environment, traffic-engineered paths will take the best available path at the time of setup. As links and/or nodes go up or down, path calculations will change and re-route into pre-defined routes. Using a make-before-break approach maintain the customer services while changes in the network made, or happen

Summary

Based on typical customer requirements, MPLS with traffic engineering and HQoS is the technology of choice in the access and aggregation networks, just as it is in the core. A network architecture that is decoupled from the service architecture and incorporates intelligent switching closer to the customer premises will optimize network resources and improve performance of customer traffic. It also enables true customizable services because quality of service parameters can be incorporated end-to-end. The Fast ReRoute (FRR) feature allows for sub 50ms service restoration which maintains the highest availability for the highest priority customers. VPLS and H-VPLS enable flexibility and customization. MPLS is manageable, scalable, and supports any legacy service. By factoring in the cost savings for bandwidth efficiency and high resiliency; and by allowing for additional revenue from customized services and improved service creation time, the business case is very strong for MPLS.

Telco Systems' Ethernet/MPLS Edge Solution

An MPLS and Ethernet solution from Telco Systems will enable you to engineer an easy to manage intelligent edge network that can scale to support new services and more customers, speed the time to service, and improve margins.

Based on a common, sophisticated network operating system that adapts easily for evolving standards, these solutions allow real network consolidation of multiple service types on one unified infrastructure which supports:

- VPLS or 802.1ad transport technologies as well as pseudowire and Circuit Emulation
- Full Ethernet OAM capabilities (IEEE 802.3ah, IEEE 802.1ag CFM, ITU-T Y.1731 and MEF Service OAM)
- Multiple redundancy technologies like MPLS reroute, ITU-T G.8031, LACP, and resilient link
- Intuitive, easy to use service management system that turns MPLS into a practical solution by
    - Automating all service provisioning
    - Providing advanced service status monitoring, including root cause analysis and service analyzer tool
    - Monitoring performance and resource usage, including notifications on threshold crossing
    - Analyzing and forecasting bandwidth capacity shortfalls
    - Planning for future network requirements

Telco Systems also offers a range of professional services to get you started in the world of traffic-engineering, which ultimately will help reduce costs and maximize current resources.
www.telco.com