No. 106
Technology
White Paper

MPLS for Metropolitan Area Networks

Abstract

The demand for bandwidth in Metropolitan Area Networks (MANs) and the increased availability of fiber has created a new class of Metro Service Provider (MSP). MSPs are deploying high-bandwidth networks that enable Application Service Providers (ASPs) and content hosting providers to distribute services within metropolitan areas, thus enabling service delivery closer to the metro subscriber. In addition, MSPs are offering customers native data services using Ethernet instead of traditional leased lines. Customers are attracted to these services by the flexibility offered, the low cost of bandwidth, and the fact that they can get these services up and running quickly. At the same time, customers expect the same level of guaranteed bandwidth and Quality of Service (QoS) they enjoy with TDM circuits. But as customers grow in number and as service providers expand across metropolitan areas, scaling these services becomes an issue.

Riverstone Network's hardware-based metro switch routers have enabled MSPs to offer many of the above-mentioned services using features such as VLANs, wire-speed policy-based routing/bridging, and rate limiting. The support for Metro-optimized Multi Protocol Label Switching (MPLS) on Riverstone's switch router platforms allows service providers to scale their existing offerings within a given metro area, extend their services across other metros, and offer new services — while providing reliable and responsive networks.

This paper looks at the challenges facing the new class of MSPs as they expand their networks, and examines how Riverstone's MPLS implementation elevates the MSP's service delivery capabilities to the next level.

River
STONE
NETWORKS™

Challenges Facing Metro Service Providers

Metro service providers offer point-to-point and point-to-multipoint connections using Virtual Leased Line (VLL) and Transparent LAN Services (TLS) with Riverstone's hardware-based switch routers. These services are based upon either 802.1Q VLANs or IP VPNs.

With the VLAN model, MSPs typically configure one VLAN per customer. They use GARP/GVRP protocols for automated provisioning and extensions to the Spanning Tree Protocol, providing faster convergence (on the order of one second) for improved resiliency. The VLAN model works well to begin with. But as the number of customers increases, the VLAN model presents several scalability problems. In the IP VPN model, L2TP tunnels can be used to carry customer traffic transparently across the MSP cloud. The number of IP tunnels that need to be handled can be daunting, as these tunnels must be provisioned manually and a pair of IP addresses must be assigned to each tunnel.

Riverstone Network's metro-optimized MPLS offers a standards-based way to scale these VLAN and IP-based services. Riverstone's MPLS allows 802.1q VLANs or IP subnets to be mapped to MPLS tunnels or Label Switched Paths (LSPs), while hierarchical MPLS labels enable the bundling of customers' LSPs into a small number of core LSPs in the MSP backbone.

Metro-optimized MPLS from Riverstone Networks

Riverstone switch routers support all standard MPLS features to deliver QoS and traffic-engineering capabilities to metro networks. More important, Riverstone's metro-optimized MPLS offers extensions that integrate existing features with MPLS — enabling MSPs to scale their existing services and offer new ones in metro areas.

Let's look at the key services enabled by metro-optimized MPLS on Riverstone switch routers.

Virtual Leased Line and Transparent LAN Services

As users and networks grow in number, MSPs offering 802.1q VLAN-based VLL and transparent LAN services face several scalability issues. First, the total number of VLANs in the entire network is limited to 4,096. This limits the number of customers in the entire metro to fewer than 4,096. Second, the total number of MAC addresses to be handled in the core of the network may become prohibitive unless rules can be defined to limit the maximum number of MAC addresses available. Third, multiple VLANs per customer are not easily managed unless one backbone VLAN is mapped to each customer VLAN, and VLAN identifiers do not collide between customers.

Riverstone's MPLS-based VLL services address these scalability issues, enabling the MSP to offer a logical pipe formed by two MPLS LSPs going in opposite directions. These LSPs with specific QoS characteristics can be statically pre-configured or dynamically established using MPLS signaling protocols. The exact route an LSP follows can be specified in order to meet specific traffic requirements.

Service providers can specify policies defining which path the packets from a specific customer should follow. For example, if a dedicated physical port P1 is assigned to a customer, the service provider can define a policy that directs all traffic from port P1 to a defined LSP L1 on port P2. Riverstone MPLS supports per-LSP rate limiting, which ensures that specified contracts are not violated. Per-LSP statistics allow MSPs to monitor service utilization and recommend bigger pipes when traffic exceeds thresholds.

This VLL service model based on MPLS LSPs scales well for the MSP. End-customer network information, such as MAC addresses and VLAN-IDs, are not exposed to the core network at all since they look only at the MPLS label. In addition, LSPs from the edge can be bundled into a small number of LSP tunnels in the core using hierarchical tunnels based on the MPLS label-stacking feature.

With Riverstone MPLS, metro service providers can extend the VLL model to offer transparent LAN services when more than two sites must be connected. Riverstone MPLS supports the extensions required for "emulating" a LAN with broadcasting and address-learning capabilities.

The following illustration shows how Riverstone switch routers can be deployed at multiple POPs to offer TLS.

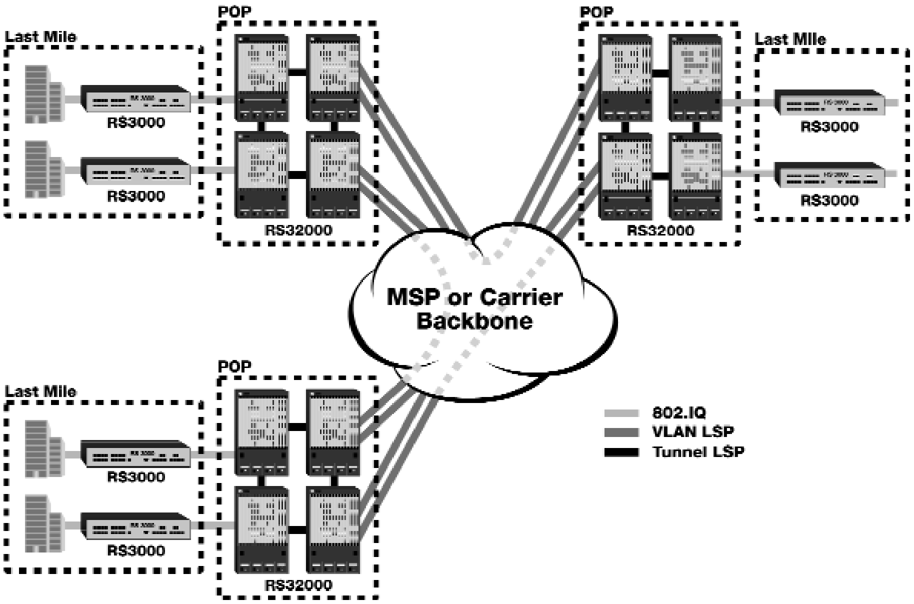

Figure 1: Riverstone switch routers can be deployed at multiple POPs to offer TLS.

In figure 1, customers' VLANs are mapped to specific "VLAN LSPs." VLAN LSPs are tunneled in the core within "Tunnel LSPs" that carry traffic between POPs. The core LSPs that must be handled are limited to a small number. These tunnel LSPs are typically signaled via RSVP-TE, since traffic engineering is often required in the core. Intra-POP LSPs that carry customer VLAN traffic do not require traffic engineering because bandwidth is generally not a concern. Thus, these LSPs are often signaled via LDP. Once a VLAN LSP has been assigned to a customer, no additional provisioning is required. In fact, this single VLAN LSP can carry all the traffic from the customer regardless of the VLAN topology configured at the customer's site.

MPLS IP VPNs

IP VPN services should grow from $200 million in 1998 to exceed $13 billion in value by 2004, according to Frost & Sullivan. The essential requirements for VPNs are security, scalability and service levels. Service providers offer VPNs using connection-oriented ATM and Frame Relay protocols or secure tunnels using IP sec. The primary issue concerning current tunnel-based technologies is that they do not scale well.

Riverstone MPLS delivers elegant and scalable VPN services with BGP extensions based on the latest draft of RFC 2547-bis. With MPLS VPNs, Metro Service Providers allocate a VPN-ID for each customer. Combinations of VPN-IDs and IP addresses are used in the forwarding tables, to make customer IP addresses unique. In the MPLS VPN, VPN information is distributed by BGP only to members of the same VPN, providing traffic separation. Traffic is forwarded using LSPs, which deliver the levels of security offered by ATM and Frame Relay networks. The forwarding table contains labels corresponding to the VPN-IP addresses.

The hierarchical-label support provided by Riverstone MPLS lets the solution scale with a small number of LSPs in the core. Metro Service Providers can offer tiered IP VPN services with MPLS based QoS and reliability commitments as described in the following sections. These simple but powerful VPN services meet users' generic VPN requirements and provide additional powerful service-delivery mechanisms for the metro service provider.

Reliability

With Riverstone's MPLS solutions, backup LSPs can be configured for fast fail-over, thus improving overall service reliability for the customer. Backup LSPs can be defined as hot-standby LSPs that have been pre-established, or can be dynamically created upon failure of the primary LSP. Another alternative is to enable the fast-reroute option when an LSP is established, leading to the creation of detour LSPs around each point of failure in the path. If a node or link fails, a local detour around the failure will be used and the ingress router will be notified. The ingress router can then decide to set up a new LSP.

When the primary LSP is restored, a user-defined option allows traffic to be switched back to the primary LSP.

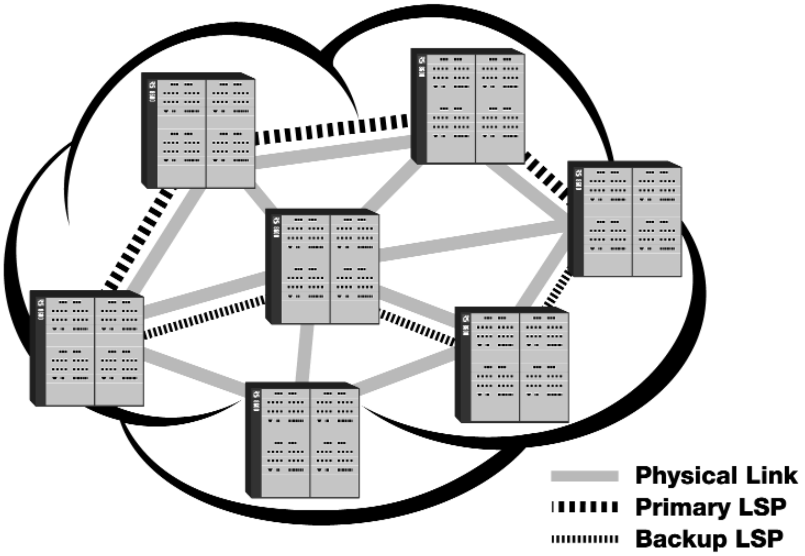

Figure 2: MSP Backbone (shown with RS 8600 Switch Routers).

Failures must be detected as quickly as possible. While the Riverstone MPLS solution is tightly integrated with the physical layers to detect local failures, it also lets users define RSVP "hello timers" to achieve a SONET-like fail-over time of 50 msec.

Quality of Service/Class of Service

Bandwidth has become such a commodity in the metro space that it's no longer sufficient for service providers to offer pure bandwidth alone. They need to differentiate their offerings and provide compelling reasons for the customer to choose their higher-grade services. Riverstone MPLS provides the tools necessary to deploy such offerings and the ability to differentiate service tiers with QoS and reliability guarantees.

Riverstone MPLS can map 802.1p or IP ToS/DSCP code points either to the MPLS Exp/CoS bits or to the LSPs for which resources have been reserved. The Exp bits in the MPLS header carry the packet's priority. Each label switch router along the path will honor the packet's priority by queueing the packet into the proper queue and by servicing the packet accordingly. When a packet is mapped to an LSP, the decision is made once and for all at the ingress point of the MPLS network. The packet will traverse a specific LSP that has been established for the corresponding packet class.

In a tiered service model, LSPs can be assigned different priorities. If a high-priority LSP fails and no resources are available to establish a new LSP, lower-priority LSPs' resources can be preempted. This guarantees that customers with higher-priority service-level agreements continue to receive service, while customers with best-effort services have no guarantee.

When a path is less than optimal, it becomes important to try to re-optimize an LSP in the background. For instance, if a high-priority LSP preempts a medium-priority LSP, a less-than-optimal medium-priority LSP might be established. Automatic re-optimization ensures that the medium-priority LSP will be reestablished with its original characteristics when resources become available.

Traffic Engineering

MPLS brings strong traffic-engineering capabilities to IP networks by providing explicit routing capabilities. IGP protocols do not take into account available bandwidth and link-state information when calculating routes. This could result in congestion on some paths in the network and in under-utilization of other paths. In nationwide IP networks, service providers have been using MPLS to conserve bandwidth by directing traffic in an optimal manner. The metro service provider's interest in traffic engineering, however, is to enable a responsive network to meet service-level agreements.

To minimize congestion in metro networks, Riverstone's MPLS supports traffic-engineering extensions to the IGP protocols OSPF-TE and IS-IS-TE. These extensions provide additional link-state information such as reserved bandwidth, available bandwidth and affinity, along with route updates. Riverstone MPLS also supports an online CSPF algorithm that dynamically computes an explicitly routed LSP.

Conclusion

Riverstone Networks, with its razor-sharp focus on the features that today's metro service provider needs, has pioneered cutting-edge MPLS technology for the metro network. Offering a wealth of service enablers, Riverstone switch routers have already proven highly successful in the MAN environment.

MPLS takes Riverstone switch routers' service-enabling functionality to the next level, enriching and expanding the MSP's ability to offer services in any given metro area and extend similar services to other metro regions — while at the same time enabling greater control over the network.

Acronyms

ACL	*Access Control List*
ANSI	*American National Standards Institute*
ASIC	*Application-Specific Integrated Circuit*
ASP	*Application Service Provider*
ATM	*Asynchronous Transfer Mode*
CBR	*Constant Bit Rate*
DS1/DS3	*Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)*
DSL	*Digital Subscriber Line*
E1/E2	*European Trunk 1/2 (2 Mbps/34.3 Mbps)*
ERP	*Enterprise Resource Planning*
HSSI	*High Speed Serial Interface*
ISP	*Internet Service Provider*
ITU	*International Telecommunications Union*
LAN	*Local Area Network*
LEC	*Local Exchange Carrier*
MAC	*Media Access Control*
MAN	*Metropolitan Area Network*
MDU	*Multiple Dwelling Unit*
MLPPP	*Multi Layer Point-to-Point Protocol*
MTU	*Multiple Tenant Unit*
OC-3/OC-12	*Optical Carrier 3/12 (155 Mbps/622 Mbps)*
POS	*Packet over SONET*
PPP	*Point-to-Point Protocol*
PVC	*Private Virtual Circuit*
QoS	*Quality of Service*
RED	*Random Early Discard*
SLA	*Service Level Agreement*
T1	*Trunk 1 (1.544 Mbps)*
TCP/IP	*Transport Control Protocol/Internet Protocol*
TDM	*Time Division Multiplexing*
UBR	*Undefined Bit Rate*
VBR	*Variable Bit Rate*
VLAN	*Virtual LAN*
VoD	*Video on Demand*
WAN	*Wide Area Network*
WDM	*Wave Division Multiplexing*
WRED	*Weighted Random Early Discard*

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

408 / 878-6500 *or* ***www.riverstonenet.com***

Printed in the USA *v 2.2 10/00*