Pipeline
Knowledge Is Power

www.pipelinepub.com Volume 4, Issue 4

VRF Visibility – The Key to MPLS Performance Assurance

by Bruce Kelley

To satisfy demands for secure, cost-effective transport of converged voice and business applications, most telecommunications providers have introduced MPLS-based VPN service offerings. Enterprises are taking advantage to break out delivery for voice and video into high-priority classes and tiered choices for their more latency-tolerant business applications. In order to meet service quality expectations, carriers are facing some new challenges:

- They need real-time, application-aware analysis of activity across the MPLS core
- They must distinguish traffic between individual customers and locations
- They have to track routing activity as an integral part of their traffic engineering tasks

In short, providers need to re-evaluate the ways in which they monitor performance as part of their assurance practices.

Network Considerations

Any performance monitoring approach needs to embrace awareness of how traffic is transported across the MPLS core. An MPLS network will adhere to the RFC 2547bis standard, which drives how VPN services are provided to customers. An MPLS network creates VPN tunnels based on MPLS routing and forwarding tables for each customer site connected to the service provider's MPLS network.

A customer site is connected to a service provider network via one or more ports, and the service provider associates each port with a specific VPN routing and forwarding identifier known as a VRF. Since each customer likely uses internal IP addresses which may be duplicates of other customers', a performance monitoring approach needs to examine and distinguish an individual customer's traffic between the PE (provider edge) and P (provider core) routers in the MPLS network. The figure below provides an illustration of a simple MPLS network deployment.

Two basic traffic flows occur in an MPLS VPN. The first is a control flow that is used for VPN route distribution and label switched path (LSP) establishment. The second is the actual data flow that is used to forward customer data traffic. Any MPLS monitoring approach will need to leverage information from both of these traffic

flows to troubleshoot and plan capacity.

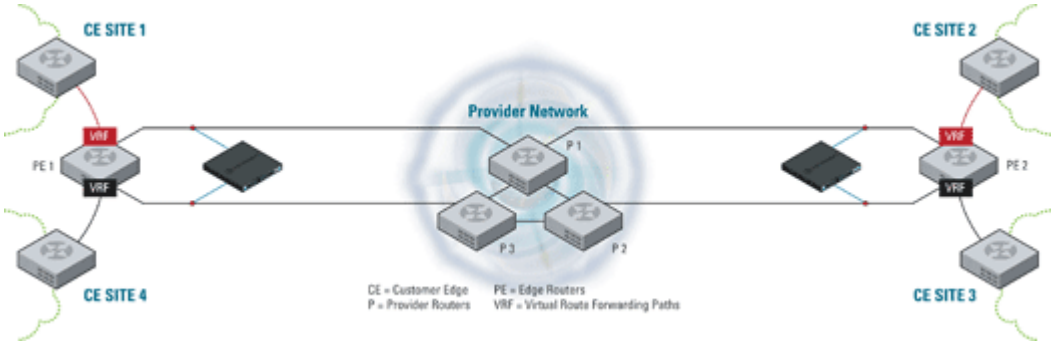

Possible Approaches

Many providers are recognizing the need for new methods of monitoring their MPLS services. First, they need to track dynamic MPLS labels in real-time, accommodate duplicate customer IP addresses, and recognize heterogeneous mixes of application traffic. This means that MPLS providers need approaches that provide detailed visibility by service and application, despite the fact that the viewpoint will be a network perspective. Also important is the ability to isolate and rapidly correct not just faults and failures, but a more challenging class of issues - service degradations. This is especially important with the advent and growth of latency-sensitive applications, such as IP voice and video conferencing.

Service performance assurance is a significant challenge for many operators because the sources of management data that they are accustomed to using do not provide the details required to meet the needs outlined above. Let's look at the most common options:

- Common monitoring technologies utilize aggregate metrics collected from network elements (routers and switches), commonly via standards-based interfaces such as SNMP. The advantage here is simplicity, but the disadvantage is also simplicity – this data does not provide details on which services or applications are active on a particular service link.
- Another approach is to use flow records, such as NetFlow, which are issued by edge routers. NetFlow can provide details on which application/service is being used and which end addresses are using it on a transaction or session basis. The disadvantages are twofold – first, NetFlow leaves out important details such as response time, and second, it creates compute load on the network elements that can impact its performance (especially during unexpected high traffic situations, such as a denial of service attack).
- Another alternative is testing agents that create synthetic traffic and measure approximate customer experience. Their advantage is that it is possible to make accurate approximations of synthetic performance and recognize degradations in the process. The major shortfall is that it does not measure actual traffic and cannot indicate why degradation is happening. It also generates non-revenue network load, and thus cannot be deployed exhaustively to monitor every combination of path and application/service type.
- Providers' traditional focus has been placed on the use of signaling traffic as

a proxy for measuring service activity and delivery. The advantage of this approach is that it provides a very complete record of sessions and service activities that have been initiated across the MPLS service links without creating load on the managed elements or adding traffic to the network. On the downside, signaling data does not provide any of the details necessary to recognize degradations, nor does it provide a basis for direct troubleshooting of performance issues. Another drawback is that all traffic in an IP network does not generate signaling traffic, thus creating significant blind spots on actual use of network resources.

- Finally, there is an approach that combines the best of all these worlds – deep packet inspection (DPI). The idea behind DPI is that dedicated instrumentation devices, attached at key traffic aggregation points, listen passively to all traffic and assemble a highly granular, complete view of all network- and application-layer transactions, including volumetric, response time, and latency measurements, along with details of all the underlying physical and virtual network constructs. This is a very complete set of data upon which to base customer-aware service performance assurance. The biggest challenge with applying DPI is identifying the optimal locations for deploying instrumentation, so that the most coverage can be established with the lightest total capital outlay.

The ultimate challenge, regardless of the approach that is chosen, is in adapting the data sets to accommodate the dynamic behaviors of MPLS and deliver a means to consistently and accurately track activity by customer.

VRF to the Rescue

In order to segregate monitoring of each customer's traffic, operators must first look at the MPLS label attached by the first Provider Edge (PE) router when the packets enter the network. Each customer and the routes their traffic traverses are both unique and dynamic, so the MPLS labels for each associated customer and PE router will be also be unique and will change as routes do. The key challenge is to be able to track and accommodate these updates in real-time.

In order to track, monitor, analyze and trend any customer data for both ingress and egress traffic that may have different labels from each end of a route, a more practical approach is to look beyond the dynamic MPLS labels to a more stable and consistent identifier. The VRF (Virtual Route Forwarding) table maintained by each PE router represents a better basis for monitoring and management. VRFs are normally assigned on a customer-by-customer or service-by-service basis, and do not change dynamically as do MPLS labels. By monitoring dynamic routing data, it is possible to establish the relationship between the MPLS traffic and VRF assignments, thus maintaining a consistent categorization of customer and service traffic.

With a real-time mapping of traffic flow metrics to VRFs, operators can much more readily monitor, characterize, and understand the activity and experience of each served customer and each active service. This capability is further enhanced by adopting a DPI-based data source, which provides the broadest and deepest information on both service volume and quality. Following are some common, everyday tasks that DPI-based VRF visibility enables:

- Identifying "bandwidth hogs" by viewing a core network segment and distinguishing between multiple individual VRFs assigned to it simultaneously
- Performing in-depth, packet level troubleshooting of an individual customer's activity to analyze and discover root cause of degradations
- Setting utilization or time-over-threshold alarms by VRF to recognize and address impending congestion problems before service quality issues impact customers
- Reporting on most utilized circuits highlighting individual VRFs to reveal that changes in bandwidth or new traffic prioritization schemes might be needed to accommodate "power user" customers or locations

There are other uses for VRF-based monitoring as well. In some provider environments, especially mobile operators that are delivering IP-based 3G+ services across an IP/MPLS core, VRFs are used to segment services and content sources. They may also be used to segment and manage Operations, Administration, and Maintenance (OA&M) traffic transiting a common core with revenue-bearing traffic. In these cases, monitoring performance and activity by VRF provides details regarding customer experience with the added benefit of revealing network and service behavior by service and/or business function.

In the end, it's all about customer assurance no matter what technologies you've deployed. The key ingredient to making this happen is to always have the appropriate visibility into the services delivered, the customers using them, and the network infrastructure delivering them. To use an analogy, not having this visibility would be the same as a pilot flying an expensive jet in the dark without instrumentation that gives the appropriate visibility -- common sense, don't you think?

If you have news you'd like to share with Pipeline, contact us at editor@pipelinepub.com.