

Regaining MPLS VPN WAN Visibility with Route Analytics

*Seeing through the
MPLS VPN "Cloud"*



Packet Design

Executive Summary

Increasing numbers of enterprises are outsourcing their backbone WAN connectivity to MPLS VPN service providers. MPLS VPN WAN services have been gaining in market traction against Frame Relay due to availability of higher bandwidth links, and their price advantage when delivering “full mesh” application traffic between many sites in the network, as opposed to simple hub and spoke. In addition, the outsourcing of the IP network backbone reduces the need to engineer those IP links. However, MPLS VPN WAN services come with some serious liabilities in terms of network management. Namely, once deployed, IT loses visibility into routing connectivity and reachability across their WAN backbone. The answer to this loss of network management information is route analytics technology, which leverages routing protocols themselves as an innovative source of network management information, delivering network analysis that can't be arrived at from any other type of technology or tool. This paper reviews Layer 3 MPLS VPNs, and outlines common network management concerns voiced by enterprise IT managers when outsourcing their WAN to an MPLS VPN provider. The paper continues with a brief explanation of route analytics technology, then illustrates how route analytics technology can be used to regain critical WAN visibility, reduce finger-pointing, keep Service Providers accountable, and more successfully deliver end-users' application traffic.

Layer 3 MPLS VPNs—A Layer 3 Network Management Challenge

Layer 3 MPLS VPNs are delivered by Service Provider IP/MPLS networks that are organized into a core of Provider or P routers, and a layer of customer-facing Provider Edge or PE routers. PE routers are configured to handle multiple VPNs through separate virtual routing and forwarding (VRF) tables. Each customer's VPN is handled by dedicated VRFs on various PEs located in different geographies across the Service Provider's network. Interconnectivity between VRFs is delivered by a mesh of MPLS tunnels, with a special extension of BGP providing control plane mapping of tunnels to VRFs.

When using MPLS-based VPN services, enterprise customers are responsible only for connections from each site to the Service Provider network, by connecting a Customer Edge (CE) router to a PE router and enabling routing, typically using the Border Gateway Protocol (BGP). The benefit of outsourcing the WAN backbone is that it offloads the management of this network to a service provider's shared infrastructure. However, this reduction of work introduces a significant network management challenge—which is that the enterprise's Layer 3 network has been effectively partitioned, and Layer 3 visibility into its WAN backbone has been taken away from network managers. This is because MPLS VPNs were designed to shield the Provider's network domain from all customer network domains, in order to preserve the privacy of each customer VPN within the shared provider backbone.

The loss of comprehensive administrative visibility into the end-to-end IP network infrastructure would be easier to cope with if the interface to the MPLS VPN were simple to manage, but in fact the opposite is true. In Frame Relay networks, the Customer-to-Provider interface was a static Layer 1/2 connection, with statically configured PVCs, which made managing the service relationship relatively simple. The main two criteria for

managing Frame Relay services were link availability (up or down), and link performance. As a result, periodic SNMP polling of Frame Relay routers, links and subinterfaces/PVCs was sufficient to gather the network management information needed to at least monitor the Service Provider's WAN links—if not the routed Layer 3 network operating over it.

By contrast, with MPLS VPNs the Customer-to-Provider interface is a highly dynamic and complex Layer 3 IP peering. Not only does the CE-to-PE link need to be up and providing correct link speeds, but it must also provide proper IP routing. Route advertisements and withdrawals are highly dynamic, so traditional multi-minute SNMP polling cycles are completely inadequate to monitor them. Furthermore, routing can be malfunctioning, even if all the links are up and providing adequate performance, making it necessary to understand new routing-based criteria to manage the service. In other words, SNMP polling could show a CE router and its interface to the PE up and running properly, while there was a complete service outage due to a routing problem.

Complicating matters is the fact that the BGP routing protocol used for CE to PE peerings is very complex, highly verbose, difficult to analyze, and relatively easy to misconfigure. In addition, many enterprise network engineers have limited experience with BGP, since it's traditionally been used by Service Providers.

Key Layer 3 Management Criteria for MPLS VPN WAN Services

Network managers tasked with managing a MPLS VPN WAN backbone must ensure that they have the management information needed to monitor and localize routing issues within the end-to-end network, and thereby determine whether the root cause is stemming from the Provider or from the internal network. Without sufficient visibility, enterprises can get caught in a classic finger-pointing exercise, while critical application availability and performance targets are missed. Following are key Layer 3 management criteria that network managers need to be concerned with beyond SNMP device and Layer 2 link monitoring:

- VPN Privacy and Integrity: “How do I monitor the routing integrity of my VPN?” Specifically, how can a network manager be sure that his network is not being mixed with another VPN customer network?
- CE Reachability: “How do I know if my CE routers have a proper routing connection to the VPN?” Specifically, even if the CE to PE Layer 2 connectivity is “up”, is the CE actually speaking BGP properly to the PE router so that it can send routes to the VPN and vice-versa? How can network managers know rapidly when things aren't working properly from a CE-to-PE routing point of view?
- Remote Site Prefix Reachability: “Even if the CE routers have a proper routed connection to the VPN, are the remote sites' routes getting advertised to the VPN?” Specifically, how can network managers track the consistency of their route reachability to all of their remote sites without having to perform brute force analysis? Is there a way to actually see – and monitor - which routes are being advertised or not?

- Forensic Analysis of Reachability Issues: “How can I see exactly what happened in the past, to solve problems and give customers forensic information?” Is there a way to retrace the event history rather than guess at what happened, or worse, just wait for it to happen again?
- Remote site IGP monitoring: “How do I get insight into potentially complex routing issues within the remote sites?” How can network engineers see deeply into the inner workings of the networks *behind* the CE router, especially if they are in large campus environments with their own IGP routing complexity?

Route Analytics—Unique Insight Through the MPLS VPN Cloud

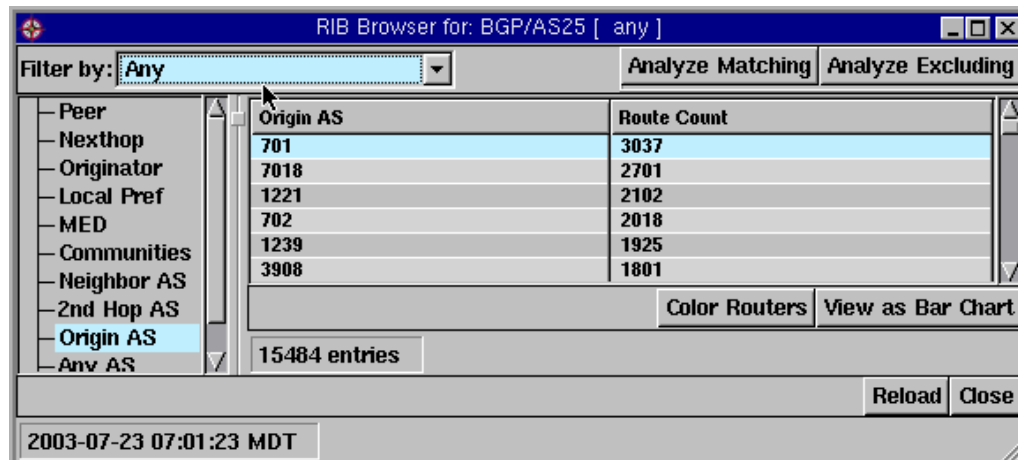
Route analytics is a technology deployed in hundreds of large enterprise, government and service provider networks worldwide, that offers an effective solution to the crippling lack of management visibility that results when using MPLS VPN WAN backbones. Route analytics solutions listen passively to routing protocol exchanges on the network and deliver a “router’s eye view” of Layer 3 connectivity and reachability, providing network engineers with previously unavailable intelligence on the real-time Layer 3 operation of an IP network. Route analytics works by forming passive (listen-only) peerings with a few key routers in the network using common routing protocols such as BGP, OSPF, IS-IS and EIGRP, recording every routing protocol update, and creating a model of the network that is as accurate as the network itself. In the case of BGP, route analytics peers with the CE routers and receives all the routing updates that the CE router receives from other CE routers via the MPLS VPN, though it never advertises routes or makes any changes to routing in the network, so it can’t adversely affect connectivity. Route analytics provides a variety of analysis tools that help network managers make sense of what is happening to their WAN. For a full introduction to route analytics technology, features and benefits, please read the Packet Design White Paper entitled “Route Analytics—Foundation of Modern Network Operations”, on Packet Design’s website: <http://www.packetdesign.com>.

Monitoring MPLS VPN Privacy

Route analytics can help network managers ensure the privacy and integrity of an enterprise’s MPLS VPN backbone by monitoring BGP Autonomous System Numbers (ASNs). When connecting to a provider’s layer 3 VPN service using BGP, each of the enterprise’s sites must have a unique Autonomous System Number (ASN), typically private ASNs assigned by the Service Provider. These ASNs in effect represent the list of VPN sites. The Service Provider’s network should never inject routes into the customer’s VPN that are from an unknown ASN, as this would indicate that another customer’s VPN has inadvertently been connected into the VPN.

One of the simplest ways to determine at a VPN site that no unintended routes are being received is to use route analytics to examine the origin Autonomous System (AS) of all of

the received routes at that site. Figure 1 below shows an example of such a summary provided by route analytics.



The screenshot shows a window titled "RIB Browser for: BGP/AS25 [any]". It features a "Filter by:" dropdown menu set to "Any", and two buttons: "Analyze Matching" and "Analyze Excluding". A tree view on the left lists various BGP attributes, with "Origin AS" selected. The main area displays a table with two columns: "Origin AS" and "Route Count". The table contains the following data:

Origin AS	Route Count
701	3037
7018	2701
1221	2102
702	2018
1239	1925
3908	1801

Below the table, there are buttons for "Color Routers" and "View as Bar Chart", and a text box showing "15484 entries". At the bottom right are "Reload" and "Close" buttons. The status bar at the bottom left shows the date and time: "2003-07-23 07:01:23 MDT".

Figure 1: Listing of ASNs seen by route analytics and their respective advertised route counts. If a foreign ASN appears in this listing, then network managers know that the privacy and integrity of their VPN service has been compromised.

If any routes are received from ASNs not owned by the enterprise, it will be immediately obvious in a summary such as this. From this summary view, network managers can drill down to historical event details to determine exactly when the “foreign” routes from improper ASNs were introduced to the VPN service, where they were introduced, view their characteristics (BGP attributes), and see when (if at all) they were withdrawn. Using this information, network managers can hold their Service Provider accountable for service privacy issues.

Route analytic enables another way to monitor the integrity of the MPLS VPN that is not dependent on AS numbers, by setting alerts on significant increases or decreases in the number of prefixes in the VPN. Since an enterprise WAN should be relatively stable in the number of its advertised prefixes, if a large number of prefixes are advertised into the network in an unexpected manner, then it is possible that the Service Provider has inadvertently mixed customer VPNs. When this sort of alert is triggered, engineers can easily get an at-a-glance view of baseline vs. active routes in the network, as shown in Figure 2.

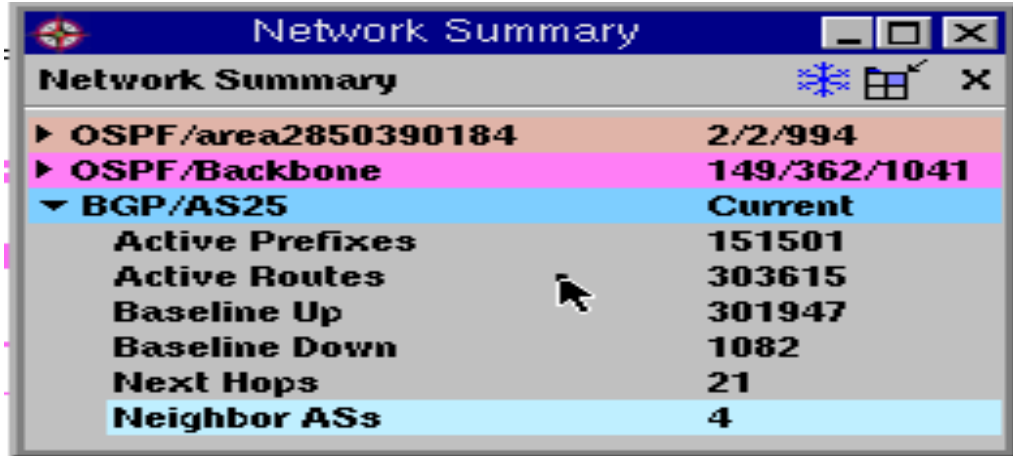


Figure 2: Route analytics provides a network summary of routing state. In the case of BGP, it shows baseline routes that are up, baseline routes that are down, and total active routes. If there is a significant difference in active routes vs. baseline up minus baseline down routes, then there may be a problem with the VPN's routing integrity within the Service Provider's network.

Once an increase of prefixes has been detected, engineers can use route analytics to look at a list of all known prefixes. Since route analytics baselines BGP routes, the table can be easily filtered to show any routes that are up but not in the baseline, as seen in Figure 3.

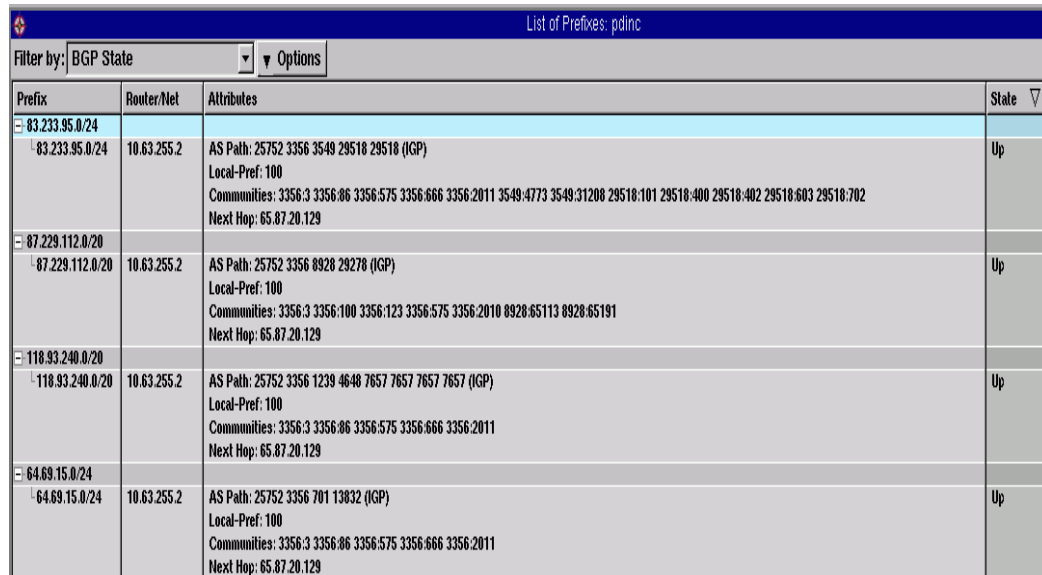


Figure 3: Route analytics baselines BGP routes and allows engineers to easily filter views of the VPN's prefixes to find any active prefixes that aren't in the baseline and are therefore suspect

If unknown prefixes are found, the engineer can also look at a histogram of prefixes in the network and find the time when the prefixes entered the network by looking for a jump in the prefix graph. From there, route analytics allows engineers to drill down to detailed tables of BGP events in a specified time range to further analyze the newly advertised routes.

MPLS VPN Reachability Visualization

Route analytics can provide a view of the MPLS VPN topology, with each of the site ASes connected to a central Service Provider AS. This visualization shows numerical counts and visual weighting of how many prefixes are carried by each routing branch. This view allows network engineers to see at-a-glance the behavior of BGP routing in the network, assess routing policies, enable evaluation of existing peerings and plan future peering or routing policies. This view is particularly valuable when assessing the status of dual VPN provider routing and reachability. Network managers can select how much detailed information they want to view, or reduce visual clutter by filtering on various parameters. An example of BGP AS-level topology visualization can be seen in the view of a large campus network shown in Figure 4.

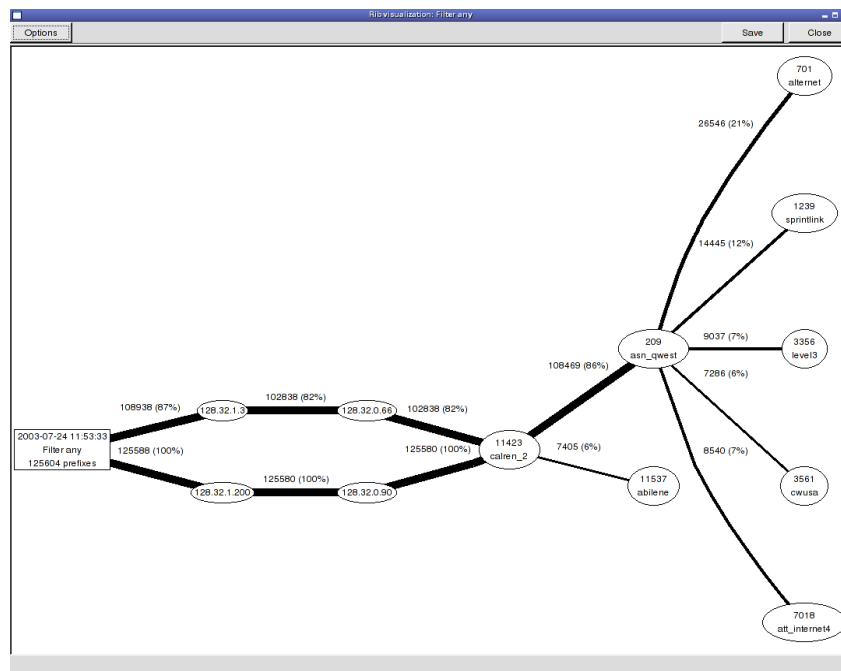


Figure 4: A multi-hop BGP topology visualized by route analytics

Monitoring CE Reachability

Route analytics can be used in conjunction with CE loopback address advertisement to provide real-time monitoring of remote CE reachability—telling engineers whether or not remote site CE routers are communicating properly via BGP to the rest of the VPN. Every CE router has an internal loopback IP address, and while they are not usually advertised into the MPLS VPN, by doing so every CE's internal address will become visible via route analytics, which hears all routing prefixes advertised across the MPLS VPN. Since route analytics is aware of all paths (including BGP AS paths) in the network, it can be configured to watch reachability to particular routes. By monitoring and sending traps on changes in the path to all remote CEs, network managers can be rapidly alerted if any remote CE loses routing connectivity to the VPN. This is particularly useful since in some cases, Layer 2 connectivity may still be up while routing has gone down. Furthermore, other monitoring methods, such as SNMP, can introduce significant delays in monitoring due to multi-minute polling intervals. By contrast, since the BGP routing protocol will immediately communicate loss of reachability to a prefix, route analytics provides extremely rapid response-time monitoring. Another method for monitoring CE reachability is by monitoring and alerting when CE loopback BGP prefix paths drop from single to zero paths.

In the case where network managers need to know whether redundancy has been lost from dual provider-connected CEs (where the CE is connected via redundant MPLS VPN provider networks), route analytics allows for monitoring BGP prefix path redundancy. Since route analytics understands multiple paths to any given BGP prefix, it can monitor whether a prefix associated with a CE router has experienced a loss of redundant reachability via one of the MPLS VPN provider networks. In these cases, route analytics can be configured to send a SNMP trap or syslog alert if prefix redundancy reduces from dual paths to a single path.

Enterprises can also purchase redundant CE connectivity to a single MPLS VPN provider's networks. This could involve implementing dual CEs at a sensitive site, or require dual links to the MPLS VPN provider's network. In either case, it is advisable to ensure that the dual connections are terminating on separate PEs. Not only is there an increase in true redundancy of the connectivity, but dual PEs also increase network management visibility. With redundant PE connectivity to a CE site, similar to the dual provider-connected case, route analytics can monitor and alert on loss of redundant paths.

Finally, loss of CE reachability can be monitored by looking at the listing of ASNs shown in Figure 1. If a known remote site ASN is dropped from the listing, then network engineers will know that the CE has lost its BGP peering with the MPLS VPN, and can start to troubleshoot the BGP event details recorded by route analytics to understand more.

Monitoring Remote Site Prefix Reachability

Even if remote CE routers are communicating properly to the MPLS VPN via BGP, it is still important to ensure that the prefixes behind the CE are reachable, since while a CE BGP failure will drop all prefixes, individual or groups of prefixes can still be withdrawn, experience flapping or other problems while the CE to PE routing is working flawlessly.

Route analytics monitors a number of conditions, helping network managers quickly localize a problem domain in the network:

- Comparing lost prefixes to the CE loopback reachability: If all prefixes from an AS are withdrawn but the CE router's loopback address is still advertised, then it means that there is a routing issue between the CE and the networks behind it, rather than a problem with the CE routing through the MPLS VPN.
- Alerting on Gained or Lost Route Thresholds: Route analytics can be configured to measure an ongoing baseline of the MPLS VPN's BGP routes and alert if the number of routes exceeds or drops below a pre-determined percentage of that baseline. In most enterprises, once a MPLS VPN service is fully deployed, there should be relatively little variation in the number of routable prefixes. In an MPLS VPN with 1000 prefixes, it may be unusual and notable if more than 20 prefixes are lost or gained at a time, so in this case the threshold could be set at 2%. While this threshold will not catch all dropped or flapping prefixes, it will show if something significant is happening.
- Key Route Lost Alerting: In addition, specific alerts can be triggered for certain key routes, such as to data centers, so that if routes are lost to those facilities, network managers are alerted in real-time.

Forensic Analysis of Reachability Issues

One of the biggest challenges with managing complex, redundant IP networks is understanding precisely what happened in the past, whether five minutes or five days ago. This is no less true of trying to troubleshoot what happened in a MPLS VPN service problem. Fortunately, route analytics continuously records all BGP events into a database and can analyze, visualize and even animate past event streams. For example, BGP root cause animations allow network managers to replay a large stream of events and visualize advertised, withdrawn, and moving routes between peerings across a multi AS-hop topology, such as in Figure 5. Once the "big picture" of the problem is understood through the animation, network managers can then step through the event detail with a greater contextual understanding to find the root cause of the issue.

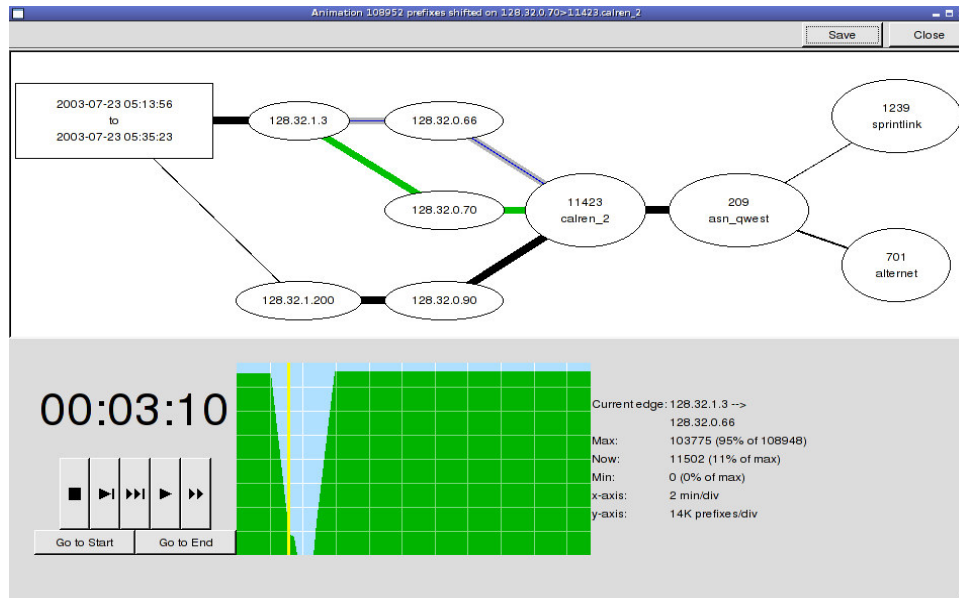


Figure 5: BGP root cause animation of a large number of advertised, withdrawn and shifting routes in a BGP event stream across multiple ASs makes forensic analysis of MPLS VPN reachability issues more intuitive

Monitoring of Remote Site IGP Routing Issues

Some CE sites have extensive IGP routed networks behind them, perhaps with multiple OSPF areas or EIGRP AS. In these cases, network managers also need to be able to get insight into routing issues within those IGP domains, especially in cases where WAN reachability issues are traced to the IGP domain behind the CE. Route analytics provides extensive OSPF, IS-IS, and EIGRP monitoring, historical analysis and even scenario modeling. For more details on how route analytics can be used for a variety of network management purposes, please visit Packet Design's white paper library at: <http://www.packetdesign.com/technology/wp.htm>

The Benefits of Route Analytics for MPLS VPN WANs

Route analytics offers IT a number of benefits when deployed to help manage MPLS VPN services:

- Faster response times due to real-time alerting on critical network events. Unlike SNMP, routing protocols operate with milli-second response times. Since route analytics "sees" network events at the same speed as routers, network managers get the benefit of real-time alerting on critical network issues such as CE to PE



peering outages or lost redundancy, lost or suspicious additions of routes to the VPN, and catastrophic loss of reachability to a remote site's routes.

- Ability to localize the network problem domain and reduce finger pointing. Rather than wasting time wondering who's to blame for a problem or waiting for the provider to respond, network managers can now proactively find the source of reachability issues that impact application delivery.
- Forensic trail to keep providers accountable. In the case where a provider has caused a reachability problem, network managers now have a complete recorded history of all routing events in the network in order to ensure that they can hold their provider accountable for violating their Service Level Agreement (SLA).

Conclusion

Route analytics provides the next-generation network management capability needed to meet the demands of MPLS VPN WANs. With network managers increasingly being "graded" on application delivery rather than just basic infrastructure availability, route analytics' Layer 3 visibility is a must-have network management capability to ensure successful MPLS VPN WAN deployments.

To learn more about Packet Design and route analytics, please:

- Email us at info@packetdesign.com
- Visit Packet Design's web site at <http://www.packetdesign.com>
- Call us at 408.490.1000