Tequipment.NET USA
205 Westwood Ave
Long Branch, NJ 07740
1-877-742-TEST (8378)
Fax: (732) 222-7088
salesteam@Tequipment.NET

White Paper

Migrating to an MPLS-based/Private IP Network: Are You Ready?

With the increased reliance on business-critical applications, enterprises are considering MPLS-based private IP networks to improve performance with class of service prioritization and any-to-any connectivity. This white paper will discuss the benefits associated with private IP networks as well as the challenges that occur when networks go from Layer 2 to Layer 3 based connectivity.

Table of contents

FLUKE networks

Today's enterprises are faced with a number of network-related needs that are challenging from both a technology and a financial perspective. Organizations are looking for proven, cost-effective solutions to:

- Create and improve disaster recovery infrastructures.
- Replace outdated hub-and-spoke architectures.
- Meet increasingly complex growing network requirements.
- Prioritize certain applications, such as VoIP.
- Maximize performance while minimizing costs.

Your enterprise might have one or many of the needs listed above and is in the midst of evaluating solutions to meet your goals and objectives. One of today's hot topics for network and application executives to help meet these growing demands is multi-protocol label switching (MPLS) based/private IP networks. Private IP networks (generally MPLS-based, but not exclusively) are rapidly growing in popularity for medium- to large-sized enterprises across a wide range of industries. Benefits for adopting these networks include the ability to prioritize bandwidth by application, the potential for automatic redundancy, and the value of a fully-meshed network with less complexity and cost.

However, before you begin migrating your existing network to a private IP network, consider these challenges of MPLS-based networks: what applications you are running, how to prioritize them, and the need for visibility into individual IP Subnet pairs.

Many white papers describe technically how MPLS works – down to the details of configuration across different customer premise equipment (CPE) types. Before you take a deep dive into the inner workings of the technology, determine if migrating to MPLS is the best decision for your enterprise. Thoroughly investigate the key business reasons your enterprise might consider a private IP network as well as the issues you must monitor and manage to make MPLS a success.

MPLS-based private IP – site-to-site IP VPNs: what's the difference?

Before discussing private IP networks in more detail, let us first appreciate the number of different flavors of IP virtual private networks (VPNs). The term "IP VPNs" has many different connotations, but in general, IP VPNs are broken into two distinct categories: site-to-site and MPLS-based/private IP VPNs.

Site-to-site IP VPNs

A site-to-site IP VPN, also known as a CPE-based VPN, is defined as using the public Internet as the core backbone of the VPN. Enterprises use dedicated Internet connectivity, cable, wireless, DSL, satellite and dial-up as different means to connect to the IP VPN. Enterprises then typically use CPE (VPN gateway, router, etc.) to create a VPN with tunnels and encryption so different remote sites can communicate with the headquarters and with other remote sites directly and securely.

The primary benefit for site-to-site IP VPNs is tied to access – low cost and ubiquity. Traditionally, Internet connectivity is much less expensive than frame relay or ATM (and the difference grows quite substantially for international locations). The ubiquitous nature of Internet connectivity is a major benefit as well – IP is IP regardless of whether it is dial-up, dedicated, domestic, or international.

While there are cost savings to access site-to-site IP VPNs, the total cost of ownership can actually be higher for enterprises because CPE is needed at every site to create the VPN. IT support costs are high, with the need to manage the additional devices and locations. The potential exposure of using the public Internet for business-critical traffic is extremely high. The Internet's architecture was not designed for business-grade applications and requirements and does not offer strong service level agreements (SLAs) for enterprises. The dangers associated with outages, over-utilization and hacking are so severe many enterprises will not risk business-grade services or applications to a public IP VPN. Enterprises have used public IP VPNs as a way to connect remote offices or satellite employees instead of building the entire business-critical infrastructure via the Internet.

MPLS-based private IP VPNs

An MPLS-based private IP VPN, also known as a network-based VPN, is defined as using a service provider's offering (usually frame relay and ATM and with more frequency, Ethernet) as the backbone of the VPN. MPLS is a solution that allows the service provider's core routers to transition from Layer 2-based connectivity to Layer 3-based connectivity.

For example, traditional frame relay requires a dedicated permanent virtual circuit (PVC) for remote sites to communicate between each other. The traffic follows the path for every transaction between two sites.

With a private IP network, the carrier's routers use MPLS-based routing to communicate between locations in the network via IP addressing and is not limited to individual PVCs. Tunneling and encryption are generally not required for the enterprise because the network is on the private carrier's backbone, not the public Internet. Additionally, the carrier's core routers create the VPN, not the enterprise's CPE.

A key advantage of a private IP VPN is the ease of migration from a frame relay or ATM network. For most service providers, there is no need for provisioning new circuits (since the frame relay and ATM is still the Layer 2 access) and routers may not need to be replaced. The complexity to the enterprise of managing the VPN is less with private IP networks because it is network based, meaning the provider, by means of its core routers, handles the creation of the VPN.

The frame relay world considers two bandwidth numbers: port size and committed information rate (CIR). The private IP world manages two criteria: committed access rate (CAR) and thresholds by class of service (CoS). The frame relay port size and private IP CAR are very similar – this is the maximum amount of bandwidth that can be used for a single circuit/location. The differences lie between CIR and CoS.

Enterprises purchase CIR per PVC as a way to "guarantee" bandwidth from the service provider. Private IP VPNs provide a CoS capability to the enterprise by allowing multiple classes of prioritization, each a percentage of CAR. CoS is more application-focused since enterprises can select which applications are the most critical or the most delay-sensitive and assign them the highest priority. The network-based VPN polices and queues the classes instead of deploying traffic shapers to prioritize the traffic.

Now let's look at the benefits of a private IP network in greater detail.

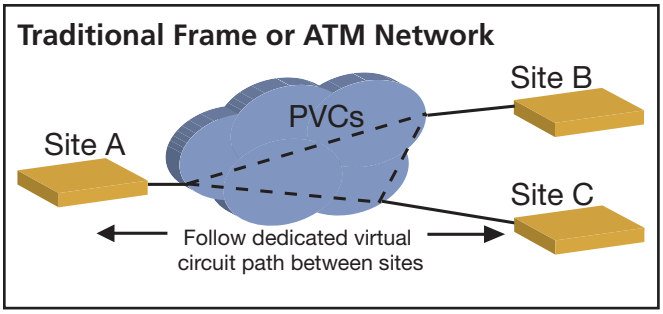

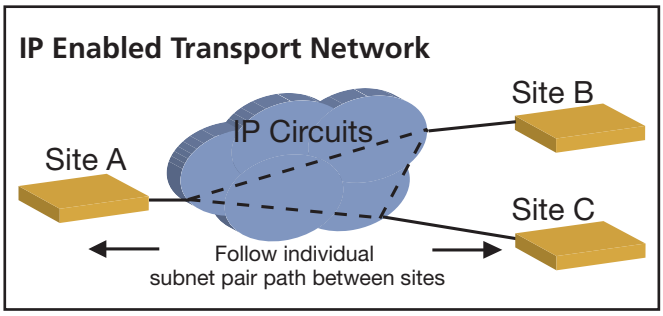

Private IP – what are the benefits?

Assuming your enterprise already has a wide area network (WAN), the question most people ask themselves is why migrate an existing network that is working fine (hopefully working fine) to a new, unknown platform or technology?

Any network manager working 10 years ago or so knows why enterprises migrated from point-to-point/private line networks to frame relay. It was cheaper! While there will likely be cost savings for access by migrating to private IP, there are several other key benefits of moving from a frame relay or ATM network to a private IP VPN. These benefits include CoS, automatic redundancy/disaster recovery, fully-meshed infrastructure and reduced complexity.

Class of Service (CoS)

Prioritizing applications by CoS has been the key driver for many enterprises rolling out private IP networks. Instead of using critical router resources to shape traffic or buying an expensive traffic shaper to prioritize your applications, the service provider's core routers do the "heavy lifting" in prioritizing different classes of service in private IP networks.

Enterprises decide which applications are the most mission-critical or delay-sensitive and assign them specific classes. Most service providers today offer between three and six different classes of service. They might have different marketing names, but the key point is they differentiate policies for the applications to travel through the network – each class has a different priority setting than the others. In addition to the number of classes, the provider also allows a certain amount of bandwidth per class (depending on total circuit size, number of classes and contract agreement). Once you know the number of classes and the amount of bandwidth assigned to each, you are ready to begin assigning priority.

For your enterprise, Voice over IP (VoIP) and Oracle might be the most important applications, so they are set for the highest priority while e-mail and Web browsing receive lower assignments. You configure your router to tag the correct class in the IP header (DiffServ or TOS bit settings) and pass the data to the service provider's edge router. The provider's router looks at the IP header for the class setting and polices the traffic. The policing from the source site (ingress) ensures the threshold of bandwidth is not exceeded for each class setting. The core router then sends the traffic across the network with prioritization.

At the destination location (or egress), the provider's router queues the prioritized traffic from each site and delivers the highest priority first. With the capability of CoS, you can easily implement a prioritization schedule across the enterprise so the most mission-critical applications receive the highest priority.

Automatic redundancy/disaster recovery

The need for infrastructure redundancy has grown substantially over the past few years as more enterprises drive revenue, reduce costs, or provide services based on applications on the WAN. If a credit card authorization site is down, a retail enterprise may lose sales because today's consumers carry less cash. If a manufacturing company cannot transmit its stocking order to a plant, the vendor might impose penalties. These two examples are common scenarios that are exacerbated when there is a single point of failure in the infrastructure. That does not always mean the network itself is the cause of the problem; it can be any portion of the infrastructure that does not allow the business transaction.

If a transaction cannot be completed at the primary location, enterprises can forward it for completion to a back-up location. In a frame relay or ATM environment, this is typically handled by provisioning additional PVCs between critical sites. While this architecture is viable, it can be extremely expensive depending on the size and scope of the network. Managing tens or hundreds of PVCs at multiple sites can be a logistical and managerial nightmare – especially at remote locations with no networking support staff. Enterprises must determine if the exposure to lost revenues and incurred additional costs exceeds the extra resources and complexity for a disaster recovery infrastructure.

A private IP network is fully redundant based on its Layer 3-based connectivity (IP subnet-to-IP subnet). Now each site in your entire network has access to every other site. Implementing a disaster recovery/back-up strategy is much easier and more cost-effective since dedicated PVCs are no longer needed for each and every location.

Fully-meshed infrastructure

In a hub-and-spoke environment, traffic from a remote site must traverse a host location before it is routed to the destination address. This architecture was adequate several years ago, but as Web-based applications or time-sensitive applications like VoIP become more prevalent, the hub-and-spoke architecture becomes more stressed.

A fully-meshed architecture (meaning every site can communicate directly with any other site without having to run through a hub/host location first) has two key benefits for most enterprises: improved site-to-site performance and less burden on host locations.

When an application such as VoIP is used between two remote sales offices, there is no benefit to "home-running" the application back to the host location. This scenario adds potential delay by adding extra steps and distance to complete the application transaction. With

a fully-meshed network, a VoIP call from between two offices in London will flow directly between the locations instead of having to go to the hub site located in Los Angeles. In a private IP network, the number of steps and physical distance alone (not even including CoS capabilities) can be greatly reduced.

In conjunction with improving parameters with a fully-meshed network, including delay or jitter, reducing host site bandwidth usage is also a major benefit. Since the organization does not bring all traffic back to the host site, this may greatly reduce the bandwidth requirements at a host location, as well as enhance the performance across the entire infrastructure.

Reduced complexity

These examples focus on performance of the network and the application. The complexity of managing a frame relay or ATM network grows exponentially with network size and the number of virtual circuits. Managing hundreds of sites and thousands of PVCs is a daunting task for many enterprises just to handle moves, adds, and changes on a daily basis.

As discussed earlier, IP subnet addressing is used to connect every site in the network. PVCs are no longer the connection between sites. Instead of managing a port for every site and tens, hundreds, or thousands of individual PVCs, the private IP network has a single port for each site and then uses IP addressing to connect to every other site.

This architecture is much less complex, meaning it is easier to administer and allows enterprises to focus limited resources on more important activities.

Private IP sounds great – do I jump on board?

Every enterprise can likely benefit in some fashion from private IP. Before deciding to migrate your existing frame relay or ATM network to private IP, consider these key factors and to achieve the maximum benefit with the minimum pain.

During the evaluation and decision process, you should consider the following:

- How applications should be prioritized
- How to manage CoS thresholds
- Measuring SLAs across multiple classes
- Optimizing bandwidth requirements
- Managing throughout the network migration

How applications should be prioritized

Most enterprises know the most business-critical applications on the network today either by management direction, trial and error or which ones get the most complaints if there are performance issues. But after the first handful of applications, many enterprises have difficulty slotting other applications in medium or low classes. Compounding the problem is a lack of knowing what applications are actually on the network. A recent study shows 75 percent of respondents said they lacked sufficient knowledge of applications on the network.[1] If you do not know what applications are on the network, it is impossible to prioritize them.

Another challenge for network managers is the addition of new, powerful applications. If the enterprise decides to implement Oracle or SAP in the next six months, how will the new high-priority application affect existing high-priority applications? Networks and applications are continuously evolving and changing, and CoS prioritization must be able to adapt hand-in-hand.

How to manage class of service thresholds

Because you are now prioritizing applications with more granularity, managing CoS thresholds becomes absolutely critical. As mentioned earlier, service providers allow a certain amount of bandwidth for class of service. As long as you stay below that usage threshold, there should not be any problems. If you exceed the threshold, the pain can be magnified. If the network is congested and exceeds your highest priority class threshold, the traffic will drop to the lowest priority or may be discarded. So now your most mission-critical applications may be at risk.

Rolling out new applications across the infrastructure compounds the problem. A new application will be assigned a class and its usage may cause a domino effect for other applications. For example, a new custom financial application might cause the highest priority class to exceed the threshold, so one of the other applications in the same class will be dropped one prioritization lower, and so on. As an enterprise with mission-critical applications, you cannot risk the highest priority applications being downgraded because the class threshold was exceeded.

Measuring service level agreements

Many enterprises have SLAs with the service provider as an attempt to ensure the network can handle its applications. Measuring SLAs from a provider's point of view is delivering a monthly report with weighted averages across the entire infrastructure for delay, throughput and availability. Private IP compounds the issue by eliminating traditional virtual circuits, as well as layering on multiple classes of service.

Before private IP, it was difficult for many enterprises to proactively monitor service level parameters from an end-to-end point of view. Now with the extra complexity, some enterprises find it extremely difficult to measure SLA parameters across IP subnets and by individual classes of service. This lack of visibility makes it more difficult to leverage the day-to-day benefits of SLAs.

Optimizing bandwidth requirements

In today's world of squeezed budgets, optimizing bandwidth is critical for many enterprises. While enterprises need to ensure they have sufficient bandwidth to meet the needs of applications and end-users, they do not have the luxury of over-provisioning every circuit in the network. But as application performance becomes more important, the enterprise cannot be so aggressive and not have sufficient bandwidth resources. Most organizations walk this fine line.

As you migrate to a private IP network, there will be a change in your bandwidth requirements across the majority of your locations. The biggest challenge is understanding the magnitude of the changes. Will you need more bandwidth at some remote locations because of classes of service? Will you be able to reduce bandwidth at a previous hub site because you are no longer on a hub-and-spoke architecture?

Managing throughout the network migration

The ease in migrating from frame relay or ATM to private IP was highlighted earlier in this paper. This migration becomes easy because organizations can leverage existing Layer 2 infrastructure and CPE. Enterprises cannot let that ease lull them into a false sense of security as they plan migration efforts.

Even though it may be a fairly simple migration path to private IP, the transition takes months. You will not be able to snap your fingers and have 75 sites converted from frame relay to private IP overnight. You must manage your network throughout the entire migration. Some of your locations may be frame relay while others are already converted to private IP. Without visibility into the performance of both, you will be managing blind. That is unacceptable for most organizations, even those in the midst of a migration.

How do I solve the private IP challenges?

Enterprises are seeing the benefits of a private IP infrastructure, but remain worried about some or all of the challenges we highlighted. You are likely weighing the benefits of private IP, but are uncertain if you should migrate from your existing infrastructure – one that is working for you.

Some people follow the old adage, "If it isn't broke, don't fix it." Sticking strictly to that mentality, you would still be using carbon paper to make copies or pen and paper to balance the ledger of a billion-dollar organization. In a networking context, it would be like building a 300-site network with purely dedicated circuits. That strategy would still work today, but would probably not be the most cost efficient or robust. You need to find a way to eliminate the challenges of migrating to private IP so you can reap its benefits.

Visual Performance Manager – solving private IP challenges

Fluke Networks has been a pioneer in developing network performance management and SLA validation solutions for more than a decade. As networks have evolved, so has our technology. Visual Performance Manager provides complete visibility for managing MPLS-based networks through distributed deployments and/or leveraging existing infrastructure. By allowing either approach or a combination, enterprises can easily monitor and manage the complex MPLS-based network within in single system.

Visual Performance Manager provides detailed visibility into the nuances of the new private IP environments. By leveraging the core Layers 1 and 2 strength and incorporating new functionality for CoS and IP subnets, it helps enterprises understand and solve the challenges of migrating to a private IP network.

Prioritizing applications

Visual Performance Manager helps you identify and solve the two critical aspects of prioritizing your applications – what applications are on the network and whether the applications are prioritized correctly. As referenced earlier, three-quarters of enterprises claim insufficient knowledge of applications within the enterprise. Visual Performance Manager auto-discovers the applications across the entire infrastructure (see Figure 1). So instead of prioritizing a subset of all the applications, you now have the knowledge of total application usage so you can completely assign the CoS settings.

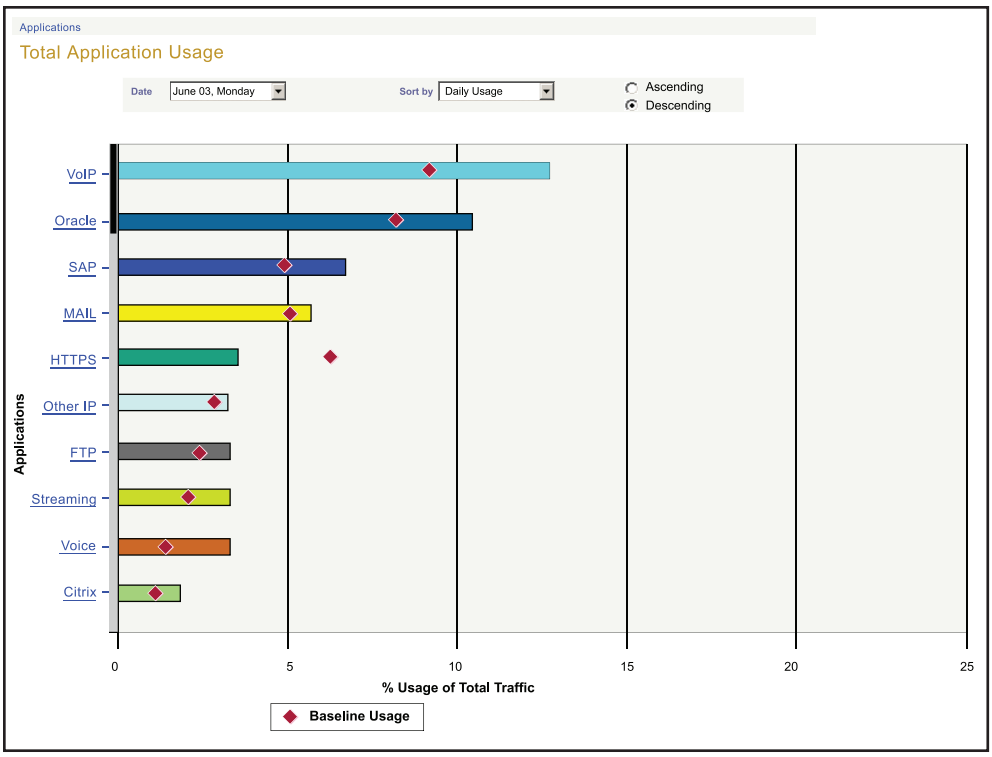

Figure 1 – Auto-discovery of applicants network wide eliminates the guesswork

Once users prioritize applications, they determine if each has been properly configured. Two potential pitfalls commonly occur during this process: IP header configurations are set incorrectly and applications are placed in the wrong class.

The main issue for many enterprises that have already migrated to private IP was application misconfiguration. For one enterprise, VoIP was deemed critical and should have received the highest-class setting. However, as the network engineer was configuring the 80 sites, he made a simple mistake when setting the DiffServ for one location and VoIP calls were impacted. Without visibility, identifying this mistake was like finding a needle in the haystack. They believed VoIP was receiving the highest priority, so they began troubleshooting other parameters including CPE.

With Visual Performance Manager, identifying misconfigured applications is extremely fast and easy by measuring utilization for each CoS across every site in the network.

By looking at the site impacted by poor VoIP quality, network managers could quickly identify that a significant amount of traffic is categorized as unknown (see Figure 2). Using deep-packet inspection and the traffic capture feature, the misconfigured application can be identified in a matter of minutes.

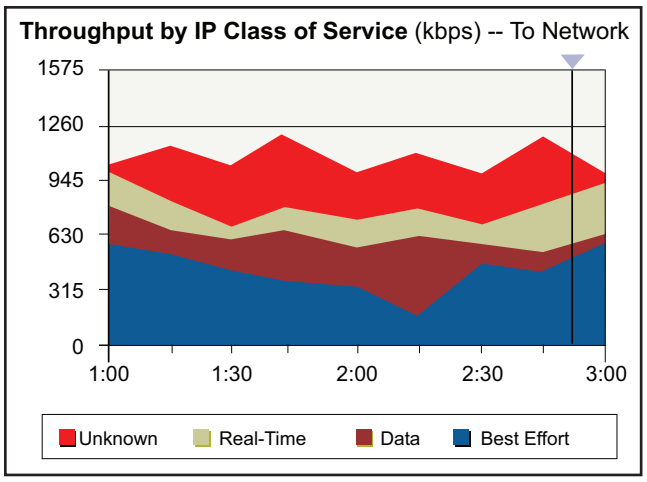

Figure 2 – Measure usage for each CoS setting including unknown/misconfigured applications.

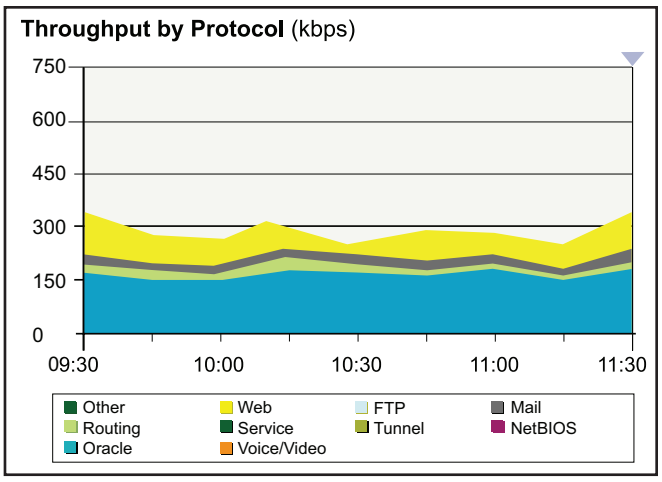

Figure 3 – See which applications are using every class of service

Now that the applications across the entire network have been identified and each IP header configuration was set correctly, you need to verify each application is set correctly, you can see that the gold (highest) setting has applications including Oracle, routing and Web traffic (see Figure 3). Network managers must take steps to ensure applications are assigned correctly. In this example, Web traffic should have been assigned a lower priority but it is consuming valuable high priority resources. Armed with this information, the end-user can correctly assign Web traffic to the bronze priority setting.

How to manage class of service thresholds

Once the applications on the network have been identified, configured, and categorized correctly, it is critical not to exceed the bandwidth thresholds allocated by the service provider. Because of the bursty nature of data traffic, the enterprise might be well below its allocation for the vast majority of the day, but at peak times, it exceeds the threshold. Consequently, mission-critical applications might be negatively impacted due to reclassification or discarding.

Visual Performance Manager monitors utilization for each individual CoS and measures whether usage is above or below threshold allocations. If you are above your threshold allocation, you need to either upgrade your circuit for additional bandwidth or move one or more applications to another CoS (see Figure 4). While there are additional considerations for managing CoS thresholds, they can be easily addressed with Visual Performance Manager.

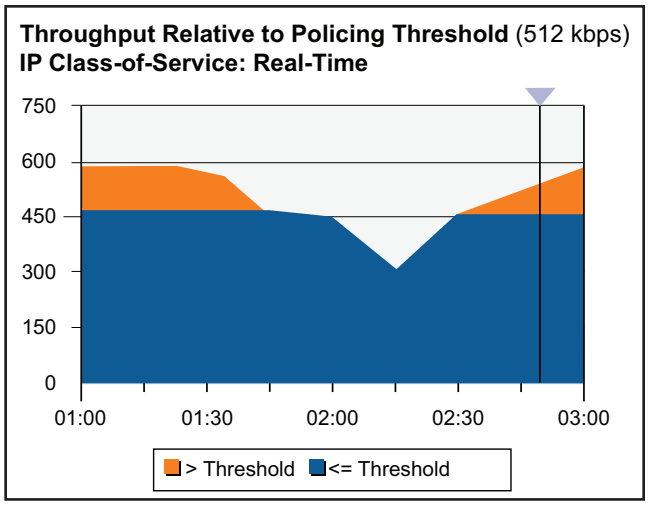

Figure 4 – Determine if your actual usage is exceeding your policing threshold

Measuring service level agreements

Visual Performance Manager has been a leader in SLA verification by providing independent analysis and measurements from an end-to-end perspective. Instead of waiting for a monthly report from your provider, you receive an up-to-the-minute view of service level parameters including availability, delay and throughput. The up-to-the-minute view is critical in isolating poor performance between the network and other aspects. For example, you can quickly determine if a long, round-trip delay is the cause of the application degradation.

As you migrate to private IP, you will implement multiple classes of service. Visual Performance Manager still provides the up-to-the-minute SLA measurements, but also breaks those down by individual CoS settings (see Figure 5). So now you can evaluate if your highest priority setting has better SLA performance than your lowest priority setting.

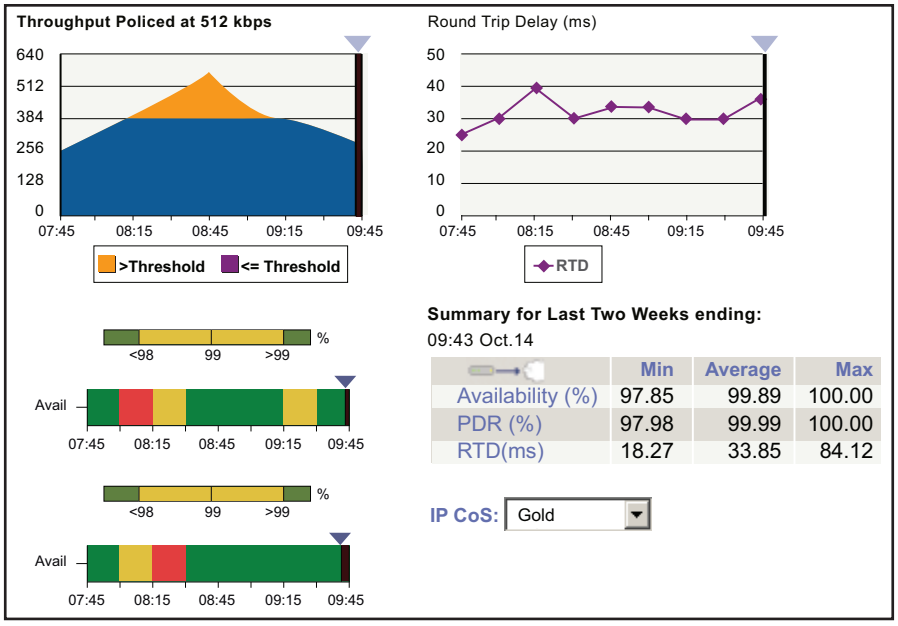

Figure 5 – Up-to-the-minute service level measures for individual CoS levels

Optimizing bandwidth requirements

The cost of bandwidth alone can comprise as much as 60 to 70 percent of total networking budgets for some enterprises.

With tighter resources, optimizing bandwidth is a key step toward maximizing network and application performance. With the new Layer 3 connectivity (IP addressing) instead of Layer 2 connectivity (PVCs), utilization will change in a private IP environment.

Adding CoS capabilities also throws a new wrinkle into the equation.

With the bursty nature of traffic, it is important to understand the impacts of usage on a private IP environment. Burst Advisor measures actual usages in one-second increments in order to properly size your circuits (see Figure 6). Consequently, the networking group has the information to make decisions such as increasing bandwidth for critical locations or moving bandwidth from less-utilized sites to over-utilized sites.

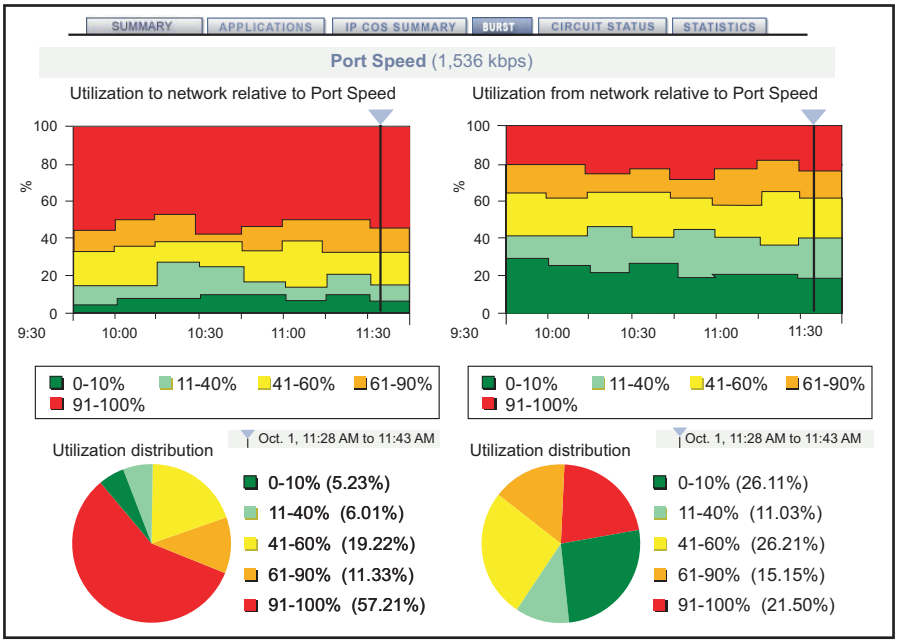

Figure 6 – Burst Advisor measures usage with one-second granularity

Managing throughout the network migration

During a private IP implementation, there will more than likely be a period of time where multiple technologies such as frame relay, ATM and private IP are being run over the network. With a mixed network, it is mandatory to maintain visibility into the performance of the network and applications. Ideally, enterprises would baseline performance based on existing technology and, as the migration progresses, IT staff would see the changes and impact of Layer 3 connectivity and CoS.

Visual Performance Manager offers an unique opportunity to manage the transition to private IP. By managing frame relay, ATM and private IP, enterprises instrument the network with frame or ATM to baseline and then have a single platform manage network application integrity as you migrate to private IP. Visual Performance Manager provides visibility into interworked environments so you can still see end-to-end performance, even if one site is frame relay and the other is private IP (see Figure 7). Now you have the ability to monitor either PVCs in a frame relay environment or the harder challenge of managing IP-subnet pairs in a private IP network.

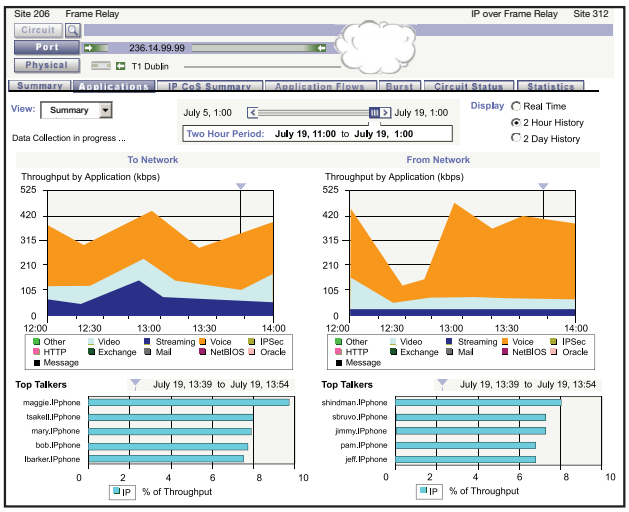

Figure 7 – End-to-end network visibility

Conclusion

As you begin or continue your consideration of MPLS-based/private IP networks for your enterprise, it is extremely important to consider the benefits of migration as well as what will change as you migrate. There are many key reasons or benefits for most enterprises to implement a private IP network, which is a key reason a growing number of enterprises have already implemented an MPLS-based network.

MPLS networks are no longer on the bleeding edge due to their growing adoption. However, there are still key hurdles that must be addressed ranging from CoS prioritization to understanding the change from PVC-based connectivity to IP subnet-based connectivity. The good news is there are tools available to bridge the challenge of migration to a private IP network.

Visual Performance Manager integrates the analysis of application performance with analysis of network performance – so IT departments can take a holistic approach to network application integrity. Now enterprises can understand what applications are on the network so they can be assigned correctly as well as monitor and manage the more complex IP subnet-to-IP subnet connectivity between each and every location in the network.

MPLS-based private IP networks provide many benefits for enterprises ranging from increased performance to less complexity to lower costs, but this also changes the infrastructure IT managers have administered for years. Visual Performance Manager leverages its history in traditional network performance management with enhanced private IP and application functionality so IT managers can confidently make the decision to migrate to an MPLS-based network if it benefits the enterprise.

Tequipment.NET USA
205 Westwood Ave
Long Branch, NJ 07740
1-877-742-TEST (8378)
Fax: (732) 222-7088
salesteam@Tequipment.NET

2006 Fluke Corporation. All rights reserved.
inted in U.S.A. 9/2007 2750400 D-ENG-N Rev B