

The Compelling ROI of Adaptive Private Networking



By Jim Metzler, Cofounder, Webtorials Editorial/Analyst Division

Executive Summary

Adaptive Private Networking (APN) is an emerging way to create virtual WANs. When compared to traditional WAN services such as Frame Relay and MPLS, APNs provide both higher levels of reliability¹, dramatic cost savings, and significantly more bandwidth. An APN is based on packet-by-packet, real-time traffic engineering that leverages the reliability and bandwidth of multiple active paths through the Internet. The reliability improvement that APN delivers by leveraging multiple active paths allows an APN to exploit the superior price/performance of consumer-oriented ISP services.

APNs offer the following benefits:

- 30 - 100 times more bandwidth per dollar than Frame Relay or MPLS services
- 40% - 90% monthly WAN expense reduction
- Payback in 4 - 9 months when migrating completely from Frame Relay or MPLS services
- Payback in 8 - 12 months for partial migrations
- End-to-End QoS functionality
- Greater availability and reliability than typical Frame Relay or MPLS services

To demonstrate the benefits that APN offers, this white paper analyzes how medium to large-sized companies can save in excess of three million dollars in WAN costs over a three-year period while increasing available bandwidth to its branch offices by up to an order of magnitude.

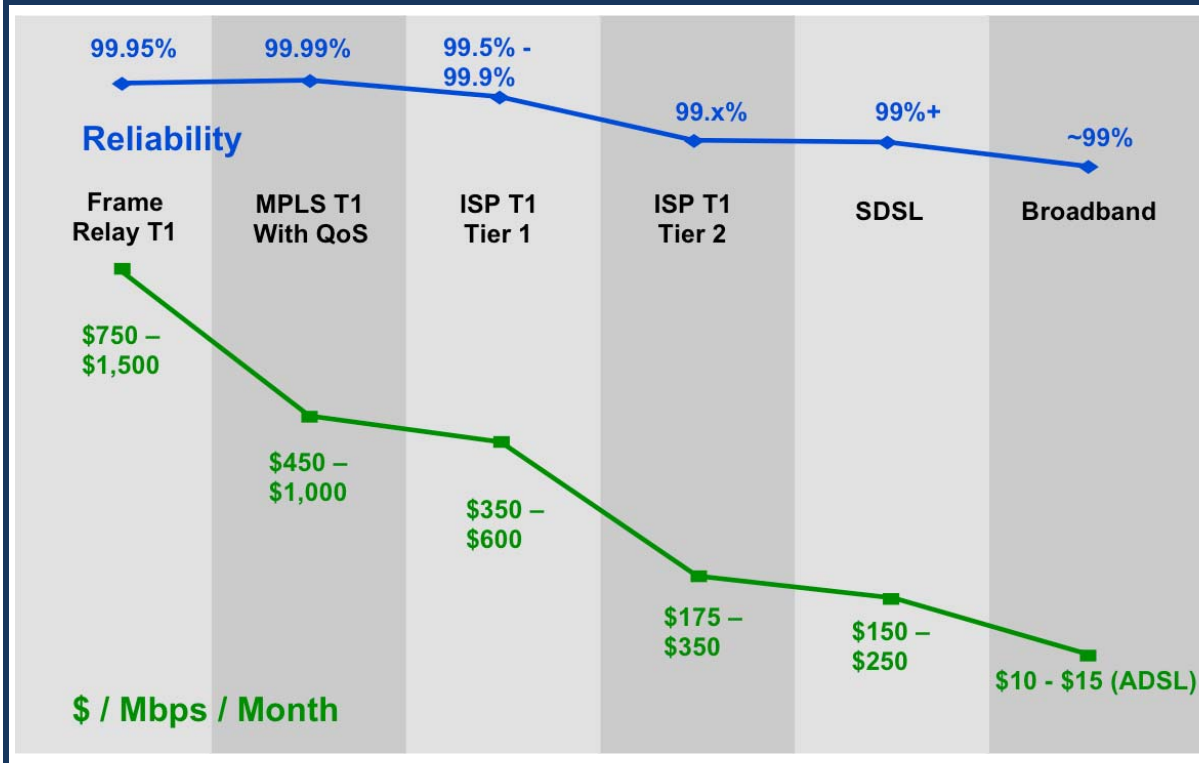
Introduction

The combination of technological advances, Moore's Law and a competitive marketplace have greatly reduced the price that IT organizations pay for most of the major components of IT including processing, memory, data storage, and LAN bandwidth. Unfortunately those factors have had little impact on the price of traditional private WAN services, such as the Frame Relay and MPLS. However, as shown in [Figure 1](#), unlike traditional WAN services, the cost of high-speed Internet connections has decreased rapidly due to both a wider set of technology choices and a more competitive environment.

The low cost of Internet access has made Internet VPNs using broadband connections, such as those that can be implemented with IPsec routers or gateways, a very attractive economic alternative to traditional WAN services. However, quality issues in the public Internet and in consumer-class ISP

¹ Throughout this white paper, the term reliability will refer to the combination of availability, packet loss, latency and jitter.

Figure 1: The Reliability and Cost of WAN Bandwidth



services generally prevent Internet VPNs from meeting the reliability standards of enterprise IT departments. As a result, Internet VPNs are most often used only as a backup connection to a primary

private WAN circuit. This is unfortunate because the shortfall in quality is fairly small when compared to the dramatic cost savings and additional bandwidth that can be realized by using broadband connections. In cases where Internet-based VPNs are deployed today, businesses typically prefer an expensive T1 for access, since a single xDSL connection often results in insufficient upstream bandwidth and a higher Mean Time to Repair (MTTR).

The potential for dramatic WAN savings has begun to attract the attention of CIOs. One such CIO is Geir Ramleth of Bechtel. In a recent article², Ramleth stated that because of the need to become significantly more efficient he intends to fundamentally change how Bechtel provides IT services. One of his goals is to get the cost that he pays for WAN services to be in line with the WAN costs of industry leaders such as YouTube. Ramleth estimates that YouTube spends between \$10 and \$15 per megabit/second for WAN bandwidth, while Bechtel is spending \$500 per megabit/second for its WAN bandwidth.

The goal of this white paper is to describe Adaptive Private Networking, an emerging technology for implementing Internet-based virtualized WANs that can either augment or replace traditional private WANs. APN greatly increases the reliability of Internet-based connections and enables IT organizations to achieve the kind of WAN savings envisioned by Ramleth. APN achieves these goals by leveraging the superior price/performance that is associated with consumer-class Internet services at remote offices, as well as competitively priced business-class Internet access at data center sites. This results in recurring monthly WAN cost savings in the range of 40% to 90%.

² The Google-ization of Bechtel, Carolyn Duffy Marsan, Network World, October 28, 2008

Traditional WAN Services

The emergence of the Internet in the early to mid 1990s was the impetus for most telecommunications carriers to begin a gradual transition away from a time division multiplexing (TDM) based infrastructure and towards a packet switched infrastructure focused on the Internet protocol suite. One of the advantages of packet switching is its ability to emulate dedicated private leased line data networks with virtual networks that are implemented over a shared infrastructure. While leased lines will continue to be used well into the future for certain applications, most enterprise WANs are increasingly dependent on some form of virtual WAN service for connecting remote sites to consolidated data centers.

The rest of this section of the white paper will describe the primary virtual WAN services currently in use.

Frame Relay

The Frame Relay services that were first deployed in the early 1990s were the natural extension of the X.25 services that were first deployed in the early 1980s. By that is meant that Frame Relay is very similar to X.25, with the primary differences being that Frame Relay operates at higher speeds in part because Frame Relay does less error checking than X.25 does.

X.25 and Frame Relay were the first significant services introduced into the virtual WAN market. Within an X.25 or Frame Relay service, traffic from multiple customers is packet switched over a shared and typically over-subscribed network of X.25/Frame Relay switches. The resultant relatively low infrastructure costs per unit of bandwidth allow carriers to offer these services at lower prices than comparable bandwidth on dedicated leased lines.

Frame Relay customers purchase services based on a desired Committed Information Rate (CIR), which is the rate the carrier guarantees will be delivered without dropped packets. Although it is typically possible for the customer to send traffic bursts up to the access line rate, any traffic in excess of the CIR is subject to packet discard when the network is congested. For traffic within the CIR, Frame Relay typically delivers error free packets 99.95% of the time.

One of the weaknesses of Frame Relay is that it does not inherently support Quality of Service (QoS) or packet prioritization well. This makes Frame Relay a sub-optimal WAN service for enterprises that want to use a converged enterprise network to support either real-time traffic (e.g., VoIP and video conferencing) or a mix of traffic - some of which is both business critical and delay sensitive and some of which is not. As indicated in Figure 1, the current cost per Mbps of Frame Relay services is typically in the range of \$750-\$1,500 per month. Referring back to Bechtel, the unit cost of Frame Relay services is approximately fifty to one hundred and fifty times higher than the goal for the unit cost of WAN bandwidth that was established by Ramleth.

MPLS VPNs

Most large carriers are deploying MPLS as a unifying network core technology that can support both legacy Layer 2 subscriber services (Frame Relay and ATM) and emerging packet-based services both at Layer 2 and at Layer 3. MPLS services offer excellent support for meshed topologies, making these services well suited for VoIP, videoconferencing, and collaborative applications requiring any-to-any connectivity.

The major suppliers of MPLS services offer a number of different classes of service (CoS) designed to meet the QoS requirements of different types of applications. Real-time applications are typically placed in what is often referred to as a Differentiated Services Code Point (DSCP) Expedited Forwarding class that offers minimal latency, jitter, and packet loss. Mission critical business applications are typically relegated to what is often referred to as a DSCP Assured Forwarding Class. Each class of service is typically associated with a service level agreement (SLA) that specifies contracted ranges of availability, latency, jitter, and packet loss. The carriers' ability to offer disparate classes of service is supported both by configuring the QoS functionality of the carrier's MPLS routers and by traffic engineering to ensure that the Label Switched Paths (LSPs) that carry high priority traffic have the appropriate characteristics to meet the targeted service level.

As noted, carriers provide an SLA for their MPLS services. Unfortunately, in most cases the SLAs are weak. In particular, it is customary to have the SLAs be reactive in focus; i.e., the computation of an outage begins when the customer opens a trouble ticket. In most cases, the carrier's SLA metrics are calculated as network-wide averages rather than for a specific customer site. As a result, it is possible for a company's data center to receive notably poor service in spite of the fact that the network-wide SLA metrics remain within agreed bounds. In addition, the typical level of compensation for violation of service level agreements is quite modest.

Pricing for MPLS Layer 3 services typically includes separate charges for the local access circuit, the speed of the port of the ingress label switch router, the CoS profile selected, and any advanced services (e.g., Internet access, multicast, network firewalls, etc.) that are part of the service. With careful network design and traffic engineering, MPLS services can deliver both better reliability (up to 99.99%) and somewhat lower cost (between \$450 and \$1,000 per Mbps per month domestically and as high as \$2,500 per Mbps per month internationally) than Frame Relay services. In order to get the best pricing for MPLS services, subscribers are required to sign long term contracts – typically three years. Referring again to Bechtel, the unit cost of MPLS services is approximately thirty to one hundred times higher than the goal for the unit cost of WAN bandwidth that was established by Ramleth.

IPsec VPNs

IPsec is a suite of protocols that provides for the encryption and authentication of IP traffic. By using IPsec in tunnel mode, enterprise routers or gateways can create a VPN by tunneling IP traffic through the public Internet. The tunnel is created by encrypting the original packet and encapsulating it within a new packet whose source and destination are the IPsec routers. Compared to a Frame Relay or MPLS VPN service, an IPsec VPN has a relatively low monthly recurring cost that is comprised entirely of Internet access fees paid to the ISPs serving each enterprise site participating in the VPN. Based on Internet access via a T1 local circuit from a Tier 1 ISP, the cost of an IPsec VPN connection is in the range of \$300-\$600 per Mbps per month, while a similar service from a Tier 2 ISP is typically in the range of \$175-\$350 per Mbps per month. The reliability of an IPsec VPN is fairly high. It can be approximately 99.5% for a Tier 1 ISP and approximately 99% for a Tier 2 ISP. Even though IPsec VPN services leverage the low cost of the Internet, the cost of these T1-based services is still approximately twelve to sixty times higher than the goal for the unit cost of WAN bandwidth that was established by Ramleth.

One of the major disadvantages of an IPsec VPN is that the Internet does not support any form of QoS. As a result, IPsec VPNs don't always provide acceptable levels of latency, jitter, and packet loss for interactive enterprise applications or real-time applications. Part of the quality issue is due to the fact that packets are routed based on what BGP (Border Gateway Protocol) determines to be the shortest path

regardless of whether the traffic traversing that path is encountering high levels of congestion resulting in unacceptable levels of latency, packet loss or jitter.

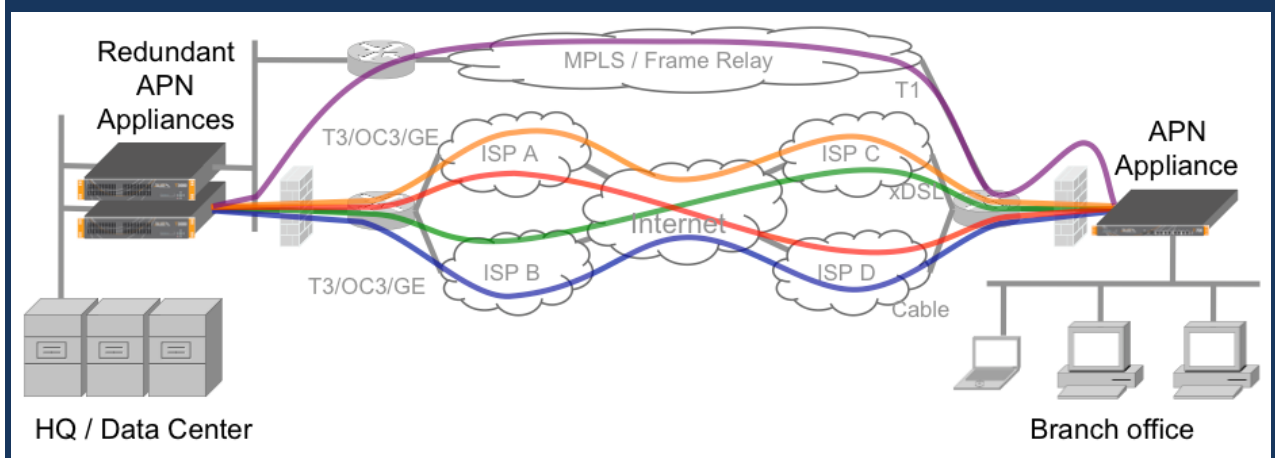
IPsec VPNs can also utilize consumer-grade Internet access using cable and/or xDSL. While these access technologies are high speed, and very low cost (\$3-\$15 per Mbps per month), the reliability is generally regarded as too low to meet enterprise quality standards for connecting branch offices to central data centers. In addition, a single such connection will frequently offer insufficient upstream bandwidth, and unacceptable MTTR. As described in the next section of this white paper, APN is an emerging virtual networking technology that solves the reliability and performance issues that are associated with the public Internet and commodity ISP access networks by using instantaneous traffic engineering to distribute traffic across a number of parallel Internet paths.

Adaptive Private Networking

One of the key concepts that underlies APN is the concept of adding intelligence to a multitude of consumer grade network connections and hence creating a low cost, highly reliable enterprise class product or service. This basic concept has been used successfully in other aspects of IT, most notably storage. For example, RAID (Redundant Array of Inexpensive Disks) arrays are created by wrapping a layer of intelligence around high volume consumer grade hard disk drives. RAID arrays have become popular because they are inexpensive, highly reliable and highly scalable.

As shown in **Figure 2**, APN is a two-ended solution and requires that appliances be placed physically or logically in-line with the Internet access routers at every site that participates in the APN. In the figure, both the data center site and the remote office sites access the Internet via two ISPs each. As a result, there are four distinct paths between the pair of sites over the Internet. As is also shown in Figure 2, it is possible to have a fifth path between the two sites using a traditional WAN service such as Frame Relay or MPLS. The appliances encrypt and encapsulate each packet traversing the Internet. Encryption is optional for packets sent over existing private WAN connections. The encapsulation header includes a timestamp, sequence number and an IP address of the destination APN appliance that corresponds to the far end ISP that will be used by this particular packet.

Figure 2: The Adaptive Private Network



The APN appliances use a combination of the capabilities described below to make multiple parallel network paths function as a single high bandwidth connection with very high reliability; e.g., high availability, and low packet loss, latency, and jitter, and thus predictable application performance.

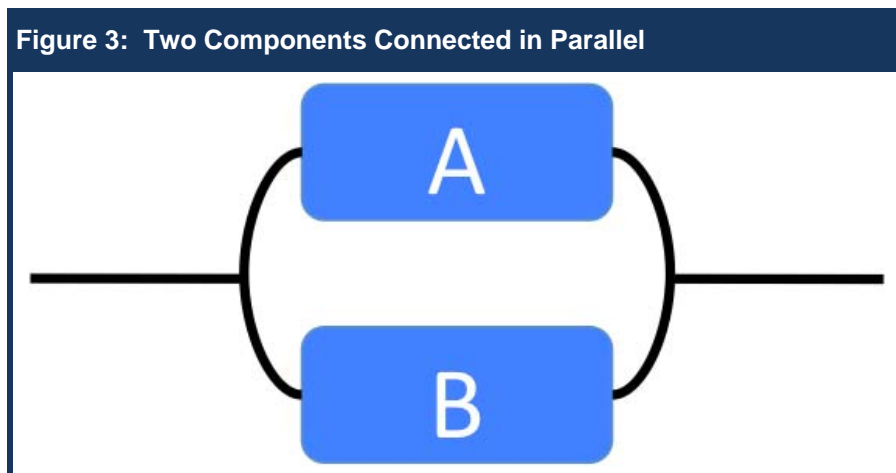
Real-time Performance Monitoring: The APN appliances continually monitor each of the parallel network paths on a packet-by-packet basis for loss, latency, and jitter. Data on these metrics is derived from the headers that encapsulate the user traffic. As a result, the monitoring does not consume any extra bandwidth and the performance characterization always reflects the current state of each path. The real-time nature of the monitoring allows even short-term congestion events to be detected, allowing subsequent packets to be re-routed around the congested path until the event is cleared.

Application Awareness: An APN appliance can recognize a wide range of application packet flows and can be configured to optimally leverage the network fabric differently for bulk data transfers, interactive traffic and real-time communication.

Resilient Multipath Connectivity: APN appliances multiplex traffic onto multiple paths for both aggregating bandwidth and providing enterprise class reliability in the face of outages, delay and packet loss. They perform stateful traffic steering based on the adaptive path selection algorithms described below. When faced with a transient loss of packets, Resilient Multipath Connectivity quickly retransmits lost packets and hence avoids the performance penalty that is associated with TCP's Slow Start and Congestion Avoidance algorithms. If the loss is not transient, Resilient Multipath Connectivity directs all traffic off of the troubled link and onto high quality links until the problem is resolved. When supporting real-time applications such as VoIP, traditional videoconferencing and telepresence, Resilient Multipath Connectivity chooses the path with the lowest delay, packet loss and jitter.

Adaptive Path Selection: Adaptive path selection algorithms are employed to make instantaneous selections of the best path for each application type on a packet-by-packet basis. If there is a failure or congestion in one of the paths, traffic can be re-directed in as little as a few hundred milliseconds. Application-aware adaptive path selection among multiple paths provides a form of *virtual QoS* for the public Internet. The path selection algorithms also provide load distribution across all paths to ensure that the maximum advantage is taken of all the available bandwidth.

Because of adaptive path selection, the availability of an APN that consists of multiple parallel paths is very high even if the availability of each component path is only moderately high. For example, [Figure 3](#) depicts a system that is composed of two components that are connected in parallel.




The system depicted in **Figure 3** is available unless both of the two components are unavailable. Assuming that each component is a diversely routed DSL or cable access line and that one of the access lines has an availability of 99% and the other has an availability of 98%, then the system has an availability of 99.98%. Alternatively, if both access lines have an availability of 99%, then the system is available 99.99% of the time. This level of availability is equal to or exceeds the availability of single vendor Frame Relay or MPLS networks.

The APN functionality described above, in conjunction with parallel paths, can provide similar benefits in terms of protecting traffic from excessive congestion-related delay, jitter and loss. In particular, the delay and jitter are reduced for real-time or interactive applications, and bandwidth is increased and loss reduced for bulk data transfers.

Network Fabric Optimization: The APN appliance can also include network fabric optimization functionality. For example, the APN device can perform packet re-ordering to ensure that packets are delivered to the end system in the correct sequence. For data applications, the APN appliance can detect a packet loss and immediately request a re-transmission by issuing a negative acknowledgement to its partner appliance. The lost packet can be recovered and placed into sequence, making it appear to the end systems that a lossless network connects them. For VoIP, APN can detect and move packets off paths experiencing excessive loss or jitter within a fraction of a second. For even greater VoIP quality, APN can leverage the ample bandwidth to minimize latency by transmitting redundant VoIP packets over multiple paths. The egress APN appliance will insert the first copy of each VoIP packet it receives into the voice stream delivered to the IP phone, discarding the duplicate packets that have encountered higher latency.

Figure 4 shows the unique value proposition of APN vs. the more traditional types of virtual WANs in terms of reliability and the price per Mbps per month. The latter metric for APN reflects the very low costs

Figure 4: Comparison of VPN Technologies in Terms of Reliability and Price/Performance



of cable and xDSL Internet access that can be leveraged at remote sites, which account for the majority of WAN access costs for most enterprise branch office networks. The access costs are low enough to allow two or more Internet connections at each remote site via different ISPs, as well as dual high-speed Internet connections via different ISPs at the data center site. The cost savings are also significant enough in comparison with the WAN alternatives to allow significant additional bandwidth to be provisioned on each link and at each site.

WAN Optimization and Adaptive Private Networking

The phrase *WAN Optimization* refers to an extensive set of techniques that organizations have deployed in an attempt to ensure both acceptable costs and acceptable application performance. Two of the primary roles that these techniques play are to:

- Reduce the amount of data that must be transmitted over a WAN;
- Mitigate the impact of inefficient protocols and applications;

One of the principal categories of WAN Optimization products focuses on improving the cost and performance of delivering applications from centralized data centers to branch office employees over networks technologies such as Frame Relay and MPLS. Products of this type are typically referred to as WAN Optimization Controllers (WOCs) and are produced by numerous vendors including Riverbed, Cisco, Blue Coat, Juniper, F5 and Citrix. WOCs are often referred to as *symmetric solutions* because they require functionality in both the data center as well as the branch office.

As noted, one of the goals of a WOC is to mitigate the need for an expensive WAN upgrade by reducing the amount of data that must be transmitted over the WAN. To support this goal, WOCs implement myriad technologies including caching, compression, and de-duplication. In many, but not all, instances these techniques can improve application performance. While there are exceptions, in most cases a WOC reduces the amount of data that must be sent over a wide area network by a factor that typically ranges between two and four. Relative to the goal of mitigating the need for an expensive WAN upgrade, WOCs and APN are competitive solutions. The advantage of APN is that it provides thirty to one hundred times the bandwidth per dollar when compared to traditional Frame Relay and MPLS services.

Another goal of a WOC is to mitigate the impact of chatty protocols and applications by implementing techniques such as request prediction and spoofing. A chatty protocol such as Common Internet File System (CIFS) requires hundreds of round trips to complete a transaction. For example, a chatty protocol could require two hundred round trips to complete a transaction. If the round trip WAN delay is seventy five milliseconds, the round trips would take fifteen seconds. Adding more bandwidth will not reduce the amount of time that the round trips require. As such, WOCs and APNs are complementary solutions in those instances in which a WOC is optimizing the performance of chatty protocols and applications.

A Return on Investment (ROI) Analysis

This section will present two case studies to demonstrate the dramatic cost savings that are associated with migrating, in part or completely, from traditional WAN services to APN. To exemplify the factors that drive the cost effectiveness of APN, the first case study will be presented in detail. However, in order to not unduly increase the length of this white paper, the second case study will be presented in summary fashion.

There are numerous metrics that can be used to measure the financial viability of deploying a new network technology such as APN. One of the most useful metrics is the payback period, which is the amount of time before the resultant savings equals or exceeds the cost of deploying a new technology or service. As was previously mentioned, migrating to APN from traditional WAN services significantly improves the price/performance ratio without sacrificing service quality. As will be demonstrated in this section, the payback period for a full migration to APN is typically between three and eight months. A more conservative approach to adopting APNs is to leave some Frame Relay or MPLS service in place at selected sites as additional parallel paths within the overall network. In this scenario, payback is typically achieved within eight months to a year.

Another useful financial metric is the internal rate of return (IRR)³. According to Wikipedia⁴, the internal rate of return is the annualized effective compounded return rate that can be earned on the invested capital. One way to look at the IRR metric is to consider an IT organization that invested \$100,000 in new technology and after three years the use of that new technology produced a hard savings⁵ of \$200,000. The IRR metric is the answer to the question “If the company had invested that \$100,000 in the bank, what annual rate of return would the bank had to have given them for the three years, in order for their \$100,000 investment to grow to \$200,000”. In this particular case, the answer to that question is 26%.

Large High-Tech Company

This case study involves a multi-national high-tech company that has headquarters in California, annual revenues of approximately two billion dollars and currently interconnects most of its sites using MPLS. Throughout the rest of the white paper, this company will be referred to as The Large High-Tech Company. **Appendix A** contains details about The Large High-Tech Company’s current MPLS network including the fact that the company has seventy-four sites located in Europe, the U.S., the Asia Pacific region, and India.

As is also shown in **Appendix A**, The Large High-Tech Company’s:

- Cost for MPLS service per location ranges from a low of \$133 per Megabit/second/month to a high of \$1,250 per Megabit/second/month
- Monthly recurring cost for MPLS services is \$216,233.
- Annual recurring cost for MPLS services is \$2,594,796.

The Large High-Tech Company’s three-year cost for its current MPLS network is \$7,784,394. This cost is based on the assumption that The Large High-Tech Company does not add any WAN bandwidth. Since it is unlikely that The Large High-Tech Company will go three years without adding WAN bandwidth, the figure of \$7,784,394 most likely understates the actual three-year costs.

³ More information on how to compute the IRR can be found at <http://www.webtorials.com/abstracts/MetzlerPresentation1.htm>

⁴ http://en.wikipedia.org/wiki/Internal_rate_of_return

⁵ Hard savings refers to a verifiable reduction in spending such as the reduction that results from canceling an MPLS service. This is in contrast to soft savings such as the savings that are associated with increasing the productivity of an organization’s employees. Soft savings, while important, are both harder to measure and more difficult to use as justification for a capital purchase

The Large High-Tech Company performed an ROI analysis relative to deploying an adaptive private network. They evaluated two scenarios that have a number of common elements. For example, both scenarios:

- Assumed a three year time horizon
- Deployed the APN hardware at the beginning of the first year at a cost of \$892,800
- Implemented between two and twelve times as much bandwidth at each site as was present in the status quo MPLS network. As such, the two scenarios provide for growth in WAN capacity that was not provided for in the analysis of the status quo MPLS network.
- Did not have any hardware maintenance costs in the first year and had hardware maintenance costs of \$107,136 in each of the second and third years.
- Had monthly WAN bandwidth costs at the remote offices that ranged from \$13 per Megabit/second to \$40 per Megabit/second.
- Had monthly WAN bandwidth costs at the three data centers that ranged from \$56 per Megabit/second to \$100 per Megabit/second⁶.

One of the two scenarios was conservative in that it called for leaving some MPLS circuits at The Large High-Tech Company's medium and large size sites in Europe, the Asia Pacific region, and India. It also called for leaving some MPLS circuits at The Large High-Tech Company's three data centers. A site-by-site analysis of the conservative APN scenario is contained in [Appendix B](#). A summary of that analysis is contained in [Table 1](#).

⁶ WAN bandwidth costs within a region typically span a range of values. The WAN bandwidth costs for both the remote offices and the data centers used in this white paper are estimates based on the high end of the range for each territory.

Table 1: Three-Year Costs in the Conservative APN Scenario

	Year 1	Year 2	Year 3
Monthly APN Bandwidth Costs	\$88,644	\$88,644	\$88,644
Annual APN Bandwidth Costs	\$1,063,728	\$1,063,728	\$1,063,728
APN Hardware Costs	\$892,800	\$0	\$0
APN Maintenance	\$0	\$107,136	\$107,136
Total Annual APN Cost	\$1,956,528	\$1,170,864	\$1,170,864
Annual Cost of the Current MPLS Network	\$2,594,796	\$2,594,796	\$2,594,796
Cumulative APN Cost	\$1,956,528	\$3,127,392	\$4,298,256
Total MPLS Costs over Time	\$2,594,796	\$5,189,596	\$7,784,394
Percentage Annual Savings: APN vs. MPLS	25%	55%	55%
Cumulative APN Savings vs. Status Quo MPLS Network	\$638,268	\$2,062,204	\$3,486,138

If The Large High-Tech Company implements the conservative APN deployment, they will reduce their monthly bandwidth costs by \$127,589 (\$216,233 - \$88,644). Since the conservative APN deployment requires a capital expenditure of \$892,800, this results in a payback period of seven months (\$892,800 / \$127,589 per month). The IRR associated with the conservative APN is 57%.

The second scenario was more aggressive in that it called for an eventual total replacement of the current MPLS network with APN. A site-by-site analysis of the aggressive APN scenario is contained in [Appendix C](#). A summary of that analysis is contained in [Table 2](#).

Table 2: Three-Year Costs in the Aggressive APN Scenario

	Year 1	Year 2	Year 3
Monthly APN Bandwidth Costs	\$45,744	\$45,744	\$45,744
Annual APN Bandwidth Costs	\$548,928	\$548,928	\$548,928
APN Hardware Costs	\$892,800	\$0	\$0
APN Maintenance	\$0	\$107,136	\$107,136
Total Annual APN Cost	\$1,441,728	\$656,064	\$656,064
Annual Cost of the Current MPLS Network	\$2,594,796	\$2,594,796	\$2,594,796
Cumulative APN Cost	\$1,441,728	\$2,097,792	\$2,753,856
Total MPLS Costs over Time	\$2,594,796	\$5,189,596	\$7,784,394
Percentage Annual Savings: APN vs. MPLS	44%	75%	55%
Cumulative APN Savings vs. Status Quo MPLS Network	\$1.153.068	\$3.091.804	\$5.030.538

If The Large High-Tech Company deploys the aggressive APN, they will reduce their monthly bandwidth costs by \$170,489 (\$216,233 - \$45,744). Since the deployment of the aggressive APN requires a capital expenditure of \$892,800, this results in a payback period of just over five months (\$892,800 /\$170,489 per month). The IRR associated with the aggressive APN is 78%.

The key results of the ROI analysis performed by The High Tech Company are depicted in [Table 3](#).

Table 3: Summary of 3 Year ROI Analysis

	Incremental Investment	Incremental Operating Cost	Three Year WAN Cost	Percentage Savings vs. Status Quo	Payback Period	IRR
Status Quo	\$0	\$0	\$7,784,394	NA	NA	NA
Conservative APN	\$892,800	\$107,136	\$4,298,256	45%	7 Months	57%
Aggressive APN	\$892,800	\$107,136	\$2,753,856	65%	5.2 Months	78%

Medium High-Tech Company

This case study involves a multi-national high-tech company that has headquarters in California and has annual revenues of approximately five hundred million dollars. The company has twenty-eight sites, nineteen of which are connected using MPLS and nine are connected using Internet access. Throughout the rest of the white paper, this company will be referred to as The Medium High-Tech Company.

The Medium High-Tech Company's:

- Cost for MPLS service per location ranges from a low of \$88/Megabit/second/month to a high of \$2,469 /Megabit/second/month.
- Recurring cost for MPLS service is \$111,058 per month and \$1,332,696 per year.
- Recurring cost for Internet access is \$2,485 per month and \$29,820 per year.
- Total recurring bandwidth cost is \$113,543 per month and \$1,362,516 per year.

If nothing changes, The Medium High-Tech Company's cost for bandwidth over the next three years will be \$4,087,548. Similar to the previous case study, The Medium High-Tech Company evaluated two scenarios for upgrading its WAN. Both scenarios:

- Assumed a three year time horizon
- Deployed the APN hardware at the beginning of the first year at a cost of \$285,640
- Implemented between two and eight times as much bandwidth as was present in the status quo MPLS network.
- Did not have any hardware maintenance costs in the first year and had hardware maintenance costs of \$34,277 in each of the second and third years.
- Had monthly WAN bandwidth costs at the remote offices that ranged from \$13 per Megabit/second to \$40 per Megabit/second.
- Had monthly WAN bandwidth costs at the three data centers that ranged from \$56 per Megabit/second to \$100 per Megabit/second.

Similar to the previous case study, one of the scenarios was conservative in that it called for leaving some MPLS circuits in place – primarily in the Asia Pacific and India regions. If The Medium High-Tech Company deploys the conservative APN, they will reduce their monthly bandwidth costs by \$64,911 (\$113,542 - \$48,632). Since the deployment of the conservative APN requires a capital expenditure of \$285,640, this results in a payback period of roughly four and a half months (\$285,640 /\$64,911 per month). The IRR associated with the conservative APN is 91%.

If The Medium High-Tech Company deploys the aggressive APN, they will reduce their monthly bandwidth costs by \$97,329 (\$113,542 - \$16,214). Since the deployment of the aggressive APN requires a capital expenditure of \$285,640, this results in a payback period of roughly three months (\$285,640 /\$97,329 per month). The IRR associated with the conservative APN is 123%.

The financial metrics that are associated with the two scenarios that The Large High-Tech Company evaluated were dramatic, but not as dramatic as the financial metrics that are associated with the two

scenarios that The Medium High-Tech Company evaluated. The reason for that is that many of the sites that The Medium High-Tech Company connects via MPLS are in India where the cost of MPLS is extremely high. For example, one of their sites is in Bangalore where the monthly cost of a 4 Mbps MPLS circuit is \$9,875. This equates to a monthly cost of a megabit/second of bandwidth of \$2,469. Because these costs are so high, the APN savings are even more significant than normal.

Summary

The two case studies demonstrate the dramatic cost savings that are associated with APN. In the status quo MPLS networks the monthly cost for a megabit per second of bandwidth ranged from a low of \$333 to a high of \$2,469. In contrast, with APN the monthly cost for a megabit per second of bandwidth ranged from a low of \$8 to a high of \$100. The payback periods ranged from three to seven months and the IRR ranged between 57% and 123%. Making these cost savings all the more compelling is the fact that the alternative APN designs also provided between two and twelve times the bandwidth than the status quo MPLS networks provided at each site.

In addition to the compelling cost savings, APN is appealing in part because it builds on Internet access technologies that are well understood. It is also appealing because it leverages a well-understood concept: adding intelligence to a multitude of consumer grade products or services and hence creating a low cost, highly reliable enterprise class product or service. This concept has been used successfully in other aspects of IT such as RAID arrays. RAID technology revolutionized storage by enabling production IT business systems to be reliably built using PC economics. RAID was a critical component in the virtualization of storage and it enables not just reliability and cost savings, but storage scalability as well

By combining traditional private WAN reliability with public Internet economics, Adaptive Private Networking technology promises similar benefits for the enterprise WAN. The ability to save money and generate a positive ROI in a short timeframe *alone* makes APN a technology to investigate in these difficult economic times. However, APN technology enables IT organizations to build virtual WANs that exhibit not only the highest levels of reliability and cost savings, but network scalability. It also provides IT organizations with leverage over their telecom suppliers. As such, APN is perhaps ultimately a more compelling technology for IT organizations relative to its *strategic* impact on designing, building, managing and evolving their wide area networks in the next several years.

Appendix A

The Large High-Tech Company Current MPLS Network

	Number of Sites	Bandwidth (Mbps)	Cost/Site	Total Cost	Cost/Mbps
Small Sales Offices	20				
US	10	1.5	\$661	\$6,610	\$441
EMEA	5	2	\$2,000	\$10,000	\$1,000
Asia Pacific	3	2	\$2,500	\$7,500	\$1,250
India	2	2	\$2,500	\$5,000	\$1,250
Medium Sales Offices	15				
US	8	4.5	\$1,983	\$15,864	\$441
EMEA	4	4	\$4,000	\$16,000	\$1,000
Asia Pacific	2	4	\$5,000	\$10,000	\$1,250
India	1	4	\$4,000	\$4,000	\$1,000
Large Sales Offices	5				
US	2	12	\$4,000	\$8,000	\$333
EMEA	1	12	\$6,500	\$6,500	\$542
Asia Pacific	1	10	\$6,000	\$6,000	\$600
India	1	10	\$6,000	\$6,000	\$600
Small Development Offices	15				
US	7	1.5	\$661	\$4,627	\$441
EMEA	3	2	\$2,000	\$6,000	\$1,000
Asia Pacific	2	2	\$2,500	\$5,000	\$1,250
India	3	2	\$2,500	\$7,500	\$1,250
Medium Development Offices	9				
US	4	4.5	\$1,983	\$7,932	\$441
EMEA	2	4	\$4,000	\$8,000	\$1,000
Asia Pacific	1	4	\$5,000	\$5,000	\$1,250
India	2	4	\$4,000	\$8,000	\$1,000
Large Development Offices	6				
US	3	9	\$3,800	\$11,400	\$422
EMEA	1	9	\$6,300	\$6,300	\$700
Asia Pacific	1	8	\$6,000	\$6,000	\$750

India	1	8	\$6,000	\$6,000	\$750
Extra Large Development Offices	4				
US	1	45	\$6,000	\$6,000	\$133
EMEA	1	45	\$8,500	\$8,500	\$189
Asia Pacific	1	45	\$8,500	\$8,500	\$189
India	1	50	\$10,000	\$10,000	\$200
MRC					
ARC				\$2,594,796	

Appendix B
The Large High-Tech Company
Conservative APN Scenario

	Number of Sites	Bandwidth (Mbps)	<i>Cost/Site</i>	Total Cost	Cost/Mbps
Small Sales Offices	20				
US	10	18	\$234	\$2,400	\$13
EMEA	5	18	\$270	\$1,350	\$15
Asia Pacific	3	18	\$144	\$432	\$8
India	2	18	\$720	\$1,440	\$40
Medium Sales Offices	15				
US	8	18	\$234	\$1,872	\$13
EMEA	4	18	\$270	\$1,080	\$15
		2	\$2,000	\$8,000	MPLS*
Asia Pacific	2	18	\$144	\$288	\$8
		2	\$2,500	\$5,000	MPLS
India	1	18	\$720	\$720	\$40
		2	\$2,000	\$2,000	MPLS
Large Sales Offices	5				
US	2	24	\$312	\$624	\$13
EMEA	1	18	\$270	\$270	\$15
		2	\$1,200	\$1,200	MPLS
Asia Pacific	1	18	\$144	\$144	\$8
		2	\$1,200	\$1,200	MPLS
India	1	18	\$720	\$720	\$40
		2	\$2,000	\$2,000	MPLS
Small Development Offices	15				
US	7	18	\$234	\$1,638	\$13
EMEA	3	18	\$270	\$810	\$15
Asia Pacific	2	18	\$144	\$299	\$8
India	3	18	\$720	\$2,160	\$40
Medium Development Offices	9				
US	4	24	\$312	\$1,248	\$13
EMEA	2	24	\$360	\$720	\$15
		2	\$2,000	\$4,000	MPLS

Asia Pacific	1	24	\$192	\$192	\$8
		2	\$2,500	\$2,500	MPLS
India	2	24	\$960	\$1,920	\$40
		2	\$2,000	\$4,000	MPLS
Large Development Offices	6				
US	3	24	\$312	\$936	\$13
EMEA	1	24	\$360	\$360	\$13
		2	\$2,000	\$2,000	MPLS
Asia Pacific	1	24	\$192	\$192	\$13
		2	\$2,500	\$2,500	MPLS
India	1	24	\$960	\$960	\$13
		2	\$2,000	\$2,000	MPLS
Extra Large Development Offices	4				
US	1	90	\$5,040	\$5,040	\$56
EMEA	1	90	\$5,850	\$5,850	\$65
		2	\$2,000	\$2,000	MPLS
Asia Pacific	1	90	\$3,150	\$3,150	\$35
		2	\$2,500	\$2,500	MPLS
India	1	90	\$9,000	\$9,000	\$100
		2	\$2,000	\$2,000	MPLS
MRC				\$88,644	
ARC				\$1,063,728	

* Indicates a backup MPLS circuit

Appendix C
The Large High-Tech Company
Aggressive APN Scenario

	Number of Sites	Bandwidth (Mbps)	Cost/Site	Total Cost	Cost/Mbps
Small Sales Offices	20				
US	10	18	\$234	\$2,400	\$13
EMEA	5	18	\$270	\$1,350	\$15
Asia Pacific	3	18	\$144	\$432	\$8
India	2	18	\$720	\$1,440	\$40
Medium Sales Offices	15				
US	8	18	\$234	\$1,872	\$13
EMEA	4	18	\$270	\$1,080	\$15
Asia Pacific	2	18	\$144	\$288	\$8
India	1	18	\$720	\$720	\$40
		2	\$2,000	\$2,000	MPLS
Large Sales Offices	5				
US	2	24	\$312	\$624	\$13
EMEA	1	18	\$270	\$270	\$15
Asia Pacific	1	18	\$144	\$144	\$8
India	1	18	\$720	\$720	\$40
Small Development Offices	15				
US	7	18	\$234	\$1,638	\$13
EMEA	3	18	\$270	\$810	\$15
Asia Pacific	2	18	\$144	\$299	\$8
India	3	18	\$720	\$2,160	\$40
Medium Development Offices	9				
US	4	24	\$312	\$1,248	\$13
EMEA	2	24	\$360	\$720	\$15
Asia Pacific	1	24	\$192	\$192	\$8
India	2	24	\$960	\$1,920	\$40
Large Development Offices	6				
US	3	24	\$312	\$936	\$13
EMEA	1	24	\$360	\$360	\$13

Asia Pacific	1	24	\$192	\$192	\$13
India	1	24	\$960	\$960	\$13
Extra Large Development Offices	4				
US	1	90	\$5,040	\$5,040	\$56
EMEA	1	90	\$5,850	\$5,850	\$65
Asia Pacific	1	90	\$3,150	\$3,150	\$35
India	1	90	\$9,000	\$9,000	\$100
MRC				\$45,744	
ARC				\$548,928	

A Word From the Sponsor - Talari Networks



With its innovative Adaptive Private Networking (APN) technology, Talari Networks is pioneering a new class of product to deliver a revolutionary approach to building enterprise WANs.

Talari's technology combines diverse, abundant, affordable IP bandwidth sources, transforming broadband links to delivery business-class reliability at consumer prices. APN delivers a network with 30 to 100 times the bits per dollar, ongoing WAN costs reduced by 40% to 90% and greater reliability than single-provider Frame Relay or MPLS services.

Launched in 2007, Talari Networks is headquartered in Cupertino, California with a major product development center in Raleigh, North Carolina. More information about Talari Networks can be found at <http://www.talari.com>

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Webtorials Briefs

Vol 3, Number 6

Published by Webtorials Editorial/Analyst Division

www.Webtorials.com

Division Cofounders:

Jim Metzler

jim@webtorials.com

Steven Taylor

taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Copyright © 2009, Webtorials

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor.

The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.