Adaptive Private Networking

By John Bartlett
September 2007

Enterprise networking is constrained by the high cost of Wide Area Networking access. Telecommunications companies own the last-mile links and are taking advantage of their position to keep costs high. Consumer broadband is much more competitive, and the resulting price differential is substantial. A T1 MPLS link may cost $1000/month while a DSL link only costs $30 to $90/month. But the lower reliability of the broadband connection usually disuades enterprises from depending on this link for business critical applications.

Adaptive Private Networking (APN) is a new technology that combines the reliability and performance of multiple WAN links into a single, high quality and highly reliable connection. By deploying APN technology in the headquarters and branch office, an enterprise can reduce WAN costs, increase link reliability and increase available bandwith, all at the same time.

Enterprises are Constrained by High WAN Cost

Today's business is on the network. Today's productive enterprise depends heavily on its network to support the applications that run the business. Over the network we share data constantly – with employees, partners and customers – and we execute daily business transactions.

But when the application leaves the enterprise campus and crosses the wide area network (WAN), it is constrained by limited bandwidth. WAN access links are predominantly controlled by a few telco providers, who are keeping costs high.

At the same time, enterprises are distributing their offices and workers across the country and the globe. To make these distributed workers productive, business managers want to run their applications across the whole enterprise network, integrating the work of all employees. The convergence of voice and video conferencing onto this same network further integrates remote workers and adds additional demands to the WAN. Server consolidation, driven by both IT productivity and regulatory concerns, means all those remote offices need to be supported from a centralized server farm. This combination of trends is causing WAN links to be strained, and application performance to suffer. Enterprises are caught between the business goals of productivity and geographic diversity, and the high cost of WAN access links.

Far less expensive WAN access is being offered to consumers and small businesses in the form of DSL and Cable-Modem broadband Internet access. Some companies are using broadband access either as a back-up link or in some cases as their sole connection to the central office. The difficulty with this solution is that commercial broadband does not provide the quality level required for business in terms of availability, congestion events, and packet loss. So existing solutions leave the enterprise in a dilemma.

Adaptive Private Networking (APN) is a new technology that breaks this log jam by combining the availability of multiple low cost network links into a single, high bandwidth, high availability, low loss and low cost WAN connection.

Achieving Low Cost *and* High Quality

To achieve low cost and high quality in manufacturing, we use a strategy of multiple sourcing. Whenever possible, each component of our manufactured product is sourced from at least two component providers. Whether it be a screw, a fan assembly, a

NetForecast Report 5090

microprocessor or a jet engine, we prefer to have at least two sources available so either supplier can satisfy our manufacturing needs.

The benefit of this approach is that we are assured of a steady supply of the product, we get good pricing because of the competitive environment, and we have the flexibility to purchase more product if demand suddenly increases. Let us see how we can apply this principle to network access.

Providing a Second Network Path

If we have a branch office currently connected to the central office via Frame Relay, ATM or MPLS, we have a single source for our network connection between these offices. If congestion, packet loss, excessive jitter or even link failure occurs, our communications are slowed or interrupted. Because we have no other path to the branch office, we are at the mercy of our single supplier.

If we add a second path between offices, however, we have created a second source for our network needs (see Figure 1).

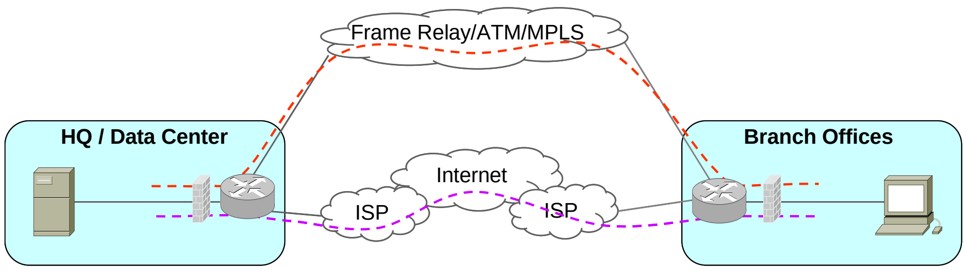

Figure 1 - Providing a second link to the branch office

This is nothing new; this is standard fault tolerant design. Many enterprises have a back-up link in place to ensure connectivity. Some use ISDN dial-up as a backup, some use a DSL or ADSL line, and some have a complete second Frame / ATM / MPLS link to ensure survivability of the connection. The routers know there is a back-up link, and if the primary link fails, traffic will be re-routed to the backup link. Careful programming of the router can even take advantage of some of the available bandwidth in the back-up link, but this gets tricky, especially when a link failure occurs.

So by adding a link in figure 1 we have decreased the chances of being without connectivity due to link failure. We have not increased our bandwidth, reduced our costs or helped applications that are slow due to congestion events on our primary connection.

Adaptive Private Networking

Adaptive Private Networking (APN) adds a new class of appliance into the network. This is a dual-ended solution, so one unit resides at the central office, and the second unit is deployed in the branch. All traffic flowing between the two sites flows through the APN appliance.

The job of the APN device is to optimize the use of all available paths between the HQ/datacenter and branch offices. The APN monitors the characteristics of each path constantly, measuring transit delay (latency), jitter, packet loss and available bandwidth. The APN then directs traffic to flow over the link that best serves that traffic. Forwarding decisions can be made on a packet-by-packet basis, allowing very fast switching of traffic when a momentary congestion event or failure occurs.

During the vast majority of time both lines will be available and working well. When this is true, the bandwidth of both lines will be used, and the enterprise gets the advantage of this additional bandwidth. When congestion occurs on one line, traffic is shunted to the other line until the congestion passes.

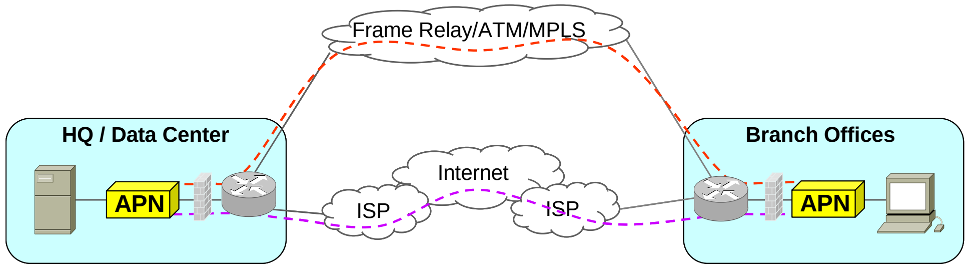

Figure 2 - APN Supporting Two Network Paths

Data applications require high availability, moderately low packet loss, and low latency. Voice and video require very low packet loss and jitter, but can sustain a higher latency if it is not too great. These differences are accommodated by the APN device. Traffic is sent across the link with the best characteristics for each data type. Rather than forcing the application to accommodate the network, the APN adapts the network to provide high quality transport for distributed applications.

The APN concept can be extended to additional paths, as shown in Figure 3. Here the branch office and the central office have contracted with two Internet service providers (ISPs). With two ISPs at each end, four paths across the Internet are provided, in addition to the original dedicated connection. Now the APN has many more choices for branch office connectivity, and the resulting quality and bandwidth are increased.

The complexity of managing multiple network paths is hidden by the APN device. The decisions made about which path is taken by which application, flow or packet are all managed by the APN. The whole connection behaves as a single, high quality, high bandwidth link.

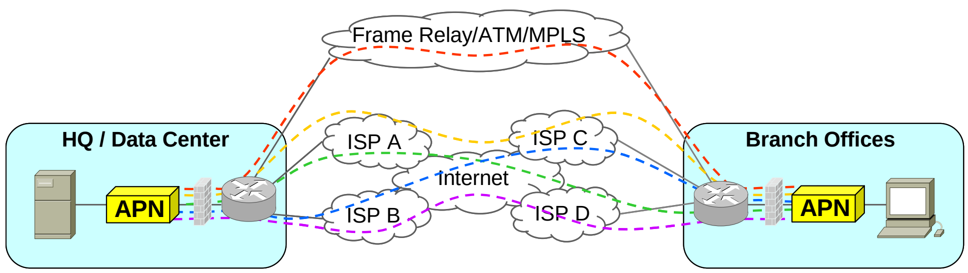

Figure 3 - APN Supporting 5 Network Paths

Benefits of the APN Approach

Combining multiple links with an APN device provides the same advantages we expect from multiple sourcing in the manufacturing environment. Availability is improved because the likelihood of all links failing at the same time is much lower than the failure rate of any individual link. Quality is improved because traffic can use the link providing the highest quality at any point in time. Traffic sensitive to certain types of quality are constantly directed towards those links where the quality is being provided. Bandwidth is increased because multiple links are in place and the bandwidth of all links can be utilized if needed. And application performance is consistent because momentary congestion events are bypassed by using an alternate path.

Cost Reduction

WAN costs may appear to be increased because the method requires the use of multiple links. If an enterprise is already using a back-up link, however, there is no cost increase, just better utilization. And if an enterprise moves solely to Internet connected paths, using inexpensive SDSL or ADSL links, overall WAN cost can be substantially reduced. More conservative enterprises will choose to keep a dedicated link (Frame, ATM or

MPLS) in place, but still may experience cost reductions since not as much bandwidth is required on the dedicated link if additional xDSL connections are used in conjunction with it. In either case, the cost per Megabit/second will drop as the lower cost links are added.

Reliability

Increased availability and quality take advantage of the redundant nature of the connections. Figure 4 shows us that although the cost of xDSL connections is much lower than high quality Frame Relay or MPLS links, xDSL quality is only slightly lower, dropping to about 98%.

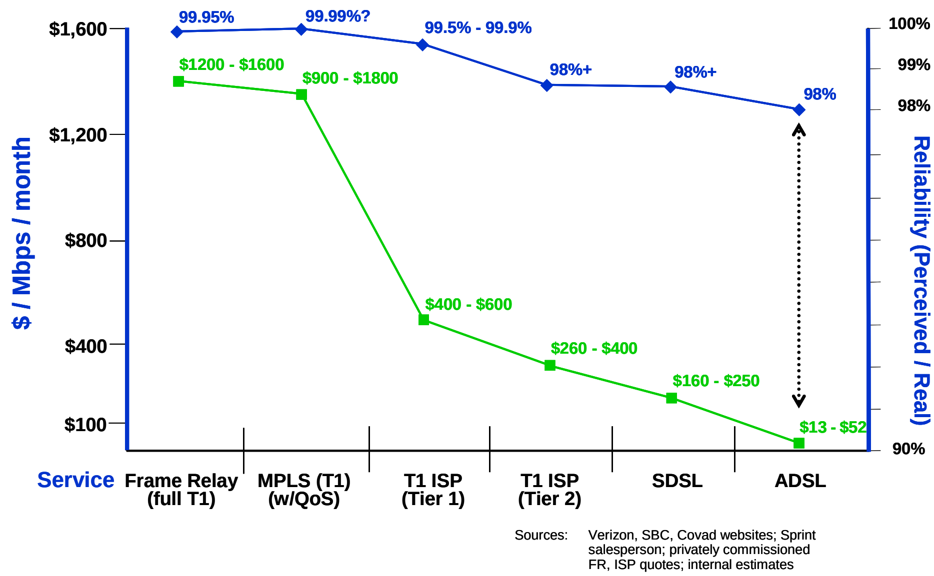

Figure 4 - Cost and Reliability Comparison

The two equations below show how reliability can be calculated as additional links are added.

Reliability for similar links	**Reliability for different links**
$A = 1 - (1 - a)^n$	$A = 1 - (1 - a)(1 - b)(1 - c)$
Where:	Where:
A = overall reliability a = reliability of each link n = number of links	A = overall reliability a = reliability of first link b = reliability of second link c = reliability of third link (etc.)

$1 - (1 - .98)^2 = .9996$ Using two links that have 98% reliability gives an overall reliability of 99.96%. This brings reliability back up to the level of a Frame Relay service. Adding additional links will further improve reliability.

Note that the built-in assumption about increasing reliability with redundant links is that failures on the separate links are independent events. This assumption is only true if the links are obtained from different carriers. If the same carrier is used for more than one link, it is much more likely that a failure can occur on both links at the same time, thus reducing the value of the multiple links from a reliability point of view.

Figure 5 charts reliability as the number of paths between a home office and branch office increase. The vertical axis is the reliability of each individual link, and the horizontal axis is the number of network paths in use. The colored areas of the graph indicate regions where increasing reliability is achieved.

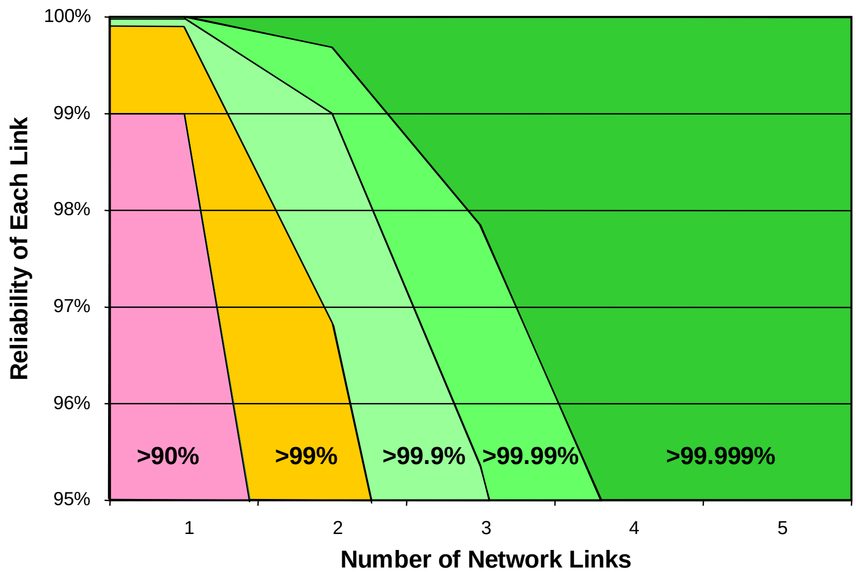

Figure 5 – Reliability Increases as Links are Added

Figure 5 shows that with two links each providing 98% reliability, our combined reliability is greater than 99.9%. I f a third link is added, reliability improves to over five nines (>99.999%). So by using multiple low-cost links, enterprises can increase bandwidth, increase network reliability and reduce costs all at the same time.

How Does it Work?

Adaptive Private Networking appliances sit in-line with the traffic both on the headquarters office and branch office ends, with all traffic between the two locations passing through both boxes. Each packet passing through the APN is encapsulated with a special header before crossing the WAN. This header includes timing and sequence information and a destination IP address. The IP address determines which far end connection this packet will use. For instance, for a packet traveling from the home office to the branch office in Figure 3, the IP address will either direct the packet towards the ISP C connection or the ISP D connection.

The timing and sequence information in the APN header allow the APN at the far end to determine how long it took that packet to cross the network. Sequence information is used to determine if any packets were lost, and to insure that packets arriving on different links are properly sequenced before being forwarded to the LAN.

By monitoring the timing and loss information on each packet, the APN can determine in real-time the quality of each network path. The APN shares this information with its counterpart at the other end of the connection so that rapid decisions can be made about which path to use for any given packet.

Part of the APN task is then to direct traffic to the best link for each type of traffic. Data applications work best on a link that has low latency and low packet loss. Real-time traffic like voice and video conferencing can sustain longer latency, but need to have virtually no packet loss, and are also sensitive to jitter, which is the variation in latency. Because the APN is continuously monitoring the quality of the links, it can redirect traffic in the span of a few hundred milliseconds as the network characteristics change. The APN takes advantage of the available resources on all available paths to insure that the combined network provides a high quality connection for application traffic at all times.

This is a significant advantage over normal IP routing protocols such as RIP, OSPF or BGP. These protocols determine which path to use based on hop count (the number of router hops along the path), but do not take into account the quality of the link. Even if a link is suffering from congestion and packet loss the routing protocol will not switch to another link until congestion is so bad that the routing updates are not forwarded.

Alternatively, typical routers could be configured to load balance across multiple links, but if load balancing is done on a per-packet basis, all applications suffer when one of the links gets congested or exhibits packet loss and out-of-order packets cause performance problems for applications. Meanwhile, per-flow load balancing in the routers would not take full advantage of the bandwidth, and still causes problems for half of the applications whenever loss or latency spikes occur.

Networks that require high availability often use dual links with fail-over switching at the layer 2 level. These backup links can be brought on line in under 50 milliseconds when a failure occurs. Although this scenario works well for link failures, it does not solve the problem of a congested link or one that has increasing packet loss. The switch-over mechanism for redundant links is triggered by a link failure, so increasing packet loss or congestion will not cause the backup link to be used. Furthermore the bandwidth of the backup link is unavailable until the active link fails, so bandwidth is always limited to the capacity of a single link.

Consistent Application Performance

One way to view the job of the APN is that it takes a collection of somewhat unreliable network connections and turns them into a single reliable nearly-lossless high bandwidth WAN connection. This means that applications being run across the wide area network will not have wildly varying performance, but will instead behave consistently throughout the work day. The APN will take advantage of the best links and all the available bandwidth to ensure that applications are well connected to their servers, with low latency and low packet loss at all times. An APN-supported WAN will work the same way at 3:00 PM when demand is high, as it does at 3:00 AM when demand is low.

The APN is not an application acceleration appliance per se. Those devices work to reduce the effect of WAN packet loss and latency on application performance. They use techniques like data compression, TCP enhancement, caching and prefetching to reduce the amount of data crossing the WAN and provide data to remote clients quickly. The APN, on the other hand, works to increase WAN bandwidth and ensure the WAN is highly available and has low loss and predictable latency at all times. Choosing between these technologies will depend on the applications being supported, the distance between the central and remote offices, and on the protocols in use. It is possible to use both technologies together to get an even better result for remote office productivity. with the application acceleration technology providing improved application performance in the average case, while the APN technology provides more bandwidth to enable many more user application flows, as well as improves reliability and application predictability by fixing “worst case” performance issues which would otherwise occur during periods of network stress resulting in poor application performance.

Voice and Video Conferencing, a Special Case

The benefits of using IP-based voice and video conferencing are pushing enterprises to converge these communications applications onto their IP networks. Both voice and video conferencing require a nearly loss-free network to sustain business quality audio and video communications.

This quality requirement is very difficult to achieve over VPN connections riding across the Internet because of the unpredictability of the Internet's quality. No standard way exists today to insure Quality of Service (QoS) is carried between service providers. Peering points between service providers often become the network congestion point where packet loss and jitter occur. So Enterprises wishing to support voice or video across their WANs have traditionally had to rely on expensive dedicated links or more recently on MPLS connections.

Adaptive Private Networking eases this burden by providing a nearly lossless connection across the Internet. Part of this is accomplished by sending real-time traffic down the cleanest link at any given time, as described above. But an APN appliance has one more trick for real-time streams.

If sufficient bandwidth is available (and if configured to do so), the APN appliance will send voice and/or video traffic down multiple links simultaneously. By sending traffic via two separate paths, the probability of at least one copy of each packet arriving is greatly increased. The APN appliance on the receiving end discards duplicates, and forwards a correctly sequenced stream of packets to the LAN. With this approach, the APN is able to tradeoff increased bandwidth use for a higher quality connection. If the cost of the additional bandwidth is low because xDSL lines are being used, this is a very cost effective way of achieving very high quality voice and video communications.

Conclusion

The enormous price gap between broadband connections (xDSL) and the commercial offerings of Frame Relay, ATM and MPLS provides an opportunity to solve the limited bandwidth WAN problem. Using this technology gives the enterprise complete control over the cost / quality / bandwidth tradeoff so the right decision for each business can be implemented. As this technology becomes deployed, businesses will change the way they work. Having additional WAN bandwidth can lead to significant productivity increases, as more decision making gets pushed out to field offices and employees working directly with customers.

Adaptive Private Networking provides a way to break the WAN bottleneck and extend enterprise applications across the WAN. This technology puts control of the WAN quality into the hands of the enterprise. Use of APN will allow geographically disperse enterprises to be tightly integrated, leading to improved productivity and competitiveness. By combining the reliability of disparate WAN connections, and by dynamically directing traffic to the right link at the right time, APN can create a high quality, highly reliable, high bandwidth WAN connection at much lower cost than is offered through traditional sources.

About the Author

John Bartlett John is Vice President at NetForecast, Inc. specializing in real-time application performance and Quality of Service implementation. NetForecast helps enterprise clients understand application performance issues across enterprise networks and the Internet. John can be reached at john@netforecast.com.

NetForecast
Assess and improve the performance of voice, video, and data applications.

Additional information is available at: www.netforecast.com