

Optimizing the Corporate Network with WAN Virtualization



A Webtorials® Thought Leadership Spotlight™



Steven Taylor, Webtorials

WAN Virtualization is a term that commonly refers to the addition of hardware and/or software to enhance WAN performance by adding hardware and/or software to intelligently manage WAN connectivity services in order to provide extremely reliable and responsive service (as one might expect from an MPLS service) at a fraction of the price (as one might expect from an Internet-based service).



Two companies - **Ipanema Technologies** and **Talari Networks** - are clearly the leaders in this market space. Both offer excellent products and provide a superb ROI. And even though the approaches vary somewhat, the end-game is the same.

Since there are no industry-wide standards for interoperability, one or the other must be ultimately chosen for implementation.

Please join my co-founder at the Webtorials Analyst Division - Jim Metzler - and me as we chat with Keith Morris from Talari Networks and Thierry Grenot from Ipanema Technologies to discuss both the common advantages and the competitive advantages of each solution.

We look forward to your participation in the discussion!



Steven Taylor, Webtorials

What types of transport networks do you support (Internet, MPLS, Private Line, Private IP, etc.), and why do you think these are the most important?



Thierry Grenot - CTO - Ipanema Technologies

Hi Steve. Actually Ipanema's Hybrid Network Unification (HNU) does not make assumptions about the transport network. It can then be Internet, MPLS, Frame relay, Ethernet, etc. In most deployed situations, we have a hybrid of MPLS+Internet but we also have pseudo hybrid situations like Internet+Internet and MPLS+MPLS.

Internet is clearly the most important use case, as HNU's main benefit is to turn the public Internet into a business-grade network.



Keith Morris - Talari

We support any kind of IP WAN as well, just as Thierry describes for Ipanema. We can support up to 8 WAN links at any given customer location. To do our reliability magic, we do require diverse WAN providers at each location. But you can mix and match as you please. You could have, e.g., 4 DSL links from the local RBOC plus a T1 from Sprint. Probably half our customers will initially combine an existing MPLS connection with a single Internet connection to get started.



Jim Metzler, Ashton, Metzler & Associates

Steve: I believe that one of the most important aspects of this discussion is the status of WAN services. Starting in the mid 1980s, we went through a 20 year period in which we saw the continued evolution of WAN services. They went from T1/E1 TDM networks to Frame Relay to ATM and now to MPLS. There is no successor to MPLS on the horizon. Given the gestation period that is associated with WAN services, that means there will not be a fundamentally new WAN service for the foreseeable future.

That fact, combined with the fact that the WAN is one of the few components of IT that does not follow Moore's law means that in order to control the cost of WAN services, IT organizations must evaluate alternatives such as those discussed in this thread.



Steven Taylor, Webtorials

What sizes of companies will benefit most from your solution - and why?



Keith Morris - Talari

WAN Virtualization technology is broadly applicable to almost any company with 5 locations or more, since over time all but perhaps the smallest companies need more bandwidth, are looking to reduce monthly costs and would like greater network reliability and application performance predictability

That said, we tend to focus right now on the Fortune 301 - 20,000 - i.e. those companies with 10 - 100 locations. This is our sweet spot at the moment. And the more international locations customers have, the more valuable the technology is, since WAN costs outside of North America are frequently much higher. Just as Riverbed and Peribit showed, when selling a two-ended WAN solution, you really need to prove yourself on the 10-20 site networks before deploying the 50 site networks before deploying 100 site networks, before you can think about doing multi-hundreds. The way you manage the solution today is really optimized around that 10 - 100 site size as well. Over time we'll add the additional scaling and especially management tools to make it easier to deploy multi-hundred site networks.



Thierry Grenot - CTO - Ipanema Technologies

Agreed with Keith. Most companies can benefit by turning the Internet into a business network. Ipanema's HNU particularly targets large networks (from 30 up to 1,000+ sites), but smaller accounts are also accessible through managed services. Both domestic (e.g. retail) and international (e.g. industry) networks use it already.



Steven Taylor, Webtorials

Both companies perform optimization by sending particular traffic types over different networks. For instance, voice might be sent over an MPLS network to ensure low loss and low latency while FTP traffic could be easily relegated to the Internet. How do you determine the traffic type? Do you do your own inspection? Why or why not?



Keith Morris - Talari

Steve, first an important clarification. We don't typically limit any given traffic flow to a single WAN connection. We make per-packet forwarding decisions, not simply per-flow.

This allows us to use all of the available bandwidth even for just a single flow and the overwhelming majority of the time when all the connections are working well. This also means that for delay and jitter-sensitive protocols like RTP or Citrix, we not only put them on the best quality network at flow initiation, we will move the packet flow to a better connection, sub-second, if congestion or link failure causes network quality to get meaningfully worse mid-flow.

Now this said, we do recognize different flows and treat them differently, of course, as does any decent middlebox. We support DSCP and ToS markings, and also support 5-tuple classification (source and destination IP addresses and ports plus IP protocol) to distinguish flows.



Thierry Grenot - CTO - Ipanema Technologies

This is one area where Ipanema and Talari diverge. Ipanema has decided to go with a per-flow decision (preserving natively the packet delivery order) rather than per-packet in order to simplify the deployment of secured environments like stateful firewalls and also to be able to work without an appliance at both ends. Application classification is one of our key techniques, and we use advanced DPI (deep packet inspection) to classify and then control each and every individual flow.



Keith Morris - Talari

Just to ensure no confusion, even though we make per-packet decisions and can and will use multiple connections even for a single flow, thus using all available bandwidth even for a single flow, we too preserve the packet delivery order, delivering packets in order to the receiving host.

We hold packets at the receiving appliance both to avoid the network monitoring nightmare of seeing a lot of out-of-order packets on your LAN, but also because while it's the case that packet loss is the biggest killer of IP application performance, if there is too much out-of-order traffic, TCP's Fast Retransmit algorithm will kick in, reducing window size, and hurt performance that way. Do note, however, that because we know the relative unidirectional latency of each of the different connections between any two locations, unless a packet is lost on the WAN, it's rare that we need to hold up delivery of packets for very long to ensure in-order delivery, because we schedule the packets on each connection to arrive at the proper time.



Steven Taylor, Webtorials

OK. We seem to have found a key difference that our community of technical folks at Webtorials will appreciate. May I ask each of you to summarize with a brief summary of why your solution is "better"?



Keith Morris - Talari

Steve, there are two basic reasons why our per-packet forwarding approach is better than per-flow. First, we can use all of the bandwidth across all links even if there is just a single large transfer. This contrasts with per-flow forwarding, where a single flow can only use a single link. Second, and in fact more importantly for delivering reliability, our per-packet decision making means that if a network path starts to perform much worse - e.g., due to packet loss, or congestion-related increases in latency/jitter--we move the flow to a better path, in less than one second. Sessions are not lost, and good network performance is maintained even in a network "brownout" (congestion-related performance problem) or complete link failure. On the other hand, per-flow forwarding approaches make decisions at flow initiation time, and therefore frequently cannot respond to link failure, and definitely cannot react to congestion-related performance problems. To leverage the "works pretty well most of the time" public Internet with any reliability, it is especially important to do the sub-second switching afforded by per-packet forwarding.

For per-packet forwarding, it's critical to measure the performance of all network paths continuously, and to mitigate the effect of lost packets and re-order packets on the receiving side to deliver them in-order to the receiving client. Absent this technology, per-flow decision making is the only sensible approach.



Thierry Grenot - CTO - Ipanema Technologies

Ipanema's HNU clearly differentiates the forwarding mechanism - which is flow based - from probing and control that decides what is the best network to use from A to B for a given application flow at a given time.

While we constantly probe all possible paths in order to get the real-time quality and bandwidth map of each way, we trust it is usually more efficient to maintain a flow on a given interface for many reasons among which:

- a) it's simpler
- b) it is stateful-firewall friendly, and
- c) if you split among several interfaces, you basically get the quality of the worst one as you have to wait for the slower packet.

This does not imply that the choice of the network must be static. Actually, depending on the customer's security architecture, we propose several modes where the outgoing network might or not be dynamically reallocated.



Steven Taylor, Webtorials

How do you differentiate your products from the numerous products that perform "application acceleration" or "application delivery" enhancement?



Keith Morris - Talari

First, let's make sure we're using the same terminology. "Application delivery" typically refers to single-ended boxes like those made by F5. We don't compete with them; and in fact, WAN Virtualization is complementary for enterprise intranet uses.

"Application acceleration" typically refers to two-ended WAN Optimization (i.e., the market in which Riverbed is the leader). While our technology overlaps with theirs by probably 10% - 15% (more on that in a second); in fact, WAN Virtualization is also complementary to WAN Optimization. About two-thirds of our customers use WAN Optimization.

While both WAN Optimization and WAN Virtualization solutions can answer the problem of more bandwidth--WAN Optimization frees up bandwidth on existing private WAN links, where WAN Virtualization allows you to aggregate existing plus inexpensive Internet bandwidth, and utilize the aggregated bandwidth more efficiently--when it comes to application acceleration, WAN Optimization and WAN Virtualization "excel" at different things.

The two most important capabilities that WAN Optimization offers are disk-based data streamlining/deduplication, and application-specific support for Microsoft CIFS, the one protocol which really is "broken" on the WAN. WAN Virtualization's bandwidth aggregation, loss mitigation, sub-second switching and unique ability to deliver high-quality, real-time support even over the public Internet, are things which WAN Optimization doesn't do.

Customers have already started to widely deploy WAN Optimization for data center consolidation projects, and it delivers "unbelievable" speed-ups for "warm" transfers--files which have previously been accessed over the WAN for which the disk-based data deduplication delivers truly LAN-speed results--and substantial speed-up for CIFS file access overall. Talari's WAN Virtualization, by contrast, accelerates getting those huge files to a remote site much faster the first time ("cold transfers") with its bandwidth aggregation capability. The loss mitigation and sub-second switching mean it's the ideal solution for the "works pretty well most of the time, but occasionally performance is terrible" problem that is almost always a network congestion issue. WAN Virtualization can improve the performance of VDI/Citrix-type applications, where the highly optimized

protocol can't be improved much by WAN Optimization. And WAN Virtualization makes the network more reliable, and can enable higher quality VoIP, videoconferencing and video streaming even when using inexpensive Internet connections.



Thierry Grenot - CTO - Ipanema Technologies

We believe that there is not "one" magic technique, and that is why Ipanema's Autonomic Networking System (ANS) integrates several technologies: application classification to understand the applications, QoS and Control to guarantee business critical application performance, WAN Optimization to accelerate apps and minimize the required bandwidth and finally WAN virtualization (HNU) to unify application performance across hybrid networks. Ipanema's core strength lies in its ability to integrate all these apparently disparate technologies in a coherent, efficient and self-managed system.



Steven Taylor, Webtorials

Does encrypted traffic present particular challenges for you? If so, in what way?



Thierry Grenot - CTO - Ipanema Technologies

Ipanema's HNU performs even with encrypted flows, like SSL for example. Application classification, Control and dynamic WAN selection work over encrypted flows without sharing any confidential information between our system and the customer's IT. For WAN optimization to combine with HNU, the customer must activate the SSL acceleration feature (soon available).



Keith Morris - Talari

Our answer is similar to Thierry's here. We work just fine with encrypted flows. We continue to do bandwidth aggregation, sub-second switching of traffic away from problem connections, etc. If priorities are marked using DSCP/ToS bits, and/or classification based on 5 tuple information is sufficient, then in fact 100% of the functionality as for unencrypted flows is maintained. Remember, we are doing only WAN Virtualization, not WAN optimization, so we have no need to see the data contents, encrypted or not. Only if the encrypted stream is a bundle of otherwise separate TCP/UDP flows is any capability at all lost. In that case, the loss mitigation works for the bundle overall, rather than for the individual flow; for all but high loss or very high bandwidth situations, this difference will be barely noticeable, and even in this case, high reliability in addition to bandwidth aggregation is still delivered.



Steven Taylor, Webtorials

What is your position on whether it is reasonable to send realtime traffic (such as voice) over Internet links?



Keith Morris - Talari

So long as you have diverse networks at each location (and 2 is completely sufficient), delivering high quality VoIP over public Internet links is completely reasonable. We have many customers doing just that. We've done it ourselves for our internal communications for more than 2 1/2 years now. Our development team in North Carolina and our outsourced QA team in Bangalore, India use this daily, and in India they have only the least expensive local, highly oversubscribed DSL links.

We have two different techniques which enable highly reliable real-time application support. The first is pretty much the same as what we do for TCP traffic: namely, in addition to putting the real-time flows on the best path between locations with the lowest loss and the lowest jitter, we will quickly switch the flow, sub-second, to a better path in the event of network congestion related severe packet loss or even just high jitter. This alone will result in high quality voice, where occasionally a word or two might be missed, much as might result from, say, a quarter second of static on a mobile voice call.

By turning on replication for such voice flows, we will actually replicate the packet stream on two different paths, as unrelated as possible in terms of first mile and last mile links, and suppress the duplicate at the other side. In this way, even if what had been the best path starts experiencing 50%+ packet loss, or say 190 ms burst of jitter, the packets show up a few milliseconds later on what had been the slower/worse path, and the application never misses a beat. We count on their being jitter buffers in the clients, but in fact that's exactly how VoIP and videoconferencing works. In this way, we deliver "platinum" quality VoIP, better than the best MPLS network - including MPLS with QoS - can deliver.

Video can also run over Internet links the same way, with one important caveat: there needs to be sufficient available bandwidth. In particular, each link at a site with only two WAN links must have sufficient bandwidth to support the desired video session, in case congestion or link failure causes a problem with the other link



Thierry Grenot - CTO - Ipanema Technologies

I'm sure you have used Skype. Even if not always perfect, Skype is not that bad most of the time, which proves that the Internet suits for voice in many situations. With Ipanema's HNU, you can even improve the quality and resiliency by selecting the best network, be it another Internet connection or the corporate

MPLS VPN. Of course, a well-designed QoS engine has to take care of protecting voice flows from other data connections and preventing jitter and delays.



Keith Morris - Talari

This is one of our biggest disagreements with Ipanema, specifically where they make the comparison with Skype. The popular perception of the Internet is true: it works pretty well most of the time. But "pretty well" is not good enough for most people, and most of the time is not good enough for almost any enterprise IT manager. Skype may work "very well" with its various optimization techniques, but it still only does this most of the time. In particular, it simply can't do anything about last mile congestion.

By measuring end-end and switching sub-second, we turn the network-of-networks, done via peering points, weakness of the Internet back into a strength. We'll always find a path that's working well, even mid-call, no matter what single point of congestion occurs. And by doing replication of real-time flows, we deliver higher quality still, better than can be obtained by the best-engineered MPLS with QoS network, and far better than Skype.



Thierry Grenot - CTO - Ipanema Technologies

In a nutshell:

- I agree with Keith that in very tough situations where transport media or networks are really bad, replication might be a solution.
- Nevertheless, it is complicated and costly. Either the switching has to be much faster than the problem it wants to escape from or the duplication will... duplicate the bandwidth.
- While I agree with Keith that last mile is the most pregnant network problem, this issue can be efficiently addressed by simpler methods like QoS and control.
- MPLS with proper QoS is able to deliver perfect quality for voice, whatever the application mix and traffic load on the line...



Steven Taylor, Webtorials

Do you recommend and/or support the use of Internet services from multiple ISPs? Similarly, how do you handle traffic when one ISP is used in one network location and not at another?



Thierry Grenot - CTO - Ipanema Technologies

For resiliency reasons, we recommend several ISPs (MPLS and/or Internet). Ideally the split should include wiring to POPs in order to limit the probability of severe simultaneous incidents. On the other hand, there is nothing that prevents a customer from having the same ISP in a branch, or different ISPs end-to-end (it is generally the case as ISPs have only a local reach): HNU will perform well and automatically deliver the best of the combined network.



Keith Morris - Talari

We can deliver our better-than-MPLS reliability and application performance predictability so long as there are two or more WAN connections at each location from at least two different providers. For MPLS augmentation, this could be as little as an existing MPLS connection plus whatever the customer is using for local Internet access or VPN backup.

For MPLS replacement or MPLS avoidance, two different ISPs at each site are sufficient. They can be the same two at each location, but do not need to be. We support up to eight WAN links per location. You could have, say, four DSL links from the local RBOC/PTT, plus a T1 from Sprint. We will of course aggregate the bandwidth from all of your links, even for a single flow. Thus, while we'll certainly support having multiple different ISPs at a given location, in reality most customers will only use two or at most three different providers per location.

As for the case where there are different ISPs at one site versus another, this is where our implementation of WAN Virtualization really shines. If you have, say, two ISPs at a remote location and two different ones at a data center, we create four paths between the two sites, using all four of the 2 x 2 possibilities. Because we do per packet rather than per flow forwarding, and are continuously measuring one-way network loss, latency, jitter and bandwidth utilization, we can switch away from a problem connection sub-second even for existing flows.

Because we can switch sub-second, we turn the major "QoS" weakness of the public Internet--namely, the peering points interconnecting the various ISP networks--into a strength. If just one of the four paths (up to four possible peering

points) in any one direction is working, we'll get your traffic through, certainly your important traffic. And just as the Internet does "hot potato" routing, we don't need to send the reverse path traffic on the same path as the forward direction, and so even if it is a different path using completely different peering points in the reverse direction, we'll ensure that flow performance is maintained and sessions are not lost or degraded.



Steven Taylor, Webtorials

Must you deploy appliances at all branch offices?



Thierry Grenot - CTO - Ipanema Technologies

Ipanema's HNU works at the flow level. Among its many different advantages, this allows for the selection of the best network interface from a site even if the remote site is not equipped by a physical (or a virtual) appliance. The flow back from the remote site will be normally routed.



Keith Morris - Talari

We are a two-ended solution, so for us to add value in terms of reliability and performance predictability, to say nothing of adding bandwidth or lowering WAN costs, we do need to have an appliance at the site to add value directly for that site.

That said, we certainly don't need to be deployed at all of a customer's locations! Our WAN Virtualization technology plays well in the sandbox across sites that have Talari appliances and those that do not. In fact, our ability to move traffic away from a congested link or path sub-second means that under times of duress, say for example congestion on an MPLS link to a data center, the sites with Talari appliances will move away from using that link until network conditions improve, thus freeing it to service sites which only have MPLS connectivity. In this way, we will add some value for all sites on the network, even those without a Talari appliance.



Steven Taylor, Webtorials

What should a typical company see as a ROI?



Keith Morris - Talari

Our solution is designed to pay for itself in less than a year versus buying more North American MPLS bandwidth or replacing your MPLS connection. The more international locations there are, the faster the time to cash breakeven, in the range of 5 - 8 months.

This is just the "hard" cash return, of course. In addition customers have a lot more bandwidth: 4 - 10x is typical. And of course greater reliability and application performance predictability. This adds "soft" ROI in terms of avoidance of downtime and lower IT costs in terms of troubleshooting, and of course the "softer" still return from ensuring that end users maintain a good application experience. Unlike other companies, because the benefits are so compelling, we're able to focus only on that "hard" ROI. We're not looking to take money from the Cisco equipment budget - nor the WAN Opt budget for that matter - but rather from the AT&T/Verizon/BT budget.



Thierry Grenot - CTO - Ipanema Technologies

Savings can come from many aspects: WAN cost savings like Keith is proposing above, but also users' productivity (no lost time due to application performance brownouts) and IT agility. Just think about cloud computing. One of our large international customers that moved to Googleapps and used Ipanema's HNU to manage both internet and intranet flows saved not only on the application side in a controlled and safe manner, but also decreased the network budget by 20% while getting 3x the previous bandwidth. A win/win/win situation.



Keith Morris - Talari

We pretty much agree with Ipanema that there are also many "soft" benefits that WAN Virtualization technology delivers to customers. After all, we've designed the WAN as a "Smart Network" in a way that Cisco has talked about for years but never really delivered. Rather than merely giving you the tools to troubleshoot WAN problems - and we do give you richer information on the WAN than is available anywhere else - we actually go and fix the problem for you, sub-second, perhaps sending you an alarm/alert/email notifying you of the issue, rather than you having to be reactive and address the problem when your users complain.

That said, ever since the economy got soft, most customers care more about hard dollar savings, and so we tend to focus on these far more than we do the softer productivity or ease of management benefits.



Steven Taylor, Webtorials

In the majority of the discussion above, it seems to have been assumed that the two services that are combined/contrasted are MPLS and Internet-based VPNs.

Over the past several years, we've also heard a lot about "Private IP" nets (without MPLS) and also about Ethernet-based VPLS.

Why are these not mentioned? Are you not seeing these services used - or are they used by a different community of customers?

Do they present any particular challenges or advantages for use with your solutions?



Keith Morris - Talari

Steve: just two small points here. We'd classify Ethernet-based VPLS as simply another variant of MPLS as far as WAN Virtualization technology is concerned. In fact, any kind of private IP networks—including point-to-point TDM connections and customers' internal privately built MPLS networks--are the effective of MPLS or Frame Relay to WAN Virtualization. We simply usually refer to service provider MPLS networks when referring to the private WAN simply because these are the most common used by our customers.



Thierry Grenot - CTO - Ipanema Technologies

Agree with Talari: what matters is end-to-end IP connectivity, whatever the media and the network service enterprises buy from their telecom operator. It is true that today we see a strong demand for combining MPLS with Internet (even MPLS+MPLS or Internet+Internet), but if/when other cases will arise, the same solution will apply. After all, clients and servers will still continue to communicate over IP.



Steven Taylor, Webtorials

So here's the bottom line.

You're in the elevator with the CTO of a company, and he's headed to a meeting to decide between your two solutions.

What's the bottom line as to why your choice should be the ultimate winner.

(You must be succinct. Pushing the button for stopping at every floor is not permitted - and would not get you any points with the CTO anyhow.)



Keith Morris - Talari

Talari's WAN Virtualization technology does for enterprise WANs what RAID did for storage--delivering a network with 30 to 100 times the bandwidth per dollar, ongoing WAN costs reduced by 40 to 90 percent, and greater reliability than existing corporate WANs--bringing Moore's Law and Internet economics to enterprise WAN buyers for the first time in more than 15 years. By continuously measuring end-end across all network paths and switching sub-second when network performance problems occur, our WAN Virtualization technology turns the weakness of the network-of-networks that is the Internet into a strength, enabling customers to cost effectively augment or replace their MPLS WANs without having to sacrifice reliability or application performance predictability.



Thierry Grenot - CTO - Ipanema Technologies

"Ipanema's WAN Governance - powered by its Autonomic Networking System - will provide you with full control and optimization of all applications over your global network, private cloud and public cloud. You will get clear KPIs about application performance, enforce business critical application SLAs, accelerate flows and optimize cost across hybrid networks - all of this with a fully self-managed solution."

And of course, I'll invite him to share a single malt right after his meeting, whatever the outcome!

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.

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 © 2011, 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.