# Hybrid Networking and WAN Optimization



# A Webtorials® Featured Discussion

Regardless of whether your current computing model is primarily premises-based or you're moving toward a cloud-based model, the WAN remains the "glue" that holds the network together.

And that glue must be optimized in several respects. On the one hand, there's performance optimization to ensure that your applications, whether they're within your data center or in the cloud, perform at a predictable and consistent level.

At the same time, they also must be "cost-optimized." That is, you need to pay for an appropriate level of service for the various applications, and all applications do not require the same level of service. You're willing to pay a premium price for services that require it. At the same time, it would be great to be able to use less expensive services for less mission-critical apps.

Ipanema Networks is emerging as a major force in this particular market. In particular, while they offer network application performance optimization products, they augment this capability with their "hybrid networking" capabilities for optimization based on using a variety of carrier services.

The foundation for the use of multiple services is discussed as an overview in the paper, <u>*Hybrid Network*</u> <u>*Unification - MPLS plus Internet-VPNs*<sup>1</sup>, and I recommend it as a supplement to this discussion.</u>

I invite you now to join me as I discuss these issues with Thierry Grenot, CTO at Ipanema Technologies.



# Steven Taylor

Welcome / Bienvenu Thierry! And thank you for joining us.

The paper referenced above mentions hybrid networking using a combination of MPLS and Internet-based VPNs. This makes a lot of sense because using these two services are at the extremes of high control, reliability, and price (MPLS) and minimal control, reliability and price.

However, there are other options available as well, such as WAN Ethernet services and even dedicated transmission services (such as T3/E3 and SONET/SDH). Can you please tell us a bit about which services are supported by your Hybrid Networking technology and how a user might go about deciding which services are optimal for their situation?



# Thierry Grenot, CTO, Ipanema Technologies

Hi Steven - Actually the transport technology does not matter. Hybrid networks can be made of any type of combinations of different networks. It happens that, as most of services are IP-based, in most cases we found MPLS and Internet combined in one 'unified network'. But we also see two MPLS, and why not Ethernet or pure virtual leased lines (like T1 or T3). At the end of the day, it is about of IP traffic between clients and servers, that's all the matter.

<sup>&</sup>lt;sup>1</sup> Please note that Webtorials <u>registration</u> required for access to some hyperlinked materials. Click <u>here to register</u> or click <u>here if you forgot your username/password</u>.



Steven Taylor

In the network diagram shown here, it appears as if you would deploy appliances using a variety of physical network interfaces and logical interfaces within the physical interfaces. (Click on the diagram to see a larger version.)

Can you please explain how many physical interfaces are supported, and also how logical interfaces are separated within the physical interfaces?



The Ipanema system is able to work with up to 3 logical interfaces. While in most cases 2 interfaces are sufficient it may happen that enterprises want to specialize one Internet access for IP/VPN and another one for Internet browsing - so if you add one for MPLS then you need 3. To be honest, having more is not an issue - we just didn't foresee the need.





### Thierry Grenot, CTO, Ipanema Technologies

Just to complement the previous information - Currenty Ipanema proposes up to 3 virtual interfaces shared on a unique Ethernet physical interface (on the WAN side of the ip|engine appliance). Multiplexing is provided either through TOS/DSCP marking or destination MAC address selection (spoofing). We will also propose soon devices with multiple WAN-side Ethernet interfaces. Cheers.



# Steven Taylor

Thank you. So, since there are two Ethernet interfaces, what is the most common place for the appliance to be placed? Between the LAN switch and the router? And then the router will choose the appropriate WAN interface(s)?

Also, since this placement could be viewed as a possible "single point of failure," do you have "fail to wire" capabilities?



# Thierry Grenot, CTO, Ipanema Technologies

You want all details, Steven, don't you? Concerning the appliance locations:

- EITHER there is only 1 router with two access lines - so we're mapping two 'virtual' interfaces on a single

Ethernet line. Note that there is no need for 'virtual routers' (VRF) – just a static policy telling /if/ TOS bit x = Y /then/ go left /if not/ then go right.

- OR there are two routers (one per line) and here two options are offered:

--> The same TOS field is used to ask the first router (usually the default gateway) to direct flows to the other one (usually through a line between them): /if/ TOS bit x = Y /then/ send flows to the other router /if not/ then take care by yourself

--> Or MAC re-writing to directly aim at the appropriate router thanks to an intermediary switch.

There is no constraint about router, networks, bandwidth, etc.

This does not raise availability features (the Ipanema appliance would fail to wire in case of pb.) but may require an additional device (low cost switch or hub). This is why we will provide a multiple WAN interface device very soon.



# Steven Taylor

Even though the diagram above only shows two network sites, I assume the capabilities scale well to a meshed and/or partially meshed network.

Can you give us a feel for the scalability of the solution, in terms of number of site supported, etc.



# Thierry Grenot, CTO, Ipanema Technologies

Correct, Steven. Application traffic is more and more meshed (just think about voice traffic and even datacenter 'deconsolidation' caused by Cloud computing and SaaS) so it is important to work in an any-to-any topology. The Ipanema system works (well :-) for very large networks having thousands of sites, with an actual average customer size of ~ 100

sites.



# Steven Taylor

Moving for a moment to optimization issues, Ipanema has used the diagram shown here rather extensively. (Again, click on the diagram to enlarge.)

Can you tell us a bit about your optimization solutions and how we should interpret this diagram?



percentage of satisfied users across the network



# Thierry Grenot, CTO, Ipanema Technologies

Do you have kids? I have, and I can tell you that when she manage to put her room in order, it is always too small and chests do not close well. When it's one of her parent, the room is neat and there is still lot of available space in drawers... That's the same with networks: when all individual flows are put in row, there is plenty of space for new one and quality is excellent - This is the power of per-flow adaptive control: all flows are continuously rearranged in real-time to fit with available end-to-end network resources.



# Steven Taylor

Am I correct in my understanding that one of the primary benefits of hybrid networking is that both links may be used under normal circumstances for load sharing), but that all traffic will move to the other link(s) in the event of a link failure?

If so, can you tell us a bit about the metrics that are used for deciding when a link is "failing"?

In particular, must one totally lose connectivity, or is there a way that you also watch for degradation (especially in the case of Internet-based services) and start shifting traffic based on poor performance?

_ 1	
	100
	( Cardon )

## Thierry Grenot, CTO, Ipanema Technologies

The fundamental benefit of Hybrid Network Unification is that you always use (on a very granular and dynamic manner) the best network. Of course, if only one network is available, it has to be this one. But while networks are pretty reliable today, their end-to-end performance is very much fluctuating because of a) core network occupancy - e.g. Internet at peak time and b) usage of access lines by the enterprise applications. So Ipanema is selecting the best WAN based on application performance objectives, i.e. according to the effective performance that can be delivered by the user. The set of criterias depends on the application: voice will likely use the fasted path (here delay metrics are key) while file transfers will probably prefer the largest access (available bandwidh is the useful metric). With thousands of users and hundreds of applications, you can easily understand the combination of cases and all the variability that should be taken into account: manually impossible, a well designed optimization algorithm can make it.



# Steven Taylor

Let's chat for a moment about the interaction between using alternate links (hybrid networking) and network optimization.

In some cases, a well-optimized Internet-based link may approach the performance of an MPLS link.

How do you decide when to use which? Is this something that must be in some way specified (such as on a perapplication basis), or does your appliance help do this - or at least provide "recommendations"?



Thierry Grenot, CTO, Ipanema Technologies

This question is actually related to the former one. It does happen that Internet is faster than MPLS, or that MPLS is less loaded than Internet. And so on. So the point is to not decide in advance: we don't know tomorrow's weather. We use a real-time sense and response principle that understand what is the performance for all networks and select the "best one" from a set of several criteria that depends on the application - what we call our 'application performance objectives'. No need to be too fast on file transfer - no need to be too large for voice. What about video?



Steven Taylor

Most optimization solutions to date have been designed primarily for use between premises. However, we are now in the midst of a massive shift in some cases to cloud-based computing.

Can you give us a brief overview of how you can provide optimization for cloud services? (Of course, this is probably a sufficiently large question to merit its own discussion at some point.)



# Thierry Grenot, CTO, Ipanema Technologies

I agree with you Steven: this is a fairly large point! In a nutshell, most of Cloud applications are over the Internet ("THE" cloud)- This radically transforms the traffic matrix, from the traditional hub-and-spoke to a weird hub-and-hub configuration. Cloud computing --> Simplifying IT + Complexifying networking! A good start to talk about control and optimization, isn't it?



### Steven Taylor

What size do networks need to be in order for your solutions to be cost-effective? Is this primarily for Enterprises, or may mid-sized businesses also find a quick Rol?



# Thierry Grenot, CTO, Ipanema Technologies

There is not particular size - it works from small up to very large networks. Actually ROI is a multi-criteria animal and unifying hybrid networks provides much more than 'only' network budget reduction. This being said, one of our largest customer recently announced that it x3 is global capacity while decreasing by -20% its network budget and improving application performance at the same time: a good win-win-win situation.

# About the Webtorials<sup>®</sup> Editorial/Analyst Division

The Webtorials<sup>®</sup> 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<sup>®</sup> Editorial/Analyst Division products, please contact Jim Metzler at <u>jim@webtorials.com</u> or Steven Taylor at <u>taylor@webtorials.com</u>.

Published by Webtorials	Professional Opinions Disclaimer
Editorial/Analyst	All information presented and opinions expressed in this publication represent the
Division	current opinions of the author(s) based on professional judgment and best available
www.Webtorials.com	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.
Division Cofounders:	
Jim Metzler	Copyright © 2010, Webtorials
jim@webtorials.com	For editorial and sponsorship information, contact Jim Metzler or Steven Taylor.
Steven Taylor	The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture
taylor@webtorials.com	of Steven Taylor and Jim Metzler.