

Sponsored by:

## NETWORKWORLD

This story appeared on Network World at  
http://www.networkworld.com/newsletters/frame/2011/050911wan1.html

# WAN Virtualization - Per-flow vs. per-packet advantages

[Wide Area Networking Alert](#) By [Jim Metzler](#) and Steve Taylor, Network World  
May 09, 2011 12:01 AM ET



After a short break to discuss forthcoming Interop sessions, Sponsored by: this is the third of a series of four newsletters where we're sharing excerpts from the Webtorials [Thought Leadership Discussion on WAN Virtualization](#). This series features a "virtual discussion" with Keith Morris from Talari Networks and Thierry Grenot from Ipanema Technologies, the two leading companies in the WAN virtualization space, and today we found an area where there's a significant and most interesting difference between the two companies and their approaches.

### [WAN Virtualization - Transport and encrypted flows](#)

We observed and asked, "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 (Talari) began the discussion with, "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 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 (Ipanema) responded, "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 (Talari) further commented, "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."

Since we found a significant difference here, we then asked each to summarize briefly why their solution is "better."

First, for Keith (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."

And from Thierry (Ipanema), "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 might not be dynamically reallocated."

It's great to have such a spirited and insightful discussion. And we'd love to have you join us with further comments at the [Thought Leadership Discussion](#).

[Read more about lans & wans](#) in Network World's LANs & WANs section.

<http://www.networkworld.com/newsletters/frame/2011/050911wan1.html>

[Steve Taylor](#) is president of Distributed Networking Associates and publisher/editor-in-chief of [Webtorials](#).  
[Jim Metzler](#) is vice president of [Ashton, Metzler & Associates](#).

All contents copyright 1995-2012 Network World, Inc. <http://www.networkworld.com>

[http://www.networkworld.com/...ww.networkworld.com/newsletters/frame/2011/050911wan1.html&site=printpage&nsdr=n\[8/19/2012 10:50:33 AM\]](http://www.networkworld.com/...ww.networkworld.com/newsletters/frame/2011/050911wan1.html&site=printpage&nsdr=n[8/19/2012 10:50:33 AM])

FATPIPE-001616



Find authenticated court documents without watermarks at [docketalarm.com](http://docketalarm.com).