

► GIGABIT ETHERNET

The pipes aren't the problem

Here in the LAN Times Testing Center we have some of the coolest gigabit Ethernet gear in the world. For example, we have this keen new gigabit Ethernet switch. This baby was designed to put my servers at the core of a gigabit Ethernet backbone.

And what a screamin' server farm this switch could build! It can support up to 2.5 gigabit Ethernet ports and route more than 50 million packets per second at Layer 3. It supports policy-based QoS (quality of service) and vLANs (virtual LANs). Its backplane is over 50Gbps wide. I even love its management application.

I also have some very slick gigabit Ethernet adapters. Not only do they deliver the highest throughput to the server, but they also have onboard processors that limit stress on the server and balance loads among several cards.

Yep, with this equipment, I could build one smokin' server farm. I would have all of the throughput anyone would need for years. I would be the ruler of time, space, and bandwidth.

There's only one problem: Performance on my network wouldn't improve one bit. And this stuff would not improve performance on anyone else's network either.

At least not right now, because while we weren't looking the network bottleneck moved again, and most people still haven't noticed.

You see, despite all the hype, the world of Intel and RISC networking already has plenty of bandwidth. Right now, with 100Mbps cheaply and readily available to nearly every computer across the land, networks have more than enough bandwidth to handle all but the most demanding applications. Therefore, only uplinks aggregating traffic from multiple ports really need pipes wider than 100Mbps. And not all of them even need that kind of speed.

So if the problem isn't bandwidth, what is it? It's internal I/O on the computing hardware platforms. There's not an Intel-based server around with I/O that can keep up with 100Base-T, much less gigabit Ethernet. In fact, in our Server

Only uplinks aggregating traffic from multiple ports really need pipes wider than 100Mbps. And not even all of them need that kind of speed.

And load balancing gigabit Ethernet NICs? They're a waste right now.

So what is the clatter all about? Why is everyone being inundated with infrastructure devices with which to widen pipes when the pipes aren't the problem? And more importantly, why are we buying all of this gear?

Since network managers are under intense pressure to improve network performance, they often find themselves

Track testing, the fastest server we've found so far has a peak throughput of a little below 10MB per second.

So we're not even close to threatening gigabit Ethernet here. In fact, you could add bandwidth from now until doomsday but never experience better network performance.

buying higher bandwidth switches. They must appear to their management to be doing *something*. However, genuine gains in performance are going to require grassroots lobbying of RAID and server vendors.


Why aren't we instead insisting on higher internal throughput at the server? Take a look at the Server Track data, and you'll see that there's a lot of room for improvement.

LOOKING FOR TRENCHES?
It's now located in the Hands On section.
See page 68

When the throughput of a \$28,000 server (the Micron Technology Inc. 9008XP) doesn't vary much from that of a

\$14,000 server (the Digital Equipment Corp. 3205R) in the same environment under the same load, you've got to figure that it's the server industry-not the switching business-that has hit the performance wall. And server makers need to figure out how to break through it.

That will let us put all these pipes to work and start worrying about insufficient bandwidth again. ■

 What do you think? E-mail us at letters@lantimes.com.



Tere Parnell