

The price of 10Gig network hardware has tumbled in the last 12 months, with stackable switches offering 10Gig ports and blades in Chassis readily available. This is leading people to create 10Gig backbones, mainly because they can. As people upgrade or contemplate upgrading we are frequently asked about the true increase in performance, surely when you install 10Gig it runs at 10Gig right? Wrong, not even close, and here is why....

## Introduction

Network Speeds have been increasing for a number of years, always in blocks of 10. However the relative increase in performance that the users see is actually getting smaller and smaller. This is complicated by the fact that whenever something runs slowly, the network is blamed and the default response has been to increase the bandwidth.

Whilst this has worked to some degree for a while, the increasing cost for diminishing returns argument is going to run out sooner or later. So if Bandwidth is not the issue, what is and why can't you get anything like the stated performance numbers in real life.

## Technology

When Ethernet was designed in the early 80s, frame sizing and timings were all based around 1, 3 and 10Mbs transmission rates, shared cable technologies, the expectation of high error rates and Terminal server size files and screens. If we fast forward to today, 1000Mbs networks are common and 10Gig networks are finally starting to take off.

These early frame formats are basically very in-efficient ways of transferring the sizes of files that are common today. Error rates are much lower now meaning more data can be sent without the fear of resends, management information takes up too big a percentage of the amount of data being sent and the need to allow for half duplex links has gone as this was finally taken out of the 1Gig standard.

However the biggest issue with the early performance of Gigabit cards was the speed at which the drivers could move the data between the NIC card and the Server hardware. It was quite common 4-5 years ago to get a Gigabit NIC realistically moving data at around 300Mbs. Today with the development of technology, Gigabit NICs in the right sort of server hardware are realistically capable of 8-900Mbs.

## So the question remains, what can we expect out of 10Gig Connections?

As a basic rule, if you take any 10Gig NIC card, install it and just use default settings, the manufacturers are claiming about 2Gig of performance.

To get anything like the 10Gig you were hoping for you have got to start "tuning" the NIC drivers and even how the drivers are handled. The first trick that you need to be aware of is that the stated performance figures on a lot of NIC card datasheets are for UDP only - which is hardly real world as apart from UDP applications like VoIP and Video almost all network traffic is TCP. The issue with TCP being that you have to wait for ACKs and these effectively throttle the usable bandwidth in latent environments as this becomes the slowest part of the transfer mechanism.

The second thing to be aware of is that most of the stated throughput figures will refer to a bigger packet size (see Jumbo frames below). This is a very effective way of increasing performance, so long as both ends and the switches in the middle support it.

If you play with these numbers then the chip sets and NIC cards tested claim to produce something between 6 and 8 Gig performance, but remember this is in the lab with very unrealistic data streams.

These numbers also avoid the biggest problem with real performance, the speed of the drivers .....

## 10Gig Performance tuning Options – or the Wizard Gap

The first thing to note is that to call something 10Gig is to refer to the serialization rate of the NIC card and bus, the rate at which it can take and receive 1s and 0s. There is nothing in this statement which refers to the performance of the drivers or how efficient Ethernet is at transferring the data, just the rate of 1s and 0s on the line.

## Jumbo Packets

The good news is that in order to get some really fast numbers out of 10Gig interfaces the standards are allowing for bigger and bigger packets to be used, sometimes referred to as Jumbo packets. The standards mention numbers up to 16,000 Bytes (the default maximum up to this point is only 1500 Bytes) although the industry seems to be settling on 9,000 Bytes as the size of choice.

The bad news here is that each end of the conversation has to support this Jumbo standard for this to work, so not an option for all the legacy devices in your network. This option also needs to be turned on, so things are not going to spring into life just because you have NIC cards with this feature.

## Interframe Spacing

There is also a growing amount of information on the Internet about performance increases available by changing the space between Ethernet Frames. In a full duplex environment there are no actual real collisions in Ethernet, it's all a pseudo collision mechanism to ensure compatibility with older types of Ethernet.

Taking the gap between frames to zero doesn't appear to be the answer, performance tests on the web show packet loss issues at this point, but improvements can be found by making the gap smaller.

## TCP off load engines

However, by far the biggest problem in getting better performance out of a 10Gig network is the speed of the drivers passing information on and off the NIC. The solution to this today has been driven by *Packet Analysis* and *Network Testing Technologies* that really do need to perform at the stated speeds.

Most of these solutions have developed custom NIC card technology which gives the role of the drivers off to a dedicated part of the NIC card chip set (off load engine!) and more specifically gets it off your end user hardware.

This technology is allowing some of the top packet analyser companies to claim performance figures for 10Gig to memory (which will fill very quickly so not very useful) and 8-9 Gig to Disk (which is very useful).

## Summary

Although the NIC cards for 1 and 10Gig can move 1s and 0s at line rate, we are a long way from a real life network seeing anything like these speeds. 10Gig technology is available out of the box today, but without a lot of work your realistic starting point is about 2Gig and with a lot of work you might get about 6Gig.

The positive point is that specialist networking tools are aware of these performance issues and in effect have developed specific technology to ensure they can operate on or near 10Gig speeds so if you are throwing a number of servers and users all down a key link, you do have a way of keeping up!