

Can my network support 10Gig?

Which is the right tool for that anyway?

The issue

An increasing amount of people are asking “Can my network support 10Gig or not?”. If the network cabling and hardware are new, then this is more of a checking mechanism than anything else, but what if your cabling is a few years old or your switches have 10Gig uplinks but not full 10Gig support? Can you now light up the network to 10Gig and expect it to work?

There is still much confusion about what 10Gig network testing really means. Some engineers see network testing as looking at the physical copper and fibre cabling, whilst others see it as loading up the switches and routers to verify that all, or most of it, gets through. So let’s take them one at a time.

In summary, there are two parts to this:-

- Is my cabling good enough?
- Can my network hardware actually run at these speeds?

10Gig Fibre Cable Testing

Historically to run 10Gig connections you had to have single-mode fibre. In recent years however the new multi-mode fibre (OM3 and above) can now support these speeds. One of the biggest issues tends to be people who installed single-mode a few years ago and now want to light it up, you need to watch out for:

- Dirt – have the dust caps been in place all the time?
- Water ingress – has some part of the duct got wet and damaged the cable?
- Stress damage – have a pile of newer cables been installed laid on top of the fibre stressing it in some way?
- Quality of the original installation – are the splices and terminations good enough to support 10Gig?

Once you are sure the fibre end faces have been cleaned properly, the starting point with any sort of fibre testing is light loss. Every standard specifies the amount of loss the receiver can take before it loses lock on the signal. With modern SFPs these numbers are improving (i.e. they can recover a weaker signal), however with speed increasing the amount of data each modulation is carrying is increasing, so the need to be ever more accurate in recovering the data.

The standards are moving towards specifying more and more “event” based testing; an event being a splice, join, termination etc. The aim here is to spot whether one poor event is affecting the whole performance of the cable and therefore can be fixed to bring it back into specification. To achieve any information at this level you will need an OTDR which sends a pulse of light and a sensitive receiver to look for tiny reflections caused by the events.

OTDRs are also the only ones to answer the second most popular question we get asked about fibre testing, “Can you do distance to break or the worst event?” The confusion here comes from the fact that a number of light loss testers claim to measure the distance of a cable. What they are actually doing is measuring the light loss of the cable and calculating the likely length based upon light being lost at a regular uniform rate along its length, but they have no mechanism for spotting a bad event or complete break in the cable.

One point to remember when using an OTDR is launch leads. The first part of the cable where the light is launched is effectively a blind spot until the whole pulse is on the cable and the light has settled down into a modulated stream.

Can my network support 10Gig?

Which is the right tool for that anyway?

The cure for this is to use a launch cable, which is not part of the run you want to measure, to give the light the chance to settle and allow more accurate measurements.

10Gig Copper Cable Testing

The 10Gig over copper standard covers a number of different cable types but the big one is Cat6a. There is a rise in installing this cable type in data centres as it provides a high density of 10Gig cabling between data cabinets.

Testing cables to standards is really about the ability of the cable and connectors to keep the signals on each pair isolated from each other. The phrase used here is Cross Talk, and there are an increasing number of different types of Cross Talk, some measured and some calculated.

There is chatter about using Cat6 cable to support 10Gig over short distances; as up to 37m Cat6 Alien Cross Talk measurements are inside 10Gig spec and between 37 to 55m they might be in spec. However this is not a supported technology for 10Gig, hence you take your chances here. If you have a data centre cabled with Cat6 and you want to use it for 10Gig then testing the Alien Cross Talk is your only way forward.

Another factor that has emerged in the last year or so is the rise of “cheap” cable and components from the Far East which have turned out to be copper coated aluminium. As the standards are trying to squeeze more performance from your existing cable infrastructure (10Gig over Cat6) performance for such inferior cables completely fall off a cliff and go nowhere.

The “other” type of copper 10Gig cabling we have been asked to test is CX4 or Infiniband. It’s a form of modern coax cable with a length of 15m and some blade chassis solutions are shipped with this as the cable type.

So in summary, if you are talking about 10Gig copper cable testing, then the important measurement is the Cross Talk and with failure rates on installation and testing seen higher than 10% accurate testing is more important than ever.

Network or Traffic Based Testing

The other side of 10Gig testing comes from the network engineers who want to know if their network hardware can handle 10Gigs worth of traffic. As soon as we put boxes on the end of cables (switches, routers, firewalls etc) they are going to have some sort of impact, the question is how much?

Throughput

The first and most common question people want to answer is “Can I actually get 10Gig from one point to another?”. To achieve this you need to generate 10Gigs worth of traffic, count it out and count it back in again. The first challenge here is that 10Gig NIC cards are better at receiving frames than they are at generating them and almost none of the commercial NIC cards are capable of operating at 100% for more than a few frames at a time, so beware of freeware tools who at the higher speeds are only testing the limits of the NIC card in your machine. Instead a dedicated 10Gig device with a special interface NIC is required to carry out this type of work.

There are standards around this type of testing (more of which below) but the first point to note is that all these tests need to be based on UDP frames. This is because TCP conversations rely on ACKs between blocks of data to

Can my network support 10Gig?

Which is the right tool for that anyway?

keep the conversation going. ACKs slow things down and part of the process of resending lost data in the testing situation are actually not relevant.

One of the most common issues we see is people “testing” links with file transfers and getting confusing results. This is because the protocol being used (FTP, TFTP, NetBIOS etc) actually has an impact on the throughput rate so to get a consistent result we need to remove the protocol.

The test we use the most here is RFC2544, which is a simple UDP packet stream with an increasing frame size. It’s simple to understand the results, packets sent and packets received.

10Gig Servers and Virtual Deployments

Part of the challenge here is that very few 10Gig connected devices can actually work at these speeds, but that doesn’t stop people believing this is the starting point for the expected performance. 10Gig NIC card performance is improving, but at the time of writing a good card can achieve a sustained 6Gig of performance. They are allowed to say it’s a 10Gig card because the bus speed means that they can serialise data onto/off the wire at the rate associated with 10,000,000,000 bits per second, but what they can’t do is keep this up across multiple frames, rather they have to wait for the buffers to fill/empty before serialising the data again.

If you then take this to a virtual server deployment your 10Gig connection supporting several servers with a reduced throughput performance becomes a very interesting question. In these situations you need some sort of software solution (no hardware deployed inside a VM) but there are options:-

- AppNet PathView Virtual testing – target an IP address
- Fluke Networks OptiView XG – has a downloadable software target

Latency

One of the factors becoming a key player in end to end performance during the last 5 years is latency (and therefore jitter). When companies started to ditch the expensive WAN links for cheap DSL style circuits which claimed the same bandwidth, most applications really struggled. The issue here was that TCP will only wait so long for the ACK frames and if they are not coming back fast enough then it resends the data anyway. In a link where the delay varies quite a lot (jitter) then this really messes up the throughput performance and causes all sorts of issues.

The issue with the RFC tests above is that UDP is not affected by jitter so the simple throughput tests are not going to spot this. To this end there is another standard, Y1564, which allows you to include measurements for Delay, Jitter and Packet Loss as well and set Pass/Fail measurements around them. There is some skill required in knowing what are good and bad numbers but depending on the products you are using there is help available.

You might think that as you are only running 10Gig in the Data Centre or Server Room then these factors are unlikely to be an issue, but experience has shown as otherwise. Typically layer 2 devices (simple switches) do perform well here as they have less to do and interestingly a lot of the switch manufacturers are putting 10Gig Layer 2 devices into their server designs. As you ask the switches to do more, VLANs, routing, ACLs etc performance really drops off quickly and it is these bottle necks that you want to find. If they work up to 3 Gig then struggle, it’s not necessarily the case you need to replace them but you do need to be aware of their limits.

Can my network support 10Gig?

Which is the right tool for that anyway?

Once you have found a link with issues the next stage is to try and identify which component has the problems.

Testing Live Circuits

The elephant in the room with all the above is that you need to flood the link with traffic to run these tests, so best done out of hours if possible. However most issues are seen during the day and that's when you are going to get the most realistic results so there is a compromise here.

One way out is not to flood the link with traffic but find a way of measuring its performance through other means. Technology exists which uses an algorithm for measuring throughput and hence only sends a burst of 30 or so packets to achieve this. This has the advantage that you can test every few minutes during the day time and plot the ups and downs of the throughput rates and see patterns emerge. The technology is also capable of showing at which hop the performance drops off hence highlighting the hardware (or configuration) that's causing you issues.

One of the great strengths of this approach is the opportunity to see performance stability and how it varies through the day/week/month and can be very revealing in showing one off issues versus regular weaknesses in the designs.

Summary

If you want to test 10Gig, the first question should be "Are we talking copper/fibre or are we talking traffic through switches and routers?" Both tests are relevant but tend to be the responsibility of different groups of people. As soon as you plug boxes on the end of cables there will be an impact, the question is how much?

If you are running high speed circuits across town or via a carrier then the questions remain the same but the impacts are likely to be higher due to the probable involvement of shared infrastructure in getting your circuit to work. The questions that need to be asked when you are using a third party carrier are "am I getting what I paid for?" and more tellingly "how consistent is the service they are providing?"

When it comes to test kit, look for a hire company or professional services if your requirement is a one off test to save on the overhead of buying product.