CABLING PLANNER

MAKING THE RIGHT CONNECTIONS

2014 ISSUE 015

A special supplement with

cnme
computer news middle east

SPONSORED BY

Nexans

NEED FOR SPEED

FIBRE BECOMES POPULAR IN DATA CENTRES

www.rdm.com/meta

R&M Middle East, Turkey & Africa

Office Cabling, Data Center, FTTx.

5 CELEBRATING 50 YEARS OF CABLING SOLUTIONS

R&M as a global Swiss developer and provider of cabling solutions for office, data center and FTTx exceed the most stringent standards in product quality. Intelligently designed for ease of use and optimal performance, our systems deliver flexible, reliable, and cost-effective networks that support even the most challenging businesses requirements.

R&M
Convincing cabling solutions

CABLING PLANNER

MAKING THE RIGHT CONNECTIONS 2013 ISSUE 013

IS YOUR ETHERNET FAST ENOUGH?

Ethernet's future is now about much more than the next top speed: The engineers charting a path for the ubiquitous networking protocol are looking at several new versions to serve a variety of applications.

To meet immediate demands in cloud data centres, there's a standard in the works for 25Gbps (bits per second). For the kinds of traffic expected in those clouds a few years from now, experts are already discussing a 50Gbps specification. And for enterprises with new, fast Wi-Fi access points, there may soon be 2.5Gbps Ethernet. That's in addition to the next top speed for carrier backbones and moves to adapt the technology for use in cars.

These efforts are all meant to serve a growing demand for Ethernet outside the traditional enterprise LANs for which it was originally designed. That means solving multiple problems instead of just how to get ever more bits onto a fibre or copper wire.

Without diving too deep into those details, here are some of the new technologies brewing in Ethernet.

1. 25-Gigabit

A 25Gbps standard may seem like a step backward, because 40-Gigabit and 100-Gigabit Ethernet already exist. But in fact, it's all about the need for more speed, specifically from servers in cloud data centres. Google and Microsoft are the biggest buyers of Ethernet now, largely because their cloud operations require so much data exchange between servers, according to Dell'Oro Group analyst Alan Weckel.

The key to 25-Gigabit Ethernet is that many of the components that could go into it are already developed: The 100-Gigabit standard is made up of four "lanes" of 25Gbps, so many of the same parts go into that high-end gear. That should mean higher production volumes for parts that go into both technologies, driving prices down.

2. 50-Gigabit

Work is also beginning on a 50Gbps specification, which could be the next speed offered for linking servers in data centres. Both servers and high-performance flash storage systems will drive a need for something more than 25Gbps in the biggest data centres in a few years, according to Dell'Oro Group.

3. 2.5-Gigabit

It may not sound very fast, but 2.5-Gigabit Ethernet might help companies fill their buildings with very fast Wi-Fi. It's being proposed specifically as a tool to help enterprises' wired infrastructure keep up with wireless access points that increasingly form the edge of those networks.

The latest Wi-Fi technology, IEEE 802.11ac, can operate at more than 1Gbps -- much more, with certain configurations. With that much traffic going over the air, the Gigabit Ethernet links that most enterprises use to connect their access points to the wired network could become a bottleneck.

Upgrading to 10-Gigabit Ethernet would give networks plenty of bandwidth, but most companies don't have the right kind of cable to do that. A 2.5Gbps version of Ethernet would work on commonly used Category 5e and Category 6 cable over the standard distance of 100 meters, so users could go beyond Gigabit Ethernet without the cost of pulling new cable.

4. 400-Gigabit

Ethernet's backers haven't given up on reaching a new top speed, either. An IEEE task group is already working on a 400-Gigabit Ethernet standard, which is currently projected for completion in March 2017. The fast links might use multiple lanes of either 50Gbps or 100Gbps. Once finished, the superfast technology would be destined for the cores of service-provider networks.

Published by

CPI MEDIA GROUP

Registered at IMPZ
PO Box 13700
Dubai, UAE

Tel: +971 4 375 1500
Fax: +971 4 447 2409

Printed by
Al Ghurair Printing & Publishing

Regional partner of

IDG
INTERNATIONAL DATA GROUP

FEATURE

OPTICAL NETWORKING

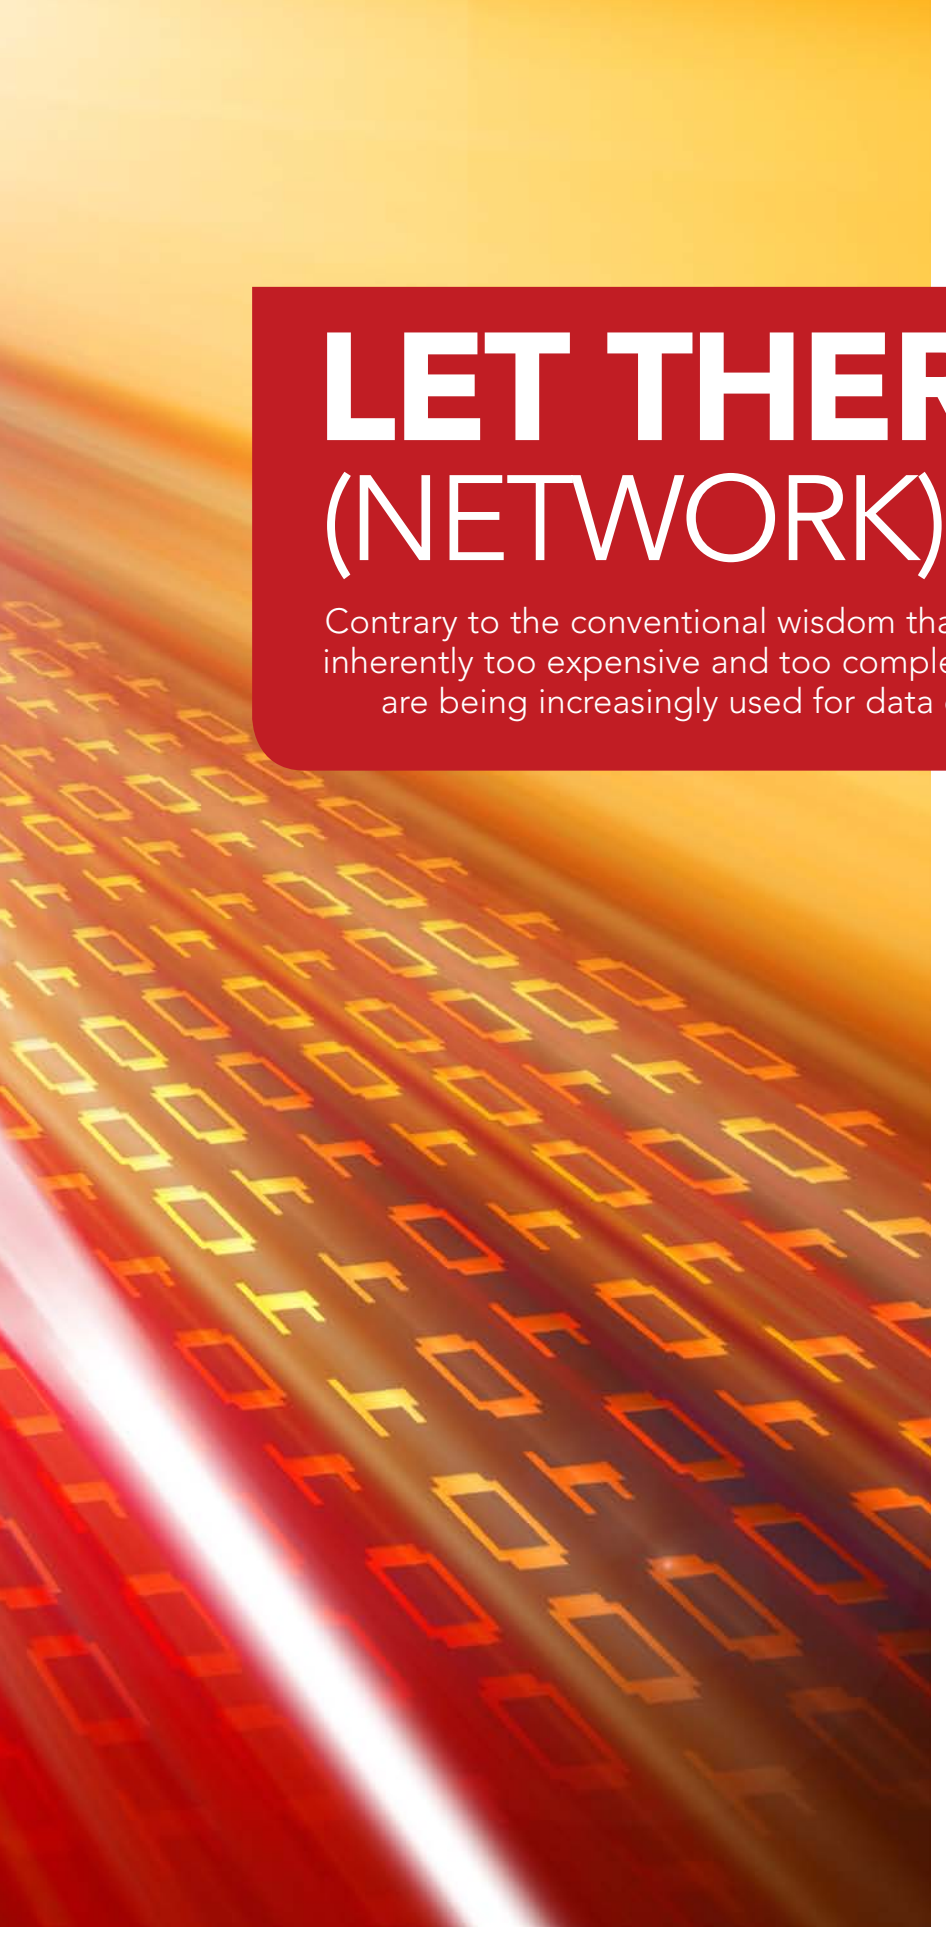

LET THERE BE (NETWORK) LIGHT

Contrary to the conventional wisdom that optical networking is inherently too expensive and too complex, fibre optic solutions are being increasingly used for data centre connectivity.

While the copper solutions were able to meet the basic needs of data centre connectivity, it is becoming hard for copper to meet the price, performance and size/weight needs for high-speed Ethernet networks. Fibre optic technology provides an alternative with higher performance, increased link distances and lower size/weight. What is more, fibre is becoming increasingly price competitive.

With the advent of SDN, many companies are now starting to rethink and redefine the fundamental approach to data centre networking.

"The adoption of enhanced network architectures is undoubtedly changing the way we look at optical fibre designs within the data centre. Data centre applications have components that are spread across many servers and performance of these applications is largely dictated by the data centre network and specifically the time it takes for different components to communicate with each other. These networks need to provide very high throughput with very low latency and this is resulting in a shift from the traditional Core/Aggregation/ Access architecture to a Leaf /Spine architecture," says Dave Hughes, Technical Director MEA, Commscope.

With a leaf/Spine design, any leaf switch can communicate with any other leaf switch by spreading traffic over multiple paths through different spine switches, so this way the bandwidth of this network can be scaled just by adding

more spine switches, he adds.

Every pair of leaf switches is just two hops away from each other, so this network design provides a very uniform and consistent latency between all pairs of switch fabric ports. Maximum number of uplink ports on a leaf switch determines maximum number of spine switches that could be used in a specific design.

Steve Morris, EMEA Product Manager, Copper and Fibre Solutions, Panduit, offers another perspective on the bold rethinking of data centre networking and its impact on cabling: "Fibre Channel over Ethernet (FCoE), Top of Rack (ToR), and fabric-based design are emerging breeds of architecture that are gaining traction as data centre (DC) strategies evolve to meet cost, scalability and agility demands.

"With DC consolidation and system aggregation omnipresent, fibre optics have become more prevalent and there's a growing need for them to support both higher bandwidths and higher port densities/fibre counts compared to legacy systems."

Alberto Zucchinali, RCDD, data centre solutions and services manager EMEA, Siemon, agrees: "With the continual requirement for expansion and scalability in the data centre, cabling infrastructures must provide reliability, manageability and flexibility. Deployment of an optical connectivity solution allows for an infrastructure that meets these requirements for current and future data rates. Switching and routing, virtualisation, convergence and high-performance computing environments are examples of where higher networking speeds will be required."

DATA CENTRE APPLICATIONS HAVE COMPONENTS THAT ARE SPREAD ACROSS MANY SERVERS AND PERFORMANCE OF THESE APPLICATIONS IS LARGELY DICTATED BY THE DATA CENTRE NETWORK AND SPECIFICALLY THE TIME IT TAKES FOR DIFFERENT COMPONENTS TO COMMUNICATE WITH EACH OTHER.

Dave Hughes, Technical Director MEA, Commscope

Tarek Helmy, Regional Director – Gulf and Middle East, South & East Africa, Nexans Cabling Solutions, adds cabling needs to support the increasing need for speed and bandwidth. Looking at bandwidth needs in data centres, we have to split data centre connectivity into two segments: "switch-to-switch" vs. "server-to switch" links. Switch-to-switch connections is fibre rich, while in server-to-switch more copper solutions are used. Parallel optics enables us to achieve up to 100GB today with multimode fibres (OM3 & OM4). Two-lane singlemode is also possible but will be much more expensive because of higher transceiver costs. Upcoming copper standard of 40Gb/s over copper (40GBASE-T) will also enable us to see more copper ports replacing the fibre ports, while new speeds over fibre are currently being tested and are not yet confirmed, he says.

"WITH THE CONTINUAL REQUIREMENT FOR EXPANSION AND SCALABILITY IN THE DATA CENTRE, CABLING INFRASTRUCTURES MUST PROVIDE RELIABILITY, MANAGEABILITY AND FLEXIBILITY. DEPLOYMENT OF AN OPTICAL CONNECTIVITY SOLUTION ALLOWS FOR AN INFRASTRUCTURE THAT MEETS THESE REQUIREMENTS FOR CURRENT AND FUTURE DATA RATES. SWITCHING AND ROUTING, VIRTUALISATION, CONVERGENCE AND HIGH-PERFORMANCE COMPUTING ENVIRONMENTS ARE EXAMPLES OF WHERE HIGHER NETWORKING SPEEDS WILL BE REQUIRED."

Alberto Zucchinali, RCDD, Data Centre Solutions and Services Manager EMEA, Siemon

What is driving the adoption of optical networking in data centres?

"It is quickly become apparent that the use of copper cables for transmission of 40/100 GbE is questionable. For example, if you use lines seven meters in length, they will pose problems but so too does the laying of ten lines parallel to each other. Given the cramped space in cable runs and the difficulty of cooling them, the use of copper does not seem feasible. Fibre optic is the solution for high throughput required in data centre backbone," says Alfred Tharwat, Head of Training and Data Centre Consultancy, R&M Middle East, Turkey & Africa.

And the number of network connections in data centres is on the rise. Data centres have to achieve ultra-high density in cabling to accommodate

all this cabling in the first place. Multi-mode fibre optics is the medium of the future for satisfying the growing need for transmission speed and data volume over short distances. Parallel optics technology is what you get if you combine both trends – cabling density and the use of fibre optics. People need more throughput, less latency and high density and optical fibre media can provide such needs effectively, he adds.

Morris says customers planning a DC network typically look for a ROI period of 5-15 year and fibre exceeds both the bandwidth requirement and reach for effective deployment of 40G/100G Ethernet, 8G fibre channel (FC), 16G FC, 32G FC and beyond.

"There is a clear migration path to 40G/100G with a fibre optic network which is attractive to those looking for a strong ROI, for example, on their 10G networks deployment. Fibre is a medium that can meet or exceed all typical 'reach' requirements of DC network architecture(s) - this includes campus, back-bone, cross-connect, interconnects and direct-connects, these factors are key drivers with regards to the proliferation of optical networks," he says.

OM4 vs. OM3

When it comes to optical cabling in data centres, IT managers have the choice between OM4 and OM3. Which one should you choose? " In the short term bandwidth and ultimately reach have not been the driving factors to get to OM4, OM3 is more than sufficient and cost effective in the vast majority of applications. Higher performance and lower loss of OM4 becomes attractive for more complex architectures like Central Cross Connects, Zone Distribution and newer Optical Shuffle Topologies, and only a small number of cases where reach and bandwidth are the limiting factors," says Paul Kish, Director of Systems and Standards, Belden.

OM4 is recommended to provide the longest reach, support the highest data rates and to provide the greatest flexibility (increased number of optical connections). OM3 can be used with reach limitations. For example, the new 100GBASE-SR4 standard supports up to 100 meters over OM4 fibre and 70 meters over OM3 fibre.

Longer term we will see a change here, longer channels (up to 500m) in mega data centres will trend toward single mode with PSM4 and for the shorter links (<100m) the cost of the cable plant becomes a dominant issue so technologies like "Ultra Wide Band" that utilise WDM techniques over multimode will definitely challenge the bandwidth requirements of fibre leading to the potential need for "OM5" higher bandwidth glass, Kish adds.

"IT IS QUICKLY BECOME APPARENT THAT THE USE OF COPPER CABLES FOR TRANSMISSION OF 40/100 GBE IS QUESTIONABLE. FOR EXAMPLE, IF YOU USE LINES SEVEN METERS IN LENGTH, THEY WILL POSE PROBLEMS BUT SO TOO DOES THE LAYING OF TEN LINES PARALLEL TO EACH OTHER. GIVEN THE CRAMPED SPACE IN CABLE RUNS AND THE DIFFICULTY OF COOLING THEM, THE USE OF COPPER DOES NOT SEEM FEASIBLE. FIBRE OPTIC IS THE SOLUTION FOR HIGH THROUGHPUT REQUIRED IN DATA CENTRE BACKBONE"

Alfred Tharwat, Head of Training and Data Centre Consultancy, R&M Middle East, Turkey & Africa

Zucchinali says the first question should be what does a company wish OM4 to do? "If it is to join two buildings together and they have a need for 10Gb Ethernet then OM4 is applicable as it can support this easily up to 550m (250m more than OM3). If a company is looking to install OM4 inside a data centre where all the links are less than 100m, then OM3 will support the same applications as OM4 at this distance. There has been a lot of hype concerning OM4 for the data centre with various manufacturers supporting this, but with inter-cabinet links longer than 60m still being a rarity, OM3 can easily cope with the current distance requirements within the majority of data centres."

Another vexing question when it comes to optical fibre cabling is choosing between single mode and multi-mode. "There is no universal answer for this question, every data centre is unique and optical source technology is evolving at an incredible pace right now making the future landscape a complex picture," says Dwayne Crawford, Product Line Manager, Belden.

Tharwat from R&M says first it is important to understand the difference between the two. Multi-mode fibre has a core size of either 50 or 62.5 microns and is commonly used for connections between telecommunications rooms within a building or campus. Preferred for most physical-security applications, multi-mode uses low-cost light source for transmission such as VCSEL. It is easier to terminate and test than single-mode. The drawback of multi-mode though is its distance limitation for Ethernet applications, which is usually a few hundred meters. Single-mode fibre, with a core size of 8 to 9 microns, typically handles longer distances, sometimes more than 60 Km and is used in high-bandwidth applications.

"The choice is depending on the applications and the required cable distances. Single mode and multimode fibre optics can in practice support the

same data transmission speeds, but there is a clear difference in the data transmission distances. In data centre environments distances are usually small, and multimode fibre is mostly sufficient. However, single mode might be required also based on the applications and systems. Cost is another consideration as we have to keep on mind that the electronics used for single mode fibre optics transmission are having much higher cost than multimode electronics," says Helmy.

Tharwat adds multi-mode is sufficient in many cases and it can provide the required throughput of 10G over LC duplex and 40/100G over MPO/MTP connectivity. When designers need a more future proof solution, they use redundant links of single-mode. "In my opinion, they don't need to use them for the moment, but as an infrastructure they will be implemented for redundancy or future requirements."

The future of connectors

With the move to 40G/100G, industry experts anticipate a change in the connector requirements. Will MPO style connector be the standard for future equipment interfaces?

"MPO is expected to be around for many years to come and is already being considered by standards organisations for 128/156G and 400G. For 100G the MPO connector requirement is expected to change, the good news is that the change will be based on using the 12 fibre ferrule MPO connector with identical lane designation and fibre utilisation as 40GBASE-SR4 (8 fibres). The new 100G reference will be 100GBASE-SR4 (4x25G Tx, 4x25G Rx) and the standard is set to be ratified/published by IEEE during 2015," says Morris.

Hughes from Commscope agrees that MPO style parallel optics is seen as the de-facto standard for data centre fibre applications. "Not only are they quick to deploy, but they offer a structured upgrade path for future applications. A need for array connectivity is triggered by various IEEE standards that rely on use of parallel optics with MPO interfaces to support high bandwidth applications."

"THERE IS NO UNIVERSAL ANSWER FOR THIS QUESTION, EVERY DATA CENTRE IS UNIQUE AND OPTICAL SOURCE TECHNOLOGY IS EVOLVING AT AN INCREDIBLE PACE RIGHT NOW MAKING THE FUTURE LANDSCAPE A COMPLEX PICTURE."

Dwayne Crawford, Product Line Manager, Belden

Looking at the current standards, as well as those under development, as bandwidth gets higher so do the number of fibres used within the MPO. However, current development proposed the distribution of the bandwidth over a fewer number of cores, as the proposed MPO-16 interface for 400GbE. This has up to 32 fibres that are fully utilised, which increases fibre efficiency over previous configurations and reduces complexity, he adds.

A similar standards development path exists for Fibre Channel standards as well. Increase in bandwidth requires using various MPO interfaces. The FC standards are also allowing separation of higher bandwidth into lower bandwidth levels. For example Gen 6th standard will support breakouts of QSFP+ to 4 x 32Gb SFP ports.

Fibre vs copper

Does the increasing use of fibre in data centre mean copper will still have a role to play in high-density data centres?

"Though it is true that optical fibre can be used in high density panels. Copper connectivity has a significant cost benefit. With the current strong increase in 10GBase-T and the upcoming 40GBase-T standard, copper will remain a very attractive solution in DCs, especially for server-to-switch connections," says Helmy.

Hughes points out that Category 8, under IEEE 802.3bq, is a standard being developed for data centre applications, up to 30 meters. This is a cost-effective solution using conventional twisted pair cable and RJ45 connectivity for 40Gb/s transmission speeds and is a good fit for Spine/Leaf network architectures.

"That is the million dollar question," says Kish. "We are seeing clear decrease in the amount of copper being used in a data centre and there is no indication that this trend will reverse. Switch to Switch is being dominated by fibre but there is still a home for copper in the server to switch segment. The biggest question will be the future of CAT 8, being a non-optical interconnect cleaning is a non-issue making copper ideal for the day to day patching work of the IT staff and the 30m reach will be more than enough for 25G or 40G interconnects in this part of the network. The economics of Cat 8 will ultimately make the outcome clear, because of electrical requirements will it continue to use the cost-effective ubiquitous RJ-45 and at 30m reach the traditional enterprise LAN market is not a potential market to help offset the cost of development for the chipset vendors which may end up impacting what has traditionally been the most cost effective medium."

Major investments in Silicon photonics is the other headwind for copper, pushing the optics on to the same die as the CPU/GPU/FPGA/ASIC is opening up a whole new paradigm for optics with Moore's Law bringing a new Economic reality with the potential to significantly reduce the cost of optical solutions. Copper still has one ace up its sleeve and that is cleaning.... until fibre connectivity can resolve this age old problem, copper still has a home in the future.

THE LINK TO **GROWTH**

Tarek Helmy, Regional Director, Gulf and Middle East, South & East Africa, Nexans Cabling Solutions, talks about the latest trends in the regional structure cabling market and growth plans of his company.

Tarek Helmy, Regional Director – Gulf and Middle East, South & East Africa, Nexans Cabling Solutions

How's the business been this year?
The business has been steady and we have maintained revenues despite the volatility in the region. If you look at the region, we ended up having only 5-6 countries out of 14 because of what's happened in Lebanon, Syria, Libya, Egypt, etc. For us, the major markets are still Saudi, UAE, and Kuwait. We have invested heavily in the Saudi market with three people on the ground and this has contributed most to our growth.

As a company, we still maintain the same focus –

we focus only on big projects. We are into value, not volume. Some of our competitors straddle both the low-end and high-end of the market but it is not easy. We target the high-end market which needs end-to-end and value-added solutions.

Which segment of the market is growing faster – LANs or data centres?

I think the LAN business still dominates though there is big growth happening in the date centre segment of the market. A couple of years of back, the definition of data centre was a bit vague in this part of the world, but now you can see governments and operators investing heavily in data centre projects. In terms of value and revenue, the data centre market is still on top of the LAN market and we expect this business to grow really big in the near future.

Within data centres, do you foresee fibre ever replacing copper?

Within the data centre, it really depends on the customer's requirement and design. Whether it is top of the rack or end-of-row design. What we see is that most customers are still looking at a mix of fibre and copper. It is cost-prohibitive to go completely fibre in a data centre because you will have to spend 100 percent extra on active electronics and interfaces.

But isn't fibre the best bet if you are moving to 40G?

With the forthcoming Cat 8 cables, 40G can run on copper, though we don't have a good sense of what this new category of cables will look like and how it will impact the data centre infrastructure. The standard is still under development and we are hoping to get some direction on this from IEEE by the first quarter of next year. There are so many conflicting marketing messages out there but we think we are in a really good position. We have developed GG45 connector Cat 7 and 7A cabling, which was selected in the latest data centre standard from ISO and we have cables that support 2Ghhz transmission speeds. Although no one knows what type of connectors will be used in Cat 8, there is a good chance that they won't be RJ45.

Do you see the increasing popularity of wireless as a threat to structure cabling?

Yes, LAN is threatened a bit by wireless, especially with the new Gigabit WLANs. But, the issue is still security. LAN will be the number one choice for government and other security-conscious customers. You will see wireless being more popular in verticals such as hospitality but I don't think it will ever replace LAN.

Do customers prefer pre-terminated solutions now for their data centres?

Absolutely. The advantage of pre-terminated solutions is that it is guaranteed that there will be no problem with quality, you get test results of every link and there is no risk of mishandling. The only disadvantage is that you have to define the length from day one of design.

> LAN WILL BE THE NUMBER ONE CHOICE FOR GOVERNMENT AND OTHER SECURITY-CONSCIOUS CUSTOMERS. YOU WILL SEE WIRELESS BEING MORE POPULAR IN VERTICALS SUCH AS HOSPITALITY BUT I DON'T THINK IT WILL EVER REPLACE LAN.

Is OM3 still good enough or do you recommend OM4?

I'd say OM3 is more than enough because the fibre length in data centres will never exceed 100-150 metres. So the decision on whether to go with OM3 or OM4 depends on the required cable distance, where OM4 supports longer data transmission distances than OM3.

Singlemode or multi-mode?

Singlemode and multimode fibre optics can in practice support the same data transmission speeds, but there is a clear difference in the data transmission distances. In data centre environments distances are usually small, and multimode fibre is mostly sufficient. However, singlemode might be required also based on the applications and systems. Cost is another consideration as we have to keep in mind that the electronics used for singlemode fibre optics transmission have a much higher cost than multimode electronics.

Is the Middle East now a de facto Cat 6A market? And is this still a UTP dominated market?

Yes, that happened a year-and-a-half ago, and we rarely come across anyone asking for Cat 6 these days. Now, it is common to see a lot of projects specifically ask for shielded solutions because customers have realised they are more economical, easy to install and have a better bending radius. In fact, there is not even a single disadvantage. What's more, when it comes to Cat 6A shielded is much cheaper than UTP!

Nexans has recently launched an FttO solution. Is this an alternative to structure cabling?

It is a fibre to the outlet solution called LANactive that combines passive cabling with active mico-switches to provide an Ethernet service via standard copper based RJ 45 technology to the devices using PoE. It is targeted at specific verticals such as utilities, airports and hotels.

What are your plans for the next 12 months?

We are planning to aggressively push the LANactive solution because it provides customers with significant cost savings, especially those having restrictions in terms of capacity and space. We will also sell our wide portfolio of data centre solutions and high-end Cat 6A and Cat 7A cables. In terms of focus markets, they are still Saudi, UAE and other GCC markets. We are also closely following the developments in Egypt.

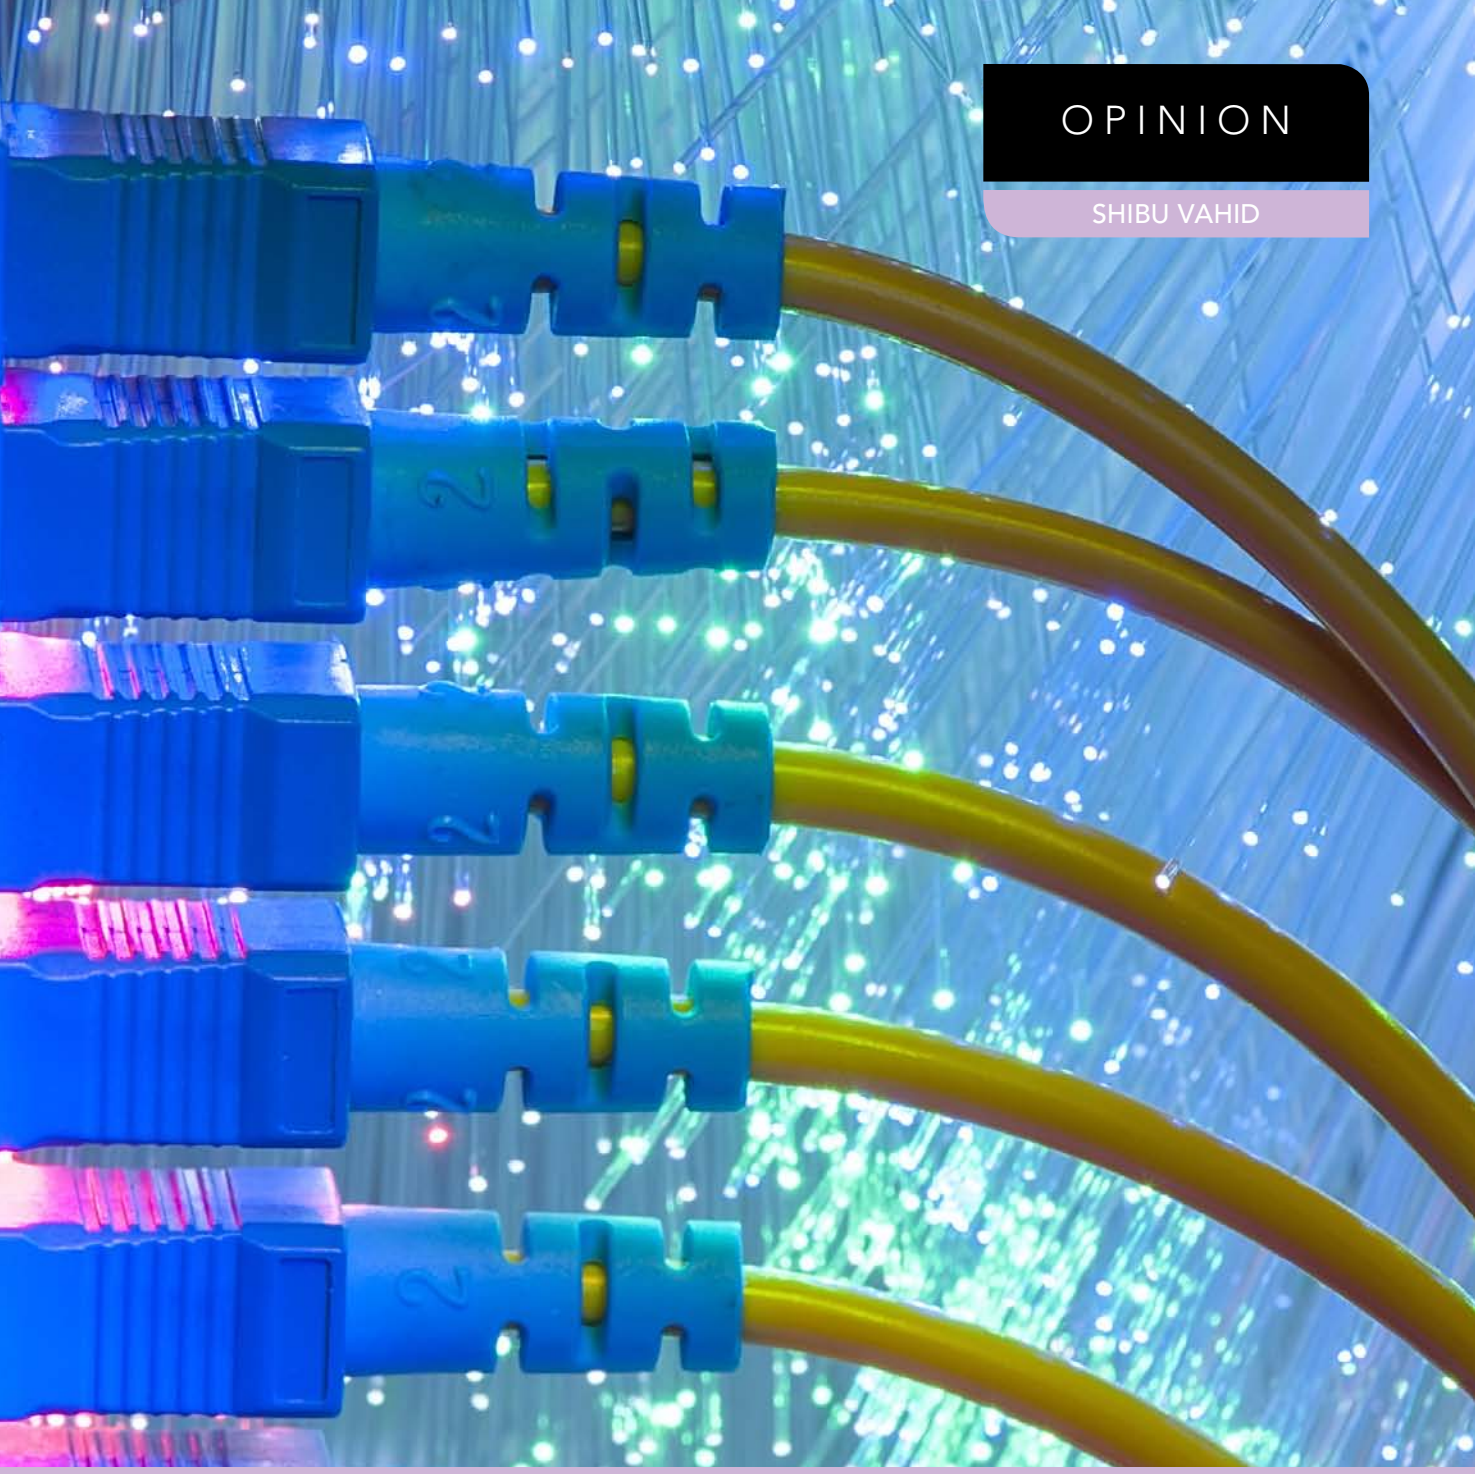

OPINION

SHIBU VAHID

TO BEND OR NOT TO BEND

Shibu Vahid, Head of Technical Operations, R&M Middle East, Turkey and Africa, writes about when does it make sense to use bend insensitive fibres.

Bend insensitive fibre can accommodate the vast amount of twisting, moving, adding and changing that takes place in confined sections of fibre rollouts. This technology can be used in very small spaces and cable cabinets, allowing data to move through cables at awkward angles without loss of quality. This reduces ownership costs, cooling, downtime and carbon footprint. However, bend insensitive fibres don't necessarily

bring an advantage to every situation. Sometimes the extra outlay outweighs the benefit. So how do you decide when to go with bend insensitive?

Last mile, SDU and MDU applications

Bend insensitive fibres can be configured and manipulated in extremely confined or irregularly shaped spaces, such as sections of cellars, basements, cabinets, pipes and more exotic channels such as sewers and air ducts. Fault tolerance and reliability are extremely high, so even in the case of severe twisting or connector misalignment, operation is maintained. In last mile FTTH cable deployments, where sharp bends are unavoidable, bend insensitive fibre offers a perfect solution. These fibres are also ideal for surface mounting, thanks to the ease of installation. The fibres can be safely bent and manipulated under practically any conditions with no loss of quality. This type of fibre is generally used in network configuration rollouts and is particularly practical for deployment in Single & Multi Dwelling Units.

One of the greatest benefits of bend insensitive fibres is the fact that the optical power budget is not changing, even when tight corners are involved. However, on a long outdoor stretch with no angles, bend insensitivity provides no extra benefits over 'regular' fibre.

Data centre applications

Besides FTTH networks, there is also a trend towards using bend insensitive fibres in data centres, for increased safety and to avoid any failure or downtime. In data centres, increases in port density and the need for greater system flexibility are leading to more and more patch cables and larger cable volumes under the floor and in the ceiling. Large numbers of cables are placed together in narrow spaces and frequently replugged in patchbays. Constant repatching and handling of 'traditional fibre' can result in failure. Network installations or moves, adds and changes (MACs) always pose some threat to cables, which can become compressed or pinched. Again, bend insensitive fibre provides a solution. Networks should be designed, built and maintained with the utmost care, but this isn't always the case. Fortunately, bend insensitive fibre is extremely forgiving although it is not a substitute for good cable housekeeping!

In a data centre, it is wise to use bend insensitive fibre only where it actually makes a difference. Cables between racks, for example, don't need to be bend-insensitive. This would only increase costs without bringing extra functionality. Conduits, pathways, trunking or cable tray can become overcrowded, but in these cases cables are generally installed according to tries-and-tested industry standards as well as strict manufacturer's guidelines describing bend radius limits. Bend Insensitive fibre cables can be spliced using the exact same methods used with 'regular' cables. Splice performance can be measured using 'standard' OTDR (Optical Time Domain Reflectometers).

In data centres, the real need for bend insensitive fibre is in patching zones, where patch cords can be tightly packed around tight bends and might be repatched repeatedly. Bad patching is a common cause of data centre failure. Pushing fibres around too many sharp bends, or corners which are tighter than the maximum bend radius, can lead to vast attenuation increases, making applications or even whole systems drop out entirely.

Shibu Vahid, Head of Technical Operations, R&M Middle East, Turkey and Africa

Some final considerations

If an application can be built entirely with bend insensitive fibres, it makes perfect sense to do so. This provides the best possible signal path and greatest ease of handling on smaller platforms such as outlets and building entry points. However, there are a few important factors to take into account. Looking at Public Networks, the most commonly used fibre types are standard fibre types, for example G.652.D. Of course, bend insensitive and 'bend sensitive' fibres are fully compatible, but there is no real advantage to using bend insensitive fibres throughout the entire network, because all the termination platforms are made for use with 'bend sensitive' fibres.

Bend Insensitive fibres are also often incompatible with regular fibres, especially when dealing with multimode. Core size measurement, NA, differential mode delay (DMD) and bandwidth had all been set before the introduction of multimode bend insensitive designs. Also, the sensitivity to bends becomes greater at higher wavelengths (such as 1625 nm).

ONE OF THE GREATEST BENEFITS OF BEND INSENSITIVE FIBRES IS THE FACT THAT THE OPTICAL POWER BUDGET IS NOT CHANGING, EVEN WHEN TIGHT CORNERS ARE INVOLVED.

DOING IT THE **RIGHT WAY**

This whitepaper from Fluke Networks examines how cloud based service can reduce costs through more effective cable test results management

Datacom installers certify cabling systems for many reasons – to support manufacturer warranty requirements, meet customer qualifications, or just to ensure quality workmanship. Managing the results generated by the testers is a critcal and time consuming part of the certification process.

In a recent Fluke Networks survey of 880 installers worldwide, the respondents reported installing an average of 1,026 links in just the prior month. The job of consolidating all those tests into a single report can be time consuming and is a source of considerable cost. In another survey, 77% of installers reported having to deal with results management issues in the prior month. Let's look at these issues.

Results management issues

Customers report a number of problems in results management that take up time. The following three are the most common:

Cable identifers in the reports don't match the specifications and need to be manually edited

Multiple test types (copper, fibre tier 1, tier two, different standards) need to be combined into a single report

Reports have to be reworked because they are incomplete

Contractors report spending 7.9 hours per month, nearly a ful day, on these isues.

This is in addition to the time required to get the results (typically stored in the tester) from the job site back to the office where the results management takes place. This also takes a tester out of circulation which may impact other jobs.

The more testers you have, the more complex it gets

For a contractor using a single tester on a single job, these problems are signifcant. But very few contractors have the luxury of dedicating a single tester to a single job at a time. To maximise eficiency, most contractors own or rent multiple testers and use them on multiple concurrent jobs.

Properly keeping track of results under these conditions become even more difficult, leading to a whole new set of issues including:

- Test results are stored in multiple testers which have to be hunted down and the results consolidated
- Tests from another job get into the report by mistake and have to be removed
- You're generating a report and you discover not all the links were tested, and a crew has to return to the site to finish the job

Installers report spending, on average, another 7.3 hours per month on these issues, for a total of nearly two days a month. And the more testers owned by a firm, the larger the scope of the reported problem. A related issue cited by customers is the lack of visibility in keeping track of job status under these circumstances.

Worst-case scenario

Other, less common, problems can be far worse. With tens of thousands of testers on the market, our tech support team sees a lot of problems that an individual customer might not consider. For example, we regularly receive calls asking how to recover inadvertently erased or corrupted memory cards which store the results. Unfortunately, in most cases, the only answer is to perform all the tests over again.

77%
of installers reported having to deal with results management issues

Testers or memory cards with results stored on them can also be lost or stolen, also requiring retesting. Thankfully, these are not common problems, because a single lost tester or erased data card can mean losing a full day's or even a week's worth of multiple techs' time – dozens of hours lost and long delays in the project.

> WITH THE ADVENT OF HIGH SPEED MOBILE NETWORKS AND LOW COST DATA STORAGE, A NEW APPROACH TO MANAGING AND STORING DATA – THE CLOUD – HAS BECOME COMMONPLACE.

Cloud services offer a solution

With the advent of high speed mobile networks and low cost data storage, a new approach to managing and storing data – the cloud – has become commonplace. By getting results out of the tester and up to the cloud, the chances of losing data are minimised. Testers don't need to be recalled from the field. Results can be automatically consolidated into the correct job.

LinkWare Live adds a cloud capability to the industry's leading reporting solution and testers. Versiv testers can be connected to the internet through a wired Ethernet or Wi-Fi connection. In instances where no Ethernet or Wi-Fi network connection is available, the Versiv unit can connect wirelesly to a smart phone configured as a personal hotspot to send results to the LinkWare Live service.

Once connected to the network, results can be uploaded with just a few keystrokes. Tests in Versiv are tied to projects, so, once they get to the cloud service, it adds them to the appropriate project automaticaly. Since the proces is so quick, it can be done many times a day, so that the impact of a card failure or stolen tester is minimised.

When the time comes to generate the report, the contractor simply downloads the results from LinkWare Live into the PC running LinkWare. From there, the status of the project can be checked, and missing or improper results immediately identified. A full report can be generated with a few keystrokes.

Since LinkWare Live understands the projects set up within the Versiv tester, it can keep track of the status of projects.

And since it's cloud-based, this information can be accessed from anywhere by tablet or smartphone. This gives the project manager or consultant a near-realtime view into the status of the project and the ability to drill down into individual test results to ensure they are done properly.

Conclusion

Managing results through a cloud based service offers significant advantages:

- Increased productivity by no longer needing to recall testers from the field just to download test results
- Reduced time by automatically consolidating al results into the correct job
- Less rework by reducing the likelihood of losing test results when testers or memory cards are lost, stolen or erased
- Instant access to results for faster troubleshooting
- Real time visibility into project status from any location
- Using a cloud based service is a sensible approach that can prevent problems and save time.

You save your work on your PC every few minutes, so why would you save your cable test results only once a day?

BLOG

PAUL KISH

HERE COMES 25G ETHERNET

Paul Kish, Director of Systems and Standards, Belden

You may have heard plenty of buzz over the past few months regarding 25 gigabit Ethernet. It's no surprise considering that several big data centre and cloud computing providers like Google, Microsoft, Broadcom, Arista and Mellanox formed the 25 Gigabit Ethernet Consortium earlier this year.

Shorter thereafter, IEEE formed a 25G Ethernet study group. And just a couple of weeks ago, Broadcom announced the availability of a new high-density 25G Ethernet switch for cloud-scale data centres.

Why 25G? It simply makes sense from a technology, cost, scalability and flexibility perspective. Let's take a closer look.

We're already there

Existing 100G standards such as 100GBASE-LR4 and 100GBASE-ER4 already are comprised of four lanes of 25G over singlemode fibre using course wavelength division multiplexing (CWDM) technology. The 100GBASE-CR4 and 100GBASE-KR4 standards for 4 X 25G over twinax and backplane were also recently ratified, and the 100GBASE-SR4 standard with 4 X 25G per lane over multimode fibre is also well underway.

As mentioned in one of my blogs earlier this year, 100 gigabit QSFP28 SR4 optical transceiver modules are already available for this application. With the same footprint as the 40G QSFP+ for four 10G lanes, the QSFP28 essentially offers 250% more density. It also supports direct-attach copper (DAC) cable assemblies, active optical assemblies and transceivers.

With the 4 X 25G per lane technology already available for 100G, the process to go to a single 25G lane is a simple one that requires very minor changes and significantly reduces cost compared to 40G using four 10G lanes.

Ideal for the Cloud

One of the key drivers for 25G is cloud computing. Not only are cloud providers looking to decrease cost and improve density, but improvements in server speeds has rendered 10G no longer fast enough.

Many of these environments have standardised on a top of rack (ToR) topology, and keeping up with these connections at 10G requires more switches in the rack and therefore higher cost.

25G also works from a scalability standpoint. With uplinks migrating to 100G using the same type of technology, it makes sense to increase the switch-server speed at the edge to 25G. With cost benefits over 40G and the technology already gaining traction, a recent five-year forecast by Dell'Oro Group predicts that 25G will take over Ethernet server port sales by 2018.

Beyond Top of Rack

While 25G over twinax DAC assemblies will fulfill the ToR server environments where a distance of 3 to 5 meters is more than adequate, there is also the need for longer distances to support middle of row (MoR) topologies to about 15m and end of row (EoR) to 30m. That's where a future 25GBASE-T application over balanced twisted-pair copper cabling has potential to fill the gap.

At the September meeting of IEEE 802.3, the Task Force developing 40GBASE-T (P802.3bq) discussed the technical feasibility and the benefits of 25GBASE-T relative to 40GBASE-T. As a result, a Call for Interest (CFI) is scheduled for the next IEEE 802.3 meeting in November to form a study group to explore 25 Gb/s BASE-T Ethernet and potentially extend the work of P802.3bq Task Force to include it. Some considerations will include maintaining backwards compatibility with 10GBASE-T and forward compatibility with 40GBASE-T through auto-negotiation, as well as keeping the power per port low and supporting Energy Efficient Ethernet.

When it comes to home run applications where servers and storage devices are consolidated into separate areas of a data centre, the distances required increase to about 50 to 100m. Multimode fibre technology based on four lanes of 25G that is used in the upcoming 100GBASE-SR4 standard supports 100m distances over OM4 fibre. A single-lane 25G multimode fibre PHY could be a significant solution to support home run applications and an overall broader range of 25G architectures.

Stay tuned

While standards groups are actively focusing on 25G Ethernet over twisted-pair, twinax and fibre, it's still too soon to tell how it will shape up. Regardless, 25G Ethernet has broad market potential in the server interconnect world because it's both economically and technically feasible—on a variety of media. In other words, 25G Ethernet is definitely coming—and it's coming fast. The experts at Belden will keep you posted as the situation develops.

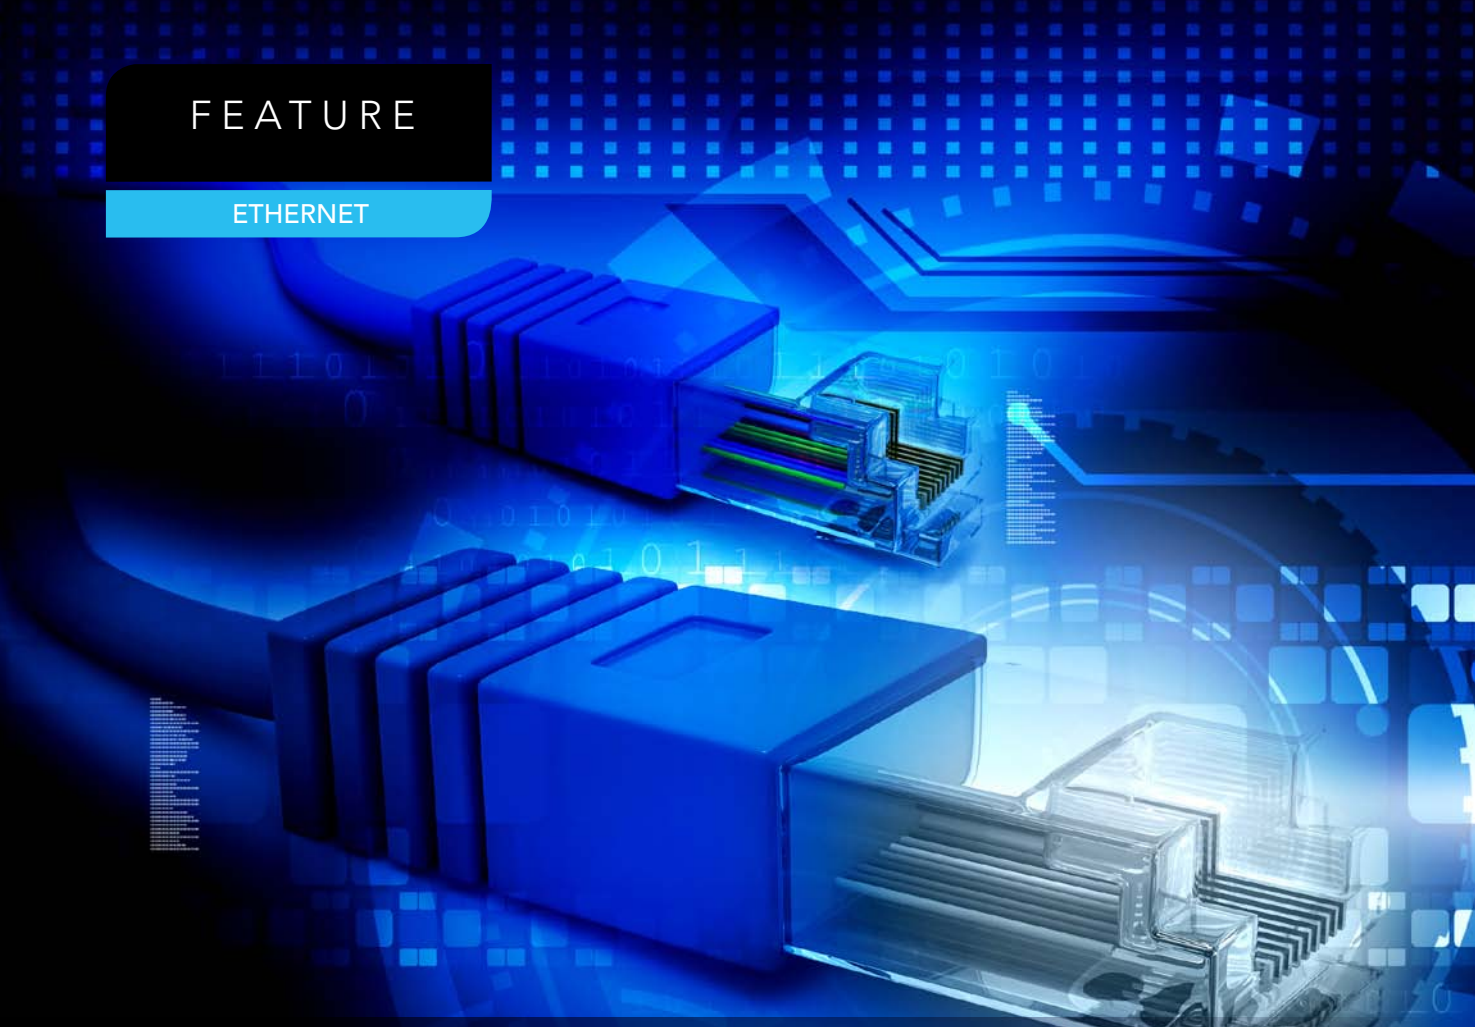

MOVING TO 40/100G ETHERNET

Advancing from 10G to 40G or 100G Ethernet not as simple as swapping out a few switches or line cards

Several factors have to be weighed, such as synchronising switch clocks for the higher-speeds, especially among multivendor equipment; ensuring latency remains at acceptable levels; keeping the network design and architecture optimal for 40/100G; and making sure the existing cabling infrastructure can accommodate the 4x to 10x increase in bandwidth.

One of the caveats that users should be aware of as they migrate from 10G to 40/100G Ethernet is the need to ensure precise clocking synchronisation between systems – especially between equipment from different vendors. Imprecise clocking between systems at 40/100G – even at 10G – can increase latency and packet loss.

The latency issue is a bigger problem than most people anticipate, industry experts say. At 10G, especially at high densities, just the smallest difference in the clocks between ports can cause high latency and packet loss. At 40G, it's an order of magnitude more important than it is for 10G.

This is a critical requirement in data centres today because a lot of the newer innovations are meant to address lower latencies.

"Where you're going to have the biggest challenges will be different latency configurations if RDMA (remote direct memory access) is used," says Shaun Walsh, Emulex senior vice president of marketing and corporate development. RDMA is a low-latency, high throughput data transfer capability where application memory is exchanged directly to and from network adapters without copying it to operating system buffers.

You see a lot more in-rack virtual switching, VM-based switching that is very application specific," Walsh says. "New line cards in new backplane architectures mean different levels of oversubscription. There'll be generational tweaks, configuration 'worrying' that has to occur. The biggest thing (testers) are running into is making sure you get the 40G you are paying for (with regard to) latency issues, hops, and congestion visibility."

Emulex acquired Endace, a developer of network performance management tools. Demand for the Endace product and the 40G capabilities of Emulex's XE201 I/O controller are picking up as more data centres and service providers upgrade from 10G to 40G.

Walsh expects 40G Ethernet to be a $700 million market in four to five years, roughly half the time it took 10G Ethernet to reach that mark. Driving it are next-gen blade server mid-plane interfaces and architectures, big data, analytics, video and data over mobile, BYOD and high frequency trading, Walsh says.

Another challenge is readying the cabling infrastructure for 40/100G, experts say. Ensuring the appropriate grade and length of fibre is essential to smooth, seamless operation. This is a big consideration for users because it could mean re-wiring a significant portion of their physical plant, if not all of it. That could be an expensive and disruptive undertaking.

At the physical layer, 40G Ethernet is essentially 4x10G "lanes." But 100G Ethernet is 4x25G lanes, which will be disruptive to the 10G and 40G infrastructure. "100G is going to be even more of a challenge because now you're dealing with a whole new layer of physical infrastructure," Walsh says. "You will have a whole new generation of optics, cables, everything will be a whole new generation at that point."

Moving bits four to 10 times faster error free is a challenge in and of itself. Making sure the higher level systems – routers and switches – deliver services and non-disruptive service quality at those speeds is equally as challenging, if not more so. Each device has to do this at one-fourth or one-tenth the time it does at 10G. For a router, it means performing all of the packet inspection, queuing, lookups, filtering, policing, prioritisation, table updating and logging while meeting SLAs by not dropping or reordering packets, or increasing latency or jitter.

"Routers aren't just forwarding packets," says Scott Register, senior director of product management for Ixia, a maker of test, measurement and visibility tools. "There's carrier grade NAT, application ID, and security processing and things like that. One of the more interesting testing is, what services can you enable at that rate before you start having problems?"

With carrier grade NAT, the problems get harder as that traffic load increases, Register says. In addition to throughput, increased session capacity and more concurrent connections are also issues as bandwidth climbs from 10G to 40/100G.

THE COMBINATION OF 40GBE DATA RATES WITH THE OPTIMISATION POSSIBLE THROUGH INTELLIGENT INFRASTRUCTURE IS A POTENT ONE. IT ENSURES MINIMAL RISK AND MAXIMUM PERFORMANCE OF IT SYSTEMS WELL INTO THE FUTURE

"You don't get more TCP or UDP ports just because you have more traffic," Register says. "How smart are you at cycling through your NAT tables? Some of the challenges that might not show up at the lower speeds show up at higher speeds. So a lot of the testing that we see is around that kind of high level stuff."

And that's pre-deployment testing. Post-deployment presents its own set of challenges, especially in enterprises with strict auditing and compliance requirements. "Increasing security requirements for compliance, recording e-mail correspondence… it's easy at 1G; at 10G or 40G or 100G it's really, really difficult," Registers says. "So if you want to do that you have to be able to do that kind of intelligent load balancing and filtering out of any unnecessary data so that your tools can keep up with that bandwidth."

That challenge is exacerbated by existing filtering and analysis tools that only run at sub-10G speeds, Register says. Tapping a 40G link would require 30 or so such tools each monitoring a slice of that 40G traffic, he says.

Ixia offers a switch that sits between the production network and the analysis tools that does just that. "They would put our switch between the production network and their tools to do the filtering – ex., seeing only the webserver traffic – to only see a very small subset," Register says. "I can take the input stream and balance it across 32 analysis tools. I can round robin or spread that traffic across a bunch of connected tools so you can preserve your investment with existing tools. There aren't many analysis tools that'll run at 40G and certainly there's nothing that runs at 100G. So a challenge we'll have is maintaining visibility into that traffic when they do their infrastructure upgrade."

And then, of course, there's always multivendor interoperability challenges at 40G. In addition to clock synchronisation between systems, adding features or applications could test the stability of that interoperability, Walsh says. "Where we're going to have problems is when people try to implement different features," he says. "Base Ethernet connectivity will work fine – but where we'll start to see challenges will be in RDMA, when you do lower latency stuff around RoCE (RDMA over Converged Ethernet). And when you load a special version of an SDN application that is trying to meet a specific need. When that SDN switch is plugged into the general Ethernet population, will it interoperate exactly right?"

Cat8
is the path forward for supporting 40GBASE-T

NEXANS DEBUTS LANACTIVE

Nexans has launched LANactive, a Fibre-to-the-Office (FttO) solution in the Middle East. The new LANactive solution offers customers a new alternative approach to office networking by using passive fibre cabling and components together with locally distributed active micro-switches to provide Ethernet services via standard copper based RJ45 technology to the device.

Nexans' LANactive solution offers significant cost savings and benefits in specific types of environment such as hospitals, universities and airports where:

- long transmission distances within the building are required
- space and/or cost restrictions limit the use of floor distributors
- cable containment has limited capacity
- refurbishment is required with minimum disruption
- redundancy at user level is required

All components including the active switches, cabling and connectivity are designed and manufactured by Nexans. The switches themselves are manufactured by Nexans in Germany which has a 20 year history of supplying a broad customer base and includes universities, hospitals and public administration facilities where the solution has been proven to demonstrate consistent reliability over many years.

R&M LAUNCHES COMBIMODULE

R&M has launched its new CombiMODULE for 19" racks which is a new scalable, easy-to-assemble distributor solution. The broad-based usage scenarios for the versatile solution include distribution projects in larger buildings of all kinds, in data centres and in exchanges.

The CombiMODULE fits in street cabinets, where it supports full coverage fibre optic distribution in cities and residential areas. The housing occupies ten units in the 19" cabinet and the module accommodates up to 1152 spliced fibre optic connections or up to 576 Small Form Factor (SFF) or 288 Big Form Factor (BFF) plug connections. Quick mounting technology simplifies mounting from the front in standard distribution frames, whether for initial installation or for existing racks.

The two-part housing has room for any combinations of splice and patch units to connect and distribute optical fibres reliably. For splice cabling, R&M offers its single circuit management (SCM) system. It consists of easy-to-use trays for the management of fibres. The carrier tray can be folded out and serves as a work table during installation and maintenance. A splice unit can hold as many as 96 SCM trays for 1152 fibre optic connections.

The patch unit is the counterpart to the splice unit. It holds up to twelve inserts to accommodate a total of 576 plug connections. All types of coupling usual in the market can be used, e.g. LC duplex, SC, E2000 SCRJ or E2000 TM* compact. R&M also has inserts for processes involving splitter applications and multiplexing such as CWDM or DWDM. R&M makes available all variations completely pre-terminated as well.

SIEMON EXPANDS ICEPACK COOLING DOOR RANGE

Siemon has extended its range of IcePack cooling doors in the Middle East to match its expanding range of data centre cabinets. The IcePack passive cooling doors are now available in options to fit Siemon's VersaPOD, V800 and V600 cabinets in both 45 and 42U heights. The thermal efficient heat exchangers replace the rear doors of Siemon's cabinets and use passive liquid cooling technology to target the hottest data centre spaces. Siemon claims that the IcePack range provides a total cooling capacity of up to 32kW per cabinet and consumes up to 80 per cent less energy than other data centre cooling options.

The IcePack system works by close-coupling a specialised fin-and-tube coil array that absorbs and cools equipment heat exhaust, thus cooling the most challenging hotspots and providing protection against future heat build-up. IcePack is a completely sealed, low pressure system and is designed to operate reliably above the dew point to avoid condensation.

FLUKE networks®

HOW DO YOU KEEP YOUR DATA FLOW UNDER CONTROL?

- ✓ Network Performance & Diagnostics
- ✓ End to end Infrastructure Visibility
- ✓ Client to Cloud Troubleshooting
- ✓ Automated Network & Application Analysis
- ✓ Wireless Security & Performance Management
- ✓ VoIP performance management

Visual TruView™ Performance Monitoring for todays IT

> Rapidly identification of network performance issues 24/7

> Application Performance Monitoring

> Network Performance Monitoring

> Distributed VoIP Performance Monitoring

For more information please contact:
meat-marketing@flukenetworks.com | www.flukenetworks.com/apps/truview

LANactive

Switch to the future

Intelligent
Fibre To The Office
Miniswitch

FTTO Active & Passive Solutions

Nexans is pleased to announce LANactive, an alternative approach to structured cabling. Using Fibre To The Office (FTTO) topology together with access switches installed near to the work place, it provides Ethernet services via standard copper based RJ45 technology to the device.

The approach offers significant cost savings and other benefits in specific circumstances:

- **Long distance transmission**
- **Eliminating the costly Floor Distribution (Space, Switching & Passive accessories)**
- **Reduced cable containment**
- **Refurbishment with minimum disruption**
- **Redundancy at user level**

Your 1GB PoE+ port at your desk

sales.middleeast@nexans.com
www.nexans.com/LANsystems

Nexans
Global expert in cables
and cabling systems