Light into Money

The Future of Fibre Optics in the Data Centre Networks

Author: **Dr Rick Pimpinella**

Contents

Dr Rick Pimpinella
Leading researcher
into fibre optics
Panduit Corporation

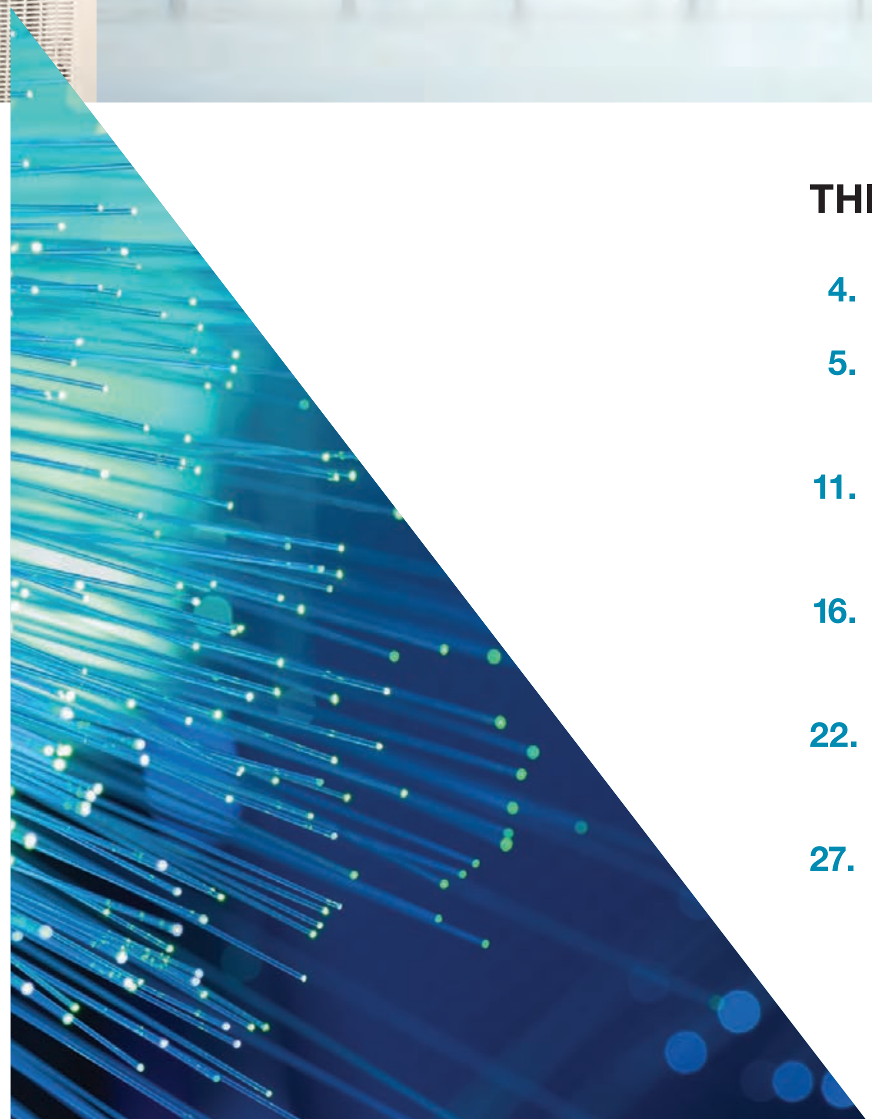

THIS WHITE PAPER CONTAINS

4. **Exec Summary**

5. **Section 1** Demand Drivers for high speed Ethernet

11. **Section 2** Traffic Growth Forecasts – Infographics

16. **Section 3** Current and Future Road maps for high speed Ethernet fibre optics

22. **Section 4** How Panduit is Future Proofing the Data Centre Network

27. **Conclusion**

Executive Summary

This business and technical White Paper is an exploration of the data centre fiber optic networking infrastructure requirements needed to meet current and future demands for data volumes and data rates. It covers how 200Gigabit (Gbps) and 400Gbps Ethernet (GE) fiber optic technologies evolved and how they should advance to 800GE and 1.6TE.

It also explores what the right technologies are for sustainable investment strategies to future proof data centre networks.

Setting the right standards is vital for infrastructure development and network design. International standards allow data centre owners to create stable platforms to use today and to plan for future migration paths. Standards are the foundation that protects investments which support speed and performance requirements while supporting new data intensive technologies and applications. These capabilities are vital for and demanded by organisations including enterprises undergoing transformation and web scale data centre operators.

Section 1:

DEMAND DRIVERS FOR HIGH SPEED ETHERNET NETWORK INFRASTRUCTURE

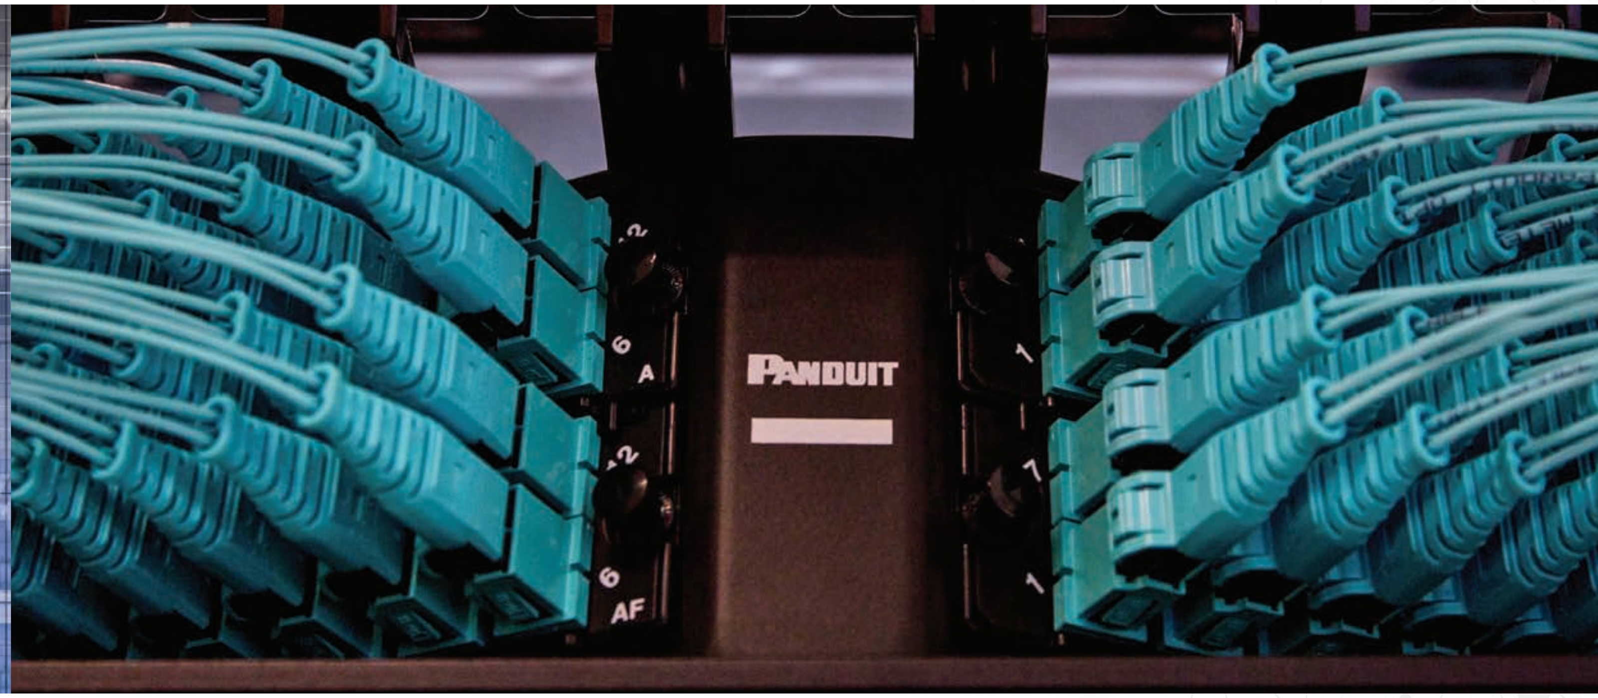

Summary

- The demand comes from:
  - Hyperscale data centre investment, generating innovation in construction, equipment and processes
  - Hyperscale data centre operations are influencing and driving changes in the network and infrastructure
  - Vertical market disruption including Telco and Media markets data centre expansion
  - Infrastructure requirements for Artificial Intelligence (AI) and Machine Learning (ML), which require access to and processing of petabyte levels of data
- New high-speed technologies including Application Specific Integrated Circuits (ASICs) and Field Progammable Gate Arrays (FPGAs) being developed for AI and ML
- Next generation network demands
- Increase in scale and operation demand faced by data centre network architects

Where demand is coming from

Super trends driving exponential data centre network expansion and demand can be found across Enterprises, Web Hyperscale giants and Communication Service Providers.

Enterprises in financial services, media, e-health, and retail continue to experience game changing digital disruption from which no industry is immune and this is spurring on data centre investment.

Data centre network architects tasked with providing physical infrastructure face exponential scale demand for bandwidth and operations. The evidence for this can be found in continued large scale networking port shipments. In line with previous years, 2018 will see the number of networking ports shipped to the enterprise market measured in the hundreds of millions.

One catalyst is that for many traditional businesses, the web scale firms represent massive market disruption and new competition. Enterprises are responding by deploying capital into digital infrastructure.

Digital growth forecasts anticipate that IoT and Edge Computing applications will generate petabytes of data on a daily basis by the year 2020. Machine to machine (M2M) data (which is data that is processed within the data centre) is expected to outstrip enterprise and personal data within five years. Data intensive technologies such as Artificial Intelligence (AI) and Machine Learning (ML) become effective when data sets reach 100s of terabytes. The bigger the data set, the better AI and ML models perform. AI models running half a petabyte of data are not unusual. See **note 1.1** Microsoft machine learning.

As scale and new technologies are being accommodated within the data centre, existing enterprise applications such as Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) and newer applications such as transactional connected commerce (mobile and fixed) are generating data at higher rates and volume then ever before.

Each of these technology advances reflect the rapid growth of the global digital economy that is creating demand for greater network speed and performance from the Internet backbone right into the core of the data centre.

Established businesses competing directly with the web scale firms cannot afford to be bound in by legacy technologies. The only response is to build new platforms and next generation Internet Protocol (IP) infrastructure. In almost every sector, digitalization of services combined with data capture at the multi-petabyte scale is driving the need for robust, high speed network infrastructure.

Hyperscale data centre investment drivers

At the web scale data centre operations level Google, Microsoft, Facebook, Amazon, Apple and the Asian giant public cloud players Alibaba and Ten Cent are deploying capital at an accelerating rate. The race is on to build fleets of hyperscale hubs as the backbone of the Internet.

By one estimate, across the world there will be more than 500 hyperscale data centres built and operational by 2020. In the last quarter of 2017 infrastructure spending reached $22bn among a group of Tier 1 cloud platform providers. See **note 1.2**.

The global hyperscale data centre market size is expected to grow from US$25.08bn in 2017 to US$80.65bn by 2022, at a Compound Annual Growth Rate (CAGR) of 26.32%. See **note 1.3**.

The networking infrastructure requirements for these facilities is dictated by the applications they serve and the high throughput they demand. (See below: The Future of Ethernet in Section 3)

The challenges for the network are ever growing demand for faster speed (40GE, 100GE today and 200GE and 400GE within three years) and the physical distances to be covered.

The server stack requiring connectivity is changing. 25GB servers are already commonly deployed in web platform data centres. As a result, cable and switching fabric must be able to handle the roll out of 50GB server technology which is already underway.

Google has said it faces a 100% increase in demand for data centre networking every 12 to 15 months. Facebook has declared that in terms of internal network traffic 'what happens inside the

Higher Performance Data Rate Sales Growth up to 2022

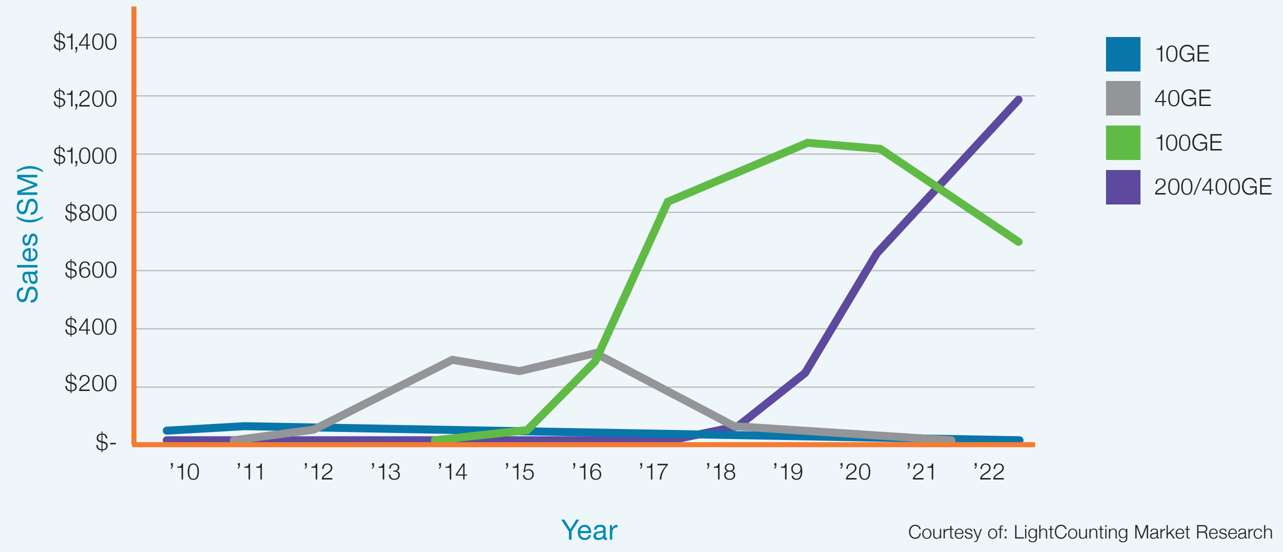

Courtesy of: LightCounting Market Research

Facebook data centres – "machine to machine" or also called East West traffic – is several orders of magnitude larger than what goes out to the Internet.' See **note 1.4**.

Wide area interconnectivity investment

Google has acquired a fully diverse dark fibre ring connecting the two main subsea landing stations in France to Paris. The creation of such a ring is a first in France, bringing multiple benefits to Google and its partners: Thanks to this ring, Google secured an initial 28Tbps of traffic, which ensures a very high reliability of Google services in Europe and better end user experience. Google has also invested similarly in the UK, Ireland and Denmark. See **note 1.5**.

More traffic means the need for more capacity requirement.

Telco and Media Transformation

In the telco sector, transformation to IP has begun in every Communication Service Provider. Many have committed to shifting away from PSTN to IP based networks by 2025. This places the data centre and its networking infrastructure, and not the old telephone exchanges, at the heart of the platform for their new services.

In a market which is becoming accustomed to price pressure and increased competition from the Over the Top (OTT) players, telcos are rolling out new core data centre infrastructure which is built on IP.

Telcos must also react to another opportunity in the shape of 5G networks. These are expected to generate a 'big bang' of bi-directional traffic as 5G upload speeds match the download speeds of today's 4G and LTE networks. 5G download speeds will be in the 1-10Gbps range with uploads

Note: 1.1: https://blogs.technet.microsoft.com/machinelearning/2017/05/04/monitoring-petabyte-scale-ai-data-lakes-in-azure/
Note: 1.2: (Synergy Research http://www.datacentreknowledge.com/cloud/hyper-scale-data-centre-spend-was-20-percent-2017-analysts-say)
Note: 1.3: http://www.transformingnetworkinfrastructure.com/news/2018/01/31/8692081.htm

Note: 1.4: https://code.facebook.com/posts/360346274145943/
Note: 1.5: https://static.googleusercontent.com/media/www.google.com/en//about/datacenters/eustory/report/Google_EU-DCs_Report.pdf

High Data Speeds Requires Optimised Infrastructure

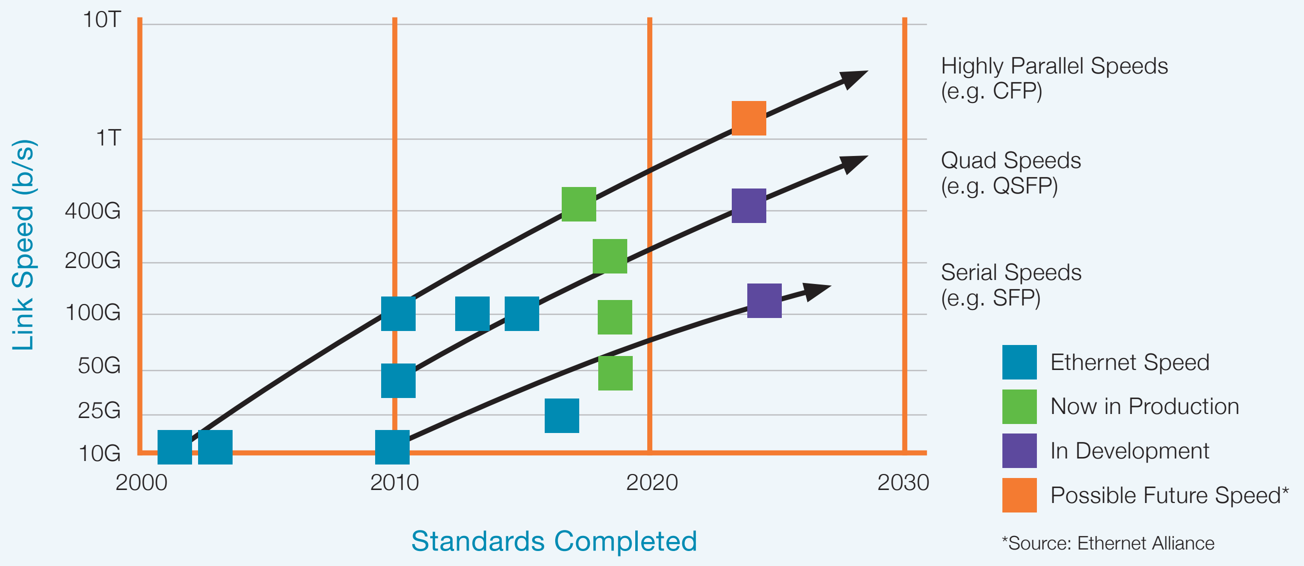

as fast as 1Gbps. Today's data centre networks are not built to transport the volumes of data that will be generated or provide the throughput and 1ms latency required.

The implications for telco operated data centres are huge.

In media, the transformation to IP as the standard platform for content creation, production, management, transport and delivery has already been agreed. This has also been ratified by The Society of Motion Picture Technology Engineers (SMPTE) and the IEEE and is rolling out across the sector.

Whether hosted in the public cloud, (Netflix is hosted on AWS) in commercial data centres or on-premise, the data volume growth for 4k and soon 8k video packets, which will be stored and moved around in media data centres, is about to explode. One hour of 4k raw uncompressed video can be as big as 250GB.

To varying degrees, all businesses including the web players themselves, are being digitally disrupted and the rate of change is accelerating.

New digital ecosystems on web platforms, in telco, media, financial services or connected healthcare are being built on high performance data centre networked infrastructure, as the foundation of new value creation for the business.

Artificial Intelligence and Machine Learning infrastructure needs

AI and Machine Learning (ML) require machine-to-machine communication at networking speeds and data volumes that have serious implications for network infrastructure topologies and connectivity. (See Spine Leaf architecture diagram below in Section 3).

By their nature AI and ML solutions rely on large data sets (preferably hundreds of terabytes and ever rising) to 'learn'. The data volumes for ML already require 50GE and the demand curve is forecast to rise sharply to 100GE and beyond.

Of its ML efforts, Facebook says "one salient feature of machine learning at Facebook is the impact of the massive amounts of data that is potentially available to train the [machine learning] models. The scale of this data has many implications that span the entire infrastructure stack."

Facebook's view is that hardware limitations can result in an unacceptable increase in overall training latency and time to convergence. Distributed training is one solution for overcoming these hardware limitations and reducing latency. "This is an active research area not only at Facebook, but also in the general AI research community… in our work on distributed training, we have found Ethernet based networking to be sufficient, providing near-linear scaling capability. The ability to scale close to linearly is closely related to both model size and network bandwidth."

"An optimization metric is network usage, where we can see the data traffic generated by machine learning training can be significant and sometimes bursty. If not handled intelligently, this can easily saturate network devices and even disrupt other services."

Facebook has already said that its machine learning developments "will result in …an increased network bandwidth for data access as well. So, significant local/nearby storage is required to allow off-line bulk data transfers from distant regions to avoid stalling the training pipelines waiting for additional example data." See **note 1.6**.

Note: 1.6: https://research.fb.com/wp-content/uploads/2017/12/hpca-2018-facebook.pdf

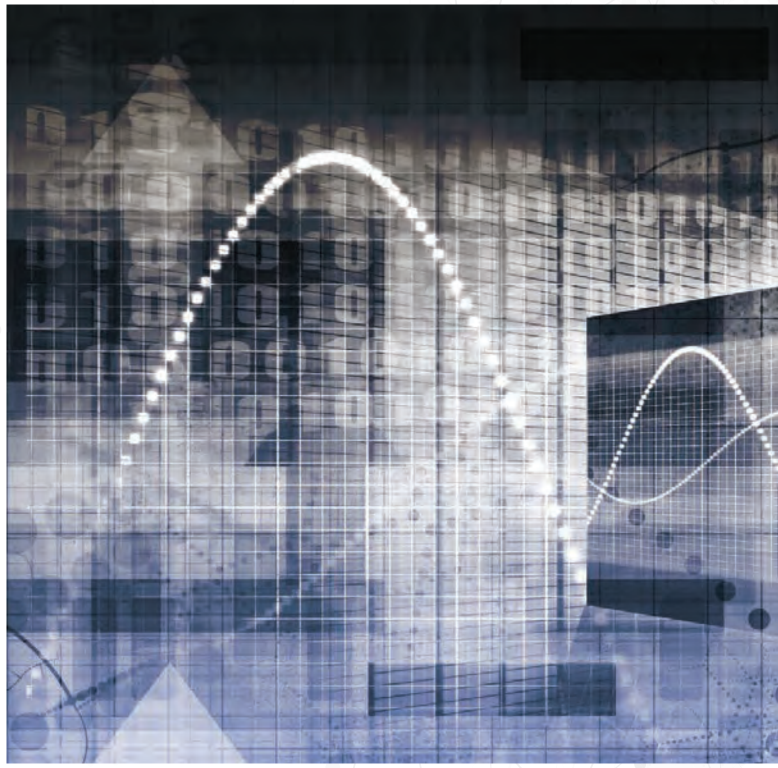

New GPUs and CPUs for Machine Learning

Already new 32GE, 50GE and 100GE ASICs (Application Specific Integrated Circuits), FPGAs (Field Programmable Gate Arrays) and GPUs (Graphics Processing Units) are being deployed for AI. 200GE, 400GE and 800GE ASICs and GPUs are already in development.

Google's efforts include Tensor Processing Units scaling to 100 petaflops (1 petaflop = 1000teraflops – floating point operations per sec). These chips are so powerful that Google said it has had to introduce liquid cooling into its data centres for the first time.

Tensor Processing Units (TPUs) are Google's custom-developed ASICs used to accelerate machine learning workloads. TPUs are designed

from the ground up with the benefit of Google's in-depth experience and leadership in machine learning. Offered as a cloud service, 'Cloud TPU' enables users to run machine learning workloads on Google's second generation TPU accelerator hardware using Tensorflow. See **note 1.7**.

Next generation network demands

All of this creates new demand for network throughput, availability, scalability, low latency and robustness.

It is also challenging data centre network architects to dynamically meet a growing diversity of different communication patterns across the network with different traffic classes having different networking requirements for throughput flows, deadline flows, and interactive flows.

The common requirement is that cable infrastructure will consist of hundreds of thousands of strands and thousands of ports to cost-effectively and exponentially scale both bandwidth and operations. Enterprise wiring market tends to have a longer lifecycle (15-20 years), whereas, data centre physical infrastructure ROI is typically 3-10 years.

What we learned:

- What is driving internal and external traffic growth
- New competition landscape thanks to the web scale disruptors – plus how much are web giants spending on big data centres
- Artificial Intelligence and Machine Learning traffic growth implications
- Disruption and vertical integration in Telco and Media

Note: 1.7: Source: https://cloud.google.com/tpu/docs/tpus. What is Tensorflow: https://opensource.com/article/17/11/intro-tensorflow

Section 2:

TRAFFIC GROWTH FORECASTS – THE INFOGRAPHICS

The traffic explosion inside and outside the data centre

The digital economy is built on data centres which can scale from a networked micro facility with a few servers to a multi-building campus. Today, these can consist of 100,000s of servers with many kilometers of cable infrastructure connecting them to storage devices, switches and networking hubs, through which they connect to the global internet backbone networks.

In every category, forecasts for volumes of data traffic are rising including:

- Internal data centre traffic
- Traffic entering and leaving the data centre
- Mobile traffic
- Consumer generated traffic
- Enterprise IT data
- Machine generated traffic
- IoT

Data centre architectures are changing with more multi fibre connectivity for high-density connectivity. 10Gbps interfaces are no longer sufficient to support bandwidth demands of the data centre. 25Gbps, 50Gbps, 100Gbps ports are expected to dominate the data centre going forward. These changes require superb quality in the physical fibre layer.

Number of hyperscale data centres

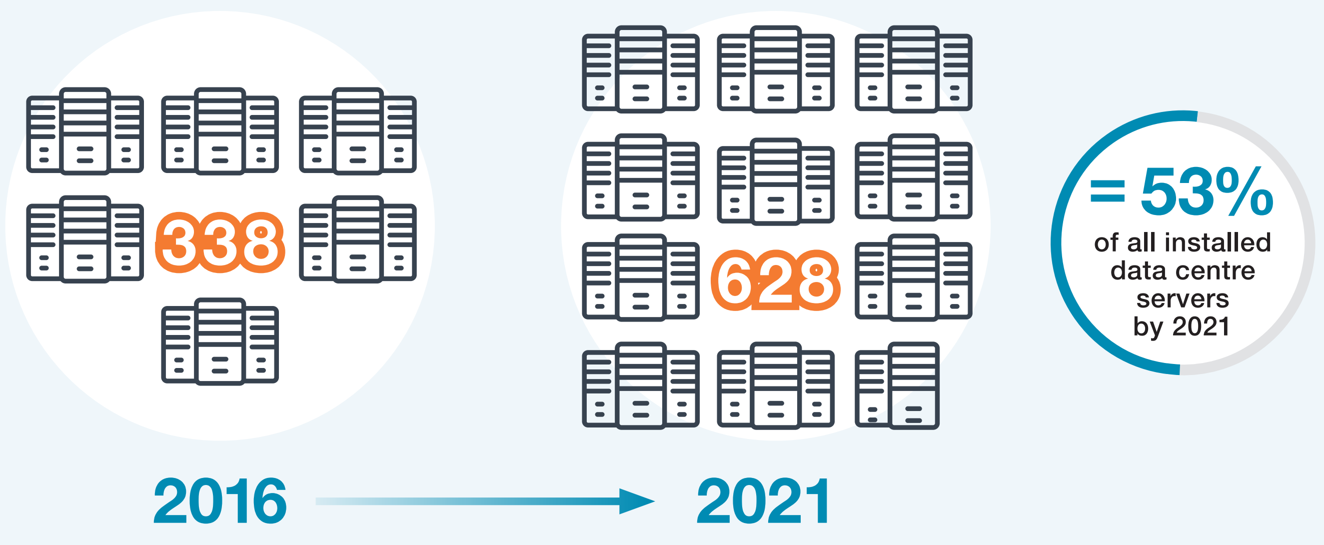

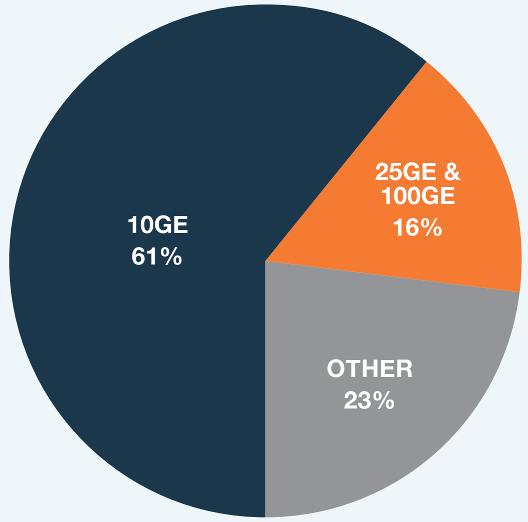

% of ports shipped in 3Q17

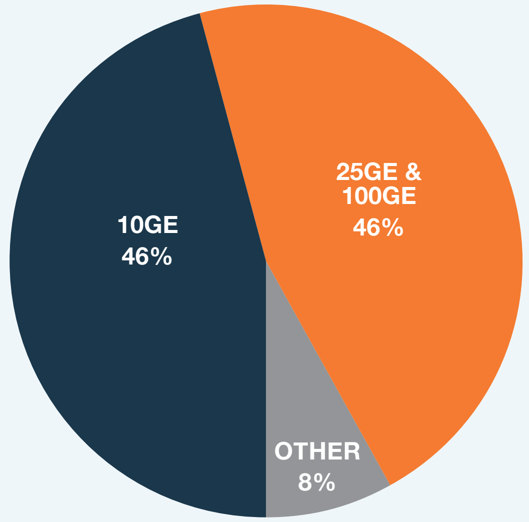

Future forecast of % of ports shipped

Forecasts on number of physical data centres and data growth

Hyperscale data centres will grow from 338 in number at the end of 2016 to 628 by 2021. They will represent 53 percent of all installed data centre servers by 2021.

Traffic within hyperscale data centres will quadruple by 2021. Hyperscale data centres already account for 39 percent of total traffic within all data centres and will account for 55 percent by 2021.

Annual global data centre IP traffic will reach 20.6 Zettabytes (ZB) (1.7 ZB per month) by the end of 2021, up from 6.8 ZB per year (568 exabytes [EB] per month) in 2016.

Global data centre IP traffic will grow 3-fold over the next 5 years. Overall, data centre IP traffic will grow at a Compound Annual Growth Rate (CAGR) of 25 percent from 2016 to 2021.

By 2021, video streaming will account for 10 percent of traffic within data centers, compared to 9 percent in 2016, but 80% of traffic from data centres to devices.

Cisco forecasts a 400% increase in global data centre IP traffic and 2.6 times growth in global data centre workloads from 2015 to 2020. See **note 2.1**.

Disruption driving data centre Ethernet

IHS Markit *Data Centre Network Equipment* market tracker, showed that worldwide data centre Ethernet switch revenue grew 12% YoY in 3Q17, reaching $2.9B.

Worldwide data centre Ethernet switch port shipments grew 24% YoY in 3Q17, reaching 12.5M.

25 and 100 Gigabit Ethernet (GE) ports shipments saw 251% and 369% YoY growth, respectively.

Note: 2.1: Source: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html

25GE and 100GE port speeds combined made up 16% of ports shipped, and 10GE comprised 61% of ports shipped in 3Q17.

25/100GE ports shipments are forecast to rise to 46% combined, whilst 10GE will decrease to 46% by 2021, as customers migrate from 10GE to 25GE server connections.

IHS Markit expects trials for 200/400GE to have begun in 2018 with production shipments occurring in 2019 and revenue to reach approximately $1Bn by 2021. See **note 2.2**.

Mobile driven traffic

Moor Insights & Strategy believes changes will drive massive IT hardware spending to the tune of $326 billion by 2025. This includes 5G data processing, storage and networking needs in the data centre and edge compute, carrier network transformation projects, and 5G modems and IP.

"The full impact of 5G on IT hardware spending will be $326 billion by 2025. This investment splits between different segments including the data centre at 56% (including public cloud), network transformation at 22%, edge computing at 3% and modems and IP at 19%." See **note 2.3**.

Note: 2.2: Source: http://techblog.comsoc.org/2018/02/15/25 -100ge-switch-port-growth-surges-200-400ge-in-2019-optical-tranceiver-market-overview/

Note: 2.3: http://www.moorinsightsstrategy.com/wp-content/uploads/2018/02/The-Full-Impact-of-5G-on-IT-Industry-Hardware-Spending-By-Moor-Insights-And-Strategy.pdf

Machine learning chip units shipped

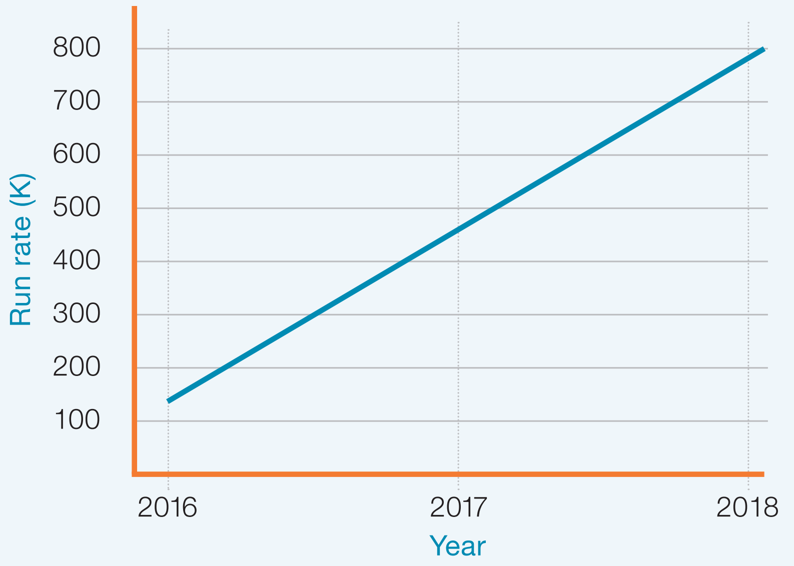

Estimated size of data centres

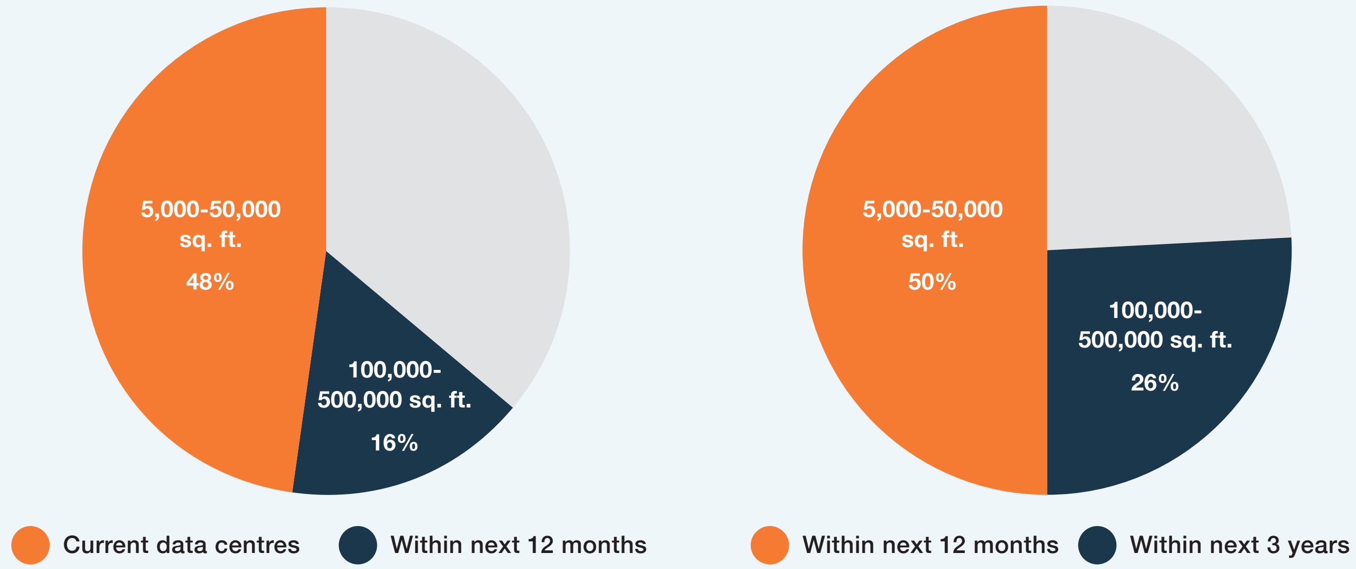

Mobile data growth forecast – Ericsson

Monthly mobile data traffic per smartphone continues to increase in all regions. Western Europe has the second highest usage, with traffic set to reach 4.1GB by the end of 2017 and 28GB by the end of 2023. Western Europe will be the region with the highest growth rate in monthly mobile data traffic per smartphone during the forecast period. This data traffic growth demonstrates the potential increase in the latest cabling solutions and the processing technology required to control and drive this demand. North America has the highest usage, and traffic is expected to reach 7.1 GigaBytes (GB) per month per smartphone by the end of the year and increase to 48GB by the end of 2023. See **note 2.4**.

Machine generated Traffic

Deloitte Global is predicting machine learning chips used in data centres will grow from a 100K to 200K run rate in 2016 to 800K this year. At least 25% of these will be -FPGA- and -ASICs-. Deloitte found the Total Available Market (TAM) for Machine Learning (ML) Accelerator technologies could potentially reach $26B by 2020. See **note 2.5**.

Types of data centres expected to be built – AFCOM report

- 48% reported that their current data centres are between 5,000 and 50,000 sq. ft.
- About 50% said their data centres will be between 5,000 and 50,000 sq. ft. over the next 12 months
- Another 16% stated that their data centres are between 100,000 and 500,000 sq. ft.
- Looking three years out, 26% answered that their data centres will be between 100,000 and 500,000 sq. ft. See **note 2.6**.

EU Data Economy

The EU data economy is expected to increase to €739 billion with an overall impact of 4% on the EU GDP under a High Growth Scenario characterised by a stronger driving role of digital innovation and higher overall ICT investments.

A study has found that 6 million people in Europe worked in data-related jobs in 2015 and 6.16 million in 2016. As far as medium-term developments are concerned, it is estimated that under a high-growth scenario, the number of data workers in Europe will increase up to 10.43 million, with a compound average growth rate of 14.1% by 2020. See **note 2.7**.

Note: 2.4: Source: https://www.ericsson.com/en/mobility-report/reports/november-2017/mobile-data-traffic-growth-outlook
Note: 2.5: Source: https://www2.deloitte.com/global/en/pages/technology-media-and-telecommunications/articles/tmt-predictions.html

Note: 2.6: Source: http://www.datacenterknowledge.com/afcom/state-data-center-industry-2018-where-we-are-and-what-expect
Note: 2.7: https://ec.europa.eu/digital-single-market/en/news/final-results-european-data-market-study-measuring-size-and-trends-eu-data-economy

Section 3:

CURRENT AND FUTURE ROAD MAPS FOR FC AND ETHERNET DEVELOPMENTS; MIGRATION PATHS TO 100GBE, 200GBE, 400GBE, 800GBE AND 1.6Tbps

Summary

- How 40/100GE is being deployed
- The latest Standards – 200GE and 400GE
- The path to 800GE and 1.6TE
- The IEEE Task Force on next generation networks

At the transport layer, there is unlikely to be a single answer to networking when it comes to future proofing an investment. But the direction of travel for speed and capacity is only going one way and the networking technology which companies trust to scale linearly is Ethernet on single-mode and multimode fibre optic cable.

With 40/100GE already being deployed and the standards for 200 and 400GE being approved and published by the IEEE, ethernet is the preferred networking platform. It is able to deliver data centre architecture and design to meet tomorrow's applications needs whilst providing low operating costs today.

Alongside booming demand within enterprise and hyperscale data centres, an unstoppable wave of new applications continue to emerge. High performance data centre and networking infrastructure demands for these applications must be met.

These factors raise key questions:

- What are the data centre network foundations that need to be built?
- How can cabling and connectivity infrastructure investments ensure performance certainty for today's applications and a clear migration path to future Ethernet and Fiber Channel speeds?

As noted above in section 1 (**AI and Machine Learning infrastructure needs**), new applications are driving the demand for speed. Given the equipment that is coming onto the market the bandwidth speed needed on the backbone to take advantage of new ML and AI solutions will become the 'de facto' baseline requirement to avoid the network becoming the bottleneck.

The ML transport solution today is the ability to transmit at rates of 40Gbps and 100Gbps across single-mode and multimode fibre. The move to 200Gbps and 400Gbps is already underway. Whether considering this from a cable, connector, switch or network architecture perspective, there are implications for all.

Networking topologies must change. Hierarchical three tier architectures will have far too much latency to accommodate AI and ML. For such high volumes of East-West traffic, the server needs to connect to a single non-blocking switch or leaf. In leaf spine architecture, every leaf is connected to every spine but the spines are not connected. This alleviates oversubscription and latency problems.

Leaf Spine Configuration Reduces Complexity of Network Architecture

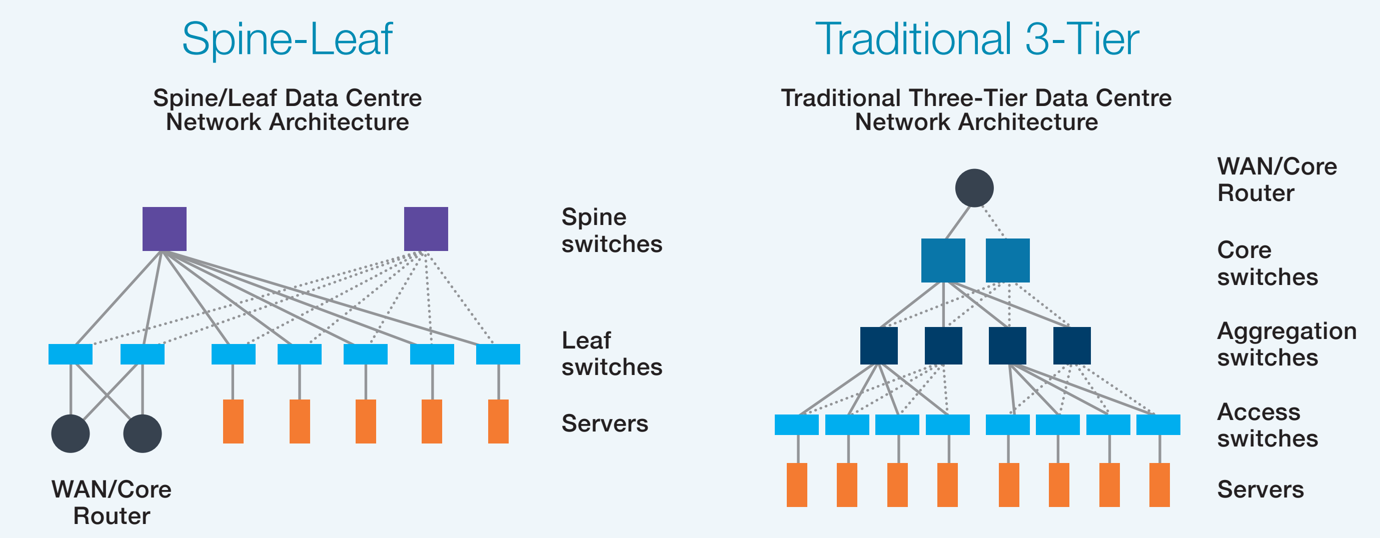

Signature Core

Today's Panduit multimode Signature Core fibre optic cabling technology provides the high throughput performance and low latency required for next generation networks. It supports 40GE bi-directional networking in the data centre. OM4 Signature Core and BiDi is far less expensive than deploying single-mode systems for 40G Ethernet.

OM5 Signature Core at 40G SWDM extends reach to 485m and maintains high performance values versus (minimum standards based) reach of 440m. At 100GE, SWDM OM5 Signature Core extends reach to 185m reach versus the standard 150m.

And it provides a clear migration path to future higher compatibility 400GE Bi-directional networks. See **note 3.1** and **note 3.2**.

Reach Table 1.5dB Connectivity

	Fibre Type				
	OM3	OM4	OM4 SC	OM5	OM5 SC
Application	**Maximum Channel Length (m)**				
10GBASE-SR	300	400	550	400	465
40GBASE-SR4	100	125	165	125	145
100GBASE-SR10	100	125	165	125	145
100GBASE-SR4	70	100	125	100	115
4GFC	380	400	500	400	445
8GFC	150	190	250	190	215
16GFC	100	125	200	125	160
32 Fibre Channel	70	100	125	100	115
128G Fibre Channel	60	85	95	85	90
Cisco 40G BiDi	100	135	200	150	175
Cisco 100G BiDi	70	100	150	130	150
40G SWDM4	240	350	420	440	485
100G SWDM4	75	100	150	150	185

Note: 3.1: http://www.panduit.com/ccurl/428/199/signature-core-d-fbfl04,0.pdf
Note: 3.2: https://fibrechannel.org/

Table 1

Data Rate Gbps	Nomenclature	Lane Rate Gbps	Number of fibre pairs	Number of wavelengths	Year Standardised
10 40	10GBASE-SR 40GBASE-SR4	10	1 4	1	2002 2015
25 100	25GBASE-SR 100GBASE-SR4	25	1 4	1	2016 2015
50 100 200	50GBASE-SR 100GBASE-SR2 200GBASE-SR4	50	1 2 4	1	2018
400	400GBASE-SR8		8		~2021
	400GBASE-SR4.2		4	2	

Table 2

Data Rate Gbps	Transceiver Form Factor	IEEE 802.3 PMD	Minimum Reach (Km)
10	SFP+	10GBASE-LR	10
25	SFP28	25GBASE-LR/ER	10/30,40
50	SFP56	50GBASE-FR	2
50	SFP56	50GBASE-LR	10
100	QSFP56	100GBASE-DR	0.5
100	QSFP28	100GBASE-LR4/ER4	10/40
200	QSFP56	200GBASE-DR4	0.5
200	QSFP56	200GBASE-FR4/LR4	2/10
400	QSFP-DD	400GBASE-DR4	0.5
400	QSFP-DD	400GBASE-FR8/LR8	2/10

200GE and 400GE Standards

The migration path from 40/100GE to 200GE and 400GE is specified in the IEEE 802.3bs Standard, (originally established in 2014) and had its 200GE and 400GE single-mode Standards approved in December 2017.

The single-mode Standards were published in March 2018 and are rapidly moving to commercial development and deployment.

The equivalent multimode Standards are defined in 802.3cd (50/100/200GE) and 802.3cm (400GE), to be published in 2018 and 2020 respectively.

Currently, 400GBASE-SR16: A multimode Standard specifying 32 fibres (16 transmit and 16 receive) with each fibre transmitting at 25 Gbps over a reach of at least 100m will be replaced with new reduced fibre pair Standards 802.3cm. Table 1 above lists the multimode fibre IEEE Standards and Table 2 (opposite) lists the single-mode Standards.

400GBASE-DR4, for reach of at least 500m over singlemode fibre. It uses four parallel fibres in each direction transmitting 100Gbps on each. It uses pulse amplitude modulation (PAM4) to double a 50GBd stream to 100Gbps.

400GBASE-FR8, uses wavelength division multiplexing (WDM) for reaches of at least 2km over a single-mode fibre in each direction.

400GBASE-LR8, reaches to at least 10 km over single-mode fibre using PAM4 to create 8 50Gbps data rates.

200GBASE-DR4 uses four pairs of single mode fibre to transmit up to 500 meters with each consisting of a 25GBd stream doubled using PAM4.

For 2km 200GBASE-FR4 uses WDM and for 10km distances 200GBASE-LR4 uses PAM4 to transmit four 50Gbps up to 10km over one single-mode pair.

The standards were endorsed by the Ethernet Alliance.

“The new IEEE 802.3 Standards represent a transformational moment in the move to next generation networks. The delivery of 200GE and 400GE is arriving just in time to meet growing needs for reliable, high-speed connectivity from a diverse array of applications and markets,” said John D’Ambrosia, chairman of the Ethernet Alliance and chairman of the IEEE P802.3bs Task Force. “The exceptional effort resulting in the completion of the 802.3bs standard is only the start of the industry’s investment in the networks of tomorrow. We’ve laid a firm foundation for 200G and 400G with our early interoperability demonstrations, but it’s time to kick things into high gear. The real work of testing and verifying multivendor interoperability begins now, and the Ethernet Alliance is ready. We look forward to building on past successes and helping accelerate 200G and 400G Ethernet’s rollout and adoption.”

See **note 3.3**.

Note: 3.3: https://ethernetalliance.org/wp-content/uploads/2016/03/EthernetRoadmap-2018-Side1-1600x1200.jpg

The Future of Ethernet

The Ethernet Alliance 2018 Future of Ethernet Map illustrates the current and future use cases for Ethernet for the years 2018-2020 and 2020 and beyond across GE, 10GE, 100GE, 400GE to TE+ (1.6Tbps).

Automotive and Industrial Automation are seen as major Ethernet Platforms. Automation applications, smart cars and sensors will develop from 100Mbps to 1GE range and are seen as a major new Ethernet platform, generating massive data flows over the coming years.

Enterprise data centres running LAN switches (100Mbps – 10GE) today and server stacks with 1GE – 10GE today will expand to 10-50GE quickly by the year 2020 and beyond.

The initial leaps to exponential speed are expected to take place in hyperscale data centres run by cloud providers which today run racks of homogenous servers at 50GE and data centre switches at 25-200GE. These will jump to running network speeds of 50-100GE and 50-400GE in less than a two-year timeframe.

Speeds beyond 400GE will see mobile and fixed telecom communication service providers require 800GE to 1.6TE. Optical transport equipment and wireless backhaul has continuously pushed Ethernet to higher rates and distances to meet wireless connectivity demands. Data loads required for high definition mobile video will continue to push Ethernet to multi terabit rates. As stated previously, one hour of 4k raw uncompressed video can be as big as 250GB.

What we learned:

- *How cabling and connectivity infrastructure investments can ensure performance certainty for today's applications and a clear migration path to future Ethernet and Fiber Channel speeds*
- *Why Ethernet over single mode and multimode fiber is the technology of choice*
- *The importance of today's 200GE and 400GE standards*
- *The future of Ethernet as it evolves to 800GE and 1.6TE*
- *Fibre Channel roadmap: https://fibrechannel.org/*

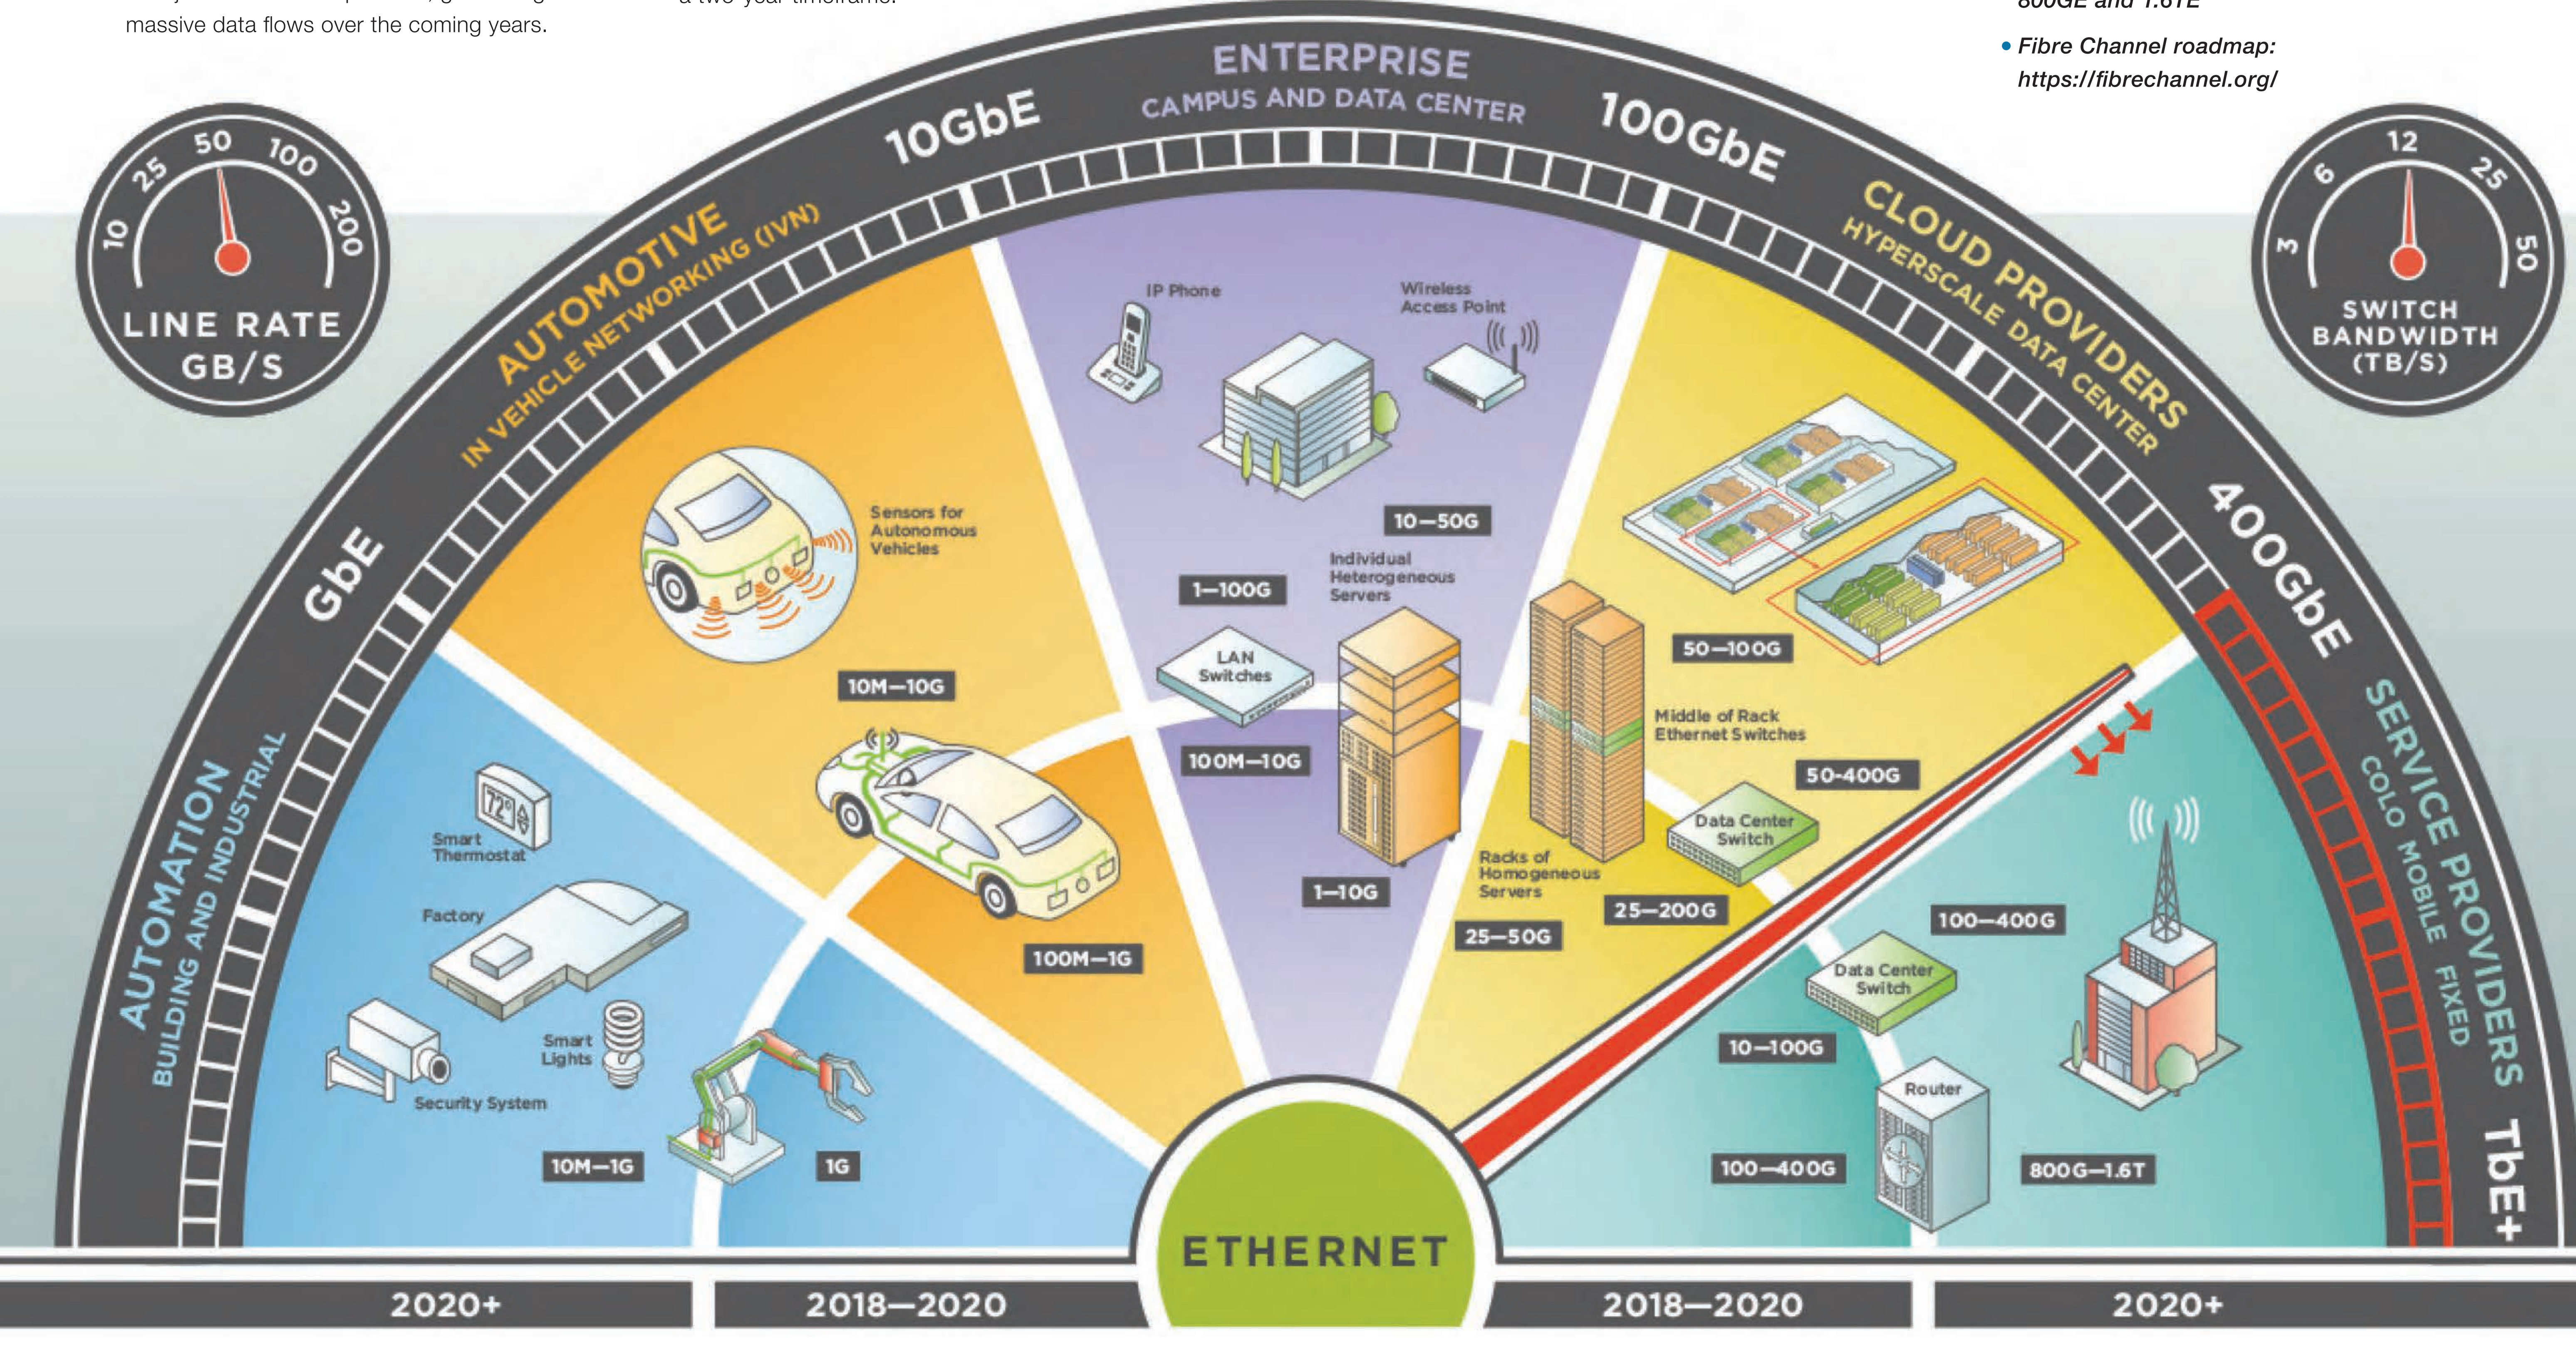

ILLUSTRATION CREDIT: ETHERNET ALLIANCE IEEE ROADMAP https://ethernetalliance.org/the-2018-ethernet-roadmap/

Section 4:

FUTURE PROOFING THE NETWORK WITH PANDUIT

Summary

- The path to 800GE and 1.6TE
- Working Group on next generation networks
- Next Generation Standards

Networking requirements for hyperscale data centres are pushing the industry to deliver faster and faster speeds. Demand for higher data rates from Microsoft, Google, and the other internet giants is driving the industry. At the same time all the communication service providers are moving from switching to packet switching and converting their data centres for this environment. The internet backbone is getting bigger and bigger at faster and faster data rates. Simply imagine how much data a company such as Netflix pushes across the internet backbone from its AWS cloud to consumers on a daily basis.

Networks continue to evolve and many data centre operators are moving rapidly to speeds of 40Gbps and 100Gbps and are already looking to speeds of 200Gbps and 400Gbps.

A key question for those with responsibility for data centre networking infrastructure revolves around how to future proof a 10-25 year capital investment with a restricted view and sometimes no view at all of the application, workload, or IT stack road maps beyond three years.

Every customer wants to know that their data centre investment will last and not contribute towards bandwidth bottlenecks, throughput or latency challenges. Future proofing requires vision. This is the responsibility of organisations such as Panduit and the work of its laboratories.

Among the questions that data centre operators need answers to are:

- What data rates will be required and what bandwidth will be needed in 2, 5, 10 and >15 years?
- Should data centre operators be aiming for 400GE today?
- How much head room will that provide?
- What are the foundational technologies which can be put in place today to provide readiness for 800GE and 1.6TE?
- What will be the best standards which will protect today's investments and create smooth migration paths?

From a cable infrastructure perspective Panduit believes that there is a way to future proof the data centre.

If looking to deploy cable infrastructure at distances of 100m (which could be classified as a medium sized facility) then opting for OM4 Signature Core or OM5 will future proof the data centre for up to the next 25 years. OM4 Signature Core and OM5 enable the reduction of fibre numbers allowing 40Gbps and 100Gbps transmission, and future 400Gbps with lower numbers of fibres using SWDM (Shortwave Wavelength Division Multiplexing) at a wavelength from 840nm to 950nm. However, SWDM transceivers are likely to cost more than their parallel fibre counterparts. A SWDM transceiver can use duplex multi-mode fibre connectivity for channel reaches of more than 100m. Parallel fibre solutions on the other hand, such as 100GBASE-SR4, will support the evolution to higher data rates over 4 and 8 pairs of multimode fibre to achieve data rates of at least 800Gbps. For short reach applications such as switch to server interconnects, multimode fibre will continue to offer lowest cost. Parallel fibre solutions also support port breakout required for high port density switching fabrics, unlike SWDM.

40G/100G SWDM Link Distance

PMD	OM3	OM4	OM5	OM4 SC	OM5 SC
40GBASE-SWDM4	240m	350m	440m	440m	485m
100GBASE-SWDM4	75m	100m	150m	150m	185m

Assumes 1.5 dB Connectivity Loss (IL) Budget

SC = Signature Core

Meeting demand for ever faster speeds

Though there are continual technological challenges to overcome, the technologies that are being developed to meet the demand for ever greater speed will push data centre networking into the terabit era.

In order to explore the future of fibre optics, an understanding of how today's speeds were reached is useful.

Today, networking speeds are averaging 10Gbps, with data centre operators already moving to 40GE and 100GE deployments.

In technology development terms, the move to 40Gbps saw the emergence of parallel optics because the VCSELS (Vertical Cavity Surface Emitting Laser) used in fibre optics transmits at no faster than 10Gbps. To achieve high data rate transmission, aggregated lanes are used. To reach 40Gbps four parallel fibres aggregate 10Gbps per fibre. To reach 100Gbps, 10 lanes of 10Gbps was initially developed, but replaced with four parallel lanes of 25Gbps which is already standardised.

The addition of lanes had an unforeseen benefit. It spurred-on switch manufacturers to increase the port density on their switches. This came about because each lane in Ethernet can be independent. Instead of building switches with 96x 10GB ports, this was condensed to 32x 40GB high density ports. This is known as breakout and is important because everything that has been worked on to develop standards to get to the PMDs of 40Gbps, 100Gbps, 200Gbps, and 400Gbps is breakout compatible.

This year 50GE, 100GE and 200GE multimode fibre standards will be ratified. Over the next three years – 100Gbps breakout at 4x 25G lanes and 200Gbps breakout at 4x 50G lanes will be widely deployed.

Road Map for Ethernet Data Rates Multimode Beyond Single-Wavelength Transmission

Data Rate Gbps	Lane Rate Gbps	No. of Fibres	No. of Wavelengths	Year Ratified
10 40	10	2 8	1	2002 2015
25 100	25	2 8	1	2016 2015
50 100 200	50	2 4 8	1	2018
400 800	50	4 8	4	2021

Currently the IEEE 802.3 task force has been established (802.3cm), in which Panduit is a key contributor, to define and propose standards. One objective is to define 400Gbps across eight separate 50Gbps lanes to be used for connecting servers and switches. There is also a specification for 400Gbps over four lanes at 100Gbps per lane, each lane carrying two wavelengths – which Panduit supports, but will require a breakout solution not to be specified in IEEE.

Moving forward to 800Gbps and beyond will also likely be breakout compatible. Panduit believes that it makes more sense to reduce 800Gbps switches down to 32 ports. This can be achieved by specifying 100Gbps serial transmission per lane. Feasibility was demonstrated at the 2018 Optical Fiber Conference (OFC) by Panduit researchers. The race is always to smaller devices with higher densities. Breakout at 800Gbps will be far more efficient in terms of heat dissipation, space, and power densities. It will also bring down the cost of manufacture. The move to 800GE will require agreement and ratification of technology standards.

The path to 800GE

An IEEE Task Force is currently specifying the use of Multimode Fibre (MMF) for rates beyond 100GE.

The "Next-Generation" 400Gbps Physical media (PHYs) over fewer MMF Pairs than Existing Ethernet Projects is underway. This new standard (802.3cm) will include two 400Gbps solutions, 400GBASE-SR8 having 8 pairs of multi-mode fibres each carrying 50Gbps, and 400BASE-SR4.2 having 4 pairs of pairs of fibres each carrying two bi-directional (BiDi) wavelengths at 50Gbps.

The formation of this IEEE task force is seen as indicative of the future of MMF use for high speeds in the data centre. Source: https://www.lightwaveonline.com/articles/2018/02/ieee-study-group-explores-future-multimode-fibre-roles.html

The next step up is already in discussion by the Ethernet Alliance and an IEEE Ad hoc group which are mapping the journey to 800GE and 1.6TE networks.

Today, networking is at 400GBASE-SR16 on 16 fibre pairs (16x 25Gbps), but rapidly moving to 400GBASE-SR8. To double this to 800G there is the question of whether MMF will be required. Due to its relatively short reach it is not clear if MMF will be used for data rates beyond 800Gbps transport for switch to switch connectivity. Panduit is proposing multiple lanes 100Gbps serial data transmission. The company believes that utilising 100Gbps serial and bi-directional SWDM2 a data rate of 1.6Tbps can be achieved.

Panduit believes there is no significant advantage of using 400GBASE-SR4.2 since it cannot be used for port breakout, nor will it offer any cost benefit. Although 400GBASE SR4.2 may be the appropriate choice for legacy (brownfield) networks requiring higher data rates.

Advantages of parallel fibres (parallel optics) Include:

- Facilitates adoption in both green field and brown field installations.
- 4 fibre pair structured cabling is used for 10/40G, 25/100G, 50/200G, and future 100/400G, 50/400G Ethernet applications.
- Works better with legacy fibre since the longer wavelength is not utilised.
- Less complex transceiver: fewer components, cheaper optical filters, less crosstalk, less need for equalisation.
- Enables breakout (400G to 50G) thereby satisfying current need for high density MoR to eight 50Gbps servers.
- Leverages well established ecosystem for optical taps, adapters, connectors, and cabling.
- Preserves current Ethernet cabling use cases
- Enables a simple upgrade path to higher speed Ethernet
- 4 fibre pair PMDs are universally accepted in industry
- A 2 or 4 wavelength solution will be more challenging to develop than parallel optics
- 2 wavelength solutions are better than 4. 4 wavelengths require tighter wavelength grid: higher control of 4 VCSEL EPI growths, tighter filter control, smaller guard bands
- No Standards specifying EMB at wavelengths other than 850 nm for OM3 and OM4
- Channel reach over OM3 & OM4 will be shorter for 4 wavelengths compared to 2

Panduit's position on next steps towards 800Gbps.

Advantages are:

- Multimode fibre will continue to offer the lowest cost connectivity for high-speed short reach data communications.
- Multimode lasers are inherently cheaper than single-mode which are required for high-speed channel reaches greater than 150 m.
- SWDM provides high-speed multimode solutions in optical cabling systems where utilising duplex fibres are the best solution.
- MPO based structured cabling systems will provide long life for MMF in the data centre.
- The transition to 800Gbps will be straight forward (just commission additional fibre pairs), compared to increasing the number of wavelengths.

Exactly which standards will be adopted and ratified for 800GbE and beyond will be the subject of debate over the next two years. Panduit does not recommend the use of multiple wavelengths since it is not cost effective, can potentially exceed eye safety limits, and does not support port breakout at 400GbE and beyond.

Next generation standards, ASICs and architecture

Elsewhere the pace of development continues in every area. At a microprocessor level there is research ongoing into 100GB per lane ASICs. Already web giants are deploying 50GB servers and the IEEE is now developing a Standard for 400GE that can be split into 8x 50GB lanes to support these roll outs. Commercial shipments of 100GB servers and switches may be some way off, but are being developed.

In Architecture terms, Panduit believes the most logical topology for high performance networking in data centres is the leaf spine (or spine leaf) architecture. As more data moves within the data centre than anywhere else, most traffic in a data centre is East-West – the goal remains to minimise the delay and latency. In a spine leaf network data only hops through one switch. A server goes to the leaf switch and every leaf switch is connected to a server and to a spine. But spines are not connected to one another, so this eliminates the added latency caused by multiple hops. For incoming traffic, a border leaf transports the data to a server in a single hop.

Panduit is at the forefront of developments to provide the networking technology interconnectivity by influencing the best standards possible. Its hundreds of patents help simplify network architectures, reduce cost of installation and operation and break the barriers to higher data rates and faster networks in data centres, commercial buildings, hospitals and every type of networked environment.

Panduit labs are a world renowned Research and Development facility dedicating millions of dollars each year to improving network infrastructure performance, resilience, efficiency and security and helping future proof data centre investments for every type of business while meeting the demand for higher and higher speed. The industry standing of the Panduit Laboratories is in no doubt as hardware and software partners and competitors utilise these facilities to develop and test their infrastructure products and systems.

What we learned:

- ***How Data centre owners can future proof the network using Panduit technologies***
- ***Why parallel optics are still the most effective***
- ***The Multimode fibre standards to be ratified***
- ***Use of multiple lanes to reach 800GE***
- ***How a data rate of 1.6Tbps can be achieved***
- ***Panduit's leadership position in ultra high speed network infrastructure***
- ***Panduit's position on future standards***

Conclusion

Fibre optic technology is truly the backbone of current and future data centres and enterprises, supporting the platforms that are expanding the data economy. It is essential that future networks are built on industry standards, which provide the platform for cohesive developments.

As the data centre and high-end enterprise IT environments push for greater and faster data throughput, the choice of future-proof cabling infrastructure is an essential one. Fibre optic cable and connectivity systems are that critical path-way that will take today's 10G – 40G systems through to 100G and 400G, and beyond.

Panduit's partnerships with leading server, switch and storage manufactures, together with hyperscalers ensures that product innovation is equal to the requirements of quality, reliability and cost, while providing additional benefits to help drive the industry forward.

Panduit's R&D expertise continues to push the limits of network performance, while providing sustainable alternative environments with opportunities of transparent migration paths.

Panduit is focused on lowering today's networking operation costs. By aligning IT and business strategy, Panduit's approach combines clear technology roadmaps with business deliverables.

Fiber science leadership: By investing 10% into R&D Panduit truly conducts the "science" behind our fiber technologies.

Panduit enclosures, connectors and fibre cabling are compelling solutions excelling in performance beyond typical fibre solutions.

Panduit has an industry leading product portfolio that pushes the limits beyond the standard to yield more future focused results, performance and compatibility for customers.

PANDUIT™

Panduit EMEA, West World, W Gate, London W5 1UD, +44 (0)20 8601 7200
For more information: Visit us at www.panduit.com