



VIAVI Solutions

The Data Center of Tomorrow

INSIGHT REPORT

viasolutions.com/hyperscale

Contents

CHAPTER 1		
Data: The World's Most Valuable Resource		3
<hr/>		
CHAPTER 2		
Servers: Less Is More		8
<hr/>		
CHAPTER 3		
Processing Power Gets Edge-y		12
<hr/>		
CHAPTER 4		
Achieving Scale and Speed: The New Breed of Hyperscale Data Centers		17
<hr/>		
CHAPTER 5		
Delivering More Speed – With Less Power		23
<hr/>		
CHAPTER 6		
Conclusion: The Challenges and Opportunities Ahead for Data Centers		27
<hr/>		

CHAPTER 1

Data: The World's Most Valuable Resource

More than the drops in the ocean

Water. Air. Oil. For decades, these types of resources have been considered the world's most valuable. But as technology becomes more pervasive in every aspect of our work and personal lives, another resource has emerged that is equally important – if not more so: data.

The world's dependence upon data is primarily driven by Internet Content Providers (ICPs) like Amazon, Google, Facebook and Tencent – companies that generate revenue through online sales, financial transaction fees, paid advertising, cloud services and a host of other business lines.

The phenomenal growth in data can also be attributed to device proliferation. [There were 35 billion IoT connections in 2020, expected to hit 83 billion by 2024.](#) The explosion of data has been [spurred on by 5G](#), which has seen adoption surge during the pandemic, [eclipsing the comparable rate of 4G adoption by a factor of four](#) and setting the stage for even bigger growth in data that needs to be managed. Little wonder, then, that this data boom has also fueled the growth in data centers, which are set to grow steadily by [a compound annual growth rate \(CAGR\) of 4.5 percent.](#)

And stars in the sky

To understand just how quickly data is proliferating, consider the ways in which we measure it.

In a short timeframe, data units have moved from being measured by kilobytes, megabytes, gigabytes and terabytes to zettabytes, or 1 billion terabytes. The [digital universe generated around 64 zettabytes of data in 2020](#) – meaning that over 60 times more bytes of data were processed than there are stars in the observable universe.

The rise of data presents multiple challenges for the organizations and infrastructure required to support it. Internet of Things (IoT) sensors that continually transmit signals are just one aspect of the data challenge. 5G, the cloud and digital transformation also demand connectedness and drive demands for data. Whether we're binge-watching the latest television series, checking social media or using online retailers, the engagement economy is also driving data growth.

5G technologies are pushing more intelligence to the network edge.

While the numbers associated with the rise of data are staggering, the hyperscale data center ecosystem is rising to the challenge of meeting this insatiable demand for memory, bandwidth, computing power, storage and speed.

Cloudification is blurring the traditional lines between networks and applications, and 5G technologies are pushing more intelligence to the network edge.

This exponential data growth is forcing data centers to be closer to their customers, resulting in more edge deployments, while also forcing data center operators to increase speed, security and efficiency at the same time as they minimize latency. Hyperscale data centers, and the cutting-edge technologies accompanying them, are turning these unprecedented challenges into opportunities.

35 billion
connected things

There were 35 billion IoT connections in 2020

83 billion
connected

It's expected there will be 83 billion connected things by 2024

64
zettabytes

The digital universe generated 60 times more bytes of data in 2020 than stars in the observable universe

73
zettabytes

By 2025, IoT devices will generate about 73 zettabytes of data

Infrastructure and Content: Two different mindsets

At one time, traditional telecom operators were the heaviest users of data centers. Operators needed to plan for high availability, so they conducted numerous tests, monitored and measured systems multiple times and over long periods.

For traditional operators, investment in network hardware is a long-term game: typical equipment lifespan can be 10 or more years, so selecting the right equipment and maintaining it is something operators take very seriously. The gold standard is five-nines availability (99.999%). To achieve that, hundreds of engineers work tirelessly 24/7 testing, measuring and maintaining networks.

The mindset of ICPs is quite different. Content is their main focus, so data center networks are just seen as the pipes to deliver that content to users. The pace of operations for ICPs is distinctive too, where much shorter timeframes are the norm. For instance, ICPs don't expect infrastructure to last more than five years. The amount of time spent in the testing phase is also shorter just to get the new data center up and running as fast as possible. ICPs are fiercely competitive, and the pressure is on to innovate faster and deliver content more quickly.

Operations teams are often smaller, relying on application programming interfaces (APIs) and software-based network automation to maximize workloads and network operation. Likewise, ICPs don't always follow standards bodies as their timelines are different and don't match the ICP business model. As a result, they frequently white box technology from multiple vendors or develop it themselves – even before industry standards have been agreed.

While ICPs are driving the data boom, it is data centers that support that growth. As such, it's imperative that next-generation network infrastructure be able to handle these increased traffic requirements. Recognizing these different approaches to data is the first step.

Faster, higher, stronger... but with less power

Managing ever-increasing data speeds and volumes is not the only challenge for data centers. Power consumption is another major concern.

Energy forecasts for data center electricity consumption are pretty shocking. By some estimates, projected usage will run between [8-21% of total electricity demand by 2030](#). Other estimates peg it at around 3% of worldwide electricity supply. However, the latest analysis comparing different methodologies, and taking into account improvements in data center power consumption, show that a more realistic figure is [around 1% of global energy use](#). While these updated figures are good news, the industry cannot afford to rest on its laurels. How much further can we push server efficiencies? How will data growth impact energy consumption?

The need to maximize performance while minimizing power consumption is attracting the brightest minds in the industry to innovate data center design. For instance, data centers are increasingly located in colder climates to ease the cooling requirements. Renewable energy sources are being harnessed to reduce the burden on the energy grid.

Infrastructure as an enabler

The seamless interconnect of data center facilities requires lightning-fast speeds, and yet ICPs have grown at such a rate that there is little time for the necessary rigorous testing.

Add the rising cost of cabling infrastructure and the array of interoperability protocols to the mix, and the scale of the challenge starts becoming clear.

Old data center infrastructure is reaching a critical point. Having driven the data boom to a large extent, ICPs now have an opportunity to create a blueprint for the next generation of data centers. Preparing for future data needs is as much about innovating data creation as it is about ensuring the supporting infrastructure is flexible, scalable and responsive.

As technologies like the IoT, artificial intelligence (AI) and machine learning (ML) mature and proliferate, the data requirements for the engagement economy will extend far beyond social media, entertainment, shopping and news. Public services, energy infrastructure and other core sectors will increasingly rely on data center performance. Test and measurement has a critical role to play in assuring and delivering next-generation networking.

CHAPTER 2

Servers: Less Is More

Serverless computing is redefining the way that software is deployed, wherein developers don't need to know anything about the hardware or operating system (OS) that code runs on.

Also known as platform-as-a-service (PaaS) or function-as-a-service (FaaS), serverless architecture offers an additional layer of abstraction to the software stack, meaning developers no longer having to worry about underlying server or capacity management.

The code is executed and fully managed by the cloud provider. Globally, [the serverless architecture market is expected to grow at a CAGR of nearly 23%](#) to reach \$21.1 billion by 2025.

For many enterprises and ICPs, serverless architecture delivers something closer to the original concept of utility computing. Enterprises that move to the cloud gain on-demand responsiveness, scalability and faster set-up – the very traits that applications like IoT and big data initiatives require to be viable. But there are negatives as well. Enterprise teams increasingly spend time maintaining the software stack, brokering capacity and managing the spiraling complexity of the cloud model.

The secret formula

Serverless computing is transforming the way complex software is developed, managed and deployed and offers agility, scalability, security and simplified billing.

Little wonder that the names driving the content boom are the same names driving the serverless computing revolution:

So what makes serverless so popular?

- **Billing:** AWS Lambda charges in 100 millisecond increments. Serverless computing is taking the on-demand computing model to a new level.
- **Reduced complexity:** It is not just the billing model that is appealing. Because teams no longer need to concern themselves with managing and provisioning servers, resources are freed up. Those resources can be redeployed in more strategic, Agile initiatives. It also significantly reduces operational overhead.
- **Scalability:** The ability to scale rapidly is another big advantage of serverless. Software and DevOps teams no longer have to prep computing resources to accommodate big spikes in usage. Likewise, applications automatically scale down when things are less busy.
- **Security:** Everyone knows the damage a distributed denial of service (DDoS) attack can cause. Serverless architectures aren't immune, but they are far less threatened by them. Serverless also goes a long way to protecting against OS-level attacks.

Software and DevOps teams no longer have to prep computing resources to accommodate big spikes in usage.

- **Edge:** Serverless was originally conceived for cloud environments, but is increasingly also being deployed at the network edge. Why? Serverless edge computing helps address the processing requirements of growing numbers of IoT and 5G applications.

Regardless of these benefits, serverless is not a “one size fits all” solution. It has gained most traction in event-driven use cases. Yet, serverless doesn't guarantee the same service-level agreements (SLAs) as more mainstream cloud computing offerings, and there are concerns over vendor lock-in.

Moreover, it's easy to be overwhelmed by the offerings in the cloud marketplace. There are hundreds of different vendors with many more cloud platform services, and yet, the market is increasingly dominated by the “big three” (AWS, Microsoft Azure and Google Cloud) who, together, account for [58% of the global marketplace](#).


What serverless means for servers

Regardless of whether it's called PaaS, FaaS or serverless computing, the concept is here to stay, and cloud models are evolving to include it. But it's also important to determine how serverless architecture is impacting underlying data center infrastructure.

The performance and reliability of existing infrastructures, such as data center interconnects, will come into sharp focus. The results may come as a surprise to ICPs, whose traditional test and measurement practices usually stop short of scrutinizing data center infrastructure. In a world where capacity demands can fluctuate wildly, delivering on these data demands efficiently and effectively becomes a competitive differentiator. ICPs need to gain a clearer understanding of how far they can push their data center interconnects (DCIs) to achieve maximum speed and capacity. Using the latest test and measurement techniques, ICPs will be able to better respond to issues or even head them off before they become a problem.

It will be more important than ever that data center monitoring and orchestration tools keep pace with these changes. Test and measurement tools need to be agile, automated and virtual to support data center infrastructure and achieve the volume and

speed of traffic flow that will be required to power on-demand computing. Serverless computing heralds an additional layer of abstraction for developers, enabling them to focus on their work. At the backend of this process, serverless also heralds a new age of always-on, endless scalability at the push of a button. In infrastructure terms, the demands of this new era means the margin for error – downtime – is even slimmer. Data center teams will need a helping hand with such a task ahead of them. That's where test and measurement plays an important role in meeting capacity demands and assuring performance.



CHAPTER 3

Processing Power Gets Edge-y

As the IoT continues to expand exponentially, the edge becomes increasingly important.

By 2025, IoT devices will generate around 73 zettabytes (ZB) of data. The IoT, enabled and amplified thanks to the commercial rollout of 5G, will soon run through almost every aspect of daily life, from national infrastructure like the electric grid, to the cars we drive (or perhaps don't drive) in the future.

Nearly half of organizations plan to increase IoT investments as a result of the pandemic. This demand for data has pushed ICPs to expand the storage and transport capacities of their data centers to meet the storage and processing needs for billions of applications supported in our "always on" environment. None of this would be possible without the sky-high processing power of the cloud.



The need for speed

In particular, hyperscale multi-tenant data centers are transformative. The data speeds and scale support the seemingly limitless possibilities of the IoT. To keep up with the demand, one area of expansion for ICPs has been in DCIs, which link up data centers around the world.

However, for IoT to reach its full potential, yet more interconnect is required, in terms of speeding the transport and processing of data that is generated. For instance, autonomous cars and eHealth are just two instances that demand a zero-tolerance approach to latency, as every nanosecond holds the potential for a life-or-death situation.

For the most part, hyperscale cloud computing takes care of the processing. Likewise, the evolution of mobile networks from 4G to 5G takes care of data speeds, especially as 5G is able to transmit data 10 times faster than the previous 4G LTE technologies.

Latency, however, is the major limiting factor, which is why ICPs have looked to mobile edge computing (MEC) for a new approach. Mobile edge computing addresses this challenge by storing and conducting the analysis at the edge – outside of the cloud – dramatically cutting data transmission times.

Everything from robotics to smart cities to drones and eHealth will depend on MEC. The principle is simple enough — huge processing requirements, including data storage and analytics, at the edge of the network are coupled with the immediacy that comes from phenomenal speeds. These are the basic tenets of the IoT.

The importance of the edge becomes clear when considering that the market was valued at \$4.7 billion in 2020 and is expected to grow by 38.4% CAGR through 2028.

Indeed, [Gartner predicts](#) that three-quarters of enterprise-generated data will be processed at the edge by 2025, up from just 10% in 2018.

Keeping it local

Strictly speaking, edge computing takes place on the device. Another term that is often used in the same breath as edge computing is fog computing.

The difference with fog is the location of the distributed processing, which takes place within the local area network (LAN), as opposed to on the device itself. However, both approaches collect and process data “locally” and relieve pressure from the cloud. The benefits include real-time analysis, lower operating costs, efficient device battery life and the all-important “immediacy” that is expected from IoT.

Both edge computing and fog computing collect and process data “locally”.

FOG COMPUTING

Location of distributed processing takes place within the LAN

Testing times for the edge

Expectations are sky – or even cloud – high when it comes to IoT. It will be crunch time for ICPs to deliver on those expectations.

As more services go online, those expectations will continue to rise. Developers talk in terms of customer experience (CX). Network operators obsess over subscriber quality of experience (QoE). Retailers worry about consumer experience. For ICPs, certainty is paramount. Test and measurement (T&M) has a critical role to play in keeping the IoT running 24/7/365, and the stakes could not be higher.

It may come as a surprise that many ICPs do not routinely test their data center interconnects. Faster troubleshooting is a huge advantage of routine testing; but possibly an even greater benefit is the power of prevention. In addition to ongoing maintenance, ensuring optimal performance of new, upgraded systems prior to turn-up is always critical.

Running DCIs at full capacity makes sense, especially as buying additional DCI connections can be expensive and resource-intensive.

Here are three takeaways to strengthen your T&M strategy and maximize uptime:

- 1 Take a structured, scheduled and consistent approach to DCI testing. Robust testing is the only way to really know if all the available capacity can be utilized.
- 2 Conduct stress tests to assess DCI connections to find potential issues before a fault actually happens.
- 3 T&M practices can pinpoint any issues in achieving full capacity. Furthermore, testing helps to resolve these issues much faster than it would otherwise be the case.

In the era of IoT, edge computing is relieving the pressure on data centers. For ICPs, keeping on top of the data center is no longer enough. ICPs must have the assurance that storage and transport capacity is performing at optimal levels – locally too. Next-generation T&M is part of the essential toolset for keeping the cloud and the edge working optimally – with a zero-tolerance approach to latency and downtime.

CHAPTER 4

Achieving Scale and Speed:

The New Breed of Hyperscale Data Centers

Data centers power the digital economy. But how do ICPs and other businesses deliver the always-up, always-on, lightning-fast processing required?

The rise of both the cloud computing model and the steadily advancing adoption of 5G are accelerating digital transformation in almost every area of our lives.

To keep pace with this rapid change, it is no surprise that data centers themselves are undergoing a major reboot — from design, to scale, to the way they are powered, organized and run. The rise of the cloud computing model has redefined the data center market as well. This has seen growth and consolidation in equal measure. The key players are building bigger data centers, in locations that suit their customers' needs, which has enabled a raft of efficiencies.

The telecom industry, meanwhile, has traditionally favored a long-game approach to data center deployment, characterized by meticulous planning and comprehensive pre-deployment testing to ensure five nines (99.999%) availability and a data center lifespan of 10+ years.

The rollout of 5G is placing unprecedented strain on data center infrastructure worldwide, with the introduction of new services, IoT verticals and the intelligent edge that powers them. Increased end-to-end network complexity is challenging the old deployment schedules, while simultaneously raising the bar for performance, efficiency and reliability.

To address these challenges, a new breed of data center continues to grow in scale and importance. Enter...the hyperscalers.

The rise of hyperscalers

Hyperscale is increasingly used to define not just the scale and size of these new data centers, but also their architecture.

Still, their size and scale is a useful place to start. [Hyperscale data centers have a minimum of 5,000 servers and at least 10,000 square feet in size.](#) In general, however, hyperscale data centers are much larger, frequently numbering tens of thousands of servers. But the main thing about hyperscalers is less to do with the proportional requirements and much more to do with their ability to scale rapidly in response to demand. Beyond the footprint and server figures, equally important is what is going on inside, where they're architected [for a homogeneous scale-out greenfield application portfolio](#) using increasingly disaggregated, high-density and power-optimized infrastructures.

There are now [more than 650 hyperscale data centers in the world](#) today, with many more being planned and regular capacity expansion ongoing at existing facilities. Amazon, Microsoft and Google account for more than half of all major data center facilities globally.

They are monuments of advanced architecture, networking and automation. They might also be seen as temples built to satisfy our appetite for data, and they are dominated by the world's greatest creators of data: ICPs.



HYPERSCALE DATA CENTER

Minimum of
5,000 servers and
at least 10,000
square feet in size

5G - The challenges of speed and scale

The intricacies of distributed, disaggregated and cloudified 5G networks are pushing data center design to the limit.

Balancing cost is one challenge: data center costs are rising just as their scale and complexity continue to escalate. There are also the performance pitfalls surrounding 5G: virtualized RAN, massive MIMO and antenna beamforming complicate the performance testing requirements within hyperscalers, introducing new spectrum analysis, demodulation and SLA challenges.

To overcome these challenges requires end-to-end network slicing to be orchestrated seamlessly for each unique vertical – in other words, hyperscalers need to have a seamless, holistic approach to network management. In order to deliver fully automated and programmable network slicing and edge computing, hyperscalers require a next-generation approach to network test and assurance, a step change from outdated, traditional data centers with their siloed network management. Critical 5G IoT use cases leave no margin for error with respect to network performance and reliability.

Staying up to speed

As hyperscale data center footprints continue to expand, traditional management practices are making way for a greater reliance on automation.

Where once it made sense to employ engineers to walk the floors, monitoring infrastructure such as power and cooling, data centers seem to be much emptier spaces these days. Their sheer scale, coupled with technological advances, means that previously manual tasks have slowly been replaced by sensor-embedded hardware that's remotely managed.

Perhaps the biggest changes in data centers are those that pave the way for speeds of 400G and beyond. Next-generation networking is not just about bigger pipes. The demand for data brings with it a corresponding requirement to manage it all efficiently and securely. The demand is not just for more capacity, but for new data services, meaning there is more riding on data centers than ever before. Hyperscale data centers need to be intelligent, flexible and automated, as well as scalable.

Keeping pace with the advances in speeds and demand for data is no walk in the park. ICPs, the foremost owners and users of hyperscale data centers, face the challenge of managing their legacy installed bases while simultaneously transitioning to new capabilities with no drop in availability.

Test and measurement plays an important role in alleviating the pressure and riskiness associated with integrating modern technologies like 100G or 400G into the data center. When the infrastructure is as complex as it is awe-inspiring, we need a means of ensuring reliability across the whole network ecosystem regardless of capacity demand or underlying technology.

This is where a comprehensive approach to deployment and maintenance comes into its own. For instance, automated testing tools can now inspect and certify fiber endfaces for faster network buildouts and test functionality for multi-fiber push-on (MPO). Likewise, advanced active optical cable (AOC) and direct attach cable (DAC) test practices can be deployed and are essential for ensuring optimum network performance and to address the challenges brought on by the growth of multi-fiber connectivity. Moreover, in such complex, heterogeneous environments, automated test scripts have an important role to enable replicable, easy testing.

Hyperscale data centers need to be intelligent, flexible and automated, as well as scalable.

The next frontier in data center management

Automated management, test and monitoring solutions herald a new range of exciting possibilities in hyperscale data centers.

Sensors are already being deployed to “listen” for signs that a fan might develop a fault or for the sound of dripping water, signaling a leak. Sensor-embedded hardware will play an increasingly important role in monitoring infrastructure, with responsibilities for temperature, noise, vibration and others. Where sensors are found a layer of intelligence is needed to make sense of the new data streams and to act on the information being generated. Advanced analytics and machine learning are being deployed alongside these sensor-based tools and will play an increasingly pivotal role well into the future.

Sensors are already being deployed to “listen” for faults.



CHAPTER 5

Delivering More Speed – With Less Power

Not only are data centers getting bigger in terms of servers and footprint, there are more of them as well, with data traveling faster and faster. Speeds at data center interconnects (DCIs) are already moving from 100G to 400G, with speeds of up to 800G planned in the near future.

This data avalanche is seemingly unstoppable. However, there is one thing that could curtail growth. As speeds and data keep rising, so too does power consumption. While data centers have successfully managed to optimize power consumption as data has grown, to what extent power efficiencies can keep up with data demand is uncertain. What is certain is that rapidly growing power consumption is unsustainable. ICPs have been phenomenally successful in serving up content that is extremely appealing, immediately available and (practically) free. They are now paying for it, however, in terms of power consumption.

Rising power consumption rates are unsustainable



Going green

ICPs have had to become inventive in their approach to power usage. For instance, some ICPs are turning to the colder climates to lend them a helping hand, reducing the need for power-hungry air conditioning and using ambient air to cool servers.

One such example is Facebook's state-of-the-art data center in northern Sweden, located in the town of Luleå. The town is less than 70 miles south of the Arctic Circle, and Facebook utilizes the cold outside air temperatures to help cool its thousands of servers. Furthermore, hydro-electric plants operate on nearby rivers, providing a reliable and renewable power source.

Facebook's state-of-the-art data center in northern Sweden utilizes the cold outside air temperatures to help cool its thousands of servers.

The efficiency paradox

There is some evidence that ICPs are having some success with curbing energy usage.

Despite the fact that data center growth has risen, growth in power consumption is flattening out. However, as the appetite for data keeps rising, it will be far from easy for ICPs to constrain their power consumption, as per the Jevons paradox. This effect, named after the 19th century English economist William Stanley Jevons, occurs when a technological process becomes more efficient, but rather than leading to less consumption (energy in this case), it actually leads to greater usage.

As global citizens of the Internet, we show no signs of tiring of data. Short of taxing, or otherwise restricting the use of the Internet, the data trend shows no signs of slowing. Increasingly, ICPs are turning to renewable energy sources to address the challenge of energy usage.

As global citizens of the Internet,
we show no signs of tiring of data.

There are other weapons ICPs can utilize to combat spiraling energy bills. For instance, ensuring data center interconnects are running efficiently and delivering anticipated speeds can be achieved as part of a rigorous monitoring and testing program. To prepare for increasing DCI speeds, ICP engineers need to run testing on 400G interfaces to pinpoint potential problems and troubleshoot them early.

While ICPs have reacted nimbly to the corresponding growth in power consumption with innovative data center designs, the challenge is far from solved. ICPs must remain open to emerging solutions that ameliorate their carbon footprint. However, there is much that can be achieved using the tools already at their disposal.



CHAPTER 6

Conclusion:

The Challenges and Opportunities Ahead for Data Centers


Content is king

As global demand for data continues to expand exponentially, it's not difficult to view data as the world's most valuable resource. Companies like Amazon, Apple, Facebook, Microsoft and Netflix are the primary drivers behind this change, and ICPs are driving the new data economy.

The sheer quantities of data being generated currently, while quantifiable, are hard to conceive. Increasingly large measures – teraflops, zettabytes and even yottabytes - are needed to denote data quantities, and we have only just begun. As 5G, IoT and cloud computing continue to pick up pace, the data-driven economy will itself go into overdrive.

But data volumes are only half the story. The promise of self-driving cars, drones and robotics, as well as more everyday examples like streaming videos, voice assistants and chat bots, all rely on ever-increasing data speeds to function optimally. It is estimated that a self-driving car could generate over 25 gigabytes per hour – that's nearly 30 times more than a HD video stream.

It is easy to appreciate why a self-driving car or a robot performing surgery requires zero-latency network performance. However, it doesn't take a life-or-death situation for latency to become a pressure point. It is not unusual for consumers to react badly to waiting even when what's at stake is less serious: a patchy VoIP call, an online search that takes too long, or a buffering cute cat video. And when consumers are made to wait or otherwise suffer poor QoE, they react with their wallets. ICPs know this.



Infrastructure and business models

The advent of spiraling data volumes and lightning-fast speeds also adds a new kind of pressure.

To deliver the expected data speeds and volumes, data center facilities need to ensure seamless DCIs. The old order in data centers is being challenged on all sides: cabling infrastructure costs are rising, interoperability protocols abound and, crucially, even the way ICPs operate is changing.

Traditional network infrastructure investment cycles – and subsequent testing and maintenance – do not match the super-fast cycles in which ICPs operate. Content is what counts, and demand for it has moved at such a pace that there has inevitably been less scrutiny of the underlying infrastructure. To continue to support the engagement economy, data centers are adopting a fresh approach to keep pace with the ICP business model.

The sheer quantities of data being generated currently are hard to conceive.

New times, new challenges

Relatively new forms of computing, such as serverless and edge, are taking some of the pressure off the data center, but they create other challenges instead.

For instance, the role of DCIs in keeping the data ecosystem moving will become increasingly important. Delivering on data demands with an optimally efficient architecture will become a competitive differentiator. Getting the most out of existing data center infrastructure is a challenge ICPs now face.

As data center footprints and capabilities grow, the issue of power consumption looms ever larger. ICPs are urgently having to rethink the relationship between data processing and power consumption. Data center design increasingly considers how ICPs can use natural resources and renewable energy to keep costs down.

Data centers are not just getting an external makeover; their design is different on the inside, as well. As their footprint expands and the available automation and monitoring technologies mature, data centers are becoming less reliant on manpower to ensure smooth operation. Instead, these tasks are being replaced by sensor-embedded hardware, as well as networking and automation software – all of which are remotely managed. As data speeds and volumes grow and data centers increasingly support a flourishing industry for data services, the margin for error – or downtime – shrinks.

Testing times for the data center

Data center monitoring and orchestration tools are now increasingly agile, automated and virtual to support the new needs of data center infrastructure.

ICPs, whose traditional focus has been on content, are waking up to the benefits of comprehensive deployment, management and optimization techniques to push their existing infrastructure further, deliver on data speeds, maintain uptime and head off issues before problems arise.

Rigorous test and measurement can also help ICPs manage change and speed network buildouts. For example, [test process automation](#) and automated inspection practices alleviate the headaches involved in integrating new technologies into the data center. As data centers upgrade their equipment to enable speeds of 400G and 800G, and beyond, ICPs need a way of ensuring reliability. And as data processing becomes distributed and abstracted across hybrid environments, ICPs also need to ensure performance across the whole ecosystem.

In short, ICPs are learning to harness the power of the data center in new ways. Where the role of test and measurement in the data center had once been somewhat overlooked by ICPs, it is becoming clear that rigorous testing across virtual, physical and cloud-based infrastructure plays an integral part in the mission to achieve 24/7/365 uptime.

Solutions for hyperscale ecosystems

VIAVI Solutions partners with [hyperscale](#) data center operators, internet content providers (ICPs), cloud service providers and those deploying robust [data center interconnects](#) (DCI) to reduce testing time, optimize optical networks, reduce latency and ensure 100% reliability that supports SLAs.

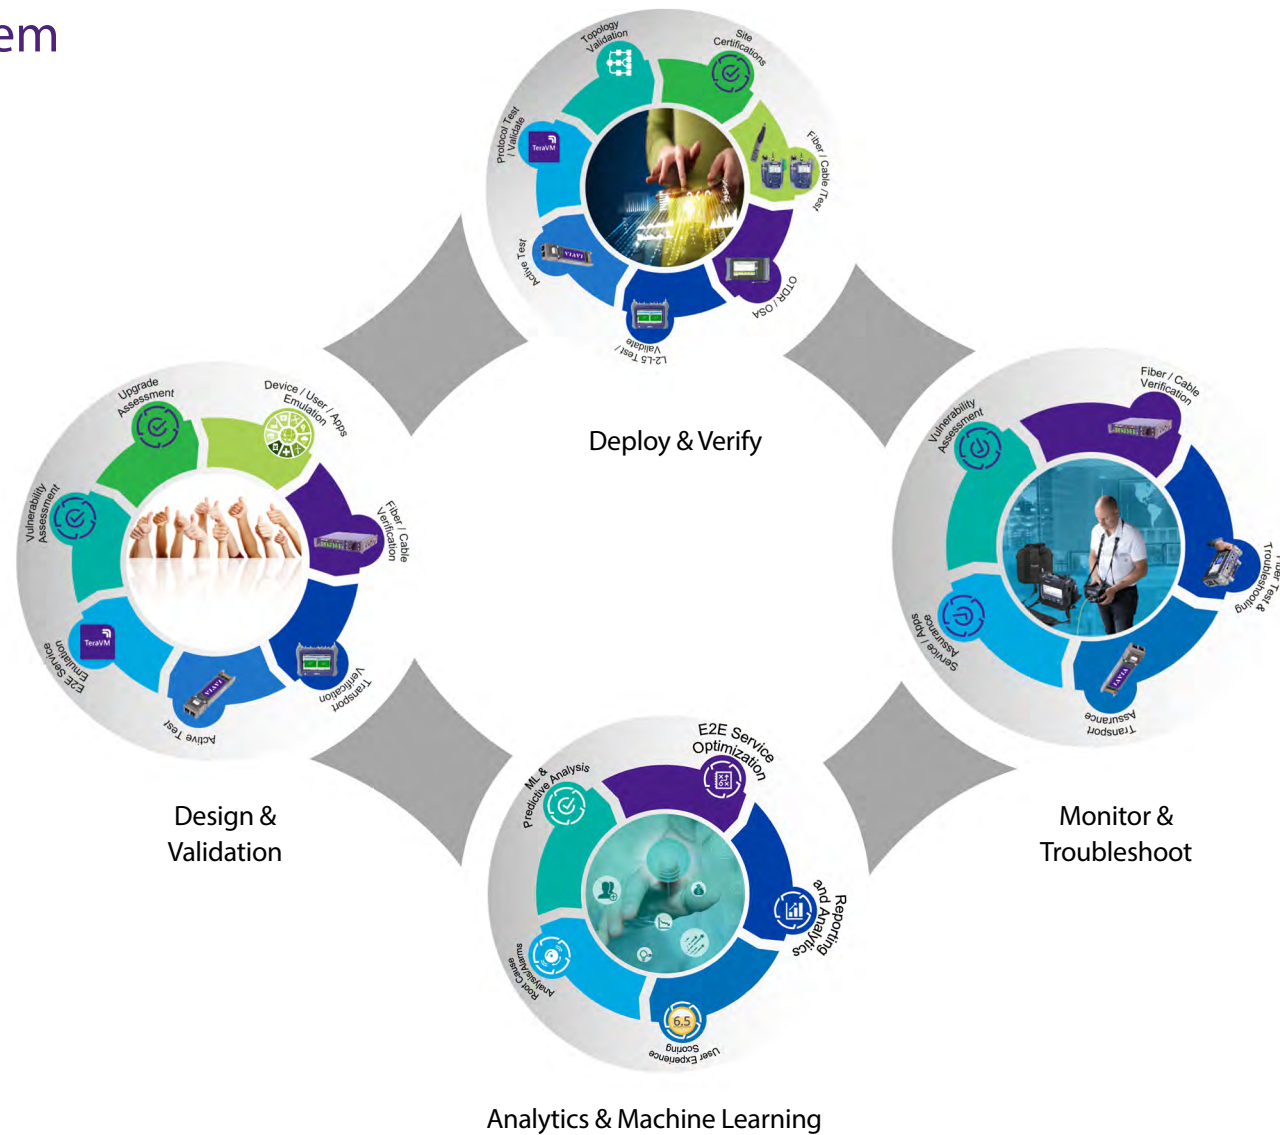
We guarantee performance of optical hardware from labs to turn-up to monitoring, including equipment that can inspect MPO connectors in 12 seconds as well as equipment that can test two 100G ports simultaneously. Because VIAVI Solutions is involved in all stages of hyperscale optical testing, we understand how you're building high speed networks up to 400G and beyond, and we have the equipment to test it all — from traditional test and measurement tools to [test process automation](#) and [cloud-based management](#) solutions.

VIAVI Solutions is an active participant in over thirty standards bodies and open source initiatives including the Telecom Infra Project (TIP). But when standards don't move quickly enough, we anticipate and develop equipment to test evolving standards. We believe in open APIs, so hyperscale companies can write their own automation code. VIAVI Solutions has been testing communications equipment for nearly 100 years.



Testing the ecosystem

VIAVI provides the multi-dimensional visibility, intelligence and insight you need to efficiently manage physical and virtual environments, in order to profitably deliver optimum service levels, transition to new technologies and launch innovative services.





Contact Us +1 844 GO VIAVI (+1 844 468 4284)

To reach the VIAVI office nearest you, visit viavisolutions.com/contacts

© 2021 VIAVI Solutions Inc.

viavisolutions.com/hyperscale

© 2021 VIAVI Solutions, Inc.
Product specifications and descriptions in this
document are subject to change without notice.
data-center-of-tomorrow-bk-opt-nse-ae
30193177 902 1221