DCD

> Telco | Supplement

Sponsored by

Schneider Electric™

INSIDE

Sponsored by

Schneider Electric

Contents

Hailing on all frequencies

Telecoms is an equal partner with data centers in the digital infrastructure which is vital to our lives. Either without the other would much less useful.

Telecoms has been seen as a slow-moving sector compared with the digital world and the cloud (see comments by OVHcloud CEO Michel Paulin, on p20, for instance).

That is now changing. The vast and rapid shifts which are continuously remaking the data center sector, are shaking up the telecoms sector, transforming the ways in which operators interact with customers, transmit data, and choose their battles.

This supplement takes a look at the biggest of those changes.

A competitive Edge
Data center resources are being directed to the edge of the network, to support emerging applications like the IoT (p26).

This move depends on having access to fast networks to transmit the data those applications need. So Edge is as much an issue for telecoms as it is for facilities.

There's an unprecedented crossover in the moves to deliver the hardware that these applications need. Standards for telecoms facilities are emerging from the data center world

5G brings a revolution
The fifth generation of mobile services (5G) is more than just a new and faster kind of network. It's data-driven, and delivered through small cells.

5G services will be fundamentally data driven, and will drive a new dependence on resources embedded into the network

Using satellites
Fiber has been seen as the best way to distribute telecoms services. After the losses, limits and power demands of copper, glass fiber sparked a revolution when it began to displace metal wires in long distance communications.

But now satellites are seeing an unprecedented comeback. It turns out that low earth orbit (LEO) constellations can take a lot of data. And even though the route up to space and back down seems like a long way to go, there are fewer hops, and the light travels faster in the vacuum of space than it does in glass (p34).

So expect to be offered satellite-based offerings for applications like back up and recovery

Funding fiber
Despite that, fiber roll-outs continue apace. As well as the obvious international submarine cable connections, projects are aiming to find new ways to fund fiber for areas where cost is a barrier (p32).

Seeing sense
Finally, this rush of exciting telecoms activity shouldn't make operators feel invulnerable. They are still recovering from a rash of over-investment in data centers in the early years of this decade (p36).

As telcos shift their data center assets to the specialists who can handle them, we are seeing an object lesson on focus.

Telcos need to think on their feet, and move smartly.

A competitive Edge for Telcos

Shifting virtual workloads into the Edge gives both fixed and mobile network operators multiple advantages. *Martin Courtney* reports

Martin Courtney
Contributor

As a new approach to architecting telecommunications networks, there is no doubt that Edge computing has significant potential to change the way that carriers and service providers deliver a range of services to their business and consumer customers. But with the technology at such an early stage of its development, specific use cases are still under development - particularly as telcos work out the best way to exploit software-defined networking (SDN) and network functions virtualization (NFV) to drive down their own infrastructure costs and streamline provisioning, configuration and management processes.

There is still debate over what the network Edge actually is, and no strict definition to clear up the confusion. Most see it as smaller data center hosting/processing facilities located closer to the end user, but others feel it could incorporate local workloads running on customer premise equipment and other points of presence with local (LANs) rather than wide area networks (WANs). Instead of a location, the Edge defines only a workload hosted at some indistinct node within the provider or customer infrastructure.

According to network and telecommunications equipment manufacturer Cisco, the point of the Edge is threefold: to deliver lower latency to the end device to benefit application performance and improve the quality of the experience; implement Edge offloading for greater network efficiency; and perform computations that augment the capabilities of devices and reduce network transport costs. To that end, much of the ongoing innovation has so far focused on the enablement of fifth generation (5G) cellular networks.

Indeed most mobile operators agree that cost efficient 5G service delivery is simply unfeasible without the deployment of some form of Edge data hosting and processing capability to override the need to transmit and crunch large volumes of information via centralized data centers and the core.

But Cisco's imperatives apply equally to applications and services delivered over wired broadband connections as they do to 5G links, as much to Edge infrastructure workloads in cable broadband and gigabit-capable passive optical networks (GPON) access as to the 5G radio access network (RAN). As such, fixed line carriers and service providers too are looking at where Edge computing solutions can help them deliver wired broadband connectivity - and the range of IP-based voice and data services which that supports - to customers previously accessed via local loop telephone exchanges.

> *"Edge computing will process data to facilitate services as close to the user as possible"*

US telco AT&T, long at the vanguard of SDN/NFV adoption, is currently working to convert some of its estimated 4,700 telephone exchanges into mini data centers. The fixed and mobile network giant is close to achieving its goal of virtualizing 75 percent of its infrastructure by 2020, having already deployed SDN enabled broadband access (SEBA) to deliver superfast fiber broadband services to consumers and businesses in US cities such as Irving and Atlanta. SEBA is a set of open networking components that virtualize the software to run optical network terminals (ONTs) and optical network units (ONUs) on fiber networks, though it can be extended to other types of network including fixed wireless and Gfast that use copper cabling.

In the UK, BT is extending core network functions currently hosted within five to ten of its exchanges to around 100 metro locations. BT's own Network Cloud will evolve to reduce data and application latency, again initially for 5G applications and services. But, once the infrastructure is in place, it can be used for a variety of different functions, including broadband, IP telephony, and unified communications as a service (UCaaS) provision to customers. BT currently has around 1,200 local exchanges in the UK which serve as a first point of aggregation, more of which could be migrated to Edge facilities to meet the needs of different cloud hosted services in the future.

5G cell towers are another proposed location for Edge compute resources in the base of 5G cell towers which can also be used to accommodate fixed line operators' equipment. Defined by the Open Network Foundation (ONF), the Central Office Re-architected as a Datacenter (CORD) initiative combines NFV, SDN and commodity clouds to bring cost efficiency and cloud agility to the Telco Central Office (UK parlance the local telephone exchange), allowing them to dynamically configure new services for residential, enterprise and mobile customers in real time.

As voice as well as data traffic becomes IP enabled, routing and switching functions can be virtualized, making them easier to provision, configure and manage remotely. It is envisaged that the reference implementation of CORD will be built from commodity servers and white-box switches defined by the Open Compute Project (OCP) which are cheaper to buy than proprietary telecommunications hardware for example, alongside disaggregated access technologies (vOLT, vBBU, vDOCSIS), and open source software (OpenStack, ONOS, XOS).

Elsewhere, the European Telecommunications Standards Institute (ETSI) multi-access Edge computing (MEC) specification was designed to promote the

convergence of mobile base stations and IT and telecommunications networking, ostensibly to support anticipated new business cases around video analytics, location services, IoT, augmented reality, data caching and optimized local content distribution (what used to be known as a content delivery network - CDN).

Those use cases were defined specifically with 5G in mind, but as software overlays devolved from the underlying network, there is no reason why they cannot be applied equally to wired broadband connections too (multi-access is included in the acronym for a reason). Similarly, Edge routers, designed to process data collected from thousands of different devices and end users, already provide various interfaces to both wired and radio-based transmission technologies and communication standards - everything from 5G and WiFi to Bluetooth and Ethernet.

Edge can compete with cloud, after a fashion. Telcos like AT&T and BT need to be able to deliver fast, reliable hosted voice and data services. They are a crucial element of commercial cloud strategies, but the cloud is delivered from centralized facilities, a sector where the telcos have failed.

IT giants like Amazon Web Services, Microsoft, Google, IBM and others have won in the enterprise space by investing heavily in building their own hyperscale facilities.

After finally admitting defeat, AT&T sold off its core data center assets to Brookfield Infrastructure for $1.1bn earlier this year, following similar divestitures by other telcos (see p36).

Having a distributed compute infrastructure at their disposal gives telcos something the cloud service providers do not, and would find very difficult to obtain for themselves: dedicated Edge hosting and processing facilities closer to the customer which are better able to support a range of latency sensitive applications for business customers. Those could include everything from infrastructure- (IaaS), platform- (PaaS), network- (NaaS) and unified communication-as-a-service (UCaaS) to industrial IoT (IIoT) and high definition video capture (e.g. CCTV surveillance and consumer retail applications), the provision of which (telcos hope) could be supplemented by value added systems integration and managed services contracts.

And building out their Edge facilities puts telcos in a prime position to make themselves indispensable to bigger cloud providers when it comes to delivering more latency sensitive services and applications to their own customers - a potential market carve up that plays to both sides' strengths and reach.

5G, the Edge and the service revolution

Vlad-Gabriel Anghel
Contributor

A lot of 5G predictions sound like hype. *Vlad-Gabriel Anghel* explains the reality

In mainstream media during 2019, the term 5G has been increasingly seen and touted as the future for mobile communications and data processing. But infrastructure industry giants have been hard at work for quite some time getting ready to tackle the challenges that come with the vast amounts of possibilities that 5G will allow.

It is tempting to see 5G as an incremental step up from 4G/LTE, but 5G is exponentially better. It is capable of reaching speeds of up to 20Gbps and supporting up to a million devices per square kilometer (that's a lot of IoT devices) while providing an alleged 1ms latency. 5G is ultimately the true foundation for the Internet of Things.

Since the emergence of IoT devices, network limitations have placed numerous boundaries in terms of real-life use cases, while fields like HPC applications have seen limitations on data handling requirements.

A ridesharing app can only reach out to an AI prediction algorithm within a data center so many times per minute, and the same applies for other types of apps that rely on data processing in the cloud. This is true regardless of where the data is obtained from: Ultimately, network constraints will not allow for a fully seamless and instantaneous end user experience if the processing is centralized in the cloud.

The digital infrastructure industry has proposed a solution to this in the form of Edge computing - a way to make distributed systems more efficient by taking out parts from a centralized core and making them available closer to the data source or the Edge.

In simpler terms, this means storage, data services and computing power being redistributed accordingly based on their service or function. The ones that benefit from lower latency are moved closer to the Edge, while the rest remain at the core.

Edge computing can reduce latency by placing critical resources close to the end users and increases resiliency as it creates alternate data transmission routes. It does, however, fragment the system and that can pose a risk in terms of both physical and logical security, and because it relies on additional hardware it requires a significant upfront investment. In practice, the current capabilities of Edge computing are far from being able to support the innovative use cases envisioned in one form or another for several decades. It all comes down to network latency and availability.

5G can go a long way in removing these constraints, because it effectively increases the capability of the network edge. The demand for data storage and processing

power there will tremendously increase.

If 5G delivers on that promise of 1ms latency and one million devices per square kilometer, it will reshape a lot of industries, right down to their best practices and design standards.

For example, the majority of distributed system architects have been limited in their design choices by bandwidth and latency considerations. If 5G brings these barriers down then, instead of monitoring 30 sensors in real time, systems might manage 1,000, if this would bring a competitive or strategic advantage.

The same goes for mobile apps. Instead of querying a cloud endpoint every minute, why not every second? As these design choices evolve, it will have a massive impact on the digital infrastructure supporting these systems, and businesses will need to embrace HPC technologies and design strategies. Essentially, 5G will rely heavily on high-performance computing elsewhere.

Current mobile networks are capable to a certain extent of providing services for technologies like autonomous cars, drones and weather forecasting, but these applications will be truly unlocked by the use of 5G.

What gives 5G network their tremendous speeds and bandwidth is their technology. They operate on what is known as millimeter waves - radio signals with a frequency between 30GHz and 300GHz (4G operates between 1GHz and 5GHz). These have less range than shorter wavelengths, so the area previously covered by one 4G transmission tower must now be covered by a multitude of smaller, inexpensive 5G antennas fixed to buildings and streetlights.

Telco Edge data centers are the first to see this change and keeping up will be tricky. Just revisiting the field of autonomous cars - with their numbers on the rise, telemetry data will be gathered by multiple 5G antennas on a continuous basis. To analyze this in real time, and keep all the cars in lane and on the road, will require high-performance computing through storage and AI predictive services.

This raises considerations about deploying the proper HPC equipment in an efficient and sustainable manner. With 4G, devices connect on a one to one basis: the cellphone connects to the telecoms tower; the tower connects to another tower and so on. 5G will allow devices to connect to multiple antennas and this presents the possibility of the utopian scenario of 100 percent reliability. However, operators will need to embrace HPC technologies like distributed file systems, in-memory data grids etc. and implement proper design methodologies as traffic scales exponentially to millions of writes per second.

The possibilities and requirements of 5G will reshape how data centers are built and operated and current service providers will need an overhaul on their infrastructure in order to keep up. The backend services supporting 5G will need to be much more scalable than the previous 4G equipment, however these services will most likely be born out of already existing cloud native tools and technologies which are the center point of building scalable cloud services.

Current mobile communication frequency (long wavelength)

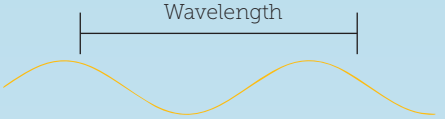

Extremely high frequency mmWave (short wavelength)

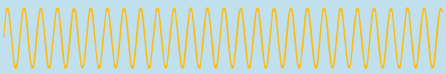

Furthermore, storage and computing power will shift towards the Edge, as close to the end users as possible. As mentioned before, these will need to be designed and deployed through an HPC methodology.

With data increasing vastly, the need for data analytics and management will also increase. This will happen again through HPC technologies like object stores, distributed databases and file systems. Operators will require these tools and technologies in order to streamline the deployment, management and scalability of larger data volumes. As the data moves from the secure core cloud data center to the network Edge, new security issues arise. The more fragmented a system is, the harder it becomes to provide proper security. This need is further underlined by the critical nature of the applications which are likely to use 5G services, such as connected traffic control systems, autonomous cars, drones and the like.

The arrival of 5G has propelled a diversification exercise throughout the data center industry, and the landscape is changing. Businesses like Vapor IO and EdgeConneX are attempting to create a new ecosystem of Edge modular data centers and have predicted tremendous growth for this sector due to 5G. Meanwhile, already established players within the data center sector could need to shift their focus towards deploying micro data centers, while making additional investments in already existing data centers and colocation facilities to keep up with the upcoming demand.

This should not be seen as blocker in the expansion of these businesses but rather a necessary step in fully releasing the potential of mobile technology and communications.

Ultimately, 5G is evolving to become what is known as a general-purpose technology (GPT) - a type of technology that has the ability to drive fundamental change across the entire global economy. Previous examples of GPTs have been the printing press, the automobile and the steam engine. As data center owners and operators come to grips with the needs and challenges of 5G and adjust their infrastructure and facilities accordingly, 5G end user numbers will soar, making it the latest and most impactful GPT to date.

With the number of IoT devices connected to the Internet expected to reach the order of billions in the near future, it is through 5G that all of these devices will be able to interconnect and exchange data, more quickly and more reliably than before.

It does, however, require data center owners to first find huge investments while managing stakeholders' expectations on short term ROI. Deployment of these methodologies will need precise and careful planning and implementation.

The industry as a whole is not there yet, still a lot needs to be done in order for 5G to become the network and service revolution it aspires to be, but it seems, at least for now, it is on the right path.

The future is around the corner!

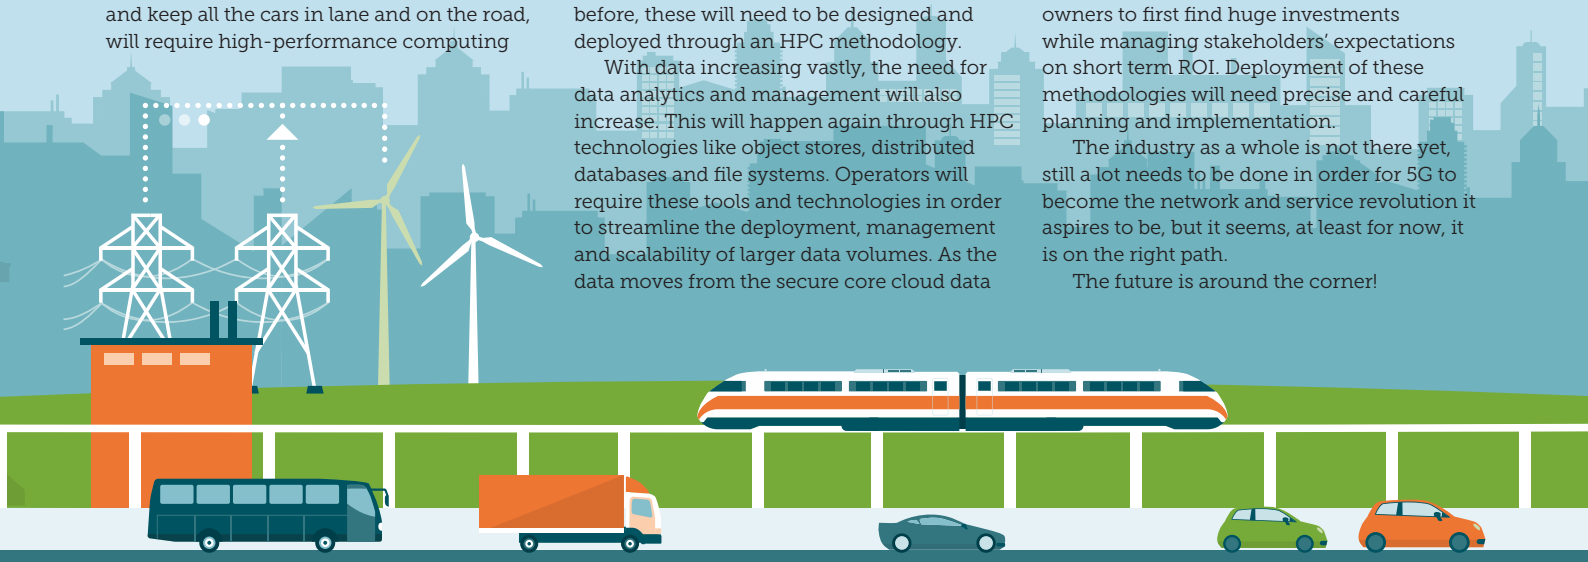

Getting up to Speed on 5G Strategy and Micro Data Centers

Greg Jones. Schneider Electric's VP of Strategy and Offer Management for the Cloud & Service Provider Segment talks to Steven Carlini, VP of Innovation and Data Center

Everywhere I go, it seems that people are talking about 5G. I wanted to find out more about the specific role of micro data centers in 5G so I sat down with Steven Carlini, our Vice President of Innovation and Data Centers. Steven is responsible for developing integrated solutions and communicating the value proposition for Schneider Electric's data center segment. I knew he'd have a wealth of timely insights on the topic and he didn't disappoint. Here is part of our conversation.

What is your definition of micro data centers? I would classify micro data centers as two or less IT racks, where a massive amount of computing power or storage can be managed. Today, we have micro data centers as small as 6U that can hang on a wall or even be put in a ceiling! I see micro data centers as a critical extension of cloud data center architectures to reduce latency and add redundancy in a hybrid cloud environment. These micro data centers are a key building block. However, because they are spread out all over the place, they do present challenges in the form of troubleshooting, maintenance, and repair. Also, energy usage becomes a critical operating expense at scale. For example, let's say you have 2,000 sites and 10 KW per site – that's 20MW of power! That equates to roughly $20 million in electric bills a year (operating at average efficiency). This is why the efficiency of micro data centers is top of mind for a lot of companies.

Is Schneider Electric investing in micro data centers and/or infrastructure changes? For Schneider Electric's Secure Power Division, edge data centers including micro data centers are a top priority. Many projections show that the market for local edge micro data centers will approach or exceed the market for hyperscale mega data centers. Historically speaking, the data center market has been cyclical between centralized and distributed. I see it coming again, but it may move to a more balanced architecture as core and edge become an integrated architecture.

Why are micro data centers important? Applications and operations are moving closer to the user or the data on the edge and micro data centers are handling many business functions. Think of the hotels that solely rely on local data centers for coding digital room keys and managing reservations. In the future, many hotels are considering facial recognition for completely automated experiences. As the world gets more automated, it will rely more and more on micro data centers.

In a 5G architecture, micro data centers are essential. And, local clusters are necessary to meet 5G performance targets. But we have the speed of light limitation of 300 million meters per second. So, with less than 1 ms of latency spec for 5G, the maximum distance is less than 200 miles round trip – and that's a theoretical best case. We are dealing with many carriers and they are laying out their clusters in circles with a much smaller radius than 100 miles, especially in densely populated areas.

Can you provide some background on 5G micro data centers? Let's start with the two main enablers of 5G: one is the new radio access network (RAN) and the other is the data center architecture. 5G uses microwaves and millimeter waves at very high spectrum so the signal does not go far. An easy way to think of it is the three-house rule: to operate 5G you need an antenna for every three houses. That is how close together they need to be. And, let's go back to the speed of light at 300M meters per second. The only way to achieve the required latency of less than 1 ms is to build local clusters that will include micro data centers. In a recent 5G test in Chicago, a 4K movie was downloaded in 20 minutes using 4G and 19.5 minutes using 5G. Why did that happen? It was because the only 5G portion on the connection was from the small cell hanging on the light pole to the phone – about 100 feet. The movie was in a data center many miles away. In the near future, micro data centers will serve the function of mobile edge computing (MEC). They will have traditional telco functions, like call routing, and also IT functions, like content delivery. For example, that movie could have been stored in the small cell and downloaded almost instantly.

In a 5G architecture, micro data centers are essential. And, local clusters are necessary to meet 5G performance targets

What will 5G micro data centers do for energy savings? This is a lively discussion topic as Schneider and our industry have been focused on energy efficiency for larger data centers over the years. Our goal is to make the micro data centers as efficient as the hyperscale data centers with no increased OpEx. Schneider, like many companies, has carbon neutral goals and edge data centers are a big part of those goals. We know that globally, millions of units will be needed to support 5G. As we

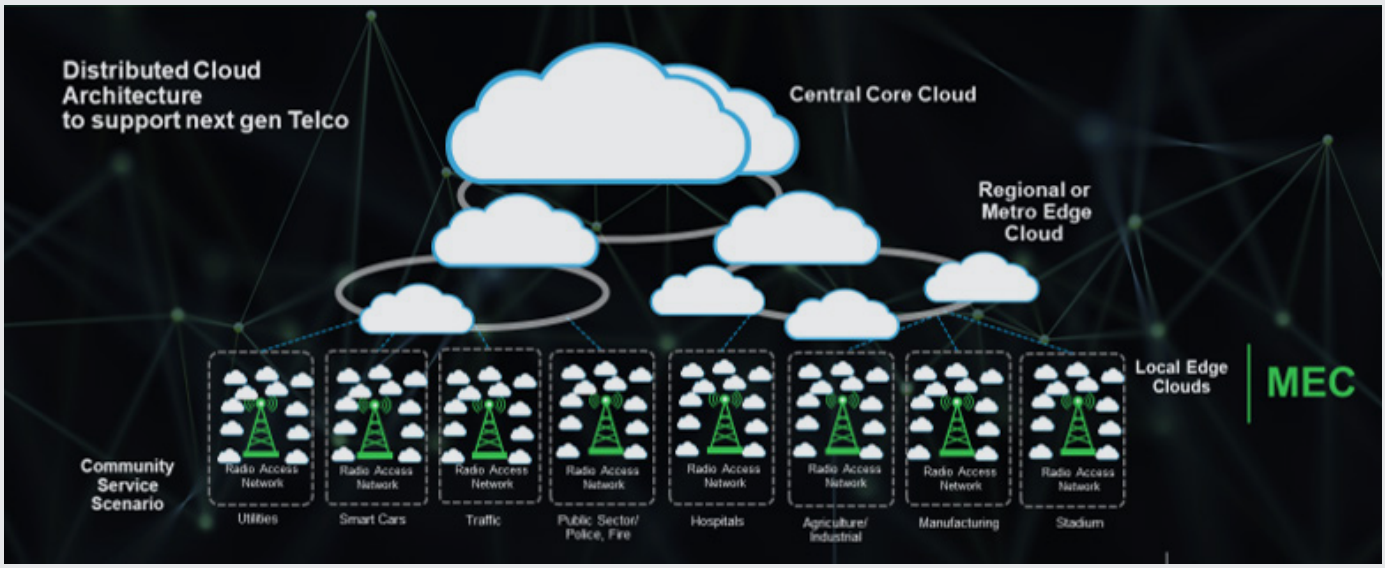

discussed, the energy use is being shifted from the core to the edge, and it's a top-of-mind issue to address energy efficiency. For 20 MW, 2,000 10kW micro data centers, a 20 percent efficiency reduction will roughly cost an extra $4 million per year. That's why it's important to make sure the designs of micro data centers include highest efficiency cooling technologies like liquid cooling, for example. And we can't forget the management and maintenance aspect. A cloud-based management system is critical for looking at thousands of sites. And accurate reporting is absolutely essential so it's clear when maintenance is needed, and it can be performed quickly and efficiently. This has been an area of focus for Schneider.

Can you provide updates on Schneider Electric Secure Power strategy for 2019, and key focus areas for next year?
Schneider is working across different stakeholders that are vying for a leadership spot in the data center architecture needed for 5G deployment. Telco used to be an exclusive country club with only a few members. But 5G, a new technology that embraces openness, is a public course where everyone from Tiger Woods to Happy Gilmore can show up with clubs. In addition to the carriers and traditional telco equipment providers, we are dealing with cable TV service providers, IT companies, and internet giants. Our strategy is to enable these companies to add value in the 5G ecosystem. We are focused on integrating the safest enclosures, latest battery technology, most innovative cooling, and cloud-based management systems. Of course, in 5G, as in golf, the best player will win.

Leveraging Cell Sites for Mobile Edge Computing

4.5-5G requires computational power in closer proximity to users - creating a unique edge computing opportunity for cell site owners. However, this transformation is not without its challenges. End-to-end solutions focused on power, cooling, enclosures, and management software are needed.

Schneider Electric
Contact Details
Steve Carlin
Steven.Carlini@se.com

Flood the world with fiber

If we can lower the cost to deploy fiber, we could connect the planet. *Sebastian Moss* reports

Sebastian Moss
Deputy Editor

Working in this industry, discussing the roll-out of 5G and the impact of a world full of connected machines, it can feel like we have already solved the basics of the digital age.

But we've left half of the planet behind us. "We've reached four and a half billion Internet users," Isfandiyar Shaheen told *DCD*. "How do we onboard the next three to four billion people on the Internet?"

This question, of how to connect billions and provide a stable connection for billions more, has consumed the lives of many. Giant corporations have sunk millions into ambitious schemes involving balloons, satellites and huge drones.

Shaheen, founder and CEO of NetEquity, believes there's a simpler approach: Fiber. "When you compare the bandwidth of fiber to everything else, whether it's a LEO satellite, microwave, millimeter wave - those things are not even in the same quadrant. It's literally, thousands of times less bandwidth than a single set of fiber."

His company believes that it is possible to reach a vast number of the disconnected or poorly connected people in the world by deploying fiber across utility networks. "The theory I'm working with is saying that if you follow the electrical grid to deploy fiber, you can fiberize 80-90 percent of cell towers in the world," Shaheen said

His hope is that he can convince utilities to deploy fiber along their grid infrastructure as "utilities need fiber to achieve substation automation," Shaheen said. "Utilities have also created a standard called IEC 61 850 that relies on fiber to do key things such as integrating renewables, integrating storage, and achieving a hell of a lot more automation to reduce their line losses, which will lower operating expenses and through which they can start lowering their capital expenditures. Because when a substation goes digital, its copper footprint goes down by 80 percent."

The utilities he hopes to convince are, however, "mostly bankrupt - they're running on state subsidies. So the solution I've come up with is to say to them 'I will build you a fiber network on my expense, and you get to lease it based on whatever the going rate for dark fiber in your market is.' It's usually about 60 cents to 70 cents a meter a year."

"That kind of lease payment makes it feasible to raise about 80 percent of the project costs as debt."

Once the network is underway, then it becomes easier to turn to telecoms companies and say "'hey I am bringing fiber to your tower, but instead of charging you whatever the market rate is for this service, let me make this an opex neutral deal for you. If you are spending $500 a month per tower on running microwaves, pay me the same $500 per tower per month, and I will give you a fiber connection,'" Shaheen said.

"But I won't give it as dark fiber, I will give it as a lit service. Because if I give it to you as dark fiber, you will lock away bandwidth. My goal is to make bandwidth abundant. So I would offer you abundant bandwidth, for the same operating costs of running microwaves. Now that combination of selling services to cell towers and getting these revenues from dark fiber leases from the utility that makes the whole business case hold."

By playing to the pain points of the utility and the telecoms company, Shaheen hopes to avoid the aggressive legal fights and turf wars that have stymied others' attempts at creating utility bandwidth.

A few examples of municipal fiber exist, with Shaheen pointing to efforts of small San Juan-based Orcas Power & Light Cooperative (OPALCO), which "deployed fiber at the cost of about $30 a meter.

"What's been cool about this cooperative is they published a very nice payback analysis in one of their board publications. And that shows a payback period of about 12 years. It shows the breakdown of the sources of savings, which came largely from saving the number of trips that their guys were making for repairs, because fiber gave them visibility across the grid."

> *"Currently, the dry powder of infrastructure funds is sitting at $180 billion"*

For Shaheen's plan to work, however, requires fiber to be deployed at a lot less than $30 a meter.

"Fiber cables are a commodity, it costs less than $1 per meter. Yet most deployments end up costing $30 to $50 a meter, and that's because all the cost is in labor and right of way. If there is a way to bypass labor and right of way, there's no reason you can't deploy fiber at $4 a meter or around that benchmark."

Shaheen believes that there is a way to avoid these cost barriers. But to achieve a radical price cut, Shaheen has to rely on an ambitious secret hardware project by another company to be pulled off without a hitch.

"What allows me to pull this stunt off is that I have this relationship with Facebook," Shaheen said. "I signed an MoU with Facebook that gives me access to some technologies that can substantially lower the cost of fiber deployments. It's not a completely ready technology, that's the risk associated with the project."

Facebook declined to detail the technology, which is currently under development. The project is part of Facebook's Connectivity division, home to several projects focused on improving Internet penetration, including an ill-fated drone project and the controversial Free Basics effort.

Shaheen met Facebook representatives at a Telecom Infra conference, advised the company as a consultant, and then became its first 'Entrepreneur In Residence.'

"I have no funding from them," Shaheen clarified. "This is important and by design, because if I were to have a deeper relationship with them, then I am subjected to their bureaucracy - they're a 60,000 person company. I don't need the entire company, I just need to talk to a couple of really solid engineers."

Those engineers are key to Shaheen's strategy, and roll-outs can't proceed without the technology working. "My bet isn't so much on Facebook the company, my bet is on who I am collaborating with."

In the meantime, Shaheen is working on "building a business around the promise," he said. "My timeline is such that for the next year to year and a half, some of the useful tech that Facebook is working on will get finalized and will get prepared. During that period, I want to generate a pipeline of deals, sign NDAs with utilities, get their grid maps, turn them into investor documents, and then start putting together investor consortiums."

So far, he says he has "made progress with utilities in Pakistan and South Africa," and is working with a large East Asian investment group to explore a partnership. If the technology works and $4 a meter fiber is possible, Shaheen is convinced he can sign up more investors: "Currently, the dry powder of infrastructure funds is sitting at $180 billion."

But to pull off that dream of connecting the world will require a large chunk of those funds. "There's 25 million route kilometers of power lines that are suitable for fiber deployments and at $4 a meter that's $100 billion; at $5 a meter, that's $125 billion."

Should he be able to line up investors, Shaheen hopes to "flood the world with fiber, make sure it's not under exclusive contracts, and create a true public utility that can that can help many billions of people access the Internet at a price that they can afford."

Routers in the sky

Terrestrial fiber has ruled the telecoms world for a generation. Now satellites are promising to match its performance, says *Doug Mohney*

Doug Mohney
Contributor

Gigabit-speed satellite broadband promising performance equal to or better than terrestrial fiber is almost here. Newcomer OneWeb and industry darling SpaceX are launching satellites in earnest by the end of 2019 to build global high-speed networks. Dozens of satellites at a time will go into space on each rocket with each mission building towards constellations of thousands of spacecraft overhead.

With full global coverage expected to be completed in 2021, OneWeb and SpaceX's Starlink project aims to extend broadband to underserved and unserved regions of the world, along with more profitable markets including aviation, maritime, government and enterprise sectors.

In contrast, LeoSat and Telesat plan to launch satellites in 2021 specifically designed for enterprise-class services - "MPLS routers in the sky" with "fiber-like" performance according to executives from both firms - with global service turnup expected by the end of 2022 to mid-2023. Looming behind them all is Amazon, with its own ambitions and needs.

"Our approach to [low earth orbit broadband] is we're building a Layer 2 satellite system," said Erwin Hudson, vice president of Telesat LEO. "We're designing our system to be compatible with MEF standards from the bottom up to provide enterprise-quality service with business class SLAs."

Appealing to the corporate IT department is a strategy start-up LeoSat is finding success with. "I can explain what we do to a data guy in five minutes," chief commercial officer Ronald van der Breggen said. "We're putting a bunch of MPLS routers in the sky, connecting them with lasers, and you can use them on any point in the world to connect. 'When is that available?' is the first question. We have a service which is unique, resonating with enterprise and governments."

Both satellite constellation designs include optical cross-links between satellites, a feature providing more speed and security over traditional fiber. Customers can choose to set up a connection between sites that exclusively rides over the satellite network, bypassing traditional network exchange points and reducing the number of hops found with terrestrial connectivity. Laser light traveling through the emptiness of space in a straight line between satellites moves faster than it does traveling through a glass fiber following a meandering path under the sea, through cities, and along railroads, highways, and gas pipelines.

Fewer ground hops mean fewer access points for interception or disruption, providing a level of resilience against the threat of backhoes and other fiber disruptions.

> *"Data centers are a key component of our network when it is fully rolled out"*

LeoSat and Telesat expect to have their respective networks ready to deliver service in the same timeframe, but the companies have very different financial paths to get there, with LeoSat facing a bigger hill in front of it.

"It's no secret, we've had difficulties in raising the equity," said van der Breggen, who left the company after our interview. "We had hoped to close the Series A right before the summer. Now we're working very hard to find additional investors to get us over the hump. There's a lot of interest from venture funds, strategic funds. We need another investor to line up with what we already have and get us over that hump."

LeoSat estimates it will take $3 billion to build and launch 90 satellites for its initial network. Satellite providers SKY Perfect JSAT and Hispasat placed early investments in LeoSat, but exact amounts have not been disclosed. The company also has logged $2 billion of customer commitments in Memorandum of Understandings (MoUs), with a diversified mix of firms including high-frequency trading markets, oil and gas, and telecom firms for good measure.

Broadband speed offerings for LeoSat start at 50Mbps and range up to 1Gbps, with the most popular requests in the MoU stack 100Mbps. Latency is expected to be in the 20 millisecond (ms) range for a simple trip up and down between ground and satellite, but van der Breggen played down fixating on the simplest of latency examples.

"It's a meaningless case and only tells you so much what an individual satellite is able to do," van der Breggen said. "The capabilities of a service are going to be far more important. You have to add in all the fiber [involved] as well as the satellite." Routing via satellite optical links between major financial centers is expected to be significantly faster than submarine fiber connections, he said, providing a competitive advantage to banks and stock trading firms.

In comparison, privately held Telesat is one of the oldest and largest satellite operators in the world. The Ottawa, Canada-based company publishes quarterly and annual financial reports since it issues publicly traded debt, providing a transparent window into its financials. Last year, Telesat closed around $680 million in revenue and had a contracted backlog for future services of nearly $2.8 billion.

Telesat plans to launch 198 satellites for global coverage in the first half of 2023, with service available in the polar regions by Q3 2022. An additional 100 satellites will be added by the end of 2023 for a total of 298 satellites in the initial constellation. User service speeds are expected to be scalable to the Gbps range with one-way latency "less than 50ms" in the same continental region.

Pricing to build the system, including satellites and ground equipment, and launch it into orbit hasn't been finalized with the constellation expected to cost "several billions of dollars," Erwin said. A primary

Source: LeoSat

contractor for the system should be selected by the end of this year, with Telesat asking the winning company to build a factory for its satellites in Canada.

"We intend to finance Telesat LEO with a combination of cash, equity and debt," said Erwin. "As a leading global satellite operator, we generate substantial cash flow and are able to provide a significant amount of funding for the LEO development ourselves."

With an established customer base and a large sales and marketing organization, Telesat also has the Government of Canada as an anchor customer, with a commitment of $500 million over 10 years as a part of efforts to expand broadband access. Canada will use Telesat to deliver backhaul services for ISPs and phone companies in under-connected communities.

Both LeoSat and Telesat see data centers as primary partners for delivering network access. "The short answer is data centers are a key component of our network when it is fully rolled out," van der Breggen stated. "A large portion of traffic will go into data centers."

Multi-tenant facilities such as Equinix are prime real estate, enabling a gigabit satellite service provider or a third-party handling the work within a physical facility to connect customers to high-speed interconnection points, existing terrestrial networks, Edge computing, and other colocated resources.

Telesat plans to establish PoPs at major Internet and cloud exchange points to interconnect with customer networks. "Our network architecture will drive more traffic to data centers on our global WAN," Erwin said.

Telesat also plans to work with data center operators as well as cloud service providers to simplify access to "cloud on-ramps" for its existing enterprise customers. "There's about a dozen network access points around the world where we aggregate three to five of our earth landing stations at a common point," said Erwin. "Regional customers can connect to our network at that point of presence."

However, data centers could be more than where LeoSat and Telesat spend money to connect into the rest of the world.

"We're not really focused on fiber replacement, but we can provide disaster recovery if [a data center] lost terrestrial connectivity," Erwin said. Telesat can deliver 1Gbps service using a 1m satellite dish, with higher speeds possible with larger dishes. Delivering 1.2Gbps to 2Gbps of connectivity using a 1.8m antenna is feasible.

While there has been skepticism in some circles about the commercial viability of LEO broadband constellations, LeoSat, OneWeb, SpaceX, Telesat received some validation and headache this spring.

Amazon announced Project Kuiper, its satellite broadband project, this spring after its ITU spectrum filings become public. The e-commerce and cloud giant wants to launch over 3,200 satellites for internal broadband use as well as to provide broadband to customers but hasn't discussed a timetable for when it will start putting hardware in the sky or offer services. Amazon may start launching satellites by 2023, but it isn't clear at this point in time if it will be a direct competitor to the enterprise-designed services of LeoSat and Telesat.

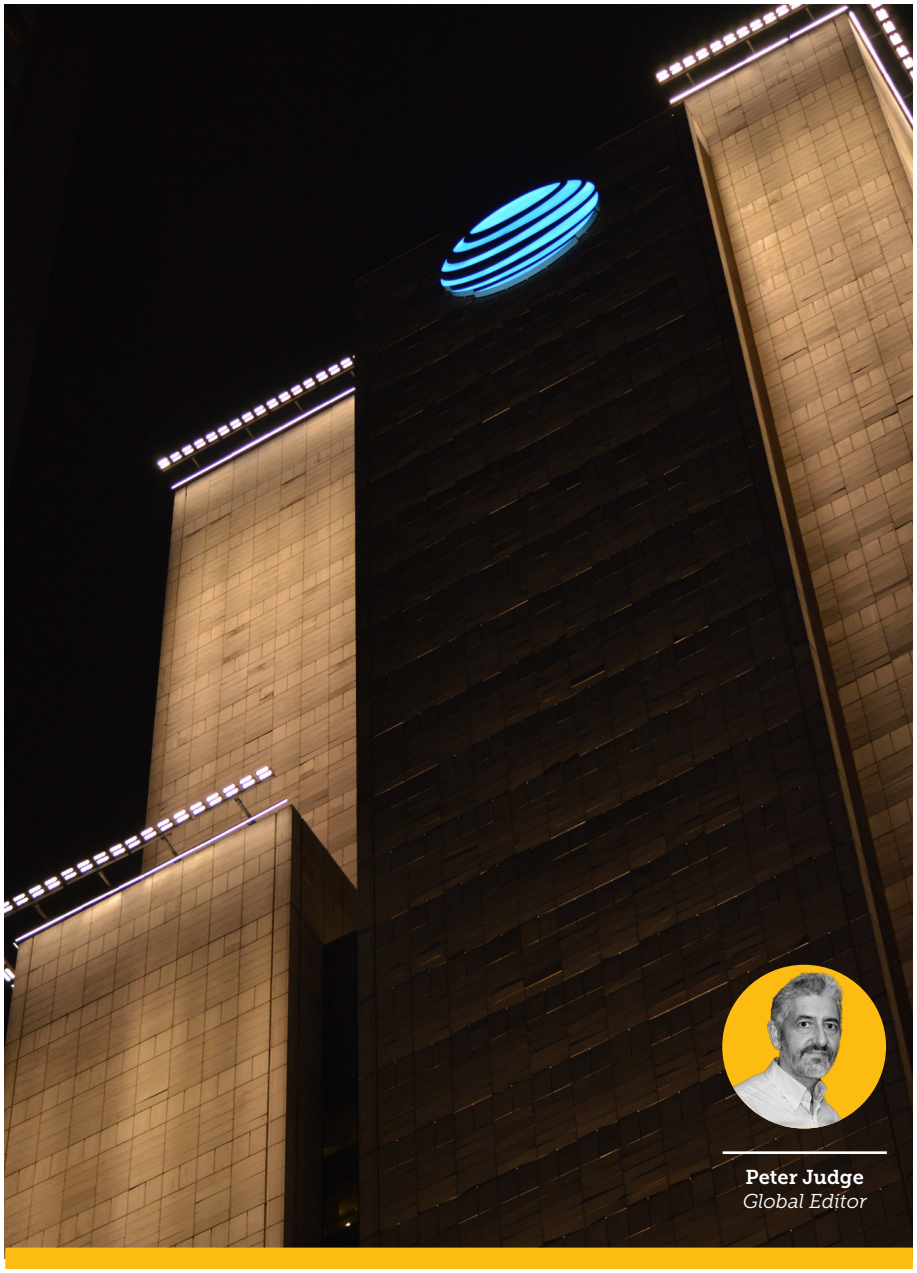

Peter Judge
Global Editor

The Telco data center sell-off

Telecoms providers thought they'd be great at colocation data center services. Now they're mostly getting out of the game, says *Peter Judge*

The news that Telecom Italia is looking to spin-off 23 of its data centers and list them on the stock market is only the latest in a series of moves, which are seeing telecoms service providers backing away from earlier plans to make a lot of money out of data center colocation.

It seemed so simple in the early years of this decade. Data centers were booming, and they are a service industry based on infrastructure hardware. To telecoms operators, it looked like a logical expansion, and many of them dived into the market. Ten years on, most of them are exiting.

"Despite many telcos making moves into the data center and cloud infrastructure markets, more and more are now realizing that they would rather concentrate on their core business and let someone else manage their data centers," says Massimo Bandinelli, marketing manager at Telecom Italia's compatriot Aruba.

Many telcos simply bought existing data center providers, often at high prices. Verizon, for instance, acquired data center provider Terremark in 2011 for $1.4 billion. Eight years later, the company decided that offering colocation services did not fit with its business model, and sold off its data centers to Equinix for $3.5 billion.

Also in the US, AT&T painstakingly accumulated a network of data centers, only to sell them off to Brookfield Infrastructure and other institutional partners for $1.1 billion in 2017. Brookfield relaunched them as a new data center provider, Evoque.

Also in 2017, CenturyLink sold 57 data centers for $2.3bn to a consortium that became another standalone data center provider, Cyxtera.

It wasn't a sudden change. *DCD* first noticed the phenomenon in 2015, when some smaller telcos unloaded their data centers. For instance, in that year Arkansas telco Windstream sold its holding of 14 data centers to TierPoint for $575m, giving that provider 179,000 sq ft (17,000 sq m) of space.

Rumors started about the imminent sales at the telco giants AT&T, Verizon and CenturyLink back in 2015, but took a couple of years to come to fruition.

The move took in telcos which had built out their own data centers, as well as those which acquired them. Telecom Italia, for instance, had at least some of its facilities built by a partner from the telecoms industry - Ericsson, which is primarily a network provider.

Also in Europe, Telefónica SA sold off its 11 colocation data centers. They went to Asterion Industrial Partners for €550m ($600m). In the UK, BT seems to have been

Photography: Sebastian Moss

selling its data centers off one-at-a-time in deals like a 2015 sale which saw a Tier III facility near Gatwick go to operator 4D.

The trend extends to younger markets as well, where the telcos' data center investment may have been much more recent. In Latin America, Mexican telco Axtel sold three data centers to Equinix for about $175m.

Indian telco Tata flipped dramatically. In 2013 and 2014, it saw a period of rapid data center expansion, building or acquiring 44 data centers in India and elsewhere in Asia. It then floated them as an independent subsidiary, Tata Communications Data Centres (TCDC). However, in 2016, Tata decided they were more of a liability than an asset, and sold TCDC to Singapore's ST Telemedia for around $650 million.

> *"The workforce required is very different from the one needed as a telco operator"*

It's easy in hindsight to think that these telcos stumbled into data centers by mistake, getting into an area they did not fully understand, where they would face more focused competition that could run rings round them.

That analysis is pretty much true, but in the latter years of this decade, telcos faced quite a few financial pressures, analysts have pointed out,

Back in 2015, *The Motley Fool*'s Adam Levy suggested US telcos needed money because they had paid heavily for wireless spectrum, but were not yet gaining huge revenues from mobile data. "Both [AT&T and Verizon] spent heavily in the FCC's AWS-3 spectrum auction, acquiring valuable airwave licenses for their wireless businesses," wrote Levy.

Heavily burdened with debt, many telcos have been looking for ways to raise money in recent years. And the data center field has been a good potential source of cash as, during a period of rapid growth, facilities are valued very highly.

While selling CenturyLink's data centers, company chief executive Glenn Post told Barclays Capital's Amir Rozwadowski that a lot of the motivation was simply that buyers were prepared to pay big money for those assets. And telecoms operators, who have been burnt in previous market crashes, have been understandably keen to get ahold of that cash, ahead of any potential future crash.

"First of all, as to why now is an opportune time... valuations are obviously good right now. They can always change, but we know the market's good." Post said, back in 2015. "We think our cashflow could be used for investments that can drive higher returns, and better shareholder value. So that's why we're looking at divesting data center assets."

But there are other reasons why data centers are not such a good fit for telcos as they might have been once. There have also been changes in the data center industry over the last several years, which have moved data centers ever further from the comfort zone of telecoms providers.

"The workforce required is very different from the one needed as a telco operator," warns Bandinelli, adding that this specialization has increased, as data centers have become more evolved and more commoditized.

There are also large investments required to keep up to date, adjusting to industry-wide regulations, adopting standards, getting certifications, and moving to renewable energy.

Alongside this, new business models are emerging like Edge resources, while at the other end of the scale, hyperscale providers are building a market for huge facilities which are not cost-effective for a services company to deliver.

When the sell-off began, Zahl Limbuwala of data center analytics company Romonet, now a subsidiary of CBRE, felt that the telcos may have taken a view that their data centers were likely to go down in value because of the investment requirements: "If your data centers are approaching 10 years old and have not had a major reinvestment, you are in for a nasty surprise," he said in 2015.

Some organizations have invested in data centers seeing them as a kind of commercial real estate with very large returns. Limbuwala pointed out that they also had high costs, with a "reinvestment time" before new investment is needed, of 10 years, about half that of mainstream commercial property.

Despite this gloomy picture, there are some telcos which are apparently exceptions to this trend. Japan's NTT has a thriving data center subsidiary, which has absorbed RagingWire in the US, NetMagic in Asia, e-shelter and Gyron in Europe, and is in the process of forming them into a single coherent unit.

However, NTT is an exception, and operates its data centers at arm's length. RagingWire CEO Doug Adams contrasts its approach with that of US telcos like Verizon and AT&T: "[The US telcos] were very short-sighted, very quarterly focused," he said in a *DCD* interview earlier this year.

"They were getting their tushes handed to them by the Equinixes, Digitals and RagingWires of the world, and they backed out. I think NTT was extraordinarily intelligent for doubling down on this business."

Whether it is standalone data center firms, or subsidiaries like NTT Data Centers operating independently, Bandinelli believes that the world is shifting towards pure-play data center providers that are able to meet market requirements and provide competent technicians at a lower cost.

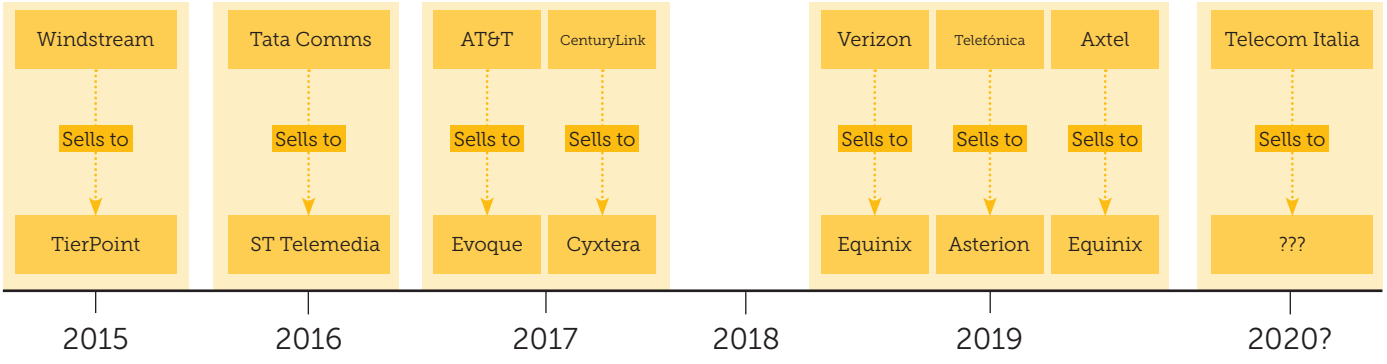

EcoStruxure™
Innovation At Every Level

Colocation providers see faster

GROWTH

with lower risk.

EcoStruxure™ for Cloud and Service Providers brings efficiency to any site, on time.

- Lower total cost of ownership with efficient, sustainable CapEx and OpEx management.
- Achieve higher operational performance while reducing energy costs up to 60%.*
- Get support throughout your data center's lifecycle — from design to construction and operation.

#WhatsYourBoldIdea

se.com/colo

** Based on previous data, 2017. This is not a guarantee of future performance or performance in your particular circumstances.*

Life Is On | Schneider Electric