DCD

The Enterprise Edge
Supplement

The reality of Edge computing

DCD

Contents

4

10

13

To the Edge and back

The rapidly changing geopolitical landscape means Edge computing deployments have taken on a new importance for many businesses.

As the impact of Donald Trump's second term as US president begins to be felt around the world, executives in companies large and small outside America find themselves grappling with issues of digital sovereignty, particularly if they are using US-based vendors in their tech stack.

Edge deployments are one way to ensure data is processed locally, in line with a particular country's rules and regulations.

OVHcloud, the European cloud provider, is keen to use this as a way to differentiate itself from the US hyperscalers. In this supplement, Georgia Butler speaks to the company about its Local Zones, Edge networks which have been rolled out across Europe and that could also prove useful to US-based customers looking to take advantage of Edge power and low latency.

Elsewhere, many consider that, years after its initial deployments, 5G remains a technology in search of a use case. Despite some success stories, the promise of high-speed, low-latency networks for industrial settings such as factories remains largely unfulfilled.

The Edge may offer an opportunity to change this by combining local processing with private 5G networks. Vendors hope that the opportunity to take data from a private network within a particular company and process it in real-time could form a compelling offer and deliver significant productivity gains for their customers.

But, as Paul Lipscombe writes in his feature on this fascinating topic, the enthusiasm across the industry is tempered by issues of cost and complexity, which will have to be resolved if companies are to be persuaded to ditch their tried-and-tested WiFi networks for more advanced 5G Edge solutions.

Space is perhaps the ultimate Edge deployment, and as the cost of launching a craft into low-Earth drops dramatically, a rapidly growing number of satellites are circling the globe.

With traffic in space getting denser and denser, the amount of debris floating above the planet is also on the rise, meaning the risk of damaging collisions is higher than ever.

Luckily help is at hand in the form of Portuguese startup Neuraspace, which is using telescopes - and Edge data centers - to analyze complex orbital data so it can provide information about dangerous debris to operators and help them avoid costly accidents. In this supplement, Jason Ma speaks to the company about the vital role Edge processing plays in its mission.

The amalgamation of cloud and the regional Edge

Georgia Butler
Senior Reporter
Cloud & Hybrid

How OVH's Local Zones are bringing Edge-like latency to its global cloud offering

The concept of cloud computing and Edge computing are, on the face of it, diametrically opposed.

One refers to a spider-web-esque network of usually massive data centers, and the other calls to mind the small containerized offerings we often write about at *DCD*.

But across the board, cloud providers have sought to close that gap. More traditional Edge solutions include the likes of Google Distributed Cloud or Amazon Web Services' (AWS) AWS Outposts, which see hardware deployed on-premises.

But another aspect of "the Edge" is being addressed through the concept of local zones: the regional Edge.

Different vendors have different ways to refer to local zones - Microsoft, for example, calls its version "Azure Extended Zones" - but the common thread is that these are smaller facilities, brought closer to metros or locations where a large region is not already established, enabling lower latency and data residency.

In terms of the big US hyperscalers, AWS first launched its Local Zones offering in 2019, initially focused on the US but with plans to later expand globally.

Microsoft joined the party a little later, opening its first Extended Zone in late 2024. Google has yet to explore a similar product for its cloud platform.

A company on the other side of the pond, however, is creating its own Local Zone empire.

European cloud provider OVHcloud started offering Local Zones in Europe in 2024, following its acquisition of Edge and hyperconverged infrastructure specialist Gridscale in 2023.

Gridscale offers multi-tenant infrastructure-as-a-service and platform-as-a-service solutions, and bases its

business model on its ability to provide public cloud services on a limited infrastructure.

At the time of acquisition, Michel Paulin, then-CEO of OVHcloud, said: "We are impressed by the solutions developed by the Gridscale teams, which we will be delighted to integrate into OVHcloud.

"The combination of our different technologies will allow us to accelerate our deployment of Edge computing solutions and our geographic expansion while optimizing our capital allocation."

The company's Local Zones offering has since been expanded to include the US.

Speaking on the Gridscale acquisition's role in OVH's Local Zones, Jeffrey Gregor, general manager of OVHcloud US, tells *DCD*: "We have leveraged some Gridscale technology to be able to quickly deploy these Edge and Local Zone offerings," adding that the company is not running them under the Gridscale name, however. "The acquisition did help us accelerate this launch."

While OVHcloud is known to many as a European cloud provider, the US business is entirely separate, incorporated in Delaware with its own board of directors.

"The reason we've structured it this way is that it is very important for OVHcloud in Europe to protect its European customers from the Cloud Act and the Patriot Act, so they don't want to be seen as a US company and doing business in the US," Gregor explains.

"That's why they, and all of their other geographic subsidiaries around the world, sell their services and operate in every market except for the US. But then we're also a global company. About a third of my customers are from outside the US because they want servers or services in our US data centers."

While this may seem like a slightly convoluted solution, such a move is prudent given the current macroeconomic and political atmosphere, the impact of which crystallized on a single day earlier this year.

The end of an era?

It was bitterly cold on January 20, 2025, in Washington, D.C., and not just for the Democrats.

> *"The choice of a cloud provider is no longer just a technical matter, but also a strategic issue,"*
>
> *>> Benjamin Revcolevschi*

longer just a technical matter, but also a strategic issue."

The concern is particularly clear between the US and the EU. The past few months alone have seen more than 100 organizations signing an open letter to European officials calling for the continent to become "more technologically independent" and saying the current reliance on hyperscalers creates "security and reliability risks."

European groups have concerns ranging from the potential for the Trump administration to use reliance on US tech companies - such as Google, Amazon Web Services, Microsoft, and others - as leverage for trade talks, to the White House's attacks on NATO, tariffs on the EU, and even threats to invade Greenland.

Hyperscalers, including Microsoft and Google have been scrambling to reassure EU customers. Google announced a series of additional elements for its sovereign

cloud in May 2025, just a few weeks after Microsoft's president, Brad Smith, penned a blog post acknowledging concerns and reaffirming the company's commitment to protecting EU customers.

"Given recent geopolitical volatility, we recognize that European governments likely will consider additional options," Smith wrote, avoiding directly addressing the cause of the volatility.

"Some of these may involve public financing to support European home-grown offerings. We recognize the importance of a diversified technology ecosystem, and we are committed to collaborating with European participants across the tech ecosystem."

But, by keeping its two businesses

Jeffrey Gregor

separate, OVH might hope that this will help protect it somewhat from the turmoil while not losing out on either side.

OVH's Local Zones are opening doors for the company, and fast. As explained by Gregor, the Local Zones, unlike OVHcloud regions, are in colocation facilities, making the launch process much simpler.

Speaking on what they look for from a colo provider, Gregor says: "It really comes down to, do we have or can we get the network connectivity? Starting up is very easy. We were just putting in handfuls of racks to start, and then we can expand from there in the areas where we see success."

Because of this, Gregor adds, sometimes the colo data centers aren't prepared for OVH's proprietary hardware, so in those cases the company instead uses "commercial third-party servers."

The Local Zones, in terms of offerings, are "much more limited."

"It's public cloud - so compute, and storage, and S3 compatible object storage, but it's sort of a minimalist key building blocks approach," says Gregor. "The focus is really on geographic dispersion and getting closer to end customers or workloads - Edge computing or lower latency."

At the time of writing, OVH has Local

> "The beauty of the system is that it is still offering public cloud. You can decide to open and close a virtual machine without any problems."
>
> >> Mathieu Colas, Starzdata

Mathieu Colas

Zones live in 23 metro areas in both the US and Europe.

The first two to launch were Madrid, Spain, and in Brussels, Belgium. An early adopter of the offering was French Software-as-a-Service provider Starzdata, a company that provides search and data enrichment services, and increasingly uses AI to do so.

"I was aware of Local Zones before the market, because we have a close relationship with OVH, so I knew when they acquired Gridscale and the existence of Local Zones, where and how they wanted to develop them," Mathieu Colas, founder of Starzdata tells *DCD*.

For Starzdata, data residency and sovereignty are key - as many of their clients are either large consultancies or private-equity owned clients in the B2C sector.

The company first started with OVH's Madrid and Rabat, Morocco, Local Zones, says Colas. "The beauty of the system is that it is still offering public cloud. You can decide to open and close a virtual machine without any problems."

With many of Starzdata's customers being "project-based," this gives the company the flexibility to open instances in new locations as needed.

But, beyond that, the benefit of a Local Zone rather than a cloud region is latency-based - and where the "Edge" really comes into play. According to Colas, the company is aiming for around 50 milliseconds, but says it depends on the rate limit and relationship with data providers.

"We rely on third-party APIs to process the data, and that means we have rate limits, and every 100 milliseconds of difference has a huge impact," Colas says. "And because we work with CxOs or a board, you have to bring a reliable answer. You can't say, 'Oh, sorry, there's a problem, let's wait a couple of days'."

Colas adds that the Local Zone enables them to get both close to their clients, but also the data providers they work with, meaning they have "complete control over the latency between ourselves and the data provider."

Starzdata is also working on penetrating the US, where it again plans to utilize OVHcloud Local Zones.

Regional Edge doesn't have to be "small"

Local Zones, beyond offering the cloud experience merged with a regional Edge solution, are enabling OVHcloud to strategically test the waters, particularly in the US.

The company, while it launched a data center in Beauharnois, Canada, in 2012, only formed a US subsidiary in 2017.

That subsidiary has two "regions" or OVH data centers - one in Vint Hill, Virginia, and another in Hillsboro, Oregon. On the Local Zone side, it has a presence in Seattle, Washington State; Palo Alto and Los Angeles, California; Salt Lake City, Utah; Denver, Colorado; Chicago, Illinois; Dallas, Texas; Atlanta, Georgia; Miami, Florida; New York; and Boston, Massachusets; and has zones set to come soon in Phoenix and St. Louis.

As Gregor puts it, they are prioritizing "geographic coverage" and, at this point, have "East, West, North, and South covered."

Actually entering these locations, OVH looked at where they already had Points of Presence (PoPs).

"If you're going to build a data center, it's got to be of a certain scale or have the ability to grow to a certain scale. We can't, or don't want, to just put those everywhere," he says.

"Local Zones, from the company's point of view, allow us to go into other markets with a very small cost or investment to get that market, and test out and see what demand is," Gregor explains, adding: "Then, maybe, we'll come back and decide to put a full OVHcloud data center there."

He continues: "We can expand and grow within these colocations along with market demand, while we find the land and suitable space to actually build a data center."

In the interim, however, OVH's intention is to potentially start adding some other OVHcloud products in these colocation facilities, until we decide to move to a full OVHcloud data center.

Specifically for the Local Zones, OVH has "some hypotheses" about the market, Gregor says.

"We really want to see some Edge gaming, because we are really strong in our private cloud the bare metal space, and some top game hosting companies already use OVHcloud." The Local Zones also offer a CDN-type service, ideal for streaming video and music.

"Or, for example, e-commerce or some kind of real-time analytics needing a small cluster with local processing of some IoT data, even that then could be sent back to some central database at a later time."

Thus far, Gregor says the company is really happy with demand from its two main US data centers in Virginia and Oregon, but that currently there isn't enough demand to have a full "central US" region in Chicago, Dallas, or St Louis. But he remains confident the firm's current strategy will enable that in the future as demand grows.

"This is going to help us test that out," he says. "Now that we've got those Local Zones, it can help us prove that business case." ●

DCIM and lessons from Jaws

When threats are headed for your IT infrastructure, stay calm and manage on, says Kevin Brown, SVP and CMO, Secure Power, Data Centers and Global Services at Schneider Electric

Kevin Brown, SVP and CMO, Secure Power, Data Centers & Global Services

The thriller movie Jaws is celebrating its 50th anniversary this year. Like many others I love this story – the water looks calm, but danger lurks under the surface as a great white shark makes local waters its home, terrorizing a beach community. Characters with different personalities approach the threat

differently, leading to chaos.

I believe this is a great analogy to how organizations may have different 'personalities' in responding to the ongoing threats to their IT infrastructure. These threats are like the shark in Jaws – you know they are out there, you generally can't anticipate them, and when they attack, they may kill you.

Unlike in Jaws, there is more than one shark lurking under the surface.

For CIOs, multiple threats are lurking

We know CIOs need to improve how they manage their sprawling, hybrid IT infrastructure and understand the number of assets they have on the network. They need to maintain resiliency and avoid outages that may disrupt business continuity, resulting in lost revenue and potential reputation damage. They are grappling with cyber threats and potential

attacks while trying to make sure firmware is up-to-date and that they are following their company's security policies. If that wasn't enough, we argue, they will be increasingly questioned on the energy consumption of their IT infrastructure and may, if global trends continue, need to report their sustainability metrics.

Many CIOs are turning to DCIM (Data Center Infrastructure Management) software, EcoStruxure™ IT specifically, for answers. I believe Jaws – the personalities of its central characters and the threat of the cinematic shark (very different from actual sharks) that made everyone scared to go in the water – offers a few lessons about managing a fearsome threat. I'm not sure spoiler alerts are needed for a film that is 50 years old but consider yourself warned.

Probing for IT infrastructure vulnerabilities to exploit

Jaws has three central characters, and they are all determined to defeat the shark, which swims around, probing for vulnerabilities to exploit. Each character is essential to the outcome of the story, with their own unique approach to finding and killing the shark.

Oceanographer Hooper is fascinated by the shark and is an intellectual. More comfortable on yachts than fishing vessels, Hooper tends to depend on technology to address the threat. He has the latest electronic equipment and, in a famous scene, he sinks into the water in a shark cage. In the book, unlike the movie, Hooper dies after going in the cage.

Quint is a modern-day Captain Ahab. He is obsessed with sharks and, in another famous scene, he details his experience on the U.S.S. Indianapolis when it sank in July 1945. Quint is 'old school'. He rejects Hooper's technology, relies on an old, trusty boat with barrels, and is hell bent on revenge of all sharks. He respects the shark but believes in, and only in, his tried-and-true approach. Quint gets eaten by the shark in both the movie and the book.

Police Chief Brody is arguably the most traditional character. He sees the chaos the shark is wreaking on the community and works to restore authority and calm from the chaos. He is concerned the community won't survive without the summer season tourism. He is pragmatic and working hard, trying to find ways to solve the problem. Drawback on Brody – he fears

Getty Images

the unknown – he doesn't know how to swim and tends to be slow to move. Yet he eventually faces these fears and brings his pragmatism to the other characters, telling Quint, in perhaps the most famous line from the movie, "you're gonna need a bigger boat". Brody learns from Quint and Hooper and becomes the hero who kills the shark and survives to fight another day.

Is your organization like Hooper, Quint, or Brody?

So, what can we learn from the approaches of the main characters – what do their personalities mean for the threat and what is a best practice? When it comes to the threats facing your IT infrastructure, are you like Hooper, Quint, or Brody?

Hooper is forward-thinking by looking to technology for answers, but he is arrogant in believing that technology alone will defeat the shark because it isn't enough. At the opposite end of the spectrum is Quint, who knows so much and has loads of experience but learns his old school ways aren't sufficient for the threat. Then there is pragmatic Brody, who evolves as the film progresses. Learning from Hooper and Quint, Brody perseveres, overcomes his fears, and lives to tell the story.

To create a well-rounded approach, I believe we need to pull from the best traits of Hooper, Quint, and Brody: embrace new technology like Hooper; respect the tried and true that has worked in the past like Quint; and be pragmatic, hardworking, and diligent like Brody. I would argue that most IT organizations swing between these three personalities, driving a multitude of debates on the best approach . . . like the chaos the shark generates in the beach community.

How DCIM can help CIOs with their IT infrastructure

DCIM is a tool that can help CIOs face the sharks lurking out there as they constantly move, feed, and look for their next victims.

Like Hooper, our team at Schneider Electric is using the latest technology in our EcoStruxure IT DCIM portfolio to address threats. We have built in new tools to help identify whether devices are in compliance with policies, we are working on AI algorithms to identify anomalies, and we have added new sustainability capabilities to help with reporting metrics. We are leaning into new technologies

Getty Images

and identifying new threats to help our customers.

Like Quint, we are building the basics into our DCIM tools. For example, we make it easy to deploy firmware updates, are improving our usability and user interfaces, and are striving to simplify our portfolio and our business. The EcoStruxure IT team is focused on getting the basics right and continually improving because having a solid boat with a bunch of barrels is part of the overall approach.

Like Brody, we recognize we need to be pragmatic. We understand the world is evolving, which is why we have invested in obtaining IEC 62443 certification for our Network Management Card. We are also ISO27001 certified for our IT Expert cloud products, EcoStruxure Data Center Expert is FIPS 140-3 Validated, Certificate #4974, and we are implementing the features needed to deploy at scale for the new challenges of the hybrid IT environment.

The combination of all three 'personalities' assures us that we are helping customers meet their IT infrastructure challenges and deal with unknown threats, so they aren't afraid to go in the water. Don't let the summer pass by without learning more about our EcoStruxure IT offer and how it can help your organization. ●

Life Is On | **Schneider** Electric

Private 5G networks and the Edge opportunity for telcos

Take-up of private 5G networks has been slow, but does the Edge offer an opportunity to change that?

Paul Lipscombe
Telecoms Editor

Getty Images

Perhaps the 5G opportunity for telecom carriers has been misunderstood all along.

While mobile carriers have struggled quite publicly with the high costs and low returns of their respective 5G mobile networks, there's a chance that private 5G networks could provide some salvation.

Private networks at the Edge, more specifically, could be a key driver for telcos finally monetizing their 5G networks.

Public or Private?

A private mobile network is based on an isolated network and designed to serve a specific purpose for a particular business or use case, such as a factory. By contrast, a public network is wireless connectivity that is available to everyone, delivered to an end user's handheld device via the mobile carriers in that particular market.

Private networks offer greater control over who can access the networks, meaning they have an added element of security for enterprises.

According to Téral Research, as of last year, just over a quarter (26.7 percent) of total private wireless worldwide was

based solely on Private 5G. The majority (52 percent) utilize LTE, considered to be a "transitional" 4G technology.

Private 5G offers faster speeds, lower latency, and higher bandwidth than LTE networks, and has been tipped to support Industry 4.0, providing data-driven automation, advanced analytics, and smart sensing to boost productivity. These private networks have been deployed to support a range of industries, including smart factories, military and defense, and mining.

> *There have been some quite dodgy forecasts out there around private networks that were predicting just a huge, huge, huge size of the market. And so there was all this kind of hype,"*
>
> *>> Rosalind Craven*

Private 5G take-off has been slow

These 5G private networks haven't taken off as anticipated, notes Stéphane Téral, founder and chief analyst at Téral Research.

"5G is coming, but in many instances, you don't need the performance of 5G or all the features of 5G," he says. "It's in factories when you have robots that need to communicate with each other, and you need very stringent latency."

Rosalind Craven, principal analyst, private networks and telecoms, STL Partners, agrees.

"If you go back maybe three years, the hype around private 5G was at its peak," Craven tells DCD. "There were some really promising deployments happening, and frustration already existed about 5G not living up to its potential and finding that sweet spot for monetization."

This combination of factors created hysteria across the industry that private 5G would provide the answer for 5G's monetization headache, Craven argues. She notes that this wasn't helped by some of the predictions around private 5G.

"There have been some quite dodgy forecasts out there around private networks

that were predicting a huge, huge, market size," Craven says.

"One of the things that has been learned slowly about private networks is that they're not just a shrunk-down version of a public network. The demands are different, the engineering is different. The things that you have to consider are different, and these factors are also different from private network to private network."

A Dell'Oro report in April suggests that the tide is turning for private networks; however,

private 4G/5G RAN sales increased by more than 40 percent last year, reaching 3-5 percent of total RAN sales.

Time for the Edge

Technologies such as 4G, more specifically LTE, and even 2G have been used to create private networks before 5G came along.

So, where does Edge computing come into play, and how can it support private 5G?

While private 5G networks provide the required low-latency wireless network for localized areas, Edge computing provides the element of data processing, bringing this closer to the devices that utilize the data.

"A few years ago, multi-Edge computing (MEC) was seen as a strong market, growing rapidly, but it did not grow as expected," notes Téral, who refers to the Covid-19 pandemic as slowing down the take-up as workforces remained at home and not at their offices.

MEC is a type of Edge compute technology centered around mobile communications.

Unlike conventional Edge, which tends to process data closer to the source, MEC leverages mobile network infrastructure to enable applications at the Edge.

"We see many people going back to the office after the pandemic," he adds. "Because of this, enterprises are seeing a rise in devices in their offices, and that requires compute power close to the Edge. Because you want to connect all these devices, these devices require compute power nearby."

There appears to be an opportunity for the Edge to play a key role in these 5G networks.

A report released by STL says that private 5G Edge can provide both enterprise and telecom companies with opportunities for monetization.

Both private 5G and Edge can be deployed separately, but there's now more of an appetite from enterprises to deploy private networks and Edge together, notes the report.

Together, they're tipped to generate revenue streams for enterprises and telcos and enable use cases around predictive maintenance, AR/VR, and technology around AGVs (Autonomous Guided Vehicles).

"If you also have Edge expertise, the two things go hand in hand," explains STL's Craven, who notes encouraging talk around private 5G and Edge being a platform that can bring AI into an operational technology (OT) world.

She says that the private 5G Edge is going to be essential to prop up some of these use cases that the enterprise demands.

"When we talk about the private 5G Edge coming into its own, it's particularly things that are leveraging computer vision, such as potential industrial use cases."

Craven says it will play a big role in supporting advanced predictive maintenance and using AI for process monitoring and quality control, where data can be analyzed in real time, and where problems can be predicted before and issues can be found immediately.

One particular use case that is a proven example of private 5G Edge, Craven points out, is around AGVs and also AMR (Autonomous Mobile Robots).

AGVs are essentially robots that follow a set path or track for material transportation, while AMRs are autonomous robots designed to navigate uncontrolled environments without the need for a fixed route.

Both have been deployed to support logistics firms and manufacturers in settings such as warehouses and factories. An example of this can be seen at an Amazon warehouse, where AGVs are used to increase efficiency and reduce manual labor.

Nokia's private Edge

One such vendor that has a heritage in the private wireless space is Nokia.

Indeed, the Finnish company was one of the first to deploy a private wireless network, doing so back in the early 2010s for mining giant Rio Tinto in Australia.

> *"As we see more and more networks deployed, the cost will go down, because right now the cost is still significant, so you'd better have very specific requirements to justify the cost of deploying a private 5G network,"*
>
> >> *Stéphane Téral*

Fast forward to 2025, and Nokia says it has deployed private wireless networks for 890 customers globally, with 24 percent of those in North America.

Téral says Nokia is a pioneer in private wireless, referring to the company's longevity and vast number of deployments.

To further support the company's private 5G and Edge ambitions, the vendor signed a partnership with digital infrastructure and smart building solutions provider Andorix in June.

"Private 5G networks are transforming how enterprises operate, and in-building connectivity is a critical piece of that puzzle," said Willie Kopp, head of enterprise campus Edge sales, North America, Nokia, at the time of the deal.

"By partnering with Andorix, we are combining 5G technology with proven in-building converged network deployment expertise to support the growing demand for private 5G connectivity within the built environment in North America."
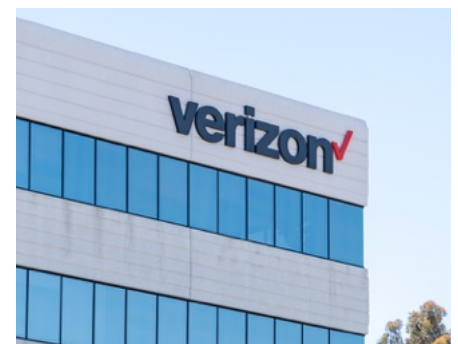
Getty Images

Getty Images

> "Private 5G networks are transforming how enterprises operate, and in-building connectivity is a critical piece of that puzzle,"
>
> >>Willie Kopp, head of enterprise campus Edge sales, North America, Nokia

Kopp elaborated further on the importance of Edge for some of Nokia's private wireless customers in an interview with *DCD*.

"Having that far Edge compute is still something that is extremely valuable to a lot of our customers," said Kopp, who highlighted the importance of Edge for the optimization of AI.

"I think AI is great when you have a bunch of Nvidia H200 [GPUs] in a datacenter that are training a model. But once that model is trained, it needs to be implemented and deployed at the Edge so that you can actually gain the insights and the values from that inference model. And that's really where we start to see Edge compute become extremely valuable and actually a requirement for a lot of customers."

Telcos also deliver on the private Edge

Mobile carriers are also taking the private 5G Edge opportunity seriously. US giant Verizon is one such telecom provider pushing hard in this area.

Hans Vestberg, CEO, Verizon claimed last year that the carrier is best placed to maximize the potential of generative AI Edge computing, noting at the time that Verizon had been talking about the opportunities around mobile Edge computing since 2018.

As it happens, some of those opportunities have centered around private networks. Last year, the carrier signed a deal with GPU giant Nvidia to develop a range of AI applications that will run over its 5G private network with private MEC.

According to Verizon, this will "deliver real-time AI services on premises for enterprise customers."

The telco said at the time that the stack is being built to handle "compute-intensive apps" such as Generative AI LLMs and Vision Language Models, video streaming, broadcast management, Computer Vision (CV), AR, VR, XR, AMR/AGV, and IoT.

Verizon has deployed a number of private 5G networks across the US, utilizing its MEC capabilities, including at a veterans' hospital in California two years ago.

Barriers to wider adoption

Despite this momentum across the industry, the deployment of private 5G Edge networks hasn't been as popular as expected. This is because multiple barriers to deployment remain, Craven says.

"Cost and complexity are both issues that we see being gradually tackled by the market. There is an awareness that for many enterprises to become open to it, costs have to come down," she says.

Should the costs come down, Craven says more new entrants could join the market, and even encourage some of the more established names to take the plunge.

A change of mindset from telcos and vendors to offer more tailored solutions should also be considered, she adds.

"The telcos and some of the big vendors who've had some success have maybe been getting in their own way in how they approach enterprises," she adds. Instead of selling a horizontal connectivity solution to the enterprise, the approach should be more vertically tailored, Craven argues.

There's also the debate around whether

private networks and WiFi can co-exist with one another, or whether it's a case of "either or."

Craven suggests that WiFi is capable of deploying AGVs, for example, and works adequately instead of Private 5G networks in certain situations.

The "private 5G versus WiFi fight" has been "problematic," she adds, explaining: "People should be working more on how to offer you the right connectivity for the right situation, sometimes it makes more sense to use WiFi, and other times private 5G."

Falling into place

While 5G has split opinion since its inception, with many, Téral included, adamant that launching the network as 5G Non Standalone (5G NSA) was a mistake, he says there are signs that the technology is falling into place, and that the number of 5G private networks will flourish.

Berg Insight, an IoT analyst firm noted that 4,700 private LTE/5G networks were deployed across the globe by the end of last year, at a market value of an estimated $1.8 billion.

Téral is, however, optimistic that further deployments will follow, despite the slow start.

"When a private network is deployed, it works and it delivers to the performance required. So this is proven," he says, but does caution that costs are still high.

"As we see more and more networks deployed, the cost will go down, because right now the cost is still significant, so you'd better have very specific requirements to justify the cost of deploying a private 5G network. ●

The Edge of the universe

Jason Ma
Junior Reporter

Portuguese startup Neuraspace is using Edge computing - and some impressive telescopes - to avoid collisions in orbit

Visualizations of Earth's space debris resemble an ever-moving, ever-growing piece of chainmail. As more satellites are sent into space, the higher the risk of an accidental collision. And with more collisions, more debris is produced. The European Space Agency (ESA) said that around 40,000 objects are currently in Earth's orbit, but their projections have shown that this number has and will continue to grow at an exponential rate.

Most debris is produced when the payload or the rocket to which a satellite is attached fragments, but satellites themselves can break under the wear and tear of the space environment. For instance, in October 2024, Intelsat-33e, operated by Luxembourgish-American operator

Intelsat, disintegrated and released over 700 pieces of debris into geostationary orbit (GEO), which is where most communications satellites are placed.

It is hard to find an earthly analogue for the kind of danger posed by space debris to spacecraft. Everything is simply more consequential in space. The ESA says collisions producing debris as small as 1cm could disable a spacecraft and potentially cause further damage. Debris larger than

10cm – the length of a small spoon – could completely destroy a piece of spacecraft.

Neuraspace, a Portuguese space software company, is using Edge computing to give satellite operators better tools

3DSculptor/Getty Images

Neuraspace

> *"An effort to make space more sustainable benefits everyone"*
>
> >> *Rob Arthurs*

company analyzes Space-Track's CDMs against data submitted by operators – Arthur said that Neuraspace currently tracks more than 450 spacecraft – alongside data from a number of companies that have optical, radar, or laser tracking facilities; and space weather nowcasts and forecasts. The company intends to integrate data from the EU Space Surveillance and Tracking platform next.

Despite being a software company, Neuraspace also gets data from its own telescopes, which is where the Edge comes in. In September 2024, Neuraspace installed its first telescope in a Portuguese military site, and delivered another to a Chilean telescope park in December that year.

But why would a software company need its own telescopes?

Arthur gave three reasons. First, there are simply no private operators that operate large networks of telescopes and lease out usage at a fixed price point. Comparing it to "climate change or any other thing where it's a tragedy of the commons," he says: "An effort to make space more sustainable benefits everyone, but if you're the only person who's putting the resources in, then it doesn't make enough difference."

Santos adds that you have to think about it like air traffic management. "When there were just a few airplanes flying around there weren't control towers, radar tracking everything, and controllers checking all the time if there was a risk of collision and so on," he says. "As the industry matured, the traffic became so intense that you started needing something that would regulate all this traffic above you."

Second, Neuraspace felt that it was necessary to have its own data. Arthur says: "If we want to use it for experimentation, testing, machine learning, there's a limit as to how much you can pay for the right data." Having its own telescope means that the company might also be able to sell that data in the future.

The third reason is supply chain management. "If you're buying a service

---

to navigate what will quickly become the world's most dangerous highway. They track their customer's satellites using containerized telescopes and publicly available data, process it, and then analyze it using machine learning powered by the cloud. Incubated in the Instituto Pedro Nunes – an innovation hub established by the University of Coimbra – and the ESA's Business Incubator initiative, the company has been operating since 2022, and has since signed up clients including on clients like American satellite operator Spire, French satellite maker U-Space, and the European Space Agency.

Feel the force

For anyone who wants to send a satellite into space, the backbone of global space traffic management is currently provided by the US Space Force, which runs a platform called Space-Track that monitors and predicts potential collisions. When the system detects a chance that two objects may be getting too close to one another,

it sends conjunction data messages to operators, who assess the threat and plan evasive manoeuvres if necessary.

But the problem, in the gingerly phrased words of Neuraspace's senior orbit dynamics engineer Rob Arthur, is that Space-Track is "non-user-friendly." CDMs consist of raw data, which operators have to parse in order to make sense of the information. A CDM is not a surefire guarantee of a collision – after all, 'conjunction', as defined in astronomy, is when two objects in space appear to be close together – and operators may get various CDMs pertaining to the same conjunction. Operators have to sift through repeat notifications in order to ascertain whether an object poses a credible threat to their satellite, which is time-consuming. Arthur adds: "Often you'll have a warning, and then you've got another one and another one."

Neuraspace's solution is to create a separate platform that does a better job of interpreting that data. This means that the

from another operator, obviously you're at the whim of whether they've got capacity and things," Arthus says, explaining that there currently aren't many suitable telescopes for hire on the market. "Right at the moment, there's just not that maturity of knowing exactly how much you have to pay and what you can get," he adds.

Processing at the Edge

All of this data is then initially processed at the Edge. Neuraspace's Edge capabilities sit in a white container located directly below the company's orb-shaped telescopes. Miguel Santos, Neuraspace's orbital dynamics engineer, says that the container is equipped with all the instrumentation that is required to operate a fully automated, remote satellite observation system.

"This includes the full optical assembly (telescope, mount, camera), meteorological instruments, environmental controls, and all IT infrastructure for local processing, monitoring and secure communications (firewalls, switches, servers)," he says.

"We process the images that we are taking in real time to generate the main product, which is the CDM."

The telescope first takes a long series of images. Some astrometry – a term describing the measurement of the positions and movements of celestial objects – is done to find the exact position of the satellite by cross-referencing it to other stars and satellites. Sensors, which are also present in the container, track information like the quality of the sky and environmental variables. The two are then combined to generate "smaller files that contain the actual positions and brightness of the satellite", which is then pushed onto the cloud for further analysis in relation to all of Neuraspace's other data sources.

But once that is finished, data is sent from the container to the cloud in order to power computationally-intensive machine learning processes that enable the data to be properly analyzed. AWS is their chosen partner, and the hyperscaler helps them sift through mountains of data – Neuraspace has two telescopes, each of which can track a maximum of 500 objects that have hundreds of relevant measurements – and cross-check their trajectories with about 20,000 other objects, other CDMs and orbital data messages.

One might wonder why the Edge might be necessary given Neuraspace's inevitable usage on the cloud. Why not just use the

cloud for everything? Santos' reply was simple: "Our telescopes operate in remote locations where network bandwidth is limited, making it impractical to transfer large volumes of raw image data in real time." The Edge also helps Neuraspace reduce data volume significantly before it is transmitted to AWS. Santos makes a point of mentioning photometric (referring to brightness) and astrometric (positional) data, which can be unwieldy if unprocessed.

Long before the ChatGPT hype-train barreled through the collective consciousness, organizations interested in space tracking like Neuraspace had already realized the benefits of AI. But Arthur says that other uses for AI are being explored.

"One is this kind of classification/ prioritization assistant, where [the AI] looks at all this conjunction data coming in, and it's finding out which ones are worrying and which ones are not," he says.

"Another one is an atmospheric model. One of the big uncertainties with satellites in low orbits is the kind of residual atmosphere and it's very hard to predict what the density is. So there are analytical, scientific-based, specific space models out there, but we're looking at whether AI can actually be used alongside those. It's incredibly complex as a human to do that – the equations and the physics to do that would be a nightmare, but maybe AI can pull out certain patterns."

> "We process the images that we are taking in real time to generate the main product, which is the CDM"
>> Miguel Santos

According to Neuraspace, interest in the product is coming from everywhere. Aside from private satellite operators, insurers might also be interested because the platform would enable them to assess and price risk in a clearer fashion. But they also see great interest from defense companies and various national air forces.

The geopolitics of space

As space becomes the newest geopolitical frontier, interest from governments and defence providers have spiked. After all, space infrastructure can be invaluable in times of crisis – Elon Musk's Starlink satellite Internet service has kept Ukraine connected during the war with Russia, despite some controversy around the role of the company in the conflict – and the rise of new threats like deliberate interference means that Neuraspace sees a demand for services that can help cut through "dodgy data".

The company has already begun moving in the direction of defense. In September 2024, it joined the European Military Integrated Space Situational Awareness and Recognition Capability project (EMISSARY). With €158m at its disposal, EMISSARY intends to develop autonomous European capabilities for Space Situational Awareness (SSA) for defence and security purposes.

Arthur says: "You've got situations where people may not want you to know what they're doing, or they might be doing something that's a bit shady. When you've got a large number of objects, and one of them suddenly does something very strange, machine learning can pick that out and flag it up and go: something weird is going on here. Then you can get your own analysts to look at it."

What the operator chooses to do with the information provided on the platform is up to them. But in the future, Arthur said that the company might work more closely with providers to interface directly with the satellite.

"I think that's definitely a longer-term ambition, and we've got maybe one or two things. One of the research projects we're looking into is looking at what we should do on board a satellite to efficiently help its operations."

But for now, Neuraspace will continue to operate outside of the satellite – and at the Edge 🟠