Edge and energy efficiency dominate

DCS talks to Giordano Albertazzi, President of Vertiv in Europe, Middle East and Africa, about the company's current focus in the data centre industry. Edge computing, energy efficiency, colocation, Cloud and security are amongst the topics discussed.

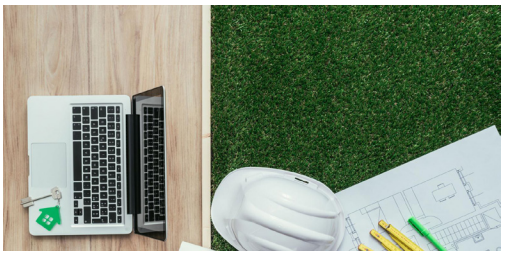

1. Edge computing is one of the major data centre developments at the present time. What work is Vertiv doing in this area?

Vertiv's edge experts, in conjunction with an independent third-party consulting firm, recently conducted a global, research-based analysis of network edge use cases. The research identified the main archetypes for edge applications and the technology required to support these. It also highlighted 24 use cases considered to have the greatest impact on businesses and end users, based on projected growth, criticality and financial impact.

The research aimed to provide more clarity on key edge use cases and the implications for the design and operation of digital infrastructure. By analysing what edge really means in all its different forms – from content distribution to autonomous vehicles – Vertiv can help its customers, partners and other stakeholders accelerate and focus their edge strategies.

As interest and demand for edge computing takes shape, we will continue to develop and deliver edge-focused products. Our Vertiv SmartCabinet and SmartMod are just a couple of examples of modular "plug and play" style solutions that are intelligent, fully integrated systems ideal for edge infrastructures.

2. Specifically, recent Vertiv research identified four specific types of edge use? Can you tell us a little bit about each, starting with the data intensive applications?

The four archetypes are: Data Intensive, Machine to Machine Latency Sensitive, Life Critical and Human-Latency Sensitive. These were determined by analysing 24 use cases for edge technology, which were identified by the research as having the greatest impact on IT infrastructure.

Data Intensive

The Data Intensive archetype encompasses use cases where the amount of data is impractical to transfer over the network directly to the cloud – or from the cloud to the point-of-use – because of data volume, cost or bandwidth issues. Think, for example, of growing IoT networks and high-definition content delivery. In the case of the latter, in 2016 video accounted for 73 percent of all IP

traffic and is expected to grow to 82 percent by 2021. Big players like Netflix are partnering with colocation providers to expand their delivery networks and bring data-intensive video streaming closer to users, as well as reduce costs and latency. With the growing demand for high-definition videos, local hubs will have to increase support for current metro hubs to reduce bandwidth costs and latency issues.

Organisations are already struggling to manage the volume of data being generated for IoT networks. Such organisations must move huge amounts of data created by devices and systems to a central location for processing.

3. And then there are the human-latency intensive applications?

The Human-Latency Sensitive archetype covers use cases where services are optimised for human consumption, with speed being the defining characteristic. Customer-experience optimisation is a good example of a use case for this archetype. Speed has a direct impact on user experience in applications. Websites which optimise for speed using local infrastructure notice a direct effect on their page views and sales. According to Google, adding a 500-millisecond delay to page response times results in a 20 percent decrease in traffic.

Another use case example within this archetype is smart retail and immersive technologies including augmented reality. Delays in delivering data directly impact a user's technology experience – and a retailer's sales and profitability. As such use cases continue to grow, so will the need for local data processing hubs.

4. And machine-to-machine latency sensitive applications?

Machine-to-Machine Latency Sensitive applications include use cases where services are optimised for machine-to-machine consumption. Speed is a characteristic of this archetype as it is needed for machines to process data to their full capabilities. The consequences for failing to deliver data at the required speeds can be even higher in this case than in the Human-Latency Sensitive Archetype.

An example of an application of this archetype would be latency sensitive systems used in automated financial transactions. Prices can fluctuate in milliseconds and if systems are lacking the latest required data and cannot optimise transactions, potential monetary gains can be converted into losses. According to Tabb Group, brokers can lose as much as $4 million in revenues per millisecond if electronic trading platforms are 5 milliseconds behind the competition.

Additionally, smart grid technologies fall into this archetype as the electrical distribution network needs to self-balance supply and demand and manage electricity use in a sustainable, reliable and economical way. It enables distribution networks to self-heal, optimise for cost and manage intermittent power sources, assuming the right data is available at the right time.

Other Machine-to-Machine Latency Sensitive applications include smart security systems that rely on image recognition and real-time analytics

5. And, finally, the life critical applications?

Life Critical use cases have serious implications for human safety. Autonomous vehicles and drones, for example, can be a great benefit when operating effectively, but can be positively dangerous to human health if they malfunction.

In the near future, self-driving cars will be on the road. If these kinds of systems don't have the

correct data, the consequences could be disastrous. The same is true of drones. We could easily be looking at a future where hundreds of delivery drones are flying over a city at any given time.

Increased use of technology in healthcare is also an example of a use case in the Life Critical archetype. Electronic health records, cyber medicine, personalised medicine (genome mapping) and self-monitoring devices are reshaping healthcare and generating huge volumes of data.

6. And how is Vertiv developing edge data centre solutions to address any/all of these?

We are building on this initial phase of research to define technology requirements for each archetype – and to help accelerate edge deployments and ensure local infrastructure provides the security, speed and availability a particular application requires.

This includes producing fully integrated solutions including power, thermal, security, along with management and software integration – all in a single package. These can be flexibly deployed to match data centre infrastructure economics at the network edge, so that operators can efficiently upgrade and migrate thousands of sites. What has changed most fundamentally is the ability to make the transition fast, simple and flexible for these network operators at hundreds or thousands of edge locations.

As mentioned previously, our Vertiv SmartCabinet and SmartMod are examples of solutions that support edge applications and we are continuing to develop our product portfolio to address the needs of our customers and anticipate the broader market trends.

7. Energy efficiency remains another major data centre issue – in general terms, what are the areas, technologies within this that interest Vertiv right now?

There is a limit that the traditional UK grid infrastructure can cope with and moving forward, we are certainly likely to see more innovative uses of renewables like solar, as well as energy storage being used to complement existing data centre power architectures. Legacy designed data centres are most at risk as they might not be able to face power usage in today's connected, mobile age, due to constraints with incoming power supplies and local infrastructure.

One area of efficiency that is often overlooked is cooling equipment. In an average data centre, cooling accounts for approximately 40 percent of total energy use – a staggering percentage given the level of equipment used in the building. To ensure availability without breaking the bank, a data centre needs a thermal management solution designed to optimise cooling efficiency while lowering the total cost of operations. This requires smart, flexible technologies that can promptly adapt to changing temperatures. Such systems can automatically respond to conditions in the data centre to optimise cooling and improve system reliability while reducing operating costs by up to 50 percent. While newer cooling units often come equipped with the latest technology and features, older units can also be upgraded to achieve similar cost-saving and efficiency benefits.

8. Vertiv recently joined the Ericsson Energy Alliance – can you talk us a little bit about this?

The Ericsson Energy Alliance is a unique complementary site solution partnership to drive cost-effective and sustainable service provider network evolution towards 5G, within the Ericsson Radio Site System.

The alliance combines Vertiv's global expertise in power, thermal and infrastructure site management solutions, and Sweden-based NorthStar's leadership in battery and energy storage solutions, into the Ericsson Radio Site System. The partnership will help establish a competitive

ecosystem and management interface, as well as help to increase the market share of the enclosure and power parts of the portfolio.

The alliance increases competitiveness and cost efficiency through the broadened portfolio and more access to new technologies. The sustainability benefits include:

- Equipment to reduce energy consumption
- The use of renewable and hybrid energy resources
- Efficient operations to reduce carbon footprint
- Remote energy management and hybrid site solutions, including solar cell panels

9. And Vertiv is working with Telefonica to offer Energy Savings as a Service – what does this involve?

Vertiv and Telefónica formed a global, long-term partnership to boost energy savings through fit-for-purpose infrastructure solutions. Under the frame agreement, Vertiv will provide Energy Savings as a Service (ESaaS) across Telefónica's core and access sites in Europe and America, covering all facets from initial site assessment to comprehensive maintenance services.

Through this agreement, Vertiv experts will carry out energy audits and deliver wide-ranging assessment reports outlining the projected KPIs as well as ROI Energy Savings Sharing for each specific site. The reports comprise a series of recommendations for optimizing the performance, capacity, availability and efficiency of critical infrastructure, ultimately increasing energy savings. Vertiv will provide total support from consultancy to execution to 24/7 monitoring and maintenance services, requiring no capital expenditure (CAPEX) from the customer, with Vertiv fully financing the project as part of the ESaaS contract.

10. Moving on to other industry 'hot topics', please can you share Vertiv's thoughts on, for example, how the colocation market seems to be developing?

In terms of the colocation market, we should most certainly be looking to the speed at which cloud adoption is happening. In many cases, cloud providers can't keep up with capacity demands and in reality, some would rather not try. They prefer to focus on service delivery and other priorities over new data centre build and will turn to colocation providers to meet their capacity demands.

With their focus on efficiency and scalability, colos can meet demand quickly while driving costs downward. The proliferation of colocation facilities also allows cloud providers to choose colo partners in locations that match end-user demand, where they can operate as edge facilities. Colos are responding by provisioning portions of their data centres for cloud services or providing entire build-to-suit facilities.

In terms of other 'hot topics,' the emergence of the Gen 4 data centre is also something that is a focus area for us. Whether traditional IT closets or 1,500 square-foot micro-data centres, organisations increasingly are relying on the edge. The Gen 4 data centre holistically and harmoniously integrates edge and core, elevating these new architectures beyond simple distributed networks.

This is happening with innovative architectures delivering near real-time capacity in scalable modules that leverage optimised thermal solutions, high-density power supplies, lithium-ion batteries, and advanced power distribution units. Advanced monitoring and management technologies pull it all together, allowing hundreds or even thousands of distributed IT nodes to operate in concert to reduce latency and up-front costs, increase utilisation rates, remove complexity, and allow organisations to add network-connected IT capacity when and where they

need it.

11. Specifically, how do you see colocation developing as Cloud Computing becomes more and more pervasive?

Vertiv and 451 Research's report, 'The Impact of Cloud and the Internet of Things on

Data Center Demand', highlights how enterprises continue to shift IT from on-premise data centres to off-premise colocation, hosted private cloud and public cloud environments. While companies are on average retaining as much as 40 percent of their workloads in-house, and up to 36 percent of workloads in non-cloud environments, most survey respondents plan to increase their use of private and public cloud over the next two years.

For providers of leased data centre space, the continued move to public clouds will drive demand under a variety of circumstances, including when:

1. Cloud providers lease data centre space rather than build it themselves.

2. Enterprises continue to shift workloads and data that are not suitable for public cloud off-premises (e.g., to a private cloud).

3. Cloud providers and enterprises seek to install points of presence in network-dense data centres to interconnect with providers, partners and customers.

While this survey focused on enterprises rather than cloud providers, separate 451 Research has found that cloud providers outside of the top three (Amazon, Microsoft and Google) have a strong tendency to lease nearly all their datacentre space. Even the top three providers, which have built very large data centre campuses, tend to lease large amounts of data centre space from specialised providers, and this tendency seems to have increased in recent years due to strong cloud take-up by enterprises and the need for cloud providers to add global infrastructure quickly.

12. And we can't talk about data centres without mentioning security! What's the Vertiv take on the recently introduced GDPR?

While compliance is vital for any organisation, IT teams must be conscious not to get forced into treating the GDPR purely as a box-ticking exercise. Firms must be smart in meeting regulatory requirements while using data - and IT holds the key to this success.

The storage, protection and handling of data are today recognised as providing a critical edge in the market. Indeed, the ownership of insights that your competition doesn't have, and the ability to mine and use data better than others, is central to the best businesses. This heightened 'data-based' competition is set within a landscape that IDC predicts will see us create 180 zettabytes of data in 2025. So irrespective of the GDPR, companies will be applying pressure on their teams to capitalise on data.

Business leadership, business units and employees themselves, therefore, need to understand the impact that adhering to changes in regulation will actually have on their business. Knowing which boxes need to be ticked by the organisation is not the answer in isolation; there needs to be a sustainable approach to competitive advantage that supports immediate compliance needs, but critically that can be deployed by IT.

Looked at another way, GDPR shouldn't be the end point – but a reset in the way that businesses use and manage their data to stay competitive.

13. The original DCIM products didn't quite deliver on the promise, do you think that the newer DCIM offerings are better equipped to deal with the challenges of managing the data centre environment?

Yes, most definitely. While DCIM is not a magical tool, it does have the capacity to give real-time insight into power, space and cooling, which helps you manage capacity, reduce risk and increase efficiency. Vertiv has evaluated the promises of DCIM and believes it can deliver real value to data centre in four key areas: thermal management, capacity planning, data centre monitoring and energy management.

When it comes to thermal management, it can be a time-consuming and costly balancing act to manage cooling capacity between IT devices and the facility if the right tools aren't available. Our solutions mean that they can analyse and monitor real-time conditions and set exhaust temperature, among other features.

Capacity planning can also be a hefty cost and it's vital to know exactly what you have so that you can maximise your resources. Our solutions allow for that, using things like accurate and up-to-date information and tracking device details to drill-down capability. Similarly, with data centre monitoring and energy monitoring, we can give real-time visibility into energy consumption and operating efficiencies, use elements such as our alarm notification system and understand the interdependencies among devices to plan for power changes.

14. And how do you see everything from IoT, AI and ML impacting on data centre management over time – will robots be running data centres any time soon?

IT workloads are expected to grow in the next two years, with some of this attributed to the increase in IoT. This will inevitably mean higher volumes of data being generated, likely leading to a shift away from on-premise data storage. One possible outcome is that the data will move to cloud hosting and colocation data centres as well as to network operator infrastructure and edge computing systems. As our work on defining edge archetypes makes clear, many edge applications, while relying on cloud computing for data archiving and deep learning, require distributed IT capacity that is located close to the source of data.