Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING WHITE PAPER

Maximizing the Value of Edge Cloud

A Heavy Reading white paper produced for Juniper Networks Inc.

AUTHOR: ROZ ROSEBORO, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

Communications service providers (CSPs) are transforming their edge infrastructure in large part to support 5G, but many also realize that they can, and should, do so much more with it. Indeed, CSPs are looking at how they can further monetize the edge by making it available to third parties. Today, their edge cloud infrastructure can support latency-sensitive applications such as the Internet of Things (IoT) and augmented reality/virtual reality (AR/VR). Besides the potential for new services, CSPs can also benefit operationally from deploying more distributed resources, including optimized network traffic and improved service performance.

Automation will play a major role in operating edge infrastructure, since manual processes cannot cost-effectively scale to support hundreds or thousands of new locations. CSPs have seen how the hyperscalers have benefitted from extreme automation, and are looking to achieve similar results. To maximize these benefits, though, CSPs must also adopt the mindset of the hyperscalers and trust that the infrastructure will deliver what the application needs.

The "edge cloud" characterizes environments, such as base stations and central offices, where space and power are constrained, and where connectivity may be intermittent. Edge cloud represents an extension of data center resources where the control and data planes have been separated, so traffic is processed closer to where it is generated. This approach both improves service performance and lowers transport costs. CSPs will face new operational challenges with the edge cloud, and will need to embrace a cloud-native approach to address them.

Edge cloud infrastructure will be different from the infrastructure deployed in centralized data centers. Space and power constraints drive the need to optimize on these dimensions without sacrificing functionality. The platforms must deliver high availability and be able to easily adapt to changing resource and connectivity conditions.

To manage their edge cloud operations, CSPs will leverage tools such as OpenStack and Kubernetes. Kubernetes will be particularly important in optimizing workload performance and service chaining. Highly scalable software-defined networking (SDN) platforms will provide the policy and control for the edge cloud, and with virtual routers, extend networking to the edge for connectivity and service chaining.

Multi-tenancy will be critical for CSPs wanting to monetize their edge cloud infrastructure. Network slicing, combined with microsegmentation, will allow CSPs to offer third parties their own environments to deliver services. To enforce security across OpenStack and Kubernetes, the edge cloud will leverage stateful firewall enforcement with centralized policy management. Intent-based, policy-driven networking, along with improved telemetry, will help drive the monitoring and analytics needed to support automation of the edge cloud.

The edge cloud will enable CSPs to achieve numerous benefits. Automation will drive faster service creation as well as more efficient operations. Improved service performance from lower latency will drive demand for new services such as AR/VR. Opening the platform to third parties can provide an additional way to monetize CSPs' edge cloud infrastructure even before they ramp up 5G – which itself will present a host of new opportunities for growth.

CSPs ARE TRANSFORMING THEIR NETWORKS

Drivers for Change

5G leads most discussions about what is driving CSPs to transform their networks. However, CSPs recognize that the edge can – and should – be used to deliver more than 5G. Indeed, to make the business case work, many are looking at what other localized and latency-sensitive services new edge infrastructure can support. Ideally, they would like to monetize their edge infrastructure by allowing enterprises and other customers to leverage it to run new applications. In advance of 5G availability, CSPs' edge networks can support IoT and AR/VR, both of which require local processing and minimal latency. Early movers will need to work through new operational challenges, but can also expect to reap the benefits of establishing a platform upon which others base their businesses.

Heavy Reading research shows that CSPs are moving more infrastructure to the edge for other reasons besides delivering new services. As shown in **Figure 1**, CSPs expect that they will be able to optimize traffic and improve performance, as well as support VNFs by distributing data center resources closer to the edge.

Figure 1: Benefits of the Distributed Edge

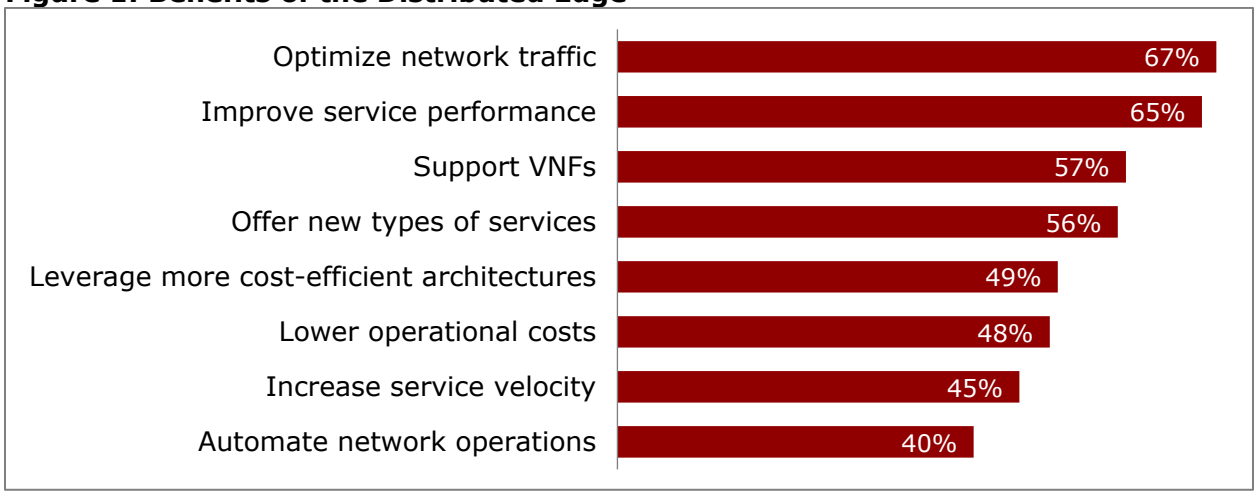

Source: Heavy Reading

Operational Considerations

The move to the edge will lead to new operational challenges. Significantly more locations, including some that may not be easily accessible, will mean changes in how infrastructure is managed on an ongoing basis. Automation will be required, as manual processes will be too costly to scale to the extent needed at the edge. Because the infrastructure could be shared across multiple constituencies, the underlying platforms need to be open, including the application programming interfaces (APIs), to ensure interoperability.

CSPs are looking at the approach taken by the hyperscalers as they architected their clouds, including technologies such as Kubernetes and processes such as DevOps and extreme automation. Importantly, though, it also must include adopting their mindset: Hyperscalers think about deploying applications without considering the infrastructure underlying it, rather than thinking about deploying boxes, software and provisioning connections. They trust that the infrastructure will provide the optimized environment required to deliver the application.

NEW APPROACH NEEDED TO SUPPORT EDGE CLOUD

Defining the Edge Cloud

CSPs will have multiple clouds to deliver different types of workloads. Functionality that once resided in a centralized data center will be more distributed, sometimes into locations that are resource-constrained. The "edge cloud" characterizes such an environment, where space and power are at a premium, and connectivity may be intermittent. The edge cloud could be in a base station, central office, or even at the customer premises. In this scenario, the system decides where the best place is to run an application, meaning ubiquity of resources will be a requirement. Resources in a multi-cloud environment need to be flexible and essentially interchangeable. While some workloads will require tighter latencies, within a given workload type, they should not be tied to any specific set of resources, e.g., through pinning.

The edge cloud is an extension of telco cloud architecture – that is, cloud infrastructure that is distributed outside of centralized data centers into edge locations, including central offices and base stations. In this paradigm, the data plane and control planes are separated, with the control plane remaining centralized and the data plane being distributed to where the traffic is generated. This provides for more efficient management and localized processing, which both improves performance and reduces transport costs. Centralized orchestration allows CSPs to define the policy once and apply it everywhere, and to manage thousands of edge locations in a seamless manner.

The challenges CSPs face is how to deploy new services in an operationally efficient manner, given the breadth, scale and diversity of locations and platforms. The only practical way to do so will be to take a cloud-native approach and embrace automation, virtualization, programmability and security.

An OpenStack working group highlighted many of the attributes of edge computing in a recently published white paper that helps define the environment of the edge cloud, including hardware limitations due to power and space constraints, cost constraints and limited or intermittent network connections. Some of the others include:

- Service agility to enable on-demand service delivery to remote locations
- Extract operational efficiencies across end-to-end distributed cloud
- Ability to perform advanced functions despite hardware, power and space limitations
- Requirements for security in isolating data and resources in a multi-tenant environment
- Consistent operating architecture across diverse, distributed infrastructure
- Zero-touch provisioning environment to avoid truck rolls

Infrastructure Requirements for the Edge Cloud

Hardware Platform Requirements

Because of the different operating environment edge cloud locations will have, the infrastructure deployed there will need to be different from that which resides in centralized data centers. In a recent white paper, AT&T described a "set of principles required to speed-up the edge computing implementations." It expects:

> "[A] Finite set of configurations with low cost, plug-and-play modular infrastructure to reduce complexity in connectivity, network, compute and peripherals, to be more scalable and lowers costs. Software abstraction-based homogeneity that will hide any hardware differences via software."

Many locations will be space- and power-constrained, so the platform needs to be optimized for such conditions without sacrificing the functionalities of a multi-tenant secure cloud. The platforms need to be engineered for five-nines availability using a highly fault-tolerant, robust design that supports load balancing, management and orchestration elasticity and rapid failure detection and recovery. It should be able to adapt to variable resource and connectivity restraints by spinning up virtual machines (VMs) and containers within the distributed cloud.

Management & Networking Requirements

To support the requirement for agile operations, the infrastructure of the edge cloud will need to leverage management and orchestration solutions such as OpenStack and Kubernetes. Kubernetes is seen as the best way to automatically repeat processes at scale. It will be the mechanism to get workloads running on the constrained edge cloud environments, by integrating it into the underlying infrastructure. Kubernetes will also be responsible for the service chaining needed to deliver new services comprised of virtualized network functions (VNFs).

The compute infrastructure of the edge cloud will have a combination of VMs and containers, so CSPs will need a seamless fabric between VM and container hosts using the same SDN cluster. It will also provide added flexibility by supporting optimized nested deployments. In addition to supporting container orchestration, Kubernetes will allow CSPs to run their platforms denser, because containers by nature have a smaller footprint and are more immediately responsive to service provisioning and introduction, because containers can be started and stopped more quickly than VMs. They will need to intelligently orchestrate and automate provisioning of virtual, physical and containerized network functions.

The edge cloud will also leverage highly scalable SDN platforms. It will need end-to-end networking policy and control for any cloud and workload from a single user interface. The SDN platform should translate abstract workflows into specific policies to simplify orchestration of virtual overlay connectivity. It can automate policies and lifecycle management of data center fabrics to allow CSPs to orchestrate both network devices and bare-metal servers. A virtual router can extend the networking to the edge to provide connectivity and service chaining for Kubernetes containers.

Multi-Tenancy Requirements

Because the edge cloud will need to support multiple tenants, network slicing must be supported. To ensure non-overlapping use of resources, including enabling resource-limiting capabilities, the edge cloud will need to support granular network segmentation. In practice, it will be the SDN controller's support for micro-segmentation that will deliver on this requirement.

Using containers and managing with Kubernetes will allow more efficient use of distributed resources, meaning more tenants can be supported on a given set of infrastructure. This will be critical for CSPs to maximize the profitability of the enterprise services delivered from the edge cloud. Open APIs and the ability to support third parties' VNFs will also be required, so CSPs can further monetize their edge cloud infrastructure by allowing others to use it to deliver their own services.

Security Requirements

To ensure security of connections all the way to the edge, the edge cloud will leverage stateful firewall enforcement with centralized policy management to enforce adaptive policies across Kubernetes and OpenStack. Compliance will be more efficient for CSPs when they can have a consistent security policy and associated tools and process for both containers and VMs. For maximum protection, firewalls should apply at Layer 4 and Layer 7. Microsegmentation will be important to ensure isolation of multi-tenant traffic, while encryption must be employed to protect data in transit and at rest.

Automation Requirements

As mentioned above, automation will be critical to managing the edge cloud. Intent-based, policy-driven networking, along with improved telemetry to drive better monitoring and analytics, will need to be a part of the automation functionality. Machine learning-based tools play an important role in monitoring edge network performance and health. Systems that manage workloads need to provide end-to-end visibility from the user to the application, as well as monitor the environment at all times, for every location and domain.

EDGE CLOUD DELIVERS WIDE-RANGE OF BENEFITS

CSPs that deploy an edge cloud should expect to realize a number of benefits. The most critical of these relate to the economic impact and monetization opportunities. With edge cloud, CSPs will be able to accelerate service creation at the edge with faster, simpler, more automated delivery that lowers the overall cost of deployment and enhances operational efficiency. The improved performance will be of particular benefit to latency-sensitive services such as AR/VR and connected car.

This edge cloud infrastructure will support new services and applications for consumers and enterprises, including IoT, allowing CSPs to generate revenue ahead of, and in addition to, 5G. The edge cloud's highly virtualized and automated capabilities conform to the economics of deploying 5G networks. Embedded security extends protection for edge devices and establishes a secure perimeter for the core network. Real-time edge network data can be used to automate resource allocation and utilization and determine conditions for meeting service reliability metrics.

CONCLUSION

The edge cloud represents a significant opportunity for CSPs to improve performance of existing services, monetize their infrastructure by making it available to third parties, and offer new latency-sensitive services, all while setting the stage for 5G. Automation will be key to managing at scale the myriad edge locations and diverse infrastructure. CSPs should adopt a cloud-native mindset, along with technologies such as OpenStack and Kubernetes. The edge cloud platforms must be optimized for space and power without compromising on functionality. Containers are key to maximizing workload density. SDN platforms will provide connectivity and policy needed to intelligently manage the network all the way to the edge. Edge cloud platforms need to provide secure multi-tenancy support to maximize profitability. CSPs should be confident that their edge cloud infrastructure can not only support 5G services in the future, but plenty of revenue-generating services today.

ABOUT JUNIPER

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business. For more information, please visit www.juniper.net.