# The Evolution of the Edge – Why Edge Compute and Networking Should Tightly Integrate into a Cable Edge Cloud

A Technical Paper prepared for SCTE by

**Idris Jafarov**
Delivery Team Leader
DriveNets
New Jersey
302-559-5952
ijafarov@drivenets.com

 

# Table of Contents

**Title** **Page Number**



## List of Figures

**Title** **Page Number**



 

## 1. Introduction

For the past decade, and within the halls of this very event (Cable-Tec Expo), technologists, researchers and experts across the compute and communications industries have been discussing and analyzing the "Edge". Even though this term is not complex in and of itself, it, nevertheless, is fairly fluid. As much as the edge is about compute, it's about communications, and even more about fusion of compute and communications for distribution of intelligence. Hence, the edge must be solved collectively by the communications providers (traditionally the cablecos and telcos) and the compute providers (the cloud), with a laser focus on consumption, treating the underlying tech merely as tools that can be swapped in and out.

Edge adoption will ultimately be driven by business value. Many organizations want their familiar cloud services brought to the edge where the data they want to process is created, on their choice of infrastructure — and with the flexibility of cloud consumption models. They want the cloud experience for the ever evolving workloads that are best suited to be at the edge. This "cloud edge" infrastructure might be in on-premises data centers, more of it will be in new edge data centers, embedded in edge devices, or even built right into the telecom infrastructure. Regardless of where it is, however, users will be able to consume services on cloud edge infrastructure the same way they consume services on traditional cloud infrastructure. For that reason, tremendous infrastructure investments are needed to support the growing device and infrastructure edge demand. Based on estimates, between 2019 and 2028, cumulative capital expenditures of up to $800 billion USD will be spent on new and replacement IT server equipment and edge computing facilities [1]. These expenditures will be relatively evenly split between equipment for the device and infrastructure edges.

Many revenue-generating applications such as gaming, healthcare, IoT, AR/VR require low-latency networks with compute resources close to end users, leading to improved user experience and service quality. This brings multiple challenges regarding where to locate edge data centers, how to scale capacity and services within space- and power-limited environments, and how to provision and operate an increasing number of data centers and resources at scale. The answer is a cloud-like, virtualized, shared infrastructure combining compute and networking resources and supporting multiple container-based applications.

In this paper, we will analyze the evolution of the edge, emerging use cases that will accelerate and drive innovation, and key players in the ecosystem. We will touch upon the infrastructure ecosystem for the edge cloud and deep dive into the Network Cloud that will enable operators to utilize network functions on a shared pool of resources.

## 2. What is the Edge?

There is a biased view that edge and cloud are competing solutions. Actually, they're part of the same continuum of putting colocation, compute, networking and storage in the most effective place, done in the most efficient way. Edge computing is comprised of combinations of systems that span a wide range of locations and conditions and support a diverse set of use cases. Certain use case might demand high-powered Graphics Processing Units (GPU) for AI (Artificial Intelligence), another one might demand low power consumption to lengthen battery life.

There's not just one answer to the question of the definition of "edge cloud". Each definition points out a unique and important concept in the world's computing infrastructure. For example, edge data centers are small data centers that are located close to the edge of a network. They provide the same hardware found in traditional data centers, but are contained in a smaller footprint, closer to end users and devices. Distributed edge of 5G where a decentralized cell network made of edge data centers can help provide low latency for use cases with high device density, is another form of edge cloud. Cable operators, also, aim to leverage the edge and in this paper we will deep dive into multiple approaches in which this can be accomplished.

### 2.1. Evolution of the Edge

Historically, all the data and applications that enterprises needed was stored and processed locally in their on-premises data centers. Due to these resources being available locally, the latency and bandwidth was typically not a concern. These issues became apparent when users connected to the corporate network via a Remote Access VPN. Then, as the cloud computing paradigm took off, some of these applications moved to the cloud compute environment while the rest remained on-premises. This hybrid cloud model certainly has its benefits, but as digital transformation is driving enterprise applications and processes to the cloud, and particularly to public clouds the value of having data centers on site diminishes. Therefore, cloud infrastructure needs to become more distributed and ubiquitous to meet future enterprise needs as demand for cloud-based compute and storage grows. The first generation of distributed edge clouds is being deployed in metro data centers to provide enterprises with a greater choice of compute locations, both within and between countries. Such clouds are serving the needs of enterprises that want to process certain types of data at the edge for regulatory compliance reasons and/or that do not want to incur the cost of backhauling large amounts of data to centralized clouds. The proximity of the cloud infrastructure to the residential customers is, also, of high importance as the ”work from home” shift is here to stay and there is a plethora of emerging uses cases like online gaming/streaming, AR/VR that are picking up the pace.

However, beyond this first wave of ‘edge clouds’, enterprises will require a more distributed, edge-native compute fabric that is available in every conceivable location, not just in a few hundred distributed cloud data centers worldwide, in order to support emerging software applications and software architecture. A new generation of data- and event-driven edge-native services will depend on the existence of a seamless fabric of edge-native clouds that can process their data very close to where it is needed. These edge-native services include but not limited to AI/ML, computer vision, autonomous mobility and applications that support collaboration across multiple contexts and devices, such as multi-player, AR gaming. Such edge-native services are being developed using a very different architectural approach to that used for today’s web applications that run in the public cloud, and they have in common the need to be deployed to specific edge cloud locations to execute with the right level of security, latency and compliance.

The new class of edge-native services will serve an increasingly hyperconnected web of devices that is producing and consuming more and more data. AI/ML applications will turn this wealth of data into intelligence far more quickly than humans could and will replace manual steps that introduce friction into processes and experiences. At the end of this decade, industry visionaries expect an ‘AI of Things’ [2] (Figure-1) to have emerged, whereby everything in our lives will be able to exchange data with everything else, thereby enabling ‘magical’ and unprecedented levels of automation. The net effect of more data, and the application of AI/ML to it, will be the better optimization of business processes and user experience.

 

**Figure 1 - AI of Things**

Edge clouds enable enterprises to run semi-autonomous, latency-sensitive operations locally, to integrate applications into site-based processes more easily and to use server resources more efficiently at individual branches or sites. Cablecos and Telcos have driven early infrastructure edge demand as they virtualize and cloudify their networks. Initially core and transport networks are being transformed with disaggregated distributed networking solutions to enable massive scale and reduction of cost per bit. On top of that, ISPs are uniquely positioned with geographically distributed network infrastructure, which is well suited for infrastructure edge cloud implementations.

## 2.2. Key Players

There are three primary infrastructure service providers that play a significant role in the buildout of the edge cloud: data center operators, last mile operators, and cloud and edge computing service providers.

Virtualized networking services (NFV, VNF) have also had an early start in edge computing. However, they did not deliver on their promise due to multiple concerns which we will touch upon later. A vast majority of the computing service providers have allocated significant capital into the development of advanced software management platforms. This enables edge computing services to be delivered on-demand and remotely just like with cloud computing. Computing service providers are the most common tenants in early generation edge deployments, which is not an accidental event: they are the ones working directly with end user customers and ultimately hosting and managing the workloads and applications that are being placed at the edge.

Also, the Content Delivery Networks (CDNs) that already have substantial data center footprint predate the modern edge cloud. When CDNs were conceived in the 1990's, it was to enable video on demand and save peering and transport costs and thereby host storage closer to users. This was one of the first applications of edge computing. Now, there is a significant demand for plethora of other applications at the edge.

The CSPs/ISPs can themselves become providers of cloud infrastructure because they will be placing edge-native clouds across geographic locations in order to effect end-to-end, cloud-based transformations of their access and transport networks. Operators can then expose such edge-native clouds to support third-party application pipelines in an infrastructure-as-a-service model (IaaS). Operators are actively partnering with the vendors of disaggregated cloud-native networking solutions to provide a more extensive edge Network-as-a-Service (NaaS) solution leveraging open hardware that can support use cases that need a high degree of mobility across geographies and can enable them to spin up network functions within a pool of shared resources [3]. We will dive deep into this partnership in a later section.

## 2.3. Use Cases

There is plenty of demand today for enhanced connectivity and edge computing in multitude of geographical locations. Here are some of the use cases that edge computing will enable and optimize. It is, however, paramount to understand that use cases are different than business cases, as edge cloud ROIs are still a hotly contested item that we are in many cases looking for.

### *2.3.1. Remote Work*

The Work-from-Home (WFH) trend is here to stay. The latest global health concerns have only underlined the importance of minimizing the need for onsite expertise in data centers of any type, from core to edge, and this has hastened the development of tools for remote monitoring, provisioning, repair and management, which could greatly reduce the cost of edge computing. Our solution allows to spin network functions remotely without having anyone on site thereby easing the requirements for on-premises presence.

Also, enterprises have to assure the security of the teleworkers' workstations and secure connectivity to the corporate digital assets. Due to the fact that majority of the resources employees need to access and work on are hosted in the cloud, and depending on the location of the workers, the latency might cause inadequate application response. The necessary workloads have to move closer to the end users for optimal performance.

### *2.3.2. Smart City Applications and Autonomous Vehicles*

The widespread adoption of autonomous vehicles is inevitable and will require a proliferation of edge computing services along roadways to ensure the cars can navigate properly. This will become particularly important in rural areas, where infrastructure will need to be purposefully built out to support autonomous cars in areas that may have no connectivity at all. Similarly, small rural towns and villages will need enhanced edge computing capabilities to implement smart city applications such as turning street lights on or off or monitoring traffic cameras [4].

### *2.3.3. OTT Content Delivery and Gaming*

One of the consumer-driven use cases for edge computing that became a crucial benefit for many during the pandemic is content delivery and online gaming/streaming. Edge computing can help OTT service providers deliver better user experiences for customers and can improve online gaming experiences by reducing latency.

### *2.3.4. Industry 4.0 Use Cases*

As a natural extension of cloud computing, the edge cloud construct is increasingly viewed as a key enabler for the "Fourth Industrial Revolution" in which the widespread deployment of the Internet of Things (IoT), the global sharing economy and the increase of zero marginal cost manufacturing deliver unprecedented communication-driven opportunities with massive economies of scale.

Despite its slow adoption rate, precision agriculture is a great example of this. Farmers using precision agriculture tools and applications will generate thousands if not millions of data points from things like measuring soil moisture and nutrient content to plant growth and stress, to temperatures and precipitation—to even operating autonomous combines [4] (Figure-2).


**Figure 2 - Precision Agriculture**

Other promising use cases are security via drone or camera across large agriculture sites, tracking herds and detecting imminent births or injuries (animal not moving for a long period), optimized feeding for dairy herds and worker safety. All of that data needs to be managed locally at the farm, for example, but it doesn't necessarily need to travel off the farm. Such applications will require connectivity, of course, but because the needs are localized, that connectivity can be served by an array of local network technologies (Figure-3).

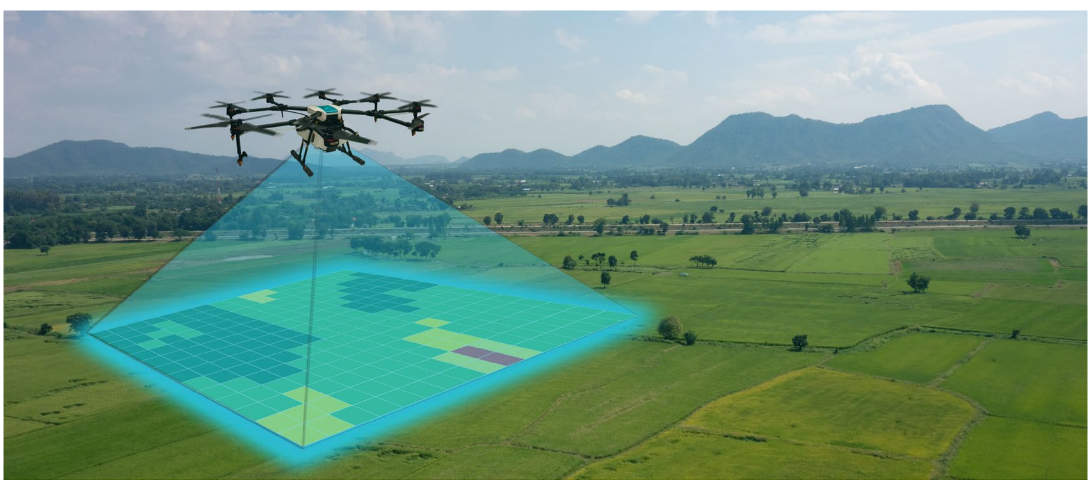

**Figure 3 - Security and Tracking via Drone**

One of the most fascinating use cases is the notion of a digital twin: a near-real-time digital image of a physical object or process that helps optimize business performance. The digital twin can allow companies to have a complete digital footprint of their products from design and development through the end of the product life cycle.

 

Until recently, the digital twin—and the enormous amounts of data it generates and processes—often remained elusive to enterprises due to limitations in digital technology capabilities as well as prohibitive computing, storage, and bandwidth costs, but with the introduction and rapid roll out of the edge cloud environments such impediments, however, have diminished dramatically [4] (Figure-4).

**Figure 4 - Digital Twin Concept as Applied to Industrial Automation**

### *2.3.5. Content Delivery Networks*

As has been mentioned earlier, CDN providers are among the pioneers who brought the edge computing capability into production networks ahead of other industries, primarily to enable fast content delivery.

Typically, major CDN provider manages and operates a large number of content servers geographically distributed across regions, countries or even continents. Besides the origin server, copies of the content are cached on different servers. Cached content offers clients a faster loading experience since their requests can be answered by geographically closer servers.

Such servers are called CDN edge/cache servers relative to the origin servers. Not only is content retrieval latency reduced, but the network traffic is more distributed and balanced without generating traffic spikes and heavy workload on the origin server. Additionally, the origin server’s exposure to cyberattacks and other threats is alleviated.

 

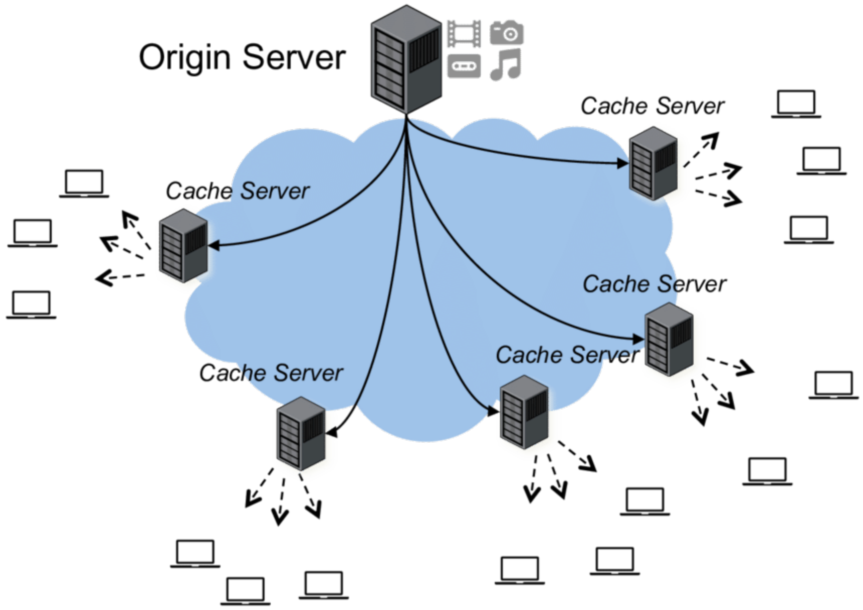


**Figure 5 – Content Delivery Network**

## 3. Infrastructure Ecosystem for the Edge

At the edge, the physical infrastructure is disparate, vendor-agnostic and comes in various shapes and sizes. The layers are built on top of each other and form a harmonious relationship. From the underlying wholesale data center to the cloud infrastructure that it houses; to the multiple sources of connectivity that connects end users and moves data from the core to the edge; to the real estate that is able to support all these complex requirements. At the edge, each critical infrastructure component is crucial in and of itself, but they work together as part of a single integrated ecosystem.

### 3.1. Data Centers

Roughly everything important to data center design comes down to three elements: carbon, real estate footprint and scale. The overwhelming tendency in the industry has been to pack increasing amounts of power into ever smaller spaces without overwhelming the cooling systems, which themselves have become progressively more efficient. The end result of all these optimizations has been to increase the density of the compute, networking and storage equipment a data center is capable of handling.

The “rack” in “rack density” refers to a standard 7-foot equipment rack, the kind found in virtually every data center around the world, while “density” refers to how much equipment can be packed into that rack. Rack density, then, becomes a way of expressing the power density of a data center and this concept is being extended to new form factors.

For instance, a typical enterprise data center might average 6-12 kW per rack, whereas hyperscale data centers can handle densities upwards of 50 kW per rack. A data center with 50 kW rack density can consume four times the power of a typical enterprise data center on a rack-by-rack basis.

At the edge, rack density becomes important because space is scarce and expensive. The higher the density, the more we can do at the edge. Several approaches exist that can eliminate the limitations posed by the constraints at the edge by hosting multiple cloud-native network functions on a shared infrastructure which we will touch upon later.

 

## 3.2. Cable Headends and Wireless Towers

Real estate is gaining more and more importance in locations adjacent to cable headends, wireless towers, and fiber aggregation points, thereby attracting landowners who will likely become players in the ecosystem. Historically, landlords were involved in building wholesale data centers on their land with cloud infrastructure providers. Now, landowners are providing strategic real estate for edge cloud, however, at a multitude of locations due to the vast amount of viable use cases. Landowners provide the real estate, micro data center operators provide the colocation facilities, and then cloud providers, networking software vendors provide the necessary equipment and applications that operate in those facilities.

Edge cloud locations will also pave a way for new forms of interconnection and peering. Traditional hosting data centers turned out to be meeting points for networks, and akin to that the micro data centers at cable headends and wireless towers that will drive edge computing are generally located at the intersection of connectivity paths. These network regions will become very attractive for local interconnection and edge exchange, enabling more optimal routes for data.

## 3.3. Networking and Edge InterConnect

The rise of edge cloud is going to require interconnection to move from its traditional centralized Internet Exchange (IX) model, typically in primary locations within major metros like Frankfurt, London, Hong Kong and etc., to an Edge Exchange (EX) model [1]. End users and devices out at the edge are far away from primary IX points and the distance it takes for traffic to travel to these locations degrades performance and also increases transport costs significantly. To solve this problem, interconnection of networks will need to happen in edge data centers near the last mile network in much closer proximity to the end user. CableCos are using multiple real estate locations (and are perfectly positioned for/to take advantage of) such as headends for aggregation and backbone, and hubs that are hosting vCMTS' and other access equipment which can be used to deploy a next-generation cloud-native solutions to enable a plethora of edge use cases.

We will see edge exchanges emerge to allow peering and data sharing at the edge without necessarily involving the core. Edge cloud is not managed in isolated and independent fashion. It will allow more traffic to remain local, but it will also become an interdependent extension of the traditional IX.

Functions of the applications at the edge will be divided between the edge (real-time processing) and the core (primary functions, analysis, archiving) and only the required data will move back and forth through interconnection services. Having an exchange point at the core that is also tied to the edge will enhance performance and drive cost efficiencies.

## 3.4. Alliances and Partnerships

The diagram below shows the range of network access points and possible topologies, with a range of expected latency and network access:

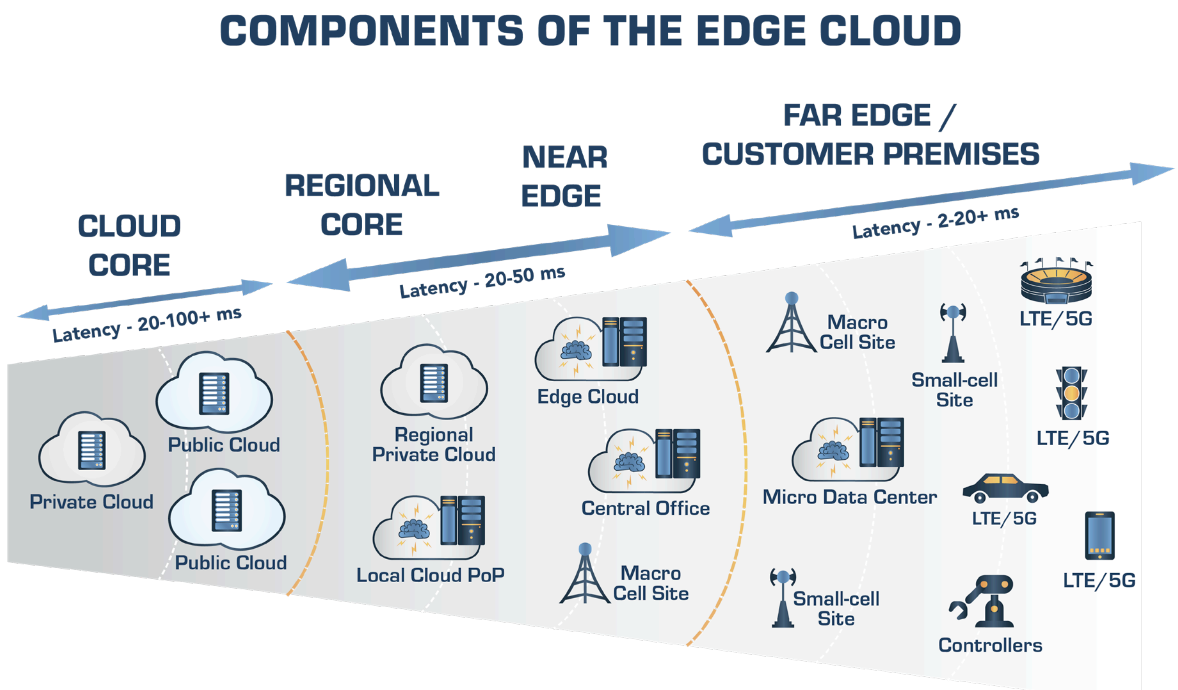


**Figure 6 - Components of the Edge Cloud**

The reality is that latency choices will depend on the applications. In some cases, it makes sense that compute functions are going to happen at the edge, whether that's a local node or even a device. In one example, retail analytics will use edge compute nodes at the location to deliver specific data or applications to the customers. Yet another example is a connect vehicle, which will process a lot of data on board while sharing telemetry and other information with the cloud. It's all about the data and where it's needed. Much of the data that's gathered at the edge has a short lifespan and really isn't economical to bring back into the core cloud.

At the end of the day, it is partnerships that will bring all these disparate pieces together. Wireless tower operators have considered owning data centers, but the most efficient way for them to get to market is to partner. Micro data center vendors are partnering with the service providers that host and manage cloud infrastructure and bringing them on as tenants. Meanwhile, micro data center operators are working with wireless tower and cable companies to access the real estate at headend/hub and tower sites.

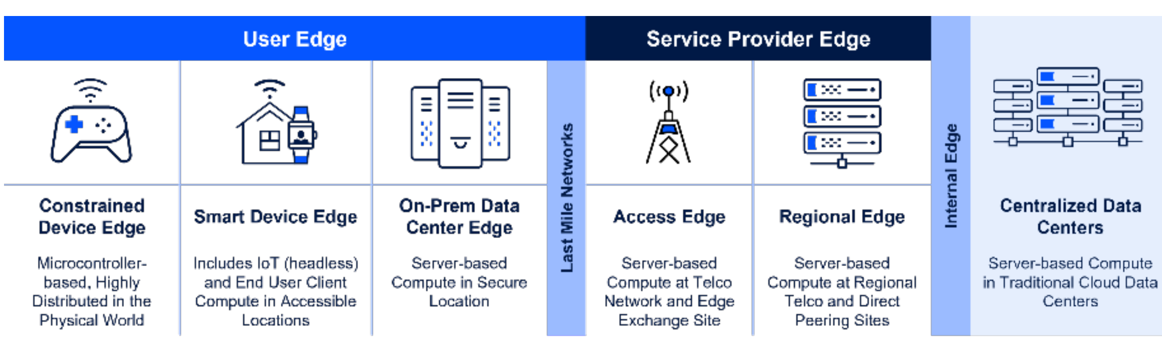


**Figure 7 - Edge Continuum**

 

An exceptional example of this kind of partnership is Open Grid Alliance (OGA), where industry leaders and innovative disruptive startups join forces to rearchitect the Internet. This alliance will work to evolve the Internet to be a global, shared platform that distributes compute, data, and intelligence to when and where it's needed, on demand [5].

# 4. Network Cloud

## 4.1. Evolution of Network Functions

A key consideration for the industry is how to shift Network Functions (NFs) from physical instances to containerized microservices and disaggregated business models. This poses new requirements for the orchestration and control of highly disaggregated and distributed capabilities.

Network Functions were originally deployed on dedicated physical appliances. Then with the widespread adoption of compute and storage virtualization Network Function Virtualization (NFV) initiative started to pick up steam and Virtual Network Functions (VNF) ran on x86 servers (on CPU) which was sub-optimal. It could scale to multiple servers, but the performance was not up to par. As microservices architecture became increasingly popular, Cloud native Network Functions (CNF) were able to run on the public cloud infrastructure and this enabled further growth. However, these NFs still ran on COTS x86 servers.

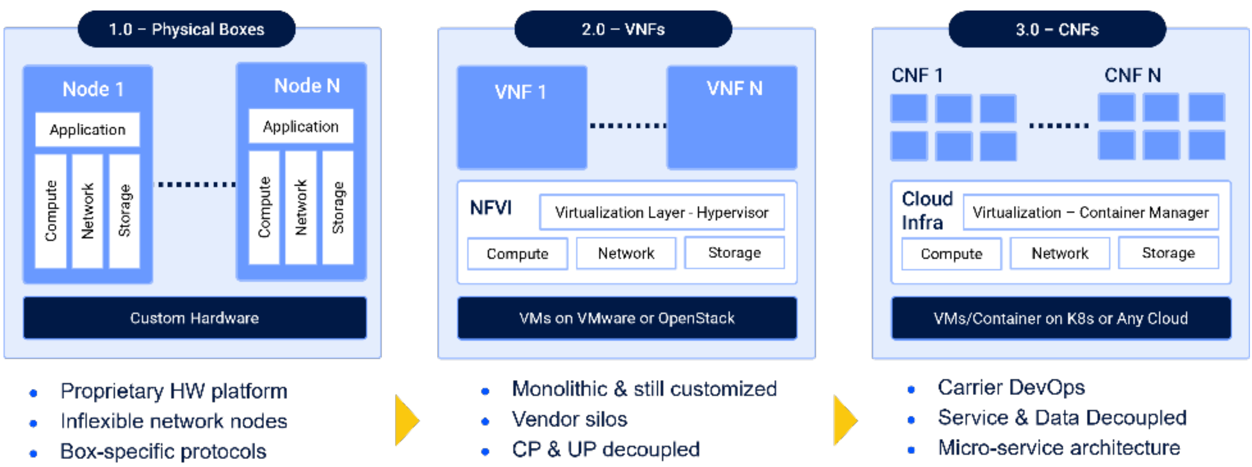


**Figure 8 - Evolving from Bare Metal to VNFs to CNFs**

Today it is commonly proposed to offload the NF to the Network Processing Unit (NPU) on a dedicated Input/Output (I/O) device. This can also be done through a "smart NIC" provider to offload the NF to the Data Processing Unit (DPU) and accelerate the NF as a result.

## 4.2. Edge Cloud Enablers

Hyperscalers have spearheaded the development and implementation of the open hardware and software data center solutions. Organizations like Open Compute Project (OCP) [6] and Telecom Infra Project (TIP) [7] have introduced Disaggregated Distributed Chassis (DDC) [8] and Disaggregated Distributed Backbone Router (DDBR) [9] architecture. It is based on whitebox architecture where each component of the standalone monolithic chassis router is distributed into individual components:

 

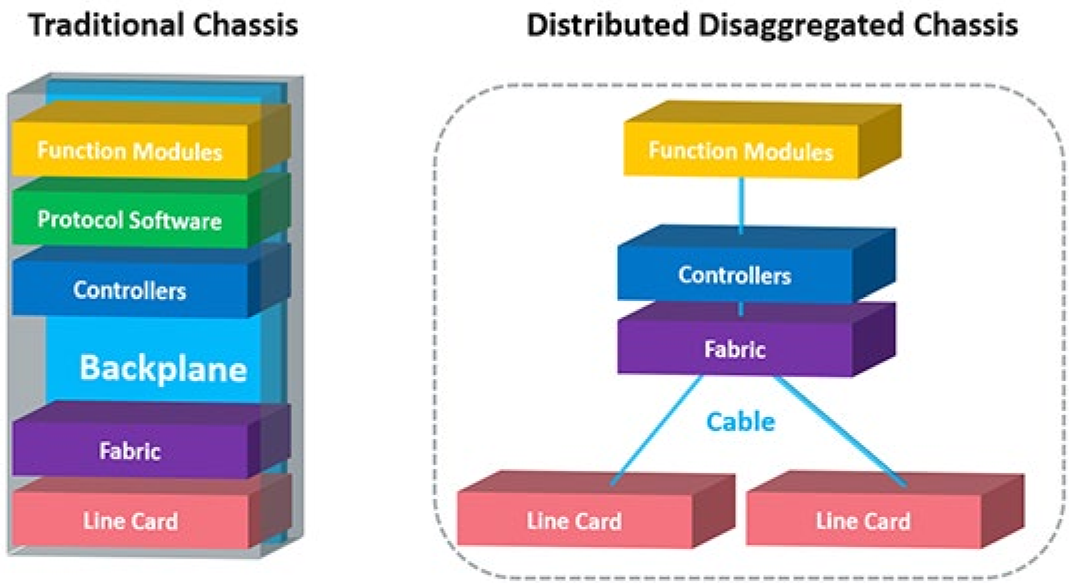


**Figure 9 - DDC/DDBR Architecture**

This allows the operators to scale out this re-architected chassis by simply adding whiteboxes that perform fabric and packet forwarding functions.

Evidently, this is an architectural improvement related to cloud-native networking, and potentially another indication of a future in which carrier networking is highly modular and disaggregated. The true value of such a disruptive model is not necessarily limited to up-front cost-savings. Rather, the value lies in the optionality that comes from building an infinitely scalable and malleable network cloud platform that is pay-as-you-go.

### *4.2.1. Edge Hardware*

Hardware deployed at the edge has historically been purpose-built for specific workloads, frequently CDNs (Content Delivery Networks) or IoT. As edge computing grows in popularity and new use cases emerge, general purpose infrastructure is also being deployed to run cloud-like workloads.

Edge systems must take the path of the hyperscale cloud buildout – low-cost commodity hardware combined with high-powered software automation and distributed orchestration.

The above described DDC/DDBR architecture is perfectly suited to be deployed at the edge. A variety of Original Design Manufacturers (ODM) [10] have introduced open hardware based on merchant silicon certified by OCP and TIP which is already carrying production traffic in the backbone and aggregation networks of major cable and telecom operators.

Due to the disaggregated nature of the architecture where control and forwarding planes run on different physical platforms, there is only a need for a data plane to be installed at the edge whereas control functions run in the core cloud.

 

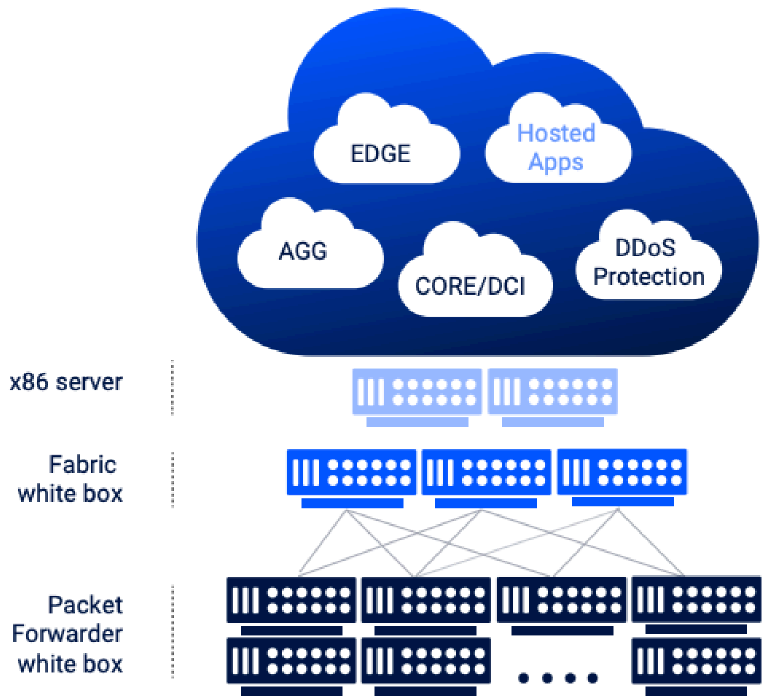


**Figure 10 - Network Cloud**

Packet and fabric forwarders (PF and FF) which comprise data plane, have very high port density (40x100GE, 36x400GE):

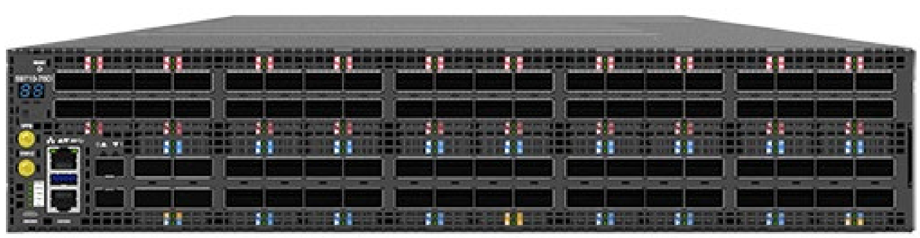

**Figure 11 - Packet Forwarder**

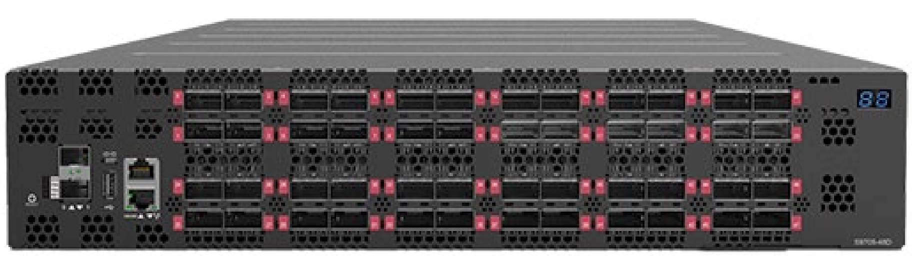

**Figure 12 - Fabric Forwarder**

By strategically placing these disaggregated clusters at the edge facilities (cable headends/hubs, wireless towers) operators can utilize this platform to run multiple CNFs in a cloud-like manner thereby eliminating the need to install separate physical appliances for each network function and utilize any port of the system for any service. The latter point is of a particular importance for the edge environment, because the cost of deployments is higher and rack density, power, cooling, scale, flexibility and even physical access to the site is very limited.

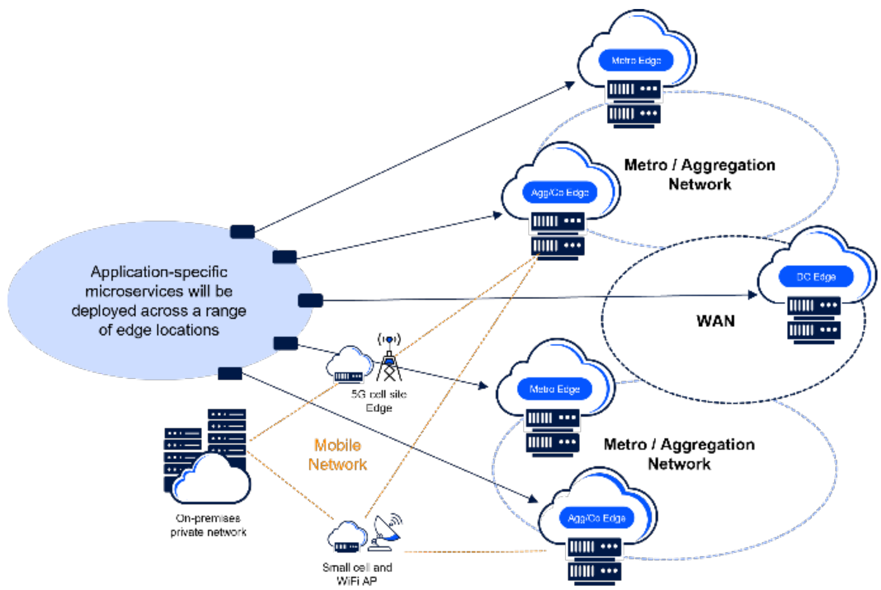


Figure 13 - Edge Locations for Network Cloud

## 4.3. Multi-Services Network Cloud Architecture for the Edge

Initially, it was believed that low-latency needs would drive edge compute, however, we believe the real advantage of network cloud edge is the flexibility, scale and automation provided in deploying new services wherever they are needed, using standardized and affordable software with cloud-based networking and automation software.

Many network edge projects have failed because of efforts to deploy proprietary platforms or complicated architectures and high costs. Following the cloud will be key: leveraging existing cloud-native technologies such as Kubernetes, APIs, and public cloud services and extending them out to the edge.

From a networking perspective, a distributed disaggregated carrier-class networking operating system predicated on cloud-native software principles and standard white-box functionality is required at the edge. This NOS should be predicated on the following cloud-native software design guidelines:

- **Natively Distributed NOS**: marks a departure from the monolithic network hardware designed around a proprietary model of equipment and separates logical routing functionality from the physical infrastructure.
- **Extensive Use of Containers**: designed from the onset to leverage Docker containers to ease development, deployment, and upgrade.

 

- **Optimized Resource Utilization**: anchored on small, highly cohesive, and loosely-coupled microservices.

There are solutions in the market that address majority of the inefficiencies at the edge and offer operators a way to significantly increase their resource utilization while gaining service & architecture flexibility, optimal scaling and software-paced innovation and time to market. Network Cloud for the edge architecture addresses many challenges by combining networking and compute resources over a shared, cloud-like infrastructure. It allows operators to put greater functionality at the network edge, even with space and power limitations.

The foundation for this significant value is the way the Network Cloud is built – in a disaggregated, cloud native architecture. This means the hardware resources of a cluster of multiple white boxes are abstracted by the NOS to a level in which it is consumed as a virtualized resource pool, based on the DDC/DDBR architecture mentioned above. Each networking function, which runs a containerized Service Instance (SI), can be allocated with its required hardware resource (Physical interfaces, NPU, CPU, TCAM, QoS etc.) out of the underlaying shared hardware infrastructure.

The following diagram illustrates the software architecture that allows this:

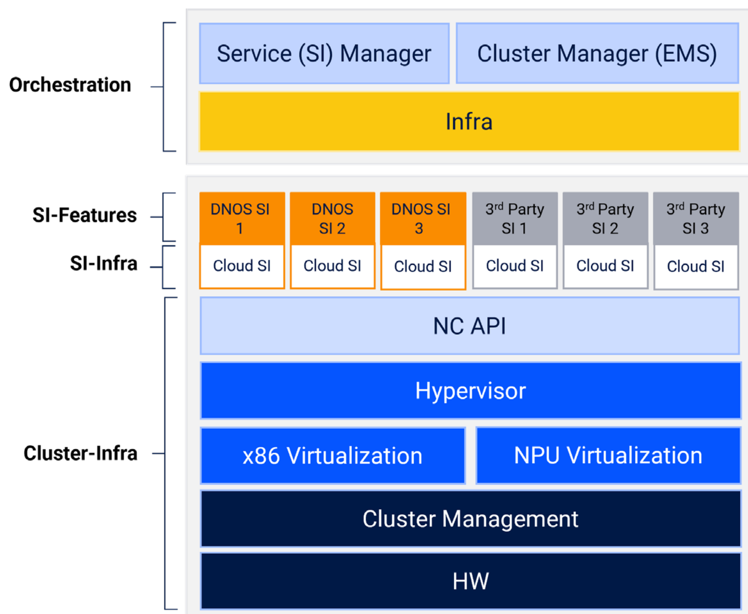


**Figure 14 - Software Architecture of the Network Cloud**

While different hardware resources are abstracted and different cloud SIs run over it, the system represents each network function as a standalone node, keeping the network manageable. This means that a single cluster can integrate multiple networking functions that are physically collocated, yet logically separated.

A multiservice architecture creates a separation of the data plane from the service plane (or control plane). Multiple services can coexist over the same virtualized physical infrastructure and as long as resources can be made available for a service to perform, the network cloud can launch this service into production through a centralized orchestration system.

 

This shared use of resources enables operators to place such scalable cluster at the edge locations while lowering footprint, power consumption, space requirements while reducing the need for engineers get on site for installation and troubleshooting activity. Additional network services can be enabled on this infrastructure on-demand in an automated fashion via an orchestrator that has a full view of the topology and hardware resources of all the components of the cluster.

In order to implement a multiservice architecture a couple of virtual entities were introduced:

- **Service Instance (SI)**: a network function (e.g., PE router, VPN Gateway), which runs independently of any other networking function co-located on the same cluster. While the service instance is logical, it is assigned with dedicated resources out of the shared pool of hardware resources in the network cloud cluster.
- **Inter Service Link (ISL)**: A logical connection between two SIs, which is represented as a (physical) link connecting the two networking nodes. This link allows control plane interconnection between the instances as well as data path features, including QoS and access lists. The ISL is totally logical, hence no cabling work is needed.

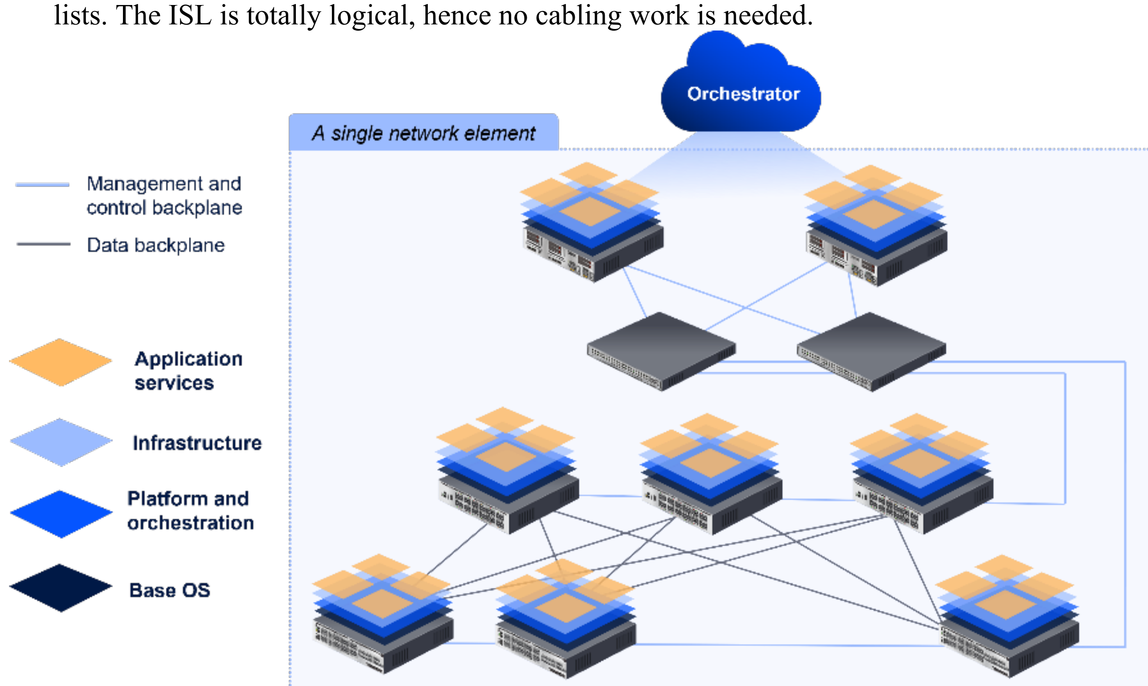


**Figure 15 - Abstraction Layer of the Network Cloud**

The multiservice solution allows for replacing the model of multiple network appliances at the edge, serving different functions and networks (e.g., Network Security, Mobile Backhaul and Broadband), with a single unified infrastructure that can support all the networking services on a software-based network and cloud-native technologies (i.e., microservices, containers).

Let's provide an example based on a TCAM allocation scheme. Three different services are allocated to the same physical ports (Yellow could be a VPN service, Orange – a Mobile Backhaul and blue – a DDoS Mitigation Application). Their TCAM requirements are different, so there might be a case in which a packet arrives to a port, under the yellow service and there are no TCAM resources available. In such an occasion, it will use a TCAM resource from a different whitebox, which has required resources at its disposal. This will be done over the cluster's fabric, as illustrated below.

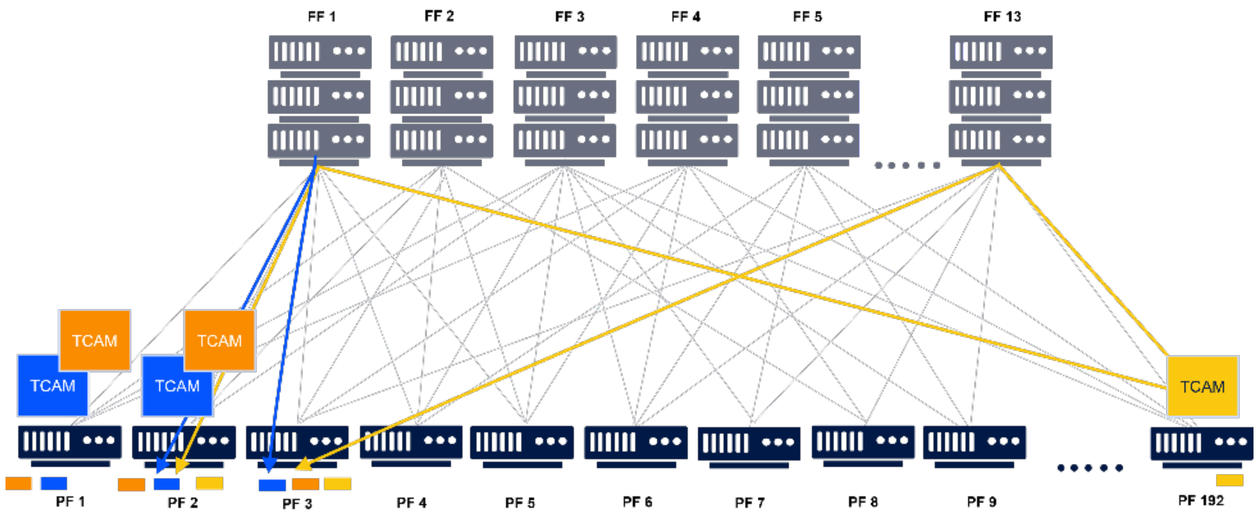


| SI | Maximum rate | Maximum TCAM size |
|---|---|---|
| Yellow | 100% rate | 100% TCAM size |
| Orange | 200% rate | 50% TCAM size |
| Blue | 200% rate | 50% TCAM size |

**Figure 16 - An example of multiple applications running over the Network Cloud**

The same edge cloud platform will also run the network's increasingly autonomous operational and management functions. This is providing operators with the opportunity to create a distributed cloud fabric that can unify multiple network domains and which they can then make available to third-party application pipelines [11].

We believe that operators who have already deployed this architecture in other places of their production environment will have an upper hand as they can utilize the same hardware and leverage new innovative features of the software like the multiservice functionality demonstrated above.

Rest assured – the new network edge is taking shape. But the most successful networking and compute edge deployments will have the characteristics of solving the challenges by simultaneously delivering new applications capabilities and radically lowering the cost of deploying these technologies.

## 5. Conclusion

The global expansion of internet infrastructure is happening at a rapid pace, and the pandemic was at the center of a junction of events that are accelerating the trends already in place. The next decade will see an explosion of activity as a result. There will be more opportunities for creation of a new infrastructure locations at the edge and in strategic and difficult-to-access places. Moving to the edge is fundamentally about getting infrastructure closer to end users, whether that be in new and emerging markets or wherever critical masses of end users cluster. The underlying infrastructure will be provided by a hyperscale platforms, internet service providers and innovative startup companies.

Operators are in a strong position because their networks are, by nature, distributed and ubiquitous across their geographic footprints. Network cloudification necessarily entails building a distributed, edge-native cloud fabric across that footprint too. Cablecos and telcos are engaged in a series of network transformations that will eventually introduce cloud into different network domains (backbone, aggregation, mobile, fixed access and transport). Leading operators have a unified vision for such a network cloud, and anticipate that over the next few years, they will put in place a common, distributed and hybrid cloud-based platform to support all software-only network functions and their management and operational systems. The network cloud must extend to the network edge in order to support access network functions, but it will also need to be instantiated in many other locations, and operators will require a holistic way of orchestrating network function workloads across thousands of cloud locations to provide a 'network'.

# Abbreviations

| | |
|---|---|
| AI/ML | Artificial Intelligence/Machine Learning |
| API | Application Programming Interface |
| AR | Augmented Reality |
| CDN | Content Delivery Network |
| CNF | Cloud native Network Function |
| CPU | Central Processing Unit |
| CSP | Communications Service Provider |
| DDBR | Disaggregated Distributed Backbone Router |
| DDC | Distributed Disaggregated Chassis |
| EX | Edge Exchange |
| FF | Fabric Forwarder |
| GPU | Graphics Processing Unit |
| I/O | Input/Output |
| IoT | Internet of Things |
| ISL | Inter Service Link |
| ISP | Internet Service Provider |
| IX | Internet Exchange |
| NaaS | Network as a Service |
| NF | Network Function |
| NFV | Network Function Virtualization |
| NIC | Network Interface Card |
| NOS | Network Operating System |
| NPU | Network Processing Unit |
| OCP | Open Compute Project |
| ODM | Original Design Manufacturer |
| OGA | Open Grid Alliance |
| OTT | Over The Top |
| PF | Packet Forwarder |
| QoS | Quality of Service |
| SI | Service Instance |
| TCAM | Ternary Content-Addressable Memory |
| TIP | Telecom Infra Project |
| vCMTS | Virtual Cable Modem Termination System |
| VNF | Virtual Network Function |
| VPN | Virtual Private Network |

| | |
|---|---|
| VR | Virtual Reality |
| WFH | Work from Home |

# Bibliography & References

[1] *The State of the Edge report 2021;* The Linux Foundation

[2] https://www.forbes.com/sites/bernardmarr/2019/12/20/what-is-the-artificial-intelligence-of-things-when-ai-meets-iot/?sh=592e3cf8b1fd; Bernard Marr, Forbes.

[3] https://about.att.com/story/2020/open_disaggregated_core_router.html

[4] *The State of the Edge report 2022;* The Linux Foundation

[5] https://www.opengridalliance.org/

[6] https://www.opencompute.org/

[7] https://telecominfraproject.com/

[8] *Hardware Specifications and Use Case Description for J2-DDC Routing System*, Tuan Duong; Open Compute Project

[9] *Distributed Disaggregated Backbone Router (DDBR) Technical Requirements Document*, Eva Rossi, Jose Angel Perez, Kenji Kumaki, Ryuji Matsunaga, Yuji Sonoki, Ahmed Hatem, Diego Marí Moretón; Telecom Infra Project

[10] https://drivenets.com/news-and-events/press-release/drivenets-partners-with-industry-leaders-broadcom-ufispace-edgecore-and-delta/

[11] https://drivenets.com/products/multiservices/

 