

Red Hat & Kaloôm™

A Unified Solution for the Distributed Edge

Introduction

In the realm of 5G, numerous applications are emerging, such as virtual and augmented reality (AR/VR), industrial robotics/controls as part of the Industrial Internet of Things (IIoT), interactive gaming, autonomous driving, and remote medicine, to name a few. These applications require a latency, cost point, service availability, and capability to work at scale that simply cannot be delivered via the use of virtual machines running in a typical public cloud infrastructure. Many of these new applications require an end-to-end latency below 10 milliseconds, however, typical public clouds are unable to fulfill such requirements. A modern, more adapted, cloud-based IT service infrastructure is required to address the real-time, high-bandwidth, latency-critical applications distributed at the edge of the network. These requirements have given rise to edge computing architectures that deliver computing power closer to the end user and the applications they depend upon. As a result, a new burgeoning class of cloud-native applications has emerged enabling Service Providers (SPs) to open their networks and reap the benefits provided by modern ecosystems and their inherent value chain. In this regard, Kaloôm and Red Hat have jointly developed a solution which solidifies Kaloôm's Cloud Edge Fabric™ product¹ for emerging 5G edge data centers. The solution is based on open networking principles and is built with Red Hat® OpenShift® Container Platform, as the foundation of a unified edge solution for network, compute, and storage.

Background

Edge computing is deemed to be an absolute necessity in the era of 5G as a paradigm shift is occurring to move computing and communications from the core network and a centralized cloud architecture and distribute them to the edge. With the densification resulting from billions of connected devices via emerging applications service providers are looking to edge computing to eliminate latency and congestion problems, as well as to improve overall application performance running on those devices. As such, SPs will need to focus on driving efficiencies into their day-to-day operations to improve their profitability. In addition, they will need to devise new 5G use cases which cannot be supported today via existing legacy cloud architectures. The next generation of applications and services will be edge-native and focused on delivering high value out of a smaller infrastructure.

At the heart of this new edge architecture is a unified solution for distributed edge computing. For the longest time, network, compute, and storage capabilities have been deployed as “silos” with each running its own control plane or operating system (OS). As much as automation has eased the deployment, configuration and management of these elements, the need for common services (i.e., zero-touch provisioning (ZTP), upgrades/downgrades, scalability on demand, etc.) has generated different OS’s, including network OS’s. As Kubernetes has become the “de-facto” standard for the orchestration of containers on compute, Kaloom has sought to apply similar methodologies to the network component and add the result to its integration of Red Hat OpenShift Container Platform and Red Hat Enterprise Linux® CoreOS into the Cloud Edge Fabric. By using the same orchestration layer for network, compute, and storage, management of resources at the edge is simplified and optimized.

Containers or VMs?

When considering how to best package emerging applications, containers or VMs, the evidence points to containers as the clear choice, with several key reasons.

For example, in the event of failure for a given VM(s), several virtual network functions (VNFs) demonstrate severe service disruption. Often, it takes several minutes to restart a VM after a crash. Such system behavior is not acceptable in latency sensitive applications.

Another example is seen in Figure 1 (on the following page), which illustrates how system resources are packaged differently in the two environments. Compared to VM packaging, containers package system resources less densely, resulting in at least twice as many applications being able to run on a given server. This advantage maximizes resource usage and brings down operating costs. As well, with agile development and testing, faster time-to-market is achieved with containers versus VMs. Therefore, with the considerations listed in Table 1, one can see how the overall cost effectiveness of deploying containers versus legacy VM based applications becomes truly compelling.

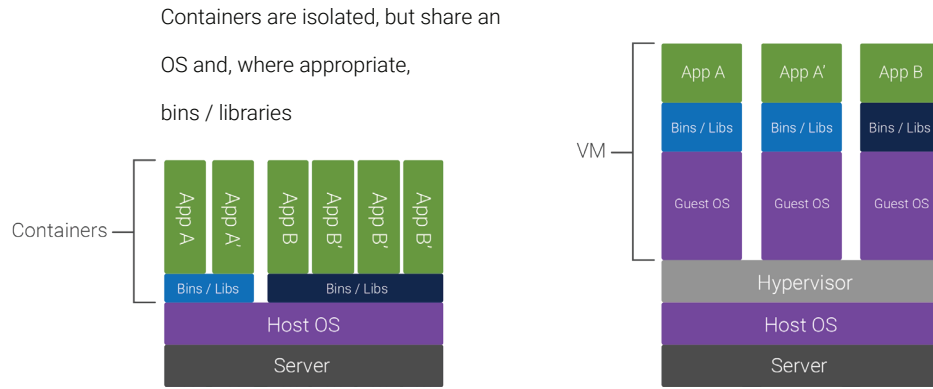


Figure 1: Comparison of Containers and VMs

Attributes	Container Benefits	Virtual Machine Benefits
Consistent Runtime Environment	X	X
Application Sandboxing	X	X
Fast Startup	X	
Low Overhead	X	

Table 1: Benefits Comparison; Containers vs VMs

The Many Edges of a Service Provider's Network

Today, SPs are grappling with the challenges of making their networks more efficient as they aim to deliver higher throughput and avoid bottlenecks. To address this requirement, they look to distribute their data centers at the edge to lower the latency and improve performance of their mission critical applications. However, as seen in figure 2 below, there are several edges.

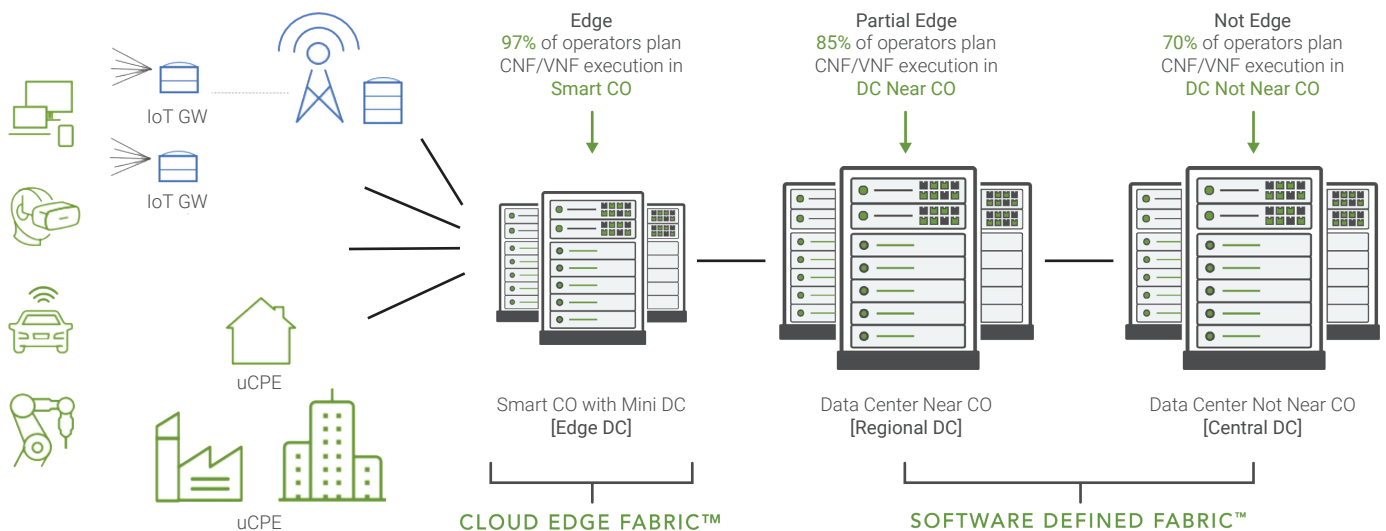


Figure 2: The Various Edges of a Network

For example, Telco edge data centers have several constraints because they must fit within an existing legacy central office and as a result they must contend with very limited space, power, and cooling as part of their infrastructure. These telco edge data centers will need to host a variety of container-based applications, namely real time interactive applications, content delivery networks, basic networking services, mobile and fixed packet core nodes, as well as IoT frameworks. To successfully meet these challenges, an infrastructure comprising Compute, Storage, and Networking is imperative.

Typical Edge Infrastructure Form Factor

The container-based applications previously mentioned can be hosted in a 20 RU edge data center infrastructure composed of Compute, Storage, and Networking nodes. Figure 3 provides a graphical representation of a typical edge data center configuration running a Kubernetes application container platform. The configuration would most likely comprise the following: management switch, fabric and network overlay controllers, fabric switches, Kubernetes masters, servers for 5G user plane function (UPF) to manage 500 gigabits of mobile traffic. This would leave limited space for application servers available as Kubernetes worker nodes. This configuration would limit the capacity for edge applications and thus opportunity for monetization.

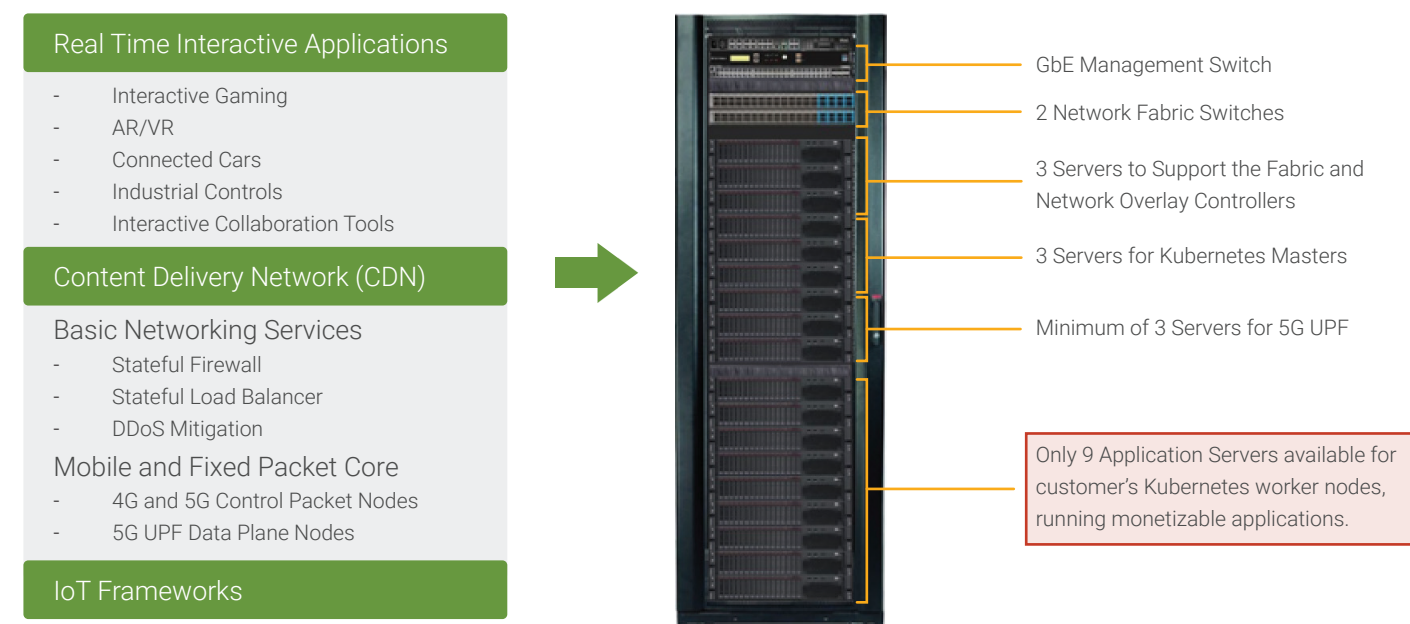


Figure 3: Typical Edge Infrastructure Form Factor

Typical Edge Versus Unified Edge with Embedded Red Hat OpenShift Container Platform

As Kubernetes has become the standard for orchestrating containers on compute, Kaloom has applied the same technology to network components by integrating Red Hat OpenShift into its Cloud Edge Fabric, streamlining, and optimizing resources at the edge.

Figure 4 on the following page compares a typical edge data center having a Kubernetes container orchestration platform to that of a Kaloom unified edge solution based on Red Hat OpenShift.

The components of the new jointly developed solution from Red Hat and Kaloom include:

- Kaloom's Cloud Edge Fabric which supports sophisticated services chaining to improve overall networking efficiency. Hybrid 4G and 5G infrastructures can be supported with cloud native network function (CNF) fabric architecture.
- Kaloom's 5G UPF - a high performance, scalable, low latency cloud native solution.
- Red Hat OpenShift and Red Hat Enterprise Linux, the world's leading enterprise Linux platform, to enable better control and better support for Kubernetes environments.

As well, the unified edge configuration incorporates switches that are Open Compute Project (OCP) accepted white boxes, each equipped with an Intel XEON processor along with a high performance and fully programmable packet processor.

The three Intel XEON CPUs on the three switches run the:

- Red Hat OpenShift supervisor nodes
- Kaloom Fabric Controller software
- Kaloom UPF control plane as a cloud native network function (CNF)

Moreover, in such a configuration Kaloom's 5G UPF data plane application is executed on the Intel Tofino ASIC and FPGA devices within the three switches. Such a solution provides a 5G UPF with much higher throughput, lower latency, and an improved cost/performance ratio when compared to solutions based solely on traditional X86. Within the same space and power constraints, the Kaloom and Red Hat unified edge solution as seen in Figure 4, having Red Hat OpenShift running over switches, application and storage servers, provides 8 additional and potentially monetizable application servers as compared to the typical edge infrastructure previously shown in Figure 3. These additional "freed-up" servers run all the required applications within the given constraints and provide:

- A lower cost by using containers rather than VMs, and
- Lower latency resulting from the deployment of container applications at the edge

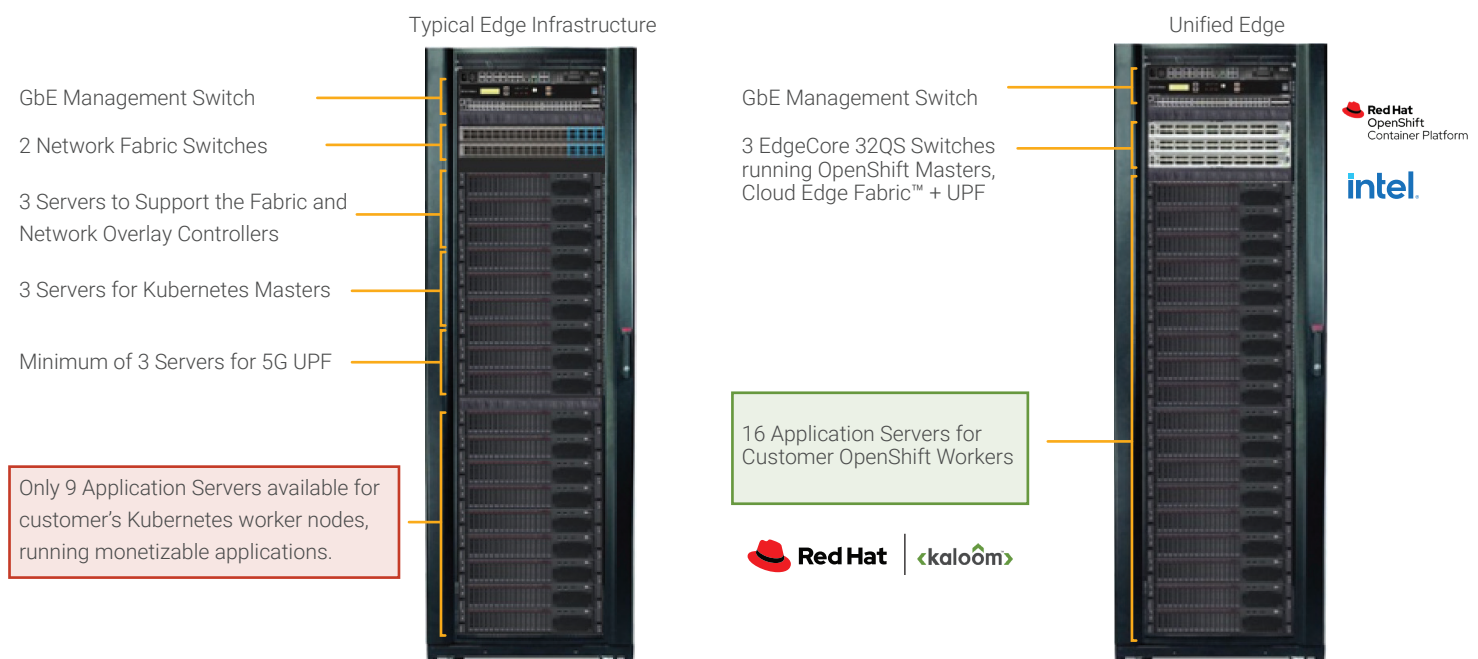


Figure 4: Typical Edge vs Kaloom and Red Hat Unified Edge Data Center Solution

5G Network Slicing

The joint unified edge solution can enable service providers to create multiple edge virtual data centers as depicted in figure 5, each with its own set of physical servers, and Kaloom's Cloud Edge Fabric native support for network slicing can keep them virtually separated.

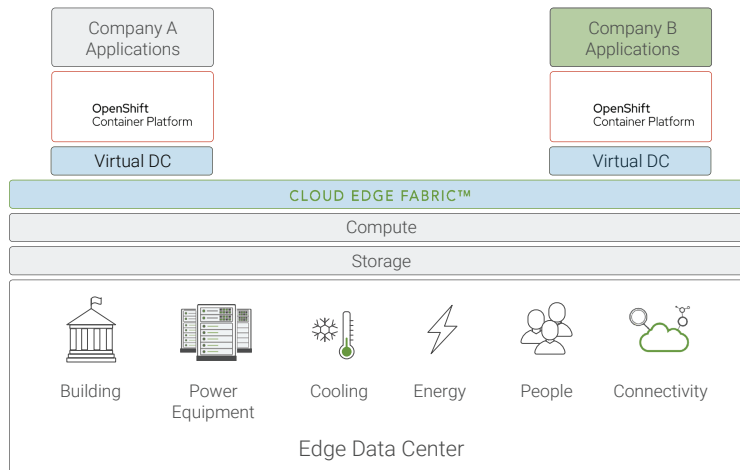


Figure 5: Data Center Virtualization

Figure 6 illustrates the creation of virtual data centers whereby an edge data center can be partitioned into multiple independent virtual data centers. Each virtual data center is provided with its own virtual fabric called a “vFabric”. Each virtual data center and its associated vFabric can be assigned to a different operator. This solution’s use of network slicing allows multiple operators to share a common distributed cloud infrastructure. Each entity enjoys full isolation down to the hardware level for better security and a better quality of experience.

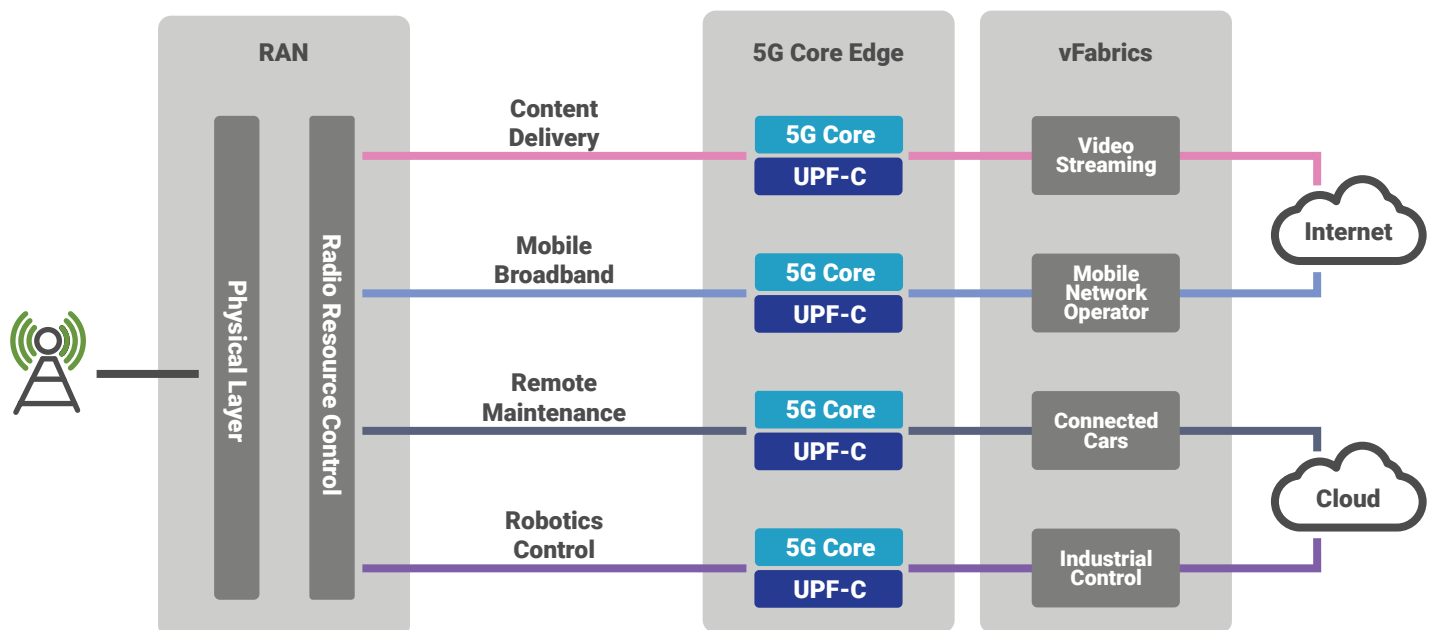


Figure 6: 5G Network Slicing

Conclusion

Service providers are embracing edge computing as a solution that reduces latency and throughput problems, enhances customer experience, and drives more efficient day-to-day operations. Legacy cloud architectures will struggle to meet these varied requirements. Building 5G networks requires a fundamental shift that emphasizes support for edge-native apps and delivery of high performance at the lowest possible cost. This new jointly developed unified solution for distributed edge computing allows network, compute, and storage nodes to share the same underlying container-based execution environment and, in turn, helps address key issues. The integration of Kaloom's Cloud Edge Fabric with Red Hat OpenShift will help service providers simplify complex next-generation networks, accelerate time to market new services while significantly reducing the total cost of ownership for the edge infrastructure.

References

¹ For more detailed information regarding Kaloom's Cloud Edge Fabric product, please refer to:
www.kaloom.com/products/cloud-edge-fabric

For more information please visit: www.kaloom.com

Red Hat, Red Hat Enterprise Linux, the Red Hat logo, and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Copyright 2020 Kaloom, Inc. The information contained herein is subject to change without notice and is correct to the best of Kaloom's knowledge at the time of publication. Kaloom shall not be liable for technical or editorial errors or omissions contained herein. Kaloom, the Kaloom logo, and Cloud Edge Fabric are trademarks of Kaloom Inc. Other product or service names may be trademarks or service marks of others. **Document Version - 1.3 Publication Date: September 27, 2021**



Headquarters

355 Rue Peel, Suite 403
Montreal, Quebec, Canada
H3C 2G9

www.kaloom.com
info@kaloom.com

