

Overview

To sustain emerging 5G applications in a business viable fashion, it is imperative that service providers address the key requirements for the Cloud Edge market. To meet these requirements, Kaloomb has developed the Cloud Edge Fabric™ that provides its customers with the following:

- Multi terabits throughput per logical UPF
- Sub-microsecond latency
- Manageability
- Reduced cost
- Future proofing through a fully programmable environment

Service providers are under pressure to ensure that their networks provide high application performance that efficiently delivers throughput and avoids bottlenecks. Doing this requires distributing data centers at the edge and lowering the latency to process traffic. This improves performance of mission critical applications such as NFV, AR/VR, IoT and financial applications. The Cloud Edge Fabric solution does exactly this - lowers latency and improves performance for the cloud edge market at a drastically reduced cost basis.

Cloud Edge Fabric is the first fully automated data center networking fabric solution specifically designed to provide native support for network slicing along with embedded 5G User Plane Function (UPF) at the data center edge for simultaneous 4G and 5G applications.

The Cloud Edge Fabric and the Software Defined Fabric™ share the following core attributes to work together:

- Programmability
- Automation
- White box networking
- Data plane offloading
- Integration of virtual components (i.e., vRouters, vSwitchs, and vGateways).

Cloud Edge Fabric and Software Defined Fabric Implementation

The Cloud Edge Fabric seamlessly builds on the Software Defined Fabric¹ solution for large scale centralized or regional data centers by expanding the range to include:

- Distributed edge
- Central office
- Virtual central office
- CORD
- Branch office applications

Figure 1 shows the cloud edge DC configuration comprised of the Cloud Edge Fabric that sits in between the base stations and regional DCs. Base Stations are typically connected to a nearby Cloud Edge Data Center via an Optical Backhaul. These Cloud Edge DCs are connected to Regional Data Centers. Regional Data Centers can be connected to National/Central Data Centers. Regional and Central Data Centers can have optimized data connections to public clouds such as AWS, Azure, GCP, etc.

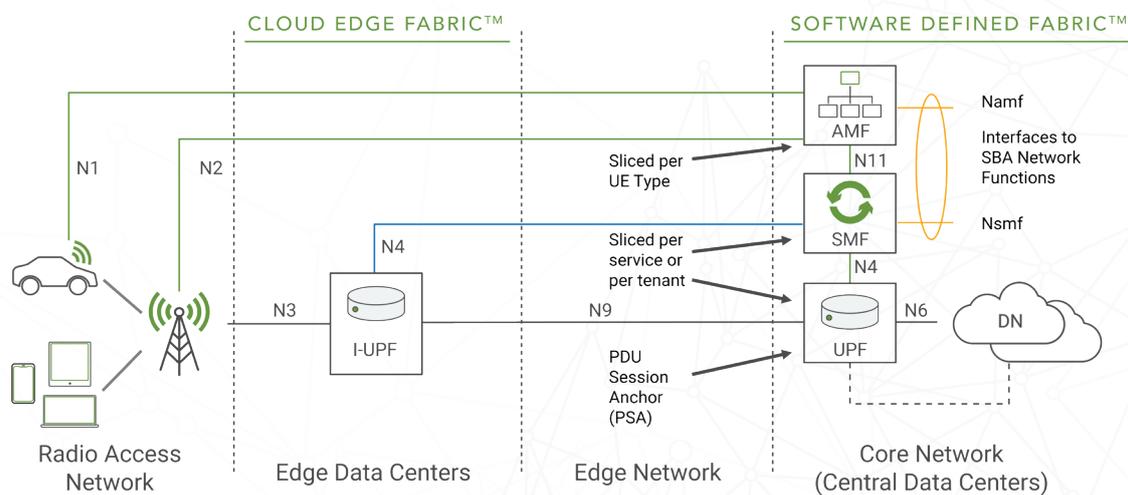


Figure 1: Cloud Edge Fabric and Software Defined Fabric Implementation

Note also Kaloom's UPF implementation within the overall 5G Mobile Packet Core. A 5G EPC comprises several control and data plane nodes such as UPF, SMF, AMF, etc. Kaloom provides the networking environment to support the execution of these control and data plane nodes as either Virtualized Network Functions (VNFs) or Cloud Native Network Functions (CNFs). The Kaloom UPF is based on the latest 3GPP specifications. It is designed to be deployed in a multi-vendor environment and is compatible with SMF, AMF and UPF nodes from other vendors. The Kaloom UPF can either be deployed as a stand-alone node or as an integrated network function inside a Kaloom Cloud Edge Fabric.

Figure 2 provides a conceptual and high-level architectural overview illustrating Kaloom's logical UPF implementation that consists of distributed functionality within a Cloud Edge Fabric. Incoming packets from the RAN or the internet are shown arriving on the edge switches of the fabric upon which a load balancing function distributes the packets to the various Leaf Switches.

The UPF control plane functions implemented at the N4 reference point toward the SMF are fully distributed over the dedicated fabric controllers. The capacity of the UPF control plane can scale according to the number of physical controllers. The UPF data plane is distributed over the leaf switches. The more leaf switches that are deployed in an edge fabric, the higher the resulting capacity of the UPF.

¹For more information on the Software Defined Fabric™, please refer to: <https://www.kaloom.com/product-collateral>

CLOUD EDGE FABRIC™

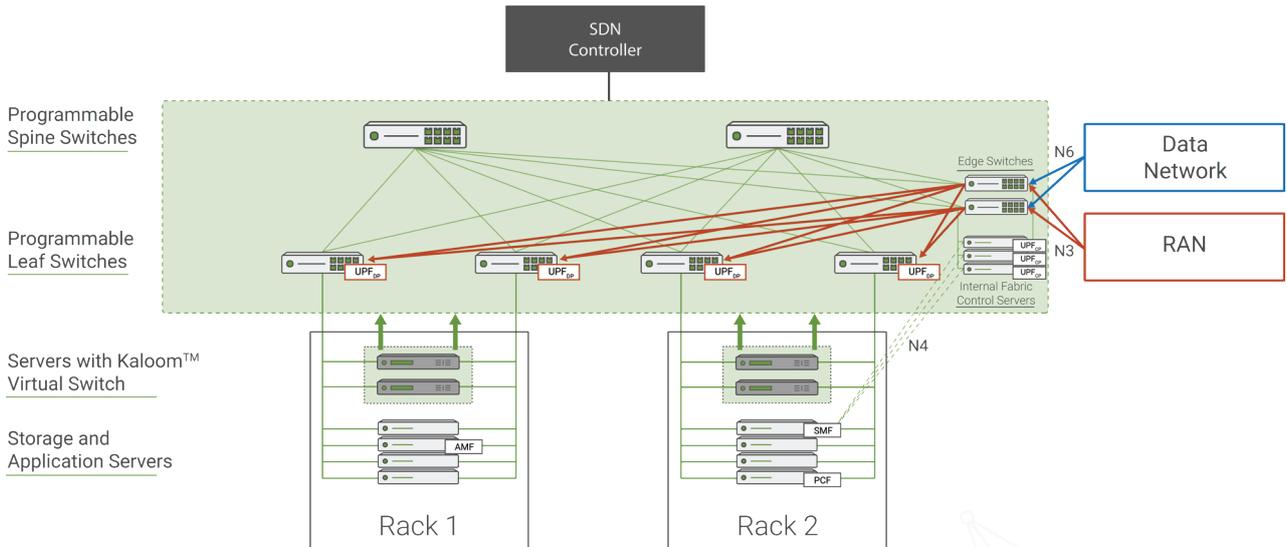


Figure 2: Edge Fabric Implementation

Network Slicing

Cloud Edge Fabric natively supports 5G network slicing whereby an edge data center can be partitioned by a Data Center Infrastructure Provider (DCIP) into multiple independent virtual data centers, with each virtual data center being provided its own virtual fabric called a “vFabric”. Each virtual data center with an associated vFabric can be assigned to a different virtual Data Center Operator (vDCO) or Cloud Service Provider (CSP) that can offer differing SLAs per cloud service user. In this regard, slicing permits multiple operators and large enterprises to share, as Cloud Service Users (CSUs), a common distributed cloud infrastructure, with each CSU enjoying full isolation down to the hardware level for better security and a better quality of experience provided to the individual Cloud Service Consumer (CSC) consuming the services/applications offered by the CSU.

Figure 3 illustrates the concept of network slicing. In Kaloom’s implementation, a slice corresponds to one vfabric which corresponds to a single logical UPF. Each vfabric provides a fully distributed and isolated domain, thereby making it possible to deploy isolated packet cores for either specific applications or large-scale and strategic customers.

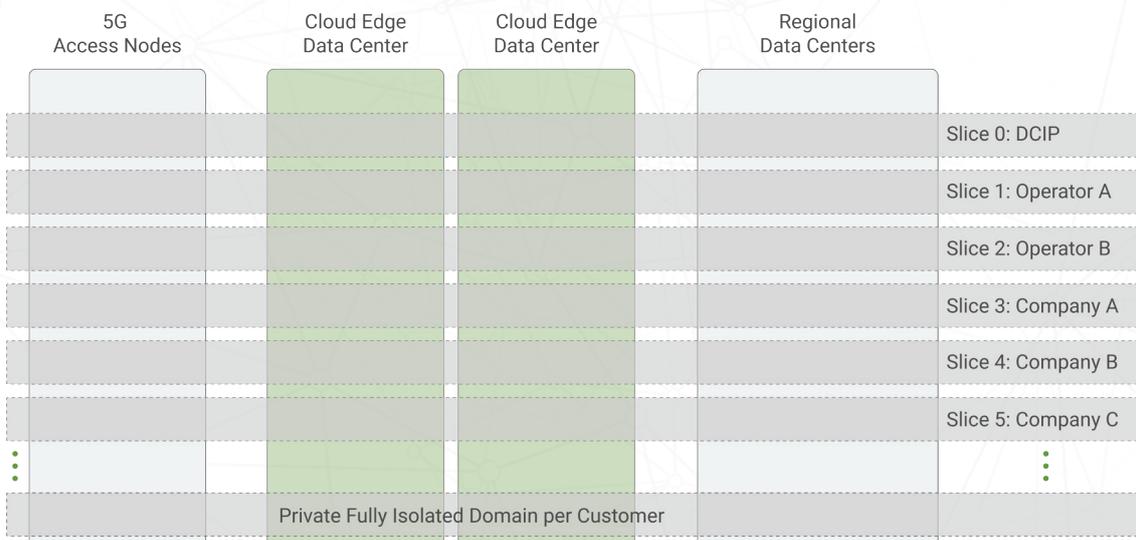


Figure 3: Network Slicing

Cloud Edge Fabric Key Features and Benefits

Software based datacenter fabric that manage spine switches, leaf switches and controllers

- Autonomous Leaf/Spine switch management using a combination of centralized and distributed control functions
- Fabric Controller for
 - Switch topology management of Spine, Leaf and controllers
 - Virtual Fabric creation, resource allocation and modification
 - External Management Interfaces

Programmable software-based data center networking fabric

- Enables the addition of new features and services in runtime without impacting traffic
- Allows developers to develop new code and drive innovation
- Avoids vendor lock-in & eliminates the need to wait for silicon upgrades
- Allows for customer programmability
- Industry standard P4 programming language

Fully automated fabric

- Advanced self-forming & self-discovery fabric
 - Automated discovery of incorrect network topology & cabling mistakes
- Zero-touch provisioning of the virtual networking and virtual components
 - Zero-touch fabric with minimal to no human intervention
- Provisions in minutes vs. hours/days with traditional solutions
- Automated software upgrades
- True network automation with autonomous spine and leaf switches
- Self-healing (automated remediation)

Scalable for distributed edge, central office, virtual central office, CORD, and branch office applications

- Native IPv6, multi POD, cloud and multi data center solutions
- Optimized solution for edge data centers running 4G and 5G applications.
- Embedded with 5G User Plane Function (UPF)

Integrated virtual network components

- vRouter, vSwitch and vGW
- Reduced complexity and reduced TCO

Optimized for virtual environments

- Provides support for both VM and container-based workloads

Advanced monitoring and segment analytics capabilities built-in

- Enables programmable in/out-of-band telemetry per flow/packet

Multivendor solution by design and no vendor lock-in

- Support for networking white boxes from multiple vendors such as Accton and Delta
- Support for Point of Delivery (POD) modules from different vendors
- Easy integration into existing solutions

Fabric Virtualization

- Provides full support for network virtualization with native 5G network slicing (see figure 4)
- Enables allocation of physical resources into multiple autonomous isolated network slices, called vfabrics
- Enables the assignment of vfabrics to different virtual DC operators
- Enables DC operators to faster provision new customers in software
- Provides full isolation between customers/tenants

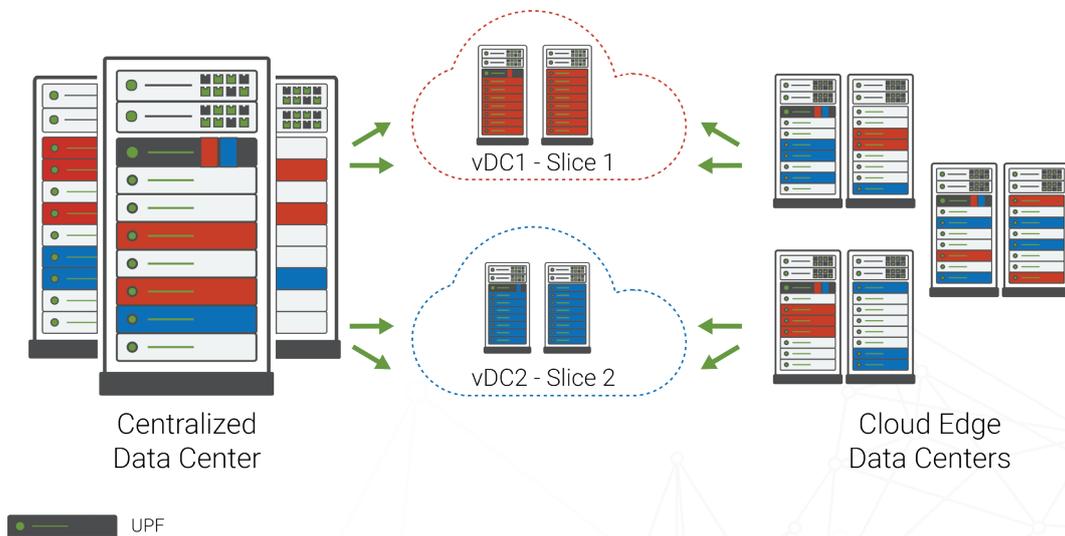


Figure 4: Creation of vFabrics

Open Networking and APIs

- Standard Linux based
 - No kernel patches
- Open APIs
 - Netconf API and Yang modeling
- SDN controller agnostic
 - Plugins for OpenDaylight
- Orchestration agnostic
 - Plugins for OpenStack, K8s, and OpenShift
- No vendor lock-in
 - Networking white box friendly
- Open-source friendly
 - Contributing improvements upstream to Linux and K8s

Significant x86 resource utilization improvements

- Delivers power efficiencies and savings
- Reduces number of CPU cores used for networking and frees them up for payload tasks
- Delivers lower latency and higher performance by offloading and improving NIC functionality into the fabric with the Kaloom Virtual Switch

Lower latency

- Provides significantly lower server and NIC latency when using the Kaloom Virtual Switch
- Improves virtual end to end latency with advanced service chaining capabilities

Standard Linux (Red Hat CoreOS)

- No hacked and/or out-of-date Linux kernels which guarantees faster updates and security fixes
- Allows standardizing on the same OS for compute, storage and networking
- Leverages security-enhanced Linux (SELinux), control groups (cgroups), and kernel namespaces to provide military-grade security
- Lightweight OS that provides the flexible & modular capabilities of Linux containers

Best of Breed Ecosystem

Open Application Environment



Open Fabric



Mainstream Linux OS



Merchant Silicon



Figure 5: Best of Breed Ecosystem

For more information please visit: www.kaloôm.com



Headquarters

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

www.kaloôm.com
info@kaloôm.com