WIND™

STARLINGX

Accelerating the Deployment of Critical Infrastructure Edge Services Through the OpenStack StarlingX Project

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

By contributing proven technology from the Wind River® Titanium Cloud™ critical infrastructure platform to the recently launched StarlingX project hosted by the OpenStack Foundation, Wind River is helping the community to accelerate the release of an OpenStack platform that delivers the uptime, performance, security, and operational simplicity that are critical for edge cloud solutions.

Technology already proven in Titanium Cloud will be key to enabling new business opportunities and innovative applications across multiple critical infrastructure market segments. Today, typical edge cloud use cases include multi-access edge computing (MEC), universal customer premises equipment (uCPE), virtual customer premises equipment (vCPE), virtual radio access network (vRAN) and the Industrial Internet of Things (IIoT). As the industry moves beyond these traditional functions, StarlingX will leverage Wind River technology to power new business opportunities such as connected vehicles, augmented reality, telemedicine, and advanced drones.

The StarlingX OpenStack project, launched by the OpenStack Foundation, Intel®, and Wind River on May 21, 2018, will comprise value-added services that enable highly reliable applications and services to be deployed at the network edge. Technology contributed from Titanium Cloud led the initial seed code release. As a project outside the OpenStack core services, StarlingX avoids any need to fork the standard OpenStack code in order to deliver edge-optimized applications. Edge deployments based on the components in StarlingX automatically and seamlessly leverage updates to the mainline OpenStack code base.

Code from the StarlingX project will also be contributed to the open source Akraino Edge Stack project, launched by the Linux Foundation in February 2018.

TABLE OF CONTENTS

NEW BUSINESS OPPORTUNITIES IN EDGE SERVICES

In the telecom market, communications service providers (CSPs) worldwide are increasingly viewing applications hosted at the network edge as compelling business opportunities. Edge applications present opportunities to sell new kinds of services to new kinds of customers, while revenue from traditional broadband and voice services is essentially flat. This trend provides the potential for increased market penetration, as well as improved average revenue per user (ARPU) for those CSPs who can be early to market with attractive offerings.

Some examples of telecom-oriented edge-hosted applications and functions that are generating wide interest are MEC, uCPE, vCPE, vRAN.

By bringing content and applications to data centers in the radio access network (RAN), MEC allows service providers to introduce new types of services that are unachievable with cloud-hosted architectures because of latency or bandwidth constraints. Specific new business opportunities enabled by MEC include applications such as:

- **Small-cell services for stadiums and other high-density locations:** By deploying applications hosted at the network edge (i.e., in the stadium itself), stadium owners and service providers can offer a wealth of integrated services that include real-time delivery of personalized content to fans' devices. In this case MEC also minimizes backhaul loading because the new traffic is both generated and delivered locally.
- **Augmented reality, virtual reality, and tactile Internet applications:** These applications are just not viable without superfast response times, local image analytics, and deterministic, low-latency communications. Use cases such as remote medical diagnostics and tele-surgery will demand millisecond response times, far quicker than those achievable by round-trip communications with a remote cloud data center.
- **Vehicle-to-everything (V2X) communication:** Whether the use case is vehicle-to-vehicle or vehicle-to-infrastructure, it requires high bandwidth, low latency, guaranteed availability, and robust security—performance requirements that are impossible to meet with a centralized, cloud-hosted compute model.
- **Mobile HD video and premium TV with end-to-end quality of experience (QoE):** Content providers have learned the hard way that video quality is critical for subscriber retention: Studies show that a one-second rebuffering event during a 10-minute premium service clip causes a 43% drop in user engagement. MEC allows service providers to optimize video content, enabling a superior user experience as well as smarter utilization of network resources, and ensuring a fast start and smooth delivery.

Applicable to both enterprise and residential scenarios, uCPE and vCPE deployments replace physical hardware appliances, traditionally located at the customer premises, that provide connectivity, security, and other functions. By deploying general-purpose compute platforms based on industry-standard servers, either locally at the customer's premises (uCPE) or in a data center (vCPE), service providers can remotely instantiate, configure, and manage functions that were once deployed on dedicated hardware platforms. This virtualization of CPE functions reduces operational expenses (OPEX) through more efficient utilization of compute resources, through increased agility in the deployment of services, and through the elimination of "truck rolls" required to update and maintain equipment at remote locations. CSPs can also grow topline revenue and increase margins by allowing customers to self-provision their services, for example by configuring higher-bandwidth connectivity, enhancing security features, or adding options such as high-end video with cloud-based digital video recording.

Other edge-hosted applications and functions bring operational cost reductions as their primary business benefit. RAN virtualization is a good example.

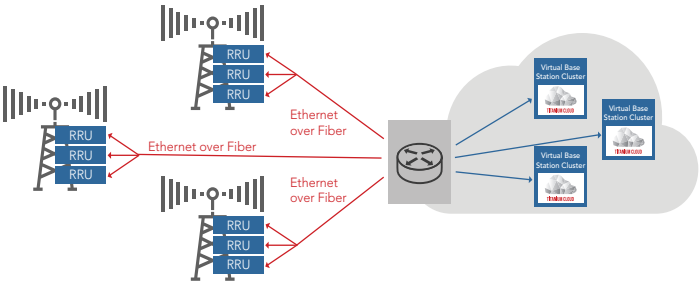

Figure 1. Virtual RAN architecture

In a vRAN architecture, the baseband units (BBUs) are virtualized rather than being located at the cell site as in physical equipment. The virtual BBUs are deployed on software platforms running on industry-standard servers and consolidated in centralized data

centers, while the remote radio units (RRUs) remain at the cell sites at the edge of the network. By leveraging standard server hardware that cost-effectively scales processor, memory, and I/O resources based on dynamic changes in demand, vRAN infuses the RAN with the capacity for application intelligence, which significantly improves service quality and reliability. With vRAN, service providers can achieve a combination of cost savings, dynamic capacity scaling, better QoE, and rapid instantiation of new services.

In the industrial market, IIoT applications represent new business opportunities for several categories of companies. Use cases such as smart cities, smart buildings, connected cars, robotics, and process control all require the aggregation of large numbers of data streams in an IIoT gateway prior to analytics performed either on-premise or in the cloud.

IIoT offers the potential to sell new kinds of services (asset monitoring, analytics, business processes, etc.) to new kinds of customers (manufacturing facilities, car dealers, city governments, hospitals, etc.). Recognizing these new business opportunities, many CSPs have established vertically oriented service delivery teams focused on these opportunities.

Many of these IIoT services are required to be hosted at the network edge, either to enable ultra-low latency connectivity (process control), or to perform on-premise analytics (patient monitoring), or to minimize backhaul traffic (video surveillance). Edge compute solutions are therefore a key requirement as companies exploit business opportunities in IIoT.

Traditional control applications also leverage edge compute solutions as critical infrastructure companies look to slash their operational costs by deploying secure, robust, flexible software-based solutions as alternatives to legacy, fixed-function hardware. Control systems installed since the 1980s present major business challenges such as increasing OPEX due to high maintenance and replacement costs plus a dwindling pool of skilled technicians, limited flexibility resulting from sole-sourced solutions with proprietary programming models, and outdated box-level security features with no provision for end-to-end threat protection or dynamic updates.

These challenges can be addressed by OpenStack-based virtualization solutions. As in the case of new IIoT applications, many of these traditional control services require edge solutions in order to guarantee ultra-low latency response times or to perform on-premise compute functions.

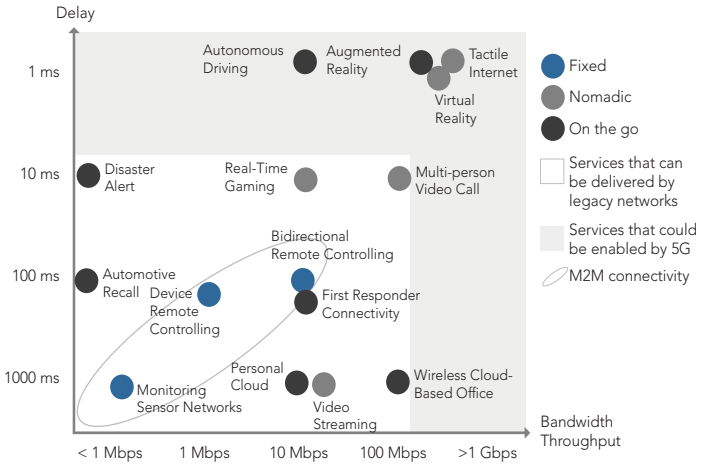

Source: GSMA Intelligence, 2015

Figure 2. Bandwidth and latecy requirements for generic applications

CHALLENGES FOR EFFECTIVE DEPLOYMENT OF EDGE SERVICES

Companies planning the deployment of edge services must select solutions that address key technical and business challenges:

- Many edge applications require ultra-low latency communications with the devices that they serve, recognizing that the "device" could be a smartphone, a tablet, a vehicle, an industrial controller, a set of virtual reality glasses, a TV, or any of a wide range of end points that are emerging as new services are planned and deployed.
- The type of edge applications mentioned above typically need to be deployed on low-cost, low-power servers that are priced appropriately for small branch offices and other remote locations.
- To avoid costly truck rolls and service calls, the infrastructure software platforms need to support automatic installation, provisioning, and maintenance, while communicating with a centralized orchestrator.
- Edge computing applications are often installed by the end users themselves, in unattended, open environments away from cloud data centers, central offices, or points of presence (PoPs). These factors present unique security risks that are greater than in the case of services hosted in the tightly controlled environment of the core network.
- Some of these use cases are targeted at consumers who may have been conditioned to expect less-than-perfect service availability. Many, however, represent either enterprise or industrial

business opportunities where guaranteed uptime is a hard requirement and there are significant financial impacts associated with any uplanned downtime, even during maintenance or update operations. Service continuity and service-level agreements (SLAs) are critical factors as service providers develop their strategies for deploying edge applications.

- To avoid any risk of vendor lock-in and to maximize their flexibility in vendor selection, critical infrastructure companies have developed a strong preference for edge infrastructure solutions that are based on open source software and that have been proven to be fully compatible with open industry standards.

Recognizing that no single OpenStack project existed to address all these challenges, on May 21, 2018, the OpenStack Foundation, Intel, and Wind River launched the StarlingX project, which will comprise value-added services that enable highly reliable applications and services to be deployed at the network edge.

StarlingX is part of the "Edge Computing Infrastructure" strategic focus area within OpenStack. In December 2017, the OpenStack Foundation began incubating strategic focus areas that represent communities focusing on a common set of problems. Currently four strategic focus areas are in incubation:

- Datacenter cloud infrastructure (existing OpenStack)
- CI/CD systems
- Container and serverless infrastructure
- Edge computing infrastructure, including StarlingX

As a project outside the OpenStack core services, StarlingX avoids any need to fork the standard OpenStack code in order to deliver edge-optimized applications. Edge deployments based on the components in StarlingX automatically and seamlessly leverage updates to the mainline OpenStack code base.

Code from the StarlingX project will also be contributed to the open source Akraino Edge Stack project, launched by the Linux Foundation in February 2018.

By extending existing work to upstream proven technology from the Titanium Cloud critical infrastructure platform to the OpenStack community, Wind River is helping to accelerate the release of an OpenStack-based platform that delivers the uptime, performance, security, and operational simplicity that are critical for edge cloud opportunities.

THE WIND RIVER PLATFORM FOR EDGE APPLICATIONS

Edge applications rely on software infrastructure platforms to ensure maximum agility in their orchestration and management as well as optimized, dynamic resource allocation. The challenges described above place severe constraints on the software platform that hosts the virtualized applications. Despite running on low-cost hardware with limited CPU, memory, and storage resources, the platform needs to support "hands off" deployment, remote management, telco-grade service reliability, ultra-low latency, and robust, end-to-end security.

Enabling critical infrastructure companies to efficiently address a wide range of edge-hosted applications, Wind River provides a portfolio of commercially suppported, fully integrated products based on the Titanum Cloud architecture.

Leveraging standard open source software such as Linux, OpenStack, KVM, DPDK, and Ceph, Titanium Cloud incorporates a wide range of enhancements to achieve the levels of performance, platform management, security, and virtual machine (VM) or container lifecycle management that are required for telecom infrastructure. As an active contributor to all these open source projects, Wind River upstreams patches back to the community, collaborating with other companies to solve important industry problems with features critical to the telecom market.

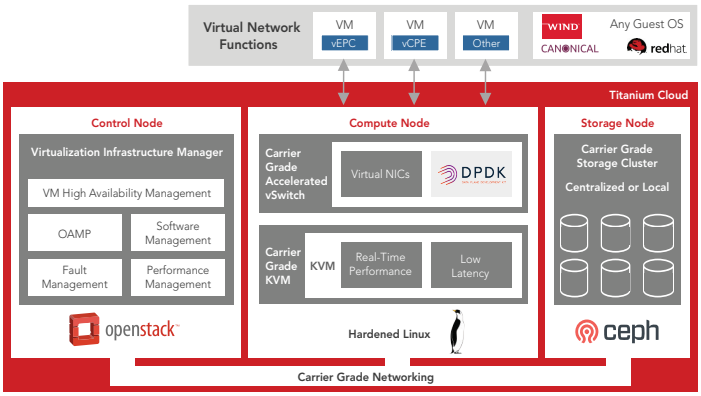

Figure3. Titanium Cloud architecture

To support low-latency edge applications and functions such as MEC, uCPE, vCPE, vRAN, and IIoT, Titanium Cloud delivers to guest VMs an ultra-low average interrupt latency of 3 µs, leveraging a low-latency configuration of the integrated KVM hypervisor. At the same time, its accelerated virtual switch (vSwitch) delivers

up to 40x the switching performance of kernel-based vSwitches, thereby minimizing the number of processor cores required for switching and maximizing the number available for running VMs. This in turn maximizes the number of users supported on a single server, which is key to reducing the overall cost per subscriber.

Titanium Cloud helps critical infrastructure companies minimize their operations costs, typically saving millions of dollars in installation, commissioning, and maintenance when compared to roll-your-own (RYO) solutions. While the platform is delivered as a single, pre-integrated image installed with no manual intervention, the intelligent orchestrated patching engine allows up to hundreds of nodes to be upgraded quickly and minimizes the length of maintenance windows.

The unique security risks of edge applications are addressed in Titanium Cloud by a comprehensive set of end-to-end security features. These include: Unified Extensible Firmware Interface (UEFI) secure boot; cryptographically signed images for host protection; virtual Trusted Platform Module (vTPM) device for highest-security VM deployments; secure API access and control; and secure keyring database for storage of encrypted passwords. Collectively, these features and others ensure that edge applications running on Titanium Cloud are protected against threats, wherever they originate.

By working with over 50 companies in the Titanium Cloud partner ecosystem and many others at industry plugfests, Wind River has already validated and demonstrated 100% compatibility both with open industry standards and with products from other industry-leading vendors. Through the StarlingX open source community, Wind River will build on these existing collaborations to accelerate the deployment of open source edge cloud solutions.

WIND RIVER AND STARLINGX

As part of the company's collaboration in the StarlingX OpenStack community, Wind River is contributing technology from the Titanium Cloud platform that delivers key capabilities and functions, including:

- Service management, REST APIs, and process monitoring
- Standalone fault management services, including extensions to the OpenStack Horizon project
- Software repository management, patching, upgrade, backup, and restore services
- Bare-metal management, a next-generation Virtual Infrastructure Manager (VIM) along with VIM helper components, the OpenStack Nova API proxy, and guest API infrastructure back-end

As a long-standing contributor to multiple OpenStack core service projects, Wind River has a proven track record of contributing blueprints that address important problems applicable to critical infrstaructure applications. Those contributions will continue in parallel with the commits to new projects that are part of StarlingX.

As one example of the business benefits that will result from this collaboration, the Wind River contributions will help the industry streamline the installation, commissioning, and maintenance of edge clouds. To date, the complexity of the operational aspects of existing solutions, along with the specialized IT knowledge required to deploy them, has impacted CSPs' ability to capitalize on the business opportunities presented by edge applications. By simplifying and automating these processes, Titanium Cloud technology contributed to StarlingX will minimize the cost, lower the risk, and reduce the staffing requirements for edge deployments.

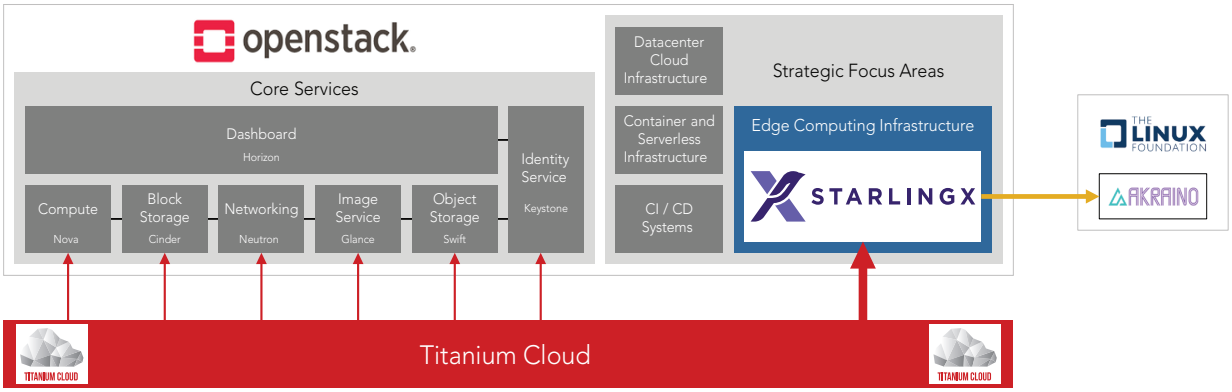

Figure 4. The StarlingX OpenStack project, including Titanium Cloud technology

CONCLUSION

Wind River is accelerating the development of the StarlingX OpenStack project by contributing technology from Titanium Cloud that has been proven to solve critical challenges associated with the reliability, performance, security, deployment, and lifetime operation of edge cloud applications. The expertise that the Wind River team has developed by addressing complex business and technical challenges represents a key resource for the StarlingX community.

For current edge cloud applications, StarlingX will enable critical infrastructure companies to boost their top-line revenue while minimizing operational costs. Looking ahead over the next few years as the industry moves beyond traditional telecom functions, companies will leverage StarlingX to capitalize on new business opportunities enabled by upcoming technologies, such as connected vehicles, augmented reality, telemedicine, and advanced drones. As a project outside the OpenStack core services, StarlingX avoids any need to fork the standard OpenStack code in order to deliver edge-optimized applications. Edge deployments based on the components in StarlingX automatically and seamlessly leverage updates to the mainline OpenStack code base.

Wind River is focused on accelerating the massive innovation and disruption at the network edge through this important industry initiative. Technology already proven in the Titanium Cloud platform will be key to enabling new business opportunities and innovative applications across multiple market segments.

WIND

Wind River is a global leader in delivering software for IoT. Its technology is found in more than 2 billion devices and is backed by world-class professional services and customer support. Wind River is accelerating digital transformation of critical infrastructure systems that demand the highest levels of safety, security, performance, and reliability.