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

HEAVY READING WHITE PAPER

Software Platform Design Strategies for Critical Infrastructure

A Heavy Reading white paper produced for Wind River Systems Inc.

WIND™

AUTHOR: JIM HODGES, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

The uptick in the adoption of software platforms is now a truly global phenomenon that is fundamentally reshaping how industries across the board manage and execute application delivery. Consequently, not only is the process of moving workloads and onboarding pure software applications to the cloud or on-premises accelerating, but the nature and scope of workloads making this shift are also expanding.

One of the most visible observations of this process is the redefinition of applications to critical applications that run on a new class of critical infrastructure, rather than bespoke, standalone network infrastructure. In short, the software compute era is now moving to a stage of growth and maturity that will see this new classification of critical applications moving beyond the telco domain to apply to functions such as connected cars and advanced, ultra-low-latency Internet of Things (IoT)-based critical applications.

While this process will unquestionably deliver scale and agility, the software platforms these applications will run upon must also evolve to conform to meet more demanding performance tolerances. Accordingly, this white paper documents how the adoption of multi-domain, critical applications is driving a holistic reassessment of the design requirements of software platforms. The white paper also presents a multi-faceted use case that illustrates the pertinent design considerations in more detail.

CRITICAL INFRASTRUCTURE; CRITICAL APPLICATIONS

Within the past five years, application delivery has transitioned from a data-center-focused architecture to a pervasive, more flexible, software-based, critical infrastructure model. In turn, underlying software platforms have evolved into a fluid and distributed software architecture to enable the delivery of cloud-native applications to the edge – thereby fulfilling the promise of edge computing. The risk of edge deployments from a security perspective is that when critical applications are deployed outside of the traditional data center, they are inherently more accessible to physical intrusion or software hacking.

At the same time, application and service templates are also changing, since any cloud, whether centralized or distributed, is now capable of delivering ultra-low-latency services across any vertical segment. Therefore, transportation, industrial, medical and telco segments are all looking to leverage these attributes to deliver high-value, critical applications.

But this service delivery model is not just about low-latency and high-bandwidth delivery. Rather, the focus is on how the cloud and edge can be commercialized to support the interworking of these cross-segment critical applications.

In recognition that the cloud and future technologies such as 5G would drive a profound service transformation by harmonizing service segments beyond the telco realms, in 2014 3GPP started to define "smarter" services (as they classified them) and their associated service design requirements.

This activity culminated in the release of 3GPP TR.22.891 V14.1.0, which for organization purposes segments critical application use-case services into five distinct service dimensions, as shown in **Figure 1**.

Figure 1: 3GPP Future Services SMARTER Service Dimensions

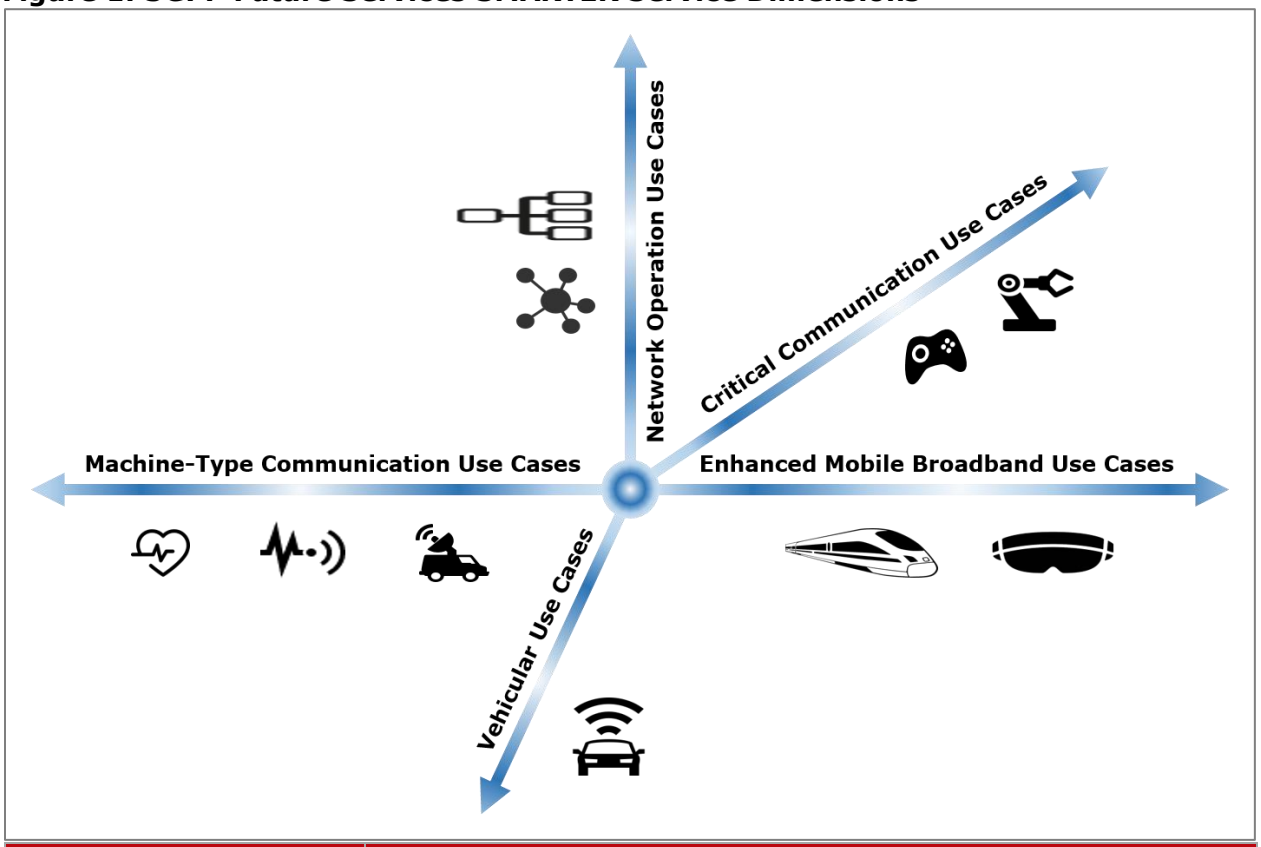

Service Dimensions	Applicable Use Cases
Enhanced Mobile Broadband	Mobile Broadband, Ultra HD/Hologram, High-mobility, Virtual Presence
Critical Communication	Interactive Game/Sports, Industrial Control, Drone/Robot/Vehicle, Emergency
Machine Type Communication	Subway/Stadium Service, eHealth, Wearables, Inventory Control
Network Operation	Network Slicing, Routing, Migration and Interworking, Energy Saving
Vehicular	Autonomous Driving, safety and non-safety aspects associated with vehicle

Source: Heavy Reading/3GPP TR 22.891 V14.2.0 (2016-09)

CRITICAL APPLICATION SECURITY REQUIREMENTS

Irrespective of how these service-driven use cases are classified, since they will run either in the cloud or on-premises, flexible approaches for securing critical applications are necessary. In turn, this means that underlying software platforms must integrate these new security design and software capabilities as well. To document these security-based design requirements on a more granular level, we break them into three categories:

- Development and Release
- Commercial Monitoring and Upgrade Lifecycle
- Operational Reach Expansion

Development & Release

Encryption is a design requirement for almost any software platform, cloud-based or non-cloud-based. However, software platforms supporting critical applications need to support additional advanced encryption techniques to ensure they can meet current and future cloud-driven security requirements.

These capabilities must be designed-in during the development phase and cannot simply be grafted on in future upgrade cycles. One approach to meet this requirement is to ensure that the software platform being designed is based on open security specifications.

An example of such a specification is the Trusted Platform Module (TPM) created by the Trusted Computing Group. TPM supports advanced encryption techniques that utilize software to authenticate devices and create trust domains. This will become critical as the number of devices grows and threat vectors adopt new approaches to compromise devices. TPM also supports the running of virtual machines on a virtual TPM (vTPM) to accommodate fully software-based implementations as scale and security policies dictate.

Critical infrastructure platforms also need to support built-in design features such as secure boot that are tested prior to release. Secure boot is particularly important because it protects the integrity of a product by ensuring that the image it boots from has not been tampered with, or in any way altered since it was originally securely delivered and installed.

Monitoring & Upgrade Lifecycle

Critical infrastructure software platforms must also support advanced monitoring and upgrade capabilities to respond to the changing security threat landscape. This means that software platform developers must be committed to monitoring the latest in community threat data and assessing potential security platform impacts.

This is vital, since the threat landscape has become so dynamic and active that monitoring must be current, so that proactive steps can be taken to upgrade platforms sooner vs. later. In addition, software platforms must support the ability to take hitless upgrades so that performance is not compromised at the expense of security upgrades. To be clear, hitless software upgrades have always been important, it's just that it is even more essential for critical infrastructure software platforms, given the extremely tight performance service tolerances and the life-and-death impacts they may have on service outcomes.

Another consideration of an effective upgrade or patching program is the incorporation of cryptographic signatures to validate the integrity and source of patches and software code when upgrades are necessary. For purposes of completeness, this process should also include attaching signatures not only to patches, but also during the initial code build of the software platform.

Operational Reach Expansion

Critical infrastructure-based applications also impose additional operational requirements on software platforms and their related control systems. Unlike non-cloud critical platforms, which didn't require a focus on expanded operational control systems, critical infrastructure platforms must support expanded operational reach. This is in large part because IoT breaks the operational model and demands operational reach expansion, since it is by design a

connected multi-domain service that will require more flexible control systems to enable IoT applications to run anywhere in the cloud, even in multiple clouds or on customer premises.

Moreover, IoT will push the boundaries and must be flexible enough to enable onboarding of new IoT applications with different service profiles and security requirements. This is because we have only begun to scratch the surface of IoT use cases and their related security requirements (see **Critical Application Slice Use Case**, below).

Consequently, software platforms must be designed with the operational flexibility to support this expanded and dynamic connectivity model. Another operational but business-related consideration is that traditional software platforms embedded rigid control systems that are more expensive to maintain and scale compared to the lower cost of ownership associated with critical infrastructure platforms, which are flexible and highly scalable in a security context.

CRITICAL INFRASTRUCTURE CONVERGENCE: THE IMPACT OF MOBILE

The current adoption of the mobile cloud leveraging virtualization to enhance 4G performance and facilitate the adoption of multi-access edge computing (MEC) represents nothing short of a sea change in application performance and delivery. And 5G will take this to the next level, because 5G is not simply delivering lower-latency performance than 4G cloud-based networks; it is designed to support and harmonize the seamless delivery of cross-segment critical applications from multiple domains.

Accomplishing this feat requires the adoption of a novel and disruptive approach – 5G network slicing. On a most basic level, as shown in **Figure 2**, network slicing works by the creation of software-based application and use-case-specific partitions, known as Network Slice Instances (NSI), in the 5G RAN (a.k.a. 5G New Radio) and core networks.

Figure 2: Network Slicing – Use-Case Relationships

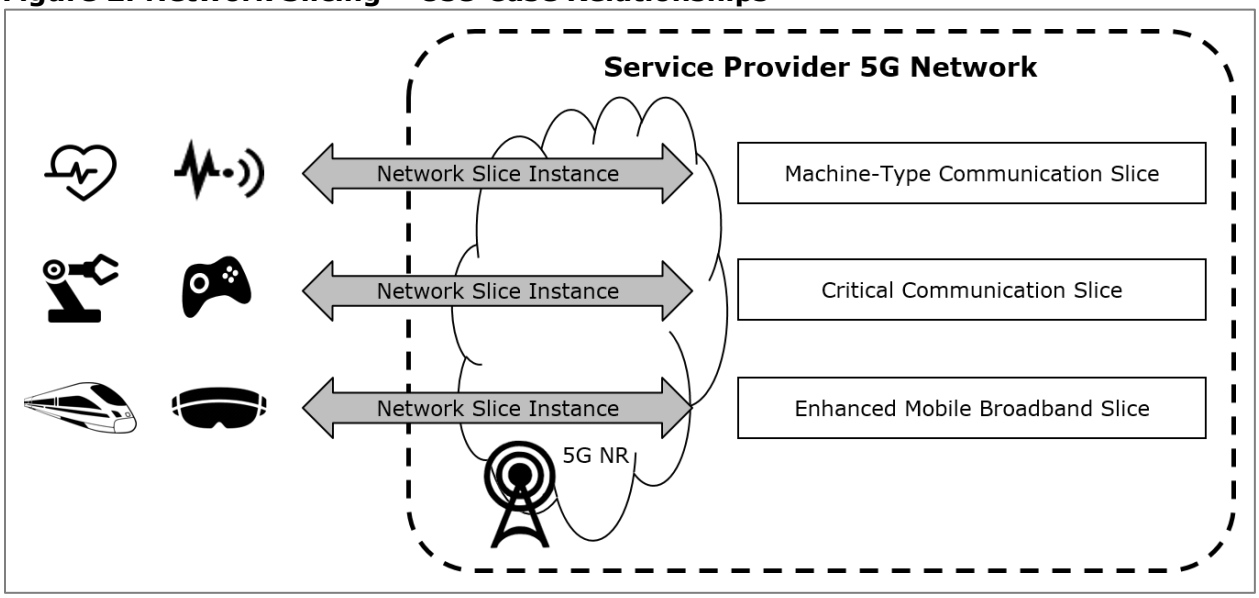

Source: Heavy Reading/3GPP TR 22.891 V14.2.0

In effect, network slicing necessitates the real-time assembly of all the logical network functions to activate and deactivate a specific use-case type, which introduces additional security requirements for critical infrastructure. While these security requirements continue to be defined, 5GPP, an EU-based organization focused on monetizing and securing 5G services, has created a 5G ENSURE reference architecture.

This architecture, as shown in **Figure 3**, extends the 3GPP security architecture from TS 33.401 and domain concept from 3GPP TS 23.101 to create flexible trust models for an individual slice instance. This is accomplished through the introduction of a Trust Anchor (TA), which supports the exchanging of user credentials to validate the identity of parties in the service path.

Figure 3: 5G ENSURE Network Slice With Trust Anchor

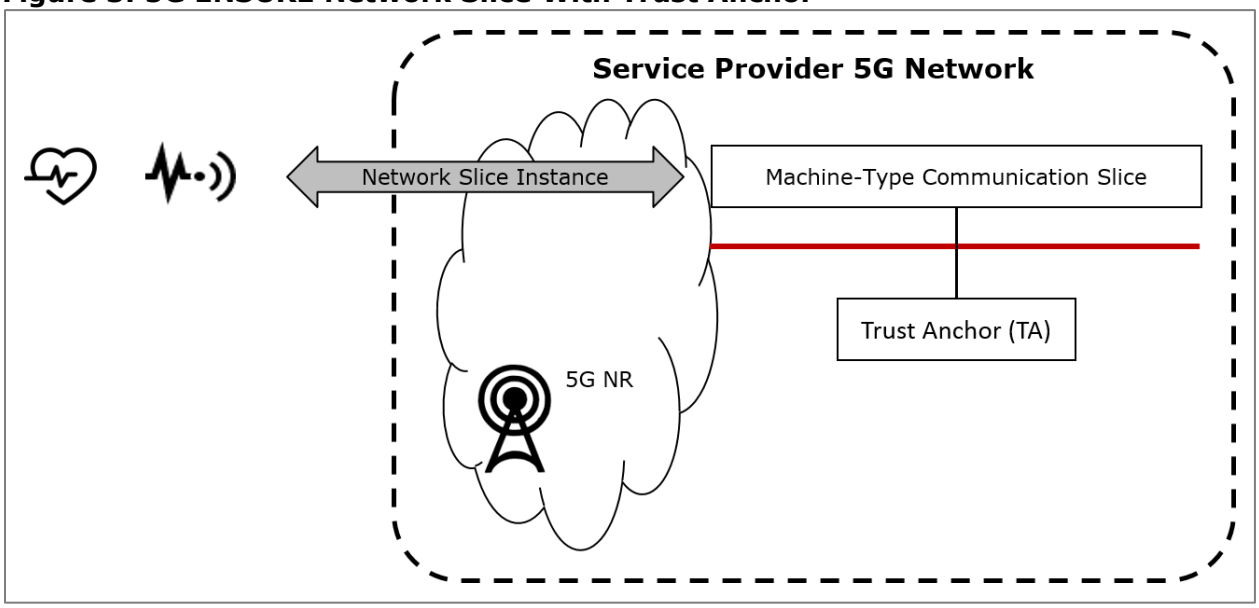

Source: Heavy Reading interpretation of 5G PPP Architecture Working Group – View on 5G Architecture (Version 2.0) – 2017-07-18 and 3GPP TR 22.891 V14.2.0

In addition to the Trust Anchor, each slice is reinforced by a security control mechanism that supports a broad range of security functions and encryption mechanisms addressing identity/trust management auditing.

This security control mechanism plays an important role, given that trust in this multi-domain world is fundamental from a security policy enforcement perspective. Therefore, it is also imperative that the design cycle of underlying software platforms incorporate a flexible, programmable, software-defined model capable of adapting to the security requirements of the individual use cases.

CRITICAL APPLICATION SLICE USE CASE

In this final section of the white paper, we present a representative use case designed to tie together the security requirements with the software requirements of critical infrastructure platforms.

The use case selected and illustrated in **Figure 4** documents the interaction of an IoT medical wearable with a 5G-equipped ambulance. This use case traverses several smart service categories, including machine-type communications, enhanced mobile broadband, critical communications and even network operations (e.g., slicing).

- In this scenario, an alert has been received from the wearable device. An ambulance is dispatched, locating a non-verbal patient based on GPS data provided by the device.
- Once the patient has been located but before treatment is started in the ambulance, first responders link the device data into the ambulance IoT interface to make an initial medical risk assessment.
- Since this patient is non-verbal, drug allergies are not necessarily known, so the ambulance control module creates a secure 5G slice to the nearest hospital database to find a medical profile based on name and registration data provided by the device. When this is not found, the ambulance searches other external linked databases and discovers a profile for the patient in another city. Alternatively, the medical wearable device could also potentially store all medical information and profile data, which could simply be downloaded by first responders.
- Based on the patient's profile, first responders commence treatment. This treatment is delivered in real time and updated in the ER database to ensure that all pertinent medical data is available to ER teams as the ambulance is in transit. In the future, this ambulance will be an autonomous, self-driving vehicle that can utilize a 5G transportation slice to calculate and execute the fastest, most expedient path to the ER.
- Once the patient is stabilized in the ER, medical data and associated case management data is provided to the patient's hospital database to be assessed by their medical team. In addition, the patient's emergency contacts are notified of the event by law-enforcement staff.

Figure 4: Critical Medical Application Use Case

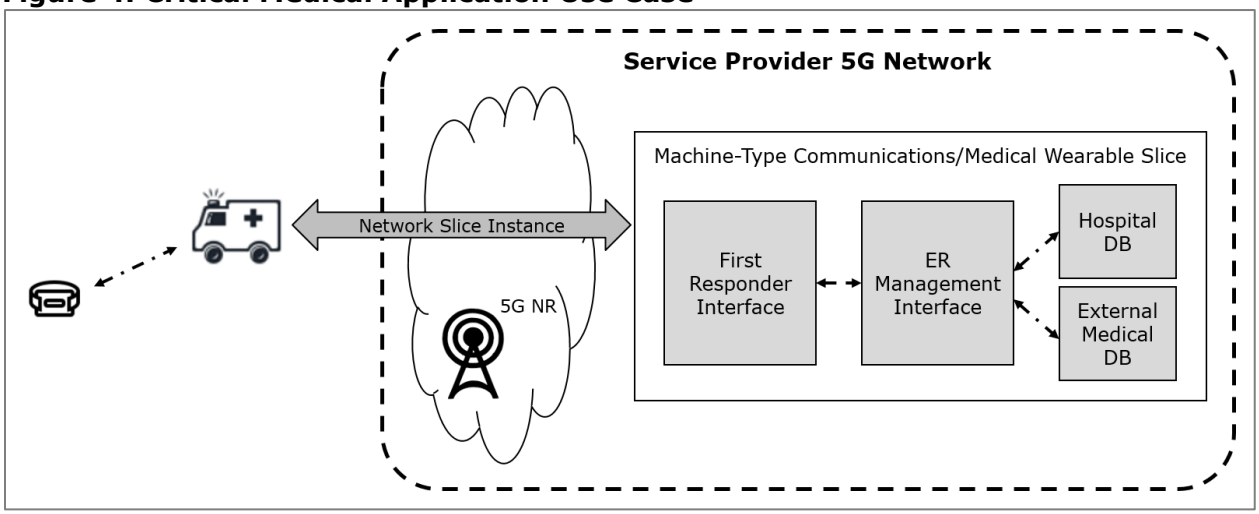

Source: Heavy Reading

As we have documented, each network slice-based critical application use case possesses unique security requirements, which must be designed into the software platforms that

support them. In this use case, the specific design-cycle requirements, mapped to the three design-related categories defined above, are as follows:

Development & Release

One key software design consideration in this use case relates to trust domain and data sharing. In this instance, since the patient is non-verbal and third-party databases are utilized, it's crucial that end-to-end encryption be supported from the wearable device itself to any of the multiple interfaces in the network slice path. Moreover, before commencing treatment it would be necessary to validate that the patient profile image was not subjected to tampering by any third party.

Monitoring & Upgrade Lifecycle

Due to the complex nature of this use case, several servers and critical access databases would be required to manage and execute the use-case slice. This applies to not only access to these databases, but also mobile devices – in this case, the wearable and the ambulance itself.

Therefore, flexible software programmability tools are mandatory to ensure successful outcomes of this current iteration – and even future iterations in which the ambulance evolves to become an autonomous vehicle and machine learning is integrated into treatment protocols. To achieve this level of secure flexibility on an end-to-end basis requires that software upgrades sourced natively or from third parties support advanced capabilities such as cryptographic signatures, to mitigate the risk of loading malicious third-party software into the upgrade cycle.

Operational Reach Expansion

This use case exemplifies the strong value proposition associated with the delivery of critical applications, as well as the security-related operational reach considerations. First, as an IoT-based application capable of running anywhere in the cloud, there must be a strong measure of operational flexibility to enable various parties to transfer and exert control as the use case progresses.

For example, in this use case, while first responders have the greatest amount of control initially, as treatment is extended utilizing the ER's more advanced medical tools, operational control intrinsically shifts – and this must be accomplished in a seamless manner. This is not as straightforward as it seems, given that prior to the cloud migration many of these functions were initially designed as standalone applications (e.g., law enforcement and first responders) with limited support for real-time interworking.

CONCLUSION

The software domain has now entered a new expansion phase – one that will transform software development into a multi-domain, application-centric architecture that provides the foundational underlay to execute critical applications across all domains. To secure these applications, a much-needed reassessment is underway to define the additional design capabilities that software platforms will need to accommodate the scale, trust and identity management attributes that critical infrastructure will mandate.

ABOUT WIND RIVER

Wind River is a world leader in delivering software for the Internet of Things. With its technology found in more than 2 billion products, Wind River offers the industry's most comprehensive edge-to-cloud software portfolio, supported by world-class global professional services and support and a broad partner ecosystem. Wind River delivers the technology and expertise that enables the innovation and deployment of safe, secure and reliable intelligent systems.

For more than 30 years, Wind River has helped the world's technology leaders power generation after generation of the safest, most secure devices in the world. Companies managing or delivering critical infrastructure services can turn to the Wind River Titanium Cloud family of products to secure their cloud environments and safeguard their ongoing business operations.

The Wind River Titanium Cloud portfolio includes the industry's only fully integrated and deployment-ready virtualization platforms that deliver the uptime, performance and security needed for communications networks, industrial applications and control services at any scale. When service uptime is critical for profitability, Titanium Cloud products ensure virtualized services run when, where and how they need to, always.