



Wireless Network Virtualization: Ensuring Carrier Grade Availability

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

The wireless industry's battle to acquire new subscribers and retain existing ones is accelerating the need for new services. Profit margins are under pressure from the increased infrastructure and operations costs required to satisfy the growing demand. Network functions virtualization (NFV), a rapidly growing initiative in telecom networks, promises to revolutionize how networks are architected and managed. It allows communications service providers (CSPs) to virtualize network functions and consolidate them on standard off-the-shelf servers.

Although 4G LTE is effective at meeting the increasing bandwidth demands at lower costs than its predecessors, its mobile base station, E-UTRAN Node B (eNB), located at the edge of the radio access network (RAN), is underutilized at certain times of the day and has grown in complexity, resulting in higher downtime and field maintenance costs. As a result, the eNB is a good candidate for NFV. Cloud-RAN (C-RAN) is the virtualization of the eNB's control plane and data plane functions, consolidating it in one or more data centers. The result is significantly higher equipment utilization, cost-efficient redundancy to achieve high availability, and lower operations and maintenance costs.

However, to realize these benefits, the foundation for virtualized eNBs must be a robust, carrier grade NFV platform that incorporates advanced fault management features. Poor implementation results in lower quality execution with excessive outage and maintenance costs. Wind River® Titanium Cloud is the industry's first NFV-ready solution that incorporates advanced carrier grade fault, security, performance, and network management features. Titanium Cloud supports a broad range of affordable redundancy configurations, giving CSPs the flexibility to meet a range of availability targets, including "six 9s," or 99.9999% uptime at the service level.

TABLE OF CONTENTS

Executive Summary	2
Network Transformation: The Next Step	3
NFV Building Block	3
C-RAN: Reliability-Related Opportunities	4
Carrier Grade Quality	5
Wind River NFV Platform	6
Conclusion	7

NETWORK TRANSFORMATION: THE NEXT STEP

The wireless carrier network is undergoing significant transformation. New multimedia broadband services and universal device connectivity continue to fuel the rapid growth in network traffic. To remain competitive, CSPs have taken advantage of multi-core-based hardware and open software platforms to reduce costs and accelerate service introduction times. The next step in this transformation, enabled by these technologies, is NFV. NFV is a carrier network architecture in which network functions are virtualized and chained together to provide communication services that can be more efficiently managed by CSPs.

NFV started in 2012 as an initiative by leading CSPs to take advantage of IT virtualization concepts and apply them to carrier networks. IT virtualization has a proven track record in reducing enterprise IT costs, but it lacks the robustness and carrier grade features required by CSPs. The carrier version addresses the IT version’s shortcomings, such as poor interrupt performance, lack of determinism, slow data path performance, and inadequate fault and network management capabilities. Using NFV, network functions can be virtualized, giving CSPs unprecedented flexibility on where and how equipment and services are deployed. NFV decouples network functions from network equipment so that a virtualized network function (VNF) can run on many virtual machines (VMs) distributed across one or more servers, which are consolidated either centrally or in multiple geo-diverse locations.

Because VNFs have access to a pool of computing and storage resources anywhere in the network, CSPs have the flexibility to deploy network functions where they are most cost-effective—only limited by latency considerations. As a result, NFV provides dynamic or “elastic” scalability, highly efficient resource utilization, automation of network management tasks, and support for any redundancy options, all requiring less equipment than legacy strategies. This means NFV networks can achieve high performance and availability at lower equipment costs.

Take-away: *NFV is transforming the network to give CSPs the flexibility to locate network functions where they are most cost-advantageous. NFV also provides a cost-efficient means to implement redundancy to achieve high availability solutions.*

The carrier version addresses the IT version’s shortcomings, such as poor interrupt performance, lack of determinism, slow data path performance, and inadequate fault and network management capabilities.

NFV BUILDING BLOCK

The NFV conceptual architecture, shown in Figure 1, consists of a virtualization layer that decouples the VNFs from the hardware so that network functions can be substantiated on one or more VMs across multiple servers. The NFV platform is the building block required by network equipment providers (NEPs) to design NFV network equipment. It is an NFV-ready software and hardware platform that extends the capabilities of IT virtualization by incorporating carrier grade features used by the platform’s middleware and virtualization layers.

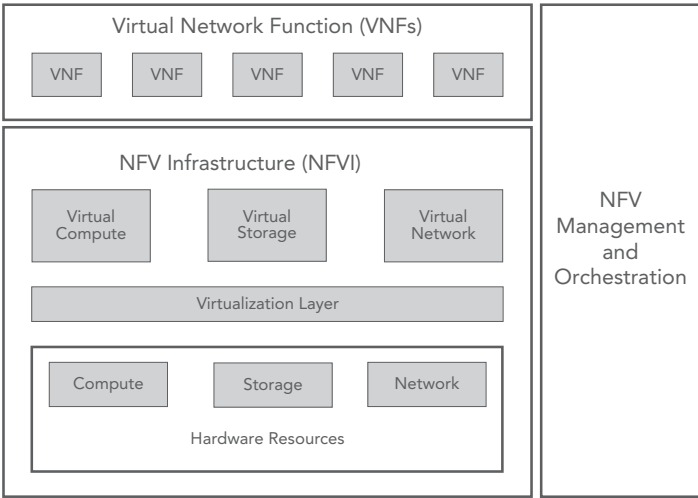


Figure 1: NFV building block

The NFV platform is also used to run the VNF management and orchestration function, which is responsible for creating, provisioning, and dynamically managing VNF-based services. This function provides system-level monitoring of the health of VNF systems, removing failed VNFs from service and reconfiguring VNFs to auto-heal services. This layer exploits software-defined networking (SDN) techniques that separate the control and data planes to virtualize network management functions in order to simplify and automate the management of the network.

NFV can be used to implement any type of redundancy, such as N load-shared, 1:1, N:1 and N:M. For example, a pool of VNFs can be configured and adequately provisioned to share the load such that if one NFV server fails, it is removed from service, and the remaining servers can then handle the traffic. When implemented using carrier grade NFV servers, the system can provide “hitless” failover protection.

Take-away: *NFV is being used to build platforms and solutions that incorporate carrier grade features required by CSPs to build robust, fault tolerant networks that are easier to operate, manage, and maintain.*

C-RAN: RELIABILITY-RELATED OPPORTUNITIES

4G LTE’s mobile base station, eNB, provides wireless coverage over a specific cell area by processing control and data information to and from the mobile device and to and from the core network via a backhaul optical network. A group of eNBs cannot share the load, so each system must be provisioned adequately to handle peak traffic loads. Consequently, the average utilization of each eNB is sub-optimal, resulting in the inefficient use of power, processing, and memory resources. By virtualizing the eNB’s control functionality, the C-RAN solution can efficiently share the traffic load, as well as provide efficient redundancy.

The control and packet processing function is stripped from the legacy eNB, deployed on NFV-based platforms, and consolidated in a data center, leaving the radio portion, referred to as the remote radio head (RRH), at the edge (Figure 2). The RRH is a much simpler device, with lower cost, failure rate, and downtime than the legacy eNB. The digital functionality, referred to as the baseband unit (BBU), is deployed across multiple NFV platforms, resulting in significantly better equipment utilization and scaling

than can be achieved by the traditional RAN approach. It also enables a cost-effective way to achieve high availability at lower maintenance costs.




Figure 2: RAN and C-RAN

Figure 3 compares the legacy RAN and the virtualized C-RAN using two reliability metrics: “cell area downtime” measures the average minutes per year in which all users in the coverage area are without service due to base station failures; and “unplanned maintenance actions” measures the rate at which CSP operations staff must respond to unplanned failures. The analysis considers only steady-state hardware and software failures for the base station functionality, and assumes a mean time-to-repair of three hours.

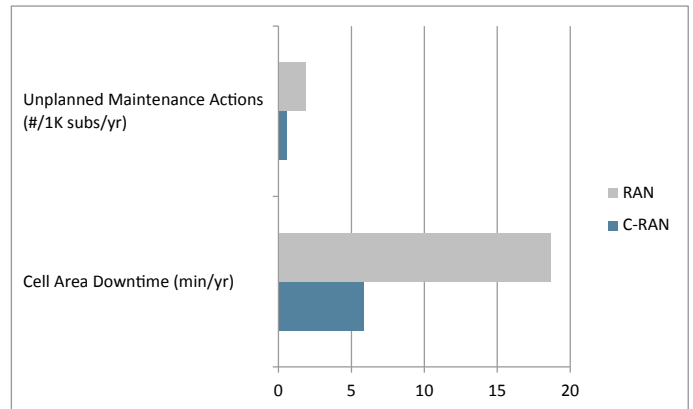


Figure 3: RAN/C-RAN Comparison

Because the RANs’ eNBs do not incorporate redundancy, the predicted unplanned downtime is high—in this case 18.8 minutes per year, compared to 5.8 minutes per year for the C-RAN. Most of the C-RAN downtime is due to the non-redundant RRH portion of 5.4 minutes per year, while only 0.45 minutes per year—better

than 99.99991% availability—is due to the BBU. This high availability is achieved because the virtualized BBUs can automatically recover from both hardware and software failures with no impact on mobile communication sessions.

Adding redundancy normally increases maintenance actions because of the extra equipment. However, the C-RAN solution achieves almost a three-fold reduction in maintenance actions, because less equipment is required for the C-RAN solution as a result of virtualization’s higher equipment utilization.

C-RAN consolidation in a data center also provides other reliability-related benefits not directly quantified in the table. Average travel time to repair failed BBUs when deployed in attended data centers is reduced, resulting in lower maintenance costs and shorter outage durations. Maintenance costs are further reduced due to the lower cost of the inventory of maintenance spares because of centralization and the smaller number of commercial off-the-shelf (COTS) board types.

Take-away: Virtualizing the eNB’s control function can achieve a three-fold improvement in cell downtime and reduction in maintenance costs over the current RAN implementation.

CARRIER GRADE QUALITY

The public network provides mission-, business-, and life-critical services, and therefore must be resilient to failures. Outage costs can range from thousands to millions of dollars per hour depending on the service and failure mode impact. Maintenance costs increase with the amount of time to diagnose and repair the failure, and with the reliability of the equipment. The following illustrates the impact of three reliability quality factors using the C-RAN example.

Reliability of design and technology determines equipment failure rates. Functional design problems such as interoperability, inadequate robustness to overloads, design errors, and component reliability problems have a significant impact on failure rates.

Fault coverage and fault isolation are the equipment’s ability to detect, contain, recover, alarm, and diagnose failures, measured as a percentage of the total failure rate. Low fault detection coverage results in more frequent and longer outages. Alarms and diagnostics with poor fault isolation increase field diagnosis time. Figure 4 illustrates the impact of fault coverage on cell area downtime.

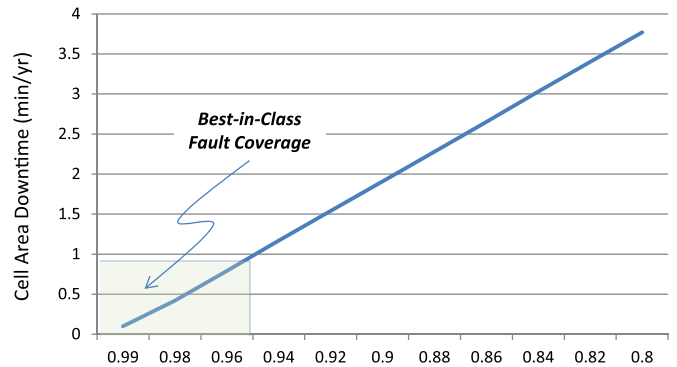


Figure 4: Impact of fault coverage

Best-in-class fault coverage is greater than 95%. Fault detection coverage in the range of 80% is not uncommon, resulting in a factor-of-four increase in downtime—accompanied by higher outage, maintenance, and product upgrade costs.

The benefit of **fault detection and failover speed** is best measured by the “probability of dropped communication session” metric. Using the C-RAN example from the previous section, Figure 5 illustrates the difference between hitless and non-hitless C-RAN designs.

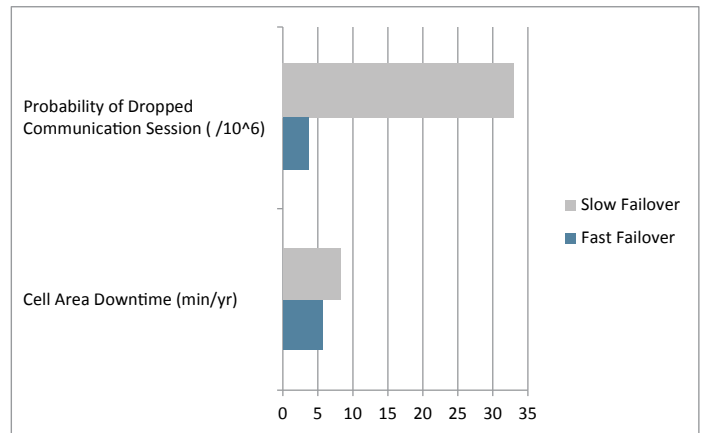


Figure 5: Hitless vs. non-hitless failover

The slower failover case results in marginally higher downtime, as the failover times are relatively small compared to repair times. But there is a ten-fold increase in dropped sessions because all successful failovers drop sessions—resulting in increased customer complaints.

Take-away: Not all carrier grade solutions are equal. Design and implementation quality of the fault management features have a significant impact on reliability-related metrics. Good implementation quality is required to meet the “always-on” expectations of end-users and minimize outage and maintenance costs.

WIND RIVER NFV PLATFORM

Wind River Titanium Cloud is the industry’s first fully integrated, open source, feature-complete NFV server (Figure 4). It has a full complement of carrier grade fault management features and employs advanced robustness techniques to form the foundation for building carrier grade products, networks, and services. Titanium Cloud supports hitless software upgrades and provides a variety of redundancy options. This means the designer can implement a redundancy scheme that is optimal for the particular VNF service. As a result, NFV products that use Titanium Cloud as their platform can achieve service availability levels from five 9s to better than six 9s, depending on the criticality of the service.

Titanium Cloud is built on carrier grade Linux, a hardened, dependable platform. It complies with the SCOPE Alliance Linux profiles requirements, while also addressing key gaps. For example, features include shared memory with the system black box, coherent user and kernel tracing framework, fast system reboots, advanced error detection and recovery, run time analysis, and low latency event monitoring.

Wind River Open Virtualization employs a real-time virtualization function that incorporates kernel preemption at very low interrupt latency to achieve deterministic, predictable performance resulting in more reliable NFV systems. This reliability is made possible by the carrier grade vSwitch, which enables high VM-to-VM communications to achieve ultra-high packet processing.

Titanium Cloud adds fault handling extensions to OpenStack, which was originally designed for IT-grade virtualization, to create a carrier grade version of OpenStack that achieves VM migration in hundreds of milliseconds rather than minutes, faster VM failure detection time (60 times faster than the IT version), more efficient VM resource management, and faster host and controller node failover times.

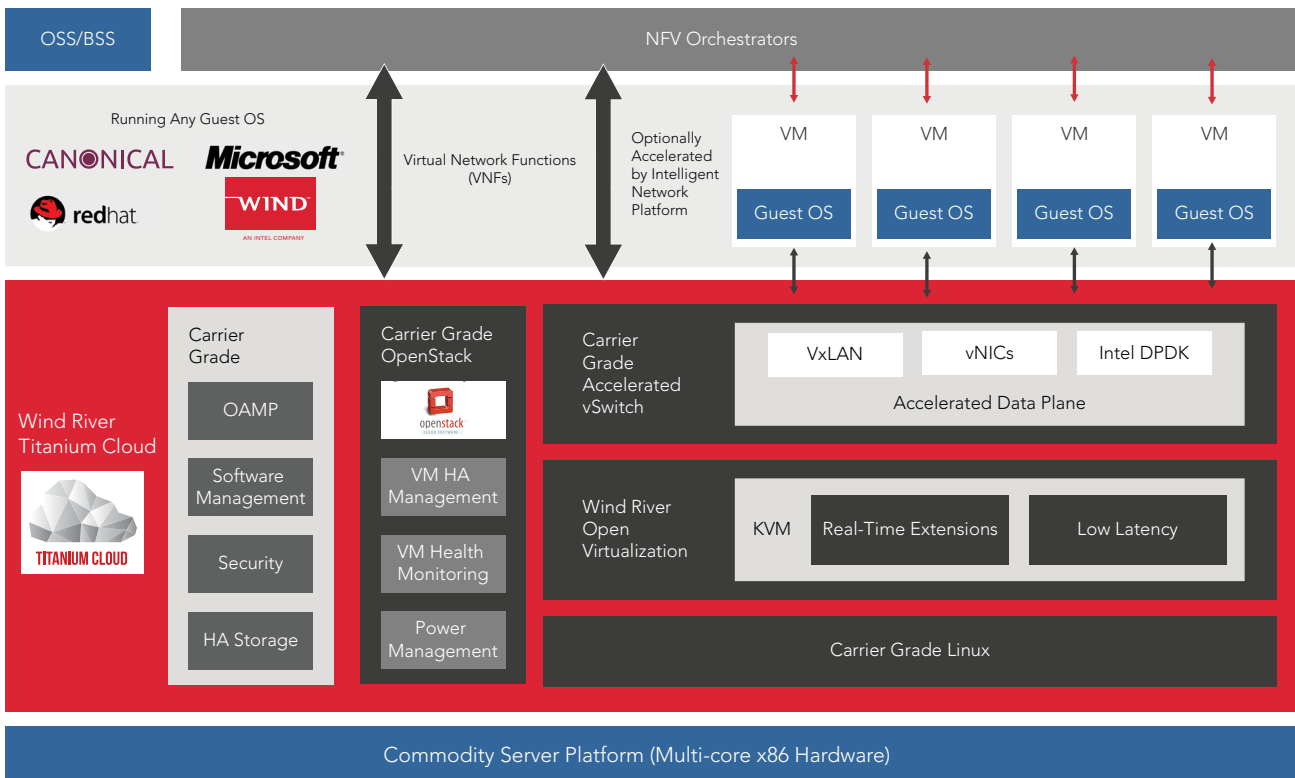


Figure 6: Wind River Titanium Cloud



The solution's carrier grade middleware, which provides VM "health" monitoring and management, supports live patching and hitless software upgrades with automatic fallback to the previous release for failed insertions. VM protection groups can be set up on different compute nodes to eliminate common failure modes. Its VM health check monitoring is integrated into application health check within the guest, resulting in higher fault coverage.

Take-away: *Wind River Titanium Cloud is the industry's first fully integrated carrier grade NFV server. It implements advanced carrier grade features to achieve C-RAN high availability at low maintenance costs.*

CONCLUSION

The virtualization of wireless carrier network functions provides CSPs with unprecedented flexibility regarding where and how equipment and services are deployed. NFV decouples network functions from network equipment so that a VNF can run on many VMs distributed across multiple servers, which, in turn, can be consolidated in a data center. This flexibility enables CSPs to optimize network designs to reduce cost yet achieve high availability.

The low equipment utilization and relatively higher downtime and maintenance costs of the RAN's eNB make it a good candidate for NFV. By virtualizing the eNB's control functions and consolidating them in data centers, CSPs can realize significant cost savings. NFV enables CSPs to incorporate the cost-efficient redundancy required to increase availability and reduce maintenance costs by three-fold over the RAN solution. However, to realize these improvements, the C-RAN must be built using carrier grade NFV solutions, as poor implementation quality erodes these benefits. Wind River Titanium Cloud is the industry's first fully integrated, open source, feature-complete NFV server with a full complement of carrier grade fault management features.

About Wind River

Wind River, a wholly owned subsidiary of Intel® Corporation (NASDAQ: INTC), is a world leader in embedded software for intelligent connected systems. The company has been pioneering computing inside embedded devices since 1981, and its technology is found in nearly 2 billion products. Wind River offers the industry's most comprehensive embedded 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. To learn more, visit Wind River at www.windriver.com/without-compromise, or call 800-872-4977.

